

## Chapter 9

# Discussion: A Mistake in Dynamic Coherence Arguments?

Brian Skyrms

### Static Coherence of Degrees of Belief

The person whose degrees of belief are being tested for coherence acts as a bookie. She posts her fair prices for wagers corresponding to her degrees of belief. Her degrees of belief are *incoherent* if a cunning bettor can make a Dutch book against her with a finite system of wagers—that is, there is a finite set of wagers individually perceived as fair, whose net payoff is a loss in every possible future. Otherwise her degrees of belief are *coherent*. De Finetti ([1937] 1980) proved the following theorem: *Degrees of belief are coherent if and only if they are finitely additive probabilities.*

Obviously, if a Dutch book can be made with a finite number of fair transactions, it can be made with a finite number of uniformly favorable transactions. The bettor pays some small transaction premium  $\varepsilon$  to the bookie for each of the  $n$  transactions where  $n\varepsilon$  is less than the guaranteed profit that the bettor gets under the Dutch book based on fair prices. Let us bear in mind that this point applies equally well in what follows.

---

Received April 1992; revised June 1992.

Send reprint requests to the author, Department of Philosophy, University of California, Irvine, CA 92717, USA.

B. Skyrms (✉)

Department of Philosophy, University of California, Irvine, CA, USA

## Dynamic Coherence for Updating Rules

The epistemologist acts as bookie. Her updating rule is public knowledge. Today she posts her fair prices, and does business. Tomorrow she makes an observation (with a finite number of possible outcomes each of which has positive prior probability) and updates her fair prices according to her updating rule. The updating rule is thus a *function* from possible observations to revised fair prices. The day after tomorrow she posts prices again, and does business. The pair consisting of the (1) her fair prices for today and (2) her updating function will be called the bookie's *epistemic strategy*.

The bookie's *epistemic strategy* is *coherent* if there is no possible bettor's strategy which makes a Dutch book against him (the bettor's strategy being a pair consisting of (1) a finite number of transactions today at the bookie's posted prices and (2) a function taking possible observations into a finite number of transactions the day after tomorrow at the prices that the bookie will post according to her epistemic strategy). Lewis (reported in Teller 1973) proves that *the epistemologist's strategy is coherent only if her degrees of belief today are finitely additive probabilities and her updating rule is Bayes's rule of conditioning*. The "only if" can be strengthened to "if and only if" (see section "The converse"). (For generalizations of this theorem see van Fraassen 1984 and Skyrms 1987, 1990.)

Notice that the relevant notions of coherence and incoherence here apply not just to the *pair* of degrees of belief for today and the day after tomorrow, but rather to an *epistemic strategy*, which is a more complicated object. A focus on the former notion leads understandably to skepticism regarding dynamic coherence, as in Hacking (1967), Kyburg (1978), and Christensen (1991).

## The Dynamic Dutch Book

Coherence of degrees of belief today is the static case. It remains to show that for any non-Bayes updating rule, there is a bettor's strategy which makes a Dutch book. Let the conditional probability of  $A$  on  $e$ , that is  $\Pr(A \& e)/\Pr(e)$ , be symbolized as usual, as  $\Pr(A|e)$ , and let the probability that the updating rule gives  $A$  if  $e$  is observed be  $\Pr_e(A)$ . If the predictor's rule disagrees with conditioning, then for some possible evidential result  $e$  and some  $A$ ,  $\Pr_e(A)$  is not  $\Pr(A|e)$ . Suppose that  $\Pr(A|e) > \Pr_e(A)$ . (The other case is similar.) Let the discrepancy be  $\delta = \Pr(A|e) - \Pr_e(A)$ . *Here is a bettor's strategy which makes a Dutch book:*

TODAY: Offer to sell the bookie at her fair price:

- 1: [\$1 if  $A \& e$ , 0 otherwise]
- 2: [\$ $\Pr(A|e)$  if not- $e$ , 0 otherwise]
- 3: [\$ $\delta$  if  $e$ , 0 otherwise]

DAY AFTER TOMORROW:

If  $e$  was observed, offer to buy [\$1 if  $A$ , 0 otherwise] for its current fair price,  $\Pr_e(A) = \Pr(A|e) - \delta$ .

Then in every possible situation, the bookie loses  $\delta \Pr(e)$ .

## The Converse

If the bookie has the strategy of updating by Bayes's rule of conditioning, then every payoff that a bettor's strategy can achieve can be achieved by betting only today (see Skyrms 1987). This reduces our case to the static case. Thus, by de Finetti's result, if the epistemologist's prior degrees of belief are finitely additive probabilities and her updating rule is Bayes's rules of conditioning, then she is dynamically coherent.

## Sequential Analysis 1: A Mistake in the Dynamic Coherence Argument?

Maher's (1992b) objection is that the bookie will see it coming and refuse to bet. This is made precise by modeling the bookie's situation as a sequential choice problem, as shown in Fig. 9.1. The bookie sees that if she bets today and  $e$  occurs, then at decision node 2, she will find the cunning bettor's offer fair according to her revised probability,  $\Pr_e(A)$ . Thus she sees that betting today leads to a sure loss. Since she prefers net gain of zero to a sure loss, she refuses to bet today—frustrating the cunning bettor who goes home unable to execute his plan.

The first thing that must be said about "Maher's objection" is that it is misleading to represent it as showing a "mistake" in the dynamic coherence theorem. Under the conditions of the theorem the bookie posts her fair prices for today and honors them. There is no provision for changing one's mind when approached by a cunning bettor who discloses his strategy, nor indeed any mention of a requirement that the cunning bettor disclose his strategy prior to the initial transaction. But Maher might be read as suggesting a different conception of dynamic coherence in this setting:

The epistemologist acts as bookie. Her updating rule is public knowledge. Today she posts her tentative fair prices, but in fact does business only with bettors who disclose their strategies in advance, and does so on the basis of sequential decision analysis. Tomorrow she makes an observation (with a finite number of possible outcomes each of which has positive prior probability) and updates her probabilities according to her updating rule. The day after tomorrow she posts prices again, and does business according to those prices.

She is *coherent* if there is no possible bettor's strategy which makes a Dutch book against her.

This is an interesting modification of the usual notion of dynamic coherence, and it merits investigation. Is it a better motivated conception of dynamic coherence? What differences does it make?

## Sequential Analysis 2: A Mistake in the Mistake?

A natural reaction to Maher's line might be to say that the redefinition unfairly prejudices the case against dynamic coherence arguments. It is therefore of some interest to see that the dynamic Dutch book still goes through under the revised scenario.

There is a gratuitous assumption in the analysis presented in Fig. 9.1. Why is it assumed that the cunning bettor will just go home if the bookie refuses to bet today? The bettor's strategy which I presented says otherwise. The bettor will make an offer the day after tomorrow if  $e$  was observed. So the branch of the decision tree where the bookie refuses transactions today cannot simply be assumed to have payoff of zero, but requires further analysis. This is done in Fig. 9.2.

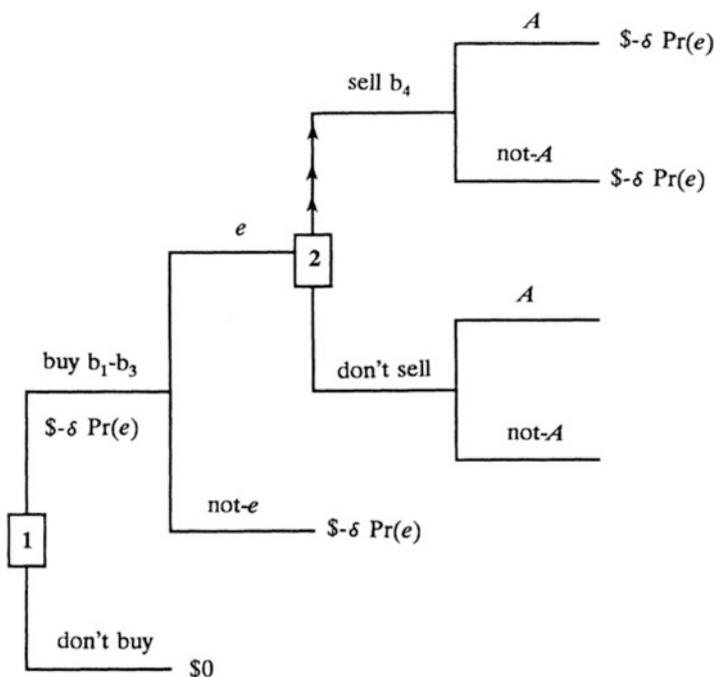


Fig. 9.1 Maher-Levi sequential analysis

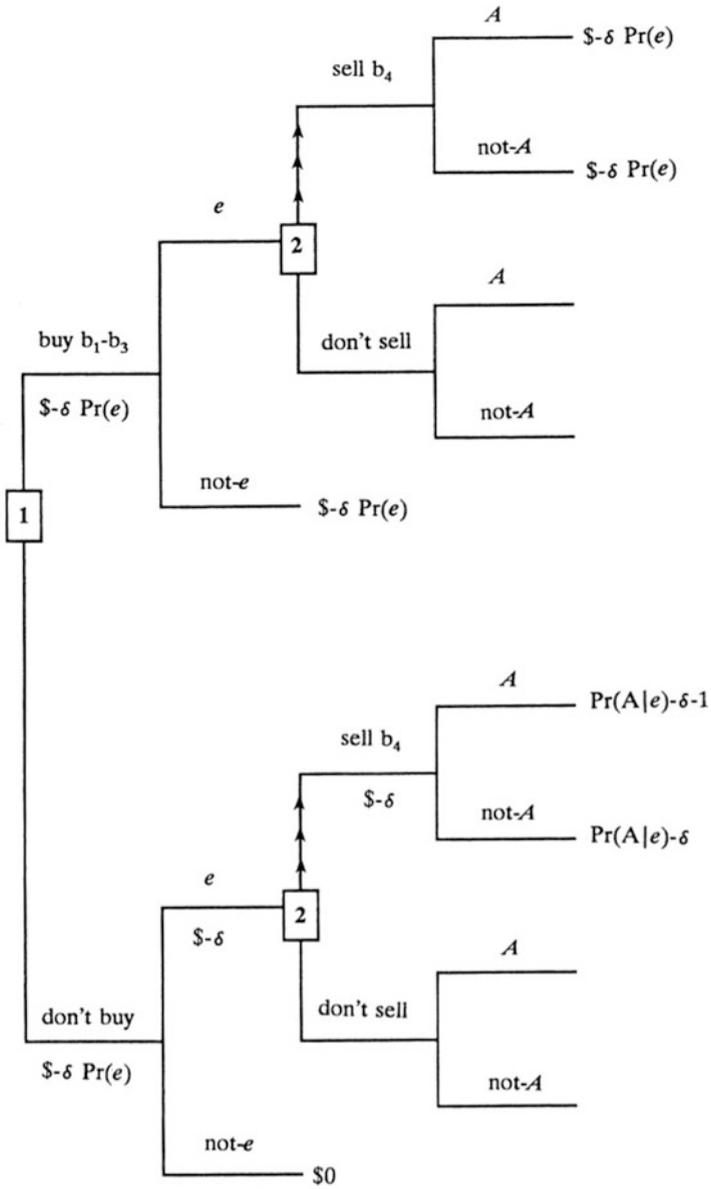


Fig. 9.2 Full sequential analysis

Note that the bookie knows that if  $e$  is observed, she will accept the offer the day after tomorrow for the same reason on the lower path as on the upper. Deciding now not to bet ever is not an option. If the offer the day after tomorrow is accepted but the offer today was not and  $e$  and  $A$  both happen, then the net payoff is the price the cunning bettor paid,  $\$Pr(A|e) - \delta$ , less the lost bet,  $\$ - 1$ , as shown. If  $e$  occurs but  $A$  does not, the net payoff is just  $\$Pr(A|e) - \delta$ . For the bookie's current analysis of this decision tree, to get the relevant expectation over  $A$  occurring or not we average using as weights her current conditional probabilities,  $Pr(A|e)$  and  $Pr(-A|e)$ . Thus the value at the node where the bookie refused to bet today and where  $e$  is observed tomorrow is

$$Pr(A|e) \$ \left[ \{Pr(A|e) - \delta\} - 1 \right] + [1 - Pr(A|e)] \$ \left[ Pr(A|e) - \delta \right] = \$ - \delta.$$

Then the value at the node where the bookie refused to bet today is not 0 but rather  $\$ - \delta Pr(e)$ . This is just the same as the value at the node where the bookie agrees to bet today.

In fact, if we consider the version of the Dutch-book strategy where the bettor adds an  $\varepsilon$  premium for each transaction, the upper branch involves four transactions and the lower branch involves only one, so the upper branch has a higher payoff than the lower branch. *Even though the bookie sees it coming, she will prefer the sure loss of the upper branch because doing so looks strictly better to her than the alternative.*

### Sequential Analysis 3: What Makes the Cunning Bettor Tick?

Why did the cunning bettor adopt a strategy of staying around if the bookie decided not to bet today? The official answer in sections “[Dynamic coherence for updating rules](#)” and “[The dynamic Dutch book](#)” is “Don’t ask”. Any bettor’s strategy which makes a Dutch book will prove incoherence. But, as Levi (1991) points out, that sort of analysis proceeds in strategic normal form rather than in extensive form. Might it be that the cunning bettor’s strategy described would have to be *sequentially irrational*? That is to say, might it not be that staying around and betting the day after tomorrow if the bookie decided not to bet today would not maximize expected utility for the cunning bettor in the belief state he would have in that case the day after tomorrow? If this could be shown, then the cunning bettor’s strategy that I have described would have to rest on a noncredible threat, and the significance of the analysis of the previous section would be called into question. (For discussion of such noncredible threats in extensive form games and of sequential rationality, see Selten 1975 and Kreps and Wilson 1982.)

But such is not the case. Suppose that the bettor is a Bayesian; that he starts out with exactly the same degrees of belief as the bookie; and that he updates by conditioning. If  $e$  is observed tomorrow—whether or not the bookie accepted the

bet today—he conditions on  $e$  and the day after tomorrow his fair price for  $b_4$  is  $\$ \text{pr}(A|e)$ . But his strategy only commits him to offering to pay the bookie’s fair price,  $\$ \text{pr}(A|e) - \delta$ , to buy back  $b_4$  for what he perceives as a net gain in expected utility of  $\$ \delta$ . This bettor’s threat to stick around and bet the day after tomorrow if  $e$ , even if the bookie declines to bet today, is perfectly credible and consistent with sequential rationality. If he is called upon to carry out the threat, he maximizes expected utility by doing so.

## Strategic Rationality

Let us explicitly model the bookie’s choice of an updating strategy. The bookie and the bettor start out with identical priors. The bettor updates by conditioning. First the bookie chooses an updating strategy. Then the bettor bets, the evidence comes in, the bookie updates according to her updating rule, and the bettor bets again. The bookie’s initial strategy is either to choose updating by conditioning or not.

If the bookie chooses the strategy of updating by conditioning, then the fair prices of the bookie and bettor agree at all times. Thus either no transactions are made, or any transactions have net change in expected utility of 0 for both players. The bookie’s expected utility of choosing the strategy of updating by conditioning is zero. If, however, the bookie chooses an updating strategy at variance with conditioning then, for the bettor, the expected utility of betting is greater than that of not betting (section “[Sequential analysis 3: what makes the cunning bettor tick?](#)”) and the net expected utility for the bookie is negative (section “[Sequential analysis 2: a mistake in the mistake?](#)”). At the first choice point the bookie is strictly better off by choosing the rule of updating by conditioning.

Thus the strategy combination in which the bookie updates by conditioning and the bettor does not bet at all is an *equilibrium* in the sense that no player will perceive it in his or her interest at any decision node to deviate from that strategy. But no strategy combination in which the bookie chooses a strategy at variance with conditioning is such an equilibrium.

## The Bottom Line

Two ways of strengthening the requirements for a dynamic Dutch book were suggested by the discussions of Levi (1987) and Maher: (1) We require the cunning bettor to disclose his strategy, and allow the bookie to use knowledge of that strategy in a sequential analysis when deciding whether to bet today or not, and (2) we require that the cunning bettor’s strategy itself be sequentially rational. The somewhat surprising result is that the additional restrictions made no difference. The bookie whose epistemic strategy is at odds with conditioning is also subject to a Dutch book in this stronger sense. “Seeing it coming” does not help. It is at

the very least a noteworthy property of the rule of conditioning that in this sort of epistemic situation, it alone is immune from a Dutch book under either the original or strengthened requirements.

## Postscript: Conditioning, Coherence and Rationality

Many of the concerns of Levi and Maher have not been addressed in the foregoing. Levi is concerned to resist the doctrine of “confirmational tenacity”, according to which the only legitimate way in which to update is by conditioning. Maher wishes to resist the doctrine that rationality requires dynamic coherence at all costs. Does the foregoing show that conditioning is the only coherent way to ever update one’s probabilities? Does it show that rationality requires coherence at all costs?

I agree with Levi and Maher in answering “no” to both questions. With regard to the first, let me emphasize that the Lewis proof takes place within the structure of a very special epistemic model. In that context it shows that the rule of conditioning is the unique dynamically coherent updating rule. It does not show that one must have an updating rule. It does not apply to other epistemic situations which should be modeled differently. The modeling of a variety of epistemic situations and the investigation of varieties of dynamic coherence in such situations is an ongoing enterprise (in which I take it that both Levi and I are engaged; see Skyrms 1990 for further discussion).

Maher is concerned that an uncritical doctrine of “dynamic coherence at all costs” could lead one to crazy belief changes and disastrous actions. Should Ulysses have changed to 1 his prior probability of safe sailing conditional on hearing the Sirens’ song so that subsequently his belief change would be in accordance with the rule of conditioning? Nothing in the foregoing implies that he should. In the first place, there is something a little odd in thinking that one achieves dynamic coherence by changing the original prior  $pr_1$  to the revised prior  $pr_2$  so that the change to  $pr_3$  will agree with conditioning. What about the change from  $pr_1$  to  $pr_2$ ? But, more fundamentally, I would agree with Maher that rationality definitely does not require coherence *at all costs*. Where costs occur they need to be weighed against benefits. There are lucid discussions of this matter in Maher (1992a, b). These things said, it remains that in the Lewis epistemic model under the “sequentialized” notion of dynamic coherence, the unique coherent updating rule is the rule of conditioning.

**Acknowledgement** I would like to thank Brad Armendt, Ellery Eells, Isaac Levi, Patrick Maher and an anonymous referee for helpful comments on an earlier draft of this note. I believe that Maher, Levi and I are now in substantial agreement on the issues discussed here, although differences in emphasis and terminology may remain.

## References

- Christensen, D. (1991). Clever bookies and coherent beliefs. *The Philosophical Review*, 100, 229–247.
- de Finetti, B. ([1937] 1980). Foresight: Its logical laws, its subjective sources, translated in H. E. Kyburg, Jr., & H. Smokler (Eds.), *Studies in subjective probability* (pp. 93–158). (Originally published as “La Prevision: ses lois logiques, ses sources subjectives”, *Annales de l’Institut Henri Poincaré*, 7, 1–68.) Huntington: Kreiger.
- Hacking, I. (1967). Slightly more realistic personal probability. *Philosophy of Science*, 34, 311–325.
- Kreps, D., & Wilson, R. (1982). Sequential equilibria. *Econometrica*, 50, 863–894.
- Kyburg, H. (1978). Subjective probability: Criticisms, reflections and problems. *The Journal of Philosophical Logic*, 7, 157–180.
- Levi, I. (1987). The demons of decision. *The Monist*, 70, 193–211.
- Levi, I. (1991). Consequentialism and sequential choice. In M. Bacharach & S. Hurley (Eds.), *Foundations of decision theory* (pp. 92–146). Oxford: Basil Blackwell.
- Maher, P. (1992a). *Betting on theories*. Cambridge: Cambridge University Press.
- Maher, P. (1992b). Diachronic rationality. *Philosophy of Science*, 59, 120–141.
- Selten, R. (1975). Reexamination of the perfectness concept of equilibrium in extensive form games. *International Journal of Game Theory*, 4, 25–55.
- Skyrms, B. (1987). Dynamic coherence and probability kinematics. *Philosophy of Science*, 54, 1–20.
- Skyrms, B. (1990). *The dynamics of rational deliberation*. Cambridge: Harvard University Press.
- Teller, P. (1973). Conditionalization and observation. *Synthese*, 26, 218–258.
- van Fraassen, B. (1984). Belief and the will. *Journal of Philosophy*, 81, 235–256.