Short Guide to Exponential Competing Hazards

In these notes $Exp(\gamma)$ refers to an Exponential random variable with parameter γ and $\exp\left[\gamma\right] = e^{\gamma}$ is the exponential function.

By: Sean L. Wu (slwu89@berkeley.edu)

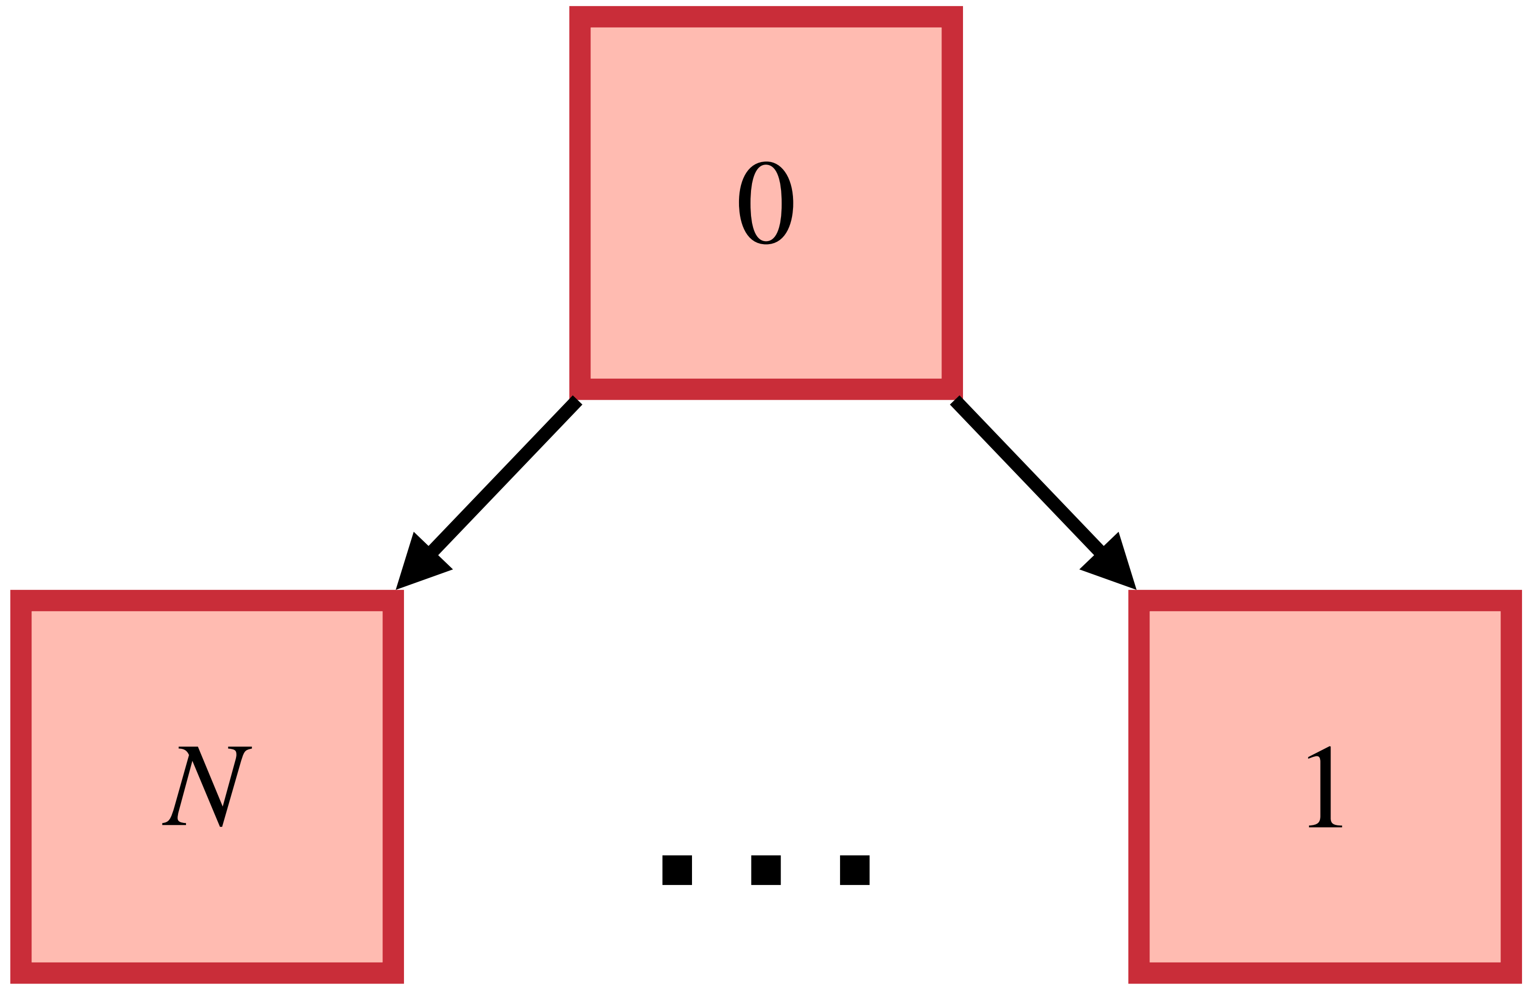

Modelling with competing hazards is useful when you are dealing with continuous time stochastic models of populations, or when doing survival analysis. The diagram to the left is the general case, where an individual starts in state 0. There are *N* different *events* (denoted by edges) which can happen to that individual, such that event *i* will cause that individual to transition from state 0 to state *i*, where *i* is in 1,…,N.

We assume that for an individual in state 0, only a single event can occur to them, and once it occurs, they will change their state. In an *SIR* model, for an individual in state *S*, there is only one event that can happen to them, the infection event which will change their state to *I*.

Each event *i* has a hazard function (sometimes called rate function). Let's assume all events occur after an Exponentially distributed amount of time, so that the hazard function is the parameter of the corresponding Exponential distribution. Call the hazard function for event *i*: λ_i. Recall that the mean of that Exponential distribution will be $1/\lambda_i$.

Let's start with the simple case with a single event, the situation for a person in state *S* in an *SIR* model. The infection event has a hazard function equal to $\lambda_{S \to I}$.

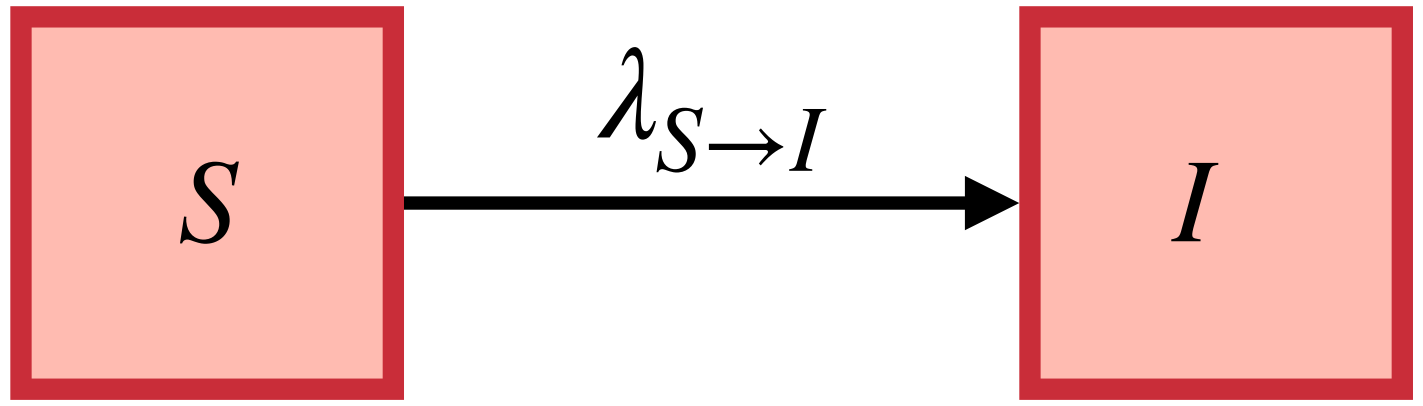

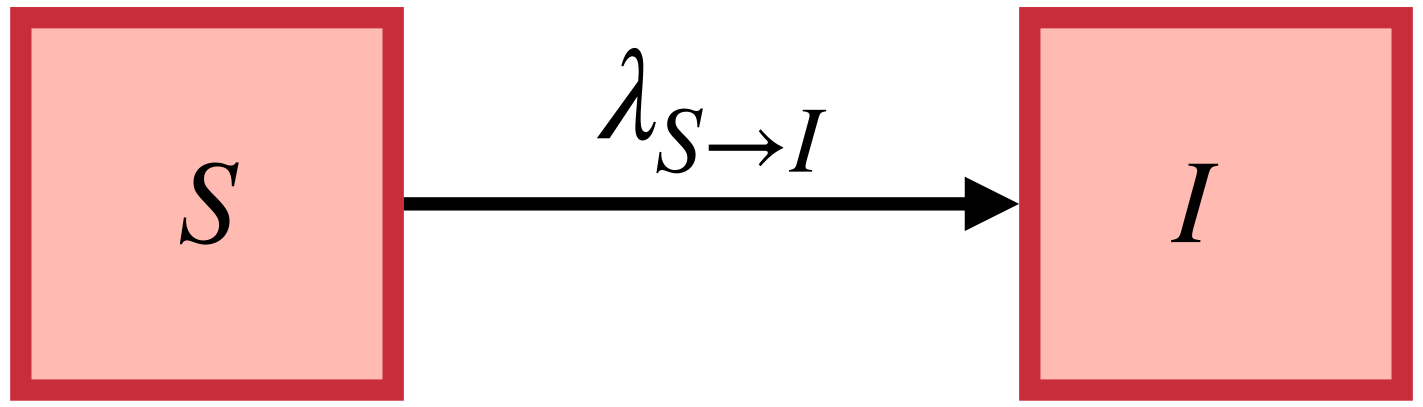

Now we know that this individual will wait for an Exponentially distributed amount of time before becoming infected. If $\tau_{S\to I}$ is the random variable that describes how long before they become infected, then $\tau_{S\to I} \sim Exp(\lambda_{S\to I})$. If you wanted to simulate this process, you'd draw an Exponential random variable to represent the time at which this person becomes infected, and update their state from *S* to *I* at that time.

Let's say, however, that we don't know the hazard $\lambda_{S\to I}$ and need to calculate it from observed probabilities. Say we observe this person over a time interval t. Because the infection event follows an exponential distribution, can compute the probability that they become infected by t, which is $P(I; t)$:

- $P(I; t) = 1 - \exp\left[-t\lambda_{S\to I}\right]$, that is, the probability the event occurred by Δt is the CDF of an exponential random variable.
- $P(S; t) = 1 - P(I; t) = \exp\left[-t\lambda_{S\to I}\right]$, that is, the probability the event hasn't occurred by Δt, which means they are still in state *S*, is the survival function of an exponential random variable.

If all we knew was one of the two probabilities $P(I; t)$ or $P(S; t)$, and t, we could compute $\lambda_{S\to I}$ by solving for it in the equations above. One key thing here is that as $t \to \infty$, for any non-zero $\lambda_{S\to I}$, $P(S; t)$ will equal 1. Intuitively this is because if only one thing can happen, and it happens at a nonzero rate, as time becomes infinite, the probability it will happen becomes 1.

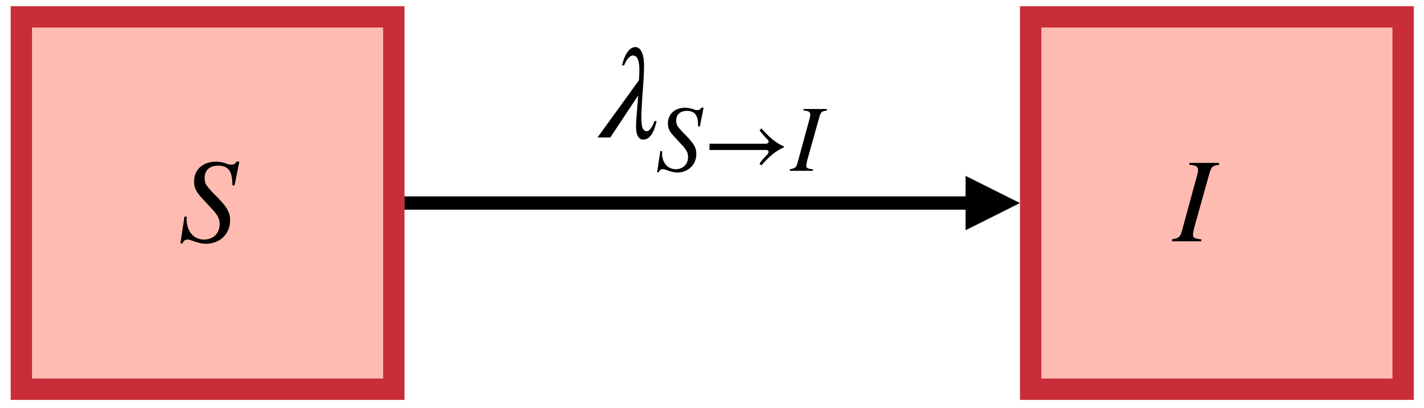

At this point we should realize that what we've made is a very simple 2 state continuous-time Markov chain. The "stochastic rate matrix" (a.k.a., infinitesimal generator matrix) is this:

$$Q = \begin{bmatrix} -\lambda_{S\to I} & \lambda_{S\to I} \\ 0 & 0 \end{bmatrix}$$

The (i, j) element of the transition probability matrix $P(t)$ gives the probability to move from state *i* to state *j* in a time interval t, and is given below:

$$P(t) = \begin{bmatrix} -\exp\left[-t\lambda_{S\to I}\right] & 1 - \exp\left[-t\lambda_{S\to I}\right] \\ 0 & 1 \end{bmatrix}$$

In a technical sense, $P(t) = e^{Qt}$. To simulate from the simple process, if someone started in *S* they'd draw an exponentially distributed time to infection $\tau_{S\to I}$. At that time, they would change state to *I* and stay there forever.

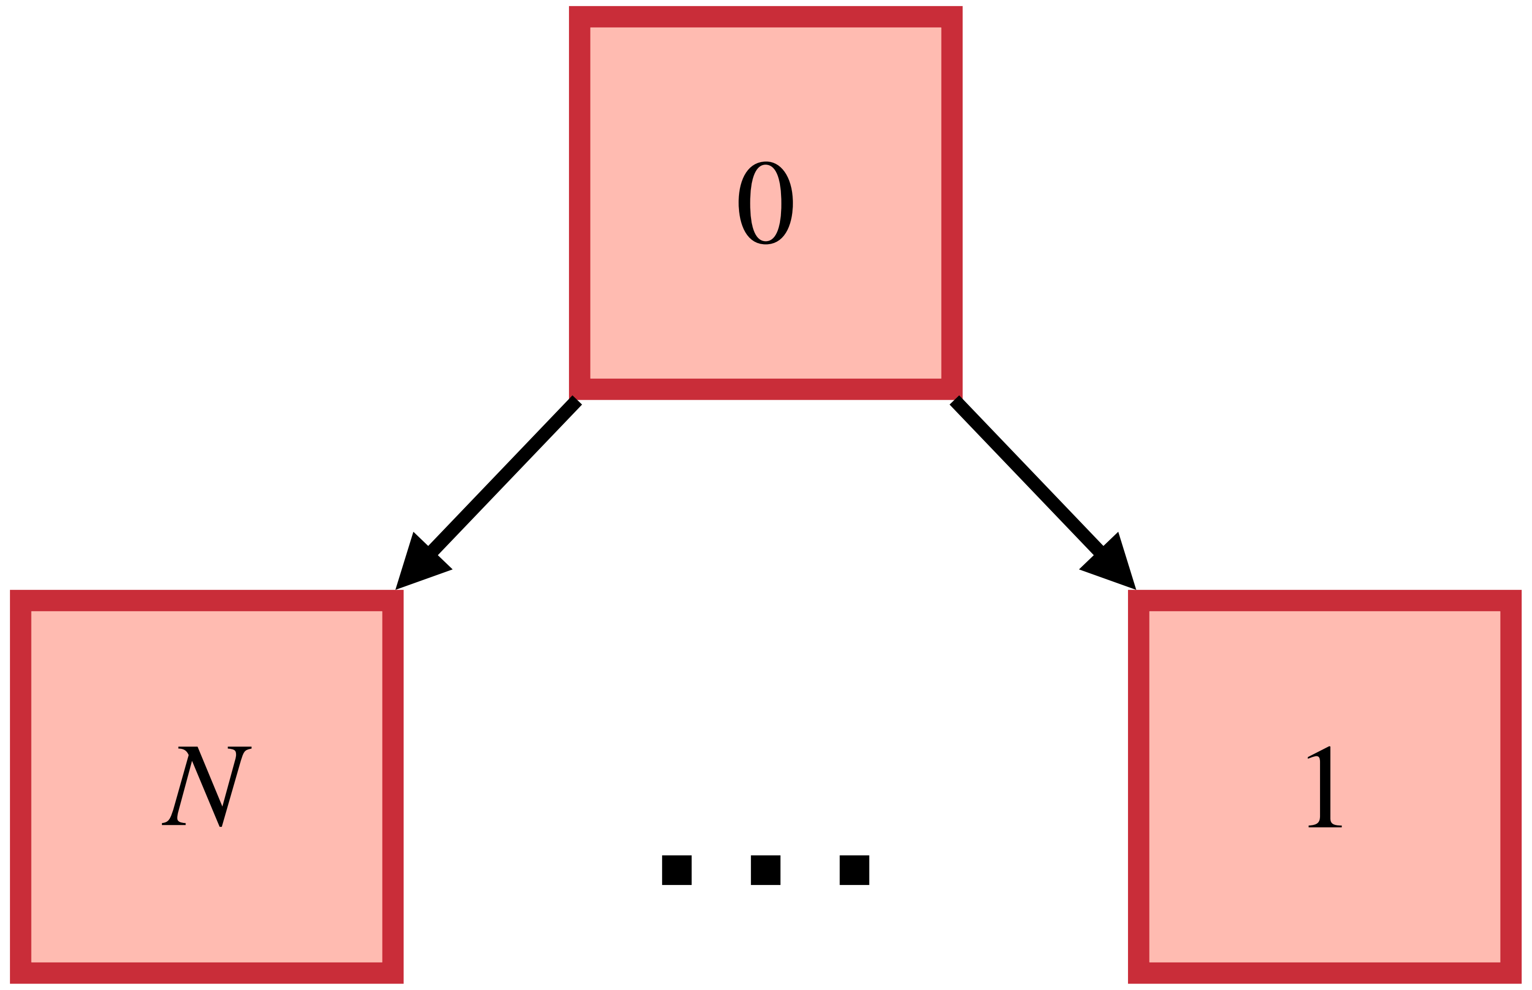

At this point it's a good time to return to the general competing hazards (sometimes called competing risks) framework, with *N* competing events. Again, to keep the math clean, we'll assume everything is Exponentially distributed, which is the same as assuming that the process is a continuous-time Markov chain.

In order to simulate from the model, we need to either specify, or calculate from other quantities, the hazard function for each event: $\lambda_1, ..., \lambda_N$. The interpretation for λ_i is this: *if* event *i* is the event which happens (which is the same as saying event *i* happens before any other event $k \neq i$), it occurs after a waiting time given by a Exponential random variable with parameter λ_i (with a mean delay of $\frac{1}{\lambda_i}$).

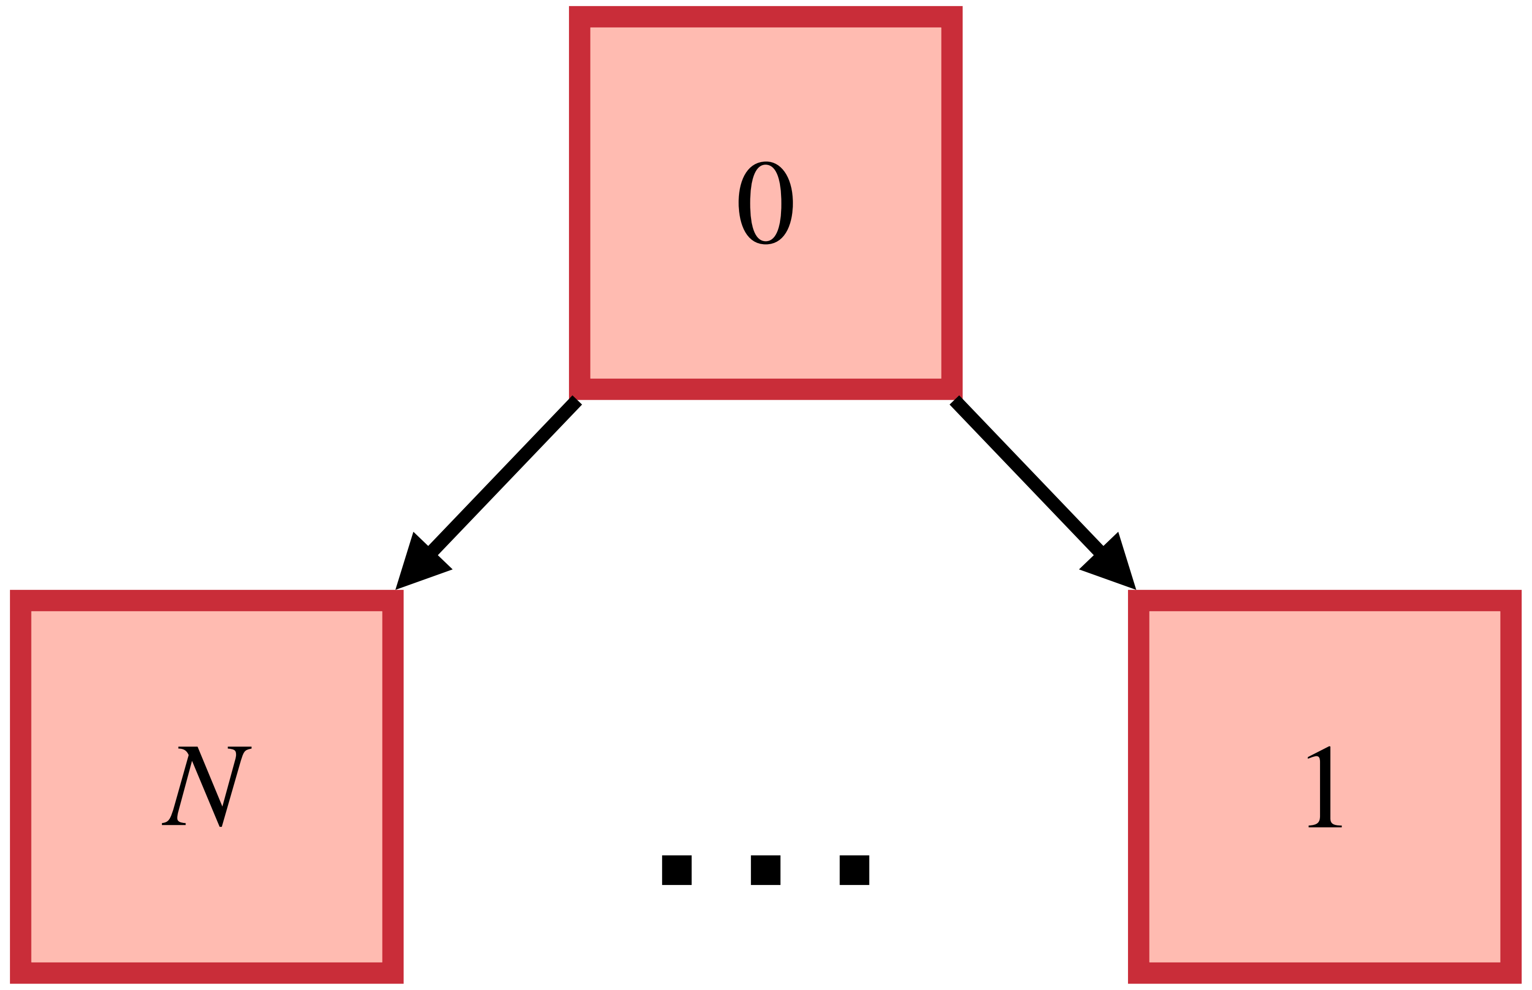

Let's work out some probabilities associated with this model. In order to simulate what happens to an individual in state *0*, we can do two things, which, as it turns out, are equivalent:

1. Sample $\tau_1, \ldots, \tau_N$ random numbers from the Exponential distributions associated with each event with hazards $\lambda_1, \ldots, \lambda_N$. Let $\tau_i = \min(\tau_1, \ldots, \tau_N)$ be the time at which that event which attains the minimum happened.
2. Realize that the previous method really only needs to sample 2 random values. τ is the time at which *anything* happens to that individual, and i is the event which happens at time τ.
   - The time at which *anything* happens to the individual will be the sum of all the possible hazards: $\lambda = \sum_i \lambda_i$ so to sample we do $\tau \sim Exp(\lambda)$. This is because the minimum of a set of exponential random variables is another exponential random variable whose rate parameter is equal to the sum of all those rates.
   - The probability that event i is the one that happened is just a discrete distribution: $P(i) = \frac{\lambda_i}{\sum_j \lambda_j}$

This may seem tangential, but it's intimately related to constructing the transition probability matrix for this process, as we'll see on the next page.

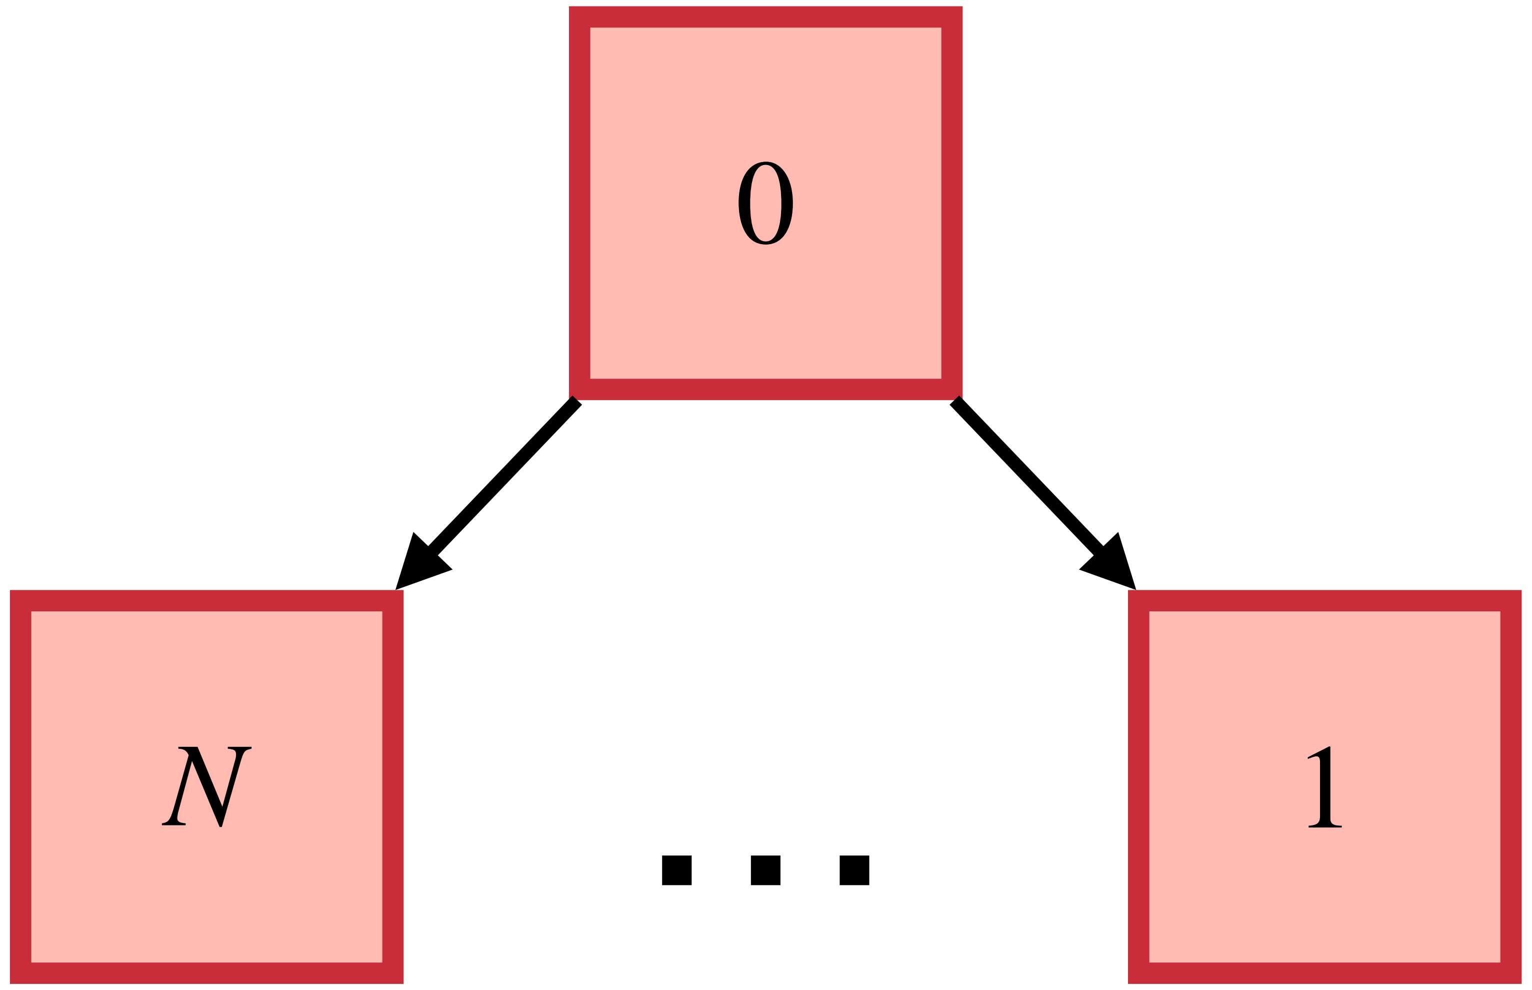

Now we have an *N+1* state (*0,1,…,N*) Markov chain. For any time t, we can compute the probability that nothing has happened yet, which, phrased another way, is the probability we are still in state *0* and time t: $P(0; t)$, easily by using the rule that hazards are additive:

- $$P(0; t) = \exp\left[-t\left(\sum_i \lambda_i\right)\right]$$
- Likewise, the probability that any event happens by time t, $P(\overline{0}; t)$ is:

$$P(\overline{0}; t) = 1 - P(0; t) = 1 - \exp\left[-t\left(\sum_i \lambda_i\right)\right]$$

This should remind us of the simple 2-state example earlier, where we've simply lumped *N* states into a single state.

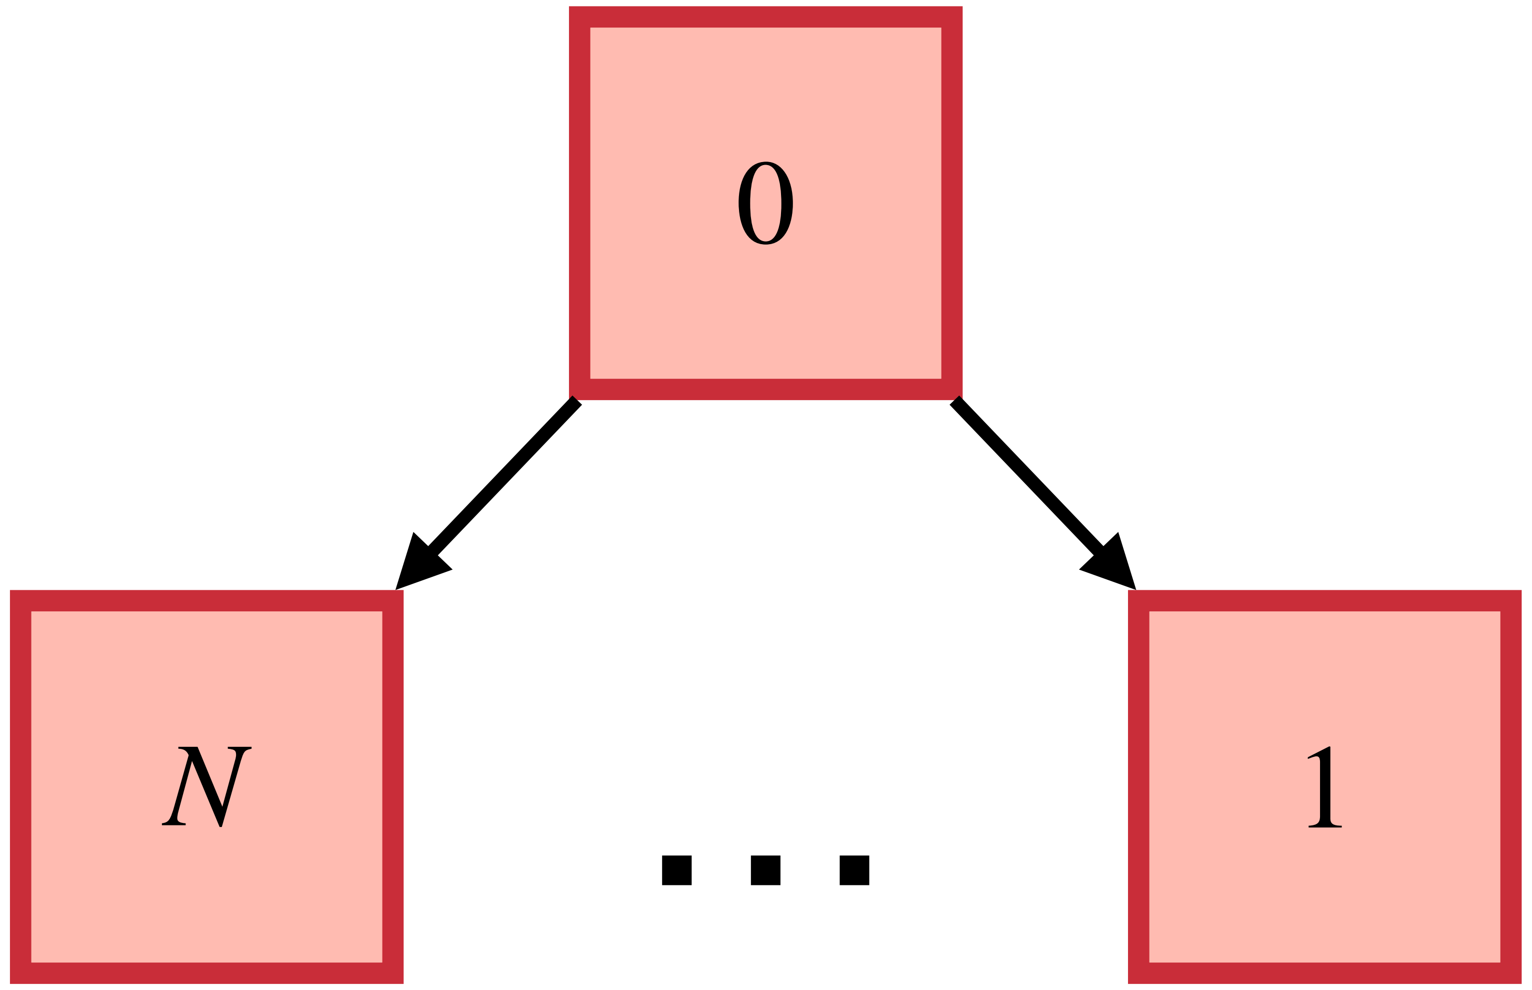

Finally we're ready for the rest of the probabilities. We'll continue to assume that we have access to the entire vector $\lambda_1, \ldots, \lambda_N$ to get the probabilities of each individual event happening, but note that if we only had their sum λ, and conditional probabilities of what happened given some event happened, $P(j \,|\, \bar{0})$, then $\lambda_j = \lambda P(j \,|\, \bar{0})$. Because $\sum_j P(j \,|\, \bar{0}) = 1$ (it is a probability distribution), $\lambda = \sum_j \lambda_j$ so the method is self-consistent.

First we give the stochastic rate matrix:

$$Q = \begin{bmatrix} -\lambda & \lambda_1 & \ldots & \lambda_N \\ 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Then the transition probability matrix:

$$P(t) = \begin{bmatrix} \exp\left[-t\lambda\right] & \left(\frac{\lambda_1}{\lambda}\right) 1 - \exp\left[-t\lambda\right] & \ldots & \left(\frac{\lambda_N}{\lambda}\right) 1 - \exp\left[-t\lambda\right] \\ 0 & 1 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 1 \end{bmatrix}$$