

OBJECTIVE BAYESIANISM, BAYESIAN CONDITIONALISATION AND VOLUNTARISM

Jon Williamson

To appear in *Synthese*
Draft of April 2, 2009

ABSTRACT

Objective Bayesianism has been criticised on the grounds that objective Bayesian updating, which on a finite outcome space appeals to the maximum entropy principle, differs from Bayesian conditionalisation. The main task of this paper is to show that this objection backfires: the difference between the two forms of updating reflects negatively on Bayesian conditionalisation rather than on objective Bayesian updating. The paper also reviews some existing criticisms and justifications of conditionalisation, arguing in particular that the diachronic Dutch book justification fails because diachronic Dutch book arguments are subject to a reductio: in certain circumstances one can Dutch book an agent *however she changes her degrees of belief*.

One may also criticise objective Bayesianism on the grounds that its norms are not compulsory but voluntary, the result of a stance. It is argued that this second objection also misses the mark, since objective Bayesian norms are tied up in the very notion of degrees of belief.

CONTENTS

§1	Introduction	2
§2	Objective Bayesianism	2
§3	Objective and Subjective Bayesian Updating	5
§4	Four Kinds of Incompatibility	7
§5	Criticisms of Conditionalisation	11
§6	A Dutch Book for Conditionalisation?	14
§7	Conditionalisation from Conservativity?	15
§8	Summary	16
§9	Voluntarism and Stances	16

§1

INTRODUCTION

Objective Bayesianism is construed here as an epistemological thesis which implies that an agent has little or no choice as to how strongly she should believe a proposition—these rational degrees of belief are forced upon her, determined by the extent and limitations of her evidence (§2). This is a controversial thesis and it has been criticised on many fronts; see [Williamson \(2009\)](#) for an overview of the challenges that face this position. The goal of this paper is to respond to two criticisms of objective Bayesianism that are connected to themes in the writings of Bas van Fraassen. First, objective Bayesianism has often been dismissed on account of differences between its form of updating and Bayesian conditionalisation (§3). Such differences are typically taken to count against objective Bayesianism and are exemplified in van Fraassen’s Judy Benjamin story ([van Fraassen, 1981](#); [van Fraassen et al., 1986](#)). In contrast I argue in §4 that such differences count against Bayesian conditionalisation and that conditionalisation can only safely be used where its results agree with those of objective Bayesian updating. In §5 I rehearse several well-known criticisms of conditionalisation. I then go on to dismiss two justifications of conditionalisation. I consider the dynamic Dutch book argument for conditionalisation in §6, and offer a reductio against dynamic Dutch book arguments in general. In §7 I examine the argument for conditionalisation that appeals to stability of conditional probabilities, but show that this argument fails to offer a justification of conditionalisation where it disagrees with objective Bayesian updating. Overall, conditionalisation compares poorly with objective Bayesian updating (§8).

The paper also addresses a second criticism of objective Bayesianism. It may be argued that one is not *bound* by the norms of objective Bayesianism but can choose whether or not to follow them and remain rational either way. Under this view, motivated by [van Fraassen \(2002\)](#), objective Bayesianism is a voluntary stance rather than a set of compulsory laws of thought. In §9 I maintain that one has to take a stance as to what one means when one talks of ‘belief’, but that once the semantic question has been agreed, all parties are bound by the consequences of the chosen interpretation. Moreover, under a sensible interpretation of belief, degrees of belief are bound by the norms of objective Bayesianism.

§2

OBJECTIVE BAYESIANISM

Objective Bayesianism incorporates the following three principles:

PROBABILITY: An agent’s degrees of belief should be representable by probabilities. Thus for example your degree of belief that it will not rain here tomorrow and your degree of belief that it will rain here tomorrow should sum to 1.

CALIBRATION: An agent’s degrees of belief should be appropriately constrained by empirical evidence. If you know just that days like today are followed by rain somewhere between 70% and 80% of the time, then your degree

of belief that it will rain tomorrow should lie somewhere in the interval $[0.7, 0.8]$.

EQUIVOCATION: An agent's degrees of belief should be as middling as these constraints permit. Given the above evidence you should equivocate as far as possible between rain and no rain; i.e., you should believe it will rain to degree 0.7.¹

The ideas behind these three tenets were already present around the turn of the 18th century in the writings of Jakob Bernoulli. Bernoulli maintained that probability and degree of certainty were to be identified: '*Probability*, indeed, is degree of certainty, and differs from the latter as a part differs from the whole' (Bernoulli, 1713, p. 211). He argued that in the absence of evidence one should equivocate, but that evidence should take precedence over equivocation:

For example, three ships set sail from port. After some time it is reported that one of them has perished by shipwreck. Which do we conjecture it to be? If I considered only the number of ships, I might conclude that misfortune might equally befall any of them. But since I remember that one of them was more eaten away by decay and age than the others, that it was badly equipped with sails and sail-yards, and also that it was commanded by a new and inexperienced skipper, I judge that it is surely more probable that this one perished than the others. (Bernoulli, 1713, p. 215)

After Bernoulli, writers on probability attempted to dispense with one or more of these principles. Thus proponents of the classical and logical views of probability—including Laplace, Keynes, and Carnap—focussed on Probability and Equivocation at the expense of Calibration. Proponents of the frequency view—e.g., Venn, von Mises, Reichenbach—focussed on the empirical probability of Calibration at the expense of the epistemological content of Probability and Equivocation. Proponents of the subjective view—e.g., Ramsey, de Finetti—focussed on Probability at the expense of Equivocation and, in de Finetti's case, Calibration.

It wasn't until the work of Jaynes and Rosenkrantz that these principles were reunited, in the form of contemporary objective Bayesianism (see, e.g., Rosenkrantz, 1977; Jaynes, 2003). This is 'Bayesianism' in the sense that probabilities are construed as degrees of belief, as in Bayes (1764), *not* in the sense that probabilities are updated by Bayesian conditionalisation: as will be explained in §3, the objective Bayesian does not need to advocate Bayesian conditionalisation, and indeed the version of objective Bayesianism presented here explicitly rejects conditionalisation. It is objective in the sense that the

¹Note that Equivocation presupposes that Calibration be properly applied first. In particular, in this example it is important that there is no information that locates the empirical probability to a point within the interval $[0.7, 0.8]$. If you were to have the further information that $[0.7, 0.8]$ is a confidence interval for the empirical probability, then that information differentiates 0.75 as the most plausible candidate for the empirical probability—in which case arguably Calibration requires that degree of belief be set to 0.75 and there is nothing for Equivocation to do. More complex knowledge about the location of the empirical probability may involve higher-order probabilities, e.g., an empirical probability distribution over the location of the empirical probability within the interval. Calibration may then require the formation of higher-order degrees of belief—this added complexity is well studied by objective Bayesian statisticians, has no bearing on what follows and will not be considered further here.

three principles outlined above strongly constrain degrees of belief, leaving little room for an agent to subjectively choose how strongly to believe a proposition.² Objective Bayesianism is currently popular in AI (especially machine learning and natural language processing), statistics, engineering and physics.

Contemporary objective Bayesianism implements the three principles as follows. First (Probability), given the agent’s language (or some formal representation of the agent’s world), the agent’s degrees of belief should be representable by a probability function over propositions expressible in that language. We will denote the set of probability functions on this domain by \mathbb{P} . Second (Calibration), the agent’s *evidence* or *epistemic background*—the set \mathcal{E} of propositions that the agent takes for granted, including her background knowledge, observations and theoretical assumptions—imposes constraints on her degrees of belief, in the sense that it narrows down a set $\mathbb{E} \subseteq \mathbb{P}$ of probability functions that are compatible with her evidence, according to the following recipe. Propositions in \mathcal{E} impose constraints χ on the agent’s degrees of belief: for example, if proposition a is in \mathcal{E} then this imposes the constraint $P(a) = 1$; if \mathcal{E} contains a narrowest-reference-class frequency of 0.7 for a then this imposes the constraint $P(a) = 0.7$; if theoretical facts in \mathcal{E} deem the physical chance of a to be 0.4 then $P(a) = 0.4$ is a constraint in χ .³ Let $\mathbb{P}_\chi \stackrel{\text{df}}{=} \{P \in \mathbb{P} : P \text{ satisfies } \chi\}$. As long as χ is consistent, the agent’s probability function $P_\mathcal{E}$ should lie within the set $\mathbb{E} = [\mathbb{P}_\chi]$, the convex hull of the set of probability functions that satisfy χ (Williamson, 2005, §5.3). Otherwise some consistency maintenance procedure needs to be invoked.⁴ Third (Equivocation), $P_\mathcal{E}$ should be as close as possible to an *equivocator* $P_=$, a function that is maximally non-committal with respect to the various outcomes. In sum, the agent’s degrees of belief should be representable by some $P_\mathcal{E} \in \{P \in \mathbb{E} : P \text{ is closest to } P_=\}$, where closeness to the equivocator is explicated in an appropriate way.

This paper will be concerned with the simplest framework for objective Bayesian epistemology, namely that in which the key features of the agent’s world are represented by a finite domain Ω of mutually exclusive elementary out-

²It is assumed by objectivists and subjectivists alike that some propositions—e.g., the proposition that it will rain here tomorrow—are sufficiently speculative that one can deliberate as to how strongly one ought to believe them, and then apportion the strengths of one’s beliefs according to the results of one’s deliberations. Of course there are other propositions for which there is no option of engaging in serious deliberation—e.g., you can hardly be less than certain of the proposition that you are reading. See Rowbottom and Bueno (2008, §3) for a discussion of doxastic voluntarism.

³ \mathcal{E} contains a narrowest-reference-class frequency of 0.7 for a if it implies that (i) the frequency of property S amongst individuals that satisfy property R is 0.7, (ii) a is equivalent to an application of S to an individual, and (iii) that individual also satisfies R ; and \mathcal{E} does not imply (i)–(iii) for any R' that holds of a strict subset of the set of individuals that satisfy R .

⁴If χ is inconsistent, i.e., if there is no probability function P that satisfies χ , then we can consider the maximal consistent subsets χ^* of χ , written $\chi^* \in \chi$, and take $\mathbb{E} = [\bigcup_{\chi^* \in \chi} \mathbb{P}_{\chi^*}]$, the convex hull of the union of the sets of probability functions that satisfy maximally consistent subsets of χ . Taking maximal consistent subsets of χ is the simplest way of handling inconsistencies in χ , but more sophisticated procedures can be used where there is further information about \mathcal{E} itself. For example, some propositions in \mathcal{E} may be more entrenched than others: these can be taken to override those of the others that give rise to inconsistencies. Here a might be considered more entrenched than b in \mathcal{E} if, were the agent to question the truth of one or other of a and b , she would be more inclined to question b than a —i.e., would be more inclined to keep granting a than b . In this paper I will not discuss exactly how the agent’s evidence might be determined or revised—this is clearly a concern of rational epistemology but has no bearing on the arguments put forward here.

comes. This finite case proceeds as follows. A proposition or event is represented by a subset of Ω . A function from subsets of Ω to the unit interval is a probability function if $P(\Omega) = 1$ and $P(a) = \sum_{\omega \in a} P(\omega)$ for each subset a of Ω . The equivocator $P_{=}$ sets $P_{=}(\omega) = 1/|\Omega|$ for each $\omega \in \Omega$. Distance is measured by cross entropy $d(p, q) = \sum_{\omega \in \Omega} P(\omega) \log P(\omega)/q(\omega)$. As long as \mathbb{E} is closed, $P_{\mathcal{E}}$ turns out to be uniquely determined and characterisable as the function in \mathbb{E} that has maximum entropy, where the entropy $H(P) = -\sum_{\omega \in \Omega} P(\omega) \log P(\omega)$. Hence we have the *Maximum Entropy Principle* (*maxent* for short): $P_{\mathcal{E}} \in \{P \in \mathbb{E} : P \text{ maximises } H\}$.

Objective Bayesianism has also been applied to countable logical languages and is often applied to uncountable domains in statistics, but in these cases the equivocators and the measures of distance are more complicated, and, while highly constrained, $P_{\mathcal{E}}$ may not be uniquely determined (Williamson, 2008a,b; Williamson, 2009, §19). The lessons of this paper carry over to these more complex situations, but for simplicity of exposition we stick with finite domains here. For simplicity we shall also focus on the case in which χ is a consistent set of affine constraints.⁵ This ensures that \mathbb{P}_{χ} is non-empty, closed and convex; hence $\mathbb{E} = \mathbb{P}_{\chi}$.

§3

OBJECTIVE AND SUBJECTIVE BAYESIAN UPDATING

Under objective Bayesianism, updating in the light of new evidence is rather straightforward. If one initially takes for granted propositions in \mathcal{E} , but this then changes to \mathcal{E}' , then one's degrees of belief should change correspondingly from $P_{\mathcal{E}}$ to $P_{\mathcal{E}'}$. On a finite domain, when granting \mathcal{E} the agent should select the $P_{\mathcal{E}} \in \mathbb{E}$ that maximises entropy, and then when granting \mathcal{E}' she should select the $P_{\mathcal{E}'} \in \mathbb{E}'$ that maximises entropy. Degrees of belief track the agent's evidence; since the agent's background \mathcal{E} strongly constrains (on a finite domain, fully constrains) her degrees of belief, changes to this background also strongly constrain changes to her degrees of belief.

As mentioned in §2, advocates of subjective probability—often called *subjective Bayesians*—advocate the Probability principle, usually advocate the Calibration principle, but do not advocate Equivocation. This leaves much weaker constraints on degrees of belief: granting \mathcal{E} , choose some $P \in \mathbb{E}$. Subject to just these constraints, one's degrees of belief are permitted to vary wildly under small (or no) changes to one's evidence \mathcal{E} . Accordingly, subjective Bayesians impose a further principle in order to more strongly constrain updated degrees of belief. This is the principle of Bayesian Conditionalisation: if on granting \mathcal{E} your degrees of belief are represented by P , and if your evidence changes to $\mathcal{E}' = \mathcal{E} \cup \{e\}$ where e is an event in your domain such that $P(e) > 0$, then your new degrees of belief should be represented by $P' = P(\cdot|e)$, i.e., by your old degrees of belief conditional on the new evidence. This principle does the job very

⁵An *affine* constraint is satisfied by every probability function $P_{\lambda} = (1-\lambda)P_0 + \lambda P_1$ that lies on a line through any two probability functions P_0, P_1 that satisfy the constraint. On a finite space, affine constraints take the form of sets of expectation constraints, $\sum_{\omega} P(\omega)f(\omega) = c$ for real valued function f and real number c . Maximum entropy methods were originally applied solely to affine constraints for computational reasons, although this restriction may be relaxed.

well—provided the conditions on e are satisfied, P' is uniquely determined.⁶

Unsurprisingly perhaps, these two different approaches to updating can lead to different recommendations as to which degrees of belief to adopt in the light of new evidence. A standard example proceeds as follows. Suppose there are two agents, Olive (who abides by objective Bayesian norms) and Sylvester (who follows the subjectivist path). They both know \mathcal{E} , namely that an experiment is being performed which has three possible outcomes. The outcome space is $\Omega = \{1, 2, 3\}$ representing the three possible outcomes. By Equivocation, Olive sets $P_{\mathcal{E}}^O(1) = P_{\mathcal{E}}^O(2) = P_{\mathcal{E}}^O(3) = 1/3$. Sylvester can choose any belief distribution he likes, but we shall suppose that his beliefs match Olive's, $P^S = P_{\mathcal{E}}^O$. Now both agents are presented with new information e , which says that $1 \times P^*(1) + 2 \times P^*(2) + 3 \times P^*(3) = x$, i.e., with respect to physical probability P^* the expected outcome E takes some particular value x in the interval $[1, 3]$. In this case $\mathbb{P}_x = \{P : 1 \times P(1) + 2 \times P(2) + 3 \times P(3) = x\}$ is the set of probability distributions that yield expected outcome $E = x$; this is non-empty, closed and convex, so $\mathbb{E}' = [\mathbb{P}_x] = \mathbb{P}_x$. Olive then chooses the distribution in \mathbb{E}' that has maximum entropy as her new belief function $P_{\mathcal{E}'}^O$. Now Sylvester can't update at all in this situation, because e is not in the outcome space. But we can ask what would happen were the outcome space enlarged to incorporate e , and were Sylvester then to update his degrees of belief by conditionalising on e . It turns out that if Sylvester's updated degrees of belief are to match Olive's, he would have to have given prior probability 1 to the outcome $E = 2$ (Shimony, 1973). This seems ridiculous, since it amounts to prior certainty that the initial probability will not be revised. Of course Sylvester is a subjectivist and is rationally permitted to believe $E = 2$ to any degree he wishes, and if he were less than certain that $E = 2$ then his degrees of belief would differ from Olive's after updating.

The Judy Benjamin problem can be interpreted as providing another example of a disagreement between the objective and subjective Bayesian versions of updating (van Fraassen, 1981; van Fraassen et al., 1986; van Fraassen, 1987, 1989). Private Judy Benjamin is dropped into a swampy area which is divided into four quadrants each occupied by a different company: Red HQ (r_1), Red 2nd (r_2), Blue HQ (b_1), Blue 2nd (b_2). Benjamin does not know which zone she is in, and initially equivocates, setting $P_{\mathcal{E}}(r_1) = P_{\mathcal{E}}(r_2) = P_{\mathcal{E}}(b_1) = P_{\mathcal{E}}(b_2) = 0.25$. She then learns evidence that imposes the constraint $P(r_1|r_1 \vee r_2) = 3/4$. Maximising entropy then gives $P_{\mathcal{E}'}(r_1) = 0.351, P_{\mathcal{E}'}(r_2) = 0.117, P_{\mathcal{E}'}(b_1) = P_{\mathcal{E}'}(b_2) = 0.266$. Grove and Halpern (1997) argue that by expanding the probability space Bayesian conditionalisation can be applied to this problem; this update differs from the objective Bayesian update, though for slightly different reasons to the previous example (Uffink, 1996, §6).⁷

⁶Conditionalisation is called a *conservative* updating rule since it tries to conserve prior degrees of belief. In contrast, objective Bayesian updating is *foundational* in the sense that objective Bayesian degrees of belief track their foundations, namely the agent's evidence. Historically the split between objective and subjective Bayesian updating has not always been clear cut—some objective Bayesians have advocated conditionalisation (Jaynes, 2003) (as have proponents of logical probability such as Carnap), while some subjectivists have only endorsed conditionalisation in a limited way (Howson, 1997). The points I make in the following sections remain: objective Bayesians *need not* advocate conditionalisation, and indeed *should not* where conditionalisation disagrees with the objective Bayesian approach presented above.

⁷Note that the Judy Benjamin problem was originally presented as a problem for mini-

FOUR KINDS OF INCOMPATIBILITY

In the literature thus far, any incompatibility between maximum entropy updating (*maxent* for short) and Bayesian conditionalisation has normally been taken to reflect negatively on the former rule rather than the latter (see, e.g., Friedman and Shimony, 1971; Shimony, 1973; Seidenfeld, 1979; Dias and Shimony, 1981; Shimony, 1985; Skyrms, 1985; Seidenfeld, 1986). Thus the standard line of argument is: maxent is incompatible with Bayesian conditionalisation, but conditionalisation is well-established and intuitive, and so maxent should be rejected.

Here I shall argue the opposite: that Bayesian conditionalisation should be rejected because it disagrees with maximum entropy updating. The argument has two steps. First I shall reformulate an old result which shows the conditions under which maxent and conditionalisation agree. Next I shall go through each of these conditions in turn and we shall see that where the condition fails, maxent is a more appropriate form of updating than conditionalisation.

As before, suppose that the agent initially takes \mathcal{E} for granted, that \mathcal{E} imposes constraints χ , and that $P_{\mathcal{E}}$ is the probability function in \mathbb{E} that has maximum entropy. Then the agent learns e to give $\mathcal{E}' = \mathcal{E} \cup \{e\}$ as her new evidence. This imposes constraints χ' and $P_{\mathcal{E}'}$ is the new maximum entropy probability function.

DEFINITION 4.1 *e is simple with respect to \mathcal{E} iff $\chi' = \chi \cup \{P(e) = 1\}$, i.e., the only constraint that e imposes in the context of \mathcal{E} is $P(e) = 1$.*

The following is a reformulation of a result, versions of which are well-known to both proponents and opponents of maximum entropy methods (Williams, 1980; Seidenfeld, 1986, Result 1):

THEOREM 4.2 *If*

1. *e is a domain event,*
2. *e is simple with respect to \mathcal{E} ,*
3. *χ' is consistent, and*
4. *$P_{\mathcal{E}}(\cdot|e)$ satisfies χ ,*

maximum cross-entropy updating, yet another form of updating that incorporates aspects of both maximum entropy updating and Bayesian conditionalisation.

It has been claimed that in this example the maximum entropy update is counter-intuitive: the new evidence does not provide any reason to stop equivocating between red and blue, yet now $P_{\mathcal{E}'}(r_1 \vee r_2) = 0.468$ and $P_{\mathcal{E}'}(b_1 \vee b_2) = 0.532$. But this objection begs the question, since it assumes that one should continue to equivocate between red and blue if the evidence permits. Maximum entropy updating is not intended to be conservative: the idea is not to preserve features of prior degrees of belief when updating, but to calculate appropriate degrees of belief afresh. This requires being as equivocal as possible with respect to the space of elementary outcomes $\Omega = \{r_1, r_2, b_1, b_2\}$, not with respect to combinations of elementary outcomes such as $r_1 \vee r_2$ and $b_1 \vee b_2$. The new evidence is equivalent to the constraint $P(r_1) = 3P(r_2)$; the function satisfying this constraint that is maximally equivocal over Ω is simply not that which is maximally equivocal over red and blue. (Of course if there were a good reason why one should continue to equivocate between red and blue then this would yield a further constraint and the resulting probabilities would satisfy this constraint.) See Keynes (1921, Chapter 4) for a discussion of why one should equivocate over the more refined partition.

then $P_{\mathcal{E}'}(a) = P_{\mathcal{E}}(a|e)$.

This result clearly states conditions under which the maximum entropy update will agree with the Bayesian conditionalisation update of $P_{\mathcal{E}}$. The following special case conveys the substantial agreement between the two forms of updating (here \mathcal{E} is empty):

COROLLARY 4.3 *If the agent learns just consistent domain events e_1, \dots, e_n and each event e_i is simple with respect to e_1, \dots, e_{i-1} then maximum entropy updating agrees with Bayesian conditionalisation.*

See Seidenfeld (1986, Result 1 and corollary) for proofs of Theorem 4.2 and Corollary 4.3.

Incompatibilities between maxent and conditionalisation can only arise if the conditions of Theorem 4.2 fail. Let us examine these conditions in turn.

CASE 1. First consider the case in which e is not an event in the original domain. In order to update, the domain must of course be extended to include e . Then the objective Bayesian will update as normal by maximising entropy on the new domain. However the situation is not so straightforward for the subjectivist: it is not possible to apply Bayesian conditionalisation because the prior $P_{\mathcal{E}}(a|e)$ is not defined. The best one can do is to assign a ‘prior’ on the extended domain post hoc, and then update by conditionalisation. This is the strategy of the Olive-Sylvester and the Judy Benjamin examples discussed above. It is not a good strategy in general, because the new ‘prior’ is assigned at the same stage as posterior, so can be engineered to fit any desired posterior. In effect, conditionalisation offers no constraint on the updating of probabilities if this post-hoc-prior strategy is pursued. To see this, suppose for example that the original domain has two possible outcomes $\{a, \neg a\}$ where a signifies *the patient has cancer*. An agent assigns prior probability $\frac{1}{2}$ to each proposition, then changes her degree of belief to 0.99 that the patient has cancer, for no good reason at all. This is quite rational under the subjective Bayesian account if the above strategy is permitted. All the agent needs to do is find some proposition e that is not in the domain but which is true (e.g., *the moon is not made of blue cheese*, or, *the first digit of π is 3*), extend her domain to include e and $\neg e$, and define a new post hoc ‘prior’ degree of belief $P(a|e) = 0.99$; then her posterior degree of belief that the patient has cancer is $P'(a) = 0.99$. The point of Bayesian conditionalisation is to ensure that changes of degrees of belief are not so wild; clearly conditionalisation achieves nothing in this case. Thus this case favours maximum entropy updating over Bayesian conditionalisation.

Note that if this post-hoc-prior strategy is implemented and then an incompatibility arises between maxent and conditionalisation, this must be due to an infringement of one of the other conditions of Theorem 4.2, so we may move on to these other cases.

CASE 2. Second, consider the case in which e is not simple with respect to \mathcal{E} . An example will help show why maxent can disagree with conditionalisation here. Suppose e says ‘ $P^*(a) = 0.8$ ’, where P^* is physical probability. Then by the Calibration principle, learning e does not merely impose the constraint $P(e) = 1$, but also the constraint $P(a) = 0.8$; hence as long as $P(a) = 0.8$ is not

already a constraint imposed by \mathcal{E} , e is not simple with respect to \mathcal{E} . Should one prefer maxent or conditionalisation in this situation? According to maxent the agent will set $P_{\mathcal{E}'}(a) = 0.8$, which seems just right, given that the evidence imposes that constraint. On the other hand if $P_{\mathcal{E}}(a|e) \neq 0.8$ then it is quite wrong to conditionalise, because the new degree of belief will conflict with the new evidence. One might try to save conditionalisation by arguing that the agent *should* have set $P_{\mathcal{E}}(a|e) = 0.8$ in the first place, and then no conflict would have arisen. But this argument fails because subjectivism permits the agent to set $P_{\mathcal{E}}(a|e)$ to any value she likes—since e is not granted at the time of setting the prior, it neither constrains the value of $P_{\mathcal{E}}(a)$ nor of $P_{\mathcal{E}}(a|e)$. Of course if e deductively entailed a (respectively $\neg a$) then $P_{\mathcal{E}}(a|e)$ would be constrained to 1 (respectively 0); but this is not the case here. (If the subjectivist would like to add a further norm that forces $P_{\mathcal{E}}(a|e) = 0.8$, that is all well and good: no disagreement between the objective and subjective forms of updating will arise in this case, and we must focus on the other cases of disagreement. Of course, subjectivism will become that much less subjective as a consequence.)

Note that Bacchus et al. (1990, §4) argue that conditionalisation goes wrong for just this reason; here is their example. Suppose a is ‘Peterson is a Swede’, b is ‘Peterson is Norwegian’, c is ‘Peterson is Scandinavian’ and e is ‘80% of all Scandinavians are Swedes’. Initially the agent sets $P_{\mathcal{E}}(a) = 0.2$, $P_{\mathcal{E}}(b) = 0.8$, $P_{\mathcal{E}}(c) = 1$, $P_{\mathcal{E}}(e) = 0.2$, $P_{\mathcal{E}}(ae) = P_{\mathcal{E}}(be) = 0.1$. All these degrees of belief satisfy the norms of subjectivism. Updating by maxent on learning e , the agent believes that Peterson is a Swede to degree 0.8, which seems quite right. On the other hand, updating by conditionalising on e leads to a degree of belief of 0.5 that Peterson is a Swede, which is quite wrong. Thus we see that maxent is to be preferred to conditionalisation in this kind of example because the conditionalisation update does not satisfy the new constraints χ' , while the maxent update does.

There is another way in which e not being simple with respect to \mathcal{E} can cause problems. In this scenario, both the conditionalisation update and the maxent update satisfy χ' , yet maxent is still to be preferred over conditionalisation. Suppose a is ‘Peterson is a Swede’, c is ‘Peterson is Scandinavian’ and e is ‘at least 50% of Scandinavians are Swedes and all politicians are liars’. \mathcal{E} is the knowledge that c is true together with the testimony from a politician that $\neg ae$ is false; so $P_{\mathcal{E}}$ satisfies the constraints $\chi = \{P(c) = 1, P(\neg ae) = 0\}$. Maximising entropy, she sets $P_{\mathcal{E}}(a) = 2/3$, $P_{\mathcal{E}}(a|e) = 1$. Then the agent learns that e is true. Now e is not simple with respect to \mathcal{E} —in particular it undermines the politician’s testimony. Plausibly, $\chi' = \{P(c) = 1, P(e) = 1, P(a|e) \geq 0.5\}$. Maximising entropy now gives $P_{\mathcal{E}'}(a) = 0.5 \neq 1 = P_{\mathcal{E}}(a|e)$. In this example $P_{\mathcal{E}}(a|e)$ does satisfy the constraints in χ' , yet intuitively it is wrong to update by conditionalisation because, given current evidence, to believe a to degree 1 is too extreme. Neither subjectivists nor objective Bayesians will want to *force* an agent to fully believe a when there is no evidence to warrant such a strong belief. If an updating rule is to be used to constrain the new degree of belief in a , then the maxent rule is much more plausible.

We see then that if e is not simple with respect to \mathcal{E} , maxent is to be preferred over conditionalisation where they differ, whether or not the conditionalisation update satisfies the new constraints χ' : if the conditionalisation update does not satisfy constraints that should be satisfied, then it is clearly inappropriate; if it does satisfy the constraints then either it agrees with the maxent update

or it is less equivocal than the maxent update, but in the latter case it hardly seems right that an update rule should *force* one to have more extreme degrees of belief than the evidence warrants. Note that in the two examples considered earlier—Olive-Sylvester and Judy Benjamin—in which the domain is extended to represent the new evidence as a domain event, the new evidence is *not* simple with respect to the background knowledge. I have suggested that the right thing to do in such a circumstance is to satisfy *all* the constraints imposed by the evidence (as maxent does), not just the constraint $P(e) = 1$ (which is essentially all that Bayesian conditionalisation does).

CASE 3. The third kind of infringement of the conditions of Theorem 4.2 is the case in which χ' is inconsistent. This implies that either χ is inconsistent or e is inconsistent with χ . In line with the restrictions imposed at the end of §2, we will not consider the former situation in any detail: here the consistency maintenance procedure will be carried out, $\mathbb{E} \neq \mathbb{P}_\chi$, and, as long as the other conditions hold, maxent will agree with conditionalisation. The latter situation is the real problem area. If e is inconsistent with χ then $P_{\mathcal{E}}(e) = 0$, and $P_{\mathcal{E}}(a|e)$ is either undefined or unconstrained, according to how conditional probabilities are interpreted. (If conditional probability is *defined* by $P(a|e) = P(ae)/P(e)$ then it is normally taken to be undefined when $P(e) = 0$; alternatively if conditional probability is taken as primitive and subject to the constraint that $P(a|e)P(e) = P(ae)$ then it is unconstrained when $P(e) = 0$.) If $P_{\mathcal{E}}(a|e)$ is undefined then Bayesian conditionalisation can not be applied at all. If $P_{\mathcal{E}}(a|e)$ is unconstrained then the agent can assign it whatever value she likes, and Bayesian conditionalisation offers no substantive constraint, in the sense that her new degree of belief in a need in no way be connected to her unconditional prior degree of belief in a, e or indeed any other proposition. This leads to crazy changes in belief being deemed rational: for example, evidence that is known to count *against* a can be invoked to *raise* the probability of a . In sum, when e is inconsistent with χ , conditionalisation is inapplicable or vacuous. In contrast, maximum entropy updating is perfectly applicable in this situation: $P_{\mathcal{E}'}(a)$ is the probability function, from all those in \mathbb{E}' , that is most equivocal. Note that here \mathcal{E}' is inconsistent so the consistency maintenance procedure will determine \mathbb{E}' (which will of course differ from $\mathbb{P}_{\chi'} = \emptyset$). Hence in this case the right way to update is via maxent and not via conditionalisation.

CASE 4. The fourth and final case is that in which $P_{\mathcal{E}}(\cdot|e)$ does not satisfy χ . Here is an example from Seidenfeld (1986, §2). A die is thrown, the outcome space being $\Omega = \{1, \dots, 6\}$. The agent knows that the expected score for this die is 3.5, so $\chi = \{\sum_{\omega} P(\omega) = 3.5\}$. Subject to just this constraint, $P_{\mathcal{E}}$ is uniform, giving $P_{\mathcal{E}}(\omega) = 1/6$ for each outcome. (The probability function from all those that satisfy the constraint which is closest to the equivocator is the equivocator itself, i.e., the uniform distribution.) The agent then learns e , that an even outcome occurred. If $\chi' = \{\sum_{\omega} P(\omega) = 3.5, P(\{2, 4, 6\}) = 1\}$ then we get $P_{\mathcal{E}'}(2) = 0.47$, $P_{\mathcal{E}'}(4) = 0.32$, $P_{\mathcal{E}'}(6) = 0.22$. This differs from the result of Bayesian conditionalisation since $P_{\mathcal{E}}(2|e) = P_{\mathcal{E}}(4|e) = P_{\mathcal{E}}(6|e) = 1/3$. Note that the solution obtained by conditionalisation does not satisfy the constraints χ , since the expectation is now 4, not 3.5. Which is right, maxent or conditionalisation? If the analysis is right, and $\sum_{\omega} P(\omega) = 3.5$ really is a constraint on

both the prior and the updated belief functions, then conditionalisation *must* be wrong—it simply doesn’t satisfy that constraint. On the other hand, suppose the analysis is wrong. If the constraint should not be in χ then there need not be any disagreement between maxent and conditionalisation after all—this will depend on whether $P_{\mathcal{E}}(\cdot|e)$ satisfies the correct χ . Alternatively, if the constraint should be in χ but should not be in χ' then e is not simple with respect to \mathcal{E} . This is the second case discussed above. There we saw that whether or not the conditionalisation update satisfies the correct χ' , it can hardly be considered the correct update where it differs from the maxent solution.

This concludes the argument. We have seen that in each of the four cases where maxent disagrees with conditionalisation, the maxent update is to be preferred over the conditionalisation update.

§5

CRITICISMS OF CONDITIONALISATION

Before turning to questions of normativity in §9 it is worth rehearsing some of the standard arguments for and against conditionalisation, in the light of the previous discussion. We will take a look at two arguments *for* conditionalisation in §§6, 7. In this section we will concentrate on arguments *against* conditionalisation. The argument of §4 is only one of many that have been pitched against Bayesian conditionalisation. It has been argued that conditionalisation is problematic for the following reasons.

For one thing, conditionalisation requires being eternally true to prior beliefs: an agent’s degrees of belief throughout time are fixed by her initial probability distribution—in particular by conditional probabilities of the form $P(a|e)$. Now while the agent may deem these conditional degrees of belief appropriate at the outset, they may well turn out to be inappropriate in retrospect. Indeed, this is typically the situation when the new evidence is not simple with respect to prior knowledge (case 2 of §4). Thus, for example, exchangeable prior degrees of belief may be appropriate in the absence of knowledge about an experimental set-up, but when the string e of evidence indicates that the data is being produced by a Markovian process, rather than a sequence of independent trials, exchangeable degrees of belief are disaster (Gillies, 2000, pp. 77–83; Williamson, 2007a); however, the agent is not permitted by conditionalisation to change her degrees of belief to be non-exchangeable.⁸ Even advocates of conditionalisation admit that this is unsatisfactory and that conditionalisation may need to be overridden when a prior is re-evaluated (see, e.g., Lange, 1999, §2; Earman, 1992). But these advocates of conditionalisation offer little in the way of a general rule for updating. Of course a vague prescription like ‘conditionalise unless an update that does not conform to conditionalisation is warranted, in which case use the latter update’ is hardly a rule at all. In the light of the discussion of §4, it appears that the best way to formulate a precise rule of this form is to say ‘conditionalise unless one of the conditions of Theorem 4.2 is infringed, in which case update by maximising entropy’. But then it is equivalent and simpler just to say ‘update by maximising entropy’.

⁸Given a sequence of outcomes, positive or negative, an agent’s prior degrees of belief are *exchangeable* iff for each number n , the agent assigns the same degree of belief to each possible sequence of outcomes that yields n positive outcomes in total.

Note that maximum entropy updating—and objective Bayesian updating in general—does not suffer from this pathology: it does not force the agent to remain eternally true to her prior beliefs. At each step the agent chooses a belief function, from all those that fit the evidence at the time, that is maximally equivocal. There is simply no room for the agent’s updated degrees of belief to be out of step with the evidence. This fact is often overlooked: it is often charged that objective Bayesianism suffers from being stuck with an unreasonable prior, namely one that renders learning from experience impossible (see, e.g., [Dias and Shimony, 1981](#), §4). Consider events b_1, \dots, b_{101} , where b_i signifies that the i ’th observed raven is black. Subject to no constraints, maximising entropy yields $P(b_{101}) = P(b_{101}|b_1 \cdots b_{100}) = 1/2$: observing a hundred black ravens does not strengthen an agent’s belief that raven 101 is black—i.e., the agent fails to learn from experience. But this objection is misplaced since $P(b_{101}|b_1 \cdots b_{100})$ does *not* represent the degree to which the agent should believe that raven 101 is black on the basis of evidence $b_1 \cdots b_{100}$. This is because the evidence $b_1 \cdots b_{100}$ is not simple. Not only does learning this evidence impose the constraint $P(b_1 \cdots b_{100}) = 1$, but also the constraint that $P(b_{101}) \simeq 1$. (Exactly how this last constraint is to be made precise is a question of statistical inference which may be solved by appealing to confidence intervals or higher-order probabilities—the details need not worry us here.) When evidence $b_1 \cdots b_{100}$ constitutes the new evidence \mathcal{E}' , we see that $P_\emptyset(b_{101}) = 1/2$ (and $P_\emptyset(b_{101}|b_1 \cdots b_{100}) = 1/2$) but $P_{\mathcal{E}'}(b_{101}) \simeq 1$. Hence this is a case in which maximum entropy updating disagrees with conditionalisation, and the agent does, after all, learn from experience.⁹

The above objection to conditionalisation is often put another way: conditionalisation requires specifying all future degrees of belief in advance, i.e., specifying from the outset $P(a|e)$ for all possible sequences e of evidence, but this is impossible from a practical point of view. Note that the Maximum Entropy Principle does not require that so many probabilities be specified. Since degrees of belief are a function of the agent’s evidence \mathcal{E} , she does not need to calculate the full function $P_{\mathcal{E}}$ at any particular time, but only that part of the function that is required for practical purposes. Thus at any stage $P_{\mathcal{E}}$ may be partially specified—because it is determined afresh on changes to \mathcal{E} , this does not prevent updating. The subjectivist may appeal by saying that she can also partially specify her belief function, and just specify the ‘prior’ $P(a|e)$ after the evidence e has been discovered. But our discussion in §4 of the post-hoc-prior strategy shows that this does not save conditionalisation. Moreover, if e is not simple with respect to prior knowledge, then for a post hoc ‘prior’ $P(a|e)$ to be reasonable, it can typically not really be construed as part of the prior at all—the reasonable value of $P(a|e)$ will in general not be the value that would have been chosen prior to the gathering of evidence—it would be the maxent update instead. Contrapositively, if the post hoc $P(a|e)$ were to match the genuine prior value, then either it would agree with maxent or it would not lead to a reasonable update.

⁹There is another, albeit less straightforward, way of replying to objection that objective Bayesianism fails to permit learning from experience. This involves arguing that it is a mistake to say that there are no prior constraints in operation. It is known that the b_i are all related in that they are observations of the colour of ravens; when the constraints imposed by this knowledge are made explicit, learning from experience can be recaptured. See [Williamson \(2007a\)](#) and [Williamson \(2008b\)](#) for this line of argument.

Closely related to this objection is the charge that Bayesian conditionalisation requires every possible future piece of evidence to be representable as a proposition in the agent's prior language. It is clearly unreasonable to demand that an agent have this amount of foresight. As we saw, the natural response for the subjectivist is to claim that one can extend the language or outcome space to incorporate the new evidence as and when it becomes known, and then to give a 'prior' value to $P(a|e)$ post hoc. But, to reiterate, this is a poor strategy because it offers no constraint on the post hoc degrees of belief. As we also saw, maxent does not succumb to this problem. If the new evidence is not in the old language then the language can simply be extended to incorporate e and maxent performed as normal (Williamson, 2005, Chapter 12). Since on the objective Bayesian account posterior degrees of belief are constrained by maxent rather than by prior degrees of belief, there is no need for evidence to be representable in the prior language.

Bayesian conditionalisation has also been criticised on account of the fact that conditionalisation must render all evidence eternally certain, even though evidence is clearly fallible. When e is learned, $P(e)$ remains 1 for all time, even if later evidence casts doubt on e . There doesn't seem to be any way for the subjectivist to uphold Bayesian conditionalisation yet get round this objection. It is sometimes argued that e should not have been considered conditionalisable evidence if it were not beyond all possible doubt. (If say $P(e) = 0.9$ then Jeffrey conditionalisation, another special case of maxent updating, could be applied instead of Bayesian conditionalisation.) But of course this response is tantamount to abandoning Bayesian conditionalisation, since practically no evidence is incontrovertible. On the other hand maxent does not render evidence eternally certain: if the veracity of e is in doubt, one simply removes e from the current evidence, and then maximises entropy as usual.¹⁰ Indeed, objective Bayesian epistemology does not require that e be certain in the first place. If e is obtained by a process with error rate ε then learning e imposes the constraint $P(e) = 1 - \varepsilon$, rather than $P(e) = 1$. Objective Bayesianism can handle such constraints with no difficulty, while Bayesian conditionalisation requires that $P(e) = 1$ when e is learned.

Howson has argued compellingly that Bayesian conditionalisation is inconsistent in some situations (Howson, 1997, §3; Howson, 2000, p. 136). Consider the following example, cast in subjectivist terms. Today you are certain of a , so $P(a) = 1$, but suppose you think that tomorrow you may be significantly less than certain about a , i.e., $P(P'(a) \leq 0.9) > 0$. Since $P(a) = 1$, $P(a|P'(a) \leq 0.9) = 1$. Then tomorrow comes and indeed $P'(a) \leq 0.9$, and you realise this by introspection. But by conditionalisation, $P'(a) = P(a|P'(a) \leq 0.9) = 1$, a contradiction. Note that here the evidence $P'(a) \leq 0.9$ is not simple, since it imposes the constraint $P'(a) \leq 0.9$ as well as the constraint $P'(P'(a) \leq 0.9) = 1$. So in the light of our preceding discussion it is no wonder that conditionalisation goes wrong. An objective Bayesian approach is clearly to be preferred here: if $P(a) \leq 0.9$ is genuinely a constraint on new degrees of belief then it will be satisfied with no contradiction; if it is not genuinely a constraint then it will not be imposed and again no contradiction will arise.

¹⁰It should be emphasised that the norms of objective Bayesianism govern neither the question of what is to be included in the agent's evidence nor the questions of whether and why the agent is rational to take for granted what is included in this background. Objective Bayesianism covers the relation between evidence and degrees of belief, not the evidence itself.

A DUTCH BOOK FOR CONDITIONALISATION?

Having set the case against conditionalisation, it is only fair to briefly consider a couple of the more widely cited arguments in favour of conditionalisation: the argument from conservativity (discussed in §7), and the diachronic Dutch book argument for conditionalisation (discussed in this section).

Paul Teller presents the Dutch book argument in Teller (1973), attributing the idea to David Lewis. The argument aims to show that if you fail to update by conditionalisation then you are susceptible to a Dutch book, i.e., if you bet according to your degrees of belief then a judicious bookmaker can make you lose money whatever happens. This appears to be a clear reason to favour conditionalisation over maxent. van Fraassen (1989, p. 174) points out that this Dutch book argument only shows that *if you specify an updating rule in advance* and that rule is not conditionalisation, then you are susceptible to Dutch book; but this is exactly the case we have here: maxent is a rule, specified in advance, that may disagree with conditionalisation. Thus this Dutch book argument appears to be a serious challenge to maxent.

We won't go through the details of the Dutch book argument here, since no doubt will be cast on its validity. Instead, I shall cast doubt on the importance of diachronic Dutch books in general, by means of a reductio. Specifically I shall show that in certain situations one can Dutch book *anyone who changes their degrees of belief at all*, regardless of whether or not they change them using conditionalisation. Thus avoidance of Dutch book is a lousy criterion for deciding on an update rule.¹¹

Suppose it is generally known that you will be presented with evidence that does not count against a , so that your degree of belief in a will not decrease. It is also known that the evidence is inconclusive, so that you will not be certain about a . It is not known in advance what this evidence is. (This is the kind of situation which might arise in a court of law, for example, when a juror hears that the prosecution will present its evidence, knowing that the prosecution is competent enough not to present evidence that counts against its case.) You set your prior $P(a) = q$. Now in the betting set-up that is used by Dutch book arguments, this is tantamount to a bet which will yield S if a occurs for a payment of qS , where S is a real-valued stake chosen by a bookmaker (note in particular that S may be positive or negative). Thus your gain is $G = (I_a - q)S$, where I_a is the indicator function for a ,

$$I_a = \begin{cases} 1 & : a \\ 0 & : \neg a \end{cases}$$

The stake-maker chooses $S = -1$. Then the new evidence arrives and you duly set $P'(a) = q'$, where $q \leq q' < 1$. This is interpreted as a second bet,

¹¹I should stress that this is only an attack on dynamic Dutch books, i.e., Dutch books involving degrees of belief at more than one time, not on synchronic Dutch books such as the Dutch book argument for the axioms of probability, which I claim (but will not argue here) remain compelling. There is a large literature concerning the cogency of Dutch book arguments. See for example Howson (1997, §7) for related worries about dynamic Dutch books, Foley (1993, §4.4) for reservations about Dutch book arguments in general, Rowbottom (2007) for a recent critique of the synchronic Dutch book argument in the subjectivist framework, and Skyrms (1987) for a defence of Dutch book arguments in general.

with stake S' . Your total gain over the prior and posterior bets on a is $G' = -(I_a - q) + (I_a - q')S'$. Now unless $q' = q$, the stake-maker can choose stake S' such that $q/q' < S' < (1 - q)/(1 - q')$ to ensure that $G' < 0$ whatever happens. Thus unless your degree of belief in a remains constant, you are susceptible to a Dutch book.

In conclusion, don't trust dynamic Dutch books: it can be more important to change degrees of belief to reflect new evidence than to avoid certain loss.¹²

§7

CONDITIONALISATION FROM CONSERVATIVITY?

Another justification of Bayesian conditionalisation proceeds along the following lines. By the axioms of probability,

$$P'(a) = P'(a|e)P'(e) + P'(a|\neg e)P'(\neg e) = P'(a|e) \text{ if } P'(e) = 1.$$

So conditionalisation is the only possible way to update if $P'(a|e) = P(a|e)$. Now plausibly any agent should be conservative in the sense that she should keep $P'(a|e) = P(a|e)$ if at all possible. But it is always possible to be conservative

¹²Note that this Dutch book argument is rather different to the argument of van Fraassen (1984) in favour of the *principle of reflection*, which holds that $P(a | P'(a) = r) = r$. That argument concerns how one should set one's prior degrees of belief, while this argument concerns posterior degrees of belief. Moreover, here we do not assume that statements about her future beliefs are in the domain of the agent's belief function. If the reductio of this section is accepted, van Fraassen's diachronic Dutch book argument for reflection should be treated with suspicion.

More than that though, the principle of reflection itself should be treated with caution. Consider the following extension of the above betting scenario. Suppose first that you are prepared to bet on your future beliefs, so that it makes sense to consider expressions like $P(P'(a) = r)$. Suppose too that you know that your degree of belief in a should either remain the same or should increase to some value r . Now by the principle of reflection,

$$q = P(a) = P(a|P'(a) = q)P(P'(a) = q) + P(a|P'(a) = r)P(P'(a) = r) = qx + r(1 - x)$$

where $x = P(P'(a) = q)$. But for this identity to hold either $x = 1$ or $r = q$. Either case implies that you should fully believe that your degree of belief in a will not change. But that seems a ridiculously strong conclusion. Is this a reductio of the principle of reflection?

The proponent of the principle of reflection might suggest that all this shows is that your current degree of belief should not be q ; perhaps it should be somewhere between q and r , the possible values of your future degree of belief in a . But in the example you know that your degree of belief should either remain the same or increase to a fixed r . The proponent's suggestion rules out the former possibility. Hence your degree of belief should increase to r and, by reflection, your present degree of belief should be r . Far from escaping the reductio, the proponent of reflection has ended up in a rather worse situation. By trying to avoid the conclusion that you should believe that your degree of belief in a should not change, we have reached the conclusion that your degree of belief in a should be greater than is independently plausible *and* should not change.

In fact the principle of reflection and Bayesian conditionalisation come unstuck for similar reasons. Reflection is intended to capture the following idea: if one were to know that one will come to believe a to degree r , then one should currently believe a to degree r . But $P(a | P'(a) = r) = r$ fails to express this idea. The conditioning proposition $P'(a) = r$ is not simple; hence this is a case where conditionalisation can break down and the degree to which one ought to believe a , were one to grant that $P'(a) = r$, should not be expected to equal $P(a | P'(a) = r)$. This analysis motivates reformulating the reflection principle to more adequately express the above idea; perhaps as $P_{\mathcal{E} \cup \{P'(a)=r\}}(a) = r$. From this new principle of reflection one can no longer derive the above reductio.

in this sense if all one learns is e . Hence, conditionalisation is the right way to update.

I think this argument fails for the following reason. The kind of conservativity it appeals to is far too strong. Plausibly, perhaps, any agent should be conservative in the sense that she should keep $P'(a|e) = P(a|e)$ *unless the evidence e indicates otherwise*. But the evidence may indicate otherwise—the evidence may not be simple with respect to the previous evidence (case 2 of §4), or $P'(a) = P(a|e)$ may not satisfy the original constraints χ (case 4 above). Of course in the case in which e is not in the event space then the above argument does not succeed either (case 1). Nor does it work if $P(e) = 0$ (case 3). The best one can say then is that updating should agree with conditionalisation if the conditions of Theorem 4.2 hold.

In sum, when the conservativity condition is rendered less dubious this justification fails to tell us anything we didn't already know on the basis of the discussion of §4.

§8

SUMMARY

I argued in §4 that if we take the trouble to look at the different scenarios in which maxent disagrees with conditionalisation, maxent is to be preferred over conditionalisation in each case. Hence the argument against maximum entropy updating on the grounds of the existence of such disagreements backfires on the subjectivist.

This analysis helps to shed some light on some of the standard complaints against conditionalisation (§5). While these remain cogent criticisms of conditionalisation, maxent does not succumb to any of these objections. On the other hand, two key justifications of conditionalisation turn out to be less than cogent. The Dutch-book justification of conditionalisation fails because diachronic Dutch books are subject to a reductio (§6). The argument from conservativity fails to offer a fully general justification of conditionalisation—it only applies where conditionalisation agrees with maxent (§7). Thus the case against conditionalisation is stronger than ever.

But without a viable form of updating, subjectivism is laissez-faire to the point of ridicule: not only is an agent rationally entitled to adopt practically any probability function as her prior, she is rational to change it in any way she wishes too. In particular, under a subjectivist account without updating if a consultant has no evidence which bears on whether a patient has cancer or not then she is not only rational to initially believe that the patient has cancer to degree 1 but also rational to change her degree of belief to 0 on a whim.

§9

VOLUNTARISM AND STANCES

One kind of subjectivist objects to objective Bayesianism on the grounds that objective Bayesian updating can disagree with Bayesian conditionalisation; this objection backfires on the subjectivist. But perhaps as awkward for the objec-

tive Bayesian is the subjectivist who *doesn't* object to objective Bayesian assignments of degrees of belief. Such a subjectivist may say something like this: the objective Bayesian satisfies all the requirements of subjective Bayesianism—give or take a few discrepancies when updating—so there is nothing to object to; if her degrees of belief are always maximally equivocal, that's fine since subjectivism is a broad church.

The question then arises as to the status of the norms of objective Bayesianism—are they a voluntary code of conduct, or is an agent somehow compelled to abide by them?

The subjectivist will take issue with Equivocation in particular. I have argued in [Williamson \(2007b\)](#) that an agent should equivocate because equivocal degrees of belief will, on average, expose her to the least risky course of action. For instance, if a consultant has no evidence which bears on whether a patient has cancer or not, and she adopts a very high or very low degree of belief that the patient has cancer, then she exposes herself to unnecessary risks—a high degree of belief may trigger unnecessary treatment that is harmful to the patient, while a low degree of belief may lead to the patient being dismissed and a cancer not being treated—whereas a middling degree of belief is likely to trigger the gathering of further evidence. Thus the subjectivist concern with Equivocation can, I think, be tackled.

But there is a deeper worry that affects objective and subjective Bayesianism alike. What compulsion is there to abide by *any* norm of rational belief? The complaint goes: you may say that my degrees of belief should be such-and-such, but I do not consider myself to be subject to your rules. This leads to what might be called *radical subjectivism* or *radical voluntarism*, the view that rules of rational belief ascription express individual inclinations or value judgements and cannot be imposed on others.¹³ Under this position, objective and subjective Bayesianism are just two stances one might take as to how to set one's degrees of belief; neither can be considered right or wrong. A radical subjectivist might grant that if you value avoiding sure loss then you should abide by the Probability norm, for otherwise you are prone to a Dutch book. But, the radical subjectivist might say, I don't take avoiding sure loss to be paramount even in the synchronic case, let alone in the diachronic case. Similarly, the radical subjectivist may disavow a close correspondence between degrees of belief and physical probability (Calibration) and may not mind exposing himself to unnecessary risks (Equivocation).

One might be inclined to bite the bullet here and accept radical subjectivism, safe in the knowledge that others are in the same camp as oneself. Thus one might express the norms of rational belief in terms of means-ends rationality: *if* your ends are the avoidance of sure loss, a fit with reality and avoiding exposure to unnecessary risks, then you had better be a good objective Bayesian. If not, not. But as a matter of fact most of us have these values or ends, so we should abide by these norms. Then objective Bayesianism speaks to those who take the required stance; others may be perfectly rational to flout its rules.

But I think we can do a bit better than this. When we talk of degrees of belief

¹³Note that van Fraassen's voluntarism is less radical inasmuch as it holds that while rational norms incorporate value judgements, they are open to some sort of debate. See [van Fraassen \(1984\)](#); [van Fraassen \(1989, pp. 174–175\)](#) and [van Fraassen \(2002\)](#). For a recent discussion of radical subjectivism see *Episteme: A Journal of Social Epistemology* 4(1), 2007, a special issue on epistemic relativism.

we talk of something practical: the strengths of our beliefs guide our behaviour. Thus Ramsey (1926, p. 183) points out that ‘all our lives we are in sense betting. Whenever we go to the station we are betting that a train will really run, and if we had not a sufficient degree of belief in this we should decline the bet and stay at home.’ We use our degrees of belief to make decisions and to draw various inferences (e.g., to make predictions) and these are practical uses. Now the ends or values connected with objective Bayesianism come part-and-parcel with these uses. When we make decisions we want them to be good decisions—we don’t want to take unnecessary risks, for instance. When we make predictions we want them to be correct, not disjoint from reality. So merely by talking of degrees of belief we buy into these values. The ends are in the meaning of belief, they are not optional.

The radical subjectivist might take issue with this by saying: when I say ‘degrees of belief’ I do not buy into these values; I do not care whether or not I make good decisions on the basis of my ‘degrees of belief’; I talk about them in isolation from their consequences. Fine, but then the radical subjectivist is talking about something else—what we might call ‘degrees of idle-belief’, not the degrees of belief that do unavoidably ground his actions. The Bayesian may then respond by admitting that her norms do not govern the radical subjectivist’s degrees of idle-belief, but insisting that they do apply to his degrees of belief.

In that sense, then, the norms of objective Bayesianism are compulsory. You can use the terminology ‘degrees of belief’ to refer to anything you wish, but what I call your ‘degrees of belief’ had better abide by Probability, Calibration and Equivocation.

Consider an analogy: constructing a *boat* rather than constructing *degrees of belief*. Boats are used to travel along the surface of water, and floating is clearly part-and-parcel of the proper functioning of a boat. A boat is still a boat if it has been holed and sinks, so floating isn’t an essential property of a boat, but if you take it upon yourself to build a boat, you should make sure it floats on water. This *Floating* norm is thus imposed on the boat-building process rather than the boat itself. The radical boat-builder might not buy into this norm, saying: I’m not interested in travelling on water when I build my boat. But then the radical boat-builder is not building a boat in the normal sense of the word. Rather, he is building something that might be called an ‘art-boat’ (if it is being built for its aesthetic properties) or an ‘idle-boat’ (if it is to have no uses at all). Determining degrees of belief is like building a boat, inasmuch as once one has decided on the thing one is constructing, one is bound to construct it in such a way that it has its proper function, otherwise one fails in one’s task.

To conclude, not only is objective Bayesian updating preferable to Bayesian conditionalisation, but in principle one can also defend norms of rational belief—such as those that characterise objective Bayesianism—from the clutches of the radical subjectivist. This is not to say that contemporary objective Bayesianism is the be all and end all of rationality; indeed it faces a number of important challenges (Williamson, 2009). It is, however, a promising first approximation.

ACKNOWLEDGEMENTS

This research was supported by a fellowship from the Leverhulme Trust. I am grateful to David Corfield, Bas van Fraassen, Darrell Patrick Rowbottom and an anonymous referee for very helpful comments.

REFERENCES

- Bacchus, F., Kyburg, Jr., H. E., and Thalos, M. (1990). Against conditionalization. *Synthese*, 84(3):475–506.
- Bayes, T. (1764). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53:370–418.
- Bernoulli, J. (1713). *Ars Conjectandi*. The Johns Hopkins University Press, Baltimore, 2006 edition. Trans. Edith Dudley Sylla.
- Dias, P. M. C. and Shimony, A. (1981). A critique of Jaynes' maximum entropy principle. *Advances in Applied Mathematics*, 2(2):172–211.
- Earman, J. (1992). *Bayes or bust?* MIT Press, Cambridge MA.
- Foley, R. (1993). *Working without a net*. Oxford University Press, New York.
- Friedman, K. and Shimony, A. (1971). Jaynes's maximum entropy prescription and probability theory. *Journal of Statistical Physics*, 3(4):381–384.
- Gillies, D. (2000). *Philosophical theories of probability*. Routledge, London and New York.
- Grove, A. J. and Halpern, J. Y. (1997). Probability update: conditioning vs. cross-entropy. In *Proceedings of the 13th Annual Conference on Uncertainty in Artificial Intelligence*, pages 208–214. Morgan Kaufmann, San Francisco, CA.
- Howson, C. (1997). Bayesian rules of updating. *Erkenntnis*, 45:195–208.
- Howson, C. (2000). *Hume's problem: induction and the justification of belief*. Clarendon Press, Oxford.
- Jaynes, E. T. (2003). *Probability theory: the logic of science*. Cambridge University Press, Cambridge.
- Keynes, J. M. (1921). *A treatise on probability*. Macmillan (1948), London.
- Lange, M. (1999). Calibration and the epistemological role of Bayesian conditionalization. *The Journal of Philosophy*, 96:294–324.
- Ramsey, F. P. (1926). Truth and probability. In Kyburg, H. E. and Smokler, H. E., editors, *Studies in subjective probability*, pages 23–52. Robert E. Krieger Publishing Company, Huntington, New York, second (1980) edition.
- Rosenkrantz, R. D. (1977). *Inference, method and decision: towards a Bayesian philosophy of science*. Reidel, Dordrecht.
- Rowbottom, D. P. (2007). The insufficiency of the Dutch Book argument. *Studia Logica*, 87:65–71.
- Rowbottom, D. P. and Bueno, O. (2008). How to change it: modes of engagement, rationality, and stance voluntarism. *Synthese*, This volume.
- Seidenfeld, T. (1979). Why I am not an objective Bayesian. *Theory and Decision*, 11:413–440.
- Seidenfeld, T. (1986). Entropy and uncertainty. *Philosophy of Science*, 53(4):467–491.
- Shimony, A. (1973). Comment on the interpretation of inductive probabilities. *Journal of Statistical Physics*, 9(2):187–191.
- Shimony, A. (1985). The status of the principle of maximum entropy. *Synthese*, 63:35–53.
- Skyrms, B. (1985). Maximum entropy inference as a special case of conditionalization. *Synthese*, 63:55–74.
- Skyrms, B. (1987). Coherence. In Rescher, N., editor, *Scientific inquiry in philosophical perspective*, pages 225–242. University Press of America, Lanham, Maryland.

- Teller, P. (1973). Conditionalisation and observation. *Synthese*, 26:218–258.
- Uffink, J. (1996). The constraint rule of the maximum entropy principle. *Studies in History and Philosophy of Modern Physics*, 27:47–79.
- van Fraassen, B., Hughes, R., and Harman, G. (1986). A problem for relative information minimisers, continued. *British Journal for the Philosophy of Science*, 37:453–475.
- van Fraassen, B. C. (1981). A problem for relative information minimizers in probability kinematics. *British Journal for the Philosophy of Science*, 32:375–379.
- van Fraassen, B. C. (1984). Belief and the will. *Journal of Philosophy*, 81(5):235–256.
- van Fraassen, B. C. (1987). Symmetries of personal probability kinematics. In Rescher, N., editor, *Scientific inquiry in philosophical perspective*, pages 183–223. University Press of America, Lanham, Maryland.
- van Fraassen, B. C. (1989). *Laws and symmetry*. Clarendon Press, Oxford.
- van Fraassen, B. C. (2002). *The empirical stance*. Yale University Press, New Haven.
- Williams, P. M. (1980). Bayesian conditionalisation and the principle of minimum information. *British Journal for the Philosophy of Science*, 31:131–144.
- Williamson, J. (2005). *Bayesian nets and causality: philosophical and computational foundations*. Oxford University Press, Oxford.
- Williamson, J. (2007a). Inductive influence. *British Journal for the Philosophy of Science*, 58(4):689–708.
- Williamson, J. (2007b). Motivating objective Bayesianism: from empirical constraints to objective probabilities. In Harper, W. L. and Wheeler, G. R., editors, *Probability and Inference: Essays in Honour of Henry E. Kyburg Jr.*, pages 151–179. College Publications, London.
- Williamson, J. (2008a). Objective Bayesian probabilistic logic. *Journal of Algorithms in Cognition, Informatics and Logic*, 63:167–183.
- Williamson, J. (2008b). Objective Bayesianism with predicate languages. *Synthese*, 163(3):341–356.
- Williamson, J. (2009). Philosophies of probability. In Irvine, A., editor, *Handbook of the philosophy of mathematics*. Elsevier, Amsterdam. Handbook of the Philosophy of Science volume 4.