4

Facing the Challenges

While according to some, abduction is the cornerstone of scientific methodology, critics have argued that if abduction is at variance with Bayes's rule, then, like any other such rule, it is to be rejected as leading to irrational belief updates. For many years, the standard argument for this claim has been the previously mentioned dynamic Dutch book argument, which purports to show that updating by any rule other than Bayes's makes one liable to sure financial losses. However, Bayesians have increasingly come to regard their erstwhile favored argument as addressing the wrong issue, to wit, whether it is rational from a *practical* rather than an *epistemic* viewpoint to deviate from Bayes's rule. This has led some theorists to pursue a different strategy in defense of Bayes's rule, a strategy that is purported to offer a distinctively epistemic argument in support of that rule, spelled out in terms of inaccuracy minimization. Roughly, the argument is that by updating via any non-Bayesian rule, one's degrees of belief are not as accurate as they would have been had one updated via Bayes's rule.

In this chapter, I begin with showing that neither of the aforesaid arguments succeeds. The first part of the chapter argues that while current developments in mainstream epistemology may help to deflect some of the criticism that the dynamic Dutch book argument has met, the argument fails nonetheless because it rests on unfounded (and unstated) premises. The second part focuses on the inaccuracy-minimization defense of Bayes's rule, arguing that there appear to be several equally legitimate ways to interpret the notion of inaccuracy minimization and using computer simulations to show that under some of them it may be abduction rather than Bayes's rule that does best with regard to inaccuracy minimization. Chapters 6 and 7 greatly expand on the latter finding.

4.1 The Dynamic Dutch Book Argument Revisited

According to Bayesians, an agent is rational exactly if, first, at any given moment her degrees of belief obey the probability axioms—this was the content of SC discussed in section 3.2.1—and second, as she acquires new evidence, she changes her degrees of belief in accordance with Bayes's rule, which was the content of DC discussed in the same section. As before, where Pr is defined to be a probability function representing a given agent's belief system at a certain point in time, $\Pr_E$ represents the belief system that results from Pr upon learning E and nothing stronger. We can thus write Bayes's rule for belief change

$$\Pr_E(H) = \Pr(H \mid E), \tag{4.1}$$

with

$$\Pr(H \mid E) = \Pr(H)\frac{\Pr(E \mid H)}{\Pr(E)}. \tag{4.2}$$

It was pointed out that (4.2) is a form of Bayes's theorem, which follows from the Kolmogorovian ratio definition of conditional probability. By contrast, (4.1) does not follow from probability theory, which merely constrains what beliefs can be rationally held at one and the same time and is silent on *change* of belief.

For many decades, the Ramsey–de Finetti Dutch book argument has been viewed as key to the Bayesian account of rationality. We saw in chapter 3 that according to this argument we are susceptible to Dutch books—collections of bets ensuring a negative net pay-off come what may—precisely if our degrees of belief violate the axioms of probability. From this, Ramsey and de Finetti concluded that rational degrees of belief are formally probabilities.

Hacking (1967) may have been the first to observe that the Ramsey–de Finetti argument in fact does nothing to justify Bayes's rule. However, a few years after the publication of Hacking's paper, Paul Teller (1973) reported a Dutch book argument—which he attributed to Lewis—aimed at justifying Bayes's rule as the only rational update rule.[1] This dynamic Dutch book argument (as it is now called) purports to show that if a person updates by some rule other than Bayes's, she can be offered a series of bets at different

1. The source of the argument reported by Teller was later published as Lewis (1999).

Table 4.1: Results of updating three bias hypotheses on two consecutive heads via EXPL and, in parentheses, Bayes's rule.

	H_1	H_2	H_3	Next toss heads
Initial probabilities	.33	.33	.33	.50
First toss heads	.00 (.00)	.30 (.33)	.70 (.67)	.85 (.83)
Second toss heads	.00 (.00)	.16 (.20)	.84 (.80)	.92 (.90)

points in time such that each bet will seem fair at the time it is offered and yet jointly the bets guarantee a financial loss. What is worse, the argument continues, the person could have seen this loss coming. This vulnerability to "dynamic" Dutch books has convinced many that non-Bayesian updating is a mark of irrationality.

Van Fraassen (1989, ch. 6) gives a particularly lucid presentation of the dynamic Dutch argument that specifically targets abduction. His presentation makes use of a specific statistical model and, although not explicitly stated, models probabilistic abduction in terms of the rule we labeled "EXPL," which, as explained, first updates on incoming evidence via Bayes's rule, then adds an explanation bonus (.1, in van Fraassen's argument) to the best-explaining hypothesis, and finally renormalizes.

There is considerable leeway in defining explanatory bestness. Simplifying van Fraassen's presentation somewhat, we look at EXPL in the context of a coin-tossing model, the hypotheses of interest all pertaining to possible biases of the coin. (For brevity of notation, and unless indicated otherwise, from here on I use the label "EXPL" to designate the instance that assigns a bonus of .1 to best explanations.) We also follow van Fraassen in taking a bias hypothesis to be the best explanation of the available evidence at a given time (the outcomes of the tosses until that time) if that bias was closest to the actually observed frequency of heads in the evidence.

Consider three hypotheses concerning the bias of a given coin C:

H_1: C has perfect bias for tails.
H_2: C is fair.
H_3: C has perfect bias for heads.

The three hypotheses are assumed to be jointly exhaustive and to have equal initial probability. We now start flipping our coin and update our probabilities as the first and second toss both come up heads. Table 4.1 gives the results when we update by means of EXPL and by means of Bayes's rule (Bayes's rule updates are given in parentheses). If the dynamic Dutch book argument is correct, then it must be that a Dutch book can be made against anyone who updates her beliefs by means of EXPL. That seems to be the case indeed.

Suppose a bookie approaches you and offers to make bets on the following hypotheses concerning tosses of C, of which you know it is either fair or has a perfect bias either for heads or for tails:

A: The first two tosses will land heads.
B: The third toss will land tails.

The bookie offers two bets:

I pays \$240 if A is false.
II pays \$3,000 if $A \wedge B$ is true.

The expected values (initial probability times pay-off) of these bets are \$140 and \$125, respectively. The bookie is happy to sell the bets at these prices and you are happy to buy them, at a total cost of \$265.

Now suppose that at least one of the first two tosses does *not* come up heads. Then you have won bet I but lost bet II, which means that you receive \$240 but nevertheless suffer a total loss of \$25 (given that you payed \$ 265 for the bets). That, it would seem, is all in the game. However, suppose the first two tosses do come up heads. Then you have lost bet I but may still win bet II. However, before the coin is tossed a third time, the bookie approaches you again and now instead of proposing to sell any bets proposes to buy the following bet from you:

III pays \$3,000 if B is true.

If we assume that you update by EXPL, your probability that B will come true is now .08, as shown in table 4.1. So you will agree to sell III for the price of \$240. However, whatever further may happen, you have now lost money. If the third toss does land tails, you receive \$3000 but have to pay the same amount; in the other case, you do not have to pay anything but will also not receive anything. Regardless, you have spent \$265 on the bets you bought and only made \$240 from selling bet III, meaning that whether A does or does not come true, you are bound to lose \$25. Were you to use Bayes's rule instead, nothing of this kind could happen. Using that rule you would in our example have assigned a probability of .1 to *B* had the first two tosses come up heads. Thus not $240 but $300 would have seemed a fair price to ask for III, in which case you would have had a gain of $35 (and in which case presumably no bookie would have proposed to buy III).

The foregoing argument is open to a number of immediate objections. For example, the betting concept of probability (see ch. 3, footnote 3), on which all Dutch book arguments ultimately rest, may be deemed psychologically unrealistic (see, e.g., Williamson, 1998). Perhaps even more unrealistic is the fact that the bookie is assumed to know not only your current probabilities but also your update rule. That is unrealistic because, typically, *we ourselves* do not exactly know how we update our belief states. Most importantly, it may be held against this argument that a loss in one respect can be outweighed by a benefit in another respect; the *net* effect of non-Bayesian updating could thus be positive, no matter the costs that also come with it. That is a main point to be argued for further on.

But note that all these objections grant (what is sometimes called) the dynamic Dutch book *theorem*—which asserts that updating by a probabilistic version of abduction such as EXPL or by any other non-Bayesian rule makes one susceptible to a dynamic Dutch book—and target the dynamic Dutch book *argument*—which infers from the theorem that non-Bayesian updating is irrational. Patrick Maher's (1992) is one of the few attempts to refute the former. He argues that if the bettor in the dynamic Dutch book argument looks ahead, she can see the sure loss coming and so will decline to bet. In terms of the above example, already when the bookie offers you bets I and II, you can calculate that, should the first two tosses indeed come up heads, you will prefer to sell back the bet on *B* to the bookie and thus be guaranteed to lose. This situation is easily avoided: simply do not buy bets I and II in the first place. We just noted that even if updating by some rule other than Bayes's should entail costs, it does not follow that non-Bayesian updating is irrational; it might have advantages that outweigh the costs. However, it seems that Maher's "look-before-you-leap principle" (as Earman, 1992, p. 49, calls it) opens up the possibility of non-Bayesian updating at no cost at all.

Unfortunately, as Brian Skyrms (1993b) points out, Maher's argument hinges crucially on the assumption that if (again in terms of our example) you refuse to buy the first two bets, the bookie will not propose to buy III, and this need not be so. If the assumption is dropped and the expected pay-off of selling III but not buying I and II is calculated, you will find that, although this time there is no sure loss, the *expected* loss is the same as when the first bets are bought. In fact, if the bookie should add some premium, however minute, for each transaction made, buying these bets will be preferable to only selling III.

Maher (1993) acknowledges Skyrms's point of critique and concludes from it that the dynamic Dutch book theorem is correct after all. As I argued in Douven (1999), however, Maher may be conceding too fast here, for there appears to be a strategy, building on the principle he proposed, that does offer the possibility of non-Bayesian updating free of charge. The basic idea underlying this strategy is that the advice to look before you leap can be made into a precise method for calculating the initial probability of a proposition that *corresponds to*, or is *symmetrical with*, the process of updating that we enter at the moment the evidence relevant to the truth of that proposition comes in. As subsequently shown, had you looked ahead in this sense, you would not have accepted the previously described initial bets I and II—not because you would have seen that in combination they make you vulnerable to a Dutch book even though separately they seem fair, but simply because the bets would *not* have seemed fair. The strategy comprises three parts.

We previously saw that, on a Bayesian model of learning, our present beliefs constrain the acquisition of new beliefs through equation (4.1). The first part of my proposal is to keep this on board but to take seriously a suggestion by Ramsey concerning how we arrive at (nontrivial) conditional probabilities, a suggestion that makes those probabilities depend on updating, specifically, on hypothetical updating. Ramsey (1929), in a famous footnote, suggests that to determine our probability of H conditional on E, we run through all the following in "simulation mode": (1) add E to our stock of beliefs; (2) make minimal adjustments to preserve consistency of our beliefs (if necessary); and (3) determine in that new belief state how probable we deem H. Various authors have argued that this procedure actually captures the notion of conditional probability better than the ratio definition does (see, e.g., Edgington, 1995). The simulation procedure has also been shown to be more descriptively adequate than the latter (Zhao, Shah, & Osherson, 2009). The proposal is now to bring your actual and hypothetical, or "pretend," updating practices into harmony. That will make your prior conditional probabilities reflect whichever update rule you actually happen to use. More exactly, the non-Bayesian should, on this proposal, adhere to the following:

$$\text{(4.3)} \qquad \Pr_E(H) \;=\; \Pr(H \mid E) \;\equiv\; \ldots$$

with the definition of the preferred rule (in the case of our example, EXPL) replacing the ellipsis.[2]

It may not be immediately clear how this could affect the problem the non-Bayesian faces, for conditional probabilities apparently had no part in the situation that led to the Dutch book presented previously. However, that applies only under a particular assumption. The Dutch book in this chapter was closely modeled on van Fraassen's more elaborate example. This included the way the initial probabilities of A and $A \wedge B$ were calculated. In particular, it was assumed that to calculate initial probabilities of n successive heads for a model with m bias hypotheses, one can use what is often called "the special multiplication rule":

$$\text{(SMR)} \qquad \Pr(E_1 \wedge \cdots \wedge E_n) \;=\; \sum_{i=1}^{m} \prod_{j=1}^{n} \Pr(H_i)\Pr(E_j \mid H_i).$$

Provided that the E_js are conditionally independent given each H_i, SMR is an admissible way to calculate initial probabilities. A more general method, which does not assume conditional independence, is

$$\text{(GM)} \quad \Pr(E_1 \wedge \cdots \wedge E_n) \;=\; \left(\sum_{i=1}^{m} \Pr(H_i)\Pr(E_1 \mid H_i)\right) \times \left(\prod_{j=2}^{n} \sum_{i=1}^{m} \Pr(H_i \mid E_1 \wedge \cdots \wedge E_{j-1})\Pr(E_j \mid H_i)\right).$$

Note that if the hypotheses and evidence statements are of the sort we considered in the dynamic Dutch book argument, the antecedents and consequents of all the conditional probabilities of any hypothesis H_i given one or more of the E_js are explanatorily dependent—the hypotheses are explanatorily relevant to the evidence statements—or so we assumed for the sake of van Fraassen's argument.

Appendix A shows the following:

2. Note that this would also allow explanationists to appropriate the Bayesian definition of confirmation (see section 3.2.1) and thus obliterates the need to come up with a definition of their own, for instance, in the manner of Conee and Feldman (2004, 2008), which has appeared to be problematic (Byerly, 2013).

Theorem 4.1 *Where $H_1, \ldots, H_n$ is a set of mutually exclusive and jointly exhaustive hypotheses such that* $\Pr(H_i) > 0$, *for all i, and $E_1, \ldots, E_m$ are independent conditional on H_i, for all i, then SMR and GM are equivalent methods for computing initial probabilities if, and only if, the conditional probabilities in GM are calculated according to Bayes's theorem.*

The second part of the strategy of Douven (1999) consists of the recommendation that in this kind of case, non-Bayesian updaters (e.g., EXPL updaters) use the more general method.

To illustrate the difference and also to show how the strategy helps you avoid the Dutch book, note that applying SMR in our example to compute the probability of A amounts to computing

$$\sum_{i=1}^{3} \Pr(H_i) \Pr(\text{first toss heads} \mid H_i) \Pr(\text{second toss heads} \mid H_i) \quad = \quad \left(\frac{1}{3}\right)(0)(0) + \left(\frac{1}{3}\right)\left(\frac{1}{2}\right)\left(\frac{1}{2}\right) + \left(\frac{1}{3}\right)(1)(1) = \frac{5}{12} \approx .417.$$

By contrast, using GM while attending to (4.3) yields the following probability for A:

$$\left(\sum_{i=1}^{3} \Pr(H_i) \Pr(\text{first toss heads} \mid H_i)\right) \times \left(\sum_{i=1}^{3} \Pr(H_i \mid \text{first toss heads}) \Pr(\text{second toss heads} \mid H_i)\right) =$$

$$\left(\left(\frac{1}{3}\right)(0)+\left(\frac{1}{3}\right)\left(\frac{1}{2}\right)+\left(\frac{1}{3}\right)(1)\right) \times \left((0)(0)+\left(\frac{3}{10}\right)\left(\frac{1}{2}\right)+\left(\frac{7}{10}\right)(1)\right) = .425.$$

Thus, had you followed the strategy recommended here, you would *not* have been willing to pay \$240 for a bet on A being false and indeed would have been guarded against dynamic Dutch books quite generally.

To see why this is so, consider that what the bookie in the foregoing example exploits is that your initial odds for $A \wedge B$ to $A \wedge \neg B$ predictably diverge from what your odds for $A \wedge B$ to $A \wedge \neg B$ will become if the first two tosses should land heads. This divergence can arise because if the first tosses indeed land heads, you will twice endow H_3 with a bonus, which is not accounted for when you calculate the initial probabilities of $A \wedge B$ and

$\neg A$. Because you already *know* that H_3 will get two bonuses under the given circumstances, an obvious idea is that, to ensure that the divergence cannot arise, you do account for these bonuses when you calculate initial probabilities. However, the method of calculating initial probabilities employed in the example leaves no room for that. In contrast, because using GM to calculate $\Pr(E_1 \wedge \cdots \wedge E_n)$ in a model with m hypotheses requires calculating each $\Pr(H_i \mid E_1 \wedge \cdots \wedge E_j)$, this procedure, when combined with (4.3), lets us take account of anything we take account of when we actually update on $E_1, \ldots, E_{n-1}$, should these obtain. As per Ramsey's suggestion, it is as if in calculating in this method the initial probability of a proposition, stating the occurrence of n events, we go through a process of updating hypothetically, "pushed into the fantasy" that we come to know, one after the other, the first $n - 1$ events that will occur if the proposition is true. In this process, each of the conditional probabilities, including those with consequents that possibly explain their antecedents, obtains exactly the value its consequent will unconditionally get if its antecedent should come true. That way, the bonuses get accounted for quite automatically, and the divergence the bookie exploited is avoided from the start.[3]

We have now seen how the non-Bayesian can avoid dynamic Dutch book vulnerability. The story cannot end here, however. The reason is that as de Finetti (1937, p. 68 ff) has shown, you will be vulnerable to a *static* Dutch book unless your degrees of belief satisfy, for all φ and ψ, the following double equation:

$$(4.4) \qquad \Pr(\varphi \wedge \psi) = \Pr(\varphi)\Pr(\psi \mid \varphi) = \Pr(\psi)\Pr(\varphi \mid \psi).$$

3. You may have noticed that I tacitly assume that the bettor will come to know the data in the order she expects. Of course this assumption may be false. In cases in which it is, "hypothetical" and "real" updating may still fail to match. Because non-Bayesian update rules are not generally permutation invariant, it can happen that when $E_1, \ldots, E_n$ all turn out to be true but the bettor learns about their truth in an order other than she had expected, her posterior probability for H does not equal her initial probability for H conditional on $E_1 \wedge \cdots \wedge E_n$. Such a divergence would be innocuous, however. To construct a dynamic Dutch book the divergence must be *predictable*, and in a case of the kind considered, that requirement is not fulfilled. (If the bookie were to know that the data will come in some particular order different from the one the bettor expects, he *can* predict the divergence; but then again, not even full-blooded Bayesians are safe from bookies who know more than they do.) Shafer (1985) shows in a general fashion that dynamic Dutch book constructions are themselves based on assumptions concerning the order in which the data relevant to the bets will (or is expected to) be revealed.

As is readily verified, in our account the non-Bayesian will *not* generally satisfy (4.4); for instance, for a proponent of EXPL who follows the advice given above it does not hold that

$$\Pr(H_3)\Pr(\text{first toss heads} \mid H_3) = \Pr(\text{first toss heads})\Pr(H_3 \mid \text{first toss heads}),$$

where H_3 is our earlier hypothesis that asserts coin C to have perfect bias for heads. As a result, it might seem that the recommended strategy merely trades one problem for another.

As argued in Douven (1999), however, is actually trades a problem that Maher's look-before-you-leap principle *cannot* handle for one it *can* handle. Commenting on Maher's argument against the dynamic Dutch book theorem, Earman (1992, p. 49) writes: "In essence, the decision-theoretic message is to look before you leap. Such advice is just as valid in the synchronic setting as in the diachronic or multitemporal setting." This was written before Skyrms had shown that, although possibly valid, the advice to look before you leap was unhelpful in the diachronic case. However, in the synchronic case it is clearly unproblematic: simply check, whenever you want to make a bet, that it does not lead to a sure loss in combination with bets already made. Quite evidently, if a bettor follows this principle, no static Dutch book can be made against her. Perhaps it is not so simple to follow the advice in practice, but because it may not be too difficult for the non-Bayesian to see under what circumstances problems could arise (for instance, the EXPL user must pay extra attention when joint bets are proposed on propositions that are explanatorily dependent or when a bet is proposed on a proposition that is explanatorily related to a proposition on which she already has a bet), it seems reasonable to suppose that more workable principles are available. So, the final part of the strategy proposed in Douven (1999) consists of Maher's look-before-you-leap principle or some similar principle.[4,5]

It merits stressing that in that paper I did *not* claim that coherent non-Bayesian updating cannot be achieved by other means than the strategy presented previously, nor that that strategy is necessarily the best—nor even that rationality requires coherent updating. Rather, the point of that paper was to show that the property of coherence pertains not to update rules in isolation but to packages of both epistemic and decision-theoretic principles; that an update rule that in combination with certain principles leads to Dutch book vulnerability need not do so when combined with other principles; and that this is enough to undermine van Fraassen's main argument against abduction.[6]

---

4. In responding to the proposal as presented in Douven (1999), Tregear (2004) objected that the look-before-you-leap principle, when applied not only to betting but to forming judgments, would lead to cognitive paralysis in that an agent would then never make *any* judgments about any of her beliefs unless she were willing to risk ending up in an inconsistent state of mind. Writes Tregear (2004, pp. 513–514): "The conditions under which judgements of fairness are not to be made are just those conditions where the set of beliefs lead into a Dutch Book. These conditions are recognised by assessing the set of beliefs and determining whether they lead into a Dutch Book. Thus, in order to decide whether a set of beliefs is inconsistent, the beliefs must be judged." However, as I said in the main text, the situations in which a Dutch book looms for the advocate of a rule like EXPL have a generic characterization (roughly, as those in which a bet is proposed on a proposition that is explanatorily connected to some other proposition or propositions on which one has bet). Correspondingly, it is possible to "assess" a set of beliefs in order to avoid a possible inconsistency in one's belief state without having to determine degrees of belief in each of the separate beliefs. As a result, the proposed method for avoiding inconsistencies does not have to lead to a total cognitive paralysis.

5. Roche and Sober (2013) have sought to buttress the conclusion of van Fraassen's dynamic Dutch book argument by arguing that explanatory considerations cannot provide the kind of probabilistic boost that they are supposed to provide according to explanationists. Specifically, they argue that the evidence E for a hypothesis H screens off from H any fact F about explanatory connections between E and H, meaning that $\Pr(H \mid E \wedge F) = \Pr(H \mid E)$. It should be noted, however, that their argument for this claim proceeds from start to end on Bayesian assumptions. If giving bonuses to best explanations makes no sense from a Bayesian perspective, why should that bother explanationists? It should also be noted that, in view of (4.3), there may be no reading of $\Pr(H \mid E \wedge F)$ and $\Pr(H \mid E)$ that is neutral in the present debate (see also McCain & Poston, 2014; Climenhaga, 2017a; Lange, 2017).

6. Some might want to note at this juncture that van Fraassen has given an argument for the claim that explanatory considerations cannot have any positive epistemic import that is more general than the argument of the bad lot and the dynamic Dutch book argument. The more general argument starts from the premise that it is part of the meaning of "explanation" that if one theory is more explanatory than another, it must be more informative than the other.

More than twenty years after that paper, it seems to me that I was too accommodating to the Bayesian back then. For instance, Maher's look-before-you-leap principle is meant to be compatible with the Bayesian principle that rational decisions maximize expected utility. I no longer believe such compatibility is required. One could, for instance, think of it as a default–interventionist principle (to borrow a term from Evans, 2007b), according to which expected utility maximization is the default but can, under circumstances (e.g., in the face of a looming Dutch book), be overridden. In fact, in view of the literature on ecological rationality that inspired some of the subsequent chapters, it is even doubtful that expected utility maximization should be assumed the default. As Herbert Simon (1982) has argued, in reality people not only *tend* to be "satisficers" rather than maximizers, they *should* be, in view of limitations on their working memory and attentional resources. And building on Simon's insights, Rakefet Ackerman and colleagues propose that a choice is rational for a given agent in a given context to the extent that it facilitates achievement of the agent's goals within the computational constraints posed by the combination of agent and context (Ackerman et al., 2020).

Be this as it may, the question of whether a response to the dynamic Dutch book argument along the foregoing lines might be able to convince Bayesians has lost some of its interest, because most of them have become

Further, according to van Fraassen (1989, p. 192), it is "an elementary logical point that a more informative theory cannot be more likely to be true" so that "attempts to describe inductive or evidential support through features that require information (such as 'Inference to the Best Explanation') must either contradict themselves or equivocate." The elementary logical point, he claims, is "most [obvious] . . . in the paradigm case in which one theory is an extension of another: clearly the extension has more ways of being false" (1985, p. 280). However, the problem with this argument is that in any other kind of case the "elementary" point is not obvious at all. In particular, it is not obvious in those cases in which we are confronted with empirically equivalent rivals. For instance, it is entirely unclear in what sense the Special Theory of Relativity "has more ways of being false" than Lorentz's version of the ether theory (which as previously stated is inconsistent with, yet demonstrably empirically equivalent to, the Special Theory of Relativity). Yet, it seems that the Special Theory of Relativity is superior, qua explanation, to Lorentz's theory. If van Fraassen were to object that the former is not really more informative than the latter, or at any rate not more informative in the appropriate sense, whatever that is, then we should certainly refuse to grant the premise that in order to be more explanatory a theory must be more informative. See also Leeds (1994) for an excellent discussion of van Fraassen's general argument against explanatoriness as a confirmational virtue.

wary of that argument, and indeed of Dutch book arguments generally, on unrelated grounds. In particular, the view has spread in Bayesian quarters that the argument addresses the wrong kind of rationality. When we are concerned with the rationality of degrees of belief as well as with the rationality of how we change our degrees of belief over time, we are concerned with questions of *epistemic* rather than *practical* rationality. Given that being vulnerable to cunning bookies seems primarily a practical liability, the Dutch book arguments have been said to be beside the point (Joyce, 1998). In response to this, some have claimed that Dutch book vulnerability does flag an underlying epistemic defect: it is a manifestation of the fact that a person deems a single bet or series of bets as both fair and not fair and thus is in an inconsistent state of mind.[7] Note, however, that even if this claim is true for the (static) Ramsey–de Finetti argument, the point does not carry over to the dynamic Dutch book argument. An agent may be susceptible to engage in the kind of betting over time that figures in that argument without at any one time holding inconsistent views on the fairness of any bets. Naturally, after a learning event she may regard a bet as unfair that previously she regarded as fair, but the same would have been true had she been a Bayesian learner.

But even if the dynamic Dutch book argument is all about pragmatic rationality, that in itself need not be a reason to abandon it. More specifically, it seems to me that Bayesians could try to get some mileage out of the pragmatic turn that a number of epistemologists have recently been taking (see, e.g., Fantl & McGrath, 2009; Rinard, 2017). The epistemic status of a belief has traditionally been thought to depend solely on matters that bear on the truth of the belief, such as the quality of one's evidence or whether or not one is reliably connected to what the belief is about. But over the past years, various authors have argued that the epistemic status of a belief is inextricably bound up with the believer's practical situation, in particular, with what is at stake for her in believing correctly or incorrectly. Bayesians wishing to maintain the integrity of the Dutch book defense may not want to buy into any particular one of the arguments that have been advanced in favor of this "pragmatic encroachment view," as it has been called. However, it may suffice for them to argue that the mere prominence in current epistemology of the debate on pragmatic encroachment is enough to call into question the existence of the clear-cut divide between the epistemic and the pragmatic that the critics of

7. This point is to be found in Skyrms (1987).

the Dutch book arguments are presupposing. The reason is that, it may be said, if there were such a clear-cut divide, contributions to this debate should have gone down like lead balloons. And, Bayesians may conclude, if there is no such divide, then little is left of the charge that Dutch book arguments address the wrong type of rationality.[8]

Be this as it may, there is a deeper problem with the dynamic Dutch book argument, one that remains even if the pragmatic encroachment view is endorsed and even if the strategy for avoiding Dutch bookability described previously is, for some reason I happen to overlook, to be rejected. In essence, the argument is that updating by means of a non-Bayesian rule can cost you money that would stay in your pocket if you updated by Bayes's rule. But note that many things cost money. By dining out you can, and typically will, lose money, and in a foreseeable manner: waiters tend to present you with a bill not long after you have finished your meal. Still, that the loss can be prevented by fixing your own dinner at home is not a compelling argument against dining out; if all goes well, you get something in return that is worth the expense. More concretely, Dutch book invulnerability is only one among, in principle, an indefinite number of practical interests that people may have. What if EXPL or some other non-Bayesian rule serves other practical interests better than Bayes's rule? Could the Bayesian in that case still maintain that, in view of the dynamic Dutch book argument, Bayes's rule is the only rational update rule? Surely there exist practical goals whose achievement would more than make up for running the risk of being fleeced by a Dutch bookie—especially in view of the fact that Dutch bookies exist only as fictional characters in philosophers' tales!

From here on, one of the main themes is that abduction has practical advantages over Bayes's rule—not only practical advantages, in fact, but these are the ones we start with—and that it is a major oversight of Bayesians to focus strictly on what the costs of non-Bayesian updating may be and never to ask whether those costs could be worth incurring, given what one may get in return.

8. They might also want to point out that, next to a possible conceptual connection between the epistemic and the practical, there is a rather mundane connection, to wit, that the kind of epistemic projects that one can pursue depends on all sorts of practicalities. Whether I am in the position to pursue any research projects at all depends on how well I can organize my life, on whether I can secure a living as a researcher, on the degree to which I succeed in protecting my research time, and so on.

4.2 Abductive Reasoning and Practical Interests

When comparing Bayes's rule with one or more probabilistic explications of abduction, I will typically do so in the kind of statistical model that van Fraassen used in his presentation of the dynamic Dutch book argument. In particular, it will be supposed that an agent or agents are receiving binomial data (i.e., independent and identically distributed binary data), the data coming in one at a time and the agent's or agents' task being to estimate the probability of "success" (e.g., the probability that the patient survives, or that the treatment is effective, or that the coin lands heads). For present purposes, let the data consist of the outcomes of tosses with the same coin, where it is antecedently known that the bias of the coin (for concreteness, the probability of heads) can take eleven possible values, ranging from 0 (certain to land tails) to 1 (certain to land heads), in increments of .1. A priori, none of these values is more likely to be the true bias than any other, so before any tosses have been observed, it is reasonable to be no more confident that the bias takes one particular value than that it takes another. In other words, where H_i is the hypothesis that the bias for heads equals $i/10$, we are looking at the set of hypotheses $\{H_i\}_{0 \leqslant i \leqslant 10}$, with each of those having initially the same probability of $1/11$.

For now, I limit myself to a comparison between Bayes's rule and EXPL (meaning, again, the $c = 0.1$ instance of the schema labeled as such in ch. 2); subsequent comparisons include other probabilistic explications of abduction. Suppose then that a Bayesian and an EXPL user are watching the same sequence of coin tosses to which the previously mentioned model pertains. In particular, both have started with a flat probability distribution over the eleven bias hypotheses. Then they start updating their probabilities for these hypotheses via their respective rules as they see the sequence unfold. It has been said that a proposition "is assertable to the extent that it has high subjective probability for its assertor" (Jackson, 1979, p. 565).[9] If so, we may ask who—the Bayesian or the EXPL user—is more likely to first be in a position

9. More recent work on assertion suggests that this can be only approximately correct (see, e.g., Williamson, 2000, ch. 11; Douven, 2006, 2009; Goldberg, 2015; Kelp, 2016), but while some of that more recent work holds that assertion requires certainty and not just high probability, empirical evidence indicates that that requirement is too strong (Kneer, 2018, 2021; Reuter & Brössel, 2018; Marsili & Wiegmann, 2021).

to assert the truth about the bias of the coin, supposing the sequence is long enough for both eventually to be in that position.

To answer this question, I ran computer simulations of sequences of coin tosses that were long enough for both the Bayesian and the EXPL user to assign a high probability to the true bias hypothesis. "High probability" is a vague term, which I somewhat arbitrarily precisified as a probability greater than .9.[10] More exactly, for each bias hypothesis I ran 1,000 sequences of 500 (simulated) tosses of a coin with the bias as specified by the hypothesis, letting the Bayesian and the EXPL user update their probabilities for the various bias hypotheses and registering who was first in assigning a probability greater than the threshold value of .9 to the true hypothesis (called "convergence" here). I was also interested in who was the first to *stably* assign to the truth a probability above that threshold, that is, who was first in assigning a probability to the truth above .9 and *staying* above .9 for the remaining tosses in the sequence of 500 tosses ("stable convergence"). The results are graphically summarized by the plots shown in figure 4.1.[11]

This figure needs a little explanation. The individual plots (i.e., the twenty colored lines) are density plots, which can be interpreted as showing the relative likelihood that the first time a probability above the threshold is assigned (in the left panel of the figure) or is stably assigned (in the right panel) to a hypothesis occurs at a given toss. For instance, in the left panel the orange line for Bayes's rule peaks at about 100 and the orange line for EXPL peaks at about 90, both peaks being quite pronounced. This indicates that in the vast majority of the simulations with a coin with bias .2, the Bayesian converged after observing about 100 tosses, whereas in the vast majority of the same simulations, the EXPL user converged after observing about 90 tosses. Looking

10. I say that this specification of the threshold for high probability is *somewhat* arbitrary because .9 is often mentioned in the literature as a threshold for acceptability; see, for instance, Kaplan (1981, p. 308), Moser and Tlumak (1985, p. 128), and Kyburg (1990, p. 64). Relatedly, it has been recommended by legal scholars (e.g., Kagehiro & Stanton, 1985) as an explication of the notion of being beyond reasonable doubt, which is supposed to be a criterion for a conviction in the practice of criminal law. Note that one could also try other arguably more sophisticated "stopping rules" here; see, for instance, Kruschke (2013a).

11. Results for biases from .6 to .9 are not displayed in figure 4.1 because they are the mirror images (barring random noise) of the results shown—as they should be, given that a bias of p for heads is equivalent to a bias of $1 - p$ for tails. For the extreme biases, convergence and stable convergence happen at a fixed toss: for a coin that always comes up heads, Bayes's rule assigns a probability to the truth that is above the threshold after 23 tosses, EXPL after 14 tosses. Obviously, convergence and stable convergence coincide in the extreme cases.

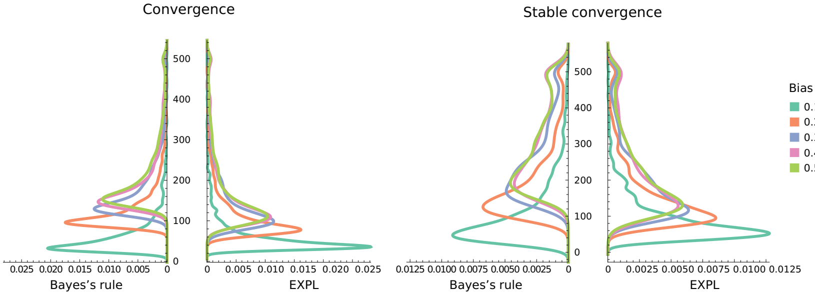

Figure 4.1: Paired density plots showing speed of convergence to the truth (*left*) and speed of stable convergence to the truth (*right*) for Bayes's rule and EXPL, for biases from .1 to .5. The *x*-axes show probability densities, and the *y*-axes show toss numbers.

in this way at the rest of the results, we see that for both convergence or at stable convergence, the peak (officially known as "the mode of the distribution") for EXPL lies typically quite a bit lower than the peak for Bayes's rule. More exactly, in a series of Mann–Whitney U tests comparing for each bias hypothesis the convergence results from Bayes's rule with those from EXPL, the latter rule each time came out as significantly faster (all ps < .0001) with the exception of the test for the bias = .1 hypothesis, in which there was no significant difference. Exactly the same outcome was obtained for the stable convergence results.

If this sounds unremarkable, imagine that the hypotheses concern some scientifically interesting quantity—such as the success rate of a medical treatment or the probability of depressive relapse—rather than the bias of a coin, and the tosses are observations or experiments aimed at determining that quantity. Which researcher would not want to use an update rule that increases her chances of being in a position to make public a scientific theory or a new medical treatment ahead of the (Bayesian) competition?

However, there is a possible downside to the apparent success of EXPL. The reason why the EXPL user so often beats the Bayesian in these simulations in assigning a high probability to the truth is that she is, in a clear sense, bolder in her responses to new information because she adds a bonus to the best explanation. It does not take much to see that this feature also makes the EXPL user more prone to assign a high probability to some false hypothesis:

a row of consecutive tosses producing a subsequence in which the relative frequency of heads starkly deviates from the probability for heads is more likely to push the EXPL user's probability for some false bias hypothesis over the .9 threshold than it is the more cautious Bayesian's. In other words, EXPL's greater speed can come at the expense of accuracy, so that even if EXPL may put one more rapidly in a position to assert truths, it may also put one more rapidly in a position that licenses the assertion of falsehoods.

To make this more formal, note that we can associate with each EXPL update a probability that the explanation bonus in its entirety is assigned to a false hypothesis and that this probability can be considerable. Suppose, for instance, that in our present model a coin that has, unbeknownst to the updater, an actual bias of .5 (so a fair coin) is tossed ten times. Then the probability that after the tenth toss EXPL will assign the bonus to the *true* bias hypothesis is the probability that there will be exactly five heads in the ten tosses, which can be calculated to be approximately .25. Consequently, the probability that after ten tosses a *false* hypothesis receives the bonus is about .75.

This probability goes down as the coin is tossed more often, and it is also lower for more heavily biased coins (whether toward heads or toward tails). For example, if the bias were .2 (equivalently, .8) and we were to consider the possibility that the bonus gets assigned to a false hypothesis after the 100th toss, the probability of that actually happening would be about .17; for the 1,000th toss, it is smaller than .0001. More generally, the probability of a wrong assignment goes to 0 (regardless of the bias of the coin) as the number of tosses taken into consideration goes to infinity.[12] But if we are interested in which rule should be used, not by idealized agents but by us ordinary mortals, then such limit behavior carries little significance.

---

12. To be entirely exact, we first note that the normal distribution gives a good approximation to the binomial distribution provided both $np \geqslant 10$ and $n(1-p) \geqslant 10$, where p is the success probability (the bias) and n the number of trials (tosses, in our case). Excluding the extreme bias hypotheses (always heads, always tails), these conditions are satisfied for even moderately large numbers of tosses. Then the probability that the bonus is assigned to a false hypothesis after n tosses of a coin with bias p (where $0 < p < 1$) is approximately

$$1 - \operatorname{erf}\left(\frac{.05}{\sqrt{2}} \cdot \sqrt{\frac{n}{p-p^2}}\right),$$

where erf is the error function, which is defined: $\operatorname{erf}(x) = 2/\sqrt{\pi}\int_0^x e^{-t^2}\,dt$. The given expression can be shown to converge to 0 as n tends to infinity.

Consider, then, an arguably fairer comparison between EXPL and Bayes's rule. A magician has a set of eleven coins in his pocket: one coin with a perfect bias for heads, one with a bias of .1 for heads, one with a bias of .2 for heads, and so on. He entices a Bayesian and an EXPL user into playing a game against one another. In this game, the magician picks one of the coins and starts tossing it, showing each outcome to the two players. A player scores a point by raising her hand and asserting the truth about the bias of the coin. However, if the player gets the bias wrong, the point goes to the other player. If both players raise their hands after the same update and both are right about the bias, they both receive a point; if they are both wrong, neither receives a point; if one is right and the other wrong, the former receives two points. Once a player has asserted a bias hypothesis—whether correctly or incorrectly—the magician puts the coin back into his pocket and again picks a (not necessarily different) coin. Then everything starts all over. This procedure is repeated 100 times, after which the player with the highest score is declared the winner. Note that this game seems to treat the possible advantages and disadvantages of EXPL updating in an evenhanded manner: because of her boldness, the EXPL user might often be first to identify the truth, but by the same token she might earn quite a few points for her opponent. Who is more likely to win the game, the EXPL user or the Bayesian? Or are their chances of winning equal?

To answer these questions, I ran 1,000 simulations of the game. In these simulations, the EXPL user *always* won and typically did so by a wide margin. To be more exact, the mean of the differences between the EXPL user's score and the Bayesian's score was 27 (± 9); the smallest difference between the two players' score at the end of a game was 5, and the greatest was 52.

This is not meant to demonstrate the superiority of EXPL. Sets of trick coins of the variety required for playing this game are hardly more common than Dutch bookies. Moreover, nothing said here excludes the possibility that there are circumstances under which Bayesian updating offers practical advantages more consequential than Dutch book invulnerability. Instead, the point is to draw attention to a hidden assumption in the Dutch book approach, to wit, that Dutch book invulnerability trumps all other practical goals we may have. Absent an argument for this assumption, the Dutch book approach does not pose a threat to EXPL or any other probabilistic explication of abduction, even granting that the pragmatic pervades the epistemic.

4.3 Abduction and Our Epistemic Goal

The relation between the pragmatic and the epistemic is immaterial to an approach to justifying Bayes's rule and more generally the tenets of Bayesianism that has emerged over the past twenty years or so. Various theorists, disconcerted by the apparently pragmatic focus of the Dutch book arguments, have sought to justify the Bayesian tenets in (supposedly) strictly epistemic terms. Specifically, they have sought to show that those tenets are most conducive to the achievement of our epistemic goal, as spelled out for graded beliefs.

This development owes much to Joyce (1998), in which a first attempt is made to formulate an epistemic goal in terms of graded beliefs and which argues that for every degrees-of-belief function that violates the probability axioms, there is a degrees-of-belief function that obeys those axioms and is closer to the epistemic goal. Joyce thought that this vindicated the synchronic part of Bayesianism, according to which degrees of belief ought to be probabilities, though in Joyce (2009) he admits that the earlier argument was not airtight. For present purposes, however, what matters is mostly Joyce's conception of our epistemic goal as pertaining to graded beliefs.

Among mainstream epistemologists, who tend to be concerned first and foremost with categorical beliefs, it is almost universally held that our epistemic goal is to believe all that is true and only what is true (see, e.g., Lehrer, 1974, p. 202; BonJour, 1985, p. 8; Foley, 1993, p. 19). Until recently, there was no equally clear conception of an epistemic goal in terms of graded beliefs. This changed when Joyce proposed an epistemic goal for graded beliefs that was explicitly meant to be analogous to the aforementioned epistemic goal for categorical beliefs. The proposal is that our system of degrees of belief ought to be *gradationally accurate*, and in fact as gradationally accurate as any other system of degrees of belief that we might adopt. The notion of gradational accuracy is defined in technical terms, in particular in terms of so-called scoring rules, briefly mentioned in section 3.5 and discussed in detail in the next chapter. But the basic intuition underlying it is clear enough, to wit, that the higher one's degree of belief in a true proposition is, the more accurate one is, *ceteris paribus*, and also the lower one's degree of belief in a false proposition is, the more accurate one is, *ceteris paribus*.

In the same vein, Hannes Leitgeb and Richard Pettigrew (2010a, 2010b) have tried to give a nonpragmatic justification not only of the synchronic part of Bayesianism but also of its diachronic part, that is, of Bayes's rule.[13] They share Joyce's view of our epistemic goal, which they put as follows (2010a, p. 202):

Accuracy: An epistemic agent ought to approximate the truth. In other words, she ought to minimize her inaccuracy.

Also like Joyce, Leitgeb and Pettigrew make the notion of accuracy precise by reference to scoring rules. They then argue that Bayesian updating minimizes expected inaccuracy. More exactly, they argue that if, and only if, an agent updates via Bayes's rule, she minimizes the expected inaccuracy of her post-update probability function, where the expectation is minimal according to her pre-update probability function (Leitgeb & Pettigrew, 2010b, pp. 249–250).[14]

What Leitgeb and Pettigrew aim to show for Bayes's rule is in an important respect *dis*analogous with what Joyce aims to show for the synchronic part of Bayesianism. While Joyce argues that an agent whose degrees of belief are not probabilities fails to minimize *actual* inaccuracy, Leitgeb and Pettigrew argue that an agent who updates by a rule other than Bayes's fails to minimize *expected* inaccuracy. Nothing that they say precludes the possibility that some non-Bayesian rule outperforms Bayes's rule with respect to actual inaccuracy minimization. More importantly, Leitgeb and Pettigrew are concerned only with the inaccuracy of the immediate post-update belief state. But an equally—if not more—legitimate question one may ask about update rules is which of them will lead to the most accurate belief state *in the long run*. Indeed, given that the epistemic goal for graded beliefs is supposed to be analogous to the epistemic goal as discussed by mainstream epistemologists, the long-run question is the more natural one to ask. After all, in mainstream epistemology the epistemic goal is mostly conceived as an *ultimate* goal to which all our epistemic endeavors are geared and in light of which they are to be assessed.[15]

13. For similar attempts at vindicating Bayes's rule, see Rosenkrantz (1992) and Greaves and Wallace (2006).

14. To be still more specific about the expectation: it is the weighted average, taken over all worlds w that the agent deems epistemically possible before the update, of the inaccuracies of the post-update probability function if w were the actual world, with the pre-update probabilities for the worlds w serving as weights.

15. See, for instance, Latus (2000, p. 29), who writes that in asking what our epistemic goal is "we are interested in what we want the overall result of our various ways of 'finding

To see just how unconnected the two are, let us for a moment set aside the idealizations commonly made in discussion about updating and note that, in science, we rarely just *happen* across useful data. Typically, we must actively search for data, and in the many areas of science that rely on experimentation even *produce* our data. Because our time is limited, as is our funding, we constantly have to make decisions as to which instruments (telescopes, microscopes, etc.) to construct, which expeditions to undertake, which experiments to run, and so on. Such decisions will be informed by which hypotheses we deem most promising. Had we deemed hypothesis H promising, and had we wanted to compare that with the hypothesis currently dominant in our field, we might have run a different set of experiments from those that we actually did, given that in fact we deemed H' more promising than H and were mainly interested in comparing H' with the received doctrine. Which hypothesis or hypotheses that we deem most promising and most worthy of spending our limited resources on, will at least in part depend on how probable they appear to us, compared with their most direct rivals. If, for example, a Bayesian update makes H more probable than H', whereas the opposite will be the case if we update via some non-Bayesian update rule, then our decision to use one of these rules may put us on a very different research path with very different downstream consequences from the path encountered if we had decided to use the other rule. Which of these paths will eventually lead us to the more accurate representation of the world will have nothing to do with which of the rules minimizes expected inaccuracy of the piece of evidence now lying before us.

The general problem for the inaccuracy-minimization approach to which this points is that Accuracy permits a number of different interpretations. For instance, it can be interpreted as demanding that every single update minimize expected inaccuracy, as Leitgeb and Pettigrew do, or that every update minimize actual inaccuracy, or that every update be aimed at realizing the long-term project of coming to a minimally inaccurate representation of the world, even if individual updates do not always minimize inaccuracy

out' . . . to be." Goldman (2010, section 2) also argues that (what he calls) epistemic systems are to be evaluated in terms of their accuracy, which he phrases in terms of degrees of truth-possession, which in turn is defined by reference to a scoring rule. However, he is very clear that what he has in mind is long-term accuracy: "Epistemic system E is better than epistemic system E^* iff conformity to E would produce (in the long run) a higher total amount of degrees of truth-possession than conformity to E^* would produce" (p. 194).

or expected inaccuracy. This is problematic because none of these can offhand be rejected as not a legitimate epistemic goal, and that it is by no means obvious that if rule R is most conducive to the realization of one goal and rule R' is most conducive to the realization of a second, then $R = R'$, or that at least the rules are equivalent in that they always yield the same output, given the same input.

In fact, there is not just the question of whether the goal of inaccuracy minimization is meant to pertain (only) to the "next update" or is rather meant to be a long-term project. Even if we assume that it is conceived as a long-term project, there are still further distinctions that can be made. One notable distinction concerns whether we should aim to have a minimally inaccurate probability function in the long run, *however* far in the future that may be, or whether it is better to have a moderately accurate probability function in the shorter run. It is certainly imaginable that, for purely epistemic reasons, one might opt for the latter, supposing that there is a choice to be made. We are, after all, *curious* about the truth (Hempel, 1965, p. 333; Foley, 1987, p. 11), and curiosity typically comes with a sense of urgency. As a result, we might prefer an update rule that is more likely to take us fairly close to the truth in a reasonably short time span over one that is more likely to take us *extremely* close to the truth in the long run but less likely to take us even fairly close to the truth in, for example, the medium-long run.[16,17]

16. Douven (2010, section 4) makes a parallel point in terms of the qualitative notion of truth approximation for update procedures that take into account peer opinions.

17. Briggs and Pettigrew (2020) present an inaccuracy-minimization argument assuming that we aim to be maximally accurate over a lifetime. They do not argue for this assumption nor say how it relates to the discussion of our epistemic goal in traditional epistemology (which Joyce took as a starting point). On its face, the assumption is implausible, if only because, realistically speaking, it would make sense to sacrifice some accuracy along the way for the benefit of having a *highly* accurate belief state in the somewhat longer run. See, for instance, Galison (1997) on how scientists with partly different backgrounds (e.g., theoretical physicists and engineers) must sometimes be willing to settle on a kind of common language (a *pidgin*, as Galison's refers to it) in order to be able to communicate at all, even though the language allows none of the collaborators to be, by their own lights, entirely accurate. All the while, the ultimate goal of the collaborative effort is that the accuracy of our representation of this or that part of the world be enhanced. Those of us who are educators will be familiar with a similar phenomenon. In explaining to freshmen concepts or theories or standpoints that are new to them, there is sometimes no other route to a first understanding than to cut some corners, skip some details, and omit some qualifications, all for the good of *ultimately* achieving an exact understanding of those concepts, theories, or standpoints some time later in the curriculum. More generally, note that at a minimum it is not a priori that aiming to be as accurate as

We could go on for some time disambiguating Accuracy in this way, if only because speed and accuracy of convergence to the truth can be traded off in an indefinite number of ways. However, our aim here is not to catalog all reasonable precisifications of Accuracy but rather to show that the inaccuracy-minimization argument poses no real threat to abduction. Leitgeb and Pettigrew argue that Bayesian updating minimizes the expected inaccuracy of our next belief state relative to our current belief state and thus may be said to be most conducive to *one* particular epistemic goal. Even granting this, it does little to undermine abduction as long as it is understood that abduction (or a particular explication of it) is most conducive to some other epistemic goal or goals, where possibly the different epistemic goals can both (or all) be regarded as explications of the same broad idea that we should strive for a maximally accurate belief system.

At this juncture, Bayesians might try to argue that expected inaccuracy minimization of our next belief state as judged from our present one trumps any other epistemic goal that we may have. The prospects of this tactic are about as bleak as the prospects of showing that Dutch book invulnerability trumps any other practical goal that we may have; bleaker still, for it would seem absurd to claim that it is epistemically more important to have an update rule that minimizes *expected* inaccuracy than to have one that minimizes *actual* inaccuracy. This seems especially true in view of Robbie Williams's (2012, p. 835) observation that expected inaccuracy minimization defenses of Bayes's rule raise the question of "why [you should] trust an *outdated* belief state to tell you how to fix your beliefs now you have new information."

---

possible overall offers the best guarantee of ending in a highly accurate belief state. And Briggs and Pettigrew make no attempt to show that there is such a guarantee. There is an even more fundamental problem with their argument. The argument is supposed to show that a non-Bayesian updater could come to see that there is an alternative prior from which she might have started from that in combination with the strict use of Bayes's rule would have made her more accurate over the course of her epistemic life than she now is, because she did *not* start from that prior and is committed to a non-Bayesian update rule. Aside from concerns about the idealizing assumptions that go into their argument, it is to be noted that for all that Briggs and Pettigrew show, adopting the alternative prior might from the person's current perspective appear to amount to "believing crazy stuff." More generally, why should it bother the non-Bayesian that she could have been more accurate than she actually is if only she had started from degrees of belief which she simply never had and which perhaps would never have seemed reasonable to her? The point of Briggs and Pettigrew's argument is that a Bayesian would never find herself in the kind of situation just sketched, but to the extent that it is unclear what is so dreadful about that kind of situation, the argument fails.

Alternatively, Bayesians might try to show that Bayes's rule does best with respect to *all* epistemic goals. This may be hard to establish in an a priori manner, supposing that among the relevant goals are ones whose realization requires actual inaccuracy minimization. How well a rule does in terms of truth approximation may plausibly depend not only on the rule but also on which world is the actual one. For example, even if the advocates of abduction have so far failed to produce a reason for believing that best explanations tend to be true, we might happen to inhabit a world in which—perhaps as a brute fact—best explanations do tend to be true. Worse yet (for the Bayesians), simulations that I performed, which compare Bayes's rule and EXPL in terms of gradational inaccuracy in the setting of our earlier statistical model, give reason to doubt that Bayes's rule is the most conducive to the realization of every epistemic goal.

In these simulations, gradational inaccuracy is measured by means of the Brier scoring rule. Using the same statistical model that was used in the previous simulations, I again ran 1,000 simulations of a sequence of 1,000 coin tosses for each possible bias. Also as previously, the Bayesian and the EXPL user updated their degrees of belief after each toss. I registered in how many of those 1,000 simulations the EXPL user incurred a lower Brier penalty than the Bayesian at various reference points. The outcomes of these simulations are displayed in table 4.2. They show that, at each reference point, EXPL is usually the winner (it incurs the lowest penalty). Hence, at least in this kind of context, EXPL appears a better choice than Bayes's rule.

Or maybe not—for I also calculated averages, taken over the 1,000 simulations for each bias value, of the penalties incurred by the EXPL user and the Bayesian at the designated reference points. Table 4.2 gives for each of these points the mean of the EXPL penalties minus the mean of the Bayes's rule penalties (in parentheses). As shown by those numbers, if there is a difference between the two means, it is always in favor of Bayes's rule. The reason for this is manifest from the spread of the simulation outcomes (not represented here), which shows that, although EXPL wins in most instances, it is typically by a relatively small margin, whereas in some of the runs in which EXPL loses, it incurs considerably greater penalties than does Bayes's rule. This in turn is due to the previously noted fact that the EXPL user reacts in a bolder fashion to the evidence than does the Bayesian and is therefore more easily led astray by rows of tosses that produce subsequences with deviating relative frequencies. Indeed, inspection of individual runs shows that such subsequences,

Table 4.2: Results of 1,000 simulations of sequences of 1,000 coin tosses for each bias value: the columns give the number of simulations in which EXPL incurred a *lower* Brier score than Bayes's rule after 100, 250, . . . , tosses; in parentheses, the mean of the EXPL penalties minus the mean of the Bayes's rule penalties, taken over the 1000 simulations (rounded to two decimal places).

Bias	100	250	500	750	1,000
.0	1,000 (.00)	1,000 (.00)	1,000 (.00)	0 (.00)	0 (.00)
.1	923 (.02)	993 (.00)	1,000 (.00)	1,000 (.00)	1,000 (.00)
.2	748 (.10)	937 (.02)	991 (.00)	1,000 (.00)	1,000 (.00)
.3	717 (.11)	918 (.03)	984 (.00)	998 (.00)	999 (.00)
.4	690 (.11)	865 (.06)	968 (.01)	990 (.00)	997 (.00)
.5	679 (.12)	879 (.05)	972 (.00)	992 (.00)	998 (.00)
.6	660 (.14)	904 (.03)	977 (.00)	992 (.00)	995 (.00)
.7	698 (.12)	907 (.03)	978 (.01)	995 (.00)	1,000 (.00)
.8	754 (.09)	947 (.02)	996 (.00)	999 (.00)	999 (.00)
.9	910 (.04)	990 (.00)	1,000 (.00)	1,000 (.00)	1,000 (.00)
1.0	1,000 (.00)	1,000 (.00)	1,000 (.00)	0 (.00)	0 (.00)

while they push both agents off the right track, impact the EXPL user much more severely than the Bayesian, and that in turn has a noticeable effect on the differences between the Brier penalties that they incur.

From these simulations, each of the two update rules under consideration comes out doing better in different respects; this further buttresses the point of this section. There seems to be no clear answer to the question of whether it is better, epistemically speaking, to use an update rule that in general achieves greater accuracy than other update rules, even if typically not *much* greater accuracy; or to use an update rule that is less likely than another update rule to ever make one *vastly* inaccurate, even though it *typically* makes one *somewhat* more inaccurate. Naturally, we would prefer a rule that offered the best of both worlds, or better yet, of *all* worlds—in other words, a rule that was most conducive to any epistemic goal that we might have. But there may be no rule that fits this bill, and in any event, the previously described simulations give reason to believe that Bayes's rule is not that rule. To be sure, these simulations do not warrant a more positive verdict about EXPL either. But the aim of this section is not to make a case for EXPL or even for the broader idea of abduction. Rather, it is to highlight and question a hidden assumption in the inaccuracy-minimization defense of Bayes's rule, to wit, that there is but one way in which an update rule can be said to minimize inaccuracy, or at least only one that matters epistemically.

Bayesians might seem to have an easy response to the foregoing. Previously in this chapter, it was argued that what makes the dynamic Dutch book argument objectionable may be not so much the fact that the argument addresses the wrong kind of rationality—practical rationality—but rather that it focuses on just one possible practical concern to the neglect of all others, apparently without good reason. If pragmatic considerations can be legitimately invoked in the present discussion, then one might try to argue that even if the goal or goals that Bayes's rule serves best are not privileged from an epistemic viewpoint, they *are* privileged from a practical one. Specifically, one might try to argue that if the results of the above simulations have some general validity and do not hold only for the particular type of statistical model considered here, then the fact that updating via Bayes's rule leads on average to a lower penalty is enough to justify this rule. And indeed, it does follow from the said fact that were one to keep betting on consecutive tosses with a given coin, each time posting betting odds in accordance with one's degrees of belief at that time, then one would maximize one's expected payoff by updating via Bayes's rule.

But it would be false to think that payoffs and scoring rule penalties must always be so strictly related as in this example. To see how they can come apart, suppose that you are the owner of one of two engineering firms in a city. The two firms are in ongoing competition for contracts from the local authorities. As a rule, a contract goes to the firm that proposes the better solution to whatever the engineering problem is that the authorities want to be solved. Given that the engineers employed by your firm are about as competent as the engineers employed by the other firm, which firm comes up with the better plan typically depends on the accuracy of the information on which they base their proposals. Note that under these circumstances, being able to base one's proposal on information that is just slightly more accurate than the information available to the competition is enough to get the contract. By contrast, if the available information is less accurate, it is immaterial whether it is only slightly less accurate or much less accurate: one will not get the contract either way; in this case, a miss is as good as a mile.[18] Again supposing the results of our simulations to have some general validity, they suggest that if you and your team update via EXPL and the competition updates via Bayes's rule, then you are likely to get the vast majority of the contracts. Once in a while—very rarely—a contract will go to the other firm. When it does, the information on which your proposal was based was probably much more inaccurate than the information on which the rival plan was based. However, this has no financial consequences beyond the fact that you miss out on the contract, which would also have happened if the information available to you had only been slightly less accurate than the information available to the other firm.

At this point, it has come to appear that there are circumstances under which it is better to update via EXPL, practically speaking, given that it greatly increases the chance of ending up with a (perhaps only slightly) more accurate belief system than one would have ended up with had one updated via Bayes's rule. Naturally, this is only one sense in which one can aim to minimize one's inaccuracy, and the foregoing considerations in fact suggest that from a pragmatic viewpoint, there may be no single best update rule. Depending on the circumstances and on what exactly one's interests are, Bayes's rule may serve one's interests best, or EXPL may do so, or perhaps yet another rule may do so. In epistemology, contextual approaches have been much in the limelight lately. I am not aware of anyone suggesting that which update rule use by may depend on context, and yet that is one of the main claims of this book.

Consider the previously described situation in which you own one of two engineering firms. There, EXPL clearly seemed the better option when the choice was between that rule and Bayes's rule. Or consider an example from Douven (2017b), in which you are wondering about the source of the persistent rumor that Mozart did not die from natural causes but was poisoned, a rumor that is almost certainly false (Zegers, Weigl, & Steptoe, 2009). Actually it was Mozart himself who started the rumor that he had been poisoned, shortly before his death (Robbins Landon, 1999, p. 148 ff). If you did not know this already and if I had not told you, you might out of sheer curiosity have wanted to find out about the source of the rumor. And if you were really curious, you might have wanted to find out quickly. But if in the course of

18. We shall see more of this in Chapter 6.

your investigations of the matter, before you eventually found out the truth, you had come to mistakenly hold that it was Mozart's wife who started the rumor or that is was one of the newspapers that first reported his death, that might have done little to no harm: accuracy does not matter much in this kind of case, in which being mistaken is unlikely to have any dire consequences, if at all. Here, too, EXPL might be the rule of choice, given how much faster than Bayes's rule it is, even if it is on average not quite as accurate as that rule.

By contrast, to use another example from Douven (2017b), suppose that you are the producer of roulette tables and want your tables to be extensively tested before they leave the factory. Suppose further that the testing proceeds by rolling a ball 1,000 times and registering whether it lands on red or on black. You do not want to sell a table with a bias, as that might damage the reputation of your brand. Here, you might want to use Bayes's rule for updating your beliefs about the fairness of a table on the basis of the outcomes of the rolls of the ball. Accuracy matters greatly here; speed is not much a concern.

This suggestion of context dependence is explored further in subsequent chapters. In fact, we go beyond the simulations presented in this chapter along two different lines. In chapter 6, we consider more realistic situations and compare in them a greater number of update rules than we have done here. In that chapter we also embed the question of which update rule to use in a broader discussion about the notion of rationality, paying special attention to recent developments in cognitive psychology and cognitive science that indeed make rationality an inherently context-bound notion.

In chapter 7, we address another limitation of the simulations described in this chapter, to wit, that they ignore that we typically learn as members of a collective. As social epistemologists have been emphasizing for some time, we will obtain an incomplete and possibly distorted picture of epistemology if we fail to recognize the extent to which our beliefs are shaped by, and even rooted in, our interactions with others. Taking the lessons from social epistemology seriously, chapter 7 looks at communities of agents who update their opinions by taking into account both the evidence that they receive directly from the world and the opinions of their epistemic peers.

4.4 Summary

We have considered the currently prevailing arguments for the claim that when it comes to rational update rules, Bayes's rule is the only game in town. We have argued that neither the dynamic Dutch book approach nor the inaccuracy-minimization approach succeeds in challenging the rationality of updating via probabilistic versions of abduction. We proceeded by laying bare some unsupported hidden assumptions of these arguments. At the most general level, the problematic (tacit) assumption is that Dutch book invulnerability and / or expected next-step inaccuracy minimization are *all* that we should care about when it comes to judging update rules. As briefly mentioned in chapter 1, if you are an epistemic consequentialist and hold that update rules should be judged by their fruits, then do not cherry-pick: judge the rules by *all* their fruits, the sweet and the sour.

With respect to the dynamic Dutch book argument: if updating via Bayes's rule has certain practical advantages in comparison with other rules, a version of EXPL may have different, possibly more important, practical advantages than does Bayes's rule. And that is to grant that non-Bayesian updating does imply Dutch book vulnerability—something that may have seemed plausible only because of another, more specific hidden assumption taken aboard by Bayesians, to wit, that update rules are to be judged on their own and not in combination with other epistemic or decision-theoretic principles with which they must interact if they are to have any consequences at all.

Equally, if Bayes's rule minimizes inaccuracy in some sense, a version of EXPL may minimize inaccuracy in another sense, in which for all anyone has shown, no single one of these senses can be said to be privileged in that it captures *the* epistemic goal. That expected next-step inaccuracy minimization trumps any other kind of inaccuracy minimization, or indeed any other possible benefit of non-Bayesian updating, is again an assumption for which no Bayesian has cared to argue.

Note that empirically the Bayesian claim of inaccuracy minimization was implausible already in light of the re-analysis of the data from Douven and Schupbach (2015a), discussed in section 3.5, which showed that Douven and Schupbach's participants tended to have more accurate degrees of belief the more weight they had given, in their updating, to explanatory considerations. We can further mention Nathan Berg, Guido Biele, and Gerd Gigerenzer's (2016), for which 125 academics were asked for their degrees of belief concerning the usefulness of prostate cancer screening. These authors found literally zero correlation between the degree to which their participants conformed to Bayesian principles and the actual accuracy of their degrees of belief. Even more ironic, they found that the participant who conformed most closely to those principles was also the one whose degrees of belief were most inaccurate.

It is one thing to argue that the criticisms leveled against non-Bayesian updating in general, and probabilistic versions of abduction in particular, are baseless, but it is another to defend specific forms of non-Bayesian updating, such as, most relevantly, updating via EXPL and kindred rules. Much of the remainder of this book is devoted to defending forms of probabilistic abduction. First, however, I would like to point out yet another hidden assumption of the inaccuracy-minimization argument. We saw that this argument was explicitly meant to be *epistemic* in nature and thereby to improve on the older defense of Bayes's rule in terms of dynamic Dutch book vulnerability, which many had come to regard as posing only a *pragmatic* problem. We also saw that the inaccuracy-minimization argument relies crucially on a scoring rule as a formal measure of inaccuracy. For the argument to be purely epistemic in nature—as it is supposed to be—scoring must be a purely epistemic matter. The next chapter challenges this assumption, arguing that there are inextricably pragmatic aspects to scoring: how best to measure inaccuracy may vary from one context to another, and the decision to use this or that scoring rule may partly depend on what our practical concerns are. This also leads me to challenge the claim, taken for granted so far, that Bayes's rule minimizes at least expected next-step inaccuracy.