RANDOM DISTRIBUTION FUNCTIONS

LESTER E. DUBINS and DAVID A. FREEDMAN
UNIVERSITY OF CALIFORNIA, BERKELEY

1. Introduction and summary

How can one choose, at random, a probability measure on the unit interval? This paper develops the answer announced in [4]. Section 1, which can be skipped without logical loss, gives a fairly full but slightly informal account. The formalities begin with section 2. All later sections are largely independent of one another. Section 10 indexes definitions made in one section but used in other sections.

A *distribution function* F on the closed unit interval I is a nondecreasing, nonnegative, real-valued function on I, normalized to be 1 at 1 and continuous from the right. To each F there corresponds one and only one probability measure $|F|$ on the Borel subsets of I, with $F(x)$ equal to the $|F|$-measure of the closed interval $[0, x]$, for all $x \in I$. Choosing at random a probability on I is therefore tantamount to choosing at random a distribution function on I.

A *random distribution function* $\mathbf{F}$ is a measurable map from a probability space $(\Omega, \mathfrak{F}, Q)$ to the space Δ of distribution functions on the closed unit interval I, where Δ is endowed with its natural Borel σ-field, that is, the σ-field generated by the customary weak* topology. The distribution of $\mathbf{F}$, namely $Q\mathbf{F}^{-1}$, is a *prior* probability measure on Δ. Of course, $\mathbf{F}$ is essentially the stochastic process $\{F_t, 0 \leq t \leq 1\}$ on $(\Omega, \mathfrak{F}, Q)$, whose sample functions are distribution functions: $F_t(\omega)$ is $\mathbf{F}(\omega)$ evaluated at t. Therefore, this paper can be thought of as dealing with a class of random distribution functions, with a class of stochastic processes, or with a class of prior probabilities. Similar priors are treated in [9], [11], [16], and [17].

Since the indefinite integral of a distribution function is convex, this paper can also be thought of as dealing with a class of random convex functions, but we do not pursue this idea.

Which class of random distribution functions does this paper study? A *base probability* μ is a probability on the Borel subsets of the unit square S, assigning measure 0 to the corners $(0, 0)$ and $(1, 1)$. Each such μ determines a random distribution function $\mathbf{F}$ and so a prior probability P_μ on Δ, which will now be described, by explaining how to select a value of $\mathbf{F}$, that is, a distribution function F, at random.

ASSUMPTION. *For ease of exposition, we assume throughout this section that μ concentrates on, that is, assigns probability 1 to, the interior of S.*

This paper was supported in part by the National Science Foundation Grant GP-2593, and in part by a grant to Freedman from the Sloan Foundation.