Verification in POMDPs

Privacy, machine teaching and other belief-related problems

Ufuk Topcu

The University of Texas at Austin

Slides originally prepared by Bo Wu.

http://u-t-autonomous.info

aUTonomous SYSTEMS GROUP

Protecting mission-critical information

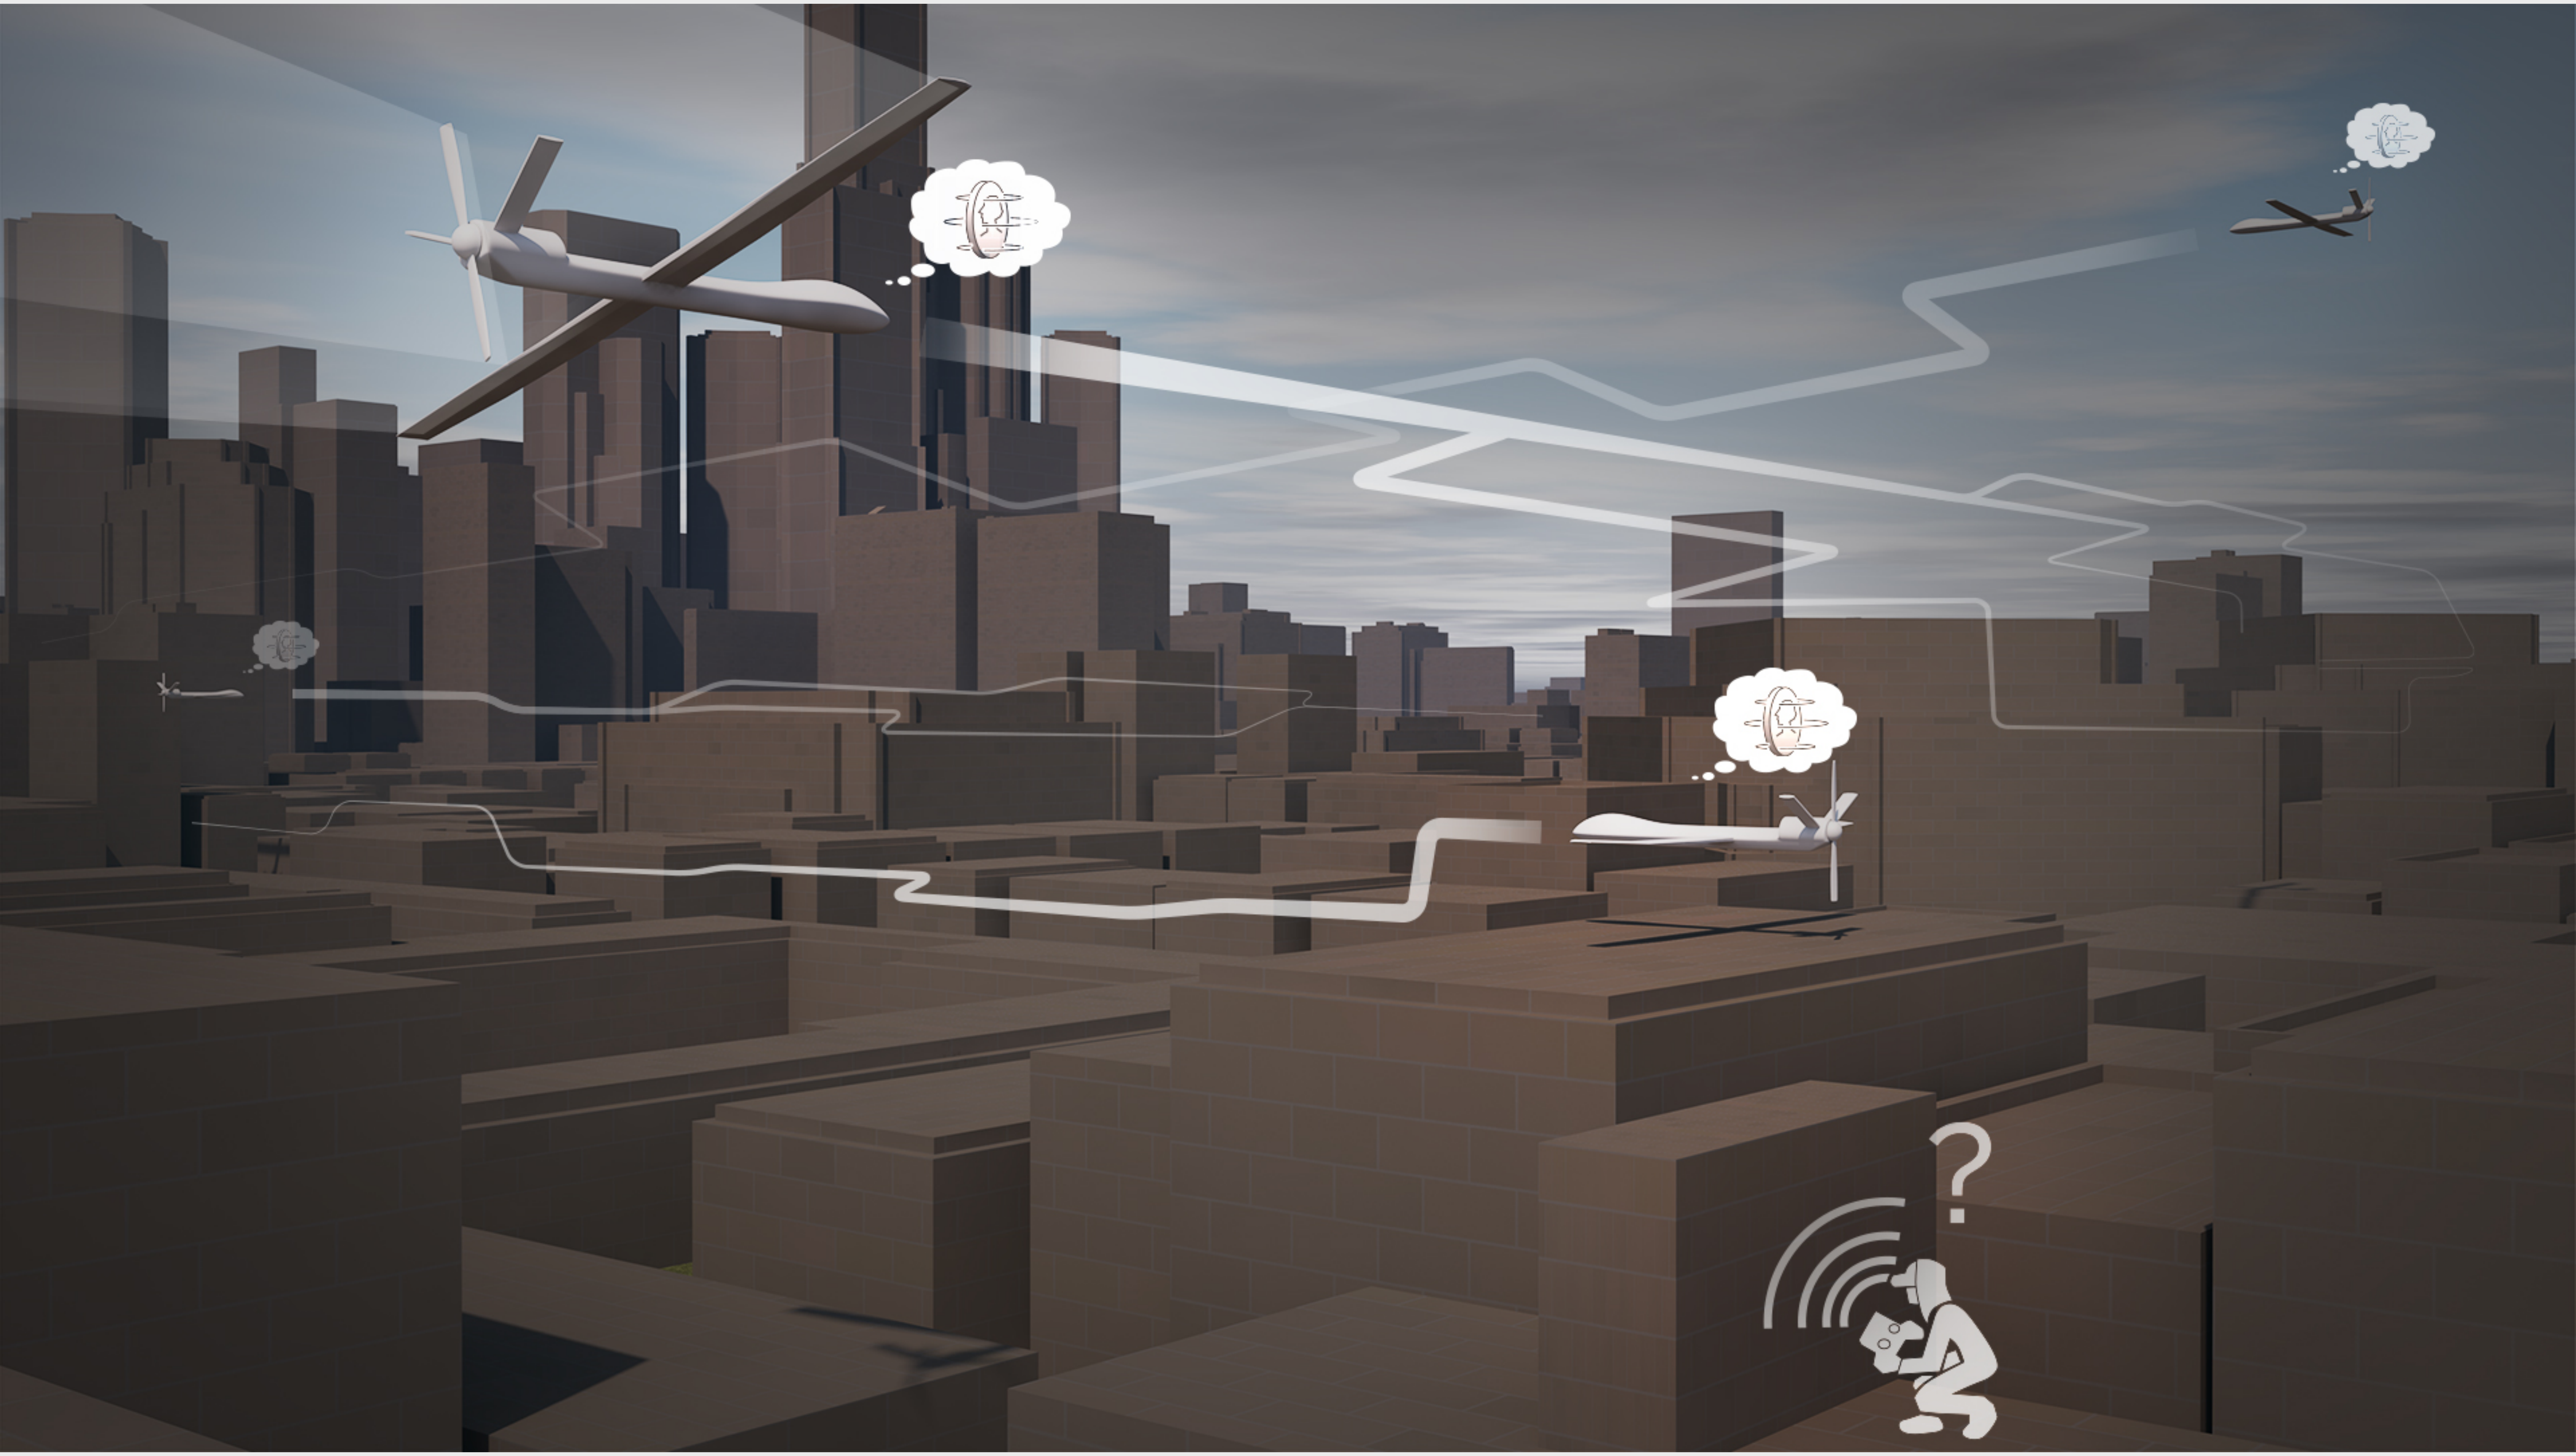

Protecting mission-critical information

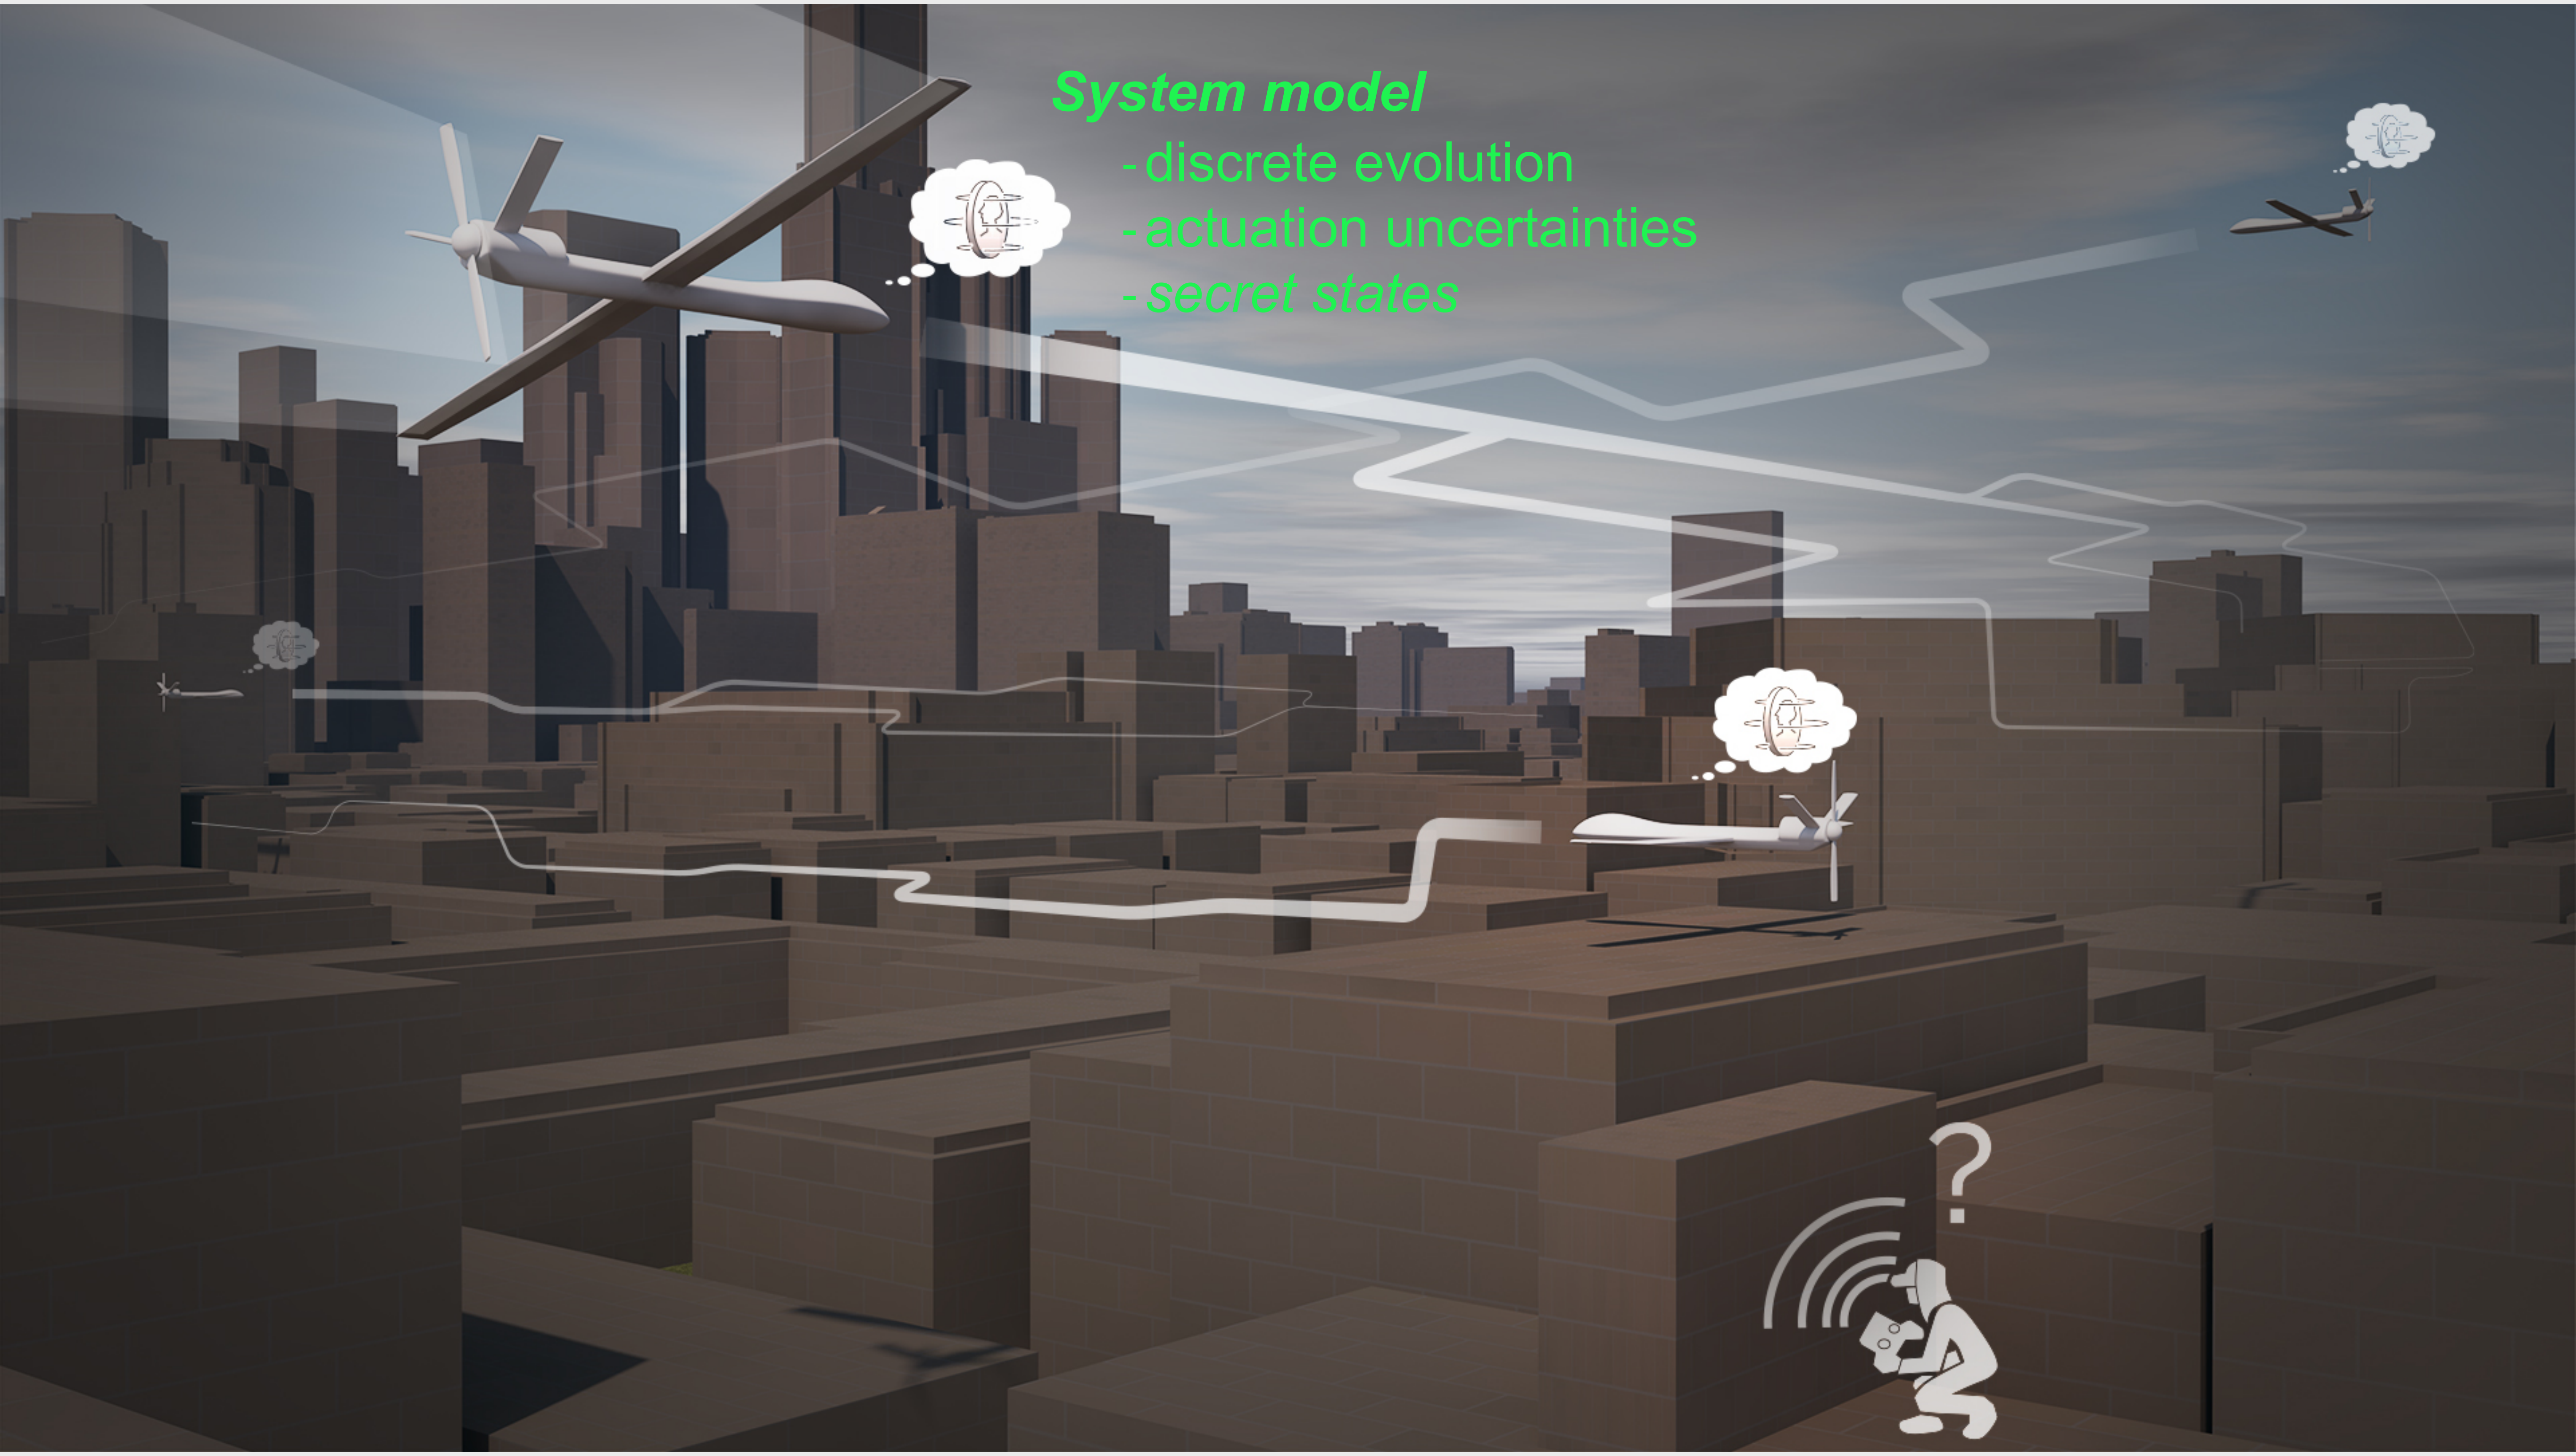

Protecting mission-critical information

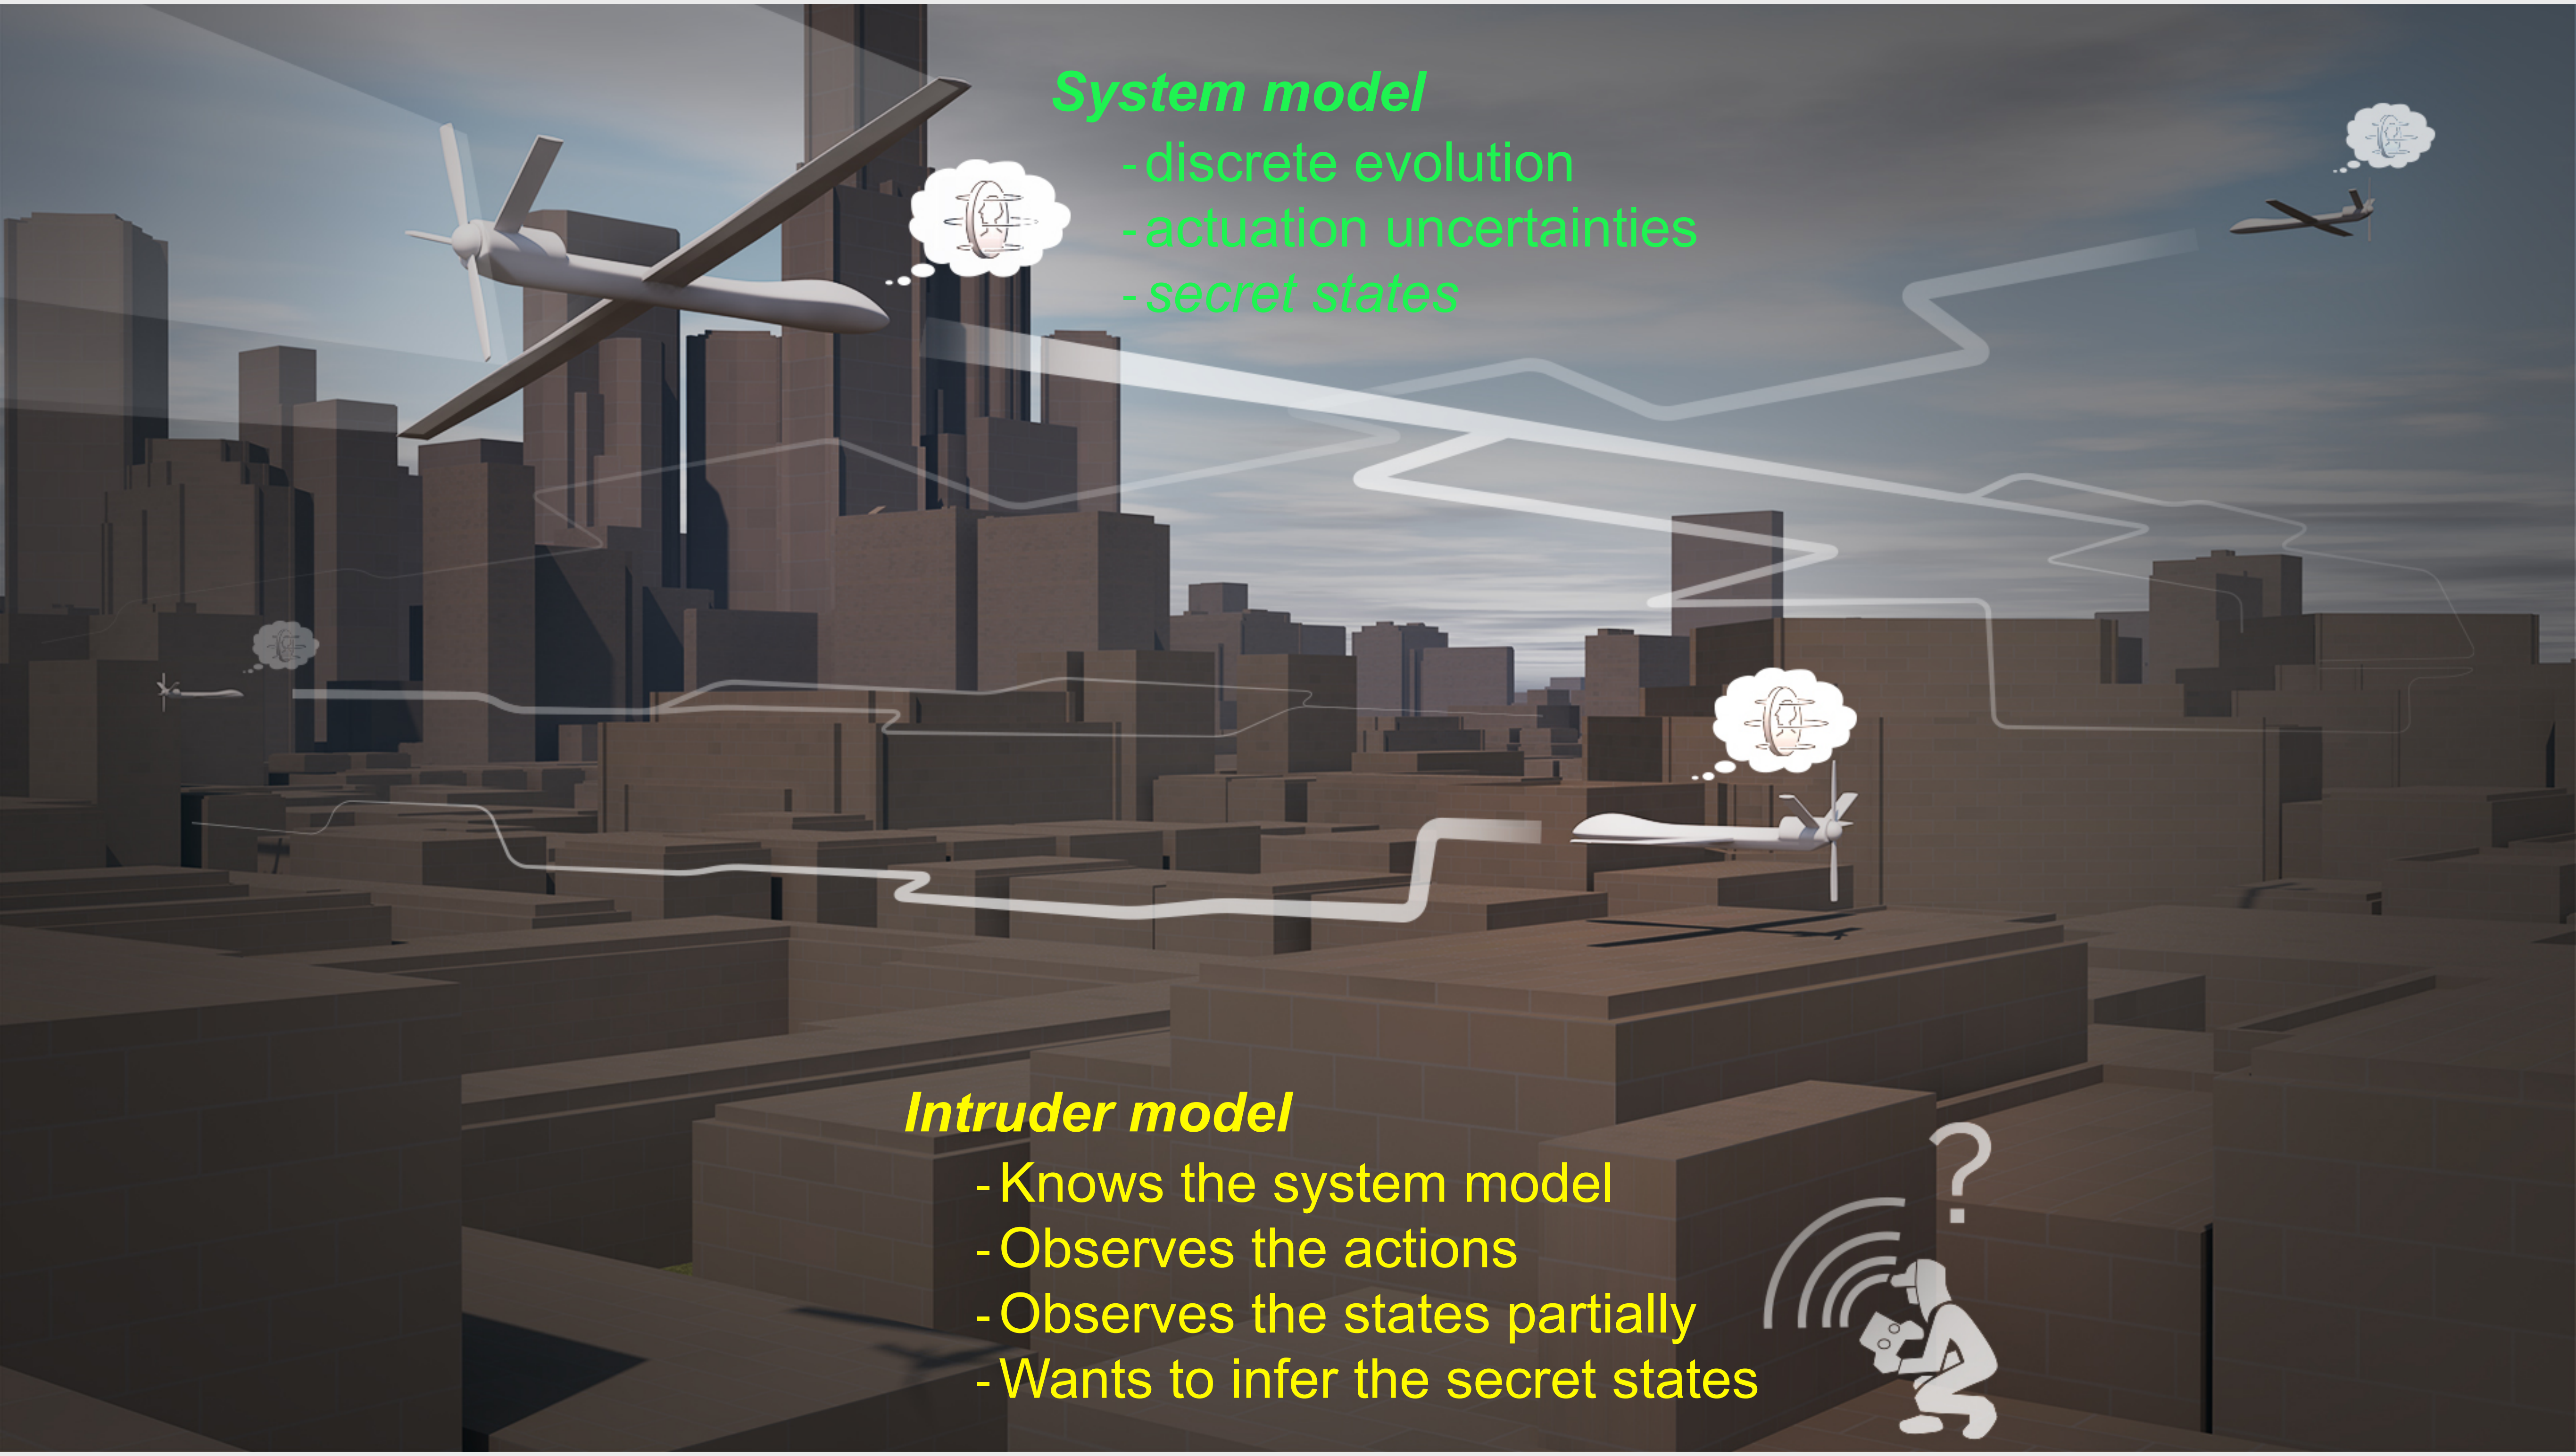

A formulation based on POMDPs

The system is modeled by a Markov decision process (MDP)

$$\mathcal{M} = (Q, \pi, A, T)$$

$Q:$ a finite set of states

$\pi : Q \to [0,1]$ initial distribution

$A:$ a finite set of actions

$T : T(q,a,q') = P(q_t = q' | q_{t-1}, a_{t_1})$

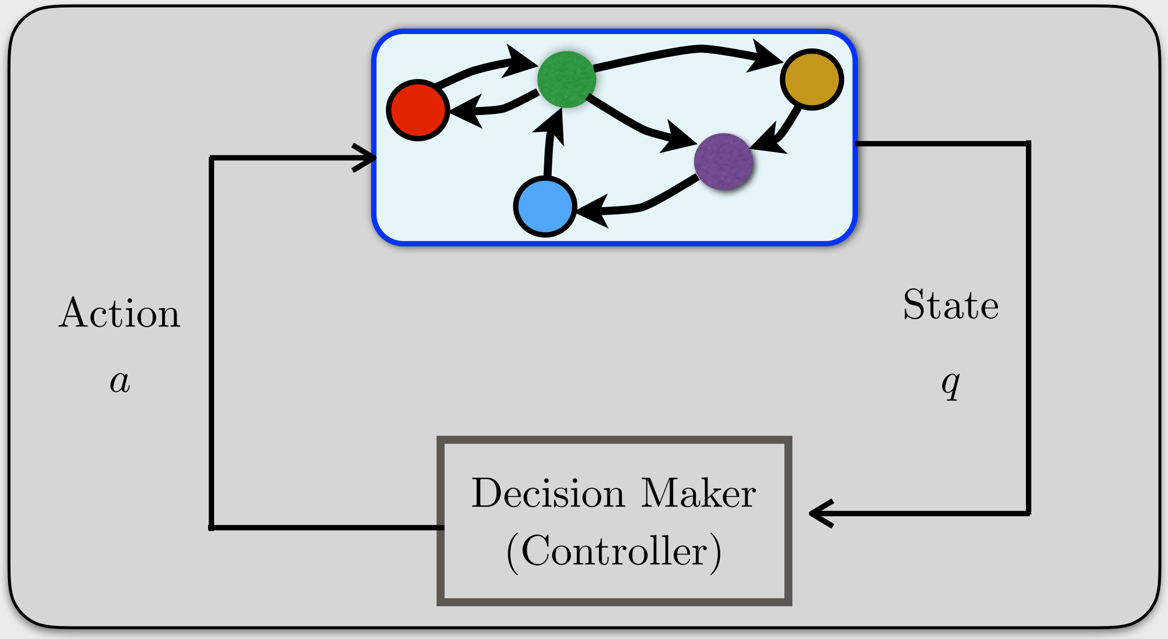

A formulation based on POMDPs

The system is modeled by a Markov decision process (MDP)

$$\mathcal{M} = (Q, \pi, A, T)$$

$Q :$ a finite set of states

$\pi : Q \rightarrow [0,1]$ initial distribution

$A :$ a finite set of actions

$T : T(q,a,q') = P(q_t = q' | q_{t-1}, a_{t_1})$

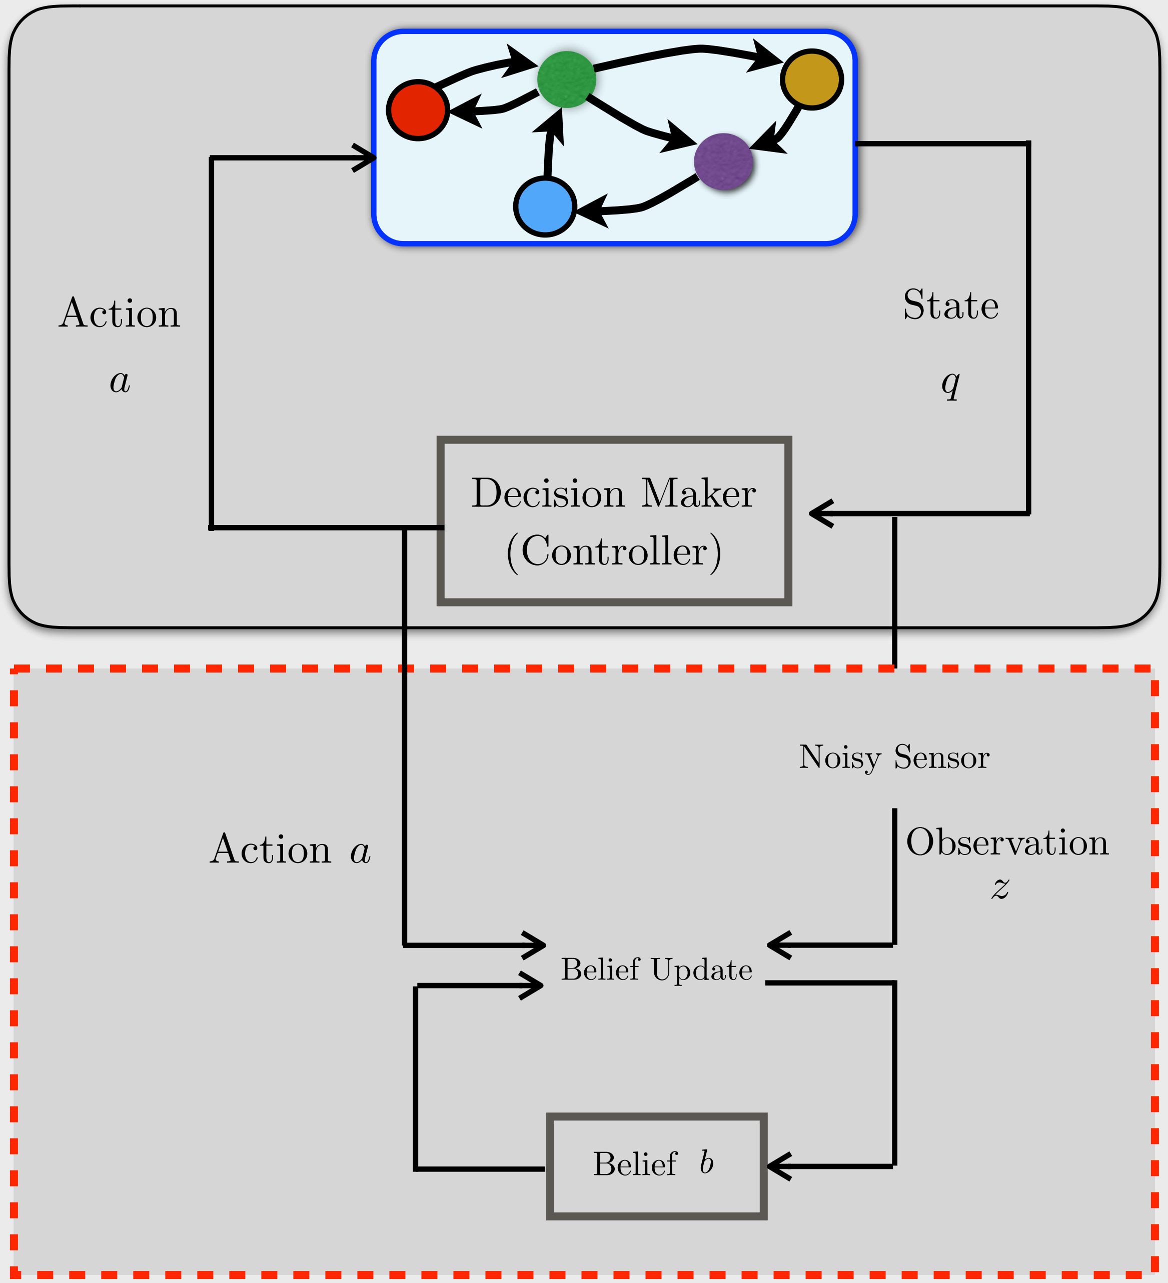

Intruder: partially observe (PO) the system

$$\mathcal{P} = (\underbrace{Q, \pi, A, T}_{\mathcal{M}}, Z, O)$$

$Z :$ a finite set of observations

$O :$ observation function where

$O(q,a,z) = P(z_t = z | q_t = q, a_{t-1} = a)$

A formulation based on POMDPs

The system is modeled by a Markov decision process (MDP)

$$\mathcal{M} = (Q, \pi, A, T)$$

$Q:$ a finite set of states

$\pi : Q \rightarrow [0,1]$ initial distribution

$A:$ a finite set of actions

$T : T(q, a, q') = P(q_t = q' | q_{t-1}, a_{t_1})$

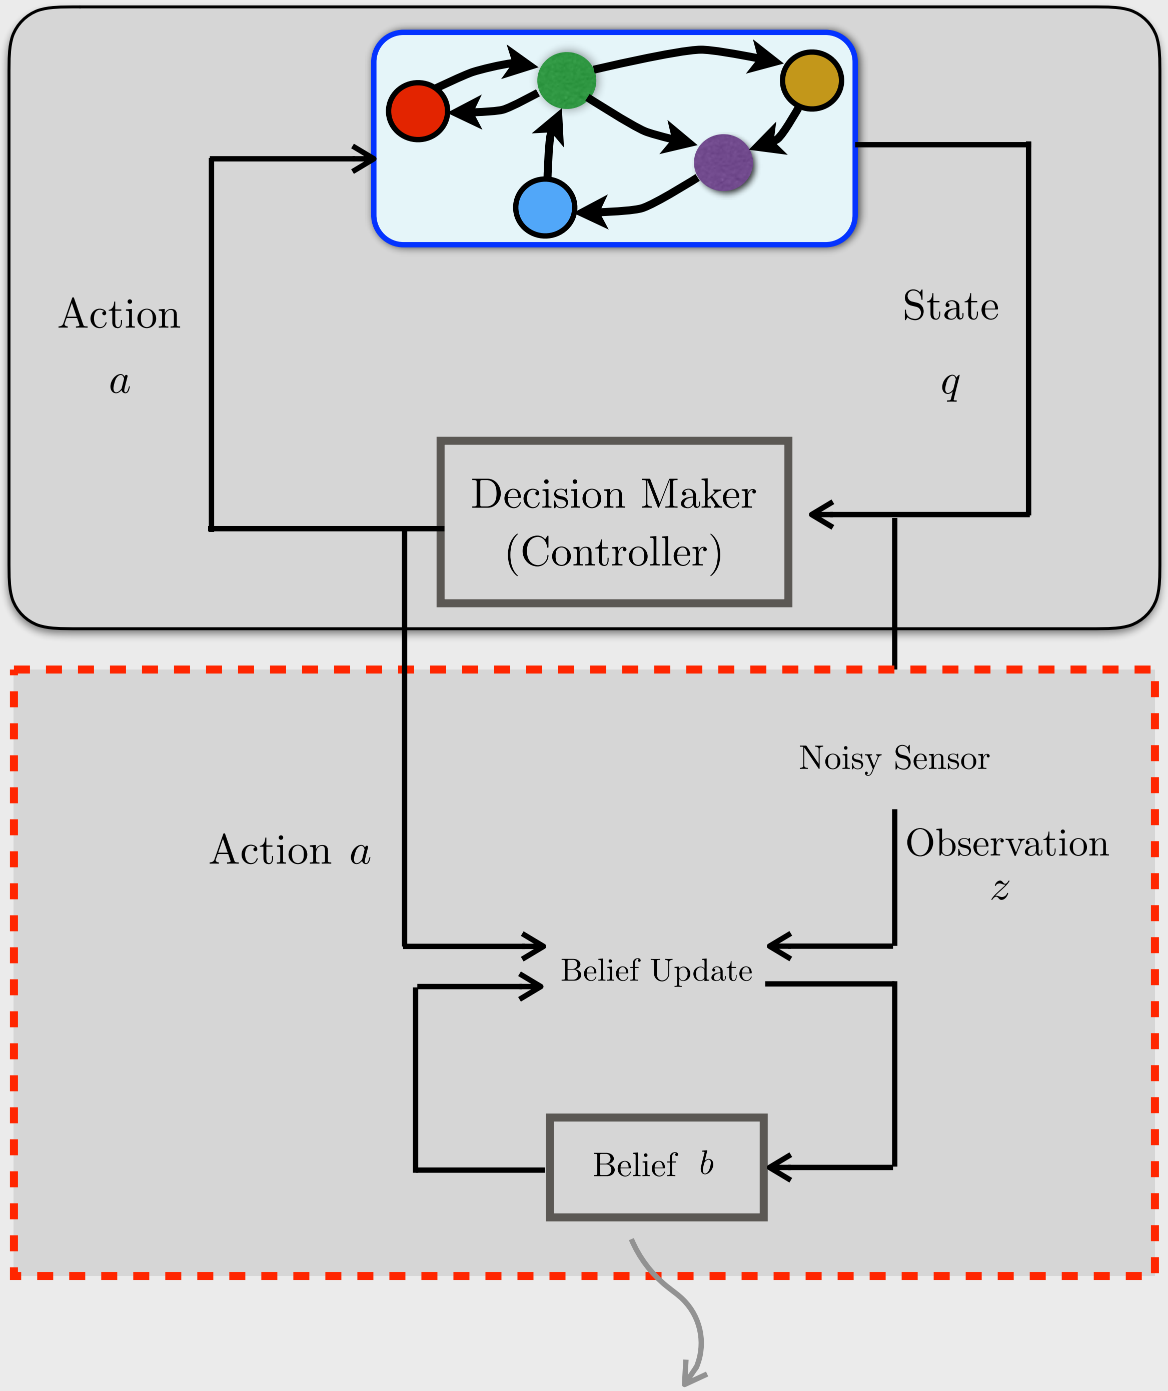

Intruder: partially observe (PO) the system

$$\mathcal{P} = (\underbrace{Q, \pi, A, T}_{\mathcal{M}}, Z, O)$$

$Z:$ a finite set of observations

$O:$ observation function where

$O(q, a, z) = P(z_t = z | q_t = q, a_{t-1} = a)$

Intruder's belief

$$b : Q \rightarrow [0,1], \sum_{q \in Q} b(q) = 1$$

Belief evolution as a switched dynamical system

$$b_t(q) = \qquad f_a^q(b_{t-1}, z_t) = P(q'|z_t, a_{t-1}, b_{t-1})$$

Belief evolution as a switched dynamical system

new belief　　old belief

$$b_t(q) = \quad f_a^q(b_{t-1}, z_t) = P(q'|z_t, a_{t-1}, b_{t-1})$$

action　　observation

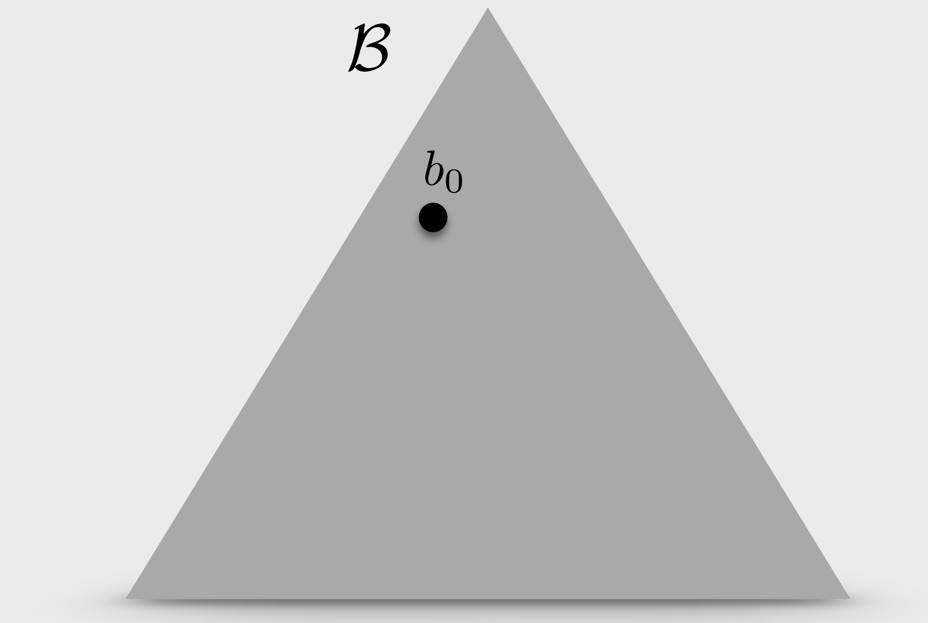

Belief evolution as a switched dynamical system

new belief

old belief

$$b_t(q) = \quad f_a^q(b_{t-1}, z_t) = P(q'|z_t, a_{t-1}, b_{t-1})$$

action

observation

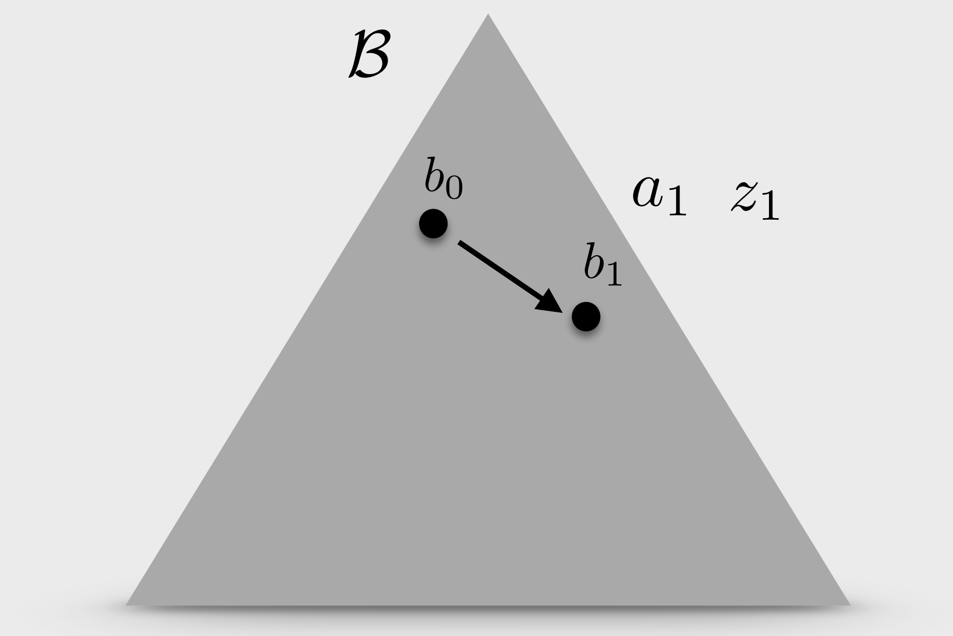

Belief evolution as a switched dynamical system

new belief, old belief

$$b_t(q) = f_a^q(b_{t-1}, z_t) = P(q'|z_t, a_{t-1}, b_{t-1})$$

action, observation

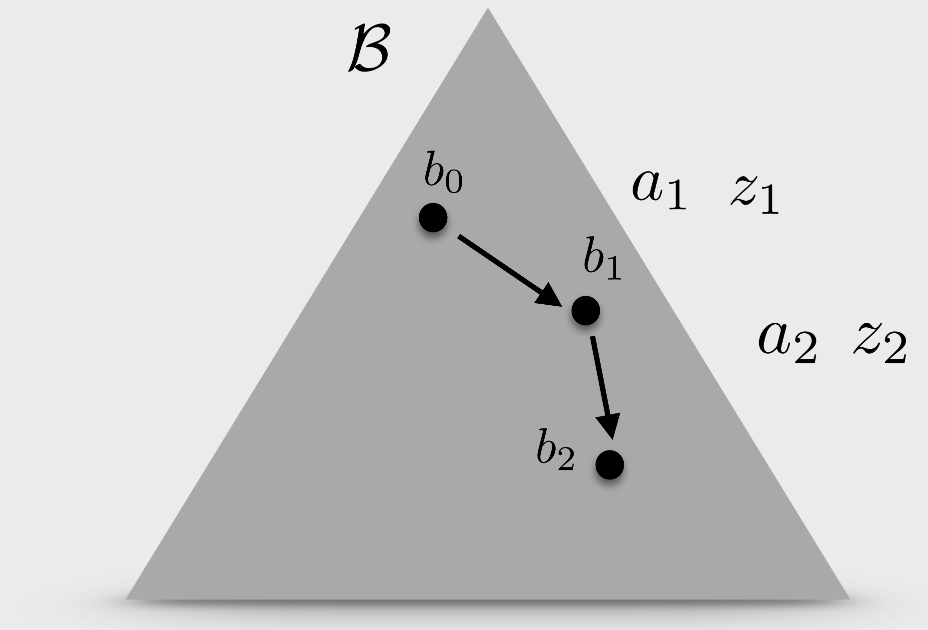

Belief evolution as a switched dynamical system

new belief — old belief

$$b_t(q) = \quad f_a^q(b_{t-1}, z_t) = P(q' | z_t, a_{t-1}, b_{t-1})$$

action — observation

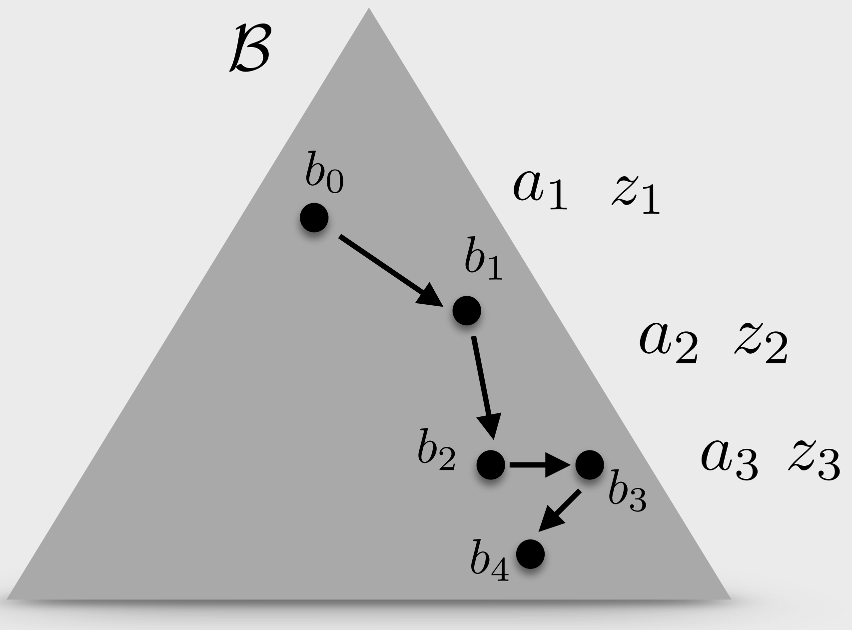

Belief evolution as a switched dynamical system

new belief

old belief

$$b_t(q) = f_a^q(b_{t-1}, z_t) = P(q'|z_t, a_{t-1}, b_{t-1})$$

action

observation

$$= \frac{O(q', a, z)\sum_{q\in Q} T(q, a, q')b_{t-1}(q)}{\sum_{q'\in Q} O(q', a, z)\sum_{q\in Q,} T(q, a, q')b_{t-1}(q)}$$

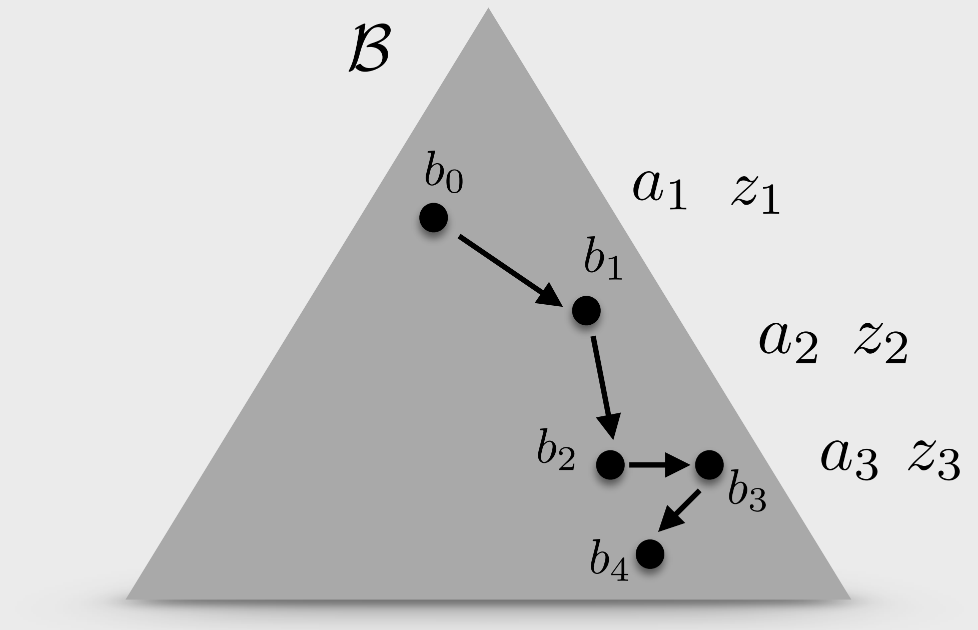

Belief evolution as a switched dynamical system

new belief — old belief

$$b_t(q) = f_a^q(b_{t-1}, z_t) = P(q'|z_t, a_{t-1}, b_{t-1})$$

action — observation

$$= \frac{O(q',a,z)\sum_{q\in Q} T(q,a,q')b_{t-1}(q)}{\sum_{q'\in Q} O(q',a,z)\sum_{q\in Q}, T(q,a,q')b_{t-1}(q)}$$

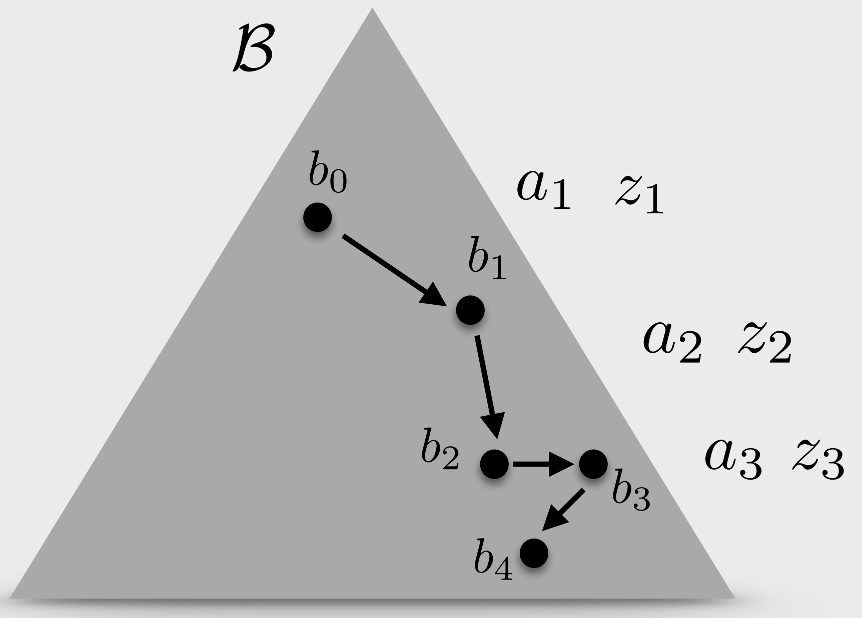

States and actions in the (PO)MDP are discrete and finite
but
the belief evolves over a continuous space

Results in a switched system with modes induced from the actions

$$b_t = f_a\left(b_{t-1}, z_t\right)$$

Privacy in terms of the belief of the intruder

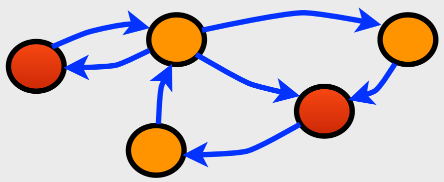

A set of secret states: $Q_s \subset Q$

to represent the information we want to keep private, e.g., intent, target, goal, preference, etc.

Privacy in terms of the belief of the intruder

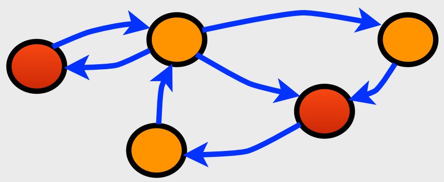

A set of secret states: $Q_s \subset Q$
to represent the information we want to keep private, e.g., intent, target, goal, preference, etc.

Privacy is breached if the intruder becomes confident in that the system is in a secret state with a probability larger than a threshold at a time t:

$$\sum_{q \in Q_s} b_t(q) > \gamma$$

(sum over the secret states)

Privacy verification

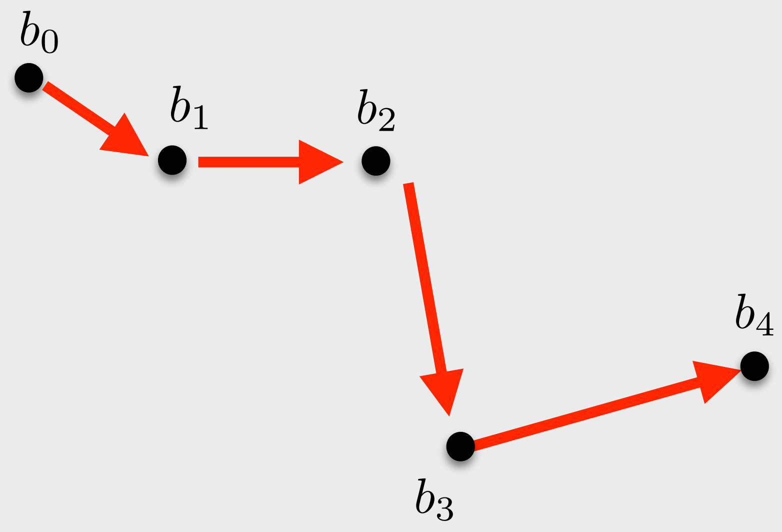

Privacy verification

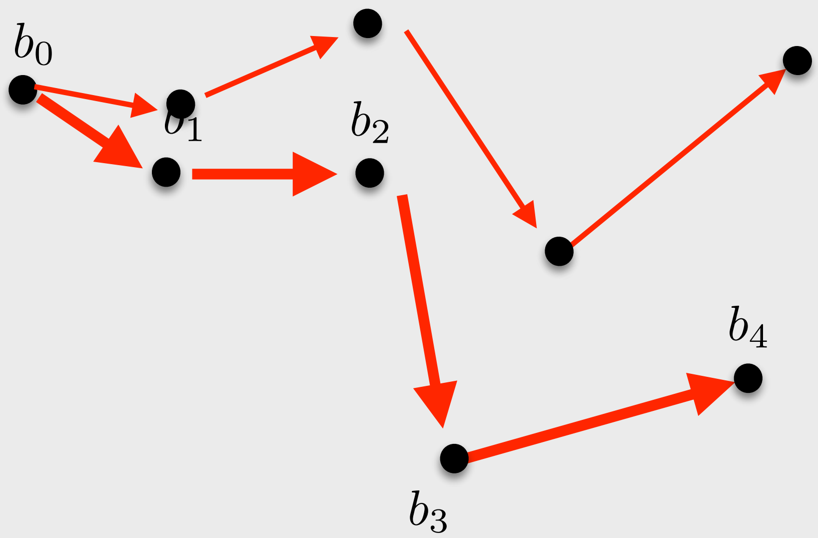

Privacy verification

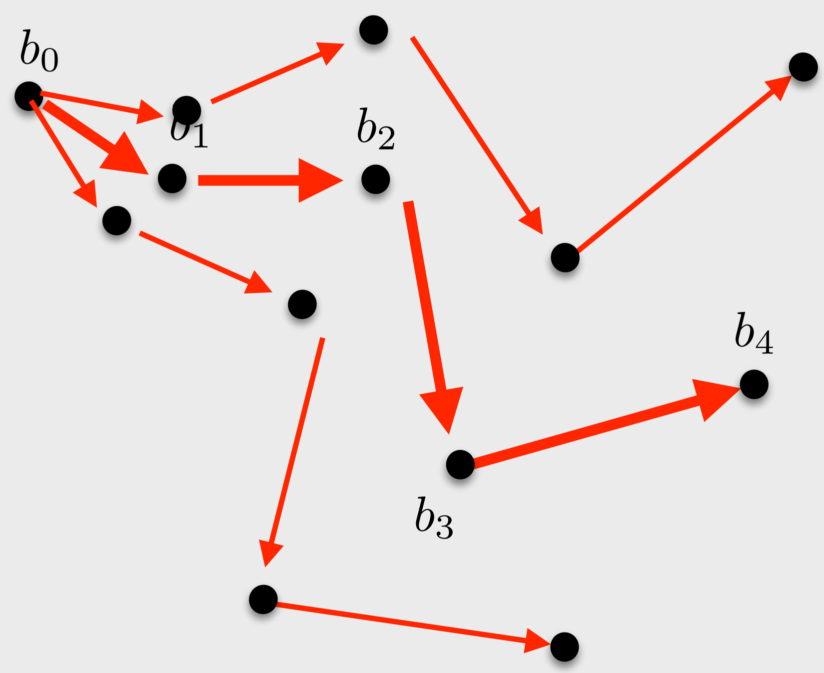

Privacy verification

$\mathcal{R}(b_0)$

Set of "reachable" belief states

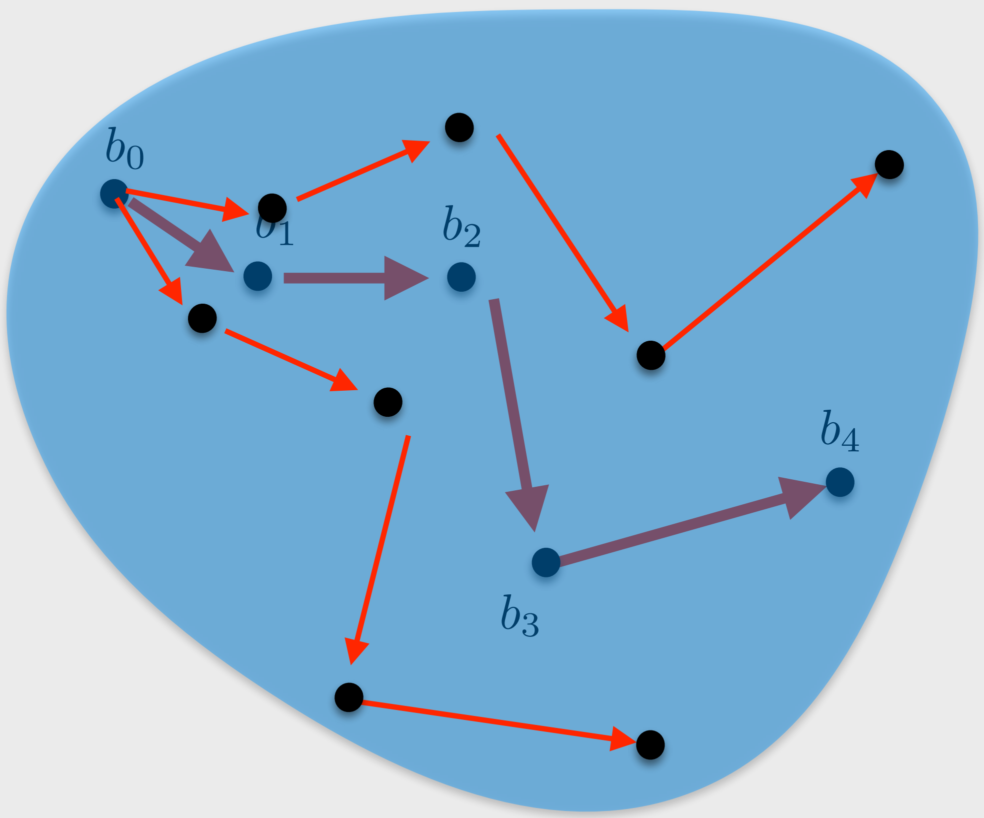

Privacy verification

$\mathcal{R}(b_0)$

Set of "reachable" belief states

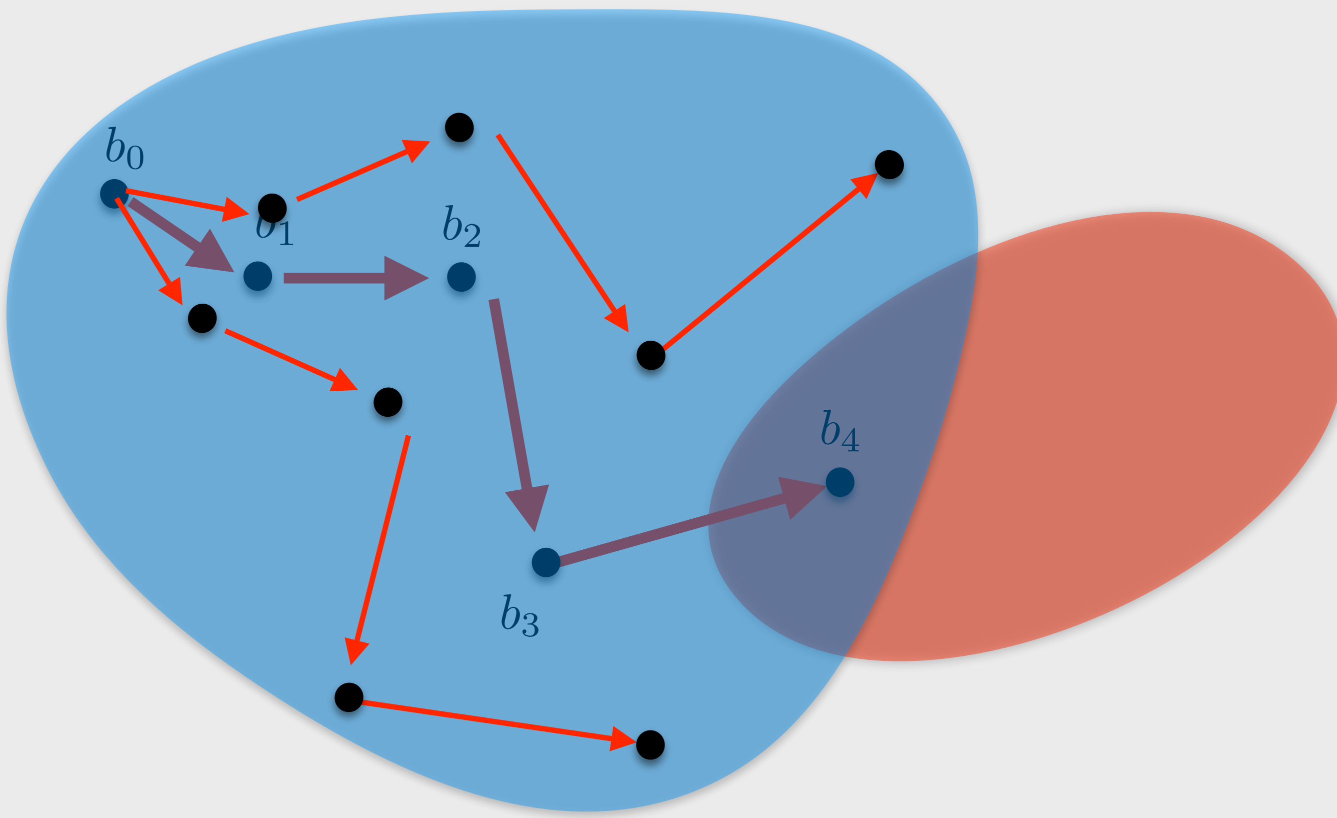

"Unsafe" belief set

$$\mathcal{B}_u = \{b | \sum_{q \in Q_s} b(q) > \gamma\}$$

- Verify whether

$$\mathcal{R}(b_0) \bigcap \mathcal{B}_u = \emptyset$$

- That is, privacy is not breached at <u>any</u> time t.

How to attempt to verify the set emptiness?

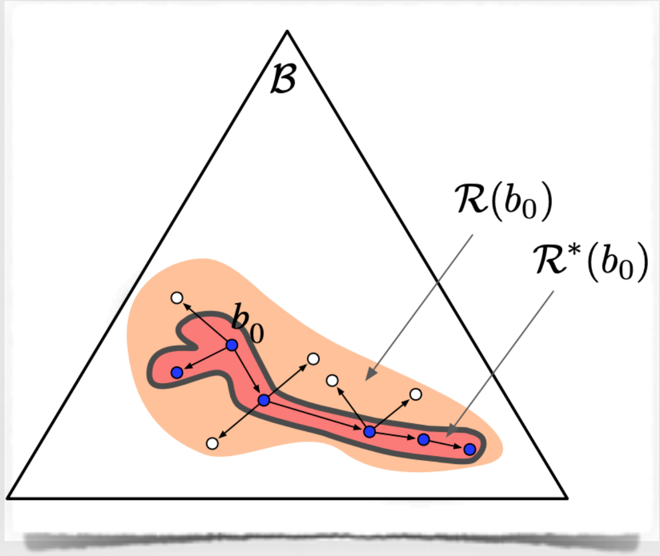

In general an undecidable problem.

POMDPs have been heavily studied in planning.
- Usually heuristics (or impractical exhaustive methods)
- Approximate results with no "guarantee"
- Hard to adapt to the problem in hand

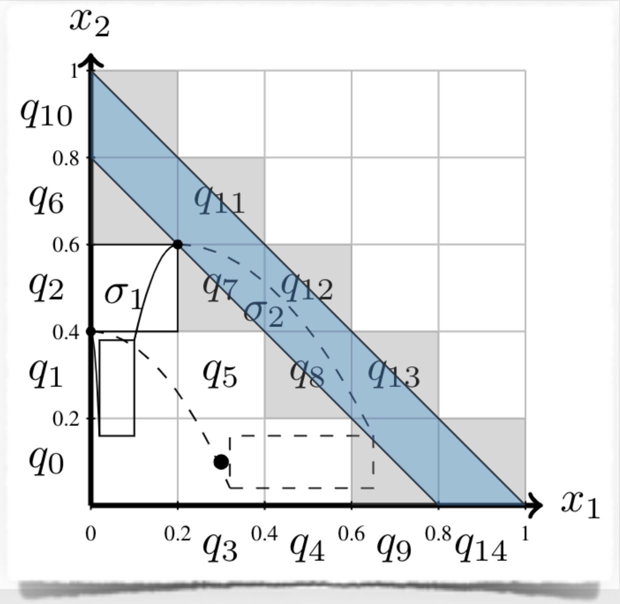

Abstraction-based techniques, i.e., discretization, have their intrinsic limitations

The proposed approach: search for algebraic certificates that witness privacy

Over-approximation of the reachable belief set

If there exists a function V such that (with additional standard technical assumptions)

$$V\left(f_a\left(b_{t-1}, z\right)\right) - V\left(b_{t-1}\right) < 0 \qquad \forall a \in A,\ z \in Z,\ b_{t-1} \in \mathcal{V}$$

$$b_0 \in \mathcal{V} = \{b \in \mathcal{B} \mid V(b) \leq 1\}$$

Then, $\mathcal{R}(b_0) \subseteq \mathcal{V}$.

$\mathcal{V} = \{b \in \mathcal{B} \mid V(b) \leq 1\}$

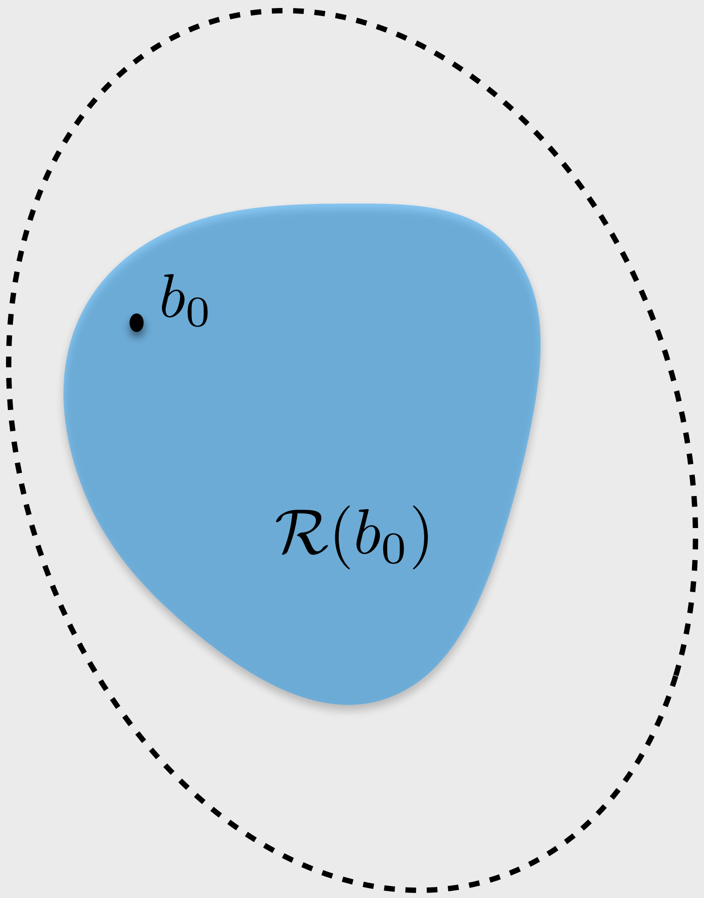

Over-approximation of the reachable belief set

If there exists a function *V* such that (with additional standard technical assumptions)

$$V\left(f_a\left(b_{t-1}, z\right)\right) - V\left(b_{t-1}\right) < 0 \qquad \forall a \in A,\ z \in Z,\ b_{t-1} \in \mathcal{V}$$

$$b_0 \in \mathcal{V} = \{b \in \mathcal{B} \mid V(b) \leq 1\}$$

Then, $\mathcal{R}(b_0) \subseteq \mathcal{V}$.

$\mathcal{V} = \{b \in \mathcal{B} \mid V(b) \leq 1\}$

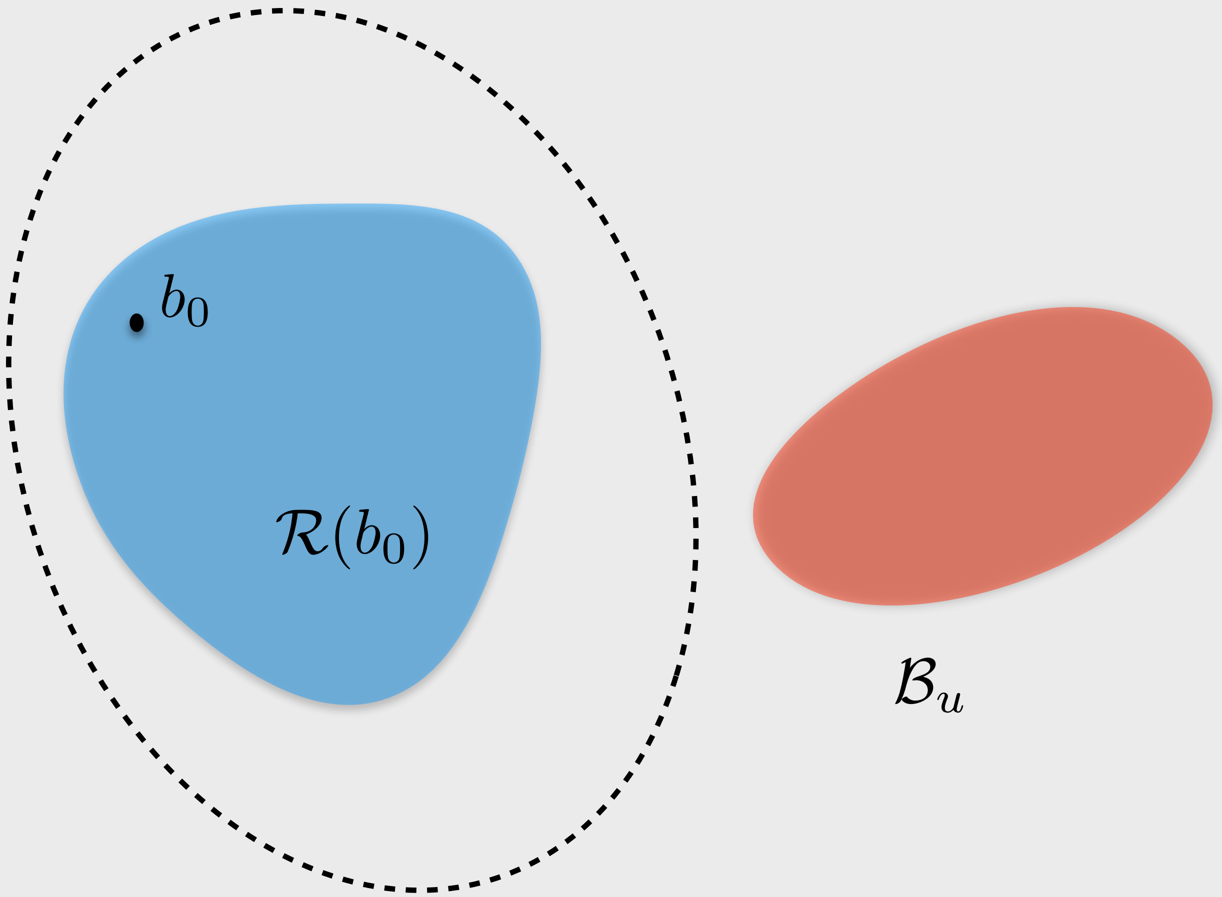

Over-approximation of the reachable belief set

If there exists a function *V* such that (with additional standard technical assumptions)

$$V\left(f_a\left(b_{t-1}, z\right)\right) - V\left(b_{t-1}\right) < 0 \qquad \forall a \in A,\ z \in Z,\ b_{t-1} \in \mathcal{V}$$

$$b_0 \in \mathcal{V} = \{b \in \mathcal{B} \mid V(b) \leq 1\}$$

Then, $\mathcal{R}(b_0) \subseteq \mathcal{V}$.

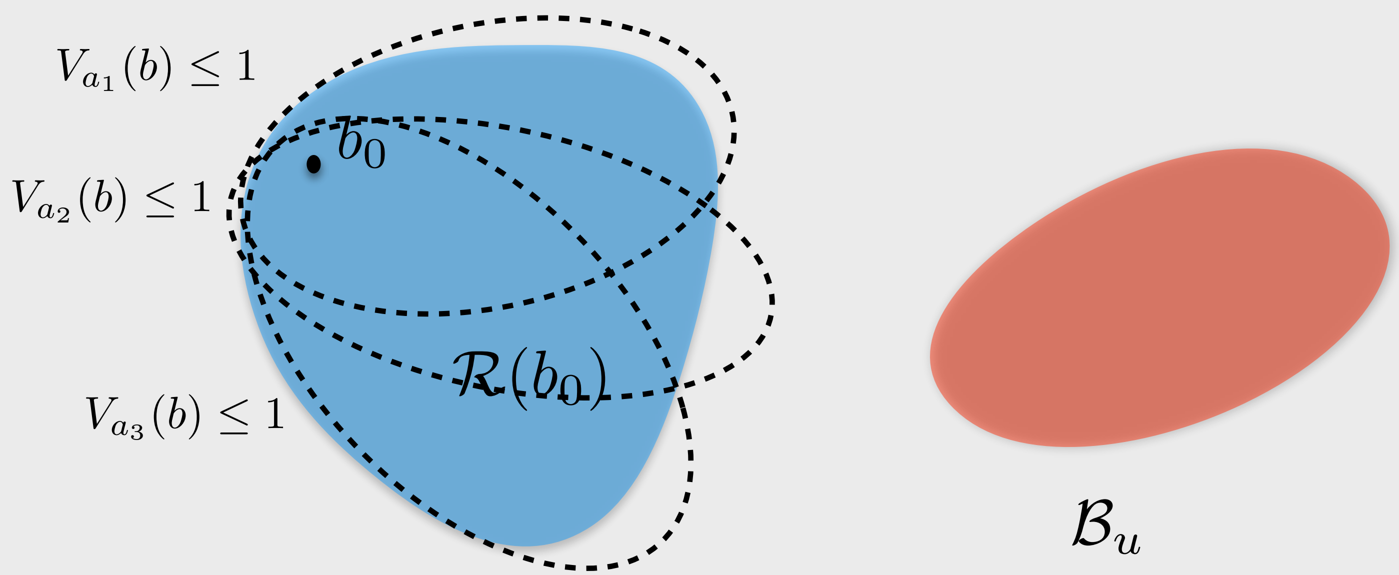

Over-approximation of the reachable belief set

If there exists a function *V* such that (with additional standard technical assumptions)

$$V\left(f_a\left(b_{t-1}, z\right)\right) - V\left(b_{t-1}\right) < 0 \qquad \forall a \in A,\ z \in Z,\ b_{t-1} \in \mathcal{V}$$

$$b_0 \in \mathcal{V} = \{b \in \mathcal{B} \mid V(b) \leq 1\}$$

Then, $\mathcal{R}(b_0) \subseteq \mathcal{V}$.

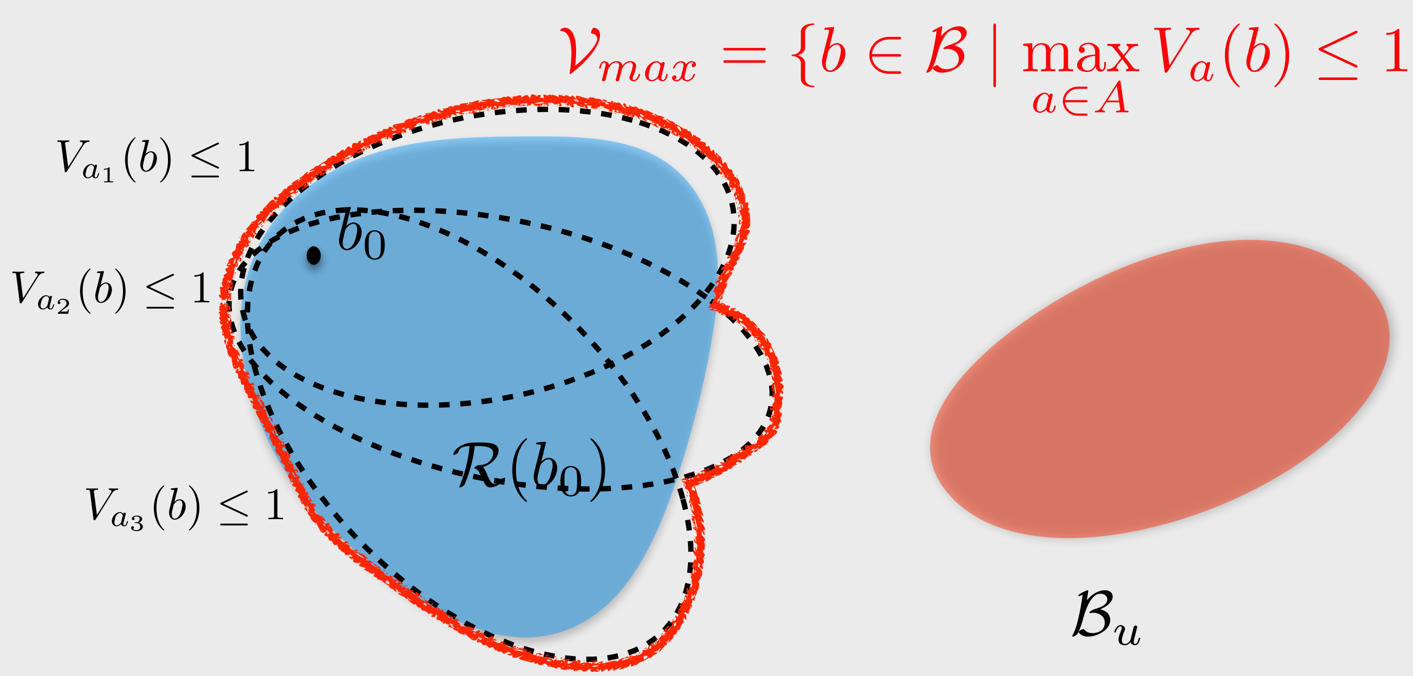

Example: over-approximation of the reachable belief set

If there exists a function V such that (with additional standard technical assumptions)

$$V\left(f_a\left(b_{t-1}, z\right)\right) - V\left(b_{t-1}\right) < 0 \qquad \forall a \in A,\ z \in Z,\ b_{t-1} \in \mathcal{V}$$

$$b_0 \in \mathcal{V} = \{b \in \mathcal{B} \mid V(b) \leq 1\}$$

Then, $\mathcal{R}(b_0) \subseteq \mathcal{V}$.

Three-state POMDP with two actions

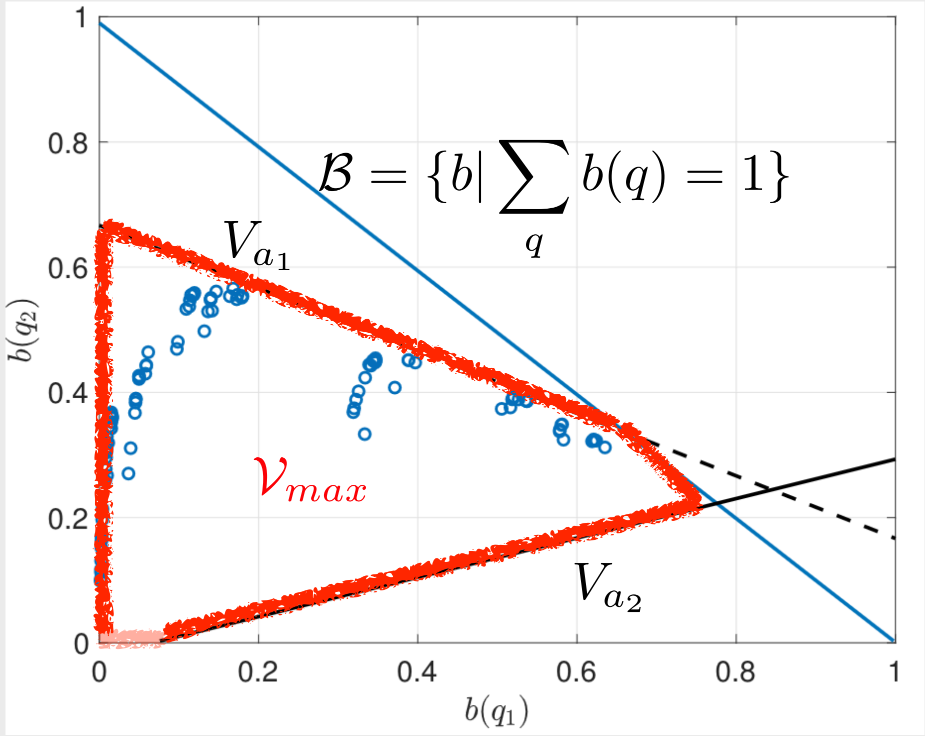

Piece-wise affine V

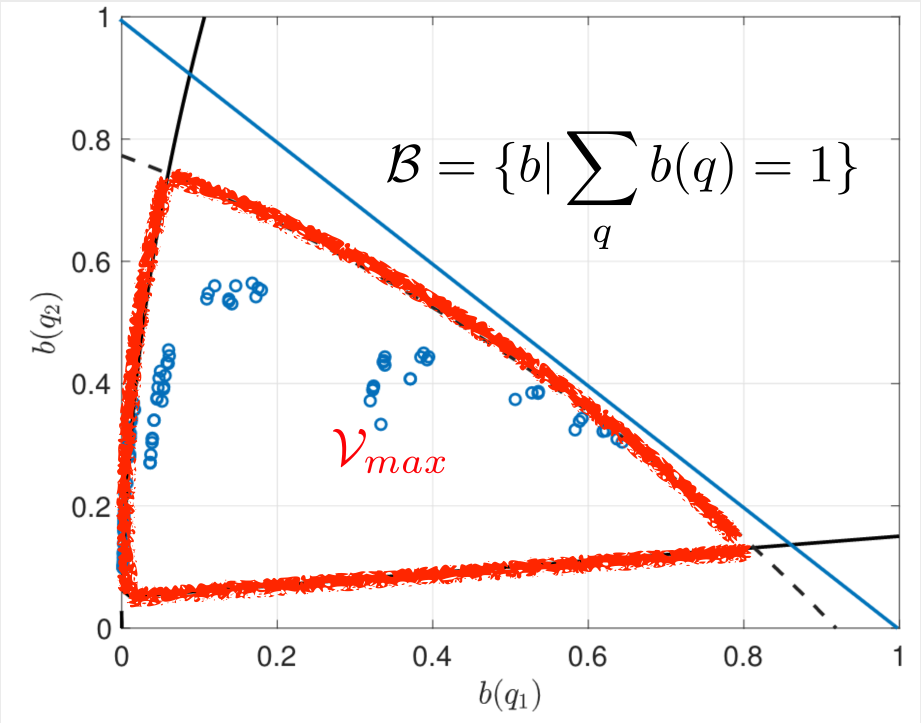

Piece-wise cubic polynomial V

Privacy verification using barrier certificates

If there exists a function *B* such that (with additional standard technical assumptions)

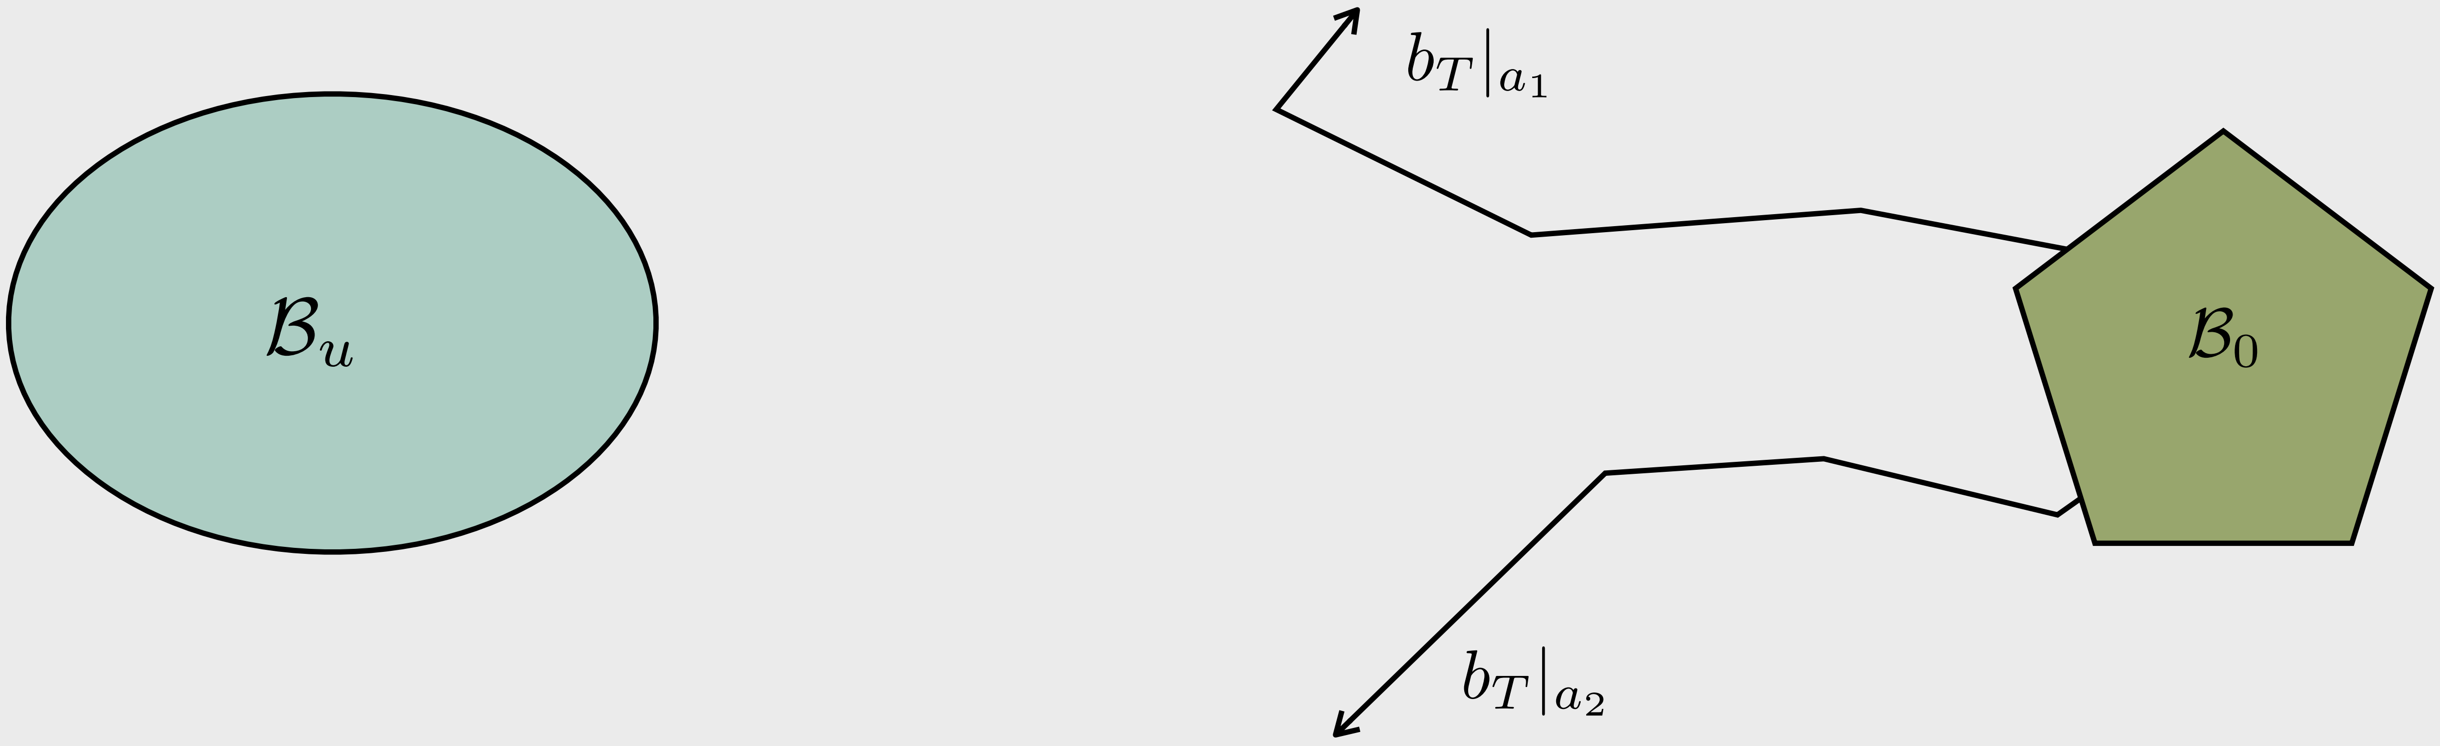

Privacy verification using barrier certificates

If there exists a function *B* such that (with additional standard technical assumptions)

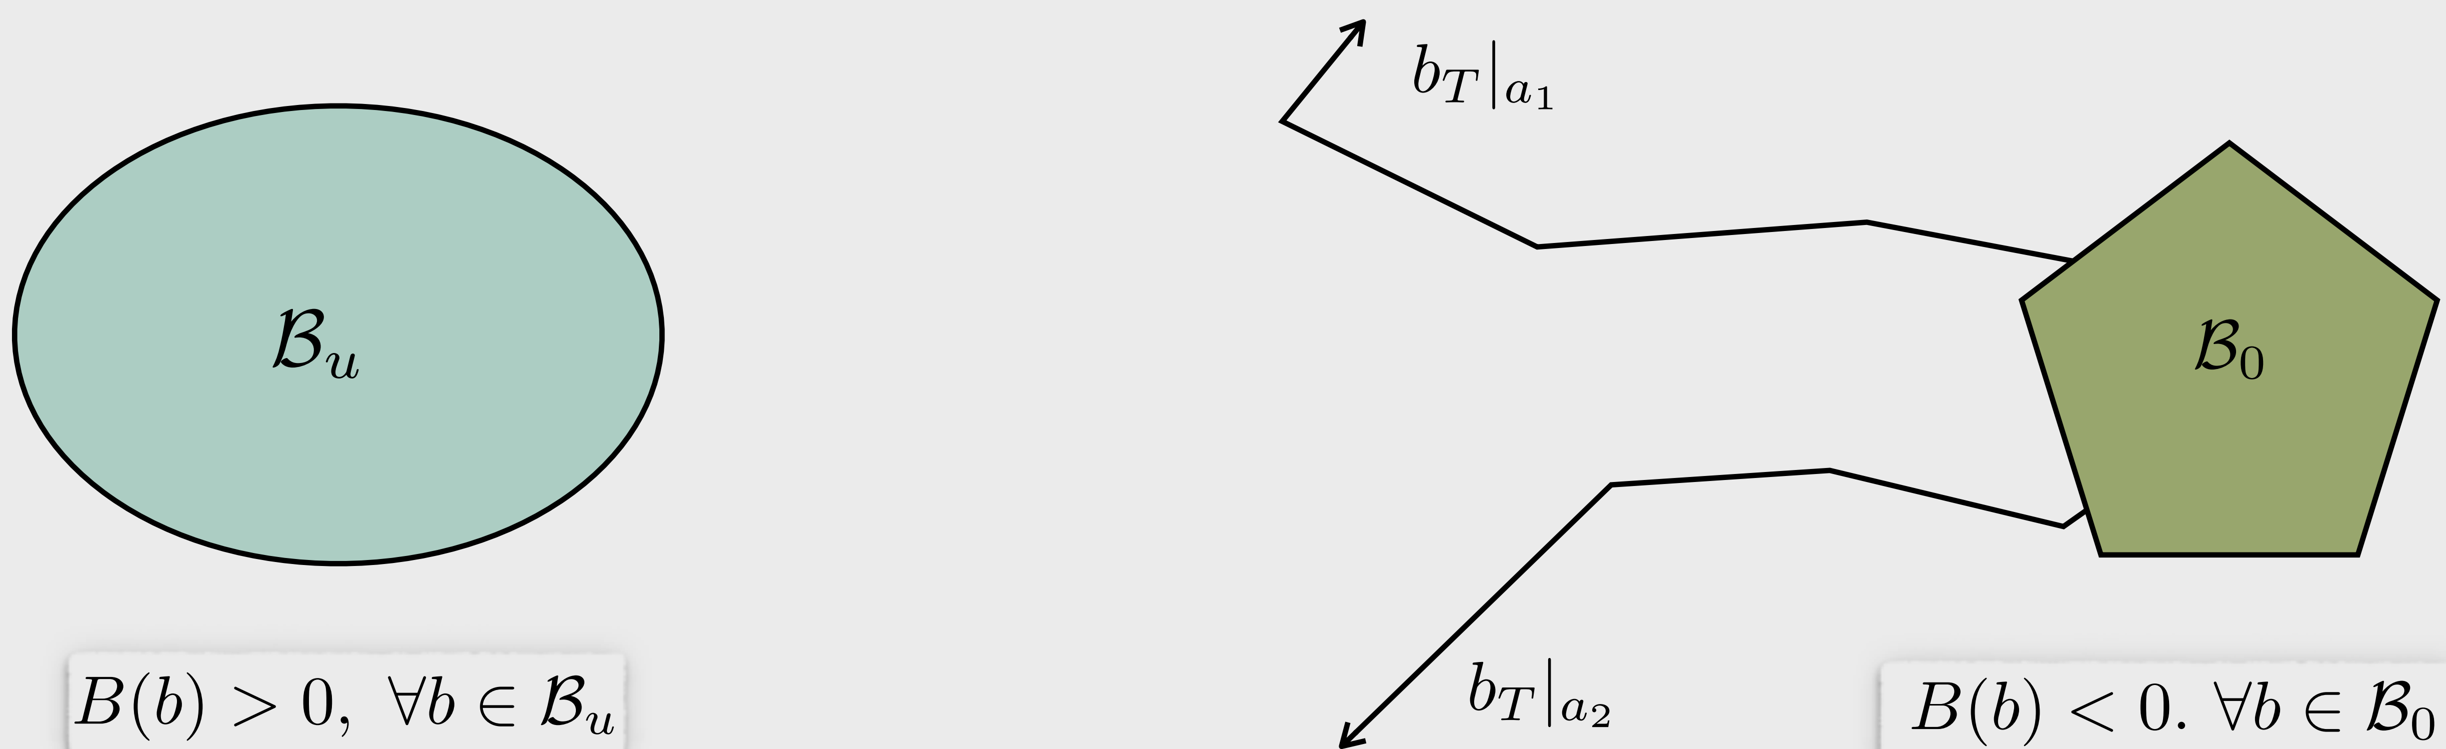

$$B(b) > 0, \ \forall b \in \mathcal{B}_u$$

$$B(b) < 0. \ \forall b \in \mathcal{B}_0$$

$$B(f_a(b_{t-1}, z)) - B(b_{t-1}, z) \leq 0 \ \forall a \in A, z \in Z, t \in \mathbb{Z}_{\geq 1}$$

Then $b_t \notin \mathcal{B}_u \ \forall t \in \mathbb{Z}_{\geq 1}$.

Privacy verification using barrier certificates

If there exists a function *B* such that (with additional standard technical assumptions)

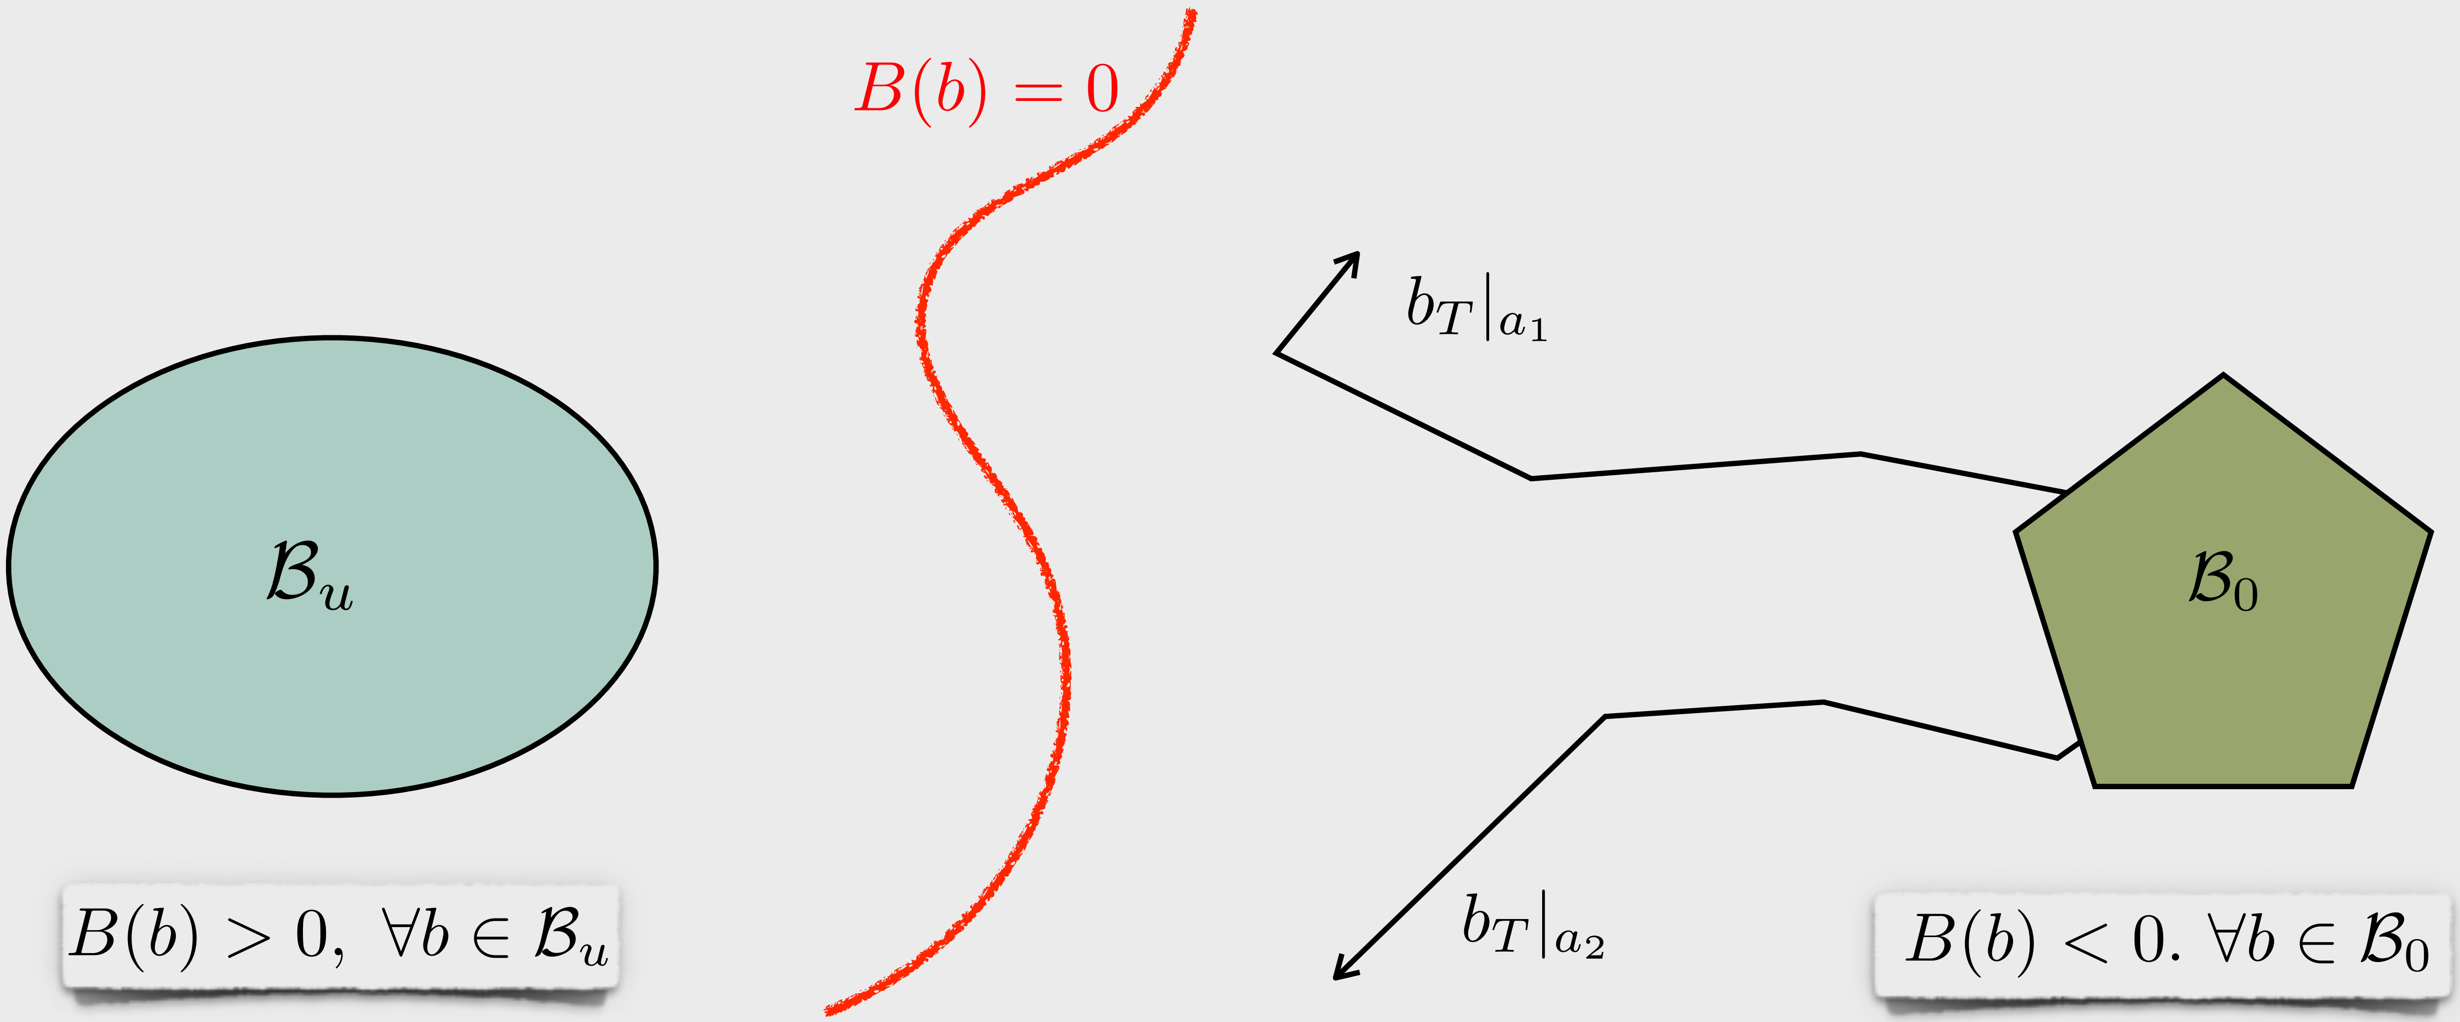

$B(b) > 0, \ \forall b \in \mathcal{B}_u$

$B(b) < 0. \ \forall b \in \mathcal{B}_0$

$$B(f_a(b_{t-1}, z)) - B(b_{t-1}, z) \leq 0 \ \forall a \in A, z \in Z, t \in \mathbb{Z}_{\geq 1}$$

Then $b_t \notin \mathcal{B}_u \ \forall t \in \mathbb{Z}_{\geq 1}$.

Privacy verification using barrier certificates

If there exists a function *B* such that (with additional standard technical assumptions)

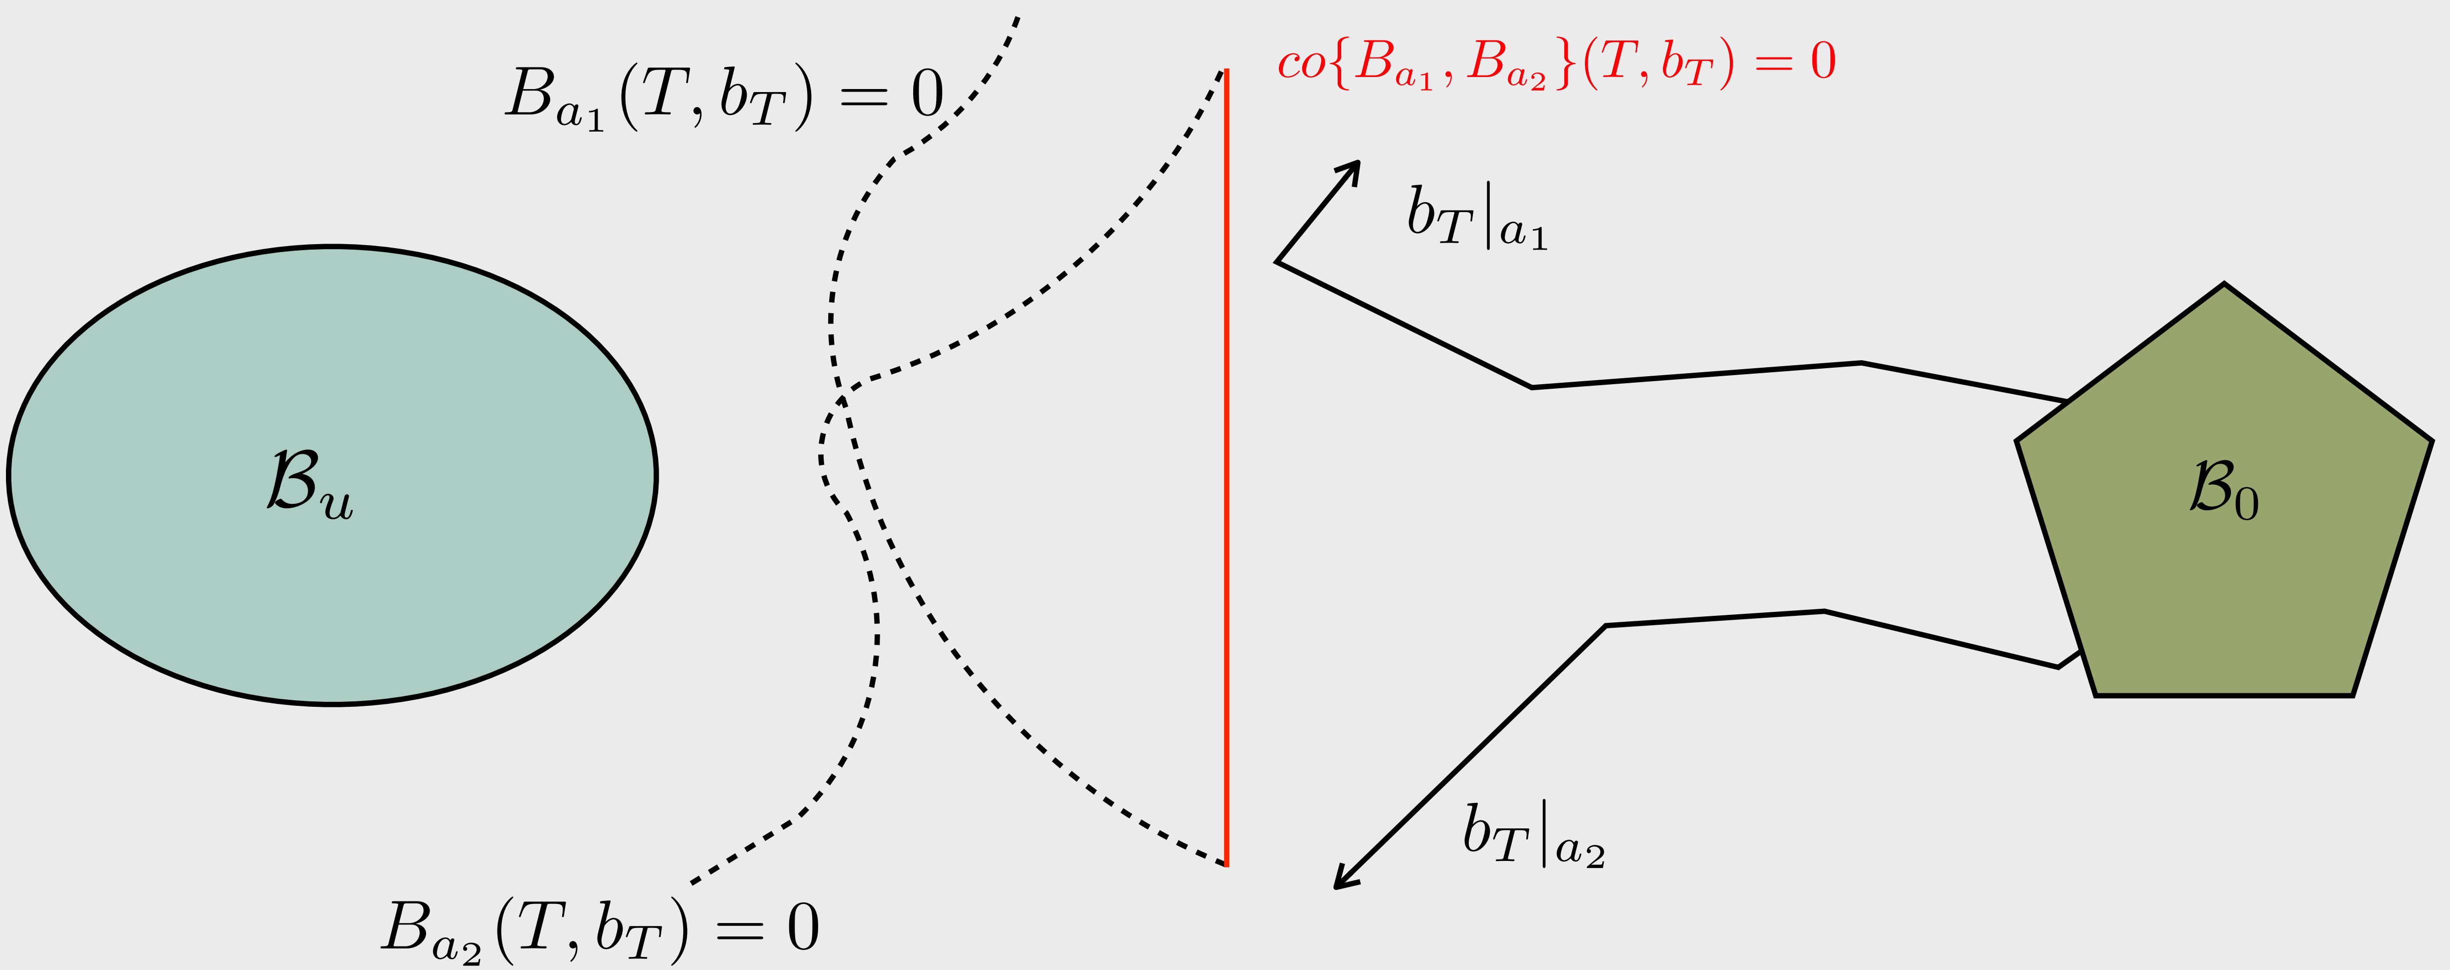

Then $b_t \notin \mathcal{B}_u \; \forall t \in \mathbb{Z}_{\geq 1}$.

How to search for V or B?

Useful features of verification problems in the belief space

- Belief dynamics are rational.
- Belief set is a unit simplex.

$$b_t(q) = \frac{O(q', a, z) \sum_{q \in Q} T(q, a, q') b_{t-1}(q)}{\sum_{q' \in Q} O(q', a, z) \sum_{q \in Q}, T(q, a, q') b_{t-1}(q)}$$

How to search for V or B?

Useful features of verification problems in the belief space

- Belief dynamics are rational.
- Belief set is a unit simplex.

$$b_t(q) = \frac{O(q', a, z) \sum_{q \in Q} T(q, a, q') b_{t-1}(q)}{\sum_{q' \in Q} O(q', a, z) \sum_{q \in Q'} T(q, a, q') b_{t-1}(q)}$$

Optimization-based search for V or B:

Restrict V or B to be polynomials of fixed, finite degree → Formulate the search as polynomial optimization → Relax to sum-of-squares optimization problems → Solve as semidefinite programming problems

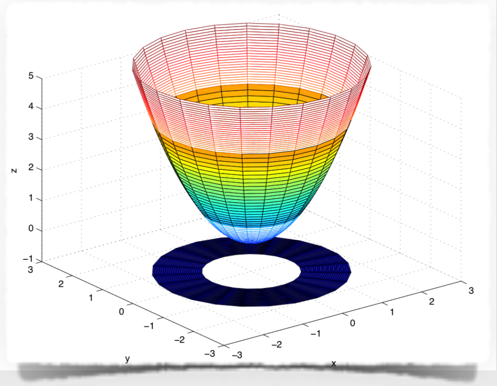

Example

Compute the minimum value γ such that

$$b_t(q_2) + b_t(q_3) \leq \gamma, \forall t$$

$$O_a(i,j) = O(q_i, a, z_j)$$

$$O_{a_1} = \begin{bmatrix} 0.7, & 0.3 \\ 0.5, & 0.5 \\ 0.8, & 0.2 \end{bmatrix}, O_{a_2} = \begin{bmatrix} 0.8, & 0.2 \\ 0.6, & 0.4 \\ 0.2, & 0.8 \end{bmatrix}.$$

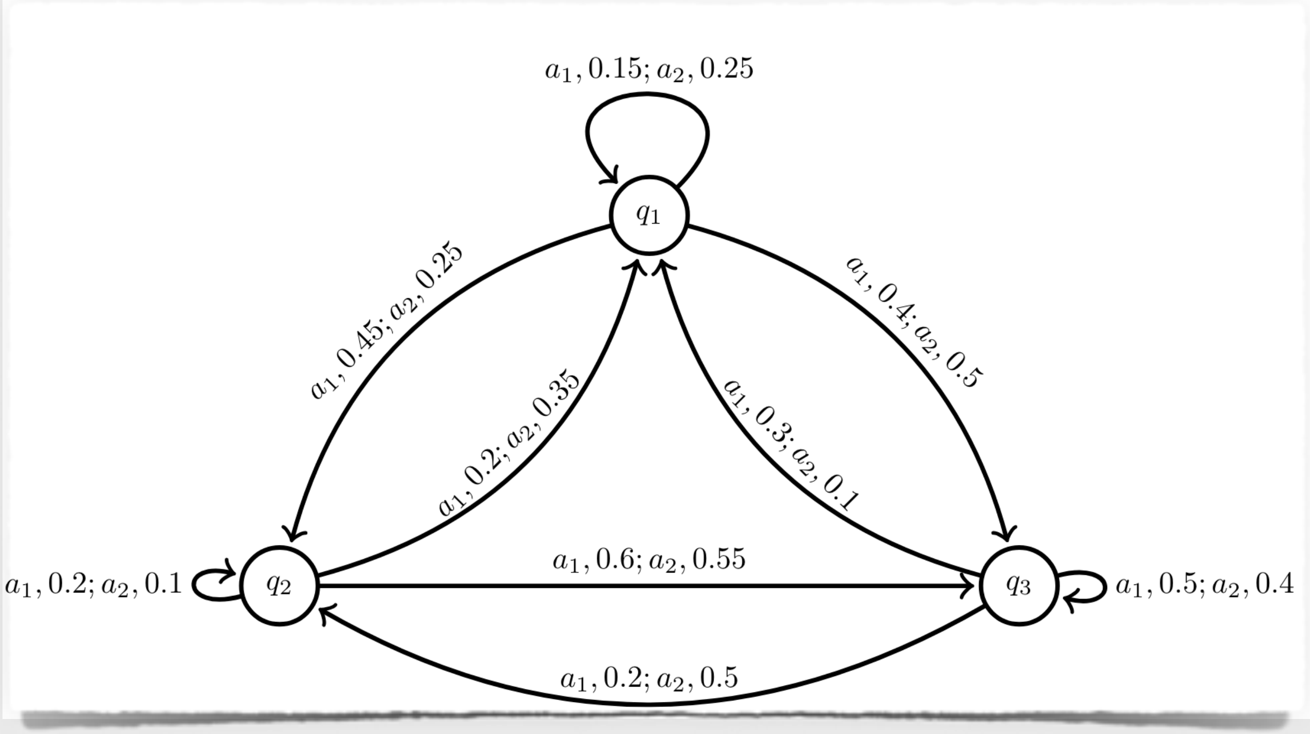

d (degree of B)	2	4	6	8	10
γ^*	0.93	0.88	0.80	0.74	0.69
Computation Time (s)	5.38	8.37	12.03	18.42	27.09

$$\begin{aligned} B(b) =& 0.1629 b(q_1)^2 - 3.9382 b(q_2)^2 + 09280 b(q_3)^2 \\ &- 0.0297 b(q_1) b(q_2) - 4.4451 b(q_2) b(q_3) - 0.0027 b(q_1) \\ &- 2.0452 b(q_2) + 9.2633. \end{aligned}$$

Wrap-up

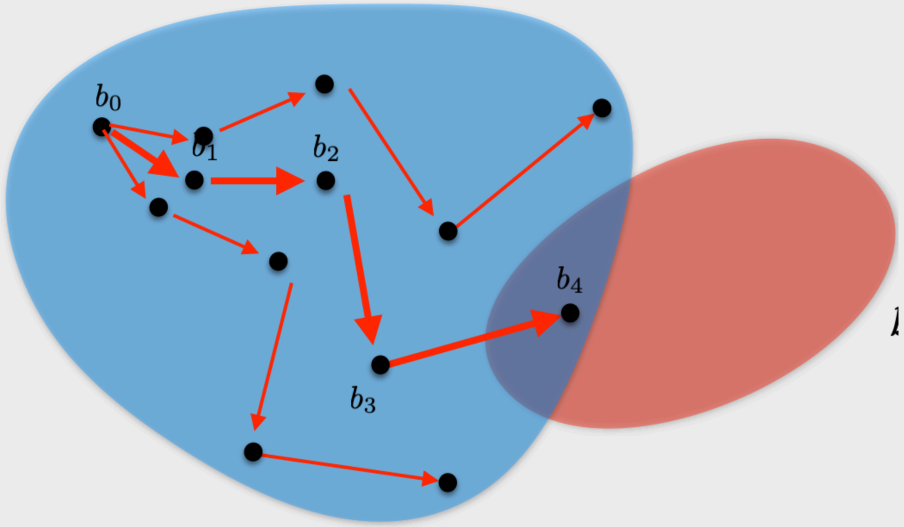

Summary: Verification in belief space as search for algebraic certificates for hybrid system dynamics

Next:
- Verification → Synthesis
- (PO)MDPs → Partial-information, two-player games

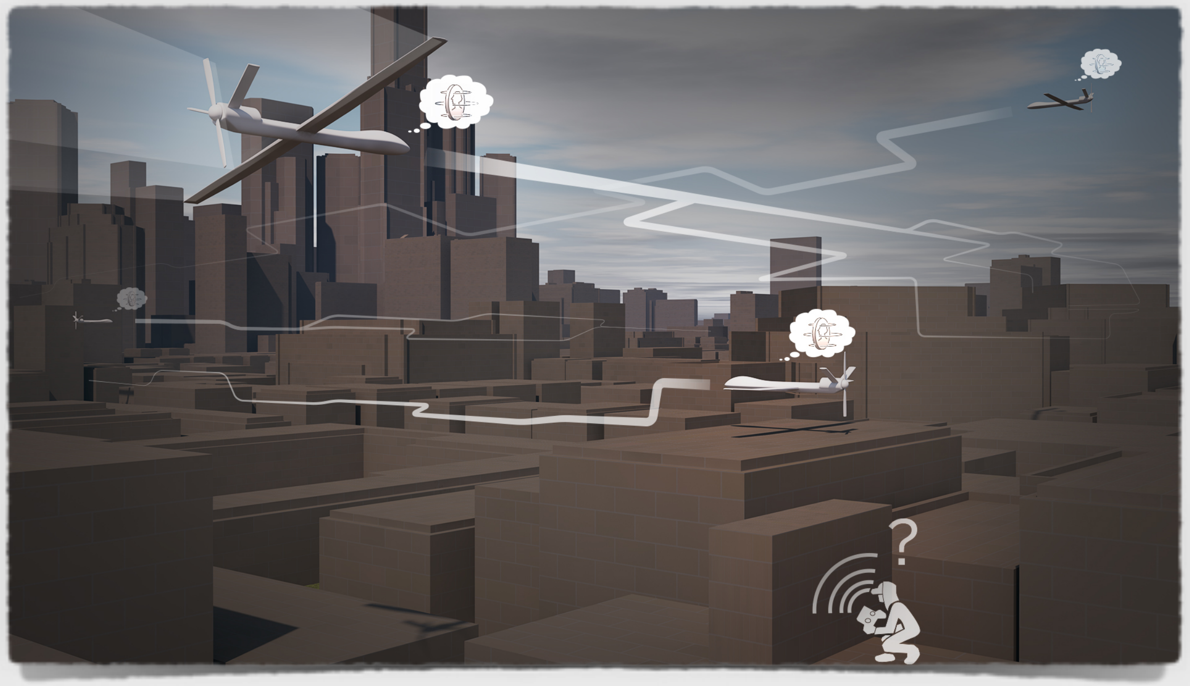

Protecting the integrity of mission-critical information

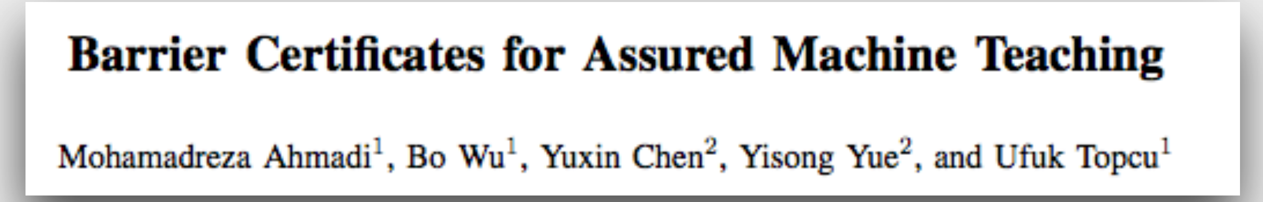

Barrier Certificates for Assured Machine Teaching

Mohamadreza Ahmadi[1], Bo Wu[1], Yuxin Chen[2], Yisong Yue[2], and Ufuk Topcu[1]

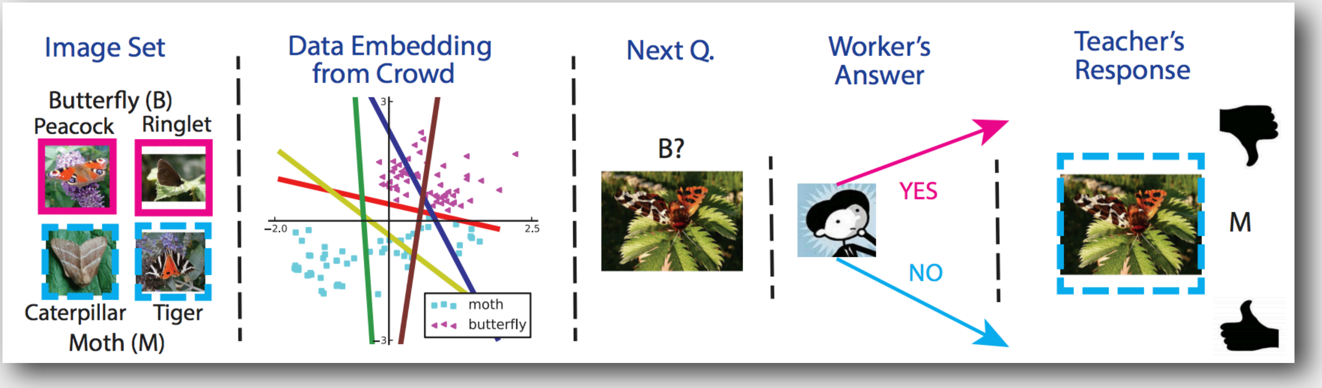

Performance prediction in "machine teaching"