

Chapter 1. Tragic Random Choices

1.1 An example of a urgent allocation: Medical treatment

1.2 Administrative allocation: What's the alternative (to coin-tossing)?

1.3 Elicitation and Prospect Theory

1.4 Public Choice Theory and General Welfare

1.5 Conclusion: Is Random Distribution the right response here?

1.1 An example of an urgent allocation: Medical treatment

The context for this example is a publicly-funded organisation dispensing a benefit, here a hospital doctor deciding which patient to treat. Because this is an emergency ward and resources are limited, the decision becomes a matter of life or death. The doctor as agent of the hospital who employs her must decide which customer-patient will get an emergency treatment. A quick decision is vital. From the earliest days of transplant surgery this dilemma has attracted the attention of social philosophers. Calabresi & Bobbit (1978) in their book *Tragic Choices* examined the implications of different procedures for rationing medical treatments in short supply. One option they suggested is to use a lottery. They discuss how a lottery might affect this tragic choice, but only in theory. Broome (1984b) refers to a single case where a formal lottery was employed to allocate medical treatment, but otherwise this is treated as a hypothetical example.

In the following example, an approach to how the random distribution of scarce medical treatment might function is taken from a paper by Anand (2002). He elicited opinions on the idea by asking:

Example: Tossing a coin to decide who should get a scarce life-preserving medical treatment

*‘Two adults arrive at casualty with a life-threatening condition that does not affect their ability to make decisions. The doctor explains that there are resources only to treat one patient and then proposes that she will decide which one is to be treated by **tossing a coin**. If you were one of the patients, would you think that a doctor’s choice based on a coin toss was a fair way of choosing which patient to treat?’*

This was one of a set of questions that was put to a random sample of the public, and reported by Anand (2001) in a paper on ‘Procedural fairness in economic and social choice’ in the *Journal of Economic Psychology*. The setting is the administrative one of an agent organising the urgent allocation of a scarce benefit to a customer-patient – in this case two patients. Anand tries to get his sample of the general public to think about these patients caught in this dilemma: Would they (the members of the public) consider it to be *fair* on the patients involved to use randomization in the form of a coin-toss to decide?

Anand found that the members of his sample did not like choices being made for the patients by the toss of a coin. They mostly thought that a coin-toss would be a ‘very unfair’ way of deciding who should be chosen to receive the scarce treatment. He then went on to make a much more general claim that ‘there is strong lay resistance to random choosing as a fair process’. Anand’s explanation for this was that random choosing deprives customers and clients of some control or ‘voice’ in the process. Earlier experiments which attempted to establish peoples’ preferences for fairness under hypothetical lotteries have been carried out by Bukszar & Knetsch (1994) and Frey & Pommerehne (1993), both referred to by Anand. The results again seem to show that the public does not like decisions which are made by a lottery.

Comment: It is understandable that a questionnaire should be short, otherwise respondents will lose the plot. In this case, however, there are many other factors which could be considered before deciding whether a coin-toss is a 'fair' procedure or not. Did the two arrive simultaneously, were their injuries equally severe, are their prospects for recovery the same, has one patient greater social responsibilities, or makes a better contribution to society than the other? One simple alternative procedure which is widely used is to treat patients on a 'first-come-first-served' basis. I will elaborate on other possible procedures for selection in the next section, and ask how this might have changed the respondents' answers. Given the actual question posed by Anand, the result he found is hardly surprising. The sample of respondents were drawn from Oxfordshire, and would almost certainly have had no first-hand experience of random distribution. Nor were they given any of the possible alternatives to random choosing. They were only invited to consider the random process, one which has a very poor image. Say 'coin-toss' or 'lottery' and gambling is the first thing that springs to mind. Editors write newspaper headlines that tell readers of a 'post-code lottery' for health treatment: This only means that facilities are not always available at every location, but the lottery reference makes it sound malevolent. The word 'random' has similar bad resonances, with 'killing at random' or 'random terror' suggesting an especial level of horror. With these negative images already imprinted on peoples' minds, and presented with an unfamiliar option, it is hardly surprising that Anand's panel of respondents immediate response was that it would be very unfair to decide by a coin-toss.

1.2 Urgent Administrative allocation: What's the alternative (to coin-tossing)?

It is all very well to say that choosing by the toss of a coin would be 'very unfair', but *some* decision has to be made. Perhaps when alternative procedures are examined, the case for random distribution might be more acceptable. Administrators, be they in publicly-funded organisations like the NHS, or private profit-oriented hospitals will almost certainly be faced with dilemmas over which patient should get a treatment which is in short supply, and is needed urgently. (Why the treatment is in short supply is another question, which will be looked at later). In other situations of excess

demand, a more deliberative approach may be taken, but not here. If the proposed coin-tossing solution to this dilemma is thought to be unsatisfactory, and other simple administrative procedures like 'first-come-first-served' are ruled out, what alternatives are there? Here I will be taking an informal look at some of the possibilities, asking if they could be implemented, and how might the patients react.

Alternative 1: Use the market: Since this is an economics study, perhaps a market-based solution might seem to be the best solution. The doctor (or more likely a hospital administrator) could explain the dilemma, and then offer the one available treatment for sale to the higher bidder. Both doctor and patients might dislike using the market, but the benefits could be substantial: The hospital would gain extra funds which it could invest in facilities, which would enable more treatment to be provided. This might mean that the next time both patients could be treated.

Such ruthless application of free-market ideas would surely be rejected by the public at large. There are some things, like kidneys for transplantation or babies for adoption that are just not for sale. Whatever some economists like Friedman (1980) might say about the benefits of marketing these commodities, most societies insist that these transactions are decided by other, non-market means. This also applies to places at universities, blood donations, which are given without recompense in most countries, and social housing, available at below-market rents. The hospital would surely not want its reputation sullied by doing something which transgresses social norms by selling treatments to the highest bidder.

Perhaps the hospital would prefer to avoid the market-based approach for another less high-minded reason: Administrative cost and convenience. Even if an auction could be organized in a short space of time, running it could be expensive. Money has to be collected with some degree of certainty. Clerical staff to deal with such collection would need to be kept on permanent standby. It may simply be less costly to let the doctor decide, and forgo the extra revenue.

*Alternative 2: The doctor uses her **clinical** expertise:* If the doctor was to take the course of action proposed by Anand, she would rightly be condemned as professionally negligent. We expect those whose judgment we value to use their skills and knowledge for our benefit. It is difficult to believe that in the scenario described above, the doctor would be unable to decide which patient had the better clinical prognosis. White & Stancombe (2003) review the literature assessing just how good are the diagnoses given by doctors, and other professionals. The conclusion is that, broadly speaking, doctors are quite good at it, but not perfect. The impact of ‘tacit’ knowledge is important, but should not be overstated. The reliability of professionals can be enhanced with ‘expert system’ techniques (Boyle, 1984), but a genuine if honest error is always possible. In the dramatic choice described by Anand, the pressure of the situation could lead to a mistaken diagnosis—that the ‘wrong’ patient would be treated. Despite this, most people would gladly put their trust in the professional judgment of these most-respected clinicians, rather than tossing a coin or auctioning to the higher bidder.

*Alternative 3: The doctor uses both **clinical** and **merit** judgment:* Again, rightly the doctor uses medical grounds to discriminate, but in the circumstances, this has to be a rapid assessment. In order to find further grounds she enquires about the character of the two patients. One may be a young parent, the other an older single criminal. With such a clear difference in social (but not medical) merit, is it right to reject the criminal? This is not an idle question: Right from the start of organ transplantation such moralistic contentions were weighed up. In Seattle the so-called ‘God committee’ was set up to make these difficult choices (reported in Calabresi & Bobbit, p187). The committee eventually found that it was too agonising to make these choices, and passed the task back to the medical practitioners. In the end it was felt that only *medical* factors should be taken into account. Even if no overt rules on social merit were in place, we should not be surprised if the doctor, genuinely uncertain on medical grounds, was to pick the ‘nicer’ of the two patients. I will be returning to the vexed subject of discrimination later in Chapter 5 which deals with both intentional and unwitting discrimination.

As to the views of the general public, I would imagine that they would be happy to take their chance with a doctor who uses both clinical and merit criteria to choose us as the winner: We all harbour a touching belief in our ability to seem nice in others' eyes.

*Alternative 4: The doctor is unable to discriminate on medical grounds, so **secretly** picks one of the patients **at random**. She then announces her decision **pretending** it to be based on medical grounds.* This is not as far-fetched as it may seem. Elster (1989) in his masterly 'Solomonic Choices' gives the example of child custody cases, where the judge is frequently unable (in his own mind) to give a clear-cut decision. Yet decide he must, so he goes ahead, dressing up the verdict with trappings of rationality. This, claims Elster, satisfies both parties, the winner praising the wisdom of the judge, the loser cursing his bias. No doubt a similar process might go on when a medical doctor decides, even if partly randomly and in secret, between her two patients: So long as both patients *believe* that their case is decided clinically by an expert, then both winner and loser may find it acceptable. The doctor herself may even be a bit cognitively dissonant—convincing herself that she is doing the right thing for the right reason, exercising judgement based on intuition rather than validated knowledge. This form of fudging may be acceptable all round, but it is fraught with dangers. If fakery is suspected, patients rapidly lose their trust in their professionals. Unwitting discrimination seems inevitable. True expertise will fail to develop unless its limits are acknowledged.

Alternative 5: The doctor truthfully explains that despite using her clinical expertise to the full she simply cannot decide between the two. She then invites the patients to decide by the toss of a coin. This is not a repeat of Anand's original idea: The doctor first uses her clinical judgment, only then explaining her dilemma. Even so, I doubt if actual patients would find this process any more acceptable than the simpler version described by Anand. As he discovered, 'coin-toss' and 'random' have negative overtones; it is far more comfortable to cling to the security blanket of 'doctor will decide'. There are a number of points to be made for and against using a simple lottery in this situation: Against a lottery is Greely (1977) who suggests that if

recipients can argue about any allocation, they feel more satisfied. Anand was also interested in what is called ‘voice’—that one of the reasons a coin-toss was thought to be unfair is that it deprived customers of a say in the decision. In favour of a visible act of coin-tossing Calabresi & Bobbit explain that it draws attention to the fact that resources are limited. Edgeworth (1888) suggested another benefit would be that the public, seeing a random drawing take place, would be alerted to the ‘aleatory’^{*} nature of the decision. Bureaucrats might not like having such attention focussed on this shortage of resources and their uncertain knowledge.

Alternative 6: The doctor explains, that in her clinical judgment, although patient A has a better chance of survival than patient B (being of a statistical turn of mind she expresses as 2:1 in favour of A), she doesn't think that B should be automatically excluded. After all, she explains, she is not infallible. So again a random event will decide, only this time it is the roll of a die: If 1, 2, 3 or 4 come up then A will be treated, if 5 or 6 appear, then B will be the winner. If this was seriously proposed to two near-death patients, they might find having to wrestle with such a complex statistical argumentation too hard to bear, and expire from mental exhaustion. Perhaps this shows that the wisest choice may be alternative 4, above: fake a clinical judgment, don't confuse the patients. Broome (1994) wrestles with some of these conflicts, asking whether we should attempt to be fair, or to try to achieve the most good. This problem arose because of the Oregon state system which tried to classify and prioritise medical prospects, and treat accordingly. Although Broome decided that it was both fair and would do the most good to treat the patient with the best prospect, he did not go on to consider their relative merits, and how they might be settled by a randomised decision as suggested here—a strange omission since Broome has written extensively elsewhere about the uses of lottery selection.

Yet it is something like this last alternative that I intend to advocate (although not necessarily in this situation). Of course professionals should use their expertise as far as it goes, but they should also recognize its limits. When expertise runs out, it is

^{*} aleatory – depending on contingencies, from the Latin *aleator* – a dicer, *alea* – a die (*Chambers English Dictionary, 1990*)

wrong to pretend. Instead, the honest answer to this lack of knowledge is, like our doctor, to toss for it. To be absolutely fair, this random selection should be weighted according to relevant objective criteria.

1.3 Elicitation and Prospect Theory

In markets, customers ‘put their money where their mouth is’, and reveal their preferences through their spending. In non-market situations it is more difficult to hear the customer’s voice. Elicitation, as in the example in the last section, is one attempt to find out directly from potential consumers which choices they prefer. ‘Prospect Theory’ is the name given to this process by Khaneman and Tversky in their 1979 paper on the risks that individuals perceived in various situations, and which risks they would accept. Their experimental technique was to give subjects a range of hypothetical choices, in the form of a questionnaire. From their answers, they deduced that customers do not always behave in ways that economic theory would predict. Elicitation in these non-market contexts is similar to the psychological approaches taken by market researchers attempting to understand the motivations of shoppers—for example Fox & Lusk (2003) on ‘Value elicitation in retail environments’.

Using elicitation, Anand claims that he has demonstrated that the public would reject Random Choosing because it would be ‘unfair’, or even ‘very unfair’. Actually, all he has done is show that in the particular circumstance of the hypothetical question posed, that the public don’t like randomness. As I have argued earlier, that this is probably as much to do with the unsavoury reputation of randomness and lotteries, rather than any deeply held and thoughtful belief. Eliciting considered judgements about an unfamiliar idea is unlikely to produce good results. In a later example (Chapter 3 Medical school entry in the Netherlands), a survey on the attitudes of those who have first-hand experience of random choosing produced a much more positive result. As a rule, questionnaires should stick to factual questions of the ‘How many...?’ and ‘How often...?’ type. Feelings can be tested, but only indirectly, again through factual questions. I give an example of a questionnaire which exhibits these

characteristics in an appendix (B), which measured the Subjective Well-Being of a group of students.

Perhaps it is these shortcomings of a single-shot questionnaire as a form of elicitation that led to the development of procedures like Citizens' Juries. The non-market choices which occur in the public sector are normally decided by elected representatives. Because it was felt that the politicians were out of touch then the authentic and considered views of a cross-section of the electorate should be canvassed in the form of Citizens' Juries. (Reported in *The Times* 1.4.96 and in *The Guardian* 20.1.99). Citizens' juries are groups of 12 to 16 people drawn randomly from the local population who are paid to spend a few days considering a particular local problem. Following presentations by experts on the options available, and led by a professional organiser, the citizens' jury will ponder the alternatives, eventually reaching a conclusion as to what is the best policy. Many of the topics discussed include medical themes: Whether Walsall should pay for a new hospice; to decide what drug treatment works best; whether a small local hospital should close. Similar 'People's Panels' are still in operation, for example in Birmingham (information at <http://www.birminghampp.org.uk/>).

They are not without their critics: The 'jury' may not be truly representative of the population which would be affected. The way in which they are presented with the issues can distort their conclusions. In any case, the results of their deliberations do not constitute any form of proof of settled preferences, such as those which emerge from the workings of a free market. According to Sutherland (2004, p28) not much has been heard lately of Citizens' Juries, so maybe they no longer represent an effective means of sounding out public policy preferences.

The central issue which elicitation tries to address remains: In particular non-market allocations, are the beneficiaries (in other contexts they would be called customers) satisfied with what they get? Also, are they happy with the *process* of allocation by which they get their benefit? (This was Anand's question). It is not always easy to produce carefully crafted questions to elicit reliable answers to these questions.

1.4 Public Choice Theory and General Welfare

Buchanan (2003), one of its founders, explains that originally Public Choice Theory (PCT) was called ‘Non-Market Decision Making’. This original title fits very well the description of the topic which I am investigating. PCT provides many useful insights and explanations which are relevant to random distribution, which I will be drawing on. In one respect, though, I have not found PCT particularly useful: The Theory of Public Choice concerns itself mainly with the choices made *for* the ‘public’, choices which are made willy-nilly by elected politicians and bureaucrats. It has not, as its title suggests, much to do with Choice *by* the Public, about how the public get to choose their preferred goods and services and whether they are satisfied with the result.

Initially, PCT was concerned with the workings of the political process, and how it might lead to influential groups, such as farmers, capturing a greater slice of the available tax-funded resources. Buchanan draws attention to the prevailing socialistic ideology, which was so pervasive in the middle of the 20th Century. Welfare Economics had identified failures in the market economy. Collective provision was the answer, but, as Buchanan suggests, it was assumed that these politicised corrections (collective provision) would work perfectly. ‘Public Choice then came along and provided analyses of the behaviour of persons acting politically, whether voters, politicians or bureaucrats.... Public Choice became a set of theories of government failures.’

Considering the topic I am investigating—the case for random distribution in non-market allocation—PCT in some ways has too broad a remit, and in another way is too narrowly focussed.

— Government decisions result in the provision of many goods. Some of these like defence are consumed collectively, others are services to individuals like health care or education. Generally in looking for reasons to adopt random selection I only consider allocations where an individual benefits. A further complication is that many of the benefits provided are classed as ‘merit

goods’—provided cheaply to encourage consumption, because of wider benefits that, for example, having an educated workforce brings.

— Many ‘non-market decisions’ do not involve government. Commercial firms also make many decisions outside the market. I will give one example of a firm distributing its product using a lottery, but the major non-market decisions which all firms make relate to employment. The bureaucratic processes used in commercial firms closely mirrors that in government, and can validly be analysed using the same PCT tools.

Public Choice starts from the assumption that the players in non-market decision making—the politicians, voters, bureaucrats—will act primarily as self-interested individuals who seek to maximise some ‘good’. A criticism, which Buchanan acknowledges is that individuals may be differently motivated when they are choosing ‘for the public’, rather than for themselves in a private capacity. He accepts that the economic model of behaviour is not the be-all and end-all of scientific explanation, but, as he points out, assuming such self-interested behaviour on the part of the players leads to strong predictions which are almost as reliable as those found in the market.

With this in mind it would be useful to think about the motivations of the players in the scarce-medical-resource allocation drama described earlier in this chapter:

The Doctor, may be quite junior and may therefore lack confidence and be risk-averse. She will operate under a professional ethos, which tells her to maximise patient benefit. But she will also, probably being a middle-class professional, have a particular set of social values, although her training should enable her to act more dispassionately. She is also an individual in a job with her own career aspirations. She may have a family to support, so is likely to practise ‘defensive medicine’— avoid taking decisions which may be criticised, or worse, lead to lawyers becoming involved. Another aspect of her career is that she would wish to be esteemed by her peers and supervisors, a major source of job-satisfaction. If she were allowed to make a decision by a coin-toss, this might lift the burden of having to decide who must die. Whether

the doctor will act like a 'Knight' and uphold public service values, or will behave in a more 'Knavish' self-interested manner depends, according to le Grand (2003) very much on the way the organisation is operated. Unlike the pessimists of Public Choice Theory, given the right environment he insists that agents can be relied upon to act like 'knights'.

The Patients and their nearest and dearest will each be attempting to get the prize for themselves. As it is valuable prize, it may be worth paying handsomely for the 'gift of life'. But since this is socialised medicine it will be free at the point of use, so their resources may be directed into other means of winning the prize. Bribery is one possibility. Strategic behaviour is another: Exaggerating the illness, concealing information which casts them in a less favourable light. A coin-toss decision should limit these forms of behaviour, by making them less profitable. A feeling of injustice at this process, or the shortage of medical equipment could lead to time spent lobbying local A.M.s or M.P.s.

The Hospital Administration: is part of an ongoing business. The first question for them should be: Why is there a shortage of resources? This may turn into an excellent opportunity to campaign for more funding and a chance to expand their empires, just as PCT would predict. They should also seek efficient means of running the business, so a coin-toss may be a cheap way of resolving dilemmas. They would also wish to avoid any scandal or legal complications which might do damage to the reputation of the hospital. The coin-toss is a two-edged solution: It might provide some legal protection against mal-practice suits, but its use, as Anand suggests, might lead to accusations of acting unfairly.

One of the core ideas of PCT is that of 'rent-seeking'. Like the patients in the example above, if the benefit they seek is provided for free, then they have an incentive to expend valuable effort in order to win that benefit. Rent-seeking as a significant factor will arise in all the examples which I will be presenting. It was Tullock who initially identified that all rent-seeking is wasteful. The value of the prize is not enhanced by

the 'rent' spent in seeking it. For the losers the expenditure is all loss. Even the winners may end up as net losers if their expenditure is too great.

Another idea associated with PCT is the 'Principal-Agent' problem: If you assume that all the actors in the allocation process are motivated solely by self-interest, how can the Principal (the person or body with a benefit to bestow) ensure that his Agent (the administrator) performs as he is required to, and does not just please himself or use the opportunity to make money at the Principal's expense? Although PCT directs its attention to publicly provided tax-funded organisations, this is problem which applies equally to both public and private sectors. Bureaucracies in both make many significant decisions which affect consumers and employees. Changing what they do, for example by introducing randomised procedures, will have to take account of their motivations if the greatest benefit for the customers is to be obtained.

The aspect of 'welfare economics' which Buchanan rails against was its concern with market failures. But welfare economics also considers how the economy produces and distributes benefits for people. It may also be a matter of welfare not just what prize is gained, but *how* it is gained. In the medical drama above, the 'how' question is not easy to answer. Attempts have been made to find a rational answer— for example using the QALY (quality adjusted life year) system or the Seattle 'God Committee' of responsible professionals, as described by Calabresi & Bobbit (1978).

Elicitation provides one method which seeks to understand how individuals value the benefits obtained through non-market allocations. It may also be possible to uncover an equivalent market value by studying consumers' behaviour (an idea developed by Tiebout). An example of this is the considerable effect a good local state school has on house prices (Cheshire & Sheppard, 2004). Both elicitation and the implied market benefit of free goods will feature in many of the examples used later.

1.5 Conclusion: Is Random Distribution the right answer to the scarce treatment dilemma?

The techniques associated with Public Choice Theory, such as the analysis of rent-seeking will provide useful insight in all the examples later on, as will Economic Welfare in non-market allocations. In the hypothetical case described by Anand, random selection is probably *not* a good idea. To toss a coin to select a patient for an emergency treatment leads to one of two results—life or death. Naturally the patients will be risk-averse, so being forced to confront such an extreme outcome would be agonising. Calabresi & Bobbit (1978) concluded that in these circumstances that a lottery is not the best procedure in these extreme situations. In other cases where the prize may still be an important one, but the alternative not too grim, a random allocation may be more acceptable. If you fail to win a place at medical school, for example, you will still be eligible for other courses. If you miss out on a job or promotion, other reasonable possibilities remain. For the life or death example in this chapter the option recommended by Elster seems the most practical: Make the decision on clinical grounds. If that is impossible, pretend it is an objective choice, while discretely using a surreptitious randomisation device (glancing at a watch with a second hand works well).