Technische Universität München

Lehrstuhl für Produktion und Supply Chain Management

Methodologies for option bundle design in the automotive industry: An operations perspective

Radu-Constantin Popa

Vollständiger Abdruck der von der Fakultät für Wirtschaftswissenschaften der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Wirtschaftswissenschaften (Dr. rer. pol.)

genehmigten Dissertation.

Vorsitzender:		Prof. Dr. Rainer Kolisch
Prüfer der Dissertation: :	1.	Prof. Dr. Martin Grunow
	2.	Prof. Dr. Matthias Holweg

Die Dissertation wurde am 14.03.2018 bei der Technischen Universität München eingereicht und durch die Fakultät für Wirtschaftswissenschaften am 15.11.2018 angenommen.

Acknowledgements

It has been a long, sometimes cumbersome road, but its end has come. At this point, I would like to thank everyone who helped me get through this journey.

I am very grateful for the professional development I experienced thanks to my supervisor, Professor Martin Grunow. His expertise, his fostering of independent work and his goal-oriented mindset have helped me pave a path through the intricacies of academia. I benefited tremendously from our research discussions, having learned how to be more focused and get my point across clearer while also broadening my theoretical knowledge. I also appreciated the support of Dr. Thomas Stäblein, who offered me the data for my research and intriguing research ideas. The talks we had gave me detailed insights in the operations of automotive manufacturers. I would also like to thank Professor Renzo Akkerman for his clear feedback during the time he was at TUM. His openenss and his unwavering availability to talk about all sorts of problems have offered me a vision of how team leadership should look like. Furthermore, I am grateful for the fruitful cooperation with Professor Matthias Holweg and for his demonstration of efficient paper-writing. I thank him and Professor Rainer Kolisch for being on the assessment committee of my thesis.

The journey would have been a lot harder without the wonderful cooperation with all my colleagues, past and present, from the group of Production and Supply Chain Management. I not only appreciated the insightful research-related discussions we had, but also the fun board game evenings and movie visits I had with some of them. I am especially grateful to Sina for the wonderful cooperation we had, starting from my time as a diploma student under her supervision. She has shown me the true meaning of attention to details and team commitment. Furthermore, I appreciated the lively professional and personal talks with Daniel.

My journey would most likely not have been as enjoyable without my family and friends. I am grateful to my parents, Alexandru and Mihaela, for facilitating my studies in Germany and lending me an open ear whenever I was confronted with uncertainties. The talks with my brother, Rares, have been invaluable and have helped me persevere even when times were tough. I appreciated a lot the time I spent with my friends Tudor, Iulia, Sabina, Alex, Teo, Coco and Marcel and their full support during my PhD, even if sometimes I was not available for meet-ups for a while.

Finally, I am most thankful to my girlfriend Nadine. Her unwaivering support, her unbreakable belief in me and her patience with me even at my low points have helped me push through the last phase of my PhD.

Radu Constantin Popa

Abstract

The automotive industry is confronted with difficult challenges ahead. A fundamental shift in customer attitude towards car ownership, more strict environmental regulations for vehicles and the disruptive impact of emerging car-related technologies jeopardize the capability of automotive manufacturers to generate satisfactory profit rates and to fund crucial investments. Product variety mitigation strategies, such as option bundling, can help automotive manufacturers reduce their costs. A reduction of product variety would also improve the accuracy of the component demand forecasts communicated to automotive suppliers.

Even though option bundling has been claimed to have a positive effect on operations due to a reduction of product variety, a thorough analysis of the impact of option bundling on product variety is lacking. Furthermore, none of the existing option bundling methodologies designs bundles such that operational objectives are considered.

This thesis introduces option bundles design approaches that not only maximize revenues, but also stabilize the demand for options. The results of our computational study indicate that a pure bundling policy inherently reduces product variety, irrespective of the car model for which bundles are designed. Option bundling can already be an effective tool for no bundle discount, if the customer base is homogeneous and the prices of the options have a wide range. For other cases, the benefits of option bundling can be attained if the bundle discount is selected carefully. Also, our computational study shows that the design of option bundles results in a trade-off between revenues and option demand stability. However, we identify bundles that simultaneously improve both measures. Bundle-specific discounts are shown to be not only an important instrument from a marketing perspective, but that they can impact the option demand variability.

Contents

Acronyms xi

Symbols xiii

1 Introduction 1

1.1 Current challenges of automotive manufacturers . . . 1
1.2 Product variety management as a cost reduction driver . . . 2
1.3 Impact of product variety on component demand forecasts . . . 3
1.4 The potential of option bundling . . . 4
1.5 Requirements for the design of option bundles . . . 5
1.6 Research objectives . . . 6
1.7 Thesis outline . . . 7
1.8 Included publications . . . 8

2 Related literature on option bundling 9

2.1 Effects of option bundles . . . 9
2.2 Option bundles design methodologies . . . 10

3 Product variety reduction through data-driven option bundling 13

3.1 Introduction . . . 13
3.2 Model description . . . 15
3.3 Branch-and-price approach for designing option bundles . . . 17
3.3.1 Column generation procedure and branching scheme . . . 18
3.3.2 Subproblem model and heuristic . . . 19
3.4 Computational study . . . 22
3.4.1 Case study data and implementation . . . 22
3.4.2 Performance of the matheuristic . . . 23
3.4.3 Analysis of the generated bundles . . . 25
3.4.4 Robustness analyses . . . 31
3.5 Conclusions . . . 37

4 Enhancing MRP-based component demand planning in a high-variety context 39

4.1 Introduction 39
4.2 The problem 41
4.2.1 Rise in supply chain complexity 41
4.2.2 Production planning at Mercedes-Benz cars 43
4.3 Component demand planning approaches: A review 44
4.3.1 MRP-based approaches 45
4.3.2 Alternative approaches 47
4.3.3 Problem synthesis 48
4.4 Model 49
4.4.1 'Data fusion' as a guiding analogy 49
4.4.2 Mathematical representation 50
4.5 Validation and implementation 54
4.5.1 Experimental design 54
4.5.2 Robustness checks 55
4.5.3 Implementation in a rolling-horizon context 57
4.6 Implementation at Mercedes-Benz Cars 60
4.7 Conclusion 61

5 Stabilizing the demand for car options by bundling **63**
5.1 Introduction 63
5.2 Problem description 64
5.3 Branch-and-price-based option bundles design 68
5.3.1 Column generation procedure and branching schemes 69
5.3.2 Subproblem models 70
5.3.3 Subproblem heuristic 75
5.4 Computational study 75
5.4.1 Computational analysis 77
5.4.2 Revenues and take-rate stability 78
5.4.3 Impact of discount flexibility 83
5.5 Conclusion 83

6 Conclusion and future research directions **87**
6.1 Summary of findings 87
6.2 Future research directions 89

A Outline of the branch-and-price approach **93**

B Outline of the subproblem heuristic subroutine 95

Bibliography 97

List of Tables

3.1 Structure of the provided past purchases data 23
3.2 Number of options, customers and months considered in the case study . 23
3.3 Computational performance of the approaches 24
3.4 Average computation time, share of optimal test instances, average and maximum optimality gap 26
3.5 Correlation of the number of car variants and revenues for the three car models 27
3.6 Number of car variants and revenues for the maximum revenue bundle design and for the first bundle design with revenues below those in unbundling 28
3.7 Performance of the bundles generated for $\rho^* = 1.1$ at an actual $\rho^* = 1.15$ 32
3.8 Performance of the bundles generated for $\rho^* = 1.1$ at an actual $\rho^* = 1.05$ 32

4.1 Evolution of product variety across Mercedes-Benz E-class product generations 42
4.2 Schedule accuracy and computational performance of the component demand planning methods 56
4.3 Schedule accuracy and computational performance of the component demand planning methods for the rolling horizon context datasets 59

5.1 Number of past customer purchases and months included in the case study 76
5.2 Average computation time, share of optimal test instances and average optimality gap for the small and the realistic test instance 77
5.3 Pareto hull surface coverage ratio 83

List of Figures

3.1 Computation time for the small dataset 25
3.2 Change of the number of car variants compared to the unbundling case . 28
3.3 Revenues change compared to the unbundling case 29
3.4 Pareto curve of the number of car variants and revenues compared to unbundling for different willingness to spend levels 31
3.5 Absolute difference spectrum for the number of car variants and revenues between the reservation-price-based model with customer-independent reservation prices and our customer behavior model 35
3.6 Absolute difference spectrum for the number of car variants and revenues between the reservation-price-based model with customer-specific reservation prices and our customer behavior model 36

4.1 Production volume, model line-up and global assembly plants at Mercedes-Benz cars, 1980-2015 42
4.2 Statistical analysis of the options-to-components dependency used to determine component-level demand for the Mercedes CLS-Class model . . . 43
4.3 A/F plot for the robustness analysis datasets 57
4.4 Absolute percentage error spectrum for the component demand planning methods in the robustness analysis datasets 58
4.5 A/F plot for the rolling horizon context 59
4.6 Stability measure spectrum for the rolling horizon context 60

5.1 Option take-rates variability for a volume car model 63
5.2 Bundles designed for model M1 for a pure bundling policy and a bundling intensity of 47% 79
5.3 Pareto frontiers for the pure bundling case 80
5.5 Ratio of surface covered by the Pareto frontiers of each bundling intensity compared to the hull 81
5.6 Pareto frontiers for the mixed bundling case 82

5.7 Ratio of surface covered by the Pareto frontiers for mixed bundling compared to the hull . . . 82

Acronyms

AIAG	Automotive Industry Action Group.
APO	Advanced Planner and Optimizer.
APS	Advanced Planning and Scheduling.
BOM	Bill of Materials.
BS1	Ryan and Foster branching scheme.
BS2	Branching scheme for the approach in chapter 5.
DF	Data fusion-based algorithm.
ERP	Enterprise Resource Planning.
LP	Linear Program.
MAD	Mean Absolute Deviation.
MAPE	Mean Absolute Percentage Error.
MRP	Material Requirements Planning.
ODETTE	Organisation for Data Exchange by Tele Transmission in Europe.
Thiel's U	Thiel Inequality Coefficient.

Symbols

B	Set of bundle candidates.
C'	Set of past orders.
C	Set of past customers.
D	Maximum bundle discount.
F_1 (also F_2)	Objective function values of the option bundles design approach in chapter 5.
F_2	*See* F_1.
K_o	Past take-rate of option o.
$LBound$	Lower bound of the branch-and-price approach.
L	Set of branching constraints of type BS1.
M'	Maximum price of a bundle that contains m options.
M_c^*	*See* M_c.
M_c (also M_c^*)	Big, customer-specific numbers.
$MaxObj2$	Subproblem-specific upper bound of the revenues generated by a single bundle.
$MaxObj$	Overall upper bound of the revenues generated by a single bundle.
$MinRed$	Minimum reduction of the revenues due to the dual variable values used for the subproblem.
Obj	Objective function value of the relaxed master problem.
O	Set of options.
P	Price of the new bundle.
$S_{c,b}$	1 if customer c would select bundle b, 0 otherwise.
T	Set of time periods.
$UBound$	Upper bound of the branch-and-price approach.
$V_{o,t}$	Planned take-rate of option o for period t.

Z	Set of components relevant for the Material Requirements Planning process.
$\Delta^{+}_{o,t}$, (also $\Delta^{-}_{o,t}$)	Auxiliary variables that represent the absolute take-rate change of option o from period $t-1$ to period t.
$\Delta^{-}_{o,t}$	*See* $\Delta^{+}_{o,t}$.
Δ^{TT}_{b}	Take-rates variability of all options included in bundle b.
Φ_c	Revenues generated by customer c due to the new bundle.
Π	Set of dual variable values corresponding to the constraints that restrict the allocation of each option to only one bundle.
Θ	Set of bundles relevant for the branching scheme BS2.
α	Multiple of the discount increment used for the bundle price.
$\beta^{+}_{b'}$, (also $\beta^{-}_{b'}$)	Deviations of the price of the new bundle from the price of bundle b'.
$\beta^{-}_{b'}$	*See* $\beta^{+}_{b'}$.
η	Bundle discounts increment.
$\iota_{z,c'}$	1 if past order c' included component z, 0 otherwise.
κ_b	Objective function value contribution of bundle b.
λ_b	1 if bundle b is offered, 0 otherwise.
$\nu_{c'}$	Frequency of past order c'.
ω	Weight for the deviation from the planned take-rates.
$\phi'_{b'}$	Price of bundle b' that should not be generated again, -1 if no bundle with the structure of bundle b' should be generated.
$\phi_{b'}$	Price of the bundle b' relevant for the branching scheme BS2.
π^{epsilon}	Dual variable value corresponding to the ϵ constraint.
π^{noBund}	Dual variable value corresponding to the constraint that limits the number of bundles.

π^{opt}	Dual variable value for option o, corresponding to the constraints that restrict the allocation of each option to only one bundle.
$\psi_{b'}$	1 if the price of the generated bundle is larger than the price of bundle b' , 0 otherwise.
$\rho'_{o,c}$	Reservation price of customer c for option o.
ρ^*	Estimated common willingness to spend for a bundle of customers.
ρ_c	Willingness to spend of customer c for a bundle: percentage of the price of the options purchased by the customer c in the past that are included in a bundle.
τ	Number of Pareto frontier points to generate for the take-rate stability objective.
$\theta'_{o,c}$	1 if customer c selects option o as part of the new bundle, 0 otherwise.
$\theta_{o,t}$	Number of customers who select option o in period t.
$\tilde{N}_t$	Planned production volume for period t.
$\zeta_{(t)}$	Set of customers who have placed their orders in period t.
$a_{c',o}$	1 if past order c' contained option o, 0 otherwise.
b'	*see* b.
b (also b')	Bundle indexes.
c'	Past orders index.
c	Customer index.
d'_b	Bundle discount of bundle b.
$d^{\text{comp}}_{z,t}$	Demand for component z in period t.
d	Common bundle discount for all bundles.
f'_o	Structure of the bundle b' that is relevant for the branching scheme BS2.
$f_{o,b}$	1 if option o is included in bundle b, 0 otherwise.
$feas^+$	Positive deviation from the required number of bundles n.
$feas^-$	Negative deviation from the required number of bundles n.

$h_{o,t}$	Take-rate of option o in period t.
l	Set of active BS1 branching constraints.
m	Maximum number of options in a bundle.
n	Number of bundles to design.
o_1	*See* o.
o_2	*See* o.
o (also o_1, o_2)	Option index.
p_o	Original selling price of option o.
q	Penalty from deviating from the desired number of bundles n.
r_b	Total revenues generated by bundle b if offered.
$s_{o,c}$	1 if option o was selected in the past by customer c, 0 otherwise.
t	Time periods index.
u	1 if the bundles are designed for mixed bundling, 0 for pure bundling.
$w_{c',t}$	Adjusted frequency of the past order c' for period t.
x_o	1 if option o is included in the new bundle, 0 otherwise.
y_c	1 if customer c acquires the new bundle, 0 otherwise.
z	Components index.

Chapter 1.

Introduction

1.1. Current challenges of automotive manufacturers

Dark clouds are looming on the horizon for the automotive manufacturers. The automotive industry has been one of the most succesful industrial sectors, having produced around 94 million cars per year (OICA, 2017a) and providing jobs for around 9 million people (OICA, 2017b). However, automotive manufacturers are confronted with many challenges that jeopardize their success: a fundamental shift in customer attitude towards car ownership, more strict environmental regulations for vehicles and the disruptive impact of emerging technologies in cars.

Increased environmental awareness and the high costs of car ownership have driven many consumers to embrace shared mobility (Efthymiou and Antoniou, 2016). The novel mobility concept enables its users to exploit the benefits of car ownership without actually owning a car. However, the steady sales increase experienced by automotive manufacturers so far could be reduced, should the popularity of shared mobility increase further. According to Gao et al., 2016 the yearly sales growth rate will most likely be reduced from 3.6% between 2012 and 2016 to 2% by 2030.

Additionally, automotive manufacturers need to cope with more strict environmental standards. Whereas the target for CO_2 emissions for vehicles imposed by the European Union in 2016 was 118.1 grams per kilometer, the automotive manufacturers need to ensure by 2021 that the vehicle emissions do not exceed 95 grams per kilometer (European Comission, 2017). To achieve these targets, automotive manufacturers must not only invest in the development of cleaner engines, but also in new engine technologies.

However, the pace at which traditional automotive manufacturers have embraced alternative drives, such as electrical or hydrogen-based, was so slow that new players, such as Tesla, managed to gain a competitive advantage and introduce attractive cars with alternative drives first to the market. Furthermore, automotive manufacturers face a stiff competition in the development of autonomous vehicles. Their biggest competitors are companies that have stronger competences in the development of artificial intelligence, such as Google or Apple.

Under these circumstances, the automotive manufacturers need to invest many resources only to keep up with their new competitors and to adhere to the strict environmental requirements. From 2006 to 2016, the investments of the 10 top automotive manufacturers in capital, research, as well as mergers and acquisitions increased on average by 4% (Parkin et al., 2017). Partly as a result of such investments, the profit margin of automotive manufacturers remained low in comparison to other industrial sectors. Whereas the annual return rates of the S&P 500 and Dow Jones Industrial Average companies were 14.8%, respectively 10.1% in the last five years, the automotive manufacturers generated only a 5.5% return on investment (Parkin et al., 2017). The higher investment requirements and the reduced revenue growth prospects jeopardize the capability of automotive manufacturers to generate satisfactory returns on investments for their shareholders. It is therefore paramount for automotive manufacturers to reduce their costs.

1.2. Product variety management as a cost reduction driver

One of the main cost drivers for automotive manufacturers is the high level of product variety they offer to their customers. When customers configure their vehicles, they express their individual preferences by choosing from a large number of pre-defined features, so-called options. The high levels of product variety that are induced by the options increase the costs throughout the life cycle of their products and reduce the responsiveness of the supply chain. Excessive product variety results in high inventory levels of unsold products, long customer lead times (ElMaraghy et al., 2013) and can even alienate the customers, cognitively overloading them with a plethora of choices (Huffman and Kahn, 1998; Lancaster, 1990; Ramdas, 2003). Costs can be reduced if the level of product variety offered to customers is lowered. According to Gertz and Haeser, 2015, a well-managed product variety could reduce service, engineering and product costs by more than 20%.

As customization is often taken for granted nowadays, the management of the resulting product variety remains a central challenge for researchers and practitioners alike. The seminal paper of Pil and Holweg, 2004 classifies methodologies aiming at a mitigation of the negative impact of product variety into four categories: modularity, product platforms, late configuration, and option bundling. Late configuration is not used by German automotive manufacturers (Staeblein and Aoki, 2015), since they cannot ensure a high quality of the assemblies done by their dealers. Modularity and product platforms in particular have been implemented by many manufacturers with the aim to reduce design complexity and manufacturing costs and to adapt to changing customer preferences. However, the two techniques did not always live up to these expectations. The design of modular products is burdened by the requirement for a deep organizational integration. In addition, many modules are only weakly integrated, which prevents the reduction of part numbers and thereby increases design and manufacturing costs (ElMaraghy et al., 2013). At the Volkswagen group, platforms led to limited product differentiation and did not achieve the desired adaptation flexibility (ElMaraghy et al., 2013). As a result, product complexity management is still viewed by 50% of the operations executives in the automotive industry as an important challenge, while 19% evaluate it even as the key challenge (Hanna and Kuhnert, 2016).

1.3. Impact of product variety on component demand forecasts

In the automotive industry, the accurate prediction of the required number of components is particularly challenging. The installation of components in vehicles is related to the options the customers choose for the cars they order. An analysis of the options interdependencies for 8,592 components of a Mercedes model has shown that the installation of 82.3% of the components in a car depended on the choice of at least one option. Under these circumstances, the accuracy of the component demand forecast is influenced by the accuracy of the forecast for the demand of options.

Automotive manufacturers would benefit from an accurate estimation of the component demand. The component demand forecasts are used in the annual sales and operations planning cycle as an input for the definition of the supply contracts and component supply plans (Jana and Grunow, 2017). Inaccuracies in the component supply plans are propagated to other planning processes such as budget planning. The inaccurate forecasts also have a detrimental impact on short-term supply process performance.

Last-minute component demand changes increase the cost of automotive manufacturers by over $ 1 billion per year. The elimination of component order cancellations alone would generate savings of up to $ 500 million (Dharmani et al., 2015).

The miscommunication of accurate component demands by the automotive manufacturers negatively impacts the cooperation between automotive manufacturers and their suppliers. Suppliers receive from their customers coarse demand plans months in advance. As more accurate customer order information is received by the automotive manufacturers, they update the component demand plans monthly. Most often, the component demand updates result in significant changes in the quantity of components ordered. As a result, the suppliers simply do not trust the forecasts communicated by their customers (Dharmani et al., 2015) and consider their relationship with automotive manufactures as 'adequate' at best (Buchholz, 2016).

Automotive manufacturers attempt to improve the accuracy of the component demand forecasts by deriving them from two types of separate forecasts. One type of forecast predicts the total volume, a second type of forecast predicts the share of vehicles containing individual options, also known as 'take-rates' (Meyr, 2004). However, the heterogeneity of the customer preferences induces a high volatility in the option take-rates.

1.4. The potential of option bundling

Option bundling, the sale of two or more options of a product as a package (Stremersch and Tellis, 2002), aims to reduce product variety without altering the design of products. It is a flexible methodology that can be used at any point during the product life cycle. Some mining equipment manufacturers sell their machines together with spare parts. Online shops like Amazon provide their customers with recommended bundles of products based on their past purchases. The computer manufacturer Dell prices complete PC systems and notebooks at a lower price than customized systems.

Bundling has widely been claimed to have a positive effect on operations due to a reduction of product variety (Fisher and Ittner, 1999; Pil and Holweg, 2004). However, a thorough analysis of the impact of option bundling on product variety is lacking. Previous literature describes the effects of option bundling on product variety only qualitatively or based on small-scale studies of limited practical relevance. Furthermore, no research so far has investigated the potential of option bundling to stabilize the demand of the options and thereby improve the accuracy of the option demand forecasts.

1.5. Requirements for the design of option bundles

The design of bundles is often done in an ad-hoc manner, based on the opinion of a small number of experts or, as recognized by Rickard, 2008, based on anecdotal evidence that customers would be interested in the purchase of bundles. None of the existing option bundling methodologies designs bundles such that the operational objectives are also considered. However, automotive manufacturers would benefit tremendously from a design methodology that manages to balance the marketing and the operational perspective.

Appealing bundles can be designed based on actual insights into customer preferences. For existing options, customer preferences can be deducted from information about past customer purchases. In contrast to the results from consumer studies, data on past customer purchases reflects a large number of actual customer decisions and is readily available. The drawback of purchase data is that often it only illustrates which option combinations were appealing to customers when they could purchase options individually. The introduction of bundles restricts customer choices in the case of pure bundling and extends customer choices in case of mixed bundling. The reaction to these changes can be predicted by means of customer behavior models, such as discrete choice models. However, in order to provide accurate reservation price estimations, these models require data sets with multiple sales prices for each option and detailed demographic information, which are often not available.

Option bundling methodologies that leverage the large number of past customer purchases available are likely to result in stable revenue and take-rate effects, even for customers not considered in the design phase. At the same time, bundles should not result in complex decisions for customers. As noted by Wendt, 2016, customers are more interested in buying if the selection process is simple and does not require too much attention.

Various option bundling policies, such as pure bundling, mixed bundling or unbundling, can be implemented (Adams and Yellen, 1976; Pierce and Winter, 1996). When pure bundling is implemented, options are sold only as part of bundles. Mixed bundling is the policy of selling options both as part of bundles as well as individually. No bundles are offered in an unbundling setting. Pure and mixed bundling can be combined with discounts to increase the attractiveness of bundles.

1.6. Research objectives

The overall aim of the thesis is to analyze the potential of option bundles to reduce product variety and improve component demand forecasts. In this regard, we also develop option bundles design approaches that could actually be used to bundle dozens of options based on the preferences of thousands of customers. Even though the problems analyzed in the thesis are relevant for an automotive setting, they can be related to other industrial sectors in which highly configurable products that require a large number of components from suppliers are manufactured.

The number of manufactured product variants has a large influence on all other operational measures. As our literature review shows, the impact of option bundles on the number of car variants has not been studied from a quantitative standpoint. Unfortunately, modeling complexities hinder the integration of the minimization of the number of car variants as an objective in option bundles design approach. However, we believe that the approaches that focus on standard objectives, such as the maximization of revenues, can actually impact product variety. Therefore, the first research question of the thesis is:

RQ1: Do option bundles result in a reduction of product variety and, if yes, which factors influence the magnitude of the reduction?

The approaches that are used by automotive manufacturers to plan the component demand rely on the classic Material Requirements Planning (MRP) logic that is not suitable for a high product variety environment. The approach of Stäblein, 2008 enhanced standard MRP approaches by 'fusing' customer preferences information and option take-rate forecasts to derive component demand plans for the suppliers. Stäblein, 2008 demonstrated the superiority of the enhanced MRP approach only for one realistic test instance. However, the robustness of the results generated by the approach and its applicability for a typical rolling horizon planning cycle were not studied. The next research question therefore addresses the gap:

RQ2: Does the enhanced MRP approach developed by Stäblein, 2008 deliver robust results and can it be embedded in a rolling horizon planning cycle?

A prerequisite for the approach of Stäblein, 2008 is that the take-rate forecasts are accurate. However, the heterogeneity of the customer preferences hinders the accurate estimation of option take-rates. We believe that the introduction of option bundles can stabilize the option take-rates enough to improve options demand forecast accuracy. However, there are no option bundle design approaches that include take-rate stability as an objective. The final research question of the thesis is then:

RQ3: How can a stabilization of the option take-rates be included in the design of option bundles and under which circumstances will bundles lead to a stabilization of the take-rates?

1.7. Thesis outline

The thesis is based on three research papers that address the research objectives stated in section 1.6. Chapter 2 provides an overview of the literature on the effects of option bundles, as well as of the existing option bundles design methodologies.

Chapter 3 quantifies the impact of option bundling on product variety for three car models of a large German automotive manufacturer. The chapter therefore addresses the first research question. Since we cannot integrate the minimization of product variety in an option bundles design methodology, an approach that only maximizes revenues is introduced. We then apply the approach to evaluate the capability of bundles to reduce product variety, as well as the robustness of the bundles to different assumptions regarding customer behavior.

Chapter 4 addresses the second research question by evaluating the performance of the component demand planning methodology of Stäblein, 2008 for multiple test instances, including for a rolling horizon planning cycle. In the paper, we first present the shortcomings of the existing MRP approaches for a high-variety context. We then present the enhanced MRP methodology developed by Stäblein, 2008 that fuses information from multiple data sources. We compare the performance of the methodology to the performance of a standard MRP approach, as well as of a state-of-the-art time-series software package.

Since the method of Stäblein, 2008 requires accurate forecasts for the demand of options as input, we investigate in chapter 5 the potential of option bundles to stabilize the option demand patterns and thereby improve the option take-rate forecast accuracy. We present the option bundles design model and the approach that balances revenues and take-rate stability. In the numerical tests we investigate the impact of the bundles on the trade-off between the two measures for a pure and a mixed bundling case.

1.8. Included publications

The chapters in this thesis are based on individual publications that are readable as individual contributions. Combined, the papers illustrate the potential of option bundling to improve operational measures, such as product variety or take-rate stability. Together with the methodology of Stäblein, 2008, option bundling has the potential to increase component demand forecast accuracy. The chapters are based on the following sources:

Chapter 2 Synthesis of the literature reviews in Popa et al., 2017 and Popa and Grunow, 2017

Chapter 3 Popa, R. C. et al. (2017). "Product variety reduction through data-driven option bundling"

Chapter 4 Stäblein, T. et al. (2016). "Enhancing MRP-based component demand planning in a high-variety context"

Chapter 5 Popa, R. C. and M. Grunow (2017). "Stabilizing the demand for car options by bundling"

Chapter 2.

Related literature on option bundling

This chapter is a synthesis of the literature reviews in Popa et al., 2017 and Popa and Grunow, 2017.

Research on option bundles has been devoted so far in two directions: the impact of option bundles and the design methodologies. We provide an overview of both literature streams.

2.1. Effects of option bundles

Since the introduction of the option bundling concept, research has mostly been devoted towards determining its impact from a marketing and an economics perspective. The perspectives on the impacts of bundling can be grouped into three categories: the impact on consumer purchases and revenues, the impact on product quality and forecasts, and the impact on buffers and costs.

The perspective in the literature regarding the impact of bundling on consumer purchases and revenues, such as Stigler, 1963, shows how bundling can increase the profit of the seller when consumer valuations are negatively correlated. Adams and Yellen, 1976 illustrate with a two-dimensional graphical framework and stylized examples that pure and mixed bundling can in some cases be more profitable than unbundling. Schmalensee, 1984 confirms the profitability of bundling and the benefits of mixed bundling over pure bundling for a negative demand correlation between the options. Behavioral customer research on bundling and the corresponding price framing are analyzed by Soman and Gourville, 2001. They explain price bundling effects by augmenting the economic model with a measure of 'how good a deal' the customer is getting through a bundle acquisition. Central to their behavioral approach is the calculation of 'gains and losses' relative to a set of reference points (which differ from the reservation prices used in economic models) and the fact that 'losses' are more detrimental than corresponding 'gains'. Cao et al., 2015 show that bundling could be especially beneficial when at least one of the products included in the bundle is in limited supply and the valuation of the products is positively correlated. Chakravarty et al., 2013 illustrate that bundling can increase the profit of a supply chain for high margin products with similar valuations and a low demand correlation if the supply chain partners coordinate to eliminate a double marginalization.

The second perspective on bundling addresses the impact of bundling on product quality and forecasts. Pil and Holweg, 2004 discuss empirical evidence from automotive manufacturers that bundling can reduce forecast error and thus the stock obsolescence risk. They find that manufacturers can also greatly simplify their whole distribution system by offering options as coherent bundles rather than offering all possible option permutations.

The third perspective highlights the potential of option bundling to reduce buffers and costs. Fisher and Ittner, 1999 recognize that bundling can reduce the required buffer capacity inside a manufacturing plant. Ringbeck et al., 1999 state that the introduction of option bundling can improve manufacturing productivity and lead to manufacturing cost savings via economies of scale. Additionally, the 'complexity costs' resulting from managing product variety in the production process are reduced. The results are confirmed by Bitran and Ferrer, 2007 for the high-tech manufacturing industry. Eppen et al., 1991 report the profit improvement potential of option bundling for the case of an automotive manufacturer.

The impact of option bundles on forecasts has been studied qualitatively, whereas there is no literature on the impact of option bundles on product variety. Whereas the literature shows that option bundling reduces costs, there are no guidelines for the integration of these effects in option bundles design methodologies.

2.2. Option bundles design methodologies

The few existing option bundling methodologies can be categorized based on the type of input required: one set of methodologies requires reservation prices, whereas the other methodologies require as inputs the 'product attractiveness'.

In Hanson and Martin, 1990 and Wu et al., 2008, the reservation prices for the products being bundled are known, whereas in Tönshoff et al., 1999 and Fuerderer et al., 1999, the reservation prices for the candidate bundles are known. All these methodologies design bundles using mathematical models. Whereas the model of Hanson and Martin, 1990 can generate any bundle, the models of Tönshoff et al., 1999 and Fuerderer et al., 1999 determine which bundles to select from a list of bundle candidates. The model of Wu et al., 2008 only specifies the number of products the customers can select to design their own customized bundle. For all methods, the prices of the bundles can be defined individually.

Chung and Rao, 2003, Bitran and Ferrer, 2007 and Cataldo et al., 2017 require as inputs the 'attractiveness' of the products included in the bundle with regards to a number of predefined attributes. These methodologies employ discrete choice models to derive the probability that customers would select certain bundles. Bitran and Ferrer, 2007 extended the method of Chung and Rao, 2003 by modelling a competitive environment. Cataldo et al., 2017 extended the methodologies of Chung and Rao, 2003 and Bitran and Ferrer, 2007 such that a pre-specified number of bundles could be generated.

There are a number of methodologies that do not maximize the profit of the bundles seller. The objective of Dixon and Thompson, 2016 is to schedule and bundle various events on different stages so that the customer satisfaction with the bundles they select is maximized. The method of Proano et al., 2012 determines bundles of vaccines and their prices such that the surplus society obtains from them is maximized.

All option bundles design methodologies except Dixon and Thompson, 2016 only focus on a single objective: on the improvement of profits, revenues, satisfaction or social surplus. However, none of the methods can integrate two conflicting objectives. The methodologies do not even incorporate operational measures as objectives. They do not highlight the influence of bundling on actual product variety. The existing option bundling approaches focus more on the pricing of the bundles rather than the composition design. With regards to input requirements, most methodologies need a considerable amount of data or data which can not be easily gathered. For Chung and Rao, 2003, Bitran and Ferrer, 2007 and Cataldo et al., 2017 the attractiveness of the products included in the bundles with regards to certain attributes needs to be estimated. Furthermore, the methods of Wu et al., 2008, Proano et al., 2012 and Dixon and Thompson, 2016 are not suitable for a manufacturing scenario. The methodology of Wu et al., 2008 can only be used for products with low variable costs, such as software. The methodology of Dixon and Thompson, 2016 is very customer-centric and cannot integrate an operational objective. The method of Proano et al., 2012 is only suitable for a pharmacoeconomical setting. Moreover, none of the option bundling approaches is capable of handling realistic test instances.

Chapter 3.

Product variety reduction through data-driven option bundling

This chapter is based on an article submitted as:

Popa, R. C. et al. (2017). “Product variety reduction through data-driven option bundling”

3.1. Introduction

The mantra of many manufacturers has been to enable extensive product customization such that customers can tailor products to their individual needs. However, offering customers an ever increasing range of choices does often not go along with increased competitiveness (Alptekinoğlu and Corbett, 2008). Dell was once the prime example of a manufacturer with high levels of product variety. However Dell has moved towards preconfigured personal computers, mainly because their online configuration system had become too complex and costly. Levi Strauss allowed its customers to customize jeans between 1993 to 2003, but abandoned the project due to the poor results. Samsung, offered a wide range of smartphone models and managed to obtain a 17% operating margin on their smartphones sales in the second quarter of 2016. Apple, in contrast, generated a 38% operating margin in the same period on their narrower range of smartphones.

The automotive industry is an example of a sector in which customers can choose between dozens of options, such as multimedia and driving assistance systems or aesthetic and safety enhancements, to customize their vehicles. The wide spectrum of available choices results in a large number of potential car variants ($2.1 \cdot 10^{20}$ for volume vehicles like the Volkswagen Golf, $4.7 \cdot 10^{24}$ for luxury vehicles such as the Mercedes C-class) (Staeblein and Aoki, 2015). The number of car variants actually built is lower than the theoretical number, but still causes significant efficiency and responsiveness losses. This is a key motivation for many automotive manufacturers to introduce option bundles in their (online) car configurators.

Despite their significant impact on revenues, the design of bundles is often done in practice in an ad-hoc manner, based on the opinion of a small number of experts. Structured option bundling approaches exist that use actual insights into customer preferences and especially into the appeal of product combinations. However, these approaches often require input data that is difficult to collect or estimate, such as the willingness to pay of customers for individual options, also known as reservation prices. Traditionally, estimates for reservation prices are obtained from conjoint analyses. The drawback of conjoint analyses is that they require costly customer studies. When the number of options is large, such consumer studies may not be feasible.

For existing options, customer preferences can alternatively be deduced from information about past customer purchases. In contrast to the results from customer studies, data on past customer purchases reflects a large number of actual customer decisions and is readily available. The drawback of purchase data is that often it only illustrates which option combinations were appealing to customers when they could purchase options individually. The customer behavior changes caused by the introduction of bundles can be predicted by means of customer behavior models, such as discrete choice models. However, these models require data sets with multiple sales prices for each option, which are often not available. Nevertheless, this paper still aims to exploit the available purchase data without reliance on customer models that have excessive input data requirements.

In this paper, we focus on pure bundling, as it is the bundling policy that has the potential to reduce product variety the most. It forces customers to decide whether they buy all options in a bundle or none at all.

Our main contribution is the quantification of the impact of option bundling on product variety for realistic settings. We develop a data-driven parallel bundling approach that is capable of generating a pre-defined number of bundles for realistic test instances. The effort connected to designing bundles is limited by using data that is often readily available: the prices of the options and past customer purchases. Manufacturers can use our approach for an efficient evaluation of the impact of the bundle number and the bundle discount rate on the number of car variants and revenues. For a case study in the automotive sector, we show that option bundling is an effective product variety mitigation strategy. Manufacturers can choose between many bundle designs that reduce the number of car variants and simultaneously increase revenues. While no bundle discounts are required for a customer base with homogeneous preferences, the discounts must be carefully selected for a customer base with heterogeneous preferences. The case study results also show that the proposed bundle design approach is robust. The performance of the determined bundles is impervious to customers showing reactions different from those predicted by the employed customer behavior model.

In 3.2 we describe the mathematical model for designing bundles that only maximize revenues given the current customer behavior model. We present the adapted branch-and-price approach for generating bundles in section 3.3. Section 3.4 contains an analysis of the effects the bundles generated by our approach on product variety and revenues for various settings. We conclude with managerial and theoretical implications of our work as well as potential extensions in section 3.5.

3.2. Model description

Considering the limitations of the existing option bundling approaches, we devised an approach that generates a required number of bundles based on past customer purchases. The approach does not necessitate the estimation of detailed customer information such as option-specific reservation prices. We describe the behavior of a customer when being offered bundles using only one parameter: the willingness to spend of a customer for a bundle. The parameter is expressed as a percentage of the price of the options that are part of the bundle and that were acquired in the past by the customer. When each option is allocated to only one bundle, the customer purchases a bundle if the willingness to spend for a bundle exceeds the bundle price. Based on our model, customers are motivated to purchase options they did not in the past, if the willingness to spend for a bundle that contains such options exceeds the bundle price. Bundle discounts can lower the price of a bundle such that the bundle becomes attractive for customers. We show in our numerical experiments that such a simplified description of customer behavior is sufficient for the purpose of bundling options. Similar results are obtained for more elaborate customer choice models that require significantly more detailed customer data.

In this paper, we allocate each option to one bundle. Since this policy has the potential to inherently reduce product variety, the bundles must be designed to maximize revenues. We therefore developed a binary linear model for designing option bundles such that revenues are maximized. An option can be assigned to multiple bundles if copies of that option are created and added to the set of options that need to be bundled. The purchase of a bundle that contains at least two options is rewarded with a discount. The discount rate is set in advance by the marketing department. The definition of its value is not within the scope of our model. We use the following notation:

Sets:	
$o \in O$	Set of options
$b \in B$	Set of bundle candidates
$c \in C$	Set of past customers
Parameters:	
n	Number of bundles to design
$f_{o,b}$	1 if option o is included in bundle b, 0 otherwise.
$S_{c,b}$	1 if customer c would select bundle b, 0 otherwise
r_b	Total revenues generated by bundle b if offered
ρ_c	Willingness to spend for a bundle: percentage of the price of the options purchased by the customer in the past that are included in a bundle ($\rho_c \geq 1$)
d	Common bundle discount for all bundles
p_o	Original selling price of option o
$s_{o,c}$	1 if option o was selected in the past by customer c, 0 otherwise
Decision variables:	
λ_b	1 if bundle b is offered, 0 otherwise.

The set of all possible bundles B consists of all possible combinations of the elements in set O. The binary matrix $f_{o,b}$ represents the bundle structure. If required, it is also possible to create a restricted set of bundles B. For example, the set could contain only the bundles with a number of options below a certain threshold.

Based on our customer model and since each option is allocated to only one bundle, it is possible to calculate in advance how much revenues each bundle generates, if it is offered. Let $S_{c,b}$ be an auxiliary parameter which takes on the value 1 if customer c would acquire bundle b, 0 otherwise. For bundles that contain at least two options, the value of the parameter is determined based on the following expression:

$$S_{c,b} = \begin{cases} 1 & \text{if } \sum_{o \in O} \rho_c \cdot p_o \cdot s_{o,c} \cdot f_{o,b} \geq \sum_{o \in O} p_o \cdot (1-d) \cdot f_{o,b} \\ 0 & \text{otherwise} \end{cases} \tag{3.1}$$

A customer c acquires a bundle b if the price of the bundle ($\sum_{o \in O} p_o \cdot (1-d) \cdot f_{o,b}$) does not exceed the predefined percentage (ρ_c) of the value of the options acquired in the past that are included in the bundle ($\sum_{o \in O} \rho_c \cdot p_o \cdot s_{o,c} \cdot f_{o,b}$).

After determining for each customer if she will acquire the bundle, the total revenues r_b generated by bundle b, should it be offered, are determined using the following expression:

$$r_b = \sum_{c \in C} (S_{c,b} \cdot \sum_{o \in O} (1-d) \cdot f_{o,b} \cdot p_o) \tag{3.2}$$

After these preprocessing steps, the allocation of options to bundles can be represented mathematically as a simple binary knapsack problem:

Objective function:

$$\text{Maximize} \sum_{b \in B} r_b \cdot \lambda_b \tag{3.3}$$

Subject to:

$$\sum_{b \in B} f_{o,b} \cdot \lambda_b = 1, \ \forall o \in O \tag{3.4}$$

$$\sum_{b \in B} \lambda_b = n. \tag{3.5}$$

Objective function (3.3) maximizes the revenues that result from the bundles offered. Constraint (3.4) ensures that each option is included in only one of the offered bundles. Constraint (3.5) requires that n bundles are selected from set B.

3.3. Branch-and-price approach for designing option bundles

It is not possible to use the model presented in the previous section for realistic settings, as the size of set B grows exponentially in the set of options to be bundled. Even when reducing set B to contain all the bundles with an acceptable number of options, the number of elements in the set is still too large to be handled by a commercial solver. We therefore can determine the option bundle designs by means of a branch-and-price approach. The approach (based on Barnhart et al., 1998) iteratively adds bundles which improve the revenues that result from a starting solution. We describe the constituents of the approach in the following subsections (outline of the approach in Appendix A).

3.3.1. Column generation procedure and branching scheme

In our column generation procedure, we use an adapted version of the model defined by expressions (3.3) - (3.5). The adaptations of the decision variables are:

- The decision variables λ_b are relaxed (i.e. $\lambda_b \in [0,1]$ instead of $\lambda_b \in \{0,1\}$) to derive the dual variable values.
- Additional positive continuous decision variables $feas^+$ and $feas^-$ are defined to ensure that there will always be a feasible solution for the relaxed master problem.

The relaxed master problem model is presented in the following:

$$\text{Maximize} \sum_{b \in B} \lambda_b \cdot r_b - (|C| \cdot \sum_{o \in O} p_o + 1) \cdot feas^+ - (|C| \cdot \sum_{o \in O} p_o + 1) \cdot feas^- \tag{3.3'}$$

subject to:

Constraint (3.4)

$$\sum_{b \in B} \lambda_b = n + feas^+ - feas^- \tag{3.5'}$$

Constraint (3.5) is extended in constraint (3.5') to allow deviations from the required number of bundles. However, the deviations are penalized in the adapted objective function (3.3') to ensure that the required number of bundles is offered if possible.

After solving the relaxed master problem, the dual variable values π^{opt} corresponding to constraint (3.4) and π^{noBund} corresponding to constraint (3.5') are derived. The values are used in the subproblem to generate bundles which lead to an improvement of the objective function (3.3'). Such bundles are obtained either by means of a heuristic or by using a mathematical model (both described in section 3.3.2). We first use the heuristic and, if it fails to find any objective-improving bundle, then we use the subproblem model. An optimal solution of the relaxed master problem is confirmed when the subproblem model cannot find any objective-improving bundle.

If the optimal solution for the relaxed master problem is also feasible for the original problem (i.e. if there is no non-binary decision variable value) and is better than the best solution found so far, we replace the best solution with the current one. However, if the solution of the relaxed master problem is not feasible for the original problem, but the objective function value of the relaxed master problem is better than the best solution found so far, then we use a branching scheme based on Ryan and Foster, 1981 (named in the following BS1) to search for an integer solution. After all branches have been visited, the best solution found by the approach is the optimal solution for the model presented in section 3.2.

The branching scheme generates a list of candidate branches L. The elements of type $(o_1, o_2, 1)$ in the list represent the branches that enforce that options o_1 and o_2 are either both allocated or are both absent in the generated bundles. The elements of type $(o_1, o_2, 0)$ in the list represent the branches in which the bundles contain either option o_1, o_2 or none of the two options. The separation of the solution space is based on the pair of options o_1 and o_2 for which the branch $(o_1, o_2, 1)$ results in the highest revenues after eliminating the illegal bundles. All other candidate branches are removed from set L. The branch $(o_1, o_2, 1)$ is explored first, since it leads to an integer solution faster.

The branch-and-price approach can be also used as a matheuristic by skipping the search for a bundle using the subproblem model. If the subproblem heuristic cannot find any bundles to improve the objective, then the matheuristic directly checks whether further branching is needed.

3.3.2. Subproblem model and heuristic

The mixed-integer linear model for the subproblem generates the bundle that can improve the objective of the relaxed master problem the most, i.e. which has the highest positive reduced costs.

We use the following additional notation for the subproblem model:

Additional parameters:

m :	Maximum number of options in a bundle
M_c	Big, customer-specific numbers
M'	Maximum price of a bundle that contains m options

Decision variables:

x_o :	1 if option o is included in the new bundle, 0 otherwise
y_c :	1 if customer c acquires the new bundle, 0 otherwise
Φ_c :	Revenues generated by customer c due to the new bundle, $\Phi_c \geq 0$

Objective function:

$$\text{Maximize } \sum_{c \in C} P - \sum_{o \in O} \pi^{\text{opt}} \cdot x_o - \pi^{\text{noBund}} \tag{3.6}$$

The objective function (3.6) represents the maximization of the reduced costs of the new bundle. The first term represents the revenues generated when offering the new bundle. The second term represents the objective reduction resulting from the options included in the bundle.

The new bundle must fulfill the following constraints:

$$\sum_{c \in C} P - \sum_{o \in O} \pi^{\text{opt}} \cdot x_o - \pi^{\text{noBund}} > 0 \tag{3.7}$$

$$\sum_{o \in O} x_o \geq 2 \tag{3.8}$$

$$\sum_{o \in O} x_o \leq m \tag{3.9}$$

$$\sum_{o \in O : s_{o,c}=1} p_o \cdot x_o \cdot \rho_c - \sum_{o \in O} p_o \cdot x_o \cdot (1-d) + M_c \cdot (1-y_c) \geq 0, \forall c \in C \tag{3.10}$$

$$P \leq (1-d) \cdot \sum_{o \in O} p_o \cdot x_o, \forall c \in C \tag{3.11}$$

$$P \leq M' \cdot y_c, \forall c \in C \tag{3.12}$$

$$x_{o_1} = x_{o_2}, \forall o_1, o_2 \in O : (o_1, o_2, 1) \in l \tag{3.13}$$

$$x_{o_1} + x_{o_2} \leq 1, \forall o_1, o_2 \in O : (o_1, o_2, 0) \in l \tag{3.14}$$

Constraint (3.7) enforces that only the bundles which improve the objective function of the relaxed master problem are generated. Also, the constraint improves the computational performance of the model. Since all 'bundles' that contain one option are available at the start of our approach, constraints (3.8) and (3.9) enforce that the generated bundles should contain between 2 and m options. Constraint (3.10) represents the bundle selection mechanism of the customers. A customer will not select a bundle if the percentage value ρ_c of the options included in the bundle and that were purchased by her in the past is less than the price of the bundle. Constraints (3.11) and (3.12) set the correct revenue generated by the new bundle for each customer. A customer will pay the discounted price of the bundle if she purchases the bundle, 0 otherwise. Constraints (3.13) and (3.14) enforce the active branching constraints.

After finding the optimal solution for the subproblem model, the new bundle b and the total revenues generated by offering it, $\sum_{c \in C} P$ are added in the relaxed master problem model.

The model is also used to determine the upper bound of the branch-and-price approach before starting it. By removing the objective function terms which include the dual variables, the model determines the maximum revenue which can be obtained by offering a bundle. Since solving this model requires a considerable amount of computation time, we use the upper bound of the revenues found during a predefined amount of time to determine the upper bound for our approach. We define $MaxObj$ as the upper bound of the revenues, Obj as the objective function of the relaxed master problem, and $MinRed$ as the minimum reduction of the objective function due to the dual variable values. The upper bound for the reduced costs is then calculated by using the following expression (Lübbecke and Desrosiers, 2005):

$$Obj + n \cdot (MaxObj - MinRed) \tag{3.15}$$

This upper bound may not in fact be attainable by means of an actual bundle, since these reduced costs are only calculated based on the dual values of individual options, without simultaneously considering possible bundle compositions. In the first iteration of the column generation approach, the upper bound is tightened by obtaining the maximum reduced costs resulting by means of a bundle. To this end, we use the upper bound found by the solver for the complete subproblem model during a predefined amount of time. Let $MaxObj2$ be the subproblem upper bound value. The second upper bound for the approach is then:

$$Obj + n \cdot MaxObj2 \tag{3.16}$$

Solving the subproblem model is computationally expensive. To speed up the column generation process, we developed a simple greedy multicore heuristic subroutine. First, it determines the bundles that contain two options and that do not violate the active branching constraints. Each of these bundles is then processed by an available CPU core. Each bundle is iteratively expanded with the options that increase the reduced costs the most and do not violate the active branching constraints. The processing of a bundle on a CPU core is terminated as soon as the maximum number of options in a bundle has been reached, all options have been added to the bundle, or the reduced costs cannot be improved by adding an option. As soon as the processing of a bundle on a CPU core is done, the core starts processing one of the remaining bundles. The outline of the subproblem heuristic is illustrated in Appendix B.

If the subproblem heuristic cannot find any objective-improving bundle, the subproblem model is used in the optimizing branch-and-price approach to find the bundle that improves the objective function of the relaxed master problem the most.

3.4. Computational study

We tested our branch-and-price approach on data made available by a large German automotive manufacturer. The purpose of our experiments was to quantify the potential of option bundles to reduce the number of car variants. We define a car variant as a unique combination of options that results from the bundles selected by the customers. In addition, we analyzed whether the generated bundles would lead to an increase of revenues, independent of the underlying customer behavior model.

In section 3.4.1 we present the data made available to us, as well as the computational hardware used for the tests. Since the optimal branch-and-price approach required a lot of time to generate bundles, we used the matheuristic version of the approach and compared it with the optimizing approach on test instances with a small number of options. The results of the comparison are presented in section 3.4.2. In section 3.4.3 we evaluate the results of the branch-and-price matheuristic on a realistic case study that includes the preferences of thousands of customers. In section 3.4.4, we determine the robustness of our approach to input inaccuracies.

3.4.1. Case study data and implementation

Data for three car models was used for the numerical computational study: a volume model (M1), a niche model (M2) and a luxury model (M3). For each model, we were provided with a list of options to be bundled, a set of past customers and their option selections, as well as the point in time when the orders were placed. Table 3.1 illustrates the structure of the past purchases data. Table 3.2 presents the number of options and customers included in our analysis, as well as the number of months over which the customer orders were spread.

Order month	Order number	Options selected
1	A1	O5; O23; O42; O63; ...
1	A2	O1; O2; O23; O42; ...
...		

Table 3.1.: Structure of the provided past purchases data

Car model	M1	M2	M3
# options	44	47	50
# customers	150,194	39,194	36,025
# months	10	9	7

Table 3.2.: Number of options, customers and months considered in the case study

To ensure that the bundles were accepted by the marketing department of the automotive manufacturer, we generated bundle designs containing a number of bundles that varied between $|O| - 1$ and $\lceil |O|/2 \rceil$. We did not limit the number of options included in a bundle, since the approach did not generate large bundles.

We split the customer dataset into a training and a validation dataset. The training dataset was used by the branch-and-price approach. We evaluated on the validation dataset how the customers would react when confronted with the generated bundles. The purpose of the split was to determine the general validity of the effects of the bundles on the number of car variants and revenues. We used the purchase data from the last three months for each car model as validation data. The purchase data from all the remaining months were used as training data.

We ran all our tests on an octacore 2.1 GHz Intel Xeon machine with 16 GB RAM. We used IBM ILOG CPLEX 12.6 to solve the relaxed master problem and the subproblem models. The branch-and-price approach was implemented in C#.

3.4.2. Performance of the matheuristic

We first compared the performance of the branch-and-price matheuristic (Heur) to the original branch-and-price approach (Opt). We used all available customer data for the comparison. We used ten test instances for each car model, each with ten randomly selected options. We generated bundle designs containing between five and nine bundles. We set a time limit of ten hours per bundle design for the approaches. If an approach

Car model	M1		M2		M3	
Approach	Opt	Heur	Opt	Heur	Opt	Heur
% # of instances that require less than 10 hours	97.91%	100%	85.71%	97.95%	97.87%	97.87%
% # of instances with identical solutions for the matheuristic and the optimal approach	93.75%		95.91%		100%	
Average gap between the optimal approach and the matheuristic for the non-identical solutions	0.96%		-0.24%		0%	

Table 3.3.: Computational performance of the approaches

required more than ten hours, it provided as output the best design found so far. We let CPLEX search one hour for the upper bound of the revenues that can be generated by a single bundle and 15 minutes for the upper bound of the reduced costs generated by a single, column-generation-iteration-specific bundle.

Table 3.3 illustrates the share of test instances for which the approaches found a solution in less than ten hours, the share of test instances in which the matheuristic found a solution at least as good as the original approach, as well as the average gap between the optimal and the matheuristic solution.

The matheuristic required significantly less time to find solutions compared to the original approach. The mean computation time for the test instances requiring less than ten hours was 414.20 seconds for the optimal branch-and-price approach, whereas for the matheuristic was 11.52 seconds. Figure 3.1 highlights the spectrum of the computation times over all test instances.

The matheuristic identified the same solution as the optimal approach in at least 93% of the test instances. The revenues for model M1 were on average 0.96% lower for the matheuristic than for the optimal approach. For model M2, in almost 5% of the test instances the matheuristic found a solution that was better than the one provided by the optimal approach. In the ten hours of computation time, not only did the optimal approach not find the optimal solution, but it was also unable to find a better solution than the matheuristic.

The results highlight the efficiency of the simple parallel subproblem heuristic. The subroutine generated multiple bundles that in most cases lead to the optimal master problem solution in a small amount of time. Analyzing the iterations of the optimal branch-and-price approach, we observed that the subproblem model most often just

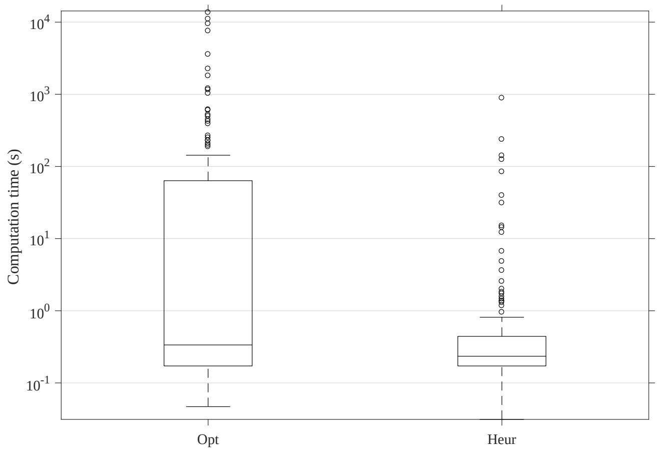

Figure 3.1.: Computation time for the small dataset

confirmed the non-existence of any bundle with positive reduced costs.

There were instances in which the branch-and-price matheuristic also had a poor computational performance. In these, the number of branches being evaluated was extremely large. However, since our matheuristic does not ever generate an optimal solution for the subproblem, we could not use dual bounds such as the one in Vanderbeck, 2011 to further prune the branches considered. We identified an upper bound only in the root node of the branch-and-price approach, since this process is computationally intensive.

3.4.3. Analysis of the generated bundles

Since the matheuristic performed very well for the small datasets, we used it for every car model to generate bundles based on all options provided. We analyzed the computational performance of our approach and the relevance of the generated bundles for real-life situations. We also identified the impact of option bundling on the number of car variants and the revenues for different discount levels and as well as willingness to spend levels.

As in section 3.4.2, we generated bundle designs with a number of bundles ranging between $\lceil |O|/2 \rceil$ and $|O| - 1$. The approach stored the best bundle design found in at most 10 hours. Based on discussions with colleagues from the automotive manufacturer, we used a common willingness to spend value $\rho_c = \rho^*$ of 110%. We ran additional tests for a ρ_c value of 105% and 115%. We generated bundles for various discount rates: 0 %, 10 % and 20 %.

Computational analysis and model validation

In Table 3.4, the average computation time, the share of test instances solved to optimality, as well as the average and the maximum optimality gap for suboptimal bundle designs are presented. The average computation time was less than 11,000 seconds. Despite the large size of the test instances, our approach generated many optimal bundle designs in less than 10 hours. Based on the upper bound values, we demonstrated that at least 50% of the generated configurations were optimal. When the approach did not guarantee optimality, the average optimality gap was lower than 0.3 %.

Car model	M1	M2	M3
∅ computation time (s)	10799.18	8419.79	5049.98
% test instances solved to optimality	56.36%	59.13%	55.20%
∅ optimality gap for suboptimal bundle designs	0.22%	0.09%	0.29%
Maximum optimality gap for suboptimal bundle designs	1.26%	0.66%	1.20%

Table 3.4.: Average computation time, share of optimal test instances, average and maximum optimality gap

We also evaluated the practical relevance of the generated bundles. It was especially important not to have a large number of options grouped together in a bundle. For a number of bundles close to the number of options, the approach generated bundles containing two options. Examples are a bundle consisting of a right comfort seat and a moonroof as well as a bundle consisting of a leather steering wheel and a 7-gear automatic gearbox. As the number of bundles was reduced, the number of options included in a bundle increased, but never exceeded seven options. Examples for larger bundles are a bundle consisting of a 4-zone air conditioning system, an ashtray, and a fire extinguisher, as well as a bundle consisting of a left comfort seat with additional support, a tire pressure monitoring system, and exterior sports styling.

Number of car variants and revenues

To analyze whether the generated bundles would also be appealing to customers not considered during the bundling process, we measured the correlation of the number of car variants and the revenues between the training and the validation datasets (shown in Table 3.5). The high correlation values suggest that the bundles have similar effects on product variety and revenues when offered to customers not considered during the

Car model	M1	M2	M3
Correlation of # car variants	0.9995	0.9983	0.9600
Correlation of revenues	0.9902	0.9848	0.8257

Table 3.5.: Correlation of the number of car variants and revenues for the three car models

design phase. The next analyses are done only based on the validation datasets, to ensure that a good performance cannot be explained by the fact that the bundles were tailored for the customers considered in the analysis.

Figure 3.2 presents the changes of the number of car variants, whereas Figure 3.3 illustrates the changes of the revenues compared to the unbundling case for the three car models. The x-axis for both figures represents the 'bundling intensity'. We define bundling intensity as the percentage of options that are included in bundles of at least two options. The bundling intensity is calculated by using the following expression:

$$1 - \frac{n}{|O|} \tag{3.17}$$

Our results confirm that pure bundling decreases the number of car variants with increasing bundling intensity, irrespective of the car model considered and the bundle discount. For a bundling intensity of 50%, the number of car variants is reduced by 63% for model M1, 30% for model M2 and 27% for model M3. The number of car variants seems to have been reduced more for model M1 than for M2 and M3 due to the higher homogeneity of the customer base and higher range of the option prices for model M1. Whereas the ratio of the number of car variants in the unbundling case to the number of customers is 22.93% for model M1, for models M2 and M3 the ratio is 48.71%, respectively 39.92%. The approach used options that many customers wanted as a catalyst for the acquisition of options for which the customer preferences were mixed.

The key advantage of our approach is that it not only reduces the number of car variants, but also improves revenues. As shown in Table 3.6, the bundle designs that generated the highest revenues increased revenues by up to 4% compared to unbundling and reduced the number of car variants by up to 9%. Figure 3.3 illustrates that revenues increased as long as the bundling intensity did not exceed 44% for model M1, 45% for model M2, and 12% for model M3. When the bundling intensity was high enough that the revenues equaled those for unbundling, the number of car variants were reduced by up to 50%.

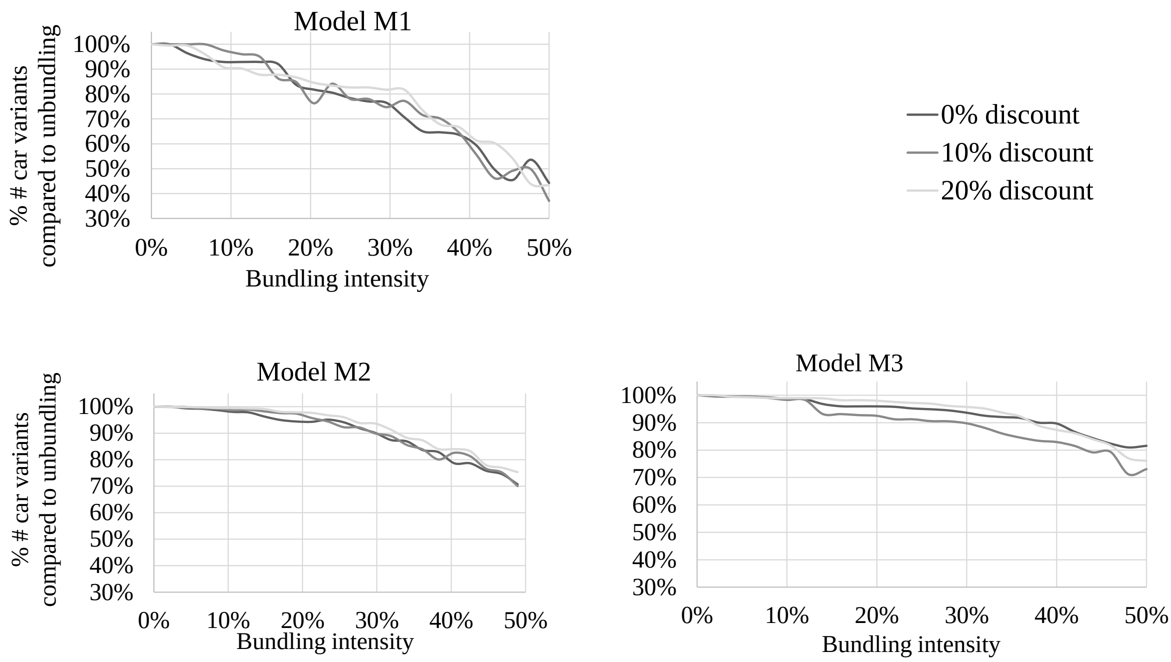

Figure 3.2.: Change of the number of car variants compared to the unbundling case

Car model	M1		M2		M3	
	# car variants	Revenues	# car variants	Revenues	# car variants	Revenues
Values for bundle design with maximum revenues	92.09%	104.60%	97.98%	104.11%	98.77%	101.34%
Values for first bundle design with revenues lower than unbundling	44.56%	99.86%	75.84%	99.14%	98.56%	98.84%

Table 3.6.: Number of car variants and revenues for the maximum revenue bundle design and for the first bundle design with revenues below those in unbundling

For low bundling intensities, the approach generated option bundles containing two options. Here, the large willingness to spend for one option was used to increase the sales of a second option. However, for a bundling intensity higher than 12%, the addition of options to bundles only decreased revenue. The sales of the remaining unbundled options were reduced when they were included in a bundle.

The revenues were increased compared to unbundling for a higher bundling intensity for models M1 and M2 than for model M3. One of the reasons for this result is the variance of the option prices for M3 together with the more heterogeneous customer base. It was more difficult to ensure for model M3 that the willingness to spend would cover the price of the bundles.

When the bundling intensity was increased, our approach sometimes attempted to avoid a revenues reduction at the cost of an increase in the number of variants by a

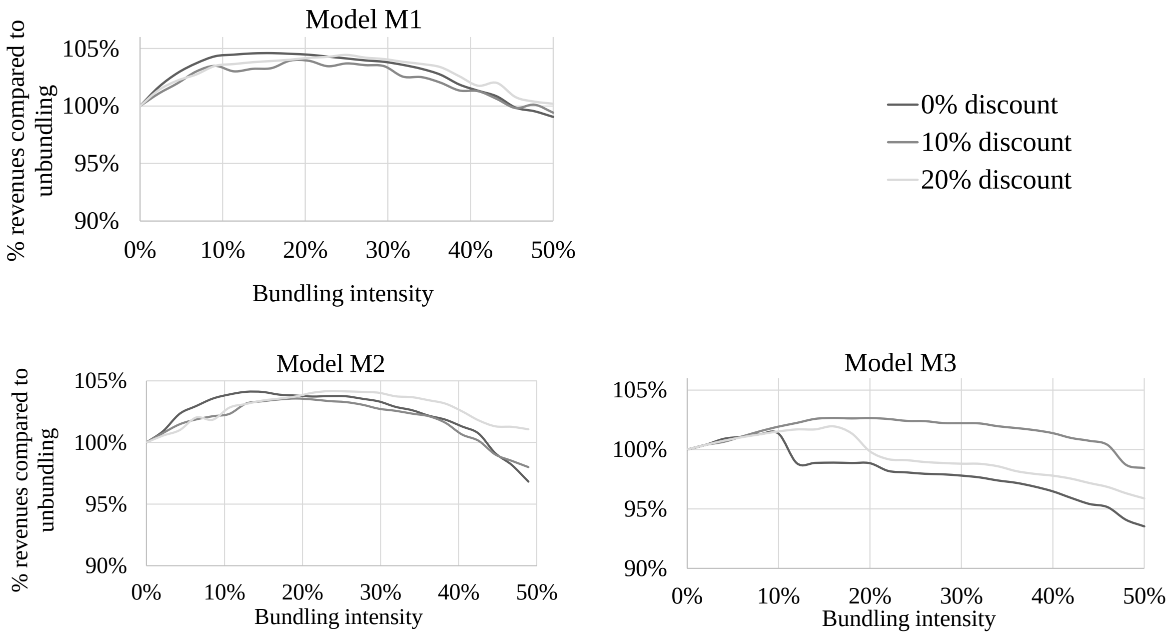

Figure 3.3.: Revenues change compared to the unbundling case

small amount. This is an effect of the revenue maximization objective pursued in our approach.

Our results show that, by simply introducing option bundles, it is possible for the current case study and for our customer model to increase revenues while simultaneously reducing the number of car variants being manufactured. These effects can already be observed for no bundle discount.

Effect of bundle discount

In practice, option bundles are often introduced at a discount to compensate customers for the loss in customization flexibility. We therefore analyzed the impact of various bundle discount levels (10%, 20%) on the number of car variants and revenues.

For models M1 and M2, discounts had only a small impact on the number of car variants and revenues. The sales numbers were higher than in the unbundling case. However, the revenue remained stable due to the discount of the option price.

For model M3, the discounts had a more significant impact. A discount rate of 10% increased the bundling intensity, up to which the revenues were higher than in the unbundling case from 12% to 46%. Here, the number of car variants was reduced by 29% compared to unbundling and by 27% compared to the no-discount case. The revenues were increased by up to 2.66% compared to the unbundling case and 1% compared to the no-discount case. The revenue maximal bundle design reduced the number of car variants by almost 7% compared to the unbundling case and 5% compared to the no-discount case. When a discount of 20% was offered, attractive, expensive options were sold unbundled to avoid the bundle discount and maintain high revenues. However, this also led to reduced sale of the bundles and a corresponding loss of the variant reduction potential.

Our results for model M3 indicate that manufacturers benefit from offering bundle discounts for a customer base with heterogeneous preferences. The choice of a moderate discount level avoids revenue neutral free-rider effects on the one hand and a loss in product variety reduction on the other.

Effect of willingness to spend

After evaluating the impact of the discounts, we then measured the impact of the willingness to spend parameter on the number of car variants and revenues for each car model. We used the discount rates than generated the highest revenues: no discount for models M1 and M2, 10% for model M3. For each car model, we generated additional bundle designs for a 105% and a 115% willingness to spend values. Figure 3.4 presents the Pareto curves of the relative number of car variants and revenues compared to unbundling for the three ρ^* levels tested: 115%, 110% and 105%. Each curve represents the revenues and number of car variants combinations for the bundle designs that have a bundling intensity higher than the revenue-maximal bundle design. The bundle designs with a bundling intensity lower than the revenue-maximal design generate a larger number of car variants compared to the designs with a higher bundling intensity than the revenue-maximal design, while the revenues are similar. Therefore the designs with an intensity lower than the revenue-maximal design are dominated and are not included in the Pareto curves.

Similar to previous results, all bundle designs reduced the number of car variants compared to unbundling. For a willingness to spend of 115%, the bundle designs of our approach always increased the revenues compared to unbundling. Even for a very conservative willingness to spend of 105%, the revenues were higher than those for unbundling for a bundling intensity of up to 34% for model M1, 27% for model M2 and 14% for model M3.

The analysis of the Pareto curves in Figure 3.4 highlighted that the number of car variants reduction had the same pattern for all car models, irrespective of the willingness to spend level. This was mainly a result of the similarity of the structure of the bundles

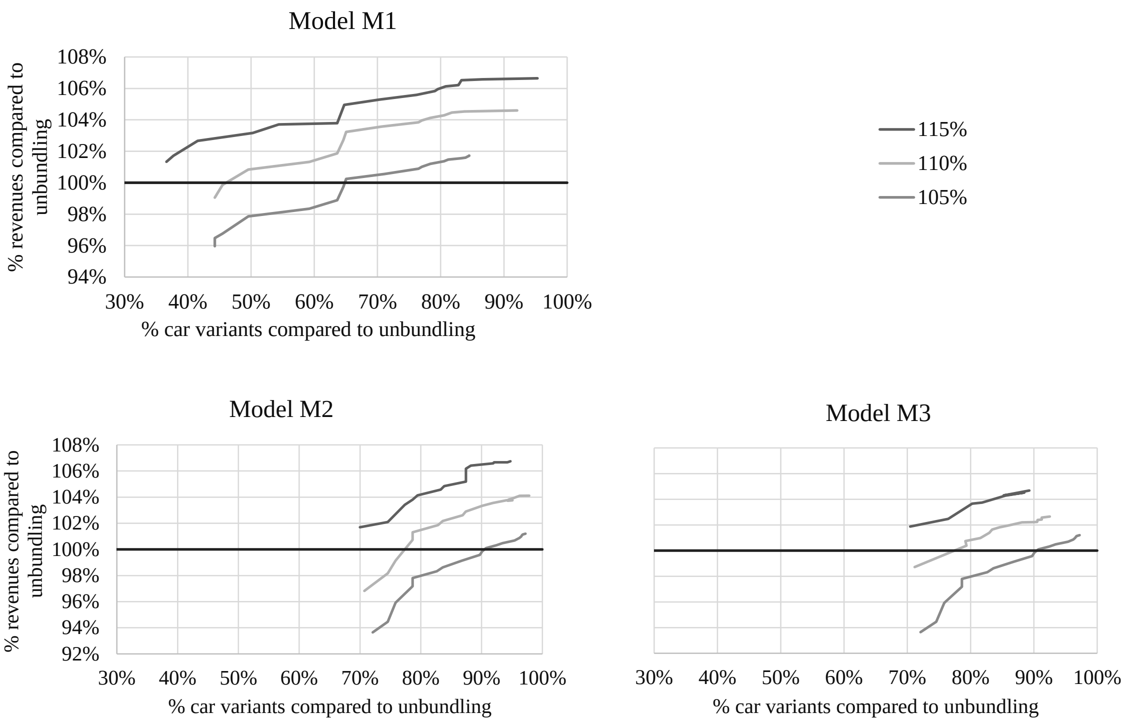

Figure 3.4.: Pareto curve of the number of car variants and revenues compared to unbundling for different willingness to spend levels

generated by the approach. Irrespective of the willingness to spend level, the approach first generated small bundles containing two options. As the bundling intensity was increased, the approach created larger bundles that resulted in a larger reduction of the number of car variants. For a willingness to spend of 115%, the reduction of the number of car variants was larger for model M1. The approach exploited the higher willingness to spend by creating larger bundles, which further reduced the number of car variants.

The willingness to spend analysis shows that our approach can efficiently exploit any level of willingness to spend. Even for a conservative willingness to spend value of 105%, our approach finds bundle designs that reduce the number of car variants and simultaneously increase the revenues.

3.4.4. Robustness analyses

In order to measure the impact of inaccurate estimations of the approach inputs, we ran two robustness analyses. We measured the impact of an inaccurate estimation of the willingness to spend parameter (section 3.4.4) and of a different customer behavior when bundles are offered (section 3.4.4).

Car model	M1		M2		M3	
	# car variants	Revenues	# car variants	Revenues	# car variants	Revenues
∅ difference between the bundles generated for $\rho = 1.1$ and the optimized bundles	-0.30% ($p = 0.76$)	-1.79% ($p = 6.69 \cdot 10^{-11}$)	-0.31% ($p = 0.15$)	-2.56% ($p = 3.3 \cdot 10^{-10}$)	-2.48% ($p = 5.20 \cdot 10^{-5}$)	-2.03% ($p = 3.49 \cdot 10^{-10}$)
∅ difference between the actual performance of the bundles generated for $\rho = 1.1$ and the expected performance	0.00% ($p = 1$)	0.00% ($p = 1$)	0.00% ($p = 1$)	0.00% ($p = 0.32$)	-0.04% ($p = 0.02$)	0.01% ($p = 0.04$)

Table 3.7.: Performance of the bundles generated for $\rho^* = 1.1$ at an actual $\rho^* = 1.15$

Car model	M1		M2		M3	
	# car variants	Revenues	# car variants	Revenues	# car variants	Revenues
∅ difference between the bundles generated for $\rho = 1.1$ and the optimized bundles	-15.23% ($p = 1.93 \cdot 10^{-9}$)	-52.09% ($p < 2.2 \cdot 10^{-16}$)	-10.98% ($p = 6.26 \cdot 10^{-12}$)	-46.42% ($p = 2.43 \cdot 10^{-16}$)	-18.01% ($p = 3.41 \cdot 10^{-10}$)	-32.92% ($p = 1.29 \cdot 10^{-14}$)
∅ difference between the actual performance of the bundles generated for $\rho = 1.1$ and the expected performance	-17.51% ($p = 8.99 \cdot 10^{-12}$)	-54.86% ($p < 2.2 \cdot 10^{-16}$)	-10.52% ($p = 1.73 \cdot 10^{-12}$)	-49.06% ($p = 2.63 \cdot 10^{-16}$)	-12.49% ($p = 2.28 \cdot 10^{-10}$)	-35.37% ($p = 5.67 \cdot 10^{-14}$)

Table 3.8.: Performance of the bundles generated for $\rho^* = 1.1$ at an actual $\rho^* = 1.05$

Willingness to spend

For the approach we developed, there are currently no methods to estimate the value of the willingness to spend parameter ρ^*. We therefore analyzed the effects of an inaccurate estimation of the parameter on the number of car variants and on the revenues. To this end, we compared the performance of the bundles generated for ρ^* equal to 1.05 and 1.15 to the performance of the bundles generated for the ρ^* value of 1.1 if offered to customers with a willingness to spend value of 1.05 and 1.15 respectively. Tables 3.7 and 3.8 illustrate the performance differences between the optimized bundles and the ones generated for a ρ^* value of 1.1. The tables also present the differences between the expected performance and the actual performance for ρ^* equal to 1.1. We ran a t-test to determine the significance of the gap between these average differences and 0. In both tables, the p value for the t-test is represented.

Our comparison shows that underestimating the willingness to spend of customers has limited negative consequences. The revenues were on average by at most 2.6% higher for the bundles optimized for the willingness to spend. The average revenues difference was statistically significant from 0 for all car models, whereas the difference of the number of car variants for the optimized bundles was not statistically significantly different from 0. The number of car variants increased on average by 2.48% for model M3. For no car model was the expected performance of the bundles very different from the one obtained for the actual willingness to spend levels.

In general, the revenues were increased for the optimized bundles. Additionally, the bundle selections for model M3 became more heterogeneous compared to those for the bundles designed for ρ^* equal to 1.1. The result stems from the larger bundles that were generated for the 5% additional willingness to spend. Many customers were able to buy the larger bundles but there were some customers who bought the cheaper options from the bundle and did not have the willingness to spend to buy the whole bundle.

As the difference between the expected and the actual performance shows, even if the willingness to spend was higher, it was often not possible with the bundles generated for ρ^* equal to 1.1 to motivate customers to buy additional bundles. In the bundles designed for ρ^* equal to 1.1, an option was usually grouped together with a comparatively cheaper one. The higher willingness to spend for ρ^* equal to 1.15 was still not sufficient to entice the customers who were interested in the cheaper option to buy the bundle.

In contrast, the overestimation of the willingness to spend can have very severe consequences on the performance of the bundles. Compared to the optimal bundles, the number of car variants was statistically significantly reduced on average by at least 10%. However, the reduction resulted from the abandonment of option purchases. On average, the revenues were statistically significantly reduced by at least 32%. Also the expected performance of the bundles was not met. Since our approach efficiently exploited the willingness to spend of the customers, a willingness to spend value of 1.05 was not high enough to finance the purchase of the bundles designed for a ρ^* value of 1.1.

The analysis highlights that it is very important for our approach not to overestimate the willingness to spend of customers. An underestimation does not reduce the performance of the bundles by a lot compared to bundles tailored for the actual willingness to spend value. Furthermore, the actual performance of the bundles was similar to the one expected for the willingness to spend value planned for.

Customer behavior modelling

Our approach generates bundles based on readily available data, namely past customer purchases. In many of the existing option bundling approaches, reservation-price-based models are used to represent customer behavior. However, these models require input data that is difficult and expensive to collect. Our goal was to evaluate whether the bundles still have a good performance if the customers actually behaved according to a customer model different from ours. We ran multiple simulations that involved a reservation-price-based customer behavior model. We then analyzed the product variety and revenues generated by the bundles designed by our approach.

In reservation-price-based customer behavior models, it is assumed that the customers have a known, predefined reservation price for each individual option and that they buy bundles in order to maximize the consumer surplus. When an option is allocated to only one bundle, then a customer buys a bundle if the sum of her reservation prices for the included options is at least as large as the bundle price.

In our approach, the decision of a customer to purchase a bundle depends on which options purchased in the past are included in the bundle, as well as the prices of these options. Our model differs from a reservation-price-based approach in the use of a single value for modeling the customer surplus, which is independent of the option and customer. Also, the parameter ρ^* does not encapsulate any information related to the reservation prices of the options that the customers did not purchase in the past.

We ensured that the settings are similar for the two customer models in our comparison. We first ran simulations in which we replicated a reservation price estimation in which customers cannot be grouped in customer segments that is based on a conjoint analysis or a discrete choice model. In such a setting, one reservation price that is representative for all the customers is derived for each option. We additionally assumed that the ρ^* parameter value used for our approach was an accurate average of the customer valuation of the options acquired in the past. We therefore drew for each option a ρ_o value from a normal distribution with the mean ρ^* and a standard deviation σ and scaled the ρ_o values such that their average was equal to ρ^*. The reservation price of a customer for an option purchased in the past was then ρ_o times higher than the price of the option.

We also ran simulations for a more realistic setting, in which each customer has a specific reservation price. We drew for each customer a ρ_c value from a normal distribution with the mean ρ^* and a standard deviation σ. We scaled the ρ_c values such that their average was equal to ρ^*. We then drew a value $\rho_{c,o}$ for each option purchased in the past by customer c from a normal distribution with the mean ρ_c and the standard deviation σ. We scaled the $\rho_{c,o}$ values for each customer such that their average was equal to ρ_c. The reservation price of a customer for an option purchased in the past was then $\rho_{c,o}$ times higher than the price of the option.

Since the reservation price of a purchased option cannot be lower than the price of the option, we set the ρ_c, $\rho_{c,o}$ and ρ_o values equal to 1 for draws lower than 1. We measured the impact of the σ parameter with two values: 1% and 2%. The 1% standard deviation represents a setting in which the ρ^* estimate is very accurate. For the 1% standard deviation, almost all reservation prices for the purchased options were between 107% and 113% of the price of these options. The 2% standard deviation represents a high uncertainty case for the actual values of the reservation prices. For the 2% standard deviation, the reservation prices of most options purchased in the past were between 104% and 116% of the price of the options. For the options the customers did not

Figure 3.5.: Absolute difference spectrum for the number of car variants and revenues between the reservation-price-based model with customer-independent reservation prices and our customer behavior model

purchase in the past, we assumed that customers would be willing to pay a conservative value of 20% of the price of the options.

For each σ value and bundles configuration generated for ρ_c equal to 1.1, we ran 100 simulations. We did not offer a discount for the bundles for models M1 and M2 and used a discount rate of 10% for model M3.

Figure 3.5 illustrates the number of car variants and revenues differences spectrum between the reservation-price-based model with customer-independent reservation prices and our model for the 1% and 2% standard deviation. The number of car variants and revenues were similar to the ones obtained with our model for the 1% standard deviation. For a standard deviation value of 2%, the revenues were reduced more than for the 1% standard deviation. Since the reservation prices were identical for all customers, the number of different bundle selections made by the customers was reduced and the decisions of customers became more homogeneous. The number of car variants was therefore lower than the one predicted using our model.

Figure 3.6 illustrates the spectrum of the differences between the average number of car variants and revenues for the reservation-price-based model with customer-specific reservation prices and our model for the 1% and the 2% standard deviation values. The number of car variants increased by up to 1% compared to the number predicted by our customer model. However, the number of car variants was always smaller compared to the unbundling case. The revenues for the reservation-price-based model were higher than for unbundling for a bundling intensity of 41% for model M1, 43% for model M2 and 46% for model M3. Even if the standard deviation was doubled, the performance of the bundles was still very good. For most bundle designs, the number of car variants was still lower compared to unbundling. The revenues were on average lower for the reservation-price-based model compared to ours. However, they were still higher

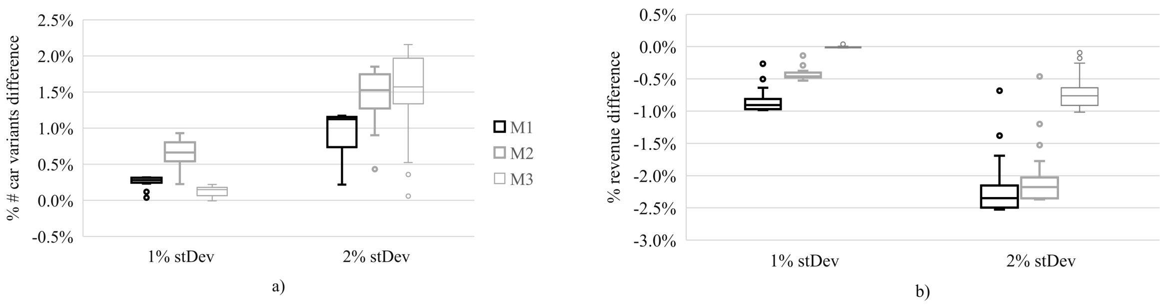

Figure 3.6.: Absolute difference spectrum for the number of car variants and revenues between the reservation-price-based model with customer-specific reservation prices and our customer behavior model

compared to unbundling for a bundling intensity of up to 34%. The analysis for both standard deviation values shows that even for a different customer model, the same trend was maintained: option bundling led to both an increase in revenues and a reduction in the number of car variants.

The revenue drop was higher for models M1 and M2 compared to model M3 mainly due to the price variability of the options. For models M1 and M2, our approach bundled the popular expensive options with cheaper options. However, the price of the cheaper options was still high enough that, for a low reservation price draw, the surplus for the options did not cover the price of the cheaper options. The homogeneity of the customer base led only to a small increase in the number of car variants. Since model M3 had similar prices for the options, the number of customers who in the simulations selected more bundles and the number of customers who purchased less bundles than predicted by our model was similar. The revenues in the simulation were therefore similar to the predicted revenues. Due to the many changes in the bundles selections and the high customer base heterogeneity, the number of car variants increased compared to our model.

Even though we proposed a completely new customer model for bundling purposes, our experiments showed for our case study that the effects on the number of car variants and revenues are similar to those obtained for a standard customer model. Even when there is a higher uncertainty regarding the value of the reservation prices, it is possible, with the bundles generated by our approach, to reduce the number of car variants while at the same time to increase revenues. For a model in which the reservation prices of the customers are identical, the results are even more indistinguishable from the ones estimated by our approach.

3.5. Conclusions

This paper systematically quantifies the effects of option bundling on product variety. For this purpose, we have developed a data-driven approach for option bundling. In contrast to conventional option bundling approaches, it does not rely on excessive data requirements. Our approach primarily requires data that is readily available: the prices of the options and past customer purchases. Due to the reduced input requirements, the effort connected to gathering and preparing the input data for our approach is limited. Our branch-and-price approach is highly relevant for practitioners, since it can deal with the preferences of thousands of customers and can generate a predefined number of bundles for a large number of options. Our approach exploits the capabilities of modern computational hardware by distributing computational tasks over the available CPU cores. For large problem instances, we designed a matheuristic that simplifies the column generation procedure of the branch-and-price approach. The matheuristic generated optimal or close-to-optimal results.

We tested our approach for three car models of a large German automotive manufacturer. We split the data on purchases of dozens of options made by thousands of customers into a training set, which we used for bundle generation, and a validation set, which we used for performance analysis. We demonstrated how the application of our methodology resulted in a definition of the relevant trade space, specifying the relationship between the number of car variants and revenues for varying numbers of bundles. For the case study, it was possible to design bundles that reduced the number of car variants and also increased revenues. For a bundle design with a neutral effect on revenues, the reduction of the number of car variants was substantial. These results were also observed for different willingness to spend and bundle discount levels, albeit at different magnitudes. No discount is required to obtain the benefits of option bundling if the customer base is homogeneous and the prices of the options have a wide range. For other cases, bundle discounts need to be carefully selected.

Even when customers made their choices not according to our behavioral model but according to the standard behavioral model that assumes a known reservation price for each option and each customer, no significant differences in the performance of the bundles were observed.

In our approach, the reduction of the number of variants is only an indirect effect. Future research could further explore the variety reduction potential by developing a multi-objective bundle design approach that in addition to revenue maximization also explicitly pursues the minimization of the number of product variants. If future research also quantified the costs related to product variety, then a new option bundling approach could even strive for a maximization of profits.

The focus of our research was to identify the impact of option bundling on product variety. For this purpose, we developed an option bundling methodology with reduced input requirements, which include the willingness to spend. We demonstrate how an analysis can be carried out for given willingness to spend values. We also investigated the impact of this parameter on the results, but its exact determination is not subject of this paper. The available datasets did not contain sufficient information to estimate customer-specific willingness to spend values. For this purpose, a consumer study would have to be performed. Unsupervised learning approaches, such as deep learning networks, could be used to enhance the estimation of the willingness to spend parameters. However, consumer studies are no longer possible for the car models in our case.

Our results are based on a case study that includes three car models. Even if the car models are designed for different customer segments and the range of option prices is model-specific, we cannot generalize our findings. Future research should incorporate tests on a larger number of car models and on products from other industries, to confirm our findings.

With our approach, practitioners have an option bundling approach at their disposal that can be used for real-life applications. The design of option bundles is no longer at the whim of a small number of experts, but results from a data-driven, well-defined process. Our methodology enables manufacturers to exploit valuable knowledge on customer preferences hidden within the trove of information lying inside their data warehouses.

Our work highlights that option bundling does live up to its promise to reduce product variety. However, for the first time, researchers now have an example of an evaluation of the actual reduction potential based on limited data. No longer will experts design bundles by blindly following their intuition regarding the reduction of the number of product variants. They now have a way to efficiently determine the impact of their designs in a small amount of time. In such a setting, the design of bundles can finally no longer be dominated by the marketing perspective alone. The design decision can now be done by including an operations-related perspective, the reduction of the number of variants.

Chapter 4.

Enhancing MRP-based component demand planning in a high-variety context

This chapter is based on an article submitted as:

Stäblein, T. et al. (2016). "Enhancing MRP-based component demand planning in a high-variety context"

4.1. Introduction

The ability to accurately plan component demand is a core capability for manufacturing firms, as the quality of future sales predictions determines not only customer service, but also inventory and supply chain performance. MRP and MRP-based scheduling modules embedded within Enterprise Resource Planning (ERP) systems have become the industry standard for production planning in discrete manufacturing contexts (Jacobs and Weston, 2007). Surveys of industrial practice consistently point to MRP usage rates of more than 74% across large manufacturing firms (Olhager and Selldin, 2003, Jonsson and Mattsson, 2006).

Despite their prominence in practice, several drawbacks of MRP systems have been identified. Of particular concern is their inherent rigidity, and in turn, inability to cope with dynamic changes of the environment in which they operate (Winters et al., 2008, Goodhue et al., 2009, Fauscette, 2013, Tenhiälä and Helkiö, 2015). As a result, a considerable risk of workarounds has emerged, potentially compromising the main advantage of ERP systems in terms of providing a reliable inter-organizational information processing capability (Gattiker and Goodhue, 2005, Tenhiälä and Helkiö, 2015). As a result, a great variability in outcomes of ERP implementations has been observed (Gattiker and Goodhue, 2005, Stratman, 2007, Hendricks et al., 2007).

In this paper we address a particular problem with MRP-based scheduling caused by the drastic increases in product variety and the globalization of manufacturing and sourcing footprints that many durable goods manufacturers are facing. Combined, these developments provide a real challenge for MRP-based scheduling systems. Due to an exponential number of possible product configurations that share underlying components and the unavailability of advance information on customer orders, the translation of planned orders into detailed component plans in MRP becomes a problem in practice. The (measurable) result has been a reduction in schedule accuracy in the supply chain, as manufacturers are no longer able to simply 'explode' the bill of materials (BOM) to derive the respective demand for the required components.

We argue that this problem can be addressed by harnessing all available planning information, past and present. To this effect we propose an enhancement to MRP systems by using the analogy of 'data fusion' (Goodman et al., 1997, Hall and MacMullen, 2004) that allows us to triangulate past order data with forward planning data for supply chain planning purposes. We demonstrate the potential of the approach by comparing it to a state-of-the-art time-series forecasting software package and to the existing MRP approach used by Mercedes. We report on the ten-year process of the development, empirical validation, and implementation of an improved algorithm at Mercedes-Benz Cars, where it has demonstrated a 15% improvement in schedule accuracy over the existing system.

The paper is organized as follows: in Section 4.2 we provide more detail on the focal problem of operating existing MRP systems in a high-variety context, before reviewing the relevant literature related to existing and emerging approaches to production planning in durable goods manufacturing in Section 4.3. In Section 4.4 we introduce the model for enhancing MRPs performance in the above context. In section 4.5 we validate our model and benchmark its performance against existing solutions. Section 4.7 presents the conclusions.

4.2. The problem

This research project was conducted in collaboration with Mercedes-Benz Cars (commonly known as 'Mercedes'), the passenger car division of Daimler AG. Mercedes vehicles are assembled in seven major assembly plants spread across four continents. Each of these plants is closely tied to a dedicated network of local and global suppliers that deliver on average 4,200 different individual components that are used to build a highly customized Mercedes product. Such a level of supply chain complexity is not unusual in the automotive industry. Premium manufacturers typically present their customers a wide array of options to customize their vehicles in terms of choices of engines, colors, interior materials, electric features or other functionalities such as sun-roofs or cup holders. These options increase external product variety exponentially due to their possible combinations (Fisher and Ittner, 1999, Pil and Holweg, 2004). Mercedes builds a very large fraction of its cars to order and regards offering a maximum level of customer choice as a primary competitive advantage. Hence, the company has elected to offer customers a choice of up to 120 options in the configuration process for its passenger cars.

4.2.1. Rise in supply chain complexity

The underlying planning problem is exacerbated by two factors: firstly, as Figure 4.1 illustrates, supply chain complexity has increased steadily over the last two decades. In 1980, Mercedes was producing just below 500,000 passenger cars in its two German plants, across four models. By 2015, the production had risen to over 2 million cars produced in seven assembly plants globally, across a total of 28 different models. The IBM MRP system in operation today was implemented in 1985, and its functionality has not evolved or been updated in any meaningful way during that time.

Secondly, product variety has been rising as durable goods manufacturers have chosen to offer a large set of choices, or options, to its customers as part of their marketing strategy. As Randall and Ulrich, 2001, p. 1603, state '[...] the ultimate success of high-variety strategies may rest not only on a supply chain's ability to physically deliver variety, but also on the ability to communicate and present options to consumers.' The higher the number of theoretical configurations of an integral product becomes, the more complex the component-level demand planning problem becomes.

The requirement for most components depends on the combination of options, and customer preferences play a central role for component-level demand scheduling. However,

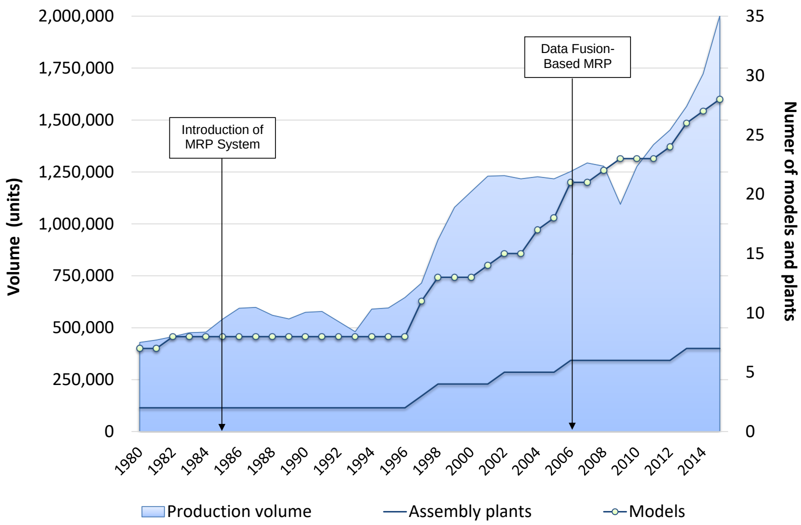

Figure 4.1.: Production volume, model line-up and global assembly plants at Mercedes-Benz cars, 1980-2015

Lifecycle	**Model name (code)**	**Body styles**	**Powertrains**	**Customization options**	**Theoretically possible product combinations**
1984-1995	200-series (W124)	2	6	16	786,432
1995-2003	E-class (W210)	2	9	41	39,582,418,599,936
2003-2009	E-class (W211)	3	15	70	53,126,622,932,283,500,000,000
2009-2016	E-Class (W212)	4	16	80	77,371,252,455,336,300,000,000,000

Table 4.1.: Evolution of product variety across Mercedes-Benz E-class product generations

due to the option induced product complexity and the high combinatorial possibilities between options, it is clearly not possible for the sales department to forecast take-rates for all possible option combinations, or even smaller subsets of option combinations. To give an example, we found that the 'CLS-Class' model can be customized with up to 80 options, which theoretically allow for $7.74 \cdot 10^{25}$ different configurations.

It is also important to note that the product variety levels offered on present models exceed the variety for earlier models by orders of magnitude. Table 4.1 shows the evolution of product variety of the E-Class platform and the constant increase of product variety at Mercedes-Benz.

To further elucidate the interlinkage between overall product variety, and the variety of components required to provide the former, we empirically analyzed the options interdependencies for all 8,592 components listed for the Mercedes CLS-Class sedan. We

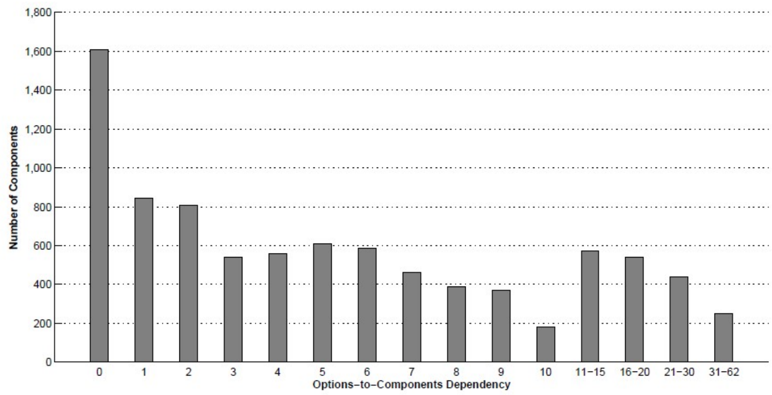

Figure 4.2.: Statistical analysis of the options-to-components dependency used to determine component-level demand for the Mercedes CLS-Class model

used an automated query of the engineering product database (which contains the bills of materials for all possible build combinations of any given model) to extract the number of options interdependencies for all components that could be used in the assembly of the CLS model. Figure 4.2 graphically illustrates the number of interdependencies of options for the components of the Mercedes CLS-Class.

While we did expect to find strong inter-dependencies, both the authors and the production scheduling team at Mercedes-Benz were surprised to learn that 81.3% (6,985) of the components incorporate at least one option in their definition, and a maximum of no fewer than 62 options. Only 18.7% (1,606) of the components were independent of any customer-driven customization (in other words they were installed in every car).

4.2.2. Production planning at Mercedes-Benz cars

At present Mercedes uses a rolling horizon cycle in their MRP-based component demand planning and communicate with their suppliers on a continuous basis. This approach does not deviate from any other 'traditional' MRP-based scheduling system commonly used by large durable goods manufacturers. In Germany, for example, a typical automotive first-tier supplier receives eight weeks' worth of demand in the form of daily requirements (updated daily), a further three months' on a weekly basis (updated weekly), and a further five months on a monthly basis (updated monthly) for capacity planning and purchase purposes. This format is a mutually agreed industry standard for delivery forecasts and secondary demand information in the automotive industry. It is a contractually binding framework that all vehicle manufacturers and component suppliers have agreed to adhere to in order to standardize and improve communication along the supply chain. The German format (VDA 4905) is compatible with the European 'DELNIS' format of Organisation for Data Exchange by Tele Transmission in Europe (ODETTE) and the 'DELFOR' format in the USA by the Automotive Industry Action Group (AIAG). Currently there are on-going activities to elaborate a common format for global use.

The 8-week and 3-month schedules are provided by the short-term MRP functions, while the five month preview is generated by the midterm planning function since not all customer orders are received yet. A central objective here is the minimization of deviations between these rolling updates and the actual call-offs, to enable suppliers to optimize and plan their operations. It is this deviation that our algorithm seeks to minimize in a medium-term horizon.

The available information in the medium-term horizon is given by master production planning, which determines the number of vehicles to produce per month in each plant in the relevant time horizon. Furthermore, the sales allocation process determines a goal estimate for dealers regarding the option customization of vehicles. The goal-setting of the take-rates for the options in sales allocation planning is also a tool for dealer and sales force organization. The production objective is to match these goals as closely as possible.

The planned take-rates can deviate from the actual demand when the customer orders are received. If a deviation is identified in the planning process, the impacts are discussed. An alignment meeting serves to solve possible conflicts between the production and sales unit and to increase the flexibility and stability of the planning process. Planned take-rates are rarely used in available state-of-the-art MRP planning systems, yet because most components depend on options and customization, we argue that it could be highly beneficial to include then in the MRP planning process.

4.3. Component demand planning approaches: A review

Conceptually, the need to convert projected sales into demand plans for component suppliers is a fundamental, recurring task in any discrete manufacturing environment (Vollmann, 2005, Stadtler et al., 2015). The calculation and sharing of accurate demand information with the supply chain partners not only supports the operational supply chain efficiency, but also reduces managerial risk for all parties involved. Prior studies highlight a potential reduction of lead-time (Cachon and Fisher, 2000), a decrease of the infamous 'bullwhip effect' (Lee et al., 2004), and improved contract decisions (Cachon and Lariviere, 2001). A manufacturer's ability to provide reliable advanced demand information for secondary demands of components and parts is a critical factor in operating an efficient supply chain (Childerhouse et al., 2008). Not surprisingly, the nature of uncertainty adversely affecting the quality of sales forecasts at finished product level and the ensuing 'nervousness' in schedules have been the subject of many studies. In this section we review how the underlying problem is being addressed within an MRP context, and outline more recent time-series and Bayesian network approaches that have been applied in durable goods manufacturing contexts.

4.3.1. MRP-based approaches

The most widely applied tools in the manufacturing industry are MRP-based algorithms, many of which are components of ERP and Advanced Planning and Scheduling (APS) systems, such as SAP's Advanced Planner and Optimizer (APO) (Stadtler et al., 2015). In a nutshell, MRP systems use a set of defined orders (e.g. actual customer orders when available, or forecast orders for finished products, also called 'anticipated', 'dummy', or 'planned' orders) and the BOM to derive the corresponding component demand. The core logic of MRP systems has been widely discussed and will not be repeated here (see Orlicky, 1975, Vollmann, 2005, Jacobs and Weston, 2007, Stadtler et al., 2015 for a complete discussion of MRP and its functionality).

As early as the mid-1980s, some of the drawbacks of MRP procedures were being discussed in the literature. Yeung et al., 1998 recognize MRP's limitations in terms of defining safety stock levels for the end items, ignoring of capacity constraints, considering very simple product structures only, and the inability of MRP to address input uncertainties. When faced with uncertainty, the rigidness of MRP systems often leads to the so-called 'MRP nervousness' in initial MRP implementations (Benton, 2007). Ho and Ireland, 1998 determine by means of simulation that forecasting errors can have a significant impact on the stability of MRP systems. Huq and Huq, 1994 note that MRP systems tend to result in an increase in inventory levels for work-in-progress items. Rom et al., 2002 recognize that the fixed lead times prevalent in MRP systems are not appropriate for some real-life situations . Jonsson and Mattsson, 2006 note that a significant percentage of MRP users make parameter selection decisions, such as lot sizes, based on experience. The literature review by Louly et al., 2008 shows that there were few studies addressing lead time uncertainties or a combination of demand and lead times uncertainties. The mentioned underlying problems fit into the wider context of the ERP systems that MRP tends to be embedded in, and more specifically, their inability to respond to changes in the context in which they operate (Tenhiälä and Helkiö, 2015).

Based on the recognition of these shortcomings, a range of studies have attempted to adapt MRP procedures to alleviate these issues. A range of different strategies have been proposed to manage and/or reduce uncertainty at component-demand level within MRP, such as rolling horizon updating of sales forecasts (Chand et al., 2002), freezing the production schedule (Blackburn et al., 1986), and responding to early sales (Fisher and Raman, 1996). Others have developed heuristics to address specific issues of MRP systems. Armentano et al., 2001 develop a heuristic to determine the lot sizes for the MRP system in a multi-stage production system so that the resulting setup, production and inventory costs are minimized while taking the production capacities into consideration. Their approach is capable of providing feasible solutions for 83.7% of the instances considered in their computation study and leads to a reduction of the resource utilization. Tang and Grubbström, 2002 present a dynamic approach for planning and re-planning the material requirements in an environment with demand uncertainty while minimizing the uncertainty and the schedule change costs. Their approach optimizes safety stock levels and improves the robustness of the MRP method with regards to forecasting errors. Xie et al., 2003 determine by means of simulations the effect of the freezing interval parameters in a constrained capacity scenario for a multi item single-level MRP system. They find that the proportion of periods which are frozen in a planning horizon leads to a trade-off between the total costs of the system, the service level and the schedule instability. Ram et al., 2006 propose a linear program that adapts MRP for flexible BOMs. The flexibility is enabled by allowing the quantities of components to take on values in ranges, instead of being fixed. The model minimizes the quantity deviations from the standard bill of material while ensuring that the planned end items demand is fulfilled.

Louly et al., 2008 develop a mathematical model and a branch-and-price algorithm to compute the order lead times that are incorporated in the MRP system when component procurement times are random. By optimizing the planned lead times, the authors ensure that the costs caused by following up the recommendations of the MRP system are reduced in an uncertain environment. Ioannou and Dimitriou, 2012 present an approach that estimates lead times more accurately by taking into consideration the existing orders and their impact on the lead times. Both approaches lead to the same performance as a computer simulation of the considered job shop. Riezebos and Zhu, 2015 develop a dynamic programming approach for MRP for a scenario in which the lead times for component replenishments depend on the point in time the order is made. Their algorithm determines the optimal order points which lead to minimal purchase, setup, holding and backordering costs. They tackle realistic instances by using three heuristics. Their work highlights the importance of considering order crossovers (i.e. the arrival of an order before another order which was requested before the first order). Their algorithm leads to a cost reduction of up to 25% compared to the heuristic not considering order crossovers.

A number of research papers provide guidelines for the integration of MRP systems with other complementary supply chain management modules. Ferrer and Whybark, 2001 illustrate the extension of a MRP system such that demand and supply decisions are integrated in a remanufacturing company. Kreipl and Pinedo, 2004 describe how to combine a MRP system with a scheduling algorithm and how the resulting process is implemented in an APS module at Tuborg. They discuss the difficulties related to the estimation of the cost parameter values for the APS module. Garcia-Sabater et al., 2009 describe a linear program (LP) for capacitated MRP adapted for an automotive manufacturer. The authors discuss implementation difficulties and the input inaccuracies prevalent in real-life scenarios. In their recent survey of automobile producers, Staeblein and Aoki, 2015 find that the classical MRP approach is still the dominant planning system in the automotive industry.

In summary, the proposed MRP extensions address many specific shortcomings of MRP systems. However, none of the approaches found in the literature is capable of addressing the particularities of an MRP system used for determining component demands in a durable goods manufacturing context, such as the automotive industry. The key difference or challenge compared to other manufacturing contexts is the utilization of a complex product structure whereby the demand for a certain component depends not only on the demand of a single end item alone, but also by combinations of end items. None of the methods found in the literature can integrate such settings.

4.3.2. Alternative approaches

Time-series approaches have long been used to forecast component demand. The objective of time-series methods is to discover patterns in the historical consumption data of each component and then extrapolate this pattern into the future. The ensuing forecast is solely based on past values of the estimated component (for a recent review of time-series techniques see Gooijer and Hyndman, 2006). Researchers have extended time-series forecasting such that it is possible to rigorously select the best forecasting approach and calibrate its parameters automatically (see for example Hyndman and Khandakar, 2008). However, these approaches are limited by not considering projected sales information. Additionally, various inputs from the master production schedule, such as production volume, model-mix constraints and planned option take-rates are not taken into account.

A more recent class of forecasting models are Bayesian network models (Gebhardt et al., 2008). Pilot implementations have been reported at Volkswagen (Detmer and Gebhardt, 2001) and Sun Microsystems (Yelland, 2010). In theory the component demand could be planned based on component take-rates derived from the independent and conditional option take-rates. However, when dealing with a large number of interrelated variables such as the options of a passenger car, the domain under consideration for conditional probabilities is growing rapidly, so scheduling based on probabilistic models requires great computational effort (Gebhardt et al., 2004). Graphical models like Bayesian networks can be used for decomposition and to represent probability distributions between options, as illustrated by Steinbrecher et al., 2008. Such models can be applied to generate a 'complete model' of the relationships between options and their combinations, which can then be used to schedule certain components, as they depend on option combinations. In practice, however, Bayesian networks have considerable drawbacks: changes to the a-priori network in a planning situation become necessary, as any new knowledge gained most likely will impact parts of the network. This task is challenging, as only incomplete knowledge of interrelated options is available, and new conditions (i.e. new desired take-rates for options) might not fit into the observed network structures. In addition, local changes can lead to network inconsistencies. Intelligent updating and revision algorithms can help identify and localize inconsistencies, but require a considerable amount of computation time and expert knowledge, rendering the approach virtually unfit for practical application in manufacturing.

4.3.3. Problem synthesis

In a context where there are an exponential number of possible configurations that share underlying components available for a vehicle, and where advance information on customer orders is not available, the translation of planned orders into detailed component plans in MRP becomes a problem. The approaches available within existing MRP systems, as well as more recent time-series and Bayesian concepts all have significant drawbacks within the context of the underlying problem. In the following we present the algorithm developed by Stäblein, 2008 that significantly enhances the ability of MRP to calculate component-level demand, by considering additional sources of forward- and backward looking planning information.

4.4. Model

4.4.1. 'Data fusion' as a guiding analogy

At present, the MRP system uses a traditional planned orders approach. Planned orders are derived from the past order history to determine component schedules, omitting relevant forward planning information, such as planned option take-rates, production volumes and advance orders. The main premise of the algorithm developed by Stäblein, 2008 is to merge all available data and forecast information for planning purposes. The author uses the analogy of 'data fusion', an approach that was originally developed for military intelligence purposes to enable a meaningful analysis of multi-sensor data. The main objective of data fusion is to use all available information when analyzing a particular problem (Hall and MacMullen, 2004). The main characteristic of data fusion is the combination of different types of information by means of mathematical methods. In its first applications different types of images were overlaid to provide an overall better understanding of the situation on the ground. Military data fusion problems include multi-target detection, object identification, and object tracking (Goodman et al., 1997). As a result of the 'fusion', complementary aspects of all available information are exploited to the overall best knowledge.

In analogy to the data fusion application, the component demand scheduling process needs to consider information which is generally not compatible in its structure or format. On the one hand, automotive manufacturers use planned sales volumes and option take-rates as inputs for the process to integrate their market knowledge about the upcoming months. Past orders should be included since they incorporate the option combinations prevalent in the market and the structural preferences of the customers. Additionally, product configuration data are used to derive the component requirements from the option combinations. These types of information have different modalities: the planned sales volumes are integer numbers, the option take-rates are defined as percentages of vehicles planned to contain each of the considered options, past orders are represented as a tuple of option combinations and their selection frequencies, and the product and component configuration data are complex Boolean relationships. To enable accurate component demand schedules we posit that it is important to consider all types of information simultaneously and combine these numerically. ‘Data fusion’ provides automotive manufacturers with the conceptual framework to achieve this meaningful combination within the context of production planning. In the following section we discuss the mathematical approach underpinning the model of Stäblein, 2008.

4.4.2. Mathematical representation

The data fusion-inspired algorithm discussed in Stäblein, 2008 consists of two steps: (1) the adjustment of past order frequencies by means of a non-linear mathematical model, which ensures that the master production schedule is adhered to, and (2), the BOM explosion based on the results of the model. By using mathematical optimization the inputs of the planning procedure are incorporated more meaningfully that any of the existing MRP heuristics used in the industry.

Our optimization model ‘learns’ from past customer preferences and assigns an adjusted selection frequency to past orders. The adjustment takes place in the light of the master production schedule information and additional forecasts. In other words, the model regularizes the past order frequencies using prior information and provides a conceptual justification and a mathematical way of combining different types of input data.

Our procedure can be used in a multi-period planning setting. Let $t \in T$ be the future periods (e.g. months) in the planning horizon for which component demands schedules need to be prepared. We assume that for each period t a certain planning and/or forecasting function (such as master production scheduling, sales allocation planning or demand forecasting) provide information regarding (1) planned production volume $\tilde{N}_t$ and (2) for each option o a planned take-rate $V_{o,t}$. Furthermore, we assume that in each planning cycle these types of information are being updated in a rolling-horizon fashion (Holweg and Pil, 2004)

For each of the future periods t requiring the computation of the component demand, the model described in the following is used.

We use some of the notation presented in chapter 3. We include the following additional set in the model:

$c' \in C'$	Set of past orders

The following parameters are incorporated in the model:

$a_{c',o}$	1 if past order c' contained option o, 0 otherwise
K_o	Past take-rate of option o
$\tilde{N}_t$	Planned production volume for period t
$\nu_{c'}$	Frequency of past order c'
ω	Weight for the deviation from the planned take-rates
$V_{o,t}$	Planned take-rate of option o for period t

The following decision variables are used in the model:

$w_{c',t}$	Adjusted frequency of the past order c' for period t

The model is provided with a set of past orders C' from a specified time interval in conjunction with the corresponding frequency of the past orders $\nu_{c'}$. The frequency $\nu_{c'}$ represents the weight of order c' in the pool of past orders. The past orders usually have a frequency of 1, since luxury automotive manufacturers often state that each order is unique. We note that the importance of a customer order based on the age of the information can be adjusted with this term, e.g. by decreasing frequencies with the increasing age of the information.

In addition, the model is provided with a set of options O that are relevant for the component demand schedule together with the take-rate of the options K_o in the past orders. The binary parameters $a_{c',o}$ describe the option composition of the past orders. The future market knowledge is included in the planned production volume $\tilde{N}_t$ and the planned option take-rates $V_{o,t}$ as a result of upstream planning and/or forecasting procedures.

The objective function and the constraints of the model are presented in expressions (4.1) through (4.6).

Minimize

$$\frac{1}{|O|} \cdot \sum_{o \in O} | \sum_{c' \in C'} \frac{a_{c',o} \cdot w_{c',t}}{\tilde{N}_t} - V_{o,t}| + \frac{\omega}{|C'|} \cdot \sum_{c' \in C'} (w_{c',t} - \nu_{c'})^2 \tag{4.1}$$

Subject to:

$$\tag{4.2}$$

$$\sum_{c' \in C'} w_{c',t} = \tilde{N}_t \tag{4.3}$$

$$V_{o,t} \geq \sum_{c' \in C'} \frac{a_{c',o} \cdot w_{c',t}}{\tilde{N}_t} \geq K_o \qquad \forall o \in O : V_{o,t} \geq K_o \tag{4.4}$$

$$K_o \geq \sum_{c' \in C'} \frac{a_{c',o} \cdot w_{c',t}}{\tilde{N}_t} \geq V_{o,t} \qquad \forall o \in O : V_{o,t} < K_o \tag{4.5}$$

$$w_{c',t} \geq 0 \qquad \forall c' \in C' \tag{4.6}$$

The model computes order frequencies $w_{c',t}$ such that the planned option take-rates $V_{o,t}$ are adhered to as well as possible. The past orders contain an inherent structure of high-dimensional option combinatorics of past customer configurations. This structure is propagated in the following months, resulting in minimal overall structural changes. Since it is not possible to forecast all different combinations for options due to their extremely large number, we do not consider the adjusted computed frequencies $w_{c',t}$ as orders that are actually going to be placed by customers. Future orders should however be seen as combinations of option blocks from the past orders, thus including groups of options the customers would desire. In light of this perspective, the computed frequencies are only used to derive the components demand. These frequencies are provided for any planning period t of the relevant planning horizon. If another aggregation level is needed for the schedules, the model can be easily adapted by incorporating the master production schedule on that aggregation level.

An analysis of the past orders has shown that the take-rates for most of the options and high-dimensional option combinations rarely vary significantly from the past ones in a timespan of a few months. Similar studies of online product configurational data analysis at Mercedes show similar results. For those cases in which they do, e.g. options with a seasonal demand or marketing promotions etc., the planners incorporate the expected changes in their planned take-rates. Therefore, the option take-rates that result from the adjusted frequencies are naturally bounded by the past option take-rates and the planned option take-rates (constraints (4.4) and (4.5). A plan in which the resulting take-rates lie outside these bounds would signal a scheduling error or inaccurate inputs.

Constraint (4.3) enforces that the sum of the adjusted frequencies of the orders is equal to the planned sales volume. The constraints (4.6) represent the non-negativity constraints.

The objective function (4.1) incorporates a trade-off between the adherence to the planned option take-rates and a minimal change of the order frequency. The adherence to the planned option take-rates is ensured by minimizing the average absolute difference between the option take-rates that result from the adjusted order frequencies $w_{c',t}$ and the planned option take-rates $V_{o,t}$. The second term of the objective function represents the average squared difference between the past order frequencies and the adjusted order frequencies. Since the term is incorporated into the objective function, we require that the frequencies of the orders are altered as little as possible. The frequency differences are squared to enforce minimal structural changes. The average frequency deviation is then weighted by parameter ω to allow the inclusion of managerial preferences with regard to the trade-off. An alternative is to select the ω parameter value that minimizes an accuracy measure such as the Mean Absolute Percentage Error (MAPE) on a training dataset (e.g. an out-of-sample forecast minimization of the most recent component demand plan). Hence, the scale of the trade-off is set so that the algorithm results in the most accurate component demand schedule under consideration of the available input data.

Due to the convex nature of the objective function, the problem can be solved efficiently by commercial solvers such as IBM ILOG CPLEX.

The component demand is calculated by means of a standard BOM explosion: The adjusted frequencies $w_{c',t}$ are combined with the component structure. Let $z \in Z$ be the set of components to be scheduled and $\iota_{z,c'}$ binary parameters with $\iota_{z,c'}$ equal to 1 if past order c' included component z, 0 otherwise. In line with VDA 4905, the component demand for the mid-term time horizon needs to be specified on the aggregation level month. The demand for a certain component z in month t, $d_{z,t}^{\text{comp}}$ is then equal to:

$$d_{z,t}^{\text{comp}} = \sum_{c' \in C'} w_{c',t} \cdot \iota_{z,c'} \tag{4.7}$$

We experimented with different approaches to model aspects of age-related influences of past order data, such as decreasing weights for past order frequencies with the increasing age of the order information. These trials did not result in any significant improvements in the forecast accuracy. Hence the model presented above uses a naive approach to model the age influence of past orders by setting the frequency of each order in the set of past orders C' to 1. We note that the pool of past order history influences the results by taking into account the basic trade-off between noise-damping and impulse-response. In practice, between five to nine months of order history (depending on model and position in its life cycle) showed the best and most robust results. More details on the size of the past orders used at Mercedes unfortunately cannot be reported due to the commercial sensitivity of this information. Still, future extensions of the model could investigate further approaches to consider age-related aspects of past orders.

4.5. Validation and implementation

The initial outline of the algorithm was proposed to the Production Planning and Control department at Mercedes' headquarters in 2006. The reception was generally positive, largely due to the fact that the shortcomings of the in-house MRP systems were widely known. At the same time a considerable degree of scepticism regarding the capability of the method to consistently provide good results in an acceptable amount of time, while coping with the rolling horizon setup used at Mercedes, prevailed. To address these concerns, the Head of Production Planning for Mercedes-Benz requested a series of tests that would compare the novel algorithm not only against the existing system, but also against other available tools. A call for proposals for a component demand plan accuracy benchmark was sent in 2006. However, only a few software vendors responded.

4.5.1. Experimental design

We validate the Data fusion-based algorithm (DF) by comparing it against an algorithm based on the existing MRP system logic (MRP) and time-series forecasting (time-series). For time-series based forecasting, we used the 'forecast' library implemented in the software R. This library can represent any type of trend and seasonality and automatically selects the best fitting forecasting parameters for the given demand pattern. The details of the procedure are described in Hyndman and Khandakar, 2008.

We measured the performance of the algorithms in terms of schedule accuracy and computation time. In this paper, schedule accuracy is defined as the error of the component demand schedule versus the actual component demand. We compared the errors by using the Mean Absolute Deviation (MAD), the MAPE, and the Thiel Inequality Coefficient (Thiel's U). We used the MAD to measure the average absolute forecasting errors and the MAPE to evaluate the relative deviations of the different components with varying demand magnitudes. The Thiel's U compares the accuracy of the algorithms to the one resulting from naive forecasting. To ensure the validity of our results, we ran an out-of-sample analysis to measure the component demand schedule accuracy following Tashman, 2000. We also measured how much computation time was needed to obtain the solutions. All tests were run on the same PC with an Intel i7 3.4 GHz 4 cores CPU and 16 GB RAM.

The real-world application of the model was tested by Stäblein, 2008 for one test instance. We therefore ran two additional sets of tests that validate the application of the model for a broader range of settings. Section 4.5.2 reports on experiments on eight fictive, more difficult datasets based on real-world data. In Section 4.5.3 the ability to operate the algorithm within a rolling horizon scheduling context is analyzed. For all the tests, we selected the ω value which leads to a minimal MAPE in the training datasets.

4.5.2. Robustness checks

To assess the performance replication possibilities for the data fusion-based method, we evaluated its schedule accuracy robustness under demand nervousness. We were provided with the whole set of past customer orders from January through September 2006 (in total 39,195 orders) for the CLS model, a premium Mercedes-Benz four-door coupe. We included the 197 options used for the configuration of vehicles. This product was selected because it was in full production and it involved a very low number of exceptional events needing to be filtered out (such as changes in product configurations, low stability in production volumes, or part changes from product development). We used the orders from January to June as common inputs for all methods.

For each algorithm, we computed the component demand plan for September and compared it to the actual component demand in that month. We modified the number of vehicles (±30%) and the option take-rates (equal, ±10%, randomly in a 15% interval), combined these settings and generated fictive orders based on these numbers, resulting in eight different datasets. In our analysis, we used 100 representative, harder to predict components. The components were selected based on three criteria: (1) high customer configuration dependency (i.e. components represented by complex options-to-components dependencies), (2) frequently installed (in at least 10% of the

Dataset no.	Method	MAD	MAPE	Thiel's U	Solving time (s)
#1 (+30% orders, same take-rates)	Time-series	670.28	18.49%	0.1889	2.68
	MRP	598.31	17.55%	0.2089	53.35
	DF	109.2	2.83%	0.0277	4.36
#2 (-30% orders, same take-rates)	Time-series	360.49	18.42%	0.1875	2.28
	MRP	319.75	17.39%	0.206	31.97
	DF	48.11	2.91%	0.028	1.99
#3 (+30% orders, -10% take-rates)	Time-series	702.91	17.94%	0.182	2.36
	MRP	637.94	16.98%	0.2005	32.92
	DF	110.67	2.82%	0.0278	5.11
#4 (-30% orders, -10% take-rates)	Time-series	383.73	18.29%	0.1872	2.35
	MRP	342.12	17.11%	0.2036	40.55
	DF	83.88	4.44%	0.0452	2.07
#5 (+30% orders, +10% take-rates)	Time-series	631.64	18.88%	0.1931	2.51
	MRP	559.23	18.07%	0.2158	58.89
	DF	164.96	5.74%	0.0584	3.96
#6 (-30% orders, +10% take-rates)	Time-series	330.13	18.14%	0.1831	2.6
	MRP	284.07	17.01%	0.1996	27.48
	DF	73.02	4.76%	0.0191	2.15
#7 (+30% orders, 15% random take-rates interval)	Time-series	692.55	18.74%	0.1882	2.41
	MRP	661.19	18.70%	0.217	31.38
	DF	141.32	3.80%	0.0367	4.25
#8 (-30% orders, 15% random take-rates interval)	Time-series	416.39	22.09%	0.2437	2.63
	MRP	460.07	24.46%	0.3145	23.07
	DF	79.46	4.21%	0.0432	1.97

Table 4.2.: Schedule accuracy and computational performance of the component demand planning methods

vehicles each month) and (3) very valuable (for example cockpits in different variations, door panels, seats, electronic equipment, etc.).

To measure the capabilities of the procedure while considering perfectly accurate inputs, the MRP algorithm and the data fusion-based algorithm were provided with the actual option take-rates in September as the planned option take-rates. MRP and DF also used the actual sales volume in September as the planned sales volume. The data fusion-based algorithm was trained by using the information from January to May as input and selected the ω value which minimized the MAPE when creating a component demand schedule for June.

The schedule accuracy and the computational performance of the three methods in each of the datasets are illustrated in Table 4.2. Figure 4.3 presents a comparison between the number of components scheduled by each approach and the actual number of components, the so-called A/F plot. The figure illustrates the results for September over all eight datasets. The comparison of the absolute percentage errors over all datasets

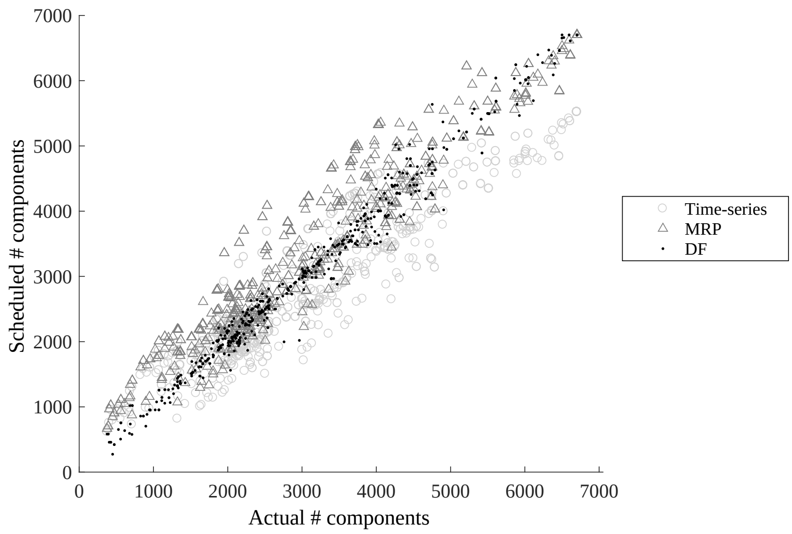

Figure 4.3.: A/F plot for the robustness analysis datasets

is presented in Figure 4.4. The data fusion-based algorithm manages to outperform the other approaches in all datasets while being almost as fast as the time-series forecasting method.

The R software package requires little time to tune the forecast model parameters and can therefore provide results very quickly. Despite the quadratic objective function of the DF model a solution is found very fast for the mathematical model due to the fact that the decision variables are real numbers. In contrast, the MRP approach spends a considerable amount of time searching for past order combinations that would reduce deviations from the planned take-rates. Despite that both DF and MRP are provided with the planned take-rates as inputs, the integrality of the frequencies enforced in MRP prevents the algorithm to find solutions that are very close to the planned take-rates and induces inaccuracies in the component demand plans.

4.5.3. Implementation in a rolling-horizon context

In the second experimental study we collected data on real-world uncertainty and replicated a rolling horizon planning context. In these tests, we also evaluated how the methods would fare when provided with inaccurate input data. We again tested time-series, MRP and DF and compared them additionally with regard to the stability of the resulting plans. We implemented the expression described by Kimms, 1998 to measure the component demand plan stability. The component demand stability measure quantified

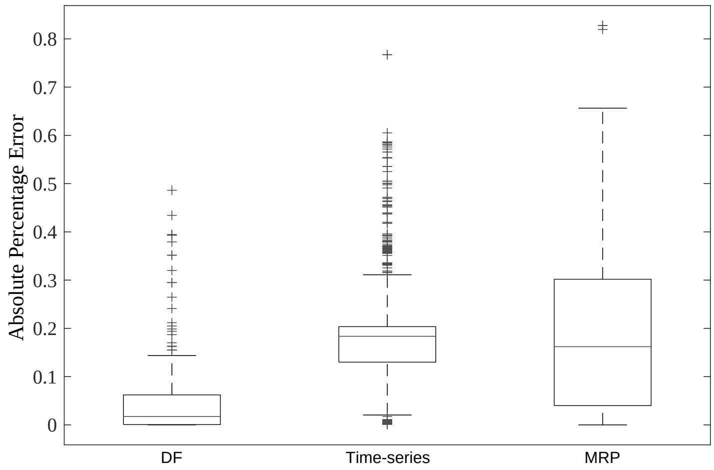

Figure 4.4.: Absolute percentage error spectrum for the component demand planning methods in the robustness analysis datasets

the magnitude of the scheduled component demand change from one planning run to the other.

We generated 40 datasets based on the real-world data. In each of these, we considered a time horizon of a year with rolling component demand updates at the end of months five, seven and nine. In each planning period, the component demand for month twelve was scheduled by using the demand information from the corresponding last five months. The same input data structure as described in section 4.5.2 was provided to the algorithms. Each scheduling run was done independently of the others i.e. previously planned component demands for the twelfth month were overwritten. We generated fictive orders by taking the average number of orders, their standard deviation, the average option take-rates, as well as their standard deviation from the real-world dataset into consideration.

To simulate the information inaccuracies inherent in a rolling horizon framework, we introduced information biases in the inputs. When integrating a bias of $x\%$ for the option take-rates, the algorithms were provided with random planned option take-rates between $(1-x)\%$ and $(1+x)\%$ of the actual take-rates in the twelfth month. To avoid the overlapping effects of two information biases, we did not consider the production volume information bias. For reasons of confidentiality, the real-world information bias could not be used. Instead, we selected a realistic bias of 12% when planning the component demand in month five, 4% in month seven and no bias in month nine. Table

Method	MAD	MAPE	Thiel's U	Average solving time (s)
Time-series	901.13	36.41%	0.4859	1.56
MRP	632.65	27.88%	0.3513	233.72
DF	182.18	8.75%	0.0949	3.15

Table 4.3.: Schedule accuracy and computational performance of the component demand planning methods for the rolling horizon context datasets

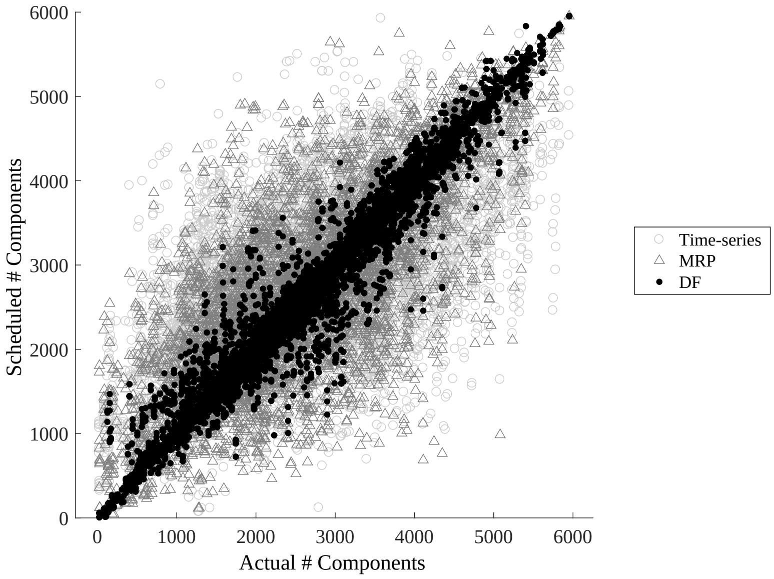

Figure 4.5.: A/F plot for the rolling horizon context

4.3 illustrates the schedule accuracy measures over all 40 datasets, as well as the average computational performance of the three methods. Figure 4.5 illustrates the A/F plot for the study.

In all cases DF outperforms MRP and time-series forecasting with regard to the accuracy measures by more than 60%. One reason is the capability of the data fusion-based method to fully incorporate the planned option take-rates and, thus, to translate the reduced information bias in subsequent months in more accurate component demand plans. Whereas the data fusion-based algorithm manages to provide very accurate component demand predictions for months seven and nine, the other two procedures result in significant deviations from the actual demand for a significant number of components.

Figure 4.6 illustrates the stability measure values for the three methods for every planning run transition. DF not only provides the most accurate component demand schedules, it also leads to minimal schedule changes from one planning run to the other.

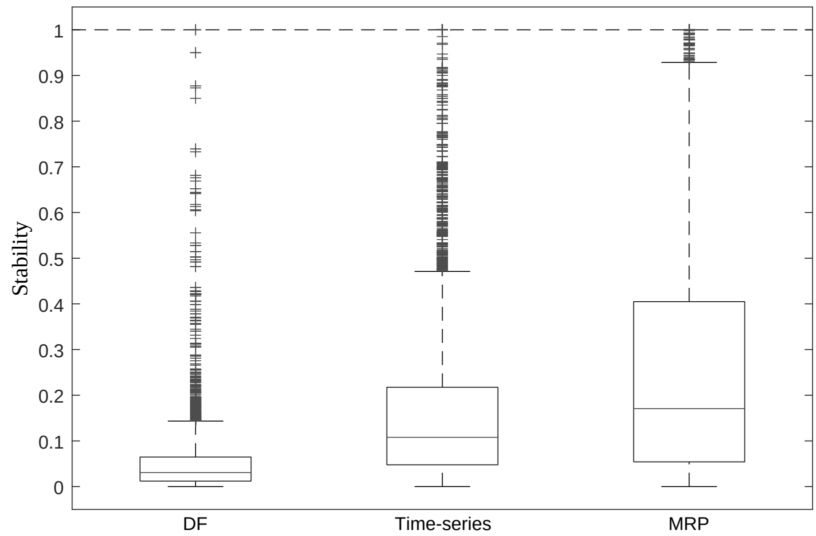

Figure 4.6.: Stability measure spectrum for the rolling horizon context

The schedules generated by the data fusion-based method are more stable than those provided by time-series forecasting or the MRP-based method for more than 80% of the components. By using the data fusion-based algorithm, Mercedes can provide suppliers with stable, reliable component demand information. None of the other approaches are capable of achieving such a performance.

Nevertheless, a comparison between Figure 4.3 and 4.5 shows that there are more components for DF that have an inaccurate component demand plan. These plans result in month five, when the planned option take-rates are inaccurate. The results highlight that the quality of component demand plans is influenced by the accuracy of the planned option take-rates used by the DF approach.

4.6. Implementation at Mercedes-Benz Cars

In preparation for a future roll-out, a project within the Sales department was initiated that investigated the possibility of integrating option take-rates in the planning process. It was at this point that the responsibility for the project was turned over from the Research department to the IT department. In collaboration with the IT department, further system requirements, exchange mechanisms with the wider IT landscape, and aspects of usability were investigated. Key requirements here were an autonomous operation (with minimal maintenance) and a user interface that did not require the planning staff to understand the core logic of the system. Hence, an automatic parameterization of the model was designed with an ex-post MAPE minimization, which runs weekly and self-adjusts the algorithm. This mechanism was tested over an extended period of time to ensure the stability of the method with a low exception rate. The specification was presented to management, but due to the deep crisis the automotive industry experienced in 2009, the roll-out was postponed.

The project was restarted in 2012 in collaboration with the IT department to draft a first complete system specification. Again, different software vendors were asked to develop the new system and prepare the global roll-out and implementation. As a corporate project it then followed the standard roll-out procedure with its testing, compatibility, documentation, and training phases. While the full details of the implementation and integration into the system landscape at Mercedes-Benz were deemed too commercially sensitive to be discussed in detail here, conceptually the algorithm remained as described in Section 4.4.

A contract with a software provider was signed in late 2012 and the change management process was initiated. Training of the 500 operators began in 2013. As of 2013 the prototype has been still part of the planning process for all new product lines at Mercedes-Benz (while existing product lines remained on the legacy MRP-system until phase-out). The complete changeover to the new procedure was completed with the last phase out of existing product lines in late 2015.

The new component demand planning approach has added the novel capability for the Mercedes-Benz planning staff to simulate different sales scenarios and their consequences at component-level, a process that the Central Planning department was only rarely able to perform in the past due to the long run-times of the existing MRP system. The planning staff had reported to us that in the past it could take up to four days to get results from such simulations within the existing MRP system. Planners are now able to compute and compare various scenarios in less than a single day.

4.7. Conclusion

MRP systems constitute the backbone of the production planning infrastructure in manufacturing, yet their inability to cope with dynamic changes of the environment in which they operate is increasingly being identified as a major area of concern (Retting, 2007, Winters et al., 2008, Goodhue et al., 2009, Fauscette, 2013, Tenhiälä and Helkiö, 2015). In this paper we have answered this call by reporting on the development, validation and implementation of the algorithm developed by Stäblein, 2008 that improves the operation of an MRP system in the context of high product variety and global manufacturing and sourcing footprints. The enhancement follows the principles of data fusion, combining information from different sources to improve the component demand plans. Our algorithm follows on from many important improvements to MRP-based systems that have been proposed, both in terms of updating its functionality (see for example: Armentano et al., 2001, Tang and Grubbström, 2002, Xie et al., 2003, Ram et al., 2006, Louly et al., 2008, Ioannou and Dimitriou, 2012, Riezebos and Zhu, 2015), and integrating it with other software systems (see for example: Ferrer and Whybark, 2001, Kreipl and Pinedo, 2004, Garcia-Sabater et al., 2009).

Our computational study has demonstrated that the enhanced MRP has a good computational performance and provides accurate component demand plans even for more difficult datasets. The approach is as fast as the more simple time-series software package, but has a higher component demand forecast accuracy due to the integration of the planned option take-rates. Additionally, the continuous decision variable values allow the approach to have a higher flexibility when combining the planned option take-rates and the past customer orders information. The flexibility also ensures that the component demand plans developed by DF are more stable when embedded in a rolling horizon context.

Our computational study could be enhanced by comparing the DF approach with state-of-the-art Bayesian component demand forecasting models. Unfortunately, no software vendor provided their software packages for such a comparison. Future research can include such methodologies in an extensive computational study.

As the rolling horizon computational study has shown, the accuracy of the component demand plans provided by DF is lower when the planned option take-rates are also inaccurate. Therefore, DF needs to be provided with accurate planned option take-rates to ensure a good accuracy for the component demand plans. Chapter 5 highlights how option bundles can be designed to stabilize the option demand such that the forecast of option take-rates becomes more accurate.

Yet, if anything, our case drastically illustrates the fact that MRP systems are embedded in a context that is rapidly changing, providing an urgent call for more research into the necessary adaptations and extensions needed to ensure MRP systems continue to be ‘fit for duty’ in modern-day manufacturing.

Chapter 5.

Stabilizing the demand for car options by bundling

This chapter is based on an article submitted as:

Popa, R. C. and M. Grunow (2017). "Stabilizing the demand for car options by bundling"

5.1. Introduction

Chapter 4 illustrated that the derivation of accurate component demand forecasts requires accurate option take-rate forecasts. However, the heterogeneity of the customer preferences induces a high volatility in the option take-rates. Figure 5.1 illustrates the take-rate variability for five options offered by a large German automotive manufacturer.

A concept discussed at automotive manufacturers to address the issue is the sta-

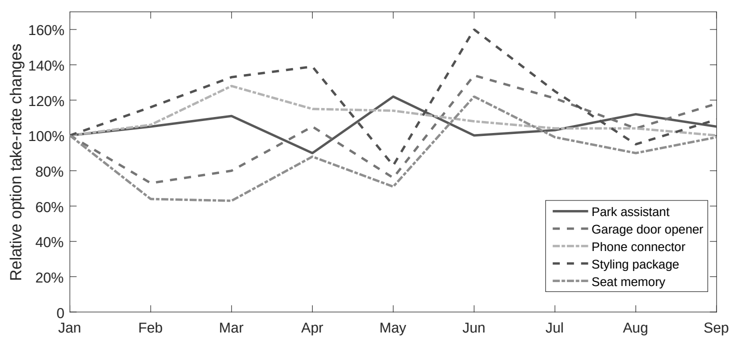

Figure 5.1.: Option take-rates variability for a volume car model

bilization of option take-rates by means of option bundles. The German automotive manufacturer could, for instance, bundle the phone connector and the seat memory system, since they have opposing demand trends. However, revenue may be lost, because customers may want to purchase individual options rather than the bundle.

Traditionally, bundles are designed to maximize revenues or profits. Manufacturers now aim to bundle options such that the higher revenues are combined with the benefits that result from take-rate stabilization. A methodology is required that balances both objectives. The academic literature neither investigates the impact of option bundling on take-rate stability nor does it offer an appropriate bundling methodology.

In contrast to previous work on bundling, the bundle design methodology developed in this paper combines an operations perspective with the more traditional sales perspective. We introduce a branch-and-price option bundle design methodology for automotive manufacturers that captures the trade-off between revenue maximization and take-rate stability. Our methodology can be used to generate a specific number of bundles for various bundling policies. This methodology is able to set an individual discount for each bundle. We tested our method on a large dataset (225,413 orders) provided by a German automotive manufacturer. Our computational study found a trade-off between revenues and take-rate stability. Hence, a sole focus on maximizing revenues has detrimental consequences on take-rate stability. Nevertheless, our methodology identifies bundle designs that simultaneously increase revenues and take-rate stability compared to the unbundling case. Our methodology helps automotive manufacturers design bundles while not only accounting for the sales perspective, but also for the operations perspective.

We present the problem settings and the mathematical model at the foundation of our methodology in section 5.2. The components of our branch-and-price methodology are described in section 5.3. The settings and results of a comprehensive computational study based on our methodology are discussed in section 5.4. We summarize the main managerial findings and address the limitations of our methodology and computational study in section 5.5.

5.2. Problem description

We devised for our problem a model that balances revenue maximization and take-rate stability maximization. The take-rate stability represents the lack of take-rate changes from one period to another.

The model designs a predefined number of bundles based on a set of options. To ensure that customers are motivated to purchase bundles, we simplify the customer choice process. To this end, we allocate each option to only one bundle. The following assumptions are used for our methodology:

a The effects of option bundling observed for past customers are also relevant for the near-future customers if a sufficiently large number of customers is considered.

b The reservation prices of the customers for the individual options are known or can be estimated. The assumption is common for option bundling research and has been used, among others, in Wu et al., 2008 and Hanson and Martin, 1990. Train, 2009 illustrates how reservation prices can be estimated based on past customer purchases using discrete choice methods.

c Each customer acts rationally and selects bundles such that the consumer surplus she achieves through her acquisition is maximized.

d The marketing department sets an upper limit for the bundle discounts.

We use the following notation in our mathematical model:

Sets:	
$o \in O$	Set of options
$b \in B$	Set of bundle candidates
$c \in C$	Set of past customers
$t \in T$	Set of time periods
$\zeta_{(t)}$:	Set of customers who have placed their orders in period t.
Parameters:	
n	Number of bundles to design
$f_{o,b}$	1 if option o is included in bundle b, 0 otherwise
m	Maximum number of options in a bundle
D	Maximum bundle discount
η	Bundle discounts increment
d'_b	Bundle discount of bundle b

p_o	Original selling price of option o
r_b	Total revenues generated by bundle b if offered
$h_{o,t}$	Take-rate of option o in period t
$S_{c,b}$	1 if customer c would select bundle b, 0 otherwise
$\rho'_{o,c}$	Reservation price of customer c for option o
Δ_b^{TT}	Take-rates variability of all options included in bundle b
u	1 if the bundles are designed for mixed bundling, 0 for pure bundling
$s_{o,c}$	1 if option o was selected in the past by customer c, 0 otherwise.
Decision variables:	
λ_b	1 if bundle b is offered, 0 otherwise.

The design of n bundles is equivalent to the selection of n bundles from the set of bundle candidates B. Since in real-life, customers are not usually interested in purchasing very large bundles, it is possible in our model to limit the number of options that a bundle can include to a maximum. We define B as the set of the bundles that contain a maximum of m options from set O. Set B also includes the pseudo-bundles with only one option. Bundles with identical structure, but different bundle discounts are included individually in set B. The bundle discounts d'_b are multiples of the bundle discount increment η, up to a value of D.

One strength of our modeling approach is that it is possible to determine if a customer would purchase a bundle at a given discount before solving the bundle selection problem:

$$S_{c,b} = \begin{cases} 1, \text{ if } (\sum_{o \in O} \rho'_{o,c} \cdot f_{o,b} \geq \sum_{o \in O} p_o \cdot (1 - d'_b) \cdot f_{o,b}) \\ \quad \wedge ((u = 0) \vee (\sum_{o \in O} (\rho'_{o,c} - p_o \cdot (1 - d'_b) \cdot f_{o,b}) \geq \sum_{o \in O} (s_{o,c} \cdot (\rho'_{o,c} - p_o) \cdot f_{o,b})) \\ 0, \text{ otherwise} \end{cases} \tag{5.1}$$

Since each option is assigned to only one bundle and a customer tries to maximize her surplus, she buys the bundle in the pure bundling case if the sum of the reservation prices of the options included in the bundle is at least equal to the price of the bundle. In the mixed bundling case, the customer purchases the bundle only if the surplus that results from the purchase of the bundle is also higher than the surplus that results from the separate purchase of the options included in the bundle that were acquired by the customer in the past.

After the identification of the bundle selections made by the customers, it is possible to determine the total revenues generated by the bundle if it is offered with the discount d'_b:

$$r_b = \sum_{c \in C} (S_{c,b} \cdot \sum_{o \in O} p_o \cdot (1 - d'_b) \cdot f_{o,b} + u \cdot (1 - S_{c,b}) \cdot \sum_{o \in O} s_{o,c} \cdot p_o \cdot f_{o,b}) \tag{5.2}$$

In the pure bundling case, a customer either buys the bundle or discards the acquisition of the bundle and thereby does not receive any of the options included in the bundle. In the mixed bundling case, she either buys the bundle or acquires the desired options separately. The second term of equation (5.2) represents the revenues that are generated through the separate acquisition of the options included in the bundle.

The take-rates that result for each option in each period can be calculated based on the bundle and option selections. The take-rate of an option in a period is the share of vehicles requested in that period that also contain the option:

$$h_{o,t} = \sum_{c \in \zeta_{(t)}} \frac{S_{c,b} \cdot f_{o,b} + u \cdot s_{o,c} \cdot (1 - S_{c,b}) \cdot f_{o,b}}{|\zeta_{(t)}|} \tag{5.3}$$

The second term of equation (5.3) includes the separate purchases of the options for the mixed bundling case in the calculation of the take-rates.

After the option take-rates have been determined, the total take-rate stability of the options included in the bundle can be calculated based on equation (5.4). The stability is defined as the negative sum of the absolute take-rate changes from one period to the next.

$$\Delta_b^{TT} = -\sum_{o \in O} f_{o,b} \cdot \sum_{t \in \{2...|T|\}} |h_{o,t} - h_{o,t-1}| \tag{5.4}$$

The design of n bundles can then be represented by the following compact binary model:

$$\text{Maximize} \begin{cases} F_1 = \sum_{b \in B} \lambda_b \cdot r_b, \\ F_2 = \sum_{b \in B} \lambda_b \cdot \Delta_b^{TT} \end{cases} \tag{5.5}$$

subject to:

$$\sum_{b \in B} \lambda_b = n \tag{5.6}$$

$$\sum_{b \in B} \lambda_b \cdot f_{o,b} = 1, \forall o \in O \tag{5.7}$$

Expression (5.5) represents the objectives of our methodology: the maximization of total revenues and option take-rate stability. Constraint (5.6) enforces exactly n bundles being offered. Constraint (5.7) requires that each option is included in one, and only one of the offered bundles at only one discount level.

We obtain the Pareto-efficient frontier for the problem by means of a standard ϵ constraint methodology. Let F_1^{max}, F_1^{min}, F_2^{max} and F_2^{min} be the payoff table values for the two objectives, τ be the number of Pareto curve points to be determined and $\epsilon_{j,2}, \cdots, \epsilon_{j,\tau_j - 1}$ be the limits for the total take-rate stability.

The model used to design bundles for the Pareto frontier point τ' is defined by the following expressions:

$$\text{Maximize } F_1 \tag{5.5'}$$

subject to:

$$\text{Constraints (5.6), (5.7)}$$

$$F_2 \geq \epsilon'_\tau \tag{5.8}$$

5.3. Branch-and-price-based option bundles design

Due to the number of decision variables, the model defined by expressions (5.5') - (5.8) cannot be used for realistic problem instances, in which dozens of options must be bundled. We therefore developed a branch-and-price methodology. The methodology can be used both for the revenue maximization, as well as the take-rate stability maximization problem required for the payoff table values. We describe the components of the methodology in the following subsections.

5.3.1. Column generation procedure and branching schemes

We use a relaxed and adapted version of the model defined by expressions (5.5') - (5.8) as part of a column generation procedure. The same relaxations described in chapter 3.3.1 are applied.

The column generation procedure starts with a set B that contains all one-option bundles. For both the revenue maximization and the take-rates stability maximization problems, the adapted model has the same structure:

$$\text{Maximize } F_1 = \sum_{b \in B} \lambda_b \cdot \kappa_b - q \cdot feas^+ - q \cdot feas^- \tag{5.5''}$$

subject to:

$$\sum_{b \in B} \lambda_b = n + feas^+ - feas^- \tag{5.6'}$$

$$\text{Constraints (5.7), (5.8)}$$

Depending on the problem that is solved, κ_b represents the revenues generated by bundle b or the total take-rate stability of the options included in the bundle. Constraint (5.6') allows deviations from the required number of bundles. These deviations are needed to always find a feasible solution for the model. However, the penalty costs q must be set high enough that the model adheres to the required number of bundles, if possible. For the revenue maximization problem, q is defined as:

$$q = |C| \cdot \sum_{o \in O} p_o + 1$$

For the take-rate stability maximization problem, q is defined as:

$$q = (|T| - 1) \cdot |O| + 1$$

When the objective function values for the payoff table are determined, constraint (5.8) is removed.

The structure of the current column generation procedure is typical. The dual variable values π^{opt} corresponding to constraint (5.7) and π^{noBund}, π^{epsilon} corresponding to constraints (5.6') and (5.8) are derived based on the relaxed master problem solution. The dual variable values are used to generate additional bundles that improve the objective function value of the relaxed master problem. We use the subproblem heuristic described in section 5.3.3 to generate bundles. If the heuristic cannot find an objective-improving bundle, a mathematical model for the subproblem is used. The problem-specific subproblem models are described in section 5.3.2.

If the column generation solution is not feasible for our problem, but better than the best solution found so far, the branching scheme Ryan and Foster branching scheme (BS1) and a problem-specific branching scheme are used. The branching scheme BS1 is used when none of the bundles with non-integer λ_b selection values have an identical structure. The second branching scheme BS2 is used when at least two bundles with the same structure, but different price levels, have a non-integer λ_b selection value.

The branching candidates of branching scheme BS2 are of type $(\{f'_o \in \{0,1\} | o \in O\}, -1)$ and $(\{f'_o \in \{0,1\} | o \in O\}, \phi_{b'})$. The set $\{f'_o \in \{0,1\} | o \in O\}$ defines the structure of the bundle based on which the branching is done. The branching candidate $(\{f'_o \in \{0,1\} | o \in O\}, -1)$ represents the branch in which no other bundle with the structure $\{f'_o \in \{0,1\} | o \in O\}$ is offered at any price other than $\phi_{b'}$. The branching candidate $(\{f'_o \in \{0,1\} | o \in O\}, \phi_{b'})$ represents the branch in which the bundle with the structure $\{f'_o \in \{0,1\} | o \in O\}$ and price $\phi_{b'}$ is not offered. The price $\phi_{b'}$ of the bundle with the highest non-integer λ_b value is used for branching. The branch $(\{f'_o \in \{0,1\} | o \in O\}, -1)$ is explored first to find an integer solution quickly.

Our branch-and-price methodology can also be used as a matheuristic. The subproblem model is used only in the root node to ensure that the optimal relaxed master problem solution has been found in the root node. The objective function value of the root node solution is the tightest upper bound that can be found for our problem.

5.3.2. Subproblem models

We developed four mathematical model versions for the design of bundles that improve the objective function value of the relaxed master problem. Two versions are used for the revenue maximization problem and two for the option take-rate stability maximization problem. For each problem there is a model version for pure bundling and one for mixed bundling. All models design the bundle that has the highest positive reduced costs.

We use the following additional notation for all model versions:

Additional sets:

l :	Set of active BS1 branching constraints

$b' \in \Theta$:	Set of bundles relevant for the branching scheme BS2
Additional parameters:	
$\phi'_{b'}$:	Price of bundle b' that should not be generated again, -1 if no bundle with the structure of bundle b' should be generated
M_c, M_c^* :	Big, customer-specific numbers
M' :	Maximum price of a bundle that contains m options
Decision variables:	
x_o :	1 if option o is included in the new bundle, 0 otherwise
y_c :	1 if customer c acquires the new bundle, 0 otherwise
$\theta_{o,t}$:	Number of customers who select option o in period t, $\theta_{o,t} \geq 0$
$\Delta^+_{o,t}, \Delta^-_{o,t}$:	Auxiliary variables that represent the absolute take-rate change of option o from period $t-1$ to period $t, \Delta^+_{o,t}, \Delta^-_{o,t} \geq 0$
P :	Price of the new bundle, $P \geq 0$
α :	Multiple of the discount increment used for the bundle price, $\alpha \in \mathbb{N}$
Φ_c :	Revenues generated by customer c due to the new bundle
$\beta^+_{b'}, \beta^-_{b'}$:	Deviations of the price of the new bundle from the price of bundle b'
$\psi_{b'}$:	1 if the price of the generated bundle is larger than the price of bundle b', 0 otherwise.

After successfully running the required subproblem model, the new bundle defined by the decision variables x_o, the corresponding objective function contribution κ_b and the take-rate stability Δ_b^{TT} are added in the relaxed master problem model.

Pure bundling model

In the following, the revenues maximization subproblem model for the pure bundling case is presented.

Objective function:

$$\text{Maximize } \kappa_b = F_{\text{rev}} = \sum_{c \in C} \Phi_c - \sum_{o \in O} \pi^{\text{opt}} \cdot x_o - \pi^{\text{noBund}} - \pi^{\text{epsilon}} \cdot \sum_{t \in \{2..|T|\}} \sum_{o \in O} (\Delta^+_{o,t} + \Delta^-_{o,t}) \tag{5.9}$$

Subject to:

$$\kappa_b \geq 0 \tag{5.10}$$

$$2 \leq \sum_{o \in O} x_o \leq m \tag{5.11}$$

$$\sum_{o \in O} (1 - D) \cdot p_o \cdot x_o \leq P \leq \sum_{o \in O} p_o \cdot x_o \tag{5.12}$$

$$P = \eta \cdot \alpha \tag{5.13}$$

$$\sum_{o \in O} \rho'_{o,c} \cdot x_o - P + M_c^* \cdot (1 - y_c) \geq 0 \quad \forall c \in C \tag{5.14}$$

$$\sum_{o \in O} \rho'_{o,c} \cdot x_o - P - M_c \cdot y_c < 0 \quad \forall c \in C \tag{5.15}$$

$$\theta_{o,t} \leq \sum_{c \in \zeta_{(t)}} y_c \quad \forall o \in O, t \in T \tag{5.16}$$

$$\theta_{o,t} \leq |\zeta_{(t)}| \cdot x_o \quad \forall o \in O, t \in T \tag{5.17}$$

$$\theta_{o,t} \geq \sum_{c \in \zeta_{(t)}} y_c - |\zeta_{(t)}| \cdot (1 - x_o) \quad \forall o \in O, t \in T \tag{5.18}$$

$$\Phi_c \leq P \quad \forall c \in C \tag{5.19}$$

$$\Phi_c \leq M' \cdot y_c \quad \forall c \in C \tag{5.20}$$

$$\frac{\theta_{o,t} - \theta_{o,t-1}}{|\zeta_{(t)}|} + \Delta^-_{o,t} - \Delta^+_{o,t} = 0 \quad \forall o \in O, t \in \{2..|T|\} \tag{5.21}$$

$$x_{o_1} = x_{o_2}, \quad \forall o_1, o_2 \in O : (o_1, o_2, 1) \in l \tag{5.22}$$

$$x_{o_1} + x_{o_2} \leq 1, \quad \forall o_1, o_2 \in O : (o_1, o_2, 0) \in l \tag{5.23}$$

$$\sum_{o \in O : f_{o,b'} = 1} x_o - \sum_{o \in O : f_{o,b'} = 0} x_o \leq \sum_{o \in O} f_{o,b'} \ \forall b' \in \Theta : \phi_{b'} = -1 \tag{5.24}$$

$$\sum_{o \in O : f_{o,b'} = 1} x_o - \sum_{o \in O : f_{o,b'} = 0} x_o - \beta^+_{b'} - \beta^-_{b'} \leq \sum_{o \in O} f_{o,b'} \ \forall b' \in \Theta : \phi_{b'} \neq -1 \tag{5.25}$$

$$P - \beta^+_{b'} + \beta^-_{b'} = \phi_{b'} \ \forall b' \in \Theta : \phi_{b'} \neq -1 \tag{5.26}$$

$$\beta^+_{b'} \leq M' \cdot \psi_{b'} \ \forall b' \in \Theta : \phi_{b'} \neq -1 \tag{5.27}$$

$$\beta^-_{b'} \leq M' \cdot (1 - \psi_{b'}) \ \forall b' \in \Theta : \phi_{b'} \neq -1 \tag{5.28}$$

The objective function (5.9) represents the maximization of the reduced costs of the designed bundle for the revenue maximization master problem. The first term represents the revenues generated by the new bundle. The second term represents the reduced costs decrease caused by the options included in the bundle. The last term represents the reduced costs decrease caused by the total take-rate variability of the bundle.

The new bundle can only improve the objective function of the relaxed master problem if it has positive reduced costs. We introduce constraint (5.10) to improve the computational performance of the model. Since all bundles that contain one option are provided at the start of the column generation procedure, constraint (5.11) hinders the subproblem model to generate discounted bundles with only one option. The constraint also enforces that the new bundle contains at most m options. Constraint (5.12) represents the bounds of the bundle price. Constraint (5.13) enforces that the bundle price is a multiple of the bundle increment η. Constraints (5.14) and (5.15) represent the bundle selection mechanism of the customers. Since each customer selects bundles such that her surplus is maximized and each option is assigned to only one bundle, a customer selects the new bundle only if the price of the bundle is at most equal to her willingness to pay for the options included in the bundle. Constraints (5.19) and (5.20) set the revenue generated by the new bundle for a specific customer equal to the price of the bundle if the customer buys the bundle, 0 otherwise. Constraints (5.16) - (5.18) determine the number of customers who selected an option in a particular month. The $\theta_{o,t}$ values are then used in equation (5.21) to calculate the absolute option take-rate changes for each option and for each time period. Constraints (5.22) and (5.23) enforce the active branching constraints that require or prohibit the inclusion of pairs of options in the generated bundles. Constraint (5.24) is used in the branches that require that a certain bundle b' is offered only with a certain price. The constraint prohibits the design of any bundle that contains the same options as bundle b'. Constraints (5.25) - (5.28) allow the generation of a bundle containing the same option as bundle b' as long as the price of the new bundle is different from the price of bundle b', $\phi_{b'}$.

For the design of the bundle that maximizes take-rate stability, only the objective function needs to be adapted:

$$\text{Maximize } \kappa_b = F_{\text{var}} = -\sum_{t \in \{2..|T|\}} \sum_{o \in O} (\Delta^+_{o,t} + \Delta^-_{o,t}) - \sum_{o \in O} \pi^{\text{opt}} \cdot x_o - \pi^{\text{noBund}} \tag{5.9''}$$

Constraints (5.10) - (5.28) are also used for the take-rates stability maximization subproblem model.

Mixed bundling model

In the mixed bundling case, the model needs to include the impact of options purchased separately by the customers. To this end, the $\theta_{o,t}$ decision variables from the pure bundling models are replaced by new decision variables:

$\theta'_{o,c}$: 1 if customer c selects option o as part of the new bundle, 0 otherwise

The adaptations of the model are presented in the following:

Objective function:

$$\text{Maximize } \kappa_b = F'_{\text{rev}} = F_{\text{rev}} + \sum_{c \in C} (\sum_{o \in O} (x_o - \theta'_{o,c}) \cdot s_{o,c} \cdot p_o) \tag{5.9*}$$

Subject to:

Constraints (5.10) - (5.13)

$$\sum_{o \in O} \rho'_{o,c} \cdot x_o - P - \sum_{o \in O} (\rho'_{o,c} - p_o) \cdot s_{o,c} \cdot x_o + M^*_c \cdot (1 - y_c) \geq 0 \quad \forall c \in C \tag{5.14'}$$

$$\sum_{o \in O} \rho'_{o,c} \cdot x_o - P - \sum_{o \in O} (\rho'_{o,c} - p_o) \cdot s_{o,c} \cdot x_o - M_c \cdot y_c < 0 \quad \forall c \in C \tag{5.15'}$$

$$\theta'_{o,c} \leq y_c \quad \forall o \in O, c \in C \tag{5.16'}$$

$$\theta'_{o,c} \leq x_o \quad \forall o \in O, c \in C \tag{5.17'}$$

$$\theta'_{o,c} \geq y_c + x_o - 1 \quad \forall o \in O, c \in C \tag{5.18'}$$

Constraints (5.19), (5.20)

$$\frac{\sum_{c \in \zeta_{(t)}} (\theta'_{o,c} + (1 - \theta'_{o,c}) \cdot s_{o,c})}{|\zeta_{(t)}|} - \frac{\sum_{c \in \zeta_{(t)}} (\theta'_{o,c} + (1 - \theta'_{o,c}) \cdot s_{o,c})}{|\zeta_{(t-1)}|}$$
$$+ \Delta^-_{o,t} - \Delta^+_{o,t} = 0 \quad \forall o \in O, t \in \{2..|T|\} \tag{5.21'}$$

Constraints (5.22) - (5.28)

In the adapted objective function (5.9*) we also include the revenues generated by the customers who purchase the options offered in the bundle separately. Constraints (5.14') and (5.15') enforce that a customer selects the bundle only if the surplus gained by purchasing the bundle is also higher than the surplus that results for a separate acquisition of the options in the bundle. Constraints (5.16') - (5.18') set $\theta'_{o,c}$ equal to 1 if customer c buys the bundle generated by the model and option o is in the bundle. In constraint (5.21'), the take-rate stability calculation also includes the number of options that the customers purchase separately.

The model for the design of the bundle for the take-rate stability maximizing problem is almost identical to the revenues maximization mixed bundling model. The only difference is that the objective function (5.9") needs to be used instead of (5.9*).

5.3.3. Subproblem heuristic

For the different subproblems, we use a greedy multi-core heuristic. This generates a large number of bundles with positive reduced costs in a reasonable amount of time. The reduced costs are determined by the objective functions of the models introduced in the previous section. The heuristic is initialized with all bundles that contain a pair of options. Then, the heuristic searches on each of the available CPU cores for whether an initial bundle can be expanded with additional options such that the reduced costs are increased. The search of an individual CPU core is terminated as soon as the expansion of a bundle does not increase the reduced costs, all options have been added to the bundle or the maximum number of options in a bundle has been reached. The heuristic can store up to 1,000 bundles with positive reduced costs per CPU core search. At the end of the search, the heuristic adds the stored bundles to the relaxed master problem.

5.4. Computational study

A large German automotive manufacturer provided us with data for three car models: a volume model (M1), a niche model (M2) and a luxury model (M3). For each of the car models, 30 options were made available for our bundling methodology. The options were electronics and seating options that could be combined without restrictions. We were also provided with customer purchase data that included the option selections made by the customers, as well as the months in which the purchases were made. Table 5.1 illustrates the number of past customer purchases included in the datasets, the number of months over which the customers were spread and the average take-rates and their stability.

Car model	Model M1	Model M2	Model M3
# customers	150,194	39,194	36,025
# months	10	9	7
∅ take-rates	23.47%	35.91%	48.17%
∅ take-rate stability	-4.24%	-4.91%	-2.65%

Table 5.1.: Number of past customer purchases and months included in the case study

Since the marketing department of the automotive manufacturer would not accept bundle designs that concentrate many options in a small number of bundles, we generated bundle designs that contained a number of bundles between $\lceil |O|/2 \rceil$ and $|O|-1$. For the same purpose we also limited the maximum number of options in a bundle to 10. We designed bundles separately for a pure and a mixed bundling policy. For each designed bundle configuration we generated a Pareto frontier that consisted of 10 points.

We ensured that the performance of our bundles was not a result of overfitting by splitting the customer data in a training and validation dataset. The customer purchases from the last three months in the datasets were used as validation data. Data from the remaining months were used as training data. The customer purchases from the training datasets were used by our branch-and-price approach to design bundles. The correlation of the take-rate stability and the revenues between the training and the validation datasets was above 0.92 for all car models for pure and mixed bundling. The very high correlation of each performance measure suggests that the bundles generated by the matheuristic do not exhibit a dataset-specific influence on the revenues and the take-rate stability. Nevertheless, the performance evaluation in the next sections is still based on the customers from the validation dataset.

We set the input parameters for our tests based on discussions with colleagues from the automotive manufacturer. Accordingly, the reservation price of a customer for an option she previously purchased was set 10% higher than the price of the option and the reservation price of a customer for an option she did not purchase in the past was set equal to 30% of the price of the option. We set the upper limit for the bundle discounts to 15% and the discount increment to 0.5%.

We ran all tests on Intel Xeon machines with 8 CPU cores and 16 GB RAM. The relaxed master problem and the subproblem models were solved with IBM ILOG CPLEX 12.7. The models were embedded in a branch-and-price framework we developed in C#.

Car model	Matheuristic small test instances		Optimizing small test instances		Realistic test instances	
	Pure	Mixed	Pure	Mixed	Pure	Mixed
∅ computation time (s)	100.30	355.88	259.35	3287.35	877.74	8,201.14
% test instances solved to optimality	100%	91.52%			27.08%	53.24%
∅ optimality gap	0%	0.01%			0.36%	0.14%

Table 5.2.: Average computation time, share of optimal test instances and average optimality gap for the small and the realistic test instance

5.4.1. Computational analysis

We first compared the performance of the branch-and-price matheuristic to the optimizing branch-and-price approach on small test instances. In these tests, we designed bundle configurations from 10 randomly selected options for each car model. 197 bundle configurations were generated by each methodology, 138 of which for pure bundling and 59 for mixed bundling. We then designed bundles with the matheuristic based on the complete set of options provided to us in realistic test instances. For these instances, we set a computation time limit of 3 hours. The matheuristic output provided the best bundle configuration found during that time. 857 bundle configurations were generated for the realistic test instances: 441 for pure bundling and 416 for mixed bundling. Table 5.2 presents an overview of the computational performance for the small, and the realistic test instances.

The overall average computation time for the matheuristic for the small test instances was substantially lower than for the optimizing approach. For the hard-to-solve mixed bundling case, the calculation time can be reduced by almost 90%. All the bundles generated by the matheuristic for pure bundling were proven to be optimal. For only less than 10% of the mixed bundle configurations, the matheuristic did not generate the optimal solution. For none of these configurations did the optimality gap exceed 0.11%.

For the realistic test instances, the matheuristic managed to always generate pure bundling configurations in less than 3 hours. Despite that not more than 27% of the bundle configurations were proven optimal, the average optimality gap was very low. The computation time for mixed bundling was a lot higher than for pure bundling, since often the matheuristic spent a lot of time in the branching phase.

5.4.2. Revenues and take-rate stability

We analyzed the impact of option bundles on revenues and take-rate stability based on the results for the realistic test instances. We first present an example of the bundle configurations generated by the matheuristic. We then provide an overview of the impact of bundling for a pure and a mixed bundling policy. To relate the results to the concentration of options to bundles, we represent the number of bundles in our analysis again by means of the bundling intensity.

Example of bundle configurations.

Figure 5.2 illustrates how option bundles induce revenue and take-rate stability changes. It presents two bundle configurations designed for model M1 for a pure bundling policy and a bundling intensity of 47%. Figure 5.2 shows only the options included in a bundle. The figure showcases the average take-rate and take-rate change between periods for each option and bundle.

The bundles that maximize revenues induce an average increase of the take-rates at the cost of take-rate stability. The increase is generated by high-take-rate options, such as 25, that increase the take-rates of the options with which they are bundled, e.g. 24 is raised from 8.64% to 57.53%. Over all options, the average take-rate is increased by 34% compared to unbundling. However, the revenue-maximizing bundles reduce the average take-rate stability. For option 24 the average take-rate change increases from 4.74% to 20.79%. Over all options, the average take-rate change increases on average by 24% compared to unbundling. Note that the high bundling intensity of 47% also requires the creation of bundles which reduce revenues, e.g. bundle (13,14). For this example, the bundling intensity that maximizes revenues lies at 30%.

In contrast, the bundles designed for the 3rd Pareto frontier point (for which there is a lower limit on the take-rate stability) result in a simultaneous increase of the revenues and take-rate stability. Option 24 is now bundled with options 22 and 23. Its average take-rate is moderately increased to 15.58%. However, its take-rate changes are reduced to an average of 3.84%. Over all options, the increases in the take-rate and in the take-rate stability are 22% and 2%, respectively. Note that revenue maximization remains the objective for bundling. Therefore some bundles still have negative effects on take-rate changes (but positive effects on take-rates).

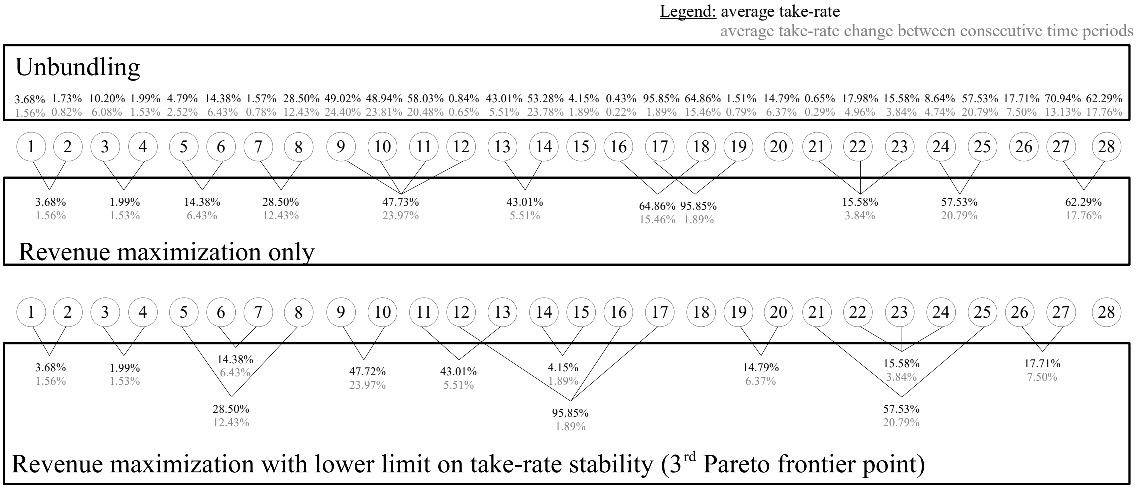

Figure 5.2.: Bundles designed for model M1 for a pure bundling policy and a bundling intensity of 47%

Pure bundling.

Pure bundling is more restrictive on the choices made available to the customers than mixed bundling. Since options can only be bought as part of bundles, a customer who wants to buy a particular option is forced to decide whether she purchases the bundle that contains the option or discards the acquisition of the option. Pure bundling can therefore be designed to efficiently steer customers such that the revenue and take-rate stability objectives are achieved.

A Pareto frontier was generated based on 10 bundle configurations for each car model and each bundling intensity. Figure 5.3 (a) illustrates the Pareto frontiers for car model M2 for three different bundling intensities. The dominant bundle configurations over all bundling intensities were used to derive an overall Pareto frontier, the 'hull'. Figure 5.3 (b) presents the hull for each car model. The revenues and take-rate stability are represented in the figure in relation to their values for unbundling. We only represent the Pareto frontier points for which the revenues were reduced by no more than 15% compared to unbundling. All points with lower revenues would be unacceptable for any automotive manufacturer.

The results highlight the trade-off between revenues and take-rate stability. Our computational study confirms that bundles designed to maximize revenues can increase revenues compared to unbundling. For the car models M1, M2, and M3 in our study, the revenue increase amounts to 10, 13 and 5%. However, our results show that the bundles

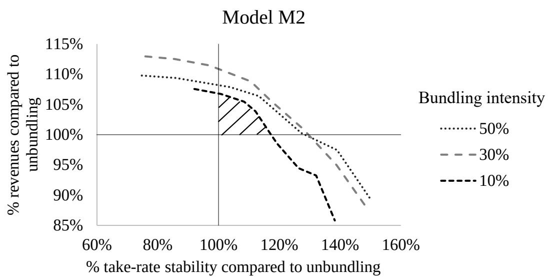

(a)

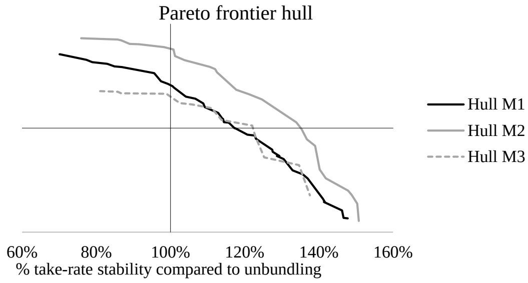

(b)

Figure 5.3.: Pareto frontiers for the pure bundling case

designed only for revenue maximization drastically reduce the take-rate stability. The loss amounts to 30% for M1, 25% for M2 and 19% for M3. The results demonstrate and quantify the drawback of a standard, marketing-oriented bundle design when viewed from an operational perspective.

However, we can design bundles that increase both revenues and take-rate stability. Bundles that do not reduce revenues compared to unbundling can increase take-rate stability by 15% (33%, 22%) for M1 (M2, M3). Hence, our methodology is able to identify a wide range of bundles with synergies between marketing and operations.

The potential of option bundling to improve revenues and take-rate stability is model-specific, as shown in Figure 5.3. The Pareto hulls of models M1 and M3 are very close, with higher revenues and take-rate stability at the edges of the Pareto frontier of model M1. Model M2 has the lowest take-rate stability for unbundling (cf. Table 5.1). Bundling for this model results in the highest revenues and take-rate stability improvements. Since model M2 has many options with a low take-rate stability, many bundles can be generated to stabilize the take-rates of these options. At the same time, some option take-rate peaks can be exploited to subsidize the purchase of other options, thereby increasing revenues.

We further analyzed the impact of bundling intensity on the bundle configurations that improve both measures. To this end, we compared the bundling intensity-specific Pareto frontiers. We calculated the ratio of the surface delimited by each Pareto frontier to the surface delimited by the Pareto frontier hull. An example of a surface used for the ratio calculation is presented in Figure 5.3 (a) as a hashed area for the 10% bundling intensity. Figure 5.5 illustrates the ratio for each car model and each bundling intensity.

For low levels of bundling intensity, an increase of the bundling intensity improves both revenues and take-rate stability. Here, an increase of the bundling intensity enables more

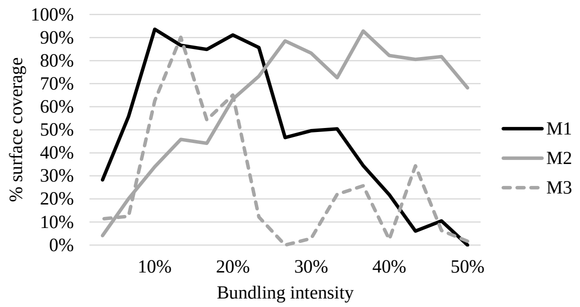

Figure 5.5.: Ratio of surface covered by the Pareto frontiers of each bundling intensity compared to the hull

flexibility in the design of bundles. Our methodology can then identify a larger diversity of bundle configurations with a stronger impact on the trade-off between revenues and take-rate stability.

The surface ratio peaks for each model at a different bundling intensity. Bundling intensities beyond this point lead to a decline of the surface ratio. While models M1 and M3 peak at low levels of bundling intensity, model M2 peaks at a high level of bundling intensity. This corresponds to the results shown in Figure 5.3 (b). For models M1 and M3, there are only few bundles with advantages. For model M2, many more bundles can be identified that have a positive impact on both performance measures. As the bundling intensity increases towards 50%, the surface ratio diminishes for all models.

We also analyzed the average discount of the bundle configurations that resulted in an improvement of both the revenues and the take-rate stability. In principle, discounts increase the appeal of bundles and thus stabilize take-rates. If set too high, however, they also reduce revenues. The resulting average bundle discount in the computational study was 7.88% for model M1, 7.56% for model M2 and 11.59% for model M3. All values are far below the discount limit of 15%. Model M3, has a high take-rate and take-rate stability in the unbundling case (cf. Table 5.1). Only high discounts can motivate customers to purchase even more options as part of bundles.

Mixed bundling.

For a mixed bundling policy, the customers also have the possibility to purchase the options they desire separately, at the original option price. The customers are then not restricted in the choices made available to them. The design space of bundles that simultaneously increase revenues and take-rate stability is therefore more restricted than for pure bundling.

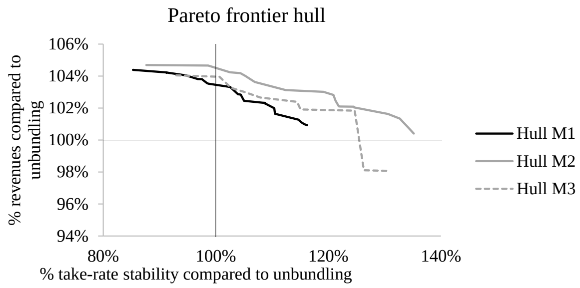

Figure 5.6.: Pareto frontiers for the mixed bundling case

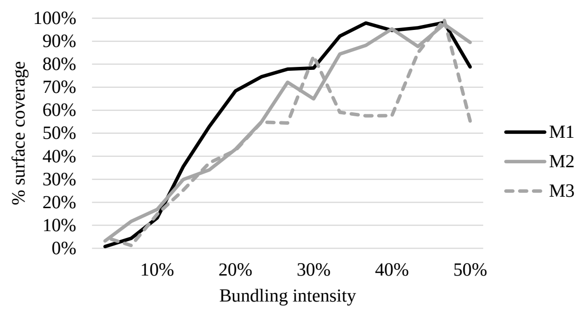

Figure 5.7.: Ratio of surface covered by the Pareto frontiers for mixed bundling compared to the hull

The Pareto frontier hulls in Figure 5.6 confirm the trade-off between revenues and take-rate stability for mixed bundling as well. As for pure bundling, the bundles that maximize revenues also reduce the take-rate stability by 15% for M1, 13% for M2 and 7% for M3 compared to unbundling. At the same time, take-rate stability can be improved by 16% for M1, 35% for M2 and 24% for M3 compared to unbundling, with at least the same revenues as in the unbundling case.

In contrast to pure bundling, the effects of option bundling for mixed bundling become more similar across car models with respect to the trade-off between revenues and take-rate stability. The results highlight that mixed bundling can generate revenue and take-rate stability improvements irrespective of the take-rates and their stability in the unbundling case. However, the improvements are also far less pronounced.

We again analyzed the impact of the bundling intensity on the range of bundle configurations that improve both revenues and take-rate stability. The comparison of the results in Figure 5.7 with 5.5 shows that mixed bundling requires a higher bundling intensity to reach the moderate performance improvement. Furthermore, for a very high bundling intensity of 50%, the surface ratio is again reduced.

Mixed bundling only works when bundle discounts are offered. Otherwise customers

	Model M1	Model M2	Model M3
Coverage for pure bundling	50.36%	55.44%	28.87%
Coverage for mixed bundling	4.88%	82.78%	79.39%

Table 5.3.: Pareto hull surface coverage ratio

would simply purchase the individual options. The average discount of the bundle configurations that result in simultaneous revenues and take-rate stability improvements is 12.99% for model M1, 13.24% for model M2 and 14.29% for model M3. The discounts are consistently high for all car models and for all Pareto frontier hull points.

5.4.3. Impact of discount flexibility

Discounts are used to exploit the consumer surplus and thereby motivate customers to purchase more products than what they originally intend. As shown in Section 2, researchers have developed approaches for option bundling to define bundle-specific discounts that enhance the extraction of consumer surplus and thereby increase revenues or profits. However, the impact of flexible discounts on operational performance indicators has not been studied. We therefore compared the effects of the bundles discussed in Section 5.4.2 with bundles designed by the matheuristic for a fixed discount of 15%. We investigated the ratio of the surface covered by the Pareto hull for the fixed discounts, when revenues and take-rate stability are both larger than in the unbundling case, to the surface of the Pareto hull for the variable discounts. Table 5.3 presents the ratio for each car model for pure and mixed bundling.

The smaller surface for fixed discounts shows that the potential of option bundling to simultaneously increase revenues and take-rate stability is reduced, irrespective of the bundling policy. In the case of model M1 for mixed bundling, the matheuristic identifies only one such design for fixed discounts. Our results show that discounts are a useful tool, not only from a marketing perspective, but also from an operations perspective.

5.5. Conclusion

Up until now, option bundles were designed from a marketing perspective, such that revenues or profits were maximized. Our bundles design approach is the first that actively pursues an operational perspective, the stabilization of the option take-rates.

Our experiments show that, contrary to previous analyses, option bundling does not lead to operational improvements if the operational effects are not taken into account during the bundle design process. A design process performed with the objective of a maximization of the revenues reduces the operational objective of take-rate stability considerably. Therefore, a trade-off exists between revenues and take-rate stability. However, our computational study shows that bundles can be designed to simultaneously increase revenues and take-rate stability.

The potential of option bundling to improve revenues and take-rate stability is best exploited for pure bundling. However, the scale of the improvement is influenced by the take-rates and their stability in the unbundling case. In comparison, the improvements generated by bundles designed for a mixed bundling policy are less pronounced than for pure bundling. Also, mixed bundling requires a higher bundling intensity to reach the moderate performance improvement.

The results highlight the potential of option bundling to increase take-rate stability without any negative impact on revenues. In our experiments the bundles that did not change revenues increased the take-rate stability by at least 15%, regardless of the bundling policy implemented.

Our computational study also shows that bundle-specific discounts are not only an important instrument from a marketing perspective. A common discount rate for all bundles results in lower revenues and take-rate stability. The results show that practitioners need to be aware that the bundle pricing decisions can also impact operational measures.

Even if our matheuristic can design bundles for realistic test instances, the computational performance of the approach could still be improved in further research. For example, branching could be enhanced by identifying stronger upper bounds that do not require an optimal subproblem solution. Alternatively, the number of bundle configurations that need to be generated could be reduced by imposing a lower bound for the revenues. We did not insert such a constraint in order to analyze the complete spectrum of the trade-off between revenues and take-rate stability . However, such a constraint is relevant in industrial settings.

Our work used the typical customer behavior models for bundle selection. Even though there are methods that can estimate reservation prices based on past customer purchases, the datasets made available to us were not suitable for this purpose. An estimation of reservation prices would, for example, require additional demographic information that can be used to identify customer segments. While the current information exchange between the car dealers and manufacturers does not allow for such analysis, future online sales will provide the required data. Future research could then integrate the reservation prices estimation process in the design of bundles.

Furthermore, the results of our computational analysis are based on a dataset with three car models. We have shown how take-rate stability can be taken into account in the design of bundles. Our approach could be applied to a more extensive dataset to draw general conclusions on the extent of operational benefits of bundling for the automotive industry in general.

The standard marketing-oriented design of option bundles has repercussions on operations. Our work provides a methodology automotive manufacturers can use to support an integrated business planning process aimed at identifying a suitable trade-off between marketing and operations.

Chapter 6.

Conclusion and future research directions

6.1. Summary of findings

In this section, we summarize the main findings from the previous chapters and we address the research questions formulated in section 1.6.

RQ1: Do option bundles generally result in a reduction of product variety and, if yes, which factors influence the magnitude of the reduction?

Option bundles design approaches that also minimize product variety are too complex to be used for realistic settings. Therefore, we analyzed whether bundles designed only from a marketing perspective would also indirectly reduce product variety. We therefore developed an option bundles design approach that uses a classic objective function, the maximization of revenues. Our approach designs bundles for pure bundling, while assuming that customers purchase a bundle if their willingness to spend for the bundle exceeds the bundle price.

Our computational analysis shows that there is a trade-off between the revenues and the levels of product variety that result from option bundles. Nevertheless, our approach identifies bundles that can simultaneously increase revenues and reduce the number of car variants. The bundles that do not change revenues compared to the unbundling case reduce the number of car variants substantially. The effects persist for different willingness to spend and bundle discount levels, albeit at different magnitudes. The effects are also robust to the choice of the behavioral model used. Even when customers make their choices according to the classic surplus maximization model, there are no major performance differences.

The results of the computational study indicate that pure bundling results in a reduction of product variety, irrespective of the car model considered. Option bundling can already be an effective tool for no bundle discount, if the customer base is homogeneous and the prices of the options have a wide range. The benefits of option bundling can be obtained for other circumstances with a proper tuning of the bundle discount.

RQ2: Does the enhanced MRP approach developed by Stäblein, 2008 deliver robust results and can it be embedded in a rolling horizon planning cycle?

The enhanced MRP approach developed by Stäblein, 2008 improved the capability of automotive manufacturers to deliver more accurate component demands to their suppliers. However, the author demonstrated the potential of the approach only for a single case study that did not simulate a rolling horizon planning cycle. Chapter 4 addresses these shortcomings by showcasing the robustness of the enhanced MRP approach, as well as the performance of the approach for a rolling horizon planning cycle. We compare the approach to a state-of-the-art time-series forecasting software package, as well as an algorithm based on the MRP logic implemented at Mercedes.

The computational study shows for eight scenarios that the approach developed by Stäblein, 2008 outperforms the other two approaches. The component demand plans are more accurate and are generated almost as fast as by the time-series forecasting software package. The comparison shows that the integration of option take-rate forecasts in MRP approaches greatly improve the accuracy of the component demand forecasts.

We demonstrate the applicability of the approach for a rolling horizon context that is typical for an automotive manufacturer. In our experiments, the forecast of the option take-rates becomes more accurate as the month for which the component demand plan is developed nears. Additionally, the classic MRP approach and the time-series methods generate unstable plans in a rolling horizon context, whereas the enhanced MRP approach stabilizes the component demand plans. Therefore, the fusion of the option take-rate forecasts with past customer preferences information in a mathematical model provides a competitive edge to the classic MRP approach.

RQ3: How can a stabilization of the option take-rates be included in the option bundles design and under which circumstances will bundles lead to a reduction of take-rate variability?

A prerequisite for the success of the enhanced MRP approach is a good forecast of the option take-rates. The accuracy of the take-rate forecasts can be improved if the demand pattern for the options is stabilized. We support automotive manufacturers in this endeavour by developing an option bundling approach that balances the revenues and the take-rate stability. We used the approach in an extensive computational study in which we designed bundles not only for a pure bundling, but also for a mixed bundling policy. The approach is also capable to define bundle-specific discount rates.

Our computational study has shown that the design of option bundles results in a trade-off between revenues and take-rate stability. However, it is possible to identify bundles that simultaneously improve both measures compared to unbundling, but at lower levels than for a design focused solely on one of the two measures. The bundles that do not change revenues compared to unbundling result in a substantial increase of the take-rate stability.

For pure bundling, our approach improves the trade-off between the two measures for different types of car models up to a moderate concentration of options to bundles. The measures can be further improved for higher concentrations of options to bundles if the take-rates and their stability in the unbundling case is low. However, as the bundling intensity approaches 50%, the matheuristic cannot generate as many bundle configurations that result in an optimal trade-off between revenues and take-rate stability.

The revenues and take-rate stability improvement potential of bundles designed for mixed bundling is not as high as for pure bundling. However, the bundles designed to maximize take-rate stability for mixed bundling do not reduce revenues as much as the bundles designed for pure bundling. Furthermore, the revenues - take-rate stability trade-off is more similar across car models and can be improved up to a higher concentration of options to bundles.

6.2. Future research directions

Even though we have developed models and branch-and-price approaches to design option bundles that can easily be adapted for different problems, the applications presented in the thesis do not address all operational aspects relevant for an automotive manufacturer. We have shown that option bundling has residual effects on product variety, but we did not develop an approach that can actually control the level of variety reduced by the bundles. The main difficulty of such an implementation lies in the complexity of the approach: the number of decision variables and constraints would be proportional to the number of variants that would be manufactured. Each bundle offered requires an additional set of constraints to be activated in the mathematical model. Future research should identify modeling approaches that can generate the relevant variants on an as-needed basis. Branch-and-price-and-cut approaches can tackle the problem, but they need to be refined to solve large-scale test instances.

The main issue for a design of bundles that also focuses on operational effects is that the effects cannot be translated in costs. Multi-objective approaches such as the ϵ constraint method generate a Pareto frontier of the revenues and the operational measure. However, the break-even point between the operational improvements and the revenue losses cannot be identified based only on the Pareto frontier. In this regard, future research should quantify the reduced costs that result from operational improvements. With such advances, option bundles could be designed to maximize profits.

In the thesis, we designed methods that require simple input data, such as past customer orders. However, the drawback of datasets that contain only past orders is that the willingness to spend of customers for options or bundles cannot be accurately estimated. Discrete choice methods can be used to obtain more accurate estimations, but there is no clear guideline for the type of data that is needed to have accurate input parameters. Future research should identify the input requirements and define a process that integrates the estimation of the input parameters and the option bundles design process.

We used past customer preferences not only in our option bundles design approach, but also in the enhanced MRP approach. Currently, past customer preferences are included in the enhanced MRP approach by associating a weight to each past customer order. The computational study of Stäblein, 2008 showed that the utilization of more complex methods based on past customer preferences, such as Bayesian networks, did not improve component demand forecast accuracy compared to the enhanced MRP method. However, recent methodologies from Artificial Intelligence research, such as Deep Learning methodologies, could be used to improve component demand forecasts even further. Future research should investigate whether the representation of customer preferences by means of neural networks could be integrated in an enhanced MRP method.

The thesis provides automotive manufacturers with a methodology that helps them determine *how* to design bundles. We did not address the question *when* bundles should be designed and introduced. A design before a model launch has the disadvantage that there is limited information on the preferences of customers for completely new options. Alternatively, bundles can be introduced after a facelift, since by then there is enough information regarding past customer preferences. Since our computational study shows that option bundling can boost revenues in general, bundles could also be introduced after the sales peak stage to maintain high revenue levels. A decision support tool that integrates the design of bundles and the timing of their introduction is needed.

Option bundling is only one of four product variety mitigation strategies. It is unclear what impact the other strategies would have on the operational measures we have investigated. Furthermore, there are methodologies for the design of platforms and modules, but there is no decision support tool that helps automotive manufacturers choose which product variety mitigation strategy to use and how to design the product based on the decision. Ideally, future research will develop a decision support tool that provides practitioners with an integrated product variety mitigation strategy selection and product design process.

Appendix A.

Outline of the branch-and-price approach

```
B ← {|O| bundles containing one option}, l = ∅
LBound = −∞, UBound = ∞
ColumnGeneration(l)

function ColumnGeneration(l)
    S_1: Solve relaxed master problem (section ??)
    Compute dual values π^opt, π^noBund and, if multiple objectives active, π^epsilon, bundle
    selection values λ_b, objective function value Obj from relaxed master problem solution
    B^+ = HEURBUND(Π, π^noBund, π^epsilon, l) (B)
    if B^+ ≠ ∅ then B = B ∪ B^+, GoTo S_1
    else
        Optimal generation of objective-improving bundle B^+
        if B^+ ≠ ∅ then B = B ∪ B^+, GoTo S_1
        else
            if l = ∅ then UBound = Obj end if
            if all λ_b integer then
                if LBound < Obj then LBound = Obj, Save B end if
            else
                if Obj > LBound then
                    Generate branching constraints L
                    for each l' ∈ L do
                        ColumnGeneration(l ∪ l'_i)
                    end for
                end if
            end if
        end if
    end if
end function
```

Appendix B.

Outline of the subproblem heuristic subroutine

function HeurBund($\Pi, \pi^{\text{opt}}, \pi^{\text{noBund}}, \pi^{\text{epsilon}}, l$)
    Sort options in ascending order according to π^{opt} in list Λ.
    $B^s = \{\{o_1, o_2\} | o_1, o_2 \in \Lambda, \Lambda^{-1}_{(o_1)} < \Lambda^{-1}_{(o_2)}\}$
    for each $OP \in B^{\text{start}}$ **do**
        repeat
            $Ok = true$
            for each $o' \in OP, o \in O : o \notin OP$ **do**
                if $(o, o', 1) \in l$ **then** $Ok = false$, $OP = OP \cup \{o\}$ **end if**
            end for
        until Ok
        for each $\{o_1, o_2\} \in OP$ **do**
            if $(o_1, o_2, 0) \in l \vee |OP| > m$ **then** Remove OP from B^s **end if**
        end for
    end for
    repeat
        for each available CPU core **do**
            $OP = B^s_{(1)}$, Remove OP from B^s, $pos = \Lambda^{-1}_{(o_2)} + 1$
            BRC = reduced costs of OP, $BSol = \emptyset$
            repeat
                $fB = false$
                for each $o \in \Lambda : (o \notin OP) \wedge (\Lambda^{-1}_{(o)} > pos)$ **do**
                    RC = Reduced costs resulting from adding o to OP
                    if $RC > BRC$ **then**
                        $BRC = RC$, $BSol = OP \cup \{o\}, fB = true$
                    end if
                end for
                if $BSol \neq \emptyset$ **then** pos = Pos. of last element in $BSol + 1$ **end if**
            until $(fB = false) \vee (pos \geq |\Lambda| + 1) \vee (|OP| \geq m)$
            if $BSol \neq \emptyset$ **then** Add $BSol$ to B^+ **end if**
        end for
    until $B^s = \emptyset$
    return B^+
end function

Bibliography

Adams, W. J. and J. L. Yellen (1976). "Commodity bundling and the burden of monopoly". *The Quarterly Journal of Economics* 90.3, pp. 475–498. DOI: 10.2307/1886045.

Alptekinoğlu, A. and C. J. Corbett (2008). "Mass customization vs. mass production: Variety and price competition". *Manufacturing & Service Operations Management* 10.2, pp. 204–217. DOI: 10.1287/msom.1070.0155.

Armentano, V. A., R. E. Beretta, and P. M. Franca (2001). "Lot-sizing in capacitated multi-stage serial systems". *Production and Operations Management* 10.1, pp. 68–86. DOI: 10.1111/j.1937-5956.2001.tb00068.x.

Barnhart, C., E. L. Johnson, G. L. Nemhauser, M. W. P. Savelsbergh, and P. H. Vance (1998). "Branch-and-price: column generation for solving huge integer programs". *Operations Research* 46.3, pp. 316–329. DOI: 10.1287/opre.46.3.316.

Benton, W. C. (2007). "Multiple price breaks and alternative purchase lot-sizing procedures in material requirements planning systems". *International Journal of Production Research* 23.5, pp. 1025–1047. DOI: 10.1080/00207548508904763.

Bitran, G. R. and J.-C. Ferrer (2007). "On pricing and composition of bundles". *Production and Operations Management* 16.1, pp. 93–108. DOI: 10.1111/j.1937-5956.2007.tb00168.x.

Blackburn, J. D., D. H. Kropp, and R. A. Millen (1986). "A comparison of strategies to dampen nervousness in MRP systems". *Management Science* 32.4, pp. 413–429. DOI: 10.1287/mnsc.32.4.413.

Buchholz, K. (2016). *GM shines brightest in latest OEM-supplier relations study (SAE white paper).* URL: http://articles.sae.org/14823/.

Cachon, G. P. and M. Fisher (2000). "Supply chain inventory management and the value of shared information". *Management Science* 46.8, pp. 1032–1048. DOI: 10.1287/mnsc.46.8.1032.12029.

Cachon, G. P. and M. A. Lariviere (2001). "Contracting to assure supply: How to share demand forecasts in a supply chain". *Management Science* 47.5, pp. 629–646. DOI: 10.1287/mnsc.47.5.629.10486.

Cao, Q., K. E. Stecke, and J. Zhang (2015). “The impact of limited supply on a firm’s bundling strategy”. *Production and Operations Management* 24.12, pp. 1931–1944. DOI: 10.1111/poms.12388.

Cataldo, A., J. C. Ferrer, and G. Bitrán (2017). “On pricing and composition of multiple bundles: A two-step approach”. *European Journal of Operational Research* 259.2, pp. 766–777. DOI: 10.1016/j.ejor.2016.11.010.

Chakravarty, A., A. Mild, and A. Taudes (2013). “Bundling decisions in supply chains”. *European Journal of Operational Research* 231.3, pp. 617–630. DOI: 10.1016/j.ejor.2013.06.026.

Chand, S., V. N. Hsu, and S. Sethi (2002). “Forecast, solution, and rolling horizons in operations management problems: A classified bibliography”. *Manufacturing & Service Operations Management* 4.1, pp. 25–43. DOI: 10.1287/msom.4.1.25.287.

Childerhouse, P., S. M. Disney, and D. R. Towill (2008). “On the impact of order volatility in the European automotive sector”. *International Journal of Production Economics* 114.1, pp. 2–13. DOI: 10.1016/j.ijpe.2007.09.008.

Chung, J. and V. R. Rao (2003). “A general choice model for bundles with multiple-category products: Application to market segmentation and optimal pricing for bundles”. *Journal of Marketing Research* 40.2, pp. 115–130. DOI: 10.1509/jmkr.40.2.115.19230.

Detmer, H. and J. Gebhardt (2001). “Markovnetze für die Eigenschaftsplanung und Bedarfsvorschau in der Automobilindustrie”. *Künstliche Intelligenz* 15.3, pp. 16–22.

Dharmani, S., D. Anand, and M. Demirci (2015). *Shifting gear: Capacity management in the automotive industry (EY white paper)*. URL: http://www.ey.com/GL/en/Industries/Automotive/ey-Capacity-management-in-the-automotive-industry.

Dixon, M. J. and G. M. Thompson (2016). “Bundling and scheduling service packages with customer behavior: Model and heuristic”. *Production and Operations Management* 25.1, pp. 36–55. DOI: 10.1111/poms.12409.

Efthymiou, D. and C. Antoniou (2016). “Modeling the propensity to join carsharing using hybrid choice models and mixed survey data”. *Transport Policy* 51, pp. 143–149. DOI: 10.1016/j.tranpol.2016.07.001.

ElMaraghy, H. et al. (2013). “Product variety management”. *CIRP Annals - Manufacturing Technology* 62.2, pp. 629–652. DOI: 10.1016/j.cirp.2013.05.007.

Eppen, G. D., W. A. Hanson, and R. K. Martin (1991). “Bundling - new products, new markets, low risk”. *Sloan Management Review* 32.4, pp. 7–14.

European Comission (2017). *Reducing CO2 emissions from passenger cars.* URL: https://ec.europa.eu/clima/policies/transport/vehicles/cars_en#tab-0-0.

Fauscette, M. (2013). *Maintaining ERP systems: the cost of change (Unit 4 white paper).* URL: http://www.unit4.com/us/resources/idc-erp-cost-change.

Ferrer, G. and D. C. Whybark (2001). "Material planning for a remanufacturing facility". *Production and Operations Management* 10.2, pp. 112–124. DOI: 10.1111/j.1937-5956.2001.tb00073.x.

Fisher, M. L. and C. D. Ittner (1999). "The impact of product variety on automobile assembly operations: Empirical evidence and simulation analysis". *Management Science* 45.6, pp. 771–786. DOI: 10.1287/mnsc.45.6.771.

Fisher, M. and A. Raman (1996). "Reducing the cost of demand uncertainty through accurate response to early sales". *Operations Research* 44.1, pp. 87–99. DOI: 10.1287/opre.44.1.87.

Fuerderer, R., A. Huchzermeier, and L. Schrage (1999). "Stochastic option bundling and bundle pricing". *Optimal bundling.* Ed. by R. Fuerderer, A. Herrmann, and G. Wuebker. Berlin and New York: Springer, pp. 61–86. ISBN: 3642084591.

Gao, P., H.-W. Kaas, D. Mohr, and D. Wee (2016). *Disruptive trends that will transform the auto industry (McKinsey white paper).* URL: https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/disruptive-trends-that-will-transform-the-auto-industry.

Garcia-Sabater, J. P., J. Maheut, and J. J. Garcia-Sabater (2009). "A capacitated material requirements planning model considering delivery constraints: A case study from the automotive industry". *International Conference on Computers & Industrial Engineering, 2009.* Ed. by I. Kacem. Piscataway, NJ: IEEE, pp. 378–383. ISBN: 978-1-4244-4135-8. DOI: 10.1109/ICCIE.2009.5223806.

Gattiker, T. F. and D. L. Goodhue (2005). "What happens after ERP implementation: understanding the impact of interdependence and differentiation on plant-level outcomes". *MIS Quarterly* 29 (3), pp. 559–585. DOI: 10.2307/25148695.

Gebhardt, J., F. Rugheimer, H. Detmer, and R. Kruse (2004). "Adaptable Markov models in industrial planning". *2004 IEEE International Conference on Fuzzy Systems.* Piscataway, N.J: IEEE, pp. 475–479. ISBN: 0-7803-8353-2. DOI: 10.1109/FUZZY.2004.1375776.

Gebhardt, J., R. Kruse, and H. Detmer (2008). "Knowledge revision in Markov networks". *Mathware & Soft Computing* 11.3, pp. 93–107.

Gertz, V. and F. Haeser (2015). *Next generation product complexity management: Are you ready to use digitisation to manage product variance (Capgemini white paper).* URL: `https://goo.gl/JXAsck`.

Goodhue, D. L., D. Q. Chen, M. C. Boudreau, A. Davis, and J. D. Cochran (2009). "Addressing business agility challenges with enterprise systems". *MIS Quarterly Executive* 8.2, pp. 73–88.

Goodman, I. R., R. P. S. Mahler, and H. T. Nguyen (1997). *Mathematics of data fusion.* Vol. 37. Theory and decision library : Series B, Mathematical and statistical methods. Dordrecht u.a.: Kluwer Acad. Publ. ISBN: 0-7923-4674-2.

Gooijer, J. G. de and R. J. Hyndman (2006). "25 years of time series forecasting". *International Journal of Forecasting* 22.3, pp. 443–473. DOI: `10.1016/j.ijforecast.2006.01.001`.

Hall, D. L. and S. A. H. MacMullen (2004). *Mathematical techniques in multisensor data fusion.* 2. ed. Artech House information warfare library. Boston Mass. u.a.: Artech House. ISBN: 1-58053-335-3.

Hanna, R. and F. Kuhnert (2016). *Reimagining automotive operations (PriceWaterhouseCoopers white paper).* URL: `http://www.pwc.com/gx/en/automotive/pdf/global-ops-survey-automotive.pdf`.

Hanson, W. and R. K. Martin (1990). "Optimal bundle pricing". *Management Science* 36.2, pp. 155–174. DOI: `10.1287/mnsc.36.2.155`.

Hendricks, K. B., V. R. Singhal, and J. K. Stratman (2007). "The impact of enterprise systems on corporate performance: A study of ERP, SCM, and CRM system implementations". *Journal of Operations Management* 25.1, pp. 65–82. DOI: `10.1016/j.jom.2006.02.002`.

Ho, C. J. and T. C. Ireland (1998). "Correlating MRP system nervousness with forecast errors". *International Journal of Production Research* 36.8, pp. 2285–2299. DOI: `10.1080/002075498192904`.

Holweg, M. and F. K. Pil (2004). *The second century: Reconnecting customer and value chain through build-to-order ; moving beyond mass and lean production in the auto industry.* Cambridge, Mass.: MIT Press. ISBN: 0-262-58262-7.

Huffman, C. and B. E. Kahn (1998). "Variety for sale: Mass customization or mass confusion?" *Journal of Retailing* 74.4, pp. 491–513. DOI: `10.1016/S0022-4359(99)80105-5`.

Huq, Z. and F. Huq (1994). "Embedding JIT in MRP: The case of job shops". *Journal of Manufacturing Systems* 13.3, pp. 153–164. DOI: `10.1016/0278-6125(94)90001-9`.

Hyndman, R. J. and Y. Khandakar (2008). “Automatic time series forecasting: The forecast package for R”. *Journal of Statistical Software* 27.3. DOI: 10.18637/jss.v027.i03.

Ioannou, G. and S. Dimitriou (2012). “Lead time estimation in MRP/ERP for make-to-order manufacturing systems”. *International Journal of Production Economics* 139.2, pp. 551–563. DOI: 10.1016/j.ijpe.2012.05.029.

Jacobs, R. F. and T. Weston (2007). “Enterprise resource planning (ERP)—A brief history”. *Journal of Operations Management* 25.2, pp. 357–363. DOI: 10.1016/j.jom.2006.11.005.

Jana, P. and M. Grunow (2017). “Integrated business planning in the automotive industry: Reference process and literature review”.

Jonsson, P. and S.-A. Mattsson (2006). “A longitudinal study of material planning applications in manufacturing companies”. *International Journal of Operations & Production Management* 26.9, pp. 971–995. DOI: 10.1108/01443570610682599.

Kimms, A. (1998). “Stability measures for rolling schedules with applications to capacity expansion planning, master production scheduling, and lot sizing”. *Omega* 26.3, pp. 355–366. DOI: 10.1016/S0305-0483(97)00056-X.

Kreipl, S. and M. Pinedo (2004). “Planning and scheduling in supply chains: An overview of issues in practice”. *Production and Operations Management* 13.1, pp. 77–92. DOI: 10.1111/j.1937-5956.2004.tb00146.x.

Lancaster, K. (1990). “The economics of product variety: A survey”. *Marketing Science* 9.3, pp. 189–206. DOI: 10.1287/mksc.9.3.189.

Lee, H. L., V. Padmanabhan, and S. Whang (2004). “Information distortion in a supply chain: The bullwhip effect”. *Management Science* 50.12_supplement, pp. 1875–1886. DOI: 10.1287/mnsc.1040.0266.

Louly, M. A., A. Dolgui, and F. Hnaien (2008). “Optimal supply planning in MRP environments for assembly systems with random component procurement times”. *International Journal of Production Research* 46.19, pp. 5441–5467. DOI: 10.1080/00207540802273827.

Lübbecke, M. E. and J. Desrosiers (2005). “Selected Topics in Column Generation”. *Operations Research* 53.6, pp. 1007–1023. DOI: 10.1287/opre.1050.0234.

Meyr, H. (2004). “Supply chain planning in the German automotive industry”. *OR Spectrum* 26.4, pp. 447–470. DOI: 10.1007/s00291-004-0168-4.

OICA (2017a). *2005-2016 Sales Statistics*. URL: http://www.oica.net/category/sales-statistics/.

OICA (2017b). *Economic impact*. URL: http://www.oica.net/category/economic-contributions/.

Olhager, J. and E. Selldin (2003). "Enterprise resource planning survey of Swedish manufacturing firms". *European Journal of Operational Research* 146.2, pp. 365–373. DOI: 10.1016/S0377-2217(02)00555-6.

Orlicky, J. (1975). *Material requirements planning: The new way of life in production and inventory management*. New York: McGraw-Hill. ISBN: 0070477086.

Parkin, R., R. Wilk, E. Hirsh, and A. Singh (2017). *2017 Automotive Trends (PwC white paper)*. URL: https://www.strategyand.pwc.com/trend/2017-automotive-industry-trends.

Pierce, B. and H. Winter (1996). "Pure vs. mixed commodity bundling". *Review of Industrial Organization* 11.6, pp. 811–821. DOI: 10.1007/BF00174408.

Pil, F. K. and M. Holweg (2004). "Linking product variety to order-fulfillment strategies". *Interfaces* 34.5, pp. 394–403. DOI: 10.1287/inte.1040.0092.

Popa, R. C. and M. Grunow (2017). "Stabilizing the demand for car options by bundling".

Popa, R. C., M. Grunow, and T. Stäblein (2017). "Product variety reduction through data-driven option bundling".

Proano, R. A., S. H. Jacobson, and W. Zhang (2012). "Making combination vaccines more accessible to low-income countries: The antigen bundle pricing problem". *Omega* 40.1, pp. 53–64. DOI: 10.1016/j.omega.2011.03.006.

Ram, B., M. R. Naghshineh-Pour, and X. Yu (2006). "Material requirements planning with flexible bills-of-material". *International Journal of Production Research* 44.2, pp. 399–415. DOI: 10.1080/00207540500251505.

Ramdas, K. (2003). "Managing product variety: an integrative review and research directions". *Production and Operations Management* 12.1, pp. 79–101. DOI: 10.1111/j.1937-5956.2003.tb00199.x.

Randall, T. and K. Ulrich (2001). "Product variety, supply chain structure, and firm performance: Analysis of the U.S. bicycle industry". *Management Science* 47.12, pp. 1588–1604. DOI: 10.1287/mnsc.47.12.1588.10237.

Retting, C. (2007). "The trouble with enterprise software". *MIT Sloan Management Review* 49, pp. 21–27.

Rickard, D. (2008). *The joy of bundling: Assessing the benefits and risks (Boston Consulting Group white paper)*. URL: https://goo.gl/4v9Gtz.

Riezebos, J. and S. X. Zhu (2015). “MRP planned orders in a multiple-supplier environment with differing lead times”. *Production and Operations Management* 24.6, pp. 883–895. DOI: `10.1111/poms.12318`.

Ringbeck, J., C.-S. Neumann, and A. Cornet (1999). “Market-oriented complexity management using the micromarket management concept”. *Optimal bundling*. Ed. by R. Fuerderer, A. Herrmann, and G. Wuebker. Berlin and New York: Springer, pp. 119–131. ISBN: 978-3-642-08459-1. DOI: `10.1007/978-3-662-09119-76`.

Rom, W. O., O. I. Tukel, and J. R. Muscatello (2002). “MRP in a job shop environment using a resource constrained project scheduling model”. *Omega* 30.4, pp. 275–286. DOI: `10.1016/S0305-0483(02)00033-6`.

Ryan, D. M. and B. A. Foster (1981). “An integer programming approach to scheduling”. *Computer Scheduling of Public Transport Urban Passenger Vehicle and Crew Scheduling*, pp. 269–280.

Schmalensee, R. (1984). “Gaussian demand and commodity bundling”. *The Journal of Business* 57.1, pp. 211–230. DOI: `10.1086/296250`.

Soman, D. and J. T. Gourville (2001). “Transaction decoupling: How price bundling affects the decision to consume”. *Journal of Marketing Research* 38.1, pp. 30–44. DOI: `10.1509/jmkr.38.1.30.18828`.

Stäblein, T. (2008). *Integrierte Planung des Materialbedarfs bei kundenauftragsorientierter Fertigung von komplexen und variantenreichen Serienprodukten*. Vol. Band 18. Innovationen der Fabrikplanung und -organisation. Aachen: Shaker. ISBN: 383226986X.

Stäblein, T., M. Holweg, and R. C. Popa (2016). “Enhancing MRP-based component demand planning in a high-variety context”.

Stadtler, H., C. Kilger, and H. Meyr (2015). *Supply chain management and advanced planning: Concepts, models, software, and case studies*. 5th ed. 2015. Springer Texts in Business and Economics. Berlin Heidelberg: Springer Berlin Heidelberg. ISBN: 978-3-642-55309-7. URL: `http://dx.doi.org/10.1007/978-3-642-55309-7`.

Staeblein, T. and K. Aoki (2015). “Planning and scheduling in the automotive industry: A comparison of industrial practice at German and Japanese makers”. *International Journal of Production Economics* 162, pp. 258–272. DOI: `10.1016/j.ijpe.2014.07.005`.

Steinbrecher, M., F. Rügheimer, and R. Kruse (2008). “Application of graphical models in the automotive industry”. *Computational intelligence in automotive applications*. Ed. by D. Prokhorov. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 79–88. ISBN: 978-3-540-79257-4. DOI: `10.1007/978-3-540-79257-45`.

Stigler, G. J. (1963). “United States v. Loew’s Inc.: A note on block-booking”. *The Supreme Court Review*, p. 152.

Stratman, J. K. (2007). “Realizing benefits from enterprise resource planning: Does strategic focus matter?” *Production and Operations Management* 16.2, pp. 203–216. DOI: 10.1111/j.1937-5956.2007.tb00176.x.

Stremersch, S. and G. J. Tellis (2002). “Strategic bundling of products and prices: A new synthesis for marketing”. *Journal of Marketing* 66.1, pp. 55–72. DOI: 10.1509/jmkg.66.1.55.18455.

Tang, O. and R. W. Grubbström (2002). “Planning and replanning the master production schedule under demand uncertainty”. *International Journal of Production Economics* 78.3, pp. 323–334. DOI: 10.1016/S0925-5273(00)00100-6.

Tashman, L. J. (2000). “Out-of-sample tests of forecasting accuracy: An analysis and review”. *International Journal of Forecasting* 16.4, pp. 437–450. DOI: 10.1016/S0169-2070(00)00065-0.

Tenhiälä, A. and P. Helkiö (2015). “Performance effects of using an ERP system for manufacturing planning and control under dynamic market requirements”. *Journal of Operations Management* 36, pp. 147–164. DOI: 10.1016/j.jom.2014.05.001.

Tönshoff, N., C. H. Fine, and A. Huchzermeier (1999). “Bundling and pricing of modular machine tools under demand uncertainty”. *Optimal bundling.* Ed. by R. Fuerderer, A. Herrmann, and G. Wuebker. Berlin and New York: Springer, pp. 87–117. ISBN: 3642084591.

Train, K. (2009). *Discrete choice methods with simulation.* 2nd ed. Cambridge, New York: Cambridge University Press. ISBN: 978-0-521-76655-5.

Vanderbeck, F. (2011). “Branching in branch-and-price: a generic scheme”. *Mathematical Programming* 130.2, pp. 249–294. DOI: 10.1007/s10107-009-0334-1.

Vollmann, T. E. (2005). *Manufacturing planning and control systems for supply chain management.* 5. ed. New York NY u.a.: McGraw-Hill. ISBN: 0-07-144033-X.

Wendt, L. (2016). *World premiere! Brain scans show WHY we prefer bundles (Linkedin blog).* URL: https://www.linkedin.com/pulse/world-premiere-brain-scans-show-why-we-prefer-bundles-dr-laura-wendt.

Winters, D. B., J. T. Lindley, S. Topping, and L. T. Lindley (2008). “The hidden financial costs of ERP software”. *Managerial Finance* 34.2, pp. 78–90. DOI: 10.1108/03074350810841277.

Wu, S.-y., L. M. Hitt, P.-y. Chen, and A. Anandalingam (2008). "Customized bundle pricing for information goods: A nonlinear mixed-integer programming approach". *Management Science* 54.3, pp. 608–622. DOI: `10.1287/mnsc.1070.0812`.

Xie, J., X. Zhao, and T. Lee (2003). "Freezing the master production schedule under single resource constraint and demand uncertainty". *International Journal of Production Economics* 83.1, pp. 65–84. DOI: `10.1016/S0925-5273(02)00262-1`.

Yelland, P. M. (2010). "Bayesian forecasting of parts demand". *International Journal of Forecasting* 26.2, pp. 374–396. DOI: `10.1016/j.ijforecast.2009.11.001`.

Yeung, J. H., W. C. Wong, and L. Ma (1998). "Parameters affecting the effectiveness of MRP systems: A review". *International Journal of Production Research* 36.2, pp. 313–332. DOI: `10.1080/002075498193750`.