More or Less True:
DCTL for CTMDPs

David N. Jansen

FORMATS 2013

A Challenge for the Bored

QEST: sound and (hopefully) complete weak simulation for substochastic DTMCs

FORMATS: More or less true: DCTL for CTMCs

A Challenge for the Bored

QEST: sound and ~~(hopefully) complete~~ weak simulation for substochastic DTMCs

Baier, Engelen, Majster-Cederbaum 2000

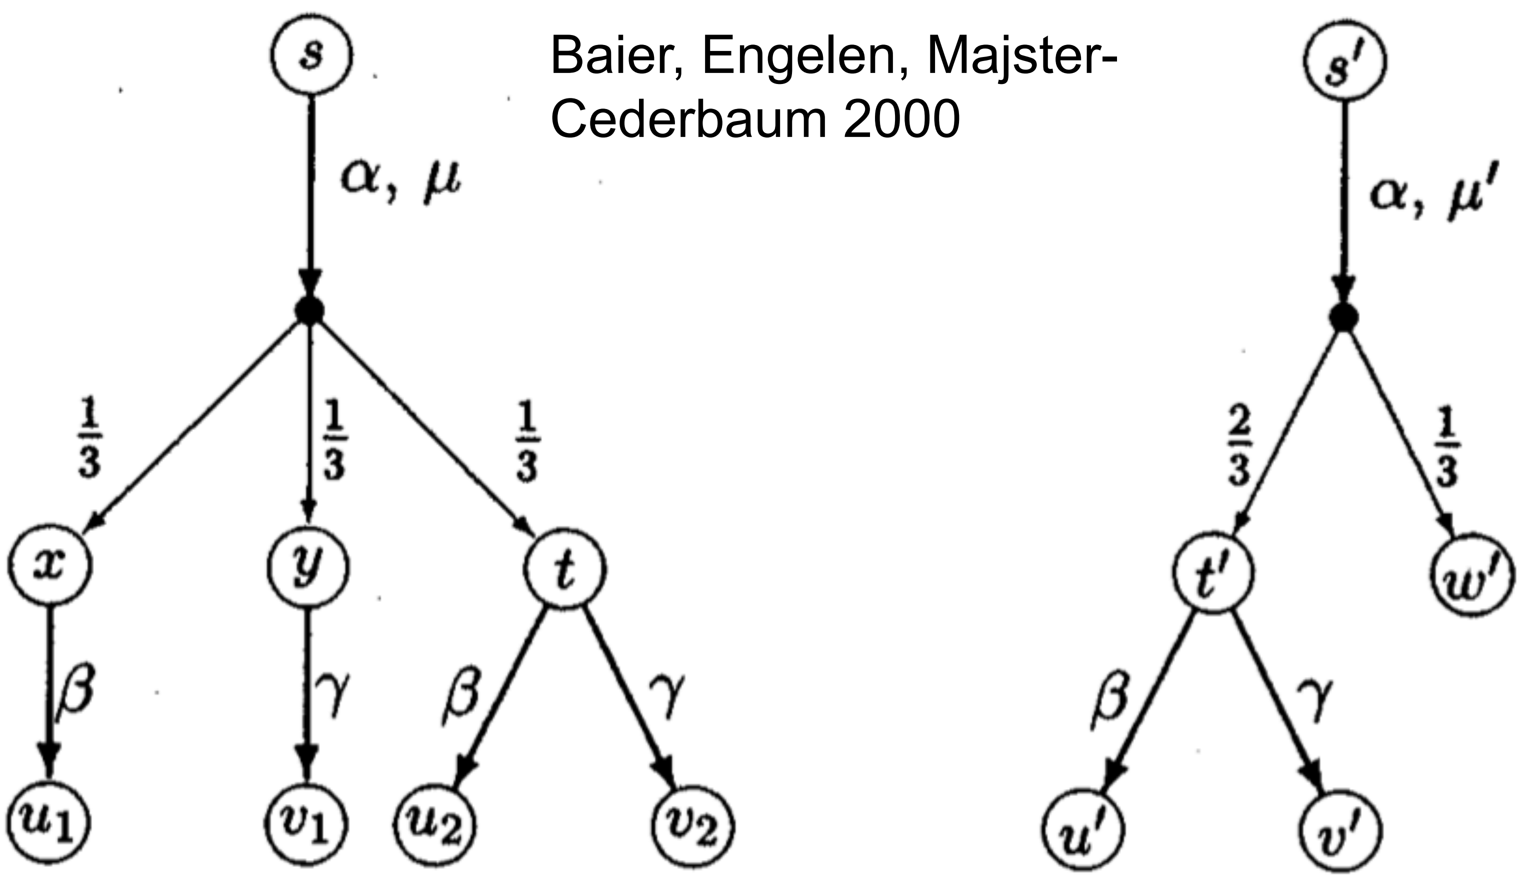

FIG. 4. $s \sqsubseteq_{\uparrow} s'$ and $s \not\sqsubseteq s'$.

More or Less True:

DCTL for CTMDPs

David N. Jansen

FORMATS 2013

Let’s Talk About the Weather

“The sun is shining.” Is this true?

10%

40%

70%

100%

Truth values: not only “false” and “true”, but full interval $[0,1] \subset \mathbb{R}$

Let's Continue Talking About the Weather

"It is going to rain." Is this true?

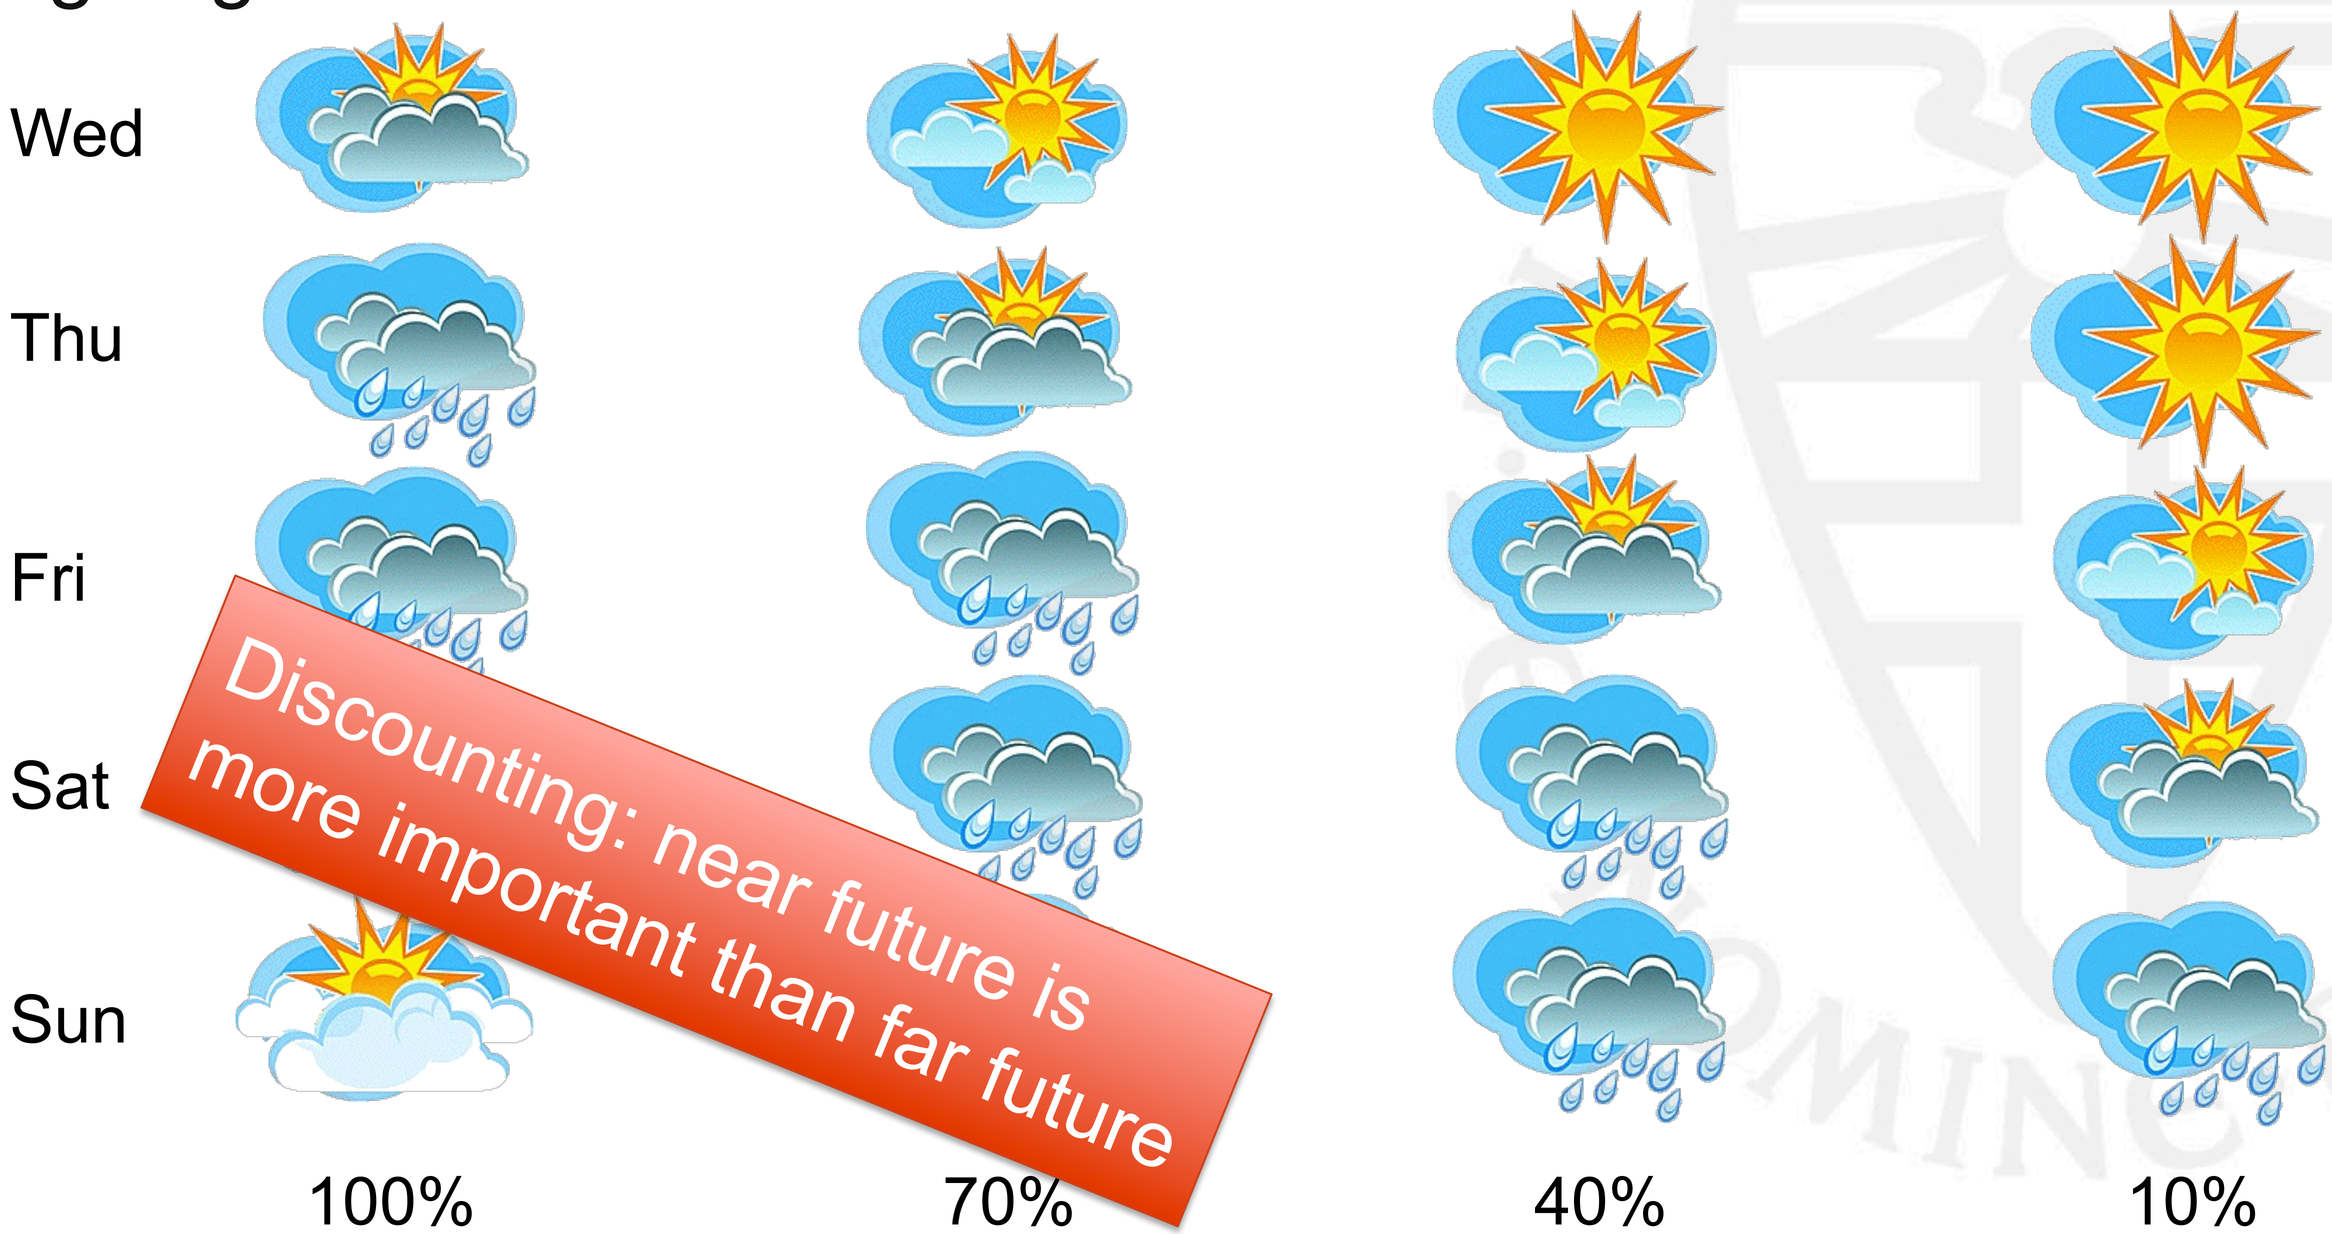

The Logic DCTL: Features

- Truth values: not only “false” and “true”,
  but full interval $[0,1] \subset \mathbb{R}$
  - e.g. express quantitative requirement on degree of sunnyness
  - more robust: Does an incidental cup of 149 ml invalidate spec “The coffee machine shall provide cups of (at least) 150 ml.”?
- Discounting: near future is more important than far future
  (in temporal formulas)
  - e.g. model impatient observer
  - different from strict deadlines in bounded-temporal CTL formulas

The Logic DCTL: History

defined for discrete-time Markov chains

de Alfaro, Faella, Henzinger, Majumdar, Stoelinga:
Model checking discounted temporal properties.
TCS, 2005.

- DCTL definition
- model checking algorithms for labelled transition systems, Markov chains and Markov decision processes

Which Coat Shall I Pack?

good against rain

good against cold

Weather Model

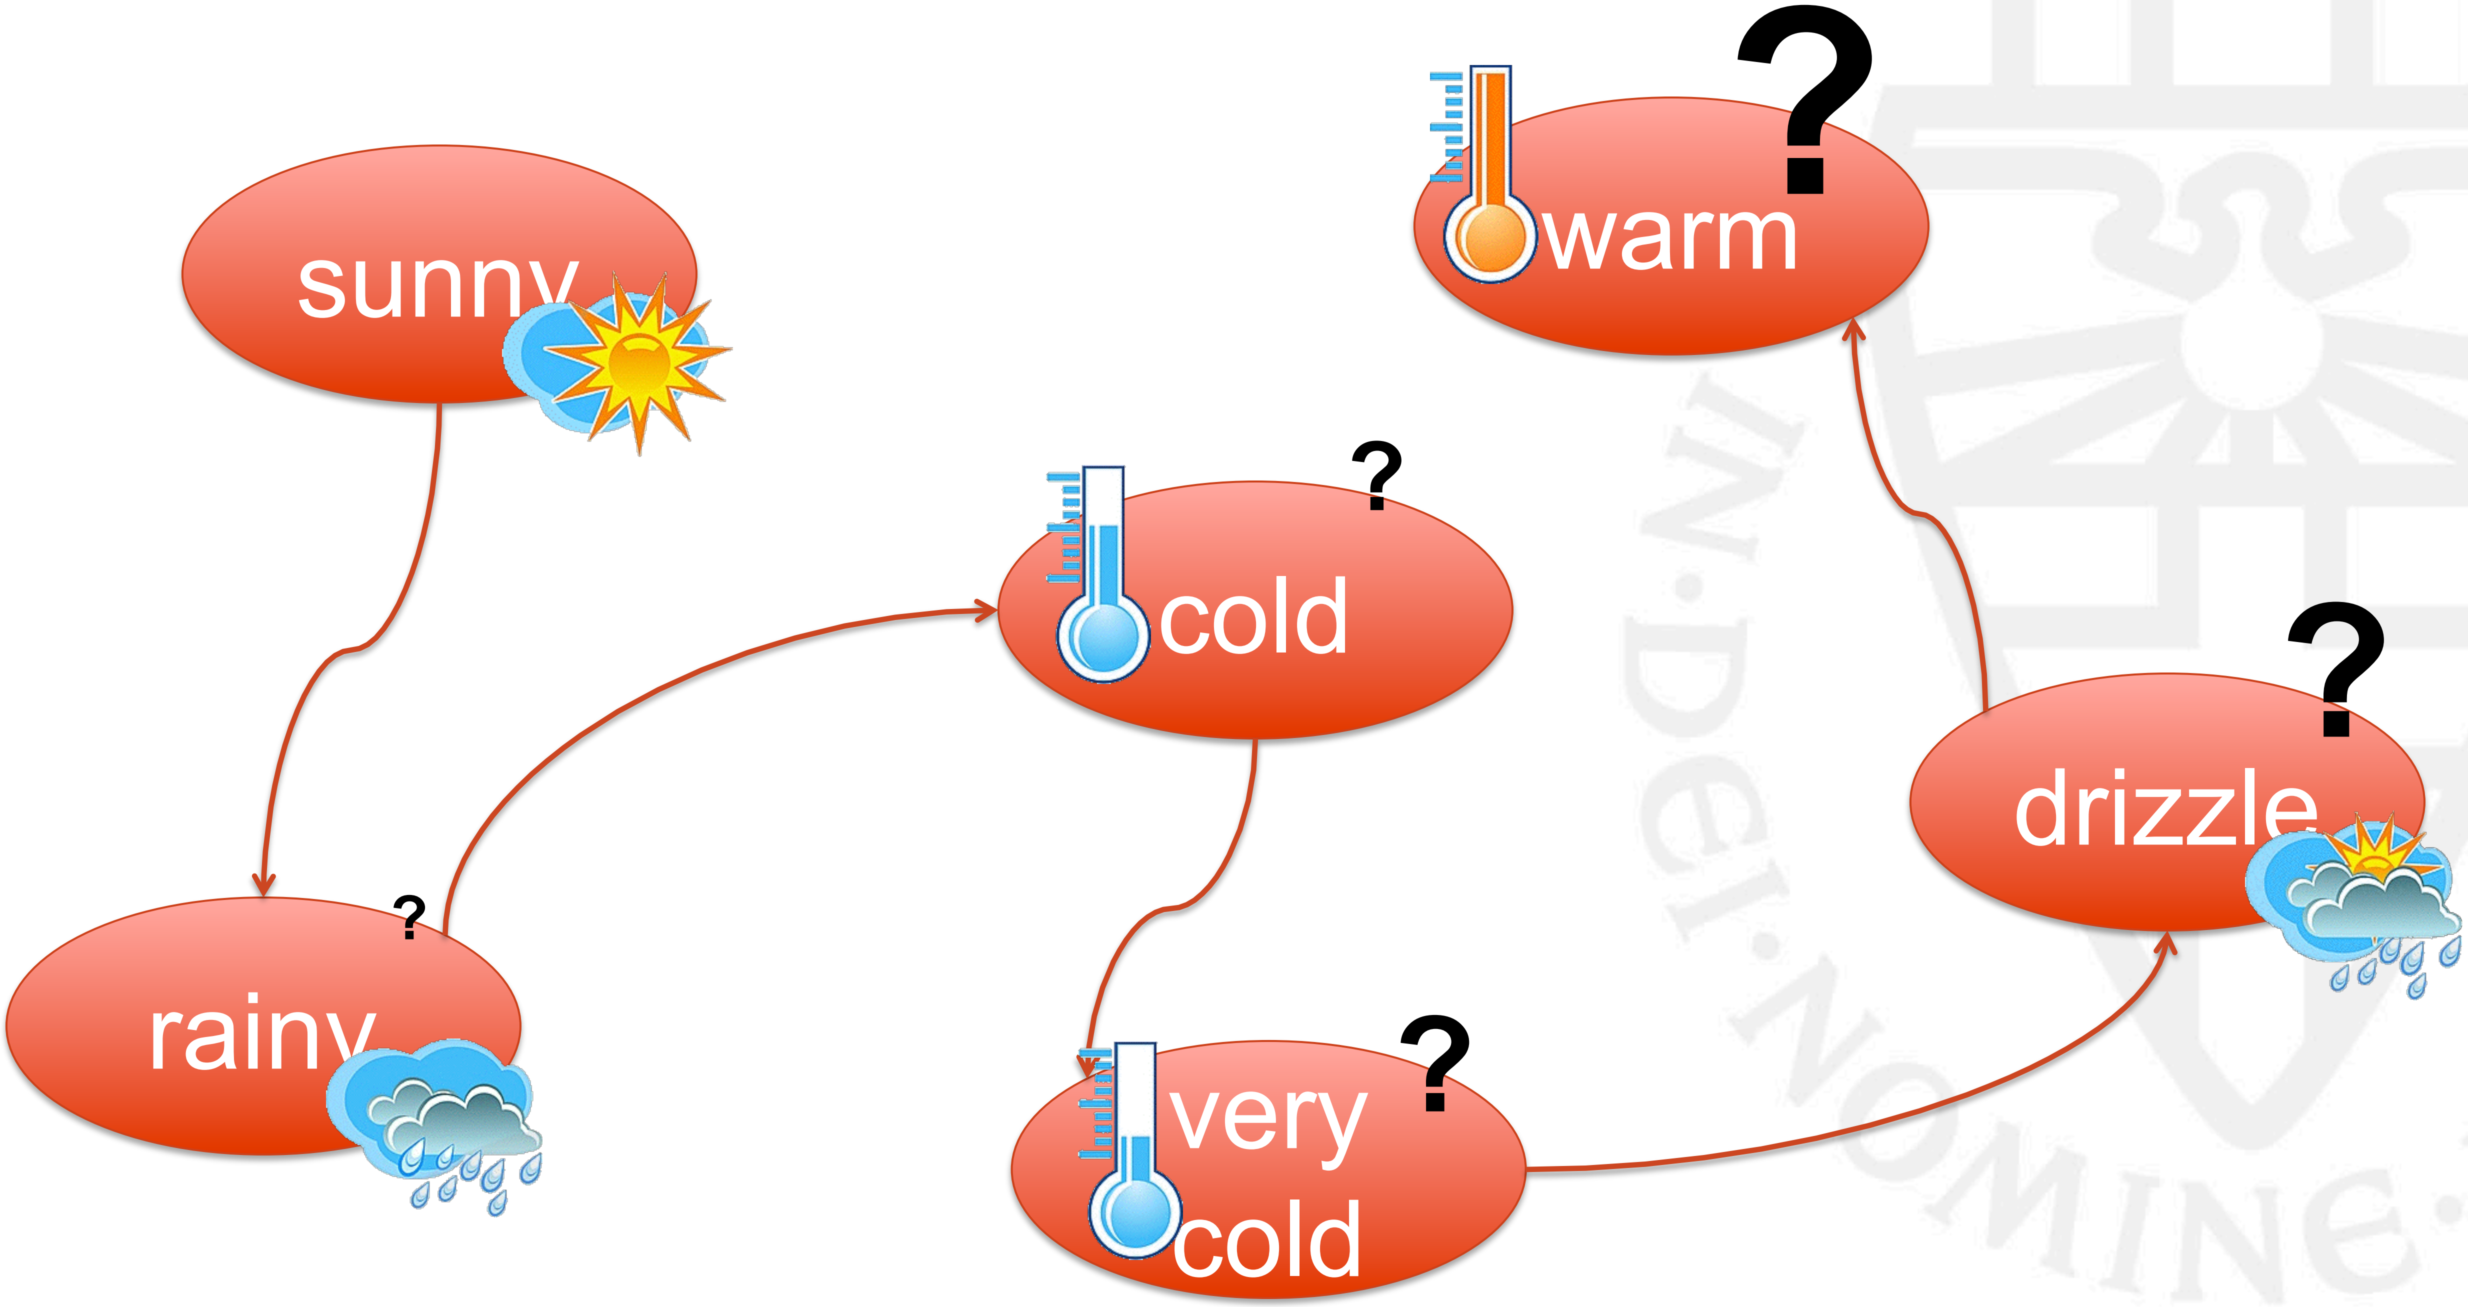

Interesting Questions...

- Should I pack my raincoat?
- Should I pack my winter coat?
- If I can only take one, which one should I pack?

The Logic DCTL: Example Formulas

weighted average between
sunnyness (70%) and warmness (30%)

$$\text{sunny} \oplus_{0.3} \text{warm}$$

pays more attention to near future

expected maximum sunnyness (for an impatient observer)

$$\forall\Diamond_{2}\ \text{sunny}$$

expected average coldness (for a very impatient observer)

$$\forall\triangle_{10}\ \neg\text{warm}$$

rate of losing patience (exponential distribution)

A Possible Application: Testing

- several potential failures of differing impact
  → interval of truth values
- failures appearing at the start of the test are more grave than errors appearing only after many passed tests
  → discounting

- gravity of fatal error
  → expected maximum

- gravity of recoverable error
  → expected average

The Logic DCTL: Syntax

- atomic proposition p
- negation $\neg\varphi$
- conjunction $\varphi \wedge \psi$
- weighted sum $\varphi \oplus_w \psi$ $w \in [0,1]$
- expected maximum $\forall\Diamond_\alpha \varphi$ $\alpha \in [0,\infty)$
- expected minimum $\forall\Box_\alpha \varphi$
- expected average $\forall\Delta_\alpha \varphi$

Let’s Play a Game

Let’s Play a Game

States

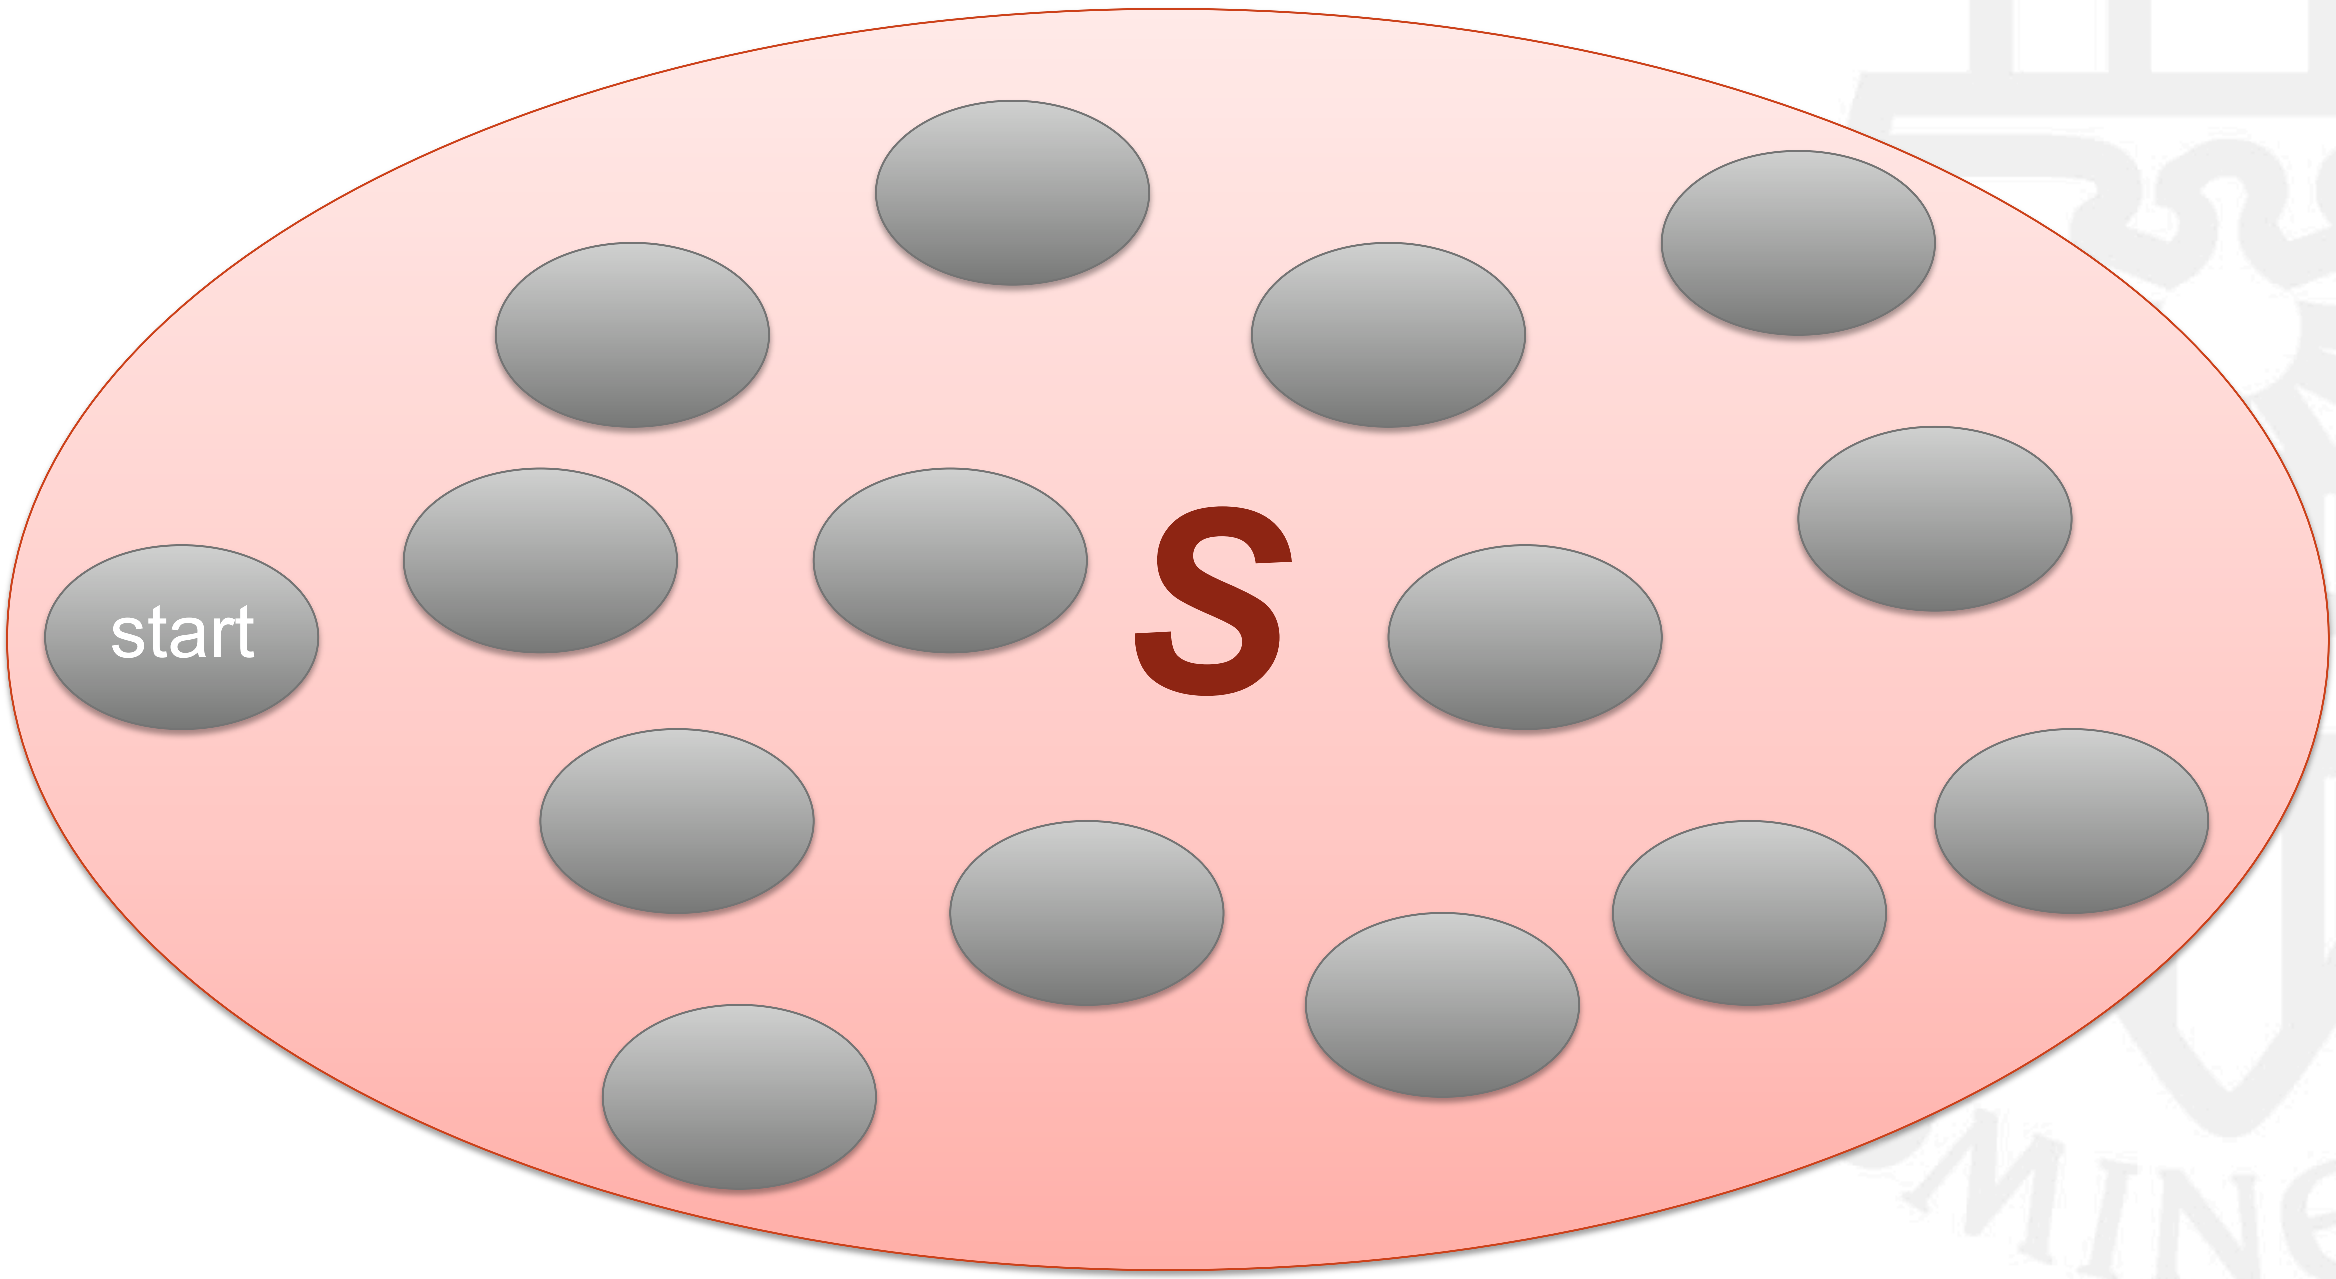

Atomic Propositions or Labels

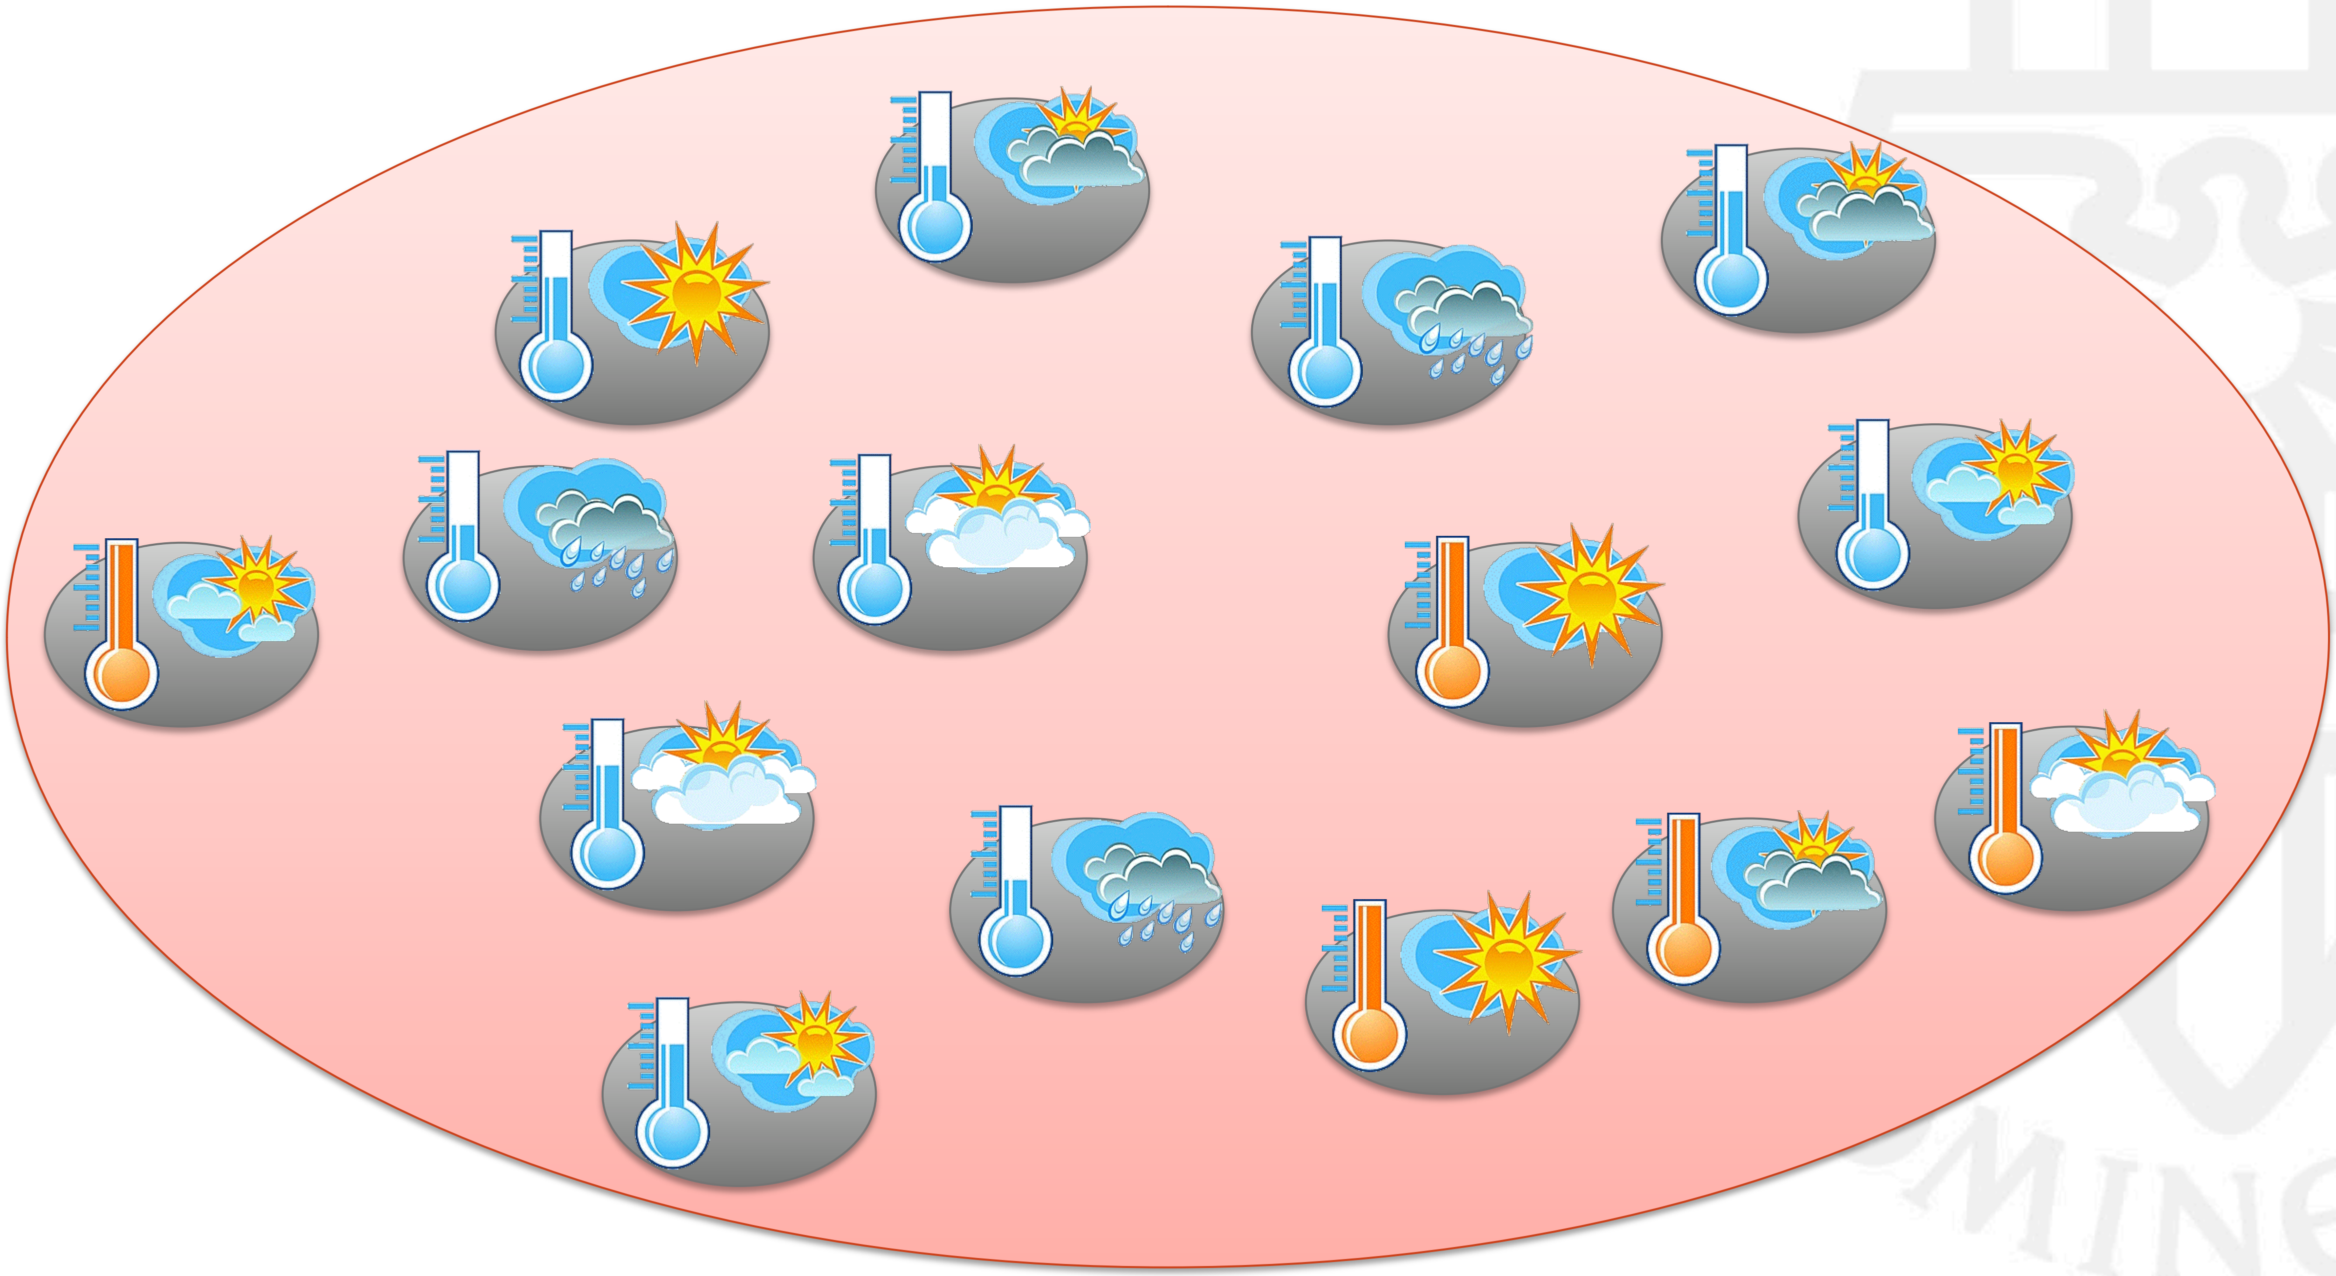

Transitions

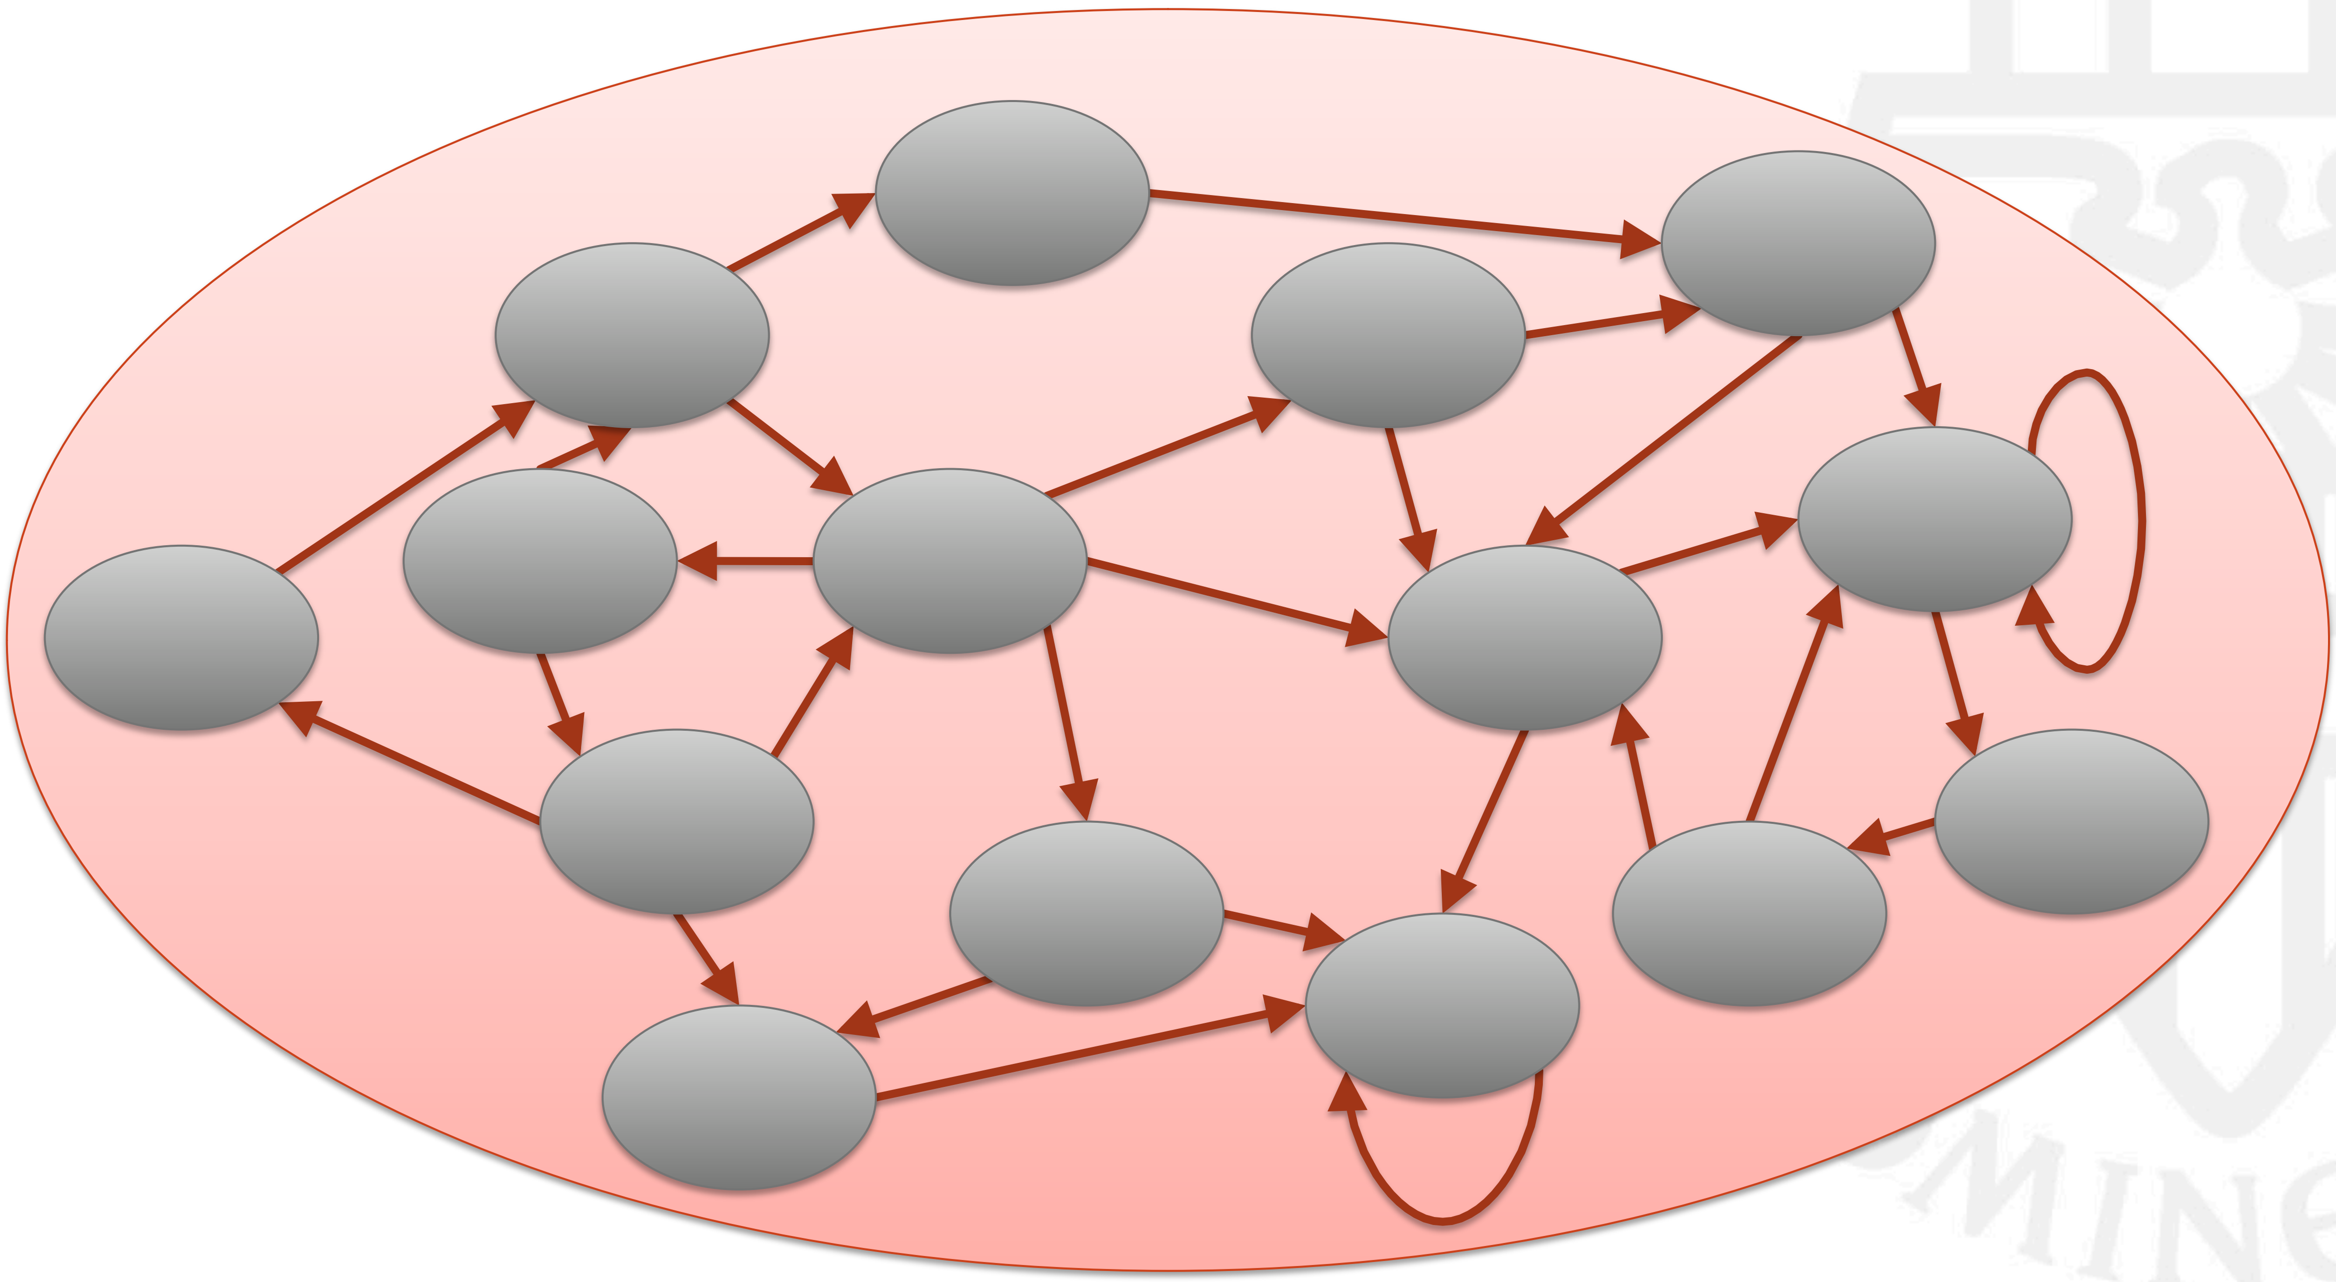

Transition: Action + Rate

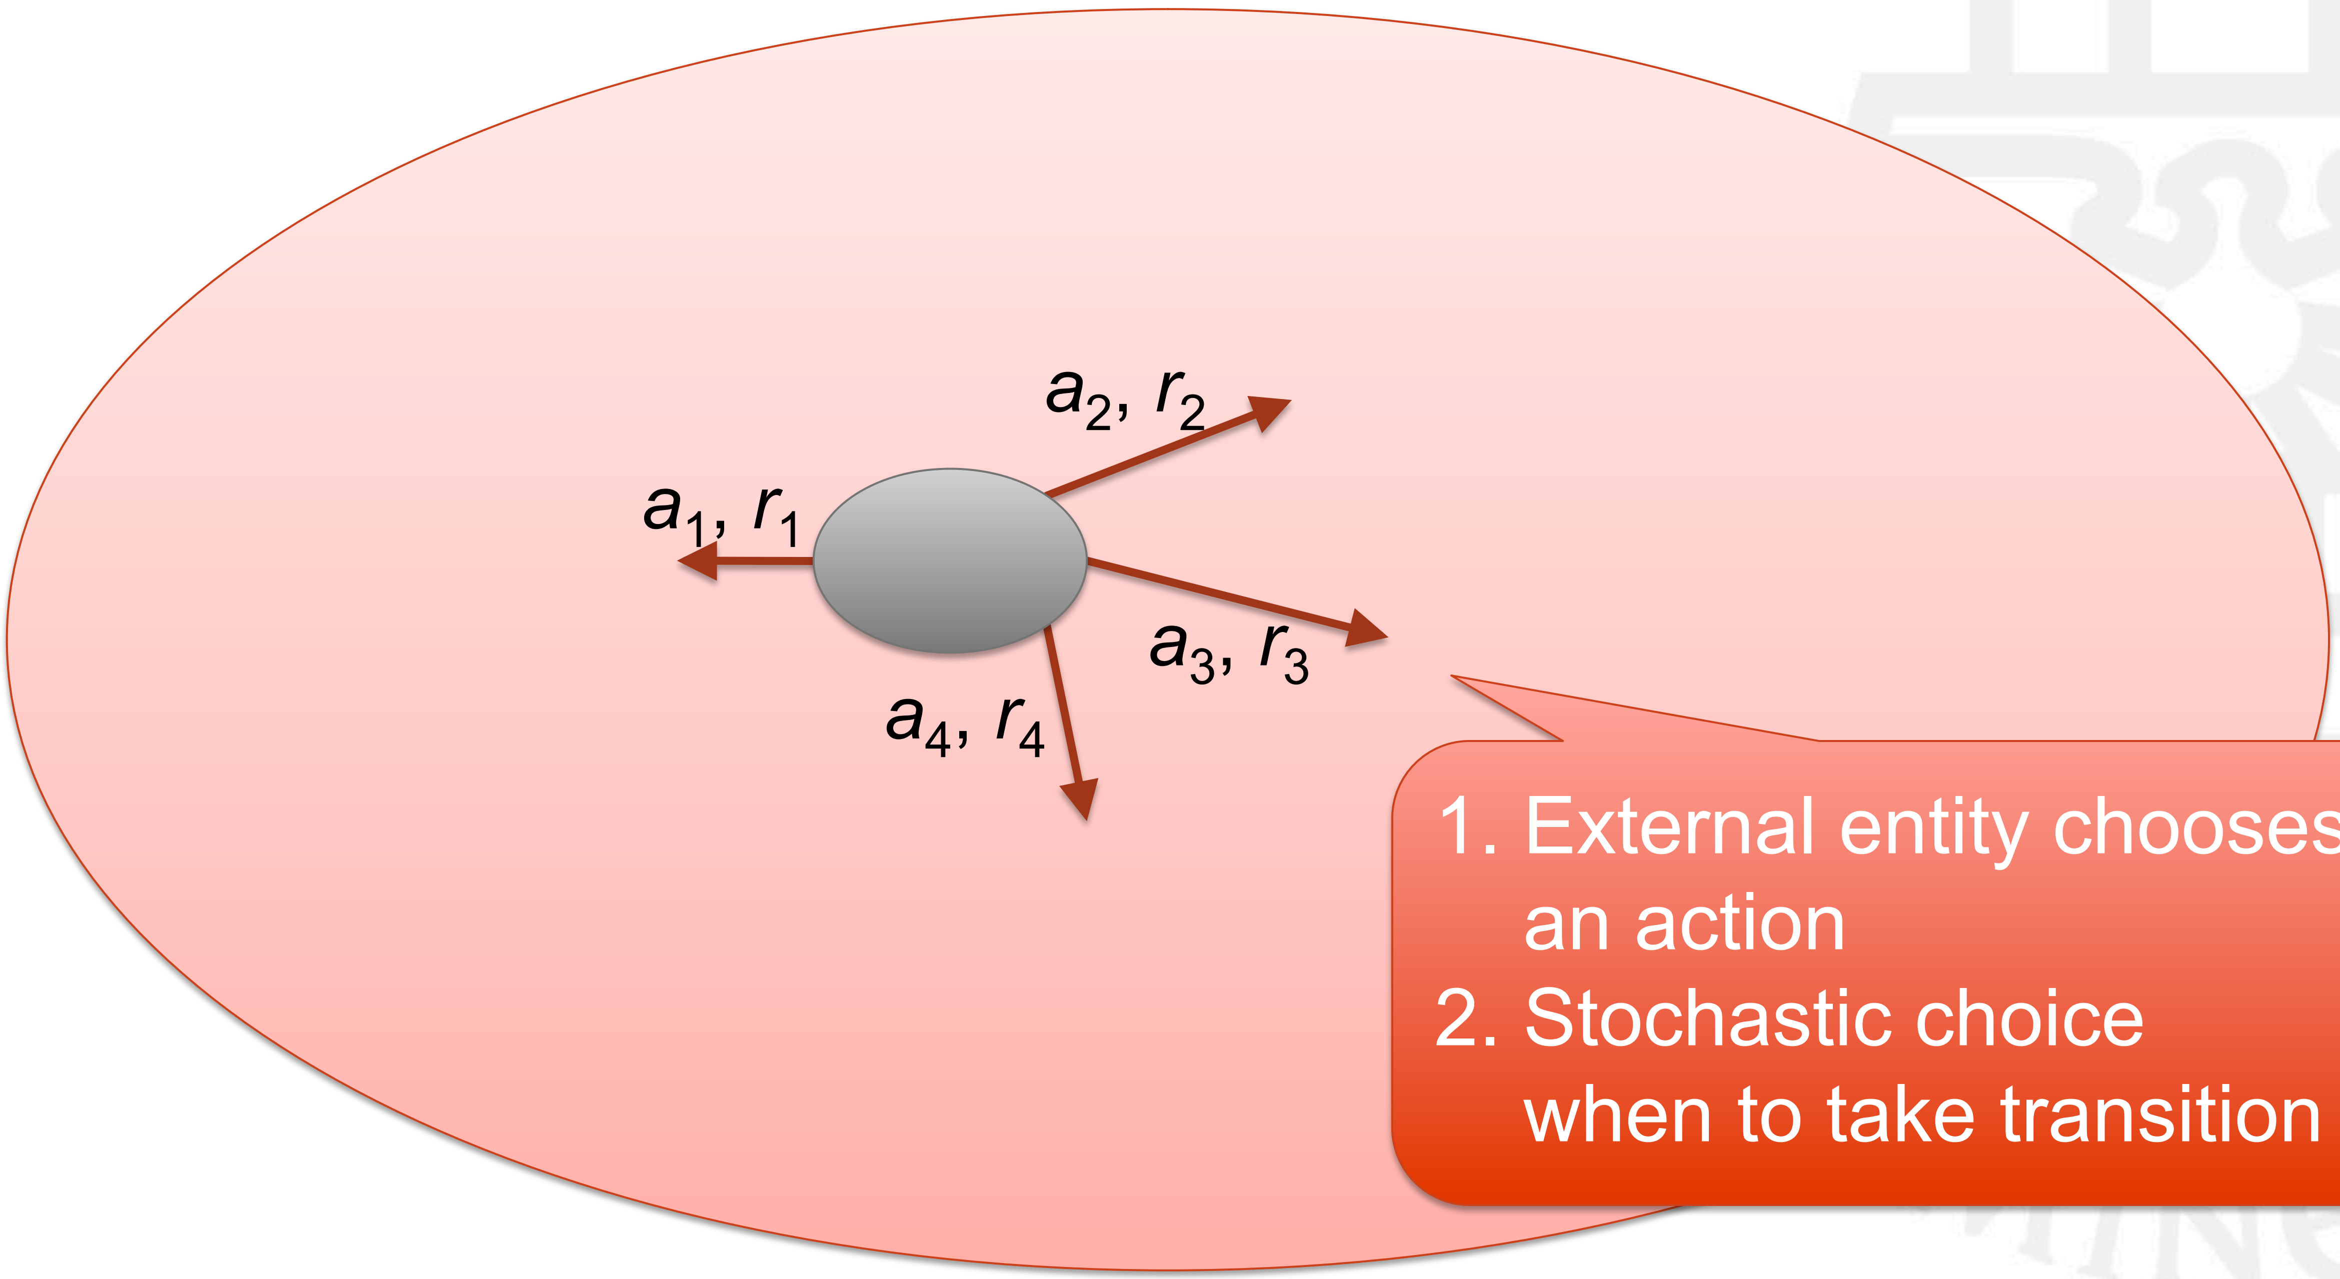

1. External entity chooses an action
2. Stochastic choice when to take transition

Continuous-Time Markov Decision Process

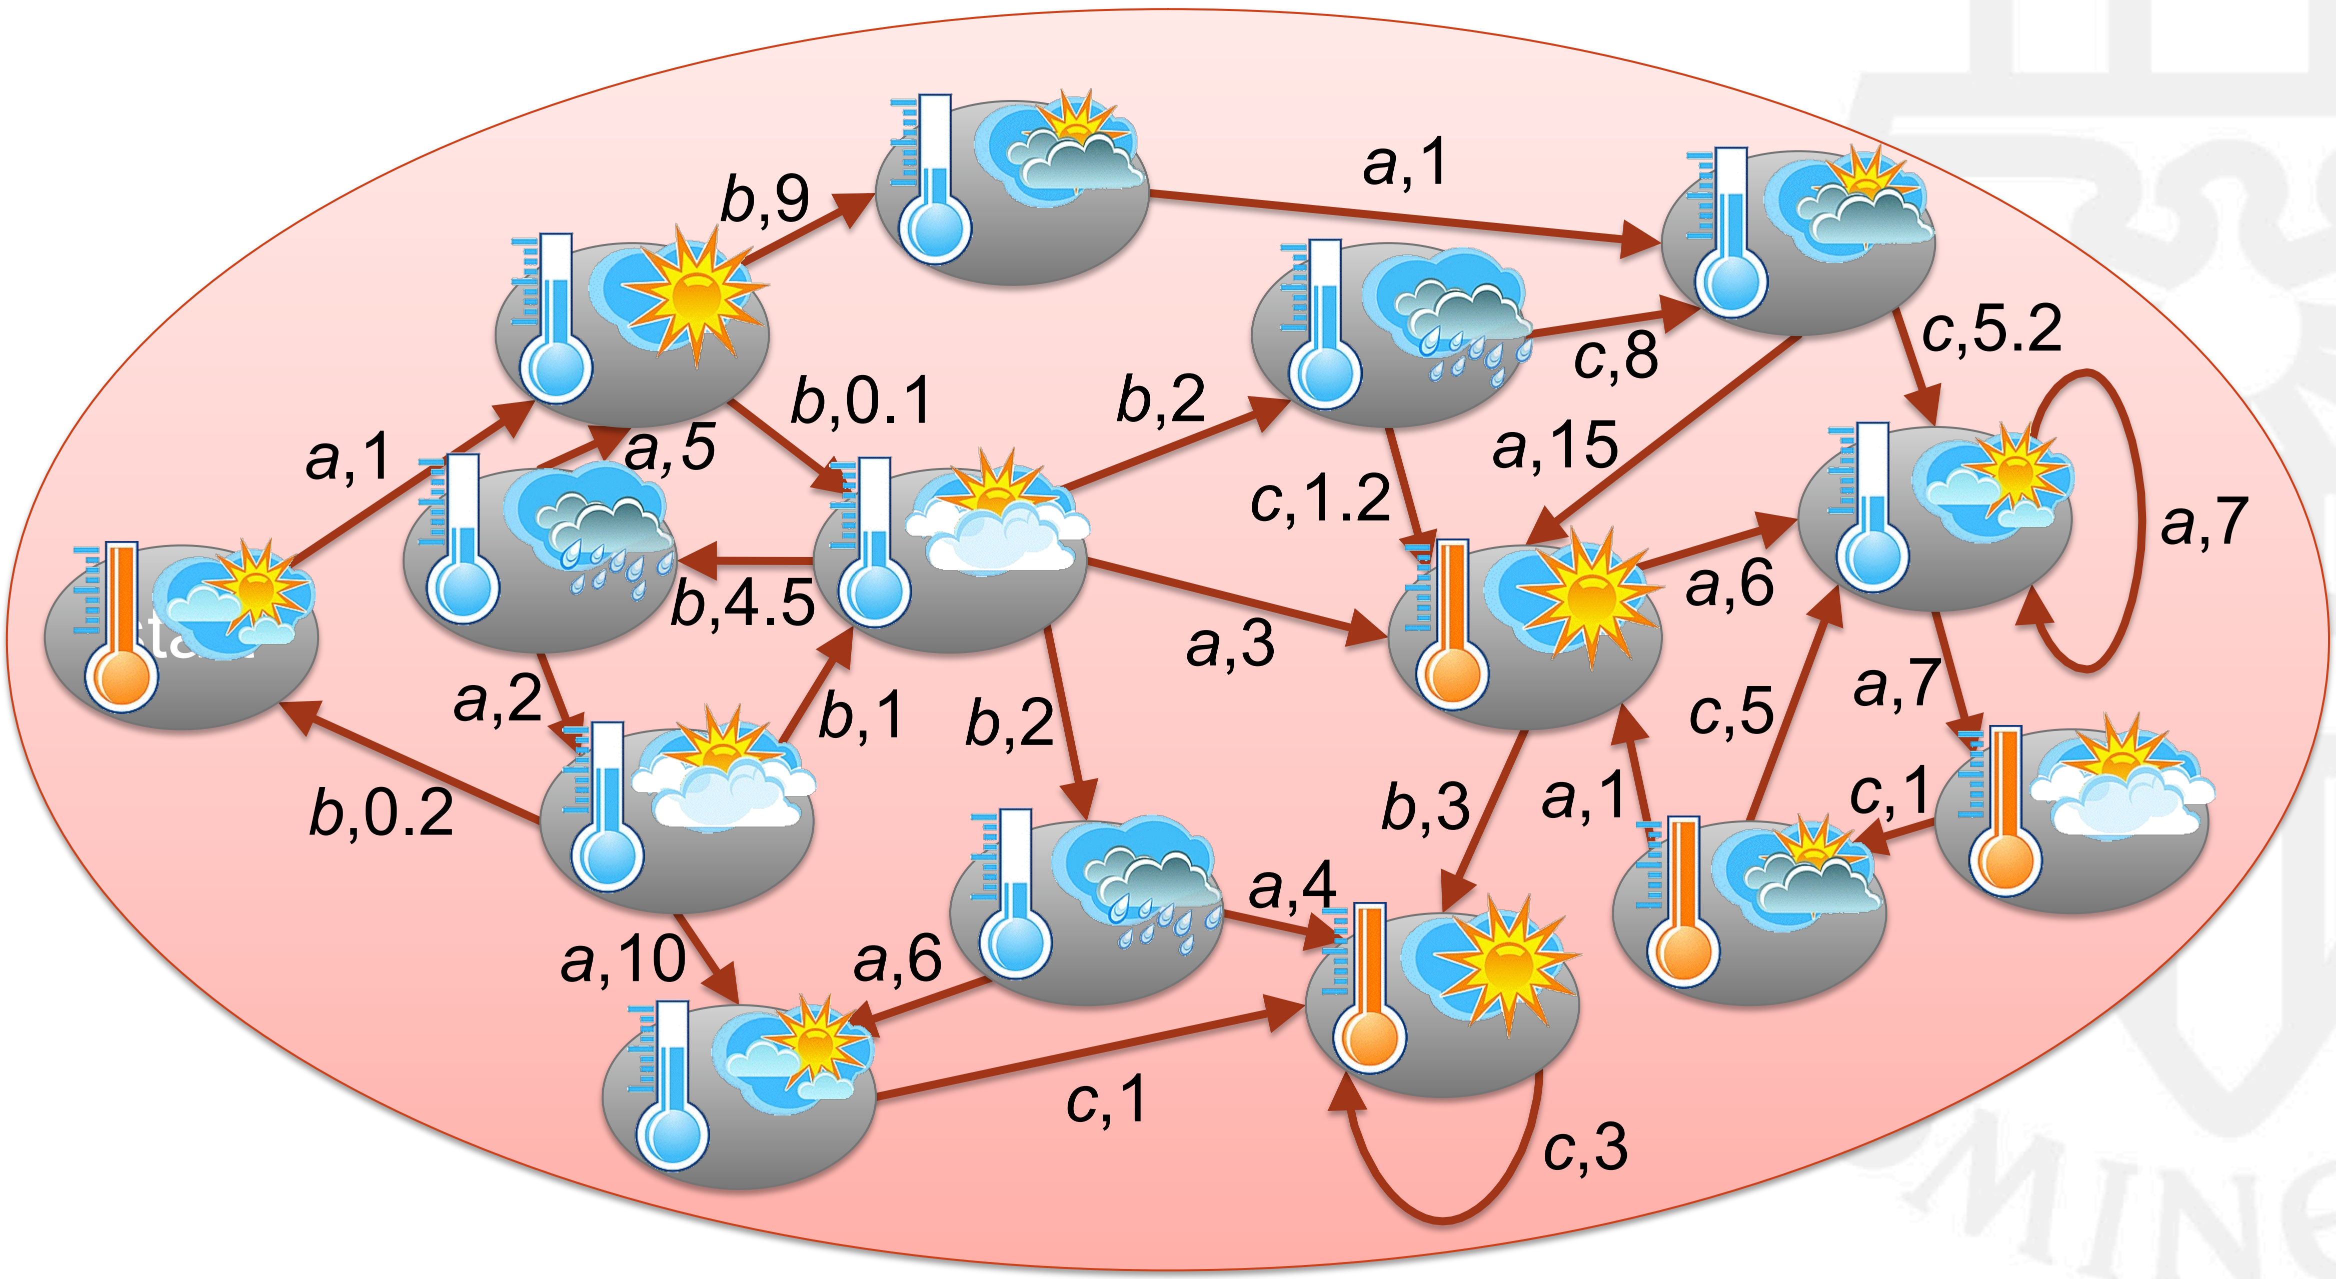

Exponential Distribution

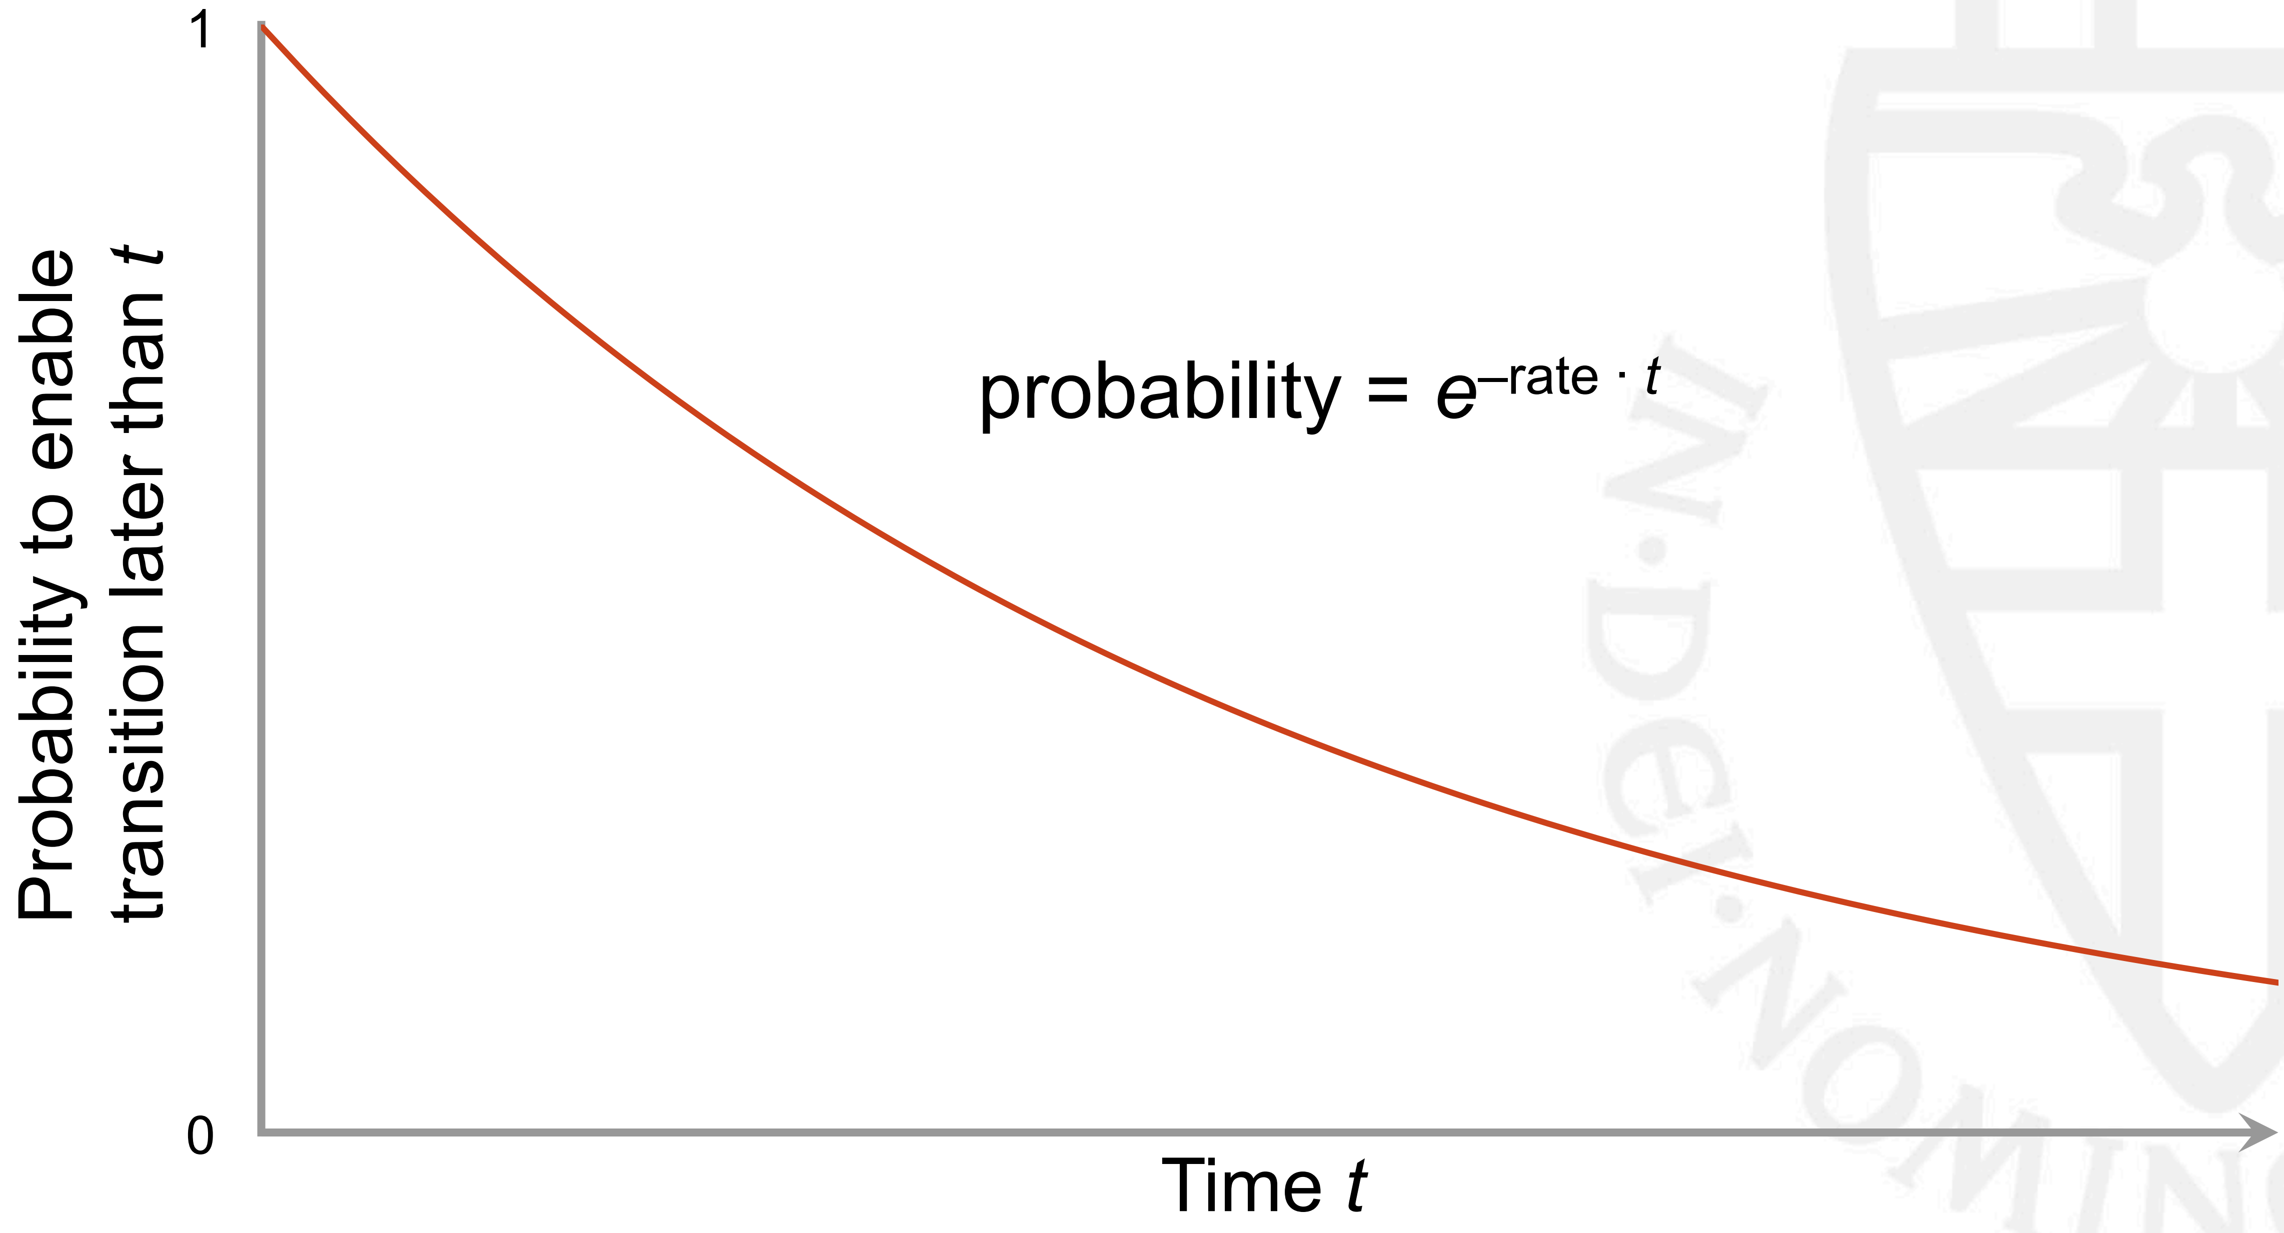

Continuous-Time Markov Decision Process

A CTMDP consists of:

- S — finite set of states
- A — finite set of actions
- $\mathbf{R}: S \times A \times S \rightarrow \mathbb{R}_{\geq 0}$ — transition rate matrix

  or

  $\mathbf{Q}: S \times A \times S \rightarrow \mathbb{R}$ — infinitesimal generator matrix

  (for all $i \in S$ and $a \in A$, $\Sigma_j \mathbf{Q}^a{}_{ij} = 0$)
- $L: S \times AP \rightarrow \cancel{\{0,1\}}$ $[0,1]$ — labelling with atomic propositions

The Logic DCTL: Semantics

interpretation of formula φ in state s is $[\![\varphi]\!](s) \in [0,1]$

- $[\![p]\!](s) = L(s,p)$
- $[\![\neg\varphi]\!](s) = 1 - [\![\varphi]\!](s)$
- $[\![\varphi \wedge \psi]\!](s) = \min \{ [\![\varphi]\!](s), [\![\psi]\!](s) \}$
- $[\![\varphi \oplus_w \psi]\!](s) = (1 - w)\, [\![\varphi]\!](s) + w\, [\![\psi]\!](s)$

The Two Semantics of $\forall\Diamond\varphi$ (in CTL)

- Fixpoint semantics: (least) solution of

$$u = \varphi \vee \forall\bigcirc u$$

$$u(\bullet) = \max \{ [\![\varphi]\!](\bullet), \min_{s \in succ(\bullet)} u(s) \}$$

- Path semantics:

$$\min_{\sigma \in Paths} \max_{n \in \{0,1,\ldots\}} [\![\varphi]\!](\sigma @ n)$$

The two semantics coincide in CTL

... but they differ in discounted setting!

The Fixpoint Semantics of $\forall\Diamond_\alpha \varphi$

(Least) solution of

$$u = \varphi \vee \forall\bigcirc_\alpha u$$

$$u(\bullet) = \max \{ [\![\varphi]\!](\bullet), \min_{a \in A} \mathbb{E}_a \, e^{-\alpha T} u(X) \}$$

- $e^{-\alpha T}$ discount for waiting until transition is taken
- T random variable for waiting time

- $\mathbb{E}_a \, e^{-\alpha T} u(X)$ discounted expectation over next state
- X random variable for next state

- u is a function $S \to [0,1]$

The Fixpoint Semantics of $\forall\Diamond_\alpha \varphi$

(Least) solution of

$$u = \varphi \vee \forall\bigcirc_\alpha u$$

$$u(\bullet) = \max \{ [\![\varphi]\!](\bullet), \min_{a \in A} \frac{1}{E^a(\bullet)+\alpha} \Sigma_{s' \in succ(\bullet)} \mathbf{R}^a(\bullet, s')u(s') \}$$

The Fixpoint Semantics of $\forall\Diamond_\alpha \varphi$

(Least) solution of

$$u = \varphi \vee \forall\bigcirc_\alpha u$$

$$u(\bullet) = \max \{ [\![\varphi]\!](\bullet), \min_{a \in A} \frac{1}{E^a(\bullet)+\alpha} \Sigma_{s' \in succ(\bullet)} \mathbf{R}^a(\bullet,s')u(s') \}$$

- can be formulated as linear program:

Minimize $\Sigma_{s \in S}\, v(s)$ subject to

– $v(s) \geq [\![\varphi]\!](s)$ for all $s \in S$

– $v(s) \geq \frac{E^a(\bullet)}{E^a(\bullet)+\alpha} \Sigma_{s' \in succ(s)} P^a(s,s')v(s')$ for all $s \in S$

same type of linear program as in DTMDPs

Model Checking the Fixpoint Semantics

- Other operators also allow reduction to discrete-time case

- Model checking algorithm:
  1. Uniformise CTMDP
     (so exit rate E no longer depends on current state + action)
  2. Reduce to discrete-time Markov chain
  3. Apply discrete-time algorithm with discount factor $E/(E+\alpha)$

The Path Semantics of $\forall\Diamond_\alpha \varphi$

Look at complete path at once:

$$\cancel{\min_{D\in Scheduler}}\ \mathbb{E}\ \sup_{t\in[0,\infty)} e^{-\alpha t}\ [\![\varphi]\!](\sigma@t)$$

– $\sup_{t\in[0,\infty)}$	supremum over all time points
– $e^{-\alpha t}$	discount at time t
– $\sigma@t$	random variable for state at time t
– ~~$\min_{D\in Scheduler}$~~	~~any scheduler class in CTMDP~~

Expected Supremum

Function of path and time

$$(\sigma, t) \mapsto e^{-\alpha t} [\![\text{sunny}]\!](\sigma @ t)$$

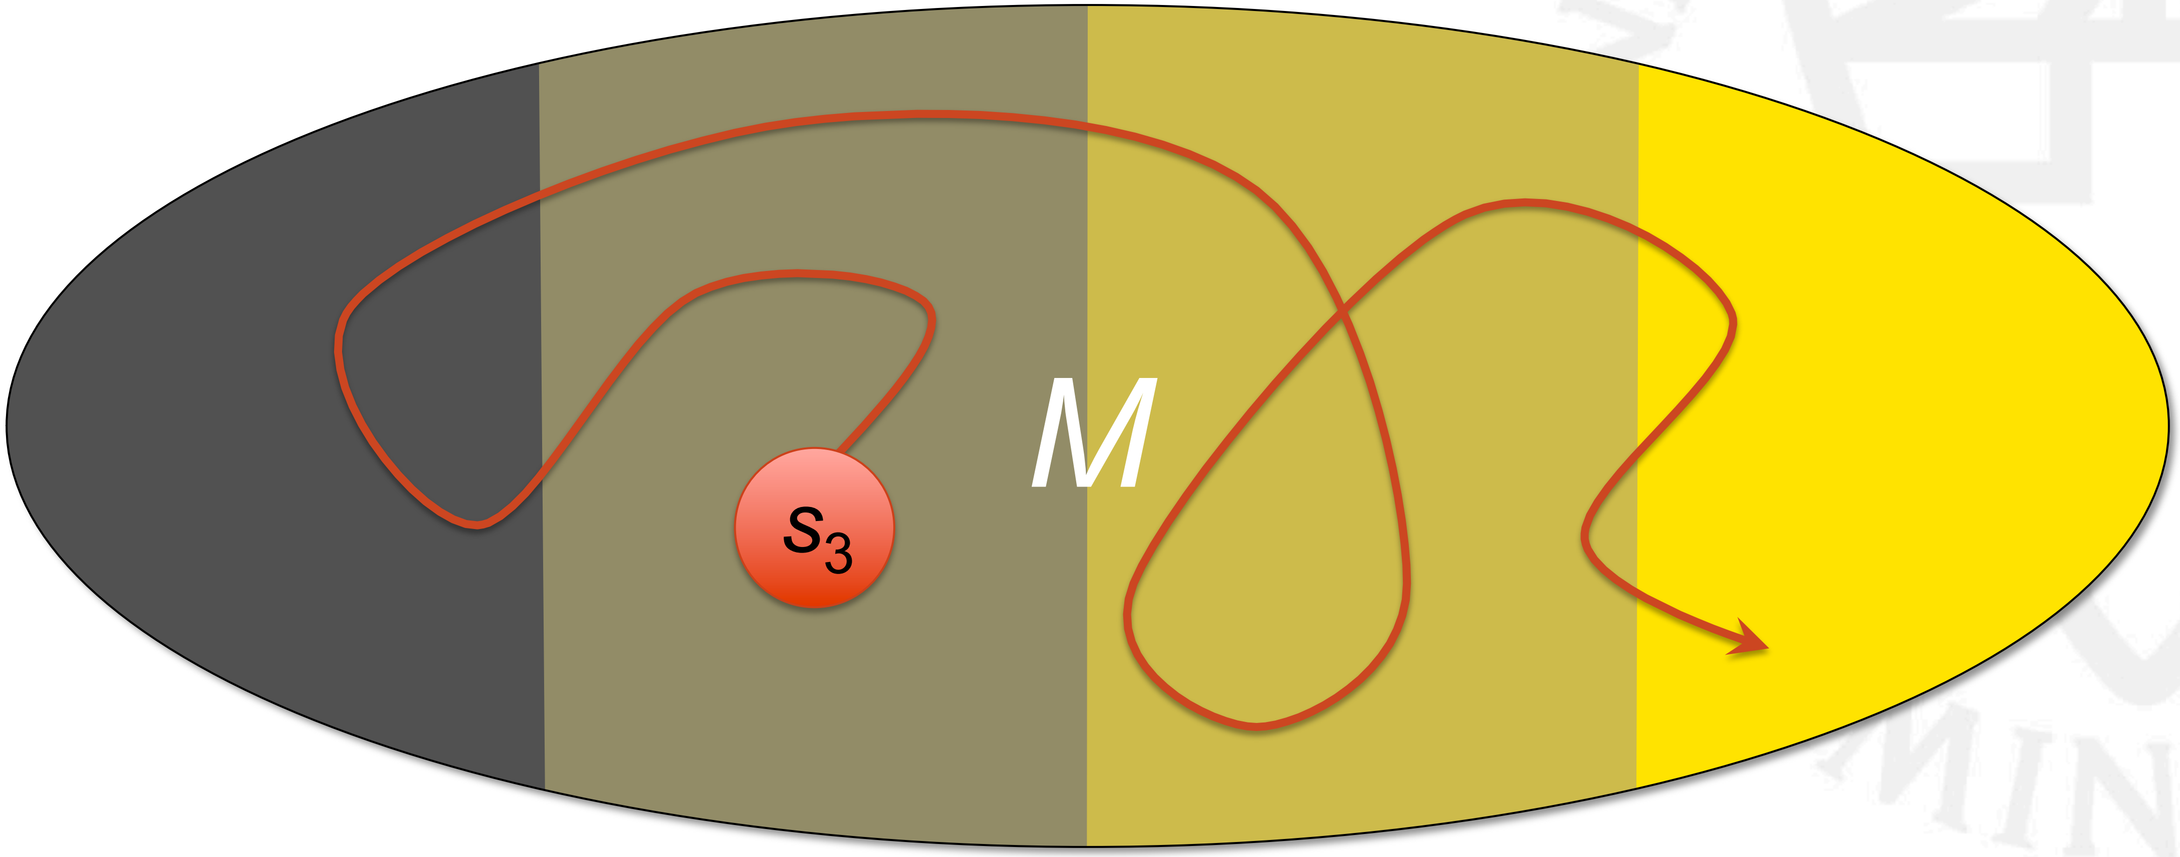

Expected Supremum

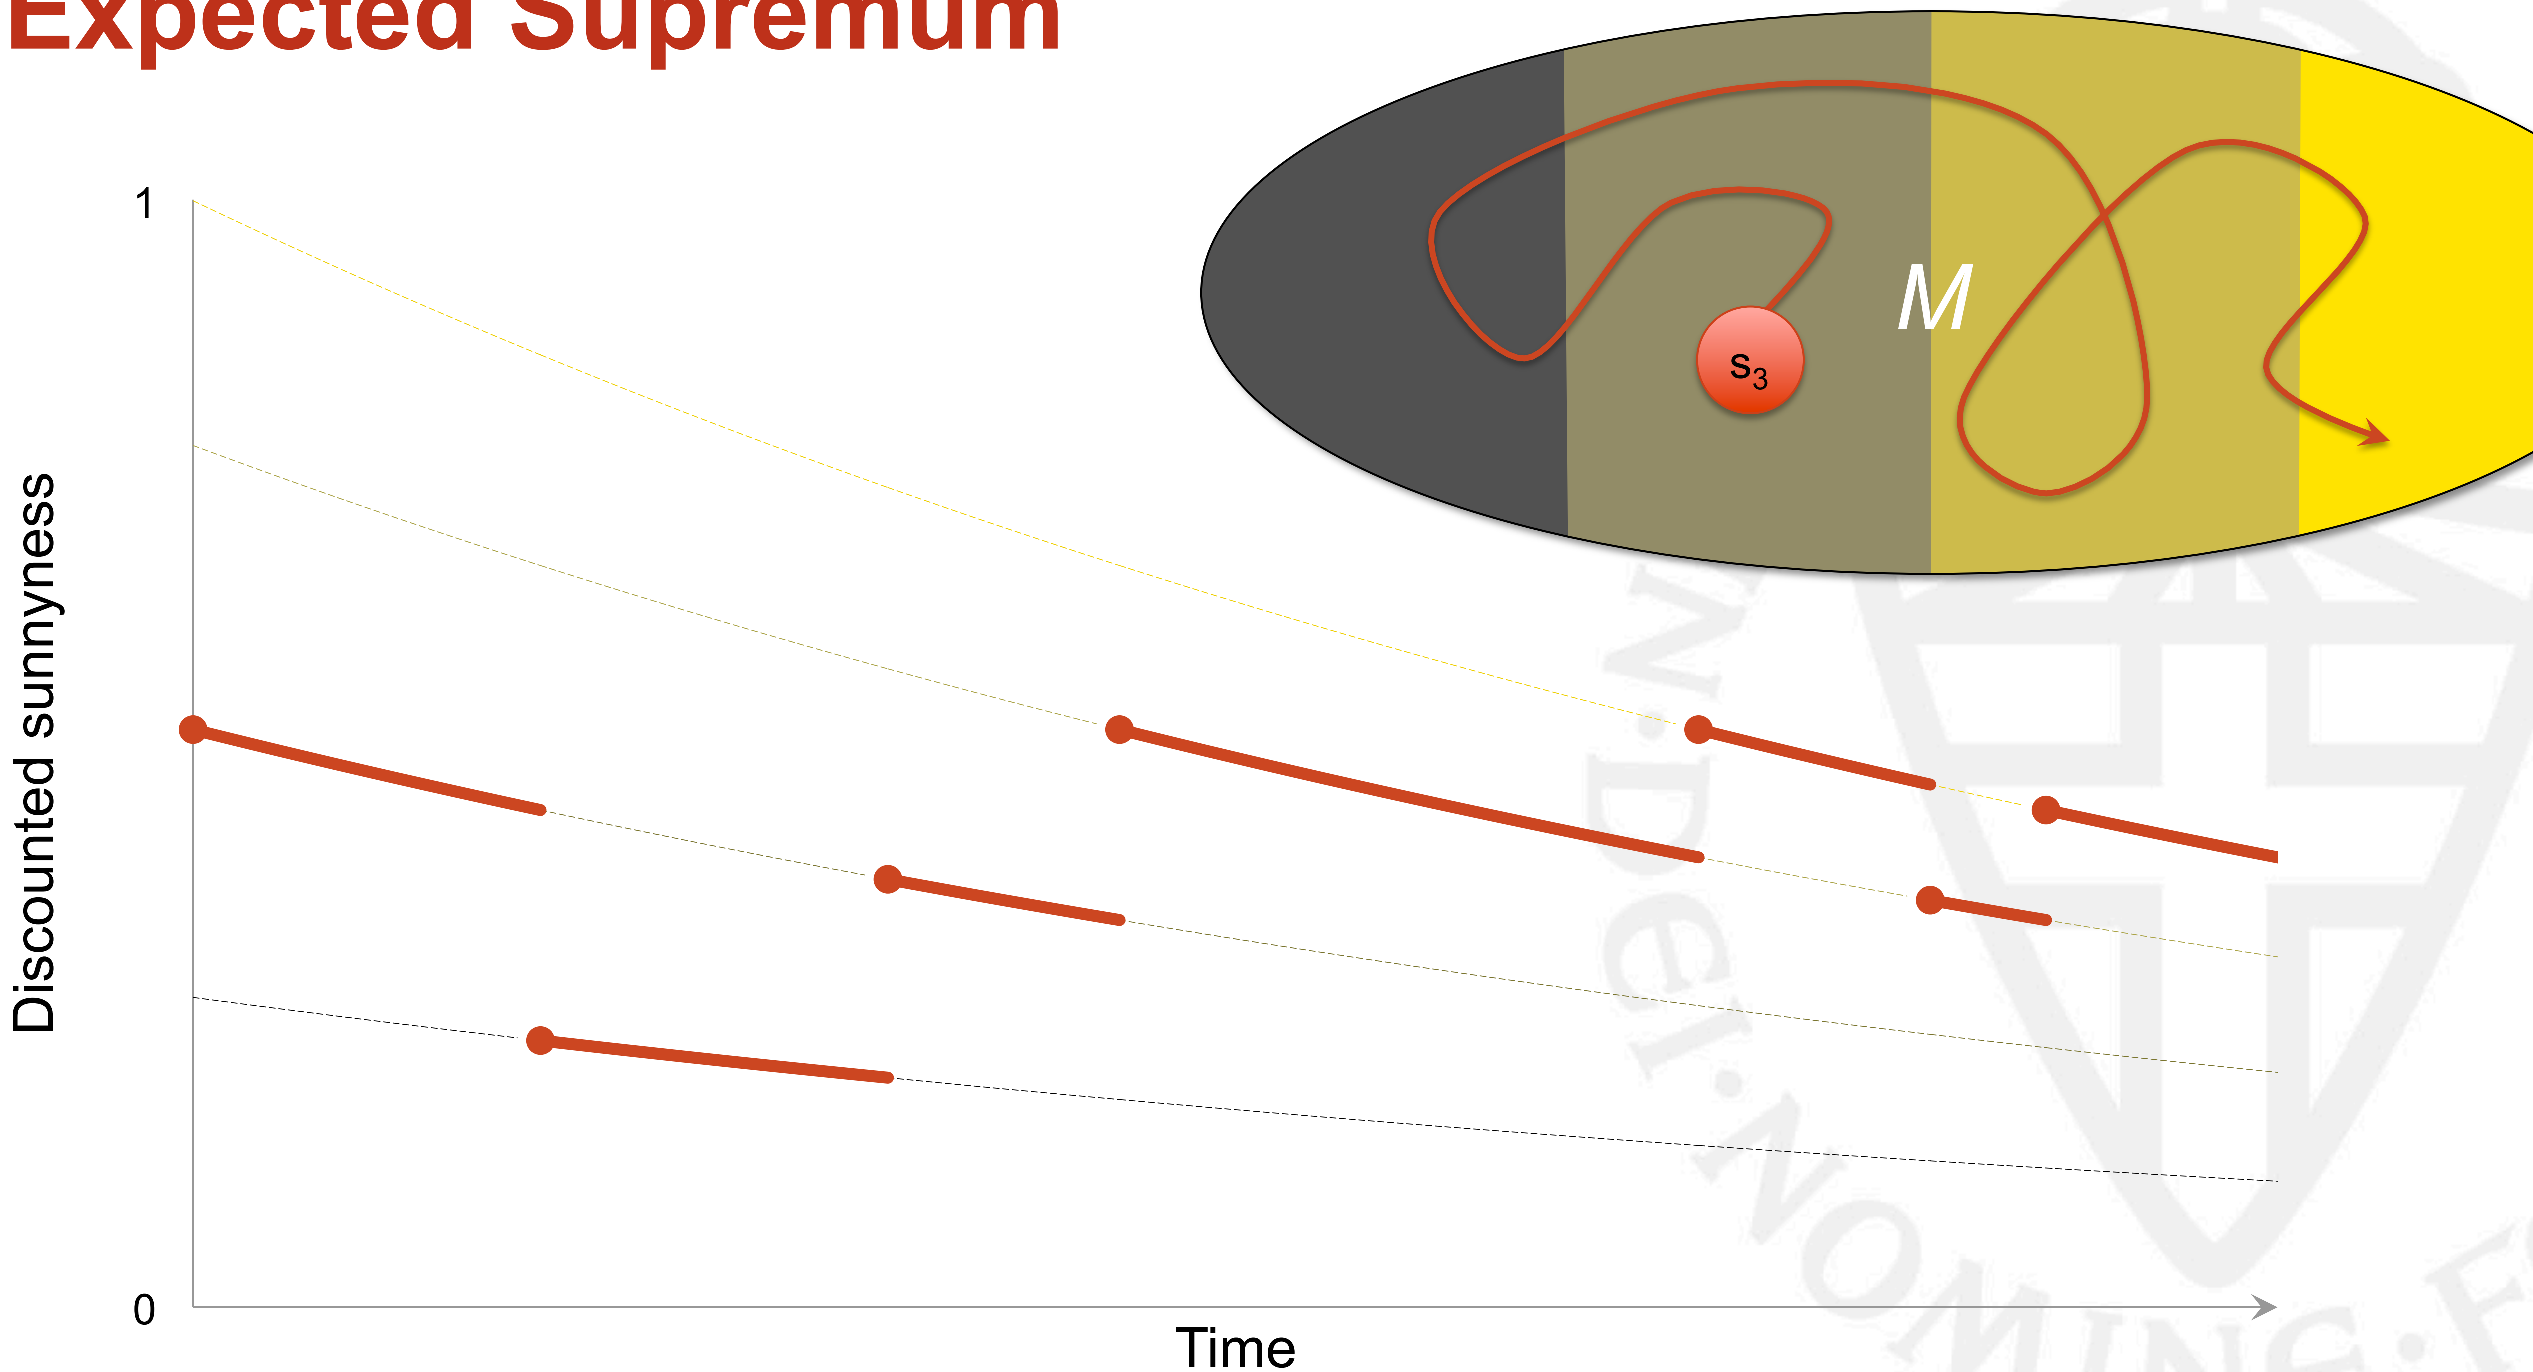

Observation

only first entry into more sunny class can improve
$[\![\forall\Diamond_\alpha \text{ sunny}]\!]^{path}$ over $[\![\text{sunny}]\!]$

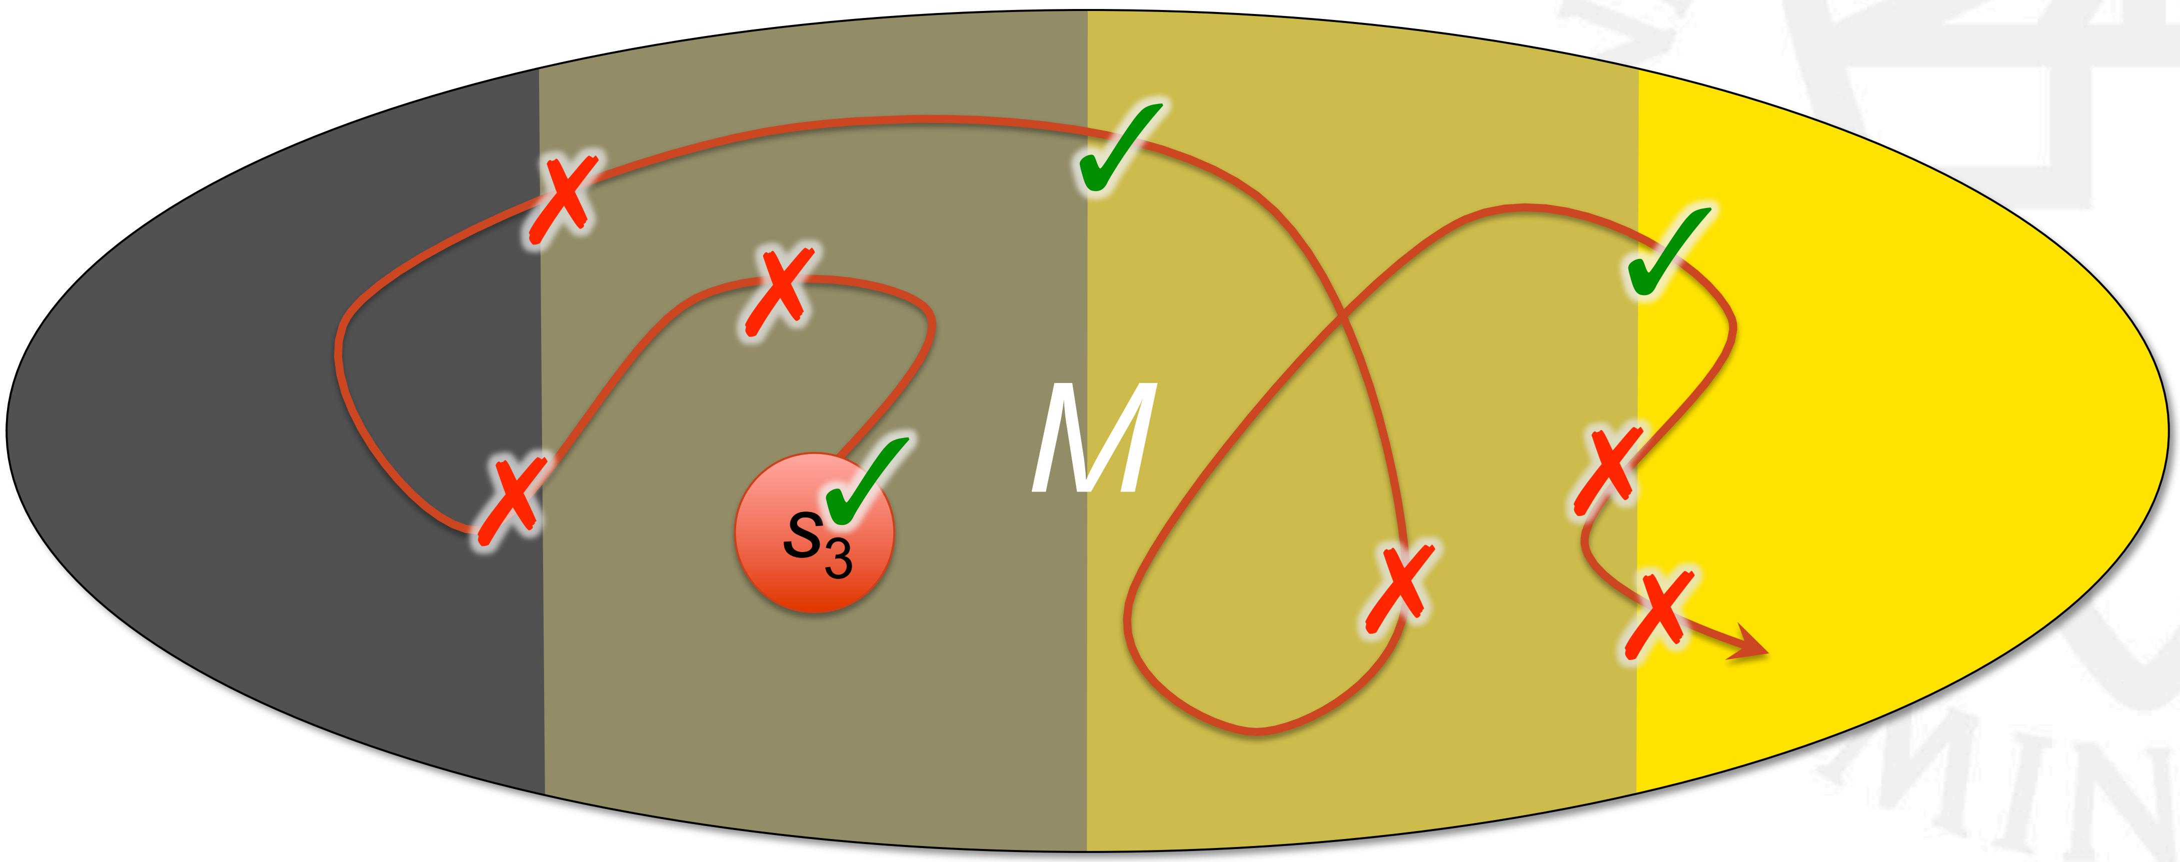

Observation

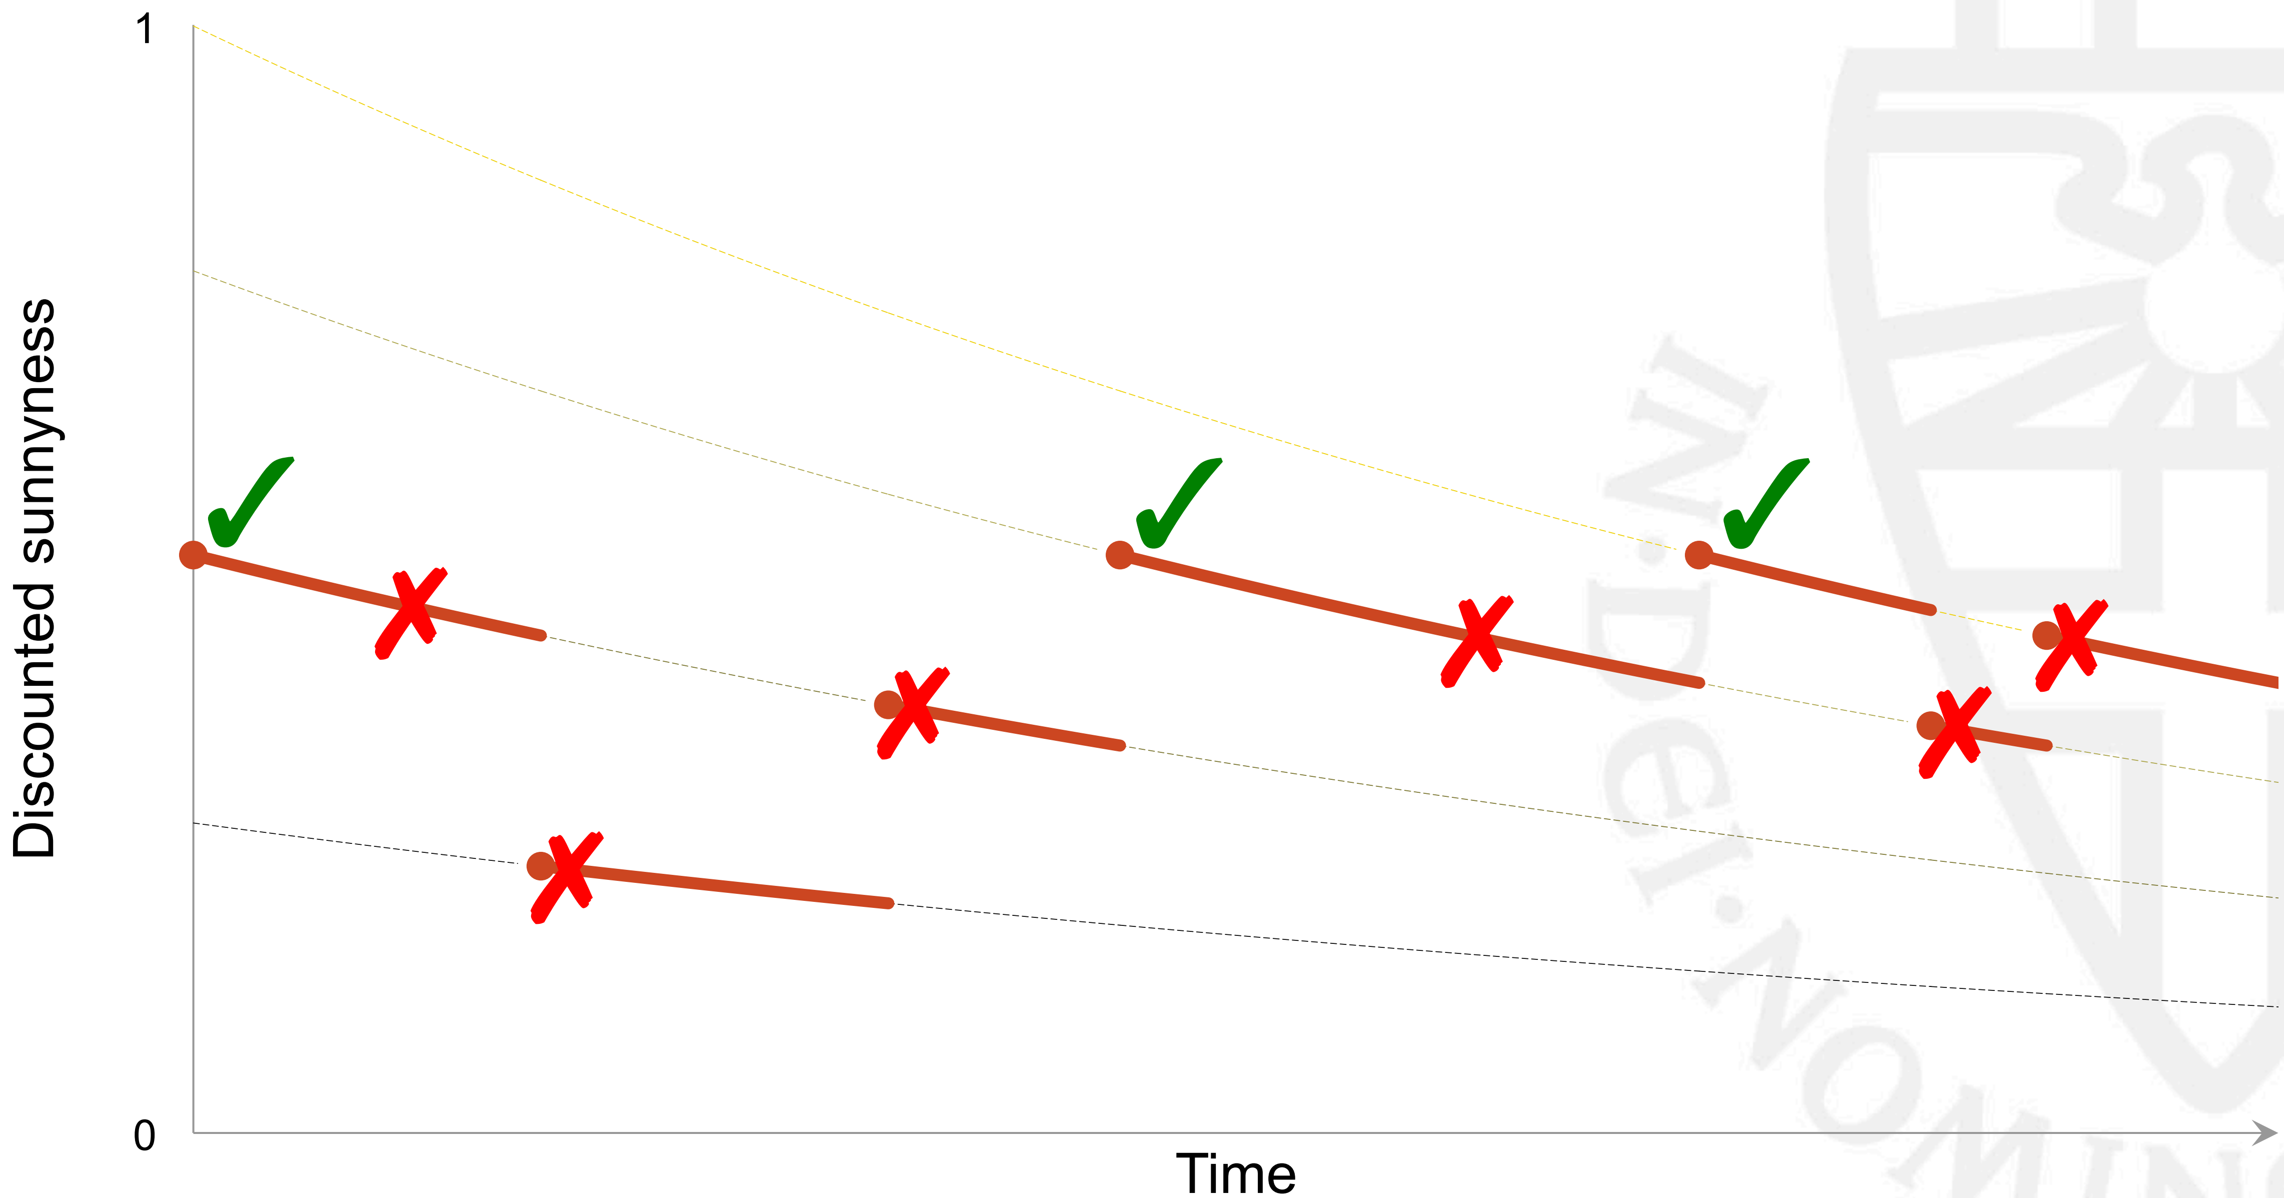

Iterative Solution

First iteration: assume all states are completely sunny

$[\![\forall\Diamond_\alpha \text{ sunny}]\!]^{\text{path}}$ is correct for sunny states

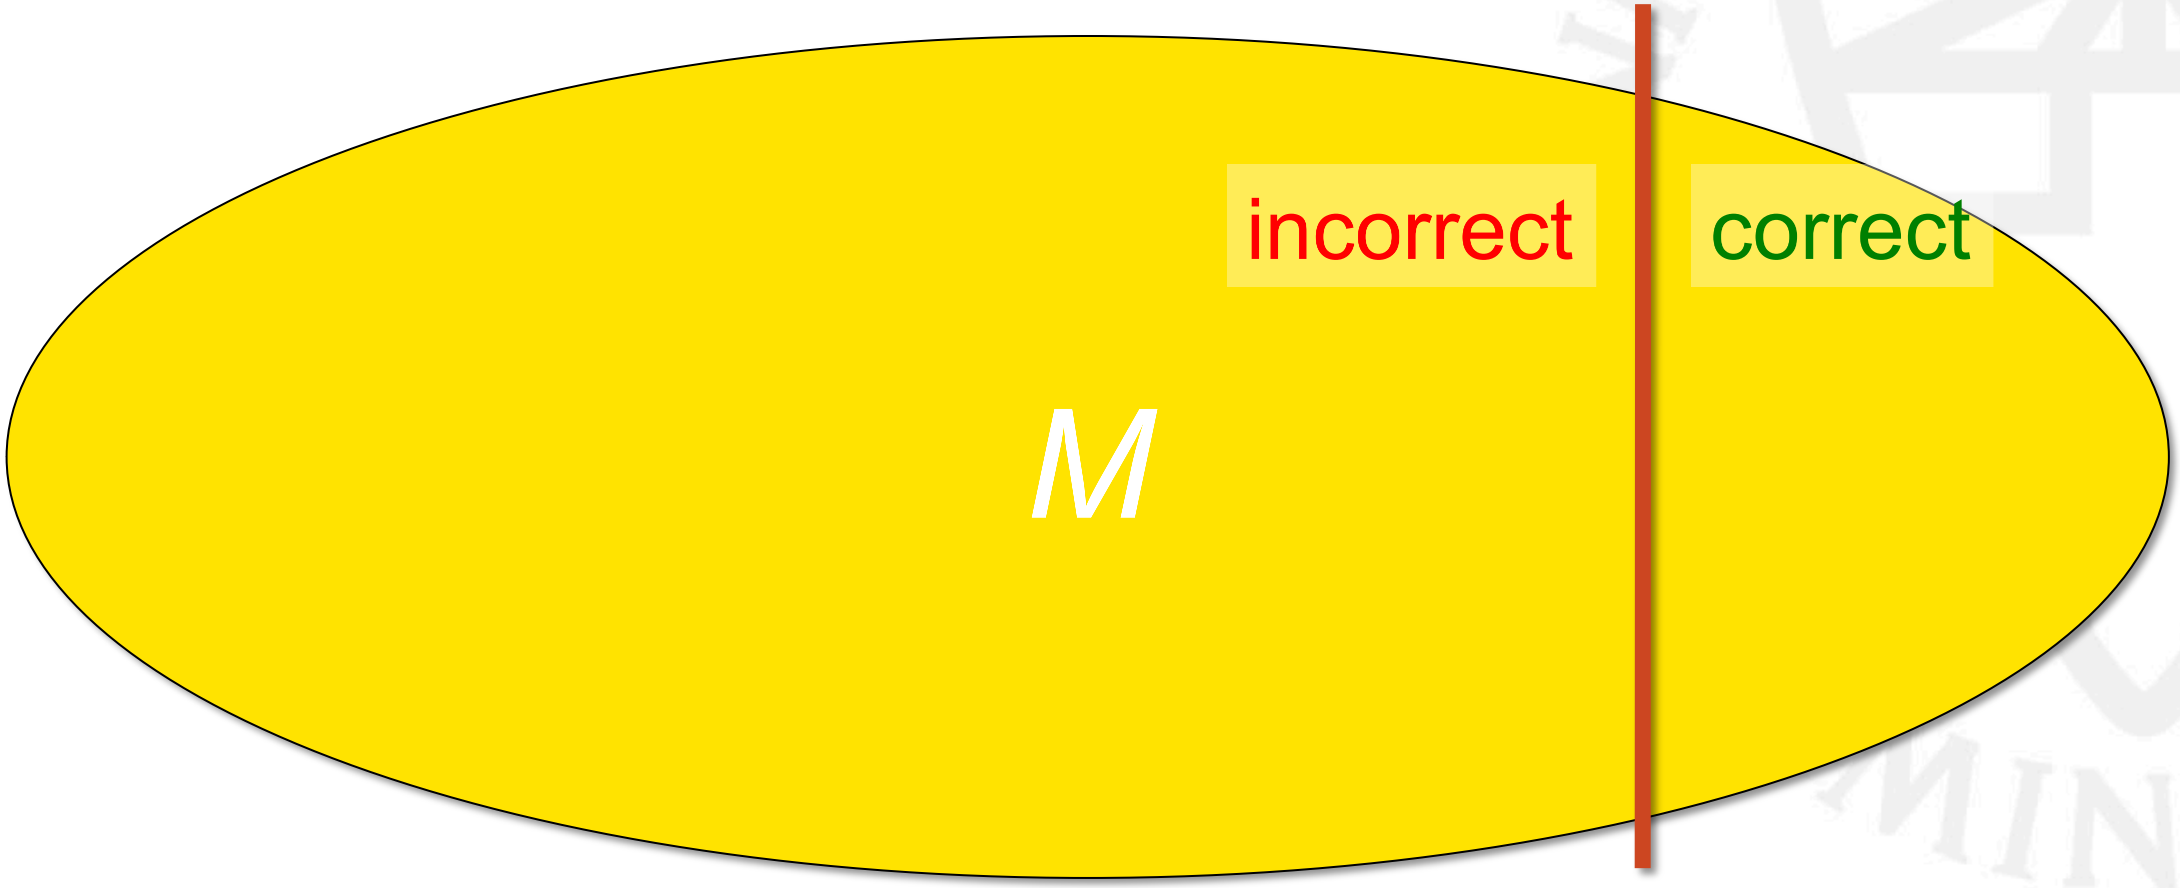

Iterative Solution

Second iteration: assume states are sunny or mostly sunny

$[\![\forall\Diamond_\alpha \text{ sunny}]\!]^{\text{path}}$ is correct for sunny and mostly sunny states

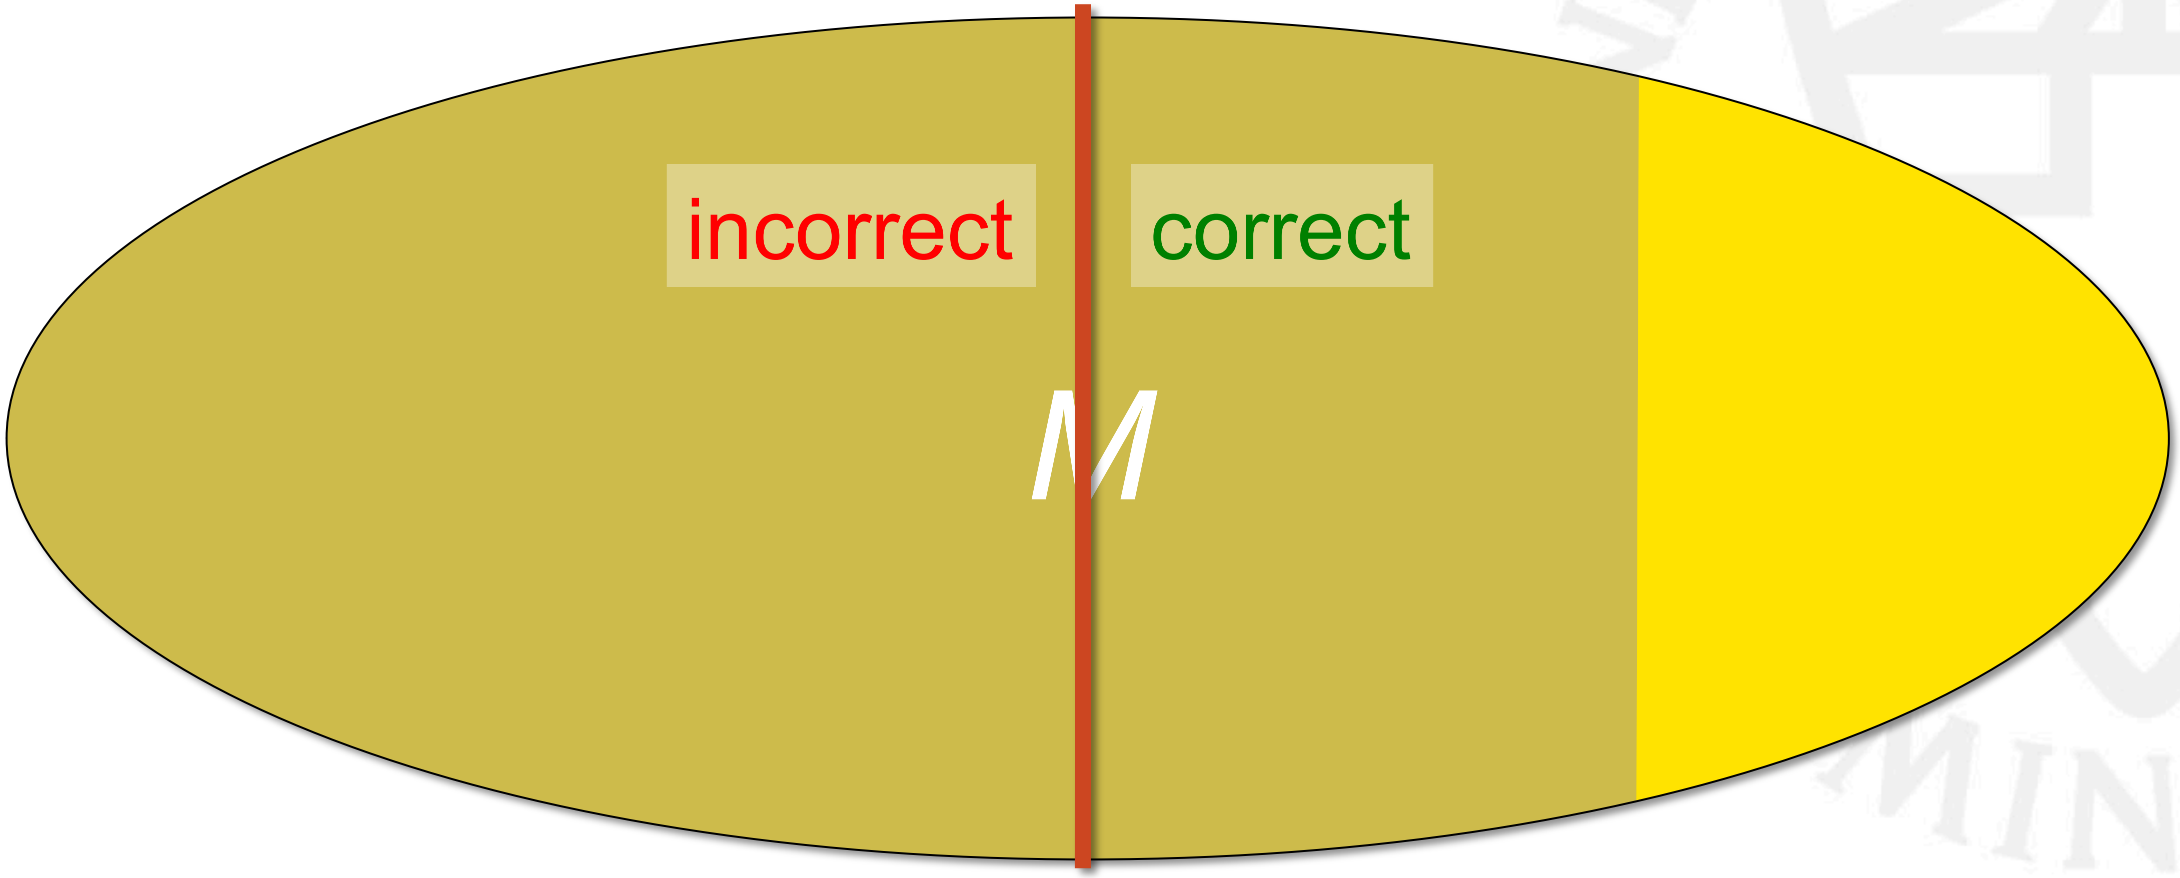

Iterative Solution

Third iteration: assume three shades of sunnyness exist

$[\![\forall\Diamond_\alpha \text{ sunny}]\!]^{path}$ is correct for three sunniest shades

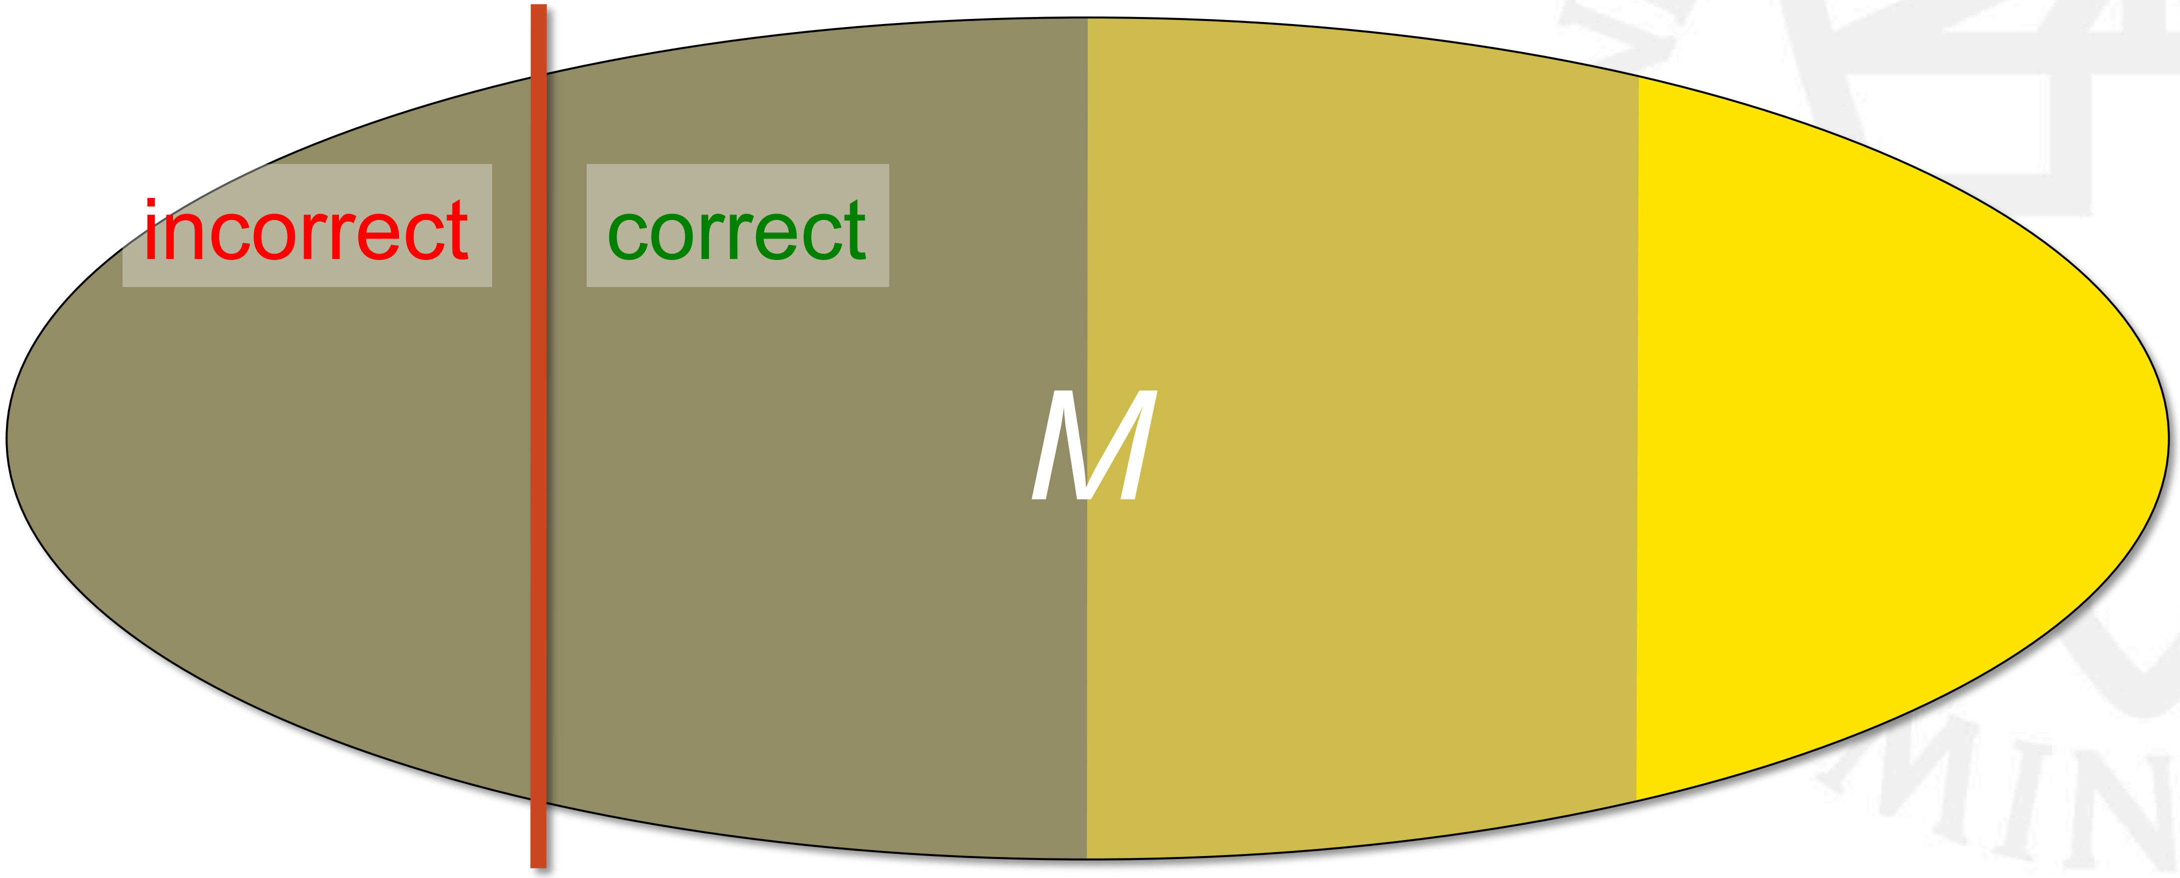

Iterative Solution

Repeat until all shades of sunnyness have passed

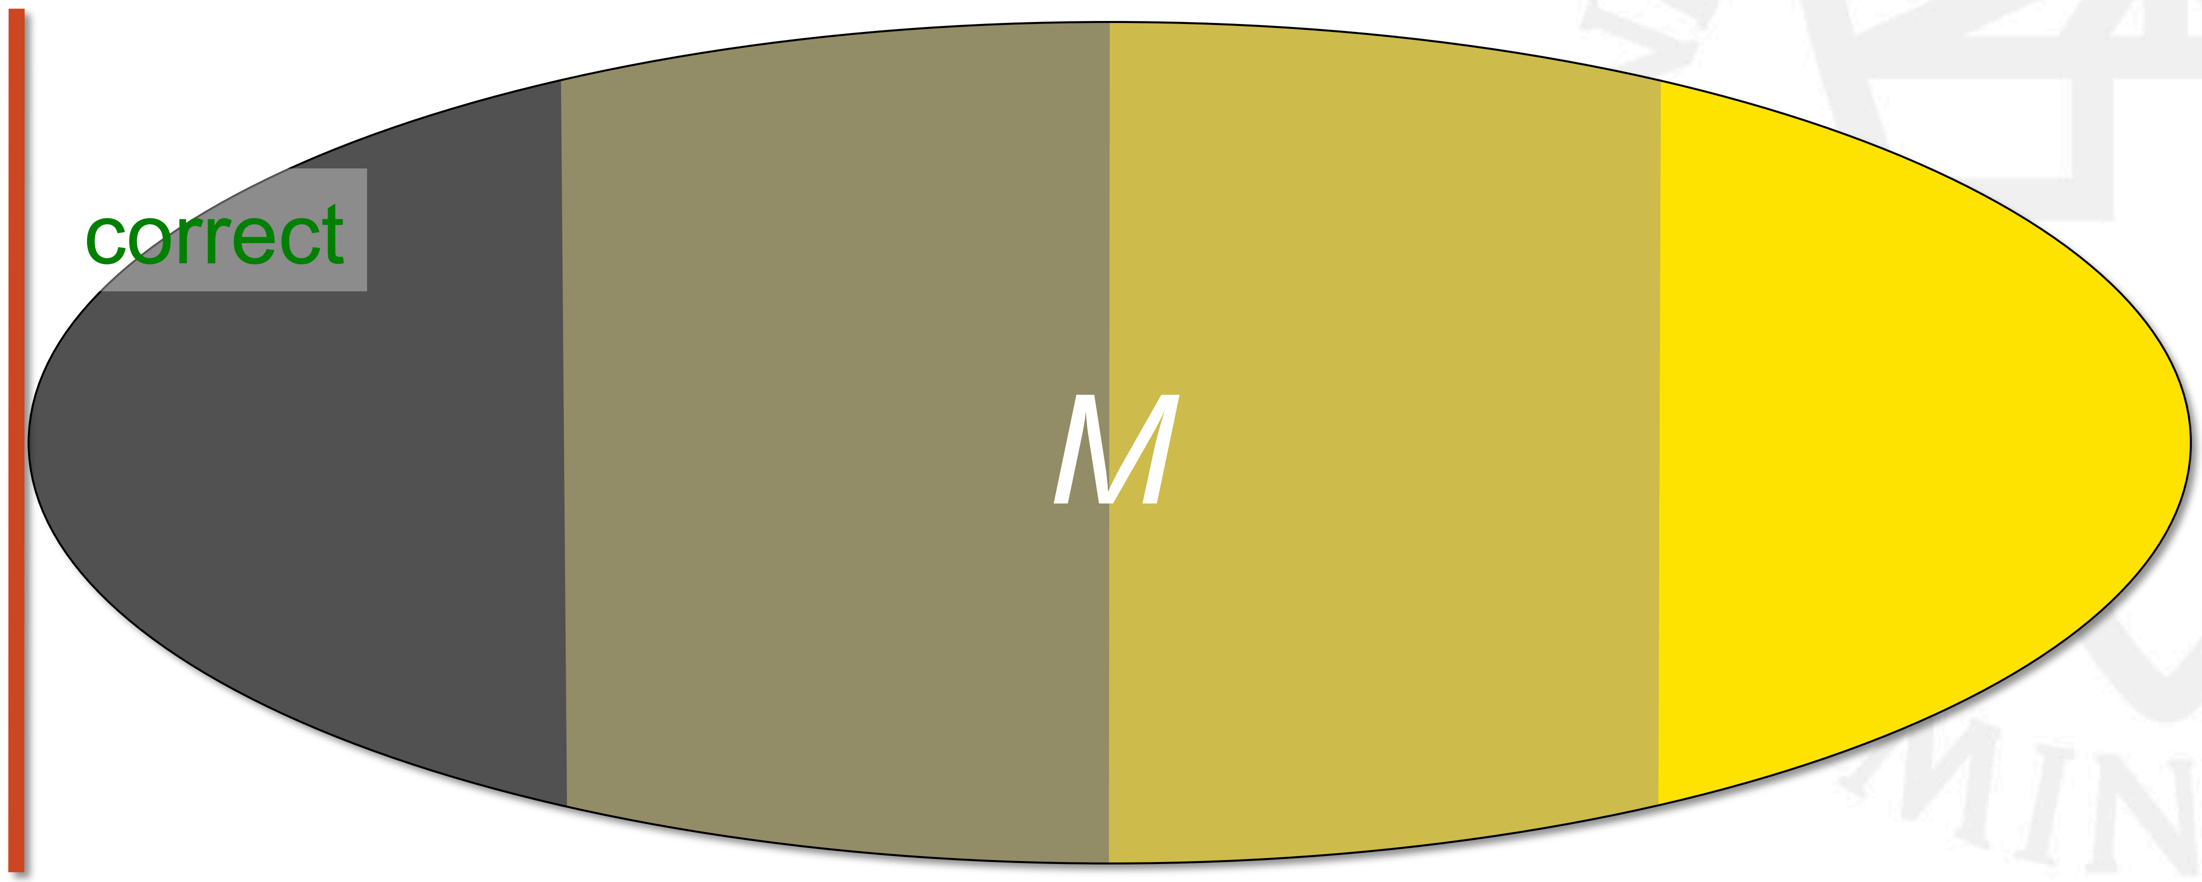

How To Take the Expectation Over Runs

many different types of runs
actually only very few cases to distinguish

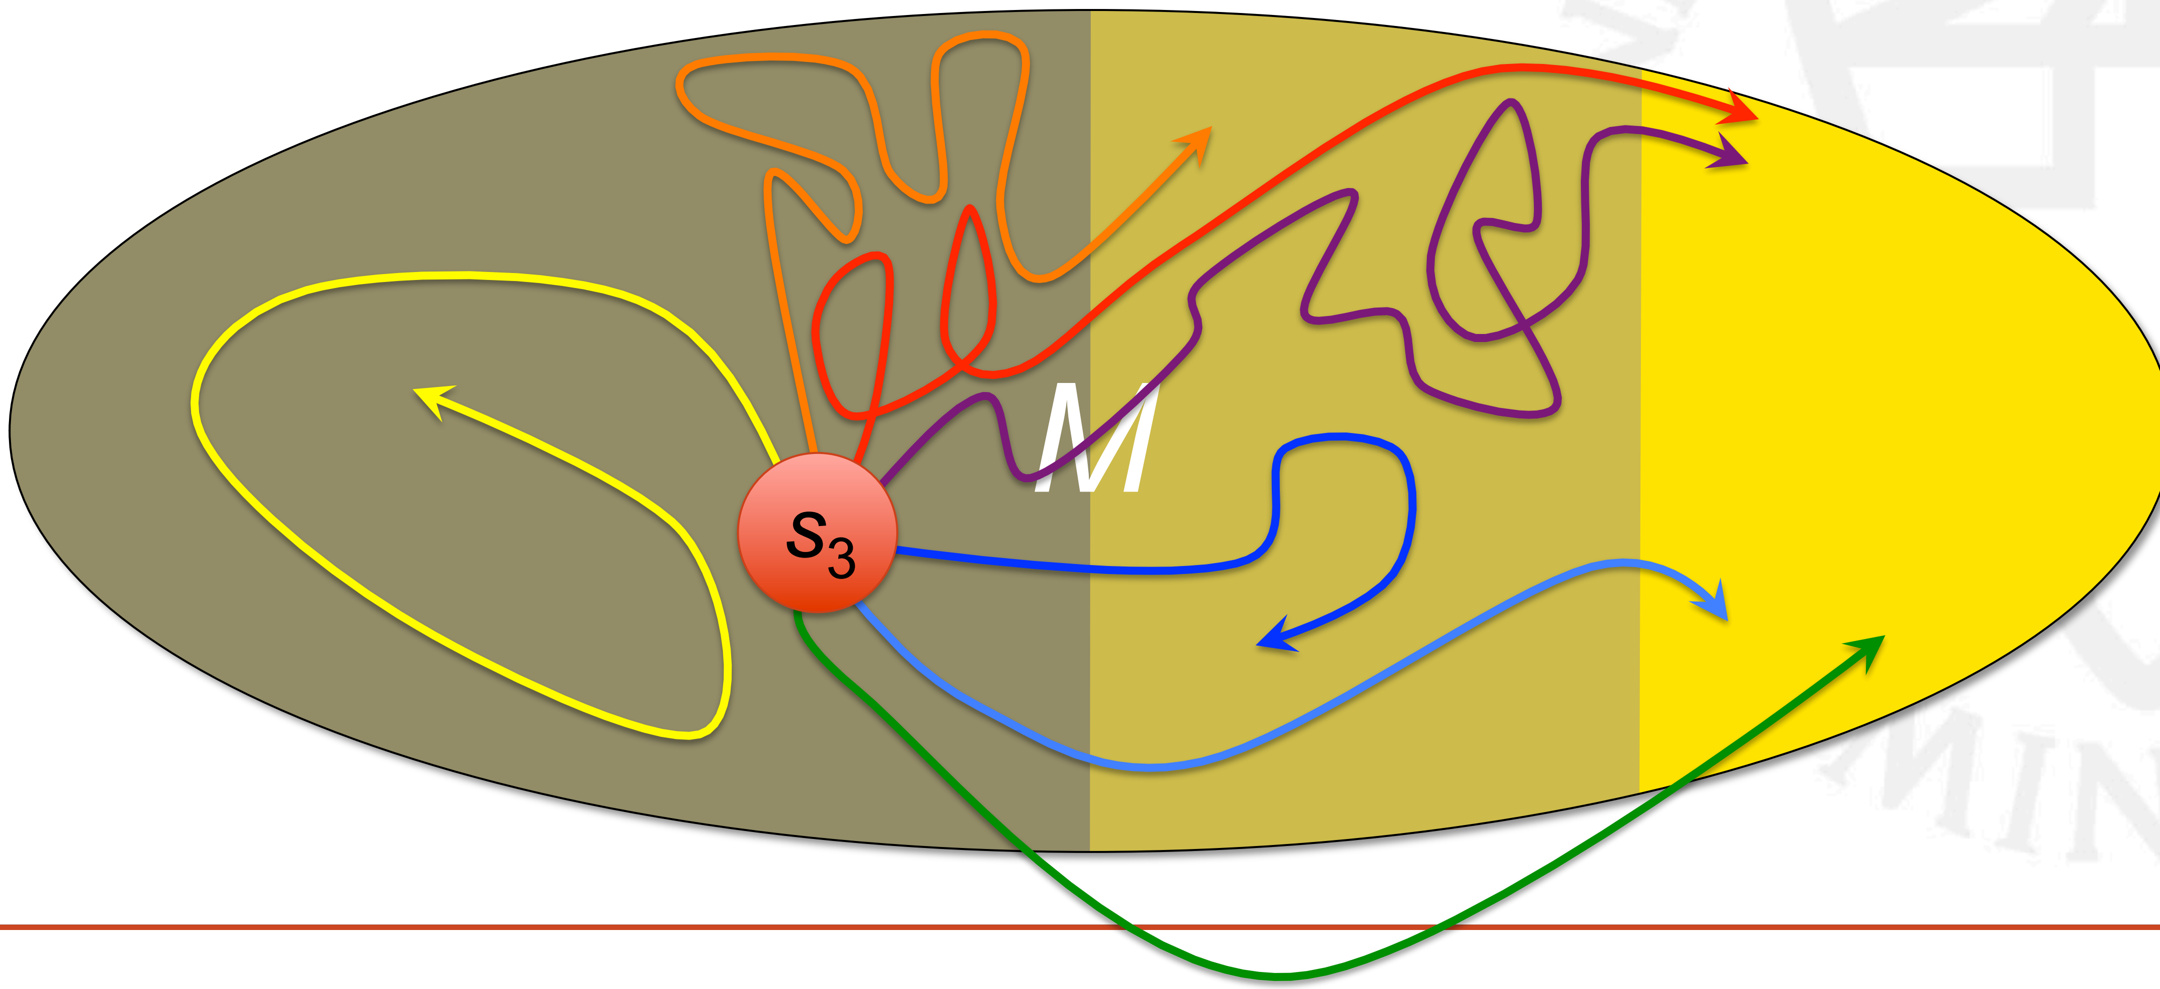

Paths That Reach a Better State Quickly

when path reaches better state,
reuse result of earlier iterations

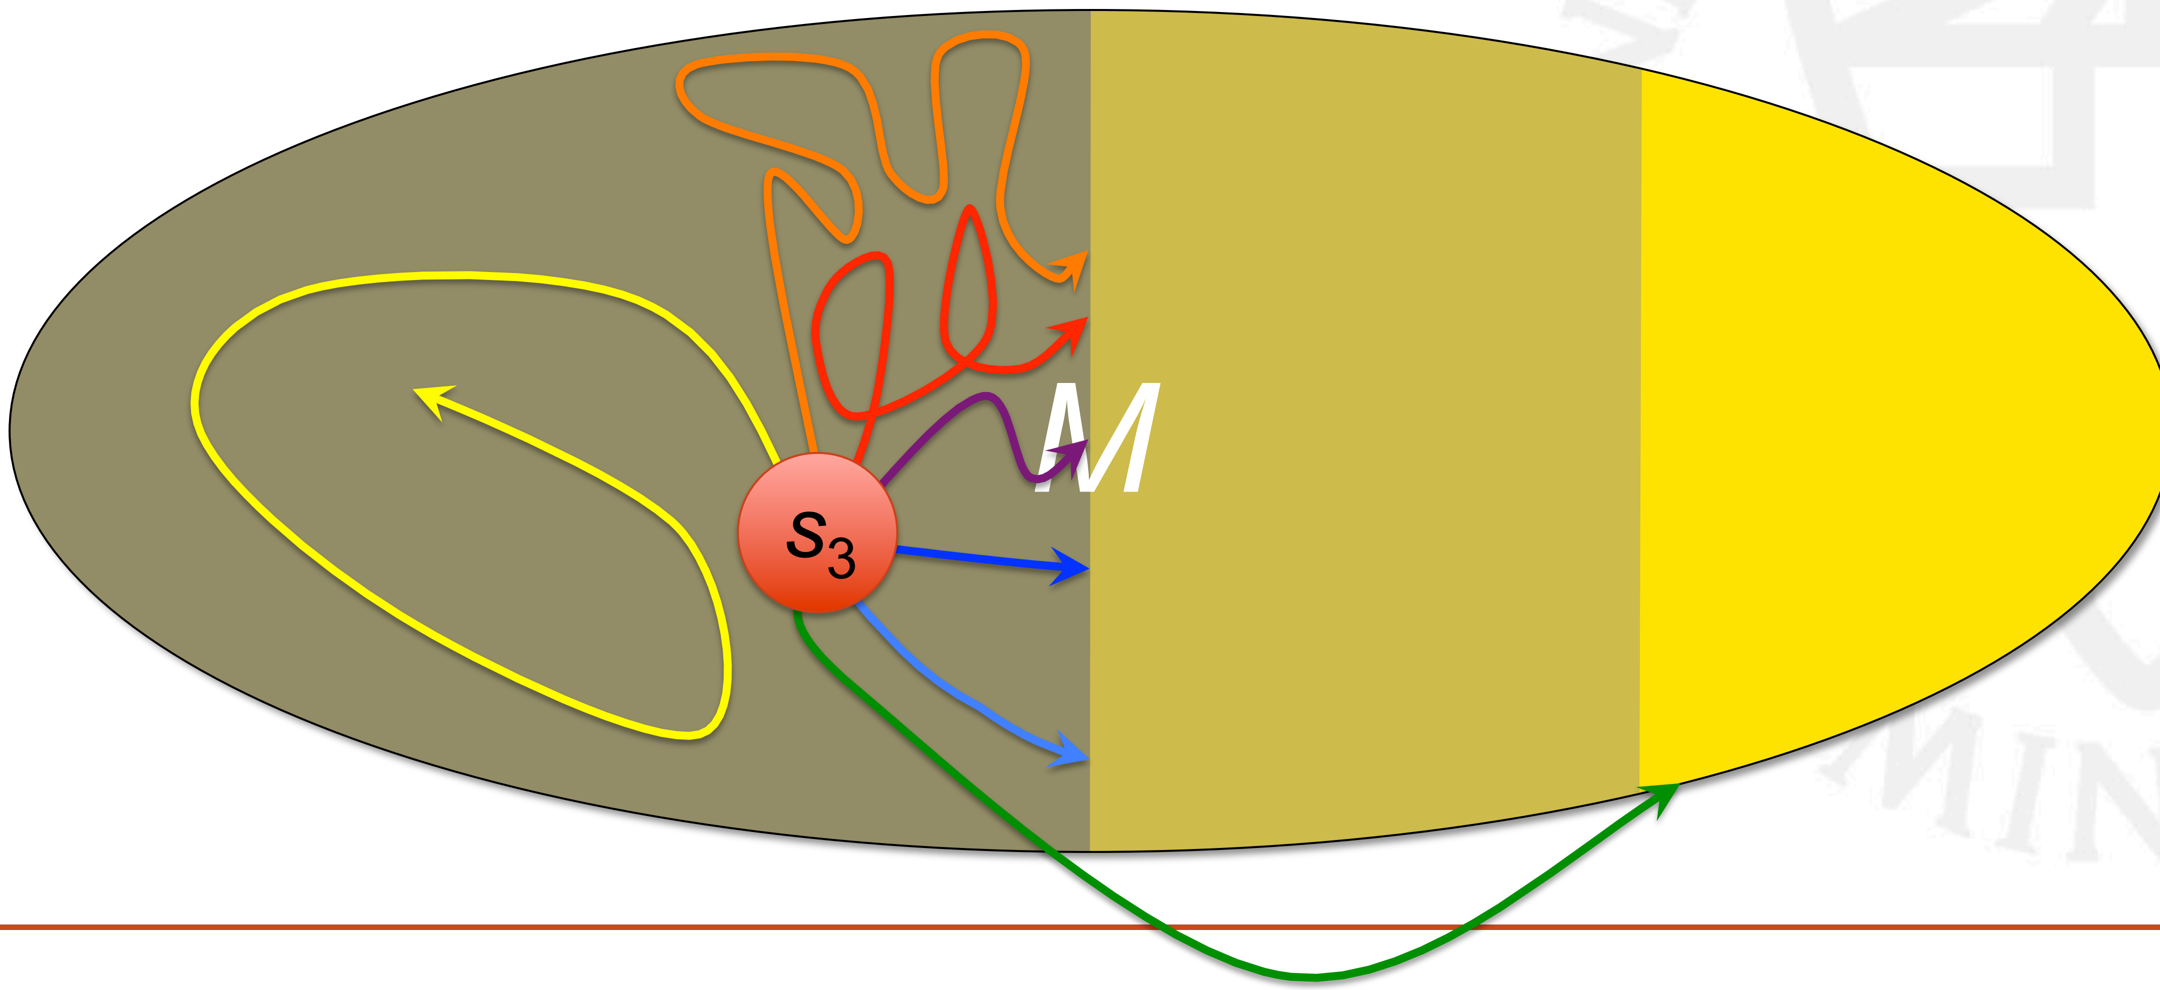

Cutoff Time

time within which a better state must be reached
otherwise, discounting compensates effect of improvement

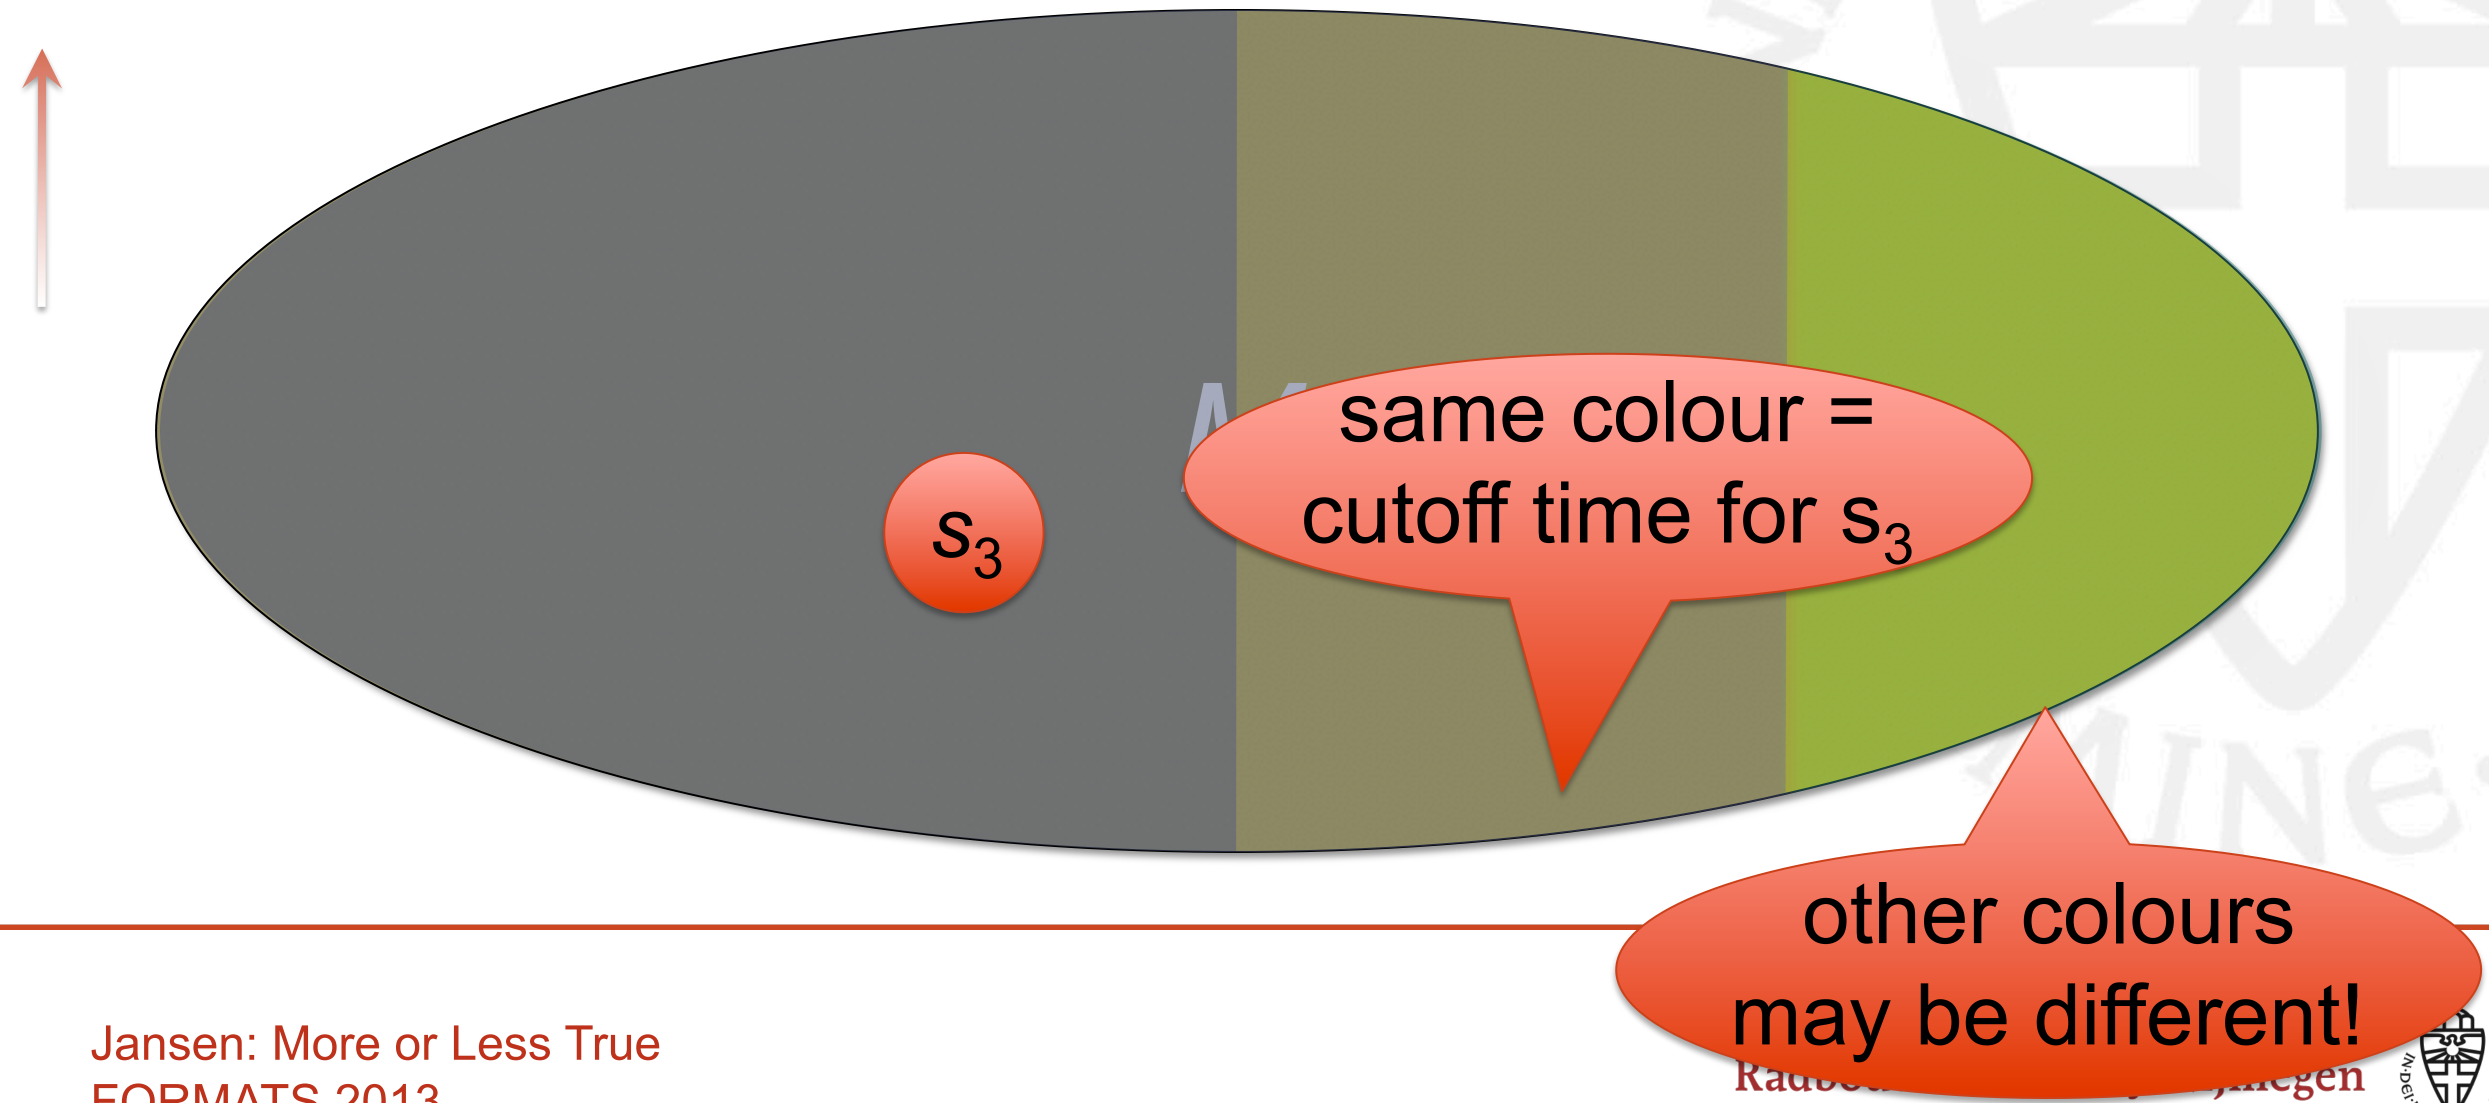

Paths that Stay in Bad States for a Long Time

at cutoff time, reuse result of earlier iterations

strictly speaking, that result was an overestimation,
but discounting until cutoff time compensates the error!

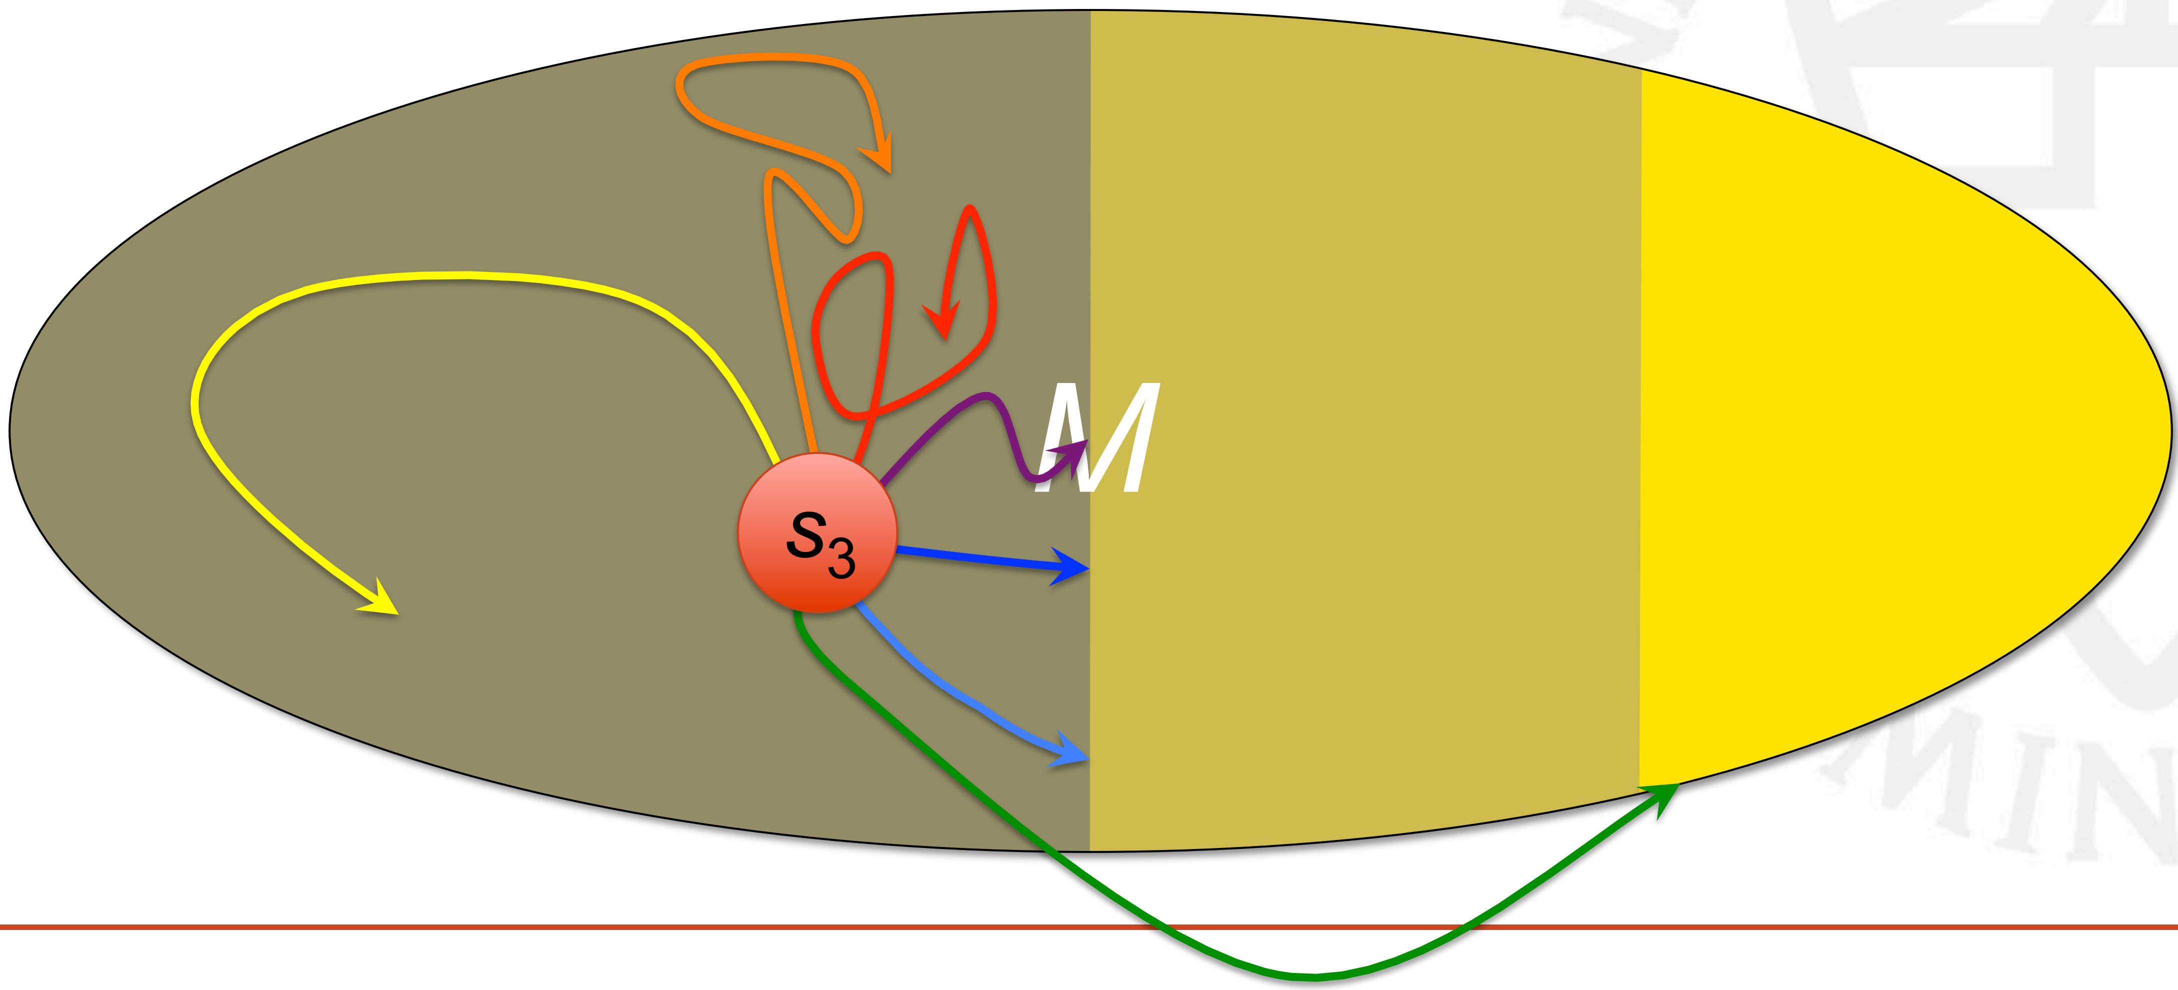

Time-bounded reach probability in CTMCs

- "How large is the probability to reach state s_2
  within time at most t_{cutoff}?"
  standard algorithms to answer this question exist

- calculating $[\![\forall\Diamond_\alpha \text{ sunny}]\!]^{\text{path}}$ reduces
  to (sequence of) time-bounded reach probability problems

Model Checking the Path Semantics

- Other operators also allow similar iteration

- Model checking algorithm for a single temporal operator:
  1 Order states according to ⟦φ⟧-ness
  2 Iterate from the most ⟦φ⟧-y to the least ⟦φ⟧-y state:
    0 In the first iteration, all states get the maximal ⟦φ⟧-ness assigned.
    1 Calculate cutoff time
    2 Calculate reach probability until cutoff time
    3 Take weighted sum over (discounted) values from earlier iteration

- Repeat this algorithm for nested formulas

Achieved results

- Extended: discounted CTL to continuous-time MCs
- Two semantics: fixpoint and path
- Model checking algorithms
  - Fixpoint: reduction to discrete-time DCTL
  - Path: reduction to time-bounded reach probability problems

A Challenge for the Bored... Resolved

QEST: sound and ~~(hopefully) complete~~ weak simulation for substochastic DTMCs

Baier, Engelen, Majster-Cederbaum 2000

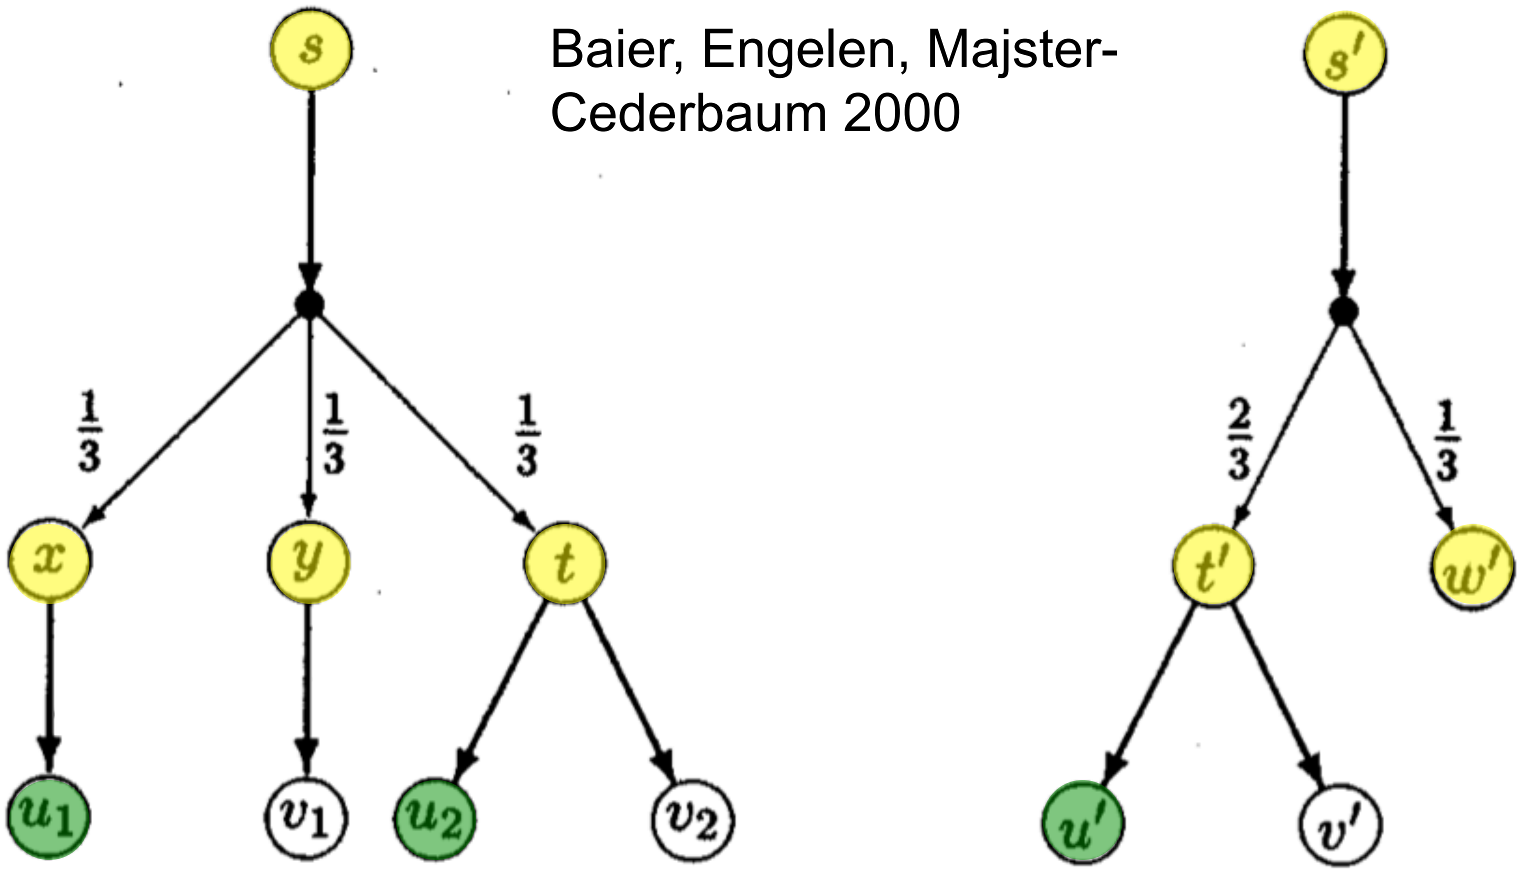

FIG. 4. $s \sqsubseteq_{\uparrow} s'$ and $s \not\sqsubseteq s'$.

A Challenge for the Bored... Resolved

QEST: sound and ~~(hopefully) complete~~ weak simulation for substochastic DTMCs

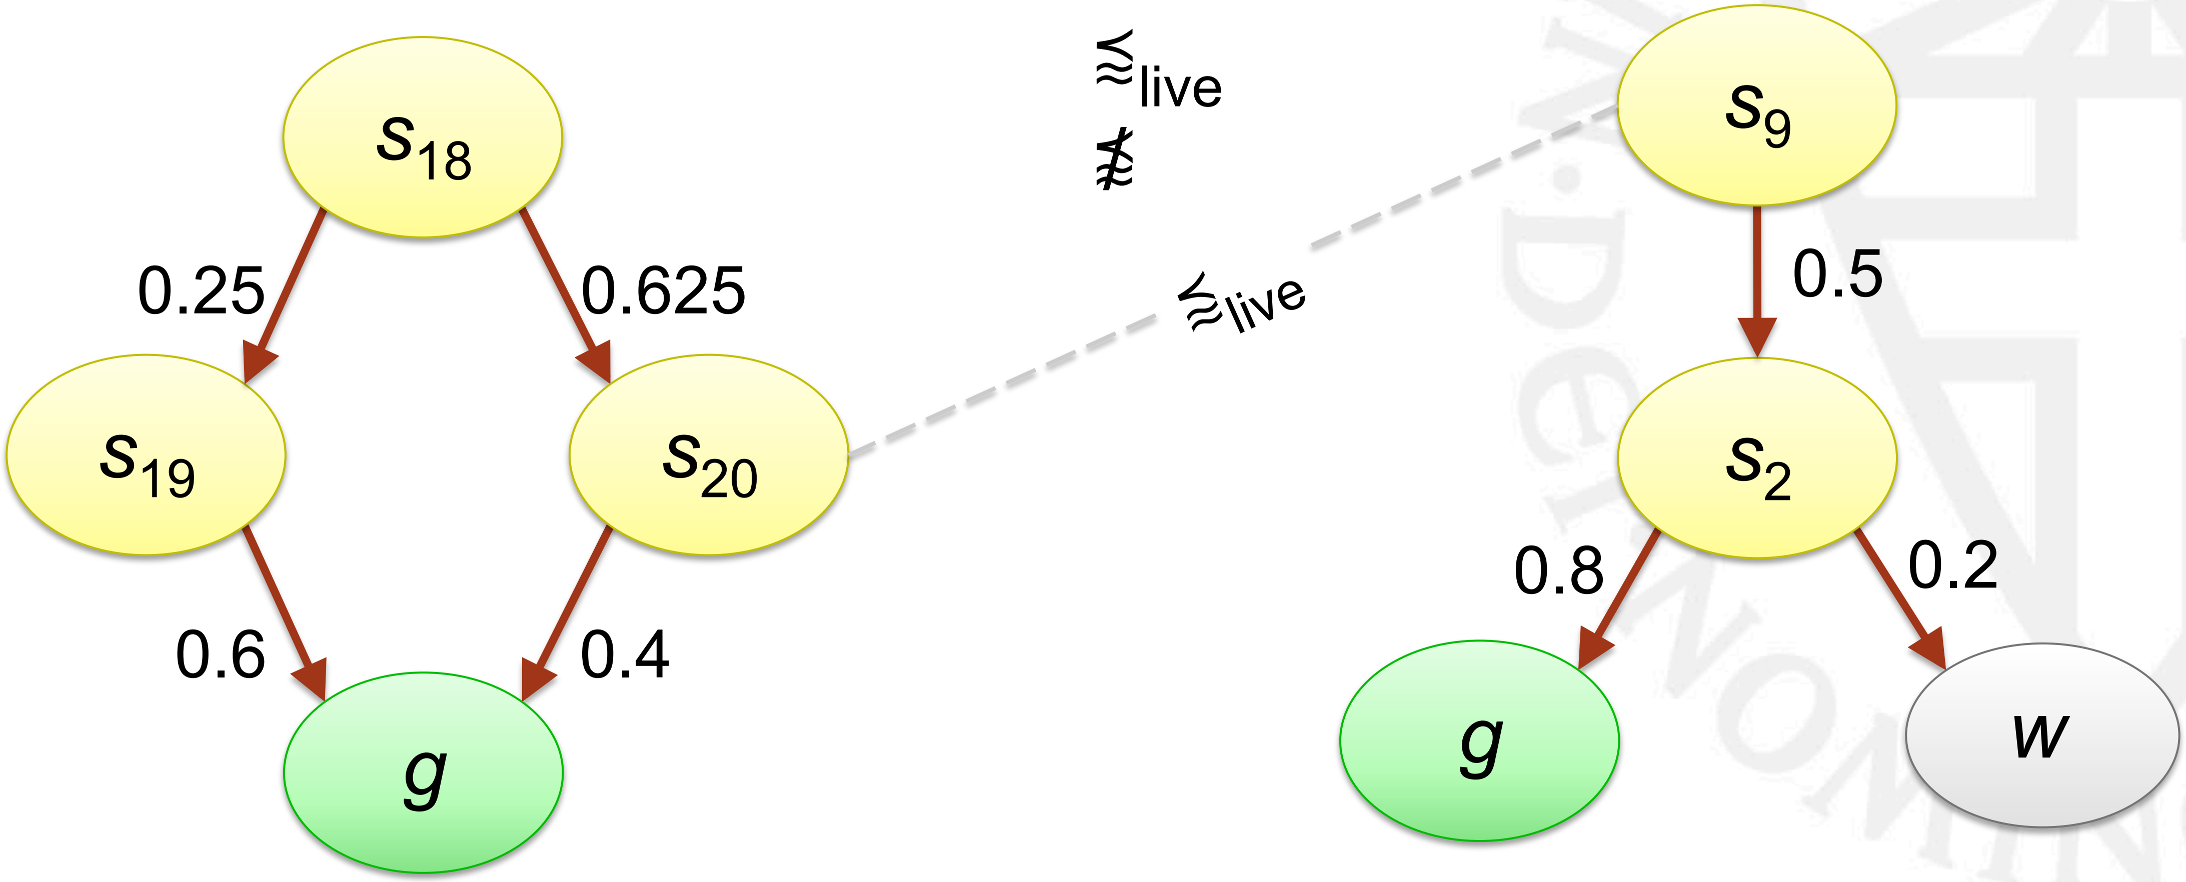

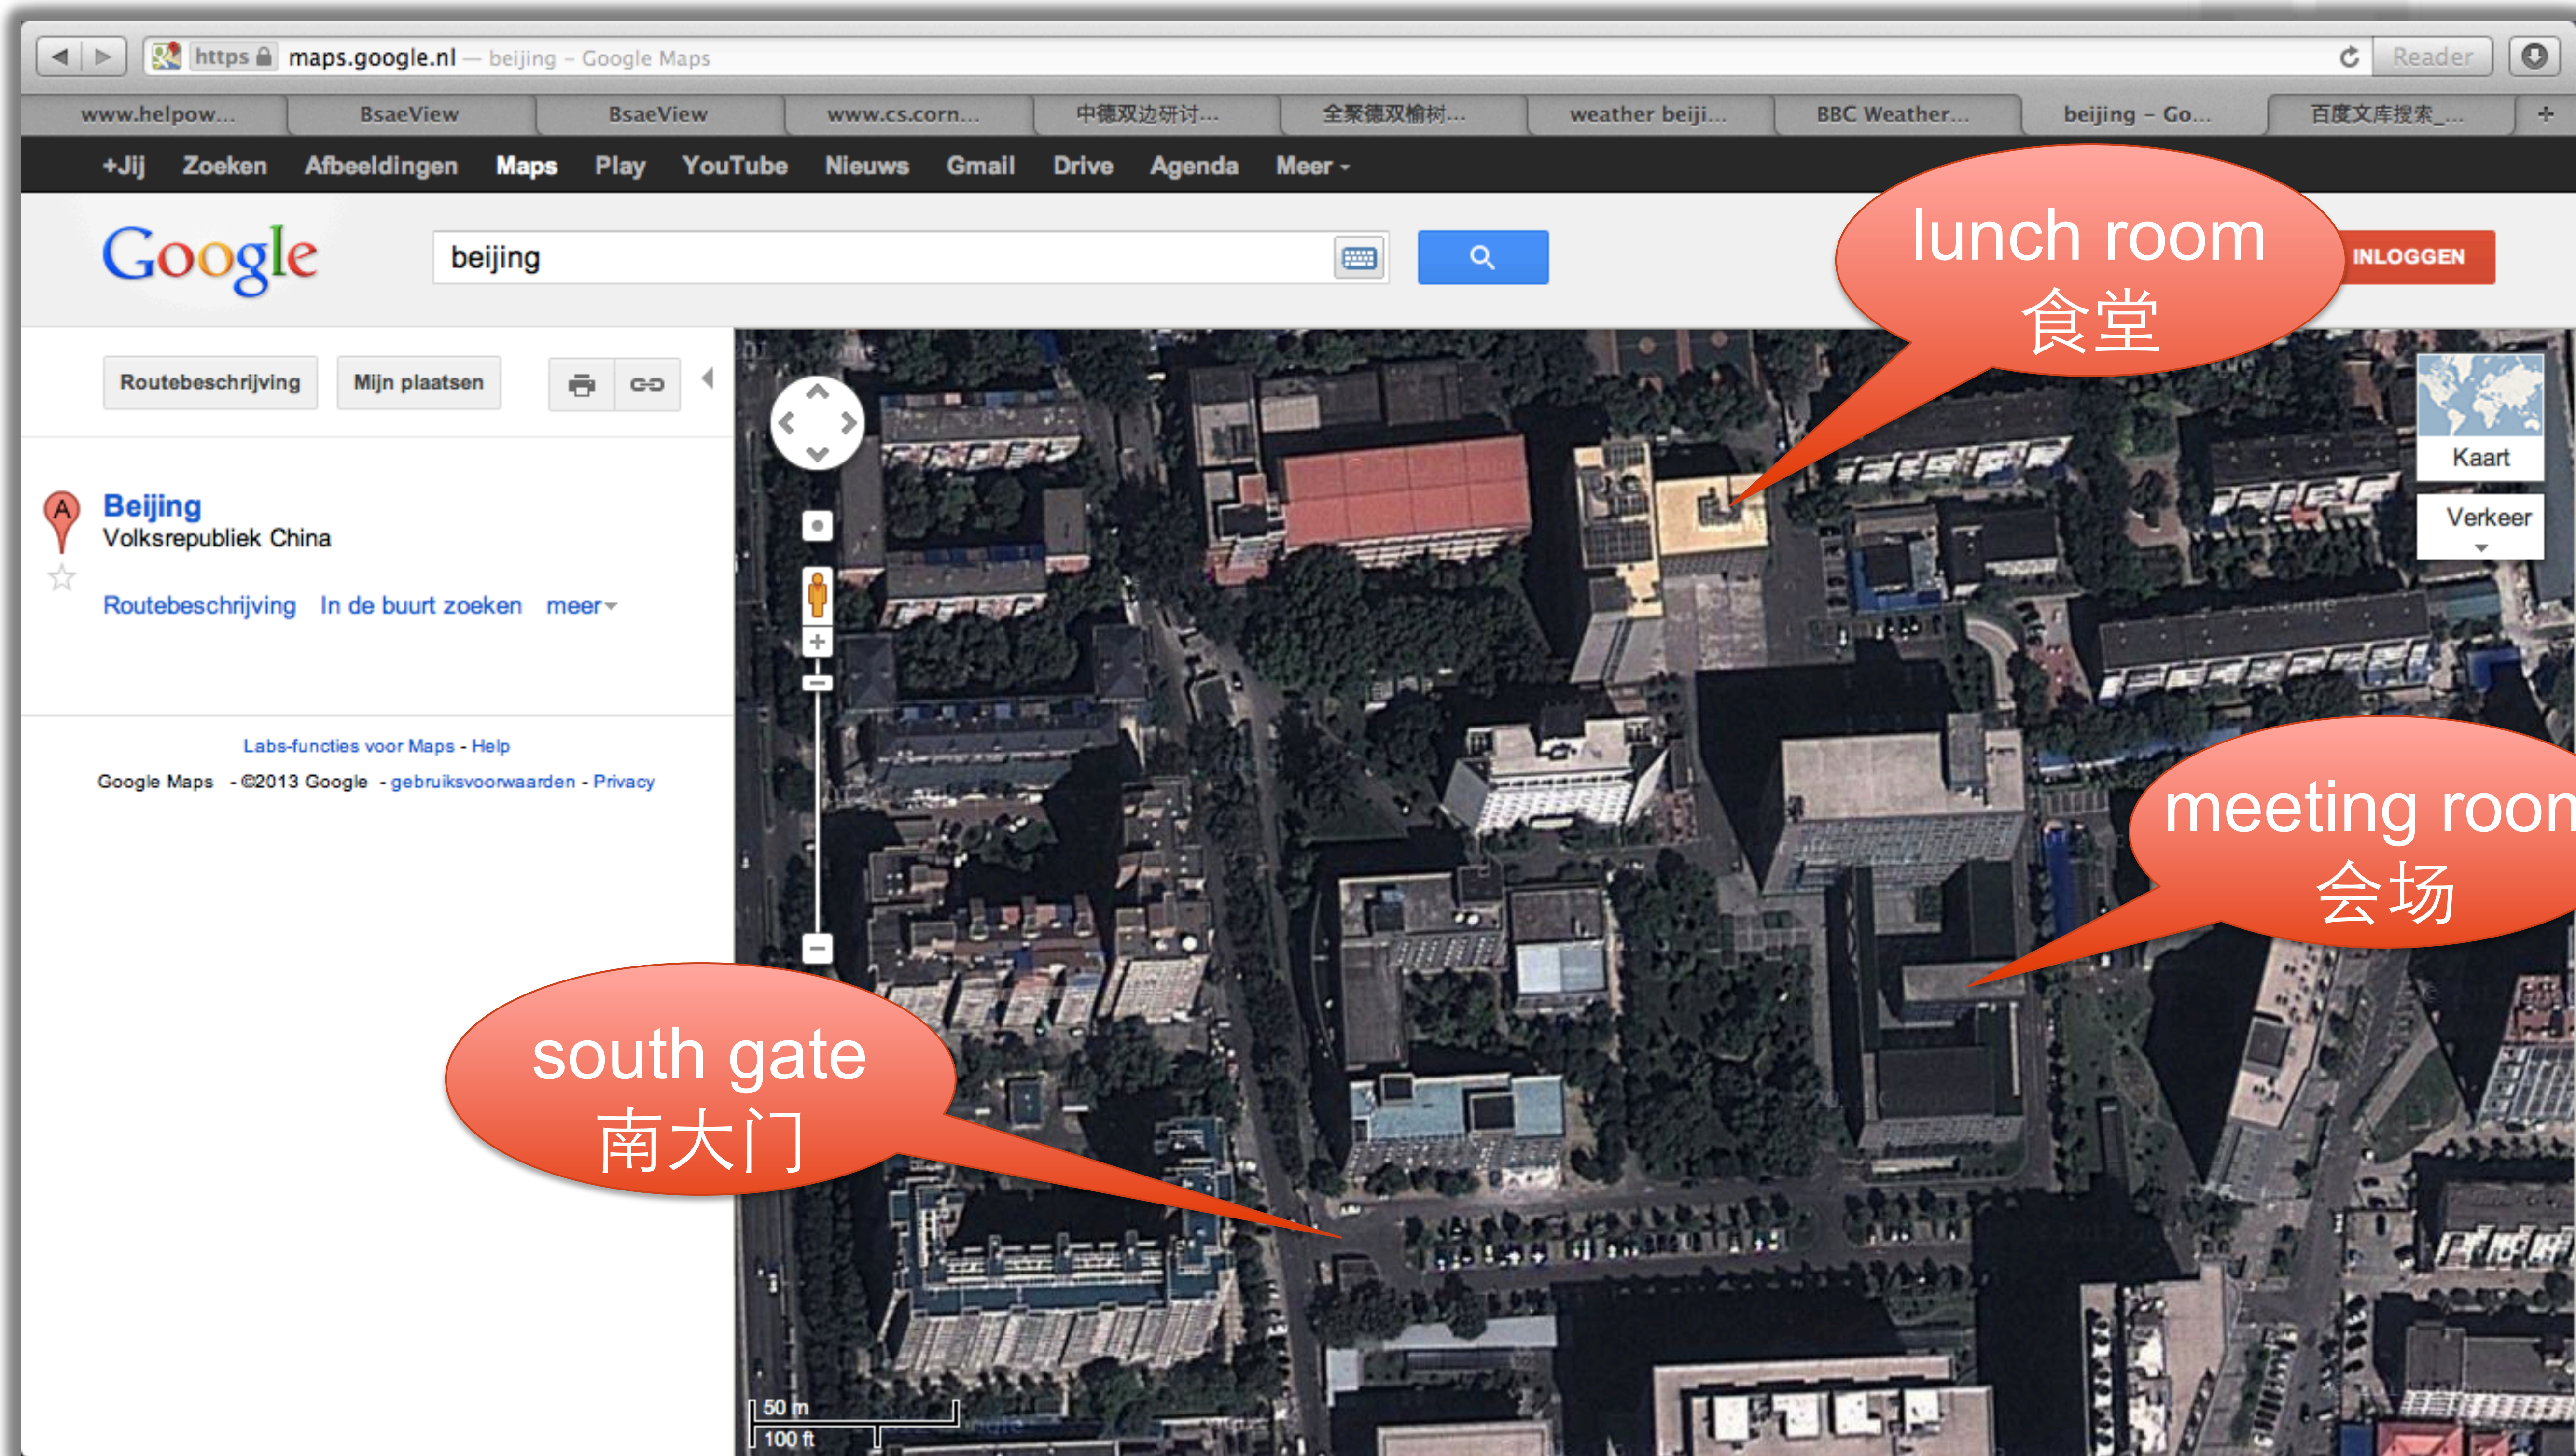
maps.google.nl — beijing – Google Maps
+Jij Zoeken Afbeeldingen Maps Play YouTube Nieuws Gmail Drive Agenda Meer
beijing
INLOGGEN
Routebeschrijving
Mijn plaatsen
Beijing
Volksrepubliek China
Routebeschrijving In de buurt zoeken meer
Labs-functies voor Maps - Help
Google Maps - ©2013 Google - gebruiksvoorwaarden - Privacy
Kaart
Verkeer
lunch room
食堂
meeting room
会场
south gate
南大门
50 m
100 ft