

Preface

How do people coordinate their intentions? How can the mechanisms that they use be classified? Are there other mechanisms that would be attractive if people knew about them? And what about the moral dimension of collective decisions? How is it possible to know what a group should do when the members of the group disagree among themselves? When, if ever, is it possible to justify a claim that individuals ought to abide by collective decisions with which they disagree? These are questions that are addressed in this inquiry into collective decisions and voting.

At one level, this book is a smorgasbord of suggestions about voting procedures and other ways of making collective decisions, suggestions that might be of interest to people who would like to change collective decision procedures in ways that would potentially be satisfying to all concerned. It provides a set of definitions, taxonomies, properties, advantages and disadvantages of a wide variety of procedures. At another level, this book is simply a report on puzzles with respect to voting that have intrigued me, and the resolutions that I have found for them. The subsequent chapters of the book develop these ideas in a logical sequence. Here I offer a more historical account of my interest in them, resulting in a sequence that is roughly backward from the sequence of chapters in the book.

My earliest recollection of an interest in voting goes back to a day when I was eight years old, in San Francisco. Looking at the election results published in the San Francisco Chronicle, I noted that nine persons had been elected to the Board of Supervisors. The method of election was that each voter was allowed to vote for four candidates (or was it five?), and the nine with the most votes were the winners. I asked my father why they did it that way. He couldn't give me a satisfying answer, and the question of how a group of persons ought to be elected stuck with me. Eventually I learned that some people thought that the Single Transferable Vote was a satisfying way of electing a group. I could see elements of merit in that procedure, but there were obvious problems as well. In the mid 1980s I spent many dozens of hours on the question of what would be a satisfying way to elect a group, and finally came up with a recipe for comparing pairs of outcomes by the Single Transferable Vote,¹ which is discussed in Chapter 15, in the context of various ways of electing a group and why one way or another might be most satisfying.

When I was 12 years old I was nominated to be treasurer of my class at school. Being a whiz a math, I thought I would be a fine treasurer and salivated for the position. Michelle was also nominated. I don't remember whether I nominated

Michelle's best friend Charlotte, or whether I seconded the nomination. But I did one or the other, and my reaction was instinctive. I knew what I was doing. In the election, Charlotte received 11 votes, Michelle received 12, and I received 13 and became treasurer. I was immediately aware that I had won the office in a way that was not objectively fair and reasonable. The thought passed through my head that the honorable thing to do would be to suggest a run-off between Michelle and me. But I wanted the office too badly. Furthermore, there is some question as to whether such an innovation in the rules is proper in the middle of an election procedure. So there was no run-off. But the event imprinted on me the problem of what procedure should be used to determine the winner of an election when there are more than two candidates. In the early 1980s, when my Department was seeking to aggregate the opinions of its members as to whom it should hire, I found that the procedure that seemed most sensible to me was the "ranked pairs" rule,² which is discussed in Chapter 13, in the context of a long list of rules for electing one candidate, their properties, advantages and disadvantages. Chapter 12 provides definitions and explanations of the properties and criteria that are employed in Chapter 13.

As an undergraduate, I learned that there was something called welfare economics, which yielded conclusions that particular economic outcomes were good. That astounded me. It suggested that when people disagreed about what ought to be done, it would be possible for one group to say to another, "We have proven that what we want is good and what you want is not. Don't give us a hard time." I could not imagine how such an assertion could be valid. I wanted to understand how the rabbit got into the hat. Eventually I learned that if one posits normative axioms, then it is possible to use accepted logic to derive normative conclusions. A proof is simply a correspondence between assumptions and conclusions developed on the basis of accepted principles of logic. It establishes that something is true only if one knows that the axioms are true. And securing agreement that normative axioms are true is very difficult. I learned that there were at least four distinct normative traditions in twentieth century economics. These are reviewed in Chapter 3. My own view of what constitutes an adequate basis for disappointing someone who desires a particular outcome is developed in Chapter 6.

As an undergraduate, I also learned that there was something called the Arrow Theorem that showed that democracy was inevitably incoherent, or some such thing. It seemed impossible that I would ever be able to comprehend such a theorem, so while I was curious about it, I left it unexplored. There were further occasional references to it in my graduate education, but no explicit treatment of it. So I was quite startled when, as a new faculty member at Harvard, one of the first things I was expected to do was to explain the Arrow theorem to students in a course in intermediate microeconomic theory. Fortunately, there were graduate students who were ready to fill me in on what I needed to know. Unfortunately, I could not comprehend the theorem to my satisfaction and was quite dissatisfied

with the explanation I gave to the students in that class. Over the years, I continued to puzzle over what was actually going on in the Arrow theorem. I looked long and hard at Arrow's proof, as well as at proofs of similar theorems. I developed an idea about what was going on in the theorem, but I was aware that my ideas were not congruous with the general consensus of economic theorists. In the late 1980s, deciding that I had to resolve the matter, I wrote letters to a number of economic theorists requesting that they engage me in dialogue on the question of what the Arrow theorem meant. One person to whom I wrote, Georges Bordes, replied with a 12-page hand-written letter. I replied with a letter half as long, and he replied with another long letter. We each learned something from the other, so we decided to collaborate on a joint paper on what the Arrow Theorem really means.³ The understandings that I thereby acquired regarding of the meaning of the Arrow Theorem and its significance for the limits of a coherent democracy are explained in Chapter 10, in the context of a detailed parsing of Arrow's proof.

The Arrow Theorem is only half of the story of the logical limits of democracy. The other half is the Gibbard–Satterthwaite Theorem. This theorem employs the Arrow Theorem to prove that no voting procedure for more than two options can motivate voters in all circumstances to report their rankings of the options truthfully. As soon as I heard about this theorem, which was first published in the mid 1970s, I realized that it was very important, but I did not expect to be able to understand the proof. The first few times I presented it to students, I understood each step, but not its overall structure. Eventually I came to understand not only the structure of the proof, but also why one should expect that the theorem would be true. I have sought to convey that understanding in the presentation and discussion of the Gibbard–Satterthwaite Theorem in Chapter 11.

The Arrow Theorem and the Gibbard–Satterthwaite Theorem are both closely tied to the phenomenon of majority-rule cycles. Thus the significance of these theorems is related to the question of how frequently cycles can reasonably be expected to occur in real voting situations. It turns out that cycles will be common if not pervasive if there are no restrictions on placing redistributive proposals on the agenda, or if there are relatively few voters or relatively many options. But in typical voting situations with relatively few options and more than several hundred voters, cycles will be quite rare and can be dealt with in a way similar to the way that ties are dealt with. This is explained in Chapter 9.

One circumstance in which cycles cannot arise is if the options under consideration lie in a one-dimensional continuum and voters have “single-peaked” preferences in that dimension. In that case there is a majority-rule equilibrium at the preference of the median voter. There is in general no majority-rule equilibrium when the options lie in a continuum of more than one dimension, but there are nevertheless some voting procedures that might be attractive when the options lie in a continuum of several dimensions. These are explored in Chapter 14.

Discussions of voting usually presume without discussion that all votes will be weighted equally. (The one exception is voting on matters affecting a corporation, where the general presumption is that all shares will be weighted equally.) In my doctoral dissertation I dealt with an issue (proposals for changes in land use) where it is arguably sensible to weight votes according to prior estimates of the relative intensities of the voters' preferences.⁴ The general question of when it might be attractive to employ weighted voting, and why, has intrigued me over the years. In Chapter 8 I explore this question, along with the question of where ordinary majority rule for choices between two options came from, and what rationales it has.⁵

While the suggestion that I made in my doctoral dissertation was for weighting according to estimates of the intensities of voters' preferences, what I would have suggested if I could have imagined it was feasible was weighting according to the actual intensities of voters' preferences. But I could not imagine that such a thing was feasible. I could not imagine it despite the fact that one of my fellow graduate students, Edward Clarke, informed me, as I worked on my dissertation in 1968, that he had found a mechanism that motivated voters to report honestly the intensities of their preferences for public goods. I was sure he was mistaken and declined to take the time to examine his claim. But Clarke's ideas kept crossing my path.

In the 1970-71 academic year I was invited to contribute a chapter to the book *Public Prices for Public Products*. I contributed a chapter on a mechanism for successive revisions in the level of a public good.⁶ Clarke offered a version of his idea in the following chapter.⁷ That same year my dissertation advisor, George Tolley, who was also Clarke's advisor, invited Clarke and me to give back-to-back presentations of our ideas in a seminar. I had no understanding of Clarke's ideas after that seminar.

In the fall of 1973 I embarked on a post-doctoral fellowship at the Center for the Study of Public Choice, at Virginia Polytechnic Institute and State University. I told Gordon Tullock, who was also at the Center, that I planned to spend my post-doctoral year reviewing ways of making collective decisions. His response was, "Well, if you are going to do that, you really ought to include Ed Clarke's ideas in your review." He said that he had published Clarke's idea in *Public Choice*⁸ because it was important if it was right, but he did not understand it. (Tullock made most editorial decisions without consulting referees.) I reluctantly agreed to include it.

Three months into my post-doctoral fellowship, I agreed to teach and postpone the postdoctoral fellowship to the following year, because of an unexpected shortage of teachers, so it was not until the spring of 1975 that I confronted the need to deal with Clarke's ideas. And I still might have let it slip by if I had not encountered Martin Bailey at the meetings of the Public Choice Society that spring. I told Bailey what I was doing, and he told me that Clarke was right. With

that endorsement and the knowledge that I had said I would deal with it, I sat down to determine what was going on in Clarke's papers.

After ten minutes or so of concentrated effort, things suddenly clicked. I saw that Clarke had indeed devised a way of motivating people to report honestly the intensities of their preferences for public goods. I quickly organized an impromptu seminar to explain the ideas to others. Five minutes into my presentation, Gordon Tullock started finishing my sentences. He said that the idea needed a new exposition, so he and I agreed to collaborate on that. We named the idea the "demand revealing process".⁹ And that is how I happened to write my least original paper, which also became my most widely cited paper. Over the years, my understanding of the demand revealing process has evolved. My current understanding is presented in Chapter 16.

So far I have described a variety of ideas about voting with only occasional connections among them. Once I had all of the ideas in hand, it was natural to look for other ways in which they were connected. The result was a description of four different dimensions in which voting situations could differ. Each of the ideas that had captured my interest could be identified with a location in this four-dimensional space. This is explained in Chapter 7.

When I first conceived the idea of this book, I realized that voting was not the only way to make collective decision, and that a thorough account of voting should include a discussion of the relationship between voting and other ways of making collective decisions. After pondering alternative classifications of ways of making collective decisions for a number of years, I realized that there was a natural two-level taxonomy. This is explained in Chapter 2.

As soon as one sees all the different ways in which collective decisions might be made, the question that naturally arises is, "Why do we use the procedures we use instead of the available alternatives?" But before this question can be addressed, one must know what a good outcome of a collective decision procedure is. Thus Chapter 4 reviews criteria for evaluating collective decision procedures, and Chapter 5 applies these criteria to the ways of making collective decisions discussed in Chapter 2.

All the concern for collective decisions caused me to think about what a collective decision really is. So in Chapter 1 I define a decision, a collectivity, and hence a collective decision. The reasoning may surprise you.

Nicolaus Tideman

Notes

- 1 Nicolaus Tideman, "The Single Transferable Vote," *Journal of Economic Perspectives* 9/1 (1995): 27–38.

- 2 T.N. Tideman, "Independence of Clones as a Criterion for Voting Rules," *Social Choice and Welfare* 4 (1987): 85–206.
- 3 Georges Bordes and Nicolaus Tideman, "Independence of Irrelevant Alternatives in the Theory of Voting," *Theory and Decision* 30 (1991): 63–86.
- 4 T. Nicolaus Tideman, "Land Use Control through Administered Compensation," *Proceedings of the Meetings of the American Real Estate and Urban Economics Association* (1969): 287–93.
- 5 Nicolaus Tideman, "A Majority-Rule Characterization with Multiple Extensions," *Social Choice and Welfare* 3 (1986): 17–30.
- 6 Nicolaus Tideman, "Efficient Provision of Public Goods," in S. Mushkin (ed.), *Public Prices for Public Products*, (Washington, 1972), pp. 111–23.
- 7 Edward H. Clarke, "Multipart Pricing of Public Goods: An Example," in S. Mushkin, (ed.), *Public Prices for Public Products*, (Washington, 1972), pp. 125–30.
- 8 Edward H. Clarke, "Multipart pricing of Public Goods," *Public Choice* 11, (1971): 17–33.
- 9 Nicolaus Tideman and Gordon Tullock, "A New and Superior Process for Making Social Choices." *Journal of Political Economy* 84 (1976), 1145–59.