

Adding Value to Promotional Forecasting through a Decision-Support Dashboard Quantifying Error Impacts

A Case Study at Picnic

Master Thesis

Frans W.A. Kraanen

Adding Value to Promotional Forecasting through a Decision-Support Dashboard Quantifying Error Impacts

A Case Study at Picnic

by

Frans W.A. Kraanen

to obtain the degree of Master of Science

at the Delft University of Technology,

to be defended publicly on Monday November 17, 2025 at 15:30 PM.

Student number:	4727282	
Project duration:	May 16, 2025 – November 17, 2025	
Thesis committee:	Dr. ir. A.J. van Binsbergen,	TU Delft, Chair
	Dr. J.H.R. van Duin,	TU Delft, Supervisor
	Ir. G. Bandini,	Picnic, Supervisor

This thesis is confidential and cannot be made public until November 17, 2026.

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.



Preface

Presented here is the final version of my Master's thesis, completed as part of the Transport, Infrastructure and Logistics programme at Delft University of Technology. When I first came into contact with Picnic, I could not have imagined that this would eventually lead me to work with their office in Paris and conduct such a fascinating case study. My ambition was to combine logistics and commerce with the opportunity to go abroad, and in that respect this project exceeded all expectations. Finding a topic that aligned with both the academic requirements of TU Delft and the practical ambitions of Picnic was at times a challenge, and I am therefore particularly grateful for the trust and flexibility shown by both parties.

The completion of this thesis marks not only the conclusion of my research project, but also the end of my studies in Delft. It has been an exciting and insightful period, during which I have had the privilege of studying alongside a talented group of students and staff working at the forefront of civil technology. I owe my sincere thanks to all those who have contributed to my academic and personal development throughout this period. In particular, I would like to express my gratitude to my thesis supervisors, Arjan van Binsbergen and Ron van Duin, for their motivation, insightful feedback, and ability to put matters into perspective. Their guidance has been invaluable in bringing this research to a successful conclusion.

The past six months, dedicated to this thesis research, have been both demanding and deeply rewarding. Every hour invested feels worthwhile, and I am proud to present the final result. This project has not only allowed me to deepen my understanding of academic research, but has also broadened my view of the practical applications of logistics and demand forecasting. For this, I am grateful to Picnic and to all colleagues who supported me throughout the process. A special word of thanks goes to my supervisor at Picnic, Gianluca Bandini, who was always accessible whenever I wished to discuss a topic. Our valuable conversations allowed me to gain knowledge and significantly improve the quality of this study.

Finally, I wish to thank my family, friends, and everyone close to me for their unwavering support and encouragement throughout my thesis and my entire studies. Their motivation and friendship have been essential throughout this journey.

It is my hope that this research may inspire future readers in the field of demand forecasting to continue seeking innovative solutions to the considerable challenges that will always persist.

*Frans W.A. Kraanen
Paris, November 2025*

Abstract

Promotional forecasting poses a significant challenge for online supermarkets, as consumer demand tend to be highly volatile. Such volatility often leads to forecasting errors, which in turn have financial, operational, and reputational consequences. This study investigates the financial and operational challenges in the context of Picnic France and addresses the knowledge gap concerning the measurement of the diverse impacts of over- and under-stocking caused by forecasting inaccuracies.

The research objective was to develop and embed an artefact that quantifies the consequences of forecasting errors in financial and operational terms. The Design Science Research Process (DSRP) was applied to structure the study. The phases were structured as follows; first the 'Problem Identification and Motivation' phase was conducted in the introduction setting out the problem, structure the research and introduce the case study. The next phase 'Objectives of a solution' is conducted by mapping the complete promotional forecasting process to understand problems and clearly pinpoint where the artefact can address the problem. In the 'Design and Development' the indicators and their calculations are defined, followed by the 'Demonstration' phase which shows the design of the artefact, a decision supportive dashboard. Afterwards, the evaluation phase presents the results for two articles and subsequently discusses the feedback received from end-users. The report finalises with the 'Communication' phase, presenting how the dashboard will be integrated in the existing process. This approach resulted in a dashboard that will be supportive in two key steps of the promotional forecasting process: the first when products are selected for promotion, and the second when Picnic's analysts review the forecasted quantities and decide whether to adjust them based on the insights provided by the dashboard.

The study offers both academic and practical contributions. Academically, it adds to the limited body of research focusing on the operational impacts of forecasting errors and on understanding how these impacts influence inventory management decisions. By integrating financial and operational cost factors into the assessment of forecasting performance, the study advances the understanding and quantification of over- and under-stocking risks. This claim is supported by a preliminary literature review, which revealed no relevant studies addressing these aspects. Practically, the developed dashboard serves as a valuable decision-support tool within Picnic, aiding in both selecting products for the promotion and reviewing the forecast quantities to be stocked for the promotion.

Contents

Preface	i
Abstract	ii
Nomenclature	viii
I Problem Identification and Motivation	1
1 Introduction	2
1.1 Background	2
1.1.1 Company Background	2
1.1.2 Case Study Background	3
1.2 Research Problem	4
1.3 Research Objective	6
1.4 Research Scope	6
1.5 Research Questions	7
1.6 Design Science Research Process	8
1.6.1 Research Approach	8
1.6.2 Research Outline	9
II Objectives of a Solution	13
2 Literature Review	14
2.1 Systematic Literature Review	14
2.1.1 Search Strategy	14
2.1.2 Selection Phase	15
2.2 Literature Review and Exploratory Field Study	16
2.2.1 Forecasting in the Online Supermarket Business	17
2.2.2 Process Flow	18
2.2.3 Forecasting Indicators	19
2.2.4 Costing Methods for Operational Indicators	20
2.2.5 Design and Implementation Considerations for Dashboards	24
2.3 Discussion on the Literature Research	25
2.4 Main Findings Chapter 2	26
3 Mapping the Promotional Forecasting Process	27
3.1 Rationale for Mapping the Promotional Forecasting Process	27
3.2 Mapping Methodology	28
3.2.1 Choice of Visualisation Methods	28
3.2.2 Approach	29
3.3 Final Result of the Promotional Forecasting Process	29
3.3.1 Elaboration on the Mapped Process	31
3.4 Main Findings Chapter 3	31

III	Design and Development	32
4	Selection and Justification of Indicators	33
4.1	Methodology for Indicator Selection	33
4.1.1	Long-list Identification and First Selection Process	33
4.1.2	Potential Indicators in Brief	34
4.1.3	Relevance-Feasibility Matrix Application	38
4.1.4	Stakeholder Validation Process & Result of the RF Matrix	39
4.2	Justification of Indicator Selection and Exclusion	39
4.2.1	Financial Indicators	40
4.2.2	Operational Indicators	40
4.2.3	Reputational Indicators	41
4.2.4	Notion on Chosen Indicators	42
4.3	Main Findings Chapter 4	42
5	Indicator Calculation Approach	43
5.1	Determination of Cost for Financial Indicators	43
5.1.1	Lost Margin	43
5.1.2	Waste Cost	45
5.1.3	Promo Extension	46
5.2	Determination of Cost for the Operational Indicators	47
5.2.1	Handling Cost	47
5.2.2	Space Cost	49
5.2.3	Maximum Economical Retention Time	51
5.3	Main Findings Chapter 5	52
IV	Demonstration	53
6	Dashboard Design	54
6.1	Rationale for the Design Approach	54
6.2	Artefact Requirements	54
6.2.1	Requirements Definition	54
6.2.2	Outline Artefact	55
6.2.3	Elicit Requirements	56
6.3	SQL Model Design	57
6.3.1	Chronological Process of the Design	57
6.3.2	General Notions on the Design	58
6.4	The Design of the Dashboard	58
6.4.1	Use of the Dashboard	58
6.4.2	Dashboard Metrics and Measurement Details	62
6.5	Main Findings Chapter 6	63
V	Evaluation	64
7	Dashboard Results and Evaluation	65
7.1	General Results Gained from the Dashboard	65
7.2	Setup of Pilot Case Study	67
7.2.1	Presentation of Pilot Case Study Articles	67
7.2.2	The Pilot Case Study	69
7.2.3	Analytical Scope and Indicator Relevance	69
7.3	Results of the Pilot Case Study	70
7.4	End-user Evaluation on the Dashboard	73
7.5	Main Findings Chapter 7	75

VI	Communication	76
8	Dashboard Integration	77
8.1	Embedding the New Artefact in the Current Process Flow	77
8.2	Communication of the Dashboard within the Organisation	80
8.2.1	Internal Communication	80
8.2.2	External Communication	80
8.3	Main Findings Chapter 8	80
9	Conclusions, Recommendations and Reflections	81
9.1	Conclusion	81
9.2	Recommendations	83
9.2.1	Short-term Recommendations	83
9.2.2	Mid-term Recommendations	84
9.2.3	Long-term Recommendations	85
9.3	Reflections	86
9.3.1	Academic Reflection	86
9.3.2	Societal Reflection	87
A	SQL Code Github Repository	92
B	Literature Research Background	93
C	Process Mapping Deliverables	96
C.1	Import-Effect Matrix	96
C.2	RASCI Process Flow Deliverables	96
C.3	New Promotional Forecasting Process Flow	99
D	Practical Deliverable: Process Bottlenecks and Recommendations	101
D.1	Identified Bottlenecks and Improvement Areas	101
D.1.1	Process Inefficiencies	101
D.1.2	Operational Risks	102
D.1.3	Communication	102
D.2	Recommendations	103
E	Indicator Selection Background	107
E.1	Relevance-Feasibility Matrix Development	107
E.2	Indicators Longlist	108
E.3	Shortlist of Excluded Indicators	110
F	Demonstration & Evaluation Phase	111
F.1	Feedback from Dutch End-Users	111
F.2	Evaluation of the Articles	112
F.2.1	Quick Approximation of Reputational Damage	112

List of Figures

1.1	Overview of research exploration level.	6
1.2	High-level overview of the stages where the dashboard can support the forecasting process at Picnic.	7
1.3	DSRP Approach used in this study with a Problem Centered Approach, adapted from Peffers et al. (2006).	9
1.4	Outline of the full research with all phases of the DSRP linked to chapters, research questions and methods.	10
2.1	Three-step protocol (Kitchenham, 2007).	14
2.2	Overview of FC0, identifying the scope of the research: the ambient zone.	21
2.3	Simple warehouse cost-tree (from Richards (2018)).	24
3.1	Approach steps process flow.	29
3.2	Swimlane diagram of complete process.	30
4.1	Research indicator scope.	34
4.2	Final RF-Matrix result.	39
5.1	Process Flow of extra handling steps in the ambient zone of FC0, located in Lille.	48
6.1	Define Requirements structure (Johannesson & Perjons, 2014).	55
6.2	Overview 'Explanation of the Dashboard' sheet	59
6.3	Overview 1 'Forecast Impact Dashboard' sheet, column A to O.	60
6.4	Overview 2 'Forecast Impact Dashboard' sheet, column P to AI.	60
6.5	Overview of the 'Graphical overview' sheet.	61
6.6	Overview 'Dashboard - Recommendations' sheet	62
7.1	Candlestick of over- and under-forecast impact per item and per campaign	66
7.2	Candlestick of each Indicator separate	67
7.3	Impact of Use Case 1 on KPIs, at both campaign and item level	71
7.4	Impact of Use Case 2 on KPIs, at both campaign and item level	73
8.1	Overview of the revised part of the process flow. <i>1. Early Fulfilment / Inbound analyst involvement, 2. Forecast review by other teams, 3. Use of the new Dashboard</i>	79
B.1	ABC implementation flow chart by Almeida and Cunha (2017).	95
C.1	Impact-Effort Matrix of the recommendations found with the process analysis.	96
C.2	RASCI process flow -10wk to -7wk.	97
C.3	RASCI process flow -3wk to +1wk.	98
C.4	Complete overview of the new Swimlane overview of the promotional forecasting process.	100
E.1	Overview of Stakeholder Modifications During RF-Matrix Validation	108

List of Tables

2.1	Paper selection criteria.	15
2.2	Categorised overview of sources used in the Literature Review	16
2.3	Overview of Process Mapping Approaches Examined in Greater Detail	19
5.1	Time window and definitions for waste categories and promo extensions.	45
5.2	Overview of FC0's Physical Dimensions and Monthly Costs. <i>Data extracted from map of FC0 and finance department.</i>	50
5.3	All used formula's in the dashboard	52
6.1	Overview of the functional and non-functional requirements.	57
7.1	Overview of the two articles used for the pilot case study	68
7.2	Overview of all indicators used in the dashboard, ranked by their relevance for product selection and forecast quantity review.	69
7.3	General Metrics of Case Article 1	70
7.4	Cumulative Impacts of KPIs of Case Article 1	71
7.5	General Metrics of Case Article 2	72
7.6	Cumulative Impacts of KPIs of Case Article 2	72
7.7	Overview of the end-users of the artefact (dashboard).	73
7.8	List of actionable feedback received on requirements, ready for the feedback-loop.	74
B.1	Overview of all preparatory literature	94
D.1	Categorised process related problems and recommendations ranked.	103
E.1	Adjustments on RF-Matrix per stakeholder	107
E.2	Longlist of Indicators	109
E.3	Non-chosen Indicators background	110

Nomenclature

Abbreviations

Abbreviation	Definition
ABC	Activity Based Costing
ABS	Activity Based Space Costing
ADR	Average Delivery Rate
API	Application Programming Interface
BBD	Best Before Date
BPMN	Business Process Model and Notation
CM	Category Manager
COGS	Cost of Goods Sold
CTE	Common Table Expression
CU	Consumer Unit
DSRP	Design Science Research Process
DWH	Data Warehouse
FC	Fulfilment Centre
FC0	Specific Fulfilment Centre located in Lille
FC1	Specific Fulfilment Centre located nearby Paris
GMI	Gross Margin Item
ITM	Intermarché
KPI	Key Performance Indicator
MERT	Maximum Economical Retention Time
RASCI	Responsible, Accountable, Supportive, Consulted, Informed
RC	Rolling Cart
RF-Matrix	Relevance-Feasibility Matrix
SF-Promo	Supplier-Funded Promotion
SQL	Structured Query Language
TIL	Transport, Infrastructure and Logistics
TU	Trading Unit



Problem Identification and Motivation

1

Introduction

This introductory chapter outlines the foundations of the study and sets the stage for the research that follows. It positions the thesis within its organisational and academic context, clarifies the problem under investigation, and specifies the objectives, scope, and research questions. In doing so, it provides the reader with a clear roadmap of how the study is structured and how its various components connect.

The remainder of this chapter introduces the research in a structured manner. Section 1.1 provides the background, outlining Picnic as a company, its international expansion, and the specific context of its French operations, before moving into the organisational details of promotional forecasting. Section 1.2 defines the research problem, followed by the research objective in Section 1.3. Section 1.4 provides the study's boundaries, and Section 1.5 formulates the main and supporting research questions. Finally, Section 1.6 presents the methodology, introducing the Design Science Research Process (DSRP) as the overarching approach guiding the research.

1.1. Background

This section provides the necessary background for understanding both Picnic as a company and the context of the case study. In Section 1.1.1, Picnic is introduced as an organisation, outlining its origins, operating model, and international growth. This is followed by a closer look at Picnic France, the specific country setting of this study. This establishes a broad yet concrete picture of the company and its operations, thereby laying the groundwork for the following section, Section 1.1.2, where additional background is provided to better understand the operational context within Picnic France and clarify the study's relevance.

1.1.1. Company Background

Picnic is a Dutch online supermarket founded in 2015 by Michiel Muller, Frederik Nieuwenhuys, Joris Beckers, and Bas Verheijen. From the outset, Picnic distinguished itself as a pioneer in the supermarket industry by adopting a purely digital model. Unlike traditional competitors, the company has no physical stores: all orders are placed via a mobile application and delivered directly to customers' homes (Vigneshwar Ravichandran, 2024).

What makes Picnic distinctive is its relentless focus on efficiency. By eliminating the costs of storefronts, employing a fleet of electric delivery vans, and increasingly automating its distribution network, the company has created a highly streamlined operation. Technology lies at the heart of this model, enabling Picnic to scale quickly while offering free delivery and competitive pricing. This vision has attracted strong backing from influential investors, including German retailer Edeka, NPM Capital, and the Bill & Melinda Gates Foundation Trust (Vigneshwar Ravichandran, 2024).

Since its inception, Picnic has expanded rapidly. By 2023 the company employed over 17,000 people and generated sales of €1.25 billion (Carter, 2024). That year marked a turning point: in the Netherlands, Picnic reached profitability (positive EBITDA), largely due to automation in its robotic distribution centres. At group level, however, the company reported a net loss of €220.6 million in 2023, reflecting

the high costs of expansion (Dayeeta Das, 2024). After consolidating its position in the Netherlands, Picnic entered Germany in 2018 and France in 2021, continuing its strategy of aggressive international growth.

Picnic France

This research is conducted in collaboration with Picnic's French office, making France the focal point of the case study. Picnic entered the French market in 2021, beginning in the Hauts-de-France region around Lille. At launch, the service was available to approximately 500,000 households, and adoption quickly reached hundreds of thousands of orders (Stefan van Rompaey, 2022).

Building on this initial foothold in the north, Picnic has since expanded into the Paris region. Its network now covers approximately 1.5 million households, with the stated goal of serving the entire Île-de-France area by 2025 (Reul, 2025). This expansion has required significant operational adaptation, including the establishment of fulfilment hubs and a reconfiguration of supply chain partnerships. Over the past few years, Picnic has shifted sourcing arrangements multiple times, from Cora, to Système U, and most recently to Intermarché, illustrating both the complexity and dynamism of operating in the French grocery market (Stefan van Rompaey, 2022).

Strategically, France holds a crucial position within Picnic's international ambitions. Alongside the Netherlands and Germany, it provides a distinct testing ground to demonstrate that the business model can succeed across very different consumer, cultural, and logistical contexts. In financial terms, Picnic France remains smaller, approximately one-sixth the size of the Dutch operations and one-fifth that of Germany, but its role is pivotal. By proving scalability in France, Picnic strengthens its claim to be not only a national disruptor, but a European one.

1.1.2. Case Study Background

Unless otherwise indicated, the information in this subsection is based on knowledge gathered during the research process, primarily through discussions with Picnic employees and the review of internal documents. Additionally, some terms and practices used within Picnic are discussed in this subsection to clarify their meaning when referenced.

Forecasting at Picnic

At the heart of Picnic's operations lies demand forecasting, a critical process that influences various aspects of the business, including inventory management and delivery scheduling. Accurate forecasting ensures that customers receive their orders on time, while also minimising food waste and optimising warehouse capacity. However, forecasting in the online supermarket sector is particularly challenging due to high demand volatility and rapidly changing consumer behaviour (Solis, 2024).

At Picnic, forecasting is a straightforward process: predictions are made using a model based on an algorithm, and the results are reviewed by a promotion analyst for validation and adjustment if necessary. However, since the model is relatively new within Picnic France, forecasts are often inaccurate and this is often not captured by the promotion analyst.

The following subsection outlines several other mechanisms that influence forecasting, which are important for the remainder of this study.

A Single Supplier: Intermarché

Picnic France works almost exclusively with one supplier, the French supermarket chain Intermarché (ITM). Although a small share of Picnic's assortment consists of private-label products, the company effectively operates with ITM as its sole supplier. This partnership is defined by a number of agreements that significantly influence Picnic's promotional operations.

Ten Weeks Upfront Ordering

Promotional planning at Picnic is carried out on the basis of a ten-week commitment. ITM shares a list of products of which Picnic can choose which will be promoted. This list is shared 10 weeks upfront, and Picnic must generate its demand forecasts accordingly. For the vast majority of products, backorders cannot be placed, meaning that under-forecasting cannot be compensated after the fact. As a result, the impact of backorders in the French market is negligible.

In/Out Products in Promotion

Picnic's assortment comprises approximately 11,000 unique consumer units (CUs), compared to more than 20,000 offered by ITM. Consequently, many products proposed by ITM for promotion are not part of Picnic's regular assortment. Picnic must then decide whether or not to include these products in its promotional portfolio.

This decision is inherently risky: there is no historical sales data for products never previously sold. To manage this uncertainty, Picnic evaluates potential substitutes in its assortment, considers their sales performance, and assesses the quality of the promotion itself. Products introduced only for the duration of a promotion and subsequently removed from the assortment are referred to as In/Out products.

Funding of Promotions

All supplier-funded promotion (SF-Promo) at Picnic France are for 100% funded by ITM. This arrangement makes running SF-promo highly attractive, especially when compared to markets such as the Netherlands where funding is organised differently. Funding applies only to the products actually sold; in cases of overstock, Picnic bears the full cost of the surplus.

Depending on the product type, Picnic applies different strategies to manage excess stock:

- **Regular assortment articles:** straight after the SF-promo an extended promotion with a 15% discount is applied until the surplus falls below an amount equal to 30 days of stock.
- **In/Out articles:** straight after the SF-promo an extended promotion with a 25% discount is applied for one week, after which unsold products are discarded or donated.

Promotional Product Selection

It is important to note that Picnic is not responsible for deciding which products are potentially eligible for promotion. This list is provided by ITM. Picnic then decides which of these products will actually be offered for promotion. Since ITM's assortment is larger than Picnic's regular assortment, some products on the list are not typically sold, and careful consideration is required before selecting them for promotion. Failure to sell these products is highly undesirable, as they must be cleared quickly after the promotion. To facilitate this, Picnic applies a self-funded promotion of 25% for one week, after which any remaining products are discarded. Consequently, selecting such products involves significant risk.

For promotional selection, Picnic always includes products from the regular assortment. Other products are selected based on estimated popularity and the quality of the promotional mechanism, with larger discounts generally increasing the likelihood of selection.

Over- and Under-Forecasting

A forecast is an estimate of expected sales generated by Picnic's forecasting model. Retrospectively, it may be observed that either too much stock was held (an overestimation) or too little (an underestimation), both of which incur costs that can be quantified. Based on such historical cost information, Picnic may choose to adjust stock levels relative to the model's forecast: increasing stock to reduce the risk of under-stocking, or decreasing stock to limit the risk of overstocking. Deliberately ordering more stock than the forecast indicates is referred to as over-forecasting, whereas ordering less than the forecast suggests is referred to as under-forecasting.

This concept is central to the present study, as the dashboard developed here is designed to support this process. By providing insights based on historical cost impacts, an analyst at Picnic may adjust the model-generated forecast to manage risks or address systematic limitations of the forecasting model. The application of such an analysis is described in detail in Section 7.2.2.

1.2. Research Problem

Promotional forecasting remains a persistent challenge for retailers, particularly in technology-driven contexts such as Picnic France. During promotions, demand volatility increases, and baseline forecasting methods often struggle to capture the uplift caused by marketing efforts, price changes, and product placement. This discrepancy leads to stock imbalances: over-stocking ties up capital, while under-stocking limits sales and damages customer satisfaction. Research shows that promotion-based

forecasting typically lags behind baseline accuracy, with errors amplified when forecasting at article level (Koottatep & Eng, 2006).

At Picnic, which relies heavily on next-day delivery, even small mispredictions can flow through the supply chain, triggering stock-outs or surplus inventory. Although Picnic France mainly uses algorithms to solve the promotional forecasting challenge, some specific promotional articles cannot be included in these algorithms due to the lack of historical data. These are the “In/Out” products, items that are temporarily added to the assortment for a promotion and removed afterwards. Currently, these products are forecasted based on spreadsheet inputs from substitutes or human judgement. However, these methods lack consistency and scalability. Moreover, not only these types of promotions suffer from high forecasting error rates; products in promotion with historical data also experience significant inaccuracies. With Picnic, a company selling products with high-frequency, mostly focusing on technical operations, a more systematic, data-driven solution is essential to align with other operations within Picnic. As Picnic France is relatively new to the field and has a promotional system and logic that differs significantly from its counterparts in the Netherlands or Germany, a tailor-made approach is needed.

In addition to the difficulties surrounding promotional forecasting, another knowledge gap emerged: the impact of forecasting itself. Specifically, what is the impact of a product being over-predicted, and conversely, under-predicted? Several factors affect the costs for Picnic, which can be categorised into three areas: financial impact, operational impact, and reputational impact, the last two also having indirect financial consequences. Ultimately, this gap in knowledge was also found to exist in the Netherlands and Germany. Therefore, this presents an interesting gap in understanding, which could also contribute to improving forecasting practices in both the Dutch and German branches of Picnic.

Finally, Picnic faces challenges in selecting the most suitable products for promotion, as outlined in Section 1.1.2. These difficulties largely stem from a lack of knowledge. Additional insights into products or comparable items would therefore be highly valuable, particularly one that highlights the impact of forecasting errors.

In conclusion, addressing the promotional forecasting problem should begin by understanding the impact of forecasting errors at the article level. These insights can then inform the development of a more effective promotional forecasting system.

Promotional forecasting remains a persistent challenge for retailers. During promotions, demand volatility increases, and baseline forecasting methods often struggle to capture the uplift caused by marketing efforts, price changes, and product placement. This discrepancy leads to stock imbalances: over-stocking ties up capital, while under-stocking limits sales and damages customer satisfaction. Research shows that promotion-based forecasting typically lags behind baseline accuracy (Koottatep 2006).

At Picnic, which relies heavily on next-day delivery, even small mispredictions can flow through the supply chain, triggering stock-outs or surplus inventory. Although Picnic France mainly uses algorithms to solve the promotional forecasting challenge, some specific promotional articles cannot be included in these algorithms due to the lack of historical data. Currently, these products are forecasted based on spreadsheet inputs from substitutes or human judgement. However, these methods lack consistency and scalability. Moreover, not only these types of promotions suffer from high forecasting error rates; products in promotion with historical data also experience significant inaccuracies.

In addition to the difficulties surrounding promotional forecasting, a knowledge gap emerged: the impact of forecasting errors. Specifically, what is the impact of a product being over-predicted, and conversely, under-predicted? Several factors affect the costs for Picnic, which can be categorised into three areas: financial impact, operational impact, and reputational impact, the last two also having indirect financial consequences. Therefore, this presents an interesting gap to do research in.

This leads to the goal of creating an artefact that is able to determine for each article in the inventory what the cost of one product would be if it was over-predicted or under-predicted to improve the forecasting accuracy. The scope of this research is only focused on the financial and operational impact of forecasting errors.

1.3. Research Objective

Building upon the challenges outlined in Section 1.2, this research focuses on on addressing the limited understanding of the complete impacts of promotional forecasting errors at Picnic. While broader improvements to the core promotional forecasting model remain important, they fall outside the scope of this study and may form the basis for future investigation.

The objective of this study is to develop and embed an artefact that assesses the effects of both over-stocking and under-stocking. By packaging this framework into a user-centred dashboard, the intention is to provide forecasting teams with clear, actionable insights into the consequences of over- and under-estimation at the article level. This addition to the forecasting process supports the selection of products for promotion and enables risk-averse adjustments to the forecasts generated by the model.

Ultimately, this study aims to demonstrate how such a framework can serve as a valuable supplementary input within Picnic’s promotional forecasting process, improving transparency and the quality of decision-making. By addressing a focused but critical gap in the understanding of forecasting impacts, this research makes a key contribution: it generates additional knowledge on the financial and, more importantly, operational consequences of forecasting errors; an aspect that is highly valued within Picnic and remains largely under-explored in the scientific literature.

1.4. Research Scope

The scope of this research, as defined in Chapter 4, is focused and clearly outlined. Promotional forecasting errors can result in three types of consequences: direct financial loss, additional operational effort, and reputational damage. These impact areas are commonly referenced in both academic literature and practical applications when assessing supply chain or inventory forecasting performance (Kerkkänen et al., 2009; Roeloffzen, 2007).

Figure 1.1 provides a visual overview of the scope considered in this study. It distinguishes between the impact dimensions that are partly monitored within Picnic (financial), those that are newly introduced in this research (operational), and those that are currently excluded from scope (reputational). This figure helps to contextualise the operational scope examined in this research and clarifies the boundaries of the current analysis.

In addition, the scope is deliberately narrowed to the context of promotional activities, as these represent a distinct and high-impact subset of forecasting decisions within Picnic. Only ambient products are included, while frozen and chilled products are excluded because their operational processes differ. Given the time constraints of this research, the analysis is further limited to a single zone of the fulfilment centre. In consultation with the Inbound Lead, it was determined that ambient products are the most relevant, as they account for the largest share of the total assortment. The research focuses specifically on Picnic France as a case study, while ensuring that the findings are scalable and transferable to other countries within Picnic’s operations. Furthermore, the analysis relies on historical data, enabling a structured and evidence-based evaluation of past forecasting performance and its consequences.

By explicitly defining which impact types are included, this study ensures a focused and actionable analysis, while also laying the groundwork for potential future research into areas such as reputational impact, which are currently outside the scope of Picnic’s operational focus.



Figure 1.1: Overview of research exploration level.

1.5. Research Questions

The main objective of the research described in Section 1.3 is to develop and embed an artefact that assesses the effects of both over-stocking and under-stocking, packaged in a dashboard based on historical data. This dashboard needs to serve as an additional decision-making input within Picnic’s promotional forecasting process, with a focus on understanding and balancing the impacts of over- and under-stocking.

Figure 1.2 provides a simplified, chronological view of the potential impact the dashboard can have within Picnic’s current forecasting process.




Figure 1.2: High-level overview of the stages where the dashboard can support the forecasting process at Picnic.

With this clear view of the potential use of the dashboard, and to provide an academic framework and structure for this research, a main research question and associated questions have been defined. Answering the main research question will directly contribute to addressing the problem described in Section 1.2. The main research question is:

How can an artefact, designed to provide insights into the impacts of over- and under-stocking, function as a decision-supportive tool for (a) selecting suitable promotional products and (b) reviewing and adjusting forecasted order quantities within Picnic France?

Research- and Sub-Questions

The main research question has been broken down into several focused research questions. Each addresses a distinct aspect of the larger issue, allowing for a more structured and manageable investigation. These research questions collectively guide the analysis, from diagnosing current practices to designing, implementing, and embedding the proposed solution.

- Q1.**
The subject of the first research question is to document Picnic’s current promotional forecasting setup by mapping workflows, responsibilities, and interactions to uncover potential bottlenecks and inefficiencies.
- *How is the current promotional forecasting process structured at Picnic, in terms of workflows, team responsibilities, and interdepartmental coordination?*
- Q2.**
The second research question aims to identify and understand the metrics that reflect the impact of forecasting errors by clarifying what Picnic measures and why these indicators matter.
- *Which key performance indicators should be used to assess the impacts of over- and under-stocking caused by promotional forecasting errors at Picnic and what insights does Picnic aim to derive from each indicator?*

Q3.

The third research question focuses on the financial impact of the operational indicators, assessing the tangible operational costs associated with incorrect forecasts to support cost-optimised decision-making.

- *How should the indicator calculations be designed to assess the cost implications of over-stocking and under-stocking resulting from forecasting errors in Picnic's promotional context?*

Q4.

The fourth research question is about conceptualising and structuring a solution: designing an artefact that transparently quantifies over- vs. under-stock impacts. This research question is divided in 2 sub-questions, to better structure the question within the research.

- *How can a user-friendly artefact be developed to assess the effects of over-stocking and under-stocking caused by forecasting errors?*

The two sub-questions that follow from this research question:

- **SQ4.1** *How should the artefact be designed to calculate and present the impacts of the indicators while serving as a supportive tool for product selection and the review of forecasted order quantities?*
- **SQ4.2** *What insights can be derived from the dashboard results, in what ways can these insights support product selection and forecast quantity review, and how do end-users evaluate its usefulness?*

Q5.

The fifth research question explores the application of this solution, ensuring that the dashboard integrates smoothly into current practices and genuinely supports decision-makers.

- *How can the dashboard be effectively integrated into Picnic's current promotional forecasting workflow and communicated across the organisation?*

1.6. Design Science Research Process

This section presents the research methodology used to address the main research question. It first introduces the DSRP as the overarching approach (Section 1.6.1), outlining its suitability for the problem at hand. Section 1.6.2 then explains how this study applies the six phases of the DSRP to design, implement, and evaluate the proposed solution. Each phase description follows a goal, question, method, output logic to maintain clarity and methodological coherence. The DSRP approach is based on the papers of Dresch et al. (2015), Johannesson and Perjons (2014), Peffers et al. (2006), and Pries-Heje et al. (2017). In addition, inspiration was drawn from the Transport, Infrastructure and Logistics master's course 'TIL Research and Design Methods' and from the design theories of Wolfert (2023).

1.6.1. Research Approach

This research adopts the DSRP, as introduced by Peffers et al. (2006), to address a concrete operational challenge: improving promotional forecasting at Picnic by developing a decision-support artefact. DSRP is designed for contexts where the objective is not only to study an existing problem but also to create, demonstrate, and evaluate an artefact that offers a practical solution. Its strength lies in combining scientific rigour with practical relevance, ensuring that outcomes are both theoretically grounded and operationally applicable.

The artefact developed in this study shows error forecasting impact on article-level in the form of an interactive dashboard. This tool quantifies the operational and financial impacts of over- and under-stocking due to forecasting errors and presents them in an accessible format for decision-makers. The research is problem-driven, solution-oriented, and iterative—the defining features of DSRP, shown in Figure 1.3.




Figure 1.3: DSRP Approach used in this study with a Problem Centered Approach, adapted from Peffers et al. (2006).

Figure 1.3 presents the DSRP model template from Peffers et al. (2006), adapted for this specific research. The red-bordered boxes indicate the adjustments made, while the overall structure remains unchanged. As shown in Figure 1.3, the model comprises six nominal phases. Although the phases suggest a sequential flow, the model retains its iterative nature. Furthermore, the DSRP allows multiple entry points, in this thesis, the process begins at the first phase; representing a problem-centred approach.

In this thesis, the process begins with identifying the core business problem, unknown impact caused by promotional forecasting errors, followed by a process mapping exercise to better understand the objectives needed for a solution. Together with the requirements, this informs the design and iterative development of the dashboard. The artefact is then demonstrated using real-world historical data from Picnic France, evaluated with stakeholder feedback and performance analysis, and prepared for integration into operational workflows.

1.6.2. Research Outline

To provide a clear and concise overview of the research structure, this section first presents a visual and after a textual outline of the DSRP phases in Figure 1.4, tailored specifically to this thesis. Each phase is linked to the relevant chapters, research questions, and applied methods, showing how the study progresses from problem identification to the final communication phase. The structure follows the DSRP model of Peffers et al. (2006), with adjustments to it described in Section 1.6.1, ensuring a coherent and logically organised methodological approach.




Figure 1.4: Outline of the full research with all phases of the DSRP linked to chapters, research questions and methods.

Phase I - Problem Identification & Motivation

The first phase of the DSRP aims to clearly articulate the core problem and build a compelling motivation for solving it (Johannesson & Perjons, 2014). In this study, the phase also focuses on defining the objective, research questions, scope, and linked methods, ultimately with the goal of designing an artefact that can address the identified shortcomings. Although not directly linked to a specific research question, this phase lays the conceptual and empirical foundation for all subsequent research activities.

Two complementary methods were employed: observation and document study. Observation involved closely examining the real-world practices and behaviours of stakeholders engaged in the forecasting process, which made it possible to identify inefficiencies, challenges, and hidden workarounds that practitioners might otherwise overlook. In parallel, the document study analysed internal process descriptions, reports containing forecasting statistical insights, and meeting notes to reveal both the fundamental structure of the process and the significance of the problem to be addressed. Alongside these two methods, preparatory literature research was conducted to build a deeper understanding of

the broader forecasting domain (Johannesson & Perjons, 2014).

The resulting insights not only confirmed the existence of structural forecasting inefficiencies but also clarified their scope and relevance, ensuring a solid and well-substantiated starting point for the subsequent phases of the DSRP.

Phase II - Objectives of a Solution

The second phase of the DSRP focuses on translating the problem understanding from Phase 1 into concrete next steps in the form of objectives for the solution. At this stage, the emphasis is on clarifying the structure of the existing forecasting process and identifying where interventions, either before or after the forecasting stage, could help increase the overall effectiveness of promotions. By mapping the current process and identifying gaps, this phase delivers a comprehensive view of the workflow around forecasted products in a promotion, from start to finish, thereby clarifying the underlying causes of forecasting challenges.

Three complementary methods were applied solely to create the process overview presented in Chapter 3. First, literature research was conducted to explore different process mapping approaches, identify the most suitable method for visualising stakeholder interactions and workflow complexity, and draw inspiration for relevant dashboard requirements (see Section 2.2.5). Second, observation and document study were used to develop an initial high-level outline of the promotional forecasting process, drawing on both first-hand insights from practice and existing organisational records. Finally, expert interviews were held with each stakeholder group involved in the process, excluding suppliers, to validate and refine the process flow, ensuring both accuracy and completeness.

The output of this phase is a validated process overview that provides the groundwork for defining the requirements in the next phase.

Phase III - Design and Development

The third phase of the DSRP, Design and Development, focuses on transforming the problem understanding and process insights from Phases I and II into a concrete, functional solution (Peppers et al., 2006). In this study, this phase results in the selection of relevant indicators and the design of the calculation approach that was implemented in the SQL code written in the next phase. This phase is linked to Research Questions 2 and 3. Research Question 2 guides the identification and definition of key indicators relevant to forecasting accuracy, while Research Question 3 addresses the quantification of operational cost implications arising from over- and under-stocking.

To address Research Question 2, a structured identification process was followed. This began with desk research combining academic literature and internal documentation to identify candidate indicators. These were then evaluated using a structured prioritisation approach based on a Relevance-Feasibility Assessment Matrix (RF-Matrix), ensuring that only the most impactful and implementable indicators were retained. The RF-Matrix was selected because relevance and feasibility were identified by Picnic's stakeholders as the most critical criteria for assessing the suitability of the indicators. The shortlist was finalised through expert validation with domain stakeholders. For Research Question 3, the development of the operational indicator calculations was achieved by gathering background knowledge internally, consulting relevant literature, and applying time-to-money methods such as Activity-Based Costing (ABC) and Activity-Based Space Costing (ABS) (see Section 2.2.4). These approaches enabled the translation of performance indicators into quantifiable cost implications, thereby linking forecasting performance directly to operational and financial outcomes.

The outcome of this phase is a list of relevant indicators, together with their associated calculations, designed to measure the cost of error forecasting impacts at both campaign and item levels. These were established through extensive reviews, including the RF-Matrix approach and expert validation.

Phase IV - Demonstration

The fourth phase of the DSRP focuses on demonstrating the development process and the functionality of the artefact. This phase complements Sub-question 4.1, which focuses on demonstration of the design of the code (calculations) and the dashboard itself, ensuring that the result remains a functional and useful decision-support tool for product selection and forecast quantity review.

The code was developed through iterative prototyping, following a step-by-step approach. Each step was reviewed by an analyst at Picnic, resulting in more than 25 development iterations, including several sub-steps, before the final version of the code was completed. This process formed the foundation for the dashboard design. The final dashboard design is visualised and clearly presented in Section 6.4, where each sheet of the Google Spreadsheet is illustrated and explained.

Phase V - Evaluation

The fifth phase of the DSRP focuses on evaluating the artefact to determine its effectiveness, relevance, and potential impact (Peppers et al., 2006). In this study, the evaluation was conducted in two parts. First, the dashboard results were analysed to assess its value and functionality. Subsequently, a real end-user evaluation was carried out, in which the predefined requirements were tested and general feedback from end-users was collected.

The data analysis was based on two articles that served as test cases, designed to demonstrate whether the dashboard successfully identified forecasting impacts, quantified operational consequences, and supported actionable recommendations. In the second part of this phase, the end-user evaluation examined how effectively the decision-support dashboard met the defined requirements and supported decision-making within the context of promotional forecasting.

Together, the iterative development process, guided by end-user feedback, and the analytical demonstration of selected articles provide a balanced perspective on the artefact's strengths and limitations. This combined assessment forms the foundation for the final phase, in which the artefact is integrated into the promotional forecasting process established in Phase II to support product selection and the review of forecasted quantities.

Phase VI - Communication

The sixth and final phase of the DSRP focuses on ensuring that the developed dashboard is effectively communicated, adopted, and integrated into Picnic's existing forecasting process. First, a clear explanation of the modifications to the current process is provided, and the role of the dashboard within this process is explicitly defined. Subsequently, the communication of the newly created artefact plays a crucial role in the overall success of the deliverable.

The goal of the communication is to ensure that all relevant stakeholders understand how and when to use the tool, recognise its benefits, and are able to apply it in their daily decision-making. The more effectively this is executed, the greater the value generated by the deliverable.

To achieve this, the communication strategy adopted in this study focuses on practical implementation. This includes presenting the dashboard's features and applications in dedicated stakeholder sessions and offering clear demonstrations to all involved parties. By tailoring the communication to the needs, terminology, and working routines of each stakeholder group, this phase maximises the likelihood that the dashboard will become a sustainable component of Picnic's promotional forecasting process.

The artefact's implementation marks an addition to the typical scientific application of the sixth phase of the DSRP method. While Dresch et al. (2015), Johannesson and Perjons (2014), Peppers et al. (2006) and Pries-Heje et al. (2017) describe only the final step of communicating the artefact and research, this study demonstrates that implementing the new deliverable and clarifying its proper use adds significant extra value to the research.

II

Objectives of a Solution

2

Literature Review

This literature review aims to build a solid theoretical and practical basis for the study that follows. It draws on the preparatory literature research carried out during the proposal phase as well as an ongoing review conducted throughout the thesis. The full details of the literature-search strategy and selection criteria can be found in Section 2.1. The approach to the writing style of the literature review is supported by the insights provided by Wee and Banister (2016).

At the outset, general literature background is provided in Section 2.2.1. After that, the best processing methods are explored in Section 2.2.2, followed by the search for potential KPIs in Section 2.2.3. In Section 2.2.4 the potential methods to monetise operational indicators are examined and Section 2.2.5 covers the final literature research topic, providing deeper knowledge on the creation of a dashboard and useful requirements to take into account upfront. The section concludes with a small discussion on the literature research in Section 2.3, reflecting on the information found and gaps identified. By the end of the chapter, the reader is expected to have a clear overview of which concepts and tools have been adopted and how these insights will shape the research design and analysis in later chapters.

2.1. Systematic Literature Review

The systematic literature review adopts the search and selection strategy guidelines of Kitchenham (2007), who propose a three-step protocol, shown in Figure 2.1. The goal of this strategy is to identify high-quality papers efficiently. Each step is explained in this section, together with its specific application in this study.




Figure 2.1: Three-step protocol (Kitchenham, 2007).

2.1.1. Search Strategy

The first step of the approach is the search strategy. For inclusion, a preliminary list of keywords was developed to ensure that the research is as comprehensive as possible. Boolean operators were used to combine search terms and refine the search results across databases. The selected papers span a publication range from 2000 to 2025, reducing the number of potential papers by applying a threshold on relevance related to date. However, studies published before 2000 that were used for the literature research were found via the snowballing process, later explained in the selection phase in Section 2.1.2. Only sources published in English were considered for inclusion. Relevant literature was retrieved using

different search engines: Scopus, Google Scholar and the TU Delft Repository. A total overview of the inclusion criteria for the paper selection is presented in Table 2.1.

Table 2.1: Paper selection criteria.

Inclusion Criteria	Description						
<i>Keywords</i>	(Promotional) forecasting	AND	Supermarket/Grocery/Retail	AND	Optimisation	AND	(Historical) data
	OR		OR		OR		OR
	Over-/under-stock		Warehouse		Operational		Waste
	OR		OR		OR		OR
	Trade-off Analysis		Dashboard		Customer Hapiness		Lost sales
<i>Document types</i>	Journal Article - Book - Thesis - Conference Paper/Report - Technical Blog/Report - Government Report						
<i>Language</i>	English						
<i>Time interval</i>	2000-2025						

2.1.2. Selection Phase

After identifying useful papers, it is important to apply clear selection criteria when deciding which studies to include. Therefore, all potentially relevant papers found in the first step of the three-step protocol by Kitchenham (2007) were reviewed. A small deviation from their approach was that the process was applied separately for each topic. After the initial search, more than 200 papers were identified. The next step involved reading the titles and abstracts in combination with predefined selection criteria, which served as a crucial filtering step in line with the protocol of Kitchenham (2007). Papers were either retained or discarded based on these criteria. The criteria used for this literature review were as follows:

1. The paper must address forecasting concepts with direct application to the retail sector.
2. The paper must examine a process flow, with a clear emphasis on explanatory approaches to process design.
3. The paper must discuss the impact of forecasting errors within the retail domain.
4. The paper must present empirical cases on the application of Activity-Based Costing (ABC) in warehouse environments.
5. The paper must discuss the design, implementation, or use of dashboards in professional contexts.

After this filtering step, around 75 papers remained. These were fully read and assessed for their relevance to the study. Following this full review, 30 papers were selected for inclusion. In addition, five further papers were identified using the snowballing technique, which involved reviewing the references of relevant articles to identify additional sources. This resulted in a total of 35 papers forming the basis of the literature review.

Finally, after selecting the papers, an overall overview was compiled. This is visualised in Table 2.2. The overview was organised by topic, paper, document type, and selection criteria.

In addition, a preparatory literature review was conducted earlier in the process. While this step generated valuable insights, its results were not directly included in this study. During this stage, 60 papers were used, and an overview of these preparatory sources is provided in Table B.1.

Table 2.2: Categorised overview of sources used in the Literature Review

Author(s), Year	Title	Document Type	Criteria
Forecasting in Online Supermarkt			
(Abolghasemi et al., 2019)	<i>Demand forecasting in supply chain: The impact of demand volatility in the presence of promotion</i>	Journal Article	1
(Bandyopadhyay, 2009)	<i>A Dynamic Model of Cross-Category Competition: Theory, Tests and Applications</i>	Journal Article	1
(Calvo-Porrá & Lévy-Mangin, 2017)	<i>Specialty food retailing: examining the role of products' perceived quality</i>	Journal Article	1
(Cooper et al., 1999)	<i>PromoCast: A New Forecasting Method for Promotion Planning</i>	Journal Article	1
(Evers et al., 2018)	<i>Demand forecast models for online supermarkets</i>	Paper / Thesis	1 and 3
(Gupta, 1988)	<i>Impact of Sales Promotions on when, what, and how Much to Buy</i>	Journal Article	1
(Heemskerk, 2024)	<i>Cold-Start Promotional Demand Forecasting with Contrastive Explanations in an E-grocery Setting</i>	Thesis	1 and 3
(Kootatpe & Eng, 2006)	<i>Promotional Forecasting in the Grocery Retail Business</i>	Thesis	1
(Ma et al., 2016)	<i>Demand forecasting with high dimensional data: The case of SKU retail sales forecasting with intra- and inter-category promotional information</i>	Journal Article	1
(Van Donselaar et al., 2016)	<i>Analysis and forecasting of demand during promotions for perishable items</i>	Journal Article	1
The Best Process Flow for Picnic			
(DeHoratius & Raman, 2008)	<i>Inventory record inaccuracy: An empirical analysis</i>	Journal Article	2
(Haefner et al., 2014)	<i>Quality value stream mapping</i>	Journal Article	2
(Kök & Fisher, 2007)	<i>Demand estimation and assortment optimization under substitution: Methodology and application</i>	Journal Article	2
(Six Sigma, 2025)	<i>Responsibility Assignment With RASCI Matrix in Lean Six Sigma Projects</i>	Webpage	2
(Stone, 2024)	<i>The Swim Lane Process: A streamlined alternative to RACI charts</i>	Webpage	2
(Von Rosing et al., 2014)	<i>Business process model and notation-BPMN</i>	Journal Article	2
Forecasting Indicators			
(Agbemadon et al., 2023)	<i>Overstock Prediction Using Machine Learning in Retail Industry</i>	Conference Paper	3
(Chen et al., 2006)	<i>Optimal pricing and inventory control policy in periodic-review systems with fixed ordering cost and lost sales</i>	Journal Article	3
(Fahimnia et al., 2024)	<i>Service-level anchoring in demand forecasting: The moderating impact of retail promotions and product perishability</i>	Journal Article	3
(Faber, 2015)	<i>Structuring warehouse management: Exploring the fit between warehouse characteristics and warehouse planning and control structure, and its effect on warehouse performance</i>	Thesis	3
(Francisca Silva, 2023)	<i>Understanding the Financial Impact Inventory Carrying Costs in Warehouses</i>	Journal Article	3
(Gustavsson et al., 2011)	<i>Global Food Losses and Food Waste-Extent, Causes and Prevention</i>	Technical Report	3
(Hill et al., 2007)	<i>A two-echelon inventory model with lost sales</i>	Journal Article	3
(Imam Raza, 2025)	<i>Retail Demand Forecasting</i>	Journal Article	3
(Jedidi et al., 1999)	<i>Managing advertising and promotion for long-run profitability</i>	Journal Article	3
(Kerkkänen et al., 2009)	<i>Demand forecasting errors in industrial context: Measurement and impacts</i>	Journal Article	3
(Lee et al., 1997)	<i>The Bullwhip Effect in Supply Chains</i>	Journal Article	3
(McColl et al., 2020)	<i>Estimating cannibalizing effects of sales promotions: The impact of price cuts and store type</i>	Journal Article	3
(Zipkin, 2006)	<i>On the Structure of Lost-Sales Inventory Models</i>	Technical Report	3
Costing Methods for Operational Indicators			
(Emmett, 2005)	<i>Excellence in Warehouse Management How to Minimise Costs and Maximise Value</i>	Book	3
(Speh, 2009)	<i>Understanding Warehouse Costs and Risks</i>	Journal Article	3
(Almeida & Cunha, 2017)	<i>The implementation of an Activity-Based Costing (ABC) system in a manufacturing company</i>	Journal Article	4
(Monti, 2009)	<i>A Warehouse Activity Based Costing Model</i>	Thesis	4
(Sofia, n.d.)	<i>Activity Based Costing in Warehousing: A Practical Example</i>	Thesis	4
Design and Implementation Considerations for Dashboards			
(Allio, 2006)	<i>Strategic dashboards: designing and deploying them to improve implementation</i>	Journal Article	5
(Bach et al., 2022)	<i>Dashboard Design Patterns</i>	Technical Report	5
(Orts, 2004)	<i>Dashboard Development and Deployment A Methodology for Success</i>	Corporate Report	5

2.2. Literature Review and Exploratory Field Study

This section presents the key literature identified through the review process outlined in Section 2.1. In addition, it incorporates insights gathered through the internal study, including discussions with experts, visits to the FC0, and reviews of internal documentation. The topics discussed are structured around the main thematic areas underpinning this research: process mapping, performance indicators, monetising operational impacts, and dashboard design principles. Together, these topics form the conceptual and methodological foundation for developing the article-level forecast impact framework.

This section presents the key literature identified through the review process outlined Section 2.1. The discussion is organised around the main thematic areas underpinning this study, including process

mapping methods, relevant performance indicators, approaches to monetising operational impacts, and dashboard design principles. Together, these topics provide the conceptual and methodological basis for developing the article-level trade-off framework.

2.2.1. Forecasting in the Online Supermarket Business

This review gathers a general understanding of promotional forecasting in online supermarkets, its challenges, common methods and the key factors that drive demand during promotions. This background was largely assembled during the preparatory literature research, to better understand the complexity and variation around this topic.

General Introduction to Demand Forecasting

Accurate demand forecasting is a critical component of supply chain management in the online supermarket industry. Forecasting helps businesses optimise inventory levels, minimise stock-outs, reduce waste, and ensure timely product availability for customers. Online supermarkets, rely heavily on data-driven forecasting methods to enhance operational efficiency. However, a set of several different indicators, poses unique challenges for accurate forecasting (Evers et al., 2018). Historical customer order data is a valuable known resource for forecasting demand, particularly for habitual products (Evers et al., 2018). However, such data is not always available or accurate, and a range of other factors further increase the complexity of forecasting in the online supermarket sector. These include high demand variability, shifting customer behaviour, product seasonality, and inconsistent availability of historical data (Heemskerk, 2024). In addition, forecasting becomes even more difficult during promotional periods due to the heightened volatility they introduce (Abolghasemi et al., 2019).

Notable Breakthroughs in Demand Forecasting

In the early stages of promotional forecasting Cooper et al. (1999) developed a 67-variable regression model using article-level data from past promotions, including factors like pricing, promotional events, and display levels. However, their model had notable limitations. The model struggled to accurately forecast when the nature of the promotion (such as display method or discount type) for an existing product changed. Additionally, it was unable to forecast demand for new products that had no prior promotional history. Koottatep and Eng (2006) stated earlier that traditional forecasting models often fail to capture these complexities, leading to either over-stocking or stock-outs. They proposed the use of more sophisticated approaches, including bias correction and hybrid models. These methods aim to improve the alignment of demand predictions with actual sales, leading to more efficient inventory management during promotional events.

Recently Heemskerk (2024) stated that predicting the effects of promotions is particularly difficult because consumer response to discounts can vary significantly, and there may be little to no historical data for certain products. This issue is compounded for new products or initial promotional events, where forecasting demand becomes even more uncertain (Van Donselaar et al., 2016). What Van Donselaar et al. (2016) also discovered that it's crucial to differentiate between routine (i.e. stable demand process) and non-routine products. Forecasting for routine categories is most effective when using data points solely from the category itself. In contrast, forecasting for non-routine categories improves when data from other categories are also incorporated.

This demonstrates that product type plays a significant role in the effectiveness of forecasting models. The use of intra-category data (within the same product category) and inter-category data (across different product categories) has proven to improve promotional forecasting accuracy. Inter-category effects refer to how changes in the price or promotion of one product can impact the sales of another product, either as substitutes or complements (Ma et al., 2016). For instance, research has shown that brand-switching within a category influences promotional response (Gupta, 1988), and complementary products from different categories, such as pasta sauce and pasta, can also affect sales (Calvo-Porrall & Lévy-Mangin, 2017).

Combining intra- and inter-category data can enhance forecasting, especially when products have both substitutive and complementary relationships. Bandyopadhyay (2009) demonstrated that using a multi-category model provides more accurate sales forecasts than using a single-category model. However,

identifying and measuring these relationships is difficult, particularly when dealing with large-scale, high-dimensional data, which complicates forecasting models (Ma et al., 2016).

2.2.2. Process Flow

After the general introduction, this section explores established methodologies for mapping and visualising business processes. The focus lies on identifying frameworks that help uncover bottlenecks, clarify team interactions, and improve communication between departments—factors that are particularly relevant in the context of forecasting and inventory management.

Before structuring a process within an organisation, it is essential to understand how process structures are typically developed. Bottlenecks in processes often originate from miscommunication between teams, which makes a clear overview of responsibilities and workflows highly valuable. Building a foundation of knowledge on proven methods to visualise and align cross-functional collaboration provides a strong basis for the process design later applied in this research.

Several studies (DeHoratius and Raman (2008), Kök and Fisher (2007)) highlight that over-stocking often results from misaligned incentives between forecasting teams and inventory managers. Forecasting models are typically designed to maximise sales accuracy, while operational teams aim to minimise logistical risk. In the context of online grocery retail, this tension may be even stronger, as demand forecasting must also consider challenges like limited warehouse space, perishable products and sharp unpredictable sales peaks caused by promotions. All kinds of factors that are not always addressed in traditional forecasting approaches.

Process Mapping Methods

To structure and visualise the promotional forecasting process within a retail company, various methods can be used to clarify stakeholder responsibilities and map tasks over time. The following subsection provides an overview of four established approaches: the RASCI matrix, Swimlane diagrams, Value Stream Mapping (VSM), and Business Process Model and Notation (BPMN).

The RASCI matrix, (Responsible, Accountable, Supportive, Consulted, and Informed) is a tool for clarifying each stakeholder's level of involvement in different steps of a process. By explicitly assigning roles, the RASCI model prevents confusion, ensures accountability, and supports coordination in cross-functional environments. It also enhances transparency by making stakeholder expectations visible from the outset, thereby helping to avoid delays caused by miscommunication or duplicated efforts. In the context of promotional forecasting, where input is often required from marketing, sales, supply chain, and finance teams, the RASCI matrix can be used to indicate the involvement of each function throughout the forecasting workflow (Six Sigma, 2025).

The Swimlane diagram offers a time-based visual representation of the process flow (Stone, 2024). Each horizontal lane represents a stakeholder or department, while activities are plotted sequentially to illustrate interactions and dependencies. This format provides visibility into task ownership and interdepartmental coordination. According to Stone (2024), Swimlane diagrams can serve as an intuitive way to illustrate timing and interaction between stakeholders, making them suitable for processes where cross-functional alignment is important, such as promotional forecasting.

Value Stream Mapping (VSM) (Haefner et al., 2014) is another approach, originating from lean management, used to analyse and optimise the flow of materials and information through a process. VSM typically maps both the current and future states of a process, highlighting the sequence of activities, information exchanges, and material flows. It helps identify inefficiencies, waste, and potential opportunities for operational improvement.

Finally, Business Process Model and Notation (BPMN) (Von Rosing et al., 2014) is a standardised modelling approach for representing complex business processes. BPMN provides a graphical notation that captures the detailed logic of workflows, including events, decision rules, and process interactions. Its flexibility allows for both high-level overviews and detailed representations, making it suitable for documenting and communicating business processes in a structured and consistent manner.

After all four methods were examined in greater detail and are summarised in Table 2.3.

Table 2.3: Overview of Process Mapping Approaches Examined in Greater Detail

Method	Description
RASCI Matrix	o A responsibility-assignment framework that defines stakeholder roles as Responsible, Accountable, Supportive, Consulted, or Informed. It clarifies task ownership, improves accountability, and prevents overlap in multi-departmental processes.
Swimlane Diagram	o A visual flowchart in which each lane represents a stakeholder or department. Activities are placed sequentially to show how tasks and responsibilities move across teams over time, highlighting interdependencies and bottlenecks.
Value Stream Mapping (VSM)	o A lean management tool used to analyse the flow of materials and information through a process. It distinguishes between value-adding and non-value-adding activities, helping identify waste and opportunities for process improvement.
Business Process Model and Notation (BPMN)	o A standardised graphical notation for modelling complex workflows. It captures detailed process logic, including events, decision rules, and conditional flows, allowing for consistent documentation and communication of business processes.

Together, these approaches illustrate the range of available methodologies for mapping business processes and provide the theoretical background for the selection discussed in Section 2.4 and further elaborated in Section 3.2.

2.2.3. Forecasting Indicators

This part of the literature research aims to gather relevant performance indicators to support the creation of a longlist of potential indicators influenced by forecasting errors.

To make accurate forecasts, it is essential to understand the consequences of both over-forecasting and under-forecasting. The impact of these forecasting errors varies per article and depends on multiple factors, making it crucial to gain a deeper understanding of their specific effects. This part of the literature research informed the creation of the longlist provided in Appendix E.2.

Impact on Forecasting Errors

When forecasting demand, online supermarkets aim for the highest possible accuracy. However, absolute precision is not achievable, which naturally introduces the risk of forecasting too much or too little. As a result, analysts frequently face a critical decision: whether to intentionally over-forecast or under-forecast relative to the model's predicted demand in order to maximise profitability.

Over-forecasting may lead to excessive stock levels, resulting in increased storage costs, potential waste (especially with perishable goods), and inefficient use of resources (Fahimnia et al., 2024; Imam Raza, 2025). On the other hand, under-forecasting can lead to stock-outs, which result in missed sales, decreased customer satisfaction, and potential damage to the supermarket's reputation. Research indicates that balancing these outcomes is essential for maintaining operational efficiency, customer trust and maximising profit (Imam Raza, 2025). Moreover, the relative weight of each factor can differ based on product type (e.g., perishables versus non-perishables), customer behaviour, and prevailing market conditions. Additionally, the trade-off becomes even more complex when accounting for elements such as promotional activity, seasonal fluctuations, and external market disruptions (Lee et al., 1997).

Under-forecasting Impact

The previous subsections already touched upon some consequences of under-forecasting. This subsection will explore these impacts in greater depth, providing a more detailed explanation of how under-forecasting affects retail operations.

Under-forecasting in the retail sector can lead to substantial operational and financial consequences, primarily driven by lost sales, poor service levels, and ineffective capacity utilisation. In competitive retail settings, unmet demand typically drives customers to rival retailers, meaning that stock-outs result in permanently lost sales rather than delayed purchases through follow-up deliveries (Hill et al., 2007). These lost sales not only reduce revenue directly but also damage customer loyalty and brand perception (Chen et al., 2006).

The cost of under-forecasting becomes even higher because of how lost-sales inventory models work. These models show that when demand varies a lot, costs go up (especially when missed sales can't be

made up later) (Zipkin, 2006). In addition, under-forecasting makes it more difficult for retailers to manage both pricing strategies and inventory levels effectively. When too little stock is available to meet actual demand, retailers may be forced to cancel or reduce planned promotional campaigns because the product is not available in sufficient quantities. This not only affects marketing plans but also limits the ability to respond flexibly to customer demand through price changes. For example, retailers may miss the opportunity to raise prices during high demand or lower them to clear stock, simply because the products are not on the shelves (Chen et al., 2006).

Empirical evidence also demonstrates that under-forecasting contributes to planning inefficiencies, such as excess re-planning and schedule instability (Kerkkänen et al., 2009). These inefficiencies can cause underutilisation of capacity or misallocation of resources, resulting in elevated operational costs and reduced responsiveness.

Over-forecasting Impact

Earlier subsections briefly discussed the effects of over-forecasting. This section now takes a closer look, offering a more detailed examination of its specific consequences for retail operations.

In retail supply chains, particularly in the context of online supermarkets, over-stocking represents a critical operational risk that can lead to both financial losses and logistical inefficiencies (indirectly leading to financial losses). While the primary commercial intention of promotional forecasting is to avoid stock-outs, and thus a loss of revenue, while the consequences of overestimating demand can be more significant.

Over-stocking results in increased costs, think about storage space, labour, insurance, and depreciation. These factors are especially acute in categories with perishable goods such as dairy, fresh fruits, vegetables or meats (Francisca Silva, 2023). According to Agbemadon et al. (2023), retailers often resort to anti-waste promotions to liquidate excess inventory approaching expiration. Although these discounts help reduce waste in the short term, they lower profit margins and can change consumer price expectations regarding pricing, making them more likely to expect discounts in the future.

In addition to cost implications, over-stocking can severely disrupt warehouse operations. Online supermarkets, which typically operate under high throughput and limited storage space, are especially vulnerable. Congested warehouses make it harder to pick and pack orders quickly, which reduces efficiency and increases the chance of mistakes (Faber, 2015). This problem is more common during promotions, when sales go up more than usual and are harder to predict compared to normal periods. The risk of over-stocking increases during promotions because demand is harder to predict and promotions can disrupt usual buying patterns. Research shows that promotions often cause temporary sales spikes, followed by a drop in demand afterward. This makes it hard to know whether people are actually buying more, or just shifting when they buy. If forecasting models only look at past promotions and ignore things like competitor actions, weather, or social trends, they can easily overestimate or underestimate demand (Jedidi et al., 1999). This phenomenon where customers purchase products during a promotion instead of at a later time, resulting in no real increase in total demand, is known as cannibalisation (McColl et al., 2020). More recently, the discussion around over-stocking has extended beyond operational metrics to include sustainability and corporate responsibility. Unsold inventory, especially perishable goods, contributes to food waste, which has environmental and ethical implications. According to Gustavsson et al. (2011), roughly one-third of all food produced globally is wasted, with retail-level over-stocking being a major contributor. Consumers are increasingly aware of such practices, and over-stock-induced waste can undermine brand trust and corporate reputation, particularly in markets with strong sustainability norms.

In conclusion bad forecasting, in this case over-forecasting has a lot of serious consequences and monetising the explicit damage is very complex.

2.2.4. Costing Methods for Operational Indicators

This section presents the literature background on the methods used to calculate the costs of operational impacts, complemented by field study insights into Picnic's operational processes. These insights were essential for understanding how activities are carried out within Picnic's fulfilment centres. The section first provides a brief explanation of Picnic's fulfilment centre operations, focusing on handling and storage activities, followed by a discussion of the two selected costing methods: Activity-Based

Costing (ABC), applied to calculate handling costs, and Activity-Based Space Costing (ABS), applied to calculate storage costs. A third indicator used in this research combines both handling and space costs to determine the Maximum Economical Retention Time (MERT)—an indicator that identifies the break-even point at which continued storage becomes more expensive than discarding stock. This indicator is further discussed in Chapter 4. The section begins, however, with the findings obtained from the field study.

Field Study on Practical Operations

This subsection examines the operational processes within Picnic France, providing the foundation required for determining the appropriate cost calculations in subsequent chapters. The analysis focuses exclusively on the operational flow of products that were overestimated in the forecast, as quantifying the costs associated with such cases constitutes a key objective of this study. The operational costs arising from underestimation are not considered, as Picnic is unable to place backorders with its supplier, ITM. The aim of this subsection is therefore to specify where handling activities take place within the fulfilment centre (FC) and where space costs are incurred. The discussion begins with the handling process.

All handling operations occur within the ambient zone of the fulfilment centre, which forms the operational scope of this research. This zone is the largest and therefore most significant area of the FC. Due to time constraints of the study, the chilled and frozen zones were not included in the analysis of operational impacts. Within the ambient zone, there are three main locations where handling activities for over-forecasted products occur: the Inbound area, the Replenishment area, and the Buffer zone. Products first arrive in the Inbound area, where they are received, sorted, and transported to the Replenishment area. Handling also occurs here when these products need to be processed. Subsequently, products are moved to the Buffer zone, where additional handling takes place as they are transferred in and out of storage before eventually returning to the Replenishment area. This sequence highlights that multiple handling activities are associated with over-forecasted products. A detailed explanation of each specific step, and how it links to the developed cost calculations, is provided in Section 5.2.




Figure 2.2: Overview of FC0, identifying the scope of the research: the ambient zone.

Regarding space costs, these are incurred only within the Buffer zone. The detailed calculation of these costs is provided in Section 5.2. Conceptually, space costs represent the expenses associated with

products occupying physical space that could otherwise be used for other purposes. During a promotion, additional space is allocated in the Replenishment area for products expected to sell in high volumes. However, when forecasts are inaccurate and excess products remain unsold, this space must be cleared for other items. The remaining stock is therefore moved to the Buffer zone, where it is stored until it can be sold at a later stage after the promotional campaign. Because the Buffer zone has limited capacity and serves multiple purposes, occupying it with over-forecasted products generates unnecessary costs, as this valuable space could otherwise be used more efficiently. In summary, space costs are incurred only when products remain stored in the Buffer zone while awaiting later sale.

Operational Costs: Handling and Space Costs

Before addressing the methods used to determine handling and space costs, a brief introduction to operational costs is provided. To understand the operational costs associated with over-stocking, it is essential to examine two key components: handling costs and space (or storage) costs. These costs represent, respectively, the expense of moving goods and the cost of holding goods in a warehouse. Together, they form a major part of the total warehousing costs and are central to the trade-off between over- and under-forecasting (Speh, 2009).

Handling costs include all activities related to the physical movement of goods within a warehouse. These typically involve receiving, putting away, picking, repackaging, and loading (Speh, 2009). Labour is the primary driver of handling costs, accounting for up to 60% of total warehouse costs in a standard operation (Emmett, 2005). In addition to labour, handling costs include the operation and depreciation of equipment such as forklifts, fuel or electricity used, and costs of consumables like pallets or wrapping materials (Speh, 2009).

To calculate a representative handling cost, a building block approach can be used: first, total direct handling costs are identified. Then, a share of operations administration (e.g. line supervision, IT systems) and general overhead (e.g. HR, finance) is allocated to these activities (Speh, 2009). This results in a "burdened handling cost," which is then divided by the number of labour hours or units handled to derive a cost per unit or per hour. In this research only the direct cost will be taken into account, as the overhead cost are not influenced by a relative small impact due to error forecasts.

Storage costs refer to the cost of keeping goods at rest. These costs are largely time-based and primarily depend on warehouse space occupancy. Rent, depreciation, property taxes, and insurance on the facility are core fixed costs, while variable costs may include lighting and security (Emmett, 2005). A typical estimate suggests space accounts for around 25% of total warehouse costs (Emmett, 2005).

Importantly, storage costs are sensitive to inventory turnover. Items that are stored for longer periods accumulate higher monthly charges. A product turning 24 times per year is significantly cheaper to store than one turning only 6 times, making inventory turnover a critical input in cost estimation (Speh, 2009). However, it is very reasonable to argue that holding costs are not actually additional costs, because the space would not have been used otherwise. This is a very interesting perspective to consider.

In over-stocking scenarios, products occupy warehouse space longer than needed and pass through the warehouse more often than necessary, increasing both storage and handling costs. These increases are not always proportional, as many warehouse costs are fixed and only become problematic when space and labour are over-committed (Emmett, 2005; Speh, 2009).

For decision-making, both costs must be calculated precisely, accurate handling and space cost calculations in the assessment of over- or under-stocking consequences is essential. It enables a better understanding of the operational burden of forecast errors and supports more informed trade-off decisions.

Activity-Based Costing and Activity-Based Space Costing

The chosen methods for calculating both handling and space costs are very similar and rely on the same principle: determining the duration of an activity and assigning a price per unit of time to monetise its effect. ABC is applied to handling costs, while ABS is applied to space costs.

In this research, the approaches are based on the academic methodologies described by Almeida and Cunha (2017), Monti (2009), and Sofia (n.d.). While this activity-based approach can be applied to an

entire set of activities, in this study it is applied separately to handling and space costs, each using the same underlying principles but with different resources and cost drivers.

Almeida and Cunha (2017) present a clear implementation framework, shown in Figure B.1, which also distinguishes between types of resources. They divide the Activity-Based method into two phases.

In the first phase, all activities within the production process and the resources they consume are identified. Resource costs are then assigned to activities using relevant cost drivers, such as space usage, energy, maintenance, insurance, labour, and equipment depreciation (Almeida & Cunha, 2017) (see the first column of Figure B.1). For the handling cost calculation, only labour is considered, as the remaining resources fall outside the scope of this indicator. The so-called resource driver for labour is the time spent on each activity (see the second column of Figure B.1). For the space cost calculation, the same approach is applied, but the relevant resource is the industrial building—for Picnic, this corresponds to their fulfilment centre (FC)—with the resource driver being the area occupied by each activity. This occupied area is the buffer zone—the only area in the FC where incorrectly forecasted articles are kept—as explained in Chapter 5

In the second phase, Almeida and Cunha (2017) describe assigning activity costs to products using activity cost drivers that reflect actual consumption, followed by calculating the total product cost by aggregating all relevant activity costs. For the handling cost indicator, this involves identifying additional handling steps resulting from forecasting errors and the corresponding extra handling time, as explained in detail in Chapter 5. For the space cost indicator, the relevant factor is the time that articles are stored in the buffer zone. For both indicators, only the main activities are included in the calculations.

Based on this setup for allocating additional handling and space costs, and using the ABC and ABS methodologies described by Monti (2009) and Sofia (n.d.), Equation (5.9) and Equation (5.13) are derived.

Space Costs per m³ per Day

After determining the amount of time a product is stored using the ABS-method, it is important to assign a realistic price to the cost of space per unit of time. Therefore, it is essential to establish which elements of warehouse expenses can reasonably be attributed to storage space and which should be excluded. A review of the literature provides relevant guidance on this matter.

Nordman (n.d.) and Speh (2009) both outline similar frameworks for analysing warehousing costs, highlighting that these can generally be divided into five categories: staff, building, building services, equipment, and information technology. Out of these, staff, handling equipment, and information technology are typically excluded when calculating pure space-related costs, as they are more directly linked to operational and handling activities rather than the physical occupation of space.

Following this approach, and in line with the overview provided by Richards (2018) (see Figure 2.3), space costs are defined as including the following elements:

- Rent or lease costs: rates, building insurance
- Utilities: electricity, gas, water
- Depreciation, repair, and maintenance of building and buffer racks
- Cleaning, security, and depreciation of non-handling equipment




Figure 2.3: Simple warehouse cost-tree (from Richards (2018)).

Finally, in Section 5.2.2, all costs listed above are determined, with the exception of depreciation, repair, and maintenance costs of the building and buffer racks, as the buffer zones at Picnic consist only of pallets and inexpensive racks, which are not considered significant enough to include.

2.2.5. Design and Implementation Considerations for Dashboards

A dashboard is an effective way to present trade-off analyses, making complex data more accessible and actionable for decision-makers. To understand its full value, it is important to explore how dashboards are used and why their design and implementation matter. This section explores these aspects and identifies key requirements to consider upfront when creating a dashboard.

Value of Dashboards

Dashboards are powerful tools that provide decision-makers with a centralised view of Key Performance Indicators (KPIs), enabling strategic guidance and timely actions (Orts, 2004). However, the development and deployment of an effective dashboard is far from trivial. It requires careful planning, iterative design, and active stakeholder involvement throughout its lifecycle.

To avoid the common pitfall of dashboards failing to align with actual strategy, organisations must begin with a clear planning phase, defining project scope, KPIs, data sources, and executive sponsorship (Orts, 2004). Dashboards should not merely reflect available data, but instead translate strategic goals into measurable indicators. Too often, dashboards focus on operational efficiency metrics, overlooking key strategic drivers such as innovation or market expansion (Allio, 2006).

An effective dashboard also depends on inclusive design. Involving cross-functional management teams in metric selection enhances accountability and ensures relevance. A strategic dashboard must be simple but not simplistic: a curated set of 10–15 indicators is ideal to avoid data overload while maintaining insight (Allio, 2006). The layout, data abstraction, and interactivity must be tailored to the audience, especially when dashboards serve different roles: strategic, tactical, or operational (Bach et al., 2022).

Finally, dashboards should be seen as evolving tools. As priorities shift, the metrics and visualisations must be reviewed and refined regularly. Adopting recognised design patterns and including iterative feedback cycles from users will increase adoption and long-term effectiveness (Bach et al., 2022).

2.3. Discussion on the Literature Research

The goal of the literature review was to build a solid theoretical and practical basis for the study. This goal was largely achieved. As each topic within the literature review contributes to a different part of the research, it is useful to walk through each section and discuss what was completed, what was missing, and the choices made. Much of the literature consulted is not explicitly outlined in this thesis because it was not directly relevant to the final report. However, methods and background information that were not ultimately included in this thesis still added value to the overall research process. Each subsection of the literature review is discussed chronologically, starting with 'Forecasting in the Online Supermarket Business'.

Forecasting in the Online Supermarket Business

A significant share of the literature review time was devoted to building a general background in forecasting. Much of this work was completed during the preparatory literature phase, before finalising the research topic. There is relatively little to critique in this section, as extensive, high-quality sources provided deep insights into the field. This report includes the evolution of forecasting because it is relevant to understanding how the practice develops over time and why it remains a challenging and important topic across industries.

Two elements were largely missing from the literature:

1. Specific approaches to promotional forecasting for products without historical data.
2. Forecasting approaches based on risk minimisation rather than accuracy.

However, the absence of the second topic reinforced the academic relevance of this research, as the gap addressed in this study concerns the incorporation of risk impacts based on forecasting errors across multiple factors.

Process Flow

Sufficient literature was available on potential methods for mapping process flows. To decide which approach best suited this research, multiple options were reviewed. Many methods were quickly ruled out after reading abstracts, as they were unsuitable for the research context. Four methods were explored in more detail: Swimlane, RASCI, VSM, and BPMN.

The Swimlane approach, combined with the more detailed RASCI method, was ultimately selected because the two techniques complement one another. Swimlane diagrams provide a clear, time-based visualisation of the process flow, illustrating when and how different stakeholders interact. However, they offer limited depth regarding the specific roles and responsibilities of each participant. The RASCI matrix, by contrast, explicitly defines stakeholder involvement in each process step, clarifying who is responsible, accountable, consulted, or informed. By integrating both methods, the process can be represented both visually and structurally—linking the sequence of activities over time (Swimlane) with precise role allocation (RASCI). This combination provides a comprehensive framework for designing, evaluating, and communicating the promotional forecasting process, ensuring clarity regarding both responsibilities and execution over time.

Furthermore, the decision to integrate the Swimlane and RASCI methods is itself a novel contribution to the field of process mapping. While existing literature often treats these techniques separately, this research demonstrates the synergistic value of their combination. This methodological innovation provides a more comprehensive and robust framework for process analysis than either method could achieve alone, highlighting a new avenue for research and application in process design.

The literature in this on process flows was adequate, and the reasoning behind selecting Swimlane and RASCI was well-founded, resulting in strong deliverables (Figure 3.2, Figure C.2 and Figure C.3).

Forecasting Indicators

This literature research was both sufficient and highly productive, providing the primary input for the longlist of indicators (Table E.2). In addition to identifying many potential indicators, reading about the different types of impacts related to forecasting errors provided inspiration and broadened the understanding of the topic.

Costing Methods for Operational Indicators

This part of the study consisted partly of a literature review and partly of an internal process review. After developing a full understanding of the operational processes in the FC, the challenge arose in aligning this understanding with academically sound methods capable of producing reliable measures of handling and space impacts.

The methods selected for quantifying these operational impact indicators needed to meet two key requirements:

- Address storage or space costs exclusively.
- Capture the operational consequences of forecasting errors within Picnic.

Despite these requirements, several potentially relevant methods were identified, with ABC and ABS ultimately selected. This choice was well supported by Picnic's available resources, which included time-related data from the Data Warehouse (DWH), product dimensions per SKU, and cost information related to warehouse rental, utilities, security, and cleaning. As all the necessary data for applying the ABC and ABS methods were available at Picnic, these approaches were quickly identified as the most suitable.

Design and Implementation Considerations for Dashboards

This literature review topic began with less clarity than the others, as dashboards vary greatly between contexts and limited academic literature addresses dashboard setup and use. The approach here was intentionally exploratory, aiming to learn as much as possible about dashboard design: what is needed, which requirements to define upfront, and other practical considerations. The information gathered, although less prescriptive than in other sections, provided valuable insights into academic perspectives on dashboard use.

Even after the literature review, some knowledge gaps remained. These were to be addressed through internal discussions with stakeholders and end-users, which helped ensure the development of a tailor-made decision-support dashboard for Picnic.

2.4. Main Findings Chapter 2

Overall, the literature review provided a strong foundation for the research, guiding the methodological choices made in each phase and shaping the conceptual development of the artefact.

Most importantly, the literature study played a crucial role in defining the direction of this thesis by revealing an interesting knowledge gap: the inclusion of the cost component of inventory management and handling resulting from forecasting errors was scarcely addressed in existing research. This limited understanding of the real impact of over- and under-estimation strongly influenced the focus of this study. A deeper investigation of both academic literature and internal company processes showed that these impacts vary significantly across product types. This finding highlighted that when incorporating potential cost impacts as risk factors in forecasting decisions, such considerations cannot be generalised but must instead be approached at the product level. Together, these insights identified a meaningful gap in the literature—one that this study takes a first step towards addressing.

As discussed, this gap helped define both the originality and the practical value of the research. By combining established theoretical frameworks with practical insights, the literature study ensured that the design process was not only evidence-based but also closely aligned with the real challenges of Picnic's promotional forecasting. This integration of theory and real-world application enhanced the artefact's credibility and strengthened the quality of the final deliverable produced in this study.

3

Mapping the Promotional Forecasting Process

This chapter provides an in-depth analysis of Picnic’s promotional forecasting process, focusing on the roles and responsibilities of different teams and how they interact throughout the process. With the goal to answer the first research question:

1. *How is the current promotional forecasting process structured at Picnic, in terms of workflows, team responsibilities, and interdepartmental coordination?*

The chapter begins with the rationale for addressing this research question in Section 3.1, followed by the methodology used in Section 3.2. Appendix D.1 and Appendix D.2 present the identified bottlenecks and corresponding recommendations. The chapter ends with a clear overview stating the main findings in Section 3.4.

3.1. Rationale for Mapping the Promotional Forecasting Process

Understanding the structure of the promotional forecasting process is critical for several reasons. First, clearly mapping out complex workflows is essential for gaining a comprehensive understanding of how the process operates within Picnic. By visualising the current process, it becomes possible to pinpoint how each step is executed, how teams interact, and why certain decisions are made. This clarity will be crucial during the analysis in later stages of the research, as it helps to contextualise the reasons behind specific actions and outcomes. Understanding the workflows also allows for an assessment of the efficiency of the current system and the identification of opportunities for optimisation.

Additionally, visualising the process helps to identify potential inefficiencies or bottlenecks that may be hindering the flow of operations. These could include delays, unclear roles, or communication breakdowns between teams. Recognising such issues at an early stage is critical to addressing them and improving the overall forecasting process.

The creation of a dashboard, which is a key deliverable in this research, further underscores the importance of mapping the process in detail. For the dashboard to be effective, it must be integrated seamlessly into the existing workflow. By understanding how the current process operates, the dashboard can be designed to address critical points in the forecasting process, enhance team coordination, and support real-time decision-making. This integration will be explored in greater detail as the research progresses, but it is essential that the process structure be well understood to ensure successful implementation.

In summary, the rationale for analysing the promotional forecasting process is based on four key points:

- Understanding complex workflows to ensure a comprehensive view of the current process.
- Supporting analysis in later stages of the research by providing context for why actions and outcomes occur.
- Identifying inefficiencies or bottlenecks that may hinder process efficiency and hinder decision-making.
- Ensuring the successful integration of the dashboard into the existing workflow, which will be a core deliverable of this research.

3.2. Mapping Methodology

This section outlines the rationale for the visualisation methods chosen to analyse the promotional forecasting process. A detailed description of the available approaches can be found in Section 2.2.2; the section focuses on the justification of the chosen methods and its application within this research.

3.2.1. Choice of Visualisation Methods

The Swimlane Diagram (Stone, 2024) and the RASCI Process Flow (Six Sigma, 2025) were selected because together they provide a comprehensive and complementary perspective on the promotional forecasting process.

As described in the discussion of Section 2.3, the Swimlane Diagram provides an intuitive, time-based visual overview of activities and stakeholder interactions, enabling the quick identification of bottlenecks and dependencies. This is complemented by the RASCI Framework, which offers a granular perspective by explicitly defining roles and responsibilities (Responsible, Accountable, Consulted, Informed) for each process step. Together, this combination ensures that the process is represented both visually and structurally, linking the sequence of activities over time with precise ownership.

Combining both methods offers a dual-layered understanding of the process: the Swimlane Diagram, shown in Figure 3.2, captures the temporal and interactive dimension of cross-team collaboration, while the RASCI Framework, provided in Appendix C.2, adds role-based precision at the task level. This combination allows for both visual accessibility and analytical depth—making the mapped process simultaneously easy to interpret and robust enough for process evaluation and improvement. The final version of the Swimlane Diagram, which incorporates the use of the new dashboard, is given in Figure C.4.




Figure 3.1: Approach steps process flow.

3.2.2. Approach

To develop and refine the Swimlane and RASCI process flow, a chronological and collaborative approach was adopted. Figure 3.1 provides a clear overview of the steps taken to create the Swimlane Diagram (Figure 3.2) and the RASCI Framework. Step 4 and step 6 in Figure 3.1 were executed as outlined for both methods in Section 2.2.2. The remaining steps are identical for both methods. The process began with a review of internal documentation related to the promotional forecasting process, which provided a basic understanding of its current structure. This was complemented by informal conversations with colleagues directly involved in the process. These discussions helped identify key stakeholders and understand the specific responsibilities of each department. Based on the information gathered, a preliminary version of the process flow was mapped out, serving as a basis for further refinement. This initial draft was then shared with the main stakeholder, an analyst from the promo team, whose end-to-end involvement and regular interaction with other departments offered valuable insights. With feedback from this stakeholder, the process flow was revised in greater detail to align more closely with actual practices. In the subsequent phase, the refined version was validated with team leads from each stakeholder department. Their feedback helped correct inaccuracies, incorporate departmental perspectives, and suggest improvements. Finally, the completed version was shared via an online collaborative platform to allow additional review and input from other team members, ensuring that the process flow is both accurate and widely supported across the organisation.

Justification for the Robustness of the Approach

The robustness of the approach lies in its combination of qualitative data collection, through interviews and informal conversations, and document-based analysis. By engaging with multiple stakeholders at different stages of the process, a comprehensive understanding of the workflow was developed from various departmental perspectives. The iterative cycle of creating, validating, and refining the process flow ensures that it closely aligns with the actual practices within the forecasting process. Furthermore, the use of an online collaborative platform for the final review enhances transparency and allows for continuous improvement based on input from relevant stakeholders.

One limitation is that the external stakeholder, ITM, was not consulted directly. Although this may slightly limit the completeness of the overview, internal stakeholders consistently validated the supplier's role in the process. Given the time constraints of this research and considering that the visualisation is primarily intended for internal use, it was deemed more efficient to rely on Picnic teams' experiences and perceptions. Therefore, the practical relevance and robustness of the final process flow remain intact for internal purposes, despite the absence of direct supplier input. This multi-step internal validation approach ensures the methodology remains accurate, comprehensive, and actionable.

3.3. Final Result of the Promotional Forecasting Process

This section explains the final result of the complete mapped process flow of promotional forecasting within Picnic, which is visualised in Figure 3.2. For more detailed insights into the specific level of involvement of each stakeholder per process step, reference can be made to the RASCI matrices presented in Appendix C.2.




Figure 3.2: Swimlane diagram of complete process.

3.3.1. Elaboration on the Mapped Process

This subsection provides a concise explanation of all the steps that form part of the promotional forecasting process. The process is divided into three stages: from ten weeks before the start of the promotion, the period between nine and seven weeks before the promotion, and the final phase from three weeks prior to the start until the beginning of the promotion itself.

Ten Weeks before Promotion

The process begins when ITM delivers a list containing hundreds of products, including the associated promotion mechanisms. From this list, the category managers select which products Picnic will actually include in the upcoming promotions (a detailed explanation of this process is provided in Section 1.1.2). In rare cases, a product may not conform to standard dimensions, in which case the Fulfilment team is consulted to confirm whether the product can feasibly be handled. Once the final product list is compiled, the analyst from the Promotion Team prepares the corresponding forecast. The output is subsequently reviewed by both the category managers and the promotion analyst. When the final quantities are agreed upon, this information is shared with the Fulfilment team and the supplier, ITM.

Nine to Seven Weeks before Promotion

During the following two-week period, ITM performs a final verification of the proposed quantities. If any forecasts appear highly irregular, ITM raises the issue with the promotion analyst. The promotion analyst then revises the quantities where necessary and confirms the definitive numbers to ITM and within Picnic.

Three Weeks before Promotion to Start of Promotion

Three weeks prior to the start of the promotion, the trading associate from the Promotion Team selects the most strategically important promotions—this process is referred to as highlighting. Simultaneously, ITM provides the final confirmation on whether all planned promotions will proceed, as deviations can still occur due to product shortages or logistical issues at the supplier's end. Once ITM confirms the final selection, the definitive overview of highlighted products is shared with the Store Campaign Team.

A discussion then follows, after which the Store Campaign Team begins configuring the online store and preparing the promotional email campaign to ensure readiness before the promotion launch. Meanwhile, the trading associate may add additional products to the campaign, which are then incorporated into the promotional page in the Picnic app. Finally, on the morning of the promotion itself, the promotional email is sent to customers, marking the official start of the campaign.

3.4. Main Findings Chapter 3

The mapping of the promotional forecasting process provided several key insights into Picnic's operational structure and cross-team collaboration. First, the rationale for mapping demonstrated that a clear visualisation of workflows is essential for understanding interdependencies, identifying bottlenecks, and supporting the later integration of the dashboard into existing forecasting activities. Second, the chosen combination of the Swimlane Diagram and the RASCI Framework proved to be an effective methodological approach, offering both a temporal and role-based perspective that ensured analytical depth and interpretability. The iterative mapping approach, involving multiple validation rounds with internal stakeholders, further strengthened the robustness and accuracy of the final process flow. The completed process map demonstrated that the promotional forecasting workflow is structured but strongly interdependent, comprising three main stages that range from product selection to promotional execution. This finding underscores the need for effective coordination between ITM, the Promotion Team, and related departments. In addition to these findings relevant to the research, the mapping exercise also uncovered several practical inefficiencies and improvement opportunities within Picnic's workflow. These operational insights, while outside the academic scope of this thesis, are elaborated in Appendix D and may serve as a valuable foundation for future process optimisation initiatives within Picnic France.

III

Design and Development

4

Selection and Justification of Indicators

In order to evaluate the consequences of promotional forecasting errors and inform targeted improvements, it is essential to identify the most relevant and actionable performance indicators. The selected indicators must offer meaningful insights while remaining feasible to measure within the constraints of data availability, analytical complexity, and the limited time-frame of this research. In short, their usefulness must justify the effort required for implementation. This leads to the following sub-question:

2. *Which key performance indicators should be used to assess the impacts of over- and under-stocking caused by promotional forecasting errors at Picnic and what insights does Picnic aim to derive from each indicator?*

Figure 4.1 provides a clear visualisation of all indicators considered in this research, distinguishing between those selected and those excluded. The selection process is outlined in Section 4.1, covering the identification process (Section 4.1.1), the Relevance–Feasibility matrix approach (Section 4.1.3), and the stakeholder validation strategy (Section 4.1.4). Figure 4.2 presents the results of the Relevance Feasibility Matrix (RF-Matrix), followed by the justification of the included and excluded indicators in Section 4.2. After all a clear table is provided in Section 4.2.4, showing the approach, goals and assumptions related to each indicator. Finally, Section 4.3 concludes the chapter by reflecting on limitations and opportunities for future refinement.

4.1. Methodology for Indicator Selection

This section outlines the structured methodology used to identify and evaluate the key operational and performance indicators relevant to assessing promotional forecasting errors at Picnic. The process begins with a long-list selection, as explained in Section 4.1.1, followed by the creation of a RF-Matrix in Section 4.1.3, and concludes with the validation approach, detailed in Section 4.1.4.

4.1.1. Long-list Identification and First Selection Process

The initial step in the indicator identification process involved an extensive exploration and collection of potential metrics. A thorough literature review was conducted to identify theoretically relevant indicators that commonly reflect the impacts of promotional forecasting errors. Subsequently, internal documentation provided insights into indicators already tracked or considered relevant within Picnic’s operational context. Additionally, informal conversations and an expert validation meeting with internal stakeholders, specifically the Inbound lead, Customer Team Lead and Promo Team lead, were carried out to gather expert insights and practical suggestions. This multi-source approach ensured comprehensive coverage of potentially relevant metrics. An overview of the longlist of all the initially identified indicators is clearly presented in Appendix E.2.

From this longlist, only the most plausible candidates were carried forward into a more in-depth selection process (shown in Section 4.1.3). Indicators deemed clearly irrelevant or infeasible were excluded based on a preliminary assessment. This assessment involved discussions with the Promo Lead at Picnic to evaluate whether each KPI could be a viable candidate, excluding considerations related to time or cost constraints. The reasoning for the excluded indicators, at this early stage, is documented in Appendix E.2.

Labelling the Chosen Indicators

Following the initial indicator identification, a categorisation step was performed, grouping the collected indicators into three main categories: financial, operational, and reputational. Financial indicators encompass direct monetary impacts resulting from forecasting inaccuracies, such as margin loss or additional costs. Operational indicators focus on logistical and supply chain-related effects, such as handling inefficiencies or space costs. Lastly, reputational indicators address impacts on customer satisfaction, brand perception, and overall customer experience due to forecasting errors.

This categorisation was strategically chosen to clarify the distinct types of consequences resulting from forecasting inaccuracies. It facilitates targeted analysis and allows for clearer communication with stakeholders who may have differing priorities or areas of responsibility. Moreover, clearly distinguishing between financial, operational, and reputational effects simplifies decision-making processes, as each indicator category aligns directly with the strategic objectives and functional domains within Picnic. The chosen indicators, their respective categories, and the scope of this study, are displayed in Figure 4.1. The scope of this research was determined based on a validation process with stakeholders, which is further explained in Section 4.1.3. The next section will briefly discuss all the indicators that were widely considered for this study.




Figure 4.1: Research indicator scope.

4.1.2. Potential Indicators in Brief

To ensure clarity, each indicator in this section is presented with a concise explanation of what it measures, how it is calculated, and why it is relevant for improving Picnic’s promotional forecasting. All calculations are based on historical data, which implies that articles without instances of overstock or under-stock will not generate any cost impact for the corresponding indicator in that specific campaign. The first indicators that will be discussed is the lost margin.

1. Lost Margin quantifies the unrealised profit resulting from product unavailability in the online store. When a customer attempts to add a product to their basket but finds it unavailable, the potential

sale is lost, despite the demonstrated purchase intent. The Lost Margin is calculated as:

$$LM = UC \cdot C \cdot CR \cdot GM \quad (4.1)$$

where

- LM = Lost Margin [eur]
- UC = Number of unavailable customers [count]
- C = Commonality factor [units/click]
- CR = Conversion ratio [%]
- GM = Article promo gross margin per item [eur/item]

Commonality Factor: average number of units of a product purchased per delivery.

Conversion Ratio: the average probability that a product added to the online shopping basket results in an actual purchase.

Goal: To quantify the financial loss from under-forecasted promotions and use that insight to improve forecast accuracy. This indicator helps Picnic prioritise product categories or promotions where the cost of underestimating demand is most severe, thereby informing both forecasting parameters (1) and campaign design (2) decisions.

Product Owner: Promo Lead

2. Waste Cost quantifies the financial loss of unsold products due to expiry or slow sales, calculated as:

$$WC = Q_{waste} \cdot COGS \quad (4.2)$$

where

$$\begin{aligned} WC &= \text{Waste Cost [eur]} \\ Q_{waste} &= \text{Quantity of waste in a campaign [units]} \\ COGS &= \text{Cost of Goods Sold [eur/unit]} \end{aligned}$$

Goal: To evaluate the financial consequences of over-forecasting by tracking the value of inventory that is ultimately discarded or donated. The aim is to reduce avoidable waste by incorporating expected waste risk more explicitly into promotional stock decisions and campaign planning. While sustainability is not the primary goal, the indicator also supports broader environmental accountability.

Product Owner: Promo Lead and Inbound Lead

3. Promo Extension captures the additional costs of unplanned discounts or clearing actions that arise from forecasting errors. When too much stock remains after a promotion, Picnic applies one of two mechanisms, depending on whether the article is part of the regular assortment or not.

Mechanism 1 - For regular assortment items, the promotion is extended until the remaining stock corresponds to fewer than 30 days of regular sales. The extended promotion cost is calculated as:

$$C_{ext} = (Q_{base} \cdot M_{reg}) - (Q_{promo} \cdot M_{promo}) \quad (4.3)$$

where

$$\begin{aligned} C_{ext} &= \text{Cost of the extended promotion [eur]} \\ Q_{base} &= \text{Expected base sales quantity [units]} \\ M_{reg} &= \text{Regular front margin per item [eur/unit]} \\ Q_{promo} &= \text{Actual sales quantity during promotion [units]} \\ M_{promo} &= \text{Promo front margin per item [eur/unit]} \end{aligned}$$

This formulation compares the foregone regular margin with the realised promotional margin, thereby quantifying the financial impact of extending the campaign.

Mechanism 2 - For non-regular assortment items, excess inventory cannot be absorbed into the baseline assortment. In these cases, a one-off clearing promotion is applied, with costs calculated as:

$$C_{\text{clear}} = -(Q_{\text{promo}} \cdot M_{\text{promo}}) \quad (4.4)$$

where

$$\begin{aligned} C_{\text{clear}} &= \text{Cost of the clearing promotion [eur]} \\ Q_{\text{promo}} &= \text{Actual sales quantity during promotion [units]} \\ M_{\text{promo}} &= \text{Promo front margin per item [eur/unit]} \end{aligned}$$

This reflects the negative margin contribution of having to clear surplus stock entirely through promotion. Further explanations on this indicators is given in Section 5.1.3.

Goal: To monitor unplanned cost overruns caused by extending promotions to clear excess inventory. This supports more accurate promo budgeting, helps identify systematic over-forecasting, and provides a cost signal that can be fed back into future campaign planning to tighten operational and commercial discipline.

Product Owner: Promo Lead

4. Handling Cost reflects the additional labour required to move surplus stock caused by over-forecasting. The calculation estimates the cost incurred when an article has excess inventory remaining after a campaign:

$$HC_{\text{total},x} = (HT_{\text{handling}} + HT_{\text{transport},x}) \cdot C_{\text{shopper}} \cdot Q_{\text{overstock}} \quad (4.5)$$

where

$$\begin{aligned} HC_{\text{total},x} &= \text{Total handling cost per CU of article } x \text{ [€]} \\ HT_{\text{handling}} &= \text{Average handling time per CU for all handling steps [sec]} \\ HT_{\text{transport},x} &= \text{Average handling time per CU of article } x \text{ for transport steps [sec]} \\ C_{\text{shopper}} &= \text{Labour cost of one shopper [€/h]} \\ Q_{\text{overstock}} &= \text{Quantity of overstocked units at the end of the promotion [units]} \end{aligned}$$

Goal: Tracking this indicator allows the fulfilment team to better understand how forecasting errors disrupt standard processes and contribute to unnecessary labour costs. The intention is to integrate this knowledge into process design and forecasting adjustments to reduce costly manual interventions.

Product Owner: Inbound Lead

5. Space Cost captures the value of storage capacity occupied as a result of over-forecasting. The cost is calculated by multiplying the storage price per m^3 per day by the volume of an article, the total overstocked quantity, and the number of days the article remains in the buffer. This is expressed as:

$$SC = TBT \cdot AV \cdot STC \cdot Q_{\text{overstock}} \quad (4.6)$$

where

$$\begin{aligned} SC &= \text{Space Cost [eur]} \\ TBT &= \text{Total Buffer Time [days]} \\ AV &= \text{Article Volume [m}^3\text{]} \\ STC &= \text{Storage Cost [eur/m}^3\text{/day]} \\ Q_{\text{overstock}} &= \text{Quantity of overstock counted [units]} \end{aligned}$$

Goal: Monitoring space costs associated with promotional items helps identify the impact of avoidable storage caused by over-forecasting. In some cases, storing excess stock may even become more costly than discarding it outright. Understanding these costs enables better-informed trade-off decisions and operational discipline. By quantifying these impacts, Picnic can refine its planning thresholds to:

- Set clear limits on acceptable stock levels per promotion;

- Determine when a promotion idea should be rejected because the associated risk or cost is too high;
- Define space constraints that must be respected in fulfilment centres;
- Avoid scenarios where storage costs outweigh the expected revenue.

Product Owner: Inbound Lead

6. Maximum Economical Retention Time identifies the break-even point when storage becomes more expensive than discarding stock, calculated as:

$$MERT = \frac{COGS - HC}{SC_{day}} \quad (4.7)$$

where

$MERT$ = Maximum Economical Retention Time [days]

$COGS$ = Cost of Goods Sold [eur]

HC = Handling Cost [eur]

SC_{day} = Space Cost per day [eur/day]

Goal: To identify a cost-based time limit for storing slow-moving promotional inventory, helping Picnic avoid operational inefficiencies and prevent negative margin events due to late removal.

Product Owner: Inbound Lead

7. Customer Satisfaction indicates how promotions meet customer expectations, measured through CSAT (Customer Satisfaction Score) surveys or ratings.

Goal: To assess whether promotional execution aligns with customer expectations and to identify areas where poor forecasting results in customer dissatisfaction, helping to prioritise improvements in communication, timing, or availability. By knowing the monetary value of the impact caused by over- or under-forecasting, this feedback can be explicitly integrated into the promotional forecasting process to improve its accuracy.

Product Owner: Customer Team Lead

8. Brand Impact evaluates the long-term effect of stock-outs and promo errors on brand trust, typically measured with perception surveys or sentiment analysis scores.

Goal: To understand the reputational consequences of recurring forecasting errors and assess whether specific promotional failures, such as frequent stockouts, reduce customer trust in Picnic's brand. This insight can support more informed forecasting decisions and strengthen campaign quality control measures.

Product Owner: Customer Team Lead

9. Complaint Rate measures the impact of promotions on forecasting-related customer complaints. To isolate the effect of promotions, the calculation compares the complaint rate during a promotion with the baseline complaint rate for the same article in the regular assortment. The difference represents the net promotional impact, which can be either positive or negative. As a promotion may also reduce the number of complaints, the corresponding cost of complaints can take on a negative value.

This is calculated as:

$$CR = \left(\frac{C_{promo}}{N_{promo}} - \frac{C_{reg}}{N_{reg}} \right) \cdot CC \cdot Q_{promosales} \quad (4.8)$$

where

CR_{promo} = Extra Complaint Rate Cost of a promotion [eur/item]

C_{promo} = Number of promo-related complaints [count]

N_{promo} = Number of promo items sold [count]

C_{reg} = Number of regular-related complaints [count]

N_{reg} = Number of regular items sold [count]

CC = Cost per complaint [eur]

$Q_{promosales}$ = Cost per complaint [eur]

Goal: To detect whether forecasting errors during promotions lead to higher complaint rates, in comparison to regular sales, and to use this signal to inform and improve promotional forecasting accuracy and related decision-making.

Product Owner: Customer Team Lead

10. Product Waste Carbon Footprint estimates the greenhouse gas emissions associated with products discarded by Picnic. The calculation multiplies the quantity of wasted products by their average emission factor per unit, expressed in kilograms of CO₂-equivalent:

$$CF = Q_{waste} \cdot EF \cdot CP \quad (4.9)$$

where

$$\begin{aligned} CF &= \text{Carbon Footprint cost [eur]} \\ Q_{waste} &= \text{Number of unsold articles} \\ EF &= \text{Emission factor per article [kg CO}_2\text{e]} \\ CP &= \text{Carbon price [eur / kg CO}_2\text{e]} \end{aligned}$$

with

$$EF = (PE + PAE + TE) \cdot PW \cdot RW \quad (4.10)$$

where

$$\begin{aligned} PE &= \text{Production emissions [kg CO}_2\text{e / kg]} \\ PAE &= \text{Packaging emissions [kg CO}_2\text{e / kg]} \\ TE &= \text{Transport emissions [kg CO}_2\text{e / kg]} \\ PW &= \text{Product weight [kg / unit]} \end{aligned}$$

Goal: To quantify the environmental impact of promotional forecasting errors and use this information to support more sustainable inventory and campaign planning decisions, particularly by clarifying the sustainability implications of over-forecasting.

Product Owner: Inbound Lead

4.1.3. Relevance-Feasibility Matrix Application

With the ten chosen indicators from the long-list, the next step in the selection process was the creation of the RF-Matrix (see Appendix E), a structured decision-making tool commonly used in management and design research to prioritise and evaluate potential metrics or interventions. The rationale for applying this methodology is described in Phase 3 of Section 1.6.2. This matrix evaluates each option based on two critical dimensions: relevance, representing how effectively an indicator supports organisational goals and decision-making; and feasibility, reflecting the practicality and ease of obtaining accurate and consistent data for the indicator (Al-Ashaab et al., 2013).

In the context of this thesis, relevance was assessed according to how clearly each indicator reflects critical consequences of promotional forecasting errors and aligns with Picnic's strategic and operational objectives. Indicators scoring high on relevance are those directly linked to decision-making quality, operational efficiency, financial impact, or customer satisfaction. Feasibility was judged based on data availability, accuracy, and the time available for the thesis. Highly feasible indicators were already tracked in Picnic's systems and easy to access. Moderately feasible ones needed extra data collection or processing but could be done within the time available. Low feasibility indicators lacked reliable data, required a lot of manual work, or needed complex analysis beyond this thesis. Some low feasibility indicators could be possible for future research or dashboard development.

To summarise, the scoring of the RF-Matrix considers two key factors:

- Relevance was assessed based on the expected business impact for Picnic. In some cases, this required estimating strategic or financial significance despite the absence of direct monetary quantification.
- Feasibility was assessed based on current data availability at Picnic, the effort required to gather or compute reliable values, the analytical complexity involved and the time-frame of the thesis.

This RF-Matrix application ensured a balanced prioritisation, integrating both strategic value and practical considerations, resulting in a focused set of indicators most appropriate and actionable for Picnic’s current needs and the constraints of this thesis.

4.1.4. Stakeholder Validation Process & Result of the RF Matrix

To ensure that the selected indicators were both strategically relevant and practically feasible, a structured validation process was conducted with three domain experts within Picnic, each representing one of the indicator domains: financial (Promo Lead), operational (Inbound Lead), and reputational (Customer Team Lead). The approach consisted of three steps, which the researcher based on the elaboration of the RF-Matrix by Al-Ashaab et al. (2013), as outlined below:

- **One-on-one Reviews:** Each stakeholder was consulted individually to review the RF-Matrix created earlier, with a focus on their respective domain, allowing for targeted feedback. Their personal adjustments were visualised on a new RF-Matrix, which are shown in Appendix E.
- **Iterative Refinement:** Feedback was incorporated into a revised RF-Matrix. In cases where small disagreements arose between stakeholders, the final RF-Matrix gave more weight to the influence of a stakeholder within their own domain, compared to the input from other stakeholders.
- **Final Validation:** The updated RF-Matrix was presented to all stakeholders for a final check. After minor adjustments, all parties confirmed the final version shown at the bottom of the overview in Appendix E.

The most suitable indicators for evaluating the consequences of promotional forecasting errors at Picnic were earlier shown in Figure 4.1. The resulting RF-Matrix, derived from the three-step approach, is shown in Figure 4.2. A visualisation of this process is provided in Appendix E.1.




Figure 4.2: Final RF-Matrix result.

As demonstrated, all financial and operational indicators are positioned near the top right of the RF-Matrix and therefore chosen to be within the scope of this research. The rationale behind their selection, and that of the non-selected indicators, is discussed in Section 4.2. It should be noted that the non-chosen indicators may still be of interest, but they were not considered suitable for the scope of this study.

4.2. Justification of Indicator Selection and Exclusion

Based on their strong position in the mainly upper-right quadrant of the RF-Matrix, indicators 1 through 6 were selected for inclusion in the dashboard and further analysis. These indicators offer both high relevance to Picnic’s decision-making and sufficient feasibility for implementation within the scope of this thesis. The financial indicators 1 and 2 are already partially tracked within Picnic and can be reliably linked to financial impact. The operational indicators (4-6) are also directly affected by stock imbalances, their inclusion is justified by their practical measurability and strategic importance in fulfilment operations. This section justifies the selection of

the chosen indicators. Further insights are presented in Section 4.2.4, which offers a comprehensive overview, including the calculation method, goal, limitations and assumptions, future improvements, and the responsible stakeholder for each indicator.

4.2.1. Financial Indicators

The selected financial indicators reflect the most direct and quantifiable consequences of promotional forecasting errors. Each of these metrics links explicitly to Picnic's profitability, and the impact of promotional forecasting. Financial indicators are prioritised not only because they are relatively straightforward to measure and monetise, but also because they align directly with core business KPIs within Picnic. These indicators enable commercial teams to assess the financial impact of forecasting decisions and optimise planning accordingly.

1. Lost Margin

This indicator captures the unrealised profit associated with under-forecasted promotions. When demand is underestimated, products run out of stock too early, causing potential sales to be missed. Since Picnic often operates with tight margins and time-limited campaigns, the ability to measure how much margin was lost due to insufficient inventory is highly valuable. This indicator is particularly relevant for this thesis, as it provides direct insight into the financial impact of under-forecasting errors. It is also feasible within the scope of this research, as Picnic already tracks this metric internally.

2. Waste Cost

This indicator shows the total value of products that cannot be sold because they passed their best-before date or do sell way too slow. This usually happens when too much is ordered. Waste is especially important for fresh or 'chilled' (cooled) products, which make up a large part of Picnic's promotions. Besides the direct financial loss from wasted products, it also hurts Picnic's sustainability goals. Although this impact is not yet measured in detail company-wide, this indicator helps to better understand Picnic's food waste and thus contributes to evaluating a part of its sustainability impact, which gives some extra relevance. It is also considered feasible within the scope of this research, as again Picnic already tracks this metric internally.

3. Promo Extension

Promo extensions are often used as a reactive measure to clear excess inventory caused by over-forecasting. These extensions may include unplanned discounts, which have direct financial implications. By monitoring these costs, Picnic can assess how frequently deviations from the original promotion plan occur and the impact on budget overruns. However, in some cases, the forced extension of promotions for certain products may positively impact profitability, providing valuable insights. This information is essential for improving promotional forecasting and pricing strategies. The indicator is highly relevant, as it directly reflects the financial consequences of forecasting errors and supports better budget management. It is also feasible within the scope of this thesis, as promotional extension activities are tracked within the company, allowing for accurate calculation of the associated costs.

4.2.2. Operational Indicators

Two operational indicators were selected based on their practical feasibility and direct relevance to fulfilment operations: Handling Costs and Space Costs. A third indicator, Maximum Economical Retention Time (MERT), is derived from the handling and space cost indicators and provides valuable insight into whether a product should be stored or discarded upfront (before storage). Due to its significant decision-making value, MERT scored highly on relevance. Furthermore, it is relatively easy to calculate once the space and handling costs have been determined. The relevance of the indicators aligns with the financial indicators, while their feasibility is slightly lower, as they are not currently tracked within Picnic. Nevertheless, the indicators are conceptually straightforward and could be approximated using available data. Therefore, their inclusion in this research could provide valuable new insights and help pave the way for formalising their tracking within operational systems in the future.

4. Handling Cost

This indicator reflects the additional labour and complexity associated with moving, managing, or relocating inventory during or after a promotion. Over-forecasting can result in surplus stock that needs to be handled multiple times. At Picnic, products are moved from replenishment to the 'buffer zone' and back to replenishment. Under-forecasting may lead to urgent replenishments or manual handling, disrupting standard workflows and reducing warehouse efficiency. Since Picnic's main supplier, ITM, does not allow back-orders for the majority of its assortment, the occurrence of these extra handling costs is very low. Due to the relatively limited frequency and lack of systematic tracking, this impact falls outside the scope of this research. However, given its clear operational relevance and practical feasibility, this indicator has been selected for inclusion in the dashboard to be developed.

5. Space Cost

This indicator captures the opportunity cost of using limited warehouse space for unsold or slow-moving promotional inventory. As shown in the final RF-matrix (Appendix E.1), it was included due to its high relevance and practical feasibility. While not currently tracked, it can be estimated using available data on inventory volumes and storage duration. Within this research, it is feasible to calculate these costs in time, offering clear insights into when storage costs outweigh potential revenue. This supports better decision-making around stock limits, planning thresholds, and the selection of space-intensive promotions.

6. Maximum Economical Retention Time

The Maximum Economical Retention Time (MERT) indicator serves as a critical tool for evaluating the cost-effectiveness of storing promotional inventory. The goal of this indicator is to identify the point in time at which the cost of holding inventory exceeds the cost of discarding it upfront. By calculating the number of days until the cost of storage surpasses the benefit of keeping the product, Picnic can optimise inventory management and avoid unnecessary operational inefficiencies. MERT is highly relevant for Picnic, particularly in preventing negative margin events. When products are not selling as expected, they tie up valuable storage space and resources, which increases handling costs and impacts warehouse efficiency. By identifying the optimal time to dispose of slow-moving items, this indicator helps mitigate these costs and avoid prolonging the shelf life of unprofitable stock. Given that MERT is derived from other operational indicators like handling and space costs, it is relatively straightforward to calculate once these base metrics are available. Although it is not currently tracked within Picnic, the feasibility of this indicator is high, and it can be easily integrated into the dashboard. Its inclusion would provide actionable insights into inventory management and significantly enhance decision-making regarding slow-moving stock, making it a valuable addition to the dashboard developed in this research.

4.2.3. Reputational Indicators

While reputational indicators can provide important insights into how promotional forecasting errors affect customer perception and brand trust, none were selected for inclusion in this research. Although some were conceptually relevant, their feasibility is limited due to the absence of existing tracking mechanisms, the complexity involved in isolating the impact of forecasting errors, and the time-intensive nature of monetising their effects. As these indicators typically require qualitative interpretation or long-term observation, making them less suitable for the time-bound and data-driven scope of this thesis. Nonetheless, they remain promising candidates for future research or internal development.

7. Customer Satisfaction

This indicator measures how satisfied customers are with specific promotions, often via post-purchase surveys or app feedback. CSAT (Customer Satisfaction Score) captures the extent to which customer expectations are met, typically through a simple rating question (Gastezzi C. et al., 2024). It is relevant because stock-outs can negatively affect customer perception. However, it was not selected due to limited feasibility: Picnic does not currently track CSAT at a promotion level, and isolating the impact of forecasting adds complexity. In addition, monetising this indicator requires extra effort. Altogether, this makes it unfeasible within the scope of this research.

8. Brand Impact

The Brand Impact indicator captures how promotional forecasting errors affect customer trust in Picnic's brand. Repeated stock-outs or inconsistent promotional execution may erode brand reliability over time. Brand impact can be measured through methods such as brand perception surveys or social media sentiment analysis. However, isolating the specific effect of forecasting errors is complex, as brand perception is influenced by many other factors. Due to this complexity, lack of direct data, and the time-intensive nature of such analysis, the indicator was considered infeasible within the scope of this thesis. While relevant, it was not selected.

9. Complaint Rate

This indicator reflects the frequency of complaints related to promotional orders and is typically calculated by dividing the number of complaints linked to promos by the total number of promotional orders. To make the metric meaningful, it would ideally be benchmarked against the complaint rate for regular sales. While conceptually clear, the relevance of this indicator is limited. Most complaint types (e.g. delivery or product quality issues) are not forecasting-related, and only a small increase in unavailability-related complaints might be observed during promotions. The feasibility is also constrained. It is challenging to create a model that directly links complaints to promotional forecasting errors. Additional data processing would be required to isolate meaningful, promo-specific insights. As such, the indicator was not selected due to both its modest added value and implementation effort.

10. Carbon Footprint

This indicator measures the environmental impact of forecasting errors, focusing on emissions from overproduction, transport, storage, and waste. It is typically calculated using carbon accounting models that apply emission factors to different supply chain activities. Despite growing attention to sustainability, this indicator was not selected due to both low feasibility and low relevance in the context of this thesis. Picnic does not currently attribute carbon emissions at a detailed, promotion-specific level, and calculating these impacts would require substantial modelling and assumptions. Additionally, the relevance of this indicator is limited because Picnic rarely discards food outright. Unsold items are typically donated to those in need through redistribution partners, meaning this is not considered true food waste and therefore does not have a major negative impact on the carbon footprint. While the difference in waste between promotional and non-promotional products could have had a significant impact, it does not. Although some handling and redundant transport do contribute a small negative impact, it remains minimal. Consequently, the direct carbon cost of waste is negligible, making this indicator less valuable compared to the more actionable metrics already selected.

4.2.4. Notion on Chosen Indicators

The formulas presented for the six selected indicators yield absolute values. In the developed dashboard, however, these indicators are calculated and displayed at two complementary levels:

1. Absolute value per campaign
2. Value per item

This dual representation enables both a strategic overview and a more granular, article-level perspective.

To test the robustness and practical relevance of the indicators, they will be applied in Chapter 7 to two illustrative cases: one article that is frequently under-forecasted and another that is both over-forecasted and under-forecasted. These contrasting examples allow the usefulness of the artefact, developed in Chapter 6 of this study, to be assessed and provide a comprehensive understanding of how each indicator behaves in practice.

It should also be noted that while reputational indicators were ultimately not included in the final selection, the conceptualisation presented in this study is not definitive. Their purpose here was to ensure a robust research framework (RF-Matrix). Future approaches may develop alternative formulations of reputational indicators, which could yield additional insights.

4.3. Main Findings Chapter 4

This chapter aimed to select the key performance indicators affected by over- and under-stocking due to promotional forecasting errors at Picnic, and to determine the specific insights Picnic aims to derive per indicator. A structured, stepwise approach was followed, beginning with the construction of a longlist of fifteen indicators, followed by the refinement into a shortlist of ten indicators, which were subsequently examined in greater depth through a review and validation process involving three key stakeholders.

The approach proved both efficient and reliable. Consensus among stakeholders was reached with minimal iterations, and the combination of internal expertise and systematic evaluation ensured that the selected indicators were both practically relevant and strategically aligned. Six indicators were ultimately chosen for implementation in the dashboard, reflecting strong relevance and feasibility within the scope of this thesis. These were financial and operational indicators, while the remaining four (reputational) indicators were excluded, mainly due to feasibility constraints. Notably, while the indicators 'Carbon footprint' and 'Complaint rate' scored moderately to low on relevance, the indicators 'Customer satisfaction' and 'Brand impact' scored high on relevance. However, the overarching reason for excluding reputational indicators was feasibility rather than lack of strategic value, as the research needed to remain a manageable challenge. It is therefore recommended to revisit these excluded indicators in future stages, when more time and improved data infrastructure may enable their use. For each selected indicator, this chapter also outlined the goal, limitations, and assumptions, ensuring a complete and transparent foundation for readers.

In sum, this chapter achieved its goals by providing a validated set of indicators that form a strong analytical basis for a better understanding of the consequences of promotional forecasting errors. These indicators not only guide the structure of the dashboard developed in the next phase, but also contribute to Picnic's broader efforts to enhance operational efficiency and decision-making. At the same time, the chapter introduces areas of ambiguity, particularly concerning the specific calculation methods and the intended level of detail for each KPI. These issues are addressed in the following chapter, Chapter 5.

5

Indicator Calculation Approach

As the justification, objectives, and simplified calculations for each indicator have already been discussed in Chapter 4, this chapter provides a detailed examination of the selected indicators. For each indicator, the calculation method is presented alongside the underlying background, assumptions, and limitations. Where relevant, potential avenues for improvement and refinement in future applications are also highlighted. In particular, this section focuses on the development work carried out in the third phase of the research, where the calculations and methodological choices were elaborated. On that basis, the following research question was formulated:

3. *How should the indicator calculations be designed to assess the cost implications of over-stocking and under-stocking resulting from forecasting errors in Picnic's promotional context?*

The chapter is structured as follows. First, the three financial indicators are discussed in Section 5.1. This is followed by a discussion of the three operational indicators in Section 5.2. The chapter concludes with final remarks and a reflection on the findings in Section 5.3.

5.1. Determination of Cost for Financial Indicators

This section elaborates on the approach taken for calculating the financial indicators. It begins with lost margin, which reflects lost sales resulting from product unavailability in the store. Next, waste cost is addressed, representing the financial loss associated with unsold products. The section concludes with promo extension cost, a cost item arising from unplanned discounts.

5.1.1. Lost Margin

The determination of lost margin accounts for the gross profit that Picnic loses when a product is unavailable in the online store (the Picnic app). In the app, when a customer selects a product, it is directly added to the shopping basket. However, if the product is unavailable, the customer is notified and cannot purchase it. Picnic tracks all customer interactions within the app, including clicks on unavailable products. These data form the basis for calculating the lost margin.

The first step is to determine the unavailable quantity per campaign. This is defined as the product of the number of unique customers clicking on unavailable products, the commonality factor, and the conversion ratio. The unique customers clicking on unavailable products is tracked in the app by Picnic. At Picnic, the conversion ratio, the percentage of clicks in the application that result in an actual purchase, is fixed at 70% across all markets. This percentage is based on an internal study conducted by Picnic employees and is applied company-wide. The commonality factor, the average number of units of a product purchased per delivery, is calculated at the article level to ensure product-specific accuracy. It accounts for the fact that an unavailability does not only represent the loss of a single unit, but also of potential multiple units that a customer might have purchased. In practice, the commonality factor reflects the average sales quantity per product per delivery, which typically ranges between 1.1 and 1.3 units. The calculation is given by:

$$UQ_y = C_{\text{unav}} \cdot C \cdot CR \quad (5.1)$$

where

$$\begin{aligned} UQ_y &= \text{Unavailable quantity per campaign } y \text{ [units]} \\ C_{\text{unav}} &= \text{Number of unique customers clicking on unavailable products [count]} \\ C &= \text{Commonality factor [units/order]} \\ CR &= \text{Conversion ratio} = 70\% = 0.70 \end{aligned}$$

The commonality factor is calculated at article level, ensuring product-specific accuracy:

$$C = \frac{Q_{\text{promo}}}{D_{\text{promo}}} \quad (5.2)$$

where

$$\begin{aligned} C_{\text{promo}} &= \text{Commonality factor [units/order]} \\ Q_{\text{promo}} &= \text{Total items sold in promotion [units]} \\ D_{\text{promo}} &= \text{Unique orders containing the promotion [count]} \end{aligned}$$

Once the unavailable quantity has been determined, the lost margin can be calculated. As outlined in Section 1.1, all promotions are funded by Picnic's supplier, ITM. The margin for Picnic on a promoted product therefore consists of two components: (1) the margin based on the promotional price, referred to as the front margin, and (2) the supplier contribution, i.e. the discount reimbursed by ITM. The calculation is as follows:

$$LM_{x,y} = UQ_{x,y} \cdot LM_x \quad (5.3)$$

$$\begin{aligned} LM_x &= FM_x + SC_x \\ LM_{x,w} &= \text{lost margin for article } x \text{ in campaign } y \text{ [€]} \\ UQ_{x,w} &= \text{unavailable quantity (missed units) of article } x \text{ in campaign } y \text{ [units]} \\ FM_x &= \text{front margin per unit of article } x \text{ [€/unit]} \\ SC_x &= \text{supplier contribution per unit of article } x \text{ [€/unit]} \end{aligned}$$

These calculations provide a clear representation of how lost margin is quantified in Picnic's dashboard. However, several limitations must be acknowledged:

- The conversion ratio is fixed at 70%, based on the Dutch market and across all product categories. This generalised approach may not reflect variations between product types or across countries (NL, DE, FR). Moreover, the conversion ratio is based on regular sales, whereas promotional sales may exhibit different behaviour.
- The effect of cannibalisation has not been considered. Promotions can affect the sales of substitute products, as demand is partly shifted from alternatives to the promoted item. Consequently, when a product becomes unavailable, the extent of cannibalisation decreases. From an analytical perspective, this implies that the additional loss from reduced sales of substitute products should be incorporated into the calculation of the total loss associated with an unavailable product.
- Conversely, the mitigating effect of substitution has also been excluded. When a product is unavailable, a certain proportion of customers may opt for an alternative product. In such cases, the actual lost margin is lower than the full lost margin estimate, as part of the demand is absorbed by substitutes. This effect, which reduces the effective loss, is not accounted for in the current framework.
- Reputational damage is out of scope. A product being unavailable in the app may negatively affect customer satisfaction and Picnic's brand image. As discussed in Section 4.2, this effect is not included in the calculation.

5.1.2. Waste Cost

The determination of waste cost is relatively complex, as it requires careful identification of whether disposed products are the result of forecasting errors in promotional sales. To achieve this, the root causes of waste are first categorised. Causes such as *broken*, *out of assortment*, *stock count* are unrelated to forecasting errors, whereas categories such as *delivery forecast*, *promotion*, *trend* are more likely to be associated with forecasting inaccuracies. However, this distinction alone is insufficient: it must also be verified that the waste originates specifically from promotional sales rather than regular sales, and that it occurs within an appropriate time window following the promotion of the article.

Before presenting the calculation, a distinction is made between three categories of promotional waste:

1. **Regular promo waste:** waste of products also sold in the regular assortment.
2. **BBD waste:** waste of products with a short best-before date (BBD).
3. **In/Out promo waste:** waste of products only available during a specific promotion and not part of the permanent assortment.

In order to determine whether waste of a given type can be attributed to over-forecasting, the occurrence of waste is examined relative to the defined start and end windows. If waste occurs within these windows, and the recorded root cause is consistent with forecasting error, it is considered valid for inclusion in the calculation. The specific time windows applied are summarised in Table 5.1.

Table 5.1: Time window and definitions for waste categories and promo extensions.

Type	Start Date	Explanation	End Date	Explanation
Waste Regular	Monday after promotion	Regular product waste is assumed to occur from the first Monday following the promotion.	Until the next purchase order inbound date	The inbound of a new product signals that any remaining stock from the previous promotion is cleared.
Waste BBD	Monday of first promotion week	BBD (best-before-date) waste may arise from the second day of the promotion.	Until the next purchase order inbound date	The inbound of a new product signals that any remaining stock from the previous promotion is cleared.
Waste In/Out	Monday after promotion	In/Out product waste is assumed to occur from the first Monday following the promotion.	Until the next promotion has taken place	As these products are not sold regularly, all waste is attributable to forecasting error until a new In/Out promotion occurs.
Promo Extension	Monday after promotion	An extended promotion always begins on the Monday after the original promotion ends.	Until the next promotion starts or purchase order is placed	A new promotion or inbound order defines the endpoint, as this indicates that any remaining stock from the previous promotion is cleared.

Once the time windows are established, the cost of waste can be calculated. Picnic’s dataset includes both the root cause of waste and the COGS associated with each article. Waste cost is therefore defined as the product of the quantity of waste within the valid window and the corresponding COGS of the article:

$$WC_{x,y} = Q_{\text{waste},x,y} \cdot COGS_x \quad (5.4)$$

where

$$\begin{aligned} WC_{x,y} &= \text{Waste cost of article } x \text{ in campaign } y \text{ [€]} \\ Q_{\text{waste},x,y} &= \text{Quantity of waste for article } x \text{ in campaign } y \text{ [units]} \\ COGS_x &= \text{Cost of goods sold for article } x \text{ [€/unit]} \end{aligned}$$

Now that the calculation has been fully explained, a number of additional remarks remain that are important to be aware of. These are outlined below:

- The COGS of each article is used as a proxy for waste cost, thereby ignoring any overhead or handling costs associated with waste processing. According to the inbound lead, this costs are relatively too small, and therefore no need to take into account.
- Waste that is donated to food banks or to alternative channels (e.g., Too Good To Go) is still included, although such flows positively influence environmental impact.
- Waste is currently recorded and categorised through an internationally standardised system based on the Dutch market. However, as Picnic France uses In/Out products, these cannot be captured by the system and are automatically classified as ‘unknown’. Consequently, a large share of waste causes in the French market remains unclear. In this study, all ‘unknown’ cases are treated as over-forecast waste, as In/Out waste predominantly results from over-forecasting.

- The potential reputational damage for Picnic as a company, linked to the environmental aspects of food waste, as well as the price for unnecessary CO₂ emissions associated with waste, falls outside the scope of this research. This decision was made due to the relatively minor impact, further justification provided in Section 4.2.

5.1.3. Promo Extension

The third financial indicator concerns the margin lost due to the application of additional discounts through promotion extensions. When a 'supplier-funded promo', a general promotion, is running, it may occur that not enough products are sold, resulting in overstock. To reduce this stock more quickly, Picnic typically extends the promotion period. The percentage of discount offered depends on the type of promotion. As explained in Section 5.1.2, two types of promotions are distinguished: regular promotions and in/out promotions. No distinction is made between BBD and non-BBD products for the purpose of promo extension, resulting in two variants:

- **Extended promo (15% discount):** continues until an acceptable stock threshold (equivalent to one month of average sales) is reached. This applies to regular products.
- **Clearing promo (25% discount):** continues for exactly one week. After this period, all remaining stock is either discarded or donated, except for certain high-value products, and therefore considered waste.

The cost of a promotion extension is based on an uplift in sales and a reduction in margin, and the outcome can either be positive or negative. Two types of promo extensions exist, extended promo or clearing promo. The discounts of the promotions for both are financed internally by Picnic, and the extra discount is treated as a lost margin cost to the company. While the principle appears straightforward, the calculation methodology differs between extended and clearing promotions. For consistency, both types are evaluated over the same time window. Although clearing promotions typically conclude within one week, exceptions occur, and therefore it is operationally more practical to apply a uniform time horizon across both cases, as the window is still correct for both types.

For **extended promotions**, two elements are considered: the cost incurred and the revenue gained. A positive effect of promo extension is the potential for increased turnover. To capture this, base sales are compared with actual promotional sales. Base sales are automatically tracked within Picnic and represent the estimated turnover under regular sales conditions. The difference between the margin under base sales and the margin realised during the promo extension indicates the net impact of the extension. A constraint is implemented such that base sales cannot fall below promotion sales. This relationship is expressed as follows:

$$C_{\text{ext}} = (Q_{\text{base}} \cdot M_{\text{reg}}) - (Q_{\text{promo}} \cdot M_{\text{promo}}) \quad (5.5)$$

where

$$\begin{aligned} C_{\text{ext}} &= \text{Cost of the extended promotion [eur]} \\ Q_{\text{base}} &= \text{Expected base sales quantity [units]} \\ M_{\text{reg}} &= \text{Regular front margin per item [eur/unit]} \\ Q_{\text{promo}} &= \text{Actual sales quantity during extended promotion [units]} \\ M_{\text{promo}} &= \text{Extended promo front margin per item [eur/unit]} \end{aligned}$$

For **clearing promotions**, the situation differs. As the product in clearance is not typically sold under regular conditions, no base sales exist. Instead, the benchmark value is zero. Consequently, every sale during a clearing promotion generates a positive margin relative to this benchmark. When expressed as costs, this results in a negative value, reflecting the revenues gained. The formula is therefore defined as:

$$C_{\text{clear}} = - (Q_{\text{promo}} \cdot M_{\text{promo}}) \quad (5.6)$$

where

$$\begin{aligned} C_{\text{clear}} &= \text{Cost of the clearing promotion [eur]} \\ Q_{\text{promo}} &= \text{Actual sales quantity during clearing promotion [units]} \\ M_{\text{promo}} &= \text{Clearing promo front margin per item [eur/unit]} \end{aligned}$$

The distinction between the two approaches lies in the benchmark applied. Extended promotions are benchmarked against the margin that would have been achieved under regular sales conditions, whereas clearing promotions are benchmarked against a zero-sales baseline.

For this KPI, several important assumptions and considerations apply, all of which were discussed with the Promo Team Lead and are outlined below:

- Cannibalisation effects, whereby promo-driven sales reduce sales of substitute products, are excluded. In case, the extended and clearing promo would have extra costs.
- Base sales are constrained not to fall below promotion sales; this condition is embedded in the code.
- The accuracy of the benchmark relies on consistently reliable base sales forecasts.
- Customer retention or goodwill effects of extended promotions are not captured.
- Self-financed promotions always generate a positive margin (according to the law). They may not be loss-making, for example to retain customers.
- No marketing costs are considered within the scope of this indicator.

5.2. Determination of Cost for the Operational Indicators

This section elaborates on the approach taken for calculating the operational indicators. It begins with handling cost, which reflects the additional labour required to move surplus stock or urgently replenish items. Next, space cost is discussed, which captures the value of storage capacity occupied as a consequence of forecasting errors in a Picnic FC. The section concludes with MERT, which identifies the break-even point at which storage becomes more costly than discarding stock.

5.2.1. Handling Cost

For the handling cost, the additional handling and transport movements caused by a forecasting error are taken into account. Therefore, it is important to understand which additional steps occur when a forecasting mistake arises. After multiple visits to the FC to observe the process and extensive discussions with both the operational manager of FC0 (Lille) and the Inbound Lead, the approach outlined in this section has been determined.

Figure 5.1, which provides a further elaboration of the overview map of FC0 presented in Section 2.2.4, illustrates both the general process and the additional steps in detail. The approach shown is based on the operations within the ambient zone of FC0. The steps are as follows:

1. **Handling:** Fill Rolling Container (RC) from replenishment to RC at aisle
2. **Transport:** Transport RC from aisle to buffer zone
3. **Handling:** Empty RC in buffer zone
4. **Handling:** Fill RC in buffer zone
5. **Transport:** Transport RC from buffer zone to aisle
6. **Handling:** Replenish products

This constitutes the complete extra loop resulting from forecasting errors. However, step 6 is excluded for a specific reason: replenishment is part of the standard process and is not repeated. Therefore, step 6 does not generate additional costs, even though it appears as the final step in the process. For the handling operations, these steps were implemented in SQL code, and the calculation shown below was then carried out for each process step.

To determine an average handling time for products, a large dataset covering the last 12 months was considered. As many factors influence handling time, an average across a one-year dataset was deemed the fairest approach, regarding the Inbound Lead. The dataset consists of approximately 129,000 rows, which is sufficiently reliable. The total time for steps 1, 3, and 4 was summed and divided by the total number of CUs handled, yielding the average handling time per CU. The formula below represents this calculation:

$$HT_{handling} = \frac{T_{total,1}}{CU_{total,1}} + \frac{T_{total,3}}{CU_{total,3}} + \frac{T_{total,4}}{CU_{total,4}} \quad (5.7)$$

where

$HT_{handling}$ = Average handling time per CU for all handling steps [sec]

$T_{total,y}$ = Total time to handle all CUs in step y [sec]

$CU_{total,y}$ = Total CUs handled in step y [units]




Figure 5.1: Process Flow of extra handling steps in the ambient zone of FC0, located in Lille.

For transport, the calculation is more straightforward. As advised by the FC Operational Manager and confirmed by the Inbound Lead, transport times (steps 2 and 5) are relatively small, making a simplified calculation sufficient. A distinction, however, must be made between the two steps: step 2 always involves a fully loaded RC, whereas step 5 on average involves an RC filled to one-third of its capacity. This was confirmed both by the FC Lead and through direct observations. The lower fill rate in step 5 arises because replenishment from the buffer to the shelf occurs immediately when products are needed; shoppers do not wait until the RC is fully loaded. For both step 2 (aisle to buffer) and step 5 (buffer to aisle), an average transport time of one minute per movement is assumed. The calculation is expressed as:

$$HT_{transport,x} = t_{2,5} \cdot \left(\frac{CU_x}{RC} + \frac{CU_x}{\frac{1}{3}RC} \right) \quad (5.8)$$

where

$HT_{transport,x}$ = Average handling time per CU of article x for transport steps [sec]

$t_{2,5}$ = Total time for transport movements, step 2 and 5 of the process [sec] = 120sec

CU_x = Total cubic volume of CU of article x [m^3]

RC = 90% of Roll container capacity [m^3] = $0.8323[m^3]$

With the average handling time per CU established for each step, a straightforward formula was sufficient to determine the average cost of handling per CU:

$$HC_{total,x} = \left((HT_{handling} + HT_{transport,x}) \cdot C_{shopper} \right) \cdot Q_{overstock,x,y} \quad (5.9)$$

where

$$\begin{aligned}
 HC_{total,x} &= \text{Total handling cost per CU of article } x \text{ [€]} \\
 HT_{handling} &= \text{Average handling time per CU for all handling steps [sec]} \\
 HT_{transport,x} &= \text{Average handling time per CU of article } x \text{ for transport steps [sec]} \\
 C_{shopper} &= \text{Cost of one shopper [€/h]} \\
 Q_{overstock,x,y} &= \text{Quantity of overstocked units of article } x \text{ in campaign } y \text{ [CU]}
 \end{aligned}$$

The following assumptions and principles were agreed upon with the operational manager of FC0 Lead and the Inbound Lead:

- The total cost of handling is based entirely on operations in the ambient zone of FC0. Processes for fresh and frozen products have not been determined; therefore, their costs are approximated using the ambient process. Operations for ambient products in FC1 are considered comparable.
- The pure handling time per item, excluding transport, is calculated as an average of the total sum. This approach is not entirely precise but was agreed to be sufficiently accurate.
- Step 6 is excluded from the handling cost, as the final replenishment action also forms part of the general process and does not constitute an additional step.
- A one-year dataset is used, which is deemed sufficient to provide reliable results.
- Step 5, in which the RC transports products to the replenishment shelf, is assumed to operate at an average fill rate of one-third.
- The average transport time per RC (steps 2 and 5 combined) is assumed to be 120 seconds, based on 60 seconds per movement.
- The average total cost per hour of a shopper at Picnic France is €xxx, this is based on data of 2025-01 to 2025-09.
- 90% of a standardised RC size of 0.9248 m^3 (Kruizinga, 2025) is used, as the RC won't be perfectly filled.
- Back-orders resulting from under-forecasting may lead to urgent manual handling and workflow disruptions. However, as back-orders are permitted for approximately only 25% of promotional products and consistent tracking data are lacking, this factor is excluded from the present research.

These formulas form the foundation of the SQL script used to extract the data for the dashboard, this script can be found, on code sharing platform 'Github', via a link in Appendix A. It should be noted that the SQL script calculates at different levels of aggregation, rather than solely at the level of the average cost per CU, which was demonstrated in this subsection.

5.2.2. Space Cost

The fifth indicator included in the dashboard is the space cost. This indicator captures the opportunity cost of storage capacity in the buffer zone that is unnecessarily occupied as a result of forecasting errors. As explained in Section 2.2.4, the cost is determined using the ABC principle. Which simply is explained as the time multiplied by the related cost. To translate the time into costs, the relevant cost drivers associated with buffer space in Picnic's fulfilment centres (FCs) were identified through a dedicated literature review (Section 2.2.4). These costs are linked specifically to the operation of the buffer zone, and form the basis of the storage cost per cubic metre per day, defined in Equation Equation (5.11).

The total buffer time (TBT) is calculated using two key principles. The first part focuses on overstock that has been stored in the buffer zone from the moment it arrives at the fulfilment centre (FC) until the final day of the SF-promo. The overstock identified on the first day after the SF-promo is considered as surplus from the point it entered the FC, and this quantity is then multiplied by the number of days the products have been stored in the FC. The second part of the TBT calculation begins after the SF-promo, where the stock in the buffer is counted daily at the article level. This stock is summed until the stock level reaches zero, with each stock count multiplied by 24 hours for subsequent days and 12 hours for day one, assuming that on average a product will be left in the buffer for half the day, to account for the time dimension. The formula for this calculation is as follows:

$$TBT = (Q_{initial,x} \cdot D_{initial,x}) + \sum_{t=0}^n Q_{x,t} \cdot d_t \quad (5.10)$$

where

$$\begin{aligned}
TBT &= \text{Total Buffer Time [hours]} \\
Q_{\text{initial},x} &= \text{Overstock quantity of article } x \text{ on the first day after SF-promo [units]} \\
D_{\text{initial},x} &= \text{Total time of article } x \text{ being in buffer from arrival in FC to final day of SF-promo [hours]} \\
Q_{x,t} &= \text{Quantity of article } x \text{ in buffer at day } t \text{ [units]} \\
d_t &= \begin{cases} 12 & \text{if } t = 0 \text{ (half-day assumption for initial stock)} \\ 24 & \text{if } t \geq 1 \text{ (full-day occupancy)} \end{cases} \\
n &= \text{Last day with } Q_{x,t} > 0
\end{aligned}$$

The resulting TBT is subsequently converted into days for use in the space cost calculation. The daily storage cost per cubic metre (STC) is derived from the relevant cost components associated with buffer zone usage. These include rental or lease expenses, utilities, depreciation and maintenance of buildings and racks, as well as cleaning and related services. This breakdown is formalised in the following equation:

$$STC = \frac{(C_{\text{rent}} + C_{\text{utilities}} + C_{\text{cleaning}} + C_{\text{security}}) \cdot \frac{S_{\text{buffer}}}{S_{\text{FC}}}}{V_{\text{buffer}}} \quad (5.11)$$

where

$$\begin{aligned}
STC &= \text{Storage Cost per } m^3 \text{ per day [eur}/m^3/\text{day]} \\
C_{\text{rent}} &= \text{Rent and lease costs, including insurance [eur/day]} \\
C_{\text{utilities}} &= \text{Electricity, gas, and water costs [eur/day]} \\
C_{\text{cleaning}} &= \text{Cleaning and other services [eur/day]} \\
C_{\text{security}} &= \text{Security by video surveillance [eur/day]} \\
S_{\text{buffer}} &= \text{Surface area of the buffer zone [m}^2\text{]} \\
S_{\text{FC}} &= \text{Total surface area of the fulfilment centre [m}^2\text{]} \\
V_{\text{buffer}} &= \text{Total buffer zone volume [m}^3\text{]}
\end{aligned}$$

Using Table 5.2, the value for the STC can be calculated. The calculation is given in Equation (5.12), which shows a value of xxx euro for the storage cost per m^3 per day. All dimensions in Table 5.2 from FC0 are based on the floor plan provided by the operational manager of FC0.

$$STC = \frac{(xx + xx + xx + xx) \cdot \frac{xx}{xx}}{xx} = xx\text{eur}/m^3/\text{day} \quad (5.12)$$

Table 5.2: Overview of FC0's Physical Dimensions and Monthly Costs.
Data extracted from map of FC0 and finance department.

FC0 - Metric	Value	Unit	Share (%)	Note
FC0 total surface area	14855	m^2	100,00	Office space in FC0 is excluded
FC0 buffer zone surface area	266	m^2	1,79	-
FC0 buffer zone max. volume	80	m^3	-	1,2m max. height, 20% of total buffer zone confirmed by Inbound Lead
Monthly rent	xxx	€	74,21	-
Monthly cleaning service	xxx	€	14,55	Waste management costs included
Monthly telesurveillance	xxx	€	7,54	No physical security at FC0
Monthly utilities	xxx	€	3,70	-

The resulting STC provides the basis for calculating the total space cost (SC) associated with surplus stock caused by forecasting errors. This is achieved by multiplying the total buffer time (TBT) of a product by its article volume (AV) and the unit storage cost (STC). The resulting formula is presented below:

$$SC_{x,y} = TBT_{x,y} \cdot AV_x \cdot STC \quad (5.13)$$

where

$$\begin{aligned} SC_{x,y} &= \text{Space Cost of article } x \text{ in campaign } y \text{ [eur]} \\ TBT_{x,y} &= \text{Total Buffer Time of article } x \text{ in campaign } y \text{ [days]} \\ AV_x &= \text{Article Volume of article } x \text{ [m}^3\text{]} \\ STC &= \text{Storage Cost [eur/m}^3\text{/day]} \end{aligned}$$

A few additional remarks are worth noting to ensure full clarity around this KPI:

- The formula of STC (Storage Cost) was derived in consultation with the Inbound Lead and the Operational Manager of FC0, and validated against literature sources.
- The current space cost is based on the buffer zone of FC0. While minor differences exist for FC1, the present calculations rely solely on the operational context of FC0.
- An assumption of 12 hours at $t = 0$ is applied, reflecting the expectation that items departing at $t = 1$ are still accounted for half a day, as they remain in storage for approximately 12 hours during the initial time step.
- A uniform space cost per m^3 allocation across all articles is assumed. According to the Operational Manager of FC0, the variation between articles is negligible.
- Space usage is treated as linear and directly proportional to both product volume and duration in the buffer zone.
- When buffer stock levels increase, handling speed decreases significantly. This effect is not captured in the current model, even though it has a substantial impact on handling costs. Consequently, indirect costs associated with reduced efficiency are not reflected in the calculations.

5.2.3. Maximum Economical Retention Time

The final indicator considered is the MERT. Unlike the other indicators, this metric is not directly derived from existing literature but is introduced in this study, as it provides a valuable decision-support tool for storage management. The principle of MERT is to compare the additional handling and space costs of storing an item against the cost implications of discarding it and subsequently repurchasing it. This yields a break-even condition that indicates whether it is economically justifiable to retain the item in storage.

$$\frac{COGS_x}{C_{\text{handling},x} + d \cdot C_{\text{space},x}} < 1 \quad (5.14)$$

Rearranging the inequality and solving for d gives the expression for the MERT:

$$d > \frac{COGS_x - C_{\text{handling},x}}{C_{\text{space},x}} \quad (5.15)$$

where

$$\begin{aligned} COGS_x &= \text{Monetary value of article } x, \text{ lost if discarded [eur]} \\ C_{\text{handling},x} &= \text{Cost of handling of article } x \text{ [eur]} \\ C_{\text{space},x} &= \text{Daily space cost of article } x \text{ [eur/day]} \\ d &= \text{Maximum Economical Retention Time (MERT) [days]} \end{aligned}$$

The value of d therefore represents the maximum duration for which storing the product remains economically preferable to discarding it. Beyond this threshold, the cumulative storage and handling costs exceed the product's recoverable value. The subsection closes by outlining a number of additional remarks for completeness.

- Delivery costs and shopper picking costs are excluded, as they relate to sales activity rather than stock retention. Delivery costs (e.g., transport and labour wages) are driven by the number of items sold and remain unaffected by how long an item is stored. Likewise, shopper picking costs are only incurred once a product is sold, making them part of regular sales operations instead of retention decisions.
- Cost components (handling and space) are assumed to be stable and accurately quantified, which may not always reflect real-time variability.
- Handling cost is treated as a one-time fixed value, independent of the length of storage.
- Space cost is assumed constant per day and per unit of volume.

- The COGS used in the MERT formula is calculated as the average of the 14 most recent observations in the dataset, with the two lowest and two highest values excluded to mitigate the effect of extreme outliers.
- Similar to the space cost, the additional handling cost associated with a rise in stock levels is not captured by this KPI.

5.3. Main Findings Chapter 5

Several assumptions and limitations have been made in order to obtain results within the time frame of this thesis. It is therefore essential that any analysis or conclusion drawn from the dashboard takes these into account. The overview of assumptions and limitations for each KPI provided in this chapter is particularly valuable in this respect. Such transparency is crucial for a correct interpretation of the outcomes.

At the same time, the identified limitations provide a useful basis for developing recommendations for improvement. These recommendations will be discussed in detail in Chapter 9, where feedback from end-users gathered during the Evaluation phase (Part V) will play an important role.

In this chapter, all KPI calculations per campaign are presented. However, the final dashboard displays these KPIs across multiple levels. Table 5.3 provides an overview of all formulas for the KPIs included in the dashboard. For additional background on the calculations, refer to Section 5.1 and Section 5.2 in this chapter.

Table 5.3: All used formula's in the dashboard

KPI	Formula	Over- / Under-forecast
Per campaign		
Lost Margin	$LM_{x,y} = UQ_{x,y} LM_x$	Under-forecast
Waste	$WC_{x,y} = Q_{waste,x,y} \cdot COGS_x$	Over-forecast
Extended promo	$C_{ext} = (Q_{base} \cdot M_{reg}) - (Q_{promo} \cdot M_{promo})$	Over-forecast
Clearing promo	$C_{clear} = - (Q_{promo} \cdot M_{promo})$	Over-forecast
Handling Cost	$HC_{total,x} = \left((HT_{handling} + HT_{transport,x}) \cdot C_{shopper} \right) \cdot Q_{overstock}$	Over-forecast
Space Cost	$SC_{x,y} = TBT_{x,y} \cdot AV_x \cdot STC$	Over-forecast
Per item		
Lost Margin	$LM_x = FM_x + SC_x$	Under-forecast
Waste	$COGS_x$	Over-forecast
Extended promo	$C_{ext} = \frac{(Q_{base} \cdot M_{reg}) - (Q_{promo} \cdot M_{promo})}{Q_{overstock}}$	Over-forecast
Clearing promo	$C_{clear} = - \frac{Q_{promo} \cdot M_{promo}}{Q_{overstock}}$	Over-forecast
Handling Cost	$HC_{total,x} = (HT_{handling} + HT_{transport,x}) \cdot C_{shopper}$	Over-forecast
Space Cost	$SC_{x,y} = \frac{TBT_{x,y} \cdot AV_x \cdot STC}{Q_{overstock}}$	Over-forecast
MERT	$d > \frac{COGS_x - C_{handling,x}}{C_{space,x}}$	Over-forecast

In conclusion, this chapter addresses the research question of how indicator calculations can be developed to assess the cost implications of forecasting errors. The next chapter, Chapter 6, will discuss the dashboard design in greater detail.

IV

Demonstration

6

Dashboard Design

This chapter presents the design of the forecasting dashboard, the chosen artefact for this research. The dashboard translates the research problem outlined in Section 1.2 into a concrete, decision-supportive tool that improves promotional forecasting. It concludes Phase III, 'Design & Development' and addresses the Sub-question 4.1:

4.1 How should the artefact be designed to calculate and present the impacts of the indicators while serving as a supportive tool for product selection and the review of forecasted order quantities?

This chapter will start with presenting the rationale for designing the artefact in Section 6.1. The chapter then defines the requirements using the DSRP method of Johannesson and Perjons (2014) in Section 6.2.1. Next, Section 6.3 outlines the SQL model design and discusses general notion and limitations of the design. Section 6.4 concludes with an introduction to the dashboard design, which is discussed in greater detail in Chapter 7.

6.1. Rationale for the Design Approach

The artefact is designed on the premise that solutions should be derived from explicitly defined requirements rather than implicit assumptions. By first specifying functional and non-functional requirements, the dashboard is aligned with the explicated problem, stakeholder needs, and operational constraints, while remaining usable and adaptable in practice. This requirements-first approach enables iterative refinement: issues uncovered during testing can be traced back to specific requirements and systematically addressed. Johannesson and Perjons (2014) explicitly advocate this approach, framing "Define Requirements" as an extended problem explication that proceeds via two sub-activities, known as Outline Artefact and Elicit Requirements, to translate root causes and stakeholder input into a defensible specification for design (Johannesson & Perjons, 2014).

6.2. Artefact Requirements

This section translates the problem presented in Section 1.2 into testable preconditions and acceptance criteria for the artefact. To ensure the requirements are grounded in actual decision-making, key stakeholders will review and validate them through use of the artefact. The key stakeholders are presented in Table 7.7 in Section 7.4.

Starting with the definition of the requirements by specifying the functional and non-functional requirements that the artefact must meet (see Table 6.1). These requirements inform the development of the model and serve as a practical checklist during evaluation and communication phases. They enable an iterative process whereby identified shortcomings can be addressed. If necessary, the process can return to the Design & Development phase (Part III), or even further back to the Objectives of a Solution phase (Part II) if fundamental discrepancies are discovered. This iterative cycle continues until the design is validated and accepted by end-users, thereby ensuring both usability and reliability.

6.2.1. Requirements Definition

Within the Design Science Research Process (DSRP) framework, introduced in Section 1.6, the second step of the design phase begins with defining the requirements (Johannesson & Perjons, 2014). In this study, this step was deliberately postponed compared to the general sequence of the DSRP framework, as developing a

well-defined overview of requirements first requires sufficient prior knowledge. The objective of this step is to determine the type of artefact capable of addressing the problem outlined in Section 1.2 and to identify the essential requirements that the artefact must fulfil in order to meet stakeholder expectations. As Johannesson and Perjons (2014), properly summarised in a sentence:

“What artefact can be a solution for the explicated problem and which requirements on this artefact are important for the stakeholders?”

In this step, the problem is reconsidered with the potential solution in mind. The goal is to understand what the solution should look like and what it needs to achieve. These needs are called requirements, important features or conditions that the solution must meet. Requirements can be about what the artefact should do (functions), how it should be built (structure), where or how it will be used (context), or what effects it should have. The outcome of this activity is a precise conceptualisation of the artefact and a list of critical requirements, informed by the earlier problem analysis, academic literature, stakeholder input, and practical experience (Johannesson & Perjons, 2014).

The activity of defining the requirements follows the approach of Johannesson and Perjons (2014), which is structured and visualised in Figure 6.1. As shown in Figure 6.1, the process consists of two main steps: (1) outlining the artefact and (2) eliciting the requirements.




Figure 6.1: Define Requirements structure (Johannesson & Perjons, 2014).

6.2.2. Outline Artefact

The chosen artefact is an instantiation, defined as a working system that can be applied in practice (Johannesson & Perjons, 2014). Specifically, it takes the form of a decision-support dashboard that utilises historical operational data and was developed in consultation with key stakeholders. The artefact translates operational information into actionable insights regarding the financial consequences of promotional forecasting errors. It allows stakeholders to evaluate the performance of past campaigns, identify over- or under-forecasting impacts, detect systematic patterns, and make better-informed decisions for future promotions.

The decision to employ a dashboard as the artefact is supported by four primary considerations, which justify its suitability in the context of this research:

1. Need for Decision-Relevant Integration

Forecasting errors have multifaceted consequences: over-forecasting generates waste, forced price reductions, handling, and space occupancy costs, while under-forecasting results in lost sales, reduced service levels, and reputational effects. Existing performance metrics within Picnic are fragmented across operational and financial systems and do not directly support article-level trade-off analysis. A dashboard is well-suited to this context because it can integrate disparate indicators into a coherent, interactive framework, enabling stakeholders to simultaneously assess the financial and operational implications of forecasting decisions.

2. Facilitating Trade-off Analysis

A core requirement of the artefact is to evaluate over- and under-forecasting in a comparable and decision-relevant manner. Assigning monetary values to forecasting errors allows direct comparison of the costs on both sides.

A dashboard is the ideal medium for such trade-off visualisation because it provides immediate, interpretable feedback, supports scenario exploration, and highlights which imbalances carry the most significant consequences. This interactive capability is essential for decision-making under uncertainty, as it allows assumptions to be adjusted dynamically and their effects on the trade-offs to be observed in real time.

3. Transparency and Actionability

A key rationale for using a dashboard is that it promotes transparency and actionability. By decomposing costs into identifiable components, such as waste, lost sales, handling, and space occupancy, decision-makers gain a clear understanding of the drivers behind each article's position on the trade-off curve. This transparency is critical for operational adoption, as it establishes trust in the indicators and facilitates evidence-based decisions. Actionability is achieved through the dashboard's interactivity: assumptions can be modified, and alternative strategies can be evaluated immediately, providing operational teams with a tool to test and implement decisions efficiently.

4. Flexibility and Extensibility

Finally, the dashboard format is chosen for its inherent flexibility. The artefact is designed to be extendable: new KPIs can be added, and existing indicators can be adapted as operations evolve or additional insights are discovered. This ensures that the tool remains relevant over time, accommodating process improvements, changes in operational policy, or expanded scope of analysis. The ability to evolve alongside operational realities strengthens the artefact's long-term utility and aligns with the research's goal of providing sustainable decision support.

In conclusion, the dashboard is an artefact that suits very well for this study because it (1) integrates fragmented performance metrics into a decision-relevant framework, (2) enables explicit trade-off analysis between over- and under-forecasting, (3) promotes transparency and actionability for stakeholders, and (4) offers flexibility for future adaptation. These characteristics ensure that the artefact not only fulfils the immediate research objectives but also provides a practical, scalable tool for ongoing operational decision support.

6.2.3. Elicit Requirements

Drawing on the earlier problem analysis, academic literature, stakeholder input, and practical experience, a well-grounded set of requirements has been established to address the identified problem, as presented in Table 6.1. These requirements are divided into two categories: functional and non-functional requirements. These requirements serve different purposes in guiding the design of an artefact. Functional requirements describe the core tasks the artefact must perform in order to solve the identified problem and meet stakeholder needs. These are often concrete and context-specific. In contrast, non-functional requirements define how the artefact should operate or be structured, regardless of its specific use cases. They are typically more generic and fall into two categories: structural requirements, which concern the internal organisation of the artefact (e.g., modular design, clean code), and environmental requirements, which deal with the relationship between the artefact and its surroundings (e.g., system compatibility, runtime performance). Together, these requirements ensure that the artefact is not only functional but also usable, maintainable, and effective in practice (Johannesson & Perjons, 2014).

Table 6.1: Overview of the functional and non-functional requirements.

#	Name and (Requirement Type)	Requirements	Reasoning	Quality
Functional Requirements				
1.	Display of Selected KPIs (Functional)	The dashboard must display the six selected key performance indicators (KPIs) as defined in Chapter 4, each reflecting the financial impact of promotional forecasting errors.	Displaying these KPIs is the primary function of the dashboard and directly answers the research sub-question. These indicators represent the core output of the analysis and provide actionable insights into the consequences of error forecasting.	Suitability
2.	Multi-Level KPI Aggregation (Functional)	The dashboard must calculate and present each KPI at two aggregation levels: o Absolute Impact KPI per campaign o KPI Impact per unit over- or under-forecasted	Multiple levels of aggregation allow for richer and more flexible interpretation of results. It enables stakeholders to view both the scale and intensity of forecasting errors, and ensures insights are meaningful across various use cases.	Effectiveness
3.	Comparison of Over- vs Under-Forecasting (Functional)	The dashboard must allow for direct comparison between over- and under-forecasting per article by calculating the total cost of over- and under-forecasting based on the identified KPIs. Thereby highlighting asymmetries in forecasting impact per article.	Forecasting errors do not affect performance symmetrically. Highlighting cost differences between over- and under-forecasting helps Picnic prioritise improvements where the operational and financial consequences are most severe.	Comprehensibility
4.	KPI Analysis at Multiple Granularities (Functional)	The dashboard must offer KPI analysis at multiple granularities, specifically at the article, promotion, and fulfilment centre levels, to support operational decision-making across different layers of the supply chain.	Due to the table structuring within Picnic the granularity; Article, Promotion, Fulfilment centres is needed to have correct output.	Customisability
Non-Functional Requirements				
1.	Modular Dashboard Architecture (Structural)	The dashboard must be built with a modular structure so that individual components (e.g. KPI calculations, filters, visual elements) can be updated or replaced independently without disrupting the entire system.	Modularity increases maintainability and enables future extensions (e.g. adding new KPIs) without significant redesign. This also aligns with software engineering best practices.	Modularity
2.	Consistent Layout and Visual Hierarchy (Structural)	The dashboard design must ensure consistent layout and visual hierarchy, allowing users to quickly identify key insights and navigate between sections intuitively.	A coherent and orderly structure enhances the user's ability to extract relevant insights efficiently. It reduces cognitive load, minimises confusion, and makes the dashboard more approachable for stakeholders across departments.	Coherence
3.	Clean and Well-Documented Code (Structural)	The system must be developed using clean, well-documented code to ensure ease of debugging, validation, and future development. It should be easily understandable by other analysts or engineers without the need for additional explanation.	Well-structured code is essential for long-term maintenance. Clear documentation enables future analysts or developers to modify or expand the system without introducing errors or relying on the original creator.	Maintainability
4.	Appropriate Visualisation Types for KPIs (Structural)	Each KPI must be visualised using the most appropriate view, for its data characteristics, to ensure clarity, readability and interpretability of results.	Proper visualisation design directly affects the usability and effectiveness of the dashboard. By tailoring visual types to data types, stakeholders can interpret results more accurately and make faster decisions.	Suitability
5.	Cross-Device and Browser Compatibility (Environmental)	The dashboard must be accessible via standard office devices (e.g., desktop or laptop) and be compatible with commonly used browsers and operating systems at Picnic.	Ensuring accessibility across common devices and platforms maximises adoption across the organisation, avoids additional infrastructure costs, and enables users to access insights without technical barriers.	Accessibility
6.	Maximum 3-Minute Execution Time (Environmental)	The code must have a maximum execution time of 3 minutes to ensure timely data processing and usability during regular operations.	Timely data output is crucial for usability, particularly in dynamic operational contexts such as promotional planning. Excessive loading times reduce user engagement and hinder decision-making speed.	Efficiency
7.	User-Friendly Navigation and Accessibility (Environmental)	The dashboard must be intuitive and easy to navigate, allowing users with limited technical background to access, filter, and interpret KPI data without requiring training or assistance.	Usability is critical to ensure adoption and correct interpretation by business users and analysts alike. Your artefact should not require technical explanation to be effective.	Usability

The defined requirements provide the foundation for designing the dashboard, ensuring that both functional capabilities and non-functional qualities are aligned with stakeholder needs. These requirements will guide the development process and later serve as reference points during the Evaluation Phase. Ultimately, they enable an iterative cycle in which end-users actively engage with the artefact, providing feedback that allows the design to be refined and validated in practice.

6.3. SQL Model Design

This section outlines the design of code within SQL and the final dashboard itself.

6.3.1. Chronological Process of the Design

The foundation of the dashboard is a codebase written in 'Structured Query Language' (SQL). Within Picnic, all dashboards are developed in this manner. Picnic stores all available data in its central Data Warehouse (DWH). These tables can be directly accessed via SQL; however, the creation of a new dashboard typically requires the derivation of additional data. For this reason, supplementary SQL code must be developed to generate the desired outputs. The model is structured in such a way that each step corresponds to a new attribute or KPI. The development of several KPIs necessitated the use of multiple Common Table Expressions (CTEs), which were therefore organised into sub-steps.

The chronological approach to the construction of this dashboard was as follows:

1. Identification of all relevant attributes and KPIs.

2. Development of the attributes:
 - (a) Establish a logical chronological order for the attributes to be developed.
 - (b) Identify which tables are required to retrieve the necessary data.
 - (c) Develop a CTE for each attribute, combining multiple attributes where feasible.
 - (d) Test each CTE to ensure the output is correct.
3. Development of the KPIs:
 - (a) Determine the order in which the KPIs should be developed.
 - (b) Identify the tables required to source the data.
 - (c) Define the method for calculating the desired metrics.
 - (d) Develop each calculation to the level of absolute costs per campaign and costs per individual article.
 - (e) Test each KPI within SQL to ensure the output is correct
4. Determine which columns should be visible in the final dashboard and incorporate these into the final SELECT statement.
5. Export the raw data from the table to Google Spreadsheets and construct a user-friendly dashboard interface.

6.3.2. General Notions on the Design

To fully understand the construction of the model, several key design considerations are discussed in this subsection. A direct link to the code can be found in Appendix A.

Granularity

The data granularity is defined at the levels of article, campaign, and site. The site dimension (i.e. fulfilment centre location) was included because several KPIs, such as extension, waste, and handling, are directly influenced by site-level operations. In the final 'SELECT' statement of the SQL code, data was aggregated at three levels:

- Article - campaign - FC0
- Article - campaign - FC1
- Article - campaign

This structure provides insight into performance differences across fulfilment centres while also enabling campaign-level analysis.

Clearly Naming and Commenting

All variables and expressions in the SQL code were named clearly, enabling subsequent users to understand the processes being executed. In addition, explanatory comments were inserted at complex sections of the code to further improve readability and transparency.

Room for Future Extensions

The development of a dashboard is, by its nature, never complete. To accommodate this, the model has been designed with flexibility in mind, allowing for future adjustments or enhancements to be implemented.

6.4. The Design of the Dashboard

This section highlights the design of the dashboard, beginning with a clear explanation of its use in Section 6.4.1, where the design principles and structure are outlined, supported by illustrations. In Section 6.4.2, the metrics and measurement details are discussed.

6.4.1. Use of the Dashboard

This section outlines the functionality and operational use of the dashboard, explaining its design, rules of usage, and its integration within Picnic's processes. The aim is to provide clarity on how the dashboard is structured, how it fits into the process, and how it can be used effectively for promotional forecasting.

Before discussing the design process, it is useful to refer back to the forecasting process mapped in Chapter 3, which is illustrated in Figure 3.2. The dashboard is designed to support two key stages within this process. First, it assists in the selection of products for promotion at the very beginning of the process. Second, it can be used to support reviewing the forecasts generated by the algorithm. All forecast quantities that appear risky or unusual can be examined using the dashboard, which visualises historical data to indicate what quantities can typically be expected. A more detailed explanation of how the dashboard is integrated into the process, as well as the broader process adjustments, is provided in Section 8.1.

Dashboard Design Principles

At Picnic, the use of dashboards is governed by the principle that no coding should be embedded within the dashboard itself. This principle ensures sustainability, as it avoids the complexity and potential errors introduced by internal calculations that are hard to track or maintain. If code were embedded within the dashboard, there would be a risk that modifications could lead to inconsistencies with the associated 'Google Spreadsheet' files or other linked resources. As such, the dashboard is designed to be straightforward and transparent, with no internal calculations present, ensuring its clarity and ease of use.

Dashboard Design Structure

The dashboard structure is intentionally simple, comprising five distinct sheets, each designed to serve a specific purpose. This subsection provides a clear discussion and visualisation of each sheet.

- **Explanation of the Dashboard:** This tab provides an overview of the dashboard's purpose and functionality. It clearly outlines how the tool should be used, ensuring that users understand its intended role within the forecasting process. The sheet also serves as a guide to the proper application of the dashboard, making it clear how the insights can be leveraged effectively. Figure 6.2 clearly represents the sheet in 'Google spreadsheet'.




Figure 6.2: Overview 'Explanation of the Dashboard' sheet

- **The Dashboard:** This is the main sheet consolidating all relevant data, supported by intuitive filters that allow users to tailor the analysis to their needs. The layout is organised into distinct sections, enabling direct comparison of how key performance indicators (KPIs) respond to different data selections.

The sections include:

- General information: date, site, brand, name
- Article attributes,
- General Metrics,
- Aggregated Cost of KPIs,
- KPIs at campaign level,
- KPIs at article level.

And so this overview provides detailed information about costs incurred or avoided due to forecasting errors, both at the campaign level and the site-specific level. It provides valuable insights into the impact of these errors, allowing for a clear understanding of the financial consequences associated with over- or under-forecasting. Figure 6.3 and Figure 6.4 offer a visual summary of the sections outlined above.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	
2		Promo week	Article_id	Brand	Article name	Site										
3						ALL										
4																
5																
6		Week	Article				Attributes				General Campaign / Article Metric					
7		Promo week	Article ID	Site	Brand	Article name	Temperature Zone	Best Before Date	Mover category	Assortment continuity	Unavailable products	Promotion sales	Promo GM	Total margin	Overstock	
8		Integer	Integer	Str	Str	Str	Str	Days	Str	Str	Integer	Integer	€	€	€	
9	5	202517	12029616	ALL	Paic	Paic-Liquide Vaisselle bi-activ citron 2 x 75 cl	ambiant	No Mover	In/Out	18	164	1.74	286.9	0		
10	10	202515	12029586	ALL	Carle Noire	Carle Noir Café moulu pur arabica 3 x 250 g	ambiant	No Mover	In/Out	1473	312	1.44	448.34	0		
11	11	202443	1339694	ALL	Révilion	Révilion Papillotes lait 360 g	ambiant	120	No Mover	Temporary	0	71	149	105.72	0	
12	12	202443	1339657	ALL	Lindt	Lindt Calendrier de l'avent Lindor 300 g	ambiant	110	No Mover	Temporary	0	46	3	138.28	0	
13	13	202517	12029615	ALL	Paic	Paic Liquide vaisselle bi-activ 2 x 75 cl	ambiant	No Mover	In/Out	0	223	74	387.29	0		
14	14	202443	1339677	ALL	Nestlé	Nestlé Mouline rennes crunch 88 g	ambiant	80	No Mover	Temporary	0	97	0.39	37.72	0	
15	15	202443	1339708	ALL	Nestlé	Nestlé Smarties moulinés péris noisettes 88 g	ambiant	80	No Mover	Temporary	89	37	14.45	0		
16	16	202511	1330553	ALL	Petit Navire	Petit Navire Thor enfariné à la bûche de pain 3 x 100 g	ambiant	140	No Mover	In/Out	136	43	0.62	39.47	0	
17	17	202520	12029541	ALL	Petit Navire	Petit Navire Thor enfariné au naturel 4 x 60 g	ambiant	No Mover	In/Out	297	262	1.2	469.71	0		
18	18	202443	1339657	ALL	Lindt	Lindt Assortiment champs lysées 127 g	ambiant	110	Slow Mover	Temporary	195	48	1.87	89.88	0	
19	19	202443	1339657	ALL	Lindt	Lindt Assortiment champs lysées 127 g	ambiant	110	Slow Mover	Temporary	195	48	1.87	89.88	0	
20	20	202512	12029512	ALL	McCaïn	McCaïn Mini mozzarella sticks 124 pièces	frozen	20	Medium Mover	Regular	264	1029	1.38	1380.24	0	
21	21	202515	12029558	ALL	M. Chapoulier	M. Chapoulier La croix des grives côtes du rhône AOC 75 cl	chilled	6	No Mover	In/Out	176	38	2.07	78.83	0	
22	22	202441	10402717	ALL	La Laiterie	La Laiterie Feuilleté de mousse vanille 4 x 57 g	chilled	6	Medium Mover	Regular	1544	970	0.24	225.56	0	
23	23	202502	10008421	ALL	Calin	Calin Fromage blanc 0% mg 8 x 100 g	chilled	8	Medium Mover	Regular	776	26	0.77	20.48	0	
24	24	202517	13396069	ALL	Hollywood	Hollywood 25cent menthe fraîche & forte 3 x 10 dragées	ambiant	260	Slow Mover	Regular	425	23	0.86	19.71	0	
25	25	202507	11817774	ALL	Isigny Ste Mère	Isigny Ste Mère Beurre de baratte sel de guérande 250 g	chilled	18	Medium Mover	Regular	618	657	0.59	386.67	0	
26	26	202517	12029610	ALL	Andros	Andros Pur jus d'orange 1.5 L	chilled	13	No Mover	In/Out	655	204	0.55	113.09	0	
27	27	202506	11244135	ALL	Dauzat	Dauzat Salade caesar 320 g	chilled	3	No Mover	Temporary	58	96	1.03	110.59	0	
28	28	202445	10089387	ALL	Purina One	Purina One Croquettes chat stérilisé saumon 1.5 Kg	ambiant	500	Fast Mover	Regular	308	1009	1.17	1128.21	0	
29	29	202515	11563458	ALL	Bini	Bini Petits feuilletés 40 unités	frozen	30	No Mover	In/Out	286	115	1.26	142.97	0	
30	30	202504	11222913	ALL	Orly	Orly Bouillon végétarien avoine 1 L	chilled	140	Medium Mover	Regular	513	467	0.7	325.58	0	
31	31	202516	10008248	ALL	Suchard	Suchard Rochers lait 7 x 35 g	ambiant	110	Medium Mover	Regular	458	677	0.78	434.86	0	
32	32	202507	10008495	ALL	Bonne Maman	Bonne Maman Petits cookies pépites chocolat 250 g	ambiant	159	Medium Mover	Regular	404	1001	0.88	832.81	0	
33	33	202511	10008451	ALL	Nestlé	Nestlé Yaos à la vanille 125 g	chilled	6	Medium Mover	Regular	294	747	1.21	890.66	0	
34	34	202504	11611554	ALL	Herta	Herta Bacon	chilled	9	Medium Mover	Regular	705	922	0.5	456.55	0	
35	35	202519	11353880	ALL	Leroux	Leroux Chic	chilled	508	Slow Mover	In/Out	251	124	1.4	173.78	0	
36	36	202449	11563460	ALL	Nestlé	Nestlé Bûche	chilled	20	Slow Mover	Temporary	298	90	1.3	105.29	0	
37	37	202514	11368987	ALL	BN	BN Biscuits	chilled	142	Medium Mover	Regular	805	719	0.44	310.2	0	
38	38	202502	10509562	ALL	Madame Loik	Madame Loik Fromage nature sel de mer 150 g	chilled	0	Medium Mover	Regular	581	307	0.8	208.7	0	
39	39	202518	11271957	ALL	Ele & Vire	Ele & Vire Crème Ruide entière 300 g	chilled	19	No Mover	In/Out	343	51	1.02	51.93	0	
40	40	202512	11246189	ALL	Lune de miel	Lune de miel Miel de fleur liquide 100 g	ambiant	260	Medium Mover	Regular	154	182	2.26	409.05	0	
41	41	202516	10353287	ALL	Oasis	Oasis Oasist tropical 100 unités	frozen	19	Medium Mover	Temporary	483	663	0.73	460.45	0	
42	42	202514	12029534	ALL	Tuliac	Tuliac Côtes de bourg ap 75 cl	ambiant	0	No Mover	In/Out	189	50	1.85	92.55	0	
43	43	202516	11271954	ALL	Florette	Florette Mâche 125 g	chilled	3	Fast Mover	Regular	53	909	0.99	989.29	0	
44	44	202512	11367441	ALL	Bledina	Bledina Morceaux caisses jaunes haricots verts patate douce	ambiant	130	No Mover	In/Out	1481	10	0.24	2.35	0	
45	45	202504	11818287	ALL	Montagne Noire	Montagne Noire Saussisse sèche 225 g	chilled	13	Medium Mover	Regular	246	1316	1.42	1655.49	0	
46	46	202514	11568195	ALL	Vico	Vico Curry doux avec fenouil 80 g	ambiant	55	Medium Mover	Regular	991	114	0.32	34.53	0	
47	47	202512	10008445	ALL	Volvic	Volvic Eau minérale 200 cl	chilled	320	Medium Mover	Regular	128	126	1.26	789.37	0	
48	48	202441	11240242	ALL	La Laiterie	La Laiterie 2% au lait hum 144 x 115 g	chilled	6	No Mover	In/Out	2276	104	0.14	13.53	0	
49	49	202508	11653920	ALL	Haribo	Haribo Tribali mega boites 1 kg	ambiant	6	No Mover	In/Out	408	99	0.85	82.44	0	
50	50	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
51	51	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
52	52	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
53	53	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
54	54	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
55	55	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
56	56	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
57	57	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
58	58	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
59	59	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
60	60	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
61	61	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
62	62	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
63	63	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
64	64	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
65	65	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
66	66	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
67	67	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
68	68	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
69	69	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
70	70	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
71	71	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
72	72	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
73	73	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
74	74	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
75	75	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
76	76	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
77	77	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
78	78	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
79	79	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
80	80	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
81	81	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
82	82	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
83	83	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
84	84	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
85	85	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
86	86	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
87	87	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
88	88	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
89	89	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
90	90	202509	11653605	ALL	BN	BN Biscuits saumon chocolat 3 unités	ambiant	6	No Mover	Regular	666	43	0.52	22.29	0	
91	91	202509	11653605	ALL</												

Figure 6.3: Overview 1 'Forecast Impact Dashboard' sheet, column A to O.

[illegible]

Figure 6.4: Overview 2 'Forecast Impact Dashboard' sheet, column P to AI.

- **Relevant Graphs for Quick Analysis:** This sheet is dedicated to the most relevant visualisations, aimed at providing quick insights during the product forecasting process. The graphs display historical data on over-forecasting and under-forecasting costs for the selected article. The accompanying candlestick charts provide essential insights into the distribution of the entire dataset, allowing users to understand the relative impact of over- or under-forecasts across the whole data history. These visualisations are

designed to assist users in understanding the financial impact of forecasting errors swiftly and clearly, thereby supporting faster decision-making.



Figure 6.5: Overview of the 'Graphical overview' sheet.

- **Recommendations to Upgrade the Dashboard:** This sheet allows users to suggest potential upgrades or modifications to the dashboard and its code. It provides a central location for users to submit their feedback, which can then be reviewed by product owners for future implementation. Users are encouraged to explain the issue they are facing and identify themselves, so that product owners can follow up and engage in discussion if necessary. Furthermore, this sheet allows product owners to track the progress of each suggestion or request, providing transparency in the improvement process. The tab is visualised in Figure 6.6

into future dashboard developments. The detailed rationale and calculation behind this approximation are provided in Appendix F.

6.5. Main Findings Chapter 6

This chapter demonstrates how the artefact, the forecasting impact dashboard, was systematically designed to calculate and visualise indicator impacts while serving as a decision-supportive tool for product selection and forecast quantity review. The design was grounded in explicitly defined functional and non-functional requirements, ensuring that the dashboard not only performs the intended calculations but also remains transparent, user-friendly, and adaptable to future developments.

By integrating financial and operational KPIs within a modular SQL-based structure, the dashboard translates complex data into clear, actionable insights. Its design allows users to assess the trade-offs between over- and under-forecasting, evaluate product suitability for promotions, and review forecasted order quantities in a structured and interpretable way.

Furthermore, the dashboard design prioritises usability and sustainability, with clear documentation, consistent layout, and flexible architecture enabling continuous improvement. While the current version operates deterministically, the chapter recognises the potential value of introducing stochastic elements to represent uncertainty in demand and operational outcomes, thereby enhancing the artefact's analytical depth in future iterations.

V

Evaluation

Dashboard Results and Evaluation

Having established the theoretical foundations, operational indicators, and dashboard presentation in the previous chapters, this chapter turns to the practical evaluation of the developed dashboard. The focus lies on analysing the insights derived from the dashboard results, examining how these insights can support decision-making in the promotional forecasting process, and evaluating the usefulness of the tool from an end-user perspective. In doing so, the chapter aims to answer the following research question:

4.2 What insights can be derived from the dashboard results, in what ways can these insights support product selection and forecast quantity review, and how do end-users evaluate its usefulness?

The first part of this chapter (Section 7.1) presents general findings derived from the dashboard, providing context on the overall distribution and behaviour of the key indicators. Subsequently, Section 7.2 describes the setup of the pilot case study. The results of the pilot case study are then discussed in Section 7.3, where the dashboard's practical application is demonstrated for two real product cases. Finally, Section 7.4 summarises the end-user evaluation, covering requirement validation, and general feedback on usability, and perceived value. The chapter concludes with the main findings and reflections in Section 7.5.

7.1. General Results Gained from the Dashboard

This section provides broader insights derived from the dashboard, offering valuable context for both product selection and forecast quantity evaluation. Such knowledge helps to interpret, for example, what constitutes a relatively high or low lost margin per item, the typical magnitude of over- and under-forecasting costs, and what can be considered a normal range for over-forecast indicator costs. The figures and results presented in the pilot case study in Section 7.2 gain meaning only when viewed in relation to the overall dataset. Therefore, this section presents aggregated insights and explains the candlestick charts shown in the figures below. These charts are also integrated into the developed dashboard (as explained in Section 6.4) to support comparative analysis and interpretation.

Candlestick: Total Under- and Over-Forecast per Item

The candlestick chart of total under- and over-forecast costs per item (shown on the left in Figure 7.1) displays a clear difference in distribution. The median value of under-forecast costs lies slightly above zero, with a relatively narrow interquartile range, indicating that for most articles, the financial impact of under-forecasting is limited. However, the upper whisker extends to around €15 per item, reflecting a few exceptional cases with substantial missed-margin costs, such as foie gras, Pommery Champagne, or Oral-B toothbrushes.

In contrast, over-forecast costs are generally close to zero for most items, yet the long lower whisker reveals that several campaigns showed negative over-forecast impacts—representing instances of profit. These cases arise from extended promotions where the additional sales generated at a reduced promotional margin outweighed the profits that would have been achieved under regular pricing margins. The interquartile range is very narrow, primarily due to the strong concentration of values close to zero, driven by the limited financial impact of the clearing, handling, and space indicators, as shown in Figure 7.2 and discussed in the paragraph about on the individual indicator costs per item.

Candlestick: Total Under- and Over-Forecast per Campaign

At the campaign level, shown on the right of Figure 7.1, the candlestick charts display a similar overall shape; however, the variation in costs becomes more pronounced. The median under-forecast cost per campaign is positive and notably higher than that of over-forecasting, confirming that underestimation generally entails a greater financial impact. While most campaigns show limited losses, the upper whisker extends beyond €2,000, indicating that in certain cases, under-forecasting resulted in substantial lost margins. Over-forecasting, by contrast, exhibits lower median costs and smaller variation, and—as discussed earlier—includes instances of negative cost impact, reflecting campaigns where the effects of a promotional extension ultimately generated profit.




Figure 7.1: Candlestick of over- and under-forecast impact per item and per campaign

Candlestick: Each Indicator Cost per Item

The chart shown on the left of Figure 7.2 provides a breakdown of the cost distribution for each KPI indicator. The Lost Margin indicator shows a slightly positive median, confirming that missed sales opportunities represent a moderate cost component. The Extended Promotion cost is the only indicator with a negative median, indicating that extended promotions can occasionally recover part of the overstock-related losses. The Clearing Promotion cost per item is very low and shows no notable outliers. Waste costs also display a relatively wide spread, with the upper whisker reaching nearly €7 per item, suggesting that in certain cases, unsold products result in substantial financial loss per item. The Handling and Space costs remain close to zero for most items, indicating that their operational impact per item is limited in relative terms.

Candlestick: MERT Days

The candlestick chart of MERT days reveals an extraordinarily wide distribution. Although the theoretical calculation yields a median MERT of around 2,100 days (approximately 5.7 years), with the interquartile range stretching from 1,300 to 4,000 days (roughly 3.5 to 10.5 years), these figures are purely academic. As this MERT is a calculation result, it fails to account for practical constraints such as physical durability, increasing market uncertainty, and operational feasibility over extended periods. Therefore, while the calculation suggests a long economically viable storage tenure for many articles, in practice at Picnic, a maximum MERT value of two years will be considered for all subsequent operational analyses.

As previously noted, the chart reveals several extreme outliers, corresponding to very small and expensive products such as foie gras or 30 ml perfume samples, which have exceptionally low daily space costs and therefore result in an upper whisker extending up to 100,000 days. This considerable variation illustrates that storage viability differs greatly across product categories and that MERT-based evaluations should therefore always be interpreted on a product-specific basis. In conclusion, only a very small number of articles have a

MERT value that is operationally relevant and thus storing products is always a good idea.

For the case study presented in Section 7.2.2, the MERT indicator will be excluded from further analysis. There are two main reasons for this decision. First, the indicator has limited relevance to the two focal areas of the promotional forecasting process, product selection and forecast quantity review, where the dashboard is intended to provide support. The MERT KPI is a more independent measure that could become valuable in future applications, particularly when evaluating whether it is more economical to store or discard excess products. Second, the distribution of the indicator’s outcomes shows that it consistently provides the same insight: in nearly all cases, storing the product in the buffer zone is economically viable. As such, incorporating this indicator into the case study analysis would add little additional value. That said, the indicator’s relevance may increase in the future with the introduction of new products or changing assortment characteristics.




Figure 7.2: Candlestick of each Indicator separate

7.2. Setup of Pilot Case Study

This section introduces the pilot case study that will be conducted in Section 7.2.2. It outlines the two selected articles and explains why they provide a representative basis for demonstrating the practical application of the dashboard. In short, these articles were chosen to illustrate the dashboard’s ability to capture contrasting forecasting results and to demonstrate the types of insights it can generate.

7.2.1. Presentation of Pilot Case Study Articles

For the pilot case study, two articles were selected: one with a history of forecasting errors with an under-estimation across multiple promotions (case 1), and one with a history of both an under-estimation in one promotion and an over-estimation error in another (case 2). The dataset used for the pilot includes all promotions that took place within Picnic between week 40 of 2024 and week 23 of 2025, covering a period of less than a year. Table 7.1 presents the key characteristics of the two case study articles.

Table 7.1: Overview of the two articles used for the pilot case study

	Article 1	Article 2
ID	XXX	XXX
Name	XXX 1	XXX 2
Brand	XXX 1A	XXX 2A
Temperature Zone	Ambient	Ambient
Best Before Date	65	200
Mover Category	Medium Mover	Medium/Fast Mover
Assortment Continuity	Regular	Regular

A brief explanation of Table 7.1 is provided here to clarify the meaning of the article-specific information and attributes presented. First, the article-specific information is described: the Article ID is simply the unique numerical identifier assigned to each article, while the name and brand are self-explanatory.

The temperature zone indicates the storage conditions within the FC, with three categories: ambient, chilled, and frozen. The operational process differs per zone, therefore difference in cost per zone will occur. However, as noted earlier in this study, operational costs are currently based on the logistical processes in the ambient zone; consequently, the focus of this research is on this category.

The BBD refers to the product's expiry date. This attribute can provide insights into waste costs or the MERT, as shorter BBDs often lead to higher waste, while the MERT reflects the economically viable duration for storing products in the buffer. In many cases, the BBD is shorter than the MERT, so the MERT may not be very relevant. Therefore, it is important to consider the BBD when making decisions about the MERT.

The mover category indicates the popularity of a product, determined by the average weekly sales over the four weeks preceding the promotion. The categories are defined as follows:

- No mover: products with an average of 0 sales in the four weeks prior to the promotion; typically in/out products.
- Slow mover: average sales of 0–10 units per week.
- Medium mover: average sales of 10–100 units per week.
- Fast mover: average sales exceeding 100 units per week.

Understanding product popularity is useful, for instance, to estimate the risk of products remaining in the buffer for extended periods following an overestimation.

The final attribute, assortment continuity, reflects the presence of a product within the assortment. Regular products are always part of the assortment; temporary products are included for a limited period, such as barbecue items in summer or Christmas products; and in/out products are included solely for the purpose of a promotion.

Article 1 – XXX 1

This article belongs to a popular brand with consistently high sales. However, the selected variant, a more unique item with a XXX flavour, proved more difficult to forecast accurately. Despite the strength of the parent brand, the distinctive nature of this product line led to forecasting errors across four separate promotions within a nine-month period. This makes the article a particularly interesting case to demonstrate how the dashboard can provide insights into these historical events and serve as a supportive tool for the next promotion.

Article 2 – XXX 2

This product is part of Picnic's regular assortment, and in principle, regular articles should be easier to forecast due to the availability of historical sales data. However, promotional effects can still make forecasting challenging, as this article demonstrates. The product falls between a medium and fast mover, with weekly sales generally around 100 units; during one promotion it behaved as a fast mover, while in another it was classified as medium. The Best Before Date (BBD) is relatively long (200 days), so the risk of waste is low, especially considering that the product is a medium/fast mover. As with Article 1, this product is stored in the ambient zone, meaning that operational cost calculations from this research are consistent with the processes applied to these types of products.

The setup of the pilot case study has now been established; the subsequent analysis is presented in Section 7.2.2.

7.2.2. The Pilot Case Study

The purpose of this section is to demonstrate the analytical potential of the newly developed dashboard by applying it to two selected articles. This section explains how the dashboard values can be interpreted and how they can support decision-making in selecting products for promotion or in evaluating the forecast quantities generated by the model.

The section concludes, in Section 7.5, with a reflection on the overall strengths, limitations, and opportunities for improvement of the dashboard, thereby laying the groundwork for the study’s final conclusions.

7.2.3. Analytical Scope and Indicator Relevance

Before presenting the results for the two articles introduced in Section 7.2.1, this subsection outlines all metrics and indicators included in the dashboard. For each metric and indicator, a brief explanation is provided, along with an assessment of its relevance for the two main purposes of the dashboard: product selection and forecast quantity review.

To clarify the differences in relevance: during product selection, the key consideration is whether a product is worth the associated risk. Metrics such as margin and gross margin per item are therefore particularly important. In addition, the impact of forecasting errors plays a crucial role—low impact levels may support taking the risk of promotion, while high impact levels would discourage it. When the dashboard is used for forecast quantity review, the most relevant indicators are those reflecting the effects of over- and under-forecasting on item level. These indicators provide a basis for adjusting forecast quantities to mitigate risk. Furthermore, general metrics such as unavailability and overstock help identify where forecast quantities have historically been inaccurate—either too high or too low. For instance, identifying products that have been consistently under-forecasted may justify increasing their order quantities to prevent unavailabilities.

The relevance levels in Table 7.2 are categorised as follows:

- High relevance
- Medium relevance
- Low relevance

Table 7.2: Overview of all indicators used in the dashboard, ranked by their relevance for product selection and forecast quantity review.

Metric Category	Metric	Description	Product Selection	Quantity Review
General Metrics	Unavailability	Number of items that could have been sold if available in the app, including conversion to purchase and commonality effects.	Medium	High
	Promotion sales	Total quantity sold during the promotion.	High	Low
	Promo GMI	Gross margin per item during the promotion.	Medium	Medium
	Total margin campaign	Absolute gross margin generated during the campaign.	High	Low
	Overstock	Quantity of products in surplus on the first day after the promotion, calculated as buffer stock on day 1 minus buffer stock two weeks before the promotion.	Medium	High
Cumulative Forecasting Impact	Total under-forecast cost per item	Lost margin per item.	Low	High
	Total over-forecast cost per item	Sum of costs per item for waste, extended promotions, handling, and space.	Low	High
	Total under-forecast cost per campaign	Lost margin per campaign.	High	Medium
	Total over-forecast cost per campaign	Sum of costs per campaign for waste, extended promotions, handling, and space.	High	Medium
	Total cost per item	Combined over- and under-forecasting costs for the item.	High	Low
	Total cost campaign	Combined over- and under-forecasting costs for the campaign.	High	Low
KPI Cost per Campaign	KPI cost under-forecast per campaign	Absolute cost of the under-forecast KPI.	High	Medium
	KPI cost over-forecast per campaign	Absolute cost of the over-forecast KPIs, reported separately for each.	High	Medium
KPI Cost per Item	KPI cost under-forecast per item	Cost per item of the under-forecast KPI.	Low	High
	KPI cost over-forecast per item	Cost per item of the over-forecast KPIs, reported separately for each.	Low	High
	MERT per article	Maximum number of days for which storing a product in the buffer remains economically viable.	Low	Low

7.3. Results of the Pilot Case Study

This subsection presents the results of the pilot case study at both campaign- and article-level granularity. The reported figures and tables combine the aggregated values of FC0 (Lille) and FC1 (Paris). Each case is discussed per metric group, with each group supported by a graphical illustration or tabular representation. For additional insights into costs at site level, reference can be made to Appendix F.2.1.

A small side note to avoid confusion: both articles show a significant increase in GMI during the most recent promotions. This effect is caused by a supplier change at Picnic around week 12 of 2025.

Use Case 1 – XXX 1

As outlined in the introduction of the use case articles in Section 7.2.1, the first article, XXX 1, was featured in multiple promotions during the period from week 202440 to week 202523. In total, four promotions were conducted for this product. In each case, the forecast proved to be inaccurate. This use case demonstrates what the indicators reveal about these promotions and how, in future promotional selections, it can be determined whether this product is worth including. If deemed worthwhile, it also illustrates how the product's historical data could inform potential adjustments to forecast quantities.

General Metrics

Starting with the product selection perspective, the high number of unavailable units is noteworthy, as it indicates strong potential demand that was not fully met in previous promotions. This suggests that there is clear market interest in the product, but that this opportunity has not yet been fully captured. In addition, the high sales volumes are positive, while the relatively low GMI has a negative influence on overall profitability. The total campaign margin can therefore be considered moderate, resulting from a combination of low GMI but high sales quantities. Taking all factors of Table 7.3 into account, this leads to a positive recommendation for selecting the product for future promotions.

Turning to the forecast quantity, an examination of the unavailable quantities and overstock levels (zero across all promotions) shows that each promotion was under-forecasted. This indicates that Picnic's forecasting model consistently underestimated demand. Consequently, the results suggest that forecasted quantities should be increased in similar future cases. The Promo GMI provides valuable insight into profitability, supporting more informed decisions on whether it is worthwhile to take a calculated risk by forecasting slightly higher quantities. The GMI, which is the same as lost margin per item, of 0.24 lies at the lower bound of Picnic's product range and is therefore not particularly positive (see Figure 7.2); the increase to 0.36 shows clear improvement but remains within the lower range. Conversely, total sales quantities and total campaign margins provide limited guidance for adjusting forecast quantities. Nevertheless, based on Table 7.3, Picnic's historical forecasting performance indicates a general tendency towards under-forecasting, implying that future forecasts should be raised. Moreover, given that high-volume products tend to carry lower risk, since unsold stock is likely to sell shortly after a promotion, slightly higher forecast quantities appear justifiable.

Table 7.3: General Metrics of Case Article 1

Promo week	Unavailable products (units)	Promotion sales (units)	Promo GMI (€)	Total margin campaign (€)	Overstock (units)
202442	238	121	0.24	29.00	0
202450	2253	480	0.24	113.31	0
202509	409	870	0.24	213.04	0
202518	314	740	0.36	263.26	0

Cumulative Forecasting Impact

From a product selection perspective, the results in Table 7.4 show that all four campaigns were affected exclusively by under-forecasting costs, with no recorded over-forecasting costs. These costs represent lost margins and can therefore be viewed as opportunity costs, indicating that the product has strong potential to generate additional profit if demand is forecasted more accurately. The Under-forecast Cost per item remains stable at €0.24 for the first three campaigns, increasing to €0.36 in the final campaign (likely due to a new supplier). At the campaign level, the lost margin ranged from €56.90 to €533.10, resulting in a total of €800.60 and an average of roughly €200 per campaign. While these insights do not provide information about whether the risk of over- or under-forecasting is greater, given the absence of over-forecasting data, the relatively low unit impact (€0.24–€0.36) suggests that under-forecasting poses limited financial risk. However, since the Picnic forecasting model has consistently underestimated demand across all four campaigns, this structural tendency weighs more heavily in determining the appropriate forecast order quantity. Combined with the considerable unrealised profit potential per campaign, it would therefore be reasonable both to select the product for future promotions and to slightly increase forecast quantities to capture its full margin potential.


Table 7.4: Cumulative Impacts of KPIs of Case Article 1

Promo week	Under-forecast Cost per item (€)	Over-forecast Cost per item (€)	Under-forecast Cost campaign (€)	Over-forecast Cost campaign (€)	Total Cost per item (€)	Total Cost of campaign (€)
202442	0.24	0.00	56.90	0.00	0.24	56.90
202450	0.24	0.00	533.10	0.00	0.24	533.10
202509	0.24	0.00	98.90	0.00	0.24	98.90
202518	0.36	0.00	111.7	0.00	0.36	111.70

KPI Costs per Campaign and Costs per Article

From a product selection perspective, the results in Figure 7.3 confirm that the only cost driver for this article is the under-forecast lost margin. As there is only one under-forecast indicator in this research, the values are similar to the cumulative results. Since no over-forecasting costs were recorded, the figures are effectively identical to the total under-forecast costs previously shown in Table 7.4. This reinforces that the product holds strong profit potential that was not fully realised in past campaigns, making it an attractive candidate for future promotional selection.

From a forecast quantity perspective, as the values mirror those of the cumulative results, the conclusions are consistent with those gained from Table 7.4. Given the limited financial risk associated with under-forecasting (€0.24–€0.36 per item) and the clear evidence of lost margin opportunities, it would be advisable to slightly increase the forecasted quantities in future promotions to capture the product's full profit potential.

**Figure 7.3:** Impact of Use Case 1 on KPIs, at both campaign and item level

Use Case 2 – XXX 2

As introduced in Section 7.2.1, the second article, XXX 2, was selected as it represents a regular assortment product that nonetheless exhibits considerable variation in forecasting accuracy across promotions. Two promotional campaigns were conducted for this article within the period from week 202450 to week 202518, each showing distinct forecasting outcomes—one under-forecast and one over-forecast. This case is therefore particularly useful for demonstrating how the dashboard captures both sides of forecasting errors and how these insights can enhance understanding of their underlying causes. Furthermore, the analysis illustrates how such historical insights can support more informed decision-making in both promotional product selection and the review of forecast quantities.

General Metrics

From the product selection perspective, the results in Table 7.5 show that the number of unavailable units was low in the first promotion (29) and absent in the second. This indicates that the product's demand was largely met, and there was no significant missed opportunity in terms of availability. The sales volumes, however, were high in both promotions (around 900 units), suggesting a stable and reliable sales potential. The Promo GMI values of 0.54 and 0.92 are relatively strong compared with the average across Picnic's assortment, particularly the latter, which indicates excellent profitability. Consequently, the total campaign margins of €468.32 and €795.78 confirm that this product delivers solid financial results. Taken together, these outcomes suggest that the product is a strong candidate to select for future promotions.

Considering the forecast quantity, the data show a shift between the two campaigns. During the first promotion,

the small number of unavailable units and the absence of overstock indicate that the forecast was accurate or slightly under-forecasted. In contrast, the second promotion shows an overstock of 323 units, suggesting that the forecast was too high in that instance. As only two promotions are available for comparison, the data are insufficient to draw firm conclusions about structural forecasting bias. Nevertheless, the overstock in the second campaign highlights a possible tendency towards overestimation when sales remain stable. Therefore, based on these metrics, no clear adjustment, upward or downward, can be recommended; rather, maintaining similar forecast quantities while monitoring future outcomes appears to be the most appropriate approach.

Table 7.5: General Metrics of Case Article 2

Promo week	Unavailable products (units)	Promotion sales (units)	Promo GMI (€)	Total margin campaign (€)	Overstock (units)
202450	29	908	0.54	468.32	0
202518	0	896	0.92	795.78	323

Cumulative Forecasting Impact

From a product selection perspective, the results in Table 7.6 show that this product experienced both under-forecasting and over-forecasting across the two observed campaigns. In week 202450, the forecast underestimated demand, resulting in an Under-forecast Cost per item of €0.54 and a total under-forecast cost of €15.80. In contrast, the later promotion (week 202518) shows an overestimation, with an Over-forecast Cost per item of €0.35 and a total over-forecast cost of €110.27. The relatively low magnitude of both cost types indicates that neither under- nor over-forecasting had a severe financial impact. Nonetheless, the product's consistent sales volumes and moderate per-item costs suggest that it remains a suitable candidate for future promotions, offering a balanced risk–return profile.

From a forecast quantity perspective, the results reveal that the financial impact of over-forecasting per item (€0.35) is lower than that of under-forecasting (€0.54). This implies that the financial risk associated with slightly higher forecast quantities is relatively limited. Therefore, adopting a modestly higher forecast could be a prudent strategy to reduce the risk of lost sales and missed margin opportunities. Additionally, examining the contextual factors behind both campaigns—such as promotion type, discount level, or timing—may provide deeper insight into the causes of the forecasting errors observed and further support more accurate future forecasts.

Table 7.6: Cumulative Impacts of KPIs of Case Article 2

Promo week	Under-forecast Cost per item (€)	Over-forecast Cost per item (€)	Under-forecast Cost campaign (€)	Over-forecast Cost campaign (€)	Total Cost per item (€)	Total Cost of campaign (€)
202450	0.54	0.00	15.80	0.00	0.54	15.80
202518	0.00	0.35	0.00	110.27	0.35	110.27

KPI Costs per Campaign and Costs per Article

The results in Figure 7.4 show that this article experienced both under-forecasting and over-forecasting across the two campaigns. In week 202450, the only cost driver was the under-forecast lost margin, whereas in week 202518, over-forecasting costs emerged, primarily driven by handling and space-related KPIs. These findings do not alter the overall recommendation to include the product in future promotions, but they provide valuable insight into which specific factors contribute to over-forecasting costs. As shown, the space cost is more than three times higher than the handling cost, indicating that prolonged storage can have a significant financial impact when products remain in the buffer for extended periods.

This observation suggests that further investigation could assess whether extending the promotion period might have reduced the buffer days and consequently lowered the overstock costs. While the insights at indicator level do not directly affect the decision on forecast quantity, just as with promotion selection, they are nonetheless valuable for understanding where operational inefficiencies arise and for exploring strategies to mitigate them in future campaigns.




Figure 7.4: Impact of Use Case 2 on KPIs, at both campaign and item level

7.4. End-user Evaluation on the Dashboard

This section outlines how the requirements defined in Chapter 6 were met and summarises the general feedback provided by end-users during the dashboard review. At the start of this report, in Section 1.6.1, the feedback loop used throughout the dashboard’s development is visualised. The subsections in this section present the feedback received from test users, structured by the requirements defined in Chapter 6, complemented by broader remarks on the code implementation and dashboard functionality.

Before turning to these subsections, a brief overview of the reviewers is provided. Table 7.7 summarises the key stakeholder groups and their anticipated level of engagement with the artefact. The international relevance of this research is further reinforced by the involvement of Picnic’s Dutch office, which has expressed an intention to adopt the dashboard as a supportive tool for forecasting. Two employees from the Dutch office, a member of the Promo Team and a buyer from the trading team, contributed by providing feedback on the dashboard. Their input is briefly presented in Appendix F, as it lies outside the primary scope of this study and was therefore not integrated into the feedback loops.

Before presenting the evaluation results, it is worth noting that the case study was conducted after the end-user evaluation. This section nevertheless presents the evaluation to illustrate the remarks and suggestions made by the reviewers, many of which have already been incorporated into the final design.

Table 7.7: Overview of the end-users of the artefact (dashboard).

Stakeholder (Group)	Quantity	Background	Excess level
Promotion Team, NL & FR	3	Promo Analyst NL & FR, and Trading Associate	Code & Dashboard
Inbound Team	1	Inbound Analyst	Code & Dashboard
Category Managers, NL & FR	4	Category Managers / Buyers	Dashboard

Requirements

When evaluating the dashboard, it is important to distinguish between functional and non-functional requirements. Functional requirements, such as the calculation of specific KPIs, are fact-based and can be objectively verified through inspection of the SQL code and dashboard outputs. As such, these requirements cannot be meaningfully reviewed by end-users, since their validity is determined by whether the implementation matches the formal definitions rather than by subjective experience. By contrast, non-functional requirements address the quality of the artefact; how it operates, is perceived, and is used in practice. These can be assessed directly by end-users, as they reflect usability, accessibility, clarity, and overall user experience.

Requirements Suitable for End-User Evaluation

With regard to the non-functional requirements listed in Table 6.1, several are suitable for end-user evaluation. The consistent layout and visual hierarchy can be readily reviewed, since users directly interact with the visual structure and can assess whether it supports intuitive navigation and efficient interpretation. Similarly, appropriate visualisation types for KPIs can be judged by end-users, who can determine whether the chosen graphs or charts effectively convey insights and support their decision-making. User-friendly navigation and

accessibility are also well-suited for end-user review, as these aspects depend on how easily users can engage with the dashboard without extensive training or technical expertise. Finally, clean and well-documented code and the maximum three-minute execution time were specifically reviewed by the Promo Lead, who assessed both the execution time and whether the codebase was sufficiently structured and transparent to allow for validation and future adjustments.

Table 7.8 provides an overview of the feedback received on the non-functional requirements, indicating which elements were directly implemented in the dashboard and which, owing to the time constraints of this study, were noted for potential future improvements.

Table 7.8: List of actionable feedback received on requirements, ready for the feedback-loop.

Requirement	Comment	Reviewer	Implemented / Noted
Consistent lay-out	o Add a possibility to filter by article name	o Trading associate	o Implemented
	o Order the promo week filter	o Inbound analyst	o Noted
	o Add period level (period of analysis)	o Category manager	o Noted
Clean code	o Add a comment to every CTE to extra clarify	o Promo lead	o Implemented
	o Add extra comments to critical calculations for better understanding	o Promo analyst	o Implemented
Appropriate visualisation	o Add unconstrained sales	o Trading associate	o Noted
	o Explanation on KPI MERT	o Inbound analyst	o Implemented
	o Having a visual graph on week level (e.g. total margin and lost sales)	o Inbound analyst	o Noted
	o Add total cost of over-forecast per campaign	o Category manager	o Implemented
	o Add an attribute with the promo mechanisms per promotion	o Category manager	o Noted
	o Add impact calculator tab	o Promo analyst	o Noted
	o Add a stochastic indicator capturing uncertainty in forecasting impacts	o Inbound analyst	o Noted
3-Minute execution time	No comment	o Analysts	N/A
User-friendly	o Unclear what cost per item over-/ under-forecasted means	o Inbound analyst	o Implemented
	o Add brand to the table	o Inbound analyst	o Implemented
	o Add FC filter	o Inbound analyst	o Implemented
	o A general guide on the Dashboard would be very useful	o Inbound analyst	o Implemented
	o Ability to sort the column from high to low or vice versa	o Category manager	o Implemented

Requirements Less Accessible to End-Users

Other requirements are less accessible to end-users. For example, modular dashboard architecture is largely invisible to them, as it concerns maintainability and developer practices rather than user interaction. Similarly, cross-device and browser compatibility can only be partially assessed. End-users may notice technical issues on specific devices or browsers, but a full evaluation of this requirement requires systematic technical testing beyond day-to-day use.

In summary, the evaluation by end-users provides valuable insights into the usability-related non-functional requirements, particularly those concerning layout, visualisation, and accessibility, while other requirements remain primarily the responsibility of developers and technical evaluators.

General Feedback

All reviewers also provided their overall assessment of the dashboard. Several recurring points emerged, which are summarised below:

- The dashboard is considered to be of clear added value in relation to their work.
- The dashboard is user-friendly and transparent, although certain improvements could still be implemented.
- The design of both the code and the dashboard is regarded as well-structured, ensuring flexibility for future extensions and modifications.
- Reviewers emphasised the importance of clear explanations and documentation, both within the code and in the dashboard interface, as this strongly supports correct use and long-term maintainability.
- The dashboard is regarded as having considerable potential. While the known back-end limitations of this study leave scope for refinement, they also open opportunities for further development, such as the integration of reputational damage or stochastic analysis. In addition, front-end enhancements were suggested, such as the inclusion of more insightful graphs and a dedicated tab within the dashboard featuring an internal model capable of calculating impacts based on selected indicators.

In summary, the dashboard is already regarded as a useful tool for multiple employees, while at the same time offering significant opportunities for future enhancement and expansion.

7.5. Main Findings Chapter 7


This chapter shows that the dashboard delivers clear, decision-oriented insights at item and campaign level and is considered useful by end-users. The two use case analyses demonstrate that the dashboard effectively provides the intended insights to support better decision-making in both product selection and the review of forecasted order quantities. Its structure allows users to interpret key cost drivers and performance indicators efficiently, translating data into actionable understanding of forecasting outcomes.

While the dashboard already offers substantial decision support, additional attributes, such as promotion duration and promotion mechanism, could enable faster and deeper interpretation of results, further increasing its analytical value. Moreover, expanding the underlying dataset would enhance the robustness of insights, as many articles currently have only a single promotional case. The examples with multiple use cases used in this research (four and two, respectively) illustrate the stronger learning potential that arises when more historical data are available.

End-user evaluation finds non-functional requirements largely met (layout, visualisations, usability, and execution time), with several improvements implemented and further enhancements noted (e.g., additional graphs and an impact-calculator tab). This is a logical outcome, as dashboards are inherently iterative tools that can be continuously improved with new features and functionalities.

VI

Communication



Dashboard Integration

This chapter addresses the integration and communication of the artefact developed and evaluated in the preceding phases. It responds to the fifth research question:

5. *How can the dashboard be effectively integrated into Picnic's current promotional forecasting workflow and communicated across the organisation?*

The purpose of this phase is to ensure that the artefact is not only conceptually understood but also practically adopted within Picnic's day-to-day promotional forecasting processes. The first section (Section 8.1) examines how the dashboard can be embedded within the process flow of Picnic's promotional forecasting, alongside several complementary adjustments aimed at enhancing the overall workflow quality. These adaptations build directly on the recommendations identified earlier during the mapping of the process flow in Phase I of this study. The recommendations can be found in Appendix D. The chapter concludes with Section 8.2, which focuses on the organisational communication of the revised process flow and the practical use of the newly developed artefact.

8.1. Embedding the New Artefact in the Current Process Flow

The DSRP method proposed by Johannesson and Perjons (2014), Peffers et al. (2006), and Pries-Heje et al. (2017) has slightly different perspectives on the final phase of the DSRP. In this study, the final phase, as outlined in Section 1.6.2, extends beyond communicating the artefact to the wider community of researchers and practitioners. It also emphasises the embedding of the artefact into the current process flow of promotional forecasting at Picnic France, ensuring that the tool is not only conceptually validated but also practically integrated into day-to-day decision-making. This section will explain what changes are made to the promotional forecasting process and how the dashboard is embedded in it.

Elaboration on the Process Changes

The promotional forecasting process was systematically mapped in order to identify its limitations, which in turn led to a prioritised list of recommendations. The artefact developed in this study serves as a supportive tool for demand forecasting and primarily enhances the early stages of the process, namely the selection of promotional products for forecasting and the generation of demand forecasts.

Drawing on the insights gained from Sub-question 1 (see Chapter 3) and recommendations identified during this process (provided in Appendix D, the process flow was adapted accordingly (see Figure 8.1). The implemented changes directly address the first, second, and sixth ranked bottlenecks/recommendations in Table D.1. It should be noted that the implementation of the dashboard and the associated recommendations has evolved iteratively through the feedback loop presented in Figure 1.3. Valuable input received during the communication of preliminary versions informed successive refinements, ultimately shaping the current approach.

The modifications made to the promotional forecasting process are illustrated in Figure 8.1. In addition to indicating where the dashboard is integrated, the figure also highlights other process adjustments introduced to further enhance overall efficiency and decision quality. A brief explanation of all process changes is provided below:

1. **Implement Early Involvement of the Fulfilment / Inbound Team:**

In the revised process, a fulfilment or inbound analyst is involved prior to the final selection of promotional products. Together with the promo analyst, this individual reviews a shortlist of critical articles (created by the CMs). By combining dashboard insights on the impact of forecasting errors with experiential knowledge, the analysts jointly determine which products from the critical shortlist should or should not be selected.

2. **Improving Forecast Validation Through Cross-Team Review:**

In the new process, a fulfilment or inbound analyst reviews the forecast quantities generated by the promotion analyst. This review goes beyond numerical accuracy to incorporate logistical considerations in the fulfilment centres. As such, the process not only strengthens the robustness of the forecast validation but also introduces an operational feasibility perspective.

3. **Improving the Efficiency and Quality of Product Selection and Forecast Quantity Reviewing:**

The development of the dashboard itself constitutes this recommendation. In practice the tool, as described in the first adaptation, is used by both promo and fulfilment/inbound analysts. However, the process begins earlier: CMs first employ the dashboard to refine the longlist provided by ITM into a confirmed list of products to be included in the upcoming promotion and a shortlist of articles requiring further review. This shortlist is then passed on to the analysts, who use the dashboard once more to make the final product selection decisions. Subsequently, the dashboard also supports the forecast review stage by providing clear insights into the potential impact of forecasting errors. These insights enable more risk-aware decision-making, helping teams balance promotional goals with practical and financial considerations.

The Overview of the New Process Flow

The modifications to the process flow are highlighted in bold in Figure 8.1. If the new step is directly a step linked to one of the three recommendations from Appendix D.1 tackled this number is provided next to the step. Since the implemented recommendations affect only the first stages of the workflow, the figure presents a snapshot of the beginning of the process. The full updated process flow is provided in Figure C.4.




Figure 8.1: Overview of the revised part of the process flow.

1. Early Fulfilment / Inbound analyst involvement, 2. Forecast review by other teams, 3. Use of the new Dashboard

8.2. Communication of the Dashboard within the Organisation

This section outlines the measures taken to ensure that the newly developed dashboard is effectively transferred within Picnic. The aim is to guarantee that all relevant stakeholders understand its purpose, functionality, and potential applications. In addition, the section briefly describes how the study will be presented externally and made publicly accessible.

8.2.1. Internal Communication

Several complementary measures have been implemented to facilitate the internal transfer of the artefact within Picnic. The more widely it is understood across the organisation, the greater its potential impact and the lower the likelihood of incorrect use.

First, the research and the resulting dashboard were presented in a dedicated session at the Paris office. During this session, the objectives of the study were explained, and the artefact was demonstrated. Emphasis was placed on the value of the tool not only for promotional forecasting and product selection, changes to the promotional forecasting process, but also for other teams. For example, the inbound and fulfilment teams can use the dashboard to gain insights into storage costs and handling costs at fulfilment centres.

Second, the dashboard itself contains a dedicated explanatory tab. This tab provides users with guidance on correct use, outlines the main objectives and functions of the tool, and illustrates potential applications. Its purpose is to encourage users to reflect on the broader utility of the dashboard and ensure that it is applied consistently and appropriately.

Third, following interest from the Dutch office, a separate one-to-one handover was organised with an analyst from the Dutch promotion team. This individual has assumed responsibility for exploring the implementation of the dashboard in the Netherlands, thereby extending the artefact's potential relevance beyond the French context.

Finally, the revised version of the promotional forecasting process flow has been uploaded to Picnic's Confluence platform. As this cloud-based platform functions as the central repository for all organisational processes, it provides an accessible and well-recognised location for communicating the updated workflow. To support this, a clear written description of the revised process was also added, ensuring transparency and accessibility for all employees.

8.2.2. External Communication

The external communication of this study has taken place through an official presentation at Delft University of Technology, where the research was presented in full and subsequently discussed during a question-and-answer session with the audience. In addition, the thesis has been archived in the TU Delft Repository, ensuring permanent accessibility for academic and professional audiences alike.

With this, the chapter is completed and the final sub-question of this phase has been addressed, allowing the thesis to proceed to Chapter 9, which presents the conclusions and recommendations and reflects on the research as a whole.

8.3. Main Findings Chapter 8

This chapter demonstrates how the developed dashboard was successfully embedded within Picnic's promotional forecasting process and communicated across the organisation. The integration improved decision quality, efficiency, and transparency by introducing earlier involvement of fulfilment analysts, cross-team forecast validation, and structured use of the dashboard during product selection and forecast review. Internally, clear communication ensured user understanding and adoption through presentations, guidance tabs, and documentation on Confluence. Externally, the research was shared through academic dissemination at TU Delft, confirming both its practical and scholarly relevance.

Conclusions, Recommendations and Reflections

This chapter presents the main conclusions of the study, outlines recommendations for future research and practice, and concludes with an academic reflection. First, Section 9.1 summarises the key findings and provides answers to the central research question and sub-questions. Building on these insights, Section 9.2 formulates recommendations for the further development of the artefact created in this research, as well as for the broader advancement of decision-support tools within the field of forecasting. Finally, the chapter concludes with an academic reflection in Section 9.3.

9.1. Conclusion

Before discussing the findings related to each research question, it is useful to first reflect on the title of this thesis:

**Adding Value to Promotional Forecasting through a Decision-Support
Dashboard Quantifying Error Impacts**

A Case Study at Picnic

The structure of the thesis closely follows the logic implied by this title. It begins with establishing a clear understanding of the current promotional forecasting process at Picnic in Chapter 3, followed by identifying the specific impacts of forecasting errors in Chapter 4. These impacts are subsequently quantified in Chapter 5, providing the analytical foundation for the development of the artefact, a decision-support dashboard, presented in Chapter 6. The value of this dashboard is confirmed by Chapter 7, which shows that the tool effectively enhances decision-making in the promotional forecasting context. Finally, Chapter 8 ensures that this added value is anchored within Picnic's existing forecasting workflow by integrating the dashboard into the process established in Chapter 3.

Next, the main research question, the question formulated in Section 1.5 is:

How can an artefact, designed to provide insights into the impacts of over- and under-stocking, function as a decision-supportive tool for (a) selecting suitable promotional products and (b) reviewing and adjusting forecasted order quantities within Picnic France?

Based on the findings of this study, the research question can be answered as follows:

"A forecasting impact dashboard, grounded in historical data, can effectively function as a decision-support tool within Picnic's promotional forecasting process by quantifying the multidimensional impacts of forecasting errors and providing actionable insights to support both promotional product selection and the review of forecasted order quantities."

The remainder of this section presents the main conclusions that underpin this answer. For clarity and structure, these are organised according to the research (sub)-questions defined in Section 1.5.

Research Question 1:

How is the current promotional forecasting process structured at Picnic, in terms of workflows, team responsibilities, and interdepartmental coordination?

To address the first research question, a comprehensive overview of Picnic's promotional forecasting process was developed through a combination of literature research, internal document analysis, observational study, and expert interviews with internal stakeholders.

- The current promotional forecasting process at Picnic is structured around the Promo Team, which holds end-to-end responsibility for generating and validating forecasts, while coordinating with CMs, the Store Campaign Team, and the Fulfilment/Inbound Team. The process relies heavily on informal communication and strong ownership by the Promo Team, but lacks systematic cross-team review, early operational input from the Fulfilment/Inbound Team, and structured supplier alignment. In addition, product selection and forecasting are conducted largely without the support of dedicated decision-support tools, leaving the process dependent on individual expertise. In summary, the process is effective in execution but remains vulnerable due to its reliance on centralised ownership and informal coordination, highlighting the need for greater transparency, formalisation, and scalability.

Research Question 2:

Which key performance indicators should be used to assess the impacts of over- and under-stocking caused by promotional forecasting errors at Picnic and what insights does Picnic aim to derive from each indicator?

The second research question was examined through an extensive literature review and an internal document study, with findings validated and refined through feedback from subject-matter experts.

- The key performance indicators affected by promotional forecasting errors at Picnic are Lost Margin, Waste Costs, and Promo Extension Costs on the financial side, and Handling Costs, Space Costs, and Maximum Economical Retention Time (MERT) on the operational side. Together, they quantify both profitability risks and fulfilment impacts of over- and under-stocking. These KPIs provide Picnic with targeted insights into the financial losses and operational inefficiencies associated with forecasting errors. Reputational indicators such as customer satisfaction and brand impact were considered highly relevant but excluded due to time constraints within the scope of this study.

Research Question 3:

How should the indicator calculations be designed to assess the cost implications of over-stocking and under-stocking resulting from forecasting errors in Picnic's promotional context?

The third research question was addressed through additional literature research and iterative prototyping of the indicator calculations, supported by continuous stakeholder feedback.

- The indicator calculations were developed as cost-based formulas that translate forecasting errors into financial and operational impacts. Lost Margin, Waste Costs, and Promo Extension Costs quantify missed sales, unsold products, and unplanned discounts, while Handling Costs, Space Costs, and Maximum Economical Retention Time (MERT) capture additional labour, storage costs, and the economic break-even point for retaining stock. Each calculation was detailed with explicit assumptions, data inputs, and limitations to ensure transparency. Together, they provide a structured, quantifiable basis for assessing the cost implications of over- and under-stocking in Picnic's promotional context.

Research Question 4:

How can a user-friendly artefact be developed to assess the effects of over-stocking and under-stocking caused by forecasting errors?

This research question was addressed in two steps, each linked to a dedicated sub-question. First, the framework was designed in the back-end through SQL-based model development. Second, its applicability was demonstrated by means of simulations and case-based testing, together with an evaluation of end-users feedback. The answers to the two sub-questions are presented below, followed by the overall conclusion to Research Question 4.

Sub-question 4.1: *How should the artefact be designed to calculate and present the impacts of the indicators while serving as a supportive tool for product selection and the review of forecasted order quantities?*

The first sub-question of the fourth research question was addressed through iterative prototyping within an SQL-based modelling environment, combining SQL simulations with repeated design iterations of the dashboard itself.

- The artefact is designed based on explicitly defined requirements that capture both current needs and future functionality. To achieve this, it was developed as a requirements-driven dashboard that (i) calculates the six selected KPIs from historical data within an SQL model, (ii) presents them across multiple granularities (article, campaign, and site) and aggregation levels (absolute values and per-unit over- or under-forecast), and (iii) visualises total over- and under-stock costs side by side to allow direct monetary trade-off analysis. Usability is ensured through a consistent layout, clear labelling, and a modular architecture, enabling future adjustments to indicators, assumptions, or visuals without disrupting the system. Together, these elements make the dashboard a practical and effective tool for both product selection and forecast quantity review.

Sub-question 4.2: *What insights can be derived from the dashboard results, in what ways can these insights support product selection and forecast quantity review, and how do end-users evaluate its usefulness?*

The second sub-question was addressed by conducting simulations and applying a case study demonstration to illustrate the framework's applicability.

- The dashboard provides clear, data-driven insights into the financial and operational impacts of forecasting errors at both item and campaign level. It supports product selection by helping identify which products carry high or low financial risk, and it assists in forecast quantity review by revealing whether historical forecasts tended to over- or underestimate demand. These insights make it easier to justify forecast adjustments and prioritise products for future promotions. End-users evaluated the dashboard as useful, transparent, and user-friendly, noting that it supports more informed and consistent decision-making while offering clear potential for further refinement and expansion.

With the two sub-questions addressed, the consolidated answer to the fourth research question is as follows:

- The artefact was developed as a user-friendly, requirements-driven dashboard that quantifies and visualises the financial and operational effects of over- and under-stocking caused by forecasting errors. Built on an SQL-based model, it calculates six key performance indicators (KPIs) derived from historical data and presents them across multiple granularities and aggregation levels, enabling clear monetary trade-off analysis between over- and under-forecasting. The dashboard translates these complex outcomes into accessible, decision-oriented insights that support both product selection and forecast quantity review. End-user evaluation confirms that the artefact is intuitive, transparent, and valuable in daily operations, providing a structured basis for more consistent and evidence-based forecasting decisions.

Research Question 5:

How can the dashboard be effectively integrated into Picnic's current promotional forecasting workflow and communicated across the organisation?

The fifth research question was addressed by integrating the artefact into practice, presenting it internally, and incorporating stakeholder feedback to ensure its effective adoption within Picnic's forecasting process.

- The dashboard can be effectively integrated into Picnic's promotional forecasting workflow by embedding it in the early product selection and forecast validation stages, where Promo, Category Management, and Fulfilment/Inbound teams jointly use its insights. This integration strengthens cross-team review, improves operational feasibility checks, and supports data-driven product selection. Communication is ensured through dedicated training sessions, explanatory materials within the dashboard, one-to-one handovers across offices, and documentation on Picnic's online platform ('Confluence'), enabling consistent adoption and organisation-wide accessibility.

9.2. Recommendations

Based on the findings in this report, recommendations are provided for further research on the problem at hand. The recommendations focus on both practical implementations within the promotional forecasting field of Picnic, and on further academical research in the field of (promotional) forecasting. The recommendations are prioritised based on an indication of the estimated implementation time.

9.2.1. Short-term Recommendations

The short-term recommendations consist of two categories: immediate improvements to the dashboard itself and refinements to the calculation of specific KPIs. All cases can be implemented with relatively low effort yet provide tangible gains in accuracy, usability, and decision-making support.

Recommendation 1: Implement Actionable Dashboard Improvements

Based on end-user feedback, a set of concrete improvements was identified (see Table 7.8). Addressing these points, ranging from enhanced filtering options to clearer visualisations, would elevate the dashboard to a higher level of usability and impact, ensuring it delivers even greater value across Picnic’s forecasting process.

Recommendation 2: Refine KPI Calculations

Several refinements are proposed to improve the accuracy and relevance of the KPIs:

- **Extended Promotion KPI:** Incorporate the “anti-waste promotion” as a third type of extended campaign. Although less frequent, including it would ensure a more complete representation of forecasting outcomes.
- **Handling Cost KPI:** (1) Handling times for all product types are based on ambient values from FC0. To improve accuracy, handling times should be calculated separately for chilled and frozen products in FC0, and likewise determined individually for all three product types in FC1. (2) The current average-per-CU method should be replaced with a TU-based calculation: first compute the average handling time per TU, then divide this by the CU/TU ratio. This approach better reflects operational reality, as handling is performed per TU rather than per CU.
- **Space Cost KPI:** Extend the calculation to include the impact of buffer stock levels on picking efficiency. At present, space costs only capture rent, security, cleaning, and utilities. Since high buffer levels can slow down picking pace, incorporating this factor would provide a more accurate estimate of the true cost of overstocking in constrained buffer situations.

9.2.2. Mid-term Recommendations**Recommendation 3: Integrate the Dashboard into the Forecasting Model**

Currently, the dashboard functions as a retrospective decision-support tool, providing valuable insights into the financial and operational consequences of past forecasting errors. However, these insights remain disconnected from the actual forecasting model, limiting their immediate applicability in real-time decision-making.

To close this gap, a complementary tool should be developed that directly links the dashboard’s outputs to Picnic’s forecasting model. This integration will provide the forecasted output along with recommended quantities to over- or under-forecast, depending on the potential impact of errors.

Implementation should proceed in three steps:

- **Data integration:** build an API or automated pipeline that feeds dashboard outputs (e.g. waste costs, lost margin estimates) directly into the forecasting model.
- **Model enhancement:** extend the forecasting model to incorporate these indicators as weighted decision inputs, allowing it to balance over- and under-forecasting risks.
- **User interface:** create an interactive layer where analysts can adjust assumptions, run scenarios, and immediately see how changes impact both forecasts and associated costs.

Feasibility is high, given Picnic’s existing SQL- and dashboard-based infrastructure. The integration requires primarily engineering capacity to establish data pipelines and adapt the forecasting model, rather than fundamental methodological innovation. Once implemented, this tool would transform the dashboard from a retrospective analysis instrument into a proactive forecasting aid, bridging the gap between insight generation and operational decision-making.

Recommendation 4: Expand Dashboard Implementation Beyond France

Once the dashboard has been fully stabilised and embedded within Picnic France, it should be implemented in other key markets, notably the Netherlands and Germany.

The dashboard addresses universal challenges in promotional forecasting, and many of the structural and coding foundations developed in this study can be reused. While not a one-to-one copy, the core framework—including KPI calculations, SQL model design, and dashboard structure—provides a strong basis for adaptation. Tailoring will be required to reflect market-specific operational contexts and data structures, but the investment effort is substantially lower than starting from scratch.

A phased roll-out is recommended, beginning with the Netherlands, where stakeholder interest is strongest, followed by Germany. This approach ensures lessons learned from France can be leveraged to accelerate adoption.

Recommendation 5: Redesign the Calculation of Unavailability

Picnic currently estimates the effect of unavailability by multiplying in-app clicks on unavailable products with an outdated conversion-to-purchase ratio and product commonality. This method introduces considerable uncertainty and may misrepresent the true financial and operational impact of stock-outs.

Accurately capturing the costs of unavailability is essential for assessing under-forecasting impacts, particularly as they directly influence key KPIs such as lost margin and reputational risk. The existing approach relies too heavily on assumptions and static ratios that may no longer reflect actual consumer behaviour.

A new methodology should be developed that bases unavailability calculations on observed sales patterns across the entire promotion period. Specifically, the share of sales achieved during availability periods should be extrapolated to estimate the sales that would likely have occurred during unavailability. This approach provides a more data-driven and transparent estimate, grounded in actual promotional performance.

The proposed method should be regarded as a starting point. Future researchers or analysts are encouraged to refine or adapt the methodology as appropriate, ensuring that the calculation remains robust and reflective of Picnic's evolving customer base and forecasting needs.

9.2.3. Long-term Recommendations**Recommendation 6: Investigate the Cannibalisation Effect**

One of the most critical next steps for Picnic is to establish an understanding of the cannibalisation effect in general. Currently, no structured estimate exists within the company, despite the fact that cannibalisation directly influences several topics within Picnic. And for this research in particular two key performance indicators (KPIs). This lack of insight creates a blind spot that undermines both financial accuracy and operational decision-making.

A dedicated study should be launched to investigate cannibalisation in a structured, sustainable manner. Such a study should:

- **Differentiate by context:** analyse cannibalisation effects by article type (e.g. staple versus niche products), promotion timing (seasonal versus off-season), and customer segments.
- **Account for influencing factors:** assess how campaign design, product similarity, and market conditions shape cannibalisation levels.
- **Enable growth over time:** ensure that the methodology is extensible and maintained as a sustainable deliverable, allowing Picnic to update and refine the analysis as influential factors can change over time. This could be realised, for example, by embedding the project into a thesis internship that builds a scalable model as its primary outcome.

From a KPI perspective, cannibalisation affects forecasting evaluation in two ways:

1. **Lost Margin:** reduced sales of substitute products during a promotion, or conversely, reduced cannibalisation when a product is unavailable.
2. **Extended Promotion Costs:** decreased sales of similar products due to cannibalisation by the extended promotion item.

Given these dual effects, Picnic should consider integrating cannibalisation as a distinct KPI within the dashboard framework. This would clarify the contribution of cannibalisation to forecasting impacts and prevent it from being obscured within the broader lost margin or extension cost calculations.

Recommendation 7: Strengthen Strategic Alignment with Key Supplier (ITM)

Invest in closer collaboration with Picnic's main supplier, ITM, through weekly joint working sessions. Current bottlenecks, such as long forecast lead times, limited system integration, and unclear inbound delivery updates, directly undermine forecasting accuracy by reducing visibility on product availability and buffer management. Stronger alignment with ITM would streamline operations and improve the reliability of Picnic's promotional forecasting.

Implementation:

- Establish weekly co-working days with ITM to align forecasts, resolve integration issues, and improve transparency.
- Formalise agreements on delivery updates and lead time flexibility, especially during peak periods.
- Prioritise system integration improvements that allow real-time visibility of inbound flows.

Regular collaboration can be implemented quickly and may yield early wins in information exchange, trust, and forecasting reliability. However, achieving the broader structural improvements intended by this recommendation will require sustained cooperation, with full benefits only materialising over the longer term.

Recommendation 8: Investigate the Real Costs of Reputational Impact

At present, reputational impacts are entirely excluded from Picnic’s forecasting evaluation, leaving a significant blind spot in understanding the broader consequences of forecasting errors. To gain a better initial understanding of this impact, Appendix F.2.1 presents a brief approximation of the potential financial cost per unavailability. This calculation only considers the damage caused by product unavailability, yet even with a rough estimate of €0.10 per unavailable item, the potential impact is already substantial. Therefore, this recommendation is well worth further investigation. To address this, additional KPIs should be developed to capture reputational impacts, supported by a dedicated study to systematically measure their real costs.

A dedicated study should therefore be initiated to systematically measure the real costs of reputational impact, with particular attention to:

- **Customer-facing dimensions:** effects of product unavailability on customer satisfaction, trust, and loyalty, as well as potential positive effects of extended promotions on customer perception.
- **Brand reputation:** how repeated forecasting failures or waste events influence Picnic’s image as a reliable and sustainable retailer.
- **Environmental consequences:** reputational costs associated with avoidable waste, including the negative perception of excess CO₂ emissions linked to discarded products.
- **Positive spillovers:** potential gains in reputation from demonstrating transparency, sustainability, and responsibility in managing forecasting outcomes.

Such a study should go beyond anecdotal evidence and employ a robust, data-driven methodology—for example, combining customer surveys, churn analysis, and external benchmarking with financial proxies for reputational events. Ideally, the deliverable would be sustainable and extensible, providing Picnic with a framework that can be updated over time and reused in future campaigns.

9.3. Reflections

This section reflects on the relevance of the research from both an academic and a practical perspective. It discusses how the study contributes to the scientific understanding of forecasting in online retail and how its outcomes create tangible value within Picnic’s operational context.

9.3.1. Academic Reflection

This research contributes to the academic understanding of forecasting within the context of online retail by introducing an approach that integrates operational cost considerations into decision-making. While existing literature primarily focuses on improving predictive accuracy or refining model architectures, this study advances the discussion by emphasising the consequences of forecasting errors. The developed framework extends beyond accuracy metrics by quantifying the operational and financial impacts of over- and under-stock, thereby linking forecasting performance to its real-world implications.

During the preparatory literature review, a notable research gap was already identified: no studies in online retail explicitly connect forecasting outcomes to their operational cost impacts. This absence served as the main inspiration for this thesis. To support this claim, additional exploratory literature searches were conducted using terms such as “forecasting using operational impact indicators” and “cost-based evaluation of forecast errors.” These searches further confirmed that the integration of operational cost indicators into forecasting evaluation remains largely under-explored in academic research.

By embedding these indicators into a decision-support artefact, this research offers a methodological contribution that bridges theory and application. It demonstrates how operational costing methods—such as Activity-Based Costing and Activity-Based Space Costing—can be systematically translated into forecasting impact indicators. Consequently, this work provides a conceptual foundation for future studies on decision-support systems that incorporate cost-based, stochastic, or risk-aware elements.

Beyond the primary contribution to cost-based forecasting, this research offers two key additions to existing methodological literature. The first of these is the extension of the established Design Science Research Process (DSRP). While the conventional DSRP cycle concludes with the communication of the artefact and research through a paper, this thesis incorporates an additional, critical action: the implementation of the new artefact into the operating system. This final step adds a new dimension to the DSRP approach, highlighting the value of systemic integration beyond purely academic communication.

Another contribution to the literature involves the methodology for process mapping. Existing literature on process mapping often treats techniques in isolation, but this research demonstrates the synergistic value of combining methods with different strengths. Specifically, this study's approach of integrating the visual and temporal clarity of the Swimlane Diagram with the role-based precision of the RASCI Framework offers a more comprehensive and robust analytical framework than either method could achieve independently. Exploratory searches using terms such as "process mapping methods combined" confirm the novelty of this integrated approach within the field.

In essence, this thesis achieves a three-fold contribution: it repositions forecasting as a decision-oriented rather than purely prediction-oriented discipline—adding a crucial layer of operational relevance to an area traditionally dominated by statistical metrics. Furthermore, it offers methodological advancements by demonstrating the synergistic value of combined process mapping techniques and by advancing the DSRP method through the inclusion of a critical implementation in the final phase.

9.3.2. Societal Reflection

From a societal and practical perspective, this research demonstrates the tangible value of translating data-driven insights into daily operational decision-making within a fast-growing online retailer. The developed dashboard has proven to be a valuable addition to Picnic's forecasting process, allowing decision-makers to visualise and quantify the operational and financial implications of forecasting errors. By doing so, it enables more transparent, evidence-based, and sustainable inventory decisions—reducing waste, improving stock management, and supporting efficient resource allocation across fulfilment centres.

The societal relevance lies in the broader applicability of this approach. As e-grocery and online retail sectors continue to expand, the balance between operational efficiency and sustainability becomes increasingly critical. A tool that helps organisations anticipate and mitigate the costs of misaligned forecasts contributes directly to reducing unnecessary production, transport, and waste—issues that extend beyond firm-level performance to environmental and societal responsibility.

Moreover, by demonstrating how academic concepts such as impact indicators and cost modelling can be operationalised in a business context, this research highlights the societal importance of bridging the gap between analytical research and practical implementation. The artefact not only enhances Picnic's internal forecasting decisions but also serves as an example of how data science and operational management can jointly advance sustainable and efficient retail logistics.


Bibliography

- Abolghasemi, M., Gerlach, R., Tarr, G., & Beh, E. (2019). Demand forecasting in supply chain: The impact of demand volatility in the presence of promotion. <http://arxiv.org/abs/1909.13084>
- Agbemadon, K. B., Couturier, R., & Laiymani, D. (2023). Overstock Prediction Using Machine Learning in Retail Industry. *2023 3rd International Conference on Computer, Control and Robotics, ICCCR 2023*, 439–444. <https://doi.org/10.1109/ICCCR56747.2023.10194060>
- Al-Ashaab, A., Golob, M. M., Busachi, A., & Al-Ashaab Matic Golob Paula Noriega Filippo Torriani Pedro Alvarez Ana Beltran Alessandro Busachi Lambert Ex-Ignotis Carlotta Rigatti Sanjiv Sharma, A. (2013, September). *Capturing the Industrial Requirements of Set-Based Design for CONGA Framework* (tech. rep.). <https://www.researchgate.net/publication/251880527>
- Allio, M. K. (2006). Strategic thinking: the ten big ideas - Strategic dashboards: designing and deploying them to improve implementation. *Strategy & Leadership*, 34(4), 4–13. <https://doi.org/10.1108/10878570610676837>
- Almeida, A., & Cunha, J. (2017). The implementation of an Activity-Based Costing (ABC) system in a manufacturing company. *Procedia Manufacturing*, 13, 932–939. <https://doi.org/10.1016/j.promfg.2017.09.162>
- Bach, B., Freeman, E., Abdul-Rahman, A., Turkay, C., Khan, S., Fan, Y., & Chen, M. (2022). Dashboard Design Patterns. <http://arxiv.org/abs/2205.00757>
- Bandyopadhyay, S. (2009). A Dynamic Model of Cross-Category Competition: Theory, Tests and Applications. *Journal of Retailing*, 85(4), 468–479. <https://doi.org/10.1016/j.jretai.2009.05.001>
- Calvo-Porrà, C., & Lévy-Mangin, J.-P. (2017). Specialty food retailing: examining the role of products' perceived quality. *British Food Journal*, 119(7), 1511–1524. <https://doi.org/10.1108/BFJ-11-2016-0567>
- Carson, R. T. (2000). Contingent valuation: A user's guide. *Environmental Science and Technology*, 34(8), 1413–1418. <https://doi.org/10.1021/es990728j>
- Carter, R. (2024). Dutch app supermarket boss eyes tech boom in online delivery. *Tech Xplore*.
- Chen, Y., Ray, S., & Song, Y. (2006). Optimal pricing and inventory control policy in periodic-review systems with fixed ordering cost and lost sales. *Naval Research Logistics*, 53(2), 117–136. <https://doi.org/10.1002/nav.20127>
- Cheng, L., Tsou, C. S., Lee, M. C., Huang, L. H., Song, D., & Teng, W. S. (2013). Tradeoff analysis for optimal multiobjective inventory model. *Journal of Applied Mathematics*, 2013. <https://doi.org/10.1155/2013/619898>
- Cooper, L. G., Baron, P., Levy, W., Swisher, M., & Gogos, P. (1999). *PromoCast: A New Forecasting Method for Promotion Planning* (tech. rep. No. 3).
- Dai, Y., & Huang, J. (2021). A Sales Forecast Method for Products with No Historical Data. *2021 IEEE 6th International Conference on Cloud Computing and Big Data Analytics (ICCCBDA)*, 229–233. <https://doi.org/10.1109/ICCCBDA51879.2021.9442603>
- Damelio, R. (2011, May). *The Basics of Process Mapping* (Second Edition, Vol. 183). Productivity Press. <https://doi.org/10.4324/9781439891278>
- Daniels, J., Werner, P. W., & Bahill, A. T. (2001). Quantitative methods for tradeoff analyses. *Systems Engineering*, 4(3), 190–212. <https://doi.org/10.1002/sys.1016>
- Dayeeta Das. (2024, August). Picnic Posts Full-Year Profit Despite Seeing Overall Losses Rise.
- DeHoratius, N., & Raman, A. (2008). Inventory record inaccuracy: An empirical analysis. *Management Science*, 54(4), 627–641. <https://doi.org/10.1287/mnsc.1070.0789>
- Dresch, A., Daniel, ., Lacerda, P., Antônio, J., & Antunes, V. (2015). *Design Science Research A Method for Science and Technology Advancement* (tech. rep.). Springer.
- Emmett, S. (2005). *Excellence in Warehouse Management How to Minimise Costs and Maximise Value* (tech. rep.).
- Evers, J. M. ; Tavasszy, L. ; Van Duin, R. ; Schott, D. ; & Gorte, F. (2018). *Demand forecast models for online supermarkets* (tech. rep.). TU Delft. Delft.

- Faber, N. (2015, September). *Structuring warehouse management: Exploring the fit between warehouse characteristics and warehouse planning and control structure, and its effect on warehouse performance* (tech. rep.). Erasmus University - Journal of Physical Distribution & Logistics Management. Rotterdam. www.b-en-t.nl
- Fahimnia, B., Tan, T., & Tahirov, N. (2024). Service-level anchoring in demand forecasting: The moderating impact of retail promotions and product perishability. *International Journal of Forecasting*. <https://doi.org/10.1016/j.ijforecast.2024.07.007>
- Farris, P. W., Bendle, N. T., Pfeifer, P. E., & Reibstein, D. J. (2014, February). *Marketing metrics: the definitive guide to measuring marketing performance*. (Second Edition). Pearson Education.
- Francisca Silva. (2023, October). Understanding the Financial Impact Inventory Carrying Costs in Warehouses.
- Fujiwara, D., & Campbell, R. (2011). *Valuation techniques for social cost-benefit analysis : stated preference, revealed preference and subjective well-being approaches : a discussion of the current issues*. HM Treasury.
- Gastezzi C., Rodríguez M., & Castillo A. (2024). Theoretical foundations on Customer Experience (customer experience, NPS, CSAT, CES, Service Balcony, Journey Map). *Journal of Business and Entrepreneurial Studies*.
- Gruen, T. W., & Corsten, D. (n.d.). *A Comprehensive Guide To Retail Out-of-Stock Reduction In the Fast-Moving Consumer Goods Industry* (tech. rep.).
- Gupta, S. (1988). Impact of Sales Promotions on when, what, and how Much to Buy. *Journal of Marketing Research*, 25(4), 342–355. <https://doi.org/10.1177/002224378802500402>
- Gupta, S., & Lehmann, D. R. (2003). Customers as assets. *Journal of Interactive Marketing*, 17(1), 9–24. <https://doi.org/10.1002/dir.10045>
- Gustavsson, J., Cederberg, C., & Sonesson, U. (2011). *Global Food Losses and Food Waste-Extent, Causes and Prevention* (tech. rep.). FOOD and AGRICULTURE ORGANIZATION OF THE UNITED NATIONS. Rome. <https://www.researchgate.net/publication/285683189>
- Haefner, B., Kraemer, A., Stauss, T., & Lanza, G. (2014). Quality value stream mapping. *Procedia CIRP*, 17, 254–259. <https://doi.org/10.1016/j.procir.2014.01.093>
- Hanemann, W. (1989, March). *Willingness to pay and willingness to accept: How much can they differ?* (Tech. rep.). American Economic Review.
- Heemskerk, T. (2024, March). *Cold-Start Promotional Demand Forecasting with Contrastive Explanations in an E-grocery Setting* (tech. rep.). Erasmus University. Rotterdam.
- Hill, R., Seifbarghy, M., & Smith, D. (2007). A two-echelon inventory model with lost sales. *European Journal of Operational Research*, 181(2), 753–766. <https://doi.org/10.1016/j.ejor.2006.08.017>
- Imam Raza. (2025, March). Retail Demand Forecasting. <https://data.folio3.com/blog/retail-demand-forecasting/>
- Jedidi, K., Mela, C. F., & Gupta, S. (1999). Managing advertising and promotion for long-run profitability. *Marketing Science*, 18(1), 1–22. <https://doi.org/10.1287/mksc.18.1.1>
- Johannesson, P., & Perjons, E. (2014, September). *Design Science* (tech. rep.). Springer. Stockholm.
- Kaplan, R. S., & Anderson, S. R. (2004). Time-Driven Activity-Based Costing. *Harvard Business Review*.
- Kerckänen, A., Korpela, J., & Huiskonen, J. (2009). Demand forecasting errors in industrial context: Measurement and impacts. *International Journal of Production Economics*, 118(1), 43–48. <https://doi.org/10.1016/j.ijpe.2008.08.008>
- Kitchenham, B. (2007). *Guidelines for performing Systematic Literature Reviews in Software Engineering* (tech. rep.). <https://www.researchgate.net/publication/302924724>
- Kök, A. G., & Fisher, M. L. (2007). Demand estimation and assortment optimization under substitution: Methodology and application. *Operations Research*, 55(6), 1001–1021. <https://doi.org/10.1287/opre.1070.0409>
- Kök, A. G., Serdar, A., İm, S., & Moore, D. (2021, March). *Variety and Inventory Trade-off in Retail: Unpacking the Role of Inventory Concentration* *Necati Tereya' go' glu* (tech. rep.). <https://doi.org/http://dx.doi.org/10.2139/ssrn.3791772>
- Koottatep, P., & Eng, B. (2006). *Promotional Forecasting in the Grocery Retail Business* (tech. rep.).
- Kruizinga. (2025, September). 2-Heks rolcontainer klemhekken + 2 spanbanden.
- Lee, H. L., Padmanabhan, V., & Whang, S. (1997). *The Bullwhip Effect in Supply Chains* (tech. rep. No. 3). <http://mitsmr.com/1phEOiM>

- Louviere, J. J., Hensher, D. A., & Swait, J. D. (2000). *Stated choice methods : analysis and applications*. Cambridge University Press.
- Ma, S., Fildes, R., & Huang, T. (2016). Demand forecasting with high dimensional data: The case of SKU retail sales forecasting with intra- and inter-category promotional information. *European Journal of Operational Research*, 249(1), 245–257. <https://doi.org/10.1016/j.ejor.2015.08.029>
- McColl, R., Macgilchrist, R., & Rafiq, S. (2020). Estimating cannibalizing effects of sales promotions: The impact of price cuts and store type. *Journal of Retailing and Consumer Services*, 53. <https://doi.org/10.1016/j.jretconser.2019.101982>
- Mitchell, R. C., & Carson, R. T. (2013, October). *Using Surveys to Value Public Goods*. RFF Press. <https://doi.org/10.4324/9781315060569>
- Monti, N. E. (2009, April). *A Warehouse Activity Based Costing Model*. TU Eindhoven. www.tue.nl/taverne
- Mupanni, V., & Adil, G. (2008). Efficient formation of storage classes for warehouse storage location assignment: A simulated annealing approach . *Omega*, 36(4), 609–618. <https://doi.org/10.1016/j.omega.2007.01.006>
- Nicholls, J., Lawlor, E., Neizert, E., & Goodspeed, T. (2012, January). *A guide to Social Return on Investment* (tech. rep.).
- Nordman, J. (n.d.). *Analyzing warehouse cost structures and pricing models* (tech. rep.).
- Orts, D. (2004). *Dashboard Development and Deployment A Methodology for Success* (tech. rep.). Noetix. www.noetix.com
- Peffer, K., Tuunanen, T., Gengler, C. E., Rossi, M., Hui, W., Virtanen, V., & Bragge, J. (2006). *The Design Science Research Process: A Model for Producing and Presenting Information Systems Research* (tech. rep.). Claremont Graduate University.
- Pries-Heje, J. R., Baskerville ; Venable, J. ; Pries-Heje, J. R., & Baskerville, J. (2017). *Choosing a Design Science Research Methodology* (tech. rep.). APA.
- Rasiklal Rakholia, K., Chandraprabha, Ramesh, R., Srinivasa Rao, K., Punitha, S., & Guru Vimal Kumar, M. (2025). Optimizing Inventory Management Through Demand Forecasting: A Data-Driven Approach for Enhanced Supply Chain Efficiency. <https://doi.org/10.2139/ssrn.5076111>
- Reul, M. (2025). Picnic continues French expansion. *Retail Detail*.
- Richards, G. (2018). *Warehouse management : a complete guide to improving efficiency and minimizing costs in the modern warehouse* (Third Edition). Kogan Page Ltd.
- Roeloffzen, M. W. H. (2007, October). *Trade-off between service and inventory costs* (tech. rep.). University of Twente. Twente.
- Six Sigma. (2025, April). Responsibility Assignment With RASCI Matrix in Lean Six Sigma Projects.
- Sofia, S. (n.d.). *Activity Based Costing in Warehousing: A Practical Example* (tech. rep.).
- Solis, D. (2024). Under the hood of Picnic’s demand forecasting model: A Deep Dive into the Temporal Fusion Transformer.
- Speh, T. W. (2009). *Understanding Warehouse Costs and Risks* (tech. rep.). Warehouse Forum - The Ackerman Company.
- Stavrova, O., & Evans, A. M. (2019). Examining the trade-off between confidence and optimism in future forecasts. *Journal of Behavioral Decision Making*, 32(1), 3–14. <https://doi.org/10.1002/bdm.2087>
- Stefan van Rompaey. (2022, December). Picnic advances towards Paris. <https://www.retaildetail.eu/news/food/picnic-advances-towards-paris/>
- Stone, L. (2024, September). The Swim Lane Process: A streamlined alternative to RACI charts.
- Stratton, R., & Warburton, R. D. (2006). Managing the trade-off implications of global supply. *International Journal of Production Economics*, 103(2), 667–679. <https://doi.org/10.1016/j.ijpe.2006.01.001>
- Then, V., Schober, C., Rauscher, O., & Kehl, K. (2017). *Social Return on Investment Analysis Measuring the Impact of Social Investment* (tech. rep.). <https://doi.org/https://doi.org/10.1007/978-3-319-71401-1>
- Van Donselaar, K. H., Peters, J., De Jong, A., & Broekmeulen, R. A. (2016). Analysis and forecasting of demand during promotions for perishable items. *International Journal of Production Economics*, 172, 65–75. <https://doi.org/10.1016/j.ijpe.2015.10.022>
- Vigneshwar Ravichandran. (2024, August). Amsterdam’s Picnic achieves its first profit in the Netherlands: Know more. <https://siliconcanals.com/amsterdams-picnic-achieves-its-first-profit/#:~:>

- text=Amsterdam%2Dbased%20Picnic%2C%20an%20online,efficiency%2C%20reports%20FD.nl.
- Von Rosing, M., White, S. A., Cummins, F., & De Man, H. (2014, December). Business process model and notation-BPMN. In *The complete business process handbook: Body of knowledge from process modeling to bpm* (pp. 429–453, Vol. 1). Elsevier Inc. <https://doi.org/10.1016/B978-0-12-799959-3.00021-5>
- Wee, B. V., & Banister, D. (2016). How to Write a Literature Review Paper? *Transport Reviews*, 36(2), 278–288. <https://doi.org/10.1080/01441647.2015.1065456>
- Wolfert, A. R. (2023, August). *Open design systems*. IOS Press. <https://doi.org/10.3233/RIDS10>
- Zhang, L.-f., Zhu, Y.-g., & Hu, H.-y. (2010). Models of Storage Cost Analysis for Warehouse Leasing and Expanding. *2010 International Conference on Optoelectronics and Image Processing*, 16–18. <https://doi.org/10.1109/ICOIP.2010.174>
- Zipkin, P. (2006). *On the Structure of Lost-Sales Inventory Models* (tech. rep.).



SQL Code Github Repository

The **SQL** code developed for this research is available in the **GitHub Gist** at <https://gist.github.com/FransKraanen/b087103a2157ef95025942497bc0b972>

B

Literature Research Background

This Appendix section has content related to, Chapter 2, the literature review. Starting with an overview of all the papers used for the preparatory literature review.

Table B.1: Overview of all preparatory literature

General Topic	Specified Topic	Paper(s)	Document Type
<i>Demand Forecasting</i>	Promotion	(Abolghasemi et al., 2019)	Journal Article
		(Cooper et al., 1999)	Journal Article
		(Gupta, 1988)	Journal Article
		(Jedidi et al., 1999)	Journal Article
		(Koottatep & Eng, 2006)	Thesis
		(McColl et al., 2020)	Journal Article
	Over-/under-forecasting consequences	(Van Donselaar et al., 2016)	Journal Article
		(Imam Raza, 2025)	Journal Article
		(Fahimnia et al., 2024)	Journal Article
		(Hill et al., 2007)	Journal Article
		(Chen et al., 2006)	Journal Article
		(Lee et al., 1997)	Journal Article
	Inventory optimisation	(Rasiklal Rakholia et al., 2025)	Journal Article
	Error impact	(Kerkk�nen et al., 2009)	Journal Article
	Cold-start Promotion	(Heemskerk, 2024)	Thesis
	No historical data	(Dai & Huang, 2021)	Conference Paper
	Online supermarket	(Evers et al., 2018)	Paper / Thesis
	High-dimensional data	(Ma et al., 2016)	Journal Article
	Multi-objective trade-offs	(Cheng et al., 2013)	Journal Article
	Temporal Fusion Transformer	(Solis, 2024)	Technical Blog
	Confidence and optimism trade-off	(Stavrova & Evans, 2019)	Journal Article
	Cross-category effects	(Bandyopadhyay, 2009)	Journal Article
<i>Machine Learning in Retail</i>	Overstock prediction	(Agbemadon et al., 2023)	Conference Paper
	Inventory optimisation	(Rasiklal Rakholia et al., 2025)	Journal Article
	Cold-start Promotion	(Heemskerk, 2024)	Thesis
	No historical data	(Dai & Huang, 2021)	Conference Paper
	Inventory Modelling	(Zipkin, 2006)	Technical report
<i>Process Mapping</i>	Process improvements	(Damelio, 2011)	Book
	Process misalignments	(DeHoratius & Raman, 2008)	Journal Article
		(K�k & Fisher, 2007)	Journal Article
	RASCI method	(Six Sigma, 2025)	Webpage
	Swimlane	(Stone, 2024)	Webpage
<i>Dashboard</i>	Strategy	(Allio, 2006)	Journal Article
	Design	(Bach et al., 2022)	Technical Report
	Development	(Orts, 2004)	Corporate Report
<i>Valuation Methods</i>	Contingent valuation method	(Carson, 2000)	Journal Article
	Customer lifetime value	(Farris et al., 2014)	Book
		(Gupta & Lehmann, 2003)	Journal Article
		(Fujiwara & Campbell, 2011)	Government Report
	Social cost-benefit analysis	(Hanemann, 1989)	Journal Article
	WTP vs. WTA	(Kaplan & Anderson, 2004)	Journal Article
	Warehouse costs valuation	(Kaplan & Anderson, 2004)	Journal Article
	Stated choice methods	(Louviere et al., 2000)	Book
	Survey-based valuation	(Mitchell & Carson, 2013)	Book
	Social Return on Investment (SROI)	(Nicholls et al., 2012)	Technical Report
		(Then et al., 2017)	Book
<i>Warehouse Mangement</i>	Operational Efficiency	(Emmett, 2005)	Book
	Holding costs	(Francisca Silva, 2023)	Corporate report
		(Mupanni & Adil, 2008)	Journal Article
		(Zhang et al., 2010)	Conference paper
		(Speh, 2009)	Journal Article
	Handling costs	(Speh, 2009)	Journal Article
<i>Retail Trends</i>	Tech in online delivery	(Carter, 2024)	News Article
	Food waste	(Gustavsson et al., 2011)	Technical Report
	Cross-category effects	(Bandyopadhyay, 2009)	Journal Article
	Perceived product quality	(Calvo-Porr�l & L�vy-Mangin, 2017)	Journal article
<i>Trade-off Analysis</i>	Quantitative evaluation	(Daniels et al., 2001)	Journal Article
	Inventory strategy	(K�k et al., 2021)	Journal Article
	Service vs. inventory cost	(Roeloffzen, 2007)	Thesis
	Global supply chain	(Stratton & Warburton, 2006)	Journal Article

Below is the ABC flow chart provided, used as inspiration for the ABC formulas used in this research.




Figure B.1: ABC implementation flow chart by Almeida and Cunha (2017).

Process Mapping Deliverables

In this Appendix the main deliverables of Chapter 3 are shown. Starting with two overviews of the promotional forecasting processes and presented using the RASCI method (Six Sigma, 2025). The first overview describes the process from ten weeks until seven weeks prior to the actual promotion, while the second visualises the process from three weeks before until the conclusion of the promotion. No changes occur in the process between these periods.

C.1. Import-Effect Matrix

For the ranking of the recommendations a supportive Impact-Effort Matrix was created.




Figure C.1: Impact-Effort Matrix of the recommendations found with the process analysis.

C.2. RASCI Process Flow Deliverables

The RASCI overviews outline the process flow of promotional forecasting. Five main stakeholders are incorporated in the RASCI framework: Fulfilment / Inbound Team, Store Campaign Team, Promo Team, Category Managers, and the Supplier (ITM). The process is divided into two phases. The first phase, covering weeks -10 to -7, focuses on the selection of products for promotion and the initial forecasting and communication of demand commitments to the supplier (ITM). The second phase, spanning weeks -3 to +1, involves the preparation of all promotional activities for the store, the re-forecasting of regular products, and, where necessary, the adjustment of commitments made during week -10.




Figure C.2: RASCI process flow -10wk to -7wk.




Figure C.3: RASCI process flow -3wk to +1wk.

C.3. New Promotional Forecasting Process Flow

Here is the full overview of the adapted process flow, with more involvement from the fulfilment / inbound team, broader forecasting review and the inclusion and use of the new artefact.





Figure C.4: Complete overview of the new Swimlane overview of the promotional forecasting process.



Practical Deliverable: Process Bottlenecks and Recommendations

This appendix chapter elaborates on the potential improvements and recommendations identified during the mapping of the promotional forecasting process in Chapter 3. As these findings do not directly contribute to answering the main research question, they are excluded from the main body of the thesis. However, the insights gained during the mapping process are considered highly valuable for Picnic France and may serve as a basis for future operational improvements.

D.1. Identified Bottlenecks and Improvement Areas

This section consolidates all relevant bottlenecks or inefficient process steps, uncovered during the analysis of the promotional forecasting process. Findings are structured into three thematic categories to increase clarity and traceability: Process Inefficiencies, Operational Risks and Communication. Each insight is backed by evidence from stakeholder, Swimlane and RASCI diagrams insights, or observed interactions during process mapping. The final outcome of this mapping exercise is the Swimlane process diagram shown in Figure 3.2 below, which provides a high-level overview of the entire workflow. More detailed diagrams, including RASCI matrices clarifying roles and responsibilities, are available in Appendix C.

D.1.1. Process Inefficiencies

This category captures structural shortcomings in the promotional forecasting process that reduce transparency, consistency, and accountability across teams.

1. Lack of Independent Forecast Review

While mapping the forecasting process and discussing it with the Promo Lead, a lack of independent review was identified. In the current process setup, the Promo Team is responsible for both the creation and validation of promotional forecasts. As observed through the Swimlane diagram and stakeholder interviews, no formal review loop with other departments is in place. The RASCI diagram shows no designated “Challenger” role for this step. This introduces a risk of internal confirmation bias, as forecasts are evaluated by the same team that generated them. Although this setup speeds up the process, it reduces the level of critical review. Without an independent review, incorrect assumptions or misjudged demand levels may persist, affecting order accuracy and ultimately customer satisfaction.

2. Unstructured Process for Highlights of Promotion Page in the App

During stakeholder discussions and analysis of the process flow, an unstructured approach to selecting promotional highlights was observed. These highlights, promotional products with the highest visibility in the Picnic app, are selected around week -2. This selection process currently takes place through informal communication channel within Picnic, as confirmed by the Swimlane diagram. However, the RASCI model reveals no formally assigned roles or responsibilities for this step. The absence of a documented process or shared quality criteria means decisions are often made ad hoc, based on subjective preference rather than strategic fit or data. During interviews, some teams were unaware of the criteria for inclusion, leading to occasional overlap or misalignment in promoted products.

3. Fragmented Ownership of In-App Promotional Content

While applying the detailed RASCI framework to the process, a fragmentation in ownership of in-app promotional content was identified. The Promo Team is responsible for configuring specific elements (such as the promotion page), while the Store Campaign Team manages the homepage and broader campaign components. Although this division aligns with each team's area of expertise, it creates a split ownership structure without a single team being accountable for the full customer-facing result. The RASCI diagram confirms the lack of overarching responsibility for visual consistency, tone of voice, and update coordination. This leads to inconsistencies in app strategy and reliance on informal communication for late-stage changes.

D.1.2. Operational Risks

These risks highlight vulnerabilities in the process design that may undermine resilience and scalability when disruptions occur.

1. No Plan B for Supplier Cancellations

As explained in Section 3.3 ITM can cancel orders two weeks before the start of the promotion. During the process mapping exploration, it became evident that when suppliers cancel promotions at the last minute, there is no predefined protocol to follow. These cancellations require the Promo Team to react under high time pressure, often scrambling for substitute products or abandoning the promotion altogether. This reactive mode of working introduces avoidable stress, reduces planning confidence, and increases the likelihood of incoherent or lower-quality promotional offerings. Furthermore, last-minute replacements often lack proper internal alignment or funding, as suppliers do not subsidise the fallback promotions.

2. Strong Ownership Promo Team

The Promo Team's strong sense of ownership in the whole promotional forecasting process was repeatedly highlighted as a strength during process related conversations with stakeholders. Their deep product knowledge and ability to quickly pivot during disruptions are valuable. However, this high degree of autonomy also introduces risks, particularly around transparency, documentation, and handover. The Swimlane diagram shows them as active across nearly all phases, but without formal external input at key stages. This setup, while effective in execution, makes the process highly person-dependent and may limit scalability.

3. Manual Product Selection and Forecast Reviewing

Through conversations with Category Managers, it became clear that the selection of 'In/Out' products, temporary items not part of the standard assortment, is currently performed manually by the Category Managers. This selection depends heavily on individual expertise and intuition. In addition, the manual review of forecasts, also based primarily on individual judgement, lacks clear guidance. These steps together were identified as a bottlenecks, particularly in terms of transparency, knowledge transfer and efficiency. While this method allows for flexibility and expert judgement, it introduces subjectivity, limits reproducibility, and complicates scaling. New team members, in particular, lack structured guidance or replicable criteria, which makes the process difficult to transfer or audit.

D.1.3. Communication

This category outlines misalignments in cross-team and supplier communication that hinder coordination, planning accuracy, and strategic coherence.

1. Lack of Input from Inbound/Fulfilment Team at Begin of the Process

While mapping the process and conducting informal conversations with members of the Fulfilment/Inbound teams, it became clear that their involvement is limited to checking forecasts that have already been made. Yet, this team possesses essential knowledge about warehouse capacity, congestion, and picking constraints; factors that could meaningfully influence product selection and forecast accuracy if considered earlier. The lack of a structured mechanism to include their input during the initial stages leads to occasional mismatches between promotional plans and what is operationally feasible.

2. Limited Strategic Alignment with Main Supplier (ITM)

Informal conversations with the Fulfilment/Inbound team and reflections on external collaboration revealed that Picnic's main promotional supplier, ITM, does not align its planning cycles with Picnic's internal timelines. This misalignment leads to recurring operational challenges during the execution of promotions. Strategic coherence between ITM's roadmap and Picnic's retail strategy is lacking, despite the high volume of promotional goods sourced through this partnership. Currently, no structured coordination mechanism exists to bridge this gap.

3. Store Campaign Team Unaware of 10-Week Lead Times

It became apparent that the Store Campaign team was unaware that certain promotional products require commitment ten weeks in advance due to supply chain constraints. This is a critical aspect of the process that all involved teams should understand. The knowledge gap, uncovered while validating the process overview with the Store Campaign lead, highlights insufficient internal communication.

D.2. Recommendations

Based on the analysis of the promotional forecasting process, a set of targeted recommendations has been developed. These are given in Table D.1. First, each problem was ranked to provide an overview of its relative severity.

This problem ranking then served as input for prioritising the recommendations. The problem ranking was established through a discussion with the main stakeholder, the promotion team lead, who is involved in all process steps and is therefore well positioned to assess the severity of each issue. His ranking was subsequently verified with other stakeholders. The final ranking of problems is presented in Appendix D.1. Afterwards, all identified problems and their associated solutions were discussed with the promotion team lead. Using an Impact–Effort Matrix, the recommendations were positioned according to their expected benefit and required effort. This ranking differs from the problem ranking, as a low-effort recommendation, for example, may generate a quick positive impact.

Each recommendation is ranked and linked to the problems identified in Appendix D.1. In this section, the rationale for the ranked place of each recommendation is outlined, together with suggested implementation steps, expected impact, and estimated effort. To support prioritisation, all recommendations are positioned within an Impact–Effort Matrix, as shown in Figure C.1. The prioritisation logic follows three principles: (1) capture near-term value and build stakeholder confidence with low-effort/high-impact “quick wins”; (2) schedule higher-complexity, high-impact “big bets” once prerequisites (data, governance, capacity) are in place; and (3) interleave incremental improvements to sustain momentum, minimise operational disruption, and validate assumptions before scaling. This sequencing reduces delivery risk, accelerates learning, and helps early benefits de-risk and fund subsequent investments. As two recommendations address the same issue, the recommendation ranking ranges from 1 to 10, while the problem ranking ranges from 1 to 9.

Table D.1: Categorized process related problems and recommendations ranked.

Categorie	Problem	Rank	Recommendation	Period*	Rank
Process	Lack of independent forecast review	5	Implement cross-team review	ST	2
Process	Unstructured page highlight coordination	9	Structure highlight coordination	ST/MT	10
Process	Fragmented ownership of in-app promotional content	8	Full responsibility to store campaign team	ST/MT	7
Operational Risk	No plan B for supplier cancellations	6	Develop a fallback plan	MT	8
Operational Risk	Strong ownership promo team	4	Create a more balanced model	MT/LT	9
Operational Risk	Manual products selection and forecast reviewing	2	1 - Provide tooling to guide decisions	1 - MT	4
			2 - Full automated selection process	2 - LT	6
Communication	Lack of early input from Fulfilment/Inbound team	3	Strengthen early involvement of the Fulfilment/Inbound team	ST	1
Communication	Strategic alignment with main supplier	1	Strengthen strategic alignment with key supplier	MT/LT	5
Communication	Store campaign unaware of 10-week lead times	7	Actively focus on process knowledge of involved teams	ST	3

*ST = Short-term, MT = Mid-term, LT = Long-term

#1 - Strengthen Early Involvement of the Fulfilment Team

Recommendation: Establish a recurring alignment mechanism through which the Fulfilment/Inbound team provides early input on the operational situation at the start of the forecasting process.

Promotional product selection and forecasting are currently carried out with limited input from the Fulfilment team, despite their critical knowledge of warehouse capacity, inbound flow pressure, and picking limitations. This disconnect has led to occasional misalignments between promotional ambitions and operational feasibility. To address this, it is recommended to implement a routine touchpoint, such as a brief weekly sync or a shared dashboard, where key fulfilment metrics (e.g. storage congestion, turnaround times, tote capacity) are shared with the Promo Team before key planning decisions are made. Implementing a regular touchpoint or shared dashboard with relevant fulfilment metrics would improve feasibility assessments and promote more realistic forecasting. A clear easy win, with almost the lowest effort and a very high impact this recommendation is ranked as number 1.

#2 - Improve Forecast Validation Through Cross-Team Review

Recommendation: Introduce a formal forecast review by another team.

At present, the Promo Team is solely responsible for both the creation and internal validation of promotional forecasts. This lack of external challenge introduces a risk of confirmation bias, where flawed assumptions may go unchecked. Introducing a structured review step, carried out by a team not involved in the initial forecast creation, would strengthen the reliability of projections. This cross-team validation could be embedded

into existing process flow, but can just be covered by another team, the category managers seems the best option. While the effort required is relatively low, especially when leveraging current alignment meetings, the impact is significant. A more robust and transparent forecasting process reduces errors and increases overall confidence in the promotional plan. This makes the measure both feasible and high-priority. This recommendation is ranked highly for similar reasons as the previous one: a lightweight adjustment with significant benefit.

#3 - Actively Focus on Process Knowledge of Involved Teams

Recommendation: Actively improve process knowledge across involved teams, particularly regarding important 'need-to-knows' of the forecasting process.

During the process validation phase, it became apparent that the Store Campaign team was unaware that certain promotional products require commitment ten weeks in advance due to supply chain constraints. This type of knowledge gap illustrates broader misalignment in shared understanding of the forecasting process. When critical constraints are not known by all parties involved, it can lead to inefficient planning, miscommunication, and quality risks in campaign execution. To address this, Picnic should strengthen cross-team process awareness by embedding key operational 'need-to-knows' into standard communication routines. Suggested steps include the development of concise onboarding materials and the introduction of an obligatory review and understanding of the newly developed process flow for all involved stakeholders. This recommendation requires relatively little effort but offers high impact. Better-informed teams will collaborate more effectively, anticipate constraints earlier, and enhance the overall robustness of the promotional forecasting process. This is the last clear easy win of the recommendations and therefore ranked as third.

#4 and #6 - Improve the Efficiency and Quality of Product Selection and Forecast Reviewing

Recommendation 1: Keep manual selection but provide optional tooling to guide decisions (#4).

Recommendation 2: Full automated selection process (#6).

The current product selection process relies heavily on individual expertise within the Promo Team, and the same applies to the review of the forecast output. While this approach yields solid outcomes due to strong internal intuition and experience, it presents challenges in terms of scalability, transparency, and onboarding. In the mid-term, the introduction of optional tooling, such as dashboards displaying historical promotional performance or lift prediction models, would support consistency and assist newer team members. These tools would not replace expert judgement but enhance it by making data more accessible and actionable. In the long-term, a shift towards a fully automated selection mechanism could be considered. Automating this step would reduce manual effort, improve reproducibility, and allow for quicker iteration. Although implementation would require a significant initial investment, the long-term benefits in terms of reliability and efficiency justify its inclusion as a strategic improvement trajectory. As this bottleneck relates to two potential recommendations, a choice must be made between a long-term solution with very high impact and a mid-term option with slightly lower impact. Nonetheless, it is evident that improving the efficiency of product selection is a high priority as it directly will effect efficiency.

#5 - Strengthen Strategic Alignment with Key Supplier

Recommendation: Initiate structured discussions with ITM to identify and implement feasible changes that improve logistical alignment, while remaining realistic and workable for both parties.

While many operational challenges can be addressed internally, recurring inefficiencies in Picnic's promotional execution stem from a lack of coordination with ITM, its primary promotional supplier. ITM's planning timelines are currently not synchronised with Picnic's internal cycles, and this misalignment has been passively accepted. As Picnic's promotional strategy becomes more sophisticated, closer strategic alignment with ITM could enhance campaign timing, product availability, and overall execution quality. To achieve this, Picnic should initiate structured dialogue to build mutual understanding of planning constraints and explore concrete adjustments that would significantly benefit Picnic while remaining manageable for ITM. Initial steps could include joint planning pilots, earlier information exchange, and quarterly alignment meetings. Though this is a longer-term effort with external dependencies, the potential impact on supply reliability and campaign effectiveness makes it a high-leverage improvement.

This strategic 'big bet' requires time, patience, and diplomacy given the importance of the supplier relationship. While implementation is not straightforward, even modest gains can significantly boost planning reliability and flexibility. The more effort invested, the greater the potential impact. This recommendation focuses on long-

term improvement and high strategic value, requiring persistence in pursuit of major gains. All in all a fifth rank is fair for this recommendation.

#7 - Fragmented ownership of in-app promotional content

Recommendation: Full responsibility to store campaign team.

Currently, responsibility for in-app promotional content is divided across multiple teams. The Promo Team selects which products will be highlighted in the app, while the Store Campaign team manages the homepage layout and broader campaign elements. This fragmentation leads to inconsistencies in overall application strategy and creates coordination inefficiencies, as informal communication is often required to align last-minute changes.

To address this, full operational responsibility for in-app promotional content should be assigned to the Store Campaign team. This would include visual strategy, configuration, and the alignment of promotional highlights and messaging across all app sections. The Promo Team would still provide input on product selection and campaign priorities, but would no longer be involved in execution. This shift would enhance coherence, clarify ownership, and reduce the risk of conflicting strategies. The required effort is moderate and mainly organisational, while the impact on campaign quality and internal alignment is expected to be high. This recommendation is ranked seventh as it has moderate impact and effort; improves ownership and coordination, though the current situation is not a major bottleneck.

#8 - Develop a Structured Fallback Plan for Supplier Cancellations

Recommendation: Create a generalised decision framework to guide responses when suppliers cancel promotions shortly before launch.

Although supplier cancellations are infrequent, they are often high-pressure events that disrupt the promotional planning process. At present, responses are handled ad hoc, with limited time for coordination and little to no supplier funding for substitute campaigns. This results in reactive decision-making, potential inefficiencies, and inconsistent campaign quality. To mitigate this, a structured fallback framework should be developed. Rather than specifying detailed replacements in advance, the framework would define key decision points, such as whether to cancel or substitute the promotion, based on factors like timing, product category, and available alternatives. It should also outline the financial implications, as fallback promotions are typically self-funded. This recommendation is relatively light in terms of effort but significantly increases organisational readiness. By offering clear guidance, it supports faster and more coherent decision-making, reduces stress during disruptions, and protects the overall integrity of promotional campaigns. However, given the limited frequency of such cancellations, the overall priority (8th) of this recommendation is lower compared to improvements that affect daily operations.

#9 - Strong Ownership Promo team

Recommendation: Develop a more balanced model with clearer structure, shared responsibilities, and cross-team support.

The Promo Team plays a central role in the entire promotional process, showing strong ownership and high execution speed. While this is generally a strength, the current setup makes the process heavily person-dependent and limits scalability. Responsibilities are not always documented or shared, making it difficult for others to step in when needed. To address this, a more balanced model could be developed. A process that formalises responsibilities, clarifies handovers, and encourages cross-team involvement during key decision points. The current process already delivers solid performance, and the issue is not particularly pressing. Since implementation would require considerable effort relative to the urgency and current impact, this recommendation is ranked second to last.

#10 - Structure Highlight Coordination

Recommendation: Replace informal feedback, via communication platform Slack, with a lightweight review board or shared file, including a checklist for selection quality.

The current process for selecting promotional highlights operates through informal conversation via communication platform 'Slack'. This limits traceability and creates ambiguity around selection criteria. There is little visibility into why specific products are chosen or which alternatives were considered. This lack of structure

increases the risk of inconsistencies and subjective decisions, especially as multiple teams are involved. A simple shared workspace, such as a collaborative file with a standardised checklist, could serve as a lightweight quality gate. Criteria might include category diversity, balance between impulse and core products, and alignment with strategic goals. Such a tool would not only improve transparency but also enable more structured feedback and informed decision-making. This recommendation is ranked lowest due to its limited impact. Although current coordination via Slack functions adequately, it lacks structure and professionalism. While formalising this process could be a future improvement, the current setup does not pose urgent issues, making this the lowest ranked priority.

Indicator Selection Background

This Appendix is related to all background deliverables of Chapter 4. Starting with the visualisation of the RF-Matrix Development in Appendix E.1, followed by the created longlist of potential KPIs in Table E.2

E.1. Relevance-Feasibility Matrix Development

This appendix section provides an overview of the iterative process that led to the construction of the RF-matrix. Table E.1 presents the modifications made by stakeholders during the validation phase. A +R5, for example, indicates that the stakeholder assigned a higher score to the 'Relevance' of the 5th indicator. These inputs were incorporated to create the final result. In the final matrix, it is evident that stakeholder influence weighed more heavily for indicators related to their own area of expertise.

Table E.1: Adjustments on RF-Matrix per stakeholder

Stakeholder	Adjustments Relevance	Adjustments Feasibility
Promo Lead	-R5, R6	-F1, -F3, -F4, -F6, -F9
Inbound Lead	-R3, + R5, +R6, +R9, -R10	-F4, -5F, -F9
Customer Team Lead	-R6, +R7, -R8, +R9, -R10	-F6, -F9

For some extra textual background read the following:

To ensure the robustness and practical relevance of the selected indicators, a validation step was conducted with three key stakeholders within Picnic, each representing one of the three indicator domains: financial, operational, and reputational. Specifically, the Promo Lead was consulted for financial indicators, the Inbound Lead for operational indicators, and the Customer Team Lead for reputational indicators. These individuals were deliberately selected for their strong decision-making authority within their respective domains and their necessary domain knowledge to assess both the relevance and feasibility of each proposed indicator.

The validation process began with the presentation of the initial RF-matrix, which had been constructed based on prior literature, internal documentation, and the researcher's understanding of data availability and Picnic's strategic context. Each stakeholder was approached individually to review the full set of indicators, regardless of category, while drawing particularly on their specific area of expertise. Rather than collaboratively developing the matrix from scratch, this approach allowed for more focused discussions. Stakeholders were asked to confirm, adjust, or challenge the assessments of both relevance and feasibility, enabling efficient and targeted refinement of the indicator set. After receiving insights from all three stakeholders, a final RF-Matrix was created. This version was presented to them for final validation, with all three confirming it after a few iterations, as shown in Appendix E. In cases of significant disagreement, the stakeholder related to the specific domain had more authority. The outcome of this process was a refined and validated indicator set, which will serve as the foundation for subsequent dashboard development and analysis.




Figure E.1: Overview of Stakeholder Modifications During RF-Matrix Validation

E.2. Indicators Longlist

The longlist of potential indicators includes a comprehensive set of financial, operational, and reputational metrics, derived from internal documentation, discussions, and literature research. Each indicator in the list is accompanied by an explanation, and where indicators were not selected for further investigation, the reasoning behind their exclusion is provided.

Table E.2: Longlist of Indicators


Indicator	Explanation	Reason to Exclude Indicator
<i>Financial</i>		
Lost Margin	Estimates the unrealised profit from under-forecasted promotions where products go out of stock too early. Helps quantify the financial cost of missed sales opportunities.	N/A
Waste Costs	Captures the value of products that must be discarded due to expiry or poor sales.	N/A
Promo Extension Costs	Measures the extra cost of extending promotions to clear excess stock, calculated as the difference in margin between total forecasted regular sales and total discounted promo sales. Indicates the financial impact of over-forecasting.	N/A
Supplier Chargebacks	Accounts for penalties, missed bonuses, or lost funding due to poor promotional execution, such as inaccurate volumes. Reflects how forecasting errors affect supplier agreements.	Within Picnic this is not the case for promotional forecasting and therefore not taken into account for this research
<i>Operational</i>		
Handling Time Costs (€)	Measures the extra time spent handling due to forecasting errors. Increases with inefficient flows caused by over- or under-forecasting.	N/A
Picking Errors	Tracks the proportion of items picked incorrectly, which may rise during irregular or last-minute changes caused by forecast deviations. Affects operational quality and customer satisfaction.	This indicator overlaps significantly with both Handling Time Costs and Customer Satisfaction. Including it would compromise a MECE (Mutually Exclusive, Collectively Exhaustive) approach, which is undesirable for the structure of this research.
Space Costs	Measures the cost of using space in the fulfilment centre over time, typically calculated in euros per cubic metre per hour (€/m³/h). This KPI helps quantify the financial impact of inefficient storage caused by over-forecasting.	N/A
MERT	Maximum Economical Retention Time - The maximum number of days a promotional product can remain in the buffer zone before its cumulative handling and space costs outweigh its financial value.	N/A
Transport Utilisation	This indicator measures how efficiently delivery capacity is used. Over-forecasting can lead to excess stock that doesn't leave the fulfilment centre as expected, resulting in partially filled delivery runs. Lower utilisation increases transport cost per item and reduces logistical efficiency.	Comment from Inbound Lead: "In general this is the case, however for promo forecasting this is not applicable."
Stock Reallocation Events	This indicator counts the number of urgent internal stock movements triggered by inventory imbalances, typically due to under-forecasting. These events often require manual intervention to move products between fulfilment locations or zones, increasing handling time and disrupting standard workflows.	This number can be interesting but the financial impact of promo-forecasting related to handling will be calculated in 'Handling Time Costs' therefore this is irrelevant
<i>Reputational</i>		
CSAT	Customer Satisfaction Score - Measures how satisfied customers are with a specific promotion, typically collected via post-purchase surveys or in-app feedback. It reflects the perceived quality of the promotional experience.	N/A
NPS	Net Promotor Score - Evaluates the extent to which promo-related issues, such as stockouts or incorrect pricing, influence customers' willingness to recommend Picnic to others. A decline in score may indicate reputational risk.	CSAT and NPS are two different ways of measuring customer satisfaction, each offering a slightly different perspective. Since both cannot be used simultaneously, a choice must be made. CSAT is selected because it is more granular, operationally actionable, and better aligned with the goal of this framework: assessing the consequences of forecasting errors at the article level.
Brand Impact	Brand Trust Impact Score - Assesses how promotional execution, especially repeated failures, affects overall trust in Picnic's brand reliability. This can inform long-term brand perception tracking.	N/A
Complaint Rate	Tracks the proportion of promo-related orders that result in complaints due to unavailability, delays, or incorrect fulfilment. It serves as a proxy for service consistency during campaigns.	N/A
Carbon Footprint of Unsold Products	Estimates the emissions linked to unsold promotional inventory, including production, transport, and disposal. It provides insight into the environmental cost of over-forecasting.	N/A

E.3. Shortlist of Excluded Indicators

The following table presents the four reputational indicators that, while relevant, were ultimately excluded in this study due to feasibility constraints.

Table E.3: Non-chosen Indicators background

#	Indicator	Comments	Limitations & Assumptions
Reputational			
7.	Customer Satisfaction	o Measures immediate customer satisfaction with a specific promotion, typically collected via surveys after purchase. Focuses on short-term experience rather than long-term loyalty.	- o Short-term focus: CSAT reflects immediate customer sentiment and may not capture the lasting effects of promo-related issues. - o Survey bias: Results can be skewed by low response rates or disproportionately Me vocal customers (either very satisfied or very dissatisfied). - o Context sensitivity: External factors (e.g. weather, delivery experience) may influence CSAT even if unrelated to the promotion itself. - o Assumes respondents provide honest and representative feedback. - o Assumes CSAT scores can be meaningfully attributed to forecasting accuracy and promotional execution.
8.	Brand Impact	o Estimates how promotional execution (e.g., repeated stockouts or poor availability) affects long-term trust in the brand. It reflects customers' confidence in Picnic's reliability over time.	- o Attribution difficulty: Brand trust is influenced by many factors beyond promotional execution, making it hard to isolate the effect of forecasting errors. - o Measurement complexity: Quantifying trust reliably requires longitudinal data and may rely on proxy indicators, which can reduce precision. - o Time lag: Changes in brand trust may take time to manifest and thus cannot be immediately linked to specific events. - o Assumes that customers perceive promo failures as reflective of overall brand reliability. - o Assumes consistent methodology is used to monitor brand trust over time.
9.	Complaint Rate	o Promo related complaints are mainly unavailability's. A small share is delivery error's (due to logistical congestion)	- o Complaint underreporting: Not all dissatisfied customers submit complaints, potentially underestimating the true issue. - o Overlap with other KPIs: Some complaints (e.g., about unavailability) may also affect CSAT or Brand Trust, risking double counting in impact assessments. - o Assumes the relative difference in complaint rates between promo and regular items reflects the specific impact of forecasting errors. - o Assumes €/complaint rate is a reasonable proxy for the operational and reputational cost per complaint.
10.	Carbon Footprint	- o EF consist of production, packaging, transport - o Carbon Price is a market-based price - o % Real waste, is the percentage of food that literally is thrown away, instead of donated, eaten by employees.	- o Data availability: Exact emission data per article is often unavailable; emission factors are likely based on averages or proxies. - o Variability in carbon pricing: Market-based carbon prices vary over time and geography, affecting comparability and monetary valuation.



Demonstration & Evaluation Phase

F.1. Feedback from Dutch End-Users

This section presents one-to-one feedback gathered from two employees at the Dutch office. While these insights are valuable for guiding the future development of the dashboard, they fall outside the formal scope of the present study.

Feedback from commercial (promotion) business analyst:

what I like:

- gives immediate insight into the additional costs of a campaign, which I can subtract from the earned gross margin to get the "actual" incremental impact of a campaign (e.g. if I earned 20,000 euro gross margin with a campaign but I also has 5,000 euro waste, the actual incremental impact is 15,000 euro -> super helpful for a more helicopter view on financial reporting of campaigns)
- supports me in making business decisions for future campaigns as well, e.g. how much do I want to intentionally over-forecast a campaign to ensure enough availability, knowing that the cost of over-stock is X cts per item (now we just over-forecast by a lot when there's no waste risk)

what I miss:

- would be super valuable if there would be an "impact calculator" tab, where the user can input all info for an upcoming campaign (e.g. article X in week Y with gross margin Z, expecting to sell 25,000 CU but intentionally over-forecasting 35,000 CU to ensure availability) and that the calculator then automatically calculates what would be the actual impact of the campaign (earned gross margin minus costs of over-forecasting). That way I can play around with the 35,000 CU forecasted volume (should it be more / less?) and see what the impact is on additional costs

Feedback from category manager (dairy):

The dashboard is intuitive and easy to use. However, the addition of further graphical representations could add value by providing deeper insights.

In the Netherlands, Category Managers (CMs) do not select promotions; this responsibility lies with the suppliers. CMs only negotiate the terms of promotions, while the actual forecasting of quantities is handled by commercial and trading analysts. For this reason, the current version of the dashboard is not directly relevant for CMs. Nevertheless, it could still provide valuable insights to support negotiations with suppliers. For example, if two promotional deals are offered by a dairy supplier, the dashboard may reveal that one of them generates relatively limited value for Picnic. This insight could guide CMs to focus future negotiations on securing a single, more attractive deal that benefits both customers and Picnic.

F.2. Evaluation of the Articles

F.2.1. Quick Approximation of Reputational Damage

Reputational damage falls outside the formal scope of this research, and precise valuation could not be achieved within the available timeframe. Nevertheless, a quick approximation has been developed to provide indicative figures of the potential cost per unavailability event. Since the number of unavailability events is already tracked, these costs can be estimated by combining the tracked data with the proposed approximation. This addition to the dashboard may serve as an extra incentive to incorporate reputational impact in the future into the dashboard with correct numbers.

Estimating Reputational Costs of Product Unavailability

This estimation is based on the study of Gruen and Corsten (n.d.).

The study indicates that 20–40% of customers confronted with an out-of-stock switch immediately to a competitor; 10–15% abandon the purchase altogether; and the remainder opt for a substitute from the same retailer. Persistent exposure to unavailable products increases the likelihood of customer churn and gradually erodes loyalty.

Assuming Picnic customers place approximately 25 orders per year with an average basket size of €50, this equates to €1,250 in annual turnover and roughly €50 in gross margin (4%) per customer.

Empirical evidence suggests that 1–3% of customer lifetime value (CLV) may be lost per negative shopping experience, particularly if such incidents occur repeatedly. A single out-of-stock event rarely leads directly to churn, but it does increase the probability. As this study considers individual incidents rather than the cumulative effect of repeated experiences, a conservative estimate assumes 10% of the 1–3% loss (on average 0.2%). Applied to the estimated €50 annual margin, this corresponds to approximately €0.10 per unavailability event.

Lower bound (€0.10): When most customers substitute within Picnic's assortment, but minor reputational damage still occurs.

Upper bound (€1.00): When a meaningful share of customers shift their purchase to a competitor or develop negative sentiment towards Picnic.

Additional Factors

For In/Out products (temporarily listed items), the reputational impact is likely lower, as customers have no strong expectation of availability. Conversely, for core products such as milk or bread, the impact is considerably higher, as unavailability undermines Picnic's central value proposition. Nevertheless, for the purposes of this analysis, a unified value of **€0.10** per unavailability event is applied throughout.