Introduction

Some great ideas are born in a flash of inspiration, perhaps announced to the world by a pathbreaking paper. R. A. Fisher's 1925 article on maximum likelihood estimation is a classic example. Nothing at all like that happened with exponential families. The theory accrued slowly over a period extending roughly between 1932 and 1970. Applications lagged behind, a turning point being the advent of logistic regression and McCullagh and Nelder's 1983 book on generalized linear models.

A salient fact is that no one person is credited with the development of exponential families, though it will be clear from these notes that Fisher's work was instrumental. The name "exponential familes" is comparatively recent. Until the late 1950s they were often referred to as "Koopman–Darmois–Pitman" families (crediting three prominent statisticians working separately in three different countries); the awkward nomenclature suggests only minor importance being attached to the ideas.

Figure 1 gives a rough schematic history of Twentieth Century statistics. The inner circle represents normal theory, the preferred venue of classical methodology. Exact inference – t tests, F tests, chi-squared statistics, ANOVA, multivariate analysis – was feasible inside the circle. Outside the circle was a general theory based on large-sample asymptotic approximations involving Taylor series, Edgeworth expansions, and the central limit theorem. A few special exact results lay outside the normal circle, relating to especially tractable distributions such as the binomial, Poisson, gamma and beta families. These are the figure's green stars. A happy surprise, though a slowly emerging one beginning in the 1930s, was that the special cases were all examples of a powerful general construction, *exponential families*, the intermediate circle in Figure 1. Within this circle, "almost exact" inferential calculations are possible, where any necessary approximations can be pictured in simple geometric diagrams. Such diagrams play a major role in what follows.

Two complementary types of mathematical development can be labeled

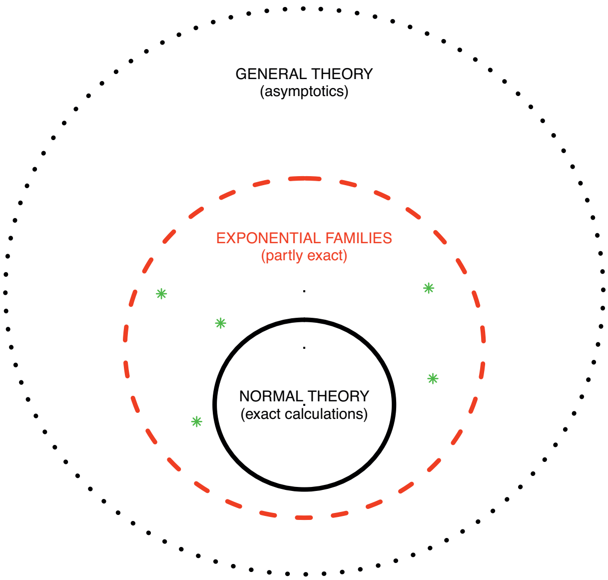

Figure 1 Three levels of statistical modeling.

as "theorem-proof" and "descriptive". The former has a worst-case aspect: "What I stated remains true even under the worst possibility of what I've allowed." Descriptive mathematics, of the kind encountered in introductions to calculus or linear algebra, is less pessimistic: disregarding pathologies, interest centers on the broad central run of useful results. This book pursues the theory of exponential families from a descriptive point of view, aiming at the parts of the theory most useful for applications.

Exponential families, used flexibly, can gracefully bridge the gap between statistical theory and its practice. These notes collect a large amount of material useful in statistical applications, but also of value to the theoretician trying to frame a new situation without immediate recourse to asymptotics. My own experience has been that when I can put a problem, applied or theoretical, into an exponential family framework, a solution is often imminent. There are almost no proofs in what follows, but hopefully enough motivation and heuristics to make the results believable if not obvious. References are given when this doesn't seem to be the case. Readers who desire a more thorough approach can look in Sundberg's excellent text, *Statistical Modelling by Exponential Families* (2019), or Brown's IMS monograph, *Fundamentals of Statistical Exponential Families* (1986).