Approximation of continuous LMPs[1]

Alexandre Bouchard-Côté
Supervisors: Prakash Panangaden, Doina Precup
Reasoning and Learning Lab, McGill University
Sponsored by: NSERC, McGill School of Computer Science.

[1] Labelled Markov Processes

Motivation: example

Continuous system:

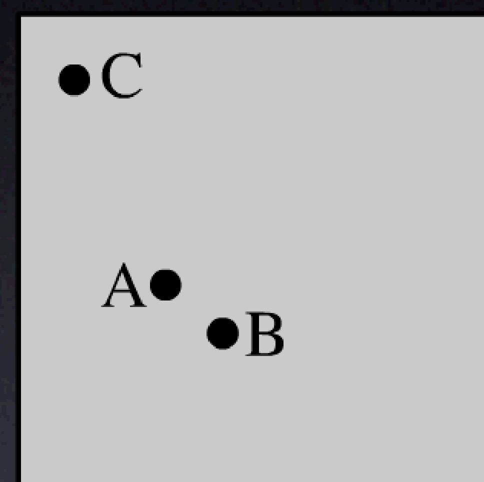

state space

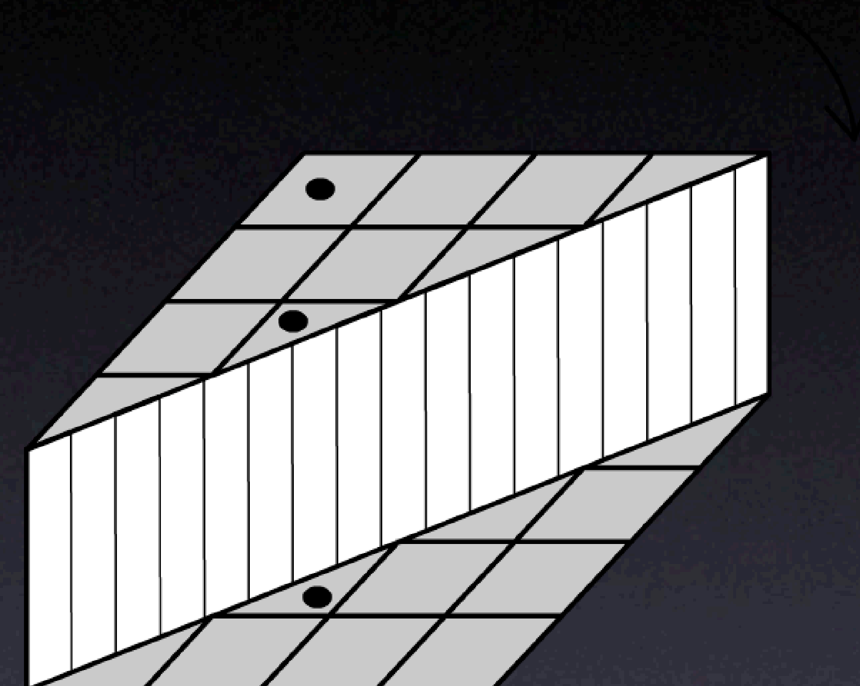

dynamics

Possible finite state approx.:

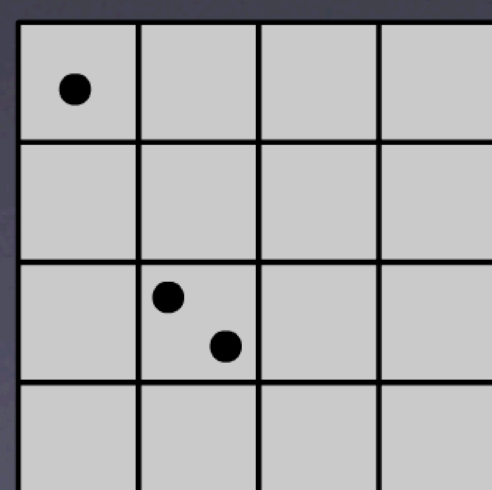

“geometric”

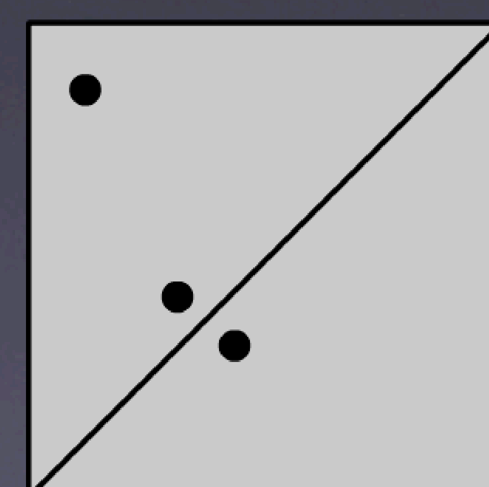

better

The approx. scheme[2]

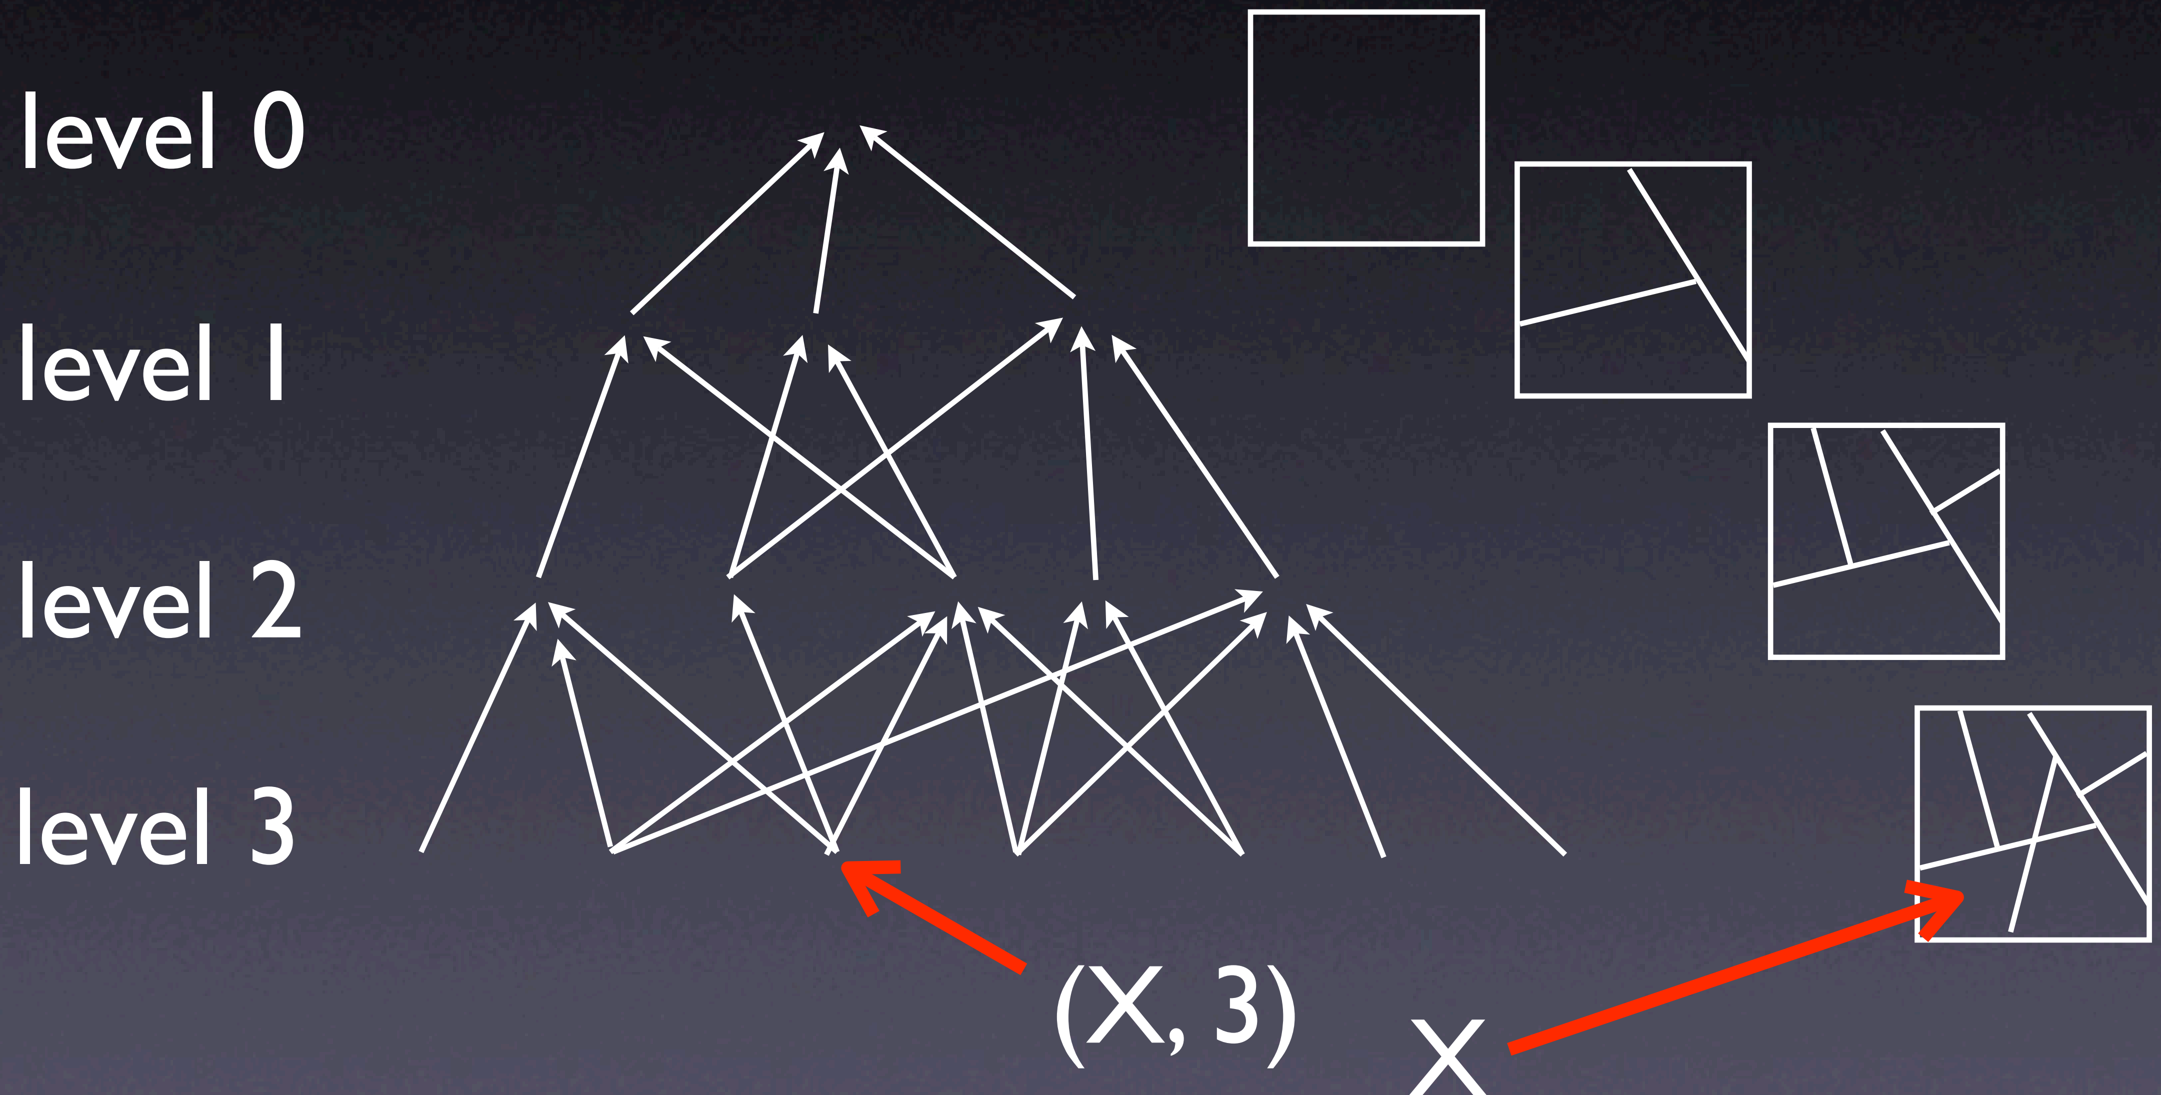

[2] J. Desharnais, V. Gupta, R. Jagadeesan, P. Panangaden. (2002). Approximating Labelled Markov Processes.

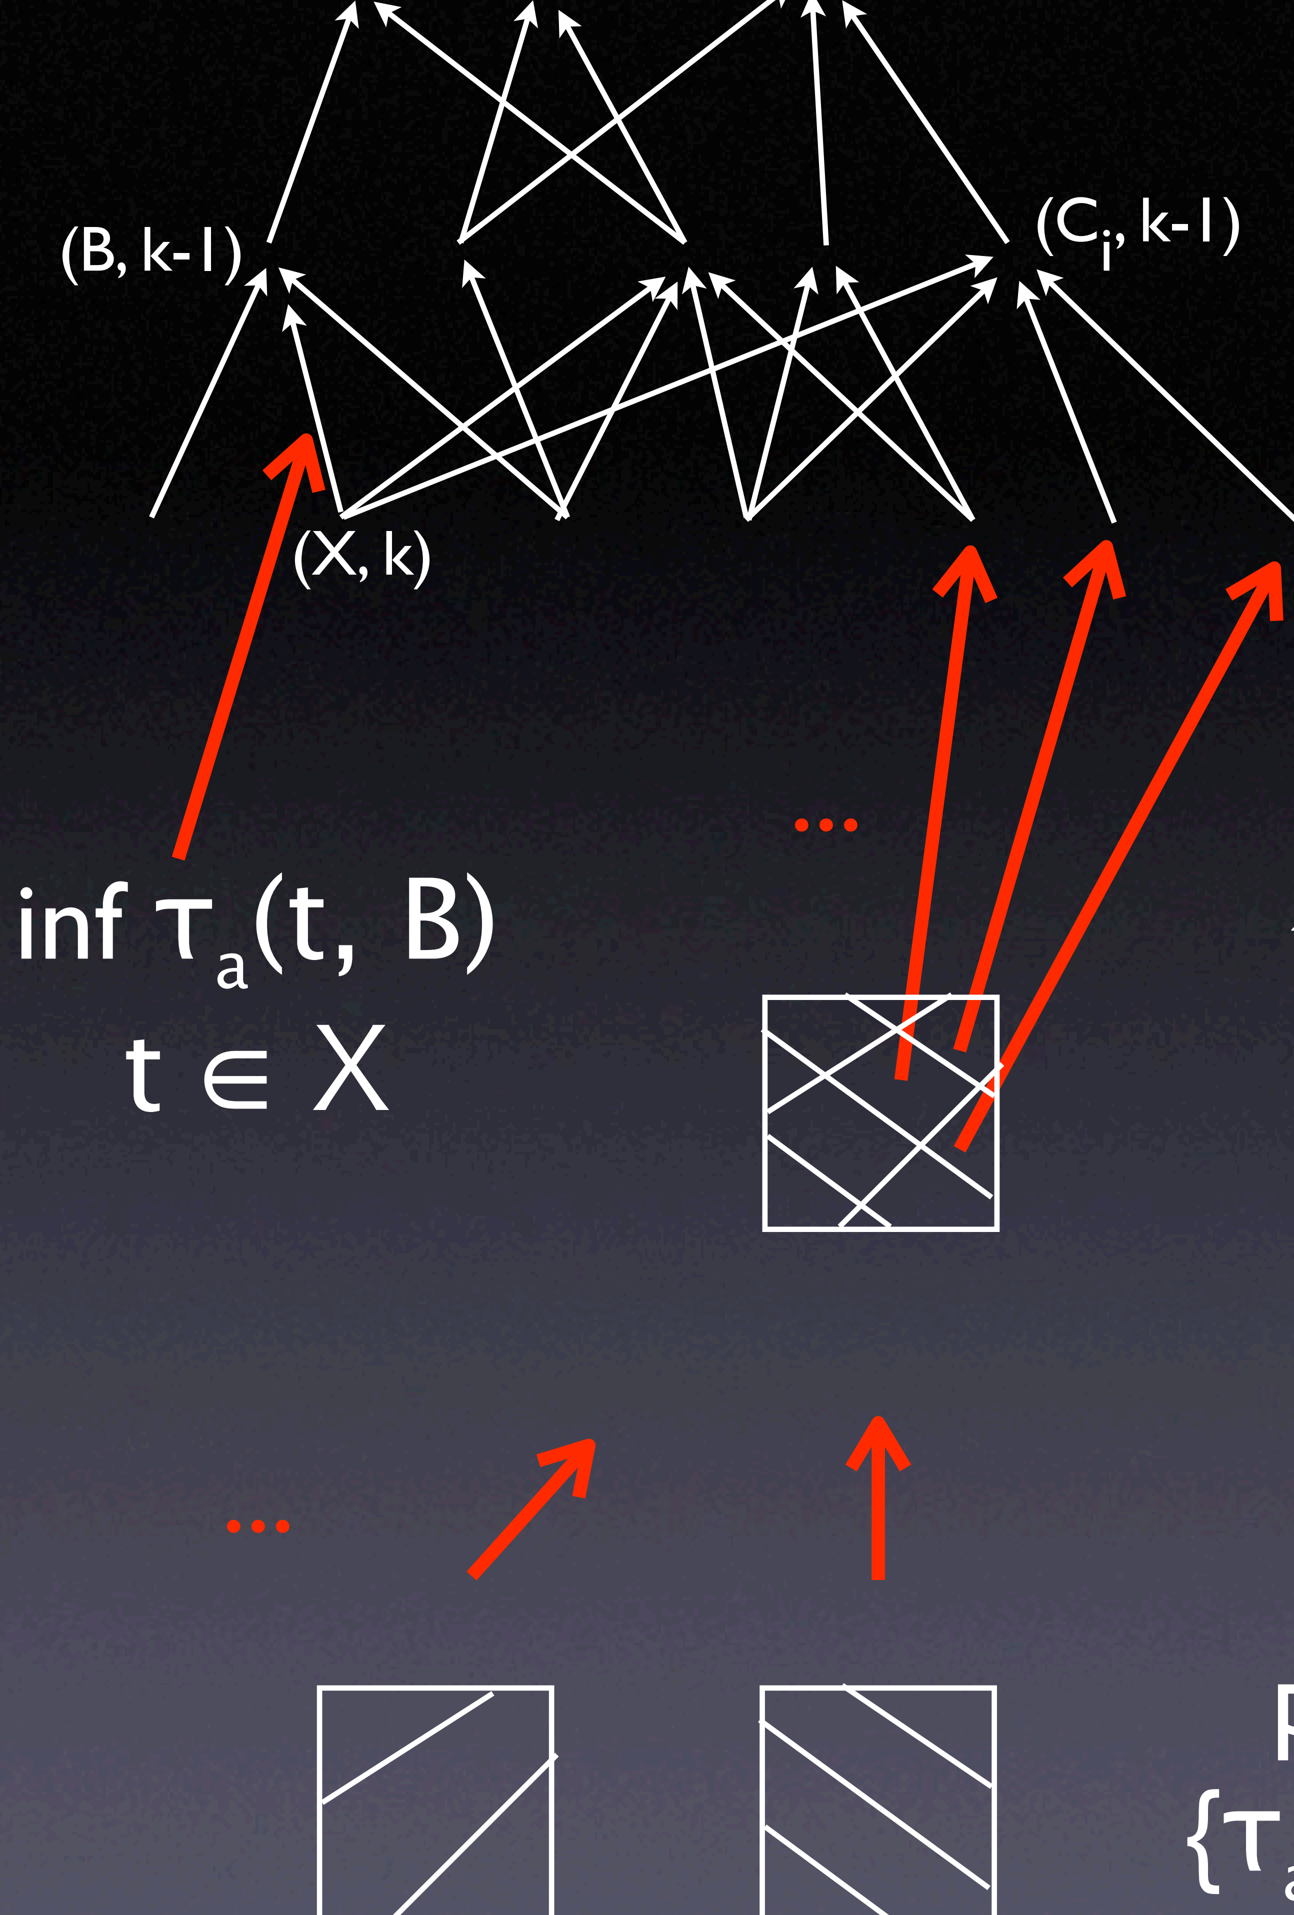

m states in the preceding level

partition of the range of the probability kernels into h-intervals of length ϵ/m

partition of the states

$\{\tau_a^{-1}(\bullet, C_i)(A_j) : A_j \in P\}$

Implementation difficulties

Infimum of measurable functions

$$\inf_{t \in X} \tau_a(t, B)$$

Generate partition (check if a set is empty)

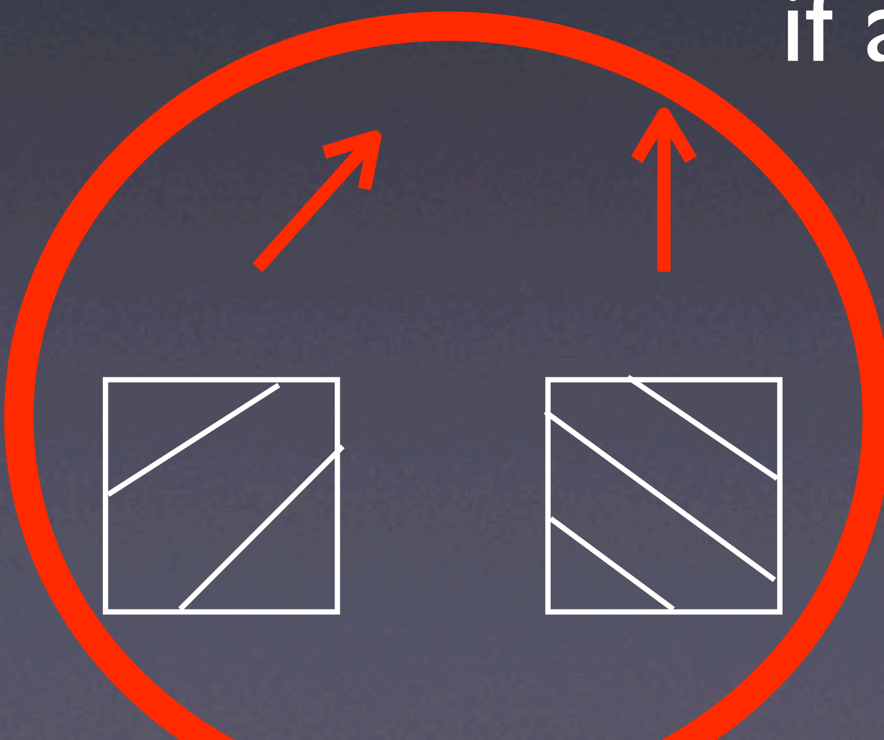

Invert a measurable function

$$\{\tau_a^{-1}(\bullet, C_i)(A_j)\}$$

How to “invert” the kernels

Representation of $\tau_a^{-1}(\bullet, C)((a,b])$

Instance’s variables:

f_a	C	(a,b]

Operations:

Check if $s_0 \in \tau_a^{-1}(\bullet, C)((a,b])$

Output true iff $\int_C f_a(s_0, x)\, d\mu(x) \in (a,b]$

Infimum

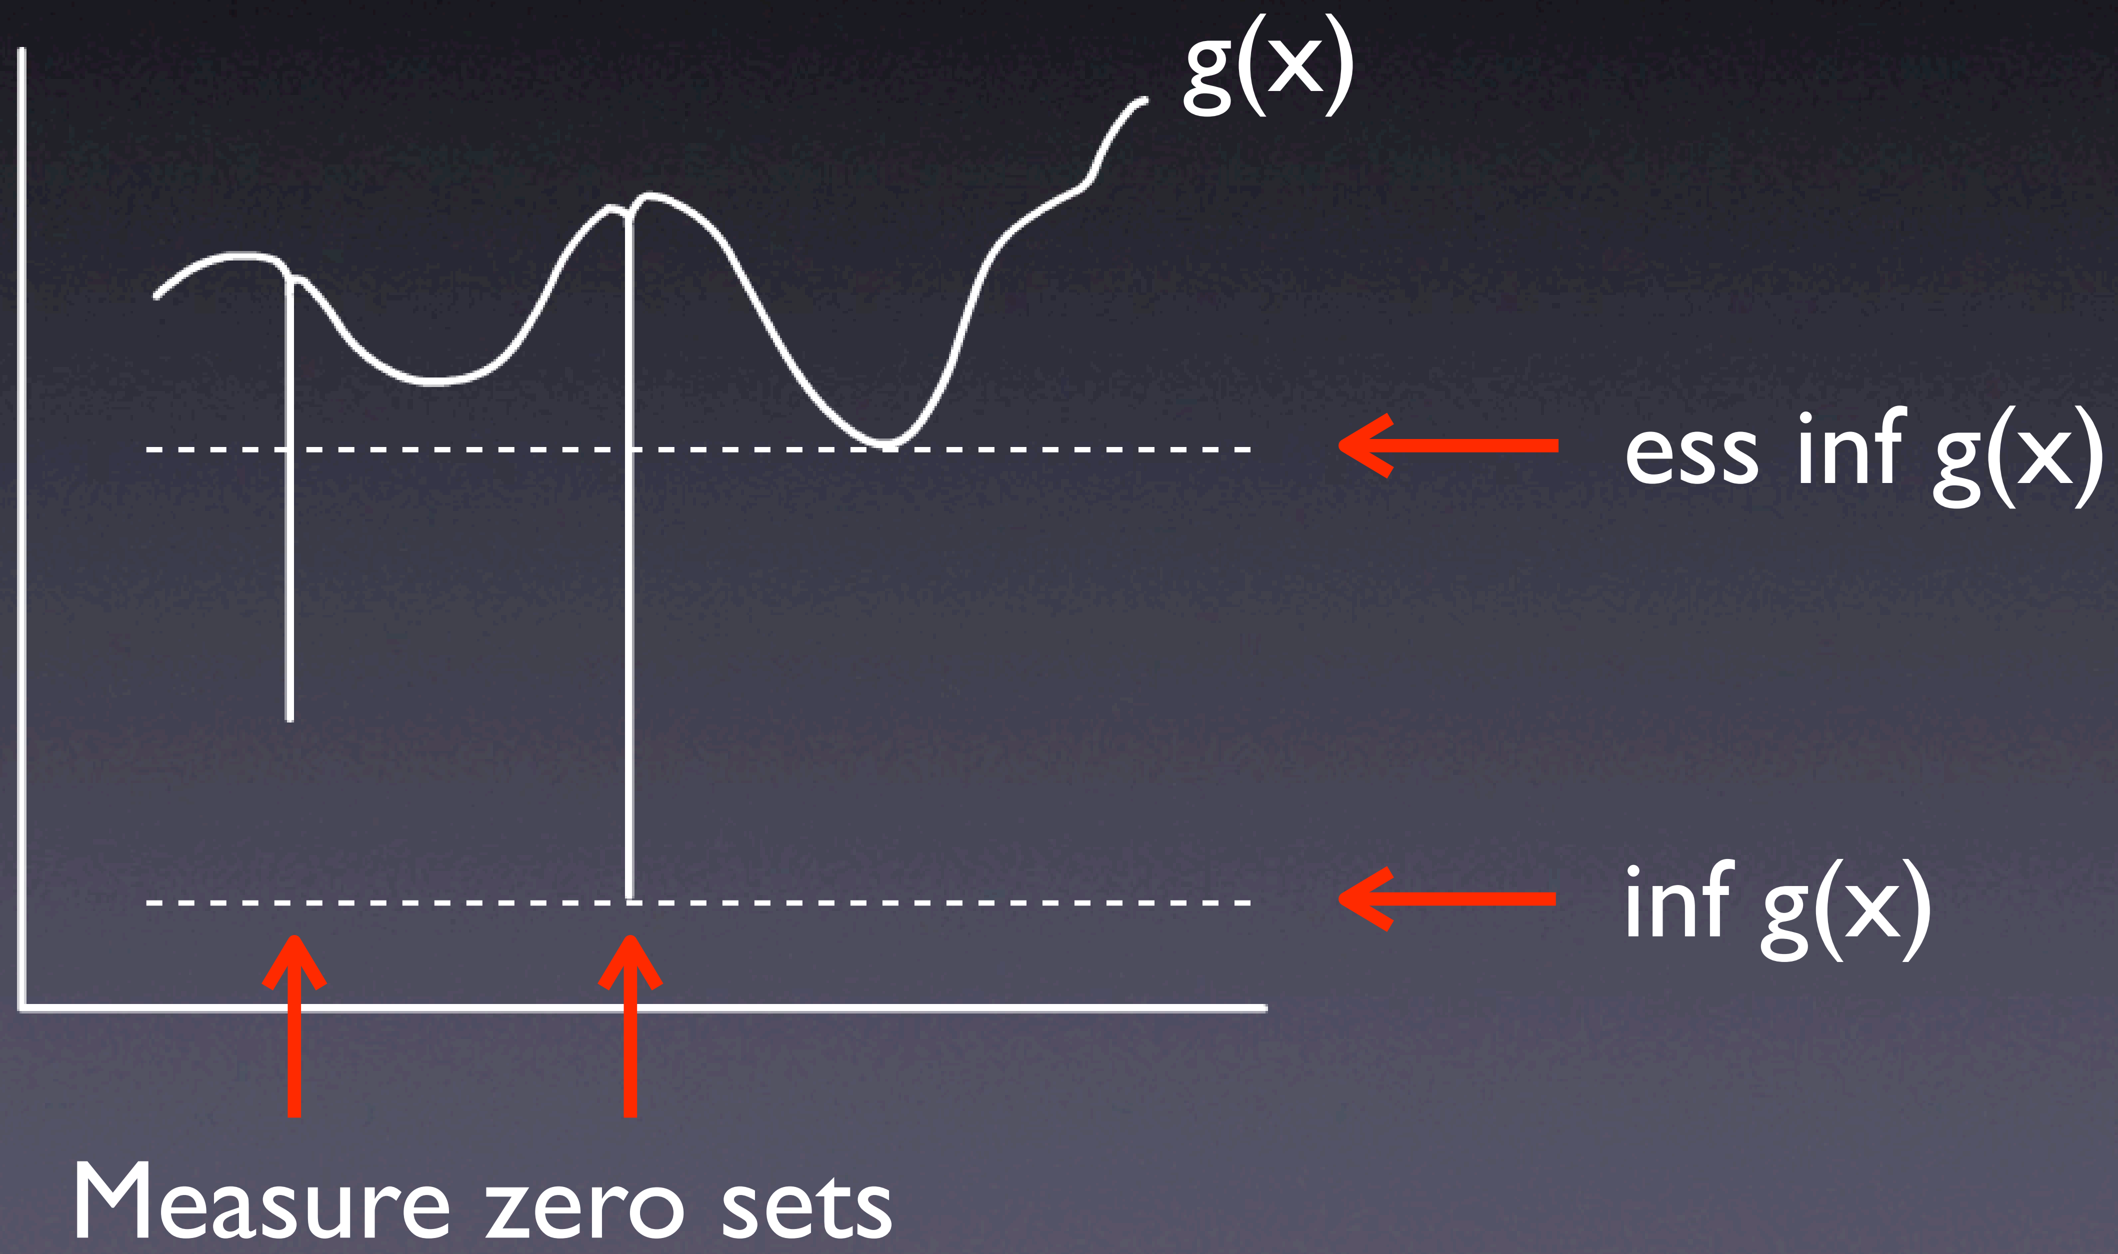

Proof of correctness (sketch)

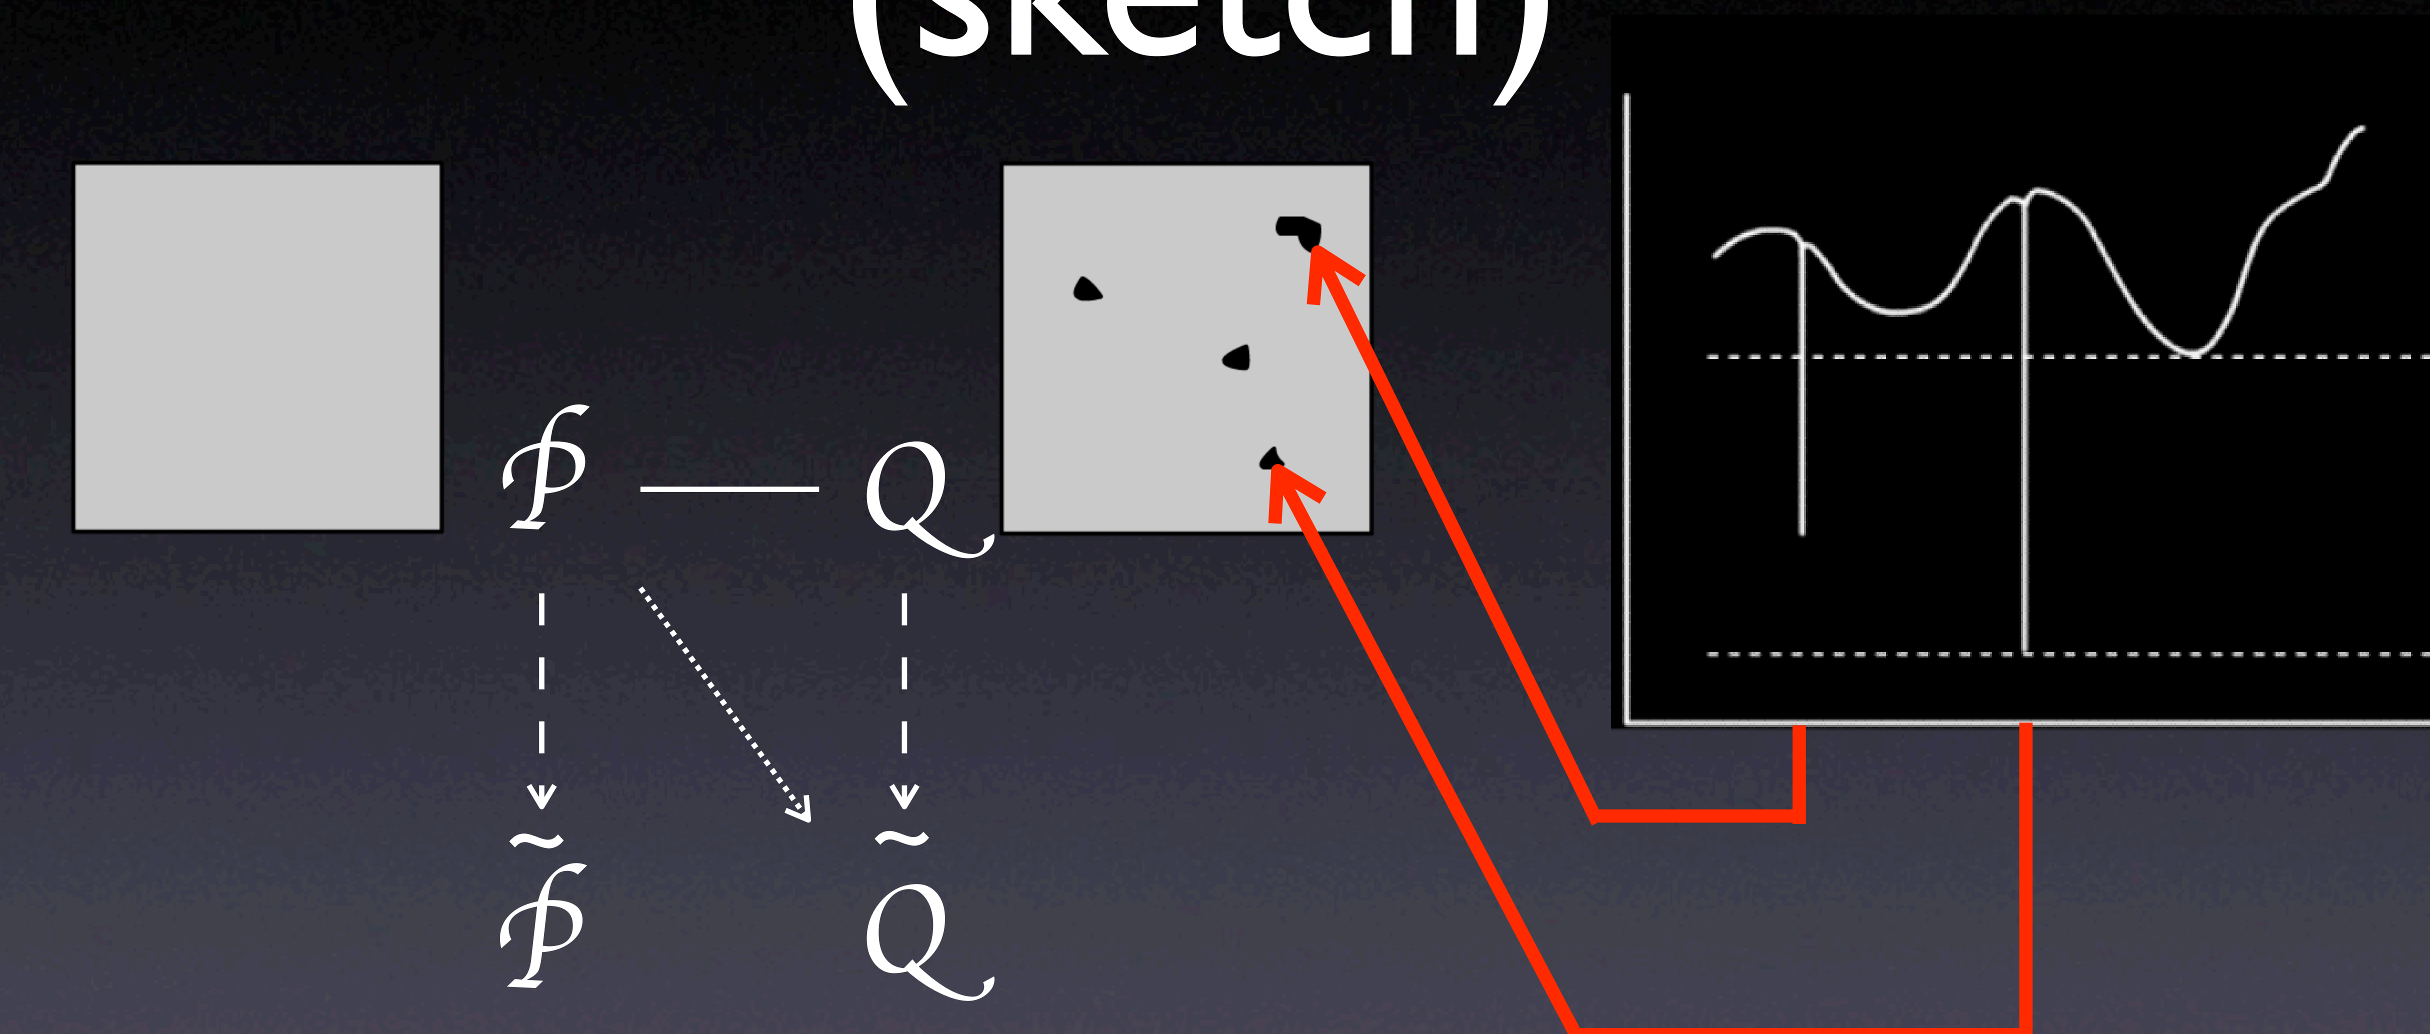

— bisimulation

- - → approximation

······→ sampling approximation

ϵ-homogeneity

M_1 ϵ-homogenous w.r.t. M if

$\exists \Phi: S \rightarrow S_1$ surj. s.t. $\forall s \in S \;\; \forall a \in A$

$$\Sigma_{s' \epsilon S} \mid P_1(\Phi(s), s', a) - \Sigma_{t \epsilon \Phi^{-1}(\{s'\})} P(s, t, a) \mid^k \leq \epsilon^k$$

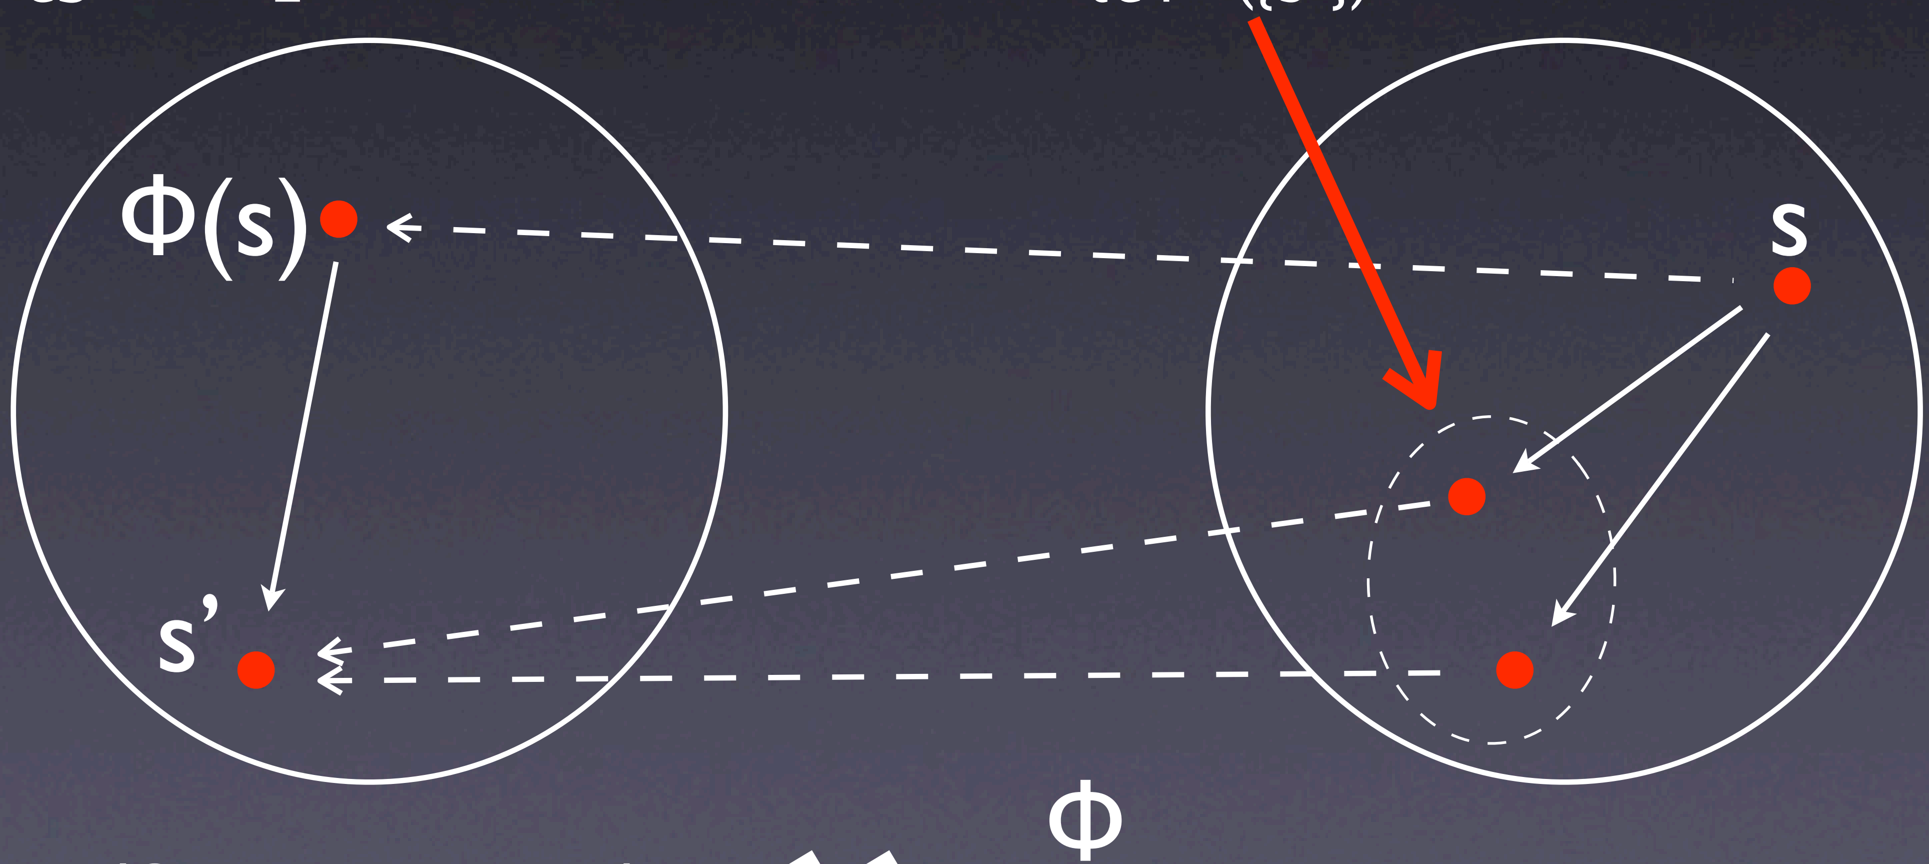

$M_1 = (S_1, A, R_1, P_1)$ ⟸ $(S, A, R, P) = M$

Link between 0-homogeneity and bisimulation

Let $R \equiv 0$, $M_1 = (S_1, A, R, P_1)$, $M = (S, A, R, P)$ be MDP's (and therefore LMP's). Then they are 0-homogenous with mapping Φ iff $\{\Phi^{-1}(\{s'\}) : s' \in S_1\}$ is a bisimulation equivalence relation on M.

Proof idea

Enough: if s_1, s_2 are s.t. $\Phi(s_1) = \Phi(s_2) = s$, then they satisfy the same formulas in $\mathcal{L}_0$.

Structural induction on $\mathcal{L}_0$.

As usual, the "hard" step is $\langle a \rangle_q \varphi$. By induction hypothesis, $[[\varphi]]$ has the form:

$$[[\varphi]] = \cup \{\Phi^{-1}(\{s'_i\})\}$$

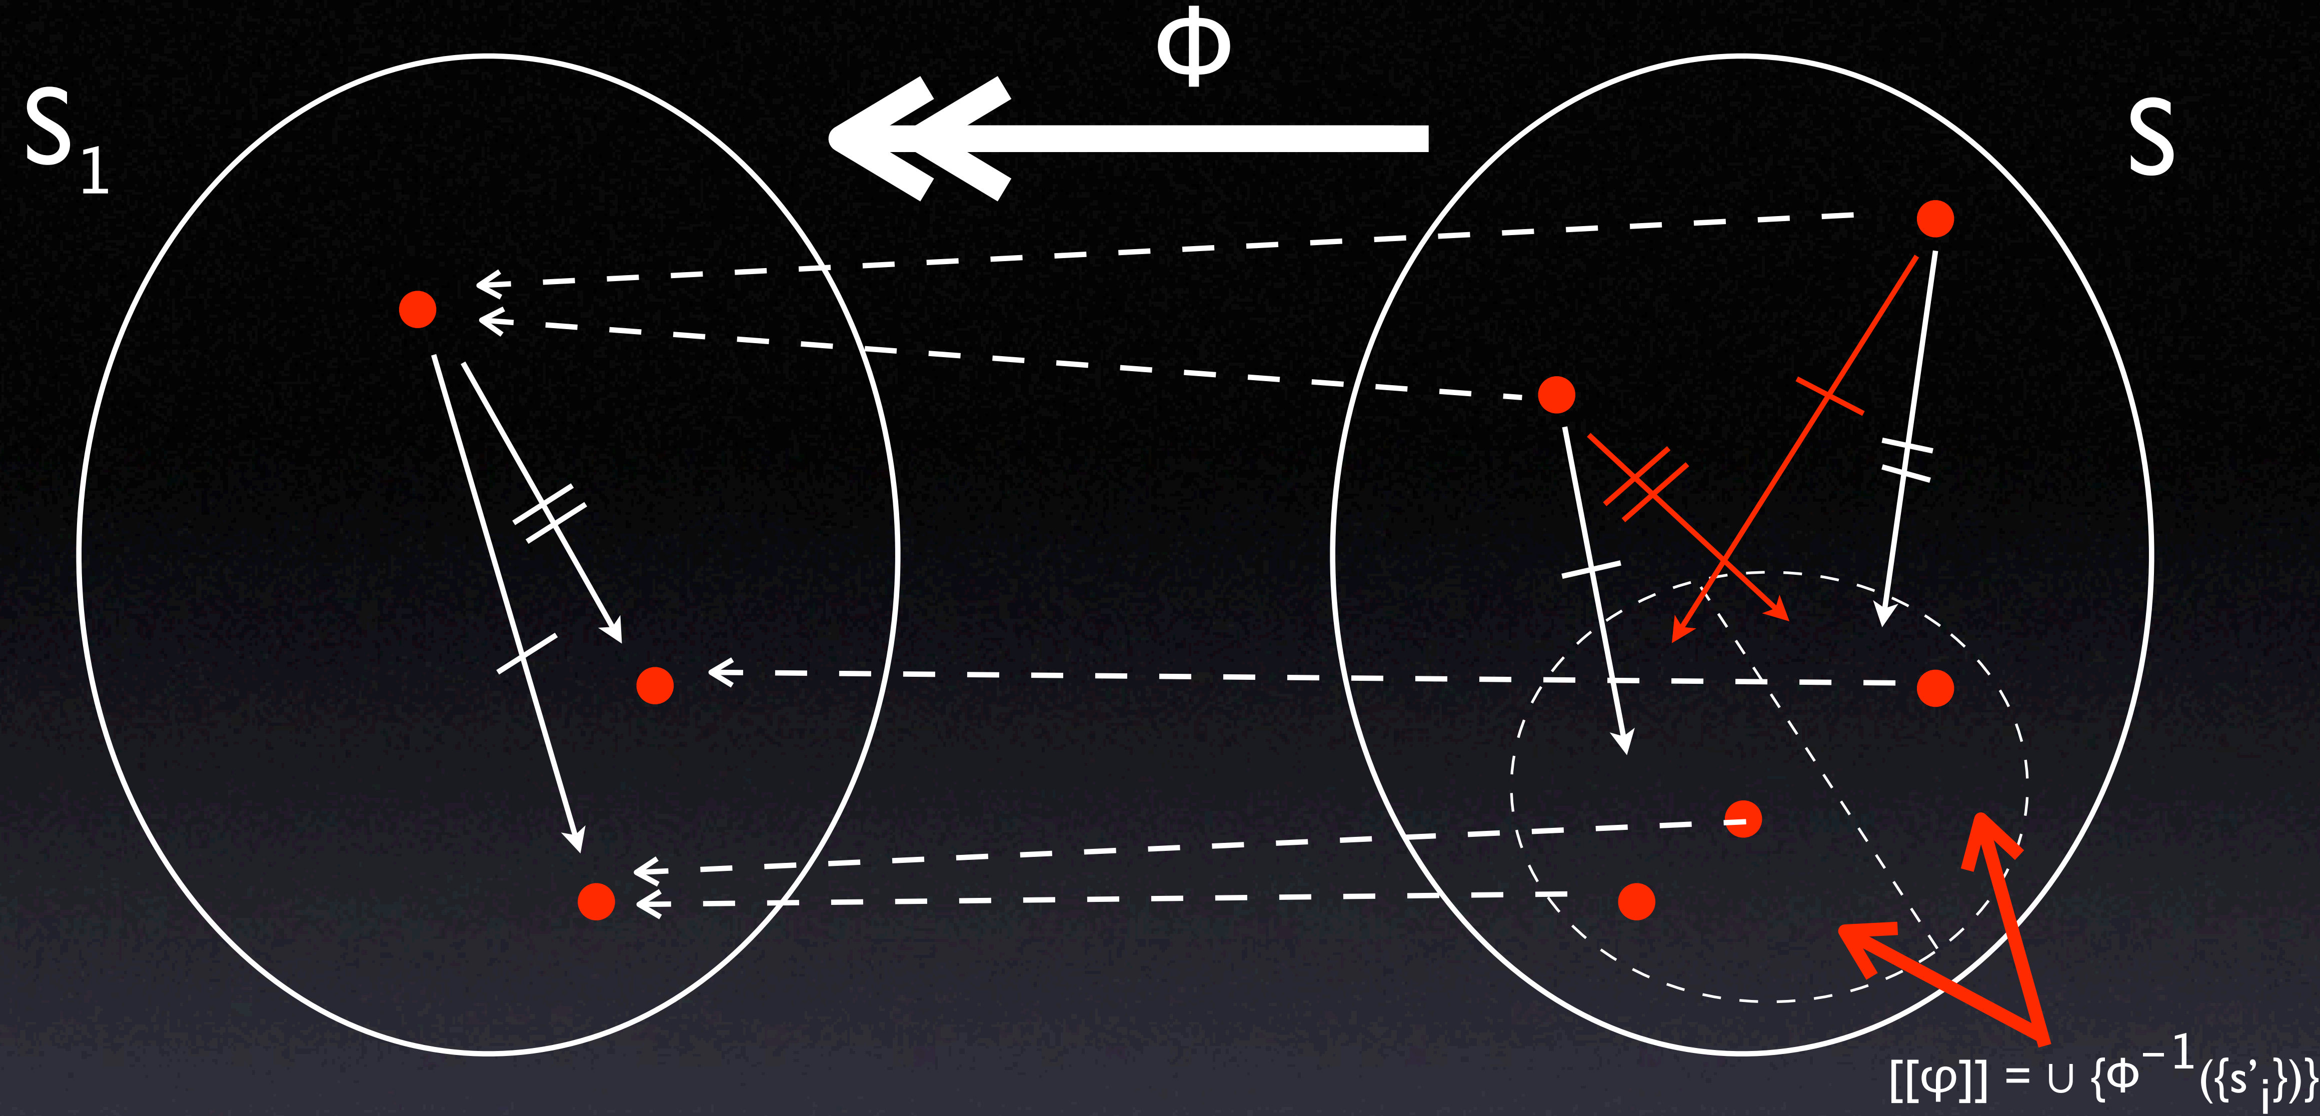

For each of these s'_i, we have, by 0-homogeneity:

$$\Sigma_{t\in\Phi^{-1}(\{s'_i\})} P(s_j, t, a) = P_1(\Phi(s_j), s'_i, a) \text{ for } j=1,2$$

$$\therefore P(s_j, [[\varphi]], a) = \Sigma_i \Sigma_{t\in\Phi^{-1}(\{s'_i\})} P_1(s, s'_i, a)$$

$$\therefore P(s_1, [[\varphi]], a) = P_1(s_2, [[\varphi]], a)$$

$$\therefore s_1 \models \langle a \rangle_q \varphi \Leftrightarrow s_2 \models \langle a \rangle_q \varphi$$