

Chapter 8. Why Random Distribution Works

'The engineer is the guy who makes for one dollar what any damn fool can make for two'
(popular saying amongst engineers)

8.1 Review

8.2 Mostly it's about designing an economic mechanism

8.3 Does it work?

8.4 Does Random Distribution work well: For whom?

8.5 Works well for whom: Efficiency

8.6 Works well for whom: Reciprocity and Inter-personal Comparison

8.7 Works well for whom: Justice and Fairness in Society

8.8 Works well overall: Stability, Accountability and Rotation

8.9 Conclusion: A plausible idea?

8.1 Review

The context for non-market distribution involving randomisation: In the previous chapters, I have introduced seven different examples of non-market distribution, all of which involved an element of randomness. Each example was chosen to illustrate a different context, some which show where random distribution works, and some where it may be inappropriate:

<i>Ch</i>	<i>Prize</i>	<i>Source of Prize</i>	<i>Recipient</i>	<i>Allocator</i>
1	Medical treatment	Public Agency(hospital)	Patient	Doctor
2	Wimbledon tickets	AELTC – commercial firm	Fan	Manager
3	Place at med. school	NL Government	Student	Govt agent
4	118 phone numbers	OfCom Govt Agency	Firm	Oftel
5	Sack from Job	Chinese Govt	Employee	Govt
6	Workplace	Mine owner/miners	Co-worker	Co-worker
7	Green Card	US Government	Anyone	INS

Only the third example—the NL medical school entry—is a weighted lottery; all the rest use a strictly equal means of random distribution. As can be seen from this list, the source of the prizes is always some form of organisation. These may be commercial firms or governmental agencies, and often the prize-winners may be uncertain which. All of the recipients, bar one are people. The one exception, in Chapter 4, is of a government agency awarding commercial firms, rather than individual people.

I introduced some ideas from economic theory with each chapter which will re-appear in the discussion in the sections that follow.

8.2 Mostly it's about designing an economic mechanism

Roth's (2002) ideas about design of economic mechanisms has provided the most important framework for judging the impact of randomised distribution. Of course, managers, politicians and sometimes even economists have long been involved in developing and implementing non-market distributional mechanisms. Generally, these practical people ask just two simple pragmatic questions about an economic mechanism (actual or proposed):

- does it work?
- does it make things better?

Economists, with their tradition of always seeking to optimise, might add a third question: Is this the best that could be done?

Although Roth provides some theoretical ideas for design of economic mechanisms, it is Binmore & Klemper (2002) (B&K) who have provided a hugely successful example. I find their description of design of economic mechanisms more succinct than Roth's. B&K describe the three main elements that went into their design: The first was experiential: They drew on the previous experience of sales and allocations for radio frequencies. To better understand the likely reaction of the players involved in actual auctions, they conducted directed experiments. And thirdly they used economic theory to explain and understand. All three—experience, experiment and

theory—can be called in aid when appropriate, or as B&K put it: ‘It’s horses for courses’.

The experience element cited by B&K has another dimension: The use of validated knowledge. Before accepting any ‘conventional wisdom’, it is worthwhile to ask if we have any evidence to support it. For example, most selection processes place great store by interviews, but are they effective? Higher examination grades suggest higher ability, but does that relationship hold over the full range of scores? There are many sources of validated knowledge, not least from the extensive literature coming out of experimental economics.

8.3 Does random distribution (RD) work?

If common-sense deems random distribution to be ‘ludicrous’, the results of surveys of public opinion do not provide much encouragement either. As was shown in Chapter 1, invariably the public feel that lotteries are an unfair way to resolve acute medical dilemmas. Even in controlled economics experiments, there is little enthusiasm for a lottery over more manipulable distribution mechanisms. Only the example in Chapter 3 of entry to Dutch medical schools reveals any public support for random distribution. The students who have experienced the system are quite positive about the value of using a lottery.

A further problem I have discovered is that even where random distribution has been used, it has not had any champion, any leading figure to promote its use. In the examples in the previous chapters, random distribution was a compromise between warring parties (Dutch medical school entry), emerged from a pre-industrial age (the Cavil in the coal mines of County Durham) or was used as a quick, cheap way of unloading booty (118 phone numbers). In the economics literature, there is little to draw on either: I have only found two articles which directly address the economics of random distribution, starting from an actual experience of its use: Boyce (1994) and the follow-up by Taylor, Tsui & Zhu (2003). Others make use of the ‘natural experiment’ presented by random distribution to pursue other questions: for example Sacredote (2001) used the random allocation of students to university accommodation

to test if being lodged with a bright student improved the grades of dimmer fellow lodgers. Cullen et al (2003) examined the effectiveness of randomly allocating schoolchildren to schools in Chicago, but only to test if their education improved. Papers such as these do not ask if random distribution was a good idea in the first place. One intriguing theoretical conclusion is given by Jeol & Laffont (1999) is that under certain asymmetric information conditions, randomisation in layoffs can be shown to be 'optimal'.

It is only by examining a number of different real-world examples of random distribution that evidence emerges that it can work, and indeed works successfully. The example in Chapter 6 of the Cavil was a random distribution mechanism which significantly affected the lives of a large group of workers. The Durham coalminers were not passive acceptors of their situation, rather they were a well-organised group who had some control over their fate. They chose to persist with the Cavil, as did their employers. The system was in use over many years, and throughout the Durham coalfield, in its time the biggest in the UK. Despite alternatives readily available, the Cavil survived as the preferred option. The evidence is overwhelming that the Cavil, an example of random distribution, really did work.

If the Cavil, a Victorian institution, might be dismissed as irrelevant to modern circumstances, the same cannot be said for the Dutch medical school entry lottery. As shown in Chapter 3, this currently existing hybrid merit and lottery selection mechanism has been in use for more than 30 years. It affects thousands of students each year, but it has not been without its critics. Crucially, it has been subjected to rigorous scrutiny by the Drenth Commission (1999) and has passed with flying colours. Indeed Drenth concluded that the evidence supported *increasing* the random element of the selection and rejection process.

Even where the results of random distribution are questionable, such as the 118 phone number lottery or the Wimbledon tennis tournament ticket ballot, the random distribution *worked*. The numbers and the tickets were distributed, whatever the ultimate outcome. Although random distribution may be rare, it cannot be said to have failed in its basic function of distribution.

8.4 Does Random Distribution (RD) work well: For whom?

Does it work well, better than alternatives? begets the follow-up Better for whom? which I want to address first. In standard market transactions, there are two players, the supplier and the purchaser. Their objectives are clear-cut and un-ambiguous: The supplier wants to maximise profit, and the customer/purchaser wishes to maximise satisfaction. This can be elaborated if firms or corporations are taken into account. Firms may also wish to remain in business, corporations may also have a reputation to maintain. Consumers continue to be seen as individualistic maximisers, seeking the best basket of goods for their money.

Non-market transactions will normally be managed by a bureaucratic 'agent' of the organisation which has the asset to bestow. The recipient may be more than just an individual 'customer', but be seen as part of a larger community. Hence those at the heart of the transfer may not share the motivations of those behind them. For simplicity I will identify four entities who are involved in any non-market transfer especially those which involve random distribution:

Principal (Organisation) >> **Agent** → **Recipient** << Community
(1) (2) (3) (4)

I will deal with each of these entities separately. Of course, in some situations, like the fishermen's cooperative in Kerala (in Chapter 6), 'organisation' and 'community' are the same people. Even in the first example (in Chapter 1), deciding who should get the scarce medical treatment, the organisation behind the doctor/agent is a hospital; this in turn may be a public governmental body, and so part of the larger community.

1. The '**principals**': who decide what mechanism of selection/rejection is to be used, can be in either the commercial or the public sector.

- 1a. Commercial Organisations (Firms):** are in business to make (long-term) profit, indeed are required to do so by the doctrine of primacy of share-holder interests. Profits can be enhanced by increasing revenue and/or reducing costs.

But maximising revenue, for example by the Wimbledon tennis tournament organisers ((Chapter 2) may not be pursued to the full. Orderly marketing, or a wish not to antagonise their fan-base, may explain their seemingly un-commercial behaviour. A cynic might argue that this does not detract from the primacy of maximising profits: That giving the appearance of being nice to their customers, nice to their employees and nice to the environment boosts long-term profits. Giving away their products cheaply using a random technique may be just such a strategy.

1b. Public organisations: include Government, politicians, and publicly funded organisations. These are supposed to act on behalf of the electorate. They too will be driven by the presentational questions that concern commercial organisations, as well as a need to constrain costs. Instead of profit, there may be a range of conflicting aims, such as widespread distribution of a service, or compensating for mal-distribution. The Theory of Public Choice reminds us that these politicians may be just as responsive to corporate influence.

2. Agents are the essential bureaucrats allocating the benefits to individuals according to the directives of their principals. Because of the structures of organisations, it may make little difference whether the organisation is commercial or public. There is a great concern that these agents can be induced to align their objectives with that of their principals (Public Choice Theory again). In addition, we should not forget that these agents are human beings, with the normal human feelings and frailties. Their welfare should not be forgotten in the design of economic mechanisms.

3. Recipients are the people who win or lose in allocations. These are the customers, tenants, pupils, parents, patients, job-hunters, employees, or any other role they may be fulfilling at that moment. There seems to be an attitude, for example in Roth (2002), that the recipients are a pesky nuisance who must be fobbed off with something, enough to stop complaints, or from trying to change their allocation. Balinski & Sonmez (1999) found similar problems in Turkish school allocations. I have described this as a ‘beggars can’t be choosers’ attitude. This, I believe is wrong. The promise of free markets is that the customer is king; is delivered the most product at the least cost. When designing or assessing non-market allocations, the ‘customer’

should be considered first. *The ultimate aim, and indeed justification of any man-made economic system should be the enhancement of the human condition.* How this is operationalised in actual allocations may be difficult to identify, let alone achieve. I will make some suggestions in the final chapter.

4. Community in relation to a non-market allocation is relative. Group size has a particular significance in judging whether a specific allocation mechanism is appropriate: At the smallest level is a group of people who can know each other on a face-to-face basis. Wider communities such as the employees of a firm or citizens in a province are the next level. Whole countries, or even humanity as a whole can also be the context for allocations. The significance of size matters because of the possible different motivations involved: Self-interested behaviour is present at all levels of course, but care for the well-being of others is more potent in a smaller group. Concern about more abstract notions of justice and fairness are more likely to be found at the wider level.

Wider social benefits may also over-ride individual merit: Consider the allocation of places for entry to medical school: The most likely to succeed (the most meritorious?) are likely to be not just the applicants with the highest A-levels, but who are also female, middle class, and white, with a previous degree (according to statistical analysis by Leslie, 2003). They may also be the candidates preferred by the medical school, but other social objectives may be required or enforced:

- Fairness and equality: would require widening access to other groups.
- Diversity can pay: Learning with a more diverse student group may be useful in a profession which requires contact with the population at large.

Criteria for deciding if RD (Random Distribution) works well

In each of the next three sections I will examine one of the criteria by which any particular mechanism involving random distribution might be judged. These criteria will be related to the objectives of each of the four ‘players’—Principals, Agents, Recipients, Community—identified above. The criteria start with the most concrete

(and least tendentious for economists)—Efficiency, especially involving rent-seeking. (8.5) Expanding the criteria will include Reciprocity and Inter-personal Comparisons. (8.6) Finally I invoke the somewhat philosophical criteria of Justice and Fairness, which are seen as very significant. (8.7)

8.5 RD works well for whom: Efficiency

(1). Efficient for the Principals: Organisations both public and private:

Efficiency for firms and organisations in both public and private sectors is broadly similar. Of course both will seek a distribution mechanism that works, and does not come apart post-allocation. I will be dealing with these system-wide considerations later in Section 8.8. Here I am concentrating on the organisation-specific efficiency considerations. For them it is always worthwhile to reduce input costs while at the same time achieving the same or greater outputs. For example, in personnel selection processes, using random allocation is generally very efficient (cheap) compared to the complex procedures required by conventional merit assessment. A large organisation may find setting up merit selection procedures a relatively minor cost, but smaller groups may struggle. In the case of the allocation of social housing, a large Local Authority can pursue the ideal of a complex points-systems. A small Housing Association, having few units to allocate to ‘deserving’ applicants should mirror the system devised by the large Local Authority (Council) housing departments, but the administrative burden would be considerable. How much easier and cheaper for them to announce simple objective entry criteria, invite applications and run a lottery if demand exceeds supply.

Reducing administrative costs is one way that random allocation can be more efficient, but what of the outcome? Commonsense decrees that the more effort that is put into the selection, especially to identify merit, the better the result. This, as Drenth (1999) was able to show in the case of selecting students for courses is a delusion. Even using the simple available measurements gives very weak predictive power. To re-iterate the argument in Chapter 3 on ‘merit’: The relationship between merit score and performance is almost flat in the likely operating zone, since all applicants have

been pre-selected to some extent. There is a great deal of ‘noise’ in the system meaning that any trend has a lot of variation about it. A top-merit candidate is almost as likely to fail as is a lower ranked one to succeed.

The same strictures apply to employment-related selection. From what is known about valid methods of identifying talent (see Chapter 5) most of what goes on can only be described as ‘dignified ritual’. Selection by a random process, with some defensible eligibility criterion will give results which are hardly worse than any form of ‘merit’ screening. There can also be positive benefits for the organisation: Cook (2003) says that some of the most elaborate selection procedures—he cites the case of the UK Foreign Office—tend to pick the usual pleasant, loquacious candidates, who mirror the characteristics of those already inside. A benefit of random selection is that it will throw up a few oddities, especially beneficial in bureaucracies which are prone to getting stuck in their ways. This is only speculation, but I believe that the ‘grit-in-the-oyster’ from random selection could enable just the right kinds of mutation to allow the organisation to survive by evolving.

Achieving more effect through randomisation? Rationally, the value placed on a 10% chance of a £100 prize should be the same as a certainty of £10. This may not be how human psychology works. Perry, Erev & Haruvy (2001) suggest that if motorists became aware that some speeding violators were to be given ‘bad lottery immediate punishment’ this would be more effective than a fixed penalty fine for everyone caught. Their results were based on experimental work. They suggest that ‘large rare punishments are stretched in effectiveness through the use of lotteries’. Another example, this time based on an actual lottery is given by Hassink & Koning (2005). They studied a Dutch firm which was trying to reduce absenteeism. To encourage attendance, regularly attending workers were entered into a lottery, with a small proportion publicly winning 75 euro. The results for the firm were spectacular: For an expenditure of 525 euro they achieved a return of 5,760 euro. As an explanation H&K speculate that ‘Workers may be intrinsically motivated to participate in the lottery, just because of fun.’ Both of these examples raise the intriguing possibility that the value of a randomised prize will be perceived as greater than its deterministic

equivalent. There may be significant potential for organisations to boost their effectiveness through randomisation.

Protecting the organisation: Organisations are already compelled through anti-discrimination legislation laws to treat people in a non-racist, non-sexist way. Other legislation may follow covering age and other categories. To comply with this requires some effort by organisations, such as staff re-training and monitoring of employees. Failure to comply even if unintended may give the organisation a major financial or reputational problem. Since random selection is inherently fair to *all* groups and classifications, it will provide a defence that no discrimination witting or otherwise has taken place. Random allocation is even proof against any future anti-discrimination legislation.

Controlling the agents Corruption is an ever-present problem in organisations, and is one reason why random selection has been used in the past—for example, distributing government posts amongst the ruling oligarchy in renaissance Venice. Since lottery results cannot be predicted, they cannot be fixed. In a modern British setting, especially in the public service, protection against corrupt behaviour may seem superfluous, but not entirely so. Lord Bancroft (1995), former head of the British Civil Service says that ‘it is natural for bureaucracies to be corrupt’. The more the distribution mechanism is determined randomly, the less possible it is to corrupt it. A fully randomised distribution is incorruptible.

Effort by the agents: Above all, the principals want their agents to exercise diligence on their behalf. Financial incentives may work in some cases, but are not usual for people-selectors like admissions tutors or housing managers. Instead it is hoped that the agents have sufficient intrinsic motivation to do a good job. I accept that superseding professional judgement by the use of random procedures may lead to demotivation. Alternatively, since selection can safely be reduced to a simple fact-checking process, lower grade staff can be employed.

(2). Efficient for the agents: What's in random allocation for them?

De-skilling, demotivation Because random allocation removes some of the need to exercise judgement by the agents, this may turn 'knights' into 'knaves' to use the labels developed by le Grand (2003). It may be personally rewarding to imagine one possesses special skills and uses them on behalf of one's principal. It is de-motivating to be told that such skills are ineffective, mere ceremonial.

For some situations, for example doctors deciding who should get a scarce treatment, there is still a great deal of clinically autonomous judgement required. For them, a random distribution would be a rare event. But in most human-resources type selection I envisage an element of random selection could be near-universal, with 'judgement' little used. This could lead these agents into a moral hazard: They may become complacently dependent on random selection, mistakes being shrugged off by comments like 'Well what do you expect? It's only a lottery'. If as suggested in the last section, lower grade staff were to be employed, they would have neither the ability or the incentive to seek out the better, yet still valid merit criteria.

Avoiding agent anguish Many of the agents' decisions are difficult because they involve inflicting losses. Deciding who should be dismissed in a redundancy, or even failing to award a job to a well-qualified candidate brings woe to both parties. Having a lottery shifts the burden to a neutral non-human arbiter. If the organisation is small-scale this intensifies the potential grief of sacking or job refusal, so random selection should be particularly helpful here.

Drive out false pride A particular delusion that some agents may harbour is that they are special, just because they are in the fortunate position to select people for some highly desired prize, such as a place at medical school, whereas in fact they are only 'rent-awards'. Pride may lead them to believe that popular equals intellectually demanding, which in turn leads them to accept only highly qualified applicants. Boosting the strength of the signal required forces the applicants to make more

(unnecessary) effort. Selection with the aid of a lottery removes this malign effect of false pride*.

When agents reach the top of the greasy pole within organisations they seem to become not just knavish, but swinish in respect of their own pay. Grossly inflated salaries are justified, in part, as a reward for exceptional talent. If agents had reached the top through a regular process involving random selection, they would have less reason to make such claims. On the evidence, for example on football managers, there is little reason to expect worse performance in the job either. Thus the salary bill for the organisation could be significantly reduced.

(3). Efficient for the Applicants and Recipients: The people on the receiving end:

Rent-seeking and signalling: Rather than the market, comparison with conventional ‘merit’ selection is a much more appropriate way to judge random allocation. Where selection is on ‘merit’ (the standard contemporary mode) and there is an excess of applicants, then a complex charade ensues: Since all have the necessary basic qualifications, then those with better grades will stand out. When everyone has top grades then this fails to give out the required signal, so secondary characteristics are invoked—an elaborate C.V. of good works might be looked for.

Investing in these extra ‘merits’ may pay off and win the prize. It may also have some benefit for all the applicants. Gaining more qualifications may benefit the economy generally. What is clear is that many applicants spend a great deal of time in order to gain these extra ‘merits’, not as an essential pre-requisite, but just as a signal, a form of wasteful rent-seeking. If the merit requirement is limited to what can be shown to be significant in predicting success, with excess demand dealt with through a lottery, such waste could be eliminated.

* A good example which seems to have taken this advice : At Huddersfield University the physiotherapy course was heavily over-subscribed. Instead of increasing the entry requirements, which they knew was an irrelevancy, they chose to award places by a lottery (personal correspondence with the V.C., Aug 2004)

Another perverse signalling effect could also be eliminated by random allocation: If a course advertises that it will only accept top-scoring students, this creates a challenge and a lure. Students who feel they might be good enough may be tempted into that field, just because they can gain entry. An honest statement of the real, if more modest entry requirements, to be followed by a lottery could act as a screening device. Applicants would then be more likely to choose an appropriate course.

(4). Efficient for Society: How it might benefit from the use of Random Distribution

If a significant fraction of the energy dissipated through rent-seeking could be saved, then a large resource could be liberated for more useful social goals. In one small example, I estimated that students were each spending, on average, two extra months of their time to gain unnecessary extra entry qualifications. (see Appendix B for details). The alternative opportunity implications for this wasted resource are obvious. So too are the potential tax-reductions brought about by public bodies becoming less wasteful through the use of randomised procedures. Commercial firms could also produce more cheaply, benefiting customers.

8.6 Works well for whom: Reciprocity and Inter-personal Comparison

Since this is about emotions, there is no need to consider *things* like organisations. In the context of a group of *people* who interact with each other in an economic transaction, the idea of reciprocity has been explored by experimental economics. From numerous cases it has been shown that individuals do not act in systematically self-interested ways. They have a care for others in their group, and feel better by being nice to them (Layard, 2003). As explained in Chapter 6, this benefit of ‘altruistic’ behaviour should not be seen as just an add-on to consumer self interest. As Frank (2004) insists, values of fairness and reciprocity are in a different domain to those of greed and self-interest. Oxoby (2003) has modelled an extension to allocation process satisfaction which includes the value placed on inter-personal comparison, which may provide a framework to measure its significance.

Compared to distribution through the market, random allocation seems to have a fundamental weakness: Benefits do not finish up with those most eager for them, as measured by willingness to pay. It is always possible to imagine that some prize-winners in a lottery would be willing to trade with others for money. But direct comparison with the market may be inappropriate. When non-market allocation is used, be it by lottery or on merit, it is reasonable to assume that collective values are involved.

Attempts have been made to modify existing lottery mechanisms to make them give a more market-like result. For example, a paper by Hyland and Zeckhauser (1979) was specifically motivated by the housing lottery for students at Harvard. To produce a more market-like result required two main assumptions: first that ‘money is not an acceptable instrument’ in this allocation (heavily qualified by a footnote), and secondly that ‘each person’s preferences are assumed to concern solely his own assignment’. They then proceed to develop a clever algorithm which would produce a more market-like result.

This seems to deny significant features which are the essence of the Harvard student housing lottery. The university authorities must have their reasons for using a non-market allocation procedure, so it seems perverse to try and impose a market structure. Also, it is surely unrealistic to assume that in a community of undergraduates they would all be indifferent as to the ‘luck’ of their fellows in housing allocation. The lottery symbolises the wish of the authorities to act fairly towards the student body. The ceremony of drawing lots which is used in some US universities could be seen as part of the process of creating that community feeling.

This good feeling of knowing that others in a group are benefiting could apply equally in isolated mining communities, as well as for ad-hoc social groups like members of an office or a faculty. Whilst individuals might prosper by ruthlessly competing with their colleagues, it leaves a bad taste. Advancement on some merit, mediated by random arbitration should promote self-esteem, willingness to co-operate and ultimately happier individuals.

Lottery: Unpopular so bad for recipients' welfare? The ultimate arbiter of the validity of an allocation system is how it works for the people it is meant to benefit. However technically satisfying a mechanism might appear, if people genuinely don't like it then it has failed. As explained before (Chapter 1) in surveys people do not like the idea of allocation by lottery. A possible reason for this, as Anand (2001) explained is that a lottery deprives the customers of a voice in the procedure. In another survey, Benz & Stutzer (2002) identified the positive effect of 'voice'. From a survey of British workers they were able to identify that having some say in the procedure for setting wages made the workers happier. This might suggest that a lottery, which would deprive customers or employees of an opportunity to haggle, will deliver less 'procedural utility'. However, where it is actually used, random allocation seems fine—witness the US student housing lottery, or how lottery is enthusiastically endorsed by Dutch student opinion. There is need to explain the benefits of random distribution to the potential recipients.

Can random distribution save workers' co-operatives? The prospect for workers' cooperatives is not promising. A paper by Kremer (1997) asks: Why are worker cooperatives so rare? They should, he claims, have the edge over shareholder firms, because of their tax advantages, greater ability to monitor the workers, and because they satisfy peoples' wish to be involved in running their own workplace. There are some well-known examples of coops: Mondragon in Spain and plywood works in the U.S., but these are indeed rare. Worker coops lose their competitive edge, says Kremer because of their democratic structure. The median (in terms of ability and effort) worker will vote to compress wage scales: the extreme example is that of legal partnerships which usually grant equal pay to all partners at the same level. This dulls incentive effects, and leads to a less efficient firm. Kremer also acknowledges that a particular problem with worker control is that it often degenerates into disputes, involving mistrust and envy, a feature which was also present in the Durham pit villages. In the Third World there is the strong social matrix of the community that binds worker cooperatives together. In more open societies, this disputatiousness may be a much more significant factor in the dissolution of worker coops, rather than any

inefficiency due to compression of wage differentials (which Fehr & Schmidt claims would actually be a beneficial characteristic).

To rescue worker cooperatives from their own inevitable inefficiency, and demise in the face of capitalistic competitors, Kremer says they need better institutions. They need less worker control, and some kind of independent arbiter to resolve disputes. This is what the cavil did for the quasi-worker cooperatives in the Durham coalfield. The most significant wage-differentiating decision was taken, not by a manager or a co-worker, but by a lottery.

8.7 Works well for whom: Justice and Fairness in Society

The ‘fairness principle’ proposed by Zajac (1995) that ‘equals be treated equally and unequals unequally in proportion to relevant similarities and differences’.

Converting this into statistical terms means accepting as equivalent anyone whose score is ‘not significantly different’ from another score. This ‘other score’ could be a fixed threshold, as with the 11+ test, or it could be the highest score found in an entry test. There remains a question of judgement which level of certainty should be applied. Conventionally 95% or 99% values are used—that unless you can be 95% certain that the score is *less* than X, then it must be accepted that they are ‘not significantly different’.

Taking the standard statistical approach may produce very wide acceptance bands—in other words, the statistical approach, because of its stringency, may give virtually everyone ‘the benefit of the doubt’. It may be more practical, and make the acceptance of random selection easier if narrower bands were to be used. In an earlier paper (Boyle, 1998) I tried to operationalise this definition, using the old 11+ IQ test as the basis. A standard procedure divided the population into two groups, pass—25% and fail—75%, based on a cut-off score of 110, (IQ scores on a Normal Distributional Mean 100, SD 15). Instead I suggested three groups:

- automatic pass—the top 5%, because they had a very high probability (90%+) of all scoring at least 110, the cut-off value,
- fail—the bottom 55%, because they had were probably (<75%+) all less than 110, and
- a border zone of the middle 40%, above and below the cut-off value.

It was this ‘border zone’ which constituted the ‘relevantly like cases’, and which I proposed should therefore be subject to a graduated lottery.

In a comment on my paper, Barbara Goodwin in her 2005 book, *Justice by Lottery* comments:(p249) ‘I suggest that the logic of Boyle's arguments should entail a lottery for *all* the candidates for a job, or in an examination, even when they appear to be 'significantly different' in terms of their score. The adverse personal circumstances and possible examiners' errors (e.g. mis-transcribing marks) which affect border-zone candidates could apply equally to those below the border zone. The objective of preventing demotivation and disillusion among the less able candidates would be achieved if *all* candidates were given at least some tickets for the graduated lottery (however few tickets). Boyle's assertion that it is salutary for borderline candidates to be aware of the chance-dependent nature of any selection process could equally be extended to those with the least chance of success. In all these respects, the border-zone candidates and those below the border zone are *not significantly different* even though their actual scores were significantly different.’ She continues later:(p250) ‘One objection to Boyle's proposal, then, is that it does not go far enough. Logically, it should extend to all the candidates in the competition below the cut-off point and, arguably, to those above it.’ This may be a valid point, but might be impractical. It comes dangerously close to the modish idea of suggesting that there are no failures only ‘deferred successes’.(Press Association, 2005).

A more systematic approach to deciding the border zone comes from Dodge & Romig who developed the theory of industrial sampling inspection for incoming batches (reviewed by Barnard, 2004). They introduced the concept of ‘Producer’s Risk’ and ‘Consumer’s Risk’. In such a contractual situation it was easy to decide that

these risks should be the same for both parties: That the Producer's risk of having a good batch rejected should be the same as the Consumer's risk of having a bad batch accepted. ('Good' means within acceptable quality level, agreed in advance. These risks are labelled Type I and Type II errors respectively). As a matter of Justice between the parties in educational selection, it could be argued that the same equality of risks should apply.

Earlier in Chapter 5 I quoted Duxbury (1999) on the uses of randomisation in the selection of personnel for employment. He reviewed the experience of Northern Ireland which positively encouraged the use of a lottery to reduce the number of applicants to a manageable size. This advice had been tested in the courts and found acceptable. But Duxbury went further: 'It is worth noting that when arguments in favour of randomized recruitment practices are advanced or accepted, it is almost invariably in relation to low-grade posts which require that employees possess no special skills. Rarely is it argued that shortlists should be determined randomly where there exists an excess of suitably qualified candidates for skilled or professional occupations.' This strikes me as profoundly *unjust*. If random sampling is acceptable for the lowest in society, it must, in fairness, apply equally to the highest. Divine (1976) argued this case in respect of academic appointments.

8.8 Works well overall: Stability, Accountability and Rotation

Stability: By definition, a random distribution means prizes are given away at below market price. The results of the draw are not what would arise from a purely market situation, so it is not Pareto-optimal. This is the economists' crucial *first* criterion for judging any form of allocation, 'the one and only uncontroversial normative argument in economic theory' according to Moulin (1995, p6). It comes as no surprise to find Pareto-optimality used as a criterion to judge allocations. A non-Paretian distribution is open to manipulation.

In the last section I quoted from Hyland & Zeckhauser (1979) who became interested in random distribution because the students at Harvard had found ways to bend the

lottery for housing in their favour. Prior to the draw some discovered that by falsely stating their *second* preference as their first, they had a better chance of at least achieving their second-best option. Once any draw has taken place, prize-winners have even more opportunity to pervert the intentions of the principals. Through trading they may swap their prize with other prize-winners, or they may sell it on. Roth (2002) describes this as ‘unravelling’, and it is clearly a consequence of the lack of Pareto optimality. In the case of the intern matching program described by Roth, both strategic behaviour and post-allocation trading were undesirable.

However in some random allocations, trading is permitted: In the case of the telephone numbers (Chapter 4), this was encouraged. Whether trading is allowed or not, there will always be pressure to engage in it. For Wimbledon tickets there are strict rules which limit the number of applications made, and forbid post-ballot trading. All of these are easy to circumvent, multiple applications under different names take place, and selling on at a premium frequently take place. It is not easy to police such behaviour.

The implication for stability—ensuring that the allocation does not unravel—which so taxes Roth (2002) and other designers of economic mechanisms, can often be dealt with through normal administrative tools. Pareto optimality, too, may be important to economists, but it is only a first, not the last requirement. In practical engineering terms, designs may be appropriate even if *potentially* prone to unravelling:

- there will always be friction in the system. This is the phenomenon of ‘liking what you already have’ or an ‘endowment effect’ (as explained by Huck et al, 2005). In Roth’s intern allocation example, if doctor A doesn’t like Chicago but got Seattle and doctor B preferred but got the opposite, they might be tempted to swap. But this has costs in time and effort, so may not be worthwhile. It is likely that once allocated Seattle, doctor A may begin to grow to like it; ditto with doctor B.
- many engineering mechanisms are designed to be stable only for conditions that are likely to be encountered (No need to earthquake-proof in Swansea). If a structure does show instability such as wobbling, then the first answer is to

add some form of stabilisation. Administrators of allocation processes can do the same, making them more stable so as to avoid unravelling. This can be done by making it costly or awkward to swap post-allocation, or to make information about potential swaps difficult to discover.

Accountability and Trust: If random allocation is to be trusted, then it is vital that the mechanics of the process be independently verified. Many current examples are deficient in this regard: The Green Card lottery is carried out by a program on a mysterious computer in a Federal building in Kentucky. The Dutch medical school entry lottery is entrusted to a lawyer who performs it in his office (so Professor Drenth tells me). In the event of any disagreement, a lottery leaves no audit trail to follow. If the draw had been rigged then this would be difficult to disprove. What is needed is that at the very least the detailed results of the lottery be published so that statisticians could test its reliability. Conducting the draw in public, using physical randomisation devices adds to the credibility. The US Military Draft famously used the ‘gold-fish bowl’ when conducting a public draw. The published results formed the basis of several academic papers (Fienberg, 1971) proving that no bias had occurred.

A public drawing can also be an opportunity. Edgeworth (1890) suggested that randomised degree classifications at Cambridge be decided by ‘a solemn conclave of the Fellows’ convened at a dignified location for the purpose. In San Marino, drawing lots to decide the Capitani Regenti for the following year is carried out in the Basilica during High Mass. (Aubert, 1959). Adding an element of ceremonial to the draw adds dignity, creates a bonding experience and indeed should be an enjoyable experience.

Add rotation? Although not intrinsic to my hypothesis, where possible, adding rotation to random allocation has its attractions. Some prizes are indivisible and cannot be rotated. A place at medical school cannot be handed over half-way through the course. But rotation is sometimes not just possible, but desirable: The fishermen’s *padu* is a rotating lottery. If it had been a one-off final settlement, it would have been far less acceptable, unstable, subject to attack by the losers. Many jobs can be held for a fixed period, just as short-term contract workers are

currently employed. This could be made more widespread, with re-entry through a random allocation process highly likely and fewer job-security worries. This would reflect the standard Athenian practice (Headlam, 1891) of all lottery-chosen posts being held for one year only, without renewal. There are advantages in job rotation: It helps job-holders to avoid becoming captive to the producer interest, which is a particular problem in the public sector; it may prevent employees getting stuck in their ways or becoming corrupt. Against rotation is the obvious insider-outsider distinction, that existing employees build up specific skills valuable to their employers, which would be lost through rotation.

8.9 Conclusions

The *efficiency* case for the use of random distribution for at least part of an allocation process looks like this to those involved:

- For the organisations there are relatively small but significant benefits in cost reduction, and control of illegal behaviour by their agents such as corruption and discrimination. This is unlikely to make them strong advocates of randomisation.
- Agents stand to lose out. On top of the insult to their imagined powers of selection, is the downgrading and maybe elimination of their ‘important’ positions.
- Collectively the recipients have most to gain. Their huge expenditure of effort in rent-seeking and signalling could be directed to more rewarding activity. But individually they are locked in a competitive situation, so a breakout is difficult.
- For society the benefits of reclaiming this wasted effort should be worthwhile, but many of the rent-seeking activities serve another purpose—social control. Keeping young peoples’ noses in textbooks keeps them off the streets. Locking individuals into a competitive struggle with each other similarly keeps them out of other kinds of mischief.

Taken as a whole, the efficiency case for the use of random distribution is valid but not overwhelming.

The *reciprocity and inter-personal comparison* case for random distribution can best be deduced from the examples presented earlier. From experimental economics we know that these feelings are valued. Random distribution supports and demonstrates these higher human values, and by so doing improve people's lot.

For *Justice and Fairness* to prevail there has to be a positive need for them. Because Random Distribution is by its essence a 'fair' procedure then it is fairness manifest. Not only is it fair, but in the drawing ceremony it can be seen to be fair and hence just.

