

*NSF-CBMS Regional Conference Series
in Probability and Statistics
Volume 2*

EMPIRICAL PROCESSES: THEORY AND APPLICATIONS

David Pollard

Yale University

Sponsored by the Conference Board of the Mathematical Sciences
Supported by the National Science Foundation
Published by the Institute of Mathematical Statistics
and the American Statistical Association

*NSF-CBMS Regional Conference Series
in Probability and Statistics
Volume 2*

**EMPIRICAL
PROCESSES:
THEORY
AND
APPLICATIONS**

David Pollard
Yale University

Institute of Mathematical Statistics
Hayward, California

American Statistical Association
Alexandria, Virginia

Conference Board of the Mathematical Sciences

*Regional Conference Series
in Probability and Statistics*

Supported by the
National Science Foundation

The production of the *NSF-CBMS Regional Conference Series in Probability and Statistics* is managed by the Institute of Mathematical Statistics: Paul Shaman, IMS Managing Editor; Jessica Utts, IMS Treasurer; and Jose L. Gonzalez, IMS Business Manager.

Library of Congress Catalog Card Number: 90-50461

International Standard Book Number 0-940600-16-1

Copyright © 1990 Institute of Mathematical Statistics

All rights reserved

Printed in the United States of America

To Gai and James

Contents

vii	PREFACE	
1	SECTION 1.	Introduction
6	SECTION 2.	Symmetrization and Conditioning
9	SECTION 3.	Chaining
14	SECTION 4.	Packing and Covering in Euclidean Spaces
21	SECTION 5.	Stability
29	SECTION 6.	Convex Hulls
35	SECTION 7.	Maximal Inequalities
39	SECTION 8.	Uniform Laws of Large Numbers
43	SECTION 9.	Convergence in Distribution and Almost Sure Representation
50	SECTION 10.	Functional Central Limit Theorems
57	SECTION 11.	Least Absolute Deviations Estimators for Censored Regressions
65	SECTION 12.	Random Convex Sets
70	SECTION 13.	Estimation from Censored Data
75	SECTION 14.	Biased Sampling
84	REFERENCES	

Preface

These notes grew from lectures I gave at the University of Iowa in July of 1988, as part of the *NSF-CBMS Regional Conference Series*. The conference was ably organized by Tim Robertson and Richard Dykstra. I am most grateful to them for giving me the opportunity to experiment on a live and receptive audience with material not entirely polished. I also appreciate the suggestions and comments of Richard Dudley. Much of the lecture material was repackaging of ideas originally due to him.

In reworking the lecture notes I have tried (not always successfully) to resist the urge to push the presentation to ever higher levels of generality. My aim has been to introduce just enough technique to handle typical nontrivial asymptotic problems in statistics and econometrics. Of course the four substantial examples that represent the applications part of the lectures do not exhaust the possible uses for the theory. I have chosen them because they cleanly illustrate specific aspects of the theory, and also because I admire the original papers.

To anyone who is acquainted with the empirical process literature these notes might appear misleadingly titled. Empirical process theory usually deals with sums of independent (identically distributed) random variables $f(\xi_i(\omega))$, with f running over a class of functions $\mathcal{F}$. However I have chosen to present results for sums of independent stochastic processes $f_i(\omega, t)$ indexed by a set T . Such a setting accommodates not only the relatively straightforward generalization to nonidentically distributed $\{\xi_i\}$, but also such simple modifications as a rescaling of the summands by a factor that depends on i and ω . It has often irked me that the traditional notation cannot handle summands such as $f(\xi_i)/i$, even though the basic probabilistic method is unaffected.

The cost of the modified notation appears in two ways. Some familiar looking objects no longer have their usual meanings. For example, $\mathcal{F}$ will now stand for a subset of $\mathbb{R}^n$ rather than for a class of functions. Also, some results, such as the analogues in Section 4 of the standard Vapnik-Červonenkis theory, become a trifle less general than in the traditional setting. The benefits include the natural reinterpretation of the Vapnik-Červonenkis property as a sort of dimensionality

concept, and the transformation of $\mathcal{L}^2(P_n)$ pseudometrics on classes of functions into the usual (ℓ_2) Euclidean distances in $\mathbb{R}^n$.

Several friends and colleagues at Yale and elsewhere have influenced the final form of the notes. Ariel Pakes provided well thought-out comments on the paper Pollard (1989), in which I tried out some of the ideas for the Iowa lectures. Probing questions from Don Andrews firmed up some particularly flabby parts of the original lecture notes. A faithful reading group struggled through the first half of the material, finding numerous errors in what I had thought were watertight arguments. Deborah Nolan tested a slightly more correct version of the notes on a graduate class at the University of California, Berkeley. (The rate at which bugs appeared suggests there might even be other embarrassing errors lying in wait to confuse future unsuspecting readers.) I thank them all.

Very recently I had the good fortune to obtain a copy of the manuscript by Ledoux and Talagrand (1990), which provides an alternative (often mathematically more elegant) treatment for some of the material in the theory part of the notes. I am grateful to those authors for enlightening me.

As always Barbara Amato well deserves my thanks and admiration for her ability to convert unruly drafts into beautiful T_EXnical documents. Paul Shaman and Jose Gonzalez contributed valuable editorial advice. The manuscript was prepared using the $\mathcal{A}\mathcal{M}\mathcal{S}$ -T_EX macros of the American Mathematical Society.

ISBN 0-940600-16-1

SECTION 1

Introduction

As it has developed over the last decade, abstract empirical process theory has largely been concerned with uniform analogues of the classical limit theorems for sums of independent random variables, such as the law of large numbers, the central limit theorem, and the law of the iterated logarithm. In particular, the Glivenko-Cantelli Theorem and Donsker's Theorem, for empirical distribution functions on the real line, have been generalized and extended in several directions. Progress has depended upon the development of new techniques for establishing maximal inequalities for sums of independent stochastic processes. These inequalities can also be put to other uses in the asymptotic theory of mathematical statistics and econometrics. With these lecture notes I hope to explain some of the theoretical developments and illustrate their application by means of four nontrivial and challenging examples.

The notes will emphasize a single method that has evolved from the concept of a Vapnik-Červonenkis class of sets. The results attained will not be the best possible of their kind. Instead I have chosen to strive for just enough generality to handle the illustrative examples without having to impose unnatural extra conditions needed to squeeze them into the framework of existing theory.

Usually the theory in the literature has concerned independent (often, also identically distributed) random elements $\xi_1, \xi_2, \dots$ of an abstract set Ξ . That is, for some σ -field on Ξ , each ξ_i is a measurable map from a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ into Ξ . For each n , the $\{\xi_i\}$ define a random probability measure on the set Ξ : the *empirical measure* P_n puts mass $1/n$ at each of the points $\xi_1(\omega), \dots, \xi_n(\omega)$. Each real-valued function f on Ξ determines a random variable,

$$P_n f = \frac{1}{n} \sum_{i \leq n} f(\xi_i(\omega)).$$

For fixed f , this is an average of independent random variables, which, under appropriate regularity conditions and with the proper standardizations, will satisfy a law of large numbers or a central limit theorem. The theory seeks to generalize these classical results so that they hold uniformly (in some sense) for f ranging

over various classes $\mathcal{F}$ of functions on Ξ .

In asymptotic problems, $\mathcal{F}$ is often a parametric family, $\{f(\cdot, t) : t \in T\}$, with T not necessarily finite dimensional. One can then simplify the notation by writing $f_i(\omega, t)$ instead of $f(\xi_i(\omega), t)$. In my opinion, this is the most natural notation for the methods that will be developed in these notes. It accommodates gracefully applications where the function f is allowed to change with i (or n). For example, in Section 11 we will encounter a triangular array of processes,

$$f_{ni}(\omega, t) = |y_i(\omega)^+ - (x'_i\theta_0 + z'_{ni}t)^+| - |y_i(\omega)^+ - (x'_i\theta_0)^+| \quad \text{for } i = 1, \dots, n,$$

generated by a reparametrization of a censored regression. The $\{z_{ni}\}$ will be constructed from the deterministic vectors $\{x_i\}$ by means of a transformation that depends on n . Such processes do not fit comfortably into the traditional notation, but their analysis depends on the same symmetrization and conditioning arguments as developed in the literature for the empirical measure P_n .

The notation also allows for transformations that depend on i , as with the $f_i(\omega, t)/i$ that will appear in Section 8. It also eliminates an unnecessary notational distinction between empirical processes and partial-sum processes, bringing both closer to the theory for sums of independent random elements in Banach space. In these notes, however, I will concentrate on problems and methods that are usually identified as belonging to empirical process theory.

The general problem to be attacked in the next six sections will be that of finding probabilistic bounds for the maximal deviation of a sum of independent stochastic processes,

$$S_n(\omega, t) = \sum_{i \leq n} f_i(\omega, t),$$

from its expectation,

$$M_n(t) = \mathbb{P}S_n(\cdot, t) = \sum_{i \leq n} \mathbb{P}f_i(\cdot, t).$$

That is, we will seek to bound $\Delta_n(\omega) = \sup_{t \in T} |S_n(\omega, t) - M_n(t)|$. In applications the f_i will often acquire a second subscript to become a triangular array. But, since most of the argument is carried out for fixed n , there is no need to complicate the notation prematurely.

For a general convex, increasing function Φ on $\mathbb{R}^+$, Section 2 will derive a bound for $\mathbb{P}\Phi(\Delta_n)$. The strategy will be to introduce a more variable process,

$$L_n(\sigma, \omega) = \sup_t \left| \sum_{i \leq n} \sigma_i f_i(\omega, t) \right|,$$

defined by means of a new sequence of independent random variables $\{\sigma_i\}$, each σ_i taking only the values $+1$ and -1 , both with probability $1/2$. We will find that $\mathbb{P}\Phi(\Delta_n)$ is less than $\mathbb{P}\Phi(2L_n)$.

With ω held fixed, L_n is a very simple process indexed by a subset of $\mathbb{R}^n$,

$$\mathcal{F}_\omega = \{(f_1(\omega, t), \dots, f_n(\omega, t)) : t \in T\}.$$

The indexing of the points of $\mathcal{F}_\omega$ by T will become irrelevant; the geometry of $\mathcal{F}_\omega$

will be all that matters. In terms of the usual inner product on $\mathbb{R}^n$,

$$L_n(\boldsymbol{\sigma}, \omega) = \sup_{\mathbf{f} \in \mathcal{F}_\omega} |\boldsymbol{\sigma} \cdot \mathbf{f}|.$$

Section 3 will establish a general inequality for processes like this, but indexed by fixed subsets of $\mathbb{R}^n$; it will be applied conditionally to L_n . The inequality will take the form of a bound on an *Orlicz norm*.

If Φ is a convex, increasing function on $\mathbb{R}^+$ with $0 \leq \Phi(0) < 1$, the Orlicz norm $\|Z\|_\Phi$ of a random variable Z is defined by

$$\|Z\|_\Phi = \inf\{C > 0 : \mathbb{P}\Phi(|Z|/C) \leq 1\},$$

with $+\infty$ as a possible value for the infimum. If $\mathbb{P}\Phi(|Z|/C_0) < \infty$ for some finite C_0 , a dominated convergence argument shows that $\mathbb{P}\Phi(|Z|/C) \rightarrow \Phi(0) < 1$ as $C \rightarrow \infty$, which ensures that $\|Z\|_\Phi$ is finite. If one identifies random variables that are equal almost everywhere, $\|\cdot\|_\Phi$ defines a norm on the space $\mathcal{L}^\Phi$ of all random variables Z for which $\|Z\|_\Phi < \infty$. (The space $\mathcal{L}^\Phi$ is even complete under this norm, a property we will not need.) In the special case where $\Phi(x) = x^p$ for some $p \geq 1$, the norm $\|\cdot\|_\Phi$ coincides with the usual $\|\cdot\|_p$, and $\mathcal{L}^\Phi$ is the usual space of random variables with finite p^{th} absolute moments. Finiteness of $\|Z\|_\Phi$ places a constraint on the rate of decrease for the tail probabilities via the inequality

$$\begin{aligned} \mathbb{P}\{|Z| \geq t\} &\leq \mathbb{P}\Phi(|Z|/C)/\Phi(t/C) \\ &\leq 1/\Phi(t/C) \quad \text{if } C = \|Z\|_\Phi. \end{aligned}$$

The particular convex function

$$\Psi(x) = \frac{1}{5} \exp(x^2)$$

would give tails decreasing like $\exp(-Ct^2)$ for some constant C . Such a rate of decrease will be referred to as subgaussian tail behavior.

The inequality in Section 3 will be for processes indexed by a subset $\mathcal{F}$ of $\mathbb{R}^n$. It will take the form of a bound on the particular Orlicz norm,

$$\left\| \sup_{\mathbf{f} \in \mathcal{F}} |\boldsymbol{\sigma} \cdot \mathbf{f}| \right\|_\Psi,$$

involving the *packing numbers* for the set $\mathcal{F}$. [The packing number $D(\epsilon, \mathcal{F})$ is the largest number of points that can be packed into $\mathcal{F}$ with each pair at least ϵ apart.] In this way we transform the study of maximal inequalities for Δ_n into a study of the geometry of the set $\mathcal{F}_\omega$.

Section 4 will make the connection between packing numbers and the combinatorial methods that have evolved from the approach of Vapnik and Červonenkis. It will develop the idea that a bounded set $\mathcal{F}$ in $\mathbb{R}^n$ that has a weak property shared by V -dimensional subspaces should have packing numbers like those of a bounded subset of $\mathbb{R}^V$. The three sections after that will elaborate upon the idea, with Section 7 summarizing the results in the form of several simple maximal inequalities for Δ_n .

Section 8 will transform the maximal inequalities into simple conditions for uniform analogues of the law of large numbers. Sections 9 and 10 will transform them into uniform analogues of the central limit theorem—functional limit theorems that

are descendants of Donsker's Theorem for the empirical distribution function on the real line. The approach there will depend heavily on the method of almost sure representation.

Section 9 will be the only part of these notes where particular care is taken with questions of measurability. Up to that point any measurability difficulties could be handled by an assumption that T is a Borel (or analytic) subset of a compact metric space and that each of the functions $f_i(\omega, t)$ is jointly measurable in its arguments ω and t . Such niceties are left to the reader.

The challenging applications will occupy the last four sections.

The key to the whole approach taken in these notes is an important combinatorial lemma, a refinement of the so-called *Vapnik-Červonenkis Lemma*. It deserves an immediate proof so that the reader might appreciate the simplicity of the foundation upon which all else rests.

In what follows, $\mathcal{S}$ will denote the set of all 2^n possible n -tuples of $+1$'s and -1 's. The pointwise minimum of two vectors σ and η in $\mathcal{S}$ will be denoted by $\sigma \wedge \eta$. The symbol $\#$ will denote cardinality of a set. Inequalities between vectors in $\mathcal{S}$ should be interpreted coordinatewise.

(1.1) BASIC COMBINATORIAL LEMMA. *For each map η from $\mathcal{S}$ into itself there exists a one-to-one map θ from $\mathcal{S}$ onto itself such that $\theta(\sigma) \wedge \sigma = \eta(\sigma) \wedge \sigma$ for every σ .*

PROOF. Replacing $\eta(\sigma)$ by $\eta(\sigma) \wedge \sigma$ if necessary, we may simplify the notation by assuming that $\eta(\sigma) \leq \sigma$ for every σ . Then for each σ in $\mathcal{S}$ we need to choose $\theta(\sigma)$ from the set $K(\sigma) = \{\alpha \in \mathcal{S} : \alpha \wedge \sigma = \eta(\sigma)\}$. For each subset $\mathcal{A}$ of $\mathcal{S}$ define

$$K(\mathcal{A}) = \bigcup_{\sigma \in \mathcal{A}} K(\sigma).$$

The idea is to prove that $\#K(\mathcal{A}) \geq \#\mathcal{A}$, for every choice of $\mathcal{A}$. The combinatorial result sometimes known as the Marriage Lemma (Dudley 1989, Section 11.6) will then imply existence of a one-to-one map θ from $\mathcal{S}$ onto itself such that $\theta(\sigma) \in K(\sigma)$ for every σ , as required.

For the special case where $\eta(\sigma) = \sigma$ for every σ , the inequality holds trivially, because then $\sigma \in K(\sigma)$ for every σ , and $K(\mathcal{A}) \supseteq \mathcal{A}$ for every $\mathcal{A}$. The general case will be reduced to the trivial case by a sequence of n modifications that transform a general η to this special η .

The first modification changes the first coordinate of each $\eta(\sigma)$. Define a new map η^* by putting $\eta^*(\sigma)_i = \eta(\sigma)_i$ for $2 \leq i \leq n$, and $\eta^*(\sigma)_1 = \sigma_1$. Let $K^*(\sigma)$ be the subset of $\mathcal{S}$ defined using η^* . We need to show that

$$\#K(\mathcal{A}) \geq \#K^*(\mathcal{A}).$$

To do this, partition $\mathcal{S}$ into 2^{n-1} sets of pairs $\{\beta^-, \beta^+\}$, where each β^- differs from its β^+ only in the first coordinate, with $\beta_1^- = -1$ and $\beta_1^+ = +1$. It is good enough to show that

$$\#[K(\mathcal{A}) \cap \{\beta^-, \beta^+\}] \geq \#[K^*(\mathcal{A}) \cap \{\beta^-, \beta^+\}]$$

for every such pair. This will follow from: (i) if $\beta^- \in K^*(\mathcal{A})$ then $K(\mathcal{A})$ contains both β^- and β^+ ; and (ii) if $\beta^+ \in K^*(\mathcal{A})$ but $\beta^- \notin K^*(\mathcal{A})$ then at least one of β^- and β^+ must belong to $K(\mathcal{A})$.

Let us establish (i). Suppose $\beta^- \in K^*(\mathcal{A})$. Then, for some σ in $\mathcal{A}$, we have $\beta^- \in K^*(\sigma)$, that is $\beta^- \wedge \sigma = \eta^*(\sigma)$. For this σ we must have

$$-1 = \min[-1, \sigma_1] = \eta^*(\sigma)_1 = \sigma_1.$$

Since $\eta(\sigma) \leq \sigma$, it follows that $\eta(\sigma)_1 = -1$ and $\eta(\sigma) = \eta^*(\sigma)$. Thus $\beta^+ \wedge \sigma = \beta^- \wedge \sigma = \eta(\sigma)$, as required for (i).

For (ii), suppose β^+ belongs to $K^*(\mathcal{A})$ but β^- does not. Then, for some σ in $\mathcal{A}$, we have $\beta^+ \wedge \sigma = \eta^*(\sigma) \neq \beta^- \wedge \sigma$. Both vectors $\beta^+ \wedge \sigma$ and $\beta^- \wedge \sigma$ agree with $\eta^*(\sigma)$, and hence with $\eta(\sigma)$, in coordinates 2 to n . Either $\eta(\sigma)_1 = -1 = (\beta^- \wedge \sigma)_1$ or $\eta(\sigma)_1 = \sigma_1 = +1 = (\beta^+ \wedge \sigma)_1$. Thus either $\eta(\sigma) = \beta^+ \wedge \sigma$ or $\eta(\sigma) = \beta^- \wedge \sigma$, as required for (ii).

We have now shown that the modification in the first coordinate of the η map reduces the cardinality of the corresponding $K(\mathcal{A})$. A similar modification of η^* in the second coordinate will give a similar reduction in cardinality. After n such modifications we will have changed η so that $\eta(\sigma) = \sigma$ for all σ . The corresponding $K(\mathcal{A})$ has cardinality bigger than the cardinality of $\mathcal{A}$, because it contains $\mathcal{A}$, but smaller than the cardinality of the $K(\mathcal{A})$ for the original η . $\square$

REMARKS. Several authors have realized the advantages of recasting abstract empirical processes as sums of independent stochastic processes. For example, Alexander (1987b) has developed general central limit theorems that apply to both empirical processes and partial-sum processes; Gaenssler and Schlumprecht (1988) have established moment inequalities similar to one of the inequalities that will appear in Section 7.

The proof of the Basic Combinatorial Lemma is based on Lemmas 2 and 3 of Ledoux and Talagrand (1989). It is very similar to the method used by Steele (1975) to prove the Vapnik-Červonenkis Lemma (see Theorem II.16 of Pollard 1984).

SECTION 2

Symmetrization and Conditioning

In this section we begin the task of bounding $\mathbb{P} \Phi(\sup_t |S_n(\cdot, t) - M_n(t)|)$ for a general convex, increasing function Φ on $\mathbb{R}^+$. The idea is to introduce more randomness into the problem and then work conditionally on the particular realization of the $\{f_i\}$. This is somewhat akin to the use of randomization in experimental design, where one artificially creates an extra source of randomness to ensure that test statistics have desirable behavior conditional on the experimental data.

As a convenience for describing the various sources of randomness, suppose that the underlying probability space $(\Omega, \mathcal{A}, \mathbb{P})$ is a product space,

$$\Omega = \Omega_1 \otimes \cdots \otimes \Omega_n \otimes \Omega'_1 \otimes \cdots \otimes \Omega'_n \otimes \mathcal{S},$$

equipped with a product measure

$$\mathbb{P} = \mathbb{P}_1 \otimes \cdots \otimes \mathbb{P}_n \otimes \mathbb{P}'_1 \otimes \cdots \otimes \mathbb{P}'_n \otimes \mathbb{P}_\sigma.$$

Here $\Omega'_i = \Omega_i$ and $\mathbb{P}'_i = \mathbb{P}_i$. The set $\mathcal{S}$ consists of all n -tuples $\sigma = (\sigma_1, \dots, \sigma_n)$ with each σ_i either $+1$ or -1 , and $\mathbb{P}_\sigma$ is the uniform distribution, which puts mass 2^{-n} on each n -tuple.

Let the process $f_i(\cdot, t)$ depend only on the coordinate ω_i in Ω_i ; with a slight abuse of notation write $f_i(\omega_i, t)$. The Ω'_i and $\mathbb{P}'_i$ are included in order to generate an independent copy $f_i(\omega'_i, t)$ of the process. Under $\mathbb{P}_\sigma$, the σ_i are independent sign variables. They provide the randomization for the symmetrized process

$$S_n^\circ(\omega, t) = \sum_{i \leq n} \sigma_i f_i(\omega_i, t).$$

We will find that this process is more variable than S_n , in the sense that

$$(2.1) \quad \mathbb{P} \Phi(\sup_t |S_n(\cdot, t) - M_n(t)|) \leq \mathbb{P} \Phi(2 \sup_t |S_n^\circ(\cdot, t)|)$$

The proof will involve little more than an application of Jensen's inequality.

To take advantage of the product structure, rewrite the lefthand side of (2.1) as

$$\mathbb{P}_1 \otimes \cdots \otimes \mathbb{P}_n \Phi \left(\sup_t \left| \sum_{i \leq n} [f_i(\omega_i, t) - \mathbb{P}'_i f_i(\omega'_i, t)] \right| \right).$$

We can replace the $\mathbb{P}'_i$ by $\mathbb{P}'_1 \otimes \cdots \otimes \mathbb{P}'_n$, then pull that product measure outside the sum, without changing the value of this expression. The argument of Φ , and hence the whole expression, is increased if we change

$$\sup_t |\mathbb{P}'_1 \otimes \cdots \otimes \mathbb{P}'_n \cdots| \quad \text{to} \quad \mathbb{P}'_1 \otimes \cdots \otimes \mathbb{P}'_n \sup_t |\cdots|.$$

Jensen's inequality then gives the upper bound

$$\mathbb{P}_1 \otimes \cdots \otimes \mathbb{P}_n \otimes \mathbb{P}'_1 \otimes \cdots \otimes \mathbb{P}'_n \Phi \left(\sup_t \left| \sum_{i \leq n} f_i(\omega_i, t) - f_i(\omega'_i, t) \right| \right).$$

The last expression would be unaffected if we interchanged any ω_i with its ω'_i , because $\mathbb{P}_i = \mathbb{P}'_i$. More formally, the $2n$ -fold product measure is invariant under all permutations of the coordinates generated by interchanges of an ω_i with its ω'_i . For each σ in $\mathbb{S}$, the $2n$ -fold expectation would be unchanged if the integrand were replaced by

$$\Phi \left(\sup_t \left| \sum_{i \leq n} \sigma_i [f_i(\omega_i, t) - f_i(\omega'_i, t)] \right| \right),$$

which, because Φ is convex and increasing, is less than

$$\frac{1}{2} \Phi \left(2 \sup_t \left| \sum_{i \leq n} \sigma_i f_i(\omega_i, t) \right| \right) + \frac{1}{2} \Phi \left(2 \sup_t \left| \sum_{i \leq n} \sigma_i f_i(\omega'_i, t) \right| \right).$$

These two terms have the same $2n$ -fold expectation. Averaging over all choices of σ , ω_i , and ω'_i , we arrive at a $(2n+1)$ -fold expectation that is equal to the righthand side of the symmetrization inequality (2.1). Notice that the auxiliary ω'_i randomization has disappeared from the final bound, which involves only $\omega = (\omega_1, \dots, \omega_n)$ and σ .

For each ω , the sample paths of the processes trace out a subset

$$\mathcal{F}_\omega = \{ (f_1(\omega, t), \dots, f_n(\omega, t)) : t \in T \}$$

of $\mathbb{R}^n$. Consolidating the product $\mathbb{P}_1 \otimes \cdots \otimes \mathbb{P}_n$ into a single $\mathbb{P}_\omega$, and reexpressing inequality (2.1) in terms of the usual inner product on $\mathbb{R}^n$, we get a neater looking bound.

(2.2) THEOREM. *For each convex, increasing Φ ,*

$$\mathbb{P} \Phi \left(\sup_t |S_n(\cdot, t) - M_n(t)| \right) \leq \mathbb{P}_\omega \mathbb{P}_\sigma \Phi \left(2 \sup_{\mathbf{f} \in \mathcal{F}_\omega} |\sigma \cdot \mathbf{f}| \right). \quad \square$$

The inner expectation, with respect to $\mathbb{P}_\sigma$, involves a very simple process indexed by a (random) subset $\mathcal{F}_\omega$ of $\mathbb{R}^n$. The fact that T indexes the points of the sets $\mathcal{F}_\omega$ now becomes irrelevant. The sets themselves summarize all we need to know about the $\{f_i(\omega, t)\}$ processes. If we absorb the factor 2 into the function Φ , the problem has now become: find bounds for $\mathbb{P}_\sigma \Phi(\sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}|)$ for various convex Φ and various subsets $\mathcal{F}$ of $\mathbb{R}^n$.

REMARKS. There are many variations on the symmetrization method in the empirical process literature. In the original paper of Vapnik and Červonenkis (1971) the symmetrized process was used to bound tail probabilities. I learned about the simplifications arising from the substitution of moments for tail probabilities by reading the papers of Pisier (1983) and Giné and Zinn (1984). Symmetrization via moments also works with more complicated processes, for which tail probabilities are intractable, as in the papers of Nolan and Pollard (1987, 1988) on U-processes. In their comments on Pollard (1989), Giné and Zinn have traced some of the earlier history of symmetrization, with particular reference to the theory of probability in Banach spaces.

SECTION 3

Chaining

The main aim of the section is to derive a maximal inequality for the processes $\sigma \cdot \mathbf{f}$, indexed by subsets of $\mathbb{R}^n$, in the form of an upper bound on the Ψ norm of $\sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}|$. [Remember that $\Psi(x) = 1/5 \exp(x^2)$.] First we need a bound for the individual variables.

(3.1) LEMMA. *For each $\mathbf{f}$ in $\mathbb{R}^n$, the random variable $\sigma \cdot \mathbf{f}$ has subgaussian tails, with Orlicz norm $\|\sigma \cdot \mathbf{f}\|_{\Psi}$ less than $2|\mathbf{f}|$.*

PROOF. The argument has similarities to the randomization argument used in Section 2. Assume the probability space is a product space supporting independent $N(0, 1)$ distributed random variables $g_1, \dots, g_n$, all of which are independent of the sign variables $\sigma_1, \dots, \sigma_n$. The absolute value of each g_i has expected value

$$\gamma = \mathbb{P}|N(0, 1)| = \sqrt{2/\pi}.$$

By Jensen's inequality,

$$\begin{aligned} \mathbb{P}_{\sigma} \exp \left(\sum_{i \leq n} \sigma_i f_i / C \right)^2 &= \mathbb{P}_{\sigma} \exp \left(\sum_{i \leq n} \sigma_i f_i \mathbb{P}_g |g_i| / \gamma C \right)^2 \\ &\leq \mathbb{P}_{\sigma} \mathbb{P}_g \exp \left(\sum_{i \leq n} \sigma_i |g_i| f_i / \gamma C \right)^2. \end{aligned}$$

The absolute value of any symmetric random variable is independent of its sign. In particular, under $\mathbb{P}_{\sigma} \otimes \mathbb{P}_g$ the products $\sigma_1 |g_1|, \dots, \sigma_n |g_n|$ are independent $N(0, 1)$ random variables. The last expected value has the form $\mathbb{P} \exp(N(0, \tau^2)^2)$, where the variance is given by

$$\tau^2 = \sum_{i \leq n} (f_i / \gamma C)^2 = |\mathbf{f}|^2 / \gamma^2 C^2.$$

Provided $\tau^2 < 1/2$, the expected value is finite and equals $(1 - 2|\mathbf{f}|^2 / \gamma^2 C^2)^{-1}$. If we choose $C = 2|\mathbf{f}|$ this gives $\mathbb{P} \Psi(\sigma \cdot \mathbf{f} / C) \leq 1$, as required. $\square$

The next step towards the maximal inequality is to bound the Ψ norm of the maximum for a finite number of random variables.

(3.2) LEMMA. *For any random variables $Z_1, \dots, Z_m$,*

$$\left\| \max_{i \leq m} |Z_i| \right\|_{\Psi} \leq \sqrt{2 + \log m} \max_{i \leq m} \|Z_i\|_{\Psi}.$$

PROOF. The inequality is trivially satisfied if the right-hand side is infinite. So let us assume that each Z_i belongs to $\mathcal{L}^{\Psi}$. For all positive constants K and C ,

$$\begin{aligned} \Psi(\max |Z_i|/C) &\leq \Psi(1) + \int_1^{\infty} \{K \max |Z_i|/C > Kx\} \Psi(dx) \\ &\leq \Psi(1) + \int_1^{\infty} \frac{\Psi(K \max |Z_i|/C)}{\Psi(Kx)} \Psi(dx) \\ &\leq \Psi(1) + \sum_{i \leq m} \int_1^{\infty} \frac{\Psi(K Z_i/C)}{\Psi(Kx)} \Psi(dx). \end{aligned}$$

If we choose $C = K \max \|Z_i\|_{\Psi}$ then take expectations we get

$$\begin{aligned} \mathbb{P} \Psi(\max |Z_i|/C) &\leq \Psi(1) + m \int_1^{\infty} \frac{1}{\Psi(Kx)} \Psi(dx) \\ &= \frac{e}{5} + m(K^2 - 1)^{-1} \exp(-K^2 + 1). \end{aligned}$$

The right-hand side will be less than 1 if $K = \sqrt{2 + \log m}$. (Now you should be able to figure out why the $1/5$ appears in the definition of Ψ .) $\square$

Clearly the last lemma cannot be applied directly to bound $\|\sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}|\|_{\Psi}$ if $\mathcal{F}$ is infinite. Instead it can be used to tie together a sequence of approximations to $\sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}|$ based on an increasing sequence of finite subsets $\mathcal{F}$. The argument, which is usually referred to as *chaining*, depends on the geometry of $\mathcal{F}$ only through the size of its *packing numbers*. To begin with, let us consider a more general—more natural—setting: a stochastic process $\{Z(t) : t \in T\}$ whose index set T is equipped with a metric d . [Actually, d need only be a pseudometric; the argument would not be affected if some distinct pairs of points were a distance zero apart.]

(3.3) DEFINITION. The packing number $D(\epsilon, T_0)$ for a subset T_0 of a metric space is defined as the largest m for which there exist points $t_1, \dots, t_m$ in T_0 with $d(t_i, t_j) > \epsilon$ for $i \neq j$. The covering number $N(\epsilon, T_0)$ is defined as the smallest number of closed balls with radius ϵ whose union covers T_0 .

The two concepts are closely related, because

$$N(\epsilon, T_0) \leq D(\epsilon, T_0) \leq N(\epsilon/2, T_0).$$

Both provide approximating points $t_1, \dots, t_m$ for which $\min_i d(t, t_i) \leq \epsilon$ for every t in T_0 . Sometimes the $\{t_i\}$ provided by D are slightly more convenient to work with, because they lie in T_0 ; the centers of the balls provided by N need not lie in T_0 . The

definition of D depends only upon the behavior of the metric d on the set T_0 ; the value of N can depend upon the particular T into which T_0 is embedded. If $T = T_0$ the ambiguity disappears. However, it is largely a matter of taste, or habit, whether one works with covering numbers or packing numbers. Notice that finiteness of all the packing or covering numbers is equivalent to total boundedness of T_0 .

For the general maximal inequality let us suppose that some point t_0 has been singled out from T . Also, let us assume that the process $Z(t) = Z(\omega, t)$ has continuous sample paths, in the sense that $Z(\omega, \cdot)$ defines a continuous function on T for each ω . For the intended application, this causes no loss of generality: clearly $\sigma \cdot \mathbf{f}$ is a continuous function of $\mathbf{f}$ for each fixed σ . [Without the continuity assumption the statement of the next lemma would have to be modified to assert existence of a version of the process Z having continuous sample paths and satisfying the stated inequality.]

(3.4) LEMMA. *If the process Z has continuous sample paths and its increments satisfy the inequality*

$$\|Z(s) - Z(t)\|_{\Psi} \leq d(s, t) \quad \text{for all } s, t \text{ in } T,$$

and if $\delta = \sup_t d(t, t_0)$, then

$$\left\| \sup_T |Z(t)| \right\|_{\Psi} \leq \|Z(t_0)\|_{\Psi} + \sum_{i=0}^{\infty} \frac{\delta}{2^i} \sqrt{2 + \log D(\delta/2^{i+1}, T)}.$$

PROOF. The inequality holds trivially if the right-hand side is infinite. So let us assume that δ and all the packing numbers are finite.

Abbreviate $\delta/2^i$ to δ_i . Construct a succession of approximations to the supremum based on a sequence of finite subsets $\{t_0\} = T_0 \subseteq T_1 \subseteq \dots$ with the property that

$$\min_{t^* \in T_i} d(t, t^*) \leq \delta_i \quad \text{for every } t \text{ in } T.$$

Such sets can be obtained inductively by choosing T_i as a maximal superset of T_{i-1} with all points of T_i greater than δ_i apart. [Notice that the definition of δ ensures that $\{t_0\}$ has the desired property for δ_0 .] The definition of packing number gives us a bound on the cardinality of T_i , namely, $\#T_i \leq D(\delta_i, T)$. Let us write m_i for this bound.

Fix, for the moment, a non-negative integer k . Relate the maximum of $|Z(t)|$ over T_{k+1} to the maximum over T_k . For each t in T_{k+1} let t^* denote a point in T_k such that $d(t, t^*) \leq \delta_k$. By the triangle inequality,

$$\max_{t \in T_{k+1}} |Z(t)| \leq \max_{t \in T_{k+1}} |Z(t^*)| + \max_{t \in T_{k+1}} |Z(t) - Z(t^*)|.$$

On the right-hand side, the first term takes a maximum over a subset of T_k . The second term takes a maximum over at most m_{k+1} increments, each of which has Ψ norm at most δ_k . Taking Ψ norms of both sides of the inequality, then applying Lemma 3.2 to the contribution from the increments, we get

$$\left\| \max_{T_{k+1}} |Z(t)| \right\|_{\Psi} \leq \left\| \max_{T_k} |Z(t)| \right\|_{\Psi} + \delta_k \sqrt{2 + \log m_{k+1}}.$$

Repeated application of this recursive bound increases the right-hand side to the contribution from T_0 , which reduces to $\|Z(t_0)\|_\Psi$, plus a sum of terms contributed by the increments.

As k tends to infinity, the set T_{k+1} expands up to a countable dense subset T_∞ of T . A monotone convergence argument shows that

$$\left\| \max_{T_{k+1}} |Z(t)| \right\|_\Psi \nearrow \left\| \sup_{T_\infty} |Z(t)| \right\|_\Psi.$$

Continuity of the sample paths of Z lets us replace T_∞ by T , since

$$\sup_{T_\infty} |Z(\omega, t)| = \sup_T |Z(\omega, t)| \quad \text{for every } \omega.$$

This gives the asserted inequality. $\square$

Now we have only to specialize the argument to the process $\sigma \cdot \mathbf{f}$ indexed by a subset $\mathcal{F}$ of $\mathbb{R}^n$. The packing numbers for $\mathcal{F}$ should be calculated using the usual Euclidean distance on $\mathbb{R}^n$. By Lemma 3.1 the increments of the process satisfy

$$\|\sigma \cdot (\mathbf{f} - \mathbf{g})\|_\Psi \leq 2|\mathbf{f} - \mathbf{g}|,$$

which differs from the inequality required by Lemma 3.4 only through the presence of the factor 2. We could eliminate the factor by working with the process $1/2\sigma \cdot \mathbf{f}$.

To get a neater bound, let us take the origin of $\mathbb{R}^n$ as the point corresponding to t_0 . At worst, this increases the packing numbers for $\mathcal{F}$ by one. We can tidy up the integrand by noting that $D(x, \mathcal{F}) \geq 2$ for $x < \delta$, and then using the inequality

$$\sqrt{2 + \log(1 + D)} / \sqrt{\log D} < 2.2 \quad \text{for } D \geq 2.$$

It has also become traditional to replace the infinite series in Lemma 3.4 by the corresponding integral, a simplification made possible by the geometric rate of decrease in the $\{\delta_i\}$:

$$\delta_i \sqrt{\log D(\delta_{i+1}, \mathcal{F})} \leq 4 \int_{\{\delta_{i+2} < x \leq \delta_{i+1}\}} \sqrt{\log D(x, \mathcal{F})} dx.$$

With these cosmetic changes the final maximal inequality has a nice form.

(3.5) THEOREM. *For every subset $\mathcal{F}$ of $\mathbb{R}^n$,*

$$\left\| \sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}| \right\|_\Psi \leq 9 \int_0^\delta \sqrt{\log D(x, \mathcal{F})} dx \quad \text{where } \delta = \sup_{\mathcal{F}} |\mathbf{f}|. \quad \square$$

The theorem has several interesting consequences and reformulations. For example, suppose the integral on the right-hand side is finite. Then there exist positive constants κ_1 and κ_2 such that

$$\mathbb{P}_\sigma \left\{ \sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}| \geq \epsilon \right\} \leq \kappa_1 \exp(-\kappa_2 \epsilon^2) \quad \text{for all } \epsilon > 0.$$

It will also give bounds for less stringent norms than the Ψ norm. For example, for each p with $\infty > p \geq 1$ there exists a constant C_p such that $|x|^p \leq \Psi(C_p x)$ for all x .

This implies that $\|Z\|_p \leq C_p \|Z\|_\Psi$ for every random variable Z , and, in particular,

$$(3.6) \quad \left\| \sup_{\mathcal{F}} |\boldsymbol{\sigma} \cdot \mathbf{f}| \right\|_p \leq 9C_p \int_0^\delta \sqrt{\log D(x, \mathcal{F})} dx, \quad \text{where } \delta = \sup_{\mathcal{F}} |\mathbf{f}|.$$

Such bounds will prove convenient in later sections.

REMARKS. The literature contains many different maximal inequalities derived by chaining arguments. The method presented in this section could be refined to produce more general inequalities, but Theorem 3.5 will suffice for the limited purposes of these notes.

I learnt the method for Lemma 3.1 from Gilles Pisier. The whole section is based on ideas exposted by Pisier (1983), who proved an inequality equivalent to

$$\mathbb{P} \sup_{s,t} |Z(s) - Z(t)| \leq K \int_0^\delta \Phi^{-1}(D(x, T)) dx$$

for general convex, increasing Φ with $\Phi(0) = 0$. This result is weaker than the corresponding inequality with the left-hand side increased to

$$\left\| \sup_{s,t} |Z(s) - Z(t)| \right\|_\Phi.$$

For the special case where $\Phi(x) = 1/5 \exp(x^2)$ the improvement is made possible by the substitution of my Lemma 3.2 for Pisier's Lemma 1.6 in the chaining argument. Ledoux and Talagrand (1990, Chapter 11) have shown how the stronger form of the inequality can also be deduced directly from a slight modification of Pisier's inequality.

Both Gaenssler and Schlumprecht (1988) and Pollard (1989) have established analogues of Theorem 3.5 for $\|\cdot\|_p$ norms instead of the $\|\cdot\|_\Psi$.

SECTION 4

Packing and Covering in Euclidean Spaces

The maximal inequality from Theorem 3.6 will be useful only if we have suitable bounds for the packing numbers of the set $\mathcal{F}$. This section presents a method for finding such bounds, based on a geometric property that transforms calculation of packing numbers into a combinatorial exercise.

The combinatorial approach generalizes the concept of a Vapnik-Červonenkis class of sets. It identifies certain subsets of $\mathbb{R}^n$ that behave somewhat like compact sets of lower dimension; the bounds on the packing numbers grow geometrically, at a rate determined by the lower dimension. For comparison's sake, let us first establish the bound for genuinely lower dimensional sets.

(4.1) LEMMA. *Let $\mathcal{F}$ be a subset of a V dimensional affine subspace of $\mathbb{R}^n$. If $\mathcal{F}$ has finite diameter R , then*

$$D(\epsilon, \mathcal{F}) \leq \left(\frac{3R}{\epsilon} \right)^V \quad \text{for } 0 < \epsilon \leq R.$$

PROOF. Because Euclidean distances are invariant under rotation, we may identify $\mathcal{F}$ with a subset of $\mathbb{R}^V$ for the purposes of calculating the packing number $D(\epsilon, \mathcal{F})$. Let $\mathbf{f}_1, \dots, \mathbf{f}_m$ be points in $\mathcal{F}$ with $|\mathbf{f}_i - \mathbf{f}_j| > \epsilon$ for $i \neq j$. Let B_i be the (V -dimensional) ball of radius $\epsilon/2$ and center $\mathbf{f}_i$. These m balls are disjoint; they occupy a total volume of $m(\epsilon/2)^V \Gamma$, where Γ denotes the volume of a unit ball in $\mathbb{R}^V$. Each $\mathbf{f}_i$ lies within a distance R of $\mathbf{f}_1$; each B_i lies inside a ball of radius $3/2 R$ and center $\mathbf{f}_1$, a ball of volume $(3/2 R)^V \Gamma$. It follows that $m \leq (3R/\epsilon)^V$. $\square$

A set of dimension V looks thin in $\mathbb{R}^n$. Even if projected down onto a subspace of $\mathbb{R}^n$ it will still look thin, if the subspace has dimension greater than V . One way to capture this idea, and thereby create a more general notion of a set being thin, is to think of how much of the space around any particular point can be occupied by

the set. The formal concept involves the collection of 2^k orthants about each point $\mathbf{t}$ in $\mathbb{R}^k$ defined by means of all possible combinations of coordinatewise inequalities.

(4.2) DEFINITION. For each $\mathbf{t}$ in $\mathbb{R}^k$ and each subset J of $\{1, \dots, k\}$, define the J^{th} orthant about $\mathbf{t}$ to consist of all those $\mathbf{x}$ in $\mathbb{R}^k$ for which

$$\begin{aligned} x_i &> t_i && \text{if } i \in J, \\ x_i &< t_i && \text{if } i \in J^c. \end{aligned}$$

A subset of $\mathbb{R}^k$ will be said to *occupy* the J^{th} orthant of $\mathbf{t}$ if it contains at least one point in that orthant. A subset will be said to *surround* $\mathbf{t}$ if it occupies all 2^k of the orthants defined by $\mathbf{t}$.

There is a surprising connection between the packing numbers of a set in $\mathbb{R}^n$ and the maximum number of orthants its lower dimensional projections can occupy. The projections that we use will differ slightly from the usual notion. For each k -tuple $I = (i(1), \dots, i(k))$ of integers from the range $1, \dots, n$, call $(x_{i(1)}, \dots, x_{i(k)})$ the I -projection of the point $(x_1, \dots, x_n)$ in $\mathbb{R}^n$, even if the integers $i(1), \dots, i(k)$ are not all distinct. Call such a map into $\mathbb{R}^k$ a k -dimensional coordinate projection. If all the integers are distinct, call it a *proper coordinate projection*.

(4.3) DEFINITION. Say that a subset $\mathcal{F}$ of $\mathbb{R}^n$ has a *pseudodimension* of at most V if, for every point $\mathbf{t}$ in $\mathbb{R}^{V+1}$, no proper coordinate projection of $\mathcal{F}$ can surround $\mathbf{t}$.

The concept of pseudodimension bears careful examination. It requires a property for all possible choices of $I = (i(1), \dots, i(V+1))$ from the range $1, \dots, n$. For each such choice and for each $\mathbf{t}$ in $\mathbb{R}^{V+1}$, one must extract a J from I such that no $\mathbf{f}$ in $\mathcal{F}$ can satisfy the inequalities

$$\begin{aligned} f_i &> t_i && \text{for } i \in J, \\ f_i &< t_i && \text{for } i \in I \setminus J. \end{aligned}$$

Clearly any duplication amongst the elements of I will make this task a triviality. Only for distinct integers $i(1), \dots, i(V+1)$ must one expend energy to establish impossibility. That is why only proper projections need be considered.

If a set $\mathcal{F}$ actually sits within an affine space of dimension V then it has pseudodimension at most V . To see this, notice that a $(V+1)$ -dimensional projection of such an $\mathcal{F}$ must be a subset of an affine subspace $\mathcal{A}$ with dimension strictly less than $V+1$. There exists a nontrivial vector $\boldsymbol{\beta}$ in $\mathbb{R}^{V+1}$ and a constant γ such that $\boldsymbol{\beta} \cdot \boldsymbol{\alpha} = \gamma$ for every $\boldsymbol{\alpha}$ in $\mathcal{A}$. We may assume that $\beta_i > 0$ for at least one i . If $\mathbf{t}$ has $\boldsymbol{\beta} \cdot \mathbf{t} \leq \gamma$ it is impossible to find an $\boldsymbol{\alpha}$ in $\mathcal{A}$ such that

$$\begin{aligned} \alpha_i &< t_i && \text{when } \beta_i > 0, \\ \alpha_i &\geq t_i && \text{when } \beta_i \leq 0, \end{aligned}$$

for these inequalities would lead to the contradiction $\gamma = \sum_i \beta_i \alpha_i < \sum_i \beta_i t_i \leq \gamma$. If $\boldsymbol{\beta} \cdot \mathbf{t} > \gamma$ we would interchange the roles of “ $\beta_i > 0$ ” and “ $\beta_i \leq 0$ ” to reach a similar

contradiction. For the pseudodimension calculation we need the contradiction only for $\alpha_i > t_i$, but to establish the next result we need it for $\alpha_i \geq t_i$ as well.

(4.4) LEMMA. *Suppose the coordinates of the points in $\mathcal{F}$ can take only two values, c_0 and c_1 . Suppose also that there is a V -dimensional vector subspace Λ of $\mathbb{R}^n$ with the property: for each $f \in \mathcal{F}$ there is a $\lambda \in \Lambda$ such that $f_i = c_1$ if and only if $\lambda_i \geq 0$. Then $\mathcal{F}$ has pseudodimension at most V .*

PROOF. We may assume that $c_0 = 0$ and $c_1 = 1$. Suppose that some proper I -projection of $\mathcal{F}$ surrounds a point $\mathbf{t}$ in $\mathbb{R}^{V+1}$. Each coordinate t_i must lie strictly between 0 and 1. The inequalities required for the projection of $\mathbf{f}$ to occupy the orthant corresponding to a subset J of I are

$$\begin{aligned} f_i &= 1 & \text{for } i \in J, \\ f_i &= 0 & \text{for } i \in I \setminus J. \end{aligned}$$

That is,

$$\begin{aligned} \lambda_i &\geq 0 & \text{for } i \in J, \\ \lambda_i &< 0 & \text{for } i \in I \setminus J. \end{aligned}$$

As shown above, there exists a J such that this system of inequalities cannot be satisfied. $\square$

The connection between pseudodimension and packing numbers is most easily expressed if we calculate the packing numbers not for the usual Euclidean, or ℓ_2 , distance on $\mathbb{R}^n$, but rather for the ℓ_1 distance that corresponds to the norm

$$|\mathbf{x}|_1 = \sum_{i \leq n} |x_i|.$$

To distinguish between the two metrics on $\mathbb{R}^n$ let us add subscripts to our notation, writing $D_1(\epsilon, \mathcal{F})$ for the ℓ_1 packing number of the set $\mathcal{F}$, and so on. [Notice that the ℓ_1 norm is not invariant under rotation. The invariance argument used in the proof of Lemma 4.1 would be invalid for ℓ_1 packing numbers.]

A set in $\mathbb{R}^n$ of the form $\prod_i [\alpha_i, \beta_i]$ is called a *box*. It has ℓ_1 diameter $\sum_i (\beta_i - \alpha_i)$. The smallest integer greater than a real number x is denoted by $\lceil x \rceil$.

(4.5) LEMMA. *Let $\mathcal{F}$ lie within a box of ℓ_1 diameter one in $\mathbb{R}^n$. If $\mathcal{F}$ contains m points, each pair separated by an ℓ_1 distance of at least ϵ , then: for $k = \lceil 2\epsilon^{-1} \log m \rceil$, there exists a point $\mathbf{t}$ in $\mathbb{R}^k$ and a k -dimensional coordinate projection of $\mathcal{F}$ that occupies at least m orthants of $\mathbf{t}$.*

PROOF. We may assume that the box has the form $\prod_i [0, p_i]$, where the p_i are nonnegative numbers summing to one. Partition $[0, 1]$ into subintervals $I_1, \dots, I_n$ of lengths $p_1, \dots, p_n$. Generate $i(1), \dots, i(k)$ and $\mathbf{t} = (t_1, \dots, t_k)$ randomly, from a set of independent Uniform $[0, 1]$ random variables $U_1, \dots, U_k$, in the following way. If U_α lands in the subinterval I_j , put $i(\alpha)$ equal to j and t_α equal to the

distance of U_α from the left endpoint of I_j . [That is, the method chooses edge i with probability p_i , then chooses t_i uniformly from the $[0, p_i]$ interval.]

Let $\mathcal{F}_0$ be the subset of $\mathcal{F}$ consisting of the m points with the stated separation property. To each $\mathbf{f}$ in $\mathcal{F}_0$ there corresponds a set of n points in $[0, 1]$: the j^{th} lies in I_j , at a distance f_j from the left endpoint of that interval. The $2n$ points defined in this way by each pair $\mathbf{f}, \mathbf{g}$ from $\mathcal{F}_0$ form n subintervals of $[0, 1]$, one in each I_j . The total length of the subintervals equals $|\mathbf{f} - \mathbf{g}|_1$, which is greater than ϵ , by assumption. If U_α lands within the interior of the subintervals, the coordinates $f_{i(\alpha)}$ and $g_{i(\alpha)}$ will be on opposite sides of t_α ; the projections of $\mathbf{f}$ and $\mathbf{g}$ will then lie in different orthants of $\mathbf{t}$. Each U_α has probability at most $1 - \epsilon$ of failing to separate $\mathbf{f}$ and $\mathbf{g}$ in this way. Therefore the projections have probability at most $(1 - \epsilon)^k$ of lying in the same orthant of $\mathbf{t}$.

Amongst the $\binom{m}{2}$ possible pairs from $\mathcal{F}_0$, the probability that at least one pair of projections will occupy the same orthant of $\mathbf{t}$ is less than

$$\binom{m}{2} (1 - \epsilon)^k < \frac{1}{2} \exp(2 \log m - k\epsilon).$$

The value of k was chosen to make this probability strictly less than one. With positive probability the procedure will generate $i(1), \dots, i(k)$ and $\mathbf{t}$ with the desired properties. $\square$

Notice that the value of k does not depend on n , the dimension of the space $\mathbb{R}^n$ in which the set $\mathcal{F}$ is embedded.

The next result relates the occupation of a large number of orthants in $\mathbb{R}^k$ to the property that some lower-dimensional projection of the set completely surrounds some point. This will lead to a checkable criterion for an $\mathcal{F}$ in $\mathbb{R}^n$ to have packing numbers that increase at the same sort of geometric rate as for the low-dimensional set in Lemma 4.1. The result is a thinly disguised form of the celebrated Vapnik-Červonenkis lemma.

(4.6) LEMMA. *A coordinate projection into $\mathbb{R}^k$ of a set with pseudodimension at most V can occupy at most*

$$\binom{k}{0} + \binom{k}{1} + \dots + \binom{k}{V}$$

orthants about any point of $\mathbb{R}^k$.

PROOF. Let $\mathcal{H}$ be a set with pseudodimension at most V . Its projection into $\mathbb{R}^k$ also has pseudodimension at most V . So without loss of generality we may assume that $\mathcal{H} \subseteq \mathbb{R}^k$. Let $\mathcal{S}$ denote the set of all k -tuples $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_k)$ with $\sigma_i = \pm 1$ for each i . Identify the 2^k orthants of $\mathbf{t}$ with the 2^k vectors in $\mathcal{S}$. The orthants of $\mathbf{t}$ that are occupied by $\mathcal{H}$ correspond to a subset $\mathcal{A}$ of $\mathcal{S}$. Suppose $\#\mathcal{A}$ is strictly greater than the asserted bound, and then argue for a contradiction.

The vectors in $\mathcal{S}$ also index the proper coordinate projections on $\mathbb{R}^k$. Let us denote by π_σ the projection that discards all those coordinates for which $\sigma_i = -1$. The orthants of $\pi_\sigma \mathbf{t}$ correspond to vectors $\boldsymbol{\eta}$ in $\mathcal{S}$ with $\boldsymbol{\eta} \leq \boldsymbol{\sigma}$: we merely ignore those

coordinates η_i for which $\sigma_i = -1$, then identify the orthants by means of the signs of the remaining η_i . For the projection $\pi_\sigma \mathcal{H}$ to occupy the orthant corresponding to $\boldsymbol{\eta}$, there must exist an $\boldsymbol{\alpha}$ in $\mathcal{A}$ such that $\alpha_i = \eta_i$ whenever $\sigma_i = +1$; that is, $\boldsymbol{\alpha} \wedge \boldsymbol{\sigma} = \boldsymbol{\eta}$.

Let $\mathbb{S}_V$ denote the set of all vectors $\boldsymbol{\sigma}$ in $\mathbb{S}$ with $\sigma_i = +1$ for at least $V + 1$ coordinates. The assumption of pseudodimension at most V means that $\pi_\sigma \mathcal{H}$ does not surround $\pi_\sigma \mathbf{t}$, for every $\boldsymbol{\sigma}$ in $\mathbb{S}_V$. Thus for each $\boldsymbol{\sigma}$ in $\mathbb{S}_V$ there exists an $\boldsymbol{\eta}(\boldsymbol{\sigma}) \leq \boldsymbol{\sigma}$ such that $\boldsymbol{\alpha} \wedge \boldsymbol{\sigma} \neq \boldsymbol{\eta}(\boldsymbol{\sigma})$ for every $\boldsymbol{\alpha}$ in $\mathcal{A}$. For definiteness define $\boldsymbol{\eta}(\boldsymbol{\sigma}) = \boldsymbol{\sigma}$ for $\boldsymbol{\sigma} \notin \mathbb{S}_V$.

Invoke the Basic Combinatorial Lemma from Section 1 to obtain a one-to-one map θ from $\mathbb{S}$ onto itself such that $\theta(\boldsymbol{\sigma}) \wedge \boldsymbol{\sigma} = \boldsymbol{\eta}(\boldsymbol{\sigma})$ for every $\boldsymbol{\sigma}$. The assumption about the size of $\mathcal{A}$ ensures that

$$\#\{\theta^{-1}(\boldsymbol{\alpha}) : \boldsymbol{\alpha} \in \mathcal{A}\} + \#\mathbb{S}_V > 2^k,$$

which implies that there exists an $\boldsymbol{\alpha}$ in $\mathcal{A}$ for which $\theta^{-1}(\boldsymbol{\alpha}) \in \mathbb{S}_V$. But then, for that $\boldsymbol{\alpha}$, we have

$$\boldsymbol{\alpha} \wedge \theta^{-1}(\boldsymbol{\alpha}) \neq \boldsymbol{\eta}(\theta^{-1}(\boldsymbol{\alpha})) = \theta(\theta^{-1}(\boldsymbol{\alpha})) \wedge \theta^{-1}(\boldsymbol{\alpha}),$$

a contradiction that establishes the assertion of the lemma. $\square$

The V in the statement of the last lemma plays almost the same role as the dimension V in Lemma 4.1, which gave the $O(\epsilon^{-V})$ bound on packing numbers. By combining the assertions of the last two lemmas we obtain the corresponding bound in terms of the pseudodimension.

(4.7) THEOREM. *Let $\mathcal{F}$ have pseudodimension at most V and lie within a box of ℓ_1 diameter one in $\mathbb{R}^n$. Then there exist constants A and W , depending only on V , such that*

$$D_1(\epsilon, \mathcal{F}) \leq A (1/\epsilon)^W \quad \text{for } 0 < \epsilon \leq 1.$$

PROOF. Fix $0 < \epsilon \leq 1$. Let $m = D_1(\epsilon, \mathcal{F})$. Choose $k = \lceil 2\epsilon^{-1} \log m \rceil$ as in Lemma 4.5. From Lemma 4.6,

$$\binom{k}{0} + \cdots + \binom{k}{V} \geq m.$$

The left-hand side of this inequality is a polynomial of degree V in k ; it is smaller than $(1 + V)k^V$. [There is not much to be gained at this stage by a more precise upper bound.] Thus

$$(1 + V) \left(\frac{1 + 2 \log m}{\epsilon} \right)^V \geq m,$$

whence

$$\frac{(1 + V)}{\epsilon^V} \geq \frac{m}{(1 + 2 \log m)^V}.$$

For some positive constant C depending on V , the right-hand side is greater than $C\sqrt{m}$, for all positive integers m . The asserted inequality holds if we take $A = (1 + V)^2/C^2$ and $W = 2V$. $\square$

For the sake of comparison with Lemma 4.1, let us see what sort of bound is given by Theorem 4.7 when $\mathcal{F}$ is contained within a V -dimensional affine subspace of $\mathbb{R}^n$. If $\mathcal{F}$ also lies within an ℓ_1 box of diameter one, the argument from the proof of Theorem 4.7 gives packing numbers that grow as $O(\epsilon^{-W})$, for $W = 2V$. We could reduce W to any constant slightly larger than V . [Use $Cm^{1-\delta}$, for some tiny positive δ , instead of $C\sqrt{m}$, in the proof.] This falls just slightly short of the $O(\epsilon^{-V})$ bound from Lemma 4.1.

Theorem 4.7 has a slightly more general version that exploits an invariance property of orthants. For each vector $\alpha = (\alpha_1, \dots, \alpha_n)$ of nonnegative constants, and each $\mathbf{f}$ in $\mathbb{R}^n$, define the pointwise product $\alpha \odot \mathbf{f}$ to be the vector in $\mathbb{R}^n$ with i^{th} coordinate $\alpha_i f_i$. Write $\alpha \odot \mathcal{F}$ to denote the set of all vectors $\alpha \odot \mathbf{f}$ with $\mathbf{f}$ in $\mathcal{F}$. At least when $\alpha_i > 0$ for every i , a trivial, but significant, property of orthants is: $\mathcal{F}$ occupies orthant J of $\mathbf{t}$ if and only if $\alpha \odot \mathcal{F}$ occupies orthant J of $\alpha \odot \mathbf{t}$. Similarly, if some coordinate projection of $\mathcal{F}$ cannot surround a point $\mathbf{t}$ then the corresponding coordinate projection of $\alpha \odot \mathcal{F}$ cannot surround $\alpha \odot \mathbf{t}$. The key requirement of the theorem is unaffected by such coordinate rescalings. We can rescale any bounded set $\mathcal{F}$ with an *envelope* $\mathbf{F}$ —that is, a vector such that $|f_i| \leq F_i$ for each $\mathbf{f} \in \mathcal{F}$ and each i —to lie within a box of ℓ_1 diameter one, and then invoke the theorem.

(4.8) THEOREM. *Let $\mathcal{F}$ be a bounded subset of $\mathbb{R}^n$ with envelope $\mathbf{F}$ and pseudo-dimension at most V . Then there exist constants A and W , depending only on V , such that*

$$D_1(\epsilon|\alpha \odot \mathbf{F}|_1, \alpha \odot \mathcal{F}) \leq A (1/\epsilon)^W \quad \text{for } 0 < \epsilon \leq 1,$$

for every rescaling vector α of non-negative constants.

PROOF. We may assume $\alpha_i > 0$ for every i . (The cases where some α_i are zero correspond to an initial projection of $\mathcal{F}$ into a lower dimensional coordinate subspace.) Apply Theorem 4.7 to the rescaled set $\mathcal{F}^*$ consisting of vectors $\mathbf{f}^*$ with coordinates

$$f_i^* = \frac{\alpha_i f_i}{2 \sum_j \alpha_j F_j}.$$

Then observe that, for vectors in $\mathcal{F}^*$,

$$|\mathbf{f}^* - \mathbf{g}^*|_1 > \epsilon/2 \quad \text{if and only if} \quad |\alpha \odot \mathbf{f} - \alpha \odot \mathbf{g}|_1 > \epsilon|\alpha \odot \mathbf{F}|_1.$$

Absorb the extra factor of 2^W into the constant A . $\square$

Sets with an $O(\epsilon^{-W})$ bound on packing numbers arise in many problems, as will become apparent in the sections on applications. The main role of the pseudo-dimension of a set $\mathcal{F}$ will be to provide such a geometric rate of growth for packing numbers of $\mathcal{F}$. It also applies to *any subclass of $\mathcal{F}$ under its natural envelope*. For subclasses with small natural envelopes, this method sometimes leads to bounds unattainable by other methods.

The added generality of an inequality that holds uniformly over all rescaling vectors allows us to move back and forth between ℓ_1 and ℓ_2 packing numbers. The bounds from Theorem 4.8 will translate into bounds on ℓ_2 packing numbers suitable for the chaining arguments in the Section 3.

(4.9) LEMMA. *For each bounded $\mathcal{F}$ with envelope $\mathbf{F}$, and each $\epsilon > 0$,*

$$D_2(\epsilon, \mathcal{F}) \leq D_1(\tfrac{1}{2}\epsilon^2, \mathbf{F} \odot \mathcal{F}) \leq D_2(\tfrac{1}{2}\epsilon^2, |\mathbf{F}|_2, \mathcal{F}).$$

PROOF. For each pair of vectors $\mathbf{f}, \mathbf{g}$ in $\mathcal{F}$,

$$|\mathbf{f} - \mathbf{g}|_2^2 \leq 2|\mathbf{F} \odot \mathbf{f} - \mathbf{F} \odot \mathbf{g}|_1 \leq 2|\mathbf{F}|_2 |\mathbf{f} - \mathbf{g}|_2.$$

The first inequality follows from the bound $(f_i - g_i)^2 \leq 2F_i|f_i - g_i|$; the second follows from the Cauchy-Schwarz inequality. $\square$

(4.10) COROLLARY. *If $\mathcal{F}$ is a bounded subset of $\mathbb{R}^n$ with envelope $\mathbf{F}$ and pseudo-dimension at most V , then there exist constants A_2 and W_2 , depending only on V , such that*

$$D_2(\epsilon|\boldsymbol{\alpha} \odot \mathbf{F}|_2, \boldsymbol{\alpha} \odot \mathcal{F}) \leq A_2 (1/\epsilon)^{W_2} \quad \text{for } 0 < \epsilon \leq 1$$

and every rescaling vector $\boldsymbol{\alpha}$ of non-negative constants.

PROOF. The set $\boldsymbol{\alpha} \odot \mathcal{F}$ has envelope $\boldsymbol{\beta} = \boldsymbol{\alpha} \odot \mathbf{F}$. Because $\boldsymbol{\beta} \odot \boldsymbol{\alpha} \odot \mathcal{F}$ has envelope $\boldsymbol{\beta} \odot \boldsymbol{\beta}$ and $|\boldsymbol{\beta}|_2^2 = |\boldsymbol{\beta} \odot \boldsymbol{\beta}|_1$, the ℓ_2 packing number is bounded by

$$D_1(\tfrac{1}{2}\epsilon^2|\boldsymbol{\beta} \odot \boldsymbol{\beta}|_1, \boldsymbol{\beta} \odot \boldsymbol{\alpha} \odot \mathcal{F}) \leq A(\tfrac{1}{2}\epsilon^2)^{-W},$$

with A and W from Theorem 4.8. $\square$

The presence of an arbitrary rescaling vector in the bound also gives us added flexibility when we deal with sets that are constructed from simpler pieces, as will be explained in the next section.

REMARKS. My definition of pseudodimension abstracts the concept of a Vapnik-Červonenkis subgraph class of functions, in the sense of Dudley (1987). Most of the results in the section are reformulations or straightforward extensions of known theory for Vapnik-Červonenkis classes, as exposited in Chapter II of Pollard (1984), for example. See that book for a listing of who first did what when.

The nuisance of improper coordinate projections was made necessary by my desire to break the standard argument into several steps. The arguments could be rewritten using only proper projections, by recombining Lemma 4.5 and Theorem 4.7. The proof of Lemma 4.6 is a novel rearrangement of old ideas: see the comments at the end of Section 1 regarding the Basic Combinatorial Lemma.

SECTION 5

Stability

Oftentimes an interesting process can be put together from simpler processes, to which the combinatorial methods of Section 4 apply directly. The question then becomes one of stability: Does the process inherit the nice properties from its component pieces? This section provides some answers for the case of processes $\sigma \cdot \mathbf{f}$ indexed by subsets of Euclidean space.

Throughout the section $\mathcal{F}$ and $\mathcal{G}$ will be fixed subsets of $\mathbb{R}^n$, with envelopes $\mathbf{F}$ and $\mathbf{G}$ and $\sigma = (\sigma_1, \dots, \sigma_n)$ will be a vector of independent random variables, each taking the values ± 1 with probability $1/2$. In particular, σ will be regarded as the generic point in the set $\mathcal{S}$ of all n -tuples of ± 1 's, under its uniform distribution $\mathbb{P}_\sigma$. The problem is to determine which properties of $\mathcal{F}$ and $\mathcal{G}$ are inherited by classes such as

$$\begin{aligned}\mathcal{F} \oplus \mathcal{G} &= \{\mathbf{f} + \mathbf{g} : \mathbf{f} \in \mathcal{F}, \mathbf{g} \in \mathcal{G}\}, \\ \mathcal{F} \vee \mathcal{G} &= \{\mathbf{f} \vee \mathbf{g} : \mathbf{f} \in \mathcal{F}, \mathbf{g} \in \mathcal{G}\}, \\ \mathcal{F} \wedge \mathcal{G} &= \{\mathbf{f} \wedge \mathbf{g} : \mathbf{f} \in \mathcal{F}, \mathbf{g} \in \mathcal{G}\}, \\ \mathcal{F} \odot \mathcal{G} &= \{\mathbf{f} \odot \mathbf{g} : \mathbf{f} \in \mathcal{F}, \mathbf{g} \in \mathcal{G}\}.\end{aligned}$$

The reader might want to skip the material in the subsection headed “General Maximal Inequalities”. It is included in this section merely to illustrate one of the more recent developments in the subject; it is based on the paper by Ledoux and Talagrand (1989). For most applications to asymptotic problems, the simpler results contained in the first two subsections seem to suffice.

Pseudodimension. This property is stable only for the formation of unions, pointwise maxima, and pointwise minima.

Suppose that both $\mathcal{F}$ and $\mathcal{G}$ have pseudodimension at most V . Then, for every $\mathbf{t}$ in $\mathbb{R}^k$ and every k less than n , Lemma 4.6 asserts that the projection of $\mathcal{F}$ can occupy at most

$$m = \binom{k}{0} + \dots + \binom{k}{V}$$

of the orthants around $\mathbf{t}$, and similarly for $\mathcal{G}$. For any two vectors α and β in $\mathbb{R}^k$, the orthants of $\mathbf{t}$ occupied by $\alpha \vee \beta$ and $\alpha \wedge \beta$ are uniquely determined by the orthants occupied by α and β . (The same cannot be said for $\alpha + \beta$ or $\alpha \odot \beta$.) Thus the projections of $\mathcal{F} \vee \mathcal{G}$ and $\mathcal{F} \wedge \mathcal{G}$ each occupy at most m^2 different orthants. It is even easier to show that the union $\mathcal{F} \cup \mathcal{G}$ occupies at most $2m$ orthants. If k could be chosen so that $m^2 < 2^k$, this would imply that none of the projections surrounds $\mathbf{t}$. So, we need to find a k such that

$$\left[\binom{k}{0} + \cdots + \binom{k}{V} \right]^2 < 2^k.$$

On the left-hand side we have a polynomial of degree $2V$, which increases much more slowly with k than the 2^k on the right-hand side. For k large enough the inequality will be satisfied. Just knowing that such a k exists is good enough for most applications, but, for the sake of having an explicit bound, let us determine how k depends on V .

Introduce a random variable X with a $\text{Bin}(k, 1/2)$ distribution. The desired inequality is equivalent to

$$\left[\mathbb{P}\{X \geq k - V\} \right]^2 < 2^{-k}.$$

Bound the left-hand side by

$$\left[9^{-(k-V)} \mathbb{P} 9^X \right]^2 = 81^{-(k-V)} 25^k,$$

then choose $k = 10V$ to make the bound less than 2^{-k} for every V . [It is possible to replace 10 by a smaller constant, but this has no advantage for our purposes.]

(5.1) **LEMMA.** *If both $\mathcal{F}$ and $\mathcal{G}$ have pseudodimension at most V , then all of $\mathcal{F} \cup \mathcal{G}$ and $\mathcal{F} \vee \mathcal{G}$ and $\mathcal{F} \wedge \mathcal{G}$ have pseudodimension less than $10V$. $\square$*

Unfortunately neither sums nor products share this form of stability.

Packing Numbers. Stability properties for packing or covering numbers follow easily from the triangle inequality: we construct approximating subclasses $\{\mathbf{f}_i\}$ for $\mathcal{F}$ and $\{\mathbf{g}_j\}$ for $\mathcal{G}$, and then argue from inequalities such as

$$|\mathbf{f} \vee \mathbf{g} - \mathbf{f}_i \vee \mathbf{g}_j|_2 \leq |\mathbf{f} - \mathbf{f}_i|_2 + |\mathbf{g} - \mathbf{g}_j|_2.$$

In this way we get covering number bounds

$$N_2(\epsilon + \delta, \mathcal{F} \square \mathcal{G}) \leq N_2(\epsilon, \mathcal{F}) N_2(\delta, \mathcal{G}),$$

where $\square$ stands for either $+$ or $\vee$ or $\wedge$. The corresponding bounds for packing numbers,

$$D_2(2\epsilon + 2\delta, \mathcal{F} \square \mathcal{G}) \leq D_2(\epsilon, \mathcal{F}) D_2(\delta, \mathcal{G}),$$

follow from the inequalities that relate packing to covering. An even easier argument would establish a stability property for the packing numbers for the union $\mathcal{F} \cup \mathcal{G}$.

Pointwise products are more interesting, for here we need the flexibility of bounds valid for arbitrary rescaling vectors. Let us show that the covering numbers for the

set $\mathcal{F} \odot \mathcal{G}$ of all pairwise products $\mathbf{f} \odot \mathbf{g}$ satisfy the inequality

$$(5.2) \quad N_2(\epsilon + \delta, \boldsymbol{\alpha} \odot \mathcal{F} \odot \mathcal{G}) \leq N_2(\epsilon, \boldsymbol{\alpha} \odot \mathbf{G} \odot \mathcal{F}) N_2(\delta, \boldsymbol{\alpha} \odot \mathbf{F} \odot \mathcal{G}),$$

which implies the corresponding inequality for packing numbers

$$D_2(2\epsilon + 2\delta, \boldsymbol{\alpha} \odot \mathcal{F} \odot \mathcal{G}) \leq D_2(\epsilon, \boldsymbol{\alpha} \odot \mathbf{G} \odot \mathcal{F}) D_2(\delta, \boldsymbol{\alpha} \odot \mathbf{F} \odot \mathcal{G}).$$

Choose approximating points $\boldsymbol{\alpha} \odot \mathbf{G} \odot \mathbf{f}_1, \dots, \boldsymbol{\alpha} \odot \mathbf{G} \odot \mathbf{f}_r$ for $\boldsymbol{\alpha} \odot \mathbf{G} \odot \mathcal{F}$, and points $\boldsymbol{\alpha} \odot \mathbf{F} \odot \mathbf{g}_1, \dots, \boldsymbol{\alpha} \odot \mathbf{F} \odot \mathbf{g}_s$ for $\boldsymbol{\alpha} \odot \mathbf{F} \odot \mathcal{G}$. We may assume each $\mathbf{f}_i$ lies within the box defined by the envelope $\mathbf{F}$, and each $\mathbf{g}_j$ lies within the box defined by $\mathbf{G}$. For an $\boldsymbol{\alpha} \odot \mathbf{f} \odot \mathbf{g}$ in the set $\boldsymbol{\alpha} \odot \mathcal{F} \odot \mathcal{G}$, and appropriate $\mathbf{f}_i$ and $\mathbf{g}_j$,

$$\begin{aligned} |\boldsymbol{\alpha} \odot \mathbf{f} \odot \mathbf{g} - \boldsymbol{\alpha} \odot \mathbf{f}_i \odot \mathbf{g}_j|_2 &\leq |\boldsymbol{\alpha} \odot \mathbf{f} \odot \mathbf{g} - \boldsymbol{\alpha} \odot \mathbf{f}_i \odot \mathbf{g}|_2 + |\boldsymbol{\alpha} \odot \mathbf{f}_i \odot \mathbf{g} - \boldsymbol{\alpha} \odot \mathbf{f}_i \odot \mathbf{g}_j|_2 \\ &\leq |\boldsymbol{\alpha} \odot \mathbf{f} \odot \mathbf{G} - \boldsymbol{\alpha} \odot \mathbf{f}_i \odot \mathbf{G}|_2 + |\boldsymbol{\alpha} \odot \mathbf{F} \odot \mathbf{g} - \boldsymbol{\alpha} \odot \mathbf{F} \odot \mathbf{g}_j|_2 \\ &\leq \epsilon + \delta. \end{aligned}$$

Inequality (5.2) fits well with the bounds from Section 4.

(5.3) LEMMA. *Suppose $\mathcal{F}$ and $\mathcal{G}$ are subsets of $\mathbb{R}^n$ for which*

$$\begin{aligned} D_1(\epsilon |\boldsymbol{\alpha} \odot \mathbf{F}|_1, \boldsymbol{\alpha} \odot \mathcal{F}) &\leq A(1/\epsilon)^W, \\ D_1(\epsilon |\boldsymbol{\alpha} \odot \mathbf{G}|_1, \boldsymbol{\alpha} \odot \mathcal{G}) &\leq A(1/\epsilon)^W, \end{aligned}$$

for $0 < \epsilon \leq 1$ and every rescaling vector $\boldsymbol{\alpha}$ of nonnegative weights. Then, for every such $\boldsymbol{\alpha}$,

$$(5.4) \quad N_2(\epsilon |\boldsymbol{\alpha} \odot \mathbf{F} \odot \mathbf{G}|_2, \boldsymbol{\alpha} \odot \mathcal{F} \odot \mathcal{G}) \leq A^2(8/\epsilon^2)^{2W} \quad \text{for } 0 < \epsilon \leq 1.$$

A similar inequality holds for the packing numbers.

PROOF. The set $\mathcal{H} = \boldsymbol{\alpha} \odot \mathcal{F} \odot \mathcal{G}$ has envelope $\mathbf{H} = \boldsymbol{\alpha} \odot \mathbf{F} \odot \mathbf{G}$, whose ℓ_2 norm,

$$|\mathbf{H}|_2 = \left(\sum_i \alpha_i^2 F_i^2 G_i^2 \right)^{1/2} = \left(|\mathbf{H} \odot \mathbf{H}|_1 \right)^{1/2},$$

provides the natural scaling factor. From inequality (5.2) and Lemma 4.9, which relates ℓ_1 and ℓ_2 packing numbers, we get

$$\begin{aligned} N_2(\epsilon |\mathbf{H}|_2, \mathcal{H}) &\leq N_2(\tfrac{1}{2}\epsilon |\mathbf{H}|_2, \boldsymbol{\alpha} \odot \mathbf{G} \odot \mathcal{F}) N_2(\tfrac{1}{2}\epsilon |\mathbf{H}|_2, \boldsymbol{\alpha} \odot \mathbf{F} \odot \mathcal{G}) \\ &\leq D_1(\tfrac{1}{8}\epsilon^2 |\mathbf{H}|_2^2, \mathbf{H} \odot \boldsymbol{\alpha} \odot \mathbf{G} \odot \mathcal{F}) D_1(\tfrac{1}{8}\epsilon^2 |\mathbf{H}|_2^2, \mathbf{H} \odot \boldsymbol{\alpha} \odot \mathbf{F} \odot \mathcal{G}). \end{aligned}$$

The set $\mathbf{H} \odot \boldsymbol{\alpha} \odot \mathbf{G} \odot \mathcal{F}$ has envelope $\mathbf{H} \odot \mathbf{H}$, which has ℓ_1 norm $|\mathbf{H}|_2^2$, and likewise for the set $\mathbf{H} \odot \boldsymbol{\alpha} \odot \mathbf{F} \odot \mathcal{G}$. With the uniform bounds on D_1 packing numbers applied to the last two factors we end up with the asserted inequality. $\square$

The results in this subsection are actually examples of a more general stability property involving *contraction maps*. A function $\boldsymbol{\lambda}$ from $\mathbb{R}^n$ into another Euclidean space is called an ℓ_2 -contraction if it satisfies the inequality

$$|\boldsymbol{\lambda}(\mathbf{f}) - \boldsymbol{\lambda}(\mathbf{g})|_2 \leq |\mathbf{f} - \mathbf{g}|_2 \quad \text{for all } \mathbf{f}, \mathbf{g} \text{ in } \mathbb{R}^n.$$

For such a map λ , it is easy to show that

$$D_2(\epsilon, \lambda(\mathcal{F})) \leq D_2(\epsilon, \mathcal{F}).$$

When applied to various cartesian products, for various maps λ from $\mathbb{R}^{2n}$ into $\mathbb{R}^n$, this would reproduce the bounds stated above.

General maximal inequalities. It is perhaps most natural—or at least most elegant—to start from the assumption that we are given bounds on quantities such as $\mathbb{P}_\sigma \Phi(\sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}|)$, for a convex, increasing nonnegative function Φ on $\mathbb{R}^+$. The bounds might have been derived by a chaining argument, based on inequalities for packing numbers, but we need not assume as much.

Without loss of generality we may assume sets such as $\mathcal{F}$ to be compact: by continuity, the supremum over $\mathcal{F}$ in each of the asserted inequalities will be equal to the supremum over the closure of $\mathcal{F}$; and the inequalities for unbounded $\mathcal{F}$ may be obtained as limiting cases of the inequalities for a sequence of bounded subsets of $\mathcal{F}$. Also we may assume that the zero vector belongs to $\mathcal{F}$.

The stability property for sums follows directly from the convexity of Φ :

$$(5.5) \quad \mathbb{P}_\sigma \Phi\left(\sup_{\mathcal{F}, \mathcal{G}} |\sigma \cdot (\mathbf{f} + \mathbf{g})|\right) \leq \mathbb{P}_\sigma \Phi\left(\sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}| + \sup_{\mathcal{G}} |\sigma \cdot \mathbf{g}|\right) \\ \leq \frac{1}{2} \mathbb{P}_\sigma \Phi\left(2 \sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}|\right) + \frac{1}{2} \mathbb{P}_\sigma \Phi\left(2 \sup_{\mathcal{G}} |\sigma \cdot \mathbf{g}|\right).$$

To eliminate the extra factors of 2 from the last two terms (or from similar terms later in this section) we could apply the same argument to the rescaled function $\Phi_0(x) = \Phi(x/2)$.

More subtle is the effect of applying a contraction operation to each coordinate of the vectors in $\mathcal{F}$. Suppose we have maps $\lambda_i : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$(5.6) \quad \lambda_i(0) = 0 \quad \text{and} \quad |\lambda_i(s) - \lambda_i(t)| \leq |s - t| \quad \text{for all real } s, t.$$

They define a contraction map on $\mathbb{R}^n$ pointwise, $\lambda(\mathbf{f}) = (\lambda_1(f_1), \dots, \lambda_n(f_n))$.

(5.7) THEOREM. *For every subset $\mathcal{F}$ of $\mathbb{R}^n$, and contraction maps λ_i ,*

$$\mathbb{P}_\sigma \Phi\left(\sup_{\mathcal{F}} |\sigma \cdot \lambda(\mathbf{f})|\right) \leq \frac{3}{2} \mathbb{P}_\sigma \Phi\left(2 \sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}|\right),$$

where $\lambda(\mathbf{f}) = (\lambda_1(f_1), \dots, \lambda_n(f_n))$. $\square$

Before proceeding to the proof, let us see how the theorem can be applied.

(5.8) EXAMPLE. We can build the class $\mathcal{F} \vee \mathcal{G}$ (or $\mathcal{F} \wedge \mathcal{G}$) using sums and contractions, based on the representation

$$f_i \vee g_i = (f_i - g_i)^+ + g_i.$$

Arguing as for (5.5) we get a bound for the set of all differences $\mathbf{f} - \mathbf{g}$. With the contraction maps $\lambda_i(s) = s^+$ we get a bound for the set of vectors with components $(f_i - g_i)^+$, which we combine with the bound for $\mathcal{G}$ using (5.5). $\square$

(5.9) EXAMPLE. If we impose the condition that $|f_i| \leq 1$ and $|g_i| \leq 1$ for all components of all vectors in $\mathcal{F}$ and $\mathcal{G}$, then we can build $\mathcal{F} \odot \mathcal{G}$ using sums and contractions, based on the representation

$$f_i g_i = \frac{1}{4}(f_i + g_i)^2 - \frac{1}{4}(f_i - g_i)^2.$$

Stability for sums (and differences) gives bounds for the sets of vectors with components $\frac{1}{2}(f_i \pm g_i)$. With the contraction map $\lambda_i(s) = \frac{1}{2} \min(1, s^2)$ we get a suitable bound for both the squared terms, which we again combine by means of inequality (5.5). $\square$

As the first step towards the proof of Theorem 5.7 we must establish a stronger result for a special case, using only elementary properties of Φ .

(5.10) LEMMA. *If $\mathcal{F}$ lies within the positive orthant of $\mathbb{R}^n$,*

$$\mathbb{P}_\sigma \Phi \left(\sup_{\mathcal{F}} |\sigma \cdot \lambda(\mathbf{f})| \right) \leq \mathbb{P}_\sigma \Phi \left(\sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}| \right)$$

for contraction maps λ_i , as in (5.6).

PROOF. It would suffice to consider the effect of the contractions one coordinate at a time. We would first show that

$$\mathbb{P}_\sigma \Phi \left(\sup_{\mathcal{F}} \left| \sum_{i < n} \sigma_i f_i + \sigma_n \lambda_n(f_n) \right| \right) \leq \mathbb{P}_\sigma \Phi \left(\sup_{\mathcal{F}} |\sigma \cdot \mathbf{f}| \right).$$

Then we could argue similarly for the $(n-1)^{st}$ coordinate—replacing the right-hand side by the quantity now on the left-hand side, and replacing f_{n-1} on the left-hand side by $\lambda_{n-1}(f_{n-1})$ —and so on.

Let us establish only the inequality for the n^{th} coordinate. Argue conditionally on $\sigma_1, \dots, \sigma_{n-1}$. To simplify the notation, write λ instead of λ_n , write $x(\mathbf{f})$ for the contribution from the first $n-1$ coordinates, and write $y(\mathbf{f})$ for f_n . Then we need to show that

$$(5.11) \quad \begin{aligned} & \Phi \left(\sup_{\mathcal{F}} |x(\mathbf{f}) + \lambda(y(\mathbf{f}))| \right) + \Phi \left(\sup_{\mathcal{F}} |x(\mathbf{f}) - \lambda(y(\mathbf{f}))| \right) \\ & \leq \Phi \left(\sup_{\mathcal{F}} |x(\mathbf{f}) + y(\mathbf{f})| \right) + \Phi \left(\sup_{\mathcal{F}} |x(\mathbf{f}) - y(\mathbf{f})| \right). \end{aligned}$$

The argument will be broken into four cases. Suppose the supremum in the first term on the left-hand side is achieved at $\mathbf{f}_0$ and for the second term at $\mathbf{f}_1$. That is, if $x_0 = x(\mathbf{f}_0)$ and so on,

$$(5.12) \quad \begin{aligned} |x_0 + \lambda(y_0)| & \geq |x(\mathbf{f}) + \lambda(y(\mathbf{f}))| \\ |x_1 - \lambda(y_1)| & \geq |x(\mathbf{f}) - \lambda(y(\mathbf{f}))| \end{aligned}$$

for all $\mathbf{f}$ in $\mathcal{F}$. For the first two cases we will need only to appeal to the facts: $\Phi(t)$ is an increasing function of t on $\mathbb{R}^+$; both y_0 and y_1 are nonnegative; and

$$(5.13) \quad |\lambda(y_i)| = |\lambda(y_i) - \lambda(0)| \leq |y_i| = y_i \quad \text{for } i = 0, 1,$$

as a consequence of the contraction property for λ .

For notational convenience, extend the function Φ by symmetry to the whole real line: $\Phi(-t) = \Phi(t)$. Then it will be enough to show that in each case at least one of the following inequalities holds:

$$(5.14) \quad \Phi(x_0 + \lambda(y_0)) + \Phi(x_1 - \lambda(y_1)) \leq \begin{cases} \Phi(x_0 + y_0) + \Phi(x_1 - y_1) \\ \Phi(x_1 + y_1) + \Phi(x_0 - y_0) \end{cases}$$

First case: if $x_0 + \lambda(y_0) \geq 0 \geq x_1 - \lambda(y_1)$, then

$$\begin{aligned} \Phi(x_0 + \lambda(y_0)) &\leq \Phi(x_0 + y_0), \\ \Phi(x_1 - \lambda(y_1)) &\leq \Phi(x_1 - y_1). \end{aligned}$$

Second case: if $x_0 + \lambda(y_0) \leq 0 \leq x_1 - \lambda(y_1)$, then

$$\begin{aligned} \Phi(x_0 + \lambda(y_0)) &\leq \Phi(x_0 - y_0), \\ \Phi(x_1 - \lambda(y_1)) &\leq \Phi(x_1 + y_1). \end{aligned}$$

At least one of the inequalities in (5.14) is clearly satisfied in both these cases.

For the other two cases, where $x_0 + \lambda(y_0)$ and $x_1 - \lambda(y_1)$ have the same sign, we need the following consequence of the convexity of Φ : if $\alpha \leq \beta$ and $\beta \geq 0$ and $0 \leq s \leq t$, then

$$(5.15) \quad \Phi(\beta + t) - \Phi(\beta) - \Phi(\alpha + s) + \Phi(\alpha) \geq 0.$$

If $s = 0$ this inequality reasserts that Φ is an increasing function on $\mathbb{R}^+$. If $s > 0$ it follows from the convexity inequality

$$\frac{\Phi(\alpha + s) - \Phi(\alpha)}{s} \leq \frac{\Phi(\beta + t) - \Phi(\beta)}{t}$$

and the nonnegativity of the ratio on the right-hand side.

Third case: if $x_0 + \lambda(y_0) \geq 0$ and $x_1 - \lambda(y_1) \geq 0$, then invoke inequality (5.15) with

$$\begin{aligned} \alpha &= x_1 - y_1, & \beta &= x_0 + \lambda(y_0), & s &= y_1 - \lambda(y_1), & t &= y_0 - \lambda(y_0) & \text{if } y_0 \geq y_1, \\ \alpha &= x_0 - y_0, & \beta &= x_1 - \lambda(y_1), & s &= y_0 + \lambda(y_0), & t &= y_1 + \lambda(y_1) & \text{if } y_0 < y_1. \end{aligned}$$

The inequalities (5.12) and (5.13) give $\alpha \leq \beta$ in each case, and the inequality $s \leq t$ follows from the contraction property

$$|\lambda(y_1) - \lambda(y_0)| \leq \begin{cases} y_0 - y_1 & \text{if } y_0 \geq y_1, \\ y_1 - y_0 & \text{if } y_0 < y_1. \end{cases}$$

Fourth case: if $x_0 + \lambda(y_0) \leq 0$ and $x_1 - \lambda(y_1) \leq 0$, then invoke (5.15) with

$$\begin{aligned} \alpha &= -x_1 - y_1, & \beta &= -x_0 - \lambda(y_0), & s &= y_1 + \lambda(y_1), & t &= y_0 + \lambda(y_0) & \text{if } y_0 \geq y_1, \\ \alpha &= -x_0 - y_0, & \beta &= -x_1 + \lambda(y_1), & s &= y_0 - \lambda(y_0), & t &= y_1 - \lambda(y_1) & \text{if } y_0 < y_1. \end{aligned}$$

The required inequalities $\alpha \leq \beta$ and $s \leq t$ are established as in the third case. $\square$

PROOF OF THEOREM 5.7. Notice that $\lambda_i(f_i) = \lambda_i(f_i^+) + \lambda_i(-f_i^-)$, because either

$$f_i \geq 0 \quad \text{and} \quad \lambda_i(f_i) = \lambda_i(f_i^+) \quad \text{and} \quad \lambda_i(-f_i^-) = \lambda_i(0) = 0,$$

or

$$f_i \leq 0 \quad \text{and} \quad \lambda_i(f_i) = \lambda_i(-f_i^-) \quad \text{and} \quad \lambda_i(f_i^+) = \lambda_i(0) = 0.$$

Convexity of Φ gives the inequality

$$\begin{aligned} \mathbb{P}_\sigma \Phi \left(\sup_{\mathcal{F}} \left| \sum_{i \leq n} \sigma_i [\lambda_i(f_i^+) + \lambda_i(-f_i^-)] \right| \right) \\ \leq \frac{1}{2} \mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \left| \sum_{i \leq n} \sigma_i \lambda_i(f_i^+) \right| \right) + \frac{1}{2} \mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \left| \sum_{i \leq n} \sigma_i \lambda_i(-f_i^-) \right| \right). \end{aligned}$$

Lemma 5.10 shows that the right-hand side increases if $\lambda_i(f_i^+)$ is replaced by f_i^+ and $\lambda_i(-f_i^-)$ is replaced by $-f_i^-$. (For $-f_i^-$, note that $\lambda(-t)$ is also a contraction mapping.) Argue from convexity of Φ and the inequality $f_i^+ = 1/2(f_i + |f_i|)$ that

$$\mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \left| \sum_{i \leq n} \sigma_i f_i^+ \right| \right) \leq \frac{1}{2} \mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \left| \sum_{i \leq n} \sigma_i f_i \right| \right) + \frac{1}{2} \mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \left| \sum_{i \leq n} \sigma_i |f_i| \right| \right),$$

with a similar inequality for the contribution from the $-f_i^-$ term. The proof will be completed by an application of the Basic Combinatorial Lemma from Section 1 to show that

$$\mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \left| \sum_{i \leq n} \sigma_i |f_i| \right| \right) \leq 2 \mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \left| \sum_{i \leq n} \sigma_i f_i \right| \right).$$

Because Φ is increasing and nonnegative, and $\mathcal{F}$ contains the zero vector,

$$\mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \left| \sum_{i \leq n} \sigma_i |f_i| \right| \right) \leq \mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \sum_{i \leq n} \sigma_i |f_i| \right) + \mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \sum_{i \leq n} (-\sigma_i) |f_i| \right).$$

The two expectations on the right-hand side are equal; it will suffice if we bound the first of them by the corresponding quantity with $|f_i|$ replaced by f_i .

To do this, let us construct, by means of the Basic Combinatorial Lemma, a one-to-one map θ from $\mathcal{S}$ onto itself such that

$$(5.16) \quad \sup_{\mathcal{F}} \sum_{i \leq n} \sigma_i |f_i| \leq \sup_{\mathcal{F}} \sum_{i \leq n} \theta(\sigma)_i f_i.$$

For each σ in $\mathcal{S}$, the compactness of $\mathcal{F}$ ensures existence of a vector $\mathbf{f}^\sigma$ for which the left-hand side of (5.16) equals

$$\sum_{i \leq n} \sigma_i |f_i^\sigma|.$$

Define the map η from $\mathcal{S}$ into itself by

$$\eta(\sigma)_i = \begin{cases} +1 & \text{if } \sigma_i = +1 \text{ and } f_i^\sigma \geq 0, \\ -1 & \text{otherwise.} \end{cases}$$

For every σ we have $\eta(\sigma) \leq \sigma$. The Basic Combinatorial Lemma gives a one-to-one map θ that has $\theta(\sigma) \wedge \sigma = \eta(\sigma)$. In particular, $\theta(\sigma)_i$ is equal to $+1$ if both $\sigma_i = +1$ and $f_i^\sigma \geq 0$, and equal to -1 if $\sigma_i = +1$ and $f_i^\sigma < 0$. Thus

$$\begin{aligned} \sum_{i \leq n} \sigma_i |f_i^\sigma| &= \sum_{\sigma_i = +1} \theta(\sigma)_i f_i^\sigma - \sum_{\sigma_i = -1} |f_i^\sigma| \\ &\leq \sum_{i \leq n} \theta(\sigma)_i f_i^\sigma \\ &\leq \sup_{\mathcal{F}} \sum_{i \leq n} \theta(\sigma)_i f_i, \end{aligned}$$

as asserted by (5.16). Because θ is one-to-one, the random vector $\theta(\sigma)$ has a uniform distribution under $\mathbb{P}_\sigma$, and

$$\mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \sum_{i \leq n} \sigma_i |f_i| \right) \leq \mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \theta(\sigma) \cdot \mathbf{f} \right) = \mathbb{P}_\sigma \Phi \left(2 \sup_{\mathcal{F}} \sigma \cdot \mathbf{f} \right),$$

as required. $\square$

REMARKS. The last subsection corresponds to a small fraction of the Ledoux and Talagrand (1989) paper. Ledoux and Talagrand (1990, Chapter 4) have further refined the method of proof. Except perhaps for the stability result for covering numbers of products, the rest of the section merely collects together small results that have been derived many times in the literature.

SECTION 6

Convex Hulls

Sometimes interesting random processes are expressible as convex combinations of more basic processes. For example, if $0 \leq f_i \leq 1$ for each i then the study of f_i reduces to the study of the random sets $\{\omega : s \leq f_i(\omega, t)\}$, for $0 \leq s \leq 1$ and $t \in T$, by means of the representation

$$f_i(\omega, t) = \int_0^1 \{s \leq f_i(\omega, t)\} ds.$$

More generally, starting from $f_i(\omega, t)$ indexed by T , we can construct new processes by averaging out over the parameter with respect to a probability measure Q on T :

$$f_i(\omega, Q) = \int f_i(\omega, t) Q(dt).$$

[This causes no measure-theoretic difficulties if there is a σ -field $\mathcal{T}$ on T such that f_i is jointly measurable in ω and t and Q is defined on $\mathcal{T}$.] Let us denote the corresponding process of sums by $S_n(\omega, Q)$, and its expectation by $M_n(Q)$. Because

$$|S_n(\omega, Q) - M_n(Q)| \leq \int \sup_t |S_n(\omega, t) - M_n(t)| Q(dt),$$

it is easy to verify that

$$(6.1) \quad \sup_Q |S_n(\omega, Q) - M_n(Q)| = \sup_t |S_n(\omega, t) - M_n(t)|.$$

Some uniformity results for the processes indexed by probability measures on T follow trivially from uniformity results for processes indexed by T .

The operation of averaging out over t corresponds to the formation of convex combinations in $\mathbb{R}^n$. The vectors with coordinates $f_1(\omega, Q), \dots, f_n(\omega, Q)$ all lie within the closed convex hull $\overline{\text{co}}(\mathcal{F}_\omega)$ of the set $\mathcal{F}_\omega$. The symmetrization analogue of the equality (6.1) is

$$\sup_{\overline{\text{co}}(\mathcal{F}_\omega)} |\sigma \cdot \mathbf{f}| = \sup_{\mathcal{F}_\omega} |\sigma \cdot \mathbf{f}|,$$

which suggests that there might be a connection between the packing numbers for $\mathcal{F}_\omega$ and the packing numbers for $\overline{\text{co}}(\mathcal{F}_\omega)$. A result of Dudley (1987) establishes such

a connection for the ubiquitous case of sets whose packing numbers grow like a power of $1/\epsilon$. Even though inequality (6.1) makes the result slightly superfluous for the purposes of these lecture notes, it is worth study as a beautiful example of a probabilistic method for proving existence theorems.

The result could be stated in great generality—for Hilbert spaces, or even for “spaces of type 2”—but the important ideas all appear for the simple case of a bounded subset of Euclidean space.

(6.2) THEOREM. *Let $\mathcal{F}$ be a subset of the unit ball in a Euclidean space. If there exist constants A and W such that*

$$D_2(\epsilon, \mathcal{F}) \leq A(1/\epsilon)^W \quad \text{for } 0 < \epsilon \leq 1,$$

then for each τ with $2 > \tau > \frac{2W}{2+W}$,

$$D_2(\epsilon, \overline{\text{co}}(\mathcal{F})) \leq \exp(C(1/\epsilon)^\tau) \quad \text{for } 0 < \epsilon \leq 1.$$

for some constant C that depends only on A , W and τ . $\square$

Note that the inequality $2 > \tau$ ensures

$$\int_0^1 \sqrt{\log D_2(x, \overline{\text{co}}(\mathcal{F}))} dx < \infty.$$

Indeed $\tau = 2$ represents the critical value at which the integral would diverge. For these notes the theorem has one major application, which deserves some attention before we get into the details of the proof for Theorem 6.2.

(6.3) EXAMPLE. Let $\mathcal{F}$ be a bounded subset of $\mathbb{R}^n$ with envelope $\mathbf{F}$. The convex cone generated by $\mathcal{F}$ is the set $\mathcal{G} = \{r\mathbf{f} : r > 0, \mathbf{f} \in \mathcal{F}\}$. Suppose $\mathcal{G}$ has the property: *for some integer V , no $(V+1)$ -dimensional coordinate projection of $\mathcal{G}$ can surround the corresponding projections of $\mathbf{F}$ or $-\mathbf{F}$.* Then Theorem 6.2 and the results from Section 4 will imply that $D_2(\epsilon|\boldsymbol{\alpha} \odot \mathbf{F}|_2, \boldsymbol{\alpha} \odot \mathcal{F}) \leq \exp[C(1/\epsilon)^\tau]$ for $0 < \epsilon \leq 1$ and all nonnegative $\boldsymbol{\alpha}$, with constants C and $\tau < 2$ depending only on V .

Without loss of generality, suppose $\boldsymbol{\alpha}$ has all components equal to one, and $\mathcal{F}$ is a subset of the positive orthant with $F_i > 0$ for each i . [The projection property of $\mathcal{G}$ still holds if we replace each $\mathbf{f}$ by the vector with coordinates f_i^+ or the vector with coordinates f_i^- .] By a trivial rescaling, replacing $\mathbf{f}$ by $\mathbf{f}/|\mathbf{F}|_2$, we may also assume that $|\mathbf{F}|_2 = 1$, so that $\mathcal{F}$ is a subset of the unit ball.

Define a new set $\mathcal{H}$ of all vectors with coordinates of the form

$$h(r, \mathbf{f})_i = F_i \{rf_i \geq F_i\},$$

where r ranges over positive real numbers and $\mathbf{f}$ ranges over $\mathcal{F}$. Certainly $\mathcal{H}$ is a subset of the unit ball. Its closed convex hull contains $\mathcal{F}$, because

$$f_i = \int_0^1 F_i \{f_i > sF_i\} ds$$

for every nonnegative f_i . We have only to check that

$$D_2(\epsilon, \mathcal{H}) \leq A(1/\epsilon)^W \quad \text{for } 0 < \epsilon \leq 1$$

then appeal to Theorem 6.2.

The geometric bound for packing numbers of $\mathcal{H}$ will follow from the results in Section 4 if we show that $(V + 1)$ -dimensional proper coordinate projections of $\mathcal{H}$ cannot surround any $\mathbf{t}$ in $\mathbb{R}^{V+1}$. Let I be the set of $V + 1$ coordinates that defines the projection. Let J be the orthant of the I -projections of $\mathbf{F}$ that the I -projection of $\mathcal{G}$ cannot occupy. Suppose, however, that the I -projection of $\mathcal{H}$ does surround some point $\mathbf{t}$. This could happen only if $0 < t_i < F_i$ for each i . For the projection of the vector $\mathbf{h}(r, \mathbf{f})$ to occupy orthant J of $\mathbf{t}$ we would need to have

$$\begin{aligned} r f_i &\geq F_i & \text{for } i \in J, \\ r f_i &< F_i & \text{for } i \in I \setminus J. \end{aligned}$$

Increasing r slightly to make these inequalities strict, we would then have found a projection of a vector in $\mathcal{G}$ occupying the orthant J . The contradiction establishes the desired projection property for $\mathcal{H}$, and hence leads to the asserted rate of growth for the packing numbers of $\mathcal{F}$. $\square$

PROOF OF THEOREM 6.2. We may as well assume that $\mathcal{F}$ is compact, because packing numbers for a set always agree with packing numbers for its closure. This makes $\overline{\text{co}}(\mathcal{F})$ the same as the convex hull $\text{co}(\mathcal{F})$, which will slightly simplify the argument.

By a succession of approximations, we will be able to construct a set with cardinality at most $\exp(C(1/\epsilon)^\tau)$ that approximates each vector of $\mathcal{F}$ within an ℓ_2 distance less than 4ϵ . With some adjustment of the constant C after replacement of ϵ by $\epsilon/8$, this would give the asserted bound for the packing numbers.

In what follows the ℓ_2 norm will be denoted by $|\cdot|$, without the subscript 2.

Let $\alpha = 2/(2 + W)$. Choose a maximal subset $\mathcal{F}_\epsilon$ of points from $\mathcal{F}$ at least ϵ apart, then let $\{\phi_1, \dots, \phi_m\}$ be a maximal subset of $\mathcal{F}_\epsilon$ with points at least ϵ^α apart. By assumption,

$$\begin{aligned} m &\leq D_2(\epsilon^\alpha, \mathcal{F}) \leq A(1/\epsilon^\alpha)^W, \\ \#\mathcal{F}_\epsilon &\leq D_2(\epsilon, \mathcal{F}) \leq A(1/\epsilon)^W. \end{aligned}$$

Notice that m is smaller than $A(1/\epsilon)^\tau$; the exponent of $1/\epsilon$ is $2W/(2 + W)$, which is less than τ . Each $\mathbf{f}$ in $\mathcal{F}$ lies within ϵ of some $\mathbf{f}^*$ in $\mathcal{F}_\epsilon$. Each finite convex combination $\sum_{\mathcal{F}} \theta(\mathbf{f})\mathbf{f}$ lies within ϵ of the corresponding $\sum_{\mathcal{F}} \theta(\mathbf{f})\mathbf{f}^*$. (Here the $\theta(\mathbf{f})$ multipliers denote nonnegative numbers that sum to one, with $\theta(\mathbf{f}) \neq 0$ for only finitely many $\mathbf{f}$.) It therefore suffices to construct approximations within 3ϵ to the vectors in $\text{co}(\mathcal{F}_\epsilon)$.

Because each vector in $\mathcal{F}_\epsilon$ lies within ϵ^α of some ϕ_i , there exists a partition of $\mathcal{F}_\epsilon$ into subsets $\mathcal{E}_1, \dots, \mathcal{E}_m$ for which

$$(6.4) \quad |\mathbf{f} - \phi_i| \leq \epsilon^\alpha \quad \text{if } \mathbf{f} \in \mathcal{E}_i.$$

Each convex combination $\sum \theta(\mathbf{f})\mathbf{f}$ from $\text{co}(\mathcal{F}_\epsilon)$ can then be reexpressed as a convex combination of vectors from the convex hulls $\text{co}(\mathcal{E}_i)$:

$$\sum_{\mathbf{f} \in \mathcal{F}_\epsilon} \theta(\mathbf{f})\mathbf{f} = \sum_{i \leq m} \lambda_i \mathbf{e}_i,$$

where

$$\lambda_i = \sum_{\mathbf{f} \in \mathcal{E}_i} \theta(\mathbf{f})$$

and

$$\mathbf{e}_i = \sum_{\mathbf{f} \in \mathcal{E}_i} \frac{\theta(\mathbf{f})}{\lambda_i} \mathbf{f}.$$

Here the vector $\boldsymbol{\lambda}$ of convex weights ranges over the m -dimensional simplex

$$\Lambda = \{\boldsymbol{\lambda} \in \mathbb{R}^m : \lambda_i \geq 0 \text{ for all } i, \text{ and } \sum_i \lambda_i = 1\}.$$

Because $\mathcal{E}_i$ lies inside the unit ball,

$$\left| \sum_{i \leq m} \lambda_i \mathbf{e}_i - \sum_{i \leq m} \mu_i \mathbf{e}_i \right| \leq \sum_{i \leq m} |\lambda_i - \mu_i|.$$

We can therefore approximate each point in $co(\mathcal{F}_\epsilon)$ within ϵ by means of a convex combination with weights $\boldsymbol{\lambda}$ chosen from a maximal subset Λ_ϵ of points from Λ at least ϵ apart in ℓ_1 distance. Notice that

$$\#\Lambda_\epsilon \leq (4/\epsilon)^m,$$

because the ℓ_1 balls of radius $\epsilon/2$ about each point in Λ_ϵ are pairwise disjoint, and their union lies within an ℓ_1 ball of radius 2.

Fix a $\boldsymbol{\lambda}$ in Λ_ϵ . Define positive integers $n(1), \dots, n(m)$ by

$$\lambda_i (1/\epsilon)^{2-2\alpha} < n(i) \leq 1 + \lambda_i (1/\epsilon)^{2-2\alpha}.$$

Let $\Phi(\boldsymbol{\lambda})$ denote the set of all convex combinations

$$\sum_{i \leq m} \lambda_i \bar{\mathbf{y}}_i$$

with $\bar{\mathbf{y}}_i$ a simple average of $n(i)$ vectors from $\mathcal{E}_i$. Its cardinality is bounded by the number of ways to choose all the averages,

$$\#\Phi(\boldsymbol{\lambda}) \leq \prod_{i \leq m} (\#\mathcal{F}_\epsilon)^{n(i)}.$$

The upper bound has logarithm less than

$$\sum_{i \leq m} n(i) \log[A(1/\epsilon)^W] \leq (m + (1/\epsilon)^{2-2\alpha}) \log[A(1/\epsilon)^W].$$

The nicest part of the argument will show, for each $\boldsymbol{\lambda}$, that each convex combination $\sum_i \lambda_i \mathbf{e}_i$ from $co(\mathcal{F}_\epsilon)$ can be approximated within 2ϵ by a vector in $\Phi(\boldsymbol{\lambda})$. Hence the union of the $\Phi(\boldsymbol{\lambda})$ as $\boldsymbol{\lambda}$ ranges over Λ_ϵ will approximate to the whole of $co(\mathcal{F}_\epsilon)$ within 3ϵ . The cardinality of this union is at most

$$(\#\Lambda_\epsilon) \max_{\boldsymbol{\lambda} \in \Lambda_\epsilon} \#\Phi(\boldsymbol{\lambda}),$$

which has logarithm less than

$$m \log(4/\epsilon) + [m + (1/\epsilon)^{2-2\alpha}] \log[A(1/\epsilon)^W].$$

The small interval between τ and

$$\frac{2W}{2+W} = \alpha W = 2 - 2\alpha$$

absorbs the factors of $\log(1/\epsilon)$, leading to the desired bound, $C(1/\epsilon)^\tau$, for an appropriately large constant C .

It remains only to prove the assertion about the approximation properties of $\Phi(\boldsymbol{\lambda})$, for a fixed $\boldsymbol{\lambda}$ in Λ_ϵ . Given $\mathbf{e}_i$ from $\text{co}(\mathcal{E}_i)$, we need to find simple averages $\bar{\mathbf{y}}_i$ of $n(i)$ vectors from $\mathcal{E}_i$ such that

$$\left| \sum_{i \leq m} \lambda_i \mathbf{e}_i - \sum_{i \leq m} \lambda_i \bar{\mathbf{y}}_i \right| \leq 2\epsilon.$$

Existence of such $\bar{\mathbf{y}}_i$ will be established probabilistically, by means of randomly generated vectors $\bar{\mathbf{Y}}_i$ for which

$$\mathbb{P} \left| \sum_{i \leq m} \lambda_i \mathbf{e}_i - \sum_{i \leq m} \lambda_i \bar{\mathbf{Y}}_i \right|^2 \leq 4\epsilon^2.$$

Some realization of the $\bar{\mathbf{Y}}_i$ must satisfy the desired inequality.

Each $\mathbf{e}_i$, as a vector in $\text{co}(\mathcal{E}_i)$, has a representation as a convex combination

$$\mathbf{e}_i = \sum_{\mathbf{f} \in \mathcal{E}_i} p_i(\mathbf{f}) \mathbf{f}.$$

Interpret $p_i(\cdot)$ as a probability distribution on $\mathcal{E}_i$. Generate independent random vectors $\mathbf{Y}_{ij}$, for $j = 1, \dots, n(i)$ and $i = 1, \dots, m$, with

$$\mathbb{P}\{\mathbf{Y}_{ij} = \mathbf{f}\} = p_i(\mathbf{f}) \quad \text{for } \mathbf{f} \in \mathcal{E}_i.$$

By this construction and inequality (6.4),

$$\begin{aligned} \mathbb{P} \mathbf{Y}_{ij} &= \mathbf{e}_i, \\ \mathbb{P} |\mathbf{Y}_{ij} - \mathbf{e}_i|^2 &\leq (\text{diam } \mathcal{E}_i)^2 \leq 4\epsilon^{2\alpha}. \end{aligned}$$

Define $\bar{\mathbf{Y}}_i$ to be the average of the $\mathbf{Y}_{ij}$ for $j = 1, \dots, n(i)$. With independence accounting for the disappearance of the crossproduct terms we get

$$\begin{aligned} \mathbb{P} \left| \sum_{i \leq m} \lambda_i (\mathbf{e}_i - \bar{\mathbf{Y}}_i) \right|^2 &= \sum_{i \leq m} \lambda_i^2 \mathbb{P} |\mathbf{e}_i - \bar{\mathbf{Y}}_i|^2 \\ &\leq \sum_{i \leq m} \lambda_i^2 4\epsilon^{2\alpha} / n(i). \end{aligned}$$

Our choice of $n(i)$ lets us bound $\lambda_i/n(i)$ by $\epsilon^{2-2\alpha}$, then sum over the remaining λ_i to end up with the desired $4\epsilon^2$. $\square$

To generalize the result to subsets $\mathcal{F}$ of more general normed linear spaces, we would need only to rejustify the last few assertions in the proof regarding the $\bar{\mathbf{Y}}_i$. Certainly the necessary cancellations are still valid for any Hilbert space. Type 2 spaces (Araujo and Giné 1980, page 158) enjoy a similar bound for $\mathcal{L}^2$ norms of sums of independent random elements, essentially by definition of the type 2 property.

REMARKS. The property of $\mathcal{F}$ introduced in Example 6.3 corresponds to the *VC major* property for classes of functions, studied by Dudley (1987). My example merely translates his result for empirical processes indexed by classes of functions to the more general setting, relaxing his assumption of bounded envelopes.

Dudley (1985) has shown that the Donsker-class property is preserved under the formation of (sequential closures of) convex hulls of classes of functions. (See the notes to Section 10 for more about Donsker classes.) This gives yet another way of handling processes representable as convex combinations of simpler processes. The same stability property is also implied by the first theorem of Talagrand (1987).

SECTION 7

Maximal Inequalities

Let us now pull together the ideas from previous sections to establish a few useful maximal inequalities for the partial-sum process S_n . To begin with, let us consider an infinite sequence of independent processes $\{f_i(\omega, t)\}$, in order to see how the bounds depend on n . This will lead us to the useful concept of a manageable triangular array of processes.

The symmetrization bound from Section 2 was stated in terms of a general convex, increasing function Φ on $\mathbb{R}^+$. The chaining inequality of Section 3 was in terms of the specific convex function given by $\Psi(x) = 1/5 \exp(x^2)$.

Section 2 related the maximum deviation of S_n from its expected value,

$$\Delta_n(\omega) = \sup_t |S_n(\omega, t) - M_n(t)|,$$

to the process $\sigma \cdot \mathbf{f}$ indexed by the random set

$$\mathcal{F}_{n\omega} = \{(f_1(\omega, t), \dots, f_n(\omega, t)) : t \in T\}.$$

If we abbreviate the supremum of $|\sigma \cdot \mathbf{f}|$ over $\mathcal{F}_{n\omega}$ to $L_n(\sigma, \omega)$, the inequality becomes

$$(7.1) \quad \mathbb{P} \Phi(\Delta_n) \leq \mathbb{P} \Phi(2L_n).$$

We bound the right-hand side by taking iterated expectations, initially conditioning on ω and averaging over σ with respect to the uniform distribution $\mathbb{P}_\sigma$.

The chaining inequality from Theorem 3.5 bounds the conditional Ψ norm of L by

$$J_n(\omega) = 9 \int_0^{\delta_n(\omega)} \sqrt{\log D(x, \mathcal{F}_{n\omega})} dx, \quad \text{where } \delta_n(\omega) = \sup_{\mathcal{F}_{n\omega}} |\mathbf{f}|.$$

Here, and throughout the section, the subscript 2 is omitted from the ℓ_2 norm $|\cdot|_2$; we will make no use of the ℓ_1 norm in this section. Written out more explicitly, the inequality that defines the Ψ norm becomes

$$(7.2) \quad \mathbb{P}_\sigma \exp(L_n(\sigma, \omega)/J_n(\omega))^2 \leq 5.$$

Because J_n is a random variable, in general we cannot appeal directly to inequality (7.1) with $\Phi(x) = \exp(x^2/2J_n^2)$, to get some sort of bound for the Ψ norm of the

partial-sum process. We can, however, combine the two inequalities to get several most useful bounds.

The simplest situation occurs when $J_n(\omega)$ is bounded by a constant K_n . As we shall see soon, this often happens when the envelopes F_i are uniformly bounded. Increasing J_n to K_n in (7.2), then taking expectations we get, via (7.1),

$$\mathbb{P} \exp(\tfrac{1}{2} \Delta_n^2 / K_n^2) \leq 5.$$

It follows that Δ_n has subgaussian tails:

$$(7.3) \quad \mathbb{P}\{\Delta_n \geq t\} \leq 5 \exp(-\tfrac{1}{2} t^2 / K_n^2) \quad \text{for all } t > 0.$$

This is not the best subgaussian upper bound; the constant K_n could be replaced by a smaller constant.

If $J_n(\omega)$ is not uniformly bounded, but instead has a finite Ψ norm, we still get an exponential bound on the tail probabilities for Δ_n , by means of the inequality

$$2L_n/C \leq J_n^2/C^2 + L_n^2/J_n^2 \quad \text{for constant } C.$$

With $C = \|J_n\|_\Psi$ this inequality implies

$$\begin{aligned} \mathbb{P} \exp(\Delta_n/C) &\leq \mathbb{P} \exp(2L_n/C) \\ &\leq \mathbb{P}_\omega [\exp(J_n^2/C^2) \mathbb{P}_\sigma \exp(L_n^2/J_n^2)] \\ &\leq 25. \end{aligned}$$

Consequently,

$$(7.4) \quad \mathbb{P}\{\Delta_n \geq t\} \leq 25 \exp(-t/\|J_n\|_\Psi) \quad \text{for all } t > 0.$$

We have traded a strong moment condition on J_n for a rapid rate of decrease of the Δ_n tail probabilities.

With weaker moment bounds on J_n we get weaker bounds on Δ_n . Remember that for each p with $1 \leq p < \infty$ there is a constant C_p such that

$$\|Z\|_p \leq C_p \|Z\|_\Psi$$

for every random variable Z . In particular,

$$\mathbb{P}_\sigma |L_n|^p \leq (C_p J_n(\omega))^p,$$

which gives

$$(7.5) \quad \mathbb{P}|\Delta_n|^p \leq (2C_p)^p \mathbb{P}J_n^p.$$

This inequality will be most useful for p equal to 1 or 2.

The preceding inequalities show that the behavior of the random variable $J_n(\omega)$ largely determines the form of the maximal inequality for the partial-sum process. In one very common special case, which is strongly recommended by the results from Section 4, the behavior of J_n is controlled by the envelope $\mathbf{F}_n(\omega)$. Let us suppose that $\lambda_n(\cdot)$ is a deterministic function for which

$$(7.6) \quad D(x|\mathbf{F}_n(\omega)|, \mathcal{F}_{n\omega}) \leq \lambda_n(x) \quad \text{for } 0 < x \leq 1 \text{ and all } \omega.$$

Because $\mathcal{F}_{n\omega}$ lies within a ball of radius $|\mathbf{F}_n(\omega)|$, we could always choose $\lambda_n(x)$ equal to $(3/x)^n$. [We can pack $D(x|\mathbf{F}_n(\omega)|, \mathcal{F}_{n\omega})$ many disjoint balls of radius $^{1/2}x|\mathbf{F}_n(\omega)|$

into the ball of radius $3/2|\mathbf{F}_n(\omega)|$.] To be of any real use, however, the λ_n function should not increase so rapidly with n . For example, if there is a fixed V such that each $\mathcal{F}_{n\omega}$ has pseudodimension V we could choose $\lambda_n(x) = Ax^{-W}$, with A and W depending only on V , which would lead to quite useful bounds. In any case, we may always assume that $\sqrt{\log \lambda_n}$ is integrable, which ensures that the function defined by

$$\Lambda_n(t) = \int_0^t \sqrt{\log \lambda_n(x)} dx \quad \text{for } 0 \leq t \leq 1$$

is well defined and finite. A simple change of variable in the integral that defines $J_n(\omega)$ now gives

$$(7.7) \quad \begin{aligned} J_n(\omega) &\leq 9|\mathbf{F}_n(\omega)|\Lambda_n(\delta_n(\omega)/|\mathbf{F}_n(\omega)|) \\ &\leq 9\Lambda_n(1)|\mathbf{F}_n(\omega)| \quad \text{because } |\mathbf{f}| \leq |\mathbf{F}_n(\omega)| \text{ for every } \mathbf{f} \text{ in } \mathcal{F}_{n\omega}. \end{aligned}$$

When expressed in terms of Λ_n the inequalities for Δ_n take a particularly simple form. Suppose, for example, the envelope functions $F_i(\omega)$ are uniformly bounded, say $F_i(\omega) \leq 1$ for each i and each ω . Then $J_n(\omega)$ is bounded by $9\sqrt{n}\Lambda_n(1)$. If $\Lambda_n(1)$ stays bounded as $n \rightarrow \infty$, the standardized processes

$$\frac{1}{\sqrt{n}}\Delta_n(\omega) = \frac{1}{\sqrt{n}} \sup_t |S_n(\omega, t) - M_n(t)|$$

will have uniformly subgaussian tails.

If instead of being uniformly bounded the random variables F_i^2 have uniformly bounded moment generating functions in a neighborhood of the origin, and if $\Lambda_n(1)$ stays bounded as $n \rightarrow \infty$, we get another useful bound on the Ψ norms of the J_n . For suppose that

$$\mathbb{P} \exp(\epsilon F_i^2) \leq K \quad \text{for all } i.$$

Then there is a constant K' , depending on K and ϵ , such that

$$\mathbb{P} \exp(s F_i^2) \leq 1 + K's \quad \text{for } 0 \leq s \leq \epsilon \text{ and all } i.$$

With $C = 9 \sup_n \Lambda_n(1)$, independence of the F_i gives, for $C' \geq C^2/n\epsilon$,

$$\begin{aligned} \mathbb{P} \exp(J_n^2/nC') &\leq \prod_{i \leq n} \mathbb{P} \exp(C^2 F_i^2/nC') \\ &\leq (1 + K'C^2/nC')^n. \end{aligned}$$

Certainly for $C' \geq K'C^2/\log 5$ the last bound is less than 5. It follows that

$$\|J_n\|_\Psi \leq K''\sqrt{n} \quad \text{for some constant } K'',$$

which guarantees a uniform exponential bound for the tail probabilities of the partial-sum processes with the usual standardization.

Finally, even with only moment bounds for the envelopes we still get usable maximal inequalities. For $1 \leq p < \infty$, inequalities (7.5) and (7.7) give

$$(7.8) \quad \begin{aligned} \mathbb{P} \sup_t |S_n(\cdot, t) - M_n(t)|^p &\leq (18C_p)^p \mathbb{P} |\mathbf{F}_n|^p \Lambda_n(\delta_n/|\mathbf{F}_n|)^p \\ &\leq (18C_p \Lambda_n(1))^p \mathbb{P} |\mathbf{F}_n|^p. \end{aligned}$$

In applications such moment bounds are often the easiest to apply, typically for p equal to 1 or 2. They show that, in some sense, the whole process is only as badly behaved as its envelope.

The special cases considered above show that maximal inequalities for Δ_n can be derived from uniform bounds on the random packing numbers $D(x|\mathbf{F}_n(\omega)|, \mathcal{F}_{n\omega})$. The concept of *manageability* formalizes this idea. To accommodate a wider range of applications, let us expand the setting to cover triangular arrays of random processes,

$$\{f_{ni}(\omega, t) : t \in T, 1 \leq i \leq k_n\} \quad \text{for } n = 1, 2, \dots,$$

independent within each row. Now $S_n(\omega, t)$ denotes the sum across the n^{th} row. To facilitate application of the stability arguments, let us also allow for nonnegative rescaling vectors.

(7.9) DEFINITION. Call a triangular array of processes $\{f_{ni}(\omega, t)\}$ *manageable* (with respect to the envelopes $\mathbf{F}_n(\omega)$) if there exists a deterministic function λ , the *capacity bound*, for which

$$(i) \int_0^1 \sqrt{\log \lambda(x)} dx < \infty,$$

$$(ii) D(x|\boldsymbol{\alpha} \odot \mathbf{F}_n(\omega)|, \boldsymbol{\alpha} \odot \mathcal{F}_{n\omega}) \leq \lambda(x) \text{ for } 0 < x \leq 1, \text{ all } \omega, \text{ all vectors } \boldsymbol{\alpha} \text{ of nonnegative weights, and all } n.$$

Call a sequence of processes $\{f_i\}$ manageable if the array defined by $f_{ni} = f_i$ for $i \leq n$ is manageable.

In the special case where $\lambda(x) = A(1/x)^W$ for constants A and W , the processes will be called *Euclidean*. Most of the the applications in the final sections of these notes will involve Euclidean processes.

The inequalities developed in this section all carry over to the more general setting. In particular, for a manageable array there is a continuous, increasing function Λ with $\Lambda(0) = 0$, for which the analogue of (7.8) holds: for $1 \leq p < \infty$ there exists a constant K_p such that

$$(7.10) \quad \begin{aligned} \mathbb{P} \sup_t |S_n(\cdot, t) - M_n(t)|^p &\leq K_p \mathbb{P} |\mathbf{F}_n|^p \Lambda(\delta_n/|\mathbf{F}_n|)^p \\ &\leq K_p \Lambda(1)^p \mathbb{P} |\mathbf{F}_n|^p. \end{aligned}$$

REMARKS. When specialized to empirical processes, the exponential inequality (7.3) is inferior to the results of Alexander (1984) and Massart (1986). By refinement of the approach in this section my inequality could be improved. However, a reader interested in better bounds would be well advised to first consult the book of Ledoux and Talagrand (1990).

SECTION 8

Uniform Laws of Large Numbers

For many estimation procedures, the first step in a proof of asymptotic normality is an argument to establish consistency. For estimators defined by some sort of maximization or minimization of a partial-sum process, consistency often follows by a simple continuity argument from an appropriate uniform law of large numbers. The maximal inequalities from Section 7 offer a painless means for establishing such uniformity results. This section will present both a uniform weak law of large numbers (convergence in probability) and a uniform strong law of large numbers (convergence almost surely).

The proof of the weak law will depend upon the following consequence of the first two lemmas from Section 3: *for every finite subset $\mathcal{F}$ of $\mathbb{R}^n$,*

$$(8.1) \quad \mathbb{P}_\sigma \max_{\mathbf{f} \in \mathcal{F}} |\boldsymbol{\sigma} \cdot \mathbf{f}| \leq C \max_{\mathbf{f} \in \mathcal{F}} \|\mathbf{f}\|_2 \sqrt{2 + \log(\#\mathcal{F})}.$$

Here $\#\mathcal{F}$ denotes the number of vectors in $\mathcal{F}$, as usual, and C is a constant derived from the inequality between $\mathcal{L}^1$ and $\mathcal{L}^\Psi$ norms.

(8.2) THEOREM. *Let $f_1(\omega, t)$, $f_2(\omega, t)$, $\dots$ be independent processes with integrable envelopes $F_1(\omega)$, $F_2(\omega)$, $\dots$. If for each $\epsilon > 0$*

(i) there is a finite K such that

$$\frac{1}{n} \sum_{i \leq n} \mathbb{P} F_i \{F_i > K\} < \epsilon \quad \text{for all } n,$$

(ii) $\log D_1(\epsilon | \mathbf{F}_n|, \mathcal{F}_{n\omega}) = o_p(n)$,
then

$$\frac{1}{n} \sup_t |S_n(\omega, t) - M_n(t)| \rightarrow 0 \quad \text{in probability.}$$

PROOF. Let us establish convergence in $\mathcal{L}^1$. Given $\epsilon > 0$, choose K as in assumption (i) and then define $f_i^*(\omega, t) = f_i(\omega, t) \{F_i(\omega) \leq K\}$. The variables

discarded by this truncation contribute less than 2ϵ :

$$\frac{1}{n} \mathbb{P} \sup_t \left| \sum_{i \leq n} (f_i - f_i^*) - \mathbb{P}(f_i - f_i^*) \right| \leq \frac{2}{n} \sum_{i \leq n} \mathbb{P} F_i \{F_i > K\}.$$

For the remaining contributions from the $f_i^*(\omega, t)$ processes, invoke the symmetrization inequality from Theorem 2.2, with Φ equal to the identity function.

$$\frac{1}{n} \mathbb{P} \sup_t \left| \sum_{i \leq n} f_i^* - \mathbb{P} f_i^* \right| \leq \frac{2}{n} \mathbb{P} \mathbb{P}_\sigma \sup_{\mathcal{F}_{n\omega}} |\boldsymbol{\sigma} \cdot \mathbf{f}^*|.$$

Given ω , find a set $\mathcal{D}_{n\omega}$ of at most $M_n = D_1(\epsilon|\mathbf{F}_n|, \mathcal{F}_{n\omega})$ many points in $\mathcal{F}_{n\omega}$ that approximate each point of $\mathcal{F}_{n\omega}$ within an ℓ_1 distance of $\epsilon|\mathbf{F}_n|_1$. By assumption (ii), the random variables $\{\log M_n\}$ are of order $o_p(n)$. The expectation with respect to $\mathbb{P}_\sigma$ on the right-hand side of the last expression is less than

$$\frac{\epsilon}{n} |\mathbf{F}_n|_1 + \frac{1}{n} \mathbb{P}_\sigma \max_{\mathcal{D}_{n\omega}} |\boldsymbol{\sigma} \cdot \mathbf{f}^*|.$$

The first of these terms has a small expectation, because assumption (i) implies uniform boundedness of $\frac{1}{n} \mathbb{P} |\mathbf{F}_n|_1$. The second term is bounded by K . By virtue of inequality (8.1) it is also less than

$$\frac{C}{n} \max_{\mathcal{D}_{n\omega}} |\mathbf{f}^*|_2 \sqrt{2 + 2 \log M_n}.$$

The square root factor contributes at most $o_p(\sqrt{n})$ to this bound. The other factor is of order $O_p(\sqrt{n})$, because, for each point in $\mathcal{F}_{n\omega}$,

$$|\mathbf{f}^*|_2^2 = \sum_{i \leq n} f_i^2 \{F_i \leq K\} \leq K \sum_{i \leq n} F_i.$$

A uniformly bounded sequence that converges in probability to zero also converges to zero in $\mathcal{L}^1$. $\square$

When the processes $\{f_i(\omega, t)\}$ are identically distributed, the convergence in probability asserted by the theorem actually implies the stronger almost sure convergence, because the random variables

$$\frac{1}{n} \sup_t |S_n(\omega, t) - M_n(t)|$$

form a reversed submartingale. (Modulo measurability scruples, the argument for empirical processes given by Pollard (1984, page 22) carries over to the present context.) Without the assumption of identical distributions, we must strengthen the hypotheses of the theorem in order to deduce almost sure convergence. Manageability plus a second moment condition analogous to the requirement for the classical Kolmogorov strong law of large numbers will suffice. The stronger assumption about the packing numbers will not restrict our use of the resulting uniform strong law of large numbers for the applications in these notes; we will usually need manageability for other arguments leading to asymptotic normality.

(8.3) THEOREM. *Let $\{f_i(\omega, t) : t \in T\}$ be a sequence of independent processes that are manageable for their envelopes $\{F_i(\omega)\}$. If*

$$\sum_{i=1}^{\infty} \frac{\mathbb{P}F_i^2}{i^2} < \infty,$$

then

$$\frac{1}{n} \sup_t |S_n(\omega, t) - M_n(t)| \rightarrow 0 \quad \text{almost surely.}$$

PROOF. Define

$$\begin{aligned} f_i^*(\omega, t) &= f_i(\omega, t) - \mathbb{P}f_i(\cdot, t), \\ Z_{k,n}(\omega) &= \sup_t \left| \frac{f_k^*(\omega, t)}{k} + \cdots + \frac{f_n^*(\omega, t)}{n} \right| \quad \text{for } k \leq n, \\ B_k(\omega) &= \sup_{i,n \geq k} Z_{i,n}(\omega). \end{aligned}$$

By the triangle inequality

$$\begin{aligned} \sup_t \left| f_1^*(\omega, t) + \cdots + f_n^*(\omega, t) \right| &\leq Z_{1,n}(\omega) + \cdots + Z_{n,n}(\omega) \\ &\leq B_1(\omega) + \cdots + B_n(\omega). \end{aligned}$$

It therefore suffices to prove that $B_n \rightarrow 0$ almost surely.

From inequality (7.10) applied to the processes $f_i^*(\omega, t)/i$ instead of to $f_i(\omega, t)$, manageability implies existence of a constant C such that

$$(8.4) \quad \mathbb{P}Z_{k,n}^2 \leq C \sum_{i=k}^n \frac{\mathbb{P}F_i^2}{i^2} \quad \text{for } k \leq n.$$

For fixed k , the random variables $Z_{k,n}$ for $n = k, k+1, \dots$ form a submartingale. By Doob's (1953, page 317) inequality for nonnegative submartingales, for each m greater than k ,

$$\mathbb{P} \max_{k \leq n \leq m} Z_{k,n}^2 \leq 4\mathbb{P}Z_{k,m}^2.$$

Letting m tend to ∞ , we deduce for each k that

$$\mathbb{P} \sup_{k \leq n} Z_{k,n}^2 \leq 4C \sum_{i=k}^{\infty} \frac{\mathbb{P}F_i^2}{i^2}.$$

The sum on the right-hand side converges to zero as $k \rightarrow \infty$. From the bound

$$B_k \leq 2 \sup_{k \leq n} Z_{k,n}$$

it follows that $\mathbb{P}B_k^2 \rightarrow 0$. Because $\{B_k\}$ is a decreasing sequence of random variables, it follows that $B_k \rightarrow 0$ almost surely, as required. $\square$

REMARKS. Theorem 8.2 is based on Theorem 8.3 of Giné and Zinn (1984). They established both necessity and sufficiency for empirical processes with independent, identically distributed summands. The direct use of inequality (8.1)

simplifies the argument, by avoiding an appeal to a chaining inequality. Except for the use of ℓ_2 packing numbers instead of ℓ_1 covering numbers, my proof is close to the proof of Theorem II.24 of Pollard (1984).

Theorem 8.3 is actually a special case of the Ito-Nisio Theorem (Jain and Marcus 1978, Section II.3). Zaman (1989) used a type 2 inequality analogous to (8.4) to reduce his proof of a uniform strong law of large numbers to the Ito-Nisio theorem. He imposed the same sort of moment condition as in Theorem 8.3.

SECTION 9

Convergence in Distribution and Almost Sure Representation

Classical limit theorems for sums of independent random vectors will often suggest a standardization for the partial-sum process, S_n , so that its finite dimensional projections have a limiting distribution. It is then natural to ask whether the standardized stochastic process also has a limiting distribution, in some appropriate sense. The traditional sense has been that of a functional limit theorem. One identifies some metric space of real-valued functions on T that contains all the standardized sample paths, and then one invokes a general theory for convergence in distribution of random elements of a metric space (or weak convergence of probability measures on the space).

For example, if $T = [0, 1]$ and the sample paths of S_n have only simple discontinuities, the theory of weak convergence for $D[0, 1]$ might apply.

Unfortunately, even for such simple processes as the empirical distribution function for samples from the Uniform $[0, 1]$ distribution, awkward measurability complications arise. With $D[0, 1]$ either one skirts the issue by adopting a Skorohod metric, or one retains the uniform metric at the cost of some measure theoretic modification of the definition of convergence in distribution.

For index sets more complicated than $[0, 1]$ there is usually no adequate generalization of the Skorohod metric. The measurability complications cannot be defined away. One must face the possibility that the expectations appearing in plausible definitions for convergence in distribution need not be well defined. Of the numerous general theories proposed to handle this problem, the one introduced by Hoffmann-Jørgensen (and developed further by Dudley 1985) is undoubtedly the best. It substitutes outer expectations for expectations. It succeeds where other

theories fail because it supports an almost sure representation theorem; with suitable reinterpretation, most of the useful results from the classical theory carry over to the new theory.

The theory concerns sequences of maps $\{X_n\}$ from a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ into a metric space $\mathcal{X}$. If each X_n is measurable with respect to the Borel σ -field $\mathcal{B}(\mathcal{X})$, convergence in distribution to a probability measure P on $\mathcal{B}(\mathcal{X})$ can conveniently be defined to mean

$$\mathbb{P}f(X_n) \rightarrow Pf \quad \text{for every } f \text{ in } \mathcal{U}(\mathcal{X}),$$

where $\mathcal{U}(\mathcal{X})$ stands for the class of all bounded, uniformly continuous, real functions on $\mathcal{X}$. If X_n has no particular measurability properties, $f(X_n)$ need not be measurable; the expectation $\mathbb{P}f(X_n)$ need not be well defined. But the outer (or inner) expectation is defined: for each bounded, real-valued H on Ω ,

$$\mathbb{P}^*H = \inf\{\mathbb{P}h : H \leq h \text{ and } h \text{ integrable}\}.$$

The inner expectation $\mathbb{P}_*H$ is defined analogously. The new definition of convergence in distribution replaces $\mathbb{P}$ by $\mathbb{P}^*$, while retaining some measure theoretic regularity for the limit P in order to exclude some unpleasant cases.

(9.1) DEFINITION. If $\{X_n\}$ is a sequence of (not necessarily Borel measurable) maps from Ω into a metric space $\mathcal{X}$, and if P is a probability measure on the Borel σ -field $\mathcal{B}(\mathcal{X})$, then $X_n \rightsquigarrow P$ (read as “ X_n converges in distribution to P ”) is defined to mean $\mathbb{P}^*f(X_n) \rightarrow Pf$ for every f in $\mathcal{U}(\mathcal{X})$.

The equality $\mathbb{P}_*f(X_n) = -\mathbb{P}^*[-f(X_n)]$ shows that the definition could be stated, equivalently, in terms of convergence of inner expectations. It could also be stated in terms of convergence to a Borel measurable random element X : one replaces Pf by $\mathbb{P}f(X)$.

In requiring convergence only for f in $\mathcal{U}(\mathcal{X})$ my definition departs slightly from the Hoffmann-Jørgensen and Dudley definitions, where f runs over all bounded, continuous functions. The departure makes it slightly easier to prove some basic facts without changing the meaning of the concept in important cases.

(9.2) EXAMPLE. Here is a result that shows the convenience of requiring uniform continuity for f in Definition 9.1. *If $\{Y_n\}$ is a sequence of random elements of a metric space $(\mathcal{Y}, e)$ which converges in probability to a constant y , that is, $\mathbb{P}^*\{e(Y_n, y) > \delta\} \rightarrow 0$ for each $\delta > 0$, and if $X_n \rightsquigarrow X$, then $(X_n, Y_n) \rightsquigarrow (X, y)$.* For if f is a uniformly continuous function on $\mathcal{X} \otimes \mathcal{Y}$, bounded in absolute value by a constant M , then, for an appropriate choice of δ ,

$$f(X_n, Y_n) \leq f(X_n, y) + \epsilon + 2M\{e(Y_n, y) > \delta\}.$$

Taking outer expectations of both sides then letting $n \rightarrow \infty$, we get

$$\limsup \mathbb{P}^*f(X_n, Y_n) \leq \limsup \mathbb{P}^*f(X_n, y) + \epsilon.$$

Uniform continuity of $f(\cdot, y)$ ensures that the right-hand side equals $\mathbb{P}f(X, y) + \epsilon$. Replacement of f by $-f$ would give the companion lower bound needed to establish the required convergence. $\square$

It has long been recognized (see Pyke 1969, for example) that many arguments involving convergence in distribution are greatly simplified by use of a technical device known as *almost sure representation*. Such a representation usually asserts something like:

*If $X_n \rightsquigarrow P$ then there exist $\tilde{X}_n$ and $\tilde{X}$ such that
 $\tilde{X}_n$ and X_n have the same distribution, $\tilde{X}$ has
distribution P , and $\tilde{X}_n \rightarrow \tilde{X}$ almost surely.*

The random elements $\tilde{X}_n$ and $\tilde{X}$ are defined on a new probability space $(\tilde{\Omega}, \tilde{\mathcal{A}}, \tilde{\mathbb{P}})$. For Borel measurable X_n , “the same distribution” is interpreted to mean that

$$\mathbb{P}g(X_n) = \tilde{\mathbb{P}}g(\tilde{X}_n) \quad \text{for all bounded, Borel measurable } g.$$

Without the measurability, it would seem natural to require equality of outer expectations. Dudley’s (1985) form of the representation theorem achieves this in a particularly strong form.

With the Dudley representation, $\tilde{\mathcal{A}} \setminus \mathcal{A}$ -measurable maps ϕ_n from $\tilde{\Omega}$ into Ω are constructed to be *perfect* in the sense that not only is $\mathbb{P}$ the image of $\tilde{\mathbb{P}}$ under each ϕ_n , but also

$$\mathbb{P}^*H = \tilde{\mathbb{P}}^*H \circ \phi_n \quad \text{for every bounded } H \text{ on } \Omega.$$

The representing random elements $\tilde{X}_n$ are defined by

$$\tilde{X}_n(\tilde{\omega}) = X_n(\phi_n(\tilde{\omega})) \quad \text{for each } \tilde{\omega} \text{ in } \tilde{\Omega}.$$

Thus $\mathbb{P}^*g(X_n) = \tilde{\mathbb{P}}^*g(\tilde{X}_n)$ for every bounded g on $\mathcal{X}$, regardless of its measurability properties. In general the outer integrals satisfy only an inequality,

$$\mathbb{P}^*H \geq \tilde{\mathbb{P}}^*(H \circ \phi_n) \quad \text{for every bounded } H \text{ on } \Omega,$$

because $h \circ \phi_n \geq H \circ \phi_n$ whenever $h \geq H$. To establish that ϕ_n is perfect it is therefore enough to prove that

$$(9.3) \quad \mathbb{P}^*H \leq \tilde{\mathbb{P}}^*g \quad \text{for all } \tilde{\mathcal{A}}\text{-measurable } g \geq H \circ \phi_n.$$

We then get the companion inequality by taking the infimum over all such g .

The Dudley representation also strengthens the sense in which the representing sequence converges. For possibly nonmeasurable random elements mere pointwise convergence would not suffice for applications.

(9.4) REPRESENTATION THEOREM. *If $X_n \rightsquigarrow P$ in the sense of Definition 9.1, and if the limit distribution P concentrates on a separable Borel subset $\mathcal{X}_0$ of $\mathcal{X}$, then there exists a probability space $(\tilde{\Omega}, \tilde{\mathcal{A}}, \tilde{\mathbb{P}})$ supporting $\tilde{\mathcal{A}} \setminus \mathcal{A}$ -measurable maps ϕ_n into Ω and an $\tilde{\mathcal{A}} \setminus \mathcal{B}(\mathcal{X})$ -measurable map $\tilde{X}$ into $\mathcal{X}_0$, such that:*

- (i) *each ϕ_n is a perfect map, in the sense that $\mathbb{P}^*H = \tilde{\mathbb{P}}^*(H \circ \phi_n)$ for every bounded H on Ω ;*
- (ii) *$\tilde{\mathbb{P}}\tilde{X}^{-1} = P$, as measures on $\mathcal{B}(\mathcal{X})$;*
- (iii) *there is a sequence of $\tilde{\mathcal{A}} \setminus \mathcal{B}[0, \infty]$ -measurable, extended-real-valued random variables $\{\delta_n\}$ on $\tilde{\Omega}$ for which $d(\tilde{X}_n(\tilde{\omega}), \tilde{X}(\tilde{\omega})) \leq \delta_n(\tilde{\omega}) \rightarrow 0$ for almost every $\tilde{\omega}$, where $\tilde{X}_n(\tilde{\omega}) = X_n(\phi_n(\tilde{\omega}))$.*

It is easy to show that if (i), (ii), and (iii) hold then $X_n \rightsquigarrow P$; the assertions of the theorem are more natural than they might appear at first glance.

A sketch of Dudley's construction will close out this section. But first an example—a revamped Continuous Mapping Theorem—to show how perfectness compensates for the lack of measurability. In my opinion, an unencumbered form of Continuous Mapping Theorem is essential for any general theory of convergence in distribution.

(9.5) EXAMPLE. Suppose $X_n \rightsquigarrow P$ with P concentrated on a separable Borel subset $\mathcal{X}_0$ of $\mathcal{X}$. Suppose τ is a map into another metric space $\mathcal{Y}$ such that

- (i) the restriction of τ to $\mathcal{X}_0$ is Borel measurable,
- (ii) τ is continuous at P almost all points of $\mathcal{X}_0$.

Then we can deduce from the Representation Theorem that $\tau(X_n)$ converges in distribution to the image measure $P\tau^{-1}$.

Fix an f in $\mathcal{U}(\mathcal{Y})$. Define $h = f \circ \tau$. We need to verify that $\mathbb{P}^*h(X_n) \rightarrow Ph$. With no loss of generality we may suppose $0 \leq h \leq 1$. Fix an $\epsilon > 0$. For each positive integer k define G_k to be the open set of all points x in $\mathcal{X}$ for which h oscillates by $> \epsilon$ within the open ball of radius $1/k$ and center x . [That is, there are points y and z with $|h(y) - h(z)| > \epsilon$ and $d(x, y) < 1/k$ and $d(x, z) < 1/k$. The same y and z will provide oscillation $> \epsilon$ for every center close enough to x .]

As $k \rightarrow \infty$ the set G_k shrinks down to a set that excludes all continuity points of τ , and thereby has zero P measure. We can therefore find a k such that $PG_k < \epsilon$.

The definition of G_k ensures that if $\tilde{X}(\tilde{\omega}) \notin G_k$ and if $\delta_n(\tilde{\omega}) < 1/k$ then

$$|h(\tilde{X}_n(\tilde{\omega})) - h(\tilde{X}(\tilde{\omega}))| \leq \epsilon.$$

Consequently,

$$h(\tilde{X}_n) \leq (\epsilon + h(\tilde{X}))\{\tilde{X} \notin G_k, \delta_n < 1/k\} + \{\tilde{X} \in G_k\} + \{\delta_n \geq 1/k\}.$$

The expression on the right-hand side is measurable; it is one of the measurable functions that enters into the definition of the outer expectation of $h(\tilde{X}_n)$. It follows that

$$\tilde{\mathbb{P}}^*h(\tilde{X}_n) \leq \epsilon + \tilde{\mathbb{P}}h(\tilde{X}) + \tilde{\mathbb{P}}\{\tilde{X} \in G_k\} + \tilde{\mathbb{P}}\{\delta_n \geq 1/k\}.$$

Measurability of δ_n and dominated convergence ensure that the last probability tends to zero. And the perfectness property lets us equate the left-hand side with $\mathbb{P}^*h(X_n)$. Passing to the limit we deduce

$$\limsup \mathbb{P}^*h(X_n) \leq Ph.$$

An analogous argument with h replaced by $1 - h$ gives the companion lower bound needed to establish the desired convergence. $\square$

OUTLINE OF A PROOF OF THE REPRESENTATION THEOREM

Step 1. The indicator function of a closed ball with zero P measure on its boundary can be sandwiched between two functions from $\mathcal{U}(\mathcal{X})$ whose expectations are arbitrarily close. If B is an intersection of finitely many such balls, the approximating functions can be combined to construct f_1 and f_2 in $\mathcal{U}(\mathcal{X})$ such that

$P(f_1 - f_2) < \epsilon$ and

$$f_1(X_n) \geq \{X_n \in B\} \geq f_2(X_n).$$

Taking outer and inner expectations, then passing to the limit, we deduce that

$$\begin{aligned} \mathbb{P}^*\{X_n \in B\} &\rightarrow PB, \\ \mathbb{P}_*\{X_n \in B\} &\rightarrow PB, \end{aligned}$$

for every such B . (We will need this result only for sets B constructed from balls with centers in $\mathcal{X}_0$.)

Step 2. If π is a partition of $\mathcal{X}$ generated by a finite collection of closed balls, each with zero P measure on its boundary, then

$$\mathbb{P}_*\{X_n \in B\} \rightarrow PB \quad \text{for each } B \text{ in } \pi.$$

This follows from Step 1, because the sets in π are proper differences of intersections of finitely many closed balls.

Step 3. For each positive integer k , cover $\mathcal{X}_0$ by closed balls of diameter less than $1/k$, with zero P measure on their boundaries. Use separability of $\mathcal{X}_0$ to extract a countable subcover, then use countable additivity of P to find a subcollection that covers all of $\mathcal{X}_0$ except for a piece with P measure less than 2^{-k} . Generate a finite partition $\pi(k)$ of $\mathcal{X}$ from this collection. All except one of the sets in $\pi(k)$ has diameter less than $1/k$, and that one has P measure less than 2^{-k} . The convergence property from Step 2 gives an $n(k)$ such that

$$\mathbb{P}_*\{X_n \in B\} \geq (1 - 2^{-k})PB \quad \text{for all } B \text{ in } \pi(k), \text{ all } n \geq n(k).$$

Step 4. Assuming that $1 = n(0) < n(1) < \dots$, define $\gamma(n)$ to equal the k for which $n(k) \leq n < n(k+1)$. For $\gamma(n) = k$ and each B_i in $\pi(k)$, find measurable A_{ni} with $A_{ni} \subseteq X_n^{-1}B_i$ and

$$\mathbb{P}A_{ni} = \mathbb{P}_*\{X_n \in B_i\}.$$

Define a probability measure μ_n on $\mathcal{A}$ by

$$2^{-\gamma(n)}\mu_n(\cdot) + \left(1 - 2^{-\gamma(n)}\right) \sum_i PB_i \mathbb{P}(\cdot \mid A_{ni}) = \mathbb{P}(\cdot).$$

The inequality from Step 3, and the inequality

$$\mathbb{P}A \geq \sum_i \mathbb{P}(A \mid A_{ni})\mathbb{P}A_{ni} \quad \text{for measurable } A,$$

ensure that μ_n is nonnegative. For each t in $[0, 1]$ and each x in $\mathcal{X}$ define a probability measure $K_n(t, x, \cdot)$ on $\mathcal{A}$ by

$$(9.6) \quad K_n(t, x, \cdot) = \begin{cases} \mathbb{P}(\cdot \mid A_{ni}) & \text{if } t \leq 1 - 2^{-\gamma(n)} \quad \text{and} \quad x \in B_i \in \pi(\gamma(n)), \\ \mu_n(\cdot) & \text{if } t > 1 - 2^{-\gamma(n)}. \end{cases}$$

The kernel K_n will provide a randomization mechanism for generating $\mathbb{P}$, starting from a t distributed uniformly on $[0, 1]$ independently of an x distributed according

to P . Specifically, if λ denotes Lebesgue measure on $[0, 1]$, then

$$\mathbb{P}A = \iint K_n(t, x, A) \lambda(dt) P(dx)$$

for each A in $\mathcal{A}$.

Step 5. Define $\tilde{\Omega}$ as the product space $[0, 1] \otimes \mathcal{X} \otimes \Omega^{\mathbb{N}}$, where $\mathbb{N} = \{1, 2, \dots\}$. Equip it with its product σ -field. For t in $[0, 1]$ and x in $\mathcal{X}$ define the probability measure $K(t, x, \cdot)$ on the product σ -field of $\Omega^{\mathbb{N}}$ as a product

$$K(t, x, \cdot) = \prod_n K_n(t, x, \cdot).$$

With λ denoting Lebesgue measure on $[0, 1]$, define $\tilde{\mathbb{P}}$ on the product σ -field of $\tilde{\Omega}$ by

$$\tilde{\mathbb{P}}(\cdot) = \lambda \otimes P \otimes K.$$

That is, for $I \in \mathcal{B}[0, 1]$ and $B \in \mathcal{B}(\mathcal{X})$ and C in the product σ -field of $\Omega^{\mathbb{N}}$,

$$\tilde{\mathbb{P}}(I \otimes B \otimes C) = \iint \{t \in I, x \in B\} K(t, x, C) \lambda(dt) P(dx).$$

Some measurability details must be checked to ensure that $\tilde{\mathbb{P}}$ is well defined.

Step 6. Define maps ϕ_n (from $\tilde{\Omega}$ into Ω), and $\tilde{X}$ (from $\tilde{\Omega}$ into $\mathcal{X}$), and $\tilde{X}_n$ (from $\tilde{\Omega}$ into $\mathcal{X}$) by

$$\begin{aligned} \phi_n(t, x, \omega_1, \omega_2, \dots) &= \omega_n, \\ \tilde{X}(t, x, \omega_1, \omega_2, \dots) &= x, \\ \tilde{X}_n(t, x, \omega_1, \omega_2, \dots) &= X_n(\omega_n). \end{aligned}$$

Use the representations from Steps 4 and 5 to verify that $\tilde{\mathbb{P}}\phi_n^{-1} = \mathbb{P}$ and $\tilde{\mathbb{P}}\tilde{X}^{-1} = P$.

Step 7. Temporarily fix a value of k . Let B_0 be the member of $\pi(k)$ that might have diameter greater than $1/k$. Define the subset $\tilde{\Omega}_k$ of $\tilde{\Omega}$ to consist of all those $\tilde{\omega}$ for which $t \leq 1 - 2^{-k}$ and $x \in B_i$ for some $i \geq 1$ and $\omega_n \in A_{ni}$ for that same i , for all n in the range $n(k) \leq n < n(k+1)$. By the construction of $\pi(k)$,

$$d(\tilde{X}_n(\tilde{\omega}), \tilde{X}(\tilde{\omega})) \leq 1/k \quad \text{for } n(k) \leq n < n(k+1) \text{ and } \tilde{\omega} \text{ in } \tilde{\Omega}_k.$$

If $n(k) \leq n < n(k+1)$, define $\delta_n(\tilde{\omega})$ to equal $1/k$ on $\tilde{\Omega}_k$ and ∞ elsewhere. By the Borel-Cantelli lemma, the δ_n sequence converges to zero almost surely, because the construction of $\tilde{\mathbb{P}}$ ensures $\tilde{\mathbb{P}}\tilde{\Omega}_k^c \leq 2(1/2)^k$.

Step 8. Prove that each ϕ_n is perfect. Let H be a bounded function on Ω , and let g be a bounded, measurable function on $\tilde{\Omega}$ for which

$$g(\tilde{\omega}) \geq H(\omega_n) \quad \text{for all } \tilde{\omega} = (t, x, \omega_1, \omega_2, \dots).$$

Establish (9.3) by finding a measurable function g^* on Ω for which $\tilde{\mathbb{P}}g \geq \mathbb{P}g^*$ and

$$g^*(\omega_n) \geq H(\omega_n) \quad \text{for all } \omega_n.$$

For fixed t , x , and ω_n , let $g_n(t, x, \omega_n)$ be the measurable function obtained by integrating g with respect to the product of all those $K_i(t, x, \cdot)$ with $i \neq n$. Then $\tilde{\mathbb{P}}g_n = \tilde{\mathbb{P}}g$ and

$$g_n(t, x, \omega_n) \geq H(\omega_n) \quad \text{for all } t, x, \omega_n.$$

Now comes the crucial argument. The kernel K_n depends on (t, x) in a very simple way. There is a finite partition of $[0, 1] \otimes \mathcal{X}$ into measurable sets D_α , and there are probability measures m_α on Ω , such that

$$K_n(t, x, \cdot) = \sum_{\alpha} \{(t, x) \in D_\alpha\} m_\alpha(\cdot).$$

Define g^* by

$$g^*(\omega_n) = \min_{\alpha} P \otimes \lambda(g_n(t, x, \omega_n) \mid (t, x) \in D_\alpha),$$

with the minimum running over those α for which $P \otimes \lambda(D_\alpha) > 0$. Finiteness of the $\{D_\alpha\}$ partition ensures that g^* is measurable. It is easy to check that it also satisfies the desired inequalities.

REMARKS. Typically measurability is not a major concern in specific problems. Nevertheless, it is highly desirable that a general theory for convergence in distribution, free from unnatural measurability constraints, should exist. Unfortunately, Hoffmann-Jørgensen's (1984) theory was presented in a manuscript for a book that has not yet been published. However, detailed explanations of some parts of the theory have appeared in the papers of Andersen (1985a, 1985b) and Andersen and Dobrić (1987, 1988).

Many measure theoretic details have been omitted from the outline of the proof of the Representation Theorem, but otherwise it is quite similar to the version in Chapter IV of Pollard (1984), which was based on Dudley's (1968) original paper. Dudley (1985) discussed the notion of a perfect map in some detail, and also showed how slippery a concept almost sure convergence can be for nonmeasurable random processes.

SECTION 10

Functional Central Limit Theorems

When does the standardized partial-sum processes converge in distribution, in the sense of the previous section, to a Gaussian process with nice sample paths? This section will establish a workable sufficient condition.

Part of the condition will imply finiteness (almost everywhere) of the envelope functions, which will mean that $S_n(\omega, \cdot)$ is a bounded function on T , for almost all ω . Ignoring negligible sets of ω , we may therefore treat S_n as a random element of the space $B(T)$ of all bounded, real-valued functions on T . The natural metric for this space is given by the uniform distance,

$$d(x, y) = \sup_t |x(t) - y(t)|.$$

One should take care not to confuse d with any metric, or pseudometric, ρ defined on T . Usually such a ρ will have something to do with the covariance structure of the partial-sum processes. The interesting limit distributions will be Gaussian processes that concentrate on the set

$$U_\rho(T) = \{x \in B(T) : x \text{ is uniformly } \rho \text{ continuous}\}.$$

Under the uniform metric d , the space $U_\rho(T)$ is separable if and only if T is totally bounded under ρ . [Notice that total boundedness excludes examples such as the real line under its usual metric.] In the separable case, a Borel probability measure P on $U_\rho(T)$ is uniquely determined by its finite dimensional projections,

$$P(B \mid t_1, \dots, t_k) = P\{x \in U_\rho(T) : (x(t_1), \dots, x(t_k)) \in B\},$$

with $\{t_1, \dots, t_k\}$ ranging over all finite subsets of T and B ranging over all Borel sets in $\mathbb{R}^k$, for $k = 1, 2, \dots$.

Let us first consider a general sequence of stochastic processes indexed by T ,

$$\{X_n(\omega, t) : t \in T\} \quad \text{for } n = 1, 2, \dots,$$

and then specialize to the case where X_n is a properly standardized partial-sum process. Let us assume that the finite dimensional projections of X_n converge in distribution. That is, for each finite subset $\{t_1, \dots, t_k\}$ of T there is a Borel probability measure $P(\cdot \mid t_1, \dots, t_k)$ on $\mathbb{R}^k$ such that

$$(10.1) \quad (X_n(\cdot, t_1), \dots, X_n(\cdot, t_k)) \rightsquigarrow P(\cdot \mid t_1, \dots, t_k).$$

Usually classical central limit theorems will suggest the standardizations needed to ensure such finite dimensional convergence.

(10.2) THEOREM. *Let $\{X_n(\cdot, t) : t \in T\}$ be stochastic processes indexed by a totally bounded pseudometric space (T, ρ) . Suppose:*

- (i) *the finite dimensional distributions converge, as in (10.1);*
- (ii) *for each $\epsilon > 0$ and $\eta > 0$ there is a $\delta > 0$ such that*

$$\limsup \mathbb{P}^* \left\{ \sup_{\rho(s,t) < \delta} |X_n(\omega, s) - X_n(\omega, t)| > \eta \right\} < \epsilon.$$

Then there exists a Borel measure P concentrated on $U_\rho(T)$, with finite dimensional projections given by the distributions $P(\cdot \mid t_1, \dots, t_k)$ from (10.1), such that X_n converges in distribution to P .

Conversely, if S_n converges in distribution to a Borel measure P on $U_\rho(T)$ then conditions (i) and (ii) are satisfied.

SKETCH OF A PROOF. The converse part of the theorem is a simple exercise involving almost sure representations.

For the direct part, first establish existence of the measure P concentrating on $U_\rho(T)$. Let $T(\infty) = \{t_1, t_2, \dots\}$ be a countable dense subset of T . Define $T(k) = \{t_1, \dots, t_k\}$. The Kolmogorov extension theorem lets us build a measure P on the product σ -field of $\mathbb{R}^{T(\infty)}$ with the finite dimensional distributions from the right-hand side of (10.1). By passing to the limit in (ii) we get

$$P \left\{ x \in \mathbb{R}^{T(\infty)} : \max_{\substack{\rho(s,t) < \delta \\ s,t \in T(k)}} |x(s) - x(t)| \geq \eta \right\} \leq \epsilon \quad \text{for every } k.$$

Letting $k \rightarrow \infty$, then casting out various sequences of negligible sets, we find that P concentrates on the set $U_\rho(T(\infty))$ of all uniformly continuous functions on $T(\infty)$. Each function in $U_\rho(T(\infty))$ has a unique extension to a function in $U_\rho(T)$; the extension carries P up to the sought-after Borel measure on $U_\rho(T)$.

To complete the proof let us construct a new probability space $(\tilde{\Omega}, \tilde{\mathcal{A}}, \tilde{\mathbb{P}})$ that supports perfect maps ϕ_n into Ω , such that $X_n \circ \phi_n$ converges to an $\tilde{X}$ with distribution P , in the strengthened almost sure sense of the Representation Theorem from the previous section. This is not the circular argument that it might appear; we do not need to assume the convergence $X_n \rightsquigarrow P$ in order to adapt some of the ideas from the proof of that theorem. Indeed, we can break into the proof between its second and third steps by establishing directly that $\liminf \mathbb{P}_* \{X_n \in B\} \geq PB$ for every B that is a finite intersection of closed balls in $B(T)$ with centers in $U_\rho(T)$ and zero P measure on their boundaries.

Such a set B has a simple form; it is defined by a pair of functions g, h in $U_\rho(T)$:

$$B = \{x \in B(T) : g(t) \leq x(t) \leq h(t) \text{ for all } t\}.$$

It has zero P measure on its boundary. For $\eta > 0$ define

$$B_\eta = \{x \in B(T) : g(t) + \eta < x(t) < h(t) - \eta \text{ for all } t\}.$$

As $\eta \rightarrow 0$, the sets B_η expand up to the interior of B . The fact that P puts zero measure on the boundary of B lets us choose η so that $PB_\eta \geq PB - \epsilon$. This inequality gives us room to approximate the paths of the X_n processes from their values on a finite subset of T .

Fix an $\epsilon > 0$. The fact that P concentrates on $U_\rho(T)$ lets us choose a $\delta > 0$ so that the set

$$F = \left\{x \in B(T) : \sup_{\rho(s,t) < \delta} |x(s) - x(t)| \leq \eta/2\right\}$$

has P measure at least $1 - \epsilon$. Condition (ii) of the theorem lets us assume that δ is small enough to ensure $\limsup \mathbb{P}^*\{X_n \in F^c\} < \epsilon$. We may also assume that both g and h belong to F , because both are uniformly continuous.

Now let $T(k) = \{t_1, \dots, t_k\}$ be a finite set that approximates within a distance δ to every point of T . For a function x in F and a t with $\rho(t, t_i) < \delta$, if $x(t_i) < h(t_i) - \eta$ then $x(t) \leq x(t_i) + \eta/2 < h(t_i) - \eta/2$. The upper bound is less than $h(t)$, because $h \in F$. A similar argument with g would give a lower bound. It follows that the set

$$\{X_n \in F : g(t_i) + \eta < X_n(\cdot, t_i) < h(t_i) - \eta \text{ for } t_i \in T(k)\}$$

is contained within $\{X_n \in B\}$, and hence

$$\mathbb{P}_*\{X_n \in B\} \geq \mathbb{P}\{g + \eta < X_n < h - \eta \text{ on } T(k)\} - \mathbb{P}^*\{X_n \in F^c\}.$$

The first term on the right-hand side may be reexpressed as

$$\mathbb{P}\{(X_n(\cdot, t_1), \dots, X_n(\cdot, t_k)) \in G\},$$

where G is the open subset of $\mathbb{R}^k$ defined by the inequalities

$$g(t_i) + \eta < x_i < h(t_i) - \eta \quad \text{for } i = 1, \dots, k.$$

From assumption (i), the $\liminf$ of the last probability is greater than

$$P\{x \in U_\rho(T) : (x(t_1), \dots, x(t_k)) \in G\} \geq PB_\eta.$$

It follows that

$$\liminf \mathbb{P}_*\{X_n \in B\} \geq PB - 2\epsilon \quad \text{for each } \epsilon > 0.$$

By copying Steps 3 through 8 in the proof of the Representation Theorem we could now complete the construction of versions of the X_n that converge in the strong sense to an $\tilde{X}$ with distribution P . The assertion of the theorem would then follow easily. $\square$

Condition (ii) of Theorem 10.2 is sometimes called *stochastic equicontinuity* or, less precisely, uniform tightness. It is equivalent to the requirement: for every sequence $\{r_n\}$ of real numbers converging to zero,

$$(10.3) \quad \sup\{|X_n(s) - X_n(t)| : \rho(s, t) < r_n\} \rightarrow 0 \quad \text{in probability.}$$

It also implies (and actually is implied by) that for every sequence of random elements $\{\sigma_n\}, \{\tau_n\}$ of T with $\rho(\sigma_n, \tau_n) \rightarrow 0$ in probability,

$$(10.4) \quad X_n(\sigma_n) - X_n(\tau_n) \rightarrow 0 \quad \text{in probability.}$$

One has only to choose r_n converging to zero so slowly that $\mathbb{P}^*\{\rho(\sigma_n, \tau_n) \geq r_n\} \rightarrow 0$ to establish the implication. Notice that (10.4) is much stronger than the corresponding assertion for deterministic sequences $\{\sigma_n\}, \{\tau_n\}$ with $\rho(s_n, t_n) \rightarrow 0$. Verification of the weaker assertion would typically involve little more than an application of Tchebychev's inequality, whereas (10.4) corresponds to a much more powerful maximal inequality.

Let us now specialize Theorem 10.2 to random processes constructed from a triangular array $\{f_{ni}(\omega, t) : t \in T, 1 \leq i \leq k_n, n = 1, 2, \dots\}$, with the $\{f_{ni}\}$ independent within each row. Define

$$X_n(\omega, t) = \sum_{i \leq k_n} \left(f_{ni}(\omega, t) - \mathbb{P}f_{ni}(\cdot, t) \right),$$

$$\rho_n(s, t) = \left(\sum_{i \leq k_n} \mathbb{P}|f_{ni}(\cdot, s) - f_{ni}(\cdot, t)|^2 \right)^{1/2}.$$

The double subscripting allows us to absorb into the notation the various standardizing constants needed to ensure convergence of finite dimensional distributions. If we also arrange to have

$$(10.5) \quad \rho(s, t) = \lim_{n \rightarrow \infty} \rho_n(s, t)$$

well defined for each pair s, t in T , then such a ρ will be an appropriate choice for the pseudometric on T . In the frequently occurring case where $f_{ni}(\omega, t) = f_i(\omega, t)/\sqrt{n}$, with the $\{f_i\}$ independent and identically distributed, we have $\rho(s, t) = \rho_n(s, t)$, and condition (v) of the next theorem is trivially satisfied.

(10.6) FUNCTIONAL CENTRAL LIMIT THEOREM. *Suppose the processes from the triangular array $\{f_{ni}(\omega, t)\}$ are independent within rows and satisfy:*

- (i) *the $\{f_{ni}\}$ are manageable, in the sense of Definition 7.9;*
- (ii) *$H(s, t) = \lim_{n \rightarrow \infty} \mathbb{P}X_n(s)X_n(t)$ exists for every s, t in T ;*
- (iii) *$\limsup \sum_i \mathbb{P}F_{ni}^2 < \infty$;*
- (iv) *$\sum_i \mathbb{P}F_{ni}^2\{F_{ni} > \epsilon\} \rightarrow 0$ for each $\epsilon > 0$;*
- (v) *the limit $\rho(\cdot, \cdot)$ is well defined by (10.5) and, for all deterministic sequences $\{s_n\}$ and $\{t_n\}$, if $\rho(s_n, t_n) \rightarrow 0$ then $\rho_n(s_n, t_n) \rightarrow 0$.*

Then

- (a) *T is totally bounded under the ρ pseudometric;*
- (b) *the finite dimensional distributions of X_n have Gaussian limits, with zero means and covariances given by H , which uniquely determine a Gaussian distribution P concentrated on $U_\rho(T)$;*
- (c) *X_n converges in distribution to P .*

PROOF. Conditions (ii) and (iv) imply (Lindeberg central limit theorem) that the finite dimensional distributions have the stated Gaussian limits.

The stochastic equicontinuity requirement of Theorem 10.2 will be established largely by means of maximal inequalities implied by manageability. Recall from Section 7 that manageability of the $\{f_{ni}\}$ means that there exists a deterministic function λ with $\sqrt{\log \lambda}$ integrable and

$$(10.7) \quad D_2(x|\alpha \odot \mathbf{F}_n|_2, \alpha \odot \mathcal{F}_{n\omega}) \leq \lambda(x) \quad \text{for } 0 < x \leq 1, \text{ all } \omega, \text{ all } \alpha, \text{ all } n.$$

For manageable arrays of processes we have the moment bounds, for $1 \leq p < \infty$,

$$(10.8) \quad \mathbb{P} \sup_t \left| \sum_i f_{ni}(\omega, t) - \mathbb{P} f_{ni}(\cdot, t) \right|^p \leq \mathbb{P} |\mathbf{F}_n|_2^p \Lambda_p(\delta_n / |\mathbf{F}_n|_2),$$

where $\delta_n^2 = \sup_t \sum_i f_{ni}(\omega, t)^2$ and Λ_p is a continuous, increasing function that depends only on λ and p , with $\Lambda_p(0) = 0$ and $\Lambda_p(1) < \infty$.

The presence of the rescaling vector α in (10.7) will allow us to take advantage of the Lindeberg condition (iv) without destroying the bound. Because (iv) holds for each fixed $\epsilon > 0$, it also holds when ϵ is replaced by a sequence $\{\epsilon_n\}$ converging to zero slowly enough:

$$\sum_i \mathbb{P} F_{ni}^2 \{F_{ni} > \epsilon_n\} \rightarrow 0.$$

We can replace f_{ni} by $f_{ni}\{F_{ni} \leq \epsilon_n\}$ and F_{ni} by $F_{ni}\{F_{ni} \leq \epsilon_n\}$ without disturbing inequality (10.7); the indicator function $\{F_{ni} \leq \epsilon_n\}$ is absorbed into the weight α_i . The same truncation has no bad effect on the other four assumptions of the theorem. We therefore lose no generality by strengthening (iv) to:

$$(iv)' \quad F_{ni}(\omega) \leq \epsilon_n \quad \text{for all } n, \text{ all } i, \text{ all } \omega.$$

Henceforth assume this inequality holds.

The idea will be to apply a maximal inequality analogous to (10.8) to the processes

$$h_{ni}(\omega, s, t) = f_{ni}(\omega, s) - f_{ni}(\omega, t),$$

at least for those pairs s, t with $\rho(s, t) < r_n$, with the aim of establishing stochastic equicontinuity in the form (10.3). The maximal inequality will involve the random variable

$$\theta_n(\omega) = \sup\{|\mathbf{h}_n(\omega, s, t)|_2 : \rho(s, t) < r_n\}.$$

We will use manageability to translate the convergence $r_n \rightarrow 0$ into the conclusion that $\theta_n \rightarrow 0$ in probability.

From the stability results for packing numbers in Section 5, the doubly indexed processes $\{h_{ni}(\omega, s, t)\}$ are also manageable, for the envelopes $H_{ni} = 2F_{ni}$, with capacity bound $\lambda(x/2)^2$. And the processes $\{h_{ni}(\omega, s, t)^2\}$ are manageable for the envelopes $\{H_{ni}^2\}$, by virtue of inequality (5.2) for packing numbers of pointwise products. The analogue of (10.8) therefore holds for the $\{h_{ni}^2\}$ processes, with envelopes $\{H_{ni}^2\}$ and the Λ_p function increased by a constant multiple. In particular,

there is a constant C such that

$$\begin{aligned} \mathbb{P} \sup_{s,t} \left| \sum_i h_{ni}(\omega, s, t)^2 - \mathbb{P} h_{ni}(\cdot, s, t)^2 \right|^2 &\leq C \mathbb{P} \sum_i F_{ni}^4 \\ &\leq C \sum_i \epsilon_n^2 \mathbb{P} F_{ni}^2 \\ &\rightarrow 0. \end{aligned}$$

Consequently,

$$(10.9) \quad \sup_{s,t} \left| |\mathbf{h}_n(\omega, s, t)|_2^2 - \rho_n(s, t)^2 \right| \rightarrow 0 \quad \text{in probability.}$$

The second part of assumption (v) implies that

$$\sup_{s,t} \{ \rho_n(s, t) : \rho(s, t) < r_n \} \rightarrow 0.$$

Together these two uniformity results give $\theta_n \rightarrow 0$ in probability.

The convergence (10.9) also establishes total boundedness of T under the ρ pseudometric, with plenty to spare. First note that assumption (iii) and the fact that $\sum_i \mathbb{P} F_{ni}^4 \rightarrow 0$ together imply that $|\mathbf{F}_n|_2$ is stochastically bounded: for some constant K there is probability close to one for all n that $|\mathbf{F}_n|_2 \leq K$. Now suppose $\{t_1, \dots, t_m\}$ is a set of points with $\rho(t_i, t_j) > \epsilon K$ for $i \neq j$. By definition of ρ and by virtue of (10.9), with probability tending to one,

$$|\mathbf{f}_n(\omega, t_i) - \mathbf{f}_n(\omega, t_j)|_2 > \epsilon K \quad \text{for } i \neq j.$$

Eventually there will be an ω (in fact, a whole set of them, with probability close to one) for which $m \leq D_2(\epsilon |\mathbf{F}_n|_2, \mathcal{F}_{n\omega})$. It follows from (10.7) that $m \leq \lambda(\epsilon)$. That is, λ is also a bound on the packing numbers of T under the ρ pseudometric.

To complete the proof of stochastic equicontinuity, invoke the analogue of (10.8) with $p = 1$ for the processes of differences $h_{ni}(\omega, s, t)$ with $\rho(s, t) < r_n$. By manageability, there is a continuous, increasing function $\Gamma(\cdot)$ with $\Gamma(0) = 0$ such that

$$\mathbb{P} \sup \{ |X_n(s) - X_n(t)| : \rho(s, t) < r_n \} \leq \mathbb{P} |\mathbf{F}_n|_2 \Gamma(\theta_n / |\mathbf{F}_n|_2).$$

For a fixed $\epsilon > 0$, split the right-hand side according to whether $|\mathbf{F}_n|_2 > \epsilon$ or not, to get the upper bound

$$\epsilon \Gamma(1) + \mathbb{P} |\mathbf{F}_n|_2 \Gamma \left(1 \wedge \frac{\theta_n}{2\epsilon} \right).$$

The Cauchy-Schwarz inequality bounds the second contribution by

$$\left[\mathbb{P} |\mathbf{F}_n|_2^2 \mathbb{P} \Gamma^2 \left(1 \wedge \frac{\theta_n}{2\epsilon} \right) \right]^{1/2}.$$

Assumption (iii) keeps $\mathbb{P} |\mathbf{F}_n|_2^2$ bounded; the convergence in probability of θ_n to zero sends the second factor to zero. Stochastic equicontinuity of $\{X_n\}$ follows. $\square$

REMARKS. The original functional central limit theorem for empirical distribution functions is due to Donsker (1952). Dudley (1978) extended the result to

empirical processes indexed by general classes of sets. He used the term *Donsker class* to describe those classes of sets (and later, also those classes of functions—see Dudley (1987), for example) for which a functional central limit theorem holds.

The literature contains many examples of such limit theorems for empirical processes and partial-sum processes, mostly for identically distributed summands. Some of the best recent examples may be found in the papers of Dudley (1984), Giné and Zinn (1984), Alexander and Pyke (1986), Ossiander (1987), Alexander (1987a, 1987b), Talagrand (1987), and Andersen and Dobrić (1987, 1988). My Theorem 10.6 extends a central limit theorem of Kim and Pollard (1990), refining the earlier result from Pollard (1982). It could also be deduced from the theorems of Alexander (1987b). The assumption of manageability could be relaxed.

Theorem 10.2 is based on Theorem 5.2 of Dudley (1985), which extends a line of results going back at least to Dudley (1966). See also Dudley (1984). My method of proof is different, although similar in spirit to the methods of Skorohod (1956).

SECTION 11

Least Absolute Deviations Estimators for Censored Regression

Suppose random variables $y_1, y_2, \dots$ are generated by a regression $y_i = x_i' \theta_0 + u_i$, with θ_0 an unknown d -dimensional vector of parameters, $\{x_i\}$ a sequence of observed vectors, and $\{u_i\}$ unobserved errors. The method of least absolute deviations would estimate θ_0 by the θ that minimized the convex function

$$\sum_{i \leq n} |y_i - x_i' \theta|.$$

Convexity in θ makes the asymptotic analysis not too difficult (Pollard 1990). Much more challenging is a related problem, analyzed by Powell (1984), in which the value of y_i is observed only if $y_i \geq 0$ and otherwise only the information that $y_i < 0$ is available. That is, only y_i^+ is observed. In the econometrics literature this is called a Tobit model (at least when the $\{u_i\}$ are independent normals).

Powell proposed an interesting variation on the least absolute deviations estimation; he studied the $\hat{\theta}_n$ that minimizes

$$\sum_{i \leq n} |y_i^+ - (x_i' \theta)^+|$$

over a subset Θ of $\mathbb{R}^d$. This function is not convex in θ ; analysis of $\hat{\theta}_n$ is quite difficult. However, by extending a technique due to Huber (1967), Powell was able to give conditions under which $\sqrt{n}(\hat{\theta}_n - \theta_0)$ has an asymptotic normal distribution.

With the help of the maximal inequalities developed in these notes, we can relax Powell's assumptions and simplify the analysis a little. The strategy will be to develop a uniformly good quadratic approximation to the criterion function, then show that $\hat{\theta}_n$ comes close to maximizing the approximation. Powell's consistency argument was based on the same idea, but for asymptotic normality he sought

an approximate zero for a vector of partial derivatives, a method that is slightly complicated by the lack of smoothness of the criterion function.

Assumptions. Let us assume that the $\{x_i\}$ vectors are deterministic. Results for random $\{x_i\}$ could also be obtained by a conditioning argument. The following assumptions would be satisfied by a typical realization of independent, identically distributed random vectors $\{X_i\}$ with finite second moments and $\mathbb{P}\{X_i'\theta_0 = 0\} = 0$ and $\mathbb{P}X_iX_i'\{X_i'\theta_0 > 0\}$ nonsingular. The assumptions on the errors $\{u_i\}$ are the usual ones for least absolute deviations estimation. They could be relaxed slightly at the cost of increased notational complexity.

- (i) The $\{u_i\}$ are independent, identically distributed random variables each having zero median and a continuous, strictly positive density $p(\cdot)$ near zero.
- (ii) For each $\epsilon > 0$ there is a finite K such that

$$\frac{1}{n} \sum_{i \leq n} |x_i|^2 \{|x_i| > K\} < \epsilon \quad \text{for all } n \text{ large enough.}$$

- (iii) For each $\epsilon > 0$ there is a $\delta > 0$ such that

$$\frac{1}{n} \sum_{i \leq n} |x_i|^2 \{|x_i'\theta_0| \leq \delta\} < \epsilon \quad \text{for all } n \text{ large enough.}$$

- (iv) The sequence of smallest eigenvalues of the matrices

$$\frac{1}{n} \sum_{i \leq n} x_i x_i' \{x_i'\theta_0 > 0\}$$

is bounded away from zero, for n large enough.

Powell required slightly more smoothness for $p(\cdot)$, and a more awkward moment condition analogous to (iii), in order to fit his analysis into the framework of Huber's method.

(11.1) THEOREM. Suppose θ_0 is an interior point of a Θ , a bounded subset of $\mathbb{R}^d$. Then, under assumptions (i) to (iv),

$$2p(0)\sqrt{n}V_n(\hat{\theta}_n - \theta_0) \rightsquigarrow N(0, I_d),$$

where V_n is the positive definite square root of the matrix from assumption (iv).

The proof of this result is quite a challenge. Let us begin with some heuristic arguments, which will develop notation and focus attention on the main technical difficulties.

Heuristics. The assumptions (ii), (iii), and (iv) control how much influence any single x_i can have over V_n . If $x_i'\theta_0 < 0$ then, for θ near θ_0 , the term $|y_i^+ - (x_i'\theta)^+|$ reduces to y_i^+ ; it should not greatly affect the local minimization; it should not have an effect on the limiting distribution of $\hat{\theta}_n$; it should not contribute to V_n . Assumption (iv) captures this idea. Assumption (ii) prevents a single very large

$|x_i|$ from dominating V_n ; the least absolute deviations criterion prevents it from having a dominating influence on the minimization. The x_i with $x'_i\theta_0 \approx 0$ are the most troublesome, because their contribution to the criterion function is nonlinear in θ , even when θ is close to θ_0 ; assumption (iii) will allow us to ignore the combined effect of all such troublesome x_i .

The assumption of boundedness for the parameter set Θ is aesthetically irksome, even if it does have little practical significance. I would be pleased to learn how to dispose of it.

As a notational convenience, reparametrize by putting $t = V_n(\theta - \theta_0)$. Then define $z_{ni} = V_n^{-1}x_i$ and $x'_i\theta_0 = \mu_i$. Define

$$f_{ni}(\omega, t) = |y_i^+ - (\mu_i + z'_{ni}t)^+| - |y_i^+ - \mu_i^+|.$$

The centering ensures that

$$|f_{ni}(\omega, t)| \leq |z'_{ni}t|,$$

and hence $f_{ni}(\cdot, t)$ has a finite expectation for each t . The centering does not affect the minimization; the standardized estimator $\hat{t}_n = V_n(\hat{\theta}_n - \theta_0)$ minimizes the process

$$G_n(\omega, t) = \frac{1}{n} \sum_{i \leq n} f_{ni}(\omega, t).$$

Assumptions (ii) and (iv) imply existence of positive constants C' and C'' for which, when n is large enough,

$$C'|x_i| \geq |z_{ni}| \geq C''|x_i| \quad \text{for } i \leq n,$$

which lets us translate the assumptions on the $\{x_i\}$ into:

(ii)* For each $\epsilon > 0$ there is a finite K such that

$$\frac{1}{n} \sum_{i \leq n} |z_{ni}|^2 \{|z_{ni}| > K\} < \epsilon \quad \text{for all } n \text{ large enough.}$$

(iii)* For each $\epsilon > 0$ there is a $\delta > 0$ such that

$$\frac{1}{n} \sum_{i \leq n} |z_{ni}|^2 \{|\mu_i| \leq \delta\} < \epsilon \quad \text{for all } n \text{ large enough.}$$

(iv)*

$$\frac{1}{n} \sum_{i \leq n} z_{ni} z'_{ni} \{\mu_i > 0\} = I_d \quad \text{for all } n \text{ large enough.}$$

For convenience, let us ignore from now on the finitely many n excluded by these three conditions.

As a first approximation, $G_n(\cdot, t)$ should be close to its expected value, $\Gamma_n(t)$. If we define

$$H_i(s) = \mathbb{P}(|y_i^+ - (\mu_i + s)^+| - |y_i^+ - \mu_i^+|) \quad \text{for } s \in \mathbb{R},$$

then

$$\Gamma_n(t) = \mathbb{P}G_n(\cdot, t) = \frac{1}{n} \sum_{i \leq n} H_i(z'_{ni}t).$$

Each H_i is expressible in terms of the function

$$M(s) = \mathbb{P}(|u_1 - s| - |u_1|).$$

Indeed, by separate consideration of the contributions from the sets $\{u_i < -\mu_i\}$ and $\{u_i \geq -\mu_i\}$ one can show

$$(11.2) \quad H_i(s) = \begin{cases} M(s) & \text{if } \mu_i \geq 0 \text{ and } s > -\mu_i, \\ M(-\mu_i) & \text{if } \mu_i \geq 0 \text{ and } s \leq -\mu_i, \\ M(s) - M(-\mu_i) & \text{if } \mu_i < 0 \text{ and } s > -\mu_i, \\ 0 & \text{if } \mu_i < 0 \text{ and } s \leq -\mu_i. \end{cases}$$

From assumption (i), the expected value $M(s)$ is increasing in $|s|$. The function M has a unique minimum at the origin, and

$$M(s) = p(0)s^2 + o(s^2) \quad \text{near the origin.}$$

Moreover, there is a constant C such that

$$(11.3) \quad H_i(s) \leq M(s) \leq C s^2 \quad \text{for all } s.$$

At least when t is small, we should be able to ignore those H_i with $\mu_i < 0$ or $\mu_i \approx 0$, to get, via (iv)*,

$$\Gamma_n(t) \approx \frac{1}{n} \sum_{i \leq n} p(0) |z'_{ni} t|^2 \{\mu_i > 0\} = p(0) |t|^2.$$

If $G_n(\cdot, t)$ lies uniformly close to its expectation, the $\hat{t}_n$ that minimizes G_n should be drawn close to zero, where the approximation to Γ_n is minimized.

With the help of a maximal inequality for $G_n - \Gamma_n$, we will even be able to force $\hat{t}_n$ into a $O_p(1/\sqrt{n})$ neighborhood of the origin. To learn more about $\sqrt{n} \hat{t}_n$ we will then need a better approximation to G_n , obtained by a sort of Taylor expansion for f_{ni} .

The random function $f_{ni}(\omega, \cdot)$ has a derivative at $t = 0$ except perhaps when $\mu_i = 0$ or $u_i(\omega) = 0$. If we ignore these cases, straightforward differentiation suggests we treat

$$\Delta_{ni}(\omega) = \left(\{u_i(\omega) < 0\} - \{u_i(\omega) \geq 0\} \right) \{\mu_i > 0\} z_{ni}$$

as the derivative of f_{ni} at $t = 0$. The difference

$$R_{ni}(\omega, t) = f_{ni}(\omega, t) - \Delta_{ni}(\omega)' t$$

should behave like the remainder in a Taylor expansion. By direct calculation for the various pairings of inequalities, one can verify that

$$|R_{ni}(\omega, t)| \leq 2 |z'_{ni} t| \left(\{|\mu_i| \leq |z'_{ni} t|\} + \{|u_i(\omega)| \leq |z'_{ni} t|\} \right).$$

The two indicator functions vanish when $|z'_{ni} t|$ is smaller than both $|\mu_i|$ and $|u_i(\omega)|$, which should happen with large probability when $|t|$ is small.

Write W_n for $1/\sqrt{n}$ times the sum of the Δ_{ni} . By the Lindeberg central limit

theorem it has an asymptotic $N(0, I_d)$ distribution:

$$\begin{aligned}\mathbb{P}\Delta_{ni} &= \left(\mathbb{P}\{u_i < 0\} - \mathbb{P}\{u_i \geq 0\}\right)\{\mu_i > 0\}z_{ni} = 0; \\ \frac{1}{n} \sum_{i \leq n} \mathbb{P}\Delta_{ni}\Delta'_{ni} &= \{\mu_i > 0\}z_{ni}z'_{ni} = I_d; \\ \frac{1}{n} \sum_{i \leq n} \mathbb{P}|\Delta_{ni}|^2 \{|\Delta_{ni}| > \epsilon\sqrt{n}\} &= \frac{1}{n} \sum_{i \leq n} \{\mu_i > 0\}|z_{ni}|^2 \{|z_{ni}| > \epsilon\sqrt{n}\} \rightarrow 0.\end{aligned}$$

Ignoring the contributions from the $\{R_{ni}\}$, we get an improved approximation to G_n :

$$\begin{aligned}G_n(\omega, t) &= \frac{1}{\sqrt{n}}W'_n t + \Gamma_n(t) + \frac{1}{n} \sum_{i \leq n} \left(R_{ni}(\omega, t) - \mathbb{P}R_{ni}(\cdot, t)\right) \\ &\simeq \frac{1}{\sqrt{n}}W'_n t + p(0)|t|^2 \quad \text{for small } |t|.\end{aligned}$$

The random vector $\hat{t}_n$ should be close to the vector $-W_n/2\sqrt{n}p(0)$ that minimizes the approximating quadratic, which leads us to the limit distribution asserted by Theorem 11.1.

Now let us make these arguments precise. The technical challenge in the proof will come from the two approximations to G_n . To obtain the necessary uniform bounds on the errors we will make use of maximal inequalities for processes with finite pseudodimension.

Behavior of Γ_n . From (11.2), it follows that $\Gamma_n(t) = \mathbb{P}G_n(\cdot, t)$ is an increasing function of $|t|$, taking its minimum value uniquely at $t = 0$. Given $\epsilon > 0$, choose K and δ according to (ii)* and (iii)*; then put $r = \delta/K$. From (11.3), the terms where $|\mu_i| \leq \delta$ or $|z_{ni}| > K$ contribute at most $2C\epsilon|t|^2$ to $\Gamma_n(t)$. For the other terms we have $|z'_{ni}t| \leq \delta$ if $|t| \leq r$. If $\mu_i \leq -\delta$ this makes $H_i(z'_{ni}t)$ zero. Within an error of $2C\epsilon|t|^2$, the expectation equals

$$\sum_i \{\mu_i > \delta, |z_{ni}| \leq K\} \left(p(0)|z'_{ni}t|^2 + o(|z'_{ni}t|^2)\right),$$

the $o(\cdot)$ being uniform in n and i . Adding back contributions from the terms where $|\mu_i| \leq \delta$ or $|z_{ni}| > K$, we then get via (iv)* that

$$(11.4) \quad \Gamma_n(t) = p(0)|t|^2 + o(|t|^2) \quad \text{uniformly in } n.$$

In particular, if r is small enough,

$$\Gamma_n(t) \geq \frac{1}{2}p(0)|t|^2 \quad \text{for all } n, \text{ all } |t| \leq r.$$

This local lower bound implies a global lower bound,

$$(11.5) \quad \liminf_{n \rightarrow \infty} \inf_{|t| \geq r} \Gamma_n(t) > 0 \quad \text{for each } r > 0,$$

because $\Gamma_n(t)$ is an increasing function of $|t|$. The last inequality together with a uniform law of large numbers will imply consistency of $\hat{t}_n$.

Manageability. We will need maximal inequalities for both $\{f_{ni}\}$ and $\{R_{ni}\}$. Let us verify that both processes generate random subsets of $\mathbb{R}^n$ with a pseudodimension determined by d . From the results in Section 4, this will imply that both processes are manageable.

Consider first the set $\mathcal{F}_{n\omega}$ of all points in $\mathbb{R}^n$ with coordinates $f_{ni}(\omega, t)$, as t ranges over the set

$$T_n = \{t \in \mathbb{R}^d : \theta_0 + V_n^{-1}t \in \Theta\}.$$

We need to find a dimension V such that, for every β in $\mathbb{R}^{V+1}$, no $(V+1)$ -dimensional coordinate projection of $\mathcal{F}_{n\omega}$ can surround β . This property is not affected if we translate $\mathcal{F}_{n\omega}$ by a fixed amount; it is good enough to establish the property for the set of points with coordinates

$$|y_i^+ - (\mu_i + z'_{ni}t)^+| = \max[y_i^+ - (\mu_i + z'_{ni}t)^+, (\mu_i + z'_{ni}t)^+ - y_i^+].$$

From the stability results in Section 5 for pseudodimension, it is good enough to treat the two terms in the maximum separately. Consider, for example, the set of points with coordinate $y_i^+ - (\mu_i + z'_{ni}t)^+$. Again we translate to eliminate the y_i^+ . We now must determine, for $I = \{i_1, \dots, i_k\}$, with k suitably large, and β a point in $\mathbb{R}^k$, whether it is possible to find for each $J \subseteq I$ a t in T_n for which

$$(\mu_i + z'_{ni}t)^+ \begin{cases} > \beta_i & \text{for } i \in J, \\ < \beta_i & \text{for } i \in I \setminus J. \end{cases}$$

The inequalities when J is the empty set show that every β_i would have to be strictly positive, so the problem is unchanged if we replace $(\mu_i + z'_{ni}t)^+$ by $\mu_i + z'_{ni}t$, which is linear in t . Even if t ranges over the whole of $\mathbb{R}^d$, the points with these linear coordinate functions can at best trace out an affine subspace of dimension d . If $k = d + 1$, it is impossible to find for each J a t that satisfies the stated inequalities. By Lemma 5.1, the set $\mathcal{F}_{n\omega}$ has pseudodimension less than $10d$. (The bound could be improved, but there is no point in doing so; it matters only that the pseudodimension is the same for all n .)

Similar arguments serve to bound the pseudodimension for the set of points $\mathcal{R}_{n\omega}$ with coordinates $R_{ni}(\omega, t)/|t|$, as t ranges over the nonzero points in T_n . Indeed, inequalities

$$R_{ni}(\omega, t)/|t| \begin{cases} > \beta_i & \text{for } i \in J, \\ < \beta_i & \text{for } i \in I \setminus J, \end{cases}$$

are equivalent to

$$|y_i^+ - (\mu_i + z'_{ni}t)^+| - |y_i^+ - \mu_i^+| - \Delta_{ni}(\omega)'t - \beta_i|t| \begin{cases} > 0 & \text{for } i \in J, \\ < 0 & \text{for } i \in I \setminus J. \end{cases}$$

Again several translations and appeals to the stability property for maxima reduces the problem to the result for affine subspaces of dimension d . The sets $\mathcal{R}_{n\omega}$ have pseudodimension less than $1000d$ (or something like that).

Maximal Inequalities. The sets T_n all lie within some bounded region of $\mathbb{R}^d$; there is a constant κ such that $|t| \leq \kappa$ for every t in every T_n . It follows that

$$|f_{ni}(\omega, t)| \leq \kappa |z_{ni}| \quad \text{for all } t.$$

The maximal inequality (7.10) for manageable processes provides a constant C for which

$$\mathbb{P} \sup_t |G_n(\cdot, t) - \Gamma_n(t)|^2 \leq \frac{C}{n^2} \sum_i |z_{ni}|^2.$$

Condition (ii)* bounds the sum on the right-hand side by a multiple of $1/n$. We deduce that

$$(11.6) \quad \sup_t |G_n(\cdot, t) - \Gamma_n(t)| = o_p(1).$$

[Actually we get $O_p(1/\sqrt{n})$, but $o_p(1)$ will suffice for our later purposes.] For the remainder terms we have a slightly more complicated envelope for t in a small neighborhood of the origin,

$$\sup_{|t| \leq r} \frac{|R_{ni}(\omega, t)|}{|t|} \leq 2|z_{ni}| \left(\{|\mu_i| \leq r|z_{ni}|\} + \{|u_i(\omega)| \leq r|z_{ni}|\} \right).$$

Maximal inequality (7.10) provides another constant C for which

$$\begin{aligned} \frac{1}{n} \mathbb{P} \sup_{0 < |t| \leq r} \left| |t|^{-1} \sum_{i \leq n} R_{ni}(\omega, t) - \mathbb{P} R_{ni}(\cdot, t) \right|^2 \\ \leq \frac{C}{n} \sum_{i \leq n} |z_{ni}|^2 \left(\{|\mu_i| \leq r|z_{ni}|\} + \mathbb{P}\{|u_i| \leq r|z_{ni}|\} \right). \end{aligned}$$

By condition (ii)* the summands where $|z_{ni}| > K$ contribute at most $C\epsilon$ to the right-hand side. The remaining summands contribute at most

$$\frac{C}{n} \sum_{i \leq n} |z_{ni}|^2 \{|\mu_i| \leq Kr\} + \frac{C}{n} \mathbb{P}\{|u_1| \leq Kr\} \sum_{i \leq n} |z_{ni}|^2.$$

Conditions (iii)* and (i) ensure that this contribution converges to zero, uniformly in n , as $r \rightarrow 0$. It follows that

$$\left| \frac{1}{n} \sum_{i \leq n} R_{ni}(\omega, t) - \mathbb{P} R_n(\cdot, t) \right| = o_p(|t|/\sqrt{n})$$

uniformly in n and uniformly over t in shrinking neighborhoods of the origin. That is,

$$\begin{aligned} G_n(\omega, t) &= \Gamma_n(t) + \frac{1}{\sqrt{n}} W'_n t + o_p(|t|/\sqrt{n}) \\ (11.7) \quad &= p(0)|t|^2 + o(|t|^2) + \frac{1}{\sqrt{n}} W'_n t + o_p(|t|/\sqrt{n}) \quad \text{uniformly,} \end{aligned}$$

where the uniformity is over all n and all t in a neighborhood $\{|t| \leq r_n\}$, for every sequence $\{r_n\}$ of positive real numbers converging to zero.

Proof of the Theorem. Drop the ω from the notation. It will be enough if we show that $\hat{t}_n = o_p(1/\sqrt{n}) - W_n/2\sqrt{n}p(0)$. First establish consistency, by means of the inequality

$$G_n(\hat{t}_n) = \inf_t G_n(t) \leq G_n(0) = 0.$$

The random point $\hat{t}_n$ lies in the range over which the $o_p(1)$ bound from (11.6) holds. Approximating G_n by Γ_n we get

$$\Gamma_n(\hat{t}_n) \leq o_p(1).$$

Using (11.5) deduce that $\hat{t}_n = o_p(1)$. If r_n tends to zero slowly enough,

$$\mathbb{P}\{|\hat{t}_n| > r_n\} \rightarrow 0.$$

This brings $\hat{t}_n$ into the range where we can appeal to (11.7) to deduce

$$G_n(\hat{t}_n) = \left(p(0) + o_p(1)\right) \left|\hat{t}_n + \frac{W_n + o_p(1)}{2\sqrt{n}p(0)}\right|^2 - \frac{|W_n|^2}{4np(0)} + o_p(1/n).$$

When $-W_n/2\sqrt{n}p(0)$ lies in T_n , which happens with probability tending to one because θ_0 is an interior point of Θ and $W_n = O_p(1/\sqrt{n})$, we can again invoke approximation (11.7) to get

$$G_n(-W_n/2\sqrt{n}p(0)) = -\frac{|W_n|^2}{4np(0)} + o_p(1/n).$$

From the comparison

$$G_n(\hat{t}_n) \leq G_n(-W_n/2\sqrt{n}p(0)),$$

deduce that

$$\left|\hat{t}_n + \frac{W_n + o_p(1)}{2\sqrt{n}p(0)}\right|^2 = o_p(1/n),$$

from which the desired approximation to $\hat{t}_n$ follows.

REMARKS. For the theory of (uncensored) least absolute deviations estimators see Bloomfield and Steiger (1983). A central limit theorem for such estimators was derived using elementary convexity arguments (which will reappear in Section 14) by Pollard (1990).

Chapter 10 of Amemiya (1985) describes many different approaches to estimation for Tobit models.

SECTION 12

Random Convex Sets

Donoho (1982) and Donoho and Gasko (1987) studied an operation proposed by Tukey for extending the idea of trimming to multidimensional data. Nolan (1989a) gave a rigorous treatment of the asymptotic theory. Essentially the arguments express the various statistics of interest as differentiable functionals of an empirical measure. The treatment in this section will show how to do this without the formal machinery of compact differentiability for functionals, by working directly with almost sure representations. [Same amount of work, different packaging.]

To keep the discussion simple, let us consider the case of an independent sample $\xi_1, \xi_2, \dots$ of random vectors from the symmetric bivariate normal distribution P on $\mathbb{R}^2$, and consider only the analogue of 25% trimming.

The notation will be cleanest when expressed (using traditional empirical process terminology) in terms of the *empirical measure* P_n , which puts mass $1/n$ at each of the points $\xi_1(\omega), \dots, \xi_n(\omega)$.

Let $\mathcal{H}$ denote the class of all closed halfspaces in $\mathbb{R}^2$. Define a random compact, convex set $K_n = K_n(\omega)$ by intersecting all those halfspaces that contain at least $3/4$ of the observations:

$$K_n(\omega) = \bigcap \{H \in \mathcal{H} : P_n H \geq \tfrac{3}{4}\}.$$

It is reasonable to hope that K_n should settle down to the set

$$B(r_0) = \bigcap \{H \in \mathcal{H} : PH \geq \tfrac{3}{4}\},$$

which is a closed ball centered at the origin with radius r_0 equal to the 75% point of the one-dimensional standard normal distribution. That is, if Φ denotes the $N(0, 1)$ distribution function, then $r_0 = \Phi^{-1}(3/4) \approx .675$. Indeed, a simple continuity argument based on a uniform strong law of large numbers,

$$(12.1) \quad \sup_{\mathcal{H}} |P_n H - PH| \rightarrow 0 \quad \text{almost surely,}$$

would show that, for each $\epsilon > 0$, there is probability one that

$$B(r_0 - \epsilon) \subseteq K_n(\omega) \subseteq B(r_0 + \epsilon) \quad \text{eventually.}$$

In a natural sense, K_n is a strongly consistent estimator. Let us not dwell on the details here, because the next argument, which gives the finer asymptotics for K_n , is much more interesting. [The almost sure representation that will appear soon would imply the “in probability” version of (12.1). This would give consistency in probability, which is all that we really need before embarking upon the asymptotic distribution theory for K_n .]

Once K_n contains the origin as an interior point it makes sense to describe its boundary in polar coordinates. Let $R_n(\theta) = R_n(\omega, \theta)$ denote the distance from the origin to the boundary in the direction θ . The consistency result then has the reformulation:

$$\sup_{\theta} |R_n(\omega, \theta) - r_0| \rightarrow 0 \quad \text{almost surely.}$$

With the help of the functional central limit theorems from Section 10, we can improve this to get convergence in distribution of a random process,

$$\sqrt{n}(R_n(\omega, \theta) - r_0) \quad \text{for } -\pi \leq \theta \leq \pi,$$

to a Gaussian process indexed by θ . [It would be more elegant to take the unit circle as the index set, identifying the points $\theta = \pi$ and $\theta = -\pi$.] Such a result would imply central limit theorems for a variety of statistics that could be defined in terms of K_n .

Heuristics. We need to establish a functional central limit theorem for the standardized *empirical process*,

$$\nu_n(\omega, H) = \sqrt{n}(P_n H - PH),$$

as a stochastic process indexed by $\mathcal{H}$. We must show that $\{\nu_n\}$ converges in distribution to a Gaussian process ν indexed by $\mathcal{H}$.

Let $H(r, \theta)$ denote the closed halfspace containing the origin with boundary line perpendicular to the θ direction at a distance r from the origin. That is, $H(r, \theta)$ consists of all points whose projections onto a unit vector in the θ direction are $\leq r$. For a given point with polar coordinates (r, θ) , the halfspace $H(r, \theta)$ maximizes PH over all H that have (r, θ) as a boundary point. The boundary point of $B(r_0)$ in the direction θ is determined by solving the equation $PH(r, \theta) = 3/4$ for r , giving $r = r_0$. Similarly, the boundary point of K_n in the direction θ is almost determined by solving the equation $P_n H(r, \theta) = 3/4$, as we will soon see. (Discreteness of P_n might prevent us from getting exact equality; and the halfspace that determines the boundary point will be rotated slightly from the $H(r, \theta)$ position.) That is, $R_n(\theta)$ is approximately determined by solving the following equation for r :

$$\frac{3}{4} \approx P_n H(r, \theta) = PH(r, \theta) + \frac{1}{\sqrt{n}} \nu_n H(r, \theta).$$

Asymptotically the right-hand side is distributed as

$$\Phi(r) + \frac{1}{\sqrt{n}} \nu H(r, \theta) \approx \Phi(r_0) + (r - r_0) \Phi'(r_0) + \frac{1}{\sqrt{n}} \nu H(r_0, \theta).$$

Thus $\sqrt{n}(R_n(\theta) - r_0)$ should behave asymptotically like $-\nu H(r_0, \theta) / \Phi'(r_0)$, which is a Gaussian process indexed by θ .

The functional limit theorem for ν_n . Define a triangular array of processes,

$$f_{ni}(\omega, H) = \frac{1}{\sqrt{n}} \{\xi_i(\omega) \in H\} \quad \text{for } H \in \mathcal{H} \text{ and } i \leq n.$$

They have constant envelopes $F_{ni} = 1/\sqrt{n}$. We will apply the Functional Central Limit Theorem of Section 10 to the processes

$$\nu_n H = \sum_{i \leq n} \left(f_{ni}(\omega, H) - \mathbb{P} f_{ni}(\cdot, H) \right).$$

It is easy to show, by an appeal to Lemma 4.4, that the processes define random subsets of $\mathbb{R}^n$ with pseudodimension 3. Every closed halfspace has the form

$$H = \{x \in \mathbb{R}^2 : \alpha \cdot x + \beta \geq 0\}$$

for some unit vector α in $\mathbb{R}^2$ and some real number β . Notice that $f_{ni}(\omega, H) = 1/\sqrt{n}$ if and only if $\alpha \cdot \xi_i + \beta \geq 0$. The points in $\mathbb{R}^n$ with coordinates $\alpha \cdot \xi_i + \beta$ trace out a subset of a 3-dimensional subspace as α and β vary.

The other conditions of the Theorem are just as easy to check. For every pair of halfspaces H_1 and H_2 , and every n ,

$$\mathbb{P}(\nu_n H_1 \nu_n H_2) = P H_1 H_2 - P H_1 P H_2,$$

and

$$\rho(H_1, H_2)^2 = \rho_n(H_1, H_2)^2 = P|H_1 - H_2|.$$

[Typically, manageability is the only condition that requires any work when the Functional Central Limit Theorem is applied to the standardized sums of independent, identically distributed processes.]

The Theorem asserts that ν_n converges in distribution, as a random element of the function space $B(\mathcal{H})$, to a Gaussian process concentrated on $U(\mathcal{H})$, the set of all bounded, ρ -uniformly continuous functions. The Representation Theorem from Section 9 provides perfect maps ϕ_n and a Gaussian process $\tilde{\nu}$ with sample paths in $U(\mathcal{H})$ such that the random processes $\tilde{\nu}_n = \nu_n \circ \phi_n$ satisfy

$$\sup_{\mathcal{H}} |\tilde{\nu}_n(H) - \tilde{\nu}(H)| \rightarrow 0 \quad \text{almost surely.}$$

We need not worry about measurability difficulties here, because the supremum over $\mathcal{H}$ is equal to the supremum over an appropriate countable subclass of $\mathcal{H}$. The representation also gives a new version of the empirical measure,

$$(12.2) \quad \tilde{P}_n H = P H + \frac{1}{\sqrt{n}} \tilde{\nu}_n H = P H + \frac{1}{\sqrt{n}} (\tilde{\nu} H + o(1)),$$

where the $o(1)$ represents a function of H that converges to zero uniformly over $\mathcal{H}$.

Asymptotics. With (12.2) we have enough to establish an almost sure limit result for $\tilde{R}_n(\tilde{\omega}, \theta) = R_n(\phi_n(\tilde{\omega}), \theta)$, which will imply the corresponding distributional result for $R_n(\omega, \theta)$. Let $\{\delta_n\}$ be a sequence of random variables on $\tilde{\Omega}$ that

converges almost surely to zero at a rate to be specified soon. Define

$$\begin{aligned} Z(\theta) &= \tilde{\nu}(H(r_0, \theta)) / \Phi'(r_0), \\ \ell_n(\theta) &= r_0 - \frac{1}{\sqrt{n}}(Z(\theta) + \delta_n), \\ u_n(\theta) &= r_0 - \frac{1}{\sqrt{n}}(Z(\theta) - \delta_n). \end{aligned}$$

If we can find δ_n uniformly of order $o(1)$ such that, eventually,

$$\ell_n(\theta) \leq \tilde{R}_n(\theta) \leq u_n(\theta) \quad \text{for all } \theta,$$

then it will follow that

$$\sqrt{n}(\tilde{R}_n(\theta) - r_0) \rightarrow -Z(\theta) \quad \text{uniformly in } \theta,$$

as desired.

Consider first the upper bound on $\tilde{R}_n(\theta)$. Temporarily write $H_n(\theta)$ for the half-space $H(u_n(\theta), \theta)$. Then

$$\tilde{P}_n H_n(\theta) = P H_n(\theta) + \frac{1}{\sqrt{n}}(\tilde{\nu} H_n(\theta) + o(1)) \quad \text{uniformly in } \theta.$$

Apply the Mean Value Theorem to approximate the contribution from P :

$$\begin{aligned} P H_n(\theta) &= \Phi(u_n(\theta)) \\ &= \Phi(r_0) + (u_n(\theta) - r_0)(\Phi'(r_0) + o(1)) \\ &= \frac{3}{4} - \frac{1}{\sqrt{n}}\left(\tilde{\nu} H(r_0, \theta) - o(1) - (\Phi'(r_0) + o(1))\delta_n\right), \end{aligned}$$

where the $o(1)$ represent functions of θ that converge to zero uniformly in θ . For the contribution from $\tilde{\nu}$ consider first the difference $|H_n(\theta) - H(r_0, \theta)|$. It is the indicator function of a strip of width $|Z(\theta) - \delta_n|/\sqrt{n}$; its P measure converges to zero uniformly in θ . Thus

$$\rho(H_n(\theta), H(r_0, \theta)) \rightarrow 0 \quad \text{uniformly in } \theta.$$

By the uniform continuity of the $\tilde{\nu}$ sample paths it follows that

$$\tilde{\nu}(H_n(\theta)) = \tilde{\nu} H(r_0, \theta) + o(1) \quad \text{uniformly in } \theta.$$

Adding the two contributions to $\tilde{P}_n H_n(\theta)$ we get

$$\tilde{P}_n H_n(\theta) = \frac{3}{4} + \frac{1}{\sqrt{n}}\left((\Phi'(r_0) + o(1))\delta_n - o(1)\right).$$

We can choose δ_n converging to zero while ensuring that the coefficient of $1/\sqrt{n}$ is always positive. With that choice, the set $H_n(\theta)$ becomes one of the half spaces whose intersection defines $\tilde{K}_n$; the boundary point in the θ direction must lie on the ray from the origin to the boundary of $H_n(\theta)$; the distance $\tilde{R}_n(\theta)$ must be less than $u_n(\theta)$.

Now consider the lower bound on $\tilde{R}_n(\theta)$. Let $\mathbf{t}_n(\theta)$ denote the point a distance $\ell_n(\theta)$ from the origin in the θ direction. It is enough if we show that $\tilde{K}_n$ contains every $\mathbf{t}_n(\theta)$.

If, for a particular θ , the point $\mathbf{t}_n(\theta)$ were outside $\tilde{K}_n$, there would exist a halfspace H with $\tilde{P}_n H \geq 3/4$ and $\mathbf{t}_n(\theta) \notin H$. By sliding H towards $\mathbf{t}_n(\theta)$ we would get an H' with $\tilde{P}_n H' \geq 3/4$ and $\mathbf{t}_n(\theta)$ on the boundary of H' . The right choice for δ_n will ensure that such an H' cannot exist.

For each θ let $H_n(\theta)$ denote the halfspace with $\mathbf{t}_n(\theta)$ on its boundary and the largest $\tilde{P}_n$ measure. (Of course this is not the same $H_n(\theta)$ as before.) The maximum of PH over all halfspaces with $\mathbf{t}_n(\theta)$ on the boundary is achieved at $H(\ell_n(\theta), \theta)$. So, uniformly in θ ,

$$\frac{3}{4} \leq \tilde{P}_n H_n(\theta) = PH_n(\theta) + O(1/\sqrt{n}) \leq PH(\ell_n(\theta), \theta) + O(1/\sqrt{n}) \rightarrow \frac{3}{4}.$$

It follows that $PH_n(\theta)$ also converges uniformly to $3/4$. This forces the boundary of $H_n(\theta)$ to orient itself more and more nearly perpendicular to the θ direction. Consequently,

$$\rho(H_n(\theta), H(r_0, \theta)) \rightarrow 0 \quad \text{uniformly in } \theta.$$

Uniform continuity of the $\tilde{\nu}$ sample paths now lets us assert

$$\tilde{P}_n H_n(\theta) = PH_n(\theta) + \frac{1}{\sqrt{n}} \left(\tilde{\nu} H(r_0, \theta) + o(1) \right) \quad \text{uniformly in } \theta.$$

Again using the fact that the maximum of PH over all halfspaces with $\mathbf{t}_n(\theta)$ on the boundary is achieved at $H(\ell_n(\theta), \theta)$, we deduce that, uniformly in θ ,

$$\begin{aligned} PH_n(\theta) &\leq PH(\ell_n(\theta), \theta) \\ &= \Phi(\ell_n(\theta)) \\ &= \Phi(r_0) + (\ell_n(\theta) - r_0) \left(\Phi'(r_0) + o(1) \right) \\ &= \frac{3}{4} - \frac{1}{\sqrt{n}} \left(\tilde{\nu} H(r_0, \theta) - o(1) + (\Phi'(r_0) + o(1)) \delta_n \right). \end{aligned}$$

With δ_n converging to zero slowly enough to cancel out all the $o(1)$ terms, plus a little bit more, we get a contradiction, $\tilde{P}_n H_n(\theta) < 3/4$ for all θ . There can therefore be no halfspace with $\tilde{P}_n H' \geq 3/4$ and $\mathbf{t}_n(\theta)$ on its boundary. The point $\mathbf{t}_n(\theta)$ must lie inside K_n . The argument for the lower bound on $\tilde{R}_n(\theta)$ is complete.

REMARKS. Nolan (1989b) has studied an estimator related to K_n , following Donoho (1982). Its analysis is similar to the arguments given in this section, but more delicate.

SECTION 13

Estimation from Censored Data

Let P be a nonatomic probability distribution on $[0, \infty)$. The cumulative hazard function β is defined by

$$\beta(t) = \int \frac{\{0 \leq x \leq t\}}{P[x, \infty)} P(dx).$$

It uniquely determines P . Let $T_1, T_2, \dots$ be independent observations from P and $\{c_i\}$ be a deterministic sequence of nonnegative numbers representing censoring times. Suppose the data consist of the variables

$$T_i \wedge c_i \quad \text{and} \quad \{T_i \leq c_i\} \quad \text{for } i = 1, \dots, n.$$

That is, we observe T_i if it is less than or equal to c_i ; otherwise we learn only that T_i was censored at time c_i . We always know whether T_i was censored or not.

If the $\{c_i\}$ behave reasonably, we can still estimate the true β despite the censoring. One possibility is to use the Nelson estimator:

$$\hat{\beta}_n(t) = \frac{1}{n} \sum_{i \leq n} \frac{\{T_i \leq c_i \wedge t\}}{L_n(T_i)},$$

where

$$L_n(t) = \frac{1}{n} \sum_{i \leq n} \{T_i \wedge c_i \geq t\}.$$

It has become common practice to analyze $\hat{\beta}_n$ by means of the theory of stochastic integration with respect to continuous-time martingales. This section will present an alternative analysis using the Functional Central Limit Theorem from Section 10. Stochastic integration will be reduced to a convenient, but avoidable, means for calculating limiting variances and covariances.

Heuristics. Write $G(t)$ for $\mathbb{P}\{T_i \geq t\}$ and define

$$\Gamma_n(t) = \frac{1}{n} \sum_{i \leq n} \{c_i \geq t\}.$$

Essentially we need to justify replacement of L_n by its expected value,

$$\mathbb{P}L_n(t) = \frac{1}{n} \sum_{i \leq n} \mathbb{P}\{T_i \geq t\} \{c_i \geq t\} = G(t) \Gamma_n(t).$$

That would approximate $\hat{\beta}_n$ by an average of independent processes, which should be close to its expected value:

$$\begin{aligned} \hat{\beta}_n(t) &\approx \frac{1}{n} \sum_{i \leq n} \frac{\{T_i \leq c_i \wedge t\}}{G(T_i) \Gamma_n(T_i)} \\ &\approx \frac{1}{n} \sum_{i \leq n} \mathbb{P} \frac{\{T_i \leq t\} \{T_i \leq c_i\}}{G(T_i) \Gamma_n(T_i)} \\ &= \mathbb{P} \left(\frac{\{T_1 \leq t\}}{G(T_1) \Gamma_n(T_1)} \frac{1}{n} \sum_{i \leq n} \{T_1 \leq c_i\} \right) \\ &= \beta(t). \end{aligned}$$

A more precise analysis will lead to a functional central limit theorem for the standardized processes $\sqrt{n}(\hat{\beta}_n - \beta)$ over an interval $[0, \tau]$, if we assume that:

- (i) the limit $\Gamma(t) = \lim_{n \rightarrow \infty} \Gamma_n(t)$ exists for each t ;
- (ii) the value τ is such that $G(\tau) > 0$ and $\Gamma(\tau) > 0$.

The argument will depend upon a limit theorem for a process indexed by pairs (t, m) , where $0 \leq t \leq \tau$ and m belongs to the class $\mathcal{M}$ of all nonnegative increasing functions on $[0, \tau]$. Treating β as a measure on $[0, \tau]$, define

$$\begin{aligned} \beta(t, m) &= \int \{0 \leq x \leq t\} m(x) \beta(dx), \\ f_i(\omega, t, m) &= \{T_i \leq t \wedge c_i\} m(T_i) - \beta(t \wedge T_i \wedge c_i, m). \end{aligned}$$

Such a centering for f_i is suggested by martingale theory, as will be explained soon. We will be able to establish a functional central limit theorem for

$$\begin{aligned} X_n(t, m) &= \frac{1}{\sqrt{n}} \sum_{i \leq n} f_i(\omega, t, m) \\ &= \sqrt{n} \left(\left(\frac{1}{n} \sum_{i \leq n} \{T_i \leq t \wedge c_i\} m(T_i) \right) - \beta(t, mL_n) \right). \end{aligned}$$

Putting m equal to $1/L_n$ we get the standardized Nelson estimator:

$$X_n(t, 1/L_n) = \sqrt{n}(\hat{\beta}_n(t) - \beta(t)).$$

The limit theorem for X_n will justify the approximation

$$X_n(t, 1/L_n) \approx X_n(t, 1/G\Gamma_n).$$

It will also give the limiting distribution for the approximating process.

Some martingale theory. The machinery of stochastic integration with respect to martingales provides a very neat way of calculating variances and covariances for the f_i processes. We could avoid stochastic integration altogether by direct, brute force calculation; but then the happy cancellations arranged by the martingales would appear most mysterious and fortuitous.

The basic fact, not altogether trivial (Dellacherie 1972, Section V.5), is that both

$$Z_i(t) = \{T_i \leq t\} - \beta(t \wedge T_i) \quad \text{and} \quad Z_i(t)^2 - \beta(t \wedge T_i)$$

are continuous parameter martingales in t . That is, both the simple jump process $\{T_i \leq t\}$ and the submartingale Z_i^2 have compensator $\beta(t \wedge T_i)$. The f_i process is expressible as a stochastic integral with respect to Z_i :

$$f_i(\omega, t, m) = \int \{0 \leq x \leq t \wedge c_i\} m(x) Z_i(dx).$$

It follows that, for fixed m , the process f_i is also a martingale in t . In particular, $\mathbb{P}f_i(\omega, t, m) = \mathbb{P}f_i(\omega, 0, m) = 0$ for every t .

From now on let us omit the ω from the notation.

Stochastic integration theory tells us how to calculate compensators for new processes derived from the martingales $\{Z_i\}$. In particular, for fixed t_1, t_2, m_1 , and m_2 , the product $f_i(t \wedge t_1, m_1)f_i(t \wedge t_2, m_2)$ has compensator

$$A_i(t) = \int \{0 \leq x \leq t \wedge t_1 \wedge t_2 \wedge T_i \wedge c_i\} m_1(x) m_2(x) \beta(dx);$$

the difference $f_i(t \wedge t_1, m_1)f_i(t \wedge t_2, m_2) - A_i(t)$ is a martingale in t . This implies that

$$\mathbb{P}f_i(t \wedge t_1, m_1)f_i(t \wedge t_2, m_2) = \mathbb{P}A_i(t) \quad \text{for each } t.$$

Put $t = \max(t_1, t_2)$, then average over i . Because each T_i has the same distribution, we get

$$\begin{aligned} \mathbb{P}X_n(t_1, m_1)X_n(t_2, m_2) &= \frac{1}{n} \sum_{i \leq n} \mathbb{P}f_i(t_1, m_1)f_i(t_2, m_2) \\ &= \mathbb{P} \int \{0 \leq x \leq t_1 \wedge t_2\} L_n(x) m_1(x) m_2(x) \beta(dx) \\ (13.1) \quad &= \beta(t_1 \wedge t_2, G\Gamma_n m_1 m_2). \end{aligned}$$

The calculations needed to derive this result directly would be comparable to the calculations needed to establish the martingale property for Z_i .

Manageability. For each positive constant K let $\mathcal{M}(K)$ denote the class of all those m in $\mathcal{M}$ for which $m(\tau) \leq K$. To establish manageability of the $\{f_i(t, m)\}$ processes, as t ranges over $[0, \tau]$ (or even over the whole of $\mathbb{R}^+$) and m ranges over $\mathcal{M}(K)$, it suffices to consider separately the three contributions to f_i .

Let us show that the indicator functions $\{T_i \leq t \wedge c_i\}$ define a set with pseudo-dimension one. Suppose the (i, j) -projection could surround some point in $\mathbb{R}^2$. Suppose $T_i \leq T_j$. We would need to be able to find t_1 and t_2 such that both pairs

of inequalities,

$$\begin{aligned} T_i &\leq t_1 \wedge c_i & \text{and} & & T_j &\leq t_1 \wedge c_j, \\ T_i &> t_2 \wedge c_i & \text{and} & & T_j &\leq t_2 \wedge c_j, \end{aligned}$$

were satisfied. The first pair would imply $T_i \leq c_i$ and $T_j \leq c_j$, and then the second pair would lead to a contradiction, $t_2 \geq T_j \geq T_i > t_2$, which would establish the assertion about pseudodimension.

For the factors $\{m(T_i)\}$ with m ranging over $\mathcal{M}(K)$, we can appeal to the result from Example 6.3 if we show that no 2-dimensional projection of the convex cone generated by $\mathcal{M}(K)$ can surround the point (K, K) . This is trivial. For if $T_i \leq T_j$ then no $r \in \mathbb{R}^+$ and $m \in \mathcal{M}(K)$ can achieve the pair of inequalities $rm(T_i) > K$ and $rm(T_j) < K$.

The argument for the third contribution to f_i is similar. For each $t \leq \tau$ and $m \in \mathcal{M}(K)$, the process $\beta(t \wedge T_i \wedge c_i, m)$ is less than $K' = K\beta(\tau)$. If, for example, $T_i \wedge c_i \leq T_j \wedge c_j$ then it is impossible to find an $r \in \mathbb{R}^+$, an $m \in \mathcal{M}(K)$, and a $t \in [0, \tau]$ such that $r\beta(t \wedge T_i \wedge c_i, m) > K'$ and $r\beta(t \wedge T_j \wedge c_j, m) < K'$.

Functional Central Limit Theorem. It is a simple matter to check the five conditions of the Functional Central Limit Theorem from Section 10 for the triangular array of processes

$$f_{ni}(t, m) = \frac{1}{\sqrt{n}} f_i(t, m) \quad \text{for } i = 1, \dots, n, \quad t \in [0, \tau], \quad m \in \mathcal{M}(K),$$

for some constant K to be specified. These processes have constant envelopes,

$$F_{ni} = K(1 + \beta(\tau))/\sqrt{n},$$

which clearly satisfy conditions (iii) and (iv) of the theorem. The extra $1/\sqrt{n}$ factor does not affect the manageability. Taking the limit in (13.1) we get

$$H((t_1, m_1), (t_2, m_2)) = \beta(t_1 \wedge t_2, G\Gamma m_1 m_2).$$

For simplicity suppose $t_1 \leq t_2$. Then, because f_{ni} has zero expected value, (13.1) also gives

$$\begin{aligned} \rho_n((t_1, m_1), (t_2, m_2))^2 &= \mathbb{P}|X_n(t_1, m_1) - X_n(t_2, m_2)|^2 \\ &= \beta(t_1, G\Gamma_n m_1^2) + \beta(t_2, G\Gamma_n m_2^2) - 2\beta(t_1, G\Gamma_n m_1 m_2) \\ &= \int \{0 \leq x \leq t_1\} G\Gamma_n (m_1 - m_2)^2 \beta(dx) + \int \{t_1 \leq x \leq t_2\} G\Gamma_n m_2^2 \beta(dx) \\ &\leq \int \{0 \leq x \leq t_1\} (m_1 - m_2)^2 \beta(dx) + \int \{t_1 \leq x \leq t_2\} m_2^2 \beta(dx). \end{aligned}$$

A similar calculation with Γ_n replaced by Γ gives

$$\begin{aligned} \rho((t_1, m_1), (t_2, m_2))^2 &= \int \{0 \leq x \leq t_1\} G\Gamma (m_1 - m_2)^2 \beta(dx) + \int \{t_1 \leq x \leq t_2\} G\Gamma m_2^2 \beta(dx), \end{aligned}$$

which is greater than the positive constant factor $G(\tau)\Gamma(\tau)$ times the upper bound just obtained for $\rho_n((t_1, m_1), (t_2, m_2))^2$. The second part of condition (v) of the Functional Central Limit Theorem follows.

The processes $\{X_n(t, m)\}$, for $0 \leq t \leq \tau$ and $m \in \mathcal{M}(K)$, converge in distribution to a Gaussian process $X(t, m)$ with ρ -continuous paths, zero means, and covariance kernel H .

Asymptotics for $\hat{\beta}_n$. We now have all the results needed to make the heuristic argument precise. A straightforward application of Theorem 8.2 shows that

$$\sup_t |L_n(t) - G(t)\Gamma_n(t)| \rightarrow 0 \quad \text{almost surely.}$$

If we choose the constant K so that $G(\tau)\Gamma(\tau) > 1/K$, then, with probability tending to one, both $1/L_n$ and $1/G\Gamma_n$ belong to $\mathcal{M}(K)$ and

$$\sup_{0 \leq t \leq \tau} \rho((t, 1/L_n), (t, 1/G\Gamma_n)) \rightarrow 0 \quad \text{in probability.}$$

From stochastic equicontinuity of $\{X_n\}$ we then deduce that

$$\begin{aligned} \sqrt{n}(\hat{\beta}_n(t) - \beta(t)) &= X_n(t, 1/L_n) \\ &= X_n(t, 1/G\Gamma_n) + o_p(1) \quad \text{uniformly in } 0 \leq t \leq \tau \\ &\rightsquigarrow X(t, 1/G\Gamma). \end{aligned}$$

The limit is a Gaussian process on $[0, \tau]$ with zero means and covariance kernel $\beta(t_1 \wedge t_2, 1/G\Gamma)$. It is a Brownian motion with a stretched out time scale.

REMARKS. As suggested by Meier (1975), deterministic censoring times $\{c_i\}$ allow more flexibility than the frequently made assumption that the $\{c_i\}$ are independent and identically distributed random variables. A conditioning argument would reduce the case of random $\{c_i\}$ to the deterministic case, anyway.

The method introduced in this section may seem like a throwback to the original proof by Breslow and Crowley (1974). However, the use of processes indexed by $\mathcal{M}(K)$ does eliminate much irksome calculation. More complicated forms of multivariate censoring might be handled by similar methods. For a comparison with the stochastic integral approach see Chapter 7 of Shorack and Wellner (1986).

I am grateful to Hani Doss for explanations that helped me understand the role of martingale methods.

SECTION 14

Biased Sampling

Vardi (1985) introduced a far-reaching extension of the classical model for length-biased sampling. He solved the problem of estimating a distribution function based on several independent samples, each subject to a different form of selection bias. Using empirical process theory, Gill, Vardi and Wellner (1988) developed the asymptotic theory for generalizations of Vardi's method to abstract settings. They showed that the general model includes many interesting examples as special cases. This section presents a reworking of the ideas in those two papers. It is part of a study carried out by me in collaboration with Robert Sherman of Yale University.

The general problem is to estimate a distribution P on some set S using independent samples of sizes n_{i+} from distributions Q_i , for $i = 1, \dots, s$, where the Q_i are related to P by means of known nonnegative weight functions $W_i(\cdot)$ on S :

$$\frac{dQ_i}{dP} = \pi_i W_i(\cdot) \quad \text{where } \pi_i = 1/PW_i.$$

Of course the normalizing constants π_i , which we must assume to be finite and strictly positive, are unknown. For example, the W_i might be indicator functions of various subdomains of S . The problem is then one of combining the different samples in order to form an estimate of P over the whole of S . The difficulty lies in deciding how to combine the information from samples whose subdomains overlap.

For the general problem, to ensure that we get information about P over the whole domain, we must assume that the union of the sets $\{W_i > 0\}$ covers S .

Vardi suggested that a so-called nonparametric maximum likelihood estimator $\hat{P}_n$ be used. This is a discrete probability measure that concentrates on the combined observations $x_1, x_2, \dots$ from all s samples. If x_j appears a total of n_{ij} times in the i^{th} sample, the combined empirical measure $\hat{Q}_n$ puts mass n_{+j}/n at x_j , where

$$n_{+j} = \sum_i n_{ij} \quad \text{and} \quad n = \sum_{i,j} n_{ij}.$$

The estimator $\hat{P}_n$ modifies $\hat{Q}_n$, putting at x_j the mass $\hat{p}_j$ defined by maximization

of a pseudo log-likelihood function: maximize

$$\sum_{i,j} n_{ij} \left[\log p_j - \log \left(\sum_k W_i(x_k) p_k \right) \right],$$

subject to the constraints

$$p_j > 0 \quad \text{for each } j, \quad \text{and} \quad \sum_j p_j = 1.$$

In this form the estimation problem involves parameters whose number increases with the sample sizes. The first part of the analysis will show how to transform the problem into an equivalent maximization involving only a fixed number of unknown parameters.

Simplify the notation by writing W_{ik} for $W_i(x_k)$. Reparametrize by substituting $\exp(\beta_j)$ for p_j . Then we need to maximize the function

$$L_n(\beta) = \sum_j n_{+j} \beta_j - \sum_i n_{i+} \log \left(\sum_k W_{ik} \exp(\beta_k) \right)$$

over all real $\{\beta_j\}$, subject to the constraint

$$\sum_j \exp(\beta_j) = 1.$$

Let $\mathbf{1}$ denote a vector of ones. The criterion function L_n is constant along the lines $\{\beta + t\mathbf{1} : t \in \mathbb{R}\}$; the constraint serves to locate a unique point on each such line.

Simple calculus shows that L_n is a concave function. Indeed, for each fixed β and δ the function $L_n(\beta + t\delta)$ has derivative

$$(14.1) \quad \sum_j n_{+j} \delta_j - \sum_i n_{i+} \left(\frac{\sum_k W_{ik} \delta_k \exp(\beta_k)}{\sum_k W_{ik} \exp(\beta_k)} \right) \quad \text{at } t = 0,$$

and second derivative

$$- \sum_{i,k} n_{i+} B_{ik} (\delta_k - \bar{\delta}_i)^2 \quad \text{at } t = 0,$$

where

$$B_{ik} = W_{ik} \exp(\beta_k) / \sum_j W_{ij} \exp(\beta_j)$$

and $\bar{\delta}_i$ is the weighted average

$$\bar{\delta}_i = \sum_k B_{ik} \delta_k.$$

Clearly the second derivative is always nonpositive; the function L_n is concave along every line. The second derivative can equal zero only if δ_k is constant over each of the subsets

$$K(i) = \{k : W_{ik} > 0\} \quad \text{for } i = 1, \dots, s.$$

Under mild connectedness assumptions about the regions $\{W_i > 0\}$, it can be shown (almost surely as the n_{i+} tend to infinity) that constancy over each $K(i)$ forces δ to

be a multiple of $\mathbf{1}$. That is, L_n is strictly concave along all directions except the $\mathbf{1}$ direction. Moreover, the connectedness assumption also forces the derivative to be strictly negative for t large enough. It follows that the constrained maximization problem eventually has a unique solution $\hat{\beta}$. (Clearly $\hat{\beta}$ depends on n , but I will omit the subscript n to avoid notational clutter.)

For a precise statement of the necessary connectedness property, see pages 1071-1072 of Gill et al. (1988). Let us assume such a property to hold from now on.

Transformation to an equivalent problem. The function L_n must have all its directional derivatives equal to zero at its maximizing point. Putting δ equal to a vector with δ_j as its only nonzero component, we get from (14.1) that

$$(14.2) \quad \exp(\hat{\beta}_j) = \frac{n_{+j}}{\sum_i (n_{i+} W_{ij} / \sum_k W_{ik} \exp(\hat{\beta}_k))} \quad \text{for each } j.$$

Notice that the s linear combinations of the $\exp(\hat{\beta}_k)$ values on the right-hand side determine all the $\hat{\beta}_j$ values. That is why we will be able to reduce the problem to one involving only s unknown parameters.

Introduce new parameters $\alpha_1, \dots, \alpha_s$. Trivially, the constrained maximization of L_n is equivalent to the problem: maximize

$$\sum_j n_{+j} \beta_j + \sum_i n_{i+} \alpha_i$$

subject to the constraints

$$\begin{aligned} \sum_j \exp(\beta_j) &= 1, \\ \exp(-\alpha_i) &= \sum_j W_{ij} \exp(\beta_j) \quad \text{for each } i. \end{aligned}$$

Equality (14.2) translates into a set of relations that the maximizing $\hat{\alpha}$ and $\hat{\beta}$ must satisfy; the maximization problem is unaffected if we add another constraint,

$$\exp(\beta_j) = \frac{n_{+j}}{\sum_i n_{i+} W_{ij} \exp(\alpha_i)} \quad \text{for each } j,$$

to the list. This allows us to eliminate the $\{\beta_j\}$ from the problem altogether, leaving a constrained maximization over the $\{\alpha_i\}$: maximize

$$\sum_i n_{i+} \alpha_i - \sum_j n_{+j} \log \left(\sum_k n_{k+} W_{kj} \exp(\alpha_k) \right),$$

subject to the constraints

$$\begin{aligned} \sum_j \frac{n_{+j}}{\sum_i n_{i+} W_{ij} \exp(\alpha_i)} &= 1, \\ \exp(-\alpha_i) &= \sum_j \frac{W_{ij} n_{+j}}{\sum_k n_{k+} W_{kj} \exp(\alpha_k)} \quad \text{for each } i. \end{aligned}$$

Just as the addition of an extra constraint did not affect the previous maximization, so will the elimination of a constraint not affect this maximization. By marvellous good luck (What is going on here?) the last set of equations corresponds exactly to the requirement that the directional derivatives of the criterion function all equal zero at its global maximizing value $\hat{\alpha}$; it can be discarded without changing the problem. The remaining constraint then serves only to locate a unique point along the lines of constancy of the criterion function.

The Vardi procedure takes a much neater form when expressed in empirical process notation. Write λ_{ni} for the proportion n_{i+}/n of observations that belong to the i^{th} sample, and $h_n(\cdot, \alpha)$ for the function $(\sum_i \lambda_{ni} \exp(\alpha_i) W_i(\cdot))^{-1}$. Then the Vardi estimator is determined by: maximize

$$M_n(\alpha) = \lambda'_n \alpha + \hat{Q}_n \log h_n(\cdot, \alpha)$$

subject to the constraint

$$\hat{Q}_n h_n(\cdot, \alpha) = 1.$$

Under the connectedness assumptions mentioned earlier, the function M_n is (almost surely, with increasing sample sizes) strictly concave along all directions except those parallel to $\mathbf{1}$, along which it is constant. [Recycled notation.] The constraint locates the unique maximizing $\hat{\alpha}$ along a line of constancy. The measure $\hat{P}_n$ is determined by putting mass

$$\hat{p}_j = \exp(\hat{\beta}_j) = \frac{n_{+j}}{n} h_n(x_j, \hat{\alpha}) \quad \text{at } x_j.$$

That is, $\hat{P}_n$ has density $h_n(\cdot, \hat{\alpha})$ with respect to the empirical measure $\hat{Q}_n$.

Heuristics. The estimator $\hat{P}_n$ is partly parametric and partly nonparametric. The $\hat{\alpha}$ is determined by a finite-dimensional, parametric maximization problem. It determines the density of $\hat{P}_n$ with respect to the nonparametric estimator $\hat{Q}_n$. Limit theorems for $\hat{P}_n$ will follow from the parametric limit theory for $\hat{\alpha}$ and the nonparametric limit theory for $\hat{Q}_n$.

To simplify the analysis let us assume that the proportions are well behaved, in the sense that $\lambda_{ni} \rightarrow \lambda_i > 0$ as $n \rightarrow \infty$, for each i . This assumption could be relaxed. Let $\hat{Q}_{ni}$ denote the empirical measure for the i^{th} sample (mass n_{ij}/n_{i+} on each observation from Q_i). We should then have

$$\hat{Q}_n = \sum_i \lambda_{ni} \hat{Q}_{ni} \rightarrow \sum_i \lambda_i Q_i$$

for some mode of convergence. Call the limit measure Q . For each integrable function f ,

$$Qf = \sum_i \pi_i \lambda_i P(fW_i);$$

the measure Q has density $G(\cdot) = \sum_i \pi_i \lambda_i W_i(\cdot)$ with respect to P . The function $h_n(\cdot, \alpha)$ converges pointwise to

$$h(\cdot, \alpha) = \left(\sum_i \lambda_i \exp(\alpha_i) W_i(\cdot) \right)^{-1}.$$

Notice that $G(\cdot) = 1/h(\cdot, \boldsymbol{\alpha}^*)$, where $\boldsymbol{\alpha}^*$ is determined by

$$\exp(\alpha_i^*) = \pi_i \quad \text{for } i = 1, \dots, s.$$

It would seem reasonable that the limiting behavior of $\hat{\boldsymbol{\alpha}}$ should be obtained by solving the limiting form of the constrained maximization problem. That is, $\hat{\boldsymbol{\alpha}}$ should converge to the $\boldsymbol{\alpha}$ that maximizes

$$M(\boldsymbol{\alpha}) = \boldsymbol{\lambda}'\boldsymbol{\alpha} + Q \log[G(\cdot)h(\cdot, \boldsymbol{\alpha})],$$

subject to the constraint

$$Qh(\cdot, \boldsymbol{\alpha}) = 1.$$

The extra factor G contributes a centering of the log term; each product $G(\cdot)h(\cdot, \boldsymbol{\alpha})$ is bounded away from zero and infinity. This ensures that $M(\boldsymbol{\alpha})$ is well defined for every $\boldsymbol{\alpha}$, without affecting the location of the maximizing value.

Calculation of first and second directional derivatives, in much the same way as before, shows that M is concave. The connectedness assumption implies strict concavity, except along the $\mathbf{1}$ direction, along which it is constant. Modulo $\mathbf{1}$, it has a unique maximizing value, determined by setting all the partial derivatives

$$\begin{aligned} \frac{\partial M}{\partial \alpha_i} &= \lambda_i - Q \left(\frac{\lambda_i \exp(\alpha_i) W_i}{\sum_k \lambda_k \exp(\alpha_k) W_k} \right) \\ &= \lambda_i - \lambda_i \exp(\alpha_i) P(W_i G h(\cdot, \boldsymbol{\alpha})) \end{aligned}$$

to zero. Since $1/G(\cdot) = h(\cdot, \boldsymbol{\alpha}^*)$, these derivatives are zero at $\boldsymbol{\alpha} = \boldsymbol{\alpha}^*$, and the constraint is satisfied:

$$Qh(\cdot, \boldsymbol{\alpha}^*) = P(G/G) = 1.$$

It follows that $\boldsymbol{\alpha}^*$ uniquely solves the limiting constrained maximization problem.

If $\hat{\boldsymbol{\alpha}}$ does converge to $\boldsymbol{\alpha}^*$ then the density $h_n(\cdot, \hat{\boldsymbol{\alpha}})$ of $\hat{P}_n$ with respect to $\hat{Q}_n$ will converge pointwise to $h(\cdot, \boldsymbol{\alpha}^*) = 1/G(\cdot)$. For a fixed integrable f we should then have

$$\hat{P}_n f = \hat{Q}_n(f(\cdot)h(\cdot, \hat{\boldsymbol{\alpha}})) \approx Q(f/G) = Pf.$$

A precise formulation of these heuristic approximations will establish a central limit theorem for $\hat{P}_n$ as an estimator for P .

Asymptotic behavior of $\hat{\boldsymbol{\alpha}}$. Decompose $\hat{\boldsymbol{\alpha}}$ into a sum $\boldsymbol{\alpha}^* + \hat{\boldsymbol{\delta}}/\sqrt{n} + \hat{\epsilon}\mathbf{1}$, where the random vector $\hat{\boldsymbol{\delta}}$ lies in the subspace $\mathcal{D}$ of vectors in $\mathbb{R}^s$ that are orthogonal to $\mathbf{1}$. Constancy of M_n along the $\mathbf{1}$ directions lets us ignore the $\hat{\epsilon}$ in the maximization; the vector $\hat{\boldsymbol{\delta}}$ maximizes the concave function

$$H_n(\boldsymbol{\delta}) = n(M_n(\boldsymbol{\alpha}^* + \boldsymbol{\delta}/\sqrt{n}) - M_n(\boldsymbol{\alpha}^*))$$

over $\boldsymbol{\delta}$ in $\mathcal{D}$. The constraint may be written as

$$(14.3) \quad \exp(\hat{\epsilon}) = \hat{Q}_n h_n(\cdot, \boldsymbol{\alpha}^* + \hat{\boldsymbol{\delta}}/\sqrt{n}).$$

Equivalently, for each integrable f ,

$$(14.4) \quad \hat{P}_n f = \frac{\hat{Q}_n f h_n(\cdot, \boldsymbol{\alpha}^* + \hat{\boldsymbol{\delta}}/\sqrt{n})}{\hat{Q}_n h_n(\cdot, \boldsymbol{\alpha}^* + \hat{\boldsymbol{\delta}}/\sqrt{n})}.$$

The denominator has the interpretation of a normalization factor needed to make $\hat{P}_n$ a probability measure.

The asymptotic behavior of $\hat{\delta}$ will be controlled by a quadratic approximation to H_n . To develop the approximation we decompose the empirical measures into deterministic parts (for which Taylor expansion to quadratic terms is appropriate) plus smaller perturbations due to the empirical processes. Define

$$\begin{aligned}\nu_{ni} &= \sqrt{n_{i+}}(\hat{Q}_{ni} - Q_i) \quad \text{for } i = 1, \dots, s, \\ \nu_n &= \sqrt{n}(\hat{Q}_n - \mathbb{P}\hat{Q}_n) = \sum_i \sqrt{\lambda_{ni}} \nu_{ni}.\end{aligned}$$

Here $\mathbb{P}\hat{Q}_n$ represents the measure $\sum_i \lambda_{ni} Q_i$, which has density

$$G_n(\cdot) = \sum_i \pi_i \lambda_{ni} W_i(\cdot) = 1/h_n(\cdot, \alpha^*)$$

with respect to P . For each f we have a decomposition

$$(14.5) \quad \hat{Q}_n f = P(G_n f) + \frac{1}{\sqrt{n}} \nu_n f.$$

If $P(Gf^2) < \infty$, the random component has an asymptotic normal distribution,

$$(14.6) \quad \nu_n f = \sum_i \sqrt{\lambda_{ni}} \nu_{ni} f \rightsquigarrow N(0, \sigma^2(f)),$$

where

$$\sigma^2(f) = \sum_i \lambda_i (Q_i f^2 - (Q_i f)^2) = P(Gf^2) - \sum_i \lambda_i \pi_i^2 (PW_i f)^2.$$

A similar multivariate central limit theorem would hold for each vector-valued function $\mathbf{f}$ with $P(G|\mathbf{f}|^2)$ finite.

Substituting for $\hat{Q}_n$ in the definition of M_n , using (14.5), we get

$$(14.7) \quad \begin{aligned} H_n(\delta) &= \left[\sqrt{n} \lambda'_n \delta + n P G_n \log(G_n h_n(\cdot, \alpha^* + \delta/\sqrt{n})) \right] \\ &\quad + \sqrt{n} \nu_n \log(G_n h_n(\cdot, \alpha^* + \delta/\sqrt{n})) \end{aligned}$$

Fix δ . Calculation of first and second derivatives, in much the same way as for L_n , shows that the deterministic contribution (the first term on the right-hand side) is of the form $-1/2 \delta' V \delta + o(1)$ as $n \rightarrow \infty$, where V equals $\text{diag}(\lambda_i)$ minus the $s \times s$ matrix whose $(i, j)^{th}$ element is $\pi_i \pi_j \lambda_i \lambda_j P(W_i W_j / G)$. Of course $V \mathbf{1} = \mathbf{0}$, but the connectedness assumption ensures that V acts as a positive definite linear transformation on the subspace $\mathcal{D}$.

The term linear in δ is contributed by the random perturbation (the second term on the right-hand side of (14.7)). Again a Taylor expansion gives

$$\log(G_n h_n(\cdot, \alpha^* + \delta/\sqrt{n})) = \frac{1}{\sqrt{n}} \delta' \mathbf{D}_n(\cdot) + \rho_n(\cdot),$$

where $\mathbf{D}_n$ is an $s \times 1$ vector of uniformly bounded functions,

$$D_{ni}(\cdot) = \frac{\pi_i \lambda_{ni} W_i(\cdot)}{G_n(\cdot)},$$

and ρ_n is a remainder function less than $|\delta|^2/n$ in absolute value. For fixed δ , the contribution to $H_n(\delta)$ from ρ_n converges in probability to zero, because

$$\begin{aligned} \text{var}(\nu_n \rho_n) &= \text{var}\left(\sum_i \sqrt{\lambda_{ni}} \nu_{ni} \rho_n\right) \\ &\leq \sum_i \lambda_{ni} Q_i(\rho_n^2) \\ &\leq |\delta|^4/n^2. \end{aligned}$$

The remainder term $\sqrt{n} \nu_n \rho_n$ is actually of order $O_p(1/\sqrt{n})$.

Collecting together these contributions to H_n we get, for each fixed δ ,

$$H_n(\delta) - \delta' \nu_n \mathbf{D}_n \rightarrow -\frac{1}{2} \delta' V \delta \quad \text{in probability.}$$

The stochastic process on the left-hand side is concave in δ . A simple modification (see Section 6 of Pollard 1990, for example) of a standard result from convex analysis (Theorem 10.8 of Rockafellar 1970) shows that such convergence automatically holds in a stronger sense:

$$(14.8) \quad H_n(\delta) = \delta' \nu_n \mathbf{D}_n - \frac{1}{2} \delta' V \delta + o_p(1) \quad \text{uniformly on compacta.}$$

The $o_p(1)$ term is a random function of δ and n whose supremum over bounded sets of δ converges in probability to zero.

Singularity of V slightly complicates the argument leading from (14.8) to an asymptotic expression for $\hat{\delta}$. A reparametrization will solve the problem. Let J be an $s \times (s-1)$ matrix whose columns span $\mathcal{D}$. Then for θ ranging over $\mathbb{R}^{s-1}$,

$$H_n(J\theta) = \theta' J' \nu_n \mathbf{D}_n - \frac{1}{2} \theta' J' V J \theta + o_p(1) \quad \text{uniformly on compacta.}$$

The $(s-1) \times (s-1)$ matrix $J' V J$ is nonsingular. A small concavity argument (as in Pollard 1990) shows that the $\hat{\theta}$ that maximizes $H_n(J\theta)$ over $\mathbb{R}^{s-1}$ must lie close to the value that maximizes the quadratic approximation, that is,

$$\hat{\theta} = (J' V J)^{-1} J' \nu_n \mathbf{D}_n + o_p(1).$$

Hence

$$(14.9) \quad \hat{\delta} = J(J' V J)^{-1} J' \nu_n \mathbf{D}_n + o_p(1).$$

Let us denote by V^- the matrix multiplying $\nu_n \mathbf{D}_n$; it is a generalized inverse of V .

For each i , the functions D_{ni} converge uniformly to

$$D_i(\cdot) = \frac{\pi_i \lambda_i W_i(\cdot)}{G(\cdot)}.$$

This allows us to invoke a multivariate analogue of (14.6) to show that

$$(14.10) \quad \nu_n \mathbf{D}_n = \nu_n \mathbf{D} + o_p(1) \rightsquigarrow N\left(\mathbf{0}, P(G\mathbf{D}\mathbf{D}') - \sum_i \lambda_i (Q_i \mathbf{D})(Q_i \mathbf{D})'\right).$$

It follows that $\hat{\delta}$ also has an asymptotic normal distribution.

It is possible to solve (14.3) to get a similar asymptotic expression for $\hat{\epsilon}$, and hence for $\hat{\alpha}$. That would lead to an asymptotic normal distribution for $\sqrt{n}(\hat{\alpha} - \alpha^*)$. Such

a calculation will be implicit in the next stage of the argument, which will apply the so-called delta method to (14.4) to derive a central limit theorem for $\hat{P}_n$.

Asymptotic behavior of $\hat{P}_n$. Yet another Taylor expansion gives an approximation that lets us capture the effect of $\hat{\delta}$ on $\hat{P}_n f$.

$$h_n(x, \alpha^* + \delta/\sqrt{n}) = \frac{1}{G_n(x)} - \frac{\delta' \mathbf{D}_n(x)}{\sqrt{n} G_n(x)} + \frac{|\delta|^2}{n G_n} R_n(x, \delta).$$

The remainder function R_n is uniformly bounded on compact sets of δ , in the sense that for each compact K there is a constant C_K such that

$$|R_n(x, \delta)| \leq C_K \quad \text{for all } x, \text{ all } n, \text{ all } \delta \text{ in } K.$$

If f is P -integrable, the contribution from the remainder term can be ignored because

$$(14.11) \quad \frac{|\delta|^2}{n} \mathbb{P} \hat{Q}_n \left| \frac{f R_n}{G_n} \right| \leq \frac{|\delta|^2}{n} P |f R_n| \leq C_K \frac{|\delta|^2}{n} P |f|.$$

Since $|\hat{\delta}| = O_p(1)$, the remainder terms will contribute only a $O_p(1/n)$ to $\hat{P}_n f$.

From (14.5), the leading term in the Taylor expansion contributes

$$(14.12) \quad \hat{Q}_n(f/G_n) = Pf + \frac{1}{\sqrt{n}} \nu_n(f/G_n),$$

which, by (14.6), is asymptotically normal if $P(f^2/G) < \infty$.

The linear term contributes

$$-\frac{1}{\sqrt{n}} \hat{\delta}' \left(P(f \mathbf{D}_n) + \frac{1}{\sqrt{n}} \nu_n(f \mathbf{D}_n/G_n) \right).$$

The ν_n part can be absorbed into the $O_p(1/n)$ term if $P(f^2/G) < \infty$, because

$$(14.13) \quad \text{var } \nu_n(f D_{nj}/G_n) \leq \text{const} \sum_i \lambda_{ni} Q_i(f^2/G_n^2) < \text{const} P(f^2/G).$$

If $P(1/G) < \infty$, similar approximations are valid for the denominator in (14.4). Consequently, if both $P(1/G) < \infty$ and $P(f^2/G) < \infty$ (which also takes care of P -integrability of f),

$$\hat{P}_n f = \frac{Pf + \left(\nu_n(f/G_n) - \hat{\delta}' P(f \mathbf{D}_n) \right) / \sqrt{n} + o_p(1/\sqrt{n})}{1 + \left(\nu_n(1/G_n) - \hat{\delta}' P \mathbf{D}_n \right) / \sqrt{n} + o_p(1/\sqrt{n})}.$$

The right-hand side simplifies to

$$Pf + \frac{1}{\sqrt{n}} \left(\left(\nu_n(f/G_n) - Pf \nu_n(1/G_n) \right) - \hat{\delta}' (P(f \mathbf{D}_n) - Pf P \mathbf{D}_n) \right)$$

plus terms of order $o_p(1/\sqrt{n})$. The coefficient of the linear term in $\hat{\delta}$ might be thought of as a covariance. Substituting from (14.9) for $\hat{\delta}$, then consolidating the lower-order terms, we get

$$(14.14) \quad \sqrt{n}(\hat{P}_n f - Pf) = \nu_n \left(\text{cov}_P(\mathbf{D}, f)' V^{-1} \mathbf{D} + f/G_n - (Pf)/G_n \right) + o_p(1).$$

The right-hand side has an asymptotic normal distribution, by virtue of the multivariate central limit theorem.

Uniformity in f . The preceding calculations are easily extended to provide a functional central limit theorem for $\hat{v}_n = \sqrt{n}(\hat{P}_n - P)$ treated as a stochastic process indexed by a class of functions $\mathcal{F}$.

Let us assume that $\mathcal{F}$ has an envelope $F(\cdot)$, that is, $|f| \leq F$ for each f in $\mathcal{F}$. If F is P -integrable, the analogue of (14.11), with f replaced by F , shows that the remainder terms are of order $O_p(1/n)$ uniformly over $\mathcal{F}$.

If both $P(1/G) < \infty$ and $P(F^2/G) < \infty$, and if the processes indexed by the classes of functions that appear in (14.13) and (14.14) are manageable in the sense of Section 7, then the maximal inequalities from that section can take over the role played by (14.13). (Here the stability results from Section 5 could be applied.) The random contribution to the linear term can again be absorbed into the $o_p(1/\sqrt{n})$, this time uniformly over $\mathcal{F}$. The $o_p(1)$ remainder in (14.14) then also applies uniformly over $\mathcal{F}$, which gives the desired uniform functional central limit theorem.

REMARKS. The concavity argument leading to the central limit theorem for $\hat{\delta}$ is adapted from similar arguments for least absolute deviations regression estimators in Pollard (1990). Almost sure convergence of $\hat{P}_n$ could be established by an even simpler concavity argument, based on pointwise application of a strong law of large numbers, somewhat in the style of Lemma 5.3 of Gill et al (1988). Concavity also explains the success of Vardi's (1985) algorithm—his procedure climbs a concave hill by successive maximizations along coordinate directions.

References

- ALEXANDER, K. S. (1984). Probability inequalities for empirical processes and a law of the iterated logarithm. *Ann. Probab.* **12** 1041-1067.
- ALEXANDER, K. S. (1987a). The central limit theorem for empirical processes on Vapnik-Červonenkis classes. *Ann. Probab.* **15** 178-203.
- ALEXANDER, K. S. (1987b). Central limit theorems for stochastic processes under random entropy conditions. *Probab. Theory Related Fields* **75** 351-378.
- ALEXANDER, K. S. and PYKE, R. (1986). A uniform central limit theorem for set-indexed partial-sum processes with finite variance. *Ann. Probab.* **14** 582-597.
- AMEMIYA, T. (1985). *Advanced Econometrics*. Harvard Univ. Press, Cambridge, Mass.
- ANDERSEN, N. T. (1985a). The central limit theorem for non-separable valued functions. *Z. Wahrsch. verw. Gebiete* **70** 445-455.
- ANDERSEN, N. T. (1985b). The calculus of non-measurable functions and sets. Various Publications Series No. 36, Matematisk Inst., Aarhus Univ.
- ANDERSEN, N. T. and DOBRIĆ, V. (1987). The central limit theorem for stochastic processes. *Ann. Probab.* **15** 164-177.
- ANDERSEN, N. T. and DOBRIĆ, V. (1988). The central limit theorem for stochastic processes. II. *J. Theoret. Probab.* **1** 287-303.
- ARAÚJO, A. and GINÉ, E. (1980). *The Central Limit Theorem for Real and Banach Valued Random Variables*. Wiley, New York.
- BLOOMFIELD, P. and STEIGER, W. L. (1983). *Least Absolute Deviations: Theory, Applications, and Algorithms..* Birkhäuser, Boston.
- BRESLOW, N. and CROWLEY, J. (1974). A large sample study of the life table and product limit estimates under random censorship. *Ann. Statist.* **2** 437-453.
- DELLACHERIE, C. (1972). *Capacités et processus stochastiques*. Springer, Berlin.
- DONOHU, D. (1982). Breakdown properties of multivariate location estimators. Ph. D. qualifying paper, Harvard Univ.
- DONOHU, D. and GASKO, M. (1987). Multivariate generalizations of the median and trimmed mean. I. Technical Report 133, Dept. Statistics, Univ. California, Berkeley.
- DONSKER, M. D. (1952). Justification and extension of Doob's heuristic approach to the Kolmogorov-Smirnov theorems. *Ann. Math. Statist.* **23** 277-281.
- DOOB, J. L. (1953). *Stochastic Processes*. Wiley, New York.
- DUDLEY, R. M. (1966). Weak convergence of probabilities on nonseparable metric spaces and empirical measures on Euclidean spaces. *Illinois J. Math.* **10** 109-126.
- DUDLEY, R. M. (1968). Distances of probability measures and random variables. *Ann. Math. Statist.* **39** 1563-1572.
- DUDLEY, R. M. (1978). Central limit theorems for empirical measures. *Ann. Probab.* **6** 899-929.

- DUDLEY, R. M. (1981). Donsker classes of functions. In *Statistics and Related Topics* (M. Csörgő, D. A. Dawson, J. N. K. Rao and A. K. Md. E. Saleh, eds.) 341-352. North-Holland, Amsterdam.
- DUDLEY, R. M. (1984). A course on empirical processes. *Ecole d'Eté de Probabilités de Saint-Flour, XII-1982. Lecture Notes in Math.* **1097** 1-142. Springer, New York.
- DUDLEY, R. M. (1985). An extended Wichura theorem, definitions of Donsker classes, and weighted empirical distributions. *Probability in Banach Spaces V. Lecture Notes in Math.* **1153** 141-178. Springer, New York.
- DUDLEY, R. M. (1987). Universal Donsker classes and metric entropy. *Ann. Probab.* **15** 1306-1326.
- DUDLEY, R. M. (1989). *Real Analysis and Probability*. Wadsworth, Belmont, Calif.
- GAENSSLER, P. and SCHLUMPRECHT, TH. (1988). Maximal inequalities for stochastic processes which are given as sums of independent processes indexed by pseudo-metric parameter spaces (with applications to empirical processes). Preprint No. 44, Mathematics Inst., Univ. Munich.
- GILL, R., VARDI, Y. and WELLNER, J. (1988). Large sample theory of empirical distributions in biased sampling models. *Ann. Statist.* **16** 1069-1112.
- GINÉ, E. and ZINN, J. (1984). Some limit theorems for empirical processes. *Ann. Probab.* **12** 929-989.
- HOFFMANN-JØRGENSEN, J. (1984). *Stochastic Processes on Polish Spaces*. Unpublished manuscript.
- HUBER, P. J. (1967). The behaviour of maximum likelihood estimates under nonstandard conditions. *Proc. Fifth Berkeley Symp. Math. Statist. Probab.* **1** 221-233. Univ. California Press.
- JAIN, N. C. and MARCUS, M. B. (1978). Continuity of sub-Gaussian processes. In *Probability in Banach Spaces* (J. Kuelbs, ed.). *Advances in Probability* **4** 81-196. Dekker, New York.
- KIM, J. and POLLARD, D. (1990). Cube root asymptotics. *Ann. Statist.* **18** 191-219.
- LEDOUX, M. and TALAGRAND, M. (1989). Comparison theorems, random geometry and some limit theorems for empirical processes. *Ann. Probab.* **17** 596-631.
- LEDOUX, M. and TALAGRAND, M. (1990). *Isoperimetry and Processes in Probability in Banach Spaces*. To appear.
- LOÈVE, M. (1977). *Probability Theory, 4th ed.* Springer, New York.
- MASSART, P. (1986). Rates of convergence in the central limit theorem for empirical processes. *Ann. Inst. H. Poincaré Probab. Statist.* **22** 381-423.
- MEIER, P. (1975). Estimation of a distribution function from incomplete observations. In *Perspectives in Statistics* (J. Gani, ed.) 67-87. Academic, London.
- NOLAN, D. (1989a). Asymptotics for multivariate trimming. Technical Report 151, Dept. Statistics, Univ. California, Berkeley.
- NOLAN, D. (1989b). On min-max majority and deepest points. Technical Report 149, Dept. Statistics, Univ. California, Berkeley.
- NOLAN, D. and POLLARD, D. (1987). U-processes: rates of convergence. *Ann. Statist.* **15** 780-799.
- NOLAN, D. and POLLARD, D. (1988). Functional limit theorems for U-processes. *Ann. Probab.* **16** 1291-1298.
- OSSIANDER, M. (1987). A central limit theorem under metric entropy with L_2 bracketing. *Ann. Probab.* **15** 897-919.
- PISIER, G. (1983). Some applications of the metric entropy condition to harmonic analysis. *Banach Spaces, Harmonic Analysis, and Probability Theory. Lecture Notes in Math.* **995** 123-154. Springer, New York.
- POLLARD, D. (1982). A central limit theorem for empirical processes. *J. Austral. Math. Soc. Ser. A* **33** 235-248.
- POLLARD, D. (1984). *Convergence of Stochastic Processes*. Springer, New York.
- POLLARD, D. (1989). Asymptotics via empirical processes (with discussion). *Statist. Sci.* **4** 341-366.
- POLLARD, D. (1990). Asymptotics for least absolute deviation regression estimators. *Econometric Theory*. To appear.

- POWELL, J. L. (1984). Least absolute deviations estimation for the censored regression model. *J. Econometrics* **25** 303-325.
- PYKE, R. (1969). Applications of almost surely convergent constructions of weakly convergent processes. *Probability and Information Theory. Lecture Notes in Math.* **89** 187-200.
- ROCKAFELLAR, R. T. (1970). *Convex Analysis*. Princeton Univ. Press, Princeton, New Jersey.
- SHORACK, G. R. and WELLNER, J. A. (1986). *Empirical Processes with Applications to Statistics*. Wiley, New York.
- SKOROHOD, A. V. (1956). Limit theorems for stochastic processes. *Theory Probab. Appl.* **1** 261-290.
- STEELE, J. M. (1975). Combinatorial Entropy and Uniform Limit Laws. Ph. D. dissertation, Stanford Univ.
- TALAGRAND, M. (1987). Donsker classes and random geometry. *Ann. Probab.* **15** 1327-1338.
- VAPNIK, V. N. and ČERVONENKIS, A. YA. (1971). On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probab. Appl.* **16** 264-280.
- VARDI, Y. (1985). Empirical distributions in selection bias models. *Ann. Statist.* **13** 178-203.
- ZAMAN, A. (1989). Consistency via type 2 inequalities: A generalization of Wu's theorem. *Econometric Theory* **5** 272-286.