

Shoot to Kill? Bentham's Defence

January 13, 2006

Chapter 1

Introduction

1.1 Context

This document contains a number of emails that appeared on the email list of the Leverhulme Evidence Programme at UCL.

The exchange arose spontaneously on the email list and was not formally organised in any way.

1.2 Notes on this compilation

Email exchanges are difficult to follow at the best of times because different strands of discussion go off at different tangents: here an eclectic approach is taken.

Chapter 2

The Exchange

2.1 Tony Gardner-Medwin opened the discussion ...

(Saturday, August 20, 2005 11:58:41 am)

Londoners and evidencers,

Y is on trial for murdering X. Y estimated at the time (based on evidence and inferences available to him) that there was a probability P (his subjective degree of belief) that Y was an armed suicide bomber, one of a set who had so far killed on average 6.5 persons each. Y confronted the exhaustive alternatives of either killing X (with estimated probability 10% of thereby detonating a bomb) or allowing X to activate any bomb, which he believed X would (if a bomber) in the circumstances intend to and be able to carry out with probability 90%.

What value of P (the probability that X was a bomber) would justify Y's decision to kill X? Should the jury convict on the grounds that Y should only have killed X if he was "sure" that X was a bomber (e.g. $P > 67\%$, 90% or 99%)? Or should it acquit on the basis that maximisation of expected utility (minimisation of the expected no. of innocent lives to be lost) was satisfied by $P > 16\%$? **

Bentham's defence of Y rests on the grounds that it was reasonable for Y (and/or his colleagues) to estimate P as $> 16\%$, even though none felt the evidence rendered P close to 100%. Are we happy with this? If not, why not?

Tony GM

** Simplifying the maths slightly: with $P = 16\%$, killing X would kill an innocent person with $P = 0.84$, while not killing X would kill 6 innocent people with $P = 0.14$, an expectation of 0.84 lost innocent lives.

2.2 Mike Redmayne

wrote (Monday, August 22, 2005 4:32:57 pm)

I think there are a number of simplifications in the argument that Y's killing X is justified where $P > 16\%$. First, as a matter of human psychology, I don't find it terribly realistic that Y will have a finely graded belief in the proposition 'X is a bomber'. I am not unsympathetic to the argument that we can grade a good number of beliefs in terms of probabilities, but when someone is acting in extremis, in response to a threat, I do suspect that beliefs are closer to yes/no than graded in continuous probability increments. But's let's ignore that, and go with the argument. A second problem is that many people think that there is a difference between killing and letting die, in other words that it is worse for someone in Y's position to kill an innocent X than to let a guilty X kill an innocent. The simple utility calculation does not take that into account. Of course, hard line utilitarians deny that there is a distinction between killing and letting die. But then it seems they have to sign up to the following: Y is sure that only one of a group of 5 people is a bomber (with probabilities of bomb detonation etc as below), but is not sure which, and has no time to find out. Obviously, Y should kill them all in order to maximise utility. Many people will find that an unpalatable conclusion.

Mike Redmayne

2.3 David Colquhoun

wrote *in response to the two preceding messages* (Tuesday, August 23, 2005 9:50:38 am)

I imagine most people's answer to Tony's penultimate question is that they were not happy at all with $P > 0.16$. As an experimentalist, and (in the absence of any real prior information) unrepentant frequentist, I would say that the reason for that is that the entirely subjective probabilities that go into the calculation cannot be justified, so garbage in, garbage out.

Even if that were not the case, there are a lot of things that are relevant to the calculation that have not been included in it. Perhaps the most obvious is the general desirability of having a police force that is allowed to shoot down citizens because they suspect (but are not sure) that their victim might be about to do harm. I'd guess that most people would think that sort of society was so undesirable that it would require virtual certainty (like seeing the bomb) before shooting a suspect was justifiable. No doubt that view could be put into numbers, but the numbers would be spurious – they would be chosen retrospectively so as to produce a conclusion that was already thought desirable on ethical grounds. In that case they can hardly then be used to justify that ethical view.

David Colquhoun

2.4 David Colquhoun

also wrote (Tuesday, August 23, 2005 7:05:11 pm)

After re-reading Tony's question, and my reply, a slightly more general point occurs to me. Clearly probabilities are involved. The problem is not so much that they are not interpretable as frequencies, but simply that they are not known. Arguments of the sort he presented for discussion are dangerous because they give a (largely spurious) air of numerical respectability to what, in the end, are arguments about ethics and opinions. Putting numbers on things is my trade, but they have to be soundly based numbers.

I can also see that calculations of the sort that Tony gave can be useful for getting a feel for a problem, exposing implicit assumptions and exploring the consequences of different assumptions. But only as long as one doesn't take the numbers seriously.

Moral. Bayesian approaches (to problems of this sort) have their place, but are best restricted to consenting adults in private.

David Colquhoun

2.5 Tony Gardner-Medwin

wrote (Tuesday, August 23, 2005 8:34:07 pm)

I haven't wanted to respond with my own views about Bentham's defence till others have had plenty of chance – which of course is still there. I've had 2 private responses I hope people may put on the list.

However, I think it's worth trying to deflect the issue from one about whether one can or can't assign good numbers to quantities involved. Obviously debates about precise numbers are futile – pretty much like for climate change. But the essential question is not about numbers, but about whether one wants a policeman (yes, Y is a policeman!) to weigh the enormity of what may happen if he doesn't shoot, to overcome a degree of belief that may be very far short of what David Colquhoun calls 'virtual certainty'. I'm not saying you do want this. But Mr. Bentham is deeply worried that if you don't, the well-being of the greatest number isn't getting much of a look in. Why exactly should it not?

I do concur rather with Mike Redmayne, that an answer may lie in considering utilities that are actually about the kind of society people want to live in and find palatable – not necessarily with a rational basis. I used a similar trick to try to reconcile liberal legal principles with utilitarian logic in my article about probability of guilt in jury trials. But in the present case, I'm not sure I want to depart from immediate utilities so far. It's hard not to imagine oneself in a tube train watching a new Y and a new X, and thinking "For God's sake pull the trigger – you don't have to be certain!"

Tony GM

2.6 Tony Gardner-Medwin

wrote (Wednesday, August 24, 2005 11:54:32 am)

Someone pointed out that there is a (fairly obvious) typo in my original email, where in the 2nd sentence it should of course say "P.. that X was an armed suicide bomber", not Y. Apologies.

TGM

2.7 Ian Dennis

wrote (Wednesday, August 24, 2005 5:37:06 pm)

Sent on behalf of Ian Dennis by Tony GM:

Dear All,

I've read this hypothetical with interest. Tony's question needs to be put into its legal context.

In this scenario for Y to be convicted of murder the jury would have to be satisfied beyond reasonable doubt (alternatively satisfied so that they were sure) that Y killed X unlawfully intending to kill him or at least to cause X serious injury. The killing appears to have been intentional, therefore the only issue is whether it was unlawful. It would be lawful only if Y has a valid defence. The only defence currently applicable in English law is the use of reasonable force in the prevention of crime.

Then, if X is in fact a suicide bomber, and Y knows this, it is for the jury to decide whether Y used reasonable force in the circumstances. For force to be reasonable it has to be necessary (ie no viable alternative) and proportionate to the threat. This is an objective question, but the judge can tell the jury that Y's own opinion as to the necessity and degree of force required is entitled to great weight. It was said in one case that the law does not require force to be measured in "jewellers' scales".

If X is not in fact a suicide bomber, but Y thinks he is, Y has made a mistake. The current law states that Y is to be treated, for the purposes of the law of reasonable force, as if the facts he believed to exist did exist. Provided his belief is genuine, the issue for the jury will then be the same – was Y's use of force reasonable in the circumstances as he believed them to be?

The hypothetical suggests that Y does not know or have a positive belief that X is a suicide bomber. However, if he is acting for a lawful purpose (preventing crime), the question for the jury will still be one of whether they think a reasonable man would have been justified in using the force Y used in the circumstances that Y believed to exist (ie that there was a probability that X was a suicide bomber). This is a question which for courts and lawyers is context-dependent and also value-laden. If Y is a police-officer who had to make a split-second decision whether to shoot a possible terrorist the judge is likely to stress the point about the importance of Y's judgment of necessity in the circumstances. I'm confident that in practice a court would discourage resort to mathematical calculations about the precise degree of risk.

However, if Y's judgment about the probability of X being a suicide bomber was formed in a way that was reckless or grossly negligent then he might find himself liable for manslaughter. The inquiry will then be focusing on different

issues of the adequacy of police information, the organisation of the operation etc.

Regards.

Ian

Ian Dennis, Professor of Law

2.8 Richard Leary

wrote *replying to Ian Dennis' message* (Thursday, August 25, 2005 9:24:25 am)

Bentham

Dear All,

I have been following this with great interest. The way the arguments have 'fermented' in an effort to 'tease' out the relevant legal, scientific, mathematical and to some degree philosophical issues is enlightening in terms of evidence and argumentation in itself. Anyway, back to the original issues.

My interest in this discussion is based upon a rather unusual series of stand-points. Firstly, as a member of the project, secondly as someone who has spent some time studying the evidential aspects of reasoning in crime investigation and probably most interestingly in this discussion as a former police officer. With regard to the latter standpoint, I have on two occasions faced making the making the very decision suggested: Whether to exercise force based on my judgement that would in all reasonable circumstances result in the death of the suspect ('legal' use of a firearm in exercising force).

The point that I believe is of great relevance right now is the 'shoot to kill' policy currently subject of news and an investigation by the PCA following the fatal shooting of the foot passenger on the London Underground. Keeping in mind issues of 'sub judice' I will limit my comments to the general issues it raises rather than the specific issues of that case.

The script Tony prepared reveals an important point: How a police officer acting in these circumstances can show necessity. This for me is the crucial issue in evidence and is a matter that must be addressed subjectively in terms of the state of the mind of the officer at the time. The variables involved in such a decision are diverse, complex and subject to different levels of reliability. They are too wide for discussion here. One can imagine many issues that need to be considered in the string of events leading upto and then immediately before the officer exercises his/her judgement. Some may say also that judgement about the degree of force exercised is also important because the officer may be able to incapacitate the suspect rather than inflict fatal wounds.

Firstly, let me introduce my experience. Never at any point in my training was I ever told to do anything other than inflict fatal force. The use of warning shots or shots to any part of the anatomy other than the chest was strongly discouraged. In short, once the decision was made to resort to force it was fatal force and nothing short of it would do. Secondly, in all my training I was trained, examined and critiqued about the use of verbal and visual warnings to the suspect that fatal force was a choice being contemplated against them and that if they did not comply with commands it would be used. This was a psychological nightmare for me because

I knew in my mind that I would never resort to such force unless I or another person was at risk of force being used against me or another person. However, the warnings were directed in such a way that left the suspect in no doubt that the officer (myself) was about to release that force. The verbal warnings were aggressive, uncompromising and unpleasant to anyone who heard them. The idea was of course to exercise strict and utter control of the suspect and the environment – not to shoot unless absolutely necessary.

Could this be described as a Shoot to Kill Policy? I think it is such and would not argue otherwise. However, it is not a 'Shoot Without Warning policy' and that in my mind is the central issue here. To exercise such a policy would in my mind require a very high degree of proof that the suspect would (without the use of such force) kill people without warning also. The real question therefore is how the officer forms such a view and why. Furthermore, does the officer resort to the use of near linear likelihoods, reasoned assessments, luck and fortune, or perhaps most importantly, following of orders from a command structure (senior officers). Subjective questions about necessity fall into a different light once these are considered. The issue of 'split second' decision making is in my view only an issue practically once the suspect is engaged. Before then the situation will be more remote and often have the luxury of time to make considered decisions. I do not subscribe to the view that these situations require split second decision making all the time. They do not.

If a suspect is being followed covertly then the real issue to be addressed is one of (a) Identity and (b) Intention. If both of those questions can be addressed properly then good decision making can follow. The problem of course is that these are difficult questions in their own right and often require the power of 'insight' as well as the evidence needed to make the decision. Furthermore, the decision point on intervention and the challenging of a suspect needs to be considered. The police officer in England and Wales has always been charged with the responsibility (and long may it be so) of preventing crimes rather than detecting them. When the police are in control, this means intervention before the crime is committed and not after it has happened. However, there are some circumstances that may require the decision to be delayed so as to discover additional evidence or intelligence that may save other lives in time. This might mean 'letting the suspect run' until co-suspects can be identified, a 'safe house' is discovered or additional evidence is uncovered of other crimes. This is a very controversial issue and one that requires careful handling. The decisions again should in my view be based on prevailing circumstances supported by evidence. Subjective tests later will only expose the officer to pressure to reveal what these were so it is best to deal with them 'real time' and as events unfold. (Record them and the decisions as Bentham would have suggested – professional people write things down! Of course I am not suggesting that the officer places a weapon on the ground or back in his

pocket whilst making copiuos notes and observations. But, technology can and will record voice, visuals and so on).

The 'Shoot to Kill Policy' should in my view be termed the 'Shoot Wthout Warning Policy'. This would create a degree of clarity that currently does not exist. It would also allow everyone involved to understand exactly what is at stake, the decision making required and the standards required to make those decisions to exercise force without warning. The evidence required to make those decisions then has a degree of clarity that at present is missing. My fear is that officers who are well intentioned become embroiled in the use of a policy that fails in clarity, is misguided in terms and results in public disquiet about what in reality are common 'real world decisions' based on evidence. I also fear that the reason it is not termed 'Shoot Without Warning' is that it is too controversial. In short, the decision making at Command Level are too difficult to justify especially when they need to be evidenced.

My own experience is relevant. On two occasions I have had cause to consider the use of a firearm to exercise force but on both of those occasions (very luckily) I exercised restraint and all turned out well. Had I acted and used force in one of those cases I would have killed the innocent occupier of a house we were protecting who decided in the dark of night to walk un-escorted around a house we were expecting to be burgled. In the other case the suspect gave up quietly after a very strong warning about the consequences of his failure to do so. Interestingly in the first case I was following the 'suspect' (the intended victim of the burglary) in an effort to discover where he was going and what he was doing. No warning was given for that reason. Eventually, and by accident, I realised it was the intended victim because I recognised him. I recall well how I felt and it was not good.

The evidential arguments about reasoning in these circumstances bring into focus the problems faced by ordinary people as decision makers. My recollection of the way I conducted myself was that I activelt sought evidence and information to make the decision

2.9 Richard Leary

wrote (Thursday, August 25, 2005 9:31:23 am)

Bentham continued....

The following paragraph should be added to my last E Mail. This paragraph failed to send:

My recollection of the way I conducted myself was that I actively sought evidence and information to make the decision. However, this was not based upon subjective numeric probability. It was the amalgamation of many sources of evi-

dence that could impact upon the unfolding scenario. My best description of what I experienced was one where I was constantly considering the reduction of risk. Prior information was involved in this but I did not and still do not recollect any notion of numeracy, subjective probability in numeric terms or the weighing of competing scenarios in any formal sense.

Richard Leary

2.10 David Lagnado

wrote in response to David Leary's message (Friday, August 26, 2005 10:13:15 am)

Your account raises many fascinating psychological issues. One distinction that may prove useful here is between implicit and explicit decision making. For example, it is possible that you weight alternatives etc. without any explicit awareness of doing so (there is a large psychological literature on this). So your recollections may not always reveal how you actually processed the information at the time of decision. This of course makes prior (naturalistic) training etc very important, as this has a strong impact on how these implicit processes develop. In fact the kind of thinking involved in these situations seems to be a complex mix of both implicit and explicit processing – both proceduralised thought (possibly inaccessible to conscious awareness) and explicit thought (that you are aware of and can later report).

cheers

Dave

2.11 John Adams

wrote (Sunday, August 28, 2005 10:59:14 pm)

Sent on behalf of John Adams by Tony GM:

I think the core of the problem is Tony's appointment of Bentham as the lawyer to defend Y the shooter. Bentham's pursuit of the greatest happiness for the greatest number has long been acknowledged as the inspiration of cost-benefit analysis. But Tony's formulation focuses exclusively on the number (of lives that might be lost or saved) to the neglect of happiness.

Let us introduce Vilfredo Pareto as posthumous attorney for X the shot. X, briefly, and all those who grieve for him are very unhappy. How should their unhappiness be weighed against the hypothetical happiness of those who were not killed, and indeed may never have been at risk?

Pareto introduced a simple rule for valuing the costs and benefits of projects that would benefit some at the expense of others: the benefits should be valued in terms of what the potential beneficiaries would be willing to pay for them (WTP) and the costs should be valued in terms of what the potential losers would be willing to accept as compensation for them (WTA).

This rule whose fairness none contest has created an insuperable problem for cost-benefit analysts ever since. For example, the Department of Transport proposes to build a motorway through my back garden, or house. They have little difficulty coming up with estimates of the value of time savings of the motorists who would benefit from the scheme. But how should they value my losses? They have to ask me what I would be willing to accept as compensation. Let's assume I am old (true), happy in my present house, and unable to imagine how I might spend millions. So I say no amount of money would compensate me for the loss of my house and its neighbourhood connections. It takes only one infinity to blow up the whole cost-benefit exercise.

When people are confronted not with the threatened loss of their houses, but risks to their lives they are even more unlikely to place a finite value on the potential loss. Spending it after you are dead is difficult.

If the "project" in question is a shoot-to-kill-suspected-suicide-bombers policy frequentists have a contribution to make. An actuary asked to comment on the relative risks of being killed by a terrorist or a motorist would conclude that just about everywhere in the world outside Baghdad you are at much greater risk from the motorist. So what might "we" be prepared to pay to reduce the small risk of being killed by a terrorist? And what might "we" accept as compensation for the lost civil liberties and increased risk of being killed by an anti-terrorist policeman? Who is "we"?

Sadly "we" is not a homogeneous group of like thinking people. "We" will

value the benefits and risks of a shoot-to-kill policy differently depending on whether we see it as a cost or a benefit. “We” therefore will defeat all attempts to devise a quantifiable, Benthamite formula to answer Tony’s question. For an elaboration of these thoughts see “What kills you matters”:

<http://www.socialaffairsunit.org.uk/blog/archives/000512.php>.

Prof. John Adams, Geography Department, UCL

2.12 Michael Joffe

wrote (Monday, August 29, 2005 12:00:32 pm)

I have been following this discussion with a lot of interest. I think the main question is whether we should analyse the situation in terms of some model that takes rationality as its basis, or whether it should be in terms of actual psychological processes. The use of probabilities is just one version of the former, but there are others, and the term utility generally means that wellbeing is being employed in a rational calculus of this sort. There are also non-numerical versions of the tendency to model behaviour “as if” it were rational. This is a very strong tradition in the social sciences, based (I believe) on the confusion that in order to be rational – as observer – one needs to attribute rational calculation to the behaviour – a confusion of epistemology with ontology. The alternative, analysis in terms of actual psychological processes, seems preferable to me – this is now becoming quite a prevalent view in economics, due to the rise of experimental economics, in which actual behavioural characteristics are regarded as important (although it will probably be a long time before the traditional modelling approach dies away).

Talking of economics, John you said none contest the fairness of “willingness to pay”. While I agree with some of the other things you say, especially to be suspicious of the word “we”, I want to challenge that. Think about Brazil or South Africa: the situation of a poor person in a barrio or township contrasted with that of a rich person in a highly defended villa. How much would they be willing to pay to avoid something – let’s take something that could happen to either the poor or rich person, which is itself difficult to think of – a cloud of highly toxic fumes coming from a chemicals plant. One could measure this, at least for the property owner, in terms of property prices, so there is not a severe measurement issue (which there usually is for willingness to pay – but I digress). I think it likely that the rich person would be willing to pay a lot more than the poor person! – and the same applies in any unequal society, although to a lesser extent than in the two countries I singled out – i.e. in almost any society. In fact this is one of the reasons why “we” would be an improper conceptual basis for thinking about an issue like this.

Coming back to the original problem, I think we need to treat both Bentham and Pareto with suspicion. The moral question needs to be phrased in terms of subjective behaviour that is understandable in intuitive terms, in the light of the situation and the information available to participants. This involves psychological knowledge, and can be illuminated by psychology, but fundamentally the gold standard for judgement is the complex understanding of human behaviour that we all have – although again “we” is misleading, as we all have a slightly differ-

ent understanding. So when we are asked why we might be uncomfortable with a Benthamite calculation, we are asked to juxtapose the two methodologies, and my personal view is that the discomfort should take precedence over the calculation. It seems from an earlier contribution that the courts (at least in England) take this view – the difficulties that the use of statistics have sometimes given rise to would seem to justify that it may be dangerous to depart from it.

At the beginning we were asked, if Bentham is wrong, why is he wrong? – I can't answer that, but I do think that he is leading us in the wrong direction. Maybe it would be possible to compare instances where some form of rational-calculation imputation, e.g. using probabilities, is used with instances where a more traditional intuitive judgement method is relied on, and see whether one or the other produced the better outcome – which would require some external criterion of the goodness of each outcome. In the legal context, miscarriages of justice seems a candidate – but this begs a lot of questions.

2.13 Martin Sewell

wrote (Monday, August 29, 2005 3:45:06 pm)

At 11:58 20/08/2005 +0100, Tony Gardner-Medwin wrote:

>[...]
 >What value of P (the probability that X was a bomber) would justify Y's
 >decision to kill X? Should the jury convict on the grounds that Y should
 >only have killed X if he was "sure" that X was a bomber (e.g. P>67%, 90%
 >or 99%)? Or should it acquit on the basis that maximisation of expected
 >utility (minimisation of the expected no. of innocent lives to be lost)
 >was satisfied by P>16%? **
 >
 >Bentham's defence of Y rests on the grounds that it was reasonable for Y
 >(and/or his colleagues) to estimate P as >16%, even though none felt the
 >evidence rendered P close to 100%. Are we happy with this? If not, why not?

At 19:05 23/08/2005 +0100, David Colquhoun wrote:

>[...]
 >Moral. Bayesian approaches (to problems of this sort) have their place,
 >but are best restricted to consenting adults in private.

At 12:00 29/08/2005 +0100, Joffe, Michael wrote:

>I have been following this discussion with a lot of interest. I think the
 >main question is whether we should analyse the situation in terms of some
 >model that takes rationality as its basis, or whether it should be in
 >terms of actual psychological processes. The use of probabilities is just
 >one version of the former, but there are others, and the term utility
 >generally means that wellbeing is being employed in a rational calculus of
 >this sort. There are also non-numerical versions of the tendency to model
 >behaviour "as if" it were rational. This is a very strong tradition in the
 >social sciences, based (I believe) on the confusion that in order to be
 >rational - as observer - one needs to attribute rational calculation to
 >the behaviour - a confusion of epistemology with ontology. The
 >alternative, analysis in terms of actual psychological processes, seems
 >preferable to me - this is now becoming quite a prevalent view in
 >economics, due to the rise of experimental economics, in which actual
 >behavioural characteristics are regarded as important (although it will
 >probably be a long time before the traditional modelling approach dies away).
 >
 >[...]

Thomas Bayes is not on trial here, probability is the **only** consistent way of dealing with uncertainty.

The question is a legal one. What utility function should be employed by the law of the land and thus upheld by the policeman?

From an evolutionary perspective, kin selection dictates that one should value one's own life as equal to that of two brothers or eight cousins (and we should care little about strangers).

Regards

Martin

2.14 John Adams

wrote (Tuesday, August 30, 2005 11:03:40 am)

Michael says

“At the beginning we were asked, if Bentham is wrong, why is he wrong? – I can't answer that, but I do think that he is leading us in the wrong direction.”

I agree that he leads in the wrong direction, because he pursued the greatest happiness for the greatest number before it became widely recognized that this objective collided with the impossibility of making meaningful interpersonal comparisons of utility (or happiness).

I have more time for Pareto, who can be used to demonstrate the futility of pursuing objective decision-making formulas that will produce right (rational?) decisions. Cost-benefit analysts label a “Pareto Improvement” as a change that leaves at least one person better off while leaving no one worse off. According to this definition a change that left one millionaire better off while leaving no one else worse off would be an improvement.

But playing the game by Pareto's rules requires the beneficiary (the millionaire) to value the benefits in terms of willingness-to-pay (WTP) and the loser (all the poor people) in terms of willingness-to-accept-compensation (WTA). Thus, to address Michael's objection to WTP being unfair because rich people can afford to pay more, poor people are empowered to demand more compensation. It is the person suffering the loss who must be the judge of how much compensation would leave them feeling as well off after the hypothesized “improvement” as before – not the person benefiting from the change. Under Pareto's rules the poor losers can demand compensation that would restore the prior relative income differential or anything else that will leave them feeling as well off after the change as before.

It is the justice of these rules that I said none contest. The practice is quite different. No effective way has been discovered to test the veracity of compensation claims. If someone says “no amount of money will fully compensate me” – and means it – that answer, under the rules, more than offsets all the benefits that people are willing to pay for. It is this practical difficulty that routinely leads cost-benefit analysts to substitute WTP for WTA and this, as Michael notes, is grossly unfair.

In building new roads governments routinely override those who refuse to sell by means of compulsory purchase. Imposing a shoot-to-kill-suspected-suicide-bombers policy, overriding the fears of those who feel threatened by it, would be a policing equivalent. In both cases the decision will be political, and indefensible by any formal calculus of human welfare of which I am aware.

More on this theme on my website at <http://tinyurl.com/b4ljz>¹.

Prof. John Adams

¹Shortened from the original hyperlink for this document.

2.15 Michael Joffe

wrote (Tuesday, August 30, 2005 6:54:07 pm)

I agree that Pareto's logic, if taken on its own terms, has elegant properties. Though I don't think that even within the theory it escapes the point I made: this can be expressed as "the marginal utility of an extra \$100 is greater for a poor person than a rich person". If this were incorporated into the calculation of utilities, I suspect it would affect the conclusions – because a poor person would find compensation of, say, \$20,000 to be quite a lot, whereas a well-off person might find this a rather paltry amount. The poor person will (on average) systematically demand less compensation than the rich person. This point about differing marginal utilities was recognised by Alfred Marshall in the 1870s, by Joan Robinson about 60 years later, and again in recent years by Richard Layard – but in each case they think of it only in relation to policy, not to the core theory of utility, which is odd.

But the deeper point is the one you make, John, that even if we accepted that the theory was not contested, it falls down in the application – as you put it, the veracity of compensation claims is problematic, as well as the specific case of infinity that you mentioned originally. I say it is deeper because the same thing applies whenever we try and model actual human behaviour by attributing rationality. So while I agree with Martin that the Bayesian approach could give a satisfactory theoretical solution, there will always be a breakdown when we try and attribute specific probabilities. The reason is that the probability, or in Pareto's case the calculation of utility, does not correspond with a real psychological entity – it is grafted on "as if" it applied.

Martin, I'm glad you mentioned evolutionary psychology, as it makes this point about "real psychology" more clearly. In fact there are two evolutionary explanations of altruism. One as you say is kin selection: it explains why people are loyal to their family members, and possibly can be extended to racial loyalty as well. But it does not explain the actually observed altruism that empirically exists. The second explanation is reciprocal altruism, which is well expressed in the strategy of "tit-for-tat", and this turns out to correspond well with actual behaviour as well as to have good properties in game theory terms. The general lesson here is that strategies have evolved, based usually on quite simple rules – sometimes they are called heuristics – that help people deal with the complexities of actual situations without doing complicated calculations.

So returning to the suspected bomber. As Tony says, the number of potential victims is clearly important (and in view of the first paragraph, I need to add that we should impute a standard set of utilities to each life, rather than let everyone choose how valuable their family member's life was to them). You also pointed out that peace compared with war is routinely taken into account – personally I

can find no logical justification for that – and so are other things. The number of people killed and injured in the London bombings was equivalent to that in similar incidents in Iraq about twice a week, and also that on the roads of Britain every week (despite our internationally excellent road safety record). The latter is presumably tolerated partly because the deaths are not intentionally caused. Of these three, only the London bombings have made any impact on people's consciousness. War/peace, distance and intention are factors, and so is routineness versus unexpectedness/newness. For me, in judging a policeman faced by a potential bomber, it comes down to what perception he could have had of the situation – what in law is I believe referred to as what “a reasonable person” would do – in terms of actual psychological processes. Then the question remains, should these other factors (war/peace, intention) be taken into account? Would the policeman be judged the same if he were preventing an accidental explosion by killing someone? Would it be different in a war situation? Why?

Mike

2.16 Martin Sewell

wrote (Tuesday, August 30, 2005 9:46:46 pm)

At 11:58 20/08/2005 +0100, Tony Gardner-Medwin wrote:

>[...]

>

>** Simplifying the maths slightly: with $P=16\%$, killing X would kill an
>innocent person with $P=0.84$, while not killing X would kill 6 innocent
>people with $P=0.14$, an expectation of 0.84 lost innocent lives.

If we treat the loss of innocent lives as a personal loss, we can apply Cumulative Prospect Theory (Tversky and Kahneman, 1992) which tells us that, in practice, given the above probabilities, one should prefer to kill X.

Killing X gives a Certainty Equivalent of -0.6744 Not killing X gives a Certainty Equivalent of -1.0076

Approx. $P=10\%$ should provide the greatest dilemma.

Regards

Martin

2.17 Martin Sewell

wrote (Wednesday, August 31, 2005 10:15:17 pm)

At 18:54 30/08/2005 +0100, Joffe, Michael wrote:

>[...] In fact there are two evolutionary explanations of altruism. One as
>you say is kin selection: it explains why people are loyal to their family
>members, and possibly can be extended to racial loyalty as well. [...]

Whilst it is true that kin selection dictates that one should behave more altruistically towards those genetically similar to ourselves, sexual selection dictates that one should avoid mating with an individual with a high probability of being genetically similar (harmful recessive genes would have deleterious effects).

If we extend Hamiltonian kin selection from family to race, should we not also extend incest avoidance from family to race?

However, empirical evidence shows us that when seeking partners, women in particular, say that they prefer to meet someone of the same ethnic background as themselves (Hitsch, Hortacsu and Ariely, 2004).

Regards
Martin

2.18 John Adams

wrote (Wednesday, September 7, 2005 12:31:47 pm)

I've been away – so a delayed response to the contributions of T G-M and MJ on 30 August.

Risk-benefit analysis is a version of cost-benefit analysis in which estimates of cost and benefit are multiplied by the probabilities that frequentists (or Bayesians) can attach to them. If all lives are of equal value, then numbers of lives at risk X probabilities ought to settle all arguments.

The principal difficulty with cost/risk-benefit analysis lies not with the problem of assigning probabilities - challenging though that often is - but with attaching values to the costs and benefits. As I noted in an earlier email, the requirement to measure benefits in terms of WTP and costs in terms of WTA invalidates the assumption of all lives being of equal value.

I accept Michael's point that the marginal utility of a given sum of money is likely to be greater for a poor person than a rich person, but there is a larger valuation problem that this point does not address. "We" do not share a common objective function. Consider a proposal to drain a scruffy wetland and convert it into a golf course. To the non-golfing twitcher/environmentalist the scruffy wetland is natural habitat beyond price. Or, another example from the real world, a proposal to build a new road to relieve congestion. The time-saving benefit releases suppressed demand and increases traffic. This extra traffic, treated as a benefit in most CBAs, is seen as a cost by environmentalists concerned with resource depletion, pollution, and growing dependence on the car.

Whenever there is a conflict of values, CBA or RBA can settle no arguments. The beauty of Tony's hypothetical example is that it embodies the most extreme value conflict imaginable. One person's freedom fighter (good) is another person's terrorist (bad), with the added complication that even the failed suicide bomber gains paradise (good). A macabre example of win-win?

Prof. John Adams

2.19 Tony Gardner-Medwin

wrote (Friday, September 9, 2005 1:40:44 pm)

Sent on behalf of John Adams:

(Tony forwarded the above message (p. 26) and added a note)

Dear John, ...I'd be interested to see your response to Michael Joffe's worries about Pareto. Surely money doesn't account in the same way for happiness or distress for different people. Ultimately I think we have to deal in some sort of utilities about what sort of world people want to live in, to balance against the immediate ones of saving/ destroying lives..

All the best,
Tony

2.20 Tony Gardner-Medwin

wrote (Sunday, September 11, 2005 1:34:20 pm)

My response to some of the recent issues raised -

The evidence list seems to have recovered from a technical problem that led to hold-up and duplication of emails. At the end of last month the Bentham defence puzzle seemed to me side-tracked to rather minor side issues: (a) how you might personally weigh the 'cost' of the deaths of different innocent people, e.g. those more or less prepared to pay to live, or who share your own genes to marginally different extents. Not really of the essence. (b) the idea that one should judge a decision maker on the basis of Kahneman & Tversky's studies of how people behave somewhat non-optimally and irrationally when they compare + and - utilities, and exaggerate small probabilities in economic decisions. Neither issue has great impact compared with the main direct utility ratio involved - the ratio of the number of people who would be killed by a bomb blast to the death of a single possibly innocent person killed by police.

John Adams has suggested that the diversity of different people's views of the utilities may involve huge and imponderable value conflicts, based on religious or political beliefs. Again, I think this is a bit of a red herring. The decision at issue (Y's decision whether to shoot X) is an ethical one based on the weighing up of potential loss of life and ideas about how security forces should act. The fact that X and his cronies may think he would right a great wrong or encounter virgins in heaven surely doesn't come into it. Neither the law, nor a jury in this country, nor surely any of us, would say X's belief in a great utility was a reason to allow him to bomb innocent people. The same goes for the argument that significant sections of the population might support a politically motivated bombing. These may be relevant utilities for X if he is a bomber, but not for Y.

I've tried to list below the possible outcomes of a decision not to shoot (A1-4) or to shoot (B1-6). Every outcome is bad or at best neutral - except for one (B5) which I suspect most of us fear was a utility that carried excessive weight in events on 22/7. Bentham's defence is that the balance in this lose-lose situation must be largely based on the immediate utilities (principally A1, B1) in which he (and I) can see no reason for not weighing all lost innocent lives equally. But I would add that it must also be based partly on the essentially political utilities at A4 and B4. I find it very disturbing that Y as a member of a security force should have to weigh these immediate and global utilities one against the other and be held responsible for a decision that in Ian Dennis' words is obviously "context-dependent and also value-laden". Poor Y may have a few seconds to weigh up an issue that a jury may take weeks to ponder in the national interest. Unless we can somehow clarify how Y should behave in a manner that we all understand and agree on and that

would be clear to a jury, I go along with Bentham's defence and acquit Y even if, based on reasonable inference from the evidence available to him, he killed X while being far from certain that X was guilty.

Of course I don't think policemen while chasing suspects should be doing numerical Bayesian inference calculations. But policemen, like cricketers and all the rest of us, run their lives by balancing risks, benefits and probabilities. A batsman who tries to hit a six has 0.5s to weigh the probability that this will cost his wicket against the utility of 6 extra runs. Not a single number enters his head (except 6, I guess!) but he weighs it up nevertheless, and can be held to account by his captain. We rely a lot on our policemen being good at cricket. Though psychologists and economists tell us that people take shortcuts and make errors in such inferences - undoubtedly true - they are the essence of neural function, and good brains do them well. It is I think the job of the courts and people such as ourselves interested in inference, to establish a standard for what 'well' means, with as much rigour as possible.

TABLE OF POSSIBLE OUTCOMES:

If I do not kill X, then one or more of these outcomes may transpire:

A1: X may detonate a bomb and kill many bystanders, probably including me.

A2: X may prove to be innocent and suffer only from the trauma of the chase

A3: X may prove to have a bomb and may fail to detonate it or refrain from doing so, and be arrested

A4: if A1, and possibly if A3, I and the security forces may be subject to condemnation for failing to take action to prevent a disaster

If I shoot X dead, then one or more of these outcomes may transpire:

B1: I will prove to have killed an innocent person or one of no immediate threat to anyone

B2: I will prove to have killed a person armed and likely to have been of immediate threat to others

B3: I may (probably posthumously) prove to have contributed to the detonation of a bomb by my action

B4: if B1, I and my force will be widely condemned for shooting to kill without warning and with insufficient or unreliable evidence, and I may lose my job and possibly go to jail.

B5: if B2, I and my force will be widely praised and I will be regarded by many people as a hero

B6: I will probably have nightmares about this event for the rest of my life

Tony GM

2.21 Martin Sewell

wrote (Sunday, September 11, 2005 9:06:44 pm)

At 11:58 20/08/2005 +0100, Tony Gardner-Medwin wrote:

>[...]

>

>What value of P (the probability that X was a bomber) would justify Y's
>decision to kill X? Should the jury convict on the grounds that Y should
>only have killed X if he was "sure" that X was a bomber (e.g. P>67%, 90%
>or 99%)? Or should it acquit on the basis that maximisation of expected
>utility (minimisation of the expected no. of innocent lives to be lost)
>was satisfied by P>16%? **

>

>Bentham's defence of Y rests on the grounds that it was reasonable for Y
>(and/or his colleagues) to estimate P as >16%, even though none felt the
>evidence rendered P close to 100%. Are we happy with this? If not, why not?

>

>[...]

Bentham's defence of killing Y if and only if P >16% is derived in a straight-forward and logical manner, yet (reflecting expected reactions) the question is phrased in such a way as to concede that "16% feels too low". I believe that this is due to various "anchors":

1) It is axiomatic in law that letting a guilty man go free is preferable to convicting an innocent man. 2) Criminal law would require P approx. 100% to convict X. 3) Civil law would require P >50% for a guilty verdict. 4) The real case that the hypothetical example is based upon resulted in an incorrect decision to kill when the probability that the victim was a bomber should have been close to zero.

Regards

Martin

2.22 Bilal Barakat

wrote (Monday, September 12, 2005 4:18:10 pm)

I believe some of the suggested ways of reasoning about this are ethically disastrous, which is obvious if they are applied in other contexts.

Consider the following scenario: Terrorists threaten to kill 6 hostages 10 minutes from now (so attempts at a rescue operation are out of the question) unless a certain Mr. Smith is shot dead by the police. Is anyone who favours shooting suspects willing to suggest compliance with the demand is a reasonable thing to do? The numbers are the same, after all.

The objection that in this case it is certain that the victims will be innocent, in contrast to the original example, is not compelling. After all, what is up for debate is not a singular decision but a strategy. If it is thought reasonable to shoot suspects with a probability greater 16% of being about to kill 6 bystanders with a bomb, then if this is applied consistently it is likewise certain that the vast majority of people shot under this rule will be innocents, by definition. If there might be 6 such decisions over the year, and the example above is modified to 36 hostages and the whole Smith family (of 5) being at stake, we'd expect to kill just as many innocent people by application of the "shoot without warning" order as we would by giving in to the hostage-takers. To spice up the example, imagine it was the prime minister rather than Mr. Smith whose head they'd want.

In any case, the suggestion that all potential victims should carry the same weight in determining the decision rule is absurd. If that were the case, there should be constellations where we are **indifferent** as to whether the police let terrorists kill x innocent people or instead the police kill x (or maybe $x-1$) innocents to prevent this. Is anyone seriously suggesting that it should not matter whether 5 or 10 or 100 random people are killed by terrorists or by the police, as long as the numbers are the same?

I think that even if there were 100 or even 1000 hostages, Mr. Smith would have a very good case for arguing that the police have no right at all to sacrifice him (unless he volunteers selflessly).

Regards, Bilal.

2.23 Tony Gardner-Medwin

wrote (Monday, September 12, 2005 7:34:04 pm)

Once more to the fray,

Bilal Barakat cogently points us at blackmail situations. But a new dimension is introduced here, and a new utility. Police or public policy to shoot innocent persons to comply with blackmail, because of a balance of immediate utilities, is usually seen as likely to encourage future blackmail or hostage taking, with a potentially overwhelming (what I call global or indirect or political) negative utility. This has of course many times been the rationale for not complying with the demands of hostage takers, even when the demand has been less drastic.

Bilal also suggests that the shooting by police of a person who is certainly innocent would be a factor making us regard this act as ethically abhorrent. I find this a more interesting challenge. One can isolate it from the hostage/blackmail issue by considering for example whether one should kill a person Z who is innocently about to do something that will trigger a disaster, or dire criminal act, or an accident (which I think someone brought up earlier). It is harder to be categorical in condemning this, and there must be many instances in war when armies have acted in such a situation and killed men of their own forces to prevent them unwittingly giving away information that could trigger military disaster. One might dream up possibilities also in civilian life. Nobody is comfortable about the ensuing court martial or trial, but I don't feel the decision should be a foregone conclusion on ethical grounds. If indirect utilities are weak, I don't think ethical decisions are too unhappily based on immediate utilities.

This brings me back to my regular mantra - that probably nobody is happy with Bentham's defence on its own. Immediate utilities have to be set against the global and political issues and utilities. As Bilal nicely puts it (more or less): we would much prefer to live in a London with 52 innocent people/ann killed by terrorists than with 51 innocent people/ann killed by hair-trigger police suspecting they were terrorists. This is fine, but how do we help the armed policeman Y to know how to balance the strength of this argument against the realities of his immediate situation? If we expect him to kill a suspect only if sure of intent and capability to the standard of a criminal court, then why do we ask courts, lawyers and juries to deliberate for days or weeks over such a conclusion when he certainly can't afford to? The only hope I can see to present meaningful guidance for someone in the heat of the moment would be to set some arbitrary heuristic, not particularly justified by logic - e.g. regardless of circumstances, would you (Y or someone giving Y an order) accept automatic dismissal if you shoot this suspect and he proves not to have been equipped to kill, and a month's bonus if he does prove to have been equipped to kill? I'm sure people can pick holes in that as a

suggestion - especially of course because it leaves out the issue of intention, and makes the critical P value dependant on the policeman's personal circumstances - but it seems better than nothing. Until you write something into the rule book, you can only expect police to act on their perception of the obvious utilities and uncertainties and it seems to me you should acquit them if, with reasonable diligence, they do so.

Tony GM

2.24 Christian Hennig

wrote (Monday, September 12, 2005 8:37:29 pm)

I think that an important aspect is missing in the discussion up to now. The ethical considerations condemning the killing of innocents by the police have been treated until now as if they are not linked to any important utility considerations.

But it may well be that there are some people (especially adolescents) who can be driven into the terrorist's camp. They may start to consider the police and the society as their enemy because of the police's killing of innocent people with which they feel united, and this process may cost the lives of many more people later on.

It may also be important to consider which effects a policy like "kill a suspect above $x\%$ of subjective probability of carrying a bomb" may have on the behaviour and thoughts of the policeman. In the long run this may encourage an attitude which also costs more innocent lives and damages our democracy seriously.

On the other hand, it could also be argued that killing some suspects from time to time is repulsive to the terrorists even if the killed people are innocent.

Could all this be formalized in terms of probability and utility? Hardly. Especially, all the mentioned long term effects will heavily depend on future developments, policies and reactions of many people on future events, and how can such a complex and self-referencing process be formalized with any reliability?

For the individual policeman in a single situation, this can only mean "garbage in, garbage out", as was argued before. It makes more sense to talk about a policy, which is a decision rule rather than a single yes-no-decision and such a policy can be a bit more complex than relying on a one-dimensional numerical assessment.

My viewpoint is that I would trust the non-formal judgement of a well trained policeman with a good deal of common sense much more than any ad-hoc computation of subjective probabilities. Formalization tends to concentrate on some isolated aspects that are more or less easy to formalize. A less formalized decision means more space for ethics and intuition, and even though it is not well understood how this exactly works, I would presume that it operates from a more holistic, broader perspective. I hope that the aspects above illustrate why it may be good and useful to shrink back from shooting a suspect because of ethical considerations even as long as there is no clue how to work out the numerical utility impact of these considerations.

Thank you for a very interesting discussion,
Christian Hennig

2.25 Bilal Barakat

wrote (Monday, September 12, 2005 9:44:28 pm)

>Once more to the fray,
>
>Bilal Barakat cogently points us at blackmail situations. But a new
>dimension is introduced here, and a new utility. Police or public
>policy to shoot innocent persons to comply with blackmail, because of
>a balance of immediate utilities, is usually seen as likely to
>encourage future blackmail or hostage taking, with a potentially
>overwhelming (what I call global or indirect or political) negative
>utility. This has of course many times been the rationale for not
>complying with the demands of hostage takers, even when the demand has
>been less drastic.
...

I for one do not believe the reason we wouldn't comply with the hostage takers' demands is the concern for encouraging future repetitions of the situation. Consider another situation: correct me if I'm wrong, but I'm sure that fire-fighters are not required to sacrifice their lives to an in an attempt to save people, regardless of whether a larger number of deaths could be avoided by such a sacrifice. If we do not ask this of professionals who took the conscious decision of taking reasonable risks on themselves, how much less can we ask this of random civilians.

Of course in war soldiers can be ordered on what might effectively be a suicide mission for the benefit of a larger number, and people are killed to prevent the presumed greater harm of military disaster. However, until London is officially under martial law, we are dealing explicitly with a civilian context, and arguments about the necessities of war do not apply.

On the point of the policeman having to make a decision in a split-second: surely the fact that there is no time to carefully weigh the consequences and probabilities involved implies that there is all the more reason to put priority on not being **directly** responsible for harm, and the heuristic should therefore be conservative. This holds even if the policeman's life is at stake too; if someone has to, it is clearly the security professional who should tolerate the risk (of being too cautious) rather than placing it on the unsuspecting public (that of falling victim to a less cautious approach).

But I think this is all moving rather far away from questions of evidence. Of course the above could be reformulated: instead of asking whether to err on the side of caution one might ask whether under pressure one should be more exacting of evidence or less so. Is anything gained by framing it in these terms? What does

the concept and field of "evidence" have to offer to the discussion of this dilemma?
These are, I think, the evidence-related questions lurking here.

Regards, Bilal.

2.26 Martin Sewell

wrote (Monday, September 12, 2005 10:32:11 pm)

At 11:58 20/08/2005 +0100, Tony Gardner-Medwin wrote:

- >Londoners and evidencers,
- >Y is on trial for murdering X. Y estimated at the time (based on evidence
- >and inferences available to him) that there was a probability P (his
- >subjective degree of belief) that X was an armed suicide bomber, [...]

Y has reason to believe that three suicide bombers of Somali or Ethiopian appearance reside in a three-storey block of flats (containing 9 separate flats).

X emerges from the communal entrance of the block whilst Y is relieving himself.

What value should Y assign to P?

Y follows X. X is of Brazilian appearance, unaware of Y and behaving normally. He is wearing a pair of jeans and a denim jacket and carries no bag.

Now what value should Y assign to P?

Regards

Martin

2.27 Tony Gardner-Medwin

wrote (Wednesday, September 14, 2005 9:37:56 am)

I'm not quite sure why DC (below) is laying into me when he also seems to be saying much the same thing: "Everything depends on what the rule book says." But his wording of the rule book "you must have seen the bomb" is not surely going to help much. Y in the witness box may quite truthfully say "I believed I saw a bomb". On cross-questioning doubtless he may concede "No, I suppose I wasn't absolutely certain I saw a bomb, and with hindsight I can see that obviously I didn't see a bomb. But at the time I was say 90% sure it was a bomb." My proposal for the rule book ("You will automatically lose your job if there is no bomb, etc..") at least has clearer meaning for the policeman.

The reason this debate is actually quite profound, I think, not just post-22/7 pub chat, is the eternal one as to the status of subjective probabilities or degrees of belief. Some anti-Bayesians challenge the very meaning of the concept. But I think we can't get away from the fact that (a) there is no self-consistent theoretical framework for making decisions that doesn't involve them, and (b) they are fundamental to how people act, and (c) they are fundamental to the processes of law. Martin Sewell challenged Bayesians yesterday to say how if you knew a few of the alleged facts from 22/7 you would arrive at a Bayesian P. It seems to me he answers his own question, by what I think he is implying by his selection of facts - which is that the P based on them is at one and the same time obviously unquantifiable in any precise way, but also obviously very low. My point is that a person acting in a reasonable manner in possession of just these facts would (or should) not have had a strong belief that X was a bomber. The way a Bayesian would establish what his/her degree of belief was would be to offer them contingent bets at variable odds, or payoffs. This is roughly what I was trying to do by making it wholly explicit in the rule book that if wrong you lose your job, while if right you only gain a months pay. It's sordid to bring in such trivia as losing your job when lives are at stake, but it helps to define what degree of belief means.

Tony GM At 22:11 13/09/2005, you wrote:

- >Tony Gardner-Medwin concludes
- >- "Until you write something into the rule book, you can only
- >expect police to act on their perception of the obvious utilities
- >and uncertainties and it seems to me you should acquit them if,
- >with reasonable diligence, they do so" This statement seems to me
- >to be ethically and politically outrageous. It is also quite
- >unnecessary, because presumably the police are given a rule book
- >before being out into a position where they have to decide.
- >Everything depends on what the rule book says, and it is hard to

>say much about that because it has been kept secret. The secrecy
 >is quite wrong because the innocent tube traveller takes the risk
 >of being shot without knowing the rules of engagement. Presumably
 >in the real case, either the rule book was not followed, or the
 >rules in it are quite unacceptably lax.

>

>The rule book should, in my opinion, say that the policemen has to
 >see the bomb before he can shoot, just as it says (or used to say)
 >that the policemen has to see a gun (not a chair leg in a bag)
 >before he can shoot. But does it? We have no idea.

>

>There are many intangibles that have not been considered at all in
 >these arguments, not the least of which is the personality of the
 >sort of person who volunteers to be a police gunman. Perhaps some
 >of them are brave altruists, doing their duty to their country. On
 >the other hand, perhaps some of them just fancy the chance to do a
 >bit of shooting (you have only to look at the personalities of the
 >people convicted of the Abu Ghraib abuses to see that suggestion is
 >only too plausible). Even if the probabilities in Tony's original
 >problem could be given sensible values (which I don't believe to be
 >>true), the calculations miss out so many relevant factors that it
 >is no wonder that they produce a silly conclusion.

>

>The only way to deal with the problem is to give the policeman very
 >clear rules. The rule book should, in my opinion, say that the
 >policemen has to see the bomb before he can shoot, just as it says
 >(or used to say) that the policemen has to see a gun (not a chair
 >leg in a bag) before he can shoot. But what do the rules say? We
 >have no idea at all.

>

>One of the nicest ways of putting it so far is "we would much
 >prefer to live in a London with 52 innocent people killed by
 >terrorists than with 51 innocent people killed by hair-trigger
 >police suspecting they were terrorists." But by keeping the rules
 >secret the police (or the home office) have denied the public, and
 >indeed denied parliament, any chance to know what policy is being
 >implemented on our behalf.

>

>David Colquhoun

2.28 Martin Sewell

wrote (Wednesday, September 14, 2005 4:20:29 pm)

At 09:37 14/09/2005 +0100, Tony Gardner-Medwin wrote:

>[...]

>Martin Sewell challenged Bayesians yesterday to say how if you knew

>a few of the alleged facts from 22/7 you would arrive at a Bayesian

>P. It seems to me he answers his own question, by what I think he

>is implying by his selection of facts - which is that the P based on

>them is at one and the same time obviously unquantifiable in any

>precise way, but also obviously very low.

>[...]

To clarify...

How people *should* behave:

1. a) Cox, de Finetti, etc. showed that the only consistent way of dealing with uncertainty is via probability theory;

b) ergo, everyone should be a Bayesian.

Explanation of Y's behaviour:

1. a) Probability in its current form has only been with us since the eighteenth century so our brains have not evolved to think in probabilistic terms,

b) ergo, we are poor Bayesians.

2. In predicting the future, people tend to get anchored by salient past events. Consequently, they underreact to new information and fail to update their degree of belief (probability).

3. Prospect theory shows us that people tend to make risk-seeking choices to avoid negative outcomes.

4. People, especially men, tend to be overconfident.

Regards

Martin

2.29 Christian Hennig

wrote (Wednesday, September 14, 2005 4:30:37 pm)

Dear Martin,

>How people **should** behave:

>1. a) Cox, de Finetti, etc. showed that the only consistent way of

>dealing with uncertainty is via probability theory;

>b) ergo, everyone should be a Bayesian.

it seems I missed something. How has it been shown that (Bayesian) probability theory is the **only** consistent way, i.e., that no other consistent way is possible? (This is different from showing that **certain** other ways are not consistent.)

And: coherent Bayesian reasoning implies that all possible future events are specified a priori. Do we want such a restriction, be it consistent or not?

Best, Christian Hennig

2.30 Martin Sewell

wrote (Thursday, September 15, 2005 2:19:49 am)

At 16:30 14/09/2005 +0100, Christian Hennig wrote:

>Dear Martin,

>

>>How people **should** behave:

>>1. a) Cox, de Finetti, etc. showed that the only consistent way of

>>dealing with uncertainty is via probability theory;

>>b) ergo, everyone should be a Bayesian.

>

>it seems I missed something.

>How has it been shown that (Bayesian) probability theory is the **only**

>consistent way, i.e., that no other consistent way is possible? (This is

>different from showing that **certain** other ways are not consistent.)

R. T. Cox (1946, 1961) showed that the probability axioms attributed to Kolmogorov constitute the only consistent extension of ordinary (Aristotelian) logic in which degrees of belief are represented by real numbers.

>And: coherent Bayesian reasoning implies that all possible future events

>are specified a priori. Do we want such a restriction, be it consistent or

>not?

One can specify events thus: event 1, event 2, ..., event n, "any other".

Regards

Martin

2.31 Christian Hennig

wrote (Thursday, September 15, 2005 10:21:40 am)

>>>How people *should* behave:

>>>1. a) Cox, de Finetti, etc. showed that the only consistent way of

>>>dealing with uncertainty is via probability theory;

>>>b) ergo, everyone should be a Bayesian.

>>

>>it seems I missed something.

>>How has it been shown that (Bayesian) probability theory is the *only*

>>consistent way, i.e., that no other consistent way is possible? (This is

>>different from showing that *certain* other ways are not consistent.)

>

>R. T. Cox (1946, 1961) showed that the probability axioms attributed

>to Kolmogorov constitute the only consistent extension of ordinary

>(Aristotelian) logic in which degrees of belief are represented by

>real numbers.

Why is the only consistent way of dealing with uncertainty the extension of ordinary (Aristotelian) logic in which degrees of belief are represented by real numbers?

What about, for example, Walley's (1991) concept of imprecise probabilities?

Christian Hennig

2.32 Martin Sewell

wrote (Sunday, September 25, 2005 1:49:46 am)

At 10:21 15/09/2005 +0100, Christian Hennig wrote:

>>>>How people *should* behave:

>>>>1. a) Cox, de Finetti, etc. showed that the only consistent way of

>>>>dealing with uncertainty is via probability theory;

>>>>b) ergo, everyone should be a Bayesian.

>>>

>>>it seems I missed something.

>>>How has it been shown that (Bayesian) probability theory is the *only*

>>>consistent way, i.e., that no other consistent way is possible? (This is

>>>different from showing that *certain* other ways are not consistent.)

>>

>>R. T. Cox (1946, 1961) showed that the probability axioms attributed

>>to Kolmogorov constitute the only consistent extension of ordinary

>>(Aristotelian) logic in which degrees of belief are represented by

>>real numbers.

>

>Why is the only consistent way of dealing with uncertainty the

>extension of ordinary (Aristotelian) logic in which degrees of

>belief are represented by real numbers?

Higher-order probabilities, for example, have problems with semantics and utility (Kyburg (1989) and Pearl (1988)).

>What about, for example, Walley's (1991) concept of imprecise probabilities?

Although this provides us with a measure of ignorance, when it comes to decision making or prediction, we require a probability.

Regards

Martin

Index

John Adams, 14, 20, 26

Bilal Barakat, 31, 35

David Colquhoun, 4

Ian Dennis, 7

Tony Gardner-Medwin, 2, 5, 6, 27, 28,
32, 38

Christian Hennig, 34, 41, 43

Michael Joffe, 16, 22

David Lagnado, 13

Richard Leary, 9, 11

Mike Redmayne, 3

Martin Sewell, 18, 24, 25, 30, 37, 40,
42, 44