Data-Driven Control of Stochastic Systems

Representation, Prediction, and Optimal Control

Dissertation (monograph) approved by the
Doctoral Degree Committee of
Hamburg University of Technology

in pursuit of the academic degree of

Doktor-Ingenieur (Dr.-Ing.)

written by
Guanru Pan

from
Zhejiang, China

2025

The chair of the examination committee: Prof. Dr.-Ing. Gerhard Bauch, Hamburg University of Technology

1. Reviewer: Prof. Dr.-Ing. Timm Faulwasser, Hamburg University of Technology
2. Reviewer: Prof. Dr.-Ing. Frank Allgöwer, University of Stuttgart
3. Reviewer: Prof. Dr. Paolo Rapisarda, University of Southampton

Date of oral examination: June 6, 2025

Summary

This thesis develops a framework for data-driven control of linear stochastic systems by combining two foundational approaches: the fundamental lemma by Jan C. Willems et al., a cornerstone of behavioral systems theory, and Polynomial Chaos Expansions (PCEs) by Norbert Wiener, a powerful tool for uncertainty representation. It addresses the challenges of representing, predicting, and controlling stochastic systems directly from data, eliminating the need for explicit parametric models.

First, this thesis establishes a formal framework for the behavioral theory of stochastic systems. It demonstrates that the behaviors of L^2 random variables and their PCE coefficients align with the underlying system dynamics. This extends the fundamental lemma by Willems et al. to stochastic Linear Time-Invariant (LTI) systems, enabling the use of past input/output/disturbance data to predict future stochastic system responses.

Another contribution is a data-driven stochastic prediction scheme for LTI systems subject to unmeasured disturbances. The proposed method consolidates the unmeasured disturbances into a residual disturbance and estimates its realizations using a least-squares approach. Furthermore, a regularized stochastic predictor is introduced to improve the robustness against estimation errors. Using the PCE framework, this approach efficiently quantifies stochastic predictions via confidence intervals based on higher-order Chebyshev's inequalities. A real-world case study using residential building data validates the effectiveness of the proposed prediction method.

This thesis also investigates stochastic optimal control for LTI systems under exogenous disturbances. We establish sufficient conditions for the equivalence between the model-based and the data-driven Optimal Control Problems (OCPs) and for the optimality of affine disturbance-feedback policies. To further enhance distributional robustness and mitigate uncertainty regarding future disturbance distributions, we work with distributionally robust chance constraints and provide their exact PCE formulations as tractable second-order cone constraints.

Furthermore, we propose a data-driven output-feedback predictive control scheme that requires knowledge of only the first two moments of the disturbance. The proposed scheme is based on a tailored data-driven OCP with a stochastic initial condition and terminal ingredients. Through online design of the initial conditions, we provide sufficient conditions to ensure recursive feasibility of the proposed output-feedback scheme. Additionally, by leveraging a data-driven design of the terminal ingredients, we establish an average asymptotic performance bound for the closed-loop system and provide a robustness analysis of the closed-loop performance.

In conclusion, this thesis bridges the gap between data-driven control and stochastic control by integrating behavioral systems theory with PCE. It introduces innovative tools for representing stochastic behavior, ensuring robust predictions, and designing optimal control policies for stochastic systems.

For my wife,
with all my love.

Acknowledgments

This thesis is the result of my doctoral studies in the Optimization and Control Group of Prof. Timm Faulwasser at the Institute of Energy Systems, Energy Efficiency and Energy Economics, TU Dortmund University, during the first three years, and at the Institute of Control Systems, Hamburg University of Technology, in the final year. I would like to express my deepest gratitude to all those who have supported me throughout this journey and made this thesis possible.

First and foremost, I wish to express my sincere appreciation to my supervisor, Prof. Timm Faulwasser, for his invaluable academic guidance, insightful discussions, and constant encouragement. His advice has not only shaped this research but also broadened my perspective beyond academia. I would also like to thank the external collaborators whose discussions have greatly enriched this work.

I am deeply grateful to my external reviewers, Prof. Paolo Rapisarda from the University of Southampton and Prof. Frank Allgöwer from the University of Stuttgart, for their careful evaluation of my dissertation and their valuable feedback.

My heartfelt thanks go to my colleagues Alex, Gösta, Ruchuan, Debbie, Jens, Julian, Maísa, Mohammad, Rahul, Oleksii, Dirk, and Kirsten for the pleasant collaboration and friendly working atmosphere. It was a great pleasure to work with you all.

Finally, I owe my deepest gratitude to my family, and especially to my wife, Boshu, for her unwavering support, patience, and companionship throughout this journey. Your love and encouragement have been my greatest strength.

Contents

Index of Notation

Abbreviations and Acronyms

LTI	Linear Time-Invariant
MIMO	Multi-Input Multi-Output
OCP	Optimal Control Problem
PCE	Polynomial Chaos Expansion
MPC	Model Predictive Control
LQR	Linear Quadratic Regulator
ARX	AutoRegressive with eXogenous input
RMSE	Root Mean Square Error
i/o	input/output
pdf	probability density function
cdf	cumulative density function

Mathematical Notations

$\mathbb{Z}$	set of all integers
$\mathbb{N}$, $\mathbb{N}^+$	set of non-negative (positive) integers
$\mathbb{R}$, $\mathbb{R}^+$	set of all (positive) real numbers
$\mathbb{I}_{[t_1,t_2]}$	set of all integers between t_1 and t_2
$(\mathbb{R}^{n_s})^{\mathbb{Z}}$	space of all trajectories $s : \mathbb{Z} \to \mathbb{R}^{n_s}$
$\mathfrak{B}$	LTI behavior, cf. Definition 2.1
$\boldsymbol{\sigma}\mathfrak{B}$	forward shifted LTI behavior
$\mathfrak{B}_{[t_1,t_2]}$	restricted behavior to interval $[t_1, t_2]$
$\mathfrak{B}_T$	set of all restricted behavior with length T
$n^\star$	minimal state of the LTI behavior $\mathfrak{B}$
$\mathcal{O}_i(A, C)$	i-th order observability matrix of $\mathfrak{B}(A, B, C, D)$, cf. (2.14)
$\mathcal{R}_i(A, B)$	i-th order reachability matrix of $\mathfrak{B}(A, B, C, D)$, cf. (2.15)
ℓ	lag of an ARX representation $\mathfrak{B}(\Xi, D)$, cf.(2.19)
$\ell^\star$	minimal lag of the LTI behavior $\mathfrak{B}$
$\mathcal{H}_N(s_{[T-1]})$	Hankel matrix of order N by the signal sequence $s_{[T-1]}$, cf. (2.24)
ω, Ω	outcome ω as elements of the sample space Ω
$\mathcal{F}$	general σ-algebra
$\sigma(V)$	σ-algebra induced by random variable V, cf. (3.1)
μ	probability measure $\mu : \mathcal{F} \to [0, 1]$
$(\Omega, \mathcal{F}, \mu)$	probability space
$L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$	set of all L^2 random variables, cf. Definition 3.3
$\langle V, \tilde{V} \rangle$, $\|V\|$	inner product and norm of random variables, cf. (3.3)
$\mathbb{E}[V]$, $\Sigma[V]$	mean and covariance of $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$, cf. (3.5)
$\mu_n(M)$	n-th central moment of $M \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$
$\{\phi^j(\xi)\}$	polynomial basis with stochastic germ $\xi \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\xi})$
$\mathcal{P}_1(\xi)$	polynomials with degree of at most 1

f_V	pdf of random variable V
F_V	cdf of random variable V
$\mathbb{P}[A]$	probability of event A
S, s, s	stochastic process, realization signal, trajectory of PCE coefficients
$\mathfrak{S}$	stochastic behavior, cf. (4.3)
$\mathfrak{C}$	behavior of PCE coefficients, cf. (4.10)
$\mathfrak{M}$	behavior of mean and covariance, cf. (4.12)
$\prod_{j\in\mathbb{N}} \mathfrak{B}$	Infinite Cartesian product of the behavior $\mathfrak{B}$, i.e., the set of all sequences $(\mathsf{s}^j)_{j\in\mathbb{N}}$ such that $\mathsf{s}^j \in \mathfrak{B}$ for all $j \in \mathbb{N}$.
$\prod_{i=1}^{n_s} \mathfrak{B}$	Finite Cartesian product of the behavior $\mathfrak{B}$, i.e., the set of tuples $(s_1, \cdots, s_{n_s})$ with $\mathsf{s}^j \in \mathfrak{B}$ fo $i = 1, \cdots, n_s$.
$\mathfrak{S}^V$	corrupted stochastic behavior with residual disturbance V, cf. (5.1)
δ^{ij}	Kronecker delta, defined as $\delta^{ij} = 1$ if $i = j$, and $\delta^{ij} = 0$ if $i \neq j$
π	control policy
$Q \succeq 0$	positive semi-definite matrix
$R \succ 0$	positive-definite matrix
$\mathrm{trace}(Q)$	trace of a positive semi-definite matrix
$\lambda_{\max}(Q)$	maximal eigenvalue of a positive semi-definite matrix
$A^\dagger$	Moore-Penrose inverse of a matrix, cf. (3.10)
$\|A\|$	Frobenius norm of matrix A, i.e., $\|A\| \doteq \sqrt{\mathrm{trace}(AA^\top)}$
$\|v\|$	2-norm of vector v, i.e., $\|v\| \doteq \sqrt{v^\top v}$
$\|v\|_Q$	Q weighted norm of vector v, i.e., $\|v\|_Q \doteq \sqrt{v^\top Q v}$
$\mathrm{colspan}(A)$	column span of a matrix
$\mathbb{G}\left((m,\Gamma),(\bar{m},\bar{\Gamma})\right)$	Gelbrich distance between moment pairs, cf. Definition 6.1
$\mathbb{W}(\mu,\bar{\mu})$	Wasserstein matrix between distributions $\mu, \bar{\mu}$, (6.19)
$\mathcal{C}(m,\Gamma)$	Cantelli ambiguity set, cf. (3.20)
$\mathcal{G}_\rho(\bar{m},\bar{\Gamma})$	Gelbrich ambiguity set, cf. (6.18)
$\mathcal{W}_\rho(\mu_{\bar{V}})$	Wasserstein ambiguity set, cf. (6.26)

1. Introduction

In many real-world applications, stochastic disturbances, such as uncertain environmental conditions and random exogenous disturbances, can significantly impact both the performance and safety of the system. In automotive applications, for instance, the uncertain behavior of other vehicles and road users must be carefully considered to ensure reliable operation (Carvalho et al., 2014). Similarly, building climate control systems have to deal with uncertain weather conditions and highly variable occupancy rates (Drgoňa et al., 2020), while wind turbine control faces the task of handling fluctuating wind speeds (Chen et al., 2010). Moreover, power systems and distributed multi-energy systems are subject to uncertainties in both energy generation and consumption (Mühlpfordt et al., 2019; Özmeteler et al., 2024).

Despite advances in sensors and machine learning techniques for predicting disturbances, these predictions remain inherently uncertain and reliable only in a probabilistic sense. In addition to such measurable and statistically modeled uncertainties—referred to as *structured stochastic disturbances*—systems are also subject to *unstructured stochastic disturbances*, which are unmeasured or statistically unmodeled.

Figure 1.1 categorizes exogenous disturbances from two perspectives:

i) Whether their realizations can be measured: measured, estimated, or unmeasured.

ii) Whether their statistics can be modeled: modeled, partially modeled, or unmodeled.

As shown in Figure 1.1, *structured* disturbances refer to the realization-measured and statistically modeled disturbances, while the others are referred to as *unstructured* disturbances. This spectrum of uncertainties, from structured to unstructured disturbances, highlights the importance of stochastic control approaches in ensuring system safety and performance under diverse stochastic conditions.

Another emerging trend in control research is the paradigm shift from model-based to data-driven approaches (Dörfler et al., 2023a). This shift is driven by the fact that, in many practical scenarios, our understanding of the underlying system dynamics is limited. A convincing example is building control systems, where their dynamics are influenced by complex interactions between HVAC devices, thermal properties, and environmental factors. The construction of a first-principles model for such systems is cumbersome and subject to inaccuracies, especially if the states are not properly chosen. Moreover, the estimation of state-space representations through subspace

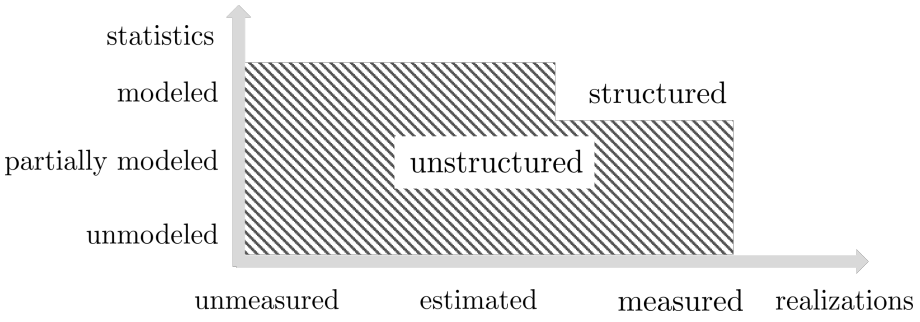

Figure 1.1.: Spectrum of exogenous disturbances

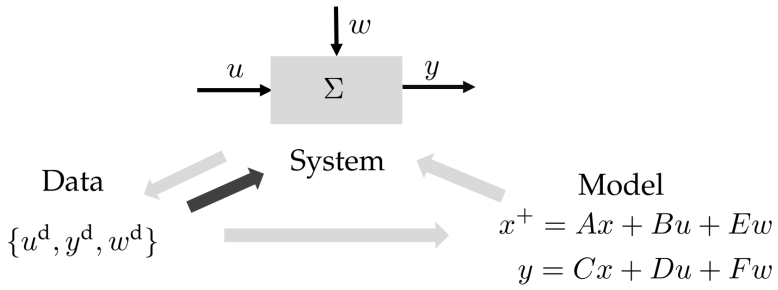

Figure 1.2.: Comparison of the model-based routine (grey arrows) and the data-driven routine (black arrows) for representing an unknown dynamical system Σ exposed to exogenous disturbance w

identification often lacks interpretability, as the identified states may not correspond to meaningful physical quantities. Furthermore, system parameters in buildings, such as thermal insulation or HVAC efficiency, tend to change over time due to aging or seasonal variations, necessitating frequent model updates. In building control, however, we often have access to input, output, and some disturbance data, such as control signals, indoor temperatures, and weather conditions. Therefore, data-driven approaches offer an effective alternative to traditional modeling, efficiently ensuring building comfort and energy efficiency (Bilgic et al., 2024; Di Natale et al., 2022; Lian et al., 2023).

As illustrated in Figure 1.2, rather than following the conventional routine of first measuring the system data, identifying a parametric model, and then using this model to represent the system (Katayama, 2005; van Overschee et al., 2012), the data-driven methodology—aligned with *the fundamental lemma* by Willems et al., 2005—suggests using the available data directly to represent the system; see also De Persis et al., 2019 and Markovsky et al., 2021 for recent reviews.

The fundamental lemma highlights a pivotal insight: for a controllable Linear Time-Invariant (LTI) system, its input/output trajectories lie exactly within the column space spanned by Hankel matrices constructed from past data, provided that the input signal is sufficiently persistently exciting. This lemma has sparked various extensions and refinements. The persistency of excitation is refined by van Waarde, 2021, while the controllability assumption is relaxed by Yu et al., 2021. In addition, Bilgic et al., 2022 and Lian et al., 2023 provide strategies for the online adaptation of Hankel matrices. There have also been recent extensions to linear descriptor systems (Schmitz et al., 2022) and to continuous-time LTI systems (Lopez et al., 2022; Schmitz et al., 2024). Moreover, generalizations in frequency domain have been explored using the discrete Fourier transform by Meijer et al., 2024 and the discrete wavelet transform by Sathyanarayanan et al., 2023. Beyond the usual LTI setting, there are tailored variants for linear parameter-varying systems (Verhoek et al., 2021), for linear time-varying ones (Nortmann et al., 2023), and for linear time-delay systems (Rueda-Escobedo et al., 2022).

Dating back to Willems, 1986, 1991, the fundamental lemma is deeply rooted in the behavioral theory of dynamical systems. Unlike the state-space approach, which represents dynamical systems primarily through *latent* state variables, behavioral theory shifts the focus to *manifest* variables—namely, the inputs, outputs, and disturbances. The *manifest behavior* of a system is defined as the set of all possible trajectories of its manifest variables. Importantly, the behavior of an LTI system is proven to form a lin-

ear subspace of the signal space, which, as shown in the fundamental lemma, can be spanned by past trajectories.

As illustrated in Figure 1.2, this paradigm shift in control involves not only the transition from model-based to data-driven approaches but also the adoption of a behavioral perspective over the state-space representation. Traditional model-based routines, such as subspace identification approaches, focus on deriving state-space models to describe system dynamics. However, since states are latent variables that are typically not directly measurable, additional steps, such as designing state observers, are required to implement state-feedback controllers.

In contrast, data-driven approaches leverage past measured external data to directly represent system behavior. This enables the prediction of future output trajectories and the design of output-feedback controllers, bypassing the need for intermediate state reconstruction (Berberich et al., 2020; Coulson et al., 2019; Yang et al., 2015). For a comprehensive overview of behavioral theory in control, see Markovsky et al., 2021.

In particular, the fundamental lemma enables data-driven predictive control with output feedback as shown by Yang et al., 2015, Coulson et al., 2019, and Berberich et al., 2020. In addition, Fiedler et al., 2021 explore the connections between data-driven predictive control and subspace predictive control considered by Favoreel et al., 1999 and Sedghizadeh et al., 2018. Applications of data-driven predictive control across various domains are discussed by Carlet et al., 2023, Bilgic et al., 2022, Wang et al., 2023, and Lian et al., 2023.

Interestingly, while behavioral systems theory has achieved significant success in data-driven control, its application to stochastic systems remains relatively less explored. This gap aligns with the open challenges highlighted by Markovsky et al., 2021.

The ultimate paper of Willems, 2012 discusses behavioral ideas for stochastic systems, focusing primarily on open stochastic static systems, their interconnections, and the construction of appropriate σ-algebras. Subsequently, Baggio et al., 2017 extend this framework to a canonical kernel representation of stochastic LTI processes, which, as demonstrated in Chapter 5, corresponds to the stochastic behavior of systems exposed to unstructured disturbances. Yet, none of these works covers data-driven representations of stochastic LTI systems. Following a different perspective, Pola et al., 2015 and Pola et al., 2016 employ behavioral ideas to study equivalence concepts for stochastic linear systems in discrete-time and continuous-time settings, without an explicit definition of the stochastic behavior. In our work (Faulwasser et al., 2023; Pan et al., 2023c), which is also revisited in this thesis, we analyze the stochastic behavior of LTI dynamical systems influenced by structured disturbances. In summary, while behavioral systems theory has been extended to stochastic settings in various ways, a comprehensive data-driven representation of stochastic LTI systems remains an open challenge, which this thesis aims to address.

1.1. Towards Data-Driven Control for Stochastic Systems

In short, the key research question of this thesis is:

How to *represent*, *predict*, and *control* stochastic systems in data-driven fashion?

An essential step in addressing this question is to establish an appropriate representation of stochastic uncertainties.

Stochastic Behaviors via Random Variables

Traditional approaches often rely on probability densities. However, this becomes increasingly cumbersome, particularly when propagating uncertainty through system dynamics. Consider, for example, two independent random variables X and Y that follow uniform distributions with support $[-0.5, 0.5]$. The sum $Z = X + Y$ requires the convolution of their probability density functions (pdfs) to determine the pdf of Z, i.e.,

$$f_Z(z) = \int_{-\infty}^{\infty} f_X(x) f_Y(z - x) \, \mathrm{d}x.$$

This involves piecewise calculations: $f_Z(z) = z - 1$ for $-1 \leq z < 0$, $f_Z(z) = 1 - z$ for $0 \leq z \leq 1$, and $f_Z(z) = 0$ otherwise. While it is straightforward in this simple case, such computations become non-trivial for more complex distributions or higher-dimensional systems. For discussions on probability density propagation with empirical distributions, see Aolaritei et al., 2023a and Aolaritei et al., 2023b.

Instead, we shift the focus from probability densities to random variables, defined as measurable functions on a common probability space. This offers two clear advantages. On the one hand, instead of describing uncertainties through individual probabilities and correlations, the common probability space encodes the probabilistic information for all uncertainties. That is, it effectively represents the uncertain environment to which the system is exposed. On the other hand, the random-variable description of the stochastic system follows the same dynamics as the deterministic behavior of the system.

These insights lead to the conceptualization of stochastic behavior as trajectories of random variables defined on a common probability space. Crucially, the stochastic trajectories of the system remain compatible with its deterministic behavior. This observation directly links the stochastic behavior to the deterministic system dynamics and the common probability space, while the latter represents the uncertain environment to which the system is exposed. This perspective allows for a unified treatment of both stochastic and deterministic behavior within a cohesive framework; see Chapter 4.

The definition of stochastic behavior proposed by Willems, 2012 and Baggio et al., 2017 differs from this perspective. Instead of considering trajectories of random variables fully aligned with deterministic behavior, these works define stochastic behaviors as probabilistic deviations from the deterministic behavior caused by unstructured (unmeasured or unmodeled) disturbances. However, this difference can be resolved through a shift in interpretation as detailed in Chapter 5. We view the deterministic behavior considered by Willems, 2012 and Baggio et al., 2017 as the "uncorrupted" deterministic behavior which represents system dynamics without unstructured disturbances. Then, the stochastic behavior therein remains compatible with the "corrupted" deterministic system dynamics. In this way, their definition represents a special case of our broader framework.

To address challenges posed by unstructured disturbances, we unify the approaches of Willems, 2012, Baggio et al., 2017, and our own work to address the representation, prediction, and optimal control of systems affected by such disturbances. Specifically,

we propose to consolidate unstructured disturbances into a residual disturbance that can be estimated and partially statistically modeled. These concepts are further developed in Chapters 5 and 7.

Tractable Representation via Wiener's Polynomial Chaos

Representing stochastic behavior in terms of random variables poses additional challenges due to their infinite-dimensional nature. A common approach is to use statistical moments–such as the expectation, covariance, skewness, and kurtosis–as finite-dimensional approximations. These representations are typically limited to simple cases such as i.i.d. Gaussian uncertainties. Beyond Gaussian assumptions, capturing higher-order moments becomes necessary, particularly for non-Gaussian random variables or for non-linear maps. However, the nonlinear relationship between higher-order moments and random variables can complicate their practical use.

Remarkably, an alternative approach to representing random variables has been proposed by Wiener, 1938. Norbert Wiener suggests using series expansions expressed in suitable polynomial bases of the underlying probability space. This method relies on the Hilbert space of L^2 random variables, a linear function space with an inner product that contains all random variables with finite variance. Specifically, Wiener, 1938 shows that all L^2 random variable can be represented by polynomials of Gaussian variables using a Hermite basis. Building on this foundation, Xiu et al., 2002 provide the framework of generalized polynomial chaos, which extends the stochastic germ (the arguments used in polynomials) to uniform, Beta-, and Gamma-distributed random variables. Additionally, Witteveen et al., 2006 propose constructing PCE bases with arbitrary random variables as the germ using the Gram-Schmidt procedure.

These approaches, now widely known as the *Wiener's chaos expansion* or *Polynomial Chaos Expansion* (PCE), provide a linear structure for representing a broad class of Gaussian and non-Gaussian random variables in contrast to the nonlinear nature of moment-based representations. PCEs have also been extensively used in systems and control, including system analysis (Ahbe et al., 2020; Kim et al., 2013; Lucia et al., 2017), stochastic optimal control (Ou et al., 2021; Ou et al., 2025), and optimization in power systems (Mühlpfordt, 2020). Polynomial basis is also used for data-driven analysis of continuous-time systems (Rapisarda et al., 2024). Therein, orthogonal polynomial expansions are not applied to L^2 random variables. Instead, the continuous input and state trajectories are decomposed into discrete sequences of expansion coefficients.

For tractable representations, PCEs are typically truncated to a finite number of terms (Mühlpfordt et al., 2017). However, exact finite-dimensional expansions of random variables can always be achieved by selecting an appropriate basis. Intuitively, each scalar L^2 random variable admits a two-term exact PCE, expressed as the sum of its mean and the standard deviation multiplied by its normalized version. In Chapter 3, we generalize this affine expansion to vector-valued random variables. Furthermore, focusing on the propagation of exogenous disturbances through LTI dynamics, we show that by concatenating independent bases induced from affine expansions of disturbances, all random variables involved in the dynamics can be represented using finite-dimensional PCEs. We remark that the constructed PCE basis requires only first-degree polynomials.

Moreover, within the framework of finite-dimensional exact PCEs, moment infor-

mation can be efficiently constructed (Lefebvre, 2020). Building on this, in Chapter 5, we leverage PCE to quantify stochastic predictions by defining confidence intervals based on the Chebyshev's inequality using fourth-order moments. Compared to confidence intervals derived from Gaussian assumptions or the Chebyshev's inequality relying only on the first two moments, the proposed approach offers a balance between achieving tighter bounds and maintaining a higher coverage rate.

Stochastic Variants of the Fundamental Lemma by Willems et al.

In the line of emerging extensions of the fundamental lemma, several works have proposed stochastic variants from different perspectives.

Considering both process noise and measurement noise, Coulson et al., 2023 introduce a quantitative notion of persistency of excitation by analyzing the minimum singular value of an input Hankel matrix. With this persistency of excitation condition, their proposed variant of the fundamental lemma guarantees lower bounds on the singular values of the data Hankel matrix, thereby enhancing the robustness of data-driven predictions against stochastic disturbances. This extension provides valuable guidance on designing input trajectories that robustly excite the system. However, it does not address the data-driven propagation of stochastic disturbances through system dynamics. Hiremath et al., 2022 propose a stochastic variant of the fundamental lemma for representing trajectories of stochastic LTI systems with additive Gaussian processes and measurement noise. However, their approach is limited to Gaussian disturbances and cannot be extended to non-Gaussian cases.

Ferizbegovic et al., 2021 propose a variant of the fundamental lemma leveraging second-order moments. Specifically, they construct Hankel matrices using estimated correlations of past inputs and outputs to predict future deterministic trajectories. Hence, it is, in fact, a deterministic variant of the fundamental lemma refined by stochastic techniques. By estimating correlation functions from past data, this approach reduces the influence of measurement noise. However, it requires the input signal to be a persistently exciting stationary process during data collection, which increases the amount of data needed.

Interestingly, Ferizbegovic et al., 2021 note that (though do not formally explore) the possibility of using Hankel matrices constructed from past data to predict future correlations between stochastic inputs and outputs. This observation is rooted in the fact that input/output correlations evolve according to the same linear dynamics as the system itself. We delve into this concept in detail as a second-order moment representation of stochastic behavior in Chapter 4.

In our work (Faulwasser et al., 2023; Pan et al., 2023c), *the fundamental lemma by Willems et al.* and *the polynomial chaos expansion by Wiener et al.* are combined to propose a stochastic variant of the fundamental lemma to enable data-driven representations of stochastic behaviors. As discussed earlier, the stochastic behavior in random variables is defined to remain compatible with the dynamics encoded in the deterministic behavior. Furthermore, the linear nature of PCE ensures that the behavior of PCE coefficients is also inherently compatible with the underlying LTI dynamics. Hence, the Hankel matrices constructed by measured deterministic data are capable of characterizing the dynamics of PCE coefficients, while the PCE framework allows capturing random variables exactly. The relations among the stochastic behavior, the determin-

istic behavior, the PCE coefficient behavior, and their data-driven representations are elaborated by Faulwasser et al., 2023. These concepts will be revisited and further explored in Chapter 4.

Towards Data-Driven Predictive Control of Stochastic Systems

While stochastic extensions of the fundamental lemma provide valuable tools for data-driven representations of stochastic system behavior, their application in predictive control requires additional considerations.

One trend advocated in the literature to handle noise in both Hankel matrix and initial condition is to use different regularizations to enhance robustness. Regularization for data-driven control is first introduced by Coulson et al., 2019, where its two-norm variant forms the foundation for robust formulations by Huang et al., 2021, 2023. Additionally, its one-norm variant inspires distributionally robust formulations by Coulson et al., 2019.

In combination with regularization techniques, sufficient conditions for recursive feasibility, stability, and robustness of output-feedback predictive control for deterministic LTI systems subject to noise-corrupted measurements are given by Berberich et al., 2020, Berberich et al., 2021, and Bongard et al., 2023. Recently, Berberich et al., 2025 offer a comprehensive overview of the advancements of theoretical guarantees for data-driven control.

Following another routine, Wang et al., 2025 apply the fundamental lemma to the innovation-form state-space model, which arises in the traditional stochastic subspace identification approach (van Overschee et al., 2012). In their work, they estimate past innovations from input and output data and then eliminate their influence. These past innovations account for both process and output noise. By projecting the Hankel matrix onto the null space of the recorded innovations, they effectively remove the influence of noise corrupted in the Hankel matrices.

However, while these works focus on mitigating the effects of noise in past data, they do not propagate the stochastic uncertainty through the system dynamics to ensure future stochastic satisfaction of constraints.

Focusing similarly on noise-corrupted past data, Yin et al., 2021 propose a maximum likelihood estimation approach for future output trajectories. This approach provides probabilistic prediction in second-order moments. Based on this, Yin et al., 2024 further include chance constraints within a stochastic data-driven predictive control. However, recursive feasibility and stability guarantees are not provided.

Köhler et al., 2022a suggest identifying a multi-step predictor rather than a state-space model. They show that by combining the identified multi-step predictors with the predictive control framework, it avoids the sequential propagation of parametric uncertainty, which can result in significant conservatism in state-space approaches. Based on the identified multi-step predictor, Balim et al., 2024 provide a stochastic predictive control approach with a conservatism-reduced chance constraint satisfaction. We remark that due to the intermediate identification step, these approaches represent a refinement of the model-based routine in Figure 1.2. Moreover, closed-loop guarantees are not addressed therein.

Kerz et al., 2023 provide a tube-based data-driven framework with closed-loop guarantees. However, it requires the measured state data to resemble the traditional routine

of the state-feedback tube-based stochastic MPC (Lorenzen et al., 2016; Mesbah, 2016). Teutsch et al., 2024 further propose a sampling-based approach following the offline sampling concept by Lorenzen et al., 2017 to provide closed-loop guarantees for an output-feedback stochastic MPC. However, their methods require the stochastic disturbances to be bounded.

Our proposed approaches (Pan et al., 2023b, 2025a,b) pave the way for output-feedback data-driven stochastic predictive control with closed-loop guarantees. Combined with the PCE framework, it accommodates a broad range of stochastic disturbances, including Gaussian, non-Gaussian, bounded, and unbounded distributions. In our prior work (Pan et al., 2025b), a *state-feedback* predictive control scheme with closed-loop guarantees is proposed relying on the concept of backup initial condition as in the model-based approaches (Farina et al., 2013; Farina et al., 2015, 2016). However, this approach relies on a PCE basis growing in dimension as the closed loop evolves. Building on this, we refine recursive feasibility in Pan et al., 2023b by considering interpolation between the measured initial condition and the backup condition as in the model-based approaches (Köhler et al., 2022b; Schlüter et al., 2023). Compared to the binary selection of initial conditions by Pan et al., 2025b, the interpolation of initial conditions avoids solving the optimization twice and hence is more efficient. Furthermore, in Pan et al., 2025a, we address the open problem of closed-loop analysis for data-driven stochastic output-feedback predictive control. Moreover, in contrast to Pan et al., 2025b, the proposed approach uses a fixed-dimension basis and establishes average asymptotic performance bounds through data-driven terminal ingredient design without requiring state information. These results will be summarized and further explored in Chapter 7.

1.2. Thesis Outline and Contributions

By combining *Wiener's polynomial chaos* with *Willems' behavioral systems theory*, this thesis proposes a unified framework for the data-driven control of stochastic systems. This approach enables the representation, prediction, and control of stochastic LTI systems using measured data without requiring a full model of the system. Figure 1.3 provides the outline of this thesis.

After Chapter 1, the remainder aims to tackle the key question:

How to *represent*, *predict*, and *control* stochastic systems in data-driven fashion?

Chapters 2–4 lay the foundation focusing on the transition from representing the deterministic behavior to *representing the stochastic behavior* of linear systems.

In Chapter 2, we view dynamical systems through the lens of behavioral systems theory. That is, we characterize system dynamics via the behavior of their input and output trajectories. Following Willems, 1986, 1991, we revisit the definition of Linear Time-Invariant (LTI) systems in the behavioral setting. Moreover, we compare various parametric representations of LTI systems and recap the non-parametric representation by the fundamental lemma.

In Chapter 3, we turn to the second key ingredient of stochastic systems: representations of stochastic uncertainties. We revisit essential concepts in probability theory to rigorously define random variables and their distributions. Furthermore, we employ

the framework of PCE to represent random variables through deterministic expansion coefficients in suitable polynomial bases. Furthermore, we revisit stochastic processes and their PCE representations, which serve as the stochastic counterpart to deterministic trajectories.

Chapter 4 presents the first main contribution of the thesis, i.e. parametric and non-parametric representations for stochastic LTI systems, based on the results by Pan et al., 2023c and Faulwasser et al., 2023. In this chapter, we rigorously define the stochastic behavior of dynamical systems by incorporating the uncertain environment into the deterministic behavior leveraging random variables. Moreover, we prove that the stochastic behavior of a LTI system remains linear and time invariant.

Since stochastic behaviors rely on infinite-dimensional stochastic processes and L^2 random variables in a Hilbert space, we introduce more tractable representations using PCE coefficients and statistical moments. To characterize the underlying relations between behaviors in terms of L^2 random variables, PCE coefficients, statistical moments, and realizations, we revisit the results of behavioral inclusion, lift, and equivalence from Faulwasser et al., 2023. Specifically, we demonstrate the equivalence between random variable behaviors and PCE coefficient behaviors while showing how realization and moment representations lead to information loss.

Section 4.3 then continues to the non-parametric representations via the fundamental lemma tailored to stochastic systems. The crucial observation is that, through the behavioral lift, trajectories of random variables, PCE coefficients, and moment pairs can

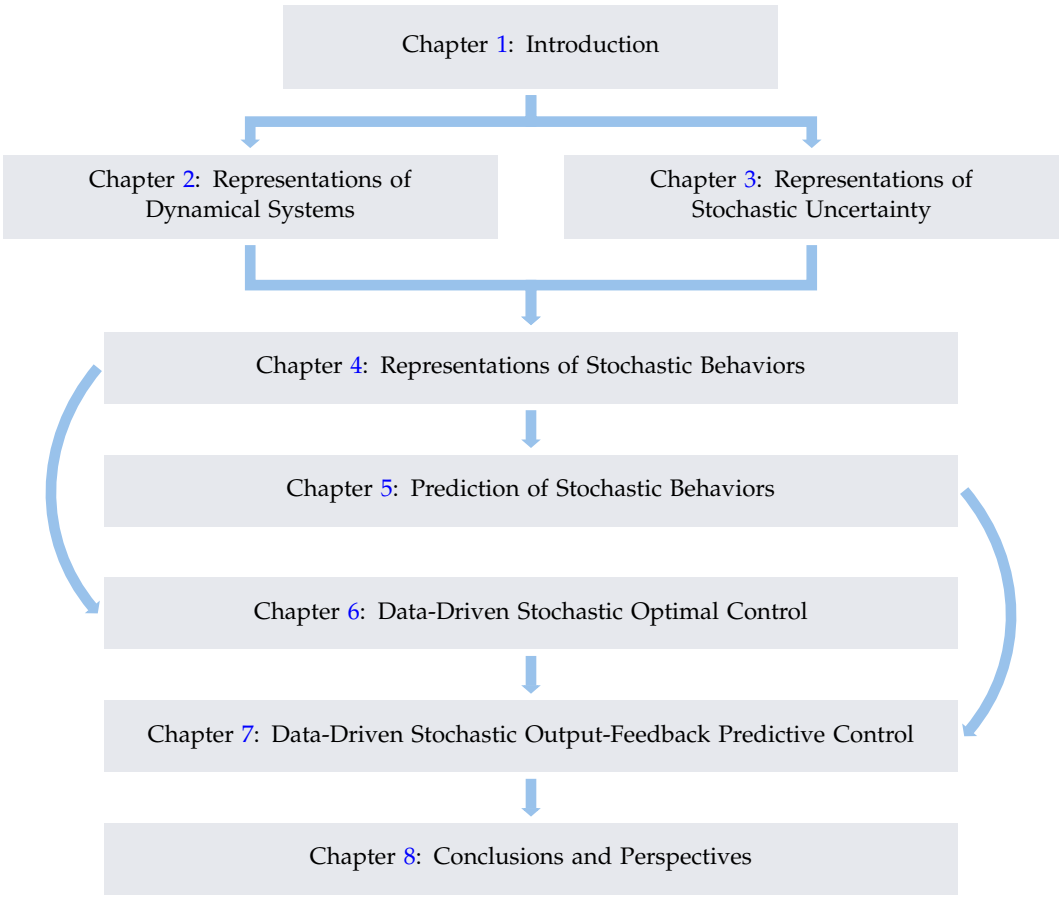

Figure 1.3.: Outline of the thesis.

all be represented using the same Hankel matrices constructed from measurements of realization data but with different column selectors to alternate between behaviors.

Chapter 5, which has not yet been published, focuses on addressing the prediction challenges posed by unstructured disturbances. We investigate the stochastic behavior of dynamical systems affected by such disturbances and, following the works by Willems, 2012 and Baggio et al., 2017, consolidate these disturbances into a residual disturbance. Furthermore, we propose a least-squares method to estimate the realizations of the residual disturbance and partially model its statistics using its empirical distribution. In addition, we introduce a regularized stochastic predictor to mitigate the effects of estimation errors.

The proposed stochastic prediction approach is validated using data from a real building application (Reinhardt et al., 2025). To quantify the accuracy of the stochastic predictions, we present PCE reformulations for constructing confidence intervals based on higher-order Chebyshev inequalities. This approach exploits the PCE representation to efficiently compute higher-order moments.

In Chapter 6, we focus on optimal control of stochastic discrete-time LTI systems with structured (measured and statistically modelled) disturbances, building on the results from Chapter 4 and extending the findings by Pan et al., 2023c. In this chapter, we first introduce the stochastic Optimal Control Problem (OCP) in a model-based framework, emphasizing the role of control policies and chance constraints, which are essential elements for *stochastic* optimal control. Furthermore, we reformulate the OCP using a data-driven approach and establish the conditions under which the model-based and data-driven OCPs are equivalent. Additionally, we present sufficient conditions under which the affine disturbance-feedback policy is proven to be optimal.

Section 6.3 extends the analysis to consider disturbances with uncertain distributions, based on the results by Pan et al., 2023a. Herein, we explore the relations between ambiguity sets induced by the distances of PCE coefficients, of L^2 random variables, of moment pairs (Gelbrich metric), of and distributions (Wasserstein metric). Based on the results by Nguyen et al., 2021, we establish the conditions under which these ambiguity sets are equivalent. Furthermore, we propose a tractable and exact second order cone reformulation of distributionally robust chance constraints using PCE coefficients. Finally, in Section 6.4, we demonstrate through numerical simulations of a two-mass spring-damper system that the proposed distributionally robust data-driven optimal control problem can be efficiently solved.

In Chapter 7, we transition from the open-loop implementation of optimized control policies from a single OCP to the closed-loop implementation of each first optimized policy from a series of OCPs. Our main contribution is the development and closed-loop analysis of a data-driven stochastic output-feedback predictive control scheme, based on the findings by Pan et al., 2023b, 2025a,b. Our approach does not require past measurements of exogenous disturbances. Instead, building on the results from Chapter 5, we estimate the residual disturbance that condenses the unstructured disturbances. Furthermore, the method requires only limited statistical information about the residual disturbance, specifically its first two moments.

In Section 7.1, we first present the conceptual predictive control scheme in a model-based form and then apply the results from earlier chapters to transform it into a tractable data-driven formulation. Additionally, we propose a data-driven output-feedback LQR approach to design the terminal ingredients of the optimal control prob-

lem. In Section 7.2, we address the open problem of closed-loop analysis for data-driven stochastic output-feedback predictive control. Specifically, we treat the initial conditions of the OCP, solved in each closed-loop step, to be stochastic such that its predicted counterpart from the last step serves as a feasibility-ensuring backup. Moreover, we establish average asymptotic performance bounds for the closed-loop system and analyze its robustness against errors resulting from the estimation of past disturbances. Crucially, the proposed scheme only requires information about the first two statistical moments of the disturbance rather than the full knowledge of its distribution. Finally, in Section 7.3, we validate the proposed approach's stability and robustness through numerical simulations of an aircraft system. This thesis ends with a brief summary and an outlook of open issues in Chapter 8.

2. Representations of Dynamical Systems

We begin by recalling some basic concepts of behavioral systems theory before briefly revisiting the fundamental lemma. We refer to Willems, 1986, 1991 for the original exposition of LTI systems in the behavioral setting, to Willems et al., 2005 for the fundamental lemma, and to Markovsky et al., 2021 for a recent survey on behavioral systems theory and its use for data-driven control of deterministic systems.

In the behavioral setting, a dynamical system Σ is specified by the tuple $(\mathbb{T}, \mathbb{S}, \mathfrak{B})$, where $\mathbb{T} \subseteq \mathbb{R}$ is the time axis, $\mathbb{S}$ is the signal space, and $\mathfrak{B}$ is a subset of $\mathbb{S}^{\mathbb{T}}$ called the *behavior* of this system. The behavior $\mathfrak{B}$ consists of those trajectories $s : \mathbb{T} \to \mathbb{S}$, which are compatible with the laws governing the dynamical system. A general continuous-time system, for instance, can be represented as the tuple $(\mathbb{T} = \mathbb{R}, \mathbb{S} = \mathbb{R}^{n_s}, \mathfrak{B})$, where the behavior is defined as:

$$\mathfrak{B} \doteq \left\{ s : \mathbb{R} \to \mathbb{R}^{n_s} \;\middle|\; s \text{ subject to } f\left(\frac{\mathrm{d}^\ell s}{\mathrm{d}t^\ell}, \cdots, \frac{\mathrm{d}s}{\mathrm{d}t}, s\right) = 0 \right\}. \tag{2.1}$$

Here, the system dynamics are governed by *differential equations*. Similarly, a general discrete-time system can be described by the tuple $(\mathbb{T} = \mathbb{Z}, \mathbb{S} = \mathbb{R}^{n_s}, \mathfrak{B})$ with

$$\mathfrak{B} \doteq \left\{ s : \mathbb{Z} \to \mathbb{R}^{n_s} \;\middle|\; s \text{ subject to } f\left(\boldsymbol{\sigma}^\ell s, \cdots, \boldsymbol{\sigma} s, s\right) = 0 \right\}, \tag{2.2}$$

where $\boldsymbol{\sigma}$ represents the forward-time shift, defined as $(\boldsymbol{\sigma} s)(k) \doteq s(k+1)$. In the discrete-time case, the system dynamics are typically expressed using *difference equations*.

In this thesis, we deal with discrete-time systems, i.e., $\mathbb{T} = \mathbb{Z}$, and with an n_s-dimensional signal, i.e., $\mathbb{S} = \mathbb{R}^{n_s}$. Specifically, to account for past signal values with $t \leq 0$, we consider the time axis $\mathbb{Z}$ rather than only the non-negative integers $\mathbb{N}$. For a detailed technical comparison of the systems with time axes $\mathbb{T} = \mathbb{Z}$ and $\mathbb{T} = \mathbb{N}$, see p. 567 of Willems, 1986.

Central to this representation of dynamics is the notion of *system behavior*. Indeed, the other three components—the time axis, the signal space, and the dynamic equations (models)—are inherently captured within the characterization of the behavior; see (2.1) and (2.2). Unlike the conventional model-centric approach, the behavioral perspective shifts the focus to the set of all trajectories that satisfy the system dynamics. In essence, while there may theoretically be infinitely many models, they all describe the same underlying behavior. This shift in focus has directly inspired advancements in data-driven representations as we will demonstrate in Section 2.3.

Furthermore, this shift can be refined as a transition from focusing on state-space models, which provide a latent description of system dynamics, to focusing on the system behavior itself, which is manifest. Specifically, we consider s—the variables that

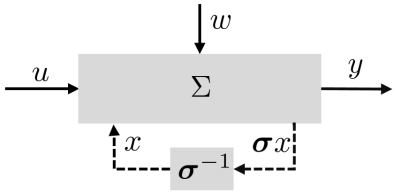

Figure 2.1.: Manifest and latent variables of the dynamical system Σ, where the input u, the disturbance w, and the output y constitute manifest variables, and the state x is a latent variable

the model aims to describe—to be *manifest* that can be directly measured externally. As shown in Figure 2.1, the manifest variables s consist of the input u, the output y, and the disturbance w. [1] *Latent* variables, in contrast, are auxiliary variables used in mathematical models to simplify representations or capture underlying dynamics. The term "latent" emphasizes that these variables are not directly observable externally or only indirectly observable through the manifest variables. In this context, the state x is considered to be latent as in Figure 2.1.

Introducing the state x allows us, for instance, reformulate $\mathfrak{B}$ from (2.2) into a state-space representation

$$\mathfrak{B} \doteq \{s : \mathbb{Z} \to \mathbb{R}^{n_s} \mid \exists x : \mathbb{Z} \to \mathbb{R}^{n_x},\ (s, x) \text{ subject to } f_{\text{latent}}(\boldsymbol{\sigma} x, x, s) = 0\},$$

which employs a first-order difference equation rather than an order-ℓ equation as in (2.2).

Denoting $\mathfrak{B}$ as the *manifest behavior* of the system Σ, we define the *latent behavior* $\mathfrak{B}^{\mathrm{L}}$ with respect to a chosen set of states x as

$$\mathfrak{B}^{\mathrm{L}} \doteq \{(s, x) : \mathbb{Z} \to \mathbb{R}^{n_s} \times \mathbb{R}^{n_x} \mid (s, x) \text{ subject to } f_{\text{latent}}(\boldsymbol{\sigma} x, x, s) = 0\}.$$

As we can see, specifying which variables are considered states inherently requires conceptualizing a model of the system. This approach contrasts with the data-driven methodology, which focuses on leveraging measured data from the manifest variables of the system without explicitly modeling latent variables. Therefore, in the following, we explore both parametric and non-parametric representations of the manifest behavior $\mathfrak{B}$, with the latter closely aligning with the data-driven methodology.

2.1. Linear Time-Invariant Systems and their Behaviors

This thesis focuses on Linear Time-Invariant (LTI) systems. Given $t_1,\ t_2 \in \mathbb{Z}$, $t_1 \leq t_2$, we define the restriction of a signal $s : \mathbb{Z} \to \mathbb{R}^{n_s}$ to the interval $[t_1, t_2]$ as

$$s_{[t_1,t_2]} \doteq \left[s^\top(t_1), \cdots, s^\top(t_2)\right]^\top.$$

Similarly, the restriction of the behavior $\mathfrak{B}$ to $[t_1, t_2]$ is defined as

$$\mathfrak{B}_{[t_1,t_2]} \doteq \{s_{[t_1,t_2]} \mid s \in \mathfrak{B}\}.$$

Definition 2.1 (LTI behaviors, cf. Section 6 of Willems, 1986):
A behavior $\mathfrak{B}$ is linear time-invariant if the following conditions hold.

i) The behavior $\mathfrak{B}$ is a linear subspace of $(\mathbb{R}^{n_s})^{\mathbb{Z}}$.

ii) The behavior $\mathfrak{B}$ is *shift-invariant* if $\boldsymbol{\sigma}\mathfrak{B} = \mathfrak{B}$, meaning $\mathfrak{B}$ remains unchanged under time shifts, i.e., $\boldsymbol{\sigma}\mathfrak{B} = \{\boldsymbol{\sigma} f \mid f \in \mathfrak{B}\}$.

iii) The behavior $\mathfrak{B}$ is *complete* if and only if $\mathfrak{B}$ is closed in the topology of pointwise convergence. That is, if $s_{[t_1,t_2]} \in \mathfrak{B}_{[t_1,t_2]}$ holds for any $t_1, t_2 \in \mathbb{Z}$, $t_1 \leq t_2$, then $s \in \mathfrak{B}$.

[1] We consider the cases where the disturbance w is unmeasured in Chapter 5.

From now on, unless otherwise specified, $\mathfrak{B}$ will refer specifically to an LTI behavior, rather than to a more general behavior. Similarly, $(\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{B})$ represents an LTI system.

The concept of shift invariance implies that restrictions of the same LTI behavior over equally long time-intervals coincide. Specifically, for any intervals $[t_1, t_2]$ and $[t_1', t_2']$ with $t_2 - t_1 = t_2' - t_1' = T \in \mathbb{N}$, it follows that $\mathfrak{B}_{[t_1,t_2]} = \mathfrak{B}_{[t_1',t_2']}$. To simplify the notation, we introduce the finite-length behavior $\mathfrak{B}_T$ with $T \in \mathbb{N}^+$ as

$$\mathfrak{B}_T \doteq \left\{ s_{[t_1,t_1+T-1]} \,|\, s \in \mathfrak{B}, t_1 \in \mathbb{Z} \right\}, \tag{2.3}$$

which denotes the set of all trajectories of length T within $\mathfrak{B}$. Due to shift invariance, $\mathfrak{B}_T$ is equivalent to $\mathfrak{B}_{[t_1,t_2]}$ for any interval of length T, i.e,

$$\mathfrak{B}_T = \mathfrak{B}_{[t_1,t_2]}, \quad \forall t_1, t_2 \in \mathbb{Z} \text{ that } t_2 - t_1 = T.$$

Shift invariance simplifies the proof of completeness for LTI behaviors. Specifically, due to shift invariance, property iii) of Definition 2.1 can be restated as follows:
iii) If every finite-length segment $s_{[0,T-1]}$ belongs to $\mathfrak{B}_T$ for all $T \in \mathbb{N}^+$, and this guarantees that the entire trajectory s is in $\mathfrak{B}$, then the behavior $\mathfrak{B}$ is complete.

In the behavioral setting, controllability of a dynamical system is defined as follows.

Definition 2.2 (Controllability): A behavior $\mathfrak{B}$ is controllable if for any two trajectories $s_1, s_2 \in \mathfrak{B}$, there exist $t_1 \geq 0$ and $s_{[0,t_1-1]} \in \mathfrak{B}_{t_1}$ such that $s' \in \mathfrak{B}$ with s' defined by

$$s'(t) = \begin{cases} s_1(t), & t < 0 \\ s(t) & 0 \leq t \leq t_1 - 1 \\ s_2(t) & t \geq t_1 \end{cases}.$$

In other words, there exists a system trajectory that connects any past trajectory to any future trajectory in finite time. We note that controllability in the behavioral setting is related but not equivalent to the classical notion of state controllability. We will elaborate this in Section 2.2.3.

2.2. Parametric System Representations

After briefly introducing LTI behaviors, we delve into their *parametric* representations. Specifically, we consider kernel representations, input/output (i/o) representations using AutoRegressive with eXtra input (ARX) models, and input/state/output representations using state-space models. Rigorous definitions of these representations are given by Willems, 1986, 1991.

2.2.1. Kernel Representations

A behavior $\mathfrak{B}$ is linear time-invariant if and only if all $s \in \mathfrak{B}$ satisfy the following autoregressive equation (Willems, 1986, Theorem 5)

$$R_\ell s(k + \ell_1) + R_{\ell-1} s(k + \ell_1 - 1) + \cdots + R_0 s(k + \ell_2) = 0, \quad \forall k \in \mathbb{Z}, \tag{2.4}$$

where R_ℓ, $R_{\ell-1}$, $\cdots$, $R_0 \in \mathbb{R}^{n_R \times n_s}$, n_R is the number of rows in R, and $\ell_1, \ell_2 \in \mathbb{Z}$ with $\ell = \ell_1 - \ell_2 > 0$. With $R_\ell \neq 0$ and $R_0 \neq 0$, the time shift ℓ is referred as the *lag* of this

specific representation. Due to the shift invariance, ℓ_1 and ℓ_2 can be arbitrarily chosen from $\mathbb{Z}$, and hence can be negative.

Since (2.4) holds for all $k \in \mathbb{Z}$, it can be rewritten in terms of the entire signal s using the shift operator $\boldsymbol{\sigma}$, i.e,

$$R_\ell \boldsymbol{\sigma}^{\ell_1} s + R_{\ell-1} \boldsymbol{\sigma}^{\ell_1-1} s + \cdots + R_0 \boldsymbol{\sigma}^{\ell_2} s = 0.$$

For a more compact representation, we define a polynomial matrix in the shift operator

$$R(\boldsymbol{\sigma}) \doteq R_\ell \boldsymbol{\sigma}^{\ell_1} + R_{\ell-1} \boldsymbol{\sigma}^{\ell_1-1} + \cdots + R_0 \boldsymbol{\sigma}^{\ell_2}, \tag{2.5}$$

with $R_\ell \neq 0$ and $R_0 \neq 0$. This polynomial matrix $R(\boldsymbol{\sigma})$ compresses the coefficient matrices $\{R_i\}_{i=\ell_2}^{\ell_1}$ by using the shift operator $\boldsymbol{\sigma}$ as its indeterminate.

Based on that, a linear time-invariant $\mathfrak{B}$ can be expressed through the following *kernel representation*:

$$\mathfrak{B}(R) \doteq \{s : \mathbb{Z} \to \mathbb{R}^{n_s} \mid R(\boldsymbol{\sigma}) s = 0\}, \quad \mathfrak{B}(R) = \ker R(\boldsymbol{\sigma}). \tag{2.6}$$

This representation highlights that $\mathfrak{B}(R)$ is a linear subspace of $(\mathbb{R}^{n_s})^{\mathbb{Z}}$, specifically the kernel of $R(\boldsymbol{\sigma})$. The term "kernel" implies the linearity required for LTI systems while shift invariance and completeness are inherently encoded in the structure of (2.6); see Theorem 5 of Willems, 1986. [2]

Polynomial Matrices in the Z-Domain

The Z-transform provides a frequency-domain representation of discrete-time signals and is particularly useful for analyzing LTI systems. By leveraging the Z-transform, time-domain operations involving the shift operator $\boldsymbol{\sigma}$ can be converted into simple algebraic manipulations in the Z-domain, simplifying the analysis of LTI systems.

Definition 2.3 (Bilateral Z-transform): The bilateral Z-transform of a discrete-time signal $s : \mathbb{Z} \to \mathbb{R}^{n_s}$ is defined as

$$\mathcal{Z}[s] \doteq \sum_{k=-\infty}^{\infty} s(k) z^{-k}, \quad z \in \mathbb{C}.$$

For the forward-shifted signal $\boldsymbol{\sigma} s$ and the backward-shifted signal $\boldsymbol{\sigma}^{-1} s$, their Z-transforms satisfy

$$\mathcal{Z}[\boldsymbol{\sigma} s] = z\mathcal{Z}[s], \quad \mathcal{Z}[\boldsymbol{\sigma}^{-1} s] = z^{-1}\mathcal{Z}[s].$$

These relations establish a direct connection between the shift operator $\boldsymbol{\sigma}$ in the time domain and the algebraic variable z in the Z-domain. In other words, by substituting $\boldsymbol{\sigma}^i$ with z^i, $i \in \mathbb{Z}$, any polynomial matrix in $\boldsymbol{\sigma}$ naturally can be transformed into a polynomial matrix in z.

To this end, we first revisit the definitions of polynomial and rational matrices in the Z-domain.

[2] We note that the term "kernel" is also used in the context of reproducing kernel Hilbert spaces in functional analysis. For discussions on the connection between reproducing kernels and the fundamental lemma, see Huang et al., 2024; Molodchyk et al., 2024.

Definition 2.4 (Polynomial matrices in the Z-domain): The set of all polynomials with real coefficients that include positive powers, negative powers, or both powers of z is denoted as $\mathbb{R}[z, z^{-1}]$.

A polynomial matrix $R(z)$, which contains polynomials as its elements, is denoted as

$$R(z) \in \mathbb{R}^{n_R \times n_s}\left[z, z^{-1}\right], \quad R : \mathbb{C} \to \mathbb{C}^{n_R \times n_s},$$

where n_R denotes the number of rows and n_s as the number of columns, and $R(z)$ can be viewed as a function mapping a scalar complex number to a complex matrix.

Rational matrices in the Z-domain will be used to represent transfer functions later.

Definition 2.5 (Rational matrices in the Z-domain):

- Let $\mathbb{R}(z)$ denote all scalar rational functions, a rational function $r(z) \in \mathbb{R}(z)$ is defined as
$$r(z) = \frac{\sum_{i=0}^{\ell_n} b_i z^i}{\sum_{i=0}^{\ell_d} a_i z^i}, \quad \{a_i\}_{i=0}^{l_d} \in \mathbb{R}^{l_d}, \quad \{b_i\}_{i=0}^{l_n} \in \mathbb{R}^{l_n}.$$
with $\ell_n \geq 0$, $\ell_d \geq 0$ as the degrees of the nominator and the denominator, respectively. A rational function $r(z)$ can be viewed as a map from $\mathbb{C}$ to $\mathbb{C}$ defined for all $z \in \mathbb{C}$ except the zeros of the denominator.
- A rational function is (strictly) proper if the degree of numerator is (smaller or) no larger than the degree of denominator. Equivalently, $r(z)$ is proper if and only if $\lim_{z\to\infty} r(z)$ is finite, and is strictly proper if and only if $\lim_{z\to\infty} r(z) = 0$.
- A rational matrix $G(z) \in \mathbb{R}^{n_y \times n_u}(z)$, typically used as the transfer function of a MIMO system, is a matrix whose elements are rational functions. It is proper if and only if $\lim_{z\to\infty} G(z) \leq \infty$, and is strictly proper if only if $\lim_{z\to\infty} G(z) = 0$.

By replacing each occurrence of $\boldsymbol{\sigma}^i$ with z^i in the polynomial matrix $R(\boldsymbol{\sigma})$ from (2.5), we derive its corresponding polynomial matrix $R(z)$ in the Z-domain

$$R(z) = R_\ell z^{\ell_1} + R_{\ell-1} z^{\ell_1 - 1} + \cdots + R_0 z^{\ell_2} \in \mathbb{R}^{n_R \times n_s}\left[z, z^{-1}\right]. \tag{2.7}$$

Similarity Transforms and Minimal Representations

Kernel representations of an LTI behavior are not unique and often contain redundancy, particularly in terms of the row number and the length of time span. A kernel representation is considered *minimal* if the row number n_R of $R(z)$ is as small as possible among all kernel representations (Willems, 1991, Definition III.3). We denote this minimal row number by $n_R^\star$. On the other hand, representations with the smallest possible *lag* are denoted as *minimal-lag* representations. The minimum lag is denoted as $\ell^\star$. Put simply, as shown in (2.4), minimal kernel representations use the fewest number of equalities to describe LTI dynamics, while *minimal-lag* representations achieve the shortest shift span.

The rank of the polynomial matrix $R(z)$ provides insights into the minimality and controllability of the system.

Lemma 2.1 (Minimality and controllability test (Willems, 1991)):
Given a kernel representation $\mathfrak{B}(R)$.

i) The representation is minimal if and only if $R(z)$ is of full row rank for all $z \in \mathbb{C}$ except a finite number of points (Willems, 1991, Proposition III.3).

ii) It represents a controllable behavior if and only if $R(z)$ is of constant rank for all $z \neq 0$ in the complex space $\mathbb{C}$ (Willems, 1991, Theorem V.2).

Obviously, if $R(z)$ is of full row rank for all $z \in \mathbb{C}$ with $z \neq 0$, then, $\mathfrak{B}(R)$ is a minimal kernel representation of a controllable system.

To systematically eliminate redundancy, we introduce similarity transforms of kernel representations. The key tool of similarity transforms is the left multiplication of square polynomial matrices, which can be further classified based on their determinants.

Definition 2.6: A square polynomial matrix $P(z) \in \mathbb{R}^{n_R \times n_R}[z, z^{-1}]$

- is non-singular if $\det(P(z)) \neq 0$ for all $z \in \mathbb{C}$ except a finite number of points where $\det(P(z)) = 0$.
- is unimodular if $\det(P(z)) = \alpha z^d$. Then, $\det(P(z)) \neq 0$ for all $z \in \mathbb{C}$ except the origin $z = 0$.

Note that while the inverse of a non-singular polynomial matrix is generally a rational matrix, the inverse of a unimodular matrix remains polynomial and unimodular, making it fundamental to similarity transformations.

Lemma 2.2 (Similarity transforms (Willems, 1986)):
Given a polynomial matrix $R(z) \in \mathbb{R}^{n_R \times n_s}[z, z^{-1}]$

i) for any polynomial matrix $R'(z) \in \mathbb{R}^{n_{R'} \times n_R}[z, z^{-1}]$, we have $\mathfrak{B}(R) \subseteq \mathfrak{B}(R'R)$,

ii) for any unimodular $U(z) \in \mathbb{R}^{n_R \times n_R}$, we have $\mathfrak{B}(R) = \mathfrak{B}(UR)$.

Proof. For any $s \in \mathfrak{B}(R)$, we have $R(z)s = 0$, which implies $R'(z)R(z)s = 0$ for any polynomial matrix $R'(z) \in \mathbb{R}^{n_{R'} \times n_R}$. Hence, we conclude that $\mathfrak{B}(R) \subseteq \mathfrak{B}(R'R)$.

Since the inverse $U(z)$ of a unimodular matrix $U(z)$ is also a polynomial matrix, we have $\mathfrak{B}(R) \subseteq \mathfrak{B}(UR) \subseteq \mathfrak{B}(U^{-1}UR) = \mathfrak{B}(R)$. Therefore, we conclude that $\mathfrak{B}(R) = \mathfrak{B}(UR)$. □

Moreover, a unimodular matrix can be further factorized as a product of elementary matrices that corresponds to the simple row operations: switching, scaling with αz^i, or adding a multiple of one row to another, see p. 524 of Antsaklis et al., 2006. The following matrix, for instance, is unimodular,

$$U(z) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & z & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & z \\ 1 & 1 & 1 \end{bmatrix}, \quad \det(U(z)) = -z.$$

Note that the row elementary operations of switching, multiplication, and addition does not alter the kernel space. Moreover, time shifts induced by row multiplications

with z or z^{-1} are permissible under shift invariance. As a result, the transformed behavior induced by left multiplication of a unimodular matrix remains equivalent to the original one.

Using similarity transforms, we can systematically reduce the *row number* and then the *lag* of a kernel representation to their minimal values.

Lemma 2.3 (Minimal representations via row and lag reductions cf. step (3) on p. 579 of Willems, 1986):
Consider a polynomial matrix $R(z) \in \mathbb{R}^{n_R \times n_s}[z, z^{-1}]$.

i) Row reduction: there exists a unimodular $U_1(z) \in \mathbb{R}^{n_R \times n_R}[z, z^{-1}]$ such that

$$U_1(z)R(z) = \begin{bmatrix} R_1(z) \\ 0 \end{bmatrix},$$

where $R_1(z)$ is of full row rank for all $z \in \mathbb{C}$ except a finite number of points. Moreover, the row number $n_R^\star$ of $R_1(z)$ is minimal.

ii) Lag reduction (left): there exists a unimodular $U_2(z) \in \mathbb{R}^{n_R^\star \times n_R^\star}[z, z^{-1}]$ such that

$$R_2(z) = U_2(z)R_1(z) = R_{2,\ell}z^{\ell} + R_{2,\ell-1}z^{\ell-1} + \cdots + R_{2,0}$$

with $R_{2,\ell}$ of full row rank and $R_{2,0} \neq 0$.

iii) Lag reduction (right): there exists a unimodular $U_3(z) \in \mathbb{R}^{n_R^\star \times n_R^\star}[z, z^{-1}]$ such that

$$R_3(z) = U_3(z)R_2(z) = R_{3,\ell^\star}z^{\ell^\star} + R_{3,\ell^\star-1}z^{\ell^\star-1} + \cdots + R_{3,0}$$

with $R_{3,0}$ of full row rank and $R_{3,\ell^\star} \neq 0$. Here, $R_3(z)$ has the minimal lag $\ell^\star$.

During the above similarity transforms, the behavior remains invariant, i.e.,

$$\mathfrak{B}(R) = \mathfrak{B}(R_1) = \mathfrak{B}(R_2) = \mathfrak{B}(R_3).$$

The first step removes redundant rows from the kernel representation, ensuring minimality while preserving the behavior. The second and third steps sequentially reduce time shifts from the left and right, ultimately yielding a minimal-lag representation. To illustrate this process, we present the following example.

Example 2.1: Consider

$$R(z) = \begin{bmatrix} 4 & 3 & 2 \\ 0 & 0 & 0 \\ 5 & 4 & 3 \end{bmatrix} z^2 + \begin{bmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix} z^1 + \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 0 \\ 1 & 2 & 2 \end{bmatrix} + \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix} z^{-1}.$$

We have the following results by applying the steps of Lemma 2.3

$$U_1(z) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & z & -1 \end{bmatrix}, \quad R_1(z) = \begin{bmatrix} 4 & 3 & 2 \\ 0 & 0 & 0 \end{bmatrix} z^2 + \begin{bmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix} z^1 + \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} z^{-1},$$

$$U_2(z) = \begin{bmatrix} z & 0 \\ 0 & z^2 \end{bmatrix}, \quad R_2(z) = \begin{bmatrix} 4 & 3 & 2 \\ 1 & 1 & 1 \end{bmatrix} z^3 + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} z^2 + \begin{bmatrix} 0 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} z + \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \end{bmatrix},$$

$$U_3(z) = \begin{bmatrix} z^{-1} & -z^{-2} \\ 0 & z^{-1} \end{bmatrix}, \quad R_3(z) = \begin{bmatrix} 4 & 3 & 2 \\ 1 & 1 & 1 \end{bmatrix} z^2 + \begin{bmatrix} -1 & -1 & -1 \\ 0 & 0 & 0 \end{bmatrix} z^1 + \begin{bmatrix} 0 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}.$$

Note that $R_1(z)$ achieves the minimal row number $n_R^\star = 2$ while $R_3(z)$ further possesses the minimal lag $\ell^\star = 2$.

2.2.2. Input and Output Representations

Beyond the LTI information encoded in the kernel representation, in the context of control, we require a more structured system representation considering inputs and outputs. This necessitates separating the signal s into inputs and outputs. In the behavioral framework, such a separation always exists for LTI dynamical systems, given a proper definition of the inputs (Willems, 1991).

Consider an LTI behavior $\mathfrak{B}$, a permutation matrix $\Pi \in \mathbb{R}^{n_s \times n_s}$, and an integer $1 \leq n_{\tilde{u}} \leq n_s$, we define $(\tilde{u}, y) = \Pi^{-1} s$ as a partitioning of the variables $s : \mathbb{Z} \to \mathbb{R}^{n_s}$ into input variables $\tilde{u} : \mathbb{Z} \to \mathbb{R}^{n_{\tilde{u}}}$ and output variables $y : \mathbb{Z} \to \mathbb{R}^{n_y}$ with $n_{\tilde{u}} + n_y = n_s$. Here, the input $\tilde{u}$ is distinct from the manipulatable input u and the disturbance w, as will be further elaborated in Chapter 4. Specifically, $\tilde{u} = (u, w)$ contains all exogenous inputs to the system.

Let $\Pi_{\tilde{u}}$ be the projection of s onto $\tilde{u}$, i.e.,

$$\Pi_{\tilde{u}} : s = \Pi\,(\tilde{u}, y) \mapsto \tilde{u}, \quad \Pi_{\tilde{u}}\mathfrak{B} \doteq \{\tilde{u} : \mathbb{Z} \to \mathbb{R}^{n_{\tilde{u}}} \mid \tilde{u} = \Pi_{\tilde{u}} s,\ s \in \mathfrak{B}\}. \tag{2.8}$$

Similarly, we define Π_y as the projection onto the output component. We define an i/o partitioning of $\mathfrak{B}$ as follows.

Definition 2.7 (i/o partitioning (Willems, 1991)): The partitioning $(\tilde{u}, y) = \Pi^{-1} s$ is an input/output partitioning of $\mathfrak{B}$ if

i) the input $\tilde{u}$ is a free variable, i.e., $\Pi_{\tilde{u}}\mathfrak{B} = (\mathbb{R}^{n_{\tilde{u}}})^{\mathbb{Z}}$.

ii) the output y does *not anticipate* the input $\tilde{u}$, i.e., given any $(\tilde{u}_1, y_1) \in \Pi^{-1}\mathfrak{B}$, then for any $\tilde{u}_2 \in (\mathbb{R}^{n_{\tilde{u}}})^{\mathbb{Z}}$ such that $\tilde{u}_1(t) = \tilde{u}_2(t)$ for all $t \leq 0$, there exists $y_2 \in (\mathbb{R}^{n_y})^{\mathbb{Z}}$ such that $y_1(t) = y_2(t)$ for all $t \leq 0$ and $(\tilde{u}_2, y_2) \in \Pi^{-1}\mathfrak{B}$.

iii) the output y *processes* the input $\tilde{u}$, i.e., given $(\tilde{u}, y_1), (\tilde{u}, y_2) \in \Pi^{-1}\mathfrak{B}$ with the same input signal, if $y_1(t) = y_2(t)$ for all $t < 0$, then $y_1(t) = y_2(t)$ for $t \geq 0$ as well.

iv) the number of inputs $n_{\tilde{u}}$ is maximal over all partitionings of $\mathfrak{B}$ that satisfy properties i)–iii).

We note that these conditions imply the following:

i) The input is imposed externally and cannot be influenced by the system dynamics.

ii) The past values of y do not contain any information about the future values of $\tilde{u}$.

iii) Once the input signal $\tilde{u}$ is given and the initial conditions (e.g. the past of y) are set, the future values of y are uniquely defined.

In the later discussion on state-space representations, these initial conditions can be further simplified to the initial state variable $x(0)$. In other words, the future of y can be uniquely determined by the future of $\tilde{u}$ and the initial condition $x(0)$.

After separating s into inputs and outputs, we are able to formulate LTI systems in i/o representations. Consider an i/o partitioning $s = \Pi\,(\tilde{u}, y)$ of the LTI system $\mathfrak{B}$.

Then, there exist $P_\ell, \cdots, P_0 \in \mathbb{R}^{n_y \times n_y}$, and $Q_\ell, \cdots, Q_0 \in \mathbb{R}^{n_y \times n_{\tilde{u}}}$ such that all trajectories $s \in \mathfrak{B}$ can be represented as

$$\sum_{i=\ell_2}^{\ell_1} P_{i-\ell_2} y(k+i) = \sum_{i=\ell_2}^{\ell_1} Q_{i-\ell_2} \tilde{u}(k+i), \quad \forall k \in \mathbb{Z}. \tag{2.9}$$

Following the idea of consolidating parameter matrices into polynomial matrices, we introduce $P(z) \in \mathbb{R}^{n_y \times n_y}[z, z^{-1}]$ and $Q(z) \in \mathbb{R}^{n_y \times n_{\tilde{u}}}[z, z^{-1}]$ as

$$P(z) \doteq \sum_{i=\ell_2}^{\ell_1} P_{i-\ell_2} z^i, \quad Q(z) \doteq \sum_{i=\ell_2}^{\ell_1} Q_{i-\ell_2} z^i.$$

Replacing z with $\boldsymbol{\sigma}$, the i/o representations of $\mathfrak{B}$ are defined as

$$\mathfrak{B}(P, Q) \doteq \{(\tilde{u}, y) \,|\, P(\boldsymbol{\sigma}) y = Q(\boldsymbol{\sigma}) \tilde{u}\}. \tag{2.10}$$

The following results illustrate the equivalence between the LTI behavior $\mathfrak{B}$ and its i/o representation $\mathfrak{B}(P, Q)$, with additional conditions on polynomial matrices P and Q to ensure the properties of an i/o partitioning.

Lemma 2.4 (LTI behavior via $\mathfrak{B}(P, Q)$): A behavior $\mathfrak{B}$ is LTI if and only if there exist a permutation matrix $\Pi \in \mathbb{R}^{n_s \times n_s}$, a polynomial matrix $P(z) \in \mathbb{R}^{n_y \times n_y}[z, z^{-1}]$, and $Q(z) \in \mathbb{R}^{n_y \times n_{\tilde{u}}}[z, z^{-1}]$ such that P is a non-singular polynomial matrix and $P^{-1}Q$ is a proper rational matrix. Then,

i) $s = \Pi(\tilde{u}, y)$ is an i/o partitioning of $\mathfrak{B}$ satisfying Definition 2.7.

ii) $\mathfrak{B}$ admits the i/o representation $\mathfrak{B}(P, Q)$ under permutation Π, i.e., $\mathfrak{B} = \Pi\mathfrak{B}(P, Q)$.

The previous result stems from Theorem 2 of Willems, 1986 and Proposition VIII.6 of of Willems, 1991.

Here, the nonsingularity of $P(z)$, cf. Definition 2.6, ensures that the input is free and that the output processes the input. In addition, the rational matrix $P^{-1}Q$ explicitly represents the transfer function from $\tilde{u}$ to y. By enforcing $P^{-1}Q$ to be a proper rational matrix, cf. Definition 2.5, we guarantee that the output does not anticipate the inputs.

Specifically, $\mathfrak{B}(P, Q)$ corresponds to a kernel representation $\mathfrak{B}(R)$ with

$$R(z) = \begin{bmatrix} Q(z) & -P(z) \end{bmatrix} \Pi^{-1}.$$

Moreover, nonsingularity of $P(z)$ further implies that there exists a $z \in \mathbb{C}$ such that $R(z) = [Q(z), -P(z)]\Pi^{-1}$ is of full row rank. Thus, the corresponding kernel representation $R(\boldsymbol{\sigma})s = 0$ is minimal; see Proposition III.3 of Willems, 1991. As a result, the minimal row number of a kernel representation equals the dimension of outputs n_y, i.e.,

$$n_R^\star = n_y.$$

2.2.3. State-Space Representations

One of the most widely used parametric descriptions for LTI systems is the state-space representation. In addition to the manifest variable $s = \Pi(\tilde{u}, y)$ that can be measured

externally from the system, state-space representations further introduce the latent state variable x, which captures the necessary information for predicting future system trajectory.

Suppose there exists a latent state signal $x : \mathbb{Z} \to \mathbb{R}^{n_x}$ such that for all manifest variables $s = \Pi(\tilde{u}, y) \in \mathfrak{B}$, the following holds

$$\boldsymbol{\sigma} x = Ax + B\tilde{u}, \tag{2.11a}$$
$$y = Cx + D\tilde{u}. \tag{2.11b}$$

See the discussion of Figure 2.1 for the distinction between *latent* and *manifest* variables. We denote the state-space representation of $\mathfrak{B}$ as

$$\mathfrak{B}(A, B, C, D) \doteq \{(\tilde{u}, y) \,|\, \exists x : \mathbb{Z} \to \mathbb{R}^{n_x},\ x, \tilde{u}, y \text{ satisfy (2.11)}\}, \tag{2.12}$$

and the *latent behavior* corresponding to a specific choice of x as

$$\mathfrak{B}^{\mathrm{L}}(A, B, C, D) \doteq \{(x, \tilde{u}, y) \,|\, x, \tilde{u}, y \text{ satisfy (2.11)}\}. \tag{2.13}$$

The inclusion of x allows the formulation of (2.11) with only a one-step time shift, unlike the longer lags required in kernel or input/output representations. Essentially, x acts as a memory component that stores and compresses past information, thereby serving as the "state" of the system, cf. the properties of state in Definition VII.1 of Willems, 1991. We have the following result.

Lemma 2.5 (Initial condition for a state-space representation): Consider a state-space representation $\mathfrak{B}(A, B, C, D)$ of an LTI behavior $\mathfrak{B}$. For any prediction horizon $N \in \mathbb{N}^+$ and an input trajectory $\tilde{u}_{[0,N-1]} \in \mathbb{R}^{Nn_{\tilde{u}}}$, there are unique trajectories $y_{[0,N-1]} \in \mathbb{R}^{Nn_y}$ and $x_{[0,N-1]} \in \mathbb{R}^{Nn_x}$ such that $(x, \tilde{u}, y)_{[0,N-1]} \in \mathfrak{B}^{\mathrm{L}}_N(A, B, C, D)$ if the initial condition $x(0) \in \mathbb{R}^{n_x}$ is specified.

Proof. The unique state and output trajectories are, for $k \in \mathbb{I}_{[0,N-1]}$,

$$x(k) = A^k x(0) + \sum_{j=0}^{k-1} A^{k-(j+1)} B u(j), \quad y(k) = CA^k x(0) + \sum_{j=0}^{k-1} CA^{k-(j+1)} B u(j) + Du(k).$$

□

The following lemma, combining the results of Theorem 1 and Theorem 3 of Willems, 1986, illustrates the equivalence of state-space representations and others.

Lemma 2.6 (Equivalence between $\mathfrak{B}$ and $\mathfrak{B}(A, B, C, D)$): A behavior $\mathfrak{B}$ is LTI if and only if there exist a permutation matrix $\Pi \in \mathbb{R}^{n_s \times n_s}$ and system matrices (A, B, C, D) such that $\mathfrak{B} = \Pi\mathfrak{B}(A, B, C, D)$.

State-space representations corresponding to an LTI behavior $\mathfrak{B}$ are not unique and generally depend on the choice of state variables x. A state-space representation is said to be *minimal* if the state dimension n_x is the smallest possible among all state-space representations of the same behavior $\mathfrak{B}$. In this thesis, we denote the minimal state dimension as $n^\star$.

In a minimal state-space representation, the pair (A, C) is required to be *state observable*, meaning that the observability matrix

$$\mathcal{O}_i(A, C) \doteq \begin{bmatrix} C^\top & (CA)^\top & \cdots & (CA^{i-1})^\top \end{bmatrix}^\top, \quad i \in \mathbb{N}^+. \tag{2.14}$$

has full rank at $i = n_x$, i.e., $\operatorname{rank}(\mathcal{O}_{n_x}(A, C)) = n_x$, cf. p. 566 of Willems1986. In such cases, all unobservable modes have been eliminated, ensuring that the internal state dimension n_x is truly minimal: $n_x = n^\star$.

However, *state controllability* of the pair (A, B) is not required for a minimal state-space representation. The pair (A, B) is defined to be *state controllable* if the reachability matrix $\mathcal{R}_i(A, B)$

$$\mathcal{R}_i(A, B) \doteq \begin{bmatrix} B & AB & \cdots & A^{i-1}B \end{bmatrix}, \; i \in \mathbb{N}^+ \tag{2.15}$$

satisfies $\operatorname{rank}(\mathcal{R}_{n_x}(A, B)) = n_x$.

This asymmetry becomes clearer when distinguishing *minimal state-space representations of behaviors* from *minimal state-space realizations of transfer functions*. As shown in Proposition VIII.8 of Willems, 1991, a manifest behavior uniquely determines a transfer function, but a transfer function corresponds to a family of behaviors that share the same controllable part. Consequently, a minimal representation of a behavior need only be state observable; it may include uncontrollable modes that do not affect the input-output map but are consistent with the behavior.

Moreover, while state controllability implies behavioral controllability, the converse does not hold. Specifically, if (A, B) is state controllable, then the underlying behavior $\mathfrak{B}$ is controllable in the sense of Definition 2.2; see Proposition VII.11 of Willems, 1991. However, for a controllable behavior $\mathfrak{B}$, it may admits state-space representations that are not state controllable.

While state controllability and observability are key properties for state-space representations, two weaker yet important notions are state detectability and state stabilizability.

Lemma 2.7 (Hautus tests for stabilizability and detectability (Hespanha, 2023)): A pair (A, B) is state stabilizable if and only if

$$\operatorname{rank}\left(\begin{bmatrix} \lambda I - A & B \end{bmatrix}\right) = n_x, \quad \forall \lambda \in \mathbb{C} \text{ that } |\lambda| \geq 1.$$

A pair (A, C) is state detectable if and only if

$$\operatorname{rank}\left(\begin{bmatrix} \lambda I - A \\ C \end{bmatrix}\right) = n_x, \quad \forall \lambda \in \mathbb{C} \text{ that } |\lambda| \geq 1.$$

Remark 2.1 (Minimal lag = Observability index): Consider a minimal state-space representation $\mathfrak{B}(A, B, C, D)$ with state-dimension $n_x = n^\star$ and a minimal-lag kernel representation $\mathfrak{B}(R)$ with lag $\ell = \ell^\star$. The observability index is defined as the smallest value of $i \in \mathbb{N}^+$ for which $\operatorname{rank}(\mathcal{O}_i(A, C)) = n^\star$. This index coincides with the minimal lag $\ell^\star$ (Willems, 1986, Theorem 6). Since the observability index is minimal by definition, we have $\ell^\star \leq n^\star$. Furthermore, since $\mathcal{O}_{\ell^\star}(A, C) \in \mathbb{R}^{n_y \ell^\star \times \ell^\star}$, we require $n^\star \leq n_y \ell^\star$ to ensure that $\operatorname{rank}(\mathcal{O}_i(A, C)) = n^\star$. In summary, we obtain the relation:

$$\ell^\star \leq n^\star \leq n_y \ell^\star. \tag{2.16}$$

2.2.4. ARX Representations

Note that the state-space representation provides a more structural insight than i/o representations when additional information about the state is available, e.g., from first-principle formulations. However, as we aim for a data-driven representation of

LTI systems, the model-based counterpart should be purely formulated with inputs and outputs, without explicitly defining states. Therefore, in this thesis, we adopt the *Autoregressive model with Exogenous inputs* (ARX) as the primary framework for analysis.

Consider an ARX model with lag ℓ

$$y(k) = -\sum_{i=1}^{\ell} P_{\ell-i} y(k-i) + \sum_{i=1}^{\ell} Q_{\ell-i} \tilde{u}(k-i) + D\tilde{u}(k), \quad \forall k \in \mathbb{Z}. \tag{2.17}$$

We remark that this ARX model is an input-output representation with

$$P(z) = I + \sum_{i=1}^{\ell} P_{\ell-i} z^{-i}, \text{ and } Q(z) = D + \sum_{i=1}^{\ell} Q_{\ell-i} z^{-i}.$$

Note that with the leading term of $P(z)$ as the identity, the properties required in Lemma 2.4 for an i/o representation of $\mathfrak{B}$ to be LTI are inherently guaranteed.

Let

$$\Xi \doteq [Q_{\ell-1}, \cdots, Q_0, -P_{\ell-1}, \cdots - P_0] \in \mathbb{R}^{n_y \times n_s \ell}, \tag{2.18}$$

we define the ARX representation of $\mathfrak{B}$ with lag ℓ as

$$\mathfrak{B}(\Xi, D) \doteq \left\{ (\tilde{u}, y) \,\middle|\, y(k) = \Xi \begin{bmatrix} \tilde{u}_{[k-\ell,k-1]} \\ y_{[k-\ell,k-1]} \end{bmatrix} + D\tilde{u}(k), \quad \forall k \in \mathbb{Z} \right\}. \tag{2.19}$$

The following lemma shows how a general i/o representation can be transferred into an ARX form.

Lemma 2.8 (Equivalence of i/o and ARX representations): Given an i/o representation $\mathfrak{B}(P, Q)$ of an LTI behavior $\mathfrak{B}$, where $P(z)$ is nonsingular and $P^{-1}Q$ is proper. Then, there exists a unimodular polynomial matrix $U(z) \in \mathbb{R}^{n_y \times n_y}[z, z^{-1}]$ such that

$$U(z)P(z) = I + \sum_{i=1}^{\ell} P_{\ell-i} z^{-i}, \quad U(z)Q(z) = \sum_{i=0}^{\ell} Q_{\ell-i} z^{-i}. \tag{2.20}$$

Then, consider Ξ as in (2.18) and $D = Q_\ell$, we have

$$\mathfrak{B}(P, Q) = \mathfrak{B}(\Xi, D).$$

Proof. This proof is achieved by applying similarity transformations to the kernel matrix of the original i/o representation. That is, for any unimodular $U(z)$, consider $R(z) = [Q(z), -P(z)]\Pi^{-1}$, we have

$$U(z)R(z) = \begin{bmatrix} U(z)Q(z) & -U(z)P(z) \end{bmatrix} \Pi^{-1}.$$

This implies $\mathfrak{B}(P, Q) = \mathfrak{B}(UP, UQ)$.

The next step is to determine a unimodular $U(z)$ that achieves (2.20). Following Step ii) of Lemma 2.3, given a non-singular polynomial matrix P, there exists a unimodular matrix $U'(x) \in \mathbb{R}^{n_y \times n_y}[z, z^{-1}]$, such that

$$U'(z)P(z) \doteq P'_\ell + P'_{\ell-1} z^{-1} + \cdots + P'_0 z^{-\ell}$$

with $P'_\ell \in \mathbb{R}^{n_y \times n_y}$ non-singular, i.e., $(P'_\ell)^{-1}$ exists. Finally, the required $U(z)$ is determined as $U(z) = (P'_\ell)^{-1} U'(z)$. □

Equivalence between ARX and State-Space Representations

For equivalence transformations between ARX models and state-space representations, Lemmas 3 and 4 by Sadamoto, 2023 offer detailed procedures for converting an ARX model to a minimal state-space representation and vice versa. In the following, we present two specific transformations between ARX models and state-space representations that provide key insights for subsequent discussions.

Given an ARX representation $\mathfrak{B}(\Xi, D)$, one can construct a (not necessarily minimal) state-space representation by considering the extended state

$$z(k) \doteq \begin{bmatrix} \tilde{u}_{[k-\ell,k-1]} \\ y_{[k-\ell,k-1]} \end{bmatrix} \in \mathbb{R}^{\ell(n_{\tilde{u}}+n_y)}. \tag{2.21}$$

That is, for all $k \in \mathbb{Z}$, we have

$$z(k+1) = \widetilde{A}z(k) + \widetilde{B}\tilde{u}(k), \tag{2.22a}$$
$$y(k) = \Xi z(k) + D\tilde{u}(k), \tag{2.22b}$$

with

$$\widetilde{A} \doteq \begin{bmatrix} \bar{A} \\ \Xi \end{bmatrix}, \ \widetilde{B} \doteq \begin{bmatrix} \bar{B} \\ D \end{bmatrix}, \ \bar{A} \doteq \begin{bmatrix} 0 & I_{(\ell-1)n_{\tilde{u}}} & 0 & 0 \\ 0_{n_{\tilde{u}}\times n_{\tilde{u}}} & 0 & 0_{n_{\tilde{u}}\times n_y} & 0 \\ 0 & 0 & 0 & I_{(\ell-1)n_y} \end{bmatrix}, \ \bar{B} \doteq \begin{bmatrix} 0 \\ I_{n_{\tilde{u}}} \\ 0 \end{bmatrix}.$$

Given that the components in $z\,(k+1)$ and $z\,(k)$ overlap from $k-\ell+1$ to $k-1$, as in (2.21), it follows that $\bar{A}$ and $\bar{B}$ possess a known structure independent of Ξ and D. Moreover, with $y(k)$ as the last component of $z(k+1)$, we have

$$y(k) = \widetilde{H}^\top z\,(k+1)\,, \quad \widetilde{H} \doteq \begin{bmatrix} 0_{n_y\times((n_y+n_{\tilde{u}})\ell-n_y)} & I_{n_y} \end{bmatrix}^\top. \tag{2.22c}$$

As discussed in Lemma 1 of Bongard et al., 2023, if the behavior is controllable, the pair $\left(\widetilde{A}, \widetilde{B}\right)$ is state stabilizable.

Lemma 2.9 (ARX $\Rightarrow$ state space): Given any LTI behavior $\mathfrak{B}$ and its ARX representation $\mathfrak{B}(\Xi, D)$, then $\mathfrak{B}\left(\widetilde{A}, \widetilde{B}, \Xi, D\right)$, i.e. (2.22), is a state-space representation of $\mathfrak{B}$. That is,

$$\mathfrak{B}(\Xi, D) = \mathfrak{B}\left(\widetilde{A}, \widetilde{B}, \Xi, D\right).$$

Proof. We show the equivalence by proving both inclusions: $\mathfrak{B}(\Xi, D) \subseteq \mathfrak{B}\left(\widetilde{A}, \widetilde{B}, \Xi, D\right)$ and $\mathfrak{B}\left(\widetilde{A}, \widetilde{B}, \Xi, D\right) \subseteq \mathfrak{B}(\Xi, D)$.

Given any trajectory pair $(\tilde{u}, y) \in \mathfrak{B}(\Xi, D)$, construct a trajectory $z \in (\mathbb{R}^{\ell(n_{\tilde{u}}+n_y)})^{\mathbb{Z}}$ as in (2.21) such that for each $k \in \mathbb{Z}$, the state $z(k)$ consists of the last ℓ inputs and outputs. Then, the triple $(z, \tilde{u}, y)$ satisfies the dynamics (2.22). By definition:

$$\mathfrak{B}(\widetilde{A}, \widetilde{B}, \Xi, D) \doteq \left\{(\tilde{u}, y) \;\middle|\; \exists z \in (\mathbb{R}^{\ell(n_{\tilde{u}}+n_y)})^{\mathbb{Z}}, \text{ such that } (z, \tilde{u}, y) \text{ satisfy (2.22)}\right\},$$

and hence, the inclusion $\mathfrak{B}(\Xi, D) \subseteq \mathfrak{B}(\widetilde{A}, \widetilde{B}, \Xi, D)$ follows.

For the other inclusion, for any trajectory pair $(z, \tilde{u}, y)$ satisfying (2.22), let

$$z \doteq \begin{bmatrix} z_1 & z_2 & \cdots & z_{\ell(n_{\tilde{u}}+n_y)} \end{bmatrix}^\top, \quad z_i : \mathbb{Z} \to \mathbb{R}, \quad i \in \mathbb{I}_{[1,\ell(n_{\tilde{u}}+n_y)]}.$$

From the structure of $\bar{B}$ and (2.22c), we have

$$z_{[(\ell-1)n_{\tilde{u}}+1,\ell n_{\tilde{u}}]} \doteq \begin{bmatrix} z_{(\ell-1)n_{\tilde{u}}+1} & \cdots & z_{\ell n_{\tilde{u}}} \end{bmatrix}^\top = \boldsymbol{\sigma}^{-1}u, \quad z_{[\ell n_{\tilde{u}}+(\ell-1)n_y+1,\ell(n_{\tilde{u}}+n_y)]} = \boldsymbol{\sigma}^{-1}y.$$

Moreover, due to block structure of $\bar{A}$, for each $i = 1, \cdots, \ell-1$

$$\begin{aligned} z_{[(i-1)n_{\tilde{u}}+1,in_{\tilde{u}}]} &= \boldsymbol{\sigma}^{-1} z_{[in_{\tilde{u}}+1,(i+1)n_{\tilde{u}}]} &&= \boldsymbol{\sigma}^{-(\ell-i+1)}u, \\ z_{[\ell n_{\tilde{u}}+(i-1)n_y+1,\ell n_{\tilde{u}}+in_y]} &= \boldsymbol{\sigma}^{-1} z_{[\ell n_{\tilde{u}}+in_y+1,\ell n_{\tilde{u}}+(i+1)n_y]} &&= \boldsymbol{\sigma}^{-(\ell-i+1)}y. \end{aligned}$$

Therefore, for all $k \in \mathbb{Z}$, the state $z(k)$ consists of the last ℓ inputs and outputs, satisfying the structure in (2.21). Combined with (2.22b), this implies $(\tilde{u}, y) \in \mathfrak{B}(\Xi, D)$, and the inclusion $\mathfrak{B}\left(\widetilde{A}, \widetilde{B}, \Xi, D\right) \subseteq \mathfrak{B}(\Xi, D)$ holds.

Hence, we conclude that $\mathfrak{B}(\Xi, D) = \mathfrak{B}(\widetilde{A}, \widetilde{B}, \Xi, D)$. □

The next lemma shows that we can construct an ARX representation with lag $\ell \geq \ell^\star$ from any minimal state-space representation. The crucial step in this transformation lies in eliminating x using a dead-beat observer gain H such that $(A - HC)^\ell = 0$.

Lemma 2.10 (Minimal state space ⇒ ARX): Given any LTI behavior $\mathfrak{B}$ and a minimal state-space representation $\mathfrak{B}(A, B, C, D)$. Then, for any $\ell \geq \ell^\star$, there exist $H \in \mathbb{R}^{n_x \times n_y}$ and an ARX representation $\mathfrak{B}(\Xi, D)$ as in (2.19) such that $(A - HC)^\ell = 0$ and

$$\Xi = C \begin{bmatrix} \mathcal{R}_\ell (A - HC, B - HD) & \mathcal{R}_\ell (A - HC, H) \end{bmatrix}. \tag{2.23}$$

Proof. The proof relies on findings from Wu, 2022 and employs a dead-beat observer gain to transform the minimal state-space representation to the ARX representation. Since $\mathfrak{B}(A, B, C, D)$ is minimal, the pair (A, C) is state observable. Hence, there exists a deadbeat observer gain H such that $(A - HC)^\ell = 0$ (Fahmy et al., 1980).

We rewrite (2.11) as

$$\begin{aligned} x(k+1) &= Ax(k) + B\tilde{u}(k) + H(y(k) - y(k)) \\ &= (A - HC)x(k) + (B - HD)\tilde{u}(k) + Hy(k). \end{aligned}$$

By recursively expressing x from $k - \ell$ to k, we derive

$$x(k) = (A - HC)^\ell x(k-\ell) + \sum_{i=1}^{\ell} (A - HC)^{i-1} \Big(Hy(k-i) + (B - HD)\tilde{u}(k-i) \Big).$$

Multiplying both sides by C and adding $D\tilde{u}(k)$, we obtain

$$y(k) = \sum_{i=1}^{\ell} C(A - HC)^{i-1} \Big(Hy(k-i) + (B - HD)\tilde{u}(k-i) \Big) + D\tilde{u}(k),$$

with $(A - HC)^\ell = 0$. Using the reachability matrix definition in (2.15)

$$\sum_{i=1}^{\ell} C(A - HC)^{i-1}(B - HD)\tilde{u}(k-i) = C\mathcal{R}_\ell(A - HC, B - HD)\tilde{u}_{[k-\ell,k-1]},$$

$$\sum_{i=1}^{\ell} C(A - HC)^{i-1} Hy(k-i) = C\mathcal{R}_\ell(A - HC, H)y_{[k-\ell,k-1]}.$$

Thus, the proof is complete. □

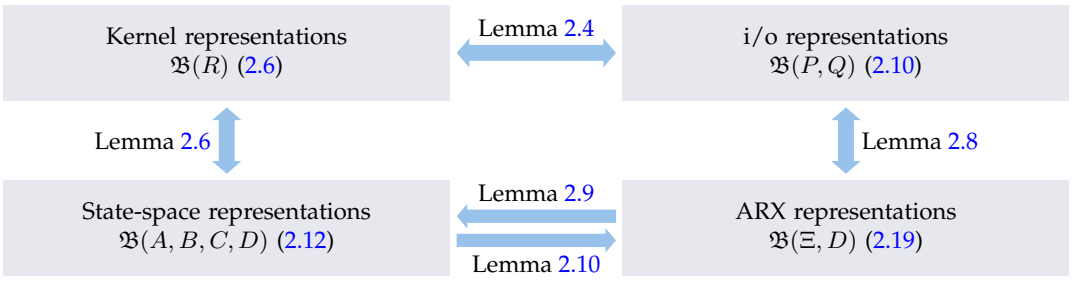

Figure 2.2.: Parametric representations for the LTI behavior $\mathfrak{B}$

Remark 2.2 (Comparison of different representations): Figure 2.2 provides a summary of the parametric representations of LTI behaviors along with lemmas demonstrating their equivalences. These representations differ in their structural characteristics. Kernel representations impose the least structure, demonstrating $\mathfrak{B}$ as a linear subspace of $(\mathbb{R}^{n_s})^{\mathbb{Z}}$. Furthermore, i/o representations explicitly define the i/o structure by identifying the input $\tilde{u}$ as the manifest components that are free, processed by, and not anticipated by the output. Additionally, state-space representations introduce the latent state variables x, which act as a memory function that condense relevant past information for predicting future system response (Willems, 1986, p.565).

Among those representations, minimal formulations are of particular interest. Specifically, minimal kernel representations use the fewest many linear equalities to characterize the behavior, while the minimal-lag kernel representations further minimize the required time span. Remarkably, i/o representations are inherently minimal kernel representations, meaning the minimal number of equations required to specify an LTI behavior equals the dimension of the output. The minimal state-space representations, on the other hand, use the smallest number of states to retain past information necessary for predicting the future. In summary, the three minimal indices $n_R^\star = n_y, \ell^\star, n^\star$ satisfy the relation given in (2.16).

2.3. Non-Parametric System Representation

Non-parametric system representations provide a data-driven alternative for modeling dynamical systems. Unlike parametric methods, which rely on predefined model structures, non-parametric approaches leverage recorded data directly to characterize system behavior, bypassing the need for explicit modeling.

2.3.1. The Fundamental Lemma by Willems et al.

The fundamental lemma by Willems et al., 2005 provides a non-parametric representation of controllable LTI systems using Hankel matrices constructed only from measured data, which carries sufficiently rich information. This richness is specified as follows.

Definition 2.8 (Persistency of excitation (Willems et al., 2005)): Let $T, N \in \mathbb{N}^+$. A signal

sequence $s_{[0,T-1]}$ is said to be persistently exciting of order N if the Hankel matrix

$$\mathcal{H}_N(s_{[0,T-1]}) \doteq \begin{bmatrix} s(0) & s(1) & \cdots & s(T-N) \\ s(1) & s(2) & \cdots & s(T-N+1) \\ \vdots & \vdots & \ddots & \vdots \\ s(N-1) & s(N) & \cdots & s(T-1) \end{bmatrix} \tag{2.24}$$

is of full row rank.

We now present the fundamental lemma by Willems et al., 2005.

Lemma 2.11 (Fundamental lemma (Willems et al., 2005)):
Let T, $N \in \mathbb{N}^+$. Consider a controllable LTI system $(\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{B})$ that satisfies Definitions (2.1) and (2.2), with an i/o partitioning $s = (\tilde{u}, y)$, and its minimal state dimension $n^\star$. For a T-length trajectory $s^{\mathrm{d}}_{[0,T-1]} = \left(\tilde{u}^{\mathrm{d}}, y^{\mathrm{d}}\right)_{[0,T-1]} \in \mathfrak{B}_T$, if $\tilde{u}^{\mathrm{d}}_{[0,T-1]}$ is persistently exciting of order $N + n^\star$, then any N-length trajectory $s_{[0,N-1]} \in \mathfrak{B}_N$ if and only if there exists $g \in \mathbb{R}^{T-N+1}$ such that

$$\mathcal{H}_N\left(s^{\mathrm{d}}_{[0,T-1]}\right) g = s_{[0,N-1]}. \tag{2.25}$$

Equivalently, we have

$$\operatorname{colspan}\left(\mathcal{H}_N\left(s^{\mathrm{d}}_{[0,T-1]}\right)\right) = \mathfrak{B}_N.$$

A key advantage of these results is that Hankel matrices can be directly constructed from measured input and output data, eliminating the need for explicit system models.

Compared to the subspace identification methods such as N4SID (van Overschee et al., 2012) and MOSEP (Verhaegen et al., 1992a,b), the fundamental lemma avoids the assumption on the persistency of excitation of the underlying states. Specifically, subspace identification methods often require the condition

$$\operatorname{rank}\begin{bmatrix}\mathcal{H}_1(x_{[0,T-N]}) \\ \mathcal{H}_N(\tilde{u}_{[0,T-1]})\end{bmatrix} = n^\star + N n_{\tilde{u}}.$$

Instead, the fundamental lemma requires the system to be controllable and the inputs to be persistently exciting of order $N + n^\star$, with $n^\star$ extra order to excite the state. Furthermore, the condition of persistently exciting can be refined (van Waarde, 2021), and the controllability assumption can be relaxed (Yu et al., 2021).

2.3.2. Data-Driven Deterministic Prediction

Predicting future trajectories of a dynamical system relies on three key elements: the future input, the system dynamics represented by $\mathfrak{B}$, and the initial condition. While the dynamics are implicitly captured via Hankel matrices in the fundamental lemma, we next determine the initial condition necessary for a prediction without requiring explicit state information.

The following lemma specifies the initial conditions required for uniquely predicting future trajectories of an LTI system based solely on past input and output data.

Lemma 2.12 (Initial condition for a finite-length behavior): Consider an LTI behavior $\mathfrak{B}$ with an i/o partitioning $s = (\tilde{u}, y)$. Given a lag $\ell \geq \ell^\star$, a prediction horizon $N \in \mathbb{N}^+$, and an input trajectory $\tilde{u}_{[0,N-1]} \in \mathbb{R}^{N n_{\tilde{u}}}$, there is a unique $y_{[0,T-1]}$ such that $(\tilde{u}, y)_{[0,T-1]} \in \mathfrak{B}_N$ if the initial condition $(\tilde{u}, y)_{[-\ell,-1]} \in \mathfrak{B}_\ell$ is specified.

Proof. By Lemma 2.6, for an LTI behavior $\mathfrak{B}$, there is a minimal state space representation of $\mathfrak{B}$, i.e. $\mathfrak{B}(A, B, C, D)$, with (A, C) observable. Furthermore, by Lemma 1 of Markovsky et al., 2008, $(\tilde{u}, y)_{[-\ell,-1]} \in \mathfrak{B}_\ell$ determines a unique $x(0) \in \mathbb{R}^{n^\star}$. By applying Lemma 2.5, we further obtain unique trajectories $y_{[0,N-1]}$, $x_{[0,N-1]}$ such that $(x, \tilde{u}, y)_{[0,T-1]} \in \mathfrak{B}^{\mathrm{L}}_N(A, B, C, D)$, which concludes that $(\tilde{u}, y)_{[0,T-1]} \in \mathfrak{B}_N$. □

Combining the above results, we arrive at a data-driven predictor for controllable LTI systems. To predict the future trajectory over a horizon N, given the initial condition $(\tilde{u}, y)_{[-\ell,-1]} \in \mathfrak{B}_\ell$ and the input trajectory $\tilde{u}_{[0,N-1]} \in \mathbb{R}^{Nn_u}$, we aim to determine a unique trajectory $(\tilde{u}, y)_{[-\ell,N-1]} \in \mathfrak{B}_{N+\ell}$.

Using Lemma 2.11, we record a T-length trajectory $(\tilde{u}^{\mathrm{d}}, y^{\mathrm{d}})_{[0,T-1]} \in \mathfrak{B}_T$ with $\tilde{u}^{\mathrm{d}}_{[0,T-1]}$ being persistently exciting of order $N + n^\star + \ell$. We partition the recorded data into two parts which correspond to the prediction of horizon $[-\ell, -1]$ and $[0, N-1]$, respectively. That is,

$$\mathcal{H}_{N+\ell}(\tilde{u}_{[0,T-1]}) \doteq \begin{bmatrix} \mathcal{H}_{\tilde{u},\mathrm{p}} \\ \mathcal{H}_{\tilde{u},\mathrm{f}} \end{bmatrix} \tag{2.26}$$

where $\mathcal{H}_{\tilde{u},\mathrm{p}}$ consists of the first ℓ block rows of $\mathcal{H}_{N+\ell}\left(\tilde{u}_{[0,T-1]}\right)$ and $\mathcal{H}_{\tilde{u},\mathrm{f}}$ consists of the remaining N block rows. Similarly, we define $\mathcal{H}_{y,\mathrm{p}}$, $\mathcal{H}_{y,\mathrm{f}}$. Let p and f denote the ranges $[-\ell, -1]$ and $[0, N-1]$ respectively, an exact data-driven output predictor can be formulated as

$$y_{\mathrm{f}} = \mathcal{H}_{y,\mathrm{f}} g, \quad g \text{ subject to } \begin{bmatrix} \mathcal{H}_{\tilde{u},\mathrm{p}} \\ \mathcal{H}_{y,\mathrm{p}} \\ \mathcal{H}_{\tilde{u},\mathrm{f}} \end{bmatrix} g = \begin{bmatrix} \tilde{u}_{\mathrm{p}} \\ y_{\mathrm{p}} \\ \tilde{u}_{\mathrm{f}} \end{bmatrix}. \tag{2.27}$$

As shown in Theorem 1 of Fiedler et al., 2021, the solution y_{f} to (2.27) is unique. This approach allows for an exact determination of future output trajectories based solely on measured data, without requiring explicit knowledge of the state variables. As a closed-form solution, we have

$$y_{\mathrm{f}} = \begin{bmatrix} \mathcal{H}_{\tilde{u},\mathrm{p}} \\ \mathcal{H}_{y,\mathrm{p}} \\ \mathcal{H}_{\tilde{u},\mathrm{f}} \end{bmatrix}^{\dagger} \begin{bmatrix} \tilde{u}_{\mathrm{p}} \\ y_{\mathrm{p}} \\ \tilde{u}_{\mathrm{f}} \end{bmatrix}. \tag{2.28}$$

A Subspace Predictor for Noise Corrupted Data

Suppose that the data $(\tilde{u}, y)^{\mathrm{d}}_{[0,T-1]}$ now are corrupted with measurement noise and some disturbances that are not specified in the manifest input $\tilde{u}$, the column span of Hankel matirces then does not represent behaviors of any LTI system. Thus, (2.27) cannot be solved with one unique solution but instead admits infinitely many solutions for y_{f}. Therefore, one may consider the least-square solution

$$\hat{y}_{\mathrm{f}} = \mathcal{H}_{y,\mathrm{f}} g^\star, \quad g^\star = \operatorname*{argmin}_g \|g\|^2 \text{ subject to } \begin{bmatrix} \mathcal{H}_{\tilde{u},\mathrm{p}} \\ \mathcal{H}_{y,\mathrm{p}} \\ \mathcal{H}_{\tilde{u},\mathrm{f}} \end{bmatrix} g = \begin{bmatrix} \tilde{u}_{\mathrm{p}} \\ y_{\mathrm{p}} \\ \tilde{u}_{\mathrm{f}} \end{bmatrix}. \tag{2.29}$$

with its closed-form solution the same as (2.28). We note that (2.29) is also known as a subspace predictor; see Favoreel et al., 1999; Fiedler et al., 2021; Sedghizadeh et al., 2018.

2.4. Summary

In this chapter, we examined various representations of deterministic LTI systems, considering both parametric and non-parametric frameworks. Parametric representations, including kernel, input/output, state-space, and ARX representations, provide diverse perspectives on system behavior, each emphasizing distinct structural characteristics. For non-parametric representations, we employed the fundamental lemma by Willems et al. to characterize controllable LTI behaviors solely based on past recorded data. This data-driven approach forms the foundation for designing predictors and controllers without relying on explicit model parameters. In the next chapter, we introduce essential concepts related to stochastic uncertainties, leading to a stochastic variant of the fundamental lemma in Chapter 4.

3. Representations of Stochastic Uncertainty

In this chapter, we explore various frameworks for representing stochastic uncertainties. In Section 3.1, we begin by revisiting essential concepts in probability theory to rigorously define random variables and their distributions. Next, in Section 3.2, we introduce the framework of Polynomial Chaos Expansion (PCE), which expresses random variables in terms of deterministic expansion coefficients in suitable polynomial bases. Section 3.3 introduces stochastic processes and their PCEs, which plays a crucial role in analyzing stochastic LTI systems in Chapter 4. Lastly, Section 3.4 addresses uncertainty quantification through probabilistic bounds. Specifically, two-sided bounds derived from Chebyshev's inequality are used to construct confidence intervals, while one-sided bounds from Cantelli's inequality are applied to chance constraints.

3.1. Random Variables and Probability Spaces

In this section, we revisit essential notions of probability theory to describe probability spaces and random variables. For rigorous definitions, we refer to the textbooks by Sullivan, 2015 and Fristedt et al., 2013.

Definition 3.1 (Probability space, cf. p.10 of Sullivan, 2015): A *probability space* $(\Omega, \mathcal{F}, \mu)$ is a triple,

i) Ω is a set, called the *sample space*. Elements of Ω, i.e. $\omega \in \Omega$ are called *outcomes*.

ii) $\mathcal{F}$ is a σ-algebra on Ω, i.e. a collection of subsets of Ω containing $\emptyset$ and closed under countable applications of union ($\cup$), intersection ($\cap$) and complementation ($\cdot^{\mathrm{C}}$) relative to Ω.

iii) Elements of $\mathcal{F}$, i.e., $E \in \mathcal{F}$, are called *measurable sets* or *events*, and $(\Omega, \mathcal{F})$ is called a *measurable space*.

iv) The function $\mu : \mathcal{F} \to [0, 1]$ with $\mu(\Omega) = 1$ is called a *probability measure* on the measurable space $(\Omega, \mathcal{F})$.

Put simply, a probability space $(\Omega, \mathcal{F}, \mu)$ models a random experiment, where Ω represents the set of all possible outcomes, $\mathcal{F}$ collects all measurable events, and μ assigns probabilities to these events, ranging between 0 and 1.

The σ-algebra $\mathcal{F}$ characterizes how "coarse" (or how "fine") the sample space Ω is measurable. In the coarsest scenario, with the trivial σ-algebra $\{\emptyset, \Omega\}$, only the empty set $\emptyset$ and the entire set Ω are measurable. In contrast, for the finest case, one may define the σ-algebra $\mathcal{F}$ as the power set 2^{Ω}, which collects all subsets of Ω. However, when Ω contains infinitely many elements, the power set is often too large and impractical for most applications. Instead, it is common to use the Borel σ-algebra $\mathcal{B}(\Omega)$ as $\mathcal{F}$, which is the smallest σ-algebra on Ω that ensures every open set (and hence every closed set, by complement) is measurable.

Between the finest and coarsest cases, when only the probabilities of certain events $\mathcal{G} \subset 2^{\Omega}$ are known due to limited information, one can define the smallest σ-algebra $\sigma(\mathcal{G})$ on Ω containing $\mathcal{G}$ as $\mathcal{F}$. Here, $\sigma(\mathcal{G})$ is called the σ-algebra on Ω generated by $\mathcal{G}$.

Lemma 3.1 (σ-algebra generated by $\mathcal{G}$, cf. Definition 1.5 of Fristedt et al., 2013): Let $\mathcal{G}$ be a collection of subsets of Ω, i.e., $\mathcal{G} \subset 2^{\Omega}$. Then, there exists a unique σ-algebra $\sigma(\mathcal{G}) \supseteq \mathcal{G}$ such that if $\mathcal{F} \supseteq \mathcal{G}$ and $\mathcal{F}$ is a σ-algebra on Ω, then $\mathcal{F} \supseteq \sigma(\mathcal{G})$. In other words, $\sigma(\mathcal{G})$ represents the smallest σ-algebra on Ω containing $\mathcal{G}$.

The following example illustrates how a random experiment can be modeled as a probability space.

Example 3.1 (Throwing a dice once): Using a single dice throw as an example for a random experiment, where the outcome ω is the number on top, we assume equal probability for each facet. This experiment can be modelled by $(\Omega, \mathcal{F}, \mu)$, where $\Omega \doteq \{1,2,3,4,5,6\}$, $\mathcal{F} \doteq 2^{\Omega} = \{\emptyset, \{1\}, \cdots, \{2,3,4,5,6\}, \Omega\}$, and $\mu \doteq \#E/6$ with $\#E$ denoting the cardinality of the event E. Here, the probability measure μ then assigns each event $E \in \mathcal{F}$ a probability value; for example $\mu(\omega \leq 4) = 0.5$.

Now, consider the scenario where the dice is damaged on some facets, but we are certain that the probability of $\omega = 1$ remains $\frac{1}{6}$. In this case, one may consider $\mathcal{F} = \{\emptyset, \{1\}, \{2,3,4,5,6\}, \Omega\}$. However, with this reduced $\mathcal{F}$, we cannot determine $\mu(\omega \leq 4)$ because the event $\{\omega \in \Omega \,|\, \omega \leq 4\}$ is not an element of $\mathcal{F}$.

Random Variables

After the outcome ω of a random experiment modeled by $(\Omega, \mathcal{F}, \mu)$ is determined, a random variable V further maps the outcome $\omega \in \Omega$ to a vector $V(\omega) \in \mathbb{R}^{n_v}$, which is known as the realization of V.

Definition 3.2 ($\mathbb{R}^{n_v}$-valued random variable, cf. p.14 of Sullivan, 2015): Given a probability space $(\Omega, \mathcal{F}, \mu)$ and a measurable space $(\mathbb{R}^{n_v}, \mathcal{B}(\mathbb{R}^{n_v}))$ with the Borel σ-algebra $\mathcal{B}(\mathbb{R}^{n_v})$.

i) A *$\mathbb{R}^{n_v}$-valued random variable* V is a measurable function $V : \Omega \to \mathbb{R}^{n_v}$ such that for every Borel subset $E \in \mathcal{B}(\mathbb{R}^{n_v})$, the pre-image

$$V^{-1}(E) \doteq \{\omega \in \Omega \,|\, V(\omega) \in E\}$$

is an element of $\mathcal{F}$. As a shorthand, we denote $V \in (\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$.

ii) With $\omega \in \Omega$, we denote $V(\omega) \in \mathbb{R}^{n_v}$ as a *realization* of V.

iii) The *distribution* of V is the probability measure $\mu_V : \mathcal{B}(\mathbb{R}^{n_v}) \to [0,1]$ defined on $(\mathbb{R}^{n_v}, \mathcal{B}(\mathbb{R}^{n_v}))$ satisfying

$$\mu_V(E) \doteq \mu\left(V^{-1}(E)\right), \quad \forall E \in \mathcal{B}(\mathbb{R}^{n_v}).$$

For compactness, we write $V \sim \mu_V$.

Here, item i) implies that the smallest σ-algebra $\sigma(V)$, which covers all events necessary to define V, is a subset of $\mathcal{F}$, i.e.,

$$\sigma(V) \doteq \sigma\big(\{V^{-1}(E) \,|\, E \in \mathcal{B}(\mathbb{R}^{n_v})\}\big) \subseteq \mathcal{F}. \tag{3.1}$$

Henceforth, we refer to $\sigma(V)$ as the σ-algebra on Ω generated by the random variable V.

In other words, (3.1) requires the σ-algebra $\mathcal{F}$ to be rich enough to define V. If, for example, $\mathcal{F}$ is the power set of Ω, then every function $V : \Omega \to \mathbb{R}^{n_v}$ is measurable. At the opposite extreme, if $\mathcal{F}$ is the trivial σ-algebra $\{\emptyset, \Omega\}$, then the only measurable functions are constant functions. Hence, for a deterministic random variable (i.e., $V = v$ almost surely), the σ-algebra it generates is simply $\sigma(V) = \{\emptyset, \Omega\}$.

Moreover, with the distribution μ_V defined in item iii), we note that the random variable V indeed maps a probability space to another, i.e. from $(\Omega, \mathcal{F}, \mu)$ to $(\mathbb{R}^{n_v}, \mathcal{B}(\mathbb{R}^{n_v}), \mu_V)$; see Proposition 2.5 of Fristedt et al., 2013.

L^2 Random Variables

In this thesis, we focus on *L^2 random variables*, which are random variables with finite expectation and covariance.

Definition 3.3 (L^2 random variable space): Let $(\Omega, \mathcal{F}, \mu)$ be a probability space. The L^2 random variable space is defined as

$$L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v}) \doteq \{V : \Omega \to \mathbb{R}^{n_v} \mid V \text{ satisfies Definition 3.2 and } \|V\| \text{ is finite}\}, \quad (3.2)$$

where the inner product and norm on the L^2 space are defined as

$$\langle V, \tilde{V} \rangle \doteq \int_\Omega V(\omega)^\top \tilde{V}(\omega)\, \mathrm{d}\mu(\omega), \quad \|V\| \doteq \sqrt{\langle V, V \rangle},\ V, \tilde{V} \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v}). \quad (3.3)$$

Consider two random variables $V, \tilde{V} \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ with $V \sim \mu_V$, $\tilde{V} \sim \mu_{\tilde{V}}$, and $(V, \tilde{V}) \sim \mu_{V,\tilde{V}}$. The equivalence $V = \tilde{V}$ is to be understood in the L^2 sense (Fristedt et al., 2013), i.e.,

$$V = \tilde{V} \quad \Leftrightarrow \quad \|V - \tilde{V}\| = 0 \quad \Leftrightarrow \quad V(\omega) = \tilde{V}(\omega), \text{ for } \mu\text{-almost all } \omega \in \Omega. \quad (3.4)$$

Here, μ-almost all outcomes exclude those in a set with measure zero under μ, meaning the probability of such outcomes occurring is zero.

The expectation and the covariance of V are defined as $\mathbb{E}[V] \in \mathbb{R}^{n_v}$ and $\Sigma[V] \in \mathbb{R}^{n_v \times n_v}$, respectively. That are

$$\mathbb{E}[V] \doteq \int_\Omega V(\omega) \mathrm{d}\mu(\omega) = \int_{\mathbb{R}^{n_v}} v\, \mathrm{d}\mu_V(v), \quad (3.5a)$$

$$\Sigma[V] \doteq \int_\Omega (V(\omega) - \mathbb{E}[V])(V(\omega) - \mathbb{E}[V])^\top \mathrm{d}\mu(\omega) \quad (3.5b)$$

$$= \int_{\mathbb{R}^{n_v}} (v - \mathbb{E}[V])(v - \mathbb{E}[V])^\top \mathrm{d}\mu_V(v).$$

In addition, we denote the covariance between V and $\tilde{V}$ as $\Sigma[V, \tilde{V}]$. Moreover, as per Definition 3.3, $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ implies that $\mathbb{E}[V]$ and $\Sigma[V]$ is finite.

Distribution Functions

The random variable $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ with $V \sim \mu_V$ is discrete if $\mu_V(V \in \mathcal{K}) = 1$ for some finite or countably infinite set $\mathcal{K} \subset \mathbb{R}^{n_v}$. It is continuous if $\mu_V(V = v) = 0$ for all $v \in \mathbb{R}^{n_v}$. We define the *cumulative distribution function* (cdf) of V to be

$$F_V : \mathbb{R}^{n_v} \to \mathbb{R}, \quad F_V(v) \doteq \mu_V(V < v),$$

where '$<$' is applied elementwise. In case that V is a discrete random variable, we define the associated *probability mass function* (pmf) as

$$p_V : \mathbb{R}^{n_v} \to \mathbb{R}, \quad p_V(v) \doteq \mu_V(V = v).$$

If V is a continuous random variable, we define the associated *probability density function* (pdf) f_V as the derivative of the cdf

$$f_V : \mathbb{R}^{n_v} \to \mathbb{R}, \quad f_V(v) \doteq \begin{cases} \nabla F_V(v), & \text{if } \nabla F_V(v) \text{ exists,} \\ 0, & \text{else.} \end{cases}$$

Example 3.2 (Discrete RV: throwing a dice twice): Consider the sample space $\Omega = \{\omega = (a, b) : a \in \{1, \cdots, 6\}, b \in \{1, \cdots, 6\}\}$, the probability measure $\mu(E) = \#E/36$, and the random variable $V(a, b) = a + b$, we can define the following events and their probabilities

$$\{V = 4\} = \{(1, 3), (3, 1), (2, 2)\}, \quad \mu_V(V = 4) = \mu(\{(1, 3), (3, 1), (2, 2)\}) = 3/36,$$
$$\{V \leq 4\} = \{(1, 1), (1, 2), (2, 1)\}, \quad \mu_V(V \leq 4) = \mu(\{(1, 1), (1, 2), (2, 1)\}) = 3/36.$$

Note that V in this case is a discrete random variable and its pmf reads

$$p_V(v) = \begin{cases} (v - 1)/36, \ v \in \{2, 3, 4, 5, 6, 7\} \\ (13 - v)/36, \ v \in \{8, 9, 10, 11, 12\} \\ 0, \text{ else.} \end{cases}$$

Definition 3.4 (Gaussian RV): The random variable $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ is Gaussian with mean $m \in \mathbb{R}^{n_v}$ and covariance $\Sigma \in \mathbb{R}^{n_v \times n_v}$ with $\Sigma \succeq 0$ if its pdf reads

$$f_V(v) = \frac{1}{(2\pi)^{n_v/2}\sqrt{\det \Sigma}} \exp\left(-\frac{1}{2}(v - m)^\top \Sigma^{-1}(v - m)\right).$$

Compactly, we denote the distribution μ_V as $\mathcal{N}(m, \Sigma)$ and thus $V \sim \mathcal{N}(m, \Sigma)$.

Remark 3.1 (Why use random variables instead of individual distributions?): Representing stochastic uncertainties through random variables links them to a single probability space $(\Omega, \mathcal{F}, \mu)$, rather than assigning each uncertainty an independent distribution. This approach inherently captures dependencies between random variables without explicitly defining their joint distribution.

Remark 3.2 (From distribution to random variable): A random variable provides more information than its distribution. However, there might be cases where the underlying probability space $(\Omega, \mathcal{F}, \mu)$ is not explicitly defined but the distribution μ_V can be estimated through realizations of V. In such cases, one may reconstruct a random variable on the probability space $(\mathbb{R}^{n_v}, \mathcal{B}(\mathbb{R}^{n_v}), \mu_V)$ by considering V as the identity function $V : v \mapsto v$. For an one-dimensional distribution $\mu_V : \mathcal{B}(\mathbb{R}) \to [0, 1]$ with cdf $F_V : \mathbb{R} \to [0, 1]$, V, an alternative construction is: Let $\Omega = [0, 1]$, $\mathcal{F} = \mathcal{B}([0, 1])$ (the Borel algebra on Ω), $\mu = \mathcal{U}([0, 1])$ (the uniform distribution on Ω). Then, V can be defined as $V : \omega \mapsto F_V^{-1}(V(\omega))$, where F_V^{-1} is the quantile function of μ_V; see Proposition 4 in Section 3.1 of Fristedt et al., 2013.

3.2. Polynomial Chaos Expansion

This section shows how an L^2 random variable can be transformed to real numbers using PCE. The key idea behind PCE is that an L^2 random variable can be expressed in a suitable polynomial basis. This concept dates back to Wiener, 1938; for a general introduction to PCE, see Chapters 11.2 and 11.3 of Sullivan, 2015.

Let $\xi_i \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$, $i \in \mathbb{I}_{[0,n_\xi]}$ with $n_\xi \in \mathbb{N} \cup \{\infty\}$ be a finite or countably infinite collection of independent random variables. Define $\xi \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\xi})$ as their vectorization. Consider $\mathcal{F}$ to be the σ-algebra generated by ξ, i.e., $\mathcal{F} = \sigma(\xi)$. Then, these random variables are referred to as *stochastic germs* since they cover all necessary information to represent other random variables in $L^2(\Omega, \mathcal{F}, \mu)$. As the generalization of the theorem by Cameron et al., 1947, using the Gram-Schmidt process, we can construct an orthogonal polynomial basis $\{\phi^j(\xi)\}_{j\in\mathbb{N}}$ that spans $L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$

$$\phi^j : L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\xi}) \to L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}),$$
$$\langle \phi^i, \phi^j \rangle = \int_\Omega \phi^i(\xi(\omega))\phi^j(\xi(\omega))\, \mathrm{d}\mu(\omega) = \delta^{ij}\|\phi^j\|^2,$$

where δ^{ij} is the Kronecker delta.

As a consequence, every random variable in $L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$ can be represented as an L^2-convergent series using the orthogonal polynomial basis $\{\phi^j(\xi)\}_{j\in\mathbb{N}}$.

Definition 3.5 (Polynomial chaos expansion): The PCE of a random variable $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$ with respect to the basis $\{\phi^j(\xi)\}_{j=0}^\infty$ is

$$V = \sum_{j=0}^{\infty} \mathsf{v}^j \phi^j(\xi) \quad \text{with} \quad \mathsf{v}^j \doteq \frac{\langle V, \phi^j(\xi)\rangle}{\|\phi^j(\xi)\|^2},$$

where $\mathsf{v}^j \in \mathbb{R}$ is called the j-th PCE coefficient.

We remark that after applying PCE componentwise, the j-th PCE coefficient of a random vector $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ reads

$$\mathsf{v}^j \doteq \begin{bmatrix} \mathsf{v}_1^j & \mathsf{v}_2^j & \cdots & \mathsf{v}_{n_v}^j \end{bmatrix}^\top \in \mathbb{R}^{n_v},$$

where v_i^j is the j-th PCE coefficient of the $i-$th component V_i. Due to the L^2 nature of V, the sequence of PCE coefficients $\mathsf{v} \doteq (\mathsf{v}^0, \mathsf{v}^1, \cdots)$ is square summable, i.e.

$$\mathsf{v} \in \ell^2(\mathbb{R}^{n_v}) \doteq \left\{ \mathsf{v} \in (\mathbb{R}^{n_v})^{\mathbb{Z}} \,\middle|\, \sum_{j\in\mathbb{N}} (\mathsf{v}^j)^\top \mathsf{v}^j < \infty \right\}.$$

In numerical implementations, the series has to be terminated after a finite number of terms which may lead to truncation errors. For details on truncation errors and error propagation see Field et al., 2004 and Mühlpfordt et al., 2017.

Definition 3.6 (Exact PCE representation): A random variable $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$ is said to admit an exact PCE with L terms in $\{\phi^j(\xi)\}_{j\in\mathbb{N}}$ if

$$V - \sum_{j=0}^{L-1} \mathsf{v}^j \phi^j(\xi) = 0,$$

where the equality holds in the L^2 sense as per (3.4).

Table 3.1.: Correspondence of random variables and underlying orthogonal polynomials.

Distribution	Support	Orthogonal basis $\{\phi^j\}_{j=0}^{\infty}$	Argument $\xi_i(\omega)$
Gaussian	$(-\infty, \infty)$	Hermite	$\mathcal{N}(0,1)$
Uniform	$[a, b]$	Legendre	$\mathcal{U}([-1,1])$
Beta	$[a, b]$	Jacobi	$B(\alpha, \beta, [-1,1])$
Gamma	$(0, \infty)$	Laguerre	$\Gamma(\alpha, \beta, (0,\infty))$

For an arbitrary L^2 random vector $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ admitting exact PCEs with L terms, we introduce the following notations for compactness. Let

$$\mathsf{v}^{[0,L-1]} \doteq \left[\mathsf{v}^{0\top}, \mathsf{v}^{1\top}, \ldots, \mathsf{v}^{L-1\top}\right]^\top \in \mathbb{R}^{n_v L}$$

denote the vectorization of the coefficients over the PCE dimensions, and let

$$\mathcal{H}_1\left(\mathsf{v}^{[0,L-1]}\right) \doteq \left[\mathsf{v}^0, \mathsf{v}^1, \cdots, \mathsf{v}^{L-1}\right] \in \mathbb{R}^{n_v \times L}$$

denote the horizontally stacked matrix of coefficients, which forms a Hankel matrix of order 1 as defined in (2.24).

Given $V, \tilde{V} \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ admitting exact PCEs of L terms, cf. Definition 3.6, the norm of the difference $V - \tilde{V}$ can be obtained from the PCE coefficients as

$$\|V - \tilde{V}\|^2 = \left\|\mathsf{v}^{[0,L-1]} - \tilde{\mathsf{v}}^{[0,L-1]}\right\|_{\mathcal{D}}^2, \quad \mathcal{D} \doteq \operatorname{diag}\left(\{\|\phi^j(\xi)\|^2\}_{j=0}^{L-1}\right). \tag{3.6}$$

Similarly, with exact PCEs, the L^2 norm of random variables can be interpreted as a weighted norm of PCE coefficients in the Euclidean space $\mathbb{R}^{n_v L}$.

Moreover, with exact PCEs, the expectation $\mathbb{E}[V] \in \mathbb{R}^{n_v}$ and the covariance $\Sigma[V, \tilde{V}] \in \mathbb{R}^{n_v \times n_v}$ can be calculated as

$$\mathbb{E}[V] = \mathsf{v}^0, \quad \Sigma[V, \tilde{V}] = \sum_{j=1}^{L-1} \mathsf{v}^j \tilde{\mathsf{v}}^{j\top} \|\phi^j(\xi)\|^2 = \mathcal{H}_1\left(\mathsf{v}^{[0,L-1]}\right) \mathcal{D} \mathcal{H}_1\left(\mathsf{v}^{[0,L-1]}\right)^\top. \tag{3.7a}$$

For all $n \in \mathbb{N}^+$, the n-th central moments of the i-th component $V_i \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$ are given by

$$\mu_n(V_i) \doteq \mathbb{E}[(V_i - \mathbb{E}[V_i])^n] = \sum_{j_1=1}^{L-1} \sum_{j_2=1}^{L-1} \cdots \sum_{j_n=1}^{L-1} \mathsf{v}_i^{j_1} \mathsf{v}_i^{j_2} \cdots \mathsf{v}_i^{j_n} \mathbb{E}\left[\phi^{j_1}(\xi)\phi^{j_2}(\xi)\cdots\phi^{j_n}(\xi)\right], \tag{3.7b}$$

where the terms $\mathbb{E}\left[\phi^{j_1}(\xi)\phi^{j_2}(\xi)\cdots\phi^{j_n}(\xi)\right]$ can be computed offline. Note that for $V_i \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$, its n-th central moments with $n > 2$ are not necessarily finite. For an in-depth discussion and computationally efficient methods, see Lefebvre, 2020.

Given an L^2 random variable with known distribution, the key to construct an exact finite-dimensional PCE is the appropriate choice of basis functions. Indeed, random variables that follow some widely used distributions admit exact finite-dimensional PCEs in suitable polynomial bases as shown in Table 3.1. Notice that one uses specific random-variable arguments $\xi_i \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$ for different polynomial basis.

On the other hand, for random variables not listed in Table 3.1 we can construct a finite dimensional basis utilizing the first two moments of random variables. As a motivating example, note that any arbitrary scalar L^2 random variable $M \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$

admits an affine expansion in terms of its mean and standard definition as follows

$$M = \mathbb{E}[M] + \sqrt{\mu_2[M]}\xi, \quad \xi \doteq \frac{M - \mu[M]}{\sqrt{\mu_2[M]}}.$$

This representation is an exact PCE with a polynomial of degree 1, where the normalized variable ξ serves as the stochastic germ. This result can be extended from scalar to vector-valued random variables.

Lemma 3.2 (Exact PCEs with given mean and covariance): Consider $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ with mean $\mathbb{E}[V] \in \mathbb{R}^{n_v}$ and covariance $\Sigma[V] \succeq 0$. For any $\mathsf{M} \in \mathbb{R}^{n_v \times n_\xi}$ satisfying $\mathsf{M}\mathsf{M}^\top = \Sigma[V]$, there exists $\xi \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\xi})$ such that

$$V = \mathbb{E}[V] + \mathsf{M}\xi, \quad \mathbb{E}[\xi] = 0 \text{ and } \Sigma[\xi] = I_{n_\xi}.$$

Consider the basis

$$\{\phi^j(\xi)\}_{j=0}^{n_\xi} = \mathcal{P}_1(\xi) \doteq \left\{1, \{\xi_i\}_{i=1}^{n_\xi}\right\} \tag{3.8}$$

comprising polynomials $\mathcal{P}_1(\xi)$ with degree of at most 1. An exact and finite-dimensional PCE of V is obtained as

$$V = \mathbb{E}[V] + \mathsf{M}\xi = \sum_{j=0}^{n_\xi} \mathsf{v}^j \phi^j(\xi). \tag{3.9}$$

with $\mathsf{v}^0 = \mathbb{E}[V]$, $\mathcal{H}_1\left(\mathsf{v}^{[1,n_v]}\right) = \mathsf{M}$.

Proof. Since $\mathsf{M} \in \mathbb{R}^{n_v \times n_\xi}$ is in general not invertible, we utilize Moore-Penrose inverses to construct ξ. We note that the Moore-Penrose inverse $\mathsf{M}^\dagger$ of $\mathsf{M} \in \mathbb{R}^{n_v \times n_\xi}$ is unique and satisfies the following properties (Penrose, 1955),

$$\mathsf{M}\mathsf{M}^\dagger\mathsf{M} = \mathsf{M}, \mathsf{M}^\dagger\mathsf{M}\mathsf{M}^\dagger = \mathsf{M}^\dagger \tag{3.10a}$$

$$\mathsf{M}\mathsf{M}^\dagger = (\mathsf{M}\mathsf{M}^\dagger)^\top, \mathsf{M}^\dagger\mathsf{M} = (\mathsf{M}^\dagger\mathsf{M})^\top. \tag{3.10b}$$

Let

$$\xi \doteq \mathsf{M}^\dagger(V - \mathbb{E}[V]) + (I_{n_\xi} - \mathsf{M}^\dagger\mathsf{M})\xi_{\mathcal{N}} \tag{3.11}$$

with $\xi_{\mathcal{N}}$ as an arbitrary random variable in $L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\xi})$ that is independent of V. Moreover, we consider $\mathbb{E}[\xi_{\mathcal{N}}] = 0$ and $\Sigma[\xi_{\mathcal{N}}] = I_{n_\xi}$. Specifically, if M is of full column rank, then (3.11) recovers $\xi = \mathsf{M}^\dagger(V - \mathbb{E}[V])$ since in this case $\mathsf{M}^\dagger\mathsf{M} = I_{n_\xi}$.

Next, we prove that ξ from (3.11) satisfies the required conditions: $\mathbb{E}[\xi] = 0$, $\Sigma[\xi] = I_{n_\xi}$, and $V = \mathbb{E}[V] + \mathsf{M}\xi$. First, we have $\mathbb{E}[\xi] = \mathsf{M}^\dagger(\mathbb{E}[V] - \mathbb{E}[V]) + (I_{n_\xi} - \mathsf{M}^\dagger\mathsf{M})\mathbb{E}[\xi_{\mathcal{N}}] = 0$ since $\mathbb{E}[\xi_{\mathcal{N}}] = 0$ by construction. Furthermore,

$$\begin{aligned}
\Sigma[\xi] &= \mathsf{M}^\dagger\Sigma[V]\mathsf{M}^{\dagger\top} + (I_{n_\xi} - \mathsf{M}^\dagger\mathsf{M})\Sigma[\xi_{\mathcal{N}}](I_{n_\xi} - \mathsf{M}^\dagger\mathsf{M})^\top \\
&= \mathsf{M}^\dagger(\mathsf{M}\mathsf{M}^\top)\mathsf{M}^{\dagger\top} + (I_{n_\xi} - \mathsf{M}^\dagger\mathsf{M})(I_{n_\xi} - \mathsf{M}^\dagger\mathsf{M})^\top \\
&= \mathsf{M}^\dagger\mathsf{M}(\mathsf{M}^\dagger\mathsf{M})^\top + I_{n_\xi} - \mathsf{M}^\dagger\mathsf{M} - (\mathsf{M}^\dagger\mathsf{M})^\top + \mathsf{M}^\dagger\mathsf{M}(\mathsf{M}^\dagger\mathsf{M})^\top \overset{(3.10b)}{=} I_{n_\xi}.
\end{aligned}$$

We note that the random-variable equality $V = \mathbb{E}[V] + \mathsf{M}\xi$ is considered as $V(\omega) = \mathbb{E}[V] + \mathsf{M}\xi(\omega)$ for μ-almost all $\omega \in \Omega$. Let $\Delta \doteq V - \mathbb{E}[V] - \mathsf{M}\xi$. This equality is equivalent to $\Delta(\omega) = 0$ for μ-almost all $\omega \in \Omega$, which is true iff $\|\Delta\|^2 = \mathbb{E}[\Delta^\top\Delta] = 0$ as shown in (3.4).

With ξ from (3.11), we have

$$\begin{aligned}\Delta &= V - \mathbb{E}[V] - \mathsf{M}\left(\mathsf{M}^\dagger(V - \mathbb{E}[V]) + (I_{n_\xi} - \mathsf{M}^\dagger\mathsf{M})\xi_{\mathcal{N}}\right)\\ &= (I_{n_v} - \mathsf{M}\mathsf{M}^\dagger)(V - \mathbb{E}[V]) - (\mathsf{M} - \mathsf{M}\mathsf{M}^\dagger\mathsf{M})\xi_{\mathcal{N}} \overset{(3.10a)}{=} (I_{n_v} - \mathsf{M}\mathsf{M}^\dagger)(V - \mathbb{E}[V]).\end{aligned}$$

Hence, we have

$$\begin{aligned}\mathbb{E}[\Delta^\top\Delta] &= \operatorname{trace}\left((I_{n_v} - \mathsf{M}\mathsf{M}^\dagger)^\top(I_{n_v} - \mathsf{M}\mathsf{M}^\dagger)\Sigma[V]\right)\\ &= \operatorname{trace}\left((I_{n_v} - \mathsf{M}\mathsf{M}^\dagger)^\top(I_{n_v} - \mathsf{M}\mathsf{M}^\dagger)(\mathsf{M}\mathsf{M}^\top)\right)\\ &= \operatorname{trace}\left((I_{n_v} - \mathsf{M}\mathsf{M}^\dagger)^\top\left(\mathsf{M}\mathsf{M}^\top - \mathsf{M}\mathsf{M}^\dagger\mathsf{M}\mathsf{M}^\top\right)\right) = 0,\end{aligned}$$

where the last equality holds since $\mathsf{M}(\mathsf{M}^\dagger\mathsf{M}\mathsf{M}^\top) \overset{(3.10a)}{=} \mathsf{M}\mathsf{M}^\dagger$.

Combining the above results, we conclude that the random variable ξ given by (3.11) satisfies the conditions stated in Lemma 3.2. □

Typical choices of M include the principal square root or the Cholesky decomposition of $\Sigma[V]$. To span the whole $L^2(\Omega, \sigma(\xi), \mu; \mathbb{R}^{n_v})$, one may construct an orthogonal basis $\{\phi^j(\xi)\}_{j\in\mathbb{N}}$ with higher-degree polynomials using the Gram-Schmidt procedure by Witteveen et al., 2006. For its efficient implementation, we refer to Oladyshkin et al., 2012; Paulson et al., 2017, which are also known as *arbitrary polynomial chaos expansion*. However, as we will show next, the finite-dimensional basis $\{\phi^j(\xi)\}_{j=0}^{n_v-1}$ with polynomials of degree at most 1 are already rich enough to exactly represent all random variables satisfying $\tilde{V} \doteq AV + b$.

Galerkin Projection

Next, we first revisit Galerkin projection for the general case with $\tilde{V} \doteq f(V)$.

Given the map $f : L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v}) \to L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_{\tilde{v}}})$ and the PCE of V with respect to the orthogonal polynomials $\{\phi^j(\xi)\}_{j\in\mathbb{N}}$, Galerkin projection can be used to obtain the PCE coefficients of the image random variable $\tilde{V} = f(V)$, cf. Ghanem et al., 2003; Mühlpfordt, 2020. It consists of the following steps:

i) Substitute V and $\tilde{V}$ with their PCEs $\displaystyle\sum_{j=0}^{\infty}\tilde{\mathsf{v}}^j\phi^j(\xi) = f\left(\sum_{j=0}^{\infty}\mathsf{v}^j\phi^j(\xi)\right)$.

ii) For all $i \in \mathbb{N}$, project onto the basis ϕ^i, i.e.,

$$\left\langle\sum_{j=0}^{\infty}\tilde{\mathsf{v}}^j\phi^j(\xi), \phi^i(\xi)\right\rangle = \left\langle f\left(\sum_{j=0}^{\infty}\mathsf{v}^j\phi^j(\xi)\right), \phi^i(\xi)\right\rangle.$$

iii) Define $\tilde{\mathsf{v}}^i \doteq \dfrac{\left\langle f\left(\sum_{j=0}^{\infty}\mathsf{v}^j\phi^j(\xi)\right), \phi^i(\xi)\right\rangle}{\|\phi^i(\xi)\|^2}$.

In case of an affine mapping $\tilde{V} \doteq AV + b$ with $A \in \mathbb{R}^{n_A\times n_v}$ and $b \in \mathbb{R}^{n_v}$, following from the affinity of f and the orthogonality of $\{\phi^j\}_{j\in\mathbb{N}}$, we have

$$\tilde{\mathsf{v}}^i = \frac{\left\langle\sum_{j=0}^{\infty}A\mathsf{v}^j\phi^j(\xi) + b, \phi^i(\xi)\right\rangle}{\|\phi^i(\xi)\|^2} = \frac{\left\langle A\sum_{j=0}^{\infty}\mathsf{v}^j\phi^j(\xi) + b, \phi^i(\xi)\right\rangle}{\|\phi^i(\xi)\|^2} = A\mathsf{v}^i + b\delta^{0i}, \tag{3.12}$$

where δ^{0i} denotes the Kronecker delta that $\delta^{0i} = 0$ for $i \neq 0$. Note that due to the linear nature of PCEs, under affine maps, the i-th order PCE coefficient $\tilde{\mathsf{v}}^i$ is only related to v^i, and not to all other v^j for $j \neq i$. However, for nonlinear mappings, as indicated in step iii) of Galerkin projection, the calculation of $\tilde{\mathsf{v}}^i$ involves all v^j for $j \in \mathbb{N}$. This finding helps us to establish the following results.

Lemma 3.3 (Exact PCEs under affine mappings): Given $\{\phi^j(\xi)\}_{j\in\mathbb{N}}$ and $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$. If V admits an exact PCE with L terms in $\{\phi^j(\xi)\}_{j\in\mathbb{N}}$, then for arbitrary $A \in \mathbb{R}^{n_{\tilde{v}}\times n_v}$ and $b \in \mathbb{R}^{n_{\tilde{v}}}$, $\tilde{V} \doteq AV + b$ admits an exact PCE with at largest L terms in $\{\phi^j(\xi)\}_{j\in\mathbb{N}}$.

Proof. As indicated in (3.12), all j-th order PCE coefficients $\tilde{\mathsf{v}}^j$ with $j \geq L$ are zero since $\mathsf{v}^j = 0$ for $j \geq L$ as per Definition 3.6. Consequently, according to Definition 3.6 the assertion holds. □

Based on the results of Lemmas 3.2 and 3.3, we have the following corollary.

Corollary 3.1: Given $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ and the finite-dimensional PCE basis $\{\phi^j(\xi)\}_{j=0}^{n_\xi}$ constructed as in Lemma 3.2, then for arbitrary $A \in \mathbb{R}^{n_{\tilde{v}}\times n_v}$ and $b \in \mathbb{R}^{n_{\tilde{v}}}$, $\tilde{V} = AV + b$ admits an exact PCE in $\{\phi^j(\xi)\}_{j=0}^{n_\xi}$ as defined in (3.8).

This result plays an essential role by ensuring that affine transformations of random variables maintain the finite-dimensional PCE structure, which simplifies the analysis of propagating stochastic uncertainties through LTI dynamics. We will elaborate on this point in Chapter 5.

3.3. Stochastic Processes and their PCEs

Another key concept in stochastic control is the stochastic process. By collecting random variables over time, we can characterize the underlying dynamics utilizing stochastic processes. This topic will be explored in more detail in Chapter 4, where we analyze the behavior of stochastic LTI systems.

Instead of describing individual and joint distributions for each time instant, we follow the structure of random variables by considering stochastic processes as the results of a single random experiment modeled by the probability space $(\Omega, \mathcal{F}, \mu)$.

Definition 3.7 (Stochastic process (Sullivan, 2015)): A function $S : \mathbb{Z} \times \Omega \to \mathbb{R}^{n_s}$ is a discrete-time stochastic process if $S(k, \cdot)$ is a $\mathbb{R}^{n_s}$-valued random variable on $(\Omega, \mathcal{F}, \mu)$ while $S(\cdot, \omega)$ is a *realization trajectory* or a *path* in $(\mathbb{R}^{n_s})^{\mathbb{Z}}$.

Henceforth, for compactness, we consider $S(k)$, $S(\omega)$, and $S(k, \omega)$ as the k-th random variable, the realization, and the realization of k-th random variable of S, respectively. As shorthand, we write $S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$, that is, $S(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})$ holds for all $k \in \mathbb{Z}$.

Remark 3.3 (Stochastic process as $(\mathbb{R}^{n_s})^{\mathbb{Z}}$-valued random variables):
The stochastic process $S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$ can be equivalently represented as a $(\mathbb{R}^{n_s})^{\mathbb{Z}}$-valued random variables $S \in \big(\Omega, \mathcal{F}, \mu; (\mathbb{R}^{n_s})^{\mathbb{Z}}\big)$. Note that although each random variable $S(k)$ lies in $L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})$, it does not implies that S lies in $L^2(\Omega, \mathcal{F}, \mu; (\mathbb{R}^{n_s})^{\mathbb{Z}})$, since the variance of S may be infinite. However, we can relax the L^2 requirement and instead consider $S \in \big(\Omega, \mathcal{F}, \mu; (\mathbb{R}^{n_s})^{\mathbb{Z}}\big)$.

PCEs of Stochastic Processes

Given a stochastic process $S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$, let the collection of stochastic germs satisfy $\sigma(S) \subseteq \sigma(\xi) \subseteq \mathcal{F}$. We derive the polynomial chaos expansion of S. Consider an orthogonal polynomial basis $\{\phi^j(\xi)\}_{j\in\mathbb{N}}$ that spans $L^2(\Omega, \sigma(\xi), \mu; \mathbb{R}^{n_s})$, each random variable $S(k)$ in the process for $k \in \mathbb{Z}$ admits a polynomial chaos expansion. That is,

$$S(k) = \sum_{j\in\mathbb{N}} \mathsf{s}^j(k)\phi^j(\xi). \tag{3.13a}$$

For the sake of compactness, we write

$$\mathsf{s} = \left(\left(\mathsf{s}^j(k)\right)_{j\in\mathbb{N}}\right)_{k\in\mathbb{Z}} \simeq \begin{array}{c} \text{expansion index} \rightarrow \\ \text{time} \downarrow \end{array} \begin{bmatrix} \vdots & \vdots & \cdots & \vdots & \\ \mathsf{s}^0(0) & \mathsf{s}^1(0) & \cdots & \mathsf{s}^i(0) & \cdots \\ \mathsf{s}^0(1) & \mathsf{s}^1(1) & \cdots & \mathsf{s}^i(1) & \cdots \\ \vdots & \vdots & \ddots & \vdots & \\ \mathsf{s}^0(k) & \mathsf{s}^1(k) & \cdots & \mathsf{s}^i(k) & \cdots \\ \vdots & \vdots & & \vdots & \ddots \end{bmatrix} \in (\ell^2(\mathbb{R}^{n_s}))^{\mathbb{Z}}, \tag{3.13b}$$

$$\mathsf{s}(k) = (\mathsf{s}^0(k), \mathsf{s}^1(k), \cdots) = \left(\mathsf{s}^j(k)\right)_{j\in\mathbb{N}} \in \ell^2(\mathbb{R}^{n_s}), \tag{3.13c}$$

$$\mathsf{s}^j = (\cdots, \mathsf{s}^j(-1), \mathsf{s}^j(0), \mathsf{s}^j(1), \cdots) = \left(\mathsf{s}^j(k)\right)_{k\in\mathbb{Z}} \in (\mathbb{R}^{n_s})^{\mathbb{Z}}, \tag{3.13d}$$

where we organize the PCE coefficients of $S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$ into an infinite-dimensional matrix $\mathsf{s} \in (\ell^2(\mathbb{R}^{n_s}))^{\mathbb{Z}}$, where each row $\mathsf{s}(k) \in \ell^2(\mathbb{R}^{n_s})$ corresponds to the PCE coefficients of $S(k)$, and each column $\mathsf{s}^j \in (\mathbb{R}^{n_s})^{\mathbb{Z}}$ represents the trajectory of the j-th PCE coefficients over time.

Filtrations

For a stochastic process, filtrations are widely used to represent the information available at a given time instant. To this end, we denote the restriction of S at any subset $\mathbb{T} \subseteq \mathbb{Z}$ as $S_{\mathbb{T}} \doteq (S(k))_{k\in\mathbb{T}}$. For $i \in \mathbb{T}$, we denote the set of all elements of $\mathbb{T}$ that are no larger than i as $\mathbb{T}_{\leq i}$.

Definition 3.8 (Filtration):
Consider the probability space $(\Omega, \mathcal{F}, \mu)$ and a stochastic process $S : \mathbb{Z} \times \Omega \to \mathbb{R}^{n_s}$.

i) A *filtration* of the sigma algebra $\mathcal{F}$ is a family $(\mathcal{F}_i)_{i\in\mathbb{T}}$ of sub-σ-algebras of $\mathcal{F}$, such that $\mathcal{F}_k \subseteq \mathcal{F}_t \subseteq \mathcal{F}$ for all $k, t \in \mathbb{T}$ and $k \leq t$.

ii) The *natural filtration* of the stochastic process S is the filtration $(\mathcal{F}_i^S)_{i\in\mathbb{T}}$ defined by

$$\mathcal{F}_i^S \doteq \sigma\left((S(k))_{k\in\mathbb{T}_{\leq i}}\right)$$

as the smallest σ-algebra covering $(S(k))_{k\in\mathbb{T}_{\leq i}}$.

iii) The restriction $S_{\mathbb{T}}$ is *adapted* to a filtration $(\mathcal{F}_i)_{i\in\mathbb{T}}$ if $\mathcal{F}_i^S \subseteq \mathcal{F}_i$ for each $i \in \mathbb{T}$.

In other words, filtrations are families of σ-algebras that are ordered non-decreasingly. Specifically, the natural filtration $(\mathcal{F}_i^S)_{i\in\mathbb{T}}$ collects the σ-algebra $\mathcal{F}_i^S$ that represents all historical information of S available at time instants $k \in \mathbb{T}_{\leq i}$. For more details on filtrations, see Fristedt et al., 2013.

3.4. Uncertainty Quantification via Probability Bounds

In this section, we aim to quantify the uncertainty of random variables using probability bounds. For a given confidence level $0 < \gamma < 1$ and a real-valued random variable $M \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$, our goal is to determine a confidence interval $I_\gamma \subseteq \mathbb{R}$ such that

$$\mathbb{P}\left[M \in I_\gamma\right] \geq \gamma.$$

With the PCE of M, we can recover its distribution by sampling the stochastic germs. However, it can be computationally expensive. Alternatively, we can efficiently obtain the moments of the distribution with PCEs as shown in (3.7). In the following, we begin by constructing one-sided and two-sided probability bounds based on the first two moments, and then proceed to refine these intervals by incorporating additional information.

Lemma 3.4 (Chebyshev's inequality for two-sided bounds (Chebyshev, 1867)): Consider a real-valued random variable $M \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$ with its mean $\mathbb{E}[M]$ and variance $\mu_2(M)$, then for a confidence level $0 < \gamma < 1$, the following probability inequality holds

$$\mathbb{P}\left[\mathbb{E}[M] - \alpha(\gamma)\sqrt{\mu_2(M)} \leq M \leq \mathbb{E}[M] + \alpha(\gamma)\sqrt{\mu_2(M)}\right] \geq \gamma, \tag{3.14a}$$

where the concentration factor $\alpha(\gamma)$ reads

$$\alpha(\gamma) = \begin{cases} \sqrt{\frac{1}{1-\gamma}}, & M \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}), \\ F_{\mathcal{N}}^{-1}(\frac{1+\gamma}{2}), & M \sim \mathcal{N}\left(\mathbb{E}[M], \mu_2(M)\right), \end{cases} \tag{3.14b}$$

where $F_{\mathcal{N}}^{-1}$ refers to the inverse cdf of a standard Gaussian distribution.

Lemma 3.5 (Cantelli's inequality for one-sided bounds (Cantelli, 1928)): Consider a real-valued random variable $M \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$ with its mean $\mathbb{E}[M]$ and variance $\mu_2(M)$, then for a confidence level $0.5 < \gamma < 1$, the following probability inequality holds

$$\mathbb{P}\left[M \leq \mathbb{E}[M] + \alpha(\gamma)\sqrt{\mu_2(M)}\right] \geq \gamma, \tag{3.15a}$$

where the concentration factor $\alpha(\gamma)$ reads

$$\alpha(\gamma) = \begin{cases} \sqrt{\frac{\gamma}{1-\gamma}}, & M \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}), \\ F_{\mathcal{N}}^{-1}(\gamma), & M \sim \mathcal{N}(\mathbb{E}[M], \mu_2(M)). \end{cases} \tag{3.15b}$$

For general L^2 random variables, Cantelli's inequality provides a tight one-sided probability bound, while Chebyshev's inequality determines a tight two-sided probability bound. Specifically, for one-sided bounds, Cantelli's inequality is an improvement over Chebyshev's inequality. Consider an arbitrary $M \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$, Chebyshev's inequality yields

$$\mathbb{P}\left[M \leq \mathbb{E}[M] + \sqrt{\tfrac{1}{1-\gamma}}\sqrt{\mu_2(M)}\right] \geq \mathbb{P}\left[|M - \mathbb{E}[M]| \leq \sqrt{\tfrac{1}{1-\gamma}}\sqrt{\mu_2(M)}\right] \geq \gamma$$

with the factor $\sqrt{\frac{1}{1-\gamma}}$ larger than the one $\sqrt{\frac{\gamma}{1-\gamma}}$ from Cantelli's inequality.

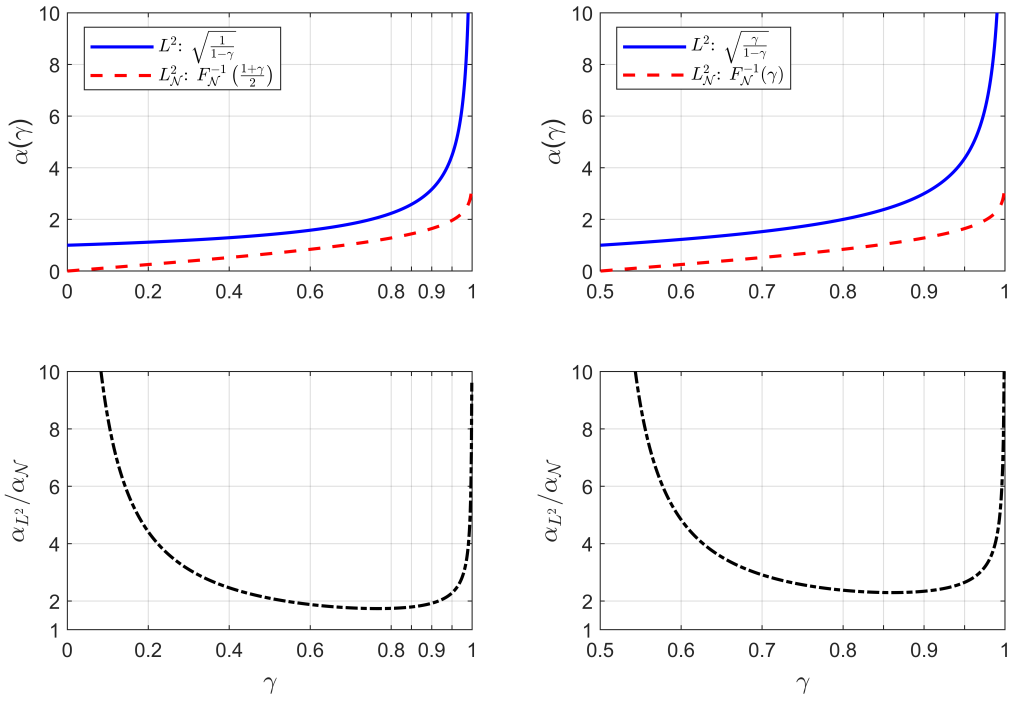

Figure 3.1.: Comparison of concentration factors in (3.14)–(3.15)

.

On the other hand, for two-sided confidence intervals, Cantelli's inequality gives

$$\begin{aligned}
&\mathbb{P}\left[|M-\mathbb{E}[M]| \leq \sqrt{\tfrac{\gamma}{1-\gamma}}\sqrt{\mu_2(M)}\right] \\
&= \mathbb{P}\left[M \leq \mathbb{E}[M]\sqrt{\tfrac{\gamma}{1-\gamma}}\sqrt{\mu_2(M)}\right] + \mathbb{P}\left[M \geq \mathbb{E}[M] - \sqrt{\tfrac{\gamma}{1-\gamma}}\sqrt{\mu_2(M)}\right] - 1 \geq 2\gamma - 1 \\
&\overset{\gamma=\frac{1+\gamma'}{2}}{\Longrightarrow} \mathbb{P}\left[|M-\mathbb{E}[M]| \leq \sqrt{\tfrac{1+\gamma'}{1-\gamma'}}\sqrt{\mu_2(M)}\right] \geq \gamma'
\end{aligned}$$

with the factor $\sqrt{\frac{1+\gamma'}{1-\gamma'}}$ larger than the one $\sqrt{\frac{1}{1-\gamma'}}$ from Chebyshev's inequality.

Given the differences between one-sided and two-sided bounds, it is important to understand when to apply which result. Two-sided bounds, like Chebyshev's inequality, are more appropriate for constructing confidence intervals in stochastic predictions. In contrast, one-sided bounds, such as Cantelli's inequality, are better suitable for reformulating individual chance constraints.

In stochastic prediction, the process typically involves specifying a target probability and then determining the corresponding confidence interval centred around the mean. Conversely, in chance constraints, the constraint is given first, and the goal is to ensure that the probability of satisfying the constraint exceeds a given threshold.

Moreover, chance constraints typically involve intervals that are not centered around the mean. For example, in the box constraint $|M| \leq m^{\mathrm{u}}$, the interval is not mean-centered, making the use of two-sided results less appropriate. Though the constraint is two-sided, violations may only occur from one side at a time, either exceeding m^{u} or dropping below $-m^{\mathrm{u}}$, but not both simultaneously. Hence, it is often more effective to decompose the constraint into two individual one-sided constraints, suggesting for the use of Cantelli's inequality.

Cantelli and Chebyshev bounds are typically loose because they rely only on the first two moments. Figure 3.1 compares the concentration factors $\alpha(\gamma)$ of both arbitrary L^2 random variables and Gaussian random variables, considering both one-sided and two-sided bounds. We observe that the concentration factor for arbitrary L^2 random variables is approximately twice as large as that for Gaussian random variables; moreover, it grows significantly faster as γ approaches 1. This indicates that the bounds for general L^2 are notably more conservative, particularly at high confidence levels.

However, assuming a non-Gaussian distribution to be Gaussian can introduce biased errors, as the actual distribution may exhibit skewness, heavier tails, or multimodality. In such cases, the Gaussian concentration factor may be overly optimistic, underestimating the probability of extreme events. Thus, while Gaussian-based bounds provide tighter intervals, they may fail to accurately capture tail probabilities in non-Gaussian settings.

To bridge the gap between the probabilistic bounds of Gaussian and arbitrary L^2 random variables, we consider the following two approaches:

i) considering higher-order moments,

ii) and using additional shape information about the distribution.

In particular, we employ approach i) to refine Chebyshev's inequality for two-sided confidence intervals and use approach ii) to improve Cantelli's inequality for one-sided chance constraints. Furthermore, we demonstrate that both approaches can be effectively implemented within the PCE framework.

3.4.1. Confidence Intervals via Higher-Order Chebyshev's Inequality

Note that the original Chebyshev's inequality typically provides rather loose bounds, since it applies to all random variables with finite first two moments. However, with the information of higher-order moments, Buot, 2006 notes that by applying Markov's inequality to the non-negative variable $(M - \mathbb{E}[M])^{2n}$, one can get a family of tail bounds.

Lemma 3.6 (Chebyshev's inequalities with higher-order moments): Consider a real-valued random variable $M \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$. For $n \in \mathbb{N}^+$, suppose its the $2n$-th central moment of M, denoted by $\mu_{2n}(M) \doteq \mathbb{E}\left[(M - \mathbb{E}[M])^{2n}\right]$, is finite. Then, for any confidence level $0 < \gamma < 1$, the following two-sided probabilistic inequality holds

$$\mathbb{P}\left[|M - \mathbb{E}[M]| \leq \sqrt[2n]{\frac{\mu_{2n}(M)}{1-\gamma}}\right] \geq \gamma. \tag{3.16}$$

Proof. For $a \in \mathbb{R}^+$ and $n \in \mathbb{N}^+$, we have

$$\mathbb{P}\left[|M - \mathbb{E}[M]| \leq a\right] = \mathbb{P}\left[(M - \mathbb{E}[M])^{2n} \leq a^{2n}\right] \geq 1 - \frac{\mu_{2n}(M)}{a^{2n}},$$

where the second inequality holds from the Markov inequality. By considering $\gamma \doteq 1 - \frac{\mu_{2n}(M)}{a^{2n}}$, we obtain (3.16). With $n = 1$, this result reduces to the original Chebyshev's inequality. □

Let the introduced confidence interval of (3.16) be

$$I_{2n,\gamma} \doteq [\mathbb{E}[M] - r_{2n,\gamma}, \mathbb{E}[M] + r_{2n,\gamma}], \quad \text{with radius} \quad r_{2n,\gamma} \doteq \sqrt[2n]{\frac{\mu_{2n}(M)}{1-\gamma}}.$$

We observe that this interval $I_{2n,\gamma}$ is tighter than the interval $I_{2,\gamma}$ induced by the original Chebyshev's inequality if the following condition holds

$$\frac{\mu_{2n}(M)}{(\mu_2(M))^n} \leq \left(\frac{1}{1-\gamma}\right)^{n-1}. \tag{3.17}$$

For $n = 2$ and $\gamma = 90\%$, this condition requires the kurtosis $\kappa = \frac{\mu_4}{\mu_2^2}$ to be less than 10. The kurtosis measures how much the shape of a distribution deviates from a normal distribution, which has $\kappa = 3$. Distributions with $\kappa \geq 10$ typically have a much heavier tails relative to a normal distribution. In other words, if condition (3.17) is violated, this indicates that the underlying distribution has a relative higher probability of extreme values, and we should use the interval $I_{2,\gamma}$ for conservatism. In an extreme case for $\kappa = \infty$, such as a t-distribution with 3 degrees of freedom, only the interval $I_{2,\gamma}$ can be used. Conversely, if the distribution of M has light tails and satisfy condition (3.17), we can opt for $I_{2n,\gamma}$ for a tighter confidence interval.

Note that the higher-order moments μ_{2n} can be directly calculated by (3.7) from the PCE coefficients of M and the expected values of the basis products. Furthermore, (3.7) can be simplified by assuming that each basis function is independently distributed, which enables the use of the property $\mathbb{E}\left[(\phi^i)^{l_1}(\phi^i)^{l_2}\right] = \mathbb{E}\left[(\phi^i)^{l_1}\right]\mathbb{E}\left[(\phi^i)^{l_2}\right]$. Suppose that M admits an exact PCE in a basis $\{\phi^j\}_{j=0}^{L-1}$ with $\phi^0 = 1$, and its terms $\{\phi^j\}_{j=1}^{L-1}$ are orthogonal and independently distributed. The PCE reformulation of its fourth-order central moment is

$$\begin{aligned}\mu_4(M) &= \sum_{j=1}^{L-1}(\mathsf{m}^j)^4\mathbb{E}\left[(\phi^j)^4\right] + \sum_{i=1}^{L-1}\sum_{j=i+1}^{L-1} 6(\mathsf{m}^i)^2(\mathsf{m}^j)^2\mathbb{E}\left[(\phi^j)^2\right]\mathbb{E}\left[(\phi^i)^2\right] \\ &\qquad + \underbrace{\sum_{i=1}^{L-1}\sum_{j\neq i} 4(\mathsf{m}^i)(\mathsf{m}^j)^3\mathbb{E}\left[\phi^i\right]\mathbb{E}\left[(\phi^j)^3\right]}_{=0,\quad \mathbb{E}[\phi^i]=0} \\ &= \sum_{j=1}^{L-1}(\mathsf{m}^j)^4\mathbb{E}\left[(\phi^j)^4\right] + \sum_{i=1}^{L-1}\sum_{j=i+1}^{L-1} 6(\mathsf{m}^i)^2(\mathsf{m}^j)^2\mathbb{E}\left[(\phi^j)^2\right]\mathbb{E}\left[(\phi^i)^2\right].\end{aligned} \tag{3.18}$$

Note that the $\mathbb{E}[\phi^i] = 0$ for all $i \in \mathbb{I}_{[0,L-1]}$ due to the orthogonality with $\phi^0 = \{1\}$ i.e. $\mathbb{E}[\phi^i \cdot 1] = 0$.

Example 3.3: Consider a scalar random variable $M = m + \sum_{j=1}^4 \xi_i$ where $\{\xi_j\}_{j=1}^4$ are i.i.d. non-Gaussian scalar random variables. Suppose that $\mathbb{E}[\xi_j] = 0$ and $\mu_2(\xi_j) = 1$ for $j = 1, 2, 3, 4$, then $\{1, \xi_1, \xi_2, \xi_3, \xi_4\}$ is a valid orthogonal PCE basis.

Suppose the fourth-order moment of ξ_j is given as $\mu_4(\xi_j) = 10$ for $j = 1, 2, 3, 4$. Then, by (3.18), we obtain

$$\mu_4(M) = \textstyle\sum_{j=1}^4 \mu_4(\xi_j) + \sum_{i=1}^4 \sum_{j=i+1}^4 6 = 76.$$

For a confidence level $\gamma = 0.95$, the radius of the confidence interval determined by the higher-order and the original Chebyshev inequalities are, respectively,

$$r_{4,0.95} = \sqrt[4]{\mu_4(M)/0.05} \approx 6.24, \quad r_{2,0.95} = \sqrt{\mu_2(M)/0.05} = \sqrt{20 \left\| \begin{bmatrix} 1 & 1 & 1 & 1 \end{bmatrix} \right\|_2^2} \approx 8.94,$$

which illustrates the tightening of the confidence interval by including higher-order moments.

3.4.2. Chance Constraints via Refined Cantelli's Inequality

When extending to higher-order moments, the result for two-sided bounds in Lemma 3.6 does not directly apply to the one-sided case. Instead, determining a tight one-sided bound using higher-order moments can be formulated as a semidefinite program, as shown by Bertsimas et al., 2005. While He et al., 2010 provide an explicit formulation for one-sided tail bounds involving fourth-order moments, it is not necessarily tight. Therefore, in this section, we turn to the other approach that incorporates additional information about the distribution's shape.

The appeal of Gaussian random variables lies in two key properties:

i) they are uniquely characterized by their first two moments,

ii) and the set of all Gaussian random variables is closed under summations of its two independent elements and linear operations on its elements.

These properties make Gaussian random variables particularly useful for uncertainty propagation under linear transformations. As alternatives, we can relax the Gaussian assumption step by step while retaining the second property: Gaussian (Definition 3.4), symmetric linear unimodal (Definitions 3.9–3.10), symmetric (Definition 3.9), and arbitrary L^2 random variables.

Definition 3.9 (RV with symmetric distribution): The distribution μ_V of a random variable $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ is symmetric if there exists $\mu \in \mathbb{R}^{n_v}$ such that for all $\tau \in \mathbb{R}^{n_v}$, $\mu_V(V \leq \mu - \tau) = \mu_V(V \geq \mu + \tau)$, where "$\geq$" and "$\leq$" hold elementwise.

We denote the set of all symmetric $\mathbb{R}^{n_v}$-valued L^2 random variables as $L^2_S(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$.

Definition 3.10 (RV with linear unimodal distribution): The distribution μ_V of a random variable $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ is linear unimodal if for all $a \in \mathbb{R}^{n_v}$ there exists $\bar{v} \in \mathbb{R}^{n_v}$ such that the cumulative distribution function $F_{a^\top V}$ of $a^\top V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$ is convex on $(-\infty, a^\top \bar{v}]$ and concave on $[a^\top \bar{v}, +\infty)$.

Similarly, the set of all symmetric linear unimodal distributed random variables in $L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ is represented by $L^2_{SU}(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$, and the set $L^2_{\mathcal{N}}(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ consists of all Gaussian distributed $\mathbb{R}^{n_v}$-valued L^2 random variables.

Example 3.4: To illustrate the differences among these three classes of distributions, as in Figure 3.2, we choose

- Gaussian distribution: A standard normal distribution $V_1 \sim \mathcal{N}(0, 1)$. This distribution is symmetric and unimodal, with a single peak at the mean.

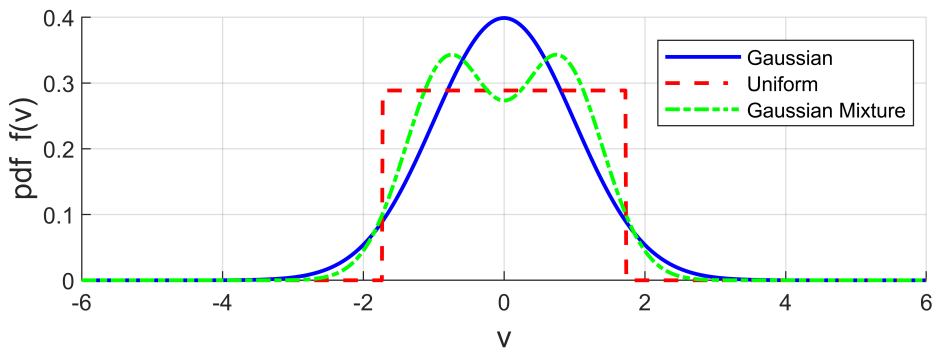

Figure 3.2.: Illustration of a Gaussian distribution, a uniform distribution, and a symmetric Gaussian mixture model

- Uniform distribution: A symmetric unimodal distribution, $V_2 \sim \mathcal{U}([-\sqrt{3}, \sqrt{3}])$, which has a constant density over the interval $[-\sqrt{3}, \sqrt{3}]$.
- Gaussian Mixture Model (GMM): A symmetric but non-unimodal distribution created by a mixture of two Gaussian distributions centered at -2 and 2 with equal weights. Specifically, $V_3 \sim 0.5 \cdot \mathcal{N}(-0.8, 0.36) + 0.5 \cdot \mathcal{N}(0.8, 0.36)$.

Note that all these classes of random variables retain property ii).

Lemma 3.7 (Closure under independent sums and linear operations (Nguyen et al., 2021; Yu et al., 2009)): For any set $L^2_{\mathcal{A}} \in \{L^2, L^2_S, L^2_{SU}, L^2_{\mathcal{N}}\}$ the following properties hold:

i) For any $n_v, n_w \in \mathbb{N}^+$, any $V \in L^2_{\mathcal{A}}(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$, and any arbitrary affine mapping with $A \in \mathbb{R}^{n_w \times n_v}$ and $b \in \mathbb{R}^{n_w}$, we have $W \doteq AV + b \in L^2_{\mathcal{A}}(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_w})$.

ii) For any independent $V_1, V_2 \in L^2_{\mathcal{A}}(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$, we have $V_1 + V_2 \in L^2_{\mathcal{A}}(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$.

Based on the seminal work by Yu et al., 2009, due to this closure property, the concentration factor $\alpha(\gamma)$ for Cantelli's inequality can be refined for L^2_S, L^2_{SU}. Specifically, for confidence level $0.5 < \gamma < 1$, we have

$$\mathbb{P}\left[M \leq \mathbb{E}[M] + \alpha(\gamma)\sqrt{\mu_2(M)}\right] \geq \gamma,\ \alpha(\gamma) = \begin{cases} \sqrt{\frac{\gamma}{1-\gamma}}, & M \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}), \\ \sqrt{\frac{1}{2(1-\gamma)}}, & M \in L^2_S(\Omega, \mathcal{F}, \mu; \mathbb{R}), \\ \frac{2}{3\sqrt{2(1-\gamma)}}, & M \in L^2_{SU}(\Omega, \mathcal{F}, \mu; \mathbb{R}), \\ F^{-1}_{\mathcal{N}}(\gamma), & M \in L^2_N(\Omega, \mathcal{F}, \mu; \mathbb{R}). \end{cases} \tag{3.19}$$

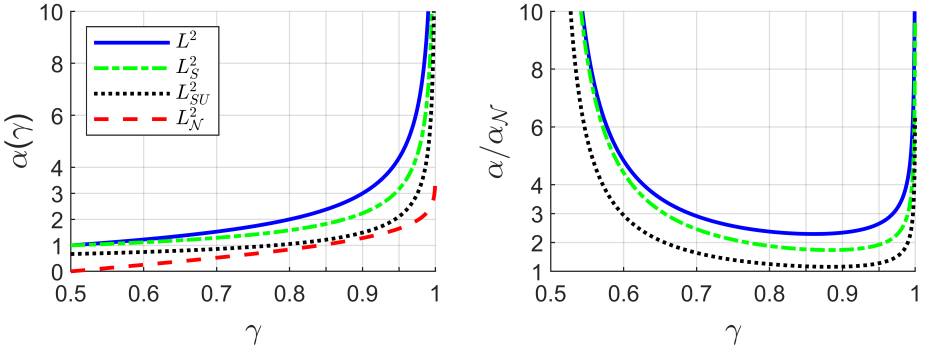

Figure 3.3.: Comparison of concentration factors in (3.19)

As illustrated in Figure 3.3, this refinement bridges the gap between the concentration factors for arbitrary L^2 random variables and Gaussians. Specifically, for symmetric unimodal distributions (black-dotted line), the concentration factor $\alpha(\gamma)$ closely aligns with that of Gaussian distributions (red-dashed line) for confidence level $0.8 \leq \gamma \leq 0.95$. For symmetric distributions, take $\gamma = 0.9$ for instance, the concentration factor improves from 3 (for general L^2 random variables) to approximately 2.24, which provides a tighter bound.

PCE Reformulation of Chance Constraints with Cantelli Ambiguity

The refined Cantelli's inequality (3.19) provides an uncertainty quantification approach that leverages both the first two moments and the shape information of the distribution. To integrate these aspects, we define the Cantelli ambiguity sets as follows.

Definition 3.11 (Cantelli ambiguity sets): For any $L^2_{\mathcal{A}} \in \{L^2, L^2_S, L^2_{SU}, L^2_{\mathcal{N}}\}$, its corresponding Cantelli ambiguity set is defined as the set of all random variables in $L^2_{\mathcal{A}}(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ that share the same mean $m \in \mathbb{R}^{n_v}$ and covariance $\Gamma \in \mathbb{R}^{n_v \times n_v}$, i.e.,

$$\mathcal{C}_{\mathcal{A}}(m, \Gamma) \doteq \left\{ V \in L^2_{\mathcal{A}}(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v}) \,\middle|\, \mathbb{E}[V] = m,\ \Sigma[V] = \Gamma \right\}. \tag{3.20}$$

Next, we demonstrate that the closure of independent sums in Lemma 3.7 enables an exact PCE formulation of chance constraints with Cantelli ambiguity sets.

Consider the PCE basis

$$\{1, \xi_1, \xi_2, \cdots, \xi_{L-1}\}, \quad \xi = \left[\xi_1, \xi_2, \cdots, \xi_{L-1}\right]^\top \in \mathcal{C}_{\mathcal{A}}\left(0, I_{L-1}\right). \tag{3.21}$$

That is, we consider a basis of polynomial degree of 1 with stochastic germs $\{\xi_j\}_{j=1}^{L-1}$ that are independent to each other (though not necessarily identical). Moreover, since all stochastic germs belongs to the same subclass, the closure of independent sums in Lemma 3.7 implies that any random variable admitting an exact PCE in the basis (3.21) also belongs to the same subclass.

Corollary 3.2 (Closure of independent sum in exact PCEs): If $V \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v}\right)$ admits an exact PCE in (3.21)

$$V = \mathsf{v}^0 + \mathcal{H}_1\left(\mathsf{v}^{[1,L-1]}\right)\xi, \quad \xi \in \mathcal{C}_{\mathcal{A}}(0, I_{L-1}), \tag{3.22}$$

then $V \in \mathcal{C}_{\mathcal{A}}\left(\mathsf{m}^0, \mathcal{H}_1\left(\mathsf{v}^{[1,L-1]}\right)\mathcal{H}_1\left(\mathsf{v}^{[1,L-1]}\right)^\top\right)$.

This result allows us to reformulate chance constraints with Cantelli's ambiguity sets using PCE.

Corollary 3.3 (PCE for chance constraints (Calafiore et al., 2006)): Consider a random variable $V \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v}\right)$ with its PCE given by (3.22). Then, a user-defined chance constraint with $a \in \mathbb{R}^{n_v}$

$$\mathbb{P}[a^\top V \leq 1] \geq \gamma$$

is satisfied if the following second-order cone condition holds

$$a^\top \mathsf{v}^0 + \alpha(\gamma)\left\|a^\top \mathcal{H}_1\left(\mathsf{v}^{[1,L-1]}\right)\right\| \leq 1, \tag{3.23}$$

where the concentration factor $\alpha(\gamma)$ is chosen from (3.19).

Proof. This result holds by directly combining Lemma 3.7 and the refined Cantelli's inequality (3.19). □

Observe that (3.23) does not explicitly dependent on the basis (3.21) and thus the knowledge of the exact distribution of $\xi \in \mathcal{C}_{\mathcal{A}}(0, I_{L-1})$ is not required.

3.5. Summary

In this chapter, we provided a foundational overview of random variables (Section 3.1), polynomial chaos expansion (Section 3.2), and stochastic processes 3.3. We explored the role of stochastic processes and their PCE representations, which are essential for describing the stochastic behaviors of LTI systems in Chapter 4. Additionally, we introduced exact PCEs and filtration techniques, which are instrumental for the propagation of uncertainties in stochastic systems, as will be further explored in Chapter 5.

For uncertainty quantification, we examined two key probability bounds:

- Higher-order Chebyshev bounds (Section 3.4.1), which provide two-sided bounds using higher-order moments. These bounds will be applied in Chapter 5 to construct confidence intervals for stochastic predictions.
- Refined Cantelli bounds (Section 3.4.2), which offer one-sided probability bounds, refined for general symmetric or symmetric unimodal distributions. They form the basis for chance constraints in stochastic optimal control in Chapter 6.

These two bounds extend the original Chebyshev and Cantelli inequalities, bridging the gap between two extreme assumptions: treating random variables as either arbitrary L^2 or Gaussian distributed. Combined with the PCE framework, they provide an efficient approach for uncertainty quantification of non-Gaussian random variables.

The hierarchical relationship between these concepts and their role in subsequent chapters is illustrated in Figure 3.4.

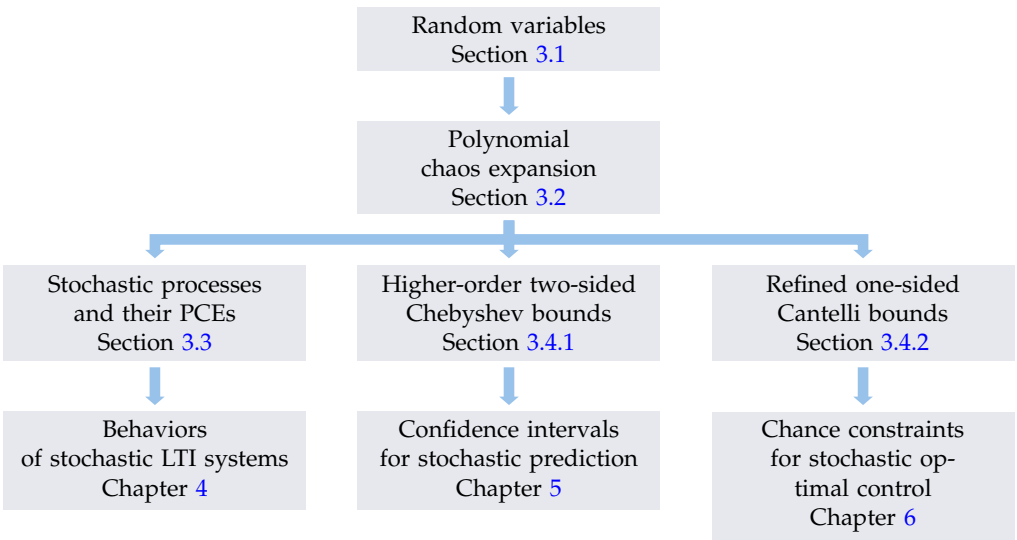

Figure 3.4.: Overview of this chapter and the link to the rest of the thesis

4. Representations of Stochastic System Behaviors

This chapter addresses the challenge of representing LTI systems under stochastic uncertainties using a behavioral approach. The cornerstone of our approach lies in extending the behavioral setting, traditionally used for deterministic LTI systems, to encompass stochastic behaviors. This extension is facilitated by considering L^2 random variables, PCEs, realizations, and statistical moments as tools to characterize stochastic behaviors.

Furthermore, we find that, subject to LTI dynamics, stochastic behaviors can be characterized purely by realization data, without the need to estimate statistical information of past trajectories. These results build upon our previous work in Faulwasser et al., 2023; Pan et al., 2023c, extending it to provide rigorous definitions of LTI stochastic behaviors in terms of L^2 random variables and their first two moments.

4.1. Uncertainty Propagation of Exogenous Disturbances

In control systems, not all inputs to a dynamical system can be directly controlled by the user. While *manipulated inputs* $u : \mathbb{Z} \to \mathbb{R}^{n_u}$ are chosen by the user, the system is also subject to *exogenous disturbances* $w : \mathbb{Z} \to \mathbb{R}^{n_w}$ that are determined by the environment. In building control, for instance, ambient temperature, electricity prices, and user thermal demand are examples of exogenous disturbances (Özmeteler et al., 2024).

Exogenous disturbances could also be considered as input components of the system, provided that they satisfy the conditions for i/o partitioning, cf. Definition 2.7. Specifically, we make the following assumptions on the disturbance signal $w \in (\mathbb{R}^{n_w})^{\mathbb{Z}}$

i) w can be chosen freely by the environment from $(\mathbb{R}^{n_w})^{\mathbb{Z}}$ and is not influenced by the system's dynamics,

ii) w is not anticipated by the output y, and

iii) to uniquely determine y, we require the knowledge of w.

Under these assumptions, the deterministic behavior $\mathfrak{B}$ can be partitioned into manipulated inputs u, exogenous disturbances w, and outputs y, such that $s = (u, w, y)$, with the manifest input $\tilde{u} \doteq (u, w)$ and $n_s = n_u + n_w + n_y$. A state-space representation of $\mathfrak{B}$ subject to exogenous disturbances reads

$$x(k+1) = Ax(k) + Bu(k) + Ew(k), \quad k \in \mathbb{Z} \tag{4.1a}$$

$$y(k) = Cx(k) + Du(k) + Fw(k), \tag{4.1b}$$

where u and w are treated symmetrically as input signals.

While deterministic control theory focuses on exact trajectories of inputs and outputs, it falls short in practical scenarios where the exact values of exogenous disturbances w are unpredictable. In such cases, we do not have access to the exact future realizations of w. Instead, we often have statistical knowledge of w as a stochastic process $W \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_w})$. Underlying this, the probability space $(\Omega, \mathcal{F}, \mu)$ in fact captures the *uncertain environment*. That is, for an outcome $\omega \in \Omega$ representing one possible choice of the environment, w can be seen as the corresponding realization of W determined by the environment, i.e., $w \doteq W(\omega)$.

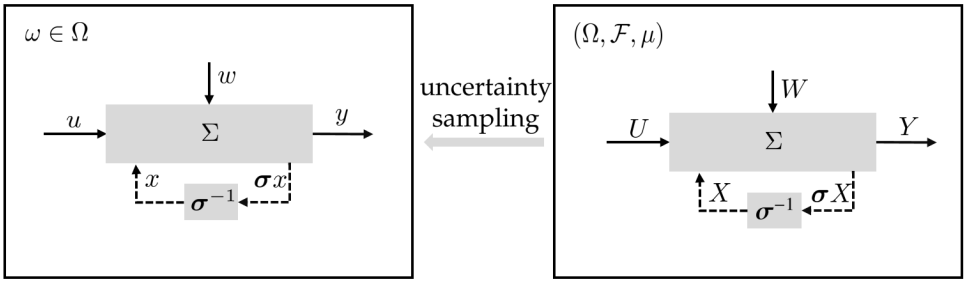

Figure 4.1.: Conceptual visualization of the deterministic behavior and the stochastic behavior of the dynamical system Σ

Moreover, exogenous disturbances introduce uncertainties into the system. In such cases, deterministic behavior $\mathfrak{B}$ is insufficient to fully characterize the system behavior. Instead, we have to introduce stochastic behavior to account for the propagation of uncertainty introduced by the disturbances to all other system components including

- the state x, which evolves based on w and u,
- the input u, which is usually determined by feedback policies dependent on x and w,
- and the output y, which reflects the cumulative effect of stochastic u and w.

This results in stochastic versions of all components:

$$X \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_x}\right)^{\mathbb{Z}}, \quad U \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_u}\right)^{\mathbb{Z}}, \quad Y \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y}\right)^{\mathbb{Z}}.$$

The shift in focus—from deterministic realizations to stochastic processes—motivates the introduction of stochastic behavior. By substituting all deterministic signals in (4.1) with their stochastic counterparts, we obtain

$$X(k+1) = AX(k) + BU(k) + EW(k), \quad k \in \mathbb{Z} \tag{4.2a}$$
$$Y(k) = CX(k) + DU(k) + FW(k). \tag{4.2b}$$

Note that while the system matrices $(A, [B, E], C, [D, F])$ remain unchanged, the propagation of uncertainties renders all components stochastic. As shown in Figure 4.1, the trajectories (u, w, y, x) of the deterministic behavior in (4.1) are a realization of the stochastic processes (U, W, Y, X) in (4.2), associated with an outcome $\omega \in \Omega$ sampled from the probability space $(\Omega, \mathcal{F}, \mu)$.

In our previous work (Faulwasser et al., 2023; Pan et al., 2023c), the stochastic behavior of LTI dynamical systems is defined using (4.2), which is essentially a state-space representation. In this chapter, we rigorously define the stochastic behavior of dynamical systems by incorporating uncertainty into the deterministic framework of the environment through the formalism of a probability space.

4.2. Stochastic LTI Behavior

The propagation of uncertainty highlights the necessity of a stochastic behavioral framework. Deterministic behavior $\mathfrak{B}$ encodes the dynamics of the system for given and

fixed trajectories of inputs and disturbances. However, to account for the inherent uncertainty introduced by exogenous disturbances, we extend this framework to stochastic behavior, which captures all stochastic processes $S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$ that are compatible with the underlying deterministic dynamics. This leads to the definition of stochastic behavior as a natural extension of deterministic behavior that incorporates probabilistic uncertainty.

Definition 4.1 (Stochastic behavior): Consider a dynamical system $\Sigma = (\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{B})$, its stochastic behavior $\mathfrak{S}$ with respect to the probability space $(\Omega, \mathcal{F}, \mu)$ is defined as the set of all stochastic processes $S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$ whose realizations $S(\omega)$ are elements of $\mathfrak{B}$ almost surely, i.e.,

$$\mathfrak{S} \doteq \left\{ S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}} \,\middle|\, S(\omega) \in \mathfrak{B} \subset (\mathbb{R}^{n_s})^{\mathbb{Z}} \text{ for } \mu\text{-almost all } \omega \in \Omega \right\}. \tag{4.3}$$

In this case, we call $\mathfrak{B}$ the *realization behavior* of $\mathfrak{S}$.

The key insight of this definition is that both deterministic behavior $\mathfrak{B}$ and stochastic behavior $\mathfrak{S}$ are governed by the same dynamical system Σ. The distinction lies in their respective perspectives: the deterministic behavior focuses on *a posteriori* observations of realized trajectories, while the stochastic behavior provides an *a priori* description of the system, incorporating the probabilistic nature of the environment as defined by the probability space $(\Omega, \mathcal{F}, \mu)$.

The following theorem reveals that the LTI property of a stochastic behavior actually stems from its deterministic counterpart.

Theorem 4.1 (LTI stochastic behavior): The stochastic behavior $\mathfrak{S} \subset L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$ of an LTI dynamical system $\Sigma = (\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{B})$ (Definition 2.1) is LTI, i.e., $\mathfrak{S}$ satisfies

i) $\mathfrak{S}$ is a linear subspace of $L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$, i.e., for arbitrary $a, b \in \mathbb{R}$ and $S_1, S_2 \in \mathfrak{S}$, $aS_1 + bS_2 \in \mathfrak{S}$ holds.

ii) $\mathfrak{S}$ is shift invariant, i.e., $\boldsymbol{\sigma}\mathfrak{S} = \mathfrak{S}$.

iii) $\mathfrak{S}$ is complete, i.e., $S_{[t_1,t_2]} \in \mathfrak{S}_{[t_1,t_2]}$ holds for all $t_1, t_2 \in \mathbb{Z}$, $t_1 \leq t_2$, implies $S \in \mathfrak{S}$.

Proof. This proof illustrates how the linearity and shift invariance of $\mathfrak{S}$ can be inherited from $\mathfrak{B}$. First, we prove the linearity. Let $a, b \in \mathbb{R}$ and $S_1, S_2 \in \mathfrak{S}$. By definition of $\mathfrak{S}$ in Definition 4.1, we have $S_1(\omega), S_2(\omega) \in \mathfrak{B}$ for μ-almost all $\omega \in \Omega$. Then, by the linearity of $\mathfrak{B}$, $(aS_1(\omega) + bS_2(\omega)) \in \mathfrak{B}$ for μ-almost all $\omega \in \Omega$. Thus, it implies $(aS_1 + bS_2) \in \mathfrak{S}$ and thus the linearity of $\mathfrak{S}$.

To show the shift invariance of the stochastic behavior $\mathfrak{S}$, we prove that $\boldsymbol{\sigma}\mathfrak{S} = \mathfrak{S}$.

We begin by expanding the definition of the left-hand side:

$$\begin{aligned} \boldsymbol{\sigma}\mathfrak{S} &\doteq \{\boldsymbol{\sigma}S \mid S(\omega) \in \mathfrak{B} \text{ for } \mu\text{-almost all } \omega \in \Omega\} \\ &\overset{\tilde{S}=\boldsymbol{\sigma}S}{=} \left\{ \tilde{S} \,\middle|\, \boldsymbol{\sigma}^{-1}\tilde{S}(\omega) \in \mathfrak{B} \text{ for } \mu\text{-almost all } \omega \in \Omega \right\}. \end{aligned}$$

Next, we note that for $S(\omega), \tilde{S}(\omega) \in (\mathbb{R}^{n_s})^{\mathbb{Z}}$ with $\tilde{S}(\omega) = \boldsymbol{\sigma}S(\omega)$,

$$\boldsymbol{\sigma}^{-1}\tilde{S}(\omega) \in \mathfrak{B} \Leftrightarrow S(\omega) \in \mathfrak{B} \Leftrightarrow \tilde{S}(\omega) \in \boldsymbol{\sigma}\mathfrak{B}. \tag{4.4}$$

Hence,

$$\boldsymbol{\sigma}\mathfrak{S} = \left\{ \tilde{S} \,\middle|\, \tilde{S}(\omega) \in \boldsymbol{\sigma}\mathfrak{B} \text{ for } \mu\text{-almost all } \omega \in \Omega \right\}.$$

Now, using the shift invariance of $\mathfrak{B}$, i.e., $\boldsymbol{\sigma}\mathfrak{B} = \mathfrak{B}$ if follows that

$$\boldsymbol{\sigma}\mathfrak{S} = \{ S \,|\, S(\omega) \in \mathfrak{B} \text{ for } \mu\text{-almost all } \omega \in \Omega \} = \mathfrak{S}$$

Therefore, we conclude that $\boldsymbol{\sigma}\mathfrak{S} = \mathfrak{S}$.

For the completeness, consider S satisfying $S_{[t_1,t_2]} \in \mathfrak{S}_{[t_1,t_2]}$ for all $t_1, t_2 \in \mathbb{Z}$, t_1, t_2. This implies $S_{[t_1,t_2]}(\omega) \in \mathfrak{B}_{[t_1,t_2]}$ for μ-almost $\omega \in \Omega$. By completeness of $\mathfrak{B}$, we have $S(\omega) \in \mathfrak{B}$ for μ-almost $\omega \in \Omega$, which concludes the proof. □

Similar as in the deterministic case, shift invariance implies $\mathfrak{S}_{[t_1,t_2]} = \mathfrak{S}_{[0,t_2-t_1]}$ for all $t_1, t_2 \in \mathbb{Z}$, $t_1 \leq t_2$. Thus, with $T \in \mathbb{N}^+$, we define the finite-length stochastic behavior as

$$\mathfrak{S}_T \doteq \left\{ S_{[t_1,t_1+T-1]} \,\middle|\, S \in \mathfrak{S}, t_1 \in \mathbb{Z} \right\},$$

which matches any restricted behavior $\mathfrak{S}_{[t_1,t_2]}$ with $t_2 - t_1 = T - 1$.

4.2.1. Parametric Representations with Random Variables

Building on the parametric representations outlined in Section 2.2, we extend these concepts to LTI stochastic behaviors. Similarly, we define the kernel representation $\mathfrak{S}(R)$.

Lemma 4.1: Consider the stochastic behavior $\mathfrak{S} \subset L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$ of an LTI dynamical system $\Sigma = (\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{B})$. Suppose the deterministic behavior $\mathfrak{B}$ admits a kernel representation $\mathfrak{B}(R)$ with $R(z) \in \mathbb{R}^{n_R \times n_s}[z, z^{-1}]$, then the stochastic behavior $\mathfrak{S}$ admits a kernel representation

$$\mathfrak{S}(R) \doteq \left\{ S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}} \,\middle|\, R(\boldsymbol{\sigma})S = 0 \right\} \tag{4.5}$$

where the equality holds in L^2 sense as defined in (3.4).

Proof. As per Definition 4.1 and (4.3), the stochastic counterpart of $\mathfrak{B}(R)$ is expressed as

$$\mathfrak{S}(R) = \left\{ S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}} \,\middle|\, R(\boldsymbol{\sigma})S(\omega) = 0 \text{ for } \mu\text{-almost all } \omega \in \Omega \right\}.$$

By replacing the term "for μ-almost all $\omega \in \Omega$" by the equality of L^2 random variables as defined in (3.4), we obtain (4.5). □

Similarly, with $P(z) \in \mathbb{R}^{n_y \times n_y}[z, z^{-1}]$ nonsingular and $Q(z) \in \mathbb{R}^{n_u \times n_y}[z, z^{-1}]$ such that $P^{-1}Q$ proper, we then defnie the i/o representation of the stochastic behavior as

$$\mathfrak{S}(P, Q) \doteq \left\{ (\tilde{U}, Y) \,|\, P(\boldsymbol{\sigma})Y = Q(\boldsymbol{\sigma})\tilde{U} \right\}. \tag{4.6}$$

Furthermore, we define the ARX representation $\mathfrak{S}(\Xi, D)$ and the state-space representation $\mathfrak{S}(A, B, C, D)$ as

$$\mathfrak{S}(\Xi, D) \doteq \left\{ (\tilde{U}, Y) \,\middle|\, Y(k) = \Xi \begin{bmatrix} \tilde{U}_{[k-\ell,k-1]} \\ Y_{[k-\ell,k-1]} \end{bmatrix} + D\tilde{U}(k), \forall k \in \mathbb{Z} \right\}, \tag{4.7}$$

$$\mathfrak{S}(A,B,C,D) \doteq \left\{ (\tilde{U},Y) \;\middle|\; \begin{array}{c} \exists X \in L^2\left(\Omega,\mathcal{F},\mu;\mathbb{R}^{n_x}\right)^{\mathbb{Z}} \\ \boldsymbol{\sigma} X = AX + B\tilde{U},\; Y = CX + D\tilde{U} \end{array} \right\}. \tag{4.8}$$

Note that (4.2) in Section 4.1 is indeed a state-space representation $\mathfrak{S}(A,[B,D],C,[D,F])$. The latent behavior of (4.8) is defined as

$$\mathfrak{S}^L(A,B,C,D) \doteq \left\{ (\tilde{U},Y,X) \;\middle|\; \boldsymbol{\sigma} X = AX + B\tilde{U},\; Y = CX + D\tilde{U} \right\}.$$

We remark that the all the above parametric representations of the stochastic behavior $\mathfrak{S}$ can be established using the same proof technique as in Lemma 4.1, by applying it to the corresponding representations of the realization behavior $\mathfrak{B}$. Consequently, the stochastic input $\tilde{U}$, output Y, and state X are defined relatively to the underlying deterministic system realization as follows.

Definition 4.2 (Stochastic input, output, and state): Consider the stochastic behavior $\mathfrak{S} \subset L^2\left(\Omega,\mathcal{F},\mu;\mathbb{R}^{n_s}\right)^{\mathbb{Z}}$ of an LTI dynamical system $\Sigma = (\mathbb{Z},\mathbb{R}^{n_s},\mathfrak{B})$.

i) We say $(\tilde{U},Y) = \Pi^{-1}S$ is an input/output partitioning of $\mathfrak{S}$, if $(\tilde{u},y) = \Pi^{-1}s$ is an input/output partitioning of $\mathfrak{B}$ satisfying Definition 2.7.

ii) We say $X \in L^2\left(\Omega,\mathcal{F},\mu;\mathbb{R}^{n_x}\right)^{\mathbb{Z}}$ is a state trajectory of $\mathfrak{S}(A,B,C,D)$, if $X(\omega) \in (\mathbb{R}^{n_x})^{\mathbb{Z}}$ is a state trajectory of $\mathfrak{B}(A,B,C,D)$ for μ-almost all $\omega \in \Omega$.

We note that the notion of a stochastic state is introduced only after a state-space representation has been specified. In contrast, in the deterministic setting, the concept of state admits an intrinsic definition within behavioral system theory (cf. Definition VII.1 in Willems, 1991). Establishing an intrinsic definition of the stochastic state remains an open problem for future research.

Proposition 4.1 (i/o partitioning of the stochastic behavior): The input/output partitioning of $\mathfrak{S}$ satisfies the following properties:

i) $\tilde{U}$ is *free*, i.e., $\Pi_{\tilde{U}}\mathfrak{S} = L^2\left(\Omega,\mathcal{F},\mu;\mathbb{R}^{n_{\tilde{u}}}\right)^{\mathbb{Z}}$.

ii) the output Y is *not anticipating* the input $\tilde{U}$, i.e., given $(\tilde{U}_1,Y_1) \in \Pi^{-1}\mathfrak{S}$, then for any $\tilde{U}_2 \in \left(L^2\left(\Omega,\mathcal{F},\mu;\mathbb{R}^{n_{\tilde{u}}}\right)\right)^{\mathbb{Z}}$ such that $\tilde{U}_1(t) = \tilde{U}_2(t)$ for all $t \leq 0$, there exists $Y_2 \in \left(L^2\left(\Omega,\mathcal{F},\mu;\mathbb{R}^{n_y}\right)\right)^{\mathbb{Z}}$ such that $Y_1(t) = Y_2(t)$ for all $t \leq 0$ and $(\tilde{U}_2,Y_2) \in \Pi^{-1}\mathfrak{S}$.

iii) the output Y *processes* the input $\tilde{U}$, i.e., given $(\tilde{U},Y_1),(\tilde{U},Y_2) \in \Pi^{-1}\mathfrak{S}$ with the same input process, if $Y_1(t) = Y_2(t)$ for all $t < 0$, then $Y_1(t) = Y_2(t)$ for $t \geq 0$ as well.

iv) the number of inputs $n_{\tilde{u}}$ is maximal over all partitionings that satisfy properties i)–iii).

Proof. i): Given $(\tilde{u},y) = \Pi^{-1}s$ as an input/output partitioning of $\mathfrak{B}$, $\tilde{u}$ is a free variable that $\Pi_{\tilde{u}}\mathfrak{B} = (\mathbb{R}^{n_{\tilde{u}}})^{\mathbb{Z}}$ by Definition 2.7. Since $\mathfrak{S}$ collects all stochastic processes S such that $S(\omega) = s \in \mathfrak{B}$ almost surely, its input projection $\Pi_{\tilde{U}}\mathfrak{S}$ has to cover all processes whose realizations in $(\mathbb{R}^{n_{\tilde{u}}})^{\mathbb{Z}}$. Therefore, $\Pi_{\tilde{U}}\mathfrak{S} = L^2\left(\Omega,\mathcal{F},\mu;\mathbb{R}^{n_{\tilde{u}}}\right)^{\mathbb{Z}}$.

ii): Let $(\tilde{U}_1,Y_1) \in \Pi^{-1}\mathfrak{S}$, meaning there exists a stochastic process $S_1 \in \mathfrak{S}$ such that $\Pi S_1 = (\tilde{U}_1,Y_1)$. For μ-almost all ω, the sample path $S_1(\omega) \doteq (\tilde{U}_1(\omega),Y_1(\omega)) \in \mathfrak{B}$.

Now, consider any $\tilde{U}_2 \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_{\tilde{u}}}\right)^{\mathbb{Z}}$ such that $\tilde{U}_2(t) = \tilde{U}_1(t)$ for all $t \leq 0$. Define a new stochastic process $S_2 \doteq \Pi(\tilde{U}_2, Y_2)$ as follows: for μ-almost all ω, let $S_2(\omega) \in \mathfrak{B}$ be such that

$$\tilde{U}_2(t,\omega) = \tilde{U}_1(t,\omega) \text{ for all } t \leq 0, \quad \text{and} \quad Y_2(t,\omega) = Y_1(t,\omega) \text{ for all } t \leq 0,$$

which is possible by the non-anticipation property of $\mathfrak{B}$ (Definition 2.7).

Replacing "μ-almost all $\omega \in \Omega$" with L^2-equality as in (3.4), we obtain $(\tilde{U}_2, Y_2) = \Pi^{-1}\tilde{S} \in \Pi^{-1}\mathfrak{S}$ and $Y_2(t) = Y_1(t)$ for all $t \leq 0$, as required.

iii): Let $(\tilde{U}, Y_1), (\tilde{U}, Y_2) \in \Pi^{-1}\mathfrak{S}$. Then there exist $S_1, S_2 \in \mathfrak{S}$ such that $\Pi S_1 = (\tilde{U}, Y_1)$, $\Pi S_2 = (\tilde{U}, Y_2)$. For μ-almost all ω, the sample paths $S_1(\omega), S_2(\omega) \in \mathfrak{B}$ and satisfy

$$\tilde{U}_1(t,\omega) = \tilde{U}_2(t,\omega), \quad Y_1(t,\omega) = Y_2(t,\omega) \text{ for all } t < 0.$$

By the processing property of $\mathfrak{B}$, it follows that $Y_1(t,\omega) = Y_2(t,\omega)$ for all $t \in \mathbb{Z}$ and for μ-almost all ω. Therefore, $Y_1(t) = Y_2(t)$ for all $t \in \mathbb{Z}$.

iv): Let $n_{\tilde{u}}' > n_{\tilde{u}}$ and suppose that there exists an input/output partitioning of $\mathfrak{S}$ with $n_{\tilde{u}}'$ inputs satisfying i)–iii). Then this would induce a corresponding partitioning of $\mathfrak{B}$ with more than $n_{\tilde{u}}$ inputs, contradicting the maximality of the original partitioning of $\mathfrak{B}$ by Definition 2.7. Hence, $n_{\tilde{u}}$ is maximal.

□

Observe that a stochastic process is considered as an input component if it can be chosen freely from $L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_{\tilde{u}}}\right)^{\mathbb{Z}}$. As elaborated in Section 4.1, we further distinguish the input process $\tilde{U} \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_{\tilde{u}}}\right)^{\mathbb{Z}}$ into two components: the manipulated process $U \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_u}\right)^{\mathbb{Z}}$ and the disturbance process $W \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_w}\right)^{\mathbb{Z}}$. Both are freely chosen from their respective spaces; however, U is chosen by the users, while W is determined by the uncertain environment.

Another property of the i/o partitioning is that the outputs processes the inputs. That is, the future outputs are only dependent on previous inputs and initial conditions. In the context of stochastic process, it can be linked to the filtration. We have the following result that is a direct corollary of Lemma 2.12.

Corollary 4.1 (Filtration of LTI stochastic behavior): If $(\tilde{U}, Y)_{[-\ell,N-1]} \in \mathfrak{S}_{N+\ell}$, then $Y_{[0,N-1]}$ is adapted to the filtration $\{\mathcal{F}_k\}_{k \in \mathbb{I}_{[0,N-1]}}$ with $\mathcal{F}_k = \sigma\left(Y_{[-\ell,-1]}, \tilde{U}_{[-\ell,k]}\right)$, $k \in \mathbb{I}_{[0,N-1]}$.

Proof. By the definition of the stochastic behavior $\mathfrak{S}$, for μ-almost all $\omega \in \Omega$, $(\tilde{u}, y) \doteq (\tilde{U}(\omega), Y(\omega)) \in \mathfrak{B}$. From Lemma 2.12, we know that if the initial trajectory $(\tilde{u}, y)_{[-\ell,-1]} \in \mathfrak{B}_\ell$ is fixed, and for $k \in \mathbb{I}_{[0,N-1]}$ the input trajectory $\tilde{u}_{[0,k]}$ is specified, then the output trajectory $y_{[0,k]}$ is uniquely determined. That is, there exists some deterministic function ϕ_k such that

$$y_k = \phi(\tilde{u}_{[-\ell,-1]}, y_{[-\ell,-1]}, \tilde{u}_{[0,k]}).$$

Translating this to the stochastic setting: since $Y(\omega) = y$ and $\tilde{U}(\omega) = \tilde{u}$ for μ-almost all ω, it follows that each $Y(k)$ is $\mathcal{F}_k$-measurable, where

$$\mathcal{F}_k \doteq \sigma\left(Y_{[-\ell,-1]}, \tilde{U}_{[-\ell,k]}\right).$$

Hence, $Y_{[0,N-1]}$ is adapted to the filtration $\{\mathcal{F}_k\}_{k \in \mathbb{I}_{[0,N-1]}}$. □

The definition of stochastic behavior in (4.3) involves L^2 random variables, which are generally infinite-dimensional objects in a Hilbert space. To facilitate a practical analysis and representation, we further characterize stochastic behavior in terms of PCE coefficients and statistical moments, which are finite-dimensional.

4.2.2. The PCE Coefficient Behavior

In this section, we first apply PCE to all stochastic processes in $\mathfrak{S}$ following approaches in Section 3.3, then analyze the dynamics and the behavior of the resulting coefficients.

We begin by constructing a suitable polynomial basis that allows stochastic processes within $\mathfrak{S}$ to admit PCEs. Note that according to Proposition 4.1, all admissible input processes within $\mathfrak{S}$ span the entire L^2 space, denoted as $\Pi_{\tilde{U}}(\mathfrak{S}) = L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_{\tilde{u}}})$. Therefore, to construct PCEs for all stochastic processes $S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$ in $\mathfrak{S}$, we require a stochastic germ $\xi \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\xi})$, $n_\xi \in \mathbb{N}^+ \cup \{\infty\}$, satisfying $\sigma(\xi) = \mathcal{F}$. Consequently, there exists an orthogonal polynomial basis $\{\phi^j\}_{j\in\mathbb{N}}$ spanning $L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$. This way, all stochastic processes $S \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s})^{\mathbb{Z}}$ in $\mathfrak{S}$ admit PCEs in the basis $\{\phi^j\}_{j\in\mathbb{N}}$, given by $S = \sum_{j\in\mathbb{N}} \mathsf{s}^j \phi^j(\xi)$ as shown in (3.13). The PCE coefficients of a stochastic process are introduced in (3.13).

We proceed by applying PCEs to all stochastic processes in $\mathfrak{S}$, particularly focusing on its kernel representation $\mathfrak{S}(R)$. By substituting all random variables in $\mathfrak{S}(R)$ with their respective PCEs and performing Galerkin projection onto the basis functions ϕ^j, $j \in \mathbb{N}$, we derive the dynamics of the PCE coefficients

$$R(\boldsymbol{\sigma})\mathsf{s}^j = 0, \quad \forall j \in \mathbb{N}. \tag{4.9}$$

Due to the linearity of PCE, the PCE coefficients for all basis dimensions $j \in \mathbb{N}$ satisfy the same linear system equations as shown in (4.9), which aligns with the realization behavior $\mathfrak{B}(R) = \mathfrak{B}$. This leads us to define the behavior of the PCE coefficients as follows.

Definition 4.3 (PCE coefficient behavior): Consider the stochastic behavior $\mathfrak{S}$ of an LTI dynamical system $\Sigma = (\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{B})$. The PCE coefficients of all stochastic processes in $\mathfrak{S}$, with respect to any orthogonal polynomial basis $\{\phi^j\}_{j\in\mathbb{N}}$ spanning $L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$, are elements of the PCE coefficient behavior

$$\mathfrak{C} \doteq \left\{ \mathsf{s} \in (\ell^2(\mathbb{R}^{n_s}))^{\mathbb{Z}} \,\middle|\, \mathsf{s}^j \in \mathfrak{B},\ \forall j \in \mathbb{N} \right\}. \tag{4.10}$$

Observe that the PCE coefficient behavior $\mathfrak{C}$ and the realization behavior $\mathfrak{B}$ satisfy the inclusion

$$\mathfrak{C} \subset \prod_{j\in\mathbb{N}} \mathfrak{B}, \quad \text{if } \mathfrak{B} \text{ is LTI.} \tag{4.11}$$

where the product $\prod_{j\in\mathbb{N}} \mathfrak{B}$ denotes the set of all sequence $\mathsf{s}^j_{j\in\mathbb{N}}$ such that each $\mathsf{s}^j \in \mathfrak{B}$.

This inclusion arises from the fact that, for all dimensions $j \in \mathbb{N}$, the PCE coefficient trajectories are compatible with the LTI dynamics encoded in $\mathfrak{B}$. As shown in (4.9), this compatibility stems from the linearity of both the PCE representation of the random variable and the system dynamics. However, we remark that the inclusion (4.11) may not hold if the system dynamics are nonlinear.

The coefficient behavior $\mathfrak{C}$ inherits desirable properties like linearity, shift invariance, and completeness from the realization behavior $\mathfrak{B}$.

Lemma 4.2 (Linear time-invariant $\mathfrak{C}$): The PCE coefficient behavior $\mathfrak{C}$ in (4.10) is

i) linear, i.e., $\mathfrak{C}$ is a linear subspace of $(\ell^2(\mathbb{R}^{n_s}))^{\mathbb{Z}}$,

ii) shift-invariant, i.e., $\mathfrak{C} = \boldsymbol{\sigma}\mathfrak{C}$,

iii) and complete, i.e. $\mathsf{c}_{[t_1,t_2]} \in \mathfrak{C}_{[t_1,t_2]}$ for all $t_1, t_2 \in \mathbb{Z}$, $t_1 \leq t_2$ implies $\mathsf{c} \in \mathfrak{C}$.

Proof. i) Linearity: Consider $\mathsf{s}_1, \mathsf{s}_2 \in \mathfrak{C}$, we have $\mathsf{s}_1^j, \mathsf{s}_2^j \in \mathfrak{B}$ for all $j \in \mathbb{N}$. Therefore, $a\mathsf{s}_1^j + b\mathsf{s}_2^j$ remains within $\mathfrak{B}$ for any real numbers a, b due to the linearity of $\mathfrak{B}$. Moreover, the square-summability of sequences $(\mathsf{s}_1^j)_{j\in\mathbb{N}}$ and $(\mathsf{s}_2^j)_{j\in\mathbb{N}}$ guarantees that $(a\mathsf{s}_1^j + b\mathsf{s}_2^j)_{j\in\mathbb{N}}$ is also square-summable. Thus, we have $a\mathsf{s}_1 + b\mathsf{s}_2 \in \mathfrak{C}$ by (4.10).

ii) Shift-invariance: by the equivalence in (4.4) and the shift-invariance of $\mathfrak{B}$, we have

$$\begin{aligned}\boldsymbol{\sigma}\mathfrak{C} &\doteq \left\{\boldsymbol{\sigma}\mathsf{s} \,\middle|\, \mathsf{s}^j \in \mathfrak{B}\right\} = \left\{\tilde{\mathsf{s}} \,\middle|\, \boldsymbol{\sigma}^{-1}\tilde{\mathsf{s}}^j \in \mathfrak{B}\right\} \\ &= \left\{\tilde{\mathsf{s}} \,\middle|\, \tilde{\mathsf{s}}^j \in \boldsymbol{\sigma}\mathfrak{B}\right\} = \left\{\tilde{\mathsf{s}} \,\middle|\, \tilde{\mathsf{s}}^j \in \mathfrak{B}\right\} = \mathfrak{C}.\end{aligned}$$

iii) Completeness: Consider s satisfying $\mathsf{s}_{[t_1,t_2]} \in \mathfrak{C}_{[t_1,t_2]}$ for all $t_1, t_2 \in \mathbb{Z}$ and $t_1 \leq t_2$. It implies that $\mathsf{s}^j_{[t_1,t_2]} \in \mathfrak{B}_{[t_1,t-2]}$ by (4.10). Furthermore, by completeness of $\mathfrak{B}$, we have $\mathsf{s}^j \in \mathfrak{B}$ for all $j \in \mathbb{N}$, which further implies $\mathsf{s} \in \mathfrak{C}$ by (4.10). □

By shift invariance, we define the finite-length PCE coefficient behavior as

$$\mathfrak{C}_T \doteq \{\mathsf{s}_{[t_1,t_1+T-1]} \,|\, \mathsf{s} \in \mathfrak{C}, t_1 \in \mathbb{Z}\},$$

which matches any restricted behavior $\mathfrak{C}_{[t_1,t_2]}$ with $t_2 - t_1 = T - 1$.

4.2.3. The Behavior of the First Two Moments

An alternative way to describe the L^2 random-variable behavior is by focusing on its first two moments, namely the mean and covariance. This second-order characterization is especially relevant in cases where the underlying processes are Gaussian.

Consider the following definition of the behavior associated to the dynamics of the first two moments (in short *moment behavior*),

$$\mathfrak{M} \doteq \left\{ (m, \Gamma) \,\middle|\, \begin{array}{c} S \in \mathfrak{S}, \quad m \doteq \mathbb{E}[S] \in (\mathbb{R}^{n_s})^{\mathbb{Z}} \\ \Gamma : \mathbb{Z} \times \mathbb{Z} \to \mathbb{R}^{n_s \times n_s},\ \Gamma(k,t) \doteq \Sigma[S(k), S(t)],\ k, t \in \mathbb{Z} \end{array} \right\}. \tag{4.12}$$

where Γ is the *covariance function* corresponding to the stochastic process $S \in \mathfrak{S}$. The restriction of $\mathfrak{M}$ to interval $[t_1, t_2]$ is denoted as $\mathfrak{M}_{[t_1,t_2]} \doteq \{(m,\Gamma)_{[t_1,t_2]} | (m,\Gamma) \in \mathfrak{M}\}$ with restriction of covariance function as $\Gamma_{[t_1,t_2]} \doteq (\Gamma(k,t))_{k,t\in\mathbb{I}_{[t_1,t_2]}}$.

It is worth noting that $\Gamma_{[t_1,t_2]}$ can be interpreted both as a two-dimensional matrix-valued sequence in $\left(\mathbb{R}^{n_s\times n_s}\right)^{\mathbb{I}_{[t_1,t_2]}\times\mathbb{I}_{[t_1,t_2]}}$ and as a matrix in $\mathbb{R}^{n_s(t_2-t_1+1)\times n_s(t_2-t_1+1)}$. Moreover, due to the positive semi-definiteness of the covariance operator, we have

$$\Gamma_{[t_1,t_2]} = \Sigma[S_{[t_1,t_2]}, S_{[t_1,t_2]}] \succeq 0,$$

which implies that $\Gamma_{[t_1,t_2]}$ is a positive semi-definite matrix for all $t_1, t_2 \in \mathbb{Z}$.

Though $\mathfrak{M}$ is not a linear subspace, its shift-invariance and completeness are inherited from the stochastic behavior $\mathfrak{S}$.

Lemma 4.3 (Shift-invariance and completeness of $\mathfrak{M}$): The moment behavior $\mathfrak{M}$ is shift-invariant, i.e. $\mathfrak{M}_{[t_1,t_2]} = \mathfrak{M}_{[t_1',t_2']}$ for all $t_1, t_2, t_1', t_2' \in \mathbb{Z}$ and $t_2 - t_1 = t_2' - t_1'$. The moment behavior $\mathfrak{M}$ is complete, i.e., if $(m,\Gamma)_{[t_1,t_2]} \in \mathfrak{M}_{[t_1,t_2]}$ for all $t_1, t_2 \in \mathbb{Z}$, $t_1 \leq t_2$, then $(m,\Gamma) \in \mathfrak{M}$.

Proof. Due to the shift invariance of $\mathfrak{S}$, we have $\mathfrak{S}_{[t_1,t_2]} = \mathfrak{S}_{[t_1',t_2']}$ for $t_2 - t_1 = t_2' - t_1'$, thus their first two moments are also equivalent. Hence, $\mathfrak{M}_{[t_1,t_2]} = \mathfrak{M}_{[t_1',t_2']}$. For the completeness of the moment behavior, by (4.12), if $(m,\Gamma)_{[t_1,t_2]} \in \mathfrak{M}_{[t_1,t_2]}$ for all $t_1, t_2 \in \mathbb{Z}$, $t_1 \leq t_2$, then there is $S_{[t_1,t_2]} \in \mathfrak{S}_{[t_1,t_2]}$ for all $t_1, t_2 \in \mathbb{Z}$, $t_1 \leq t_2$. By completeness of $\mathfrak{S}$, we have $S \in \mathfrak{S}$, which implies its first two moments $(m,\Gamma) \in \mathfrak{M}$ by (4.12). □

By shift invariance, we define the finite-length moment behavior as

$$\mathfrak{M}_T \doteq \{(m,\Gamma)_{[0,T-1]} \,|\, (m,\Gamma) \in \mathfrak{M}\}. \tag{4.13}$$

In the following discussion, we present an alternative representation of the moment behavior (4.12) by linking it to the underlying realization behavior $\mathfrak{B}$.

Fixing one time argument $t \in \mathbb{Z}$ in the covariance function Γ, the mapping $\Gamma(\cdot, t)$ defines a sequence of covariance matrices $(\Sigma[S(k), S(t)])_{k\in\mathbb{Z}}$, each in $\mathbb{R}^{n_s \times n_s}$. Owing to the bilinearity and time-invariance of the covariance operator, we will show that $\Gamma(\cdot, t)$ inherits the LTI structure from $\mathfrak{B}$.

To formalize this, we introduce the *stacked behavior* $\mathfrak{B}^{n_s} \doteq \prod_{i=1}^{n_s} \mathfrak{B}$. We interpret $\Gamma(\cdot, t)$ as a matrix-valued trajectory whose columns lie in $\mathfrak{B}$, and hence $\Gamma(\cdot, t) \in \mathfrak{B}^{n_s}$. Moreover, by the symmetry of the covariance function, $\Gamma(k,t) = \Gamma(t,k)^\top$, it follows that the row-wise slices $\Gamma(k,\cdot)$ also lie in $\mathfrak{B}^{n_s}$.

Proposition 4.2 (Moment behavior of LTI systems): Given a stochastic LTI behavior $\mathfrak{S}$, its realization behavior $\mathfrak{B}$, and its moment behavior $\mathfrak{M}$ defined in (4.12). For a given horizon length $T \in \mathbb{N}^+$, the finite-length moment behavior $\mathfrak{M}_T$ can be equivalently represented as

$$\mathfrak{M}_T = \left\{ (m,\Gamma)_{[0,T-1]} \;\middle|\; \begin{array}{c} m_{[0,T-1]} \in \mathfrak{B},\ \Gamma : \mathbb{I}_{[0,T-1]} \times \mathbb{I}_{[0,T-1]} \to \mathbb{R}^{n_s \times n_s}, \\ \Gamma_{[0,T-1]} \succeq 0,\ (\Gamma(k,t))_{k\in\mathbb{I}_{[0,T-1]}} \in \mathfrak{B}_T^{n_s},\ \forall t \in \mathbb{I}_{[0,T-1]} \end{array} \right\}. \tag{4.14}$$

Proof. For the equivalence between (4.13) and (4.14), concerning the behavior of mean trajectories, with $\phi^0 \equiv 1$, it coincides with the behavior of the 0-th PCE coefficients. Hence, based on the property of the PCE coefficient behavior as in Definition 4.3, we conclude that $\mathbb{E}[S_{[0,T-1]}] \in \mathfrak{B}_T$ for $S_{[0,T-1]} \in \mathfrak{S}_T$.

Next, we investigate the behavior of the covariance function Γ. To this end, we first prove (4.12) $\Rightarrow$ (4.14). Define $\bar{S}_{[0,T-1]} \doteq S_{[0,T-1]} - \mathbb{E}[S_{[0,T-1]}]$. Given $S_{[0,T-1]} \in \mathfrak{S}_T$ and $\mathbb{E}[S_{[0,T-1]}] \in \mathfrak{B}_T$, it follows that $\bar{S}_{[0,T-1]} \in \mathfrak{S}_T$. Thus, considering the kernel representation $\bar{S}_T \in \mathfrak{S}_T(R)$, we obtain

$$R(\boldsymbol{\sigma})\bar{S} = 0 \Rightarrow \textstyle\sum_{i=0}^{\ell} R_i \bar{S}(k+i) = 0, \quad \forall k \in \mathbb{I}_{[0,T-1]}.$$

For all $k, t \in \mathbb{I}_{[0,T-1]}$, multiplying the above equation from the right by $\bar{S}(t)^\top$ and taking the mean over the whole equation, we have

$$\mathbb{E}\left[\textstyle\sum_{i=0}^{\ell} R_i \bar{S}(k+i)\bar{S}(t)^\top\right] = 0 \Rightarrow \textstyle\sum_{i=0}^{\ell} R_i \mathbb{E}\left[\bar{S}(k+i)\bar{S}(t)^\top\right] = 0$$

$$\Rightarrow \sum_{i=0}^{\ell} R_i \Gamma(k+i,t) = 0, \quad \forall k, t \in \mathbb{I}_{[0,T-1]}.$$

Compactly, the above implies $R(\boldsymbol{\sigma})\Gamma(\cdot,t) = 0, \forall t \in \mathbb{I}_{[0,T-1]}$. In other words, we have $\Gamma(\cdot,t) \in \mathfrak{B}^{n_s}$. Combining the positive semi-definiteness of the covariance matrix

$$\Gamma_{[0,T-1]} \doteq \begin{bmatrix} \Gamma(0,0) & \cdots & \Gamma(0,T-1) \\ \vdots & \ddots & \vdots \\ \Gamma(T-1,0) & \cdots & \Gamma(T-1,T-1) \end{bmatrix},$$

we establish (4.14). Moreover, it implies that $\Gamma_{[0,T-1]} \in (\mathfrak{B}_T^{n_s})^T$.

(4.14) $\Rightarrow$ (4.12): We aim to prove that for all $(m,\Gamma)_{[0,T-1]}$ satisfying (4.14), there is a stochastic process $S_{[0,T-1]} \in \mathfrak{S}_T$ such that $\mathbb{E}[S_{[0,T-1]}] = m_{[0,T-1]}$ and $\Sigma[S_{[0,T-1]}] = \Gamma_{[0,T-1]}$.

Applying Cholesky decomposition to $\Gamma_{[0,T-1]} \succeq 0$, we obtain a unique lower-triangular matrix L_T with $n_s T - n_T$ columns of zeros such that $\Gamma_{[0,T-1]} = L_T L_T^\top$, where n_T is the rank of $\Gamma_{[0,T-1]}$. Since each column of $\Gamma_{[0,T-1]}$ is an element of $\mathfrak{B}_T$, we have

$$\operatorname{colspan}(\Gamma_{[0,T-1]}) = \operatorname{colspan}(L_T) \subseteq \mathfrak{B}_T.$$

Thus, each column of L_T (including zeros) lies in $\mathfrak{B}_T$ as well, i.e., $L_T^j \in \mathfrak{B}_T, \quad j \in \mathbb{I}_{[1,n_s T]}$.

Let $\xi_{[0,T-1]} \in \mathbb{R}^{n_s T}$ be an independent distributed stochastic process with zero means and unit covariance, we can construct $S_{[0,T-1]} \doteq m_{[0,T-1]} + L_T \xi_{[0,T-1]}$. By considering $\{\phi^j\}_{j=0}^{n_s T} = P_1(\xi_{[0,T-1]})$ similar as defined in (3.8), we have

$$S_{[0,T-1]} = \sum_{j=0}^{n_s T} \mathsf{s}^j_{[0,T-1]} \phi^j(\xi), \quad \mathsf{s}^j_{[0,T-1]} = \begin{cases} m_{[0,T-1]}, j = 0 \\ L_T^j, j \in \mathbb{I}_{[1,n_s T]} \end{cases} \Rightarrow \mathsf{s}_{[0,T-1]} \in \mathfrak{C}_T.$$

Due to the behavioral equivalence that will be proven in Theorem 4.3, we have $S_{[0,T-1]} \in \mathfrak{S}_T$. □

For the relation between $\mathfrak{M}$ and $\mathfrak{B}$, by (4.14), we have

$$\mathfrak{M}_T \subset \mathfrak{B}_T \times (\mathfrak{B}_T^{n_s})^T, \tag{4.15}$$

where the relation holds with equality if positive semi-definiteness is imposed on elements of $(\mathfrak{B}_T^{n_s})^T$.

4.2.4. Behavioral Lift, Reconstruction, and Equivalence

The stochastic behavior $\mathfrak{S}$, the PCE coefficient behavior $\mathfrak{C}$, and the first two moment behavior $\mathfrak{M}$, as defined in (4.3), (4.10), and (4.14), are all linked to the realization behavior $\mathfrak{B}$. The following theorem by Faulwasser et al., 2023 further establishes the mappings among these behaviors.

Theorem 4.2 (Behavioral lift (Faulwasser et al., 2023)):

i) The linear map

$$\Phi : \mathfrak{C} \to \mathfrak{S}, \ \Phi(\mathsf{s}) \doteq \sum_{j \in \mathbb{N}} \phi^j(\xi) \mathsf{s}^j \tag{4.16a}$$

is bijective, with the inverse given by

$$\Phi^{-1} : \mathfrak{S} \to \mathfrak{C}, \ \Phi^{-1}(S) = \mathsf{s}, \quad \mathsf{s}^j(k) = \frac{\langle S(k), \phi^j(\xi) \rangle}{\|\phi^j(\xi)\|^2}, j \in \mathbb{N}, \ k \in \mathbb{Z}. \tag{4.16b}$$

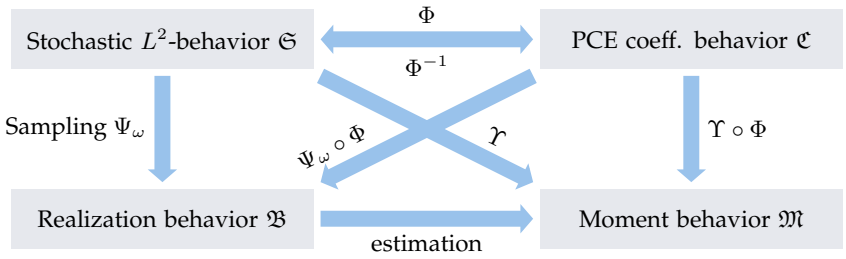

Figure 4.2.: Relations and maps between the behaviors $\mathfrak{B}, \mathfrak{C}, \mathfrak{S}$, and $\mathfrak{M}$ (Faulwasser et al., 2023)

ii) For fixed $\omega \in \Omega$, the linear map

$$\Psi_\omega : \mathfrak{S} \to \mathfrak{B}, \;\; \Psi_\omega(S) \doteq S(\omega)$$

is surjective. The concatenation

$$\Psi_\omega \circ \Phi : \mathfrak{C} \to \mathfrak{B}, \;\; (\Psi_\omega \circ \Phi)(\mathsf{s}) = \sum_{j \in \mathbb{N}} \phi^j(\omega) \mathsf{s}^j$$

is also surjective.

iii) The map

$$\Upsilon : \mathfrak{S} \to \mathfrak{M}, \;\; S \mapsto \Upsilon(S) \doteq (\mathbb{E}[S], \Sigma[S]) \,,$$

is surjective. The concatenation

$$\Upsilon \circ \Phi : \mathfrak{C} \to \mathfrak{M}, \;\; \Upsilon \circ \Phi(\mathsf{s}) = \left(\mathsf{s}^0, \textstyle\sum_{j=1}^{\infty} \mathsf{s}^j \mathsf{s}^{j\top} \right)$$

is also surjective.

See Faulwasser et al., 2023 for a detailed proof. The relations between the behaviors and the maps derived in Theorem 4.2 are sketched in Figure 4.2.

Conceptual Reconstruction and Behavioral Equivalence

As shown in Figure 4.2, the relations $\mathfrak{S} \Rightarrow \mathfrak{B}$, $\mathfrak{C} \Rightarrow \mathfrak{B}$, $\mathfrak{S} \Rightarrow \mathfrak{M}$, and $\mathfrak{C} \Rightarrow \mathfrak{M}$ are one-directional. That is, the mappings from left to right are surjective, which leads to information loss. In the following, we show that conceptual reconstructions from right to left only recover subsets of the original set on the left.

For the reconstruction from $\mathfrak{B}$ to $\mathfrak{S}$, we note that any deterministic trajectory $s \in \mathfrak{B}$ can be interpreted as a stochastic process that is almost surely constant at each time instant. Therefore, the stochastic behavior defined in (4.3) includes its realization behavior,

$$\mathfrak{B} \subset \mathfrak{S}. \tag{4.17}$$

Furthermore, the reconstruction from $\mathfrak{B}$ to $\mathfrak{C}$ can be achieved using the inverse map $\Phi^{-1} : \mathfrak{S} \to \mathfrak{C}$, since $\mathfrak{B}$ is a subset of $\mathfrak{S}$. That is,

$$\Phi^{-1}(\mathfrak{B}) \subset \Phi^{-1}(\mathfrak{S}) = \mathfrak{C}.$$

In particular, for any deterministic trajectory $s \in \mathfrak{B}$, we can construct the corresponding PCE coefficient trajectory $\mathsf{s} = \Phi^{-1}(s) \in \mathfrak{C}$, with

$$\mathsf{s}^j = \begin{cases} s, & \text{if } j = 0, \\ 0, & \text{if } j \in \mathbb{N} \setminus \{0\}. \end{cases}$$

The reconstruction from the moment behavior $\mathfrak{M}$ to the PCE coefficient behavior $\mathfrak{C}$, and subsequently to the random-variable behavior $\mathfrak{S}$, relies on covariance decomposition.

Corollary 4.2 ($\mathfrak{M}_T \to \mathfrak{S}_T$, $\mathfrak{M}_T \to \mathfrak{C}_T$): Let $\xi_{[0,T-1]} \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s T}\right)$ be an arbitrary random variable with zero mean and unit covariance. For $(m, \Gamma)_{[0,T-1]} \in \mathfrak{M}$, consider a Cholesky decomposition $\Gamma_{[0,T-1]} = L_T L_T^\top$, there are $S_{[0,T-1]} \in \mathfrak{S}_T$ and $\mathsf{s}_{[T-1]} \in \mathfrak{C}_T$ with

$$S_{[0,T-1]} \doteq m_{[0,T-1]} + L_T \xi_{[0,T-1]}, \quad \mathsf{s}^0_{[0,T-1]} \doteq m_{[0,T-1]}, \, \mathsf{s}^j_{[0,T-1]} = L^j_T,$$

where L^j_T is the j-th column of L_T with $j \in \mathbb{I}_{[1,n_s T]}$.

Due to the lower triangular structure of L from the Cholesky decomposition, we ensure that L_T is a principal submatrix of L_{T+1}, i.e.,

$$\Gamma_{[0,T]} = L_{T+1} L_{T+1}^\top, \quad L_{T+1} \doteq \begin{bmatrix} L_T & 0_{n_s T \times n_s} \\ & \ell_T \end{bmatrix}, \quad \ell_T \in \mathbb{R}^{n_s \times (T+1)}.$$

This structure facilitates the recursive construction of $S_{[0,T]}$ as

$$S_{[0,T]} \doteq m_{[0,T]} + L_{T+1} \xi_{[0,T]} = \begin{bmatrix} S_{[0,T-1]} \\ S_T \end{bmatrix}, \quad S_T = m_T + \ell_T \xi_{[0,T]}.$$

Furthermore, the selection of ξ provides an additional degree of freedom in constructing S. By choosing the germ ξ to be Gaussian, one can bijectively recover LTI Gaussian processes from the moment behavior.

Note that the relation in $\mathfrak{S} \Leftrightarrow \mathfrak{C}$ is bilateral as a consequence of the bijective mapping Φ. This leads to the equivalence of behaviors.

Theorem 4.3 (Behavioral equivalence (Faulwasser et al., 2023)): Let $S \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s}\right)^{\mathbb{Z}}$ with its corresponding PCE coefficients $\mathsf{s} \in (\ell^2(\mathbb{R}^{n_s}))^{\mathbb{Z}}$ and for $\omega \in \Omega$ its realizations $S(\omega) \in (\mathbb{R}^{n_s})^{\mathbb{Z}}$. Then, the following statements are equivalent:

i) $S(\omega) \in \mathfrak{B}$ for μ-almost all $\omega \in \Omega$,

ii) $S \in \mathfrak{S}$,

iii) $\mathsf{s} \in \mathfrak{C}$,

iv) $\mathsf{s}^j \in \mathfrak{B}$ for all $j \in \mathbb{N}$.

Proof. The equivalence between i) and ii) as well as between iii) and iv) follows by definition. The equivalence of ii) and iii) follows with the behavioral lift, Theorem 4.2 i). □

By restricting the elements of $\mathfrak{B}$, $\mathfrak{S}$, and $\mathfrak{C}$ to a finite time horizon T, one obtains the equivalences for $\mathfrak{B}_T$, $\mathfrak{S}_T$, and $\mathfrak{C}_T$.

4.2.5. Discussion

As illustrated in Figure 4.2, stochastic LTI systems can be described through four distinct behavioral representations, each with its own characteristics, advantages, and limitations:

- Stochastic behavior $\mathfrak{S}$ with L^2 random variables.

  The L^2 random-variable framework is general and extends beyond standard linear-Gaussian settings. It conceptually enables exact forward propagation of uncertainty through dynamics. However, in practice, L^2 random variables are defined within infinite-dimensional Hilbert spaces, which makes numerical implementation challenging. Approximations have to be employed to deal with this infinite-dimensional nature.

- Realization behavior $\mathfrak{B}$ with sampled trajectories.

  The realization behavior leverages sampled trajectories and thus makes it inherently finite-dimensional. However, uncertainty propagation in this setting has certain limitations: It either assumes unrealistically perfect knowledge of future disturbances, or it relies on sampling-based methods, which introduce approximation errors and often scale poorly in non-Gaussian settings.

- The first two moments behavior $\mathfrak{M}$.

  Representing the system via its statistical moments, especially the first two moments, is a standard approach in control theory. This is sufficient for the Gaussian setting. However, in non-Gaussian settings, restricting the representation to the first two moments can lead to significant information loss. Moreover, existing approaches for moment propagation often assume a fixed feedback policy, such as $U(k) = KX(k)$, to simplify covariance pair characterizations across variables (e.g., U, Y, X) and across time instances.

- The PCE coefficient behavior $\mathfrak{C}$.

  Due to the linearity of PCE, the PCE coefficient behavior aligns naturally with the linear dynamics encoded by $\mathfrak{B}$. Unlike the realization behavior, PCE does not require perfect knowledge of future disturbances; instead, it relies only on the statistical information about them. Moreover, PCE behavior enables an exact propagation of uncertainties and can determine moment behavior through the map $\Upsilon \circ \Phi$. While the PCE coefficient behavior is initially introduced with an infinite PCE dimension, subsequent chapters demonstrate how Lemma 3.2 can be employed to obtain a finite-dimensional representation. This finite-dimensional framework enables efficient stochastic predictions in Chapter 5 and facilitates computationally tractable optimal control approaches in Chapters 6–7.

4.3. Non-Parametric Stochastic Representation

The behavioral lift and equivalence, as established in Theorems 4.2 and 4.3, respectively, provide a foundation for formulating a version of the fundamental lemma tailored for stochastic systems. In particular, leveraging these results allows us to lift the non-parametric representation of realization behavior to its stochastic counterparts.

4.3.1. Column Spaces Equivalence and Inclusion

Notice that the PCE coefficients of the random variables are determined concerning their known distributions, and thus, the PCE coefficient behavior $\mathfrak{C}$ is *deterministic*. Therefore, it admits the conceptual application of the usual LTI fundamental lemma.

Lemma 4.4 (Fundamental lemma for PCE coefficients (Pan et al., 2023c)):
Consider a controllable LTI system $\Sigma = (\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{B})$, its input-output partitioning $(\tilde{u}, y) = s$, and its minimal state dimension $n^\star$. Suppose $\mathfrak{S}$ represents the stochastic behavior in (4.3), and $\mathfrak{C}$ denotes the corresponding PCE coefficient behavior in (4.10).

Consider $T, N \in \mathbb{N}^+$. Let $S^{\mathrm{d}}_{[0,T-1]} \in \mathfrak{S}_T$ with corresponding expansion coefficients $\mathsf{s}^{\mathrm{d}}_{[0,T-1]} = (\tilde{\mathsf{u}}^{\mathrm{d}}, \mathsf{y}^{\mathrm{d}})_{[0,T-1]} \in \mathfrak{C}_T$ be such that $\tilde{\mathsf{u}}^{\mathrm{d},j}_{[0,T-1]}$ is persistently exciting of order $N + n^\star$ for all $j \in \mathbb{N}$. Then, $\mathsf{s}_{[0,N-1]} \in \mathfrak{C}_N$ if and only if there is $\mathsf{g}^j \in \mathbb{R}^{T-N+1}$ such that

$$\mathcal{H}_N\left(\mathsf{s}^{\mathrm{d},j}_{[0,T-1]}\right)\mathsf{g}^j = \mathsf{s}^j_{[0,N-1]} \tag{4.18}$$

holds for all $j \in \mathbb{N}$.

The proof follows from Lemma 2.11 and is thus omitted.

While Lemma 4.4 is relatively straightforward, measuring or estimating the PCE coefficients of a stochastic LTI system is non-trivial. Therefore, the previous result may appear to be of limited practical use. However, as will be demonstrated below, the structural similarity of the PCE coefficient behavior, the realization behavior, and the original stochastic behavior enables further valuable insights.

Lemma 4.5 (Column-space equivalence and inclusion (Pan et al., 2023c)):
Consider a controllable LTI system $\Sigma = (\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{B})$, its input-output partitioning $(\tilde{u}, y) = s$, and its minimal state dimension $n^\star$. Suppose $\mathfrak{S}$ represents the stochastic behavior in (4.3), and $\mathfrak{C}$ denotes the corresponding PCE coefficient behavior in (4.10).

For $T \in \mathbb{N}^+$, let $S_{[0,T-1]} \in \mathfrak{S}_T$ with corresponding expansion coefficients $\mathsf{s}_{[0,T-1]} = (\tilde{\mathsf{u}}, \mathsf{y})_{[0,T-1]} \in \mathfrak{C}_T$ and $s_{[0,T-1]} = (\tilde{u}, y)_{[0,T-1]} \in \mathfrak{B}_T$.

i) Assume that both the realization $\tilde{u}_{[0,T-1]}$ and the coefficients $\tilde{\mathsf{u}}^j_{[0,T-1]}$ for all $j \in \mathbb{N}$ are persistently exciting of order $N + n^\star$. Then, for all $j \in \mathbb{N}$

$$\operatorname{colspan}\left(\mathcal{H}_N\left(\mathsf{s}^j_{[0,T-1]}\right)\right) = \operatorname{colspan}\left(\mathcal{H}_N\left(s_{[0,T-1]}\right)\right) = \mathfrak{B}_N. \tag{4.19a}$$

ii) Assume that only $\tilde{u}_{[0,T-1]}$ is persistently exciting of order $N + n^\star$. Then,

$$\operatorname{colspan}\left(\mathcal{H}_N\left(\mathsf{s}^j_{[0,T-1]}\right)\right) \subseteq \operatorname{colspan}\left(\mathcal{H}_N\left(s_{[0,T-1]}\right)\right). \tag{4.19b}$$

Proof. The proof of (4.19a) in Part i) follows directly from the observation that the realization behavior $\mathfrak{B}$ and the j-th PCE coefficient of $\mathfrak{C}$ share the same dynamics for all $j \in \mathbb{N}$. Thus, by applying the fundamental lemma (Lemma 2.11) to them respectively, we prove assertion i). Statement ii) follows by relaxing the persistency of excitation of the PCE coefficients trajectories. □

It should be noted that from the perspective of the applications, statement i) of the previous lemma has significant limitations due to the assumption of persistency of excitation for all expansion indices $j \in \mathbb{N}$ in (4.19a). By relaxing this assumption, statement ii) illustrates the inclusion of column spaces between the trajectories of PCE coefficients and realization trajectories.

4.3.2. A Stochastic Fundamental Lemma

Based on the result of column inclusion, we can now formulate the fundamental lemma for stochastic LTI systems as suggested by Faulwasser et al., 2023; Pan et al., 2023c.

Lemma 4.6 (Stochastic fundamental lemma):
Consider a controllable LTI system $\Sigma = (\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{B})$, its input-output partitioning $(\tilde{u}, y) = s$, and its minimal state dimension $n^\star$. Denote by $\mathfrak{S}$ the stochastic behavior in (4.3), and $\mathfrak{C}$ the corresponding PCE coefficient behavior in (4.10).

Let $s^{\mathrm{d}}_{[0,T-1]} = \left(\tilde{u}^{\mathrm{d}}, y^{\mathrm{d}}\right)_{[0,T-1]}$ be an element of $\mathfrak{B}_T$ with $\tilde{u}^{\mathrm{d}}_{[0,T-1]}$ persistently exciting of order $N + n^\star$. Then, the following statements hold:

i) $\mathsf{s}_{[0,N-1]} \in \mathfrak{C}_N$ if and only if there is $\mathsf{g} \in \ell^2(\mathbb{R}^{T-N+1})$ such that

$$\mathcal{H}_N\left(s^{\mathrm{d}}_{[0,T-1]}\right) \mathsf{g}^j = \mathsf{s}^j_{[0,N-1]} \tag{4.20a}$$

holds for all $j \in \mathbb{N}$.

ii) $S_{[0,N-1]} \in \mathfrak{S}_N$ if and only if there is $G \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{T-N+1}\right)$ such that

$$\mathcal{H}_N\left(s^{\mathrm{d}}_{[0,T-1]}\right) G = S_{[0,N-1]}. \tag{4.20b}$$

iii) $(m, \Gamma)_{[0,N-1]} \in \mathfrak{M}_N$ if and only if there exist $m^G \in \mathbb{R}^{T-N+1}$ and $\Gamma^G \in \mathbb{R}^{(T-N+1)\times(T-N+1)}$ such that

$$\mathcal{H}_N\left(s^{\mathrm{d}}_{[0,T-1]}\right) m^G = m_{[0,N-1]} \tag{4.20c}$$

$$\mathcal{H}_N\left(s^{\mathrm{d}}_{[0,T-1]}\right) \Gamma^G \mathcal{H}_N\left(s^{\mathrm{d}}_{[0,T-1]}\right)^\top = \Gamma_{[0,N-1]}, \quad \Gamma^G \succeq 0. \tag{4.20d}$$

Proof. Statement i) builds on the column space equivalence of Lemma 4.5 and directly follows from the fundamental lemma for PCE coefficients (Lemma 4.4). We now demonstrate that g^j is squarely summable for all $j \in \mathbb{N}$, i.e., $\mathsf{g} \in \ell^2(\mathbb{R}^{T-N+1})$.

Since the linear equation (4.20a) is underdetermined, for a given coefficient trajectory $\mathsf{s}_{[0,N-1]} \in \mathfrak{C}_N$, the least-square solution g is

$$\mathsf{g}^j = \mathcal{H}_N\left(s^{\mathrm{d}}_{[0,T-1]}\right)^\dagger \mathsf{s}^j_{[0,N-1]}, \; \forall j \in \mathbb{N}.$$

With bounded $\left\| \mathcal{H}_N\left(s^{\mathrm{d}}_{[0,T-1]}\right)^\dagger \right\|_F \leq \infty$, the square summability of this g^j for all $j \in \mathbb{N}$ then is inherited from square summability of $\mathsf{s}^j_{[0,N-1]}$ for all $j \in \mathbb{N}$.

Statement ii) follows by combining Statement i) and the behavioral lift established in Theorem 4.2. The relation between $G \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{T-N+1}\right)$ and $\mathsf{g} \in \ell^2(\mathbb{R}^{T-N+1})$ is given by

$$G = \sum_{j \in \mathbb{N}} \phi^j(\xi) \mathsf{g}^j.$$

The third statement is derived by evaluating the mean and covariance of (4.20b), where

$$m^G = \mathbb{E}[G], \quad \Gamma^G = \Sigma[G].$$

□

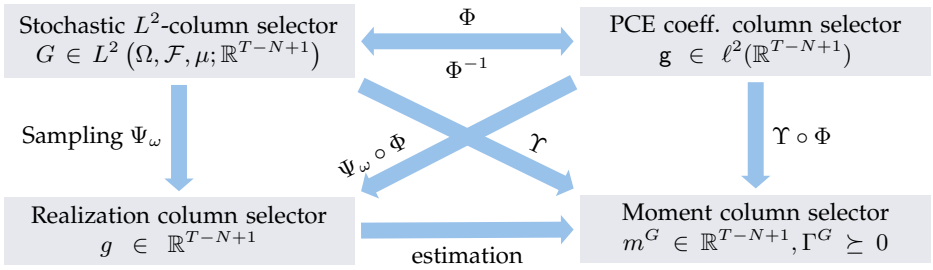

Figure 4.3.: Relations and maps between column selectors appeared in Lemma 4.6.

The standard form of the fundamental lemma states that the columns of the Hankel matrix constructed from past trajectories span the linear subspace of all potential trajectories of the behavior. However, obtaining the past trajectories of $\mathfrak{S}$, $\mathfrak{C}$, and $\mathfrak{M}$ is challenging, as it typically requires a large number of realization trajectories for accurate estimation. In contrast, the equivalence and inclusion properties of column spaces established in Lemma 4.5 provide a crucial link: one can safely use a Hankel matrix of realizations to represent trajectories in terms of random variables, PCE coefficients, and moment pairs.

While the data-driven representations of $\mathfrak{S}$, $\mathfrak{C}$, and $\mathfrak{M}$ share the same Hankel matrices derived from realization trajectories, the column selectors in (4.20) serve the role of behavioral lifting. As shown in Figure 4.3, the relations between the column selectors in (4.20) exhibit the same structural patterns as those among the behaviors shown in Figure 4.2. Specifically, g, g^j, and (m^G, Γ^G) represent the realization, the PCE coefficients, and the first two moments of the random-variable column selector G, respectively.

In addition, as demonstrated in Proposition 4.2, every column—and, by symmetry, every row—of the covariance matrix $\Gamma_{[0,T-1]}$ belongs to the realization behavior $\mathfrak{B}_T$. This property aligns with the structure of $\mathcal{H}_N \Gamma^G \mathcal{H}_N^\top$ in (4.20d), as the column space of the Hankel matrix satisfies $\operatorname{colspan}(\mathcal{H}_N) = \mathfrak{B}_T$.

Non-iff Formulation of the Stochastic Fundamental Lemma

In the stochastic setting, if the Hankel matrix is constructed directly from random variables, the standard form of the fundamental lemma does not necessarily hold. In other words, this alternative version of the stochastic fundamental lemma does not admit an *iff* condition.

Corollary 4.3 (Non-iff formulation (Pan et al., 2023c)): Let $S_{[0,T-1]} = (\tilde{U}, Y)_{[0,T-1]} \in \mathfrak{S}_T$ with corresponding expansion coefficients $\mathsf{s}_{[0,T-1]} = (\tilde{\mathsf{u}}, \mathsf{y})_{[0,T-1]} \in \mathfrak{C}_T$ be such that $\tilde{\mathsf{u}}^j_{[0,T-1]}$ are persistently exciting of order $N + n^\star$ for all $j \in \mathbb{N}$. Additionally, let the realization trajectory $\mathsf{s}_{[0,T-1]} = (\tilde{u}, y)_{[0,T-1]} \in \mathfrak{B}_T$ be such that $\tilde{u}_{[0,T-1]}$ is persistently exciting of order $N + n^\star$.

i) For arbitrary $g \in \mathbb{R}^{T-L+1}$, there exists $G \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{T-L+1})$ such that

$$\mathcal{H}_N\left(S_{[0,T-1]}\right) g = \mathcal{H}_N\left(s_{[0,T-1]}\right) G. \tag{4.21a}$$

ii) If there exists $g \in \mathbb{R}^{T-N+1}$ such that

$$\mathcal{H}_N\left(S_{[0,T-1]}\right) g = S'_{[0,N-1]}, \tag{4.21b}$$

then $S'_{[0,N-1]} \in \mathfrak{S}_N$.

Proof. Part i): Note that we have the following reformulation

$$\mathcal{H}_N(S_{[0,T-1]})g = \mathcal{H}_N\left(\sum_{j=0}^{\infty} \mathsf{s}^j_{[0,T-1]}\phi^j(\xi)\right) g$$
$$= \sum_{j=0}^{\infty} \mathcal{H}_N\left(\mathsf{s}^j_{[0,T-1]}\phi^j(\xi)\right) g = \sum_{j=0}^{\infty} \phi^j(\xi)\mathcal{H}_N\left(\mathsf{s}^j_{[0,T-1]}\right) g.$$

Here, the second equality holds since $\mathcal{H}(\cdot)$ and the summation are both linear operations. Moreover, the basis function $\phi^j(\xi)$ is a scalar polynomial of $\xi \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\xi})$, i.e., $\phi^j(\xi) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$. Therefore, the third equality holds by pulling $\phi^j(\xi)$ out from $\mathcal{H}_N(\cdot)$.

Then, using the column space equivalence (4.19a), for all $j \in \mathbb{N}^+$ and any $g \in \mathbb{R}^{T-N+1}$, we can find $\mathsf{g}^j \in \mathbb{R}^{T-N+1}$, such that $\mathcal{H}_N\left(\mathsf{s}^j_{[0,T-1]}\right) g = \mathcal{H}_N\left(s_{[0,T-1]}\right) \mathsf{g}^j$. This leads to

$$\mathcal{H}_N(S_{[0,T-1]})g = \sum_{j=0}^{\infty} \phi^j(\xi)\mathcal{H}_N\left(\mathsf{s}^j_{[0,T-1]}\right) g = \mathcal{H}_N(s_{[0,T-1]}) \sum_{j=0}^{\infty} \phi^j(\xi)\mathsf{g}^j.$$

The assertion follows with $G \doteq \sum_{j=0}^{\infty} \phi^j(\xi)\mathsf{g}^j$.

Part ii): Note that this statement only asserts that $g \mapsto S'_{[0,N-1]}$. With statement i), we have that for any $g \in \mathbb{R}^{T-N+1}$ there exists $G \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{T-N+1}\right)$ such that (4.21a) holds. Furthermore, Part ii) of Lemma 4.6 gives that G determines a random variable trajectory $S' \in \mathfrak{S}$ by (4.20b). □

We conclude our discussion with a simple example illustrating this observation.

Example 4.1: Consider the scalar stochastic system $X(k+1) = X(k) + U(k)$ with past data given by the PCEs

$$\begin{aligned} X(0) &= 0\phi^0 + 0\phi^1, \quad U(0) = 0\phi^0 + 1\phi^1, \\ X(1) &= 0\phi^0 + 1\phi^1, \quad U(1) = 1\phi^0 + 0\phi^1, \\ X(2) &= 1\phi^0 + 1\phi^1, \quad U(2) = 1\phi^0 + 1\phi^1. \end{aligned}$$

Note that the PCE coefficients of $U_{[0,2]}$ satisfy the persistency of excitation required by Part ii) of Lemma 4.6. We aim to find g in (4.21b) to represent $X'(0) = 0\phi^0 + 1\phi^1$ and $U'(0) = 0\phi^0 + 1\phi^1$. We obtain (4.21b) as

$$\begin{bmatrix} X(0) & X(1) & X(2) \\ U(0) & U(1) & U(2) \end{bmatrix} g = \begin{bmatrix} X'(0) \\ U'(0) \end{bmatrix}. \tag{4.22}$$

After applying Galerkin projection onto the basis functions and stacking the projected equations we obtain $Mg = c$ with

$$M = \left[\begin{array}{ccc} 0 & 0 & 1 \\ 0 & 1 & 1 \\ \hline 0 & 1 & 1 \\ 1 & 0 & 1 \end{array}\right] \quad c = \left[\begin{array}{c} 0 \\ 0 \\ \hline 1 \\ 1 \end{array}\right]$$

where the upper block corresponds to ϕ^0 and the lower one to ϕ^1.

By the Rouché–Capelli theorem, $Mg = c$ admits a solution g if and only if the block matrix $[M \ c]$ has the same rank as M. Observe that in the example above $\operatorname{rank}(M) = 3$ and $\operatorname{rank}([M \ c]) = 4$. Thus, we conclude that (4.22) does not admit solutions $g \in \mathbb{R}^4$.

4.4. Summary

In this chapter, we have studied the stochastic behavior of dynamical systems. Stochastic behavior encompasses all stochastic processes whose realizations are compatible with the dynamics encoded by the deterministic behavior, while the uncertain environment is represented by the underlying probability space. Furthermore, the random-variable stochastic behavior can be further characterized by its PCE coefficients and the first two moments.

A key observation in our development is that the random-variable behavior $\mathfrak{S}$ (4.3), the realization behavior $\mathfrak{B}$ (2.6), and the PCE coefficient behavior $\mathfrak{C}$ (4.9) are governed by the same dynamics. These behaviors share identical LTI kernels $R(\boldsymbol{\sigma})$, as proved in Theorem 4.2. This structural similarity results from the synergy of the linearity of the system structures and the linearity of the PCEs of the random variables.

In contrast, the moment behavior deviates from this pattern. It focuses on covariance matrices $\Sigma\left[S(k), S(t)\right]$ at time pairs or, with input/output partitioning, on input-output covariances $\Sigma\left[\tilde{U}(k), Y(t)\right]$, which require tracking dependencies between multiple variables over time. Especially in control applications, this results in complex parametric forms and increased computational burden. On the other hand, the PCE approach expands the random variables into a set of independent coefficients. Because of its linearity, PCE behavior shares the same parametric representation as realization behavior. Moreover, it enables the reconstruction of moment behavior through the mapping $\Upsilon \circ \Phi$, offering a more computationally efficient alternative.

The PCE approach allows us to extend the standard fundamental lemma to stochastic behaviors. Crucially, these extensions are built on Hankel matrices constructed using realizations of stochastic variables, i.e., they rely on measurement data only. Moreover, in the later development of data-driven stochastic prediction and control, PCE serves as a unifying tool to bridge the gap between representing stochastic behavior, extracting moment information, and sampling realizations.

5. Prediction of Stochastic System Behaviors

In the previous chapter, the stochastic behavior of a dynamical system is defined by considering the system trajectory S to be stochastic. As shown in Section 4.1, the stochastic uncertainties arise from treating the exogenous disturbance W as stochastic. However, many disturbances cannot be measured, or their statics are not modeled, such as occupancy rates in buildings, door openings, or other unobserved events. While additional sensors can be deployed to measure new disturbance variables, there will always remain unstructured (unmeasured or unmodeled) ones. This highlights the need for the representation and the prediction of stochastic behaviors influenced by both structured and unstructured disturbances.

In this chapter, we follow the routine established by Willems, 2012 and Baggio et al., 2017 to define stochastic behaviors corrupted by unstructured disturbances using a residual disturbance framework. By introducing the ARX representation of the corrupted stochastic behavior, we explicitly characterize key properties of these behaviors. Building on this, we extend the stochastic fundamental lemma to construct a data-driven representation of the corrupted behavior. Furthermore, we demonstrate that past realizations of residual disturbances can be estimated using recorded input and output data. With these estimated disturbances, we propose a regularized stochastic predictor which mitigates the effects of estimation. Finally, in Section 5.3, we examine the effectiveness of the proposed approach using measured data from a real-world building.

5.1. Stochastic Behavior with Residual Disturbance

As illustrated in Figure 5.1, in addition to the manifest inputs $\tilde{U} = (U, W)$, which include the manipulated input U and the structured (measured and modeled) exogenous disturbances W, we introduce a disturbance process $V \in (L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y}))^{\mathbb{Z}}$ to capture all unstructured (unmeasured or unmodeled) disturbances. We refer to W as the *structured disturbance* and V as the *residual disturbance*. We collectively describe all unmeasured or unmodeled disturbances as the *unstructured disturbance*, denoted by W'. Later, we will demonstrate how consolidating all unstructured disturbances into a single residual term allows for the estimation of $V \in (L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y}))^{\mathbb{Z}}$ and partial modeling of its statistical properties.

Suppose that the stochastic behavior without the residual disturbance can be de-

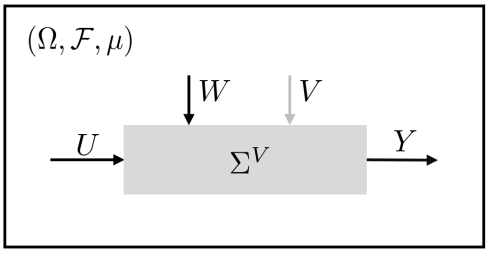

Figure 5.1.: A corrupted stochastic dynamical system $\Sigma^V = (\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{S}^V)$ with the structured disturbance $W \in (L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_w}))^{\mathbb{Z}}$, the residual disturbance $V \in (L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_x}))^{\mathbb{Z}}$, and the input $U \in (L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_u}))^{\mathbb{Z}}$

scribed by $\mathfrak{S}(R)$ with $R(z) \in \mathbb{R}^{n_y \times n_s}$ of full row rank. Following the framework proposed by Willems, 2012 and Baggio et al., 2017, we define *the corrupted stochastic behavior* for a given residual disturbance $V \in (L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y}))^{\mathbb{Z}}$ in a kernel representation as

$$\mathfrak{S}^V(R) \doteq \left\{ S \in (L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_s}))^{\mathbb{Z}} \,\middle|\, R(\boldsymbol{\sigma})S = V \right\}. \tag{5.1}$$

Finally, we define the corrupted stochastic dynamical system as $\Sigma^V = (\mathbb{Z}, \mathbb{R}^{n_s}, \mathfrak{S}^V)$.

5.1.1. Parametric Representations

Consider the input and output partition with $S = (\tilde{U}, Y)$, the kernel representation (5.1) can be reformulated as

$$Y(k) = \Xi \begin{bmatrix} \tilde{U}_{[k-\ell,k-1]} \\ Y_{[k-\ell,k-1]} \end{bmatrix} + D\tilde{U}(k) + V(k), \ \forall k \in \mathbb{Z}. \tag{5.2}$$

For a specific $V \in (L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y}))^{\mathbb{Z}}$, we define its ARX representation as

$$\mathfrak{S}^V(\Xi, D) \doteq \left\{ (\tilde{U}, Y) \in (L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y+n_{\tilde{u}}}))^{\mathbb{Z}} \,\middle|\, (\tilde{U}, Y, V) \text{ satisfy (5.2)} \right\}. \tag{5.3}$$

Lemma 5.1 (Equivalence of kernel and ARX representations): Consider the corrupted stochastic behavior $\mathfrak{S}^V$ with its kernel representation $\mathfrak{S}^V(R)$ with $R(z) \in \mathbb{R}^{n_y \times n_s}[z, z^{-1}]$ of full row rank. Then, there exists a unimodular polynomial matrix $U(z) \in \mathbb{R}^{n_y \times n_y}[z, z^{-1}]$ such that

$$U(z)R(z) = \begin{bmatrix} \sum_{i=0}^{\ell} Q_{\ell-i} z^{-i} & I + \sum_{i=1}^{\ell} P_{\ell-i} z^{-i} \end{bmatrix}$$

Then, consider Ξ as in (2.18), $D = Q_\ell$, and $V' = U(z)V$, we have

$$\mathfrak{S}^V(R) = \mathfrak{S}^{V'}(\Xi, D). \tag{5.4}$$

Proof. The equivalence between the kernel representation and the ARX representation follows directly from Lemma 2.4 and Lemma 2.8. Moreover, since $U(z)$ is unimodular, its inverse $U^{-1}(z)$ exists and is also unimodular. This allows the transformation of an ARX representation into a kernel representation use $U^{-1}(z)$, thereby establishing the equivalence stated in (5.4). □

Note that the ARX model (5.2) serves as the input/output counterpart of $\boldsymbol{\sigma}X = AX + BU + V$. Specifically, the disturbance V represents the residual term that accounts for deviations from the corrupted behavior to the ideal behavior $\mathfrak{S}(\Xi, D)$. Thus, V captures the influence of all disturbances not included in the structured disturbances W. The following result highlights this point.

Consider there are unmeasured or unmodeled disturbances $W'(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_{w'}})$ alongside the structured disturbance $W(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_w})$. We denote $W'(k)$ as *unstructured disturbances*. We use an observer-inspired technique to consolidate $W' \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_{w'}})^{\mathbb{Z}}$ into a single term $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y})^{\mathbb{Z}}$.

Proposition 5.1 (Compression of unstructured disturbances): Consider any minimal state-space representation with unstructured disturbance process $W' \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_{w'}})^{\mathbb{Z}}$

$$\begin{aligned} X(k+1) &= AX(k) + B\tilde{U}(k) + EW'(k), \quad k \in \mathbb{Z} \\ Y(k) &= CX(k) + D\tilde{U}(k) + FW'(k). \end{aligned} \tag{5.5}$$

There exists a dead-beat observer gain H such that $(A-HC)^\ell = 0$ for ℓ not smaller than the minimal lag of system (5.5). Moreover, for any input-output trajectory $(\tilde{U}, Y)_{[0,N-1]}$ of (5.5), there exists $V_{[0,N-1]} \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{Nn_y}\right)$ such that $(\tilde{U}, Y)_{[0,N-1]}$ is an N-length input-output trajectory of $\mathfrak{S}^V(\Xi, D)$ with Ξ as given in (2.23).

Proof. This proof follows similarly to the proof of Lemma 2.10. Since (5.5) is state observable, for any $\ell \in \mathbb{N}^+$ not smaller than the minimal lag of (5.5), there exists a dead-beat observer gain H such that $(A - HC)^\ell = 0$ (Fahmy et al., 1980). Furthermore, leveraging the observer gain H, we can construct the following dynamics mimic an observer

$$\hat{X}(k+1) = A\hat{X}(k) + B\tilde{U}(k) + H\left(Y(k) - \hat{Y}(k)\right), \quad \hat{X}(0) = X(0) \tag{5.6a}$$

$$\hat{Y}(k) = C\hat{X}(k) + D\tilde{U}(k). \tag{5.6b}$$

Define $V(k) \doteq Y(k) - \hat{Y}(k)$, we have

$$\begin{aligned} \hat{X}(k+1) &= (A - HC)\hat{X}(k) + (B - HD)\tilde{U}(k) + HY(k) \\ Y(k) &= C\hat{X}(k) + D\tilde{U}(k) + V(k). \end{aligned}$$

Similar as the proof of Lemma 2.10, recursively expressing $\hat{X}$ from $k-\ell$ to k by the above equation, we obtain the following

$$\hat{X}(k) = (A-HC)^\ell \hat{X}(k-\ell) + \sum_{i=1}^{\ell} (A-HC)^{i-1}\Big(HY(k-i) + (B-HD)\tilde{U}(k-i)\Big). \tag{5.7}$$

Note that with the dead-beat observer gain H, the term $(A - HC)^\ell \hat{X}(k-\ell) = 0$.

By multiplying C and adding $D\tilde{U}(k)$ and $V(k)$ on both sides of the above equation, we arrive at

$$Y(k) = \sum_{i=1}^{\ell} C(A - HC)^{i-1}\Big(HY(k-i) + (B - HD)\tilde{U}(k-i)\Big) + D\tilde{U}(k) + V(k).$$

Then, following the proof of Lemma 2.10, we have the above equation equivalent to the ARX expression (5.2) with Ξ the same as (2.23). □

Since we focus on data-driven representations, the system matrices A, B, C, D, Ξ, and H are used here solely for the convenient explanation of how V captures all unstructured disturbances. These matrices are neither required to be known nor computed in our approach.

Observe that (5.6) can be reformulated as follows

$$\hat{X}(k+1) = A\hat{X}(k) + B\tilde{U}(k) + HV(k), \tag{5.8a}$$

$$Y(k) = C\hat{X}(k) + D\tilde{U}(k) + V(k), \tag{5.8b}$$

which closely resembles the state-space representation in innovation form, commonly used in stochastic subspace identification problems (van Overschee et al., 2012). The main difference lies in the observer gain: the dead-beat observer gain H is replaced by the Kalman filter gain K. However, with K, the residual term $(A - KC)^\ell \hat{X}(k-\ell)$ in (5.7) is not zero. As a result, the innovation-form representation with K is not strictly

equivalent to the ARX model. For a recent development on using innovation form for data-driven control, we refer to Wang et al., 2025.

Leveraging only the input and output data, we bypass the step of subspace identification of the innovation-form representation. Instead, we directly represent the stochastic behavior with residual disturbances $\mathfrak{S}^V$ through the stochastic fundamental lemma.

To this end, we first extend the result of Lemma 2.9 to transform the ARX model (5.2) to a non-minimial state-space representation. We recall that Z composed of past manifest inputs $\tilde{U} = (U, W)$ and outputs Y is denoted as the extended state

$$Z(k) \doteq \begin{bmatrix} \tilde{U}_{[k-\ell,k-1]} \\ Y_{[k-\ell,k-1]} \end{bmatrix} \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_z}\right), \quad n_z = \ell(n_y + n_{\tilde{u}}). \tag{5.9}$$

Using the extended state $Z\left(k\right)$, we construct a state-space representation of (5.2),

$$Z\left(k+1\right) = \widetilde{A}Z\left(k\right) + \widetilde{B}\tilde{U}\left(k\right) + \widetilde{H}V\left(k\right), \quad \forall k \in \mathbb{Z} \tag{5.10a}$$

$$Y(k) = \Xi Z(k) + D\tilde{U}(k) + V\left(k\right) \tag{5.10b}$$

with $\widetilde{A} = \begin{bmatrix} \bar{A} \\ \Xi \end{bmatrix}$, $\widetilde{B} = \begin{bmatrix} \bar{B} \\ D \end{bmatrix}$, $\widetilde{H} = \begin{bmatrix} 0 \\ I_{n_y} \end{bmatrix}$ as given in (2.22). Moreover, since $Y\left(k\right)$ is the last component of $Z\left(k+1\right)$, it follows that

$$Y\left(k\right) = \widetilde{H}^\top Z\left(k+1\right). \tag{5.10c}$$

Note that (5.10) corresponds to a stochastic behavior $\widetilde{\mathfrak{S}}\left(\widetilde{A}, [\widetilde{B}, \widetilde{H}], \Xi, D\right)$. Here, we denote $\widetilde{\mathfrak{S}}$ as *the extended stochastic behavior* with input component expanded from $\tilde{U}$ to $(\tilde{U}, V)$.

Combining previous results, we establish the following relation. For given $V \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y}\right)^{\mathbb{Z}}$, we have

$$\mathfrak{S}(\Xi, D) = \mathfrak{S}^{V=0}(\Xi, D), \quad \mathfrak{S}^V(\Xi, D) = \Pi_{(\tilde{U},Y)}\tilde{\mathfrak{S}}\left(\widetilde{A}, \left[\widetilde{B}, \widetilde{H}\right], \Xi, D\right) \tag{5.11}$$

which shows the relations among

- $\mathfrak{S}$, the stochastic LTI behavior without the residual disturbance V,
- $\mathfrak{S}^V$, the behavior corrupted by the residual disturbance V, and
- $\widetilde{\mathfrak{S}}$, the extended behavior that consider V as part of its input component.

Furthermore, by explicitly considering the residual disturbance V as part of the input component, we derive the following result.

Lemma 5.2: In the state-space representation (5.10), the pair $\left(\widetilde{A}, \left[\widetilde{B}, \widetilde{H}\right]\right)$ is controllable regardless of the ARX system matrices Ξ and D.

Proof. We examine the controllability of the pair $\left(\widetilde{A}, \left[\widetilde{B}, \widetilde{H}\right]\right)$ via the rank of the reachability matrices $\mathcal{R}_{n_z}\left(\widetilde{A}, \widetilde{B}\right)$ and $\mathcal{R}_{n_z}\left(\widetilde{A}, \widetilde{H}\right)$ per (2.15).

Recall from (2.22) that

$$\widetilde{A} = \begin{bmatrix} \bar{A} \\ \Xi \end{bmatrix}, \widetilde{B} = \begin{bmatrix} \bar{B} \\ D \end{bmatrix}, \bar{A} = \begin{bmatrix} 0 & I_{(\ell-1)n_{\tilde{u}}} & 0 & 0 \\ 0_{n_{\tilde{u}} \times n_{\tilde{u}}} & 0 & 0_{n_{\tilde{u}} \times n_y} & 0 \\ 0 & 0 & 0 & I_{(\ell-1)n_y} \end{bmatrix}, \bar{B} = \begin{bmatrix} 0_{(\ell-1)n_{\tilde{u}} \times n_{\tilde{u}}} \\ I_{n_{\tilde{u}}} \\ 0_{(\ell-1)n_y \times n_{\tilde{u}}} \end{bmatrix}.$$

Observe that for $i \in \mathbb{I}_{[0,\ell-1]}$,

$$\widetilde{A}^i \widetilde{B} = \begin{bmatrix} 0_{(\ell-1-i)n_{\tilde{u}} \times n_{\tilde{u}}} \\ I_{n_{\tilde{u}}} \\ 0_{n_{\tilde{u}} i \times n_{\tilde{u}}} \\ \star_{\ell n_y \times n_{\tilde{u}}} \end{bmatrix}, \quad \widetilde{A}^i \widetilde{H} = \begin{bmatrix} 0_{\ell n_{\tilde{u}} \times n_y} \\ 0_{(\ell-1-i)n_y \times n_y} \\ I_{n_y} \\ \star_{n_y i \times n_y} \end{bmatrix},$$

where $\star$ indicates matrix blocks depending on Ξ and D. Observe that when multiplied by $\widetilde{A}^i$, the identity matrix blocks $I_{n_{\tilde{u}}}$ in $\widetilde{B}$ and I_{n_y} in $\widetilde{H}$ are shifted upwards by $n_{\tilde{u}} i$ rows and $n_y i$ rows, respectively. As a result,

$$\mathcal{R}_{n_z}(\widetilde{A}, \widetilde{B}) = \begin{bmatrix} I_{\ell n_{\tilde{u}}} \\ \star \end{bmatrix}, \quad \mathcal{R}_{n_z}(\widetilde{A}, \widetilde{H}) = \begin{bmatrix} 0_{\ell n_{\tilde{u}} \times n_y} & 0 & 0 \\ I_{n_y} & 0 & 0 \\ \vdots & \ddots & 0 \\ \star & \star & I_{n_y} \end{bmatrix}.$$

Therefore, $\mathcal{R}_{n_z}\left(\widetilde{A}, \left[\widetilde{B}, \widetilde{H}\right]\right)$, whose row rank is the same as $\left[\mathcal{R}_{n_z}\left(\widetilde{A}, \widetilde{B}\right) \quad \mathcal{R}_{n_z}\left(\widetilde{A}, \widetilde{H}\right)\right]$, is of full row rank. Hence, $\left(\widetilde{A}, \left[\widetilde{B}, \widetilde{H}\right]\right)$ is a controllable pair. □

Remark 5.1 (Residual disturbance vs. measurement noise): Note that the ARX model (5.2) inherently incorporates the residual disturbance $V(k)$ but it does not explicitly consider measurement noise. To include measurement noise $M(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y})$, one can modify the state-space realization (5.10) with the output channel (5.10c) replaced by $\hat{Y}(k) = \widetilde{H}^\top Z(k+1) + M(k)$. The distinction between $\hat{Y}(k)$ and $Y(k)$ is important since the output $Y(k)$ contributes directly to the system dynamics as it is part of $Z(k+1)$, while $M(k)$ acts on the feedback channel. In this thesis, we focus on exogenous disturbances while a detailed analysis of noisy measurements is left for future work. We conjecture that the approaches proposed by Berberich et al., 2020; Coulson et al., 2019 offer a promising direction for addressing this issue.

Remark 5.2 (Residual disturbance vs. model residual): In addition to accounting for unstructured disturbances, the residual $V(k)$ can also be used to capture plant nonlinearity. In this context, $V(k)$ in (5.2) is commonly referred to as *the model residual* and is expected to have a nonlinear dependence on the extended state $Z(k)$. This nonlinear dependence can be modeled using data-driven techniques, e.g., Gaussian process regression by Hewing et al., 2020a. However, incorporating such a nonlinear mapping into online optimization remains challenging. Hence, in this work, we do not explicitly consider the influence of model residual in $V(k)$ and instead assume that $V(k)$ is independent of $Z(k)$.

5.1.2. Non-Parametric Representation

By applying the stochastic fundamental lemma (Lemma 4.6) to the extended behavior $\tilde{\mathfrak{S}}$, we arrive at the non-parametric representation for stochastic behaviors with the residual disturbance.

Corollary 5.1 (Extended stochastic fundamental lemma (Pan et al., 2025a)): Consider the extended stochastic behavior $\widetilde{\mathfrak{S}}\left(\widetilde{A}, [\widetilde{B}, \widetilde{H}], \Xi, D\right)$, its realization behavior $\widetilde{\mathfrak{B}}$, and its PCE coefficient behavior $\widetilde{\mathfrak{C}}$. For $T \in \mathbb{N}^+$, let $(\tilde{u}^{\mathrm{d}}, v^{\mathrm{d}}, y^{\mathrm{d}})_{[0,T-1]} \in \widetilde{\mathfrak{B}}_T$ with $(\tilde{u}^{\mathrm{d}}, v^{\mathrm{d}})_{[0,T-1]}$ persistently exciting of order $N + \ell(n_{\tilde{u}} + n_y)$. Then,

i) $(\tilde{u}, v, y)_{[0,N-1]} \in \widetilde{\mathfrak{B}}_N$ if and only if there is $g \in \mathbb{R}^{T-N+1}$ such that

$$\begin{bmatrix} \mathcal{H}_N\left(\tilde{u}^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_N\left(v^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_N\left(y^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix} g = \begin{bmatrix} \tilde{u}_{[0,N-1]} \\ v_{[0,N-1]} \\ y_{[0,N-1]} \end{bmatrix} \tag{5.12}$$

holds.

ii) $(\tilde{\mathsf{u}}, \mathsf{v}, \mathsf{y})_{[0,N-1]} \in \widetilde{\mathfrak{C}}_N$ if and only if for all $j \in \mathbb{N}$, there is $\mathsf{g} \in \ell^2(\mathbb{R}^{T-N+1})$ such that

$$\begin{bmatrix} \mathcal{H}_N\left(\tilde{u}^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_N\left(v^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_N\left(y^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix} \mathsf{g}^j = \begin{bmatrix} \tilde{\mathsf{u}}^j_{[0,N-1]} \\ \mathsf{v}^j_{[0,N-1]} \\ \mathsf{y}^j_{[0,N-1]} \end{bmatrix} \tag{5.13}$$

holds for all $j \in \mathbb{N}$.

iii) $(\tilde{U}, V, Y)_{[0,N-1]} \in \widetilde{\mathfrak{S}}_N$ if and only if there is $G \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{T-N+1}\right)$ such that

$$\begin{bmatrix} \mathcal{H}_N\left(\tilde{u}^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_N\left(v^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_N\left(y^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix} G = \begin{bmatrix} \tilde{U}_{[0,N-1]} \\ V_{[0,N-1]} \\ Y_{[0,N-1]} \end{bmatrix} \tag{5.14}$$

holds.

This corollary is a direct consequence of Lemma 4.6. By considering V as an additional input component, the persistency of excitation is imposed for $(\tilde{u}^{\mathrm{d}}, v^{\mathrm{d}})_{[0,T-1]}$. Moreover, thanks to Lemma 5.2, the extended lemma does not require any additional controllability assumption.

The key difference between (5.12) and (5.14) lies in the treatment of future residual disturbances. In (5.14), the future random variable $V_{[0,N-1]}$ is required. As will be shown later, $V_{[0,N-1]}$ can be estimated. In contrast, (5.12) assumes knowledge of the future realizations of residual disturbance, which is generally impractical. Next, we show that with the extended dynamics (5.10), the past $v^{\mathrm{d}}_{[0,T-1]}$ and the future $V_{[0,N-1]}$ used in (5.14) can be estimated from recorded $(\tilde{u}, y)^{\mathrm{d}}_{[0,T-1]}$.

Estimation of Past Residual Disturbances

Consider a realization trajectory $(z, \tilde{u}, v, y)^{\mathrm{d}}_{[0,T-1]}$ of (5.10), the measured data satisfy the following proposition.

Proposition 5.2 (Consistency of data): Any realization trajectory $(z, \tilde{u}, v, y)^{\mathrm{d}}_{[0,T-1]}$ of (5.10) satisfies

$$\left(\mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) - \mathcal{H}_1\left(v^{\mathrm{d}}_{[0,T-1]}\right)\right)\left(I_T - \begin{bmatrix}\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right)\\ \mathcal{H}_1\left(\tilde{u}^{\mathrm{d}}_{[0,T-1]}\right)\end{bmatrix}^{\dagger}\begin{bmatrix}\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right)\\ \mathcal{H}_1\left(\tilde{u}^{\mathrm{d}}_{[0,T-1]}\right)\end{bmatrix}\right) = 0, \tag{5.15}$$

where I_T denotes an identity matrix of size T, and $\cdot^{\dagger}$ denotes the Moore-Penrose inverse.

Proof. By horizontally stacking (5.10b) for $k \in \mathbb{I}_{[1,T]}$, we have

$$\mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) = [\Xi, D]\begin{bmatrix}\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right)\\ \mathcal{H}_1\left(\tilde{u}^{\mathrm{d}}_{[0,T-1]}\right)\end{bmatrix} + \mathcal{H}_1\left(v^{\mathrm{d}}_{[0,T-1]}\right).$$

Let $M \doteq \left[\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right)^{\top}, \mathcal{H}_1\left(\tilde{u}^{\mathrm{d}}_{[0,T-1]}\right)^{\top}\right]^{\top}$, its Moore-Penrose inverse $M^{\dagger}$ exists regardless of the row rank of M. Moreover, $M^{\dagger}$ is unique and satisfies $MM^{\dagger}M = M$. Thus, we obtain

$$\mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) = [\Xi, D]\, MM^{\dagger}M + \mathcal{H}_1\left(v^{\mathrm{d}}_{[0,T-1]}\right).$$

After substituting $[\Xi, D]M$ with $\mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) - \mathcal{H}_1\left(v^{\mathrm{d}}_{[0,T-1]}\right)$, we have

$$\mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) = \left(\mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) - \mathcal{H}_1\left(v^{\mathrm{d}}_{[0,T-1]}\right)\right)M^{\dagger}M + \mathcal{H}_1\left(v^{\mathrm{d}}_{[0,T-1]}\right),$$

which is equivalent to (5.15). □

Note that (5.15) admits infinitely many solutions of $v^{\mathrm{d}}_{[0,T-1]}$. Since we do not have further information about the statistics of $v(k)$, we rely on the least-squares estimate

$$\hat{v}^{\mathrm{d}}_{[0,T-1]} = \underset{v^{\mathrm{d}}_{[0,T-1]}}{\operatorname{argmin}} \left\|v^{\mathrm{d}}_{[0,T-1]}\right\|^2, \text{ subject to (5.15)}, \tag{5.16}$$

which admits a closed-form solution

$$\mathcal{H}_1(\hat{v}^{\mathrm{d}}_{[0,T-1]}) = \mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right)\left(I_T - \left(\begin{bmatrix}\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right)\\ \mathcal{H}_1\left(\tilde{u}^{\mathrm{d}}_{[0,T-1]}\right)\end{bmatrix}\right)^{\dagger}\begin{bmatrix}\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right)\\ \mathcal{H}_1\left(\tilde{u}^{\mathrm{d}}_{[0,T-1]}\right)\end{bmatrix}\right).$$

Proposition 5.3: Consider a realization trajectory $(z, \tilde{u}, v, y)^{\mathrm{d}}_{[0,T-1]}$ of (5.10). Then, for any estimated $\hat{v}^{\mathrm{d}}_{[0,T-1]}$ satisfying (5.15), there exist $\widehat{\Xi} \in \mathbb{R}^{n_y \times \ell(n_{\tilde{u}}+n_y)}$ and $\widehat{D} \in \mathbb{R}^{n_y \times n_{\tilde{u}}}$ such that $(z, \tilde{u}, \hat{v}, y)^{\mathrm{d}}_{[0,T-1]}$ satisfies

$$y^{\mathrm{d}}(k) = \widehat{\Xi} z^{\mathrm{d}}(k) + \widehat{D}\tilde{u}^{\mathrm{d}}(k) + \hat{v}^{\mathrm{d}}(k). \tag{5.17}$$

Proof. Let $M = \left[\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right)^{\top}, \mathcal{H}_1\left(\tilde{u}^{\mathrm{d}}_{[0,T-1]}\right)^{\top}\right]^{\top}$, $\left(I_T - M^{\dagger}M\right)$ represents the orthogonal projection onto the rightkernel of M, i.e.,

$$M\left(I_T - M^{\dagger}M\right) = 0.$$

Thus, $(z, \tilde{u}, \hat{v}, y)^{\mathrm{d}}_{[0,T-1]}$ satisfying (5.15) implies that its projection onto the rightkernel of M is zero; in other words, it lies in the rowspan of M. Thus, for each tuple $(z, \tilde{u}, \hat{v}, y)^{\mathrm{d}}_{[0,T-1]}$ satisfying (5.15), there exists a matrix $\Gamma \in \mathbb{R}^{n_y \times \ell(n_{\tilde{u}}+n_y)}$ such that

$$\mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) - \mathcal{H}_1\left(\hat{v}^{\mathrm{d}}_{[0,T-1]}\right) = \Gamma M.$$

Then, for each row of the above equation, we have

$$y^{\mathrm{d}}(k) = \Gamma \begin{bmatrix} z^{\mathrm{d}}(k) \\ \tilde{u}^{\mathrm{d}}(k) \end{bmatrix} + \hat{v}^{\mathrm{d}}(k), \quad \forall k \in \mathbb{I}_{[0,T-1]}.$$

By splitting Γ into $[\widehat{\Xi}, \widehat{D}]$ the assertion follows. □

As shown in Proposition 5.3, any $(z, \tilde{u}, \hat{v}, y)^{\mathrm{d}}_{[0,T-1]}$ satisfying (5.15) implicitly determines an LTI system. Thus, the usual Hankel matrix equations stated in Corollary 5.1 hold for $(z, \tilde{u}, \hat{v}, y)^{\mathrm{d}}_{[0,T-1]}$.

Prediction of Future Residual Disturbances

With the estimated past disturbance $\hat{v}^{\mathrm{d}}_{[0,T-1]}$, we then predict its future stochastic trajectory $V_{[0,N-1]}$ required in (5.14). For the sake of simplicity, we assume $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y})^{\mathbb{Z}}$ to be an i.i.d. stochastic process and independent of the extended state Z as discussed in Remark 5.2. Then, the identical distribution of $V(k)$ can be estimated by the empirical distribution of the past samples.

Let $\hat{V} \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y})^{\mathbb{Z}}$ be an estimate of $V \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y})^{\mathbb{Z}}$, we consider

$$\hat{V}(k) \sim \mu_{\hat{V}}, \ \forall k \in \mathbb{Z}, \quad \mu_{\hat{V}} = \frac{1}{T}\sum_{i=0}^{T-1} \delta_{\hat{v}^{\mathrm{d}}(i)},$$

where $\delta_{\hat{v}}$ is the Dirac measure at the point $\hat{v} \in \mathbb{R}^{n_y}$. The mean and covariance of $\mu_{\hat{V}}$ are given by

$$\hat{m}_v = \frac{1}{T}\sum_{i=0}^{T-1} \hat{v}^{\mathrm{d}}(i), \quad \hat{\Sigma}_v = \frac{1}{T-1}\sum_{i=0}^{T-1} \left(\hat{v}^{\mathrm{d}}(i) - \hat{m}_v\right)\left(\hat{v}^{\mathrm{d}}(i) - \hat{m}_v\right)^{\top}.$$

To construct exact PCEs of $\hat{V}(k)$, we reformulate $V(k)$ as

$$\hat{V}(k) = \hat{m}_v + \hat{\Sigma}_v^{\frac{1}{2}} \xi(k), \quad \xi(k) \sim \frac{1}{T}\sum_{i=0}^{T-1} \delta_{\xi^{\mathrm{d}}(i)}, \quad \xi^{\mathrm{d}}(i) = \hat{\Sigma}_v^{-\frac{1}{2}}\left(\hat{v}^{\mathrm{d}}(i) - \hat{m}_v\right)$$

where $\xi(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y})$ follows the normalized empirical distribution with zero mean and unit covariance. By Lemma 3.2, $\hat{V}(k)$ admits an exact PCE in the basis

$$\mathcal{P}_1(\xi(k)) \doteq \{1, \xi_1(k), \cdots, \xi_{n_y}(k)\},$$

where $\xi_i(k)$, $i \in \mathbb{I}_{[0,n_y]}$ are elements of $\xi(k)$.

Suppose a prediction of the future $W_{[0,N-1]}$ is available, and that $W_{[0,N-1]}$ admits an exact PCE in the basis $\{\phi_W\}_{j=0}^{L_w-1}$. We can construct a joint basis $\{\phi\}_{j=0}^{L-1}$ that allows exact PCEs for both $W_{[0,N-1]}$ and $\hat{V}_{[0,N-1]}$ as the union of ϕ_W and $\mathcal{P}_1(\xi(k))$

$$\{\phi\}_{j=0}^{L-1} \doteq \{\phi_W^j\}_{j=0}^{L_w-1} \bigcup \{\xi_1(k), \cdots, \xi_{n_y}(k)\}_{k=0}^{N-1}, \quad L = L_w + Nn_y. \tag{5.18}$$

Furthermore, considering $U_{[0,N-1]}$ as admitting an exact PCE in (5.18) is equivalent to determining $U_{[0,N-1]}$ through affine disturbance-feedback policies, which will be discussed in Chapter 6.

Based on this construction, we formalize the following assumption.

Assumption 5.1 (Exact PCEs for $\tilde{U}$ and $\hat{V}$): For a prediction horizon N, we assume that the manifest input $\tilde{U}_{[0,N-1]}$ composed of the manipulated input $U_{[0,N-1]}$, the structured disturbance $W_{[0,N-1]}$, and the estimated residual disturbance $\hat{V}_{[0,N-1]}$ admit exact PCEs in a finite-dimensional orthogonal basis $\{\phi^j(\xi)\}_{j=0}^{L-1}$ with $\xi \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\xi})$. Moreover, the exact PCEs are

$$\begin{aligned} U_{[0,N-1]} &= \textstyle\sum_{j=0}^{L-1} \mathsf{u}^j_{[0,N-1]} \phi^j(\xi), \quad & W_{[0,N-1]} &= \textstyle\sum_{j=0}^{L-1} \mathsf{w}^j_{[0,N-1]} \phi^j(\xi), \\ \tilde{U}_{[0,N-1]} &= \textstyle\sum_{j=0}^{L-1} \tilde{\mathsf{u}}^j_{[0,N-1]} \phi^j(\xi), \quad & \hat{V}_{[0,N-1]} &= \textstyle\sum_{j=0}^{L-1} \hat{\mathsf{v}}^j_{[0,N-1]} \phi^j(\xi). \end{aligned}$$

Remark 5.3 (Non-i.i.d. future residual disturbances): Following Remark 5.2, the residual disturbance $V(k)$ can be decomposed into two components

$$V(k) = g(Z(k)) + V_{\mathrm{d}}(k).$$

Both components may be non-i.i.d. The fist term $g(Z(k))$ corresponds to the model mismatch and depends on the extended state $Z(k)$, which can lead to potential dependencies across time instants. The second term $V_{\mathrm{d}}(k)$ captures residual disturbances and can be dominated by some non-i.i.d. factors, such as windows opening or changes in occupancy. We conjecture that machine learning techniques, e.g. LSTM neural networks (Hochreiter, 1997), could improve the prediction of non-i.i.d. future residual disturbances. However, by assuming that $V(k)$ is i.i.d., we will demonstrate the effectiveness of our approach through a real-world example in Section 5.3.

5.2. Data-Driven Stochastic Prediction

By estimating past residual disturbances $\hat{v}^{\mathrm{d}}_{[0,T-1]}$ and predicting future $\hat{V}_{[0,N-1]}$, we formulate a causal data-driven stochastic predictor. Given a T-length measured trajectory $(\tilde{u}^{\mathrm{d}}, y^{\mathrm{d}})_{[0,T-1]}$ of (5.10), we estimate $\hat{v}^{\mathrm{d}}_{[0,T-1]}$ by (5.16). Assume $(\tilde{u}^{\mathrm{d}}, \hat{v}^{\mathrm{d}})_{[0,T-1]}$ is persistently exciting of order $N + n_z + \ell$. Let p and f denote the ranges $[-\ell, -1]$ and $[0, N-1]$, respectively. Adopt the same partitioning of Hankel matrices as in (2.26). With ℓ past measurements $(\tilde{u}, y)_{[-\ell,-1]}$, the predicted $(\tilde{U}, \hat{V})_{[0,N-1]}$ and their PCEs per Assumption 5.1, a causal data-driven stochastic predictor reads

$$\hat{Y}_{\mathrm{f}} = \mathcal{H}_{y,\mathrm{f}} \sum_{j=0}^{L-1} \mathsf{g}^j, \quad \mathsf{g}^j \text{ s.t. } \begin{bmatrix} \mathcal{H}_{\tilde{u},\mathrm{p}} \\ \mathcal{H}_{y,\mathrm{p}} \\ \mathcal{H}_{\tilde{u},\mathrm{f}} \\ \mathcal{H}_{v,\mathrm{f}} \end{bmatrix} \mathsf{g}^j = \begin{bmatrix} \delta^{0j} \tilde{u}_{\mathrm{p}} \\ \delta^{0j} y_{\mathrm{p}} \\ \tilde{\mathsf{u}}^j_{\mathrm{f}} \\ \hat{\mathsf{v}}^j_{\mathrm{f}} \end{bmatrix}. \tag{5.19}$$

By Proposition 5.3 and Corollary 5.1, $(\tilde{U}, V, \hat{Y})_{[0,N-1]}$ is a trajectory of a stochastic LTI system. Hence, the input components $\tilde{U}$ and V are non-anticipated by the predicted output $\hat{Y}$. In other words, (5.19) is a causal prediction. Moreover, similar to (2.28),

(5.19) admits a unique solution in closed form as

$$\hat{Y}_{\mathrm{f}} = \mathcal{H}_{y,\mathrm{f}} \begin{bmatrix} \mathcal{H}_{\tilde{u},\mathrm{p}} \\ \mathcal{H}_{y,\mathrm{p}} \\ \mathcal{H}_{\tilde{u},\mathrm{f}} \\ \hat{\mathcal{H}}_{v,\mathrm{f}} \end{bmatrix}^{\dagger} \begin{bmatrix} \tilde{u}_{\mathrm{p}} \\ y_{\mathrm{p}} \\ \tilde{U}_{\mathrm{f}} \\ \hat{V}_{\mathrm{f}} \end{bmatrix}.$$

We note that the estimation error between the estimated disturbances and the true disturbance introduces a bias. This discrepancy arises in the context where we impose the equality constraint

$$\hat{\mathcal{H}}_{v,\mathrm{f}} \mathrm{g}^j - \hat{\mathrm{v}}_{\mathrm{f}}^j = 0.$$

However, due to estimation errors, the true past disturbances and predictions satisfy

$$\mathcal{H}_{v,\mathrm{f}}^{\mathrm{true}} \mathrm{g}^j - \mathrm{v}_{\mathrm{f}}^{j,\mathrm{true}} \neq 0.$$

To hedge against the estimation error, we propose the following robust predictor

$$\min_{\mathrm{g}^j} \max_{\|[\Delta\ \delta]\|_F \leq \beta^j} \left\| (\hat{\mathcal{H}}_{v,\mathrm{f}} + \Delta)\mathrm{g}^j - (\hat{\mathrm{v}}_{\mathrm{f}}^j + \delta) \right\|^2, \tag{5.20a}$$

$$\text{subject to} \begin{bmatrix} \mathcal{H}_{\tilde{u},\mathrm{p}} \\ \mathcal{H}_{y,\mathrm{p}} \\ \mathcal{H}_{\tilde{u},\mathrm{f}} \end{bmatrix} \mathrm{g}^j = \begin{bmatrix} \delta^{0j} \tilde{u}_{\mathrm{p}} \\ \delta^{0j} y_{\mathrm{p}} \\ \tilde{\mathrm{u}}_{\mathrm{f}}^j \end{bmatrix}, \tag{5.20b}$$

where $\Delta \in \mathbb{R}^{Nn_y \times n_g}$ and $\delta \in \mathbb{R}^{Nn_y}$ denote the perturbations affecting the estimation of past disturbance and the prediction of future disturbance, respectively. Moreover, $\|[\Delta\ \delta]\|_F$ denotes the Frobenius norm, and $\beta^j \in \mathbb{R}^+$ is the bound on $\|[\Delta\ \delta]\|_F$. The above equation minimizes the worst-case deviation with respect to all possible perturbations on $\hat{\mathcal{H}}_{v,\mathrm{f}}$ and $\hat{\mathrm{v}}_{\mathrm{f}}^j$. Hence, the minimizer $\mathrm{g}^{j,\star}$ is robust to estimation errors.

The following result shows how the minimizer of (5.20) can be obtained with a quadratic regularization of g^j.

Proposition 5.4 (Robust prediction by regularization): Let $\mathrm{g}^{j,\star} \in \mathbb{R}^{n_g}$ be the minimizer of

$$\min_{\mathrm{g}^j} \|\hat{\mathcal{H}}_{v,\mathrm{f}} \mathrm{g}^j - \hat{\mathrm{v}}_{\mathrm{f}}^j\|^2 + \lambda_g \|\mathrm{g}^j\|^2, \quad \mathrm{g}^j \text{ subject to (5.20b)} \tag{5.21}$$

then $\mathrm{g}^{j,\star}$ also minimizes (5.20) with

$$\beta^j = \begin{cases} \dfrac{\lambda_g \sqrt{\|\mathrm{g}^j\|^2 + 1}}{\|\hat{\mathcal{H}}_{v,\mathrm{f}} \mathrm{g}^{j,\star} - \hat{\mathrm{v}}_{\mathrm{f}}^j\|} & \text{if } \hat{\mathcal{H}}_{v,\mathrm{f}} \mathrm{g}^{j,\star} \neq \hat{\mathrm{v}}_{\mathrm{f}}^j \\ \lambda_g \sqrt{\|\mathrm{g}^{j,\star}\|^2 + 1} & \text{otherwise.} \end{cases} \tag{5.22}$$

Moreover, if $\mathrm{g}^{j,\star} \neq 0$, β^j (5.22) is strictly monotonically increasing with λ_g from (5.21).

Proof. We recall that the robust optimization problem in (5.20) can be interpreted as a robust least-squares estimation with uncertainty in the regressor matrix as discussed by El Ghaoui et al., 1997. Specifically, the formulation allows for uncertainty in $\mathcal{H}_{v,\mathrm{f}}$, bounded in Frobenius norm, and seeks to minimize the worst-case residual.

Huang et al., 2021 showed that the robust solution of (5.20) can be obtained by simply adding a quadrativ regularization term. The regularization weight is determined by

the uncertainty bound and the norm of the solution. Applying Corollary 2 from that work, we obtain the regularized form in (5.21), and the corresponding regularization weight β^j as in (5.22).

To make this explicit, let us denote the solution to (5.21) by $\mathrm{g}^{j,\star}$. If $\mathrm{g}^{j,\star} \neq 0$, the regularization parameter β^j adjusts to compensate for the uncertainty, maintaining equivalence to the robust formulation. The specific form of β^j given in (5.22). follows directly from the robust-to-regularized equivalence shown by Huang et al., 2021. The monotonic dependence of β^j on λ_g follows from the convexity and scaling properties of the regularization term.

This completes the proof.

□

Observe that (5.21) is a quadratic optimization problem subject to linear constraints. Its optimal solution can be obtained compactly via solving a linear system ,i.e. Karush–Kuhn–Tucker (KKT) system, as

$$\begin{bmatrix} 2(\hat{\mathcal{H}}_{v,\mathrm{f}}^\top \hat{\mathcal{H}}_{v,\mathrm{f}} + \lambda_g I_{n_g}) & \mathcal{H}^\top \\ \mathcal{H} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathrm{g}^{j,\star} \\ \eta^\star \end{bmatrix} = \begin{bmatrix} 2\hat{\mathcal{H}}_{v,\mathrm{f}}^\top \hat{\mathrm{v}}_{\mathrm{f}}^j \\ \delta^{0j}\tilde{u}_{\mathrm{p}} \\ \delta^{0j} y_{\mathrm{p}} \\ \tilde{\mathrm{u}}_{\mathrm{f}}^j \end{bmatrix}, \quad \mathcal{H} = \begin{bmatrix} \mathcal{H}_{\tilde{u},\mathrm{p}} \\ \mathcal{H}_{y,\mathrm{p}} \\ \mathcal{H}_{\tilde{u},\mathrm{f}} \end{bmatrix},$$

where $\eta^\star$ is the optimal Lagrange multiplier corresponding to the equality constraints in (5.20b). We have

$$\begin{bmatrix} M \\ \bar{M} \end{bmatrix} \doteq \begin{bmatrix} 2(\hat{\mathcal{H}}_{v,\mathrm{f}}^\top \hat{\mathcal{H}}_{v,\mathrm{f}} + \lambda_g I_{n_g}) & \mathcal{H}^\top \\ \mathcal{H} & \mathbf{0} \end{bmatrix}^\dagger, \quad \mathrm{g}^{j,\star} = M \begin{bmatrix} 2\hat{\mathcal{H}}_{v,\mathrm{f}}^\top \hat{\mathrm{v}}_{\mathrm{f}}^j \\ \delta^{0j}\tilde{u}_{\mathrm{p}} \\ \delta^{0j} y_{\mathrm{p}} \\ \tilde{\mathrm{u}}_{\mathrm{f}}^j \end{bmatrix},$$

which provides the closed-form solution of the stochastic predictor based on (5.21) as

$$\hat{Y}_{\mathrm{f}} \doteq \mathcal{H}_{y,\mathrm{f}} \sum_{j=0}^{L-1} \mathrm{g}^{j,\star} = M \begin{bmatrix} 2\hat{\mathcal{H}}_{v,\mathrm{f}}^\top \hat{V}_{\mathrm{f}} \\ \tilde{u}_{\mathrm{p}} \\ y_{\mathrm{p}} \\ \tilde{U}_{\mathrm{f}} \end{bmatrix}.$$

Comparisons of Predictors

First, note that the subspace predictor (2.29) does not propagate the uncertainty caused by the unstructured disturbances. If the future inputs and the structured disturbances are considered deterministic, the output response is also deterministic. In contrast, by estimating the past residual disturbances and then including them in the Hankel-matrix equalities, (5.19) and (5.21) take the influence of unstructured disturbances into account.

Second, the regularized predictor (5.21) serves as an interpolation between the causal predictor (5.19) and the subspace predictor (2.29). When $\lambda_g = 0$, the regularized predictor (5.21) recovers the causal predictor (5.19). Conversely, as $\lambda_g \to +\infty$, it converges to the subspace predictor (2.29). Note that the λ_g is proportional to the robustness bound β^j. If λ_g is too small, the prediction heavily depends on the estimation, making

it sensitive to estimation quality. Conversely, if λ_g is too large, the uncertainty of the unstructured disturbance is not effectively propagated to the output. In practice, λ_g can be gradually increased from $\lambda_g = 0$ to optimize the performance of the regularized predictor (5.21).

Among the three predictors, (5.19) is causal. This follows from Proposition 5.3 and Corollary 5.1, which establish that $(\tilde{U}, \hat{V}, Y)_{[0,N-1]}$ is a trajectory of a stochastic LTI system. By Proposition 4.1, the output of a stochastic LTI system is non-anticipated by input. In contrast, the predictions from (5.20) and (2.29) are not necessarily causal.

5.3. Case Study: A Real-World Residential Building

To demonstrate the effectiveness of the proposed stochastic prediction method, we apply it to a dataset from a real-world residential building located in Norway, operated by Norwegian University of Science and Technology (NTNU). For more details, see Reinhardt et al., 2025. The results presented in this section are based on a dataset provided by NTNU (Reinhardt et al., 2025) and a data preprocessing code developed by Kaupmann, 2023 as part of his master's thesis.

In cold climates, heating systems contribute significantly to total energy consumption. Efficient management of these systems is therefore crucial for achieving energy savings in residential buildings. However, modeling heating systems is challenging due to the complex interactions among various components and the stochastic nature of disturbances. Many of these disturbances, such as occupant behavior and certain weather conditions, are often not measured or statistically modeled.

To encourage widespread adoption, prediction tools for the thermal response of residential buildings need to be cost-effective. The proposed data-driven stochastic prediction method, combined with residual disturbance estimation, effectively addresses this need. It relies solely on measurements of the heating system's inputs and outputs, eliminating the need for additional sensors or exogenous disturbance predictions. Moreover, its data-driven nature bypasses the system identification step for individual residential buildings, which is typically time-consuming. Importantly, the method enables the prediction of the stochastic behavior of the thermal response, which further allows for the use of stochastic control tools.

5.3.1. Real-World Dataset and Preprocessing

The residence in the NTNU dataset is divided into four distinct volumes, each with different thermal capacities and different thermal losses to the environment. Each volume has a temperature sensor and a heating device connected to heat pumps. Heating power is controlled indirectly by setting the unit to ON/OFF ($O_i \in \{0, 1\}$), selecting a target temperature ($T_i^{\text{set}} \in \mathbb{I}_{[16,31]}$ [°C]), and adjusting the fan speed level ($F_i \in \mathbb{I}_{[1,7]}$) to promote air circulation. These variables form the manipulable input vector

$$u \doteq (O, T^{\text{set}}, F)_{[1,4]} \in \mathbb{R}^{12}. \tag{5.23}$$

In addition, the temperatures $T_i \in \mathbb{R}^+$ [°C] in the four volumes are considered as the output

$$y \doteq [T_1, \cdots, T_4]^\top \in \mathbb{R}^4. \tag{5.24}$$

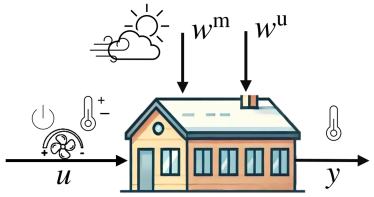

Figure 5.2.: Conceptual block diagram of a residential building with input u as in (5.23), output y as in (5.24), measured disturbance w^{m} as in (5.25), and unmeasured disturbance w^{u}

Moreover, the solar azimuth angle θ_a [°], solar zenith angle θ_z [°], outdoor air temperature $T^{\mathrm{out}} \in \mathbb{R}^+$ [°C], wind speed v_w [m/s], cloud opacity τ_c [-], and global horizontal irradiance I_G [W/m^2] are measured disturbances. Specifically, we denote the measured disturbance as

$$w^{\mathrm{m}} \doteq [\theta_a, \theta_z, T^{\mathrm{out}}, v_w, \tau_c, I_G]^\top, \tag{5.25}$$

and the unmeasured disturbance as w^{u}. The dataset provides the real measurements of inputs, outputs, and the measured disturbances from May 11^{th}, 2021, 00:00, to December 29^{th}, 2021, 00:00, with a sampling time of 15 minutes.

Figure 5.2 presents a conceptual block diagram of the residential building, illustrating the relations among input u, output y, measured disturbance w^{m}, and unmeasured disturbance w^{u}.

We employ the stochastic data-driven predictors outlined in Section 5.2 to forecast future thermal response based on historical data. At each prediction instant, we use the data from the most recent 30 days ($T = 2880$) to estimate the past realizations $\hat{v}^{\mathrm{d}}$ and the future prediction $\hat{V}$ of the residual disturbance as described in Section 5.2. Combining the measured data and the estimated residual disturbances, we predict the stochastic thermal response for the upcoming 24 hours (prediction horizon $N = 96$) by (5.20) with different λ_g and ℓ.

We assess the impact on choosing different structured disturbances. Specifically, we consider three schemes:

i) no structured disturbance $w = []$,

ii) only solar angles as the structured disturbance $w = [\theta_a, \theta_z]^\top$,

iii) and all measured disturbances as the structured disturbance $w = w^{\mathrm{m}}$.

The process of disturbance estimation and output prediction is recursively applied at time instants from November 8, 2021, 00:00 to December 27, 2021, 00:00, every 6 hours ($\Delta = 24$), resulting in a total of $N_s = 200$ test scenarios during the typical winter period.

At each prediction instant $t\Delta$, we denote the predicted random variables at time $t\Delta + k$ as $M(k \,|\, t\Delta)$ for $M \in \{U, W, \hat{V}, \hat{Y}\}$. In this section, we focus on the propagation of uncertainties arising from unstructured disturbances. To simplify the exposition, we assume future inputs $U(k \,|\, t\Delta)$ and structured disturbances $W(k \,|\, t\Delta)$ are deterministic. In other words, we consider that the future structured disturbances are perfectly predicted. Hence, the basis (5.18) can be further simplified with $\phi_W = \{1\}$

$$\{\phi\}_{j=0}^{L-1} \doteq \{1\} \bigcup \{\xi_1(k), \cdots, \xi_{n_y}(k)\}_{k=0}^{N-1}, \quad L = 1 + Nn_y. \tag{5.26}$$

Then the input components $\hat{V}$, U, and W admits exact PCEs in this basis as follows

$$\hat{V}(k\,|\,t\Delta)=\sum_{j=0}^{L-1}\hat{\mathrm{v}}^{j}(k\,|\,t\Delta)\,\phi^{j}(\boldsymbol{\xi}),\quad \hat{\mathrm{v}}^{0}(k\,|\,t\Delta)=\hat{m}_{v},\ \mathcal{H}_{1}\left(\hat{\mathrm{v}}^{[(k-1)n_y,\,kn_y]}(k\,|\,t\Delta)\right)=\left(\hat{\Sigma}_{v}\right)^{1/2},$$
$$\tilde{U}(k\,|\,t\Delta)=\sum_{j=0}^{L-1}\tilde{\mathrm{u}}^{j}(k\,|\,t\Delta)\,\phi^{j}(\boldsymbol{\xi}),\quad \tilde{\mathrm{u}}^{j}(k\,|\,t\Delta)=\delta^{0j}\tilde{u}(k\,|\,t\Delta),\quad \forall k\in\mathbb{I}_{[0,N-1]},$$

with $\tilde{U}=(U,W)$, $\tilde{\mathrm{u}}^{j}=(\mathrm{u}^{j},\mathrm{w}^{j})$, and $\tilde{u}=(u,w)$. In addition, the predicted outputs $\hat{Y}(k\,|\,t\Delta)$ also admits an exact PCE in (5.26) according to Lemma 3.3.

Basis (5.26) consists of independently distributed terms, specifically $\{1\}$ and i.i.d. $\xi(k)$. The fourth-order moments of $\hat{Y}(k\,|\,t\Delta)$ can then be determined using the simplified calculation (3.18). Consequently, offline computation only requires evaluating the second and fourth moments of $\xi(k)$.

5.3.2. Numerical Validation

At each prediction instant $t\Delta$, the predicted output at time $t\Delta+k$ is denoted as $\hat{Y}(k\,|\,t\Delta)$, with its corresponding predicted PCE coefficients given by $\hat{\mathrm{y}}^{j}(k\,|\,t\Delta)$. Note that $\hat{\mathrm{y}}^{0}(k\,|\,t\Delta)$ corresponds to the predicted mean of $\hat{Y}(k\,|\,t\Delta)$. To quantify the accuracy of the predicted means, we compute the Root Mean Square Error (RMSE) between the true output $y(t\Delta+k)$ and its predicted mean $\hat{\mathrm{y}}^{0}(k\,|\,t\Delta)$ across all samples

$$\overline{\mathrm{RMSE}}\doteq\sqrt{\sum_{t=0}^{N_s-1}\sum_{k=0}^{N-1}\frac{\|\hat{\mathrm{y}}^{0}(k\,|\,t\Delta)-y(t\Delta+k)\,\|^{2}}{NN_s}}.\tag{5.27}$$

Additionally, to assess the accuracy of the stochastic prediction, we compute the percentage of ground-truth temperature samples that fall within the predicted 90% confidence intervals, denoted as the averaged coverage rate $\bar{c}_{90\%}$. With the indicator function $I(p)$ defined as $I(p)=1$ if statement p is true, and $I(p)=0$ otherwise, the averaged coverage rate is given by

$$\bar{c}_{90\%}\doteq\sum_{t=0}^{N_s-1}\sum_{i=1}^{n_y}\sum_{k=0}^{N-1}\frac{I\left(|y_i(t\Delta+k)-\hat{\mathrm{y}}_{i}^{0}(k\,|\,t\Delta)|\le r_{i,90\%}(k\,|\,t\Delta)\right)}{NN_sn_y}.\tag{5.28a}$$

The predicted 90% confidence interval for $y_i(t\Delta+k)$ is centered around the predicted mean $\hat{\mathrm{y}}_{i}^{0}(k\,|\,t\Delta)$ with the radius chosen from three variants as introduced in Section 3.4.1: the original Chebyshev bound, the Chebyshev bound incorporating fourth-order central moments, and the bound for Gaussian distributions. The radius is calculated as follows

$$r_{i,90\%}(k\,|\,t\Delta)\doteq\begin{cases}\sqrt{\frac{\mu_2\left(\hat{Y}_i(k\,|\,t\Delta)\right)}{1-90\%}}\approx 3.16\sqrt{\mu_2}, & \text{(Chebyshev }\mu_2\text{)}\\ \sqrt[4]{\frac{\mu_4\left(\hat{Y}_i(k\,|\,t\Delta)\right)}{1-90\%}}\approx 1.78\sqrt[4]{\mu_4}, & \text{(Chebyshev }\mu_4\text{)}\\ F_{\mathcal{N}}^{-1}\left(\frac{1+\gamma}{2}\right)\mu_2\left(\hat{Y}_i(k\,|\,t\Delta)\right)\approx 1.64\sqrt{\mu_2} & \text{(Gaussian).}\end{cases}\tag{5.28b}$$

Additionally, we define the averaged interval radius as

$$\bar{r}_{90\%}\doteq\sum_{t=0}^{N_s-1}\sum_{i=1}^{n_y}\sum_{k=0}^{N-1}\frac{r_{i,90\%}(k\,|\,t\Delta)}{NN_sn_y}.\tag{5.28c}$$

Table 5.1.: Comparison of the prediction results for different ℓ with $\lambda_g = 0$ and $w = []$.

ℓ	$\overline{\text{RMSE}}$	Chebyshev μ_2		Chebyshev μ_4		Gaussian	
		$\bar{c}_{0.9}$ [−]	$\bar{r}_{0.9}$ [°C]	$\bar{c}_{0.9}$ [−]	$\bar{r}_{0.9}$ [°C]	$\bar{c}_{0.9}$ [−]	$\bar{r}_{0.9}$ [°C]
2	0.9663	96.29%	2.377	92.35%	1.793	81.67%	1.236
4	0.9124	96.99%	2.426	93.63%	1.824	84.78%	1.262
6	0.8850	97.22%	2.422	94.20%	1.819	85.93%	1.260
8	0.8740	97.30%	2.407	94.13%	1.807	85.78%	1.252
10	0.8586	97.29%	2.360	94.34%	1.772	85.96%	1.228
12	0.8745	97.11%	2.324	93.77%	1.745	85.01%	1.209
14	0.8574	97.19%	2.273	93.67%	1.706	84.90%	1.182
16	0.8481	97.19%	2.241	93.66%	1.682	84.74%	1.166
18	0.8468	97.07%	2.178	93.30%	1.635	83.95%	1.133
20	0.8457	96.97%	2.143	92.95%	1.610	83.39%	1.115

Table 5.2.: Comparison of the prediction results for different λ_g with $\ell = 10$ and $w = []$.

λ_g	$\overline{\text{RMSE}}$	Chebyshev μ_2		Chebyshev μ_4		Gaussian	
		$\bar{c}_{0.9}$ [−]	$\bar{r}_{0.9}$ [°C]	$\bar{c}_{0.9}$ [−]	$\bar{r}_{0.9}$ [°C]	$\bar{c}_{0.9}$ [−]	$\bar{r}_{0.9}$ [°C]
0	0.8586	97.29%	2.360	94.34%	1.772	85.96%	1.228
0.01	0.8583	97.28%	2.358	94.33%	1.770	85.93%	1.227
0.1	0.8563	97.24%	2.339	94.21%	1.756	85.73%	1.217
1	0.8431	96.61%	2.181	93.19%	1.638	83.70 %	1.134
2	0.8369	95.97%	2.050	91.99%	1.540	81.69%	1.066
5	0.8363	94.06%	1.784	88.41%	1.340	76.11%	0.928
10	0.8497	90.49%	1.515	82.85%	1.138	68.48%	0.788
100	0.9763	45.96%	0.526	36.42%	0.364	26.10 %	0.273
$+\infty$	1.0676	0	0	0	0	0	0

The coverage rate and confidence interval radius offer two complementary perspectives for evaluating stochastic predictions. A higher coverage rate indicates a more robust prediction, as it means that more actual values fall within the predicted intervals. However, a wider confidence interval corresponds to a more conservative prediction, as it covers a broader range of potential outcomes. Thus, there is a trade-off between robustness and conservatism: increasing the coverage rate generally leads to a wider confidence interval.

Table 5.1 compares the aforementioned performance metrics for the stochastic predictions (5.20) with different ℓ, which reflects the system order of the underlying ARX models. In this comparison, the regularization factor is fixed at $\lambda_g = 0$ and no structured disturbances are considered ($w = []$). As shown in Table 5.1, when ℓ increases from 2 to 10, $\overline{\text{RMSE}}$ decreases by approximately $0.1°C$. For $\ell \geq 10$, further increases in ℓ yield only slight improvements in the mean prediction accuracy. Additionally, the coverage rates and confidence interval radius remain relatively constant across different values of ℓ. Among the variants of (5.28b), the fourth-order Chebyshev approach generally yields a tighter confidence bound than the original Chebyshev's approach and achieves a higher coverage probability than the Gaussian setting. Thus, the fourth-order Chebyshev approach effectively balances robustness, as indicated by the coverage probability, and conservatism, as reflected in the interval radius, in this specific example.

Table 5.2 presents the prediction results of the regularized predictor (5.20) for differ-

Table 5.3.: Comparison of the prediction results for different choices of structured disturbances with $\lambda_g = 0$ and $\ell = 10$.

w	$\overline{\text{RMSE}}$	Chebyshev μ_2		Chebyshev μ_4		Gaussian	
		$\bar{c}_{0.9}$ [−]	$\bar{r}_{0.9}$ [°C]	$\bar{c}_{0.9}$ [−]	$\bar{r}_{0.9}$ [°C]	$\bar{c}_{0.9}$ [−]	$\bar{r}_{0.9}$ [°C]
[]	0.8586	97.29%	2.360	94.34%	1.772	85.96%	1.228
$[\theta_a, \theta_z]^\top$	0.8016	96.90%	2.067	93.40%	1.558	84.02%	1.075
$[\theta_a, \theta_z, T^{\text{out}}, v_w, \tau_c, I_G]^\top$	0.7846	95.97%	1.784	91.11%	1.351	78.62%	0.928

ent values of λ_g. Note that when $\lambda_g = 0$, the regularized predictor (5.20) reduces to the causal predictor (5.19), whereas for $\lambda_g = \infty$, (5.20) simplifies to the subspace predictor (2.29). As λ_g increases, $\overline{\text{RMSE}}$ decreases slightly until $\lambda_g = 5$, beyond which it begins to increase. This indicates that tuning λ_g can enhance prediction accuracy. In addition, the coverage rate decreases as λ_g increases, consistent with the discussion in Section 5.2 that the predictor tends to propagate less uncertainty as λ_g increases. Especially for $\lambda_g = \infty$, the subspace predictor (2.29) is a deterministic prediction and fails to account for uncertainties induced by unmeasured disturbances.

Table 5.3 presents a comparison of the results for different choices of structured disturbances. The solar azimuth angle θ_a and the zenith angle θ_z can be perfectly predicted for any given time and location. Including the solar angles as structured disturbances improves $\overline{\text{RMSE}}$ by approximately $0.05°C$, while the average radius of the confidence intervals decreases by about 0.2–$0.3°C$. However, the outdoor air temperature T^{out}, wind speed v_w, cloud opacity τ_c, and global horizontal irradiance I_G depend on weather conditions, making them costly to measure and difficult to forecast accurately. Even with perfect forecasts of these values, incorporating additional weather conditions improves $\overline{\text{RMSE}}$ by only about $0.02°C$, suggesting that the potential benefits may not justify the effort.

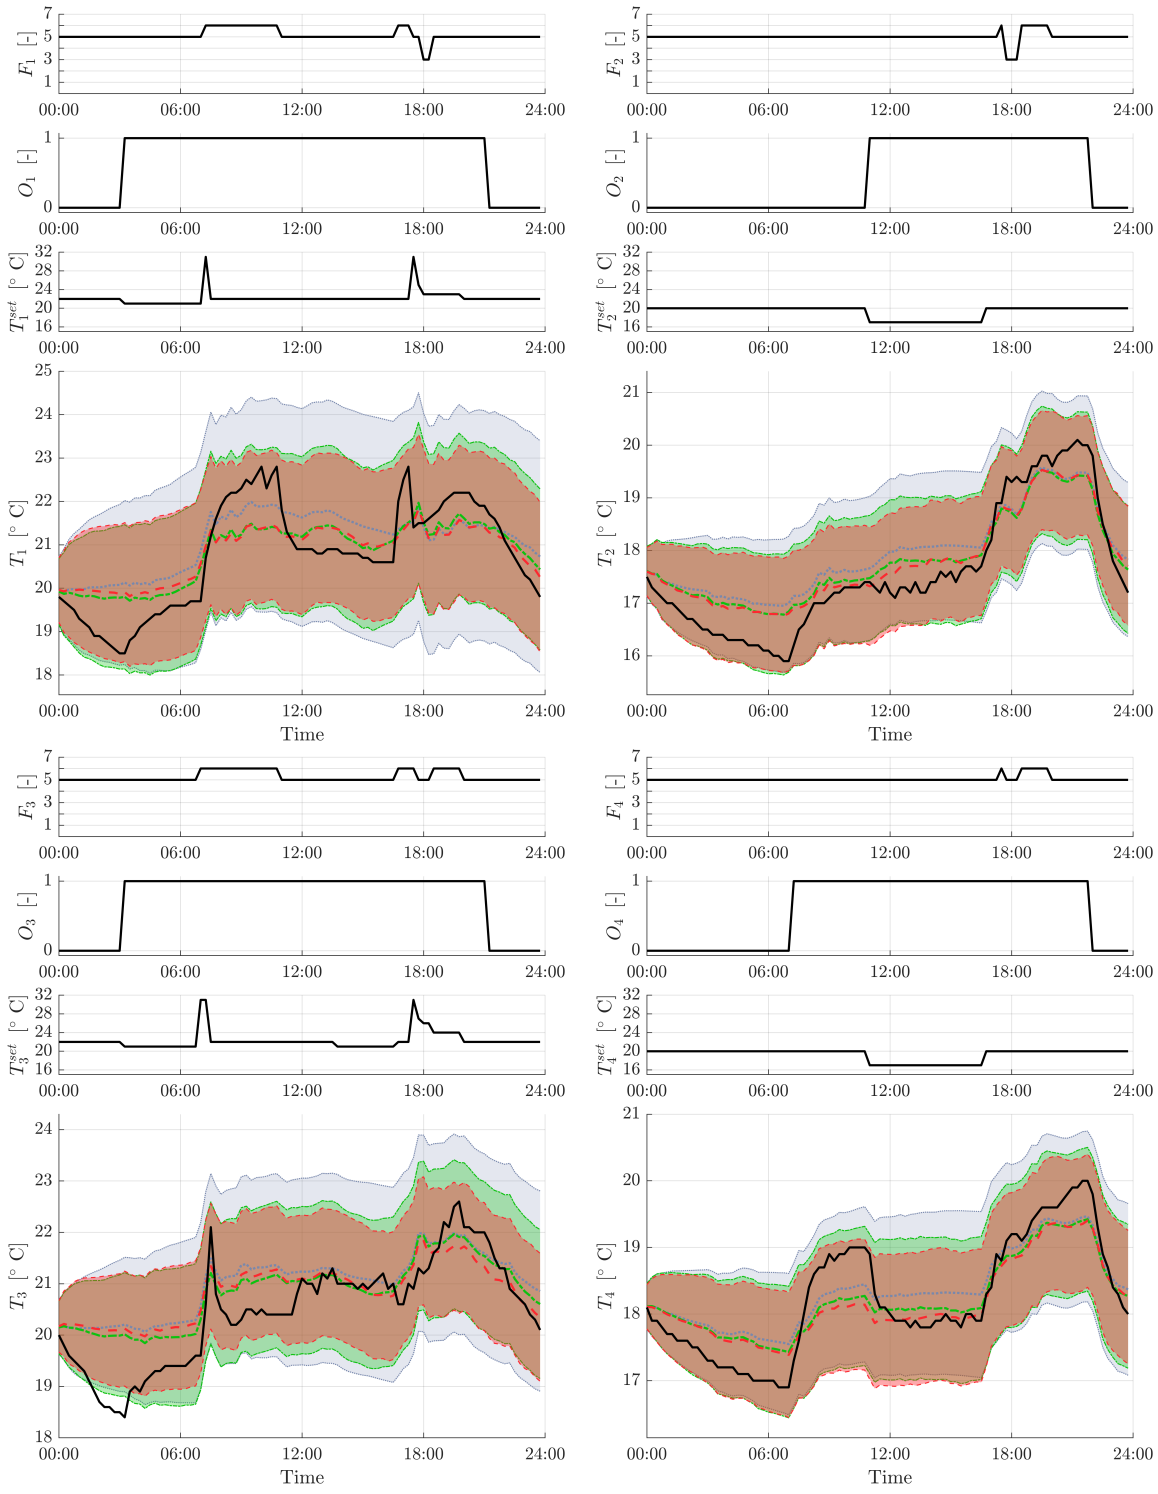

Figure 5.3.: Comparison of the true and the predicted thermal responses for November 17, 2021. Black-solid lines show measured input and output trajectories. Blue-dotted, green dash-dotted, and red-dashed lines represent predicted mean trajectories with 90% confidence intervals for $w = []$, $w = [\theta_a, \theta_z]^\top$, and $w = [\theta_a, \theta_z, T^{\text{out}}, v_w, \tau_c, I_{\text{G}}]^\top$, respectively.

Figure 5.3 illustrates the comparison of the true and predicted output responses for different choices of structured disturbances. Specifically, in black solid lines, we plot the true input and output trajectories of the four volumes for the whole day of November 17, 2021. Furthermore, we use blue-dotted, green dash-dotted, and red-dashed lines to show the predicted mean trajectories for $w = []$, $w = [\theta_a, \theta_z]^\top$, and $w = [\theta_a, \theta_z, T^{\text{out}}, v_w, \tau_c, I_{\text{G}}]^\top$, respectively. The 90% confidence intervals, calculated using fourth-order moments, are shown as shaded areas around the mean predictions.

Overall, incorporating solar azimuth and zenith angles as structured disturbances (green lines) significantly improves prediction accuracy and reduces the confidence interval radius. In comparison, adding weather conditions as structured disturbances (red lines) does not significantly improve the prediction, which aligns with the results in Table 5.3.

5.4. Summary

In this chapter, we classified exogenous disturbances into two categories: structured disturbances, which have known external sources and are therefore supported by past measurements and statistical forecasts; and unstructured disturbances, which originate from unspecified sources and are consequently unmeasured or unmodeled. For the latter, we condensed them as one residual disturbance and proposed a least-squares approach to estimate its past realizations and the associated statistics. Furthermore, we developed both parametric and non-parametric representations for the stochastic behavior of dynamical systems subject to residual disturbances.

Based on the stochastic fundamental lemma, we introduced a regularized data-driven stochastic predictor to propagate uncertainties from unknown disturbances to system outputs. By incorporating regularization techniques, the regularized predictor enhances robustness against errors in estimating past disturbances. Additionally, leveraging PCE representations, we computed the fourth-order moments of the stochastic predictions. This enabled us to quantify the uncertainty through confidence intervals derived from the fourth-order Chebyshev inequality. As demonstrated in the example, the fourth-order Chebyshev interval generally provides a tighter bound than the original Chebyshev interval. Moreover, the example using real building data demonstrates the practical applicability and effectiveness of the proposed techniques.

As discussed in Remarks 5.1, 5.2, and 5.3, future work may extend these methods to account for measurement noise, nonlinear model mismatch, and non-i.i.d. residual disturbances.

6. Data-Driven Stochastic Optimal Control

The previous chapter focused on data-driven propagation of uncertainties in stochastic LTI systems, laying the groundwork for predicting stochastic behaviors based on historical data. In Chapters 6–7, we transition from representation and prediction to control, focusing on optimal control in the presence of exogenous disturbances. In Chapter 6, we address the optimal control of stochastic discrete-time LTI systems with structured disturbances W, building on the results from Chapter 4. In Chapter 7, we consider stochastic systems with residual disturbance V, as introduced in Chapter 5.

In this chapter, Section 6.1 first presents the stochastic OCP of interest in a model-based framework. We discuss control policies and chance constraints, which are essential components of *stochastic* optimal control. In Section 6.2, we provide a data-driven reformulation of the proposed OCP and discuss the conditions under which the model-based and data-driven OCPs are equivalent, based on our previous paper (Pan et al., 2023c). Additionally, we establish sufficient conditions for the optimality of the affine disturbance-feedback policy and propose a simplified Toeplitz-structured policy to reduce computational complexity. These results constitute a novel contribution and have not yet been published.

Section 6.3 extends the analysis to disturbances with Gelbrich ambiguity, enhancing distributional robustness against uncertainty in future disturbance predictions. This result builds upon and extends the paper (Pan et al., 2023a). Finally, in Section 6.4, we demonstrate the effectiveness of the proposed data-driven optimal control scheme through numerical simulations.

6.1. Model-Based Stochastic Optimal Control

In this chapter, we consider the stochastic discrete-time LTI systems with structured disturbances W

$$X(k+1) = AX(k) + BU(k) + EW(k), \quad X(0) = x_{\text{ini}} \tag{6.1a}$$
$$Y(k) = CX(k) + DU(k) + FW(k), \tag{6.1b}$$

with state $X(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_x})$, input $U(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_u})$, output $Y(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y})$, and disturbance $W \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_w})$ for all $k \geq 0$. Here, we assume that $(A, [B, E])$ is a controllable pair and that (A, C) is an observable pair. We suppose that the matrices $(A, [B, E], C, [D, F])$ are unknown. Throughout this chapter, we consider that the future stochastic disturbances $W(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_w})$ for $k \geq 0$ are not necessarily identically distributed, independent, or Gaussian.

At time instant 0, we consider all past inputs, outputs, and disturbances have been realized, i.e., $U(k) = u(k)$, $Y(k) = y(k)$, and $W(k) = w(k)$ almost surely for $k < 0$. In addition, we consider $X(k) = x(k)$ almost surely for $k \leq 0$. For future inputs $U(k)$, $k \geq 1$, only historical information of stochastic uncertainties is available. Therefore, we assume that the input process is adapted to a filtration $(\mathcal{F}_k)_{k\in\mathbb{N}}$, cf. Definition 3.8, with

$$U(k) \in L^2(\Omega, \mathcal{F}_k, \mu; \mathbb{R}^{n_u}), \quad \mathcal{F}_k = \begin{cases} \sigma(x(0)), & k = 0, \\ \sigma(W_{[0,k-1]}), & k \geq 1, \end{cases} \tag{6.2}$$

where $\mathcal{F}_k$ represents the σ-algebra generated by all unrealized disturbances up to time k. Since $X(0)$ has been realized for $i = 0$ and the σ-algebra generated by a realized random variable is trivial, we have $\sigma(x(0)) = \{\emptyset, \Omega\}$. This assumption reflects a causality condition: the control input at time k depends only on information available up to time $k-1$. Such a setting is typical in real-time control scenarios, where future disturbances are not observable in advance.

Moreover, by Corollary 4.1, the state X and the output Y are also adapted to the filtration (6.2). Thus,

$$X(k) \in L^2(\Omega, \mathcal{F}_k, \mu; \mathbb{R}^{n_x}), \quad Y(k) \in L^2(\Omega, \mathcal{F}_{k+1}, \mu; \mathbb{R}^{n_y}), \quad k \in \mathbb{N}.$$

We begin our analysis with the following stochastic OCP in the conventional model-based setting. At time instant 0, given the realized initial state x_{ini} and a stochastic prediction of future disturbances $W_{[0,N-1]}$, we consider the following OCP

$$\min_{\pi, U, Y, X} \sum_{k=0}^{N-1} \mathbb{E}\left[\|Y(k)\|_Q^2 + \|U(k)\|_R^2\right] \tag{6.3a}$$

$$\text{subject to } \forall k \in \mathbb{I}_{[0,N-1]},$$

$$X(k+1) = AX(k) + BU(k) + EW(k), \quad X(0) = x_{\text{ini}}, \tag{6.3b}$$

$$Y(k) = CX(k) + DU(k) + FW(k), \tag{6.3c}$$

$$U(k) = \pi_k\left(W_{[0,k-1]}\right), \tag{6.3d}$$

$$\mathbb{P}[a_{u,i}^\top U(k) \leq 1] \geq \gamma_u, \quad \forall i \in \mathbb{I}_{[1,N_u]}, \tag{6.3e}$$

$$\mathbb{P}[a_{y,i}^\top Y(k) \leq 1] \geq \gamma_y, \quad \forall i \in \mathbb{I}_{[1,N_y]}. \tag{6.3f}$$

The objective function is the expected value of a quadratic form with $Q \succeq 0$ and $R \succ 0$. Moreover, the filtration (6.2) requires that the input is determined from its past disturbances. Hence, we consider a causal policy in (6.3d) with $\pi_k : L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{kn_w}\right) \to L^2(\Omega, \mathcal{F}_k, \mu; \mathbb{R}^{n_u})$. Instead of optimizing over deterministic future inputs, we optimize over stochastic input policies that provide feedback maps based on future disturbance realizations. After solving the OCP, the input can be determined in real-time using the measurements of realized disturbances, i.e.,

$$u(k) = \pi_k^\star(w_{[0,k-1]}), \quad k \in \mathbb{I}_{[0,N-1]}.$$

Chance constraints are imposed as individual half-space constraints using $a_{u,i} \in \mathbb{R}^{n_u}$ for $i \in \mathbb{I}_{[1,N_u]}$ and $a_{y,i} \in \mathbb{R}^{n_y}$ for $i \in \mathbb{I}_{[1,N_y]}$, with respective probability levels γ_u and γ_y, as given in (6.3e)-(6.3f).

Causal Affine Disturbance-Feedback Policy

Instead of general nonlinear policies, we consider causal and affine disturbance-feedback policies

$$U(k) = \bar{u}(k) + \sum_{j=0}^{k-1} K_{k,j} W(j). \tag{6.4a}$$

Consider a finite horizon $N \in \mathbb{N}^+$, the above policies can be rephrased in a condensed form

$$U_{[0,N-1]} = \bar{u}_{[0,N-1]} + \mathbf{K}_N W_{[0,N-1]}, \quad \mathbf{K}_N = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ K_{1,0} & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \ddots & \mathbf{0} & \vdots \\ K_{N-1,0} & \cdots & K_{N-1,N-2} & \mathbf{0} \end{bmatrix}. \tag{6.4b}$$

In Section 6.2.2, we will show that restricting to affine functions still preserves optimality for the considered optimal control problem.

Lemma 6.1 (Equivalence of affine parametrizations (Goulart et al., 2006)): Consider the stochastic dynamics (6.1). For any parametrization (6.4) with given $\bar{u}(k), K_{k,j}, \forall k \in \mathbb{N}, j \in \mathbb{I}_{[0,k-1]}$, there exists $\bar{u}^X(k), K^X_{k,j}, \forall k \in \mathbb{N}, j \in \mathbb{I}_{[0,k]}$, satisfying

$$U(k) = \bar{u}^X(k) + \sum_{j=0}^{k} K^X_{k,j} X(j), \tag{6.5}$$

which yields the same state, input, and output sequences for all realizations of the disturbance sequence $(W(k,\omega))_{k\in\mathbb{N}}$ and for all $\omega \in \Omega$.

Though the two parametrizations are equivalent in the sense of generating the same system trajectories, their numerical properties are different. Since $\bar{u}^X(k)$, $K^X_{k,j}$, and $X(j)$ are all decision variables, (6.5) is bilinear, whereas (6.4) is linear in terms of the decision variables $\bar{u}(k)$ and $K_{k,j}$. Therefore, the affine disturbance feedback policy (6.4) allows for further construction of convex optimal control problem. Another reason for not using state-feedback policies is that, in the data-driven setting, state information is often unavailable.

We note that one special case of (6.5)

$$U(k) = \bar{u}^X(k) + K^X_{k,k} X(k) \tag{6.6}$$

is usually considered in the literature to induce an explicit propagation of the state covariances, cf. Farina et al., 2013,

$$\Sigma[X(k+1)] = (A + BK^X_{k,k})\Sigma[X(k)](A + BK^X_{k,k})^\top + E\Sigma[W(k)]E^\top.$$

Compared to (6.4) and (6.5), (6.6) has a relatively limited parameterization, which may lead to suboptimality. Similar to (6.5), since both $K^X_{k,k}$ and $X(k)$ are decision variables, the resulting optimization is also non-convex (Farina et al., 2013). Therefore, in tube-based stochastic optimal control, all $K^X_{k,k}$ are typically chosen as fixed gains before optimization, which further introduces suboptimality (Mesbah, 2016).

Chance Constraints

By applying the refined Cantelli inequality (3.19), we obtain a conservative reformulation of the chance constraints (6.3e)-(6.3f)

$$a_{u,i}^\top \mathbb{E}[U(k)] + \alpha(\gamma_u)\sqrt{\operatorname{trace}\left(a_{u,i}^\top \Sigma[U(k)] a_{u,i}\right)} \leq 1, \quad \forall i \in \mathbb{I}_{[1,N_u]}, \tag{6.7a}$$

$$a_{y,i}^\top \mathbb{E}[Y(k)] + \alpha(\gamma_y)\sqrt{\operatorname{trace}\left(a_{y,i}^\top \Sigma[Y(k)] a_{y,i}\right)} \le 1, \quad \forall i \in \mathbb{I}_{[1,N_y]}. \tag{6.7b}$$

We note that, by Lemma 3.7, the sets of all real-valued random variables with Gaussian distribution (Definition 3.4), symmetric distributions (Definition 3.9), or with symmetric linear unimodal distributions (Definitions 3.9–3.10) are closed under linear operations. Note that, due to the linear dynamics (6.1) and the affine policy (6.4), both $a_{u,i}^\top U(k)$ and $a_{y,i}^\top Y(k)$ can be obtained by applying linear operations to $W_{[0,N-1]}$. Hence, if $W_{[0,N-1]}$ possess one of these distribution types, then $a_{u,i}^\top U(k)$ and $a_{y,i}^\top Y(k)$ will also adhere to the same distribution type. As a result, as shown in (3.19), the Cantelli factor $\alpha(\gamma)$ in (6.7a)-(6.7b) can be improved compared to the arbitrary L^2 case.

6.2. Data-Driven Stochastic Optimal Control

Next, we apply the results from Chapter 4 to stochastic optimal control, providing a data-driven reformulation of OCP (6.3).

6.2.1. Reformulation via PCE

We first recall the PCE formulation of (6.1) with the following corollary of Lemma 3.3.

Corollary 6.1 (Exact PCEs): Suppose $W_{[0,N-1]}$ admits an exact PCE with the first L terms of $\{\phi^j(\xi)\}_{j=0}^\infty$, consider the stochastic LTI system (6.1) with $U_{[0,N-1]}$ determined from the affine causal policy (6.4), then (U, Y, X) admits exact PCEs with the first L terms of $\{\phi^j(\xi)\}_{j=0}^\infty$.

Using the basis $\{\phi^j(\xi)\}_{j=0}^\infty$, the dynamics of the PCE coefficients are given by

$$\mathsf{x}^j(k+1) = A\mathsf{x}^j(k) + B\mathsf{u}^j(k) + E\mathsf{w}^j(k), \quad \mathsf{x}^j(0) = \delta^{0j} x_{\text{ini}}, \ \forall j \in \mathbb{I}_{[0,L-1]}, \tag{6.8a}$$

$$\mathsf{y}^j(k) = C\mathsf{x}^j(k) + D\mathsf{u}^j(k) + F\mathsf{w}^j(k), \quad \forall k \in \mathbb{I}_{[0,N-1]}. \tag{6.8b}$$

Assumption 6.1 (Data availability): Consider a given realization trajectory $(u, w, y)_{[0,T-1]}^{\mathrm{d}}$ of (6.1). We suppose that $(u, w)_{[0,T-1]}^{\mathrm{d}}$ is persistently exciting of order $n_x + N + \ell$ with ℓ not smaller than the system lag of (6.1).

Let p and f denote the ranges $[-\ell, -1]$ and $[0, N-1]$, respectively, and follow the same partitioning of Hankel matrices as in (2.26). Let the stacked Hankel matrices be $\mathcal{H}_{\mathrm{p}} \doteq [\mathcal{H}_{u,\mathrm{p}}^\top, \mathcal{H}_{y,\mathrm{p}}^\top, \mathcal{H}_{w,\mathrm{p}}^\top]^\top$ and $\mathcal{H}_{\mathrm{f}} \doteq [\mathcal{H}_{u,\mathrm{f}}^\top, \mathcal{H}_{y,\mathrm{f}}^\top, \mathcal{H}_{w,\mathrm{f}}^\top]^\top$. Applying the stochastic fundamental lemma (Lemma 4.6), we obtain the following data-driven reformulation of OCP (6.3) in terms of PCE coefficients:

$$\min_{\bar{u}, \mathbf{K}_N, \mathsf{u}, \mathsf{y}, \mathsf{g}} \sum_{k=0}^{N-1} \sum_{j=0}^{L-1} \left(\|\mathsf{y}^j(k)\|_Q^2 + \|\mathsf{u}^j(k)\|_R^2\right) \|\phi^j\|^2 \tag{6.9a}$$

$$\text{subject to } \mathcal{H}_{\mathrm{p}} \mathsf{g}^j = \delta^{0j} [u_{\mathrm{p}}^\top, y_{\mathrm{p}}^\top, w_{\mathrm{p}}^\top]^\top, \quad \forall j \in \mathbb{I}_{[0,L-1]}, \tag{6.9b}$$

$$\mathcal{H}_{\mathrm{f}} \mathsf{g}^j = [\mathsf{u}_{\mathrm{f}}^{j\top}, \mathsf{y}_{\mathrm{f}}^{j\top}, \mathsf{w}_{\mathrm{f}}^{j\top}]^\top, \quad \forall j \in \mathbb{I}_{[0,L-1]}, \tag{6.9c}$$

$$\mathsf{u}_{\mathrm{f}}^j = \delta^{0j} \bar{u}_{\mathrm{f}} + \mathbf{K}_N \mathsf{w}_{\mathrm{f}}^j, \quad \forall j \in \mathbb{I}_{[1,L-1]}, \tag{6.9d}$$

$$a_{u,i}^\top \mathsf{u}^0(k) + \alpha(\gamma_u) \sqrt{\sum_{j=1}^{L-1} \|a_{u,i}^\top \mathsf{u}^j(k)\|^2 \|\phi^j\|^2} \le 1, \quad \forall i \in \mathbb{I}_{[1,N_u]}, k \in \mathbb{I}_{[0,N-1]}, \tag{6.9e}$$

$$a_{y,i}^{\top}\mathsf{y}^{0}\left(k\right)+\alpha(\gamma_{y})\sqrt{\sum_{j=1}^{L-1}\|a_{y,i}^{\top}\mathsf{y}^{j}\left(k\right)\|^{2}\|\phi^{j}\|^{2}}\leq 1,\quad\forall i\in\mathbb{I}_{[1,N_{y}]},k\in\mathbb{I}_{[0,N-1]},\tag{6.9f}$$

where δ^{0j} is the Kronecker delta, $\mathbf{K}_N$ collects all feedback gains $K_{k,i}$ as in (6.4b). Moreover, the stochastic fundamental lemma, i.e. Lemma 4.6, justifies the data-driven representation of the PCE dynamics (6.8) in (6.9b)–(6.9c). The Kronecker delta δ^{0j} in (6.9b) ensures that the PCE coefficients of the initial condition are zero for $j>0$, implying a deterministic initial condition. Causality and affiness of polices in (6.4) are stated in (6.9d). The chance constraints (6.7a)–(6.7b) are reformulated in (6.9e) by substituting the mean and covariance terms with PCE coefficients.

Observe that (6.9) is a second-order cone program. After eliminating the equality constraints, using interior-point methods to solve (6.9) typically requires at most $O(\sqrt{n_c})$ iterations, with a per-iteration complexity of $O(n_c n_d^2)$ (Lobo et al., 1998). Here, the per-iteration complexity

$$O(n_c n_d^2)\sim O(N^5),\quad n_c=N(N_u+N_y)\text{ and }n_d=Nn_u+n_un_wN(N-1)/2\tag{6.10}$$

scales up to $O(N^5)$, with n_c and n_d representing the number of second-order cone constraints and the number of decision variables in $(\bar{u}_{\mathrm{f}},\mathbf{K}_N)$, respectively.

To reduce this complexity, one may opt for a simplified feedback

$$U\left(k\right)=\bar{u}\left(k\right)+\sum_{j=0}^{k-1}K_{k-j}W\left(j\right),\quad k\in\mathbb{I}_{[0,N-1]},\quad\Rightarrow\quad U_{\mathrm{f}}=\bar{u}_{\mathrm{f}}+\mathbf{K}_N'W_{\mathrm{f}}\tag{6.11a}$$

with the condensed feedback matrix $\mathbf{K}_N'$ in a Toeplitz structure

$$\mathbf{K}_N'=\begin{bmatrix}\mathbf{0}&\mathbf{0}&\mathbf{0}&\cdots&\mathbf{0}\\K_1&\mathbf{0}&\mathbf{0}&\cdots&\mathbf{0}\\K_2&K_1&\mathbf{0}&\cdots&\mathbf{0}\\\vdots&\ddots&\ddots&\ddots&\mathbf{0}\\K_{N-1}&\cdots&K_2&K_1&\mathbf{0}\end{bmatrix}.\tag{6.11b}$$

This modification reduces the per-iteration complexity from $O(N^5)$ to $O(N^3)$, i.e.,

$$O\left(n_c(n_d^s)^2\right)\sim O(N^3),\quad n_c=N(N_u+N_y)\text{ and }n_d^s=Nn_u+(N-1)n_un_w,\tag{6.12}$$

where the number of decision variables in $(\bar{u}_{\mathrm{f}},\mathbf{K}_N')$, i.e. n_d^s, is substantially reduced due to the Toeplitz structure. Theoretically, this simplification introduces some degree of suboptimality compared to the full policy (6.4). However, as demonstrated in Section 6.4, this suboptimality is relatively minor, whereas the simplification significantly reduces computational complexity.

Remark 6.1 (Multiple-shooting implementation): Compared to the state-space model in (6.3), the equality constraint (6.9c) increases the computational burden due to the large dense Hankel matrices. To overcome this issue, O'Dwyer et al., 2023 suggests segmenting the prediction horizon into shorter intervals and using Hankel matrices of smaller dimensions. Furthermore, the solution pieces in consecutive intervals are coupled by continuity constraints. This idea resembles the classic concept of multiple shooting by

Bock et al., 1984 in the data-driven setting. Ou et al., 2023 further tailor this concept to the data-driven stochastic OCP (6.9). Moreover, combined with a moment matching strategy, one can show that the dimension of the PCE basis and the number of decision variables can be reduced substantially. For further details, we refer to Ou et al., 2023.

Remark 6.2 (Measurement noise compensation via regularizations): Considering that the outputs are corrupted by measurement noise, we can employ the regularization techniques from Coulson et al., 2019; Huang et al., 2021 to improve robustness in optimality. Specifically, in the original objective (6.9a), we can incorporate quadratic regularizations of the column selectors g^j for $j \in \mathbb{I}_{[0,L-1]}$ and introduce a slack variable $s \in \mathbb{R}^{\ell n_y}$ to account for measurement noises. Specifically, we consider

$$\sum_{k=0}^{N-1}\sum_{j=0}^{L-1}\left(\|\mathsf{y}^j(k)\|_Q^2 + \|\mathsf{u}^j(k)\|_R^2\right)\|\phi^j\|^2 + \lambda_g\sum_{j=0}^{L-1}\|\mathsf{g}^j\|^2 + \lambda_y\|s\|^2,$$

where $\lambda_g \in \mathbb{R}^+$ and $\lambda_y \in \mathbb{R}^+$ are the regularization parameters. Furthermore, the initial condition is relaxed by introducing s, which compensates for measurement noise in the initial outputs y_p

$$\mathcal{H}_\mathrm{p}\mathsf{g}^j = \delta^{0j}[u_\mathrm{p}^\top,\, y_\mathrm{p}^\top + s^\top,\, w_\mathrm{p}^\top]^\top.$$

As demonstrated in Huang et al., 2021, the quadratic regularization is equivalent to a min-max formulation, minimizing the worst-case cost for a bounded noise set that affects the data Hankel matrices and the initial trajectory. We refer to our recent work by Özmeteler et al., 2024 on data-driven optimal control of multi-energy systems, which simultaneously account for exogenous disturbances and measurement noise.

6.2.2. Equivalence and Optimality

Recall that, to ensure causality, the stochastic inputs $U_{[0,N-1]}$ must be adapted to the filtration $\{\mathcal{F}_k\}_{k\in\mathbb{I}_{[0,N-1]}}$ in (6.2). Although choosing affine parametrizations, as in (6.4), simplifies the numerical complexity of the OCP implementation, a remaining question is whether this choice sacrifices optimality for all possible parametrizations allowed in (6.2). Beyond establishing the equivalence between (6.3) and (6.9), the following results further confirm the optimality of affine parametrizations.

To this end, let $\mathbb{L}$ be the set of all PCE orders $\mathbb{I}_{[0,L-1]}$, we define the sets of the optimal input-output trajectories for the two OCPs as

$$\mathcal{S}_1 \doteq \left\{ (U_\mathrm{f}^\star, Y_\mathrm{f}^\star) \;\middle|\; \begin{array}{c} \text{given } x_\mathrm{ini}, W_{[0,N-1]},\ \exists\, X_\mathrm{f}^\star, \pi_\mathrm{f}^\star \text{ such that} \\ (X_\mathrm{f}^\star, U_\mathrm{f}^\star, Y_\mathrm{f}^\star, \pi_\mathrm{f}^\star) \text{ is optimal in OCP (6.3)} \\ \text{with the chance constraints (6.3e)–(6.3f)} \\ \text{replaced by (6.7a)–(6.7b)} \end{array} \right\}, \tag{6.13a}$$

$$\mathcal{S}_1' \doteq \left\{ (U_\mathrm{f}^\star, Y_\mathrm{f}^\star) \;\middle|\; \begin{array}{c} \text{given } x_\mathrm{ini}, W_{[0,N-1]},\ \exists\, X_\mathrm{f}^\star, \pi_\mathrm{f}^\star \text{ such that} \\ (X_\mathrm{f}^\star, U_\mathrm{f}^\star, Y_\mathrm{f}^\star, \pi_\mathrm{f}^\star) \text{ is optimal in OCP (6.3)} \\ \text{with the chance constraints (6.3e)–(6.3f)} \\ \text{replaced by (6.7a)–(6.7b)} \\ \text{and the control policy (6.3d) replaced by (6.4)} \end{array} \right\}, \tag{6.13b}$$

$$\mathcal{S}_2 \doteq \left\{ \left(\mathsf{u}_{\mathsf{f}}^{\star,\mathbb{L}}, \mathsf{y}_{\mathsf{f}}^{\star,\mathbb{L}} \right) \;\middle|\; \begin{array}{c} \text{given } (u,w,y)_{\mathsf{p}}, \tilde{\mathsf{w}}_{\mathsf{f}}^{\mathbb{L}}, k \in \mathbb{I}_{[0,N-1]}, \exists \mathsf{g}^{\star,\mathbb{L}}, \mathbf{K}_N^\star, \bar{u}^\star \text{ such that} \\ (\mathsf{u}_{\mathsf{f}}^{\star,\mathbb{L}}, \mathsf{y}_{\mathsf{f}}^{\star,\mathbb{L}}, \mathsf{g}^{\star,\mathbb{L}}, \mathbf{K}_N^\star, \bar{u}^\star) \text{ is optimal in OCP (6.9)} \end{array} \right\}. \quad (6.13c)$$

Note that $\mathcal{S}_1$ and $\mathcal{S}_1'$ are subsets of $L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{(n_u+n_y)N}\right)$ whereas $\mathcal{S}_2$ is a subset of $\mathbb{R}^{(n_u+n_y)NL}$. Given the PCE basis $\{\phi^j(\xi)\}_{j=0}^\infty$ on which OCP (6.9) is based, the PCE coefficients and L^2 random variables can be bijectively transformed through the linear maps Φ and Φ^{-1}, as defined in (4.16).

Theorem 6.1 (Equivalence stochastic OCPs): Consider the sets S_1, S_1', and S_2 defined in (6.13). Let Assumption 6.1 hold. Then, for any given initial condition $(u,y,w)_{[-\ell,-1]}$ for OCP (6.9), which is a ℓ-length realization trajectory of (6.1), there exists $x_{\text{ini}} \in \mathbb{R}^{n_x}$ for OCP (6.3) such that

$$\Phi\left(\mathcal{S}_2\right) = \mathcal{S}_1' = \mathcal{S}_1,$$

where $\Phi(\mathcal{S}_2)$ is the element-wise image of $\mathcal{S}_2$.

Proof. The proof consists of two parts: first, we prove the equivalence $\Phi\left(\mathcal{S}_2\right) = \mathcal{S}_1'$. Then, we prove the optimality of the affine and causal policy (6.4) by showing $\mathcal{S}_1' = \mathcal{S}_1$.

Since the system is observable and the measurements (u,w,y) are exact, $(u,y,w)_{[-\ell,-1]}$ uniquely determine the initial state x_{ini} in OCP (6.3) given that ℓ is not smaller than the minimal system lag. For any basis $\{\phi^j(\xi)\}_{j=0}^\infty$, where $W_{[0,N-1]}$ admits exact PCEs with L terms, by Corollary 6.1, all random variables in OCP (6.3) admit exact PCEs with at most L terms. Replacing all random variables with their PCEs, the objective (6.3a) is equivalent to (6.9a) due to the orthogonality of the basis $\{\phi^j(\xi)\}_{j=0}^\infty$. With Assumption 6.1, (6.9b)-(6.9c) exactly capture the stochastic LTI dynamics, cf. the stochastic fundamental lemma in Lemma 4.6. Moreover, (6.9d) exactly expresses the causal and affine policies (6.4) in PCE coefficients. The reformulation of chance constraints (6.7a)–(6.7b) to (6.9e)–(6.9f) is exact by expressing moments in PCE, cf. (3.7). Since the reformulation of OCP (6.3), where (6.3d), (6.3e), and (6.3f) are replaced by (6.4), (6.7a), and (6.7b), respectively, to obtain OCP (6.9) is exact, it follows that $\mathcal{S}_1' = \Phi(\mathcal{S}_2)$. Since the reformulation of OCP (6.3)—in which (6.3d), (6.3e), and (6.3f) are replaced by (6.4), (6.7a), and (6.7b), respectively—to obtain OCP (6.9) is exact, it follows that $\mathcal{S}_1' = \Phi(\mathcal{S}_2)$.

Next, we prove the optimality, i.e., $\mathcal{S}_1' = \mathcal{S}_1$. To establish that the affine and causal policies (6.4). are optimal among all causal policies, we use a contradiction argument. Specifically, we assume that the optimal causal policies of OCP (6.3) are not affine. We then show a contradiction by constructing a feasible trajectory that achieves a smaller value for the objective function.

To this end, we first construct an infinite-dimensional basis $\{\phi^j(\xi)\}_{j=0}^\infty$, which spans the space of all L^2 random variables that admit affine or nonlinear mappings from $W_{[0,N-1]}$. Assume that there exists a covariance decomposition $\Sigma[W_{[0,N-1]}] = \mathsf{M}_w\mathsf{M}_w^\top$, where $\mathsf{M}_w \in \mathbb{R}^{n_\xi \times N n_w}$ is of full column rank. Then, by Lemma 3.2, there exists a zero-mean, unit-covariance random variable $\xi \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\xi}\right)$ such that

$$W_{[0,N-1]} = \mathbb{E}[W_{[0,N-1]}] + \mathsf{M}_w\xi.$$

This implies $W_{[0,N-1]}$ admits an exact PCE in $\mathcal{P}_1(\xi)$ with $L \doteq 1 + n_\xi$ terms. Based on $\mathcal{P}_1(\xi)$, we can then construct an infinite-dimensional orthogonal polynomial basis $\{\phi^j(\xi)\}_{j=0}^\infty$ that covers all polynomial degrees by the Gram-Schmidt process (Witteveen

et al., 2006). Without loss of generality, we assume that the polynomials of degree not larger than 1, i.e. $\mathcal{P}_1(\xi)$, are the first L terms of $\{\phi^j(\xi)\}_{j=0}^{\infty}$.

Consider the optimal solution $(X_{\mathrm{f}}^\star, U_{\mathrm{f}}^\star, Y_{\mathrm{f}}^\star, \pi_{\mathrm{f}}^\star)$ of S_1 in (6.13). Suppose there exists $\tilde{k} \in \mathbb{I}_{[0,N-1]}$ such that $U^\star\left(\tilde{k}\right) = \pi_{\tilde{k}}^\star(W_{[0,\tilde{k}-1]})$, where $\pi_{\tilde{k}}^\star : L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{\tilde{k}n_w}\right) \to L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_u}\right)$ is a nonlinear function. As a result, $U\left(\tilde{k}\right)$ admits a PCE in the constructed $\{\phi^j(\xi)\}_{j=0}^{\infty}$ with

$$U^\star\left(\tilde{k}\right) = \sum_{j=0}^{\infty} \mathsf{u}^{\star,j}\left(\tilde{k}\right)\phi^j(\xi), \quad \mathsf{u}^{\star,L_1}\left(\tilde{k}\right) \neq 0, \quad \text{for some } L_1 \geq L. \tag{6.14}$$

This result similarly applies to Y and X, as stated in Corollary 6.1.

Furthermore, we introduce the projection $\Pi^{\mathbb{L}}$ of a random variable $Z \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_z}\right)$ expressed in the basis $\{\phi^j(\xi)\}_{j=0}^{\infty}$ onto a reduced basis containing a subset of functions $\mathbb{L} \subseteq \mathbb{N} \cup \{\infty\}$ as $\Pi^{\mathbb{L}} : L^2 \to L^2$

$$\Pi^{\mathbb{L}} : \; Z = \sum_{j=0}^{\infty} \mathsf{z}^j\phi^j \mapsto \bar{Z} = \sum_{j\in\mathbb{L}} \mathsf{z}^j\phi^j. \tag{6.15}$$

Now, consider truncating $(X^\star, Y^\star, U^\star)$ to the first L PCE terms via Π, with $\mathbb{L} \doteq \mathbb{I}_{[0,L-1]}$ from (6.15),

$$(\bar{X}, \bar{U}, \bar{Y})_{[0,N-1]} \doteq \Pi^{\mathbb{L}}\left((X^\star, U^\star, Y^\star)_{[0,N-1]}\right).$$

Crucially, since the first L basis functions are $\mathcal{P}_1(\xi)$, $\bar{U}_{[0,N-1]}$ follows an affine mapping from ξ and, consequently, from $W_{[0,N-1]}$. Moreover, this truncation preserves the causality of $\pi^\star$, and hence implicitly determines an affine and causal policy $\{\bar{\pi}_k\}_{k=0}^{N-1}$ from the nonlinear policy $\{\pi_k^\star\}_{k=0}^{N-1}$.

Since the trajectory tuple $(\bar{\mathsf{x}}^j, \bar{\mathsf{y}}^j, \bar{\mathsf{u}}^j)_{[0,N-1]}$, $k \in \mathbb{I}_{[0,N-1]}$ satisfies the PCE dynamics (6.8) for every $j \in \mathbb{I}_{[0,L-1]}$, it follows that $(\bar{X}, \bar{Y}, \bar{U})_{[0,N-1]}$ satisfies (6.3b)-(6.3c). Moreover, with the PCEs of $U^\star$, the PCE reformulations of the chance constraint (6.7a) read

$$a_{u,i}^\top \mathsf{u}^{\star,0}(k) + \alpha(\gamma_u)\sqrt{\sum_{j=1}^{\infty} \|a_{u,i}^\top \mathsf{u}^{\star,j}(k)\|^2 \|\phi^j\|^2} \leq 1, \quad \forall i \in \mathbb{I}_{[1,N_u]}.$$

Note that

$$\sum_{j=1}^{L-1} \|a_{u,i}^\top \bar{\mathsf{u}}^j(k)\|^2 \|\phi^j\|^2 = \sum_{j=1}^{L-1} \|a_{u,i}^\top \mathsf{u}^{j,\star}(k)\|^2 \|\phi^j\|^2 \leq \sum_{j=1}^{\infty} \|a_{u,i}^\top \mathsf{u}^{j,\star}(k)\|^2 \|\phi^j\|^2. \tag{6.16}$$

This implies that truncating the higher-order PCEs does not result in a violation of the chance constraints. Likewise, it holds for the constraints of Y. Therefore, the truncated variables $(\bar{U}, \bar{Y})$ still satisfy the chance constraints. Thus, $(\bar{X}, \bar{U}, \bar{Y})$ is a feasible trajectory of OCP (6.3).

In addition, the PCE reformulation of objective function (7.5a) satisfies

$$\mathbb{E}\left[\sum_{k=0}^{N-1}\left(\|U^\star(k)\|_R^2 + \|Y^\star(k)\|_Q^2\right)\right] = \sum_{k=0}^{N-1}\sum_{j=0}^{\infty}\left(\|\mathsf{y}^{\star,j}(k)\|_Q^2 + \|\mathsf{u}^{\star,j}(k)\|_R^2\right)\|\phi^j\|^2$$

$$= \sum_{k=0}^{N-1}\sum_{j=0}^{L-1}\left(\|\mathsf{y}^{\star,j}(k)\|_Q^2 + \|\mathsf{u}^{\star,j}(k)\|_R^2\right)\|\phi^j\|^2 + \sum_{k=0}^{N-1}\sum_{j=L}^{\infty}\left(\|\mathsf{y}^{\star,j}(k)\|_Q^2 + \|\mathsf{u}^{\star,j}(k)\|_R^2\right)$$
$$\geq \mathbb{E}\left[\sum_{k=0}^{N-1}\left(\|\bar{U}(k)\|_R^2 + \|\bar{Y}(k)\|_Q^2\right)\right] + \|\mathsf{u}^{\star,L_1}\left(\tilde{k}\right)\|_R^2,$$

since $\|\mathsf{u}^{\star,L_1}\left(\tilde{k}\right)\|_R^2 > 0$ holds for $R \succ 0$, cf.(6.14). Thus, $(\bar{X},\bar{U},\bar{Y})_{[0,N-1]}$ achieves a smaller objective value than $(X,U,Y)_{[0,N-1]}$, leading to a contradiction that proves the statement. □

6.3. Distributionally Robust Reformulation

Previously, we assumed that the random-variable prediction of future disturbances, W_f, is exact. However, exact knowledge of the underlying distribution of future disturbances W_f is often not realistic. To hedge against the uncertainty surrounding the disturbance statistics, distributionally robust formulations optimize over an *ambiguity set* of possible disturbance distributions ensuring robust satisfaction of equality and inequality constraints (Wiesemann et al., 2014).

There are two main data-driven approaches to distributionally robust optimal control: data-based synthesis of ambiguity sets to capture the distributional ambiguity in the disturbances while requiring explicit knowledge of a system model (Coppens et al., 2021; Fochesato et al., 2022; Lu et al., 2020) and robustness analysis of data-driven system descriptions with respect to the distributional ambiguity of the measurement noise (Coulson et al., 2022). However, uncertainty propagation through dynamics without explicit knowledge of the system model and considering distributional uncertainty of the disturbance is still an open problem. In this section, we address this gap by extending the data-driven description of stochastic linear systems from the previous chapters to incorporate uncertainty surrounding the disturbance distribution.

6.3.1. Distributional Ambiguity Sets

Instead of requiring exact knowledge, we can model the unknown true distribution as an element of a predefined ambiguity set. The most commonly used ambiguity sets employ the Wasserstein metric. However, tractable reformulations of Wasserstein ambiguity sets are often restricted to specific empirical distributions (Aolaritei et al., 2023a; Wiesemann et al., 2014) or Gaussian-based ambiguity sets (Nguyen et al., 2023). As an alternative, *Gelbrich ambiguity sets* include all distributions with moments that closely match a given nominal pair $(\bar{m},\bar{\Gamma})$ based on the Gelbrich distance.

Definition 6.1 (Gelbrich distance (Givens et al., 1984)): Consider two tuples of mean vectors and covariance matrices (m,Γ) and $(\bar{m},\bar{\Gamma})$, their Gelbrich distance is

$$\mathbb{G}\left((m,\Gamma),(\bar{m},\bar{\Gamma})\right) \doteq \sqrt{\|m-\bar{m}\|^2 + \mathrm{trace}(\Gamma+\bar{\Gamma}-2(\bar{\Gamma}^{\frac{1}{2}}\Gamma\bar{\Gamma}^{\frac{1}{2}})^{\frac{1}{2}})}. \tag{6.17}$$

Using the Gelbrich distance, we define the Gelbrich ambiguity set as the collection of all random variables whose first two moments are within a given radius $\rho \in \mathbb{R}^+$ of a specified nominal moment pair. Specifically, the Gelbrich ambiguity set is given by

$$\mathcal{G}_\rho\left(\bar{m},\bar{\Gamma}\right) \doteq \left\{W_\mathrm{f} \in \mathcal{C}(m,\Gamma) \,\middle|\, \Gamma \succeq 0,\ \mathbb{G}\left((m,\Gamma),(\bar{m},\bar{\Gamma})\right) \leq \rho\right\}. \tag{6.18}$$

The 0-radius Gelbrich set corresponds to the Cantelli ambiguity set defined in (3.20), i.e., $\mathcal{G}_0(m,\Gamma) = \mathcal{C}(m,\Gamma)$. Notably, the Gelbrich ambiguity set serves as an outer approximation of the corresponding Wasserstein set (Givens et al., 1984), as discussed below.

Next, we compare various ambiguity sets to position the Gelbrich ambiguity within the spectrum of distributional uncertainties. Consider two random variables $V, \bar{V} \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$ whose distributions are μ_V and $\mu_{\bar{V}}$, respectively. Suppose their mean-covariance pairs are (m,Γ) and $(\bar{m},\bar{\Gamma})$, respectively. We quantify the distances between distributions, moment pairs, and random variables via the Wasserstein metric $\mathbb{W}(\mu_V,\mu_{\bar{V}})$, the Gelbrich distance $\mathbb{G}\big((m,\Gamma),(\bar{m},\bar{\Gamma})\big)$, and the L^2 distance $\|V-\bar{V}\| \doteq \sqrt{\mathbb{E}[(V-\bar{V})^\top(V-\bar{V})]}$, respectively. The Wasserstein metric is defined as

$$\mathbb{W}(\mu_V,\mu_{\bar{V}}) \doteq \inf_{V_1,V_2} \|V_1 - V_2\|, \quad \text{subject to} \quad V_1 \sim \mu_V,\ V_2 \sim \mu_{\bar{V}}. \tag{6.19}$$

According to Gelbrich, 1990, the Gelbrich distance defined in (6.17) is equivalent to

$$\mathbb{G}\left((m,\Gamma),(\bar{m},\bar{\Gamma})\right) = \inf_{V_1,V_2} \|V_1 - V_2\|, \quad \text{subject to} \quad V_1 \in \mathcal{C}(m,\Gamma),\ V_2 \in \mathcal{C}(\bar{m},\bar{\Gamma}). \tag{6.20}$$

Note that the feasible region in (6.20) contains that of (6.19). Thus, the optimization formulations in (6.19) and (6.20) lead to the following relations

$$\|V-\bar{V}\| \geq \mathbb{W}(\mu_V,\mu_{\bar{V}}) \geq \mathbb{G}\left((m,\Gamma),(\bar{m},\bar{\Gamma})\right). \tag{6.21}$$

This result shows that the Wasserstein distance is upper-bounded by the L^2 distance and lower-bounded by the Gelbrich distance. Computing the Wasserstein distance between two probability distributions is generally #P-hard, as stated in Theorem 2.2 of Taşkesen et al., 2023. In contrast, both the L^2 and Gelbrich distances can be computed efficiently. Thus, in practice, we can estimate the Wasserstein distance by utilizing these bounds.

The bounds in (6.21) are tight in certain cases of practical relevance.

Lemma 6.2 (Equivalence of distances (Nguyen et al., 2021)): Consider random variables $V, \bar{V} \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v})$, if

$$\bar{\Gamma} \succ 0, \quad V = A\bar{V} + b,\ A \succeq 0, \tag{6.22}$$

then there exists a PCE basis $\{\phi^j(\xi)\}_{j=0}^{L-1}$ in which V and $\bar{V}$ admit exact PCEs, and the following equivalence of distances holds, i.e.,

$$\left\|\mathsf{v}^{[0,L-1]} - \bar{\mathsf{v}}^{[0,L-1]}\right\|_{\mathcal{D}}^2 = \|V-\bar{V}\| = \mathbb{W}(\mu_V,\mu_{\bar{V}}) = \mathbb{G}\left((m,\Gamma),(\bar{m},\bar{\Gamma})\right), \tag{6.23}$$

with $\mathcal{D} \doteq \operatorname{diag}\left(\{\|\phi^j(\xi)\|^2\}_{j=0}^{L-1}\right)$. The equivalence condition (6.22) holds if any of the following statements holds.

i) For given $A \succeq 0$, $b \in \mathbb{R}^{n_v}$, $\bar{\Gamma} \succeq 0$, $\bar{m} = \bar{\mathsf{v}}^0 \in \mathbb{R}^{n_v}$, and $\bar{\mathsf{v}}^{[1,L-1]} \in \mathbb{R}^{n_v(L-1)}$,

$$\Gamma = A\bar{\Gamma}A^\top, \quad \mathsf{v}^0 = m = A\bar{m} + b, \quad \mathcal{H}_1(\mathsf{v}^{[1,L-1]}) = A\mathcal{H}_1(\bar{\mathsf{v}}^{[1,L-1]}). \tag{6.24}$$

ii) For given $\bar{\Gamma} \succeq 0$, $\bar{m} \in \mathbb{R}^{n_v}$, $\Gamma \succeq 0$, $m \in \mathbb{R}^{n_v}$, $\bar{m} = \bar{\mathsf{v}}^0 \in \mathbb{R}^{n_v}$, and $\bar{\mathsf{v}}^{[1,L-1]} \in \mathbb{R}^{n_v(L-1)}$,

$$A = \bar{\Gamma}^{-\frac{1}{2}}\left(\bar{\Gamma}^{\frac{1}{2}}\Gamma\bar{\Gamma}^{\frac{1}{2}}\right)\bar{\Gamma}^{-\frac{1}{2}},\ b = m - A\bar{m},\ \mathsf{v}^0 = m,\ \mathcal{H}_1(\mathsf{v}^{[1,L-1]}) = A\mathcal{H}_1(\bar{\mathsf{v}}^{[1,L-1]}). \tag{6.25}$$

Proof. This result follows from Theorem 2 of Nguyen et al., 2021, which proves that (6.22) implies $\|V - \bar{V}\| = \mathbb{W}(\mu_V, \mu_{\bar{V}}) = \mathbb{G}\left((m, \Gamma), (\bar{m}, \bar{\Gamma})\right)$. Thus, we focus on proving the equivalence to the weighted norm of the PCE coefficients

By Lemma 3.2, there exists a PCE basis $\{\phi^j(\xi)\}_{j=0}^{L-1}$ in which V admits an exact PCE. Moreover, since $V = A\bar{V} + b$, Lemma 3.3 ensures that V also admits an exact PCE in the same basis. Hence,

$$\mathsf{v}^0 = A\bar{\mathsf{v}}^0 + b, \quad \mathcal{H}_1(\mathsf{v}^{[1,L-1]}) = A\mathcal{H}_1(\bar{\mathsf{v}}^{[1,L-1]}), \quad \left\|\mathsf{v}^{[0,L-1]} - \bar{\mathsf{v}}^{[0,L-1]}\right\|_{\mathcal{D}}^2 = \|V - \bar{V}\|,$$

where the L^2 norm reformulation using PCE follows from (3.6).

Furthermore, regarding the equivalence of statements i)–ii) to (6.22), the implication from statement i) is straightforward, while that from statement ii) follows from the proof of Theorem 2 by Nguyen et al., 2021. □

In addition to the distances appearing in the inequality (6.21), equality (6.23) also incorporates the weighted Euclidean distance of the PCE coefficients. A comparable condition for the equivalence of the PCE-based distance and the Gelbrich distance has been proposed by Pan et al., 2023a.

Using these distances, we define their corresponding ambiguity sets as follows. First, the Gelbrich ambiguity set $\mathcal{G}_\rho(\bar{m}, \bar{\Gamma})$, defined in (6.18), consists of all random variables whose moment pairs are within a Gelbrich distance ρ of the nominal pair $(\bar{m}, \bar{\Gamma})$. Second, the Wasserstein set $\mathcal{W}_\rho(\mu_{\bar{V}})$ contains all random variables whose distributions are within a Wasserstein distance ρ of the nominal distribution $\mu_{\bar{V}}$. That is,

$$\mathcal{W}_\rho(\mu_{\bar{V}}) \doteq \left\{V \sim \mu_V \,\middle|\, \mathbb{W}(\mu_V, \mu_{\bar{V}})^2 \leq \rho^2\right\}. \tag{6.26}$$

Third, the ambiguity set $\mathcal{U}_\rho(\bar{V})$, induced by the L^2 norm, contains all random variables close to the nominal $\bar{V}$,

$$\mathcal{U}_\rho(\bar{V}) \doteq \left\{V \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_v}\right) \,\middle|\, \|V - \bar{V}\|^2 \leq \rho^2\right\}.$$

Finally, restricting V and $\bar{V}$ in $\mathcal{U}_\rho(\bar{V})$ to random variables that admit exact PCEs in the basis $\{\phi^j\}_{j=0}^{L-1}$, we define the ambiguity set corresponding to the distance between PCE coefficients as

$$\mathcal{P}_\rho\left(\bar{\mathsf{v}}^{[0,L-1]}\right) \doteq \left\{V = \sum_{j=0}^{L-1} \mathsf{v}^j \phi^j \,\middle|\, \left\|\mathsf{v}^{[0,L-1]} - \tilde{\mathsf{v}}^{[0,L-1]}\right\|_{\mathcal{D}}^2 \leq \rho^2\right\}.$$

Then, the inequality of distances in (6.21) implies

$$\mathcal{P}_\rho\left(\bar{\mathsf{v}}^{[0,L-1]}\right) \subseteq \mathcal{U}_\rho\left(\bar{V}\right) \subseteq \mathcal{W}_\rho\left(\bar{\mu}\right) \subseteq \mathcal{G}_\rho\left(\bar{m}, \bar{\Gamma}\right). \tag{6.27}$$

Given the set $\mathcal{S}\left(\bar{V}\right) \doteq \{V = A\bar{V} + b \,|\, A \succ 0,\, b \in \mathbb{R}^{n_v}\}$ and assuming $\bar{\Gamma} \succ 0$, we obtain

$$(\mathcal{P}_\rho \cap \mathcal{S}) = (\mathcal{U}_\rho \cap \mathcal{S}) = (\mathcal{W}_\rho \cap \mathcal{S}) = (\mathcal{G}_\rho \cap \mathcal{S}). \tag{6.28}$$

These results show that the Gelbrich ambiguity set provides an outer approximation of the Wasserstein ambiguity set. However, the approximation can be exact if the distributional ambiguity is limited to affine transformations with positive semi-definite matrices.

6.3.2. Distributionally Robust Chance Constraints

Note that the chance constraints (6.3e)–(6.3f) are imposed on scalar random variables $a_{u,i}^\top U(k)$ and $a_{y,i}^\top Y(k)$, both of which are affinely transformed from W_f. To investigate the distributionally robust formulation of (6.3e)–(6.3f), without loss of generality, we consider a general scalar random variable M generated from W_f via an affine mapping, i.e.,

$$M = a^\top W_\mathrm{f} + b, \tag{6.29}$$

where $a, b \in \mathbb{R}^{n_w N}$ are determined from the optimization decision variables $\bar{u}$ and $\mathbf{K}_N$.

Consider the distributionally robust formulation of chance constraints

$$\forall W_\mathrm{f} \in \mathcal{G}_\rho(\bar{m}, \bar{\Gamma}), \quad \mathbb{P}[M \le 1] \ge \gamma, \tag{6.30}$$

which requires the chance constraints to hold for all possible W_f within the Gelbrich ambiguity set $\mathcal{G}_\rho(\bar{m}, \bar{\Gamma})$.

Lemma 6.3 (Exact reformulation of Gelbrich constraints (Nguyen et al., 2021)): Consider a scalar random variable M satisfying (6.29) with $W_\mathrm{f} \in \mathcal{G}_\rho(\bar{m}, \bar{\Gamma})$. Then, the distributionally robust chance constraint (6.30) is equivalent to the following reformulation

$$a^\top \bar{m} + b + \alpha(\gamma)\sqrt{a^\top \bar{\Gamma} a} + \rho\sqrt{1+\alpha(\gamma)^2}\|a\| \le 1, \tag{6.31}$$

where $\alpha(\gamma)$ is the concentration factor in (3.19), which can be further reduced by restricting the distribution type of W_f.

The above formulation requires an explicit expression for the map $M = a^\top W_\mathrm{f} + b$. However, in the case of $M \in \{a_{u,i}^\top U(k), a_{y,i}^\top Y(k)\}$ within OCP (6.9), these maps are implicitly defined through the data-driven propagation (6.9c) and the affine input policies (6.9d). Instead of explicitly constructing these maps, we can extract a by leveraging the PCE coefficients of W and M. Consequently, the following proposition provides an exact reformulation of (6.30) in terms of PCE coefficients.

Proposition 6.1 (Exact PCE reformulation of Gelbrich constraints): Consider the PCE of a nominal future disturbance trajectory $\bar{W}_\mathrm{f} \doteq \sum_{j=0}^{L-1} \bar{\mathsf{w}}_\mathrm{f}^j \phi^j$ with and a nominal $\bar{M} \doteq \sum_{j=0}^{L-1} \bar{\mathsf{m}}^j \phi^j$ satisfying $\bar{M} = a^\top \bar{W}_\mathrm{f} + b$.

Consider the true M satisfying (6.29) with $W_\mathrm{f} \in \mathcal{G}_\rho\left(\mathbb{E}[\bar{W}_\mathrm{f}], \Sigma[\bar{W}_\mathrm{f}]\right)$. Then, the distributionally robust chance constraint (6.30) is equivalent to the following reformulation in terms of PCE coefficients

$$\bar{\mathsf{m}}^0 + \alpha(\gamma)\sqrt{\sum_{j=1}^{L-1} \|\bar{\mathsf{m}}^j\|^2 \|\phi^j\|^2} + \rho\sqrt{1+\alpha(\gamma)^2}\left\|\left(\left(\mathcal{H}_1(\bar{\mathsf{w}}_\mathrm{f}^{[1,L-1]})\right)^\dagger\right)^\top \bar{\mathsf{m}}^{[1,L-1]}\right\| \le 1. \tag{6.32}$$

This can be equivalently reformulated as the following two second-order cone constraints

$$\begin{aligned} &\bar{\mathsf{m}}^0 + \alpha(\gamma)\sqrt{\sum_{j=1}^{L-1} \|\bar{\mathsf{m}}^j\|^2 \|\phi^j\|^2} + t \le 1, \\ &t \ge 0, \quad \rho\sqrt{1+\alpha(\gamma)^2}\left\|\left(\left(\mathcal{H}_1(\bar{\mathsf{w}}_\mathrm{f}^{[1,L-1]})\right)^\dagger\right)^\top \bar{\mathsf{m}}^{[1,L-1]}\right\| \le t. \end{aligned} \tag{6.33}$$

Proof. Instead of explicitly specifying the nominal moment pair $(\bar{m}, \bar{\Gamma})$ as in Lemma 6.3, we consider a nominal random variable $\bar{W}_\mathrm{f}$. With $\bar{M} = a^\top \bar{W}_\mathrm{f} + b$, the chance constraint reformulation in Lemma 6.3 can be rewritten as

$$\mathbb{E}[\bar{M}] + \alpha(\gamma)\sqrt{\mu_2(\bar{M})} + \rho\sqrt{1+\alpha(\gamma)^2}\|a\| \leq 1, \tag{6.34}$$

since $\mathbb{E}[\bar{M}]$ and $\mu_2(\bar{M})$ can be directly obtained from the PCE coefficients of $\bar{M}$, the remaining challenge is to obtain a using the PCE coefficients of $\bar{W}_\mathrm{f}$ and $\bar{M}$.

By plugging the PCE coefficients into $\bar{M} = a^\top \bar{W}_\mathrm{f} + b$, we obtain

$$\bar{\mathsf{m}}^0 = a^\top \bar{\mathsf{w}}_\mathrm{f}^0 + b, \quad \bar{\mathsf{m}}^j = a^\top \bar{\mathsf{w}}_\mathrm{f}^j.$$

By horizontally stacking the PCE coefficients from $j = 1$ to $L-1$, we have

$$\mathcal{H}_1\left(\bar{\mathsf{m}}^{[1,L-1]}\right) = a^\top \mathcal{H}_1\left(\bar{\mathsf{w}}_\mathrm{f}^{[1,L-1]}\right). \tag{6.35}$$

Moreover, let $\mathsf{M}_w \doteq \mathcal{H}_1(\bar{\mathsf{w}}_\mathrm{f}^{[1,L-1]})$, we note that M_w is of full row rank since $\Sigma[\bar{W}_\mathrm{f}] = \mathsf{M}_w\mathsf{M}_w^\top$ is positive definite. As a result, the Moore-Penrose inverse $\mathsf{M}_w^\dagger$ serves as the right inverse, i.e., $\mathsf{M}_w\mathsf{M}_w^\dagger = I$. Multiplying both sides of (6.35) by $\mathsf{M}_w^\dagger$, we obtain

$$a^\top = \mathcal{H}_1\left(\bar{\mathsf{m}}^{[1,L-1]}\right)\mathsf{M}_w^\dagger.$$

Since $\mathcal{H}_1\left(\bar{\mathsf{m}}^{[1,L-1]}\right) = (\bar{\mathsf{m}}^{[1,L-1]})^\top$, it follows that $a = (\mathsf{M}_w)^{\dagger\top}\bar{\mathsf{m}}^{[1,L-1]}$. Finally, substituting the PCE reformulations of $\bar{M}$ and a into (6.34) leads to (6.32) and its equivalent reformulation in (6.33). □

This result appeared previously as Proposition 1 of Pan et al., 2023a with $\rho = 0$.

Remark 6.3 (Construction of $\bar{W}_f$): The preceding analysis assumes that the nominal disturbances $\bar{W}_\mathrm{f}$ are represented as random variables. However, in cases where only the nominal moment pair $(\bar{m}, \bar{\Gamma})$ is available, with $\bar{\Gamma} \succ 0$, we can proceed as follows. By Lemma 3.2, we can construct a nominal random variable $\bar{W}_\mathrm{f} = \bar{m} + \mathsf{M}_w\xi$ where $\mathsf{M}_w \in \mathbb{R}^{Nn_w \times n_\xi}$ is a full row-rank decomposition of $\bar{\Gamma}$, satisfying $\mathsf{M}_w\mathsf{M}_w^\top = \bar{\Gamma}$, with $n_\xi > Nn_w$. The distribution of $\xi \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\xi})$ is unspecified, but its elements $\{\xi_j\}_{j=1}^{n_\xi}$ are independently distributed with zero mean and unit variance. As a result, the basis $\mathcal{P}_1(\xi) = \{1, \xi_1, ..., \xi_{Nn_w}\}$ forms a valid orthonormal PCE basis, allowing for an exact PCE representation of $\bar{W}_\mathrm{f}$. Specifically, $\bar{W}_\mathrm{f} = \sum_{j=0}^{n_\xi} \bar{\mathsf{w}}_\mathrm{f}^j \phi^j$ with $\mathcal{H}_1(\bar{\mathsf{w}}_\mathrm{f}^{[1,n_\xi]}) = \mathsf{M}_w$. This construction ensures that $\bar{W}_\mathrm{f}$ has the prescribed first- and second-order moments, even in the absence of an explicit probability distribution.

Consider the PCE of a nominal future trajectory of disturbances $\bar{W}_\mathrm{f} = \sum_{j=0}^{L-1} \bar{\mathsf{w}}_\mathrm{f}^j \phi^j$ is given. We denote the shorthand notation for (6.33) as $(\bar{\mathsf{m}}^{[0,L-1]}, t) \in \mathrm{DRC}_\rho(\bar{W}_\mathrm{f})$. This leads us to the distributionally robust extension of OCP (6.9) as follows

$$\min_{\bar{u}, \mathbf{K}_N, \mathsf{u}, \mathsf{y}, \mathsf{g}, t} \sum_{k=0}^{N-1}\sum_{j=0}^{L-1}\left(\|\mathsf{y}^j(k)\|_Q^2 + \|\mathsf{u}^j(k)\|_R^2\right)\|\phi^j\|^2 \tag{6.36a}$$

$$\text{subject to } \mathcal{H}_\mathrm{p}\mathsf{g}^j = \delta^{0j}[u_\mathrm{p}^\top, y_\mathrm{p}^\top, w_\mathrm{p}^\top]^\top, \quad \forall j \in \mathbb{I}_{[0,L-1]}, \tag{6.36b}$$

$$\mathcal{H}_\mathrm{f}\mathsf{g}^j = [\mathsf{u}_\mathrm{f}^{j\top}, \mathsf{y}_\mathrm{f}^{j\top}, \bar{\mathsf{w}}_\mathrm{f}^{j\top}]^\top, \quad \forall j \in \mathbb{I}_{[0,L-1]}, \tag{6.36c}$$

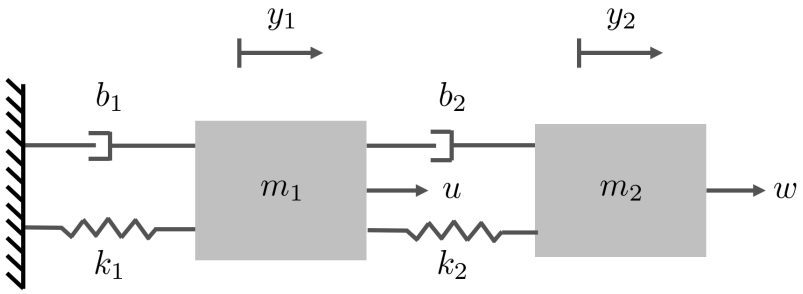

Figure 6.1.: Schematic of a mass-spring-damper system

$$\mathsf{u}_{\mathrm{f}}^{j} = \delta^{0j}\bar{u}_{\mathrm{f}} + \mathbf{K}_N \bar{\mathsf{w}}_{\mathrm{f}}^{j}, \quad \forall j \in \mathbb{I}_{[1,L-1]}, \tag{6.36d}$$

$$\left(a_{u,i}^{\top}\mathsf{u}^{[0,L-1]}(k), t_{u,i}\right) \in \mathrm{DRC}_{\rho}(\bar{W}_{\mathrm{f}}), \quad \forall i \in \mathbb{I}_{[1,N_u]}, k \in \mathbb{I}_{[0,N-1]}, \tag{6.36e}$$

$$\left(a_{y,i}^{\top}\mathsf{y}^{[0,L-1]}(k), t_{y,i}\right) \in \mathrm{DRC}_{\rho}(\bar{W}_{\mathrm{f}}), \quad \forall i \in \mathbb{I}_{[1,N_y]}, k \in \mathbb{I}_{[0,N-1]}. \tag{6.36f}$$

Aside from the distributionally robust chance constraints (6.36e)–(6.36f), the remaining components of (6.36) are expressed in terms of the PCE coefficients of the nominal disturbances $\bar{W}_{\mathrm{f}}$. Hence, (6.36) reduces to OCP (6.9) when $\rho = 0$. In practice, the parameter $\rho \in \mathbb{R}^{+}$ can be incrementally tuned from 0 to enhance the distributionally robust satisfaction of chance constraints compared to (6.9).

6.4. Case Study: Mass-Spring-Damper System

6.4.1. Simulation Setup

We consider a mass-spring-damper system as shown in Figure 6.1. The system consists of two masses $m_1 = 2\,\mathrm{kg}$ and $m_2 = 0.5\,\mathrm{kg}$, two springs with constants $k_1 = 6\,\mathrm{N/m}$ and $k_2 = 1.5\,\mathrm{N/m}$, and two dampers with coefficients $b_1 = 3\,\mathrm{Ns/m}$ and $b_2 = 1\,\mathrm{Ns/m}$. An external input force, denoted as $u(t)\,[\mathrm{N}]$, is applied to m_1, while a disturbance force, represented by $w(t)\,[\mathrm{N}]$, acts on m_2. The displacements of m_1 and m_2 are considered as the system outputs, denoted as $y_1(t)\,[\mathrm{m}]$ and $y_2(t)\,[\mathrm{m}]$, respectively.

The system can be modeled by the following equations. For m_1, we have

$$m_1\ddot{y}_1(t) = u(t) - k_1 y_1(t) - b_1\dot{y}_1(t) - k_2\left(y_1(t) - y_2(t)\right) - b_2\left(\dot{y}_1(t) - \dot{y}_2(t)\right).$$

For m_2, we have

$$m_2\ddot{y}_2(t) = w(t) - k_2\left(y_2(t) - y_1(t)\right) - b_2\left(\dot{y}_2(t) - \dot{y}_1(t)\right).$$

By considering the state variable $x(t) = \begin{bmatrix} y_1(t) & \dot{y}_1(t) & y_2(t) & \dot{y}_2(t) \end{bmatrix}$ and a sampling instant $0.4\mathrm{s}$, we obtain the following state-space representation

$$x(k+1) = \begin{bmatrix} 0.7863 & 0.2564 & 0.0331 & 0.0273 \\ -0.8797 & 0.3279 & 0.1105 & 0.1068 \\ 0.1219 & 0.1090 & 0.8296 & 0.2616 \\ 0.3760 & 0.4270 & -0.7030 & 0.3610 \end{bmatrix} x(k) + \begin{bmatrix} 0.0301 \\ 0.1282 \\ 0.0081 \\ 0.0545 \end{bmatrix} u(k) + \begin{bmatrix} 0.0081 \\ 0.0545 \\ 0.1217 \\ 0.5232 \end{bmatrix} w(k),$$

$$y(k) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} x(k).$$

Since we are exploring data-driven control design, the system matrices mentioned above are not directly accessible to the controller; they are only used as part of the simulated environment.

The disturbance force $w(k)$, $k \in \mathbb{N}$, is modeled as a realization of a *random-walk* process, given by

$$w\,(k) \doteq W(k,\omega),\; W(k{+}1) = W(k){+}\frac{2\sqrt{3}}{30}\xi(k{+}1),\; W(0) = \frac{2\sqrt{3}}{30}\xi(0),\; \xi(k) \sim \sqrt{3}\mathcal{U}([-1,1])$$

with the stochastic germ $\xi(k)$ to be i.i.d. uniform distributed over $k \in \mathbb{N}$. Moreover, we have $\mathbb{E}[\xi(k)] = 0$ and $\mu_2(\xi(k)) = 1$ for all $k \in \mathbb{N}$. Note that the random-walk process W has increments $(W(k+1) - W(k)) \sim \mathcal{U}(-0.2, 0.2)$ to be i.i.d. and uniformly distributed. Consequently, W itself is neither i.i.d. nor Gaussian.

In addition, we consider a system lag of $\ell = 2$, a control horizon $N = 10$, and weighting matrices $Q = 10I_2$ and $R = 1$. The chance constraints require that $U(k) \leq 0.5$, $U(k) \geq -0.5$, $Y_1(k) \leq 5$, and $Y_1(k) \geq -5$ are satisfied individually with a probability of no less than $\gamma = 90\%$ for $k \in \mathbb{I}_{[0,N-1]}$.

To construct exact PCEs, we consider the basis

$$\{\phi^j\}_{j=0}^{L-1} = \mathcal{P}_1(\xi_{[0,N-1]}) = \{1, \xi\,(0)\,, \cdots, \xi\,(N-1)\}$$

with $L = 1 + N$. Since $\xi(k)$ is i.i.d. and has unit variance, the basis is orthonormal. We note that $(U, Y, W)_{[0,N-1]}$ admit exact PCEs in $\{\phi^j\}_{j=0}^{L-1}$, with $k \in \mathbb{I}_{[0,N-1]}$,

$$\mathsf{w}^j\,(k) = \begin{cases} (k+1)\mu, & j = 0, \\ \sigma, & 1 \leq j \leq k+1, \\ 0, & \text{else}. \end{cases} \tag{6.37}$$

In other words, $(U, Y, W)_{[0,N-1]}$ can be mapped from $\xi_{[0,N-1]}$ via affine mappings. Furthermore, since the uniform distribution is symmetric and linear unimodal, and $\xi\,(k)$ is i.i.d. uniform, it follows from Lemma 3.7 that $U(k)$ and $Y(k)$ are also symmetric linear unimodal for all $k \in \mathbb{I}_{[0,N-1]}$. Hence, the concentration factor in (6.9e) is chosen accordingly as $\alpha(\gamma) = 2/\left(3\sqrt{2 \cdot 0.1}\right) \approx 1.491$.

To construct OCPs (6.9) and (6.36) using measured data, we first apply $T = 160$ random inputs $u^{\mathrm{d}}_{[0,T-1]}$ to the system and record the output responses $y^{\mathrm{d}}_{[0,T-1]}$ as well as the realized disturbances $w^{\mathrm{d}}_{[0,T-1]}$. Then, we use this data to construct Hankel matrices and estimate the mean and variance of the increment $W(k+1) - W(k)$. With the recorded disturbance trajectory $w^{\mathrm{d}}_{[0,T-1]}$, we compute the empirical mean and variance of the increment as follows

$$\bar{\mu} = \frac{1}{T-1}\sum_{k=0}^{T-2}(w^{\mathrm{d}}(k+1) - w^{\mathrm{d}}(k)), \quad \bar{\sigma}^2 = \frac{1}{T-2}\sum_{k=0}^{T-2}(w^{\mathrm{d}}(k+1) - w^{\mathrm{d}}(k) - \bar{\mu})^2.$$

Specifically, the estimates yield $\bar{\mu} \approx -0.0176$ and $\bar{\sigma}^2 \approx 0.0121$, in contrast to the true values $\mu_{\text{true}} = 0$ and $\sigma^2_{\text{true}} \approx 0.0133$. By replacing μ_{true} and σ_{true} with $\bar{\mu}$ and $\bar{\sigma}$ in (6.37), we obtain the estimated PCE coefficients $\bar{\mathsf{w}}^j(k)$ for $k \in \mathbb{I}_{[0,N-11]}$ and $j \in \mathbb{I}_{[0,L-1]}$. Note

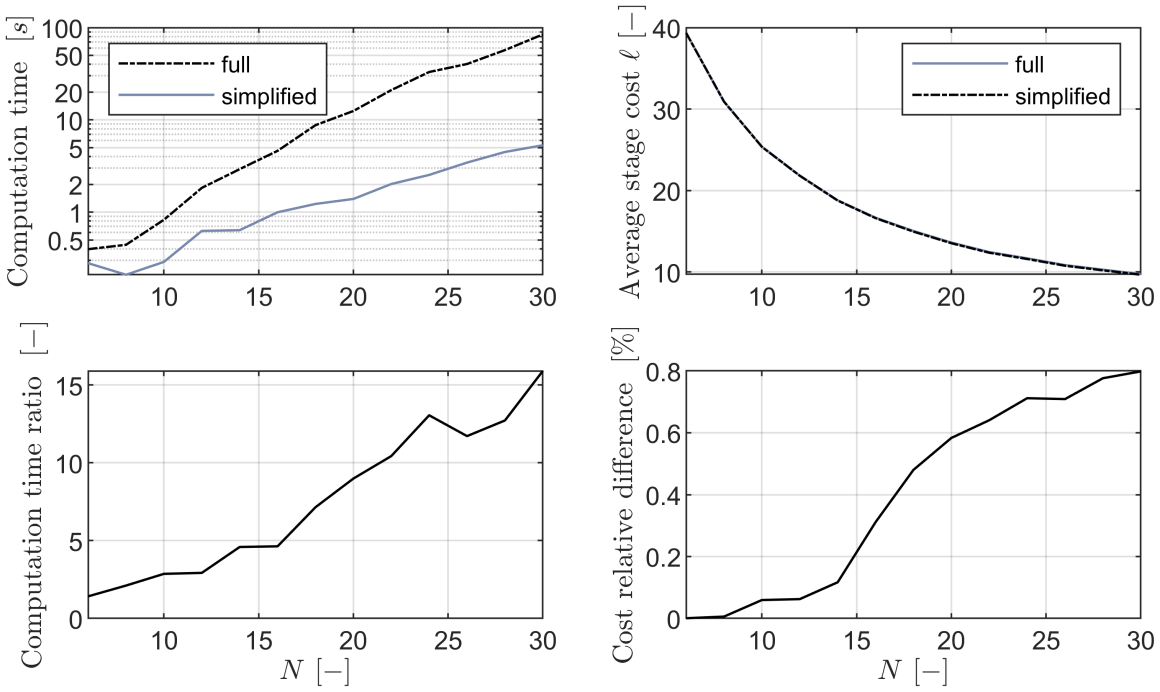

Figure 6.2.: Comparison of computation time and open-loop performance for OCP (6.9) with the full control policy (6.4) and the simplified control policy (6.11) for different control horizons N

that the estimated disturbance trajectory $\bar{W}_\mathrm{f} = \sum_{j=0}^{L-1} \bar{\mathsf{w}}_\mathrm{f}^j \phi^j$ can be linked to the true one $W_\mathrm{f} = \sum_{j=0}^{L-1} \mathsf{w}_\mathrm{f}^j \phi^j$ via an affine transformation, i.e.,

$$\bar{W}_\mathrm{f} = \frac{\sigma_\mathrm{true}}{\bar{\sigma}} I_N W_\mathrm{f} - \frac{\sigma_\mathrm{true}\bar{\mu}}{\bar{\sigma}}. \tag{6.38}$$

Since $\frac{\sigma_\mathrm{true}}{\bar{\sigma}} I_N$ is positive definite, it follows from Lemma 6.2 and the orthonormality of the basis that the Gelbrich distance between the moment pairs of W_f and $\bar{W}_\mathrm{f}$ is the same as the distance between their PCE coefficients

$$\mathbb{G}\left(\left(\mathbb{E}[W_\mathrm{f}], \Sigma[W_\mathrm{f}]\right), \left(\mathbb{E}[\bar{W}_\mathrm{f}], \Sigma[\bar{W}_\mathrm{f}]\right)\right) = \left\| \mathsf{w}_{[0,N-1]}^{[0,L-1]} - \bar{\mathsf{w}}_{[0,N-1]}^{[0,L-1]} \right\| \approx 0.564 \text{ for } N = 14. \tag{6.39}$$

6.4.2. Simplified Policy vs. Full Policy

We first examine the influence of using the simplified control policy in (6.11) instead of the full policy in (6.4). To this end, we solve (6.9) with the initial condition

$$\begin{bmatrix} u_\mathrm{p}^\top & w_\mathrm{p}^\top & y_\mathrm{p}^\top \end{bmatrix} = \begin{bmatrix} 0.0715 & 0.0019 & -2.7489 & -2.8328 & 4 & 4 & 3.2568 & 3.4721 \end{bmatrix} \tag{6.40}$$

and the PCE coefficients of the exact W_f using policies with different control horizons $N = 6, 8, \cdots, 30$. Moreover, we eliminate the equalities in OCP (6.9) by using the explicit solution to the Hankel representation, similarly to (2.27).

As shown in Figure 6.2, the computation time (top-left plot) grows significantly as N increases, particularly for the full policy in (6.4) (dashed black line), compared to the simplified policy (6.4b) (solid blue line). This observation aligns with the complexity analyses (6.10) and (6.12). In the bottom-left plot, the computation time ratio (full/simplified) indicates the relative computational expense of the full policy compared to the simplified one. This ratio steadily increases and exceeds 15 for horizons

Table 6.1.: Comparison of the computation times, the average costs, and the number of violations for 1000 realized open-loop trajectories. The results are presented for OCP (6.9) using both true and estimated PCE coefficients, as well as for OCP (6.36) with estimated PCE coefficients for different values of $\rho \in \mathbb{R}^+$.

Case	Computation time	ℓ	$\max_k \mathbb{P}^{\mathrm{u}}(k)$	$\max_k \mathbb{P}^{\mathrm{l}}(k)$
true	0.623	18.80	1.7%, $k=5$	7.4%, $k=9$
estimated	0.647	18.80	3.6%, $k=5$	14.9%, $k=9$
$\rho = 0.01$	0.958	18.80	0	13.7%, $k=9$
$\rho = 0.1$	1.131	18.80	0	5.8%, $k=9$
$\rho = 0.2$	1.157	18.81	0	0.8%, $k=8$
$\rho = 0.4$	1.457	18.81	0	0
$\rho = 0.6$	1.535	18.81	0	0

larger than 30. These two plots highlights the growing computational burden of the full control policy compared to the simplified one.

The top-right and bottom-right plots both examine the open-loop performance of the full and simplified control policies. To this end, we define the average stage cost ℓ as

$$\ell(N) = \frac{1}{N_s} \sum_{i=0}^{N_s-1} \sum_{k=0}^{N-1} \|u^i(k)\|_R^2 + \|y^i(k)\|_Q^2, \tag{6.41}$$

where (u^i, y^i) represents the realization trajectory corresponding to the disturbance realization w^i, $i \in \mathbb{I}_{[0,N_s-1]}$, and $N_s = 1000$ samples are evaluated. In the top-right plot, the average stage cost $\ell(N)$ decreases as the control horizon N increases. Both policies show almost identical performance in reducing the cost. The bottom-right plot presents the relative cost difference $\left(\ell^{\mathrm{simp}}(N) - \ell^{\mathrm{full}}(N)\right)/\ell^{\mathrm{full}}(N)$. As N increases, the cost difference remains small but gradually grows to around 0.8%. This indicates that while the simplified policy leads to a slightly higher cost, it remains very close in performance to the full policy.

In summary, these plots demonstrate that while the computational cost of the full policy increases significantly with N, the simplified policy offers a more efficient alternative with only a minor trade-off in performance.

6.4.3. Distributionally Robust Implementation

In the following, we investigate the impact of incorporating distributionally robust chance constraints when using the estimated PCE coefficients, $\bar{\mathsf{w}}$, instead of the true values. To do this, we compare the optimal policies obtained by

i) solving (6.9) with the PCE coefficients of the true W_{f},

ii) solving (6.9) with the PCE coefficients of the estimated $\bar{W}_{\mathrm{f}}$ from (6.38),

iii) and solving the distributionally robust OCP (6.36) with the estimated PCE coefficients for different values of the Gelbrich radius, $\rho \in \mathbb{R}^+$.

In all scenarios, we use the initial condition from (6.40), apply the simplified policy in (6.11), and set the control horizon to $N = 14$.

Next, we generate 1000 different disturbance realization sequences, each of length $N = 14$, and compute the corresponding open-loop responses by applying the obtained optimal policies. The computation times, the average stage cost, and the highest empirical relative frequency of constraint violations $\max_k \mathbb{P}^{\mathrm{l}}(k)$, $\max_k \mathbb{P}^{\mathrm{u}}(k)$ are summarized in Table 6.1. Here, the empirical relative frequency of constraint violation $\mathbb{P}^{\mathrm{l}}(k)$ represents the percentage of time instants where $u(k) < -0.5$ over all 1000 realizations at time instant $k \in \mathbb{I}_{[0,N-1]}$, and $\mathbb{P}^{\mathrm{u}}(k)$ is defined similarly for $u(k) > 0.5$. We note that there are no violations of the output constraints in this example across all scenarios.

As shown in Table 6.1, while the average stage costs remain consistent across all approaches, using estimated coefficients leads to more frequent constraint violations compared to using true coefficients. Specifically, without including distributional robustness, the maximum violation rate for the estimated coefficients exceeds the allowed 10% violation rate. However, incorporating distributionally robust constraints reduces violations. In particular, for $\rho = 0.1$, the maximum violation rate drops below drops below 10%, and for $\rho \geq 0.4$, there are no violations. However, this robustness comes at the cost of increased computation time due to the additional second-order cone constraints in OCP (6.36), compared to OCP (6.9).

It is important to note that although $\rho \geq 0.564$ is required to ensure that the Gelbrich

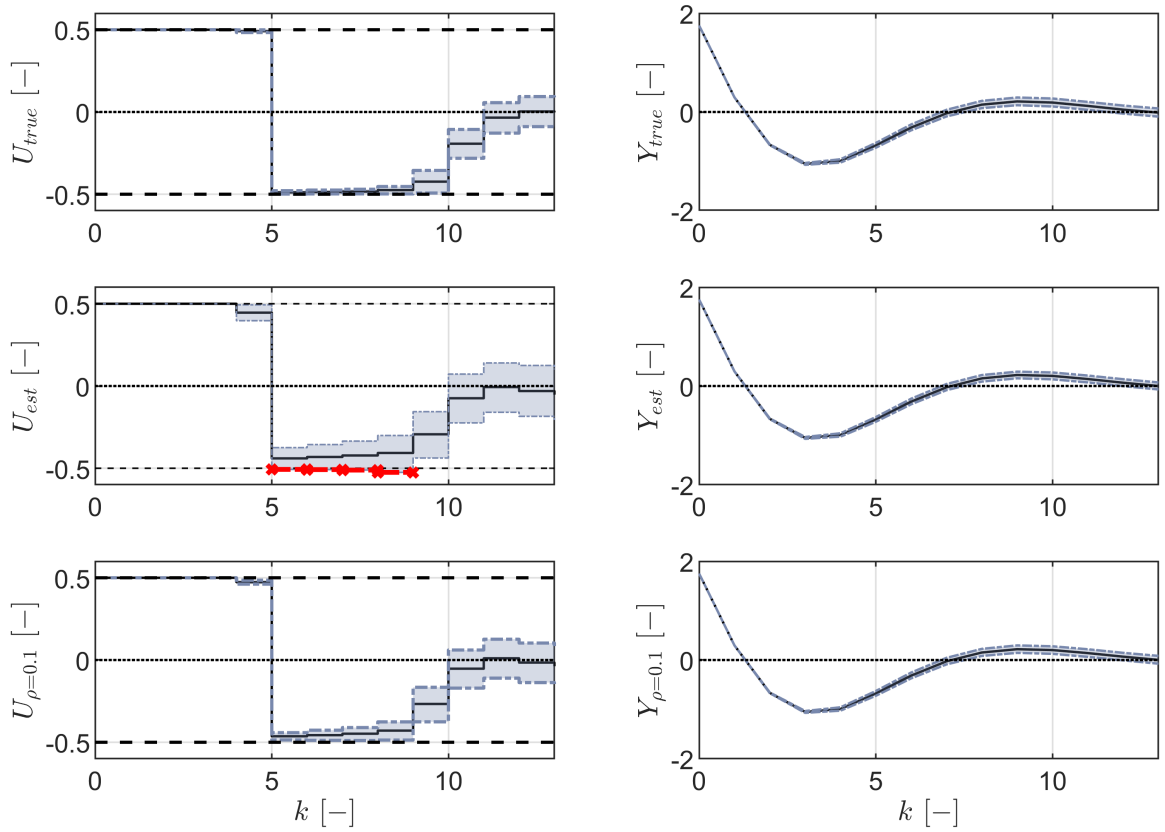

Figure 6.3.: Input and output response for 1000 disturbance sequences across three cases: using the true PCE coefficients, the estimated coefficients, and the estimated coefficients with ambiguity radius $\rho = 0.1$. Dark-solid lines represent the mean trajectories of the response. Grey dash-dotted lines represent the 90% and 10% empirical quantile paths. Red crosses highlight the points where chance constraints are violated, i.e., where quantile values exceed the defined constraints (dark-dash line) .

radius is larger than the Gelbrich distance in (6.39), we can incrementally tune ρ from $\rho = 0$ and opt for the minimal value that ensures the empirical violation rate to be below the allowed threshold. This tuning approach prevents unnecessary conservatism while maintaining robustness.

Figure 6.3 compares the input and output responses for the 1000 disturbance sequences across the "true", "estimated", and "$\rho = 0.1$" cases from Table 6.1. Here, dark-solid lines represent the mean response, and grey dash-dotted lines show the 10-th and 90-th empirical quantile values of the response trajectories. In the middle plot (with estimated coefficients), red crosses highlight time points where the 10-th quantile of the inputs violates the lower constraint $U = -0.5$. This indicates the violations of the chance constraint $\mathbb{P}[U(k) \geq -0.5] \geq 0.9$ in the absence of distributional robustness. In contrast, no such violations are observed when using true PCE coefficients or when applying the distributionally robust chance constraint with $\rho = 0.1$. Overall, the figure illustrates that incorporating distributional robustness effectively mitigates constraint violations while maintaining control performance comparable to the true PCE case.

6.5. Summary

In this chapter, we addressed the challenges of optimal control for stochastic discrete-time LTI systems under exogenous disturbances. We began by formulating a model-based stochastic OCP that incorporated disturbance-feedback control policies and chance constraints. Next, we introduced a data-driven reformulation of the OCP using PCE and established conditions for the equivalence between model-based and data-driven OCPs. We also provided a proof demonstrating the optimality of affine disturbance-feedback policies. This part of the results is based on Pan et al., 2023c.

Additionally, we extended the analysis to cases involving Gelbrich ambiguity, which enhances distributional robustness against uncertainty in future disturbance predictions. This extension builds upon the results presented in Pan et al., 2023a. Finally, we concluded with numerical simulations that demonstrated the efficacy of the proposed data-driven stochastic control approach. Specifically, we observed that using simplified control policies and a properly tuned Gelbrich radius enables distributionally robust optimal control while maintaining low computational effort.

7. Data-Driven Stochastic Output-Feedback Predictive Control

In this chapter, we extend the data-driven optimal control framework developed in the previous chapters to address stochastic output-feedback predictive control. While our earlier discussions focused on open-loop control, we now shift our attention to predictive control strategies that ensure closed-loop guarantees despite the presence of stochastic uncertainties. This chapter builds upon the results presented by Pan et al., 2023b, 2025a,b.

7.1. Stochastic Output-Feedback Predictive Control

In this chapter, we consider the input-output dynamics of stochastic discrete-time LTI systems with the residual disturbance V,

$$Y(k) = \Xi Z(k) + DU(k) + V(k), \quad Z(0) = \bar{z}(0), \tag{7.1}$$

which is the same as (5.2). Here, we do not explicitly consider exogenous disturbances W; instead, these disturbances are captured within the residual disturbance V. Moreover, instead of having exact knowledge of the disturbance $V(k)$, $k \in \mathbb{N}$, we assume the following.

Assumption 7.1 (Disturbance with known first two moments): For all $k \in \mathbb{N}$, we consider i.i.d., possibly non-Gaussian, disturbances with zero mean $\mathbb{E}[V(k)] = 0$ and covariance $\Sigma[V(k)] = \Sigma_V$.

In practice, the first two moments—mean and covariance—of the disturbance can be obtained through empirical estimation from previously measured or estimated disturbances. We present our main results assuming that the disturbances are i.i.d. with zero mean. Extensions to non-i.i.d. settings are discussed in Remark 7.1.

Recall that the extended state $Z(k)$ collects the last $\ell \in \mathbb{N}^+$ inputs and outputs,

$$Z(k) \doteq \begin{bmatrix} U_{[k-\ell,k-1]} \\ Y_{[k-\ell,k-1]} \end{bmatrix} \in L^2(\Omega, \mathcal{F}_{k-1}, \mu; \mathbb{R}^{n_z}). \tag{7.2}$$

The state-space representation of (7.1) with $Z(k)$ as the state is given by

$$Z(k+1) = \widetilde{A}Z(k) + \widetilde{B}U(k) + \widetilde{H}V(k), \quad Y(k) = \widetilde{H}^\top Z(k+1). \tag{7.3}$$

with $\widetilde{A} = \begin{bmatrix} \bar{A} \\ \Xi \end{bmatrix}$, $\widetilde{B} = \begin{bmatrix} \bar{B} \\ D \end{bmatrix}$, $\widetilde{H} = \begin{bmatrix} 0 \\ I_{n_y} \end{bmatrix}$ and $\bar{A}$, $\bar{B}$ from (2.22). As shown by Berberich et al., 2021; De Persis et al., 2019; Sadamoto, 2023, the benefit of considering (7.2) is that it allows us to leverage data-driven state-feedback design concepts to achieve output feedback control. To this end, we make the following assumption.

Assumption 7.2: The pair $\left(\widetilde{A}, \widetilde{B}\right)$ from (7.3) is stabilizable.

This assumption ensures that system (7.3) is stabilizable when $V = 0$. In the presence of stochastic disturbances, we remark that $\left(\widetilde{A}, [\widetilde{B}, \widetilde{H}]\right)$ remains controllable regardless of the system matrices (Ξ, D) as shown in Lemma 5.2.

The focus of this chapter is then to design a feedback law $u(k) = \kappa_k(z(k))$ via a data-driven stochastic predictive control approach. Instead of estimating the internal state x, we utilize past inputs and outputs to construct the extended state z. Hence, the resulting control strategy is output-feedback. In addition, the influence of the disturbances $V(k)$ is implicitly handled via state recursion, eliminating the need for direct measurement of the residual disturbance.

For a specific uncertainty outcome $\omega \in \Omega$, we denote the realization of $V(k)$ as $v(k) \doteq V(k,\omega)$. Likewise, the realizations of input, output, and extended state are written as

$$u(k) \doteq U(k,\omega), \quad y(k) \doteq Y(k,\omega), \quad z(k) \doteq Z(k,\omega).$$

For a specific outcome, the stochastic system (7.1) induces the *realization dynamics*, i.e., the dynamics satisfied by path-wise sampled trajectories,

$$y(k) = \Xi z(k) + Du(k) + v(k), \quad z(0) = \bar{z}(0). \tag{7.4}$$

Assumption 7.3 (Data availability): While the system matrices $\Xi \in \mathbb{R}^{n_y \times n_z}$ and $D \in \mathbb{R}^{n_y \times n_u}$ are considered to be unknown, throughout this chapter, inputs and outputs, i.e., $u(k)$ and $y(k)$ are available at time k through measurements.

To present the main theoretical results, we initially assume perfect knowledge of the disturbance realizations $v(k)$. Then, we perform a robustness analysis of the proposed scheme when using the estimated realizations from (5.16).

7.1.1. Conceptual Model-Based Formulation

Next, we present a conceptual framework for stochastic output feedback predictive control. We begin with a model-based exposition and then discuss the challenges of transitioning to a data-driven setting, focusing on ensuring closed-loop guarantees.

At time instant k, we denote the predicted input, output, disturbance, and the extended state at time $k+t$ for $t \in \mathbb{I}_{[0,N-1]}$ as $U(t \mid k)$, $Y(t \mid k)$, $V(t \mid k)$, and $Z(t \mid k)$, respectively. For a prediction horizon $N \in \mathbb{N}^+$, we consider the following predicted dynamics

$$\begin{aligned} Z(t+1 \mid k) &= \widetilde{A} Z(t \mid k) + \widetilde{B} U(t \mid k) + \widetilde{H} V(t \mid k), \quad Y(k) = \widetilde{H}^\top Z(t+1 \mid k) \\ Z(0 \mid k) &= \bar{Z}(k). \end{aligned}$$

For shorthand, we use subscript f to denote the predictions over the horizon $[k, k+N-1]$ at time k. Here, we consider a stochastic initial condition $\bar{Z}(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_z})$ instead of simply using the measurement $z(k) \in \mathbb{R}^{n_z}$. In Section 7.2, we will show that the design of $\bar{Z}(k)$ is the key to establishing the recursive feasibility of our algorithm. Since the first two moments of the disturbance are given in Assumption 7.1, we assume that the disturbance sequence V_f resides in the Cantelli ambiguity set $\mathcal{C}(0, I_N \otimes \Sigma_V)$, where $\otimes$ denotes the Kronecker product. Additionally, we allow for an ambiguous initial condition $\bar{Z}(k) \in \mathcal{C}(m(k), \Sigma(k))$.

With the predicted dynamics, the conceptual model-based OCP to be solved at time step k reads

$$\min_{Z_\mathrm{f}(k), U_\mathrm{f}(k), Y_\mathrm{f}(k), \mathbf{K}_N(k), \bar{u}_\mathrm{f}(k)} J_N(U_\mathrm{f}(k), Y_\mathrm{f}(k), Z(N \mid k), z(k)) \tag{7.5a}$$

$$\text{subject to } \forall \bar{Z}(k) \in \mathcal{C}(m(k), \Sigma(k)),\ \forall V_{\mathrm{f}} \in \mathcal{C}(0, I_N \otimes \Sigma_V), \forall t \in \mathbb{I}_{[0,N-1]}$$

$$Z(t+1 \,|\, k) = \widetilde{A} Z(t \,|\, k) + \widetilde{B} U(t \,|\, k) + \widetilde{H} V(k+t), \quad Y(t \,|\, k) = \widetilde{H}^{\top} Z(t+1 \,|\, k), \tag{7.5b}$$

$$Z(0 \,|\, k) = \bar{Z}(k), \tag{7.5c}$$

$$U(t \,|\, k) = \bar{u}(t \,|\, k) + K(t \,|\, k)\, \bar{Z}(k) + \textstyle\sum_{j=0}^{t-1} K(t, j \,|\, k)\, V(k+j), \tag{7.5d}$$

$$\mathbb{P}\left[a_{u,i}^{\top} U(t \,|\, k) \leq 1\right] \geq \gamma_u,\ \forall i \in \mathbb{I}_{[1,N_u]}, \tag{7.5e}$$

$$\mathbb{P}\left[a_{y,i}^{\top} Y(t \,|\, k) \leq 1\right] \geq \gamma_y,\ \forall i \in \mathbb{I}_{[1,N_y]}, \tag{7.5f}$$

$$\mathbb{E}\left[Z(N \,|\, k)\right] \in \mathbb{Z}_{\mathrm{F}}, \quad \mathbb{E}\left[\|Z(N \,|\, k)\|_{\Gamma}^2\right] \leq \tau. \tag{7.5g}$$

As the knowledge of the initial condition and of the disturbance is limited to the first two moments, we consider a distributionally robust formulation where the initial condition and the disturbance sequence are taken from Cantelli ambiguity sets (3.20).

We consider the objective functional J_N from (7.5a) by

$$J_N\left(U_{\mathrm{f}}(k), Y_{\mathrm{f}}(k), Z(N \,|\, k), z(k)\right) \doteq$$
$$\mathbb{E}\left[\sum_{t=0}^{N-1} (\|Y(t \,|\, k)\|_Q^2 + \|U(t \,|\, k)\|_R^2) + \|Z(N \,|\, k)\|_P^2 \,\middle|\, \bar{Z}(k, \omega) = z(k)\right].$$

The objective functional represents the sum of conditional expected values of the predicted stage costs with weight matrices $R = R^{\top} \succ 0$, $Q = Q^{\top} \succeq 0$, and $P = P^{\top} \succeq 0$. Specifically, the matrix $P = P^{\top} \succeq 0$ specifies the terminal cost associated with the predicted extended state $Z(N \,|\, k)$.

Rather than considering the expectation over all outcomes $\omega \in \Omega$, we focus on the conditional expectation for the outcomes $\{\omega \,|\, \bar{Z}(k, \omega) = z(k)\}$. Here, the realization $z(k) \in \mathbb{R}^{n_z}$ is constructed by stacking past ℓ measured inputs and outputs following the structure of (7.2). Similar to Hewing et al., 2020b, this approach allows including the measured $z(k)$ to enhance performance while considering the stochastic $\bar{Z}(k)$ to ensure feasibility. As will be shown in Section 7.2.2, the stochastic initial condition $\bar{Z}(k)$ can be designed to guarantee recursive feasibility, while $z(k)$ in (7.5a) is used to establish average performance decay.

Following the open-loop optimal control strategy from the previous chapter, we impose causal and affine feedback policies using (7.5d). Moreover, due to the stochastic nature of the initial condition $\bar{Z}(k)$, we introduce $K(t \,|\, k)$ as an additional feedback gain from $\bar{Z}(k)$ to $U(t \,|\, k)$. Compactly, we define $\mathbf{K}_N(k) \in \mathbb{R}^{Nn_u \times (n_z + Nn_w)}$ as

$$\mathbf{K}_N(k) \doteq \begin{bmatrix} K(0 \,|\, k) & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ K(1 \,|\, k) & K(1, 0 \,|\, k) & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ K(N-1 \,|\, k) & K(N-1, 0 \,|\, k) & \cdots & K(N-1, N-2 \,|\, k) & \mathbf{0} \end{bmatrix},$$

which transforms (7.5d) into the following compact form

$$U_{\mathrm{f}}(k) = \bar{u}_{\mathrm{f}}(k) + \mathbf{K}_N(k) \begin{bmatrix} \bar{Z}(k) \\ V_{\mathrm{f}}(k) \end{bmatrix}.$$

After solving (7.5), the closed-loop input applied in each time instant is

$$u(k) = \bar{u}^{\star}(0 \,|\, k) + K^{\star}(0 \,|\, k)\, z(k) \tag{7.6}$$

where $z(k)$ is the current measured extended state, and is also the realization of the stochastic initial condition $\bar{Z}(k)$. Thus, $u(k)$ represents the realized value of $U(0 \mid k)$.

Although OCP (7.5) utilizes disturbance-feedback policies, its closed-loop implementations do not require explicit knowledge of disturbance realizations. Instead, as shown in (7.6), the implementation results in an *output-feedback* policy since the extended state $z(k)$ is constructed from past inputs and outputs.

Additionally, we impose distributionally robust chance constraints on $U(t \mid k)$ and $Y(t \mid k)$ for all $t \in \mathbb{I}_{[0,N-1]}$ using (7.5e)–(7.5f), ensuring that these constraints hold for all disturbances in the Cantelli ambiguity set. Specifically, we define N_u input constraints as individual half-space constraints with probability γ_u, where $a_{u,i} \in \mathbb{R}^{n_u}$ for $i \in \mathbb{I}_{[1,N_u]}$. Similarly, N_y output constraints are included.

Similar to the terminal ingredients design in state-feedback stochastic MPC (Farina et al., 2013; Pan et al., 2025b), we consider terminal constraints in (7.5g). These constraints require that the expected value of $Z(N \mid k)$ lies within the set $\mathbb{Z}_{\mathrm{F}} \subseteq \mathbb{R}^{n_z}$, and the terminal covariance weighted by $\Gamma \in \mathbb{R}^{n_z \times n_z}$ is bounded by $\tau \in \mathbb{R}^{+}$.

Furthermore, we impose that the terminal ingredients—including the terminal cost weight P used in the objective functional (7.5a) and the terminal constraints (7.5g)—are not predefined but are designed to satisfy the following assumption.

Assumption 7.4 (Terminal ingredients): Consider the system matrices $\widetilde{A} \in \mathbb{R}^{n_z \times n_z}$, $\widetilde{B} \in \mathbb{R}^{n_z \times n_u}$, and $\widetilde{H} \in \mathbb{R}^{n_z \times n_w}$ from (7.3). We assume that there exist matrices $P = P^\top \succeq 0 \in \mathbb{R}^{n_z \times n_z}$, $\Gamma = \Gamma^\top \succeq 0 \in \mathbb{R}^{n_z \times n_z}$, $K \in \mathbb{R}^{n_u \times n_z}$, and a positive real number $\tau \in \mathbb{R}^{+}$, such that $A_K = \widetilde{A} + \widetilde{B}K$ is Schur stable and

$$A_K^\top P A_K - P = -K^\top R K - A_K^\top \widetilde{H} Q \widetilde{H}^\top A_K, \tag{7.7a}$$

$$A_K^\top \Gamma A_K - \Gamma = -I_{n_z}, \tag{7.7b}$$

$$\tau = \lambda_{\max}(\Gamma) \cdot \operatorname{trace}\left(\Sigma_V \widetilde{H}^\top \Gamma \widetilde{H}\right). \tag{7.7c}$$

Moreover, we assume that there exists a set $\mathbb{Z}_{\mathrm{F}} \subseteq \mathbb{R}^{n_z}$ such that for all $Z(N \mid k)$ satisfying (7.5g), we have $A_K \mathbb{E}[Z(N \mid k)] \in \mathbb{Z}_{\mathrm{F}}$, $U(N \mid k) = KZ(N \mid k)$ satisfies (7.5e), and $Y(N \mid k) = \widetilde{H}^\top A_K Z(N \mid k) + V(N \mid k)$ satisfies (7.5f).

Utilizing the not necessarily minimal state-space representation (7.3), Assumption 7.4 modifies the terminal ingredients of the model-based state-feedback case (Farina et al., 2013) to the output-feedback case. Specifically, (7.7a) is a Lyapunov equality tailored to the objective function (7.5a) with inputs and outputs; (7.7b)–(7.7c) ensure that the Γ-weighted norm of the covariance of the terminal extended state remains bounded in the presence of stochastic disturbances. Note that (7.7a)–(7.7b) are Lyapunov equalities. Thus, the solutions P and Γ exist provided that $A_K = \widetilde{A} + \widetilde{B}K$ is Schur stable and that $K^\top R K + A_K^\top \widetilde{H} Q \widetilde{H}^\top A_K$ is positive definite. In Section 7.2.2, we illustrate the crucial role that the Lyapunov equalities (7.7a)–(7.7b) play in ensuring closed-loop properties.

The main challenge of adapting (7.5) to a data-driven predictive output-feedback setting is threefold:

i) The data-driven tractable reformulation of (7.5) involves infinite-dimensional L^2 random variables and chance constraints with distributional ambiguity. We address this issue in Section 7.1.2 using ideas from Section 6.3.

ii) The design of a predictive control algorithm with recursive feasibility guarantees can be achieved via the design of the initial condition; see Section 7.2.

iii) Additionally, the terminal ingredients need to be designed without explicit knowledge of the system matrices; see Section 7.1.3.

7.1.2. Data-Driven Reformulation in PCE Coefficients

Given $\bar{Z}(k) \in \mathcal{C}(m(k), \Sigma(k))$ and $V_{\mathrm{f}} \in \mathcal{C}(0, I_N \otimes \Sigma_V)$, we construct their bases and PCEs using covariance decomposition as per Lemma 3.2. With the decompositions $\mathsf{M}_k \mathsf{M}_k^\top = \Sigma(k)$ and $\mathsf{M}_v \mathsf{M}_v^\top = \Sigma_V$, with $\mathsf{M}_k \in \mathbb{R}^{n_z \times n_\xi}$ and $\mathsf{M}_v \in \mathbb{R}^{n_y \times n_\eta}$, the bases and PCEs are constructed as follows

$$\begin{aligned} &\mathcal{P}_1(\xi(k)), \quad \bar{Z}(k) = m(k) + \mathsf{M}_k \xi(k), t \in \mathbb{I}_{[0,N-1]}. \\ &\mathcal{P}_1(\eta(k)), \quad V(k+t) = \mathsf{M}_v \eta(k+t), t \in \mathbb{I}_{[0,N-1]}. \end{aligned}$$

Moreover, uncertainty propagation through the predicted dynamics (7.5b)–(7.5c) with affine policies (7.5d) can be achieved using exact PCEs in the following orthonormal basis

$$\{\phi^j(\boldsymbol{\xi}(k))\}_{j=0}^{L-1} \doteq \mathcal{P}_1(\boldsymbol{\xi}(k)), \quad \boldsymbol{\xi}(k) \doteq [\xi^\top(k), \eta^\top(k), ..., \eta^\top(k+N-1)]^\top, \tag{7.8}$$

where $L = Nn_\eta + n_\xi + 1$, cf. Lemma 3.3. Due to the distributionally robust formulation of OCP (7.5), we do not require the exact information of $\xi(k)$ and $\eta(k)$ for OCP (7.5). Instead, we consider all possible $\xi(k) \in \mathcal{C}(0, I_{n_\xi})$ and $\eta(k+t) \in \mathcal{C}(0, I_{n_\eta})$, $t \in \mathbb{I}_{[0,N-1]}$, cf. (3.20). In other words, the constraints of OCP (7.5) are required to hold for all $\boldsymbol{\xi}(k) \in \mathcal{C}(0, I_{Nn_\eta + n_\xi})$.

Applying Galerkin projection onto the basis (7.8) yields the dynamics of the PCE coefficients, for all $j \in \mathbb{I}_{[0,L-1]}$ and $t \in \mathbb{I}_{[0,N-1]}$,

$$\mathsf{z}^j(t+1 \,|\, k) = \widetilde{A}\mathsf{z}^j(t \,|\, k) + \widetilde{B}\mathsf{u}^j(t \,|\, k) + \widetilde{H}\mathsf{v}^j(t \,|\, k), \quad \mathsf{y}^j(t \,|\, k) = \widetilde{H}^\top \mathsf{z}^j(t \,|\, k+1) \tag{7.9a}$$

$$\mathsf{z}^j(0 \,|\, k) = \bar{\mathsf{z}}^j(k), \tag{7.9b}$$

where the PCE coefficients of $\bar{Z}(k)$ and $V_{\mathrm{f}}(k)$ for $\mathrm{f} = [0, N-1]$ are

$$\begin{aligned} \bar{\mathsf{z}}^0(k) \doteq m(k), \, \mathcal{H}_1\left(\bar{\mathsf{z}}^{[1,n_\xi]}(k)\right) \doteq \mathsf{M}_k, \, \bar{\mathsf{z}}^j(k) \doteq 0, \forall j \in \mathbb{I}_{[n_\xi+1,L-1]}, \\ \mathsf{v}_{\mathrm{f}}^0(k) \doteq 0, \, \mathcal{H}_1\left(\mathsf{v}_{\mathrm{f}}^{[1,L-1]}(k)\right) \doteq I_N \otimes \mathsf{M}_v, \end{aligned} \tag{7.10}$$

Note that the distributional ambiguity surrounding $V_{\mathrm{f}}(k)$ and $\bar{Z}(k)$ leads to a situation where the predictions obtained through (7.5b)–(7.5d) are characterized by ambiguous distributions. However, representing all random variables by their PCEs in the basis (7.8), we obtain deterministic coefficients that correspond to the mean and the covariance decomposition. In other words, the distributional ambiguities are encoded in the basis $\{\phi^j(\boldsymbol{\xi}(k))\}_{j=0}^{L-1}$ with $\boldsymbol{\xi}(k) \in \mathcal{C}(0, I_{Nn_\eta + n_\xi})$, cf. (3.20).

Before presenting the data-driven reformulation of OCP (7.5), we first obtain the PCE reformulation of the objective J_N in (7.5a). Recall that (7.5a) is a conditional expectation for the outcomes $\{\omega \,|\, \bar{Z}(k, \omega) = z(k)\}$. Using the PCE coefficients of $\bar{Z}(k)$ in (7.10), we reformulate the set of corresponding outcomes accordingly

$$\left\{\omega \,\middle|\, \bar{\mathsf{z}}^0(k) + \mathcal{H}_1\left(\bar{\mathsf{z}}^{[1,n_\xi]}(k)\right)\xi(k, \omega) = z(k)\right\}. \tag{7.11}$$

At this point, there are two principal options: either determine $\xi(k,\omega)$ prior to optimization or treat it as an additional decision variable. For the further developments, we temporarily assume that $\xi(k,\omega)$ is known, while we show how it can be determined in Section 7.2.1.

Given the realization $\xi(k,\omega)$, the reformulation of (7.5a) reads

$$
\begin{aligned}
J_N &= \sum_{t=0}^{N-1} \left(\ell_{u,t}(k) + \ell_{y,t}(k)\right) + \ell_{z,N}(k), \qquad (7.12)\\
\ell_{u,t}(k) &= \mathbb{E}\left[\| \textstyle\sum_{j=0}^{L-1} \mathsf{u}^j(t \mid k)\, \phi^j(\boldsymbol{\xi}(k)) \|_R^2 \,\middle|\, \xi(k,\omega) \right] \\
&= \|\mathsf{u}^0(t \mid k) + \mathcal{H}_1\left(\mathsf{u}^{[1,n_\xi]}(t \mid k)\right)\xi(k,\omega)\|_R^2 + \textstyle\sum_{j=n_\xi+1}^{L-1} \|\mathsf{u}^j(t \mid k)\|_R^2, \\
\ell_{y,t}(k) &= \|\mathsf{y}^0(t \mid k) + \mathcal{H}_1\left(\mathsf{y}^{[1,n_\xi]}(t \mid k)\right)\xi(k,\omega)\|_Q^2 + \textstyle\sum_{j=n_\xi+1}^{L-1} \|\mathsf{y}^j(t \mid k)\|_Q^2, \\
\ell_{z,N}(k) &= \|\mathsf{z}_N^0(k) + \mathcal{H}_1\left(\mathsf{z}^{[1,n_\xi]}(N \mid k)\right)\xi(k,\omega)\|_P^2 + \textstyle\sum_{j=n_\xi+1}^{L-1} \|\mathsf{z}^j(N \mid k)\|_P^2.
\end{aligned}
$$

Now, consider a realization trajectory of length T, $(u^{\mathrm{d}}, y^{\mathrm{d}}, v^{\mathrm{d}})_{[0,T-1]}$, from (7.4). Let p and f denote the ranges $[-\ell,-1]$ and $[0,N-1]$, respectively, and follow the same partitioning of Hankel matrices as in (2.26). Using the extended stochastic fundamental lemma from Corollary 5.1, we are ready to state the data-driven reformulation of OCP (7.5). At time instant k, given the PCE coefficients $\bar{\mathsf{z}}^{[0,L-1]}(k)$ and $\mathsf{v}_{\mathrm{f}}^{[0,L-1]}(k)$ as in (7.10), we solve

$$
\begin{aligned}
\min_{\bar{u}_{\mathrm{f}}(k),\mathbf{K}_N(k),\mathsf{u}_{\mathrm{f}}(k),\mathsf{y}_{\mathrm{f}}(k),\mathsf{g}} \quad & J_N\left(\mathsf{u}_{\mathrm{f}}(k), \mathsf{y}_{\mathrm{f}}(k), \mathsf{z}^{[0,L-1]}(N \mid k), \xi(k,\omega)\right) && (7.13a)\\
\text{subject to } & \forall j \in \mathbb{I}_{[0,L-1]},\ \forall t \in \mathbb{I}_{[0,N-1]} \\
& \begin{bmatrix} \mathcal{H}_{u,\mathrm{f}} \\ \mathcal{H}_{y,\mathrm{f}} \\ \mathcal{H}_{v,\mathrm{f}} \end{bmatrix} \mathsf{g}^j = \begin{bmatrix} \mathsf{u}_{\mathrm{f}}^j(k) \\ \mathsf{y}_{\mathrm{f}}^j(k) \\ \mathsf{v}_{\mathrm{f}}^j(k) \end{bmatrix} && (7.13b)\\
& \begin{bmatrix} \mathcal{H}_{u,\mathrm{p}} \\ \mathcal{H}_{y,\mathrm{p}} \end{bmatrix} \mathsf{g}^j = \bar{\mathsf{z}}^j(k) && (7.13c)\\
& \mathsf{u}_{\mathrm{f}}^j(k) = \delta^{0j}\bar{u}_{\mathrm{f}}(k) + \mathbf{K}_N(k) \begin{bmatrix} \bar{\mathsf{z}}^j(k) \\ \mathsf{v}_{\mathrm{f}}^j(k) \end{bmatrix} && (7.13d)\\
& a_{u,i}^\top \mathsf{u}_{t|k}^0 + \sigma(\varepsilon_u)\|a_{u,i}^\top \mathcal{H}_1\left(\mathsf{u}^{[1,L-1]}(t \mid k)\right)\| \le 1, \forall i \in \mathbb{I}_{[1,N_u]} && (7.13e)\\
& a_{y,i}^\top \mathsf{y}_{t|k}^0 + \sigma(\varepsilon_y)\|a_{y,i}^\top \mathcal{H}_1\left(\mathsf{y}^{[1,L-1]}(t \mid k)\right)\| \le 1, \forall i \in \mathbb{I}_{[1,N_y]} && (7.13f)\\
& \begin{bmatrix} \mathsf{u}^j_{[N-\ell,N-1]}(k) \\ \mathsf{y}^j_{[N-\ell,N-1]}(k) \end{bmatrix} = \mathsf{z}^j(N \mid k), \quad \mathsf{z}_N^0(k) \in \mathbb{Z}_{\mathrm{F}}, \quad \textstyle\sum_{j=1}^{L-1} \|\mathsf{z}_N^j(k)\|_\Gamma^2 \le \tau. && (7.13g)
\end{aligned}
$$

Here, $\mathsf{u}_{\mathrm{f}}(k)$ denotes the shorthand for the collection of all PCE coefficients of $U_t(k)$, i.e. $\mathsf{u}^j(t \mid k)$ for $j \in \mathbb{I}_{[0,L-1]}$ and $t \in \mathbb{I}_{[0,N-1]}$. Similarly, we define $\mathsf{y}_{\mathrm{f}}(k)$ and g.

The equivalence of OCP (7.5) and OCP (7.13) relies on several factors similar to those in Theorem 4.3. First, the PCE representations of all random variables are exact by the construction of basis (7.8). Second, by Corollary 5.1, the data-driven reformulation of the system dynamics from (7.5b)–(7.5c) to (7.13b)–(7.13c) is exact if the disturbance realization $v^{\mathrm{d}}_{[\ell,T-1]}$ is known exact. Third, the reformulation of the distributionally robust chance constraints from (7.5e)–(7.5f) to (7.13e)–(7.13f) is also exact by Proposition 6.1 with $\rho = 0$. Note that (7.13) does not explicitly depend on the basis $\{\phi^j(\boldsymbol{\xi})\}_{j=0}^{L-1}$ and thus the knowledge of the exact distribution of $\boldsymbol{\xi} \in \mathcal{C}(0, I_{Nn_\eta+n_\xi})$ is not required.

7.1.3. Data-Driven Design of Terminal Ingredients

We now turn to the data-driven design of terminal ingredients. Consider a realization trajectory $(u^{\mathrm{d}}, y^{\mathrm{d}}, v^{\mathrm{d}})_{[0,T-1]}$ of (7.4), with given initial condition $(u^{\mathrm{d}}, y^{\mathrm{d}})_{[-\ell,-1]}$ and disturbance realizations $v^{\mathrm{d}}(k)$ for $k \in \mathbb{I}_{[0,T]}$. Here, the subscript "d" indicates recorded data. We have the following assumption regarding the data.

Assumption 7.5: The block matrix $\begin{bmatrix} \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix} \in \mathbb{R}^{(n_z+n_u)\times T}$ is of full row rank.

For LTI systems without disturbance, a sufficient condition for Assumption 7.5 to hold is $\ell \cdot n_y = n_x$, where n_x is the minimal state dimension; see Lemma 13 of Berberich et al., 2021. This condition requires the system to admit a specific structure, and hence renders Assumption 7.5 relatively restrictive in disturbance-free scenarios. However, when disturbances are present, the pair $\left(\widetilde{A}, \left[\widetilde{B}, \widetilde{H}\right]\right)$ in (7.3) is controllable, as shown in Lemma 5.2. This implies that a sufficient condition for Assumption 7.5 is that $(u, v)^{\mathrm{d}}_{[0,T-1]}$ is persistently exciting of order $n_z + 1$ as shown by Corollary 2 of Willems et al., 2005. Therefore, the full-rank condition in Assumption 7.5 becomes less restrictive in the presence of disturbances.

Next, we provide a data-driven surrogate for $A_K = \widetilde{A} + \widetilde{B}K$ as an extension of the state-feedback design by De Persis et al., 2019.

Lemma 7.1 (Data-driven surrogate of A_K): Let Assumption 7.5 hold. Then, for any given $K \in \mathbb{R}^{n_u \times n_z}$, there exists a $\Psi \in \mathbb{R}^{T \times n_z}$ such that

$$\begin{bmatrix} \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix} \Psi = \begin{bmatrix} I \\ K \end{bmatrix}. \tag{7.14}$$

The data-driven reformulation of $A_K \doteq \widetilde{A} + \widetilde{B}K$ reads

$$A_K = S\Psi, S \doteq \begin{bmatrix} \bar{A}\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) + \bar{B}\mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) - \mathcal{H}_1\left(v^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix} \in \mathbb{R}^{n_z \times T}.$$

Proof. Observe that (7.14) holds given the full row rank condition of Assumption 7.5. Horizontally stacking the data of the realization dynamics in (7.4), we have

$$\mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) = \begin{bmatrix} \Xi & D \end{bmatrix} \begin{bmatrix} \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix} + \mathcal{H}_1\left(v^{\mathrm{d}}_{[0,T-1]}\right).$$

We arrive at $A_K = \begin{bmatrix} \widetilde{A} & \widetilde{B} \end{bmatrix} \begin{bmatrix} I_{n_z} \\ K \end{bmatrix} = \begin{bmatrix} \bar{A} & \bar{B} \\ \Xi & D \end{bmatrix} \begin{bmatrix} \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix} \Psi = S\Psi.$ □

A key difference between this result and the approach of De Persis et al., 2019 is that it leverages the known matrix components $\bar{A}$ and $\bar{B}$ from (2.22), whereas De Persis et al., 2019 treat the system matrices $[\widetilde{A}\ \widetilde{B}]$ as unknown.

Using the data-driven surrogate $S\Psi$ for A_K, we now focus on constructing K, P, and Γ. Observe that the Lyapunov equality (7.7a) is equivalent to

$$A_K^\top \left(P + \widetilde{Q}\right) A_K - \left(P + \widetilde{Q}\right) = -\left(K^\top R K + \widetilde{Q}\right)$$

with $\widetilde{Q} \doteq \widetilde{H} Q \widetilde{H}^\top \succeq 0$. Since this structure resembles the closed-loop Lyapunov equation of an LQR design, we select K as the LQR feedback gain for the state-space model (7.3).

For stabilizable $(\widetilde{A}, \widetilde{B})$, Dörfler et al., 2023b show that the LQR design for (7.3) is equivalent to

$$\min_{\widetilde{P} \in \mathbb{R}^{n_z \times n_z}, K \in \mathbb{R}^{n_u \times n_z}} \operatorname{trace}\left(\widetilde{Q}\widetilde{P} + K^\top R K \widetilde{P}\right) \tag{7.15a}$$

$$\text{subject to} \quad \widetilde{P} \succeq \widetilde{H}\widetilde{H}^\top, \quad (\widetilde{A} + \widetilde{B}K)\widetilde{P}(\widetilde{A} + \widetilde{B}K)^\top - \widetilde{P} + \widetilde{H}\widetilde{H}^\top \preceq 0. \tag{7.15b}$$

Replacing $\widetilde{A} + \widetilde{B}K = S\Psi$ and applying the convexification procedure outlined by Dörfler et al., 2023b, we obtain the following data-driven and tractable reformulation of (7.15)

$$\min_{X_1 \in \mathbb{R}^{n_u \times n_u}, X_2 \in \mathbb{R}^{T \times n_z}} \operatorname{trace}(\widetilde{Q}\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) X_2) + \operatorname{trace}(X_1) \tag{7.16a}$$

$$\text{subject to} \quad \begin{bmatrix} \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) X_2 - \widetilde{H}\widetilde{H}^\top & SX_2 \\ \star & \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) X_2 \end{bmatrix} \succeq 0, \tag{7.16b}$$

$$\begin{bmatrix} X_1 & R^{\frac{1}{2}}\mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) X_2 \\ \star & \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) X_2 \end{bmatrix} \succeq 0, \tag{7.16c}$$

with $K = \mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) X_2 \left(\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) X_2\right)^{-1}$ and $\Psi = X_2 \left(\mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) X_2\right)^{-1}$.

Using K and Ψ, we compute P and Γ by solving

$$\Psi^\top S^\top (P + \widetilde{Q}) S\Psi - P = -K^\top R K, \tag{7.17a}$$

$$\Psi^\top S^\top \Gamma S \Psi - \Gamma = -I_{n_z}. \tag{7.17b}$$

We compute τ using (7.7c). Finally, we define the terminal constraint on the expected value $\mathbb{Z}_{\mathrm{F}}$ as a sublevel set of the terminal cost: $\mathbb{Z}_{\mathrm{F}} \doteq \{z \in \mathbb{R}^{n_z} \mid \|z\|_P^2 \leq \varepsilon_z\}$ where $\varepsilon_z \in \mathbb{R}^+$ is chosen sufficiently small to ensure that Assumption 7.4 holds.

7.2. Closed-Loop Guarantees

Next, we introduce our output-feedback stochastic data-driven predictive control scheme based on OCP (7.13). To this end, we design a stochastic initial condition to ensure recursive feasibility and then conduct a closed-loop analysis of the proposed predictive control scheme.

7.2.1. Recursive Feasibility

Given the measured past ℓ input and output realizations, i.e. the current realization $z\,(k)$ of the extended state $Z\,(k)$, one can set the initial condition of (7.13) as

$$\bar{\mathsf{z}}^{j}\,(k) = \delta^{0j} z\,(k)\,, j \in \mathbb{I}_{[0,L-1]} \tag{7.18}$$

with δ^{0j} as the Kronecker delta. In other words, we construct a deterministic initial condition by assigning the measured value to the 0-th PCE coefficient while setting higher-order PCE coefficients, which represent stochastic uncertainties, to zero. Henceforth, we refer to (7.18) as the *measured initial condition*.

However, at time step k, OCP (7.13) may become infeasible, for example, if the disturbance V has infinite support and a large realization $v\,(k-1)$ occurs at step $k-1$. If so, we rely on a so-called *backup initial condition* to ensure recursive feasibility. For the underlying concept from model-based stochastic predictive control, we refer to Farina et al., 2013; Hewing et al., 2018.

Similar to deterministic predictive control, one way for constructing a backup initial condition $\bar{Z}\,(k)$ is to consider its predicted value based on the optimal solution of the last OCP solved, i.e. $\bar{Z}\,(k) = Z^{\star}\,(1\,|\,k-1)$. However, the one-step prediction of $Z^{\star}\,(1\,|\,k-1)$ by (7.3) with the last initial condition $\bar{Z}\,(k-1)$ involves the stochastic uncertainty of $V\,(k-1)$. This leads to the growth of the PCE basis; see, e.g. Pan et al., 2025b. Here, using the insights of Lemma 3.2, we update the basis (7.8) at each time instant to keep its dimensionality constant.

For all $k \in \mathbb{N}$, we can define a feasibility-preserving initial condition for OCP (7.13) in this updated basis. At time instant $k-1$, suppose OCP (7.13) is feasible with the initial condition $\bar{\mathsf{z}}\,(k-1)$, where $\bar{\mathsf{z}}^{j}\,(k-1) = 0$ for all $j \in \mathbb{I}_{[1+n_\xi,L-1]}$. The optimal solution of the OCP determines a prediction

$$Z^{\star}\,(1\,|\,k-1) = \mathsf{z}^{0,\star}\,(1\,|\,k-1) + \mathcal{H}_1\left(\mathsf{z}^{[1,n_\xi+n_\eta],\star}\,(1\,|\,k-1)\right)\begin{bmatrix}\xi\,(k-1)\\ \eta\,(k-1)\end{bmatrix}.$$

Note that the above represents an exact PCE of $Z^{\star}\,(1\,|\,k-1)$ with dimensionality $n_\xi + n_\eta$. To avoid any dimension growth of the PCE basis, we design an orthonormal projection matrix $T\,(k) \in \mathbb{R}^{(n_\xi+n_\eta)\times n_\xi}$, i.e., $T^{\top}\,(k)\,T\,(k) = I_{n_\xi}$.

Proposition 7.1 (Recursive feasibility with updated basis (Pan et al., 2025a)): Let Assumptions 7.3–7.4 hold. At time instant $k-1$, suppose OCP (7.13) is feasible in the basis $\{\phi^{j}(\boldsymbol{\xi}\,(k-1))\}_{j=0}^{L-1}$ from (7.8). Then, at time instant k, for any orthonormal matrix $T\,(k) \in \mathbb{R}^{(n_\xi+n_\eta)\times n_\xi}$ satisfying $T^{\top}\,(k)\,T\,(k) = I_{n_\xi}$, considering $\{\phi^{j}(\boldsymbol{\xi}\,(k))\}_{j=0}^{L-1}$ with

$$\xi\,(k) \doteq T^{\top}\,(k)\begin{bmatrix}\xi\,(k-1)\\ \eta\,(k-1)\end{bmatrix},\ \boldsymbol{\xi}\,(k) \doteq [\xi^{\top}\,(k)\,,\eta^{\top}\,(k)\,,...,\eta^{\top}\,(k+N-1)]^{\top}, \tag{7.19}$$

OCP (7.13) is feasible with the initial condition $\bar{\mathsf{z}}\,(k)$ with $\bar{\mathsf{z}}^{j}\,(k) \doteq 0$ for all $j \in \mathbb{I}_{[1+n_\xi,L-1]}$ and

$$\bar{\mathsf{z}}^{0}\,(k) \doteq \mathsf{z}^{0,\star}\,(1\,|\,k-1)\,,\ \mathcal{H}_1\left(\bar{\mathsf{z}}^{[1,n_\xi]}\,(k)\right) \doteq \mathcal{H}_1\left(\mathsf{z}^{[1,n_\xi+n_\eta],\star}\,(1\,|\,k-1)\right)T\,(k)\,. \tag{7.20}$$

The detailed proof is given in Section A.1. Observe that the updated $\boldsymbol{\xi}\,(k)$ still satisfies $\boldsymbol{\xi}\,(k) \in \mathcal{C}(0, I_{Nn_\eta+n_\xi})$, ensuring that the chance constraint reformulation remains valid.

Design of the Projection $T(k)$

Given the backup initial condition (7.20) in terms of PCE coefficients, its random-variable counterpart is expressed as

$$\begin{aligned}\bar{Z}(k) &= \bar{\mathsf{z}}^0(k) + \mathcal{H}_1\left(\bar{\mathsf{z}}^{[1,n_\xi]}(k)\right)\xi(k)\\ &= \mathsf{z}^{0,\star}(1\,|\,k-1) + \mathcal{H}_1\left(\mathsf{z}^{[1,n_\xi+n_\eta],\star}(1\,|\,k-1)\right)T(k)\,T^\top(k)\begin{bmatrix}\xi(k-1)\\ \eta(k-1)\end{bmatrix}.\end{aligned} \tag{7.21}$$

Note that the orthonormal matrix $T(k) \in \mathbb{R}^{(n_\xi+n_\eta)\times n_\xi}$ satisfies

$$T^\top(k)\,T(k) = I_{n_\xi}, \quad T(k)\,T^\top(k) \neq I_{n_\xi+n_\eta}.$$

Thus, using the projection $T(k)$ to prevent the growth of the PCE basis, we do not achieve the equality $\bar{Z}(k) = Z^\star(1\,|\,k-1)$ almost surely, as considered by Pan et al., 2025b. Instead, we design $T(k)$ to ensure that the realizations of both the prediction $Z^\star(1\,|\,k-1)$ and the backup initial condition $\bar{Z}(k)$ coincide with the current measurement $z(k)$. This leads to a conditional equality

$$\bar{Z}(k,\omega) = Z^\star(1\,|\,k-1,\omega), \quad \forall\omega \in \{\omega\,|\,Z^\star(1\,|\,k-1,\omega) = z(k)\}. \tag{7.22}$$

To enforce (7.22), we treat $\xi(k) \in L^2(\Omega,\mathcal{F},\mu;\mathbb{R})$ in (7.21) as a scalar-valued random variable with realization

$$\xi(k,\omega) = \left\|\begin{bmatrix}\xi(k-1,\omega)\\ \eta(k-1,\omega)\end{bmatrix}\right\| \in \mathbb{R}. \tag{7.23a}$$

Moreover, we use the projection

$$T(k) = \frac{[\xi(k-1,\omega), \eta^\top(k-1,\omega)]^\top}{\|[\xi(k-1,\omega), \eta^\top(k-1,\omega)]^\top\|} \in \mathbb{R}^{(1+n_\eta)\times 1} \tag{7.23b}$$

ensuring that

$$\xi(k,\omega) = T^\top(k)\begin{bmatrix}\xi(k-1,\omega)\\ \eta(k-1,\omega)\end{bmatrix}, \quad T(k)\,T^\top(k)\begin{bmatrix}\xi(k-1,\omega)\\ \eta(k-1,\omega)\end{bmatrix} = \begin{bmatrix}\xi(k-1,\omega)\\ \eta(k-1,\omega)\end{bmatrix}.$$

Substituting this into (7.21) ensures the conditional equality (7.22). Furthermore, recall that the current measurement $z(k)$ is the realization of the prediction $Z^\star(1\,|\,k-1)$, i.e.,

$$Z_1^\star(k-1,\omega) = \mathsf{z}_1^{0,\star}(k-1) + \mathcal{H}_1\left(\mathsf{z}_1^{[1,1+n_\eta],\star}(k-1)\right)\begin{bmatrix}\xi(k-1,\omega)\\ \eta(k-1,\omega)\end{bmatrix} = z(k).$$

Since $\mathcal{H}_1\left(\mathsf{z}_1^{[2,1+n_\eta],\star}(k-1)\right) = \widetilde{H}\mathsf{M}_v$ by (7.9)–(7.10), we compute $\eta(k-1,\omega)$ using the least-square solution to the above equality

$$\eta(k-1,\omega) = \mathsf{M}_v^\dagger\widetilde{H}^\top\left(z(k) - \mathsf{z}_1^{0,\star}(k-1) - \mathsf{z}_1^{1,\star}(k-1)\,\xi(k-1,\omega)\right). \tag{7.23c}$$

Online Selection of Initial Conditions

To distinguish the initial conditions, we refer to (7.18) as the *measured initial condition* and to (7.20) as the *backup initial condition*. Moreover, we denote the optimal value function of OCP (7.13) obtained with the measured initial condition (7.18) as V_N^{m} and to its counterpart obtained with the backup initial condition (7.20) as V_N^{b}. The notion *backup initial condition* reflects the principle of using measured data as much as possible, as it provides feedback. Eventually, we consider an initial condition selection strategy similar to Farina et al., 2013:

i) At each time instant $k \in \mathbb{N}$, initialize $V_N^{\mathrm{b}} = +\infty$ and $V_N^{\mathrm{m}} = +\infty$.

ii) If the OCP (7.13) with the measured initial condition (7.18) is infeasible, then re-solve with the backup one (7.20). The optimal cost V_N^{b} is updated correspondingly.

iii) Otherwise, if the OCP (7.13) with the measured initial condition (7.18) is feasible, update the resulting optimal cost V_N^{m} and compare it to the performance $\bar{J}_N(k)$ obtained from a feasible candidate solution, whose construction is detailed in (A.4). If $V_N^{\mathrm{m}} \leq \bar{J}_N(k)$ holds, apply the optimal input obtained with the measured initial condition.

iv) Else, i.e., whenever $V_N^{\mathrm{m}} > \bar{J}_N(k)$ holds, we re-solve OCP (7.13) with the backup initial condition (7.20). Note that due to optimality the resulting V_N^{b} satisfies $V_N^{\mathrm{b}} \leq \bar{J}_N(k)$.

It is easy to see that this procedure ensures that the resulting optimal cost at time instant k satisfies

$$V_N(k) \doteq \min\left\{V_N^{\mathrm{m}}, V_N^{\mathrm{b}}\right\} \leq \bar{J}_N(k), \quad k \in \mathbb{N}. \tag{7.24}$$

Online Interpolation of Initial Conditions

Instead of the selection strategy described above, as proposed in Pan et al., 2023b, we consider an approach that interpolates between the measured and the backup initial conditions. In model-based stochastic predictive control this has been considered by Köhler et al., 2022b; Schlüter et al., 2023, 2022.

Specifically, we define the *interpolated initial condition* as

$$\begin{aligned} \bar{\mathsf{z}}^0(k) &= \mu z(k) + (1-\mu)\mathsf{z}^{0,\star}(1 \,|\, k-1), \quad 0 \leq \mu(k) \leq 1 \\ \mathcal{H}_1\left(\bar{\mathsf{z}}^{[1,n_\xi]}(k)\right) &= (1-\mu)\mathcal{H}_1\left(\mathsf{z}^{[1,n_\xi+n_\eta],\star}(1 \,|\, k-1)\right) T(k). \end{aligned} \tag{7.25}$$

Here, $\mu(k) \in [0,1]$ is an additional decision variable governing the interpolation between the PCE coefficients of the measured and the backup initial conditions. Due to optimality, this approach also ensures the resulting optimal cost to be minimized, i.e.

$$V_N(k) \leq \bar{J}_N(k), \forall k \in \mathbb{N}. \tag{7.26}$$

7.2.2. Closed-Loop Analysis

We are ready to present the main result, i.e., a data-driven stochastic output-feedback predictive control scheme with closed-loop guarantees based on OCP (7.13). The scheme

Algorithm 1 Output-feedback stochastic data-driven predictive control algorithm with OCP (7.13)

Input: T, N, $\bar{J}_N(0) \leftarrow +\infty$, $V_N^{\mathrm{m}} \leftarrow +\infty$, $V_N^{\mathrm{b}} \leftarrow +\infty$, $k \leftarrow 0$

Offline data collection:

1: Randomly sample $u^{\mathrm{d}}(k) \in \mathbb{U}$ for $k \in \mathbb{I}_{[0,T-1]}$
2: Apply $u^{\mathrm{d}}_{[0,T-1]}$ to system (7.1) under disturbances $v^{\mathrm{d}}_{[0,T-1]}$ for given $(u^{\mathrm{d}}, y^{\mathrm{d}})_{[-\ell,-1]}$, record $y^{\mathrm{d}}_{[0,T-1]}$
3: Record $v^{\mathrm{d}}_{[0,T-1]}$ or estimate $\hat{v}^{\mathrm{d}}_{[0,T-1]}$ via (5.16)
4: Determine terminal ingredients (K, P, Γ, τ) with $(u^{\mathrm{d}}, v^{\mathrm{d}}, y^{\mathrm{d}})_{[0,T-1]}$
5: Construct (7.13) with $(u^{\mathrm{d}}, v^{\mathrm{d}}, y^{\mathrm{d}})_{[0,T-1]}$ and P, Γ and τ

Online optimization with selection of initial conditions

1: Measure $z(k) = (u, y)_{[k-\ell,k-1]}$
2: **if** $k \geq 1$ **then** update $T(k)$, $\xi(k,\omega)$, $\eta(k-1,\omega)$ and $\bar{J}_N(k)$
3: **end if**
4: Solve OCP (7.13) with the measured initial condition (7.18)
5: **if** (7.13) is feasible with optimal cost $V_N^{\mathrm{m}} \leq \bar{J}_N(k)$ **then** continue
6: **else**
7: Solve (7.13) with the backup initial condition (7.20), update V_N^{b}
8: **end if**
9: Apply $u(k)$ from (7.6) to system (7.4)
10: $V_N(k) \leftarrow \min\{V_N^{\mathrm{m}}, V_N^{\mathrm{b}}\}$
11: $k \leftarrow k+1$, $V_N^{\mathrm{m}} \leftarrow +\infty$, $V_N^{\mathrm{b}} \leftarrow +\infty$ go to Step 1

Online optimization with interpolation of initial conditions

1: Measure $z(k) = (u, y)_{[k-\ell,k-1]}$
2: **if** $k = 0$ **then**
3: Solve OCP (7.13) with the measured initial condition (7.18)
4: **else** update $T(k)$, $\xi(k,\omega)$, $\eta(k-1,\omega)$
5: Solve (7.13) with the interpolated initial condition (7.25) with $0 \leq \mu(k) \leq 1$ as an additional decision variable
6: **end if**
7: Apply $u(k)$ from (7.6) to system (7.4)
8: $k \leftarrow k+1$, go to Step 1

consists of an offline data collection phase and an online optimization phase. We consider two variants of the online phase that includes different initial condition design strategies. The overall scheme is summarized in Algorithm 1.

In the offline phase, a random input trajectory $u^{\mathrm{d}}_{[0,T-1]}$ is generated to obtain $y^{\mathrm{d}}_{[0,T-1]}$ under disturbances $v^{\mathrm{d}}_{[0,T-1]}$. With the recorded input and output, we obtain the estimated disturbance data $\hat{v}^{\mathrm{d}}_{[0,T-1]}$. These recorded trajectories, along with the estimated disturbances, are then used to determine the terminal ingredients as proposed in Section 7.1.3 and to construct the Hankel matrices required for OCP (7.13).

In the online optimization phase, we assume that the OCP (7.13) is initially feasible with the measured initial condition at $k = 0$. At each time instant k, before applying the initial condition selection strategy, we update $\eta(k-1,\omega)$, $T(k)$, $\xi(k,\omega)$ using (7.23), based on the measured $z(k)$, the predicted $\mathsf{z}^\star(1 \mid k-1)$, and the previous $\xi(k-1,\omega)$. Subsequently, we update $\bar{J}_N(k)$ by evaluating (7.12) with $\xi(k,\omega)$ and the feasible candidate (A.4). After solving OCP (7.13) with initial condition selection, we obtain the

optimal policy $\bar{u}_0^\star(k)$ and $K_0^\star(k)$. Then, we obtain the input $u(k)$ from (7.6), which is a feedback of $z(k)$.

Alternatively, Algorithm 1 includes an online optimization phase that utilizes the interpolation strategy. By introducing an additional interpolation decision variable, this approach not only eliminates the need to solve the OCP twice, as required by the selection strategy, but also simplifies implementation by removing the need to store and adapt the parameters V_N^{m}, V_N^{b}, and $\bar{J}_N(k)$.

Stability Analysis via Averaged Performance Bound

For a given initial condition $\bar{z}(0)$ and a disturbance realization trajectory $\{v(k)\}_{k\in\mathbb{N}}$, we obtain a realization trajectory $(u, v, y, z)_{k\in\mathbb{N}}$ of (7.4) by applying the feedback input $u(k)$ computed from Algorithm 1. Moreover, Algorithm 1 generates a sequence $V_N(k) \in \mathbb{R}$, $k \in \mathbb{N}$, corresponding to the optimal value functions arising in the closed loop.

Accordingly, for a probabilistic initial condition $\bar{Z}(0)$ and disturbance $V(k)$, $k \in \mathbb{N}$, we determine a stochastic trajectory $(Z, U, V, Y)_{k\in\mathbb{N}}$ of (7.1) by setting the realization of $U(k)$ as $u(k)$. Similarly, we define the probabilistic optimal cost $\mathcal{V}_N(k) \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R})$ as $\mathcal{V}_N(k, \omega) \doteq V_N(k)$. Thus, the tuple $(Z, U, V, Y, \mathcal{V}_N)_{k\in\mathbb{N}}$ represents the probabilistic collection of all closed-loop realizations.

The next theorem summarizes the closed-loop properties of the proposed scheme.

Theorem 7.1 (Closed-loop properties (Pan et al., 2025a)): Let Assumptions 7.3–7.4 hold and let the closed-loop stochastic trajectory $(Z, U, V, Y, \mathcal{V}_N)_{k\in\mathbb{N}}$ be determined by Algorithm 1 with perfectly measured disturbance realizations $v^{\mathrm{d}}_{[0,T-1]}$. Suppose that at time instant $k = 0$, OCP (7.13) is feasible under the measured initial condition (7.18). Then, for

$$\beta \doteq \operatorname{trace}\left(\Sigma_V(Q + \widetilde{H}^\top P \widetilde{H})\right) \in \mathbb{R}^+ \tag{7.27a}$$

the following statements hold:

i) For all $k \in \mathbb{N}$, OCP (7.13), with the selection (or interpolation) of initial conditions from Algorithm 1, is feasible.

ii) The probabilistic optimal costs at consecutive time instants satisfy

$$\mathbb{E}\left[\mathcal{V}_N(k+1) - \mathcal{V}_N(k)\right] \leq -\mathbb{E}\left[\|U(k)\|_R^2 + \|Y(k)\|_Q^2\right] + \beta. \tag{7.27b}$$

iii) The average asymptotic cost is bounded from above by

$$\lim_{k\to\infty} \frac{1}{k} \sum_{i=0}^{k} \mathbb{E}\left[\|U(i)\|_R^2 + \|Y(i)\|_Q^2\right] \leq \beta. \tag{7.27c}$$

The detailed proof is given in Section A.2.

We note that recursive feasibility of the OCP (statement i)) does not guarantee the recursive satisfaction of chance constraints in closed loop, which can be interpreted as the recursive satisfaction of one-step conditional probability constraints, i.e., $\mathbb{P}[a_y^\top Y(k) \leq 1 \,|\, k] \geq \gamma_y$. While one-step chance constraints hold when measured initial conditions are used, with backup initial conditions only $\mathbb{P}[a_y^\top Y(k) \leq 1, |, k-t] \geq \gamma_y$ can be guaranteed, where $k-t$, $t \in \mathbb{I}_{[0,k-1]}$, is the last time a measured initial condition was applied. We refer the readers to Farina et al., 2016, Remark 1 for a comprehensive discussion.

Moreover, similar to Cannon et al., 2009; Hewing et al., 2020b, statements ii)-iii) utilize the average asymptotic performance bound as a notion of stability for the stochastic closed-loop dynamics. We note that β stands for the asymptotic average cost calculated under $U(k) = KZ(k)$ with K corresponding to P via (7.7a). Thus, the proposed scheme guarantees either improved or equal asymptotic average cost compared to the feedback policy $U(k) = KZ(k)$ (Hewing et al., 2020b). The equivalence holds when K is determined by a linear-quadratic Riccati design.

Remark 7.1 (Extension to non-i.i.d. disturbances): If the disturbances have non-zero mean, one can consider an offset-free data-driven scheme to ensure recursive feasibility and performance guarantees (Lazar et al., 2022). Conceptually, this offset-free scheme considers the stage cost $\|Y(k)\|_Q^2 + \|U(k) - U(k-1)\|_R^2$ to compensate the mean part of the disturbance. For cases of zero-mean disturbance, if $V(k+t)$ are independent but not identically distributed for different $t \in \mathbb{I}_{[0,N-1]}$, one can still use the basis (7.8). However, one will obtain different PCE coefficients for each $V(k+t)$. To ensure the closed-loop properties of Theorem 7.1, one can choose Σ_V to be element-wise larger than all $\Sigma[V(k)]$, $k \in \mathbb{N}$, in this case. However, handling dependent disturbances in the closed-loop analysis remains an open problem.

Robustness Analysis with Estimated Disturbances

The previous discussions are based on the assumption of exact measurement of past disturbances. Next, we analyze the closed-loop behavior of Algorithm 1 when past disturbances are estimated using the method described in (5.16).

Lemma 7.2 (Data-driven one-step representation): Let Assumption 7.5 hold. Then, system (7.4) has the following equivalent representation

$$y(k) = \left(\mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) - \mathcal{H}_1\left(v^{\mathrm{d}}_{[0,T-1]}\right)\right) \begin{bmatrix} \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix}^{\dagger} \begin{bmatrix} z(k) \\ u(k) \end{bmatrix} + v(k). \tag{7.28a}$$

Moreover, replacing v^{d} with its least-squares estimate $\hat{v}^{\mathrm{d}}$ from (5.16), we have the one-step prediction

$$\hat{y}(k) = \mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) \begin{bmatrix} \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix}^{\dagger} \begin{bmatrix} z(k) \\ u(k) \end{bmatrix} + v(k). \tag{7.28b}$$

Proof. The proof of (7.28a) follows arguments similar to those used in the state-feedback case in Theorem 1 of De Persis et al., 2019. Moreover, with v^{d} replaced by its least-squares estimate $\hat{v}^{\mathrm{d}}$, we have

$$\begin{aligned} \hat{y}(k) &= \left(\mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) - \mathcal{H}_1(\hat{v}^{\mathrm{d}}_{[0,T-1]})\right) \begin{bmatrix} \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix}^{\dagger} \begin{bmatrix} z(k) \\ u(k) \end{bmatrix} + v(k) \\ &\overset{(5.16)}{=} \mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right) \begin{bmatrix} \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix}^{\dagger} \begin{bmatrix} z(k) \\ u(k) \end{bmatrix} + v(k). \end{aligned}$$

□

Notice that the Hankel matrix equation (7.13b) in OCP (7.13) holds for $(u, y, \hat{v})^{\mathrm{d}}_{[0,T-1]}$ when combined with the estimated disturbances $\hat{v}^{\mathrm{d}}_{[0,T-1]}$. They characterize the system dynamics of (7.28b). We define the underlying model mismatch between (7.28a) and (7.28b), which arises due to disturbance estimation error, as

$$\Delta \doteq \mathcal{H}_1\left(v^{\mathrm{d}}_{[0,T-1]}\right)\begin{bmatrix}\mathcal{H}_1\begin{pmatrix}z^{\mathrm{d}}_{[0,T-1]}\end{pmatrix}\\ \mathcal{H}_1\begin{pmatrix}u^{\mathrm{d}}_{[0,T-1]}\end{pmatrix}\end{bmatrix}^{\dagger} \in \mathbb{R}^{n_y\times(n_z+n_u)}. \tag{7.29}$$

Assumption 7.6 (Bounded optimal solution): For all $\bar{\mathsf{z}}^{[0,L-1]}(k)$ for which OCP (7.13) is feasible, the set of optimal solutions $(\mathsf{u}^\star(k), \mathsf{y}^\star(k))$ of OCP (7.13) is non-empty and bounded.

Note that this assumption can be satisfied by imposing compact constraints on the decision variables.

Proposition 7.2 (Robustness of performance (Pan et al., 2025a)): Let Assumption 7.6 and the conditions of Theorem 7.1 hold. Consider OCP (7.13) with $v^{\mathrm{d}}_{[0,T-1]}$ replaced by $\hat{v}^{\mathrm{d}}_{[0,T-1]}$ from (5.16). At time instant $k = 0$, suppose that OCP (7.13) is feasible with the measured initial condition (7.18). Then, statement i) of Theorem 7.1 holds, and there exists $C \in \mathbb{R}^+$ such that

$$\mathbb{E}[\mathcal{V}_N(k+1) - \mathcal{V}_N(k)] \leq -\mathbb{E}\left[\|U(k)\|^2_R + \left\|Y(k) - \Delta\begin{bmatrix}Z(k)\\ U(k)\end{bmatrix}\right\|^2_Q\right] + \beta + \mathbb{E}\left[C\left\|\Delta\begin{bmatrix}Z(k)\\ U(k)\end{bmatrix}\right\|\right]. \tag{7.30}$$

The detailed proof is given in Section A.3. Note that (7.30) reduces to (7.27b) when there is no model mismatch, i.e., when $\Delta \to 0$.

7.3. Case Study: Aircraft Dynamics

The results of this section have been published in Pan et al., 2025a.

We consider an LTI aircraft model, exactly discretized with sampling time $t_s = 0.5$ s, taken from Maciejowski, 2002. The ARX matrices are

$$\Xi = \begin{bmatrix}-0.019 & -1.440 & -0.201 & 0.256 & 0.050 & 0.160 & -0.256 & 0.086\\ 0.711 & -1.800 & -4.773 & 3.688 & 0.650 & 2.982 & -2.688 & 1.707\\ 1.444 & -26.922 & -15.746 & 12.898 & 2.319 & 10.461 & -12.897 & 5.171\end{bmatrix},$$

and $D = 0_{3\times 1}$ with $n_y = 3$, $n_u = 1$, $\ell = 2$, and thus $n_z = 8$. A minimal state-space representation with $n_x = 4$ is given in Pan et al., 2023c. We consider the ARX dynamics with data as above, where $V(k), k \in \mathbb{N}$, are i.i.d. random variables following a Gaussian mixture model, used as the simulated plant. The underlying Gaussian mixture model is a mixture of

$$V \sim 0.5\mathcal{N}\left(\begin{bmatrix}-0.005\\ -0.5\\ -0.05\end{bmatrix}, \begin{bmatrix}0.0001 & 0 & 0\\ 0 & 1 & 0\\ 0 & 0 & 0.01\end{bmatrix}\right) + 0.5\mathcal{N}\left(\begin{bmatrix}0.005\\ 0.5\\ 0.05\end{bmatrix}, \begin{bmatrix}0.0001 & 0 & 0\\ 0 & 1 & 0\\ 0 & 0 & 0.01\end{bmatrix}\right)$$

with zero mean. Despite its name, Gaussian mixture models do not correspond to a single Gaussian distribution. They admit infinite support, which renders recursive feasibility and robust constraint satisfaction challenging. The corresponding

$$\mathsf{M}_v = \begin{bmatrix} 0.011 & 0.002 & 0.002 \\ 0.002 & 1.118 & 0.020 \\ 0.002 & 0.020 & 0.110 \end{bmatrix}$$

in (7.9) is determined as the principal square root of the disturbance covariance matrix. Note that Y^j denotes the j-th element of Y. We impose chance constraints $\mathbb{P}[Y^1\,(i \,|\, k) \geq -1] \geq \gamma_y$ and $\mathbb{P}[Y^1\,(i \,|\, k) \leq 1] \geq \gamma_y$ for all $i \in \mathbb{I}_{[0,N-1]}$, ensuring that the output remains within the specified bounds with probability at least γ_y. We choose $\gamma_y = 0.9$ which corresponds to $\alpha(\gamma_y) = 3$. The weighting matrices in the stage cost are given by $Q = \mathrm{diag}([1,1,1])$ and $R = 1$.

We compare two schemes:

I) Algorithm 1 with exact disturbance data.

II) Algorithm 1 with estimated disturbance data.

In the data collection phase of Scheme I), we record an input-output-disturbance trajectory of 1000 steps; while in Scheme II), only the input-output trajectory is recorded and subsequently used to estimate the disturbance realizations via (5.16).

We use the first 22 recorded input-output data pairs along with the measured or estimated disturbances to determine the terminal ingredients P, Γ, τ, and $\mathbb{Z}_{\mathrm{F}}$ for both schemes using the data-driven procedure outlined in Section 7.1.3. Note that $\ell \cdot n_y = n_x$ does not hold for this system, as $2 \cdot 3 \neq 4$, yet the generated data still satisfies Assumption 7.5. For the sake of comparison, we also generate a 22-step input-output trajectory with the same inputs and initial condition without disturbance. We observe that for this data Assumption 7.5 is not satisfied. This indicates that in the presence of disturbances, the data requirement condition of Assumption 7.5 is less restrictive than in the disturbance-free case.

We use `MOSEK` ApS, 2023 to solve (7.16) for K and H. Compared to the feedback $K^\star$ obtained by directly solving the usual discrete-time algebraic Riccati equation for the exact system matrices of (7.3), the relative difference is $\|K - K^\star\|_2^2 \cdot (\|K^\star\|_2^2)^{-1} = 0.088$ for Scheme II), while it is $1.01 \cdot 10^{-3}$ for Scheme I). Note that for both cases A_K is Schur stable. Finally, we calculate P and Γ solving (7.17) and thus obtain τ via (7.7c).

We implement Algorithm 1 with a prediction horizon $N = 10$ using `IPOPT` (Wächter et al., 2006). The computations are done on a virtual machine with an AMD EPYC processer with 2.8 GHz, 32GB of RAM in `julia`. Similar to before, in the data collection phase we use the first 90 recorded inputs-outputs and measured/estimated disturbances to construct the Hankel matrices. Applying Lemma 3.2 we obtain the PCE representation for each component of $V\,(k)$. Moreover, the dimension of the overall PCE basis is $L = 1 + n_\xi + Nn_\eta = 32$ and the number of decision variables for the considered OCP is 3488.

We sample 1000 different disturbance realization sequences of length 31 each and then compute the corresponding closed-loop responses of the two schemes using distributed computing in `julia`. The computation times, the average asymptotic costs, and the highest empirical relative frequency of constraint violations $\max_k \mathbb{P}^{\mathrm{l}}\,(k)$ and

Table 7.1.: Comparison of the computation times, the average costs, and the number of violations for 1000 realized closed-loop trajectories with a length of 31 steps for measured and estimated disturbances for Schemes I) and II).

Scheme	Computation time		$\bar{\ell}$	$\max_k \mathbb{P}^{\mathrm{l}}(k)$	$\max_k \mathbb{P}^{\mathrm{u}}(k)$
	Mean [s]	SD [s]			
I	0.66	0.050	537.99	0.3%, $k = 11$	3.1%, $k = 3$
II	0.68	0.049	543.54	0.2%, $k = 7$	4.1%, $k = 3$

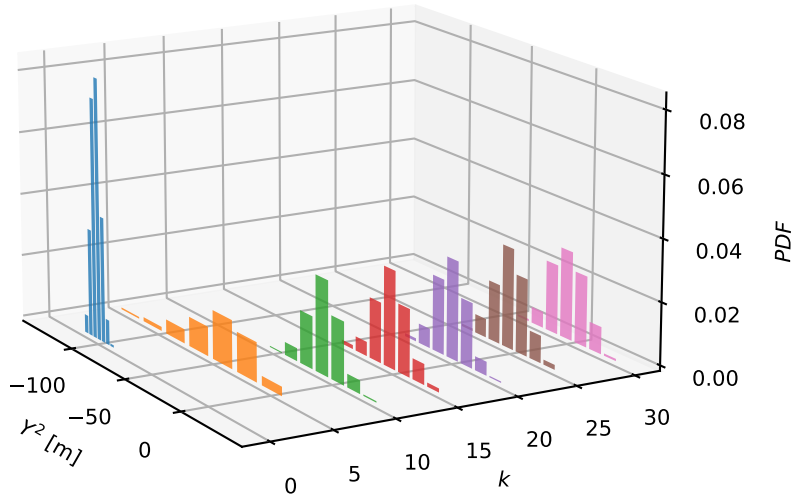

Figure 7.1.: Histograms of the output Y^2 from 1000 closed-loop realization trajectories of Scheme II)

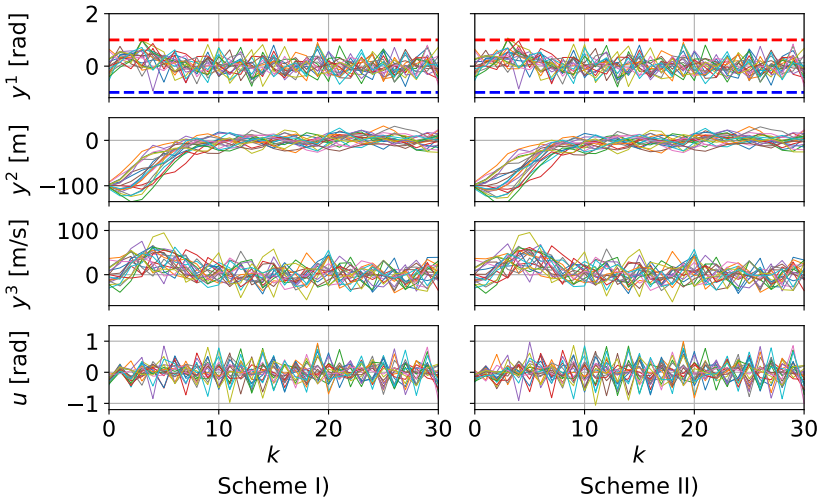

Figure 7.2.: 20 different closed-loop realization trajectories of Algorithm 1. Blue-dashed line: $y^1 = -1$; red-dashed line: $y^1 = 1$. Left: Scheme I) with measured disturbances; right: Scheme II) with estimated disturbances.

$\max_k \mathbb{P}^{\mathrm{u}}(k)$ are summarized in Tab. 7.1. Here, the empirical relative frequency of constraint violation $\mathbb{P}^{\mathrm{l}}(k)$ represents the percentage of time instants where $y^1(k) < -1$ over all 1000 realizations at time instant k, and $\mathbb{P}^{\mathrm{u}}(k)$ is defined similarly for $y(k)^1 > 1$.

Due to the distributionally robust formulation of (7.13), the highest empirical relative frequencies of constraint violation for Schemes I) and II) is substantially lower than the allowed violation probability of 10%. The two schemes exhibit similar results when compared based on the mean and the Standard Deviation (SD) of the computation time of each time step evaluated in the closed loop.

Moreover, evaluating the closed-loop realization trajectories over time, we obtain the average cost

$$\bar{\ell} \doteq \mathbb{E}\left[\sum_{k=11}^{30} \ell(k)\right]/20, \quad \ell(k) \doteq \|u\,(k)\,\|_R^2 + \|y\,(k)\,\|_Q^2.$$

The two schemes achieve similar average costs. Moreover, since K is determined from data-driven LQR design as in Section 7.1.3, the theoretical value of the average asymptotic cost for Scheme I) is $\beta = 540.94$ by Theorem 7.1, which is consistent with the simulation results of Scheme I) with $\bar{\ell} = 537.99$ and Scheme II) with $\bar{\ell} = 543.54$.

Additionally, for Scheme II), Figure. 7.1 depicts the time evolution of the (normalized) histograms of the output realizations y^2 at $k = 0, 5, 10, 15, 20, 25, 30$, where the vertical axis represents the (approximated) probability density of Y^2. As one can see, the proposed scheme controls the system to a stationary distribution of Y^2 centered at 0. With 20 different sampled sequences of disturbance realizations, we show the corresponding closed-loop realization trajectories of Schemes I) and II) in Figure 7.2. Both schemes exhibit similar closed-loop responses.

7.4. Summary

This chapter has presented and analyzed a data-driven output-feedback predictive control scheme for stochastic LTI systems, building on the results of Pan et al., 2023b, 2025a,b. Using a stochastic variant of the fundamental lemma—which relies on polynomial chaos expansions—we have presented sufficient conditions for recursive feasibility and characterized the closed-loop performance. Moreover, we have discussed the data-driven design of terminal ingredients that satisfy the proposed conditions.

It deserves to be noted that the proposed approach requires only knowledge of the first two moments, without needing exact information about the underlying distribution. Thus, the proposed scheme incorporates a weak form of distributional robustness and is applicable to a broad class of both Gaussian and non-Gaussian settings. A numerical example has demonstrated the efficacy of the proposed scheme in handling Gaussian mixture models. Future work will focus on enhancing robustness, extending the approach to nonlinear systems, and incorporating regularization techniques for measurement noise from Berberich et al., 2020; Coulson et al., 2019.

8. Conclusions and Perspectives

In this thesis, we have explored the following research question
How to *represent*, *predict*, and *control* stochastic systems in data-driven fashion?

In Chapter 4, we introduced the stochastic behavior of dynamical systems as all stochastic processes that are compatible with the dynamics encoded in the deterministic behavior. Moreover, the probability space on which the stochastic behavior is defined represents the uncertain environment to which the systems are exposed. Furthermore, we showed that the stochastic behavior of LTI systems is linear, time-invariant, and complete. By applying PCE and moments evaluation to stochastic processes, we further obtained the other two variants of behavioral characterizations for stochastic LTI systems, i.e., the behavior of PCE coefficients and the behavior of mean-covariance moment pairs. We demonstrated that these behaviors are all connected to the deterministic behavior of the realization system. Consequently, in the stochastic fundamental lemma, we showed that these behaviors of stochastic LTI systems can all be represented by selecting columns from the same Hankel matrices, constructed from realization data, but with different column selectors.

In Chapter 5, we investigated the stochastic behavior of dynamical systems influenced by unstructured (unmeasured or statistically unmodeled) disturbances by consolidating them into a residual disturbance that can be estimated and partially statistically modeled. We then proposed a data-driven stochastic prediction method. A simulation study using real residential building data shows that while prediction accuracy can be improved by identifying additional disturbances through costly measurements and forecasts, the proposed approach—based on the estimation of the residual disturbance—already provides reasonable predictions.

In Chapters 6–7, we addressed optimal control problems for stochastic systems with structured or unstructured disturbances. For structured (measured and modeled) disturbances, we formulated data-driven stochastic optimal control problems and their distributionally robust extensions. We provided sufficient conditions for the optimality of the proposed affine disturbance-feedback policies. For systems with unstructured disturbances, we leveraged the residual disturbance as in Chapter 5. Then, we developed an output-feedback predictive control scheme with closed-loop guarantees that relies only on the first two moments of the residual disturbance.

Finally, this thesis paves the way for several promising avenues for future research, as outlined below.

Towards Fast Computation of Data-Driven Stochastic Optimal Control

The data-driven stochastic optimal control problems considered in this thesis, such as OCP (6.9), can be formulated as distributed optimization problems. This interpretation arises from the linearity of polynomial chaos expansion, which ensures that the equality constraints induced by system dynamics are imposed independently for each PCE order, and the quadratic objective functional is decomposed into the sum of performance evaluations for each order. Thus, (6.9) can be reformulated as follows

$$\min_{\mathsf{x}^{[0,L-1]}} \sum_{j=0}^{L-1} f^j(\mathsf{x}^j)$$

$$\text{subject to } A\mathsf{x}^j = b^j,\ \mathsf{x}^{[0,L-1]} \in \mathrm{SOC},$$

where x^j represents all decision variables for the j-th order PCE coefficients, $f^j(\mathsf{x}^j)$ captures their corresponding contribution to the performance, $A\mathsf{x}^j = b^j$ summarizes the equalities in (6.9b)–(6.9c), and SOC represents all second order cone constraints by (6.9e)–(6.9f).

This reformulation highlights that the problem can be decomposed into L subproblems, each corresponding to a specific PCE order, coupled by second-order cone constraints. We can tailor the Alternating Direction Method of Multipliers (ADMM) (Boyd et al., 2011) to efficiently solve the problem, especially addressing the coupling through second-order cone constraints. We remark that this approach effectively leverages the problem's distributed structure to achieve computational efficiency, and holds promising potential for real-world applications.

Towards Data-driven Propagation of Epistemic Uncertainty

In this thesis, we focus on *aleatoric uncertainty*, which refers to exogenous disturbances affecting the system. This type of uncertainty is associated with the inherent variability and unpredictability of the environment. On the other hand, we also acknowledge the presence of *epistemic uncertainty*, which arises from the lack of knowledge or imprecise understanding of the system itself; see Hüllermeier et al., 2021 for the discussion of aleatoric and epistemic uncertainties.

In model-based approaches, *epistemic uncertainty* is typically represented as parametric uncertainty stemmed from system identification, for example:

$$\begin{aligned} X(k+1) &= A(\Theta)X(k) + B(\Theta)U(k) \\ Y(k) &= C(\Theta)X(k) + D(\Theta)U(k), \end{aligned}$$

where $\Theta \in L^2(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_\Theta})$ models the *epistemic* uncertainty surrounding system matrices. The propagation of such uncertainty using PCE has been studied extensively, as seen in Fisher et al., 2008; Wan et al., 2023, 2021.

In contrast, in data-driven approaches, we focus on non-parametric models derived from collected data. Here, the *epistemic uncertainty* stems from the estimation of past realizations $v^{\mathrm{d}}_{[0,T-1]}$ of the unstructured disturbances V. Furthermore, in Chapter 5, we hedge against the estimation error appeared during the estimation of past disturbance $\hat{v}^{\mathrm{d}}_{[0,T-1]}$ by using a regularized stochastic predictor in Chapter 7.

However, since the estimation itself is uncertain, we can treat the estimated values as random variables. Leveraging the consistency of data (Proposition (5.2)), the uncertain disturbance realization can be represented as

$$\mathcal{H}_1(V^{\mathrm{d}}_{[0,T-1]}) = \mathcal{H}_1\left(y^{\mathrm{d}}_{[0,T-1]}\right)\left(I_T - M^\dagger M\right) + \Theta M, \quad M = \begin{bmatrix} \mathcal{H}_1\left(z^{\mathrm{d}}_{[0,T-1]}\right) \\ \mathcal{H}_1\left(u^{\mathrm{d}}_{[0,T-1]}\right) \end{bmatrix}$$

where $\Theta \in L^2\left(\Omega, \mathcal{F}, \mu; \mathbb{R}^{n_y \times (\ell n_y + (\ell+1)n_u)}\right)$ represents the *epistemic uncertainty* in the data-driven setting. Moreover, we can include the stochastic past disturbance into Hankel

matrix to propagate epistemic uncertainty through dynamics

$$\begin{bmatrix} \mathcal{H}_N(y^{\mathrm{d}}_{[0,T-1]}) \\ \mathcal{H}_N(u^{\mathrm{d}}_{[0,T-1]}) \\ \mathcal{H}_N\left(V^{\mathrm{d}}_{[0,T-1]}(\Theta)\right) \end{bmatrix} G = \begin{bmatrix} Y \\ U \\ 0 \end{bmatrix}. \tag{8.1}$$

As future work, a comparison between the data-driven formulation and its model-based counterpart could be considered to assess their respective strengths and effectiveness in handling uncertainty.

Towards a Data-Driven Representation for Nonlinear Stochastic Systems

PCE has been applied to nonlinear MPC with parametric uncertainty, as shown by Fagiano et al., 2012 and Mesbah et al., 2014; see a recent review by Mishra et al., 2024. However, when it comes to additive disturbances, the problem remains open. Furthermore, the development of data-driven approaches for nonlinear stochastic systems is also an unsolved challenge.

A promising direction for data-driven representation of deterministic nonlinear systems is to mimic kernel regression methods, as explored by Huang et al., 2024 and Molodchyk et al., 2024. To extend this framework to stochastic nonlinear systems, a potential approach is to combine kernel regression with collocation methods, as proposed by Xiu, 2007 and Lefebvre et al., 2017. Future work could focus on enabling efficient and accurate data-driven representations of nonlinear stochastic systems while addressing both theoretical and computational challenges.

Bibliography

Ahbe, E., A. Iannelli, and R. S. Smith (2020). "Region of attraction analysis of nonlinear stochastic systems using polynomial chaos expansion". In: *Automatica* 122, p. 109187 (cit. on p. 5).

Antsaklis, P. J. and A. N. Michel (2006). *Linear Systems*. Birkhäuser, Boston (cit. on p. 18).

Aolaritei, L., M. Fochesato, J. Lygeros, and F. Dörfler (2023a). "Wasserstein tube MPC with exact uncertainty propagation". In: *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 2036–2041. DOI: 10.1109/CDC49753.2023.10383526 (cit. on pp. 4, 93).

Aolaritei, L., N. Lanzetti, and F. Dörfler (2023b). "Capture, propagate, and control distributional uncertainty". In: *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 3081–3086. DOI: 10.1109/CDC49753.2023.10383424 (cit. on p. 4).

ApS, M. (2023). *The Mosek optimization toolbox for Julia manual Version 10.1.17*. URL: https://docs.mosek.com/latest/juliaapi/index.html (cit. on p. 120).

Baggio, G. and R. Sepulchre (2017). "LTI stochastic processes: a behavioral perspective". In: *IFAC-PapersOnLine* 50.1, pp. 2806–2811. DOI: 10.1016/j.ifacol.2017.08.631 (cit. on pp. 3–4, 10, 67–68).

Balim, H., A. Carron, M. N. Zeilinger, and J. Köhler (2024). "Stochastic data-driven predictive control: chance-constraint satisfaction with identified multi-step predictors". In: *IEEE Control Systems Letters* 8, pp. 3249–3254. DOI: 10.1109/LCSYS.2024.3523238 (cit. on p. 7).

Berberich, J. and F. Allgöwer (May 2025). "An overview of systems-theoretic guarantees in data-driven model predictive control". In: *Annual Review of Control, Robotics, and Autonomous Systems* 8.1, pp. 77–100. ISSN: 2573-5144. DOI: 10.1146/annurev-control-030323-024328 (cit. on p. 7).

Berberich, J., J. Köhler, M. A. Müller, and F. Allgöwer (2020). "Data-driven model predictive control with stability and robustness guarantees". In: *IEEE Transactions on Automatic Control* 66.4, pp. 1702–1717 (cit. on pp. 3, 7, 71, 122).

– (2021). "On the design of terminal ingredients for data-driven MPC". In: *IFAC-PapersOnLine* 54.6, pp. 257–263 (cit. on pp. 7, 105, 111).

Bertsimas, D. and I. Popescu (2005). "Optimal inequalities in probability theory: A convex optimization approach". In: *SIAM Journal on Optimization* 15.3, pp. 780–804 (cit. on p. 45).

Bilgic, D., A. Harding, and T. Faulwasser (2024). "Data-driven predictive control of bilinear HVAC dynamics–An experimental case study". In: *IEEE Control Systems Letters* 8, pp. 3009–3014. DOI: 10.1109/LCSYS.2024.3519224 (cit. on p. 2).

Bilgic, D., A. Koch, G. Pan, and T. Faulwasser (2022). "Toward data-driven predictive control of multi-energy distribution systems". In: *Electric Power Systems Research* 212, p. 108311. DOI: 10.1016/j.epsr.2022.108311 (cit. on pp. 2–3).

Bock, H. and K. Plitt (1984). "A multiple shooting algorithm for direct solution of optimal control problems". In: *IFAC Proceedings Volumes* 17.2, pp. 1603–1608 (cit. on p. 90).

Bongard, J., J. Berberich, J. Köhler, and F. Allgöwer (2023). "Robust stability analysis of a simple data-driven model predictive control approach". In: *IEEE Transactions on Automatic Control* 68.5, pp. 2625–2637. DOI: 10.1109/TAC.2022.3163110 (cit. on pp. 7, 25).

Boyd, S., N. Parikh, E. Chu, B. Peleato, and J. Eckstein (2011). "Distributed optimization and statistical learning via the alternating direction method of multipliers". In: *Foundations and Trends® in Machine Learning* 3.1, pp. 1–122. DOI: 10.1561/2200000016 (cit. on p. 124).

Buot, M. (2006). *Probability and computing: randomized algorithms and probabilistic analysis* (cit. on p. 43).

Calafiore, G. C. and L. E. Ghaoui (2006). "On distributionally robust chance-constrained linear programs". In: *Journal of Optimization Theory and Applications* 130.1, pp. 1–22 (cit. on p. 47).

Cameron, R. H. and W. T. Martin (Apr. 1947). "The orthogonal development of non-linear functionals in series of Fourier-Hermite functionals". In: *Annals of Mathematics* 48.2, p. 385. ISSN: 0003-486X. DOI: 10.2307/1969178 (cit. on p. 35).

Cannon, M., B. Kouvaritakis, and X. Wu (2009). "Model predictive control for systems with stochastic multiplicative uncertainty and probabilistic constraints". In: *Automatica* 45.1, pp. 167–172 (cit. on p. 118).

Cantelli, F. P. (1928). "Sui confini della probabilità," in: *Atti del Congresso Internazionale dei Matematici* 6.47-59 (cit. on p. 41).

Carlet, P. G., A. Favato, R. Torchio, F. Toso, S. Bolognani, and F. Dörfler (2023). "Real-time feasibility of data-driven predictive control for synchronous motor drives". In: *IEEE Transactions on Power Electronics* 38.2, pp. 1672–1682. DOI: 10.1109/TPEL.2022.3214760 (cit. on p. 3).

Carvalho, A., Y. Gao, S. Lefevre, and F. Borrelli (2014). "Stochastic predictive control of autonomous vehicles in uncertain environments". In: *12th International Symposium on Advanced Vehicle Control*. Vol. 9 (cit. on p. 1).

Chebyshev, P. L. (1867). "Des valeurs moyennes". In: *Journal de Mathématiques Pures et Appliquées* 2.12, pp. 177–184 (cit. on p. 41).

Chen, P., P. Siano, B. Bak-Jensen, and Z. Chen (2010). "Stochastic optimization of wind turbine power factor using stochastic model of wind power". In: *IEEE Transactions on Sustainable Energy* 1.1, pp. 19–29 (cit. on p. 1).

Coppens, P. and P. Patrinos (2021). "Data-driven distributionally robust MPC for constrained stochastic systems". In: *IEEE Control Systems Letters* 6, pp. 1274–1279 (cit. on p. 93).

Coulson, J., J. Lygeros, and F. Dörfler (2019). "Data-enabled predictive control: In the shallows of the DeePC". In: *2019 18th European Control Conference (ECC)*. IEEE, pp. 307–312 (cit. on pp. 3, 7, 71, 90, 122).

– (2022). "Distributionally robust chance constrained data-enabled predictive control". In: *IEEE Transactions on Automatic Control* 67.7, pp. 3289–3304. DOI: 10.1109/TAC.2021.3097706 (cit. on p. 93).

Coulson, J., H. J. van Waarde, J. Lygeros, and F. Dörfler (2023). "A quantitative notion of persistency of excitation and the robust fundamental lemma". In: *IEEE Control Systems Letters* 7, pp. 1243–1248. DOI: 10.1109/LCSYS.2022.3232303 (cit. on p. 6).

De Persis, C. and P. Tesi (2019). "Formulas for data-driven control: Stabilization, optimality, and robustness". In: *IEEE Transactions on Automatic Control* 65.3, pp. 909–924. DOI: 10.1109/tac.2019.2959924 (cit. on pp. 2, 105, 111, 118).

Di Natale, L., Y. Lian, E. T. Maddalena, J. Shi, and C. N. Jones (2022). "Lessons learned from data-driven building control experiments: Contrasting gaussian process-based MPC, bilevel DeePC, and deep reinforcement learning". In: *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 1111–1117. DOI: 10.1109/CDC51059.2022.9992445 (cit. on p. 2).

Dörfler, F., J. Coulson, and I. Markovsky (2023a). "Bridging direct and indirect data-driven control formulations via regularizations and relaxations". In: *IEEE Transactions on Automatic Control* 68.2, pp. 883–897. DOI: 10.1109/TAC.2022.3148374 (cit. on p. 1).

Dörfler, F., P. Tesi, and C. De Persis (2023b). "On the certainty-equivalence approach to direct data-driven LQR design". In: *IEEE Transactions on Automatic Control* 68.12, pp. 7989–7996. DOI: 10.1109/TAC.2023.3253787 (cit. on p. 112).

Drgoňa, J., J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, and L. Helsen (2020). "All you need to know about model predictive control for buildings". In: *Annual Reviews in Control* 50, pp. 190–232 (cit. on p. 1).

El Ghaoui, L. and H. Lebret (1997). "Robust solutions to least-squares problems with uncertain data". In: *SIAM Journal on Matrix Analysis and Applications* 18.4, pp. 1035–1064 (cit. on p. 76).

Fagiano, L. and M. Khammash (2012). "Nonlinear stochastic model predictive control via regularized polynomial chaos expansions". In: *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*. IEEE, pp. 142–147 (cit. on p. 125).

Fahmy, M. M., A. A. R. Hanafy, and M. F. Sakr (1980). "On the discrete-time optimal regular control problem". In: *Information and Control* 44.3, pp. 223–235. DOI: 10.1016/S0019-9958(80)90148-5 (cit. on pp. 26, 69).

Farina, M., L. Giulioni, L. Magni, and R. Scattolini (2013). "A probabilistic approach to model predictive control". In: *2013 52nd IEEE Conference on Decision and Control (CDC)*. IEEE, pp. 7734–7739. DOI: 10.1109/cdc.2013.6761117 (cit. on pp. 8, 87, 108, 113, 115).

Farina, M., L. Giulioni, L. Magni, and R. Scattolini (2015). "An approach to output-feedback MPC of stochastic linear discrete-time systems". In: *Automatica* 55, pp. 140–149. DOI: 10.1016/j.automatica.2015.02.039 (cit. on p. 8).

Farina, M. and R. Scattolini (2016). "Model predictive control of linear systems with multiplicative unbounded uncertainty and chance constraints". In: *Automatica* 70, pp. 258–265 (cit. on pp. 8, 117).

Faulwasser, T., R. Ou, G. Pan, P. Schmitz, and K. Worthmann (2023). "Behavioral theory for stochastic systems? A data-driven journey from Willems to Wiener and back again". In: *Annual Reviews in Control* 55, pp. 92–117. DOI: 10.1016/j.arcontrol.2023.03.005 (cit. on pp. 3, 6–7, 9, 49–50, 58–60, 63).

Favoreel, W., B. De Moor, and M. Gevers (1999). "SPC: Subspace predictive control". In: *IFAC Proceedings Volumes* 32.2, pp. 4004–4009. DOI: 10.1016/s1474-6670(17)56683-5 (cit. on pp. 3, 29).

Ferizbegovic, M., H. Hjalmarsson, P. Mattsson, and T. B. Schön (2021). "Willems' fundamental lemma based on second-order moments". In: *2021 60th IEEE Conference on Decision and Control (CDC)*, pp. 396–401. DOI: 10.1109/CDC45484.2021.9683632 (cit. on p. 6).

Fiedler, F. and S. Lucia (2021). "On the relationship between data-enabled predictive control and subspace predictive control". In: *2021 European Control Conference (ECC)*, pp. 222–229. DOI: 10.23919/ECC54610.2021.9654975 (cit. on pp. 3, 29).

Field, R. V. and M. Grigoriu (2004). "On the accuracy of the polynomial chaos approximation". In: *Probabilistic Engineering Mechanics* 19.1-2, pp. 65–80 (cit. on p. 35).

Fisher, J. and R. Bhattacharya (2008). "On stochastic LQR design and polynomial chaos". In: *Proc. 2008 American Control Conference (ACC)*. IEEE, pp. 95–100 (cit. on p. 124).

Fochesato, M. and J. Lygeros (2022). "Data-driven distributionally robust bounds for stochastic model predictive control". In: *2022 IEEE 61st Conference on Decision and Control*. IEEE, pp. 3611–3616 (cit. on p. 93).

Fristedt, B. E. and L. F. Gray (2013). *A Modern Approach to Probability Theory*. Springer Science & Business Media (cit. on pp. 31–34, 40).

Gelbrich, M. (1990). "On a formula for the L2 Wasserstein metric between measures on Euclidean and Hilbert spaces". In: *Mathematische Nachrichten* 147.1, pp. 185–203 (cit. on p. 94).

Ghanem, R. G. and P. D. Spanos (2003). *Stochastic Finite Elements: A Spectral Approach*. Revised. Springer New York (cit. on p. 38).

Givens, C. R. and R. M. Shortt (1984). "A class of Wasserstein metrics for probability distributions." In: *Michigan Mathematical Journal* 31.2, pp. 231–240 (cit. on pp. 93–94).

Goulart, P. J., E. C. Kerrigan, and J. M. Maciejowski (2006). "Optimization over state feedback policies for robust control with constraints". In: *Automatica* 42.4, pp. 523–533. DOI: 10.1016/j.automatica.2005.08.023 (cit. on p. 87).

He, S., J. Zhang, and S. Zhang (2010). "Bounding probability of small deviation: A fourth moment approach". In: *Mathematics of Operations Research* 35.1, pp. 208–232 (cit. on p. 45).

Hespanha, J. P. (2023). *Linear Systems Theory*. Princeton University Press. DOI: 10.23943/9781400890088 (cit. on p. 23).

Hewing, L. and M. N. Zeilinger (2018). "Stochastic model predictive control for linear systems using probabilistic reachable sets". In: *2018 IEEE CDC*, pp. 5182–5188 (cit. on p. 113).

Hewing, L., J. Kabzan, and M. N. Zeilinger (2020a). "Cautious model predictive control using gaussian process regression". In: *IEEE Transactions on Control Systems Technology* 28.6, pp. 2736–2743. DOI: 10.1109/TCST.2019.2949757 (cit. on p. 71).

Hewing, L., K. P. Wabersich, and M. N. Zeilinger (2020b). "Recursively feasible stochastic model predictive control using indirect feedback". In: *Automatica* 119, p. 109095 (cit. on pp. 107, 118).

Hiremath, S. A., V. K. Mishra, and N. Bajcinca (2022). "Learning based stochastic data-driven predictive control". In: *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 1684–1961. DOI: 10.1109/CDC51059.2022.9993198 (cit. on p. 6).

Hochreiter, S. (1997). "Long short-term memory". In: *Neural Computation MIT-Press* (cit. on p. 75).

Huang, L., J. Lygeros, and F. Dörfler (2024). "Robust and kernelized data-enabled predictive control for nonlinear systems". In: *IEEE Transactions on Control Systems Technology* 32.2, pp. 611–624. DOI: 10.1109/TCST.2023.3329334 (cit. on pp. 16, 125).

Huang, L., J. Zhen, J. Lygeros, and F. Dörfler (2021). "Quadratic regularization of data-enabled predictive control: Theory and application to power converter experiments". In: *IFAC-PapersOnLine* 54.7, pp. 192–197. DOI: 10.1016/j.ifacol.2021.08.357 (cit. on pp. 7, 76–77, 90).

– (2023). "Robust data-enabled predictive control: Tractable formulations and performance guarantees". In: *IEEE Transactions on Automatic Control* 68.5, pp. 3163–3170 (cit. on p. 7).

Hüllermeier, E. and W. Waegeman (2021). "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods". In: *Machine Learning* 110.3, pp. 457–506 (cit. on p. 124).

Katayama, T. (2005). *Subspace Methods for System Identification*. Vol. 1. Springer London. ISBN: 9781846281587. DOI: 10.1007/1-84628-158-x (cit. on p. 2).

Kaupmann, M. (2023). "Data-driven modeling for building control". MA thesis. Institute of Energy Systems, Energy Efficiency and Energy Economics, TU Dortmund University (cit. on p. 78).

Kerz, S., J. Teutsch, T. Brüdigam, M. Leibold, and D. Wollherr (2023). "Data-driven tube-based stochastic predictive control". In: *IEEE Open Journal of Control Systems* 2, pp. 185–199. DOI: 10.1109/OJCSYS.2023.3291596 (cit. on p. 7).

Kim, K. K., D. E. Shen, Z. K. Nagy, and R. D. Braatz (2013). "Wiener's polynomial chaos for the analysis and control of nonlinear dynamical systems with probabilistic uncertainties [Historical Perspectives]". In: *IEEE Control Systems Magazine* 33.5, pp. 58–67. DOI: 10.1109/MCS.2013.2270410 (cit. on p. 5).

Köhler, J., K. P. Wabersich, J. Berberich, and M. N. Zeilinger (2022a). "State space models vs. multi-step predictors in predictive control: Are state space models complicating safe data-driven designs?" In: *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 491–498. DOI: 10.1109/CDC51059.2022.9992373 (cit. on p. 7).

Köhler, J. and M. N. Zeilinger (2022b). "Recursively feasible stochastic predictive control using an interpolating initial state constraint". In: *IEEE Control Systems Letters* 6, pp. 2743–2748. ISSN: 2475-1456. DOI: 10.1109/lcsys.2022.3176405 (cit. on pp. 8, 115).

Lazar, M. and P. C. N. Verheijen (2022). "Offset–free data–driven predictive control". In: *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 1099–1104. DOI: 10.1109/CDC51059.2022.9992453 (cit. on p. 118).

Lefebvre, T. (2020). "On moment estimation from polynomial chaos expansion models". In: *IEEE Control Systems Letters* 5.5, pp. 1519–1524 (cit. on pp. 6, 36).

Lefebvre, T., F. De Belie, and G. Crevecoeur (2017). "Polynomial chaos explicit solution of the optimal control problem in model predictive control". In: *2017 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*. IEEE, pp. 1762–1767 (cit. on p. 125).

Lian, Y., J. Shi, M. Koch, and C. N. Jones (Nov. 2023). "Adaptive robust data-driven building control via bilevel reformulation: An experimental result". In: *IEEE Transactions on Control Systems Technology* 31.6, pp. 2420–2436. ISSN: 2374-0159. DOI: 10.1109/tcst.2023.3259641 (cit. on pp. 2–3).

Lobo, M. S., L. Vandenberghe, S. Boyd, and H. Lebret (1998). "Applications of second-order cone programming". In: *Linear Algebra and Its Applications* 284.1-3, pp. 193–228 (cit. on p. 89).

Lopez, V. G. and M. A. Müller (2022). "On a continuous-time version of Willems' lemma". In: *2022 IEEE 61st conference on decision and control (CDC)*. IEEE, pp. 2759–2764. DOI: 10.1109/cdc51059.2022.9992347 (cit. on p. 2).

Lorenzen, M., F. Dabbene, R. Tempo, and F. Allgöwer (2016). "Constraint-tightening and stability in stochastic model predictive control". In: *IEEE Transactions on Automatic Control* 62.7, pp. 3165–3177 (cit. on p. 8).

Lorenzen, M., F. Dabbene, R. Tempo, and F. Allgöwer (2017). "Stochastic MPC with offline uncertainty sampling". In: *Automatica* 81, pp. 176–183 (cit. on p. 8).

Lu, S., J. H. Lee, and F. You (2020). "Soft-constrained model predictive control based on data-driven distributionally robust optimization". In: *AIChE Journal* 66.10, e16546 (cit. on p. 93).

Lucia, S., J. A. Paulson, R. Findeisen, and R. D. Braatz (2017). "On stability of stochastic linear systems via polynomial chaos expansions". In: *2017 American Control Conference (ACC)*. IEEE, pp. 5089–5094 (cit. on p. 5).

Maciejowski, J. M. (2002). *Predictive Control with Constraints*. Pearson Education (cit. on p. 119).

Markovsky, I. and F. Dörfler (2021). "Behavioral systems theory in data-driven analysis, signal processing, and control". In: *Annual Reviews in Control* 52, pp. 42–64. ISSN: 1367-5788. DOI: 10.1016/j.arcontrol.2021.09.005 (cit. on pp. 2–3, 13).

Markovsky, I. and P. Rapisarda (2008). "Data-driven simulation and control". In: *International Journal of Control* 81.12, pp. 1946–1959. DOI: 10.1080/00207170801942170 (cit. on p. 29).

Meijer, T. J., S. A. N. Nouwens, K. J. A. Scheres, V. S. Dolk, and W. P. M. H. Heemels (2024). "Frequency-domain data-driven predictive control". In: *IFAC-PapersOnLine* 58.18, pp. 86–91 (cit. on p. 2).

Mesbah, A., S. Streif, R. Findeisen, and R. D. Braatz (2014). "Stochastic nonlinear model predictive control with probabilistic constraints". In: *2014 American Control Conference (ACC)*. IEEE, pp. 2413–2419 (cit. on p. 125).

Mesbah, A. (2016). "Stochastic model predictive control: An overview and perspectives for future research". In: *IEEE Control Systems Magazine* 36.6, pp. 30–44 (cit. on pp. 8, 87).

Mishra, P. K., J. A. Paulson, and R. D. Braatz (2024). "Polynomial chaos-based stochastic model predictive control: an overview and future research directions". In: *arXiv:2406.10734* (cit. on p. 125).

Molodchyk, O. and T. Faulwasser (2024). "Exploring the links between the fundamental lemma and kernel regression". In: *IEEE Control Systems Letters* 8, pp. 2045–2050. DOI: 10.1109/LCSYS.2024.3406053 (cit. on pp. 16, 125).

Mühlpfordt, T., R. Findeisen, V. Hagenmeyer, and T. Faulwasser (2017). "Comments on quantifying truncation errors for polynomial chaos expansions". In: *IEEE Control Systems Letters* 2.1, pp. 169–174. DOI: 10.1109/LCSYS.2017.2778138 (cit. on pp. 5, 35).

Mühlpfordt, T. (2020). "Uncertainty quantification via polynomial chaos expansion – methods and applications for optimization of power systems". PhD thesis. Karlsruher Institut für Technologie (KIT) (cit. on pp. 5, 38).

Mühlpfordt, T., L. Roald, V. Hagenmeyer, T. Faulwasser, and S. Misra (2019). "Chance-constrained AC optimal power flow: A polynomial chaos approach". In: *IEEE Transactions on Power Systems* 34.6, pp. 4806–4816 (cit. on p. 1).

Nguyen, V. A., S. Shafiee, D. Filipović, and D. Kuhn (2021). "Mean-covariance robust risk measurement". In: *arXiv:2112.09959* (cit. on pp. 10, 46, 94–96).

Nguyen, V. A., S. Shafieezadeh-Abadeh, D. Kuhn, and P. Mohajerin Esfahani (2023). "Bridging Bayesian and minimax mean square error estimation via Wasserstein distributionally robust optimization". In: *Mathematics of Operations Research* 48.1, pp. 1–37 (cit. on p. 93).

Nortmann, B. and T. Mylvaganam (2023). "Direct data-driven control of linear time-varying systems". In: *IEEE Transactions on Automatic Control* 68.8, pp. 4888–4895. DOI: 10.1109/TAC.2023.3276909 (cit. on p. 2).

O'Dwyer, E., E. C. Kerrigan, P. Falugi, M. Zagorowska, and N. Shah (2023). "Data-driven predictive control with improved performance using segmented trajectories". In: *IEEE Transactions on Control Systems Technology* 31.3, pp. 1355–1365. DOI: 10.1109/TCST.2022.3224330 (cit. on p. 89).

Oladyshkin, S. and W. Nowak (2012). "Data-driven uncertainty quantification using the arbitrary polynomial chaos expansion". In: *Reliability Engineering & System Safety* 106, pp. 179–190 (cit. on p. 38).

Ou, R., M. H. Baumann, L. Grüne, and T. Faulwasser (2021). "A simulation study on turnpikes in stochastic LQ optimal control". In: *IFAC-PapersOnLine* 54.3. 16th IFAC Symposium on Advanced Control of Chemical Processes ADCHEM 2021, pp. 516–521 (cit. on p. 5).

Ou, R., G. Pan, and T. Faulwasser (2023). "Data-driven multiple shooting for stochastic optimal control". In: *IEEE Control Systems Letters* 7, pp. 313–318. ISSN: 2475-1456. DOI: 10.1109/lcsys.2022.3185841 (cit. on p. 90).

Ou, R., J. Schießl, M. H. Baumann, L. Grüne, and T. Faulwasser (Apr. 2025). "A polynomial chaos approach to stochastic LQ optimal control: error bounds and infinite-horizon results". In: *Automatica* 174, p. 112117. ISSN: 0005-1098. DOI: 10.1016/j.automatica.2025.112117 (cit. on p. 5).

Özmeteler, M. B., D. Bilgic, G. Pan, A. Koch, and T. Faulwasser (Nov. 2024). "Data-driven uncertainty propagation for stochastic predictive control of multi-energy systems". In: *European Journal of Control* 80, p. 101066. ISSN: 0947-3580. DOI: 10.1016/j.ejcon.2024.101066 (cit. on pp. 1, 49, 90).

Pan, G. and T. Faulwasser (2023a). "Distributionally robust uncertainty quantification via data-driven stochastic optimal control". In: *IEEE Control Systems Letters* 7, pp. 3036–3041. DOI: 10.1109/LCSYS.2023.3290362 (cit. on pp. 10, 85, 95, 97, 103).

Pan, G., R. Ou, and T. Faulwasser (2023b). "Data-driven stochastic output-feedback predictive control: Recursive feasibility through interpolated initial conditions". In: *Proceedings of The 5th Annual Learning for Dynamics and Control Conference*. Ed. by N. Matni, M. Morari, and G. J. Pappas. Vol. 211. Proceedings of Machine Learning Research. PMLR, pp. 980–992 (cit. on pp. 8, 10, 105, 115, 122).

– (Oct. 2023c). "On a stochastic fundamental lemma and its use for data-driven optimal control". In: *IEEE Transactions on Automatic Control* 68.10, pp. 5922–5937. ISSN: 2334-3303. DOI: 10.1109/tac.2022.3232442 (cit. on pp. 3, 6, 9–10, 49–50, 62–64, 85, 103, 119).

– (May 2025a). "On data-driven stochastic output-feedback predictive control". In: *IEEE Transactions on Automatic Control* 70.5, pp. 2948–2962. ISSN: 2334-3303. DOI: 10.1109/tac.2024.3494394 (cit. on pp. 8, 10, 72, 105, 113, 117, 119, 122).

– (Aug. 2025b). "Towards data-driven stochastic predictive control". In: *International Journal of Robust and Nonlinear Control* 35.7, pp. 2588–2610. ISSN: 1099-1239. DOI: 10.1002/rnc.6812 (cit. on pp. 8, 10, 105, 108, 113–114, 122).

Paulson, J. A., E. A. Buehler, and A. Mesbah (2017). "Arbitrary polynomial chaos for uncertainty propagation of correlated random variables in dynamic systems". In: *IFAC-PapersOnLine* 50.1, pp. 3548–3553. DOI: 10.1016/j.ifacol.2017.08.954 (cit. on p. 38).

Penrose, R. (1955). "A generalized inverse for matrices". In: *Mathematical Proceedings of the Cambridge Philosophical Society* 51.3, pp. 406–413. DOI: 10.1017/S0305004100030401 (cit. on p. 37).

Pola, G., C. Manes, and M. D. Di Benedetto (2016). "On external behavior equivalence of continuous-time stochastic linear control systems". In: *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 6583–6588. DOI: 10.1109/CDC.2016.7799282 (cit. on p. 3).

Pola, G., C. Manes, A. J. van der Schaft, and M. D. Di Benedetto (2015). "On equivalence notions for discrete-time stochastic control systems". In: *2015 54th IEEE Conference on Decision and Control (CDC)*, pp. 1180–1185. DOI: 10.1109/CDC.2015.7402371 (cit. on p. 3).

Rapisarda, P., H. J. van Waarde, and M. K. Çamlibel (2024). "Orthogonal polynomial bases for data-driven analysis and control of continuous-time systems". In: *IEEE Transactions on Automatic Control* 69.7, pp. 4307–4319. DOI: 10.1109/TAC.2023.3321214 (cit. on p. 5).

Reinhardt, D., W. Cai, and S. Gros (2025). "Data-driven domestic flexible demand: observations from experiments in cold climate". In: *Energy Systems Integration for Multi-Energy Systems*. Springer Nature Switzerland, pp. 691–728. ISBN: 9783031690150. DOI: 10.1007/978-3-031-69015-0_25 (cit. on pp. 10, 78).

Rueda-Escobedo, J. G., E. Fridman, and J. Schiffer (2022). "Data-driven control for linear discrete-time delay systems". In: *IEEE Transactions on Automatic Control* 67.7, pp. 3321–3336. DOI: 10.1109/TAC.2021.3096896 (cit. on p. 2).

Sadamoto, T. (2023). "On equivalence of data informativity for identification and data-driven control of partially observable systems". In: *IEEE Transactions on Automatic Control* 68.7, pp. 4289–4296. DOI: 10.1109/TAC.2022.3202082 (cit. on pp. 25, 105).

Sathyanarayanan, K. K., G. Pan, and T. Faulwasser (2023). "Towards data-driven predictive control using wavelets". In: *IFAC-PapersOnLine* 56.2, pp. 632–637 (cit. on p. 2).

Schlüter, H. and F. Allgöwer (2023). "Stochastic model predictive control using initial state and variance interpolation". In: *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 6700–6706. DOI: 10.1109/CDC49753.2023.10383711 (cit. on pp. 8, 115).

– (2022). "Stochastic model predictive control using initial state optimization". In: *IFAC-PapersOnLine* 55.30. 25th International Symposium on Mathematical Theory of Networks and Systems MTNS 2022, pp. 454–459. DOI: 10.1016/j.ifacol.2022.11.095 (cit. on p. 115).

Schmitz, P., T. Faulwasser, P. Rapisarda, and K. Worthmann (2024). "A continuous-time fundamental lemma and its application in data-driven optimal control". In: *Systems & Control Letters* 194, p. 105950. DOI: 10.1016/j.sysconle.2024.105950 (cit. on p. 2).

Schmitz, P., T. Faulwasser, and K. Worthmann (2022). "Willems' fundamental lemma for linear descriptor systems and its use for data-driven output-feedback MPC". In: *IEEE Control Systems Letters* 6, pp. 2443–2448. DOI: 10.1109/LCSYS.2022.3161054 (cit. on p. 2).

Sedghizadeh, S. and S. Beheshti (2018). "Data-driven subspace predictive control: Stability and horizon tuning". In: *Journal of the Franklin Institute* 355.15, pp. 7509–7547 (cit. on pp. 3, 29).

Sullivan, T. J. (2015). *Introduction to Uncertainty Quantification*. Vol. 63. Springer (cit. on pp. 31–32, 35, 39).

Taşkesen, B., S. Shafieezadeh-Abadeh, and D. Kuhn (2023). "Semi-discrete optimal transport: Hardness, regularization and numerical solution". In: *Mathematical Programming* 199.1, pp. 1033–1106 (cit. on p. 94).

Teutsch, J., S. Kerz, D. Wollherr, and M. Leibold (2024). "Sampling-based stochastic data-driven predictive control under data uncertainty". In: *arXiv:2402.00681* (cit. on p. 8).

van Overschee, P. and B. De Moor (2012). *Subspace identification for linear systems: Theory-Implementation-Applications*. Springer Science & Business Media (cit. on pp. 2, 7, 28, 69).

van Waarde, H. J. (2021). "Beyond persistent excitation: Online experiment design for data-driven modeling and control". In: *IEEE Control Systems Letters*. DOI: 10.1109/lcsys.2021.3073860 (cit. on pp. 2, 28).

Verhaegen, M. and P. Dewilde (1992a). "Subspace model identification part 1. The output-error state-space model identification class of algorithms". In: *International Journal of Control* 56.5, pp. 1187–1210. DOI: 10.1080/00207179208934363 (cit. on p. 28).

– (1992b). "Subspace model identification part 2. Analysis of the elementary output-error state-space model identification algorithm". In: *International Journal of Control* 56.5, pp. 1211–1241. DOI: 10.1080/00207179208934364 (cit. on p. 28).

Verhoek, C., H. S. Abbas, R. Tóth, and S. Haesaert (2021). "Data-driven predictive control for linear parameter-varying systems". In: *IFAC-PapersOnLine* 54.8. 4th IFAC Workshop on Linear Parameter Varying Systems LPVS 2021, pp. 101–108 (cit. on p. 2).

Wächter, A. and L. T. Biegler (2006). "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming". In: *Mathematical Programming* 106.1, pp. 25–57 (cit. on p. 120).

Wan, Y., D. Shen, S. Lucia, R. Findeisen, and R. D. Braatz (2023). "A polynomial chaos approach to robust H_infinity static output-feedback control with bounded truncation error". In: *IEEE Transactions on Automatic Control* 68.1, pp. 470–477. DOI: 10.1109/TAC.2022.3140275 (cit. on p. 124).

Wan, Y., D. E. Shen, S. Lucia, R. Findeisen, and R. D. Braatz (2021). "Polynomial chaos-based H2 output-feedback control of systems with probabilistic parametric uncertainties". In: *Automatica* 131, p. 109743 (cit. on p. 124).

Wang, J., Y. Zheng, K. Li, and Q. Xu (2023). "DeeP-LCC: Data-enabled predictive leading cruise control in mixed traffic flow". In: *IEEE Transactions on Control Systems Technology*, pp. 1–17. DOI: 10.1109/TCST.2023.3288636 (cit. on p. 3).

Wang, Y., K. You, D. Huang, and C. Shang (2025). "Data-driven output prediction and control of stochastic systems: An innovation-based approach". In: *Automatica* 171, p. 111897. DOI: 10.1016/j.automatica.2024.111897 (cit. on pp. 7, 70).

Wiener, N. (1938). "The homogeneous chaos". In: *American Journal of Mathematics* 60.4, pp. 897–936 (cit. on pp. 5, 35).

Wiesemann, W., D. Kuhn, and M. Sim (2014). "Distributionally robust convex optimization". In: *Operations Research* 62.6, pp. 1358–1376 (cit. on p. 93).

Willems, J. C., P. Rapisarda, I. Markovsky, and B. L. M. De Moor (2005). "A note on persistency of excitation". In: *Systems & Control Letters* 54.4, pp. 325–329. DOI: 10.1016/j.sysconle.2004.09.003 (cit. on pp. 2, 13, 27–28, 111).

Willems, J. C. (1986). "From time series to linear system—Part I. Finite dimensional linear time invariant systems". In: *Automatica* 22.5, pp. 561–580. DOI: 10.1016/0005-1098(86)90066-x (cit. on pp. 2, 8, 13–16, 18–19, 21–23, 27).

– (2012). "Open stochastic systems". In: *IEEE Transactions on Automatic Control* 58.2, pp. 406–421. DOI: 10.1109/tac.2012.2210836 (cit. on pp. 3–4, 10, 67–68).

– (1991). "Paradigms and puzzles in the theory of dynamical systems". In: *IEEE Transactions on automatic control* 36.3, pp. 259–294. DOI: 10.1109/9.73561 (cit. on pp. 2, 8, 13, 15, 17–18, 20–23, 53).

Witteveen, J. A. S. and H. Bijl (2006). "Modeling arbitrary uncertainties using Gram-Schmidt polynomial chaos". In: *44th AIAA Aerospace Sciences Meeting and Exhibit*. Ed. by N. J. Pfeiffer. United States: American Institute of Aeronautics and Astronautics Inc. (AIAA), p. 896 (cit. on pp. 5, 38, 91).

Wu, L. (2022). "Equivalence of SS-based MPC and ARX-based MPC". In: *arXiv:2209.00107* (cit. on p. 26).

Xiu, D. and G. E. Karniadakis (2002). "The Wiener–Askey polynomial chaos for stochastic differential equations". In: *SIAM Journal on Scientific Computing* 24.2, pp. 619–644. DOI: 10.1137/S1064827501387826 (cit. on p. 5).

Xiu, D. (2007). "Efficient collocational approach for parametric uncertainty analysis". In: (cit. on p. 125).

Yang, H. and S. Li (2015). "A data-driven predictive controller design based on reduced hankel matrix". In: *2015 10th Asian Control Conference (ASCC)*. IEEE, pp. 1–7. DOI: 10.1109/ascc.2015.7244723 (cit. on p. 3).

Yin, M., A. Iannelli, and R. S. Smith (2021). "Maximum likelihood estimation in data-driven modeling and control". In: *IEEE Transactions on Automatic Control* 68.1, pp. 317–328. DOI: 10.1109/tac.2021.3137788 (cit. on p. 7).

– (2024). "Stochastic data-driven predictive control: Regularization, estimation, and constraint tightening". In: *IFAC-PapersOnLine* 58.15, pp. 79–84. DOI: 10.1016/j.ifacol.2024.08.508 (cit. on p. 7).

Yu, Y., Y. Li, D. Schuurmans, and C. Szepesvári (2009). "A general projection property for distribution families". In: *Advances in Neural Information Processing Systems* 22 (cit. on p. 46).

Yu, Y., S. Talebi, H. J. van Waarde, U. Topcu, M. Mesbahi, and B. Açıkmeşe (2021). "On controllability and persistency of excitation in data-driven control: Extensions of Willems' fundamental lemma". In: *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, pp. 6485–6490 (cit. on pp. 2, 28).

A. Proofs

A.1. Proof of Proposition 7.1

This proof is organized into two main steps. As a preparatory step, we give a technical lemma that demonstrates the recursive feasibility of OCP (7.13) using an extended PCE basis. Then, we turn to the proof of Proposition 7.1 by using projection $T(k)$ to keep the dimension of the basis constant.

Lemma A.1 (Recursive feasibility with extended basis): Let Assumptions 7.3–7.4 hold. At time instant $k-1$, given the basis $\{\phi^j(\boldsymbol{\xi}(k-1))\}_{j=0}^{L-1}$ as in (7.8), suppose that OCP (7.13) is feasible with the initial condition $\bar{\mathsf{z}}^{[0,L-1]}(k-1)$ where $\bar{\mathsf{z}}^j(k-1)=0, \forall j \in \mathbb{I}_{[1+n_\xi,L-1]}$.

At time instant k, considering the extended basis

$$\left\{\phi^j\left(\boldsymbol{\xi}(k-1), \eta(N-1\,|\,k)\right)\right\}_{j=0}^{L+n_\eta-1}, \text{with } \{\phi^j\}_{j=L}^{L+n_\eta-1} \doteq \mathcal{P}_1(\eta(N-1\,|\,k)), \tag{A.1}$$

OCP (7.13) is feasible with the initial condition

$$\bar{\mathsf{z}}^j(k) \doteq \begin{cases} \mathsf{z}^{j,\star}(1\,|\,k-1), & j \in \mathbb{I}_{[0,n_\xi+n_\eta]} \\ 0, & j \in \mathbb{I}_{[n_\xi+n_\eta+1,L+n_\eta-1]} \end{cases} \tag{A.2}$$

where $\mathsf{z}^{j,\star}(1\,|\,k-1)$ is the predicted optimal solution of $\mathsf{z}^j(k)$ at time instant $k-1$.

Proof. Consider the optimal solution at $k-1$, i.e. $(\mathsf{u},\mathsf{y},\mathsf{z})^{j,\star}(t\,|\,k-1)$ for $j \in \mathbb{I}_{[0,L-1]}$ and $t \in \mathbb{I}_{[0,N-1]}$. Similar to the proof of recursive feasibility in deterministic MPC, we construct a candidate feasible solution of OCP (7.13) at time instant k by shifting the last optimal solution based on the original basis (7.8). For all $j \in \mathbb{I}_{[0,L-1]}$, we have

$$\{\tilde{\mathsf{u}}^j(t\,|\,k)\}_{t=0}^{N-2} \doteq \{\mathsf{u}^{j,\star}(t\,|\,k-1)\}_{t=1}^{N-1}, \quad \tilde{\mathsf{u}}^j(N-1\,|\,k) \doteq K\mathsf{z}^{j,\star}(N\,|\,k-1) \tag{A.3a}$$

$$\{\tilde{\mathsf{y}}^j(t\,|\,k)\}_{t=0}^{N-2} \doteq \{\mathsf{y}^{j,\star}(t\,|\,k-1)\}_{t=1}^{N-1}, \quad \tilde{\mathsf{y}}^j(N-1\,|\,k) \doteq \widetilde{H}^\top \tilde{\mathsf{z}}^j(N\,|\,k) \tag{A.3b}$$

$$\{\tilde{\mathsf{z}}^j(t\,|\,k)\}_{t=0}^{N-1} \doteq \{\mathsf{z}^{j,\star}(t\,|\,k-1)\}_{t=1}^{N}, \quad \tilde{\mathsf{z}}^j(N\,|\,k) \doteq A_K\mathsf{z}^{j,\star}(N\,|\,k-1) \tag{A.3c}$$

and for all $j \in \mathbb{I}_{[L,L+n_\eta-1]}$ we have

$$\{\tilde{\mathsf{u}}^j(t\,|\,k)\}_{t=0}^{N-1} \doteq 0,\ \{\tilde{\mathsf{y}}^j(t\,|\,k)\}_{t=0}^{N-2} \doteq 0,\ \{\tilde{\mathsf{z}}^j(t\,|\,k)\}_{t=0}^{N-1} \doteq 0, \tag{A.3d}$$

$$\tilde{\mathsf{y}}^j(N-1\,|\,k) \doteq \mathsf{v}^j(N-1\,|\,k), \quad \tilde{\mathsf{z}}^j(N\,|\,k) \doteq \widetilde{H}\mathsf{v}^j(N-1\,|\,k). \tag{A.3e}$$

Note that the last disturbance involves new independent stochastic uncertainty with the corresponding input coefficients being set to 0 in (A.3d).

Next, we prove the feasibility of (A.3) in OCP (7.13) with respect to the extended basis (A.1). Thanks to the shift construction of (A.3) and Assumption 7.4, constraints (7.13b)–(7.13f) are satisfied. For the terminal constraints (7.13g), we see $\tilde{\mathsf{z}}^0(N\,|\,k) = A_K\mathsf{z}^{0,\star}(N\,|\,k-1) \in \mathbb{Z}_\mathrm{F}$ holds for $\mathsf{z}^{0,\star}(N\,|\,k-1) \in \mathbb{Z}_\mathrm{F}$, cf. Assumption 7.4. For the terminal constraint on the covariance, we have

$$\begin{aligned}
&\textstyle\sum_{j=1}^{L+n_\eta-1} \|\tilde{\mathsf{z}}^j(N\,|\,k)\|_\Gamma^2 \\
=&\sum_{j=1}^{L-1} \mathsf{z}^{j,\star\top}(N\,|\,k-1) \underbrace{A_K^\top \Gamma A_K}_{\overset{(7.7b)}{=}\Gamma-I_{n_z}} \mathsf{z}^{j,\star}(N\,|\,k-1) + \sum_{j=L}^{L+n_\eta-1} \mathsf{v}^{j\top}(N-1\,|\,k)\,\widetilde{H}^\top\Gamma\widetilde{H}\mathsf{v}^j(N-1\,|\,k) \\
\leq& \left(1-\tfrac{1}{\lambda_{\max}(\Gamma)}\right) \textstyle\sum_{j=1}^{L-1} \|\mathsf{z}^{j,\star}(N\,|\,k-1)\|_\Gamma^2 + \operatorname{trace}(\Sigma_V \widetilde{H}^\top\Gamma\widetilde{H}) \\
\leq& \left(1-\tfrac{1}{\lambda_{\max}(\Gamma)}\right)\tau + \operatorname{trace}(\Sigma_V \widetilde{H}^\top\Gamma\widetilde{H}) \overset{(7.7c)}{=} \tau.
\end{aligned}$$

Thus, we conclude that at time instant k, OCP (7.13) constructed with the extended basis $\{\phi^j\}_{j=0}^{L+n_\eta-1}$ is feasible when (A.2) is used as initial condition, and (A.3) is a feasible solution. □

Now, we give the proof of Proposition 7.1.

Proof. We construct a feasible solution candidate using the projection matrix $T(k)$. According to Corollary 5.1, for each $j \in \mathbb{I}_{[0,L+n_\eta-1]}$, there exists $\tilde{\mathsf{g}}^j$ corresponding to $\{(\tilde{\mathsf{u}},\tilde{\mathsf{y}},\tilde{\mathsf{z}})^j(t\,|\,k)\}_{t=0}^{N-1}$ as shown in (A.3) such that the Hankel equality constraint (7.13b) holds. Since the initial conditions are related by the projection $T(k)$ in (7.20), we consider

$$\bar{\mathsf{g}}^0 \doteq \tilde{\mathsf{g}}^0,\ \bar{\mathsf{g}}^j \doteq \tilde{\mathsf{g}}^{j+n_\eta}, \forall j \in \mathbb{I}_{[n_\xi+1,L-1]},\quad \mathcal{H}_1\left(\bar{\mathsf{g}}^{[1,n_\xi]}\right) \doteq \mathcal{H}_1\left(\tilde{\mathsf{g}}^{[1,n_\xi+n_\eta]}\right) T(k).$$

Next, we exploit Corollary 5.1 again. For each $j \in \mathbb{I}_{[0,L-1]}$, the resulting $\{(\bar{\mathsf{u}},\bar{\mathsf{y}},\bar{\mathsf{z}})^j(t\,|\,k)\}_{t=0}^{N-1}$ of evaluating (7.13b) with $\bar{\mathsf{g}}^j$ is an input-output trajectory of the PCE coefficient dynamics. For all $\mathsf{h} \in \{\mathsf{u},\mathsf{y},\mathsf{z}\}$ and all $t \in \mathbb{I}_{[0,N-1]}$, they satisfy

$$\begin{aligned}\bar{\mathsf{h}}^j(t\,|\,k) = \tilde{\mathsf{h}}^j(t\,|\,k),\ \bar{\mathsf{h}}^j(t\,|\,k) = \tilde{\mathsf{h}}^{j+n_\eta}(t\,|\,k), \forall j \in \mathbb{I}_{[n_\xi+1,L-1]},\\ \mathcal{H}_1\left(\bar{\mathsf{h}}^{[1,n_\xi]}(t\,|\,k)\right) = \mathcal{H}_1\left(\tilde{\mathsf{h}}^{[1,n_\xi+n_\eta]}(t\,|\,k)\right) T(k).\end{aligned} \tag{A.4}$$

In other words, the resulting $\{(\bar{\mathsf{u}},\bar{\mathsf{y}},\bar{\mathsf{z}})^j(t\,|\,k)\}_{t=0}^{N-1}$ satisfies (7.13b)–(7.13d). Turning to the chance constraints (7.13e)–(7.13f), we note that for all $i \in \mathbb{I}_{[1,N_u]}$, $t \in \mathbb{I}_{[0,N-1]}$, we have

$$\begin{aligned}&\|a_{u,i}^\top \mathcal{H}_1\left(\bar{\mathsf{u}}^{[1,L-1]}(t\,|\,k)\right)\|^2 - \|a_{u,i}^\top \mathcal{H}_1\left(\tilde{\mathsf{u}}^{[1,L+n_\eta-1]}(t\,|\,k)\right)\|^2\\ &= -\left\|a_{u,i}^\top \mathcal{H}_1\left(\bar{\mathsf{u}}^{[1,n_\xi+n_\eta]}(t\,|\,k)\right)\right\|_M^2 \le 0,\quad M \doteq \left(I_{n_\xi+n_\eta} - T(k)\,T^\top(k)\right) \succeq 0.\end{aligned}$$

Note that M is positive semi-definite since $T(k) \in \mathbb{R}^{(n_\xi+n_\eta)\times n_\xi}$ is orthonormal. This holds similarly for the PCE coefficients of outputs. Therefore, applying the projection of the feasible solution (A.3) by (A.4), the left-hand sides of (7.13e)–(7.13f) do not increase. Thus, the chance constraints are satisfied with $\{(\bar{\mathsf{u}},\bar{\mathsf{y}},\bar{\mathsf{z}})^j(t\,|\,k)\}_{t=0}^{N-1}$. Similarly, applying (A.4) the left-hand side of the terminal constraint (7.13g) also does not increase. Thus, we conclude the feasibility of OCP with respect to the basis (7.8) and the initial condition specified per (7.20). □

A.2. Proof of Theorem 7.1

Proof. Note that assertion i) regarding recursive feasibility directly follows from Proposition 7.1 and the procedure of initial condition selection. In the following, we first prove the closed-loop cost decay condition (7.27b) (assertion ii)) and then use it to establish the average asymptotic performance bound (7.27c) (assertion iii)).

Assertion ii): To show that the closed-loop cost decay between time instant k and $k+1$ results from the selection (or interpolation) of the initial condition, we rely on the optimal solution at time instant k, i.e. $(\mathsf{u},\mathsf{y},\mathsf{z})_{[0,N-1]}^{[0,L-1],\star}(k)$ and its shifted trajectories $(\tilde{\mathsf{u}},\tilde{\mathsf{y}},\tilde{\mathsf{z}})_{[0,N-1]}^{[0,L+n_\eta-1]}(k+1)$ from (A.3). Furthermore, using (7.23) and (A.4), we obtain

the candidate solution $(\bar{\mathrm{u}}, \bar{\mathrm{y}}, \bar{\mathrm{z}})_{[0,N-1]}^{[0,L-1]}(k+1)$ for the OCP with the backup initial condition (7.20). For clarity, we simplify $(i \,|\, k+1)$ and $(i \,|\, k)$ to (i) whereever this does not introduce ambiguity.

We evaluate the candidate solution's resulting performance using (7.12)

$$\bar{J}_N(k+1) \doteq \sum_{t=0}^{N-1} (\bar{\ell}_{y,t} + \bar{\ell}_{u,t}) + \bar{\ell}_{z,N}.$$

We show that $\bar{J}_N(k+1)$ can be expressed using the shifted solution $(\tilde{\mathrm{u}}, \tilde{\mathrm{y}}, \tilde{\mathrm{z}})$. We first link the part $\bar{\ell}_{y,t}$ in (7.12) to $\tilde{\mathrm{y}}$. With $n_\xi = 1$, we have

$$\begin{aligned}
\bar{\ell}_{y,t} &\doteq \|\bar{\mathrm{y}}^0(i) + \bar{\mathrm{y}}^1(i)\,\xi(k+1,\omega)\|_Q^2 + \sum_{j=2}^{L-1} \|\bar{\mathrm{y}}^j(i)\|_Q^2 \\
&\overset{(A.4)}{=} \left\| \tilde{\mathrm{y}}^0(i) + \mathcal{H}_1\left(\tilde{\mathrm{y}}^{[1,1+n_\eta]}(i)\right) T(k)\, T(k)^\top \begin{bmatrix} \xi(k,\omega) \\ \eta(k,\omega) \end{bmatrix} \right\|_Q^2 + \sum_{j=2+n_\eta}^{L+n_\eta-1} \|\tilde{\mathrm{y}}^j(i)\|_Q^2 \\
&\overset{(7.23)}{=} \left\| \tilde{\mathrm{y}}^0(i) + \mathcal{H}_1\left(\tilde{\mathrm{y}}^{[1,1+n_\eta]}(i)\right) \begin{bmatrix} \xi(k,\omega) \\ \eta(k,\omega) \end{bmatrix} \right\|_Q^2 + \sum_{j=2+n_\eta}^{L+n_\eta-1} \|\tilde{\mathrm{y}}^j(i)\|_Q^2.
\end{aligned}$$

Note that the above evaluation requires the realization of $\eta(k,\omega)$ at time instant k. Recall $\eta(k)$ is the stochastic germ corresponding to the disturbance $V(k)$, whose realization is not available at time instant k. However, its mean $\mathbb{E}[\eta(k)] = 0$ and covariance $\Sigma[\eta(k)] = I_{n_\eta}$ are known. Hence, we can determine the expected value of $\bar{\ell}$.

Denote $\mathbb{E}[\bar{\ell}_{y,t} \,|\, k]$ as the conditional expectation of $\bar{\ell}_{y,t}$ at time instant k given $\xi(k,\omega)$ but with unknown $\eta(k,\omega)$. Let $\tilde{y}(t) \doteq \tilde{\mathrm{y}}^0(t) + \tilde{\mathrm{y}}^1(t)\,\xi(k,\omega)$, we have

$$\begin{aligned}
\mathbb{E}[\bar{\ell}_{y,t} \,|\, k] &= \mathbb{E}\left[\|\tilde{y}(t) + \mathcal{H}_1(\tilde{\mathrm{y}}^{[2,1+n_\eta]}(t))\eta(k)\|_Q^2\right] + \sum_{j=2+n_\eta}^{L+n_\eta-1} \|\tilde{\mathrm{y}}^j(i)\|_Q^2 \\
&= \|\tilde{y}(t)\|_Q^2 + \operatorname{trace}\left(\mathcal{H}_1(\tilde{\mathrm{y}}^{[2,1+n_\eta]}(t))^\top Q \mathcal{H}_1(\tilde{\mathrm{y}}^{[2,1+n_\eta]}(t)) \Sigma[\eta(k)]\right) \qquad \text{(A.5)} \\
&\quad + 2\tilde{y}^\top(t)\, Q \mathcal{H}_1\left(\tilde{\mathrm{y}}^{[2,1+n_\eta]}(t)\right) \mathbb{E}[\eta(k)] + \sum_{j=2+n_\eta}^{L+n_\eta-1} \|\tilde{\mathrm{y}}^j(t)\|_Q^2 \qquad \text{(A.6)} \\
&= \|\tilde{\mathrm{y}}^0(t) + \tilde{\mathrm{y}}^1(t)\,\xi(k,\omega)\|_Q^2 + \sum_{j=2}^{L+n_\eta-1} \|\tilde{\mathrm{y}}^j(t)\|_Q^2 \doteq \tilde{\ell}_{y,t}.
\end{aligned}$$

Similarly, the above also holds for the other parts of $\bar{J}_N(k+1)$, i.e., $\mathbb{E}[\bar{\ell}_{u,t} \,|\, k] = \tilde{\ell}_{u,t}$ and $\mathbb{E}[\bar{\ell}_{z,N} \,|\, k] = \tilde{\ell}_{z,N}$. As a result, we have

$$\mathbb{E}[\bar{J}_N(k+1) \,|\, k] = \sum_{t=0}^{N-1} (\tilde{\ell}_{y,t} + \tilde{\ell}_{u,t}) + \tilde{\ell}_{z,N}.$$

Moreover, from (7.24) and (7.26), we have $V_N(k+1) \leq \bar{J}_N(k+1)$ for both strategies of selection or interpolation of initial conditions. Thus, we have

$$\mathbb{E}[\mathcal{V}_N(k+1) \,|\, k] \leq \mathbb{E}[\bar{J}_N(k+1) \,|\, k].$$

Now, we have the upper bound for $\mathbb{E}[\mathcal{V}_N(k+1) \mid k] - V_N(k)$ as

$$\begin{aligned}&\mathbb{E}[\mathcal{V}_N(k+1) \mid k] - V_N(k) \leq \mathbb{E}[\bar{J}_N(k+1) \mid k] - V_N(k)\\&= \tilde{\ell}_{y,N-1} + \tilde{\ell}_{u,N-1} + \tilde{\ell}_{z,N} - \ell^\star_{u,0} - \ell^\star_{y,0} - \ell^\star_{z,N}\end{aligned}$$

where the identical terms $\{\tilde{\ell}_{u,t}, \tilde{\ell}_{y,t}\}_{t=0}^{N-2}$ and $\{\ell^\star_{u,t}, \ell^\star_{y,t}\}_{t=1}^{N-1}$ are cancelled. Furthermore,

$$\begin{aligned}&\tilde{\ell}_{y,N-1} + \tilde{\ell}_{u,N-1} + \tilde{\ell}_{z,N} - \ell^\star_{z,N}\\&= \|\tilde{\mathsf{y}}^0(N-1) + \tilde{\mathsf{y}}^1(N-1)\,\xi(k,\omega)\|_Q^2 + \textstyle\sum_{j=2}^{L-1} \|\tilde{\mathsf{y}}^j(N-1)\|_Q^2 && (i)\\&+ \|\tilde{\mathsf{u}}^0(N-1) + \tilde{\mathsf{u}}^1(N-1)\,\xi(k,\omega)\|_R^2 + \textstyle\sum_{j=2}^{L-1} \|\tilde{\mathsf{u}}^j(N-1)\|_R^2 && (ii)\\&+ \|\tilde{\mathsf{z}}^0(N) + \tilde{\mathsf{z}}^1(N)\,\xi(k,\omega)\|_P^2 + \textstyle\sum_{j=2}^{L-1} \|\tilde{\mathsf{z}}^j(N)\|_P^2 && (iii)\\&- \|\mathsf{z}^{0,\star}(N) + \mathsf{z}^{1,\star}(N)\,\xi(k,\omega)\|_P^2 - \textstyle\sum_{j=2}^{L-1} \|\mathsf{z}^{j,\star}(N-1)\|_P^2 && (iv)\\&+ \textstyle\sum_{j=L}^{L+n_\eta-1} \left(\|\tilde{\mathsf{y}}^j(N-1)\|_Q^2 + \|\tilde{\mathsf{u}}^j(N-1)\|_R^2 + \|\tilde{\mathsf{z}}^j(N)\|_P^2\right). && (v)\end{aligned}$$

Exploiting (A.3), we rewrite the terms $(i)+(ii)+(iii)$ and (v)

$$(i)+(ii)+(iii) = \|\mathsf{z}^{0,\star}(N) + \mathsf{z}^{1,\star}(N)\,\xi(k,\omega)\|_\Xi^2 + \sum_{j=2}^{L-1} \|\mathsf{z}^{j,\star}(N-1)\|_\Xi^2,$$

$$\Xi = K^\top R K + A_K^\top \widetilde{H}^\top Q \widetilde{H} A_K + A_K^\top P A_K \overset{(7.7a)}{=} P,$$

$$(v) \overset{(A.3)}{=} \sum_{j=L}^{L+n_\eta-1} \|\mathsf{v}^j(N)\|_{Q+\widetilde{H}^\top P \widetilde{H}}^2 = \operatorname{trace}(\Sigma_V(Q + \widetilde{H}^\top P \widetilde{H})) = \beta.$$

Hence, $(i)+(ii)+(iii)+(iv) = 0$ and $(v) = \beta$.

Combining the previous results, we arrive at

$$\begin{aligned}&\mathbb{E}[\mathcal{V}_N(k+1) \mid k] - V_N(k) \leq -\ell^\star_{u,0} - \ell^\star_{y,0} + \beta\\&= -\left(\|\mathsf{u}^{0,\star}(0 \mid k) + \mathsf{u}^{1,\star}(0 \mid k)\,\xi(k,\omega)\|_R^2 + \|\mathsf{y}^{0,\star}(0 \mid k) + \mathsf{y}^{1,\star}(0 \mid k)\,\xi(k,\omega)\|_Q^2\right) + \beta\\&= -\|u(k)\|_R^2 - \mathbb{E}\left[\|Y(0 \mid k)\|_Q^2\right] + \beta.\end{aligned}$$

As the above inequality holds for each closed-loop implementation, we lift $V_N(k+1) - V_N(k)$ to its random variable counterpart and obtain (7.27b).

Assertion iii): Recursively using (7.27b) from 0 to k and let $k \to \infty$, we have

$$\begin{aligned}0 \leq \lim_{k\to\infty} \frac{1}{k} \mathbb{E}\left[\mathcal{V}_N(k) - \mathcal{V}_N(0)\right] &= \lim_{k\to\infty} \frac{1}{k} \sum_{i=0}^{k} \left(\mathbb{E}\left[\mathcal{V}_N(i+1) - \mathcal{V}_N(i)\right]\right)\\ \leq \beta - \frac{1}{k} \lim_{k\to\infty} \sum_{i=0}^{k} &\mathbb{E}\left[\|U(i)\|_R^2 + \|Y(i)\|_Q^2\right].\end{aligned}$$

Thus, we obtain the average cost condition (7.27c) for the closed-loop system in random variables (7.1). □

A.3. Proof of Proposition 7.2

Proof. At time instant k, given $z(k)$, $u(k)$, and $e(k)$, the model mismatch between (7.28a) and (7.28b) leads to a prediction error of $y(k)$ as $\delta(k) = y(k) - \hat{y}(k) = \Delta \begin{bmatrix} z(k) \\ u(k) \end{bmatrix}$.

By appending $y(k)$ and $\hat{y}(k)$ to previously recorded inputs and outputs, we obtain

$$z(k+1) = \left[u_{[k-\ell+2,k+1]}^\top, y_{[k-\ell+2,k]}^\top, y^\top(k)\right]^\top, \quad \hat{z}(k+1) = \left[u_{[k-\ell+2,k+1]}^\top, y_{[k-\ell+2,k]}^\top, \hat{y}^\top(k)\right]^\top,$$

which leads to different $\eta(k,\omega)$ and $\hat{\eta}(k,\omega)$ by (7.23c).

Note that $\mathbb{E}[\hat{\eta}(k)] = 0$ and $\Sigma[\hat{\eta}(k)] = I_{n_\eta}$, and

$$\begin{aligned}
\eta(k) - \hat{\eta}(k) &= \mathsf{M}_v^\dagger \widetilde{H}^\top (z(k+1) - \hat{z}(k+1)) = \mathsf{M}_v^\dagger \delta(k), \\
\mathbb{E}[\eta(k)] &= \mathbb{E}[\hat{\eta}(k)] + \eta(k) - \hat{\eta}(k) = \eta(k) - \hat{\eta}(k) = \mathsf{M}_v^\dagger \delta(k), \\
\Sigma[\eta(k)] &= \Sigma[\hat{\eta}(k) + \eta(k) - \hat{\eta}(k)] = \Sigma[\hat{\eta}(k)] = I_\eta.
\end{aligned}$$

Due to $\mathbb{E}[\eta(k)] \neq \mathbb{E}[\hat{\eta}(k)] = 0$, (A.5) in the proof of Theorem 7.1 does not hold but

$$\mathbb{E}[\bar{\ell}_{y,t} \,|\, k] = \tilde{\ell}_{y,t} + 2\tilde{y}^\top(t) Q \mathcal{H}_1 \left(\tilde{\mathsf{y}}^{[2,1+n_\eta]}(t)\right) \mathsf{M}_v^\dagger \delta(k).$$

Note that the shifted solution $(\tilde{\mathsf{u}}, \tilde{\mathsf{y}})$ is constructed from the optimal solution of the last time instant. Given the optimal solution lives in a bounded set due to Assumption 7.6, there exists $C_{y,t} \in \mathbb{R}^+$ such that

$$\mathbb{E}[\bar{\ell}_{y,t} \,|\, k] \leq \tilde{\ell}_{y,t} + C_{y,t} \|\delta(k)\|.$$

Subsequently, there exist $C_{y,t}, C_{u,t} \in \mathbb{R}^+$ for $i \in \mathbb{I}_{[0,N-1]}$, $C_{z,N} \in \mathbb{R}^+$, and $C \in \mathbb{R}^+$ such that

$$\begin{aligned}
&\mathbb{E}[\bar{J}_N(k+1) \,|\, k] \leq \textstyle\sum_{i=0}^{N-1} (\tilde{\ell}_{y,i} + \tilde{\ell}_{u,i}) + \tilde{\ell}_{z,N} + \left(\sum_{i=0}^{N-1} (C_{y,t} + C_{u,t}) + C_{z,N}\right) \|\delta(k)\| \\
&\doteq \tilde{J}_N(k+1) + C\|\delta(k)\|.
\end{aligned}$$

Finally, we arrive at

$$\begin{aligned}
&\mathbb{E}[V_N(k+1) \,|\, k] - V_N(k) \leq \mathbb{E}[\bar{J}_N(k+1) \,|\, k] - V_N(k) \leq \tilde{J}_N(k+1) - V_N(k) + C\|\delta(k)\| \\
&= \beta + C\|\delta(k)\| - \|u(k)\|_R^2 - \mathbb{E}\left[\|\hat{Y}(k)\|_Q^2 \,\middle|\, k\right] \\
&= \beta + C\|\delta(k)\| - \|u(k)\|_R^2 - \mathbb{E}\left[\|Y(k) - \delta(k)\|_Q^2 \,\middle|\, k\right].
\end{aligned}$$

As the above inequality holds for all $\omega \in \Omega$, we lift it to its random variable counterpart and thus arrive at (7.30). □