

Math 525: Lecture 18

March 15, 2018

Consider a gambler playing roulette. He has decided that he will place a bet on 17 at time τ . We can model the game of roulette as a stochastic process with X_n being the result of the n -th spin. In a perfect world, the gambler would choose

$$\tau = \inf\{n \geq 0: X_{n+1} = 17\},$$

the time just before the ball lands on 17. Of course, this is cheating! It would mean that the gambler can look into the future. Indeed, the gambler can base his decision to bet at time $\tau = n$ only on *information available at time n* (e.g., the results of the spins at the previous times X_0, X_1, \dots, X_n). But how can we express the statement “information available before time n ” mathematically?

1 Filtrations

For our discussion, let $(X_n)_{n \in T}$ be a Markov chain on a countable state space $S \subset \mathbb{R}$. We will not assume that the state space S is finite, nor that the Markov chain is stationary. Note that if T is finite, S^T is countable and hence we can take $(\Omega = S^T, \mathcal{F} = 2^\Omega, \mathbb{P})$ to be the probability space (for some \mathbb{P}).

Definition 1.1. Let \mathcal{F} be a σ -algebra and T be an index set. A *filtration* is a sequence $(\mathcal{F}_t)_{t \in T}$ of σ -algebras such that $\mathcal{F}_s \subset \mathcal{F}_t \subset \mathcal{F}$ for all $s \leq t$.

While the above definition is general enough for continuous index sets, we will focus on the case when T is discrete (e.g., $T = \{0, 1, \dots\}$ or $T = \{0, 1, \dots, N\}$) since we are motivated by Markov chains. The following illustrates that \mathcal{F}_t can be thought of as all information accrued up to time t .

Example 1.2. A coin is flipped twice. This corresponds to the sample space

$$\Omega = \{HH, HT, TH, TT\}.$$

Since this a finite sample space, we can take $\mathcal{F} = 2^\Omega$ to be our σ -algebra. Define

$$\begin{aligned}\mathcal{F}_0 &= \{\emptyset, \Omega\} \\ \mathcal{F}_1 &= \{\emptyset, \Omega, \{HH, HT\}, \{TH, TT\}\} \\ \mathcal{F}_2 &= 2^\Omega.\end{aligned}$$

Note that $(\mathcal{F}_n)_{n \in \{0,1,2\}}$ defines a filtration since $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \mathcal{F}_2 = \mathcal{F}$. Now, let $\omega = (\omega_1, \omega_2) \in \Omega$ be the result of the coin flips. After we have observed the first coin flip, we are able to answer whether

$$\omega \in \{HH, HT\}$$

(i.e., whether the first coin flip was heads). However, we cannot answer whether

$$\omega \in \{HH\}$$

since this would require intimate knowledge of the second coin flip. More generally, at time n , we can answer any question of the form

$$\omega \in F \quad \text{where } F \in \mathcal{F}_n.$$

In this sense, \mathcal{F}_0 is the “dumb” sub σ -algebra, since it carries no information whatsoever!

As you might have noticed in the above example, it is quite cumbersome to write out filtrations by hand (and our example only had two coin flips!) There is a better way:

Definition 1.3. We say a random variable X is *measurable* with respect to some σ -algebra \mathcal{G} (or simply “ \mathcal{G} measurable”) if

$$X^{-1}(B) \in \mathcal{G} \quad \text{for } B \in \mathcal{B}(\mathbb{R}).$$

Let $\{X_\alpha\}_\alpha$ be family of random variables on the same probability space. Then *the σ -algebra generated by the family* is the smallest σ -algebra \mathcal{G} such that for each α , X_α is \mathcal{G} measurable. We will often denote this σ -algebra by $\sigma(\{X_\alpha\}) \equiv \mathcal{G}$.

If the family $\{X_0, \dots, X_n\}$ is finite, we’ll forgo the curly braces and write $\sigma(X_0, \dots, X_n)$.

Example 1.4. Returning to the coin flipping example, let

$$X_n = \begin{cases} 1 & \text{if the } n\text{-th coin flip is heads} \\ 0 & \text{otherwise.} \end{cases}$$

Let B be a Borel set. Note that there are four possibilities:

- $1 \in B$ and $0 \notin B$ in which case $X_1^{-1}(B) = \{HH, HT\}$,
- $0 \in B$ and $1 \notin B$ in which case $X_1^{-1}(B) = \{TH, TT\}$,
- $0, 1 \in B$ in which case $X_1^{-1}(B) = \Omega$.
- $0, 1 \notin B$ in which case $X_1^{-1}(B) = \emptyset$, and

Therefore,

$$\sigma(X_1) = \mathcal{F}_1 = \{\emptyset, \Omega, \{HH, HT\}, \{TH, TT\}\}.$$

Similarly, you can check that $\sigma(X_1, X_2) = \mathcal{F}_2$.

Definition 1.5. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and $\mathbb{F} = (\mathcal{F}_t)_{t \in T}$ be a filtration on \mathcal{F} where T is some index set. Then, $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ is called a *filtered probability space*.

2 Stopping times

For the remainder of this section, we will assume a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t \in T}, \mathbb{P})$ with index set $T = \{0, 1, \dots\}$ or $T = \{0, 1, \dots, N\}$ (though the ideas extend easily to general index sets). We will also assume a Markov chain on this space (i.e., $(X_n)_{n \in T}$).

Armed with the notion of a filtration, we can try to outline some fair rules that prevent our greedy gambler from being able to “look into the future”.

Definition 2.1. τ is a *stopping time* if

1. it is a random variable taking values in $T \cup \{\infty\}$ and
2. for each $n \in T$, $\{\tau = n\} \equiv \{\omega : \tau(\omega) = n\}$ is a set in \mathcal{F}_n .

The first property just tells us that τ is a random time (in the roulette example, ∞ corresponds to the gambler choosing to never place his bet!). The second property tells us that the decision to “stop” at time n (e.g., place the bet at time n) can only depend on the information witnessed up to time n .

Proposition 2.2. A random variable τ taking values in $T \cup \{\infty\}$ is a stopping time if and only if

$$\{\tau \leq n\} \in \mathcal{F}_n \quad \text{for each } n \in T. \quad (1)$$

Proof. Suppose τ is a stopping time. Note that

$$\{\tau \leq n\} = \bigcup_{m=0}^n \{\tau = m\}.$$

By definition, $\{\tau = m\} \in \mathcal{F}_m$ and $\mathcal{F}_m \subset \mathcal{F}_n$ whenever $m \leq n$. Therefore, $\{\tau \leq n\}$ is a finite union of sets in \mathcal{F}_n , and hence $\{\tau \leq n\} \in \mathcal{F}_n$.

Suppose (1). Note that

$$\{\tau = n\} = \{\tau \leq n\} \setminus \{\tau < n\} = \{\tau \leq n\} \cap \{\tau < n\}^c.$$

If $n = 0$, $\{\tau < 0\}^c = \emptyset^c = \Omega$. Otherwise, $\{\tau < n\} = \{\tau \leq n-1\} \in \mathcal{F}_{n-1} \subset \mathcal{F}_n$. Either way, the above is the intersection of elements of \mathcal{F}_n . \square

Intuitively, the statement $\{\tau \leq n\} \in \mathcal{F}_n$ says that the decision to stop before or exactly at time n can be divined knowing all information up to time n . Indeed, at time n , we can approach the gambler and ask, “have you placed your bet yet?” If the gambler says yes, then $\tau \leq n$ (this is not the same as, “did you place your bet for the next spin?”)

Exercise 2.3. Use the properties of σ -algebras to show that if τ is a stopping time, then for each n , the sets $\{\tau < n\}$, $\{\tau \geq n\}$, and $\{\tau > n\}$ are in \mathcal{F}_n .

Let’s see some examples of stopping times:

Example 2.4.

1. The constant random variable $\tau = n$ is a stopping time since

$$\{\tau = m\} = \begin{cases} \Omega & \text{if } m = n \\ \emptyset & \text{if } m \neq n. \end{cases}$$

2. Let $\mathcal{F}_n = \sigma(X_0, \dots, X_n)$ for $n = 0, 1, 2, \dots$ where $(X_n)_n$ is a sequence of random variables on the same probability space. Let $B = (-1, 1)$. Then,

$$\tau_B = \inf \{n \geq 0 : X_n \notin B\}$$

where we use the convention $\inf \emptyset = \infty$. Then, τ_B is a stopping time, called the *first exit time*. B can be, more generally, any Borel set. τ_B is a stopping time since

$$\begin{aligned} \{\tau_B = n\} &= \{X_0 \in B, \dots, X_{n-1} \in B, X_n \notin B\}. \\ &= \{X_0 \in B\} \cap \dots \cap \{X_{n-1} \in B\} \cap \{X_n \notin B\} \\ &= X_0^{-1}(B) \cap \dots \cap X_{n-1}^{-1}(B) \cap X_n^{-1}(B^c) \\ &= \underbrace{X_0^{-1}(B)}_{\in \mathcal{F}_0} \cap \dots \cap \underbrace{X_{n-1}^{-1}(B)}_{\in \mathcal{F}_{n-1}} \cap \underbrace{(X_n^{-1}(B))^c}_{\in \mathcal{F}_n}. \end{aligned}$$

3. If τ is a stopping time and $m \geq 0$ is an integer, $\tau + m$ is a stopping time. To see that this is a stopping time, note that

$$\{\tau + m = n\} = \{\tau = n - m\}.$$

4. If τ_1 and τ_2 are stopping times, so are $\tau_1 \wedge \tau_2 \equiv \min(\tau_1, \tau_2)$ and $\tau_1 \vee \tau_2 \equiv \max(\tau_1, \tau_2)$ (left as an exercise).

Given a stopping time τ , we would like to define the value of the Markov chain $(X_n)_{n \in T}$ at time τ . We do this below.

Definition 2.5. Let τ be a random variable taking values in T (not necessarily a stopping time). Define $X_\tau : \Omega \rightarrow \mathbb{R}$ by

$$X_\tau(\omega) = X_{\tau(\omega)}(\omega).$$

Proposition 2.6. X_τ defined above is indeed a random variable.

Proof. This follows from the fact that

$$X_\tau^{-1}(B) \equiv \{\omega : X_\tau \in B\} = \bigcup_{n \geq 0} (\{\tau = n\} \cap \{X_n \in B\}).$$

□

There is a potential problem: a stopping time τ is allowed to take infinite values. However, we can still use the above to define the value of the Markov chain at time τ . For example,

$$Y = X_\tau I_{\{\tau < \infty\}}$$

is a perfectly legitimate definition (other ways to define this exist, and you may be inclined to use another definition depending on the context).

Similarly, given a stopping time τ , we would like to define the information up to time τ . Of course, we cannot simply set this to be $\sigma(X_0, \dots, X_\tau)$ since this expression is meaningless: τ is not a positive integer, it is a random variable. Instead, we do the following:

Definition 2.7. Given a stopping time τ , let

$$\mathcal{F}_\tau = \{\Lambda \in \mathcal{F} : (\Lambda \cap \{\tau \leq n\}) \in \mathcal{F}_n \text{ for } n \in T\}.$$

Proposition 2.8. *Equivalently,*

$$\mathcal{F}_\tau = \{\Lambda \in \mathcal{F} : (\Lambda \cap \{\tau = n\}) \in \mathcal{F}_n \text{ for } n \in T\}.$$

Proof. To see this, note that

$$\Lambda \cap \{\tau \leq n\} = \bigcup_{m=0}^n \Lambda \cap \{\tau = m\}$$

and

$$\Lambda \cap \{\tau = n\} = (\Lambda \cap \{\tau \leq n\}) \setminus (\Lambda \cap \{\tau < n\}).$$

□

Proposition 2.9. *Let τ and τ' be stopping times. Then,*

1. \mathcal{F}_τ is a σ -algebra.
2. Both τ and $X_\tau I_{\{\tau < \infty\}}$ are \mathcal{F}_τ measurable (assuming X_n is \mathcal{F}_n measurable for each n).
3. If $\tau \leq \tau'$, then $\mathcal{F}_\tau \subset \mathcal{F}_{\tau'}$.

Proof.

1. Clearly, \emptyset and Ω are in \mathcal{F}_τ . Now, suppose $\Lambda \in \mathcal{F}_\tau$. Then, for each $n \in T$,

$$\Lambda^c \cap \{\tau \leq n\} = (\Lambda \cup \{\tau > n\})^c \in \mathcal{F}_n.$$

Lastly, let $\Lambda_1, \Lambda_2, \dots \in \mathcal{F}_\tau$. Then

$$\bigcup_{n \geq 0} (\Lambda_n \cap \{\tau \leq n\}) = \left(\bigcup_{n \geq 0} \Lambda_n \right) \cap \{\tau \leq n\} \in \mathcal{F}_n.$$

2. This follows from

$$\{X_\tau I_{\{\tau < \infty\}} = i\} \cap \{\tau = n\} = \{X_n = i\} \cap \{\tau = n\}$$

and

$$\{\tau \leq m\} \cap \{\tau \leq n\} = \{\tau \leq m \wedge n\} \in \mathcal{F}_{m \wedge n} \subset \mathcal{F}_n.$$

3. Since $\tau \leq \tau'$, if $\tau'(\omega) \leq n$, then $\tau(\omega) \leq n$. Therefore,

$$\Lambda \cap \{\tau' \leq n\} = \underbrace{\Lambda \cap \{\tau \leq n\}}_{\in \mathcal{F}_n} \cap \underbrace{\{\tau' \leq n\}}_{\in \mathcal{F}_n}.$$

□