Construction & Extension of Measures

For any finite set $\Omega = \{\omega_1, ..., \omega_n\}$, the "power set" $\mathfrak{P}(\Omega)$ has $|\mathfrak{P}| = 2^n$ elements; it can also be identified with the set of all possible *functions* $a : \Omega \to \{0, 1\}$ by the relation $A = \{\omega : a(\omega) = 1\}$. Set theorists denote the power set by $\mathfrak{P}(\Omega) = \{0, 1\}^\Omega$ or more simply by 2^Ω, even for infinite sets Ω. Last time we considered a number of properties classes of sets $\mathcal{A} \subset 2^\Omega$ might have. A class $\mathcal{A}$ of subsets of Ω is called a:

FIELD if $\Omega \in \mathcal{A}$, $E^c \in \mathcal{A}$ whenever $E \in \mathcal{A}$, and $E_1 \cup E_2 \in \mathcal{A}$ whenever $E_1, E_2 \in \mathcal{A}$;

SIGMA FIELD if $\Omega \in \mathcal{A}$, $E^c \in \mathcal{A}$ whenever $E \in \mathcal{A}$, and $\cup_{i=1}^\infty E_i \in \mathcal{A}$ whenever $E_i \in \mathcal{A},\ i \in \mathbb{N}$;

π-SYSTEM if $E_1 \cap E_2 \in \mathcal{A}$ whenever $E_1, E_2 \in \mathcal{A}$;

λ-SYSTEM if $\Omega \in \mathcal{A}$, $E^c \in \mathcal{A}$ whenever $E \in \mathcal{A}$, and $\cup_{i=1}^\infty E_i \in \mathcal{A}$ whenever $E_i \cap E_j = \emptyset$ and $E_i \in \mathcal{A}$ for all $i \neq j \in \mathbb{N}$.

Note that if $\mathcal{A}_\alpha$ is a (F, σF, $\pi-$S, resp. $\lambda-$S) for each α in any index set (even an uncountable one), then $\cap_\alpha \mathcal{A}_\alpha$ is also a (F, σF, $\pi-$S, resp. $\lambda-$S) (Exercise: show that this is not true for even finite unions). Since also 2^Ω is a (F, σF, $\pi-$S, resp. $\lambda-$S), it follows that for any collection $\mathcal{A}_0 \subset 2^\Omega$ there exists a *smallest* (F, σF, $\pi-$S, resp. $\lambda-$S): namely, the intersection of all (F, σF, $\pi-$S, resp. $\lambda-$S)'s containing $\mathcal{A}_0$. We denote the smallest (F, σF, $\pi-$S, resp. $\lambda-$S) containing $\mathcal{A}_0$ by $\mathcal{F}(\mathcal{A}_0)$, $\sigma(\mathcal{A}_0)$, $\pi(\mathcal{A}_0)$, and $\lambda(\mathcal{A}_0)$, respectively.

For example, if Ω is arbitrary and $\mathcal{A}_0 = \{\{\omega\} : \omega \in \Omega\}$, the singletons, then $\mathcal{F}(\mathcal{A}_0) = \sigma(\mathcal{A}_0) = 2^\Omega$ if Ω is finite, but $\mathcal{F}(\mathcal{A}_0)$ is the finite and co-finite sets, $\sigma(\mathcal{A}_0)$ the countable and co-countable sets if Ω is infinite. $\pi(\mathcal{A}_0)$ and $\lambda(\mathcal{A}_0)$ are just $\mathcal{A}_0$ itself!

For probability and measure theory we need probabilities to be defined for all sets in a sigma field $\mathcal{F}$, so we can compute probabilities for countable unions and intersections; we'd like the luxury of defining the measure on a much smaller collection, either a field $\mathcal{F}_0$ or a collection of sets $\mathcal{A}$ that generates a field $\mathcal{F}_0 = \mathcal{F}(\mathcal{A})$. To do this we need to know that, subject to some obvious consistency conditions, we can always *extend* a pre-measure μ_0 defined only on a field $\mathcal{F}_0$ to *some* measure μ on the sigma field $\mathcal{F} = \sigma(\mathcal{F}_0)$, and we need to prove that this μ is unique—i.e. that, if μ_1 and μ_2 are two measures on $\mathcal{F}$ such that $\mu_1(F) = \mu_2(F)$ for $F \in \mathcal{F}_0$, then also $\mu_1(F) = \mu_2(F)$ for $F \in \mathcal{F}$, i.e., μ_1 and μ_2 agree on the entire sigma field.

It turns out to be easier to show that μ_0 extends uniquely to the λ-system $\lambda(\mathcal{A}_0)$ than it is to show unique extension to the sigma field $\sigma(\mathcal{A}_0)$; luckily, when $\mathcal{A}_0$ is a *field* (or even just a π-system), these are the same:

Theorem (Dynkin's π-λ Theorem). *Let $\mathcal{F}_0$ be a π-system; then $\lambda(\mathcal{F}_0) = \sigma(\mathcal{F}_0)$. (Sketch proof).*

How can we specify μ_0 on a field $\mathcal{F}_0$? Two examples:

1. $\mathcal{A} = \{\{\omega\}\}$: Given any $\{\omega_i\}$ and $\{p_i \geq 0\}$ with $\sum_i p_i = 1$, set $\mu_0(A) = \sum[p_i : \omega_i \in A]$. In fact, this is also μ; it's the only kind of discrete measure there is, and the only kind on a finite or countable set Ω.
2. $\Omega = (-\infty, \infty)$, and $\mathcal{A} = \{(-\infty, b]\}$ for $b \in \mathbb{Q}$. Now $\mathcal{F}_0 = \mathcal{F}(\mathcal{A})$ consists of finite disjoint unions of left-open rational intervals $(a, b]$, including semi-infinite intervals of the form $(-\infty, b]$ and (a, ∞), and $\Omega = (-\infty, \infty)$. The sigma field $\sigma(\mathcal{A})$ is *not* just countable unions of such sets; it is called the "Borel sets" in the real line, and includes all open and closed sets, the Cantor set, and many others. It can be constructed explicitly by transfinite induction (!), but is not easily described. It is *not* every possible subset of $\mathbb{R}$, but it includes every set of real numbers we'll need in this course.

Given any DF $F(x)$ (i.e., right-continuous non-decreasing function on $\mathbb{R}$ with $F(x) \to 0$ as $x \to -\infty$, $F(x) \to 1$ as $x \to +\infty$), we can define a pre-pm μ_0 on $\mathcal{A}$ by setting $\mu_0((-\infty, b]) \equiv F(b)$. If $F = F_d$ is purely discontinuous this just assigns probability $p_i = F(x_i) - F(x_i-)$ to each x_i where $F(x)$ jumps; if $F(x) = F_{ac} = \int_{-\infty}^{x} f(t)\,dt$ is absolutely continuous this just assigns probability $\mu(A) = \int_A f(t)\,dt$ to A (and in fact this is the usual *definition* of that integral!)

How does the extension idea work? Suppose μ_0 is defined on a field $\mathcal{F}_0$, and $\mathcal{F} = \sigma(\mathcal{F}_0)$. Define two new set functions μ^* and μ_* on all subsets of Ω, i.e. on 2^Ω, by:

$$\mu^*(E) \equiv \inf\Big[\sum_{i=0}^{\infty} \mu_0(F_i) : E \subset \bigcup_{i=0}^{\infty} F_i,\ F_i \in \mathcal{F}_0\Big] \qquad \mu_*(E) \equiv 1 - \mu^*(E^c)$$

On reflection it's clear that $\mu_*(E) \le \mu^*(E)$ for each set $E \in 2^\Omega$, and $\mu_*(E) = \mu_0(E) = \mu^*(E)$ for each set $E \in \mathcal{F}_0$; thus there is an obvious well-defined extension of μ_0 to a set function on the μ-*completion*, $\overline{\mathcal{F}}^\mu = \{E \in 2^\Omega : \mu_*(E) = \mu^*(E)\} = \{E \in 2^\Omega : \mu^*(E) + \mu^*(E^c) = 1\}$. It remains to show that: (1) The extension μ is nonnegative and countably additive on $\overline{\mathcal{F}}^\mu$ (an $\epsilon/2^n$ argument); and (2) The σ field $\mathcal{F} = \sigma(\mathcal{F}_0)$ is contained in $\overline{\mathcal{F}}^\mu$ (just show that $\overline{\mathcal{F}}^\mu$ is a σF containing $\mathcal{F}_0$); and (3) The extension to $\mathcal{F}$ is unique (show that for any two extensions μ_1 and μ_2, $\{E \in \mathcal{F} : \mu_1(E) = \mu_2(E)\}$ is a λ-S containing $\mathcal{F}_0$). For details, see Billingsley (1995), pp. 38–41.

Examples:

Let $\Omega = \mathbb{N}$ be the natural numbers $\{1, 2, 3, ...\}$, E and E^c the even and odd ones respectively, and set

$$F = \cup_{k=0}^{\infty}\{2^{2k}+1, ..., 2^{2k+1}\} = \{2, \quad 5, ..., 8, \quad 17, ..., 32, \quad 65, ..., 128, \quad , 257, ..., 512, \quad ...\}$$

and notice that:

1. For $n = 2^{2k}$, the ratio $P_n(F) = \#[F \cap \{1, ..., n\}]/n$ is exactly $P_n(F) = (n-1)/3n$, approximately $1/3$, while for $n = 2^{2k+1}$ it is $P_n(F) = (2n-1)/3n$, approximately $2/3$; thus $P_n(F)$ cannot possibly converge as $n \to \infty$;
2. The even and odd portions of F and F^c, respectively, $A \equiv F \cap E$ and $B \equiv F^c \cap E^c$, both have relative frequencies ranging from $1/6$ to $1/3$, which also cannot converge— in fact, $A = F \cap E$ is exactly the same as the set $2 * (F^c)$, while $B = F^c \cap E^c$ is exactly the same as the set $\{1\} \cup (2 * F - 1)$;
3. $C \equiv (A \cup B)$ however DOES have an asymptotic frequency— in fact, $P_{2n}(C) = 1/2 + 1/2n$ for every n, so $P_n(C) \to 1/2$ as $n \to \infty$;
4. Thus E and C both have well-defined asymptotic frequencies (both are $1/2$), but $A = E \cap C$ does not. Thus, the collection of sets S for which $\lim_{n\to\infty} P_n(S)$ converges is not a field.

In Example 2. above we constructed a measure μ on the σ-algebra $\mathcal{F} = \sigma(\mathcal{F}_0)$ generated by a field $\mathcal{F}_0$ of subsets of the real line $\Omega = \mathbb{R}$. The same approach works more generally, starting with a set assignment μ_0 on a field $\mathcal{F}_0$ or (slightly more generally) on a "semi-algebra" $\mathcal{A}$, a π-system containing Ω for which the complement A^c of each $A \in \mathcal{A}$ can be expressed as a finite disjoint union $\mathcal{A}^c = \cup B_j$ of elements $B_j \in \mathcal{A}$ (example: intervals $(a, b] \subset \mathbb{R}$, with $a < b \in \mathbb{Q}$; rectangles $(a, b] \times (c, d] \subset \mathbb{R}^2$ or, more generally, parallelapipeds $\prod_j (a_j, b_j] \subset \mathbb{R}^n$). Any set function $\mu_0 : \mathcal{A} \to \mathbb{R}$ satisfying (1) $\mu_0(A) \ge 0$, (2) $\mu_0(\Omega) = 1$, and (3) $\mu_0(\cup A_j) = \sum \mu_0(A_j)$ if $A_j \in \mathcal{A}$, $A_i \cap A_j = \emptyset$, *and* $\cup A_j \in \mathcal{A}$, has a unique extension to a probability measure μ on $\sigma(\mathcal{A})$.

In particular this lets us construct Lebesgue measure $m(dx)$ on the unit cube in $\mathbb{R}^n$, so we can explore some of its properties.