

Journal of Philosophy, Inc.

Is Ignorance Bliss?

Author(s): Joseph B. Kadane, Mark Schervish and Teddy Seidenfeld

Source: *The Journal of Philosophy*, Vol. 105, No. 1 (Jan., 2008), pp. 5-36

Published by: [Journal of Philosophy, Inc.](#)

Stable URL: <http://www.jstor.org/stable/20620069>

Accessed: 08/02/2014 09:41

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Journal of Philosophy, Inc. is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*.

<http://www.jstor.org>

THE JOURNAL OF PHILOSOPHY

VOLUME CV, NO. 1, JANUARY 2008

IS IGNORANCE BLISS?*

“...where ignorance is bliss, 'tis folly to be wise.”

—Thomas Gray

We are surrounded by homilies supporting the idea that human progress is intimately tied to the increase of our knowledge. Research is justified that way, as is our enormous investment in education at all levels, as is our support of public libraries of all sorts. Generally, we are willing, as a society, to pay a lot to make information available to ourselves and others.

It is easy to see that there may be some kinds of information that one might rationally pay not to have, if paying that price would prevent others from learning the information as well. For example, if there were a cheap and effective means of killing all living beings, we would have no use for that knowledge ourselves, but would pay a high price to prevent others who might want to use it from gaining the knowledge of how to do such a thing. This paper focuses on a simpler problem, whether there is information a rational person would pay not to have for that person's private use.

The paper is organized as follows: Section I reviews the standard Bayesian theorem about the non-negative expected worth of cost-free sample information. Section II describes a more detailed Bayesian model that permits a more subtle understanding of what “cost-free” means in this context. Section III gives formal definitions of the value of information and cost-free information based on the formulation in section II. Section IV shows by examples that a Bayesian might pay not to receive cost-free information if the prior is improper or not countably additive. Section V reviews three extensions of the

* Research partially supported by NSF grant DMS-0139911.

0022-362X/08/0501/5-36

© 2008 The Journal of Philosophy, Inc.

Bayesian idea to sets of probabilities, and shows that each requires or permits paying not to receive sample information. Section VI concludes with a general discussion of whether refusing to pay to be shielded from information is a criterion that should be imposed on a reasonable decision theory, and, if so, what its consequences would be. Finally, the Appendix discusses the sense of Bayesian updating used in this paper.

I. THE EXPECTED VALUE OF SAMPLE INFORMATION

There is an appealing intuitive argument supporting the idea that one should not pay not to receive cost-free information. Informally stated, it goes as follows. Suppose that one faces a decision at time t_0 , to be made under uncertainty. In the absence of an offer to acquire new information, one is prepared to make decision d at t_0 with expected utility u_d . One might imagine acquiring new cost-free information x at a later time t_1 , but then still making decision d regardless of x . From the perspective of the initial choice point at time t_0 , this is the same as choosing d , and has expected utility u_d . But it might be that at time t_1 , given x , there is a better decision than d to choose. Then from the perspective of the initial choice point, at t_0 , this more complicated plan has a higher expected utility than does the choice of d . Thus, having the information x can not be harmful to one's expected utility, and it might be helpful. From the perspective of the initial choice, at t_0 , the expected worth of the new information x cannot be negative.

This argument can be made rigorous as follows. Let $U(d, \theta)$ be your utility function, which depends on both your decision d and on $\theta \in \Theta$, the unknown state of the world. You have a distribution that jointly describes your probabilities for the data $x \in \mathcal{X}$ and for θ , with joint density $p(x, \theta)$. Assume that there is no cost associated with learning x . Without the data x , you would choose d to maximize

$$(1) \quad \int_{\mathcal{X}} \int_{\Theta} U(d, \theta) p(x, \theta) d\theta dx.$$

If you were to learn the data x , you would maximize your utility with respect to your conditional distribution $p(\theta|x)$, that is, maximize

$$(2) \quad \int_{\Theta} U(d, \theta) p(\theta|x) d\theta,$$

which has expectation, with respect to the unseen value of x ,

$$(3) \quad \int_{\mathcal{X}} \left[\max_d \int_{\Theta} U(d, \theta) p(\theta|x) d\theta \right] p(x) dx,$$

where $p(x)$ is the probability distribution of the data x . The intuitive argument above suggests that (3) is no smaller than (1).

To show this, let d^* be a maximizer of (1). [The argument works just as well, if such a d^* does not exist, for d^* to be an ε -maximizer of (1).] Then for each x in (2),

$$(4) \quad \max_d \int_{\Theta} U(d, \theta) p(\theta | x) d\theta \geq \int_{\Theta} U(d^*, \theta) p(\theta | x) d\theta.$$

Integrating both sides of (4) with respect to x , yields

$$(5) \quad \int_X \left[\max_d \int_{\Theta} U(d, \theta) p(\theta | x) d\theta \right] p(x) dx \geq \int_X \int_{\Theta} U(d^*, \theta) p(\theta | x) d\theta p(x) dx \\ = \int_X \int_{\Theta} U(d^*, \theta) p(\theta, x) d\theta dx \\ = \max_d \int_X \int_{\Theta} U(d, \theta) p(\theta, x) d\theta dx,$$

as claimed. This result is a familiar one in Bayesian theory, with a long history.¹

II. BAYESIAN MODELS INCORPORATING PAYING NOT TO LEARN

I hold a ticket to a mystery play, and take a taxi to the theater. The taxi driver knows who-done-it and offers (threatens?) to inform me unless the tip is sufficiently large. Does it make sense to pay to avoid learning this information?

It is clear that there are circumstances in which I might pay, and that the result rehearsed in section I seems to contradict this. To understand the issue better, we introduce a richer decision theory that incorporates knowledge into the utility function.

In a typical decision problem, we are interested in the results at a particular time T in the future after any decisions have been made and after enough of the state of nature has been learned to determine the impact of our decisions. In order to be flexible regarding, among other relevant matters, a “small world” versus a “grand world” framing of a decision problem, we shall let the state of nature ω describe no more than we need to make sense out of each individual decision problem. We shall divide the collection of future outcomes into three parts.

- The “state of nature” ω will consist of those quantities about which a decision maker (DM) is uncertain at some time during the decision problem. The collection of all states of nature is Ω . States of nature can include unchanging facts as well as stochastic processes. They can

¹See I. J. Good, “On the Principle of Total Evidence,” *British Journal of the Philosophy of Science*, xvii (1967): 319–21; Howard Raiffa and Robert Schlaifer, *Applied Statistical Decision Theory* (Graduate School of Business Administration, Harvard University, Division of Research, 1961). See also Frank Ramsey, “Weight or the Value of Knowledge,” *British Journal for the Philosophy of Science*, xli (1990): 1–4.

include amounts of wealth that might result from gambles. They can also include inputs to decisions that might be made at later stages in a decision problem. We assume that the state of nature stays the same over time, but our knowledge and/or beliefs about it can change. In particular, if the state of nature includes a random variable whose value becomes known later, we interpret this as a change in our knowledge about the random variable rather than a change in the random variable.

- The “knowledge base” is the collection of all information known by the DM. If more than one decision maker is being contemplated, each DM has his/her own knowledge base. The knowledge base might include the knowledge that certain choices are available now or for future decision problems as well as the various random variables that might be available for observation and when they might become available. The knowledge base can, and often will change over time. We will denote the knowledge base sequence $\Psi = \{\psi_t : 0 \leq t \leq T\}$, where the subscript t denotes time. We use $t = 0$ to stand for “now” when the DM starts thinking about the decision problem, and $t = T$ stands for the time at which all relevant information becomes available and the DM has experienced the impact of the decision. The impact of the decision can be cumulative over a time interval, and it is possible that one might want to let $T = \infty$ to indicate that the impact never stops accumulating. We assume that, at each time t , the DM has no uncertainty about the corresponding value of the knowledge base ψ_t . We do not, however, require that the knowledge base be nondecreasing over time. That is, the theory allows that a DM might forget something that was known earlier. We will use Ψ with various sub- and superscripts to denote sets of possible values for ψ_t at various times.
- The set of acts, gambles, choices, or decisions could easily depend on what information ψ_t the DM has at each time t . For example, one cannot choose a decision that explicitly depends on a random variable that one will not have observed at the time that the decision must be implemented. So, we denote the set of decisions available at time t when the knowledge base is ψ_t as $\mathcal{D}(\psi_t)$. In a normal-form sequential decision problem, $\mathcal{D}(\psi_0)$ consists of sequences of choices for each stage in the problem where each choice can be a function of what will be known at that stage. At time t , $\mathcal{D}(\psi_s)$ for $s > t$ might be random because it could depend on things that might be learned between times t and s . In an extensive form decision problem, we treat each stage as a separate decision problem with the same T and the same utility function but with a changing information base sequence.

EXAMPLE 1. Suppose that all I care to think about now is where I will be one hour from now. Also, I care only to distinguish three possible descriptions of where I will be, namely at work, at home, or elsewhere. We could let Ω consist of the three places I could be, and we could let $\mathcal{D}(\psi_0)$ consist of three choices, namely to stay at work, go home, or go

somewhere else. In this case, we assume that my decision will determine the state of nature with probability 1. At the time that the choice of decision must be made, I cannot have a probability distribution over states because the state is a function of my decision. However, suppose that I do not have to decide which action to take for 1/2 hour. Then, I could have a marginal probability over Ω for 1/2 hour until it becomes time to decide. During that 1/2 hour, I might learn things that help to influence which choice I make.

At each time t , let Ψ_t be the set of all possible knowledge bases. Then the utility at time t is a bounded function U defined on

$$\Omega \times \bigcup_{\psi_t \in \Psi_t} [\{\psi_t\} \times \mathcal{D}(\psi_t)].$$

That is, the utility at time t is a bounded function of the state, the knowledge base, and the decision, denoted $U(\omega, \psi_t, d)$. As the knowledge base is allowed to (but not required to) accumulate the DM's experiences over time, the utility at time t can depend on what has happened in the past. Since the knowledge base is allowed to include anticipated experiences in the future, the utility at time t can depend on what the DM anticipates might happen in the future. If $\mathcal{D}(\psi_t)$ is empty, that is, there are no decisions to be made, we shall fill the third argument with the symbol \square .

EXAMPLE 2. (SIMPLE DECISION PROBLEM) Consider a statistical decision problem with action space \aleph , parameter space Θ , and loss function L , where $L(\theta, a)$ stands for the loss that results from choosing action a when θ is the parameter. Assume that the loss function does not change with time. Assume that no relevant data will be observed by time T , so that the knowledge base is $\psi_t = \psi_0$ for $0 \leq t \leq T$. Let $\mathcal{D}(\psi_0)$ be the set of randomized rules, that is probability distributions over \aleph . Assume that \aleph is a nice enough space so that each randomized rule δ can be constructed as the distribution of a function f_δ of a single random quantity R independent of everything else and taking values in a set \mathcal{R} . (All Polish spaces are of this type.) In particular, a non-randomized rule $\delta_a(A) = I_A(a)$ for $a \in \aleph$ and $A \subseteq \aleph$ corresponds to $f_{\delta_a}(r) = a$ for all r . Let $\Omega = \Theta \times \mathcal{R}$, where the last coordinate will hold the value of R . Let Q be the DM's prior distribution over Θ . Then, for each nonrandomized choice a , we can denote the utility at time t by $U((\theta, r), \psi, \delta_a) = -cL(\theta, a) + b$ for some $c > 0$ and real b . For randomized δ , $U((\theta, r), \psi, \delta) = -cL(\theta, f_\delta(r)) + b$. The expectation of this, over the distribution ν of R , involves the usual loss of a randomized rule:

$$\int U((\theta, r), \psi, \delta) \nu(dr) = \int_{\aleph} [-cL(\theta, a) + b] \delta(da) = -cL(\theta, \delta) + b.$$

To arrive at this point in the analysis of the problem, one does not make any use of the prior distribution over Θ . Hence, this much of the analysis is equally suitable for both Bayesians and frequentists.

EXAMPLE 3. (STATISTICAL DECISION PROBLEM) Expand Example 2 to include data X taking values in a set \mathcal{X} . In this example, let $\mathcal{D}(\psi_0)$ be a set of deterministic functions from \mathcal{X} to \aleph so that only non-randomized rules are available. Let $\Omega = \Theta \times \mathcal{X}$. In the typical statistical decision problem, we would have

$$(6) \quad U((\theta, x), \psi_t, \delta) = -cL(\theta, \delta(x)) + b,$$

for all x and all t , so that utility depends on the data $X = x$ only through the value of $\delta(x)$ and does not change with time. Up to this point, the analysis does not depend on any probability distributions over Θ . Assume now that the DM has a subjective distribution Q over Θ and a regular conditional distribution $P(\cdot|\theta)$ over \mathcal{X} given θ for each $\theta \in \Theta$. Let $Q(\cdot|x)$ represent the posterior distribution over Θ given $X = x$. From (6) it follows that a Bayesian will, for each $x \in \mathcal{X}$, choose $\delta(x)$ equal to that a that minimizes $\int_{\mathcal{X}} L(\theta, a) Q(d\theta|x)$, the usual posterior risk. A non-Bayesian will compute

$$(7) \quad \int_{\mathcal{X}} U((\theta, x), \psi_t, \delta) P(dx|\theta) = -cR(\theta, \delta) + b,$$

where R is the classical risk function. For a maximin solution, one would compute for each δ the infimum of (7) over θ and then choose that δ with the largest infimum.

EXAMPLE 4. (SEQUENTIAL DECISION PROBLEM IN NORMAL FORM) Consider a normal form sequential decision problem in which the decision maker gets to decide how much data to observe. Suppose that the n th observation costs $c_n > 0$ and that the terminal loss is $L(\theta, a)$, where a is the terminal action and $\theta \in \Theta$ is an unknown parameter. To allow for a terminal decision before any data are observed, let $c_0 = 0$. Let X be the potential data sequence taking values in a sequence space \mathcal{X} . Let $\mathcal{D}(\psi_0)$ consist of decision rules $\delta = (N, d)$. Here N is either a positive-integer-valued function of data sequences $x \in \mathcal{X}$ such that $\{N = n\}$ is measurable with respect to the first n coordinates of x or $N = 0$ with probability 1 (meaning that no data will be observed). Also, $d = (d_0, d_1, d_2, \dots)$ is a sequence of functions such that each d_n for $n \geq 1$ is a function of the first n coordinates of x and takes values in \aleph and d_0 is some element of \aleph . The interpretation of such a rule is the following. If the DM decides to observe no data, d_0 will be the terminal action. Otherwise the DM observes $N(X) > 0$ observations and takes terminal action $d_{N(X)}(X)$. Now, $\Omega = \Theta \times \mathcal{X}$. To

express the usual sequential decision problem in the present framework, we assume that

$$U((\theta, x), \psi_t, \delta) = -c \left[L(\theta, d_{N(x)}(x)) - \sum_{n=0}^{N(x)} c_n \right] + b,$$

for some real b and $c > 0$. Once again, we have assumed that utility depends on the value of $X = x$ only through $\delta(x)$.

EXAMPLE 5. (SEQUENTIAL DECISION PROBLEM IN EXTENSIVE FORM) Consider the same general situation as in Example 4, but this time, assume that the DM gets to reconsider her choice after each coordinate of X is observed. That is, the sample size N need not be a predetermined measurable function of X such that $\{N = n\}$ is a function of the first n coordinates of X . It is more convenient to express an extended form problem as a sequence of one-stage decision problems in which the utility at each stage still depends on the utilities at other stages. To be precise, let t_0 be now and let $t_1 < t_2 < \dots$ be a sequence of times at which additional data might be collected depending on what decisions are made at earlier times. Of course $t_0 < t_1$. At each time t_i ($i = 0, 1, \dots$) at which a terminal decision has not yet been made, the DM has the option of choosing to observe more data X_{i+1} at time $i + 1$ or of making a terminal decision at the present time. If we let N denote the number of observations that will eventually be chosen, then N is random until a terminal decision is made. Similarly, we will let A denote the terminal action eventually chosen so that A is random until that time when a terminal action is finally chosen. For the decision problem at each time t_i , the set of possible decisions available is either empty (if a terminal decision was made in an earlier decision problem) or it consists of the terminal actions \aleph plus the decision to continue sampling. Using the notation of Example 4, the utility for the decision problem at each time t_i prior to a terminal decision is

$$U((\theta, x), \psi_t, \delta) = \begin{cases} -c \left[L(\theta, a) - \sum_{n=0}^i c_n \right] + b & \text{if } \delta \text{ chooses terminal action } a, \\ -c \left[L(\theta, A) - \sum_{n=0}^N c_n \right] + b & \text{if } \delta \text{ chooses to observe more data.} \end{cases}$$

In the framework of Examples 4 and 5, the pre-data value of a particular decision rule is the same whether the decision problem is put into normal form or into extensive form. In the normal form there is only one decision point among rules and each rule involves potentially many contingent directions for what to do based on the outcomes of subsequent observations. In extensive form, there are many decision points, each one corresponding to a change in avail-

able information. Bayesian analysis yields the same assessment for a decision rule regardless in which form the decision is cast. However, for the decision theories discussed in section v, where uncertainty is represented by sets of probability distributions, the two forms are not equivalent.²

III. THE COST AND VALUE OF INFORMATION

Section I discusses the value of cost-free information without formally defining cost-free information. To assess the cost of observing some information or the value of the information, one needs to be able to imagine the possibility of observing (and possibly using) the information as well as the possibility of not observing it. Hence, we set up decision problems in which there is at least one stage at which the DM can choose whether or not to observe the information. Since the cost and/or value of information can change over time (imagine observing a relevant sample after having to make a decision), we attach a time of observation to any definition of cost or value. Furthermore, the cost or value of information is relative to a specific decision problem.

The value of any information depends on what one does with it, so any definition will have to take into account how information is used. For example, in a Bayesian analysis, one tries to maximize expected utility, while others might try to maximize the infimum of the utility function. We will assume that the DM wants to maximize some nondecreasing functional h of the utility function. Examples of h include the integral with respect to a probability over ω and the minimum over all ω . Different functionals h correspond to different theories of decision making. To be precise, let h be a real-valued function defined on the space of all bounded functions $g : \Omega \rightarrow \Re$ that satisfies $h(g_1) \leq h(g_2)$ whenever $g_1(\omega) \leq g_2(\omega)$ for all ω . We assume that the DM with knowledge base sequence ψ chooses a decision rule δ in order to maximize $h(U(\cdot, \psi, \delta))$ for some time t . There are principles of decision making that do not fit this description, but many popular ones do.

To distinguish cost from value, we think of the net value of information as some measure of the change in utility one would achieve by observing the information and using it however one saw fit as opposed to not observing it. We would like to define the cost to be the

²See Teddy Seidenfeld, "When Normal and Extensive Form Decisions Differ," in Dag Prawitz, Brian Skyrms, and Dag Westerstaahl, eds., *Logic, Methodology and Philosophy of Science IX* (New York: Elsevier, 1994), pp. 451–63, for in-depth discussion of how normal and extensive form decision problems differ.

change in utility one would achieve by observing the information but then ignoring it in making decisions. This is problematic because it may not be possible to ignore some information after observing it. For this reason, we take a less ambitious approach and merely define cost-free information in section III but do not quantify the cost of information that is not cost-free.

III.1. The Net Value of Information. Suppose that a DM might be able to observe some information X at a time $t_1 \geq 0$ and is faced with a decision problem as described earlier. The DM wants to say how valuable is the information X at time t_1 in the decision problem. To do this, we embed two versions of the original decision problem in a 2-stage problem in normal form (called the extended decision problem) in which the first stage involves a decision at time 0 of whether or not to observe X at time t_1 . To be clear, we imagine the DM deciding at time 0 whether or not to observe X , but the actual observation of X may occur later at time t_1 . The second stage of the problem will be one of the two versions of the original decision problem, one in which X is observed (version 1) and one in which X is not observed (version 0). The versions of the original decision problem may themselves be sequential, but this will be immaterial for the current discussion. The set of decision rules available in the extended decision problem at time $t = 0$ is $\mathcal{D}(\psi_0) = (N, d_0, d_1)$ where $N = 0$ means to not observe X , $N = 1$ means to observe X , and d_i is a decision rule for version i ($i = 0, 1$). This matches the notation used in Example 4. To make it easier to discuss the two versions together, let $\psi = \{\psi_{it}\}_{t=0}^T$ be a knowledge base for which it is known that X is not to be observed at time t_1 , and let $\psi^* = \{\psi_{it}^*\}_{t=0}^T$ denote the alternative knowledge base in which it is known that X will be or has been observed at time t_1 . We assume that ψ_0^* is the same as ψ_0 aside from the knowledge that X will be observed, so that the two versions start with essentially the same information before the decision is made whether or not to observe X .

We assume that the DM is able to determine what she would do if X were observed and what she would do if X were not observed. Formally, we assume that both version 0 and version 1 of the original decision problem can be solved by the DM. This could be possible even in cases in which the DM already knows that she will observe X but still knows what she would have done if X were not to be observed. The net value of the information X is then defined to be the difference

$$h(U(\cdot, \psi_1^*, \delta_1)) - h(U(\cdot, \psi_1, \delta_0)),$$

where δ_i is the decision rule that the DM would use if forced to solve version i of the original decision problem for $i = 0, 1$ at time t .

EXAMPLE 6. (MAXIMIN DECISION MAKER) Suppose that the DM chooses decision rules using a maximin criterion. To be more specific, consider a traditional decision problem such as the ones described in Example 3 where U is given by (6). Suppose that a decision must be made at time $t \geq t_1$. A maximin decision maker, in such a problem, uses the functional

$$(8) \quad h(g) = \inf_{\theta \in \Theta} \int_{\mathcal{X}} g(\theta, x) P(dx | \theta).$$

Substituting the utility from (6) into (8) we obtain

$$h(U(\cdot, \psi_t^*, \delta)) = \inf_{\theta \in \Theta} \int_{\mathcal{X}} [-cL(\theta, \delta(x)) + b] P(dx | \theta),$$

for each rule $\delta \in \mathcal{D}(\psi_t^*)$. In version 0, all decision rules δ have to be constant as a function of x . Hence, we have

$$h(U(\cdot, \psi_t, \delta)) = \inf_{\theta \in \Theta} [-cL(\theta, \delta) + b],$$

for each $\delta \in \mathcal{D}(\psi_t)$. So, the net value of X would be

$$\sup_{\delta \in \mathcal{D}(\psi_t^*)} \inf_{\theta} (-cR(\theta, \delta) + b) - \sup_{\delta \in \mathcal{D}(\psi_t)} \inf_{\theta} (-cL(\theta, \delta) + b).$$

where $R(\theta, \delta)$ is the classical risk function

$$R(\theta, \delta) = \int_{\mathcal{X}} L(\theta, \delta(x)) P(dx | \theta).$$

As a specific example, consider the following augmentation of the example presented by Leonard Savage³ in which we give a specific data distribution. Let $\Omega = \Theta \times \mathcal{X}$ where the parameter space $\Theta = \{1, 2\}$ and the sample space $\mathcal{X} = \{1, 2, 3\}$. Let the space of terminal decisions $\mathfrak{K} = \{1, 2\}$. At time t_0 the DM must decide whether to observe X taking values in \mathcal{X} . At time $t_1 > t_0$, the DM has to choose one of the two terminal actions from \mathfrak{K} . Let X have the following conditional distribution given θ :

$$\Pr(X = i | \theta = j) = \begin{cases} 0.5 & \text{if } i = j = 1 \text{ or } i = j = 2 \text{ or } i = 3, \\ 0 & \text{otherwise.} \end{cases}$$

The random variable X identifies θ with probability 0.5 (when $X \neq 3$) and it is irrelevant with probability 0.5 (when $X = 3$). Let the utility

³ Savage, *The Foundations of Statistics* (New York: Wiley, 1954), p. 170.

have the form of (6) with $c = 1$, $b = 0$ and consider the following loss function regardless of whether X is observed:

$$L(\theta, a) = \begin{cases} 1 & \text{if } a = 1, \\ 10 & \text{if } \theta = 1, a = 2, \\ -1 & \text{if } \theta = 2, a = 2. \end{cases}$$

That is,

$$U((\theta, x), \psi_t, \delta) = -L(\theta, \delta(x)),$$

for all t , θ , and x . Intuition suggests that X has value for the DM unless θ is known with certainty. If X is not observed, then the maximin decision is to choose terminal action $a = 1$. In this case $h(U(\cdot, \psi_b, 1)) = -1$. If X is to be observed, the maximin rule is any Bayes rule with respect to the least favorable prior. The least favorable prior is $\Pr(\theta=1)=1$, and the collection of Bayes rules is the set of rules δ_α where, for each α , $\delta_\alpha(1) = \delta_\alpha(3) = 1$ and $\delta_\alpha(2)$ randomizes between 1 and 2 with $\Pr(\delta_\alpha(2) = 1) = \alpha$ for arbitrary $\alpha \in [0, 1]$. The risk functions of these rules are

$$R(\theta, \delta_\alpha) = \begin{cases} 1 & \text{if } \theta = 1, \\ \alpha & \text{if } \theta = 2. \end{cases}$$

For all of these rules, $h(U(\cdot, \psi_t, \delta_\alpha)) = -1$, so the maximin decision maker assigns 0 net value to the information X , contrary to intuition.

Savage claims that maximin decision making is “utterly untenable for statistics” based on the fact that the seemingly relevant data in Example 6 has 0 value (*ibid.*, p. 170). (We will see in section III.2 that the X in that example is cost-free.) Savage contrasts the maximin rule with maximin-regret, which we describe next.

EXAMPLE 7. (MAXIMIN-REGRET DECISION MAKER) A variation on maximin decision making is to shift the utility function so that the maximum value is the same in every state, typically 0. In symbols, replace U by

$$(9) \quad U'(\omega, \psi_t, \delta) = U(\omega, \psi_t, \delta) - \max_{\delta} U(\omega, \psi_t, \delta).$$

The maximin-regret decision maker then behaves the same as a maximin decision maker whose utility is U' instead of U . In particular, for the calculations that appear in Examples 3 and 6, we replace $L(\theta, \cdot)$ by $L(\theta, \cdot) - \min_a L(\theta, a)$ everywhere and set $b = 0$ wherever it occurs. The functional h is still (8).

Now, reconsider the example of Savage (*ibid.*, p. 170) from this point of view. The modified loss function is

$$L'(\theta, a) = \begin{cases} 9 & \text{if } \theta = 1, a = 2, \\ 2 & \text{if } \theta = 2, a = 1, \\ 0 & \text{otherwise.} \end{cases}$$

If the DM does not observe X , the maximin-regret decision is the randomized rule δ that chooses $a = 1$ with probability $9/11$ and chooses $a = 2$ with probability $2/11$. For this rule, $h(U(\cdot, \psi_t, \delta)) = -18/11$. However, if X is observed, the risk function of the general randomized rule δ_β that chooses $a = 1$ with probability β_i when $X = i$ is observed equals

$$R'(\theta, \delta_\beta) = \begin{cases} \frac{9}{2}(2 - \beta_1 - \beta_2) & \text{if } \theta = 1, \\ \beta_2 + \beta_3 & \text{if } \theta = 2. \end{cases}$$

The maximin rule corresponds to $\beta_1 = 1$, $\beta_2 = 0$, and $\beta_3 = 9/11$ which has $h(U(\cdot, \psi_t, \delta_\beta)) = -9/11$. So, the maximin-regret decision maker assigns net value $(-9/11) - (-18/11) = 9/11$ to the information X .

Savage suggests that maximin-regret decision making is not as prone to the objection of assigning zero value to relevant data (*ibid.*, p. 200). However, Giovanni Parmigiani⁴ shows that there are examples in which maximin decision making gives positive net value to data that are assigned zero net value by maximin-regret.

EXAMPLE 8. (BAYESIAN DECISION MAKER) Suppose that the DM wants to choose the decision rule δ at time t that maximizes expected utility. Let the DM's distribution (now) for ω be P . Then $h(g) = \int_{\Omega} g(\omega) P(d\omega)$ and

$$h(U(\cdot, \psi_t, \delta)) = \int_{\Omega} U(\omega, \psi_t, \delta) P(d\omega).$$

Next, consider the special case of Example 3 in which $\omega = (\theta, x)$. Then the net value of the information X is the difference between the extreme terms in (5) if the following two conditions hold:

- the utility has the simple form $U((\theta, x), \psi_t, \delta) = U(\delta(x), \theta)$ for all t , and
- (X, θ) has a joint density $p(x, \theta)$ with respect to Lebesgue measure.

In Savage's example, suppose that the Bayesian DM starts with $\Pr(\theta = 1) = p$. Without observing X , it is easy to see that a Bayesian DM will choose $a = 1$ if $p > 2/11$ and will choose $a = 2$ if $p < 2/11$. Either action can be chosen if $p = 2/11$. Call this Bayes decision δ_p . Then

$$(10) \quad h(U(\cdot, \psi_t, \delta_p)) = \begin{cases} 1 - 11p & \text{if } p < 2/11, \\ 1 & \text{if } p \geq 2/11. \end{cases}$$

If X is observed, then the Bayesian DM chooses $a = 1$ if $X = 1$ and chooses $a = 2$ if $X = 2$. If $X = 3$, the DM uses the same δ_p she would

⁴Parmigiani, "Minimax, Information, and Ultrapessimism," *Theory and Decision*, xxxiii (1992): 241–52.

have used if X had not been observed. For this rule δ_p^* , $h(U(\cdot, \psi_b, \delta_p^*))$ equals one-half of (10) (corresponding to $X = 3$) plus $0.5 - p$ (corresponding to $X \neq 3$). The net value of X for the DM who believes $\Pr(\theta = 1) = p$ a priori is

$$(11) \quad \begin{cases} \frac{9}{2}p & \text{if } p < 2/11. \\ 1-p & \text{if } p \geq 2/11. \end{cases}$$

Notice that the largest possible net value occurs at $p = 2/11$ and equals $9/11$, the same net value assigned by the maximin DM. The least favorable distribution in the maximin-regret problem is $\Pr(\theta = 1) = 2/11$ for both version 0 and version 1.

EXAMPLE 9 (Γ -MAXIMIN DECISION MAKER) We will discuss Γ -maximin decision rules in more detail in section v. For now, we can understand Γ -maximin decision making as a variant on Bayesian decision making in the following sense. One replaces the single distribution P over states used by the Bayesian decision maker in Example 8 by a closed convex set \mathcal{P} of distributions. Then, one ranks decision rules by

$$(12) \quad h(U(\cdot, \psi_t, \delta)) = \inf_{P \in \mathcal{P}} \int_{\Omega} U(\omega, \psi_t, \delta) P(d\omega).$$

In Savage's example, suppose that the Γ -maximin DM chooses $\mathcal{P} = [p_1, p_2]$. If $p_1 \leq 2/11 \leq p_2$, then the DM behaves just like the maximin-regret DM because the infs in (12) will occur at the least-favorable distribution in both version 0 and version 1. If $p_2 < 2/11$, the DM will behave like the Bayesian in Example 8 with $p = p_2$. If $p_1 > 2/11$, the DM will behave like the Bayesian with $p = p_1$. The net value of X can be found for these last two cases in (11).

III.2. Cost-Free Information. This section formally defines what we mean by saying that observing information X at a specific time t_1 is cost-free. We will do this in a manner that does not depend on the decision principles to which one adheres. In order for information at a specific time to be cost free for every mode of decision making, we require several conditions. First, we require that it is possible not to observe X at time t_1 . Second, we require that the information be *decision-theoretically ignorable*, meaning that every decision rule that is available without the information is also available with the information so that one could ignore the information when choosing a decision rule. To make this more precise, let $\psi = \{\psi_t\}_{t=0}^T$ be a knowledge base sequence for which it is known that X is not to be observed at time t_1 , and let $\psi^* = \{\psi_t^*\}_{t=0}^T$ denote the alternative knowledge base sequence in which it is known that X will be or has been observed at time t_1 . To say that observing X at time t_1 is decision-theoretically

ignorable, we mean that for each decision time t and each $\delta \in \mathcal{D}(\psi_t)$ there is a $\delta^* \in \mathcal{D}(\psi_t^*)$ which does exactly what δ would do under all states of nature that are still possible, essentially ignoring X to the extent possible. To put it another way, learning X at time t_1 does not eliminate any decision rules from consideration. There is a subtle issue that should be addressed at this point. Suppose that, even if the DM does not observe X at time t_1 , she will nevertheless observe it at some later time. A simple example will illustrate the issue.

EXAMPLE 10. Let \aleph consist of two points $\aleph = \{a_0, a_1\}$. Let X be a Bernoulli random variable that we believe to be independent of everything that is relevant in this decision problem. Let $t_1 = 0$, and let $0 < t_2$. Suppose that X will be observed by time t_2 under all circumstances, but we have the option of observing it at time t_1 . Suppose also that, at time t_2 , we must choose between two decision rules $\delta_0 \equiv a_0$ and $\delta_1(X) = a_X$. That is, we can either pick action a_0 by choosing δ_0 or we can choose δ_1 which randomizes between a_0 and a_1 by using X : if $X = 0$ then we will pick a_0 and if $X = 1$ we will pick a_1 . So, the terminal action is not finally decided until time t_2 when X will definitely be known. If we choose to observe X at time t_1 and $X = 0$ is observed, then the only terminal action that will be available to us is a_0 regardless of whether we choose decision rule δ_0 or δ_1 . Nevertheless, the decision rule δ_1 is still available at time t_2 . It just happens that we already know that $\delta_1(X) = a_0$.

In Example 10, X is decision-theoretically ignorable because the decision rule δ_1 is still available when it comes time to choose regardless of the fact that we might already know that it will make the same choice as δ_0 . When it comes time to make our decision, we are in the same situation regardless of whether we observed $X = 0$ at time t_1 or at some later time.

Information X that is decision-theoretically ignorable will be called *cost-free* if, for every state of nature ω , every knowledge base sequence ψ , every time t , and every $d \in \mathcal{D}(\psi_t)$,

$$(13) \quad U(\omega, \psi_t, \delta) = U(\omega, \psi_t^*, \delta^*),$$

where δ^* was defined above as the decision rule in $\mathcal{D}(\psi_t^*)$ that makes the same choices as δ under all circumstance. In other words, if the DM ignores X , she will achieve the same utility as if X had not been observed. In particular, there can be no fee for merely observing X .

EXAMPLE 11. (RETURN TO EXAMPLE OF SAVAGE) This example is set up in the second half of Example 6. In the notation of this section, let $\mathcal{D}(\psi_{t_1})$ contain all randomized decision rules that are just probabilities over $\aleph = \{1, 2\}$, and let $\mathcal{D}(\psi_{t_1}^*)$ consist of all randomized decision rules that are functions of X . It is easy to see that for each rule δ_p in

$\mathcal{D}(\psi_{t_1})$ that chooses $a = 1$ with probability p and chooses $a = 2$ with probability $1 - p$, there is a $\delta_p^* \in \mathcal{D}(\psi_{t_1}^*)$ such that $\delta_p^*(x) = \delta_p$ for each $x = 1, 2, 3$. Hence, X is decision-theoretically ignorable. It is also easy to see that $U(\omega, \psi_b, \delta_p) = U(\omega, \psi_t^*, \delta_p^*)$ for all p, ψ , and t . Hence, the information X is cost-free in this example.

Now return to the example at the start of section II.

EXAMPLE 12. (THE TAXI DRIVER) Recall that I hold a ticket to a mystery play and take a taxi to the theater. The taxi driver knows who-done-it and offers to inform me unless the tip is sufficiently large. Does it make sense to pay to avoid learning this information? Assume that I am committed to attending the mystery so that I need not consider the possibility of doing something else. This makes the information offered by the taxi driver decision-theoretically ignorable because I have no decision points after I choose the tip. For this reason, we will write the utility as $U(\omega, \psi_b, \square)$ to indicate that no decisions are available in either version 0 or version 1 of the problem. If we determine that the information is not cost-free, then we are faced with the extended decision problem of whether to pay not to observe it. For simplicity, suppose that I am contemplating, at time 0, only two different tips, 0 and $v > 0$, where v is sufficiently large to prevent the taxi driver from telling me who-done-it. At present, I believe that attending the mystery will have a random (that is, unknown) effect on several of my emotions at certain times in the future with various probabilities that reflect my opinion of how well the mystery is produced and what I know at the time of the play. In particular, the play offers the possibility of suspense and surprise. If I do not know who-done-it, the probabilities of high levels of suspense and surprise are larger than if I do know. Similarly, if I do not know who-done-it, the probability of boredom is smaller than if I do know. A poorly written/produced play might cause boredom even if I do not know who-done-it. We can interpret the information about who-done-it to be data that might be added to my knowledge base at or before time T , when the mystery will be finished. The taxi driver has offered me the choice of which tip to give, which is assumed to be equivalent to whether or not to obtain the information of who-done-it. No further decision point is available after this decision is made.

To be explicit, let $\Omega = \{1, \dots, \ell\} \times \mathcal{Y}$, where ℓ is the number of characters in the play, and \mathcal{Y} is a space each of whose elements specifies a set of values for the various aspects of the quality of the play. Suppose that I have a distribution μ over the space \mathcal{Y} that gives my opinion of the play. Let ψ_0 be my current knowledge base, and assume that I will not learn anything else before I have to choose a tip at time t_1 , so that $\psi_{t_1} = \psi_0$. For each $i = 1, \dots, \ell, t > t_1$, and $x > 0$ let $\psi_{t,i,x}$

be the augmented knowledge base at time t assuming that I learned that character i done-it before attending and assuming that I have paid x as a tip. Let $\psi_{t,x}$ be the corresponding knowledge base if I do not learn who-done-it but pay a tip of x . To avoid certain philosophical difficulties, suppose that, prior to getting in the taxi I have already considered the possibility that I might learn who-done-it before seeing the play and that I had already considered the possibilities of giving various tips even if I had never dreamed that the two might be related. Assume, for simplicity, that neither $U((i,y), \psi_{t,x}, \square)$ nor $U((i,y), \psi_{t,i,x}, \square)$ depends on i for fixed $y \in \mathcal{Y}$ and fixed x . These assumptions merely mean that I have no particular interest in which character did it. Fix $t > t_1$ and $x > 0$ and let $U_x = \int_{\mathcal{Y}} U((i,y), \psi_{t,i,x}, \square) \mu(dy)$ and $U^x = \int_{\mathcal{Y}} U((i,y), \psi_{t,x}, \square) \mu(dy)$. These are respectively the expected utilities that I would have calculated at time t after giving a tip of x under the assumptions of learning and not learning who-done-it before seeing the play.

It makes sense that $U_x < U^x$ for all x , that is, it is better not to know who-done-it, regardless of what tip I might give. Also, $U_x > U_w$, and $U^x > U^w$ whenever $x < w$, meaning that my utility goes down with the size of the tip, both in the case of learning and of not learning who-done-it. The expected utility of choosing tip v is then U_v while the expected utility of choosing tip 0 is U_0 . The fact that $U_v < U^v$ makes the information not cost-free. Which tip we should give hinges on whether or not $U^v > U_0$, the only comparison not fixed so far by the problem description. This comparison is between the disutility of the large tip v , on the one hand, to the disutility of ruining the experience of the play on the other.

Examples 4, 5, and 12 are cases in which the data are not cost-free.

IV. THE VALUE OF COST-FREE INFORMATION

Assume that X is cost-free. Then a DM who maximizes expected utility will not want to pay to avoid learning X . To see this, argue as follows. Let $P(dx, d\omega)$ stand for the joint distribution of X and ω . Even though we have assumed that X is a function of ω , this joint distribution still makes sense as a measure. Also, let $P(d\omega | x)$ stand for the posterior distribution of ω given $X = x$. Let δ^* be the decision that the DM would make without observing X , and let $\delta^{*'}$ be the corresponding rule that observes X but ignores it. Then for each t

$$(14) \quad \int_{\mathcal{X}} \int_{\Omega} U(\omega, \psi_t, \delta^*) P(dx, d\omega)$$

is the DM's expected utility at time t without observing X . If $X = x$ were observed, the DM would choose the decision δ to maximize

$$(15) \quad \int_{\Omega} U(\omega, \psi_{t,x}, \delta) P(d\omega | x).$$

For each x , we use (13) to generalize (4) as

$$(16) \quad \max_{\delta} \int_{\Omega} U(\omega, \psi_{t,x}, \delta) P(d\omega | x) \geq \int_{\Omega} U(\omega, \psi_{t,x}, \delta^*) P(d\omega | x) \\ = \int_{\Omega} U(\omega, \psi_t, \delta^*) P(d\omega | x).$$

The rest of the argument is virtually identical to that in section 1.

The result above (as well as the one in section 1) applies only to joint probability distributions on the data and state of nature that are countably additive and proper, as the following examples show. Both of these examples assume that the available data are cost-free, that is, the utility function is the same before and after observing the data. What fails from the result above and that of section 1 is that finitely additive and improper joint distributions do not always factor into marginals and conditionals the way that countably additive joint distributions do.

EXAMPLE 13. (FINITELY, BUT NOT COUNTABLY ADDITIVE DISTRIBUTIONS) Suppose that there are two states of the world, A and B , each of which has probability 0.5 in your current opinion. Imagine that you can observe a positive integer N . If A is true, the integer N has a geometric distribution, as follows:

$$P(n|A) = (1/2)^n,$$

for each integer $n > 0$. However, if B is the case, the integer N is uniformly distributed on the integers in your opinion. There are many such finitely additive distributions,⁵ but they all have the property that

$$P(n|B) = 0$$

for each integer $n > 0$.

If a particular integer, say $N = 3$, is observed, an easy application of Bayes's Theorem shows that

$$P(A|N=3) = 1$$

and in fact, this is true whatever value n of N is observed. Hence you are in the peculiar state of belief that although your prior is even between A and B , you know that conditional on the observation of N , regardless the value of N , you believe with certainty that A is true.

⁵ See Joseph B. Kadane and Anthony O'Hagan, "Using Finitely Additive Probability: Uniform Distributions on the Natural Numbers," *Journal of the American Statistical Association*, xc (1995): 626–31; also Oliver Schirokauer and Kadane, "Finitely Additive Distributions on the Natural Numbers: Shift Invariance," *Journal of Theoretical Probability*, xx, 3 (2007): 429–41.

Which then is your prior, what you believe now, or what you know you would believe if you could observe N ?⁶

Suppose that you currently hold a ticket that pays \$1 if B is true, and $-\$1$ if A is true. Currently your expected winnings are \$0. However, you know that if you were to obtain the integer N , your expected winnings would be $-\$1$. Would you pay \$0.50 not to receive the data? It seems that you would have to.

Where does the proof given above fail for finitely but not countably additive distributions? The proof fails because conditional probability does not work the same way for finitely additive and countably additive distributions. Technically, it concerns the failure of conglomerability,⁷ that is

$$P(A|N=n) > P(A) \text{ for all } n > 0.$$

EXAMPLE 14. (IMPROPER DISTRIBUTIONS) There are Bayesians who use improper distributions (that is, those that integrate to infinity) as a way of modeling “ignorance”; sometimes these priors are called “reference” or “objective.” The simplest case of these arise as apparent results of limits of proper priors. Suppose for example that the datum x is normally distributed with mean W and precision (inverse variance) 1. Suppose also that W has a normal prior with mean 0 and precision r . Then it is well known that the posterior distribution of W given x is again normal with mean the precision-weighted average of x and 0, and precision equal to $r + 1$. Now imagine that r approaches 0, so the variance of the prior goes to infinity. Then the posterior of W given x approaches a normal distribution with mean x and precision 1. This is the same answer that would have been obtained from using an improper uniform distribution on W , and doing a formal calculation of the posterior. The hope is that because the improper distribution is the limit of proper posteriors, nothing “bad” happens in the limit.

Now suppose that the precision R of the normal observation is not known, but that the conjugate prior is imposed, that is, W given $R = r$ is normal with mean 0 and precision τr , and R has a $\Gamma(\alpha, \beta)$ distribution with $\alpha > 1$. The general form of the posterior, given $X = x$ is then⁸ as follows: the posterior on W given $R = r$ is normal as before,

⁶We explored this phenomenon in Kadane, Mark Schervish, and Seidenfeld, “Reasoning to a Foregone Conclusion,” *Journal of the American Statistical Association*, xci (1996): 1228–35.

⁷For more on this, see Bruno de Finetti, *Theory of Probability*, Volume 1 (New York: Wiley, 1974).

⁸See Morris Herman DeGroot, *Optimal Statistical Decisions* (New York: McGraw-Hill, 1970) pp. 167ff.

with mean the precision-weighted average of x and 0 with precisions r and τ respectively, and the posterior of R is $\Gamma(\alpha + 1/2, \beta + \tau x^2 / [2(\tau + 1)])$. As τ approaches zero, the conditional posterior distribution of W given $R = r$ is as before, normal with mean x and precision r , and the posterior of R is $\Gamma(\alpha + 1/2, \beta)$.

Now, we set up a decision problem in which one would pay to avoid learning X . Let $\sigma^2 = 1/R$, the usual variance parameter. In the notation used earlier, with $\omega = (W, R)$, let $U(\omega, \psi_b, \delta) = \sigma^2$. Note that the utility depends only on the value of ω , not on any decisions. In this case, before any data are observed,

$$(17) \quad E[U(\omega, \psi_0, \delta)] = E(\sigma^2) = \beta / (\alpha - 1),$$

and after $X = x$ is observed say at time t_1 ,

$$E[U(\omega, \psi_{t_1}, \delta)] = (\sigma^2 | X = x) = \frac{\beta}{\alpha - 1/2},$$

the same for all x , which is strictly smaller than (17).

A similar result holds no matter how many observations one contemplates. For example, suppose that $X = (X_1, \dots, X_n)$ is a vector of n observations that are conditionally independent and identically distributed given (W, R) with the same distribution described above. In this case, (1) is the same as above, namely $\beta / (\alpha - 1)$. With n observations, the posterior distribution of R derived from the improper prior is $\Gamma(\alpha + n/2, \beta + \sum_{i=1}^n (x_i - \bar{x})^2 / 2)$, where $\bar{x} = \sum_{i=1}^n x_i / n$. Using the same utility $U(\omega, \psi_b, \delta) = \sigma^2$, we see that (2) equals

$$(18) \quad \frac{\beta + \frac{1}{2} \sum_{i=1}^n (x_i - \bar{x})^2}{\alpha + \frac{n}{2} - 1}.$$

To calculate (3), recall that the distribution of X is defined by first conditioning on (W, R) and then integrating with respect to the distribution of (W, R) . The integral of (18) with respect to the conditional distribution of X given $(W, R) = (w, r)$ is easily computed as

$$\frac{\beta + \frac{(n-1)\sigma^2}{2}}{\alpha + \frac{n}{2} - 1}.$$

Integrating this with respect to the marginal distribution of (W, R) yields

$$\frac{\beta + \frac{1}{2} \frac{(n-1)\beta}{\alpha-1}}{\alpha + \frac{n}{2} - 1},$$

which is then (3) in this example. A little algebra shows that (3) is still smaller than (1). Indeed, the difference is $-\beta/[2(\alpha - 1)(\alpha - 1 + n/2)]$. One would pay, say, half of this amount to avoid seeing the data.

The issue in Example 14 is that the posterior mean of σ^2 differs from the prior mean by an amount that has negative prior mean. That is, even before collecting the data, we expect the posterior mean to be smaller than the prior mean. Just as in Example 13, it is fair to ask what your prior mean of σ^2 really is, what you think now or the mean of what you know you would think if you were to observe X . Just as before, you can be confronted with the prospect of paying not to see the data X , because it would change your mean of σ^2 in a way unfavorable to you. Hence even in the apparently simple case of a sample of observations from a univariate normal distribution, with conjugate priors and flattening on the mean, the same issue arises of potentially paying not to see the data.

V. GENERALIZATIONS OF BAYESIAN DECISION THEORY INVOLVING SETS OF PROBABILITIES

There is recent interest in decision theories involving sets of probabilities. The intuitive idea behind them is to relax the requirement of knowing one's opinion about $p(x, \theta)$, but doing so in a way different from the "reference" - "objective" - "non-informative" prior school. We study three of the leading such theories here. In all of this discussion, the joint distributions will be countably additive and all data will be cost-free.

Of the three decision rules we discuss, perhaps the most familiar one is Γ -Maximin.⁹ This rule requires that the decision maker ranks a gamble by its lower expected value, taken with respect to a closed, convex set of probabilities, \mathcal{P} , and then to choose an option from \mathcal{A} whose lower expected value is maximum.¹⁰ The Γ -Maximin decision rule creates a preference ranking of options independent of the alternatives available in A : it is context independent in that sense. But

⁹ When outcomes are cast in terms of a (statistical) loss function, the rule is then Γ -Minimax: rank options by their maximum expected risk and choose an option whose maximum expected risk is minimum.

¹⁰ This decision rule was given a representation in terms of a binary preference relation over Anscombe-Aumann horse lotteries by Itzhak Gilboa and David Schmeidler, "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, xviii (1989): 141-53, has been discussed by, for example, James O. Berger, *Statistical Decision Theory and Bayesian Analysis* (New York: Springer, 1985, 2nd ed.), section 4.76. And recently by Peter D. Grunwald and A. Philip Dawid, "Game Theory, Maximum Entropy, Minimum Discrepancy, and Robust Bayesian Decision Theory" (Technical Report 223, University College, 2002) who defend it as a form of Robust Bayesian decision theory.

Γ -Maximin corresponds to a preference ranking that fails the so-called (von Neumann-Morgenstern's) "Independence" or (Savage's) "Sure-thing" postulate of SEU (Subjective Expected Utility) theory.

The second decision rule that we consider, called *E*-admissibility ('*E*' for "expectation"), was formulated by Isaac Levi.¹¹ *E*-admissibility constrains the decision maker's admissible choices to those gambles in A that are Bayes for at least one probability $P \in \mathcal{P}$. That is, given a choice set A , the gamble f is *E*-admissible on the condition that, for at least one $P \in \mathcal{P}$, f maximizes subjective expected utility with respect to the options in \mathcal{A} .¹² Savage¹³ defends a precursor to this decision rule in connection with cooperative group decision making. *E*-admissibility does not support an ordering of options, real-valued or otherwise, so that it is inappropriate to characterize *E*-admissibility by a ranking of gambles independent of the set \mathcal{A} of feasible options. However, the distinction between options that are and are not *E*-admissible does support the "Independence" postulate. For example, if neither option f nor g is *E*-admissible in a given decision problem \mathcal{A} , then the convex combination, the mixed option $h = \alpha f \oplus (1 - \alpha)g$ ($0 \leq \alpha \leq 1$) likewise is *E*-inadmissible when added to \mathcal{A} . This is evident from the basic SEU property: the expected utility of a convex combination of two gambles is the corresponding weighted average of their separate expected utilities; hence, for a given $P \in \mathcal{P}$ the expected utility of the mixture of two gambles is bounded above by the maximum of the two expected utilities. The assumption that neither of two gambles is *E*-admissible entails that their mixture has *P*-expected utility less than some *E*-admissible option in \mathcal{A} .

The third decision rule we consider is called *Maximality* by Peter Walley,¹⁴ who appears to endorse it (*ibid.*, p. 166). If a gamble $f \in \mathcal{A}$ is such that there is no other gamble $g \in \mathcal{A}$ strictly preferred to f , then f is admissible by *Maximality*. That is, f is a *Maximal* choice from \mathcal{A} provided that there is no other element $g \in \mathcal{A}$ that, for each $P \in \mathcal{P}$, carries greater expected utility than f does. *Maximality* (under differ-

¹¹ Levi, "On Indeterminate Probabilities," this JOURNAL, LXXI, 13 (July 18, 1974): 391–418; and Levi, *The Enterprise of Knowledge* (Cambridge: MIT, 1980).

¹² Levi's decision theory is lexicographic, in which the first consideration is *E*-admissibility, followed by other considerations, for example, what he calls a Security Index. Here, we attend solely to *E*-admissibility.

¹³ Savage, *op. cit.*, section 7.2. Savage's analysis of the decision problem depicted by his Figure 1, p. 123, and his rejection of option b , p. 124 is the key point.

¹⁴ Walley, *Statistical Reasoning with Imprecise Probabilities* (New York: Chapman and Hall, 1990). There is, for our discussion here, a minor difference with Walley's formulation of maximality involving null-events. Walley's notion of maximality requires, also, that an admissible gamble be classically admissible, that is, not weakly dominated with respect to state-payoffs.




Figure 1: Expected utilities for three acts in Example 15. The thicker line indicates the surface of Bayes solutions.

ent names) has been studied, for example, in Hans G. Herzberger,¹⁵ Levi,¹⁶ Seidenfeld,¹⁷ and Amartya Sen.¹⁸ Evidently, the E -admissible gambles in a decision problem are a subset of the Maximally admissible ones.

The three rules have different sets of admissible options. Here is a heuristic illustration of that difference.

EXAMPLE 15. Consider a binary-state decision problem, $\Omega = \{\omega_1, \omega_2\}$, with three feasible options. Option f yields an outcome worth 1 utile if state ω_1 obtains and an outcome worth 0 utiles if ω_2 obtains. Option g is the mirror image of f and yields an outcome worth 1 utile if ω_2 obtains and an outcome worth 0 utiles if ω_1 obtains. Option h is constant in value, yielding an outcome worth 0.4 utiles regardless whether ω_1 or ω_2 obtains. Figure 1 graphs the expected utilities for these three acts. Let $\mathcal{P} = \{P : 0.25 \leq P(\omega_1) \leq 0.75\}$. The surface of Bayes solutions is highlighted. The expected utility for options f and g each

¹⁵ Herzberger, "Ordinal Preferences and Rational Choice," *Econometrica*, xli (1973): 187–234.

¹⁶ Levi, "On Indeterminate Probabilities," and "Imprecise and Indeterminate Probabilities," in *Proceedings of the 1st International Symposium on Imprecise Probabilities and Their Applications* (1999).

¹⁷ Seidenfeld, "Comment on 'Probability, Evidence and Judgement'," (by A.P. Dempster), in J.M. Bernardo, ed., *Bayesian Statistics 2: Proceedings of the Second Valencia International Meeting, 1983* (New York: North Holland, 1985), pp 127–29.

¹⁸ Sen, "Social Choice Theory: A Re-Examination," *Econometrica*, xlv (1977): 53–89.

has the interval of values $[0.25, 0.75]$, whereas h of course has constant expected utility of 0.4. From the choice set of these three options $\mathcal{A} = \{f, g, h\}$ the Γ -Maximin decision rule determines that h is (uniquely) best, assigning it a value of 0.4, whereas f and g each has a Γ -Maximin value of 0.25. By contrast, under E -admissibility, only the option h is E -inadmissible from the trio. Either of f or g is E -admissible. And, as no option is strictly preferred to any other by expectations with respect to \mathcal{P} , all three gambles are admissible under *Maximality*.

We link this observation to the debate about the value of new information by considering a sequential decision problem in which the decision maker has the opportunity to postpone a terminal decision in order to learn the outcome of a *mixing* variable, a variable used to convexify the option space. Let the mixing variable α equal 1 or 0 as a *fair* coin lands Heads up or Tails up on a toss, so that $P(\alpha = 1) = P(\alpha = 0) = .5$. Assume, also, that α is independent of the states, Ω , over which the *pure* options are defined, so that each $P \in \mathcal{P}$, $P(\alpha, \omega) = .5P(\omega)$.

As a modification of Example 15, consider the mixed options m , and n , defined as follows

$$m = \alpha f \oplus (1 - \alpha)g$$

$$n = \alpha g \oplus (1 - \alpha)f.$$

Thus, m is the mixed act that uses the fair coin to bet on ω_1 if Heads (H) and to bet on ω_2 if Tails (T). Likewise, n is the dual mixed act that uses the same fair coin to bet on ω_2 if Heads and to bet on ω_1 if Tails. Note that m and n each carry a constant expected utility, .50. That is, these are options of determinate *risk*, despite the fact that uncertainty is with respect to the convex set of probabilities \mathcal{P} . Last, let the option *Status quo* denote no change in expected utility, with constant value 0.

The next two examples, sequential decision problems 1 and 2, are cast in extensive form. The three decision rules we discuss with these examples, each a generalization of Bayesian decision making when uncertainty is represented by a set of probability distributions rather than a single probability distribution, yields different results when extensive form decisions are changed to their normal form. Thus, in what follows we have a different situation than we saw in Examples 4 and 5. In those examples uncertainty is represented by a single probability distribution. Then, each of these three rules yields the same Bayesian analysis either in extensive form or in normal form.¹⁹

¹⁹ See Seidenfeld, "When Normal and Extensive Form Decisions Differ" (*op. cit.*).

Sequential Decision Problem 1: The sequential decision problem, depicted in Figure 2, offers the decision maker one of three terminal options: $m-.40$, $n-.40$, and *Status quo*. The sequential option offers a .05 bonus to the decision maker to postpone this choice in order to learn the outcome of the mixing variable, α , and then to choose among these same three terminal options. Of course, given $\alpha = 1$, act m is equivalent to act f , and act n is equivalent to act g . Likewise, given $\alpha = 0$, act m is equivalent to act g , and act n is equivalent to act f .

Let us reason through this decision problem using Γ -Maximin as the decision rule. At either of the terminal choice nodes that might be reached under the sequential option, that is, after observing the coin flip, the decision problem reduces to a case formally analogous to the nonsequential problem of Example 15. That is, the constant option (worth .05) is uniquely Γ -Maximin admissible at either of these choice points. Using this as the value assigned at the initial choice point to the sequential option, that is, by backward induction, we find that Γ -Maximin rejects the sequential option in favor of either $m-.40$ or $n-.40$, which are indifferent and worth .10 each. Hence, Γ -Maximin mandates a negative value for the new information. The Γ -Maximin

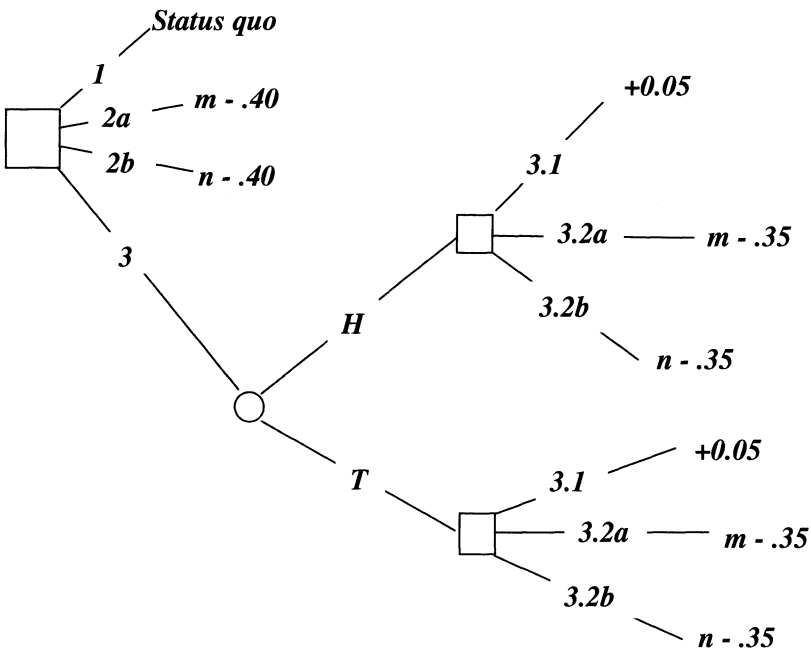


Figure 2: Sequential Decision Problem 1.

decision maker is advised, in effect, to pay negative tuition: pay not to learn!

By contrast, using E -admissibility as the decision rule, we find that the constant option is uniquely inadmissible at each of the choice points that might be reached under the sequential option. Since, for example, $m-.40$ is inadmissible when $m-.35$ is available, by backward induction, the decision maker who uses E -admissibility will not assign a negative value to the new information available in this sequential decision problem.

We are unsure what conclusion to draw about what *Maximality* recommends in this case. At the two terminal choice points that might be reached under the sequential option, all three choices are Maximal. Then, it is permissible for the decision maker who decides by Maximality to choose the constant (worth .05) if the sequential option is taken at the initial choice node. But .05 is an inadmissible option at the initial choice node, since either $m-.40$ or $n-.40$ is strictly preferred to .05. Hence, it appears that in this sequential decision problem, Maximality does not require a decision maker to assign a non-negative value to potential cost-free information.

We conclude this brief discussion of the value of new information by noting that even E -admissibility requires only that an admissible solution be Bayes in a *local*, but not in a *global* sense. That is, in a sequential decision problem, where there are hypothetical decision nodes to consider that are not the initial choice node, E -admissibility does *not* require that there be one, common Bayes model for all combinations of E -admissible choices across these hypothetical decision nodes. This is illustrated in Figure 3 by the following sequential decision problem, which is a variant of the previous one.

Sequential Decision Problem 2: Here, at each of the two terminal choice nodes that might be reached under the sequential option, both options are E -admissible. Then, it is permissible for the decision maker who decides by E -admissibility to choose the constant (worth .05) if the sequential option is taken at the initial choice node. But .05 is an inadmissible option at the initial choice node, since $m-.40$ is strictly preferred to .05. Hence, it appears that in this sequential decision problem, E -admissibility does not require a decision maker to assign a non-negative value to potential cost-free information. Of course, since, E -admissible options are always a subset of the options permissible by Maximality, to the extent that this phenomenon is a problem with E -admissibility as the decision rule, it is only more so of a problem for decision makers using Maximality.

The examples of this section all involve countably additive probabilities and decisions about cost-free information. Nevertheless, the

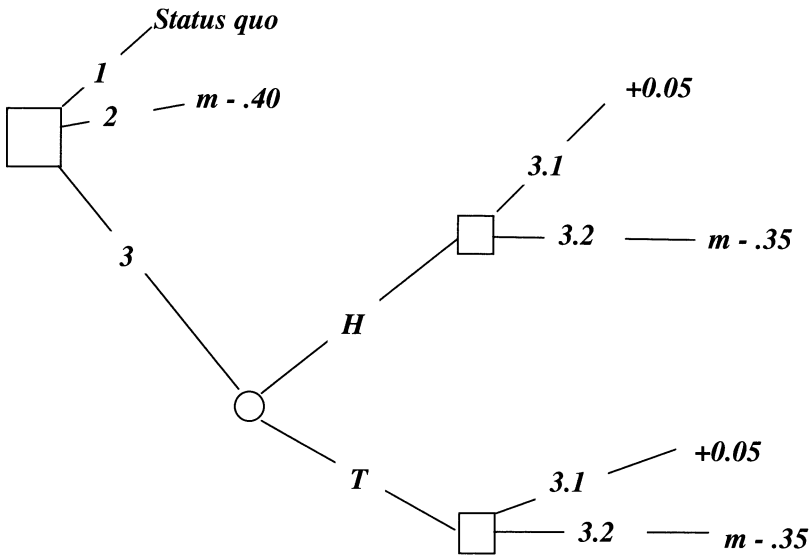


Figure 3: Sequential Decision Problem 2.

theories of E -admissibility, maximality, and Γ -maximinity either allow or mandate assigning negative value to this cost-free information. We are unable to see how these theories can justify assigning negative value to the cost-free information in these cases.

VI. IS IGNORANCE BLISS?

What would a world be like in which rational actors would pay not to receive sample information of certain kinds? It would take some effort to discover, for each candidate for such a service, just what information he or she would pay not to receive. Suppose there were two companies that would offer this service, one for \$0.50 and one for \$0.25. To be effective, the consumer has to deal with both of them, and in fact with any other company also offering not to tell. Indeed, the provider of this service could be most persuasive in guaranteeing to the customer that the data will not be provided, by not possessing it themselves. This offers a new world of entrepreneurial activity of no apparent social value whatsoever.

But are there instances in which a reasonable person would pay not to have certain information? Suppose you were offered the opportunity to learn the date and circumstances of your own death. Would you pay not to have this information? For many of us, the answer is "yes." Such knowledge would change much about what it means to be human, and would make the experience of life rather

different, in a profound way, from all of those who have gone before us. To restructure the way we think about life, and its uncertain time of end, would be profoundly unsettling and costly, and hence it would be reasonable to pay not to have to undergo that process. On the other hand, absent professional hitmen, none of us are likely to receive such an offer.

The issue is more real, however, for those who have to decide whether to be tested for an incurable genetic disease such as Huntington's Disease. This disease generally strikes in people of middle ages, with neurological and psychological symptoms, ultimately resulting in a period of helplessness followed by an early death. There are limited treatments that can help to mitigate some of the symptoms, but no cure as yet. Those who are genetic candidates for the disease have a 50% probability of having the gene. Knowing that one is a candidate, is it reasonable not to have the test, if it were free? One reason not to have the test is the response of others to the results: "people lose insurance and jobs."²⁰ But there are deeper reasons having to do with one's ability to cope and enjoy life that influence this decision. We can not label either decision irrational, given all the implications.

How can we defend the idea of refusing to pay not to see information and yet endorse a decision not to be tested for Huntington's disease by someone who has a parent who had the disease, and hence has a 50% probability of carrying the gene? The insurance issue is serious, since most life and health insurers ask about known pre-existing conditions. Hence someone who knows they have the gene could easily believe they would be denied coverage on that account. However, one could imagine a strategy of buying insurance first, and then being tested. Hence, although we take insurance to be an important practical problem, we do not take it as definitive. The real reasons, we guess, lie deeper.

EXAMPLE 16. (INHERITED DISEASE) There is a certain peace of mind associated with knowing that, aside from catastrophe, one has the ability to make plans and choices for a relatively long future. There is also value to an individual of having the prospect of making choices and enjoying results in the future. By having a short horizon placed on our ability to make choices and use our assets, we reduce the overall level of utility. In the terminology of Sen,²¹ the knowledge

²⁰ Feldman, "One Family's Story: A Quandry over Genetic Testing," *Houston Chronicle* (2003), <http://www.johnworldpeace.com/e030112e.htm>.

²¹ Sen, "Markets and Freedoms: Achievements and Limitations of the Market Mechanism in Promoting Individual Freedoms," *Oxford Economic Papers*, XLV (1993): 519–41, on p. 522.

that one has the disease reduces one's opportunity-freedom. In symbols, let ψ be the knowledge base sequence assuming that I will not learn the results of the test. Let ψ' be the knowledge base sequence assuming that I will learn (at time t_1) whether or not I have the disease. Let $X = 1$ if I have the disease and $X = 0$ if not. Let $\Omega = \Theta \times \{0,1\}$, where the second coordinate gives the value of X . Let $\theta \in \Theta$, and let δ be a decision that I could make at time $t > t_1$ regardless of the value of X . It might be reasonable that

$$(19) \quad U((\theta,1),\psi',\delta) < U((\theta,1),\psi,\delta) < U((\theta,0),\psi,\delta) < U((\theta,0),\psi',\delta)$$

for all such θ and δ . That is, knowing that I have the disease makes everything worth less than merely supposing that I have the disease, which in turn makes everything worth less than supposing that I do not have the disease, and knowing that I do not have the disease makes everything worth the most. If I am not going to learn the results of the test, the expected utility now of choosing δ at time t is

$$(20)$$

$$E[U((\theta,X)\psi,\delta)] = 0.5E[U((\theta,0),\psi,\delta) | X = 0] + 0.5E[U((\theta,1),\psi,\delta) | X = 1].$$

On the other hand, if I am going to learn X in the future, the expected utility now of choosing δ at time t

$$(21) \quad 0.5E[U((\theta,0),\psi',\delta) | X = 0] + 0.5E[U((\theta,1),\psi',\delta) | X = 1].$$

If the first inequality in (19) is satisfied by a much wider amount than the third inequality and is so satisfied uniformly in θ and δ , then the expected utility in (21) will be smaller than the one in (20) for all δ , hence one would prefer not to learn X .

In this connection, it is interesting to note that this is the kind of information the poet Thomas Gray (1747) has in mind. The quote at the start of this paper comes from a poem, "Ode on a Distant Prospect of Eton College." The first nine stanzas are about the careless optimism of the students at Eton, and the troubles they are likely to face in life. The concluding tenth stanza reads:

To each his sufferings: all are men
 Condemn'd alike to groan?
 The tender for another's pain,
 Th' unfeeling for his own.
 Yet, ah! why should they know their fate,
 Since sorrow never comes too late
 And happiness too swiftly flies?
 Thought would destroy their Paradise.
 No more; where ignorance is bliss
 'Tis folly to be wise.

So the question remains of whether it is reasonable to impose the requirement on a theory of rational decision making that it not require or permit paying not to see cost-free data. If it is, the only such theory known to us is Bayesian decision theory with a single countably-additive proper prior. Each of the weakenings of this theory—allowing finitely additive priors, or improper ones, or allowing sets of probabilities with any one of the three decision rules studied here violates this principle.

JOSEPH B. KADANE

MARK SCHERVISH

TEDDY SEIDENFELD

Carnegie Mellon University

APPENDIX

A.1 Conditioning and Learning. In this appendix we discuss two related matters for understanding how we use conditional probability in connection with changes in evidence:

- (i) what conditional probability means in connection with changing bodies of evidence, and
- (ii) what it is that is *given* in a conditional probability.

(i) The first issue about conditional probability arises both with randomized (so-called “mixed”) options and in sequential decisions as we present those in section v. We use coin flips that are stipulated to be independent of the other events of interest. That is, for each joint probability distribution we consider, the conditional probability for the event of interest given the coin flip equals its unconditional probability. In each example, we argue that, with respect to each probability distribution, given a result of the coin flip, the available options are valued the same as they are valued marginally under that distribution. Such an argument is an implicit use of conditional probability. There are at least three distinct interpretations of conditional probability, and this appendix attempts to identify which interpretation we have used in this paper.

Let A and B be events in the algebra over which probability is defined, and let I_A and I_B be their indicator functions.²² The conditional probability $P(A|B)$, as it relates to degrees of belief

²² What we say here about conditional probability is not restricted to indicator random variables, which we use for simplicity of exposition only.

and decision making, has been interpreted in at least the following three ways:

- (1) Conditional probability and *called-off preferences*. In the theories of both DeFinetti and Savage, where probability is grounded on unconditional preferences, conditional probability is reduced to called-off preferences.
 - (a) Specifically, in DeFinetti, $P(A|B)$ is that prevision p the decision maker offers in order to make the following bet “fair” for all values of $\alpha : I_B\alpha(I_A - p)$. The bet is “called-off” because there is neither a loss nor a gain if event B fails to occur. DeFinetti shows that, in order for your called-off previsions to be coherent, that is, immune to a “Book,” they must be the values of a (finitely additive) conditional probability function.
 - (b) In Savage, conditional probability is identified with the decision maker’s preferences over pairs of acts that are modified to agree with each other on states that comprise B^c , when B fails. Thus, these pairs of acts conform to the requirements of being “called off” in the sense that, for states comprising B^c , the decision maker of Savage’s theory receives the same regardless which “called-off” act is chosen.

It is important to note that in each of these two theories there is no relation between conditional probability and updating beliefs by new evidence. That is, neither of these theories goes beyond the statics of preference at a time. Neither of these theories offers a dynamic interpretation of conditional probability that makes it into a learning rule.

- (2) Conditional probability and *temporal updating*. It is a commonplace interpretation of Bayesian theory to offer the (current) conditional probability $P(A|B)$ as a “posterior” probability for what the investigator’s degree of belief will be in A on the condition that he/she learns that B obtains. Some attempts have been made to extend deFinetti’s static “Book” argument to cover such cases, for example, to make conditional probability into a dynamical, learning rule. See, for example, Levi,²³ for sound reasoning as to why this is an overreaching of the “Book” argument. Levi refers to the dynamic interpretation as *temporal credal conditionalization*.
- (3) Conditional probability and *hypothetical reasoning*. Yet a third sense of conditional probability, the sense used in this paper, is to understand $P(A|B)$ as the decision maker’s current, hypothetical degrees of belief in A on the condition that he/she were to ac-

²³ Levi, “The Demons of Decision,” *The Monist*, LXX (1987): 193–211.

cept B as true. This sense of conditional probability relates to the agent's current hypothetical preference were he/she to accept B as true. In this sense, the agent reasons from a hypothetically augmented body of knowledge that includes B as certain. It differs from the first sense of $P(\cdot|B)$ in that the space of serious possibility is changed by hypothesizing B as a certainty, equating B^c with the impossible event. It differs from the second sense of conditional probability in that the decision maker does not predict how he will update his/her degrees of belief at some later time, when B might be accepted as true. (In Levi, this is referred to as *confirmational conditionalization*.)

In short, these three senses of conditional probability $P(\cdot|B)$ differ in their consequences for decision making by addressing, respectively:

- (1) current, unconditional called-off preferences, called off if B^c obtains;
- (2) future, unconditional preferences, upon learning that B obtains;
- (3) current, hypothetical preferences obtained by positing that B is true.

Throughout this paper, we use conditional probability to relate to decision making in the third sense, and we use the formalism of sequential decision making to highlight the difference between the first and third senses of conditional probability. We do not use (or advocate) the second sense of conditional probability.

(ii) The second matter involving our use of conditional probability in this paper is made salient in Example 12, the "Taxi Driver" problem. There is an important distinction between conditioning on who-done-it and adding the information to one's knowledge base. Think of D as a random variable taking the value i if the i th character in the mystery committed the crime. Conditioning on $D(\omega) = i$ for $i = 1, \dots, \ell$ does not affect my expected utility for any of the decisions that I contemplate choosing in this decision problem. On the other hand, augmenting my current knowledge base either with the information that I know $D(\omega) = i$ or with the information that I will know $D(\omega) = i$ before getting to watch the play, has a pronounced effect on the value of going to the play. That knowledge affects my choices in the same, negative way regardless of the value of D that I (will) learn.

Thus, in our use of conditional probability, we need to distinguish conditioning from learning, and in two ways: From the first issue (i), we use conditional probability to model hypothetical changes in knowledge, rather than requiring conditional probability to be a learning rule. This takes us beyond merely the use of conditional probability to model called-off wagers. But it leaves us short of mandating a single dynamics by which rational agents learn.

From the second issue (ii), we must attend to the event on which we condition. We recognize that it is *not* generally the same to condition on E occurring, as it is to condition on the more inclusive event that you learned that E occurred by some particular method. The distinction between conditioning and learning has important consequences in the Taxi-Driver Example 12 and in the Inherited Disease Example 16. We believe that the distinction does not matter in the other examples discussed in this paper.