OXFORD

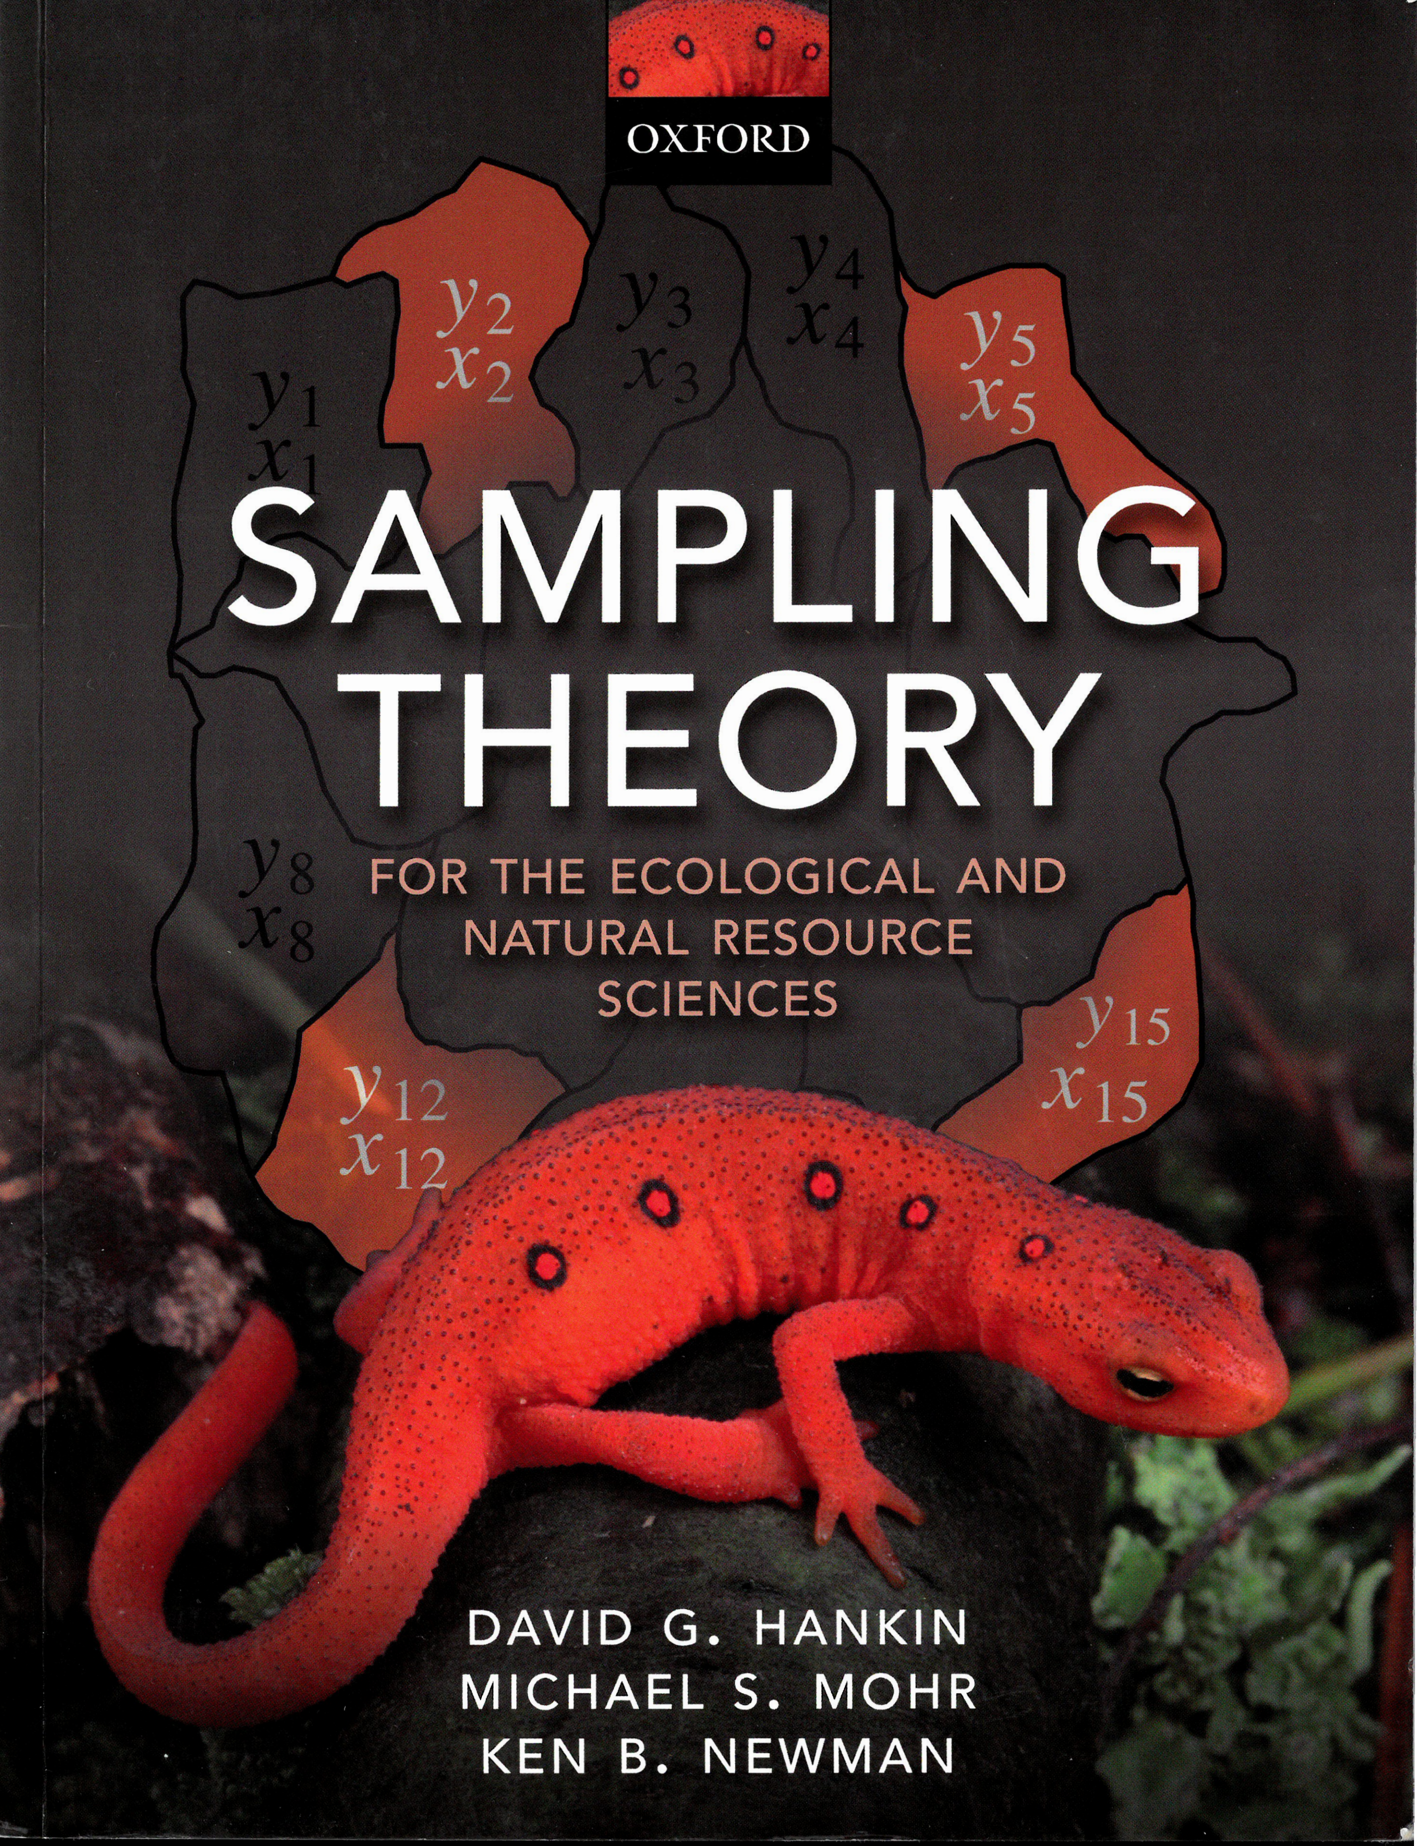

SAMPLING THEORY

FOR THE ECOLOGICAL AND NATURAL RESOURCE SCIENCES

DAVID G. HANKIN
MICHAEL S. MOHR
KEN B. NEWMAN

CHAPTER 1

Introduction

Fig. 1.1 Terrestrial form of the Pacific (coastal) giant salamander, *Dicamptodon tenebrosus*, Redwood National Park, and the original inspiration for this book's cover. Photo credit: D. Hankin.

In this text we attempt to present a rigorous but understandable introduction to the field of sampling theory. Sampling theory concerns itself with development of procedures for random selection of a subset of units, a sample, from a larger finite population, and with how to best use sample data to make scientifically and statistically sound inferences about the population as a whole. The inferences fall into two broad categories: (a) estimation of simple descriptive population parameters, such as means, totals, or proportions, for variables of interest associated with the units of finite populations, and (b) estimation of uncertainty associated with estimated parameter values. Although the targets of estimation are few and simple, estimates of means, totals or proportions see important and often controversial use in management of natural resources and in fundamental ecological research. For example, estimates of total population size and

Sampling Theory: For the Ecological and Natural Resource Sciences. David G. Hankin, Michael S. Mohr, and Ken B. Newman, Oxford University Press (2019). © David G. Hankin, Michael S. Mohr, and Ken B. Newman. DOI: 10.1093/oso/9780198815792.001.0001

associated trends in abundance play key roles in development of harvest policy for exploited species in fisheries and wildlife management settings, and in status reviews and development of recovery strategies for species listed under the Endangered Species Act. Estimates of species abundance or species proportions and associated measures of uncertainty also provide key input values for state variables, such as the true abundances or proportions, in *state-space* models of population dynamics that explicitly separate but also link a *state process* (say, trend in actual abundance of a species) and an *observation process* (process for estimating the value of a state variable) (Newman et al. 2014). Estimates of mean weights of individuals (e.g., of juvenile salmon in a small stream, or of young polar bears foraging off an ice shelf) may provide critically important data for energetic models of growth that may help predict the impacts of climate change on growth and survival of species. A temporal sequence of accurate estimates of mean water quality parameters in a lake may prove key for understanding the seasonal impacts of agricultural runoff and may lead to changes in land use practices. There are thus a wealth of practical and important applications of estimates of simple descriptive finite population parameters. Given the potentially controversial use of such estimates, being able to provide measures of the degree of uncertainty associated with these estimates is critical.

1.1 The design-based paradigm

We stress the classical *design-based* approach to sampling theory in this introductory text. For the design-based approach, a chance randomization scheme is used to select a sample of size n from a finite population consisting of N units. Each unit in the population is recognized to have fixed variable (attribute) values associated with it (e.g., population units may consist of individuals of a given species of tree and each tree, at a given point in time, has fixed diameter at breast height, total height, volume, etc.). Typically, one variable is of primary interest and we refer to this variable as the *target* variable, y. For unit i, we label its unit-specific fixed value as y_i, and we define the population parameters $\mathcal{T}_y = \sum_{i=1}^{N} y_i$ (total) and $\mu_y = \mathcal{T}_y/N$ (mean). We may also take advantage of values of an *auxiliary* variable associated with population units, denoted by x, that can be used to improve accuracy of estimation of target variable population parameters, either through direct incorporation in estimation formulas or indirectly through their influence on the randomization process used to select samples.

In design-based sampling, the randomized procedure(s) used to select sample units determines the *sample space*, the set of all possible samples of size n that can be selected (with or without replacement) from the population of size N. The set of associated sample probabilities that emerge from the randomized selection procedure, not necessarily equal, along with the unit labels that appear in the samples, can be used to calculate first and second order inclusion probabilities—the probabilities that unit i or that units i and j, respectively, appear in a sample of size n selected according to the randomized selection procedure. Associated with each possible sample s is a sample-specific estimate of a population parameter, e.g., $\hat{\mathcal{T}}_y(s)$, calculated by substituting the sample y values (and possibly also x values and/or first order inclusion probabilities) into equations (estimators) used to calculate estimates of target population parameters.

Uncertainty of estimation is measured by the variation in sample estimates over the sample space, variation induced by the interaction of the randomization scheme and estimator with the set of variable values (y, and sometimes also x) associated with population units. If the range of sample estimates is small and the average value of estimates is close to the target population parameter, then an estimator has good accuracy, but if the range of sample estimates is large, or the average value of estimates is far from the target population parameter, then an estimator has poor accuracy. Design-based sampling theorists have developed methods to estimate the variability in sample estimates over the full sample space—the *sampling variance* of an estimator—from just a single selected sample, and they use such estimates to characterize the uncertainty associated with an estimated descriptive population parameter.

We stress this design-based approach in the text for several reasons. First, in many controversial resource management contexts, *objectivity* of estimation is an important virtue of design-based sampling theory (Särndal et al. 1992 p. 21). No model assumptions are made regarding the distribution of the fixed y values over the population units and a chance randomization scheme is used to select sample units. Design-based estimators are therefore *robust* in the same sense that non-parametric statistical methods make no or mild assumptions regarding the distributions of random variables (Brewer and Gregoire 2009) and cannot be faulted by competing researchers or conflicting agencies who might allege deliberate bias resulting from model choice or purposive selection of sample units. Second, from Overton and Stehman (1995): "in finite sampling the populations being sampled and about which inferences are being made are real...and are subject to enumeration and exact description. In contrast, the populations of conventional statistics are hypothetical and represented by mathematical models". Treating the y and x associated with finite population units as fixed values has an undeniable conceptual validity in a finite population. Third, the distribution of x and y over the finite population units will often be poorly represented by any simple probability distribution. For example, asserting that the volume of timber on 1 ha lots follows a normal distribution may be a gross misrepresentation of reality. Fourth, we stress the design-based approach because we believe that it would be difficult to do justice to both design-based estimation and the model-based prediction (see below) approach in a relatively brief text designed for a first course in sampling theory.

We recognize, of course, that the use of models for inference in finite populations is an approach that is consistent with parametric statistics in general, that it has its own logic and validity, and that it may provide an alternative and sound approach for inference. In *model-based prediction* in finite populations, the actual variable values associated with population units are assumed to be realizations of random variables, and the finite population inference problem is to estimate the expected values (i.e., predict the values) of random variables in those $N-n$ units that *did not appear in the sample*, conditioned on the realized values of the random variables that have been observed in the sample of size n. Randomized selection methods used to select sample units are typically not relevant to model-based inference. Predicted values are based on assumed models that may include a dependence of the target variable on one or more auxiliary variables. When the distribution of variable values across units and the relationship between target and auxiliary variables are close to those assumed by models used for prediction, then model-based prediction can generate estimates with errors that may be smaller than those for design-based estimators. However, if the assumed models do not accurately characterize the actual distribution of the target variable and/or the relationship between the target and auxiliary variables, then errors may be large and expected values of model-based predictors may be biased when compared to the fixed population parameters that are the targets of estimation. In model-based prediction, emphasis is put on the value of *balanced* samples, for which the mean values of auxiliary values in the sample are purposively set to equal the mean values in the population (Valliant et al. 2000) to help ensure that model-based prediction is more robust to violations of assumed model structure.

While our emphasis is on design-based inference, we do provide a brief introduction to best linear unbiased (BLU) estimators and model-based prediction of finite population parameters in Chapter 7. Our intention is to provide readers with an appreciation of the fundamentally different perspectives that the design-based and model-based approaches take to estimation of finite population parameters. Both approaches rest on the same basic mathematical foundations of probability and statistics that we summarize in Appendix A. For readers desiring to learn more about model-based prediction, we recommend Valliant et al. (2000), Chambers and Clark (2012), and also Brewer's (2002) text which considers the benefits of combining design-based and model-based approaches.

1.2 Text content and orientation

Material is presented in this text at a level appropriate for and accessible to undergraduate seniors, beginning graduate students (not enrolled in a statistics program), and natural resources or environmental/ecological sciences professionals. Most material presented in the text can be fully understood and appreciated with modest mathematics and statistics training: considerable facility in algebra, including experience with multiple summation notation; at least one semester of calculus, including some experience with partial differentiation; an introductory course in statistical methods, along with an additional unspecified statistics class (including an exposure to analysis of variance); and some prior exposure to, though not necessarily programming proficiency in, R (R Core Team 2018), the statistical/programming/graphics language and environment used to generate the numerical results and figures presented in this text. Depth of appreciation for the subject matter generally would be greatly enhanced by formal background in probability theory, but we do not assume that readers have such background. Our intention is to foster interest in readers to develop or hone the necessary mathematical skills to follow through derivations of important results, especially if these skills are not in their formal backgrounds or are "rusty with lack of use", so that users of this text will be able to develop and successfully apply their own sampling strategies in their own unique application settings.

We believe that the most natural and successful order for working through the text is to directly follow the order of chapters as they appear in the text. Chapter 2 provides a very gentle and largely non-quantitative introduction to important sampling theory definitions and concepts (sampling frame, sampling design, sampling strategy) and properties of estimators (expected value, bias, sampling variance, mean square error) when sampling from finite populations, including justification of randomized selection of units as compared to judgment or purposive selection of units. Important concepts or terms appear in boldface type at first usage to emphasize their importance to readers. Sophistication of presentation and complexity of mathematical notation and sampling strategies gradually increases through the next several chapters. Chapter 3 provides an introduction to equal probability selection methods, with and without replacement, and Chapters 4–6 illustrate application of these equal probability selection methods in the contexts of systematic, stratified, and equal size cluster sampling, respectively. Chapter 7 introduces the explicit use of auxiliary variable information in estimators (ratio and regression estimation), while also providing a brief introduction to BLU estimators and model-based prediction of finite population parameters. Chapter 8 introduces unequal probability selection methods, with and without replacement, and associated estimators (notably the famous Hansen–Hurwitz and Horvitz–Thompson estimators). Selection of units with unequal probabilities, with probabilities based on auxiliary variable values, represents an alternative way to take advantage of auxiliary variable information. Material covered through Chapter 8 is all devoted to sampling from what we call *simple frames*, where the samples are selected directly from the individual units in the populations. (Stratified sampling is a special case where a simple frame is used for independent sample selection from within disjoint sets or strata of population units.) Sampling strategies with complex frames are considered in the following two chapters. Chapter 9 is devoted primarily to two-stage implementation of multi-stage sampling strategies, but we also provide the general framework for many designs with more than two stages of sampling. Chapter 10 is devoted to two-phase implementation of multi-phase sampling. The material presented in the chapters thus far mentioned comprises the "basic stuff" of classical design-based sampling theory and is covered, in various ways, in most sampling theory texts available today.

1.3 What distinguishes this text?

We hope to distinguish our text from previously published sampling theory texts in several different ways. First, we have tried to make the material presented in the text accessible to the reader who may not be "homozygous for the math gene". Toward that end, though we do provide derivations of many important sampling theory results, we also provide small sample space illustrations of the behavior of most of the sampling strategies (selection method and estimator applied to a sampling frame) that we consider. This allows readers to see numerically, rather than imagine in thought, what the sample space actually is. For individuals "heterozygous (or recessive) for the math gene", the small sample space examples often make abstract concepts like sampling variance come alive and have tangible and very real meaning. Sample space examples have been used relatively rarely in existing texts, but we acknowledge their previous effective use in the very small but very useful and repeatedly reissued (1962, 1964, 1968, 1976, 1984) text by Stuart (1984), *The Ideas of Sampling*. Sample space examples also allow us to avoid presentation of numerous fictitious examples for observed data from a single sample; the sample spaces provide an abundance of implicit calculations that can be worked through and confirmed by students and readers if they wish.

Second, we have presented some material that has not frequently or routinely been covered in existing texts. Appendix A provides a review of the mathematical foundations of design-based sampling theory including counting techniques, basic principles of probability theory, properties of discrete random variables, key discrete probability distributions and their clear links to design-based or probability sampling, as well as an overview of the use of Lagrange multipliers and the delta method. Chapter 11 covers adaptive sampling, a topic covered in Thompson (1992, 2002, 2012) and Thompson and Seber (1996) but rarely in other sampling theory texts. More unusual is our coverage of spatially balanced sampling methods [Generalized Random Tesselation Stratified sampling (GRTS) and Balanced Acceptance Sampling (BAS)] in Chapter 12, and our coverage of designs and estimators for sampling through time (including rotating panels and dual frame sampling) in Chapter 13.

Third, we have tried to make extensive use of graphs and figures to convey important sampling theory concepts and especially to help readers understand how the performances (sampling variance or net relative efficiency) of competing sampling strategies may be compared to one another. We also use graphs to portray the shapes of the sampling distributions of design-based estimators, which are often not approximately normal until $n > 20$ or so.

Fourth, we have tried to develop and use a notation that is more consistent than in many previously published sampling theory texts. We use Greek letters to represent population parameters ($\mathcal{T}$ = total, μ = mean, π = proportion, and σ^2 = finite population variance, with a divisor of $N-1$); we always use a caret or *hat* to indicate an estimator or estimate (as in $\hat{\mu}_y$ rather than $\bar{y}$); we use lowercase italic y and x to denote the target and auxiliary population variables, respectively, and uppercase letters to denote random variables; we use uppercase italic S to indicate the population units selected as the random outcome of a sample selection process (*sampling experiment*, see Section A.6.1), and $\hat{\theta} = \sum_{i \in S} f(y_i, x_i)$ as a generic representation of a design-based estimator; we use lowercase italic s to identify the set of units in a particular realized sample of size n and $\sum_{i \in s} y_i$ to indicate the sum of the y values in this sample, thereby avoiding the *re-labeling* of population units implicit in the more commonly used $\sum_{i=1}^{n} y_i$; and we use combinations of uppercase and lowercase letters to define properties relating two finite population variables and properties of estimators (e.g., $\mathrm{Cov}(x,y)$ for covariance, R for the ratio of the means of y and x, $V(\hat{\theta})$ for sampling variance of an estimator $\hat{\theta}$).

Fifth, in the text we provide occasional brief reference to simple R functions or packages that should prove useful for working through problems or for implementing computer-intensive sample selection methods. We do not, however, as a rule present code for what we believe to be relatively simple calculations for which students and other readers should learn to write their own code. Instead, we provide web-based access (www.oup.co.uk/companion/hankin) to either existing R packages or code that we have written ourselves for tasks that we believe to be beyond the level of programming proficiency and/or statistical sophistication that one can reasonably expect from students or professionals using this text as a first course in sampling theory. Examples of such R programs include several methods of selection of unequal probability without replacement samples and for calculation of associated first and second order inclusion probabilities, and programs for implementing the spatially balanced sampling methods described in Chapter 12.

The natural resources/environmental/ecological orientation of the text will be most evident from the problem exercises associated with those chapters for which we believe reasonable "classroom" problems can be developed. Many of these problems have roots in real-world natural resource settings in which the authors have worked. We hope that these problem exercises will help students and readers learn how to answer what is often the most difficult question: "How can we apply these ideas in practice?" We also have frequent references to natural resources/environmental settings/applications within the text of chapters and in the *Chapter comments* section at the conclusion of each chapter, and we begin each chapter with a relevant photograph, usually illustrating a natural world setting where sampling theory can address a topic of interest.

1.4 Recommendations for instructors

If the text were used to teach a one semester sampling theory class to a group consisting of senior undergraduates and beginning graduate students, we recommend using Appendix A, the review of mathematical foundations, as a *reference* and that it only be worked through in detail if students in the course have unusually strong backgrounds and interests in mathematics and probability theory. Instead, we recommend that when concepts or techniques covered in Appendix A are used or noted in other chapters, then the instructor should refer students to Appendix A and perhaps cover a few of the relevant concepts or techniques at that time. It is important that students do not think "this is a math class". Material through Chapter 11 can be easily covered in a one semester course. Content of Chapters 12 and 13 should at least be mentioned and practical importance discussed, but could be covered in some detail if material in other chapters was covered very rapidly.

We highly recommend addition of a weekly (two hour) laboratory/practicum session in addition to lecture/discussion presentation of the material presented in this text. These lab sessions can be used to illustrate computer implementation and evaluation of sampling strategies (using R), especially for unequal probability and spatially balanced selection methods, as well as to illustrate use of R to solve assigned problems (thereby assisting students in developing programming and graphics skills). Lab sessions can also be used as effective class exercises to contrast the performance of simple random sampling and judgment sampling (see agate population exercise described in Section 2.6), and for student presentations of their experiences in development and execution of sampling strategies for estimating simple parameters such as the total number of preserved specimens in a university's ichthyology (fish) collection; the average cost (and average quality rating) of a slice of pizza sold within the local county; the proportion of plants in an herbarium that are in bloom at the time of a survey; the fraction of trees of a given species in a well-defined plot that have diameter at breast height exceeding 1 m; the number of books in a library; the numbers of pieces of trash (carefully defined) per unit area or length in several well-defined areas (neighborhoods) of a small town or city; and numerous other ideas that students will dream up if they are given enough freedom. There is no exercise that is a more effective practical learning experience than identifying a finite population (or study area) of interest, picking a well-identified target of estimation, thinking through how measurements can be made, what sampling frames might make most sense, what selection methods might prove most effective, what estimators might be used, and then executing a small-scale survey, ideally using two alternative sampling strategies, and reporting back to classmates on their experiences.

Finally, after working through this text or taking a class using this text, we hope that all readers or students will develop a deep appreciation for the following joke about statisticians in a wildlife setting.

Three statisticians set out on a bow hunt for deer and hope to make a successful kill of a buck. They are fortunate to locate a fine buck at a decent distance. Statistician 1 lets loose his arrow, but it misses the buck by exactly 6.1 feet to the east of the buck's heart. Statistician 2 lets loose his arrow, but it misses the buck by exactly 6.1 feet to the west of the buck's heart. In excitement, Statistician 3 exclaims "We got him!"[1]

1.5 Sampling theory: A brief history

Sampling theory is a narrowly specialized but important area of the larger body of statistical theory. In Kendall and Stuart's three volume *Advanced theory of statistics* (Kendall and Stuart 1977, 1979, Kendall et al. 1983), the subject of sampling theory occupied just

[1] We thank Mathew Krachey, former sampling theory student and now Ph.D. statistician, for passing this joke on to us many years back.

84 of 2,000 pages, about 4%. But work in this small area of statistical theory has remained very active. Volume 6 (Sampling) of the Handbook of Statistics (Krishnaiah and Rao 1988) was 594 pages long, and Volumes 29A and 29B of the Handbook of Statistics (Pfeffermann and Rao 2009a,b) contain 1,340 pages combined devoted to sampling theory.

Numerous published articles have reviewed the roots of sampling theory. Bellhouse (1988), Hansen (1987), and Hansen et al. (1985), in particular, provide excellent reviews of the history and development of the design-based or probability sampling theory that is stressed in this text, along with references for many other previously published histories of sampling theory. Bellhouse (1988) argues persuasively that the development of sampling theory can be seen to have had several paradigm shifts (Kuhn 1970) which were attributable to novel new sampling ideas but also to aggressive promotion of these ideas by their proponents. In the brief description of the historical development of design-based sampling theory that follows, we rely heavily on Bellhouse (1988).

The historical origins of sampling theory can be traced to desires of government agencies for knowledge of population size, production of agricultural crops, average income, and other descriptive population parameters that are important for description and management of a country's economy and social structure. Kiaer (1897, as cited in Bellhouse 1988), Director of the Norwegian Central Bureau of Statistics, was a vigorous early advocate of *representative sampling* which was an alternative to the prohibitively expensive or logistically unfeasible complete enumerations that were judged necessary prior to the acceptance of sampling as an alternative way to obtain (acceptably accurate) estimates of population parameters. At the time, a representative sample was thought to be a sample of units that was an approximate *miniature* of the population, a "correct representation of the whole". (Kruskal and Mosteller (1979a,b,c, 1980) provide historical reviews of the rich, varied, often incorrect and inconsistent meanings that the term *representative sampling* has had in scientific, non-scientific, and statistical settings.) By about 1925, the basic idea of representative sampling (instead of complete enumeration) had been widely accepted, but there had not yet been any agreement that randomized rather than purposive selection was required to achieve a representative sample. From about 1925 to 1935, the British statistician Arthur Bowley became a staunch advocate for equal probability randomized selection of samples. Bellhouse (1988) provides direct quotations from Bowley that suggest that he understood the advantages of systematic sampling as compared to simple random sampling (SRS), but Bellhouse also points out that Bowley's (1926) monograph presented theoretical justifications and suggested methods for calculating uncertainty of estimation for both randomized and purposive selection of units.

All accounts of the history of design-based sampling theory give Neyman's (1934) paper preeminent importance in the paradigm shift that lead to widespread adoption of the randomized selection, design-based approach to generating sample estimates from finite populations. First, Neyman argued (and graphically illustrated) that a then recent purposive large-scale Italian survey had invoked tenuous assumptions of linearity in relations among some variables, had been purposively balanced on some variables, but had generated sample estimates that were inconsistent with available census counts for other variables. Second, Neyman laid out the theory for optimal allocation (ignoring costs) in stratified sampling (*Neyman allocation*, Chapter 5). In optimal allocation, units from different strata typically have different probabilities of inclusion in the overall sample, the first time this idea had been shown to be optimal. Third, Neyman characterized the notion of a *representative sample* as one that would allow estimation of accuracy "irrespectively of the unknown properties of the population studied", based on the behavior of an estimator over all possible randomized sample selections. Fourth, he defined the concept of confidence intervals, based on the notion of repeated sampling within the design-based framework. A few years later, Neyman (1938) developed the theory of double sampling, including the first joint use of cost and variance functions to optimize performance of a given sampling strategy (frame + design + estimator).

The paradigm shift to widespread adoption of randomized design-based sampling strategies for estimation of descriptive population parameters generated an explosive proliferation of important articles and texts. Particularly important were the papers concerning estimation based on unequal probability with and without replacement selection by Hansen and Hurwitz (1943) and Horvitz and Thompson (1952), respectively. Among the many fine sampling theory texts that were produced prior to 1970, we highlight those written by Deming (1950), Cochran (1953), Hansen et al. (1953a,b), Yates (1960), Murthy (1967), and Raj (1968). We stop this listing of texts at 1970 because that year may date another paradigm shift among a minority of statisticians concerned with estimation in finite populations.

Bellhouse (1988) credits Godambe (1955) with generating a paradigm shift that stimulated interest in the theoretical foundations of sampling theory (e.g., Cassel et al. 1977). Godambe's (1955) proof of the non-existence of a unique minimum variance design-based unbiased estimator of the finite population mean stimulated exploration of the use of models in finite population inference. Royall (1970) advanced the *model-based prediction* approach for estimation of finite population parameters. He showed that, if there were a linear model relation between auxiliary and target variables, passing through the origin with variation increasing with x, then the ratio estimator was the best model-based predictor of the finite population total. According to this perspective, randomized selection of units was irrelevant, or even foolish. Indeed, minimum variance of this estimator, given the assumed model, was achieved if the sample were purposively selected to include those n units with largest auxiliary variable values. Given this sample of units, predictions would need to be made only for units with a smaller value of x, for which errors would be smaller because, under the model, variation is less at smaller values of x. (Brewer (1963) had earlier presented similar findings for model-based ratio estimation in finite populations, including the finding that accuracy would be best if the sample contained the units with the largest values of x, but his work was not cited by Royall (1970) and may not have been aggressively promoted.)

Royall's (1970) paper and his subsequent vigorous advocacy for model-based prediction in finite populations generated heated exchanges between practicing proponents of design-based sampling theory and model-based prediction advocates (see, e.g., the paper, comment, rejoinder by Hansen et al. (1983a), Royall (1983), and Hansen et al. (1983b), respectively) for at least the next two to three decades. At least one prominent sampling theorist was first persuaded by the new model-based prediction paradigm but then, on further reflection, proclaimed himself a staunch design-based advocate [contrast Smith (1994) with Smith (1976)], much to the chagrin of Royall (Royall 1994). Publication of Särndal et al.'s (1992) text on model-assisted sampling represents one very important response to the model-based prediction paradigm. This text (and many earlier published papers by Särndal and colleagues) illustrates how models might be used to guide development and adoption of competing sampling strategies, but then estimation remains firmly within the design-based framework. Särndal (2010) provides a recent review of the role of models in sampling theory. Among other things, he notes that the notion of purposive selection of sample units that are *balanced* on a set of auxiliary variables, to achieve *model-robust* estimation within the model-based prediction framework (Valliant et al. 2000), has a new parallel in the design-based setting—the *cube method* (Deville and Tillé 2004, Tillé 2011)—where randomization is used to achieve balancing on auxiliary variables but inference remains within the design-based framework. Thus, one might argue that sampling theory has come *full circle* from the purposively balanced samples that had been recommended prior to Neyman's (1934) landmark paper, but thereafter rejected, to a contemporary recognition in both model-based prediction and design-based estimation that balancing is desirable in both designed-based and model-based settings. Current methods for selecting spatially balanced samples (Chapter 12) are consistent with this notion of balancing used within a design-based framework.

Finally, it is important to distinguish between sample surveys that have a strictly *descriptive* objective—estimation of descriptive population parameters—as compared to those that have an *analytic* or *causal* objective (Hansen 1987). In analytic or causal surveys, primary interest lies not in estimation of the descriptive population parameters themselves, but instead on analysis of the possible *relationships* among survey variables. There is no alternative to conjecture of hypothetical model relations in the context of analytic surveys, but statistical analysis of these relations should acknowledge the probability structure of the sampling design under which the survey data were collected (Korn and Graubard 1999, Lohr 2010 Chapter 10).