Multi-agent Coordination in Uncontrolled Airspaces

Sarah H.Q. Li
Control, Coordination, and Competition under Uncertainty (C3U) Lab
Aerospace Engineering, Georgia Institute of Technology

Georgia Tech

Daniel Guggenheim School of Aerospace Engineering

C3U

Agenda

- Coordination in Air Traffic
- Audio-based Coordination in Uncontrolled Airspaces
- Commitment-breaking Stackelberg Dynamic Games
- Insights Gained from Breaking Commitments

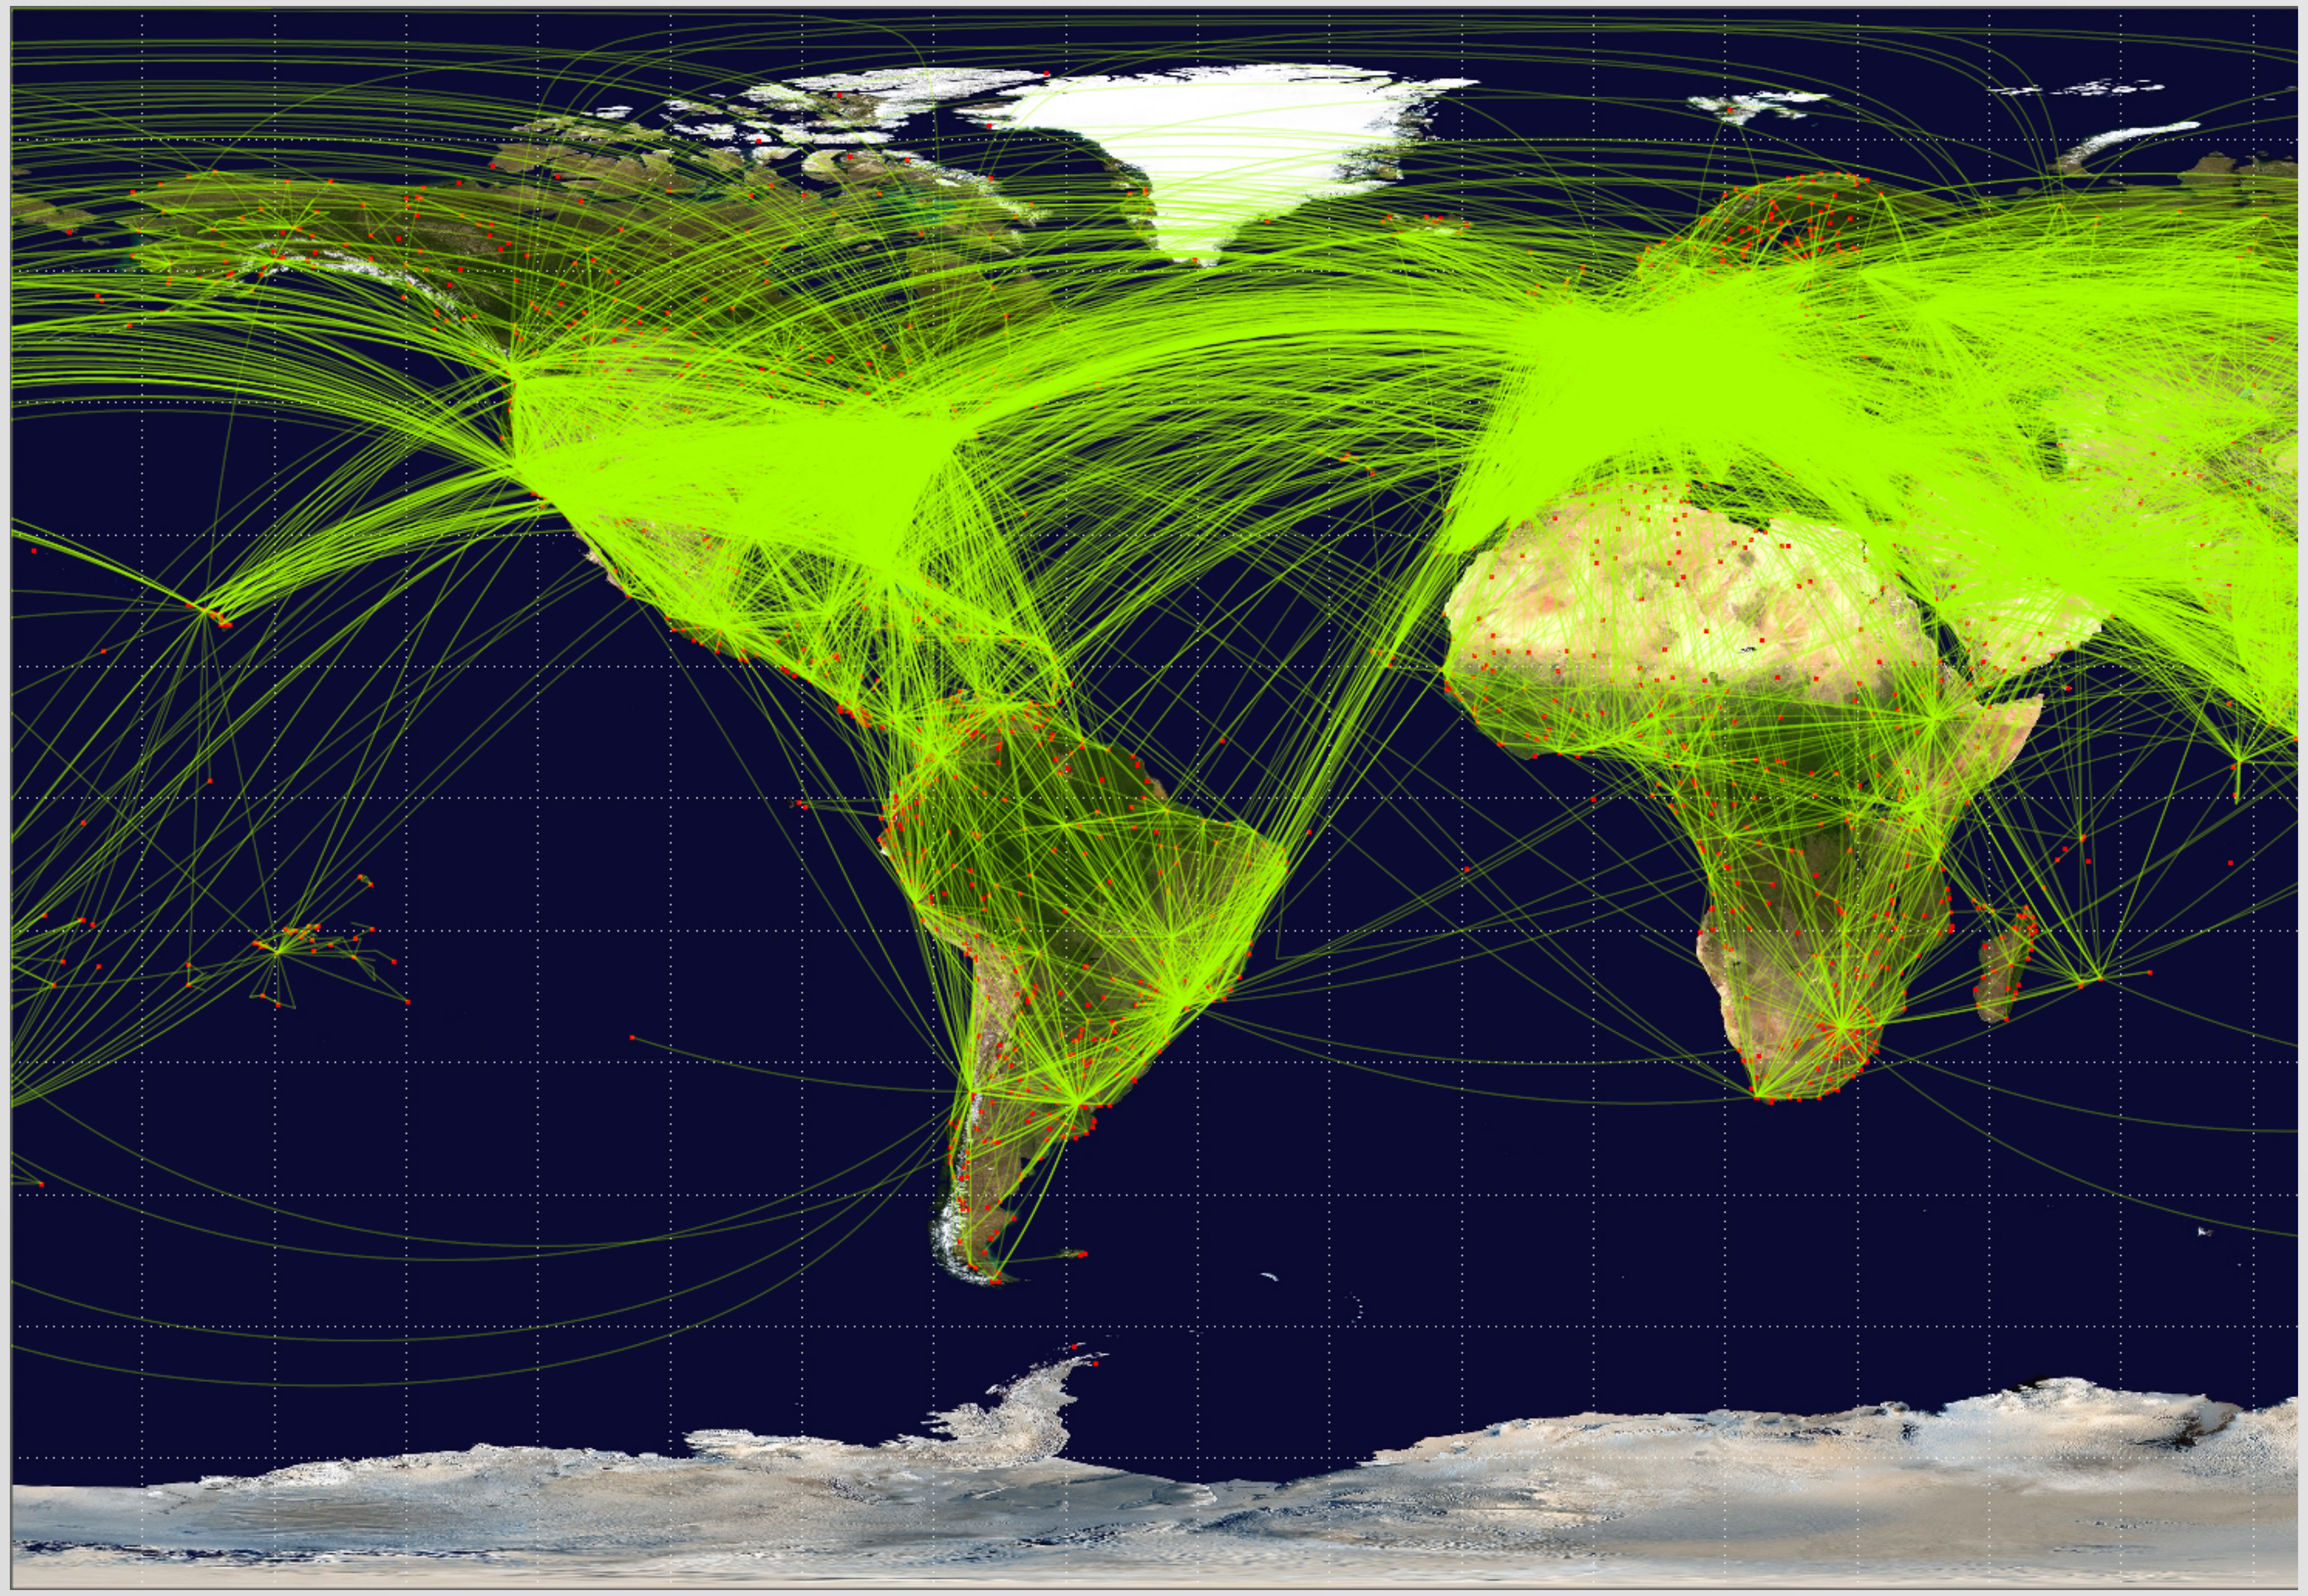

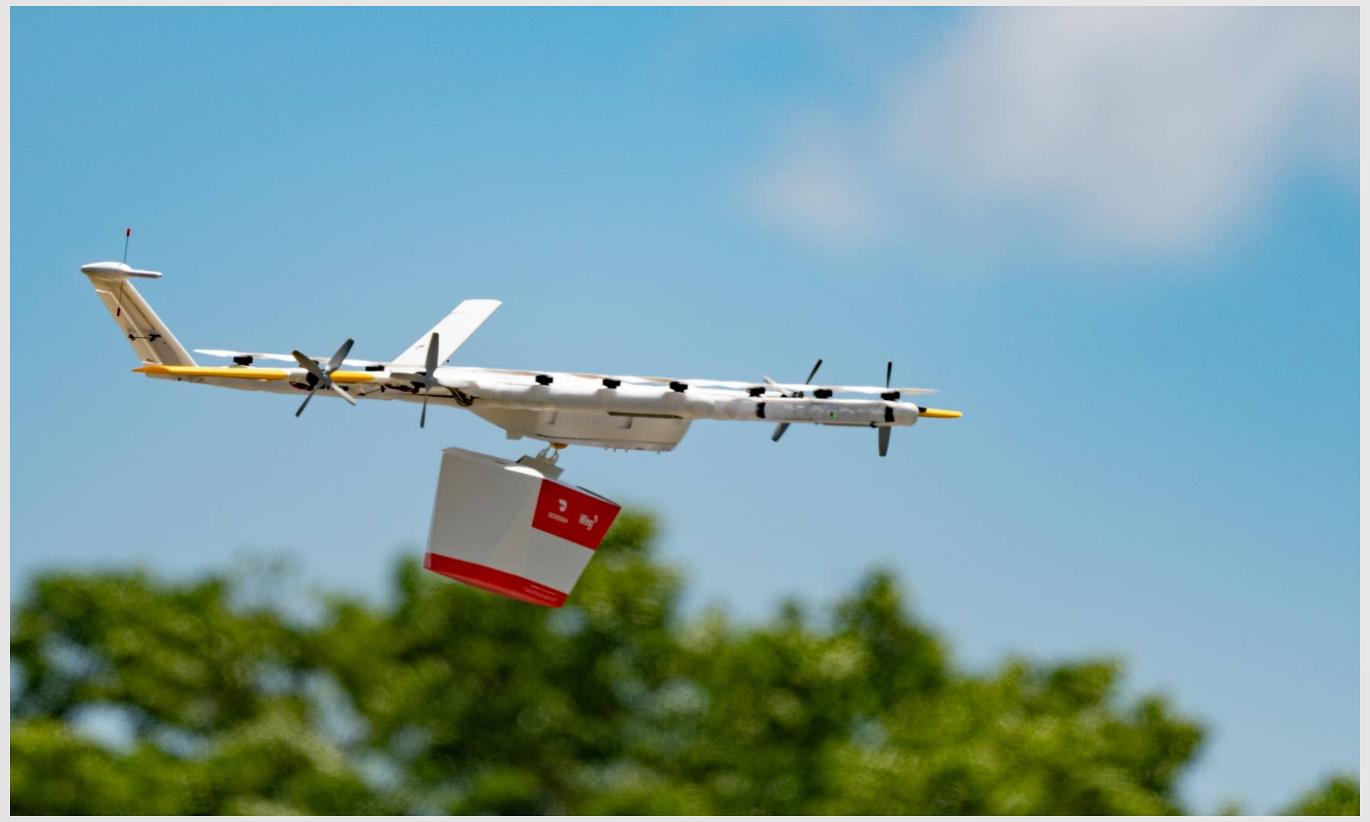

Interaction in Multi-agent Aerial Systems

Centralized coordination through routing and traffic management
Decentralized coordination through air traffic controllers
Real-time coordination via intent exchange (uncontrolled airspaces)
All strictly protocoled for different scenarios, timescales, and safety needs

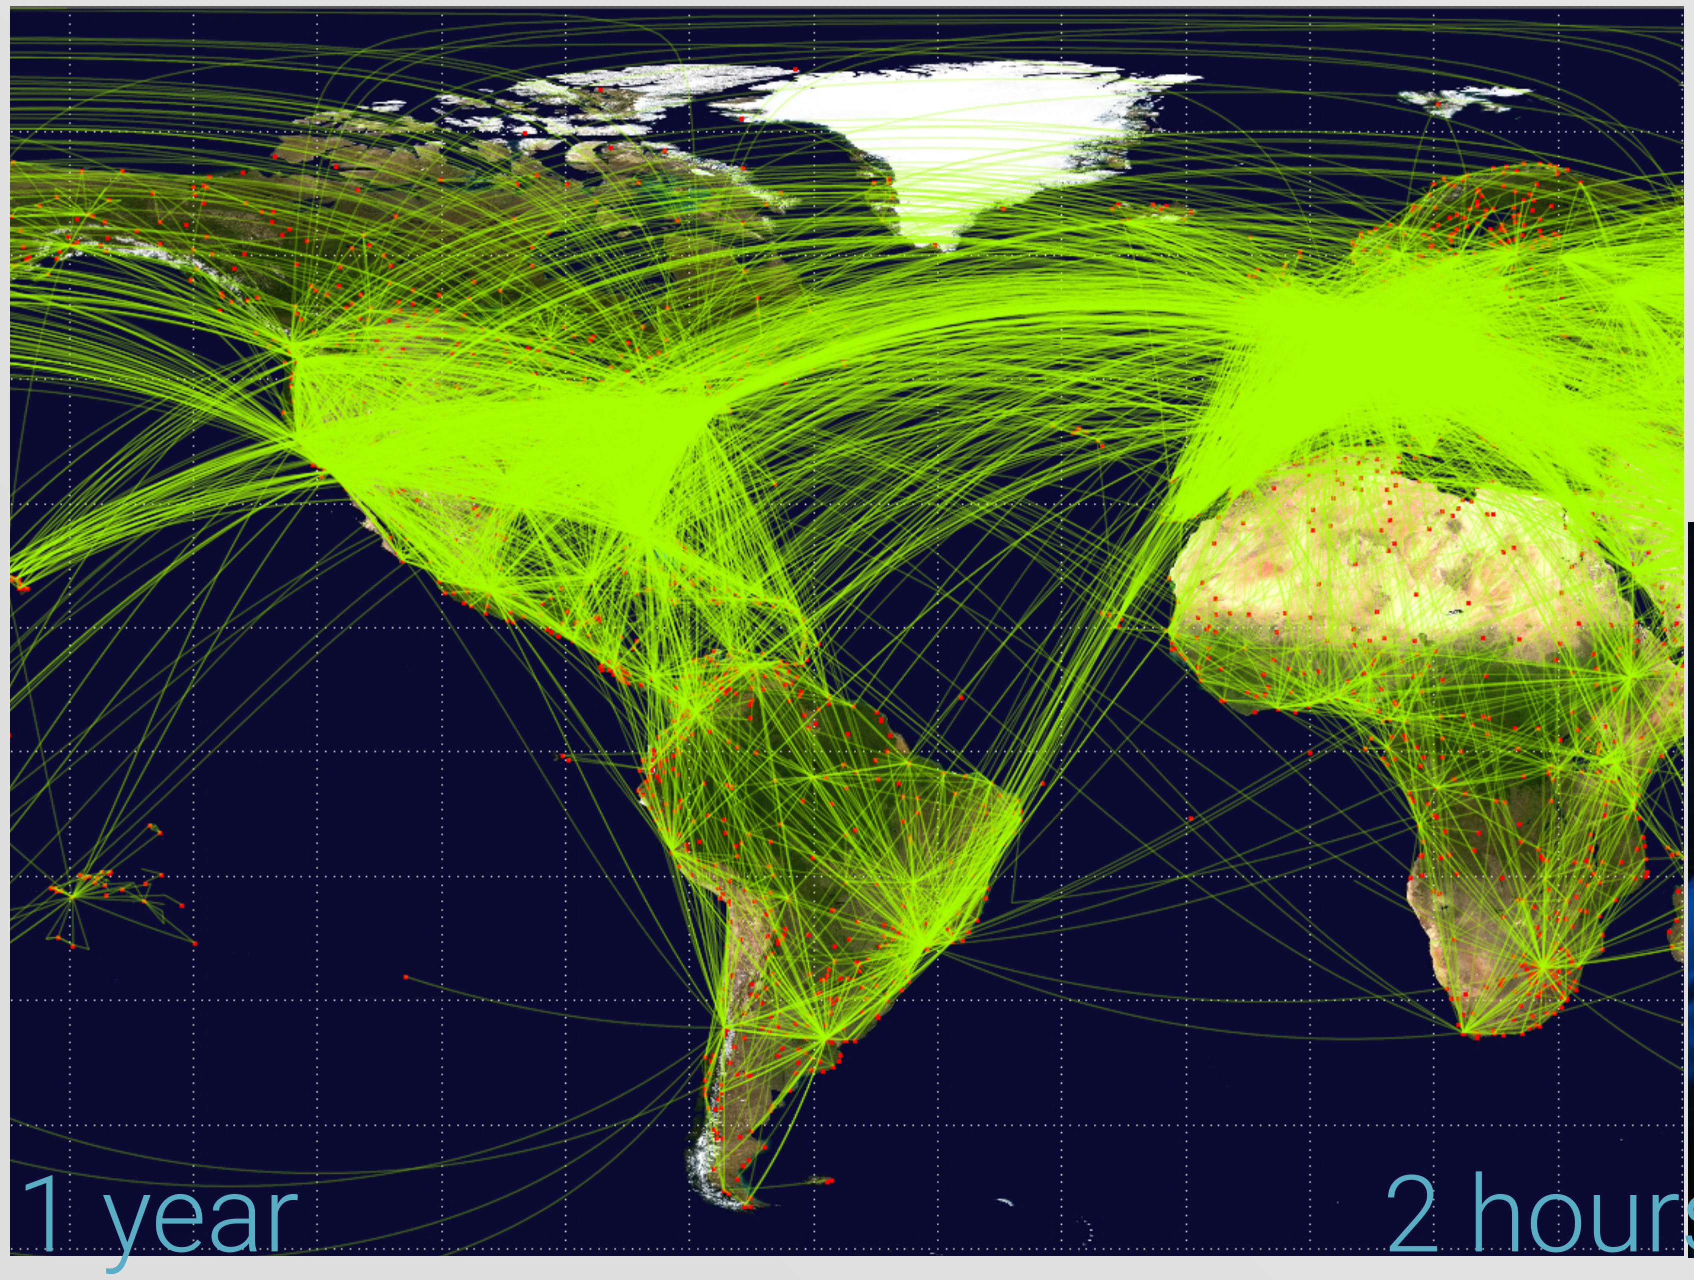

1 year

2 hours

Right upwind runway 8

10 minutes

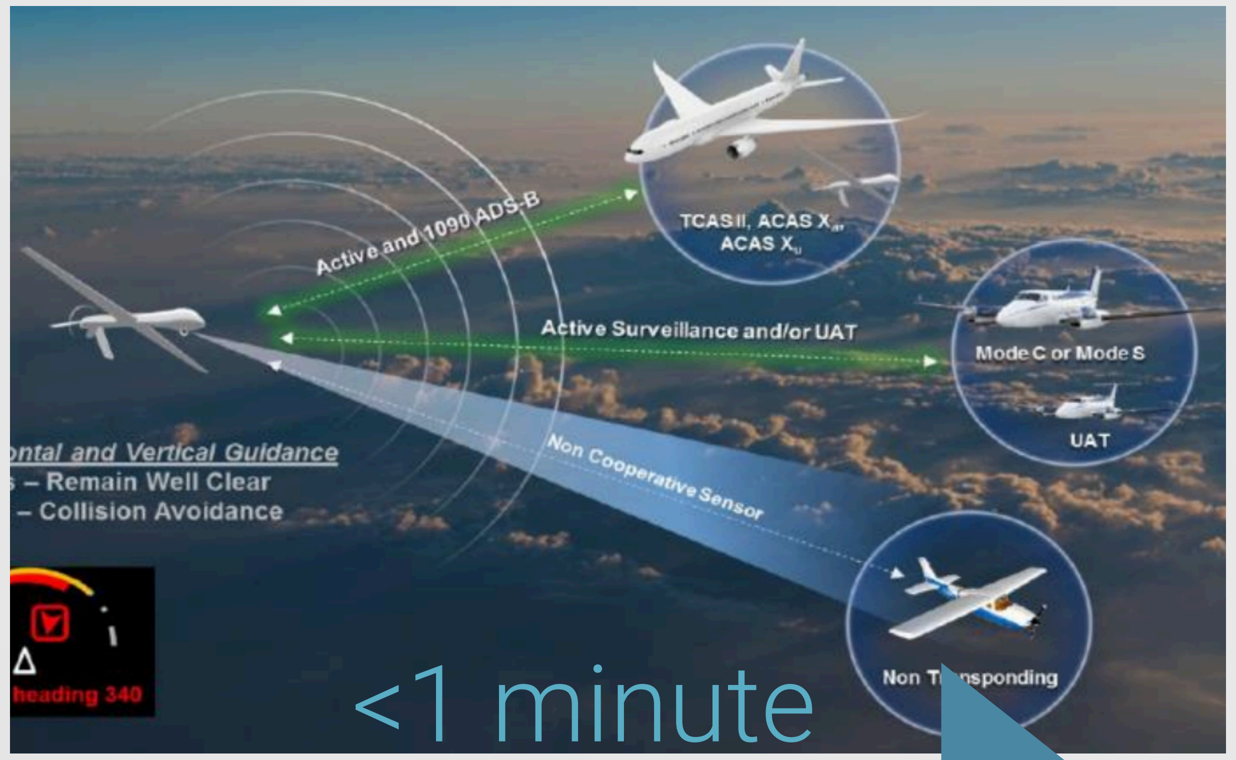

<1 minute

Centralized coordination,
Local policy

Local coordination,
Local policy

No coordination
Local policy

Coordination Objectives at Different Timescales

Longer timescales focus on efficiency
Tactical time scale (2 hours) focus on safety + efficiency
Sub-minute scale purely maximize safety

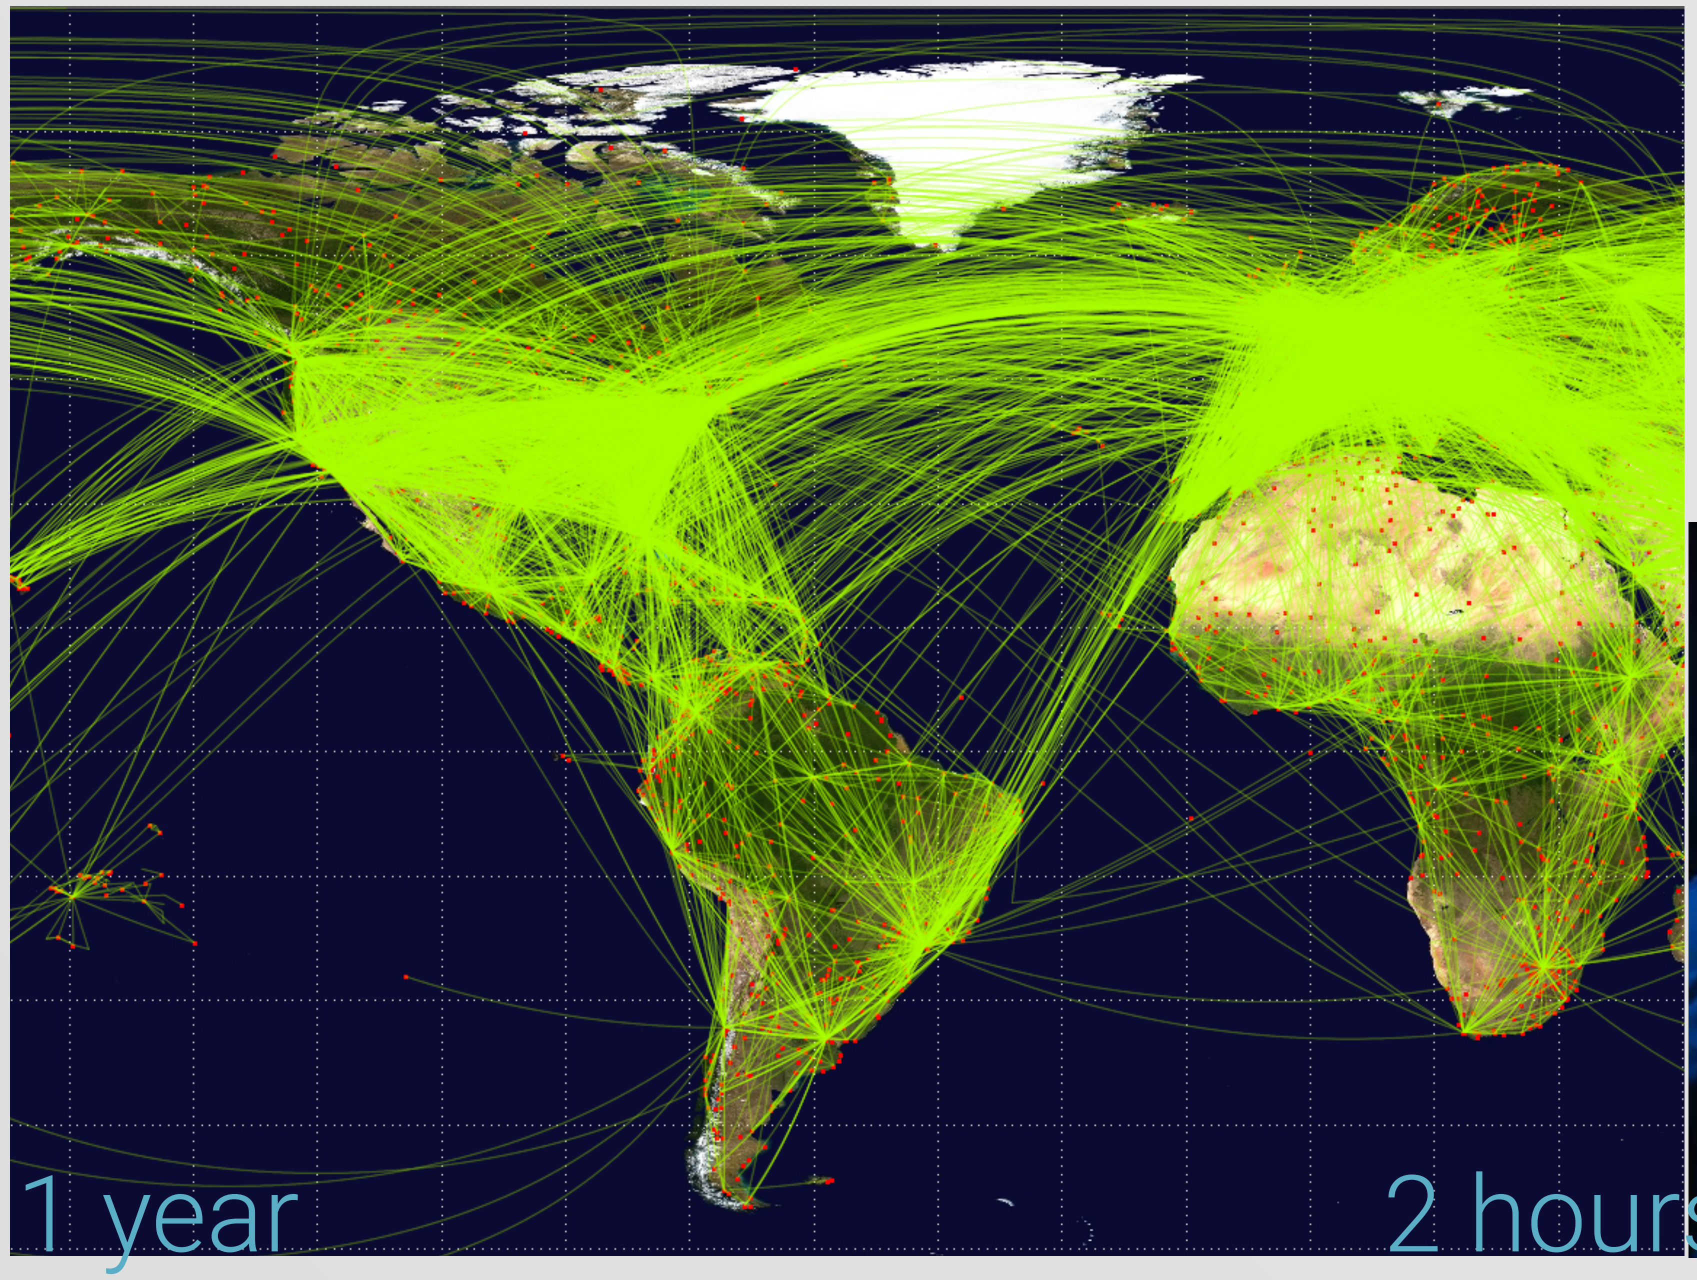

Right upwind
runway 8

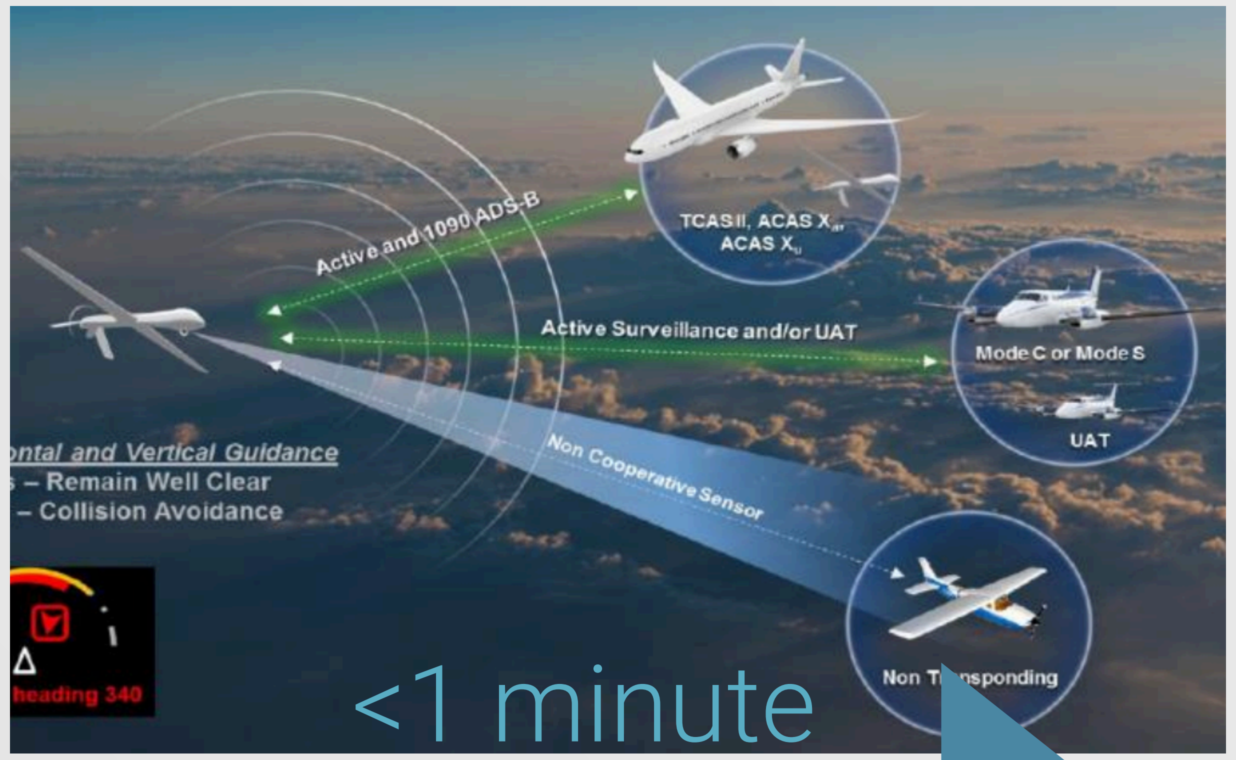

1 year — 2 hours — 10 minutes — <1 minute

Centralized coordination,
Local policy

Local coordination,
Local policy

No coordination
Individual policy

Coordination at different time scales

Longer timescales focus on efficiency
Tactical time scale (2 hours) focus on safety + efficiency
Sub-minute scale purely maximize safety
Aircraft collision avoidance is predominantly achieved without close-range cues

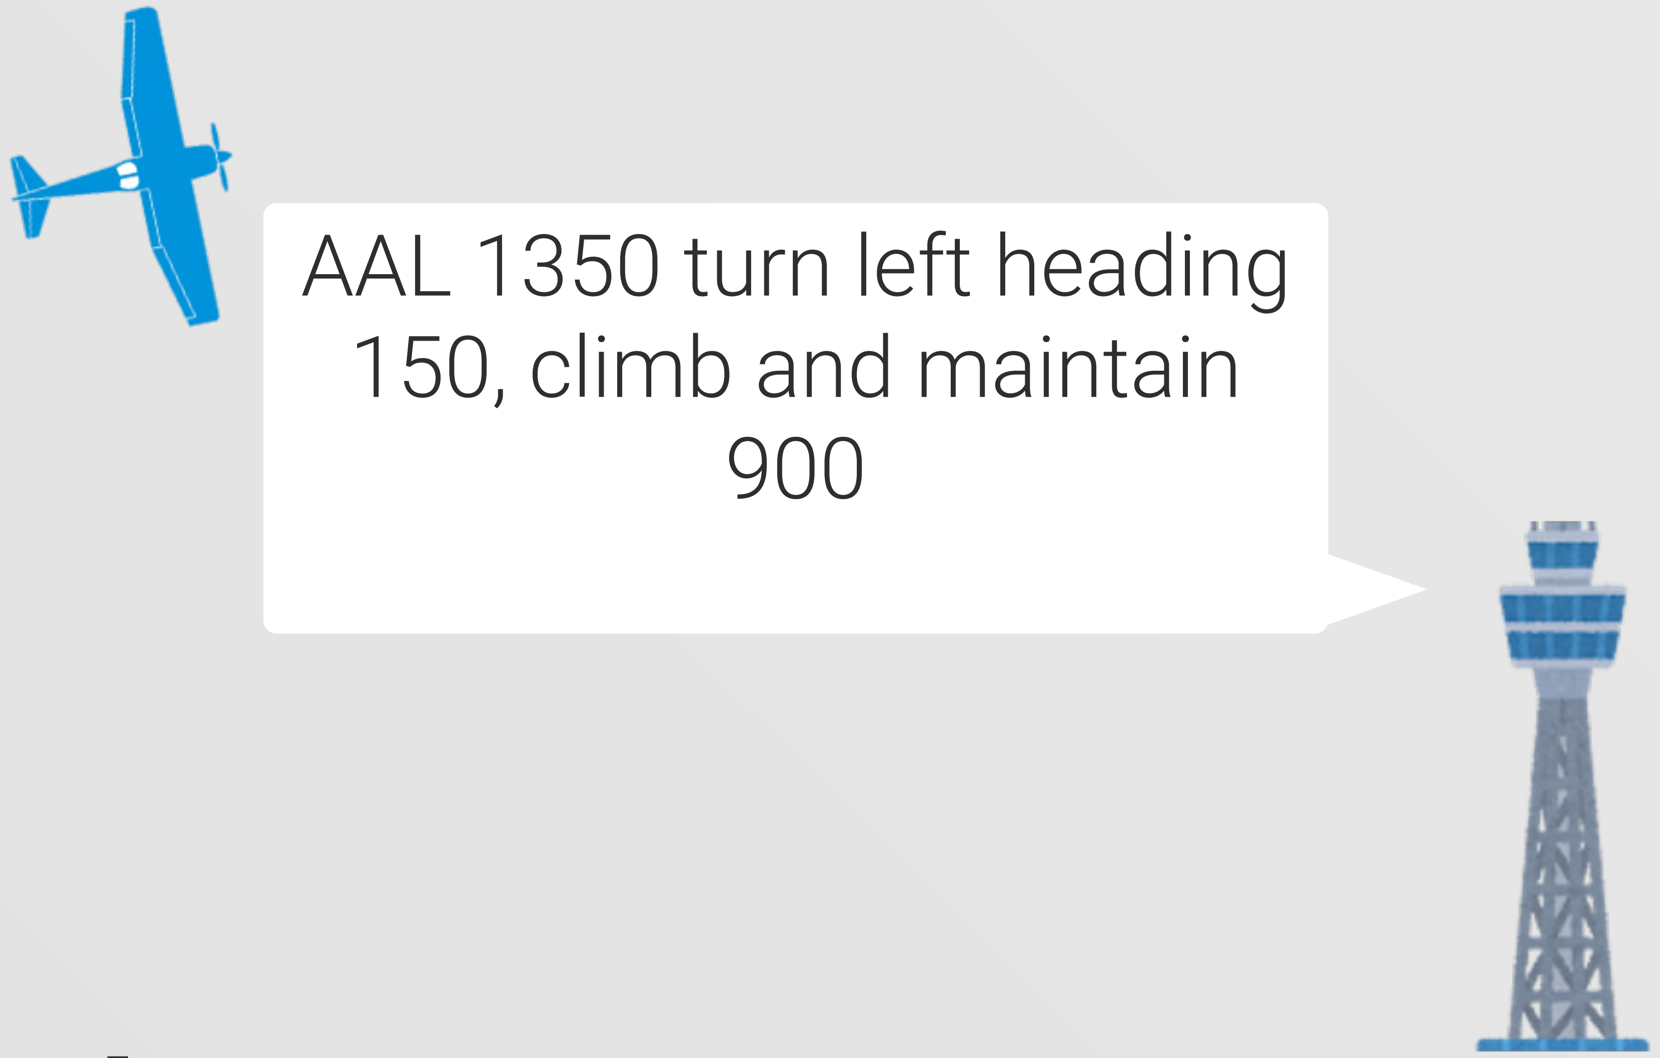

Aircraft Coordination in Controlled Airspaces

Aircraft follows local air traffic controller’s **audio** instructions to avoid conflict with other vehicles

Source: Reliable Robotics

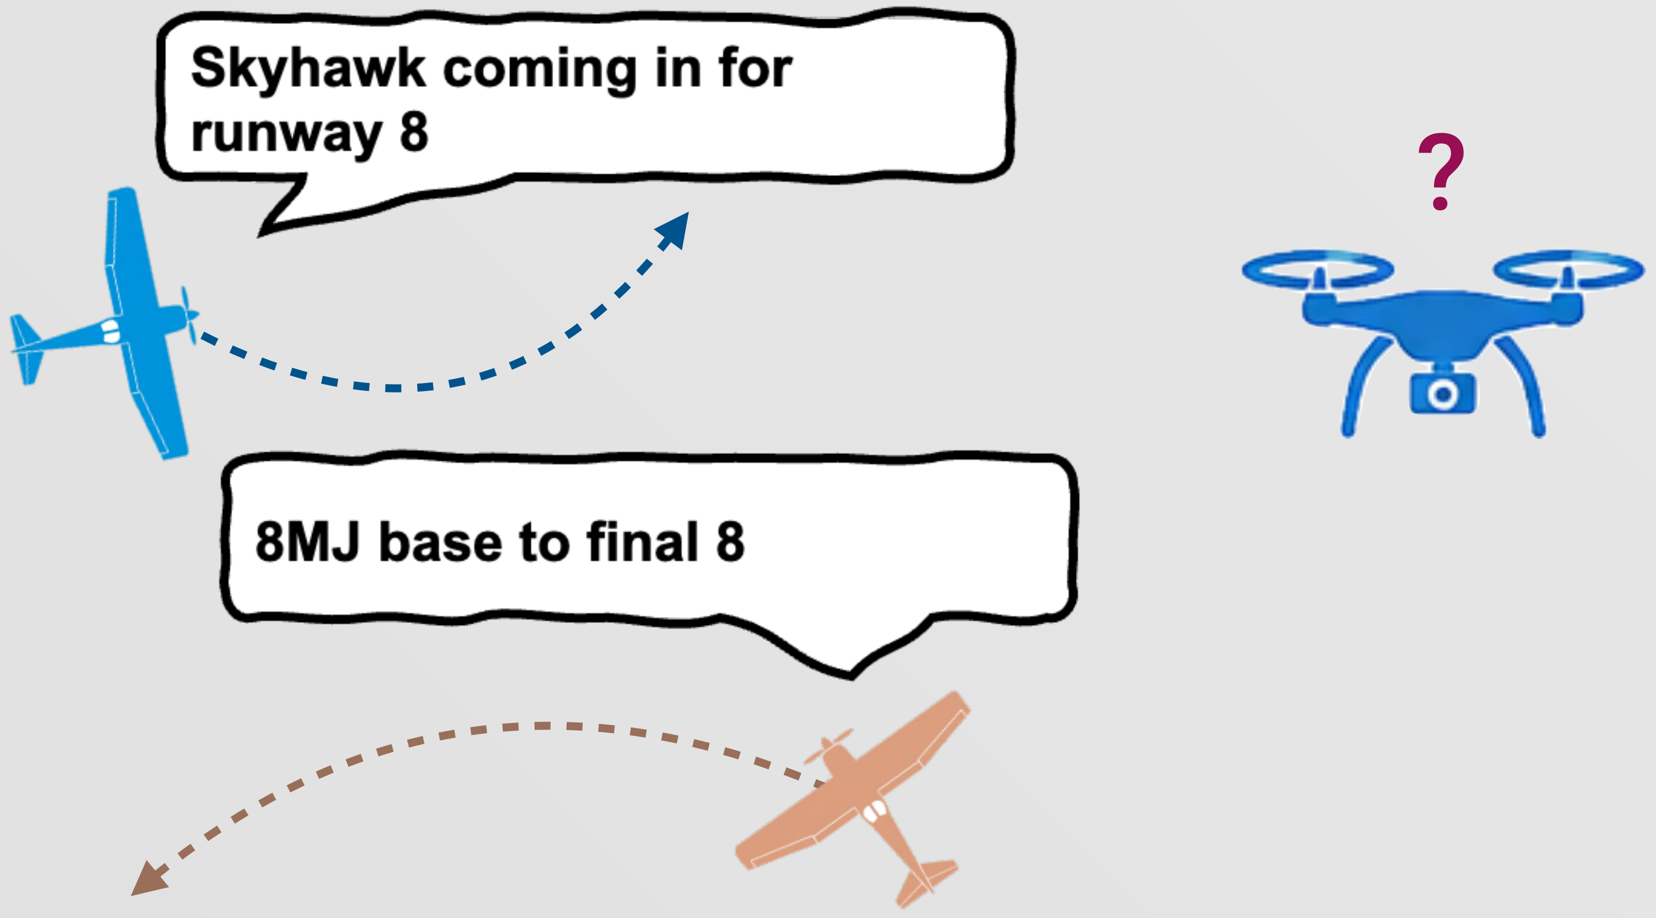

Majority of airports in the US do not have control towers

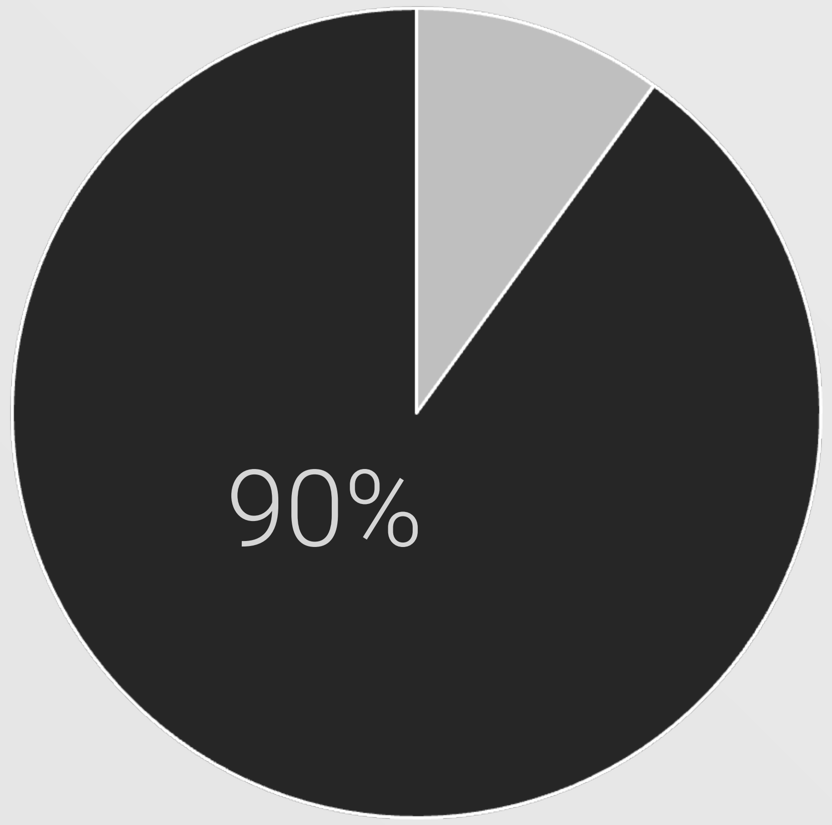

Aircraft Coordination in Uncontrolled Airspaces

Instead of **visual cues**, human pilots rely on **audio and language** to infer the relative position and intent of other aircraft in its vicinity

Can autonomous vehicles do the same?

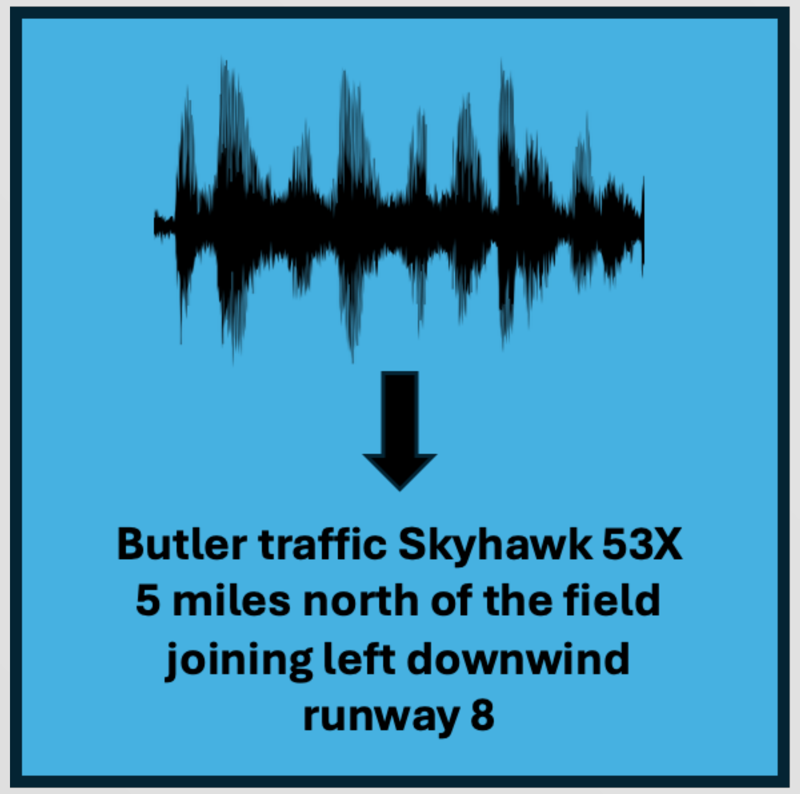

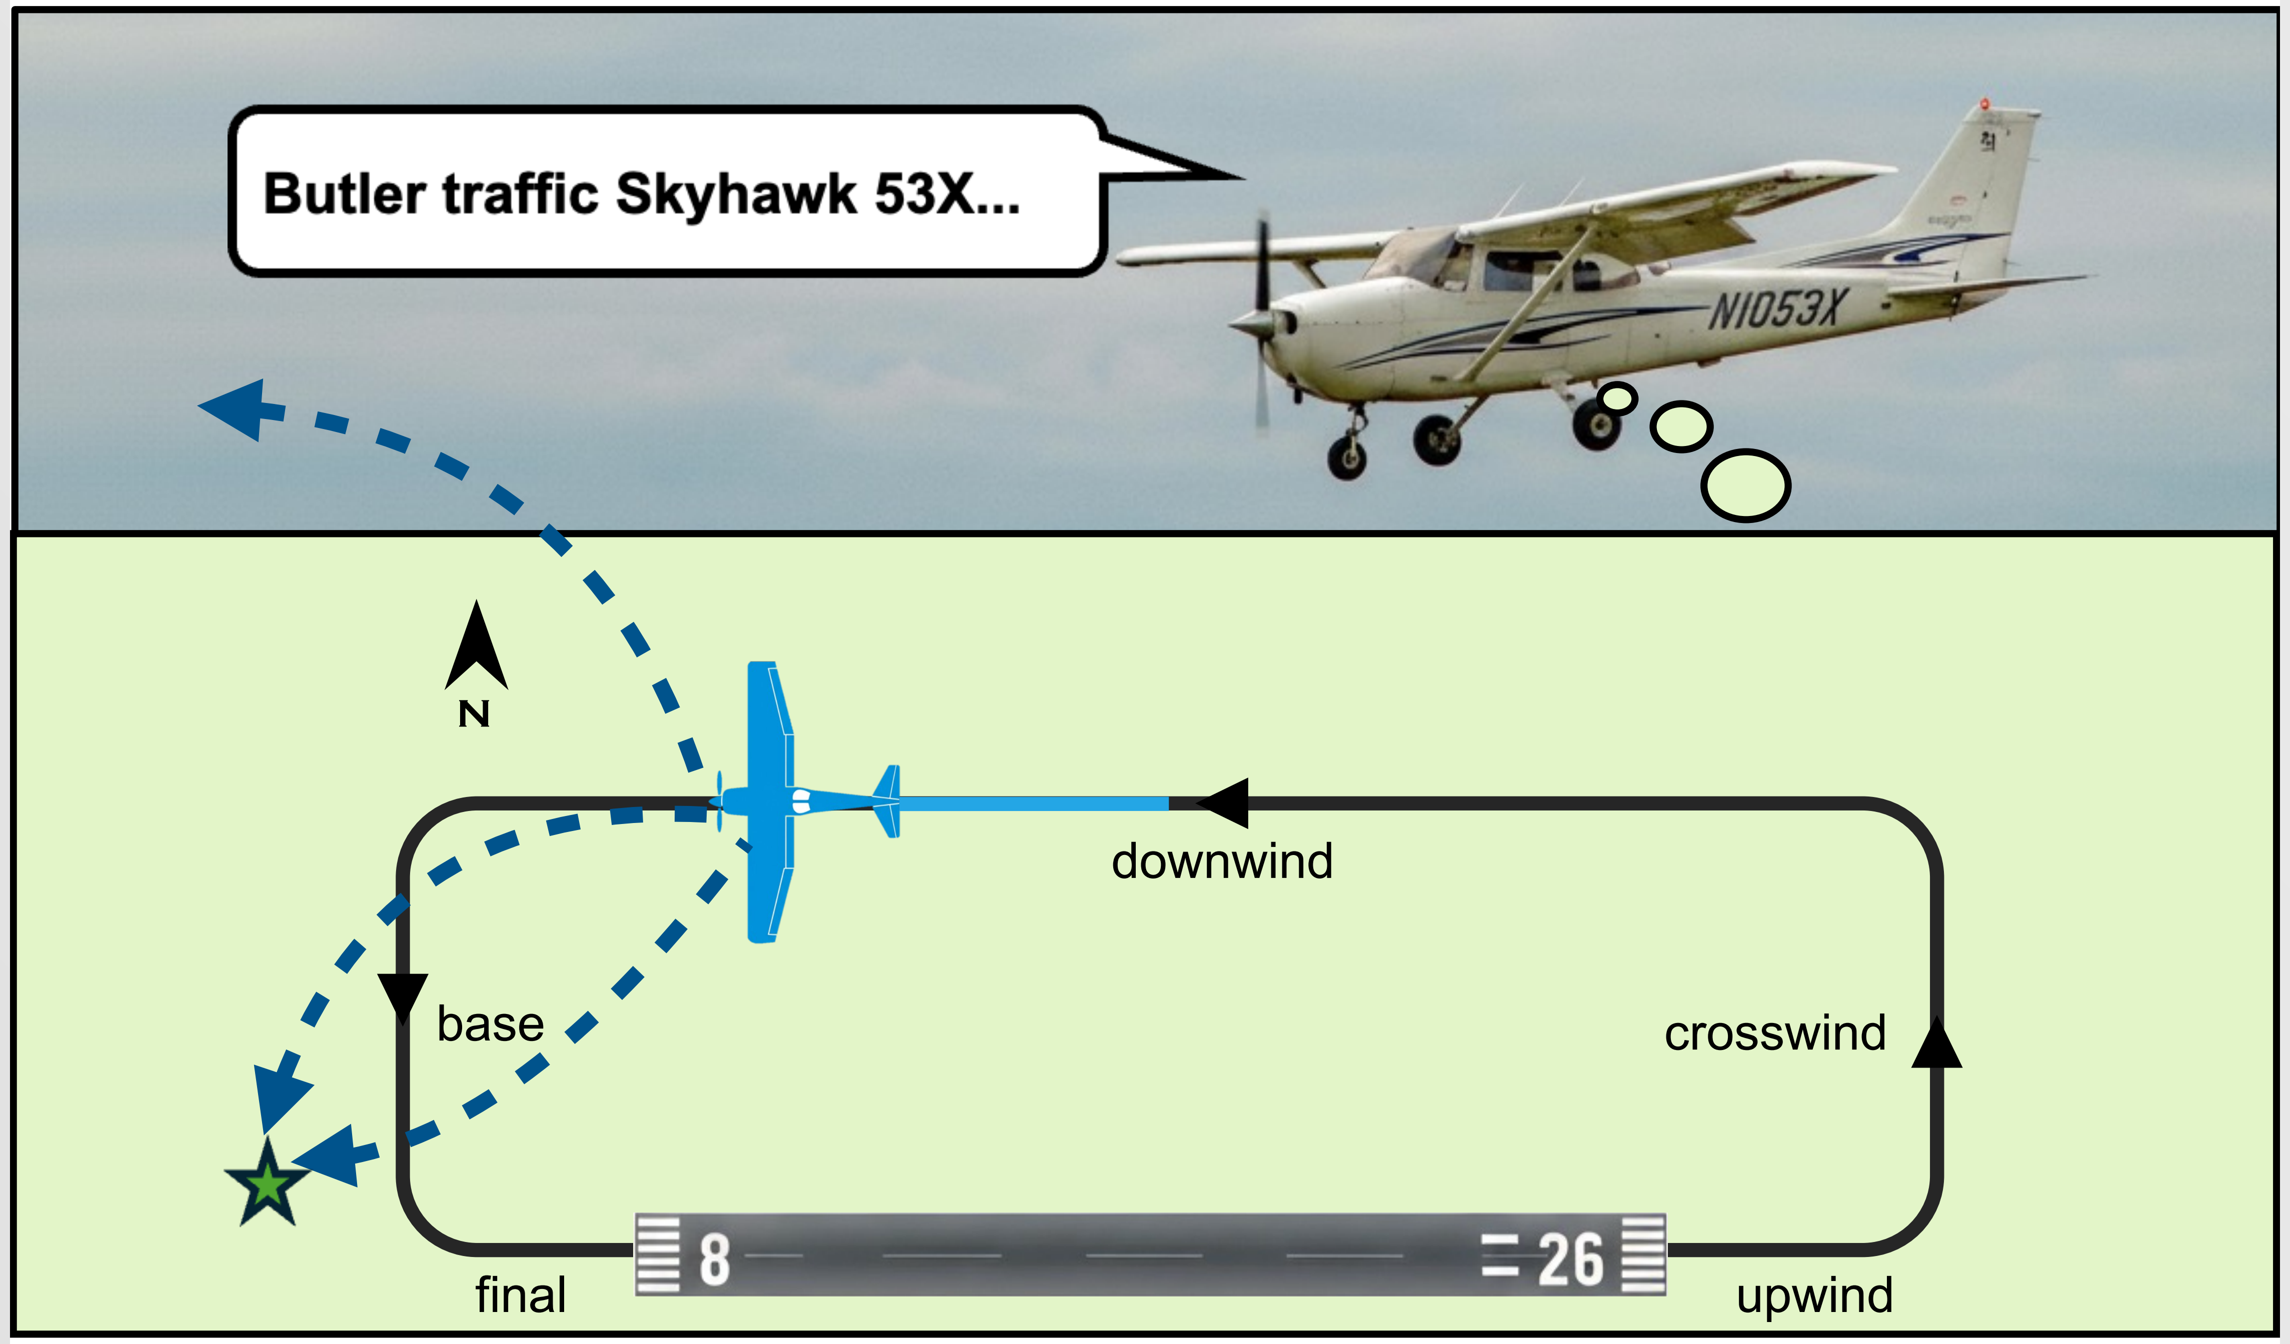

Pilot Audio for Trajectory Prediction - Problem Definition

Can audio from human pilots improve the accuracy of aircraft trajectory prediction? Specifically, we seek to use LLMs to infer a **physical goal location** ★ based on pilot audio, and use it to guide trajectory prediction algorithms.

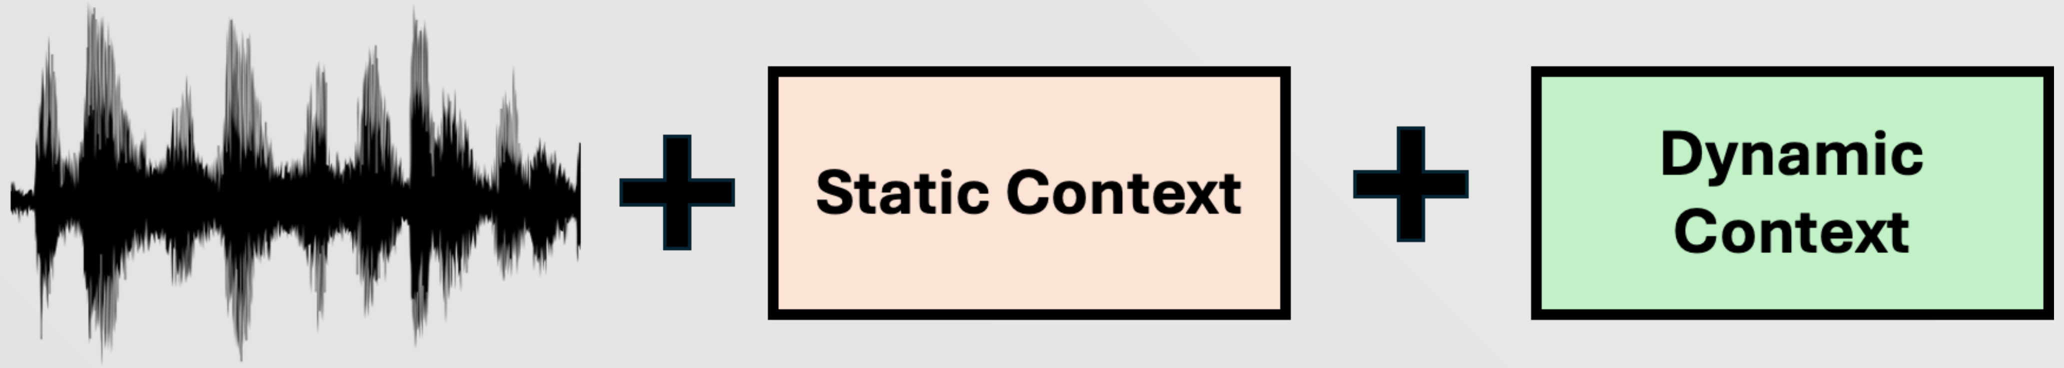

Static Context

- **Domain Vocabulary**
- **Runway Numbers**
- **Airport Name**
- **Traffic Pattern Segments**

Context Source: FAA standards

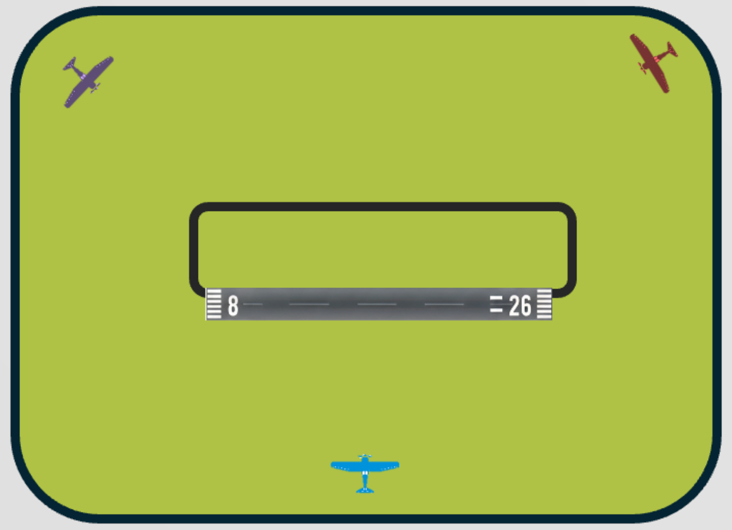

Context Source: ADS-B transponder data

Dynamic Context

N41A	**Mooney**	**1 mile S**
N72C	**Piper**	**3 miles NW**
N53X	**Skyhawk**	**5 miles NE**

25

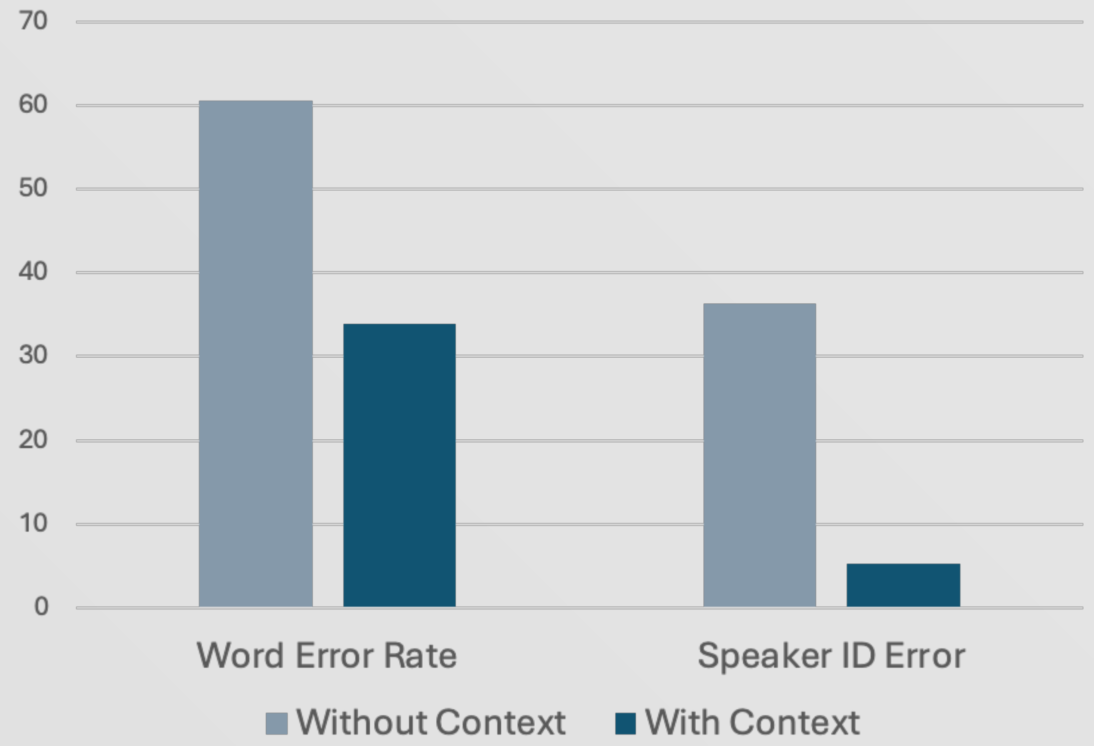

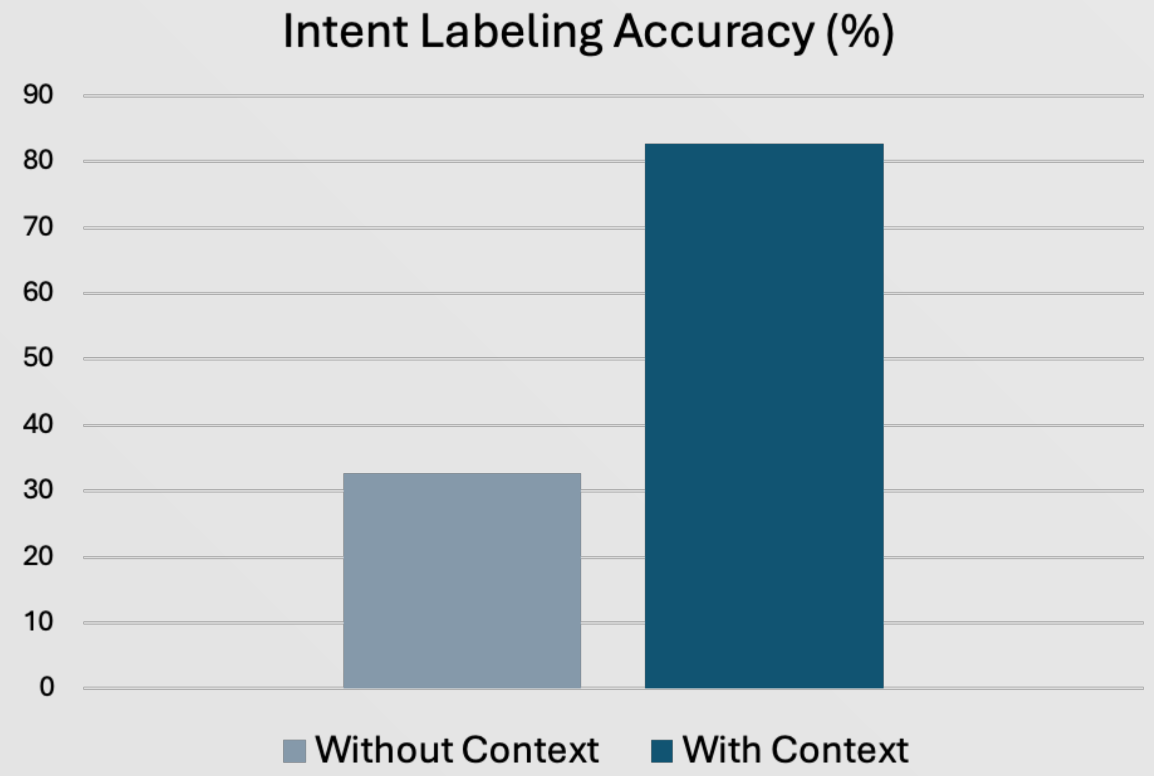

Technical challenges

- Providing LLMs with static contexts through prompting improved accuracy with no additional training
- We handcraft a small, discrete intent vector space structured around traffic patterns used at non towered airports
- Static + dynamic context significantly improved speech-to-intent performance

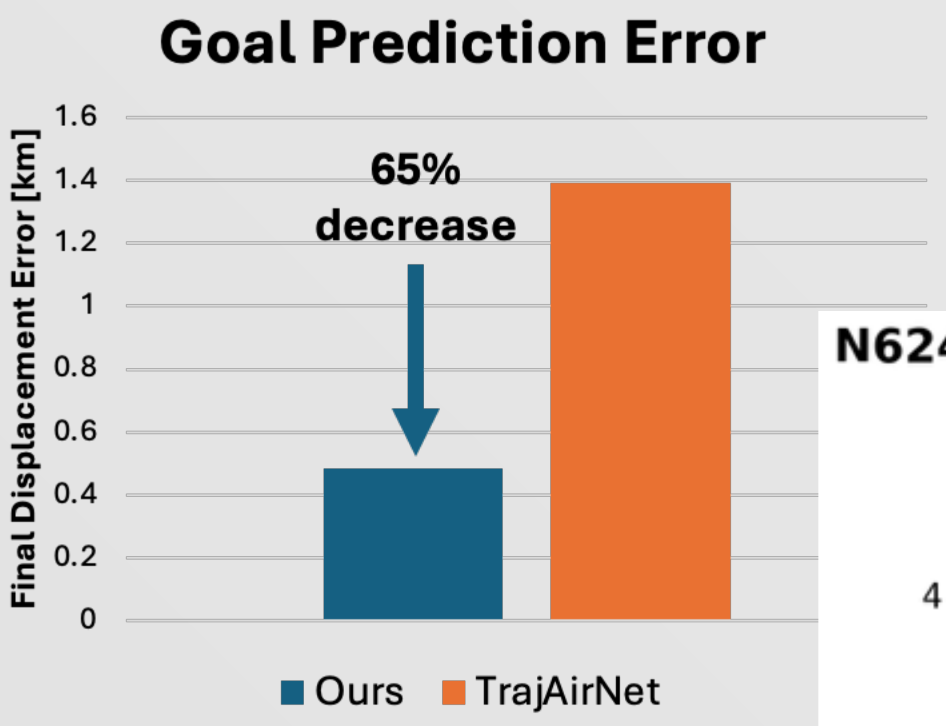

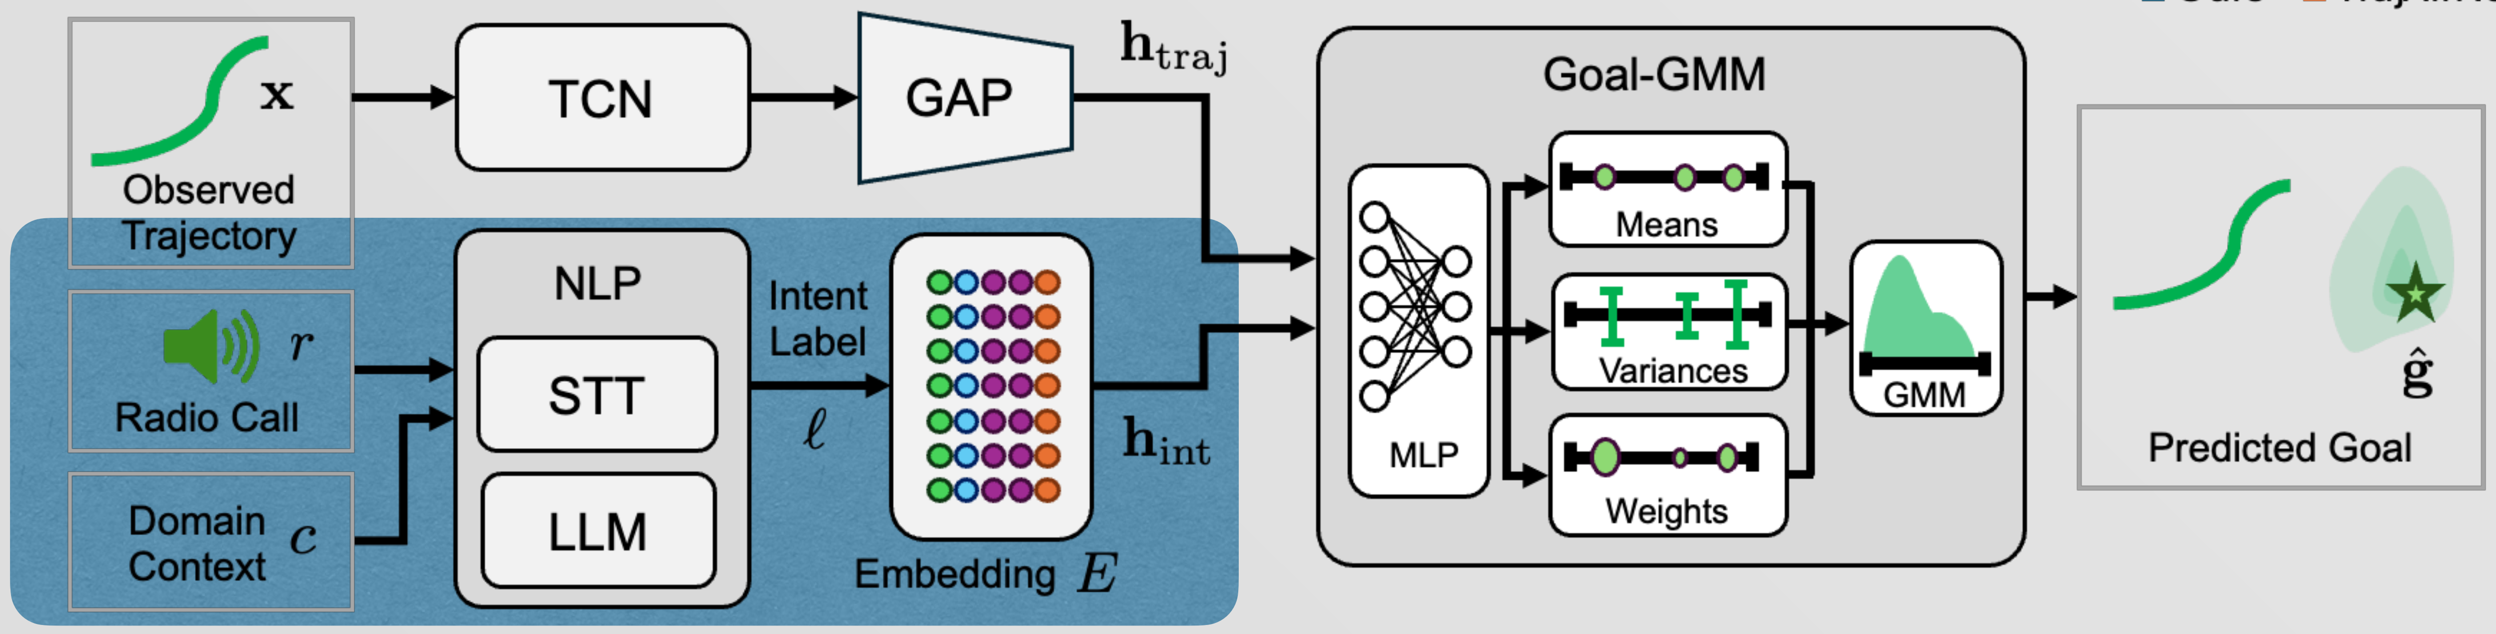

TrajAirNet (Patrikar'22)

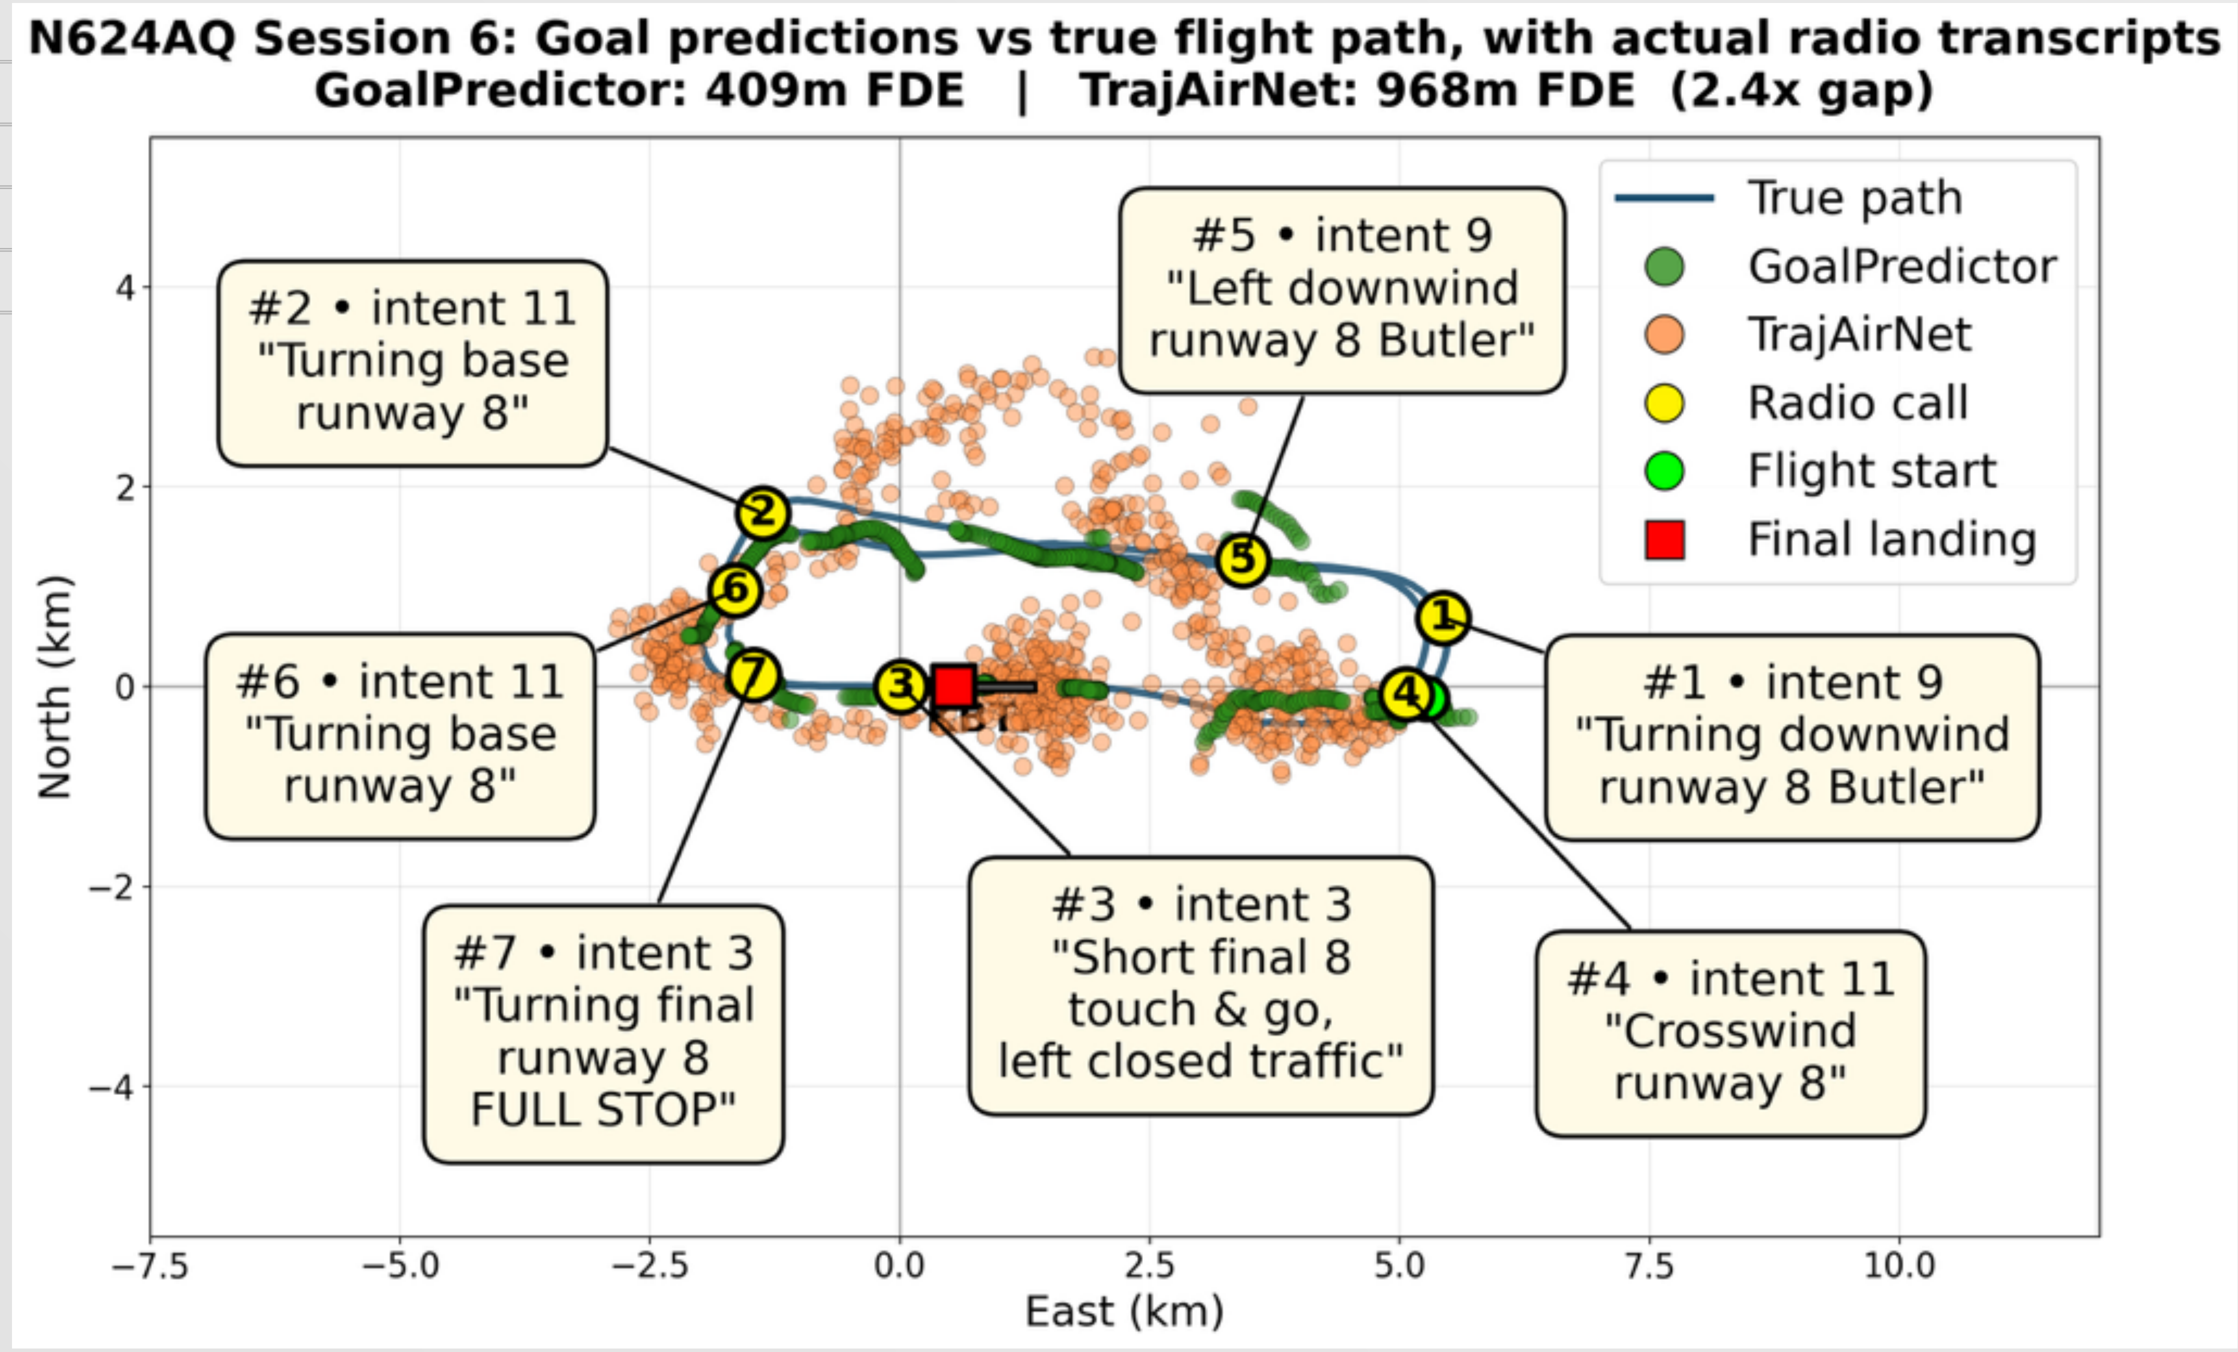

Audio-enabled Aircraft Trajectory Prediction

- Validated on goal prediction error (FDE $= \frac{1}{A} \sum_{i \in [A]} \|g_{a,j} - \hat{g}_{a,j}\|_2$) over a seven days audio + trajectory dataset
- Pilot audio can guide autonomous trajectory prediction in airspaces without control towers
- Audio can act as **a coordination signal between aircraft** without air traffic controller support

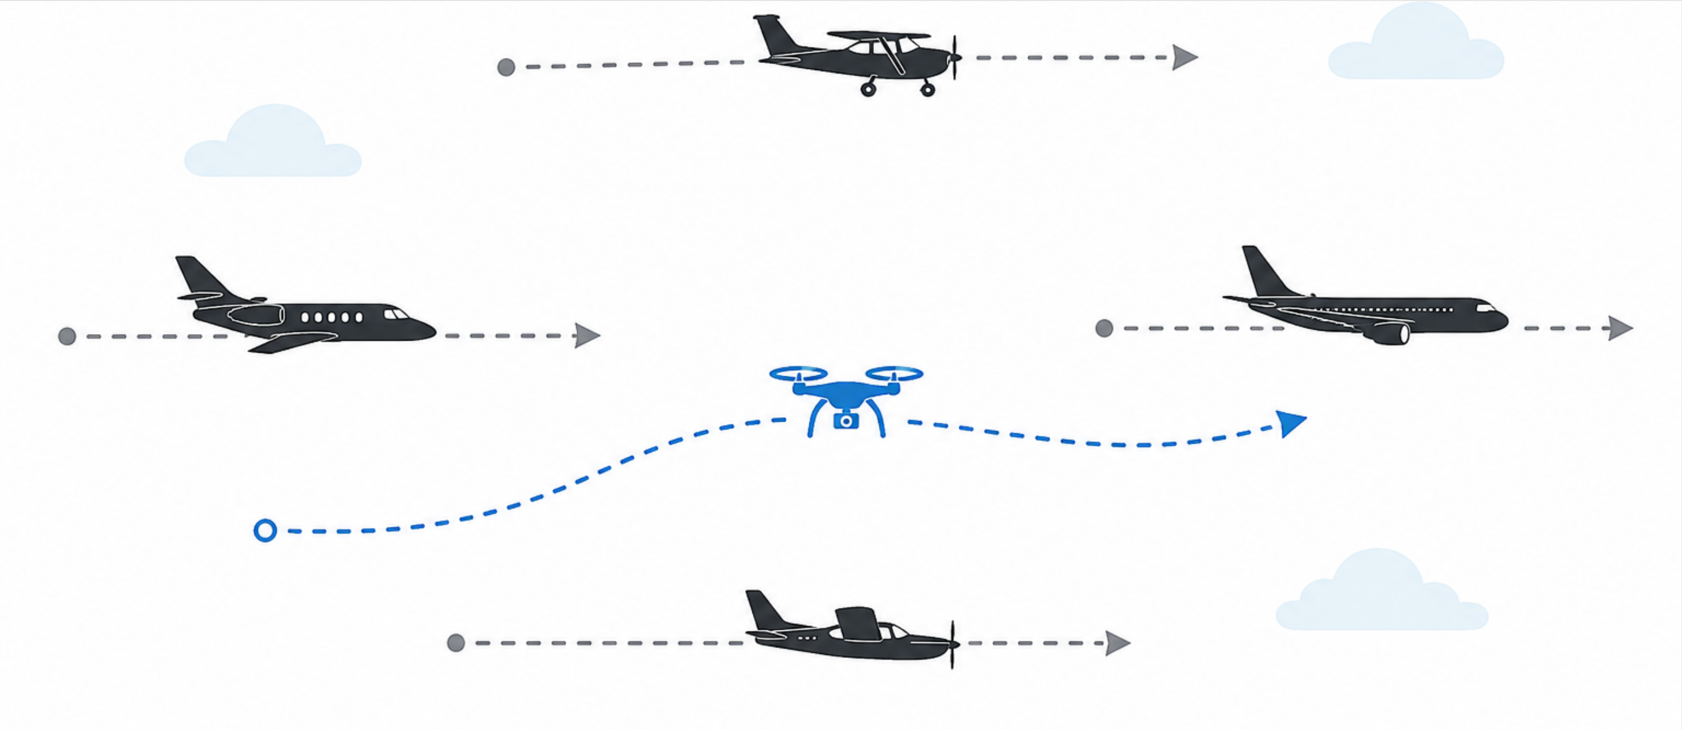

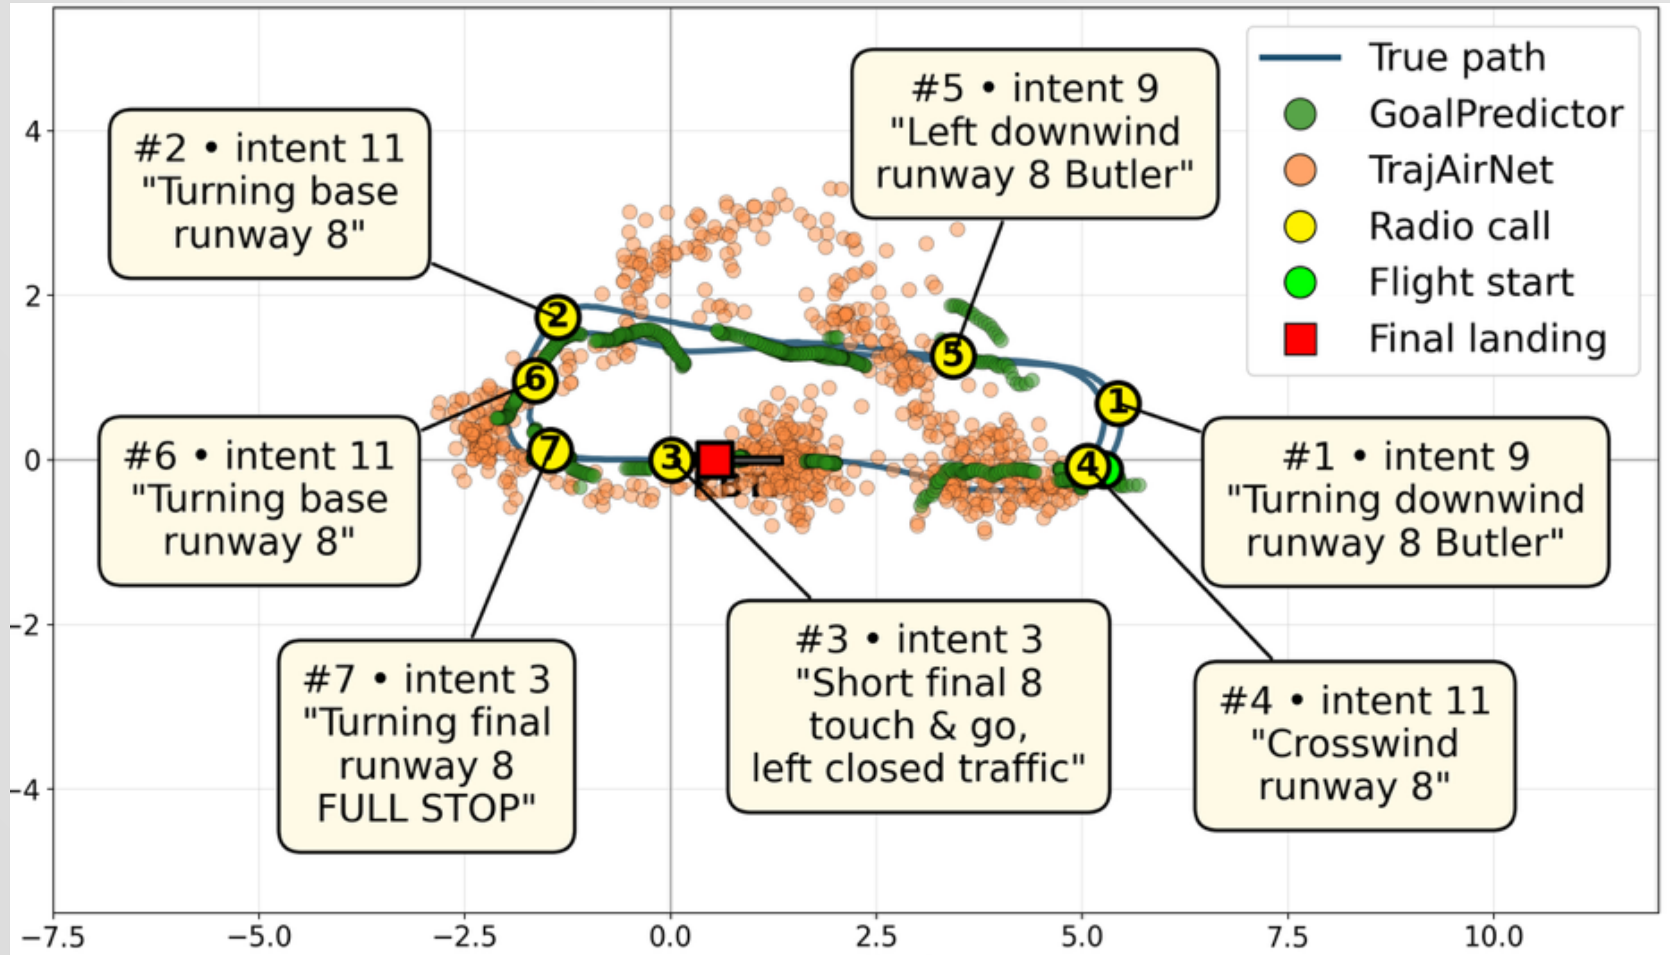

Autonomous aircraft can use audio to coordinate with human pilots

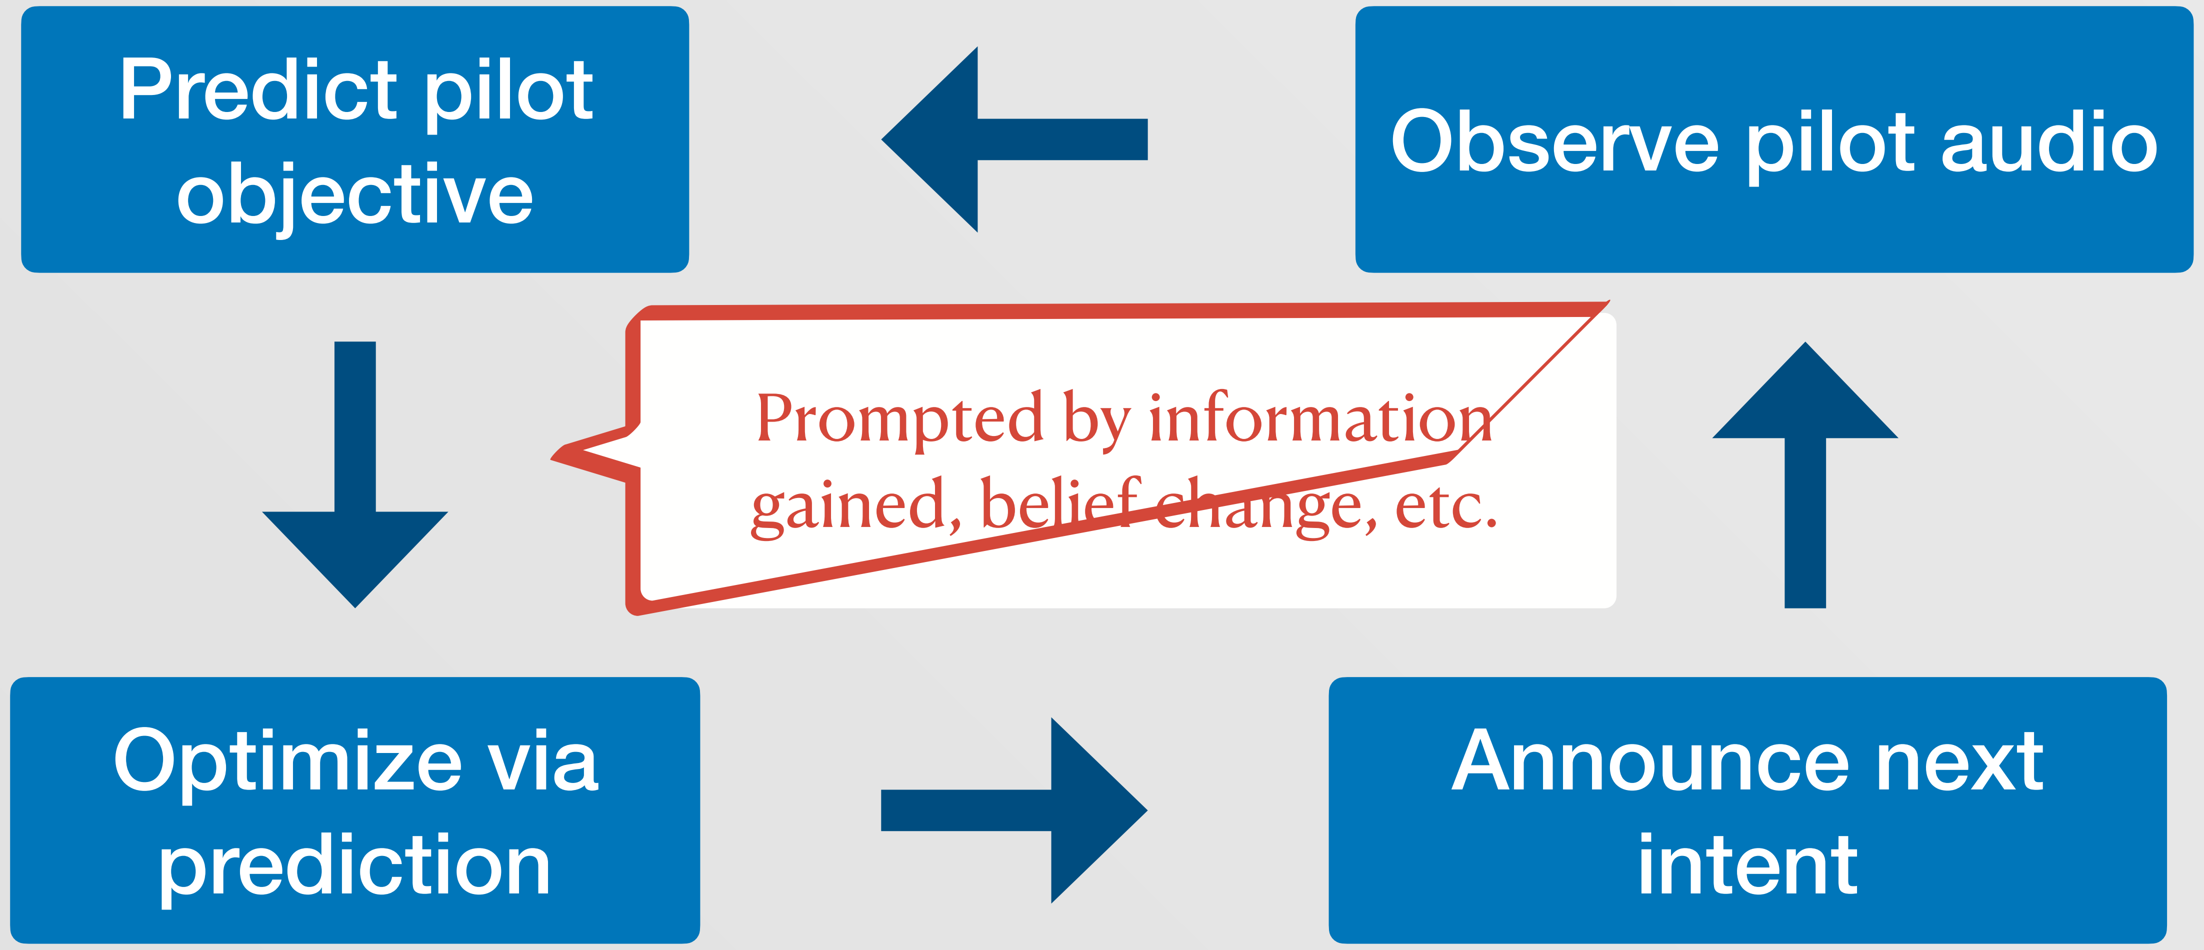

Mixed Human-Autonomous Aircraft Operations

- Pilot audio can guide autonomous trajectory prediction in airspaces without control towers
- Audio can act as **a coordination signal between aircraft** without air traffic controller support

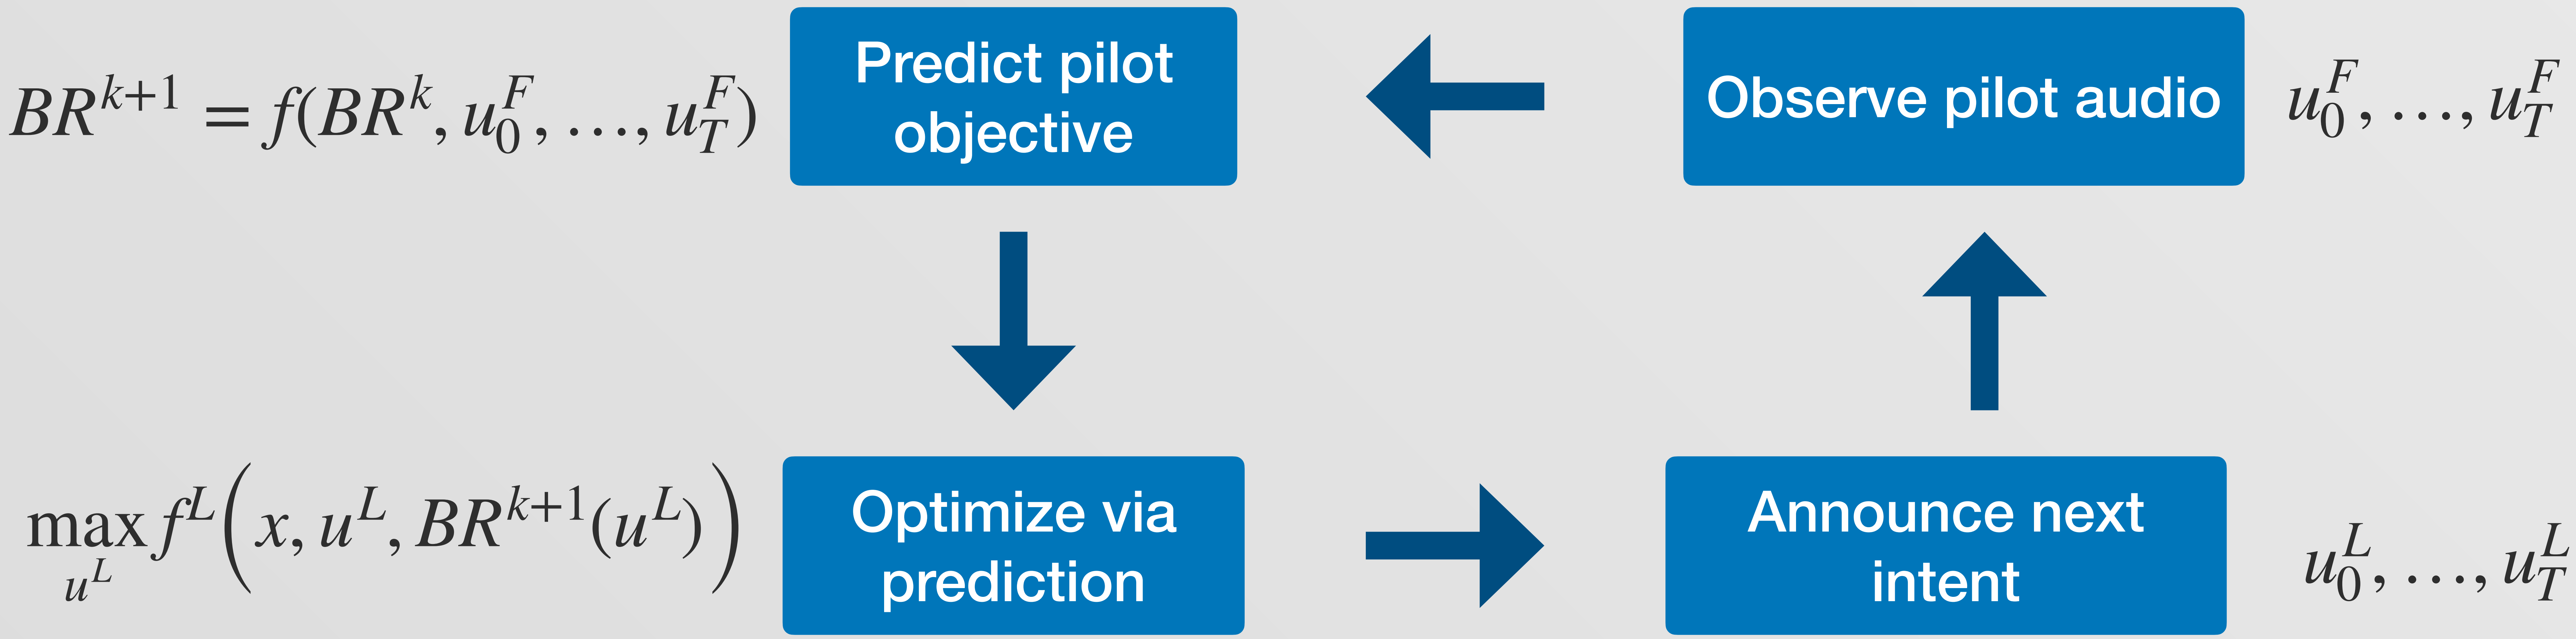

A Stackelberg Game Approach for Coordination

- Autonomous aircraft predict pilot responses to its open loop trajectories and use this model to achieve its objectives
- **Is Stackelberg's open loop solution always optimal under recomputation?**

- We consider a discrete-time LTI system with dynamics

Leader (L)

$$x_{t+1} = Ax_t + B^L u_t^L + B^F u_t^F, \forall t \in [T-1]$$

$$x_{0:T} = Hx_0 + G^L u_{0:T-1}^L + G^F u_{0:T-1}^F$$

Follower (F)

- Players have open-loop information structure

$$\Pi_t^i = (u_0^i, \ldots u_{T-1}^i), i = L, F$$

open loop (OL)

- The player's cost is differentiable, strictly convex on x_t, u_t^L, u_t^F (respectively), and stage-decomposable

$$J_{0:T}^i(x_{0:T}, u_{0:T-1}^L, u_{0:T-1}^F) = J_T^i(x_T) + \sum_{t=0}^{T-1} J_t^i(x_t, u_t^L, u_t^F), \quad i = L, F$$

Stackelberg Dynamic Game

- The follower's BR is given by

$$\begin{aligned} b^\star(\boldsymbol{\pi}^L) \in \arg \min_{\boldsymbol{\pi}^F \in \Pi_{0:T-1}^F} \quad & J_{0:T}^F(\mathbf{x}, \boldsymbol{\pi}^L, \boldsymbol{\pi}^F) \\ \text{s.t.} \quad & \mathbf{x} = Hx_0 + G^L \boldsymbol{\pi}^L + G^F \boldsymbol{\pi}^F \end{aligned}$$

(BR)

- And the leader solves the problem

$$\begin{aligned} \min_{\boldsymbol{\pi}^L \in \Pi_{0:T-1}^L} \quad & J_{0:T}^L(\mathbf{x}, \boldsymbol{\pi}^L, b^\star(\boldsymbol{\pi}^L)) \\ \text{s.t.} \quad & \mathbf{x} = Hx_0 + G^L \boldsymbol{\pi}^L + G^F b^\star(\boldsymbol{\pi}^L) \end{aligned}$$

(L)

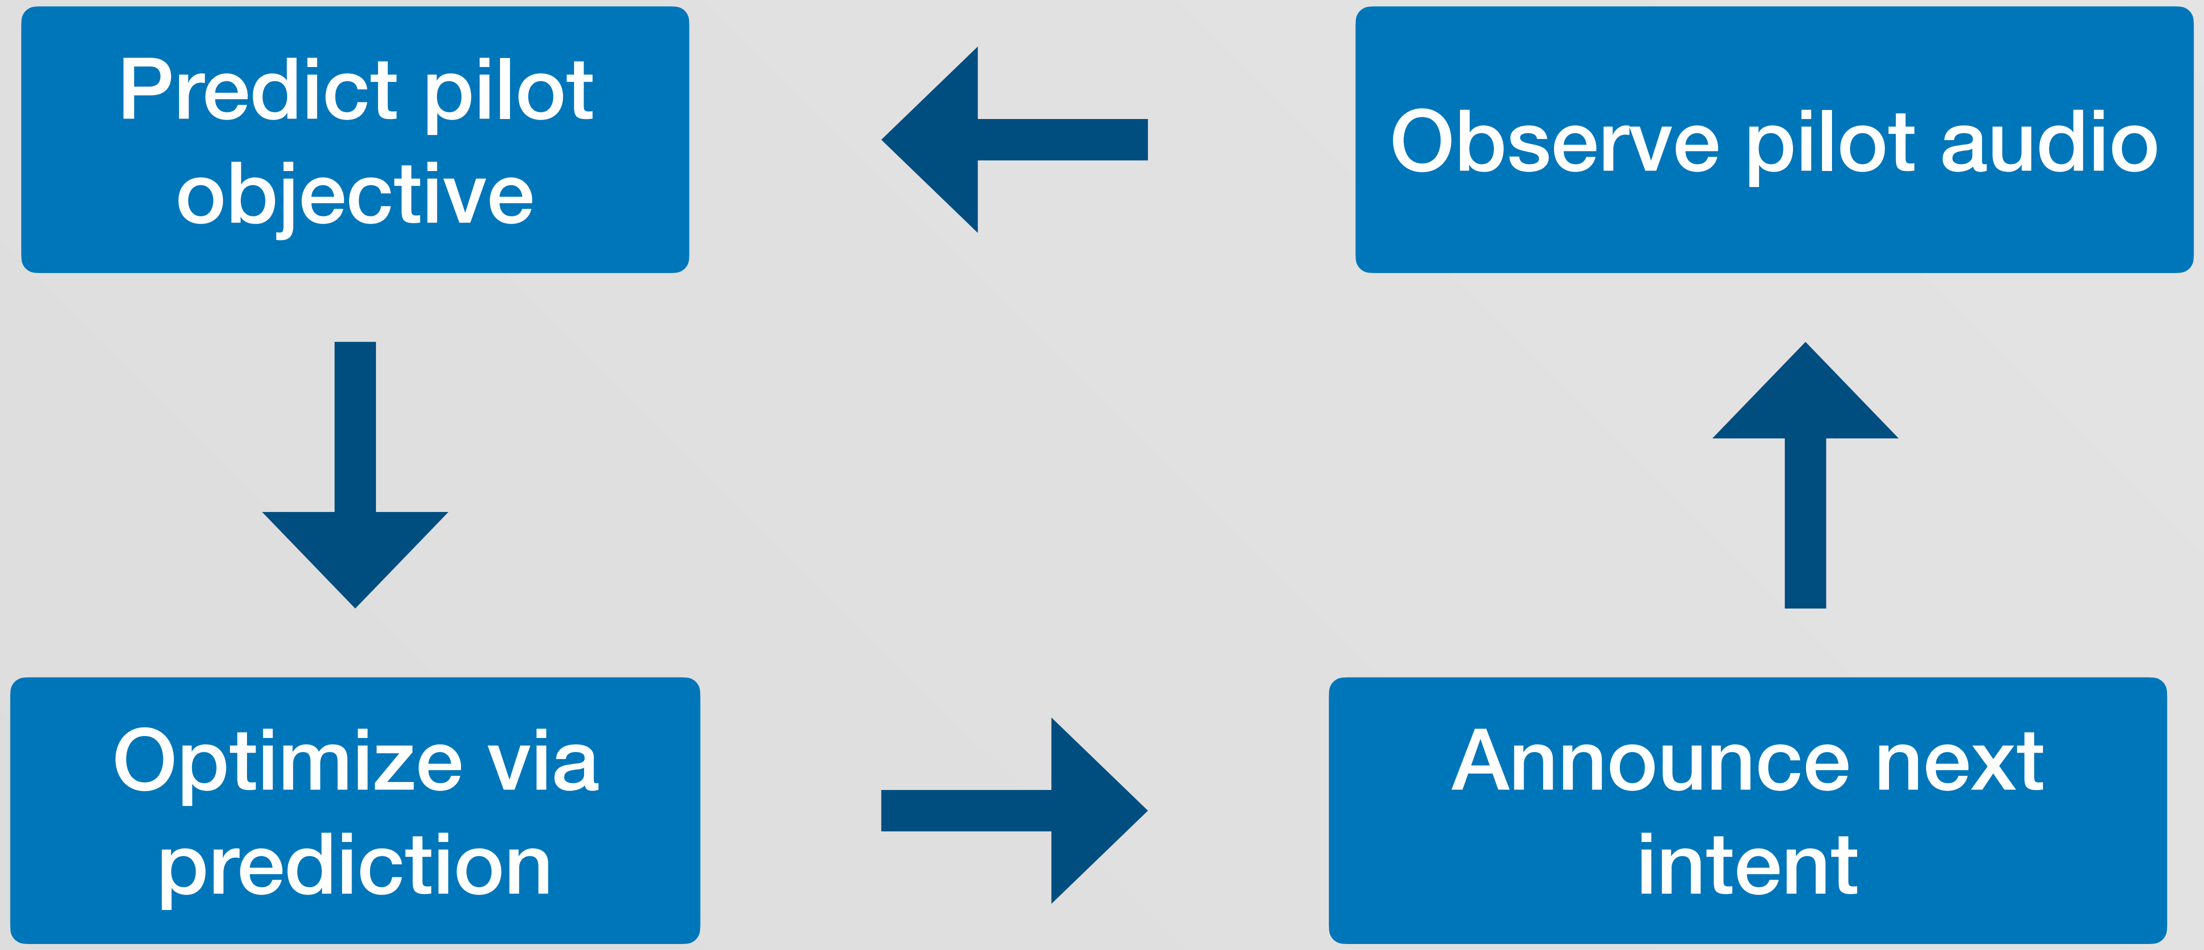

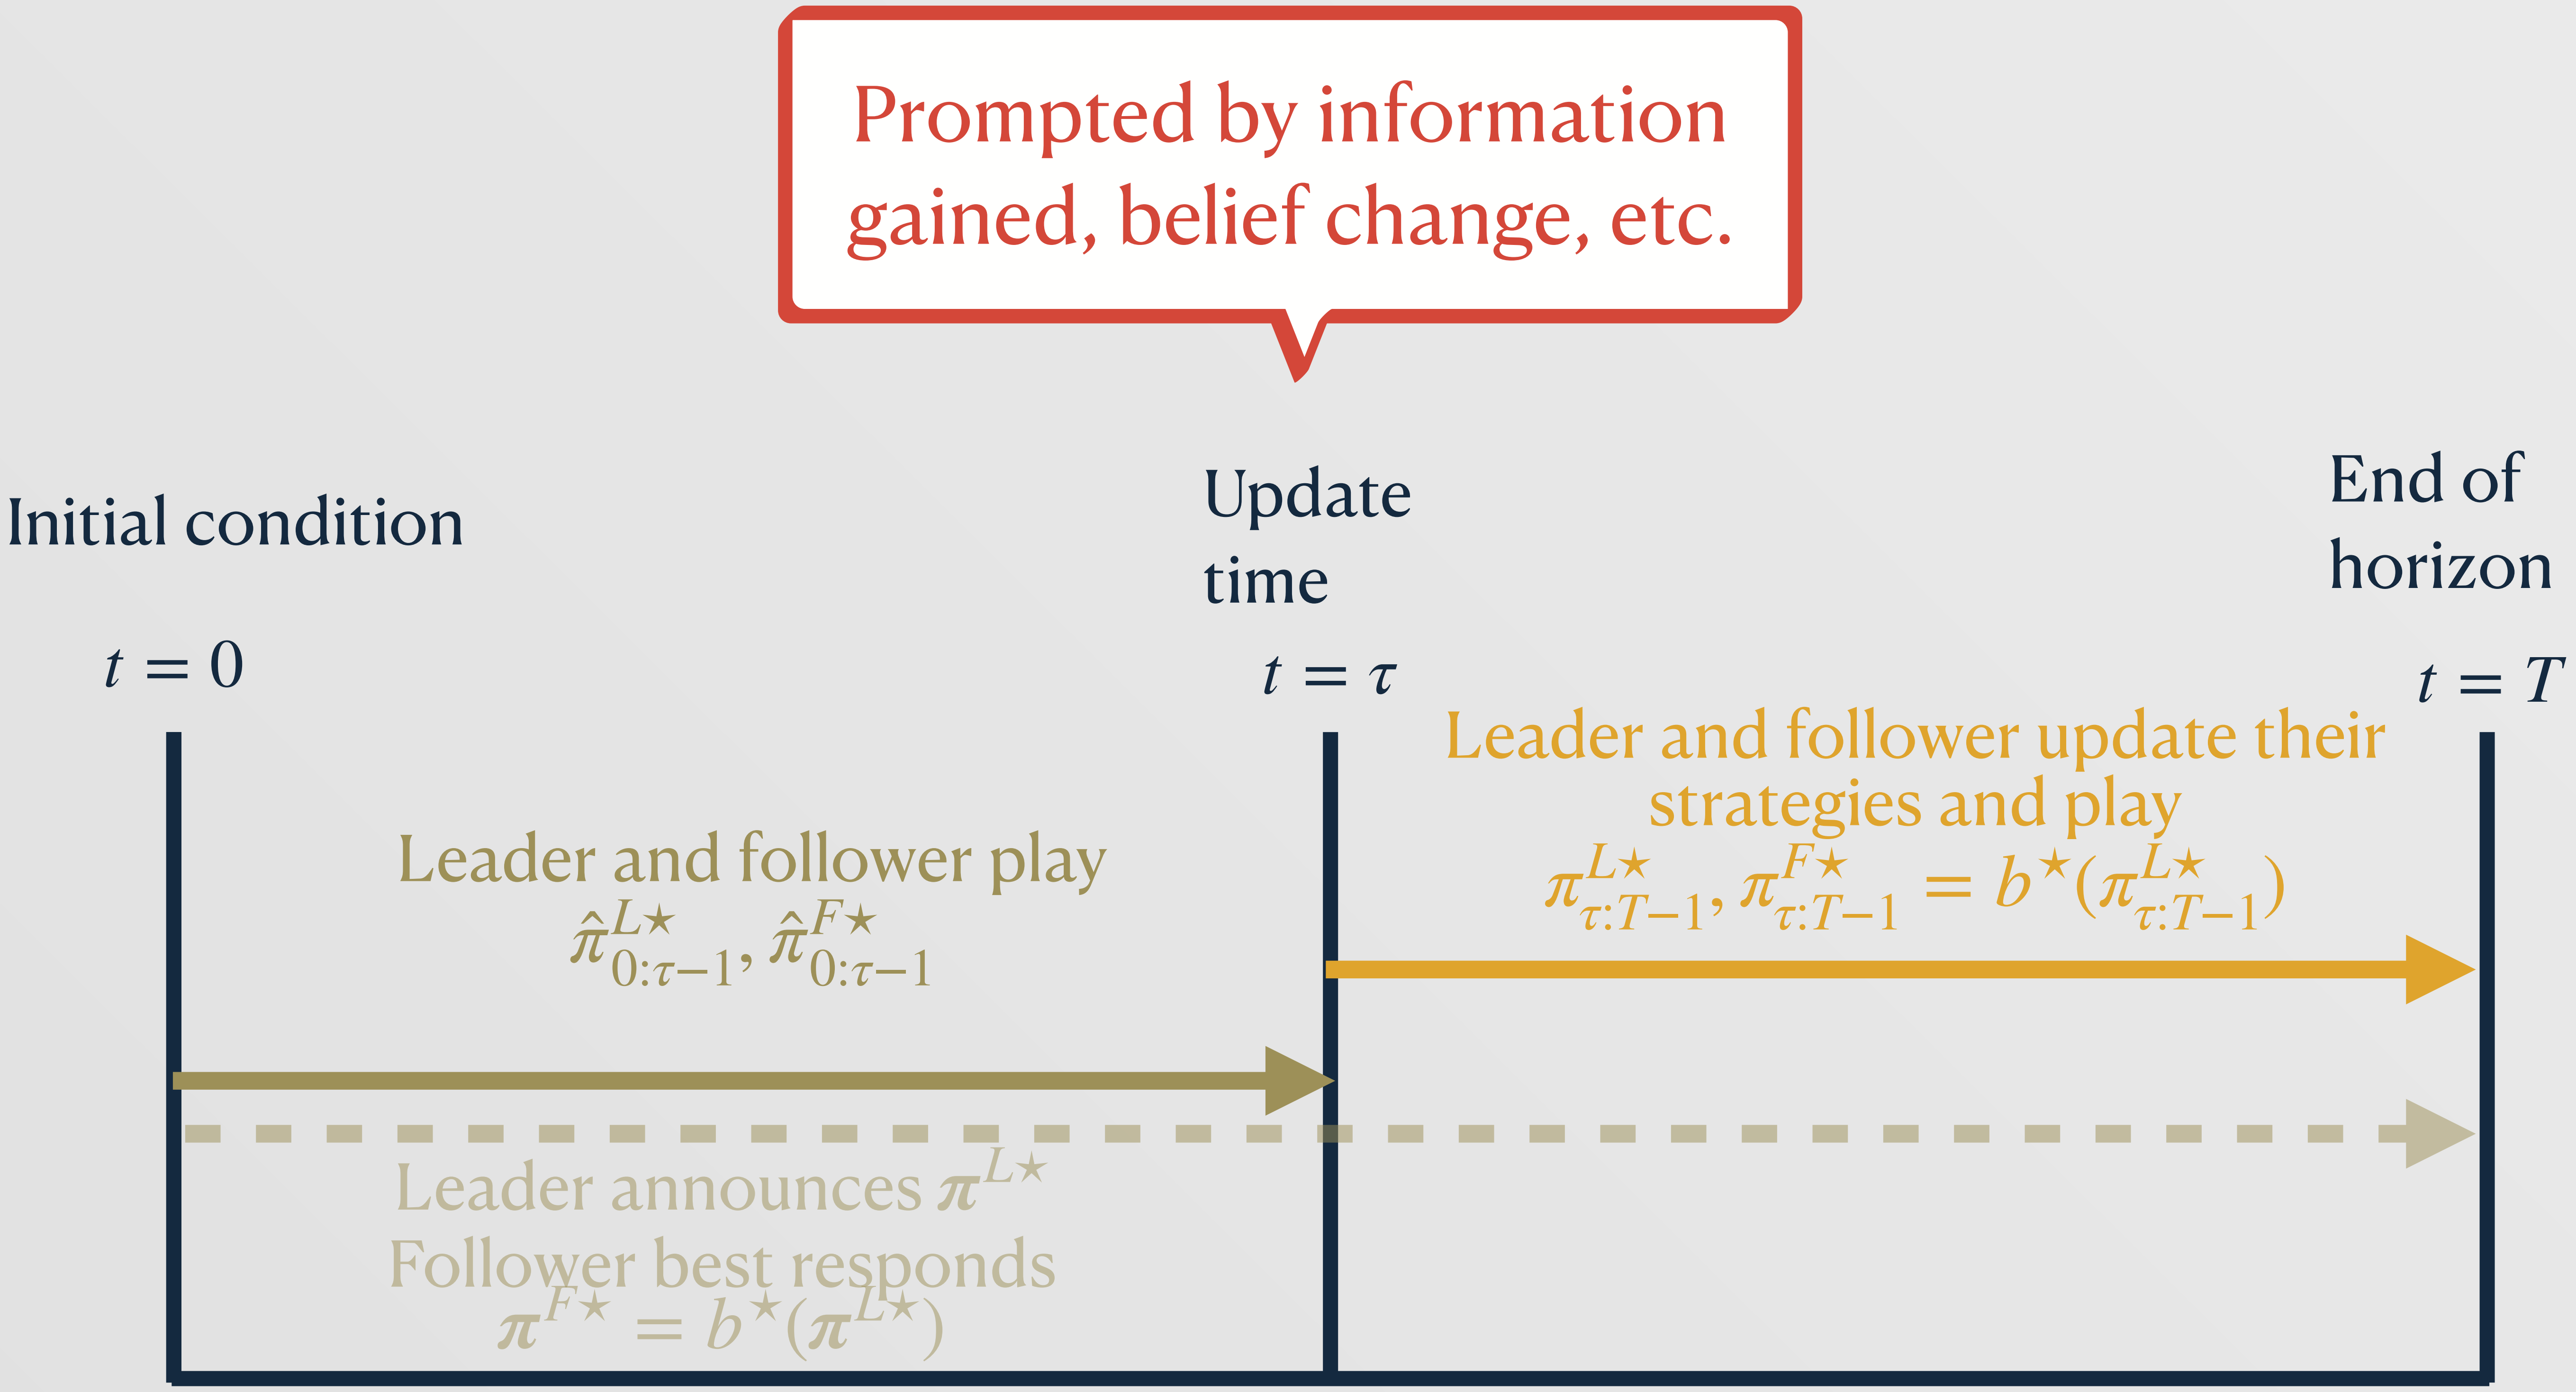

Stackelberg Dynamic Game

- At update time τ, is $\pi^{L\star}_{\tau:T-1}$ still optimal for the remaining time steps?
- No (time-inconsistency), the leader can announce a different strategy and achieve a lower cost

Open-loop Stackelberg Equilibrium (OLSE)

$J^L_{0:T} = 3.61$

At $t = 0$ leader announces $\hat{u}^L_{0:T-1} =$ $[-0.82, -0.34, -0.15, -0.06, -0.02, -0.01]$

At $t = \tau$ leader announces $u^L_{\tau:T-1} = [-0.14, -0.06, -0.02]$

OL strategy with update at tau

$J^L_{0:T} = 2.64$

At $t = 0$ leader announces $\hat{u}^L_{0:T-1} =$ $[-0.32, 0.1, 0.21, 3.52, 2.23, 1.08]$

At $t = \tau$ leader announces $u^L_{\tau:T-1} = [0.1, 0.04, 0.02]$

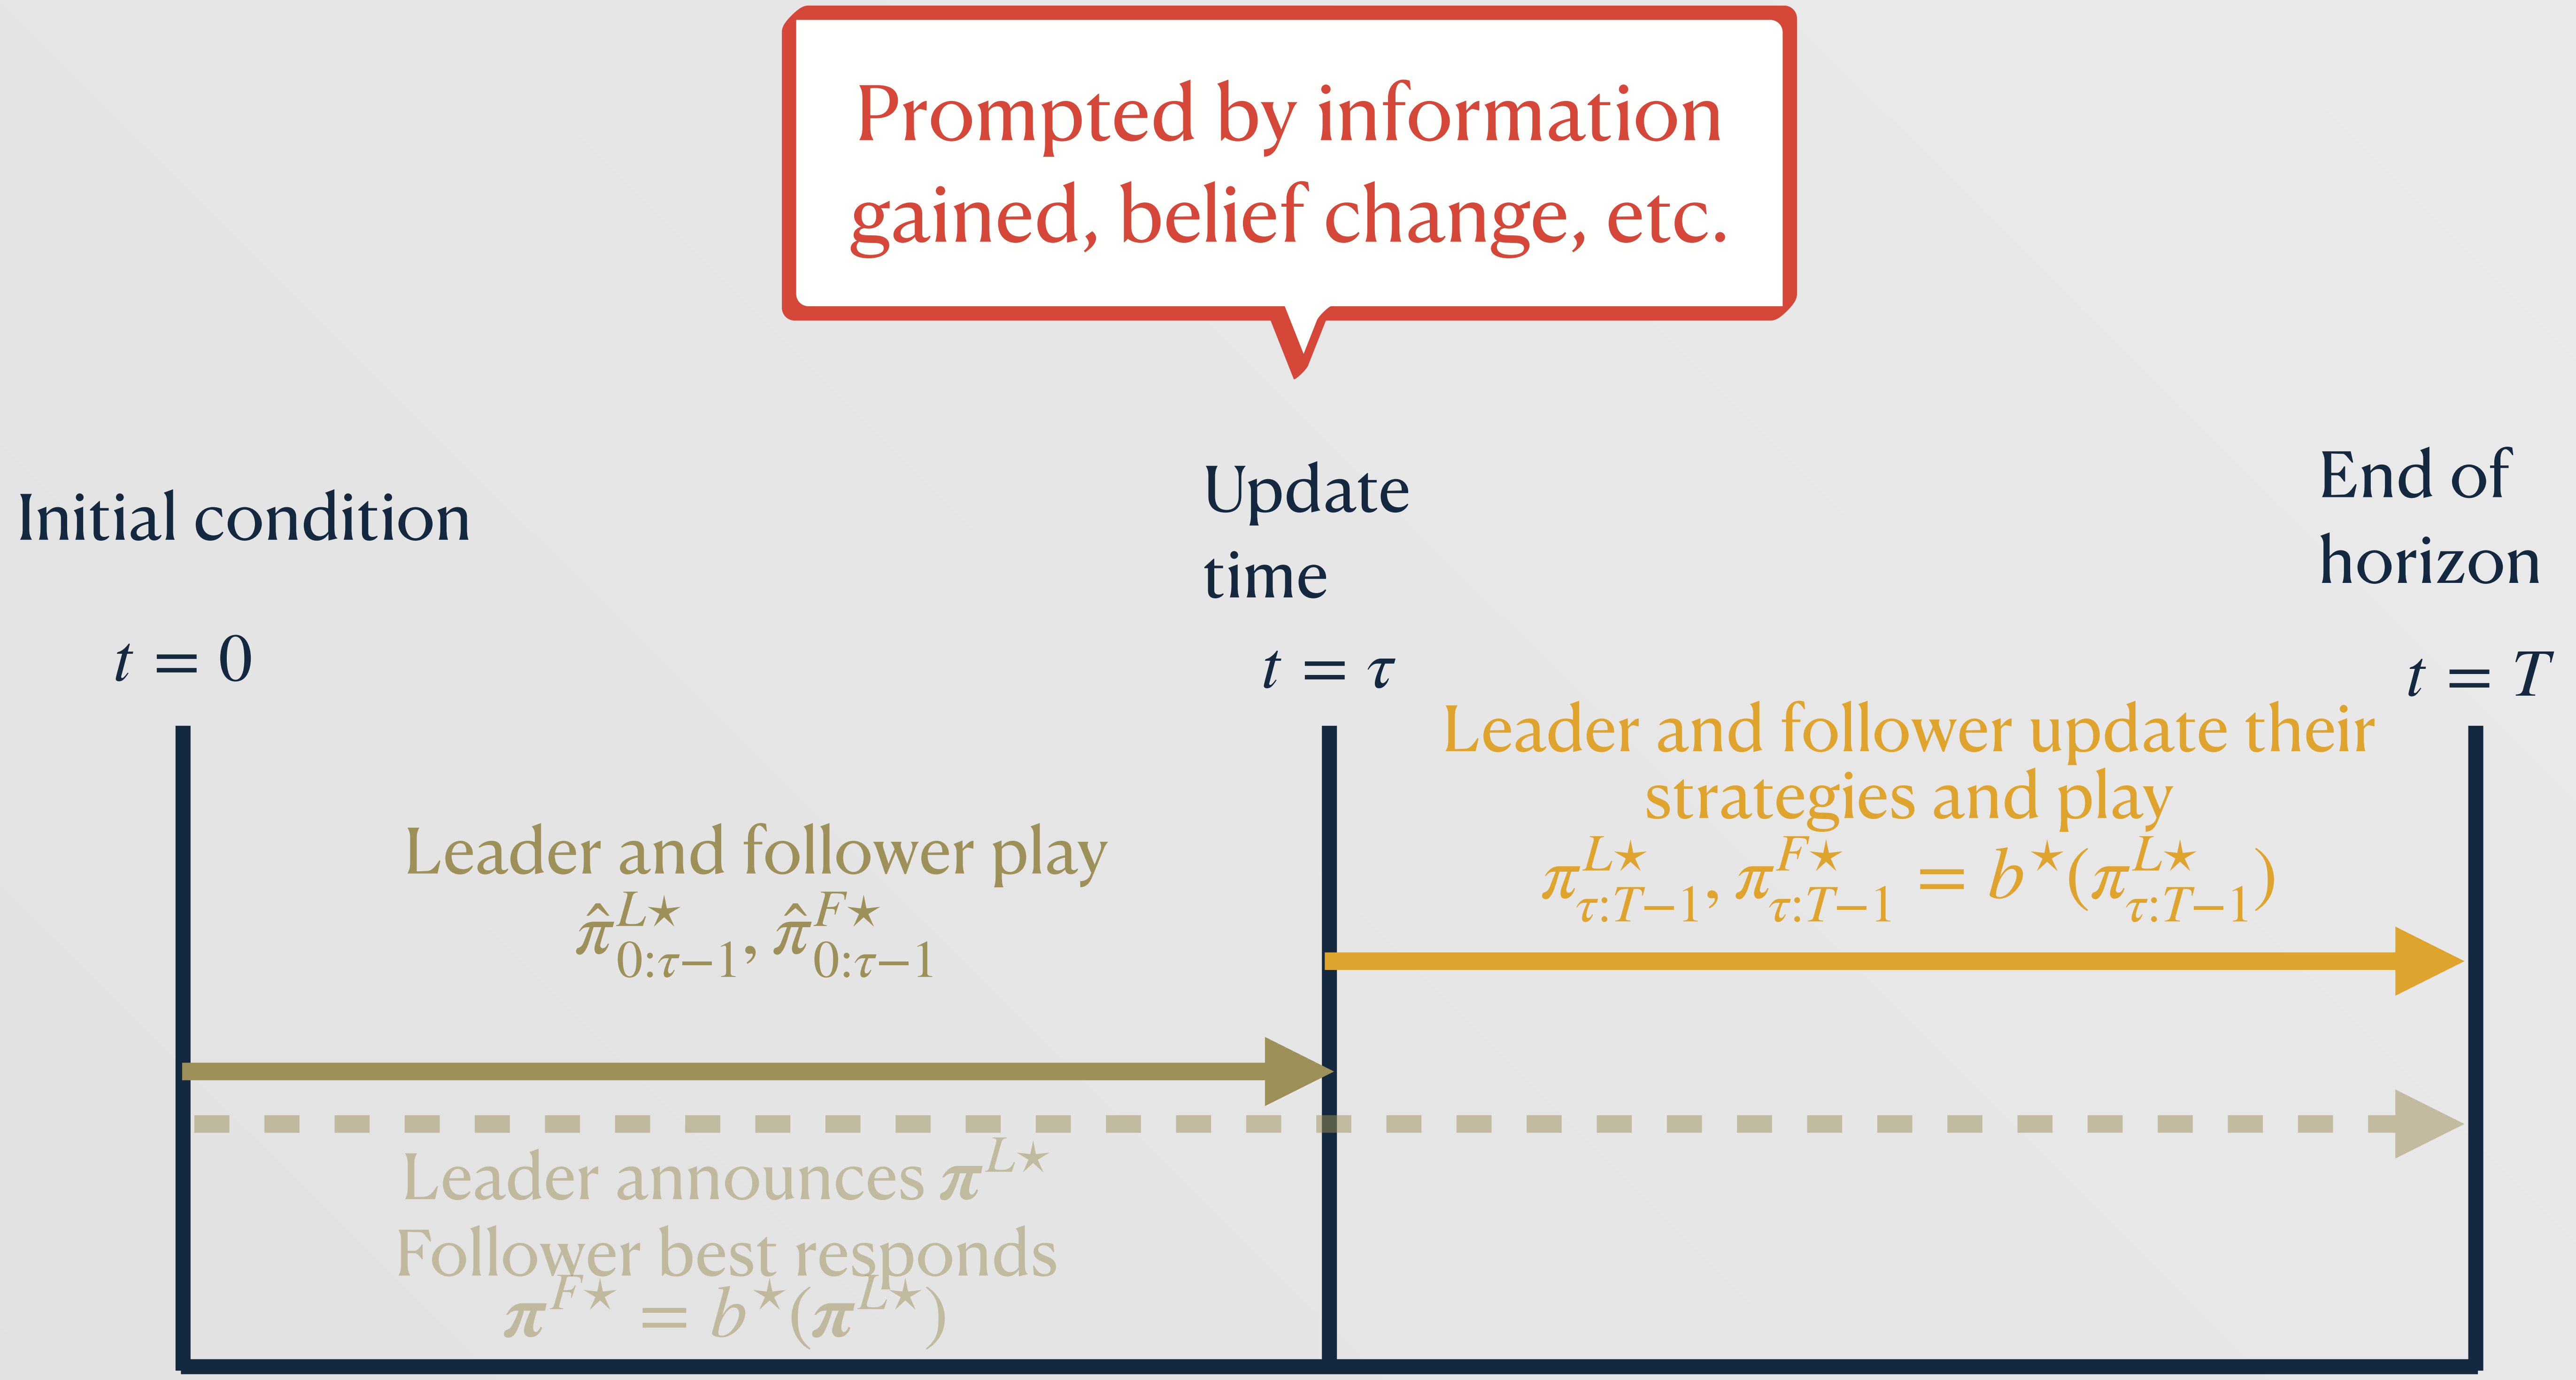

1D linear quadratic (LQ) Stackelberg Game

- 6 time steps, update time is 3, 1-D dynamics

$$x_{t+1} = x_t + 0.5\, u^L_t + 0.5\, u^F_t, \quad x_0 = 1, \quad T = 6, \ \tau = 3$$

- Player cost functions are:

$$J^L_{0:T} = \sum_{t=0}^{T} x_t^\top 2 x_t + \sum_{t=0}^{T-1} u_t^{L\top} u_t^L \qquad J^F_{0:T} = \sum_{t=0}^{T} x_t^\top 0.5 x_t + \sum_{t=0}^{T-1} u_t^{L\top} 2 u_t^L$$

LQ collision avoidance example

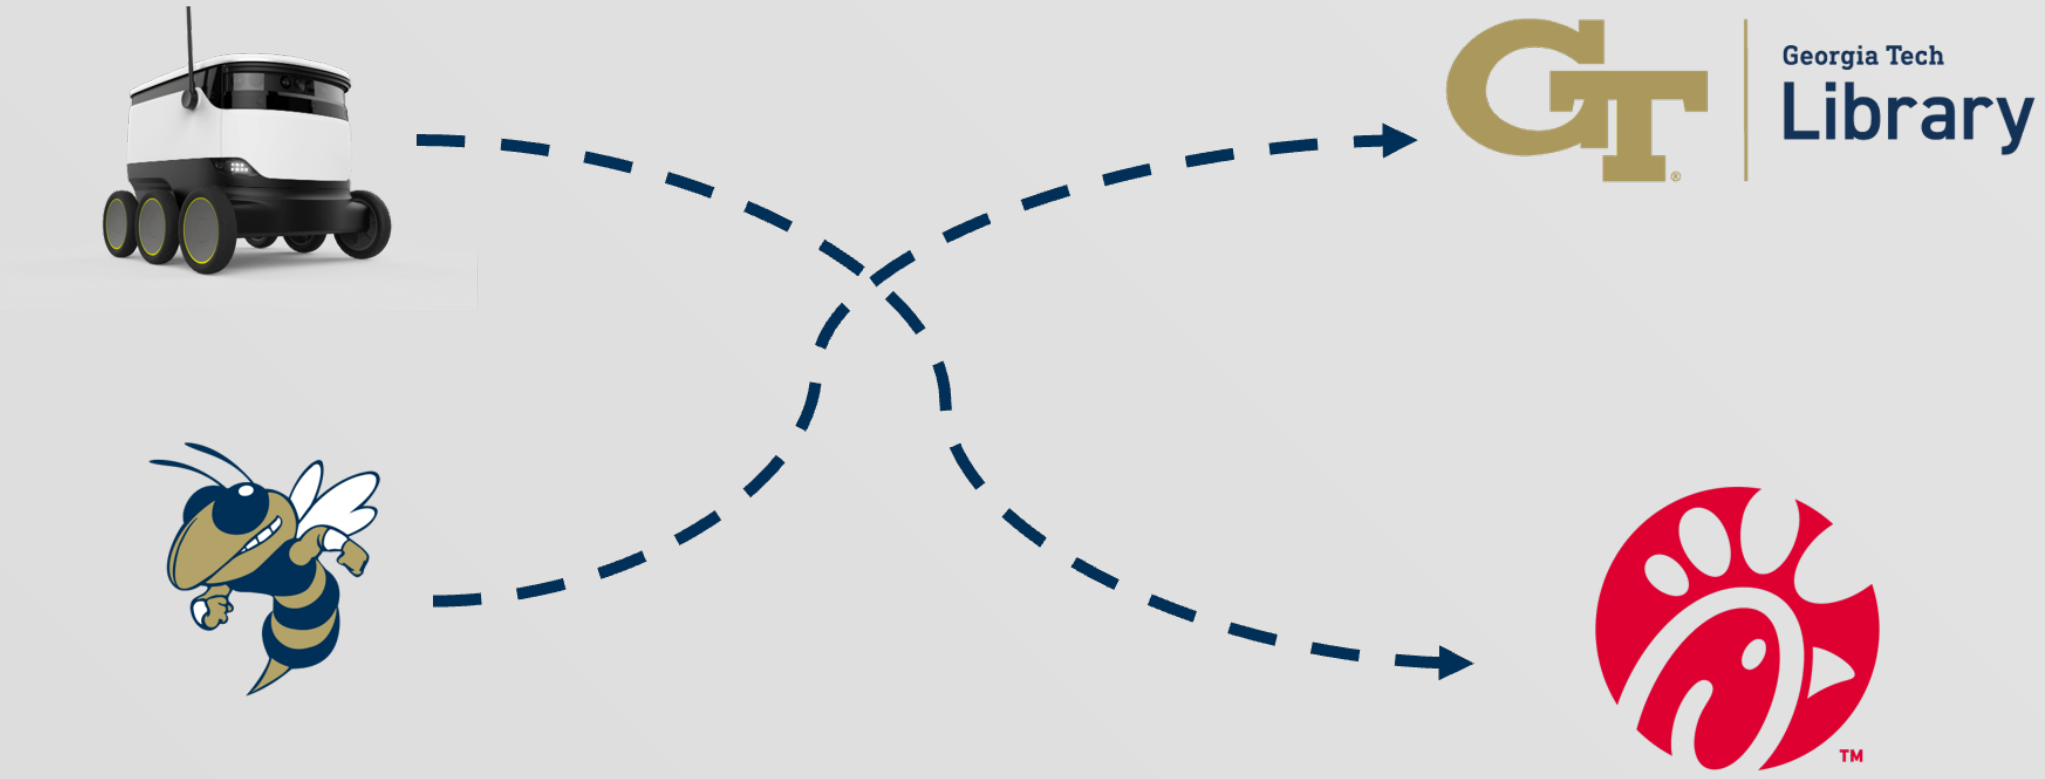

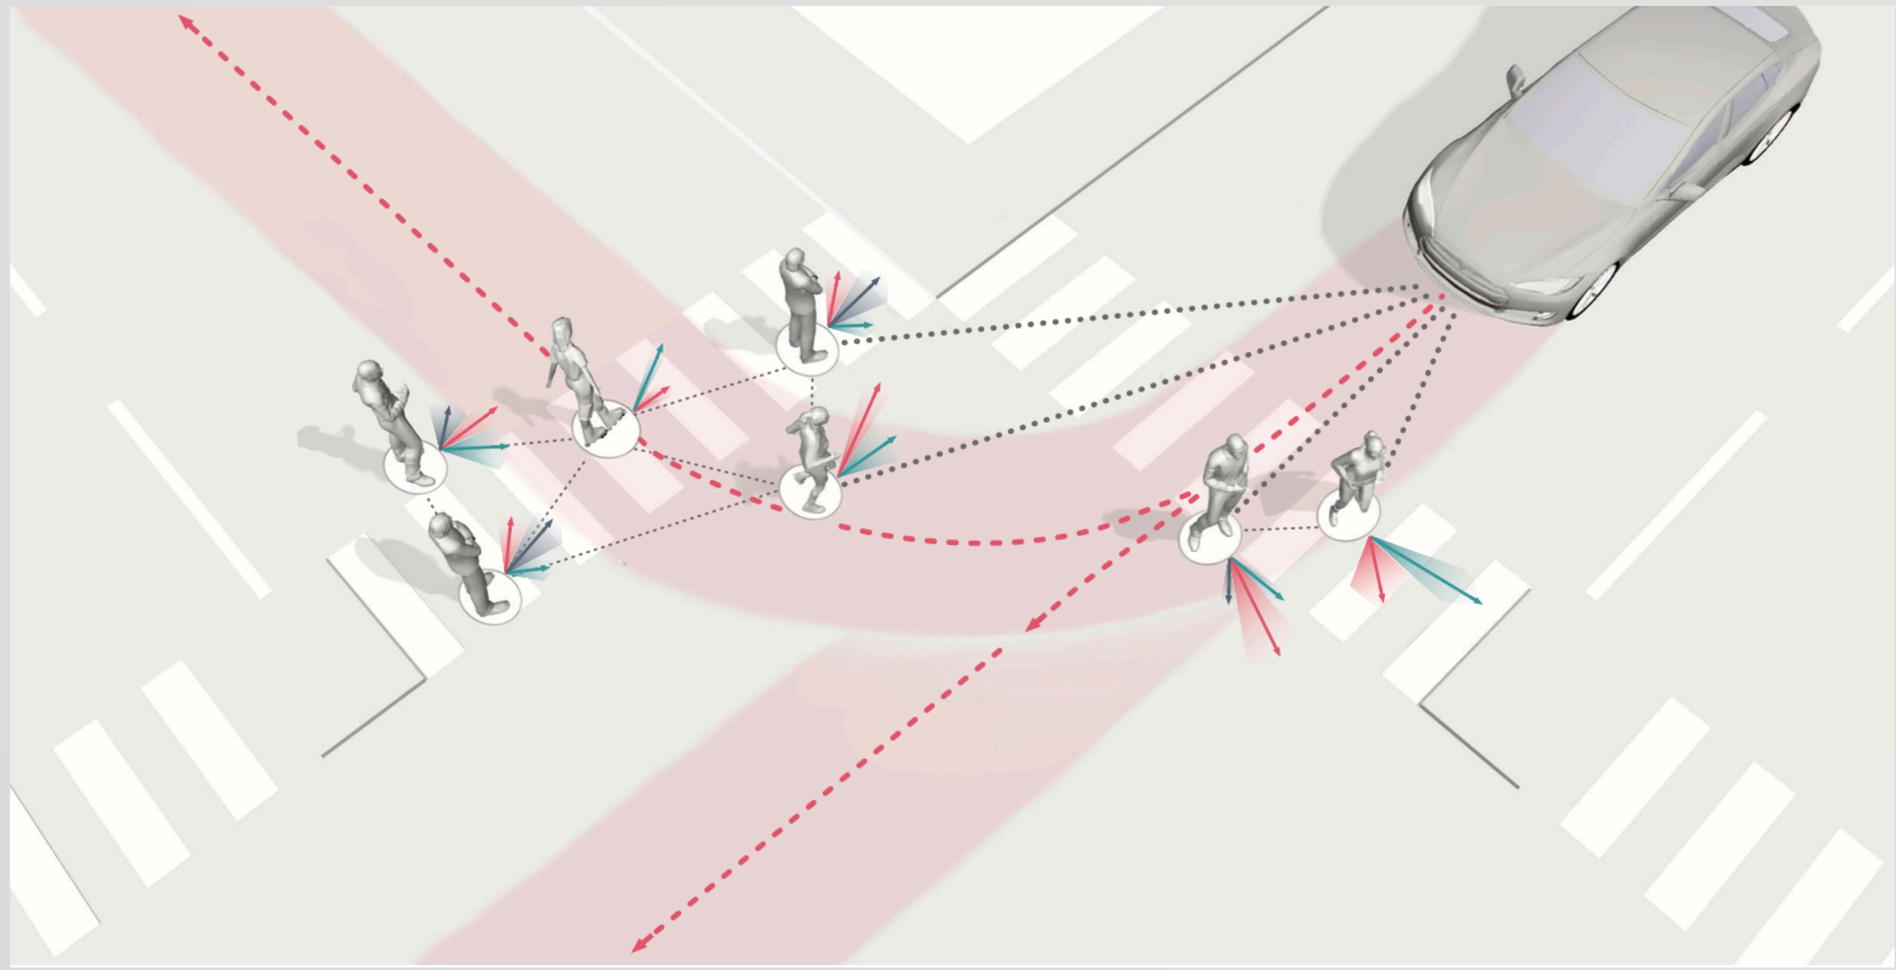

Setup

- Each player has double integrator dynamics and some reference trajectory
- The follower has three possible intentions: tracking (T) the leader, indifferent (I) or avoiding (A).
- The leader either assumes the one of the three follower intentions or an adaptive (Ad) belief given by a multi model estimator
- Randomize over initial conditions and reference trajectories

LQ collision avoidance example

Optimality of wrong beliefs

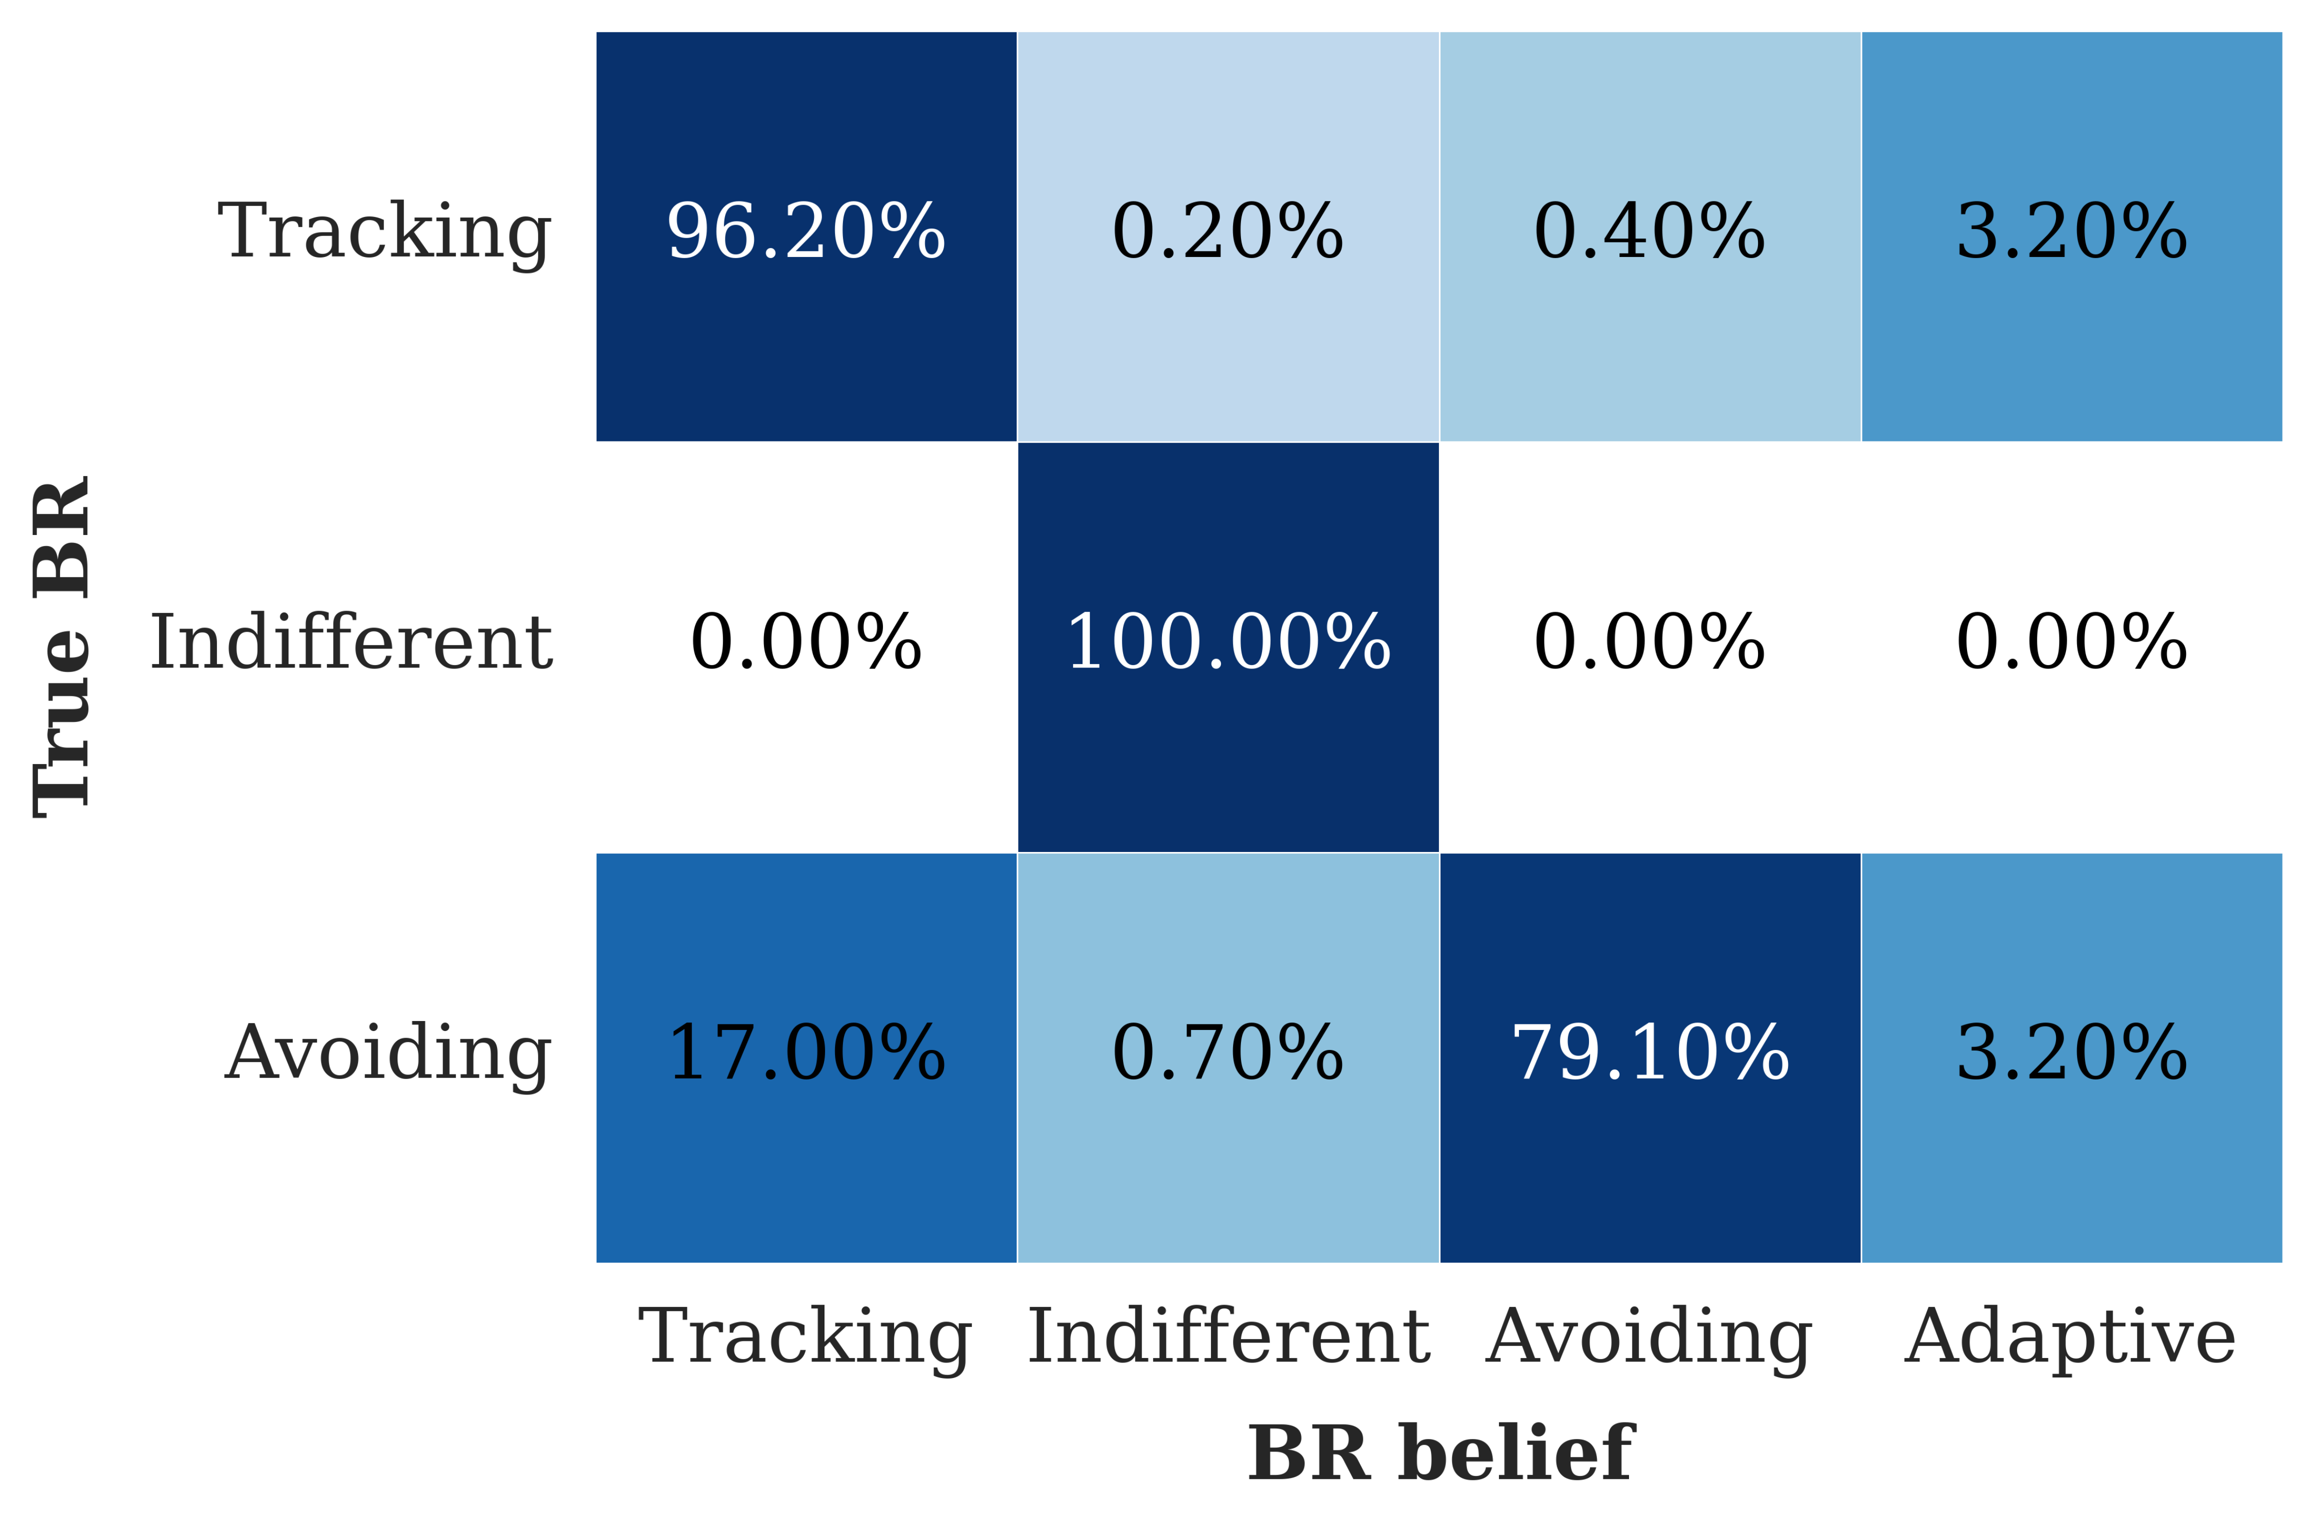

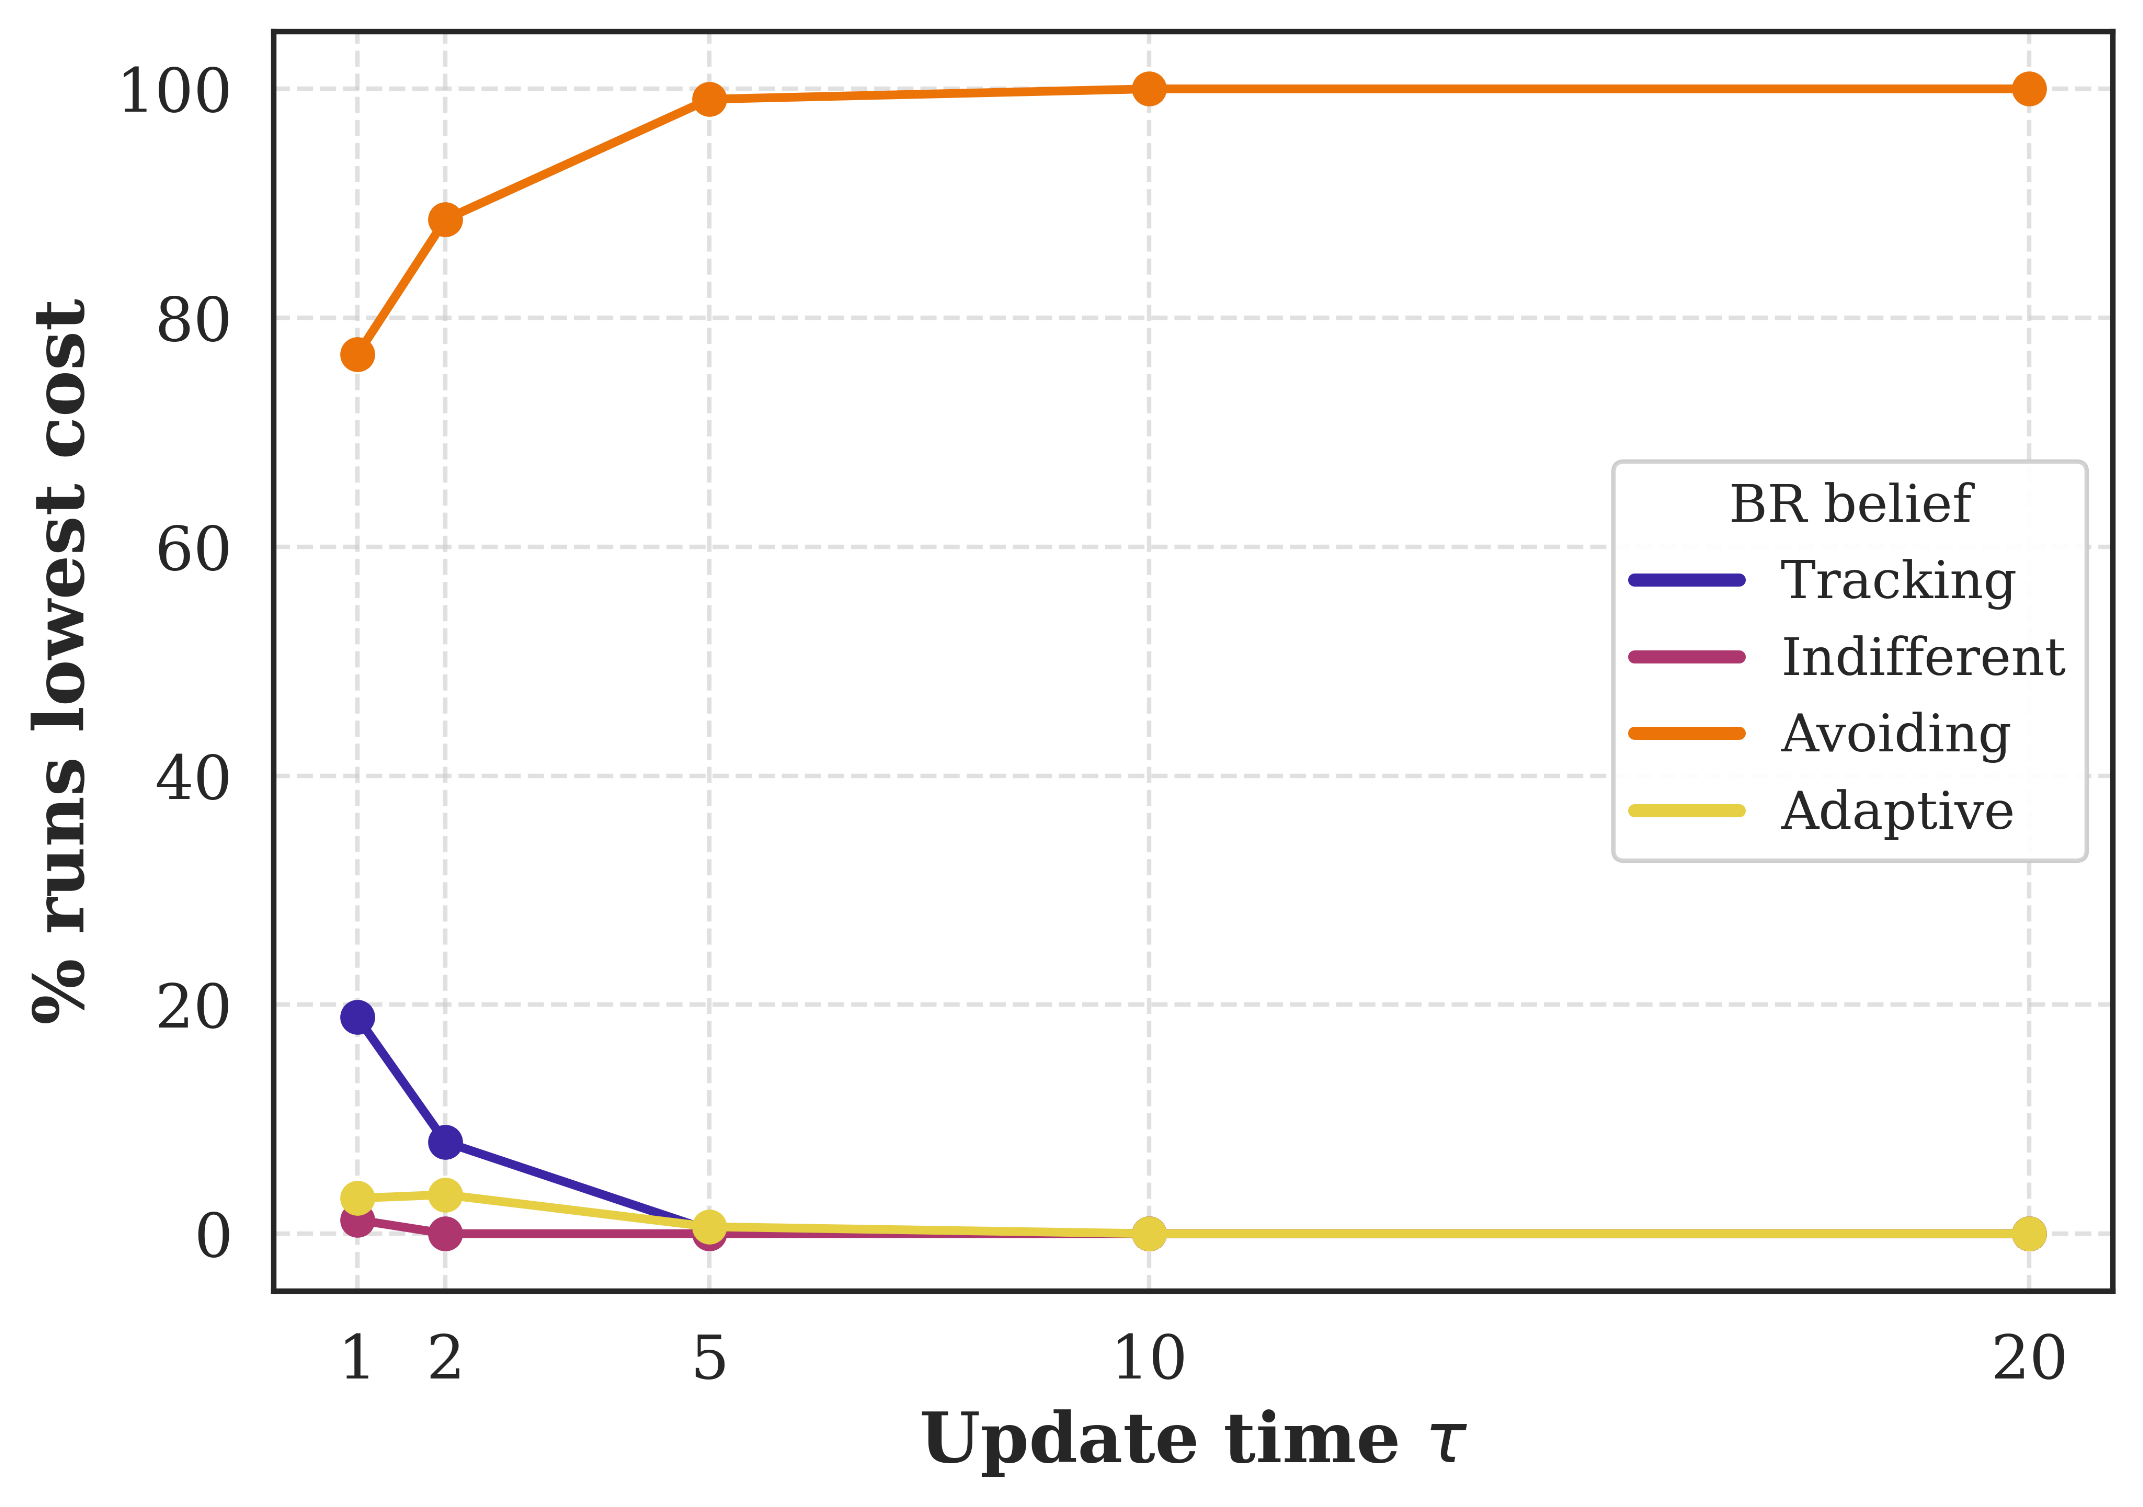

Key insight:

- OLSE is not optimal under leader strategy updates

What is optimal?

- Prior work on time-inconsistent control do not achieve time-consistent optimality
- Common in economics, where follower trust is measured against leader's commitment-breaking behavior

Illustrative example

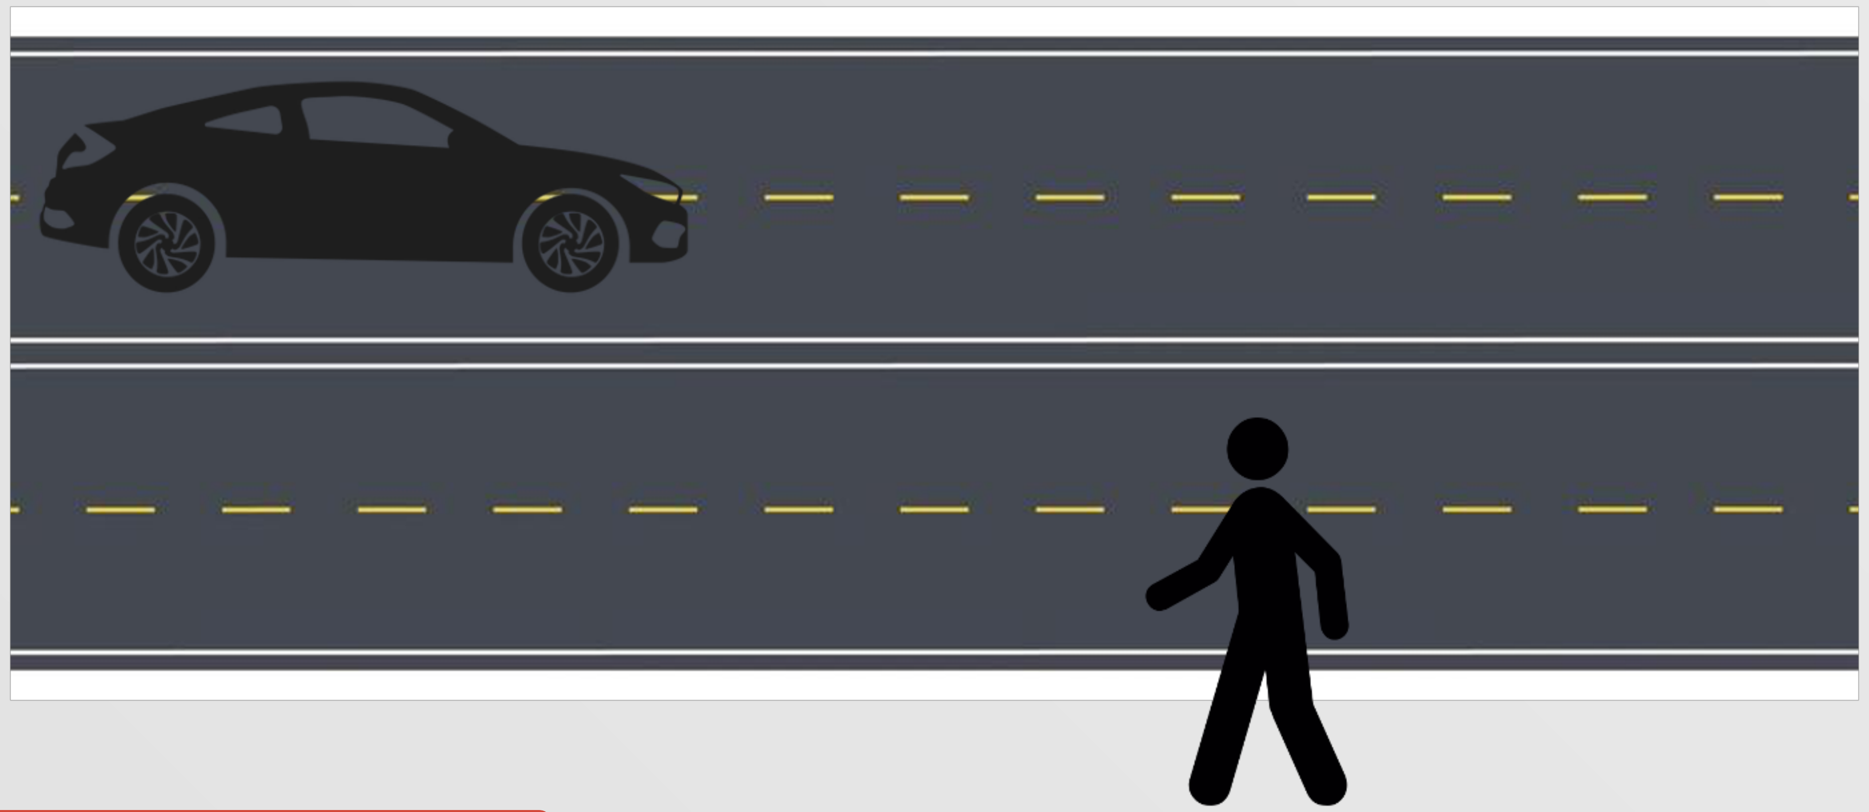

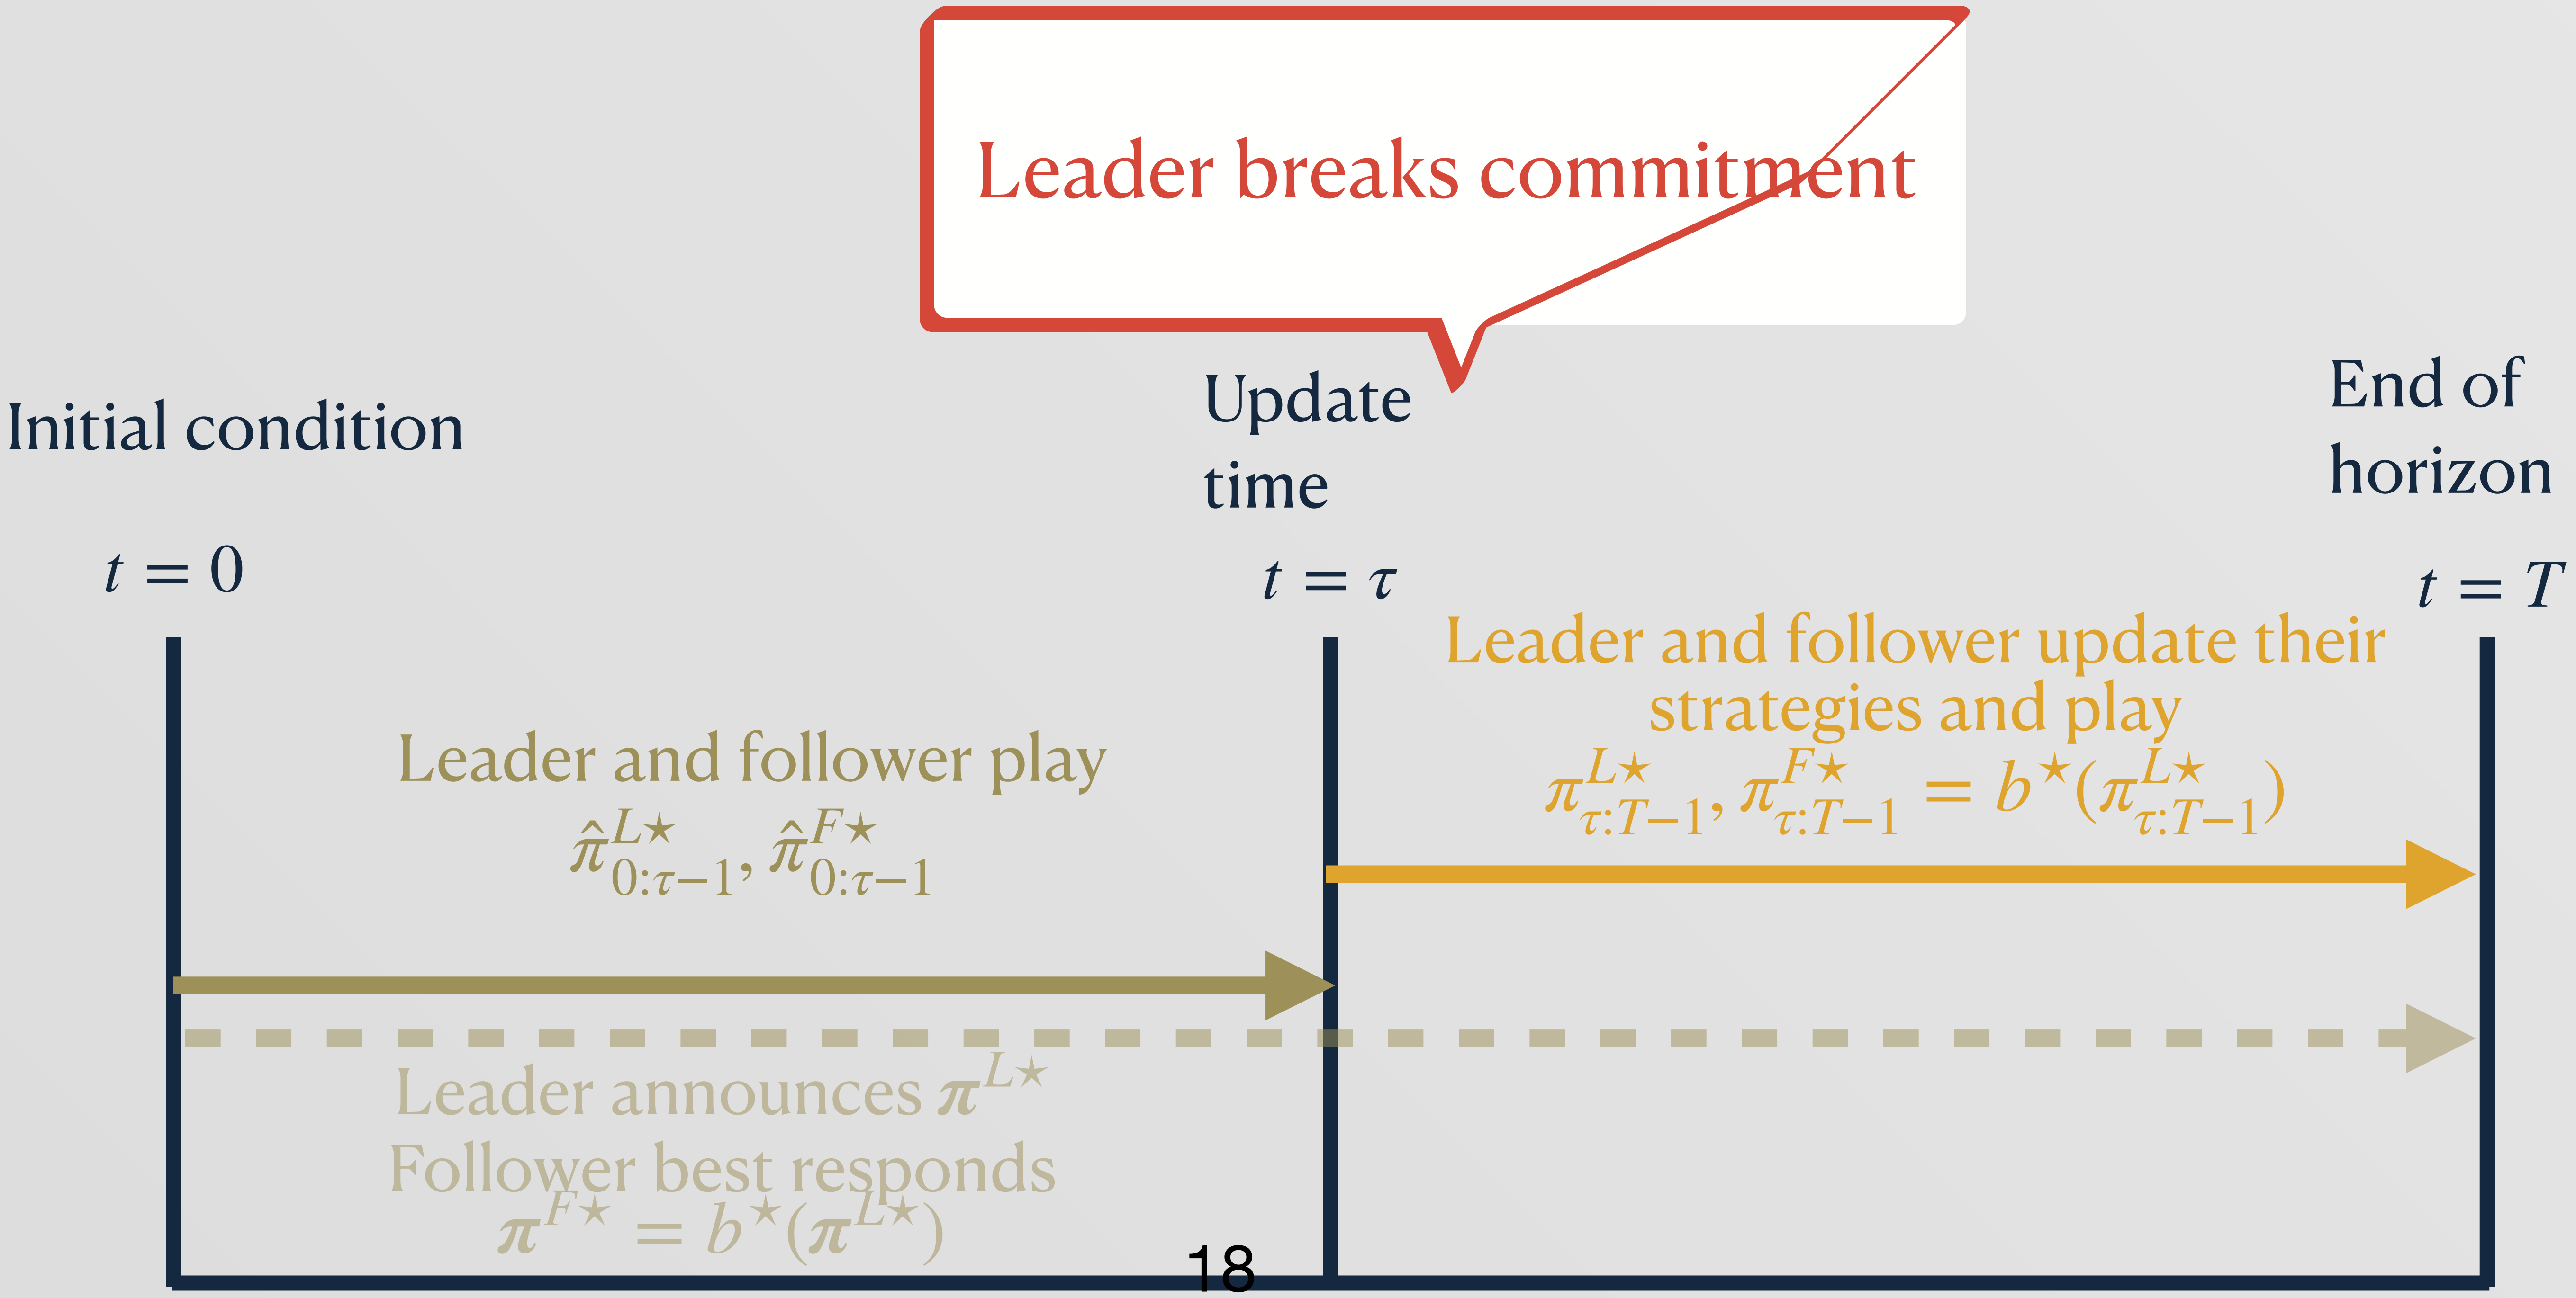

Def: commitment-breaking Stackelberg strategy: $\hat{u}^L_{0:T-1} \cup u^L_{\tau:T-1}$

At $t = 0$ we optimize over

2T $-$ τ **control** elements

First T elements: announced strategy at $t = 0$

Last $T - \tau$ *elements:* updated strategy at $t = \tau$

$$s.t \quad x_{0:T} = Hx_0 + G^L \begin{bmatrix} \hat{u}^L_{0:\tau-1} \\ u^L_{\tau:T-1} \end{bmatrix} + G^F u^F_{0:T-1}$$

$$u^F_{0:T-1} = \begin{bmatrix} u^F_{0:\tau-1} \\ u^F_{\tau:T-1} \end{bmatrix} = \begin{bmatrix} [I_{\tau-1} \quad 0](\hat{G}^F \hat{u}^L_{0:T-1} + \hat{H}^F x_0) \\ \hat{G}^F_{\tau:T-1} u^L_{\tau:T-1} + \hat{H}^F_{\tau:T} x_\tau \end{bmatrix}$$

The announced strategy affects the response of the follower over $[0, \tau - 1]$

Stackelberg Dynamic Game with Updates

If the leader knows when it will break commitment next, we can explicitly model leader's **entire strategy profile**.

Def: commitment-breaking Stackelberg strategy: $\hat{u}^L_{0:T-1} \cup u^L_{\tau:T-1}$

At $t = 0$ we optimize over **$2T - \tau$ control** elements

First T elements: announced strategy at $t = 0$

Last $T - \tau$ elements: updated strategy at $t = \tau$

time-consistent by construction

$$\min_{\hat{u}^L \in \mathcal{U}^L_{0:T-1}} \left[J^L_{0:\tau-1}(x_{0:\tau-1}, \hat{u}^L_{0:\tau-1}, u^F_{0:\tau-1}) + \min_{u^L_{\tau:T-1} \in \mathcal{U}^L_{\tau:T-1}} J^L_{\tau:T}(x_{\tau:T}, u^{L'}_{\tau:T-1}, b^\star(u^{L'}_{\tau:T-1})) \right]$$

$$s.t \quad x_{0:T} = H x_0 + G^L \begin{bmatrix} \hat{u}^L_{0:\tau-1} \\ u^L_{\tau:T-1} \end{bmatrix} + G^F u^F_{0:T-1}$$

$$u^F_{0:T-1} = \begin{bmatrix} u^F_{0:\tau-1} \\ u^F_{\tau:T-1} \end{bmatrix} = \begin{bmatrix} [I_{\tau-1} \quad 0](\hat{G}^F \hat{u}^L_{0:T-1} + \hat{H}^F x_0) \\ \hat{G}^F_{\tau:T-1} u^L_{\tau:T-1} + \hat{H}^F_{\tau:T} x_\tau \end{bmatrix}$$

The announced strategy affects the response of the follower over $[0, \tau - 1]$

Stackelberg Dynamic Game with Updates

If the leader knows when it will break commitment next, we can explicitly model leader's **entire strategy profile**.

Thm [Rodriguez'26]: the optimal solution $(\hat{u}^{L\star}, u^{L\star})$ to the minimization problem below achieves the minimal leaders cost across all commitment-breaking Stackelberg strategies

time-consistent by construction

$$\min_{\hat{u}^L \in \mathcal{U}^L_{0:T-1}} \left[J^L_{0:\tau-1}(x_{0:\tau-1}, \hat{u}^L_{0:\tau-1}, u^F_{0:\tau-1}) + \min_{u^L_{\tau:T-1} \in \mathcal{U}^L_{\tau:T-1}} J^L_{\tau:T}(x_{\tau:T}, u^{L'}_{\tau:T-1}, b^\star(u^{L'}_{\tau:T-1})) \right]$$

$$\text{s.t} \quad x_{0:T} = Hx_0 + G^L \begin{bmatrix} \hat{u}^L_{0:\tau-1} \\ u^L_{\tau:T-1} \end{bmatrix} + G^F u^F_{0:T-1}$$

$$u^F_{0:T-1} = \begin{bmatrix} u^F_{0:\tau-1} \\ u^F_{\tau:T-1} \end{bmatrix} = \begin{bmatrix} [I_{\tau-1} \quad 0](\hat{G}^F \hat{u}^L_{0:T-1} + \hat{H}^F x_0) \\ \hat{G}^F_{\tau:T-1} u^L_{\tau:T-1} + \hat{H}^F_{\tau:T} x_\tau \end{bmatrix}$$

The announced strategy affects the response of the follower over $[0,\tau-1]$

Stackelberg Dynamic Game with Updates

Key insight:

- Leader can commit to a strategy $\hat{u}^L$ that incurs high cost, knowing that the strategy won't be followed through

LQ collision avoidance example

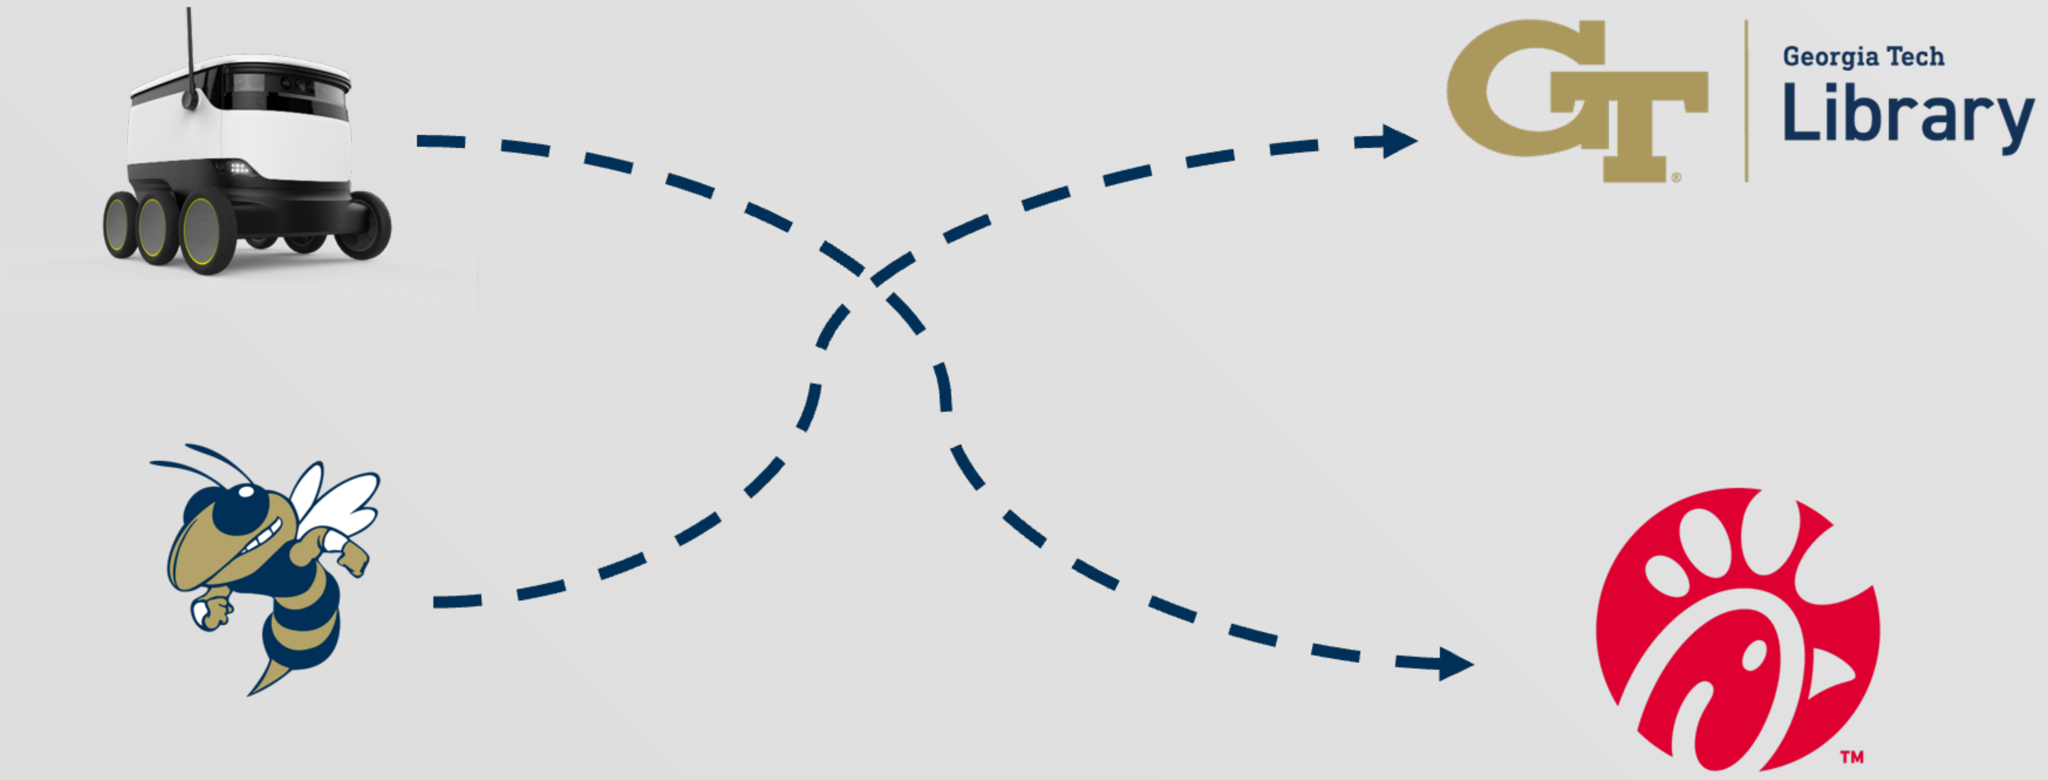

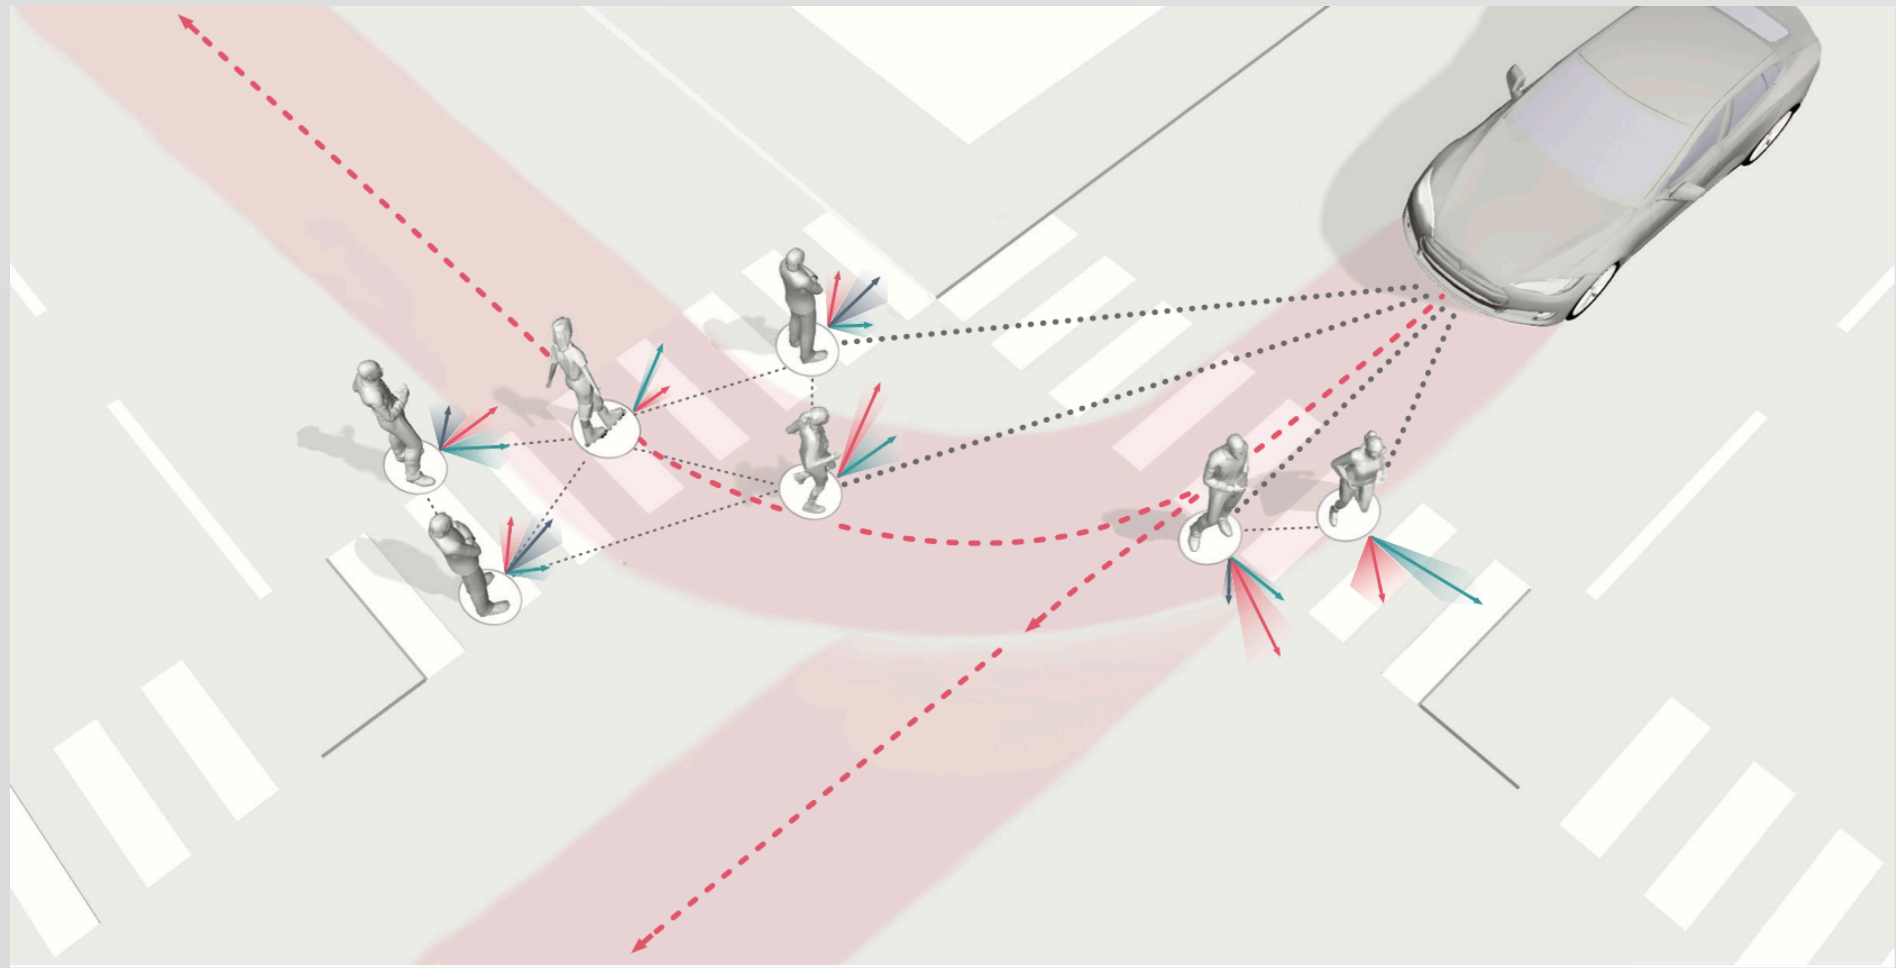

Setup

- Each player has double integrator dynamics and some reference trajectory
- The follower has three possible intentions: tracking (T) the leader, indifferent (I) or avoiding (A).
- The leader either assumes the one of the three follower intentions or an adaptive (Ad) belief given by a multi model estimator
- Randomize over initial conditions and reference trajectories

LQ collision avoidance example

Optimality of commitment-breaking Stackelberg strategy (CBSS)

True BR \ BR belief	Tracking	Indifferent	Avoiding	Adaptive
Tracking	96.20%	0.20%	0.40%	3.20%
Indifferent	0.00%	100.00%	0.00%	0.00%
Avoiding	17.00%	0.70%	79.10%	3.20%

True BR \ Control played	Tracking BR belief	Indifferent BR belief	Avoiding BR belief	Adaptive BR belief	Commitment breaking
Tracking	0.00%	0.00%	0.00%	0.00%	100.00%
Indifferent	0.00%	0.00%	0.00%	0.00%	100.00%
Avoiding	0.00%	0.00%	0.00%	0.00%	100.00%

Commitment breaking Stackelberg strategy

CBSS vs. OLSE update vs. OLSE no update

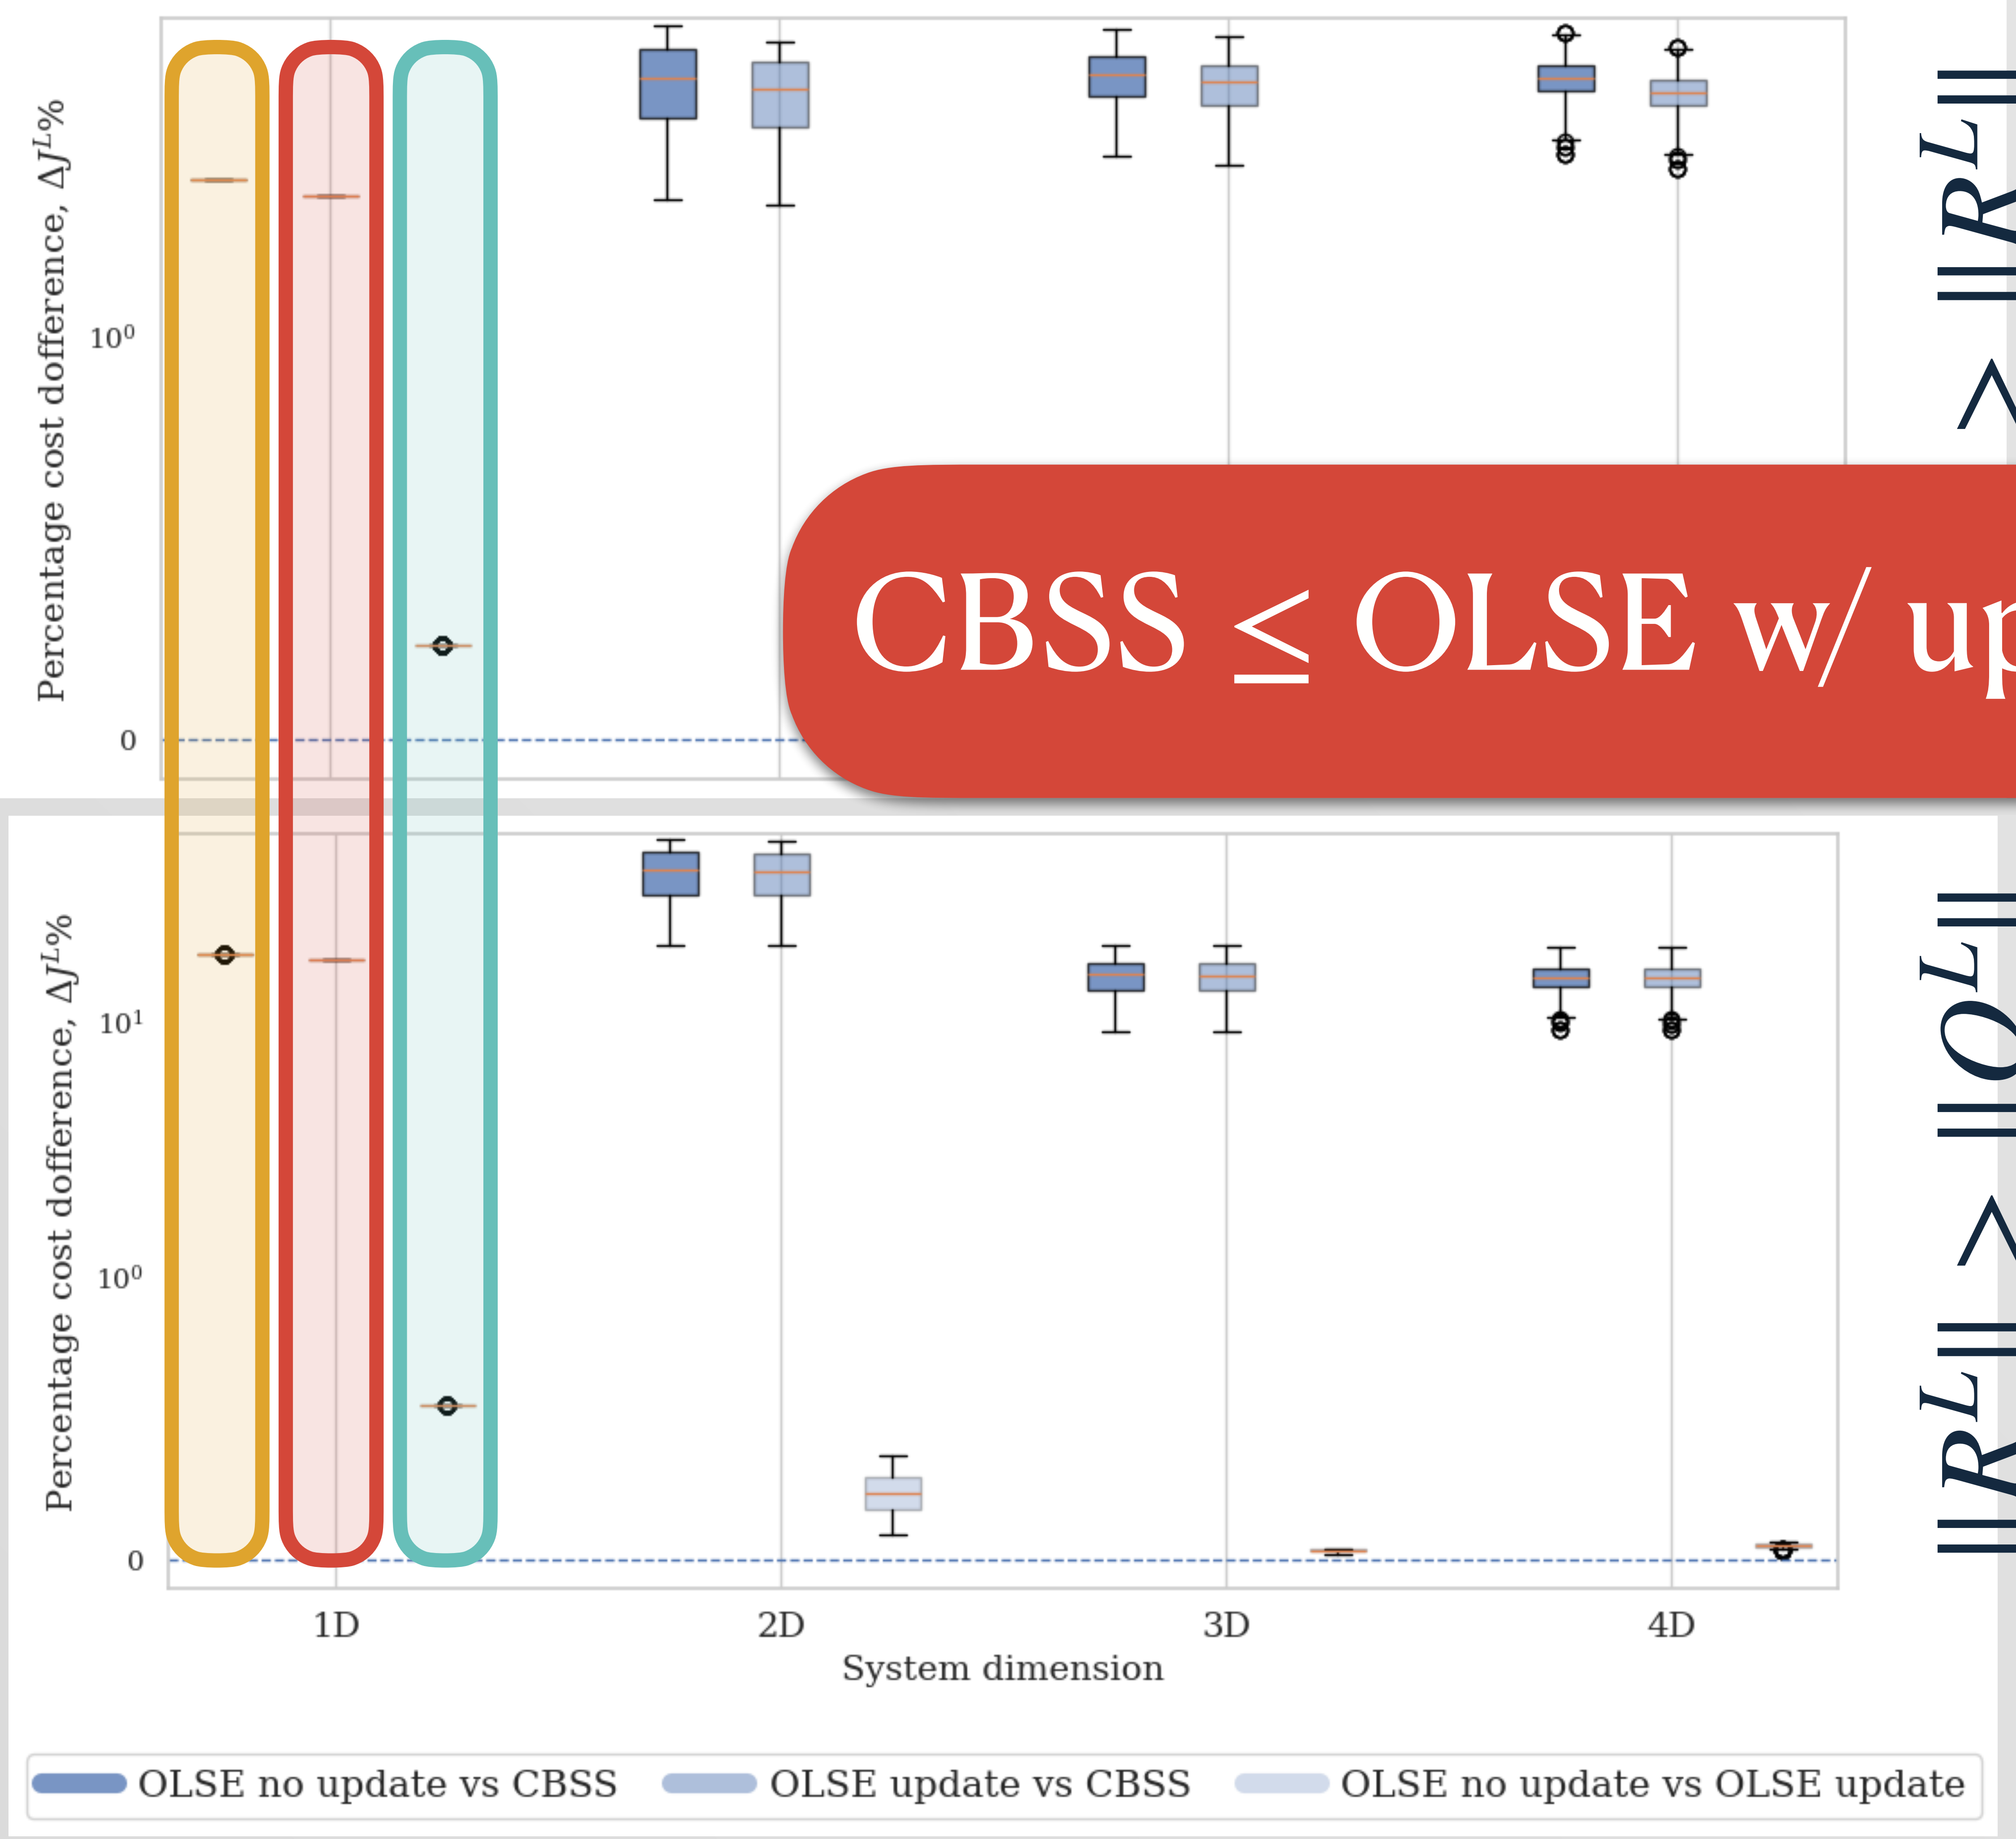

$\|R^L\| > \|R^L\|$

$\|R^L\| > \|Q^L\|$

OLSE without updating vs. CBSS

"Naive" OLSE with update vs. CBSS

OLSE with no update vs "naive" OLSE with update

CBSS ≤ OLSE w/ update ≤ OLSE w/o update

$|J_{0:T}^{L,CBSS}|$

- For each state dimension, dynamics and cost matrices are fixed

$i = \{\text{OLSE w/o update, OLSE w/ update}\}$

- $\|Q^L\|, \|Q^F\|_\infty$ remain constant

Third box:

- Initial state x_0 randomized

$$\frac{J_{0:T}^{L,\text{OLSE w/o update}} - J_{0:T}^{L,\text{OLSE w/ update}}}{|J_{0:T}^{L,\text{OLSE w/ update}}|} \cdot 100$$

Commitment breaking Stackelberg strategy

Effect of update time on CBSS performance

If the leader is playing CBSS, what is the best update time?

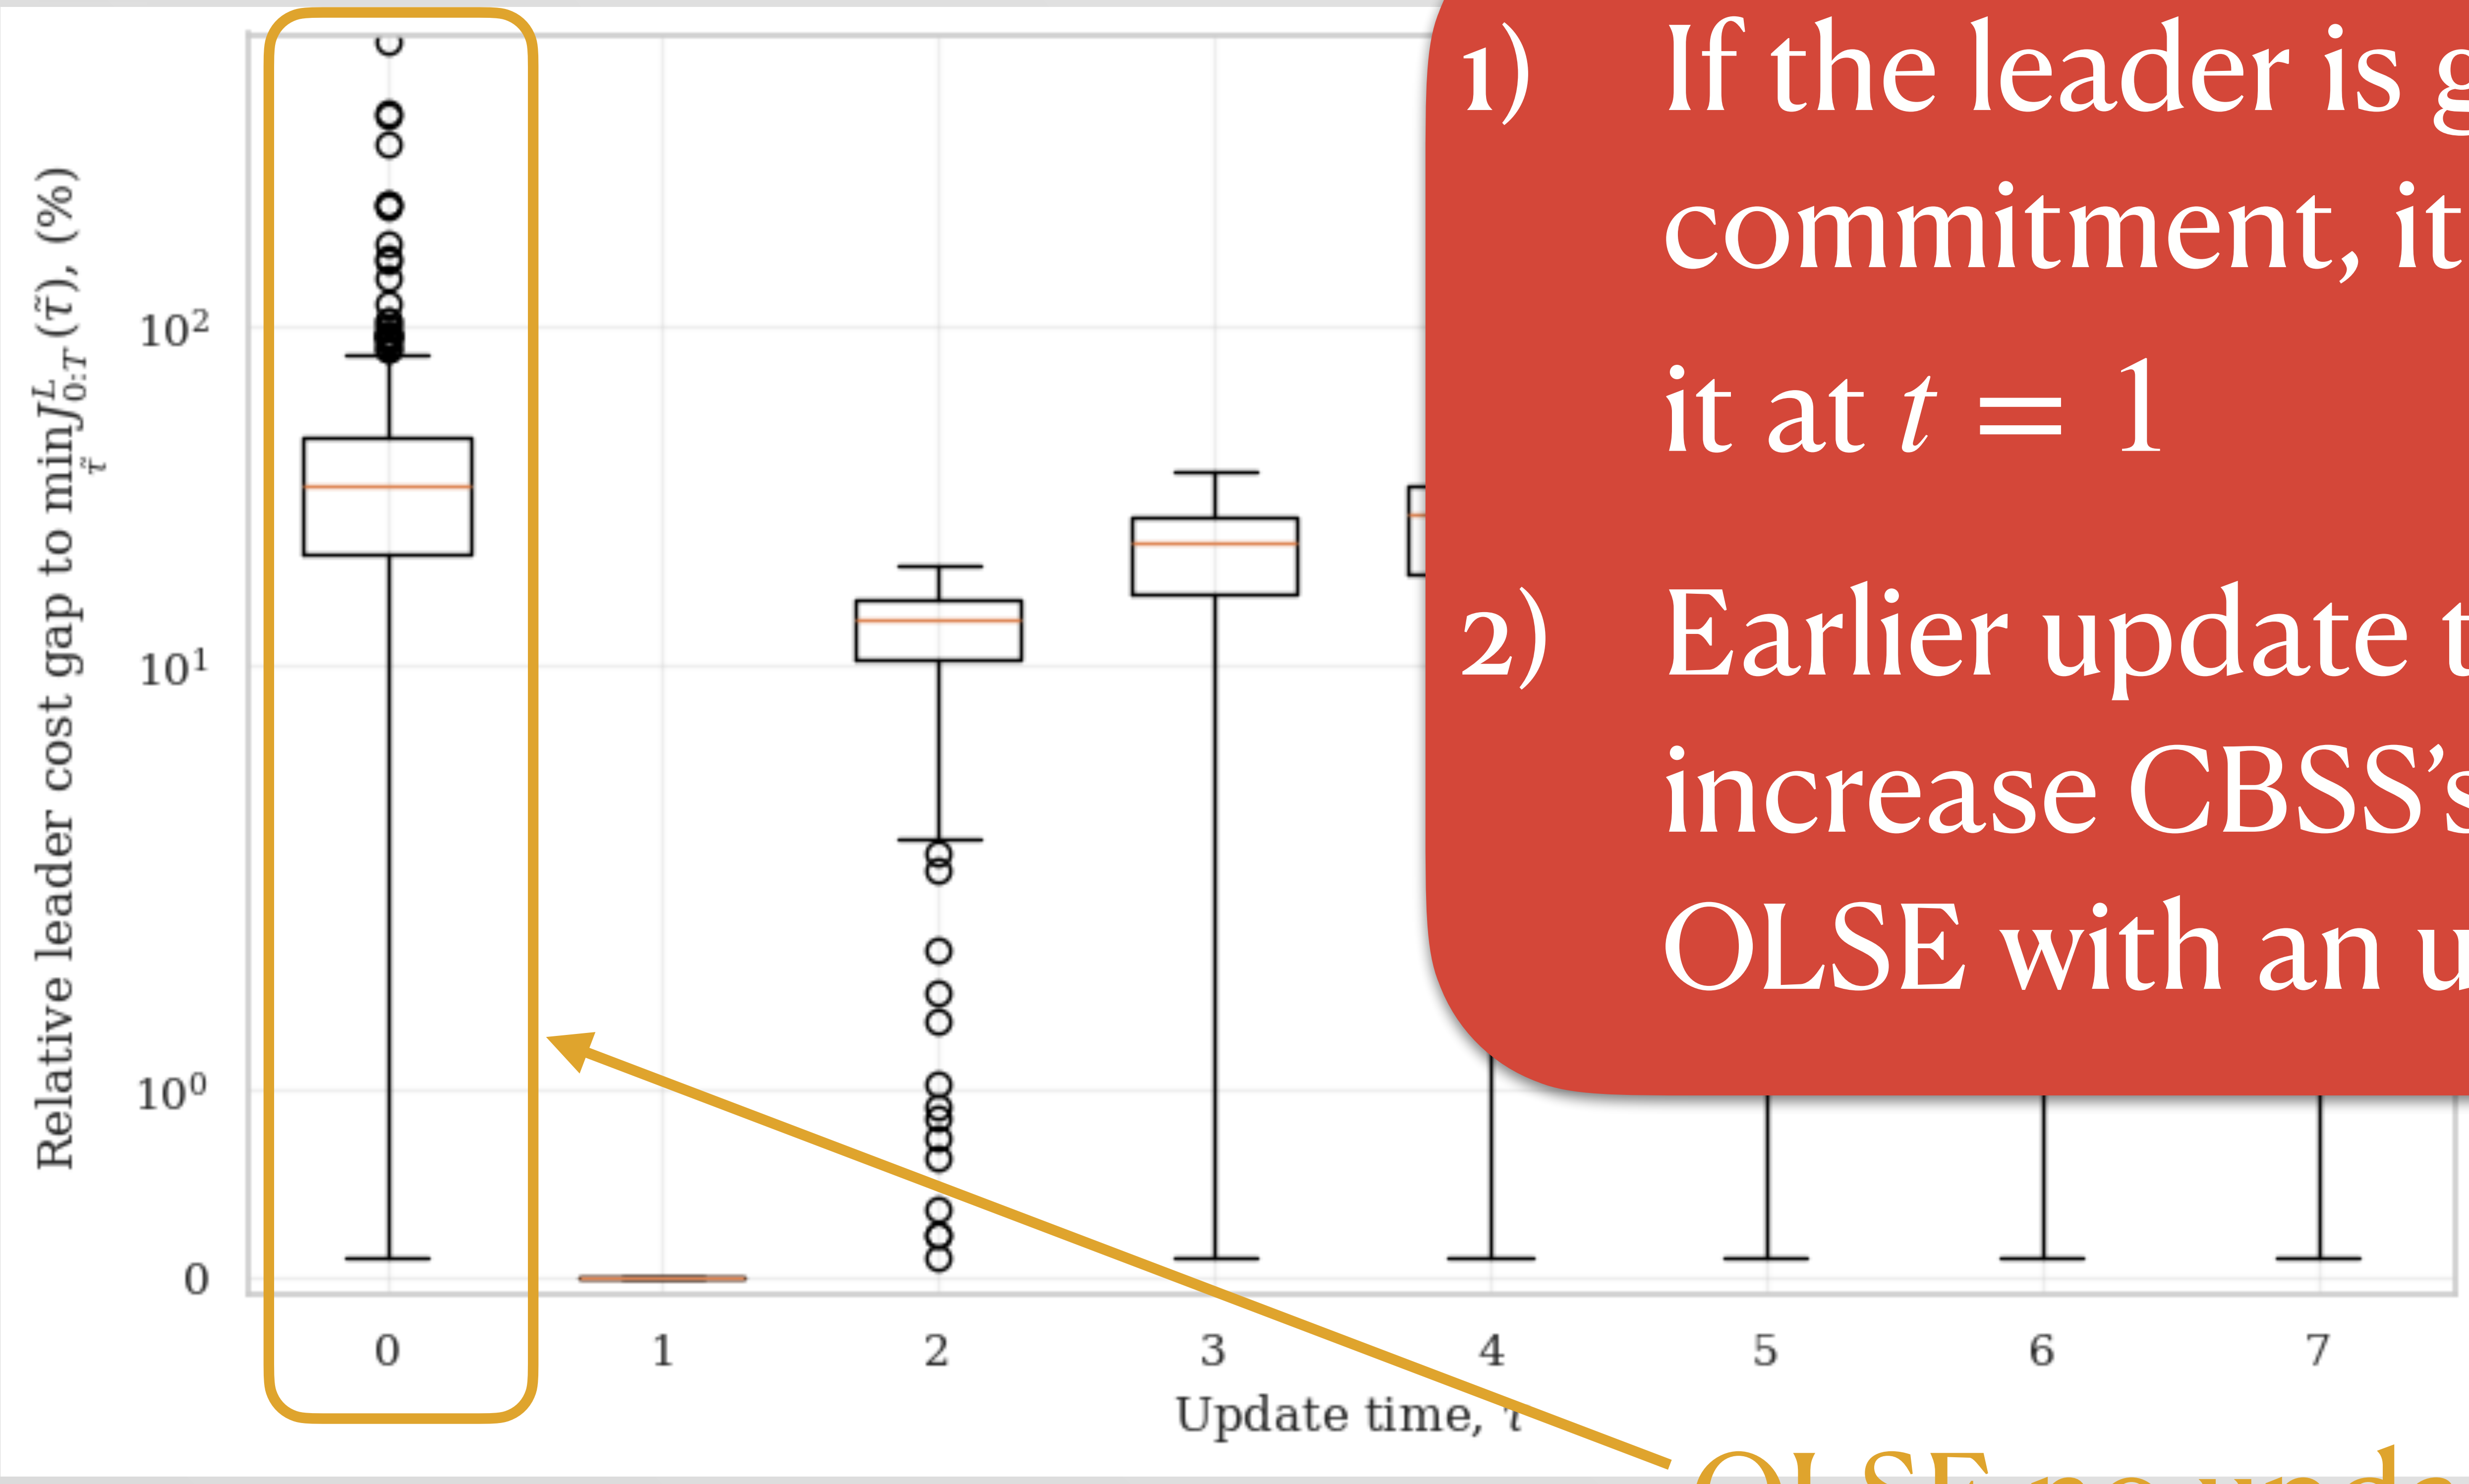

$$\frac{J^L_{0:T}(\tau) - \min_{\tilde{\tau}} J^L_{0:T}(\tilde{\tau})}{|\min_{\tilde{\tau}} J^L_{0:T}(\tilde{\tau})|} \cdot 100$$

$$\frac{J^{L,OLSE}_{0:T} - J^{L,CBSS}_{0:T}}{|J^{L,CBSS}_{0:T}|} \cdot 100$$

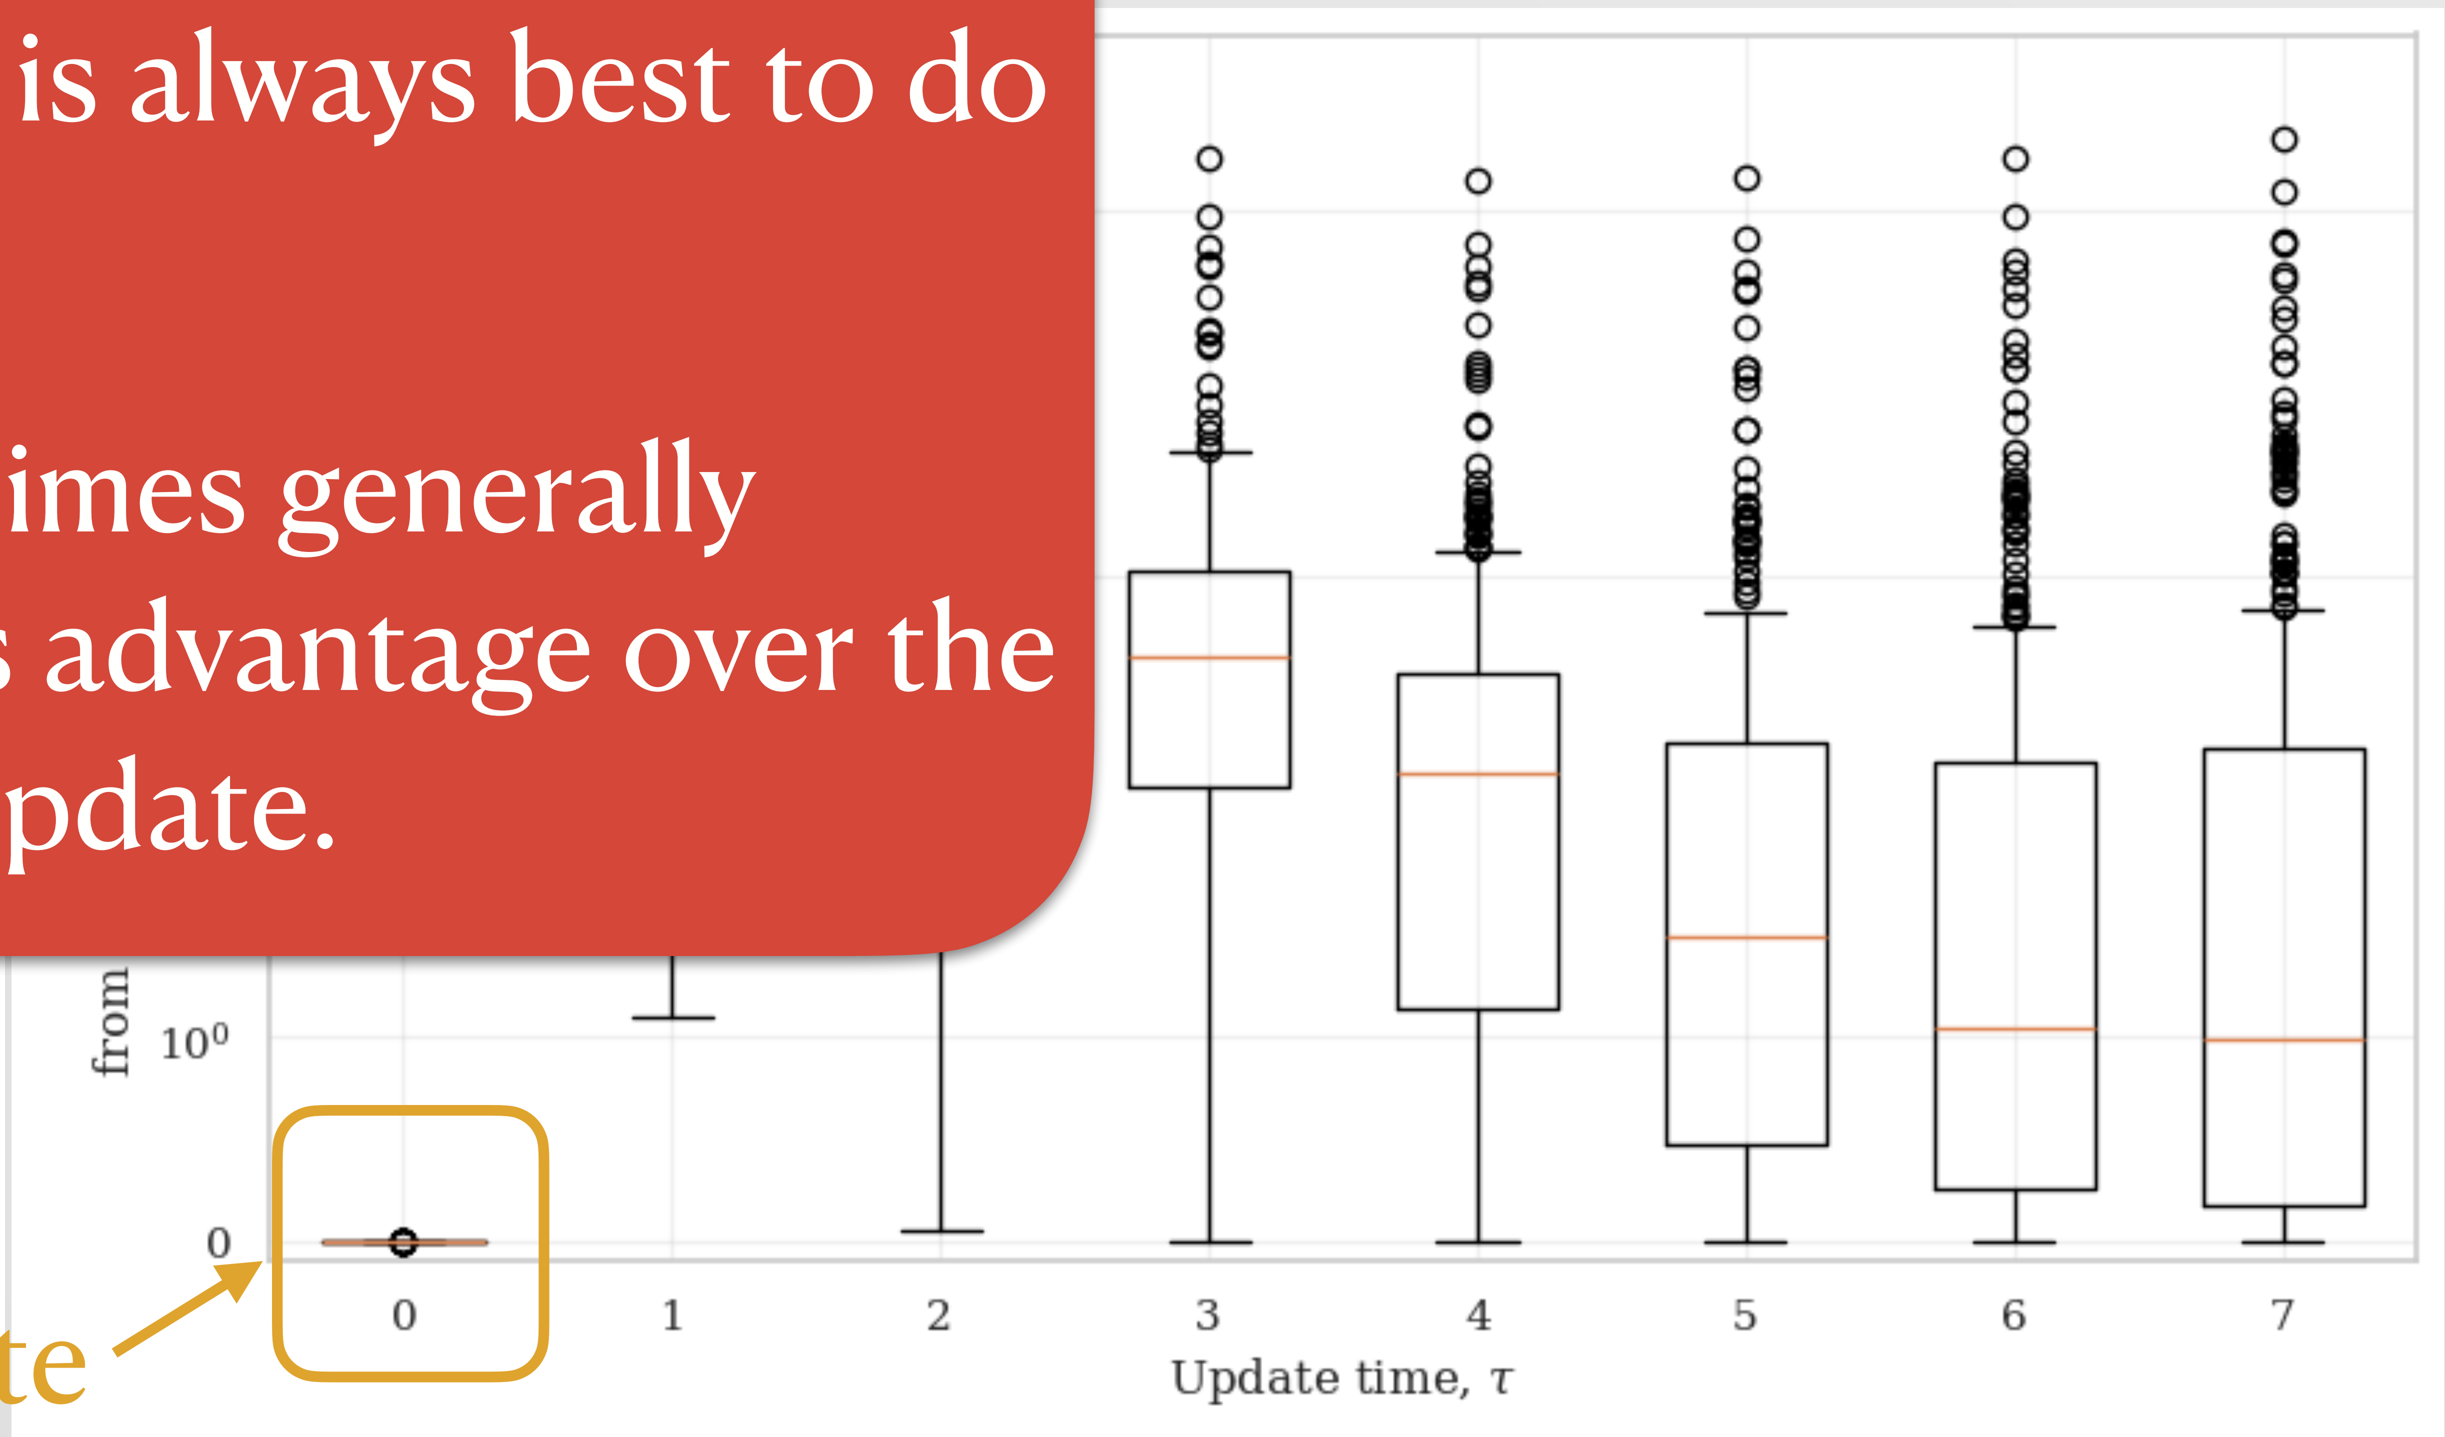

1) If the leader is going to break commitment, it is always best to do it at $t = 1$
2) Earlier update times generally increase CBSS's advantage over the OLSE with an update.

OLSE no update

What update time gives the greatest advantage to CBSS against naively solving OLSE and updating?

Commitment breaking Stackelberg strategy

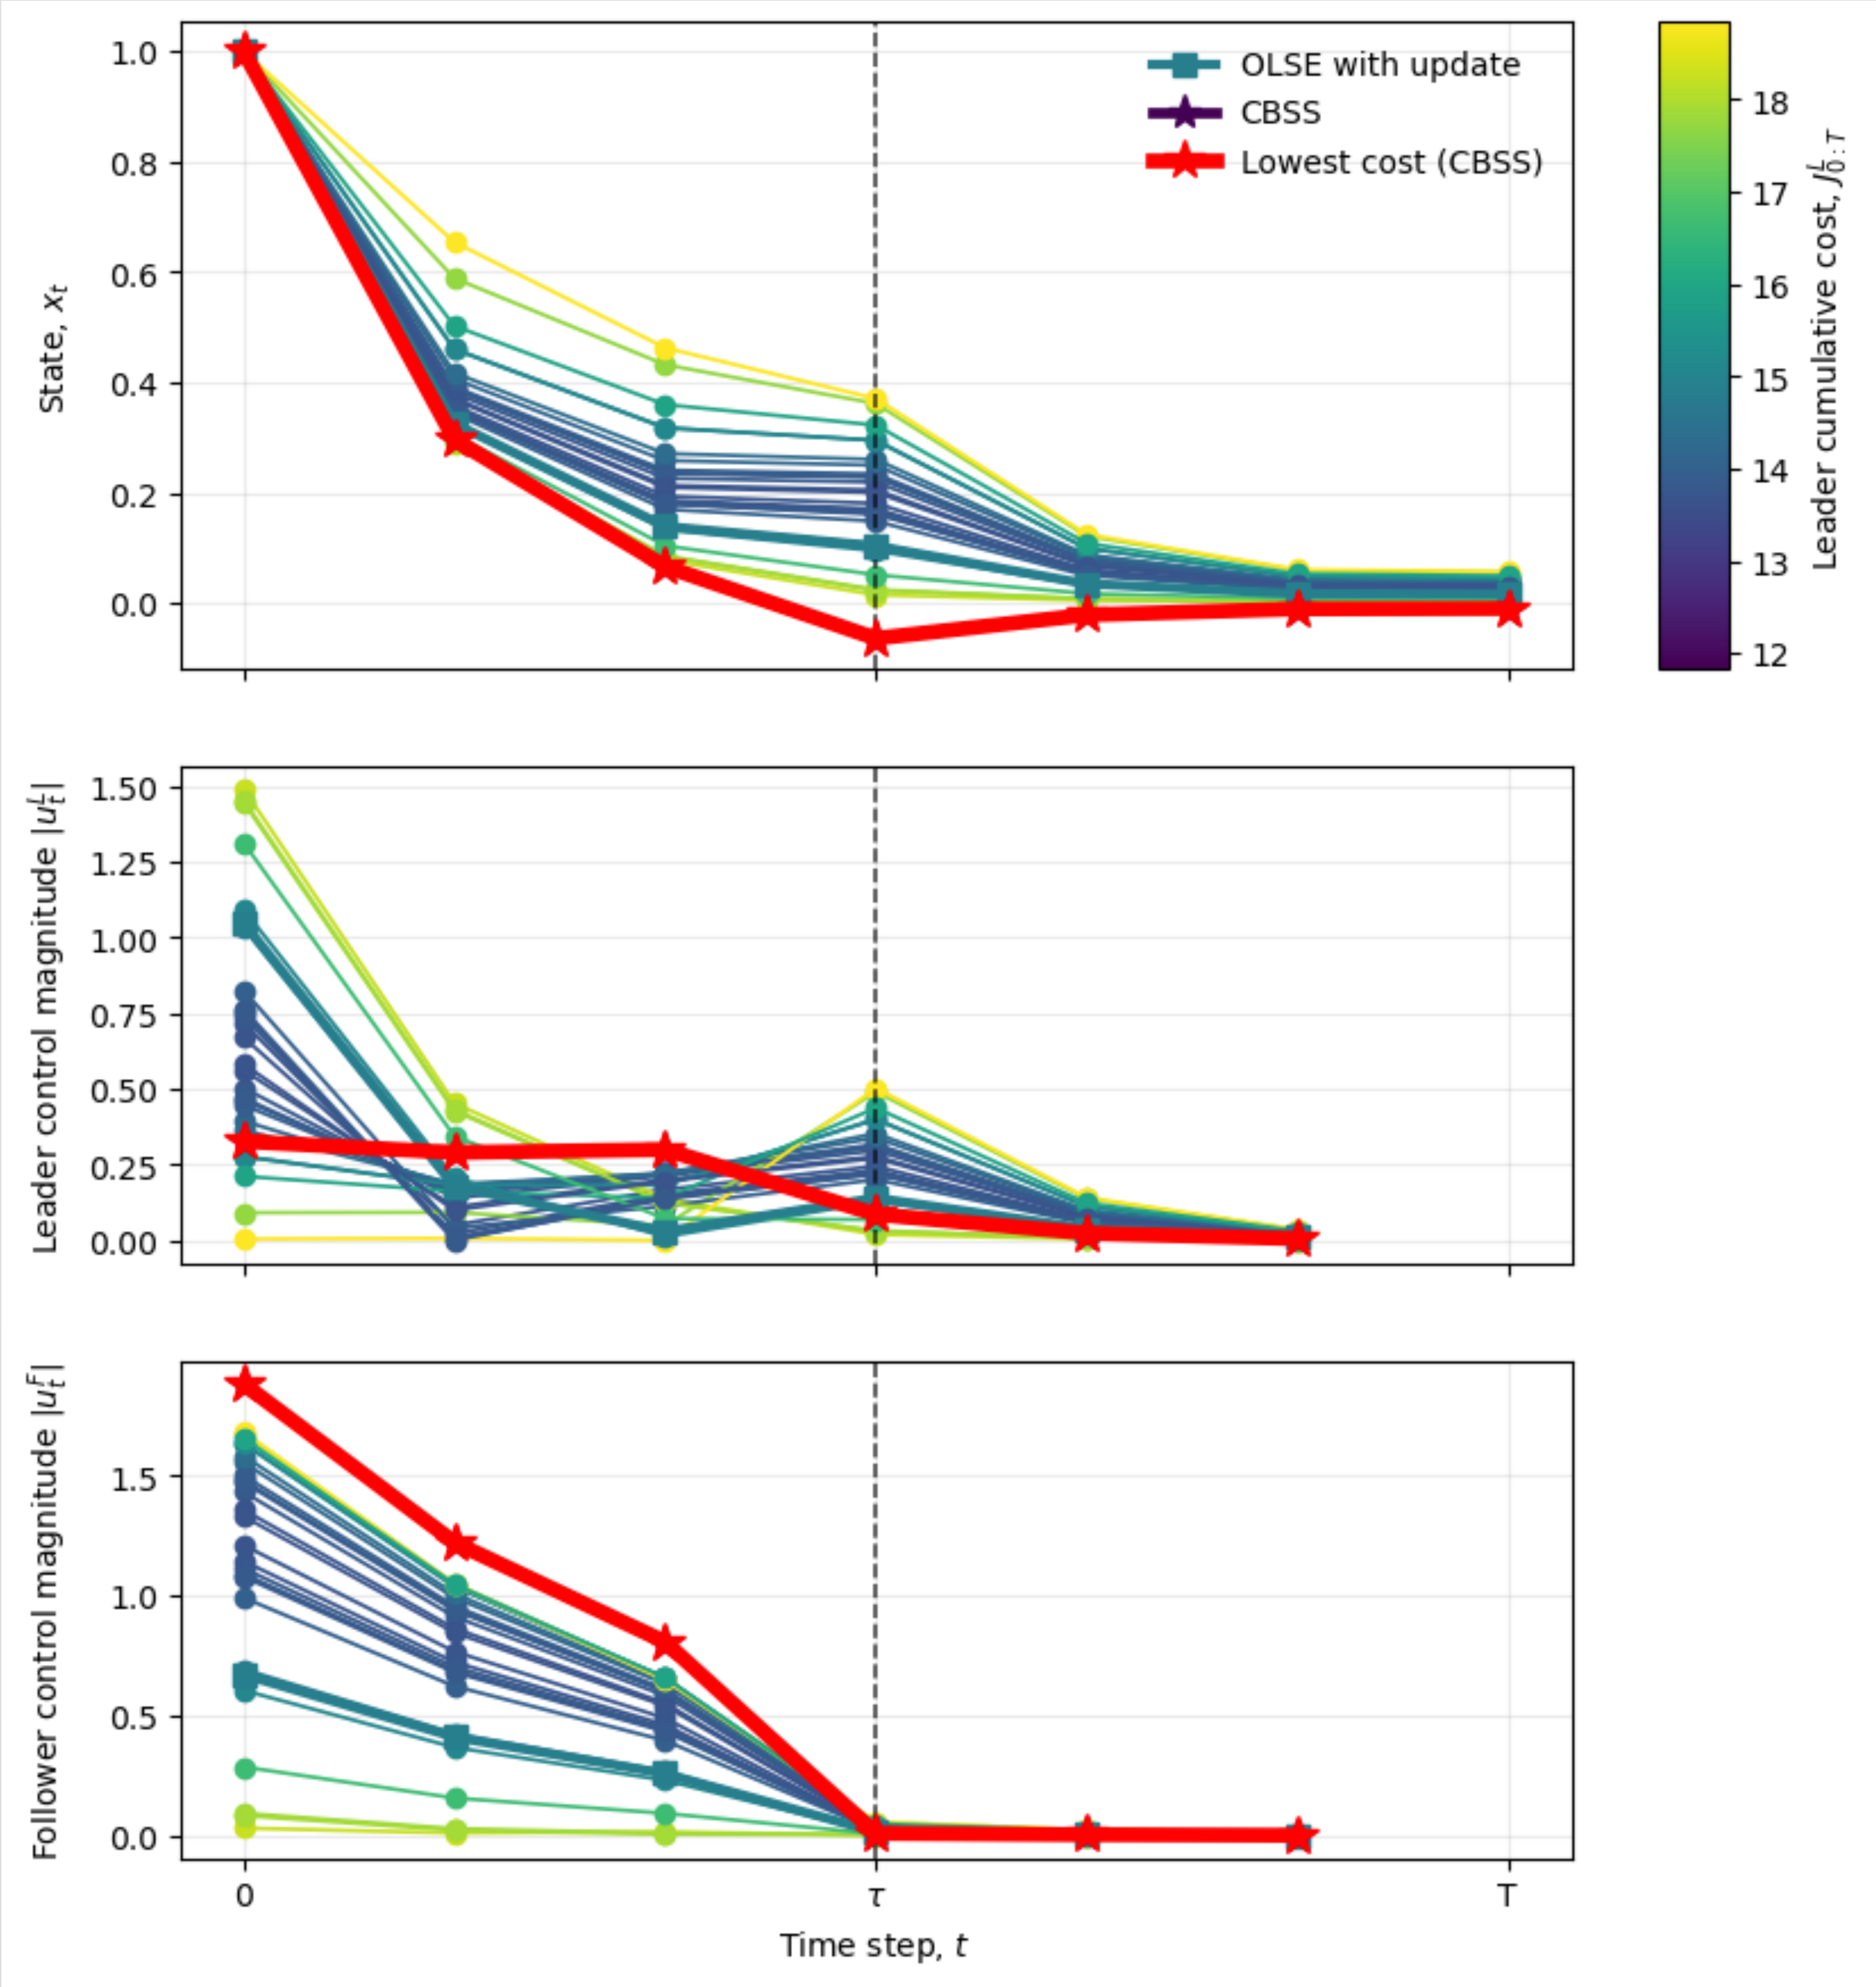

Benchmarking state behavior with CBSS

Acknowledgments/references

Language Conditioning Improves Accuracy of Aircraft Goal Prediction in Untowered Airspace
To appear in ICRA 2026, Vienna

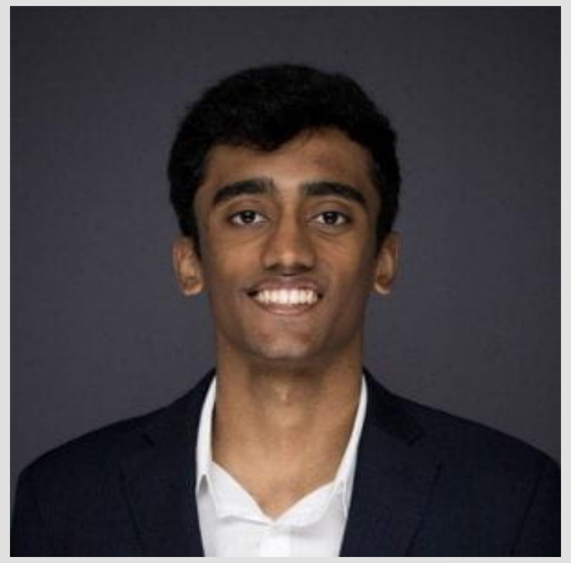

Sundhar Sangeetha (MS)

Owen Daum (High school)

Shreyas Kousik (PI)

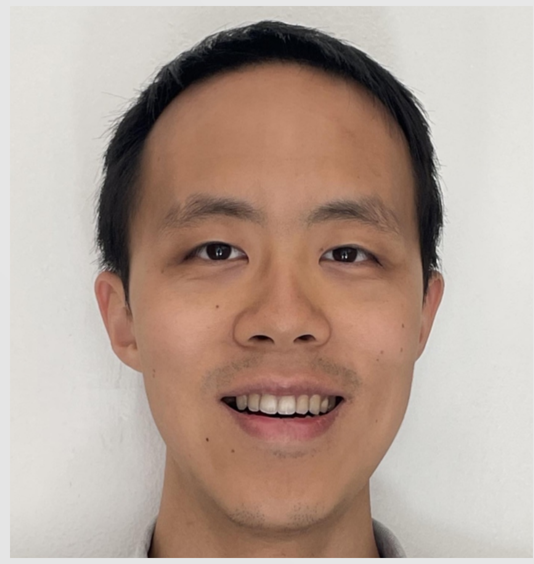

Frank (C.Y.) Chiu (PI)

When the Correct Model Fails: The Optimality of Stackelberg Equilibria with Follower Intention Updates
To appear in ECC 2026, Reykjavik, journal in preparation

Cayetana Salinas Rodriguez (MS)

Jonathan Rogers (PI)