Chapter 11
Stopping to Reflect

Mark J. Schervish, Teddy Seidenfeld, and Joseph B. Kadane

Our note is prompted by a recent article by Frank Arntzenius, "Some Problems for Conditionalization and Reflection".[1] Through a sequence of examples, that article purports to show limitations for a combination of two inductive principles that relate current and future rational degrees of belief: *Temporal Conditionalization* and *Reflection*:

(i) *Temporal Conditionalization* is the rule that, when a rational agent's corpus of knowledge changes between *now* and *later* solely by learning the (new) evidence, B, then coherent degrees of belief are updated using conditional probability according the formula, for each event A,

$$P_{later}(A) = P_{later}\,(A|B) = P_{now}\,(A|B)$$

(ii) *Reflection*[2] between *now* and *later* is the rule that current conditional degrees of belief defer to future ones according to the formula that, for each event A,[3]

$$P_{now}\,(A|P_{later}(A) = r) = r.$$

[1] *The Journal of Philosophy* Vol C, Number 7 (2003), 356–370.

[2] See B.van Frasseen's "Belief and the Will," this *Journal*, 81 (1984), 235–256. van Fraassen's *Reflection* has an antecedent in M.Goldstein's "Prevision of a Prevision," *JASA* 78 (1983): 817–819.

[3] Here and through the rest of this note 'r' is a standard designator for a real number – this in order to avoid Miller-styled problems. See, D.Miller's "A Paradox of Information," *Brit. J. Phil. Sci*. 17 (1966):144–147.

M.J. Schervish (✉) • T. Seidenfeld • J.B. Kadane
Carnegie Mellon University, Pittsburgh, PA, USA

H. Arló-Costa et al. (eds.), *Readings in Formal Epistemology*, Springer Graduate Texts in Philosophy 1, DOI 10.1007/978-3-319-20451-2_11

We will use the expression "*Reflection* holds with respect to the event A." to apply to this equality for a specific event A.

It is our view that neither of these principles is mandatory for a rational agent.[4] However, we do not agree with Arntzenius that, in the examples in his article, either of these two is subject to new restrictions or limitations beyond what is already assumed as familiar in problems of stochastic prediction.

To the extent that a rational person does not know *now* exactly what she or he will know in the future, anticipating one's future beliefs involves predicting the outcome of a stochastic process. The literature on stochastic prediction relies on two additional assumptions regarding states of information and the temporal variables that index them[5]:

(iii) When $t_2 > t_1$ are two fixed times, then the information the agent has at t_2 includes all the information that she or he had at time t_1.[6] This is expressed mathematically by requiring that the collection of information sets at all times through the future form what is called a *filtration*.

Second, since the agent may not know *now* the precise time at which some specific information may become known in the future, then future times are treated as *stopping times*. That is:

(iv) For each time T (random or otherwise) when a prediction is to be made, the truth or falsity of the event $\{T \leq t\}$ is known at time t, for all fixed t. Such (random) times T are called *stopping times*.

In this note, we apply the following three results[7] to the examples in Arntzenius' article. These results, we believe, help to explain why the examples at first appear puzzling and why they do not challenge either *Temporal Conditionalization* or *Reflection*. Result 11.1 covers the ordinary case, where *Reflection* holds. Results 11.2 and 11.3 establish that *Reflection* will fail when one or the other of the

---

[4]We have argued, for example, that when (subjective) probability is finitely but not countably additive, then there are simple problems where (*i*) is reasonable, but where (*i*) precludes (*ii*). See our "Reasoning to a Foregone Conclusion," *JASA* 91 (1996): 1228–1236. Also, Levi argues successfully, we think, that (*i*) is not an unconditional requirement for a rational agent. See his "The Demons of Decision," *The Monist* 70 (1987): 193–211.

[5]See, for example, section 35 of P.Billingsley, *Probability and Measure* 3*rd* edition, J.Wiley, 1995.

[6]Here and through the rest of this note 't' is a standard designator for a real number for time. More precisely, we use the subscripted variable, e.g. 't_1' to denote a specific time as the agent of the problem is able to measure it. We presume that the agent has some real-valued "clock" that quantifies a transitive relation of "is later than." Subtleties about the differences between how time is so indexed for different observers is relevant to one of Arntzenius' puzzles, to wit, the Prisoner's Problem.

[7]Proofs for these three are given in the "Appendix". In this note, we assume that all probability is countably additive.

two additional assumptions, (*iii*) and (*iv*) fail. It is not hard to locate where these assumptions are violated in the examples that Arntzenius presents.

Result 11.1 When "*later*" is a *stopping time*, when the information sets of future times form a *filtration*, and assuming that degrees of belief are updated by *Temporal Conditionalization*, then *Reflection* between *now* and *later* follows.

Result 11.2 When the information known to the agent over time *fails* to form a filtration, not only is *Temporal Conditionalization* vacuously satisfied (as its antecedent fails), but then *Reflection* fails unless what is forgotten in the failure of filtration becomes practically certain (its probability becomes 0 or 1) in time for future predictions, *later*.

Result 11.3 However, if the information known to the agent over time forms a filtration and *Temporal Conditionalization* holds, but "*later*" is not a *stopping time*, then *Reflection* between *now* and *later* holds for the specific event A, i.e., $P_{now}\,(A|P_{later}(A) = r) = r$, subject to the necessary and sufficient condition, (11.1), below.

Let H_t be the event "$t = later$." When *later* is not a stopping time, the event H_t is news to the agent making the forecasts. The question at hand is whether this news is relevant to the forecasts expressed by *Reflection*. To answer that question, concerning such forecasts about the event A, define the quantity $y_t(A)$ by

$$y_t(A) = \frac{P_{now}\,(H_t|P_t(A) = r\&A)}{P_{now}\,(H_t|P_t(A) = r)}.$$

The quantity $y_t(A)$ is an index of the current conditional dependence between A and H_t, given that $P_t(A) = r$. For example, $y_t(A) = 1$ if and only if A and H_t are conditionally independent for the agent, *now*, given that $P_t(A) = r$. In other words, by symmetry of conditional independence, $y_t(A) = 1$ if and only if the agent's current conditional probability of A given that $P_{later}(A) = r$, is unchanged by the added information H_t.

Reflection holds for A between *now* and *later*, $P_{now}\,(A|P_{later}(A) = r) = r$ if and only if, given $P_{later}(A) = r$, the conditional expected value $y_T(A) = 1$. Specifically, if and only if

$$1 = \sum\nolimits_t y_t(A) P_{now}\,(H_t|P_{later}(A) = r) \tag{11.1}$$

Thus, *Reflection* is satisfied between *now* and *later* if and only if (11.1) holds for each A.

Next, we illustrate the second and third results with examples that show how *Reflection* may fail.

Example 11.1 (Illustrating Result 11.2)

Suppose that the agent will observe a sequence of coin tosses, one at a time at a known rate, e.g. one toss per minute. Let $X_n = 1$ if the coin lands heads up on toss n, and let $X_n = 0$ otherwise. The agent does not know how the coin is loaded, but believes that it is *fair* (event A) with personal probability ½, and that with personal probability ½ it is biased with a chance of ¾ for landing tails (event A^c). Also, he believes that tosses are conditionally independent given the loading, i.e., given that the coin is *fair* or given that it is biased ¾ for tails.

Time is indexed for the agent by the number of the most recent coin toss. The time "*now*" occurs after the first toss ($t = n = 1$), and "*later*" denotes the time ($t = n = 2$) just after the second toss. Unfortunately, at each time t, the agent knows that he can remember only the most recent flip, X_t, though he knows which numbered toss it is because, e.g., he can see a clock. Suppose that the first toss lands heads up, which is the event $C = \{X_1 = 1\}$. The information that will be available to the forgetful agent later (at $t = 2$) will be only that either $B_1 = \{X_2 = 1\}$ or $B_0 = \{X_2 = 0\}$. He will not recall C because of his predictable memory lapse, and he knows all this. It is straightforward to compute:

$$P_{later}\,(A|B_1) = 2/3 \text{ and } P_{later}\,(A|B_0) = 2/5.$$

However, at $t = 1$, the agent's conditional probability for A, given event B_1 occurring at $t = 2$, satisfies $P_{now}(A \mid B_1) = 4/5$. Similarly, if *now* he conditions on event B_0 occurring at $t = 2$, his conditional probability will satisfy $P_{now}(A \mid B_0) = 4/7$.

Of course, *Temporal Conditionalization* holds vacuously at the *later* time, since the information sets available to the agent do not form a filtration. *Reflection* fails in this setting, as the agent does not remember at the *later* time what happened *now*, and he knows this all along. If B_1 occurs then $P_{later}(A) = P_{later}\,(A|B_1) = 2/3$, and if B_0 occurs then $P_{later}(A) = P_{later}\,(A|B_0) = 2/5$. Hence,

$$P_{now}\,(A|P_{later}(A) = 2/3) = 4/5$$

and

$$P_{now}\,(A|P_{later}(A) = 2/5) = 4/7. \qquad \square$$

Example 11.2 (Illustrating Result 11.3 when condition (11.1) fails and then *Reflection* fails too)

Modify Example 11.1 so that the agent has no memory failures and updates his degrees of belief by *Temporal Conditionalization*. Also, change the time "*now*" to denote the minute prior to the first toss, i.e., *now* is $t = n = 0$. Define the time "*later*" to be the *random time, T, just prior to the first toss that lands heads up*. From the point of view of the agent, the quantity T is not an observable random variable up to and including time T, and it is not a *stopping time* either. It is observable to the agent starting with time $T+1$, of course, as by then he will have seen when the first head occurs.

With probability 1 the possible values for T are $T = 0, 1, 2, \ldots$. It is straightforward to verify that: $P_{later}(A) = [1 + (3/2)^n]^{-1}$, when $T = n$, for $n = 0, 1, 2,$

. . . . Notice that $P_{later}(A) \le ½$, no matter when T occurs, and $P_{later}(A) < ½$ for $T > 0$, since if $T > 0$, the initial sequence of tosses that the agent observes all land tails up. However, from the value of $P_{later}(A)$ and knowing it is this quantity, one may calculate T exactly and thus know the outcome of the $n + 1$st toss, which is a heads. But when the agent computes $P_{later}(A)$ at the time *later*, he does not then know that *later* has arrived. Thus, *later*, he is not in a position to use the extra information that he would get from knowing when T occurs to learn the outcome of the $n + 1$st toss. To repeat the central point, T is not a stopping variable.

It is evident that *Reflection* fails, $P_{now}(A \mid P_{later}(A) = r) \neq P_{later}(A)$. The extra information, namely that $P_{later}(A) = r$ rather than merely that $P_t(A) = r$ where t is the time on the agent's clock, is information that is relevant to his *current* probability of A, since it reveals the outcome of the next toss. Even *now*, prior to any coin tosses, when he computes $P_{now}(A \mid P_{later}(A) = r)$, the conditioning event reveals to him the value of T, since n is a function of r. In this case, the conditioning event entails the information of n and when the first heads occurs, namely, on the $n + 1$st toss. Then *Reflection* fails as

$$P_{now}\left(A|P_{later}(A) = [1 + (3/2)^n]^{-1}\right) = \left(1 + 3^n/2^{n+1}\right)^{-1}.$$

It remains only to see that (11.1) fails as well. Consider the quantity $y_t(A)$ used in condition (11.1). $y_t(A) = \frac{P_{now}(H_t|P_t(A)=r\&A)}{P_{now}(H_t|P_t(A)=r)}$. Given $P_t(A) = r$, the added information that A obtains is relevant to the agent's current probability when *later* occurs. Specifically, as $P_t(A) = [1 + (3/2)^n]^{-1}$ entails that $t = n$,

$$\begin{aligned} P_{now}\left(H_t|P_t(A) = [1 + (3/2)^n]^{-1}\right) &= P_{now}\left(X_{t+1} = 1|P_t(A) = [1 + (3/2)^n]^{-1}\right) \\ &= (1/2)\,[1 + (3/2)^n]^{-1} + (1/4)\,(3/2)^n[1 + (3/2)^n]^{-1} < \tfrac{1}{2} \\ &= P_{now}\left(X_{t+1} = 1|P_t(A) = [1 + (3/2)^n]^{-1}\&A\right) \\ &= P_{now}\left(H_t|P_t(A) = [1 + (3/2)^n]^{-1}\&A\right). \end{aligned}$$

Thus, $y_t > 1$.
Hence, $1 < \sum_t y_t(A) P_{now}\,(H_t|\,P_{later}(A) = r)$. □

Example 11.3 (Illustrating Result 11.3 when (11.1) obtains and *Reflection* holds even though *later* is not a *stopping time*)

In this example, consider a sequence of three times, $t = 0$, 1, and 2. *Now* is time $t = 0$. The available information increases with time, so that the information sets form a filtration, and the agent updates his degrees of belief by *Temporal Conditionalization*. Let the random time *later* be one of the two times $t = 1$, or $t = 2$, chosen at random, but which one is not revealed to the agent. Let the event H_i be that *later* $= i$, ($i = 1$, 2) and suppose that the occurrence of H_i (or its failure) while not known to the agent at any of the three times is independent of all else that the agent does know at all three times. In this case, for each event A (even for $A = H_i$) Eq. (11.1) is satisfied. That is, by the assumptions of the problem, either

$y_t(A) = \frac{P_{now}(H_t|y_t(A)=r\&A)}{P_{now}(H_t|P_t(A)=r)} = 1$, or if $A = H_i$ then $y_t(A) = \frac{P_{now}(H_t|P_{later}(A)=r)}{P_{now}(H_t|P_t(A)=r)} = 1$. Thus, $P_{now}(A \mid P_{later}(A) = r) = r$. That is, even though *later* is not a stopping time, *Reflection* holds in this case since, given that $P_{later}(A) = r$ no new (relevant) evidence about A is conveyed through knowing that *later* has arrived, H_i. □

We note that Result 11.2 applies to the *Sleeping Beauty*[8] *Shangri La*, and *Duplication* examples of Arntzenius' article, where known failures of memory are explicit to the puzzles. Result 11.3 applies to explain the failure of *Reflection* in the two versions of the "Prisoner" example where the *local time* in the story, as measured by an ordinary clock (e.g., "11:30 PM" in John Collins's example) is not a *stopping time* for the Prisoner.

It is our impression of Collins's *Prisoner* example that the reader is easily mistaken into thinking that the *local time*, as measured by an ordinary clock in the story, is a *stopping time* for all the characters in the story. Then *Reflection* holds for each of them, in accord with Result 11.1. In Collins' example, the *local time*, e.g., 11:30 PM, is a *stopping time* for the Jailor (and also for the reader), but *not* for the Prisoner. For the Prisoner, time is measured by real-valued increments over the starting point, denoted by "now." Increments of *local time* are *stopping times* for the Prisoner. This is because the Prisoner does not know at the start of the story which of two *local times* equals his time *now*. For subsequent times, he does know how much *local time* has elapsed since *now*. But that information is not equivalent to knowing the *local time*. That difference in what is a *stopping time* for different characters is what makes this a clever puzzle, we think.

Acknowledgements Our research was carried out under NSF Grant DMS 0139911. We thank Joseph Halpern for alerting one of us (*T.S.*) to the *Sleeping Beauty* problem, independent of Arntzenius' article.

[8]See also J.Y.Halpern's "Sleeping Beauty Reconsidered: Conditioning and Reflection in Asynchronous Systems," Dept. of Computer Science, Cornell University. September, 2003. We agree with Halpern that, in our words, coherence of a sequence of previsions does not require that they will be *well calibrated* – in a frequency sense of "well calibrated." That is, we think it is reasonable for Sleeping Beauty to give a prevision of ½ to the event that the known fair coin landed heads on the flip in question, each time she is woken up. What complicates the analysis is that the repeated trials in Sleeping Beauty's game do not form an independent sequence, and her mandated forgetfulness precludes any "feedback" about the outcome of past previsions. When repeated trials are dependent and there is no learning about past previsions, coherent previsions may be very badly calibrated in the frequency sense. For other examples and related discussion of this point see, e.g., Seidenfeld, T. (1985) "Calibration, Coherence, and Scoring Rules," *Philosophy of Science* 52: 274–294.

Appendix

Proof of Result 11.1[9] Assume that when X is a random variable and C is an event, the agent's expected value $\mathrm{E}_{\mathrm{P}}(X)$ and conditional expected value $\mathrm{E}_{\mathrm{P}}(X|\ C)$ exist with respect to the probability P. Let A be an event and let $X = P(A|Y)$ be a random variable, a function of the random variable Y. Then, as a consequence of the law of total probability, with C also a function of Y,

$$P\,(A|C) = E_P\,[X|C]\,. \tag{11.2}$$

Assume that the agent's degrees of belief *now* include his *later* degrees of belief as objects of uncertainty. That is, future events such as "$P_{later}(A) = r$" and "$P_{later}(A \mid C) = q$" are proper subjects, *now*, of the agent's current degrees of belief. Suppose that, *now*, the agent anticipates using (i) *Temporal Conditionalization* in responding to the new evidence $Y = y$ that he knows he will learn at the stopping time, *later*. For example, Y might be the result of a meter reading made at the *later* time, with a sample space of m many possible values $Y = \{y_1, \ \ldots, \ y_m\}$. Thus, by ($i$), for whichever value y of Y that results,

$$P_{later}(A) = P_{later}\,(A|Y = y) = P_{now}\,(A|Y = y)\,. \tag{11.3}$$

Then, by (i) and (11.2), for C also a function of Y, the agent *now* believes that

$$P_{now}\,(A|C) = E_{P_{now}}\,[P_{later}(A)|C]\,. \tag{11.4}$$

Let C be the event, "$P_{later}(A) = r$," which we presume is a possible value for $\mathrm{P}_{later}(A)$ from the agent's current point of view. (This C is function of Y.) Then, because *later* is a stopping time,

$$P_{now}\,(A|P_{later}(A) = r) = E_{P_{now}}\,[P_{later}(A)|P_{later}(A) = r)]. \tag{11.5}$$

As

$$\mathrm{E}_{P_{now}}\,[P_{later}(A)|P_{later}(A) = r)] = r, \tag{11.6}$$

[9] van Fraassen (1995) "Belief and the Problem of Ulysses and the Sirens," *Phil. Studies* 77: 7–37, argues (pp. 17–19) that *Temporal Conditionalization* implies *Reflection*. His argument (pp. 18–19) has an additional, tacit assumption that the time t at which conditioning applies for *Reflection* is a stopping time.

therefore

$$P_{now}\,(\mathrm{A}|P_{later}(A) = r) = r, \tag{11.7}$$

i.e., then *Reflection* holds as well. □

Proof of Result 11.2 To show that *Reflection* fails, consider two times $t_1 < t_2$. Call an event *forgotten* if its truth or falsity is known at time t_1 but not at time t_2. From the assumption that these times do not form a filtration, let E be forgotten between t_1 and t_2 and allow that at t_1 this is known to happen at t_2. Since $P_{t_1}(E) \in \{0, 1\}$, conditioning will not change this value, i.e.,

$$P_{t_1}(E) = P_{t_1}\,(E|P_{t_2}(E) = r) \tag{11.8}$$

for a set of r-values of probability 1 under P_{t_1}. But, since it is known at t_1 that E will be forgotten at t_2, $P_{t_1}\,(0 < P_{t_2}(E) < 1) = 1$. Hence *Reflection* fails as $0 < r < 1$ in (11.8). □

Proof of Result 11.3 Assume that the agent's information sets form a filtration over time and that *Temporal Conditionalization* holds between *now* and *later* but that *later* is not a stopping time for the agent. Let H_t be the event "*later* $= t$" for the specific time t. That is, assume that $0 < P_{later}(H_t) < 1$, when *later* occurs at t.

Later is the future time we will focus on in calculating whether *Reflection* holds, i.e. we will inquire whether for each event A, $P_{now}(\mathrm{A} \mid P_{later}(A) = r) = r$, or not. We calculate as follows.

$$P_{now}\,(A|P_{later}(A) = r)$$

$$= \sum\nolimits_t P_{now}\,(A\&H_t|P_{later}(A) = r)$$

by the law of total probability.

$$= \sum\nolimits_t P_{now}\,(A|P_{later}(A) = r\&H_t)\,P_{now}\,(H_t|P_{later}(A) = r)$$

by the multiplication theorem

$$= \sum\nolimits_t \frac{P_{now}\,(H_t|P_t(A) = r\&A)}{P_{now}\,(H_t|P_t(A) = r)} P_{now}\,(A|P_t(A) = r)\,P_{now}\,(H_t|P_{later}(A) = r)$$

by Bayes' theorem and the equivalence of

$$(P_{later}(A) = r\&H_t) \;\; and \;\; (P_t(A) = r\&H_t)$$

$$= r\sum_t \frac{P_{now}\,(H_t|P_t(A) = r\&A)}{P_{now}\,(H_t|P_t(A) = r)} P_{now}\,(H_t|P_{later}(A) = r)$$

as $P_{now}\,(A|P_t(A) = r) = \text{r}$ by Result 11.1.

$$= r\sum_t y_t(A) P_{now}\,(H_t|P_{later}(A) = r).$$

by the definition of $y_t(A)$

Hence, $P_{now}\,(A|P_{later}(A) = r) = r$ if and only if $\sum_t y_t(A) P_{now}(H_t|P_{later}(A) = r) = 1$, which is condition (11.1). □