University of Cyprus

Linear Quadratic Tracking Control of Hidden Markov Jump Linear Systems Subject to Ambiguity

I. Tzortzis, C.N. Hadjicostis, and C.D. Charalambous
University of Cyprus
Department of Electrical and Computer Engineering

60th CONFERENCE ON DECISION AND CONTROL, DECEMBER 13-17, 2021

OUTLINE

- Introduction and Motivation
- Problem Formulation
- MinMax Stochastic Control
- Numerical Example
- Concluding Remarks

INTRODUCTION - MOTIVATION

Introduction - Motivation

High-Level Objective

Develop an LQ tracking control approach applicable to hidden MJLS with time-varying transition p.d., capable of capturing and restricting the influence of uncertainty of the underlying MC on the performance of the optimal controller

A Short Review

- Linear dynamical systems subject to abrupt changes in their operating modes, are known as Markov Jump Linear Systems (MJLS)
  - Tracking control systems, networked control systems, fault-tolerant control systems, etc
- In practice, the designer of MJLS faces two different issues:
  - Information structure available to the controller
  - Model and Markov uncertainty
- Our focus is for the scenario in which: (a) the Markov states are not accessible to the controller, and (b) the underlying MC is uncertain

Introduction - Motivation

Methodology

- State estimation problem addressed by employing the one-step forward Viterbi algorithm
  - Implemented as soon as a new observation of the state is obtained
- Robust control problem is addressed via minimax optimization
  - Using a Dynamic Programming approach
- The influence of uncertainty is addressed through an ambiguous stochastic optimization modelling approach
  - Formed with all p.d. whose discrepancy to a nominal p.d. is within a TV distance metric

Main Contributions

- A systematic modelling approach to deal with uncertainty in the underlying MC
- A maximizing, time-varying, probability distribution
- An optimal control policy with some desired robustness properties

Problem Formulation

Problem Formulation

Model Description

Discrete-time control system:

$$x_{k+1} = A_k(\theta_k)x_k + B_k(\theta_k)u_k, \quad x_1 = x \tag{1}$$

- $x_k \in \mathscr{X}_k$: state process
- $u_k \in \mathscr{U}_k$: control process
- $\{\theta_k, k \geq 1\} \in \Theta$: finite-state, non-homogeneous MC
- $A_k(\theta_k), B_k(\theta_k)$: dynamics and input matrices of compatible dimensions

Assumptions.

- Complete observations of system's state
- $A_k(\theta_k), B_k(\theta_k)$ known for each $\theta_k \in \Theta$ with bounded and measurable entries
- MC state $\{\theta_k, k \geq 1\}$ is not observable

Problem Formulation

Hidden Markov Model

A 5-tuple HMM $= (\Theta, Z, q, e, \pi)$ consisting of:

- A set of hidden states $\{\theta_k, k \geq 1\} \in \Theta$
- A set of observed variables $\{z_k, k \geq 1\} \in Z$
- A time-varying transition p.d. $q_{ij}(k) = P_k(\theta_{k+1} = j \,|\, \theta_k = i)$, $i, j \in \Theta$
- A time-varying output/emission p.d. $e_{is}(k) = P_k(z_k = s \,|\, \theta_k = i)$, $s \in Z$, $i \in \Theta$
- An initial p.d. $\pi(i) = P(\theta_1 = i)$, $i \in \Theta$

Estimation Problem

For an HMM and an observation sequence $z_{1:N} = \{z_1, z_2, \ldots, z_N\}$, we wish to compute the most probable sequence of hidden states, that is,

$$\theta^*_{1:N} = \arg\max_{\theta_{1:N}} P(\theta_{1:N} \,|\, z_{1:N})$$

Problem Formulation

Viterbi Algorithm

Define the recursions:

$$\mu_{k+1}(\theta_{k+1}) \triangleq \max_{\theta_{1:k}} P(\theta_{1:k+1}, z_{1:k+1}) = \max_{\theta_k} P(z_{k+1} \mid \theta_{k+1}) P(\theta_{k+1} \mid \theta_k) \mu_k(\theta_k) \qquad (2a)$$

$$\mu_1(\theta_1) \triangleq P(\theta_1, z_1) = P(z_1 \mid \theta_1) P(\theta_1), \quad k = 1,2,\ldots,N \qquad (2b)$$

Remark.

- Total maximizing probability obtained by $\max_{\theta_N} \mu_N(\theta_N) \triangleq \max_{\theta_{1:N}} P(\theta_{1:N}, z_{1:N})$
- $\theta^*_{1:N}$ obtained by simply backtracking the sequence which led to total maximizing probability
- The state estimate $\hat{\theta}_t$ for each time-step $t = 1,2,\ldots,N$ is obtained by defining:

$$\hat{\theta}_t \triangleq \arg\max_{\theta_t} \mu_t(\theta_t) \qquad (3)$$

Problem Formulation

Nominal Control Policies

- Markov feedback control policies: $g_k : \mathcal{X}_k \times \hat{\Theta}_k \longmapsto \mathcal{U}_k$ such that $u_k^g \triangleq g_k(x_k, \hat{\theta}_k)$
  - Complete state information about x_k
  - Available information about $\hat{\theta}_k \in \hat{\Theta}$
- Nominal transition p.d.: $p_{ij}^0 \triangleq P(\hat{\theta}_{k+1} = j \,|\, \hat{\theta}_1, \ldots, \hat{\theta}_k = i) = P(\hat{\theta}_{k+1} = j \,|\, \hat{\theta}_k = i), \quad i, j \in \hat{\Theta}$
  - As an approximation of the true transition p.d. of the estimator (in general, non-Markov)
  - Given $\hat{\theta}_{1:N}$, then p_{ij}^0 is evaluated by employing ML estimation

Ambiguity Class

All possible, time-varying, transition p.d. modelled by

$$\mathbb{B}(i,k) \triangleq \left\{ p_{i\bullet}(k) \in \mathbb{P}_k(\Theta \,|\, i) : \sum_{j \in \Theta} |p_{ij}(k) - p_{ij}^0| \leq R_{TV}(i,k) \right\} \tag{4}$$

- $R_{TV}(i,k) \in [0,2]$: preselected by the decision maker

Problem Formulation

Stochastic Control Problem

Define the N-stage expected cost by:

$$J_N(g,p) \triangleq \mathbb{E}_{x,i}^{g,p}\Big[(x_N-\bar{x}_N)^T Q_N(\hat{\theta}_N)(x_N-\bar{x}_N)+\sum_{k=1}^{N-1}\Big((x_k-\bar{x}_k)^T Q_k(\hat{\theta}_k)(x_k-\bar{x}_k)+u_k^T R_k(\hat{\theta}_k)u_k\Big)\Big] \quad (5)$$

- $\mathbb{E}_{x,i}^{g,p}[\,\cdot\,]$: indicates the dependence of the expectation operation on feedback Markov control policy, induced by $p_{i\bullet}(k)\in\mathbb{B}(i,k)$, for fixed initial data $x_1 = x$ and $\hat{\theta}_1 = i$

MiniMax Stochastic Control Problem

Find an optimal control policy $g^* \in G$, and a maximizing transition p.d. $p_{i\bullet}^*(k)\in\mathbb{B}(i,k)$ within the total variation distance ambiguity set which solve

$$J^* = J_N(g^*,p^*) \triangleq \min_{g\in G}\ \max_{\substack{p(k)\in\mathbb{B}(k)\\ k=1,2,\dots,N}} J_N(g,p)$$

MinMax Stochastic Control

MinMax Stochastic Control

Dynamic Programming

- Apply dynamic programming to obtain:

$$V_N(x_N, \hat{\theta}_N) = (x_N - \bar{x}_N)^T Q_N(\hat{\theta}_N)(x_N - \bar{x}_N)$$

$$V_k(x,i) = \min_{u_k \in \mathcal{U}} \max_{p_{i\bullet}(k) \in \mathbb{B}(i,k)} \mathbb{E}^{g,p}_{x,i}\Big[(x - \bar{x}_k)^T Q_k(i)(x - \bar{x}_k) + u_k^T R_k(i) u_k + V_{k+1}(x_{k+1}, \hat{\theta}_{k+1})\Big]$$

$$= \min_{u_k \in \mathcal{U}} \Big\{ (x - \bar{x}_k)^T Q_k(i)(x - \bar{x}_k) + u_k^T R_k(i) u_k + \max_{p_{i\bullet}(k) \in \mathbb{B}(i,k)} \mathbb{E}^{g,p}_{x,i}\big[V_{k+1}(x_{k+1}, \hat{\theta}_{k+1})\big] \Big\}, \quad x_k^g = x,\ \hat{\theta}_k = i$$

- *Induction hypothesis:* Suppose that:

$$V_t(x,i) = x^T P_t(i) x + x^T f_t(i) + r_t(i) \quad \text{holds for} \quad t = k+1, k+2, \ldots, N$$

By characterising the solution of the inner and outer optimization of the minimax dynamic programming equation, show that:

$$V_k(x,i) = x^T P_k(i) x + x^T f_k(i) + r_k(i), \quad P_k(i) \geq 0,\ f_k(i) \in \mathbb{R}^n,\ r_k(i) \in \mathbb{R} \tag{6}$$

MinMax Stochastic Control

Dynamic Programming

- Define the sequence by $\ell_k(x_k, \hat{\theta}_k, \hat{\theta}_{k+1}, u_k) \triangleq \int_{\mathcal{X}_{k+1}} V_{k+1}(x_{k+1}, \hat{\theta}_{k+1}) P^g(x_{k+1} \,|\, x_k, \hat{\theta}_k)$. Then, the dynamic programming equation becomes

$$V_k(x, i) = \min_{u_k \in \mathcal{U}} \left((x - \bar{x}_k)^T Q_k(i)(x - \bar{x}_k) + u_k^T R_k(i) u_k + \max_{p_{i\bullet} \in \mathbb{B}(i,k)} \sum_{j \in \Theta} \ell_k(x, i, j, u_k) p_{ij}(k) \right) \qquad (7)$$

Inner Optimization

- Identified partition: $\mathcal{P}(k, \hat{\theta}_k) \triangleq \{\Theta^0(k, \hat{\theta}_k), \Theta_0(k, \hat{\theta}_k), \ldots, \Theta_r(k, \hat{\theta}_k)\}$
- Inner optimization:

$$\begin{aligned} \max_{p_{i\bullet} \in \mathbb{B}(i,k)} \mathbb{E}_p^{\mathcal{P}}\left[\ell_k(x_k = x, \hat{\theta}_k = i, \hat{\theta}_{k+1}, u_k)\right] &= \ell_k\left(x, i, \hat{\theta}_{k+1} \in \Theta^0(k, i), u_k\right) \sum_{j \in \Theta^0(k,i)} p_{ij}^*(k) \\ &+ \ell_k\left(x, i, \hat{\theta}_{k+1} \in \Theta_0(k, i), u_k\right) \sum_{j \in \Theta_0(k,i)} p_{ij}^*(k) + \sum_{s=1}^{r} \ell_k\left(x, i, \hat{\theta}_{k+1} \in \Theta_s(k, i), u_k\right) \sum_{j \in \Theta_s(k,i)} p_{ij}^*(k) \end{aligned} \qquad (8)$$

MinMax Stochastic Control

Theorem: The maximizing, time-varying, transition p.d. $p^*_{i\bullet}(k) \in \mathbb{B}(i,k)$, for each $\hat{\theta}_k = i \in \hat{\Theta}$ and $k = 1,2,\ldots,N$ is given by

$$\sum_{j\in\Theta^0(k,i)} p^*_{ij}(k) = \sum_{j\in\Theta^0(k,i)} p^0_{ij} + \frac{\alpha_i(k)}{2}$$

$$\sum_{j\in\Theta_0(k,i)} p^*_{ij}(k) = \Big(\sum_{j\in\Theta_0(k,i)} p^0_{ij} - \frac{\alpha_i(k)}{2}\Big)^+$$

$$\sum_{j\in\Theta_s(k,i)} p^*_{ij}(k) = \Big(\sum_{j\in\Theta_s(k,i)} p^0_{ij} - \Big(\frac{\alpha_i(k)}{2} - \sum_{z=1}^{s}\sum_{j\in\Theta_{z-1}(k,i)} p^0_{ij}\Big)^+\Big)^+ \text{ for } s = 1,2,\ldots,r$$

where $\alpha_i(k) = \min\Big(R_{TV}(i,k), 2\big(1 - \sum_{j\in\Theta^0(k,i)} p^0_{ij}\big)\Big)$, and $(x)^+ \triangleq \max\{0,x\}$.

Remark.

- The solution is described by a water-filling effect
- By appropriately adjusting $R_{TV}(i,k)$ we can control (i) the size of the ambiguity set, and (ii) the degree of conservatism of the optimization problem

MinMax Stochastic Control

Outer Optimization

- Dynamic programming equation becomes

$$V_k(x,i) = \min_{u_k \in \mathcal{U}} \left((x-\bar{x}_k)^T Q_k(i)(x-\bar{x}_k) + u_k^T R_k(i) u_k + \sum_{j \in \Theta} \ell_k(x,i,j,u_k) p_{ij}^*(k) \right) \qquad (9)$$

- Differentiating with respect to u_k and setting the derivative equal to zero, we obtain $u_k^* = -L_k(i)x - s_k(i)$ for $x_k = x$ and $\hat{\theta}_k = i$
  - Preserves its linearity similar to the classical case
  - The evaluation of the feedback gain matrices, and the Riccati equations is performed based on the maximizing transition p.d.

- Minimax stochastic control problem optimal cost given by $J^* = \mathbb{E}^{g^*,p^*}[V_1(x_1,\hat{\theta}_1)]$

Numerical Example

Numerical Example

Input Data

- Number of states $n = 2$, number of different operating modes $n_\theta = 2$
- Final time $N = 200$, initial conditions $x_1 = [2 \quad 0]^T$
- Reference trajectory signal: $\bar{x}_k = \left[10(1 - \exp(-0.05k)) \quad 0\right]^T$ for $k = 1,2,\ldots,N$
- Dynamics and input matrices:

$$A_k(\theta_k = 1) = \begin{bmatrix} 0.99 & 0.53 \\ -0.10 & 1.15 \end{bmatrix}, \quad B_k(\theta_k = 1) = [0.0013 \quad 0.0539]^T$$

$$A_k(\theta_k = 2) = \begin{bmatrix} 1 & -0.59 \\ -0.025 & 0 \end{bmatrix}, \quad B_k(\theta_k = 2) = [0.02 \quad 0.10]^T$$

- State and cost matrices:

$$Q_k(\theta_k = 1) = Q_k(\theta_k = 2) = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.05 \end{bmatrix}, \quad R_k(\theta_k = 1) = R_k(\theta_k = 2) = 0.01$$

Numerical Example

HMM

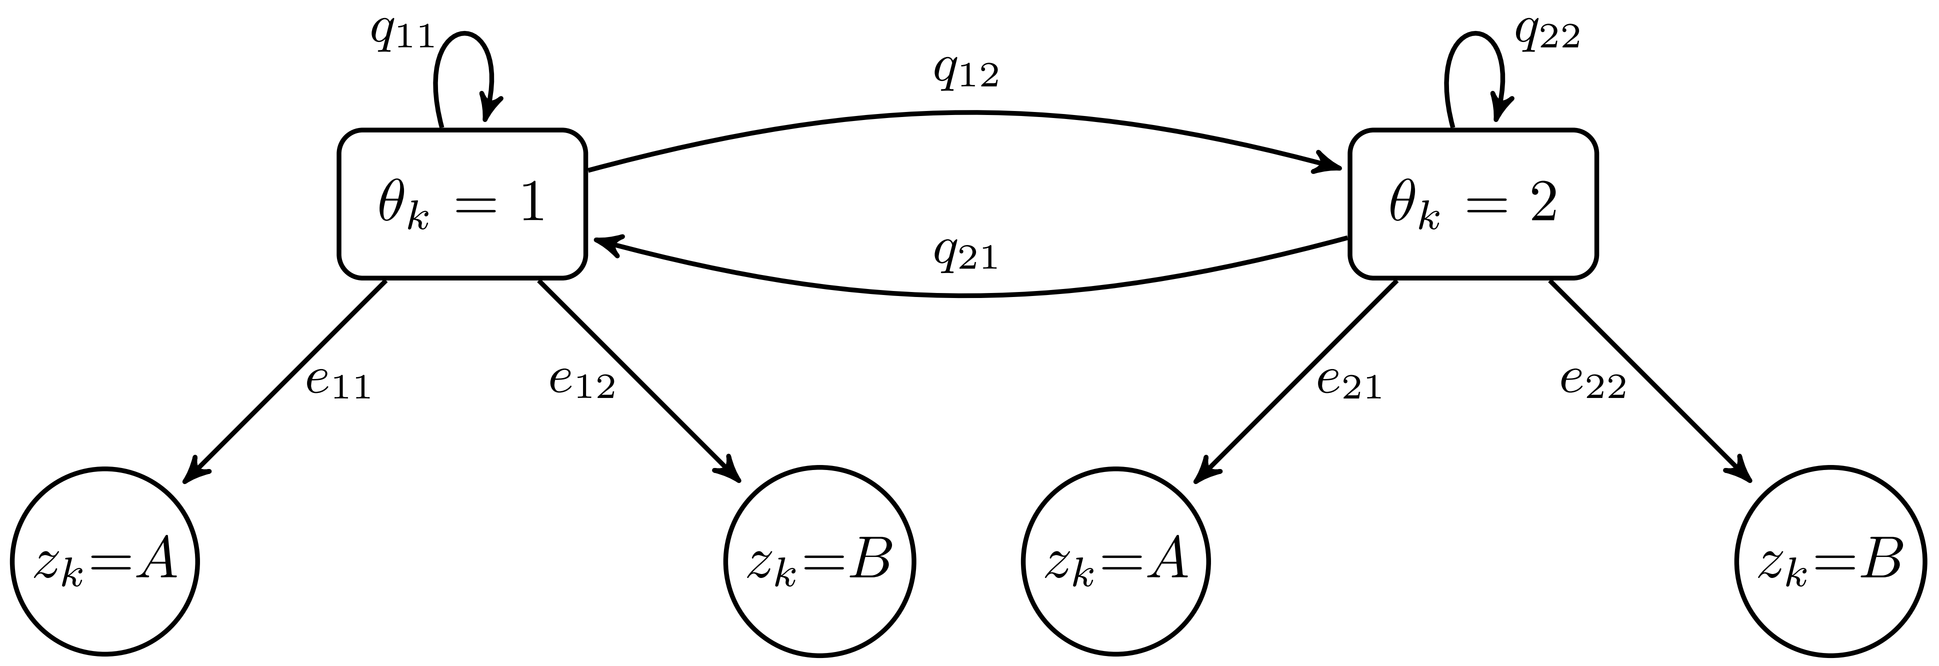

Hidden States: $q_{ij}(k) = \begin{bmatrix} 0.45 & 0.55 \\ 0.6 & 0.4 \end{bmatrix}$

Observed Variables: $e_{is}(k) = \begin{bmatrix} 0.1 & 0.9 \\ 0.8 & 0.2 \end{bmatrix}$

True HMM:

$$q_{ij}^{true}(k) = \begin{pmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{pmatrix}, \quad \text{for } k = 1,2,\ldots,N/2$$
$$q_{ij}^{true}(k) = \begin{pmatrix} 0.7 & 0.3 \\ 0.4 & 0.6 \end{pmatrix}, \quad \text{for } k = N/2+1,\ldots,N-1. \tag{10}$$

- **Total Variation Distance:**

$$R_{TV}(1,k) = R_{TV}(2,k) = \alpha, \quad k = 1,2,\ldots,N/2,$$
$$R_{TV}(1,k) = R_{TV}(2,k) = \alpha/2, \quad k = N/2+1,\ldots,N-1 \tag{11}$$

Numerical Example

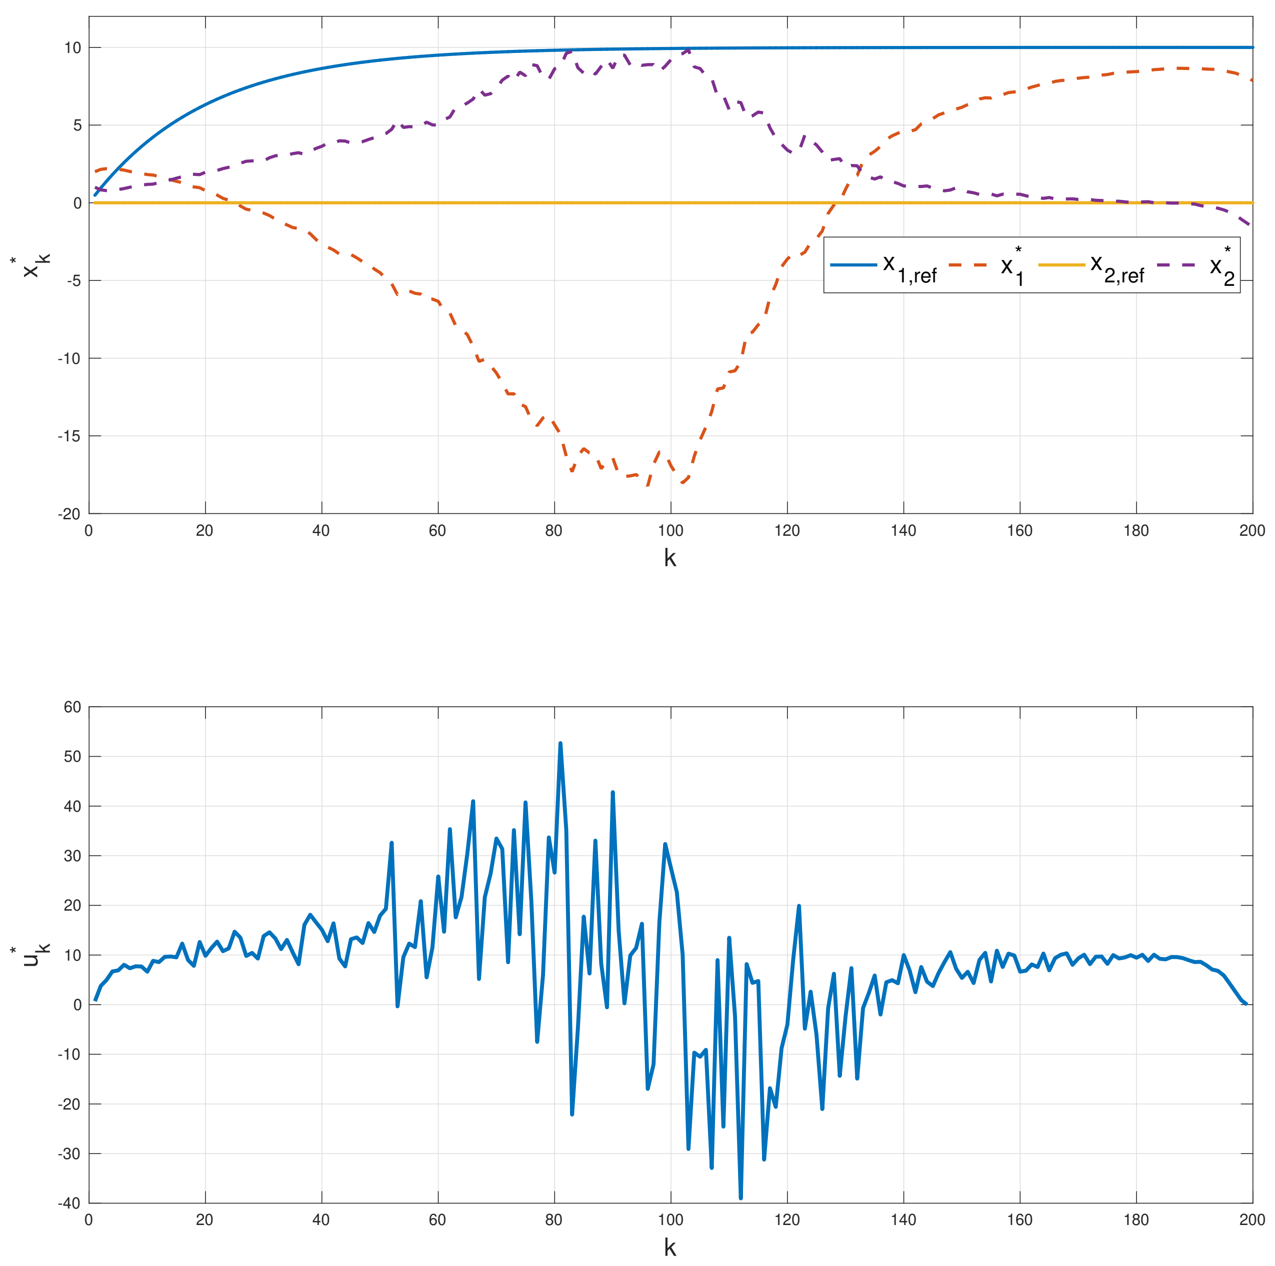

$\alpha = 0$

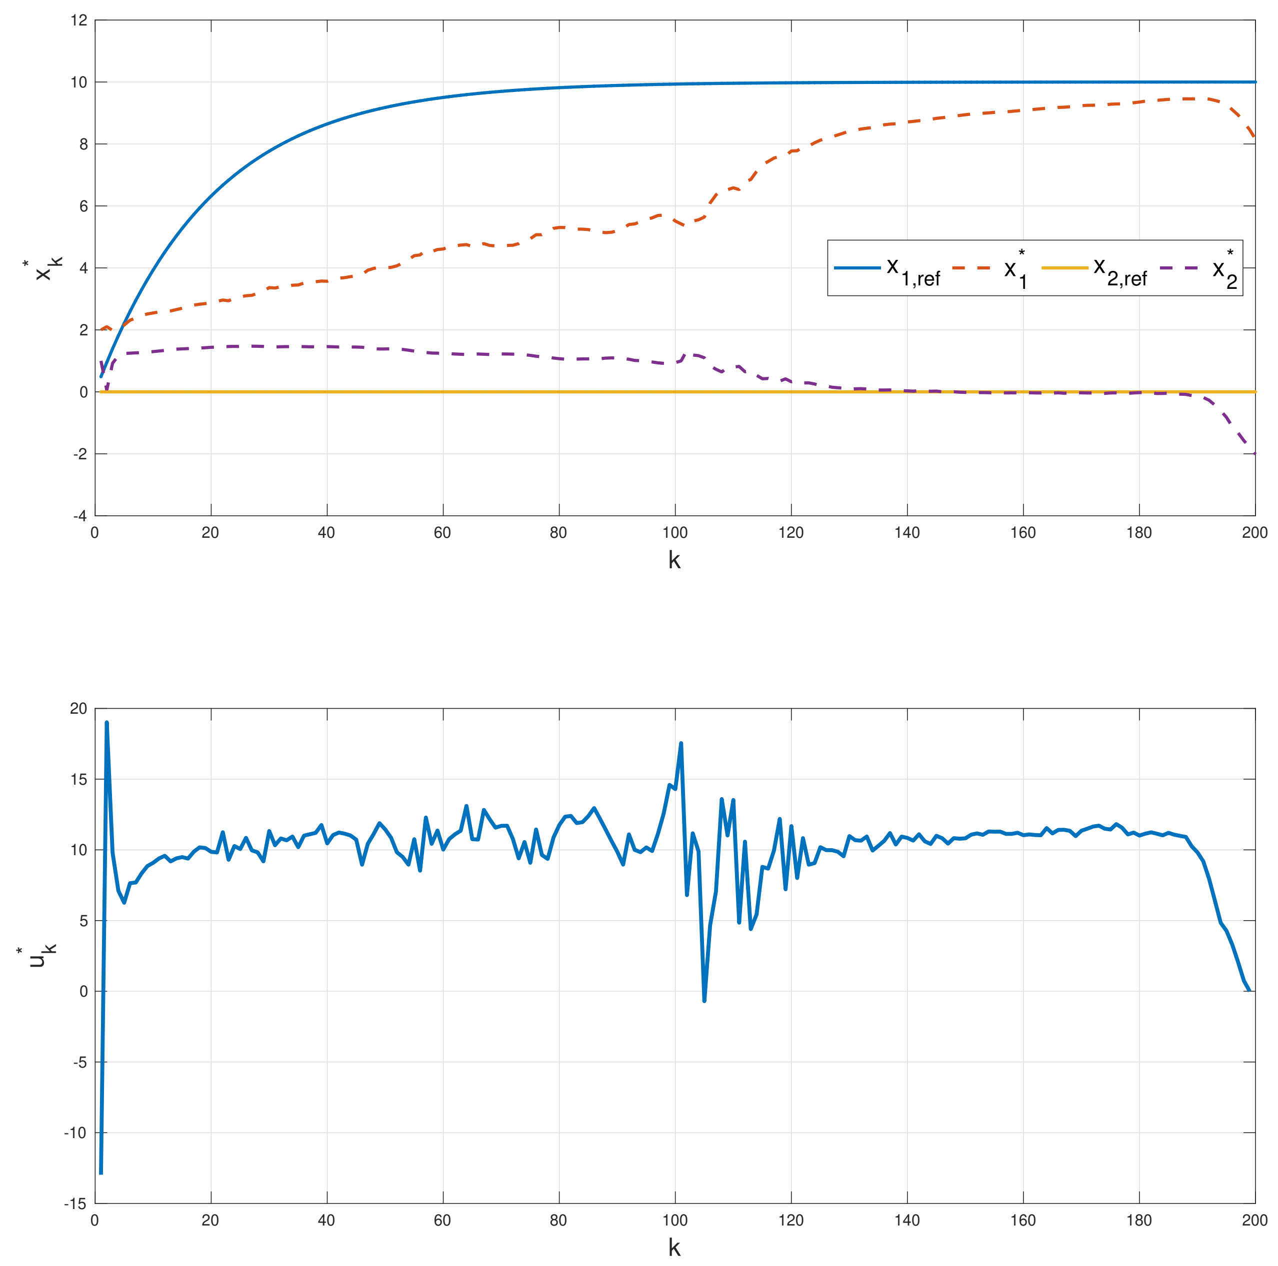

$\alpha = 0.3$

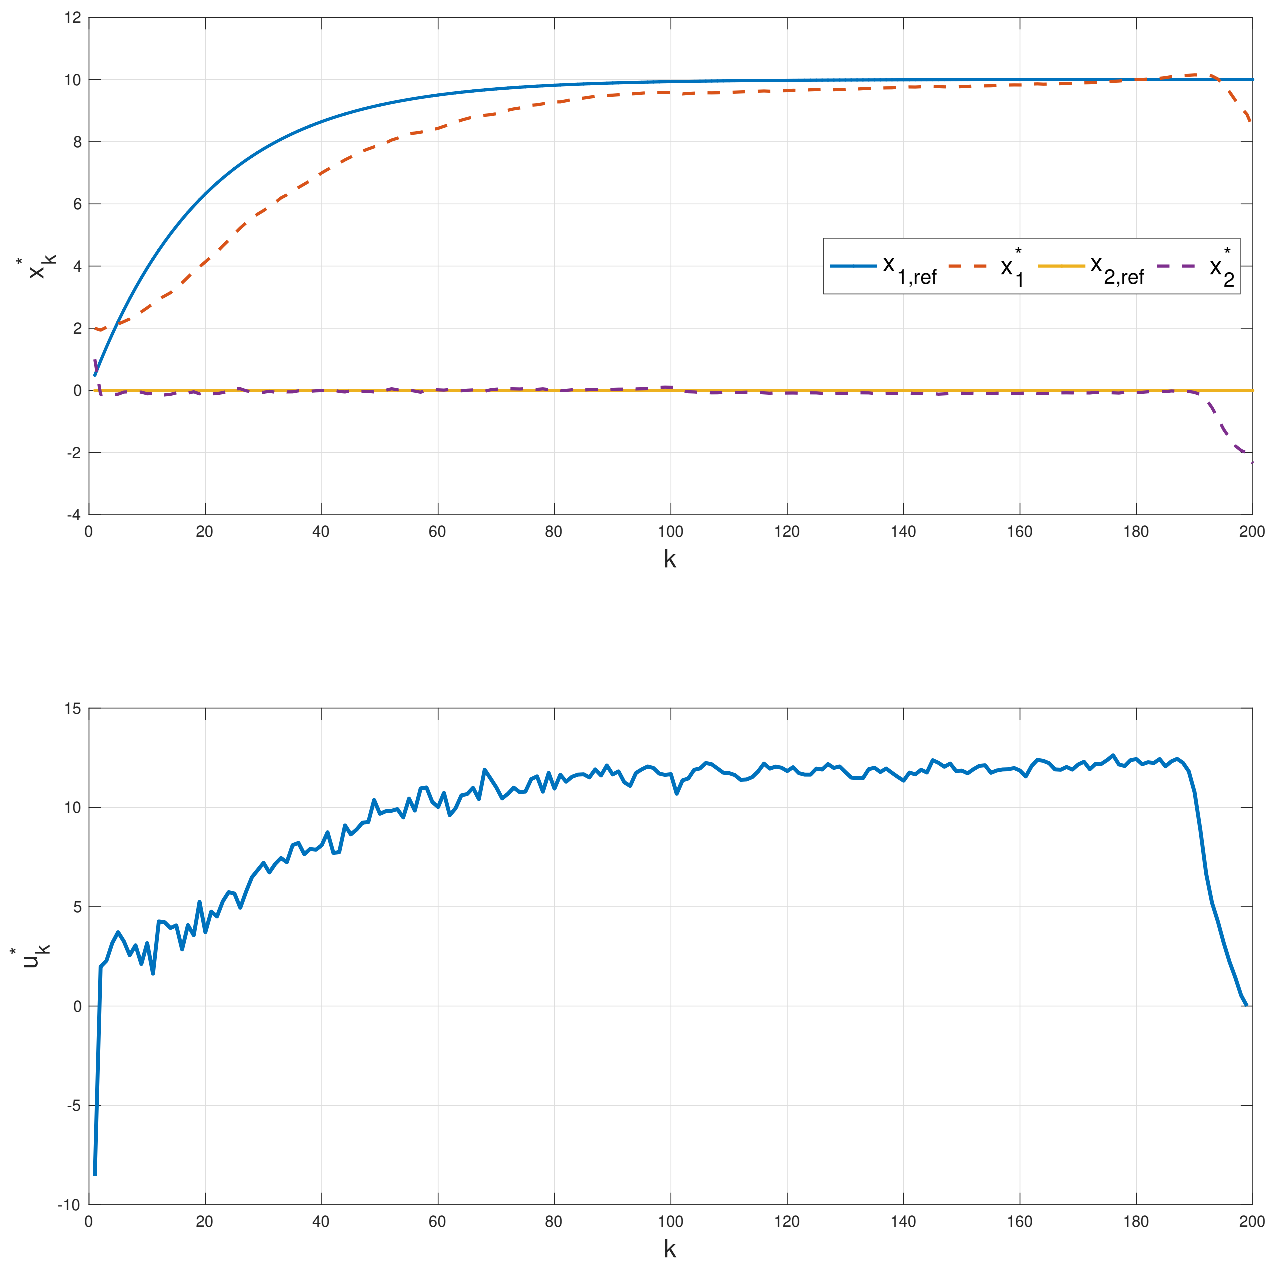

$\alpha = 0.9$

Two scenarios are considered:

(S1) $\alpha = 0$: no uncertainty, and hence, nominal HMM is correct

(S2) $\alpha > 0$: nominal HMM is uncertain

Concluding Remarks

Concluding Remarks

- LQ tracking control problem is analysed for:
  - MJLS with hidden states and uncertain conditional distribution
- State estimation and robust control problems are addressed, resulting in:
  - A maximizing, time-varying, conditional distribution
  - An optimal control policy with some desired robustness properties
- The proposed solution illustrated through an example:
  - Validates the capability on restricting the influence of uncertainty
  - Ensures the performance of the LQ tracking controller

Concluding Remarks

- LQ tracking control problem is analysed for:
  - MJLS with hidden states and uncertain conditional distribution
- State estimation and robust control problems are addressed, resulting in:
  - A maximizing, time-varying, conditional distribution
  - An optimal control policy with some desired robustness properties
- The proposed solution illustrated through an example:
  - Validates the capability on restricting the influence of uncertainty
  - Ensures the performance of the LQ tracking controller

Acknowledgement:
This work was partially funded by the European Regional Development Fund and the Republic of Cyprus through the Cyprus Research and Innovation Foundation
(Project: POST-DOC/0916/0139)

RESEARCH
& INNOVATION
FOUNDATION

Structural Funds
of the European Union in Cyprus