Journal of Applied Logic 11 (2013) 373–385

Contents lists available at ScienceDirect

Journal of Applied Logic

www.elsevier.com/locate/jal

The sure thing principle, dilations, and objective probabilities ☆

Haim Gaifman

Columbia University, New York, NY 10027, United States

ARTICLE INFO

Article history:
Available online 13 August 2013

Keywords:
Sure Thing Principle
Axiom P2
Imprecise probabilities
Dilation
Partial ordering
Objective probabilities
Savage's system
Partial acts

ABSTRACT

The common theme that unites the four sections is STP, the *sure thing principle*. But the paper can be divided neatly into two parts. The first, consisting of the first two sections, contains an analysis of STP as it figures in Savage's system and proposals of changes to that system. Also possibilities for partially ordered acts are considered. The second, consisting of the last two sections, is about imprecise probabilities, dilations and objective probabilities. Variants of STP are considered but this part is self-contained and can be read separately. The main claim there is that dilations, which can have extremely counterintuitive consequences, can be eliminated by a more careful analysis of the phenomenon. It outlines a proposal of how to do it. Here the concept of objective probabilities plays a crucial role.

1. The sure thing principle, Savage's system, and partial acts

The sure thing principle, henceforth abbreviated STP, was introduced by Savage in his seminal work [10] by means of the following story:

> A businessman contemplates buying a. certain piece of property. He considers the outcome of the next presidential election relevant to the attractiveness of the purchase. So, to clarify the matter for himself, he asks whether he would buy if he knew that the Republican candidate were going to win, and decides that he would do so. Similarly, he considers whether he would buy if he knew that the Democratic candidate was going to win, and again finds, that he would do so. Seeing that he would buy in either event, be decides that he should buy, even though he does not know which event obtains, or will obtain, as we would ordinarily say. It is all too seldom that a decision can be arrived at on the basis of the principle used by this businessman, but except, possibly for the assumption of simple ordering, I know of no other extralogical principle governing decisions that finds such ready acceptance.
>
> [10, p. 21][1]

☆ Thanks are due to many people: to Jeff Helzner, for organizing the conference, editing the volume, and patiently waiting for my final draft; to the participants – for making this an enjoyable high quality conference; to Teddy Seidenfeld – for his observations, in long email exchanges, which forced me to think harder on the subject (he also pointed out to me additional works that are relevant to the points I am discussing, but which I could not include in the present paper and may have to address in future works on these topics); to Anubav Vasudevan, Sidney Felder, Liu Yang, Rush Stewart and all the participants of the Columbia Formal Epistemology workshop, for the occasion to present some of these ideas and to get their useful reactions; and to Isaac Levi for fruitful exchanges. I would also like to mention Horatio Arlo-Costa, no longer with us; his instructive observation on related material were quite helpful at an earlier stage of this work.

E-mail address: hg17@columbia.edu.

[1] All page numbers for Savage's book refer to the 1972 revised edition.

1570-8683/$ – see front matter
http://dx.doi.org/10.1016/j.jal.2013.06.003

Underlying the scenario is the obvious assumption that the acts in question – buying the property or not buying it – have no effect whatsoever on the result of the election. They do not change the likeliness of any possible outcome.

Let us recall that Savage's framework is based on: (i) A set S of *states of the world* and a Boolean algebra of *events* (or propositions) that are sets of states, (ii) A set of *acts*, where each act is a function, f, that associates with each state, s, an *outcome*, $f(s)$ – the result of the agent's performing the act in a world in state s, and (iii) The agent's preference relation, which is a simple ordering (i.e., total preorder) of the acts.[2,3] There is a crucial assumption that the obtaining of any state is an objective fact, which is unaffected by the agent's act. We can make it even more obvious, by modifying a little the story above: Assume that the election took place already, and the businessman does not know the result – say he is on a trip in a remote region, but he has the option of sending a message to his office instructing his staff to carry out the transaction. Given STP, he can do so without waiting to find which candidate won.

The above scenario concerns strict preference for buying the property, but the same logic applies to weak preference: if the businessman decides that, in each possible outcome of the election, buying is at least as good as not buying, then, without knowing the outcome, he can conclude that buying is at least as good as not buying.

The force of STP does not depend on there being only two options. It would apply equally well if there were many; say, he could buy or not buy any of a number of lots. If, under either of the assumptions, that the Republican will win and that the Democrat will win, buying Lot 1 is preferable (or weakly preferable) to buying Lot 2, then, without any assumption, buying Lot 1 is preferable (or weakly preferable) to buying Lot 2. STP concerns two acts under explicitly stated assumptions; its force derives only from these explicit assumption, and it remains the same, whatever properties the preference relation has, be it a total ordering, as Savage assumes, or only a partial one as is the case in other systems developed in the last forty years.

Let $\leqslant$ be the weak preference relation (of a rational agent): $f \leqslant g$ means that g is at least as good as f.[4] The strong ordering (g is better than f) can be then defined by: $f < g \Leftrightarrow_{\mathrm{Df}} f \leqslant g$ and $g \nleqslant f$. Savage's first postulate, P1, states that this is what he calls a "simple ordering" and which in current terminology is known as a total preorder: it is a reflexive transitive relation, and its being total means that for every f and g, either $f \leqslant g$ or $g \leqslant f$. In a preorder the conjunction $f \leqslant g$ & $g \leqslant f$ defines an equivalence relation, $\equiv$, which need not be the identity relation (as it is in total orderings); we get an induced ordering of the equivalence classes. Since different acts can be equi-preferable, "ordering" in the context of preference means a preorder, and this is the way I shall use "ordering" in the paper. That the ordering is total means that for any f, g exactly one of the three possibilities holds: either $f < g$, or $g < f$, or $f \equiv g$. In the last case we say that the agent is *indifferent* between f and g. Savage, who adopts STP without reservations, prefers however not to introduce it directly, since it would require additional machinery not included in his system (I shall return to this point shortly). He uses it rather as a guiding principle that leads him to his second postulate P2.

P2, and other postulates, can be stated more transparently than in Savage's book, by using partial acts: Let S be the set of all states, then a *partial act* is a function, r, defined over some event $E \subseteq S$, $E = \mathrm{Dom}(r)$, which assigns outcomes to all states in $\mathrm{Dom}(r)$ and is undefined otherwise.[5] An "act" in Savage's terminology (which is also the current adopted terminology) is *total*: it is a partial act whose domain is S. For a (total) act f and an event B, let $f|B =$ *the restriction of* f to B; it is the partial act with domain B, such that $f|B\,(s) = f(s)$, for all $s \in B$. Obviously, $f|B \cup h|B'$ is a partial act iff f and h agree on $B \cap B'$. In particular it is a partial act if $B \cap B' = \emptyset$; In that case let us rewrite $f|B \cup h|B'$ as:

$$f|B + h|B'$$

We shall use this notation as implying that $B \cap B' = \emptyset$. $f|B$ can be further restricted to C, where $C \subseteq B$; obviously $(f|B)|C = f|C$.[6] Unless stated otherwise, 'f', 'g', ... range over total acts. Hence '$f|A$', '$f|B$', ..., '$g|A$', '$g|B$', ... range over partial acts. Let $\overline{B} = S - B =$ the complement of B.[7]

The setup presupposes some non-empty set of total acts, call it Act. Savage assumes implicitly *closure under cut-and-paste*:

- If f, $g \in$ Act and B is any event, then $f|B + g|\overline{B} \in$ Act.

The second postulate can be now stated as follows:

P2 For every event B: $f|B + h|\overline{B} \leqslant g|B + h|\overline{B} \Rightarrow f|B + h'|\overline{B} \leqslant g|B + h'|\overline{B}$

[2] In possible-world lingo, the states are the possible worlds. They represent however only those aspects that are relevant to the problems at hand and, most importantly, the agent's choice of act is external to the possible world, not part of it.

[3] Savage used bold face letters $\mathbf{f}, \mathbf{g}, \ldots$ for the acts, and $f, g, \ldots$ for the associated functions, although acts with the same associated functions are regarded equal. He also used 'f' 'g' ambiguously, to denote constant outcomes (the same in all states), in the cases where the act yields the "same" outcome in all states. We diverge from his notation also in other details.

[4] Savage uses a special small-or-equal notation for acts; a special notation, such as '$\preccurlyeq$', is often used in the literature. I have preferred the simpler more economical choice. As long as we can read from the context the intended relation, we can help ourselves to '$\leqslant$'. The same considerations apply to the use of '$+$' below.

[5] Partial acts are a technical device, and are used under the assumption that the state is in the act's domain. For more on partial acts and their comparison with other related kinds see [11].

[6] Also other obvious equalities hold, for example: if $C \subseteq B$ then $(f|B + h|B')|C = f|C$.

[7] In Savage's book this is denoted as $\sim B$.

Due to obvious symmetry, we can replace '$\Rightarrow$' by '$\Leftrightarrow$' at no additional cost. Assuming (P2), Savage defines:

(SD) $f \leqslant g$, *given* $B \Leftrightarrow f^* \leqslant g^*$, for some f^*, g^* that respectively agree with f and g on B, and with each other on $\overline{B}$.

P2 guarantees that the inequality $f^* \leqslant g^*$ does not depend on the choice of f^* and g^*. (The existence of such f^* and g^* is guaranteed via cut-and-paste; e.g., $f^* = f$, $g^* = g|B + f|\overline{B}$.) The argument for P2 is based on heuristic considerations, which are plausible but far less compelling than the businessman story, which motivates the sure thing principle. I shall later argue that P2 is actually stronger than what the businessman story implies.

Savage's explanation for not including a formalization of STP in his system is:

> The sure-thing principle cannot appropriately be accepted as a postulate in the sense that Pl is, because it would introduce new undefined technical terms referring to knowledge and possibility that would render it mathematically useless without still more postulates governing these terms. It will be preferable to regard the principle as a loose one that suggests certain formal postulates well articulated with Pl.
>
> [10, p. 22]

This is inaccurate, for it makes it appear as if STP has to do with counterfactual knowledge. It does not.[8] STP is based on straightforward suppositional reasoning in which one draws conclusion from a given assumption (or hypothesis). For example, in Savage's story, the point regarding a Republican win can be put thus: "Assuming that the Republican candidate will win, it is preferable to buy the property". This of course is the agent's subjective judgment (the setup as a whole describes a subjective view under rationality constraints). There is no need for the businessman to ask "whether he would buy if he knew that the Republican candidate were going to win"; he can simply ask whether, assuming that the Republican candidate will win, buying the property is preferable. The details of the reasoning in Savage's scenario can, for example, be filled thus: "If the Republican candidate wins, there will be fewer government restrictions on real estate developers, the property will be more in demand and its value will increase, hence buying it now is preferable. On the other hand, if the Democrat candidate wins, the government will invest in the region's infrastructure, which will raise the value of the property; hence also in that case buying it now is preferable." There is no appeal to knowledge.[9] (Other forms of STP that do involve such an appeal will be discussed later.)

Knowledge may however enter into the choice of act; because knowing, or not knowing, may itself be a given assumption. Suppose that one of two sealed envelopes, e_1 and e_2, contains a large sum of money and the other is empty. Jane can choose one of the two and, since she has no clue where the money is, she decides by tossing a coin.[10] Now, on each of the assumptions, that e_1 contains the money and that e_2 contains the money, the coin-tossing method is not the preferable act. But given Jane's state of knowledge, choice via coin-tossing is not irrational. The example also shows that one should be careful how to apply STP. The rejection of an act, under each of two mutually exclusive and exhaustive assumptions, is not sufficient for its rejection. What is needed is an agreement on what one *can rationally do* in each of the cases. And here there is no such agreement: on one assumption, Jane should reject the coin-tossing method and choose e_1, and on the other assumption Jane should reject the coin-tossing method and choose e_2.

Savage motivates P2 by the need to express within the system "the idea that **f** would be preferred to **g**, if B were known to obtain". But what we have to express is rather the idea that under the assumption B, f is preferable to g. Savage then suggests the technique of modifying two acts so as to agree with one another outside B, leaving them unchanged over B; it is supposed to reveal what the preference would be "if B were known". Again, the language is misleading, but his formal terminology "$\mathbf{f} \leqslant \mathbf{g}$, given B" (p. 22) is right on the mark, and I shall adopt it, for my own purpose, in the form:

$$(f \leqslant g)|B$$

Let us take this as a primitive for expressing *conditional ordering*: g is weakly preferable to f, under the assumption B (and here we do not appeal to P2). I shall also refer to $(f \leqslant g)|B$ as the *relativization* of $f \leqslant g$ to B. The conditionalization on B determines an ordering, $\leqslant_B$, of the partial acts defined over B:

$$f|B \leqslant_B g|B \Leftrightarrow_{\mathrm{Df}} (f \leqslant g)|B$$

For convenience I shall sometimes rewrite '$f|B$' as $f_{B,}$, so that the left-hand side of the last definition becomes: $f_B \leqslant_B g_B$, or, with our notational convention: $f_B \leqslant g_B$.

Taking an additional step, let $\leqslant$ be the union of all $\leqslant_B$'s. It is a relation extending our previous $\leqslant$, to partial acts, but it relates partial acts only if they have the same domain. Our previous $\leqslant$ is now $\leqslant_S$ (where S is the set of all states).

8 Savage tends in general to phrase things in terms of possible knowledge, when in fact only hypothetical reasoning is involved. In his discussion of constant acts and state-independent outcomes, on p. 27, he observes that "knowledge of an event cannot establish a new preference among consequences or reverse an old one". Actually, it is the *occurrence* of an event that cannot change the preferences with respect to state-independent outcomes.

9 It has of course have implications for knowledge: if one accepts that, under the assumption B, g is preferred to f, then if one *knew* that B one should prefer g to f, provided that no other relevant piece of information is known. But this is a different matter.

10 This scenario is easily incorporated into the Savage system by including the outcome of the coin-toss as a parameter that is part of the state of the world – a parameter whose value is known to Jane.

Similarly we get the strong preference relation for partial acts with the same domain, expressed as: $f|B < g|B$, as well as the indifference relation: $f|B \equiv g|B$. Of course, conditionalization need not preserve truth-value: $f \leqslant g$, does *not*, in general, imply $f|B \leqslant g|B$.

Let a *total-act* system, be one that, like Savage's system, is based on a set, Act, of total acts. Let a *partial-act* system be one that is based directly on the corresponding set, Act*, of partial acts.

The extension of Savage's system to a partial-act system requires of course a suitable extension of the axioms, in order to cover partial acts as well. This is rather a straightforward matter: any postulate P, of the kind considered by Savage, has a natural generalization to a corresponding partial-act axiom, P*. It consists of all relativizations of P to D, where 'D' ranges over the non-empty events; in each relativization, D plays the role of S. Formally: rewrite every 'f', 'g' ... in P as 'f_S', 'g_S', ... and replace the subscript 'S' by 'D' (all events that are mentioned in the axiom are now sub-events of D). In particular, P1* states that $\leqslant_D$ is a total preorder; i.e., for each non-empty D and each f and g, either $f|D \leqslant g|D$ or $g|D \leqslant f|D$.

We shall *not* however generalize P2 in this way, because one of the goals of setting up an Act*-based system is to state STP directly, in terms of partial acts, and avoid the roundabout way that Savage adopts. As we shall see, this leads to a finer analysis of STP, and it brings to the fore implicit assumptions that come with the use of P2. Consider the following most obvious cases of STP:

(STP$_{<,<}$) $f|B < g|B$ and $f|\overline{B} < g|\overline{B} \Rightarrow f < g$
(STP$_{\leqslant,\leqslant}$) $f|B \leqslant g|B$ and $f|\overline{B} \leqslant g|\overline{B} \Rightarrow f \leqslant g$

In the Act*-based system, each comes with its relativizations to non-empty events. The adoption of (STP$_{<,<}$) means that whenever $D = B' \sqcup B''$, where '$\sqcup$' means disjoint union, we also have:

$$f|B' < g|B' \quad \text{and} \quad f|B'' < g|B'' \quad \Rightarrow \quad f|D < g|D$$

and similarly for (STP$_{\leqslant,\leqslant}$). But these relativizations can be ignored as far as the analysis below is concerned.

In the following I shall compare different sets of postulates of the Act*-based system with each other, and also with P2. P1* is thus assumed. But I shall not consider the later postulates of the Savage system.[11] Claims of non-implications are provable by considering a space based on two states s_1, s_2, and four acts: f, g, $f|s_1 + g|s_2$, $f|s_2 + g|s_1$; the proofs are left to the reader. Note that the definition (SD) presupposes P2. But in a partial-act based system (SD) has the status of an axiom; such a system can satisfy P2 but violate (SD).

None of the above two forms implies the other. Obviously (STP$_{\leqslant,\leqslant}$) implies, in general (i.e., also for partial orderings) the following:

(STP$_{\equiv,\equiv}$) $f|B \equiv g|B$ and $f|\overline{B} \equiv g|\overline{B} \Rightarrow f \equiv g$

The conjunction (STP$_{<,<}$) & (STP$_{\equiv,\equiv}$) does not imply (STP$_{\leqslant,\leqslant}$), even under total ordering. Both of course hold in Savage's system; that is, they are implied by P2 & (SD). It would seem natural to adopt both (STP$_{<,<}$) and (STP$_{\leqslant,\leqslant}$). But this would still fall short of P2 (in one counterexample P2 does not hold, in another P2 holds but (SD) fails). Savage's way amounts to adopting also the following version of STP[12]:

(STP$_{<,\equiv}$) $f|B < g|B$ and $f|\overline{B} \equiv g|\overline{B} \Rightarrow f < g$

Under the total ordering requirement, (STP$_{<,<}$), (STP$_{\leqslant,\leqslant}$), and (STP$_{<,\equiv}$) together imply P2 as well as (SD). Moreover, (STP$_{<,\equiv}$) and (STP$_{\equiv,\equiv}$) are sufficient because they imply both P2 and (SD). Hence (STP$_{<,\equiv}$) is the additional and crucial postulate that is needed in order to get P2 and to justify Savage's definition of $f \leqslant g$, *given* B.

Although (STP$_{<,\equiv}$) seems quite plausible, the kind of argument exemplified in the businessman story, which justifies the other postulates, does not extend to it. The left-hand side of (STP$_{<,\equiv}$) implies that, in each of the cases, $f \leqslant g$. But in order to get $f < g$ the following consideration is needed: If the possibility that B occurs cannot be ignored, then there is a real possibility that g is better; hence altogether $f < g$. But what if B is ignorable? This possibility is ruled out if we agree that over an ignorable event any two acts are equivalent; this would imply that $f|B < g|B$ only if B is not ignorable. The argument for (STP$_{<,\equiv}$), with its appeal to the notion of "ignorable events" is weaker than the arguments for the other postulates. This shows what is being smuggled by the use of P2.

Ignorable events are introduced by Savage under the name *null events* (p. 24). Eventually, when null events turn out to be the events of probability 0, all this fits nicely together. But this requires additional postulates.

[11] In [10] definitions and postulates are interlaced. Definitions are introduced on the basis of postulates; when later postulates are phrased in terms of these definition, they make sense only if the previous postulates are assumed. A fuller comparison that extends beyond P1 and P2 requires a more detailed analysis of possibilities beyond the scope of this paper.

[12] Actually it requires a somewhat weaker version, in which '$f|\overline{B} \equiv g|\overline{B}$' is replaced with strict identity '$f|\overline{B} = g|\overline{B}$'. But this difference can be ignored since the weaker version implies (STP$_{<,\equiv}$).

2. The sure thing principle in the context of partial orderings

The generalization of Savage's system to the case of partially ordered acts (by which I mean preorders) is well motivated by the difficulty encountered when one tries to impose a total order, in situations in which the likeliness of the events cannot be assessed on the basis of robust statistical phenomena. One's uncertainty about one's subjective probability gives rise to uncertainty about the ordering of acts. Another source of uncertainty is the uncertainty about the values of the outcomes. In this paper I shall be concerned only with the difficulty of assessing the probabilities. I shall assume that all outcomes are real numbers representing money payoffs, in a range in which utility is linear in money. The difficulties of assigning subjective probabilities are well-known, and a subject of a huge literature; I need not go into them here.

Savage himself considered, in a short paragraph (p. 21), the "temptation to explore the possibilities of analyzing preference among acts as a partial ordering", one that admits incomparable acts and that "would seem to give expressions to sensations of indecision or vacillation, which we may be reluctant to identify with indifference". This he thinks is unpromising ("a blind alley"), but he adds that "only an enthusiastic exploration could shed real light on the question".

During the last 50 years there have been works and research programs leading, in this way or other, to partial orderings of acts. The above mentioned difficulty of assigning subjective probabilities motivated proposals to replace a single real-valued probability, by an interval-valued probability, or by a set of probabilities. If each probability determines a total ordering of acts, then a set of probabilities leads to a partial ordering in which $f \leqslant g$ iff the inequality holds under every probability in the set. Early works [7,15,5,2], were followed by [8] – a forerunner of what later became the *imprecise probabilities* framework – in which the credal state of an agent is represented by a set of probability distributions [17]. The distinguishing mark of Levi's proposal in [8] is that he does not provide an account of what the probabilities in the set are, but takes the whole set as a primitive whose interpretation is determined implicitly by its role as a guide for choosing acts. This feature constitutes, I think, a weakness of the approach, since an independent interpretation of the probabilities can be quite important. On the other hand, the focus on choosing acts endows the framework directly with operational meaning.

The works just mentioned start with sets or probabilities and use them to derive preferences for acts. Savage's approach goes in the opposite direction: from totally ordered preferences to probabilities (and utilities), via a representation theorem. Results of the last kind, from partial ordering to sets of probabilities (and utilities) have been obtained in [13]; it yielded a representation theorem in terms of sets of probability functions and sets of utilities. (A later generalization of this kind, which is based on choice functions, is given in [14].) These generalizations, however, concern dualistic systems of the Anscombe–Aumann type [1], which build on [16]. Their system is dualistic since it presupposes both objective probabilities and a subjective preference relation between acts; in [1] these two basic elements are represented by two kinds of lotteries, "roulette lotteries" and "horse lotteries". Savage's system, on the other hand, can be characterized as purely subjectivist; it does not have an objective-probability component. There are some well-known parallels between the two approaches; in particular, Savage's P2 corresponds to what later came to be known as the "independence axiom" of [16]. Nonetheless the extension of Savage's system to the case of partial orderings is a separate task, which as far as I know, is still open. Here I shall only touch on some points concerning STP in the context of a partial ordering.

A partial ordering results if we include the possibility that both $f \nleqslant g$ and $g \nleqslant f$. In this case we say that f and g are *incomparable*, to be denoted here as: $f \bowtie g$. It represents the state of an agent who is unable to rank the two with regard to preference. Note that if $f \leqslant g$, then the acts are comparable; in that case there are two possibilities: either $g \nleqslant f$ – implying that $f < g$, or $g \leqslant f$ – implying that $f \equiv g$. The relation $\equiv$ is sometimes referred to as the *indifference* relation, a name that can mislead, since in this case it does not indicate lack of belief, but, on the contrary, belief that the two acts are equally good. Incomparability, however, means that the agent lacks grounds for deciding on the ordering. Indifference must be an equivalence relation, but incomparability is highly unrestricted; for example, we can have: $f \bowtie h$ and $g \bowtie h$ and $f < g$.

So far we have overlooked the possibility that the agent believes that $f \leqslant g$, but remains undetermined as to whether $f < g$ or not. We could have introduced another relation that corresponds to this case, at a cost of complicating considerably the setup. This fine tuning is not worth the cost – at least not at the present stage of investigation.

The STP reasoning in the businessman scenario derives only from the details relating to the acts of buying and not buying the property. It does not depend on whether the ordering that involves other acts is total or not. Therefore, in an Act*-based system ($\text{STP}_{<,<}$) and ($\text{STP}_{\leqslant,\leqslant}$) should be adopted. Also the plausibility of ($\text{STP}_{<,\equiv}$) remains the same. But the possibility of incomparable acts requires additional STP postulates

The relativization to an event, B, makes sense also for incomparability statements; $(f \bowtie g)|B$ means: $f \bowtie g$, *given B*. As usual, we write this also as $f|B \bowtie g|B$. Since $f \bowtie g \Leftrightarrow f \nleqslant g \;\&\; g \nleqslant f$, we can relativize the right-hand side in order to define $(f \bowtie g)|B$. Moreover, hypothetical reasoning applies also here; the agent can infer that, in the case of B, there are no grounds for ordering the acts. Incomparability is not to be construed merely as a psychological reluctance to decide, but a rational judgment that in such and such circumstances there is no basis for comparing the acts.[13] Judgments of this type are a kin to the judgments that distinguish between risk and uncertainty. Consider a modification of Savage's example: The Republican candidate supports a reduction of real estate regulations; but he has also promised to develop and preserve national parks, and one of these parks is near the property in question. Assuming that the Republican candidate wins, it is impossible to predict, or even to assess, how these two factors will play out, and which factor is going to be more decisive

[13] To be sure this requires a somewhat idealized agent, but idealization is required in any case.

for determining the value of the property: his anti-regulation stance or the need to keep some promises. If it is the first, the value will increase (better possibilities for real estate developers), if it is the second it will decrease (restrictions due to the proximity to the park). Hence, assuming that the Republican candidate wins, the two acts, buying and not buying, are incomparable.

Note that the framework is meant to be applied in all hypothetical scenarios in which an agent has to choose an act belonging to an arbitrary finite set. If "doing nothing" is a possibility, then it must be included among the options. Suppose that, being offered some bets I reject all: I do not bet. This means that I choose the vacuous bet, whose payoff is 0 on each outcome, and it should be included among the options. But the do-nothing act need not be included, if we assume a situation that rules it out. Thus, in Savage's omelet making example, the agent is supposed to finish making an omelet that was began by his wife, and he has to choose one of three ways of handling an unopened egg that might be bad. Abandoning the omelet-making task is not an option. Or, if for sake of a more "realistic" picture we want to include it, then we can give the corresponding outcome (a big quarrel with the wife) a negative utility that will practically rule it out. The theory should therefore determine, given any finite set of acts, which act can be rationally chosen. (The choice need not be unique, even under a total ordering, given that two acts can be equivalent.) The theory thus determines a subset of the given options, those acts that *can* be rationally chosen. Let us call these acts *acceptable*.

If we apply STP in this context, then, assuming some set of acts out of which the agent has to choose, we have, for all non-empty B:

($\mathrm{STP_A}$) If f is acceptable, given B, and f is acceptable given $\overline{B}$, then f is acceptable.

Here 'A' is a mnemonic for 'Acceptable' The relativized form of ($\mathrm{STP_A}$) is:

($\mathrm{STP^{\#}{}_A}$) If $\{B_1, B_2\}$ is a partition of D, and if, for each $i = 1, 2$, f is acceptable given B_i, then f is acceptable given D.

It is easily seen that ($\mathrm{STP^{\#}{}_A}$) implies its own generalization: If $B_1, \ldots, B_k$ is a partition of D and, for each $i = 1, \ldots, k$, f is acceptable given B_i, then f is acceptable given D.

Consider now a choice from $\{f, g\}$, where $f \bowtie g$; the theory is supposed to determine which act is acceptable and at least one act must be, since these are all the options. If only one is, then the acceptable act would be strongly preferable to the other. Hence both must be acceptable. Therefore ($\mathrm{STP_A}$) implies that if $f|B \bowtie g|B$ and $f|\overline{B} \equiv g|\overline{B}$ then both f and g are acceptable; this rules out $f < g$ as well as $g < f$. Hence, either $f \bowtie g$ or $f \equiv g$. If there is a non-ignorable possibility that $f \bowtie g$, then the same line of thought that leads to ($\mathrm{STP}_{<,\equiv}$) can lead us to:

($\mathrm{STP}_{\bowtie,\equiv}$) $f|B \bowtie g|B$ and $f|\overline{B} \equiv g|\overline{B} \Rightarrow f \bowtie g$

Note that ($\mathrm{STP_A}$) is a sure thing principle of a different kind, than the principles of the previous section. It does not rely on straightforward suppositional reasoning about the facts, but also on the "fact" that we have a method which prescribes what we can to do in various hypothetical situations. The reasoning underlying ($\mathrm{STP_A}$) can be summed up thus: *Given that, either in case B or in case $\overline{B}$, our method classifies f as acceptable, we might as well classify f as acceptable* (*without waiting to find out whether B or $\overline{B}$*). A more detailed discussion of this kind of STP is postponed to the next section.

In an Act-based system postulate P2 makes sense for partially ordered, as well as for totally ordered acts. Adopting it we can define $f|B \leqslant g|B$ via (SD) as before. The other relativized cases follow by expressing $<$ and $\bowtie$ in terms of $\leqslant$; in particular: $f|B \bowtie g|B \Leftrightarrow_{\mathrm{Df}} f|B \nleqslant g|B$ & $g|B \nleqslant f|B$. The beauty of this approach is that we can adopt P2 for partial orders without any change. A system based on ($\mathrm{STP}_{<,<}$), ($\mathrm{STP}_{\leqslant,\leqslant}$),($\mathrm{STP}_{<,\equiv}$) and ($\mathrm{STP}_{\bowtie,\equiv}$) is equivalent to the system based on P2 and (SD). The equivalence holds in the same sense that it holds for totally ordered acts, as explained in the previous section.

All this requires a slight emendation. We should add the following postulate in order to prevent a pathological counterexample:

($\mathrm{STP}_{\bowtie,\bowtie}$) $f|B \bowtie g|B$ and $f|\overline{B} \bowtie g|\overline{B} \Rightarrow f \bowtie g$

Obviously, it should hold. But it is not implied by the above postulates, hence also not by P2 & (SD).

There is another version of P2, call it $\mathrm{P2}_<$, which is formulated by using $<$, instead of $\leqslant$, as a basic relation. Under total orderings each of $<$ and $\leqslant$ is definable in terms of the other and the two versions of P2 are equivalent. But under partial orderings $\mathrm{P2}_<$ is considerably weaker. This has to do with the fact that, in partial orderings, $<$ is definable in terms of $\leqslant$ but not vice versa. The weaker version is used in [13]. Prima facie I find it hard to accept P2 for total ordering and at the same time reject it for partial ordering and replace it with the weaker $\mathrm{P2}_<$. It seems that the intuitive arguments for $\mathrm{P2}_<$ extend to P2. But $\mathrm{P2}_<$ has been used in the context of dualistic systems, where the situation might be different.

3. Imprecise probabilities and the paradox of dilation

The term 'imprecise probabilities' has come to denote a framework in which credal states are represented by a family of probability distributions, defined over some Boolean algebra, rather than by a single probability – as it is done in a Bayesian setup. The name is somewhat misleading. The better term is, as Levi suggested, *indeterminate* probabilities[14]; but since the name has been adopted by researchers and forms the title of the book dedicated to the subject [17] I shall go along with it.

I shall be concerned with a particular phenomenon, known as *dilation*, which arises in the framework of imprecise probabilities. The technical aspects have been thoroughly analyzed in [12]. Here I shall be concerned with a paradigmatic very simple case. First, some notation: let $\mathcal{B}$ be a Boolean algebra of events (or propositions), construed as subsets of the set S.[15] As in the previous sections, 'A', 'B' 'C' 'D' range over $\mathcal{B}$. We let 'P', 'P_1', 'P'', ... etc., range over probabilities defined over $\mathcal{B}$. If $\mathcal{P}$ is a family of probabilities let $\mathcal{P}(A) =_{\mathrm{Df}} \{P(A)\colon\ P \in \mathcal{P}\}$. If $P(B) > 0$, then $P_B = P|B =$ the conditional probability P, given B (defined as usual). In the imprecise probabilities framework, conditionalization is generalized to families of probabilities by putting:

$$\mathcal{P}_B = \mathcal{P}|B =_{\mathrm{Df}} \left\{P_B\colon\ P \in \mathcal{P} \text{ and } P(B) > 0\right\}$$

A *gamble* is an act, f, defined over a partition $\{D_1, D_2, \ldots, D_k\}$, $D_i \in \mathcal{B}$, such that $f(D_i)$ is a real number representing money. To do f is to accept the gamble: the agent's payoff, if D_i occurs, is $f(D_i)$. A *bet* over A is a gamble over $\{A, \overline{A}\}|$. A bet *on A at odds $x : y$* is a bet in which $f(A) \geqslant 0$, $f(\overline{A}) \leqslant 0$ and $|f(A)|/|f(\overline{A})| = x : y$. (If $f(\overline{A}) = 0$, and $f(A) > 0$ then the odds are ∞. If $f(A) = f(\overline{A}) = 0$ this is the vacuous bet and the odds are 1:1.)

If $\mathcal{B}' \subseteq \mathcal{B}$ is a subalgebra of $\mathcal{B}$, then $P|\mathcal{B}'$ is the *restriction* of the probability P to $\mathcal{B}'$. If $\{D_1, \ldots, D_k\}$ is a partition, Then $P \,||\, \{D_1, \ldots, D_k\}$ is the restriction of P to the subalgebra generated by $D_1, \ldots, D_k$. This probability is uniquely determined by the vector $\langle P(D_1), \ldots, P(D_k)\rangle$; vice versa any vector of non-negative reals whose sum is 1 determines a probability. Put:

$$\mathcal{P} \,||\, \{D_1, \ldots, D_k\} =_{\mathrm{Df}} \left\{P \,||\, \{D_1, \ldots, D_k\}\colon\ P \in \mathcal{P}\right\}$$

For a partition $\{A, \overline{A}\}$, this boils to down to the set $\mathcal{P}(A)$.

An imprecise probability (i.e., family of probabilities) is supposed to provide guidance in the choice of acts. This is its operational meaning and it usually comes with a *locality assumption*:

> The recommendation for acts over a partition $\{D_1, \ldots, D_k\}$ depend only on $\mathcal{P} \,||\, \{D_1, \ldots, D_k\}$. In particular, the recommended bets over A depend only on $\mathcal{P}(A)$.

The assumption is obvious, but not always justified. It is obvious, given the goal of providing a systematic application of probabilities that is based on the explicit relevant factors of the problems in question. Other factors may involve additional elements, such as utilities. But the contribution of the probabilistic element should amount to the probabilities over the algebra $\mathcal{B}$ generated by the events that figure in the statement of the problem. If additional events are relevant then they should be included in the algebra and their contribution made explicit. This holds if the system is based on a single subjective probability function, and we expect it to hold if we consider a family of probabilities. Sometimes it does not. Dilations are cases in which an additional event not figuring in the statement of a problem can be included, it then "does some work" and disappears, leaving us with a counter-intuitive result. I shall use the following as the a paradigmatic example.

3.1. Basic example

Consider a Boolean algebra containing two events A and B. The subalgebra generated by them is the algebra generated by the partition $\{A \cap B,\ A \cap \overline{B},\ \overline{A} \cap B,\ \overline{A} \cap \overline{B}\}$. Let $\mathcal{B}$ be this subalgebra. As usual $[0, 1] = \{x\colon\ 0 \leqslant x \leqslant 1\}$. For $\lambda \in [0, 1]$, let P_λ be defined by:

$$P_\lambda(A \cap B) = \lambda \cdot \frac{1}{2} \qquad P_\lambda(A \cap \overline{B}) = (1 - \lambda) \cdot \frac{1}{2} \qquad P_\lambda(\overline{A} \cap B) = (1 - \lambda) \cdot \frac{1}{2} \qquad P_\lambda(\overline{A} \cap \overline{B}) = \lambda \cdot \frac{1}{2}$$

Each of the equalities is implied by the other three. P_λ is also characterized by the equalities:

$$P_\lambda(A) = P_\lambda(B) = \frac{1}{2}, \quad P_\lambda(A|\mathrm{B}) = \lambda, \quad \text{which imply } P_\lambda(A|\overline{B}) = 1 - \lambda$$

Now assume that $\mathcal{P} = \{P_\lambda\colon\ \lambda \in [0, 1]\}$. Then $\mathcal{P}(A) = \{\frac{1}{2}\}$ and $(\mathcal{P}|B)(A) = (\mathcal{P}|\overline{B})(A) = [0, 1]$.

Initially, without any information regarding B, the agent is in the happy situation of assigning to A the single sharp probability $\frac{1}{2}$. But, conditional on B, the probability has the maximal range of all possible values; and also conditional on $\overline{B}$ it has the maximal range. This is the simplest case of dilation.

An example that falls under the above scheme is given in [17, pp. 289–299], who uses a scenario based on objective probabilities. Walley draws from the story the following moral:

---

[14] 'Imprecise' means imprecision in the values of a function. But here we concerned with a family of functions, each of which is fully precise.

[15] I use 'S' for reason of continuity with the previous sections. Usually 'Ω' is used in this role.

> This shows that receiving extra information can sometimes be a bad thing, in the sense that it is certain to produce indeterminacy and indecision.

He goes on to observe that the effect is quite common when artificial randomization is involved and when knowledge of the randomly drawn sample may have a bad effect. Since this is a problem regarding objective probabilities it will be handled in the next section, where it will turn out that the relevance of the new information is ambiguous and that the agent has latitudes in this matter. For the moment, I only observe that, in general, excessive information can have the bad effect of "cloging the machine"; we cannot process it efficiently, or even process it at all. To decide what is relevant and what is not and to pursue the enormous amount of suggested possibilities is waste of time and may be practically impossible.[16]

In [12], where the term *dilation* is introduced, Seidenfeld and Wasserman illustrate the problem of dilation thus:

> To emphasize the counterintuitive nature of dilation, imagine that a physician tells you that you have probability $\frac{1}{2}$ that you have a fatal disease. He then informs you that he will carry out a blood test tomorrow. Regardless of the outcome of the test, if he conditions on the new evidence, he will then have lower probability 0 and upper probability 1 that you have the disease. Should you allow the test to be performed? Is it rational to pay a fee not to perform the test?
>
> [12, p. 1140]

The underlying assumption of such scenarios is that, upon getting to know that B, one updates one's credal probability – or, in our case, the family $\mathcal{P}$ – by conditionalizing on B. As a principle of rationality, this requires some qualifications. The obvious assumption is, of course, that B constitutes all the new information. Yet, in practice, obtaining new information always involves some additional items (say, receiving email), which are ignored because they are considered irrelevant. Furthermore, the new information might suggest to the agent some new possibilities not envisaged before. These obstacles always arise when a theoretical framework is applied in practice. The practical applications of a theory are subject to *ceteris paribus* clauses, "all other things being equal," "in normal circumstances", and so on. The best way to handle practical updating by conditionalization is to prescribe it as the *default procedure*: the agent is supposed to update, by conditionalizing on the new information; failures to do so require explanation and the onus of explaining is on the agent.

Now in the case of the Basic Example it is difficult to see what the explanation of not conditionalizing could be; everything is clear and we know ahead what the updating yields. So the only recourse of those who accept dilation as a fact of life, but reject its counterintuitive consequences, is refusal to being informed. You know that either B or $\overline{B}$ has occurred, you know to the last detail the effect of the conditionalization, and that it will be the same in each of the alternatives. But in order to avoid conditionalizing you refuse to be given the information.

Let us recast the situation in terms of rational choice. The imprecise probability (the family of probability functions) decides your choice of acts. In general, the theory you adopt will prescribe which of your options are acceptable and which should be rejected. Let $\mathbf{Act}_{\mathcal{P}}(X)$ be the prescription of some theory, or set of rules, where $\mathcal{P}$ is your imprecise probability and X is the set of your options. The function $X \mapsto \mathbf{Act}_{\mathcal{P}}(X)$ gives operational meaning to $\mathcal{P}$. Consider the Basic Example, and let X range over sets of bets over A – the options you can choose from. The sure thing principle can be applied in the following form. Either B is the case or $\overline{B}$ is the case; in each case this is a given fact, independent of what you know and do. Let $\mathcal{P}' = \mathcal{P}|B$, $\mathcal{P}'' = \mathcal{P}|\overline{B}$. Then, if you knew that B is the case, your rational choice would have been determined by $\mathbf{Act}'_{\mathcal{P}}(X)$. If you knew that $\overline{B}$ is the case, your rational choice would have been determined by $\mathbf{Act}''_{\mathcal{P}}(X)$. You know that $\mathcal{P}''(A) = \mathcal{P}'(A)$, therefore your rational choice should be determined by $\mathbf{Act}'_{\mathcal{P}}(X)$. This can be summarized by:

$$(\text{STP}^*) \quad (\mathcal{P}|B)(A) = \mathcal{P}^*(A) \text{ and } (\mathcal{P}|\overline{B})(A) = \mathcal{P}^*(A) \Rightarrow \mathbf{Act}_{\mathcal{P}}(X) = \mathbf{Act}_{\mathcal{P}^*}(X)$$

(STP*) is based on counterfactual reasoning about what one should do if one knew that such and such is the case. As argued in Section 1, it is of a different type than the STP used by Savage. Yet it is quite convincing: if what you should do if you knew that B is the case is the same as what you should do if you knew that B is not the case, then go ahead and do it. Resistance to this type of reasoning sometime appeal to the difficulty of evaluating "what if I knew that...". But here everything is as explicit and clear as it can be, open and shut. The trouble is that the outcome of applying (STP*) in the case of our paradigmatic example is extremely counterintuitive. Suppose you are offered a bet on A at odds 6:5 (you win \$60 if A, you lose \$50 if $\overline{A}$), you assign to A probability ½. Your adopted theory reflects your risk aversion, but since the betting odds are very good your theory recommends that you should definitely accept. But (STP*) mandates that you should treat the situation as one in which your imprecise probability is $[0, 1]$. In *that* case your theory recommends that you reject the bet, since the expectation is negative whenever A's probability is $< 5/11$.[17] It is moreover far from clear why B should be even considered relevant in this situation.

If one thinks that one's choices should be determined by $\mathcal{P}(A) = \{\frac{1}{2}\}$, one can maintain the position by refusing to obtain any information concerning B. But this, I think, is a council of despair. Imagine, you are willing to accept a bet on A,

[16] The problem of limited resources cannot be ignored, and I addressed it in [3].

[17] On Levi's proposal [8] both bets are admissible, but the vacuous bet is to be preferred because it guarantees a smaller loss. The difference between $\{\frac{1}{2}\}$ and $[0, 1]$ is so great that, for every plausible theory, there is a set of options that gives rise to different prescriptions.

thinking that it is the best option, and you hold a folded piece of paper on which either 'B' or '$\overline{B}$' is written. You know that if you take a look, then, no matter what you see, it will imply that for rational reasons you should reject the bet. So you refuse to look in order to avoid the conflict. I think this is incoherent, or, if you want, it reveals a fault line in the theory.

Note that dilation by itself does not lead to paradox, if we think that the dilated values are a better guide to choosing acts. The paradox arises only when (i) we judge the sharp value to be better, but (ii) our theory does not provide any means not to dilate, once we get the information. This is a shortcoming of the theory, which turns the conditionalization into an automatic irrational act; we refuse the information in order to avoid a bad knee-jerk reaction, over which we have no control.

Refusal to be informed in the case of dilation should be sharply distinguished from refusals of different kinds, which are legitimate and well understood. Such is the refusal to be told the solution of a riddle you are trying to solve, or to be told the surprising end of a thriller, because that will spoil the movie for you, or the refusal to get depressing news in situations in which you can do little about,[18] or the wish not to be informed in order to have a deniability option, or because you do not trust yourself not to reveal a secret, or because to know what some colleagues said about you might harm your ability to maintain formally correct relations with them, and so on and so forth.[19]

Refusal of getting information in the case of dilation should be also distinguished from refusal of getting information that is likely to be biased; for all you know the information is true, but it has been cherry picked by someone who is out to prove a particular claim, and you doubt that you can correct for the bias. Finally, the case of dilation should be distinguished from the case where the information is refused because it is "noisy" or distractive. The information might be even relevant, but its relevance is unclear and it may send you on endless inquiries. This is the reason that even a Bayesian may refuse additional information concerning the sample obtained in a statistical test.[20] In the case of dilations the information is as unbiased as it can be and it has a sharp clear effect: If you knew that B is true, you (as a rational agent) would choose such and such an act, and if you knew that B is false, you (as a rational agent) would choose this same act.

The discussion so far indicates that imprecise probabilities that enable dilations should be ruled out on grounds of incoherence. The results of [12] show that this might constitute a severe restriction, but it will still leave us a sizable interesting family of imprecise probabilities.

Imprecise probabilities with dilations can however be obtained by using objective probabilities, that is, probabilities defined in terms of random physical phenomena: drawings from urns, lotteries, and such. These cases are, so to speak, forced upon us and they cannot be ruled out on grounds of incoherence. To be sure, in order to generate a family containing more than a single probability, there should be a parameter, whose possible values are not probabilistically distributed. But this can be plausibly assumed; for example, consider drawing from an urn containing black-or-white balls in unknown proportion of color. A Bayesian might impose the use of a single prior distribution in all cases. But the distinction between risk and uncertainty is not on the side of Bayesians of this stripe. Here is a protocol that realizes the Basic Example, by means of two coins, where 1 and 0, serve as "Heads" and "Tails". Coin_1 is a fair coin; which is used to define the event A. Coin_2 yields 1 with probability λ, 0 – with probability $1-\lambda$.

Protocol 1A.

(1) Toss Coin_1, let $X =$ outcome of the toss, and let $A \Leftrightarrow_{\mathrm{Df}} X = 1$
(2) Toss Coin_2, let $Z =$ outcome of the toss, and let $C \Leftrightarrow_{\mathrm{Df}} Z = 1$
(3) Put $Y = X \cdot Z + (1-X) \cdot (1-Z)$ and let $B \Leftrightarrow_{\mathrm{Df}} Y = 1$

Note that (3) is equivalent to putting: $B = A \cap C \cup \overline{A} \cap \overline{C}$, or, equivalently, $B \Leftrightarrow X = Z$. The assumption is that λ can be any number in $[0,1]$ and the agent does not employ any subjective probability distribution for λ. The use of real numbers involves a certain idealization, but there is a discrete version of paradox, in which λ ranges over the fractions k/n, where n is a fixed even number > 0 and $k = 0, \ldots, n$. Instead of coin tossing we use drawings from urns; in each of the two urns there are n black-or-white balls. In Urn_1 there are $n/2$ black balls, in Urn_2 their number is k which can be any number from 0 to n.

Indeed, dilations would not have been such an interesting phenomenon, were it not for the fact that they can be easily generated by using lottery devices.

4. Objective probabilities

As their name implies objective probabilities are probabilities whose values have the status of objective facts. In particular, one can be wrong about them. You might think that the probability that the coin landed "Heads" is 0.5 ± 0.05 whereas

[18] It might not indicate strength of character, but it is understandable.

[19] In [6] the authors list cases where information has negative value. However they do not distinguish between the different kinds. The movie scenario is on their list but most of their cases are quite technical and I cannot go into them here.

[20] The points raised in this paragraph are of a pragmatic nature; they derive from the need to come to a decision using the limited resources we have. Each merits a much longer discussion than what I can afford within the bounds of this paper.

in fact, the coin is biased and the probability is 0.6 ± 0.05. These are probabilities that can be measured by performing experiments, and which can be confirmed or disconfirmed, by empirical data. The data reveals certain regular patterns, which, as a matter of brute fact, obtain in our world. The regularity can be temporal – appearing in sequences of repeated trials or continuous processes, or spatial – appearing in a distribution of certain phenomena over spatial regions. In physics the regularities can be displayed in distributions over phase spaces. In the sequel I shall be only concerned with regularities of sequences generated by repeated trials

We do apply objective probabilities to particular events, but these applications involve a presupposed background: an actual or possible pattern into which the events are embedded. The big attraction of subjective probabilities is that they are much simpler.

The assignment of objective probabilities involves judgments of relevance, which play an explicit role if the probabilities are imprecise. Consider the Basic Example of the previous section. If it is implemented via Protocol 1A (and this is all the information we have) the value of Z is completely irrelevant as far as the probability of A is concerned; for it is the tossing outcome of an unrelated coin. The toss of Coin_1, which determines the truth or falsity of A comes first, and then the truth or falsity of B is determined via A and Z. Hence the information regarding B is irrelevant (at the least this is a plausible view). If the conditional probability has to serve as a guide for action, then we might do better not to conditionalize. But the force of this observation remains the same after we are informed that B is the case. The agent can therefore opt for what I shall call *vacuous conditionalization*, to be denoted by $|^0$:

$$\big(\mathcal{P}|^0 B\big)(A) = \mathcal{P}(A) = \{0.5\}$$

Vacuous conditionalization is adopted when the usual formal conditionalization of the imprecise probability changes it, but the evidence is regarded irrelevant. All this is however a partial picture. We shall shortly see that, in the context of repetitive tossing, both options: vacuous and non-vacuous, can make sense.

Note that the problem does not arise in a Bayesian framework with a single probability function. In that case, the agent assigns a single value $P(C)$ to C and, since A and C are to be independent, $P(A \cap C) = P(A) \cdot P(C)$; this implies, by elementary algebra, that A and $A \cap C \cup \overline{A} \cap \overline{C}$ are independent, hence conditionalizing on B does not change $P(A)$. The single probability function does all this work for us. But in the imprecise probability framework, vacuous conditionalization is required.

The imprecise probability of Example 1 can however be generated by a different protocol, obtained from Protocol 1A by switching A and B:

Protocol 1B.

(1) Toss Coin_1, let $X =$ outcome of the toss, and let $B \Leftrightarrow_{\mathrm{Df}} X = 1$
(2) Toss Coin_2, let $Z =$ outcome of the toss, and let $C \Leftrightarrow_{\mathrm{Df}} Z = 1$
(3) Let $Y = X \cdot Z + (1 - X) \cdot (1 - Z)$ and let $A \Leftrightarrow_{\mathrm{Df}} Y = 1$

This is equivalent to putting: $A = B \cap C \cup \overline{B} \cap \overline{C}$, or equivalently: $A \Leftrightarrow X = Z$. There is a formal symmetry: $A = B \cap C \cup \overline{B} \cap \overline{C} \Leftrightarrow B = A \cap C \cup \overline{A} \cap \overline{C}$. But in Protocol 1B the causal story goes in the other direction: B has partial effect on A. Here the agent might opt for the standard conditionalization of imprecise probabilities, which I shall denote by $|^1$:

$$\big(\mathcal{P}|^1 B\big)(A) = \big\{(P_\lambda|B)(A)\colon \lambda \in [0,1]\big\} = [0,1]$$

Again, this is a partial picture. The complete story requires that we consider repeated trials. First, let me provide some further general clarifications concerning on objective probabilities and repeated trials.

The applications of probability theory to the repeated-trial setup involves an idealized model, which is based on the space of all infinite sequences of possible outcomes – conceived as possible results of a repeated trial that goes on indefinitely – and on a probability distribution, say P, defined over this space.[21] That an infinite sequence, $w = w_1, w_2, \ldots, w_i, \ldots$, accords with P means that it passes the statistical tests associated with the probability P.[22] In principle, every associated test, T, determines a subset, say U_{T}, of infinite sequences, such that $P(U_{\mathrm{T}}) = 0$; the sequence w passes

[21] The underlying probability is σ-additive, though the issue of finite versus σ-finite additivity can be bypassed as far as this paper is concerned. One can show that σ-additivity underlies the earliest applications of probability, going back at least to the works of Huygens, where it is an implicit but essential condition.

[22] This notion was first proposed in [9] for the case of Bernoulli trials and it was generalized in [4] to arbitrary probabilities and an arbitrary notion of "test". Martin-Löf's approach which drew directly on statistical practice, was based on the notion of passing a test at a certain significance level, which meant that the sequence did not belong to a certain measurable subset of small enough probability. The framework proposed in [4] applies to arbitrary probabilities and to any tests, which are defined in some powerful formal language. The generalization carries over to a space of possible worlds, where the worlds are models for a sufficiently expressive empirical language. Each test determines a set of probability 0 and a sequence passes it iff it belongs to its complement. The definition in terms of sets of probability 0, rather than the complements which have probability 1, is an inheritance from the statistical method that focuses on rejecting the null hypothesis.

the test iff $w \notin U_{\mathrm{T}}$.[23] So far this is a purely mathematical model. No trial is repeated an infinite number of times, and even if it were, we could not check whether the infinite outcome sequence passes the test, unless we were endowed with some infinite capacity.

The empirical content of a hypothesis that links the generated sequences with P comes from a testing method, which is based on finite "long enough" segments $w_1, w_2, \ldots, w_k$. The test T determines a set $U_{\mathrm{T},k}$ of small probability, $P(U_{\mathrm{T},k}) < \varepsilon$, and the sequence passes the test at significance level ε, if it is *not* in $U_{\mathrm{T},k}$. The empirical success of this methodology, in the areas in which it succeeds, endows statistical practice with its normative status. And it underlies the preference for making decisions under risk (when the chances are known) to making them under uncertainty.

In the most elementary cases, the numeric probability we assign to the outcome of a single trial can be identified with the frequency limit of such an outcome (which exists, with probability 1), hence, the so-called "frequency interpretation" of probability. The term is misleading because in less elementary cases the probability of an event cannot be so identified – unless we consider higher dimensional arrays and take great care to define the relevant sequence. It is also problematic when the objective probabilities are derived via some empirical theory and applied to events that are necessarily one-shot events, say the explosion of a given star (also, in everyday situations, we do want sometime to assign chances to the event of being killed). The so-called "propensity interpretation" avoids these problems, but has problems of its own; what is "propensity" a propensity of? It is not always clear. An assignment of objective probabilities is conceivable when an event involves many players, say the collision of two stars. But whose propensity is it: the two stars, the two stars with other material bodies in their "vicinity"? the galaxy? the cosmos? The propensity interpretation is a convenient metaphor, which emphasizes the objective brute-fact aspect of the situation. But in many cases it hides more than it reveals.

Coming back to Protocol 1A, consider its repetitive applications. For $i = 1, 2, \ldots$, let X_i and Z_i be, respectively, the outcomes of the ith toss of Coin_1 and Coin_2. Note that these are two sequences of *iid* (independent identically distributed) random variables, which are independent of each other. The two tosses of our previous procedure are now the first in an ongoing repetition described by the following protocol. Our previous A, B and C, are identical to A_1, B_1 and C_1.

Protocol 1A*. At time i:

(1) Toss Coin_1 and let $A_i \Leftrightarrow_{\mathrm{Df}} X_i = 1$
(2) Toss Coin_2 and let $C_i \Leftrightarrow_{\mathrm{Df}} Z_i = 1$
(3.i) Put: $Y_i =_{\mathrm{Df}} X_i \cdot Z_i + (1 - X_i) \cdot (1 - Z_i)$
(3.ii) Let $B_i \Leftrightarrow_{\mathrm{Df}} Y_i = 1$

We get $\mathcal{P}(A_i) = \{0.5\}$ $\mathcal{P}(B_i) = \{0.5\}$, $\mathcal{P}(C_i) = [0, 1]$ But the meaning of the conditionalization on B_1, $(\mathcal{P}|B_1)(A_1)$ is still undetermined. I have argued above for giving the agent the option to conditionalize vacuously. For repeated trials, it means that the process is left unchanged and at each time point we conditionalize vacuously. This gives us the sequence:

$$(\mathcal{P}|^0 B_i)(A_i) = \bigl(\mathcal{P}|^0 (Y_i = 1)\bigr)(X_i = 1) = \{0.5\}, \quad i = 1, 2, \ldots$$

Call this conditionalization *weak*. The non-vacuous conditionalization, which we shall call *strong*, is obtained by considering only the subsequence in which the random variable in question has the required value. In the case of Protocol 1A*, B_1 is the event $Y_1 = 1$, hence we have to consider the subsequence for which $Y_i = 1$. Let $I_1 = \{i\colon\ Y_i = 1\}$ and let $j(1), j(2), \ldots, j(k), \ldots$ be the enumeration of I_1 in increasing order. Then the new random variables should be $X'_i = X_{j(i)}$, $Z'_i = Z_{j(i)}$, $Y'_i = Y_{j(i)}$. This can be done, though it is less simple than it looks, since I_1 is itself a random variable (indeed, we appeal to the fact that $Y_i = 1$ holds infinitely often, with probability 1). An easier way of achieving the same effect is to use the following modified protocol, which outputs the required subsequence. The algorithm runs repeatedly and outputs only those values in which it "succeeds" getting $Y_i = 1$. This involves a loop. To avoid cluttered notation, I use the same symbols that I used above for the random variables and the events, although they are now defined by a different protocol, and are in fact different.

Protocol 1A#. At time i:

(1) Toss Coin_1 and let $A_i \Leftrightarrow_{\mathrm{Df}} X_i = 1$
(2) Toss Coin_2 and let $C_i \Leftrightarrow_{\mathrm{Df}} Z_i = 1$
(3.i) $Y_i =_{\mathrm{Df}} X_i \cdot Z_i + (1 - X_i) \cdot (1 - Z_i)$
(3.i#) if $Y_i = 1$ **go to** (3.ii), otherwise **go to** (1)
(3.ii) Let $B_i \Leftrightarrow_{\mathrm{Df}} Y_i = 1$

$(\mathcal{P}|^1 B_1)(A_1)$ is defined as $\mathcal{P}(A_1)$, *where* A_1 is *defined via* Protocol 1A#. For weak conditionalization we have: $\mathcal{P}(A_1) = (\mathcal{P}|^0 B_1)(A_1) = \{0.5\}$. But strong conditionalization means that we employ Protocol 1A#; in this case, the imprecise prob-

[23] More generally, the hypothesis in question need not be associated with a particular P, it can be a hypothesis about a property of P, e.g., that the repeated events are independent. The associated set U_{T} has probability 0 under each P that has the property.

ability of A_1 is $[0,1]$ right from the beginning, because this is its value once it is defined via the protocol. In either case there is no dilation. Another version of the protocol outputs only the cases in which $Y_i = Y_1$. It can be used to define strong conditionalization on B_1 as well as on $\overline{B}$.

In the same way, Protocol 1B gives rise to Protocol 1B* for repeated trials; it is obtained from 1A* by switching 'A_i' and 'B_i'. But the protocol that defines strong conditionalization can be extremely simple, reflecting the fact that here the causal arrow goes from B to A. To conditionalize on B_1 we have to consider all the cases in which $X_i = 1$, and this is done simply by setting the value of X_i to 1.

Protocol 1B#. Toss Coin$_2$ and let $C_i \Leftrightarrow_{\mathrm{Df}} Z_i = 1$
Put $Y_i = Z_i$ (the result of $X_i = 1$)
Let $A_i \Leftrightarrow_{\mathrm{Df}} Y_i = 1$

The events produced by this protocol satisfy $A_i = C_i$ and therefore the imprecise probability of A_1 is $[0,1]$.

My own intuition, which, in this case, derives from a causal picture, is that weak conditionalization fits better Protocol 1A* and strong conditionalization fits better Protocol 1B*. But the causal picture introduces an element that should be avoided, if we can derive the required conditionalization from the acts that are offered as options. Take for example the bet discussed in the previous section; here one should opt for weak conditionalization, no matter how the causality works. But suppose that your option is a called-of bet, which is conditional on B (the payoffs obtain, if B occurs; and are canceled, otherwise). In that case, risk aversion, combined with uncertainty avoidance, mandates that you should reject the bet. The recommendation is stronger in a long range scenario, where repeated trials are accompanied with repeated bets. Roughly, the unconditional repeated bet mentioned above promises a very high return with very little risk; but repetitions of the called-off bet might yield enormous gains but also enormous losses, depending on the unknown value of a parameter. The more repetitions, the stronger the recommendation is. Numeric details can be worked out via elementary methods of estimating the standard deviations of averages for Bernoulli trials.

I conclude with another way of implementing the Basic Example, which is symmetric and which does not involve a causal picture. It is the well-known case of two random variables with known distributions but unknown correlation, which arises from a double classification of some population. Consider an urn model, where the items are classified into P's and non-P's and into Q's and non-Q's. The proportion of P's and the proportion of Q's are known; but the proportion, λ, of P's among the Q's is unknown. If the population is large enough we can pretend that λ can be any number in $[0,1]$. Suppose that the proportion of P's is $\frac{1}{2}$ and the proportion Q's is $\frac{1}{2}$ and let A be the event that the drawn object is a P, and B – that the drawn object is a Q. This implements our Basic Example.

For the sake of testing everyday intuitions, let us consider an elementary case: A bag contains six balls, marked with the numbers $1,\dots,6$; classify the balls into *odd* and *even* according to their number. The balls are colored, half are white and half are black. Let λ be the proportion of odd balls among the black ones. It is known that at least one black ball is odd and at least one is even. This leaves two possibilities: (i) $\lambda = \frac{2}{3}$, two black balls are odd and one is even (ii) $\lambda = \frac{1}{3}$, one black ball is odd and two are even.[24] A ball is drawn at random. Let A be the event that it is odd, let B be the event that it is black. The imprecise probability is $\mathcal{P} = \{P_{\frac{1}{3}}, P_{\frac{2}{3}}\}$ and the following holds:

$$\mathcal{P}(A) = \mathcal{P}(B) = \left\{\tfrac{1}{2}\right\} \qquad (\mathcal{P}|B)(A) = (\mathcal{P}|\overline{B})(A) = \left\{\tfrac{1}{3}, \tfrac{2}{3}\right\}$$

Assume that this time you, the agent, are willing to bet according to the expected payoff, if the stake is not too high. So you stake \$10 on A at odds 1:1. A ball is drawn and, before you know whether it is odd or even, you are informed that it is black. If you conditional on the information you get $\{\frac{1}{3}, \frac{2}{3}\}$. Your bet, at even odds, now looks as if it were one of two bets: a very bad bet (on an event of probability $\frac{1}{3}$), or a very good one (on an event of probability $\frac{2}{3}$). You hate uncertainty and in such a two-bet situation your policy is not to bet. But are you in such a situation? If you are you should cancel the bet, if you can, and even agree to pay something for doing so. On my suggestion you should be given the rational option to go on with the bet, even if you can cancel, because you do not consider yourself as being in a two-bet situation. This is what weak conditionalization will let you do.

The protocol for repeated random drawing of balls (with replacements) is simple:

Protocol 2. At time i:

(0) Shake the bag
(1) Draw a ball from the bag and put:
(1.i) $X_i = 1$, if the ball is odd; $X_i = 0$, if the ball is even $A_i \Leftrightarrow_{\mathrm{Df}} X_i = 1$

[24] There are 3 possible distributions of black balls among the even ones which realize each of the cases, $\lambda = \frac{2}{3}$ and $\lambda = \frac{1}{3}$. For each of the cases $\lambda = 0$ and $\lambda = 1$, there is only one distribution. I have excluded the latter for the sake of complete symmetry.

(1.ii) $Y_i = 1$, if the ball is black; $Y_i = 0$, if the ball is white; $B_i \Leftrightarrow_{\mathrm{Df}} Y_i = 1$
(2) Replace the ball.

The events A and B of our single-shot scenario are now A_1 and B_1. Mathematically, the rest of the story is the same as the mathematical story in the previous protocols. Weak and strong conditionalization is defined as before. But the difference in the protocols – which is the difference in the physical implementations – can, as we saw, make a difference. Consider strong conditionalization in the case of Protocol 2. It means that we take into account only the subsequence of the draws in which the drawn ball is black. The following protocol does it:

Protocol 2$^{\#}$. At time i:

(0) Shake the bag
(1) Draw a ball from the bag and put:
(1.i) $X_i = 1$, if the ball is odd; $X_i = 0$, if the ball is even $A_i \Leftrightarrow_{\mathrm{Df}} X_i = 1$
(1.ii) $Y_i = 1$, if the ball is black; $Y_i = 0$, if the ball is white $B_i \Leftrightarrow_{\mathrm{Df}} Y_i = 1$
(1.ii$^{\#}$) If $Y_i = 0$ **go to** (1), otherwise **go to** (2)
(2) Replace the ball.

The upshot here is equivalent to what we get if we draw randomly from the black balls. Now consider the scenario of betting on A and being informed that the ball is black before knowing the outcome. Weak conditionalization means that you consider the drawing as the first in a sequence of random drawings from the six balls, in which the ball *happens to be black*. Strong conditionalization means that you consider the drawing to be the first in a sequence of random drawings from the black balls. Both are legitimate.

References

[1] F. Anscombe, R. Aumann, A definition of subjective probability, Ann. Math. Stat. (1963) 199–205.
[2] A.P. Dempster, Upper and lower probabilities induced by a multivalued mapping, Ann. Math. Stat. 38 (2) (1967) 325–339.
[3] H. Gaifman, Reasoning with limited resources and assigning probabilities to arithmetical statements, Synthese 140 (2004) 97–119.
[4] H. Gaifman, M. Snir, Probabilities over rich languages testing and randomness, J. Symb. Log. 47 (3) (1982) 495–548.
[5] I.J. Good, Subjective probability as the measure of a non-measurable set, in: Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress, Stanford University Press, 1962, pp. 319–329.
[6] J. Kadane, M. Schervish, T. Seidenfeld, Is ignorance bliss?, J. Philos. 105 (1) (2008) 5–36.
[7] B. Koopman, The bases of probability, Bull. Am. Math. Soc. 46 (10) (1940) 763–774.
[8] I. Levi, On indeterminate probabilities, J. Philos. 71 (13) (1974) 391–418.
[9] P. Martin-Löf, The definition of random sequences, Inf. Control 9 (1966) 602–619.
[10] L. Savage, The Foundations of Statistics, 2nd revised edition, Dover Publications, 1972, Originally published by John Willey and Sons, 1954.
[11] T. Seidenfeld, The independence postulate, hypothetical and called-off acts: a further reply to rabinowicz, Theory Decis. 48 (June 2000) 319–322.
[12] T. Seidenfeld, L. Wasserman, Dilation for sets of probabilities, Ann. Stat. 21 (3) (1993) 1139–1154.
[13] T. Seidenfeld, M. Schervish, J. Kadane, A representation of partially ordered preferences, Ann. Stat. 23 (6) (1995) 2168–2217.
[14] T. Seidenfeld, M. Schervish, J. Kadane, Coherent choice functions under uncertainty, in: The 5th International Symposium on Imprecise Probability and Applications, Prague, Czech Republic, 2007.
[15] C. Smith, Consistency in statistical inference and decision, J. R. Stat. Soc. B 23 (1) (1961) 1–37.
[16] J. von Neumann, O. Morgenstern, The Theory of Games and Economic Behavior, 3rd edition, John Wiley & Sons, Inc., 1964.
[17] P. Walley, Statistical Reasoning with Imprecise Probabilities, Chapman and Hall, New York, 1991.