Environmental Sustainability as a Culturally Invariant Value

John H Dreher, Associate Professor of Philosophy, University of Southern California

Abstract

Various disciplines, including philosophy, have sought culturally invariant values because they transcend cultural differences and thereby reduce the probability of discord, tension and even war. Environmental sustainability, especially of the climate, appears to be culturally invariant and necessary for human survival, even apart from its potentially stabilizing social, political and cultural effects. However, policies that promote sustainability can be expected to involve sacrifices, and as a result it has proved difficult to reach anything like consensus about rational responses to environmental challenges.

This paper explores *epistemological* principles that may be helpful in justifying, promoting and implementing public policies that justify local sacrifices for the global good of a sustainable, healthy environment. The principal contribution this paper hopes to make is to explain why it is that assigning unreasonably high prior probability measures to mere suppositions about the environment tends to undermine constructive discussion about sustainability, especially in the popular media.

Environmental Sustainability as a Transformative Value

Everyone is familiar with the difference between taking dramatic, immediate measures to reach an important goal and transforming habits or even a way of life to reach long term objectives. It is one thing to lose weight before an important social event or a physical examination; it is quite another to adopt a life style that stabilizes weight at a lower, healthier level. It is one thing to mobilize an army to confront an immediate threat; it is quite another to maintain a standing army that is sufficient to deter threats in the hope avoiding war altogether.

Call a value that is meant to replace an established value a **transformative value. Transformative measures** are the changes transformative values require and justify. Long term positive changes, for example, in the production of electricity or in commutes from the exurbs, almost always involve modifying or replacing established social values in ways that better serve our purposes.[1] Short term measures do not necessarily involve a change in entrenched values, but they sometimes do. For example, increasing recycling by improving waste disposal practices involves small changes, but in the aggregate make a significant difference in the way in which we live.

In a previous paper I suggested that a good environmental goal for now would be to keep the global environment from further degradation. This is not to say that stricter short term measures to protect the environment would not be more desirable; it is merely to suggest that what we need now is a policy that can be persuasively defended on a rational basis and that offers hope of implementation by avoiding immediate, widely dreaded reductions in living standards. (XXXXXX, 2011, p. 10) Whatever

[1] This is a point developed in (Bonner, 2011, pp. 1 - 14), where he draws our attention to reconfigurations of 'political and economic arrangements' and even 'basic social relations' that transformative environmental measures involve.

merit this proposal may have, it is only a beginning because there are substantial pressures deeply embedded in the world economy that will make it more and more difficult to keep the environment from further degradation without changing the ways in which we live. First, there are obvious demographic pressures. The growth of the human population can be predicted, and it is obvious that a growing human population in the decades ahead will require increased production of material goods just to maintain the current standard of living. Secondly, many countries with huge populations are reasonably looking to improve their living standards, which in turn will require even more resources. Thirdly, it is reasonable to suppose that the extraction of resources will itself consume additional resources and create new environmental risks. For this reason, new methods of recovering natural gas and deep water drilling for oil are already controversial.[2]

On the positive side, it appears that technological improvements in both the production and distribution of natural resources offer hope of keeping up with the demand-curve for consumer goods. In addition, it is reasonable to think that industrially advanced societies might reconsider just how much consumption is required for a good life. If so, reducing pollution by industrial gases and general waste to levels that are consistent with the long-term development of human culture will require changes in the way in which we live, including those imposed by technological advances and necessary conservation measures. Guaranteeing the sustainability of the environment is different from taking short term measures, however dramatic they may be, that are intended to save the environment from further decline at least in the near term. All this shows that environmental sustainability is itself a **transformative** value; it requires changes in the way in which we produce and consume the material resources necessary for a good and rational life in the global context of a growing human population.

Rational Procedures for Decisions Taken under Conditions of Uncertainty

Although the measures necessary to achieving sustainability are obvious when described at a high level of abstraction, e.g., technological advances and conservation, matters become immensely more complicated as soon as we try to be specific about what to do. Transformative measures are meant to promote environmental sustainability through technological advances and reductions in consumption, including new commitments to recycle waste, to generate electricity by 'clean' methods, and to reduce fuel consumption by vehicles as well as the emission of industrial gases. Transformative measures cannot succeed unless they are embraced by willing, cooperative populations and their governments.[3] Unnecessary constraints on consumption or unrealistic expectations of technology are sure to undermine the proposals meant to guarantee environmental sustainability. The heart of the problem is that there is a vast array of risk factors, and it is difficult to assess their environmental costs, although some costs are obviously intolerable, like Fukushima. Other risk factors are less threatening at least in the short run; for example, plastic bags and bottles. Surely they are aesthetically repugnant as they accumulate on beaches and parks. More to the point, however, are the environmental costs of producing and recycling the plastic itself. Although we probably could live indefinitely with *some* relatively minor assaults to sustainability like plastic bags and bottles, we cannot live indefinitely with them *all*, however minor they may be.

[2] These themes are widely discussed in the literature. See, for example, (Tecle, 2010, pp. 419 – 32).
[3] Persuading populations of the need for change, however, may need to focus at least in part upon local issues. This point is developed in (Justice, Creek, and Buckman, 2011, pp. 2, 11).

Disaster may well be brought about by an extremely large number of bad choices, none of which make much difference in the short run. In the long run it may be that all the minor assaults on the environment are more threatening than catastrophes if only because collective action is more likely to be taken to avoid immediate and unquestionable disasters rather than to ameliorate the cumulative effects of relatively minor assaults on the environment. This paper focuses on strategies for dealing with the accumulation of minor risk factors, but its conclusions also apply to potentially catastrophic events that require immediate attention.

Assessing the Dangers and Costs of Relatively Minor Risk Factors to Sustainability

It is tempting to think that that we can best proceed by ***prohibiting*** behaviors that pose even small risks to the environment. Yet, the difficulties attending this strategy are obvious and important when it comes to justifying the sacrifices that environmental sustainability may require. For example, the addition of small amounts of smoke from wood burning fires on rainy days may not be significant pollutants because they are washed out of the air almost immediately. On the other hand, fires like the 2010 fire in the Angeles National Forest, which burned tens of thousands of acres and created a blanket of dense acrid smoke for days, may be significant health hazards and adversely affect the weather (and perhaps even the climate) for a time. Yet, some areas in California now discourage or prohibit the use of wood burning fireplaces, although virtually nothing has been done to reduce the level of combustible material in its vast forests. The point is that transformative measures to sustain the environment need to be implemented by a cooperative population. Requiring annoying sacrifices of minimal value while ignoring more significant threats may be a self-defeating strategy.

Minor risk factors tend to involve dangers that are incommensurable and difficult to assess. For example, the United States instituted incentives, so-called 'cash for clunkers,' designed to induce people to trade older, less efficient automobiles for new more efficient models. The older vehicles were destroyed. Yet the productive lives of the older vehicles had not been exhausted (indeed, the point of the program was to get older vehicles still in use off the road; therefore, only *working* vehicles were accepted for exchange). However well intentioned, it may not be that there is any gain at all from an environmental standpoint in replacing an older, working vehicle with a new vehicle if the older one can be made to run efficiently.

Policies like 'cash for clunkers' involve political considerations, and it is hardly surprising that the underlying logic supporting those policies is not incisive. In the case of 'cash for clunkers' it may very well be that the entire policy was proposed primarily to 'stimulate' the economy by supporting the collapsing United States automobile industry. In any case, the two examples above illustrate the need for rational principles for assessing the costs and benefits of transformative proposals and the need to promote those proposals honestly and effectively through rational discourse within established political processes. But what form would rational discourse take? Obviously we are looking for something akin to a 'cost/benefit analysis,' that is, for a decision procedure to determine the likelihood of a probable outcome with a certain value. The expected utility of the decision would then be determined by the product of the **value** of the outcome and the **probability** that the outcome would be achieved. The proposed policy could then be compared to the costs of alternatives, as well as the cost of doing nothing at all. Costs would be calculated by determining product of the probability of the burden that the policy imposes and its (negative) value. To illustrate the underlying logic, the expected utility of correctly

picking a card from a deck of 52 is one fifty-second of the value of the bet. Thus, the bet will be fair if a correct pick is rewarded by a payout of the product of 52 and the amount bet. It sounds simple, but it becomes complicated as soon as we consider real-world applications of the theory.

Calculating Conditional Probabilities on the Bayesian Model

The leading method for evaluating hypotheses (as opposed to calculating chances) is based upon the Bayesian definition of conditional probability, which is that support given to a hypothesis H by evidence E is equal to the probability of the hypothesis on the evidence divided by the probability of the evidence alone; which is expressed as :

$P_E(H) = P(H\&E)/P(E)$, assuming that P(H & E) and (PE) are defined and the $P(E) > 0$.[4]

According to Bayes' definition of conditional probability, the probability of an outcome needs to be revised as each new piece of evidence becomes available. The textbook cases generally used to illustrate how to revise probability assessments on new information are appealing because they are simple and the theoretical rationale underlying the basic idea is clear and compelling. (Joyce, 2008, §1)

Let us see how all this might work. Suppose that I am holding a coin produced in a certain country C. The coin is large, with beautiful images, and is composed primarily of gold alloyed with silver and copper. It turns out that there were real problems about counterfeiting in C with the type of coin that I have, say the (coveted) '5C.' Due to corruption at C's mints, some of the coins were hollowed out and a combination of base metals replaced the precious metal extracted so that the counterfeits weigh exactly as much as the genuine coins. There isn't a way to distinguish the counterfeit coins without marring the beautiful images, which would destroy the considerable numismatic value of the coins.

Suppose that I want to calculate the probability that any given '5C' is counterfeit and that I *know* that 1B (billion) such coins have been produced and that of the billion, it is has been *conclusively established* that 100M (million) are counterfeit. Let H be the hypothesis that any given 5C is counterfeit. The 'unconditional' probability of H is:

$P(H) = 100M/1B = .1.$

Now suppose that additional evidence becomes available about the coins, which is that of the 1B '5C' coins produced, 250M of them bear the mint mark 'T' (the '5C-T' coins), but that only 5M of them are counterfeit. $P_E(H)$ is the probability that any 5C is counterfeit, *given that it bears mint mark T*. Now, P(H & E) is the probability that any given 5C is *counterfeit and bears mint mark T*, which is

$P(H \& E) = 5M/1B = .005.$

P(E) is the probability that any given 5C bears mint mark T. That is:

$P(E) = 250M/1B = .25.$

[4] For an excellent introduction to the technical details of the Bayesian model. Bayes' definition of conditional probability and similar issues: see (Joyce, 2008, §1)

According to the Bayesian definition of conditional probability, the probability that any given 5C is counterfeit, *given that it bears mark* T, is

$P_E(H) = [(P(H \& E)/P(E] = [(.005/.25) = .02]$.

Intuitively, this makes good sense. Without the information E, the chance of a counterfeit 5C is .1. But with the information E, the situation changes because only 5M of the 5C-Ts are counterfeit.

Now, suppose that the coin in *my* hand is worn, and the mint mark is not discernible. I can tell that *my* coin is a 5C but not that it is a 5C-T. How does that affect my epistemological situation? In the first place, it has no effect upon my knowledge that *conditional probability* of *any given* 5C with mint mark T being counterfeit is .02. My problem is that I do not know that the 5C in *my* hand is a 5C-T. So, I am in a relatively weak epistemological position, because all I know is that the coin in my hand is a 5C, which justifies my belief based upon H, the unconditional probability that any 5C is counterfeit is .1, and hence my belief that the probability that the coin in *my* hand is counterfeit, which I know to be a 5C, is also .1.[5]

[5] To complicate the analysis of my knowledge state further, suppose that the coin in my hand is completely worn so that I cannot even tell whether or not it is a 5C. Now, *I am still justified in believing that the unconditional probability of any given 5C being counterfeit is .1*, and I am still j*ustified in believing that the conditional probability of any given 5C being counterfeit, given that it is a 5C-T, is .02*. But I am not in a position yet to form a conviction on the basis of that information about the coin in *my* hand, because I do not know that the coin in my hand is a 5C or, if it is, whether or not it is a 5C-T.

This raises an interesting question about the relation of the unconditional probability P(H) to the conditional probability $P_E(H)$. Suppose that we know that 8 billion coins were produced in C, i.e., that there are 8B Cs. Suppose further that 3B of the 8B are counterfeit. As before, suppose that 1B 5Cs were minted; of them only 100M are counterfeit; 250M of the 1B 5Cs were 5C-Ts and of the 250M 5C-Ts, only 5M are counterfeit. We know the unconditional probability H*(of any given C being counterfeit) is

$P(H^*) = 3B/8B = .375$

Now, E* is the probability of any given coin C being a 5C-T:

$P(E^*) = 250M/8B = .03125$

$P(H^* \& E^*)$ is the probability that any given coin C is a counterfeit 5C-T. Since we know there are a total of 5M counterfeit 5C-Ts, we know that

$P(H^* \& E^*) = 5M/8B = .000625$.

Hence, the conditional probability that any given C is counterfeit given that it is a 5C-T is:

$P_{E^*}(H^*) = P(H^* \& E^*)/P(E^*) = .000625/.03125 = .02$.

Here again, the Bayesian result is intuitively correct. Without the knowledge that a coin is a 5C-T, the probability of any given coin being counterfeit is .375. Since only 5M of 5C-Ts are counterfeit, the probability that a coin is counterfeit on the condition that it is a *5C-T* is .02. As in the previous case, which was *not noted above to keep the main discourse as simple as possible*, the *total number of coins produced* (1B in the first case; 8B in the second) actually drops out of consideration in the algebraic manipulation, as it should intuitively. That is because what is relevant to the assessment of the conditional probability ultimately depends upon the *proportion* of counterfeit 5C-Ts to the total 5Cs or in the second case, the counterfeit 5Cs to the total Cs. In each case the total appears in the conditional probability as a divisor of the numerator and the divisor of the denominator; so, the pertinent fraction is of the form ((X/Total)/(Y/Total)). Also notice, to return to the status of P(H) and P(H*), that the conditional probability $P_E(H) = [(P(H \& E)/P(E] = [(.005/.25) = .02]$ remains the same, as it should, whether or not we *compare the counterfeit 5C-Ts with the 5Cs or with the Cs.* But something does change in that case, which is that 5C-Ts look like a better investment when compared to the whole population of Cs rather than to the total population of 5Cs. Of course the fact that they *look like* a better investment has *no* bearing on whether they really *are* a better investment, which actually explains much advertising about investments.

In the second place, suppose that I have some evidence that my coin did bear mint mark 'T,' for example, that a relative told me that the coin now in my hand is the 5C-T that our grandfather (and his father before him) carried around for good luck, and that our grandfather's father was said by long since deceased relatives to have bought the coin personally from mint T and validated its authenticity from workers he knew were trustworthy. That information does not at all affect the calculation of $P_E(H)$. What it does affect is whether or not I am in a position to rationally believe that *my* coin actually falls in the class of 5C-Ts. We know P(E), which states that the probability that any given 5C is a 5C-T is .25. In the case just described, I do not know that the probability that *my* 5C is counterfeit is a mere .02, because I do not know that my 5C is a 5C-T. But perhaps I have some confidence on the basis of the testimony of my relative that *my* 5C really is a 5C-T. Say the 'subjective' probability (confidence) I assign to my relative's testimony is .7. It would be a mistake to think that I can now simply incorporate the subjective probability .7 in some way or other in the calculation of the conditional probability $P_E(H) = [(P(H \& E)/P(E]$. Whether or not *my* 5C is a 5C-T has no bearing whatever on the proportion of 5Cs that are 5C-Ts since it has no bearing on the total number of 5Cs that bear mint mark T. (250 M 5C-Ts were produced one way or the other. Who possesses them has nothing to do with how many were actually produced.)[6] How *confident* I am that *my* coin is a 5C-T does not affect the conditional probability that *any given* 5C is counterfeit given that it is a 5C-T, which is the probability we have already determined to be .02 on the basis of the statistical data about the number of 5C-Ts actually produced.

Rational Revisions of Prior Probabilities on a Bayesian Model

By '**prior probability'** or '**prior**' decision theorists refer to the probability that an individual attaches to an outcome at the beginning of a decision procedure. This initial probability might be one that is statistically grounded, or it may be based (as in the previous example) upon some 'information' and a subjective assessment of the likelihood of its truth, or it may be completely irrational – not based upon any experience or calculation.[7] In the example concerning 5C-T, the proposition that *my* coin is a 5c-T was assigned by me a prior probability of .7 on the basis of stories about my family history. Whether or

[6]Since I know that the probability that *any given* 5C-T is counterfeit is .02, it is indeed tempting to conclude that I should hold the *belief* that *my* coin is counterfeit with *confidence* of .02*.7 =.014, but this is at least awkward (though not inconsistent). As we have seen above, my ignorance cannot reasonably justify a belief that my 5C is less likely to be counterfeit than any given 5C because I am only confident to degree .7 that I have a 5C-T. (Perhaps this is why many tend to think, incorrectly, that the probability of bad outcomes can be reasonably decreased by ignoring risk factors, which in turn may explain why it is that so many inconvenient risk factors are swept under the rug, including, of course, risk factors to the environment. Even clearer illustrations of this point concern risk factors for health. One's risk factor for lung cancer cannot be changed by reducing one's estimates of the years that one actually breathed asbestos; the risk factor depends upon the years one actually did breathe asbestos, -- not one's beliefs about it.)

So, how I should describe *my* epistemological state with respect to *my* coin, if not as a belief with confidence of .014? I suggest that at this point it is not reasonable to draw any conclusion about my epistemological state from the fact that the probability that $P_E(H) = .02$, where my *confidence* that *my* 5C is a 5C-T = .7. That is because the two calculations are incommensurable; one is based upon Bayesian theory about the confirmation of hypotheses; the other is based on commonsense generalization about the reliability of my family history. But there isn't a wider theory about how to integrate the two, and it is clearly unreasonable to presuppose what a wider theory would show.

[7] For a thorough, high-level technical discussion of these issues see (Howson, 2000, pp. 168 – 238, especially, pp. 214-15).

not factual reports, like the reports about my family, should be deemed will depend upon the *context* in which they are given. This is a familiar point in the literature that is aptly stressed by Howson:

> A large book found in a street is by itself not evidence that Jones killed Smith, but given the further information that Smith was killed by a blow to the head with a large object, in that particular street, and that the book was damaged, and had Smith's blood and Jones's fingerprints on it, it is. Evidence issues in the enhancement or diminution of the credibility of a hypothesis, and this capacity will be determined only in the context of some specified ambient body of information. (Howson, 2000, p. 179)

In our example, the 'evidence' supporting the belief that *my* 5C is a 5C-T is a function of the report of my relatives and my confidence in their reliability.

Our success in reasoning with probabilities depends to a large extent upon the way we treat priors, and I do not think the significance of this point is fully appreciated. Perhaps that is because even if an initial prior probability is completely unfounded, it appears that systematic recalculation of the initial probability on new evidence will correct the original misapprehension, which might make it seem that the initial prior makes little or no difference to correct reasoning. This in turn might lead us to think that in taking positions about transformative environmental measures, we can reasonably start from *any prior* with the assurance that subsequent experience will always enable us to correct an 'irrationality ' in the prior itself. Although that is true in principle and certainly holds for textbook examples about, say, coin tosses, we shall see that rational revisions of priors about hypotheses characterizing complicated natural systems will be elusive at best. This strongly suggests that it is wise to take care in assigning priors to hypotheses about the need and utility of transformative measures concerning the environment.

To illustrate the issues involved, I want to begin analyzing *commonsense* revisions of priors. The purpose of this is to try to simulate the way in which ordinary, intelligent, well intentioned people might deal with an unreasonable prior. I think the way that we deal with priors is interesting from point of view of environmental sustainability because it turns out to be an important issue in moving toward consensus about the need for and usefulness of the transformative measures. Reaching consensus is often blocked, in my view, by unreasonable priors that are resistant to revision. The literature on transformative measures for sustainability calls for improved communication and education about sustainability, but it does not seem to me to address the main issue, which is to take into account the ways that ordinary people handle unreasonable priors on a commonsense basis.

I begin with a story in which we see how a prior might be revised on a completely *commonsense* basis, ignoring the details of technical decision theory. Suppose that someone, A, at a party offers to take bets on the outcome of coin tosses. Another, B, 'shrewdly' offers the opinion that A's coin is unfair, and attaches a high probability to the proposition that the coin will come up heads nearly all the time; let's say 9 times out of 10. B's evidence is that he recently saw A in a party shop out of town, where fake coins are sold and are said by 'reliable sources' to be fixed to turn up 'heads' every time. B draws the conclusion that A intends to fix the outcomes of bets made on the toss. Recognizing that the coin may not come up heads every time, even if it is 'fixed,' B attaches a high probability (but not certainty) to the proposition that the coin will turn up heads on any given toss. That means, say, that he expects that in the

long run, out of every 10 tosses, 'heads' will turn up 9 times. B's 'prior' is based upon the 'evidence' provided by his observation of A in the party shop, B's estimation of A's character and thus A's likely disposition to fix bets. But just how strong is that 'evidence'? *That is purely subjective matter*. Suppose that B thinks that the evidence provided by the party shop observation plus his 'knowledge' of A's character should count *just as strongly as if* he had *just seen* an actual, fair coin turn up heads in 9 out of 10 tosses. This expresses B's *extraordinarily high degree of confidence* in his prior (and thus, the '**weight**' assigned by him to the prior), inasmuch as B believes, as any sensible person would, that it is extremely improbable that a fair coin will turn up heads in 9 of ten 10 tosses.

Suppose that the party has begun. After the first 10 actual tosses, the coin has turned up heads 6 times. B recognizes that rationality requires him to take this new evidence into account. How might B modify his prior? Let's suppose that B reasons that *in effect* there have been a total of 20 tosses, 10 virtual tosses (of which 9 turned up heads just as is assumed by B's prior), and 10 actual tosses (of which 6 came up heads). Thus the ratio of favorable to unfavorable outcomes is {[(9 + 6)/(10 + 10)] = .75}, which is now B's revised prior.[8] This surely would come as something of a surprise to B; nevertheless, it would not completely demolish the hypothesis that the coin is a fake, and faked in favor of 'heads.' The party game continues, and the coin is tossed another 10 times. Suppose that the outcome is 4 heads. The new evidence for H now consists of 9 heads of 10 virtual tosses supposed in the prior, plus 6 heads of the first 10 actual tosses and 4 heads of the next 10 actual tosses. Following his original line of thinking, B recalculates his 'prior.' The new result is {[(9 + 6 + 4)/(10 + 10 + 10)] = .663}, still in favor of 'heads.'

Of course, there have been only 20 *actual* tosses, of which 10 were heads. That would lead any 'sensible' person to abandon the original prior, concluding that something had gone wrong. What went wrong, of course, was assigning enormous weight to the very high 'prior,' which is that 9 of 10 virtual tosses came up heads, a calculation based solely on seeing A at a party store where, it is said by 'reliable sources,' that fake coins are sold. Suppose that B nonetheless holds stubbornly to his prior, hoping that further tests will vindicate his confidence. B soldiers on with repeated testing; after a thousand tosses it is very likely, though not absolutely certain (the statistical details are not important), that the initial probability (B's prior) will have been revised to (509/1010), meaning that the initial prior would count as 9 favorable outcomes out of 10 virtual tosses, and the next 1000 tosses would come out 500 heads and 500 tails, reflecting the fact that the coin is actually fair.[9] By now we may assume that any sensible, intelligent and well-intentioned person would have long since concluded that the initial prior was just mistaken and that the original evidence on which the prior was based should also be revised downward,

[8] This is different from the Bayesian conditional probability calculated in the previous examples, because the relation of the second (and third) sequences of ten tosses to each other and to the virtual sequence does not reflect information about a *dependency* of one class of data upon another (as, for example, the mint mark did in the previous example), but reflects information *merely about repetitions of tests of the same form about the fairness of the coin*, which is to say simply the ratio of the number of heads to the total number of tosses. The sequences of tosses are independent of each other (as are the tosses within the sequences). More succinctly, we *cannot define a conditional probability without a condition*.

[9] The price of dogmatism is high indeed. See: (Howson, p. 185 – 89, especially p. 189) for a discussion of just how high the price can get, which is the refusal to deem any possible evidence sufficiently strong to dislodge a problem prior.

though perhaps not all the way to zero.[10] In the long run, it appears, irrational priors can be corrected by repeated testing.[11]

Indeed, this may seem to show that the prior assigned to an event does not make very much difference, when it comes to making predictions about complex systems like the environment. But nothing could be farther from the truth. That is because in dealing with complex systems we do not have the opportunity to perform *innumerable independent tests* under *exactly the same conditions* that are *virtually indistinguishable* with *unmistakable* outcomes. Coin tosses can be repeated indefinitely under virtually identical conditions that do not materially affect the outcome of other tosses, where the outcomes are unquestionably identifiable as heads or tails. When it comes to complex eco-systems, these cooperative testing conditions do not hold. This suggests that in the analysis of the complex systems of the natural world it will be more important than ever to hold *tentatively* to *cautious* priors.

In fact, even in the cases of coin tosses, where experimental conditions are favorable, it is easy to see how a cautious approach to priors pays off. Suppose that B had decided -- devaluing but not ignoring the conversation at the party store -- that he would assign an initial, unconditional probability of slightly over .5, say .55, to the proposition that heads would come up on any given toss. Assume further that he reasonably supposed that his initial probability should have the same weight in any recalculation of probabilities as, say, 5 actual tosses. In this case after 20 actual tosses with 10 positive outcomes, factoring in the 5 virtual tosses of which .55 would come out heads (in the long run), the revised prior would have been {[((.55 * 5) + 10)/25] = (12.75/25) = .51}. On the other hand, suppose that B incorrectly assumed that the coin was fair, assigning the weight to his prior of .5 heads out of 5 tosses. Suppose further that of the first 20 tosses, 18 turned up heads. B's prior, .5, would then be revised to {[((.5*5) + 18)/(5 + 20)] = .82}, rapidly correctly the mistaken prior.

All this shows that commonsense reasoning that assumes high priors (or even worse, high confidence in high priors) is very risky,[12] where confidence is represented by the number of virtual tosses, as in the previous example. We have also seen that the less conservative the prior, the greater the likelihood of error. Howson describes the situation in this way:

---

[10] After all, there is still the 'evidence' from the party shop that induced the unconditional probability, and B may believe that it is possible that the coin has been cleverly fabricated to suppress its inherent unfairness during the initial tosses so that potential gamers would become unreasonably confident that the coin is fair.

[11] The matter may be complicated in a further way. We have assumed that B's confidence in the prior was proportioned to a scenario in which 9 of 10 tosses came up heads. But suppose that the prior had been assigned a 'weight' of 15 tosses. Given the same results, the prior would have been recalculated after 1000 actual tosses as 514/1015: Hardly different from the first result based upon a confidence factor, viz., 'weight,' of 10 tosses! That is what I meant by a previous remark that large differences in priors *or* errors in assessing the strength of 'evidentiary' support for them can be reduced to almost nothing in the event of long-term testing.

[12] We have discussed the issues of revising priors on a commonsense basis, but what does probability theory itself have to say about the matter? Actually, calculating the probability of sequences of favorable outcomes to total outcomes, where the probability of a favorable outcome varies, is a subtle problem The key to the solution is to compartmentalize the data into sequences of n-tuples, and then to measure the frequencies of n-tuples according to occurrences of favorable outcomes in each. The resulting proportions of favorable outcomes in the sequences will indicate the distribution of outcomes in a given collection of n-tuples. This difficult technical problem is addressed in (Howson and Urbach, 1993, ch. 13 and summarized in Howson, 2000, pp 233-38.) Their discussion shows that shaping data so that it can be addressed by technical probability theory is a difficult problem in itself, even when it comes to relatively simple data structures.

> Inductive reasoning is justified to the extent that it is sound, given appropriate premises. These consist of initial assignments of positive probability; they cannot themselves be justified in any absolute sense…no theory of rationality that is not entirely question-begging can tell us what it is rational to believe about the future, whether based upon what the past has displayed or not. This is not to say that evidence tells us nothing. The trouble is that what it does tell us cannot be unmixed from what we are inclined to let it tell us. Increasing observational data certainly, provably, reinforces some hypotheses at the expense of others, but only if we let it by a *suitable assignment of priors*. (my emphasis) (Howson, 2000, p. 239f)

Epistemological Weight Assigned to Priors

Rational risk assessment will take care to compare the weights assigned to priors across a range of comparable data, that is, where priors were assigned similar measures on comparable evidence. In fact, suppose that B determines that observations in *situations resembling* those of the out of town party shop were misleading in many previous cases; say in 70% of the cases. This suggests that relatively little weight should be attached to the out of town party store observation. It is a matter of great theoretical complexity and controversy as to exactly how much weight to assign to a prior; however it is perhaps clear that statistical information is needed about when and how we have been misled in assigning priors in various types of situations. In other words, we can imagine classifying priors in order to assess the likelihood of error according to type.[13]

As far as I know, Hume was the first person to see the importance of this point. In a much maligned[14] section of *Treatise of Human Nature*, Hume writes:

> In every judgment, which we can form from probability, as well as concerning knowledge, we ought always to correct the *first judgment*, deriv'd from the nature of the object, by *another* judgment, deriv'd from the nature of the *understanding. (my emphasis)* (Hume, Selby-Bigge/Nidditch, 1978, p. 181f) and the same passage in (Norton and Norton, 2006, §1.4.1.5/p.122)

Hume's view marks him as a Bayesian, because he is essentially saying that in assigning any prior, we need to consider past experience with similar priors, where our 'consideration' depends on the *understanding*. But just how does the understanding operate? The strength of the 'evidence' for a prior is not something that can be straightforwardly measured like coin tosses, even though it is a matter of

[13] This reasoning could also be applied to the previous example where confidence that a coin was a 5C-T was based upon the testimony of relatives. It would have paid to consider their record in similar cases.

[14] Hume's point is much maligned not on the intrinsic merit of the observation that we must always be open to reassessment our arguments, both with respect to structure and content, but rather because he drew the conclusion that as we repeatedly try to correct our confidence in our arguments, our confidence in our judgment must decrease to zero, since the arguments on which we rely in assessing the merits of arguments will obviously be arguments themselves, which in their turn will be vulnerable to reassessment by yet further arguments. For discussion see (Garrett, 1997, pp. 222-28)

rationality and *commonsense*. In the party store case, it would have been wise for B to reflect carefully on how many times he had gone wrong in assuming the worst about someone's intentions on the basis of a chance encounter in an underdetermined context. Hume is suggesting that the revision of priors must be based upon a methodology that evaluates epistemological principles by which successful revisions of priors may be rationally evaluated; in other words, classifying priors so that they can be evaluated according to type. Ideally, those types would be conceived so that they could be incorporated in Bayesian conditional probabilities, meaning that we could quantify the types in question so that we could calculate the probability that something will be of a certain type *given that it is another type*.

When High Priors Are Warranted

It may seem that the argument has now taken a very conservative turn, casting doubt on any assignment of high priors. But this is not so, and it is important to see just when we can proceed with confidence. Consider, for example, Newton's Second Law, F = m*a (henceforth 'the second law'). This generalization is so deeply entrenched in physics that it is virtually impossible to imagine any evidence dislodging it. The reason for its secure position at the center of natural science is that it figures essentially in innumerably many correct predictions made over hundreds of years in a *wide variety* of contexts. Thus, even in the face of contrary evidence, we would be very unlikely to retreat from the second law. As Quine famously observed long ago, we quite rationally make 'adjustments' elsewhere in our theory to accommodate new data without sacrificing very well confirmed, deeply entrenched beliefs, like the second law. (Quine, 1964, p. 42) Indeed, it may seem that the weight assigned to the evidence supporting the second law should render it beyond revision, but that is not so. Even highly confirmed scientific laws are not beyond revision. Who would have thought before 1905 that Euclid's Fifth Postulate would be shown to be false of physical space?

None of this justifies high priors for theories that are *at the periphery* rather *than in the core of established science*. Tossing off worries about climate change or arguing that Armageddon is upon us are more like B's extravagant party store conjecture and less like the proper respect shown for core principles like the second law. Even in tightly controlled, virtually ideal situations like coin tosses we found that modesty in assigning priors is prudent. How much more important it is to be cautious when we do not have the advantage of repeated low cost experiments that can correct extravagant assumptions! Assignment of high priors on the basis of unjustifiably inflated evidence obviously does not have the authority of established science. This fact explains just why it is that in political contexts each group will try to place itself on the side of confirmed science, and thereby draw upon its established authority. It also shows how important it is to respect established scientific accomplishment and not to squander its credibility for illusory advantage in argument.

Revising Priors on Expert Testimony

It is time to return to the main issue: How to form rational beliefs about policies that will maximize the probability of promoting environmental sustainability at the lowest cost. Schemes to ensure environmental sustainability depend upon rational assessments of the distribution of positive and negative

environmental outputs on the basis of proposed transformative measures. The weight to be assigned to priors offered by theorists concerning environmental sustainability will depend upon their qualifications as well as the depth and breadth of their research. Of course, the opinions of experts whose priors are most credible should have outsized influence on the development of public policy. [15] Nevertheless, successfully implementing transformative measures even in relatively minor matters will depend crucially on the cooperation of the mass of the population, and the wider population cannot be expected to accept sacrificial measures unless they believe that those sacrifices are actually necessary. Unfortunately, we cannot count on the wider population for rational priors.[16]

The belief that sacrifices are warranted will initially depend upon the prior attached by the great mass of people to the proposition that the environment is in danger, and upon the prior that they attach to the estimated costs of proposed transformative measures. Much attention has been given to problems dealing with public education in academic literature focusing on transformative measure. Batteen emphasizes that sometimes the media exaggerate risks to the environment, for example in recent discussion concerning the possible Arctic ice melt, which caused many to worry unnecessarily about the 'imminent' extinction of the polar bear. (Batteen, *et. al*. p. 87)

Of course, it is more common for public perception to err by *underestimating* the danger of environmental threats.[17] No doubt priors assigned to possible environmental threats will differ considerably from person to person and sub-culture to sub-culture. As we have seen, revising unreasonable priors involve repeated tests, which are simple to contrive and analyze when we are worried about mere coin tosses, but not when dealing with risks to eco-systems. It is obviously difficult or impossible to devise simple tests or experiments that simulate Bayesian or other technical models that establish standards by which priors can be evaluated. When it comes to the environment, each argument will turn upon indefinitely many priors concerning proposed transformative measures that bear upon a large number of possible risk factors to the environment. It is not clear what would or could *count* as an analogue of a simple coin toss that would warrant change in a prior about a proposed transformative measure concerning the environment.

To be sure, in cases involving immediate and potentially catastrophic risks to sustainability there is greater promise of coming to reasonable agreement about risk factors and the policies needed to deal with them if only because society is forced to provide the resources needed to come up with an immediate, rational response to potential disaster. Even in those cases, however, there are only a limited number of 'experiments' that can be performed. To take an admittedly extreme but pressing example, think of nuclear disasters: Just how many more 'experiments' can we afford? Moreover, testing different

[15] This is a point that has been repeatedly emphasized in the literature, and I therefore see no need to pursue it here in detail. For examples, see: (Leclerc, *et. al*., 2010, pp. 351 – 74), and (Snow, *et. al*., 2006, pp. 381 – 87).

[16] It takes us somewhat beyond the scope of this paper, but the points made here about priors are substantiated by their applicability to disagreements in nearly all subjects. For example, assigning very high priors to tenaciously held religious beliefs means that rational belief revision will be virtually impossible if only because people who put themselves in dogmatic positions will not budge except on the basis of overwhelming contrary evidence. But in religion -- as in so many value-oriented fields -- it is difficult to fashion experiments or observations that even bear upon dogma. Of course this is not offered as an argument against religion, but rather as an argument for lower priors when it comes to religion.

[17] For more misconceptions about priors in assessing environment dangers, see: (Leung, 2010, pp. 39 – 46) and (Thorpe, 2010, pp. 239 – 46).

concepts in the construction of nuclear reactors or the sequestration of carbon gases is expensive and time-consuming. To complicate matters further, risk assessments need to be interpreted in light of local conditions, which may not be widely applicable. When it comes to transformative measures dealing with relatively minor issues, like forest fires in the western regions of the United States or the generation of smog in cities or plastic bags and bottles, there will be even more uncertainty, because there are limited opportunities for rational revision of priors that are applicable over a wide range of cases.

The point of all this is that the presence of considerable *entrenched* disagreement about the seriousness of the environmental change shows in itself how far we are from being able to persuade people to revise priors on a rational basis and thereby to reach the sort of reasoned accord that is necessary to support transformative measures. In fact, there is intense disagreement about the urgency of sustainability precisely *because* there isn't an easy way to induce people to revise stubbornly held priors about sustainability on a rational basis.

Implications for the Formation and Justification of Policy Proposals Concerning Sustainability

These reflections are offered in the hope of stimulating debate about steps that might increase the credibility of forecasts about sustainability. The first step, at least when it comes to analyzing relatively minor inputs, is to compartmentalize, that is, to expand and to shape our data bases.[18] There are at least two dimensions to the process. The first is to divide the subject matter into manageable units of investigation. Many of the modules are obvious; they might well include recycling vegetable and animal waste, recycling or eliminating industrial gases, and recycling or eliminating nuclear waste. Beyond that there is the matter of identifying regions of interest. Some nations and/or regions will be dramatically affected by certain forms of environmental degradation while others will remain relatively unaffected. The results of multi-disciplinary research concerning regions need to be integrated at a higher level with results from other regions. That way we shall have a better idea of just how great *global* threats to the environment really are; where they figure prominently; and exactly what has worked in dealing with them. Ideally, the data will be shaped so that hierarchies of Bayesian conditional probabilities can be properly defined. By a hierarchy I mean an expansion of the scope of conditional probabilities by incorporating them into the analyses of wider populations, as previously illustrated by the examples about counterfeit Cs

Unfortunately, there is considerable doubt among the public that sacrificial measures to ensure environmental sustainability are necessary. For many it is difficult to believe that an environment that has sustained life forms with complicated neurological structures for so long can be easily threatened by mere human activity. This is just another way of saying that *many have attached a high prior to the proposition that earth is resilient to environmental insults; in fact so resilient that we really need not worry much about environmental degradation.* Yet even a little reflection on mass extinctions should give us pause in counting heavily on the resilience of the environment. Because there aren't straightforward ways (like coin tosses) to revise entrenched priors, some have dramatized dire environmental outcomes in the attempt to awaken others. Yet those attempts can backfire if they go too far, because the

[18] This echoes one of the main points developed in (Snow and Snow, 2010, p. 387).

dramatizations that they involve may not be based upon solid scientific evidence (and, in fact, may exaggerate the weight of the evidence that we do have).

The central problem encountered in promoting environmental sustainability is tenaciously held, high priors that cannot be recalculated on a reasonable basis in a relatively short period of time with little expense. In order to draw reasonable conclusions about the best course of action in dealing with environmental threats, one would hope that enhanced research would put scientifically qualified opinion-makers in a position to speak with one voice on the issue. That way we shall be justified in assigning high priorities to carefully crafted policies that address genuinely dangerous practices, thus enabling us to forestall environmental damage.

This way of dealing with situation is likely to require costly empirical research and integration of results on a scale that is greater by orders of magnitude than anything that we have seen so far. On the brighter side, an integrated global effort directed by the intellectual community will likely have the effect of reducing contentious disagreement among the wider public, which in turn will facilitate a world-wide response to environmental threats. This is the usual salutary effect of successful scientific investigation. Moreover, it is the sort of experience of cooperative effort in the service of a cause greater than any region or time that has the potential to promote environmental sustainability as a culturally invariant value, -- invariant geographically, over the regions of Earth, and diachronically, over the generations.

Conclusion

This paper has argued that at this point standard textbook Bayesian models for assessing probabilities cannot be applied in any straightforward way to issues concerning environmental sustainability, principally due to their immense complexity. Furthermore, commonsense ways of revising priors will be helpful only under extremely unusual, simplified conditions. The principal contribution this paper hopes to make is to explain how it is that tendencies to exaggerate the epistemological weight of evidence and to assign unreasonably high priors undermine constructive discussion about sustainability, especially in the popular media. That is because there aren't easy experiments to correct the misalignment of priors or the 'evidence' on which those priors are based. This helps explain why it is that there are polarizing, entrenched positions on the environment that frustrate attempts to form and implement rational policies for sustainability. Although popular dramatizations are undoubtedly helpful in drawing attention to the issues, they really do not address the essential problem, which is to dislodge careless priors that stubbornly resist revision. Encouraging *modest* priors is the first step in moving toward consensus concerning sustainability. Although modest priors are necessary, they are not sufficient. Sensible starting points in thinking about sustainability must be accompanied by increasingly detailed analyses of the threats to the environment and the costs of possible transformative measures, whether short or long term. This calls for disciplined research and concomitant investment in database management on a scale that is barely imaginable. The call for further research is not a demand for re-examination of the long-established – for example, 'additional confirmation' of Galileo's law of freely falling bodies by dropping more cannonballs from the leaning tower in Pisa. The augmented research that is needed is multi-disciplinary research that can precisely measure the effects of transformative measures over wide range of disparate environmental venues, and shape the results in hierarchies of interdependent types or classes that can be structured as Bayesian conditional probabilities. The costs are great, but the potential reward is

also great, because it holds the promise of moving toward consensus about transformative measures that will promote environmental sustainability as a culturally invariant value.

References

1. Batteen, Stanton, Maslowski, 'Climate Change and Sustainability: Connecting Atmospheric, Ocean and Climate Science with Public Literacy,' (in Reck, 2010).
2. Bonner, Charles, 'Images of Environmental Disaster, Information and Ontology, *Forum on Public Policy*-VOL 2011, No. 2.
3. XXXXXXXXXXXXX, 'Evolution and the Goal of Environmentalism,' *Forum on Public Policy-e*, VOL 2011, No. 2.
4. Garrett, Don, *Cognition and Commitment in Hume's Philosophy*, Oxford/New York, Oxford University Press, 1997.
5. Leclerc, Holland, Foken, and Pingintha, 'Sustainability of Gaia: A Question of Balance' (in Reck, 2010).
6. Howson, Colin, *Hume's Problem, Clarendon press, Oxford, 2000.*
7. Howson and Urbach, *Scientific Reasoning: The Bayesian Approach*, 2nd edition, LaSalle/Chicago, Open Court, 1993.
8. Hume, D., *A Treatise of Human Nature* (originally published 1739), in Selby-Bigge and Nidditch, eds., A *Treatise of Human Nature*, second edition, Oxford, at the Clarendon Press 1978/1985); also, Norton and Norton, *Hume, A Treatise of Human Nature*, Oxford, at the Clarendon Press, 2006).
9. Joyce, James, 'Bayes' Theorem,' *Stanford Encyclopedia of Philosophy* {\url\{plato,Stanford.edu /archives/fall 2008 /entries.bayes theorem}}
10. Justice, Creek and Buckman, 'Ideological impacts upon environmental problem perception, *Forum on Public Policy-e*, VOL 2011, No. 2.
11. Leung, Solomon, 'Global Environment Sustainability: From Developed to Developing Countries' (in Reck, 2010), pp, 39 – 46.
12. Quine, Willard, 'Two Dogmas of Empiricism' (in*: From a Logical Point of View*, Cambridge, MA, Harvard University Press, 1964) pp 20 – 46.
13. Reck, Ruth ed., *Climate Change and Sustainable Development*, Yarton, Oxon (UK), Linton Atlantic Books, Ltd, 2010.
14. Snow and Snow, 'Climate Change and Challenge for Costal Communities' (in Reck, 2010).
15. Tecle, Aregai, 'Sustainable Management of Natural Resources in an Era of Global Climate Change' (in Reck, 2010).
16. Thorpe, H. R., 'Habitat Restoration: Aspect of Sustainable Management' (in Reck, 2010).

Published by the Forum on Public Policy