1 Introduction

In November 2009, voters in Maine were asked to vote on seven propositions. The topics varied: from school districts' consolidation (confirmed), to tax relief for new or clean energy cars (rejected), to state dispensaries for marijuana devoted to medical use (approved). Two propositions attracted intense campaigning: question 1 on the ballot concerned the repeal of a 2009 state law allowing same-sex marriage (repealed), and question 4 proposed submitting increases in state taxes to popular referendum (rejected). Many voters may have felt that these last two issues indeed dominated the ballot, but it seems likely that not everybody did–chemotherapy patients and their families, groups opposed to cars' excise taxes who had spent their resources to put the initiative to a vote, families in isolated communities without schools in close proximity. Even among voters focused on questions 1 and 4, priorities probably differed: question 1 concerned social rules, question 4 economics, and many may have considered one much more important to them than the other. Now, suppose that in addition to a regular vote on each question, each voter had at his disposal a number of "bonus votes": extra votes that could be cast on any proposition as the voter saw fit. The decision would then be taken according to the majority of all votes cast, including the bonus votes. A voter who felt that the school district question was central to his family's well-being could cast all bonus votes on that question; anti-tax voters would presumably reserve most of their bonus votes for question 4, possibly casting a few towards reducing the excise tax on cars; and voters who came to the poll mostly in reaction to the same-sex marriage question could concentrate all their bonus votes on that. Would such a scheme be a good idea? Before dismissing it out of hand, read this book.

The logic behind the idea is easy to see. Everybody is given the same number of bonus votes, and thus everybody is treated identically. Yet, a minority can

prevail on an issue if it casts more bonus votes than the majority does. But *only* if it casts more bonus votes than the majority, and thus only if voters on the minority side care more about the issue than voters on the majority side. This is exactly the case of "intense minority" and "apathetic majority" that in political theory casts doubts on the desirability, if not on the legitimacy, of majority rule. And note that the larger the difference in size between the two groups, the larger the discrepancy in the use of bonus votes must be for the minority to prevail. A large enough majority cannot be beaten. Unless of course it chooses not to go to the polls, an outcome that can in itself be interpreted as evidence of weak preferences but that in any case seems less likely to occur when the majority feels less secure, as it would in the presence of bonus votes.

On the other hand, matters may go wrong. Voters will not want to spend bonus votes on proposals they are sure to win or to lose anyway; thus they will need to forecast others' voting behavior and condition their own votes on what they expect others to do, with the result that strategic considerations may interfere with casting the bonus votes according to one's priorities. The added complexity of the decision may confuse the electorate, or at least a part of the electorate, presumably the weaker part which we particularly want to protect. The budget of bonus votes creates a link between the different proposals–what is spent on one cannot be spent on another–and may increase the importance of controlling the agenda. Proposals could be put on the agenda with the sole purpose of draining bonus votes from specific groups of voters. And if the bonus votes decision depends on forecasts of others' voting, distorting the forecasts could be advantageous. Finally, the effect of the bonus votes is to make minority victories possible, indeed the *only* effect, since otherwise the voting rule would be identical to simple majority voting. But the possibility of minority victories must be considered with great caution: every time the majority loses,

the greater part of the voters will have to live with decisions they oppose.

I hope it is becoming clear why a whole book, a short book, was necessary to begin analyzing these questions.

2 Storable Votes and Majority Rule

I call the voting rule I just sketched *Storable Votes* because the central idea is the possibility to shift one's own votes from one contest to another, to *store* votes not spent on decisions that are low priorities for use over decisions that matter more. The details can vary: the electorate may be large, as in the case of a popular vote, or small, as in the case of specialized committees; the different proposals may be presented simultaneously, as in the example of multiple propositions submitted to referendum, or be voted upon over time, as in the case of a committee that meets at regular dates– for example the Board of a Central Bank convening monthly to decide interest rate policy; the agenda may be known, as in the case of multiple referendums or of the Central Bank Board, or unknown, if the committee's role is sufficiently wide that new questions can arise. But two conditions are required. First, each decision has only two alternatives. There are of course cases where voters are asked to choose between two candidates only, but in the choice of representatives a field limited to two candidates is not the norm. Storable Votes are better suited to committee decision-making and to direct democracy, where each proposal either passes or fails. Second, the group faces several decisions, and these decisions are separable–each can be judged on its own merits and pass or fail, independently of how other decisions are resolved. The two conditions limit the applicability of the voting scheme, but, as the example of the Maine propositions shows, can still be satisfied in realistic and in fact important scenarios. For example, looking more widely at the November 2009 election, a total of 26 propositions were on state ballots; one

of them was presented alone to New Jersey voters; all others were in bundles of several propositions (7 in Maine, 11 in Texas, 3 in Ohio, 2 in New York, and 2 in Washington State), and in all cases the propositions on the same ballot were unrelated. If we take the agenda as given, an important point to which I will return below, the two conditions are routinely satisfied by committees of very different types, from corporate boards to school boards, from professional associations to faculty committees in universities. Indeed, if we take the agenda as given the group of voters could be a city council, a government committee, a legislature.

When a group needs to take a single decision choosing between two alternatives only, majority voting–defined here as enacting the alternative preferred by the larger number of voters– is a particularly desirable rule. It guarantees that individuals have no incentive to misrepresent their preferences, regardless of how others are voting. It treats the two alternatives identically (as opposed, for example, to supermajority requirements to overcome the status quo). More importantly, it treats every voter identically, and responds non-pathologically to shifts in voters' preferences: an increase in the number of voters favoring one alternative can only increase the likelihood that the alternative wins. Not only does majority rule have these properties, but it is the *only* voting rule that does.[1] In fact, there is only one intuitive difficulty that in the case of two alternatives majority rule cannot solve–what democratic theory has called the *Problem of Intensity*: "What if a minority prefers an alternative much more passionately than the majority prefers a contrary alternative? Does the majority principle still make sense?" (Dahl, 1956, p.90)

It is a single difficulty, but an important one. On ethical grounds, if majority rule runs counter to our intuitive moral sense in some circumstances, its entire

[1] The observation that misrepresenting one's preferences can never be advantageous is straightforward; the other properties are known as May's theorem (May ,1972).

foundation seems fragile. On pragmatic grounds, the one failing is at the origin of much that is decried about committee decision-making: the private deals, the horse-trading, the backrooms, at least metaphorically smoky. More dangerously, if the composition of the passionate, losing minority remains constant over many decisions, the minority is effectively disenfranchised, and the stability of the political system threatened.

A large literature has debated whether the very concept of intensity is logically rigorous. When applied to a single individual, intensity relies on a cardinal measure of preferences that acquires meaning from the comparison to some unspecified numeraire: for example, I can say that I prefer alternative A to alternative B four times more intensely than to alternative C if I am willing to devote four times more of my time to defeat alternative B than to defeat alternative C. But when applied to the legitimacy of the majority's choice, the question becomes one of interpersonal comparisons. We can say that Camilla prefers B to A three times more intensely than to C–but we have learnt nothing about the relative intensity of preferences between Camilla and me in the choice of A versus B. If intensity cannot be measured across individuals, the goal to protect "intense" minorities cannot be defined, let alone achieved. In these terms, the objection has no good answer. And yet a more modest and more pragmatic approach in my opinion is more productive. In Dahls' words: "We shall continue to believe not only that we can guess intelligently but that we must guess intelligently about such things" (Dahl, 1956, p.100). Storable Votes are an attempt to attack the problem of intensity minimizing the amount of guesswork involved. They are not a fundamental answer to the difficulties of social choice, but a practical suggestion for group decision-making in specific circumstances.

3 Objectives

Storable Votes work by linking multiple binary decisions through a single budget constraint, the limited number of votes that individuals can shift across the different decisions. Because casting more votes translates, on average, into exercising more influence, individuals choose the use of the votes at their disposal according to the relative strength of their preferences, and the number of votes cast becomes the observable measure of intensity. There are complications, but notice first what Storable Votes attempt to do.

First, by allowing each individual to cast more votes on decisions that matter to him more, Storable Votes also allow him to increase the probability of winning decisions he considers important, at the cost of a lower probability of winning where it matters to him less. Thus in a body of voters homogeneous enough that all expect to be in the minority with similar frequency, an individual will expect to fare better, on average, than he would under majority rule, even though he also expects to be on the losing side more often than he would with majority voting. Second, the possibility of minority victories arises even though each single decision is taken according to the majority of votes cast, and, what is most important, every individual is treated identically: every voter has the same total number of votes, over the full set of decisions, and every vote has the same weight, regardless of the identity of the voter. Over any one decision different voters may well choose to cast a different number of votes–indeed if they do not the voting rule reverts to majority rule–but this is a choice that any of them makes freely. In a market, two consumers will choose to buy different quantities of a given good, but if their resources are equal and they face the same price the different purchases are not evidence of unequal treatment. Finally, Storable Votes are a plausible tool for addressing the problem posed by systematic minorities, groups of voters who consistently find themselves on the

minority side of most decisions. With majority rule, such groups are effectively disenfranchised. Social choice theorists have suggested reverting in these cases to rules of *fair division*, rules that recast the decision-making power as proportional to the relative size of the two groups: for example, establishing turns in dictating the outcome, with frequency proportional to size; or assigning the power to control the outcome with probabilities again proportional to size.[2] The rules are clever and please our intuitive sense of justice, but remain a theoretical suggestion. Unless they are adopted always, routinely and mechanically allowing for minority rule, they demand the judgement that an exceptional situation has been reached requiring a deviation from the standard voting system. But who makes the judgement? And why should the majority agree?

Storable Votes, on the other hand, do not require a different voting rule to respond to minority and majority's interests. Even a systematic minority can prevail occasionally, if it cumulates its bonus votes on few decisions only, and if the majority has different priorities. The voting rule that works well when individuals find themselves changing roles–sometime on the minority's side, more often on the majority's–works well when the roles are much more rigid. Clearly Storable Votes cannot protect minority interests when the electorate is consistently polarized: the same minority always disagrees with preferences that the majority holds strongly. But in this case, the union of the group seems deeply fragile, possibly unwise.

[2] **Reference to Lani Guinier? References to Brams' work, et al.** . Gerken (2005)'s concept of *second order diversity*–diversity across representative bodies, as opposed to within bodies–and its identification in existing US institutions can be read as an empirical counterpart to theories of fair division. The focus on legal thought and jurisprudence is neatly complementary to the abstract arguments of social choice scholars and philosophers.

4 Assumptions.

The theory of Storable Votes relies on methodological assumptions that should be discussed up-front. First of all, because it addresses the problem of intensity, it has not choice but to view voters' preferences as cardinal: voters are faced with multiple decisions, and their preferences specify the precise relative importance that each of them assigns to the different decisions. It is not enough to say that I care about proposal A more than about proposal B, and more about B than about C–an ordinal ranking; cardinality requires that I be able to say that I care three times more for A than for B and five times more for B than for C–a cardinal ranking. Note that the absolute measure of intensity is irrelevant and indeed undefined; it is the *relative* intensity that matters. An intuitive, if a bit loose, interpretation is to suppose that voters divide the different proposals into classes of importance. For example, voters may think in terms of four classes–*very important, important, somewhat important, not important*–and assign a minimal weight of 1 to proposals in the lowest category, increasing the weight by 1 with each jump of category. The weight, or the importance assigned to each proposal represents the *intensity* of a voter's preferences. Of course nothing prevents a finer classification.

Where cardinal intensity really becomes problematic is in interpersonal comparisons, but the focus on multiple decisions helps. With multiple decisions it becomes natural to concentrate not on the intensity with which any one decision is felt, but on the *range* of intensities that a voter experiences over the whole set of proposals. And if the analysis is to apply beyond the specificity of a given set of proposals, the essential ingredient becomes the distribution of intensities–the range and the relative probabilities– over the full set of potential decisions, the universe of decisions that fall under the competence of the committee. It is this measure of intensities that allows the voter, and us, to form expectations about

the functioning of the voting rule and its properties independently, better, *before* the exact realization of the proposals being debated at any precise meeting. Interpersonal comparisons of intensities then take the form of assumptions about the distributions of intensities associated with different individuals. Faced with the impossibility of taking a reasoned stance on individual differences, we can restrict all distributions to be identical, without in any way imposing identical realizations of intensities over any specific set of proposals. The assumption instead amounts to requiring that all voters evaluate the importance of different decisions using the same scale: for example, returning to the 4-category classification described earlier, all voters might classify 5 percent of all possible decisions as "very important", 20 percent as "important", 25 percent as "somewhat important" and 50 percent as "non important".[3] *Which* decisions belong to which class is free to vary across individuals, and any realization of a specific set of proposals need neither cover the full set of categories for any individual, nor be ranked similarly across individuals in any systematic fashion. Only over a large number of realizations are the frequencies expected to be realized. Restricting all distributions to be identical remains an arbitrary assumption; it rules out the possibility that committee members have different views about the overall importance of the issues under the committee's jurisdiction, relative to other aspects of their life, or the possibility that differences in personality be reflected in objectively different intensities of feelings. It need not be the correct assumption, but seems a safer choice than any other alternative.

One benefit is that it allows to evaluate the performance of the voting rule through its effect on a utilitarian welfare criterion, giving each individual the

[3] Distributions of intensities differ from distributions of preferences, because the latter include not only the intensity but also the preferred direction of a vote (whether the proposal should pass or fail). In the case of systematic minorities, for example, distributions of preferences differ across individuals, because members of the minoritiy always, or at least more frequently, agree among themselves and disagree with the majority. But the distributions of intensities may well remain identical, and indeed do so in this book (see chapter 3).

same weight. I will not argue here in general defense of a utilitarian criterion; it is the meter that fits most naturally the methodological approach of this book. But utilitarian welfare functions depend on interpersonal comparisons of intensities and would be more difficult to accept without the assumption of identical distributions of intensities. The intensity of a voter's preferences maps naturally into the utility the voter receives if his preferred option is chosen, suggesting ex ante individual welfare as measure of how well the voter expects to fare. Ex ante individual welfare is evaluated "under the veil of ignorance", before specific proposals are brought to the table and specific preferences realized. It is a voter's estimate of his probability to prevail over different decisions and of the utility he derives if he does, together with the probability that the different decisions will be called for a vote. If the distribution of intensities is identical for all committee members, each individual faces identical prospects, and ex ante individual welfare is identical for all voters. It thus becomes the natural welfare criterion for the group as a whole. To see what the problem could be, suppose this were not the case and, to take an extreme example, suppose that Camilla, a skeptical character, considered all decisions "not important", assigning to them an intensity of 1, while I, a more sanguine personality, consider them all "very important" and assign to them an intensity of 4. A utilitarian criterion that were to sum our ex ante individual expected welfares would be strictly identical to a criterion that assigned to us the same intensities but weighted my utility four times as much as Camilla's. The distinction between individual weights in the utilitarian welfare function and individual distributions of intensities cannot be made. Again, assigning equal distributions and equal weights need not be the correct assumption, but outside of specific examples seems the least problematic.

Finally, even within the restricted methodological approach of this book, Storable Votes are too simple a mechanism to have ambitions of optimality.

Other decision-making rules exist that on average a voter would prefer, but whether these rules can be implemented through a voting scheme, however, is a different matter. From the theory of Mechanism Design we know that, under some conditions, an appropriate system of taxes and transfers can lead individuals to reveal their preferences and can select the policy alternative that maximizes utilitarian welfare. But the correct transfers are complex, their design requires knowledge of the voters' distributions of values, and if the budget is to be balanced, individuals may need to be forced to participate in the system. Mechanisms without transfers would be preferable: procedures that ask individuals to send messages and select the winning alternative on the basis of those messages alone–in other words, voting rules. But if we focus on voting rules, even when the choice is between two alternatives only, optimality results have been obtained so far only in limit cases, as either the number of decisions or the number of voters becomes arbitrarily large. And even in these cases, the design of the correct rule again requires the knowledge of the distribution of preferences of the voters, opening the way to two plausible objections. First, while application to a very large number of voters seems practically important, the knowledge of the voters' distributions of preferences in such a case makes voting redundant: a planner who knows how to design the correct voting rule also knows the choice that maximizes utilitarian welfare. Second, the voting rule needs to reflect the distributions of preferences, and thus needs to change when circumstances change, making its political acceptance extremely unlikely.[4]

[4] I am referring to two schemes proposed by the literature. Jackson and Sonnenschein (2007) study an environment with multiple binary decisions, as in this book, and suggest asking each individual not only his preferred alternative, but also the importance assigned to each decision, restricting the reports to mimic the known distribution of intensities. For example, returning to the 4-category classification described earlier, if it is known that the probability of a decision with intensity 4, a "very important" decision, is only 5 percent, a voter is not allowed to report an intensity of 4 for more than 5 percent of the decisions he faces. As the number of decisions beomes very large, individual reports mimic, more and more closely, the actual distirbution of preferences and the outcome approaches the utilitarian optimum. When a single binary decision is at stake, Ledyard and Palfrey (2002) show how a simple yes/no referendum approaches the utilitarian optimum as the number of voters becomes very

Storable Votes have a more modest but more practical purpose: they are a simple, intuitive rule that can be implemented in practice. The immediate question then is not whether Storable Votes are the best possible rule, but whether they are preferable to existing alternatives. With binary decisions, the concrete relevant alternative is simple majority voting, and it is to majority voting that Storable Votes are compared throughout this book. The final conclusion, built over different models and different possible applications, is that Storable Votes are not *always* superior, but they are superior with sufficient frequency and in a sufficiently systematic manner to justify, in my opinion, the attention this book devotes to them and, hopefully, some further study with an eye to actual implementation.

5 Storable Votes versus other voting rules.

Comparison to majority voting is required because majority voting is the existing default rule in binary choices. But how do Storable Votes fare relative to other possible voting systems? If we suspend for a moment ethical considerations, the most natural reference is to a market for votes. Suppose first that voters all had the same budget and were allowed to buy or sell votes, either for money, or in exchange for votes over future decisions; as in the case of Storable Votes, the system leads voters to express the intensity of their preferences by varying the number of votes that an individual casts on different decisions, and it is on this basis that it has at times been deemed desirable (**refs**.). Is it then equivalent to Storable Votes? The short answer is *no*: Storable Votes allow voters to trade intertemporally with themselves only, as opposed to trading with

large, if the threshold for acceptance is chosen correctly, reflecting the electorate's distribution of intensities.

others. The difference has a number of important implications. First, contrary to Storable Votes, it is remarkably difficult to predict how a market for votes would function. As noted in the literature, the difficulty is in finding the correct price, either in terms of money or of future votes, at which current votes would trade. Consider for example a group of three voters who disagree over whether to approve or not a specific proposal, and suppose that the decision is taken by majority but voters are allowed to buy or sell votes for money. If the minority voter succeeds in buying one vote–and there is no reason to buy more than one–the remaining vote on the majority side has zero value, and its owner then prefers to sell it too. Indeed, he would want to sell it at any price, including a lower price than the one paid by the minority voter in the original transaction. Thus, if any trade takes place, at any positive price the market has excess supply. But if the price is zero, and majority voters have no interest in selling, the minority voter is certainly interested in buying and raising the price: the market has excess demand. There is no price that clears the market.[5] The example considers trading votes for money, but the difficulty remains in the case of log-rolling, where votes are exchanged for votes. In practice, this means that any market would result in rationing: there would be voters interested in engaging in trade at the current terms but shut out of all transactions. Who these voters are depends on the specific rules through which traders are chosen, and thus an immediate consequence is that these auxiliary rules assume a crucial role in the functioning of the market. In fact, unless traders are selected through the correct mechanism, the outcome is unlikely to be desirable. The problem is that by shifting the balance of votes, any exchange affects everybody's welfare, but the two parties involved in a trade will not take these effects into consideration; only if the right to trade is allocated in such a manner as to force them to internalize the consequences on others, will the trade be beneficial in terms

[5] This observation is made most clearly in Philipson and Snyder (1996).

of overall (utilitarian) welfare. For example, if the voters actually trading are selected randomly among those interested in doing so, a market where votes are exchanged for money will generate outcomes that are typically inferior to simple majority, in terms of expected utilitarian welfare.[6] Finally, if we consider how a market for votes is likely to function in practice, two additional considerations deserve a remark. First, the rationing rules add a further layer of possible manipulations: the instinctive reluctance to accept vote markets as legitimate institutions mirrors in part our unease with the lack of transparency and oversight through which deals are concluded. Second, whether trade occurs through money or through future votes, differences among individuals in terms of monetary or political power will introduce constraints in trading that are unrelated to strength of preferences. Not only do these constraints appear unethical, but they will also interfere with the welfare performance of the market.

Storable Votes cannot be assimilated to vote markets because they do not rely on interpersonal exchange. By allowing each voter to cumulate his influence on a subset of decisions, they have more in common with a system of limited vetoes, where each voter is granted a fixed number of vetoes to spend over the different decisions. But Storable Votes allow voters not only to block proposals but also to pass them over the others' opposition; a more fitting parallel then is a scheme we can call "rotating dictatorship", where individuals take turns in determining the outcome of different decisions. Once again, it is clear that the core of the scheme must be the mechanism through which individual voters are matched to the specific decisions they will control. A mechanism that accounts for aggregate utilitarian welfare and produces outcomes that are preferable to majority voting can presumably be designed, but to my knowledge has not been proposed and seems likely to reintroduce the complex set of transfers and taxes that voting rules are designed to avoid.

[6] This is the conclusion of Casella, Llorente-Saguer and Palfrey (2009).

A voting system does exist which Storable Voters resemble closely, and it is to be found among semi-proportional systems used for candidates elections in multi-member districts. When several representatives are to be elected, Cumulative Voting grants each voter a specified number of votes, and lets him cumulate as many as he sees fit on any individual candidate. Suppose for example that five seats are to be filled from a field of twelve candidates. Voters are endowed with multiple votes–commonly five, if there are five seats–and allowed to cast them as they wish, "cumulating" the votes on as few candidates as a single one, or spreading them over five, or choosing any intermediate option. Cumulative Voting is designed to protect minority interests: if voters were limited to a single vote per candidate, a party representing more than 50 percent of the voters and fielding five candidates would control all five appointments; with Cumulative Voting, a group gathering as little as 17 percent of the electorate and coordinating on a single candidate is guaranteed to have him elected. With a long historical tradition–it was the voting system electing the Illinois House of Representatives from 1870 to 1982[7]–and a current presence in both local jurisdictions and corporate elections, Cumulative Voting is emerging as a desirable alternative to the creation of majority-minority single-member districts. Advocated by civil rights scholars, most prominently Lani Guinier, it has been mandated by the courts as remedy to violations of the Voting Rights Act. Although the unfamiliar voting system is often resisted initially, follow-up studies suggest that it works indeed as expected, leading not only to the election of minority candidates, often for the first time, but substantially increasing the expenditure in public goods in minority neighborhoods.[8]

[7] It fell in 1982 in a voters-approved constitutional amendment, apparently as stand-by victim of a campaing to reduce the size of the House, after Representatives had voted themselves a large and unpopular salary increase.

[8] Pildes and Donoghue (1995) is a fascinating case study of the introduction of Cumulative Voting in Chilton County, Alabama, and takes the reader through initial reactions, parties' strategies, voters' education campaign, and final outcomes. Bowler, Donovan, and Brockington (2003) describe more widely the use of Cumulative Voting in local elections. The final

Storable Votes differ from Cumulative Voting because the two rules apply to different types of elections. Cumulative Voting is an instrument of indirect democracy: it is an *electoral system* that applies to the choice of representatives in a single election where multiple seats are to be filled. It is related to Limited Voting and Single Transferable Vote, other semi-proportional systems designed to give voice to minorities. Storable Votes are instead an instrument of direct democracy: *a decision-making rule* resolving disagreement over multiple proposals, each of which is presented to voters in a separate election with two alternatives only. They are designed to protect minorities in yes/no decisions.

As a result of having different fields of application, the two systems function differently. If there are five seats, Cumulative Voting selects five winners, the five candidates with more votes overall; Storable Votes select one winner in each contest, and the votes the winner receives are compared only to the votes received by its competing alternative. Thus a well-organized minority who coordinates its votes and has sufficient size can guarantee itself one seat with Cumulative Voting, but cannot guarantee itself any victory with Storable Votes. Because in this latter case all contests are binary, if both minority and majority assign highest importance to the same proposal and favor different outcomes, the majority will win. And if the minority has concentrated its bonus votes on that one proposal, the majority will win all contests. But in the logic of Storable Votes, such an outcome is in fact desirable: if majority and minority agree on the relative importance of the different proposals, then the majority "should" win, on the strength of its larger size alone. The goal of Storable Votes is to solve disagreements taking into account intensity of preferences, not to guarantee representation to the minority *per se*, as in the case of Cumulative Voting.

report of the Illinois Assembly on Political Representation and Alternative Electoral Systems (2001) discusses in detail the experience with Cumulative Voting in Illinois. The background paper prepared by Richard Briffault for the New Jersey Property tax Convention Tax Force (2004) is a short and very clear compendium of the US experience with Cumulative Voting, Limited Voting and Single Transferable Vote.

Minority victories, when they occur, are the by-product of the expression of the intensity of preferences that Storable Votes make possible.

Because Storable Votes apply to different decisions, voting occurs over time, again contrary to Cumulative Votes. The horizon could be very short, as in the case of a single meeting where several proposals are voted upon in sequence. But it could also be longer: chapter 2 in this book is motivated by the example of the board of the European Central Bank, deciding its interest rate policy on the first Thursday of each month, and suggests the possibility of endowing members with bonus votes valid for a calendar year. In this latter case, the budget of bonus votes creates a tie among decisions taken over the course of many months, and the intertemporal dimension of Storable Votes is very visible. The rights of individual voters are safeguarded by granting each voter the same budget of bonus votes, and thus treating all identically over the full time horizon, although over each individual decision, after the first, different voters in general will have a different number of votes at their disposal. Storable Votes build on what Adam Cox calls the "temporal dimension of voting rights."[9] The concept seems logically unexceptional, but with the exception of Cox' work, to my knowledge has not been invoked, let alone implemented, elsewhere.

Pildes and Donoghue (1995) report the first puzzled reactions to Cumulative Voting in Chilton County, Alabama: "[T]he silliest thing..." (p.9) , "I just know it's unconstitutional." (p.14). The resistance was caused by the very idea of casting multiple votes in a single election, which the public thought *must* violate the principle of one man one vote. And yet each voter entered the voting booth with the same number of votes. It seems likely that Storable Votes, with a conception of voting rights that relies on considering the full time horizon, would encounter even more resistance. But before worrying about public reactions, we need to ask whether in fact Storable Votes *should* be applied in practice. This

[9] Cox (2006). See also the response by Richard Pildes

is the purpose of the book.