**QUEEN'S
UNIVERSITY
BELFAST**

Risk-averse model predictive control

Sopasakis, P., Herceg, D., Bemporad, A., & Patrinos, P. (2019). Risk-averse model predictive control. *Automatica*, 100, 281. <https://doi.org/10.1016/j.automatica.2018.11.022>

Published in:
Automatica

Document Version:
Peer reviewed version

Queen's University Belfast - Research Portal:
[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

© 2018 Elsevier Ltd.

This manuscript version is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>, which permits distribution and reproduction for non-commercial purposes, provided the author and source are cited

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Risk-averse model predictive control [★]

Pantelis Sopasakis ^a, Domagoj Herceg ^b, Alberto Bemporad ^b, Panagiotis Patrinos ^c

^aUniversity of Cyprus, Department of Electrical and Computer Engineering, KIOS Research Center for Intelligent Systems and Networks, 1 Panepistimiou Avenue, 2109 Aglantzia, Nicosia, Cyprus.

^bIMT School for Advanced Studies Lucca, Piazza San Francesco 19, 55100 Lucca, Italy.

^cKU Leuven, Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, Kasteelpark Arenberg 10, 3001 Leuven, Belgium.

Abstract

Risk-averse model predictive control (MPC) offers a control framework that allows one to account for ambiguity in the knowledge of the underlying probability distribution and unifies stochastic and worst-case MPC. In this paper we study risk-averse MPC problems for constrained nonlinear Markovian switching systems using generic cost functions, and derive Lyapunov-type risk-averse stability conditions by leveraging the properties of risk-averse dynamic programming operators. We propose a controller design procedure to design risk-averse stabilizing terminal conditions for constrained nonlinear Markovian switching systems. Lastly, we cast the resulting risk-averse optimal control problem in a favorable form which can be solved efficiently and thus deems risk-averse MPC suitable for applications.

Key words: Risk measures; Nonlinear Markovian switching systems; Model predictive control.

1 Introduction

1.1 Background and Contributions

There exist two main ways to deal with uncertainty in model predictive control (MPC), namely, the *robust* and the *stochastic* approaches. In *robust MPC*, modeling errors or disturbances are modeled as unknown-but-bounded quantities and the performance index is minimized with respect to the worst-case realization (min-max approach) [18]. However, such worst-case events are unlikely to occur in practice and render robust MPC severely conservative since all statistical information, typically available from past measurements, is ignored.

On the other hand, in *stochastic MPC* we assume that the underlying uncertainty is a random vector following some probability distribution [14] and minimize the expectation of a performance index; such formulations

are significantly less conservative. The driving random process is often taken to be normally and independently identically distributed [12] or it is assumed that it is a finite Markov process [17] and in *scenario-based MPC*, filtered probability distributions are estimated from data [11]. However, not always can we accurately estimate a distribution from available data, nor does it remain constant in time. Stochastic MPC will guarantee mean-square stability of the closed-loop system only with respect to the nominal probability distribution, therefore, errors in the estimation of that distribution may lead to bad performance or even instability.

The theory of *risk measures* [26] allows to interpolate between these two extreme cases. Roughly speaking, risk measures quantify the importance and effect of the right tail of a distribution of losses, that is, the impact of the occurrence of *extreme events*. As such they offer a mathematically elegant tool to tackle problems where we seek to avoid *high effect low probability* (HELP) events and can be readily used in various applications.

The first steps to risk-averse formulations can be traced back to linear-exponential-quadratic Gaussian control [13] and the study of stochastic control problems under inexact knowledge of the underlying probability distribution which is often termed *distributionally ro-*

[★] This is a preprint; please, cite the published version of the paper as follows: P. Sopasakis, D. Herceg, A. Bemporad and P. Patrinos, Risk-averse model predictive control, *Automatica*, 100:281–288, 2019. DOI: 10.1016/j.automatica.2018.11.022.

Email address: sopasakis.pantelis@ucy.ac.cy (Pantelis Sopasakis).

bust [10]. Distributionally robust control methodologies have been proposed for linear systems with probabilistic constraints assuming knowledge of some moments of the distribution [27]. The same problem was also recently addressed for Markov decision processes with uncertain transition probabilities [28].

Risk-averse MPC formulations for Markov jump linear systems (MJLS) are studied in [6, 7]. In [6] the authors formulate an MPC optimization problem employing a coherent risk measure of an uncertain cost as an objective function and give conditions under which the MPC control law is stabilizing, albeit for a system with no state-input constraints. This is extended in [7] assuming ellipsoidal state-input constraints. Building up on these results, we further improve on the state of the art by studying nonlinear systems and proposing a computationally favorable formulation for risk-averse optimization problems which leads to low computation times.

In the optimization and operations research communities, the solution of multistage risk-averse optimal control problems has been considered prohibitive as only bundle and cutting-plane methods are currently used [2, 5, 8]. Reported results are limited to short prediction horizons and linear stage cost functions. An alternative solution approach solves the dynamic programming (DP) problem using multiparametric piecewise quadratic programming [16], but its applicability is limited to systems with few states and small prediction horizons [15]. In a 2017 paper, Rockafellar proposed an algorithmic scheme for solving multistage problems using a non-composite (not nested) risk measure recognizing the difficulty of solving problems with nested risk mappings [19]. Indeed, the difficulty lies in that the cost function is written as a series of compositions of typically nonsmooth operators. In Section 5 we present a computationally tractable approach for the solution of multistage risk-averse problems by disentangling this series of compositions. This formulation renders risk-averse MPC suitable for embedded applications.

In this paper we formulate multistage risk-averse optimal control problems using Markov risk measures in a DP setting and derive Lyapunov-type risk-averse stability conditions. We study risk-averse MPC formulations for nonlinear Markovian switching systems under generally nonconvex joint state-input constraints and propose a controller design procedure for nonlinear systems with smooth dynamics and Lipschitz-continuous gradient. Lastly, we provide simulation examples to demonstrate the applicability of the proposed approach.

1.2 Notation

Let $\overline{\mathbb{R}} := \mathbb{R} \cup \{+\infty\}$ be the set of extended-real numbers, $\mathbb{N}_{[k_1, k_2]}$ the integers in $[k_1, k_2]$, for $z \in \mathbb{R}^n$ let $[z]_+ = \max\{0, z\}$ (where the max is taken element-wise).

We denote by 1_n the vector in $\mathbb{R}^n$ with all coordinates equal to 1. We denote the sets of n -by- n symmetric positive definite (semidefinite) matrices as $\mathcal{S}_{++}^n$ ($\mathcal{S}^n$). For two n -by- n symmetric matrices M_1, M_2 , $M_1 \succcurlyeq M_2$ means that $M_1 - M_2 \in \mathcal{S}_{++}^n$. We denote the transpose of a matrix A by $A^\top$ and the identity matrix by I . For a $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$, its *Jacobian matrix* is the mapping $Jg : \mathbb{R}^n \rightarrow \mathbb{R}^{m \times n}$ defined as $Jg(x) = (\partial g_i(x)/\partial x_j)_{i,j}$, provided that the partial derivatives exist. For $\epsilon \geq 0$ we define $\mathcal{B}_\epsilon = \{x \mid \|x\| \leq \epsilon\}$. For a set $C \subseteq \mathbb{R}^n$, we define its *indicator function* as $\delta_C(x) = 0$ if $x \in C$ and $\delta_C(x) = \infty$ otherwise. The *domain* of an extended-real-valued function $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is $\text{dom } f = \{x \in \mathbb{R}^n \mid f(x) < \infty\}$. An extended-real-valued function $f : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ is called *proper* if its domain is nonempty; it is called lower semi-continuous (lsc) if its lower level sets are closed. An $\ell : \mathbb{R}^n \times \mathbb{R}^m \ni (x, u) \mapsto \ell(x, u) \in \overline{\mathbb{R}}$ is called *level bounded in u locally uniformly in x* if for each $x_0 \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$, there is a neighborhood U_{x_0} of x_0 along with a bounded set $B \subseteq \mathbb{R}^m$ such that $\{u \mid \ell(x, u) \leq \alpha\} \subseteq B$ for all $x_0 \in U_{x_0}$. The *effective domain* of a set-valued mapping $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ is defined as $\text{dom } F = \{x \in \mathbb{R}^n \mid F(x) \neq \emptyset\}$. For a nonempty set E and a finite set $\mathcal{N}$ we define $\text{fcns}(E, \mathcal{N}) = \{V : E \times \mathcal{N} \rightarrow \overline{\mathbb{R}} \mid V(x, i) \geq 0, V(0, i) = 0, \text{ for all } x \in E, i \in \mathcal{N}\}$.

2 Risk-averse optimal control

2.1 Measuring risk

Let $\mathcal{N} = \mathbb{N}_{[1, n]}$ be a discrete sample space. A probability measure thereon can be identified by a probability vector $p \in \mathbb{R}^n$ with $\sum_{i=1}^n p_i = 1$, $p_i \geq 0$ for $i \in \mathcal{N}$. Let $Z : \mathcal{N} \rightarrow \mathbb{R}$ be a real-valued random variable on $\mathcal{N}$ which represents a random cost; for $i \in \mathcal{N}$ let $Z_i = Z(i)$. The vector $(Z_i)_{i \in \mathcal{N}}$ identifies the random variable Z .

The *expectation* of a random variable Z with respect to the probability vector p is defined as

$$\mathbb{E}_p[Z] \equiv \mathbb{E}_p[Z(i); i] = \sum_{i \in \mathcal{N}} p_i Z_i. \quad (1)$$

The notation $\mathbb{E}_p[Z; i]$ is to emphasize that the expectation is taken with respect to i .

A *risk measure* on $\mathbb{R}^n$ is a mapping $\rho : \mathbb{R}^n \rightarrow \mathbb{R}$. It is called *coherent* if it satisfies the following properties [26, Sec. 6.3] for $Z, Z' \in \mathbb{R}^n$, $\alpha \geq 0$, $\lambda \in [0, 1]$

- A1. *Convexity.* $\rho(\lambda Z + (1-\lambda)Z') \leq \lambda \rho(Z) + (1-\lambda)\rho(Z')$,
- A2. *Monotonicity.* $\rho(Z) \leq \rho(Z')$ whenever $Z \leq Z'$,
- A3. *Translation equivariance.* $\rho(c1_n + Z) = c + \rho(Z)$,
- A4. *Positive homogeneity.* $\rho(\alpha Z) = \alpha \rho(Z)$.

Trivially, the *expectation* is a coherent risk measure and so is the *essential maximum* $\text{essmax}[Z] := \max\{Z_i \mid$

$p_i > 0, i \in \mathcal{N}$. A popular risk measure is the *average value-at-risk*, also known as *conditional value-at-risk* or *expected shortfall*, which is defined as

$$\text{AV@R}_\alpha[Z] = \begin{cases} \min_{t \in \mathbb{R}} \{t + \alpha^{-1} \mathbb{E}_p[Z - t]_+\}, & \alpha \in (0, 1] \\ \text{essmax}(Z), & \alpha = 0. \end{cases}$$

As a result of assumptions A1–A4, coherent risk measures can be written in the following dual form [26, Thm. 6.5]

$$\rho[Z] = \max_{\mu \in \mathcal{A}(p)} \mathbb{E}_\mu[Z], \quad (2)$$

where $\mathcal{A}(p) \subseteq \mathbb{R}^n$ is a compact convex set of probability vectors containing p which we shall call the *ambiguity set* of ρ . We may think of a coherent risk measure as the worst-case expectation with respect to a probability distribution taken from a set of probability vectors. We call ρ a *polytopic* risk measure if $\mathcal{A}(p)$ is a polytope, i.e., it can be described by $\rho(Z) = \max\{\mu^\top Z \mid 1_n^\top \mu = 1, F(p)\mu \leq b(p)\}$ for some $F(p) \in \mathbb{R}^{q \times n}$ and $b(p) \in \mathbb{R}^q$. The expectation, the essential maximum and AV@R_α are polytopic risk measures. The ambiguity set of AV@R_α for $\alpha \in [0, 1]$ is the polytope $\mathcal{A}_\alpha(p) = \{\mu \in \mathbb{R}^n \mid \sum_{i=1}^n \mu_i = 1, \mu_i \geq 0, \alpha \mu_i \leq p_i\}$. The ambiguity set $\mathcal{A}_0(p)$ is the whole probability simplex. Apparently AV@R_α is a polytopic risk measure. AV@R_α interpolates between the risk-neutral expectation operator ($\text{AV@R}_1 = \mathbb{E}_p$, with $\mathcal{A}_0(p) = \{p\}$) and the worst-case essential maximum ($\text{AV@R}_0 = \text{essmax}$).

2.2 Markovian switching systems

In this work we consider Markovian switching systems

$$x_{k+1} = f(x_k, u_k, i_k), \quad (3)$$

driven by the random parameter i_k which is a time-homogeneous Markov chain with values in a finite set $\mathcal{N} = \mathbb{N}_{[1,n]}$ with transition matrix $P = (p_{ij}) \in \mathbb{R}^{n \times n}$, that is $\mathbb{P}[i_{k+1} = j \mid i_k = i] = p_{ij}$ [9]. We call the states of this Markov chain, the *modes* of (3). We denote the *cover* of each mode by $\text{cov}(i) := \{j \in \mathcal{N} \mid p_{ij} > 0\}$. We assume that at time k we measure the full state x_k and the value of i_k . As the probabilistic information available up to time k is fully described by the pair (x_k, i_k) , the control actions u_k may be decided by a causal control law $u_k = \kappa_k(x_k, i_k)$. This formulation aligns with that of the classic textbook [9], but there exist formulations where i_k is not known at time k and the control law is a function of x_k only [7].

Each $f(\cdot, \cdot, i) : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$, $i \in \mathcal{N}$, is assumed to be continuous and satisfy $f(0, 0, i) = 0$. MJLS are a special case of (3) with $f(x, u, i) = A_i x + B_i u$, $i \in \mathcal{N}$. System (3) is subject to the joint state-input constraints

$$(x_k, u_k) \in Y_{i_k}, \quad (4)$$

and we shall assume that for all $i \in \mathcal{N}$, Y_i are nonempty, closed sets containing the origin.

2.3 Markov risk measures

Consider the space of pairs (i, j) in $\Omega := \mathcal{N} \times \mathcal{N}$ equipped with the σ -algebra $\mathcal{F} = 2^\Omega$ and the probability measure $\mathbb{P}[\{(i, j)\}] = p_{ij}$. The conditional probability conditioned by the knowledge of i can be identified with the probability vector P_i — the i -th row of P . For a random variable $Z : \Omega \rightarrow \mathbb{R}$, the conditional expectation of Z conditioned by i , denoted as $\mathbb{E}_i[Z; j]$, is a random variable on $\mathcal{N}$, that is $\mathcal{N} \ni i \mapsto \mathbb{E}_i[Z; j] \in \mathbb{R}$, with

$$\mathbb{E}_i[Z; j] := \mathbb{E}_{P_i}[Z; j] = \sum_{j \in \mathcal{N}} p_{ij} Z(i, j). \quad (5)$$

We may extend this definition to define conditional variants of risk measures. Following (5), we give the following definition

Definition 1 (Markov risk measure) *Given a coherent risk measure ρ with ambiguity set $\mathcal{A}$ and a probability transition matrix P of a Markov chain, we define the Markov risk measure $\rho_i[Z; j]$ as*

$$\rho_i[Z; j] = \max_{\mu \in \mathcal{A}(P_i)} \underbrace{\sum_{j \in \mathcal{N}} \mu_j Z(i, j)}_{\mathbb{E}_\mu[Z; j]}, \quad (6)$$

for all random variables $Z : \Omega \rightarrow \mathbb{R}$.

This definition falls into the general framework of [21]. This way, with every i we associate the coherent risk measure $\rho_i[Z; j]$. As with the expectation, the notation $\rho_i[Z; j]$ is to emphasize that the risk is computed with respect to j .

2.4 Risk-averse optimal control and dynamic programming

Consider a *stage cost* function $\ell \in \text{fns}(\mathbb{R}^{n_x} \times \mathbb{R}^{n_u}, \mathcal{N})$ and a *terminal cost* $\ell_N \in \text{fns}(\mathbb{R}^{n_x}, \mathcal{N})$. Functions ℓ are extended-real-valued, therefore, they can encode constraints such as (4) by taking $\text{dom } \ell(\cdot, \cdot, i) = Y_i$, $i \in \mathcal{N}$. Likewise, ℓ_N can encode constraints on the terminal state of the form $x_N \in X_N^f$ by taking $\text{dom } \ell_N(\cdot, i) = X_i^f$, $i \in \mathcal{N}$, where X_i^f contain the origin in their interiors. We may now introduce the following finite-horizon risk-averse optimal control problem

$$\begin{aligned} \underset{u_0}{\text{minimize}} \quad & \ell(x_0, u_0, i_0) + \rho_{i_0} \left[\inf_{u_1} \ell(x_1, u_1, i_1) \right. \\ & + \rho_{i_1} \left[\inf_{u_2} \ell(x_2, u_2, i_2) + \dots \right. \\ & \left. \left. + \rho_{i_{N-1}} [\ell_N(x_N, i_N); i_N] \dots ; i_2 \right] ; i_1 \right], \end{aligned} \quad (7)$$

where $x_{k+1} = f(x_k, u_k, i_k)$, for all $k \in \mathbb{N}_{[0, N-1]}$. As it will become evident in what follows, each one of the infima at stage k in (7) is parametric in x_k and i_k , that is, the minimization takes place over causal control laws $u_0, \dots, u_{N-1}$. Note that under assumptions A1 and A2, we may interchange the Markov risk measures with the infima [26, Prop. 6.60] leading to risk-averse multistage formulations discussed in [26, Sec. 6.8.4].

Problem (7) can be described by a DP recursion. Inspired by [26, Sec. 6.8], for a $V \in \mathbf{fns}(\mathbb{R}^{n_x}, \mathcal{N})$ we define the DP operator $\mathbf{T} : \mathbf{fns}(\mathbb{R}^{n_x}, \mathcal{N}) \rightarrow \mathbf{fns}(\mathbb{R}^{n_x}, \mathcal{N})$ so that

$$\begin{aligned} (\mathbf{TV})(x, i) &= \inf_u \{ \ell(x, u, i) + \rho_i [V(f(x, u, i), j); j] \} \\ &= \inf_u \ell(x, u, i) + \max_{\mu \in \mathcal{A}(P_i)} \sum_{j \in \mathcal{N}} \mu_j V(f(x, u, i), j). \end{aligned}$$

Let $(\mathbf{SV})(x, i)$ be the corresponding set of minimizers for the optimization problem involved in the definition of $(\mathbf{TV})(x, i)$. This defines the following DP recursion

$$V_{k+1}^* = \mathbf{TV}_k^*, \quad (8a)$$

$$\mathcal{U}_{k+1}^* = \mathbf{SV}_k^*, \quad (8b)$$

for $k \in \mathbb{N}_{[0, N-1]}$ with $V_0^*(x, i) := \ell_N(x, i)$, $i \in \mathcal{N}$. For $C = \{C_i\}_{i \in \mathcal{N}}$ with $C_i \subseteq \mathbb{R}^{n_x}$, we define the mode-dependent predecessor operator $R(C) = \{R_i(C)\}_{i \in \mathcal{N}}$ with $R_i(C) = \{x \in \mathbb{R}^{n_x} \mid \exists u \in \mathbb{R}^{n_u}, (x, u) \in Y_i, f(x, u, i) \in \bigcap_{j \in \mathbf{cov}(i)} C_j\}$. Next, we present some fundamental properties of the DP operator $\mathbf{T}$.

Proposition 2 *If $\ell_N(\cdot, i)$ are proper, lsc and $\ell(\cdot, \cdot, i)$ are proper, lsc and level bounded in u locally uniformly in x for all $i \in \mathcal{N}$, then for all $i \in \mathcal{N}$:*
(i) $\mathbf{TV} \in \mathbf{fns}(\mathbb{R}^{n_x}, \mathcal{N})$ for $V \in \mathbf{fns}(\mathbb{R}^{n_x}, \mathcal{N})$,
(ii) $V_k^(\cdot, i)$ are lsc, (iii) $\mathbf{dom} V_k^*(\cdot, i) = \mathbf{dom} \mathcal{U}_k^*(\cdot, i) \neq \emptyset$,*
(iv) $\mathcal{U}_k^$ is compact-valued, (v) $\mathbf{dom}(V_{k+1}^*) = R(\mathbf{dom}(V_k^*))$.*

PROOF. *The proof goes along the lines of [17, Thm. 11a] using [20, Prop. 1.17, Prop. 1.26(a)].*

We may easily verify the monotonicity property $\mathbf{TV} \leq \mathbf{TV}'$, for all V, V' with $V \leq V'$, following [4]. An observation that will prove useful in what follows is that if $\mathbf{T}\ell_N \leq \ell_N$, then $V_{k+1}^* \leq V_k^*$. The above risk-averse optimal control problem leads naturally to the statement of a risk-averse MPC problem where control actions are computed by a control law $\kappa_N^*(x, i) \in \mathcal{U}_N^*(x, i)$. In Section 3 we state an appropriate risk-based notion of stability and provide conditions on ℓ_N for the MPC-controlled system $x_{k+1} = f(x_k, \kappa_N^*(x_k, i_k), i_k)$ to be stable.

3 Risk-averse stability

Consider the following Markovian switching system which is controlled by some control law $u_k = \kappa(x_k, i_k)$

$$x_{k+1} = f^\kappa(x_k, i_k) := f(x_k, \kappa(x_k, i_k), i_k), \quad (9)$$

subject to the constraints $(x_k, i_k) \in X^\kappa := \{(x, i) \mid (x, \kappa(x, i)) \in Y_i\}$. For convenience, we introduce the notation $X_i^\kappa = \{x \mid (x, \kappa(x, i)) \in Y_i\}$, for $i \in \mathcal{N}$. Let $i_{[k]} = (i_0, i_1, \dots, i_k)$ denote an admissible path of length k of the Markov chain $\{i_t\}_{t \in \mathbb{N}}$, that is, $i_{t+1} \in \mathbf{cov}(i_t)$ for $t \in \mathbb{N}_{[0, k-1]}$. For a given initial state x_0 , the solution of (9) at time k is denoted as $\phi(k, x_0, i_{[k-1]})$.

In order to be able to define risk-based notions of stability, we must first introduce an appropriate notion of invariance for Markovian switching systems [17].

Definition 3 (Uniform invariance) *Let $X = \{X_i\}_{i \in \mathcal{N}}$ be a collection of nonempty closed subsets of $\mathbb{R}^{n_x}$ and $X_i \subseteq X_i^\kappa$. X is called uniformly invariant (UI) for (9) subject to constraints $x \in X_i^\kappa$ if $f^\kappa(x, i) \in \bigcap_{j \in \mathbf{cov}(i)} X_j$, whenever $x \in X_i$ for all $i \in \mathcal{N}$.*

For the controlled system (9), the predecessor operator is now defined as $R_i(C) = \{x \in X^\kappa \mid f^\kappa(x, i) \in \bigcap_{j \in \mathbf{cov}(i)} C_j\}$. We have that C is UI if and only if $C_i \subseteq R_i(C)$ for all $i \in \mathcal{N}$ [17].

Given a coherent risk measure ρ and a random variable $\psi(i_0, i_1, \dots, i_k)$, let $\bar{\rho}_1[\psi] = \rho_{i_0}[\psi(i_0, i_1, \dots, i_k); i_1]$ and recursively define $\bar{\rho}_k = \bar{\rho}_{k-1} \circ \rho_{i_{k-1}}[\cdot; i_k]$, that is $\bar{\rho}_k[\psi] = \rho_{i_0}[\rho_{i_1}[\dots \rho_{i_{k-1}}[\psi(i_0, i_1, \dots, i_k); i_k] \dots; i_2]; i_1]$ [26, Sec. 6.8.2].

We may now give the following stability notion [6].

Definition 4 (Risk-square exponential stability)

We say that the origin is risk-square exponentially stable (RSES) for system (9) over a set $X = \{X_i\}_{i \in \mathcal{N}}$ if X is UI and for $x_0 \in X_{i_0}$

$$\bar{\rho}_{k-1}[\|\phi(k, x_0, i_{[k-1]})\|^2] \leq \lambda \beta^k \|x_0\|^2,$$

for all $k \in \mathbb{N}$, for some $\beta \in [0, 1)$, $\lambda \geq 0$.

RSES entails that the origin is exponentially mean-square stable for system (9) not only for the nominal probability distribution, but also for those probability distributions in the ambiguity set of the risk measure. In the unconstrained case, RSES corresponds to the notion of uniform global risk-sensitive exponential stability which is defined using the notion of dynamic risk measures [6]. If the underlying risk measure is the expectation operator, then RSES reduces to mean-square

exponential stability, whereas, if it is the essential supremum operator, it yields the definition of robust exponential stability. Additionally, since all coherent risk measures are lower bounded by the expectation, RSES is a stronger notion of stability compared to mean-square stability. The following lemma provides Lyapunov-type stability conditions for RSES.

Lemma 5 (RSES conditions) *Suppose there is a $V \in \mathbf{fcns}(\mathbb{R}^{n_x}, \mathcal{N})$, proper, lsc function such that*

- (i) $\mathbf{dom} V$ is a UI set
- (ii) $\rho_i [V(f^\kappa(x, i), j); j] - V(x, i) \leq -c\|x\|^2$, for some $c > 0$ for all $(x, i) \in \mathbf{dom} V$.

Then, $\bar{\rho}_k [\sum_{t=0}^{k-1} \|\phi(t, x_0, i_{[t-1]})\|^2]$, is uniformly bounded in k for $(x_0, i_0) \in \mathbf{dom} V$. If, additionally,

- (iii) *for all $(x, i) \in \mathbf{dom} V$, $\alpha_1 \|x\|^2 \leq V(x, i) \leq \alpha_2 \|x\|^2$, for some $\alpha_1, \alpha_2 > 0$,*

then, the origin is RSES for system (9) over $\mathbf{dom} V$.

PROOF. *The proof can be found in the appendix.*

The uniform boundedness condition in Lemma 5 is reminiscent of the notion of stochastic stability in [9, Sec. 3.3.1]. In fact, if the risk measure in Lemma 5 is the expectation operator, then the uniform boundedness condition is equivalent to mean-square stability [9, Thm. 3.9(6)].

We call a function $V \in \mathbf{fcns}(\mathbb{R}^{n_x}, \mathcal{N})$ which satisfies all requirements of Lemma 5, a (mode-dependent) *risk-averse Lyapunov function*. We may now state conditions on the stage cost ℓ and the terminal cost ℓ_N which entail RSES for the risk-averse MPC-controlled system.

4 Risk-averse MPC

4.1 Risk-averse MPC stability

Theorem 6 (RSES of MPC) *Suppose that (i) $c\|x\|^2 \leq \ell(x, u, i)$ for some $c > 0$ for all $(x, u) \in Y_i$, $i \in \mathcal{N}$ (ii) $\ell_N(x, i) \leq d\|x\|^2$, for some $d > 0$ for all $x \in X_i^f$, (iii) X_i^f contain the origin in their interiors (iv) V_N^* is locally bounded over its domain, that is, for every compact set $\bar{X} \subseteq \mathbf{dom} V_N^*$, there is an $M \geq 0$ so that $V_N^*(x, i) \leq M$ for all $(x, i) \in \bar{X}$ and*

$$\mathbf{T}\ell_N \leq \ell_N. \quad (10)$$

Then, the origin is RSES for the risk-averse MPC-controlled system $x_{k+1} = f(x_k, \kappa_N^(x_k, i_k), i_k)$ over all compact uniformly invariant subsets of $\mathbf{dom} V_N^*$.*

PROOF. *The proof can be found in the appendix.*

In Thm. 6 we show that V_N^* is a mode-dependent risk-averse Lyapunov function over compact uniformly invariant subsets of $\mathbf{dom} V_N^*$. We shall use this result in the following sections to design risk-averse stabilizing MPC controllers for MJLS as well as nonlinear Markovian switching systems. Note that Condition (iv) in Thm. 6 holds if the following assumption is satisfied (see [18, Prop. 2.15])

Assumption 7 (Local boundedness of V_N^*) *For all $i \in \mathcal{N}$, functions $\ell(\cdot, \cdot, i)$ and $\ell_N(\cdot, i)$ are continuous on their domains, and the sets $U_i(x) := \{u \in \mathbb{R}^{n_u} \mid (x, u) \in Y_i\}$ are compact and bounded uniformly in x .*

Additionally, because of the monotonicity property of $\mathbf{T}$ and since $\mathbf{T}\ell_N \leq \ell_N$, condition (10) implies $V_{k+1}^* \leq V_k^*$, thus $\mathbf{dom}(V_k^*) \subseteq \mathbf{dom}(V_{k+1}^*) = R(\mathbf{dom} V_k^*)$ (Prop. 2), thus $\mathbf{dom} V_k^*$ is UI.

4.2 Risk-averse MPC design for MJLS

Here we provide RSES conditions and design guidelines for risk-averse MPC of MJLS [9], that is $f(x, u, i) = A_i x + B_i u$, using a quadratic stage cost $\ell(x, u, i) = x^\top Q_i x + u^\top R_i u + \delta_{Y_i}(x, u)$, with $Q_i \in \mathcal{S}_+^{n_x}$, $R_i \in \mathcal{S}_+^{n_u}$ and Y_i are polytopes with the origin in their interiors. The terminal cost function is taken to be $\ell_N(x, i) = x^\top P_i^f x + \delta_{X_i^f}(x)$ with $P_i^f \in \mathcal{S}_+^{n_x}$ and X_i^f . We shall derive conditions on P_i^f and X_i^f so that the stabilizing conditions of Thm. 6 are satisfied. Condition $\mathbf{T}\ell_N \leq \ell_N$ is equivalent to

$$\min_u \{x^\top Q_i x + u^\top R_i u + \rho_i [x^\top P_j^f x; j]\} \leq x^\top P_i^f x, \quad (11a)$$

$$\mathbf{dom}(\mathbf{T}\ell_N) \supseteq \mathbf{dom} \ell_N \Leftrightarrow R(X^f) \supseteq X^f, \quad (11b)$$

where $x^+ = f(x, u, i)$ and the minimization in (11a) is over the space of admissible causal control laws $u = \kappa(x, i)$ so that $(x, i) \in X^\kappa$. An upper bound to the left hand side of (11a) is obtained by parametrizing $u = K_i x$. We introduce the shorthand notation $\bar{A}_i = A_i + B_i K_i$ and $\bar{Q}_i = Q_i + K_i^\top R_i K_i$, for $i \in \mathcal{N}$. Condition (11b) means that X^f is a UI set for the system $x_{k+1} = (\bar{A}_{i_k} + B_{i_k} K_{i_k}) x_k$ under the prescribed constraints. Such a set can be determined by the fixed-point iteration $\mathcal{O}_{k+1} = R(\mathcal{O}_k)$ with $\mathcal{O}_0 = \{(x, i) \mid (x, K_i x) \in Y_i\}$. If this iteration converges in a finite number of iterations — a sufficient condition for which is given in [17, Lem. 21] — to a set $\mathcal{O}_\infty$, this is a *polytopic* UI set.

Assuming that ρ is a polytopic Markov risk measure with ambiguity set $\mathcal{A}(P_i) = \mathbf{conv}\{\mu_i^{(l)}\}_{l \in \mathbf{N}_{[1, s_i]}}$ and using its dual representation, condition (11a) becomes $\bar{Q}_i +$

$\sum_{j \in \text{cov}(i)} \mu_{ij}^{(l)} (\bar{A}_i^\top P_j^f \bar{A}_i) \preceq P_i^f$ for all $i \in \mathcal{N}$ and $l \in \mathbb{N}_{[1, s_i]}$. This condition can be cast as a linear matrix inequality (LMI) by a change of variables $(P_i^f)^{-1} = M_i$, $K_i = Y_i M_i^{-1}$, $F_i^l = [\sqrt{\mu_{i1}^{(l)}} I \dots \sqrt{\mu_{in}^{(l)}} I]$ and $M = \text{blkdiag}(M_1, \dots, M_n)$:

$$\begin{bmatrix} M_i (A_i M_i + B_i Y_i)^\top F_i^l & M_i Q_i^{1/2} & Y_i^\top R_i^{1/2} \\ * & M & 0 \\ * & * & I \\ * & * & 0 \\ * & * & I \end{bmatrix} \succcurlyeq 0, \quad (12)$$

for all $i \in \mathcal{N}$ and $l \in \mathbb{N}_{[1, s_i]}$. The left hand side of (12) is a symmetric matrix, therefore, we show only its upper block triangular part and replaced the lower block triangular part by asterisks (*) to simplify the notation. Solving this LMI for $M_i \in \mathcal{S}_+^{n_x}$ and $Y_i \in \mathbb{R}^{n_u \times n_x}$ yields the linear gains K_i and the cost matrices P_i^f . LMI (12) has to be solved once offline to determine matrices P_i^f .

4.3 Risk-averse MPC design for nonlinear Markovian switching systems

For nonlinear systems, an obvious choice for the terminal cost function would be $\ell_N(x, i) = \delta_{\{0\}}(x)$ — meaning, $X_i^f = \{0\}$ for $i \in \mathcal{N}$ — but that would lead to a very conservative design. Here we exploit the system linearization at the origin to determine a terminal cost function and terminal constraints which render the MPC-controlled system RSES. We shall first draw the following assumption for the nonlinear dynamics:

Assumption 8 For each $i \in \mathcal{N}$, $f(\cdot, \cdot, i)$ is differentiable with L_i -Lipschitz Jacobian.

We use a parametric controller of the form $\kappa(x, i) = K_i x$ and define the associated closed-loop function $f^\kappa(x, i) = f(x, K_i x, i)$, $i \in \mathcal{N}$. Function $f^\kappa(\cdot, \cdot, i)$ can be written as a composition of $f(\cdot, \cdot, i)$ with the linear mapping $W_i : (x, u) \mapsto (x, K_i x)$, therefore, its Jacobian matrix will be Lipschitz-continuous with Lipschitz constant $L_i \|W_i\|^2$ which is bounded above by

$$\beta_i := L_i(1 + \|K_i\|^2). \quad (13)$$

The linearization of the nonlinear system at the origin is an MJLS $x_{k+1} = \hat{f}(x_k, u_k, i_k) := A_{i_k} x_k + B_{i_k} u_k$ with A_{i_k} and B_{i_k} given by the Jacobian matrices, with respect to x and u respectively, of f at the origin. That is, $A_i = J_x f(0, 0, i)$, $B_i = J_u f(0, 0, i)$. For notational convenience, we define the following quantities

$$\begin{aligned} \hat{f}^\kappa(x, i) &:= (A_i + B_i K_i)x, \\ \mathcal{L}\ell_N(x, i) &:= \rho_i [\ell_N(\hat{f}^\kappa(x, i), j); j] - \ell_N(x, i), \\ \mathcal{L}'\ell_N(x, i) &:= \rho_i \left[\ell_N(\hat{f}^\kappa(x, i), j); j \right] - \ell_N(x, i), \\ \Delta(x, i) &:= \mathcal{L}\ell_N(x, i) - \mathcal{L}'\ell_N(x, i). \end{aligned}$$

Fig. 1. (Left) A Markov chain with three modes and the corresponding transition probabilities, (Right) The corresponding tree with $i_0 = 1$.

The objective is to design the terminal cost and terminal constraints for the risk-averse MPC problem using $\mathcal{L}'\ell_N$ to yield an LMI. While our design will be based on the linearized dynamics, we need to account for the linearization error. To this end, we shall derive a quadratic upper bound for $|\Delta(x_k, i_k)|$ in a neighborhood of the origin.

Theorem 9 Suppose that Assumptions 7 and 8 hold and

$$\mathcal{L}'\ell_N(x, i) \leq -x^\top (\bar{Q}_i + m_i I)x, \quad (14)$$

for $i \in \mathcal{N}$, $m_i > 0$, ℓ , ℓ_N and X^f satisfy the requirements of Thm. 6 with $X_i^f \subseteq \mathcal{B}_{\delta_i}$ for some $\delta_i > 0$ for all $i \in \mathcal{N}$,

$$\sigma_i := \max_{j \in \text{cov}(i)} \|P_j^f\| \left(\frac{\beta_i^2 \delta_i^2}{4} + \beta_i \|\bar{A}_i\| \delta_i \right) < m_i, \quad (15)$$

and $\ell(x, K_i x, i) \leq x^\top (\bar{Q}_i + (m_i - \sigma_i)I)x$. If X^f is a UI set for (9), then the origin is RSES for the MPC-controlled system $x_{k+1} = f(x_k, \kappa_N^*(x_k, i_k), i_k)$ over the compact UI subsets of $\text{dom } V_N^*$.

PROOF. The proof can be found in the appendix.

According to Thm. 9, one first needs to select $m_i > 0$ for each $i \in \mathcal{N}$ such that (14) holds true. In the most common case where ℓ and ℓ_N are quadratic functions, this is precisely an LMI of the form (12) with $Q_i + m_i I$ in place of Q_i solving which we obtain matrices K_i and P_i^f and determine the constants β_i and find $\delta_i > 0$ so that (15) holds. The last step is to determine a UI set X^f for the nonlinear system $x_{k+1} = f^\kappa(x_k, i_k)$. We may cast the nonlinear system as a linear one with bounded additive disturbance $x_{k+1} = \bar{A}_{i_k} x_k + e(x_k, i_k)$ — indeed, as we show in the proof of Thm. 9, $\|e(x, i)\| \leq \beta_i/2 \|x\|^2$. We may follow the approach of [25] in order to determine a polytopic robustly invariant set.

5 Computationally tractable formulation of risk-averse optimal control problems

Starting from an initial state x_0 and initial mode i_0 and computing control actions according to a causal control law u_k , the future states of the Markovian system, up to some future time N , span a *scenario tree* — a tree-like structure such as the one shown in Fig. 1. Note that the state at a node ι , the input and mode leading to that node are denoted as x^ι , u^ι and i^ι respectively.

The possible realizations of the system state at time k define the *nodes* of the tree. The set of all nodes at stage k defines the set Ω_k . The set of nodes in Ω_{k+1} which are reachable from a node $\iota \in \Omega_k$ is called the set of *children* of ι and is denoted by $\mathbf{ch}(\iota)$ which is a subset of Ω_{k+1} . The space $\mathbf{ch}(\iota)$ becomes a probability space with $P[\{\eta\}] = p_{i^\iota, \eta}$ for $\eta \in \mathbf{ch}(\iota)$. As illustrated in Fig. 1, the system dynamics on the scenario tree is described by $x^\eta = f(x^\iota, u^\eta, i^\eta)$, for $\eta \in \mathbf{ch}(\iota)$ and $x^0 = x_0$, $i^1 = i_0$.

On the scenario tree, we define a process Φ as follows: for $\iota \in \Omega_N$ we define $\Phi^\iota := \rho_{i^\iota}[\ell_N(x^\iota, i^\iota); \eta] = \max_{\mu^\iota \in \mathcal{A}(P_{i^\iota})} \sum_{\eta \in \mathbf{ch}(\iota)} \mu_\eta^\iota \ell_N(x^\iota, i^\iota)$. Moreover, $\ell_N(x, i) = \inf_{\tau \in \mathcal{L}_N(x, i)} \tau$. When the underlying risk measure is polytopic with $\mathcal{A}(p) = \{\mu \in \mathbb{R}^n \mid \sum_{i=1}^n \mu_i = 1, F(p)\mu \leq b(p)\}$ with $b(p) \in \mathbb{R}^q$, then

$$\begin{aligned} \Phi^\iota &= \max_{\mu^\iota \in \mathcal{A}(P_{i^\iota})} \inf_{\substack{\ell_N(x^\iota, i^\iota) \leq \tau_\eta^\iota, \\ l \in \mathbf{ch}(\iota)}} \sum_{\eta \in \mathbf{ch}(\iota)} \mu_\eta^\iota \tau_\eta^\iota \\ &= \inf_{\substack{\ell_N(x^\iota, i^\iota) \leq \tau_\eta^\iota, \\ l \in \mathbf{ch}(\iota)}} \max_{\mu^\iota \in \mathcal{A}(P_{i^\iota})} \sum_{\eta \in \mathbf{ch}(\iota)} \mu_\eta^\iota \tau_\eta^\iota \\ &= \inf_{\substack{\tau^\iota, y^\iota \geq 0, \lambda^\iota \in \mathbb{R}, \\ \ell_N(x^\iota, i^\iota) \leq \tau_\eta^\iota, l \in \mathbf{ch}(\iota), \\ \tau^\iota = F(P_{i^\iota})^\top y^\iota + \lambda^\iota \mathbf{1}_q}} b(P_{i^\iota})^\top y^\iota + \lambda^\iota, \end{aligned}$$

where in the first equation we interchanged **max** with **inf** using [3, Prop. 2.6.4] using the fact that the level sets of the mapping $\tau^\iota \mapsto \max_{\mu^\iota \in \mathcal{A}(P_{i^\iota})} \sum_{\eta \in \mathbf{ch}(\iota)} \mu_\eta^\iota \tau_\eta^\iota$ are bounded because $\mathcal{A}(P_{i^\iota})$ is compact. The last equality is because of LP duality. Traversing indices k from $N-1$ back to 1, we define $\Phi^\iota := \rho_{i^\iota}[\ell(x^\iota, u^\eta, i^\eta) + \Phi^\eta; \eta]$, which boils down to

$$\Phi^\iota = \inf_{\substack{\tau^\iota, y^\iota \geq 0, \lambda^\iota \in \mathbb{R}, \\ \ell(x^\iota, u^\eta, i^\eta) + \Phi^\eta \leq \tau_\eta^\iota, l \in \mathbf{ch}(\iota), \\ \tau^\iota = F(P_{i^\iota})^\top y^\iota + \lambda^\iota \mathbf{1}_q}} b(P_{i^\iota})^\top y^\iota + \lambda^\iota,$$

for $\iota \in \Omega_k$. This formulation allows us to deconvolve the nested Markov risk measures. Indeed, $V_N^*(x_0, i_0)$ is the

optimal value of the following minimization problem

$$\begin{aligned} &\textbf{minimize}_{x, u, y \geq 0, \lambda, \tau} \ell(x_0, u^1, i_0) + b(P_{i_1})^\top y^1 + \lambda^1 \\ &\textbf{subject to} \quad \ell_N(x^\iota, i^\iota) \leq \tau_\eta^\iota, \eta \in \mathbf{ch}(\iota), \iota \in \Omega_N, \\ &\quad \tau^\iota = F(P_{i^\iota})^\top y^\iota + \lambda^\iota \mathbf{1}_q, \\ &\quad \ell(x^\iota, u^\eta, i^\eta) + b(P_{i^\eta})^\top y^\eta + \lambda^\eta \leq \tau_\eta^\iota, \\ &\quad x^\eta = f(x^\iota, u^\eta, i^\eta), \\ &\quad \eta \in \mathbf{ch}(\iota), \iota \in \Omega_k, k \in \mathbb{N}_{[0, N]}. \end{aligned}$$

Note that this formulation does not require the enumeration of the vertices of $\mathcal{A}(p)$ which, for instance, in the case of AV@R_α increases exponentially with the number of modes. The above optimization problem is solved at every time instant with x_0, i_0 being the current state and mode of the system. Solving this problem yields the optimal control actions u^{1*} at each node of the scenario tree. The first value, u^{1*} , defines the risk-averse MPC controller $\kappa_N^*(x, i) = u^{1*}(x, i)$. Note that in the particular case of an MJLS where stage-wise and terminal costs are quadratic and the constraints are polyhedral and/or ellipsoidal, we obtain a quadratically constrained quadratic program (QCQP) which can be solved very efficiently online as we show in Section 6. The above reformulation can be applied to risk measures whose ambiguity set is described by a set of conic inequalities (using conic duality) such as the entropic value-at-risk [1].

6 Illustrative example

Here we demonstrate the design of stabilizing risk-averse MPC controllers for a nonlinear system. We consider the following nonlinear Markovian switching system with three modes:

$$\begin{bmatrix} x_{k+1} \\ y_{k+1} \end{bmatrix} = A_{i_k} \begin{bmatrix} x_k \\ y_k \end{bmatrix} + c_{i_k} \begin{bmatrix} 1 - e^{y_k} \\ 1 - e^{x_k} \end{bmatrix} + B_{i_k} u_k. \quad (16)$$

The system matrices are

$$A_1 = \begin{bmatrix} 1 & 0.1 \\ 0.2 & 0.5 \end{bmatrix}, \quad A_2 = \begin{bmatrix} 0.1 & -0.5 \\ -0.5 & 0.5 \end{bmatrix}, \quad A_3 = \begin{bmatrix} 0.1 & -0.6 \\ 0.6 & 0.1 \end{bmatrix},$$

$$B_1 = \begin{bmatrix} 1.6 \\ 0.6 \end{bmatrix}, \quad B_2 = \begin{bmatrix} 0.1 \\ 0.9 \end{bmatrix}, \quad B_3 = \begin{bmatrix} 1 \\ 0 \end{bmatrix},$$

and parameters $c_1=0.2$, $c_2=-0.1$, $c_3=-0.3$. Stage-wise cost matrices are $Q_i = I$ and $R_i = 100 \cdot i$ for $i \in \{1, 2, 3\}$. The nominal and actual transition matrices are given by

$$P = \begin{bmatrix} 0.4 & 0.0 & 0.6 \\ 0.6 & 0.0 & 0.4 \\ 0.4 & 0.6 & 0.0 \end{bmatrix}, \quad P' = \begin{bmatrix} 0.33 & 0.0 & 0.67 \\ 0.56 & 0.0 & 0.44 \\ 0.33 & 0.67 & 0.0 \end{bmatrix}.$$

The nonlinear system is constrained to be inside the box $Y_1 = [-2.5, 2.5]^2 \times [-0.5, 0.5]$ for all three modes. Using $m = 0.5$ we compute the controller design parameters of Thm. 9 which are shown in Table 1. We take the terminal sets to be ellipsoidal $X_i^f = \{x^\top P_i^f x \leq r_i\}$. Finally, we simulate the system for different values of parameter α of AV@R_α after we formulate the problem as described

in Section 2.4, with initial condition $x_0 = (2, -2)$ and $i_0 = 1$. Resulting system trajectories are reported in Fig. 2. The proposed methodology successfully stabilizes the nonlinear system in the presence of uncertainty in the Markov transition matrix.

Table 1
Controller design parameters

i	β_i	δ_i		
		$\alpha = 1.0$	$\alpha = 0.9$	$\alpha = 0.5$
1	0.4421	0.2407	0.1783	0.1563
2	0.2210	0.3775	0.4121	0.3556
3	0.6631	0.1668	0.1130	0.0973

A similar effect is observed when inspecting the distribution of $\ell(x_k, u_k, i_k)$ for three MPC controllers. MPC controllers with higher α (closer to stochastic MPC) allow for higher costs, albeit with low probability. On the other hand, the risk-averse controller with $\alpha = 0.5$ (closer to minimax MPC) tends to produce cost distributions with shorter right tails. Interestingly, the point x_0 is not feasible for the worst case controller ($\alpha = 0$). The cost distributions are shown in Fig. 3.

Fig. 2. Trajectories of the closed-loop system with risk-averse MPC for $N = 6$ with (Left) $\alpha = 0.9$ and (Right) $\alpha = 0.5$. The green lines correspond to 1000 random simulations.

Fig. 3. Distribution of $\ell(x_k, u_k, i_k)$ estimated using 1000 randomly generated switching sequences. The cost of trajectories corresponding to higher α values are more spread out compared to $\alpha = 0.5$ and have a noticeably longer right tail.

7 Conclusions

We proposed a control methodology for constrained nonlinear Markovian switching systems. The proposed stability analysis framework hinges on dynamic programming and leads to the formulation of risk-based

Lyapunov-type conditions. These conditions can be translated into an LMI when the dynamics is linear, while, when the system is nonlinear a design methodology was proposed. In the case of MJLS, the resulting optimization problem can be formulated as a QCQP and can be solved efficiently online enabling its use in embedded applications.

We believe that risk-averse problems possess a favorable structure which can be further exploited to lead to parallelizable implementations akin to ones already developed for stochastic optimal control problems [22–24]. We plan to investigate risk-constrained formulations where we impose acceptable risk of violating the constraints instead of hard state/input constraints. This has a potential to make the overall design much less conservative.

Acknowledgements

This work was supported by the EU-funded H2020 project DISIRE, grant agreement No. 636834, the KU Leuven Research Council under BOF/STG-15-043, the Ford-KU Leuven Research Alliance under project No. KUL0023 and by Research Foundation Flanders, FWO, under project No. G086318N.

References

- [1] A. Ahmadi-Javid. Entropic value-at-risk: A new coherent risk measure. *J. Optimiz. Theory App.*, 155(3):1105–1123, 2012.
- [2] T. Asamov and A. Ruszczyński. Time-consistent approximations of risk-averse multistage stochastic optimization problems. *Math. Prog.*, 153(2):459–493, 2015.
- [3] D. Bertsekas, A. Nedić, and A. Ozdaglar. *Convex analysis and optimization*. Athena Scientific, 2003.
- [4] D.P. Bertsekas. *Dynamic Programming and Optimal Control*, volume II. Athena Scientific, 4th edition, 2012.
- [5] S. Bruno, S. Ahmed, A. Shapiro, and S. Stree. Risk neutral and risk averse approaches to multistage renewable investment planning under uncertainty. *Eur. J. Oper. Res.*, 250(3):979 – 989, 2016.
- [6] Y.-L. Chow and M. Pavone. A framework for time-consistent, risk-averse model predictive control: Theory and algorithms. In *ACC*, pages 4204 – 4211, Portland, USA, 2014.
- [7] Y.-L. Chow, S. Singh, A. Majumdar, and M. Pavone. A framework for time-consistent, risk-averse model predictive control: Theory and algorithms. *arXiv preprint arXiv:1703.01029v1*, 2017.
- [8] R.A. Collado, D. Papp, and A. Ruszczyński. Scenario decomposition of risk-averse multistage stochastic programming problems. *Ann. Oper. Res.*, 200:147–170, 2012.
- [9] O.L.V. Costa, M.D. Fragoso, and R.P. Marques. *Discrete-time Markov Jump Linear Systems*. Springer, 2005.
- [10] J. Goh and M. Sim. Distributionally robust optimization and its tractable approximations. *Oper. Res.*, 58(4-part-1):902–917, Aug 2010.
- [11] C.A. Hans, P. Sotasakis, A. Bemporad, J. Raisch, and C. Reincke-Collon. Scenario-based model predictive operation control of islanded microgrids. In *IEEE CDC*, Osaka, Japan, Dec 2015.

- [12] P. Hokayem, E. Cinquemani, D. Chatterjee, F. Ramponi, and J. Lygeros. Stochastic receding horizon control with output feedback and bounded controls. *Automatica*, 48(1):77–88, 2012.
- [13] D.H. Jacobson. Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Trans. Aut. Control*, 18(2):124–131, 1973.
- [14] A. Mesbah. Stochastic model predictive control: An overview and perspectives for future research. *IEEE CSM*, 36(6):30–44, Dec 2016.
- [15] P. Patrinos and H. Sarimveis. An explicit optimal control approach for mean-risk dynamic portfolio allocation. In *ECC*, pages 3364–3370, Jul 2007.
- [16] P. Patrinos and H. Sarimveis. Convex parametric piecewise quadratic optimization: theory and algorithms. *Automatica*, 47(8):1770–1777, 2011.
- [17] P. Patrinos, P. Sopasakis, H. Sarimveis, and A. Bemporad. Stochastic model predictive control for constrained discrete-time Markovian switching systems. *Automatica*, 50(10):2504–2514, Oct 2014.
- [18] J.B. Rawlings and D.Q. Mayne. *Model Predictive Control: Theory and Design*. Nob Hill Publishing, Madison, 2009.
- [19] R.T. Rockafellar. Solving stochastic programming problems with risk measures by progressive hedging. *Set-Valued and Variational Analysis*, Jul 2017.
- [20] R.T. Rockafellar and Roger J.-B. Wets. *Variational analysis*, volume 317. Springer, 2011.
- [21] A. Ruszczyński. Risk-averse dynamic programming for Markov decision processes. *Math. Progr.*, 125(2):235–261, Oct 2010.
- [22] A.K. Sampathirao, P. Sopasakis, A. Bemporad, and P. Patrinos. Distributed solution of stochastic optimal control problems on GPUs. In *IEEE CDC*, Osaka, Japan, Dec 2015.
- [23] A.K. Sampathirao, P. Sopasakis, A. Bemporad, and P. Patrinos. Proximal quasi-Newton methods for scenario-based stochastic optimal control. In *IFAC WC*, 2017.
- [24] A.K. Sampathirao, P. Sopasakis, A. Bemporad, and P. Patrinos. GPU-accelerated stochastic predictive control of drinking water networks. *IEEE TCST*, 26(2), 2018.
- [25] R.M. Schaich and M. Cannon. Robust positively invariant sets for state dependent and scaled disturbances. In *IEEE CDC*, pages 7560–7565, Dec 2015.
- [26] A. Shapiro, D. Dentcheva, and Ruszczyński. *Lectures on stochastic programming: modeling and theory*. SIAM, 2014.
- [27] B.P. Van Parys, D. Kuhn, P.J. Goulart, and M. Morari. Distributionally robust control of constrained stochastic systems. *IEEE TAC*, 61(2):430–442, 2016.
- [28] P. Yu and H. Xu. Distributionally robust counterpart in Markov decision processes. *IEEE TAC*, 61(9):2538–2543, 2016.

A Appendix

Proof of Lemma 5. Define $V_k := V(x_k, i_k)$ and, for fixed $x_0 \in \text{dom } V_N^*(\cdot, i_0)$ let $x_t := \phi(t, x_0, i_{[t-1]})$. We have

$$\begin{aligned}
 & \bar{\rho}_k \left[V_k - V_0 + c \sum_{t=0}^{k-1} \|x_t\|^2 \right] = \bar{\rho}_k \left[\sum_{t=0}^{k-1} V_{t+1} - V_t + c \|x_t\|^2 \right] \\
 & \leq \sum_{t=0}^{k-1} \bar{\rho}_k [V_{t+1} - V_t + c \|x_t\|^2] \\
 & = \sum_{t=0}^{k-1} \bar{\rho}_t [V_{t+1} - V_t + c \|x_t\|^2], \tag{A.1}
 \end{aligned}$$

where the inequality is because of the subadditivity of ρ (A1 and A4) and the last equality is because $V_{t+1} - V_t + c \|x_t\|^2$ is independent of $i_{t+1}, \dots, i_k$. In light of Cond. (ii) and given that $\bar{\rho}_{t+1}[V_{t+1} - V_t + c \|x_t\|^2] = \rho_{i_0} \circ \dots \circ \rho_{i_t} [V_{t+1} - V_t + c \|x_t\|^2; i_{t+1}] = \rho_{i_0} \circ \dots \circ \rho_{i_t} [V(f^{\kappa}(x_t, i_t), i_{t+1}) - V(x_t, i_t) + c \|x_t\|^2; i_{t+1}] \leq 0$, and because of (A.1) and property A2 we have that $\bar{\rho}_k [-V_0 + c \sum_{t=0}^{k-1} \|x_t\|^2] \leq \bar{\rho}_k [V_k - V_0 + c \sum_{t=0}^{k-1} \|x_t\|^2] \leq 0$. Using properties A3 and A4, $\bar{\rho}_k [\sum_{t=0}^{k-1} \|x_t\|^2] \leq v_0/c$ which proves the first part of Lemma 5.

By Cond. (ii), $\rho_{i_k} [V_{k+1} - V_k; i_{k+1}] \leq -c \|x_k\|^2 \leq -c \alpha_2^{-1} V_k \leq -\eta V_k$ for some $\eta \in (0, 1)$, so $\rho_{i_k} [V_{k+1}; i_{k+1}] \leq \beta V_k$ with $\beta := 1 - \eta \in (0, 1)$. We have $\rho_{i_0} [V_1; i_1] \leq \beta V_0$ and $\rho_{i_1} [V_2; i_2] \leq \beta V_1$, so $\rho_{i_0} [\rho_{i_1} [V_2; i_2]; i_1] \leq \beta \rho_{i_0} [V_1; i_1] \leq \beta^2 V_0$. Then, $\bar{\rho}_2 [V_2] \leq \beta^2 V_0$ and recursively

$$\bar{\rho}_k [V_k] \leq \beta^k V_0. \tag{A.2}$$

By the left hand side of Cond. (iii), $\|x_k\|^2 \leq 1/\alpha_1 V_k$ and applying $\bar{\rho}_k$ and using (A.2) and, subsequently the right hand side of Cond. (iii), $\bar{\rho}_k (\|x_k\|^2) \leq \bar{\rho}_k (V_k/\alpha_1) \leq \frac{1}{\alpha_1} \bar{\rho}_k (V_k) \leq \frac{1}{\alpha_1} \beta^k V_0 \leq \frac{\alpha_2}{\alpha_1} \beta^k \|x_0\|^2$. $\square$

Proof of Theorem 6. Let $\bar{X} \subseteq \text{dom } V_N^*$ be a compact UI set. By (8), $V_N^*(x, i) = \rho_i [V_{N-1}^*(f^{\kappa_N}(x, i), j); j] + \ell(x, \kappa_N^*(x, i), i)$. Then, for $(x, i) \in \bar{X}$,

$$\begin{aligned}
 & \rho_i \left[V_N^*(f^{\kappa_N}(x, i), j); j \right] - V_N^*(x, i) \\
 & = \rho_i \left[V_{N-1}^*(f^{\kappa_N}(x, i), j); j \right] - \ell(x, \kappa_N^*(x, i), i) \\
 & \quad - \rho_i \left[V_{N-1}^*(f^{\kappa_N}(x, i), j); j \right] \\
 & \leq -\ell(x, \kappa_N^*(x, i), i) \leq -c \|x\|^2.
 \end{aligned}$$

The first inequality is because $V_N^* \leq V_{N-1}^*$ and property A2. We have that $V_N^*(x, i) \leq \ell_N(x, i) \leq d \|x\|^2$ for all $x \in X_i^f$. Because of Cond. (iii), we may find $\epsilon > 0$ such that $\mathcal{B}_\epsilon \subseteq X_i^f$, for $i \in \mathcal{N}$. By Cond. (iv), there is an

$M > d\epsilon^2$. Then, for all $x \in \bar{X}_i \setminus X_i^f$, $\frac{M}{\epsilon^2}\|x\|^2 \geq M \geq V_N^*(x, i)$. Because of Cond. (i) and the definition of $\mathbf{T}$, we have that $V_t^*(x, i) \geq c\|x\|^2$ for all $(x, i) \in \mathbf{dom} V_t^*$ for $t \in \mathbb{N}_{[1, N]}$. The proof is complete since $V = V_N^* + \delta_{\bar{X}}$ satisfies all conditions of Lemma 5. $\square$

Proof of Theorem 9. Define $e(x, i) = f^\kappa(x, i) - \hat{f}^\kappa(x, i)$. By Assumption 8 and since $f^\kappa(0, i) = 0$ for all $i \in \mathcal{N}$, $\|e(x, i)\| \leq \beta_{i/2}\|x\|^2$. It is $\Delta(x, i) = \rho_i[f^\kappa(x, i)^\top P_j f^\kappa(x, i); j] - \rho_i[x^\top \bar{A}_i^\top P_j \bar{A}_i x; j]$. Since $\rho_i[\cdot]$ is convex and monotone, it is nonexpansive with respect to the infinity norm [26, p. 302], thus for $x \in X_i^f$

$$\begin{aligned} |\Delta(x, i)| &\leq \max_{j \in \mathbf{cov}(i)} |f^\kappa(x, i)^\top P_j f^\kappa(x, i) - x^\top \bar{A}_i^\top P_j \bar{A}_i x| \\ &= \max_{j \in \mathbf{cov}(i)} |e(x, i)^\top P_j e(x, i) + 2x^\top \bar{A}_i^\top P_j e(x, i)| \\ &\leq \max_{j \in \mathbf{cov}(i)} \|P_j\| \left(\frac{\beta_i^2}{4} \|x\|^4 + \beta_i \|\bar{A}_i\| \|x\|^3 \right) \leq \sigma_i \|x\|^2 \end{aligned}$$

Therefore, $\mathcal{L}\ell_N(x, i) = \mathcal{L}'\ell_N(x, i) + \Delta(x, i) \leq -x^\top (\bar{Q}_i + (m_i - \sigma_i)x) \leq -\ell(x, \kappa(x, i), i)$, for all $x \in X_i^f$ and since X^f is UI, $\mathbf{T}\ell_N \leq \ell_N$. The assertion follows from Thm. 6. $\square$