Forecasting and decision-making for empty container repositioning

Sommer, Benedikt

Publication date:
2023

Document Version
Publisher's PDF, also known as Version of record

Ph.D Thesis
Doctor of Philosophy

DTU Management
Department of Technology, Management and Economics

Forecasting and decision-making for empty container repositioning

Benedikt Sommer

Academic Advisors:
Dario Pacino, Technical University of Denmark
Klaus Kähler Holst, A. P. Møller-Mærsk
Pierre Pinson, Imperial College London
Submitted: June 30, 2023

DTU

DTU Management
Department of Technology, Management and Economics
Technical University of Denmark

Akademivej
Building 358
2800 Kongens Lyngby, Denmark
Phone +45 2548 00
www.man.dtu.dk

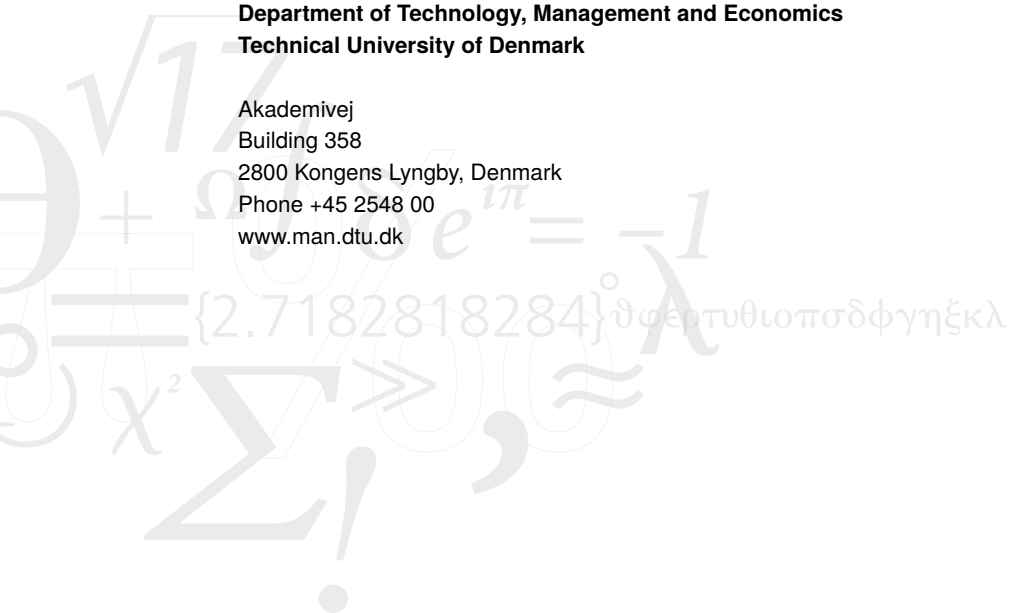

Light weight, baby.

Ronnie Coleman

Preface

This thesis is prepared at the Department of Technology, Management and Economics of the Technical University of Denmark in partial fulfillment of the requirements for acquiring the degree of Doctor of Philosophy in Engineering. The Ph.D. project was funded by the InnovationsFonden Danmark (grant number 9065-00021B) and was conducted in collaboration with A. P. Møller-Mærsk.

This dissertation summarizes the work carried out in the period 01.01.2020 to 30.06.2023. The thesis is composed of six chapters and three attached scientific papers, two of which are under revision, whereas the remaining one is currently under review.

Benedikt Sommer
Copenhagen, June 2023

Acknowledgements

First of all, I would like to thank Klaus Holst for granting me the opportunity of pursuing this industrial PhD project. I owe you the greatest gratitude for your academic supervision, leadership, and mentorship. You have been the best manager that I could have asked for during this project and your contributions had a tremendous impact on my academic and professional development. Next, I would like to thank Pierre Pinson for his academic and personal guidance since 2018. You have been there when it mattered the most. Merci beaucoup. I would also like to thank Dario Pacino for taking over the academic supervisor role at DTU after Pierre's departure to London. Finally, I want to thank Trine Krogh Boomsma for her valuable contributions towards our paper about empty container inventory management at depots.

I am also grateful to all my colleagues in the former research team at Maersk. In particular, I want to thank Baptiste Coutton for his friendship and comments that improved this thesis. I would also like to thank Sangmin Lee and Annick Verhoeven for their valuable comments.

Especially, I want to express my appreciation to all my friends. Your open ears about the hardships of this PhD project and support during the past months have been indispensable. To my girlfriend Lina, thank you for keeping up with my worries towards the end of this PhD project. Your love and care have been sturdy pillars.

Last, I am deeply indebted to my sister and brother-in-law for their constant support. Your love throughout the best part of my life means a lot to me.

Abstract

Containerized shipping serves as the backbone of our interconnected global economy, enabling low-cost and secure trade of goods between importers and exporters across long distances. Each containerized trade journey begins with exporters loading freight into empty containers, which are provided by a shipping company. It continues with the transportation of laden containers to importers, and ends when the importers unload the freight. Shipping companies, as the primary owners of containers, face the consequences of regional trade imbalances, where empty containers accumulate when the demand for receiving freight from importers exceeds the demand from exporters to ship their goods. Conversely, additional empty containers are needed in regions with a higher number of exporters than importers. This thesis focuses on the challenges shipping companies encounter as a result of these trade imbalances, better known as the empty container repositioning problem. Our contributions are based on a comprehensive problem description, that enables us to identify crucial elements of a decision system to support container shipping companies in their efforts to reduce empty container repositioning costs.

First, we propose a new inventory model for inland container depots, where empty containers are stored when not being used for shipping. We identify the inventory management problem as a sequential decision process, for which we obtain optimal decisions under realistic assumptions. Specifically, we address lead times for transporting empty containers between ports and inland depots, as well as the impact on depot operating costs when important aspects of container deliveries to exporters and returns from importers are not appropriately modeled. A case study using real-world data supports our findings and advises future inventory modeling to pay close attention to the stochastic processes describing exporter and importer behavior.

Second, we develop a novel forecasting model for future empty container demand from exporters and returns from importers. These forecasts inform shipping companies about container shortages in advance, and are therefore crucial for repositioning decisions. Our proposal extends the well-known state space modeling framework to direct multi-step ahead forecasting, which is a strategy that is commonly applied in forecasting problems with non-stationary time series. The models are derived under the theoretical properties of direct multi-step ahead forecasting and are therefore readily applicable to forecasting problems in other domains. Third, we investigate the applicability of modern hybrid forecasting models in the context of the empty container repositioning problem. We critically examine the practical usability of a recently proposed model that aims to combine the strengths of neural networks and state space models. Our investigation reveals a critical assumption for the time series data that needs to be fulfilled by a forecasting problem. Through simulation studies, we verify our claims and demonstrate the negative effects on forecast accuracies when a forecasting problem does not meet the assumptions of the hybrid forecasting model. Our results are general and expected to guide future developments in hybrid forecasting models.

Dansk resumé

Containerskibsfart spiller en afgørende rolle i den globale økonomi ved at muliggøre billig og sikker handel med varer mellem importører og eksportører over lange afstande. En containerrejse starter med, at eksportørerne laster fragt i tomme containere, som de modtager fra et containerskibsfartselskab. De fyldte containere transporteres herefter til importørerne, og endeligt afsluttes rejsen, når importørerne losser fragten. Skibsfartsselskaberne står over for konsekvenserne af regionale handelsubalancer, hvor tomme containere akkumuleres, når efterspørgslen efter at modtage fragt fra importørerne overstiger efterspørgslen fra eksportørerne for at sende deres varer. Omvendt er der behov for yderligere tomme containere i regioner med flere eksportører end importører. Denne afhandling fokuserer på de udfordringer, som skibsfartsselskaber står over for som følge af disse handelsubalancer, kendt som Empty Container Repositioning problemet. Vores bidrag er baseret på en omfattende problem beskrivelse, der gør det muligt for os at identificere afgørende elementer i et beslutningssystem, der kan støtte containerrederier i at reducere omkostningerne ved omplacering af tomme containers.

For det første foreslår vi en ny lager-model til containerdepoter, hvor tomme containere opbevares, når de ikke bruges til at sende fragt. Vi identificerer lagerstyringsproblemet som en sekventiel beslutningsproces, hvor vi opnår optimale beslutninger under realistiske antagelser. Specifikt håndterer vi leveringstider for transport af tomme containere mellem havne og indlandsdepoter. Vi illusterer også konsekvenserne på driftsomkostningerne, når vigtige aspekter af containerleverancer til eksportører og tilbageleveringer fra importører ikke er korrekt modelleret. Et case-studie understøtter vores resultater og anbefaler, at lagermodeller bør nøje tage hensyn til de stokastiske modeller der beskriver eksportørernes og importørernes adfærd.

For det andet udvikler vi en ny og innovativ prognosemodel for fremtidige behov for tomme containere fra eksportører og tilbageleveringer fra importører. Disse prognoser informerer containerskibsfartselskabet om containermangel på forhånd og er derfor afgørende for beslutninger om omplacering. Vores forslag udvider den velkendte ramme for state space modeller til direkte flertrinsprognoser, hvilket er en strategi, der ofte anvendes til prognoseproblemer med ikke-stationære tidsrækker. Modellerne er udledt under de teoretiske egenskaber ved direkte flertrinsprognoser og er derfor let anvendelige til prognoseproblemer inden for andre felter.

For det tredje undersøger vi anvendelsen af moderne hybrid-prognosemodeller i forbindelse med prognosemodellering for omplacering af tomme containere. Vi undersøger kritisk anvendeligheden af en nyligt foreslået metode, der sigter mod at kombinere styrkerne ved neurale netværk og state space modeller. Vores undersøgelse afslører, at hvis en afgørende antagelse for tidsserie-data ikke er tilstede, kan det forventes at prognosemodellen forringes væsentligt. Gennem simulationsstudier verificerer vi vores påstande og demonstrerer de negative påvirkninger på prognosepræcisionen, når et prognoseproblem ikke opfylder antagelserne i hybrid-prognosemodellen. Vores resultater er generelle og forventes at vejlede fremtidige udviklinger inden for hybrid-prognosemodeller.

Contents

Introduction

1.1 Empty container repositioning

The usage of standardized steel containers revolutionized the way goods are transported around the world, making it faster and more efficient to move large quantities of cargo over long distances by sea. At the core of containerized trade is the delivery of empty containers to export customers. Once the *sender*, or exporter, has loaded cargo into a container, a container shipping company transports the container to its destination. Upon arrival at the location of the cargo receiver, or importer, the freight is unloaded, and the container is returned back to the shipping company. The empty container is typically delivered to a depot, where it is stored until it is needed for a new export shipment request. One of the major challenges faced by global containerized shipping is the mismatch between the availability of returned containers and the demand of exporters. Regions that have a surplus of cargo exports, such as large parts of China and South-East Asia, often experience a year-round deficit of empty containers. To address this issue, shipping companies need to transport empty containers between regions to compensate for trade imbalances of laden containers. Without such repositioning efforts, exporting regions would suffer from a shortage of empty containers. Container deficits and surpluses may also vary over time, such as in countries with harvesting seasons (high demand) and regular demand periods. The task, that is known as *empty container repositioning* (ECR), requires at a global scale the anticipation of demand and supply imbalances ahead of time due to the slow transportation speed of large-scale container vessels.

The reliable provision of empty containers to exporters requires significant investments from shipping companies, amounting to billions of US dollars for the management of their empty container fleets. Modern fleets cover a diverse range of container specifications to cater for the needs of customers to ship goods with different requirements. Standard 20ft or 40ft long steel containers are the most frequently used types, of which refrigerated versions are available for shipping perishable products. Notteboom et al. (2021) report that all global shipping companies combined spend about $110 billion per year for the management of their container assets (purchase, maintenance and repairs), of which $16 billion is attributed to ECR by sea. Addressing inefficiencies in empty container management processes is not only financially beneficial for shipping companies but also carries environmental and operational advantages. Trucks, the primary mode of transportation for inland repositioning, cause emissions and congestions in port areas (Jula et al., 2006; Boile et al., 2008; Lee and Song, 2017). Inefficient repositioning decisions can further affect port operations due to the accumulation of empty containers (Song and Dong, 2022). Reducing container movements by sea and on land can thus reduce emissions and improve port operations.

The severe economical, environmental and societal impacts of ECR activities have led to extensive research in this field. Improving decision-making for ECR of a shipping company cannot be addressed in isolation, as it is interconnected with decisions within the broader empty container management context. Increasing the container fleet size provides more flexibility for repositioning decisions, but it entails higher capital investments and associated costs for

maintenance and storage (Imai and Rivera, 2001). Other solutions, such as using foldable containers (Lee and Moon, 2020), strategically selecting depots near transportation hubs and customers, or designing shipping networks considering empty and laden container flows, fall under the same category. Efficient decision-making across all planning levels is a subsequent challenge, given their distinct time resolutions and horizons (Braekers et al., 2011).

This thesis focuses on operational decisions for repositioning empty containers between depots with look-ahead times up to several months. Our scope is limited to a shipping company which operates an ocean and inland network, with container vessels connecting a set of ports. Ports are connected to dry ports of the inland network by trains, barges and trucks. Owned and long-term leased containers are stored at depots, where empty containers are returned to after an importer unloaded their freight from a received container. However, the main responsibility of a container depot is to provide empty containers to nearby exporters. This necessitates maintaining an adequate inventory of containers, reducing the risk of future shortages as a result of uncertain imbalances of container returns and demand. An ECR decision-making system of a shipping company relies on two principal types of decision models. A class of inventory models determine whether additional containers are needed at depots as well as a class of routing models to identify cost-effective routes for empty containers between depots. An essential element of the inventory models are forecasts of future imbalances between empty container demand of exporters and returns of importers. Developing an accurate ECR decision-making system is extremely challenging because the dynamic operations of a container shipping company induce a sequential dependency for present and future repositioning decisions, whereby all present decisions are made under uncertainty. This is best known as sequential decision-making under uncertainty.

1.2 Challenges and research directions

In the following, we explore different challenges for the development of decision and forecasting models for ECR, where we focus on mathematical modeling only. Even though we consider their challenges subsequently in isolation, we follow Goltsos et al. (2022) in that the integration and interaction of forecasting and decision-making models ultimately determine the overall value of a decision system. Based on the presented challenges, we introduce three research directions (**RD**s) for this thesis.

The inability to formulate and solve a single mathematical model that accurately represents the container inventory and routing decision-making processes for real-world shipping networks has led to the emergence of two research streams. The first research stream focuses primarily on addressing the empty container routing problem between depots, utilizing network flow models. However, the sequential nature of this decision problem presents a significant challenge, as multi-stage stochastic programs become computationally infeasible for realistic problem sizes due to the exponential growth of the number of decision variables relative to the number of planning stages (Shapiro and Nemirovski, 2005). Two-stage stochastic network flow models have emerged to reduce the problem complexity by assuming all first stage parameters to be deterministic (Crainic et al., 1993; Cheung and Chen, 1998; Shu and Song, 2014; Erera et al., 2009; Long et al., 2012). The applicability of these models remains computationally challenging for large real-world networks and flow models with only deterministic parameters may be considered instead. Examples for repositioning problems between sea ports include the proposals in Song and Dong (2011); Brouer et al. (2011) and in Shintani et al. (2010); Olivo et al. (2013) for inland operations. Optimal methods based on stochastic dynamic programming (SDP) have been developed for two depot systems in Song (2007); Song et al. (2010); Ng

et al. (2012). SDP proves advantageous for sequential decision problems since the number of decisions grows linearly with the number of planning stages. However, SDP suffers from its own curse of dimensionality, with exponentially increasing complexity as the number of states expands. Consequently, approximation methods, such as approximate dynamic programming, have been proposed for multi-depot systems in Lam et al. (2007).

The second research stream places more emphasis on the multi-stage nature of the ECR problem. Rather than directly determining repositioning decisions between depots, this research stream develops inventory control models to determine the appropriate quantity of empty containers to be moved out (out-positioning) or brought in (in-positioning). Decisions are made sequentially and aim to maintain balanced inventory levels in the face of uncertain future demand and returns of empty containers. Optimal control policies are developed most commonly for single depots (Li et al., 2004; Song and Zhang, 2010; Zhang et al., 2014; Song and Dong, 2022). The overall ECR problem is effectively decomposed by using inventory control policies to determine the desired number of in- and out-positioned containers for each depot, operating in conjunction with a routing model for repositioning empty containers among all depots. Simple threshold-type control policies for depots are commonly combined with heuristics that determine repositioning decisions (Li et al., 2007; Dang et al., 2012, 2013; Zhang et al., 2014). Heuristics that incorporate detailed information about the transportation network, such as vessel departure times, are developed in Song and Dong (2008); Dong and Song (2009). Additionally, the combination of inventory models with network flow models has been explored in Chou et al. (2010); Lee et al. (2012); Epstein et al. (2012). The reliance on accurate inventory control policies to determine the flow of empty containers within this decomposition approach serves as motivation for our first research direction:

[RD1] Inventory control policies for container depots

Indeed, the ability of a depot to effectively in- and out-position empty containers relies on various factors, including transportation capacities and inventory levels at other depots. For instance, empty containers can be repositioned between oversea ports for moderate prices but with long lead times (Dang et al., 2012). The lead times for in-positioning containers at deficit ports, which are in need of containers, are dependent on transportation times from surplus ports, which, in turn, are influenced by the shipping network (Dang et al., 2012). Incorporating wrong assumptions about transportation times into the inventory model can have negative effects when actual times deviate from their assumed values. A particular case are in-positioning options with falsely assumed very short lead times. Larger than expected costs can accrue for deficit ports due to container shortages when orders are made but assumed lead times are exceeded in real-world operations. Deriving control policies also requires a description of the stochastic process for exporter container demand and importer returns. A common assumption in the ECR control literature is that both quantities are independent and identically distributed, cf. Li et al. (2004); Song and Zhang (2010); Zhang et al. (2014); Song and Dong (2022). While this assumption allows for the analytical derivation of optimal policies, it is likely too simplistic to accurately capture the dynamics of real-world demand and return processes. The impact of serially dependent demand in inventory systems has been explored outside the ECR literature in Ray (1980); Graves (1999). Safety stock levels are found to be underestimated when positive autocorrelation of a demand process is undetected, and overestimated for undetected negative autocorrelation. Therefore, it is essential to account for these dependencies to avoid incurring additional costs in inventory management. Crainic et al. (1993) emphasized in this regard that accurate forecasts are critical for successful ECR planning. In this thesis, we adopt the view that forecasting can be employed to find an appropriate approximation of the unknown empty container demand and return processes from historical data.

A broad research body has evolved for forecasting unknown future freight rates, container throughput, vessel arrival times, container flows and market sizes. Albeit few publications consider forecasting within the ECR context, we find that the related publications apply widely-used methods without providing methodological contributions on the basis of the forecasting problem's fundamental characteristics. To the best of our knowledge, we find Martius et al. (2022) to be the exception with a bottom-up approach that utilizes the wealth of container movement data. Instead of following the prevailing approach of directly forecasting total container in- and outflows at depots, their method predicts the arrival and departure of individual containers. As Martius et al. (2022) note themselves, their method is challenged by the impact of rare events, such as the recent COVID-19 pandemic, general consumption pattern changes and the reduction of the global container fleet's size on empty container demand and returns. It is therefore essential to recognize time-varying customers dynamics to be one of the factors that imply non-stationary empty container demand and return processes. However, we find that none of the existing ECR forecasting publications address the challenges that are induced by non-stationarities. This motivates our next research direction:

[RD2] Forecasting non-stationary time series processes

Differencing has been traditionally applied to forecast time series with unit roots. Over time, the field has evolved to model time series processes with various non-stationarity patterns, such as mean or volatility changes due to structural breaks, where differencing is dispensable. Recent approaches involve local stationarity concepts (Dahlhaus, 2012) to theoretically motivate the estimation of stationary models on short time intervals. Models can adapt to changes in the time series process by re-estimating the model parameters only on most recent data, without considering the full history of the time series. The class of time-varying coefficient models is closely related and assumes the model coefficients to be time-inhomogeneous, hence implicitly considers the unknown stochastic process to be non-stationary (Grenier, 1983). Although each method has its own merits and caveats, the scale and complexity of the ECR problem poses a challenge to all of them. First, thousands of forecasts must be issued simultaneously for container depot and various container type combinations. Second, the large time series corpus is heterogeneous. While some time series share common patterns, such as seasonalities, volatility, trends and non-stationarities, other time series are unrelated.

The traditional *local* strategy is to estimate a set of univariate candidate models for each of the time series, followed by model selection to identify the most suitable for each series. Developing suitable candidate models requires knowledge about the time series processes, and the model tuning and selection problem can become computationally demanding when the number of time series is large. Since each forecasting model is estimated for a time series independently, no advantage is gained from the relatedness of some time series within the dataset. *Global* methods aim to overcome these limitations and estimate a single forecasting model for all time series to exploit potential similarities between them. The method commonly employs flexible non-parametric machine learning models, such as a recurrent neural network in Salinas et al. (2020), to enable adaption to complex time series processes in heterogeneous datasets without explicit assumptions on the unknown processes (Montero-Manso and Hyndman, 2021). Cross-validation is essential for flexible machine learning models to detect overfitting, and thus validating the generalizability of the model on unseen data. The risk of overfitting short time series can be alleviated for locally estimated models by imposing restrictions through structural assumptions about the time series processes, i.e. estimate parsimonious parametric forecasting models. Our last research direction explores the applicability of recently emerged hybrid forecasting models. These hybrids aim to combine "the best of both worlds" by using a flexible globally estimated machine learning model to predict the parameters of locally applied

parametric time series models:

[RD3] Fusing machine learning models with parametric time series models

Hybrid models that combine aspects of local and global strategies have been proposed for time series panel forecasting to overcome limitations of the individual strategies. A recently evolved literature strain proposes to incorporate structural assumptions about the stochastic process of individual time series into a global forecasting model (Lim and Zohren, 2021). Rangapuram et al. (2018) follow such a strategy and use a globally estimated machine learning model to predict the unknown parameters of parametric time series models which are locally applied to each time series of the panel. The specification of a local sub-model allows the forecaster to incorporate structural assumptions about seasonal, trend or other domain-specific time series properties into a global modeling framework. This class of hybrid models are believed to have several merits over flexible global models in forecasting problems with limited amount of data. The time series shortness is indeed a challenge in the ECR forecasting problem since automated data collection has not always been a priority to container shipping companies. On one hand, the local sub-model induces restrictions on the otherwise too flexible global machine learning, hence there is the possibility of alleviating overfitting. On the other hand, learning can become more efficient when the sub-model is correctly specified and multiple time series have shared patterns that the global model can learn.

1.3 Research contributions

Motivated by these developments, the main objective of this thesis is to develop decision and forecasting models for the ECR problem. We cover a broad range of topics in the decision-making and forecasting disciplines, and explore the requirements of a proposed decision model towards the forecasting model. First, this thesis explores the practical challenges and limitations of automated ECR decision-making for global shipping companies. Our detailed description of the ECR problem justifies to decompose the single large-scale ECR optimization problem into several sub-problems. On the basis of our findings, we design an inventory control model for a single depot. The inventory model is formulated as a sequential decision process with uncertain future container demand of exporters and returns from importers. Second, we propose a novel forecasting method for non-stationary time series. Our contribution is fundamental in that we extend the state space modeling framework to direct multi-step ahead forecasting. We improve previous works by demonstrating the occurrences of estimation biases if state space model parameterizations do not account for the inherent serial correlation of direct multi-step ahead forecasting errors. Our third and last contribution provides a critical view on a popular hybrid method (Rangapuram et al., 2018) for forecasting time series panels which uses a global recurrent neural network to parameterize local state space models. We discuss limitations that primarily stem from imposing the same state space model process for all time series and using a flexible recurrent neural network. Our findings draw attention to the fact that the hybrid model is only suited for a narrow range of forecasting problems. Although our first contribution is tailored to the ECR problem, the second and third contributions are more fundamental. The proposed direct multi-step ahead state space models can be readily applied to other forecasting problems and our analysis of a neural network and state space hybrid provide guidance for the future design of hybrid models.

Towards the first research objective regarding inventory control policies for container depots, we extend existing Markov Decisions Processes (MDPs) in two directions that allow us to estimate optimal inventory policies for container depots. In [**Paper A**], we present a capacitated

multiple supplier periodic-review inventory model for daily ECR planning of an inland container depot. Our MDP formulation improves on previous approaches by incorporating the different in-positioning options available to an inland container depot, considering varying transportation times and costs. The capacity of in-positioning options can reflect the unavailability of receiving empty containers within a specified time due to container shortages at other depots. Since these capacities naturally fluctuate over time based on variations in customer-based empty container demand and return behavior, we formulate a time-inhomogeneous MDP. The common assumption to date has been conversely that delivery times and cost for receiving empty containers are constant (Li et al., 2004; Song and Zhang, 2010; Young Yun et al., 2011; Song and Dong, 2022). On one hand, this prevents the estimation of policies that find optimal trade-offs between slower but cheaper, and faster but more expensive, modes of transportation. Trade-offs exist because requesting empty containers with greater lead times adds additional uncertainty on short-term inventory developments due to unknown future empty container demand and returns. On the other hand, only modeling constant short lead times can be restrictive for real-world operations when container shortages cause lead times to be greater.

Similar to the MDPs of Li et al. (2004); Song and Dong (2022), our model assumes that empty container returns from importers have a one decision period delay before they can be reused, hence container returns cannot satisfy demand within the same period. Incorporating this assumption in an MDP requires the specification of a stochastic process for exporter container demand and importer returns, instead of modeling their imbalance. An additional contribution of [**Paper A**] is the extension of stochastic empty container demand and returns to more realistic processes on the basis of real-world data. We investigate the effect of serial and cross-sectional (correlation between demand and returns within the same period) dependencies on the estimated in- and out-positioning policies. Our contribution is the demonstration of additional accrued depot operation costs when policies are estimated with an MDP for which serial and cross-sectional dependencies are misspecified. The results of [**Paper A**] emphasize the need to accurately model the joint demand and return process. The common assumption in the literature to estimate policies under serially and cross-sectionally independent demand and return processes is shown to lead to substantially higher operational costs than estimated. Our MDPs are formulated such that we can apply backwards dynamic programming to estimate optimal policies, which allows us to isolate the effects of our methodological contributions. However, the curse of dimensionality imposes restrictions on the complexity of the demand and return processes. Therefore, the following contributions should be viewed independently of the requirements in [**Paper A**] to formulate a stochastic process, such that backwards dynamic programming remains computationally tractable.

Towards the second research direction, we combine state space modeling and direct multi-step ahead forecasting to predict future observations of non-stationary time series. Direct multi-step strategies may be favored for non-stationary time series because they avoid error accumulation of iterated strategies due to model misspecification, which likely occurs in non-stationary environments (Chevillon, 2007). Forecasting empty container returns has another property that guides us towards direct multi-step ahead forecasting. Namely, present booking data of customers contains partial information about future empty container returns to container depot. It is possible to derive covariates on the basis of a container shipment's destination and expected arrival time, with the distinct property that the covariates's values decrease with an increasing forecasting horizon. This behavior is due to the fact that not all bookings that lead to future container returns are presently observed. Applying the iterated multi-step strategy is challenged by the fact that the estimated 1-step ahead relationship between the covariate and container returns does not apply for longer forecast horizons. Thus, the forecast accuracy is likely to deteriorate for multiple steps ahead. [**Paper B**] is first at describing this property for

forecasting empty container returns. Another property of this forecasting problem is that many time series are non-stationary. Level and volatility changes are present, possibly due to variations in customer behavior or network adaptions of the shipping company. Furthermore, the global COVID-19 pandemic affected supply chains, hence empty container returns of customers.

To this end, we propose to use the state space framework (Durbin and Koopman, 2012) to model empty container return time series and extend it to direct multi-step ahead forecasting. We choose state space models because the modeling framework allows us to estimate a broad range of time series models, such as time-varying coefficient models to account for non-stationarities, in a probabilistic setting. Our proposal in [**Paper B**] is to explicitly model the serial correlation of multi-step ahead forecast errors as a latent moving average process. Serial correlation exists for multi-step ahead forecast errors even if the forecast model is correctly specified (Harvey et al., 1997), hence our contribution is universal and not an artefact of model misspecification. The main methodological contribution of [**Paper B**] is the exposition of estimation biases that occur when state space models are parameterized for direct multi-step ahead forecasting but the serial correlation of the forecast errors is not accounted for. Our simulation studies present the estimation biases for time-varying coefficient autoregressive models, such as being used in state space form in Poncela et al. (2013). The results demonstrate that estimation biases can be substantial for long forecast lead times and strongly autocorrelated time series. In addition, we propose a flexible innovation process with time-varying coefficients and consider the Unscented Kalman filter (Julier and Uhlmann, 2004) for approximate state estimation of a consequently non-linear state space model.

The ambition towards the third and last research objective has been to extend [**Paper B**] into a hybrid forecasting model based on the Deep State Space Model (DSSM), which has been proposed in Rangapuram et al. (2018). The model uses a globally estimated recurrent neural network to predict time-varying parameters of local linear Gaussian state space models. Estimating the parameters of the recurrent neural network jointly is expected to extract features and learn complex temporal patterns from raw time series data, whereas incorporating structural assumptions through a state space model can alleviate overfitting (Rangapuram et al., 2018). In [**Paper B**], we propose to use the Unscented Kalman filter to estimate non-linear latent innovation processes when autocorrelation functions of the forecast errors vary over time. However, the remaining state space model remained linear and Gaussian. Our objective has been to adopt DSSM to predict time-varying moving average coefficients to linearize the non-linear innovation processes in [**Paper B**]. The premises has been to apply the computationally faster Kalman filter, while retaining the ability to model adaptive innovation processes. However, a closer look at DSSM lead us to the identification of significant limitations of this class of hybrid forecasting models.

To this end, [**Paper C**] explores these limitations on the basis of DSSM. Our first contribution is to clearly state the underlying assumption of DSSM, in that all series within the globally modeled time series panel must follow a stochastic process that is implied by the same parametric state space model. This is a limiting restriction for many real-world problems, where it is unlikely that all time series can be sufficiently well approximated by a single parsimonious state space model. Our second contribution in [**Paper C**] is the presentation of practical limitations as a result of using a recurrent neural network to predict the state space model parameters. Here we note that hyper-parameter tuning requires cross-validation, which limits the applicability of DSSM when time series are short. Moreover, we demonstrate that the employed state space models do not provide a universal safeguard against overfitting, since time-varying parameters are predicted. Based on the identified methodological and practical limitations, we briefly discuss the merits of hybrid models which first apply time series models locally and subsequently estimate the parameters of a global model on the residuals. Our last contribution in [**Paper C**]

is the verification of the experimental results in Rangapuram et al. (2018).

1.4 Thesis outline

The remainder of this thesis is organized as follows. In Chapter 2, we introduce important container shipping concepts to provide additional context for the methodological contributions of this thesis. On the basis of these preliminaries, Chapter 3 discusses the challenges of container inventory management at depots and ECR planning between depots. Our discussion provides the justification for decomposing the global ECR problem into sub-problems with various planning horizons. Chapters 4 summarizes the contribution towards the first research direction, and Chapter 5 presents our contributions towards the second and third research direction. Finally, we conclude the key findings of this thesis and discuss future research directions in Chapter 6. The scientific articles that this thesis is based on are provided as appendices.

1.5 List of publications

The following publications form the basis of this thesis:

[**Paper A**] B. Sommer, T. Krogh Boomsma, K. Kähler Holst, "Inland empty container inventory management with Markov decision processes", Submitted (first round of revision) to *Transportation Research Part E: Logistics and Transportation Review*, 2023.

[**Paper B**] B. Sommer, K. Kähler Holst, P. Pinson, "Direct multi-step ahead forecasting with state-space models", Submitted (first round of revision) to *International Journal of Forecasting*, 2022.

[**Paper C**] B. Sommer, K. Kähler Holst, P. Pinson, "A critical look at deep state space models for time series forecasting", Submitted to *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.

Container shipping preliminaries

In this chapter, we introduce essential elements and procedures in containerized shipping, while focusing specifically on the ECR problem. Nevertheless, for a more comprehensive understanding of various aspects within the extensive domain of container shipping, we recommend the following selected books. The historical development of container logistics from its infancy in the 1950s until the end of the 20th century is presented in Levinson (2006). Lee and Meng (2015) focus on technical topics of container shipping and modern supply chains, including optimizations methods. For an overview on port related topics we refer to Notteboom et al. (2021), for maritime economics to Cullinane (2010) and for the broader view on intermodal transportation systems to Rodrigue (2020). Last, we refer the interested reader to Song and Dong (2022) for a general introduction to empty equipment logistics and sequential decision-making models in the ECR context.

In Section 2.1 we introduce the essential components of a global container shipping network. The separation between ocean and inland networks, and their different transportation modes, subsequently highlights several challenges for ECR planning. A brief introduction to the lifecycle of a container shipment is presented in Section 2.2. The objective of the section is to illustrate the basic steps of a container shipment, from the initial booking request until the return of an empty container back to the shipping company. The presented processes form the basis of the covariate that we derive from customer booking data to forecast empty container returns in [**Paper B**].

2.1 The network of container logistics

Container vessels connect ports across the globe and are the foundation of containerized shipping. We begin our exposition of a modern intermodal container shipping network by introducing the key components of an ocean network. Ports are the entry point to the inland network, which we subsequently introduce.

2.1.1 Ocean network

Today, most major container shipping companies deploy their vessels in large hub and spoke networks (Song and Dong, 2015). The network structure originates from the concept of economy of scale, where the unit container transportation costs decrease with an increasing container vessel capacity. Larger container vessels (*mother vessels*) are used to connect strategically important ports (*hubs*) over long distances. These ports are often geographically central to a region and their infrastructure supports modern container vessels with capacities exceeding 20000 TEUs (*twenty-foot equivalent unit*). Hubs are connected to smaller ports (*spokes*) through container vessels (*feeders*) with significantly lower capacity than mother vessels. A hub and spoke network utilizes feeder vessels to move containers intra-regionally from spokes to hubs, where the containers are loaded onto mother vessels for their longer distance journey. The procedure

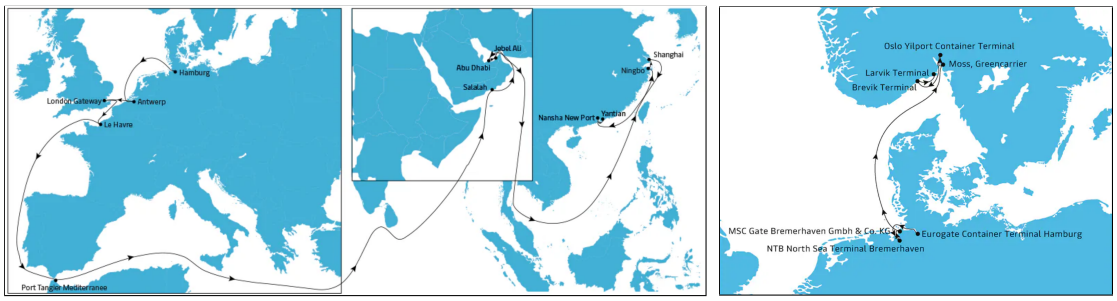

(a) Mother vessel service (Maersk AE7 Eastbound) (b) Feeder service (Maersk 49H)

Figure 2.1: Selected inter- and intra-continental services of Maersk (available as of June 2023).

of moving a container from one vessel to another is known as *transshipment*, a key feature of modern networks that connect hundreds of ports.

Each vessel in the network is deployed to a service route that serves a set of ports in a scheduled sequence. Multiple vessels, often with similar cargo capacities, are deployed to the same service route to increase the frequency for customers to use the service. This is in essence close to the operations of a public bus network, where it is also possible to redeploy buses from one service to another when demand patterns change. A real-world example of these concepts are presented in Figure 2.1 for two different Maersk services. Shipping a laden container from Larvik (Norway) to Yantian (China) may use a feeder service from Larvik to Hamburg and transship the container to a mother vessel which sails to Yantian. An important aspect of shipping networks is that sailing times of mother vessels are long. About seven weeks pass in the above example between the departure of the mother vessel in Hamburg and its arrival in Yantian.

2.1.2 Inland network

The network of a container shipping company expands further on land, with ports being the natural connection point between the ocean and inland network. Ports may generally also serve demands of other ships, such as bulk carriers or oil tankers. However, our exposition is naturally limited to ports with container terminals, which are facilities to transship containers from one to another transportation vehicle. Terminals equally exist at inland locations, such as rail and barge terminals, which we subsequently explore.

Container depots An essential component of the inland network are storage facilities, also known as container depots, to store empty containers when they are not used by customers. Indeed, shipping companies are the primary asset owners of containers and companies operate large container fleets to avoid container unavailabilities that lead to lost sales. Idling times when containers are not used are consequently long. The primary locations for storing empty containers are either ports, facilities close to ports or at strategic inland locations. All of these locations are also used by laden containers as temporary storage while they wait for a following transportation service to arrive during transshipments. Storage capacities at ports are often limited and costly, and thus shipping companies operate depots outside of ports. Shipping companies may not necessarily own a container depot, such as for depots inside ports, and instead use a third party depot. In this case, the shipping company pays storage and handling fees, such as gate-in and gate-out fees for delivering and picking up a container. Furthermore,

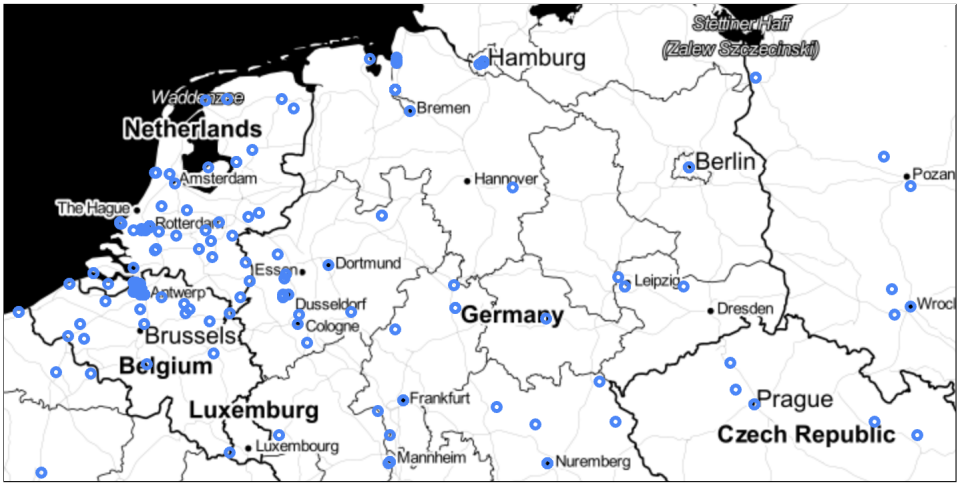

Figure 2.2: Location of Maersk container depots in parts of Western and Central Europe.

container depots are the locations where empty containers are cleaned and minor repairs are performed in between two consecutive shipments.

If circumstances allow, depots are located close to ports to reduce transportation times and costs, but also to enable more flexible operations since not all ports operate around the clock. Hence, a container can be delivered to a depot if it arrives by land outside the port operating hours. Inland container depots are operated by shipping companies to reduce the distances to their customers. This is essential to serve customer needs in a timely and cost-effective manner, which may not be possible when trucks cover large distances between customers and container depots. Figure 2.2 gives an account of Maersk's depot network in parts of Western and Central Europe, where a substantial amount of inland container depots are operated. The majority of depots in Germany are located in densely populated industrialized areas, distant from the country's largest ports in Bremerhaven and Hamburg. Many inland depots have intermodal terminals with transshipment capabilities for rail and barge, which allows using other transportation modes than truck to transport containers between inland depots and ports. Most basic depots have conversely no intermodal terminals, and thus only trucks can deliver and pick up containers.

Transportation modes Trucks are used to transport containers to and from customer sites. Exceptions where truck transportation is dispensable exist, for example, when cargo loading (*stuffing*) and unloading (*stripping*) of containers is carried out inside terminals. Container transportation is then carried out by train, barge or container vessel. However, we subsequently consider the common case where an exporter requests an empty container and performs stuffing at their location. A truck has then to pick up an empty container at a depot and deliver it to the customer site. Indeed, first and last mile deliveries generally rely on truck transportation, whereby one exception has been mentioned previously. Alternative transportation modes exist for inter-depot container movements when rail and barge infrastructures are available. In our context we use barges to refer to the general class of small vessels that transport containers on inland waterways, without further specifying whether a barge is self-propelled. Container shipping companies rarely own any inland transportation assets and must consequently negotiate transportation rates with different vendors. Contract specifications with vendors vary between

transportation modes and across countries, albeit we may generally assume that procuring capacity on the spot market for trucking services is more common than for trains and barges. Most contracts include volume-dependent transportation rates, such as procured capacity on train and barge services. Deviating from the procured volumes incurs premium costs and shall thus be avoided when planning container transports.

Trains and barges are attractive for cost optimization because their bulk carrying capabilities can reduce transportation costs between inland and port terminals. Transportation costs for larger container volumes over longer distances are additionally cheaper than truck because of labor costs. A caveat of trains and barges is a reduced flexibility compared to truck transportation because container transportation is bound to a fixed rail and inland waterway network. In addition, both modes generally follow a fixed schedule, whereas trucking services can be purchased from third-party vendors at short notice. It follows from these characteristics that shipping companies aim to benefit from the strengths of each transportation mode. That is, barges and trains transport larger container volumes over long distances, whereas trucks are dominantly used for first and last mile container transportation.

2.2 Lifecycle of a container shipment

The ambition of the following description is to provide a holistic view on container booking processes for the ECR context of this thesis. Thus, many details of the complex shipment processes are not presented. We refer the interested reader to Lee and Meng (2015) for additional information about the relations between customers and shipping companies. Let us consider the previous example of shipping cargo from Larvik (Norway) to Yantian (China) to illustrate the effect of laden container movements on ECR planning. To begin with, let us assume that a wholesaler in Yantian (import customer) wants to import frozen seafood from a producer in Norway (export customer). Both exporter and importer agreed on the number of required refrigerated containers, maximum shipping prices, and timelines for the containers to arrive in China. They proceed by contracting a supply chain management company who does not own any transportation assets (third party logistics), which we refer to as the *customer* in the following, to fulfill their request. The supply chain management company subsequently interacts with the shipping company.

Booking request Given the constraints for container quantities, prices and timelines, the customer contacts the shipping company several weeks before the planned cargo departure times to receive quotas for several shipping services. A service constitutes various aspects of the container journey, such as who is responsible for inland transportation of the container and which route the containers should preferably take. Both aspects are interlinked because a customer can contract a third-party trucking service (*merchant haulage*) to pick up an empty container at a depot and deliver a subsequently laden container to a port. In this example it may be more suitable for a customer to use its own trucking service to deliver a laden container to Hamburg than using the shipping company's feeder service from Larvik to Hamburg as shown in Figure 2.1. The responsibility for these inland services is with the shipping company when a customer chooses *carrier haulage*, which in essence is a more complete service. The customer informs the shipping company about the export and import locations, and desired arrival time of the empty container at the export site for stuffing. Based on this elementary information, the shipping company proceeds and organizes the inland container transportation, including which depot an empty container should be delivered from. For merchant haulage, the shipping company and customer agree during the booking process instead on the pick-up and return

times at specified container depots. The processes between merchant and carrier haulage are therefore different since merchant haulage bookings require a mutual agreement on the pick-up and return location of empty containers at the time of booking.

Quoted prices for all services will reflect the anticipated empty container availability at depots near the export location at the time of booking. Container unavailabilities close to the export location can lead to a rejection of a booking request. Unanticipated equipment unavailabilities can in the worst lead to cancellations of already accepted bookings, which can harm the reputation of the shipping company. An itinerary is created once a booking is created. It includes information about the assigned container depots and timings for empty container pick-ups and returns, and routing information for laden containers through the inland and ocean network of the shipping company.

Booking execution The booking execution process is initiated when the customer requests empty containers. There can be a significant time gap, ranging from several days to multiple weeks, between the acceptance of a booking request and the collection of an empty container from a depot. Once stuffing is performed at the export location, the laden container is transported to a terminal for its further transport. The terminal can either be an inland terminal, such as rail or barge, or a container terminal at a port, depending on the booking itinerary. Upon the arrival of the laden container at its destination terminal, a truck transports the container to the import customer site, where cargo is stripped. The empty container is then either picked up by the carrier haulier or delivered by the merchant haulier directly to the designated container depot. Upon arrival at the depot, containers are cleaned and small repairs are made, before the container is re-used for a new export shipment. Exceptions to this standard procedure exist when containers are *triangulated* by delivering the container immediately to an export customer instead of returning it to a depot (Furió et al., 2013). The benefits are that trucking distances can be reduced and depot fees, such as gate-in and gate-out fees, are avoided.

The duration between accepting a booking and the return of an empty container can extend to several months for shipments between Europe and Asia. During this time several factors influence the return time and location of empty containers. Delays are likely for bookings with multiple transshipments due to the likelihood of missing a scheduled connection. Similar to the airline industry, shipping companies also overbook their container vessels. Thus, empty container returns are delayed when containers must be "rolled" to another vessel. Other deviations from the expected return time can occur when customers request to keep the empty container, for example when used as additional storage, longer than agreed on at the time of booking. Similarly, customers with booked merchant haulage can request to return the empty container to a different depot when the new location helps the haulier to optimize its operations. From these dynamics it becomes apparent that utilizing booking data is beneficial for forecasting future empty container returns at depots because the data contains useful information about where and when containers are likely returned in the future. However, it is also evident that several pre-processing steps are necessary to account for variations of expected return time and location during the lifecycle of a booking.

Booking cancellation and amendment Additional uncertainty for empty container pick-ups and returns are induced by booking cancellations and amendments. Customers may place bookings for future departure dates of container vessels when prices are low, but cancel at a later stage when the booked services are not required anymore or better prices were offered by a competitor in the spot market. Moreover, customers can request amendments to their bookings, such as previously described for the return time and location of empty containers. Other amendments include requested changes for delivery times and locations of empty containers,

but also container type changes. For example, a customer can initially request one 40ft container and ask at a later stage for two 20ft containers instead. Requests to receive empty containers earlier than agreed during booking time can occur when customers require additional storage space because goods where manufactured earlier than expected. In conclusion, all booking adjustments affect the ECR planning processes in that they reduce the certainty at which empty container deliveries and returns occur.

Empty container repositioning

The introduction of the booking process illustrates the fundamental relevance of ECR for the profitability of a container shipping company. ECR overcomes at a global scale structural trade imbalances that cause empty container deficits in regions with greater laden container exports than imports. Repositioning empty containers ensures therefore the availability of containers in a deficit region and consequently the ability to generate revenue by exporting laden containers. Important examples in this regard are China and South-East Asia, since both areas have significant deficits of most container types due to greater laden export than import demands. Repositioning decisions must be made well in advance due to inherently long transportation times from European and North American to Asian ports. Unless a shipping company can lease empty containers on a short notice, repositioning too few containers can cause container unavailability and consequently lost sales when long transportation times are prohibitive to receive additional empty containers in a timely manner.

In this chapter, we explore the ECR problem in greater depth. We begin our detailed exploration by introducing the regional empty container repositioning problem. Our description of the problem in Section 3.1 highlights the importance of maintaining sufficiently large container inventory levels at inland container depots. Next, we discuss the inventory management problem of a container depot in Section 3.2, where we describe the problem as a sequential decision-making process under uncertainty. In Section 3.3, we explore the differences between global and regional repositioning due to different transportation times. Our exposition provides insights into empty container routing decisions for long-, medium- and short-term planning horizons. Last, Section 3.4 provides insights into the design of an ECR decision-making system, and highlights how the publications that were prepared over the course of this PhD project can contribute as individual components to a larger system.

3.1 Regional empty container repositioning

Similar to the global scale, regional ECR equally concerns the provision of empty containers to export customers and the revenue generated from laden export shipments. However, the planning processes differ due to the shorter transportation times within inland networks. The distribution of Maersk depots shown in Figure 2.2 and their relatively short distances from each other highlight an important aspect. Namely, equipment unavailability in a single depot does not necessarily result in lost sales, unless there is a container shortage affecting the larger geographical region. If a stock-out occurs at the depot closest to an export customer, the shipping company can arrange transportation from a farther depot to meet the customer's needs. However, it is worth considering that transporting empty containers over longer distances to an export customer site on short notice may incur additional costs, thereby reducing the profitability of the export shipment. In practice, the closest depot may not always be the preferred choice for a container shipping company to deliver an empty container to an export customer. Throughout this chapter, we will discuss how shipping companies can optimize their

operations by selecting alternative depots when it improves efficiency and reduces costs. By considering various factors, such as container availability and transportation costs, shipping companies strategically manage their container depots in conjunction with the regional inland and feeder transportation networks to meet customer demand.

Inland transportation rates for containers vary between regions and depend on contractual agreements between the shipping company and inland transportation vendors. The availability of transportation capacities also affects rates, which may be unfavorable for the shipping company due to a short notice. Shipping companies strategically utilize container depots in combination with the regional transportation network to ensure efficient supply of empty containers to exporters. Indeed, the essence of ECR is to resupply an empty container of a completed import shipment to an exporter at low costs. Regional repositioning fulfills therefore the purpose to increase the profitability for export shipments by transporting empty containers at low costs to inland depots, from where trucks can carry out the first and last mile transportation.

3.2 The inventory management problem

Container inventory management strategies play a crucial role in repositioning strategies as they assess whether current container stocks at depots, along with outstanding repositioning orders, can satisfy future customer needs. An inventory system can support the decision-making process to request more empty containers by assessing the future container availability at a depot. Projecting future container unavailability can inform decision makers to request more containers to be delivered to this depot, whereas a projection of future container surpluses is informative to reduce inventory levels presently. In what follows, we provide a comprehensive introduction to container inventory management at an inland depot in the context of the ECR problem. While real-world container depots store various types of containers, we will simplify our analysis by considering a depot that handles and manages standard 40ft steel containers exclusively. This allows us to focus on the core aspects of inventory management within the context of a container depot. Another important consideration in real-world operations is the condition of containers. Containers may vary in their condition, and damaged containers typically require repairs before they can be used for future shipments. However, for the purpose of our analysis, we will assume that all containers in the depot are in sound condition.

3.2.1 Empty container in- and outflows

Container in- and outflows occur at a depot due to various factors, as depicted in Figure 3.1. Empty container returns from import customers and deliveries to export customers are among the main factors that cause empty container stock variations. These processes are controllable for the coming days when customer bookings already exist. Within a set of constraints, the shipping company may request an import customer who just received a laden container to return the container earlier than is specified in the customer's booking documents. Similarly, the shipping company can select the container depot from where an empty container is delivered to an export customer. There will generally be one preferred depot from where a shipping company aims to deliver an empty container to an export customer. Vice versa, there is a preferred depot where to return an empty container of a completed import shipment. A more detailed description of a preferred depot is presented in the following and in Section 3.3, where we introduce the routing problems that shipping companies need to solve.

To ensure cost-effective container availability for export customers, empty containers are strategically repositioned to depots, allowing the preferred depot to offer them at lower costs.

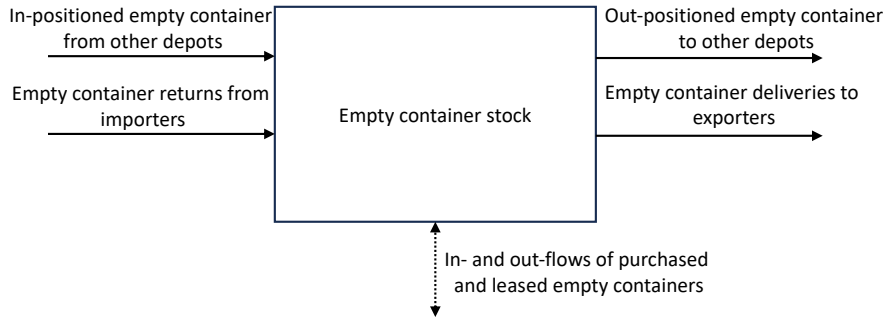

Figure 3.1: Schematic description of the empty container in- and out-flows of a container depot.

Inland depots must maintain an adequate stock of empty containers to minimize truck travel distances when delivering or collecting containers from inland customer sites. The further transport between inland locations and ports can benefit from the cost-effectiveness of trains and barges, if trucks transport containers to inland multi-modal terminals. Shipping companies reposition containers in order for depots to fulfill this task in periods when empty container deliveries exceed returns within a depot's region. Decisions to reposition containers from another depot need to be made in advance for two primary reasons. Firstly, there may not always be a nearby depot that can provide the container in a timely manner. Secondly, a short notice makes it generally more difficult to transport the container at low costs. Empty container delivery and return requests are unknown for future weeks since new bookings are accepted continuously. Bookings, and therefore also future empty container delivery and return requests, are largely uncontrollable because they represent the demand of import customers to receive goods from export customers. Empty container repositioning requests with lead times are therefore made under uncertainty.

3.2.2 Decision-making under uncertainty

The uncertainty of future events, such as unknown empty container demand and return but also delays of outstanding repositioning orders, exposes a decision maker of depots to the risks of requesting too many or too few containers. Having requested too few containers in the past can cause low inventory levels at a depot, which subsequently requires requesting more empty containers on a short notice to avoid stock outs. This can either accrue additional transportation costs to the shipping company or cause stock outs if the containers cannot be repositioned due to the short notice. Additional costs can also accrue if too many containers have been requested in the past and the estimated demand for the containers did not materialize. Thus, containers have arrived at a depot where they are not required anymore. Before containers idle at their current location, they may be repositioned to another depot. Both cases of ordering too many or too few containers illustrate the sequential dependence of current inventory levels on past repositioning decisions.

Figure 3.2: Empty container depots in the area around Frankfurt and Mannheim in Germany. The dashed line close to Rüsselsheim, Worms and Mannheim is the Rhine river. The Main river flows through the city of Frankfurt and enters the Rhine west of the city.

3.2.3 Inventory management system design

Finding the right balance between minimizing repositioning costs and reducing the likelihood of stock outs is challenging for a large network of depots. Decisions about decreasing or increasing stock levels must reflect several factors. Among them are forecasts about future empty container delivery and return imbalances, the current stock level and outstanding repositioning orders. It is further necessary to consider the connectivity of a depot to nearby depots, the remaining transportation network and costs of different inland transportation modes. Many trucking vendors charge a base rate for their services in combination with variable costs that dependent among other factors on driving distances, fuel consumption and total service times. Given short distances between depots and customers, the additionally accrued costs of using a depot for empty container pick-ups or returns that is not the closest to the customer are comparatively small. This has important implications for inventory management strategies in regions with high densities of container depots. To emphasize this, let us consider the depots in the region around Frankfurt and Mannheim in Southern Germany (cf. Figure 3.2). All depots are located at barge terminals, except for Frankfurt, where the depot is located inside a barge and rail terminal. The intermodal terminals are connected to the big European ports in Rotterdam, Antwerp and Hamburg. The driving distances by truck between the depots are short. Thus, all depots can serve customers within the shown region. The shipping company can design an inventory management system for the group of depots, which we refer to as a pool subsequently, on the basis of the previously introduced responsibilities of a depot and container transportation costs.

Pooling container depots

An inventory model for a pool of container depots can be used to make empty container inventory decisions with longer planning horizons. A longer view on future container demand and return imbalances is necessary to optimize repositioning costs when imbalances vary over time. Efficient repositioning strategies build up stock levels slowly by using cheaper transportation options in anticipation of seasonal variations, for example due to higher demand during harvesting periods. Preferred transportation modes do not necessarily have to be train or barge, but can also be trucks with attractive transportation rates during low demand periods. For the shown container depots in the region around Frankfurt and Mannheim in Figure 3.2, containers are mostly transported to and from the large ports in Rotterdam, Antwerp and Hamburg. The shipping company can adopt their repositioning planning for inland container pools in response to the potentially large uncertainty of future empty container demand and returns in the coming weeks. Inventory planning can allocate a request for additional empty containers to the con-

tainer pool for the upcoming weeks, without yet specifying to which depot the containers should be transported. Daily operational decisions will organize the transport of empty containers to each of the container depots while respecting the overall repositioning plan of the container pool.

There are two important benefits of using an inventory model at this aggregation level. Firstly, using the inventory model to plan ECR decisions becomes easier because repositioning decision-making for short and medium long planning horizons are decoupled from near-term operational decisions. One of the major tasks of a container shipping company is to route containers from an origin to their destination. Planning the transportation in inland networks is better known as container drayage, and concerns primarily optimal routing of trucks to satisfy laden and empty container deliveries and pick-ups. Daily route scheduling can use information about a higher level pool allocation plan to organize the transportation of empty containers accordingly. If an allocation plan foresees to reduce stock levels in the pool, empty containers are transported from import customers to intermodal terminals, from where containers are transported by rail or barge to a larger port. Daily decisions can also be used to short-term balance inventory levels of nearby container depots, as in the case of the depots around Frankfurt and Mannheim in Figure 3.2. An important aspect is that operational decisions are executed often immediately, whereas allocation plans are adjusted when new information becomes available.

Secondly, the task of forecasting empty container demand and return imbalances simplifies. Forecasts are essential for empty container allocation and repositioning plans because empty container demand and return imbalances can vary over time. Naturally, the objective is to forecast demand and return imbalances in regions where a shipping company does business, such that the company can reposition empty containers to prevent future shortages. The problem that arises for ECR planning is how to form the responsibility regions in which depots are responsible to deliver empty containers to export customers. Forming a responsibility region for a container pool has several benefits. Generally, fewer forecasts need to be produced because pools are groups of container depots. Aggregating the demand and returns of multiple depots has the additional benefit that historical time series, which are the foundation of data-driven forecasting models, become less erratic. The modeling implications and additional time series properties are discussed in greater details in Section 5.1, which are the foundation for our forecasting model in [**Paper B**]. These modeling and computational aspects can generally be addressed accordingly. However, a significant challenge is imposed by the necessity to forecast unconstrained demand for empty container inventory decision-making.

Forecasting unconstrained demand

Unconstrained demand refers to the demand of export customers that would materialize if enough containers are available. Observing this "true" demand is infeasible in practice because bookings that are cancelled by customers induce a censoring effect. Measuring the explicit reason why a customer decided to cancel a booking is one necessary component to derive a historical unconstrained demand time series for a region. Several challenges arise, with one being that a booking may be cancelled because a customer decides to ship at lower costs with a competitor. It is unclear whether this booking accounts as demand.

The procedure that is followed for the inventory model in [**Paper A**] is to use historical pick-ups and returns as proxies for the true unconstrained quantities. Evidently, this also induces a censoring effect because no demand is observed historically when bookings were rejected as a consequence of container shortages. Nevertheless, using historical deliveries and returns avoids the treatment of complicated censoring effects for historical bookings. A caveat is that Maersk only provided historical deliveries and returns for container depots and not for geographical

regions. A challenge that arises is that pick-ups and returns at a single depot can be controlled during operational decision-making as we have previously discussed. Thus, feedback-loops exists between forecasts, past operational routing decisions and inventory management decisions at the granularity of a single depot.

These effects can be exemplified by considering the delivery of an empty container to an export customer when the closest depot has no equipment available. Carrying out the delivery from a close-by depot subsequently assigns the demand proxy incorrectly to a depot that is not the closest to the export customer. Choosing another than the closest depot can also happen due to route optimization of trucks, hence the dependence on past routing decisions. The feedback-loops can lead to the over- or under-estimation of the true demand of customers who are closest to a depot, which will in subsequent steps affect inventory management decisions. Producing forecasts for a pool of close-by depots mitigates these effects to a certain extent when several pools are sufficiently far apart. This is based on the rational that the majority of customers within an area are served by the same pool, albeit the previous feedback effects may still persist for customers who are located in between pools. Forecasting demand for pools does evidently not circumvent the challenges that are induced by the censoring effect when no containers were available, although export customers were willing to shop freight.

3.3 The empty container routing problem

We have outlined the properties of an inventory model for managing empty container stock levels at inland container depots. Based on the introduced design, decisions are made to adjust stock levels for a group of container depots, a so-called container pool. In the following, we will outline how an inventory model for inland pools can be integrated into the ECR planning system of a globally operating shipping company. We analyze the planning horizons of varying lengths in relation to the distances that empty containers are planned to cover as part of a repositioning decision. This examination encompasses long-term, medium-term, and short-term planning perspectives. The combination of all planning horizons gives insights into the design of a global data-driven ECR planning and decision-making system.

Long-term It can take up to two months to transport empty containers from surplus regions in Europe to container deficit regions in Asia, where empty container deliveries exceed returns. It is essential that the planning horizon reflects these long transportation times to ensure that ECR plans can be established between far apart regions. The planning horizon may consequently extend up to three months, which induces several challenges. First and foremost, uncertain future empty container demand and returns limit the accuracy of detailed repositioning plans between depots. Consider that it takes about seven weeks for a vessel on Maersk service AE7 Eastbound (cf. Figure 2.1) to sail from Hamburg to Yantian. This strictly means that a repositioning decision for loading empty containers on a vessel in Hamburg must be taken at least seven weeks before the scheduled arrival in Yantian. However, more information about future container demand and return imbalances in Yantian can be obtained after the vessel departure in Hamburg. Our forecasting method that we develop in [**Paper B**] can aid the decision process for repositioning empty containers over long distances because it addresses the challenge of forecasting empty container demand and returns multiple weeks ahead in the future. When the vessel is closer to its arrival in Yantian, it may occur that the requested empty containers are not required anymore because revised forecasts become available. This highlights the second important challenge of long-term ECR planning. That is, planning processes are sequential and new information arrives dynamically.

Information about vessel arrival times, container availability in depots, bookings, inland transportation rates and vessel capacities all evolve dynamically over time. Therefore, new information becomes continuously available that allows to revise repositioning orders of moving empty containers. Moreover, current repositioning decisions have an effect on future decisions. The decision to load empty containers on a vessel takes up capacity and thus limits the vessel's ability to carry additional containers in the future. Transporting laden containers are prioritized over empty containers in practice. As a consequence, empty containers can be unloaded at a different than their scheduled destination. These dynamic aspects require the planning model to revise all active repositioning orders, otherwise containers will arrive in locations where they are not needed anymore.

A third and last challenge that we would like to highlight is the granularity of the planning model. Oversea repositioning naturally requires container vessels. Therefore, long-term planning has to reflect the vessel departure schedules at ports, hence the ocean transportation network. To which extent the inland transportation network should be represented in long-term planning is not clear. Trucks for example follow a flexible schedule, whereas trains and barges generally follow a fixed schedule. Moreover, shipping companies can control variable costs associated to owning and operating container vessels, which is contrary not the case for transportation modes in inland networks. Since repositioning decisions are affected by prices, which can vary in inland networks between vendors and transportation modes, it may be difficult to reflect this information accurately in a long-term planning model. Overall, the planning model must account for various container types, which all share the same capacities on vessels, barges and trains.

Medium-term The challenges in long-term oversea planning can motivate additional planning processes with shorter horizons to optimize empty container flows in smaller geographical regions. The design of a modern shipping network as a hub and spoke network lends itself naturally to a decomposition with long- and medium-term planning horizons. A shipping company can follow a strategy of long-term planning to transport empty containers to strategically important hubs. Medium-term ECR planning further distributes empty containers from a hub to a region with feeder services and inland transportation services. The shorter planning horizon allows representing an inland network more accurately in terms of schedules, capacities and freight rates, which may need to be forecasted when inland transportation capacities are only procured on the spot market. Planning horizons may extend to several weeks to cover the previous aspects of ECR between ports and inland pools. Our proposed empty container allocation model in [**Paper A**] can be integrated into such a regional ECR planning model, since it generates requests for reducing or increasing the inventory for an inland depot.

Decomposing this planning problem between regions is a challenging problem for a global shipping company. A natural direction to follow is to formulate repositioning models for each continent. However, feeder services may exist that connect two continents, such as parts of Southern Europe and Northern Africa. A second challenge is related to the scale of the repositioning problem and concerns the planning granularity. Repositioning plans are required for all container types, which share capacities on feeder vessels, trains and barges. Moreover, it is unlikely that a problem can be solved at a global scale for thousands of inland depots. Thus, the introduced aggregation strategy may be applied to reposition empty containers between ports and container pools. The planning problem remains dynamic and sequential due to the traveling times of several weeks for feeder vessels, which must be addressed accordingly. That is, new information becomes continuously available once a container is loaded on a feeder vessel, which may require to revise the repositioning decision.

Short-term Operational, or short-term, planning primarily concerns empty and laden container transportation between terminals, depots and customers. The routing task in inland networks is also known as container drayage. Two tasks need to be fulfilled jointly. The first task is the delivery of laden containers from terminals to import customer sites, and the subsequent pick-up of an empty container once freight has been removed. The second task is the delivery of empty containers to export customers and the subsequent transport of laden containers to terminals. Drayage operations often concern the management of a fleet of trucks that is tasked to minimize the transportation costs for picking up and delivering containers at fixed origins and destinations. Optimal routing decisions are difficult to obtain for real-world problems because the problem has to be formulated as an integer problem. Thus, a large combinatorial optimization has to be solved when many orders are handled simultaneously, which often relies on heuristics. In fact, drayage operations can utilize container triangulation (Furió et al., 2013) to minimize the transportation costs of empty containers to export customers. A challenge is that the combinatorial problem for drayage planning becomes more complicated to solve. Optimizing container triangulations creates additional possible routes for trucks because the return location of empty container is not determined beforehand.

Drayage optimization may additionally address the requirement to balance empty container inventories of nearby container depots. Formulating a drayage optimization problem with triangulation and empty container inventory balancing faces the following challenge. On one hand, the identification of triangulation to satisfy customer needs for receiving laden and empty containers can be limited to a single day. On the other hand, optimizing empty container inventories requires considering future information as we have discussed in detail. Thus, triangulation is a static problem that can be solved every day independently of the next day, whereas inventory decisions are sequential. Optimizing both decisions in a single optimization problem is challenging due to the inherently large computational complexity, thus unlikely to be practical for real-world problems. As a consequence, inter-depot balancing of empty containers is planned independently by determining the return location of empty containers that cannot be triangulated a priori. Thus, container drayage operations aim to fulfill a regional ECR plan.

3.4 An ECR decision-making system

Based on the described inventory management and container routing challenges, a shipping company can design an ECR decision-making system that follows a top-down approach. Long-term planning ensures that sufficient amounts of empty containers are available in strategically located hubs. Planning decisions with horizons up to several weeks distribute empty containers from hubs to inland depots and smaller regional ports. The distribution of containers to inland depots ensures that daily drayage operations can be carried out without being subject to container unavailabilities that interrupt operations, thus accrue additional costs. Each of the planning steps addresses the inherently long lead times of transporting empty containers over long distances, and therefore the related uncertainties due to unknown future empty container demand and returns. Maintaining sufficient empty container stock levels in strategic locations is essential in this top-down approach. It allows postponing the transport of empty containers to inland depots until more information about future container imbalances become available. Although the planning approach is considered to be top-down, decisions are executed bottom-up. That is, drayage operations are made to fulfill regional ECR plans, and consequently global long-term ECR plans. If a regional ECR plan aims to reduce empty container stocks, then drayage operations are tasked to deliver empty containers to ports, from where they can be further repositioned.

Decomposing the global empty container inventory and repositioning problem is difficult because all planning levels are interlinked with each other. Finding appropriate separations between different temporal and spatial planning levels of the ECR problem is largely influenced by the inventory and routing aspects that we have discussed throughout this chapter. Any decomposition of the global ECR problem will evidently have caveats, which eventually lead to short-comings of an ECR decision-making system. Nevertheless, the planning problem has to be decomposed in some shape or form because formulating and solving a single mathematical model to optimize empty container inventory and routing decisions jointly is infeasible. The optimization model would need to be a sequential and stochastic decision-making model for balancing inventories of multiple container types between thousands of depots. Optimizing drayage operations would require accurate representations of customer locations, whereas planning the transportation of empty containers overseas requires to represent the large-scale ocean network and the inherent uncertainties due to delays. Essentially, a single ECR planning model would require to represent the global operations of a shipping company in a single mathematical model, which evidently is infeasible.

Depending on the specific design of a shipping company's global ECR decision-making system, several forecasting, routing and inventory models are required. The empty container inventory model in [**Paper A**] can be used by a shipping company to support their inland depot operations. In [**Paper B**], we develop a forecasting model that addresses many of the real-world challenges when forecasting empty container demand and returns in the upcoming weeks. The same methodology can be extended to forecast empty container demand and returns for larger geographical regions and longer horizons, thus can benefit a shipping company to make long-term repositioning decisions. The forecasting methodology in [**Paper C**] has likely no immediate benefit for ECR planning. None of the major forecasting problems that a shipping company faces during ECR planning fulfills the requirements of the hybrid forecasting methodology of [**Paper C**].

Sequential decision-making under uncertainty for container depots

The previous chapter explored challenges for empty container inventory management and routing. Throughout this chapter, we focus our attention on the inventory management problem of an inland container depot. The assessment whether current container stock levels and outstanding empty container repositioning orders are sufficient to satisfy future container demand builds the cornerstone of the inventory problem. Decisions to receive empty containers from other depots are made under uncertainty due to unknown future demand and must be made in foresight due to transportation lead times. Inventory management decisions have been described to have both immediate and long-term consequences. On one hand, ordering too few containers can lead to future container shortages and additional costs if repositioning additional containers on a short notice requires to pay premium costs. On the other hand, ordering too many containers now can cause containers to idle in locations where they are not needed if demand has been overestimated. Therefore, present decisions will affect the future ability of the depot to provide empty containers to exporters at low costs.

In this chapter, we present our contributions towards the first research direction; inventory control policies for container depots. Section 4.1 introduces the key elements for describing the sequential decision process for managing the inventory of a container depot. We provide a brief introduction to Markov Decision Processes (MDPs) and backward dynamic programming in Section 4.2. Our contributions of [**Paper A**] are summarized in Section 4.3.

4.1 The sequential decision process of a container depot

A depot manager, which in our case is assumed to be the decision maker, faces the task of controlling the inventory of an inland container depot. The decision maker can exchange empty containers with ports and nearby inland depots to balance the inventory. The costs and lead times for repositioning containers depend on the transportation mode and container availability at other depots. Costs accrue to the decision maker for receiving and sending empty containers, as well for operating the depot where storage costs accrue. The main responsibility of the decision maker is to provide empty containers to exporters and accept container returns from importers who unloaded their freight. Future demand of exporters and returns of importers are unknown to the decision maker. A decision to in-position empty containers ensures that the inventory can satisfy the demand of exporters in periods when demand exceeds the returns from importers. Containers are conversely out-positioned to reduce the inventory, hence storage costs, when container returns exceed demand. Both decisions are made daily at a fixed time point, and no additional decision to either receive or send containers can be made before the next day. The depot operates all year round and no definite date determines the end of this planning process. The objective of the decision maker is to minimize the long-run operating costs of the depot.

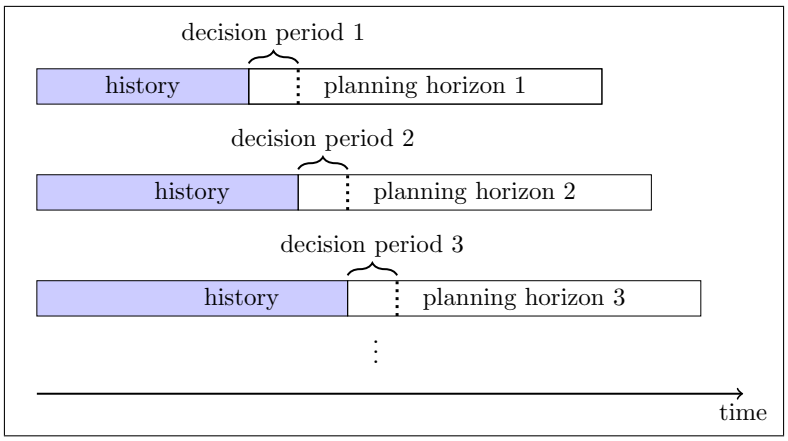

Figure 4.1: Schematic illustration of a sequential planning process. Decisions are made at the beginning of each decision period, which spans the time when additional information becomes available to the decision maker before another decision is made.

Figure 4.1 illustrates the sequential planning process of the depot operator schematically over time. An important aspect of the decision process is that the system evolves dynamically over time, which is indicated by the growing history in Figure 4.1. At this point, let the system refer in the broadest sense to a set of depots, the transportation network and customers. Today's decision to in- or out-position containers are chosen at the beginning of a decision period. Once a decision is made, time passes and the system evolves. When it is time to make a new decision during the next day to in- or out-position containers, the decision maker finds the inventory of the depot and the rest of the system in a new state. The new state can provide the decision maker with additional information about future container demand and returns, updated transportation rates as well as inventory levels at other container depots. A characteristic of sequential decision problems is that decisions have immediate and long-term effects on the system. The decision maker considers for this reason a planning horizon to assess the long-term impact of a decision, thus aims to assess whether the current decision allows to minimize future depot operating costs. For instance, ordering no container now reduces immediate costs because no costs accrue for receiving empty containers. However, long-term costs are affected when container shortages occur and costly last minute orders must be placed to receive additional containers.

Throughout the remainder of this section, we introduce the common terminology to formulate mathematical models for sequential decision processes. The mathematical formalism enables the estimation of decisions for the depot operator in complex systems. We use *decision epochs* to refer to the fixed time points when in- and out-positioning decisions are made at the beginning of every *decision period*.

States At each decision epoch, the decision maker observes the state of the system. In the broadest sense, the state provides the information basis for making decisions. The inventory level of the depot is one of many states that describe the system, as it determines how many containers can be out-positioned. Outstanding in-positioning orders provide additional information to the decision maker. A current ordering decision is naturally affected by the knowledge that more containers will arrive during the next days. Similarly, states that describe the broader system can improve a decision maker's ability to take present decisions for minimizing the depot operating costs in the long-run. For instance, a state can provide additional information about future

empty container demand and returns.

Actions The decision maker takes a decision, or action, in every epoch to out- or in-position empty containers once the system state is observed. Actions are selected from a finite set, as only a countable number of containers can be out- and in-positioned. Moreover, the present state determines the available actions. For the described inventory system, it is not possible to out-position more containers than are presently available at the depot.

Exogenous processes The transition of a state from the present to the next epoch depends commonly on two factors. First, an action will affect the states that we aim to control. That is, the inventory of the container depot, but also states of outstanding empty container orders. Second, exogenous processes can affect states. Satisfying the demand of exporters will decrease the inventory, whereas picking up empty containers from importers will conversely increase the inventory. Both quantities are stochastic and unknown for the current as well future periods when the decision maker takes a decision at the current epoch. In addition, empty container demand and returns are to a large extent uncontrollable, hence exogenous. Moreover, other exogenous processes can be responsible for the delayed arrival of outstanding repositioning orders.

Transition functions Stochastic exogenous information is only revealed for the present decision period once the decision maker has chosen an action. Transition functions are mathematical formulations that describe the evolution from the present to the next epoch as result of a selected action and realized exogenous variables. For instance, the inventory of the depot changes as a consequence of the number of out- and in-positioned containers as well as empty container deliveries to exporters and returns from importers. Modeling a delay of outstanding repositioning orders may not require a detailed physical description of the transportation process. Instead, a probabilistic model can describe the likelihood of orders being delayed.

4.2 Markov Decision Processes

The former elements are basic components of mathematical formulations that describe sequential decision processes. An objective function that specifies how present decisions should be made with respect to an uncertain future is missing. Objective functions include three important elements. First, there are *contribution* or cost functions. Costs accrue every decision period as a consequence of a decision maker's action and realization of exogenous information. Costs for depot operations are mostly associated with in- and out-positioning decisions, storage costs but also accrued costs for unsatisfied exporter demand as a consequence of container shortage. Second, an objective function includes a planning horizon (cf. Figure 4.1). Short planning horizons can lead to myopic behavior if the impact of a decision to minimize the depot operating costs in the future is underestimated. Conversely, the planning problem is becomes complex when the impact of present decisions on the evolution of the system state must be evaluated for longer horizons. The third component of an objective function specifies a criterion to account for the uncertain system transition over future periods. The criterion determines how present decisions should be taken with respect to their downstream impact for minimizing future costs. Most objective functions consider expected value problems, where a present action is selected to minimize future costs over the planning horizon in expectation. However, other risk-averse objective functions also exist.

4.2.1 Motivating Markov Decision Processes

Optimizing the objective function is challenging because decisions at the beginning of the planning horizon affect the downstream ability of later decisions to minimize costs. Finding optimal solutions with respect to the objective function requires, as a consequence of the sequential dependencies, to evaluate sequences of decisions over the whole planning horizon. The objective function value represents therefore the value of a decision at the beginning of the planning horizon to reduce future costs. A major challenge is that the number of decision variables grows exponentially with the planning horizon length. Several algorithmic strategies evolved in various communities of stochastic optimization to address this curse of dimensionality. We refer the interested reader for a review of the extensive field to Powell (2019).

One common approach for long planning horizons is to formulate the sequential problem as a Markov Decision Process (MDP). A decision process is Markovian when the state transitions and the accrued costs depend on the past of the system and the decision maker's decisions only through the current state and selected action. The benefit of imposing the Markovian assumption is that the objective function of expected value problems can be solved recursively. Instead of estimating the decisions for all periods simultaneously, it can be shown that the optimal decision sequences can be found by solving a sub-problem for each decision period. Thus, our aim is in the following to describe the depot operations as a MDP to estimate theoretically optimal decisions.

However, in our case decisions are only theoretically optimal because a MDP is only an approximation of a complex real-world system. It is therefore important to understand the differences between the true decision processes at a container depot and a specified MDP. Real-world depot operations are complex and repositioning decisions are subject to the continuously evolving information set of a non-stationary environment. In particular, any decision in Figure 4.1 relies on forward-looking information throughout the whole planning horizon. Among other factors, this includes information about the evolution of transportation rates and capacities. In addition, formulating a MDP for the container inventory management problem requires assumptions about the empty container demand and return behavior of customers. In Chapter 5, we discuss that these processes are frequently non-stationary. That is, the parameters that describe the stochastic processes vary over time. A common approximation for non-stationary processes is to assume them to be locally stationary, where parameters are constant for shorter time windows. We follow this approach and assume empty container demand and return processes to be stationary throughout the planning horizon in [**Paper A**]. Once the parameters that describe the stochastic process change, the optimal decision sequence is re-optimized.

4.2.2 Finite-horizon and discrete Markov Decison Processes

Our brief introduction to MDPs is based on the text book of Puterman (1994), where all proofs of the following standard results are found. We present a finite-horizon MDP in discrete time for greater generality without considering for now an application to the depot operation problem. Thus, let $\mathcal{T} = \{0, 1, 2, \ldots, T\}$ denotes the set of decision epochs, with $T < \infty$ being the finite planning horizon. At each decision epoch $t \in \mathcal{T}$, the decision maker observes a discrete-valued state s from a finite state space $\mathcal{S}$. Upon observing the state, the decision maker chooses an action a from the set of possible actions in state s, $\mathcal{A}_s$. The set of all actions $\mathcal{A} = \cup_{s \in \mathcal{S}} \mathcal{A}_s$ constitutes the action space of the MDP. Actions are discrete-valued, hence $\mathcal{A}$ is a finite discrete space.

Once an action a is chosen in state s and epoch t, the system evolves into a new state s' in epoch $t + 1$ and the decision maker receives a reward. Note that negative rewards are equal to costs. The transition under action a from the current state s to a state s' in the next epoch follows

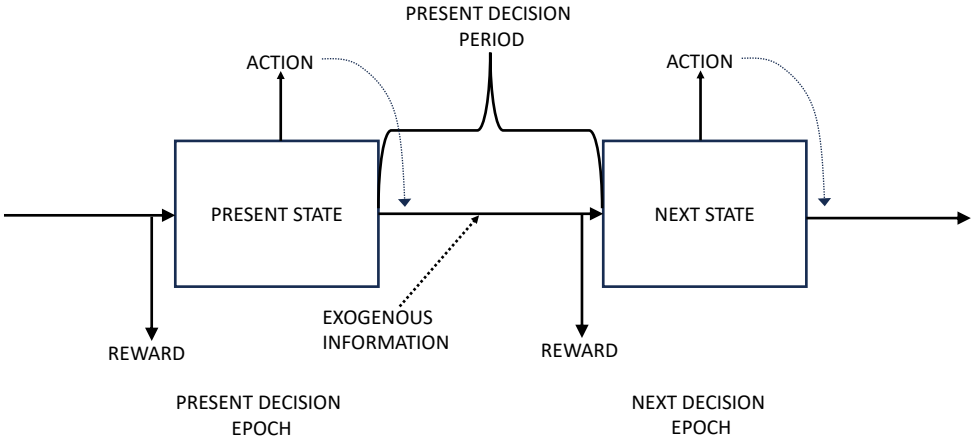

Figure 4.2: Symbolic description of a sequential decision process in discrete time. The illustration has been adopted from Puterman (1994). The dashed lines indicate that actions and exogenous information affect the transition of the state from the present to the next decision epoch.

the conditional probability $P_t(s'|s,a)$, with $\sum_{s' \in \mathcal{S}} P_t(s'|s,a) = 1$. The transition probabilities depend on the state transition functions, and may therefore depend on exogenous processes. However, we omit for now the explicit dependence of state transitions on exogenous processes to simplify our exposition. We further let $r_t(s,a,s')$ denote the real-valued reward function that assigns a contribution to the decision maker's objective in period t as a consequence of the transition from s to s' under action a. The reward functions are required to be known before choosing an action. In addition, we require that the collected reward in period t is not effected by future actions. Figure 4.2 provides a schematic illustration of this sequential decision process.

Due to uncertainty, decisions may be taken in expectation of the current period's reward

$$r_t(s,a) = \mathrm{E}(r_t(S,A,S') \mid S = s, A = a) = \sum_{s' \in \mathcal{S}} r_t(s,a,s') P_t(s'|s,a)$$

to account for the stochastic transition from state s to future states under action a. Throughout the remainder of this chapter, we follow the convention that lower case characters (s) denote realizations of stochastic variables (S). The essence of a sequential decision problem is that the ability to generate an immediate reward has to be weighted against the possibility to generate future reward when the system evolves into a new state. As we will show in the following, the reward functions and transition probabilities contain all necessary information to make optimal decisions when maximizing the sum of total expected reward over the finite planning horizon T. Finally, a MDP is defined by the tuple

$$\{\mathcal{T}, \mathcal{S}, \mathcal{A}, P_t(\cdot|s,a), r_t(s,a,s')\}.$$

4.2.3 Expected total reward problems

Let $s_0 \in \mathcal{S}$ denote the observed state of the system at the beginning of the planning horizon, $t = 0$. The decision maker is tasked to find the sequence of decision functions $d_t : \mathcal{S} \to \mathcal{A}_s$ that

maximizes the expected total reward

$$V_0^\pi(s_0) = \mathrm{E}^\pi \left[\sum_{t=0}^{T-1} r_t(S_t, d_t(S_t), S_{t+1}) + r_T(S_T) | S_0 = s_0 \right] \tag{4.1}$$

over the planning horizon T. We use upper case S_t to express the stochasticity of states at future epochs $t = 1, \ldots, N$. The transition between states is described by their transition probability functions. The expectation is with respect to the distribution that is induced by the policy $\pi = (d_0, d_1, \ldots, d_{T-1}) \in \Pi$, which denotes a sequence of decision functions. In addition, we assume that no decisions are made at the last decision epoch T. Instead, the decision maker receives a terminal or salvage value $r_T(s_T)$ for being in a terminal state s_T. $V_0^\pi(s_0)$ is known as the value function and determines the expected reward when following the policy π from the initial state s_0 at $t = 0$ onwards. An optimal policy $\pi^\star$ is obtained by solving the maximization problem

$$\pi^\star = \underset{\pi \in \Pi}{\operatorname{argmax}} V_0^\pi(s_0). \tag{4.2}$$

Maximizing $V_0^\pi(s_0)$ directly requires to explicitly obtain the joint distribution of states that is induced by a policy. Moreover, optimizing the policy for all stages simultaneously becomes computationally intractable when T is large due to the exponential growth of the number of decision variables. Both challenges can be avoided by first recognizing that Equation (4.1) can be written as

$$V_t^\pi(s_t) = \mathrm{E}^\pi \left[r_t(S_t, d_t(S_t), S_{t+1}) + V_{t+1}(S_{t+1}) | S_t = s_t \right], \tag{4.3}$$

which is due to the Markovian assumption for the state transition and expected rewards. The expression decomposes the value of being in period t in state s_t into the expected immediate reward and future rewards when following the policy π in the current and future decision epochs. The proof for Equation (4.3) uses the linearity of the expectation operator. Thus, other objective functions that may use risk measures do not necessarily follow the same recursive structure. This prohibits the application of the following algorithm to estimate optimal policies.

4.2.4 Backwards dynamic programming

The inductive scheme of the value function lends itself to solve the objective function (4.2) recursively. The standard result finds an optimal policy $\pi^\star$ to satisfy

$$V_t^{\pi^\star}(s_t) = \max_{a_t \in \mathcal{A}_{s_t}} \mathrm{E}\Big[r_t(S_t, a_t, S_{t+1}) + V_{t+1}^{\pi^\star}(S_{t+1}) | S_t = s_t \Big],$$

where $\mathcal{A}_{s_t}$ is the space of allowable actions in state s_t. The expression is known as the Bellman equation and provides the foundation of dynamic programming. Instead of maximizing the expected total reward in Equation (4.1) over all T periods directly, we subsequently solve $T - 1$ one-period maximization problems. Under the assumption that no decisions are made in the final epoch T, we obtain

$$V_T(s_T) = r_T(s_T)$$

as the value for being in state s_T in the final epoch. We then step backwards in time and solve the Bellman equation

$$\begin{aligned} V_t(s_t) &= \max_{a_t \in \mathcal{A}_{s_t}} \mathrm{E}[r_t(S_t, a_t, S_{t+1}) + V_{t+1}(S_{t+1}) | S_t = s_t] \\ &= \max_{a_t \in \mathcal{A}_{s_t}} \sum_{s' \in \mathcal{S}} \left(r_t(s_t, a_t, s') + V_{t+1}(s') \right) p_t(s' | s_t, a_t) \Big), \end{aligned}$$

since the value of being in state s' in epoch $t+1$ is known from the previous step. Obtaining the transition probability functions $P_t(\cdot|s_t, a_t)$ is one of the difficulties when formulating a MDP in many practical applications, as we illustrate for a stochastic inventory problem in the following. The optimal policy is tabulated since any decision function $d_t^\star(s_t)$ in $\pi^\star$ returns the optimal decision $a_t^\star$ for state s_t in epoch t. The described algorithm is better known as backwards dynamic programming and suffers from a curse of dimensionality. It is important to note that commonly we are only interested to find an optimal decision for the beginning of the planning horizon, that is $d_0^\star(s_0)$.

4.2.5 Curse of dimensionality

Backward dynamic programming finds a solution to the Bellman equation during each epoch for every state in the state space $\mathcal{S}$. The algorithm scales linearly in the number of decision epochs T, and is therefore attractive for solving planning problems with long horizons. However, dynamic programming suffers from the curse of dimensionality for large state spaces. Powell (2011) considers the inventory management problem for M different products to illustrate the exponential growth of the state space. If each product ranges from 0 to $P-1$ items, then the state space will have P^M unique states, where a single state is a M-dimensional vector of an unique inventory position. Backwards dynamic programming becomes computationally unattractive for large state spaces, since a solution to the Bellman equation is required for each state. A related curse of dimensionality exists for large action spaces because an enumeration of all actions is required to find the maximum to the right-hand side of the Bellman equation. Reinforcement learning, also known as approximate dynamic programming, addresses these curses of dimensionality with algorithmic strategies to learn policies for high-dimensional state and action spaces. The application of relevant techniques is out of the scope of this thesis and is considered to be future work. We refer the reader to Powell (2011) for an introduction of approximate dynamic programming.

4.3 Optimal control for container inventory management

The curses of dimensionality represent a significant challenge when it comes to estimating optimal decisions for sequential processes. In the following, we use common assumptions for container depot operations that allow us to retain a compact state space. We apply backward dynamic programming to estimate theoretically optimal inventory management policies, which can serve as a starting point for developing approximate dynamic programming algorithms in the future. In addition, using the exact algorithm provides a better understanding about the dynamics of the inventory system under various specifications. The remainder of this section is our contribution towards the first research direction and summarizes [**Paper A**].

4.3.1 Multiple supplier selection problem

Let us consider that the inland depot is exclusively responsible for satisfying the demand of exporters inside a wider region. Trucks are used to deliver empty containers from the depot to customer sites. The depot manager is responsible for maintaining sufficient inventory levels to provide this service to exporters, and therefore repositions empty containers with a port before shortages occur. Containers are transported between the port and inland depot with multiple modes of transportation. Each mode has different delivery times and costs per container. Figure 4.3 summarizes the description of this conceptual inland network. Based on our description of the inventory management problem for depots in Section 3.2, the following

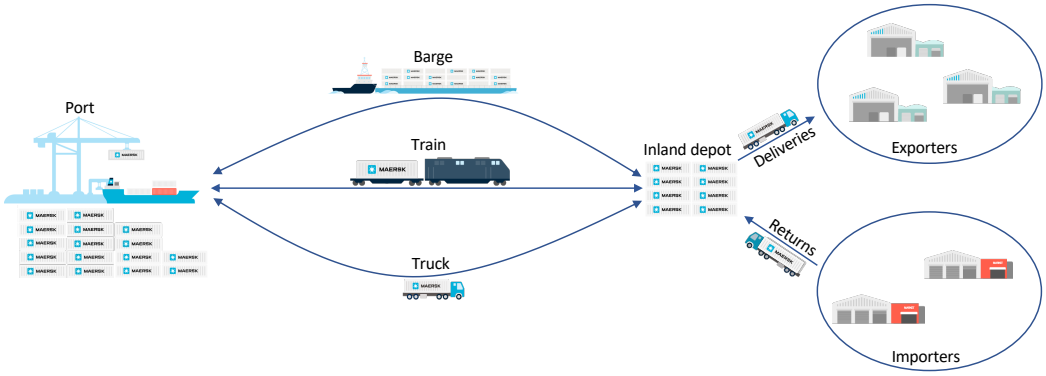

Figure 4.3: Conceptualized transportation network with a single port and inland depot. Taken from **[Paper A]**.

exposition extends also to a group of nearby depots. For instances, we may use the MDP to control the joint inventory of depots in the Frankfurt-Mannheim region in Figure 3.2, as most depots are connected with barge and train services to the port of Rotterdam.

Actions We consider that the inventory consists of a single container type. The decision maker uses out- and in-positioning decisions control the inventory, where all decisions are made once every day for a finite planning horizon T. Let $A_t^- \in \mathcal{A}^- = \{0, 1 \dots, \bar{A}^-\}$ denote the decision to out-position empty containers from the depot in decision epoch t. We follow the notational convention of **[Paper A]** and use superscripts $-$ and $+$ to denote empty container out- and inflows, respectively. In addition, we use the bar notation to denote capacities. Thus, $\bar{A}^-$ denotes the available capacity for out-positioning empty containers to the port. Empty containers can be in-positioned with lead times $l = 0, 1, \dots, L$ to represent different modes of transportation, with L representing the longest transportation time. Let $A_{l,t}^+ \in \mathcal{A}_l^+ = \{0, 1, \dots, \bar{A}_l^+\}$ denote the number of empty containers that are requested at epoch t to be delivered to the depot at the end of decision period $t + l$. This implies that $A_{0,t}^+$ requested containers will only arrive at the end of the current period. We assume that all orders are deterministic, thus all requested containers arrive without delay. Finally, we let $A_t \in \mathcal{A} = \mathcal{A}^- \times \mathcal{A}_0^+ \times \mathcal{A}_1^+ \cdots \times \mathcal{A}_L^+$ denote the full decision vector, where $\times$ denotes the Cartesian product. The problem is viewed as a multiple supplier selection problem because the decision maker chooses from various supply options. Each supply option has different lead times for receiving the additional containers from the port. Extending the depot inventory management problem to multiple lead times is one contribution of **[Paper A]**.

States The inventory state of the container depot is denoted by $I_t \in \mathcal{I} = \{0, 1, \dots, \bar{I}\}$, where $\bar{I}$ denotes the finite storage capacity. In addition, we use $O_{l,t} \in \mathcal{O}_l$, $l = 0, 1, \dots, L-1$ to denote the states of outstanding repositioning orders. These *on-order states* evolve as

$$O_{l-1,t+1} = O_{l,t} + A_{l,t}^+, \; l = 1, \dots, L-1,$$

and

$$O_{L-1,t+1} = A_{L,t}^+,$$

where our notation for $O_{l,t}$ should be read as the number of previously ordered empty containers (before epoch t), that will arrive at the end of decision period $t + l$. From the transition

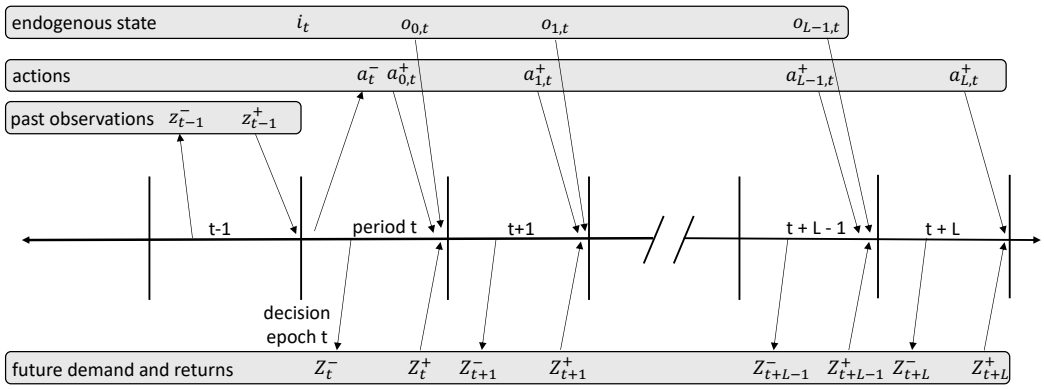

Figure 4.4: Schematic illustration of the MDP at epoch t. Arrows indicate when empty containers arrive or leave the depot. Adopted from [**Paper A**].

functions we obtain $\mathcal{O}_l = \{0, 1, \ldots, \sum_{i=l+1}^{L} \bar{A}_l^+\}$ as the finite space of the on-order state $O_{l,t}$. We denote the complete endogenous state by $S_t^\dagger = (I_t, O_{0,t}, O_{1,t}, \ldots, O_{L-1,t})$, with $S_t^\dagger \in \mathcal{S}^\dagger = \mathcal{I} \times \mathcal{O}_0 \times \mathcal{O}_1 \times \ldots \times \mathcal{O}_{L-1}$. The inventory of the depot evolves as

$$I_{t+1} = \min(\max(I_t - A_t^- - Z_t^-, 0) + A_{0,t}^+ + O_{0,t} + Z_t^+, \bar{I}),$$

where Z_t^- and Z_t^+ denote the empty container demand of exporters and returns from importers in period t, respectively. Both variables are discrete-valued random variables and unknown to the decision maker when action A_t is selected at epoch t. The inner max operator incorporates the assumption that inventory can at no point fall below zero containers. Moreover, we assume that out-positioned empty containers leave the depot immediately. The outer min operator accounts for the finite capacity of the storage and represents our assumption that empty container returns have a one decision period lead time before they can be reused to satisfy demand. The delay before returned containers can be reused necessitates to specify the joint process of Z_t^- and Z_t^+, instead a process for their difference. Figure 4.4 summarizes the state and action dynamics for this MDP.

Exogenous processes An important contribution of [**Paper A**] is the exploration of serial and cross-sectional dependent empty container demand and return processes. At this point, let us assume that both processes are exogenous, thus cannot be controlled by the decision maker. Moreover, demand cannot be backlogged. We consider the stochastic process of Z_t^- and Z_t^+ to be induced by the conditional joint distribution

$$P(z_t^-, z_t^+ | z_{t-1}^-, \ldots, z_0^-, z_{t-1}^+, \ldots, z_0^+, x_t) = P(z_t^-, z_t^+ | x_t),$$

which is conditionally independent of the process history, given an exogenous state $X_t \in \mathcal{X}$. We follow the definition of a state variable and let X_t only include information that is available at decision epoch t. In fact, [**Paper A**] considers the demand and return processes to be Markovian, hence $X_t = (Z_{t-1}^-, Z_{t-1}^+)$. Markovian demand and return processes imply that the demand and return realizations of the previous decision period $t-1$ provide useful information for choosing A_t in epoch t. Thus, the complete state vector of the inventory system becomes $S_t = (S_t^\dagger, X_t) \in \mathcal{S} = \mathcal{S}^\dagger \times \mathcal{X}$.

4.3.2 Objective function

The objective of the decision maker is the minimization of expected future total costs

$$V^{\pi}(s_0) = \mathrm{E}^{\pi}\left[\sum_{t=0}^{T-1} C_t(S_t^{\dagger}, d_t(S_t), Z_t^{-}, Z_t^{+}) + C_T(S_T^{\dagger})\middle| S_0 = s_0\right],$$

where $C_t(S_t^{\dagger}, d_t(S_t), Z_t^{-}, Z_t^{+})$ is the cost function in period t and $s_0 \in \mathcal{S}$ is the initial state of the system. The function is defined to include costs for out- and in-positioning empty containers ($d_t(S_t)$), penalty costs for not satisfying demand and for empty container returns exceeding the storage capacity, as well as holding costs. $C_T(S_T^{\dagger})$ denotes the terminal cost function. All cost parameters are time-homogeneous and known throughout the planning horizon. For a detailed description of both cost functions we refer to **[Paper A]**. We employ backward dynamic programming to estimate the optimal policy as $\pi^{\star} = \mathrm{argmin}_{\pi\in\Pi} V^{\pi}(s_0)$, where

$$\begin{aligned}
V_t(s_t) &= \min_{a_t\in\mathcal{A}_{s_t}} \mathrm{E}\big[C_t(S_t^{\dagger}, A_t, Z_t^{-}, Z_t^{+}) + V_{t+1}(S_{t+1})|S_t = s_t, A_t = a_t\big] \\
&= \min_{a_t\in\mathcal{A}_{s_t}} \mathrm{E}\big[C_t(s_t^{\dagger}, a_t, Z_t^{-}, Z_t^{+})|X_t = x_t\big] + \mathrm{E}[V_{t+1}(S_{t+1})|S_t = s_t, A_t = a_t] \\
&= \min_{a_t\in\mathcal{A}_{s_t}} \sum_{z_t^{-}, z_t^{+}} P(z_t^{-}, z_t^{+}|x_t) C_t(s_t^{\dagger}, a_t, z_t^{-}, z_t^{+}) + \sum_{s_{t+1}\in\mathcal{S}} P(s_{t+1}|s_t, a_t) V_{t+1}(s_{t+1})
\end{aligned}$$

is the Bellman equation for epoch t. The state-dependent action space $\mathcal{A}_{s_t}$ enforces the constraint that no more than currently available containers can be out-positioned. Moreover, in-positioning decisions are restricted not to exceed the maximum admissible value of each on-order state. That is, $A_{l,t}^{+} + o_{l,t} \leq \max \boldsymbol{O}_{l-1}$, with $o_{l,t}$ being the already requested containers to arrive at the end of period $t + l$. The transition probability can be shown to factorize as $P(s_{t+1}|s_t, a_t) = P(s_{t+1}^{\dagger}|x_{t+1}, s_t, a_t)P(x_{t+1}|x_t)$. It shows that we must obtain $P(x_{t+1}|x_t)$ for the exogenous state X_t if it provides useful information to the decision-maker, thus describe the evolution of the exogenous state over time.

4.3.3 Simulation studies

In **[Paper A]**, simulation studies are used to investigate how different assumptions impact optimal in- and out-positioning decisions. The storage capacity of the depot is assumed to be $I = 80$. There are three daily transportation options to in-position empty containers from the port, with lead times of zero, one, and two periods. The in-positioning costs per container decrease as the lead times increase to reflect that slower modes of transportation are cheaper. The penalty for unsatisfied demand (1,000) is large compared to the penalty for exceeding the storage capacity (50). Holding costs are small (1).

Exogenous processes and policy evaluation

All marginal demand and return processes are assumed to follow negative binomial autoregressive (NBAR) processes

$$Z_t|Z_{t-1} \sim P(z_t|z_{t-1}) = NB\left(\frac{\mu_t^2}{\sigma_t^2 - \mu_t}, \frac{\alpha}{\alpha + \mu_t}\right),$$

where the superscripts to denote demand (Z_t^-) and return (Z_t^+) variables are omitted. $NB(\cdot,\cdot)$ denotes a negative binomial distribution with mean and variance

$$\begin{aligned}\mu_t &:= \mathrm{E}[Z_t|Z_{t-1}] = \exp(c + \theta \log(Z_{t-1}+1)) \\ \sigma_t^2 &:= \mathrm{Var}[Z_t|Z_{t-1}] = \mu_t + \frac{1}{\alpha}\mu_t^2\end{aligned} \quad . \tag{4.4}$$

The autoregressive parameter θ controls the serial dependence strength, c is a drift term and a dispersion parameter α induces overdispersion. A Clayton copula

$$C_\kappa(u,v) = \max\left(\left[u^{-\kappa} + v^{-\kappa} - 1\right]^{-1/\kappa}, 0\right)$$

is applied to the marginal conditional CDFs $F(z_t^-|z_{t-1}^-)$ and $F(z_t^+|z_{t-1}^+)$ of two independent NBAR processes to generate a cross-sectional dependent process for container demand and returns. A parameter $\kappa \in [-1,\infty]\backslash\{0\}$ controls the dependence strength for the joint process. It is not straightforward to obtain analytical expressions for $P(z_t^-, z_t^+)$ when marginal processes are serially dependent and a copula is used to induce cross-sectional dependence. Therefore, $P(z_t^-, z_t^+)$ is approximated with Monte Carlo sampling. More details for the sampling method are presented in [**Paper A**].

The simulation studies investigate to the most part the effects of estimating policies under misspecified demand and return processes. For instance, a policy is estimated under the assumption that demand and return processes are serially and mutually independent. However, the true process can conversely be serially and cross-sectionally dependent. The effects of this misspecification between the assumed and true exogenous process are estimated with

$$\mathrm{E}\left[V^{\pi^\star}(S_0^\dagger, X_0)|S_0^\dagger = s_0^\dagger\right],$$

where the expectation is evaluated with respect to the true demand and return process, but the policy $\pi^\star$ has potential been estimated with a different process.

A dependence on x_0 is omitted to remove its effect on the estimated true costs of a policy. To approximate the expectation we perform three steps. First, 10,000 trajectories of length T are sampled from the true demand and return process. Second, starting for $t = 0$ in $s_0^\dagger$, the value of an estimated policy $\pi^\star$ is obtained by following the policy throughout the planning horizon while using demand and return samples to obtain the costs in each period. Third, an approximation of the expectation is obtained by averaging all 10,000 policy values. We follow this procedure to quantify the bias in the objective function, i.e. value function at $t = 0$, as a consequence of a misspecified exogenous process. However, a decision maker in real-world problems would only implement the optimal decision for the beginning of the planning horizon at $t = 0$ (cf. Figure 4.1).

An important aspect is that $\pi^\star$ denotes an optimal policy for a specified MDP, thus for an assumed demand and return process. The policy is numerically optimal because backward dynamic programming finds an exact solution for the objective function. However, we emphasize that a policy may not always be optimal in the sense of lowest operating costs as a consequence of misspecifications. The misspecification costs for policy $\pi_1^\star$ are obtained by evaluating the ratio

$$R(\pi_1^\star, \pi_2^\star) := R(\pi_1^\star, \pi_2^\star|s_0^\dagger) = \frac{\mathrm{E}\left[V_{\pi_1^\star}(S_0^\dagger, X_0)|S_0^\dagger = s_0^\dagger\right]}{\mathrm{E}\left[V_{\pi_2^\star}(S_0^\dagger, X_0)|S_0^\dagger = s_0^\dagger\right]}, \tag{4.5}$$

Table 4.1: Serial and cross-sectional dependence of the exogenous process models for the MDP formulations. Taken from [**Paper A**].

MDP name	demand & return dependencies	exogenous state
AR	$P(z_t^-\|z_{t-1}^-)P(z_t^+\|z_{t-1}^+)$	$X_t = (Z_{t-1}^-, Z_{t-1}^+)$
IID	$P(z_t^-)P(z_t^+)$	none
ARcross	$P(z_t^-, z_t^+\|z_{t-1}^-, z_{t-1}^+)$	$X_t = (Z_{t-1}^-, Z_{t-1}^+)$
IIDcross	$P(z_t^-, z_t^+)$	none

with $\pi_2^\star$ being estimated under the true exogenous process, whereas $\pi_1^\star$ is not. The ratio is also known as relative regret. The initial state $s_0^+ = (i_0, o_0) = (\bar{I}/2, 0) = (40, 0)$ is considered in the following experiments. That is, all policies are applied in a state where no previous containers have been ordered. Finally, a planning horizon is $T = 100$.

Varying in-positioning lead times

The first experiment in [**Paper A**] investigates the benefits of modeling in-positioning options with greater lead times and lower costs. Policies are estimated for MDPs that each assume different available in-positioning options. The maximum lead time is two decision periods and all in-positioning capacities are equal, i.e. $\bar{A}_0 = \bar{A}_1 = \bar{A}_2$, but costs decrease as lead times increase. Thus, there is a financial reward for using in-positioning options with greater lead times. The first MDP considers only the availability of the fastest option, thus $\bar{A}_1 = \bar{A}_2 = 0$. The remaining two MDPs assume maximum in-positioning leads times of one and two periods. These MDPs are paired with five different exogenous processes (serially and mutually independent negative binomial processes) to investigate the effect of greater demand variance on the benefits of using slower in-positioning options. The results in [**Paper A**] compare policies among each other on the basis of the proposed policy value ratio in Equation (4.5). Each ratio quantifies the additional expected depot operating costs when only fast and more expensive modes of transportation are considered. Our results show that the total cost differences decrease as the demand variance increases. We therefore conjecture that the benefits of using slower modes of transportation diminishes as a consequence of greater future uncertainty which makes planning multiple periods ahead more difficult. Nevertheless, lower expected operating costs are found for all considered demand processes when slower and cheaper modes of transportation can be used.

Misspecification of the exogenous process

Four classes of MDPs are proposed to investigate the effects of serial and cross-sectional dependence misspecification. Each MDP in Table 4.1 is labeled according to the dependence assumption of the exogenous demand and return process. The most complex process is assumed for the *ARcross* MDP, where empty container demand and returns are serially and cross-sectionally dependent. An *IID* MDP considers conversely a serial and mutual independent demand and return process. The true exogenous process follows two marginal NBAR processes, where dependence between container demand and returns is induced with a Clayton copula. [**Paper A**] describes the procedure for generating samples for this process and how the distributions in Table 4.1 are numerically approximated for the MDPs.

Misspecified serial dependence Our observations in [**Paper A**] follow the established knowledge in the inventory management literature, in that undetected positive serial dependence of the demand process leads to understocking, hence inventory is on average too low. To illustrate

Table 4.2: Relative regrets for serial and cross-sectional dependence misspecification for the *IID*, *IIDcross* and *AR* policies. Kendall's τ reports the cross-dependence strength, with $\tau = \kappa/(2+\kappa)$ for the applied Clayton copula. Adopted from [**Paper A**].

θ^-	c^-	θ^+	c^+	τ	$R(\pi^\star_{AR}, \pi^\star_{ARcross})$	$R(\pi^\star_{IIDcross}, \pi^\star_{ARcross})$	$R(\pi^\star_{IID}, \pi^\star_{ARcross})$
-0.35	1.5	-0.35	1.5	0.5	1.07	1.03	1.25
-0.35	1.5	-0.35	1.5	0.25	1.02	1.05	1.13
-0.35	1.5	-0.35	1.5	0	1.00	1.06	1.06
-0.35	1.5	-0.35	1.5	-0.25	1.01	1.07	1.03
-0.35	1.5	-0.35	1.5	-0.5	1.02	1.07	1.02
0.35	0.7	0.7	0.5	0.50	1.08	1.02	1.01
0.35	0.7	0.7	0.5	0.25	1.03	1.11	1.04
0.35	0.7	0.7	0.5	0	1.00	1.17	1.17
0.35	0.7	0.7	0.5	-0.25	1.02	1.21	1.39
0.35	0.7	0.7	0.5	-0.5	1.06	1.21	1.58
0.35	0.7	0	0.5	0	1.00	1.05	1.05
0.35	0.7	-0.7	0.5	0	1.00	1.01	1.01

the occurrence of the effect, consider that a large demand realization in the previous period reduced the inventory level at the present decision epoch. Positive serial dependence implies a greater likelihood for another larger demand observation to occur in the coming decision period. In a presently low inventory state, this informs the decision maker to order more containers than it would otherwise order if the demand in the previous period was low. This information is lost for serial independence assumptions, where the in-positioning decision only depends on the endogenous state. The reverse effect occurs for negative serially dependent demand, where overstocking occurs. Thus, inventory levels are on average too high.

[**Paper A**] extends these results to stochastic return processes. In agreement with the demand process, undetected positive serial dependence leads to understocking, whereas undetected negative dependence to overstocking. In addition, our results show interactions between the misspecification of demand and return processes. To explore these, let us consider the results in Table 4.2. The relative regret is obtained for $\pi^\star_{ARcross}$, which is the policy that is estimated under the true demand and return process, as indicated by the parameters in the first column. Each relative regret measures the increase of expected total depot operating costs during the planning horizon as a consequence of misspecifications. The highlighted rows summarize two key findings for mutually independent processes. First, misspecifications amplify if serial dependencies are of the same sign. Table 4.2 highlights this for serially dependent demand. The relative regret $R(\pi^\star_{IID}, \pi^\star_{AR})$ of the *IID* policy is 1.05, thus 5% greater depot operating costs than the *AR* policy, when empty container returns are serially independent. However, the regret increases to 1.17 when empty container returns are also positively serially dependent. Second, individual misspecifications are mitigated if serial dependencies are of opposite signs. The last row in Table 4.2 shows that the undetected negative dependence of container returns (overstocking) compensates for the undetected positive dependence of demand (understocking). Therefore, the relative regret of the *IID* policy is found to be comparatively small.

Misspecified cross-sectional dependence Similar effects can be reported when cross-sectional dependence is misspecified. The dependence strength is measured with Kendalls's $\tau \in [-1, 1]$, with $\tau > 0$ indicating positive dependence between empty container demand and returns in period t. When negative cross-sectional dependence is misspecified, understocking occurs. Sim-

Table 4.3: Inventory system response when either serial or cross-sectional dependence are misspecified. The table should be read as; misspecified positive (row) serial dependence (column) leads to understocking (cell).

	serial dependence	cross-sectional dependence
positive	understocking	overstocking
negative	overstocking	understocking

ilar to misspecified positive serial dependence, this occurs because missing information causes the inventory system to perform not optimal. In this case, information is lost that fewer container returns tend to follow larger demand in the same period. As a consequence of the absence of this information, too few containers are in-positioned, which is unfavorable because the inventory of the next epoch will be low. Overstocking conversely occurs for misspecified positive cross-sectional dependence. Table 4.3 summarizes the effects for each individual misspecified serial and cross-sectional dependence in our experiments. The relative regrets in Table 4.2 demonstrate additional interactions between misspecified serial and cross-sectional dependencies. The results show that the misspecification of each dependence interact with each other, where we identify the former explored mitigation and amplification effects.

Implications All misspecification costs are subject to defined cost parameters of the MDP. Understocking becomes consequently more costly if the penalty for unsatisfied demand is large, whereas overstocking is more costly when storage costs are high. However, our findings are general and should raise awareness to the importance of investigating the existence of serial and cross-sectional dependencies. The results in [**Paper A**] show that the common assumption of serial and mutual independent container demand and returns can have a severe impact on the container availability of a depot.

4.3.4 Empirical results

An empirical study with real-world data confirms the existence of dependencies for the empty container demand and return processes of a shipping company. A dataset of 2192 daily empty container deliveries, which we use as a proxy for demand, and return observations of a depot has been provided. The unknown stochastic process of container demand and returns are approximated on the basis of the historical time series data. All marginal processes are based on negative binomial regression models and cross-sectional dependence is induced with a Clayton copula. Similar to the previous study, we consider four combinations of serial and cross-sectional dependencies (cf. Table 4.1). The unknown model parameters are estimated with maximum likelihood and are found to imply positive serial and cross-sectional dependence.

Each of the four time series models are used to specify the exogenous process of container demand and returns of a MDP. Once policies are estimated for the MDPs, the historical delivery and return time series are used to evaluate them. The evaluation procedure is identical to the simulation study, albeit the planning horizon is reduced to 49 days. Thus, 44 demand and return trajectories are used to approximate expected depot operating costs over the planning horizon of 49 days. Our results find the misspecification of the positive serial dependence to be more harmful for depot operations than undetected cross-sectional dependence. The expected total operating costs of the policy with serial and mutual independent demand and returns are found to be 7% greater than for the policies which account for serial dependence (*AR* and *ARcross*). The costs when following a policy that is estimated under serial independent demand and returns,

but cross-sectional dependence, are found to be 9% greater. The larger costs, even though cross-sectional dependence is modelled, are due to the discussed mitigation effects. The policy that is estimated under the serial and mutual independence assumption has lower costs since the misspecified positive cross-sectional dependence (overstocking) mitigates the understocking effect of misspecified positive serial dependence.

Forecasting empty container deliveries and returns

In Chapter 3, we provided a comprehensive analysis of the ECR problem, emphasizing the critical need for accurate forecasts of empty container demand and returns at various spatial and temporal levels to facilitate repositioning decisions. In what follows, our focus shifts to the long-term repositioning problem from container surplus to deficit regions. The ECR decision system of Maersk requires for this planning granularity forecasts of weekly empty container demand and returns for container pools, which represent groups of nearby depots. In a previous section (Section 3.2), we highlighted the importance of forecasting unconstrained demand, which refers to demand that occurs regardless of the equipment availability at container depots. These forecasts play a key role for repositioning containers to ensure that the projected demand can be met proactively. However, accurately forecasting unconstrained demand presents significant challenges, as we discussed earlier. To address this issue, we adopt the approach presented in Section 3.2 and treat the historical empty container delivery and return time series as proxies for the true unobserved demand and returns. The historical observations at container pools are subsequently used for forecast model building and verification. The ECR system at Maersk requires uncertainty estimates for future empty container volumes. Therefore, all forecasts must be probabilistic. Furthermore, the extended planning horizon of the system requires forecasts to be made up to 13 weeks ahead. Additionally, distinct forecasts are required for various container types.

This chapter presents an introduction to this forecasting problem and our contributions towards the research directions of the forecasting problem of non-stationary time series, as well as the benefits of fusing machine learning models with parametric time series models. We restrict our attention to the empty container return forecasting problem to simplify this exposition. Nevertheless, the description of the forecasting problem in Section 5.1 also extends to the empty container delivery forecasting problem. Section 5.2 presents a primer on state space models, which are used in [**Paper B**] and [**Paper C**]. Our exposition is brief and covers only the most important aspects for state and parameter estimation of linear Gaussian models. The remaining two sections 5.3 and 5.4 summarize the contributions of [**Paper B**] and [**Paper C**], respectively.

5.1 Forecasting problem characteristics

Empty container return forecasts must be produced for container pools and various container type combinations. Container pools are distributed across all continents, with the exception of Antarctica. The location of all depots that form pools are shown in Figure 5.1, which exemplifies the operations of Maersk across the world.

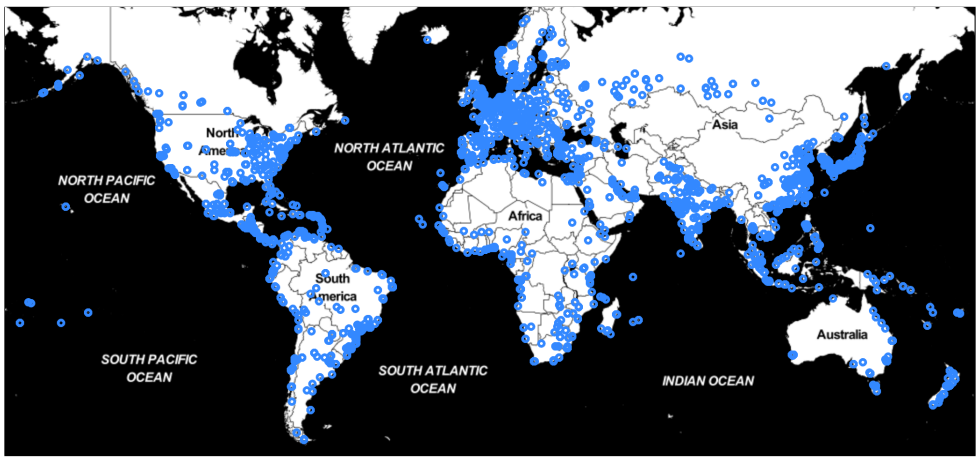

Figure 5.1: Depots that Maersk can use to store empty containers. Not all depots are always in use due to seasonal demand variations.

5.1.1 Time series heterogeneities

In each of the countries where Maersk operates, public holidays exist that can affect empty container returns due to a temporary reduction of industrial outputs. The return of containers is delayed because raw materials and components inside laden containers are not processed. These effects can be intensified when public holidays extend over several consecutive days, as is the case with Chinese New Year. The associated public holiday typically spans a full working week and significantly impacts empty container returns to many Chinese container depots. Figure 5.2 provides an illustration of the effect of Chinese New Year and other holidays on empty container returns.

It is important to note that the effect of public holidays on empty container returns is limited to the region where they are celebrated. For instance, as shown in Figure 5.2, container returns to a Japanese container pool decrease during the Golden Week public holidays in calendar week 18. This holiday has evidently no visible effect on container returns to the pools in China and Great Britain. In fact, there exists no public holiday that is universally shared across the entire world. Regional public holidays are one of many factors that explain why the properties of empty container return time series vary across different locations. For a single pool, additional variation exists between container returns of different types. The majority of dry and refrigerated cargo is transported in standardized 20ft or 40ft steel containers, whereas a smaller fraction of cargo is transported in specialized containers. Demand for specialized containers is commonly infrequent and can vary strongly across locations, hence many return time series are intermittent or lumpy. The observations of most time series for standard container types are conversely greater than zero. The forecasting problem is consequently considered heterogeneous due to the varying time series properties across locations and container types.

5.1.2 Exogenous covariates

The impact of public holidays on container returns suggest to derive covariates that enhance the forecasting models. By having knowledge of past and future dates of public holidays, the forecaster can utilize this information as the basis for deriving covariates. In this context, the available holiday information to the forecaster is deterministic since the occurrence of holidays

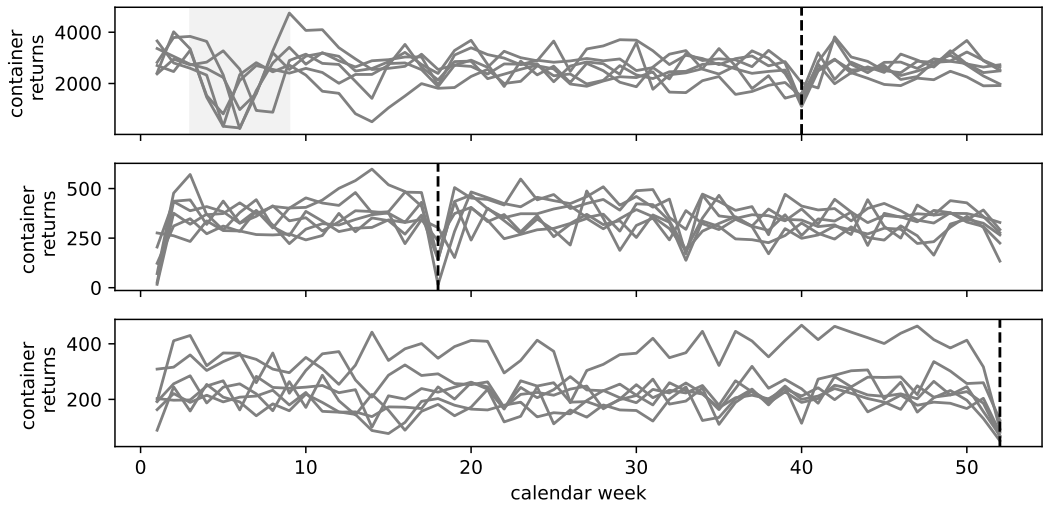

Figure 5.2: Year-to-year (2017-2022) plot for empty container (standard 40ft) returns to pools in China (top), Japan (centre) and Great Britain (bottom). The dashed vertical lines indicate regionally important public holidays that affect empty container returns. The effect of Chinese New Year is visible in the top row during calendar weeks three to nine (shared region).

is known for the forecast horizon. However, the information would become stochastic if there is a possibility for a public holiday not to occur, resulting in regular businesses to remain open. Indeed, the description of the booking process in Section 2.2 serves as the foundation for deriving stochastic covariates based on import customers's booking data. When creating forecasts for the upcoming 13 weeks, the forecaster has access to a set of active bookings, also referred to as booking data. The aim is to improve the forecasts by utilizing destination and expected arrival time information for each shipment in the booking data. However, two challenges arise for the forecaster. Firstly, the destination of a shipment does not correspond to a specific container pool, as laden containers are delivered to customer sites. Secondly, the time between the expected laden container arrival and the return of the empty container can vary due to factors such as delays during the transportation of laden containers and customers extending the permitted detention time for the empty container.

Maersk addresses both challenges by generating covariates for each pool and week within the forecast horizon. Each covariate represents the expected empty container returns for a specific pool and future week based on the current booking data. Since the return location and time of an active booking are uncertain, each covariate is stochastic. Furthermore, the covariates evolve dynamically over time as new bookings are continuously accepted while existing bookings may be canceled or modified. The shipping company has more visibility on bookings that are expected to return empty containers within the next two weeks due to the commonly long transportation times between origin and destination ports. That is, empty containers that are to be returned soon are already traveling as laden containers or have already been delivered to an importer. In contrast, fewer bookings are currently known to the shipping company for empty containers that will be returned far in the future. Figure 5.3 illustrates these effects and demonstrates how the booking covariate decreases in value as the forecast lead time increases.

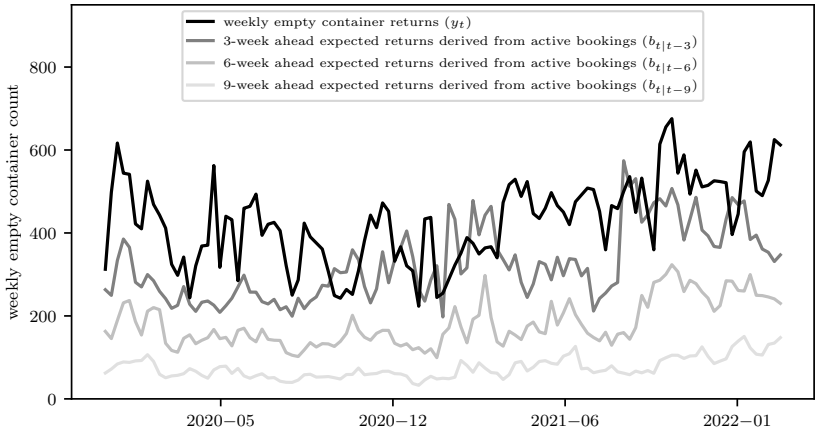

Figure 5.3: Example empty container return time series and associated booking covariates from the dataset that is considered in the empirical analysis in [**Paper B**]. The illustration is adopted from [**Paper B**].

5.1.3 Non-stationarities

Several time-varying internal and external factors can have an impact on empty container returns to a container pool. Internal effects are related to the operational decisions made by a shipping company. For example, the company can adjust its empty container return policy to reduce the detention time for importers who keep a container after it has been delivered laden. As a result, it is expected that customers will return empty containers earlier to avoid incurring additional fees. While the total number of empty container returns is unlikely to be affected by this policy change, the booking covariates will reflect the trend of earlier returns. Thus, there is a time-varying interaction between the covariate and empty container returns. Additionally, a shipping company's decision to modify its ocean network can also impact empty container returns. Introducing a new service connection to a port can increase the number of empty container returns as additional customers begin utilizing it. Conversely, if a service is removed, container returns may decrease as a result. These network changes have a direct effect on the flow of containers, thus also influence empty container returns to depots.

External factors, such as the recent COVID-19 pandemic or the blockage of the Suez Canal, are beyond the control of shipping companies. The pandemic had a significant impact on supply chains, resulting in port closures, reduced trucking capacities due to driver shortages, and factory closures leading to laden containers not being emptied. Consequently, empty containers were returned later than usual during this period. These effects persisted for several weeks, causing a decrease in empty container returns and affecting the booking covariates as well. The variability of global trade is another factor that influences empty container returns. Shipping demand is closely tied to the state of the global economy, which naturally fluctuates over time. The global recession following the financial crisis in the early 21st century had a widespread impact, although different countries were affected to varying degrees. Some countries experienced a more severe recession, resulting in a greater reduction in the demand for sending and receiving freight. Figure 5.4 illustrates three time series where non-stationarities are caused by one or multiple of the aforementioned internal and external factors.

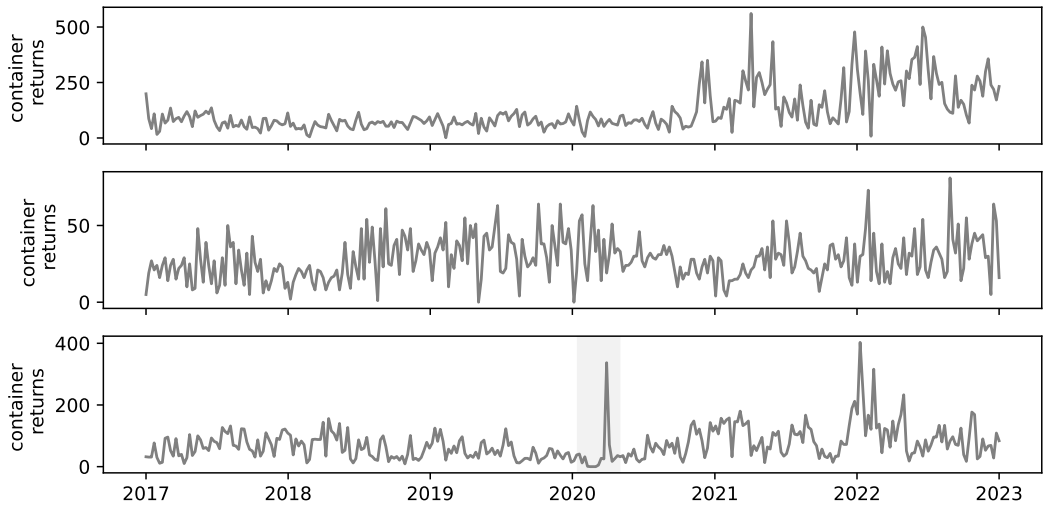

Figure 5.4: Selected non-stationary time series of empty container returns (standard 40ft) to pools in Asia. All time series show signs of volatility and level changes. The shaded region (bottom row) highlights the period when COVID-19 affected the empty container returns to this pool. Weeks with almost no returns (due to regional lockdowns) were followed by a single week with many returns.

5.1.4 Problem dimensions

The problem at hand is characterized as high-dimensional due to the need to generate empty container return forecasts for approximately 4,000 time series. Figure 5.5 shows that the majority of these time series exhibit low weekly empty container return values. This observation is attributed to the sporadic demand for specialized containers in numerous locations. Furthermore, the time series are relatively short, as reliable data is only available from 2017 onwards. Approximately 340 historical observations are available for each time series in June 2023.

5.2 State space models

We begin the description of a solution for this forecasting problem by considering an existing model selection framework to tackle the heterogeneity of the time series data. Our objective is to contribute a forecasting model to a candidate set that is composed to predict the empty container returns for a single pool and container type combination. In particular, we concentrate our attention on time series characterized by substantial weekly empty container returns at strategically important locations, specifically for standard container types. These are the container pool and type combinations situated in the right tail of the distribution of weekly empty container returns shown in Figure 5.5. We justify our focus on these time series by highlighting the critical importance of accurate forecasts for locations with significant weekly empty container return volumes. Large forecast errors in these cases can have detrimental effects on ECR decisions, emphasizing the need for precise predictions.

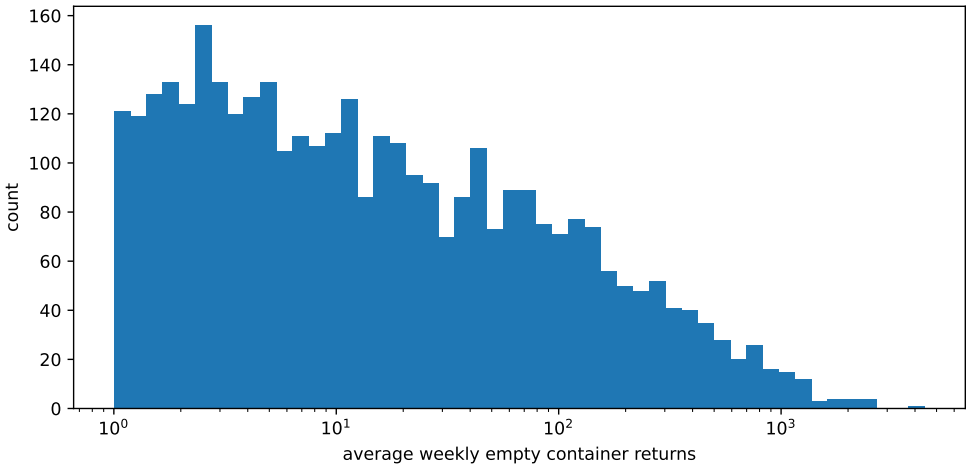

Figure 5.5: Distribution of average weekly empty container returns to pools from 2017-2023 for various container types. Combinations for pools and container types with on average less than one returned container per week are removed.

5.2.1 Motivating state space models

The presence of non-stationarities and the need to generate probabilistic forecasts provide strong motivations for employing the state space modeling framework in this problem. State space models assume that a time series process consists of underlying latent factors that evolve as a Markov process over time. The associated frameworks provide treatments for inferring the underlying latent states from potentially imprecise time series observations. This explains their popularity in many signal processing applications, where the true signal can only be measured imprecisely. The versatility of state space models also extends to forecasting tasks. The framework allows for the estimation of various widely-used time series models, including exponential smoothing (Hyndman et al., 2008) and ARIMA (Durbin and Koopman, 2012). Additionally, it facilitates the estimation of structural time series models (Durbin and Koopman, 2012), which incorporate multiple structural components such as trends, seasonalities, and cycles into the time series process. These features make state space models a powerful tool for addressing the complexities of the forecasting problem at hand.

Forecasts generated within the state space modeling framework inherently possess a probabilistic nature because the models explicitly describe stochastic processes. One of the strengths of this framework is its ability to estimate both stationary and non-stationary time series models within the same framework. For the state space models presented in [**Paper B**], non-stationary state processes are employed to capture and adapt to changes in the dynamics of the time series. A different approach is taken in [**Paper C**], where state space models with time-varying model parameters are estimated. Additionally, the state space modeling framework readily accommodates the inclusion of covariates in the parameterization of the models. The interaction between the covariates and the state variables can be explicitly modeled, providing a powerful way to capture the dynamics and dependencies within the data. In summary, the state space modeling framework strikes a balance between flexibility and structural assumptions in time series modeling. The use of latent states allows for the representation of non-stationary processes, with the ability to adapt to changes in the observed time series. At the same time, structural assumptions can be incorporated into the models, such as seasonal components or covariate information,

enabling a comprehensive and robust representation of the underlying time series dynamics.

5.2.2 Linear Gaussian state space models

The popularity of state space modeling dates back to the 1960s and is strongly associated with the derivation of the Kalman filter (Kalman, 1960). Since then numerous contributions have been made to the field, particularly in the domain of non-linear systems. Today, a large body of text books is dedicated to state space modeling. The treatment and analysis of univariate and multivariate time series can be found in Durbin and Koopman (2012), Hamilton (1994) and Harvey (1989). The following expositions are restricted to univariate linear Gaussian state space models, which constitute the foundations of the models in [**Paper B**] and [**Paper C**]. We adopt the notation of [**Paper C**] and let $z_t \in \mathbb{R}$ denote the real-valued time series observation at discrete time point t, with $t = 1, 2, \ldots, T$.

A linear Gaussian state space model assumes a latent state vector $\boldsymbol{l}_t \in \mathbb{R}^m$ to follow a Markov process

$$\boldsymbol{l}_t = \boldsymbol{F}_t \boldsymbol{l}_{t-1} + \boldsymbol{R}_t \boldsymbol{\eta}_t, \ \boldsymbol{\eta}_t \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{Q}_t), \tag{5.1}$$

where $\boldsymbol{F}_t$ and $\boldsymbol{R}_t$ are parameter matrices of appropriate dimensions. The evolution of the latent state over time is driven by Gaussian noise $\boldsymbol{\eta}_t$, with covariance matrix $\boldsymbol{Q}_t$. The observation process

$$z_t = \boldsymbol{a}_t^\top \boldsymbol{l}_t + b_t + \epsilon_t, \ \epsilon_t \sim \mathcal{N}(0, \sigma_t^2) \tag{5.2}$$

maps the unobserved state vector to the time series observation, where $\boldsymbol{a}_t$ is a parameter vector and σ_t^2 is the variance of the Gaussian observation noise ϵ_t. The parameter b_t may be a time-varying intercept or a function of covariates, i.e. $b_t = \boldsymbol{\beta}_t^\top \boldsymbol{x}_t$ for linear effects with covariates $\boldsymbol{x}_t$ and parameter vector $\boldsymbol{\beta}_t$. We assume throughout the remainder of this section that the covariates are deterministic. Moreover, all expectations are implicitly conditional on the set of covariates to simplify notations. We note that similar regression effects can also be incorporated into the state process. The parameterization of a linear Gaussian state space model is completed with the specification of a Gaussian prior $\boldsymbol{l}_0 \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{P}_0)$ for the initial state vector. Two important problems arise when the presented model is used for forecasting. First, a forecast of a new observation z_{T+1} depends on $\boldsymbol{l}_{T+1}$ and leads to the state estimation problem. Second, at least some of the parameters $\boldsymbol{F}_t, \boldsymbol{R}_t, \boldsymbol{Q}_t, \boldsymbol{a}_t, b_t$ and σ_t^2 are unknown. This leads to the parameter estimation problem.

State estimation

The inference of the unobserved state vector $\boldsymbol{l}_t$ based on the observation sequence $z_{1:t} = (z_1, z_2, \ldots, z_t)$ up to time t is better known as filtering. The filtering task addresses the estimation of the conditional probability density function $p(\boldsymbol{l}_t | z_{1:t})$. We subsequently show that this density is multivariate Gaussian and derive the Kalman filter to estimate the conditional mean and variance of the distribution recursively.

We begin our derivation by showing that the joint density $p(\boldsymbol{l}_{1:t}, z_{1:t})$ is multivariate Gaussian. First, we note that the marginal density

$$p(\boldsymbol{l}_{1:t}) = p(\boldsymbol{l}_1) \prod_{i=2}^{t} p(\boldsymbol{l}_i | \boldsymbol{l}_{1:i-1}) = p(\boldsymbol{l}_1) \prod_{i=2}^{t} p(\boldsymbol{l}_i | \boldsymbol{l}_{i-1}) \tag{5.3}$$

is multivariate Gaussian. This can be verified in a few steps. First, the factorization in Equation (5.3) is due to the Markov property of the state process (5.1). From the linearity of the state

process follows next that the conditional density $p(\boldsymbol{l}_t|\boldsymbol{l}_{t-1})$ is multivariate Gaussian. Last, we use the standard result that the product of two Gaussian densities is also Gaussian. The density $p(\boldsymbol{l}_{1:t})$ is consequently Gaussian because $p(\boldsymbol{l}_1)$ is Gaussian, which follows from the selected prior distribution $\boldsymbol{l}_0 \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{P}_0)$ and the linear state process (5.1). Second, we note that the conditional density

$$p(z_{1:t}|\boldsymbol{l}_{1:t}) = p(z_t|\boldsymbol{l}_{1:t}) \prod_{i=2}^{t} p(z_i|z_{1:i-1}, \boldsymbol{l}_{1:t}) = p(z_1|\boldsymbol{l}_1) \prod_{i=2}^{t} p(z_t|\boldsymbol{l}_t)$$

is also multivariate Gaussian. The observation process (5.2) shows that $p(z_t|\boldsymbol{l}_t)$ is Gaussian, which immediately verifies the result. From Bayes rule follows that $p(\boldsymbol{l}_{1:t}, z_{1:t})$ must be Gaussian. Since the density $p(\boldsymbol{l}_{1:t}, z_{1:t})$ is Gaussian, $p(\boldsymbol{l}_t, z_{1:t})$ is Gaussian too.

Lemma 1. *Suppose that $\boldsymbol{X}$ and $\boldsymbol{Y}$ are jointly Gaussian distributed random vectors with*

$$\mathrm{E}\left[\begin{pmatrix}\boldsymbol{X}\\ \boldsymbol{Y}\end{pmatrix}\right] = \begin{pmatrix}\boldsymbol{\mu}_x\\ \boldsymbol{\mu}_y\end{pmatrix} \quad \mathrm{Var}\left[\begin{pmatrix}\boldsymbol{X}\\ \boldsymbol{Y}\end{pmatrix}\right] = \begin{pmatrix}\boldsymbol{\Sigma}_{xx} & \boldsymbol{\Sigma}_{xy}\\ \boldsymbol{\Sigma}_{xy}^\top & \boldsymbol{\Sigma}_{yy}\end{pmatrix},$$

where $\boldsymbol{\Sigma}_{xx}$ is a non-singular matrix. The conditional distribution of $\boldsymbol{Y}$ given $\boldsymbol{X}$ is also Gaussian with conditional mean given by

$$\mathrm{E}[\boldsymbol{Y} \mid \boldsymbol{X} = \boldsymbol{x}] = \boldsymbol{\mu}_y + \boldsymbol{\Sigma}_{xy}^\top \boldsymbol{\Sigma}_{xx}^{-1} \left(\boldsymbol{x} - \boldsymbol{\mu}_x\right)$$

and conditional covariance matrix

$$\mathrm{Var}[\boldsymbol{Y} \mid \boldsymbol{X} = \boldsymbol{x}] = \boldsymbol{\Sigma}_{yy} - \boldsymbol{\Sigma}_{xy}^\top \boldsymbol{\Sigma}_{xx}^{-1} \boldsymbol{\Sigma}_{xy}.$$

The result extends to non-Gaussian joint distributions, where $\boldsymbol{\mu}_y + \boldsymbol{\Sigma}_{xy}^\top \boldsymbol{\Sigma}_{xx}^{-1} \left(\boldsymbol{x} - \boldsymbol{\mu}_x\right)$ becomes the Best Linear Unbiased Estimator (BLUP).

Proof. See Durbin and Koopman (2012).

In the next steps we use Lemma 1 multiple times. First, applying the lemma to the density $p(\boldsymbol{l}_t, z_{1:t})$ verifies that the filter density $p(\boldsymbol{l}_t|\boldsymbol{z}_{1:t})$ is also multivariate Gaussian. Similar to Durbin and Koopman (2012), we note that the joint conditional density function

$$p(\boldsymbol{l}_t, z_t|z_{1:t-1}) = p(\boldsymbol{l}_t|z_{1:t}) p(z_t|z_{1:t-1})$$

is multivariate Gaussian. This is verified by using the state (5.1) and observation (5.2) processes to write

$$p(z_t|z_{1:t-1}) = p(\boldsymbol{a}_t^\top(\boldsymbol{F}_t \boldsymbol{l}_{t-1} + \boldsymbol{R}_t \boldsymbol{\eta}_t) + b_t + \epsilon_t|z_{1:t-1})$$

for the 1-step ahead forecasting density. The density is Gaussian due to the linearity and the independence assumption of the Gaussian innovations $\boldsymbol{\eta}_t$ and ϵ_t. We apply Lemma 1 a second time to obtain recursions for the filter density's $p(\boldsymbol{l}_t|z_{1:t})$ mean vector $\boldsymbol{l}_{t|t}$ and covariance matrix $\boldsymbol{P}_{t|t}$ under serially independent innovations. With $\boldsymbol{l}_{t|t-1}$ and $\boldsymbol{P}_{t|t-1}$ denoting the mean vector and covariance matrix of $p(\boldsymbol{l}_t|z_{1:t-1})$, we apply Lemma 1 to $p(\boldsymbol{l}_t, z_t|z_{1:t-1})$ and obtain

$$\begin{aligned}
\boldsymbol{l}_{t|t} &= \boldsymbol{l}_{t|t-1} + \mathrm{Cov}[\boldsymbol{l}_t, z_t|z_{1:t-1}]\mathrm{Var}[z_t|z_{1:t-1}]^{-1} \left(z_t - \mathrm{E}[z_t|z_{1:t-1}]\right)\\
\boldsymbol{P}_{t|t} &= \boldsymbol{P}_{t|t-1} - \mathrm{Cov}[\boldsymbol{l}_t, z_t|z_{1:t-1}]\mathrm{Var}[z_t|z_{1:t-1}]^{-1}\mathrm{Cov}[\boldsymbol{l}_t, z_t|z_{1:t-1}]^\top.
\end{aligned}$$

From the state process (5.1) and the serial independence assumption of the process noise $\boldsymbol{\eta}_t$ follow the recursive expressions

$$\begin{aligned}\boldsymbol{l}_{t|t-1} &\coloneqq \mathrm{E}[\boldsymbol{l}_t|z_{1:t-1}] = \boldsymbol{F}_t \boldsymbol{l}_{t-1|t-1} \\ \boldsymbol{P}_{t|t-1} &\coloneqq \mathrm{Var}[\boldsymbol{l}_t|z_{1:t-1}] = \boldsymbol{F}_t \boldsymbol{P}_{t-1|t-1} \boldsymbol{F}_t^\top + \boldsymbol{R}_t \boldsymbol{Q}_t \boldsymbol{R}_t^\top,\end{aligned}$$

which are also known as *prediction equations*. The covariance

$$\mathrm{Cov}[\boldsymbol{l}_t, z_t|z_{1:t-1}] = \boldsymbol{P}_{t|t-1} \boldsymbol{a}_t$$

is due to the independence between $\boldsymbol{l}_t$ and ϵ_t as well as the serial independence assumption for ϵ_t. The conditional mean of the 1-step ahead forecast density $p(z_t|z_{1:t-1})$ is obtained as

$$\hat{z}_{t|t-1} \coloneqq \mathrm{E}[z_t|z_{1:t-1}] = \boldsymbol{a}_t^\top \boldsymbol{l}_{t|t-1}$$

and the corresponding variance as

$$\hat{\sigma}^2_{t|t-1} \coloneqq \mathrm{Var}[z_t|z_{1:t-1}] = \boldsymbol{a}_t^\top \boldsymbol{P}_{t|t-1} \boldsymbol{a}_t + \sigma_t^2.$$

Finally, we obtain the filter equations

$$\begin{aligned}\boldsymbol{l}_{t|t} &= \boldsymbol{l}_{t|t-1} + \boldsymbol{K}_t \left(z_t - \hat{z}_{t|t-1}\right) \\ \boldsymbol{P}_{t|t} &= \boldsymbol{P}_{t|t-1} - \boldsymbol{K}_t \boldsymbol{a}_t^\top \boldsymbol{P}_{t|t-1}^\top,\end{aligned}$$

where $\boldsymbol{K}_t = \boldsymbol{P}_{t|t-1} \boldsymbol{a}_t / \hat{\sigma}^2_{t|t-1}$ is also known as the Kalman gain. The prediction and filter equations constitute the Kalman filter recursions. The filter is initialized by using the prior state distribution $\boldsymbol{l}_0 \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{P}_0)$ to set $\boldsymbol{l}_{0|0} = \boldsymbol{\mu}_0$ and $\boldsymbol{P}_{0|0} = \boldsymbol{P}_0$ for the prediction equations at $t = 1$.

Parameter estimation

In most forecasting problems, at least some state space model parameters are unknown. Thus, let $\boldsymbol{\theta}$ include all unknown parameters in $\boldsymbol{F}_t, \boldsymbol{R}_t, \boldsymbol{Q}_t, \boldsymbol{a}_t, b_t$ and σ_t^2 for all $t = 1, \ldots, T$. In addition, we assume that the mean $\boldsymbol{\mu}_0$ and variance $\boldsymbol{P}_0$ of the initial state distribution are known. The standard procedure to obtain parameter estimates for a linear Gaussian state space model is to maximize the likelihood

$$p(z_{1:T}; \boldsymbol{\theta}) = p(z_1; \boldsymbol{\theta}) \prod_{t=2}^{T} p(z_t|z_{1:t-1}; \boldsymbol{\theta}), \tag{5.4}$$

which has an analytically tractable form because the 1-step ahead forecast distributions are Gaussian. Thus, the log-likelihood is obtained as

$$\log p(z_{1:T}; \boldsymbol{\theta}) = -\frac{T}{2} \log(2\pi) - \frac{1}{2} \sum_{t=1}^{T} \left(\log(\hat{\sigma}^2_{t|t-1}) + \frac{(z_t - \hat{z}_{t|t-1})^2}{\hat{\sigma}^2_{t|t-1}} \right),$$

where $\hat{z}_{t|t-1}$ and $\hat{\sigma}^2_{t|t-1}$ are obtained from the Kalman filter. Parameter estimates can be obtained by maximizing the log-likelihood directly with the aid of numerical optimization methods. Employing the Expectation-Maximization algorithm is a common alternative to direct maximization approaches, which enjoys good convergence properties during earlier iterations. We refer the reader to Durbin and Koopman (2012) for a detailed discussion about parameter estimation techniques in state space models and the treatment of unknown parameters in the initial state distribution.

Forecasting

In the last part of this exposition we consider producing forecasts for $t = T+1, \ldots, T+\tau$, with $\tau \in \mathbb{N}_0$ being the forecast horizon, when the last time series observation is obtained at $t = T$. To achieve this, it is necessary to determine all the parameters of the state space model for $t = 1, \ldots, T+\tau$. However, it is important to note that the coefficients in $\boldsymbol{\theta}$ have only been estimated using the observed sequence $z_1, \ldots, z_T$. Thus, the estimated coefficients must parameterize the evolution of the state space model parameters over time. Once the parameters are known for $t = T+1, \ldots, T+\tau$, forecasts can be generated using the Kalman filter, treating all observations beyond $t = T$ as missing. The procedure effectively iterates the latest 1-step ahead state prediction $\boldsymbol{l}_{T+1|T}$ forward in time on the basis of the state process (5.1). This allows for an efficient calculation of the mean and variance of the Gaussian forecast distributions $p(z_{T+h}|z_{1:T})$ by following a recursive multi-step ahead forecasting strategy. For the treatment of missing observations we refer the reader to Durbin and Koopman (2012).

5.3 Direct multi-step ahead forecasting with state space models

Multi-step ahead forecasting, where the forecast horizon τ is greater than one, is a long-standing challenge. Several methods, which may be categorized as single-output and multiple-output (Ben Taieb et al., 2012), emerged over time. Single-output methods produce predictions for each lead time $h = 1, \ldots, \tau$ independently, whereas the sequence of predictions for the whole forecast horizon is simultaneously produced by multi-output methods. The selection of a multi-step ahead forecasting strategy is mostly problem- and model-specific, since there is no universally superior method when forecasting models are misspecified. The state space modeling framework lends itself naturally to follow a single-output strategy and produce multi-step ahead forecasts recursively. Many widely used time series models, such as the common autoregressive model (AR(p)) with Gaussian innovations

$$z_t = \sum_{i=1}^{p} \theta_i z_{t-i} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2)$$

have standard parameterizations in state space form. The estimated model parameters describe a *1-step* ahead relationship between present information and the time series observation at the next time point. By following the introduced forecasting procedure for state state models, forecasts for the next $h = 1, \ldots, \tau$ observations are produced recursively. In case of the former AR(p) model, obtaining a point forecast $\hat{z}_{T+2|T}$ requires using the 1-step ahead point prediction $\hat{z}_{T+1|T}$ since z_{T+1} is unknown when the forecast is made at time point T. However, there are forecasting problems where it is desirable to model the relationship between a time series observation at time t and information at a lagged time point $t-h$ directly, where $h > 1$.

5.3.1 Direct versus iterated multi-step ahead forecasting

The recursive (iterated) strategy produces optimal multi-step ahead forecasts if the model is correctly specified for the unknown time series process. However, forecast errors accumulate for misspecified models, which has been one of the main reasons for the development of other multi-step ahead forecasting strategies. Chevillon (2007) noted in this regard that model misspecification may occur for non-stationarity processes due to unnoticed unit roots or non-stationary covariates. Iterated strategies can face additional difficulties when stochastic exogenous covariates are used to forecast z_t with information that is available at time point $t-1$.

A common approach is to use historical observations of the stochastic covariates and estimate the 1-step ahead relationship between z_t and the covariate information that is known at time point $t-1$, say $x_{t|t-1}$. Producing genuine forecasts for $t = T+1, \dots, T+\tau$ as a next step faces the challenge that the most present stochastic covariate $x_{T|T-1}$ is observed at time $t = T$. Thus, the covariates themselves must be forecasted, i.e. $\hat{x}_{T+1|T}, \dots, \hat{x}_{T+\tau|T}$ are required. The iterated forecasting strategy can accumulate additional forecast errors as a consequence of inaccurately predicted covariates, which is particularly likely if the predicted covariates are biased. Indeed, this the case for the introduced booking covariates when the 1-step ahead booking covariate $b_{t|t-1}$ is used to estimate the forecasting model parameters. The multi-week ahead covariates $b_{T+2|T}, \dots, b_{T+\tau|T}$ must be used for genuine forecast since the 1-step ahead covariate is only available for $T+1$. However, from the design of the covariate we know that the multi-week ahead covariates become systematically smaller as the forecast horizon increases (cf. Figure 5.3), which is likely to bias the multi-step ahead predictions of empty container returns if an iterated strategy is followed.

Direct multi-step ahead forecasting strategies circumvent the former limitations. A separate model is used for each lead time $h = 1, \dots, \tau$ to directly target the relationship between the observation z_t and the forecast information that is available at $t-h$. An autoregressive model with exogenous inputs (AR(p)-X) can be formulated for the task of direct h-step ahead forecasting as

$$z_t = \sum_{i}^{p} \theta_i z_{t-i+(1-h)} + \boldsymbol{\beta}^\top \boldsymbol{x}_{t|t-h} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2), \tag{5.5}$$

where $\boldsymbol{x}_{t|t-h}$ are exogenous covariates that are lagged by h steps to predict z_t. At $t = T$, all information is known to forecast a new observation at $t = T+h$ as a result of directly targeting the h-step ahead relationship between forecast information and target variable z_t. The direct strategy has not always been found in empirical studies to be superior, even though it avoids error accumulation of the iterated strategy (Ben Taieb and Atiya, 2016). Indeed, the selection of a multi-step ahead strategy is problem and model specific. For instance, the iterated strategy adds additional computational complexity, since τ models need to be estimated. Nevertheless, the direct strategy is the natural choice for the empty container return forecasting problem due to the booking covariates.

5.3.2 Serially correlated multi-step ahead forecast errors

An often overlooked property of direct multi-step ahead forecasting is that the forecast error sequence of the optimal h-step ahead forecast is serially correlated. In fact, the h-step ahead errors of any well-conceived set of forecasts are expected to follow a moving average (MA) process of order $h-1$ (Harvey et al., 1997). We verify in [**Paper B**] that the serial correlation of forecast errors stems from the autocorrelation of the time series process that we aim to forecast. It follows that a direct multi-step ahead model, such as the AR-X model in Equation 5.5, is misspecified if serially independent errors are assumed. Minimizing the sum of squared errors of the misspecified model will provide unbiased parameter estimates, albeit the standard least squares estimator is not efficient due to the ignored serial dependence of the innovation process (Hyndman and Athanasopoulos, 2021).

[**Paper B**] demonstrates that the misspecification is more costly when time-varying coefficient regression models are estimated in state space form. Time-varying coefficient models are one of many approaches that have been developed to model non-stationary time series. The coefficients evolve over time, which makes the models implicitly non-stationary, thus suitable for the empty container return forecasting problem. Several different approaches, including the state space

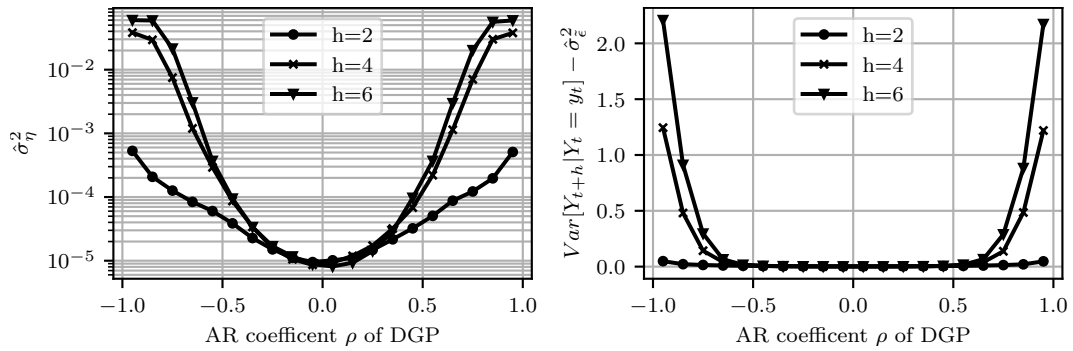

Figure 5.6: Average estimate $\frac{1}{|\mathcal{I}_\rho|}\sum_{i\in|\mathcal{I}_\rho|}\hat{\sigma}^2_{\eta,i}$ (left) and average bias $\frac{1}{|\mathcal{I}_\rho|}\sum_{i\in|\mathcal{I}_\rho|}\mathrm{Var}\left[Y_{t+h}|Y_t = y_t\right] - \hat{\sigma}^2_{\tilde{\epsilon},i}$ (right) for the stationary AR(1) process when the model (5.7) assumes the observation noise to be Gaussian white noise. Taken from **[Paper B]**.

modeling framework, exist to date that model the coefficient's evolution through time. One of the contributions in **[Paper B]** is the demonstration of biased maximum likelihood estimates for time-varying regression models in state space form that ignore the serial dependence of multi-step ahead forecast errors. Our results are significant because previous publications (cf. **[Paper B]**) have followed a similar state space approach for direct multi-step ahead forecasting without being aware that harmful consequences of misspecified innovation processes exist.

5.3.3 Estimation biases

In the following, we summarize the main results of **[Paper B]** on the basis of the conducted Monte Carlo experiment. To begin with, let a time series follow a stationary AR(1) process

$$z_t = \rho z_{t-1} + \zeta_t, \quad \zeta_t \sim \mathcal{N}(0,1), \tag{5.6}$$

for $t = 1,\ldots,2000$ and with $z_0 \sim \mathcal{N}(0, 1/(1-\rho^2))$. We generate 1000 replicates of this process for a range of autoregressive coefficients $\rho \in (-1,1)$, and let $\mathcal{I}_\rho$ denote the set of time series for a distinct value of ρ. Moreover, we define the time-varying regression model in state space form

$$\begin{aligned} \theta_t &= \theta_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0,\sigma^2_\eta) \\ z_t &= \theta_t z_{t-h} + \tilde{\epsilon}_t \\ \theta_\tau &\sim \mathcal{N}(\mu_\tau, \sigma^2_\tau) \end{aligned} \tag{5.7}$$

for the purpose of direct multi-step ahead forecasting, with lead time h and $t \geq \tau = h+1$. The innovation term $\tilde{\epsilon}_t$ is first assumed to be Gaussian white noise, i.e. $\tilde{\epsilon}_t \sim \mathcal{N}(0,\sigma^2_{\tilde{\epsilon}})$. The state space model in Equation (5.7) is purposely misspecified with respect to the true time series process. It assumes the regression coefficient to be time-varying as well the innovation to be serially independent. We proceed to estimate the unknown parameters $\boldsymbol{\psi} = (\sigma^2_\eta, \sigma^2_{\tilde{\epsilon}})$ with the first 1000 observations of each simulated time series through maximum likelihood, where we use $\mu_\tau = \rho^h$ and $\sigma^2_\tau = 0.0025$ to reduce the effect of the state initialisation on the parameter estimates. Figure 5.6 shows the bias of the variance parameter estimates as a function of the true autoregressive coefficient ρ for selected lead times. We expect from the chosen model parameterization that the estimates of $\sigma^2_{\tilde{\epsilon}}$ will be in the neighbourhood of the conditional

variance $\mathrm{Var}\left[Y_{t+h}|Y_t = y_t\right] = \sum_{i=0}^{h-1} \rho^{2i}$ of the data generation process. The results show instead that the variance is visibly underestimated for $h > 2$ and $\left|\rho\right| > 0.7$.

The innovation process should follow a MA(h-1) process when a model is parameterized for direct h-step ahead forecasting. However, we parameterized a state space model with serially independent innovations, which greater deviates from the true innovation process as the lead time h and $\left|\rho\right|$ increase. Ignoring the serial dependence leads to biased state estimates as the Kalman filter is derived under the assumption of white observation noise. The assumption is used to show that the 1-step ahead density $p(z_t|z_{1:t-1})$ is Gaussian, which requires the observation process noise ϵ_t in Equation (5.2) to be independent of $z_{1:t-1}$. The biased state estimates are consequently biasing the maximum likelihood estimates. Ultimately, we find the misspecification to negatively affect the forecast accuracy for longer lead times, given that a time series process is significantly autocorrelated with $\left|\rho\right| > 0.7$.

5.3.4 Modeling solutions

Our proposal to reduce the estimation bias is to model the h-step ahead innovations as a MA(h-1) process. The innovation process for the state space model in Equation (5.7) becomes

$$\tilde{\epsilon}_t = \sum_{i=1}^{h-1} \phi_i \epsilon_{t-i} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2), \tag{5.8}$$

where $\phi_1, \ldots, \phi_{h-1}$ are the coefficients of the moving average process. Incorporating the MA innovation process in the state space model is achieved by augmenting the state vector $\boldsymbol{l}_{t-1}$ with the lagged innovations $\epsilon_{t-1}, \ldots, \epsilon_{t-h+1}$. A detailed description of the augmentation and alternative approaches for estimating state space models with finite lag autocorrelated observation noise are presented in [**Paper B**]. Repeating the maximum likelihood estimation for the simulated AR(1) processes demonstrates the removal of the estimation biases. This improves the forecast accuracy to the optimal h-step ahead forecast with perfect information about the true AR(1) data generating process. It is straightforward to verify that the state space model in Equation (5.7) with MA($h-1$) observation noise is correctly specified for this time-invariant AR(1) process. However, the innovation process Equation (5.8) becomes misspecified in case where the autoregressive coefficient of data generating mechanism varies with time.

To extend our modeling framework to a wider range of time series processes, we subsequently consider MA processes with time-varying coefficients

$$\tilde{\epsilon}_t = \sum_{i=1}^{h-1} \tanh(\phi_{i,t}) \epsilon_{t-i} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2), \tag{5.9}$$

where each MA coefficient is a latent state that follows a random walk

$$\phi_{i,t} = \phi_{i,t-1} + \xi_{i,t}, \quad i = 1, \ldots, h-1.$$

The innovations $(\xi_{1,t}, \ldots, \xi_{h-1,t})$ are assumed to be multivariate Gaussian and serially independent. Parameterizing a state space model with this innovation process faces the challenge that both MA coefficients $\phi_{1,t}, \ldots, \phi_{h-1,t}$ and innovations $\epsilon_{t-1}, \ldots, \epsilon_{t-h+1}$ are latent state variables. The innovation process (5.9) requires therefore to model non-linear state interactions, which cannot be achieved by the Kalman filter. Our solution is to use the Unscented Kalman Filter for approximate state inference. We refer to [**Paper B**] for details about the filter and our proposal for maximum likelihood estimation of the unknown model parameters with a Gaussian surrogate likelihood.

5.3.5 Empirical results

A Monte Carlo simulation demonstrates the ability of the non-linear state space model to estimate time-varying MA coefficients. However, our empirical results for the task of empty container return forecasting finds the linear Gaussian state space model with time-invariant MA processes to produce more accurate forecasts. We expect this to be due to the introduced approximation errors of the Unscented Kalman filter. The results of our study can be summarized by the following findings. First, parameterizing an MA($h-1$) innovation process for direct h-step ahead forecasting with time-varying coefficient regression models improves forecast accuracies on average, when the models are parameterized in state space form. Here, we measure improvements with respect to the state space models with serially independent observation noise. Second, the accuracy improvements are greater for point than for probabilistic forecasts. Third, the improvements can be substantial if the innovation process misspecification leads to positively biased variance parameters that determine the variability of the time-varying regression coefficients. Overfitting occurs because the regression coefficients become too adaptive to changes in the time series process due to the positive estimation bias. A detailed discussion of the case study design and results are presented in [**Paper B**].

5.4 Fusing neural networks with state space models

A property of the empty container return forecasting problem is that time series are comparatively short, yet show complex patterns. Modeling complex time series with parameter-rich forecasting models can cause high variance for parameter estimates as a consequence of insufficiently long training data. This can lead to undesirable effects due to potentially large parameter variations if the forecasting model is re-estimated when a new observation is obtained. Pooling has been proposed in this context to reduce the variance of parameters by estimating a common forecasting model for a group of time series. Estimating a single model makes the implicit assumption that all series follow the same unknown data generating process with identical parameter values. The fulfillment of the assumption leads to a reduced parameter estimation variance since the parameters are estimated on a larger time series dataset. Following this strategy can therefore be attractive for forecasting related time series when it is challenging to estimate the effect of rare events for a single series in isolation. For instance, pooling may benefit the estimation of public holiday effects on container returns for a group of container pools within the same country.

5.4.1 Global forecasting models

Early applications of pooling can be found in forecasting problems with short panel data of related time series. For instance, Garcia-Ferrer et al. (1987) forecast future annual output growth rates for nine countries, where only 30 years of prior observations are available. The lack of historical training data poses a significant challenge for the estimation of a single model for each time series. On one hand, there is the risk of overfitting. On the other hand, it is challenging to perform model validation and selection when time series are this short. Pooling time series by estimating a parsimonious forecasting model reduces the risk of overfitting, albeit only mildly improving the model validation task. A shortcoming is that imposing the same parsimonious model for all time series in the panel can lead to estimation biases if the time series processes substantially differ from each other. That is, the panel is heterogeneous, although the pooled model imposes a homogeneity assumption. Several approaches emerged in response to this limitation, all aiming to retain the ability of learning jointly from a time series panel. Most

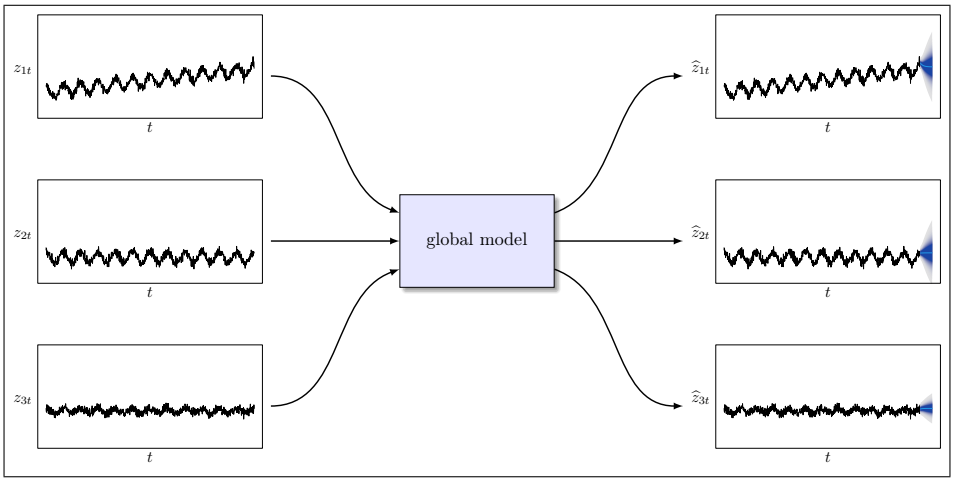

Figure 5.7: Schematic illustration for using a globally estimated model for forecasting. The global model parameters have been estimated with all three time series.

methods try to identify homogeneous sub-groups within the panel, for which a pooled model is subsequently estimated.

Many recent machine learning models for time series forecasting have been proposed in the context of forecasting related time series. Although a single model is estimated for a group of time series, there are fundamental differences to the former application of pooling. Whereas pooling has been applied in the past to reduce the variance of a parsimonious model when limited training data is available, flexible machine learning models are contemporary applied to large time series datasets. Machine learning models are flexible, and often over-parameterized, to learn complex time series patterns with limited or no assumptions about the data generating processes. The flexibility of the models additionally allows modelling heterogeneous datasets, circumventing to partition related time series into sub-groups. Cross-validation is quintessential for hyper-parameter tuning and overfitting prevention, which requires time series to be of sufficient length. Figure 5.7 illustrates schematically the application of pooling to a time series dataset, which is contemporary known as global modeling. Learning model parameters for each series in isolation is conversely known as local modeling (Montero-Manso and Hyndman, 2021), and is illustrated Figure 5.8

The inherent risk of overfitting in low data regimes as a consequence of learning complex patterns completely data-driven motivated the development of hybrid models. The principal idea of many hybrid models is to combine statistical time series models with machine learning models. On one hand, using a statistical model allows to incorporate domain knowledge about the time series processes into the hybrid model. Depending on the hybrid model design, this can induce regularization effects to mitigate the risk of overfitting. On the other hand, a hybrid model tries to benefit from the flexibility of a machine learning model to learn complex data patterns, ideally from a larger set of related time series. Several hybrid models have been proposed to date, among some following the model design to use a globally estimated machine learning model to predict parameters of locally applied statistical models. Figure 5.9 provides a schematic illustration of this hybrid strategy for time series forecasting.

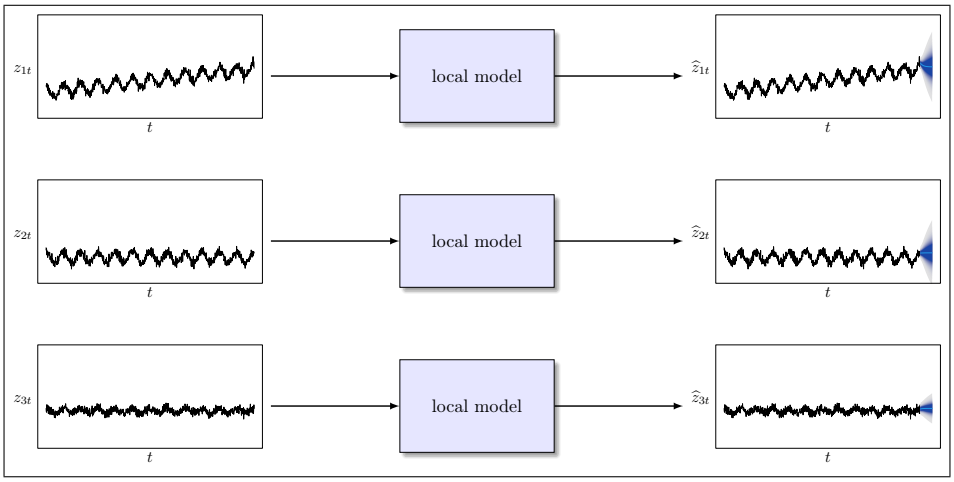

Figure 5.8: Schematic illustration for using a locally estimated models for forecasting. Each model has been estimated individually for each of the three time series.

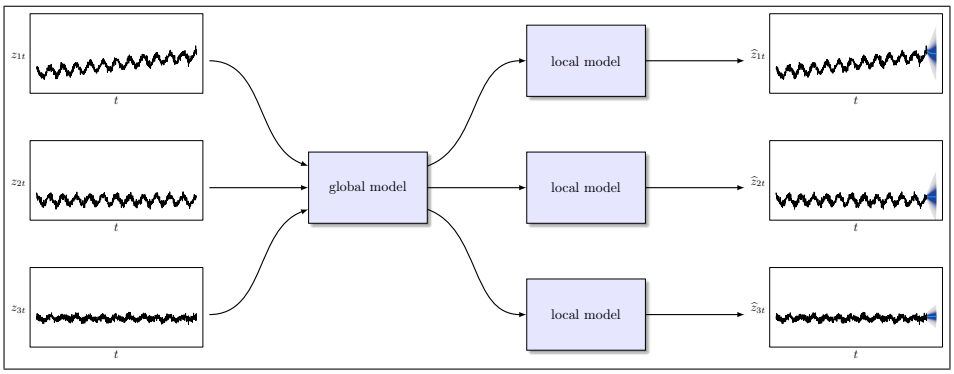

Figure 5.9: Schematic illustration of a hybrid forecasting model. A globally estimated model predicts the parameters of locally applied time series models. The input of the time series into the local models is omitted for simplicity.

5.4.2 Deep state space models

A Deep State Space Model (DSSM) follows such a strategy and uses a globally learned recurrent neural network to predict the parameters of locally applied linear Gaussian state space models. In the following, we briefly introduce DSSM to allow us to subsequently highlight important limitations of hybrid models that follow similar designs. The interested reader finds a detailed description of the model in Rangapuram et al. (2018) and a briefer summary in [**Paper C**].

Let us consider a panel of N univariate time series indexed by $i \in \mathcal{I} = \{1, \dots, N\}$. All time series are real-valued $z_t^{(i)} \in \mathbb{R}$, with $t = 1, \dots, T$, which conversely to [**Paper C**] implies equal length of all time series. We adopt the previous notation and let $z_{1:T}^{(i)} = (z_1^{(i)}, \dots, z_T^{(i)})$ denote the historical observations of time series i. For each time series, there exists a sequence of D known exogenous covariates $\boldsymbol{x}_{1:T+\tau}^{(i)} = (\boldsymbol{x}_1^{(i)}, \dots, \boldsymbol{x}_{T+\tau}^{(i)})$, where $\tau \in \mathbb{N}_{>0}$ denotes the forecast horizon. The notation omits the previous bar notation and follows the convention that $\boldsymbol{x}_t^{(i)}$ are covariates for

forecasting $z_t^{(i)}$.

DSSM tackles a probabilistic forecasting problem by assuming each time series in the panel to follow a linear Gaussian state space model. Each state space model follows an identical parameterization that is specified by the state (5.1) and observation (5.2) processes, albeit the values of unknown model parameters are allowed to vary between time series. In addition, a diffuse state prior $\boldsymbol{l}_0 \sim \mathcal{N}(\mathbf{0}, \kappa \boldsymbol{I}_m)$, with κ being large and m being the number of states, is assumed for each state space model. Let $\boldsymbol{\theta}_t^{(i)}$ denote the P-dimensional vector of unknown state space model parameters for time series i at time point t. A common approach is to use a mapping

$$\boldsymbol{\theta}_t^{(i)} = \Psi^{(i)}(\boldsymbol{x}_t^{(i)}; \boldsymbol{\phi}^{(i)})$$

to specify the evolution of the state space model parameters through time. The unknown time-homogeneous coefficients $\boldsymbol{\phi}^{(i)}$ of time series i can be obtained by maximizing the conditional likelihood

$$p(z_{1:T}^{(i)} | \boldsymbol{x}_{1:T}^{(i)}; \boldsymbol{\phi}^{(i)}) = p(z_{1:T}^{(i)}; \boldsymbol{\theta}_{1:T}^{(i)}),$$

which is computationally attractive because the likelihood can be evaluated with the Kalman filter. The difference between this and the likelihood of the linear Gaussian state space model in Equation (5.4) is that we explicitly assume the model parameters to vary over time. A DSSM extends the former idea for estimating time-inhomogeneous state space model parameters to a global modeling setup by maximizing the pseudo likelihood

$$p(z_{1:T}^{(1)}, \ldots, z_{1:T}^{(N)} | \boldsymbol{x}_{1:T}^{(1)}, \ldots, \boldsymbol{x}_{1:T}^{(N)}; \boldsymbol{\phi}) = \prod_{i=1}^{N} p(z_{1:T}^{(i)} | \boldsymbol{x}_{1:T}^{(i)}; \boldsymbol{\phi}) = \prod_{i=1}^{N} p(z_{1:T}^{(i)}; \boldsymbol{\theta}_{1:T}^{(i)}),$$

with $\boldsymbol{\theta}_t^{(i)} = \boldsymbol{\Psi}(\boldsymbol{x}_t^{(i)}; \boldsymbol{\phi})$, where the likelihood factorizes due to an imposed mutual independence assumption between the time series in $\mathcal{I}$. The global mapping $\boldsymbol{\Psi}$ is composed of a recurrent neural network and a final layer to output the state space model parameters of dimension P, where more details are provided in [**Paper C**]. The unknown parameters $\boldsymbol{\phi}$ are learned globally from all time series, whereas state space model parameters $\boldsymbol{\theta}_{1:T}^{(i)}$ are predicted locally for each series.

5.4.3 Limitations

Estimating the parameters of the recurrent neural network jointly is expected to extract features and learn complex temporal patterns from raw data, whereas incorporating structural assumptions through the state space model can alleviate overfitting (Rangapuram et al., 2018). The main contribution of [**Paper C**] is a critical analysis of the methodological and practical limitations of this approach.

Methodological limitations

By imposing the same state space model structure for all time series, DSSM follows the original concept of pooling to assume a homogeneous data generating process for all series. However, DSSM permits parameter value heterogeneities since covariates of time series i predict the state space model parameters of the respective series. An assumption of homogeneous data generating processes is in our view a major limitation for many real-world problems where different stochastic processes are likely to be present. Choosing a flexible state space model parameterization mitigates the misspecification risks, albeit defeating the purpose of using

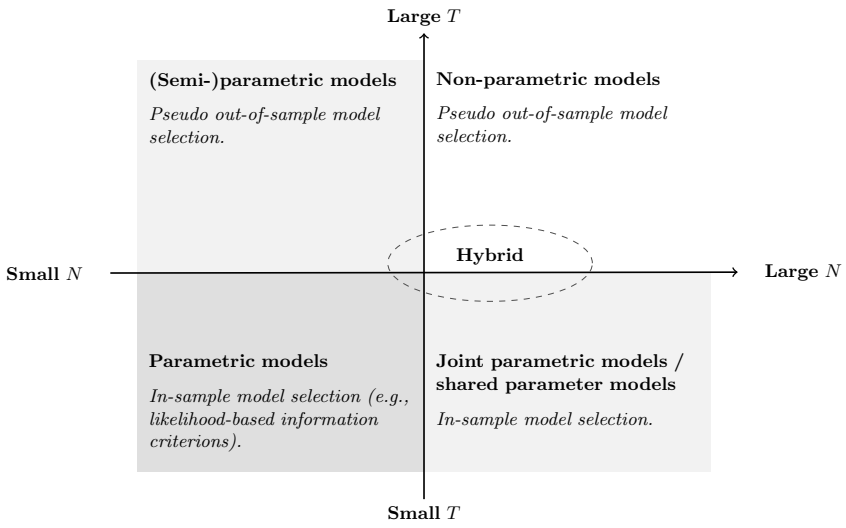

Figure 5.10: Classification of forecasting problems based on the number of time series, N, and the time series length, T. Taken from [**Paper C**].

a parametric model to mitigate the risk of overfitting. One criterion for successfully applying DSSM is therefore that all time series follow approximately the same data generating mechanism. Other criteria are the number of available time series N and time series length T. Local models are preferred when N is small and global models when N is large. Holding out valuable training observations for pseudo out-of-sample forecast evaluation is undesirable when T is small, whereas several (cross-)validation methods exists for moderately large T. Figure 5.10 summarizes these aspects and provides guidances for the applicability of a hybrid forecasting method. In [**Paper C**], we argue that DSSM is suited for forecasting problems with moderate T, as pseudo out-of-sample model forecast evaluation is required for model evaluation and selection, and moderate N of time series from a single data generation process.

Practical limitations

The imposed assumption of DSSM that each time series can be approximated by the same state space model requires model selection because of the misspecification risk. Our Monte Carlo study in [**Paper C**] highlights the negative consequences for forecast accuracies when the specified state space model cannot approximate the true data generation process well. The applicability to real-world problems with similar characteristics as the introduced empty container return forecasting task is thus limited. Moreover, [**Paper C**] demonstrates that imposing a linear Gaussian state space process for each time series does not provide a general safeguard against overfitting. Therefore, a successful application of DSSM relies on the same cross-validation methods to detect overfitting as other machine learning models. Evidently, it is possible to parameterize rigid state space models where overfitting is extremely unlikely. However, this reverts back to the loss of forecast accuracy when the state space model is significantly misspecified. Cross-validation is additionally required to tune network and optimization hyper-parameters, which is computationally costly. In particular, we note that DSSM has higher computational costs than a standard recurrent neural network since each time series is processed by the Kalman filter.

5.4.4 Empirical results

Our real-world data benchmarks in [**Paper C**] re-evaluate the forecast accuracy of DSSM on the publically available *electricity* and *traffic* datasets. The experiment follows the same design as in Rangapuram et al. (2018), albeit we additional consider forecasts of locally estimated state space models with time-homogeneous parameters. Moreover, in our experiment we replace the recurrent network of DSSM, which is subsequently denoted by DSSM-RNN, with a feedforward neural network (DSSM-FFN). The purpose is to assess the benefits of globally estimating the time-homogeneous parameters of the state space models which we use as local benchmarks for each time series. Our results for the *electricity* dataset show insignificant forecast accuracy improvements of both hybrid models in comparison to the locally estimated time-homogeneous state space models. Improvements are conversely found for the *traffic* dataset, with larger improvements for DSSM-RNN than DSSM-FFN. However, our discussions provide strong evidence that the improvements are due to misspecified state space models. The effects are less severe for DSSM-RNN since the predicted time-inhomogeneous state space model parameters, thus the ability to model heteroscedasticities, mitigate biases due to the misspecified state process.

Conclusions

This thesis presented contributions to three research directions in the field of empty container repositioning, focusing on decision and forecast model building. The first research direction involved the development of inventory control policies for container depots, with a focus on the stochastic sequential decision-making problem of an inland depot for managing container inventory. Next, the forecasting problem of non-stationary time series was studied, which has significant relevance for repositioning planning since empty container demand and return processes are often non-stationary. Last, practical applicability of a recently proposed class of hybrid forecasting models was explored.

6.1 Key findings

Towards the first research direction (**RD1**), we contributed with a novel capacitated multiple supplier periodic-review inventory model for daily ECR planning of an inland container depot. The model is formulated as a time-inhomogeneous Markov decision problem, for which we apply backward dynamic programming to estimate optimal policies that control the inventory by out- and in-positioning containers. A Monte Carlo study confirmed the ability to estimate policies that utilize slower and cheaper modes of transportation for in-positioning additional containers rather than faster but more expensive modes. The cost reduction of using slower modes of transportation were thereby found to decrease as the variance of future empty container demand increased. This verified the intuition that planning multiple periods ahead is more difficult as uncertainty increases, which makes the usage of in-positioning options with greater lead times less attractive. A second Monte Carlo study investigated the consequences for depot operational costs if serial and cross-sectional dependencies of the empty container demand and return processes were misspecified. Our results showed that undetected dependencies can lead to substantially greater operational costs than expected during policy estimation. That is, a misspecification leads to worse inventory management decisions. The implications of our Monte Carlo studies were substantiated with the existence of serial and cross-sectional dependencies in a real-world dataset of historical empty container deliveries and returns. We therefore concluded that the common assumption of serially and mutually independent empty container demand and return processes is too restrictive for real-world problems.

In relation to the problem of forecasting non-stationary time series (**RD2**), we contributed a novel methodology for direct multi-step ahead forecasting with state space models. Our methodology is based on the known property that direct multi-step ahead errors of even well-specified forecasts are finite lag serially correlated. Our main findings were estimation biases that can occur for time-varying coefficient regression models in state space form. The models were parameterized for direct multi-step ahead forecasting, but assumed serially independent observation noise. The maximum likelihood estimates of the model parameters were shown to become significantly biased if the forecast lead time was long and the true time series process was strongly autocorrelated. Parameterizing an MA($h-1$) innovation process, where h is the forecast lead time, removed the bias in our Monte Carlo simulations. We additionally introduced

a more flexible innovation process with latent MA coefficients to model a wider class of time series processes. The state space model became consequently non-linear, and we proposed to apply the Unscented Kalman filter for approximate state estimation. A Monte Carlo simulation showed the general efficacy of this approach, albeit the forecasted variances were found to be slightly biased as a consequence of the approximate state estimation. The latter approximation error was also conjectured to cause inferior forecast accuracy of the model with non-linear innovation processes at the task of forecasting empty container returns. However, our results showed the general efficacy of parameterizing a time-invariant coefficient MA($h-1$) innovation process for direct multi-step ahead forecasting. The forecasts were on average more accurate than the reference state space model with serially independent innovations.

Towards the third and last research direction (**RD3**), we conducted an assessment of Deep State Space Models to demonstrate the critical limitations of this hybrid forecasting model. The model uses a globally estimated recurrent neural network to predict the parameters of locally applied linear Gaussian state space models. We identified various methodological and practical limitations for this class of hybrid models which employ a flexible global model to predict parameters of local parsimonious sub-models that incorporate domain knowledge. The main methodological limitation stems from imposing the same state space model for all time series of the dataset. This is found to be a restriction for most real-world problems where stochastic processes can differ substantially between time series. Regarding this, we demonstrated poor forecast accuracy for the partition of a larger dataset for which the imposed sub-model could not approximate the time series well. This necessitates performing model selection for the sub-models, which limits the applicability of this hybrid model to time series with sufficiently long history for pseudo out-of-sample forecast accuracy evaluation. Moreover, we demonstrated that employing a sub-model with time-inhomogeneous parameters does not provide a general safeguard against overfitting. Even though it may reduce the speed at which the recurrent neural network overfits, the reliance on cross-validation to detect when overfitting occurs remains. Finally, our results demonstrated on two publicly available datasets that the hybrid model produced only minor forecast accuracy improvements over locally estimated parametric state space model. However, the hybrid required computationally more demanding hyper-parameter tuning schemes.

6.2 Perspectives for future research

The investigations in this thesis open up various directions for future research. To begin with, our proposed MDP to model depot operations makes several assumptions that could be relaxed in future work. The model assumes that all future transportation capacities and costs are known at the beginning of the planning horizon. This may not be the case for real-world operations where capacities and costs may only be known with certainty for the coming days. Similarly, we considered all the cost parameters of the depot to be known, even though the estimation of a lost sales penalty is challenging in practice. One direction to follow could be to use current freight rates to determine lost sales costs. However, future rates are unknown and the associated penalty would ignore the possibility of receiving additional containers for a premium cost on short notice. In regard to this, a lost sales assumption can be too restrictive when decisions are made daily, as exporters can wait in most cases another day to receive an empty container. Backlogging empty container delivery orders for one or two days could be a more realistic assumption, albeit the state space of our model would further increase. The necessity to keep the state space sufficiently compact is a limitation of our proposal to find optimal policies with backward dynamic programming. A methodological future research direction is therefore the

development of approximate dynamic programming (ADP) algorithms.

Three future directions are interesting. The first direction can utilize our MDP in conjunction with backward dynamic programming as a baseline for designing efficient ADP methods for a single depot. Once the properties of the approximation are understood, the MDP can be extended by relaxing the former assumptions. In line with this direction, we would further investigate the compatibility with routing models, particularly the effects arising when requested empty containers cannot be delivered or arrive delayed. A second research direction explores other objective functions towards risk-aware decision making. Other risk measures can consider a quantile of future total costs, or minimax objectives to minimize the maximum total costs during the planning horizon. As a third direction, we propose to investigate the application of ADP to estimate repositioning decisions in a network of depots. In all cases, using ADP tackles the curse of dimensionality of large state spaces. Depending on the ADP algorithm, it may still be necessary to parameterize a stochastic process for empty container demand and returns, which is particularly challenging for a network of depots. Therefore, future research should investigate the applicability of arbitrary probabilistic forecasts to estimate container inventory management decisions. An immediate challenge is thereby that an MDP requires the full description of the exogenous stochastic process. Thus, it is unclear how direct multi-step ahead forecasts can be used.

The exploration of Model Predictive Control (MPC) methods is in this regard an interesting future research direction. MPC solves a sequential decision problem over a specified planning horizon, but only implements the decision at the initial epoch. Once the system has evolved into a new state after the decision has been taken, MPC re-estimates the decision for the initial epoch with newly obtained information. That is, MPC is applied as a rolling horizon procedure. We may therefore consider estimating policies for our formulated MDP in a rolling horizon procedure, where the latest probabilistic forecasts for future empty container demand and returns are used instead of the fully specified stochastic process. The forecasts are assumed to be serially independent as the estimation of the joint forecast distribution over the planning horizon is likely infeasible for long horizons. MPC is only one approach that uses approximations to the sequential decision problem to permit the usage of arbitrary probabilistic forecasts. Independently of the approximation method, we see the relaxation of parameterizing the exogenous stochastic process for the MDP as a requirement to benefit from the other methodological contributions of this thesis towards the forecasting literature.

For our developed direct multi-step ahead forecasting methodology with state space models, we used the Unscented Kalman filter for approximate state estimation when non-linear innovation processes were applied. Future research should address whether more suitable methods can be found to perform state estimation at similar or less computational costs, while reducing the observed approximation errors in our simulation and empirical studies. Similarly, we may want to relax the necessity of using a Gaussian surrogate likelihood for parameter estimation. This can extend the applicability of our method to time series that are poorly approximated by a Gaussian likelihood. A limitation of our proposal is that the number of unknown MA coefficients in the innovation process grows with the forecast lead time. We therefore consider investigating whether more parsimonious innovation processes, for instance when using fewer MA terms, can prevent the observed estimation biases to a similar degree.

Other research in this direction can explore the applicability of multi-input multi-output strategies for multi-step ahead forecasting with state space models. Instead of forecasting each lead time in isolation, we can consider a multivariate state space model to forecast all or a subset of lead times simultaneously. A benefit could thereby be the ability to share parameters of the innovation processes across lead times, which can lead to more parsimonious parameterizations. Similarly, we can explore whether sharing innovation process parameters across

related time series is sufficient to reduce estimation biases. For a fixed forecast lead time, we could parameterize a multivariate state space model for a set of related time series, and share latent moving average coefficients among all innovation processes. However, it is likely that the latent state space grows too large as a growing number of lagged innovations must be included. It is additionally not clear how to perform in-sample model selection between univariate and multivariate models. Moreover, further research should investigate in-sample model selection between direct and iterated multi-step ahead parameterizations. To the best of our knowledge, it is not clear if it is possible to perform in-sample model selection based on information criteria between our direct multi-step and standard iterated multi-step ahead state space model parameterizations.

Using pseudo out-of-sample forecast evaluation remains conversely a necessity for model selection with hybrid forecasting models, which imposes the data requirement for having sufficiently long time series. Future research should determine whether the reduced overfitting risk due to the imposed parametric sub-model has practical benefits. A shortcoming of our research is in this regard a missing comparison of the hybrid model against a standard recurrent neural network. We conjecture that a sufficiently regularized and globally estimated model will perform in our experimental study equally well or better. We particularly expect improvements for the *traffic* dataset, since we have found the employed state space model to be misspecified to capture weekend effects. Evidently, the selection of a cross-validation strategy plays an important role in this comparison, as the pure neural network is more prone to overfitting. Additional future research should in this regard investigate whether recurrent neural networks and related hybrid models are suited for problems with short time series. A shortcoming is that each time series must be processed sequentially, which reduces the range of possible cross-validation strategies. Using boosted trees or feedforward neural networks, hence solving the forecasting as a regression problem, allows using K-fold cross-validation techniques.

Our research provided some insights into the forecasting problem characteristics that are favorable for the application of hybrid forecasting models. However, additional research is required to verify whether the ability to jointly learn from related time series has the desired effects when parametric sub-models are used. In the space of empty container demand and return forecasting, future research should investigate the possibilities for learning the effect of public holidays jointly from a set of time series. A challenge in this regard is to find appropriate parameterizations. Regression effects of holidays are likely to enter the state space model in the observation process as a time-varying parameter that is predicted by the global neural network. This may cause identifiability issues for state space model parameterizations, while at the same time exposing the model to a greater risk of overfitting. Finally, we propose future research to compare this class of hybrid models against bottom-up models, where parametric models are first estimated locally and a flexible model is estimated globally on the residuals. A benefit is that bottom-up models are less exposed to model misspecifications. Empirical research is required to investigate whether the greater overfitting risk of bottom-up models is a limitation for problems where we expect hybrid models to perform well.

Bibliography

Ben Taieb, S. and Atiya, A. F. (2016). A Bias and Variance Analysis for Multistep-Ahead Time Series Forecasting. *IEEE Transactions on Neural Networks and Learning Systems*, 27(1):62–76.

Ben Taieb, S., Bontempi, G., Atiya, A. F., and Sorjamaa, A. (2012). A review and comparison of strategies for multi-step ahead time series forecasting based on the NN5 forecasting competition. *Expert Systems with Applications*, 39(8):7067–7083.

Boile, M., Theofanis, S., Baveja, A., and Mittal, N. (2008). Regional Repositioning of Empty Containers: Case for Inland Depots. *Transportation Research Record: Journal of the Transportation Research Board*, 2066(1):31–40.

Braekers, K., Janssens, G. K., and Caris, A. (2011). Challenges in Managing Empty Container Movements at Multiple Planning Levels. *Transport Reviews*, 31(6):681–708.

Brouer, B. D., Pisinger, D., and Spoorendonk, S. (2011). Liner Shipping Cargo Allocation with Repositioning of Empty Containers. *INFOR: Information Systems and Operational Research*, 49(2):109–124.

Cheung, R. K. and Chen, C.-Y. (1998). A Two-Stage Stochastic Network Model and Solution Methods for the Dynamic Empty Container Allocation Problem. *Transportation Science*, 32(2):142–162.

Chevillon, G. (2007). Direct multi-step estimation and forecasting. *Journal of Economic Surveys*, 21(4):746–785.

Chou, C.-C., Gou, R.-H., Tsai, C.-L., Tsou, M.-C., Wong, C.-P., and Yu, H.-L. (2010). Application of a mixed fuzzy decision making and optimization programming model to the empty container allocation. *Applied Soft Computing*, 10(4):1071–1079.

Crainic, T. G., Gendreau, M., and Dejax, P. (1993). Dynamic and Stochastic Models for the Allocation of Empty Containers. *Operations Research*, 41(1):102–126.

Cullinane, K. (2010). *International Handbook of Maritime Economics*. Edward Elgar Publishing, Northhamption, MA.

Dahlhaus, R. (2012). *Locally Stationary Processes*, volume 30 of *Handbook of Statistics*, pages 351–414. Elservier, USA, NY, New York.

Dang, Q.-V., Nielsen, I. E., and Yun, W.-Y. (2013). Replenishment policies for empty containers in an inland multi-depot system. *Maritime Economics & Logistics*, 15(1):120–149.

Dang, Q.-V., Yun, W.-Y., and Kopfer, H. (2012). Positioning empty containers under dependent demand process. *Computers & Industrial Engineering*, 62(3):708–715.

Dong, J. and Song, D. (2009). Container fleet sizing and empty repositioning in liner shipping systems. *Transportation Research Part E: Logistics and Transportation Review*, 45(6):860–877.

Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*. Oxford University Press, Oxford, UK, 2nd edition.

Epstein, R., Neely, A., Weintraub, A., Valenzuela, F., Hurtado, S., Gonzalez, G., Beiza, A., Naveas, M., Infante, F., Alarcon, F., Angulo, G., Berner, C., Catalan, J., Gonzalez, C., and Yung, D. (2012). A Strategic Empty Container Logistics Optimization in a Major Shipping Company. *Interfaces*, 42(1):5–16.

Erera, A. L., Morales, J. C., and Savelsbergh, M. (2009). Robust Optimization for Empty Repositioning Problems. *Operations Research*, 57(2):468–483.

Furió, S., Andrés, C., Adenso-Díaz, B., and Lozano, S. (2013). Optimization of empty container movements using street-turn: Application to Valencia hinterland. *Computers & Industrial Engineering*, 66(4):909–917.

Garcia-Ferrer, A., Highfield, R. A., Palm, F., and Zellner, A. (1987). Macroeconomic Forecasting Using Pooled International Data. *Journal of Business & Economic Statistics*, 5(1):53.

Goltsos, T. E., Syntetos, A. A., Glock, C. H., and Ioannou, G. (2022). Inventory – forecasting: Mind the gap. *European Journal of Operational Research*, 299(2):397–419.

Graves, S. C. (1999). A Single-Item Inventory Model for a Nonstationary Demand Process. *Manufacturing & Service Operations Management*, 1(1):50–61.

Grenier, Y. (1983). Time-dependent ARMA modeling of nonstationary signals. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 31(4):899–911.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press, Princeton, NJ, 1st edition.

Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press, Cambridge, UK, 1st edition.

Harvey, D., Leybourne, S., and Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2):281–291.

Hyndman, R. J. and Athanasopoulos, G. (2021). *Forecasting: Principles and Practice*. OTexts, Melbourne, Australia, 3rd edition.

Hyndman, R. J., Koehler, A. B., Ord, J. K., and Synder, R. D. (2008). *Forecasting with Exponential Smoothing: The State Space Approach*. Springer Series in Statistics. Springer, Berlin Heidelberg.

Imai, A. and Rivera, F. (2001). Strategic fleet size planning for maritime refrigerated containers. *Maritime Policy & Management*, 28(4):361–374.

Jula, H., Chassiakos, A., and Ioannou, P. (2006). Port dynamic empty container reuse. *Transportation Research Part E: Logistics and Transportation Review*, 42(1):43–60.

Julier, S. and Uhlmann, J. (2004). Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3):401–422.

Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 80(1):35–45.

Lam, S.-W., Lee, L.-H., and Tang, L.-C. (2007). An approximate dynamic programming approach for the empty container allocation problem. *Transportation Research Part C: Emerging Technologies*, 15(4):265–277.

Lee, C.-Y. and Meng, Q., editors (2015). *Handbook of Ocean Container Transport Logistics: Making Global Supply Chains Effective*, volume 220 of *International Series in Operations Research & Management Science*. Springer International Publishing, Cham.

Lee, C.-Y. and Song, D. (2017). Ocean container transport in global supply chains: Overview and research opportunities. *Transportation Research Part B: Methodological*, 95:442–474.

Lee, L. H., Chew, E. P., and Luo, Y. (2012). Empty Container Management in Multi-port System with Inventory-based Control. *International Journal on Advances in Systems and Measurements*, 5(3 & 4):164–177.

Lee, S. and Moon, I. (2020). Robust empty container repositioning considering foldable containers. *European Journal of Operational Research*, 280(3):909–925.

Levinson, M. (2006). *How the Shipping Container Made the World Smaller and the World Economy Bigger*. Princeton University Press, Princeton, NJ.

Li, J.-A., Leung, S. C., Wu, Y., and Liu, K. (2007). Allocation of empty containers between multi-ports. *European Journal of Operational Research*, 182(1):400–412.

Li, J.-A., Liu, K., Leung, S. C., and Lai, K. K. (2004). Empty container management in a port with long-run average criterion. *Mathematical and Computer Modelling*, 40(1-2):85–100.

Lim, B. and Zohren, S. (2021). Time-series forecasting with deep learning: A survey. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2194):1–14.

Long, Y., Lee, L. H., and Chew, E. P. (2012). The sample average approximation method for empty container repositioning with uncertainties. *European Journal of Operational Research*, 222(1):65–75.

Martius, C., Kretschmann, L., Zacharias, M., Jahn, C., and John, O. (2022). Forecasting worldwide empty container availability with machine learning techniques. *Journal of Shipping and Trade*, 7(1):19.

Montero-Manso, P. and Hyndman, R. J. (2021). Principles and algorithms for forecasting groups of time series: Locality and globality. *International Journal of Forecasting*, 37(4):1632–1653.

Ng, C. T., Song, D.-P., and Cheng, T. C. E. (2012). Optimal Policy for Inventory Transfer Between Two Depots With Backlogging. *IEEE Transactions on Automatic Control*, 57(12):3247–3252.

Notteboom, T., Pallis, A., and Rodrigue, J.-P. (2021). *Port Economics, Management and Policy*. Routledge, London, UK, 1st edition.

Olivo, A., Di Francesco, M., and Zuddas, P. (2013). An optimization model for the inland repositioning of empty containers. *Maritime Economics & Logistics*, 15(3):309–331.

Poncela, M., Poncela, P., and Perán, J. R. (2013). Automatic tuning of Kalman filters by maximum likelihood methods for wind energy forecasting. *Applied Energy*, 108:349–362.

Powell, W. B. (2011). *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley Series in Probability and Statistics. John Wiley & Sons, Hoboken, NJ, 2nd edition.

Powell, W. B. (2019). A unified framework for stochastic optimization. *European Journal of Operational Research*, 275(3):795–821.

Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. John Wiley & Sons, Hoboken, NJ, 1st edition.

Rangapuram, S. S., Seeger, M. W., Gasthaus, J., Stella, L., Wang, Y., and Januschowski, T. (2018). Deep state space models for time series forecasting. In *Advances in Neural Information Processing Systems (NIPS), Montreal, Canada*, pages 1–13.

Ray, W. D. (1980). The Significance of Correlated Demands and Variable Lead Times for Stock Control Policies. *The Journal of the Operational Research Society*, 31(2):187.

Rodrigue, J.-P. (2020). *The Geography of Transport Systems*. Routledge, New York, NY, 5th edition.

Salinas, D., Flunkert, V., Gasthaus, J., and Januschowski, T. (2020). DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3):1181–1191.

Shapiro, A. and Nemirovski, A. (2005). On Complexity of Stochastic Programming Problems. In Jeyakumar, V. and Rubinov, A., editors, *Continuous Optimization*, pages 111–146. Springer US, New York, NY.

Shintani, K., Konings, R., and Imai, A. (2010). The impact of foldable containers on container fleet management costs in hinterland transport. *Transportation Research Part E: Logistics and Transportation Review*, 46(5):750–763.

Shu, J. and Song, M. (2014). Dynamic Container Deployment: Two-Stage Robust Model, Complexity, and Computational Results. *INFORMS Journal on Computing*, 26(1):135–149.

Song, D. (2007). Characterizing optimal empty container reposition policy in periodic-review shuttle service systems. *Journal of the Operational Research Society*, 58(1):122–133.

Song, D. and Dong, J. (2008). Empty Container Management in Cyclic Shipping Routes. *Maritime Economics & Logistics*, 10(4):335–361.

Song, D. and Dong, J. (2011). Flow balancing-based empty container repositioning in typical shipping service routes. *Maritime Economics & Logistics*, 13(1):61–77.

Song, D. and Dong, J. (2015). Empty Container Repositioning. In Lee, C.-Y. and Meng, Q., editors, *Handbook of Ocean Container Transport Logistics*, pages 163–208. Springer, Zürich, Switzerland.

Song, D. and Dong, J. (2022). *Modelling Empty Container Repositioning Logistics*. Springer International Publishing, Cham, Switzerland.

Song, D., Dong, J., and Roe, M. (2010). Optimal container dispatching policy and its structure in a shuttle service with finite capacity and random demands. *International Journal of Shipping and Transport Logistics*, 2(1):44–58.

Song, D. and Zhang, Q. (2010). Optimal inventory control for empty containers in a port with random demands and repositioning delays. In Cullinane, K., editor, *International Handbook of Maritime Economics*, pages 301–321. Edward Elgar Publishing, Cheltenham, UK.

Young Yun, W., Mi Lee, Y., and Seok Choi, Y. (2011). Optimal inventory control of empty containers in inland transportation system. *International Journal of Production Economics*, 133(1):451–457.

Zhang, B., Ng, C. T., and Cheng, T. C. E. (2014). Multi-period empty container repositioning with stochastic demand and lost sales. *Journal of the Operational Research Society*, 65(2):302–319.

Paper A

Inland empty container inventory management with Markov decision processes

Publication details: Submitted to *Transportation Research Part E: Logistics and Transportation Review*.

Inland empty container inventory management with Markov decision processes

Benedikt Sommer[1,2], Trine Krogh Boomsma[3], Klaus Kähler Holst[1]

[1]A. P. Møller-Mærsk, Copenhagen, Denmark
[2]Technical University of Denmark, Kongens Lyngby, Denmark
[3]University of Copenhagen, Copenhagen, Denmark

June 30, 2023

Abstract Container logistics companies store empty containers at depots from where they are delivered to export customers to satisfy their demand for shipping freight. Depots further satisfy the needs of import customers who received freight and must now return an empty container. A crucial problem is that demands and returns are often imbalanced, requiring the logistics companies to reposition empty containers between depots. In this paper we formulate the single depot empty container allocation model as a capacitated multiple supplier periodic review inventory control problem with lost sales and deterministic lead-times. The developed discrete-time Markov Decision Process extends previous allocation models by explicitly accounting for varying lead-times and costs of different transportation modes for receiving empty containers from other depots. Optimal policies for in- and decreasing the depot's container stock level are learned for serial and cross-sectional dependent demand and return processes. We quantify the costs of ignoring these dependencies in simulation studies and demonstrate their existence on real-world data.

Keywords Inventory control, Dependent exogenous processes, Multiple supplier sourcing, Optimal policy learning, Time-inhomogeneous Markov decision process

1 Introduction

The redistribution of empty containers from import- to export-dominant regions is an essential task of a container shipping company. Generally known as empty container repositioning (ECR), the redistribution of empty containers is unavoidable since structural trade imbalances let empty containers accumulate in import-dominant regions while export-dominant regions experience container shortages. The processes are complex and costly, with estimated annual repositioning costs exceeding $10 billion for the global container shipping market Song et al. (2005); Notteboom et al. (2021). Players in the shipping market seek cost reductions through innovative repositioning strategies to increase profitability in a competitive environment Zhang et al. (2014). Of particular interest for cost optimization are inland operations where multiple options are available for the transportation of containers in between ports and inland depots. Other than road transportation by truck, many regions have existing rail and waterway infrastructure with great optimization potential for the utilization of less flexible, but more cost-effective and environmentally friendly, transportation options.

This paper contributes to the ECR literature by proposing a novel empty container allocation model for the dynamic operations of an inland container depot. The depot is assumed to be connected to a sea port through multiple transportation modes, each mode following its own schedule. Our main contribution is the extension of the existing empty container allocation models by accounting for the different cost and delivery time parameters of each transportation mode. We formulate the problem, which is a particular instance of a multiple supplier inventory control problem (Svoboda et al., 2021), as a discrete-time Markov Decision Process (MDP) and demonstrate that an optimal policy can be obtained for realistic problem sizes under reasonable assumptions. As an additional contribution, justified by the operational data of a container shipping company, we investigate the existence of serial and cross-sectional dependence of the empty container demand (by export customers) and

return (from import customers) processes. To the best of our knowledge, the investigation of cross-sectional dependencies are new to the ECR literature.

Our work builds on the understanding that formulating and solving a single multi-modal, multi-commodity, stochastic and dynamic model is infeasible when thousands of depots are used for empty container storage Braekers et al. (2011). This explains why the ECR problem is typically handled by approximations, e.g. replacing stochastic quantities by deterministic estimates (Song and Dong, 2015), or decomposition methods, e.g. first solving a global repositioning problem between ports and then formulating a second problem that distributes empty containers from ports to inland depots (Braekers et al., 2011). Indeed, the decomposition into ocean and inland systems can be viewed as a natural consequence of the different transportation modalities and their characteristics. The connection between these systems is established by sea ports in which containers can be transferred between container vessels, trucks, trains and barges as illustrated by the stylized system in Figure 1. Another interpretation of the landside operations is that ports serve as sinks for outflowing empty containers and sources for inflowing containers. Throughout the paper, we follow this decomposition strategy with the assumption that global repositioning planning ensures sufficient empty container stock levels at ports in export-dominant regions and sufficient storage capacities in import-dominant regions. For a more complete description of the ECR problem and solutions at a global scale, we refer the interested reader to Song and Dong (2015).

Modeling the operations for only an isolated multi-modal inland network of a container logistics company remains challenging. Complexities are primarily induced by the stochastic nature of future empty container demand of export customers, future container returns by import customers, transportation delays, capacity constraints, equipment failures and the model's time resolution and planning horizon. To alleviate the related computational challenges, the operational ECR inland planning can further be decomposed into two separate optimization problems (Braekers et al., 2011). An empty container allocation model is used initially to determine the distribution of containers across the depots in the network, followed by a routing model that minimizes transportation costs subject to the fulfillment of the previously estimated container allocations. Examples that follow this decomposition for inland operations are found in Crainic et al. (1993); Olivo et al. (2005). The allocation and routing decomposition is also applied on the global scale, with examples for ocean repositioning presented in Feng and Chang (2008); Chou et al. (2010); Zhang et al. (2014).

The allocation model often relies on inventory theory to derive the ordering and releasing quantities for a single container depot, i.e. the number of containers to be exchanged with ports and other inland depots. The non-standard inventory control problem, in which positive (export customer requests) and negative (import customer returns) demand occurs, has been solved using several modelling techniques such as optimal control theory, dynamic programming or simulation (Song and Dong, 2015). An important aspect of the allocation-routing decomposition is that sufficient information about available transportation options with ports and other inland depots should be incorporated in the allocation model. It follows that the routing model cannot fulfil container requests in time if the depot's respective allocation model assumes transportation times too short to be accommodated by the inland transportation network.

It is precisely this understanding in combination with the different transportation times of truck, train and barge in inland networks that requires an extension of existing empty container allocation models. In this paper, a MDP captures the dynamic operations of an inland container depot. The Markov process allows us to accommodate a fine time resolution and a long planning horizon while accounting for uncertainty of serial and cross-sectional dependent future demands and returns. The paper is structured as follows. Section 2 illustrates how the dynamic depot operations can be viewed as a non-standard multiple supplier inventory control problem. By reviewing existing inventory models within and outside the ECR context, we identify the research gaps. Section 3 develops the MDP for the introduced inventory system and presents propositions to remove states and actions which are not visited or selected under an optimal policy, respectively. In Section 4 we first illustrate the reductions in total operating costs obtained from modelling transportation modes with longer delivery times but lower costs. We proceed to examine the impact on operating costs when policies are estimated with misspecified processes, where serial- and cross-sectional dependencies for the exogenous empty container demands and returns go undetected. The existence of these dependencies and consequences of their misspecification are showcased for real-world data in Section 5. The paper is concluded in Section 6.

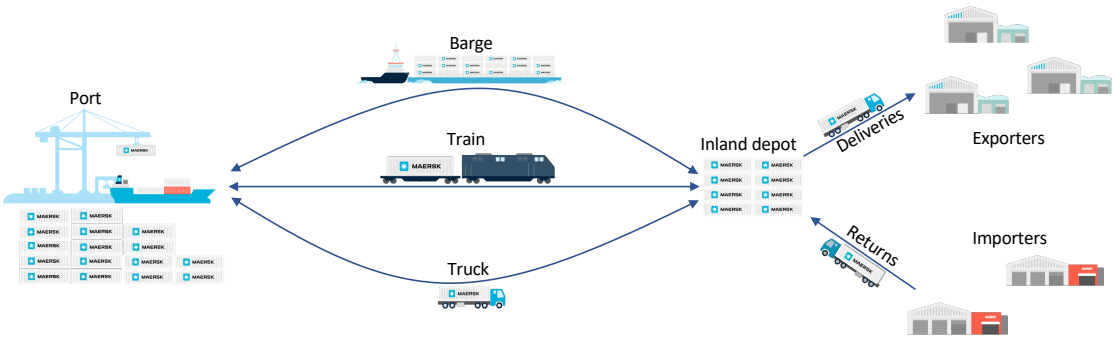

Figure 1: Empty container flows in an inland network with a single port and inland depot. The empty container flows between the port and depot are bi-directional, whereas the empty container exchange with customers is uni-directional with import customers returning empty containers and export customers receiving empty containers.

2 An inventory model for the dynamic depot operations

We demonstrate the applicability of our inventory control model by considering the stylized inland transportation network in Figure 1. The network consists of geographically distant empty container storage facilities at an inland depot and sea port, and import and export customers in the region surrounding the inland depot. The inland depot manager, subsequently referred to as the *decision maker*, is responsible for simultaneously satisfying the empty container demand of export customers and accommodating the empty container returns of import customers. Future empty container demand and returns are uncontrollable and unknown to the decision maker, hence are stochastic exogenous variables. In line with the majority of the ECR literature, we assume that a single container type is returned and demanded by customers. The objective of the decision maker is to satisfy demand and manage returns subject to minimizing costs, with penalties occurring for lost sales and exceedance of the depot's storage capacity.

To optimize operating costs, the decision maker can request empty containers from the port and thereby avoid lost sales when empty container demand exceed returns. On the contrary, inventory holding costs and storage exceedance penalties are reduced by sending empty containers from the inland depot to the port. We refer to decisions that increase inventory as in-positioning, and let out-positioning refer to decision that reduce inventory. Both types of decisions are made periodically at discrete time points. For all points in time, only past empty container demand and returns are known. Thus, the decision maker can place repositioning orders in anticipation of future empty container demand and returns, although with knowledge about current inventory stock levels and previous in-positioning orders that have not yet arrived.

As an illustrative example, the decision maker can choose to reposition containers via truck, train and barge. Trucks are the fastest transportation mode with the highest unit costs, whereas barges are the slowest but also the least expensive. We assume that the available transportation capacities and related transportation times are deterministic and known during the planning period. Furthermore, we assume that the decision maker cannot adjust previous empty container orders that have not yet arrived, i.e. it is not possible to order many containers with slow transportation modes only to later reduce such orders. Due to the difference in transportation times, the in-positioning decision problem can be viewed as a supplier selection problem. Supply options with longer delivery times may generally be preferred because of their lower costs. However, the longer delivery times of empty containers induce a higher risk of running out of equipment (lost sales) or receiving more returns than required.

2.1 Multiple supplier inventory control

Depot operations can be described as a capacitated multiple supplier periodic review inventory control problem (Minner, 2003) with lost sales and deterministic lead times for requested empty containers. The inventory varies dynamically due to exogenous demand and supply, i.e. empty container deliveries and returns, and is controlled through in- and out-positioning decisions. An out-positioning

decision can be viewed as a disposal option, and the in-positioning decisions are replenishment options with varying lead times. The described problem is non-standard because stochastic variables both increase and decrease the inventory level at the depot. Nevertheless, our problem shares some similarities with remanufacturing or product recovery systems in which product returns represent another mode of supply in addition to replenishment options (Govindan et al., 2015). The inventory decision problem is sequential due to the dependence of cost-optimal future inventory and decision trajectories on the current decision. With this sequential nature of the problem, the multistage stochastic optimization problem is best formulated as an MDP. Exact analytical solutions have only been obtained for special cases of the multiple supplier inventory problem with convex cost functions and monotonic inventory transition functions. For example, Whittemore and Saunders (1977) consider the standard inventory problem without disposal options and prove that policies with a base stock structure are optimal for two supplier systems with consecutive lead times, serially independent demand and backlogging. The authors proceed to consider systems with non-consecutive ordering lead times and show that the optimal ordering policies are no longer simple functions of the requested inventory replenishment items. For inventory systems with lost sales, analytical solutions are challenged by having to keep track of the available inventory and previously ordered items that have not yet arrived, enlarging the state space. Thus, the analytical solutions reviewed in Svoboda et al. (2021) rely on stylized assumptions such as backlogging or serially independent demand.

To avoid such stylized assumptions, we solve the multistage stochastic optimization problem numerically. We apply stochastic dynamic programming (SDP) and benefit from the linearly increasing complexity with respect to the number of decision epochs. SDP may, however, suffer from the curse of dimensionality, i.e. the algorithm's exponential complexity in the number of states and actions. As previously discussed, any multiple supplier inventory model with lost sales has to track the on-order items of previous placed orders. The size of the state space increases as a consequence exponentially with the maximum lead time among all supply options (Svoboda et al., 2021). Larger problem instances therefore often require approximation (Fang et al., 2013) or heuristic methods (Zipkin, 2008) to obtain policy estimates. Our decision to proceed with an exact solution algorithm is twofold. On one hand, we provide an algorithm that can solve small and medium-sized problems to optimality. On the other hand, our algorithm can serve as a starting point for approximation and heuristic methods for larger, unsolvable, instances.

A property of many inventory systems is for demand to be serially dependent. It is well understood that ignoring the presence of positive autocorrelation will lead to underestimated safety stock levels, while ignored negative autocorrelation will cause overestimation Ray (1980); Graves (1999). Indeed, forecasts are an essential component in many inventory control systems, where past demand observations provide information on future values. While there are many ways to incorporate demand forecast information into inventory control models (Goltsos et al., 2022), for the MDP we follow the common approach of including an exogenous stochastic demand process. This approach can likewise account for the serial dependence of empty container returns, although at the expense of further enlarging the state space. The following literature review of ECR problems shows that accounting for serial dependence has so far received little attention.

2.2 Research gaps

Few papers formulate the empty container inventory management problem of a single container depot as an MDP. To the best of our knowledge, Li et al. (2004) are the first to model the dynamic port storage operations with a discrete-time stochastic inventory model. The main differences to our work are that in-positioning decisions materialize immediately with zero lead time and that short-term container leasing is possible with infinite capacity, which effectively allows for penalized backlogging. Empty container demand and returns are assumed to be serially and mutually independent. These assumptions allow the authors to show the existence of an optimal two threshold-type policy for the discounted finite-horizon cost minimization problem. Clearly, there are limitations to the proposed model. As noted by Zhang et al. (2014), short-term leasing options with known future capacities and prices are rather limited in reality. The recent COVID-19 pandemic and the related global container shortage exemplifies this, cf. (Toygar et al., 2022). Zhang et al. (2014) replace the possibility of short-term leasing by the occurrence of lost sales if demand exceeds the inventory that remains upon out-positioning. While it is standard to assume that empty containers leave the inventory immediately for out-positioning decisions, there are several variations for in-positioning decisions. For example, Young Yun et al. (2011) assume a fixed four week lead time for in-positioning decisions in their inventory model of a deficit port. Moreover, the MDP formulation of Song and Zhang

(2010a) extends the model in Li et al. (2004) by assuming a fixed one decision period delay of in-positioned empty containers. The MDP formulation, however, is not generalized to longer lead times reflecting the varying distances to surplus regions or the different transportation times of distinct transportation modes.

A common limitation in Li et al. (2004); Song and Zhang (2010a); Zhang et al. (2014) is that empty container demands and returns are serially independent. This assumption is appealing from a modeling perspective, but may not reflect reality. Indeed, Crainic et al. (1993) noted that accurate forecasts of the complex empty container demand and return processes are critical for the successful application of the dynamic and stochastic empty container allocation model. Approximations to the stochastic processes has, with a few exceptions, only received little attention in the empty container allocation literature. For example, Dang et al. (2012) considered a discrete-time inventory control model, assuming that the empty container demands and returns follow two mutually independent AR(1) processes. For the continuous-time empty container allocation model of a single port, Song and Zhang (2010b) assume that the net flow, i.e. difference between empty container demand and returns, follows a two-state Markov chain. In both cases, modeling assumptions are made without support from real-world data. In contrast, we estimate the stochastic processes from real-world data in Section 5. Finally, existing works do not consider cross-sectional dependence between demands and returns, in spite of the assumption that returned empty containers have a one period lead time before they can be reused. Positive cross-sectional dependence can reduce target stock levels because a high number of returns tend to follow large empty container demands, whereas negative dependence requires higher stock levels. As we demonstrate in the simulation study of Section 4, the estimated policies may perform poorly if cross-sectional dependence goes undetected.

3 Markov decision process framework

This section formulates the MDP for the inventory problem, following the terminology and notation of Puterman (1994). In most cases, we make the assumptions of the reviewed MDPs for empty container inventory management, but highlight the assumptions that are unique to our MDP. The notation follows the convention that stochastic variables are denoted by upper-case characters. Realizations of stochastic variables and known parameters of the MDP are denoted by lower-case characters. We use the short-form $P(x|y)$ to denote the conditional probability $P(X = x|Y = y)$ of the random variables X and Y. Similarly, $F(x|y)$ denotes the conditional cumulative probability $P(X \leq x|Y = y)$. We further use the superscripts $+$ and $-$ to denote actions and stochastic variables that increase and decrease the inventory level, respectively.

Decisions are made at decision epochs that correspond to the beginning of a time period. At each decision epoch the decision maker has information about the current inventory position, all previously ordered but not yet arrived empty containers and all previous empty container demands and returns. The decision maker places out- and in-positioning orders, and out-positioned containers leave the depot immediately. All in-positioning orders arrive at the end of a decision period and have lead times of at least one decision period, i.e. empty containers ordered at the current decision epoch will at the earliest be available at the following epoch. Transportation capacities and lead times for the in- and out-positioning decisions are assumed to be known. Empty container returns also have one decision period lead time, and thus, demand in the current decision period has to be satisfied with the inventory that remains after the out-positioning decision has been taken. The decision maker accrue costs according to the chosen in- and out-positioning decisions, holding costs and penalty fees for loosing sales and exceeding the finite storage capacity.

3.1 Decision epochs and planning horizon

The periodic inventory review scheme involves a discrete-time decision process, and the inland transportation in- and out-positioning options suggest a daily time-resolution due to the corresponding transportation times. Hence, in- and out-positioning decisions are taken at decision epochs $t = 0, 1, \ldots, T-1$, with t being a day within the planning period of finite length T. Following Puterman (1994), no decision is made at epoch T and a terminal cost is accrued. A finite horizon is chosen to allow for policy estimation in a time-inhomogeneous MDP by backwards dynamic programming.

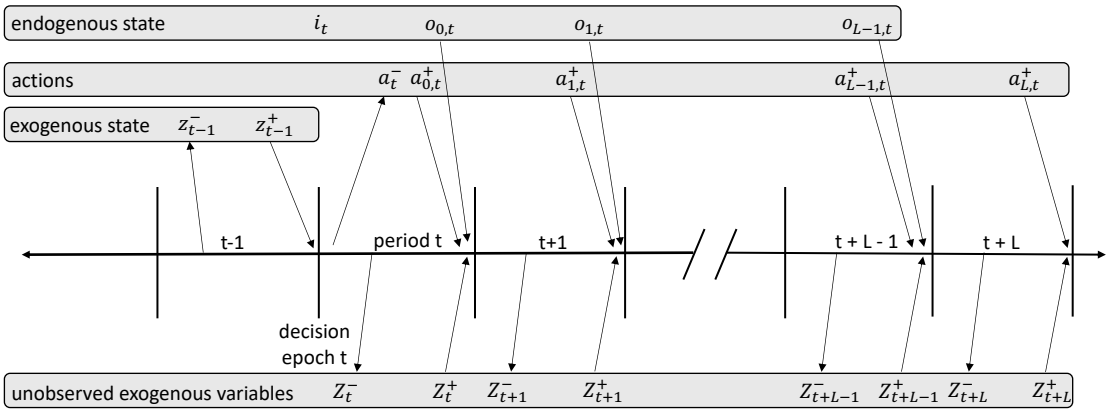

Figure 2: Illustration of states and actions at epoch t. The exogenous state is illustrated with $X_t = (Z_{t-1}^-, Z_{t-1}^+)$.

3.2 Action space

At each decision epoch, the decision maker can choose among two complementary in- and out-positioning options. The first option is to out-position empty containers $A_t^- \in \mathcal{A}^- = \{0, 1 \ldots, \bar{A}^-\}$ that immediately leave the depot, where $\bar{A}^-$ denotes the out-position transportation capacity. The second option is to place empty container in-positioning orders $A_{l,t}^+ \in \mathcal{A}_l^+ = \{0, 1, \ldots, \bar{A}_l^+\}$ with varying lead times $l = 0, 1, \ldots, L$. The convention for $A_{l,t}^+$ is that containers that are requested at epoch t will arrive at the end of decision period $t + l$. A lead time of 0 means that the requested containers will only arrive at the end of the current decision period. Contrary to Li et al. (2004); Zhang et al. (2014), we do not allow for emergency orders that arrive immediately. The full decision vector at time t is denoted by $A_t \in \mathcal{A} = \mathcal{A}^- \times \mathcal{A}_0^+ \times \mathcal{A}_1^+ \cdots \times \mathcal{A}_L^+$, and the dynamics of all decisions are illustrated in Figure 2.

3.3 State space

We first introduce endogenous states. Subsequently, we discuss exogenous states that are required for serially dependent empty container demand and return processes.

Endogenous state space The endogenous state space includes two components. The first component is the empty container inventory level I_t that is known at decision epoch $t \in \mathcal{T}$. We assume lost sales and a finite storage capacity $\bar{I}$ such that $I_t \in \mathcal{I} = \{0, 1, \ldots, \bar{I}\}$. The second component keeps track of the previously ordered but not yet arrived empty containers. With in-positioning order lead times $l = 0, 1 \ldots, L-1$, we let $O_{l,t} \in \mathcal{O}_l$ denote the total number of previously ordered containers at time t to arrive at the end of decision period $t + l$, where a lead time of 0 implies that the containers arrive at the end of the current decision period. Since the maximum lead time is L periods, we only keep track of previously ordered empty containers that arrive in the coming $L-1$ periods. Throughout the remainder of the paper we refer to these states as *on-order states* and their arrival dynamics are depicted in the process diagram in Figure 2. Their transition through time is deterministic and is governed by

$$O_{l-1,t+1} = O_{l,t} + A_{l,t}^+, \; l = 1, \ldots, L-1, \tag{1}$$

and

$$O_{L-1,t+1} = A_{L,t}^+. \tag{2}$$

From the transition functions, it follows that each on-order state takes values in $\mathcal{O}_l = \{0, 1, \ldots, \sum_{i=l+1}^{L} \bar{A}_l^+\}$. Together with the inventory state, we let $S_t^\dagger = (I_t, O_{0,t}, O_{1,t}, \ldots, O_{L-1,t})$ denote the endogenous state which takes values in $\mathcal{S}^\dagger = \mathcal{I} \times \mathcal{O}_0 \times \mathcal{O}_1 \times \ldots \times \mathcal{O}_{L-1}$, where $\times$ denotes the Cartesian product.

The inventory transition function depends on the on-order state $O_{0,t}$ and is defined as

$$I_{t+1} = \min(\max(I_t - A_t^- - Z_t^-, 0) + A_{0,t}^+ + O_{0,t} + Z_t^+, \bar{I}), \tag{3}$$

where Z_t^- and Z_t^+ are the stochastic discrete-valued empty container demands and returns that occur in decision period t, i.e., they are unknown when the decision A_t is made at epoch t. The inner max operator reflects the lost sales assumption, that out-positioned containers reduce the available inventory immediately and that returned empty containers have a one decision period lead time before they can be reused. The outer min operator accounts for the finite capacity assumption of the storage.

Exogenous state space The exogenous states contain the information about serially dependent variables that affect the transition probabilities of the inventory state due to the dependence of I_{t+1} on Z_t^- and Z_t^+. We assume that the joint distribution of empty container demand Z_t^- and returns Z_t^+ in decision period t depends on a subset of the variables $Z_0^-, Z_0^+, \ldots, Z_{t-1}^-, Z_{t-1}^+$, which we denote by the vector X_t. Without further specifying the dependence structure of the stochastic process on its own history, we let $\mathcal{X}$ denote the exogenous state space that contains the required lags of empty container demands and returns. Since Z_t^- and Z_t^+ are unknown when the decision at epoch t is made, the exogenous state X_t can only include previous observations up to and including Z_{t-1}^- and Z_{t-1}^+. The complete state vector of the systems becomes $S_t = (S_t^\dagger, X_t) \in \mathcal{S} = \mathcal{S}^\dagger \times \mathcal{X}$.

3.4 Immediate costs

We assume fixed in- and out-position ordering costs λ_l^+, λ^-. The stochastic transition costs from endogenous state $S_t^\dagger$ to $S_{t+1}^\dagger$ under action A_t are defined for $t = 0, \ldots, T-1$ as

$$\begin{aligned} C_t(S_t^\dagger, A_t, Z_t^-, Z_t^+) =& \lambda^- A_t^- + \left(\sum_{l=0}^{L} \lambda_l^+ A_{l,t}^+\right) \\ &+ \lambda_l \max(Z_t^- - (I_t - A_t^-), 0) + \lambda_h \max(I_t - A_t^- - Z_t^-, 0) \\ &+ \lambda_p \max(\max(I_t - A_t^- - Z_t^-, 0) + Z_t^+ + A_{0,t}^+ + O_{0,t} - \bar{I}, 0), \end{aligned}$$

where λ_l, λ_p and λ_h are the lost-sales penalty, storage exceedance penalty and holding costs. Lost sales $\max(Z_t^- - (I_t - A_t^-), 0)$ reflect the assumption that demand has to be satisfied with the inventory that remains after the out-positioned containers leave the depot immediately. Similarly, holding costs have to be paid for inventory that remains after demand has been satisfied, i.e. $\max(I_t - A_t^- - Z_t^-, 0)$. The last term reflects penalty costs for arriving containers that exceed the storage capacity $\bar{I}$. The costs are stochastic since Z_t^- and Z_t^+ are not yet observed when choosing action A_t.

For a fixed action $A_t = a_t$ that is selected while being in state $S_t = s_t$, the expected costs are

$$\begin{aligned} &\mathrm{E}\Big[C_t(S_t^\dagger, A_t, Z_t^-, Z_t^+) | S_t = s_t, A_t = a_t\Big] = \\ &\lambda^- a_t^- + \left(\sum_{l=0}^{L} \lambda_l^+ a_{l,t}^+\right) + \lambda_p \mathrm{E}\big[\max(\max(i_t - a_t^- - Z_t^-, 0) + Z_t^+ + a_{0,t}^+ + o_{0,t} - \bar{I}, 0) | X_t = x_t\big] \\ &+ \mathrm{E}\big[\lambda_h \max(i_t - a_t^- - Z_t^-, 0) + \lambda_l \max(Z_t^- - (i_t - a_t^-), 0) | X_t = x_t\big], \end{aligned}$$

where we use that Z_t^- and Z_t^+ may depend on some exogenous variables X_t, but are otherwise independent of the endogenous states and selected actions.

The terminal costs that accrue in the final period $t = T$ when the controllable part of the system occupies state $S_T^\dagger$ are a design choice. We propose

$$\begin{aligned} C_T(s_T^\dagger) =& \mathrm{E}\Big[\lambda_p \max(\max(I_T - Z_T^-, 0) + Z_T^+ - \bar{I}, 0) | S_T^\dagger = s_T^\dagger\Big] \\ &+ \mathrm{E}\Big[\lambda_h \max(I_T - Z_T^-, 0) + \lambda_l \max(Z_T^- - I_T), 0) | S_T^\dagger = s_T^\dagger\Big] \end{aligned} \tag{4}$$

to assign appropriate costs to inventory positions that dependent on whether the depot experiences an empty container deficit or surplus in the long-run, hence the costs are in expectation of Z_T^- and Z_T^+. The expectation is, however, not conditional on the exogenous state X_T because the value of policies with and without exogenous state spaces are compared in Section 4. Comparability is here ensured by only conditioning on the endogenous state. Due to an additional comparison of policies with varying transportation modes, hence different on-order state spaces, the terminal costs are also independent of the on-order states. Eventually, one recognizes the terminal costs to mirror the immediate costs that accrue when no previous ordered empty containers arrive and no in- and out-positioning decisions are made.

3.5 Transition probabilities

The transition functions for the endogenous states and the exogenous process model for the empty container demands and returns fully determine the one-step transition probability $P(s_{t+1}|s_t, a_t)$ from state s_t to s_{t+1} under action a_t. The transition probability has the factorization

$$\begin{aligned} P(s_{t+1}|s_t, a_t) &= P(s^\dagger_{t+1}, x_{t+1}|s_t, a_t) \\ &= P(s^\dagger_{t+1}|x_{t+1}, s_t, a_t)P(x_{t+1}|s_t, a_t) \\ &= P(s^\dagger_{t+1}|x_{t+1}, s_t, a_t)P(x_{t+1}|x_t), \end{aligned}$$

where $P(x_{t+1}|s_t, a_t) = P(x_{t+1}|x_t)$ is due to the exogeneity assumption for the empty container demand and return process. The exogenous transition probability $P(x_{t+1}|x_t)$ is obtained from the demand and return process. The deterministic transition function of the on-order states can be used to express the endogenous factor as

$$\begin{aligned} P(s^\dagger_{t+1}|x_{t+1}, s_t, a_t) &= P(i_{t+1}|x_{t+1}, s_t, a_t) \prod_{l=0}^{L-1} P(o_{l,t+1}|x_{t+1}, s_t, a_t) \\ &= P(i_{t+1}|x_{t+1}, s_t, a_t) \cdot 1\{o_{L-1,t+1} = a^+_{L,t}\} \cdot \prod_{l=0}^{L-1} 1\{o_{l-1,t+1} = o_{l,t} + a^+_{l,t}\}, \end{aligned}$$

where $1\{\cdot\}$ denotes the indicator function. The factor simplifies to

$$\begin{aligned} P(s^\dagger_{t+1}|x_{t+1}, s_t, a_t) = 1\{i_{t+1} = \max(i_t - a^-_t - z^-_t, 0) + a^+_{0,t} + o_{0,t} + z^+_t, \bar{I})\} \\ \cdot 1\{o_{L-1,t+1} = a^+_{L,t}\} \cdot \prod_{l=0}^{L-1} 1\{o_{l-1,t+1} = o_{l,t} + a^+_{l,t}\}, \end{aligned}$$

if x_t includes the empty container demands z^-_{t-1} and returns z^+_{t-1} of the previous period, since the conditioning set of $P(i_{t+1}|x_{t+1}, s_t, a_t)$ contains all the information for i_{t+1}.

3.6 State-dependent action space

The action space is naturally constrained by the physical properties of the inventory system. In state S_t, the out-positioning decision A^-_t is constrained by the available inventory I_t, i.e. $0 \leq A^-_t \leq I_t$, because it is not possible to out-position more containers than are currently at the depot. We further introduce the mutual exclusivity constraint,

$$A^-_t A^+_{0,t} = 0,$$

i.e., $A^-_t = 0$ if $A^+_{0,t} > 0$ and $A^+_{t,0} = 0$ if $A^-_t > 0$, to prohibit the estimation of policies for which empty containers are out-positioned to only receive some containers again at the end of the decision period. We exclude such actions because they are unreasonable in practice with a daily time-resolution, but also to reduce the action space. The joint action space for A^-_t and $A^+_{0,t}$ becomes

$$\mathcal{A}_{0,S_t} = \{(a^-_t, a^+_{0,t}) \mid a^-_t a^+_{0,t} = 0, a^-_t \in \{0, 1, \ldots, \min(I_t, \bar{A}^-)\}, a^+_{0,t} \in \mathcal{A}^+_0\}.$$

The endogenous state space introduces additional restrictions on the in-positioning actions because each on-order state $O_{l,t}$ takes values in $\mathcal{O}_l$ for $l = 0, 1, \ldots, L-1$. From the on-order state transition functions (1) and (2) follows

$$\mathcal{A}^+_{l,S_t} = \{a^+_{l,t} | a^+_{l,t} + O_{l,t} \in \mathcal{O}_{l-1}, a^+_{l,t} \in \mathcal{A}^+_l\}, \text{ for } l = 1, \ldots, L-1,$$

and

$$\mathcal{A}^+_{L,S_t} = \{a^+_{L,t} | a^+_{L,t} \in \mathcal{O}_{L-1}, a^+_{L,t} \in \mathcal{A}^+_l\}$$

The action space for $A^+_{0,t}$ remains unchanged because the action does not affect any on-order state. The full state-dependent action space is thus $\mathcal{A}_{S_t} = \mathcal{A}_{0,S_t} \times \mathcal{A}^+_{1,S_t} \times \mathcal{A}^+_{2,S_t} \times \cdots \times \mathcal{A}^+_{L,S_t}$.

3.7 Objective function and optimality equations

Following Puterman (1994) we let $d_t(S_t) : \mathcal{S} \to \mathcal{A}_{S_t}$ denote the decision rule for selecting an action in each state under policy $\pi = (d_1, d_2, \ldots, d_{T-1}) \in \Pi$, with the standard assumption of Π being the set of Markovian and deterministic policies. The objective of the decision maker is to find a policy π that minimizes the expected costs over the planning horizon

$$V_\pi(s_0) = \mathrm{E}_\pi \left[\sum_{t=0}^{T-1} C_t(S_t^\dagger, d_t(S_t), Z_t^-, Z_t^+) + C_T(S_T^\dagger) \middle| S_0 = s_0 \right],$$

with initial state $s_0 \in \mathcal{S}$. Here the expectation is taken with respect to the distribution induced by the policy π. The optimal policy is obtained as

$$\pi^\star = \operatorname*{argmin}_{\pi \in \Pi} V_\pi(s_0)$$

and is numerically calculated by minimizing (5) with backwards dynamic programming, starting in

$$V_T(S_T) = C_T(S_T^\dagger),$$

and solving the Bellman equation

$$\begin{aligned} V_t(s_t) &= \min_{a_t \in \mathcal{A}_{s_t}} \mathrm{E}\Big[C_t(S_t^\dagger, A_t, Z_t^-, Z_t^+) + V_{t+1}(S_{t+1}) | S_t = s_t, A_t = a_t \Big] \\ &= \min_{a_t \in \mathcal{A}_{s_t}} \mathrm{E}\Big[C_t(s_t^\dagger, a_t, Z_t^-, Z_t^+) | X_t = x_t \Big] + \mathrm{E}[V_{t+1}(S_{t+1}) | S_t = s_t, A_t = a_t] \\ &= \min_{a_t \in \mathcal{A}_{s_t}} \sum_{z_t^-, z_t^+} p(z_t^-, z_t^+ | x_t) C_t(s_t^\dagger, a_t, z_t^-, z_t^+) + \sum_{s_{t+1} \in \mathcal{S}} p(s_{t+1} | s_t, a_t) V_{t+1}(s_{t+1}) \end{aligned}$$

recursively for each decision epoch $T-1, T-2, \ldots, 0$.

3.8 Pruning state and action spaces

The following propositions can be used to reduce the state and action spaces of the formulated MDP and improve the efficiency of the backward dynamic programming algorithm. All proofs are found in the appendix.

Proposition 1 reduces the endogenous state space for on-order states by considering the accrued penalty when previously ordered containers arrive and exceed the storage capacity. The proof verifies the intuition that no optimal solution could have more ordered containers arriving at the end of a decision period than can be stored at the depot.

Proposition 1. *Let $\lambda_p > 0$. Any on-order states that satisfy*

$$O_{t,l} > \bar{I}, \quad \textit{for } l = 0, 1 \ldots, L-1$$

cannot occur in an optimal solution.

As a result, the on-order state spaces can be restricted to $\mathcal{O}_l = \{0, 1, \ldots, \min(\bar{I}, \sum_{i=l+1}^{L} \bar{A}_l^+)\}$.

The number of possible inventory states can be reduced when the exogenous state includes the empty container returns of the previous period. Proposition 2 is a result of the inventory transition function (3) and states that the previous period's empty container returns can generally not exceed the current epoch's inventory position. An exception applies, however, to the greatest inventory position $I_t = \bar{I}$. In this inventory state all previous returns that exceed the storage capacity are necessarily truncated to respect the capacity constraint.

Proposition 2. *If $Z_{t-1}^+ \in X_t$, then states S_t with*

$$I_t < Z_{t-1}^+ < \bar{I}$$

cannot occur in any solution.

Propositions 3 and 4 demonstrate how to reduce the action space for A_t^- and $A_{0,t}^+$ when the system is in state S_t. The first proposition relies on the lost sales costs and the demand process and finds the largest out-positioning decisions that may be optimal. All greater out-positions, which are shown to

Table 1: Cost and capacity parameters for in- and out-positioning decisions.

decision variable	capacity	unit cost
A_t^-	35 ($\bar{A}^-$)	25 (λ^-)
$A_{0,t}^+$	35 ($\bar{A}_0^+$)	100 (λ_0^+)
$A_{1,t}^+$	35 ($\bar{A}_1^+$)	75 (λ_1^+)
$A_{2,t}^+$	35 ($\bar{A}_2^+$)	50 (λ_2^+)

cause lost sales with probability one, are found to be suboptimal. The second proposition extends this line of thought to empty container inflows and identifies the largest in-positioning decision. All greater in-positions are shown to be suboptimal due to the finite storage capacity and the accrued penalty for exceeding it. The identified in- and out-positioning can be removed from $\mathcal{A}_{S_t}$ as a consequence of their suboptimality.

Proposition 3. *Let*

$$\mathcal{Z}_{x_t}^- = \{z_t^- \in \mathbb{N}_0 \mid P(z_t^-|x_t) > 0\}$$

denote all possible empty container demand realizations that occur with non-zero probability in decision period t. For $\lambda^- > 0$ and $\lambda_l \geq 0$, any out-positioning decision A_t^- that satisfies

$$A_t^- > I_t \ - \min \mathcal{Z}_{x_t}^- \geq 0,$$

in state S_t cannot be optimal.

Proposition 4. *Let*

$$\mathcal{Z}_{x_t}^+ = \{z_t^+ \in \mathbb{N}_0 \mid p(z_t^+|x_t) > 0\}$$

denote all possible empty container return realizations that occur with non-zero probability in decision period t. For $\lambda_0^+ > 0$ and $\lambda_p \geq 0$, any in-positioning decision $A_{0,t}^+$ that satisfies

$$A_{0,t}^+ > \bar{I} \ - \big(\max(I_t - \max \mathcal{Z}_{z_t}^-, 0) + \min \mathcal{Z}_{z_t}^+ + O_{0,t}\big) \geq 0,$$

in state S_t cannot be optimal.

4 Simulation studies

This section presents two simulation studies illustrating how different assumptions impact the optimal in- and out-positioning decisions. The first study demonstrates the potential for reducing the costs of an inland container depot when in-positioning options with varying lead times and costs are modelled explicitly. The study further investigates conditions under which more expensive but faster in-positioning options are favorable. The second study confirms the importance of accounting for serial and cross-sectional dependence in the empty container demand and return process. So far, we have assumed this process to be known. However, this assumption is violated in most real-world applications. We consider different degrees of misspecifications between the true and the assumed demand and return processes and learn the optimal policies under these assumptions. Our results illustrate that a policy's true cost can deviate significantly when the serial and cross-sectional dependence structure is not captured correctly.

Network parameters Both studies consider the inland network in Figure 1, where the inland depot can only exchange empty containers with a port. The storage capacity of the depot is assumed to be $\bar{I} = 80$ and there exist 3 daily transportation options with the port, for which lead times vary between 0, 1 and 2 days. Table 1 displays cost and capacity parameters. The cost parameters are chosen to incentivize the usage of slower in-positioning options. The capacity parameters are selected such that $A_{0,t}^+$ alone can cover all the in-positioning demand of the depot when $\bar{A}_1^+ = \bar{A}_2^+ = 0$. The out-positioning costs are selected to be lower as the in-positioning costs to incentivize stock level reduction through out-positioning at times of a container surplus, a behavior that is expected from real-world depot operations. For the cost parameters of the storage in Table 2, we choose the lost sales penalty to be in the neighborhood of current global freight rates for 20ft containers. The penalty for exceeding the storage capacity is naturally smaller since the arriving empty containers can either be redirected to a different depot or additional space at the depot can be used with a premium on costs. The holding costs are small because it is assumed that container handling costs are included in the costs for in- and out-positioning.

Table 2: Depot parameters

capacity ($\bar{I}$)	lost-sales penalty (λ_l)	exceedance penalty λ_p	holding cost λ_h
80	1,000	50	1

Demand and return processes Throughout the remainder of this section it is assumed that the marginal processes for empty container demands and returns are negative binomial autoregressive (NBAR) processes

$$Z_t|Z_{t-1} \sim P(z_t|z_{t-1}) = NB\left(\frac{\mu_t^2}{\sigma_t^2 - \mu_t}, \frac{\alpha}{\alpha + \mu_t}\right),$$

where the superscripts to denote demand (Z_t^-) and return (Z_t^+) stochastic variables are omitted. $NB(\cdot,\cdot)$ denotes a negative binomial distribution with mean and variance

$$\begin{aligned} \mathrm{E}[Z_t|Z_{t-1}] &= \mu_t = \exp(c + \theta \log(Z_{t-1} + 1)) \\ \mathrm{Var}[Z_t|Z_{t-1}] &= \sigma_t^2 = \mu_t + \frac{1}{\alpha}\mu_t^2 \end{aligned} . \tag{5}$$

The autoregressive parameter θ controls the strength of serial dependence, c is a drift term and the dispersion parameter, with $\alpha = 10$ for all simulated processes if not otherwise stated, induces overdispersion. The experimental study with real-world data in Section 5 illustrates the fit to the data. We utilize a copula to induce negative and positive cross-sectional dependence between demands and returns. Copulas have the benefit over multivariate stochastic processes, such as frailty models, that the dependence strength is controlled in a simple manner. In this study we apply the Clayton copula

$$C_\kappa(u, v) = \max\left(\left[u^{-\kappa} + v^{-\kappa} - 1\right]^{-1/\kappa}, 0\right)$$

to the conditional CDFs $F(z_t^-|z_{t-1}^-)$ and $F(z_t^+|z_{t-1}^+)$ of two independent NBAR processes, with copula parameter $\kappa \in [-1, \infty]\backslash\{0\}$. We subsequently use the Kendall rank correlation coefficient $\tau \in (-1, 1)$ to report dependence strengths between demands and returns. Kendall's τ measures the ordinal association between two quantities and relates to the Clayton copula parameter κ through $\tau = \kappa/(2 + \kappa)$ (Nelsen, 2006). Similar to a Pearson correlation coefficient, large negative τ indicate strong negative correlations and $\tau = 0$ implies no rank correlation. Other bivariate copulas could have been applied, such as the Frank (uniform dependence across the variables' domain) or Gumbel (stronger dependence in both tails) (Nelsen, 2006). However, we consider the Clayton (stronger dependence in the left tail for low demands and returns) copula because it is found to best describe the cross-sectional dependence structure of the real-world data in Section 5.

A caveat of our approach to model cross-sectional dependent processes is that no closed form expression exists for the unconditional distribution $P(z_t^-, z_t^+)$ when demands and returns are serially dependent. However, we require this distribution to learn policies under misspecified serial dependence structures. As a solution we propose to approximate $P(z_t^-, z_t^+)$ with a trajectory of 10^7 demand and return samples. Starting in $t = 1$, demand and return pairs are sampled from the copula $C_\kappa(F(z_t^-|z_{t-1}^-), F(z_t^+|z_{t-1}^+))$, where z_{t-1}^- and z_{t-1}^+ are the sampled values of the previous time step and initialized with $z_0^- = 0, z_0^+ = 0$. The first 10^4 samples of the simulated trajectory are removed to reduce the effect of this initialization. The distribution $P(z_t^-, z_t^+)$ is eventually obtained with the observed frequencies of the remaining samples.

Dynamic programming parameters To compare policies with different in-positioning transportation modes and demand and return processes, and hence, different state spaces, we use the terminal cost function (4). The distribution of Z_T^- and Z_T^+, over which the expectation is calculated, is specified by the experiment. We set $T = 100$ to reduce the effect of the terminal value on the policies during early decision epochs.

Policy evaluations and comparisons Consider the case where a policy is learned under a demand and return process that deviates from the true process. For a system occupying an endogenous state $s_0^\dagger$, a policy's costs under misspecification are calculates as

$$\mathrm{E}\Big[V_{\pi^\star}(S_0^\dagger, X_0)|S_0^\dagger = s_0^\dagger\Big],$$

where the expectation is evaluated with respect to the true demand and return process. The initial exogenous state x_0 is omitted in the conditional expectation to remove its effect on the estimated

Table 3: Policy value ratios of the models with maximum in-positioning transportation lead times of 0, 1 and 2 decision periods (denoted by a policy's subscript). The first column reports the variance of the demand process, which follows a negative binomial distribution.

$\mathrm{Var}\left[Z_t^-\right]$	$R(\pi_0^\star, \pi_1^\star)$	$R(\pi_0^\star, \pi_2^\star)$
20	1.29	1.81
40	1.27	1.73
60	1.24	1.61
80	1.20	1.47
100	1.16	1.35

costs. Monte-Carlo integration is applied to compute the expectation since an exact value is only obtainable from a policy's value function when the true and assumed exogenous processes align. 10,000 trajectories of length T are sampled from the true demand and return process to approximate the expectation by applying a learned policy forward from the initial $s_0^\dagger$ until the end of the planning horizon. The estimated values of two policies are compared by evaluating the ratio

$$R(\pi_1^\star, \pi_2^\star) := R(\pi_1^\star, \pi_2^\star | s_0^\dagger) = \frac{\mathrm{E}\Big[V_{\pi_1^\star}(S_0^\dagger, X_0) | S_0^\dagger = s_0^\dagger\Big]}{\mathrm{E}\Big[V_{\pi_2^\star}(S_0^\dagger, X_0) | S_0^\dagger = s_0^\dagger\Big]},$$

which is a relative regret if $\pi_2^\star$ is learned under the true exogenous process whereas $\pi_1^\star$ is not. Throughout the remainder of this section we assume $s_0^\dagger = (i_0, o_0) = (\bar{I}/2, 0) = (40, 0)$ in all experiments. All initial on-order states o_0 are set to zero because policies are learned for MDPs with different in-positioning lead times in our initial experiment. Non-zero on-order states would otherwise be disadvantages for a policy with maximum lead time $L = 0$, for which the endogenous state space reduces to $S_t^\dagger = I_t$.

4.1 Varying in-positioning lead times

In this experiment, we learn policies of 3 different MDPs to verify that the availability of slower in-positioning transportation modes reduces the operational costs of the depot. The first MDP does not model in-positioning options with lead times greater than 0, i.e. $\bar{A}_1^+ = \bar{A}_2^+ = 0$, but takes $\bar{A}_0^+$ as in Table 1. Similarly, the second model takes $\bar{A}_2^+ = 0$ and the third model is unrestricted with the originally proposed capacity parameters. The policies of all three models are labeled according to the maximum available lead time of the transportation modes, i.e. π_0, π_1 and π_2.

All policies are estimated for five exogenous demand and return processes. Each demand distribution is taken to be negative binomial with mean 10, i.e. $c = log(10)$ and $\theta = 0$ in the conditional expectation 5 of the NBAR process, but with varying variances (20, 40, 60, 80 and 100) that we control through the dispersion parameter α. The return variables are also negative binomial with $c = log(5), \theta = 0$ and $\alpha = 10$ in Equation 5. Both variables are independently and identically distributed, hence the exogenous state space of all models is empty. We assume that the exogenous processes of the MDP models are correctly specified.

Results The results in Table 3 show the effectiveness of the slower transportation modes to reduce costs. If the depot has initially 40 empty containers in stock, no containers are on order and the demand variance is 20, then the policy with maximum lead time of 0 has 29% and 81% larger expected operating cost during the planning period of 100 epochs than the policies with maximum transportation lead times of 1 and 2 days, respectively. It is observed, though not shown in the paper, that no in- or out-positioning decisions of the optimal policies are at their capacity bounds. Thus, the reported cost ratios are not due to capacity shortages of the policies with smaller total in-positioning capacities. The results further indicate that the usage of slower in-positioning transportation modes is subject to the demand process variance. As indicated by the declining policy value ratios in Table 3, the relative benefit of using slower in-positioning modes diminishes with increasing demand variance. This follows naturally since planning multiple periods ahead becomes more difficult as the uncertainty of future empty container demands and returns increases.

Table 4: Serial and cross-sectional dependence of the exogenous process models for the MDP formulations.

MDP name	demand & return dependencies	exogenous state
AR	$P(z_t^-\|z_{t-1}^-)P(z_t^+\|z_{t-1}^+)$	$X_t = (Z_{t-1}^-, Z_{t-1}^+)$
IID	$P(z_t^-)P(z_t^+)$	None
ARcross	$P(z_t^-, z_t^+\|z_{t-1}^-, z_{t-1}^+)$	$X_t = (Z_{t-1}^-, Z_{t-1}^+)$
IIDcross	$P(z_t^-, z_t^+)$	None

4.2 Misspecifications of the exogenous process

In this study, we examine the effects on the operating costs of the exogenous process differing between policy estimation and evaluation. We begin by investigating the consequences of not accounting for serial dependence. We proceed to examine the misspecification introduced by ignoring cross-sectional dependence between empty container demands and returns.

The policies of this study stem from four different MDPs. The second column of Table 4 lists the MDP's stochastic process for demands and returns reflecting whether serial or cross-sectional dependence is considered, and the third column shows the corresponding exogenous state. The two marginal distributions $P(z_t^-)$ and $P(z_t^+)$ of the *IID* model are approximated through Monte-Carlo sampling from two specified NBAR processes, whereas the joint distribution $P(z_t^-, z_t^+)$ of the *IIDcross* model is approximated by the trajectory sampled from the Copula $C_\kappa(F(z_t^-|z_{t-1}^-), F(z_t^+|z_{t-1}^+))$. The sample spaces, i.e. the observation pairs $x_t = (z_{t-1}^-, z_{t-1}^+)$ that are sampled during the Monte-Carlo procedure, are used as the exogenous state space of the *AR* and *ARcross* models to avoid the inclusion of exogenous states that are never visited according to the approximated unconditional distributions. To further limit the computational requirements, we remove exogenous states which are rarely visited by removing samples with low probability from the space in increasing order until a maximum of 0.0005 of the total mass is removed. The transition probabilities are corrected accordingly by redistributing missing mass in the conditional distributions equally among each of their remaining samples. A consequence of this procedure is that a pair (z_{t-1}^-, z_{t-1}^+) which has not been part of the exogenous state space during policy estimation may be sampled during policy evaluation. We propose to select the action from the exogenous state x_t in the estimated policy with the smallest Euclidean distance to the observed pair. To ensure comparability between all policies, we use the approximated unconditional distribution of mutually independent demands and returns for calculating the terminal cost in Equation (4). Thus, all MDPs in Table 4 have the same terminal costs. The remaining model parameters for the depot and transportation modes are taken as introduced in Table 1 and 2, except that the in-positioning transportation mode with the lead time of 2 days is excluded to avoid excessive computations, i.e. $\bar{A}_2^+ = 0$.

4.2.1 Misspecified serial dependence

To begin the examination of misspecified exogenous processes, we learn the policies of the *IID* and *AR* MDPs for a range of mutually independent NBAR processes for empty container demands and returns. The *IID* model is misspecified since it ignores the serial dependence of the exogenous process. Thus, to estimate the true expected costs under this policy, we use Monte-Carlo integration with samples from the mutually independent NBAR processes. We also use Monte-Carlo integration for the *AR* model with the correctly specified exogenous processes to integrate out the effect of the initial exogenous state X_0 on the costs during the planning horizon of 100 epochs.

Table 5 reports the relative regret $R(\pi_{IID}^\star, \pi_{AR}^\star)$ for 25 different demand and return processes. As before, we take the initial endogenous state to be $s_0^\dagger = (40, 0)$. The top rows with $\theta^- > 0$ and $\theta^+ = 0$ highlight the misspecification effect for demand processes with positive serial dependence and independently distributed returns. We observe that the misspecification leads to understocking, which is in concordance with known results for standard inventory systems without stochastic returns, where the understocking effect becomes more pronounced with increasing autocorrelation. Even though holding costs are comparatively lower for the *IID* policy, overall costs are higher due to the larger lost sales penalty costs that are caused by operating the depot at lower than optimal inventory levels. Therefore, the observed cost differences are highly dependent on the lost sales penalty λ_l. The higher costs under the *IID* policy will diminish with decreasing penalities because lost sales due to understocking will have smaller effects on depot operating costs. The results in the bottom rows for $\theta^- < 0$ further verify that undetected negative autocorrelated demand causes

Table 5: Relative regret $R(\pi^\star_{IID}, \pi^\star_{AR})$ for misspecified serial dependence. The first column reports the parameters of the two mutually independent NBAR processes with conditional mean and variance defined in Equation 5, and dispersion parameter $\alpha = 10$ for both processes. The autoregressive parameters θ^- and θ^+ control the serial dependence strength, for which $\theta^- = \theta^+ = 0$ implies serial independence. The two last columns list the average accrued lost sales penalty and holding cost differences for the 10,000 Monte Carlo trajectories of length T=100.

θ^-	c^-	θ^+	c^+	$R(\pi^\star_{IID}, \pi^\star_{AR})$	$\mathrm{E}\left[V_{\pi^\star_{IID}}(S^\dagger_0, X_0) \mid S^\dagger_0 = s^\dagger_0\right]$	lost sales cost (IID - AR)	holding cost (IID - AR)
0.7	0.5	0.7	0.5	1.32	22,868	8,313	-482
0.7	0.5	0.35	0.7	1.15	34,001	6,254	-832
0.7	0.5	0	1.5	1.19	25,803	6,146	-702
0.7	0.5	-0.35	1.5	1.14	35,587	6,081	-852
0.7	0.5	-0.7	2.5	1.19	24,795	5,806	-623
0.35	0.7	0.7	0.5	1.17	12,811	2,446	-817
0.35	0.7	0.35	0.7	1.06	5,264	525	-489
0.35	0.7	0	1.5	1.05	6,047	503	-412
0.35	0.7	-0.35	1.5	1.02	4,769	279	-128
0.35	0.7	-0.7	2.5	1.01	5,894	46	24
0	1.5	0.7	0.5	1.15	10,452	1,414	-983
0	1.5	0.35	0.7	1.0	10,092	37	-10
0	1.5	0	1.5	1.0	4,584	-2	8
0	1.5	-0.35	1.5	1.0	11,574	-58	49
0	1.5	-0.7	2.5	1.03	4,505	-87	486
-0.35	1.5	0.7	0.5	1.04	11,155	608	-186
-0.35	1.5	0.35	0.7	1.01	3,993	-50	59
-0.35	1.5	0	1.5	1.02	5,477	-83	266
-0.35	1.5	-0.35	1.5	1.06	3,450	-93	676
-0.35	1.5	-0.7	2.5	1.08	5,960	-120	693
-0.7	2.5	0.7	0.5	1.03	9,237	134	-178
-0.7	2.5	0.35	0.7	1.04	11,847	-322	532
-0.7	2.5	0	1.5	1.06	5,111	-241	459
-0.7	2.5	-0.35	1.5	1.03	13,487	-339	591
-0.7	2.5	-0.7	2.5	1.09	4,896	-234	744

Table 6: Relative regrets for serial and cross-sectional dependence misspecification for the *IID*, *IIDcross* and *AR* policies. Kendall's τ reports the cross-dependence strength, with $\tau = \kappa/(2+\kappa)$ for the applied Clayton copula.

θ^-	c^-	θ^+	c^+	τ	$R(\pi^\star_{AR}, \pi^\star_{ARcross})$	$R(\pi^\star_{IIDcross}, \pi^\star_{ARcross})$	$R(\pi^\star_{IID}, \pi^\star_{ARcross})$
-0.35	1.5	-0.35	1.5	0.50	1.07	1.03	1.25
-0.35	1.5	-0.35	1.5	0.25	1.02	1.05	1.13
-0.35	1.5	-0.35	1.5	0.00	1.00	1.06	1.06
-0.35	1.5	-0.35	1.5	-0.25	1.01	1.07	1.03
-0.35	1.5	-0.35	1.5	-0.50	1.02	1.07	1.02
0.35	0.7	0.70	0.5	0.50	1.08	1.02	1.01
0.35	0.7	0.70	0.5	0.25	1.03	1.11	1.04
0.35	0.7	0.70	0.5	0.00	1.00	1.17	1.17
0.35	0.7	0.70	0.5	-0.25	1.02	1.21	1.39
0.35	0.7	0.70	0.5	-0.50	1.06	1.21	1.58

overstocking when returns are serially independent. However, the cost ratios are smaller compared to the reported values in the top rows for undetected positive dependence. This follows from the smaller contribution of holding costs to total operating costs, which are largely determined by lost sales penalities and in-positioning costs. Hence, the reported results will vary for different holding cost parameters λ_h, with smaller differences for decreasing λ_h. In spite of the seemingly smaller impact of undetected negative serial dependence on a depot's operating costs, the consequences are more far-reaching because costs of an inflated container fleet size accrue when hundreds of depots experience negatively autocorrelated demand processes.

While the previous results are known for inventory systems with stochastic demands, we observe the same effects for misspecifications of stochastic returns. The center of the table shows for serially independent demand that undetected positive serial dependence in the return process leads to understocking, whereas undetected negative dependence leads to overstocking. In the special cases where demand is positively and returns negatively dependent, but also in the reverse case, we observe that the individual misspecification effects act favorable in case of the *IID* policy. The results show that undetected positive demand dependence has less severe consequences when negative serial dependence of returns also goes undetected. The reverse is true for positive return dependence that generally adds to the understocking effect of undetected positive demand dependence. In real-world applications it is unknown whether the misspecifications of demand and return processes act favorable, and thus, we emphasize the importance of accurately accounting for serial dependence.

4.2.2 Misspecified cross-sectional dependence

In this final experiment, we learn policies for all four MDPs to examine the effect of model misspecification under serial and cross-sectional dependence. We follow the steps of the previous experiment and evaluate all learned policies for $s_0^\dagger = (40, 0)$. The results with comparison to the correctly specified MDP model *ARcross* are reported in Table 6, where Kendall's τ reflects the strenght of cross-sectional dependence induced by the Clayton copula. Two distinct autocorrelated demand and return processes have been selected.

The first 5 rows of the table correspond to a balanced depot since the marginal NBAR processes are identical for the demands and returns. Both variables are negatively serially dependent, for which we have shown in the previous study that overstocking occurs when the serial dependence goes undetected. The consequently higher operating costs (6%) are shown in row number 3 for the mutually independent reference case with $\tau = 0$. With respect to the policy of the *ARcross* model, we observe that the *IID* policy of the last column performs worse for positive cross-sectional dependence and better for negative dependence, where it is on par with the *AR* policy for $\tau = -0.5$. While this seems at first counterintuitive, it can be explained by the same effects that were observed in the previous experiment when undetected serial dependence of demands and returns acted favorable for the *IID* policy.

To begin with, consider the case of negative cross-sectional dependence for serially independent demands and returns. If this dependence goes undetected, inventory levels are generally too low, hence understocking occurs. This follows from information being ignored that large demands and few returns are more likely to occur jointly. Under the occurrence of such an event the inventory system

moves into an unfavorable state since the opening inventory level of the following decision epoch may drop too low. More lost sales are likely to follow in the next decision period. Rows 4 and 5 in Table 6 report that the misspecification of negative cross-sectional dependence, hence understocked inventory levels, acts against the overstocking effect of misspecified serial dependence. The relative regret $R(\pi^\star_{\text{AR}}, \pi^\star_{\text{ARcross}})$ increases due to the stronger misspecified cross-sectional dependence, hence less than optimal inventory levels occur under the AR policy. On the contrary, $R(\pi^\star_{\text{IID}}, \pi^\star_{\text{ARcross}})$ decreases because the otherwise too high inventory levels for the misspecified negative serially dependent demand and return process are beneficial when large demands and small returns occur in the same decision period.

The same dynamics are observed for the relative regrets in the first two rows with positive cross-sectional dependence. Overstocking occurs if this dependence goes undetected because information is ignored that large demands and returns are more likely to occur jointly, which has the favorable effect of ensuring sufficient inventory levels after large quantities of demand has been satisfied. The shown relative regrets in the first two rows demonstrate that accounting for the cross-sectional dependence is essential to prevent additional overstocking when negative serial dependence of demands and returns goes undetected. It is further shown that under strong positive cross-sectional dependence ($\tau = 0.5$) it is less detrimental to misspecify the serial dependencies as long as the cross-sectional dependence is modelled correctly. This may lead to the following question. Which dependencies should be addressed more carefully? The answer is both because in real-world applications it is unclear which dependence structures are present before any investigations are conducted. The last two rows of Table 6 should serve as a warning in this regard. Learning a policy under the convenient assumption of serial and cross-sectional independent demand and returns can have catastrophic consequences if the true process deviates unfavorably.

5 Experimental study

This section demonstrates the existence of serial- and cross-sectional dependence in the real-world data of a global container shipping company, and the consequences of misspecifying these dependencies. The dataset contains daily empty container deliveries to export customers and container returns from import customers for an inland depot that is in deficit of empty containers. We consider observations from 01.01.2014 but only preceding 01.01.2020 due to the induced non-stationarities of the COVID-19 pandemic on global trade. The 2192 observations measured in 20ft containers are converted to 40ft containers (division by two and rounding down to the nearest integer number) to reduce the state space. Except for a weekend effect with lower deliveries and returns on Saturdays and Sundays, no seasonalities are detected for the time series.

Demand and return processes The true stochastic process of empty container demands and returns is unknown. We therefore use the time series data to estimate statistical models that determine the exogenous process in our MPDs. The statistical models are learned under the assumption that historical empty container deliveries to export customers are a proxy for the unknown empty container demand (Z_t^+). Similarly, historical returns are a proxy for customer initiated empty container returns (Z_t^-). In total we learn four statistical models with varying serial and cross-sectional dependencies. We use a variation of the naming convention in Table 4 to highlight the dependence assumptions of the MDPs. A minor adaptation is applied because the weekend effect induces a time-dependence in the distributions for demand and returns. In decision period t the stochastic variables are not identically distributed because their distributions differ between weekdays and weekends. Thus, *ID* and *IDcross* indicate the time-inhomogeneous MDPs with serially independent exogenous processes.

The exogenous process of the MDPs with mutual independent demands and returns are obtained through negative binomial (NB) regression. The conditional mean in Equation (5) of the NBAR process is adapted to

$$\mathrm{E}[Z_t|Z_{t-1}] = \exp(c + \beta u_t + \theta \log(Z_{t-1} + 1)), \tag{6}$$

which accounts for the weekend effect with a regression coefficient β. A deterministic indicator variable u_t is employed to differentiate between weekdays ($u_t = 0$) and weekends ($u_t = 1$). The demand and return process of the AR model is obtained by estimating the unknown coefficients in Equation (6) and dispersion parameter α of the NBAR process from the historical time series. For the exogenous process model of the *ID* MDP, $\theta = 0$. A Clayton copula with NBAR conditional

Table 7: Maximum likelihood estimates of the NB regression and copula models for empty container returns (Z_t^-) and demand (Z_t^+) with marginal distributions parameterized as in Equation (6). Kendall's τ in the second last column is obtained as $\tau = \kappa/(2+\kappa)$.

	c^-	θ^-	β^-	α^-	c^+	θ^+	β^+	α^+	κ (τ)	AIC
ARcross	1.60	0.33	-0.28	22.60	1.45	0.31	-0.19	5.31	0.18 (0.08)	24,259
AR	1.56	0.35	-0.28	22.65	1.44	0.32	-0.20	5.27	0 (0)	24,329
IDcross	2.38	0	-0.28	14.98	2.07	0	-0.14	4.48	0.22 (0.1)	24,751
ID	2.38	0	-0.29	14.78	2.08	0	-0.15	4.41	0 (0)	24,867

marginals (6) is estimated for the joint process of historical deliveries and returns to obtain the process for the cross-sectional dependent *ARcross* and *IDcross* MDPs.

The coefficients of all four statistical models are estimated by maximizing the conditional likelihood. As noted in Trivedi and Zimmer (2017), the employed autoregressive and weekday regressors mitigate the identification concerns of copulas for discrete outcomes. The estimated coefficients in Table 7 show mild positive serial dependence and weak cross-sectional dependence. The Akaike Information Criterion (AIC) identifies the copula model with autoregressive regressors as the best fitting statistical model for the unknown stochastic process (Akaike, 1974). We also estimated the Frank and Gumbel copula to the data (Nelsen, 2006). However, none attained a lower AIC than the Clayton copula. The simulation study's procedure to obtain the exogenous state spaces for the *AR* and *ARcross* model is applied. In particular, $P(z_t^-, z_t^+)$ is approximated by 10^7 Monte-Carlo samples, and removing samples with low probability from the approximated distribution. Samples are removed in increasing order until a maximum of 0.0005 of the total mass has been removed.

Network and dynamic programming parameters The depot parameters are taken as listed in Table 2. The exogenous state spaces of the MDPs with serially dependent processes are larger compared to the previous simulation study. To avoid excessive computations we reduce the in- and out-positioning capacities to $\bar{A}^- = 10, \bar{A}_0^+ = 5, \bar{A}_1^+ = 5$ and $\bar{A}_2^+ = 10$. All cost parameters are the same as in Table 1. The planning horizon is 7 weeks ($T = 49$), and planning begins on a Monday ($t = 0$) and ends on a Sunday (t=49) when the terminal costs as defined in Equation (4) accrue. With this assumption, we omit the inclusion of a state variable for the dummy variable u_t since its value can be inferred from the decision epoch t. The two marginal distributions $P(z_t^-|u_t = 1)$ and $P(z_t^+|u_t = 1)$ that are estimated for the *ID* MDP are used for the terminal cost calculation of all other MDPs to ensure identical terminal costs, and hence allows a fair comparison between all policies because cost differences accrue during $t = 0, \ldots, T-1$.

Policy evaluations and comparisons The policy values with respect to the unknown stochastic process for empty container demands and returns is estimated by using the historical time series samples to approximate the expectation (6), with $s_0^\dagger = (i_0, o_0) = (40, 0)$ as before. Each sample trajectory contains $T = 49$ observations, thus a total of 44 trajectories are obtained from the 2192 time series observations. The first observation of each trajectory corresponds to a Monday to align with the assumption that planning begins on a Monday. The exogenous state x_0 of the *AR* and *ARcross* model are thus the empty container deliveries and returns of the preceding Sunday. Since the earliest observation is available for Wednesday 01.01.2014, the first sample of the first trajectory is the delivery and return observation pair of Monday 06.01.2014 and x_0 is the observation pair of Sunday 05.01.2014.

Results The results in Table 8 agree with the misspecification effects of the simulation study. We find the *ID* and *IDcross* policies to have an average 7% and 9% higher costs than the *ARcross* reference policy. When ignoring positive serial dependencies for demand and returns in the *ID* model, inventory stock levels are comparatively lower, hence holding and in-positioning cost are lower, than the policies with serial dependent exogenous processes. Figure 3 exemplifies the difference in inventory positions for a selected sample trajectory. Eventually, the lower inventory levels imply higher lost sales, and thus, higher overall costs. Higher average out-positioning cost for policies with autoregressive exogenous processes further indicate their superiority. Indeed, the modelled serial dependence provides additional information such that containers are only out-positioned in periods when they are not needed. This is in contrast to general understocking behavior under the *ID* and *IDcross* policies. The results additionally show the interaction effects of misspecified serial and cross-sectional dependencies. The average holding costs during the planning period of the

Table 8: Estimated policy value for historical empty container delivery and return time series. Parenthesized values in the first column refer to the ratio $R(\pi^\star, \pi^\star_{ARcross})$ where $\pi^\star$ is the policy of the respective row. The remaining columns report the average accrued cost for the planning horizon of 49 days, fx. the *ARcross* policy accrued an average 682 lost sales cost for each of the 44 sample trajectories.

	$\mathrm{E}\left[V_{\pi^\star}(S_0^\dagger, X_0) \mid S_0^\dagger = s_0^\dagger\right]$	lost sales cost	holding cost	storage excess cost	in- & out-positioning cost
ARcross	9,194	682	1,723	228	6,355 & 35
AR	9,156 (1.0)	523	1,759	243	6,428 & 38
IDcross	10,051 (1.09)	1,977	1,587	218	6,012 & 6
ID	9,808 (1.07)	1,659	1,626	230	6,072 & 8

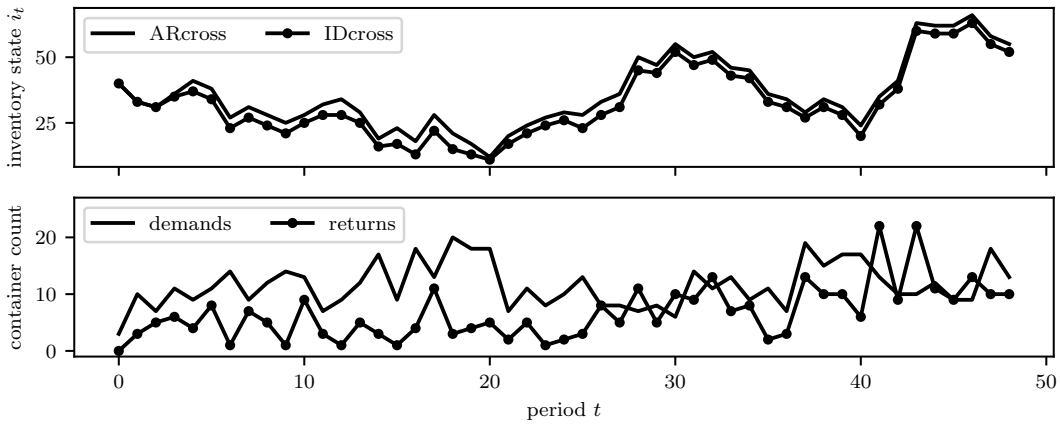

Figure 3: Attained inventory states (top) of the *ARcross* and *IDcross* policies for a selected empty container delivery and return trajectory (bottom) when starting in $s_0^\dagger = (i_t = 40, o_{0,t} = 0)$. The first delivery and return observation pair corresponds to Monday 10.05.2015.

ID (1,626) and *IDcross* (1,587) policies exemplify this effect. It is observed that the overstocking effect of undetected positive cross-sectional dependence under the *ID* policy compensates for the understocking effect of undetected positive serial dependence.

Pruning state and action spaces Finally, we would like to emphasize the effectiveness of the propositions in Section 3.8 for pruning state and action spaces. The unpruned state space of the *ARcross* MDP contains about 12.7 million states, and is reduced with the application of Proposition 2 by roughly 18%. The selection of in-positioning capacities causes Proposition 1 to be ineffective because no on-order state exceeds the storage capacity. Clearly, this will change if in-positioning capacities increase.

The maximum observable demand is below the storage capacity, which allows the application of Proposition 4 to reduce the action space for $A^+_{0,t}$. However, the resulting exogenous processes of the estimated statistical models imply a minimum demand of zero containers for each decision period. Proposition 3 is consequently ineffective since it is equivalent to the constraint $A^-_t > I_t$. The combined effect of pruning states with Proposition 2 and reducing the state-dependent action spaces with Proposition 4 is that the total number of evaluated actions per epoch reduces from approximately 12.9 to 10.9 billion, a reduction of 15%.

6 Conclusions

In this paper, we have extended existing empty container allocation models by accounting for varying lead times of transportation modes that are available in inland logistics networks. A simulation study confirmed the cost reductions due to a usage of slower but less expensive transportation modes for repositioning empty containers between a port and an inland depot. Furthermore, we investigate serial and cross-sectional dependent empty container demand and return processes, and the effect on policy values when the dependencies go undetected. The study extended known results for the misspecification of serial dependence for demand processes in inventory systems to the empty container return process. To the best of our knowledge, we are the first to investigate the consequences of ignoring cross-sectional dependence between container demands and returns for empty container allocation models. This is particular important due to the common assumption that empty container returns have a one period lead time before they can be reused to satisfy demand. A real-world dataset of delivered and returned empty containers of a depot demonstrated the existence of both types of dependence structure. Policy evaluation with historical data confirmed the reliance of our MDP on exogenous process models that accurately describe the serial and cross-sectional dependence of empty container demands and returns. A caveat of our MDP is that the estimation of an optimal policy becomes intractable for larger state spaces.

Future work should thus address policy approximation methods for the applicability of the MDP to larger problem instances with more in-positioning transportation modes and more complex demand and return processes. Our backwards dynamic programming algorithm could provide reference for how well such approximation methods perform on the problem instances considered in this paper. A next step could be assess the benefits for repositioning decisions of a routing model. To incorporate even more information about the inland network, one can investigate the effect of stochastic transportation lead times, prices and capacities on in- and out-positioning decisions. The relaxation of the fixed one period lead time for returned empty containers is likewise of practical interest. In practise, the state of returned containers varies, such that returned damaged containers have a greater lead time before they can be reused for a new shipment.

Acknowledgments

This work has been supported by "InnovationsFonden Danmark" [grant number 9065-00021B]. The authors owe a special thanks to Gitte Rasmussen for her valuable contributions at the early stages of this research project.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.

Braekers, K., Janssens, G. K., and Caris, A. (2011). Challenges in Managing Empty Container Movements at Multiple Planning Levels. *Transport Reviews*, 31(6):681–708.

Chou, C.-C., Gou, R.-H., Tsai, C.-L., Tsou, M.-C., Wong, C.-P., and Yu, H.-L. (2010). Application of a mixed fuzzy decision making and optimization programming model to the empty container allocation. *Applied Soft Computing*, 10(4):1071–1079.

Crainic, T. G., Gendreau, M., and Dejax, P. (1993). Dynamic and Stochastic Models for the Allocation of Empty Containers. *Operations Research*, 41(1):102–126.

Dang, Q.-V., Yun, W.-Y., and Kopfer, H. (2012). Positioning empty containers under dependent demand process. *Computers & Industrial Engineering*, 62(3):708–715.

Fang, J., Zhao, L., Fransoo, J. C., and Van Woensel, T. (2013). Sourcing strategies in supply risk management: An approximate dynamic programming approach. *Computers & Operations Research*, 40(5):1371–1382.

Feng, C.-M. and Chang, C.-H. (2008). Empty container reposition planning for intra-Asia liner shipping. *Maritime Policy & Management*, 35(5):469–489.

Goltsos, T. E., Syntetos, A. A., Glock, C. H., and Ioannou, G. (2022). Inventory – forecasting: Mind the gap. *European Journal of Operational Research*, 299(2):397–419.

Govindan, K., Soleimani, H., and Kannan, D. (2015). Reverse logistics and closed-loop supply chain: A comprehensive review to explore the future. *European Journal of Operational Research*, 240(3):603–626.

Graves, S. C. (1999). A Single-Item Inventory Model for a Nonstationary Demand Process. *Manufacturing & Service Operations Management*, 1(1):50–61.

Li, J.-A., Liu, K., Leung, S. C., and Lai, K. K. (2004). Empty container management in a port with long-run average criterion. *Mathematical and Computer Modelling*, 40(1-2):85–100.

Minner, S. (2003). Multiple-supplier inventory models in supply chain management: A review. *International Journal of Production Economics*, 81–82:265–279.

Nelsen, R. B. (2006). *An Introduction to Copulas*. Springer Series in Statistics. Springer, New York Berlin Heidelberg, second edition.

Notteboom, T., Pallis, A., and Rodrigue, J.-P. (2021). *Port Economics, Management and Policy*. Routledge, London, first edition.

Olivo, A., Zuddas, P., Di Francesco, M., and Manca, A. (2005). An Operational Model for Empty Container Management. *Maritime Economics & Logistics*, 7(3):199–222.

Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley, Hoboken, New Jersey, first edition.

Ray, W. D. (1980). The Significance of Correlated Demands and Variable Lead Times for Stock Control Policies. *The Journal of the Operational Research Society*, 31(2):187–190.

Song, D. and Dong, J.-X. (2015). Empty container repositioning. In Lee, C.-Y. and Meng, Q., editors, *Handbook of Ocean Container Transport Logistics*, volume 220 of *International Series in Operations Research & Management Science*, pages 163–208. Springer, Zürich, Switzerland.

Song, D., Zhang, J., Carter, J., Field, T., Marshall, J., Polak, J., Schumacher, K., Sinha-Ray, P., and Woods, J. (2005). On cost-efficiency of the global container shipping network. *Maritime Policy & Management*, 32(1):15–30.

Song, D. and Zhang, Q. (2010a). Optimal inventory control for empty containers in a port with random demands and repositioning delays. In Cullinane, K., editor, *International Handbook of Maritime Economics*, pages 301–321. Edward Elgar Publishing, Cheltenham, UK.

Song, D.-P. and Zhang, Q. (2010b). A Fluid Flow Model for Empty Container Repositioning Policy with a Single Port and Stochastic Demand. *SIAM Journal on Control and Optimization*, 48(5):3623–3642.

Svoboda, J., Minner, S., and Yao, M. (2021). Typology and literature review on multiple supplier inventory control models. *European Journal of Operational Research*, 293(1):1–23.

Toygar, A., Yildirim, U., and İnegöl, G. M. (2022). Investigation of empty container shortage based on SWARA-ARAS methods in the COVID-19 era. *European Transport Research Review*, 14(1):1–8.

Trivedi, P. and Zimmer, D. (2017). A Note on Identification of Bivariate Copulas for Discrete Count Data. *Econometrics*, 5(1):1–10.

Whittemore, A. S. and Saunders, S. C. (1977). Optimal Inventory Under Stochastic Demand with Two Supply Options. *SIAM Journal on Applied Mathematics*, 32(2):293–305.

Young Yun, W., Mi Lee, Y., and Seok Choi, Y. (2011). Optimal inventory control of empty containers in inland transportation system. *International Journal of Production Economics*, 133(1):451–457.

Zhang, B., Ng, C. T., and Cheng, T. C. E. (2014). Multi-period empty container repositioning with stochastic demand and lost sales. *Journal of the Operational Research Society*, 65(2):302–319.

Zipkin, P. (2008). Old and New Methods for Lost-Sales Inventory Systems. *Operations Research*, 56(5):1256–1263.

7 Appendix

The following proofs relate to the propositions in Section 3.8 for reducing the state and action spaces of the introduced MDP.

Proof of Proposition 1. Assume that there exists a t such that $O_{0,t} > \bar{I}$. Given $I_t, O_{1,t}, \ldots, O_{L-1,t}, X_t$, let $\hat{O}_{0,t} = \bar{I}$.
If $\lambda_p > 0$, then

$$\begin{aligned}&\lambda_p \left(\max\left(\max(I_t - A_t^- - Z_t^-, 0) + Z_t^+ + A_{0,t}^+ + O_{0,t} - \bar{I}, 0\right)\right)\\ &> \lambda_p \left(\max\left(\max(I_t - A_t^- - Z_t^-, 0) + Z_t^+ + A_{0,t}^+ + \hat{O}_{0,t} - \bar{I}, 0\right)\right)\end{aligned}$$

for all Z_t^+, Z_t^- and $A_{0,t}^+, A_t^-$ and for all $t = 1, \ldots, T-1$. All other costs are the same, as they do not depend on $O_{0,t}$.
The assumption $O_{0,t} > \bar{I}$ implies $I_{t+1} = \min\left(\max(I_t - A_t^- - Z_t^-, 0) + A_{0,t}^+ + O_{0,t} + Z_t^+, \bar{I}\right) = \bar{I}$ for all Z_t^+, Z_t^- and $A_{0,t}^+, A_t^-$ and for all $t = 1, \ldots, T-1$, and similarly for $\hat{O}_{0,t}$. Because V_{t+1} does not depend on $O_{0,t}$, $V_t(I_t, O_{0,t}, \ldots, O_{L-1,t}, Z_t) > V_t(I_t, \hat{O}_{0,t}, \ldots, O_{L-1,t}, Z_t)$. If C_T is decreasing in $O_{0,T}$, $C_T(O_{0,T}) \geq C_T(\hat{O}_{0,T})$. □

Proof of Proposition 2. The proposition follows immediately from the inventory transition function (3) □

Proof of Proposition 3. Given a state S_t with $I_t, O_{0,t}, X_t$ that satisfies $\bar{A}^- > I_t - \min \mathcal{Z}_{x_t}^- > 0$, where $\mathcal{Z}_{x_t}^-$ is the set of all demand realizations that occur with positive probability in period t as defined in Proposition 3. Assume that $A_t^- > I_t - \min \mathcal{Z}_{x_t}^-$ and $\hat{A}_t^- = I_t - \min \mathcal{Z}_{x_t}^-$. Then $\lambda^- > 0$ implies $\lambda^- \hat{A}_t^- < \lambda^- A_t^-$. Also, if $\lambda_l \geq 0$, then

$$\lambda_l \max(Z_t^- + \hat{A}_t^- - I_t, 0) \leq \lambda_l \max(Z_t^- + A_t^- - I_t, 0)$$

for all Z_t^-. All other costs are the same for out-positioning decisions A_t^- and $\hat{A}_t^-$, as $A_t^- > I_t - \mathcal{Z}_{x_t}^- \geq I_t - Z_t^-$, and so, $\max(I_t - Z_t^- - A_t^-, 0) = 0$ for all Z_t^-, and similarly for $\hat{A}_t^-$, which further implies equal transition probabilities for both decisions. Thus, $\hat{A}_t^-$ is a better solution than A_t^-. □

Proof of Proposition 4. Given a state S_t with $I_t, O_{0,t}, X_t$ that satisfies

$$\bar{A}_0^+ > \bar{I} - \big(\max(I_t - \max \mathcal{Z}_{x_t}^-, 0) + \min \mathcal{Z}_{x_t}^+ + O_{0,t}\big) \geq 0,$$

where $\mathcal{Z}_{x_t}^-$ and $\mathcal{Z}_{x_t}^+$ are respectively the empty container demand and return realizations in period t with positive probability as earlier defined in Proposition 3 and 4. Assume the in-positioning decisions with 0 lead time $A_{0,t}^+ > \bar{I} - \big(\max(I_t - \max \mathcal{Z}_{x_t}^-, 0) + \min \mathcal{Z}_{x_t}^+ + O_{0,t}\big)$ and $\hat{A}_{0,t}^+ = \bar{I} - \big(\max(I_t - \max \mathcal{Z}_{x_t}^-, 0) + \min \mathcal{Z}_{x_t}^+ + O_{0,t}\big)$. Then $\lambda_0^+ > 0$ implies $\lambda_0^+ \hat{A}_{0,t}^+ < \lambda_0^+ A_{0,t}^+$. Also, if $\lambda_p \geq 0$, then

$$\begin{aligned} &\lambda_p \max\Big(\max(I_t - Z_t^-, 0) + Z_t^+ + \hat{A}_{0,t}^+ + O_{0,t} - \bar{I}, 0\Big) \\ &\leq \lambda_p \max\big(\max(I_t - Z_t^-, 0) + Z_t^+ + A_{0,t}^+ + O_{0,t} - \bar{I}, 0\big) \end{aligned} \tag{7}$$

for all Z_t^+, Z_t^-. All other costs are the same for in-positioning decisions $A_{0,t}^+$ and $\hat{A}_{0,t}^+$. Moreover, both decisions have equal transition probabilities since $I_{t+1} = \bar{I}$ for all Z_t^+, Z_t^-.Thus, $\hat{A}_{0,t}^+$ is a better solution than $A_{0,t}^+$. □

Paper B

Direct multi-step ahead forecasting with state-space models

Publication details: Submitted to *International Journal of Forecasting*.

Direct multi-step ahead forecasting with state-space models

Benedikt Sommer[a,1], Klaus Holst[a], Pierre Pinson[b]

[a] *A. P. Møller-Mærsk, Copenhagen, Denmark*
[b] *Technical University of Denmark, Kongens Lyngby, Denmark*

Abstract

The selection of a multi-step ahead forecasting strategy is a long-standing problem. While it is known that the errors of an optimal h-step ahead forecast are serially correlated up to lag $h-1$, a common assumption for direct multi-step ahead forecasting models is to let the model innovations be serially uncorrelated. In this paper we show that this assumption can lead to estimation biases, and ultimately to forecast accuracy deteriorations, when time-varying coefficient models are estimated in state-space form for direct multi-step ahead forecasting. To enable direct multi-step ahead forecasting with state-space models we propose to explicitly model the latent serially correlated innovation process of a direct h-step ahead forecast as a MA$(h-1)$ process. We show on a real-world dataset that our proposed methodology produces on average more accurate probabilistic forecasts than the corresponding state-space models with serially uncorrelated innovations.

Keywords: State-space models, Direct multi-step ahead forecasting, Non-stationary time series, Probabilistic forecasting, Demand forecasting

1. Introduction

Forecasts are an integral component of many modern data-driven decision-making processes. Today it is widely accepted that optimal decision-making must reflect the inherent uncertainty of the future and therefore, forecasts should be thought of and issued within a probabilistic framework (Gneiting and Katzfuss, 2014). Besides being probabilistic, forecasts are often required for multiple lead times to enable decision-making for multiple steps in the future. Forecasters can choose from a variety of multi-step ahead forecasting strategies, with broad categorizations into single-output and multiple-output methods. The widely utilized iterated and direct strategies, and their less frequently applied variations, belong to the first category since the models output a single value for each lead time. Multiple-output strategies instead output forecasts for the whole forecast horizon simultaneously to account for the inter-dependencies of the forecasts for different lead times. Since none of the available strategies is superior in all conditions, it is problem- and model-specific conditions that determine multi-step ahead forecasting strategy to be applied. For example, while the

Email address: benedikt.sommer@maersk.com (Benedikt Sommer)

accumulation of forecast errors under model misspecification may discourage the usage of the iterated strategy from a theoretical standpoint, empirical studies have not collectively found the error accumulation avoiding direct strategy to be superior (Taieb and Atiya, 2015). Indeed, even if the direct strategy seems to bypass the error accumulation problem, there is a latent and serially correlated innovation process that is disregarded. We will hence explore here a more flexible approach within a state-space framework, allowing to perform direct multi-step ahead forecasting while aiming to accommodate the latent innovation process component.

In addition, a frequently encountered challenge in many forecasting problems is that time-series show non-stationary behaviour. Traditionally, stationary time series models have been applied to differenced series or to small time intervals which are assumed to be locally (or approximately) stationary (Dahlhaus, 2012). Time-varying coefficient models take a different approach by assuming the model coefficients to be time-dependent, hence the models are implicitly non-stationary (Grenier, 1983; Chen and Tsay, 1993). For example, time-varying coefficient ARMA-X models allow modelling processes with both time-varying effects of exogenous inputs and autocorrelation functions. The coefficient evolution through time can generally be accommodated in many ways. Examples include local regression methods (Chen and Tsay, 1993; Cai and Tiwari, 2000), recursive estimation techniques (Moulines et al., 2005; Messner and Pinson, 2019) and proposals within a state-space framework (Priestley, 1980; Durbin and Koopman, 2012). The latter framework offers a unified treatment for a wide range of non-stationary time series with missing values, structural breaks and seasonalities (Davis et al., 2021) – it is hence preferred here. Within the same framework it is possible to model univariate and multivariate time series based on different distributional assumptions (West et al., 1985).

We propose a novel direct probabilistic multi-step ahead forecasting methodology with state-space models for non-stationary time series. Our proposal relies on the knowledge that the errors of a direct multi-step ahead forecast with a fixed lead time but varying forecast origin are serially correlated, even when the mean structure of the forecast model is correctly specified (Harvey et al., 1997). To the best of our knowledge, this is the first time that this deficiency of dynamic models for direct multi-step ahead forecasting is discussed. For example, Poncela et al. (2013) used time-varying coefficient AR models in a state-space form for direct multi-step ahead forecasting, though overlooking the potential misspecification of the innovations process. Other selected examples from energy applications where the serial correlation in the multi-step ahead forecast errors is ignored during the modelling process include time-varying AR-X models that are either parameterized in state-space form (Sanchez, 2006) or recursively estimated (Bacher et al., 2009). We expect the application of dynamic models for multi-step ahead forecasting with potentially misspecified innovation processes not to be limited to these examples. Therefore, our main contribution is to expose an estimation bias in the coefficient estimates of dynamic models that do not account for the serial correlation in the multi-step ahead forecast errors. For state-space models we show based on simulation studies that the innovation process misspecification may cause a significant bias in the maximum likelihood estimates that ultimately degenerates the forecast performance.

Our work is inspired by the past success of modelling non-stationary time series with time-varying coefficient models (Karakatsani and Bunn, 2008; Song et al., 2011; Dangl and Halling, 2012). We exploit the flexibility of the state-space modelling framework to estimate a broad range of model types (Durbin and Koopman, 2012; Hyndman et al., 2008; Harvey, 1990) and propose a modular approach to multi-step ahead forecasting. It is based on the explicit modelling of the multi-step ahead forecast errors as a latent moving average process. A key contribution is then the proposal of relevant parameterizations and their evaluation in an empirical analysis with real-world data. We consider our approach to be modular since our methodology is generally compatible with already existing state-space model parameterizations, e.g., the time-varying AR and AR-X model parameterizations in Poncela et al. (2013) and Sanchez (2006), respectively.

The paper is structured as follows. Firstly, Section 2 motivates in detail our proposal while building on related work. The exposition of the main theoretical results on direct multi-step ahead forecasting for non-stationary time series is then provided in Section 3. Section 4 presents suitable model parameterizations and describes the associated state and parameter maximum likelihood estimation procedures. The proposed models and associated estimation framework are validated based on simulation experiments in Section 5. Section 6 gathers empirical evidence for the proposed framework before Section 7 concludes this paper by discussing the main findings and perspectives for future work.

2. Motivating multi-step ahead forecasting within a state-space framework

Multi-step ahead forecasting strategy selection is a long-standing challenge in time series forecasting. Attempts of theoretical comparisons between the iterated and direct strategy are presented in Atiya et al. (1999), Chevillon (2007) and Taieb and Atiya (2015) among others, with the finding that the direct strategy is preferred under model misspecification, i.e., when the set of candidate models does not contain the true model. In addition, Chevillon (2007) explained that direct multi-step ahead forecasting is likely superior in non-stationary environments because model misspecification may occur owing to unnoticed unit roots or non-stationary regressors. On contrast, short time series and forecasting problems with short horizons are favorable conditions in which the iterated strategy can produce superior forecasts (Ben Taieb et al., 2012).

Since there is no overall superior strategy it is often the specific characteristics of a forecasting problem that influence the selection of a multi-step ahead forecasting strategy. Consider a forecasting problem where domain knowledge suggests that the functional relationship between some regressors and the response variable depend on the forecast lead time. This may be the case for forecasting the weekly number of delivered products for a company which accepts product orders with future delivery times (Bartezzaghi et al., 1999). The functional relationship of the response variable with the regressors, which are derived from already received orders, will vary with the lead time under the assumption that the number of received orders decreases with the increased distance in delivery time. Throughout the remainder of this paper we use the term horizon-dependent to refer to this class of regressors.

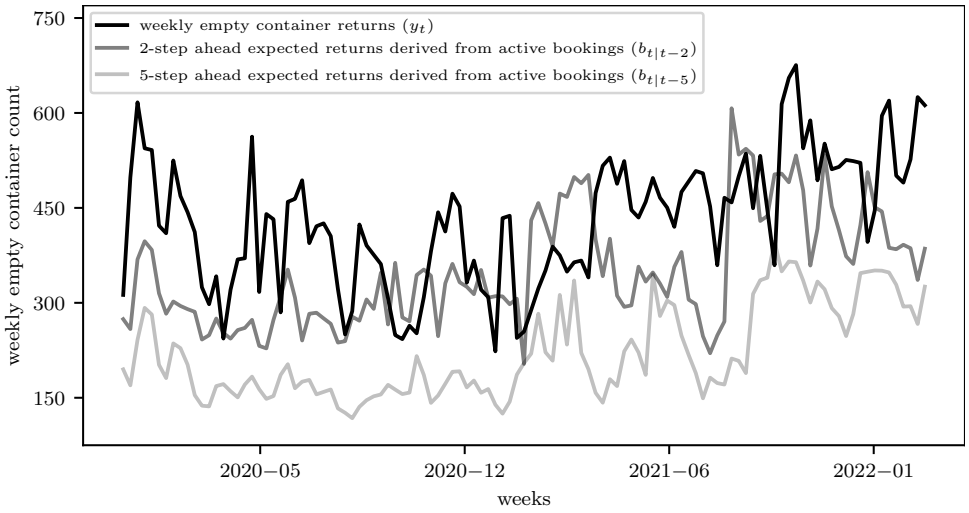

Figure 1: Example target variable time series and horizon-dependent regressors from the dataset that is considered in the empirical analysis of Section 6. The time series are aligned such that the observations y_t match with the regressors $b_{t|t-h}$, where h is the lead time measured in weeks.

Another example of a forecasting problem with horizon-dependent regressors is presented in the empirical analysis of Section 6. The objective is to forecast future empty container returns from customers to the container storage facilities of a large container shipping company. For this problem it is possible to compute highly predictive regressors by utilising the booking information of travelling full containers. The exemplary shown data in Figure 1 suggests that the functional relationship between the regressors and the response variable changes with the look ahead time. The reason for the seen behaviour is that bookings, which return the associated empty containers further ahead in the future, have not yet been observed at the forecast origin. Due to the changing relationship between the regressors and response variable it is therefore difficult to justify the application of the iterated strategy for producing multi-step ahead forecasts.

This forecasting problem has further similarities with many other demand forecasting problems where the time series often show non-stationary behaviour of changing mean, autocorrelation and variance. Other domains, where differencing or model coefficient re-estimation on small time intervals has traditionally been applied to non-stationary time series, are economics and energy forecasting applications. Many recent proposals in the forecasting literature take a different approach by avoiding stationarity considerations altogether. The time-invariant parameters of global models (Januschowski et al., 2020) are estimated on large data sets to allow parameter sharing among related time series even though individual time series may show different non-stationarity patterns. Successful methods based on deep neural networks and boosted trees are presented in Rangapuram et al. (2018); Salinas et al. (2020); Lim et al. (2021) and Ma and Fildes (2020), respectively. Non-stationarities may be tackled within the global modelling framework by parameter re-estimation on small time intervals. However, this approach faces difficulties in situations where individual time series need different interval lengths.

A benefit of local modelling techniques in this context is that the level of parameter adaptivity to address non-stationarities can be optimized for each time series individually. An example is the extension in terms of time-varying autoregressive coefficients of a stationary AR model to non-stationary time series (Grenier, 1983). While several estimation procedure exist today for time-varying coefficient models, we will focus in this paper on the state-space framework. A parameterization of a time-varying coefficient AR model in state-space form usually assumes the AR coefficients to be latent state variables that follow a first order Markov process (Durbin and Koopman, 2012). Efficient inference on the latent coefficients is often possible via Kalman-type filters and parameter estimation is conducted via maximum likelihood.

Model parameterizations are certainly not limited to these model types. The state-space framework generally allows to parameterize parsimonious structural time series models (Song et al., 2011; Durbin and Koopman, 2012) for forecasting problems with available domain knowledge and potentially little data. At the same time it is also possible to parameterize flexible non-linear models where sequential Monte Carlo methods are required for inference (Svensson and Schön, 2017). State-space models are considered extensions of generalized linear models to time series data. Thus they allow for appropriate model parameterizations of demand time series with positive integer counts where the Gaussian assumption is severely violated if the time series are zero-inflated (West et al., 1985; Seeger et al., 2016; Davis et al., 2021). An additional appealing property of state-space models and a deciding factor of their applicability in real-world forecasting problems is that the parameter estimation via maximum likelihood is efficient and fast for reasonably complex models. For example, Seeger et al. (2016) demonstrated the ability of state-space models to forecast hundred thousands of time series.

The state-space framework is conceptually designed to perform multi-step ahead forecasting by following the iterated strategy where unobserved regressors are replaced by their forecasts. Our novel methodology for direct multi-step ahead forecasting with state-space models utilises that the error sequence of the optimal h-step ahead forecast follows a moving average process of order $h-1$ (Harvey et al., 1997). Further, this result can be expected to hold approximately for any reasonably well-conceived set of forecasts (Harvey et al., 1997). A direct multi-step ahead model with independently distributed innovations is therefore clearly misspecified in its innovation process. Interestingly, one finds frequently applications of the direct multi-step ahead forecasting strategy where the serial correlation of the forecast errors is ignored. Examples for energy forecasting Sanchez (2006); Bacher et al. (2009); Poncela et al. (2013). As we show in this paper, time-varying coefficient models in state-space form are susceptible to estimation biases when the serial correlation of the forecast errors is grossly misspecified. An important aspect is that time-varying AR coefficients induce a time-dependence in the autocorrelation function of the h-step ahead forecast errors. The generality of estimating time-varying coefficient models with complex, serially correlated innovation processes sets the state-space framework apart from the referenced estimation methods.

Possible solutions to estimate models with serially correlated innovations in the state-space framework can be broadly categorized as follows. One can select filtering algorithms

that are derived under the assumption of serially correlated noise (Li et al., 2013; Sun et al., 2016). The coefficients of the autocorrelation function are static model coefficients and thus, it is not straightforward to model processes with time-varying autocorrelation functions. Alternatively, one augments the state-space model with a parametrized innovation process such that the resulting auxiliary dynamical system is driven by white noise (Bryson and Johansen, 1965). Standard filtering algorithms can then be used for state inference because the inputs to the system are white. We prefer this approach because it allows to model a wide range of time-varying autocorrelation functions in a flexible, yet parsimonious way. It is further compatible with state-space models that are already in use because one must only replace the observation noise process, while the other model components remaining unchanged. As shown in the following section, the flexibility is particularly needed for non-linear time series since the sequence of direct h-step ahead forecast errors experience complex autocorrelations.

3. On the serial correlation of multi-step ahead forecast errors

In this section we verify that the errors of a well-conceived h-step ahead forecast follows an approximate MA$(h-1)$ process. We show this on the basis of a general non-stationary autoregressive process before providing the generalisation to autoregressive processes with exogenous inputs. Thus, let us for now assume that a univariate time series $y_1, \dots, y_T$ is generated by the non-stationary autoregressive process

$$y_t = f_{\boldsymbol{\theta}_t}(\boldsymbol{x}_{t-1}) + \epsilon_t \,, \tag{1}$$

with $\boldsymbol{x}_{t-1} = (y_{t-1} \ \dots \ y_{t-d})^\top$ and where for generality f is assumed to be a non-linear function in the vector $\boldsymbol{x}_{t-1}$ of lagged endogenous variables. The additive innovations $\{\epsilon_t\}$ are assumed to be white noise, albeit the following results immediately extent to processes with serially correlated innovations. The sequence of model coefficients $\{\boldsymbol{\theta}_t\}$ that parameterizes the function f is restricted such that the stochastic process (1) is locally stationary (Dahlhaus, 2012). For the assumed data generating mechanism we now show that the errors of the optimal 2-step ahead forecast are serially correlated up to lag 1, followed by the generalization to h-step ahead forecasts and subsequently data generating mechanisms with exogenous inputs.

3.1. Properties of the 2-step ahead forecast errors

We begin by writing the stochastic process for y_{t+2} as a function of the available information set at time instance t and the future innovations ϵ_{t+1} and ϵ_{t+2}. By expansion of the data generating mechanism (1) we obtain

$$\begin{aligned} y_{t+2} &= f_{\boldsymbol{\theta}_{t+2}}(\boldsymbol{x}_{t+1}) + \epsilon_{t+2} \\ &= f_{\boldsymbol{\theta}_{t+2}}(y_{t+1}, \dots, y_{t-d+2}) + \epsilon_{t+2} \\ &= f_{\boldsymbol{\theta}_{t+2}}(f_{\boldsymbol{\theta}_{t+1}}(\boldsymbol{x}_t) + \epsilon_{t+1}, y_t, \dots, y_{t-d+2}) + \epsilon_{t+2}, \end{aligned} \tag{2}$$

where the first argument of $f_{\boldsymbol{\theta}_{t+2}}$ depends on the innovation term of time $t+1$. With a first order Taylor series expansion in the function's first argument around the point $f_{\boldsymbol{\theta}_{t+1}}(\boldsymbol{x}_t)$ we obtain

$$\begin{aligned} y_{t+2} &= f_{\boldsymbol{\theta}_{t+2}}(f_{\boldsymbol{\theta}_{t+1}}(\boldsymbol{x}_t), y_t, \dots, y_{t-d+2}) + \epsilon_{t+1}\frac{\partial f_{\boldsymbol{\theta}_{t+2}}(f_{\boldsymbol{\theta}_{t+1}}(\boldsymbol{x}_t))}{\partial x_1} \\ &\quad + \epsilon_{t+1}^2\frac{\partial f_{\boldsymbol{\theta}_{t+2}}(f_{\boldsymbol{\theta}_{t+1}}(\boldsymbol{x}_t))}{2\partial x_1^2} + o(\epsilon_{t+1}^3) + \epsilon_{t+2} \\ &= f_{\boldsymbol{\theta}_{t+2}}(f_{\boldsymbol{\theta}_{t+1}}(\boldsymbol{x}_t), y_t, \dots, y_{t-d+2}) + \tilde{\epsilon}_{t+2|t}, \end{aligned} \tag{3}$$

where $\frac{\partial}{\partial x_1}$ denotes the partial derivative with respect to the first argument in $f_{\boldsymbol{\theta}_{t+2}}$ and $o(\epsilon_{t+1}^3)$ denotes the higher order terms of the Taylor series expansion. The expansion shows that the sequence $\{\tilde{\epsilon}_{t+2|t}\}$ is at most serially correlated up to lag 1 due to the dependence of $\tilde{\epsilon}_{t+2|t}$ on the innovations ϵ_{t+1} and ϵ_{t+2}. Then, since $\tilde{\epsilon}_{t+2|t}$ additional depends on the time-varying autoregressive coefficients it follows that the autocorrelation structure of $\{\tilde{\epsilon}_{t+2|t}\}$ is time-varying too. To familiarize the reader with the previous exposition we exemplify both conditions in Example 1 for a time-varying coefficient AR(1) process and 2-step ahead forecasting.

Example 1 (2-step ahead forecast errors of a non-stationary AR(1) process). Let the time series $y_1, \dots, y_T$ follow the time-varying coefficient AR(1) process

$$y_t = \theta_t y_{t-1} + \epsilon_t, \tag{4}$$

where $\{\epsilon_t\}$ is Gaussian white noise with variance σ^2 and where $\{\theta_t\}$ follows a process such that the time series is locally stationary and for simplicity it is assumed that $y_0 = 0$. For the specified data generation mechanism we can write

$$\begin{aligned} y_{t+2} &= \theta_{t+2}\theta_{t+1}y_t + \theta_{t+2}\epsilon_{t+1} + \epsilon_{t+2} \\ &= \theta_{t+2}\theta_{t+1}y_t + \tilde{\epsilon}_{t+2|t}. \end{aligned} \tag{5}$$

from which, under the assumption that the sequence of model coefficients $\{\theta_t\}$ is known, we obtain

$$\text{Var}\left[\tilde{\epsilon}_{t+2|t}|\theta_{t+2}\right] = (\theta_{t+2}^2 + 1)\sigma^2, \tag{6}$$

and

$$\text{Cov}\left[\tilde{\epsilon}_{t+2|t}, \tilde{\epsilon}_{t+1|t-1}|\theta_{t+1}, \theta_{t+2}\right] = \theta_{t+2}\sigma^2. \tag{7}$$

Clearly, both quantities are time-varying due to the dependency on θ_{t+2}. The implication for the conditional variance is that uncertainty measures, such as prediction intervals, of the optimal forecast are necessarily time-varying. For this constructed example the forecast uncertainty is large when $|\theta_{t+2}|$ is large and small for a small autoregressive coefficient. The conditional covariance in Equation (7) verifies that the lag 1 autocorrelation of the optimal 2-step ahead forecast errors is also time-varying. Last, it is straightforward to verify that $\text{Cov}\left[\tilde{\epsilon}_{t+2|t}, \tilde{\epsilon}_{t+2-\tau|t-\tau}|\theta_1, \dots, \theta_{t+2}\right] = 0$ for $|\tau| > 1$.

3.2. Properties of the h-step ahead forecast errors

The generalisation to greater temporal differences between the forecast origin t and lead time $t+h$ can be derived via the following recursion. First write

$$y_{t+2} = \tilde{y}_{t+2|t} + \tilde{\epsilon}_{t+2|t}, \tag{8}$$

where $\tilde{y}_{t+2|t} = f_{\boldsymbol{\theta}_{t+2}}(f_{\boldsymbol{\theta}_{t+1}}(\boldsymbol{x}_t), y_t, \ldots, y_{t-d+2})$ denotes the term in Equation (3) which at time t is independent of the future innovations ϵ_{t+1} and ϵ_{t+2}. With the same notational convention we can write

$$\begin{aligned} y_{t+h} &= f_{\boldsymbol{\theta}_{t+h}}(y_{t+h-1}, \ldots, y_{t+h-d}) + \epsilon_{t+h} \\ &= f_{\boldsymbol{\theta}_{t+h}}(\tilde{y}_{t+h-1|t} + \tilde{\epsilon}_{t+h-1|t}, \ldots, \tilde{y}_{t+2|t} + \tilde{\epsilon}_{t+2|t}, f_{\boldsymbol{\theta}_{t+1}}(\boldsymbol{x}_t) + \epsilon_{t+1}, y_t \ldots, y_{t+h-d}) \\ &\quad + \epsilon_{t+h}, \end{aligned} \tag{9}$$

where $h < d$ is assumed for the sake of notational brevity. We continue as before with the first order Taylor series expansion in the function's first $h-1$ arguments around the point $(\tilde{y}_{t+h-1|t}, \ldots, \tilde{y}_{t+2|t}, f_{\boldsymbol{\theta}_{t+1}})$ to obtain

$$\begin{aligned} y_{t+h} &= f_{\boldsymbol{\theta}_{t+h}}(\tilde{y}_{t+h-1|t}, \ldots, \tilde{y}_{t+2|t}, f_{\boldsymbol{\theta}_{t+1}}(\boldsymbol{x}_t), y_t, \ldots, y_{t+h-d}) + \tilde{\epsilon}_{t+h|t} \\ &= \tilde{y}_{t+h|t} + \tilde{\epsilon}_{t+h|t}. \end{aligned} \tag{10}$$

Because all $\tilde{\epsilon}_{t+j|t}$ for $j = 2, \ldots, h-1$ in Equation (9) are respectively functions of $\epsilon_{t+1}, \ldots, \epsilon_{t+j}$ and $\boldsymbol{x}_t$, it follows that $\tilde{\epsilon}_{t+h|t}$ is also functionally dependent on $\boldsymbol{x}_t$ and $\epsilon_{t+1}, \ldots, \epsilon_{t+h}$. Consequently $\tilde{\epsilon}_{t+h|t}$ can be at most serially dependent up to lag $h-1$, with the dependence structure varying with time due to a functional dependence on $\boldsymbol{\theta}_{t+1}, \ldots, \boldsymbol{\theta}_{t+h-1}$.

The serial dependence structure for $\{\tilde{\epsilon}_{t+h|t}\}$ remains generally intact as long as the data generating mechanism is autoregressive. For example, let the time series $y_1, \ldots, y_T$ be generated by the stochastic process

$$y_t = f_{\boldsymbol{\theta}_t}(\boldsymbol{x}_{t-1}, \boldsymbol{u}_{t-1}) + \epsilon_t, \tag{11}$$

where $\boldsymbol{x}_{t-1}$ and $\{\epsilon_t\}$ as before. The non-linear data generating function f depends now additional on a vector $\boldsymbol{u}_t = (u_{1,t} \ \ldots \ u_{k,t})^\top$ of exogenous variables. Following the previous steps we can write again

$$\begin{aligned} y_{t+h} &= f_{\boldsymbol{\theta}_{t+h}}(\tilde{y}_{t+h-1|t}, \ldots, \tilde{y}_{t+2|t}, f_{\boldsymbol{\theta}_{t+1}}(\boldsymbol{x}_t, \boldsymbol{u}_t), y_t, \ldots, y_{t+h-d}, \boldsymbol{u}_{t+h-1}) + \tilde{\epsilon}_{t+h|t} \\ &= \tilde{y}_{t+h|t} + \tilde{\epsilon}_{t+h|t}, \end{aligned} \tag{12}$$

where $\tilde{\epsilon}_{t+h|t}$ is now a more complicated construction that additionally also depends on $\boldsymbol{u}_t, \ldots, \boldsymbol{u}_{t+h-2}$. However, the dependence on $\epsilon_{t+1}, \ldots, \epsilon_{t+h}$ remains and thus, the previous statements on the serial dependence of $\{\tilde{\epsilon}_{t+h|t}\}$ also remain valid.

An implication of the serial dependency up to lag $h-1$ is that $\tilde{\epsilon}_{t+h|t}$ follows a MA $(h-1)$ process, where the MA coefficients are complicated functions of observed data and unobserved model coefficients when f is non-linear. In practice however this circumstance is often ignored when estimating a h-step ahead forecast model to the unknown data generating mechanism by assuming the additive innovations $\tilde{\epsilon}_{t+h|t}$ to be independently, and

frequently also identically, distributed. Examples where the innovations of h-step ahead dynamic models are assumed to serially independent can be found in Sanchez (2006); Bacher et al. (2009); Poncela et al. (2013); Messner and Pinson (2019). As we will show later on in Section 5 for simulated data, the consequences of the serial independence assumption can be severe when using state-space models for direct h-step ahead forecasting.

4. State-space modelling

Up until this point we have used the previous data generating mechanism as a mean to explore theoretical properties of the h-step ahead forecasts. Because the data generating mechanism remains unknown to a forecaster we propose in the following a general methodology for the estimation of dynamic models for multi-step ahead forecasting in non-stationary environments. It is assumed that the non-stationary data generating mechanism of the observed time series $y_1, \dots, y_T$ can be sufficiently well approximated by a time-varying coefficient model. Moreover, the forecaster has access to a set of forecast horizon-dependent exogenous regressors $\boldsymbol{u}_{t|t-h}$, which are assumed to be available for the respective lead time $T+h$. The notational change from $u_{t+h|t}$ to $u_{t|t-h}$ is intentional to let the presented state-space model parameterizations follow the standard notational conventions within the state-space literature (Harvey, 1990; Hyndman et al., 2008; Durbin and Koopman, 2012). Genuine h-step ahead probabilistic forecasts, outside the state-space model parameterizations, for lead time $t+h$ based on the information set that is available at the forecast origin t are denoted by $\hat{F}_{t+h|t}$.

Under these assumptions we propose to approximate the unknown stochastic process for y_t as a function of at time $t-h$ available endogenous and exogenous regressors with the general state-space model

$$\boldsymbol{\theta}_{t+1} = f_h(\boldsymbol{\theta}_t, \boldsymbol{\eta}_{t+1}), \tag{13a}$$

$$y_t = g_h(\boldsymbol{\theta}_t, \boldsymbol{x}_{t-h}, \boldsymbol{u}_{t|t-h}) + \tilde{\epsilon}_{t|t-h}, \tag{13b}$$

where $t = \tau, \dots, T$ due to the unavailability of endogenous regressors $\boldsymbol{x}_{t-h} = (y_{t-h} \ \dots \ y_{t-h-d})^\top$ for $t < \tau = d + h$. The underlying idea of the model is that f_h parameterizes together with $\boldsymbol{\eta}_{t+1}$ a process for the evolution of the unobserved time-varying (regression) coefficients $\boldsymbol{\theta}_t$ through time. The observation model (13b) is responsible for mapping the latent *state* vector $\boldsymbol{\theta}_t$ into the observational space of y_t, which is decomposed into a system model function $g_h(\cdot)$ of endogenous and exogenous regressors, and an additive error term $\tilde{\epsilon}_{t|t-h}$. We assume that both functions $g_h : \mathbb{R}^m \times \mathbb{R}^d \times \mathbb{R}^k \to \mathbb{R}$ and $f_h : \mathbb{R}^m \times \mathbb{R}^r \to \mathbb{R}^m$ are smooth and potentially non-linear. Last, the initial state vector is assumed to follow $\boldsymbol{\theta}_{\tau-1} \sim p(\boldsymbol{\theta}_{\tau-1})$, with $p(\boldsymbol{\theta}_{\tau-1})$ being known and potentially degenerate.

While the evolution through time of the state vector is assumed to be driven by the white noise sequence $\{\boldsymbol{\eta}_t\}$, it has been shown that this assumption is invalid for $\tilde{\epsilon}_{t|t-h}$ when $h > 1$. We utilise instead the previous results and assume that $\tilde{\epsilon}_{t|t-h}$ follows an approximate MA$(h-1)$ model conditional on the system model $g_h(\cdot)$ and the state process (13a) being appropriate approximations to the unknown data generating mechanism. Moreover, since

we identified the MA coefficients to be time-varying for data generating mechanisms with time-varying autoregressive coefficients we consider the h-step ahead innovations to follow

$$\tilde{\epsilon}_{t|t-h} = \epsilon_t + \sum_{j=1}^{h-1} \phi_{t,j}\epsilon_{t-j}, \tag{14}$$

where $\{\epsilon_t\}$ is white noise. For generality the MA coefficients are also assumed to be latent and to follow a first order Markov process

$$\boldsymbol{\phi}_{t+1} = f_h^{\text{MA}}(\boldsymbol{\phi}_t, \boldsymbol{\xi}_{t+1}), \tag{15}$$

where $\boldsymbol{\phi}_t = (\phi_{t,1} \ \dots \ \phi_{t,h-1})^\top$ and $\{\boldsymbol{\xi}_t\}$ being white noise. The MA coefficient time progression model $f_h^{\text{MA}}(\cdot)$ is a smooth function which may be non-linear in its arguments, $f_h^{\text{MA}} : \mathbb{R}^{h-1} \times \mathbb{R}^s \to \mathbb{R}^{h-1}$. While the innovations $\boldsymbol{\xi}_t$ and $\boldsymbol{\eta}_t$ that drive the evolution of the latent states may be cross-correlated, we assume that ϵ_t of the MA process (14) is at all times independent of the state innovations.

Generally, the assumed process for $\tilde{\epsilon}_{t|t-h}$ is a modelling choice. For example, one may consider the time-invariant coefficient MA process

$$\tilde{\epsilon}_{t|t-h} = \epsilon_t + \sum_{j=1}^{h-1} \phi_j\epsilon_{t-j}, \tag{16}$$

as a less flexible, but posibbly computationally more robust, alternative. The additive decomposition of the observation model (13b) allows immediately for direct multi-step ahead forecasting. It is merely necessary to choose a parameterization for $\tilde{\epsilon}_{t|t-h}$, which in many cases is independent of the choices for $f_h(\cdot)$ and $g_h(\cdot)$. In the remainder of this section we will present the general estimation treatment for the considered state-space model under the most general assumption that $\tilde{\epsilon}_{t|t-h}$ follows the time-varying coefficient MA process (14).

4.1. Estimation

It follows from the observation model (13b) that forecasting a future dependent variable Y_{T+h} requires the estimation of the latent, dynamically evolving states $\boldsymbol{\theta}_{T+h}$ and $\boldsymbol{\phi}_{T+h}$ conditional on the information set on the stochastic process at time instance T. Optimal and computational efficient inference via recursive filters with closed form analytical expressions is often only possible in special cases. For example, the Kalman filter (Kalman, 1960) is optimal for the class of stationary linear state-space models with Gaussian innovations. It is computationally efficient since it recursively estimates the sufficient statistics of the latent state's conditional distribution given the current information set on the stochastic process.

For non-linear and non-Gaussian models it is rarely possible to derive analytical expressions for the targeted conditional distributions. Inference for these models is usually performed by approximating the filter distributions with sequential Monte Carlo methods (Doucet et al., 2001) or by placing assumptions on the filter distributions to allow for their efficient approximation. While Monte Carlo methods allow for arbitrary close approximations of the respective conditional distributions by increasing their sample sizes, they are

difficult to apply in practical forecasting problems of reasonable size due to their computational complexity. We therefore decide to trade off accuracy losses due to approximations with desirable computational properties and propose in Section 4.1.1 the application of the Unscented Kalman filter (Julier and Uhlmann, 2004) for state inference.

A second inference task involves the unknown model coefficients of the functions $f_h(\cdot)$, $f_h^{\text{MA}}(\cdot)$ and $g_h(\cdot)$, and distributional coefficients of the white noise sequences $\{\boldsymbol{\eta}_t\}, \{\boldsymbol{\xi}_t\}$ and $\{\epsilon_t\}$. We propose to estimate the unknown coefficients of the state-space model by maximizing the observed data likelihood. We consider the σ-field $\mathcal{F}_t^h = \sigma(Y_1, \dots, Y_{t-h}, U_{1|1-h}, \dots, U_{t|t-h})$ and we assume that conditional on this reduced filtration the stochastic process $\{Y_t\}_{t=1,\dots,T}$ is distributed according to the state-space model described by equations (13a), (13b), (14) and (15), with (partial) likelihood given by

$$\begin{aligned} L_h(y_1, \dots, y_T, \boldsymbol{u}_{1|1-h}, \dots, \boldsymbol{u}_{T+h|T}) = p_h(y_1, \dots, y_{\tau-1} | \boldsymbol{u}_{1|1-h}, \dots, \boldsymbol{u}_{h|\tau-h}) \\ \times \prod_{t=\tau}^{T} p_h(y_t \mid y_1, \dots, y_{t-1}, \boldsymbol{u}_{1|1-h}, \dots, \boldsymbol{u}_{t+h|t}), \end{aligned} \tag{17}$$

where the subscript h indicates that the density is implied by the filtration $\{\mathcal{F}_t^h\}_{t=1,\dots,T}$. Finding an analytical expression for this likelihood is generally challenging when the underlying state-space model is non-linear and non-Gaussian. Our proposal in Section 4.1.2 is therefore to estimate the unknown state-space model coefficients by maximizing a Gaussian likelihood as an approximation to (17).

4.1.1. State estimation

The general state filtering problem concerns the estimation of the latent states conditional on the current available information set on the stochastic process. As we showed before the stochastic process that is obtained by conditioning on the reduced filtration is assumed to follow a state-space model which is described by Equations (13a), (13b), (14) and (15). The implied state-space model assumption on the stochastic process allows us subsequently to apply standard filtering techniques by first defining the full information set on the stochastic process at time t as the σ-field

$$\begin{aligned} \mathcal{F}_t &= \sigma(Y_\tau, \dots, Y_t, \boldsymbol{X}_{\tau|\tau-h}, \dots, \boldsymbol{X}_{t|t-h}, \boldsymbol{U}_{\tau|\tau-h}, \dots, \boldsymbol{U}_{t+h|t}) \\ &= \sigma(Y_1, \dots, Y_t, \boldsymbol{U}_{\tau|\tau-h}, \dots, \boldsymbol{U}_{t+h|t}), \end{aligned} \tag{18}$$

which follows from $\boldsymbol{X}_{t|t-h} \coloneqq (Y_{t-h}\ Y_{t-h-1}\ \dots\ Y_{t-h-d+1})^\top$ and the assumption of the exogenous regressors $\boldsymbol{u}_{t+h|t}$ to be known at time instance t for lead time $t+h$. The filtering problem for the specified model concerns then the estimation of the conditional probability density $p(\boldsymbol{a}_t|\mathcal{F}_t)$, where $\boldsymbol{\alpha}_t = (\boldsymbol{\theta}_t^\top\ \boldsymbol{\phi}_t^\top)^\top$ is the complete state vector of the system. It is generally recognized that it is impractical, or at times infeasible, to derive an analytical expression with a finite number of coefficients for the filter density $p(\boldsymbol{\alpha}_t|\mathcal{F}_t)$ when the observation and state processes are non-linear, and the innovation processes are non-Gaussian. In this paper we use the Unscented Kalman filter (UKF) and approximate the unknown conditional distribution $p(\boldsymbol{\alpha}_t|\mathcal{F}_t)$ by its first two moments. For the applicability of the filter we assume in

the following that the noise sequences $\{\boldsymbol{\eta}_t\}, \{\boldsymbol{\xi}_t\}$ and $\{\epsilon_t\}$ are mutually independent mean zero white noise sequences with finite, known variance.

The filter is computationally attractive because the mean $\mathrm{E}\left[\boldsymbol{\alpha}_t|\mathcal{F}_t\right]$ and variance $\mathrm{Var}\left[\boldsymbol{\alpha}_t|\mathcal{F}_t\right]$ of the distribution are approximated with a small number of carefully chosen sample points, the so-called σ-points. A deterministic sampling approach is used to generate these points, which distinguishes the UKF from sequential Monte Carlo methods where random samples from the system noise distributions are required. In this paper we follow the sampling procedure of Algorithm 1, which is found in Julier and Uhlmann (1997).

Algorithm 1 (σ-point sampling and transformation procedure). The m-dimensional random variable $\boldsymbol{x}$ with mean $\bar{\boldsymbol{x}}$ and covariance $\boldsymbol{P}_{xx}$ is approximated by $2m+1$ weighted points given by

$$\begin{aligned} \boldsymbol{\mathcal{X}}_0 &= \bar{\boldsymbol{x}} & W_0 &= \kappa/(m+\kappa) \\ \boldsymbol{\mathcal{X}}_i &= \bar{\boldsymbol{x}} + \left(\sqrt{(m+\kappa)\boldsymbol{P}_{xx}}\right)_i & W_i &= \kappa/2(m+\kappa) \\ \boldsymbol{\mathcal{X}}_{i+m} &= \bar{\boldsymbol{x}} - \left(\sqrt{(m+\kappa)\boldsymbol{P}_{xx}}\right)_i & W_{i+m} &= \kappa/2(m+\kappa) \end{aligned}$$

where $\kappa \in \mathbb{R}$ is a tuning parameter and $\left(\sqrt{(m+\kappa)\boldsymbol{P}_{xx}}\right)_i$ is the ith column of the matrix square root of $\sqrt{(m+\kappa)\boldsymbol{P}_{xx}}$, which throughout this paper is obtained via the Cholesky decomposition. With W_i being the weight associated with the ith point, the transformation procedure is given by:

1. Instantiation of sigma points through the system model to obtain the transformed sigma points,
$$\boldsymbol{\mathcal{Y}}_i = f(\boldsymbol{\mathcal{X}}_i)$$
2. The mean and covariance of the transformed points are obtained as
$$\begin{aligned} \bar{\boldsymbol{y}} &= \sum_{i=0}^{2m} W_i \boldsymbol{\mathcal{Y}}_i \\ \boldsymbol{P}_{yy} &= \sum_{i=0}^{2m} W_i \left(\boldsymbol{\mathcal{Y}}_i - \bar{\boldsymbol{y}}\right)\left(\boldsymbol{\mathcal{Y}}_i - \bar{\boldsymbol{y}}\right)^\top \end{aligned}$$

Subsequently, we let $\hat{\boldsymbol{\alpha}}_{t|t}$ denote the UKF estimate of $\mathrm{E}\left[\boldsymbol{\alpha}_t|\mathcal{F}_t\right]$ and $\hat{\boldsymbol{\alpha}}_{t|t-1}$ be the estimate of $\mathrm{E}\left[\boldsymbol{\alpha}_t|\mathcal{F}_{t-1}\right]$. To elaborate on the failure of the filter for direct h-step ahead forecasting under model misspecification we will in the following only consider the time propagation of the state mean estimates. For the detailed exposition of the filter the interested reader is referred to Julier and Uhlmann (2004) and Wan and Van Der Merwe (2000).

In line with the Kalman filter, the UKF updates the state mean estimate $\hat{\boldsymbol{\alpha}}_{t|t-1}$ with the information set on the process at time instance t as

$$\hat{\boldsymbol{\alpha}}_{t|t} = \hat{\boldsymbol{\alpha}}_{t|t-1} + \boldsymbol{K}_t\left(y_t - \hat{y}_t\right), \tag{19}$$

where $\boldsymbol{K}_t$ is the gain term of the UKF and $\hat{y}_t$ denotes the approximation to the filter's internal 1-step ahead prediction $\mathrm{E}\left[y_t|\mathcal{F}_{t-1}\right]$, which is not to be mistaken for a genuine h-step prediction of the respective state-space model. Following the derivations in Julier and Uhlmann (2004) and Wan and Van Der Merwe (2000) we note that the UKF is derived for state-space models which assume the process and observation noise sequences to be white. An immediate implication of the white observation noise sequence is that the sequence of internal 1-step ahead prediction errors $v_t = y_t - \hat{y}_t$ must also be serially uncorrelated for an approximately correctly specified state-space model to the unknown data generating mechanism. As we showed before, the proposed state-space models for direct h-step ahead forecasting do only well approximate the data generation mechanism when the innovations term $\tilde{\epsilon}_{t|t-h}$ is well specified, i.e. follows an MA$(h-1)$ process with unknown coefficients. We demonstrate empirically in Section 5 that it is precisely this misspecification of the innovations process for state-space models with a correctly specified system model $g_h(\cdot)$ which introduces an estimation bias.

Since the UKF assumes by design white observation noise, a practical challenge is the estimation of the proposed models with finite lag serial correlated observation noise. At this stage we propose to use the state augmentation approach which was originally proposed in Bryson and Johansen (1965) for continuous-time dynamical systems. The idea is to augment the state vector $\boldsymbol{\alpha}_t = (\boldsymbol{\theta}_t^\top \ \boldsymbol{\phi}_t^\top)^\top$ with the vector of the autocorrelated observation noise $\boldsymbol{\epsilon}_{\star,t} = (\tilde{\epsilon}_{t-h|t} \ \epsilon_{t-1} \ \dots \ \epsilon_{t-h+1})^\top$ such that the evolution through time of the resulting auxiliary dynamical system is driven by white noise. The observation equation of the reparameterized model becomes then noise free because the augmented state vector $\boldsymbol{\alpha}_{\star,t} = (\boldsymbol{\theta}_t^\top \ \boldsymbol{\phi}_t^\top \ \boldsymbol{\epsilon}_{\star,t}^\top)^\top$ includes the autocorrelated noise $\tilde{\epsilon}_{t|t-h}$ of the original parameterization. The noise free observation equation of the reparameterized model is exemplary shown for a linear and Gaussian state-space model in Example 2.

Example 2 (State augmentation method for finite lag autocorrelated observation noise.). Consider the linear and Gaussian state-space model

$$\begin{aligned} \boldsymbol{\theta}_{t+1} &= \boldsymbol{T}\boldsymbol{\theta}_t + \boldsymbol{R}\boldsymbol{\eta}_{t+1} \\ \tilde{\epsilon}_{t+1} &= \phi\epsilon_t + \epsilon_{t+1} \\ y_t &= \boldsymbol{z}^\top\boldsymbol{\theta}_t + \tilde{\epsilon}_t, \end{aligned} \tag{20}$$

where $\{\boldsymbol{\eta}_t\}$ and $\{\epsilon_t\}$ are mutually independent Gaussian white noise sequences with known variances. The initial state vector $\boldsymbol{\theta}_0$ is assumed to be multivariate Gaussian with known mean $\boldsymbol{\mu}_0$ and variance $\boldsymbol{P}_0$. The MA coefficient ϕ, the transition matrix $\boldsymbol{T}$, the process noise selection matrix $\boldsymbol{R}$ and the emission vector $\boldsymbol{z}$ are assumed to be known. For the MA(1) process we can write

$$\begin{pmatrix} \tilde{\epsilon}_{t+1} \\ \epsilon_{t+1} \end{pmatrix} = \underbrace{\begin{pmatrix} 0 & \phi \\ 0 & 0 \end{pmatrix}}_{:=\boldsymbol{T}_{MA}} \underbrace{\begin{pmatrix} \tilde{\epsilon}_t \\ \epsilon_t \end{pmatrix}}_{:=\boldsymbol{\epsilon}_{\star,t}} + \underbrace{\begin{pmatrix} 1 \\ 1 \end{pmatrix}}_{:=\boldsymbol{r}_{MA}} \epsilon_{t+1} \tag{21}$$

which is taken to be the auxiliary dynamical system in the augmented state-space model

with noise free observation equation

$$\begin{aligned}\begin{pmatrix}\boldsymbol{\theta}_{t+1}\\ \boldsymbol{\epsilon}_{\star,t+1}\end{pmatrix} &= \begin{pmatrix}\boldsymbol{T} & \mathbf{0}\\ \mathbf{0} & \boldsymbol{T}_{MA}\end{pmatrix}\begin{pmatrix}\boldsymbol{\theta}_{t}\\ \boldsymbol{\epsilon}_{\star,t}\end{pmatrix} + \begin{pmatrix}\boldsymbol{R} & \mathbf{0}\\ \mathbf{0} & \boldsymbol{r}_{MA}\end{pmatrix}\begin{pmatrix}\boldsymbol{\eta}_{t+1}\\ \epsilon_{t+1}\end{pmatrix}\\ y_t &= \begin{pmatrix}\boldsymbol{z}^\top & \boldsymbol{z}_{MA}^\top\end{pmatrix}\boldsymbol{\alpha}_t,\end{aligned} \tag{22}$$

where $\boldsymbol{\alpha}_t = \begin{pmatrix}\boldsymbol{\theta}_t^\top & \boldsymbol{\epsilon}_{\star,t}^\top\end{pmatrix}^\top$ is the unobserved state vector and $\boldsymbol{z}_{MA} = \begin{pmatrix}1 & 0\end{pmatrix}$.

4.1.2. Model coefficient estimation and model selection

To this end we have assumed that the state-space model coefficients are known, which is generally not the case for real world forecasting problems. Performing maximum likelihood estimation for non-linear and non-normal state-space models often relies on approximations because exact analytical likelihood expressions are difficult to derive. In this paper we assume that the true model likelihood (17) is sufficiently well approximated by a Gaussian likelihood which is factorized as

$$\tilde{L}(\boldsymbol{Y}_T, \boldsymbol{U}_T; \boldsymbol{\psi}) = \prod_{t=\tau}^{T} p(y_t|\boldsymbol{Y}_{t-1}, \boldsymbol{u}_{t|t-h}), \tag{23}$$

where $\boldsymbol{Y}_T = (y_1 \ \dots \ y_T)^\top$, $\boldsymbol{U}_T = (\boldsymbol{u}_{\tau|\tau-h}^\top \ \dots \ \boldsymbol{u}_{T|T-h}^\top)^\top$ and $\boldsymbol{\psi}$ denotes the unknown coefficients of the state-space model. The factorization follows from $\boldsymbol{u}_{t|t-h}$ being purely exogenous. We propose this approximation because the mean and variance of the likelihood terms $p(y_t|\boldsymbol{Y}_{t-1}, u_{t|t-h})$ are estimated by the UKF during the latent state estimate propagation. With $\hat{y}_t$ being the UKF's estimate of $\mathrm{E}\left[y_t|\mathcal{F}_t\right]$ and $\hat{\sigma}_t$ being the corresponding estimate of $\mathrm{Var}\left[y_t|\mathcal{F}_t\right]$, we obtain

$$\log \tilde{L}(\boldsymbol{Y}_T, \boldsymbol{U}_T; \boldsymbol{\psi}) \propto -\frac{1}{2}\sum_{t=\tau}^{T}\left(\log \hat{\sigma}_t + \frac{(y_t - \hat{y}_t)^2}{\hat{\sigma}_t}\right). \tag{24}$$

In our experiments we estimate the unknown state-space model coefficients $\boldsymbol{\psi}$ by numerically maximizing Equation (24). We show empirically in Section 5 and Section 6 that the approximation of the true likelihood by this Gaussian likelihood in its prediction error decomposition works sufficiently well in practice because the point forecasts and its uncertainty estimates are jointly considered.

For the general case when a set of candidate state-space models are available we propose to perform model selection based on the attained maximum values of (24). This allows among other possible criteria to perform either likelihood ratio tests for nested models or base the selection on penalized likelihood scores such as the AIC or BIC for arbitrary sets of candidate models. Basing the model selection on the log-likelihood value is particular appealing for our proposed models since each state-space model in its direct multi-step ahead parametrization is an approximation to the true data generating mechanism. Therefore, models with misspecified observation innovations processes are generally expected to attain

smaller likelihood values. The model selection can further be supplemented by testing the filter's internal 1-step ahead prediction error sequence $\{y_t - \hat{y}_t\}_{t=\tau,\dots,T}$ for serial correlation. As beforehand discussed, the error sequence should be white for well fitting models since each state-space model aims to approximate the true stochastic process. For additional state-space model diagnostics we refer the interested reader to Durbin and Koopman (2012).

4.2. Forecasting

With the information set specified in Equation (18) we are interested in obtaining the distributional forecast as the estimation of $F_{T+h|T} = \mathcal{L}(Y_{T+h}|\mathcal{F}_T)$. In line with the previous expositions we may assume that this conditional distribution is sufficiently well described by its first two moments. Following the discussion in Durbin and Koopman (2012) for obtaining forecasts with the KF, we obtain estimates of $\mathrm{E}\left[Y_{T+h}|\mathcal{F}_T\right]$ and $\mathrm{Var}\left[Y_{T+h}|\mathcal{F}_T\right]$ by continuing the UKF beyond time instance $t = T$ while treating the observations $Y_{T+1}, \dots, Y_{T+h-1}$ as missing. When following this approach, prediction intervals can be obtained via distributional assumptions on $\hat{F}_{T+h|T}$.

Alternatively, one may use Monte Carlo sampling to obtain a non-parametric estimate of $\mathcal{L}(Y_{T+h}|\mathcal{F}_T)$. To this end the requirement for the application of the UKF has been that the noise sequences $\{\boldsymbol{\eta}_t\}, \{\boldsymbol{\xi}_t\}$ and $\{\epsilon_t\}$ are mean 0 white noise sequences with finite, known variances. To apply ancestral sampling we make a formal distributional assumption for each innovation term. A set of Monte Carlo samples for $\mathcal{L}(\boldsymbol{\alpha}_{T+h}|\mathcal{F}_T)$ is obtained by sampling from the innovations distributions for $t = T+1, \dots, T+h$ and subsequently projecting the latest state vector estimate $\hat{\boldsymbol{\alpha}}_{T|T}$ forward in time via the specified process model (Equations (13a), (14) and (15)). Using these samples together with the observation model $g_h(\cdot)$ provides one with samples of the forecast distribution $\hat{F}_{T+h|T}$. The benefit of this approach is that no assumptions about the forecast distribution are required to derive prediction intervals.

5. Monte Carlo experiments

We use simulated time series in this section to demonstrate that a large misspecification in the innovations process of direct h-step ahead state-space model parametrizations induces a bias in the maximum likelihood estimates of the unknown model coefficients. To begin with, we consider stationary time series for which we estimate time-varying coefficient AR models. We show that the estimation bias is significant for long forecast horizons and when the generated time series exhibit large sample autocorrelations. As a consequence of the large estimation biases the misspecified models achieve poor empirical forecast accuracy. It is then shown that the model parametrizations with MA$(h-1)$ innovations processes regain the accuracy of the oracle forecast. The results are then extended to a non-stationary AR(1) data generation process where the true AR coefficient is time-varying.

5.1. Stationary data generation process

The purpose of this experiment is to support the previous claims that state estimates, and ultimately the maximum likelihood estimates, are biased when the innovations process

of a state-space model is misspecified for direct multi-step ahead forecasting. To showcase this empirically we consider each time series in the set $\mathcal{I}_\rho$ to follow the stationary AR process

$$y_t = \rho y_{t-1} + \zeta_t, \quad \zeta_t \overset{iid}{\sim} \mathcal{N}(0,1), \tag{25}$$

where $t = 1, \ldots, 2000$ and $y_0 \sim N(0, \frac{1}{1-\rho^2})$. We generate 1000 replicates from this process for a range of $\rho \in (-1, 1)$, i.e. $\left|\mathcal{I}_\rho\right| = 1000$ where the subscript indicates the respective AR coefficient. The first 1000 observations of each time series are used for maximum likelihood estimation of the unknown state-space model coefficients while the last 1000 observations are withhold to evaluate the forecast accuracy later on. For demonstrating the estimation bias we misspecify a state-space model for this data generation mechanism by considering

$$\begin{aligned} \theta_{t+1} &= \theta_t + \eta_{t+1}, \quad \eta_{t+1} \overset{iid}{\sim} \mathcal{N}(0, \sigma_\eta^2) \\ y_t &= y_{t-h}\theta_t + \tilde{\epsilon}_{t|t-h} \\ \theta_\tau &\sim \mathcal{N}(\mu_\tau, \sigma_\tau^2) \end{aligned} \tag{26}$$

where h is the look-ahead time of the model and $t \geq \tau = h + 1$.

The model is purposely misspecified by letting the latent AR coefficient follow a first-order random walk. This parameterization is chosen because it is often applied in practice since the latent coefficient can move freely under the random walk assumption, thus capture potential temporary and permanent shifts in the true coefficient. The model is further misspecified by letting the observation noise $\tilde{\epsilon}_{t|t-h}$ be Gaussian white noise with variance $\sigma_{\tilde{\epsilon}}^2$. The model formulation is completed by taking $\mu_\tau = \rho^h$ and $\sigma_\tau^2 = 0.0025$, which will reduce the effect of the initial state value on the maximum likelihood estimates. Since the model is linear and Gaussian we use the Kalman filter for state inference and numerically maximize the Gaussian likelihood (24) to obtain estimates for the variance coefficients $\boldsymbol{\psi} = (\sigma_\eta^2, \sigma_{\tilde{\epsilon}}^2)$.

For the given data generation process one would expect the estimates of σ_η^2 to be at the boundary of the coefficient space at 0. Furthermore, the estimates of $\sigma_{\tilde{\epsilon}}^2$ are expected to be in the neighbourhood of the conditional variance $\text{Var}\left[Y_{t+h}|Y_t = y_t\right] = \sum_{i=0}^{h-1} \rho^{2i}$ of the data generation process. However, due to the misspecified autocorrelation for $\{\tilde{\epsilon}_{t|t-h}\}$ when $h > 1$, the maximum likelihood estimates are biased. Figure 2 presents for selected lead times the average estimate of σ_η^2 and the average estimation bias for $\sigma_{\tilde{\epsilon}}^2$ across all 1000 simulated time series as a function of the true AR coefficient ρ. The figure depicts that the estimates $\hat{\sigma}_\eta^2$ are negatively biased and $\hat{\sigma}_{\tilde{\epsilon}}^2$ are positively biased. The estimation biases increase non-linearly with the absolute value of ρ and the lead time h, with being significantly visible for $\sigma_{\tilde{\epsilon}}^2$ when $|\rho| > 0.7$. This dependency is expected from the previous expositions in Section 3.1 where it is shown that the residuals of the optimal direct h-step ahead forecasts follow a MA$(h-1)$ process with the MA coefficients being a function of ρ. Therefore, the iid observation noise assumption of model (26) is severely violated when $|\rho|$ and h are large.

We investigate the forecast accuracy of the misspecified model (26) by performing genuine rolling out-of-sample forecasting on the remaining 1000 observations of each time series while keeping the variance estimates fixed at their maximum likelihood values. The misspecified

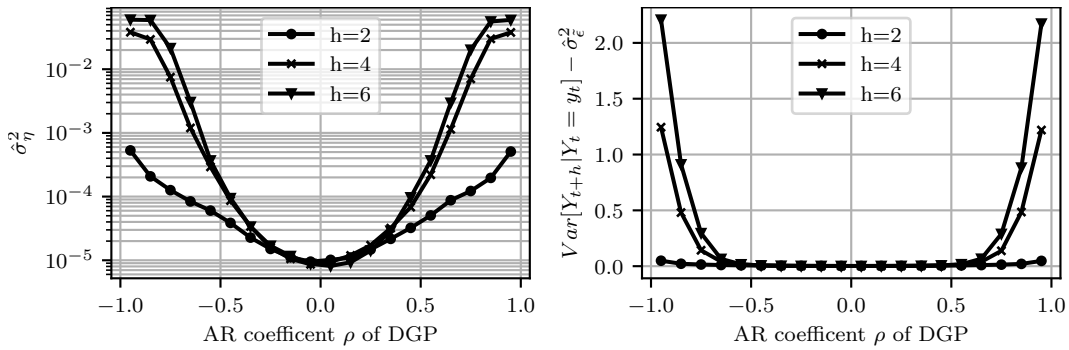

Figure 2: Average estimate $\frac{1}{|\mathcal{I}_\rho|}\sum_{i\in|\mathcal{I}_\rho|}\hat{\sigma}^2_{\eta,i}$ (left) and average bias $\frac{1}{|\mathcal{I}_\rho|}\sum_{i\in|\mathcal{I}_\rho|}\text{Var}\left[Y_{t+h}|Y_t = y_t\right]-\hat{\sigma}^2_{\tilde{\epsilon},i}$ (right) for the stationary data generation process (25) when the model (26) assumes the observation noise to be Gaussian white noise.

model is compared against the model which assumes the correct innovation process

$$\tilde{\epsilon}_{t|t-h} = \sum_{i=1}^{h-1}\phi_i\epsilon_{t-i} + \epsilon_t, \quad \epsilon_t \overset{iid}{\sim} \mathcal{N}(0,\sigma^2_\epsilon), \tag{27}$$

which follows from the data generating mechanism. The state-space model (26) remains linear and Gaussian when using the presented state augmentation approach to account for the serial correlation in $\tilde{\epsilon}_{t|t-h}$, thus we continue applying the KF for state inference. An example of the state augmentation approach for linear models is illustrated in Example 2. The unknown coefficients of the corresponding direct h-step ahead model $\boldsymbol{\psi} = (\sigma^2_\theta, \sigma^2_\zeta, \phi_1, \ldots, \phi_{h-1})$ are obtained via maximum likelihood estimation while taking $\mu_1 = \rho^h$ and $\sigma^2_1 = 0.0025$ as before.

The h-step ahead forecast accuracy of both models are compared via the skill scores

$$SS_{h,model} = \frac{1}{|\mathcal{I}_\rho|}\sum_{i\in|\mathcal{I}_\rho|} 1 - \frac{S_{(h,model),i}}{S_{(h,oracle),i}}, \tag{28}$$

where $S_{(h,oracle),i}$ is the average out-of-sample Continuous Ranked Probability Score (CRPS) (Gneiting and Katzfuss, 2014) of the $h-$step ahead oracle forecast for the i^{th} time series in the set $\mathcal{I}_\rho$. Analogously $S_{(h,model,i)}$ denotes the average out-of-sample CRPS of a proposed direct $h-$step ahead state-space models. The model $S_{(h,model,i)}$ and oracle $S_{(h,oracle),i}$ scores are both calculated as the average CRPS over the last 1000 observations of each time series. The oracle forecast is given as the conditional distribution $\mathcal{L}(Y_{t+h}|\mathcal{F}^o_t)$, where the oracle forecaster is assumed to know the data generation process. With the information set $\mathcal{F}^o_t$ including the time series observations up to time t and the true AR coefficient ρ, the h-step ahead oracle forecast is obtained as $\mathcal{L}(Y_{t+h}|\mathcal{F}^o_t) = N(\rho^h y_t, \sum_{i=0}^{h-1}\rho^{2i})$. We follow this definition of an oracle forecast to illustrate in the following that the state-space model with a correctly specified innovations process is almost optimal in the proposed accuracy measurement.

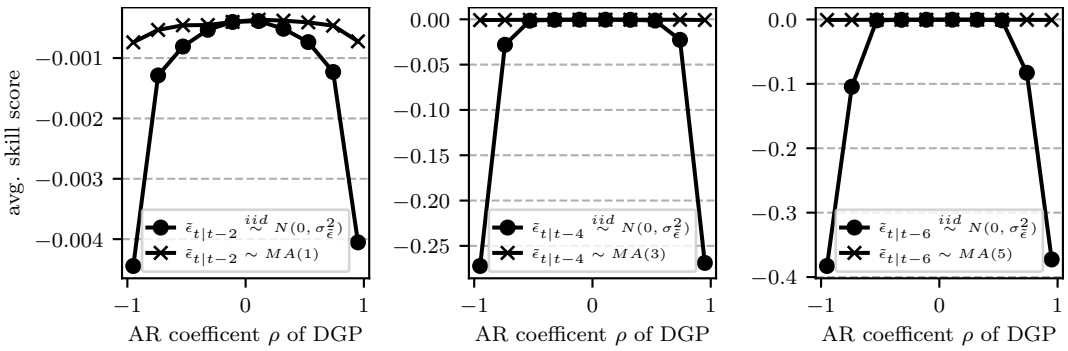

Figure 3: Averaged CRPS skill scores for the direct h-step ahead time-varying coefficient AR model (26) which takes the observation noise $\tilde{\epsilon}_{t|t-h}$ to be (i) $\tilde{\epsilon}_{t|t-h} \overset{iid}{\sim} N(0, \sigma^2_{\tilde{\epsilon}})$ and (ii) $\tilde{\epsilon}_{t|t-h} = \sum_{i=1}^{h-1} \phi_i \epsilon_{t-i} + \epsilon_t, \quad \epsilon_t \overset{iid}{\sim} N(0, \sigma^2_{\epsilon})$. The skill scores are calculated for the stationary data generation process (25) and the reference CRPS is the oracle forecast. With abuse of notation we let $\tilde{\epsilon}_{t|t-h} \sim \text{MA}(h-1)$ indicate that the observation noise follows on MA process of order h-1.

Following these definitions, Figure 3 shows the skill score $SS_{h,model}$ for the previous chosen true AR coefficient and lead time ranges. The presented skill scores verify that the biased variance estimates of the misspecified model result in poor forecast accuracy when $|\rho|$ and h are large. The direct $h-$step ahead models with the parametrized $\text{MA}(h-1)$ observation noise processes achieve scores that are comparable to the oracle forecasts. This is expected since the observation noise processes are correctly specified and therefore the maximum likelihood estimates $\hat{\sigma}^2_{\eta}$ are close to the parameter space boundary at 0 and the estimates $\hat{\sigma}^2_{\zeta}$ are close to the true value of 1.

5.2. Non-stationary data generation process

This experiment considers a different data generation process in form of the time-varying coefficient AR process

$$y_t = \rho_t y_{t-1} + \zeta_t, \quad \zeta_t \overset{iid}{\sim} \mathcal{N}(0, 1), \tag{29}$$

where $t = 1, \ldots, 2000$ and $y_0 = 0$. We take

$$\rho_t = \begin{cases} 0.9 & \text{for } t = 1, \ldots, 1000 \\ 0 & \text{for } t = 1001, \ldots, 2000 \end{cases} \tag{30}$$

to simulate a significant structural break that is expected to be captured by the proposed state-space model and generate 1000 time series replicates from this process. Letting the oracle forecaster again having access to the data generation process (29) and defining $\mathcal{F}^o_t$ as the information set to include the time series observations up to time t and the AR coefficients up to time instance $t+h$, the oracle forecast is obtained as

$$\mathcal{L}(Y_{t+h}|\mathcal{F}^o_t) = \mathcal{N}\left(\left(\prod_{i=1}^{h} \rho_{t+i}\right) y_t, 1 + \sum_{i=1}^{h-1} \left(\prod_{j=1+i}^{h} \rho^2_{t+j}\right)\right). \tag{31}$$

As it has been shown, the resulting residual of this oracle forecast follows an MA($h-1$) process, where the MA coefficients are functions of the time-varying AR coefficient ρ_t. We utilize this and demonstrate the generality of our proposed framework by considering the state-space model

$$\begin{aligned}
\begin{pmatrix} \theta_{t+1} \\ \phi_{1,t+1} \\ \vdots \\ \phi_{h-1,t+1} \end{pmatrix} &= \begin{pmatrix} \theta_t \\ \phi_{1,t} \\ \vdots \\ \phi_{h-1,t} \end{pmatrix} + \begin{pmatrix} \eta_{t+1} \\ \xi_{1,t+1} \\ \vdots \\ \xi_{h-1,t+1} \end{pmatrix}, \qquad \begin{pmatrix} \eta_{t+1} \\ \xi_{1,t+1} \\ \vdots \\ \xi_{h-1,t+1} \end{pmatrix} \overset{iid}{\sim} \mathcal{N}(0, \boldsymbol{Q}) \\
\tilde{\epsilon}_{t+1|t+1-h} &= \sum_{j=1}^{h-1} \tanh(\phi_{i,t}) \epsilon_{t-j+1} + \epsilon_{t+1}, \qquad \epsilon_{t+1} \overset{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon^2) \\
y_t &= y_{t-h} \theta_t + \tilde{\epsilon}_{t|t-h}
\end{aligned} \tag{32}$$

as an approximation to the given data generation mechanism. The latent coefficient θ_t of the model follows a random walk to allow for tracking the time-varying product term $\prod_{i=1}^{h} \rho_{t+i}$ of the oracle's forecast mean. By letting each of the model's MA coefficients $\phi_{1,t}, \ldots, \phi_{h-1,t}$ follow a random walk we achieve that the variance of $\tilde{\epsilon}_{t|t-h}$ is time-dependent, hence the model is expected to track the time-varying variance of the oracle forecast. To obtain a parsimonious model it is assumed that all MA coefficient process noises are mutually independent and share a single variance coefficient σ_ξ^2. The AR coefficient process noise η_{t+1} has variance σ_η^2 and is assumed to be independent from the remaining process noises. Therefore, the covariance matrix $\boldsymbol{Q}$ is diagonal, with $(\sigma_\eta^2, \sigma_\xi^2, \ldots, \sigma_\xi^2)$ being the h-dimensional coefficient vector of the main diagonal. The hyperbolic tangent transformation is applied to avoid model identifiability challenges due to the invertibility property of MA processes when estimating the unknown variance coefficient σ_ζ^2.

For state inference we apply the state augmentation approach, which results in a non-linear state process due to the interactions of the latent MA coefficients and innovation terms in the parameterization of $\tilde{\epsilon}_{t|t-h}$. The model formulation is completed by assuming $\boldsymbol{\alpha}_{\star,0} \sim N(\boldsymbol{0}, \boldsymbol{I})$ for the initial state vector of the augmented system. The unknown model coefficients $\boldsymbol{\psi} = (\sigma_\eta^2, \sigma_\xi^2, \sigma_\zeta^2)$ are estimated by numerically maximizing the Gaussian likelihood (24), where the mean and variance of each likelihood term are estimated with the UKF. The σ-points of the UKF are generated with Algorithm 1 while taking $\kappa = 3 - m$, with m being the length of the augmented state vector. The value is heuristically chosen under the assumption that the state vector is approximately Gaussian (Julier and Uhlmann, 1997).

The proposed state-space model shows generally a good trackability for the time-varying product term $\prod_{i=1}^{h} \rho_{t+i}$ and the oracle's forecast variance. The average estimates across all 1000 Monte Carlo replicates are shown in Figure 5.2, where a small bias in the predicted variance is observed for 4- and 6-step ahead forecasts. The bias results from the UKF only providing approximations to the filter distributions of the non-linear state-space model, which itself is only an approximation to the true data generation mechanism. Moreover, the estimated statistics are seen to lag behind the respective true values because the proposed maximum likelihood estimation jointly accounts for adaptivity and variability in

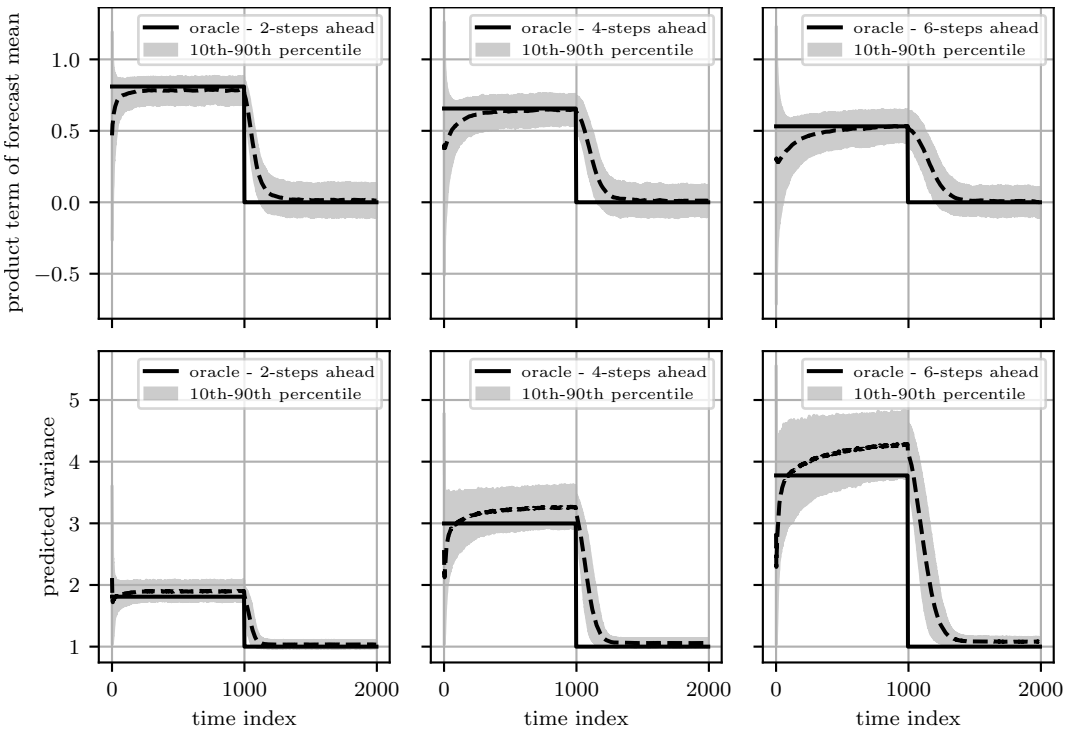

Figure 4: Averaged estimates (dashed line) across all 1000 Monte Carlo replicates of the oracle's forecast variance (bottom row) and time-varying product term $\prod_{i=1}^{h} \rho_{t+i}$ (top row) for selected lead times h.

the corresponding state estimates. This bias-variance tradeoff avoids therefore the costly consequences in terms of poor forecast accuracy due to overfitting of more adaptive, but highly volatile state estimates. The in-sample CRPS skill scores in Figure 5 confirm that the impact on the forecast accuracy of the observed estimation bias and time-delay is small.

6. Empirical analysis

In this empirical analysis we investigate the effectiveness of the proposed state-space modelling approach for direct multi-step ahead forecasting on a real-world dataset from a container shipping company. The foundation of the shipping company's business is to provide empty containers to customers, which are then being filled with freight before the container is moved to an ocean port. Once a container arrived at its destination, the freight is discharged, and the empty container is being returned to the shipping company. Providing empty containers at a low cost is difficult because global trade imbalances let empty containers accumulate in regions with negative containerized trade balances.

Empty containers are contrary scarce in regions which export more containers than they import. To satisfy demand it is therefore an essential task of the shipping company to reposition empty containers from surplus to deficit regions. The repositioning decisions must be made well in advance before empty container demand and supply mismatching occurs since it takes several weeks for containers to be transported on the major transport routes.

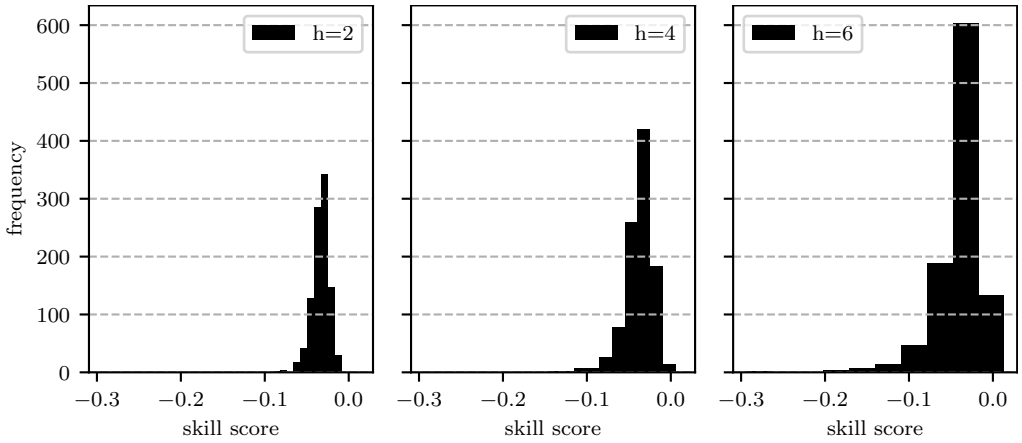

Figure 5: Empirical distribution across all 1000 Monte Carlo replicates of the in-sample CRPS skill score ($SS_{(h,model),i} = 1 - \frac{\text{CRPS}_{(h,model),i}}{\text{CRPS}_{(h,oracle),i}}, i = 1, \ldots, 1000$) for selected lead times h between the oracle forecast and estimated state-space model.

The interested reader finds a general introduction to the empty container repositioning problem in Song and Dong (2022). The associated shipping company of this study has implemented an optimization framework with the objective of finding optimal repositioning decisions. The framework requires forecasts of future weekly empty container deliveries from container storage facilities to customers, and forecasts of container returns from customers to storage facilities. The forecasts are additionally required to be probabilistic.

This forecasting problem has the property that partial information about future empty container returns is well in advance available before the actual returns materialise. Consider that it normally takes multiple weeks before a travelling full container is being returned to one of the shipping company's storage facilities. It is then possible to utilise the related booking information about the container's final destination and expected arrival time to derive an important predictor. The predictor is stochastic because the shipping company grants its customers the possibility to return the container within a loose time window after its full arrival. Customers may also request to return the container in a different container storage location than agreed upon at the time of booking, adding additional uncertainty. An important characteristic of the predictor, which we will refer to as the *booking information regressor*, is that its value depends on the lead time of the associated forecast. This is due to the circumstance that the expected empty container returns based on the available booking information decrease with the forecast lead time since bookings with returns far ahead in the future have not been observed yet at the forecast origin. This is depicted in Figure 1 for a time series from the available dataset.

6.1. Data

The associated shipping company provided two anonymized datasets (available in the supplementary material). The first dataset contains 100 time series of historical weekly empty container returns, with records dating back to week 11 of 2017. All time series are short with only 259 observations and many time series experience non-stationarities in terms

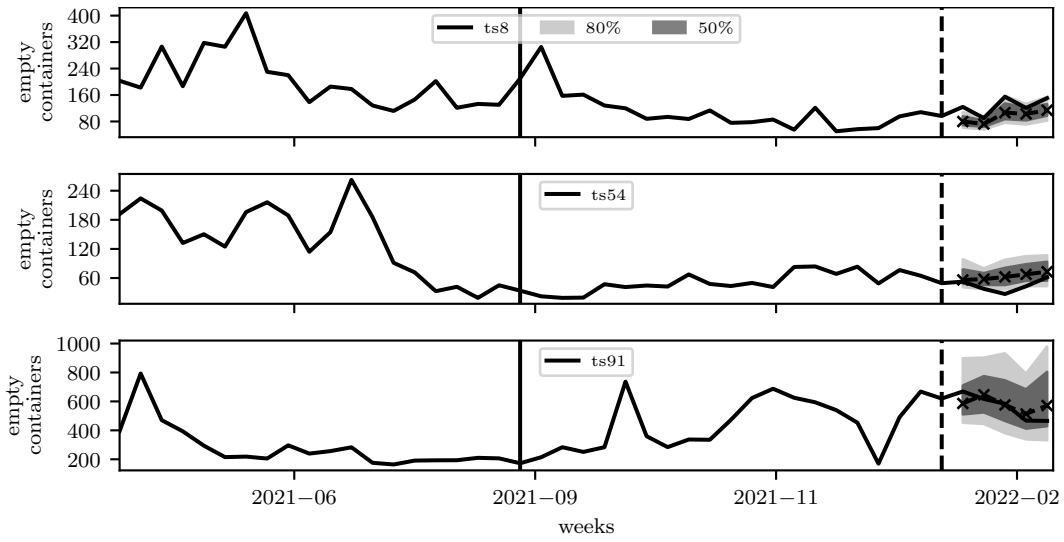

Figure 6: Truncated example time series from the available anonymised dataset. The shown predictions (origin indicated by dashed vertical line) are direct 1- to 5-step ahead predictions of tvARX-MA models, where the model parameters have been estimated on the training data (indicated by the solid line). The markers indicate the median of the predictive distributions, whose central 50% and 80% prediction intervals are overlaid.

of changing levels or variances. An example of changing levels due to the effects of the global COVID-19 pandemic on containerized trade is depicted for 3 time series in Figure 6. Similar level and volatility variations are visually observable before 2020, which may have been caused by changes in the shipping company's network or of regional economical conditions.

The second dataset contains the associated booking information regressor for all 100 time series, where each regressor is available for 1- to 5-step ahead forecasts. Since all observations and regressors are strictly greater than 0, we proceed to log transform both variables to align better with the Gaussian assumption of the following state-space models. The booking information regressor is also transformed to keep it on the scale of the response variable.

6.2. State-space model parameterization

The following base model parameterization is considered to investigate the importance of accounting for the serial correlation in the multi-step ahead forecast errors. All direct $h-$step ahead state-space models in this study are based on variations of the time-varying coefficient AR-X model

$$\begin{aligned} \begin{pmatrix} \mu_{t+1} \\ \theta_{t+1} \\ \beta_{t+1} \end{pmatrix} &= \begin{pmatrix} \mu_t \\ \theta_t \\ \beta_t \end{pmatrix} + \boldsymbol{\eta}_{t+1} \quad \boldsymbol{\eta}_{t+1} \overset{iid}{\sim} \mathcal{N}(0, \boldsymbol{Q}) \\ y_t &= \mu_t + y_{t-h}\theta_t + b_{t|t-h}\beta_t + \tilde{\epsilon}_{t|t-h} \end{aligned} \tag{33}$$

where μ_t is a local level state, θ_t and β_t are time-varying regression coefficients for the lagged response variable and the forecast horizon-dependent booking information regressor $b_{t|t-h}$, respectively. It should be noted that the index i that references to a distinct time series within the considered dataset is omitted to improve the readability.

The model aims to capture some of the observed non-stationarities by allowing the intercept, autoregressive and regression coefficient state to vary with time. The adaptivity of each state is at this controlled by the covariance matrix $\boldsymbol{Q}$, for which many design choices are available. However, since the time series in this study are short, we follow the parsimonious principle and take $\boldsymbol{Q}$ to be a diagonal matrix, where $(\sigma_\mu^2, \sigma_\theta^2, \sigma_\beta^2)$ are the coefficients on the main diagonal.

6.2.1. Serially uncorrelated innovations

The reference model in this analysis assumes $\tilde{\epsilon}_{t|t-h} \overset{iid}{\sim} N(0, \sigma_{\tilde{\epsilon}}^2)$, which ignores the serial correlation in the residuals of the multi-step ahead forecasts. State inference via the Kalman filter can be performed routinely because the state-space model is linear and Gaussian. Models with this innovations process are the most parsimonious in this analysis since only 4 unknown coefficients $\boldsymbol{\psi} = (\sigma_\mu^2, \sigma_\theta^2, \sigma_\beta^2, \sigma_{\tilde{\epsilon}}^2)$ must be estimated from data.

6.2.2. MA(h-1) innovations

This model assumes the h-step ahead innovations to follow the MA$(h-1)$ process

$$\tilde{\epsilon}_{t|t-h} = \sum_{i=1}^{h-1} \phi_i \epsilon_{t-i} + \epsilon_t, \quad \epsilon_t \overset{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon^2), \tag{34}$$

that has already been used in the simulation study of Section 5.1. Since the observation innovations are serially correlated, we employ the state augmentation technique to obtain a state-space model with an innovation free observation equation. The state-space model remains linear and the standard Kalman filter can also be applied for state inference. The vector of unknown model coefficients increases to $\boldsymbol{\psi} = (\sigma_\mu^2, \sigma_\theta^2, \sigma_\beta^2, \phi_1, \ldots, \phi_{h-1}, \sigma_\epsilon^2)$ and grows with the forecast lead time h.

6.2.3. MA(h-1) innovations with time-varying coefficients

The most flexible innovation process in our analysis extends the previous MA$(h-1)$ innovations parameterization to

$$\tilde{\epsilon}_{t|t-h} = \sum_{i=1}^{h-1} \tanh(\phi_{i,t-1}) \epsilon_{t-i} + \epsilon_t, \quad \epsilon_t \overset{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon^2), \tag{35}$$

with

$$\begin{pmatrix} \phi_{1,t+1} \\ \vdots \\ \phi_{h-1,t+1} \end{pmatrix} = \begin{pmatrix} \phi_{1,t} \\ \vdots \\ \phi_{h-1,t} \end{pmatrix} + \boldsymbol{\eta}_{\phi,t+1} \quad \boldsymbol{\eta}_{\phi,t+1} \overset{iid}{\sim} \mathcal{N}(0, \boldsymbol{Q}_{\boldsymbol{\phi}}), \tag{36}$$

where Q_ϕ is a diagonal matrix with $(\sigma^2_\phi, \ldots, \sigma^2_\phi)$ on the main diagonal, i.e. all MA coefficient states share the same innovation variance parameter. We expect this to work sufficiently well under the assumption that all MA coefficient states vary with the same rate of change, yet are uncorrelated. In addition we prefer this low-dimensional parameterization due to the limited amount of data in this study.

The hyperbolic tangent transformation is employed to avoid model identification challenges related to the invertibility property of MA processes. We further take $\boldsymbol{\eta}_{\phi,t+1}$ and the innovations $\boldsymbol{\eta}_{t+1}$ of the base parameterization (33) to be mutually independent. The state augmentation approach is employed, resulting in a non-linear state space model for which the UKF is used for state inference. The σ-points are generated with Algorithm 1 while taking $\kappa = 3 - m$, with m being the length of the augmented state vector. The value is heuristically chosen under the assumption that the state vector is approximately Gaussian (Julier and Uhlmann, 1997). Last, the vector of unknown model coefficients $\boldsymbol{\psi} = (\sigma^2_\mu, \sigma^2_\theta, \sigma^2_\beta, \sigma^2_\phi, \sigma^2_\epsilon)$ is independent of the forecast lead time given that $h > 1$.

6.3. State initialisation and model coefficient estimation

The non-stationary states of all models are initialised diffusely by taking the initial state mean to be 0 (Durbin and Koopman, 2012). For the level state we assign a variance of 20 and a smaller variance of 0.5 to the regression and MA coefficient states, with the expectation of the latter state estimates to be in the neighbourhood around 0. The remaining states of the augmented system for models with the time-invariant coefficient MA innovations process are initialised with their corresponding stationary distribution (Gardner et al., 1980). A larger initial variance of 20 is assigned to $\tilde{\epsilon}_{t|t-h}$ for the model with time-varying MA coefficients, reflecting that the state is non-stationary due to the dependency on the non-stationary MA coefficient states. The remaining serially uncorrelated innovation states of the augmented system are initialised with their stationary distribution.

A train-test split is performed to assess the genuine out-of-sample forecast performance of the considered models. The first 233 observations, approximately 90% of the total time series length, are used for parameter estimation. The remaining 26 observations are used for forecast evaluation with the model parameters fixed to the estimates that have been obtained on the training set. Even though we perform batch learning on the variance and MA coefficient estimates, the state estimates are allowed naturally to vary on the hold-out set. The model parameters are estimated by numerically maximizing the negative Gaussian log likelihood (24) as proposed in Section 4.1.2. The contribution of the first 15 log likelihood terms is set to 0 to mitigate the effect of the state estimation burn-in period on the model parameter estimates that results from the chosen initialisation procedure. To reduce convergence issues we propose to start the numerical maximization procedure of each model log likelihood in 10 different starting values. Each parameter starting value set is chosen randomly within a fixed interval in which the resulting coefficient estimates are expected to lie. The model fit with the largest log likelihood within each subgroup is then considered in the following analysis.

Table 1: Declaration of the state-space models that are being used in the empirical analysis. *tv* indicates that the corresponding regression and moving average coefficients are time-varying. All models assume $\epsilon_t \overset{iid}{\sim} N(0, \sigma_\epsilon^2)$ and h denotes the forecast lead time.

	system model $g_h(\cdot)$	innovations process $\tilde{\epsilon}_{t\|t-h}$	filter type
tvAR-iid	$\mu_t + y_{t-h}\theta_t$	ϵ_t	KF
tvAR-MA	$\mu_t + y_{t-h}\theta_t$	$\sum_{i=1}^{h-1} \phi_i \epsilon_{t-i} + \epsilon_t$	KF
tvAR-tvMA	$\mu_t + y_{t-h}\theta_t$	$\sum_{i=1}^{h-1} \tanh(\phi_{i,t-1})\epsilon_{t-i} + \epsilon_t$	UKF
tvARX-iid	$\mu_t + y_{t-h}\theta_t + b_{t\|t-h}\beta_t$	ϵ_t	KF
tvARX-MA	$\mu_t + y_{t-h}\theta_t + b_{t\|t-h}\beta_t$	$\sum_{i=1}^{h-1} \phi_i \epsilon_{t-i} + \epsilon_t$	KF
tvARX-tvMA	$\mu_t + y_{t-h}\theta_t + b_{t\|t-h}\beta_t$	$\sum_{i=1}^{h-1} \tanh(\phi_{i,t-1})\epsilon_{t-i} + \epsilon_t$	UKF

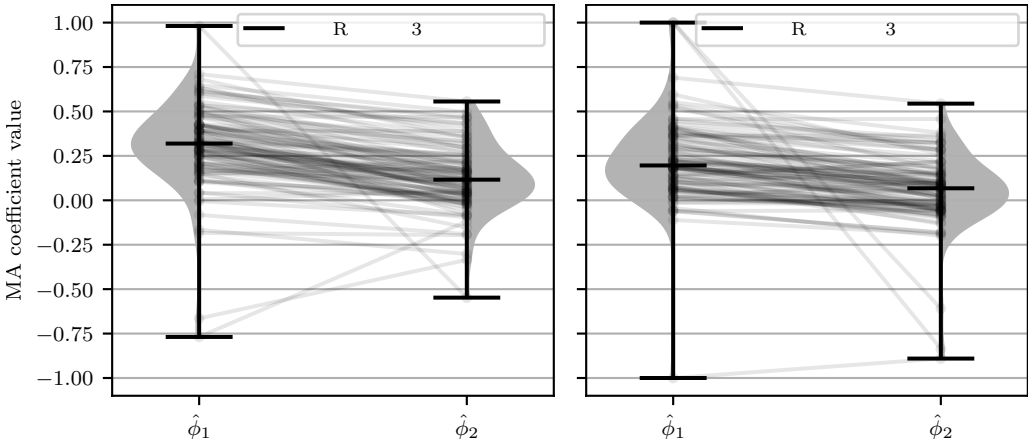

Figure 7: Estimates of MA coefficients for 3-step ahead models, tvAR-MA (left) and tvARX-MA (right), for all 100 time series. The central horizontal black line indicates the median of all 100 coefficient estimates.

6.4. In-sample results

We first analyse the in-sample properties of the proposed modelling framework before an investigation of their true out-of-sample predictive performance is presented. Throughout the following two sections we differentiate between AR and AR-X models to show that the inclusion of the booking information regressor reduces the negative effects of the misspecified innovations process. Table 1 provides a summary of all the models that are being considered in this study. We further restrict ourselves to 3- and 5-step ahead forecasts and like to remind the reader that the order of all MA innovation processes is set to $h-1$, where h denotes the forecast lead time.

6.4.1. MA coefficient estimates

To begin with, it can be reported that the vast majority of the MA coefficient estimates of the tvAR-MA and tvARX-MA models are positive and greater than 0. Figure 7 shows this for the 3-step ahead forecast models. The same observations can be reported for the 5-step ahead forecast models (not shown here). The statistical significance of the non-zero MA coefficient estimates can be estimated with a likelihood-ratio test since the tvAR(X)-

MA and tvAR(X)-iid models are nested. The null hypothesis of the restricted tvAR-iid model fitting the observed data better than the tvAR-MA model is rejected for 66 (3-steps ahead) and 57 (5-steps ahead) times series at the significance level $\alpha = 0.05$. Less frequent null hypothesis rejections of 56 (3-steps ahead) and 41 (5-steps ahead) can be reported for the tvARX-iid model. These results indicate the effectiveness of the considered models with MA innovations processes to account for the serial correlation in the h-step ahead residuals. The effectiveness is further validated below by evaluating the in-sample forecast accuracy.

6.4.2. In-sample forecast accuracy

The probabilistic forecasts are obtained in this study by Monte Carlo sampling as introduced in Section 4.2, where each forecast is made of 200 samples. The first 15 forecasts are ignored when calculating accuracy metrics to remove the effects of the state estimation burn-in period. Furthermore, we take the exponential of all Monte Carlo samples to calculate the forecast metrics on the response variable's original scale. To compare the effectiveness of our approach we calculate the scaled skill scores

$$SS_{g_h,inno} = 100\frac{S_{g_h,iid} - S_{g_h,inno}}{S_{g_h,iid}}, \tag{37}$$

where $g_h \in \{\text{tvAR}, \text{tvARX}\}$ and $inno \in \{\text{MA}, \text{tvMA}\}$ are the system model and innovation process parameterizations as defined in Table 1, respectively. With S being a forecast accuracy summary measure for a distinct time series $i = 1, \dots 100$ and look ahead time $h \in \{3, 5\}$, such as the average CRPS or MSE, it follows that the skill score assesses the difference in forecast accuracy between models with the same system model g_h, but with different innovations processes. For the negative oriented CRPS and MSE it follows from the definition in Equation (37) that a positive skill scores indicates superior forecasts of the models with serially correlated innovations.

Table 2 summarizes the skill scores on the training set when using the CRPS to evaluate the probabilistic forecasts of the considered state-space models. The CRPS scores for each lead time, time series and model are obtained as the average

$$\text{CRPS}_{g_h,inno} = \frac{1}{|\mathcal{T}_{\text{train}}|}\sum_{t\in\mathcal{T}_{\text{train}}} \text{CRPS}(y_{t+h}, \hat{F}_{t+h|t,(g_h,inno)}), \tag{38}$$

where $\mathcal{T}_{\text{train}}$ is the set of training observations indices for which the h-step ahead forecasts $\hat{F}_{t+h|t,(g_h,inno)}$ from the indicated model are available. The presented percentiles show that the greatest improvement over the corresponding model with *iid* innovations is achieved by the tvAR-MA model. For 50 time series the in-sample CRPS is reduced by more than 1.8%, with reductions greater than 8.6% for 10 time series. The table shows further that the tvAR-tvMA model performs marginally worse than the tvAR-MA model. This may be attributed to the combination of model misspecification and the deficiencies in the proposed state estimation procedure with the UKF, which has already been discussed in the simulation study of Section 5.2.

The improvements over the iid model parameterization of the tvARX models are generally found to be smaller. The smaller improvements are expected on the basis that including

Table 2: Percentiles of the scaled training CRPS skill scores across all 100 time series of the models with MA innovations processes. The skill scores, as defined in Equation (37), are calculated with respect to the corresponding models with iid innovations.

	3-steps ahead					5-steps ahead				
	10th	25th	50th	75th	90th	10th	25th	50th	75th	90th
tvAR-MA	0.1	0.5	1.8	4.4	8.6	-0.4	0.8	2.5	5.8	11.9
tvAR-tvMA	-2.2	-0.7	0.5	3.1	7.4	-5.2	-1.3	0.9	4.1	8.3
tvARX-MA	-0.4	-0.1	0.6	2.3	7.8	-1.0	-0.3	0.9	2.6	9.5
tvARX-tvMA	-4.6	-2.2	-0.3	1.5	7.8	-7.9	-3.4	-0.9	1.2	4.9

Table 3: Percentiles of the scaled training MSE skill scores across all 100 time series of the models with MA innovations processes. The skill scores, as defined in Equation (37), are calculated with respect to the corresponding models with iid innovations.

	3-steps ahead					5-steps ahead				
	10th	25th	50th	75th	90th	10th	25th	50th	75th	90th
tvAR-MA	0.2	1.3	4.4	10.3	18.7	-1.0	1.7	6.7	15.3	26.8
tvAR-tvMA	-4.1	-1.8	1.5	8.6	16.3	-13.0	-3.5	2.0	10.3	18.7
tvARX-MA	-1.0	-0.1	1.5	5.6	14.9	-2.2	-0.4	1.0	5.4	19.7
tvARX-tvMA	-9.3	-3.9	-0.3	3.8	13.8	-17.3	-7.1	-0.9	3.6	9.7

the booking information regressor accounts for additional serial correlation which the tvAR models do not capture, eventually causing the serial correlation in the h-step ahead residuals to be comparatively smaller for the tvARX models. The parameter estimation bias of the iid model parameterizations, and hence the forecast accuracy deterioration, are consequently also smaller since the residual process becomes better approximated by an iid process.

Moreover, the trend of the 5-step ahead skill scores to be slightly greater can be reported. This is in line with the results of the simulation study in Section 5 where it has been observed that the improvements over the iid model are a function of the forecast lead time and the autocorrelation function of the data generation mechanism. Last, the same observations, but generally with greater improvements, can be observed for the MSE skill scores in Table 3.

6.5. Forecasting results

We investigate the out-of-sample forecast accuracy with the same skill score as defined in Equation (37), but where we take the forecast accuracy summary metric to be the average CRPS

$$\text{CRPS}_{g_h,inno} = \frac{1}{|\mathcal{T}_{\text{test}}|} \sum_{t \in \mathcal{T}_{\text{test}}} \text{CRPS}(y_{t+h}, \hat{F}_{t+h|t,(g_h,inno)}), \tag{39}$$

where $\mathcal{T}_{\text{test}}$ are the observation indices of the test set for which observation and forecast pairs are available. Table 4 and Table 5 provide the summaries over the skill score distributions for the CRPS and MSE, respectively. The tables show generally similar results with respect to the previous in-sample scores. That is, the tvAR and tvARX models with time-invariant MA innovations processes improve on average the most over their corresponding iid model

Table 4: Percentiles of the scaled test CRPS skill scores across all 100 time series of the models with MA innovations processes. The skill scores, as defined in Equation (37), are calculated with respect to the corresponding models with iid innovations.

	3-steps ahead					5-steps ahead				
	10th	25th	50th	75th	90th	10th	25th	50th	75th	90th
tvAR-MA	-2.3	-0.5	1.8	5.9	10.6	-4.0	-0.2	2.3	7.6	15.1
tvAR-tvMA	-6.8	-2.5	0.6	4.5	9.6	-8.6	-4.7	0.6	5.3	11.4
tvARX-MA	-4.1	-0.3	0.6	2.8	6.1	-4.8	-1.7	0.4	2.8	7.3
tvARX-tvMA	-8.4	-4.7	-1.0	1.9	5.2	-8.5	-5.5	-2.7	0.9	6.3

Table 5: Percentiles of the scaled test MSE skill scores across all 100 time series of the models with MA innovations processes. The skill scores, as defined in Equation (37), are calculated with respect to the corresponding models with iid innovations.

	3-steps ahead					5-steps ahead				
	10th	25th	50th	75th	90th	10th	25th	50th	75th	90th
tvAR-MA	-6.8	-0.4	5.4	12.5	20.2	-7.5	-0.8	4.6	15.8	28.2
tvAR-tvMA	-13.4	-5.3	0.7	8.8	19.2	-22.6	-9.1	-0.6	11.3	25.0
tvARX-MA	-7.5	-0.3	2.0	6.3	16.3	-9.1	-3.1	1.8	6.3	21.9
tvARX-tvMA	-17.0	-8.4	0.2	4.4	14.2	-23.9	-10.4	-4.2	2.1	11.7

parameterizations. The observed improvements are in the order that is expected from the simulation study in Section 5.1 where large improvements over the iid models are dependent on the sample autocorrelation and degree of model misspecification. It is further interesting to point out that, compared to the training skill scores, the interpercentile-range increases virtually for all models and look-ahead times. Large negative skill scores in particular are in this study observed for time series with significant level changes that occur after training. For such time series, as illustrated in Figure 8 for the 3-step ahead predictions of the tvAR models on *ts54*, a positive estimation bias in $\hat{\sigma}_\mu^2$ is beneficial because it allows the tvAR-iid model to adapt better to the changed level. The faster and better adaptivity over the tvAR-MA model is seen as the steeper decline in $\hat{\mu}_{t|t}$ that starts around time step 220.

An example where the positive estimation bias in $\hat{\sigma}_\mu^2$ of the tvAR-iid model results in poor forecast accuracy is shown in Figure 9 for the 3-step ahead predictions on *ts55*. The forecasts for the last 8 observations are poor because the local level estimates $\hat{\mu}_{t|t}$ of the tvAR-iid model are too adaptive, causing the estimates to be dragged down with the decrease in the dependent variable around time step 243.

The reliability diagrams are depicted for the tvAR models and 3-step ahead forecasts in Figure 10. Overall, the probabilistic forecasts from the various models are well-calibrated for the majority of time series. Almost identical reliability diagrams are obtained for the 5-step ahead forecasts (not shown here), albeit slightly more time series are found above or below the shown $5\% - 95\%$ consistency bars. The same general statements also apply to the reliability of the tvARX models.

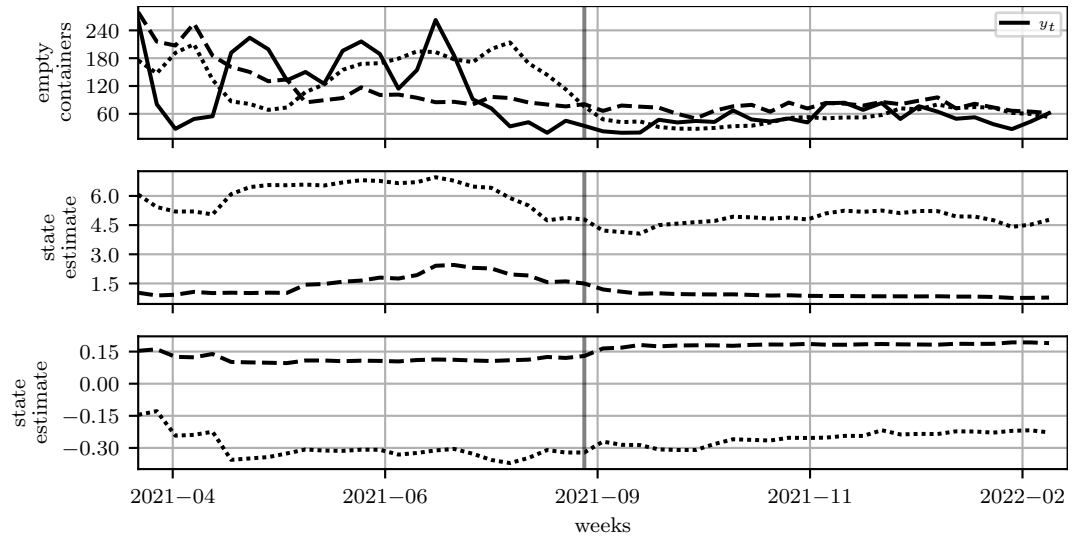

Figure 8: 3-step ahead point predictions (top), local level estimates $\hat{\mu}_{t|t}$ (middle) and autoregressive coefficient $\hat{\theta}_{t|t}$ (bottom) estimates of the tvAR-MA (dashed line) and tvAR-iid (dotted line) for time series *ts54*. The vertical line separates the training and testing sets. The top plot shows that the point predictions of the tvAR-iid model adapt better to the level change of the time series, ultimately producing a lower test set CRPS and MSE than the tvAR-MA model.

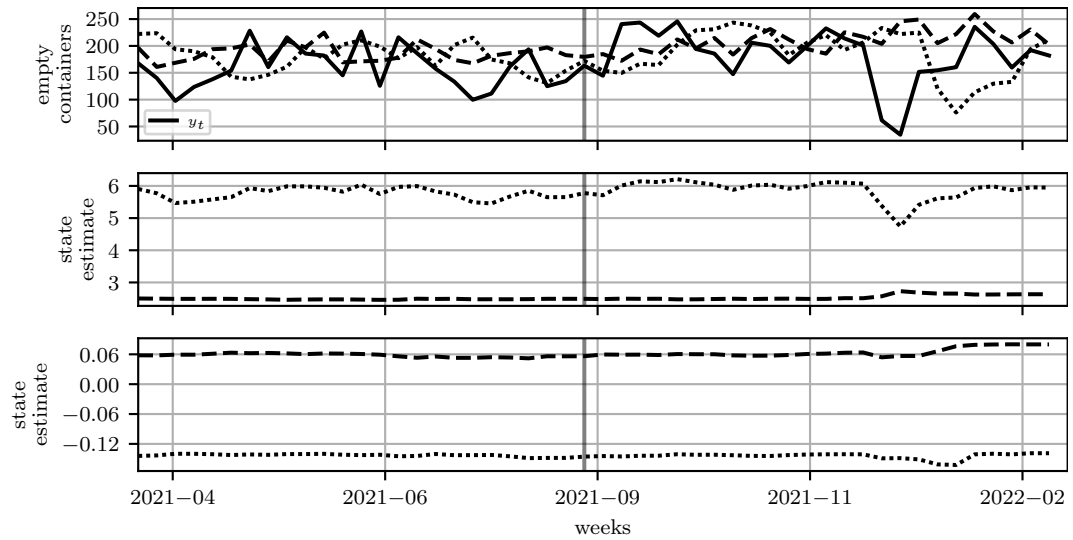

Figure 9: 3-step ahead point predictions (top), local level estimates $\hat{\mu}_{t|t}$ (middle) and autoregressive coefficient $\hat{\theta}_{t|t}$ (bottom) estimates of the tvAR-MA (dashed line) and tvAR-iid (dotted line) for time series *ts55*. The vertical line separates the training and testing sets. The top plot shows that the point predictions of the tvAR-iid model follow the lagged observations too closely, ultimately producing a higher test set CRPS and MSE than the tvAR-MA model.

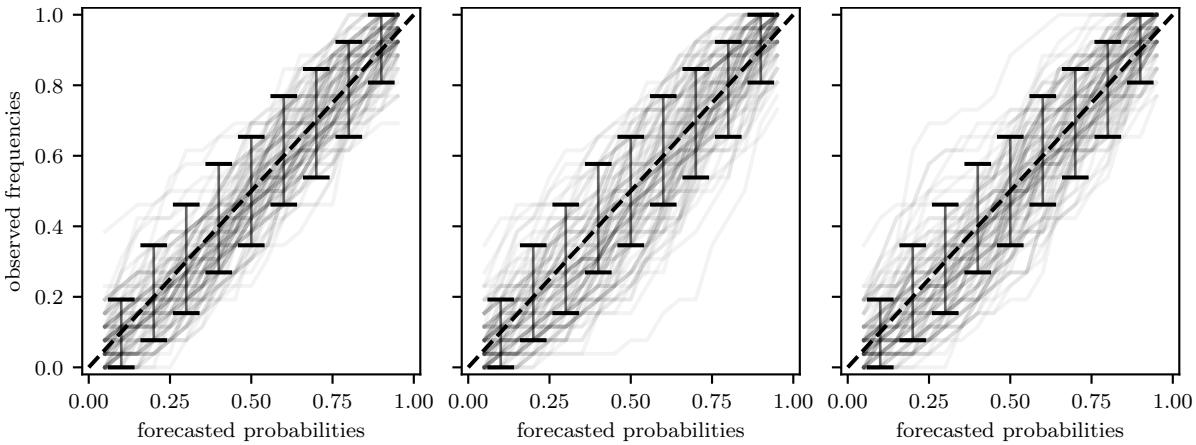

Figure 10: Out-of-sample reliability diagram for all 100 time series and 3-step ahead forecasts. The dashed diagonal line indicates a perfectly calibrated forecast and the vertical bars are the $5\% - 95\%$ consistency bars for selected nominal levels derived using the bootstrap method of Bröcker and Smith (2007) (with a number of replicates $N_{\text{boot}} = 5000$).

Table 6: Percentiles of scaled test CRPS skill score ($100\frac{\text{CRPS}_{tvAR,inno} - \text{CRPS}_{tvARX,inno}}{\text{CRPS}_{tvAR,inno}}$) over all 100 time series.

	3-steps ahead					5-steps ahead				
	10th	25th	50th	75th	90th	10th	25th	50th	75th	90th
iid	-8.2	11.4	24.6	41.3	55.4	-5.4	8.3	21.6	35.0	49.9
MA	-6.8	7.5	22.6	40.4	50.7	-13.3	3.9	17.7	34.5	47.3
tvMA	-8.5	5.1	25.0	39.8	52.7	-14.4	3.6	19.6	32.2	47.8

6.5.1. Improvements of the ARX models

Up until this point it has only been claimed that the booking information regressor is highly predictive. We validate this subsequently by calculating the improvements of the tvARX over the tvAR models for all 3 innovations processes $inno \in \{\text{iid}, \text{MA}, \text{tvMA}\}$ with

$$SS_{inno} = 100\frac{S_{tvAR,inno} - S_{tvARX,inno}}{S_{tvAR,inno}}, \tag{40}$$

where $S_{tvAR,inno}$ is a forecast accuracy summary measure on the test set as before. The CRPS skill score percentiles for the 100 considered time series and two forecast lead times are presented in Table 6, where it can be seen that the improvements are significant, and consistent, for all three innovation processes. Moreover, it can be seen that the improvements are smaller for the 5-step ahead models. This is expected from the previous introduction of the booking information regressor where it has been emphasized that the available number bookings with future empty container return dates decrease with the forecast lead time.

7. Conclusions

In this paper, we have shown that care must be taken when performing direct multi-step ahead forecasting with time-varying coefficient models in state-space form. We showed empirically that there exists an estimation bias for state-space models which incorrectly assume the innovations of a h-step ahead model to be serially uncorrelated. The severity of the bias was identified to depend on the autocorrelation function of the data generating mechanism and the lead time of the forecasting model. To enable direct multi-step ahead forecasting with state-space models we proposed to explicitly model the latent serially correlated innovation process of the h-step ahead forecasts as a MA($h-1$) process. Subsequently, we introduced two alternative time-variant and time-invariant coefficient MA process parameterizations along with their required state and model coefficient estimation procedures.

An empirical analysis of a real-world dataset of 100 time series confirmed the effectiveness of our modelling and estimation framework. The analysis showed that the state-space models with the MA innovation processes produce generally more accurate direct multi-step ahead probabilistic forecasts than the corresponding models with serially independent innovations. However, we also observed that the non-linear state-space models with the time-varying coefficient MA processes were inferior to the linear models with time-invariant coefficient MA processes. We expect this to be due to the introduced state approximation error of the Unscented Kalman filter, which ultimately propagates through to the proposed model coefficient estimation procedure based on a surrogate Gaussian likelihood.

Future work should thus address the state approximation error for non-linear state-space models by investigating other filtering methods. Of additional practical interest are in addition the identification of model selection strategies. While it is appealing to use in-sample likelihood-based criteria, it is not apparent how to perform model selection between different model classes, such as a potentially an ETS model and our proposed state-space models. This further extends to the selection of a multi-step ahead forecasting strategy when expert knowledge does not provide strong evidence for the prevalence of one of the candidate strategies. Another interesting research direction is to investigate how to exploit and share dependencies between related time series, similar to global forecasting methods.

Acknowledgements

The corresponding author is supported by an Innovation Fund Denmark grant, Case no. 9065-00021B.

References

Atiya, A., El-Shoura, S., Shaheen, S., El-Sherif, M., 1999. A comparison between neural-network forecasting techniques-case study: river flow forecasting. IEEE Transactions on Neural Networks 10, 402–409.

Bacher, P., Madsen, H., Nielsen, H.A., 2009. Online short-term solar power forecasting. Solar energy 83, 1772–1783.

Bartezzaghi, E., Verganti, R., Zotteri, G., 1999. A simulation framework for forecasting uncertain lumpy demand. International Journal of Production Economics 59, 499–510.

Ben Taieb, S., Bontempi, G., Atiya, A.F., Sorjamaa, A., 2012. A review and comparison of strategies for multi-step ahead time series forecasting based on the NN5 forecasting competition. Expert Systems with Applications 39, 7067–7083.

Bryson, A., Johansen, D., 1965. Linear filtering for time-varying systems using measurements containing colored noise. IEEE Transactions on Automatic Control 10, 4–10.

Bröcker, J., Smith, L.A., 2007. Increasing the reliability of reliability diagrams. Weather and Forecasting 22, 651–661.

Cai, Z., Tiwari, R.C., 2000. Application of a local linear autoregressive model to BOD time series. Environmetrics: The Official Journal of the International Environmetrics Society 11, 341–350.

Chen, R., Tsay, R.S., 1993. Functional-coefficient autoregressive models. Journal of the American Statistical Association 88, 298–308.

Chevillon, G., 2007. Direct multi-step estimation and forecasting. Journal of Economic Surveys 21, 746–785.

Dahlhaus, R., 2012. Locally Stationary Processes. Elservier, USA, NY, New York. volume 30 of *Handbook of Statistics*. pp. 351–414.

Dangl, T., Halling, M., 2012. Predictive regressions with time-varying coefficients. Journal of Financial Economics 106, 157–181.

Davis, R.A., Fokianos, K., Holan, S.H., Joe, H., Livsey, J., Lund, R., Pipiras, V., Ravishanker, N., 2021. Count time series: A methodological review. Journal of the American Statistical Association 116, 1533–1547.

Doucet, A., Freitas, N., Gordon, N. (Eds.), 2001. Sequential Monte Carlo Methods in Practice. Springer New York, New York, NY.

Durbin, J., Koopman, S.J., 2012. Time Series Analysis by State Space Methods. Oxford statistical science series. 2 ed., Oxford Univ. Press.

Gardner, G., Harvey, A.C., Phillips, G.D.A., 1980. Algorithm AS 154: An algorithm for exact maximum likelihood estimation of autoregressive-moving average models by means of Kalman filtering. Applied Statistics 29, 311–322.

Gneiting, T., Katzfuss, M., 2014. Probabilistic forecasting. Annual Review of Statistics and Its Application 1, 125–151.

Grenier, Y., 1983. Time-dependent ARMA modeling of nonstationary signals. IEEE Transactions on Acoustics, Speech, and Signal Processing 31, 899–911.

Harvey, A.C., 1990. Forecasting, Structural Time Series Models and the Kalman Filter. Cambridge University Press.

Harvey, D., Leybourne, S., Newbold, P., 1997. Testing the equality of prediction mean squared errors. International Journal of Forecasting 13, 281–291.

Hyndman, R.J., Koehler, A.B., Ord, J.K.O., Snyder, R.D., 2008. Forecasting with Exponential Smoothing: The State Space Approach. Springer Series in Statistics, Springer.

Januschowski, T., Gasthaus, J., Wang, Y., Salinas, D., Flunkert, V., Bohlke-Schneider, M., Callot, L., 2020. Criteria for classifying forecasting methods. International Journal of Forecasting 36, 167–177.

Julier, S., Uhlmann, J., 1997. New extension of the Kalman filter to nonlinear systems. Proc. SPIE 3068, Signal Processing, Sensor Fusion, and Target Recognition VI , 182–193.

Julier, S., Uhlmann, J., 2004. Unscented filtering and nonlinear estimation. Proceedings of the IEEE 92, 401–422.

Kalman, R.E., 1960. A new approach to linear filtering and prediction problems. Transactions of the ASME–Journal of Basic Engineering 82, 35–45.

Karakatsani, N.V., Bunn, D.W., 2008. Forecasting electricity prices: The impact of fundamentals and time-varying coefficients. International Journal of Forecasting 24, 764–785.

Li, F., Zhou, J., Wu, D., 2013. Optimal filtering for systems with finite-step autocorrelated noises and multiple packet dropouts. Aerospace Science and Technology 24, 255–263.

Lim, B., Arık, S.O., Loeff, N., Pfister, T., 2021. Temporal fusion transformers for interpretable multi-horizon time series forecasting. International Journal of Forecasting 37, 1748–1764.

Ma, S., Fildes, R., 2020. Forecasting third-party mobile payments with implications for customer flow

prediction. International Journal of Forecasting 36, 739–760.
Messner, J.W., Pinson, P., 2019. Online adaptive lasso estimation in vector autoregressive models for high dimensional wind power forecasting. International Journal of Forecasting 35, 1485–1498.
Moulines, E., Priouret, P., Roueff, F., 2005. On recursive estimation for time varying autoregressive processes. The Annals of Statistics 33, 2610–2654.
Poncela, M., Poncela, P., Perán, J.R., 2013. Automatic tuning of Kalman filters by maximum likelihood methods for wind energy forecasting. Applied Energy 108, 349–362.
Priestley, M., 1980. State-dependent models: A general approach to non-linear time series analysis. Journal of Time Series Analysis 1, 47–71.
Rangapuram, S.S., Seeger, M.W., Gasthaus, J., Stella, L., Wang, Y., Januschowski, T., 2018. Deep state space models for time series forecasting, in: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (Eds.), Advances in Neural Information Processing Systems, Curran Associates, Inc.. pp. 1–13.
Salinas, D., Flunkert, V., Gasthaus, J., Januschowski, T., 2020. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. International Journal of Forecasting 36, 1181–1191.
Sanchez, I., 2006. Short-term prediction of wind energy production. International Journal of Forecasting 22, 43–56.
Seeger, M.W., Salinas, D., Flunkert, V., 2016. Bayesian intermittent demand forecasting for large inventories, in: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (Eds.), Advances in Neural Information Processing Systems, Curran Associates, Inc.. pp. 1–9.
Song, D.P., Dong, J., 2022. Modelling Empty Container Repositioning Logistics. Springer International Publishing.
Song, H., Li, G., Witt, S.F., Athanasopoulos, G., 2011. Forecasting tourist arrivals using time-varying parameter structural time series models. International Journal of Forecasting 27, 855–869.
Sun, S., Tian, T., Lin, H., 2016. Optimal linear estimators for systems with finite-step correlated noises and packet dropout compensations. IEEE Transactions on Signal Processing 64, 5672–5681.
Svensson, A., Schön, T.B., 2017. A flexible state–space model for learning nonlinear dynamical systems. Automatica 80, 189–199.
Taieb, S.B., Atiya, A.F., 2015. A bias and variance analysis for multistep-ahead time series forecasting. IEEE Transactions on Neural Networks and Learning Systems 27, 62–76.
Wan, E., Van Der Merwe, R., 2000. The unscented kalman filter for nonlinear estimation, in: Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373), IEEE, Lake Louise, Alta., Canada. pp. 153–158.
West, M., Harrison, P.J., Migon, H.S., 1985. Dynamic generalized linear models and bayesian forecasting. Journal of the American Statistical Association 80, 73–83.

Paper C

A critical look at deep state space models for time series forecasting

Publication details: Submitted to *Conference on Neural Information Processing Systems (NeurIPS)*.

A critical look at deep state space models for time series forecasting

Benedikt Sommer[1,2], **Klaus Kähler Holst**[2], **Pierre Pinson**[3]
[1]Technical University of Denmark, Kongens Lyngby, Denmark
[2]A. P. Møller-Mærsk, Copenhagen, Denmark
[3]Imperial College London, London, United Kingdom
benedikt.sommer@maersk.com

Abstract

We explore the limitations of Deep State Space Models (DSSM), a hybrid learning and forecasting approach that fuses parametric state-space models with recurrent neural networks to forecast panels of related time series. DSSM uses a *global* (jointly-learned) neural network to learn associations between related time series whereas *locally* (per-time-series) learned state space models allow to incorporate structural assumptions about the panel process, hence providing a guard against overfitting. We revisit the derivation of the method and identify several design choices that cause distinct challenges in real-world forecasting problems. Our claims are supported in two parts. First, experiments based on Monte Carlo simulations highlight the identified limitations. Second, we use the same public data sets from the original DSSM publication to demonstrate that the method does only produce minor forecast accuracy improvements over parametric state space models. Finally, we extend our findings to the broader class of hybrid models with similar design principles.

1 Introduction

Motivated by decision problems within retail (multiple products), load management (multiple customers) and traffic management (multiple sections), recent research has been dedicated to forecasting large panels of time series simultaneously. Methods that forecast a set of time series are frequently classified as *local* or *global* to distinguish between two extreme cases (Montero-Manso and Hyndman, 2021). Local methods estimate the unknown parameters of a forecasting model for each time series in isolation, whereas the parameters of a global model are estimated jointly from all time series in the panel. Flexible non-parametric machine learning (ML) models are usually employed as global methods in larger data regimes to learn associations between features and time series observations fully data-driven. Parametric models are conceptually more suitable for panels with shorter time series when structural assumptions are necessary to estimate models that generalize well. Several hybrid methods have emerged in-between both extremes to combine "the best of both worlds", namely the ability of global models to learn jointly from all time series and the robustness of parametric models due to the incorporation of structural assumptions about the unknown data generation process.

A popular forecasting method within this paradigm is Deep State Space Models (DSSM) (Rangapuram et al., 2018), a hybrid method that fuses linear Gaussian state space models (LGSSMs) with recurrent neural networks (RNNs). The ambition of this paper is to give an independent first-hand view on the limitations of this hybrid method. Contrary to the original paper, we explore in greater detail the underlying assumptions of this approach and identify them as a limitation for many real-world forecasting problems. Our experimental studies, using datasets that have been used extensively as benchmarks in other deep learning publications, demonstrate only marginal forecast

Preprint. Under review.

accuracy improvements of DSSM over parametric state space models whose parameters are estimated independently for each time series through maximum likelihood estimation.

The remainder of the paper is organized as follows. Related forecasting methods are discussed in Section 2. Notation and models are introduced in Section 3. Section 4 elaborates on the shortcomings of DSSM. We support our claims with quantitative and qualitative experiments in Section 5. The paper is concluded in Section 6.

2 Related work

State space models offer a flexible framework to estimate a broad class of time series models, such as exponential smoothing (Hyndman, 2008), ARIMA and structural component models (Durbin and Koopman, 2012). The principal design of DSSM is based on pooling univariate LGSSMs. If the number of time series observations is small, then it is common to pool all time series and estimate a common model (Fröhwirth-Schnatter and Kaufmann, 2008). Indeed, this is a global forecasting model, however, it is parsimonious to prevent over-fitting on short time series. For heterogeneous panels, where data generating processes differ substantially between series, this will induce a bias due to the insufficient flexibility of the common model. DSSM overcomes the stringent assumption of homogeneous data generating processes with identical parameter values. A global RNN learns the mapping from a feature to parameter space, while maintaining parametric state space models for each individual time series to alleviate the risk of over-fitting. By training the global model end to end for all time series the authors aim to share statistical strength for parameter estimates between related time series.

For a review of forecasting groups of time series in panels we refer the reader to (Duncan and Szczypula, 2001). In-between group heterogeneity is discussed in Bandara et al. (2020) and dealt with by first assigning time series to groups, either by expert judgement or time series clustering, and then estimating a pooled model for each group. Unobserved heterogeneity within each group can be additionally modeled with random effects (Coakley et al., 2006). Contrary methods utilize model-based clustering to jointly estimate group assignments and model parameters in each of the groups (Fröhwirth-Schnatter and Kaufmann, 2008; Duncan and Szczypula, 2001). Exogenous features, such as geography or item category, can be incorporated to estimate group assignments. DSSM has strong similarities with model-based clustering, however, the method never estimates an explicit grouping. Identical to the work of Garcia-Ferrer et al. (1987); Baltagi (2005), DSSM assumes that all time series in the panel can be modelled with a single parametric model. Related time series are grouped implicitly because an RNN predicts the parameter values of each parametric forecasting model from a set of time series dependent features. This approach follows the rational that statistical strength can be borrowed between related time series, where relatedness is estimated non-parametrically with a neural network. Hence, DSSM aims to preserve the benefits of parametric forecasting models for short time series while utilizing a flexible global model to find associations between time series to reduce parameter estimation errors.

Different designs for blending parametric and non-parametric ML forecasting models have been proposed in the past. Bottom-up approaches are followed in Smyl (2020); Bandara et al. (2021) by first estimating a parametric model for each local time series, followed by training a global RNN on the residuals to account for remaining non-linearities that are potentially shared among related time series. Both contributions extend the hybrid model of Zhang (2003) where a feedforward network is trained on the residuals of an ARIMA fit. However, the forecasting method in Zhang (2003) is local, hence over-fitting is likely for short time series. DSSM follows a top-down approach by parameterizing a LGSSM for each time series in the panel with a global RNN. Benidis et al. (2022) determines two major limitations of DSSM, namely that of a Gaussian noise assumption and linear state dynamics. The Normalizing Kalman Filters of de Bézenac et al. (2020) is an immediate extension of DSSM to multivariate and non-Gaussian time series. The method utilizes normalizing flows for tractable inference of non-Gaussian and non-linear SSMs, however, a linear and Gaussian state processes is retained. This restriction is addressed in Kurle et al. (2020) and Ansari et al. (2021) by utilizing switching LGSSMs to introduce non-linearities in the model by modeling regime changes as stochastic processes. While DSSM also allows to detect different parameter regimes across time and different series, the employed global RNN performs this deterministically (Kurle et al., 2020).

3 Background

We adopt the notation of Rangapuram et al. (2018) throughout the remainder of the paper to facilitate comparability between their and our presentation of DSSM. Let there be a panel of N time series indexed by $i \in \mathcal{I} = \{1, 2, \dots, N\}$. All real-valued time series are univariate and we let $z_t^{(i)} \in \mathbb{R}$ denote the observation for series i at equally spaced time points $t \in \{1, 2, \dots, T_i\}$. The full history of series i is denoted by $z_{1:T_i}^{(i)} = (z_1^{(i)}, z_2^{(i)}, \dots, z_{T_1}^{(i)})$ where we assume for notational brevity that there are no missing values even though the state space model framework handles them in a straightforward manner (Durbin and Koopman, 2012). There exists further a sequence of D exogenous time-varying features $\boldsymbol{x}_{1:T_i+\tau}^{(i)} = (\boldsymbol{x}_1^{(i)}, \boldsymbol{x}_2^{(i)}, \dots, \boldsymbol{x}_{T_i+\tau}^{(i)})$ that are associated with the respective observation at time t, i.e. $\boldsymbol{x}_t^{(i)}$ are the features to forecast $z_t^{(i)}$. With $\tau \in \mathbb{N}_{>0}$ denoting the forecast horizon, we aim to estimate

$$p(z_{T_{i+1}:T_{i+\tau}}^{(i)} | \mathcal{F}_{T_i}^{(i)})$$

for each time series in $\mathcal{I}$, where $\mathcal{F}_{T_i}^{(i)}$ denotes the σ-field of the full information on the unknown stochastic process for series i that is available to the forecaster at the forecast origin T_i. Following the global learning paradigm, we allow $\mathcal{F}_{T_i}^{(i)}$ to include past observations and features from all other series in $\mathcal{I} \backslash \{i\}$ with the restriction that all T_i are associated with the same absolute point in time, e.g. a particular year for annual data. This assumption prevents data leakage during training by avoiding forecasting based on future observations of other time series.

3.1 State space models

One of the two major design principles of DSSM is the assumption that each series i follows a stochastic process that is induced by a LGSSM. To formalize this idea let $\boldsymbol{l}_t^{(i)} \in \mathbb{R}^L$ denote the latent state of series i that evolves over time as a linear Gaussian Markov process

$$\boldsymbol{l}_t^{(i)} = \boldsymbol{F}_t^{(i)} \boldsymbol{l}_{t-1}^{(i)} + \boldsymbol{R}_t^{(i)} \boldsymbol{\eta}_t^{(i)}, \; \boldsymbol{\eta}_t^{(i)} \sim N(\boldsymbol{0}, \boldsymbol{Q}_t^{(i)}), \tag{1}$$

The state process allows to incorporate structural assumptions about the unknown data generation process by appropriate choices of the transition $\boldsymbol{F}_t^{(i)}$, selection $\boldsymbol{R}_t^{(i)}$ and covariance matrix $\boldsymbol{Q}_t^{(i)}$. The unobserved states are mapped to the time series observations through the linear observation process

$$z_t^{(i)} = \boldsymbol{a}_t^{(i)\top} \boldsymbol{l}_t^{(i)} + b_t^{(i)} + \epsilon_t^{(i)}, \; \epsilon_t^{(i)} \sim N(0, {\sigma_t^{(i)}}^2), \tag{2}$$

where $b_t^{(i)}$ is a scalar to model transient effects, such as promotions in retail, or deterministic short-term trends. Effects that persist over time are better modeled by including a similar coefficient in the state process. The parameterization of the LGSSM is completed by specifying a prior distribution for the initial state $\boldsymbol{l}_0^{(i)}$. Contrary the proposed estimation of an isotropic Gaussian distribution in Rangapuram et al. (2018), where the variance parameters are difficult to estimate, we suggest the diffuse initialization $\boldsymbol{l}_0 \sim N(\boldsymbol{0}, \kappa \boldsymbol{I}_L)$. This initialization is preferred by us since it is more parsimonious and produces almost identical parameter and state estimates for time series of sufficient length. The selection of κ as a non-tuneable parameter is discussed in Section 5.1. We refer the reader to Durbin and Koopman (2012) for a detailed introduction to the state space framework, including exact treatments of the initialization problem when states are non-stationary, and relevant LGSSM parameterizations.

3.2 Parameter estimation

In practice, parameters in $\boldsymbol{F}_t^{(i)}, \boldsymbol{R}_t^{(i)}, \boldsymbol{Q}_t^{(i)}, \boldsymbol{a}_t^{(i)}$ or $b_t^{(i)}$ and ${\sigma_t^{(i)}}^2$ must be estimated from data. Thus, let $\boldsymbol{\theta}_t^{(i)}$ be the P-dimensional vector of unknown model parameters at time t. The standard procedure is to maximize the likelihood

$$p(z_{1:T_i}^{(i)}; \boldsymbol{\theta}_{1:T_i}^{(i)}) = p(z_1^{(i)}; \boldsymbol{\theta}_1^{(i)}) \prod_{t=2}^{T_i} p(z_t^{(i)} | z_{1:t-1}^{(i)}; \boldsymbol{\theta}_{1:t}^{(i)}), \tag{3}$$

which follows from the chosen parameterization of the state (1) and observation (2) process. Analytical expressions exist for LGSSMs and likelihood evaluation is performed efficiently with the Kalman

filter (Durbin and Koopman, 2012). Estimating $P \times T_i$ parameters of the presented LGSSM without additional restrictions is numerically challenging and does not provide estimates for the forecast range $\{T_{i+1}, T_{i+2}, \dots, T_{i+\tau}\}$. This is circumvented by applying

$$\boldsymbol{\theta}_t^{(i)} = \Psi^{(i)}(\boldsymbol{x}_t^{(i)}; \boldsymbol{\phi}^{(i)}) \tag{4}$$

as a map between features and LGSSM parameters, where $\boldsymbol{\phi}^{(i)}$ are learnable parameters. $\Psi^{(i)}$ can be parametric, non-parametric, or a mixture of both since a different map can be applied to each parameter in $\boldsymbol{\theta}_t^{(i)}$. For example, temporal smoothness of parameters can be modelled with splines. It is then customary to maximize the conditional likelihood

$$p(z_{1:T_i}^{(i)} | \boldsymbol{x}_{1:T_i}^{(i)}; \boldsymbol{\phi}^{(i)}) = p(z_{1:T_i}^{(i)}; \boldsymbol{\theta}_{1:T_i}^{(i)}),$$

to obtain an estimate for $\boldsymbol{\phi}^{(i)}$, where we use the mapping (4) implicitly for notational brevity. (Rangapuram et al., 2018) extends precisely this idea into a global modelling framework. DSSM uses a global non-parametric ML model $\boldsymbol{\Psi}$ to map features of each time series to the respective SSM parameter values. For real-valued features $\boldsymbol{x}_t^{(i)} \in \mathbb{R}^D$, $\boldsymbol{\Psi}$ is a stacked model with RNN layers and a linear output layer of dimension equal to the number of unknown SSM parameters P. (we propose to use GRU layers (Cho et al., 2014) as they provide a good trade-off between performance and computational efficiency (Chung et al., 2014).

Since the output of the last layer is real-valued, an affine mapping is inserted into $\boldsymbol{\Psi}$ to transform parameter estimates to their respective domains, e.g. $\mathbb{R}^+$ for variance parameters. When categorical features are available, we suggest using an embedding layer and concatenate the output with the remaining real-valued features before feeding them into the input GRU layer. The learnable parameters $\boldsymbol{\phi}$ of $\boldsymbol{\Psi}$ are jointly estimated by maximizing the pseudo likelihood

$$p(z_{1:T_1}^{(1)}, \dots, z_{1:T_N}^{(N)} | \boldsymbol{x}_{1:T_1}^{(1)}, \dots, \boldsymbol{x}_{1:T_N}^{(N)}; \boldsymbol{\phi}) = \prod_{i=1}^{N} p(z_{1:T_i}^{(i)} | \boldsymbol{x}_{1:T_i}^{(i)}; \boldsymbol{\phi}) = \prod_{i=1}^{N} p(z_{1:T_i}^{(i)}; \boldsymbol{\theta}_{1:T_i}^{(i)}),$$

with $\boldsymbol{\theta}_t^{(i)} = \boldsymbol{\Psi}(\boldsymbol{x}_t^{(i)}; \boldsymbol{\phi})$. The pseudo likelihood results from a cross-sectional independence assumption among the time series in the panel since the full likelihood is intractable. It remains that the likelihood of a series $p(z_{1:T_i}^{(i)}; \boldsymbol{\theta}_{1:T_i}^{(i)})$ can be efficiently evaluated with the Kalman filter. DSSM is thus better viewed as a vehicle that aims to achieve two tasks. First, it imposes temporal smoothness on SSM parameter estimates with the recurrent structure of $\boldsymbol{\Psi}$. Due to the recurrence, the predicted parameters $\boldsymbol{\theta}_t^{(i)}$ are a function of $\boldsymbol{x}_{1:t}^{(i)}$. Second, using an RNN as a single global mapping $\boldsymbol{\Psi}$ between features and LGSSM parameters allows learning associations between related time series, hence perform data-driven pooling to estimate parameters that are similar for related time series.

3.3 Forecasting

Forecasting the future observations of time series i is carried out straightforwardly by applying the learned mapping $\boldsymbol{\Psi}$ to the feature sequence $\boldsymbol{x}_{1:T_i+\tau}^{(i)}$ to obtain $\boldsymbol{\theta}_{1:T_i+\tau}^{(i)}$. The joint distribution $p(z_{T_i+1:T_i+\tau}^{(i)} | z_{1:T_i}^{(i)}, \boldsymbol{\theta}_{1:T_i+\tau}^{(i)})$ can be approximated by Monte Carlo forward sampling (Rangapuram et al., 2018). The marginals $p(z_{T_i+k}^{(i)} | z_{1:T_i}^{(i)}, \boldsymbol{\theta}_{1:T_i+\tau}^{(i)}),\ k = 1, \dots, \tau$ are obtained more efficiently by applying the Kalman filter for $t = 1, \dots, T_i + \tau$, while merely treating the observations beyond T_i as missing (Durbin and Koopman, 2012).

4 Limitations

4.1 Methodological limitations

The original method of Rangapuram et al. (2018) restricts itself to parameter heterogeneity because identical LGSSM parameterizations are considered for all time series in the panel. We identify this to be a major limitation for real-world problems with heterogenous stochastic processes. An immediate remedy would be to choose a flexible LGSSM parameterizations that is able to approximate a wide range of stochastic processes. Under the assumption that all time series can be modelled as an ARMA(p, q) process, one may choose p and q large and expect the

global model to predict parameters close to 0 for processes with lower autoregressive and moving average orders. Data-driven pooling for time series with similar p and q through exogenous features in the global model is essential to achieve benefits over local maximum likelihood estimates. Preventing overfitting for shorter time series due to flexible LGSSM parameterizations requires then to perform regularization at the global model level, hence requires careful hyper-parameter tuning. However, this goes against the initial motivation to use parsimonious statistical models to alleviate the risk of overfitting with a flexible global model (Rangapuram et al., 2018).

From this example it becomes apparent that there must be a sweet spot for DSSM to be preferred over its alternatives. To provide some guidance let us classify forecasting problems based on the number of time series N and time series length T (see Figure 1). Time series model selection is challenging in problems with small T, since holding out valuable training observations for pseudo out-of-sample forecast evaluation reduces an already short training dataset even further. Instead, one should perform model validation and selection based on in-sample scores (Diebold, 2015). However, DSSM requires a validation set for hyper-parameter tuning or an implementation of early stopping criteria to prevent overfitting. Model selection is contrary less challenging for problems with large T for which several (cross-)validation strategies have emerged for time series data (Bergmeir et al., 2018; Lainder and Wolfinger, 2022). However, longer time series also allow to use more flexible models that do not impose the same state space process for each series. Local methods are applied in problems with few time series when cross-learning is unlikely to improve forecast accuracy or when time series are short. Furthermore, manual model and feature engineering remains computational and work-wise feasible. Global non-parametric models may contrary be preferred for large N problems because the complexity of model selection among a set local methods is reduced. When T is also large, one may choose a global non-parametric model. Here we refer to the experimental study in Rangapuram et al. (2018), where DeepAR (Salinas et al., 2020) outperforms DSSM on the *electricity* and *traffic* dataset for rolling one day-ahead predicitons when more training data becomes available. DSSM is conceptually thus suited for problems with moderate T, due the necessity of using pseudo out-of-sample model validation, and moderate N of time series from a single data generation process due to the restrictive assumption of model homogeneity.

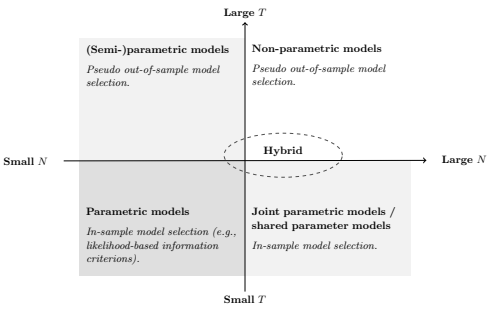

Figure 1: Classification of forecasting problems based on the number of time series, N, and the time series length, T.

4.2 Practical limitations

To apply DSSM successfully in real-world forecasting problems it is essential to choose a LGSSM that can approximate all time series in the panel. Indeed, using a DSSM as an automated forecasting method is challenging because domain expertise is required to determine the structure of the LGSSM. This is in contrast to DeepAR, which can be deployed fully automatically when time series are sufficiently long to perform hyper-parameter tuning. Choosing a LGSSM parameterization depends mostly on time series properties, such as the type of seasonalities, trends and cycles. Further considerations are based on suspected non-stationarities, i.e. properties such as mean and variance of the time series evolving over time, and the forecast horizon. Depending on the time series process, using autoregressive models may be recommended for short horizons and less preferred for longer horizons due to prediction error accumulation when forecasts are produced iteratively (Ben Taieb et al., 2012). All these considerations must be incorporated into the LGSSM and exposes DSSM to the consequences of model misspecification.

This might suggest to use a fairly rich LGSSM as a starting point, but we note that some care must be made to avoid model identifiability issues (Hamilton, 1994), and that a too flexible LGSSM structure may eliminate the possible advantages of the hybrid model approach due to increased complexity and potential optimization challenges. Hence, it is apparent to us that a selection of LGSSMs must be performed even when there is enough evidence that the time series panel is homogeneous. The

selection of LGSSMs becomes significantly more laborious for heterogeneous panels. Our first experiment in Section 5.1 demonstrates the inflexibility of DSSM due to the imposed stochastic process of the chosen LGSSM when the panel consists of a mixture of different data generation processes. The results demonstrate that a DSSM produces poor forecasts when the true stochastic processes cannot be approximated by the LGSSM. A potential remedy is to perform model selection locally. Multiple global models, each with a different LGSSM, are trained for the whole dataset. The final forecasts are then issued from the global model which performed best for each of the series on their respective test partition. However, using conventional train, validation, test splits is challenging for short time series.

Additional practical challenges stem further from the fact that input normalization is recommended to efficiently train neural networks. Exponential smoothing models, which can be parameterized in state space form (Hyndman, 2008), are used in the bottom-up approach of Smyl (2020) to perform data-driven input normalization. Indeed, Smyl (2020) can be viewed as a reversed variant of DSSM since the method first applies exponential smoothing locally and learns a global RNN on the residuals. This method does not suffer to the same extend from a misspecification of the parametric models because the global RNN in Smyl (2020) can learn from any signal that is left in the residual processes. Model selection can therefore be omitted to reduce the computational complexity. Another advantage of Smyl (2020) over DSSM is that their method is not restricted to Gaussian forecast distributions. Both methods have instead in common that hyper-parameter tuning is essential to produce accurate forecasts. We agree that imposing restrictions through the LGSSM in DSSM slows down overfitting. Nonetheless, our second experiment in Section 5.1 demonstrates that DSSM still relies on early stopping strategies because the global RNN can overfit the time-varying parameters of the LGSSM. Thus, successfully applying DSSM in real-world forecasting problems relies on the same techniques to monitor under- or over-fitting as other neural network forecasting methods. Alleviating overfitting with an imposed LGSSM structure has consequently little practical value.

Lastly but not least, DSSM and the variations hereof that we mention in Section 2 all suffer from high computational costs due to the loss-function evaluation which relies on each time series being processed by the Kalman filer sequentially. This leads not only to much slower training compared to non-hybrid methods, but also to high memory usage during back-propagation for long time series. In combination with the strong requirement of proper hyper-parameter tuning and model selection this severely restricts the practical applicability of DSSM.

5 Experiments

5.1 Monte Carlo experiments

In our first experiment, we demonstrate the rigidness that DSSM imposes with a parametric LGSSM and the consequences when the true data generation processes are not approximated well by the LGSSM. To verify our claims we consider a heterogeneous time series panel of 250 simulated time series that is a mixture of stationary and non-stationary processes. The stationary time series are generated from an ARMA(1, 1) process with autoregressive coefficient $\theta = 0.8$, moving average coefficient $\phi = 0.5$ and observation noise variance $\sigma_\epsilon^2 = 5$. Non-stationary time series are generated from a local level model with signal-to-noise ratio $\sigma_\eta^2/\sigma_\epsilon^2 = 1$ and measurement noise variance $\sigma_\epsilon^2 = 1$. We refer to the long version of this paper for more details on the parameterizations. For an introduction to both processes we refer the reader to Durbin and Koopman (2012). We generate panels that vary in the number of time series lengths $T \in \{120, 220, 420\}$ and mixture decomposition. The true data generation process of the ith time series is given by the ARMA(1,1) process with probability $\text{expit}(\mu + \boldsymbol{\beta}^\top \boldsymbol{X}_i)$ with the features $\boldsymbol{X}_i \sim \mathcal{N}(0, \boldsymbol{I}_5)$, and otherwise follows the local level process. The regression coefficient $\boldsymbol{\beta}$ is drawn randomly from a multivariate Gaussian distribution. We control the mixture of both processes by varying $\mu \in \{1.5, 0, -1.5\}$, which corresponds to approximately $25\%, 50\%$, and 75% of local level processes. The last 20 observations of each time series are reserved for rolling 1-step ahead out-of-sample forecast evaluation, hence 100, 200 and 400 observations are respectively available for training and hyper-parameter tuning of the DSSM.

Models To quantify the misspecification costs we use a DSSM with a local level model, i.e. a global RNN predicts the two time-varying variance parameters of a local level state space model. The RNN has a single GRU layer with hidden dimension being 15. We use the static standard normal

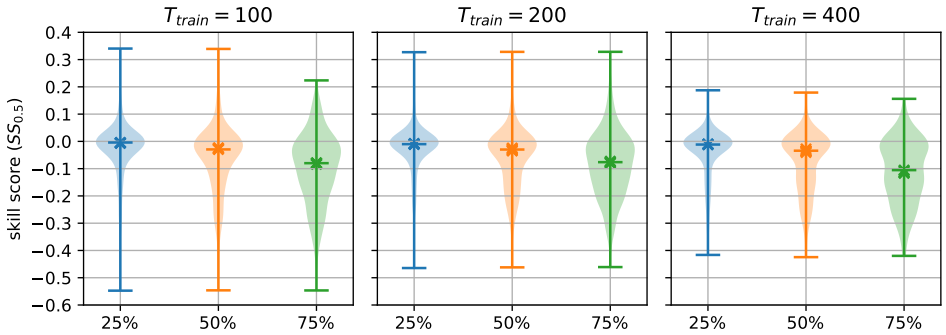

Figure 2: Test set skill score distribution of 0.5-quantile losses between a DSSM (losses averaged across 5 training runs) and locally estimated LGSSMs for the Monte Carlo experiment with mixture processes. The mixture ratio between ARMA(1,1) and local level processes increases from 25% to 75% of ARMA(1, 1) processes. Negative skill scores refer to superior forecast accuracy of a locally estimated LGSSM over DSSM. Markers indicate the median of the skill score distribution of each DSSM training run to represent the stochasticity of the training procedure.

features $\boldsymbol{X}_i$ and the time series ID as features. An embedding layer with dimension 50 is used for the categorical ID feature. We use the Adam optimizer of PyTorch (Paszke et al., 2019) to estimate the RNN and embedding parameters, and only tune the learning rate. Early stopping is implemented by tuning the number of training epochs. All hyper-parameters were tuned manually by using a conventional training and validation split (10% of the observations that remain after removing 20 test observations) as described in the Appendix B of the long version. We compare the global DSSM model against locally estimated ARMA(1, 1) and local level models. Maximum likelihood estimation is performed for each model on the full training length and the AIC (Akaike, 1974) is used for in-sample model selection. All time series are normalized based on their means and variances during the training range, which excludes the validation set for the DSSM to prevent data leakage. The normalization allows us to perform a diffuse state initialization $\boldsymbol{l}_0 \sim \mathcal{N}(0, 10\boldsymbol{I}_L)$ for all time series, irrespective of their original scales.

Results We assess the forecast accuracy of DSSM and the selected parametric LGSSM on the basis of $\rho-$quantile losses

$$P_\rho(z_{t+\tau}^{(i)}, \hat{z}_{t+\tau|t}^{(i)}) = \begin{cases} \rho(z_{t+\tau}^{(i)} - \hat{z}_{t+\tau|t}^{(i)}) & \text{if } z_{t+\tau}^{(i)} > \hat{z}_{t+\tau|t}^{(i)}, \\ (1-\rho)(\hat{z}_{t+\tau|t}^{(i)} - z_{t+\tau}^{(i)}), \end{cases}$$

where $\hat{z}_{t+\tau|t}^{(i)}$ is the estimated ρ-quantile of the forecast distribution issued at time t for lead time $t+\tau$. The forecast accuracy between both methods is compared for each time series with skill scores, $SS_\rho^{(i)} = 1 - \frac{S_{\rho,comp}^{(i)}}{S_{\rho,ref}^{(i)}}$, where we use $S_{\rho,ref}^{(i)}$ to denote the sum of $\rho-$quantile losses from the reference forecast model over the test set. Letting the parametric LGSSMs be the reference model, negative skill scores subsequently imply superior forecast accuracy over the competing DSSM because quantile losses are negative oriented scores. To assess the 1-step ahead forecast accuracy we use $S_\rho^{(i)} = \sum_{j=0}^{19} P_\rho(z_{T_{train}+j+1}^{(i)}, \hat{z}_{T_{train}+j+1|T_{train}+j}^{(i)})$. Figure 2 shows the obtained $\rho = 0.5$ skill score distribution for all 9 panels.

The results confirm the methodological limitation of a DSSM when the imposed LGSSM is unable to approximate the local stochastic process sufficiently well. Indeed, we choose a local level LGSSM for the DSSM purposely since it can only approximate an ARMA(1,1) process when $|\theta| \in [0, 1)$ is at either of the parameter space's bounds and $\phi = 0$. A growing number of negative skill scores is consequently observed for increasing proportions of ARMA(1, 1) processes in the panel, which is due to a decreasing number of time series being approximated well by the imposed local level model of the DSSM. The results further show that the skill score differences increase with increasing time series length. This effect stems from the decreasing parameter estimation variance due to increasing training data for the correctly specified LGSSMs. The DSSM benefits less from

increasing training data because the local level model is misspecified for at least 25% of the time series. Hence, obtaining more observations of a process that the local level cannot approximate is superfluous. The forecast accuracy of the DSSM would evidently improve if a ARMA(1, 1) LGSSM is used since $\theta \to 1$ will approximate a local level process. Thus, it is essential to perform model selection for DSSM at least at the global level to find a LGSSM that works well for a particular panel.

Circumventing the selection of a parsimonious LGSSM by choosing a richer parameterization comes at the price of higher computational costs and greater risk of overfitting. We subsequently use a panel with 250 ARMA(1, 1) processes of length 100 to verify that DSSM is equally exposed to overfitting as other neural network forecasting methods. The time series are assumed to have hourly frequency and we use the hour of day and day of the week as dynamic real valued features, where we use a cos and sin transformation to preserve the cyclic nature of both features. Figure 3 presents the evolution of training and validation losses for 5 different training runs. The global RNN starts overfitting after approximately 15 epochs when 90 observations of each series are used for training. Overfitting occurs due to the predicted time-varying autoregressive, moving average and variance parameters.

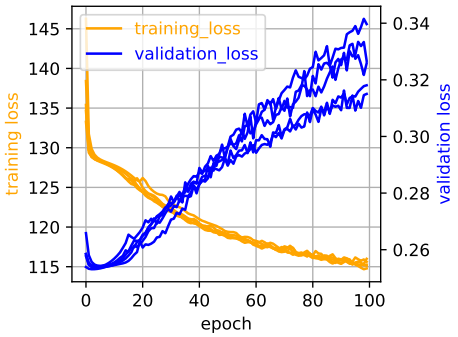

Figure 3: Training (negative log-likelihood) and validation (sum of rolling 1-step ahead weighted quantile losses) losses of DSSM (with a ARMA(1, 1) LGSSM) for a panel of 250 ARMA(1,1) processes.

5.2 Real-world data benchmarks

In our last experiment we re-evaluate the performance of DSSM on the publicly available *electricity*[1] and *traffic*[2] datasets. We employ the same LGSSM as in Rangapuram et al. (2018) with hourly and daily seasonal components. Rolling 1-week ahead (168 steps) forecasting is performed for 3 consecutive weeks, while using the previous 3 weeks for training and hyper-parameter tuning. All hyper-parameters were manually tuned on the last 10% of the 3 week long training time series, where we found the learning rate and number of training epochs to matter most. A detailed description of the hyper-parameter tuning procedure and results are presented in the Appendix C. The reason for not using an expanding training window stems from the increasing computational costs that do not result in additional accuracy improvements, which has been also observed in Rangapuram et al. (2018). We compare DSSM against locally estimated seasonal LGSSMs to investigate the benefits of learning the state space model parameters jointly on real-world homogenous datasets. A description of the seasonal LGSSM can either be found in Rangapuram et al. (2018) or in Appendix A. The local parametric LGSSM and the LGSSM of the DSSM are identical except that the parameters of the local LGSSM are time-invariant. Following Rangapuram et al. (2018), the time series ID is used as a static feature and the hour of the day and day of the week are used as dynamic real value features after they have undergone a cos and sin transformation. We additionally train a global feedforward neural network (FFN) to predict the time-invariant parameters of the same seasonal LGSSM to separate the effects of time-varying parameters and cross-learning. The FFN uses only the time series ID as a static feature, which is necessary since otherwise the global model predicts equal parameters for all time series.

Results The skill scores in Table 1 show no forecast accuracy improvements of the global hybrid over the locally estimated seasonal LGSSMs. The results are consistent across all forecast periods and hybrid training runs, which verifies the robustness of the results. Accuracy improvements of the global methods are contrary observed for the *traffic* dataset, for which we present a summary of the skill scores in Table 3 in the Appendix. The original DSSM method achieves the largest improvements, whereas minor improvements are obtained for the global FNN. We expect these results to be due to varying volatility on weekdays and weekends, an effect that does not seem to

[1] `https://archive.ics.uci.edu/ml/datasets/ElectricityLoadDiagrams20112014`
[2] `https://archive.ics.uci.edu/ml/datasets/PEMS-SF`

DSSM-RNN				
Quantile-loss	**Period**	$\overline{p_{10}}$ **(s.d.)**	$\overline{p_{50}}$ **(s.d.)**	$\overline{p_{90}}$ **(s.d.)**
0.5	1	-0.08 (0.008)	0.012 (0.005)	0.149 (0.007)
	2	-0.082 (0.007)	0 (0.003)	0.103 (0.007)
	3	-0.153 (0.109)	-0.003 (0.019)	0.141 (0.058)
0.9	1	-0.146 (0.021)	0.039 (0.005)	0.248 (0.014)
	2	-0.178 (0.012)	0.013 (0.005)	0.21 (0.011)
	3	-0.526 (0.475)	-0.025 (0.046)	0.254 (0.131)
DSSM-FFN				
0.5	1	-0.037 (0.008)	0.001 (0.002)	0.053 (0.003)
	2	-0.026 (0.003)	0 (0.001)	0.048 (0.009)
	3	-0.032 (0.003)	0.001 (0.001)	0.098 (0.007)
0.9	1	-0.13 (0.02)	0.001 (0.006)	0.147 (0.022)
	2	-0.104 (0.018)	0.006 (0.004)	0.143 (0.013)
	3	-0.086 (0.011)	0.001 (0.002)	0.198 (0.018)

Table 1: Test dataset skill scores for the *electricity* dataset. Negative scores indicate superior forecast accuracy of the local parametric LGSSM over the respective hybrid model with RNN and FFN architectures. Skills scores are shown for quantile losses 0.5 and 0.9 with a rolling forecast over 3 periods. The DSSM training is repeated 5 times and the average and standard deviation of the quantiles (10%, 50%, 90%) of the skill scores are reported.

be dominant in the *electricity* dataset. The effect can be captured by the RNN due to the prediction of a time-varying observation noise variance. The global DSSM-FFN contrary predicts a single time-invariant variance parameter. Furthermore, we observe the seasonal LGSSM to be slightly more misspecified for the *traffic* dataset than for the *electricity*. This observation is also supported by the overall larger quantile losses that are reported in Rangapuram et al. (2018). As a consequence, the local maximum likelihood estimates are more exposed to noise than the global models. Hence, we observe minor improvements of the global DSSM-FNN hybrid.

6 Conclusions

In this paper we identify substantial limitations in the design of DSSM (Rangapuram et al., 2018), a hybrid forecasting method that fuses deep learning with parametric state space models to forecast panels of related series. The method's major limitation stems from the assumed model homogeneity that is induced by adapting a single state space model parameterization for all time series in the panel. Our Monte Carlo simulation shows that DSSM produces inaccurate forecasts when the true stochastic process of a time series cannot be sufficiently well approximated by the imposed state space model structure. We further demonstrate that the parsimonious state space model structure does not guard against overfitting when its time-varying parameters are predicted by a global RNN. DSSM requires therefore the same techniques that prevent standard deep learning methods from overfitting, which makes it challenging to apply in forecasting problems with short time series. Our benchmarks for the *electricity* and *traffic* datasets show only minor forecast accuracy improvements of DSSM over locally estimated state space models with the same respective parametric structure. Overall, we find that the added computational and methodological complexity of using a RNN to parameterize local statistical time series models can only pay off in very specific forecasting problems. In case of DSSM, it requires a panel that is sufficiently long to perform hyper-parameter tuning and model selection, and where all time series follow approximately the same stochastic process. Other hybrid methods, such as the proposed extensions for DSSM in Kurle et al. (2020); de Bézenac et al. (2020); Ansari et al. (2021), aim to extend the applicability of DSSM to a broader class of forecasting problems by removing restrictions of the parametric state space models in DSSM. Ultimately, each of the methods operate under similar constraints as DSSM in terms of computational complexity and modelling assumptions, which requires a careful assessment of a methods applicability for a forecasting problem.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.

Ansari, A. F., Benidis, K., Kurle, R., Turkmen, A. C., Soh, H., Smola, A. J., Wang, Y., and Januschowski, T. (2021). Deep Explicit Duration Switching Models for Time Series. In *Advances in Neural Information Processing Systems*, volume 34. Curran Associates, Inc.

Baltagi, B. (2005). *Econometric analysis of panel data*. Wiley, 3. ed. edition.

Bandara, K., Bergmeir, C., and Hewamalage, H. (2021). LSTM-MSNet: Leveraging Forecasts on Sets of Related Time Series With Multiple Seasonal Patterns. *IEEE Transactions on Neural Networks and Learning Systems*, 32(4):1586–1599.

Bandara, K., Bergmeir, C., and Smyl, S. (2020). Forecasting across time series databases using recurrent neural networks on groups of similar series: A clustering approach. *Expert Systems with Applications*, 140:112896.

Ben Taieb, S., Bontempi, G., Atiya, A. F., and Sorjamaa, A. (2012). A review and comparison of strategies for multi-step ahead time series forecasting based on the NN5 forecasting competition. *Expert Systems with Applications*, 39(8):7067–7083.

Benidis, K., Rangapuram, S. S., Flunkert, V., Wang, Y., Maddix, D., Turkmen, C., Gasthaus, J., Bohlke-Schneider, M., Salinas, D., Stella, L., Aubet, F.-X., Callot, L., and Januschowski, T. (2022). Deep learning for time series forecasting: Tutorial and literature survey. *ACM Computing Surveys*, 55(6):1–36.

Bergmeir, C., Hyndman, R. J., and Koo, B. (2018). A note on the validity of cross-validation for evaluating autoregressive time series prediction. *Computational Statistics & Data Analysis*, 120:70–83.

Cho, K., van Merrienboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014). Learning phrase representations using rnn encoder–decoder for statistical machine translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Chung, J., Gülçehre, Ç., Cho, K., and Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *In NIPS 2014 Workshop on Deep Learning, December 2014*, abs/1412.3555.

Coakley, J., Fuertes, A.-M., and Smith, R. (2006). Unobserved heterogeneity in panel time series models. *Computational Statistics & Data Analysis*, 50(9):2361–2380. Statistical signal extraction and filtering.

de Bézenac, E., Rangapuram, S. S., Benidis, K., Bohlke-Schneider, M., Kurle, R., Stella, L., Hasson, H., Gallinari, P., and Januschowski, T. (2020). Normalizing Kalman Filters for Multivariate Time Series Analysis. In *Advances in Neural Information Processing Systems*, volume 33. Curran Associates, Inc.

Diebold, F. X. (2015). Comparing Predictive Accuracy, Twenty Years Later: A Personal Perspective on the Use and Abuse of Diebold–Mariano Tests. *Journal of Business & Economic Statistics*, 33(1):1–1.

Duncan, George T.and Gorr, W. L. and Szczypula, J. (2001). *Forecasting Analogous Time Series*, pages 195–213. Springer US, Boston, MA.

Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*. Oxford University Press.

Fröhwirth-Schnatter, S. and Kaufmann, S. (2008). Model-Based Clustering of Multiple Time Series. *Journal of Business & Economic Statistics*, 26(1):78–89.

Garcia-Ferrer, A., Highfield, R. A., Palm, F., and Zellner, A. (1987). Macroeconomic Forecasting Using Pooled International Data. *Journal of Business & Economic Statistics*, 5(1):53.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press, Princeton, NJ, 1st edition.

Hyndman, R. J. (2008). *Forecasting with Exponential Smoothing: The State Space Approach*. Springer Series in Statistics. Springer, Berlin Heidelberg.

Kurle, R., Rangapuram, S. S., de Bezenac, E., Günnemann, S., and Gasthaus, J. (2020). Deep Rao-Blackwellised Particle Filters for Time Series Forecasting. In *Advances in Neural Information Processing Systems*, volume 33. Curran Associates, Inc.

Lainder, A. D. and Wolfinger, R. D. (2022). Forecasting with gradient boosted trees: Augmentation, tuning, and cross-validation strategies. *International Journal of Forecasting*, 38(4):1426–1433.

Montero-Manso, P. and Hyndman, R. J. (2021). Principles and algorithms for forecasting groups of time series: Locality and globality. *International Journal of Forecasting*, 37(4):1632–1653.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.

Rangapuram, S. S., Seeger, M. W., Gasthaus, J., Stella, L., Wang, Y., and Januschowski, T. (2018). Deep state space models for time series forecasting. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.

Salinas, D., Flunkert, V., Gasthaus, J., and Januschowski, T. (2020). DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3):1181–1191.

Smyl, S. (2020). A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting. *International Journal of Forecasting*, 36(1):75–85.

Zhang, G. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing*, 50:159–175.

A State space model details

A.1 Local level model

A local level model is parameterized in Durbin and Koopman (2012) by the state and observation equations

$$\begin{aligned} l_t^{(i)} &= l_{t-1}^{(i)} + \eta_t^{(i)}, & \eta_t^{(i)} &\sim \mathcal{N}(0, {\sigma_{\eta,t}^{(i)}}^2) \\ z_t^{(i)} &= l_t^{(i)} + \epsilon_t^{(i)}, & \epsilon_t^{(i)} &\sim \mathcal{N}(0, {\sigma_{\epsilon,t}^{(i)}}^2), \end{aligned}$$

where $\boldsymbol{\theta}_t^{(i)} = ({\sigma_{\eta,t}^{(i)}}^2, {\sigma_{\epsilon,t}^{(i)}}^2)$ are the unknown parameters that are predicted by DSSM.

A.2 ARMA(1, 1)

An ARMA(1,1) process is parameterized in Durbin and Koopman (2012) by the state process

$$\begin{pmatrix} \tilde{z}_t^{(i)} \\ \phi_t^{(i)} \epsilon_t^{(i)} \end{pmatrix} = \begin{pmatrix} \rho_t^{(i)} & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \tilde{z}_{t-1}^{(i)} \\ \phi_{t-1}^{(i)} \epsilon_{t-1}^{(i)} \end{pmatrix} + \begin{pmatrix} 1 \\ \phi_t^{(i)} \end{pmatrix} \epsilon_t^{(i)}, \quad \epsilon_t^{(i)} \sim \mathcal{N}(0, {\sigma_{\epsilon,t}^{(i)}}^2),$$

with autoregressive $\rho_t^{(i)}$ and moving average $\phi_t^{(i)}$ coefficients, respectively. The observation process

$$z_t^{(i)} = \tilde{z}_t^{(i)}$$

is noise-free since the state process includes the serially correlated MA(1) process. DSSM predicts the unknown parameters $\boldsymbol{\theta}_t^{(i)} = (\rho_t^{(i)}, \phi_t^{(i)}, {\sigma_{\epsilon,t}^{(i)}}^2)$.

A.3 Seasonal component model

The seasonal model with hourly and daily components is described in Rangapuram et al. (2018) by the state process

$$\boldsymbol{l}_t^{(i)} = \boldsymbol{I}_{31} \boldsymbol{l}_{t-1}^{(i)} + \eta_t^{(i)} \boldsymbol{a}_t, \quad \eta_t^{(i)} \sim \mathcal{N}(0, {\sigma_{\eta,t}^{(i)}}^2),$$

where $\boldsymbol{a_t} = (1_{\{\text{hour}(t)=j\}_{j=1}^{24}}, 1_{\{\text{day}(t)=j\}_{j=1}^{7}}/24)^\top$ is an indicator vector. All seasonal components are allowed to vary over time due to the employed random walk, where the state innovation variance ${\sigma_{\eta,t}^{(i)}}^2$ controls the level of adaptivity to changes in the time series process. The proposed state process is parsimonious in that a single variance parameter is shared between hourly and daily states. A more flexible parameterization would contrary estimate two separate variances independently to circumvent model misspecification for processes where either of the seasonal components is constant over time. The division of the daily indicator vector $1_{\{\text{day}(t)=j\}_{j=1}^{7}}$ by 24 follows from employing a single shared state innovation variance. The chosen state process implies that a daily state is "updated" 24 times during a 24 hourly cycle. A division by 24 consequently distributes the total daily state variation evenly across its hourly cycle, such that the overall variability of hourly and daily states due to a shared ${\sigma_{\eta,t}^{(i)}}^2$ is the same. The state space model is completed by the observation process

$$z_t^{(i)} = \boldsymbol{a}_t^\top \boldsymbol{l}_t^{(i)} + \epsilon_t^{(i)}, \quad \epsilon_t^{(i)} \sim \mathcal{N}(0, {\sigma_{\epsilon,t}^{(i)}}^2),$$

and DSSM predicts the unknown parameters $\boldsymbol{\theta}_t^{(i)} = ({\sigma_{\eta,t}^{(i)}}^2, {\sigma_{\epsilon,t}^{(i)}}^2)$. The presented model can be viewed as an extended local level model for which each seasonal component follows a random walk.

B Monte Carlo simulations

B.1 Hyper-parameter tuning

We employ a conventional train, validation and test split in this experiment. Given a panel with 220 observations per series, we initial remove the 20 most recent observations of each series to perform out-of-sample forecast accuracy evaluation for a trained DSSM. The remaining 200 observations are

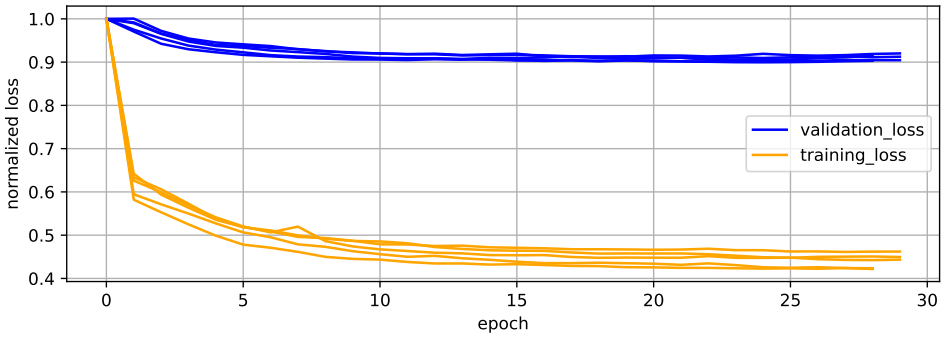

Figure 4: DSSM training and validation loss (5 runs) for the Monte-Carlo simulation (Section 5.1) with a panel of 50% local level processes and $T_{train} = 400$. The losses are normalized by their respective values after the first training epoch.

partitioned into a training (180 observations) and validation (20 observations) set. For each panel we tune the number of GRU layers and the hidden dimensions of DSSM, and keep the output dimension of the embedding layer fixed at 50. We only tune the learning rate of the Adam optimizer and use the default values of the PyTorch (Paszke et al., 2019) implementation for the remaining parameters. Early stopping is implemented by monitoring the validation loss over the number of training epochs. We further implement a total training time budget for each DSSM fit of 1 hour to limit computing time. The validation loss is calculated by performing rolling 1-step ahead forecasting.

Overall, we find the validation loss to be fairly insensitive with respect to the RNN hyper-parameters, given a sufficiently small learning rate ($< 5 \times 10^{-2}$) is employed. Figure 4 shows the evolution of the training and validation loss for a RNN with one GRU layer and 15 hidden units, and a learning rate of 3×10^{-3}. The validation loss converges after approximately 15 epochs, while the training loss decreases further at a very slow rate due to the flatness of the log-likelihood function for the imposed local level sub-model. An identical behavior is observed for the other eight panels. Thus, we proceed with a learning rate of 3×10^{-3} and stop training after 20 epochs.

B.2 Overfitting experiment - additional information

We generate 250 ARMA(1, 1) time series with $\theta = 0.8$, $\phi = 0.5$ and $\sigma_\epsilon^2 = 5$. Each time series is 100 observations long and an hourly frequency is assumed. The RNN has a single GRU layer, 15 hidden units and we employ a learning rate of 10^{-3}. An embedding layer with output dimension equal to 50 is used to encode the categorical time series ID. We further use time-varying input features for the RNN. Let $x_t^{(hour)} \in \{1, \dots, 24\}$ be the hour during which observation $z_t^{(i)}$ has been recorded. We avoid one-hot-encoding by employing the transformations

$$x_{sin,t}^{(hour)} = \sin(x_t^{(hour)}), \quad x_{cos,t}^{(hour)} = \sin(x_t^{(hour)})$$

to obtain two real-valued features. The same transformations are applied to the day of the week feature $x_t^{(day)} \in \{1, \dots, 7\}$.

C Real-world data benchmarks

C.1 Dataset partitioning

Electricity dataset The electricity load dataset contains load measurements for 370 customers that are recorded at 15 minute intervals from 2011-01-01 00:15:00 until 2015-01-01 00:00:00. We convert the time series to an hourly frequency by summing the respective four observations of a given hour. The last observation of the dataset is thus available for the one hour period that ends at 2015-01-01 00:00:00. The last 6 weeks are used in our experiment for rolling 1-week ahead forecast evaluation. Customer *MT_223* is removed from the dataset because consumption has been zero since 2013. Period 1 in Table 1 uses therefore the first three weeks for training and hyper-parameter tuning

	electricity	*traffic*
GRU layers	1	1
hidden units	40	30
learning rate	0.0025	0.0015
epochs	20	2

Table 2: Selected hyper-parameters for real-world data benchmarks.

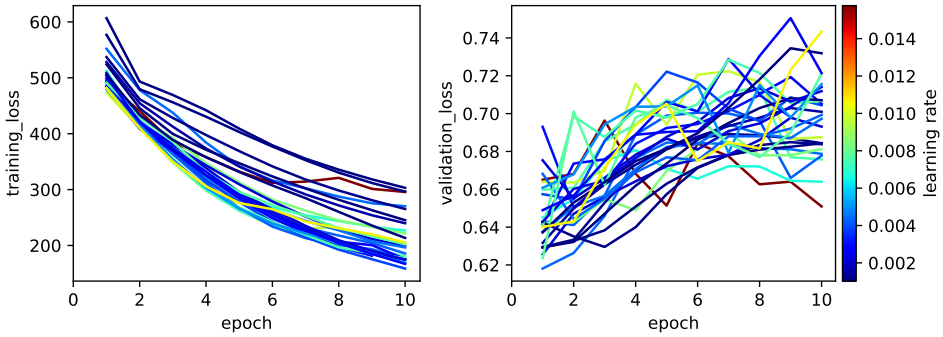

Figure 5: Validation and training losses for *traffic* dataset with varying learning rate (shown in color bar) and number of hidden units (not shown). The number of GRU layers is 1. More severe overfitting occurs for RNNs with two GRU layers.

and predicts week 4. Due to the rolling window approach we use weeks 2-4 for training in period 2 and predict week 5. Last, weeks 3-5 are used in period 3 for training and forecasts are made for week 6, hence the last 168 observations of the hourly time series.

Traffic dataset The traffic dataset requires additional pre-processing steps due to missing data. The dataset in its original form contains 10 minute averages for the occupancy rate of 963 care lanes of the San Francisco bay area. We obtain hourly time series by averaging six respective observations of a given hour. All time series span a period of 15 months and measurements are available from the 1st of January 2008. The dataset source states that public holidays are removed from the dataset. Therefore, to avoid the treatment of missing values we select a six week long period without any public holidays. The period starts with the 73th day of the original dataset and we impute the start of the the obtained time series to be Monday, the 3rd of March 2008. We then repeat the steps of the *electricity* dataset to partition the dataset into training and test sets. That is, period 1 uses the first three weeks for training and hyper-parameter tuning and forecasts are produced for week 4.

C.2 Hyper-parameter tuning

We tune the same hyper-parameters as described for the Monte Carlo simulation in Section B. 10% of the training data, i.e. 50 hours of the 3 week long training data, are used to validate a set of hyper-parameters with a pseudo out-of-sample forecasting scheme. All 50 steps of the validation set are predicted from the end of the training data in accordance with the 1-week ahead forecasting task for the test set, i.e. we do not perform rolling 1-step ahead forecasting as in the Monte Carlo simulation study. The hyper-parameters are only optimized once for the first training period of each dataset. All applied hyper-parameters are summarized in Table 2. Figure 5 verifies that overfitting occurs early for the *traffic* dataset, thus the number of training epochs is chosen to be comparatively small.

C.3 Additional results

Due to a page limit we report skill scores for the *traffic* dataset in Table 3. At this point we would like to provide additional insights into the forecast accuracy improvements of DSSM-RNN and DSSM-FFN for this dataset. Figure 6 shows two clear distinctions between both considered datasets. First, there is significantly greater hour to hour variation in the *traffic* dataset due to greater road usage

DSSM-RNN				
Quantile-loss	**Period**	$\overline{p_{10}}$ **(s.d.)**	$\overline{p_{50}}$ **(s.d.)**	$\overline{p_{90}}$ **(s.d.)**
0.5	1	-0.069 (0.018)	0.066 (0.02)	0.327 (0.024)
	2	-0.166 (0.016)	0.102 (0.011)	0.374 (0.009)
	3	-0.227 (0.041)	0.075 (0.034)	0.319 (0.023)
0.9	1	-0.123 (0.055)	0.09 (0.022)	0.449 (0.042)
	2	-0.209 (0.052)	0.072 (0.021)	0.377 (0.023)
	3	-0.213 (0.082)	0.178 (0.032)	0.416 (0.043)
DSSM-FFN				
0.5	1	-0.027 (0.006)	0.022 (0.001)	0.345 (0.007)
	2	-0.198 (0.012)	0.039 (0.005)	0.304 (0.002)
	3	-0.193 (0.015)	0.012 (0.009)	0.269 (0.011)
0.9	1	-0.078 (0.01)	0.069 (0.006)	0.501 (0.002)
	2	-0.279 (0.046)	0.014 (0.012)	0.343 (0.01)
	3	-0.366 (0.049)	0.027 (0.039)	0.325 (0.02)

Table 3: Test dataset skill scores for the *traffic* dataset. Negative scores indicate superior forecast accuracy of the local parametric LGSSM over the respective hybrid model with RNN and FFN architectures. Skills scores are shown for quantile losses 0.5 and 0.9 with a rolling forecast over 3 periods. The DDSM training is repeated 5 times and the average and standard deviation of the quantiles (10%, 50%, 90%) of the skill scores are reported.

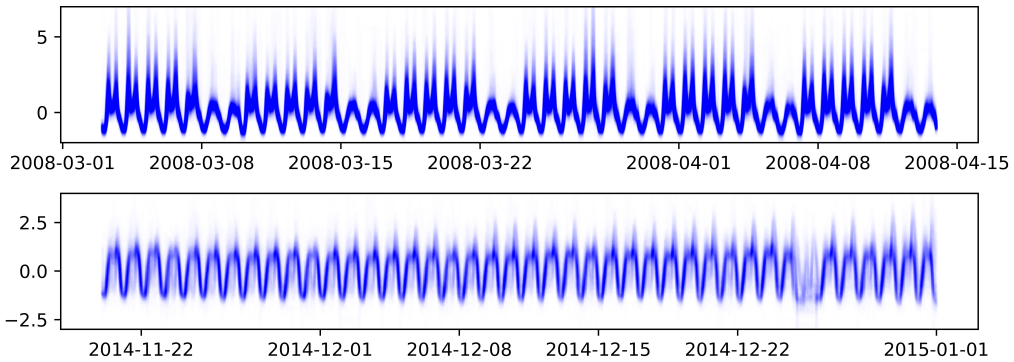

Figure 6: Standardized *traffic* (top row) and *electricity* (bottom row) datasets. The time series subsets are created by following the steps in Section C.1. Each time series is standardized by subtracting the mean and subsequently divide by the standard deviation.

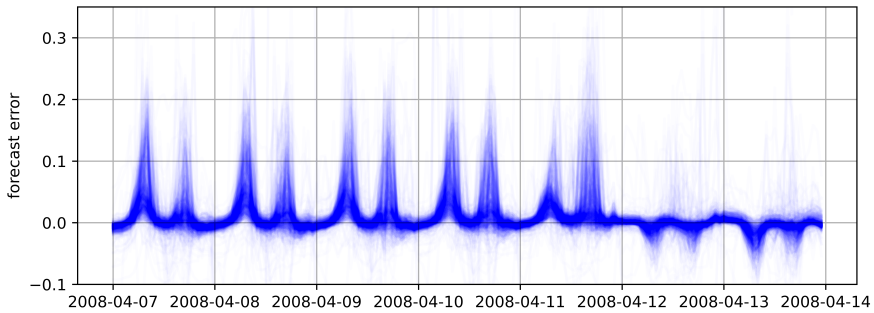

Figure 7: Forecast errors (observation - forecast distribution median) of DSSM-RNN (one training run) for the 3rd period of the *traffic* dataset. Each line represents the errors of one time series in the set over the prediction horizon of 168 hours.

during commuting hours in the morning and afternoon than during night times. This consequently implies forecast heteroscedasticity, which is not modeled by DSSM-FFN and the local parametric LGSSMs since a time-invariant observation noise variance is estimated. Second, the *traffic* data shows a clear weekend effect in terms of missing peak hours during mornings and afternoons, albeit occupancy rates of the remaining weekend hours being approximately the same as during weekdays.

From the presentation of the seasonal LGSSM in Section A it becomes apparent that the model cannot capture this effect because the forecasts are the sum of the hourly and daily state estimates (see observation equation). This parameterization is restrictive since it only models weekend effects where all hours on Saturdays and Sundays are shifted up- or downwards compared to other weekdays. Each local and global model has therefore to find a trade-off for this misspecification and choose between larger forecast errors that either occur on weekdays or weekends. Figure 7 shows this trade-off for DSSM-RNN, where the peak hours are predicted systematically too low (positive errors) on weekdays and slightly too high (negative errors) on weekends. This deficiency reemphasizes our discussion in Section 4 on the limitations of DSSM and the requirement for LGSSM selection. A more suitable LGSSM would contrary estimate hourly states for weekdays and weekends without the using any states to capture daily variations. The described effect is also visible in the forecasts (right column in Figure 5) that are presented in the long-version of Rangapuram et al. (2018). Eventually, we find DSSM-RNN to produce more accurate forecasts because modelling heteroscedasticity with time-varying innovation variances can reduce the state estimation bias of the misspecified state process. We find this to be the case because DSSM-RNN predicts large forecast variances for peak hours, which consequently makes the hourly state estimates less subjective to variations due to large forecast errors during peak hours.

C.4 Visualization of forecasts

Figures 8 and 9 show out-of-sample forecasts of DSSM-RNN for the test set of both public datasets.

D Repository

The material to reproduce our results is provided on GitHub at `https://github.com/benesom/critical-dssm`.

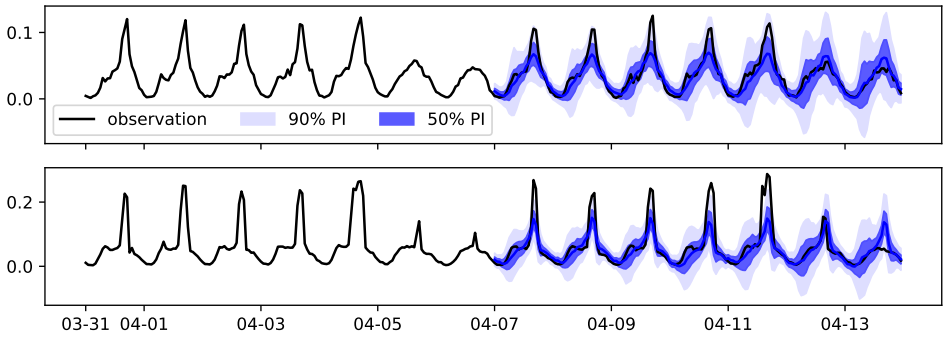

Figure 8: 90% and 50% prediction intervals (PIs) for the 3rd period of the *traffic* dataset. The forecasts are generated by DSSM-RNN and the two time series have been randomly selected.

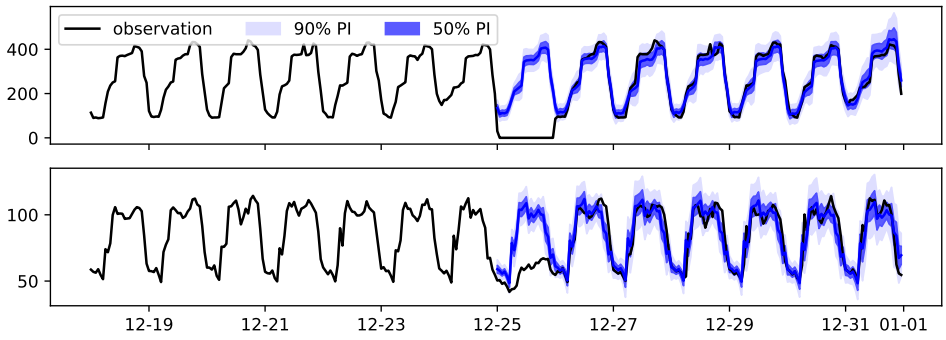

Figure 9: 90% and 50% prediction intervals (PIs) for the 3rd period of the *electricity* dataset. The forecasts are generated by DSSM-RNN and the two time series have been randomly selected.