Chapter 19

Markov transitions

A Markov system is a stochastic process for which the state of the system at any future time depends only on the present state. Such processes are called *Markov processes*. In the language of conditional expectation of random variables, a Markov process is a family $\{X_t\}$ of random variables (indexed by *time*) with the property that for any measurable function f, any t, and any $h > 0$, $E\big(f(X_{t+h})|X_s,\ s \leqslant t\big) = E\big(f(X_{t+h})|X_t\big)$. This defines a family of *transition functions* relating the distribution of the process at time t to the probability distribution of the process at time $t + h$. The process is *stationary* if such transition functions do not depend on t. Markov processes are generally considered to belong to the realm of probability theory, but some useful results can be derived by purely analytic methods. The main idea is to abstract from the random variables and look at the transition function as a mapping from states to probability measures on the set of states.

The traditional approach uses *stochastic kernels*, which are real-valued functions defined on the Cartesian product of a state space S and a σ-algebra $\mathcal{A}$ of subsets of a possibly different space X, that define probability measures on X for each state s. That is, a kernel k maps $S \times \mathcal{A}$ into $\mathbb{R}$. Presumably this approach is adopted because real-valued functions are not intimidating. We believe that in many applications it is conceptually more natural to think in terms of what we call *Markov transitions*, which are functions from S into $\mathcal{P}(X)$. For instance, in stochastic dynamic programming problems we are given a set S of states and a set A of actions. The probability of tomorrow's state is determined jointly by today's state and action. We are led naturally to consider mappings from $S \times A$ to $\mathcal{P}(S)$, the space of probabilities on the state space S. There are other, independent, reasons to study such mappings. For instance, we may have a model of commodity differentiation, as in A. Mas-Colell [241], in which allocations are mappings from a set S of *traders* to measures on a space of *attributes*. Or we may have a game-theoretic framework in mind, where we are interested in mappings from a space of *players* to *mixed strategies*. One advantage of thinking in terms of mappings from S into $\mathcal{P}(X)$ is that it is easier to generalize to *transition correspondences*, or set-valued transitions. The tradeoff is that it is more work to deal with probability measures than real numbers.

We confine our attention to the reasonably well behaved case where S and X are separable metric spaces. Since most of us are addicted to working with countably additive probabilities, this potentially creates some technical difficulties. Namely, the topological dual of the space of bounded continuous functions on a general separable metric space contains purely finitely additive probabilities. To avoid these, we work with the σ-order continuous dual, which is the space of (countably additive) Borel measures. In the important special case of compact metric spaces, the space of (countably additive) measures is also the norm dual.

Each transition function P gives rise to a *transition operator* $\boldsymbol{P}$ that maps real functions on X to functions on S. The value of $\boldsymbol{P}f$ at s is the expected value of f next period given that today's state is s. This association is reversible: Given a transition operator on functions, we can recover a transition function that generates it (Theorem 19.10). The truly abstract approach to Markov systems studies only these operators. Indeed a *Markov operator* is a positive operator between AM-spaces with units that maps the unit of the domain onto the unit in the range. Section 19.1 presents the most basic results on this class of operators and their study resumes in Section 19.9.

To tie the operator theoretic treatment to our more concrete model with transition functions, we show that a transition P is Borel measurable as a function from the metric space S into the metric space $\mathcal{P}(X)$ if and only if its associated transition operator $\boldsymbol{P}$ carries $B_b(X)$ into $B_b(S)$ (Theorem 19.7). A transition P is continuous if and only if it has the *Feller property*, that is, if and only if its transition operator $\boldsymbol{P}$ carries $C_b(X)$ into $C_b(S)$ (Theorem 19.14). The adjoint $\boldsymbol{P}'$ of the operator $\boldsymbol{P}$ (either norm adjoint or σ-order continuous adjoint, as appropriate) maps probability measures on S to probability measures on X (Theorem 19.9). The adjoint $\boldsymbol{P}'$ is continuous if and only if P is continuous (Theorem 19.14).

When $S = X$ is a compact metrizable space, then $\boldsymbol{P}'$ has fixed points in $\mathcal{P}(S)$ (Theorem 19.18). These are called *invariant measures*. (Even if S is not compact, $\boldsymbol{P}'$ has fixed points in the space of charges (Theorem 19.4). The charge may be purely finitely additive though; see Examples 19.16 and 19.17.) Given an invariant measure μ, a function f is μ-invariant if $f = \boldsymbol{P}f$ μ-a.e. The set of invariant measures is compact and convex. An invariant measure μ is *ergodic* if the only μ-invariant functions are constant μ-a.e. The ergodic measures comprise the set of extreme points of the set of invariant measures (Theorem 19.25).

In many applications (see, e.g., L. E. Blume [49], D. Nachman [258] and D. Duffie, J. Geanakoplos, A. Mas-Colell, and A. McLennan [105]) it is natural and useful to consider set-valued transition functions, or transition correspondences. We use Strassen's Sublinearity Theorem 18.35 to prove the basic result on the existence of selectors (that is, transition functions) that have ergodic measures (Theorem19.31).

We also prove the nice result that continuous transitions (with full support pointwise) correspond in a natural way to random functions. That is, given a continuous Markov transition P from S to $X = [0, 1]$, there is a Borel probability