Computational Modelling
Continuous Time Markov Chains

Luca Bortolussi[1]

[1] Dipartimento di Matematica e Geoscienze
Università degli studi di Trieste

Office 328, third floor, H2bis
luca@dmi.units.it

Trieste, Summer Semester, 2017

$T_i, \quad i \in I \qquad T_i \sim \mathrm{Exp}(\lambda_i)$ RACE CONDITION

$$\left.\begin{aligned} T &= \min T_i \\ K &= \underset{j}{\arg\min}\, T_j \end{aligned}\right\} \quad \left.\begin{aligned} &T, K \text{ are independent} \\ &T \sim \mathrm{Exp}\Big(\sum_j \lambda_j\Big) \\ &\Pr(K = k) = \frac{\lambda_k}{\sum_j \lambda_j} \end{aligned}\right\}$$

Outline

1. Continuous Time Markov Chains
   - Main concepts
   - Poisson Process
   - Time-inhomogeneous rates
2. Population Continuous Time Markov Chains
3. Simulation
   - SSA
   - Next Reaction Method
   - τ-leaping

CTMC: DEFINITION

Let S be finite or countable. A continuous-time random process $(X_t)_{t\geq 0} = \{X_t \mid t \geq 0\}$, with values in S, is a family of random variables $X_t : (\Omega, \mathcal{S}) \to (S, 2^S)$ that are *right-continuous* w.r.t. t.

Therefore, X_t (or $X(t)$) has *cadlag* sample paths.

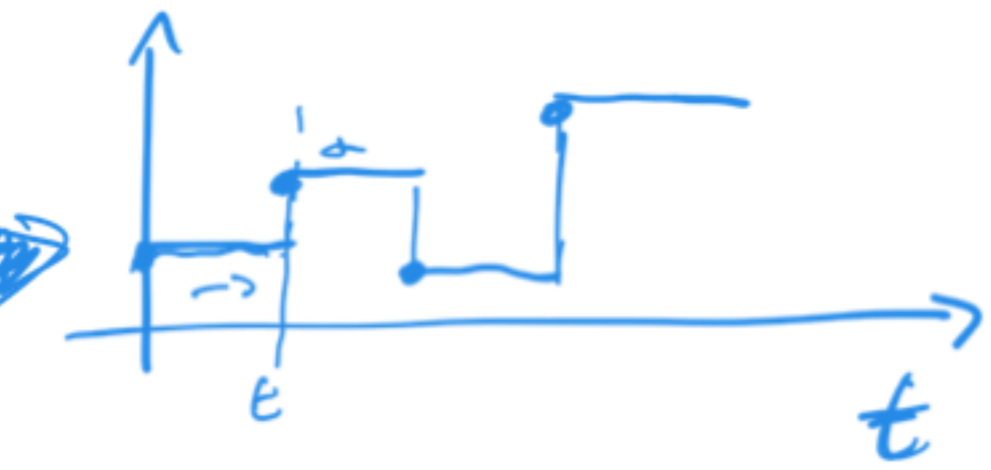

Right continuous processes are determined by their *finite-dimensional distributions*. $\forall t_0, \dots, t_n \in \mathbb{R}_{\geq 0}$ $(X_{t_0}, \dots, X_{t_n})$

A Continuous Time Markov Chain is a right-continuous continuous-time random process satisfying the memoryless condition: for each n, t_i and s_i:

$t_0 \quad t_1 \quad \dots \quad t_{n-1} \quad t_n$

$$\mathbb{P}(X_{t_n} = s_n \mid X_{t_0} = s_0, \dots, X_{t_{n-1}} = s_{n-1}) = \mathbb{P}(X_{t_n} = s_n \mid X_{t_{n-1}} = s_{n-1}).$$

$\forall n \geq 0, \forall t_0, \dots, t_n \in \mathbb{R}_{\geq 0}$

CTMC: RACE CONDITION

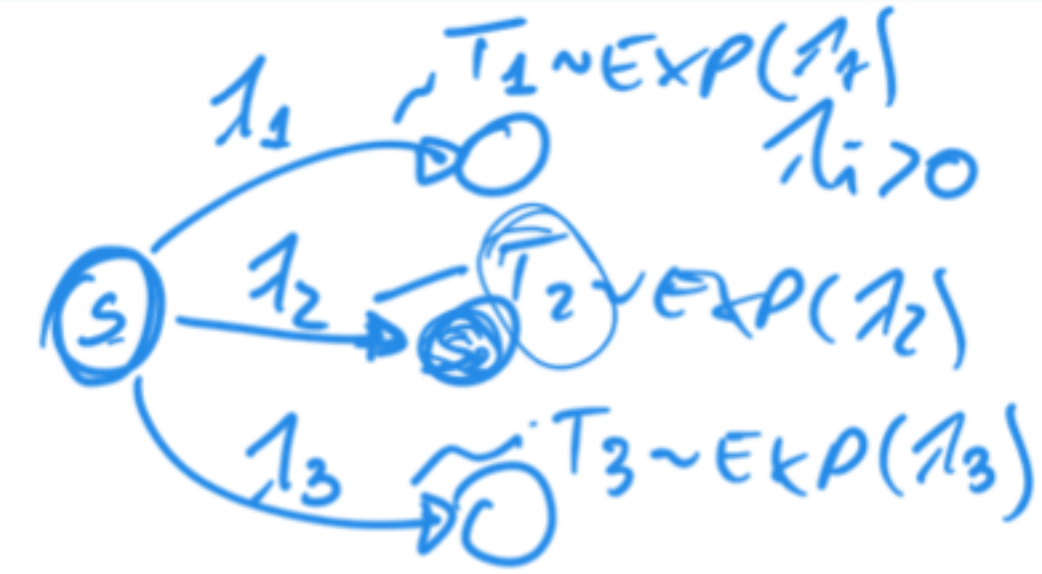

A CTMC on a state space S can be seen as a labelled graph. Each edge takes some time to be crossed, exponentially distributed with the rate labelling the edge.

In each state, there is a race condition between the different exiting edges: the fastest is traversed.

The memoryless property follows from that of the exponential distribution.

INFINITESIMAL GENERATOR

A Q-matrix is the $|S| \times |S|$ matrix such that:

1. $q_{ij} \geq 0$, $i \neq j$ is the rate of the exponential distribution giving the time needed to go from state s_i to state s_j
2. $q_{ii} = -\sum_{j \neq i} q_{ij}$ is the opposite of the exit rate $:= \sum_{j \neq i} q_{ij}$ from state i.

Therefore, each row of the Q-matrix sums up to zero.

A SIMPLE EXAMPLE: THE MOOD CHAIN

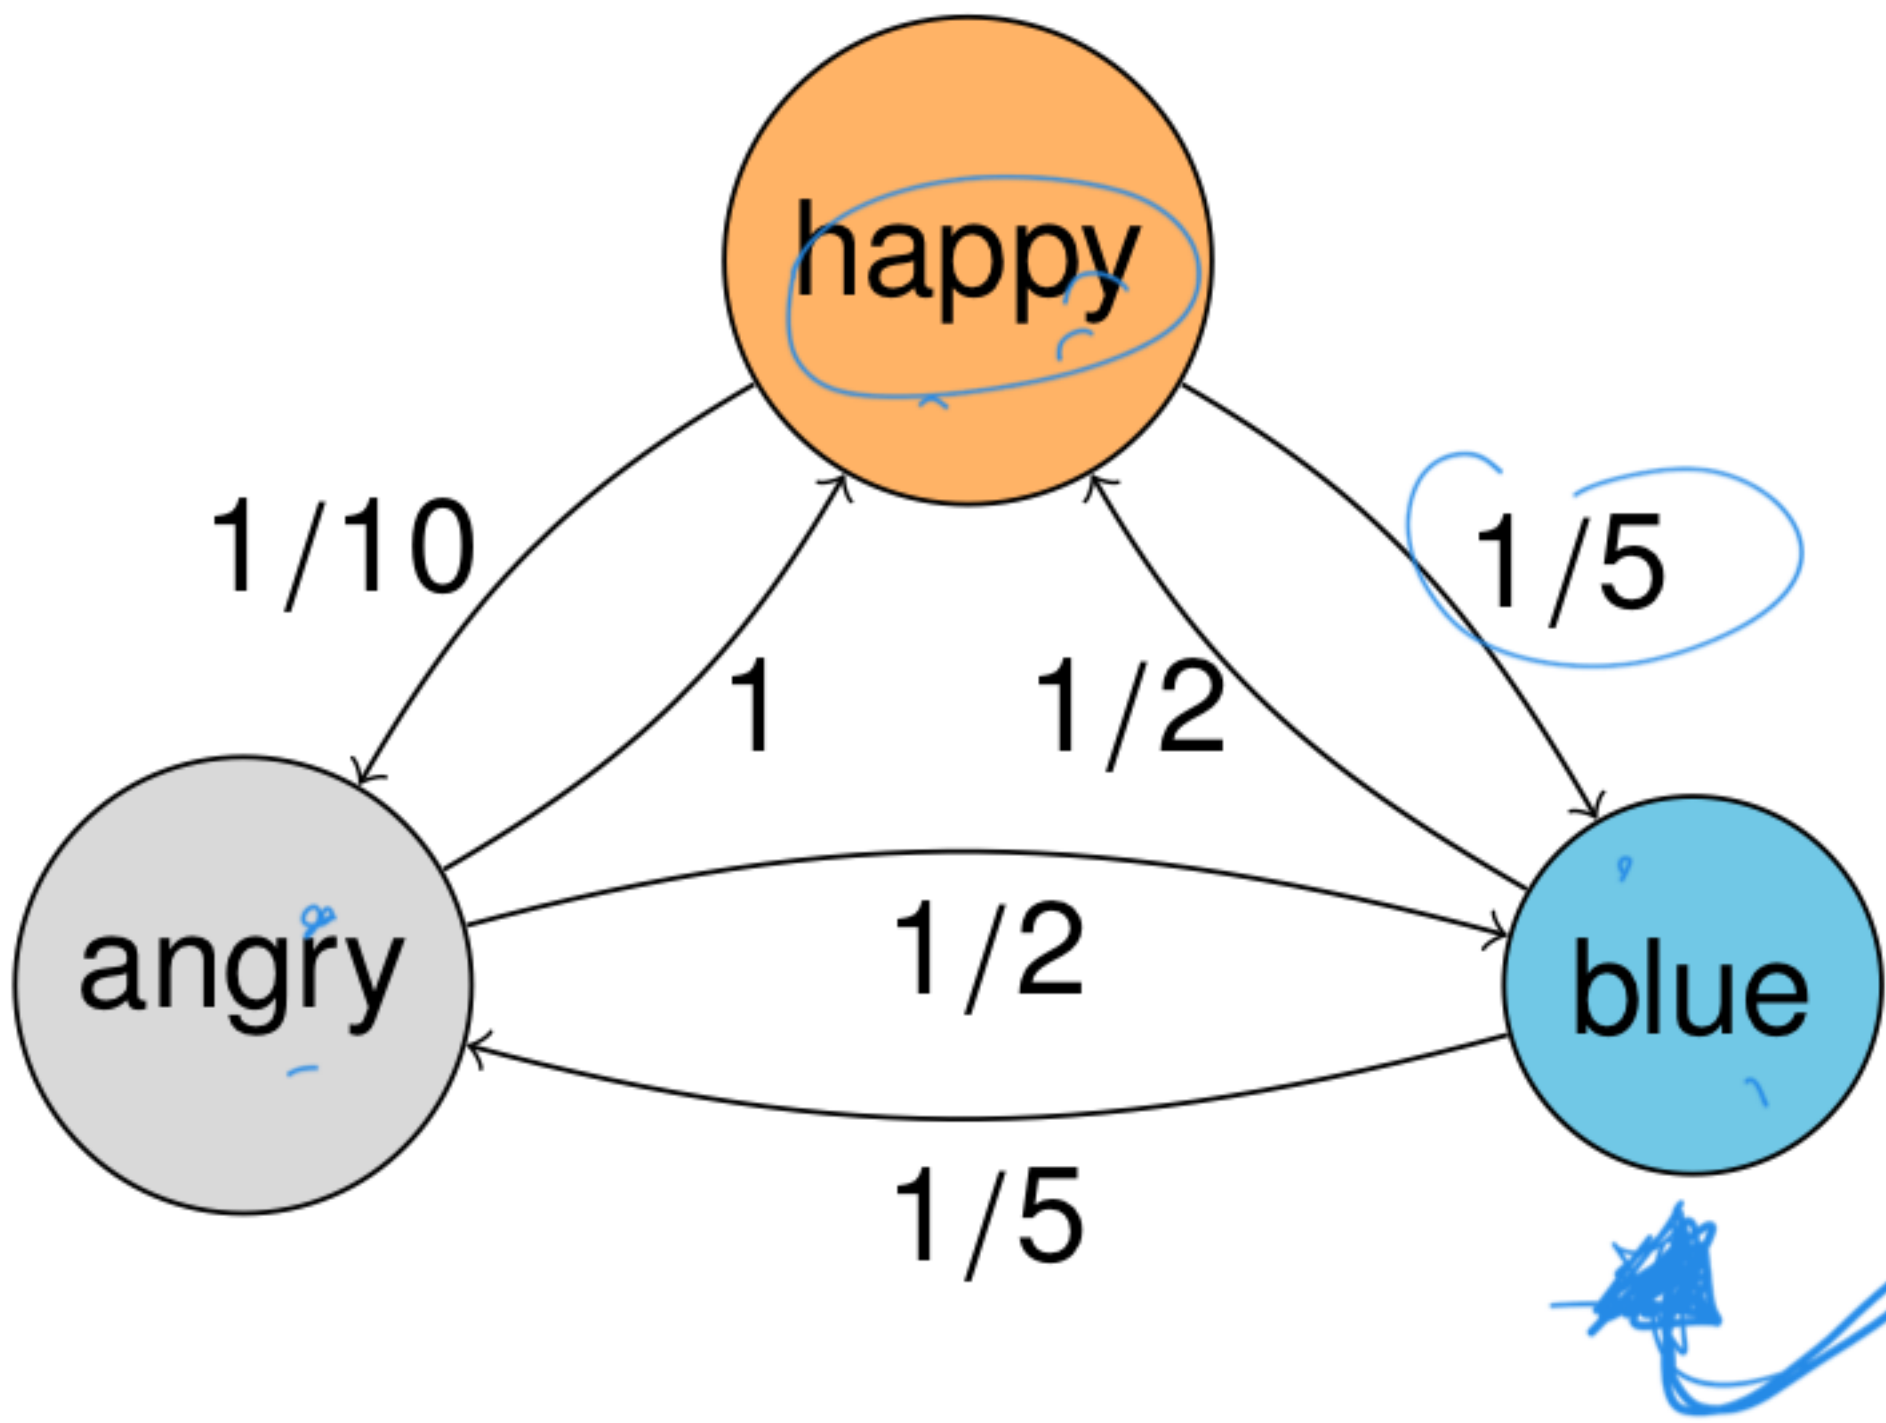

$$S = \{happy, blue, angry\}$$

$$Q = \begin{pmatrix} -\frac{3}{10} & \frac{1}{5} & \frac{1}{10} \\ \frac{1}{2} & -\frac{7}{10} & \frac{1}{5} \\ 1 & \frac{1}{2} & -\frac{3}{2} \end{pmatrix}$$

JUMP CHAIN AND HOLDING TIMES

In each state i, we have a race condition between k transitions, each exponentially distributed with rate q_{ij}. Hence, the time spent is $T = \inf T_{ij}$.

exit rate

By the properties of the exponential distribution, we know that T has rate $q_i = \sum_j q_{ij}$, and that the transition that fires is independent from T and the next state j is chosen with probability q_{ij}/q_i.

We can therefore factorize $X(t)$ into

Jump chain

- a DTMC Y_n, with probability matrix Π, defined by $\pi_{ij} = \frac{q_{ij}}{-q_{ii}}$, if $i \neq j$, and $\pi_{ii} = 0$;
- a sequence of jump times τ_n, where τ_n is the time of the n-th jump. Letting q_i the jump rate from state s_i, then $T_n = \tau_n - \tau_{n-1}$, the n-th holding time, is distributed exponentially with rate q_{Y_n}.
- Y_n and each T_i are independent.
- Hence $X(t) = Y_n$ for $\tau_n \leq t < \tau_{n+1}$.

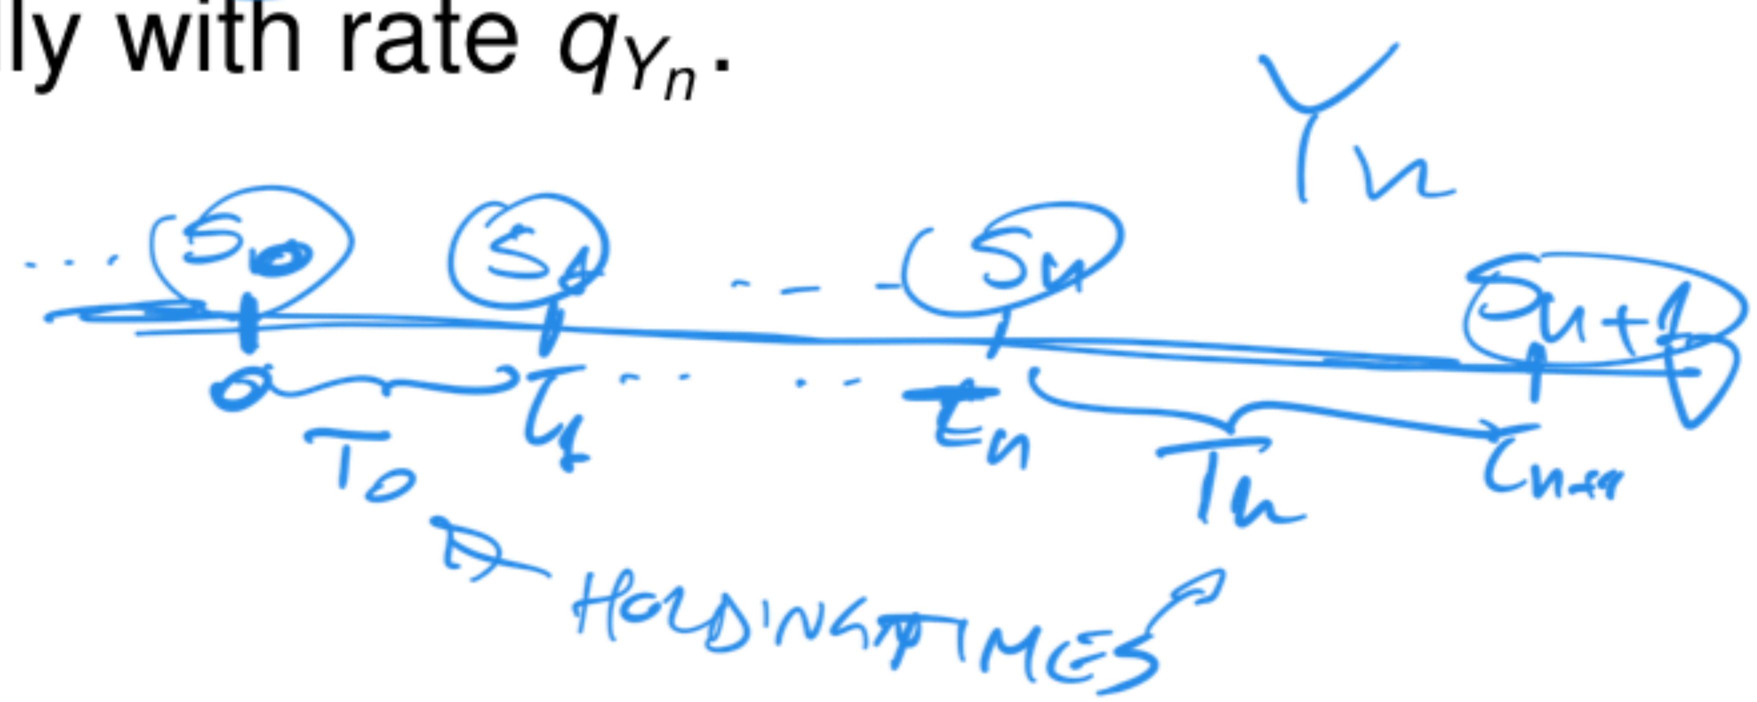

A SIMPLE EXAMPLE: THE MOOD CHAIN

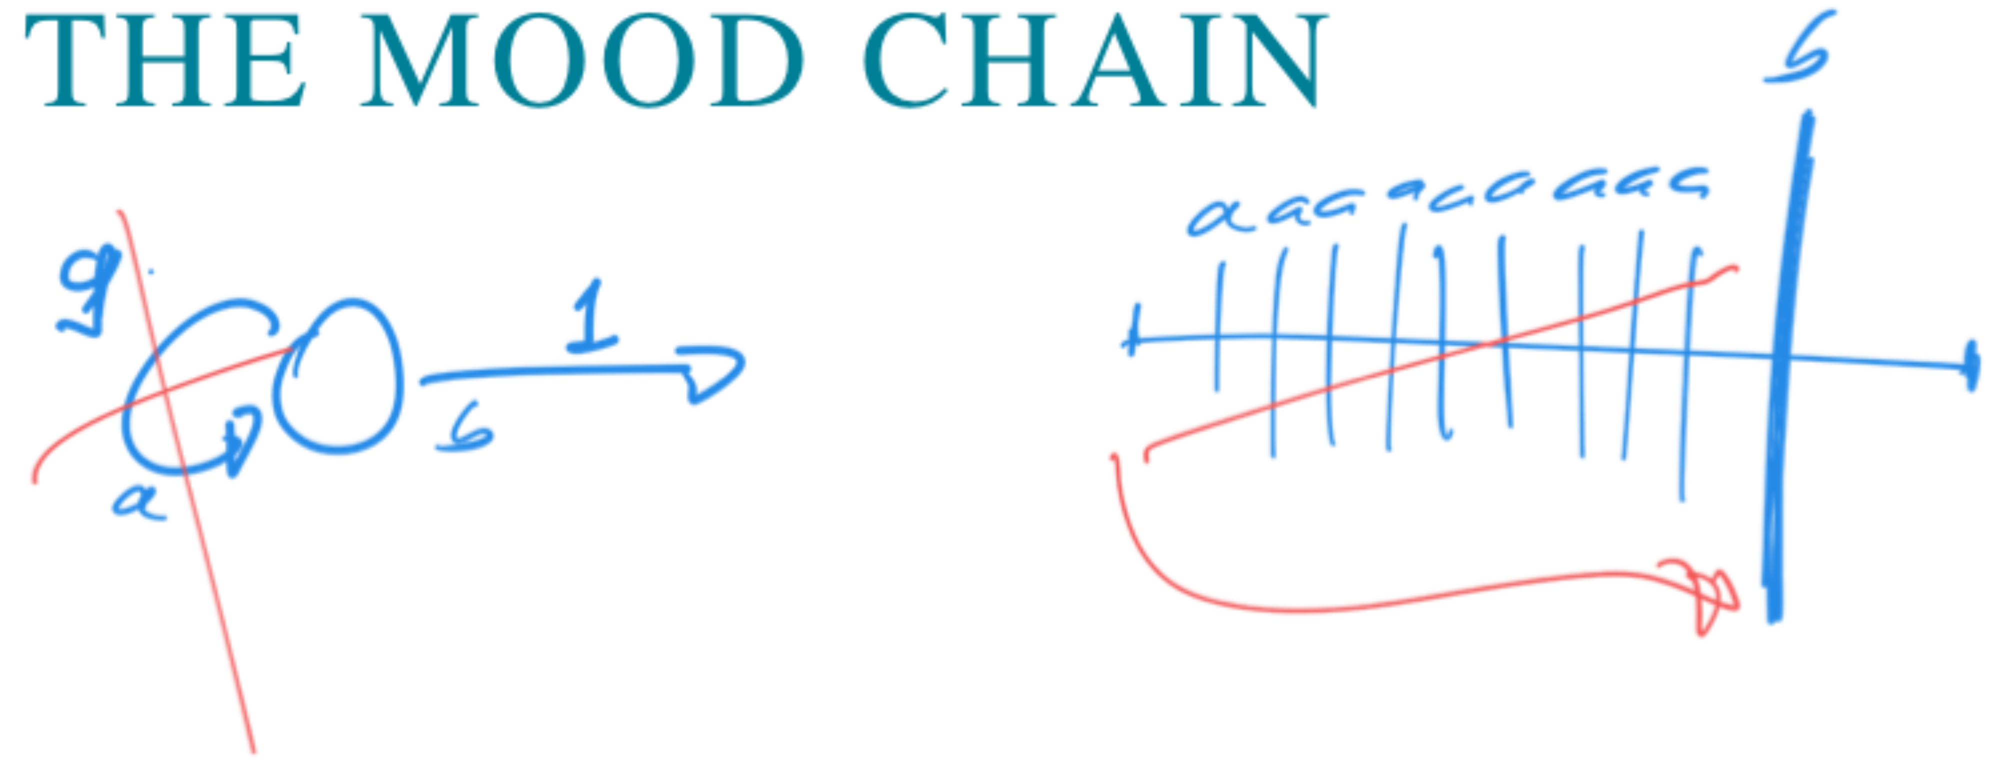

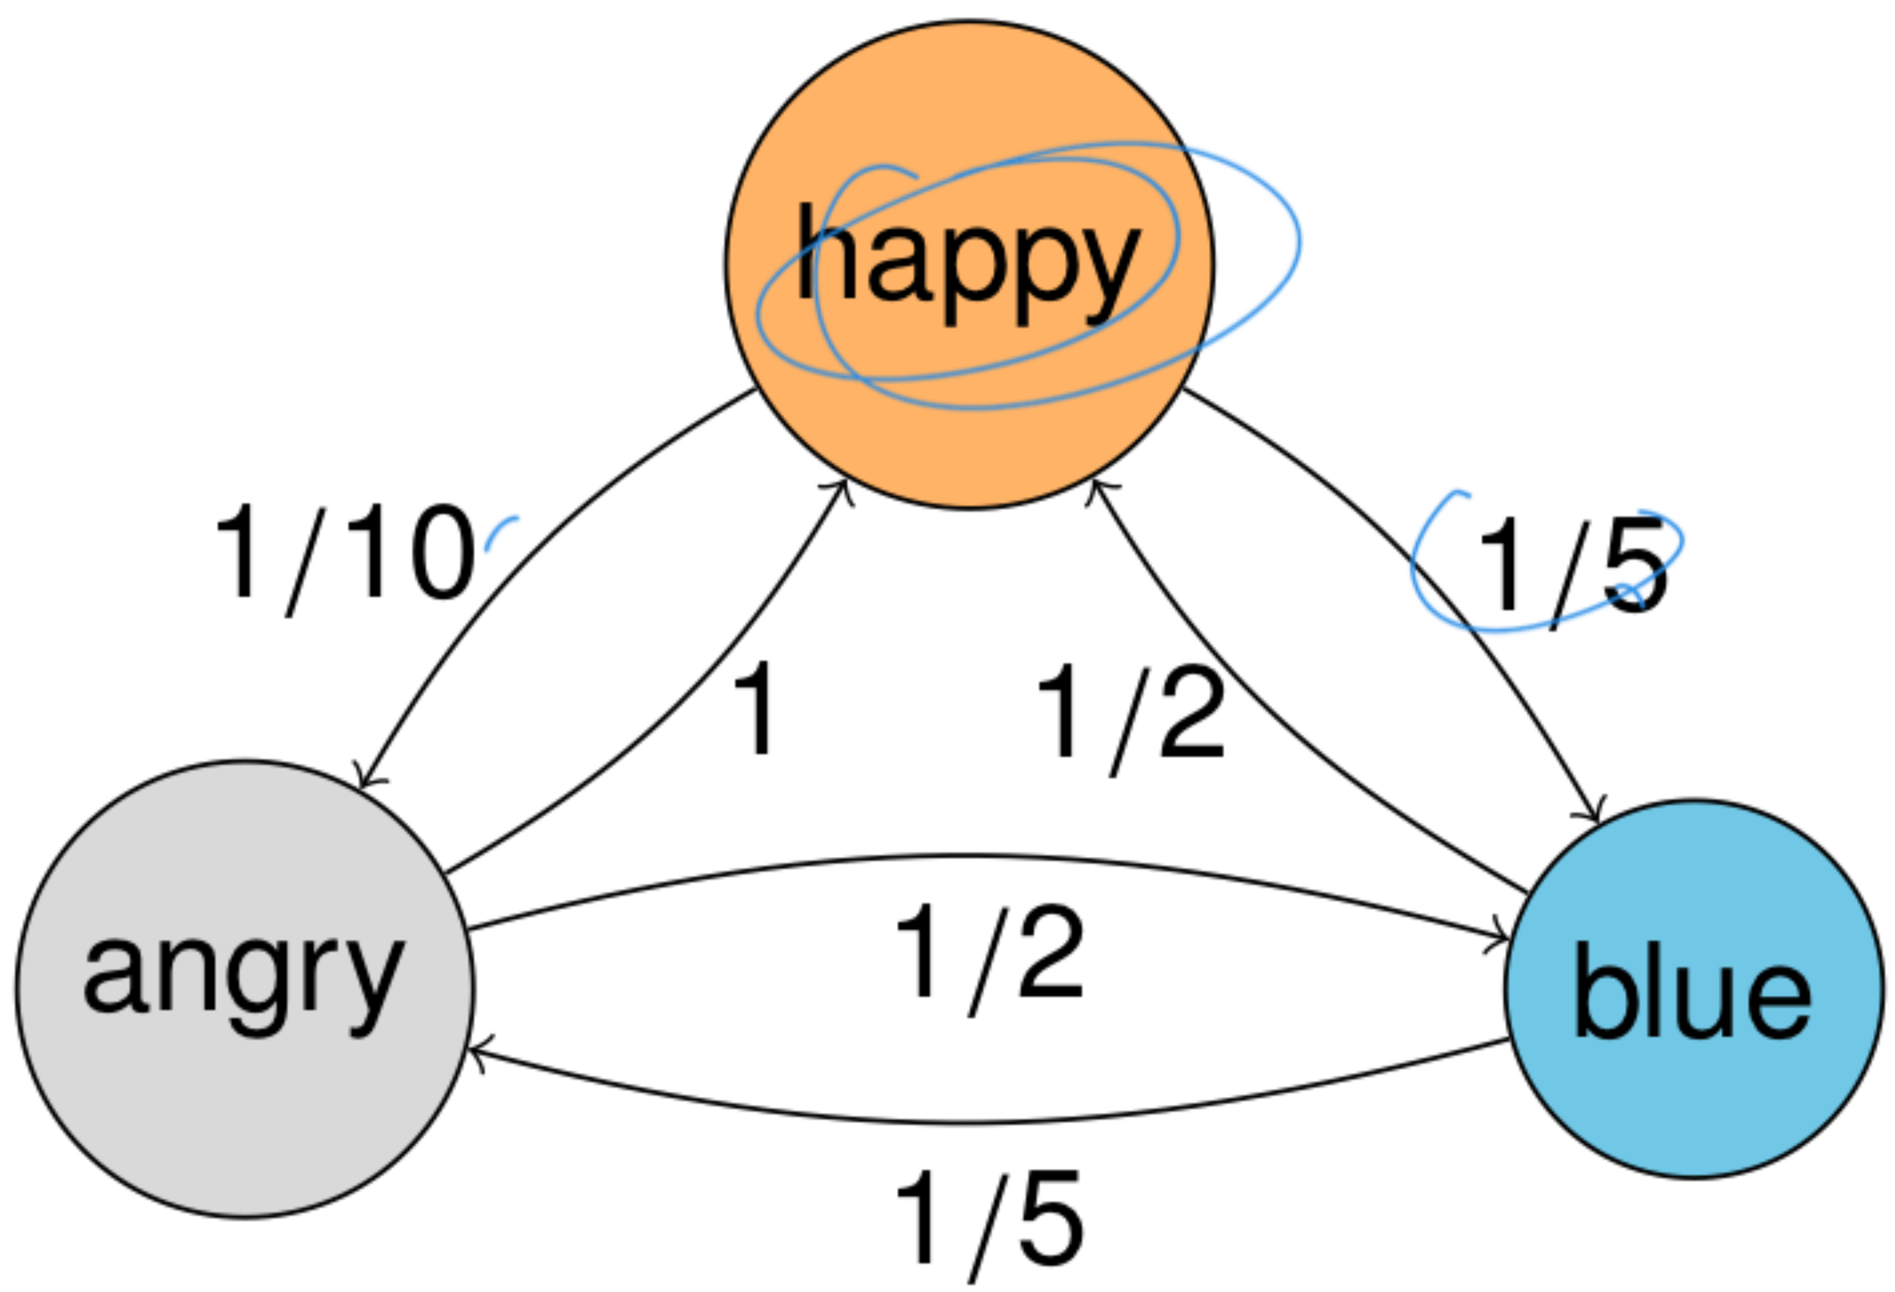

$$S = \{happy, blue, angry\}$$

Jump chain

$$\Pi = \begin{pmatrix} 0 & \frac{2}{3} & \frac{1}{3} \\ \frac{5}{7} & 0 & \frac{2}{7} \\ \frac{2}{3} & \frac{1}{3} & 0 \end{pmatrix}$$

Exit rates

$$q = \left(\frac{3}{10}, \frac{7}{10}, \frac{3}{2}\right)$$

Chapman-Kolmogorov equations

$P(t) = (P_{ij}(t))_{i,j \in S}$

$P(a,b) = P(a|b)P(b)$

$P(a,b|c) = P(a|b,c) \cdot P(b|c)$

Let $\boxed{P_{ij}(t)} = \mathbb{P}\{X(t) = s_j \mid X(0) = s_i\}$. Then

$$
\begin{aligned}
P_{ij}(t+s) &= \mathbb{P}\{X(t+s) = s_j \mid X(0) = s_i\} \\
&= \sum_k \mathbb{P}\{X(t+s) = s_j, X(t) = s_k \mid X(0) = s_i\} \\
&= \sum_k \mathbb{P}\{X(s) = s_j \mid X(0) = s_k\}\mathbb{P}\{X(t) = s_k \mid X(0) = s_i\} \\
&= \sum_k P_{ik}(s)P_{kj}(t).
\end{aligned}
$$

$P(X(t+s)=s_j \mid X(t)=s_k, X(0)=s_i)\, P(X(t)=s_k \mid X(0)=s_i)$

$P(X(t+s)=s_j \mid X(t)=s_k) =$

$= P(X(s)=s_j \mid X(0)=s_k)$

Hence $P(t)$, as a matrix, satisfies

$$\boxed{P(t+s) = P(t)P(s) = P(s)P(t),}$$

which is the semigroup property, also known as Chapman-Kolmogorov equations.

KOLMOGOROV EQUATIONS

Using properties of the exponential, we can compute $P(dt)$:

- $P_{ij}(dt) = q_{ij}dt$, for $i \neq j$;
- $P_{ii}(dt) = 1 - \sum_{j \neq i} q_{ij} dt = \boxed{1 + q_{ii} dt}$

$P(t)P(dt) = P(t)[I + Q dt]$

Hence $P(dt) = I + Qdt$

From the CK equations: $P(t + dt) = P(t) + P(t)Qdt$, from which

$\lim_{dt \to 0} \frac{P(t+dt) - P(t)}{dt} = P(t) \cdot Q$

$$\frac{dP(t)}{dt} = \boxed{P(t)Q},$$

IT'S A LINEAR ODE!

which is the forward Kolmogorov equation.

Using CK the other way round: $P(t + dt) = P(t) + QP(t)dt$, so

$P(dt + t) = P(dt)P(t)$

$$\frac{dP(t)}{dt} = \boxed{QP(t)},$$

which is the backward Kolmogorov equation.

A simple example: the mood chain

$$\frac{d}{dt}P(t) = P(t)\cdot Q \qquad P(0) = I$$

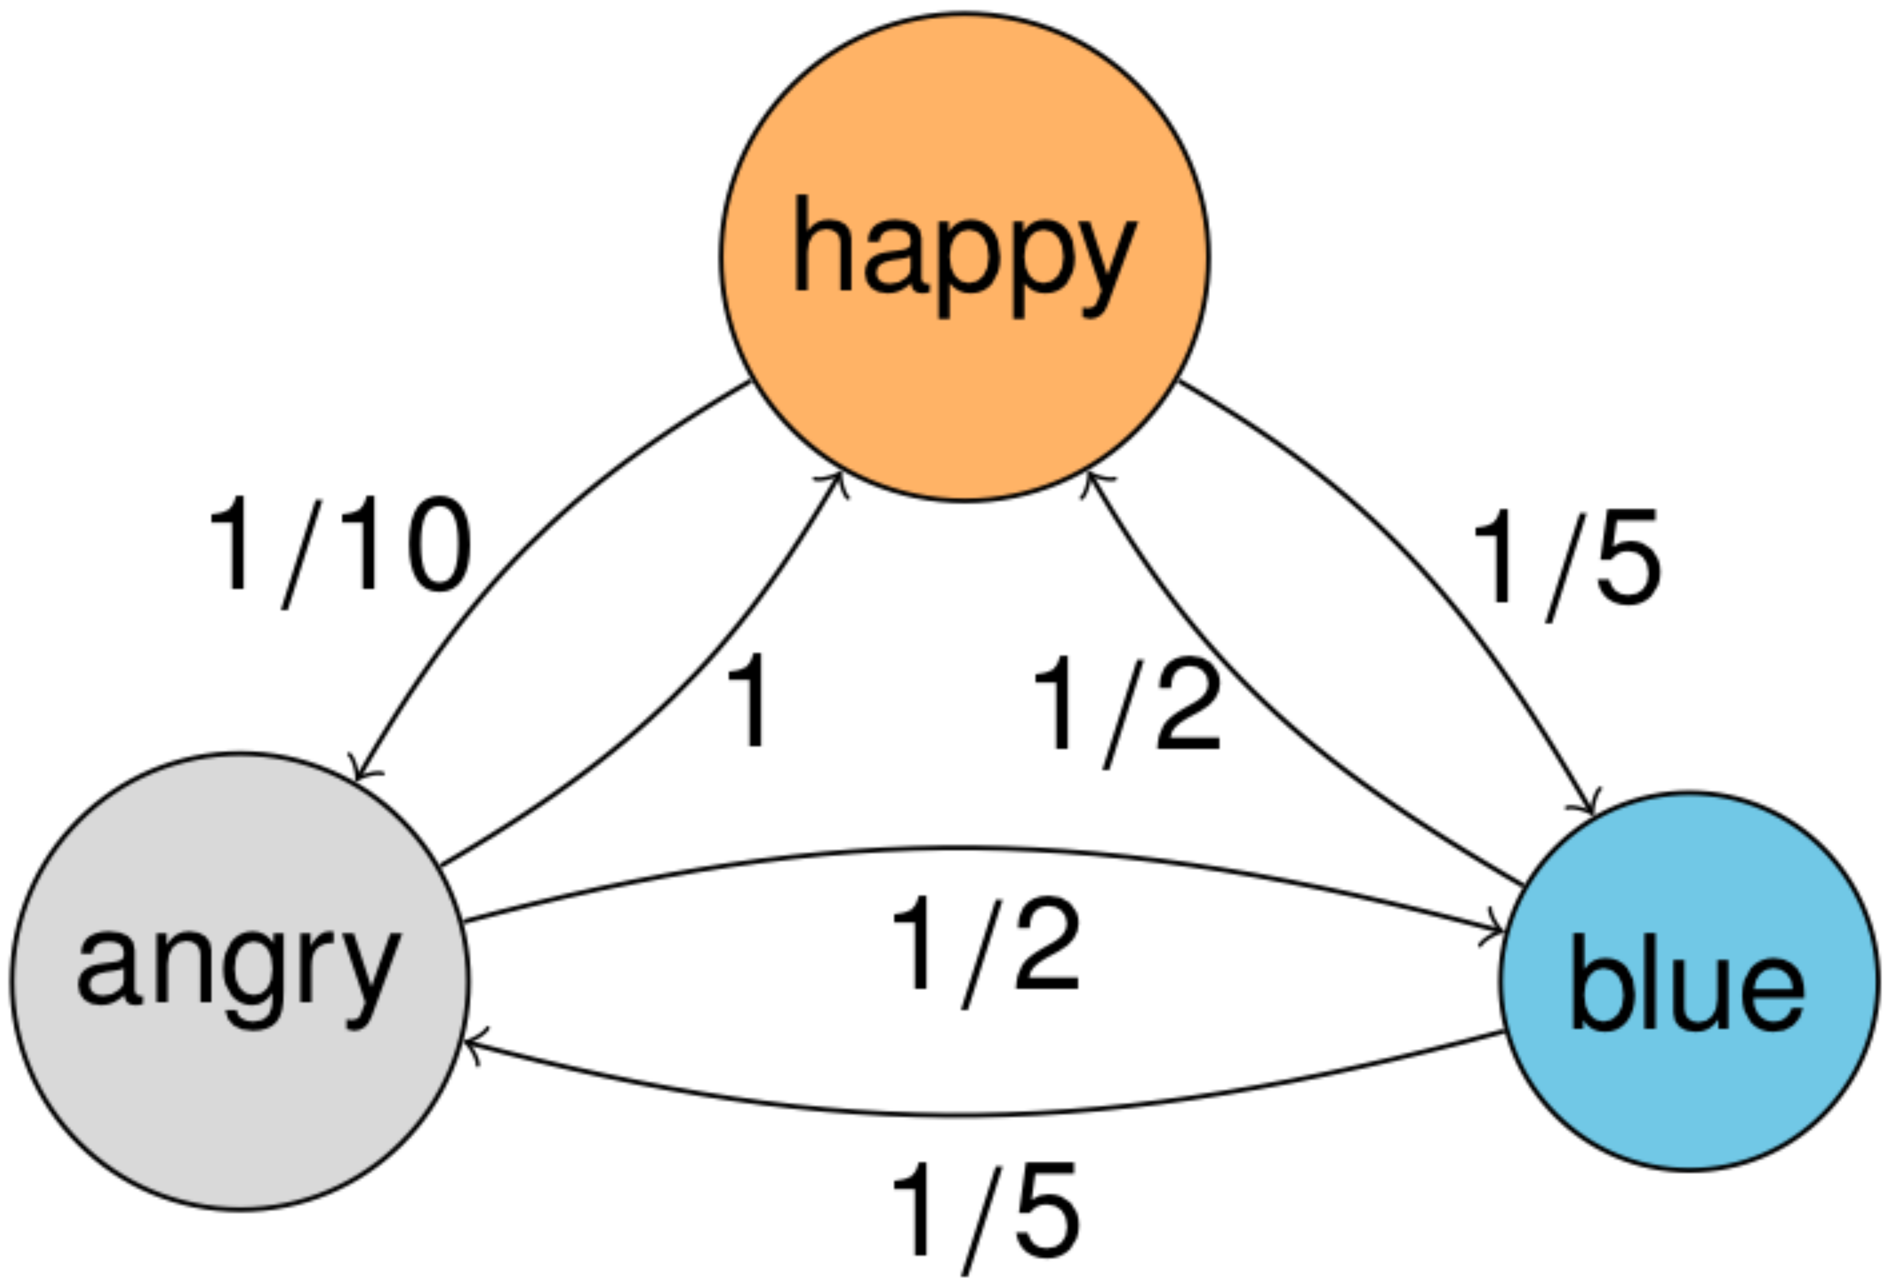

$S = \{happy, blue, angry\}$

$p_0 = (0,1,0) \quad p(t) = p_0 P(t)$

$p_0 \frac{d}{dt}P(t)$

$$\frac{d}{dt}p_0 P(t) = p_0 P(t) Q \Rightarrow \frac{d}{dt}p = pQ$$

$$\frac{d}{dt}p(t) = p(t)\,Q \qquad p(0) = p_0$$

initial distr.

A SIMPLE EXAMPLE: THE MOOD CHAIN

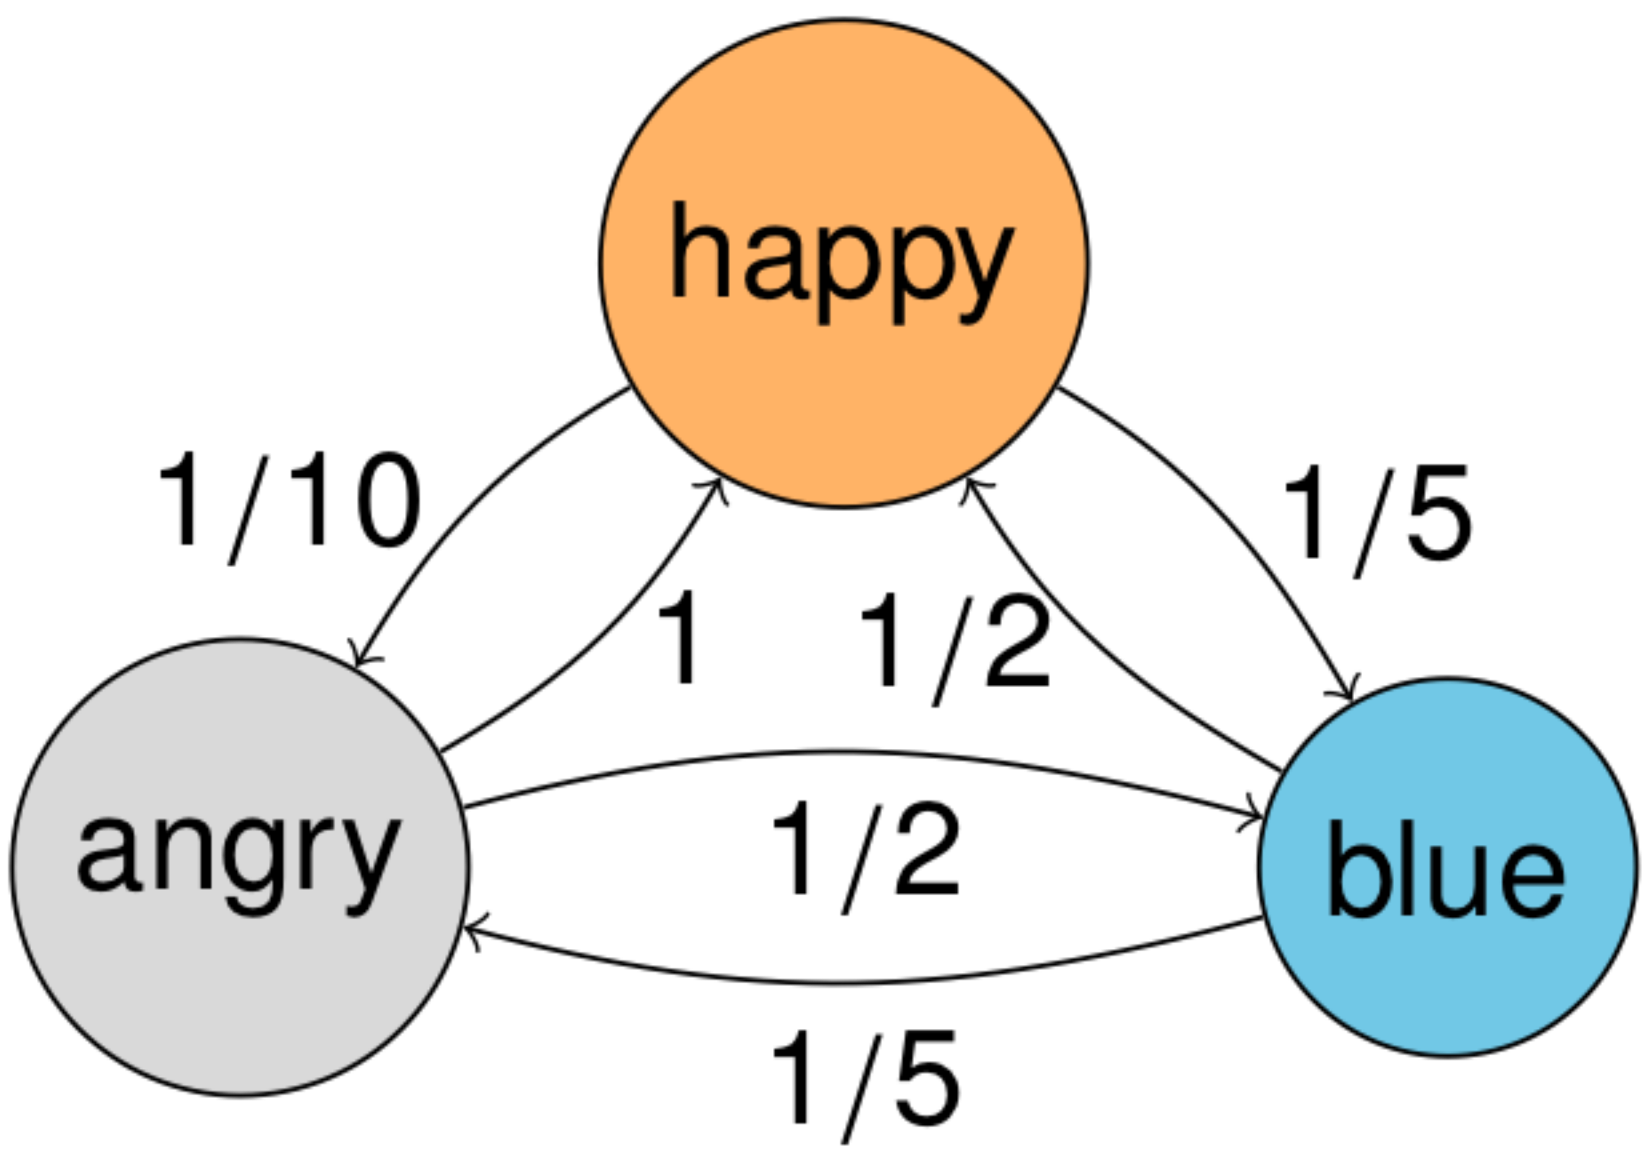

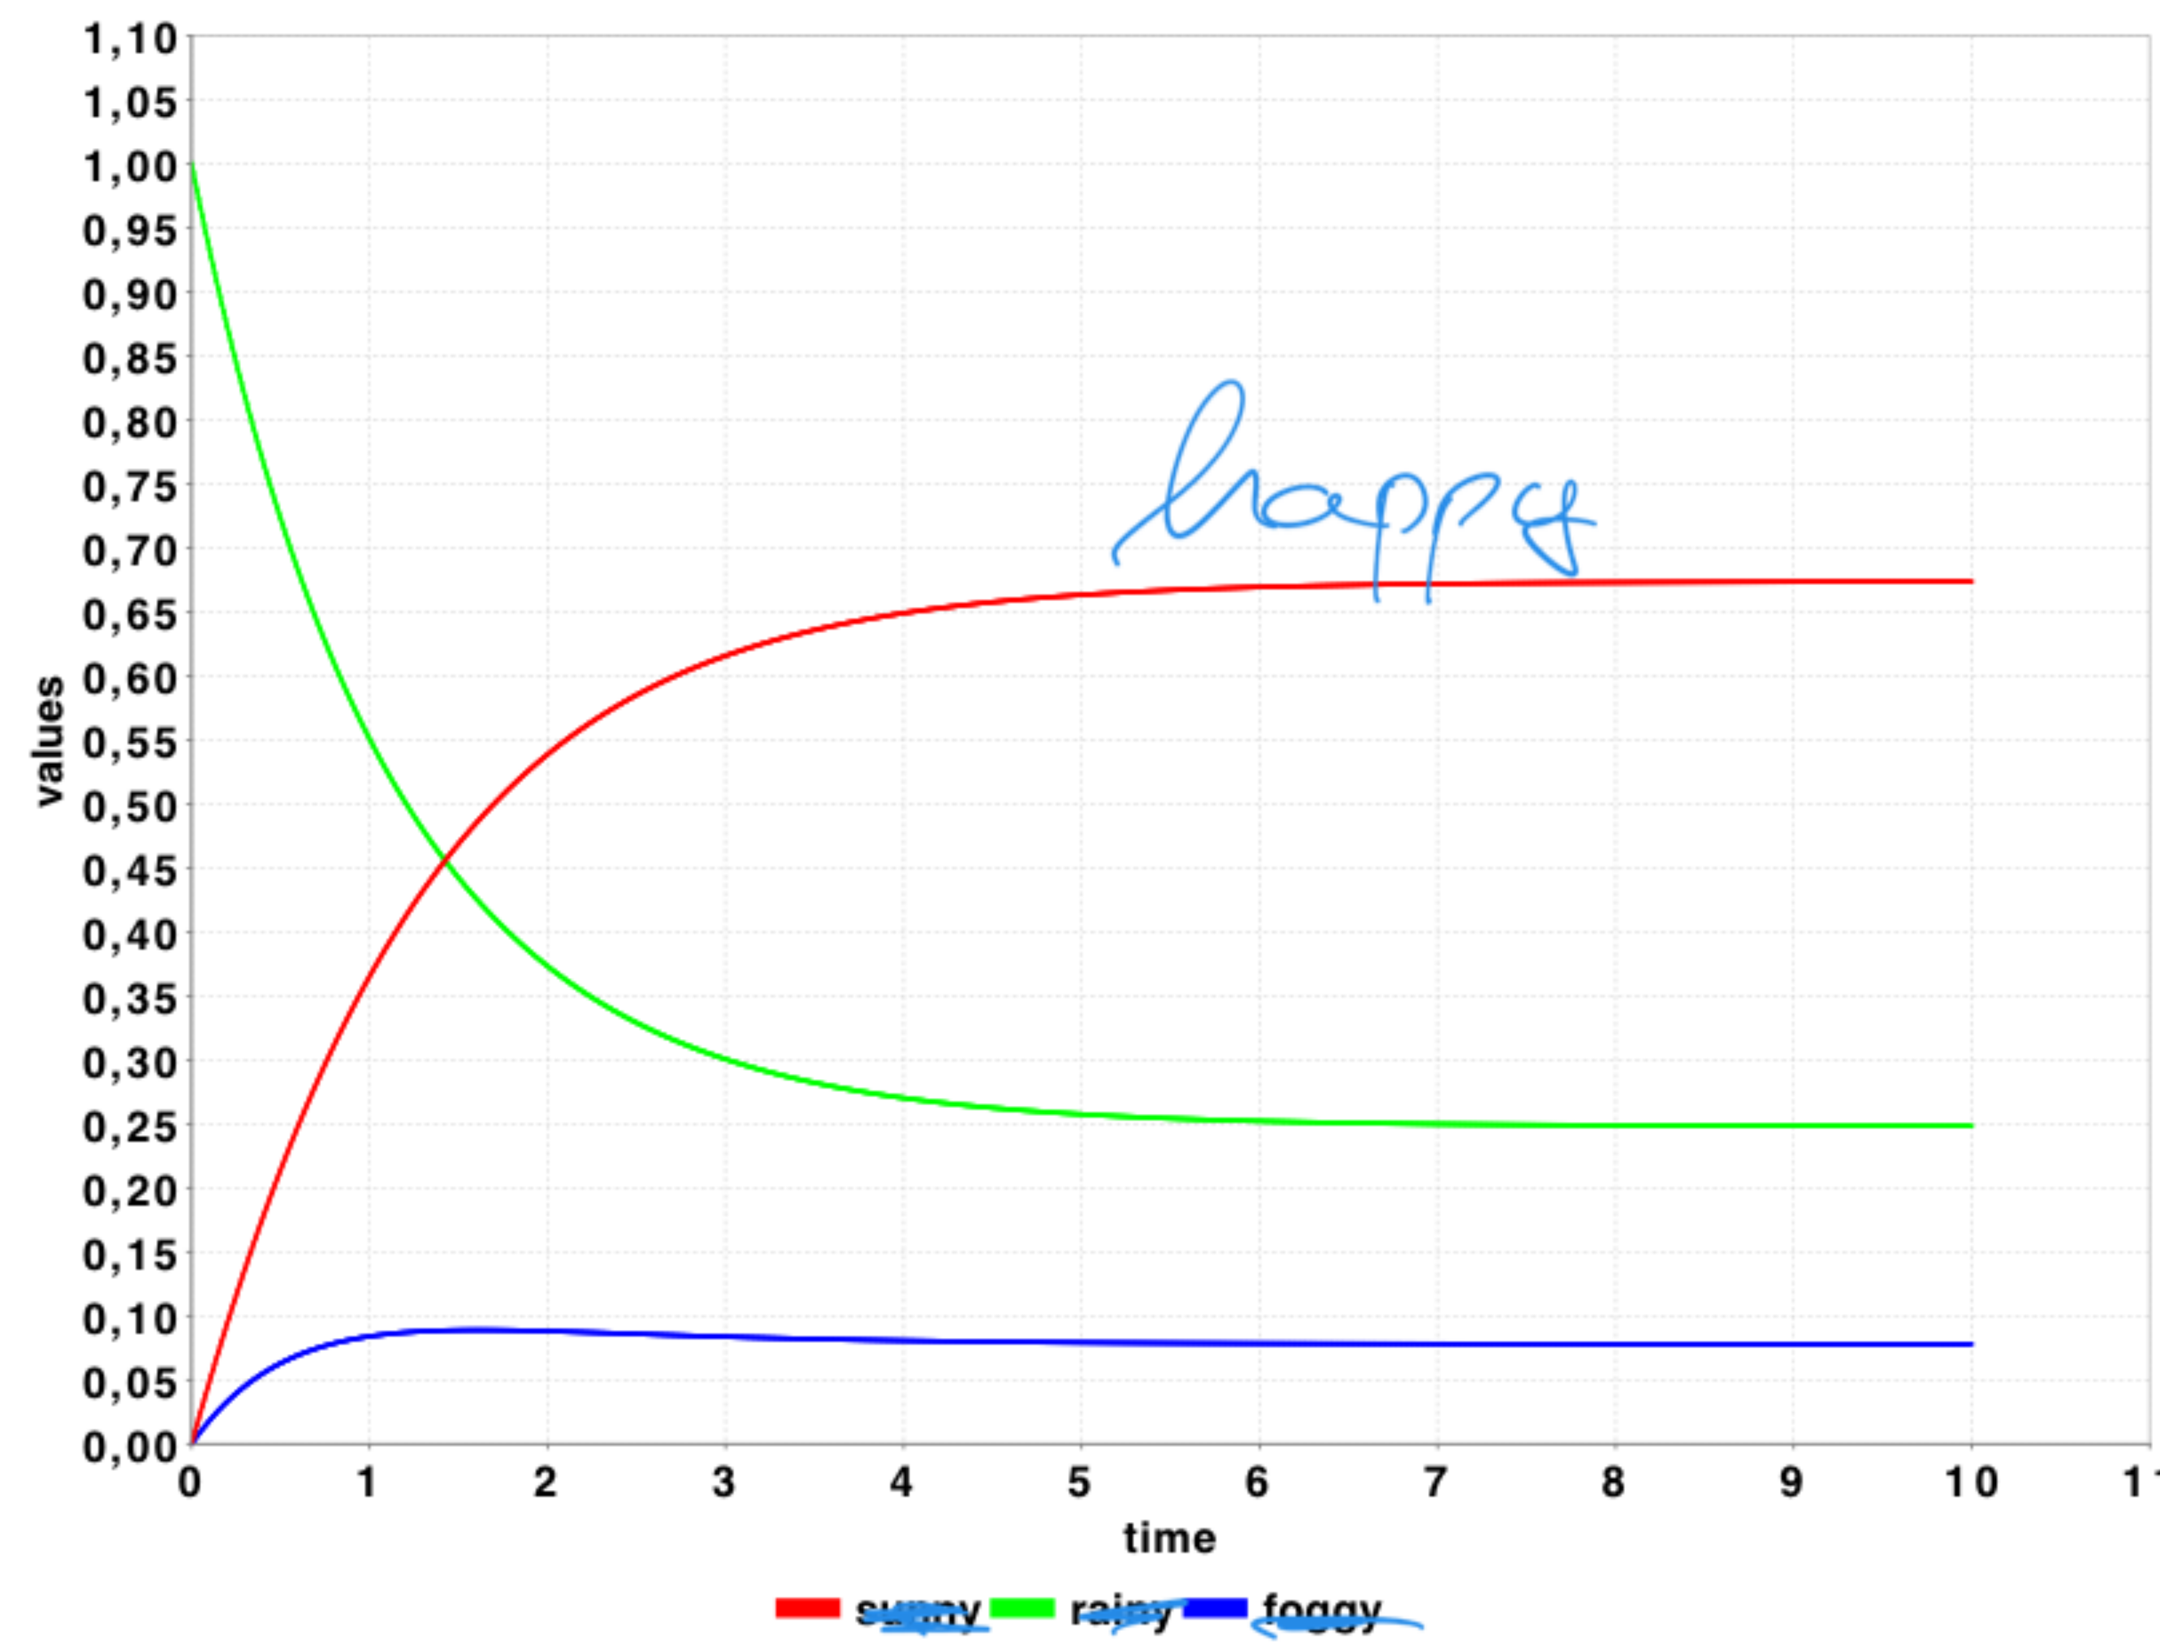

Outline

1. Continuous Time Markov Chains
   - Main concepts
   - Poisson Process
   - Time-inhomogeneous rates
2. Population Continuous Time Markov Chains
3. Simulation
   - SSA
   - Next Reaction Method
   - τ-leaping

POISSON PROCESS: DEFINITION

A Poisson process $N_\lambda(0, t)$ with rate λ is a process that counts how many times an exponential distribution with rate λ has fired from time 0 to time t.

$$Q = \begin{pmatrix} -\lambda & \lambda & & \\ & -\lambda & \lambda & 0 \\ 0 & & \ddots & \\ & & & -\lambda \ \lambda \end{pmatrix}$$

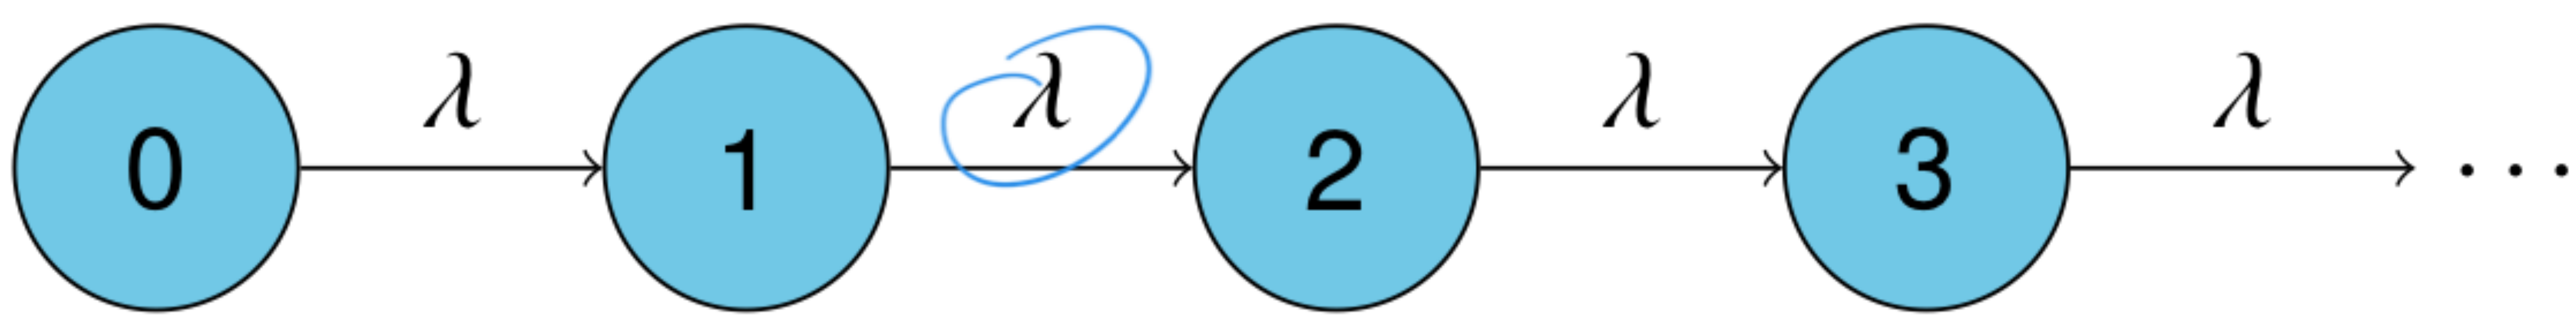

It can be seen as a CTMC on the state space $S = \mathbb{N}$, with rate matrix Q given by $\boxed{q_{i,i+1} = \lambda}$, and zero elsewhere.
It's a very common process. For instance, it is the simplest model of job arrivals in a queue.

POISSON PROCESS: BASIC PROPERTIES

A Poisson random variable $\mathcal{Y}(\lambda)$ with rate λ ($\mathcal{Y}(\lambda) \sim Poisson(\lambda)$) is a r.v. on $\mathbb{N}$ with probability distribution $\mathbb{P}\{\mathcal{Y}(\lambda) = n\} = \frac{e^{-\lambda}\lambda^n}{n!}$.

Its generating function is $G(z) = \mathbb{E}[z^{\mathcal{Y}(\lambda)}] = e^{\lambda(z-1)}$.

The distribution of $\mathcal{N}_\lambda(0,t)$ is $\boxed{Poisson(\lambda t)}$.

We show that $G_t(z) = \mathbb{E}[z^{\mathcal{N}(0,t)}] = e^{\lambda t(z-1)}$.

By the Markov property, $\mathcal{N}(0, t+s) = \mathcal{N}(0,t) + \mathcal{N}(t,s)$, and the two processes on the right are independent.

Then $G_{t+dt}(z) = \mathbb{E}[z^{\mathcal{N}(0,t)}]\mathbb{E}[z^{\mathcal{N}(t,t+dt)}]$. But $\mathbb{E}[z^{\mathcal{N}(t,t+dt)}] = (1-\lambda dt)z^0 + \lambda dt z^1$, hence $G_{t+dt}(z) = G_t(z) + \lambda(z-1)G_t(z)dt$, and so

$$\frac{dG_t(z)}{dt} = \lambda(z-1)G_t(z),$$

which has solution $G_t(z) = e^{\lambda t(z-1)}$, as $\mathcal{N}_\lambda(0,0) = 0$ with probability 1.

INVARIANT MEASURES AND STEADY STATE

P_0 $\quad p(t) \rightsquigarrow \frac{d}{dt}p(t) = p(t)Q$

invariant $\Rightarrow$ $p \cdot Q = 0$

Consider a CTMC with rate matrix Q and finite state space S. An invariant measure for the CTMC is a probability distribution π satisfying

$$\boxed{\pi Q = 0.}$$

If Q is irreducible (has a strongly connected graph), then it has a unique invariant measure.

Consider an irreducible CTMC with rate matrix Q and finite state space S, and let π be its invariant probability distribution. Then, for each $s_i, s_j \in S$,

$$\lim_{t\to\infty} P_{ij}(t) = \pi_j.$$

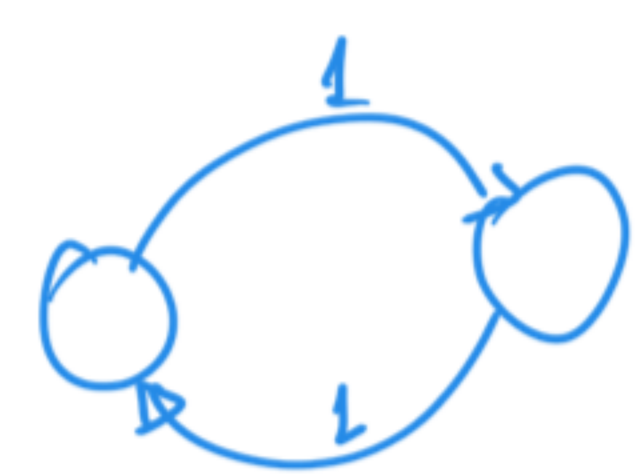

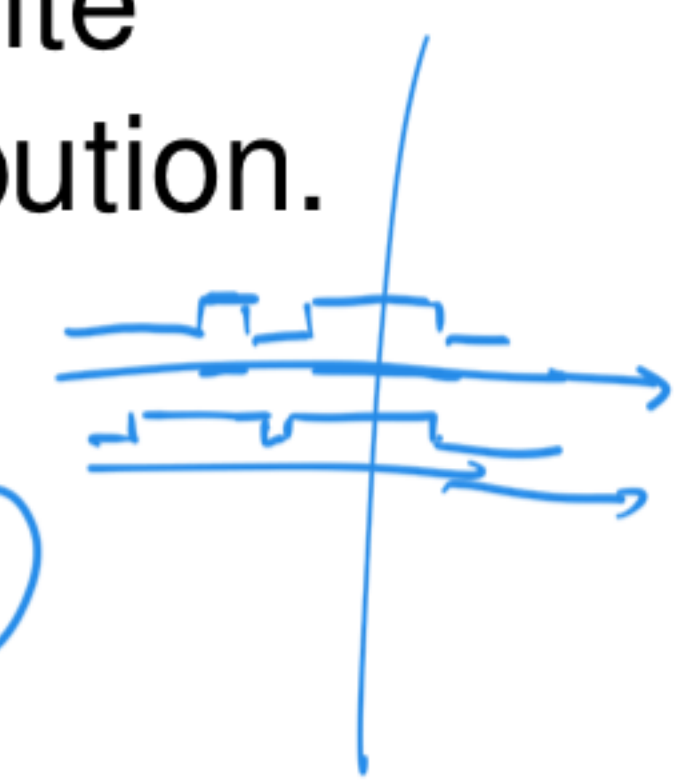

Notice that $\boxed{\text{aperiodicity}}$ is not required. Why?

case countable S

↳ POSITIVE RECURRENCE

EXAMPLE: BIRTH-DEATH PROCESS

A birth-death process is a CTMC on $S = \mathbb{N}$ with birth rate λ_i (from i to $i+1$) and death rate μ_i (from i to $i-1$).

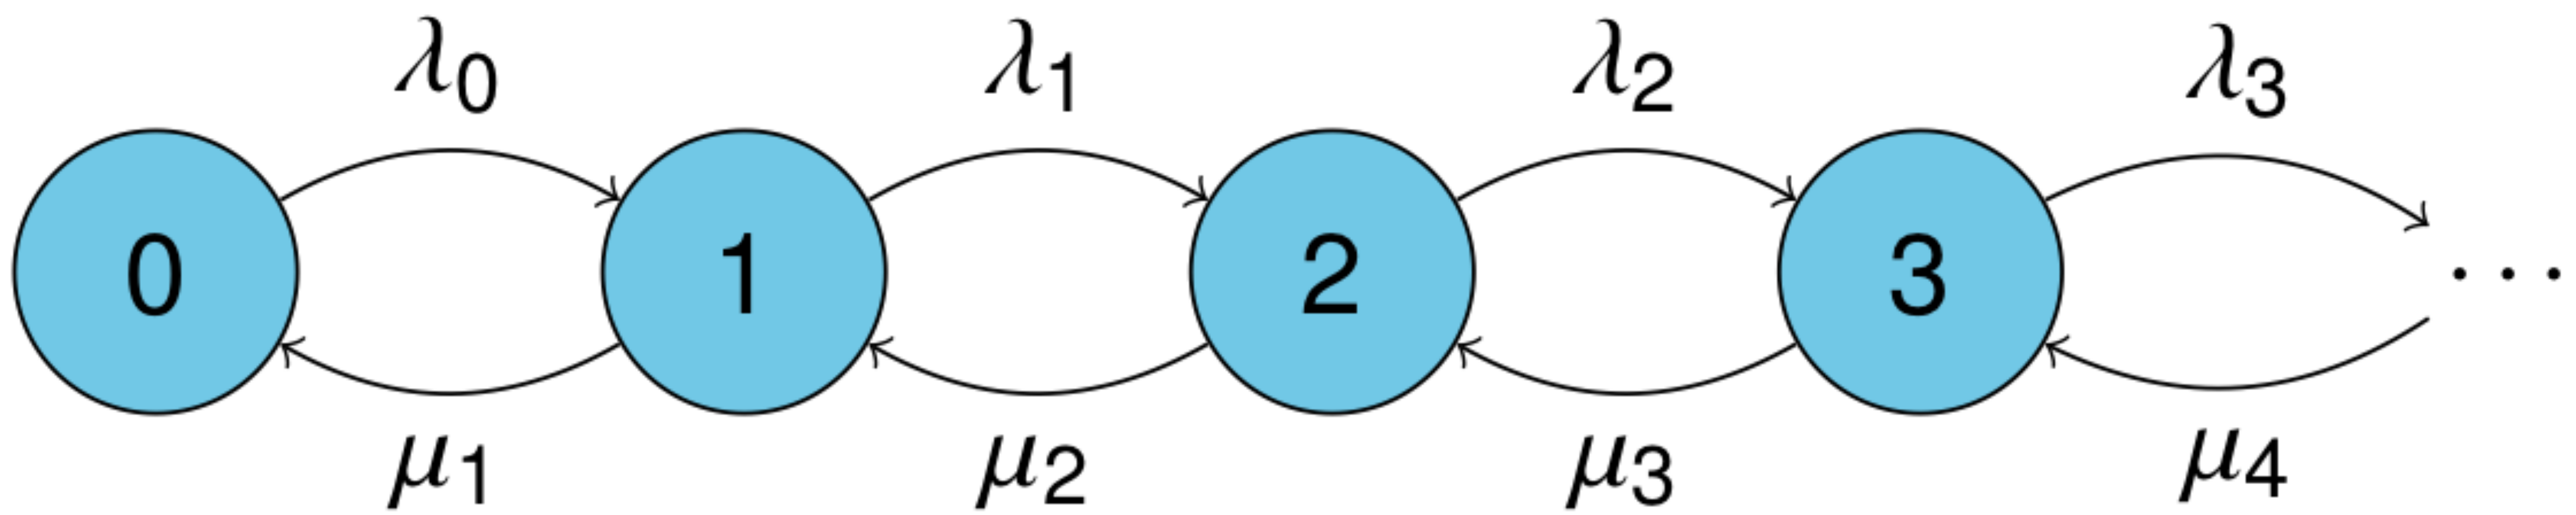

To derive the steady state π, we can use the fact that the net flow along each cut must be zero (why?):

$$\pi_i \lambda_i = \pi_{i+1} \mu_{i+1}$$

Hence we get:

$$\pi_k = \prod_{i=0}^{k-1} \frac{\lambda_i}{\mu_{i+1}} \pi_0; \qquad \pi_0 = \left(1 + \sum_{k=1}^{\infty} \prod_{i=0}^{k-1} \frac{\lambda_i}{\mu_{i+1}} \right)^{-1}$$

EXAMPLE: BIRTH-DEATH PROCESS

Consider a birth-death process with constant birth rate λ and constant death rate μ. It is the model of an M/M/∞ queue.

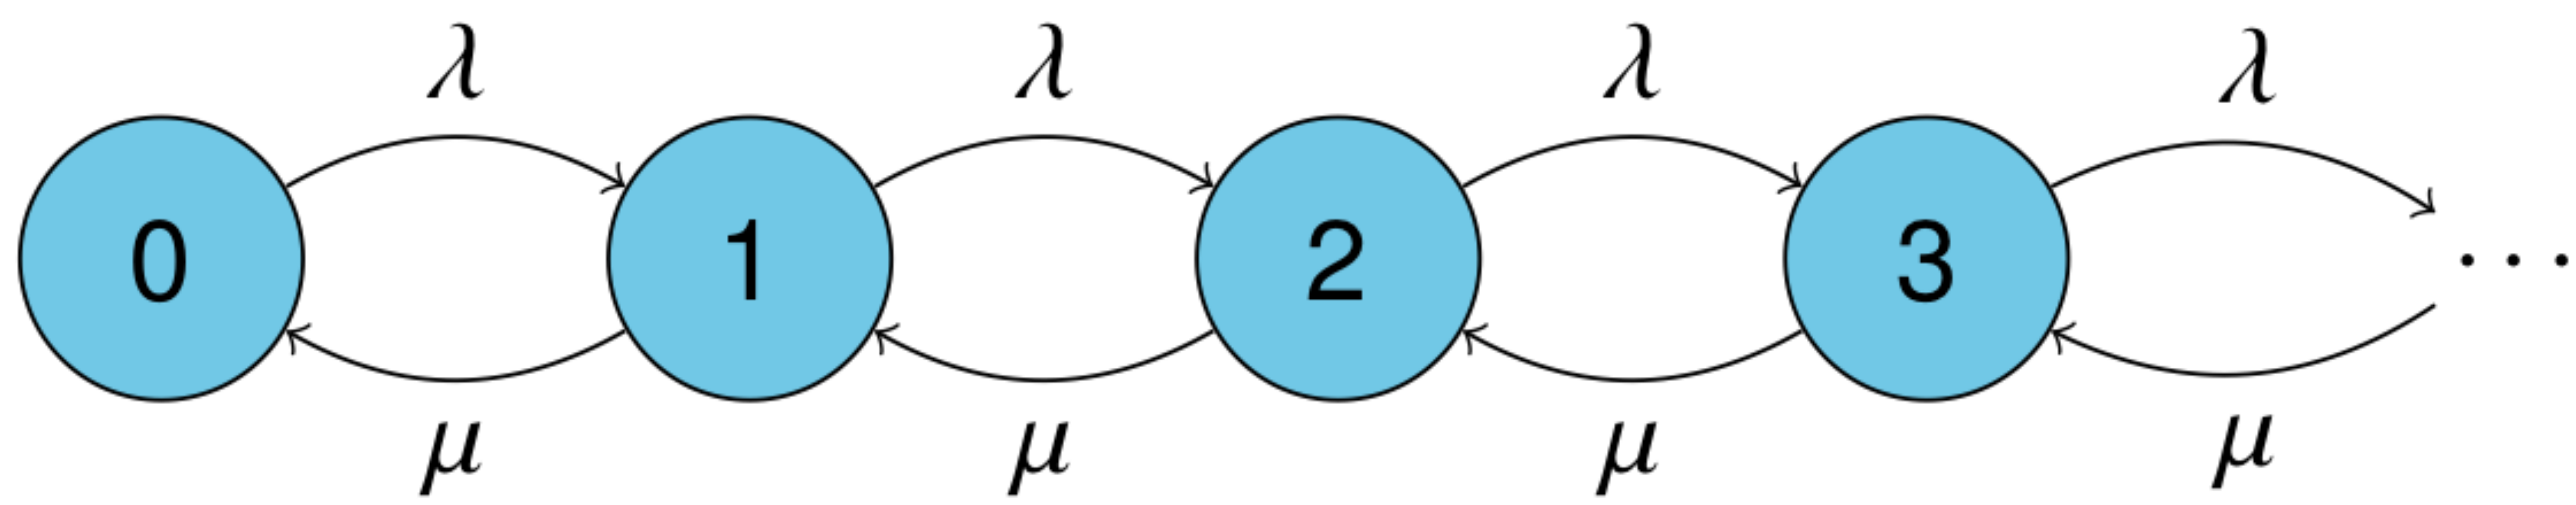

$$\pi_k = \left(\frac{\lambda}{\mu}\right)^k \pi_0; \qquad \pi_0 = \left(1 + \sum_{k=1}^{\infty} \left(\frac{\lambda}{\mu}\right)^k\right)^{-1}$$

- If $\lambda \geq \mu$, then $\pi_0 = 0 = \pi_k$. No state is positive recurrent, there is no invariant measure. The chain escapes to infinity.
- If $\lambda < \mu$, then $\pi_0 = \frac{1-\lambda/\mu}{2-\lambda/\mu}$ and $\pi_k = \left(\frac{\lambda}{\mu}\right)^k \frac{1-\lambda/\mu}{2-\lambda/\mu}$

EXAMPLE: BIRTH-DEATH PROCESS

If $\lambda < \mu$, then $\pi_0 = \frac{1-\lambda/\mu}{2-\lambda/\mu}$ and $\pi_k = \left(\frac{\lambda}{\mu}\right)^k \frac{1-\lambda/\mu}{2-\lambda/\mu}$

Assume $\lambda = 1$, $\mu = 2$.

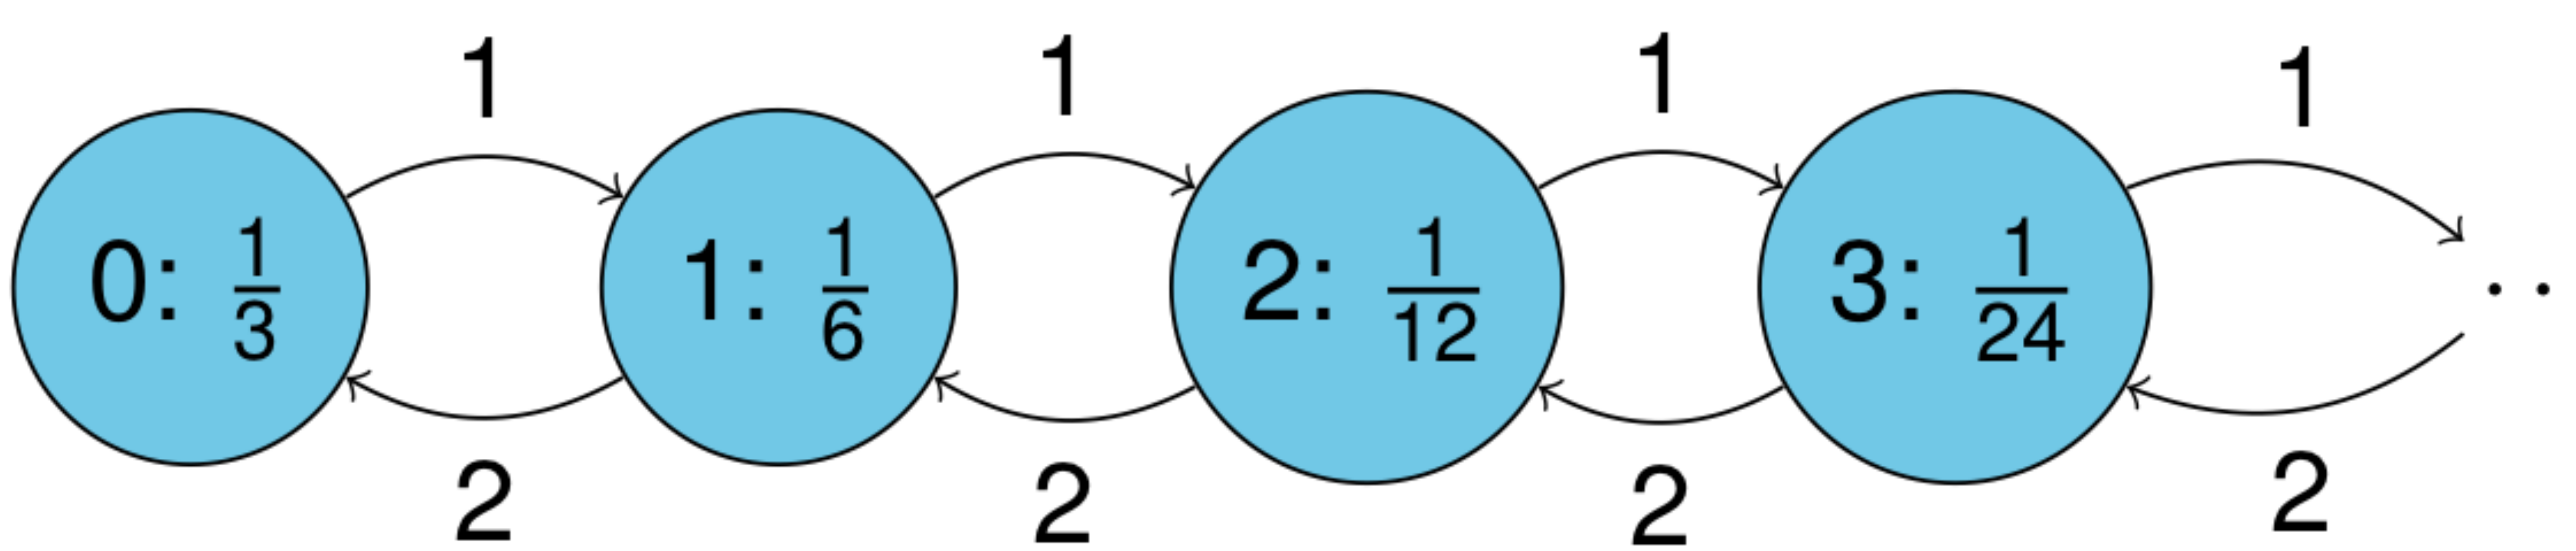

MATRIX EXPONENTIAL

The solution of the forward Kolmogorov equation $\frac{dP(t)}{dt} = P(t)Q$, for a generic CTMC, can be given in terms of the matrix exponential

$$P(t) = e^{Qt} = \boxed{\sum_{n=0}^{\infty} \frac{t^n Q^n}{n!}}.$$

$$\exp\left(\log \frac{t^n}{n!}\right) Q^n =$$

$$= \exp\Big(\underbrace{n \log t + \sum_i \lg i}_{\text{evita overflow}}\Big) Q^n$$

However, numerical computation of the series expansion is numerically unstable.

UNIFORMIZATION

A more efficient strategy is to solve the uniformized CTMC.
Let $\lambda \geq \max_i\{-q_{ii}\}$.
Then one considers a CTMC with jump chain $Y(n)$ with matrix

$$\Pi = I + \frac{1}{\lambda} Q,$$

and uniform exit rate λ.
The number of fires of this CTMC up to time t is a Poisson process $N_\lambda(0,t)$, and so

$$\boxed{X(t)} = \boxed{Y_{N(0,t)}} = \boxed{Y_{\mathcal{Y}(\lambda t)}}.$$

$X(t)=s$

$\underbrace{P(Y(n)=s \mid \mathcal{Y}(\lambda t)=n)}\ \underbrace{P(\mathcal{Y}(\lambda t)=n)}$

It follows that

$P(X(t)=s) = \sum_n P(X(t)=s, \mathcal{Y}=n)$

$$P(t) = \sum_{n=0}^{\infty} \boxed{\frac{e^{-\lambda t}(\lambda t)^n}{n!}} \Pi^n,$$

ADAPTIVE UNIF.
(FINITE PROJ METHOD)

which can be truncated above (and below) by bounding the Poisson r.v.

A SIMPLE EXAMPLE: THE MOOD CHAIN

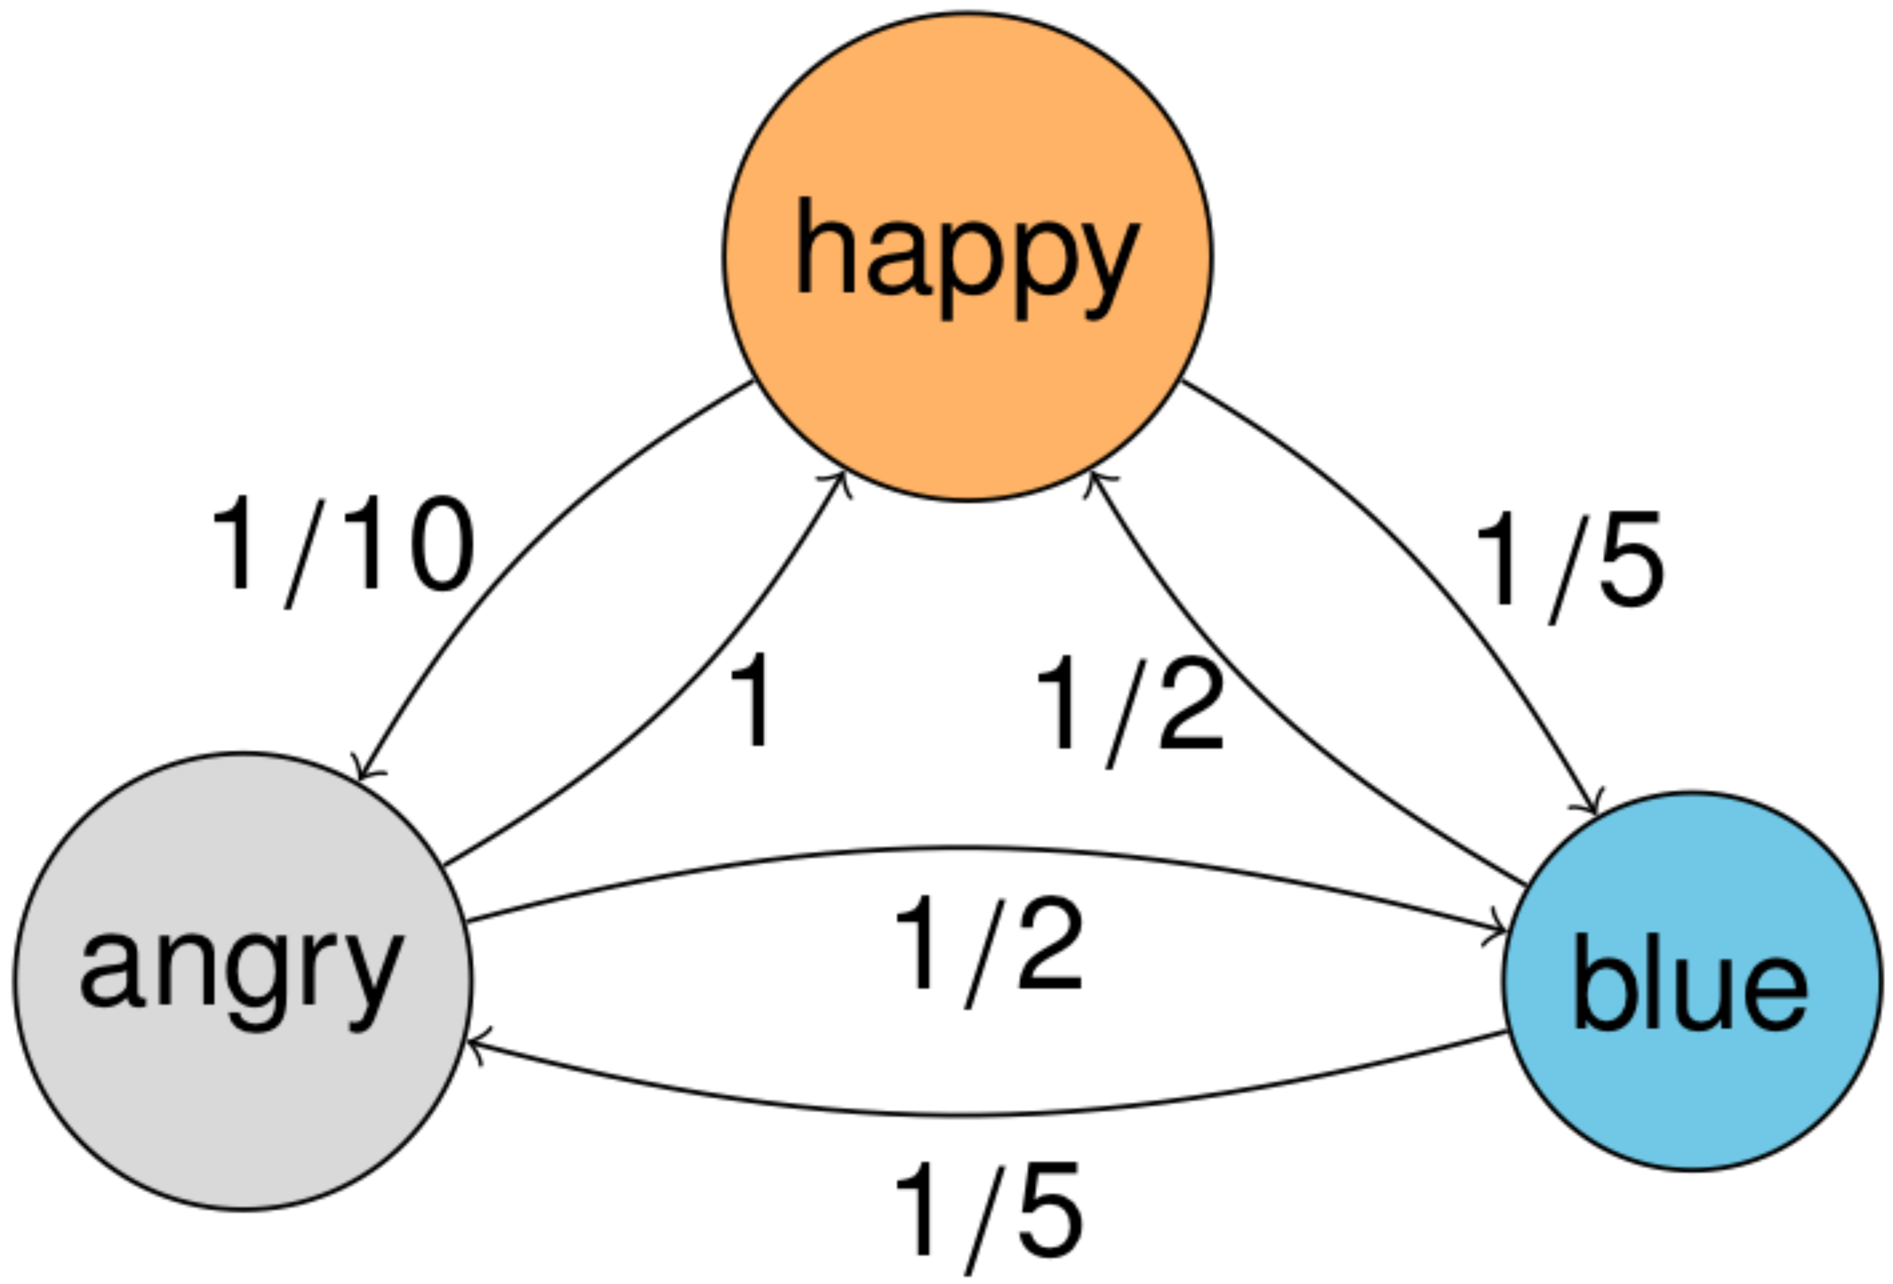

Upper bound on exit rate: 2

$$P(t) = \sum_{n=0}^{\infty} \frac{e^{-2t}(2t)^n}{n!}\Pi^n$$

$$\Pi = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} + \frac{1}{2}\begin{pmatrix} -\frac{3}{10} & \frac{1}{5} & \frac{1}{10} \\ \frac{1}{2} & -\frac{7}{10} & \frac{1}{5} \\ 1 & \frac{1}{2} & -\frac{3}{2} \end{pmatrix} = \begin{pmatrix} \frac{17}{20} & \frac{2}{20} & \frac{1}{20} \\ \frac{5}{20} & \frac{13}{20} & \frac{2}{20} \\ \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}$$