



# **Forecast Support Tool for Sales Promotion Campaigns;**

*Design, Implementation and Evaluation*

**Master Thesis Economics & Informatics**  
**Economics & ICT Program**  
Erasmus School of Economics

**Author: A. Mondom, BSc**

**Student Number: 305376**

**Supervisor: Dr. T. Tervonen**

**Co-reader: Dr. N. Basturk**

**Date: July 8, 2013**





**Author:** A. Mondom, BSc

**EUR Supervisor:** Dr. T. Tervonen

**EUR Co-reader:** Dr. N. Basturk

**SCA Supervisor:** P. Massuger

## **Acknowledgements**

This work would not have been possible without the involvement and support of many individuals whom I would like to show my sincere appreciation. Most of those individuals have frequently asked me "when will you finish the thesis?", which may very well have been one of the most difficult questions to answer. It is a question I have regularly asked myself and has been the source of some torment, if not despair. However, hope does not solely consist of the believe that things will turn out for the best, but is rather a mental state that can only be found through action, by doing something and creating a vision. For example, when my first supervisor left the university I stood up to find a new research topic and a new supervisor. The second supervisor soon also left the university, and I had to change my supervisor yet again. Passing these and other obstacles was not always easy, but I took action, had a certain vision. Together, this created hope. I could not have always done this by myself and I am therefore thankful for all the people that showed understanding and support.

First and foremost, I would like to express my deepest gratitude to my supervisor, Tommi Tervonen, for his advice, guidance, constructive criticism and attention to detail throughout this research. After briefly explaining my research project, he accepted to supervise me and immediately suggested useful literature and a suitable research approach. Briefly explaining a problem or difficulty to him always quickly yielded a comprehensible answer and a practical solution. I am also grateful to Nalan Basturk for being my co-reader. Her input and comments are greatly appreciated.

My journey of writing this thesis would not have been as fruitful and motivational without my friends with whom I had the luxury to casually discuss my research. These discussions frequently led to critical thinking, new inspiration, and new insights. Their interest motivated me even more and in certain ways they contributed to this thesis, even though they may not realize this. I would like to thank in particular Alexander Hogenboom and Frederik Hogenboom, two friends who were willing to discuss obstacles I may have encountered and provide advice how to deal with them. They were also the ones who suggested me for this research project, and thus this thesis would not have been realized without them.

This research project is part of a greater and truly groundbreaking nationwide Dinalog project called 4C4More. I would like to thank all the people and organizations involved for facilitating a pleasant and informal environment at the plenary meetings, and access to a diverse business network. I really appreciate their great enthusiasm for sharing information, experiences and research results. I want to thank Jan van Dalen for introducing me to this project, inviting me to participate, and referring me to SCA.

I want to thank SCA for their hospitality and for giving me with the opportunity to conduct this research. In particular I want to thank Patrick Massuger, Riny Strik and Paul Keijzers. They provided a pleasant environment to work in and showed great interest and active involvement, all of which has been of great value to this thesis.

Last, but certainly not the least, I would like to thank my family and in particular my mom and dad who have provided me with wisdom, constant support and unconditional love. Starting from a young age, they have offered me more opportunities than one could wish for. When my peers' parents chose to impose decisions, my parents strived to provide full disclosure about any pros and cons and allowed me to make my own decision. They have supported my every decision, offered me advice, and were always prepared to think along, even when they didn't entirely agree with me. I am therefore very grateful for always having had this voice of reason.

Andrej Mondom

Rotterdam, July 2013

*"To spend too much time in studies is sloth; to use them too much for ornament, is affectation; to make judgment wholly by their rules, is the humor of a scholar." ~ Francis Bacon*

*"Where there is no vision, there is no hope." ~ George Washington Carver*

## **Abstract**

Consumer demand forecasting is a typical task deeply rooted in practically every supply chain and it is therefore conspicuous that little studies examine how businesses should incorporate that task in their information systems. The knowledge base mainly focuses on regular demand and little on special events like product promotions that disturb regular demand patterns. In current market conditions with increasing promotion intensity, promotional demand is becoming increasingly more important for manufacturers of Fast Moving Consumer Goods (FMCG) as most of their revenue is generated by promotional activities.

This thesis attempts to close to this gap for SCA by designing a tool that enables them to manage and analyze previous promotions in terms of how promotional attributes affect sales and to use that information to plan future promotions (the business problem). We accomplish this through a design science framework that guides our research process by first analyzing and structuring the business problem to gain more insight and to help us generate possible solutions. Data is collected at SCA and an appropriate regression model is selected that incorporates promotional attributes and sales for different products and retailers. A web-based tool is then implemented that is aligned with the business process and presents analysis results in a comprehensible fashion. Finally, we perform extensive evaluations to demonstrate utility, quality, and efficacy of the tool.

Unlike the current situation, the tool encourages building and maintaining a centralized database containing all promotions which enables easy data analysis. This is an important aspect because more (adequate) data should be gathered to make it possible to analyze promotions on a more detailed level like specific product variation. It is further recommended to gather data on additional promotional attributes and other relevant domains, like consumer behavior. More data is also likely to improve model performance.

# Table of Contents

<b>Acknowledgements</b>	<b>I</b>
<b>Abstract</b>	<b>III</b>
<b>1. Introduction</b>	<b>1</b>
1.1 Problem Statement	1
1.2 Purpose of the Study	2
1.3 Rationale of the Study	3
1.4 Scope	3
1.5 Thesis Structure	4
<b>2. Related Work</b>	<b>5</b>
2.1 Promotional Campaigns	5
2.2 Decomposing Promotional Sales Effects	7
2.3 Forecasting Promotional Sales Effects	12
2.4 Conclusion	17
<b>3. Methodology</b>	<b>18</b>
3.1 Problem Definition	18
3.2 Design Science	19
3.3 Design Science Framework	20
3.4 Evaluation Methods	22
3.5 Promotional Campaign Approach	24
3.6 Forecast Support Tool	25
<b>4. Study Design</b>	<b>26</b>
4.1 Problem Analysis	26
4.1.1 Analysis	26
4.1.2 Problem Statement	27
4.2 Problem Structure	28
4.3 Data Description	29
4.4 Data Analysis Model	31
4.4.1 Preliminary Analysis & Multicollinearity	32
4.4.2 Ridge Regression Analysis	33
4.5 System Architecture	35
4.5.1 Implementation Choice	35
4.5.2 User Interface	35
<b>5. Evaluation</b>	<b>40</b>
5.1 Testing	40
5.1.1 White-box testing	40
5.1.2 Black-box testing	41
5.2 Descriptive Evaluation	42
5.3 Model Evaluation	44
<b>6. Conclusion</b>	<b>48</b>
6.1 Conclusions	48
6.2 Research Limitations	50
6.3 Managerial Implications	51
6.4 Future research & Recommendations	52
<b>References</b>	<b>54</b>
<b>Appendix A</b>	<b>A1</b>
<b>Appendix B</b>	<b>B1</b>

## List of Figures

Figure 1: Overview Literature on Promotional Campaigns	6
Figure 2: Promotional Planning Structure	29
Figure 3: Retailer and Product Structures	30
Figure 4: Ridge Trace	34
Figure 5: Dashboard	36
Figure 6: Add Promotion Form	37
Figure 7: Regression Results Overview	38
Figure 8: Activity Diagram	39
Figure 9: Search & Filter Functionality	43
Figure 10: Change Log	43
Figure 11: Model Performance Brand A at Retailer I	46
Figure 12: Model Performance Brand C at Retailer I	46

## List of Tables

Table 1: Design Evaluation Methods (Hevner et al., 2004)	23
Table 2: Correlation Matrix for the predictor variables	32
Table 3: Coefficients of Determination	44
Table 4: Coefficients of Determination without competitive outliers	45
Table 5: Overview used software	A1
Table 6: Structural Testing Python methods	B1
Table 7: Structural Testing JavaScript Methods	B2
Table 8: Functional Testing	B4

## **1. Introduction**

Information systems are designed and implemented to help control the performance of business processes (O'Brien, 2007). In order to achieve this, it is required to align information and communication technology (ICT) with the business. Designing suitable systems is a challenge because people are the end-user and therefore involves human behavior, but it also involves technology which has to be adequately administered to support the business process. This research will aim to design, implement and evaluate a tool that is aligned with a specific business process and moreover has the objective to solve an identified organizational problem related to marketing. This chapter first introduces the business problem and then presents the purpose of this study. Next the importance of the study is illustrated, and finally the scope.

### *1.1 Problem Statement*

This research will examine the forecast process for sales promotion campaigns at the consumer goods manufacturer SCA Hygiene Products Nederland BV (from now on addressed as "SCA"). The objective of this business process is to accurately predict sales of future promotional events that are held at different retailers throughout The Netherlands. A promotion is a mean by which a company like SCA can promote its brands and attract new customers by principally offering price incentives and/or advertising. These events generally elicit a significant boost in sales numbers and are fundamentally different from baseline sales because they tend to occur erratically, are relatively infrequent and are further characterized by their various attributes (e.g. promotion type, prices, package size, advertising, and display type). Due to the large sale volumes during a promotion, it is appealing to find out what causes (variation in) the sales boost and try to forecast its magnitude.

Currently available software systems mostly focus on regular sales and offer limited to no support for these (special) events and little incentive even for the creation of a detailed centralized record of past special events and their results (Lee et al., 2007). This is also the case for SCA and has therefore led to current circumstances in which existing software does not offer a well established and structured method of supporting the forecasting process in an effective way. More precisely, the data currently possessed by SCA cannot be automatically retrieved, processed and used for marketing analyses. Due to the lack of these features, employees are compelled to rely on their own judgment when making decisions. This expresses itself through the following business problem: within the current environment it is difficult to assess the impact of promotions and to identify what causes the sales boost during promotional campaigns, or to identify the nature of the relationship between those causes and the promotional sales.

## *1.2 Purpose of the Study*

There are many Enterprise Resource Planning (ERP) systems on the market that, among others, incorporate management information for manufacturing, marketing, sales and service, and are integrated across an entire organization. The primary focus of this study is on designing and evaluating a tool that supports specific events in the field of marketing, in our case sales promotion campaigns. The study was conducted as part of an internship at SCA Hygiene Products Nederland BV, a subsidiary of the Swedish mass manufacturer of consumer goods SCA. SCA's business can be categorized into three groups: Personal Care, Tissue and Forest Products. The Dutch subsidiary 'Hygiene Products' covers Personal Care and Tissue for the Dutch market.

Promotions cause certain short-term and long-term effects due to change in consumer behavior and may provoke a response from competitors. The kind of change in consumer behavior, and hence promotional sales, depends on factors like the product category, product characteristics and promotion design. Historical realized promotional sales have shown relative large variations over time which cannot yet be clearly explained and result in large forecast errors.

SCA products are bulky but light in nature which means a relatively small forecast error will induce a surplus occupying a relatively large space. Therefore it is highly profitable to lower inventory costs due to large-scale overproduction and to lower the opportunity costs due to stock outs. Thus, understanding the dynamics involved with promotions and the effects thereof on sales and use that knowledge to decrease the forecast error is of high value. This is a typical task that is deeply rooted in practically every supply chain, and it is conspicuous that very little studies examine how businesses should incorporate it in their information systems. This study attempts to contribute to this gap in theory and practice by developing a forecast support tool. An important requirement for its design is that it is aligned with the existing business process and aids the forecaster in an effective and user-friendly way. In order to accomplish this, the problem solving design science research method for information systems will be applied.

The tool can analyze past promotions based on their attributes and their effect on realized sales. Individual retailers and specific product categories or brands can be analyzed to look for and make relationships between cause and effect explicit. The tool will make use of a centralized database with historic promotions and provide a web user interface developed with Django, an open source Python web framework. Existing libraries like SciPy, NumPy and SciKit will be employed which are open source implementations of algorithms and mathematical functions that can be used for analysis. In our case a causal model like regression analysis is appealing because it capable to discover relationships between explanatory variables and the dependent variable (Armstrong & Green, 2011).

### *1.3 Rationale of the Study*

This study is part of the 4C4More (4C: Cross Chain Control Center) research project which is an intensive co-operation between 6 companies, 3 logistics providers, 5 academic institutions and Dinalog (Dutch Institute for Advanced Logistics)<sup>1</sup>. Together they work to realize the goal of 4C4More which is to increase efficiency in Dutch logistics by combining knowledge, resources and better co-operation within same and different supply chains. SCA is one of the participants and is one of the manufacturers of FMCG in The Netherlands.

SCA and other FMCG manufacturers have experienced a significant increase in promotional intensity, which means there is an increase in promotional sales relative to baseline sales. Depending on the product category, this increase lies between 10 and 20% for the past 4 years. This study is relevant not just for SCA. One of the largest marketing research firms in the world, Gesellschaft für Konsumforschung (or in short "GfK"), reports an increase in promotional intensity in Dutch retail of 25.6% between the first halves of 2008 and 2009 (GfK, July 2010). In those same periods the share of promotional sales was respectively 11.3% and 13.6% of total sales, and had risen to 17% in 2011 (GfK, April 2012). The implication of this trend is that manufacturers of FMCG's, like SCA, are now more focused on promotions and improving the forecast accuracy thereof with ICT.

However, current information systems do not sufficiently support promotional sales and research conducted on designing such systems is rare. Therefore this research is required to develop a tool which provides a better understanding of the effects of promotions and how and what factors (e.g. marketing mix) affect the promotional sales. The latter is strategy that attempts to attract attention to a product and to change consumer behavior (e.g. loyalty, usage rate, stockpiling) and therefore fits the field of marketing.

### *1.4 Scope*

This research is conducted at SCA as the case company and will focus on hygiene products available on the Dutch retail market. Because of our specific area of interest, we only focus on promotions and not the regular sales. Data is gathered from internal available sources, but in some instances may originate from third parties. Due to limited nature and completeness of the data we are only able to take into account a limited amount of variables with regard to promotions which are mainly product characteristics. This is in contrast to literature that for instance also accounts for consumer behavior. Nevertheless, our tool provides a good basis to expand the database and add additional variables in the future. The goal is to study the effects

---

<sup>1</sup> List of 4C4More participants: [http://www.dinalog.nl/en/projects/r\\_d\\_projects/4c4more](http://www.dinalog.nl/en/projects/r_d_projects/4c4more)

on promotional sales, which is a dependent variable in our model. This is used to identify explanatory variables and type of relationship with promotional sales.

### *1.5 Thesis Structure*

This thesis begins by elaborating on work related to analyzing promotions and studying their effects (Chapter 2). This provides a wide perspective on and insight into this marketing-related topic and identifies possible causes that affect promotional sales. It also offers an overview of methods that have been previously used and that we could potentially utilize in this study. Chapter 3 presents our research methodology, which is the design science method and discuss the framework that we have utilized. Chapter 4 discusses our study design by first analyzing and structuring the business problem, followed by addressing the components of our forecast support tool, and finally the system architecture is presented. Chapter 5 is dedicated to the evaluation methods that we used to ensure quality and utility of the tool. Finally, the conclusion, limitations, managerial implications and recommendations are presented in Chapter 6.

## 2. Related Work

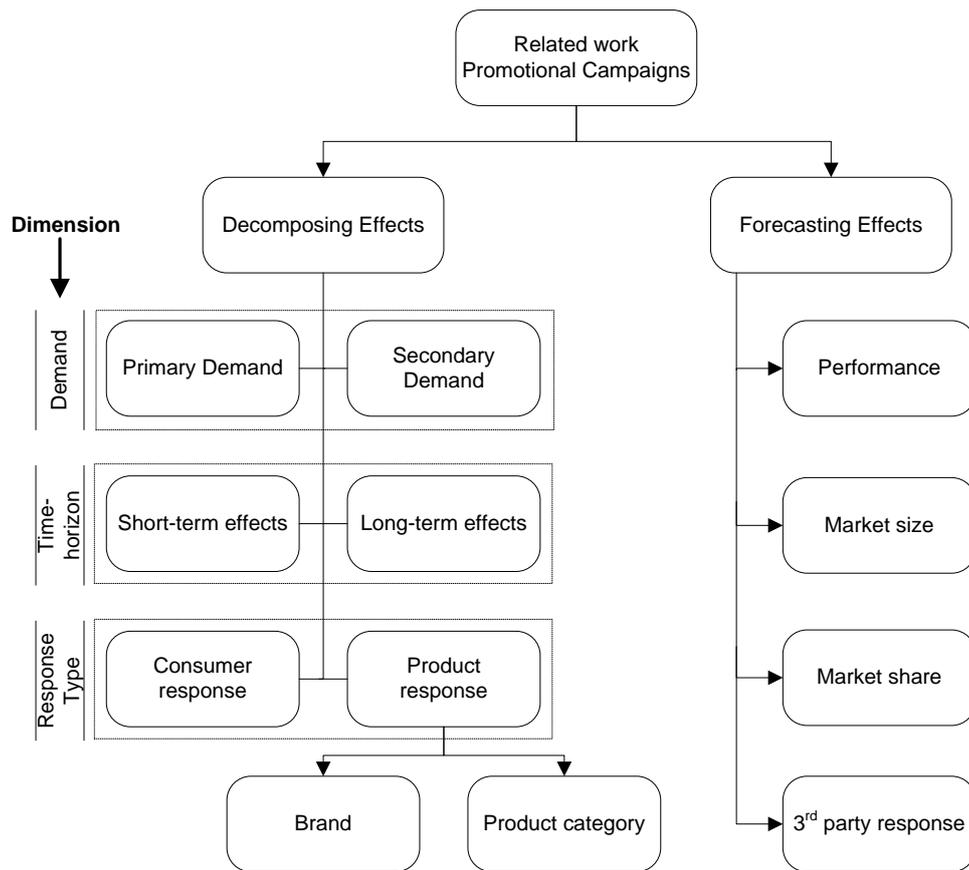
This chapter provides an overview of related work with regard to promotions. First a general overview is presented of the different types of studies on promotional campaigns. Then work will be discussed that examines the composition of promotional sales like how and what factors affect consumer behavior and therefore retail sales. This is followed by a discussion on work that focuses on methods that can be used in forecasting promotion sales. Finally we present our conclusion based on these discussions of related work.

### 2.1 Promotional Campaigns

Sales promotion campaigns, advertising, personal selling, public relations and direct marketing are the five aspects of a promotional mix (Harrell, 2010). This study addresses the sales promotion campaign aspect. In marketing, these are used extensively because they generate consumer response in the form of an increase in sales (Bell et al., 1999). The basic characteristics of a promotional campaign are that it is a departure from the norm, introducing additional factors in the market, altering perceptions, plans, and that it is not sustained for a great length of time (Pole et al., 1991). Different techniques can be deployed like price promotions, coupons, loyal reward program and rebates.

In the case of SCA mainly price promotions are organized that are accompanied by advertisement, but the latter is usually arranged by the retailer. Only in some rare instances advertisement-only campaigns are run. Price promotions are temporary price reductions offered to the consumer (Blattberg et al., 1995) and can affect the market in various ways. Figure 1 provides an overview of the different types of studies related to promotional campaigns. We divide the literature in two main parts: to the studies which examine the composition of promotion effects, and to those which examine how promotion effects can be forecast.

The literature on decomposing promotional effects is concentrated on three different dimensions: demand, time-horizon and response type. First, promotions affect demand in two ways: the sales increase could be due to consumers accelerating their purchases (primary demand) or brand switching (secondary demand). Second, a distinction can be made between the immediate effects of promotions (short-term) and extended effects after the promotion has ended (long-term). Third, short- and long-term demand effects are driven by (individual-level) consumer response (Guadagni and Little, 1983; Neslin et al., 1985) and product response (Bolton, 1989; Raju, 1992; Narasimhan et al., 1996). The latter can be further divided into an emphasis on specific brands (Bemmar & Mouchoux, 1991; Christen et al. 1997) and on product categories (Nijs et al. 2001; Poh et al., 1998). The decomposition is discussed in Section 2.2.



**Figure 1: Overview Literature on Promotional Campaigns**

Examining the composition of the promotional effect provides insight into the dynamics involved with promotions, and allows managers to draw better guidelines for policies, aligning campaigns and to anticipate the moves and countermoves of their rivals (Bell et al. 1999). The knowledge can then be utilized by forecasting the magnitude of different promotional effects which can be divided into four main elements: promotion performance, market size, market share, and third party response.

Performance is defined by the successfulness of the promotion in terms of sales and profitability. Market size can be measured in for example the total revenues, quantity sold or consumption in a given market. Market share is the portion of the market size controlled by a company, product or brand. Third party response is the reaction from competitors caused by the promotion. This research focuses only on the promotional sales performance of SCA, and this literature is discussed in Section 2.3.

The next section presents an overview of the work related to decomposing the promotional sales effects.

## 2.2 *Decomposing Promotional Sales Effects*

In his work, Gupta (1988) attempts to answer the question what is causing the sales “bump” during promotional campaigns by examining the impact of promotion on consumer decision of when, what, and how much to buy. The study investigates the impact of marketing variables on brand choice (what), interpurchase time (when) and purchase quantity (how much). The goal is to decompose the sales “bump” into sales increase due to brand switching, purchase time acceleration, and stockpiling, which implies investigating decisions at the household level. This knowledge can then be used to compare promotion strategies and help determine the most suitable and effective promotion. The brand choice is modeled using the multinomial logit model that is based on behavioral utility theory, allows explanatory variables, and accounts for competition. The interpurchase time is modeled using an Erlang-2 distribution and a set of explanatory variables. The purchase quantity is modeled using an ordered regression model. Each model utilizes its own set of variables, extracted from the collected data.

The data used by Gupta (1988) for model calibration and validation originates from IRI (Information Resources, Inc.) scanner panel data for regular caffeinated ground coffee and covers about 2000 households for a two-year period (1980-1982). The data contains variables like the brand and quantity bought, time and date, where bought, and store records provide information on prices and promotions for all coffee brands in all stores in Pittsfield, MA. The promotional instruments studied are Feature-and-Display (F-and-D), Feature-or-Display (F-or-D), and price cuts. The brand choice and interpurchase time models are estimated using the PAR program of the statistical package BMDP and the purchase quantity model is estimated using the LOGIST program of SAS. The results suggest that specifically promotion variables have the greatest influence on consumers’ brand choice behavior.

Further elasticity analysis is performed to decompose the total sales increase due to promotion, which results’ show that more than 84% of the increase can be accounted for by brand switching, 14% or less by purchase time acceleration, and less than 2% by stockpiling. Additionally, more than 98% of the sales increase due to price cut comes from brand switching, because unlike F-and-D and F-or-D, price cut does not affect consumers’ purchase time decisions. In this case stockpiling remains a negligible phenomenon for ground coffee. This can be explained because storing coffee for a long time may destroy its freshness, coffee cans come in large sizes, and the high promotion intensity in the marketplace ensures that some brand is almost always on promotion. Gupta (1988) further argues that, unlike coffee, stockpiling could be more important for other product categories like paper towels for which freshness is not important.

In their study, Bell et al. (1999) extend Gupta's research by further decomposing the effects of promotional campaigns to determine the extent to which the emphasis on primary and secondary effects varies across product categories. In contrast to previous studies which put the emphasis either on how promotions affect consumers, or brands, Bell et al. look at empirical regularities in promotional response across both brands and consumer behavior. Via meta-analysis, the study also tries to explain most of the variance found in primary and secondary demand elasticities. The data covers 173 brands across 13 product categories from which 519 price elasticities are generated and includes basket scanner panel data from 250 households.

The conceptual framework developed by Bell et al. (1999) captures the consumer's view of the characteristics of individual brands and product categories, and includes three sets of exogenous variables: (1) category factors affect the budget allocation process of buying a particular product category, (2) brand factors are about the brand's position in the market, and (3) consumer factors that include demographic characteristics of consumers. For the meta-analysis an iterated Generalized Least Squares method (Montgomery and Srinivasan, 1996) is used which allows controlling for heterogeneity in the meta-analysis model rather than in the underlying choice models. The advantage is that it accounts for an elasticity estimate that varies by brand and category.

Their study confirms that the main impact of a promotion is on secondary demand and with regard to the decomposition it shows that primary and secondary demand effects vary substantially across brands and categories and that storability is an important factor. The framework manages to explain up to 70% of the variance in promotional response, where category-specific factors explain most of the variance, followed by brand-specific factors. Consumer factors have relatively little influence. In some cases some variables do not affect total elasticity, but do affect individual components, which is especially the case for brand-specific factors.

Finally, promotional dynamics appear to vary systematically across categories, and this has to do with the effect on consumption. Some categories (e.g. bacon, potato chips, soft-drinks and yogurt) show an increase in consumption, while other categories (e.g. bathroom tissues, coffee, detergent, and paper towels) show stockpiling effects consistent with forward-buying. Wansink and Deshpandé (1994) show that stockpiling can greatly affect the rate of consumption if usage-related thoughts are salient: if products are perishable and the consumer is aware that the product is perishable and must be used, are more versatile in terms of potential usage, needs refrigeration, or are often in sight.

Nijs et al. (2001) underline that price promotions have different effects across different product categories and further make a distinction between category-demand effects in the short run and the long run. They also examine in what way marketing intensity, competitive reactivity, the competitive structure and covariates affect the category demand. These four dimensions are considered to have a moderating effect on category demand, unlike price promotions which are considered to be the main effect. The marketing intensity dimension knows two types: price promotion intensity and advertising intensity. The first consists of two components: promotional frequency which depicts how frequent consumers are exposed to price promotions and promotional depth which depicts the average size of the promotions. This distinction is made as Raju (1992) suggests the two components have a different effect on brand level.

The covariates contain product perishability, private-label share, the introduction of a new product, a deterministic-trend component and supermarket coverage. These variables have been chosen because previous research suggests they have an impact on price-promotional effectiveness. Data on e.g. volume sales, prices, feature and display activity was provided by IRI/Europanel and accounts for supermarket chains or clusters of chains. The BBC research agency provided advertising data. This covers a total of 560 frequently purchased consumer goods (FPCG) categories in The Netherlands over a 4-year period. FPCG is an alternative term for FMCG.

Using a series of unit-root tests it is determined that long-run category-demand effects of promotions are an exception and that mainly new product introductions revitalize markets. Any gradual demand effects seem to be caused by exogenous factors which are not part of brand-level marketing. Competitive reactions are uncommon which is in line with the stationary category-demand patterns, but if they occur their effects are mainly noticed in the short run. VARX models are used to estimate price-promotion elasticities of which the result shows a strong post-promotion cancellation effect for especially the tuna fish and toilet tissue categories. This is consistent with the results as reported by Van Heerde et al. (2000), who also report store switching for tuna or tissue is unlikely. Overall, persistent effects on the long-term are rare.

The results of second-stage regression analyses identify promotional frequency as the key driver of short-term price-promotion effectiveness, but their effect canceled out in the long run. Advertising intensity seems to affect both short- and long-run demand in a negative way as advertising emphasizes non-price purchase motivations and reduces price sensitivity (Mela et al., 1998). Also, stronger promotional effectiveness is perceived in perishable categories because of an increase in usage rate (Ailawadi and Neslin, 1998) and in markets with fewer industry members.

With regard to perishable and storable product categories, Pauwels et al. (2002) provide evidence that purchase quantity is more important than brand choice for storable products. The study examines the price-promotion effects on category incidence, brand choice, and purchase quantity for a storable product (canned soup) and a perishable product (yogurt). Quantity effects are short-lived for perishables because of consumer stockpiling limitations, while storable product purchase quantities are higher during a price promotion. Stockpiling can be caused because of cross-period effects where primary demand is borrowed from other time periods (van Heerde et al. 2004), which implies that it will take more time for a consumer to return to the store for storable product than for a perishable product. Another reason for stockpiling can be increased consumption.

An earlier work by Assunção and Meyer (1993) showed that product storability affects the rational consumers' decision on purchase quantity, which is also influenced by the promotional pattern (e.g. promotion frequency). In other words, consumers incorporate price expectations, based on promotional strategy, in their buying behavior. They anticipate price promotions and adjust their need of stockpiling. Further, stock pressure increases with bulky, storable products, and in markets characterized by lower promotion frequencies, but is also affected by usage rate. Higher stock pressure, which results in a higher inventory can increase consumption due to fewer stockouts (Gupta 1988, 1991; Neslin and Stone, 1996, p.89) and as mentioned earlier a faster usage rate (Wansink and Deshpandé, 1994).

Bemmaor and Mouchoux (1991) present an experiment that studies the price promotion effect on brand sales of nonperishable goods across different price discount values and levels of support. Using the Ordinary Least Squares method, the price deal elasticities turn out to be positive and large, with cat litter being the least responsive because of its bulkiness and weight. The effect of price promotions differs not only across product categories, but also strongly between brands.

The promotion effect is larger for cheaper brands and the leading brands are less sensitive to price deals, which suggests a relationship between market share and price deal. The latter has also been found by Blattberg and Wisniewski (1987). Exceptions are bathroom tissue and flour that appear to be more affected by brand switching. Nevertheless, advertising has a strong effect on sales during promotions. The study also suggests it may be useful for retailers and manufacturers to cluster products with similar price elasticity when planning promotions. Limitations are that the study does not account for dynamic effects of price promotions, the nature of the effects (primary/secondary) and different levels of advertising support.

A more recent study by Van Heerde et al. (2004) decomposes the price promotion effects in more detail and further argues that most of the decomposition studies fail to further divide primary demand into pure stockpiling and/or other effects such as consumption increases. This is why in their study they split and quantify primary demand effects into cross-period and category-expansion effects, which is also the primary goal of their study. Because the promotion effect is determined by the level of support (Lemon and Nowlis, 2002), they further analyze price promotions with different support conditions by using four separate predictor variables: without support, with feature-only support, with display-only support, and with a feature-and-display support. Additionally, they take into account the magnitude of the price discount. These variables are similar to Gupta (1988).

The data comes from two American (tuna, tissue) and two Dutch (shampoo, peanut butter weekly, store-level, and scanner datasets, provided by ACNielsen. Depending on the product, the datasets cover one to almost 3 years, 24 to 48 stores and include multiple brands of each product. Compared to the other categories, the tissue dataset (52 weeks) knows a relative small number of price promotions which in the end caused large standard-deviations.

Using the Ordinary Least Squares method the increase in sales for each brand is decomposed into three types of effects: cross-brand effect, cross-period effect and category-expansion effect. Then each effect is further decomposed for each combination of product and different support condition. Results show that the cross-period effect is much stronger for feature-only than for display-only, which implies that the first is more likely to induce stockpiling as the consumer is able to anticipate the price promotion and plan ahead. Displays have a strong impact on brand-choice and category-expansion. On average, all support conditions have an equal effect on the own-brand sales. Another interesting result is that compared to the other three products, peanut butter has the lowest average cross-period effect, implying it is less appealing for stockpiling. This can be explained because peanut butter is the only perishable product.

To show the dynamic effects of price promotions, the effects are further decomposed for different levels of price discounts using local polynomial regression. Except for display-only, with higher discounts the cross-brand and cross-period effects decrease. In general, the more support, the stronger the own-brand sales effect will be. Higher discounts result in a higher category-expansion effect, while the highest discounts have a much stronger absolute effect than lower discounts.

Most of these studies mainly focus on the short-term and immediate effect of promotions, which is comprehensible as most of the effect is noticeable during the period of a sales promotion. For example Nijs et al. (2001) and Pauwels et al. (2002) have demonstrated that for the products

they studied long-term effects are rare and usually canceled out. Nevertheless the studies do mention certain characteristics of sales promotion like promotion intensity may have long-term effects.

DelVecchio et al. (2006) therefore performed a meta-analysis of 51 empirical studies to examine what can increase or decrease preference for a brand after the promotion has ended. They conclude that on average sales promotions do not statistically affect brand preference, but under certain conditions there may be a positive or negative effect on preference. Significant results are that unannounced promotions have a strong negative effect on preference, while coupons or premiums showed a positive effect. Promotions larger than 20% and products with a lower number of competing products have a negative effect on post-promotion brand preference.

Besides consciously deciding on a promotion type, managers must also make sure to offer clear temporary promotions and make a tradeoff between the benefit of a sales boost of a large promotion and the long-term risk. Managers should also be aware of the nature of the market (competition) they operate in and the product type, because with fewer competitors and for durables the negative effect is larger. The latter occurs because promotions for frequently purchased packaged goods seem to be more accepted.

### *2.3 Forecasting Promotional Sales Effects*

Decomposing the promotional sales bump helps us identify and understand the most important factors that cause the increase in sales during a promotional campaign. The next step is to use this information in order to help forecast future promotional sales. Goodman & Moody (1970) try to answer the question how promotions can be made more profitable and how to predict promotional sales. They primarily focus on measuring the quantity which actually reaches the consumer, the sell-through quantity, and not the quantity that has been sold to distributors and/or retailers. This is because some distributors and retailers like to buildup inventories of a product at a reduced price, also known as forward buying. This inventory is then sold by the retailer for the regular price after the promotion, which results in a loss for the manufacturer. This is a reason for limiting the shipments to retailers and therefore requires an accurate forecast of the sell-through quantity. Also, stock-outs incur opportunity costs.

The sell-through quantity is assumed to be determined by a certain set of independent variables such as price, advertising expenditures, length of promotion, normal sales volume, and selling effort. The multiple regression method is used to determine coefficients for each variable. Using data from a large houseware item manufacturer the model was able to explain about 85% of the past variation. The standard error is used to define the forecast uncertainty.

Next an economic analysis is conducted in order to determine the optimal selling quantity to gain the highest profits from the promotion. In an ideal situation, a manufacturer wishes to ship a product quantity equal to the sell-through quantity. This prevents a marginal opportunity cost of shortage (stock-out) and that of inventory (forward buying). The model assumes a normal distribution of the estimation error of the sell-through quantity. In order to find the optimal sell-through quantity, the economic break-even probability has to be determined which is the probability of selling through the last unit shipped. This also requires the variable production cost. Once the probability has been calculated, the optimal sell-through can be found from a table of the standard normal distribution.

In another study, Reinmuth & Geurts (1972) view the phenomenon of promotional campaigns, price adjustments, and mergers as atypical behavior in a time series and argue that traditional modeling approaches are unfit for use because they ignore such unexpected situations. They develop a model based on a Bayesian approach to help analyze atypical situations of which the outcome is difficult to predict and therefore viewed as a chance event with an assigned probability distribution to possible outcomes.

An important prerequisite and aid for creating reliable probability distributions, a journal must be kept of atypical events and their effects. The journal holds records of different causes and their values for each event, and the outcome. By comparing the difference in events and their impact on the outcome, different situations can be assessed and a probability distribution can be estimated by the management. It can also be useful to collect sample information about the events by tracking the behavior of regular buyers. This information is incorporated using Bayes' Law to compute a probability distribution on proportionate change. Next, this continuous probability distribution is discretized with the grouped approximation approach to group different equally likely representative values together.

This is tested on sales data of a variety of perishable frozen food sold throughout the United States. The data is a time series with seasonal patterns and contains several price promotions as atypical situations. What is notable is that during the promotion there is an underforecast and the period after the promotion there is an overforecast. Using the probability distribution, the model gives a sales forecast much closer to the actual sales than the company's forecast. When only management assessment is used the model is likely to overforecast, while in the case of also using empirical information from buyers the model is likely to underforecast. This indicates management tends to be optimistic.

Pole et al. (1991) also take a Bayesian approach in the case of promotional campaigns in the context of forecasting retail sales of a soft drinks product (a FMCG). While usually campaigns last for a relative short period, the campaigns studied last up to six months, making them rather unusual. Promotions incur a change, departure from the norm, which needs to be critically examined and therefore a traditional model for routine forecast can't be used. They argue the process of forecasting promotional sales consists of adjusting routine sales forecasts in anticipation of the promotion and examining the promotion effects on sales. Objectives are preventing overstocking and stock-outs. The forecast methodology is to ground and improve forecasts in a continuous process based on the outcomes of previous similar promotions. This process should also make important explanatory variables explicit.

The model presented is an extension of the dynamic linear model (DLM) and Bayesian framework, specifically designed for analyzing promotional campaigns. It consists of a steady trend, regression on the price, and a seasonal component. For each promotion an additional explanatory component is added in order to allow effects on the time series can be isolated and used for forecasting. This splits the outcome into two parts: the expected outcome had there been no promotion and the part of the outcome incurred and affected by the promotion.

In the case of the soft drinks product, the change incurred by the promotion can be predicted by examining previous similar campaigns and compare them on a certain set of independent variables. The model breaks down information and enables a clear estimation of individual effects and allows a continuous forecast process as new data becomes available. Because the duration of promotions is six months, actual sales from the first month can be recorded and used to update estimates for the following months. Nevertheless, this should be done with caution as for example reducing the supply or shortening the promotion too early will lead to relative large forecast errors. Also, there may be a certain sales pattern because of the relative long duration of the promotion.

Another study (Caruana, 2001) which examines carbonated soft drink brand sales uses a similar approach by utilizing seasonal regression and also dividing the time series observation into four main elements: trend, seasonality, cycle effect and residual error. Time series and causal modeling are identified as the two main forecasting techniques. Time series are primarily used to find patterns over time and extrapolate them into the future. This is only useful if what has happened in the past will continue to happen in the future. When there are other variables than just time like sales promotions and advertising, the causal model is more appropriate to determine the cause and effect relationship.

The study, however, was not able to identify the required leading indicators in the 4-year period data in order to use a causal technique, and therefore uses the seasonal regression technique. In the additive model, a twelve-month period is used for seasonality and a one-month period is used for the trend. The data also contains outliers, caused by sales promotion campaigns from competitors, and were replaced using linear interpolation. The model is able to predict over 90 percent of the variation, but the residuals are not normally distributed over time. This is because in cases of high predicted sales the variance is higher than in cases of low predicted sales. This indicates heteroscedasticity (Robie & Ryan, 1999) and is the reason why the weighted least squares technique is applied by Caruana (2001), giving more weight to the more precise observations and less weight to the highly variable observations. Results show that the trend and seasonality would be overestimated had the ordinary regression method been used.

An alternative approach to the one in the previous discussed research is presented by Shao (1997) that applies multiple intervention analysis to sales data. Unlike traditional time-series intervention analysis, the case of promotion campaigns includes multiple (many) interventions and uncertainty about those interventions. An additional complicating factor for forecasting sales is that a future promotion may not always be identical in composition to previous promotions. Hence, the objective of the proposed model is to be able to cope with that.

The data covers almost 4 years and has a total of 18 observations (interventions) and a total of 15 different promotional strategies. Although this is a time series, the data is difficult to analyze because it's very dynamic with many spikes in the data which are also mutually different. The company had been applying human judgment for forecasting future sales which caused large forecast errors because the promotions tend to change a lot each time. Different forecasters would also estimate significant different forecasts.

The intervention model assumes the sales consist of a constant, a response function and a noise term. The problem with estimating the parameters in this case is the large number of parameters (15) which is why this number should be reduced by clustering similar interventions. A major advantage is that this increases the number of similar interventions that can be compared to a future intervention, which in turn can be assigned to an existing cluster. After clustering there are now five groups and the relationship between the different promotion characteristics and sales can now be easier estimated.

After evaluating the results of the model, what stands out is that both the company and the model over- or underforecast at the same time. The only difference is that the model's estimate usually has a significant lower forecast error. In one instance the company's estimate did

perform better than the model, which could imply there are still certain unaddressed factors or a new promotion strategy was deployed.

Leone (1987) also applies intervention analysis and uses it in the context of price deals and the impact they have on promotional sales. Some advantages of intervention analysis is that it doesn't assume independent observation data and constant variance in that data. Besides looking at quantitative changes, the study also looks at qualitative changes such as an introduction of a new promotion strategy. In both cases, the goal is to examine the market dynamics and both provide a description of the effect of the changes, and forecast based on those effects. With regard to the descriptive perspective, this could be short- and long-term promotion effects.

The interventions in the context of a price deals are viewed as pulses that immediately, though temporarily, increase sales. The presented model assumes that sales levels after the promotion will immediately drop below the mean sales level due to a stockpiling effect, but afterwards slowly stabilize to an equilibrium level. Daily scanner data of two competing toothpaste brand sales are used to gather enough data points. The data includes changes in advertising price deals. On the long-term, the model is able to effectively capture the causes (advertisements) and effects on sales and movements in market share of the two brands. As for the short-term effects, the results confirm the previous assumption about the stockpiling effect.

McIntyre & Achabal (1993) state that the historical data on past promotional campaigns suffer from a limited number of and erratic observations and different set of promotional attributes making each promotion more or less a unique event. Due to these limitations most retailers rely on solely judgmental and subjective forecasts, which often results in large forecast errors of typically more than 25%, and not uncommon are errors of more than 50%. Also, because forecasters base their judgment on experience rather than on an explicit approach based on well-established rules, when a forecaster leaves the organization so does his expertise.

In their study they therefore try to capture that experience in a Case-Based Reasoning (CBR) system by the use of analogies. The purpose of such a system is to help forecast sales for a future promotion by extracting knowledge from past similar promotions. Each promotion is defined by certain attributes like for example price, type of advertisement, the participating stores and number of ads. For each new planned promotion analogies are selected based on the similarity of first the product and then the rest of the promotion attributes. Additionally, based on an adjustment table containing generalizations of attribute values in categories, more analogies are selected that are not identical but similar to the planned case.

Through regression the coefficients for each attribute can be re-estimated as new promotions are added. The results indicate an improvement of forecast accuracy for one product category, but not the other. This is mainly caused by the limited data set and the large variation in past promotions characteristics. As the amount of historical data increases and with the number of potential usable analogies, so will the forecast accuracy. Overall, the system creates structure, is able to capture and share expertise, makes the judgmental process more explicit for everybody, and saves time. These advantages also make coordination between different department easier and more efficient.

Lee et al. (2007) conducted an experiment on this topic of using analogies and the effect is has on forecast accuracy. They offer the forecaster three different levels of support. Level 1 only provides memory support for a number of recent promotions for a given product. Level 2 adds similarity support by providing a number of most similar promotions from the database, based on the attributes. Level 3 adds adaptation support that calculates the difference in effect as the mean ratio for all pairs of promotions that only differ on one attribute.

Although being a laboratory experiment with simulated promotion effects, real world data was used of promotional sales of 12 products. The results show that Level 1 and 2 support do not significantly improve forecast accuracy. Level 3 support was significantly more accurate than the two other levels, but only under the conditions of low noise. Since it is not a standard facility in most current software systems, the study recommends creating a centralized and structured record of past promotional campaigns and their results.

#### *2.4 Conclusion*

We can observe that there exists a substantial knowledge base on promotional campaigns and that different studies focus on different marketing dimensions and elements. The choices therein mainly depend on what data is available and the research goal. Although some studies focus on developing models to forecast promotional sales, promotions have mostly been studied in terms of decomposing the sales "bump". Considering the fact that demand forecasting is a typical task that is deeply rooted in practically every supply chain, it is conspicuous that very little studies examine how businesses should incorporate those models in the design of their information systems. We therefore attempt to contribute to this gap in knowledge by selecting a research framework in the next chapter that enables us to design an appropriate tool that will incorporate such a model for our business problem.

### **3. Methodology**

The goal of this research is to design a prototype tool that aids forecasters in their judgment of sales promotion campaigns at SCA. It will allow for better management of the promotional campaigns and help gain a better understanding of the relationship between explanatory variables and the promotional sales. Data and sufficient understanding about the current forecast business process will be acquired through an internship at the company. All data is proprietary and therefore no real company or product names will be mentioned.

#### *3.1 Problem Definition*

SCA is a company that produces FMCG's which are sold at retailers throughout The Netherlands. The products range from kitchen paper and toilet tissue to female hygiene. These are generally light and bulky products. Most of the revenue is made through price promotions, or in other words by attracting customers through discounts. Such promotions attract consumers and results in much higher sales than during a non-promotion period, increasing sales by a couple of hundred percent – not unusual is a thousand percent or more. Unlike baseline sales, the promotional sales are known to fluctuate a lot from promotion to promotion, which causes problems for SCA: how do we know how much volume to order at our factories?

Because of the nature of the products (bulky, mostly air), a small overforecast results in a relative large extra volumes which incur high inventory and transportation cost, while an underforecast results in stockouts. This is why SCA wishes to predict the sales for a promotion as accurately as possible. Currently forecasting relies on (expert) judgment and in agreement with Customer Service, Customer Marketing, Account Managers and the retailer. There two types of retailers: those that place an order for a promotion and ships back the unsold goods, and those that don't have the liberty of doing this. This provokes a different strategy when a specific retailer places orders: the first is likely to order too much, while the second is more likely to be cautious and order just enough or too little.

SCA manufactures a broad range of products which also come into different packaging sizes and product variants. Its customers – the retailers – differ in size and public reach. Each promotion is carefully planned and designed together with the retailer during negotiations by agreeing on promotion variables like the product type, discount, the packaging size, the duration of the promotion campaign and number of selling locations. These and other variables can affect performance in terms of promotional sales, but these relationship have yet to be determined. This is why a tool is needed that can make those relationships and their nature explicit.

The tool allows the user to create a new planned promotion and then offers the opportunity to analyze similar previous promotions based on for example the selected retailer and product. The historical data is extracted from the different information systems at SCA. The analysis results are presented using graphs for which the user has the option to adjust variable values and immediately witness the effects. After a planned promotion has passed and the realized sales have been entered, the promotion is recorded with the rest of the historical data and used in future analyses. These are the most important features to be included in the tool, which should furthermore be user-friendly and aligned with the current business process.

### *3.2 Design Science*

For our study we will apply a problem solving research method which is known as Design Science for Information Systems (IS). IS research has a lot of focus on the topic of aligning IT with the business goals. An information system collects, processes, analyzes, and disseminates information for a specific purpose (Rainer & Cegielski, 2012). If implemented correctly, an information system is able to effectively and efficiently support business processes within an organization. Many factors contribute to which extent this is actually realized.

Design science theory informs researchers and practitioners of the interactions among people, technology, and organizations that must be managed if an information system is to achieve its stated purpose, namely improving the effectiveness and efficiency of an organization. Therefore, much of the work performed by IS practitioners, and managers in general (Boland, 2002), deals with design. Design is both a process (set of activities) and a product (artifact) (Walls et al., 1992). It describes the world as acted upon (processes) and the world as sensed (artifacts).

IS research focuses on IT artifacts which are concrete prescriptions that enable IT researchers and practitioners to understand and address the problems inherent in developing and successfully implementing information systems within organizations (March and Smith, 1995; Nunamaker et al., 1991a). Design science on the other hand, creates and evaluates IT artifacts intended to address and solve identified organizational problems, and hence is a problem solving paradigm (Hevner, 2004). An important prerequisite to accomplish that is to develop a design that incorporates all the relevant aspects that together are able to solve organizational problems that have not been solved before.

In order to accomplish that, it is required to collect information from the interaction of people, organizations, and technology and have that qualitatively assessed to yield an understanding of the phenomena adequate for theory development or problem solving (Klein and Myers, 1999). By the process of constructing and exercising, innovative IT artifacts enable design science

researchers to understand the problem addressed by the artifact and the feasibility of their approach to its solution (Nunamaker et al., 1991a).

The design process starts by modeling the problem in order to gain a better understanding, to define the design problem and its solution space (Simon, 1996). Hevner et al. (2004) argue that the problem model also enables the design researcher to explore the effects of design decisions and changes in the real world. This will aid to develop better ideas and solutions for the problem when it's applied in the real world. This is done by acquiring feedback on the quality of the artifact through evaluation. Then the researcher uses the improved insight, new ideas and solutions to further improve the artifact. This step is a process of iteration of which the outcome is used by the design researcher to further improve the artifact until the best possible level of quality and effectiveness has been reached (Hevner et al., 2004; Markus et al., 2002).

For this research we go through a sequence of expert activities (processes) and use it to build a tool (product) that addresses the identified business problem. We begin by exploring the problem space (environment) and the existing knowledge base. The problem space consists of the people, organization and existing technology (Silver et al., 1995) and by examining that we gain insight in the business needs which consists of goals, tasks, problems, and opportunities. We do this by visiting the company, interviewing people from different relevant departments and gathering data. By reviewing literature related to sales promotion campaigns we gather ideas for modeling our business problem and building a first tool version.

The next step is evaluation to test the utility, quality and efficacy of the product by using different evaluation methods. This step provides us with new information in the form of a better understanding of the problem and new ideas that can be used to further improve the artifact. Evaluation, together with constructing the artifact, is iterative and may be repeated multiple times until we arrive at a satisfactory end result.

### *3.3 Design Science Framework*

As distinct from research that defines research questions or hypothesis and respectively tries to answer or prove those, this research will produce a viable artifact. Hevner et al. (2004) provided the design science research community with guidelines for design science in information systems research that assist researchers, reviewers, editors and readers to understand the requirements for effective design science research. Other research frameworks exist, like the two-dimensional research framework by March & Smith (1995) which defines research outputs as constructs, model, method and instantiation and defines research activities as build, evaluate, theorize and justify.

Purao (2002) and Rossi & Sein (2003) suggest an additional research output: better theories. Then there are also Walls, Widmeyer & El Sawy (1992) and Gregor & Jones (2007) that use the term design theory instead of design framework, and the latter propose eight components of design theory. There are also arguments on how design science research seems to be very similar to Action Research in terms of paradigmatic assumptions (Burstein and Gregor, 1999; Cole et al., 2005; Järvinen, 2007), while others argue differently (e.g. Iivari & Venable, 2009). In short, within the IS research community there is still no consensus on one widely accepted design framework.

This research will apply the design science research framework as proposed by Hevner et al. (2004) as it is considerably more cited and respected in the design science field. Solving a relevant problem, making a significant contribution, designing an effective artifact, and evaluating this artifact, are some general categories the total of seven guidelines address. We next discuss these in more detail and how they apply to this research.

The selected framework for this research first states design science research must produce a viable artifact in the form of a construct, model, method, or instantiation. In this case an instantiation will be produced in the form of a tool that deals with an identified business problem. This fulfills the second guideline which requires design science research to develop technology-based solutions to important and relevant business problems. The third guideline demands the artifact to be evaluated in order to test the utility, quality and efficacy thereof. Hevner et al. (2004) present five different methodologies that can be used for design evaluation, which will be discussed in the next section.

The fourth guideline states that in order for design science research to be effective it must contribute to the relevant field. The fifth guideline mentions the importance of applying rigorous methods in both the construction and evaluation of the design artifact. The sixth guideline characterizes design as a search process in which all available means are utilized in order to reach desired ends while satisfying (typically uncontrollable) laws in the problem environment. The seventh and final guideline asks the researcher to present the work to both technology-oriented (researchers) and management-oriented (practitioners) audiences. This way the researchers gain knowledge about the development process of the design artifact, while practitioners gain insight into the application and problem-solving capabilities of the design artifact.

### 3.4 Evaluation Methods

Design evaluation has been mentioned in the previous section as the fourth guideline of the design science framework as proposed by Hevner et al. (2004). This guideline depicts different evaluation methods that can be used to improve and demonstrate the utility, quality and efficacy of a design artifact. The evaluation of the artifact is based on requirements which are dictated by the business environment and its technical infrastructure it consists out of. That's why evaluation comprises out of the artifact's integration within that technical infrastructure.

When evaluating an IT artifact, appropriate metrics need to be defined and the appropriate data needs to be gathered and analyzed. The evaluation process looks at quality attributes such as functionality, completeness, usability and fit with the organization. Evaluation also provides the researcher with feedback to the construction phase, the quality of the design process, and the design product, which fits the iterative nature of design. A design artifact is complete when it satisfies the requirements and constraints of the problem it was meant to solve. What starts out as simplified conceptualizations and representations of problems, may develop into an end-result in which certain prior assumptions have become invalid due to changes in available technology and/or organizational environment. Bearing these aspects in mind during the research process is therefore important (Hevner et al., 2004).

Selecting evaluation methods requires a match with design artifact and the evaluation metrics, as certain methods may not always be suitable. Selecting the right evaluation methods is crucial to demonstrate the goodness and efficacy of the artifact (Basili, 1996; Kleindorfer et al., 1998; Zelkowitz and Wallace, 1998). Table 1 provides a review of five evaluation methods as presented by Hevner et al. (2004) and the last column describes if and how they would be suitable for this research.

Because we view evaluation as a process, there are multiple moments in the development process where we subject our work to an evaluation. Our artifact is studied within the SCA business environment during the development process, while a field study is infeasible because we have only access to historical data and do not deploy the tool at this stage. Analytical method types are not relevant at this point, but would become useful once the tool is supposed to be deployed and for example performance, usability and accessibility become important aspects.

Simulations are not suitable because we require real data and artificial data is not useful. Controlled experiments are also continuously conducted during the development process as part of testing. We further use functional and structural testing with every iteration and finally construct a scenario to demonstrate utility, supported by informed arguments where applicable. The evaluation is addressed in more detail in chapter 5.

**Table 1: Design Evaluation Methods (Hevner et al., 2004)**

<b>Method</b>	<b>Subtype</b>	<b>Suitability</b>
1. Observational	Case Study: Study artifact in depth in business environment	Included in the building process through evaluation as a process.
	Field Study: Monitor use of artifact in multiple projects	Only historical data is used from previous promotions for testing, thus infeasible.
2. Analytical	Static Analysis: Examine structure of artifact for static qualities (e.g. complexity)	Mostly existing software packages are used, a relative small part of code may be subjected to this method.
	Architecture Analysis: Study fit of artifact into technical IS architecture	N/A at this point; focus is on stand-alone tool. Could be interesting once other parties start using it.
	Optimization: Demonstrate inherent optimal properties of artifact of provide optimality bound on artifact behavior	N/A There are no other existing tools available for this specific business problem to compare with.
	Dynamic Analysis: Study artifact in use for dynamic qualities (e.g. performance)	N/A not relevant at this point, not related to development method and requires benchmarking suites.
3. Experimental	Controlled Experiment: Study artifact in controlled environment for qualities (e.g. usability)	Also already included through evaluation as a process to inspect usability within the company.
	Simulation: Execute artifact with artificial data	Data may not represent reality; not related to development method (Zelkowitz and Wallace, 1998).
4. Testing	Functional (Black Box) Testing: Execute artifact interface to discover failures and identify defects	Feasible and required for feedback on any faults in what the tool is supposed to do.
	Structural (White Box) Testing: Perform coverage testing of some metric (e.g. execution paths) in the artifact implementation	Feasible and useful for exercising different paths and determining appropriate output, uncover faults.
5. Descriptive	Informed Argument: Use of information from the knowledge base (e.g. relevant research) to build a convincing argument for the artifact's utility	Feasible, provides link to prior research, inexpensive, but risk of selection bias and treatments may differ (Zelkowitz and Wallace, 1998).
	Scenarios: Construct detailed scenarios around the artifact to demonstrate its utility	Feasible and useful to demonstrate utility.

### 3.5 Promotional Campaign Approach

The related work discussed in Chapter 2 presents multiple methods and approaches that can be considered for identifying and examining the relationship between explanatory variables and promotional sales. Figure 1 depicts how literature that addresses the effects caused by promotional campaigns can be divided into decomposing promotional sales effects and forecasting promotional sales effects. Our study will mainly focus on decomposing the promotion effect and use that to aid the forecaster. First we define the decomposition dimensions and forecasting elements that are appropriate for our business problem.

On the demand dimension we do not have sufficient market data for studying the secondary demand effects like brand switching. We therefore focus on primary demand effects in terms of sales increase. We do not have the required data to make a distinction between sales increase due to forward-buying (stockpiling) or increased consumption (usage-rate). However the literature suggests that the SCA products (personal care & bathroom tissue) belong to a category which generally does not show consumption increases, but rather exhibits stockpiling effects (Bell et al., 1999), because they are easy to store and nonperishable.

Only short-term effects are considered because the largest variations are found in the immediate effects of promotions while baseline sales remain relatively smooth and constant. Nijs et al. (2001) and Van Heerde et al. (2000) also demonstrate a strong post-promotion cancelation effect for toilet tissue categories. As for the response type, we have no market- or household-level data to investigate consumer behavior. Instead, we focus on the product response by looking at how individual brands respond to promotions in terms of sell-through quantity which is also the forecasting element we address as promotion performance. Insufficient data is available to study market size and share, and third-party response.

The literature mainly uses regression or time series methods to decompose the promotional sales effect. Our objective is to identify explanatory variables and make their relationships to the promotional sales explicit. The variables can be controlled by SCA, and we are also interested in (estimating) the effect size. According to Armstrong & Green (2005), a causal model like multiple regression would be useful in this situation. The advantage is that it can incorporate more information about the time series than its past observations (Alon et al., 2001).

Finally, we provide the forecaster with an assessment of uncertainty by adding a confidence interval rather than a test of statistical significance, because prior research has indicated the latter is of no value for analyzing data (Schmidt & Hunter, 1997; Ziliak & McCloskey, 2008) and has made no contribution to forecasting (Armstrong, 2007). Using hold-out data as a benchmark, we can assess the predictive validity of the model (Armstrong, 2001).

Regarding software, different forms of regression analysis are widely available in different existing proprietary and open source software packages, both in the form of bespoke and off-the-shelf solutions. We have chosen to implement our tool as a web application using the free open source Django web framework and will therefore use existing Python libraries that offer a wide range of regression analysis methods.

### 3.6 Forecast Support Tool

Our research is the result and a sub-project of 4C4More, with whom SCA is affiliated and has indicated the need for a demand forecast support solution. After examining this need through preliminary interviews on the (currently unsupported) business process, we have decided in agreement with different department managers to design a forecast support tool. This is a form of Forecast Support System (FSS) and a particular type of a decision support system (DSS) (Fildes et al., 2006) where the manager's judgment *plus* the model embedded in the system can provide a more effective solution than either alone (Keen & Scott Morton, 1978).

The forecast support tool is meant to be used to support forecasters in their decision-making and not to make the decisions for them. An example is provided by Fildes et al. (2006): *"statistical models are adept at finding patterns in large volumes of data, while managers can take into account the effects of special events like future product promotions. Such special factors are often treated as noise by statistical models, because the reliable estimation of their effects is precluded by a paucity of data."*

The tool's main goal is to examine special events (promotions) and to identify and explain the relationship between explanatory variables and promotional sales. To make this information explicit, an overview of variables is presented that affect promotional sales and the nature of that relationship. This causal information can then be interpreted by the forecaster and included in his judgment when planning a new promotions. Prior research has demonstrated this practice improves forecast accuracy (Lim & O'Connor, 1996; Lee et al., 2007). Because the forecaster is able to adjust variable values and immediately witness the effects, he can examine the dynamics of the promotion.

The downside of regression analysis is that the model is informed or supported by the extent to which they fit past observations and they may be inapplicable to conditions that will apply in the future (Goodwin & Wright, 2010). Further, promotions are unique in nature and demand special attention as not all information can always be quantified and used in a statistical model (Sanders & Ritzman, 2001). This is why a forecast support tool should be able to record judgment involved with each promotion (Flores et al, 1992; Webby et al., 2005). In our case this would be relevant but no such data exists for previous promotions.

## 4. Study Design

This chapter begins by analyzing the problem area (Section 4.1.1) and defining a problem statement as posed by SCA (Section 4.1.2). In our case, information should be provided on the relationship between the different promotion attributes and sales based on historical data. This can help SCA make better informed and conscious decisions when planning future promotions and supply. We continue by structuring and modeling this problem in order to help us understand the problem better. We conducted semi-structured interviews with five managers at the start of the research process and during the research we also interviewed employees of the involved departments and performed observations to gain insight in the business operations. Adding structure in these operations is useful for exploring the solution space and enabled us to propose a set of design principles that can help solve the business problem (Section 4.2).

We have chosen to address the problem by developing a forecast support tool. Because this can be classified as an FSS which is a particular type of a DSS (Fildes et al., 2006), we use the DSS development framework proposed by Sprague (1980). This framework divides a DSS in three components: a database, a model base, and a user interface for linking the user to each of the previous two. We therefore begin by describing the data that is at our disposal (Section 4.3), followed by a data analysis model to address the model base component (Section 4.4). We have selected a model that has been widely used in similar marketing problems, and therefore has proven itself to be effective. Finally, we address the user interface and the system architecture of our forecast support tool to show how it could be implemented (Section 4.5).

### 4.1 Problem Analysis

The purpose of this section is to address the business problem by defining the problem statement after having analyzed the problem area.

#### 4.1.1 Analysis

SCA is in the business of selling FMCG to different retailers, who in turn sell the goods to consumers. As is generally the case for FMCG, the goods manufactured by SCA have a low profit-margin, but a high stock-turnover. That's why a substantial amount of earnings is realized through sales promotion campaigns which are organized in agreement with the retailer. During these promotions products are sold in much higher volumes than is the case with regular sales, and despite the significant lower prices they constitute substantial cumulative profits. Because of these sales peaks, and limited production and storage capacity, SCA is required to forecast sales in order to prevent (significant) over- or underproduction.

Forecasts are made both for baseline sales as for promotional sales, but the latter is far more important as it's an important mean to generate sales, make profits, but also to create brand awareness and attract new customers. Promotions are characterized as irregular events in terms of time interval sizes between two promotions and are rather unique in nature because its attributes may vary over time. The attributes basically shape a promotion by defining for instance the price, promotion type (e.g. buy one get one), availability (e.g. limited stores) and form of advertising (e.g. folder). Every campaign is carefully planned in advance and its attributes are determined through a negotiation process with the retailer.

It is a challenge to predict sales for each promotion because every negotiation process may yield a different set of attribute value combinations, and furthermore there are different retailers to organize a promotion with. Experience has shown that promotional sales tend to vary substantially over time and that promotional attributes are responsible. However, it is a problem when it comes to discovering the relationship between different attributes and their realized sales, because firstly there's no central repository of all historic promotions that can provide a structured overview of previous promotions. This has resulted in a situation in which forecasters mainly have to rely on their own memory and experience. Secondly, this shortcoming prevents for instance statistical analysis of historic data in order to acquire insight in what has affected promotional sales and in what way. Identifying these relationships is therefore based on human judgment.

#### *4.1.2 Problem Statement*

The fact that decisions are based on human memory and experience entails certain implications. First of all, the human's limited cognitive ability means that only a small sample of past cases may be recalled (Lee et al., 2007) and because promotions are rather unique and different from each other it is more likely specific details will be forgotten (Schank, 1982). Secondly, experience is intangible and difficult to transmit to between individual. If a person is not available or leaves the organization, so does his experience. Thirdly, even if a central repository were available for memory-support, it is a difficult and time-consuming task to analyze all the information in order to identify relationships in the data. This leads to the following problem statement, as posed by SCA:

- Is it possible to design a forecast support tool that can manage promotions and their attributes, and provide forecasters with insight on the nature of the relationship between different promotional attributes and promotional sales?

## 4.2 Problem Structure

In order to gain insight into the business problem, we conducted semi-structured interviews with five managers at the start of the research process. During research, other involved employees were also interviewed and observations were performed to understand the business operations. Using this knowledge, the addressed problem is structured and modeled to help develop solutions. Below, the background of the operations is briefly introduced, followed by an explanation of the business operations, and a set of design principles is presented.

Promotions are recorded in contractual agreements, e.g. in terms of the number of promotions per year, and are planned ahead and in agreement with the retailer. This can be a large retailer that operates in whole of The Netherlands, or a smaller, regional or local store that operates in only a part of the country. A large retailer has more capacity and flexibility in its logistic network and sells more during a promotion. A group of smaller regional retailers has established a strategic partnership which allows them to order large volumes together and take advantage of a large volume discount. However, large volumes require an adequate lead-time for the factories and therefore a strict deadline for fixing a promotion, which is set at six weeks prior delivery. All the promotion characteristics should be known and fixed before this deadline.

- ✓ Design principle 1: The tool should track planned promotions and notify the user about approaching deadlines.

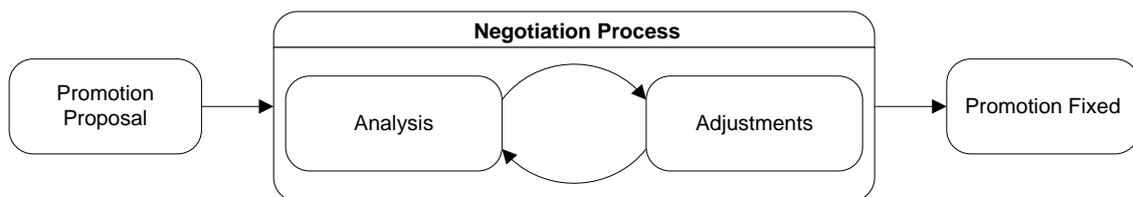
The promotion characteristics are determined through a process of negotiations between the retailer as an external party, and the internal departments of customer marketing, customer service, account manager and the factories. Each party analyzes the proposal based on their own interests and determine whether it's appropriate and feasible. When a retailer requests a promotion at the designated account manager, it is registered as a prospective promotion. Initially, the most important elements are the delivery date and product, after which the promotion is designed in greater detail by determining other characteristics like price, packaging size, position and advertising type. Throughout the negotiation process those values are subject to change, but should eventually be fixed on time for the deadline. Finally, last but not least, a sales forecast is determined, also in agreement with the retailer.

- ✓ Design principle 2: The tool should allow the user to add a new planned promotion to a central database, in a structured way, and allow adjustments and changes to be made to take into account the negotiation process.

The process of determining the sales forecast is currently mainly influenced by the large retailers that have a good bargaining position, and SLA's that require on-time delivery of the requested product volume. This has resulted in structural overproduction. SCA has insufficient insight in what effects promotional demand, but could use that understanding to gain a better bargaining position if it could argue that, and justify why the retailer is too optimistic. This can be realized by comparing a planned promotion to similar previous promotions and provide insight into the factors that affect demand and their relationship to demand.

- ✓ Design principle 3: The tool should analyze previous promotions and provide an accessible and comprehensible overview of identified explanatory variables and how they relate to demand.

To summarize, Figure 2 depicts the promotional planning structure, beginning with the promotion proposal, followed by the negotiation process, and concluding with fixing the promotion characteristics.



**Figure 2: Promotional Planning Structure**

In order to implement the analysis part of the tool we need data on previous promotions, which is discussed in the next section.

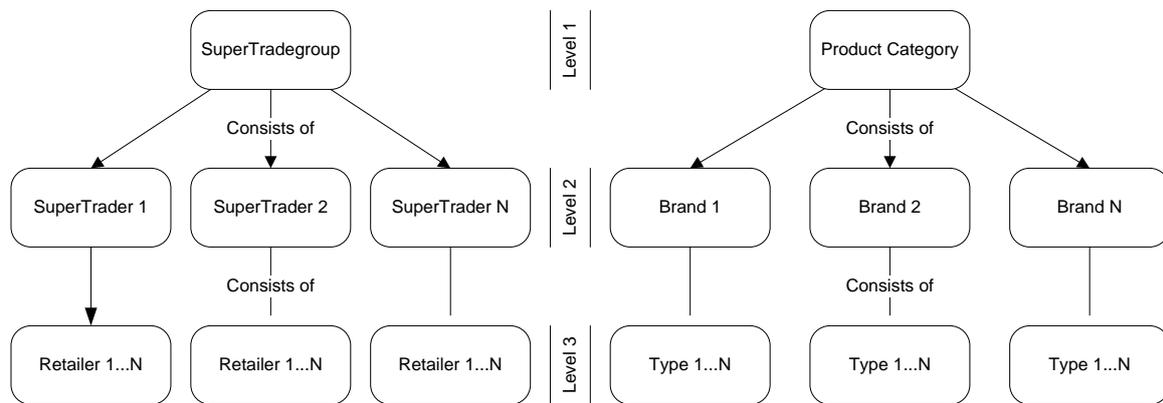
#### 4.3 Data Description

Data were available on three major SCA FMCG brands in the Tissue and Personal Care product categories. These brands consist of sub-brands and sold by retailers in The Netherlands. The data we used originated from four different sources, two of which are provided by a third party marketing research companies. First, PI Web is a business intelligence tool from Publi Info which contains information about historical sales promotion campaigns based on product (brand & product structure), advertiser (chain structure), insert properties (prices) and leaflet (spreading, type and circulation)<sup>2</sup>. Publi Info gathers this information from every form of printed retail – end-user communication like folder, flyers, door-to-door inserts, national press advertisements, and mailings in the Benelux<sup>2</sup>.

---

<sup>2</sup> Source: [www.publi-info.com](http://www.publi-info.com)

The tool provided data on promotion date, the involved retailer, selected, category, brand and sub-brand, packaging size, promotion quantity, discount structure, folder circulation & share, media type, feature position, and folder & reference price. Using the latter two the discount can be calculated. Figure 3 presents an overview of the organization and product structure on three different levels. A supertradegroup is a holding company that can own multiple supertraders under different names. A supertrader may have different retail formulations (e.g. small, regular, large size). Smaller retailers can be local and are not always part of a supertradegroup or supertrader. The products are structured by category, in which multiple brands may be managed, usually divided into sub-brands with a variation in product types.



**Figure 3: Retailer and Product Structures**

We are only interested in Dutch market and 638 promotions were available over a four-year period. This is enough to examine the supertrader and brand-level, but once more data is available they can be examined on a more detailed level. Comparing this to SCA's internal records, we found a small number of promotions to be absent from PI Web, but were able to reconstruct some of those using Excel files in which promotions are documented.

Second, ACNielsen Buy provided weekly sell-through levels for the same four-year period. There are two remarks with regard to this data: (1) it also includes sales during non-promotion weeks and (2) sales are not only subdivided into sub-brand variants but also packaging size. Using the promotion dates obtained from PI Web, we were able to filter the data leaving behind only sales for the promotional weeks, but still divided into variants. This was addressed by cross-referencing the packaging size obtained from PI Web with the variant description in the ACNielsen data. The resulting sales numbers can now be matched to every specific sales promotion campaign as identified in PI Web.

Third, SAP is used for planning purposes and contains information such as forecast numbers (produced by employees) on both baseline and promotional sales, and the customer order quantity and delivered quantities. Our focus is on promotional sales, and we can use previous

forecasts to validate our model in a later stage. Some retailers may return unsold quantities after the promotion has ended and therefore customer order and delivered quantities do not accurately state what quantity was actually sold to consumers during the promotional campaign. We are only interested in the quantities sold to the consumer, which is the sell-through quantity we already have obtained from the ACNielsen data.

Fourth, various Excel files provided complementary data to the previous sources and because they are more concentrated on internal planning they also provided documentation for the business process. For instance, products are also accompanied by base material numbers and in some instances financial information is present. What is conspicuous is that different departments often work with different worksheet formats which can also vary by brand. This is because some employees started keeping their own documentation to make an overview based on their own interests. Sales is mainly interested in financial performance, marketing in product presentation, customer service in service level, and account managers in contractual agreements.

All four sources have one major disadvantage: due to restrictions by the third-party software it is not possible to automatically access the three information systems, while the Excel files lack structure and standardization. Nevertheless, it was possible to export our data of interest into Excel worksheets which we then all merged into one single structured database. This step allows our tool to retrieve all the relevant data from one central location. The model used for data analysis is presented and discussed in the next section.

#### *4.4 Data Analysis Model*

Because we want our tool to analyze the relationship between several independent variables (promotion characteristics) and a dependent variable (promotional demand), we deem regression analysis to be appropriate (Armstrong & Green, 2005). Armstrong (1985) further argues that simple models should be used in terms of the number of equations, number of variables and the functional form.

Several studies that have been discussed in Section 2.3 have mentioned multicollinearity problems may be present in marketing models in the context of promotional demand forecasting (Leone, 1987; Pole & West, 1991; Shao, 1997). This phenomenon has been widely recognized and addressed in other marketing research studies (Mahajan et al., 1977; Timmermans, 1981; Olsen & Granzin, 1990; Moorman et al., 1993; Christen et al., 1997; Malthouse, 1999; Grewal et al., 2004; Baltagi, 2005; Osinga et al., 2010; Hair et al., 2011). We therefore first conducted a preliminary analysis which confirmed the presence of multicollinearity and then apply ridge regression analysis to solve this problem.

#### 4.4.1 Preliminary Analysis & Multicollinearity

Often the explanatory variables in marketing models are highly correlated with each other, which is a condition referred to as multicollinearity (Farrar & Glauber, 1967). The ordinary least squares method assumes independence between explanatory variables, and using the method results in inflated estimates with high variance and whose signs may reverse with negligible change in data (Hoerl & Kennard, 1970). It further prevents the disentangling of the separate influences of the explanatory variables (Kmenta, 1971). There are various methods available to detect and diagnose this condition, three of which are discussed below.

First, the correlations among the predictor variables can be calculated. A high correlation between two or more variables could be a first indication multicollinearity is present, however this is not necessary (Kmenta, 1971). Table 2 contains correlations between seven variables that are included in the model. Due to proprietary data we do not name the variables at this stage, but refer to them as  $X_1, X_2, \dots, X_n$  presented in an arbitrary order. Several strong and moderate correlations exist, which could point to the presence of multicollinearity.

**Table 2: Correlation Matrix for the predictor variables**

	$X_1$	$X_2$	$X_3$	$X_4$	$X_5$	$X_6$	$X_7$
$X_1$	1.000	0.972	0.881	0.265	0.146	0.309	0.227
$X_2$		1.000	0.439	0.537	0.183	0.251	0.789
$X_3$			1.000	0.861	0.994	0.407	0.286
$X_4$				1.000	0.472	0.294	0.240
$X_5$					1.000	0.728	0.133
$X_6$						1.000	0.254
$X_7$							1.000

Multicollinearity need not necessarily cause any problems in inference (Maddala, 1992), which is why now we examine the sensitivity of the structure of the data by adding or removing either (1) variables or (only a few) (2) observations, and investigate whether large shifts in the estimated regression coefficients occur (Hoerl & Kennard, 1970; Vinod, 1978). This is indeed the case in most instances and demonstrates very unstable parameter estimates. Timmermans (1981) states that "this leads to a situation in which it is hard and often impossible to obtain stable and precise estimates of all the individual parameters" and forms "a serious limitation when a researcher wishes to use the estimated coefficients for forecasting". We conclude that this is the objective of our research and therefore turn to ridge regression analysis.

#### 4.4.2 Ridge Regression Analysis

One of the techniques to overcome the multicollinearity problem and produces more stable estimates, is ridge regression analysis, introduced by Hoerl and Kennard (1970). Ridge regression is an alteration to ordinary least squares regression by introducing a biasing (regularization) parameter. This results in more efficient regression coefficient estimations with a smaller variance (more stable). We briefly demonstrate this alteration by considering a standard linear multiple regression model in which  $\gamma$  and  $X$  are related:

$$\gamma = X\beta + \varepsilon \quad (1)$$

where  $\gamma$  is an  $n \times 1$  vector of observations on the dependent variable.  $X$  is an  $n \times p$  matrix of full rank of observations on  $p$  explanatory variables;  $X'X$  is a zero order correlation matrix.  $\beta$  is a vector of true coefficients, and  $\varepsilon$  is an  $n \times 1$  vector of uncorrelated error terms, with  $\sigma^2$  as population variance of  $\varepsilon$ . The least squares estimate of  $\hat{\beta}$  for this model is:

$$\hat{\beta} = (X'X)^{-1}X'\gamma \quad (2)$$

and the variance-covariance matrix:

$$V(\beta') = \sigma^2(X'X)^{-1} \quad (3)$$

This results in unbiased parameters. By including a nonnegative biasing parameter, the ridge parameter  $k$  ( $k > 0$ ) in equation 2, we obtain the ridge regression estimation procedure:

$$\hat{\beta} = (X'X + kI)^{-1}X'y \quad (4)$$

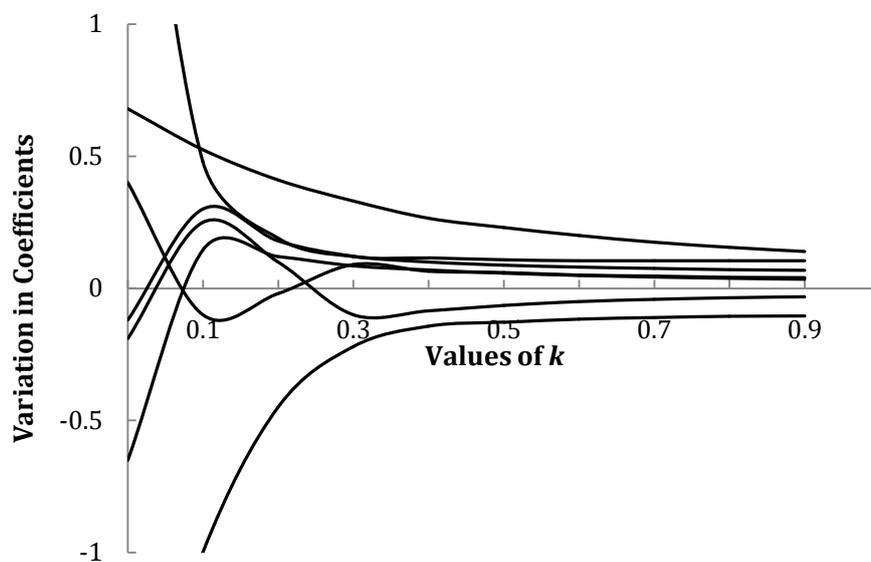
The variance-covariance matrix of the ridge regression estimate is:

$$V(\beta') = \sigma^2(X'X + kI)^{-1}X'X(X'X + kI)^{-1} \quad (5)$$

$I$  is the identity matrix through which together with  $k$ , large values of  $\beta$  are penalized, while smaller values are preferred. As  $k$  increases, the ridge estimator decreases towards zero. This forms the bias, but where  $k = 0$  or  $k \rightarrow 0$ , respectively no or little bias is present and the estimates are (almost) equal to the estimates as provided by equation 2. As  $k$  increases, so does the bias. The bias is:

$$bias(\hat{\beta}) = -(X'X + kI)^{-1}\beta \quad (6)$$

In ridge regression analysis a smallest value of  $k$  is selected where all the coefficients first begin to exhibit a stable pattern. This is known as the ridge trace method (Equation 4), in which the estimated ridge estimators are plotted against  $k$  ( $0 > k > 1$ ). Figure 4 shows such a plot and we can observe that the coefficients are relative stable between approximately  $k = 0.3$  and  $k = 0.4$ . Since ridge regression includes a bias parameter  $k$  and minimizes the mean squared error (MSE) of the estimators, Hoerl and Kennard (1970), Theobald (1974) and Vinod and Ullah (1981) have demonstrated the existence theorem prescribing a  $k > 0$  exists such that the MSE of ridge is less than the MSE of the ordinary least squares. The objective is obviously to find this value of  $k$ , while preserving stable regression coefficients.



**Figure 4: Ridge Trace**

Although Figure 4 forms a good illustration for the problem and provides insight into the effects for a different value of  $k$ , an automatic selection is preferred for our tool. The reason for this is that the optimal ridge solution may be different for every retailer-product combination, yielding a different  $k$ , and may furthermore change over time. Different methods have been presented in the literature but we will use Generalized Cross Validation (GCV) which is a computationally more efficient form of Leave-One-Out cross validation (Golub et al., 1979). This is one of the most common implemented methods used to choose a ridge parameter in statistical packages like SAS, S-PLUS and Matlab. We will use the open-source SciKit-Learn module that has automatic selection of  $k$  implemented using GCV.

Although we can use ridge regression to cope with the data, Maddala (1992) points out that weak data and inadequate information are likely the sources of the multicollinearity problem. For future purposes, it is therefore helpful to gather more data, but first it should be examined what additional data will be most helpful. Chapter 2 may provide some guidance on this issue.

## 4.5 System Architecture

This section describes the system architecture of the forecast support tool and will first briefly address the high-level implementation choices, followed by a description of the various other aspects of the user interface.

### 4.5.1 Implementation Choice

In order to develop the forecast support tool we have selected the Django web application framework, written in Python, and is a good choice for a database-driven application like ours. Because it is platform-independent, it can be deployed in a production environment on any web server that supports the Python programming language. An advantage of this client-server architectural model is that a user only needs a web browser with in our case also JavaScript support. Another advantage is that it can be used by multiple users simultaneously which suits our particular purpose as different departments are involved in the business process. The built-in authentication system controls permissions for specific users or departments.

Django also offered a major advantage for the development process of the tool that entails the availability of numerous software packages, libraries and extensions. These are existing modules, developed by the community, that offer certain functionalities or features and are readily available to be implemented with Django. It was necessary to use several of these modules, for instance SciKit-Learn that together with SciPy/NumPy enables estimation of statistical models. A complete overview of these and other modules that were used and their description can be found in Appendix A: Table 5.

A possible alternative to this choice of implementation is integrating or linking the tool to an existing information system or database. Due to technical restrictions and company policy this was not possible now, but could be considered in the future.

### 4.5.2 User Interface

The user interface is the component of the tool that links the user to the database and model. Because our tool is web-based, the graphical user interface consists of web pages. To make the application user-friendly, we have chosen to implement a dashboard which is the first thing a user sees when the application is accessed. This can be labeled as the beginning of the main workflow trajectory as depicted in the activity diagram in Figure 8. The dashboard presents an overview of two sections with information: on top is an overview of planned promotions and below an overview of previous promotions with missing realized sales data. Figure 5 shows a screenshot of the dashboard with dummy data and the retailer names and product codes concealed. Using the buttons on top, columns of interest can be selected while hiding the rest.

This list incorporates a warning system for approaching deadlines, which is an important aspect with regard to the required production lead-time. If a deadline is due within 4 weeks, promotions will be highlighted orange if no forecast has been determined yet. In case of 2 weeks or less, it will be highlighted red. The bottom list notifies and urges the user to provide the realized sales for the promotions. This is important because once this information is provided, the promotion can be included in future analyses. This helps build the promotion database.

Another aspect that helps build the database is the function to add events (Figure 6) by filling out a form with a set of required fields. This is depicted as the second user-initiated activity shown in Figure 8. Some may have a predefined set of options that are directly acquired from the database, which is supposed to minimize the risk of false user input. If for the other fields invalid user input is detected, an error message is displayed next to the affected field with instructions. This could for instance be a wrong character type (e.g. letters instead of numbers) or data formatting (e.g. date). Once all fields have been filled in correctly, the promotion is created and added to the list with planned promotions.

**List of forthcoming events**

Today is April 7, 2013 (week 201314)  
 Second warning week is 201322  
 First warning week is 201325

**Toggle Columns**

Week(YYYYWW) retailer product Quantity (#SKU) Free Quantity Face Value Folder Circulation Folder Share media Folder Price Reference Price Reduction %  
 PRODUCT PLUS

Week(YYYYWW)	retailer	product	Quantity (#SKU)	Free Quantity	Face Value	Folder Circulation	Folder Share	media	Folder Price	Reference Price	Reduction %	PRODUCT PLUS
201324	A003	ESO TP4 12	2	1	25% - 45%	1100000.0	828.3086	RETAIL MAGAZINE - Second Half/Left	5.59	8.46	29.2	Edit event Analyze event
201340	A001	ESO TP4 12	2	1	25% - 45%	6200000.0	12.2859	NATIONAL PRESS - Front Cover	8	12	33.33	2+1 Edit event Analyze event
201318	B002	EFR TP3 24	1	0	10% - 25%	6300000.0	4629.3126	FOLDER - Second Half/Left	6.79	8.57	38.82	Edit event Analyze event

**List of past events**

Week(YYYYWW)	retailer	product	Quantity (#SKU)	Free Quantity	Face Value	Folder Circulation	Folder Share	media	Folder Price	Reference Price	Reduction %	PRODUCT PLUS
200710	C010	ESO TP4	2	1	2.5% - 5%	40000.0	12.2859	REGIONAL PRESS - Front Cover	5.18	7.77	33.33	2+1 Edit event Analyze event
200710	B001	ESO TP4	2	1	2.5% - 5%	4300000.0	2740.8823	FOLDER - Second Half/Right	5.19	7.77	33.2	2+1 Edit event Analyze event
201251	A004	ESO TP4 4	2	1	0.75% - 1.25%	20000.0	11.3612	FOLDER - Second Half/Left	5.18	7.77	33.33	2+1 Edit event Analyze event
200724	B001	EFR TP3 8	2	1	0% - 0.75%	1000000.0	550.7195	REGIONAL PRESS - Front Cover	5.18	7.77	33.33	2+1 Edit event Analyze event
200724	B001	EFR TP3 8	2	1	2.5% - 5%	20000.0	12.7483	REGIONAL PRESS - Second Half/Right	5.18	0	0	2+1 Edit event Analyze event

Figure 5: Dashboard

**Figure 6: Add Promotion Form**

The need to maintain the database originates from the fact that promotions are subject to change during the negotiation process (Section 4.2), and therefore we also implemented a function to edit a promotion. This step is not shown in Figure 8 because it is facilitated by the same form as for adding a new promotion (Figure 6). Any adjustments are tracked and kept in the event log with information about which attribute has been changed and by whom. Additionally, optional comments can be added to the log by the user to explain any committed changes. Over time, this will permit a new feature to be implemented that can analyze the change logs to find out whether or not adjustments have contributed to a more accurate forecast. In order to provide memory (Hoch & Schkade, 1996) and similarity support (analogies) (Green & Armstrong, 2004), the user can use the search and filter function to find similar promotions.

So far we have discussed the part of the tool that concentrates on building and maintaining the database. The next part is the analysis functionality, facilitated by SciKit-Learn and other supporting Python modules. It aims to provide insight into the dynamics of promotions at different retailers and for different products. Analysis takes place after a certain promotion has been selected from the dashboard or from the search results. This is the third user-initiated step shown in Figure 8. At this point the aggregation level for retailer and product can be selected (Section 4.2) based on which the system extracts an appropriate reference class. Next the coefficient estimates and confidence intervals are calculated by the regression model and a graph is generated for each attribute to present its relationship to the realized sales (Figure 7).

We have also chosen to enable user interaction by adding a dynamic component in the form of sliders through which the attribute values can be adjusted. After adjusting a value, the model is automatically recalculated at run-time and the user can immediately witness the effects of his actions in the graphs. For example, if the folder price is adjusted it will have an effect on the relationship of the other attributes. This allows the user to explore the dynamics of the promotions and experiment with different values without the need to repeat the whole process again. Once the user has made a decision he can edit the promotion by providing a forecast or leave this for another time. After the promotion has passed, it will be moved on the dashboard from the upper list to the bottom list and will remain there until the realized sales are added. This data is provided by ACNielsen, but needs to be entered manually as it is not possible (yet) to directly link the tool to the ACNielsen database or to distinguish between promotional and non-promotional observations.



**Figure 7: Regression Results Overview**

Apart from our choice to implement our tool in Django/Python (Section 4.5.1), there are various alternatives or extensions that can be considered within the tool. For example a different statistical module could be used or additional models added to provide users with more options to analyze the data and to compare the results from different models. Also, additional data could be added that captures additional explanatory variables from other dimensions like consumer response or in-store presentation (Section 2.2) and also contribute to the promotional effect. Literature suggests this would result in a model that can better explain the variance, but at this stage we only had a limited number of variables available and have tried to utilize these as best as possible while providing a platform that can be expanded in the future.



## 5. Evaluation

As noted in chapter 3, the evaluation of an artifact is an important part of design science in information systems research and is considered to be a process. This chapter discusses the evaluation methods that we applied during and after the development process. Our development process consisted of four iterations, and after every iteration we performed testing (Section 5.1) to assess to which extent the identified problem has been solved, but may also uncover mistakes, faults, failures and errors. In order to guarantee quality, functional and structural testing methods were performed. We further perform a descriptive evaluation in order to demonstrate the utility of the tool (Section 5.2) and finally assess the predictive performance of the tool.

### 5.1 Testing

An evaluation method appropriate for our research is software testing which can be classified as the 'box approach' and divided into black- and white-box testing. Because we are testing a tool that we developed ourselves, we can use both box approaches and the testing type is classified as development testing. The other testing types are functional (which is part of black-box testing) and destructive testing (deliberately providing invalid input).

#### 5.1.1 White-box testing

White-box (or: structural) testing is used to assess to what extent the programming code complies with the design specifications, by looking whether the internal (data) structures and workings of software contain faults, failures and errors. Because this type of testing requires programming skills and insight in the internal workings of the software architecture, it is the developer who is most capable of performing this task. We conduct white-box tests during development and is therefore classified as development testing.

Testing can take place on unit, integration and system level, but structural testing is most often applied on unit-level. Unit-level testing among others provides information on the existence of dead (unused) code and the correctness of used logic and algorithms. There are different techniques to design a test case, which include path testing and control flow. In path testing first all possible paths through the software are identified and then selected in such a way that every statement/method has been executed at least once. The tester can then examine each path. Control flow testing involves assessing if methods and method calls are executed in the right order, which means it looks if the right choices are made during execution in what in the end will constitute a followed path.

Only code we have written ourselves is subjected to tests, as we do not possess the required programming knowledge of the software packages and libraries we have used that were developed by other parties. An overview of the third-party software is found in Appendix A: Table 5. For the code we do subject to tests, we make a distinction between Python methods and JavaScript methods. The Python code forms the inner working of the tool and interacts with the database, performs the calculations, and prepares the data which is supposed to be presented on the user-interface (web-page). The JavaScript code is utilized to render the dynamic (graphical) features of the web-page. An overview of the tested methods, their purpose and supplemental information can be found in Appendix B: Table 6 & Table 7.

We performed structural testing initially and mainly on the unit-level (mainly iteration 1 & 2) to examine the code and workings of each method individually and to ensure it does what it is intended to do. Using test code, methods are called with accompanying parameters, utilizing all possible paths, and the return value is compared to the expected value. As we constructed new methods that call previously created methods, or methods from existing software, we moved on to integration-level testing (mainly iteration 2 & 3). On this level we verified whether the new code interfaces and interacts correctly with other components.

### *5.1.2 Black-box testing*

Black-box (or: functional) testing is used to assess to what extent (implemented) software functionality complies with the design specifications, without looking at the internal software structures. This method therefore does not require knowledge about programming or the internal (data) structure, but does require knowledge about what the software is supposed to do. We performed unit- and integration-level testing with white-box testing, as for black-box testing we performed system-level testing to test the tool as a whole. Due to a different mindset between a software developer and end-user, this type of testing requires collaboration between the two if the first party wishes to satisfy the needs of the second party. An end-user is able to provide certain insights based on expert knowledge and experience which the developer lacks.

Testing proceeds by first developing test cases which mainly consist out of the identified function that is being tested, the input and the expected output. Nonetheless, non-functional tests can also be used. Next the test case is executed and the result is compared with the expected output. In our case not every function is also tested in terms of structural testing. Input validation is an example because it is already an integrated in Django so we test it as a function, but not the code itself with our structural tests. A summarized overview of the tested functions, their purpose and supplemental information can be found in Appendix A: Table 8.

## 5.2 *Descriptive Evaluation*

Descriptive evaluation is useful in demonstrating the utility of software, which is mainly a combination of behavior and opinion based measures. We begin this evaluation by briefly justify the implementation choice, assisted by information from the knowledge base. We then depict a scenario and reflect back to the design principles (challenges) as proposed in Section 4.2, supported by information from prior relevant research. The scenario is based on real company data, but due to proprietary data, numbers, company and product names have been concealed.

According to Sprague (1980) and Power (2004) we can classify our tool as a model-driven, task-specific, web-based FSS that targets internal users. Integrating the model into the tool can help gain a competitive advantage (Thearling, 1998). The choice for a web-based tool is based on advantages it has for both the development stage and deployment of the tool in a business environment. The advantage of developing with Django is that numerous modules already exist and can be used in our project, which saves time compared to programming everything ourselves. Because these modules are open-source and usually the result of community collaboration, they have already been extensively tested and updated on a regular basis.

The advantage of deploying a web-based application is that it can be accessed from any client computer with a web-browser, hence a thin-client is sufficient (Power, 2004). Another advantage is that it enables the business to centralize and control information while allowing simple and easy access to the application over LAN, intranet or the internet allowing access to multiple users simultaneously (Power & Kaparathi, 2002). Centralization of the data is an important aspect that enables marketing data manipulation, transformation, and analysis to discover usable patterns and relationships (Rygielski et al., 2002; Lee et al. (2004). Accessibility is also important because different departments are involved in the forecast task.

For our scenario, we add a new planned promotion to the database using the form as shown in Figure 6. Some fields (e.g. product, retailer) are always predefined and therefore implemented as a dropdown list in order to minimize the risk for input errors (Asimakopoulos et al., 2009). The options are retrieved from the database and can be managed through the admin interface. The structured nature of the form helps build and maintain the database which in time results in growing reference class that will enable more extensive analysis methods in the future. After the form has been filled out correctly, it can be saved and added to the database (see Figure 9). Previous promotions can also be looked up for memory support (Schank, 1982; Hoch & Schkade, 1996) or criteria can be provided for a filter to look up similar previous promotions for similarity support (Green & Armstrong, 2004).

Week	Sales	Retailer	Product	Quantity (SKU)	Free Quantity	Face Value	Folder Circulation	Folder Share	Media	Folder Price	Reference Price	Reduction %
200730	(None)		EFR TP3 24	1	0	70% - 100%	6200000.0	4811.984	FOLDER - First HalfLeft	5.79	8.37	30.82
200730	(None)		EFR TP3 24	1	0	10% - 25%	20000.0	13.673	REGIONAL PRESS - Second HalfRight	5.79	8.37	30.82
200735	(None)		EFR TP3 24	1	0	70% - 100%	6200000.0	4525.3126	FOLDER - Second HalfLeft	5.79	8.37	30.82
200745			ESO TP4 16	1	0	70% - 100%	6200000.0	2282.6563	FOLDER - Second HalfLeft	5.69	8.37	32.02
201117			EFR TP3 24	1	0	25% - 45%	6200000.0	4381.9769	FOLDER - Second HalfRight	5.00	7.28	31.32
201213			EFA TP3 32	1	0	10% - 25%	1100000.0	904.6001	FOLDER - Front Cover	7.49	11.28	33.13
201217			ESO TP4 16	1	0	45% - 70%	6200000.0	4381.9769	FOLDER - Second HalfLeft	5.99	8.93	32.92
201227			ESO TP4 16	1	0	10% - 25%	6200000.0	2119.3206	FOLDER - Second HalfRight	5.99	8.93	32.92
201227			EFA TP3 24	1	0	10% - 25%	6200000.0	2119.3206	FOLDER - Second HalfRight	5.99	8.93	32.92
201227			ESO TP4 16	1	0	5% - 10%	1000000.0	399.6226	REGIONAL PRESS - Front Cover	5.99	8.93	32.92
201227			EFA TP3 24	1	0	5% - 10%	1000000.0	399.6226	REGIONAL PRESS - Front Cover	5.99	8.93	32.92
201318	(None)		EFR TP3 24	1	0	10% - 25%	6200000.0	4525.3126	FOLDER - Second HalfLeft	5.79	8.37	30.82

**Figure 9: Search & Filter Functionality**

We continue by examining the dashboard which now also includes the promotion that was just added. We can observe one promotion highlighted orange whose deadline is in less than four weeks, and another promotions highlighted red whose deadline is in less than two weeks. Highlighting is a design issue that provides better insight to the user (Asimakopoulos et al., 2009). We can therefore conclude design principle 1 has been satisfied. To serve the negotiation process, promotions can be edited using the same form as in Figure 10 and any adjustments are recorded for each user. This allows changes to be tracked and will enable analysis of these changes in the future (Lee et al., 2004; Fildes et al., 2006). Reflecting to the ability to create new events that are added to the centralized database, these aspects satisfy design principle 2.

**Change history: Event 201324 - ESO TP4 12**  
 Choose a date from the list below to revert to a previous version of this object.

Date/time	User	Comment
April 7, 2013, 8:38 a.m.	testuser	Initial version.
April 7, 2013, 10:18 a.m.	testuser	Changed folder_pr, ref_pr and commen Folder Price from 5.99 to 8 Reference Price from 8.46 to 12

**Figure 10: Change Log**

Finally, we address the analysis component by using the promotion that was just created as an example. This is an automatic process without required any further user-input to produce the results as shown in Figure 7. The page provides information about the relationship between different attributes (predictor variables) and sales (dependent variables) (Fildes et al., 2006). We chose to visualize this using graphs with a regression line and a confidence interval in order to provide better insight in the relationships (Keim, 1996). For instance, the higher the reference price, the higher the promotional sales. An explanation could be that more expensive products are preferred over cheaper products during a promotion. Another, stronger relationship, is the quantity in one consumer unit: a higher quantity results in higher sales.

Because of the dynamic user interface (Shaw et al., 2001), attribute values can be adjusted in runtime using the sliders which automatically adjusts the graphs according to the model. This enables the user to further examine the relationships by adjusting one or more values. Unlike in Lee et al. (2004), where only one variable could be adjusted, our tool provides a more extensive form of adaptive support (Kleinmutz, 1990; Hoch & Schkade, 1996) by allowing adjustments to be performed on multiple variables at once. This enables the user to further explore the dynamics and patterns involved with promotions and we can therefore conclude design principle 3 has been satisfied.

### 5.3 Model Evaluation

Using previous evaluation methods we have determined our tool has been successfully implemented with sound code and provides the required functionality based on the proposed design principles (Section 4.2). We have selected and determined that ridge regression is the appropriate model to analyze the data currently present at the case company (Section 4.4) but it is interesting to also examine model performance. Common critique for this model is that it introduces bias, violating regular linear model theory, but in return we acquire a more stable system, coefficients that are closer to their true value and carry the proper sign, and the residual sum of squares is not inflated (Hoerl and Kennard, 1970).

For this research observations were relative scarce and the set of explanatory variables available covered only a part of all the potential useful variables as suggested in Section 2.2. The implication of this can be perceived through the coefficient of determination ( $R^2$ ) that expresses how much variance is accounted for by the model. Because we used data from the case company that deals with three different retailers (supertraders) and three different brands that respectively (may) exhibit different behavior and properties, a separate model is used for each of the possible retailer-brand combinations. This at the same time is the key for the tool when selecting an appropriate reference class. An overview of the  $R^2$  for each of these retailer-brand combinations is given by Table 3.

**Table 3: Coefficients of Determination**

		Retailer		
		I	II	II
Brand	A	0.67	0.62	0.59
	B	0.39	0.43	0.48
	C	0.41	0.44	0.45

The values are relatively low, but it should be taken into account that, compared to the potential variables presented by the literature, only a limited number of variables were available for this study. This is likely what has caused the multicollinearity problem, and introducing additional adequate data like for instance consumer behavior and in-store presentation that cover the marketing mix more extensively, might improve model fit. Furthermore, our (highest)  $R^2$  values are comparable to those reported by for instance Doyle & Saunders (1985) and Bell et al. (1999) that account for 63% and (up to) 70% of the variance, respectively, for the same product categories. The low fit for brands B and C can be explained by a significant lower amount of observations compared to brand A and the fact that we had access to a less explanatory variables than the mentioned studies. This suggests that gathering these currently unaccounted for variables should improve model fit overall, but especially brands B and C.

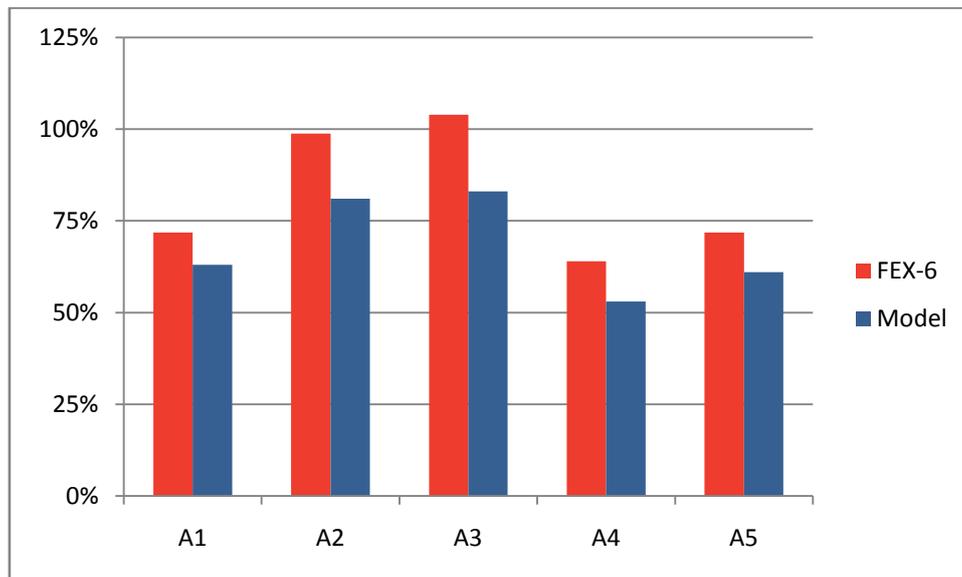
A factor that is known to affect promotional performance are promotions that are organized by competitors shortly before the promotion of the case company. Affected promotions can be classified as outliers within the model as sales substantially suffers from the stockpile effects caused by the previous promotion. When these outliers are removed,  $R^2$  shows overall increases, though in different magnitude (Table 4). The reason is that brand A has had by far the most number of promotions, almost two times more than brands B and C combined, and is therefore more likely to have been affected by a competitive promotion that has taken place in the few weeks preceding the own promotion. Because the case company's lead-time is 7 weeks, by the time a competitive promotion is perceived it's already too late and therefore it is not realistic to include this in the model but should be regarded as a managerial implication.

**Table 4: Coefficients of Determination without competitive outliers**

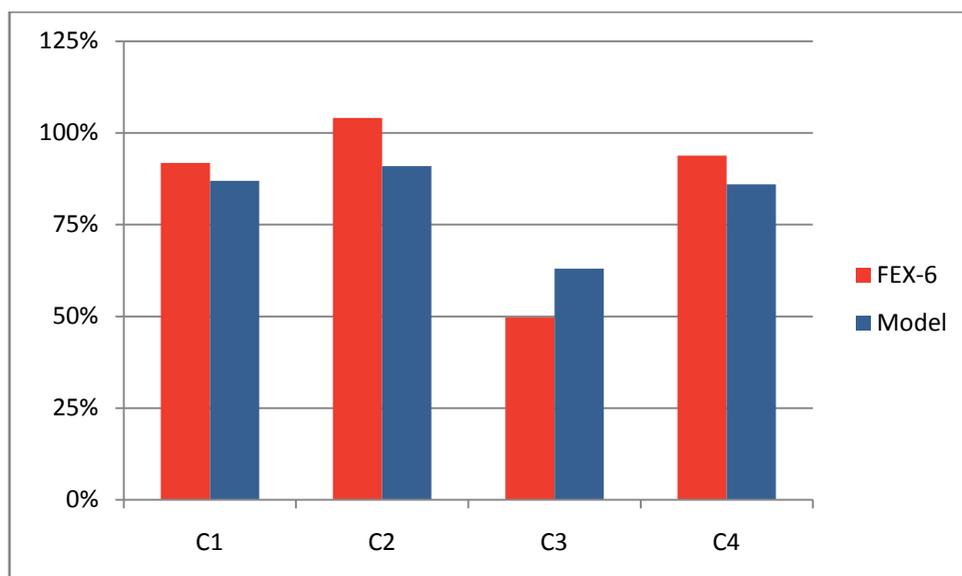
		Retailer		
		I	II	II
Brand	A	0.71	0.65	0.64
	B	0.40	0.46	0.50
	C	0.44	0.46	0.48

We can test the model performance by taking a hold-out data sample that has not yet been 'seen' by the model. Because of the relative low frequency of promotions this sample is too small to apply conventional validation methods. However, ridge regression in our case already implies generalized cross-validation (Section 4.4.2) that should always guarantee the most optimal model fit. Nevertheless, we try to provide insight in model performance by taking several promotions that took place after the latest promotion that was included in the dataset to train the model.

We compare model output to the case company's internal forecast measure, FEX-6, which expresses performance as a percentage error between the forecast six weeks prior to delivery and the actual ordered quantity in the week of delivery. Cases in which the date of the promotion has been moved after an initial forecast has been made, immediately disrupt the FEX-6 by pushing it to excessively high levels. It is not fair to compare these cases to the model output and have therefore been omitted. Also, brand B has too little overall observations to produce reliable results and no useful hold-out sample. In order to illustrate model performance, we report the results of the best and worst performing brands, which are respectively brand A (Figure 11) and C (Figure 12) at retailer I.



**Figure 11: Model Performance Brand A at Retailer I**



**Figure 12: Model Performance Brand C at Retailer I**

We can observe in Figure 11 that the model outperforms the FEX-6 in all instances, but that the difference in error seems to decrease as the overall error rate decreases. It should be noted that A1 and A5 have had almost an identical promotion structure, which is why their results are very similar. As for brand C, the difference in performance between the model and FEX-6 is relatively small, although the model performs better in three out of the four instances. This is likely caused by the small amount of observations in the training dataset and/or that certain important explanatory variables are missing in the model. C3 has a remarkably better FEX-6 performance overall, but also compared to the model which performance worse. This could indicate that certain information was available to the manager but was not accounted for by the model.

Two major assumptions exist for FEX-6. First, it is assumed that the promotion attribute values have not been adjusted by the case company or retailer after the forecast has been determined. Any changes made thereafter are likely to affect the eventual promotional sales and disturb the forecast and FEX-6. We did not have the resources to retrieve all of this information and therefore the model does not account for any possible adjustments that may have happened.

Second, FEX-6 is based on a forecast of, and the order quantity itself (by the retailer) and not the actual sell-through quantity. This practice seems reasonable since the case company has to comply to certain Service Level Agreements (SLA), but it may be a complicating factor because the two quantities may be significantly different as demonstrated by e.g. Goodman & Moody (1970). The question remains where the balance lies between service level and forecast level, and whether the current performance measure is appropriate. Furthermore, all of the studies discussed in Section 2.2 and 2.3 use actual store sales numbers (sell-through quantity) as the dependent variable. In many cases this data is provided by ACNielsen, which is also the data source used in our model.

## 6. Conclusion

This research has followed the design science framework for IS research as presented by Hevner et al. (2004), which includes seven guidelines. This chapter first discusses each one and how they apply to this research. Then the managerial implications are presented, followed by addressing the research limitations, and finally recommendations for future research.

### 6.1 Conclusions

As required by the first guideline, design science research should produce a viable artifact in the form of a construct, a model, a method, or an instantiation. Conform to the first guideline, this research resulted in a tool (instantiation) that provides insight to forecasters about promotional dynamics. The second guideline requires the development of a technology-based solution to an important and relevant business problem. The business problem addressed by this research originates from SCA and implied the desire to be able to analyze and present the dynamics of sales promotion campaigns so that it can be used for forecasting purposes.

The limitations within the (technical) environment and infrastructure concern the different data sources that cannot be linked. This is why the tool incorporates a new structured and centralized database which contains all previous promotions in terms of their attributes and result. This solution allows easy analysis, despite the stated limitation. The user interface facilitates easy management of promotions, new promotions can be added, and based on the attributes the tool analyzes previous promotions and presents the results terms of relationship between promotion attributes and realized sales.

The third guideline demands the artifact to be evaluated in order to gain better understanding of the business problem and test and improve the utility, quality and efficacy of the artifact. Different methods can be applied for this evaluation purpose, which is considered an iterative process. Testing in terms of functionality and structure was involved during the development process of the tool. As for functionality, this type of evaluation method was able to uncover software errors and faults, while structural testing provided information on the internal workings of methods. Over the different iterations, these types of testing enabled us to appropriately address different problems.

In order to demonstrate utility it was necessary to deploy a scenario as part of descriptive evaluation. The managers and employees that were involved in this research were pleased that it was finally possible to get insight in the nature of relationships between different promotion attributes and promotional sales. Other advantages are that it suits the business process, the centralized database, the visual output, and the possibility to adjust attribute values and immediately witness the effects.

The fourth guideline dictates design science needs to provide an academic contribution in the areas of the design artifact, design foundations, and/or design methodologies, but also to the business environment. Within the field of marketing a lot of research has been conducted on sales demand in general, but studies on specifically sales promotion campaigns are rare and so are software systems that adequately support them. Hence, this research contributed to the field of marketing by specifically focusing on promotions and also contributed to the business environment by how a tool can be deployed to appropriately support forecasters that deal with promotions. We have not found other studies that address this problem by using the type of regression analysis and web-based software we used. Considering the limited existing knowledge base, these aspects can be useful and a motive for future research projects.

Research rigor is addressed by the fifth guideline and covers the way in which research is conducted, that is the application of rigorous methods in both the construction and evaluation of the designed artifact. For this research regression analysis was used, which is a method used in forecasting and especially finding (and understanding) relationships between a dependent variable and several independent variables. As distinct from most other types of analysis, regression does not require time series and is therefore suitable for this business problem. Potential evaluation techniques, as discussed in the design science framework by Hevner et al. (2004), have been presented and depicted in chapter three and those with sufficient relevance were then applied in chapter five.

The sixth guideline depicts design as a search process in which all available means should be utilized in order to reach the ends, while satisfying laws and constraints in the environment. For this research software was developed by ourselves, although existing packages and libraries were incorporated to implement certain features and functionality. SCA provided data, expert knowledge and an environment for development and testing. The amount of data available on past promotions constituted an important constraint and formed analysis complications. This expressed itself by having less attributes and promotional events available than initially anticipated, multiple incompatible data sources, and multicollinearity in the data. Nevertheless, solutions were identified and utilized to successfully our objective: for instance. a new centralized database was created and a suitable analysis method was found and implemented.

The seventh guideline asks the researcher to publish the study results to both technology-oriented and management-oriented audiences. As for the first group this research provides a notion on promotions in particular and how the tool provides support. Because repeatability is of importance, a detailed description of the tool is provided in chapters four. From this technical point of view, future research may be able to deliver further extensions to the knowledge base.

As for management-oriented audiences, the tool currently serves the purpose of being an example of how a forecast support tool can be designed and implemented to effectively support promotion forecasters. The tool is a steppingstone and has the potential and flexibility to be further extended and improved based on additional or other business needs and environment.

## 6.2 *Research Limitations*

Some of the limitations to this research had been anticipated, while most turned up during the research process. It was known in advance there would be a time constraint that would put limitations on the degree of implemented (potential) functionality and features. An unexpected limitation which had a significant impact, has been the availability and usability of data due to various factors.

First of all the data pool was sparse which had to do with the scarce nature of promotions, then the number of attributes were limited, and finally values were often missing. Because of this, there were not as many promotion attributes available as provided by literature to include in our research. We should take into consideration there are unrecorded attributes that could be useful as explanatory variables and some may be better than the ones that are used now.

Retail data were also unavailable because sharing information between SCA and retailers is just in its infancy and when it does occur it's only to a limited extent. This constrained the level of detail that could be applied and led to generalizing promotions on product and retailer group, instead of respectively specific brand or product variant and retailer. Also, some products and retailers had to be evaded entirely due to a too small number of observations.

Because of the limited scope and amount of data, it is unclear how additional data would affect the current analysis results. Currently, mainly product characteristics are used as variables, but literature addresses a various other potential variables that also affect demand. If for example market data like (individual-level) consumer behavior was available, additional variables could be used, resulting in new insights and possibly different but more accurate analysis results. The existing knowledge base provides good examples of how the selected data, focus and level of aggregation can affect analysis results: e.g. Gupta (1988) versus Van Heerde et al. (2003).

Data entry could be appointed as a limitation, because the data comes from existing data sources and therefore incites manual work. Though not feasible at the moment, some form of automated data entry could be introduced. Finally, a different analysis method choice could lead to different results as every method is bound to its own distinctive set of assumptions. We selected a method which is able to cope with the (currently) limited and poor data – though more robust, it does introduce bias.

### 6.3 *Managerial Implications*

The lack of good data was a great limitation and therefore we suggest management to take appropriate steps to gather more data and do that in a structured way. Our tool is a starting point to accomplish this, but there is other potential data that has not been explored yet. Literature for example suggests that store-level and consumer data could provide more explanatory variables and new insights in the effects of promotions. This would lead to more opportunities to analyze different dimensions, e.g. effects on the short- and long-term and consumer behavior. To accomplish this, more emphasis has to be put on data collection which should be a more pro-active task, instead of (solely) waiting for third-party data.

In addition to the poor data and the multicollinearity problem, management could also experiment with promotion design by dropping the conventional designs that have been used thus far. Creating different setups and combinations of promotion attributes that have not yet, or rarely, been used, attributes could be isolated what may eliminate multicollinearity. These experiments could also be conducted on a smaller scale to minimize risk, rather than on a national scale, and saves time if organized in parallel.

We have chosen to only measure promotion performance in terms of realized sales, while it could also be a good idea to look at financial aspects, especially on the long run. According to literature, the implication of promotions is that demand is borrowed from the future, thus what would have been turned over anyway, is being sold for a lower price. If no other reasonable gains are realized on market size and market share, the effectiveness and purpose of such promotions should be questioned. In our case, the specific market is considered turbulent and governed by high and increasing promotional intensity. Nevertheless, examining effectiveness could result in new strategies with regard to promotions.

Finally, management should realize forecast support systems only provide support and there will always be need for human judgment. Such systems are able to capture key data and offer cognitive support to forecasters for what would be otherwise impossible tasks, like memory support or analysis of great sums of data. The focus should therefore mainly be on forecasting as a management process, and control systems should be implemented before selecting forecasting software, preferably in a single forecasting infrastructure (Moon et al., 1998). Whatever software is used, training should be provided on how to work with the software and to understand the used methods, which could also help gain acceptance.

#### 6.4 *Future research & Recommendations*

Reflecting to the research limitations, and adjacent to the managerial implications, there are possibilities for improvements and extensions in the form of recommendations that could be useful in future research.

Little studies have been conducted in the field of marketing on the topic of forecasting promotional sales, especially because empirical data demonstrates increasing promotional intensity and importance to manufacturers. This suggests more focus should be put on this topic to find suitable methods and techniques and how to deploy them in operational environments.

Also, existing software systems little, or do not provide effective support for special events like promotions. Often little incentive is provided for even the creation of a centralized record of promotions (Lee et al., 2007). This is the reason why decisions are mainly based on judgment, but even with quantitative methods there is room for judgmental adjustment. The data we have used is quantitative, but human judgment is difficult to quantify. Other forms of qualitative data can also be considered, though more research is required into how to capture and utilize such information.

Due to the limited data, we had to aggregate on retail and product level, while more data would allow to slice down to specific retailers, geographic regions, and specific product variants. This could also result in various shifts in results, but should be considered if more insight is desired on a more detailed level. This could lead to detection of different behavior for different regions, as some promotions become more attractive and successful than others. Additional information on consumer response could identify changes in consumer behavior over time. Having that information, could also lead to development of models that anticipate and target such behavior based on the business strategy.

The downside of the ridge regression analysis is that it introduces bias in the estimated coefficients but in return we acquire a more stable system, coefficients that are closer to their true value and carry the proper sign, and the residual sum of squares is not inflated (Hoerl and Kennard, 1970). Although ridge regression is helpful in our situation, weak data and inadequate information are likely the sources of the multicollinearity problem. This is why it is recommended to first examine what prior and additional information will be most helpful and then acquire that data. Prior studies in the field of marketing, like those discussed in Chapter 2 and decompose the promotion effect, offer insight to what other variables may be useful and often also present an underlying theoretical background for their significance.

Though forecasting is a small part of the whole supply chain process, it is considered to be of great importance by the case company due to its big impact on performance. Future research should therefore look at how our approach can be incorporated into the chain, with regard to the enterprise resource planning system. Combining this with marketing mechanisms and organizational objectives, it could provide improved anticipation and more conscious decision-making on promotional attributes in the run-up to promotions. Linking the tool with a financial information system could allow further analysis of the financial performance and effectiveness of promotions.

Our tool provides a friendly user interface and is able to deliver an analysis within a few clicks and by filling out one form. We did have to face the problem that it was not technically possible and unsupported to directly connect to existing data sources and automatically extract the required data. We solved this by exporting data into worksheets, which we then used to assemble our own (standalone) database. This is not necessarily a problem since our tool offers the opportunity to manage its database via forms and search function, but this does induce data redundancy as the same data can also be found at other data sources. A recommendation would be to examine the possibilities of interconnecting these data sources to enable automatic data access and extraction. Also, it should be questioned whether our tool is best to be used as a separate instantiation, or if it should be integrated in an existing information system like a customer relationship management system which was recently introduced in the organization.

## References

- Ailawadi, K.L., Neslin, S.A. The effect of promotion on consumption: Buying more and consuming it faster. *Journal of Marketing Research*, pp. 390-398, 1998.
- Alon, I., Qi, M., & Sadowski, R. J. Forecasting aggregate retail sales:: a comparison of artificial neural networks and traditional methods. *Journal of Retailing and Consumer Services*, 8(3), pp. 147-156, 2001.
- Armstrong, J. S. *Crystal Ball to Computer*. Wiley, New York, 1985.
- Armstrong, J.S., Green, K.C. Demand forecasting: Evidence-based methods. Working paper, Monash University, Department of Econometrics and Business Statistics 24, 2005.
- Asimakopoulos, S., Fildes, R., Dix, A. Forecasting software visualizations: an explorative study. In: *Proceedings of the 23rd British HCI Group Annual Conference on People and Computers, Celebrating People and Technology*, pp. 269-277, 2009.
- Assunção, J.L., Meyer, R.J. The rational effect of price promotions on sales and consumption. *Management Science* 39.5, pp. 517-535, 1993.
- Baltagi, B. (2008). *Econometric analysis of panel data*. Wiley.
- Basili, V. The Role of Experimentation in Software Engineering: Past, Current, and Future. In: Maibaum & Zelkowitz (eds.), *Proceedings of the 18<sup>th</sup> International Conference on Software Engineering*, Berlin, Germany, 1996.
- Bell, D.R., Chiang, J., Padmanabhan, V. The decomposition of promotional response: An empirical generalization. *Marketing Science* 18.4, pp. 504-526, 1999.
- Bemmaor, A.C., Mouchoux, D. Measuring the short-term effect of in-store promotion and retail advertising on brand sales: a factorial experiment. *Journal of Marketing Research* 28, pp. 202-214, 1991.
- Blattberg, R.C., Briesch, R., Fox, E.J. How promotions work. *Marketing Science* 14 (3,2), pp. G122-G132, 1995.
- Blattberg, R.C., Wisniewski, K.J. How retail promotions work: empirical results. Working paper No. 42, Center for Research in Marketing, University of Chicago, 1987.
- Boland, R.J. Design in the Punctuation of Management Action. In: Boland (ed.) *Managing as Designing: Creating a Vocabulary for Management Education and Research*, Stanford University Press, 2002.
- Bolton, R. N. The relationship between market characteristics and promotional price elasticities. *Marketing Science*, 8(2), pp. 153-169, 1989.
- Burstein, F. Gregor, S. The Systems Development or Engineering Approach to Research in Information Systems: An Action Research Perspective, In: Hope & Young (eds.), *Proceedings of the 10<sup>th</sup> Australian Conference on Information Systems*, Wellington, New Zealand, pp. 122-134, 1999.
- Caruana, A. Steps in forecasting with seasonal regression: a case study from the carbonated soft drink market. *Journal of Product and Brand Management* 10(2), pp. 94-102, 2001.

- Christen, Markus, Sachin Gupta, John C. Porter, Richard Staelin, Dick R. Wittink Using market-level data to understand promotion effects in a nonlinear model. *Journal of Marketing Research* 34 (3) pp. 322-334, 1997.
- Cole, R., Purao, S., Rossi, M., Sein, M.K. Being proactive: Where action research meets design research, In: *Proceedings of the 26<sup>th</sup> International Conference on Information Systems*, Las Vegas, NV, USA, pp. 325-336, 2005.
- DelVecchio, D., Henard, D.H., Freling, T.H. The effect of sales promotion on post-promotion brand preference: A meta-analysis. *Journal of Retailing* 82.3, pp. 203-213, 2006.
- Farrar, D.E., Glauber, R.R. Multicollinearity in Regression Analysis: The Problem Revisited. *The Review of Economics and Statistics*, 49(1), pp. 92-107, 1967.
- Fildes, R., Goodwin, P., Lawrence, M. The design features of forecasting support systems and their effectiveness. *Decision Support Systems*, 42(1), pp. 351-361, 2006.
- Flores, B. E., Olson, D. L., Wolfe, C. Judgmental adjustments of forecasts: A comparison of methods. *International Journal of Forecasting*, 7, pp. 421-433, 1992.
- Frawley, A., Thearling, K. Increasing customer value by integrating data mining and campaign management software. *Database Management*, 61(10), pp. 49-53, 1999.
- GfK Panel Services Retail NL, *Supermarktkenetallen april 2012* (Powerpoint slides), Dongen, 2012. <http://www.webcitation.org/6HxVVKowU>
- GfK Panel Services Retail NL, *Supermarktkenetallen juli 2010* (Powerpoint slides), Dongen, 2010. <http://www.webcitation.org/6HxWMUz63>
- Goodman, D. A., Moody, K.W. Determining optimum price promotion quantities. *The Journal of Marketing*, pp. 31-39, 1970.
- Goodwin, P., & Wright, G. The limits of forecasting methods in anticipating rare events. *Technological forecasting and social change*, 77(3), pp. 355-368, 2010.
- Green, K. C., Armstrong, J. S. Structured analogies for forecasting. *Monash Econometrics and Business Statistics Working Papers*, number 17/04, 2004.
- Gregor, S., Jones, D. The Anatomy of a Design Theory. *Journal of Association for Information Systems*, Volume 8, Issue 5, pp. 312-335, 2007.
- Guadagni, P.M., Little, J.D.C. A logit model of brand choice calibrated on scanner data. *Marketing science* 2.3, pp. 203-238, 1983.
- Gupta, S. Impact of sales promotions on when, what, and how much to buy. *Journal of Marketing Research* 25 (4), pp. 342-355, 1988.
- Gupta, S. Stochastic models of interpurchase time with time-dependent covariates. *Journal of Marketing Research*, pp. 1-15, 1991.
- Hair, J. F., Ringle, C. M., Sarstedt, M. PLS-SEM: Indeed a silver bullet. *The Journal of Marketing Theory and Practice*, 19(2), pp. 139-152, 2011.
- Harrell, G.D. *Marketing: Connecting with Customers*, ed. 12, *Chicago Education Press*, Chicago, 2010.

- Hevner, A.R., March, S.T., Park, J. and Ram, S. Design Science in Information Systems Research. *MIS Quarterly, Volume 28, No. 1*, pp. 75-105, 2004.
- Hoch, S. J., & Schkade, D. A. A psychological approach to decision support systems. *Management Science*, 42, pp. 51–64, 1996.
- Hoerl, A.E. Kennard, R.W. Ridge Regression: Applications to Nonorthogonal Problems. *Technometrics 12(1)*, pp. 69-82, 1970.
- Iivari, J., Venable, J. Action Research and Design Science Research: Seemingly Similar But Decidedly Dissimilar. In: *Proceedings of the 17<sup>th</sup> European Conference on Information Systems*, Verona, Italy, 2009.
- Järvinen, P. Action research is similar to design science, *Quality & Quantity, Volume 41*, pp. 37-54, 2007.
- Keen, P. G., Scott Morton, M.S. Decision support systems: An organizational perspective. *Addison-Wesley*, Reading, 1978.
- Keim, D. A., Kriegel, H. P. Visualization techniques for mining large databases: A comparison. *Knowledge and Data Engineering, IEEE Transactions on*, 8(6), pp. 923-938, 1996.
- Klein, H.K., Myers, M.D. A Set of Principles for Conducting and Evaluating Interpretive Field Studies in Information Systems. *MIS Quarterly, Volume 23, Issue 1*, pp. 67-94, 1999.
- Kleindorfer, G., O'Neill, L., Ganeshan, R. Validation in Simulation: Various Positions in the Philosophy of Science. *Management Science, Volume 44, Issue 8*, pp. 1087-1099, 1998.
- Kleinmuntz, B. Why we still use our heads instead of formulas: Toward an integrative approach. *Psychological Bulletin*, 107, pp. 296–310, 1990.
- Kmenta, J. Elements of Econometrics. *MacMillan Publishing*, New York, 1971.
- Lee, W.Y., Goodwin, P., Fildes, R., Nikolopoulos, K., Lawrence, M. Providing support for the use of analogies in demand forecasting tasks. *International Journal of Forecasting 23.3*, pp. 377-390, 2007.
- Lemon, K.N., Nowlis, S.M. Developing synergies between promotions and brands in different price-quality tiers. *Journal of Marketing Research*, pp. 171-185, 2002.
- Leone, R.P. Forecasting the effect of an environmental change on market performance: An intervention time-series approach. *International Journal of Forecasting 3.3*, pp. 463-478, 1987.
- Lim, J. S., & O'Connor, M. (1996). Judgmental forecasting with time series and causal information. *International Journal of Forecasting, 12(1)*, 139-153.
- Maddala, G. S., & Lahiri, K. Introduction to econometrics. ed. 2, *MacMillan Publishing*, New York, 1992.
- Mahajan, V., Jain, A.K., Bergier, M. Parameter Estimation in Marketing Models in the Presence of Multicollinearity: An Application of Ridge Regression. *Journal of Marketing Research, 14(4)*, pp. 586-591, 1977.
- Malthouse, E.C. Ridge Regression and Direct Marketing Scoring Models. *Journal of Interactive Marketing 13(4)*, pp. 10-23, 1999.

- March, S.T., Smith, G. Design and Natural Science Research on Information Technology. *Decision Support Systems, Volume 15, Issue 4*, pp. 251-266, 1995.
- Markus, M.L., Majchrzak, A., Gasser, L. A Design Theory for Systems that Support Emergent Knowledge Processes. *MIS Quarterly, Volume 26, Issue 3*, pp. 179-212, 2002.
- McIntyre, S.H., Achabal, D.D., Miller, C.M. Applying case-based reasoning to forecasting retail sales. *Journal of Retailing* 69.4, pp. 372-398, 1994.
- Mela, C.F., Gupta, S., Jedidi, K. Assessing long-term promotional influences on market structure. *International Journal of Research in Marketing* 15.2, pp. 89-107, 1998.
- Montgomery, D.B., Srinivasan, V. An improved method for meta-analysis: with application to new product diffusion. Working Paper, Graduate School of Business, Stanford University, 1996.
- Moon, M. A., Mentzer, J. T., Smith, C. D., Garver, M. S. Seven keys to better forecasting. *Business Horizons*, 41(5), pp. 44-52, 1998.
- Moorman, C., Deshpande, R., Zaltman, G. Factors affecting trust in market research relationships. *The Journal of Marketing*, pp. 81-101, 1993.
- Narasimhan, C., Neslin, S.A., Subrata, K.S. Promotional elasticities and category characteristics. *The Journal of Marketing*, pp. 17-30, 1996.
- Neslin, S.A., Henderson, C., Quelch, J. Consumer promotions and the acceleration of product purchases. *Marketing Science* 4.2, pp. 147-165, 1985.
- Neslin, S.A., Stone, L.G.S. Consumer inventory sensitivity and the postpromotion dip." *Marketing Letters* 7.1 pp. 77-94, 1996.
- Nijs, V.R., Dekimpe, M.G., Steenkamps, J.B.E., Hanssens, D.M. The category-demand effects of price promotions. *Marketing science*, 20(1), pp. 1-22, 2001
- Nunamaker, J. Chen, M., Purdin, T.D.M. Systems Development in Information Systems Research. *Journal of Management Information Systems, Volume 7, Issue 3*, pp. 89-106, 1991a.
- O'Brien, J A. Introduction to information systems: essentials for the e-business enterprise. *McGraw-Hill*, Boston, MA, 2003.
- Olsen, J. E., Granzin, K. L. Economic development and channel structure: A multinational study. *Journal of Macromarketing*, 10(2), pp. 61-77, 1990.
- Osinga, E. C., Leeflang, P. S., Wieringa, J. E. Early marketing matters: A time-varying parameter approach to persistence modeling. *Journal of Marketing Research*, 47(1), pp. 173-185, 2010.
- Pauwels, K., Hanssens, D.M., Siddarth, S. The long-term effects of price promotions on category incidence, brand choice, and purchase quantity. *Journal of Marketing Research*, pp. 421-439, 2002.
- Poh, H.L., Yao, J., Jasic, T. Neural networks for the analysis and forecasting of advertising and promotion impact. *International Journal of Intelligent Systems in Accounting, Finance & Management* 7(4), pp. 253-268, 1998.
- Pole, A., West, M., Harrison, J. *Forecasting and Assessing Structural Change: Campaign Effects*. Duke University, 1991.

- Power, D. J. Specifying An Expanded Framework for Classifying and Describing Decision Support Systems. *Communications of the Association for Information Systems*, 13(1), pp. 158-166, 2004.
- Power, D. J., Kaparathi, S. Building Web-based decision support systems. *Studies in Informatics and Control*, 11(4), pp. 291-302, 2002.
- Purao, S. Design Research in the Technology of Information Systems: Truth or Dare. *GSU Department of CIS Working Paper*, Atlanta: Georgia State University, 2002.
- Rainer, R. K., Cegielski, C.G.. Introduction to information systems: enabling and transforming business (4th ed.). *Wiley*, 2010.
- Raju, J.S. The effect of price promotions on variability in product category sales. *Marketing Science* 11.3, pp. 207-220, 1992.
- Reinmuth, J.E., Geurts, M.D. A Bayesian Approach to Forecasting the Effects of Atypical Situations. *Journal of Marketing Research*, pp. 292-297, 1972.
- Robie, C., Ryan, A.M. Effects of nonlinearity and heteroscedasticity on the validity of conscientiousness in predicting overall job performance. *International Journal of Selection and Assessment*, 7(3), pp. 157-169, 1999.
- Rossi, M., Sein, M. Design Research Workshop: A Proactive Research Approach. *Proceedings of the 26<sup>th</sup> Information Systems Research Seminar*, Haikko, Finland, 2003.
- Sanders, N. R., & Ritzman, L. P. Judgmental adjustment of statistical forecasts. In J. S. Armstrong (Ed.), *Principles of Forecasting* pp. 405–416, *Kluwer Academic*, Boston, 2001.
- Schank, R. C. *Dynamic theory: A theory of reminding and learning in computers and people*. *Cambridge University Press*, Cambridge, 1982.
- Schmidt, F. L., & Hunter, J. E. Eight common but false objections to the discontinuation of significance testing in the analysis of research data. In: *What if there were no significance tests?* Harlow, L.L., Mulaik, S.A., Steiger, J.H. Eds London: Lawrence Erlbaum, pp. 37-64, 1997.
- Shao, Y. E. Multiple intervention analysis with application to sales promotion data. *Journal of Applied Statistics* 24(2), pp. 181-192, 1997.
- Shaw, M. J., Subramaniam, C., Tan, G. W., Welge, M. E. Knowledge management and data mining for marketing. *Decision support systems*, 31(1), pp. 127-137, 2001.
- Silver, M. S., Markus, M. L., & Beath, C. M. The information technology interaction model: a foundation for the MBA core course. *MIS quarterly*, 19(3), pp. 361-390, 1995.
- Simon, H.A. *The Sciences of the Artificial* (3<sup>rd</sup> ed.), *MIT Press*, Cambridge, MA, 1996.
- Sprague, R. H. A framework for the development of decision support systems. *MIS quarterly*, 4(4), pp. 1-26, 1980.
- Timmermans, H.J.P. Multiattribute shopping models and ridge regression analysis. *Environment and Planning A* 13, pp. 43-56, 1981.
- Van Heerde, H.J., Gupta, S., & Wittink, D.R. Is 75% of the sales promotion bump due to brand switching? No, only 33% is. *Journal of Marketing Research*, pp. 481-491, 2003.
- Van Heerde, H.J., Leeflang, P.S.H., Wittink, D.R. Decomposing the sales promotion bump with store data. *Marketing Science* 23.3, pp. 317-334, 2004.

- Van Heerde, H.J., Leeftang, P.S.H., Wittink, D.R. The estimation of pre-and postpromotion dips with store-level scanner data. *Journal of Marketing Research* 37.3, pp. 383-395, 2000.
- Vinod, H.D. A survey of ridge regression and related techniques for improvements over ordinary least squares. *The Review of Economics and Statistics*, 60(1), pp. 121-131, 1978.
- Walls, J.G., Widmeyer, G.R., El Sawy, O.A. Building an Information System Design Theory for Vigilant EIS. *Information Systems Research, Volume 3, Issue 1*, pp. 36-59, 1992.
- Wansink, B., Deshpandé, R. "Out of sight, out of mind": Pantry stockpiling and brand-usage frequency. *Marketing Letters* 5.1, pp. 91-100, 1994.
- Webby, R., O'Connor, M., Edmundson, B. Forecasting support systems for the incorporation of event information: An empirical investigation. *International Journal of Forecasting*, 21, pp. 411-423, 2005.
- Zelkowitz, M., Wallace, D. Experimental Models for Validating Technology. *IEEE Computer, Volume 31, Issue 5*, pp. 23-31, 1998.
- Ziliak, S. T., & McCloskey, D. N. The cult of statistical significance: How the standard error costs us jobs, justice, and lives. *University of Michigan Press*. Ann Arbor, MI, 2008.

## Appendix A

The tool is developed using the Django web framework (v1.4.0), implemented in the programming language Python (v2.7.3). The following is a list of all the additional packages, extensions and libraries that have been used and are required to deploy and run the tool:

**Table 5: Overview used software**

Name	Version	Description
Daterange Filter	0.1.1	Allows to filter by a custom date range (From date, To Date) [ <a href="http://pypi.python.org/pypi/django-daterange-filter/0.1.1">pypi.python.org/pypi/django-daterange-filter/0.1.1</a> ]
Grappelli	2.4.2	Alternative Django Admin Page to manage data, grid-based look & feel. [ <a href="http://www.grappelliproject.com">www.grappelliproject.com</a> ]
jQuery 1.9.1	1.9.1	JavaScript library for event handling, animation and Ajax. [ <a href="http://www.jquery.com">www.jquery.com</a> ]
MPTT	0.5.2	Modified Preorder Tree Traversal' technique for storing hierarchical data in a database. [ <a href="http://django-mptt.github.com/django-mptt">django-mptt.github.com/django-mptt</a> ]
NumPy	1.6.2	Library extension to Python, supports large, multi-dimensional arrays and matrices and manipulations. [ <a href="http://www.scipy.org">www.scipy.org</a> ]
Reversion	1.6.3	Version control; roll back to historic point, recover old models, track changes. [ <a href="http://www.github.com/etianen/django-reversion">www.github.com/etianen/django-reversion</a> ]
Scikit-learn	0.11	Utilizes NumPy & SciPy to provide machine learning algorithms for data mining and analysis. [ <a href="http://www.scikit-learn.org">www.scikit-learn.org</a> ]
SciPy	0.11.0	Library depending on NumPy with modules for e.g. statistics, optimization, numerical integration & linear algebra. [ <a href="http://www.scipy.org">www.scipy.org</a> ]
SetupTools	0.6c11	Enables loading reusable, pluggable, egg-based applications without changing Django source code. [ <a href="http://pypi.python.org/pypi/setuptools">pypi.python.org/pypi/setuptools</a> ]
Statsmodels	0.4.3	Provides complement to SciPy statistical computations [ <a href="http://statsmodels.sourceforge.net/">http://statsmodels.sourceforge.net/</a> ]

## Appendix B

**Table 6: Structural Testing Python methods**

#	Method	Class	Operation description	Iteration #			
				1	2	3	4
1	forthcom_events()	Event	Generates list of scheduled events	✓	✓	✓	✓
2	forthcom_events()	Event	Highlight soon approaching events YELLOW or RED	✗	✓	✓	✓
3	past_events()	Event	Generates list of previous promotions with missing realized sales data	✓	✓	✓	✓
4	save()	Event	Calculates and saves price reduction	✓	✓	✓	✓
5	response_add()	Eventadmin	Adds new event from form	N/A	✓	✓	✓
6	response_change()	Eventadmin	Applies changes to event from form	N/A	✗	✓	✓
7	lookups()	NumRangeListFilter	Sets a user-defined numerical filter	N/A	✓	✓	✓
8	queryset()	NumRangeListFilter	Returns filtered queryset	N/A	✗	✓	✓
9	init()	DateRangeForm	Constructs form for date range filter	N/A	✓	✓	✓
10	get_form()	DateRangeFilter	Retrieves form parameters	N/A	✓	✓	✓
11	queryset()	DateRangeFilter	Returns filtered queryset	N/A	✗	✓	✓
12	get_context_data()	EventAnalytics	Defines reference class for selected retailer and product	✓	✓	✓	✓
13	getattr_related()	EventAnalytics	Retrieves attribute values for selected reference class	✓	✓	✓	✓
14	prepare_event_dict()	EventAnalytics	Only selects events with sales for analysis	✗	✓	✓	✓
15	prepare_features()	EventAnalytics	Prepares matrix for analysis	✓	✓	✓	✓
16	make_regression()	EventAnalytics	Performs analysis and returns parameters	✗	✓	✓	✓
17	combine_attr_coef()	EventAnalytics	Matches regression results to attributes	N/A	✗	✓	✓

N/A= Not Applicable/Not yet implemented, ✗ = not functioning properly, ✓ = successfully implemented

**Table 7: Structural Testing JavaScript Methods**

#	Method	Operation description	Iteration #			
			1	2	3	4
1	get_numeric_features()	Retrieves numeric attributes	N/A	✓	✓	✓
2	get_feature_by_name()	Retrieves textual attributes	N/A	✓	✓	✓
3	get_feature_coefs()	Retrieves coefficients array	N/A	✓	✓	✓
4	get_attr_val()	Retrieves attribute values	N/A	✓	✓	✓
5	get_feature_val()	Alternative to get_attr_val for complex attributes, or else calls get_attr_val()	N/A	✗	✓	✓
6	get_features_vals_arr()	Retrieves array with attribute values, calls get_feature_val()	N/A	✗	✓	✓
7	get_plot_points()	Determines plot points for graph, calls get_feature_by_name(), get_feature_coefs(), get_features_vals_arr()	N/A	✗	✓	✓
8	plot_feature()	Renders graphs, calls get_feature_by_name(), get_plot_points()	N/A	✗	✗	✓
9	Slider/Graph Constructor	Renders sliders, calls plot_feature()	N/A	✗	✗	✓

N/A= Not Applicable/Not yet implemented, ✗ = not functioning properly, ✓ = successfully implemented

## Supplemental information to Table 6 & Table 7

### Iteration #1

#### Python:

For the first iteration the focus lies on making the dashboard partially functional and start implementing analysis functionality. On the dashboard events are not always highlighted when supposed to. Saving or editing promotions is not implemented yet, and neither are the search, filter and graphical functions. As for the analysis part, most methods functioning properly except data is not filtered (14) to exclude promotions without realized sales and the regression analysis (16) is giving errors.

#### JavaScript:

For the first iteration no JavaScript methods have yet been implemented.

### Iteration #2

#### Python:

For the second iteration the dashboard is now fully functional, except for when editing a promotion (6) which is giving an error. The search and filter functions are implemented, but in some cases return wrong data sets (8, 11). The methods regarding regression analysis seem to function properly, except for matching regression results to attributes (17).

#### JavaScript:

All JavaScript methods are implemented, but only 1-4 are functioning properly. When testing we get errors for 5 and 6, but because 7 and 8 (indirectly) call these methods they also give errors. We have to examine whether the errors thrown in 7 and 8 are caused by internal code or because of the faulty methods 5 & 6, 9 calls 8, and therefore also throws errors.

### Iteration #3

#### Python:

All Python methods are now functioning properly.

#### JavaScript:

Only 8 & 9 are still giving errors and/or not functioning as we wish. Only one graph is rendered and the sliders are not functioning.

### Iteration #3

#### Python:

All Python methods are still functioning properly.

#### JavaScript:

All JavaScript methods are now functioning properly. The tool as a whole is functioning well.

**Table 8: Functional Testing**

#	Function	User Input	Desired Result	Iteration #			
				1	2	3	4
1	List scheduled events on Dashboard	N/A	Presents a list of forthcoming promotions on the Dashboard	✓	✓	✓	✓
2	List previous events w/o sales on Dashboard	N/A	Presents a list of previous promotions with missing realized sales data	✓	✓	✓	✓
3	Highlight event	N/A	Highlights promotion YELLOW or RED if deadline is within 4 or 2 weeks, respectively	✗	✓	✓	✓
4	Remove event once sales is added	Add realized sales to promotion	After user adds realized to a promotion, it is removed from the list (see #2) flag as 'usable'	✗	✓	✓	✓
5	Add event	Fill out form	New promotion is created and added to the database, flag as 'proposed'	N/A	✓	✓	✓
6	Edit event	'Edit' button; edit values in form	Carries through changes to promotion values in database	N/A	✗	✓	✓
7	Input validation error	Wrong data type input	Returns error in case of invalid user input	N/A	✗	✓	✓
8	Filter promotions	Select criteria for filter	Searches for promotions based on provided criteria	N/A	✗	✓	✓
9.1	Run Analysis	'Analyze' button	Runs analysis for selected promotion	✗	✓	✓	✓
9.2	Present Analysis results	'Analyze' button	Presents analysis results (see #9.1) in visual graphs	N/A	✓	✓	✓
10.1	Adjust values in analysis results	Adjust sliders	Adjusts attribute values	N/A	✗	✓	✓
10.2	Adjust graphs	Adjust sliders	Adjusts graphs accordingly (see #10.1)	N/A	N/A	✗	✓

N/A= Not Applicable/Not yet implemented, ✗ = not functioning properly, ✓ = successfully implemented

## **Supplemental information to Table 8**

### ***Iteration #1***

For the first iteration the dashboard is implemented, which is the first web page a user sees when he browses to the tool's location on a web server. The corresponding functions are 1, 2, 3 and 4. Functions 1-3 require no user input, but are executed when loading the page. The latter two do not function properly: promotions are not always appropriately highlighted and promotions are not commended to the database as 'usable' after the sales have been provided. 9.1 is implemented but analysis is not fully functional yet as for now only a dummy page exists.

### ***Iteration #2***

During the second iteration 3, 4 & 9.1 have been repaired, and we successfully added new functions 5 and 9.2. Other functions were also added, but not entirely with good results. Editing events resulted in an error message (6), input validation was not present for all the fields (7), not all possible filter criteria were present (8) and adjusting attribute values using the JavaScript sliders didn't function properly (10.1).

### ***Iteration #3***

During the third iteration all the previous problems were repaired and everything was functioning properly except for the graphs (10.2). The graphs were implemented in the previous iteration for the first time and don't adjust well and lack a confidence interval.

### ***Iteration #4***

All functions now operate appropriately and yield the desired result. The tool as a whole is also functioning well.