Chapter Two

Majority rule

Given the probability, even after informed public deliberation, of some level of political disagreement, no theory of democratic collective decision-making can afford to neglect the problem of selecting a decision-rule. If the people are to decide, then the people must vote, but if votes are cast for radically incompatible options, there must be a rule for aggregating votes cast and identifying the option to be implemented on the basis of the votes cast for the different options. At the beginning of the first chapter we looked in an informal and unsystematic way at a few of the available decision-rules in the situation of electing a single individual to a single position, such as the presidency of a country. What decision-rule would aggregate preferences and determine outcomes in the fairest, most rational and maximally democratic fashion? In this chapter we will begin answering this question by examining the claims of that decision-rule which has a time-honoured position in the theory and practice of democracy, and is still taken by many to be almost the self-evident embodiment of the democratic creed, namely the principle of majority rule. We are going to begin by looking at what lies behind the seeming obviousness of the claim that, in the last analysis, if political disagreement persists, majority rule is the only democratically legitimate decision-rule. We will examine the logic of the enduring appeal of the majority principle. Right from the start, certain problems with the majority principle will become evident, and in the second half of the chapter the serious weaknesses of the principle will be explored, weaknesses which are especially damaging if the principle is thought of as a universally applicable single criterion of democratic legitimacy of collective decision-making.

The arguments grounding the majority rule principle begin with the assumption that no matter how complex and multi-dimensional an issue is, an approach to a determinate final decision can always be made via a sequence of binary option choices. To illustrate this we will take two examples, one from the area of technology/ecology and one dealing with a moral social issue. Suppose that it is agreed (at this level we have literal unanimity and hence no decision-rule problems) that a new electrical generating station is required. There *is* disagreement along several different dimensions, the most important being the type of generating technology (nuclear, coal-burning, hydro etc.)

size and location. It would be possible to decide the matter first by deciding the generating technology issue, and to decide that by posing first the question, say; nuclear or non-nuclear. Given a clear majority, for instance, against nuclear power, the non-nuclear possibilities could be dealt with in the same way. When majority preference on the specific technology is determined, we can move on to the other dimensions, dealing with specific options in the same way, until type, size and location have been decided.

It is not, of course, simply technological means-ends matters that can be decided in this way. Suppose it is accepted that society requires some legal framework defining marital rights and obligations. A fundamental question that then must be decided is whether a legally defined marital relationship should be indissoluble or not, in simple terms whether divorce should be legal or not. Again, this is a complex multi-dimensional question; if divorce is allowed, what are the legitimate reasons to be, after how many years of separation, under what conditions or restrictions, with what consequent rights and obligations? But a final determinate decision could be reached through a series of binary choices. Suppose, as might seem logical, it was to be decided first whether marriages were to be indissoluble or not. If not, the issue of reasons for dissolution, conditions etc., could be decided in similar stepwise fashion. At each stage it would seem almost self-evident that the decision should accord with the expressed preferences of the majority. But why? The common-sense reply is that, at least from the perspective of democratic equality, when there is a clear majority-minority split, it is clearly better that the majority will prevails rather than having the *minority* will prevail. We can go a little further than this by noting that the case for implementing the will of the majority (rather than the will of the minority) can be argued both on an aggregate and an individual level.

One aspect of the basic egalitarian principle underlying democracy, Dahl's principle of equal intrinsic worth, is that each person is equally entitled to have their interests taken into consideration, have those interests weighed equally in the balance and, if possible, all are entitled to equal interest satisfaction. From the point of view of equal intrinsic worth, the ideal situation would be one in which everybody's interests were in fact equally satisfied. But that might not always be possible. In situations of serious political disagreement a choice might have to be made between the option favoured by the majority and the option favoured by the minority. If the majority-favoured option is chosen not everyone is satisfied, but *more are satisfied* than not, and that is a closer approximation to the ideal of everybody being satisfied, than would occur if the minority preferred option were to be chosen. It is unfortunate in this instance for the members of the minority but, from an aggregate social point of view, it would surely be more unfortunate if more people were unsatisfied rather than less, and from a democratic perspective it would surely be even more unfortunate still if the will of a minority always prevailed over the will of the majority.

The argument operating on the level of the individual is, if anything, even more persuasive. We start, not by looking at outcomes and how they might affect people's interests or accord with their preferences, but by highlighting the basic principle of democratic political equality interpreted as implying that, in determining authoritatively binding decisions, each individual is entitled to an equal share of power. The entitlement to equality of power is represented by everyone's having a vote and every vote counting equally – as the slogan has it: "one person, one vote; one vote, one weight". But literally rather than metaphorically, what does "one weight" mean? The equality of each vote is ensured by assigning a numerical value to each vote, and evidently this is an equal number, the same value. For most purposes this can be the simplest cardinal number, viz. the number one. The number representing each vote is then added to the number representing the votes cast for the same options. The aggregated numbers have only a single characteristic in terms of which they can be compared – whether they are quantitatively greater or less. Not only would it seem perverse to operate a rule which selected the option represented by a lesser number, but, in so far as the numbers are meant to represent aggregated individual *equal* powers, such a procedure would offend against individual equality. Voting procedures do, of course, operate by means of counting, adding and arithmetic comparison, but the argument can be made without the mathematics. If the will of the minority prevails over the will of the majority, the only way that this could be possible is if individuals constituting the minority are given greater power than individuals constituting the majority. And that is directly contrary to the democratic political equality of each individual. From both the aggregate and the individual perspective, it does look as though when there is political disagreement, when votes are cast for different options "the most preferred option is the one preferred by most" and any other rule will be contrary to political equality. (Dahl, preface to democratic theory, pp.).

Persuasive as these arguments might seem, democratic theorists have long known of situations in which their validity appears weakened, if not completely undermined. In fact, when Dahl constructed his model of "populist democracy", from which we quoted the equation of "most preferred" and "preferred by most", it was not to endorse it, but rather to confront it with difficulties, the first of which was the problem of intensity of preference. Suppose we had a political community deciding between two alternatives, A and B. Although there is a majority in favour of A over B, that majority is rather slender, say 51% for A as against 49% for B. Furthermore, although our 51% do indeed prefer A to B, they see B as a very close second best, whereas the 49% who prefer B believe (let us assume, to give an edge to the situation, correctly) that B is of the utmost importance to their vital interests, whereas A would be totally disastrous for them. One response to this kind of scenario might be that the 51% majority, given that they see B as a nearly indistinguishable second best, should really be prepared to concede B to the 49%. The real point here, however, is not what the groups in question should do, but whether in this kind of situation it still seems evident that the most preferred, from an aggregate point of view, is the option preferred by most. A very plausible case could be made for the conclusion that B, the top preference of 49% and the close second preference of the rest of the community, is "more preferred" than A, which, though it has a 51% top preference support, is ranked at absolute zero by almost half the community. Although one could not pretend that such a conclusion was strictly "provable", given that what we mean by "most preferred on the aggregate level" is probably a fairly indeterminate concept, it gets some support when we look at the two types of argument cited above in favour of the majority principle. Implementing B would secure for the 49% something they intensely preferred *and* it would be giving to the rest of the community something that they considered to be a close second best. Even from the perspective of the politically equal power of each individual, it might be argued that if B were to be chosen because it was not only the top preference of 49% but also seen as second best by 51%, members of the 51% were still exercising power.

There are no doubt problems with these considerations. In the first place, simple intensity of preference may not in itself be a factor relevant to selecting a decision-rule.

It is often postulated that intense preference for A over B *might* derive from a total lack of serious thought about the possible values of B, when those giving more careful consideration might concede some positive merit to options that may not be their first preference. In such a situation, taking intensity into account might be simply rewarding thoughtlessness and bigotry. Secondly, even if we could distinguish justified from unjustified intensity, it is difficult to imagine an institutionalised mechanism that would even approximately measure it with any reliability. In constructing the example above, we postulated by definition that B was an almost indistinguishable second best for the 51% who voted for A. But how could we possibly know this from a vote?

A final consideration, introducing a small amount of formal probabilistic analysis, is relevant here. The arguments we have been analysing are motivated by what we could call a democratic sympathy for the minority losers. The defeated minority do not get what they wanted, and this seems all the more unfortunate when the winning majority only mildly preferred their top option and would have been reasonably satisfied with what the minority voted for. It can be argued, however, that this is too episodic a way of looking at political decision-making, majority winners and minority losers. In the normal course of events, a political decision in the real world is one of a sequence of such decisions, and in that context, two basic theorems of probability theory become relevant to the assessment of (particularly) the plight of the losers in any one instance. Firstly, given that, by definition, if a decision is made by a simple majority, then “winners” outnumber “losers”, it follows that for any individual selected randomly there will be a greater chance of that individual being in the winning group rather than the losing group. Secondly, if a type of event has an assignable probability, in an ongoing sequence the actual frequency of the event will almost certainly approximate to the assigned probability. Putting these two theorems together, it follows that over an extended series of decisions, any individual is more likely to be a winner than a loser more of the time. And this is true equally of *each* decision-maker. Generalising the conclusion, each decision-maker is more likely to be a winner rather than a loser equally with each other decision-maker. And this looks like a good example of sequential proportional political equality, with an equal greater probability for each of winning rather than losing. Given that if some specific group were to win more times than the average, this would be at the expense of others, the equal proportional effectiveness of each is as high as it could be compatibly with the equal degree of effectiveness of others. Everybody is equally reasonably advantaged and nobody is more disadvantaged than anyone else. If it is argued that winning in trivial cases is no compensation for losing on big issues, it should be pointed out that the conclusions about winning in general can be applied to "winning in big cases". The probabilities are the same in specific cases as in the general case.

Important as these considerations are, they do have their limitations when applied to the real world. To begin with, the theorem about actual frequency approximation to assigned probability needs fairly large numbers of cases for anything like accuracy, numbers over 1,000. If we sub-divide our cases into, say, big issues, moderately important and relatively trivial, it is not at all obvious that in real-life situations there would be enough "big issues" for the implications of the theorems to be significantly applicable. Secondly, the theorems only work if the events in the sequence are "independent"; to be independent, the probability of one type of event at any given time should not be affected by the occurrence of a particular type of event earlier in the sequence. Now this is what you would normally expect with series of events like coin tosses. If tails is the result of the first toss, that is not likely to affect the results of later tosses and, hence, the sequence will be of independent events. Political decisions *may* be like this, but in many instances they are not. Take an extreme example to illustrate what is at stake. Suppose a decision is being taken on a proposal to restrict the franchise, that would determine who is entitled to vote. Imagine that a group currently entitled to vote loses and is disenfranchised. Subsequent decisions are obviously not "independent events". This is admittedly an extreme case, but there are many more subtle ways in which the effects of earlier decisions can themselves affect the likely effectiveness of different groups of people in later decisions.

Finally, preferences over the possible outcomes of the sequence of decisions have themselves to be independent of each other and distributed randomly in the group of decision-makers. To illustrate what this means, let us take a negative example where, that is, preferences over outcomes are not random or independent. Suppose a group is divided into two sub-groups, call them the Reds and the Blues; however these are defined we assume that if one Red is in favour of X, all Reds will be in favour of it and all Blues will be against it. Clearly, preferences are not distributed randomly in the group, they are distributed systematically according as one is Red or Blue. The implications for a sequence of choices are equally plain; preferences in the sequence are not independent because if Reds had one preference in the first issue area and one Red had a particular preference in another issue area, so will *all* Reds. Given our assumptions, the Reds always agree with each other and so do the Blues, disagreeing with the Reds. Now it is evident that if, say, the Reds are more numerous than the Blues, the Reds will always outvote the Blues. The conditions for the application of our probability theorems do not apply. Whatever about *abstract* probabilities, being a Blue loser in one instance will mean being a Blue loser in other instances and there will, of course, be serious deviation from anything like equal effectiveness over the community as a whole. Quite apart from the probability theorems, this situation of a "permanent minority" (and hence, a permanent majority) has been identified by many (Lively, Dahl, Hyland) as creating an even greater problem for simple majority rule than the intensities of preference problem.

Permanent minorities in the real world are usually the result of conflictual religious/cultural/political/ethnic/linguistic divisions, often stacked on top of each other. That such situations generate major problems for democratic political equality is obvious; in a sequence of simple majority votes, the permanent majority (nearly) always wins, with the permanent minority nearly always losing. But if the same individuals always lose, not only do they not enjoy equal political power, they enjoy *no* political power at all, and this seems to contravene the very definition of democracy as political equality. Some theorists see the nearly total lack of impact on political outcomes that occurs in the case of really persisting permanent minorities, such as Catholic nationalists in Northern Ireland, as being so serious as to amount to effective disenfranchisement. Formally, such minorities have full entitlement to participate in the political process, but their position as minorities has, in terms of affecting outcomes, the same effect as disenfranchisement. Countering such claims as wild exaggeration, other theorists such as Brian Barry argue that they are based on a misunderstanding of the notion of power. According to Barry, the concept of power is the concept of the potential to have an impact. If several individuals have the same equal potential, but use it in opposing directions, and if the combined potential of one group is greater (because the group is more numerous) than the combined potential of the other group, this does *not* imply that each member of the small group had no power. It simply implies, as stated, that the combined power of the first group was greater.

On the one hand, then, we can argue that when a simple majority rule principle is used, individuals who find themselves in the minority, and particularly individuals who find themselves *permanently* in the minority, can complain that not only do they not have equal political power, they do not exercise, as we put it above, any power at all. On the other hand, it can be claimed that *individuals* are the ultimate possessors of power, and if each individual does have equal potential, but those equal potentials are combined in different numbers and in opposite directions so that some will be winners and others losers; that, however, is a *result* of individuals having equal power, not something that contradicts it. Hence it is perfectly democratic, though perhaps an unfortunate consequence for the minority losers of the realities of political disagreement.

Is there any way out of this impasse? One conceptually sophisticated treatment of the subject, that does advance the argument, is that provided by Jack Lively in his 1975 book *Democracy*. The basis of Lively's approach is his distinction between *prospective* and *retrospective* political equality. Prospective political equality is initial formal equality of opportunity, where no participant in the decision-making process suffers any institutionally derived disadvantages. It is what we think of as being guaranteed by "one person, one vote: one vote, one weight", each "anonymous" vote being treated equally. Retrospective equality is achieved (if it *is* achieved) when, after an outcome has been decided, it could be said of all participants that their vote contributed equally to the determination of the outcome. Prospective political equality is relatively easy to guarantee, being dependent only on anonymous and fair procedures. A procedurally fair lottery would satisfy the conditions of prospective equality and if the lottery winner's vote on some issue were to decide the matter, we would have an illustration of maximal prospective equality without any retrospective equality at all, since the votes of everyone other than the single lottery winner would play no role in determining the outcome.

After introducing the distinction between the two types of political equality, Lively goes on to make two crucial points. Ideally, we would like to achieve complete prospective and retrospective equality, *but*, though as we saw prospective equality is easily realised, full retrospective equality *can never be* institutionally guaranteed. We can appreciate the force of Lively's argument by looking at the matter from two angles. Firstly, if there *is* irreconcilable political disagreement and votes are cast for incompatible outcomes, the votes cast against the outcome finally chosen and implemented are, in this case, ineffective; they do not carry retrospective power, there is not full retrospective equality amongst the voters. Secondly, consider a case where there is retrospective equality. Suppose that a group uses a literal unanimity rule; every single vote must be for a proposal before that proposal can be adopted and implemented. Suppose further that in some particular case there is complete unanimity. That *would* be an instance of complete retrospective equality; it would be true of each voter that had his/her vote been different the outcome would have been different and, hence, each had really determined equally the actual outcome. This kind of scenario only illustrates, however, the fragility of total retrospective equality achieved through a unanimity rule. The reason for this is that from the situation as described it is evident that if even a single voter changed her/his vote, with all the other votes remaining unchanged, the single vote would determine the outcome, and not all the others, which would be ineffective. In fact (a point we will return to shortly) if we attempted to guarantee retrospective equality by requiring unanimity, we actually increase the probability of the outcome being determined by a dissenting minority, possibly a minority of one.

Given the impossibility of institutionally guaranteeing full retrospective equality, what we should do, according to Lively, is to select a voting rule or vote aggregation procedure that guarantees the closest approximation to our goal. To achieve this we need to analyse the likely levels of retrospective power of different types of rule, just as we analysed the unanimity rule in the preceding paragraph. Lively claims that, though there can be many different voting rules, they all fall into one of four basic categories, two majority rules and two minority rules, viz. simple and stipulated majority rules and simple and stipulated minority rules. If at its simplest a choice is between two options, when votes have been cast and all votes for the same option aggregated, we have two bundles. The minority rule procedures would, in the case of the simple minority, specify that the option with the lesser number of votes was to be selected; in the case of a stipulated minority criterion the rule would specify a particular minority, such as twenty per cent, and then state: implement the option that has twenty per cent or less votes. Of course the minority rules are intuitively undemocratic. In fact Lively's analysis tells us *why* they are undemocratic; they institutionally guarantee a high deviation from retrospective equality. Furthermore, Lively is perfectly well aware that unless voters were deceived about the voting rule under which they were operating, rational voters would vote *against* the proposal that they favoured, hoping to push the vote across the specified minority threshold. The minority procedures would then, in fact, become equivalent to their majority rule mirror images. Therefore we need not give much further consideration to the minority rules.

The simple and stipulated majority rules are, of course, well known. The simple majority rule specifies the selection of an alternative if it achieves at least 50% of the vote plus one. Whereas stipulated (sometimes called qualified or super-majority rules) majority systems identify some percentage higher than 50% which an alternative must achieve before it is selected. It is worth noting in passing that the unanimity rule is simply the limiting case of stipulated majority, specifying the percentage to be achieved as 100%. Stipulated majority rules are often used in practice, particularly where, as in changes to constitutions, it is thought desirable that a large (rather than a bare, simple) majority should be positively in favour of the alternative to be implemented. And at first glance, stipulated majority rules do seem to achieve this, and in achieving it, guarantee high levels of retrospective equality. If we required, say, a minimum of 75% before a proposal was accepted, then whatever was selected for implementation would have at least 75% support, at least 75% of the votes would be contributing to the actual outcome and this is a closer approximation to full retrospective political equality than if only 52% of the votes cast determined the outcome. To believe that high stipulated majority thresholds guarantee consistently higher levels of retrospective effectiveness would, however, be a mistake. As we saw in the case of the unanimity rule which sets the threshold at 100%, if the threshold is reached, retrospective equality is achieved, but the probability is that the threshold will not be reached, and in that instance it could be that the votes of a tiny minority of dissenters who prefer the status quo determine the outcome. Even with a threshold of 75% it could easily happen that a very large majority of, say, 74% was denied any effectiveness by a relatively small 26% minority. Being foolish enough to think that we can guarantee closer and closer approximations to full retrospective equality by requiring higher and higher stipulated majority thresholds fails to take into account the fact that the higher the threshold the less likely it is to be achieved and, hence, the more likely it is that a progressively smaller minority of votes determines the outcome. *Only the simple majority system guarantees that the outcome has more voters in favour of it than against it*. If the majority happens to be large, well and good; a close approximation to full retrospective equality occurs; but, under the simple majority system, a smallish dissenting minority cannot be effective at the expense of a large majority. For a dissenting minority to become effective it must become a majority. Hence, as we said above, the group determining the outcome will always be larger, there will always be more retrospective equality rather than less.

There was, in fact, a much earlier formalised version of basically the same set of considerations in Kenneth O. May's famous 1952 Econometrica paper *A set of Independent Necessary and Sufficient Conditions for Simple Majority Decision*. Referring to Kenneth J. Arrow's seminal work *Social Choice and Individual Values*, May says: " the problem of the relation between group decision and individual preferences has been stated by Kenneth J. Arrow in terms of a 'social welfare function' that gives group choice as a function of the preferences of the individuals making up the group. One of the conditions he puts on this function is that group choice concerning a set of alternatives must depend only upon individual preferences concerning alternatives in that set. In particular, group choice in the presence of just two alternatives depends only upon individual preferences with respect to this pair of alternatives. Since it follows that the pattern of group choice may be built up if we know the group preference for each pair of alternatives, the case reduces to the problem of two alternatives." (*Econometrica*, October 1952, p. 680)

May proceeds to lay down four "weak" conditions for "group decision-rule", meaning, presumably, that these are conditions that one would, in normal circumstances, require of any putatively rational and "democratic" procedures. The four conditions are: decisiveness, anonymity, neutrality and positive responsiveness. Decisiveness does not quite mean what we might usually take it to mean; it does not require that the procedure *select* one or other of the alternatives as the outcome to be implemented. What it requires is that a definite assessment results, even if the assessment is that each alternative has equal support and, hence, there is "group indifference". Anonymity is the requirement that each vote contributes equally, irrespective of *whose* vote is being considered; it is our well-known condition of one vote, one weight. The neutrality condition is an equal chance condition as applying to the alternatives being decided on. The procedure used must not favour one of the alternatives rather than the other. It should be noted that these first three conditions are, indeed, very weak in the sense that many voting rules other than the simple majority rule can easily satisfy them, though if the alternatives are a positive proposal or maintenance of the status quo and a stipulated majority is required for the positive proposal, then the procedure favours the status quo option, as we saw above. It is really the positive responsiveness condition that drives the argument to the simple majority conclusion.

Strong positive responsiveness is defined as follows: a procedure is strongly positively responsive if, when there is a tie between two alternatives X and Y and then a single vote migrates from, say, Y to X, the procedure selects the alternative for which the total number of votes has positively increased. It does seem absolutely obvious that, in general, there could be little logic, never mind democratic logic, in having a procedure that is negatively responsive, one, that is, which selects the alternative for which support is *decreasing*. Acceptance of the condition of strong positive responsiveness is what lies behind the rejection of the stipulated majority procedures. With, for example, a 75% stipulated threshold, particularly where X is the positive proposal and Y the status quo, votes could continue migrating from Y to X, but Y would continue being selected right down to the stipulated level of 25% plus one vote. The procedure would not be responding positively to the positive increases in the number of votes for X. There does seem to be an ineluctable logic to the argument that, if a tie between X and Y results in an impasse in which neither is selected, it would be irrational for a procedure to select Y rather than X if the only change in the circumstances was a *decrease* in support for Y and an increase in the support for X.

The ineluctability of the logic, however, is seriously undermined when we reveal the extremely shaky nature of the background assumptions of arguments such as Lively's or May's. The dubious assumptions in question have to do with the nature of political conflict, the nature of the decision situation, the type of preferences relevant to determining the outcome and the type of decision procedures available for the elicitation of these preferences. We could start our exploration of these matters by looking at an assumption that May makes explicit and that seems implicit in Lively's analysis as well, namely, that decision situations are fundamentally binary option choices. Now suppose we take the kind of situation that voters are probably most familiar with in real-life democracies, namely, choosing a single person for a single position from a list of candidates; to make things a little, but not too, complex, we will assume that we have five candidates, A, B, C, D and E. The voters are asked to indicate the candidate of their choice by putting an X next to the name of their favoured candidate. The first obvious point is that we are not dealing here with a straightforward binary option choice. Secondly, because it is not a binary option choice, the simple majority rule of selecting the option with at least 50% of the votes plus one is not necessarily applicable. It may happen that no candidate gets anywhere near 50%. To examine the kind of difficulties that arise in such very normal situations when we try to apply the Lively criterion of maximising retrospective equality, imagine support for the different candidates as follows:

A	B	C	D	E
22%	21%	19%	19%	19%

Table 1

What would the logic of simple majoritarianism imply in these slightly altered non-binary option circumstances? Anyone familiar primarily with a plurality "first past the post" voting system such as is used in U.K. general elections might be inclined to say that in the absence of the assurance of a single candidate winning an overall majority, the next best thing is to select the candidate with the largest number of votes. This, it might be argued, is not a very large deviation from the higher retrospective equality achievement of the majority principle in the binary option case. A's winning margin is admittedly narrow, but one or other of the candidates has to be selected and if it is not to be A it could hardly be B, C, D or E, for their case is clearly worse. We could even apply May's positive responsiveness criterion to this situation. Suppose a tie between all candidates, and then the migration of a single vote so that one candidate's vote increases, while support for all the others decreases or remains the same, the candidate whose vote increases over the votes for the others should be selected. Any other rule would be perverse, irrational and undemocratic. Such an argument, however, would be much too hasty, the obviousness of the conclusion being an illusion generated by a feature of the simple kind of plurality system that is, in fact, quite arbitrary. The feature in question concerns the type of preference relevant to the determination of the outcome and the preferences elicited by the voting procedure.

Our voters in the above example were asked to select their most preferred candidate. Suppose, however, each was also asked which candidate they ranked last, which they *least* preferred. To keep things simple, imagine that the A supporters all disliked B the most, but that B, C, D and E supporters all disliked A the most. We could expand Table 1 as follows:

	A	**B**	**C**	**D**	**E**
First	**22%**	**21%**	**19%**	**19%**	**19%**
	B	**A**	**A**	**A**	**A**
Last	**22%**	**21%**	**19%**	**19%**	**19%**

Table 2

It follows that though A has a 22% top preference support, a massive 78% rank A as their *least* favoured candidate. People used to the simple plurality, first past the post system might, initially, think that asking voters to indicate the least preferred candidate was itself quite arbitrary. One might also raise the point of whether a negative vote against should count as much as a positive vote for, and in addition, how all this information was to be weighted and aggregated. A very simple point needs to be made immediately. Lively claimed for the simple majority procedure that at least it guaranteed that the option implemented had more in favour than against. In the kind of situation described in Table 2, the election of A because A had the highest plurality vote would *not* satisfy Lively's criterion; 78% would be against the option implemented, with only 22% in favour. One very simple way of incorporating the "least preferred" information into the decision-making procedure, and of avoiding the "metaphysical" question of the relative importance of a positive and negative vote, would be to use a two-round voting system such as, for example, is used in the French presidential elections. In the first round, voters are asked to indicate their first choice. If no candidate achieves an overall majority, a second round of voting (a run-off) takes place between the two candidates who received the highest number of votes in the first round. In our example this would be A and B. Assuming that Table 2 describes voters' preferences accurately, and that preferences do not change between round one and round two, we can compute the result of round two quite simply; A (the plurality winner) would lose by 22% as against a *positive* vote of 78% for B. One conclusion should be drawn immediately from this analysis; accepting Lively's criterion for a more or less democratic procedure as consisting in close approximation to full retrospective equality, the plurality system, as compared to the run-off system, is extremely undemocratic. (We will be noting drawbacks with the run-off system itself in the next chapter).

The commonsense support for the plurality system does not assume, of course, that choices are always single binary-option choices. What it does assume is worth stating explicitly and exploring critically in more detail. It assumes (a) that the choice is always a choice between one of a finite set of options and (b) that the choice is to be made by eliciting voters' first preference choices for one or other member of that set. The outcome is then one of these first preference options. This itself is based on a false conceptualisation of the nature of political conflict. Almost all political disagreement is zero-sum. Informally, what that means is that the more one side gets of what it wants, the less the other side gets. Where there are opposing opinions over the distribution of a quantifiable resource, such as money, the more-or-less can be strictly quantified, with the consequence that when the positive gains are combined with the negative losses they sum to zero. A very specific kind of zero-sum conflict, in fact the limiting case of such conflict, is called a "winner-takes-all" situation. As the name indicates, winner-takes-all situations are those in which one side gets everything that it wants, while those with opposing preferences get nothing. Even as modified by second-round voting, the selection of a candidate for a single position as described above still results in a winner-takes-all outcome. The modification to the straight plurality procedure consisted in eliciting second preference choices of those whose first preference choices were in the lowest minority. But at the end of the day, those who voted for B got B; those voting for A got nothing. Now it might be thought that at least in that kind of situation, voting to select representatives, say, the situation *is* a winner-takes-all one and the voting rule will reflect this. That would be to take too narrow a view, as the following example will illustrate.

A group of 100,000 people are to elect ten representatives to a legislative assembly. The representatives will come from one of three political parties, A, B and C. Firstly, the supporters of these parties are extremely "partisan" and ideally the supporters of any one party would like to see all the ten representatives coming from their favoured party. Furthermore, the voting procedure asks each voter the same simple question: from which political party, A, B or C, would you prefer the ten representatives to come? Obviously there are only three possible outcomes, if the choice is to be determined by the number of votes cast for each option with the option supported by most being implemented. Suppose that in fact we do not face a less-than-absolute majority problem in that support for the parties is as follows:

A	**B**	**C**
60%	**30%**	**10%**

The A's have it by absolute majority and the logic of simple majority positive responsiveness would indicate implementing their preference rather than those of the B's or C's. This derives, however, not from the nature of the choice situation itself but from the question asked and the voting procedure. Far from there being only three possible outcomes, each one corresponding to the top preference of one or the other of the groups there are, in fact, **66 possible outcomes consisting of the distribution of the members of the three parties over the ten seats.** What is more, there are numerous voting procedures that would select outcomes other than one of the winner-takes-all top preferences of each group outcome. There is, for example, the simple proportionality rule: distribute seats to parties in proportion to the number of votes cast for each. This would be particularly easy to apply with our hypothetical percentages, the proportion 6:3:1 resulting in party A getting 6 seats, B getting 3 and C 1 seat. That particular combination may not have been the top preference of any individual, but it is the result of weighting each individual vote equally in a procedure that delivers final effectiveness proportional to group size, ensuring *equal effectiveness for each individual vote*.[1] Some sort of procedure like this can very closely approximate (and in this case

[1] Full equal effectiveness in this case is the result of the particular percentages postulated. A generally applicable procedure would have to specify a quota and have subsidiary rules dealing with groups of

actually reach) full retrospective equality even when there is sharp partisan disagreement. Lively's criterion is a maximisation of the majority in a winner-takes-all conflict. What the above example suggests is an alternative criterion that we call the principle of generic proportionality. What this implies in practice is the construction of voting systems and procedures that will confer actual effectiveness on as many individual votes as possible, resulting in the determination of the outcome in a manner that reflects the proportional support for different options. We call it "generic" because only sometimes can the proportionality be actually quantitative, qualitatively accurate and mechanically achievable by selecting the outcome on the basis of a mathematical function of the individual votes cast. But it is the generic principle of proportionality that is important as undermining the claim of a majoritarian winner-takes-all principle to be the unique self-evident criterion of democratic legitimacy.

Given that in these early chapters we are primarily concerned with voting procedures as such, rather than with electoral systems, we will explore further the difficulties facing the simple majority principle by considering a final example in which a group choice is being made directly on a matter of substance: the distribution of a resource such as money. Imagine that a government, finding itself in the fortunate position of having a £100,000,000 budget surplus, takes the unusual step of asking the people to decide what to do with this surplus. Suppose that just prior to the generation of the surplus there were two complaints in particular being loudly voiced: many were complaining about the disastrous underfunding of the health service and, quite unrelatedly, many were complaining that the government tax on alcohol was far too high. These being the currently salient issues, the government offers the people a choice: will we use the money to upgrade the health service or will we return the money to the people by lowering the tax on alcohol? Readers who have been following the argument so far will immediately conclude that though the people in our example have been *offered* a binary option choice and when a group has, of necessity, to make a binary option choice there is greater democratic logic to selecting the option that has more rather than less support, the decision situation is not itself ineluctably binary. Rather than there being, in the

---

votes exceeding or not reaching quotas. And very small minorities might still be excluded from final effectiveness.

nature of the case, just two possible outcomes, if we allowed alternative combinations of the money, down to the smallest unit of currency (i.e. one penny) the number of possible outcomes is ten thousand million, ranging, obviously, from all the money being spent in one of the areas and none in the other through all possible divided distributions. An even more important point is that, not only might some outcome other than one of the winner-takes-all ones presented represent a more accurate synoptic picture of the group's preferences, it is perfectly possible that neither of the winner-takes-all options is *anybody's* top preference. Not only are the people *forced* into a binary choice by the way the question is posed rather than by the nature of the case, but the particular two options may be nowhere near anybody's ideal outcome. Finally, although ten thousand million possible outcomes is large enough, it pales into insignificance when we introduce what, almost certainly, would be a further complexity of such a situation in the real world. The possible outcomes as described in the example range over the different distributions of the whole £100,000,000 between the health system and the alcohol tax reduction. In the real world however, some people may prefer that some of the surplus be spent in other areas, such as improving the educational system, the road network, childcare facilities etc. Contemporary formal theorists of politics conceptualise this in terms of a multi-dimensional "policy space" in which there are an *indefinite number* of points, each representing some possible combination of, in our example, financial resource distributions. We will return to this astronomical multi-dimensional multiplicity later in this chapter; the crucial point here is that it is almost never the case in situations of real-world political disagreement that a group is *inevitably* faced with a single-dimensional binary option choice. This applies to most political issues, not just to economic distributional ones. Take the example of divorce legislation cited at the beginning of this chapter. What has to be considered is not at all some single binary choice – divorce or no divorce – but an enormous multiplicity of possibilities, defined along several dimensions such as allowable cause, procedure for establishing cause, conditions (mutual consent? years of separation?), consequent rights and responsibilities. While the "distances" between the possible positions on these issues are not arithmetically commensurable, as differential amounts of money would be (this was why the terminology of "generic" rather than mathematical proportionality was introduced earlier), it is not the case that there is some small finite number of possible outcomes, each being the top preference of some group, with the consequence that one or other of these winner-takes-all options must be selected. In the vast majority of situations of political disagreement, matters are much more complex. The maximisation of retrospective equality should be interpreted as the maximisation of generic proportionality. Sometimes, as in electing representatives of distributing economic resources, very close approximations to full proportionality can be guaranteed by relatively simple institutional procedures. Often, however, the institutional design of procedures intended to effectively empower in appropriate ways individuals and groups who might otherwise be rendered powerless by exclusionary majoritarian procedures turns out to be an extremely complex task that can never be perfectly fulfilled. The essential point of the present analysis, however, has been to demonstrate the fatal weakness of those arguments that were meant to establish the unique democratic legitimacy of the majority rule principle.

Before leaving the majority rule principle, there are two further arguments that we should examine, both also seriously weakening the case for that principle. Both have to do with what we could call radical indeterminacies in the very notion of "the will of the majority". If we artificially confined ourselves to thinking about a strict binary option choice between X and Y, the idea of what *the* majority preference was would be perfectly clear (except in the case of a tie). As we have seen repeatedly, however, once the number of possible options to be decided between is more that two it does not make determinate sense to say: do what the majority decided. Even if we are thinking along general majoritarian lines, we need to ask to which majority winner, according to which voting rule, are we referring. We have not as yet been given a systematic account of the range of voting and electoral rules normally used, but enough has already been said to demonstrate that, as was the case with our very first example at the beginning of Chapter One, different rules can quite easily give different majority winners. Hence, unless we have specified a particular voting rule, talking about *the* will of the majority is totally misleading. This would make the task of a democratic audit of voting rules a matter of urgency, even if one were still operating under the spell of the myth of majority rule.

The final argument introduces considerations of much more profound significance; it derives from the problems of preference aggregation created by "cyclical" preferences.

It was the Marquis de Condorcet who first discovered the possibility of voting cycles. The problem was rediscovered by Charles Dodgson (ref.) and, most famously, by Kenneth Arrow in his "Impossibility Theorem". The present argument, however, is based on the relatively recent discovery, by theorists such as Plott and McKelvey, of the chronic pervasiveness of cyclical preferences. We can begin to explain what is at stake here by examining the very simplest voting cycle. Suppose that three voters X, Y and Z are to choose among three options A, B and C. In order to get a complete comparative assessment, a sequence of pairwise votes are taken: A against B, B against C, and A against C, each of the options, that is, being tested against all of the others. Let us assume that X, Y and Z rank the options in order of preference as follows:

X	**Y**	**Z**
A	**B**	**C**
B	**C**	**A**
C	**A**	**B**

Table 3

An examination of the above set of preferences will show that in a pairwise vote between A and B, A will win with a 2:1 majority, X and Z clearly preferring A to B. In a vote between B and C, B would in that instance win over C with a 2:1 majority. So A beats B and B beats C. It looks as though C has to be the clear loser. But what happens if A is compared directly to C? If we look at the table again we will see that the somewhat surprising answer is that C will beat A with a 2:1 majority. Using Arrow's technical terminology for such situations, aggregated social preferences are not transitive. In the case of an individual person, we would expect that if A was preferred to B, and B was preferred to A, then A would be preferred to C, just as if A was larger than B and B larger than C, A would be larger than C. but, as Table 3 demonstrates, this type of transitivity does not necessarily occur when we are aggregating the preferences of a group of people. How such a situation creates a serious problem for majoritarian legitimacy can be easily explained. Supposing we were to claim that B should be implemented, given that B was preferred to C and C to A. An immediate majoritarian objection is possible; we cannot implement B when there is a clear majority for A against B. So, why not implement A? The majoritarian legitimacy of selecting A is obviously undermined by the fact that a clear majority prefers C to A. but selecting C could not be right because a majority prefers B. And we are back where we started. Whichever option we think of selecting, there is another option for which there is a majority preference. As Iain McLean (Held & Pollitt, *New Forms of Democracy*) said "we might think that the possibility of voting cycles is just a fairly trivial curiosity, generating a not very difficult to deal with practical problem". And on one level this is true. In an even-numbered group voting on two alternatives a *tie* is possible; if a tie actually occurs, this is taken to be a case of aggregate indifference; there is equal support for each option. Especially when, as in our example, the majority in each case is actually equal, we could be justified in interpreting the voting cycle as an indication of aggregate social indifference. Given equal support for each option, the practical problem could be resolved by random selection. One problem in differentiating cyclical majorities from even ties between two options stems from the issue of probability. Obviously, the possibility of a tie depends upon exactly even numbers, and the probability of a tie decreases with the increase in the absolute number of voters. The situation with voting cycles is the reverse; the probability increases with the number of voters, and increases dramatically with the number of alternatives. With an indefinite number of votes and seven options, Iain McLean gives the probability of voting cycles as 36.9%. Still, we might say, keeping the social indifference interpretation in mind, having to select randomly in XXXX cases might be a little worrying, but surely not disastrous. [2] The really serious threat to majoritarian legitimacy derives not from the probable frequency of actual voting cycles when groups are asked to vote on a largish number of options, but from the demonstrable pervasiveness of cyclical preferences once preferences are conceptualised as points in multi-dimensional policy space.

[2] If the reader is wondering why voting cycles seem not to occur in real life with anything like this predicted probability, one reason is that many actual voting procedures are not as rigorous as Condorcet's pairwise comparison in eliciting all the relevant information. Suppose, for example, with preferences as in Table 3, an elimination procedure was used. In a spirit of neutrality between options, the first head-to-head vote is decided by random selection. Suppose the "draw" pitted B against C; B would obviously win that round. C would be eliminated. In the next round A would beat B; and the fact that there is a potential cycle with A being beaten by C would never emerge.

As an example take the case, favoured as an illustration, of the potential differential distributions of money between guns and butter. The possible allocations of different amounts of money between the two areas can be represented by a simple two-dimensional graph:

(graph goes in here)

There are an indefinite number of points on this graph, each representing a different allocation of resources between military spending and subsidising individual consumption. Assume a minimum of three possible voters, A, B and C, each having preferences for different allocations, as in our illustration. A's "ideal" preference point represents a rather large amount spent equally on both guns and butter; B's preference is for an equal distribution of a smaller amount, while C would like more to be spent on butter than either A or B, and less on guns than A. Imagine a point P^0, equidistant from the ideal preference points of A, B and C. One might naively think that, being equidistant, this would be an equilibrium point that all would agree on. This would be illusory, assuming that each would ideally prefer a combination more closely approximating her/his ideal. Evidently, there are numerous (in fact indefinitely many) points that are, for example, closer to B and C than P^0. Just as an example, though there are many more, any point on the line P^0-P^1 would be closer to B and C than P^0, though farther away from A. Given binary choice between P^0 and any of those other options, a majority of B and C would vote for the alternative, with a minority A voting against. But there is also an indefinite number of *other* points closer to A and B than P^0, for any of which as against P^0 there would be a majority coalition of A and B; and similarly for A and C. What we are dealing with here is not the *probability* that some structures of preferences might be cyclical, but the certainty that with most ordinary differences of preferences, if there is a majority for any point in policy space, there will always be another point for which there would also be a majority, ad infinitum.

If all this seems a little too abstract, consider another example, constructed this time without graphs and policy space. Rational choice analyses of these types of situation often assume that voters are straightforwardly self-interested; this assumption is not, however, necessary and our example, adapted from that given by Brennan and Lomasky, will posit three groups who are wholly altruistic, though they radically disagree with each other over the appropriate allocation of resources. The cause to which our three groups are altruistically committed is the reduction of Third World poverty. Our first group, A, adopts a strategy of Christian charity; B would like to invest money in indigenously run economic self-help enterprises, while C are Marxist revolutionaries supporting guerrilla uprisings against capitalist/imperialist domination. If we assume that the highest priority for each group is to maximise its own share of the relevant resources, the following type of scenario would occur if allocations were to be determined by voting, and the voting power of the three groups was equal. Imagine an initial equal distribution (although this assumption is not necessary for the argument) of £1,000 each:

A	**B**	**C**
£1,000	**£1,000**	**£1,000**

The Marxists, better versed in competitive strategy than the others, realise that they can better their position by proposing an alternative distribution that, while favouring themselves most, would also be better for B than the initial status quo. The proposed reallocation being:

A	**B**	**C**
£500	**£1,100**	**£1,400**

Clearly, both B and C would vote in favour. But group B learn tactical gamesmanship quickly and make a new proposal:

A	**B**	**C**
£800	**£1,500**	**£700**

C are not very pleased with the halving of their resources, but A's position is somewhat improved and they join with B in voting for the proposal. It should not take too much reflection to realise that there are indefinitely many alternative reallocations, any one of which would secure majority support.

Formal theorists worry about the consequences of these cycles of preferences for political stability, but the normative implications are critical.

If for any policy for which there is a majority, there is another possible policy for which there would also be a majority, then saying that there is a majority in favour of a policy says nothing that would justify or legitimise the implementation of that policy. None of this implies that if it makes sense at all to talk about a fair compromise between different positions, there is no point which would constitute a fair compromise. Suppose for the sake of argument we adopted what could be called a numerically weighted egalitarian assumption: individuals should be equal, but greater numbers of individuals agreeing should have a proportionately greater effect in determining an outcome. In the Third World poverty example given above, granted that we accepted the legitimacy of the claims of the three groups on the resources in question, the initial equal distribution would be the one that satisfied our criterion. Going back to the distribution of the 10 seats among three parties with 60%, 30% and 10% support, the 6:3:1 allocation would meet the criterion. And in these very simple cases there is a collective choice mechanism that would produce these results. Continual majoritarian cycling would occur, not because there is no point of fair compromise, but because assuming that people will choose something as close as possible to their top preference, the majority-rule procedure will enable ever-changing majority coalitions to emerge. In conclusion, we should be clear about what is being claimed and what is not being claimed in this chapter. We are not claiming that there is some group decision function that will unerringly select the fairest, best outcome in all situations of political disagreement. What we are claiming is that the belief that the simple binary-option majority-rule procedure has unique democratic legitimacy is deeply flawed. There are numerous types of collective decision-making procedure, and the task of assessing their democratic credentials and their drawbacks is a serious and complex business. The next chapter will approach this task by systematically examining the main types of pure voting procedure.