I

The Simple Logic of Storable Votes

Storable votes reflect a simple idea: if you are free to distribute your votes among different proposals, casting more votes on decisions you consider important increases the probability of obtaining your preferred alternative exactly where it matters. A fully rational voting strategy will take into account not only your priorities, but also others' priorities and their expected voting strategies. The final reasoning will be more complex, and future chapters will discuss these aspects in detail. However, it is good to begin by making the central idea as transparent as possible. This chapter describes a very simple scenario in which the logic of the voting scheme appears clearly, strategies that are likely to emerge as rules of thumb in more realistic situations are optimal here, and a number of results can be proved rigorously.

1.1 The Theory

Consider a committee of voters, meeting to decide whether to implement a number of distinct proposals, each of which can pass or fail. The proposals constitute the known, fixed agenda of the meeting, and the committee votes over the proposals sequentially. Think of the committee as the board of a corporation or an assembly of share holders, or think of it as a government body: a city council, a school board, a legislature, an international organization in which members are country representatives, or in fact as any group faced with a series of collective decisions. What is important is that each vote is a simple yes/no vote, and, for now, that the set of proposals is known and fixed before the actual votes are called. Although there are several proposals on the agenda, each stands on its own. There is no aggregate constraint that puts a limit on how many proposals can pass or that ties the approval of a proposal to the approval, or rejection, of another.

1.1.1 *Preferences and information*

Each member of the committee has a preference over whether any specific proposal should pass or fail and attributes a particular value to the outcome he desires. Thus, each person's preference has two features: the *direction* in which the person wishes the vote to go, and the *intensity* with which the preference is held. The intensity is summarized by the value a person attaches to a proposal. Think of the value as precisely defined: it could be in monetary terms – the amount of money the person is willing to pay to obtain the desired outcome, as opposed to its opposite; or it could be in terms of other goods foregone – for example, the time and effort a person is willing to devote to see his preferred outcome prevail. In either case, each member of the committee is able not only to rank the importance he assigns to the different proposals, but also to state whether one decision is twice as important as another one, three times more important, or only 10 percent more important. In economic-theory terms, the problem is phrased in terms of *cardinal values*: not only the ranking, but also the precise relative values, matter.

At the end of the meeting, the total value accruing to each committee member – the *utility* he derives from the meeting – is the sum of the values he attaches to all proposals that are decided in his preferred direction. Note that a concrete definition of value requires specifying the reference point, which here is the value of winning a proposal *as opposed to losing it*. Thus, by construction, the member receives zero value from outcomes that are opposite to his preferences. If values were defined with respect to an intermediate point – a tie, or possibly the status quo, if the status quo is not one of the two alternatives – the model would involve costs of losing, in addition to benefits from winning. The difference is semantic only. The only variable that matters is the difference in utility between winning and losing – simply called *value* or *intensity* in the first formulation, adopted here, it amounts to the value of winning *plus* the avoided cost of losing in the second.[1]

Every person knows his own preferences over all proposals on the agenda but does not know the preferences of the other committee members. However, everybody knows the statistical properties of these preferences. To keep matters as simple as possible, suppose that the proposals are unrelated, and preferences are independent both across proposals and across individuals: for any given proposal, the direction and the value of one person's preferences carry no information about the direction and the value of another person's, and for any given person, being in favor of one proposal carries no information about the preferred direction of a different proposal, and the value attributed to one proposal carries no information about the value attributed to a different one. Suppose also that each person is equally likely to be in favor of or against any proposal, an assumption that captures two different aspects of the environment. First, before the agenda is specified and concrete proposals are put on the table, no committee member expects to be systematically biased toward a favorable or a negative position. Second, once the proposals are known, each member is unable to predict the preferences of the others, and in the absence of any known bias, estimates that they are most likely to be equally split. Although individuals do not know others' preferences, these regularities are known to all. Finally, individuals know how strong and how dispersed preferences are on average. More precisely, everybody knows the probability distribution from which preferences are drawn, which is

1. Suppose that two proposals are on the table. In the first formulation, you derive, for example, a value of 1 from winning the first proposal rather than losing it and a value of 0.5 from the second (the units per se have no meaning). In the second formulation, you derive a value of 0.5 from winning the first proposal, relative to the status quo, and incur a loss of 0.5 from losing it; for the second proposal, the gain from winning is 0.25, and the loss from losing is 0.25. Clearly, you would rather win both proposals. Suppose that only one can be won. In the second formulation, if you win the first proposal and lose the second, your utility is $(0.5 - 0.25) = 0.25$; if you lose the first and win the second, your utility is $(-0.5 + 0.25) = -0.25$. Hence, you prefer the first scenario to the second because $0.25 > -0.25$, or $0.5 > 0$. This is exactly what the first formulation tells immediately: if only one proposal can be won, you prefer winning the first, because $1 > 0.5$, or $0.5 > 0$. It is true that the number summarizing your total utility is different in the two formulations, but remember that the units per se have no meaning.

the frequency with which a random voter's value is expected to assume different magnitudes.

The statistical properties of the voters' preferences play an important role because they summarize the environment in which the decisions have to be taken – the type of proposals submitted to the committee and the heterogeneity in the members' preferences and priorities. I will return to them below.

1.1.2 The voting rule

The proposals are voted upon sequentially, and each committee member has one regular vote to cast on each proposal. Each member also has one single *bonus vote* that can be cast, in addition to the regular vote, on any proposal. Think of the regular votes as red votes, and of the bonus vote as a single blue vote. In principle, the blue vote need not count exactly as one red vote – it could count more and be equivalent to several red votes, or it could count less, and it would then take multiple blue votes to equalize a single red vote. Whether smaller, equal, or larger, the blue vote is indivisible, and each committee member can cast it on a single proposal of his choice. In more general designs of the voting scheme, the voters will have multiple bonus votes, but, again, it helps intuition to keep matters as simple as possible. The total number of votes cast on each proposal is expressed in terms of regular votes: if the value of the bonus vote is different, the appropriate exchange rate is applied. A proposal passes if there are more votes in favor than against, and it fails in the opposite case. For simplicity, ties are resolved randomly. The outcome and the total number of votes cast on each proposal are known publicly.

The model is designed to ask whether the addition of the bonus vote changes matters, relatively to simple majority voting, and, if so, how. More precisely, imagine a preliminary constitutional stage, preceding the workings of the committee and held in ignorance of the specific agendas that will be voted upon later. The members can choose whether the committee will take decisions according to storable votes or to simple majority. What will they choose?

Box 1.1 The Formal Model. Version 1

Consider a committee of n voters, meeting to decide whether to implement T distinct proposals, $\{P_1, ..P_T\}$, each of which can either pass or fail. The T proposals constitute the known, fixed agenda of the meeting, and the committee votes on the proposals sequentially. Each member i of the committee has preferences over each proposal P_t summarized by a cardinal valuation v_{it}. Valuations v_{it} are independently and identically distributed both across proposals and across individuals according to the distribution $G(\mathrm{v})$ with support $[-1, 1]$, symmetric around 0. A negative valuation means that the voter is against the proposal; a positive valuation means that he is in favor. Call voter i's absolute valuation $v_{it} \equiv |\mathrm{v}_{it}|$ i's *intensity* (or *value*) over proposal P_t. For each proposal P_t, voter i receives utility $u_{it} = v_{it}$ if the decision goes in his preferred direction, and 0 otherwise. Voter i's utility from the entire meeting, U_i, equals $\sum_{t=1}^{T} u_{it}$.

In addition to a regular vote over every proposal, each committee member is given one indivisible bonus vote equivalent to B regular votes, with $B > 0$. Each proposal is decided in the direction that receives more votes; in case of a tie, a coin is flipped. When voting on any proposal, a member knows $\{\mathrm{v}_{i1},..\mathrm{v}_{iT}\}$, his own valuations over all proposals, but does not know others' valuations over any of the proposals. The independence of the valuations and the distribution G from which they are drawn are common knowledge. After voting, the total number of votes cast for and against a proposal is commonly known, and thus so is the number of voters still endowed with bonus votes, a number denoted by l_t. The state of the game at t is summarized by (l_t, t). The focus is on symmetric Perfect Bayesian equilibria in weakly undominated strategies.

Proposition 1.1: Equilibrium. (1) For all T, n, G, and B, there exists an equilibrium in which each voter i casts his bonus vote on P_t if and only if $v_{it} = \max\{v_i\}$. (2) For some T, n, G, and B, no other symmetric equilibrium in weakly undominated strategies exists.

Call Ev the expectation of the absolute valuation, or the expected intensity, and $Ev_{(T)}$ the Tth order statistic among each voter's T absolute valuations, or the expected highest-intensity draw. Finally, call EV_0 ex ante utility with storable votes, and EW_0 ex ante utility with simple majority voting, both evaluated before valuations are drawn, or equivalently before the agenda is specified, but in the knowledge of T, n, G, and B, and of equilibrium strategies. Then:

Proposition 1.2: Efficiency. For all T, n, and G, there exists a range of strictly positive values for B, $\mathcal{B}(n)$, such that for all $B \in \mathcal{B}(n)$: (i) If n is even, $EV_0 > EW_0$. (ii) If n is odd, $EV_0 > EW_0$ if $Ev_{(T)}/Ev > [T(n+1)]/[T(n-1)+2]$.

1.1.3 The voting decision

To evaluate whether storable votes are preferable to majority voting, individuals need to forecast how the bonus vote will be used, by themselves and by others. First, can voters ever benefit from misrepresenting their preferences and voting against their preferred alternative? Second, on which proposal will they cast their bonus vote? The answer to the first question is simple: there is no reason to misrepresent one's preferences. Because the proposals are unrelated, the direction of past votes holds no information about the direction of future votes and thus cannot be used to manipulate others' voting strategies. Voting against one's interest on any given proposal can only reduce the probability of obtaining the preferred outcome. Even if individuals are strategic, they will vote sincerely.

Less clear is when to cast the bonus vote. The choice must depend not only on how intensely a voter feels about the different proposals – the values he attaches to them – but also on the probability that casting the bonus vote will alter the outcome, relative to what it would have been without it – the *pivot probability* of the bonus vote. The pivot probability itself must depend on how others are using their bonus votes. Imagine being one of the committee members, and suppose, for example, that the first proposal is the one you care most about but there are others on the agenda over which you have strong preferences. As the first proposal is voted upon, you may hesitate to cast your bonus vote, on the argument that its impact could be higher on later proposals, when fewer bonus votes are likely to remain in circulation. The reasoning seems plausible, but it is mistaken: it is true that fewer bonus votes will be available to the other voters in future proposals, but there is no reason to expect that more will be cast on the first proposal than on any other. In fact, if every committee member casts the bonus vote on his highest-priority proposal, then he is equally likely to cast it on any of the proposals on the agenda. When the voting decision on the first proposal must be made, all proposals look identical in terms of the expected voting behavior of the other committee members: thus, if the first proposal is the one you care most about, it is by spending your bonus vote on this proposal that you can expect to benefit the most. Of course, once the votes on the first proposal are made public, typically fewer bonus votes will be left. All future proposals are, however, equally likely to receive bonus votes, and again, any voter left with the bonus vote should cast it on the proposal to which he assigns highest value.

The optimal voting strategy is very simple and corresponds to the intuitive rule of thumb of casting the bonus vote on one's highest-priority proposal. It is the only strategy that remains optimal for every voter, regardless of committee size and number of proposals on the agenda. It is also robust, in the sense that in many cases a voter should follow it even if other voters do not. Suppose that some cast their bonus vote randomly, or on their second, or third, or fifth priority. Because any proposal is just as likely to be chosen randomly, or to be someone's second, third, or fifth priority as it is to be somebody's first priority, nothing changes from the perspective of the remaining committee members, and their optimal strategy is unaffected.[2]

Note the good incentive properties: voters reveal their priorities truthfully through their own voting choices, and the bonus vote gives them weight. This a central feature of storable votes and stems directly from the budget constraint – the limited number of bonus votes – and hence from the links among the different proposals created by the budget constraint. The intensities with which preferences over the different proposals are held must be compared, and relative intensities are revealed sincerely and taken into account by the voting rule. In the case of a single bonus vote, what is revealed is the most strongly felt decision; with multiple bonus votes, the information revealed would be richer.

2. The optimal strategy changes only if some of the voters condition their voting choices on the *order* of proposals, as opposed to the intensity of their preference. I will return to this point in Chapter 4.

From the point of view of an individual voter, casting the bonus vote on a given proposal can only increase the probability that the preferred alternative wins. Clearly, using more votes cannot decrease the weight that an individual's preferences receive. The complication is that others are using their bonus votes, too. For example, if everyone casts the bonus vote on the same proposal, then every proposal is decided by simple majority, exactly as if no bonus votes were available. Storable votes make a difference if priorities differ across voters; if they do not, the bonus vote causes no damage, but it does not help, either. Suppose, then, that priorities do differ across voters, and bonus votes are cast on different proposals. In this case, the probability of obtaining one's preferred alternative is higher than with simple majority in the proposal on which the bonus vote is cast, but it is lower in the others: just as the presence of some "regular" voters works to increase the influence of those spending the bonus vote, the presence of "bigger" voters works to reduce the influence of those voters who cast their regular vote only.

To make matters concrete, suppose the bonus vote has the same value as any other vote, and consider two examples: an agenda of three proposals and an agenda of six. With each proposal equally likely to receive bonus votes from the other committee members, a voter's *probability of winning* any of the proposals – the probability of obtaining his preferred alternative – can be read from the Figure 1.1.[3]

The variable on the horizontal axis is the number of voters in the committee. The black dots plot the probability of winning when adding the bonus vote to one's regular vote, and the gray dots when casting the regular vote only, both relative to the corresponding probability with simple majority, when no bonus votes are allowed. Thus, for example, a dot at value 1.1 indicates a 10 percent increase in the probability of winning relative to simple majority voting, and a dot at 0.9 a 10 percent decline. In both figures, all black dots are above 1, and all gray dots are below 1, as expected. When the committee is small, the winning probabilities are very sensitive to odd or even numbers of members, and the dots in the graph "jump" with the addition of a single member. As the size of committee increases, the dots stabilize. In the case of three proposals, at larger committee sizes the probability of winning when

3. All calculations used to derive the figures are in Part II.

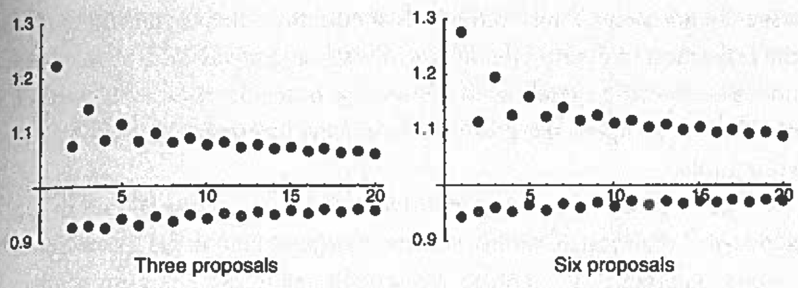

FIGURE 1.1 The effect of the bonus vote on winning probabilities. The horizontal axis is the number of voters. The black dots plot the probability of winning when casting the bonus vote, and the gray dots when not casting the bonus vote, both relative to the probability of winning with simple majority. The bonus vote is equivalent to a regular vote.

using the bonus vote is approximately 5 percent higher than with simple majority, and the probability of winning when not using the bonus vote is approximately 5 percent smaller. The direction of the change does not depend on the number of proposals on the agenda but the magnitude does, as shown by the second figure, with six proposals. As the number of proposals increases, the bonus votes become more dispersed; as a result, casting the bonus vote is more effective, and not casting it less penalizing (always relative to simple majority) than with three proposals. As the size of the committee increases, the probability of winning stabilizes around a 10 percent gain when using the bonus vote and a 3 percent loss when not using it.

The reason why storable votes can be preferable to simple majority voting is that the increase in the probability of winning is achieved over the decision that the voter considers most important, while the decline affects decisions about which the voter feels less strongly. Do storable votes then improve the committee's welfare?

1.1.4 *Welfare*

1.1.4.1 *The welfare criterion*

When we ask whether storable votes are desirable, from the point of view of the voters, the term "desirable" must be made precise. The normative evaluation must be separate from the personal benefit or loss derived in a specific instance. Thus, individuals should form their evaluation *before* the realization of their preferences is known or, equivalently, by referring not to the specific agenda or to the specific outcome of one meeting but to what they expect these to be in a typical meeting governed by the voting rule.[4]

The crucial element from which all other information can be deduced is the probability distribution that describes voters' preferences before uncertainty is resolved. As mentioned earlier, this probability distribution summarizes the environment in which decisions are taken. For example, if the committee's responsibilities are wide, a voter is likely to assign low intensities to most proposals, with few exceptions that matter to the voter greatly: the scenario is summarized by a probability distribution with large probability mass at low value realizations, and small probability mass at high values. If, instead, the committee is very specialized, a voter is likely to find all proposals of similar value, a scenario described by a probability distribution with most mass concentrated in a small interval.

As the Introduction discusses, in principle, the probability distribution of realized values could differ across different voters. In evaluating the group's welfare, however, on what basis can the preferences of one voter be recognized as systematically more intense than the preferences of another? The normative imperative of equal treatment translates into a normalization requiring that all committee members assign the same *average* value to the universe of possible proposals that come under the committee's jurisdiction. A stronger restriction, imposed in most of this book, is that the full probability distributions be identical across voters: it means forcing the voters to adopt an equal scale and to organize the different decisions according to a fixed ranking. Concretely, before the specific agenda is realized, each member expects that a typical proposal submitted to the committee will be of greatest importance to him (say, in the top 10 percent of possible values) with some probability. The model requires that such probability be equal across all members and extends

4. The observation that moral judgment requires a "veil of ignorance" over individuals' own positions in society was originally formalized by Harsanyi (1953, 1955) and recurs in welfare economics and in political philosophers' analyses of justice (most noticeably, Rawls (1971)).

this requirement to each interval of possible values. Imposing the same probability distribution does not amount to imposing the same *realizations* of intensities: *which* decisions will fall into the different intervals and, thus, which specific values will be realized typically will differ across different voters.

The evaluation of the voting rule should then be based on the expectation of future proposals, of their values, and of the voting outcomes that will follow, taking into consideration how voters will use their bonus votes. The expected utility measure that a voter can construct has no interpretation in itself: it is a number whose magnitude depends on the number of proposals and on the arbitrary scaling of the values distribution – whether, for example, the maximum possible value attached to each proposal is normalized to 1 or to 100. Any meaningful reading of the measure must be *relative*: it requires that it be compared to the equivalent measure under a different voting rule or a different decision-making procedure. An obvious possibility is to compare expected utility with storable votes to expected utility with simple majority voting; a still better alternative is to express both relative to expected utility under some ideal decision-making rule, thus being able to evaluate how well the two voting rules work, in some absolute sense.

As in the case of the two voting rules, the optimal decision-making rule specifies whether any future proposal should pass or fail, depending on the future realized preferences of the committee members, cannot discriminate according to the voter's identity, and must be chosen *ex ante*, before the agenda is specified and preferences are known. Under these constraints, the rule that every voter prefers, the rule that maximizes every voter's long-term payoff, is one that for each proposal favors the side with higher total value. A moment of thought makes the intuition clear. A rule that always favors the larger side (simple majority) means that a voter can expect to win more than half of the time. This is certainly better than deciding with a coin toss (in which the voter would expect to win half of the time), and better than a rule that favored the smaller side. It is the rule that maximizes the frequency with which any voter expects to win, but it is still not ideal because it does not recognize that some decisions matter a lot and some do not. A rule that favors the side with higher total value maximizes the total utility of the committee as a whole, and because everyone is identical before preferences are realized, that same rule must maximize the expected utility of any of the voters. Note that, although optimal, the rule cannot be used in practice because it depends on voters' intensity of preferences, a variable that is not observable and that individuals have an incentive to misrepresent: every voter would say that he cares very strongly about all decisions.

If the optimal decision-making rule provides a ceiling, it is also useful to consider a floor on the expected utility that an individual can achieve. In measuring how well a decision rule fares, a plausible floor is expected utility when decisions are made by flipping a coin, ignoring voters' preferences. The difference in expected utility between the ceiling (the optimal rule) and the floor (the equally probable random outcome) is the *available surplus* – the room for improvement over flipping a coin that is available to a decision rule. The measure of welfare most often used in this book is the share of available surplus that a voting rule is expected to appropriate: that is, how much of the improvement that the optimal scheme achieves over flipping a coin can be achieved through storable votes or majority voting.

1.1.4.2 Welfare comparisons

To begin evaluating the welfare properties of storable votes, some concrete examples will help. Even if in practice evaluations are likely to be more qualitative, precise numerical analyses discipline thought and clarify the role of different assumptions. Suppose that the bonus vote is equivalent to a regular vote. Consider first a situation in which voters have very little information, the committee's range of actions is diverse, and, as a result, voters believe that future proposals are equally likely to be of any degree of importance. For both storable votes and simple majority, the two diagrams in Figure 1.2 plot a voter's expected utility before the specific agenda is on the table, for the two cases of three and six proposals, expressed as the share of available surplus. Before the agenda is realized, all voters are identical, and a voter's expected utility summarizes the rule's welfare properties for all voters. A dot in the diagram at 0.9, for example, means that the voting rule is expected to capture 90 percent of the utility improvement of the optimal rule relative to flipping a coin. As in Figure 1.1, the variable on the horizontal axis is the number of voters in the committee. The black dots refer to storable votes, and the gray dots refer to simple majority voting. With storable votes, the figure

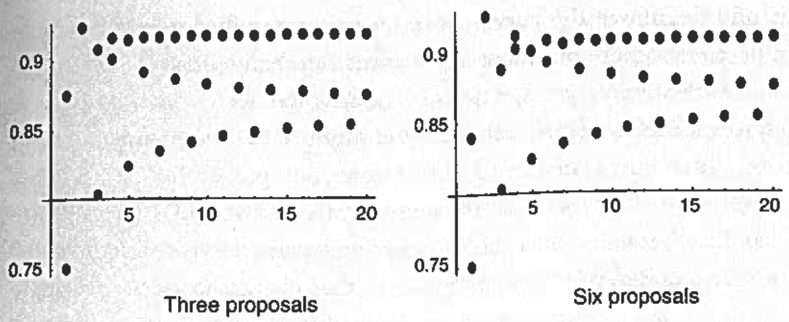

FIGURE 1.2 Expected surplus, as share of maximal surplus. The horizontal axis is the number of voters. The black dots refer to storable votes, the gray dots to simple majority, both as share of the optimal rule. The bonus vote is equivalent to a regular vote. Proposals are equally likely to be of any value.

is constructed under the expectation that all voters cast their bonus vote on their highest-intensity proposal, as indeed is rational.

The number of proposals makes very little difference. In both diagrams, the gray dots corresponding to simple majority voting trace two different curves: a lower curve, when the number of voters is even, and a higher curve, when the number is odd, whereas the black dots tend to be much closer to one another. Indeed, the most striking aspect of the figure is the stability of the storable votes' welfare measure, as opposed to the sensitivity of simple majority to odd or even numbers, especially when there are few voters.

With majority voting, when the number of voters is even and small, the probability of a tie is high, and there is no natural criterion guiding how to make the decision. Here I suppose that either side is chosen with equal probability, but other plausible outcomes (for example, favoring the status quo, if there is one) give very similar results: the voting system is inefficient because the decision criterion cannot reflect how strongly the two different sides feel about the proposal in question. When the number of voters is odd and small, on the other hand, majority voting performs well, because not only is a tie ruled out, but the majority is always so large, relative to the minority – twice as large in the case of three voters, at least 50 percent larger in the case of five – that only rarely would accounting for strength of preferences recommend the opposite outcome. The curve traced by the gray dots when the number of voters is odd lies above the curve traced when the number is even, and the difference is more pronounced when the number is small.

With storable votes, a tie occurs not when the number of *voters* is equal on the two sides, but when the number of *votes* is equal. This is very different: because of the bonus votes, the probability of a tie does not change much, whether the number of voters is odd or even. More than that, because bonus votes are used to express the intensity of voters' preferences, storable votes tend to favor the side that feels more strongly about the issue in question, as the optimal rule does. The figure shows that the two voting schemes are equivalent in committees of three voters (the highest point in both diagrams, where the black and the gray dots are superimposed), but storable votes outperform simple majority for all other committee sizes. Predictably, the difference is larger when the number of voters is even and small.

What happens when the number of proposals on the agenda increases? Both the optimal decision rule and majority voting are unaffected by changes in the number of proposals, because in both cases each decision is taken independently of all others. With storable votes, instead, decisions are tied by the votes' budget constraint: there is a single bonus vote to be cast on only one of the proposals. In this specific example, the relative performance of storable votes is slightly weakened by moving from three to six proposals and eventually stabilizes to a gain over simple majority of approximately 3 percent of possible surplus, as opposed to a gain of approximately 5 percent in the case of three proposals. The conclusion is not general, however. When the number of proposals increases, there are more proposals over which a voter cannot cast the bonus vote and over which his influence is less than with simple majority. But the value he expects to associate with the decision he will care most about increases, too, and so does the gain from having more influence over this one decision. Which of the two effects dominates depends on how large the expected differenc in value between the voter's highest priority and all others is, and thus, on the values' probability distribution, reflecting the committee' responsibilities.

It is probably unrealistic to think that voters can describe full the nuances of values they expect to attribute to future proposals. coarse classification, rather than a continuous distribution, seems mo likely. Suppose, then, that voters classify proposals into four classes of importance: "not important," "somewhat important," "important," and "very important," with these labels corresponding to a partition of the range of possible intensities into four intervals of equal size. Depending on the type of committee, members assign probabilities to the event that a random future proposal will correspond to each of the four classes. Consider the three examples in Figure 1.3.

The height of each column corresponds to the probability assigned to a typical proposal falling within each class of importance. The gray columns represent the example just discussed, in which voters believe that a typical proposal brought in front of the committee is equally likely to be "not important," "very important" or anything in between. The white columns correspond to a situation in which most decisions are expected to be strongly felt, with an increasing fraction falling into increasingly higher importance: only 5 percent of proposals are considered "not important" and close to 45 percent "very important." Finally, the black columns correspond to the opposite case, in which a full 50 percent of all

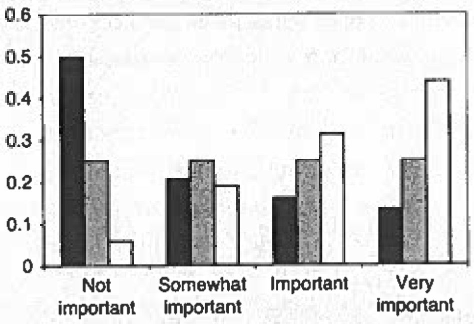

FIGURE 1.3 An example: Three coarse assessments of the probable importance of a random proposal. Each column denotes the probability of ranking a proposal in the corresponding class of importance, for three different distributions of values. The category "not important" corresponds to the lowest 25 percent of values; "somewhat important" to the second lowest quarter; "important" to values in the second highest, and "very important" to values in the highest quarter.*

* The columns are discrete approximations of continuous probability distributions, described by the cumulative distribution function $F(v) = v^b$ with different values of the parameter $b : b = 1/2$ (black columns); $b = 1$ (gray columns) and $b = 2$ (white columns).

future decisions are expected to be "not important," with the expected proportion belonging to each class declining as the attributed importance increases, until only about 10 percent of all proposals are expected to be "very important."

A bonus vote is most valuable when it allows a voter to select the rare proposal that matters to him, in contrast to the great majority that do not – the scenario represented by the black columns. Here, storable votes work particularly well: relative to simple majority, their performance is better than in Figure 1.2 for any number of proposals and all sizes of the committee. With three proposals, for example, the percentage gain over simple majority in the share of available surplus stabilizes at approximately 13 percent, at committee sizes of approximately twenty members or more.

On the other hand, consider a committee whose authority is exercised almost exclusively on matters of real consequence to all members, the scenario represented by the white columns in Figure 1.3. Now, storable votes perform less well because the possibility to increase one's voting power on a single issue, at the cost of reducing it on the others, is less valuable. With three proposals, storable votes continue to compare favorably to simple majority if the number of voters is either even or large enough, but their welfare properties are worse if the number of voters is odd and small. With a committee of three voters, storable votes are associated with a loss in expected welfare of 3 percent, relative to simple majority. As the size of the committee increases, the loss declines, reaching 1 percent for a committee of nine members and eventually becoming a gain of approximately 1 percent when the committee size reaches twenty-five or more.

These results are summarized in the top panels of Figure 1.4 (Figure 1.4a), representing the ratio of expected utility with storable votes to expected utility with simple majority, with three proposals. A dot at a vertical value of 1, for example, indicates that the two voting rules are equivalent. The horizontal axis is the number of voters, and, to keep the figures easy to read, the panel on the left refers to odd numbers; the panel on the right to even numbers. The color of the dots refers to the color of the columns in Figure 1.3.

I have used the example of three proposals to be concrete, but, in general, a committee will evaluate a different number of proposals at

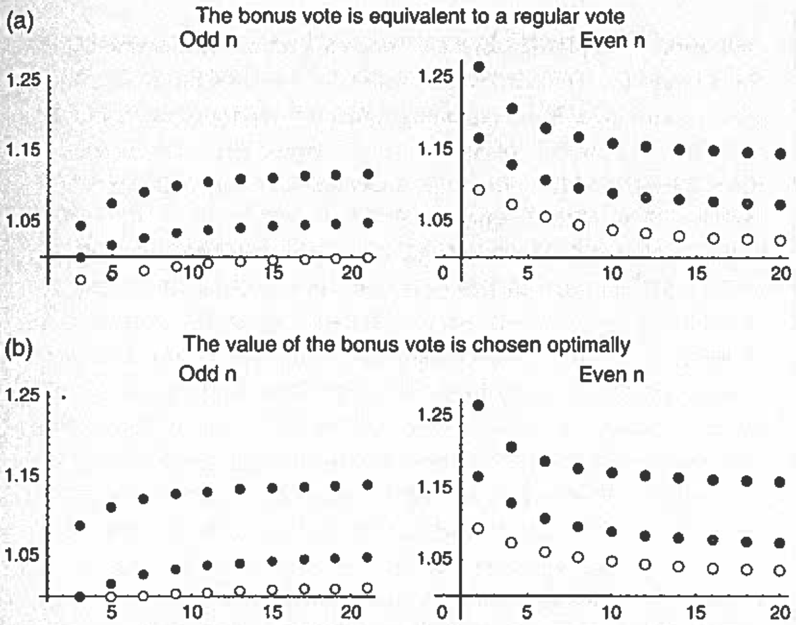

FIGURE 1.4 Expected utility with storable votes, relative to simple majority, for different value distributions. Three proposals. The horizontal axis is the number of voters, odd in the figures on the left, even in the figures on the right. The different colors correspond to the three different value distributions approximated by the columns of matching color in Figure 1.3. In Figure 1.4b, the value of the bonus vote is 2 for the black dots, 1 for the gray dots, and 0.5 for the white dots.

each meeting. The analysis can be extended immediately by considering the probability of each number of proposals. There is no reason why the calculations should not be transparent, and Box 1.2 describes one example in some detail.

Box 1.2 The welfare calculations: An example

Alice, Barbara, Caroline, and Dorothea – A, B, C, and D – form a committee that will meet regularly to vote on several proposals at each session. Consider A's reasoning in evaluating whether storable votes are desirable, relative to majority voting, keeping in mind that the evaluation takes place before any concrete proposal is identified, and the voting rule, if chosen, will be maintained over all future meetings of the committee. If the reasoning seems too sophisticated, read it as a measure of how each of A, B, C, and D would fare on average over many meetings of the committee, under either voting rule. The bonus vote is equivalent to a regular vote.

Consider first a meeting with three proposals. Alice begins by estimating the impact that storable votes will have on her probability of winning any proposal. She knows and takes into account that preferences are not correlated within the group and that, like herself, B, C, and D are equally likely to favor or oppose any proposal. If no bonus votes are allowed, A wins if B, C, and D all agree with her (with probability $(1/2)^3$); if two of the three agree with her (with probability $3(1/2)^3$), or if one of the three agrees with her and A's group wins the tie break (with probability $3(1/2)^3(1/2)$). The total probability of these events is 11/16, or 69 percent. In the presence of bonus votes, A conjectures that each of the other will cast her bonus vote on her most important proposal. With intensities of preferences uncorrelated across the four voters, any proposal has probability 1/3 of being a voter's priority, regardless of how intensely others feel about it. With these elements, A can calculate her probability of winning. Suppose, first, that A casts her bonus vote. If all three others also cast their bonus vote (an event with probability $(1/3)^3$), everyone casts the same number of votes, the scenario is identical to simple majority, and A wins with probability 11/16. If none of the others casts a bonus vote (with probability $(2/3)^3$), A wins if at least one of them agrees with her, an event occurring with probability $1 - (1/2)^3 = 7/8$. The other scenarios can be computed in the same manner: if only one of the others casts her bonus vote (with probability $3(1/3)(2/3)^2$), A wins with probability 3/4, and if two do so (with probability $3(2/3)(1/3)^2$), A again wins with probability 3/4. Thus, A's probability of winning a proposal when casting her bonus vote is as follows: $(2/3)^3(7/8) + (1/3)^3(11/16) + 3(1/3)(2/3)^2(3/4) + 3(2/3)(1/3)^2(3/4)$, or 78.5 percent. Following the same logic, A's probability of winning when *not* casting her bonus vote can be calculated as 64 percent. Thus, relative to simple majority voting, storable votes allow a voter to win her first priority decision more often (an increase of almost 10 percent) but also cause her to win the other two decisions less often (a decrease of 5 percent). Whether the trade-off is advantageous depends on the importance A assigns to her first priority, relative to the other decisions on the table.

Before any concrete agenda is specified, this relative importance must be estimated. Suppose that A estimates that any proposal is equally likely to be of any degree of importance. In different, but equivalent, terms, A estimates that there is a 10 percent chance that she will judge a proposal in the highest 10 percent range of possible importance, a 20 percent chance that she will judge it in the highest 20 percent range, etc. It is not difficult to verify that these beliefs imply that A expects to feel twice as strongly about her first priority proposal than, on average, she will feel about the others. Thus, together the two less important decisions will matter to Alice, on average, as much as the most important one. Since the increase in the probability of winning the latter is almost twice as large as the decrease in the probability of winning the two less important issues, expected utility with storable votes must be higher than expected utility with majority voting.

Assigning specific values to the boundaries of possible intensities yields specific numbers, summarizing the welfare evaluations. The exact numbers differ with different boundaries, but the ratios of expected utilities with different decision-making rules do not. Here, the interval of intensities is normalized as lying between 0 at the minimum and 1 at the maximum. On average, then, each of A, B, C, and D will assign to a random proposal an intensity of 0.5; the most strongly felt of the three proposals has an expected intensity of 0.75 and the weakest an expected intensity of 0.25. Expected utility with simple majority then equals (0.69)(0.5)(3) – the expected utility associated with each proposal is equal, and there are three proposals. Expected utility with storable votes on the other hand is (0.785)(0.75) + (0.64)(0.5) + (0.64)(0.25). The numbers are 1.03 and 1.067: relative to simple majority, storable votes yield a percentage gain in expected utility just below 4 percent, a result that would be unchanged had the normalization been different. If decisions were always resolved in favor of the side with higher total value, expected utility can be calculated to be 1.10. Thus, a committee member expects to obtain, on average, 97 percent of maximum expected utility with storable votes, and slightly below 94 percent with simple majority voting. If decisions are taken randomly, expected utility equals 0.75 – or (1/2)(0.5)(3), because the probability of winning any decision is 1/2. The percentage of available surplus over random decision-making that storable votes are expected to capture is then: (1.067 – 0.75)/(1.10 – 0.75), or 90.8 percent, versus 80.4 percent with majority voting.

It may well be that A, B, C, and D do not know exactly how many proposals will be decided at each meeting. For example, they may guess that there will be between three and ten proposals, but that this will vary across meetings. The same reasoning can be replicated for each case, and an average can be computed. If any number of proposals between three and ten proposals is equally probable, on average, expected utility with storable votes is 2.5 percent higher than with simple majority, and just above 96 percent of expected utility under the most-efficient rule (versus 93.7 percent for simple majority voting). In terms of available surplus, the percentage appropriated on average is 88 percent with storable votes, versus 80.4 percent with simple majority. The decline in the advantage enjoyed by storable votes, relative to majority, depends on the specific distribution of values – the assumption that any proposal is equally likely to be of any degree of importance – and may or may not occur with other distributions.

Once we understand what makes storable votes more or less desirable, we can think of improving the rule. First of all, when a large fraction of proposals are important, a single bonus vote does not allow a voter to represent his full priorities correctly. Why not grant *multiple* bonus votes? This is a question addressed in the next chapter.

Second, the impact of the bonus vote on the probability of winning a decision depends on the bonus vote's value, relative to a regular vote – whether it is worth one regular vote, as supposed so far, or two, or ten, or one-half. The larger its value, the stronger the effects: the larger the increase in the probability of winning the proposal on which the bonus vote is cast, but the larger the decrease in all other proposals.

Thus, ideally, the value of the bonus vote should depend on the relative importance assigned to the highest-priority proposal versus all others. Suppose that the probability of proposals of different importance is represented by the three cases in Figure 1.3. In the case of the black columns, having a bonus vote worth two regular votes, as opposed to one, improves the welfare properties of the voting mechanism, at all committee sizes and for all numbers of proposals: having a "big" bonus vote to cast on the occasional important decision is valuable, and the loss in influence over the other proposals is not very harmful. In the case of the white columns, on the other hand, storable votes perform better if the weight of the bonus vote is *reduced*. If the bonus vote is worth half of a regular vote (that is, if two bonus votes are needed to counter a regular vote), then expected utility with storable votes is always at least as high as with simple majority. Here most decisions are likely to be important, and losing influence over all proposals that are not one's single highest priority is costly. The bottom panels of Figure 1.4 show these results.[5]

The value of the bonus vote is part of the design of the voting scheme – the bonus vote is a special additional vote, and there is nothing particularly artificial in making it different from a regular vote. Is it *always*, however, possible to find a value of the bonus vote such that storable votes outperform simple majority? The answer is "almost always." More precisely, the answer is in Proposition 1.2 of Box 1.1, the main conclusion of the formal analysis. A value of the bonus vote that improves over simple majority is guaranteed to exist if one of the following conditions is satisfied: either the number of voters is even, or it is large enough, or the difference in expected intensity between a voter's highest priority and the average proposal is large enough. The proposition says a bit more: given the distribution of intensities that voters expect to assign to the proposals, it identifies the minimum committee size such that the storable votes scheme can be superior to simple majority voting. For the three examples of Figure 1.3, the only cases for which a strict improvement in expected welfare cannot be guaranteed are committees with three or five voters when most decisions are important (when the probability distribution is described by the white columns).

5. In the intermediate case, in which an equal fraction of all proposals is expected to belong to each class of importance (the gray columns), the bonus vote should be equivalent to a regular vote.

The conclusion, then, is that the welfare properties of storable votes depend on the probability distribution of preference intensities and on the value of the bonus vote, which itself should be chosen according to the distribution of intensities – large if low intensities are more likely, small if high intensities are more likely. This said, in practical applications, setting the bonus vote at the optimal value seems difficult. First, it requires the knowledge, shared by all, of the probability distribution of intensities. Second, in general, the value of the bonus vote should be different for different committees, opening the voting scheme to confusion and manipulation. It seems more reasonable to keep the system as simple, as unambiguous, and as constant as possible. For most of this book, the value of the bonus vote, or bonus votes when more than one is granted, will be kept constant. The numerical examples discussed in this chapter, the intuitions they convey, and the formal results in Box 1.1 all suggest the same lesson: it is prudent to set such a value not too high. The conservative recommendation is that the bonus vote should have a value not larger, and maybe smaller, than a regular vote.

1.2 The Experiment

The theoretical properties of the mechanism are one thing, but the use of the bonus vote in practice may be another. Laboratory experiments provide a controlled environment in which the scenario studied by the theory can be replicated, and deviations in behavior detected. They are the natural first stage in testing a new mechanism and are used extensively throughout this book. With a single bonus vote, a voter's decision problem is simple, and the optimal strategy is very intuitive. The one remaining concern is the possible temptation to hoard the bonus vote and use it disproportionately on later decisions.

Seven experimental sessions were run in 2006 and 2007 in the Center for Experimental Social Science (CESS) at New York University and studied how voters used the bonus vote.[6] Because both the practical organization of the experiment and its structure are common to all

6. The experiment was designed, programmed, and run with the help of Raj Advani, Bogachen Celen, and Shuky Ehrenberg.

experiments in the book, I describe them in some detail. Subjects were registered students who enrolled in the experiment by logging into CESS web site, and each of them participated in one experimental session only. After entering the laboratory, the subjects were seated randomly in booths separated by partitions and were assigned ID numbers corresponding to their computer terminals. When everyone was seated, the experimenter read aloud the instructions and answered all questions publicly. Each experimental session consisted of multiple *rounds*, varying in number from a minimum of five to a maximum of twenty. Each round consisted of the same voting game: subjects were matched randomly into committees of either three or four members and voted over three consecutive proposals, casting one regular vote for each proposal and a single bonus vote, which in all sessions had value equal to a regular vote.

Laboratory experiments are designed to control two variables that more than any others are difficult for an outsider to evaluate outside the laboratory: preferences and information. Preferences are not elicited, with the difficulty of evaluating the sincerity of the answers and comparing intensities across subjects; rather, they are induced, in the form of payoffs that are actually earned if the proposals are decided in the direction that the computer specifies for each subject.[7] To avoid adding subjective, and thus unknown, reactions to specific issues, the proposals are kept abstract and are distinguished by a neutral identifier, which was color in this experiment (green, blue, or red) and in the experiments described in other chapters. Information is always communicated publicly, ensuring that every subject not only shares the same information but is aware that everyone else does, too.

At the start of each voting game, each subject was told whether he was in favor of or against each of the three proposals (the direction of his preferences) and how much he would earn if each of the three proposals went in his preferred direction (the intensity of his preferences). For example, a subject would read on his computer monitor, "You are against the Green proposal and will earn 22 points if the Green

7. Factors remain that cannot easily be controlled – the value attached to money, the risk-aversion of specific subjects. They are not problematic in the experiment described here, in which the working assumption is simply that subjects prefer a higher payoff to a lower one, and, thus, everything else equal, prefer to cast more votes on the proposal that is associated with the largest gain.

proposal does not pass; you will earn 0 points if the Green proposal passes," and correspondingly for the other two proposals, typically with different preferred directions and different values. At the end of the experimental session, accumulated points were exchanged into dollars at a preannounced rate, yielding average hourly earnings of approximately \$20. Subjects were not told the preferences of the other committee members, but during the instructions period everybody was informed of the statistical process followed by the computer in assigning all preferences: each subject was equally likely to be in favor or against any proposal, and each subject's value was equally likely to be any number between 1 and 100. Thus, any proposal had equal probability of being associated with any admissible value, and a coarse approximation into four classes of importance would look like the gray columns in Figure 1.3.

The order of proposals, announced at the start of each voting game, was chosen randomly by the computer, independently of the proposals' color. After being shown his values, a subject chose whether to cast his bonus vote on the first proposal. The computer screen then showed the number of votes cast by each member of the subject's committee, whether the proposal had passed or not, and the subject's own payoff from the vote on the first proposal (ties were broken randomly by the computer). The committee then proceeded in the same manner to the second proposal, and then to the third and last, when all remaining votes were automatically cast. The vote over the third proposal concluded the voting game, the round. Subjects were then assigned a new budget of votes, and a new round started where three proposals were again voted upon in sequence. Four of the seven experimental sessions included a second treatment, described in Chapter 4, that required that the composition of the committees be kept fixed for the duration of the experiment. Because subjects' values were independent across both proposals and subjects, there should have been neither the incentive nor the possibility of coordinating voting choices, even with fixed committees. To verify that this was indeed the case, two experimental sections (identified as Sections 4 and 5 in Table 1.1) were designed with random rematching of subjects into new committees at the end of each round. The behavior of the subjects in these two sections is indistinguishable from all others. As for all

Table 1.1 Experimental Design

Session	Committees Size	No. of Subjects	Rounds
1	3	12	20
2	3	12	10
3	3	15	10
4	4	12	12
5	4	12	12
6	4	12	8
7	4	16	5

experiments described in the book, a sample of the instructions is in the Appendix.[8]

The experimental design is summarized in Table 1.1.

The results of the experiment were very close to the theory. The 91 subjects played a total of 974 rounds and in each decided which of the three proposals was to receive the bonus vote. In 974 rounds, there were 76 instances in all in which the bonus vote was not cast on the subject's highest value, and of these 76 strategic mistakes, a third (25) were due to only three subjects. Figure 1.5 reports the frequency of strategic mistakes and the corresponding number of subjects. The bins on the horizontal axis are frequencies of errors, as percentages of the total number of rounds in which a subject participated; the vertical axis is the number of subjects who fall into each bin. The figure says that 58 of the 91 subjects, close to two thirds of all subjects, never made any mistake; 12 made mistakes not more than 10 percent of the time; and 13 made mistakes not more than 20 percent of the time.

Did subjects benefit from storable votes? For the two committee sizes used in the experiments, Table 1.2 reports the expected share of available surplus that the theory associates with storable votes and simple majority voting. With the parameters used in the experiment, storable votes are expected to be equivalent to simple majority voting when committees

8. All storable-votes experiments, including those described in later chapters, were programmed with the Multistage Game software package developed jointly between the Social Science Experimental Laboratory (SSEL) at Caltech and the California Social Science Experimental Laboratory (CASSEL) at UCLA. This open-source software can be downloaded from http://software.ssel.caltech.edu/

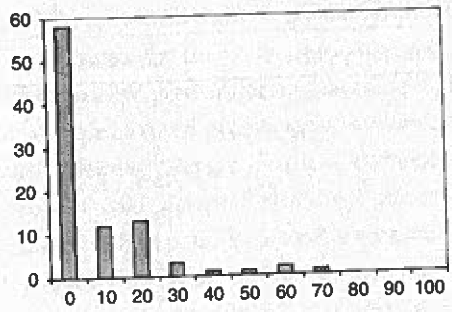

FIGURE 1.5 Experimental results: Individual error rates (91 subjects). The bins on the horizontal axis are percentages of errors over all rounds played; the vertical axis is the number of subjects in each corresponding bin.

Table 1.2 Theoretical Predictions

Committees Size	Surplus SV	Surplus Maj. Voting
3	92.3	92.3
4	90.8	80.4

are composed of three voters and superior when they are composed of four voters.

Figure 1.6 reports realized aggregate experimental earnings in each session and compares them both to the theoretical predictions (Figure 1.6a) and to the earnings that subjects would have had if decisions had been made with simple majority voting (Figure 1.6b). Each dot corresponds to an experimental session; black dots represent three-voter committees, and white dots represent four-voter committees.

Consider first Figure 1.6a. The vertical axis measures aggregate experimental earnings in each session as share of maximum possible total earnings; the horizontal axis measures the share of maximal earnings that subjects would have earned, had they made no strategic mistakes. All measures are calculated using the values assigned to subjects' preferences by the computer in each experimental session. For example, suppose that in a given round the assigned preferences for the Blue proposal in a three-person committee were $\{-27, 91, -32\}$. Then, over that proposal in that round, maximum aggregate earnings were 91 – the side favoring the proposal had higher value than the

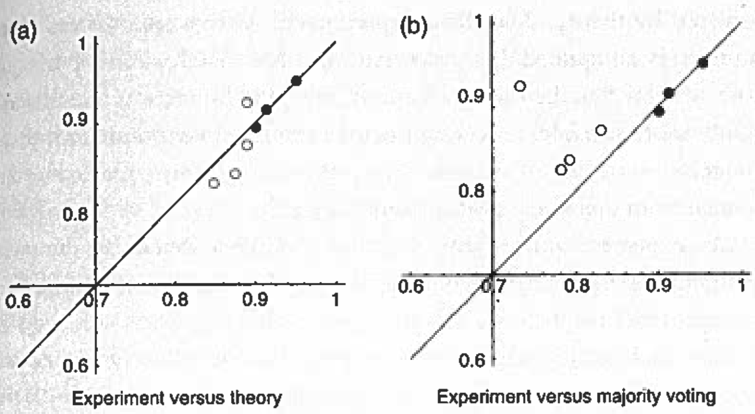

FIGURE 1.6 Experimental surplus, as share of maximal surplus. The vertical axis reports total experimental earnings per session, as share of maximal earnings. The horizontal axis reports the share of maximal earnings that subjects would have earned in each session in the absence of any strategic mistake (Figure 1.6a), or if decisions had been made through simple majority (Figure 1.6b). Black dots are sessions with three-voter committees; white dots with four-voter committees. All earnings are measured as gains over earnings with random decision-making.

side opposing it $(27 + 32 = 59)$. Aggregate experimental earnings were the earnings realized in the experiment, in which the winning side was the side casting more votes, including bonus votes. Aggregate earnings in the absence of strategic mistakes would be computed taking into account the subjects' assigned values over the other two proposals, and supposing that each subject cast his bonus vote on the proposal with highest value. In the figure, all earnings are measured as gains over expected earnings when the proposal is decided by flipping a coin $((1/2)(91) + (1/2)(59) = 75$ for the for the Blue proposal in this example), and thus the measures are immediately comparable to the numbers in Table 1.2.

The vertical axis shows that realized earnings ranged between 82 and 92 percent of maximal earnings for committees of four voters, and between 89 and 95 percent for committee of three voters. As predicted, the share tends to be higher for groups of three voters than for groups of four. The effect of strategic mistakes appears in the vertical distance between each dot and the 45-degree line: by construction, dots on the 45-degree lines are sessions in which the aggregate payoff is exactly as predicted by theory. The three experimental sessions with three-voter committees all yielded aggregate earnings that match almost perfectly those predicted by the theory. For four-voter committees, the results are slightly more variable, reflecting the lower number of rounds and, thus, the higher variability of average intensities and aggregate payoffs, but the distance from the 45-degree line remains small.

As expected from Table 1.2, the difference between the two treatments emerges clearly when experimental earnings are compared to earnings under majority voting. In Figure 1.6b, the vertical axis is again the share of maximal earnings captured in each session, but the horizontal axis reports the corresponding share under majority voting, again with the realized experimental draws – in the example above, aggregate earnings with majority voting would equal 59, since the majority opposes the proposal. As in Figure 1.6a, all earnings are measured as gains over random decision-making, and can be compared to the predictions of Table 1.2. Dots above the 45-degree line indicate experimental sessions in which aggregate earnings were higher than they would have been with majority voting.

As shown in Table 1.2, the theory predicts that three-voter committees should fare equally well under storable votes as under simple majority voting, while four-voter committees should fare better. Indeed, now all four white dots are above the 45-degree line, one particularly so, while the three black dots continue to lie on the 45-degree line. Again, the theoretical predictions are borne out. The disparity in results between the two treatments deserves emphasis: the four-voter committees' results show that storable votes are indeed capable of making a difference.

The experiment was very simple and the positive results not too surprising. They are, however, encouraging. The most basic model of storable votes has not been rejected. The next chapters can begin to enrich it.

1.3 Conclusions

Individuals differ in the strength with which they hold preferences over the alternatives offered to their choice. Any single individual will consider some choices important and some trivial, and different individuals will typically rank the importance they assign to the various options differently. This is true for private goods, and it is the basis of the functioning of markets through prices. It is also true for public decisions, and a voting system that is responsive to the intensity of people's preferences can exploit and mirror this difference, allowing voters to come closer to their preferred outcome when it matters to them most. Storable votes is a voting system with just such a property: by granting a budget of votes to cast freely over a set of decisions, it allows voters to express the intensity of their preferences, exactly as prices do in a market system, and to increase the probability of obtaining the outcome they desire in those decisions in which they are willing to "spend" more of their budget of votes.

This chapter has studied the simplest model of storable votes, in which, in addition to a regular vote for each decision, voters are granted a single extra vote to cast freely over a set of binary decisions, and preferences are not correlated either across voters or across decisions. Each decision is taken according to the majority of votes cast. The single bonus vote reduces the analysis to its essence and clarifies the intuitions on which the system rests. In contrast to simple majority voting, the addition of a single bonus vote is sufficient to stabilize the performance of the voting rule, regardless of whether the number of voters is odd or even. Majority voting performs well when the number of voters is small and odd, but especially if the bonus vote is of value comparable or somewhat smaller to a regular vote, storable votes outperform it for all or almost all numbers of voters. The result is unambiguous if decisions are either equally likely to be considered of any level of importance, or if most decisions are judged to be low priorities, with the occasional outlier that can then benefit from the bonus vote. Confirming the straightforward, intuitive logic behind storable votes, the theoretical predictions were almost exactly verified in a laboratory experiment.