Introduction and Summary

INTRODUCTION

Behavioral Social Choice Research

This book develops conceptual, mathematical, methodological, and empirical foundations of *behavioral social choice research*. Behavioral social choice research (or, more briefly, *behavioral social choice*) encompasses two major interconnected paradigms: the development of *behavioral social choice theory* and the evaluation of that theory with empirical data on social choice behavior.

The fundamental purpose of a behavioral theory of social choice processes is the development of descriptive models for real actors' social choice behavior and the statistical evaluation of such models against empirical data. Our notion of behavioral social choice research builds on and, at the same time, complements much of classical social choice theory in the tradition of leading figures such as the Marquis de Condorcet, Duncan Black, Kenneth Arrow, and Amartya Sen. Most classic approaches follow an axiomatic, normative line of reasoning. They formulate desirable properties of "rational" social choice and provide numerous "possibility" or "impossibility" theorems that classify groups of such axioms into whether or not they lead to 'feasible' aggregation procedures, given various theoretical assumptions about the nature, domain, and distribution of individual preferences (McLean and Urken, 1995). A principal task of behavioral social choice research is to evaluate such normative benchmarks of rational social choice against empirical evidence on real world social choice behavior. Consistently throughout this book we attempt to evaluate

our models against a wide range of empirical evidence drawn from large-scale real-world data sets from three different countries. To the extent that classical/normative theories fail to be descriptive of observed social choice behavior, they motivate and inspire the development of (alternative) behavioral theories that complement classical approaches by descriptively capturing the social choice behavior of real actors.

We see our work as building on the pioneering literature that integrates formal models with the analysis of real world social choice data (e.g., Chamberlin et al., 1984; Felsenthal et al., 1986, 1993; Felsenthal and Machover, 1995; Laver and Schofield, 1990; Niemi, 1970; Riker, 1958). We provide a general probabilistic modeling and statistical sampling and inference framework for the descriptive theoretical and empirical investigation of social choice behavior of real-world decision makers, but we place a major emphasis on majority rule decision making (Condorcet, 1785). Our general framework is formulated in terms of an extremely broad domain of permissible preference representations and it is applicable to an extremely broad range of empirical rating, ranking, and choice paradigms.

Six Major Contributions

While we conceptualize behavioral social choice theory as encompassing a very broad spectrum of research paradigms,[1] we focus here exclusively on the foundations for such a theory. Our main contributions are sixfold:

1. We argue for the limited theoretical relevance and demonstrate the lack of empirical evidence for cycles in mass electorates by replacing "value restriction"[2] and similar classic domain restriction conditions, as well as the "impartial culture" assumption, with more realistic assumptions about preference distributions.
2. We expand the classical domains of permissible preference states by allowing for more general binary preference relations than linear or weak orders and by considering probabilistic representations of preference and utility.
3. We develop methodologies to (re)construct preference distributions from incomplete data, that is, data which do not provide either complete rankings or complete sets of pairwise comparisons.

[1] For example, in addition to the study of committee voting and mass election processes, we see behavioral social choice theory as encompassing the empirical study of coalitions, of information pooling (such as occurs in juries), and of a wide variety of other collective choice processes.

[2] A definition of this (and related) terms is provided later in the text.

4. We highlight the dependence of social choice results on assumed models of preference or utility.
5. We develop a statistical sampling and Bayesian inference framework that usually places tight upper and lower bounds on the probability of any majority preference relation (cycle or not).
6. We demonstrate that in situations where sampling may be involved, misestimation (i.e., erroneous evaluation) of the majority preferences is a far greater (and much more probable) threat to democratic decision making than majority cycles.

Conceptually, our work is heavily influenced by the foundations of behavioral economics and behavioral decision theory (see, e.g., Akerlof, 1984; Allais, 1953; Camerer et al., 2004; Harless and Camerer, 1994; Kahneman and Tversky, 1979; Kahneman et al., 1982; Luce, 1992, 2000; Luce and Suppes, 1965; Luce and von Winterfeldt, 1994; Plott and Levine, 1978; Shleifer, 2000; Simon, 1955; Smith, 1976, 1994; Suppes, 1961; Thaler, 1993a,b; Tversky, 1969; Tversky and Kahneman, 1974, 1981). Similar to much theoretical work in those fields, our approach to behavioral social choice theory is descriptive, yet mathematically formal. Also, like those fields, our approach draws theoretically, conceptually, and methodologically on mathematical psychology and statistics. In particular, we seek to build upon the early integrative perspective to the decision sciences of two outstanding theorists, Duncan Luce and Patrick Suppes (see esp. Luce and Suppes, 1965).

We recognize that, by concentrating on foundational work here, we omit other important and, in our opinion, 'higher order' aspects of a full-fledged behavioral theory, such as issues of strategic behavior that are so central to much ongoing work on social choice.[3] While we do not investigate the behavioral ramifications of game theoretic models here, we do believe that future descriptive work on strategic social choice behavior can build on the general foundations that we lay here.

We now briefly elaborate on our six major contributions.

1 Majority Cycles in a New Light. Majority rule has played an important role in the history of social choice theory. We believe it is fair to say that majority rule continues to be broadly viewed as the most important (or at least most influential) benchmark of rational social choice, while at the same time being put into question by important theoretical classical

[3] Since most of our data analyses use survey data on mass electorates, we do not believe that our substantive empirical conclusions are likely to be affected by our implicit assumption that the data are sincere reflections of the voter preferences.

work. In particular, our work complements three major theoretical developments in classical social choice theory that bear heavily on the study of majority rule decision making: Arrow's impossibility theorem and two subsequent strands of research that were motivated by Arrow's result, namely the literature on domain restrictions and the literature that draws on the impartial culture assumption.

ARROW'S THEOREM AND CYCLES. Arrow's famous "impossibility theorem" (1951) eliminates majority rule because one can easily construct hypothetical preference distributions under which majority rule violates one of Arrow's axioms, namely transitivity.[4] (However, see Saari's recent work, 1999, 2001b, 2001c for a novel theoretical perspective on Arrow's theorem and on majority cycles in particular.) In fact, the possibility of majority cycles continues to be a major reason why so many social choice scholars (and those influenced by them) argue that majority rule decision making is flawed. Yet, Arrow's approach requires an ideal social choice procedure to satisfy a certain set of axioms under all *possible* distributions of preferences (i.e., all possible preference profiles) over a given domain of possible preferences states, and proves that such a procedure fails to exist.[5] Behavioral social choice research can bring a new perspective to Arrow's theorem if it demonstrates that *actual* (voting) data are such that majority rule is overwhelmingly transitive. One interpretation of such an empirical result is that, for real data, one does not need to assume Arrow's condition that the social welfare function (choice procedure) is defined over all possible voter profiles. To phrase it differently, in this approach the feasibility of Arrow's ideal is no longer a theoretical question alone, but rather it becomes primarily an empirical and pragmatic one. If we take the pragmatic view that democratic decision making needs to be feasible only for actual preference distributions observed in the real world (i.e., the domain of the social choice function consists only of those distributions that have actually been observed), rather than for all conceivable

[4] The simplest example of the paradox of cyclical majorities, also called the "Condorcet paradox," occurs with three voters choosing among three alternatives. Label the alternatives as A, B, and C, assume that voter 1 has preference order ABC, voter 2 has preference order BCA and voter 3 has preference order CAB. Then majorities prefer A to B and B to C, which might lead us to expect that a majority prefers A to C. But this is not the case. Instead, a majority prefers C to A, and therefore these majority preferences form a cycle, i.e., group preferences are not transitive even though the individual preferences are transitive.

[5] One can avoid the impossibility result either by generalizing the approach to *nondeterministic* methods, leading to a *probabilistic dictator* (Pattanaik and Peleg, 1986; Tangiane, 1991), or by considering an infinite set of voters (Fishburn, 1970a).

distributions, then majority rule remains a strong contender to Arrow's challenge. In this book, we consider a number of real world preference distributions and find no conclusive empirical evidence that would substantiate the existence of majority cycles. One of our theoretical aims is thus to develop descriptively adequate constraints on preference distributions that explain the absence of majority cycles, that is, to answer Gordon Tullock's (1981) question "Why so much stability?"

DOMAIN RESTRICTIONS AND CYCLES. The literature on domain restrictions such as Sen's "value restriction" and Black's "single peakedness" (Black, 1958; Sen, 1966, 1970; see Gärtner, 2001, for a nice overview of that and related work) studies ways in which the domain of feasible individual preferences can be constrained in such a way that, for example, majority cycles are eliminated, regardless of the distribution of preferences over that restricted domain. Behavioral social choice research places a premium on the analysis of real world data. We present and analyze survey data from national election studies in three countries.[6] Although we discuss noteworthy exceptions that have to do with extremely homogenous national subpopulations, we find these empirical distributions to violate every imaginable domain restriction condition because of the fact that every permissible preference state (say, strict weak order) is reported by a very large number of respondents. Therefore, in our approach, rather than restricting the *domain* of permissible preferences, we restrict the *distribution* of preferences (and, in fact, vastly enlarge the domain of permissible preferences). We expect that in general, at least for mass electorates, every preference state that a voter is allowed to report will be reported by some (and possibly many) voters, no matter how 'strange' or 'irrational' we might consider that preference state to be. Nonetheless, we also expect that the *distribution* of such preference states will be such that majority rule is transitive.

THE IMPARTIAL CULTURE AND CYCLES. The literature on the impartial culture and related distributions (DeMeyer and Plott, 1970; Gehrlein and Fishburn, 1976b; Gehrlein and Lepelley, 1999, 2001; Jones et al., 1995; Klahr, 1966) investigates the probability that majority rule leads to a cycle, when preferences are randomly sampled from a given probability distribution over a given domain of permissible preferences, as one varies

[6] We refer to the distributions of choices (votes) in these surveys as "realistic" because, while they are likely to closely resemble the actual population distributions that the surveys were drawn from, our arguments hold without having to claim that the surveys accurately reflect the exact distribution of preferences in all details.

the number of candidates and/or the number of voters. An *impartial culture* is a uniform distribution over linear or weak orders and consists, itself, of a complete majority tie among all candidates. We call such distributions *cultures of indifference*. While random samples drawn from cultures of ndifference generate majority cycles with high probability, we show that samples drawn from any other culture will have majority cycles with probability approaching zero as the sample size increases, unless the culture itself has a cycle built in. One of the policy reversals suggested by this change of perspective is that high turnout (not, as is commonly claimed, low turnout) is desirable when using majority rule. In our empirical work, we show that none of our empirical preference distributions even remotely resembles a random sample from an impartial culture.[7]

Classical approaches fail to describe empirical data on preference distributions for mass electorates, and are often construed as suggesting overly pessimistic policy implications regarding the feasibility of democratic decision making. To summarize our discussion of cycles in somewhat simplistic terms, all three sets of results, Arrow's theorem, Sen's value restriction, and the impartial culture assumption, place constraints that are too strong, each in its own way. Arrow requires a procedure that not only 'works' in practice, but that 'works' under all conceivable circumstances. Sen's value restriction and similar domain restrictions rule out preference states that will invariably be held by some large number of people.[8] The impartial culture requires extreme symmetry on the distribution of preferences and the slightest violation of that symmetry completely upsets the policy implications one would draw. We believe that a behaviorally adequate theory of majority rule requires an extremely broad domain of permissible preference states and a descriptively adequate theory of real-world preference distributions, as well as adequate probabilistic modeling and statistical analysis tools to investigate empirical data. Much of this book is dedicated to that task. In particular, we replace classical assumptions by conditions stated in terms of "net preferences" over a very broad domain of permissible preferences to obtain results (in particular, about cycles) that we consider much more behaviorally realistic and that fare well in their evaluation against empirical data.

2 Generalizing Majority Rule. Our behavioral approach dictates that we reach beyond the deterministic linear or weak order individual

[7] None of them looks like a random sample from an intransitive culture either.
[8] In particular, we conjecture that such observations cannot simply be attributed to measurement error or other noise.

preferences, commonly assumed in classical theories, to include a broad range of mathematical representations of preference or utility that have been proposed and used in the decision sciences.

The various representations of "preference" that have been studied in the decision sciences include linear, weak, partial orders and semiorders, as well as probability distributions over such orders. All of these are special cases of binary relations and probability distributions over binary relations. Most mathematical representations of "utility" (as opposed to preference) rely on real-valued functions that map objects into their (possibly vector-valued) utility values. A very general conceptual framework to represent and quantify the variability of utilities is provided by random utility theory. Here, the utility of an object is (the value of) a random variable (or a random vector). Preference relations are in close correspondence with utility functions, and probability distributions over preference relations are in close correspondence with utility random variables, that is, random utility models. Just as probability distributions over preference relations generalize and include deterministic preference relations, so do random utilities generalize and include real-valued deterministic utility functions. In this book we formulate and investigate majority rule (for finitely many candidates/choice alternatives) in terms of arbitrary binary preference relations, probabilistic binary preference relations, arbitrary real (possibly vector) valued utility functions, and arbitrary real (possibly vector) valued random utility representations. In so doing, we integrate deterministic with probabilistic representations of preference, and we allow for multiple possible representations of utility, including random utility models.

We make no assumptions about where the randomness comes from. Probabilities may capture random error, random sampling, probabilistic mechanisms inside the decision maker's head, or they may simply quantify the 'proportion' of the population that satisfies some property. In particular, we require no independence assumptions. Since people interact and communicate, we allow individual preferences to be interdependent and/or systematically biased in the following sense: In a probability distribution over preference relations, interdependencies can be quantified through setting the probabilities of certain preference orders very high or very low; in the random utility framework, the interdependent nature of utilities is captured and quantified through the joint distribution of the utility random variables. We also make no assumptions about that joint distribution.

Our very general and unifying mathematical framework also establishes a close link between social choice theory and the other areas of the decision sciences, especially individual choice theory. Furthermore, this

approach allows us to derive results at a level of mathematical generality that includes and combines many known social choice theoretic results as special cases of a much broader framework.

3 (Re)constructing Preference Distributions. We argue that majority rule decision making has received limited empirical investigation because hardly any empirical data provide the input that is technically required to compute majority rule outcomes. Therefore, 'empirical' majority outcomes are usually hypothetical in nature to the extent that their computation from most available data requires various simplifying assumptions. We use a probabilistic approach to the measurement and inference of preference distributions from incomplete and/or randomly sampled empirical data in such a fashion as to encompass a multitude of empirical choice, rating, and ranking paradigms. In particular, our approach allows us to use the kinds of ballot data that are available from many real elections (e.g., plurality bloc voting, the single transferable vote, the alternative vote, and approval voting) and the kinds of data that are frequently collected in surveys (e.g., thermometer ratings and proximity data in national election studies) to make inferences about the distribution of underlying preferences. We illustrate our general approach with a particular emphasis on "feeling thermometer" survey data and on "approval voting" (i.e., "subset choice") election ballot data.

4 Model Dependence. We discuss the fundamental problem of model dependence of social choice results. As mentioned above, theoretical results about majority rule outcomes may dramatically change as one moves from one model of preferences and their distributions to another. For instance, our theoretical view of behaviorally appropriate domains of preference relations and preference distributions minimizes the likelihood of cycles and reverses a common policy recommendation about voter turnout. In the empirical domain we show how the analysis of empirical data may crucially depend on the implicit or explicit modeling assumptions that enter the analysis. When we analyze the same set of data with multiple competing models, we find that the inferred preferences and preference distributions can be dramatically different across models. Nonetheless, by and large, the nature of the majority preference relation is not dramatically affected. Moreover, virtually all our analyses of all data sets in this book share the common conclusion that majority preferences are transitive. (The exceptional analyses, i.e., those that do not rule out a cycle, do not provide strong evidence for the presence of a cycle either.)

5 A General Sampling and Inference Framework. Classical social choice theory relies on statistics largely in work on sampling from the impartial culture or other cultures of indifference. We develop a statistical sampling and Bayesian inference framework for the theoretical and empirical investigation of majority rule decision making in samples drawn from practically any distribution over any family of binary relations (over finitely many candidates). This allows us to place upper and lower bounds on the probability of any majority outcome (cycle or not) in a sample given almost any population, or in the population given almost any sample. Our method provides, by and large, very tight bounds on these probabilities. In particular, it allows us to place tight bounds on the probability of a cycle in a fashion that reaches far beyond the traditional cultures of indifference used in classical sampling work. More importantly, using this method, as long as we know the pairwise preference margins among each pair of candidates in a random sample, we can place upper and lower bounds on the probability of any conceivable majority preference relation in the population from which this sample was drawn. We believe that this constitutes a major milestone in our ability to study majority rule outcomes, both theoretically and empirically. To our knowledge, this is also the first full-fledged statistical framework for the investigation of classical social choice concepts such as, in this book, majority rule decision making. At a practical level, this book provides tools that can be applied to a variety of real world data sets in order to address a range of important issues.

6 Majority Preference Misestimation. There is, however, one aspect of our approach and findings that raises an important caveat about the processes of democracy. We argue that one must think of ballot casting and counting as noisy processes, which correctly record a given voter's current preference with probability less than one.[9] "Preference misestimations" (by which we mean erroneous estimates of the population majority preference) can also arise in estimating preferences from incomplete voting or survey sample data.[10] While we are not the first to note these issues,

[9] Also, we can think of the choices of individual voters (e.g., whether or not to vote, and, if to vote, whom to vote for) as nondeterministic (i.e., probabilistic) processes. Furthermore, we believe that many voters experience uncertainty about candidates' utilities and/or their own preferences, and thus the 'correct' preference of a voter may have to be conceptualized either as a statistic of a probability distribution over possible preference states, or as being a random draw from such a distribution.

[10] Note that the term "preference misestimation," as we use it, should not be confused with either strategic misrepresentation of preferences, as used in game theory, or with

we are the first to suggest, and empirically illustrate, that preference misestimations may pose a far greater threat to democratic decision making than majority cycles. Our analytic results show how to precisely assess both the likelihood of cycles and of majority preference misestimations using real-world data for which voter preferences between any pair of alternatives are known or can be estimated.

In illustration of these latter points, consider the uncertainties of the 2000 United States presidential election that transfixed America and the world from election day in November until early December 2000. In this context, we observe that, clearly, nobody was ever concerned about the possibility of a majority cycle among Bush, Gore, and Nader in that election. "Who won Florida?", "Who won the majority of the electoral college votes?" were the questions. The issue was whether or not the outcome of the election might 'accidentally' misestimate the electorate's 'true' preferences by reversing the top two choices. Hence, the central questions regarding that election concern accuracy of the assessment of majority preferences. With only two strong candidates, the possible occurrence of majority cycles was simply not an issue.[11]

SUMMARY

Throughout the book, we assume that the set of candidates/options/alternatives, that voters/survey respondents rate/rank/choose from, is finite.[12] We begin each chapter with a chapter summary, which we keep as informal and as nonmathematical as possible. In principle, each chapter can be read independently of the others. Whenever material from another chapter plays a critical role, we refer back to the chapter (definition,

Althaus' (1998, 2003) misrepresentation of public opinion due to "information effects." According to Althaus, the uneven social distribution of political knowledge in mass publics may distort the assessment of collective opinion in surveys.

[11] This is the concern which underlies the statutes mandating election recounts in various American states whenever the winning candidate's margin of victory falls below some specified threshold.

[12] Scholars interested in cycles have also studied what happens when the alternatives can be thought of as points in an n-dimensional space, giving us an infinite set of possible alternatives. The key result is that, in two or more dimensions, we can expect that the entire space of alternatives will be in a majority rule cycle (McKelvey, 1976, 1979). However, if we impose some constraints, e.g., voting one issue dimension at a time (Shepsle and Weingast, 1981) or ruling out pairwise votes involving alternatives that are highly similar to one another, or setting limits on how many alternatives may be considered (Feld and Grofman, 1996; Feld et al., 1989; Miller et al., 1989), we can avoid cycling (always, in the first case; almost always in the second and third cases).