


Centre for Philosophy of Natural and Social Science**Discussion Paper Series**

DP 74/05

Transparency, Accountability and Rationality |Jeremy Hardie
LSE

Editor: Max Steuer

Transparency, Accountability and Rationality

Jeremy Hardie

Abstract

For full public accountability to work, two apparently banal but tough conditions need to be met. First, the decision must be made transparent by setting out a clear account of how it was arrived at. Second, the public must be capable of making an assessment, in the light of that account.

The first of these conditions is met when the problem has been given a structure in terms of key concepts such as objectives, facts, options. Problems which can and have been sorted out in this way are called here 'well-behaved'. The process of sorting them out is called here the operation of 'Rationality'. A person who has operated this way on a particular problem is called here the 'expert'. The second requires that the public, the 'client', be sufficiently expert to use that account to assess the decision competently.

Often accountability fails for one or other or both of two reasons. First, the problem may be intractable in the sense of not capable of being made well-behaved, for example, when you cannot in advance say what the objective really is, or what constitutes a fact relevant to the generation of options. Second, the client is insufficiently expert to assess the decision, for example, unable to see whether the judgments which the expert made in sorting the problem out, e.g. what will really constitute success, are good, because there are no and cannot be any shared set of criteria for that assessment.

If these conditions are not met, then the assessment of a decision cannot be based on the operation of Rationality, but on something else, called here 'trust'. And attempts to use full accountability when those conditions are not or are poorly met is not a second best, but counterproductive as compared with using trust.

Comments to: hardiejeremy@hotmail.com

Transparency, Rationality and Accountability

It looks like a truism to say that, for a person faced with a decision, it is best to approach the problem systematically – for example, by identifying what it is that you want to achieve, what ways there are for achieving it, and choosing the best. There are many versions of this account of deciding, some formal, some less so. The formal versions are found typically in academic and theoretical literature, and include decision theory, microeconomics, and cost benefit analysis. Members of this family share a vocabulary which includes such notions as optimization, maximization, options, trade off, objective function. The less formal versions are found in business, the military, public service and in ordinary life. Their vocabulary includes ends, means, objectives, policy, execution. The differences between these various ways of thinking about how to decide, or what to do, can be important for certain analytic and practical purposes. But for this paper I will lump them all together under the word Rationality.

This term is here meant to refer to nothing more than the systematic and well-behaved, in a sense which I hope to make clear later, features of these techniques for deciding. The main purpose of this paper is to show that the ambition, widespread today on public life and elsewhere, that decisions should be arrived at transparently and hence accountably, relies for its realism very heavily on the hopes that a Rational account of decisions can typically be provided, meaning again by Rational no more than what is suggested above: and that if it is available, communicating it openly to the public will achieve the purposes of accountability. Where those hopes are not fulfilled, and they won't be in many cases, accountability through transparency cannot be successful. If, as it is, a major aim of accountability and of transparency is to create public confidence in decisions made by others (such as the Government) on their behalf, then confidence cannot always, and perhaps typically, be created in that way alone.

If we are in a Rational world, the decision maker (still at this stage an individual) is, or at least can get to be, on top of the problem. He is not bewildered, or not bewildered for long. It is central to the required notion of Rationality that I can, for example, identify my desires, or wishes, or objectives, or passions: and that I can inspect the world, see what the facts are, see what facts are relevant in what way to my desires, or wishes, or objectives or passions: generate options: say something that is clear enough about various trade offs between objectives, and, if they matter to me, between routes to those objectives: and so on. A great deal of work has been done, for example in the decision-making literature, in economics, and in the literature of practical reason, to impart rigour to the imprecise members of this unstructured list. An example is, in public policy, the literature on cost-benefit analysis. But here only two points need to be made. First, it is not interesting for the present purpose to snipe at Rationality by

pointing out unhelpfully that facts are doubtful, probabilities uncertain, trade offs unclear, time limited, and so on. What follows does not depend on arguments such as that full Rationality is unachievable and people do not decide like that. All that is needed here is that Rationality in its various manifestations exhibit sufficiently, not perfectly, its familiar features, that deciding is a matter of taking hold of a problem, structuring it systematically by organizing its components under categories (such as objectives, facts, options), and generating a result by the application of predetermined rules. A problem which is organized in that way is well-behaved in my sense, and I hope that it is plain enough what is meant. The second point is that few problems present themselves as well-behaved. The main task of this kind of decision-making is to get a problem into the shape needed to make it possible to make a decision. The decision maker must be competent, not bewildered.

Thus far, the account has been in terms of an individual decision maker, deciding on his own, by getting the problem into shape for himself, and by himself. There need be nobody else around, and the notion of accountability does not arise. Nor does the problem of understanding why someone else has decided as he did. But if I do decide in a Rational way, say via cost benefit analysis, accountability and transparency come into the picture very easily, when other players appear. This is because if I am competent, and have succeeded in getting the problem into shape, and then decided, I can explain my decision afterwards by sharing an account publicly of what I have done and why. In particular cases, it may be complicated, but it is of the essence of Rationality that I can say what my objective was, what the options were for achieving it, why I preferred this option to that. In cost-benefit analysis, this is done most powerfully by reducing everything to the common yardstick of money, just as in business the company explains what it does by showing how it will lead (or typically should have led) to maximum profit (or so the story goes).

As a paradigm of how accountability can and should work, consider the case of professional advice. The lawyer has to decide, or recommend for decision, what should be done. I am the client. I go to him with a complex problem, and a month later I visit him to get his advice. As I arrive, he gives it to me in a single sentence – stop fighting the case, give up. I can then, perhaps disappointedly, ask why he comes to that conclusion. As I am his client, he has to tell me, and he may hand over to me the written account of how he came to that conclusion (at the English bar, called his opinion). It is quite long and complicated, but it is Rational and the problem has been made well-behaved. I can go through it, its workings are transparent, he is accountable to me for why he has said what he says. I do not just have to accept the phrase ‘stop fighting the case’. I can go behind it and understand where it came from. I can question what he says, and we have a common structure, the Rational structure of the law, within which to debate our differences.

This account of how accountability can work is not idealized – it can and does happen – but it requires the satisfaction of a number of rigorous conditions. Up to a point, it does not matter if those conditions are not fully met – that means only that the accountability process does not work as well as it might. But in many cases, the conditions are either not met at all, or are met so poorly that accountability in those terms makes no sense at all, and insisting on accountability through transparency is not just flawed as a route to giving the client, or the public, confidence, but counterproductive i.e. it destroys confidence.

The conditions required for full strength accountability to work include something like the following:

- For the legal analysis to start, the facts have to be clarified. They are then taken as given.
- The questions are clear – e.g. does this constitute fraud?
- There is a clear way of dealing with the question – apply the law – and the law is well set out in the law books
- Thus the problem can, in our terms be made well-behaved and hence clearly set out.

It makes sense that these characteristics of clarity should help to solve problems effectively. But for accountability and transparency, clarity in this sense is not enough. What matters is, clear to whom? Everything works well if, for example:

- The client understands the law
- They share a common understanding of facts like these, what constitutes a relevant fact, and how it is relevant
- They share a common understanding of how to use the manual, the law book

Consequently:

- They know how to carry out the ex post audit, by going through the analysis using their shared and mutually understood skill
- Hence the legal analysis, what the reasons were for reaching the conclusion to give up, is clear to them
- So the process is transparent, and the lawyer is successfully accountable to the client

So it is not just that the structure of the problem and of the process be capable of clear analysis. What is needed for accountability is that the analysis be clear to both the principal and the agent. Transparency requires clarity for the relevant people, not just clarity in the abstract.

The characteristic problems which arise when the client is not sufficiently expert can be seen by varying the example a little. When the lawyer gives the crisp

advice not to fight the case, the client just accepts it. But the advice turns out to have been wrong by the simple outcome test that the case is lost, at great expense to the client. He then out of curiosity or maybe recrimination sets about unpicking the original advice to fight, by going through the lawyer's reasoning, by holding him to account. He wants a transparent account of why he lost.

Up to a point, and what follows reflects the distinction made above, the task of providing him with an answer to that question looks tractable. After all, or let us suppose that the lawyer can indeed explain where it went wrong, or at least could explain it to a legal colleague. It may be that new facts emerged, or that some of the facts turned out differently. Perhaps they now see that they made a mistake in the analysis. Perhaps none of these things went wrong, but the judge just decided the other way. That might be because he made a mistake, or because there are two ways of looking at the facts, or the analysis, and he chose the other side's. In any case, the lawyer – who in this story is honest and self aware – can talk effectively about what went wrong.

But for accountability, that is not enough. The client cannot take part as an equal in that post mortem. Or rather, the post mortem will only work for him if:

- The lawyer owns up to having made a mistake – that always works even if the client does not understand the mistake: or
- What went wrong is easy to understand –e.g. a witness central to the case changed his story in crucial respects.
- The lawyer was out to serve the client's interests in the first place, not somebody else's – like his own.

Otherwise it will not work because he cannot distinguish between

- The case when the lawyer's allegedly transparent account is self serving, either because they made a genuine mistake in their reasoning or a deliberate one to serve their own interests – they say it was just that the judge disagreed, but actually they should have argued it differently: and
- The case where their explanation is honest and true.

It is quite easy to see why the client might find it hard to hold the expert to account when he is insufficiently expert to understand the language, or the concepts, which the agent uses. It is really no use for me to be taken through the complex reasoning of a legal opinion if I do not know what the word tort means. It is not enough for me to be given the dictionary definition of tort. That will be nothing like enough to know what it means, in the sense of understanding what how the concept operates. So a legal opinion may display as much good behaviour as you like, so that the solution of the problem is well structured and one part of the argument follows inexorably from the other, but that is no good for accountability, because that structure of consistent clarity cannot be displayed

transparently to me, only to an audience which is made up of people sufficiently expert to be sufficiently the equal of the expert who is being brought to account.

The analysis leading to the answer to a well-behaved problem is easy to follow, if you are expert enough to follow it, because the reasoning is tight. One thing leads to another, the reasons support the conclusion, and once the problem has been structured properly by the expert at least in the ideal case the answer just emerges from the analysis. If you specify the objective function clearly, and know the preferences, including attitudes to risk and uncertainty, the answer cannot be otherwise. You can of course get a different answer by varying the facts or the preferences or the objective. But that new answer will have the same robust characteristics.

A more difficult case, meaning one where it is harder to give a rigorous account of what has gone wrong, is where the problem cannot, even by experts, be made well-behaved, in the sense of well structured with one thing following inexorably from the other. This arises as follows. When a problem is first identified, it does not present itself in an orderly shape. The first and maybe the most important task in deciding in accordance with

Rationality is to give the problem structure, to make it well-behaved. This requires some basic work. For a start, what would constitute success? Often our objectives are initially couched in vague and abstract terms – for instance, we want “what is best” -- and we can’t just read off whether any more concrete option is a case of this or not. Usually not even the expert can make this judgment at the start, let alone the non-expert client. It takes working through the problem to begin to get a sense of how to formulate more concrete objectives and in the working through, usually the initial desires and aims of the client aren’t enough...they are just too vague. New more concrete aims must be settled on and at this stage the initial guidance gives out – it just isn’t clear enough to dictate what the new range of aims should be. It is not clear enough for the expert to decide, nor for the client to; worse, by this stage the client often can’t understand what the options really mean.

We have just the same kinds of problems with the facts. What do we guess are the facts, and which facts do we guess are relevant to this problem? What facts are relevant will depend on the more concrete, detailed formulation of the aims. But often what are even starting options for a more detailed specification of the aims will depend on the facts. Usually we don’t have the facts and they are costly to gather, sometimes extremely costly. So we are pushed to settle on the choice of a concrete set of aims to consider, to gather facts about and to compare, long before we are in a position to know if these aims really should be our aims. And for many many important decisions knowledge of the facts just isn’t there to be had, even if we know what we need to know.

It is central to the analysis of this paper that the task of sorting the problem out, of making operational what in this context are the objectives to be served, what

might constitute success, what then might be a relevant fact, is the essential skill of the expert, and sorting the problem out in this way requires often very many decisions *en route* which cannot be made by the application of rules predetermined by discussion with, briefing by, the client, unless that client is expert too, meaning could have done the task himself e.g. if he had time. But experts exist precisely to deal with cases where the client cannot do it himself, has not got the skill to take those decisions. The notion of Rationality, exemplified by for example cost-benefit analysis, aims to provide answers to these questions which result via a very tightly argued process to a conclusion which follows in a rigorous way from the analysis. Ex post, that ambition may appear often to be achieved, in that if you set up the problem in that way, that is indeed the answer, and the answer is hard to deny given the way the problem has been set up. But very often in the setting up, decisions will have been made on the facts, on what are we trying to do, on what would happen if, which are inescapably matters of judgment. That means at least that if at a particular fork in the analysis the expert decides to go left not right, there are no knock down arguments which justify or explain that choice, and in the end he has to say that although there are excellent reasons for going right, his bet, his judgment, is that left is better, and he can't say anything much more about why.

Sometimes this is analogous to the master chess player who just thinks that this move is better than that, but cannot explicitly give reasons why that will maximize his chances of victory. Even if as happens in such cases an apparently successful attempt is made ex post to give an account of the expert's decision in terms of reasons which led to the conclusion, *led* does not mean anything like *led inexorably*. It just means something like *led up to*. He could always as he will agree have chosen other reasons, but he just thought that his worked better.

Sometimes the situation is far worse. It isn't just a matter that the expert might have skills we can often rely on in refining the problem in the right way, or a 'right enough' way. Often there really isn't anything in the initial setting of the problem, anything in the client's starting preferences and desires and constraints, that settles the issue. A judgment must be made, and it is a pure judgment, no matter how well informed. And very often it will have to be a judgement by the expert since by this stage the client may well not be able to understand the technicalities and refinements at stake.

It remains true that where this is the case, the expert – so long as he continues to try to serve the client's interests – is better at deciding than the ordinary bewildered person just because he is better at guessing. Guessing is a skill which some people have and some don't. The skill is usually specific to a particular class of problem, like engineering, medicine and indeed personal relations. It is not generic meaning reducible to a skill called problem solving or logic or reason which can be applied *mutatis mutandis* to any problem – the skill of the traditional British classics educated civil servant with a fine mind which he can allegedly turn to defence today and education tomorrow.

These issues can be illustrated by the example of vCJD. When Professor Southwood was asked in 1988 whether BSE presented a danger to public health, there was no satisfactory medical answer. The facts about BSE were obscure, and its relationship to scrapie, a similar and in human terms harmless, disease of sheep, unclear. There was no sound science, no incontrovertible advice, to be had. In this case, it is not just that the public is not able to assess the expert's advice - the MMR case. It is that the expert's advice is not wholly supported by auditable reasons. There are no facts available which via the application of a manual can lead to the right result. And there isn't a manual.

There are two ways of using experts. One is as the gateway to the orthodoxy. If you ask the epidemiologist about foot and mouth you may be getting access to the uncontroversial agreed view of the epidemiological profession. You have our problem that if in an accountable fashion they explain how the facts lead incontrovertibly to the conclusion, you get lost. But in principle, meaning e.g. that if you spent time and became an epidemiologist you could get to read the book that shows you that they are right. The other way is when there is no orthodoxy, meaning not only that there is no orthodox answer to the question (e.g. because it has not been asked before) but that there is no clear way of setting about dealing with it. That was the problem with vCJD. In that case, the medical expert is not the gateway to the orthodoxy but the person who is (you believe on some grounds or other) best able to guess for you, whose judgment you prefer to rely on in preference to your own.

The words *guess* and *judgment* mean something. If the advice is incontrovertible, it means that the result follows inexorably from the facts and the arguments. Inexorably means that you are not guessing or exercising judgment (I neglect whether this can ever be true). It is of the essence of the exercise of judgment that when pressed the expert has to acknowledge that he could have advised otherwise, but that he went the way he did, and there are no knockdown reasons for having done so.

In this sense, it is unavoidably indeterminate what advice the expert gives. This indeterminacy is very like that which emerges with the notion of the right decision. If the criterion of a decision being right is that it leads to the best outcome, then the outcome has to be compared with other outcomes which would have resulted from another decision. This is unanswerable even with hindsight. Take Churchill's decision not to negotiate with Hitler in 1940. Was that a good decision? To resolve that requires an unmanageably open ended examination of possibilities such as would the US have come into the war? Would Pearl Harbor have happened? It is impossible to analyze this decision in those terms. Similarly, it is impossible to analyze the analogous recommendations of experts.

Often, the public is quite content with the decisions which government or other such agents make on their behalf. When they are content, it is sometimes because full accountability has worked well, the decision making process has been made transparent and they have been able to check it to their satisfaction. Sometimes, they are not content, and in some of those cases it is indeed because they have looked at the process in the same way and not been satisfied.

But when they are satisfied, they are often so without having been presented with a transparent account of the decision, or, if they have been so presented, without having been able to check it properly, for the reasons I have discussed.

Therefore, when they are not satisfied, it does not follow that the lesson is that in future we must have more transparency and accountability. That may not be possible, because the two tough conditions may not be capable of being met. It may not be necessary, because satisfaction does not require accountability in the sense used – thank heaven: for the conditions are so often not met that if it were necessary we would have very little satisfaction. When things go right, it may not be that we have succeeded by approximating to the paradigm case. It may be that something else is going on. If in short hand we call the paradigm case the use of Rationality, if only because it relies for its success on the open and systematic analysis of reasons, then what is required for the lay client to be satisfied may not be that, but rather trust, a word which at this stage does no more than signal that something else is needed, and it is not Rationality-like.

A final example. Take the judgment of the US Supreme Court in 2000 that the counting of votes in Florida should go no further. Many believe that to have been a bad decision, because the voting was five conservative justices on one side, against four more liberal. But why is that relevant? Surely all you have to do to assess the judgment is to look at it and see whether the legal analysis is right. But try that e.g. by reading Professor Dworkin's commentary in the New York Review of Books at the time, and you will soon realize three things. First, you don't really understand the arguments. Second that as likely as not any of least two analyses could be supported: there are no knock down arguments. Third, that the best reason for doubting the conclusion is indeed that it may have been politically based. You cannot given your lack of competence make a Rational assessment. You have to use something else. Or perhaps you just follow Professor Dworkin. But why?