Modelling consumer directed substitution

Vaagen, Hajnalka and Wallace, Stein W. and Kaut, Michal

in: International Journal of Production Economics. See also BibTeX entry below.

BibTeX:

```
@article{VaagenEA08,
  author = {Vaagen, Hajnalka and Wallace, Stein W. and Kaut, Michal},
  title = {Modelling consumer directed substitution},
  journal = {International Journal of Production Economics},
  year = {2011},
  volume = {134},
  pages = {388--397},
  number = {2},
  note = {special issue on Robust Supply Chain Management},
  doi = {10.1016/j.ijpe.2009.11.012}
}
```

Modelling consumer directed substitution

Hajnalka Vaagen[1], Stein W. Wallace[2], Michal Kaut[3]

Abstract

We discuss the challenges and difficulties arising when approaching and modelling the consumer directed substitution problem in quick response supply chains. Further, we propose heuristic solutions suited for large problems with complex uncertainty and dependency patterns. Despite the single-period newsvendor model we use, our substitution process is an approximation of the dynamic product choice. Substitution fraction estimation and inventory/assortment optimisation are discussed simultaneously, having a starting point in decision-independent substitution preferences, reflecting qualitative understanding of the market drivers; this, in order, leading to increased robustness in assortment planning. Factual substitution is an outcome of the optimisation process, constrained by the available substitutes and unfulfilled demand.
Despite being unable to fully describe the dependencies among the substitute choice possibilities, our substitution approach, together with the modelling process, allows handling the most important dependencies, such as negatively correlated substitute choice possibilities and positively/negatively correlated first and second choice possibilities.

Keywords: Assortment planning; Substitution estimation: Multi-item newsvendor; Stochastic programming; Simulation; Correlations

[1] Hajnalka Vaagen, Molde University College, P.O.Box 2110, NO-6402 Molde, Norway, Phone: +4771214000, Fax: +4771214100, e-mail: hajnalka.vaagen@himolde.no

[2] Stein W. Wallace, Department of Systems Engineering and Engineering Management, Chinese University of Hong Kong, Shatin, N.T., Hong Kong. Phone: (852)-2609-8318, Fax: (852)-2603-5505, e-mail: swallace@se.cuhk.edu.hk

[3] Michal Kaut, Molde University College, P.O.Box 2110, NO-6402 Molde, Norway, Phone: +4771214000, Fax: +4771214100, e-mail: michal.kaut@himolde.no

1. Problem statement and literature review

Supplier and retail assortment planning includes simultaneous decisions on the items to include in the product portfolio and the corresponding inventory levels, so as to maximize the portfolio profit, or minimize profit risk in the more recent literature. Assortment planning is especially difficult in volatile environments with substantial levels of demand uncertainty combined with high product variety – such as fashion apparel and sports equipment and apparel. A commonly applied strategy to hedge against uncertainty in future sales is to offer substitutable products. The idea of substitution originated in the economics literature, with a focus on how to price to avoid undercutting competitors until equilibrium prices equalled marginal costs; the Bertrand paradox. However, in highly dynamic and volatile environments, the value offered to customers is largely qualitative in nature, defined by for example trends and brand name strengths. Products are more or less homogeneous with regard to price and technical attributes. Substitution, hence, manifests structural characteristics that differ from those observed in stable environments. Therefore, focus on demand uncertainty and substitution caused by the less measurable qualitative factors is required. *This paper discusses how to model consumer directed substitution in quick response environments, while focusing on the uncertainties and dependencies arising from qualitative aspects of the problem at hand.*

Consumers' purchasing decisions are affected by the variants in the portfolio as well as their inventory levels (provided they plan to purchase more than one unit). Cost, selling prices, and technical attributes are rather homogeneous within a group where substitution naturally arises (such as watersport jackets for men from a particular supplier); the differences among variants are minimal. Non-measurable qualitative aspects, on the other hand, are heterogeneous, but difficult to quantify. Across groups the products differ significantly and we assume no substitution (such as jackets do not substitute trousers). The retailer chooses inventory levels before observing demand. Heterogeneous customers are then visiting the store in random order, looking for their first preference. When the first preference (i) is carried by the store but it is out of stock at the moment of purchase *or* (ii) not carried by the store, the customer might settle for the available substitute giving highest utility or choose not to purchase at all. The choice process is repeated until there are no substitutes with utility higher than that of not buying. End customer arrivals are random and independent over time; hence stockout and substitution can take place at any point in time, and the problem is dynamic in nature.

In the remaining of this section we discuss relevant work in the area. In Section 2 we present our approach and method for estimating substitution shares. The substitution problem is modelled and arising challenges emphasized in Section 3. Test cases are given. We conclude in Section 4.

Literature review

There is a great amount of work on newsboy models with substitutable demand. Kök et al. (2006) provide an extensive literature review and outline of industry practice on assortment planning in retailing; with substitution as an important strategy in this setting. For single period two-item newsvendor models with substitution see Parlar and Goyal (1984), Pasternack and Dresner (1991), Gerchak et al. (1996), Bitran and Dasu (1992), and Khouja et al. (1996). Multi-item newsvendors with substitution are discussed by for example Bassok et al. (1999), Rajaram and Tang (2001), and Netessine and Rudi (2003). Numerical tests to the problem are few, and mostly limited to two items. Rajaram and Tang (2001) provide a multi-item example; however, the problem complexity is reduced by approximating correlations and substitution fractions with average values, and assuming normality where the data actually indicates non-normal distributions.

Before discussing some central papers on modeling the problem and estimating substitution shares, we point out that in much of this literature there is inconsistency in the treatment of consumer directed and manufacturer directed substitution; inconsistency in what the authors claim to formulate and solve and what they actually do. For the distinction between *manufacturer directed* and *consumer directed substitution* we refer to Mahajan and Van Ryzin (2001): While the manufacturer or supplier may choose to fill demand for one product with inventory of another product to avoid stock-out, substitution in retail settings is not directed by the focal part; decisions are made by a large number of independently-minded and self interested consumers. A retailer can only *indirectly* affect his end-consumers' decisions by his own inventory and marketing policy. Given Mahajan and Van Ryzin's distinction between producer and consumer directed substitution, we point out that in several industrial arenas the substitution pattern is a mix of these major patterns. For example, in the highly uncertain fashion and sports equipment and apparel industry, the steadily more dominating *lean retail ordering strategy* (Abernathy et al., 2004) implies reduced supplier power to influence customer choices, and hence more consumer directed substitution. Previous studies on retailer reactions to supplier stockouts have indicated that a non-negligible portion of retail customers seem to switch to other suppliers or to cancel the purchase all together (Campo et al., 2004), when the first preference is out of stock. However, when the first preference is out of stock, the supplier can, to some extent, control which customers should be served first, based on the retailer customer's strategic importance and the supplier's brand name strength; indicating *partially consumer directed substitution*. This particular case is further discussed later in the paper.

Khouja *et* al. (1996) use a Monte Carlo simulation to identify an optimal solution to a two-item newsboy problem with correlated individual item demands and substitutability. Bassok *et* al. (1999) develop a full downward substitution model for multiple classes of products. A greedy allocation policy is used to obtain an optimal solution when demand is realized at the start of the planning period. The individual demands are assumed to be independently distributed. Rajaram and Tang (2001) develop a service rate heuristic to estimate the *effective demand* D_i^e under substitution (consisting of the original demand D_i for the product and the substitution demand *from* other items), and to solve the extended problem. This heuristic is further evaluated by approximating the upper bound profit using a Lagrangian dual based procedure. We observe that (1) the authors do not limit the sum of substitution fractions (the portion of unmet demand for a product *j* that can be satisfied from leftover inventory of product *i*.) over a particular item to one, and hence allow for several items to "fully" substitute one particular item. Further, (2) the sales for item *j* coming from own demand *j plus* all *i* sales generated from substituting the unmet demand for *j,* are not constrained to the available demand for *j*. These two observations together can result in total sales generated by item *j* (both direct sales and sales arising from unsatisfied demand for *j*) exceeding the demand for item *j*. It is unclear to us whether this actually happens, as the authors only state the model for two items where the issue does not arise. The most natural extension $D_i^e = D_i + \sum_{j \neq i} a_{ji}(D_j - Q_j)^+$ does exhibit the above problem, where Q_j and the term $(D_j - Q_j)^+$ define order quantities and the unsatisfied demand respectively for item *j*. Netessine and Rudi (2003) solve a similar multi-item problem under centralized and decentralized management strategies, and analytically confirm the results provided by Rajaram and Tang. They limit the sum of substitution fractions over a particular item to one, and develop first-order conditions to estimate optimal order levels.

Consumer choice models, such as the multinomial logit (MNL) and locational choice models (logit, probit) have been used by researchers in recent years to plan the optimal assortment with substitutable products. See among others Mahajan and Van Ryzin (2001) and Gaur and Honhon (2006). Despite these models' ability to capture the dynamic customer choice process, they have some severe limitations which make them less useful in agile environments facing complex uncertainties and dependencies. The most important is the IIA property of MNL models (Independence of Irrelevant Alternatives), which requires that the choice possibilities are independent of other alternatives. Consumer choice models also assume "to know" which items will become popular (see also Van Ryzin and Mahajan, 1999); a rather unrealistic assumption in quick response environments. Gaur and Honhon (2006) recognize that knowledge about the most preferred choice makes the substitution problem less difficult, and suggest generalizing their model so as to relax this assumption. However, the extension is not stated. Detailed discussion on the effect of this assumption can be found in Vaagen and Wallace, 2008. Further, locational choice models assume that total demand for an assortment that covers the entire attribute space is the same regardless of the number of products in the assortment; similar in spirit with the assortment risk model developed by Vaagen and Wallace (2008) and the present paper. In contrast, demand always increases with variety in the MNL model. For further shortcomings of the MNL models we refer to Kök et al. (2006).

Given the underlying assumptions, Mahajan and Van Ryzin (2001) achieve nearly optimal solution to the single-period newsboy-like model where heterogeneous retail customers dynamically substitute among available variants if the most preferred variant is out of stock. Customer decisions are based on maximizing utility and a sample path gradient algorithm is applied to compute inventory levels. Poisson customer arrivals are simplified by normal approximations. The Mahajan and Van Ryzin problem is similar in spirit with that of Noonan's (1995), although the model formulations and analysis differ. Noonan assumes primary demand realized in the "first stage" and substitution demand in the "second stage". The work is limited to only a few variants and only allows for one substitution attempt; in contrast to Mahajan and Van Ryzin with unlimited number of attempts.

Available methods for estimating substitution probabilities are based on inventory and sales transaction data. For a review see Kök et al. (2006). Anupindi et al. (1998) describe a method for estimating stock-out based substitution proportions based on inventory-transactions data. They show that the timing of the stockouts and the sales volumes before and after those times are actually sufficient statistics, and tracing each transaction is unnecessary. DeHoratius and Raman (2004) show that empirical inventory data may not be accurate, and that some retailers do not even track inventory data. Hence, sales data may be the only available transaction data in some situations. For a method of estimating substitution rates by the use of sales transaction data see Kök and Fisher (2007). The substitution fractions established by these methods then enter the single-period optimization models as exogenous parameters. Numerous multi-period newsvendor models are based on information updates between the stages; however, due to the problem complexity they are not connected to substitution. In rapidly changing environments optimization usually happens at the beginning of the assortment/inventory planning season. For a large share of industrial actors, waiting for sales data required for reliable demand and substitution forecasts is not compatible with the supply chain flexibility. In addition, existing assortment planning models, including those with substitutable demand, are still static without changes in portfolio as time goes by (see a review on assortment planning by Kök et al., 2006). Even so, the methods for estimating substitution probabilities require real-time inventory or sales transaction data, and a stockout usually occurs in the selling season at a later point in time than the optimisation. As far as we can understand from the literature mentioned above, the estimation of substitution probabilities is not directly connected to the newsvendor based inventory and assortment planning models. This, in order, might lead to inconsistency between the static nature of the models, and the *dynamic* data required for estimating the substitution shares. As opposed to the academic environment, innovative high-fashion companies like Zara and Mango from Spain or World Co. from Japan have managed to create highly responsive supply chains that even allow for quick assortment updates, *based on qualitative understanding of the market drivers.* It is however unclear how such important qualitative knowledge can be quantified and integrated in the existing analytical and numerical models.

2. Substitution approach and estimation of substitution probabilities

In the literature part we pointed out possible inconsistencies between the assortment planning models' static nature and the dynamic data required to estimate substitution shares by the available methods. To avoid this, it is our view that inventory and assortment optimization models should be discussed simultaneously with substitution fraction estimation. In the present work *we attempt to diminish the distance between the academics and industry practice, by (**i**) approaching substitution estimation and inventory planning simultaneously, (**ii**) pointing out the difficulties that arise when doing so, and (**iii**) suggesting models to integrate qualitative understanding and knowledge.*

We approach the problem on market level and provide an exogenous demand model. See Kök et al. (2006) for a review of the three demand model approaches to the assortment planning problem; multinomial logit MNL, locational choice models, and exogenous demand models. Contrary to MNL models, individual customer preferences are not directly approached by us but an aggregated product level demand is established.

Further, we define a substitution measure that is decision independent and integrates qualitative understanding of market drivers. We define the *a priori substitutability* $\alpha_{ij} \in [0,1]$, as the portion of customers willing to replace item *j* with item *i*. We see a priori substitutability as the grade of similarity between the products with regard to the trend drivers, and it can be defined by a multidisciplinary team consisting of designers, product developers, trend analysts and operational analysts. This measure is different from the *factual substitution*; a decision dependent outcome of our optimisation process, constrained by unsatisfied demand and the variants available at the moment of purchase. So we assume that the retailer offers *n* items, all being potential substitutes for each other. The "retailer" in our consumer directed substitution setting can be a retailer (facing end users as customers) or a supplier/manufacturer (facing retailers as customers) responsible for development, production and marketing activities.

For illustration, consider five identical jackets in colours Red, Black, Marine, White and Turquoise. The a priori probabilities are illustrated by Table 1. According to the table, if the first preference is Red, the colours Black, Marine, White and Turquoise are possible substitutes, with probabilities 0.7, 0.4, 0.4, and 0.1, respectively. That is, substitutability and non-substitutability probabilities from Red to Black are 0.7 and 1−0.7 = 0.3 respectively. The substitutability matrix is not required to be symmetric. As an example, consider White and Turquoise as standard and high fashion colours, respectively. It is not unrealistic to assume that customers having Turquoise as first preference will see White as a fairly good substitute. However, customers wanting a subdued look are less likely to substitute to a high fashion colour from the "safe" White one. See highlights in Table 1.

Table 1 Substitutability probabilities α_{ij}

	Red	Black	Marine	White	Turquoise
Red	1	0.7	0.4	0.4	0.1
Black	0.5	1	0.4	0	0
Marine	0.5	0.4	1	0	0
White	0.4	0	0	1	0.2
Turquoise	0.4	0.4	0.4	0.8	1

Observe that the way we have presented the a priori substitutability probabilities, we say nothing about how they relate to each other. We shall discuss these difficulties in more detail in Section 3. Note that individual item demand dependencies – captured by the underlying demand distributions – are not to be confused with dependencies between the substitution possibilities.

Given our approach, the process of the substitution problem defined in Section 1 is the following. (**1**) The retailer estimates the substitutability matrix with α_{ij} values; a measure that 'makes sense' to practitioners and that is consistent with assortment and inventory planning in highly uncertain and dynamic environments. (**2**) Given demand distributions and the substitutability matrix, the retailer chooses assortment; products to include in the portfolio as well as their inventory levels. (**3**) *After* observing inventory (in particular, which substitutes are available at a given point in time), each individual customer, one at a time, chooses first preference *or* settles for a substitute available on stock, or chooses not to purchase at all. To reflect our market level approach, individual customer choices are presented on aggregated level. The outcome of this process is the *factual substitution*, illustrating the share of customers that actually substitute from item *j* to item *i*.

Note that this process describes both the retailer's and its end customers' actions; actions that happen at different points in time and with different information levels. While a retailer's information at the time of planning is given by α_{ij}, the end customers' decisions are also based on the dynamically changing actual inventories. As an illustration of the dynamic consumer choice process, we introduce the *intermediate substitution probability* $\beta_{ij} \in [0,1]$. By summing the intermediate substitution probabilities *from* item *j* to all *available* items *i*, plus the no-substitution option, we end up with 1. Hence, as such, the β's are proper probabilities. Note that while the values of substitutability α_{ij} are maintained, the decision dependent β_{ij} values changes dynamically during the choice process.

As already mentioned, when defining the α's we failed to describe the dependencies among the individual elements. If the first choice is Red, but it is out of stock, and the only available possible substitutes are Marine and White (see Table 1), we do not know if the 40% accepting White as a substitute and the 40% accepting Marine are the same 40%

of the customers, or if those are two distinct groups. Given our idea of how the substitutability matrix comes about, it is not reasonable to try to specify all conditional preference probabilities. We shall discuss this more in Section 3. When defining the *intermediate substitution probability* β_{ij}, on the other hand, we shall assume that the alternative choices are *independent.* That would mean that the probability that a certain customer accept a certain substitute is independent on what other substitutes he accepts.

α_{ij} = substitutability probability from j to i

$I_j(x) =$ set of available substitutes for item j; depends on the inventory of x

$\overline{\beta_j} = \prod_{i \in I_j(x)} (1-\alpha_{ij}) =$ probability of no substitution from item j (independence is used here)

$\beta_{ij} = \dfrac{\alpha_{ij}}{\sum_{i \in I_j(x)} \alpha_{ij}} (1-\overline{\beta_j}) =$ intermediate probability to substitute demand for j by i

Unlike the alphas, the beta probabilities sum up to one:

$$\sum_{i \in I_j(x)} \beta_{ij} + \overline{\beta_j} = (1-\overline{\beta_j}) \sum_{i \in I_j(x)} \frac{\alpha_{ij}}{\sum_{i' \in I_j(x)} \alpha_{i'j}} + \overline{\beta_j} = 1 - \overline{\beta_j} + \overline{\beta_j} = 1$$

For a numerical example, consider the substitutability matrix in Table 1, with Red as the first choice. Assume that Red is not on stock and the only available substitutes are Black and Marine. The purchase/no-purchase substitution probabilities, then, are the following:

$$\Pr\,(\overline{Black \cup Marine}) = (1-0.7)*(1-0.4) = 0.18$$
$$\Pr(Black \cup Marine) = 1 - 0.18 = 0.82$$

The intermediate substitution probabilities on individual item level are then:

$$\Pr(Black) = \frac{0.7}{0.7+0.4} * 0.82 = 0.52,\ \Pr(Marine) = \frac{0.4}{0.7+0.4} * 0.82 = 0.3,$$
$$\Pr(NoPurchase) = 0.18$$

Observe that if Black is the only available substitute, the intermediate substitution probability is the same as the a priori substitutability probability of 0.7.

3. Modelling consumer directed substitution

In this section we shall discuss challenges and difficulties arising from our need to model the factual substitution. In highly dynamic and uncertain environments, demand information often comes at a point in time when changing assortment and production plans is not compatible with existing supply chain capabilities. Besides the market level objectives, substitution is a strategy introduced to hedge against uncertainty in demand; a strategy arising from recognizing the limited value of postponing assortment and inventory decisions until a later point in time. Accordingly, in this work we choose to perform the optimization at the beginning of the retail planning season. Further, we assume that the available information with regard to the substitution is given by α_{ij}.

The substitutability matrix (Table 1) does not fully describe substitution. A total description would require full knowledge about all conditional substitution probabilities. This is from a practical point of view a rather unreasonable requirement. Further, it could only be used properly in a dynamic model where inventories were continuously updated as sales took place. This would be beyond any newsboy-like model. As we shall make a static model, and as we find the substitutability matrix a very reasonable starting point, we need to define a way to find good solutions within such a framework. Waiting with optimisation until the information required to establish inventory-dependent substitution measures becomes available might make sense in more stable environments. However, it is unreasonable in our rapidly changing uncertain context. There are two main reasons for that. First, in industrial settings such as ours, some important decisions must still to be taken in light of uncertainty, with no or limited amount of data available on inventory and sales transactions. We intend, as also pointed out in Section 2, to avoid potential inconsistencies observed in the literature, between a substitution approach that needs dynamic data and the static nature of assortment planning models. Secondly, substitution in our context is an assortment planning strategy to hedge against uncertainty in future sales. If decision makers could wait for information to establish decision dependent substitution shares, they could also establish potentially better direct forecasts; this, potentially diminishing the value of substitution as a hedging strategy.

Despite the single-period model we intend to use, our substitution process is an approximation of the dynamic choice process. The outcome is one single substitution or lost sale; however, it is incorrect to call it a "one substitution attempt" model. As described in the previous section, we reflect individual choices on aggregated level, where each customer chooses first preference or a substitute available on stock, or chooses not to purchase at all. Also, the chance of finding an available substitute increases with the number of "similar" products. Kök and Fisher (2007) refer to the single-attempt model's similarity to the multinomial logit utility based dynamic model established by Mahajan and Van Ryzin (2001), where the substitution rate is inventory dependent. They say that "*Although δ is fixed in our model, if a consumer cannot find her second favourite either, the sales is lost; and that is equivalent to less-frequent substitution when the set of available items is smaller*"; where δ is the fixed substitution rate in their single-attempt model.

3.1 Stochastic programming (SP) formulation

First, we present a simple two-stage stochastic program, basically taking the classical newsvendor model and adding variables for substitution. We maximize expected assortment profit (over all defined scenarios of the uncertain demand) coming from ordinary and substitution sales, and salvage of the leftover inventory. The first stage consists of the scenario independent production decisions, while the second stage, i.e. the stage after demand has been observed, includes variables for the direct and substitution sales. For more information about two-stage stochastic programming, see for example Kall & Wallace (1994). In this particular two-stage SP formulation we directly use the decision independent α_{ij} substitutability values as input parameters, defining the substitution pattern. The intermediate substitution probability values β_{ij} imply a dynamic procedure we cannot model in the context of a two-stage program. Also, this would be beyond any newsboy-like model. The observed substitution is the outcome of our model.

Sets:
S – set of demand scenarios;
I – set of items in the reference group portfolio
If we do not state otherwise, we use indices with the following meaning $i, j \in I \text{ and } s \in S$

Variables
x_i = production of item i
y^s_i = sale for item i in scenario s
z^s_{ij}= substitution sale of item i, satisfying excess demand of item j in scenario s
zt^s_i = substitution sale of item i, satisfying excess demand from all j's in scenario s
w^s_i = salvage quantity for item i in scenario s

Parameters
d^s_i = demand for item i in scenario s
p^s = probability of scenario s
v_i = selling price for item i
c_i = purchasing cost for item i
g_i = salvage value for item i
α_{ij} = substitutability probability; the probability that item *j* can be replaced by item *i,* given that *i* is the only available substitute (on stock) for *j*; $\alpha_{ij} \in [0,1]$

$$Maximize\ Expected\ profit = \sum_{s \in S} p^s \sum_{i \in I} \left(-c_i x_i + v_i y^s_i + v_i\, zt^s_i + g_i w^s_i \right) \qquad (1)$$

Subject to:

$$y^s_i + \sum_{j \in I, j \neq i} z^s_{ji} \leq d^s_i \quad \forall i \in I, s \in S \qquad (2)$$

$$z_{ij}^s \le \alpha_{ij}\left(d_j^s - y_j^s\right) \quad \forall i, j \in I, i \ne j, s \in S \tag{3}$$

$$zt_i^s = \sum_{j \in I, j \ne i} z_{ij}^s \qquad \forall i \in I, \ s \in S \tag{4}$$

$$w_i^s = x_i - (y_i^s + zt_i^s) \quad \forall i \in I, s \in S \tag{5}$$

$$x_i \ge 0 \quad \forall i \in I \tag{6}$$

$$y_i^s \ge 0 \quad \forall i \in I, s \in S \tag{7}$$

$$z_{ij}^s \ge 0 \quad \forall i, j \in I, i \ne j, s \in S \tag{8}$$

$$zt_i^s \ge 0 \quad \forall i \in I, s \in S \tag{9}$$

$$w_i^s \ge 0 \quad \forall i \in I, s \in S \tag{10}$$

We maximize expected assortment profit from ordinary sales, substitution sales and salvage, over all items and all scenarios, using Equation (1). Equations (2) state that item i sales – coming from primary demand for i plus sales of all j that substitute unmet demand for i – cannot exceed total demand for item i. Note that this constraint can be re-organized as (2a) $\sum_{j \in I, j \ne i} z_{ji}^s \le d_i^s - y_i^s \quad \forall i \in I, \forall s \in S$

stating that substitution sales from item i cannot exceed available unsatisfied demand for i. (3) defines the upper bound for substitution sale of item i for item j; that is, excess demand for item j with given substitutability probability α_{ij}. (4) Calculates substitution sale of item i from all j's. Equation (5) gives the salvage quantity, the quantity of item i left after satisfying primary demand and substitution demand from all j. (6), (7), (8) and (9) define the non-negativity of the variables. The constraints (5) and (10) imply that the substitution sale of item i is limited to the remaining supply of the item; that is, $zt_i^s \le x_i - y_i^s \quad \forall i \in I, s \in S$

The model has been implemented in AMPL with CPLEX as the underlying solver, for a test case of 15 items; see http://www.ilog.com/ for details on both systems. Since this is a linear-programming model of a moderate size, the solution times are negligible.

PROBLEMS ARISING FROM THE SP FORMULATION

A. Optimal allocation between direct and substitution sales

The presented formulation leads to optimal allocation between *direct* and *substitution sales*; manifesting manufacturer directed substitution. In other words, the optimal solution implies that we can tell the customers how much "first preference" and how much "second preference" they can buy. If we imagine customers arriving independently to a store, one by one, it is obvious that this is not the case. However, this is actually the outcome of numerous existing substitutable multi-item newsboy problems. Netessine and Rudi (2003) enforce direct sales as first decision. However, an important assumption behind this is that there is only one substitution attempt; an assumption which we find unreasonable for our problem. Further, the authors disregard the dynamics of the problem. If we think dynamically, using the model proposed by Netessine and Rudi (2003), we still may run into the problem of getting substitution while primary demand is coming. We discuss this in more detail in Section 3.3.

Hence, to be able to apply our optimization model in practice, we would either have to be able to directly decide which customers should switch their demand, or at least to be able to identify the different types of customers and be willing to not sell an item to a customer, even if we have it in stock. If neither of these strategies is possible, the solution is most likely unrealistic and the objective function would form an upper bound on the true profit of the consumer directed substitution; with an objective value similar to that of the manufacturer directed substitution model.

B. *The effect of the aggregated market level approach in our SP formulation*

In the presented model, the substitutability probability values are interpreted as the share of customers accepting a particular item as second choice, and the model is unable to distinguish among individuals. Hence, as long as there are available substitutes and unsatisfied demand, our model provides an optimistic approach by treating customers *as if they were different with regard to preferences*. For illustration, take the following example from Table 1: the colours Marine and White are substitutes for Red, both with probability 0.4. Assume that there is some unsatisfied Red demand. In our market level approach we are unable to identify whether the same 40% of customers or two different groups prefer these two substitutes, and our model will assign 40% of the unsatisfied demand to the colour Marine and the another 40% to the colour White; covering 80% of the unfulfilled demand, despite the fact that this is maybe not possible in the actual population of customers. This assumption obviously leads to an upper bound on the true expected profit.

C. *Dependencies between the substitution possibilities in our SP formulation. The connection between individual item correlations and substitutability*

Our approach does not directly capture the dependencies between the substitutes, nor the connection between substitution and demand correlations. However, some of the dependencies can actually be captured by the individual α_{ij} values and by the modelling process. Below we discuss these issues.

(i) Positively correlated first and substitute preferences

Take the example of two jackets in different styles but in the same high-fashion colour turquoise. Assume that "colour" is an important trend driver, with about 50% impact on customer choice. Due to the colour, hence, the two items can replace each other with substitutability probability about 0.5. Further, if turquoise becomes a trendy colour the demand goes up for both items. This dependency can be expressed by a positive correlation between the two demands. Accordingly, there is a logical connection between correlations and substitutability fractions, and this understanding can be partially captured by the chosen substitutability values.

Our SP formulation limits the substitution sales to the unsatisfied demand and leftover inventory. When the colour turquoise becomes unpopular, substitution cannot be leveraged on. Despite the high substitution rate, both items face decreased demand with increased probability for leftover inventory; there will be no unsatisfied demand. If the colour becomes popular, both items might face stockout and there will be no possibility for substitution.

(ii) Negatively correlated first and substitute preferences

Consider the following example with white and black shirts. At the time of production decisions, there is no information about the items' behaviour, but we know that only one of them becomes popular. Hence the individual demands are negatively correlated. Further, assume that black turns out to be popular. We know that to some extent the two colours can replace each other, and hence additional whites can be sold. Our question is: *How high substitution fraction is it reasonable to assume?* Recall that only one of them becomes popular. Accordingly, the substitutability probability cannot be high. The unsatisfied demand is more likely to be satisfied from competitors' products. We believe that in the fashion apparel environment, and in other similar cases where trend plays an important role and the products compete for popularity, it is unrealistic to assume high negative correlation paired with high substitution, as assumed in the numerical example provided by Rajaram and Tang (2001). Negatively correlated items with high substitution fraction can be observed for rather standard/household articles, like offering bundles versus individual products; such as shampoo and conditioner offered paired in a package or separately, see Ernst and Kouvelis (1999).

(*iii*) *Negatively correlated second choice possibilities*

Consider two identical jackets in colours Black and Marine, each facing two possible states of the world; popular or unpopular. At the time of production decisions we do not know which one becomes popular. We describe this by negatively correlated individual demands; for example -0.6. Further, assume that whenever one of them is popular, it is also a good substitute for the first preference Red, for example with substitutability probability 0.5. Now, if there is unsatisfied demand for Red, according to our model, the colors Marine and Black face equally high chances of substitution sales in each scenario. This is an unrealistic situation, as one of them turns out to be unpopular with insignificant demand. There might be some substitution to an unpopular item; however, assuming high probability is incorrect. If this was a realistic situation, the problem of uncertainty around the items' popularity could have been solved by offering substitutes that are competing for popularity.

The following *solution* is proposed to this particular case. We let the substitutability probability matrix depend on the state of the world, and define one matrix for each state. The changes in values in the matrices only apply to the substitution choices manifesting the specific behaviour; everything else is equal.

3.2 Stochastic Mixed Integer Programming (MIP) formulations

One way of improving the SP model above, with regard to problem A in Section *3.1*, is to enforce the condition that substitution takes place only after as much as possible of direct demand has been satisfied. This solution, however, has several problems: *firstly*, it does not mitigate the fact that the producer is directing the substitution—while the producer would no longer have the possibility to force customers to buy a substitute instead of their first choice, the choice between the available substitutes is still decided by the producer. *Secondly*, such a condition would require introducing binary variables into the model, making it more difficult to solve. *Thirdly*, and most importantly, it is not obvious that the condition is a good approximation of reality. While it might work in a setting where all the demand is revealed and cleared initially, it does not work in a situation where customers arrive to a retailer and buy one unit of a product at a time. In such a setting, the items can run out of stock at different times, so customers would start switching to other items, even if there are still customers that will want these items as their first choices.

This can be illustrated on a three-item example where we assume that the demand is uniformly distributed over a time period of length 1. Assume the following substitutability probabilities α_{ij}

$$\begin{bmatrix} - & 0.2 & 0.1 \\ 0 & - & 0 \\ 0.1 & 0.1 & - \end{bmatrix}$$

and that we have decided the production x to be 100 for all three items. Assume we are facing scenario demands of 100, 0 and 200 for item 1, 2 and 3, respectively. Then we can see that item 3, on average, will run out of stock at time $t = 0.5$. By then we would have sold (50, 0, 100) units. After that, 10% of the demand for item 3 will be converted to item 1, so the total demand for item 1 during a time interval of Δt would be $100\Delta t$ + $0.1\text{x}200\Delta t$. This will last until we run out of item 1, i.e. $\Delta t = 50/120 = 0.42$. By this time, we would have sold 91.67 of item 1 to satisfy demand for item 1, leaving it with an unsatisfied demand of 100 – 91.67 = 8.33. Of this demand, 0.2x8.33 = 1.67 would get converted into item 2, leaving us with 6.67 of unsatisfied demand. This means that out of the 100 units of demand for item 1, 1.67 were substituted by item 2 and 6.67 units were lost, even though the supply of item 1 was 100 units.

The question is if we can model the above dynamics in a stochastic-programming model. It turns out that the answer is positive, though it means introducing a large number of extra constraints and variables. In particular, it means introducing $|S| \times |I|^2$ indicator (binary) variables, where $|S|$ is a cardinality of set S. This makes the problem tractable only for very small examples; a test case based on Rajaram and Tang (2001), with 7 items and 40 scenarios, needs 1960 binary variables. CPLEX 9 could not solve this problem in two days, after which we aborted the execution. The MIP formulation is complex and it cannot be applied to large cases; hence, we do not present it here.

3.3 A simulation-based optimization approach

While the above model discussion suggests that the problem is a very difficult one, there is a reason to believe that it should be solvable; after all, with 15 items in our test case, we have a problem with only 15 “true” decision variables—all the rest are used only to model the objective function. In other words, all the complications arise from the fact that we try to approximate the dynamics of the second stage (the customer choice) by a mathematical program. It follows that we can solve the problem using a simple simulator of the second stage, i.e. a `black box’ that takes the production vector x as an input and returns the expected profit, calculated using the step-by-step procedure presented in the previous section. What we then need is a so-called *derivative-free solver*, i.e. a solver that can solve problems using only objective function values, not derivatives. This approach has the additional advantage of being able to model the intermediate probabilities β_{ij} presented in Section 2, and hence partially avoid the problem arising from our market level approach (given in B, Section 3.1). Due to strong (and possibly unrealistic) independency assumptions between these β values we can only talk about an approximate solution. The biggest disadvantage of the simulation-based approach is that generally there is no guarantee to find the optimal solution, plus the fact that these solvers only work in moderate dimensions. However, for the problems analyzed in this paper, convergence seems to be in order.

In our experiment, we implemented the simulator in C++ and solved the problem using GANSO, a C++ library that “implements several methods of global and nonsmooth nonlinear optimization”—see http://www.ganso.com.au/ for more information. The algorithm is started from a solution where the production of each item is equal to its expected demand.

For our main test case of 15 items with 200 scenarios, the local-optimization method of GANSO takes about 55 seconds to find a solution on a 3 GHz PC, using 26124 simulator calls (objective-function evaluations). When we try one of the global optimizers, it takes about nine minutes and ends up with the same solution. This means that the method is usable for a 15-item case, but most likely will not work if we have a significantly bigger portfolio, say 30 items or more. However, these moderate sizes are reasonable for the problem at hand.

3.4 Adjusted SP model from section 3.1

Below we present an approximation that is suited for complex real assortment problems with complicated dependencies between the individual item demands. We recognize that by a simple heuristic, a “two-stage” formulation (similar in spirit to the one presented by Noonan, 1995) can be enforced in the SP formulation from 3.1, where maximum possible direct sales are ensured in the “first stage”. Unfulfilled demand is then converted into substitution demand by the producer, given the a priori substitutability probabilities between the items and given excess inventory.

Although substitution sales are not discounted in a real situation or in the existing analytical newsboy models, by introducing a *discounted-substitution-factor,* we enforce maximum possible direct sales for the original demand of an item. To better comprehend the proposed solution, consider the modelling similarity between discounted substitution and salvage of leftover inventory; both ‘discounted’ entities, enforced after satisfying direct demand. As long as discounted substitution ensures higher profit shares than salvage, substitution is enforced before salvage. Furthermore, while discounted substitution of unsatisfied demand to leftover inventory happens with predefined allocation pattern (i.e. substitutability values), salvage is only constrained by the leftover inventory. Having the original (undiscounted) objective function, we can then calculate the true expected profit corresponding to the found solution. However, the choice between the available substitutes is still decided by the optimisation model; and hence by the focal company, (which can be the producer or the retailer).

The following changes are made, relative to the original SP formulation from Section 3.1. The parameter *discounted-substitution-factor* is introduced. The notation $q \in [0,1]$ is used, and $q * v_i$ defines the discounted value of the substitution sale. By an appropriate choice of q we enforce direct sales before substitution sales. Changes in the model only apply to the objective function (1) from Section 3.1. We denote it (1*), and maximize expected profit from ordinary sales, discounted substitution sales and salvage, over all items and all scenarios:

$$\textit{Maximize Expected profit} = \sum_{s \in S} p^s \sum_{i \in I} \left(-c_i x_i + v_i y_i^s + q v_i \, zt_i^s + g_i w_i^s \right) \qquad (1^*)$$

The true value is then found by putting the optimal solution using (1*) into (1). The model is implemented in AMPL with CPLEX as the underlying solver; see http://www.ilog.com/ for details on both systems.

In our main 15-item test-case, the best discounted-substitution-factor is about q = *0.6*. Although finding the range of the factor value – within which we have robust model solutions – is data dependent, the time consumed for this task is negligible. For q values lower than about 0.6, the individual item production quantities are not sensitive; hence the model solution is robust. Small fluctuations in production quantities (up to 4 % of the total production) are observed though; they are coming from the newsvendor model's nature of simultaneously optimizing over direct sales, substitution sales and salvages, *and* some price heterogeneity. However, these fluctuations do not seriously affect our results.

3.5 Conclusion on Section 3 and test cases

Above, we presented the modelling difficulties, and proposed solution heuristics; the exact problem is intractable. Due to the complexities arising from the dynamics of the problem, the outcome of many consumer generated substitutable newsvendor models is actually optimal allocation between direct and substitution sales; hence, structurally similar to the manufacturer directed substitution. A good example can be found in Rajaram and Tang (2003). They assume consumer directed substitution, but what they solve is producer directed. We manage to enforce direct sales before substitution sales. The MIP suggestion principally captures the "correct" customer arrival process; however, it is only suited for problems with very few items, it makes strong assumptions on independence among the substitution probabilities, and it can only handle few scenarios. The adjusted stochastic programming formulation from Section 3.4 is an approximation of the MIP, able to handle large problems with complex dependency patterns. The approximation, however, requires a setting where all the demand is revealed and cleared at the same moment, and that the substitution choice is directed by the optimisation model. Although, this solution is practically impossible to achieve by a retailer, it closely describes the situation of fashion and sports apparel suppliers with *partial consumer directed substitution*, as discussed in Section 1.

For the problem arising from our market level approach discussed in B in Section 3.1, it seems to be impossible to achieve an exact solution. Although, by integrating the intermediate substitution probabilities β_{ij} in our simulation-based optimisation we could partially solve the problem, these values are still just approximations with strong independency assumptions. However, this is the solution we find appropriate to apply, given that we actually do not know how else to approach substitution shares in a way that makes sense to decision makers, and is also consistent with the optimisation process. To define a multi-period model with the inventory-dependent dynamic substitution choices we would need to describe the conditioned a priori substitutability values; a task which, in addition to its difficulty, would lead to an overly complex problem. And finally, recall that we actually manage to handle some of the important dependencies by our α values; as described in C, Section 3.1.

TEST CASES

Our purpose was to analyse modelling difficulties and suggest appropriate model formulations to the substitutable newsvendor problem. Detailed discussion on the numerical aspects and analyses on the impact of different distributional and substitutional assumptions is important and need to be focused on, more than we have space for doing it in the present paper. Accordingly, here we only present some model tests to illustrate the behaviour of our models. For further discussion on the numerical aspects we refer to the working paper of Vaagen et al. (2008).

The following test cases are used: 1) Two- and seven-item problems provided by Rajaram and Tang (2001). Note that there are no individual correlation and substitution fraction values given by the authors, but values that represent averages. Also note that normality is assumed for the individual item demands with *low*, *medium*, *high* and *mixed*-variations. With the given data, the *low*-variation case is the only one that doesn't have a reasonably high probability of negative demands. In our tests we remove the normality assumption; although for the low-variation case we perform the analyses also with assumed normal distribution, denoting it *LowN*. 2) A real assortment problem with 15 items, facing complex uncertainties and dependencies. Here we use individual correlation and substitutability values; no averages or other simplifying assumption are made. For a detailed description of the latter and estimation of the parameters not discussed in this work we refer to Vaagen and Wallace (2008).

For clarity, the following notation is used: SP^M represents the stochastic programming formulation from 3.1, indicating the underlying manufacturer directed substitution behaviour; SP^C the adjusted stochastic program from 3.4, indicating the underlying partial consumer directed behaviour; MIP the mixed integer program suggested in 3.2 but not presented; the simulator from 3.3 is denoted by SB^{β}, indicating the use of the intermediate β_{ij} probabilities.

Illustrative numerical results are summarized in
Table 2, with expected profit levels and computation times. The two- and seven-item profit levels illustrated by the table are for the low-variation test-case, with correlation values of 0.5 and 0.4, and average substitutability values 0.2 and 0.1, respectively.

Observe that the expected two-item profit levels are equal when solving the manufacturer and partial consumer directed substitution problems, obtained by SP^M and SP^C. The results confirm that the possible problem arising from the Rajaram and Tang formulation, discussed in the literature review section –that of *total sales generated by a particular item j (both direct sales and sales arising from unsatisfied demand for j) might exceed the demand for item j*–, actually does not arise in the two-item setting. Hence, it illustrates that it is difficult to generalize the conclusions of Rajaram and Tang, without actually having stated the seven-item model.

Further, we illustrate that under the two-item mix-variation case, Rajaram and Tang seem to obtain higher profit levels than we do with our SP^M and SP^C formulations (about 10% higher); see Figure 1 for our results. For the results obtained by Rajaram and Tang (2001) we refer to Figure 4 in their work. Although this observation does not prove the possible problem of the Rajaram and Tang formulation, it is a strong indication. Figure 2 gives the corresponding order quantity levels. The total order quantity for the portfolio is increasing/ decreasing in correlation values depending on the actual parameter values; as also indicated by Netessine and Rudi (2003). The monotonous increase in order quantity for Rajaram and Tang (we refer to Figure 3 in Rajaram and Tang, 2001) can, again, be an indication of the above mentioned problem.

Table 2 Numerical test results —Profit levels and computation times CPU, for the MIP, SP^M, SP^C, and SB^{β} model solutions.

		MIP		*SP^M*	*SP^C*	*SP^{β}*	
test case	*# scen.*	*profit*	*time (s)*	*profit*	*profit*	*profit*	*time (s)*
2 items R&T, *c*=0.5, *a*=0.2	1000	8 948	0.05*	8 948	8 948	8 948	10.8
7 items R&T, *c*=0.5, *a*=0.2	250	n/a	n/a	31 881	31 734	31 691	115
15 items, our test case	200	n/a	n/a	2 166 684	1 777 850	1 821 631	753

**For the MIP, we only used 10 scenarios. Solution of the 7-items case was aborted after two days.*

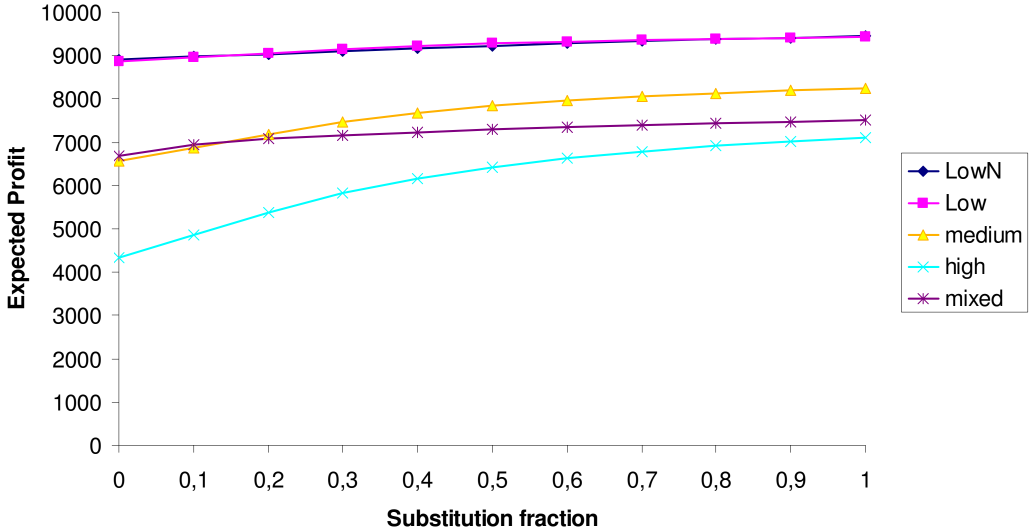

Figure 1 Expected profit versus substitution for different levels of demand variation for the two-item case —identical solution obtained by the SP^M and SP^C formulations.

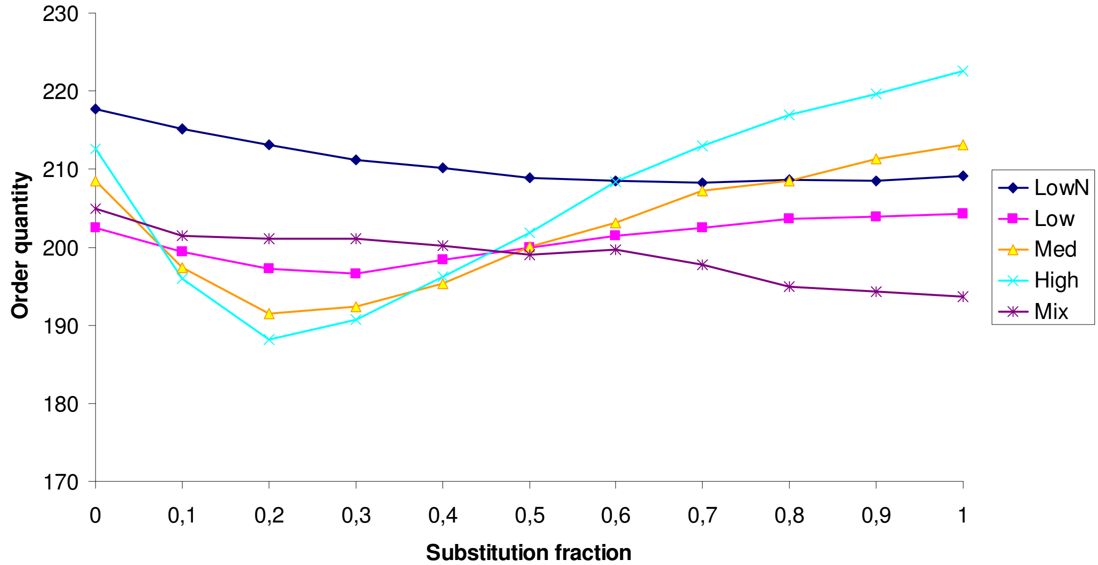

Figure 2 Order quantities versus substitution for different levels of demand variation for the two-item case —identical solution obtained by the SP^M and SP^C formulations.

Further, in our seven- and 15-item tests, SP^C results in lower profit levels than SP^M; 18% reduction for the 15-item case. Although SP^C only partially handles consumer directed substitution, thereby providing an optimistic solution to the real problem, the results indicate the potential error of solving the manufacturer directed substitution problem when the real problem is consumer directed. The simulator SB^{β} is only an approximation (due to the independence assumptions) but it provides a solution closer to reality, as it actually models the dynamic choice process described in Section 3.2. Note that customer arrival is uncertain and the modelled process provides an 'average' to this; with possible substitution while there is still unsatisfied direct demand left. As such, the profit achieved by SB^{β} can be under/over the profit indicated by SP^{C}; this despite reduced control over customer choices in SB^{β}. For the 15-item example, the simulator results in 19% less profit than SP^{M}, which solves the manufacturer directed substitution problem. A slight increase of 2% is indicated from the profit obtained by SP^{C} solving the partial consumer directed problem. We conclude that although the real problem is intractable, the partial solution with optimal substitution plan after enforcing direct sales before substitution sales is a good approximation.

Conclusions

In this work we discussed how to model consumer directed substitution, and further, we proposed heuristic approaches. The research question arose from the recognition that there is inconsistency in many articles discussing assortment and inventory problems with substitution; inconsistency in what the authors claim to solve and what they actually solve, *and* inconsistency between the static nature of the optimisation models and the dynamic data required to establish the model-input substitution parameters. Also, we have found existing work limited in value with regard to complex applications. Although the exact analysis of the problem is difficult, we believe that our approach is appealing for several reasons.

Firstly, substitution fraction estimation and assortment/inventory optimisation is discussed simultaneously, providing reasonable consistency throughout the dynamic choice process.

Secondly, despite the single-period newsvendor model we use, our substitution process is an approximation of the *dynamic* choice process. We achieve this by defining the inventory-independent substitutability matrix –with each first preference having the possibility to be replaced by several "similar" items–, *and* by constraining the factual substitution by unfulfilled demand and substitutes available on stock. We approached the problem on market level; with an attribute view, where consumers are interested in particular attributes rather than products. Our substitution measures describe the "similarity" between the items with regard to trend driver attributes, like colour and design; and as such, they "make sense" to practitioners in fashion and sports apparel environments. Understanding the trend drivers is emphasized to be the key for success in these contexts.

A limitation of our market level approach is that we are unable to fully describe the substitution. A total description would require full knowledge of all conditional substitution probabilities. From a practical point of view, this is a rather unreasonable requirement. Further, it could only be used properly in a dynamic model where inventories were continuously updated. That said, some important dependency patterns are actually built into our models. For example, we can handle negatively correlated second choice possibilities, when the negative correlation is caused by which item becomes popular. The latter is not known at the time of production decisions.

An important advantage of our market level approach, however, is that we avoid the strong (and in our highly volatile supply chain setting unreasonable) assumptions required for applying classical utility maximisation multinomial logit MNL models. In particular, we avoid the IIA property, requiring that the choice possibilities are independent of other alternatives.

Finally, our stochastic programming and simulation-based optimisation heuristics are computationally efficient. Particularly, we do not need simplifying distributional assumptions on demand, and can handle complex correlation and substitution matrices in multi-item settings.

We have tested the model numerically, and showed the potential errors of using producer directed substitution, when the problem is actually consumer directed. Even *partial consumer directed substitution*, achieved by the stochastic programming formulation, enforcing direct sales in a “first stage” but providing an optimal substitution plan, contributes substantially to a reduction in the profit estimation error. Besides the fact that this formulation is suited for large complex assortment problems, it also closely describes the substitution problem faced by fashion and sports apparel suppliers, where the retail customers’ first choice cannot be directed but in case of stockout the supplier can, to some extent, control which retailers should be served first.

With simulation-based optimisation we provide a solution “closer” to the real consumer directed substitution problem. However, due to strong independency assumptions between the substitution measures we can only talk about partially modelling the inventory-dependent dynamic substitution process.

Our results largely support previous findings on the effects of substitution on inventory decisions, with correlated individual item demands; see conclusions provided by Rajaram and Tang (2001), analytically confirmed by Netessine and Rudi (2003). The substitutable assortment profit is decreasing in correlation value; the decrease is largest when there is high variation in demand for all items. Profit gains are highest under negatively correlated demands. Assortment profit increases in any substitution value.

Acknowledgements

All three authors are partially supported by The Research Council of Norway under grant 171007/V30.

References

Anupindi, R., Dada, M., Gupta, S., 1998. Estimation of consumer demand with stockout based substitution: An application to vending machine products. Marketing Science, 17, 406-423.

Bassok, Y., Anupindi, R., Akella, R., 1999. Single-period multi-product inventory models with substitution. Operations Research, 47 (4), 632-642.

Bitran, G., Dasu, S., 1992. Ordering policies in an environment of stochastic yields and substitutable demands. Operation research, 40(5), 999-1017.

DeHoratius, N., Raman, A., 2004. Inventory record inaccuracy: an empirical analysis. Working paper. Graduate School of Business. University of Chicago.

Campo, K., Gijsbrechts, E., Nisol, P., 2003. The impact of retailer stockouts on whether, how much, and what to buy. International J. Res. Marketing, 20, 273–286.

Ernst, R., Kouvelis, P., 1999. The Effects of Selling Packaged Goods on Inventory Decisions. Management Science, 45, 1142-1155.

Gaur, V., Honhon, D., 2006. Assortment Planning and Inventory Decisions Under a Locational Choice Model. Management Science, 52(10), 1528-1543.

Gerchak, Y., Tripathy, A., Wang, K., 1996. Co-production models with random functionality yields. IIE transaction, 28(5), 391-403.

Kall, P., Wallace, S.W., 1994. Stochastic programming. John Wiley & Sons, Chichester.

Kaut, M., Høyland, K., Wallace, S.W., 2003. A heuristic for moment-matching scenario generation. Computational Optimization and Applications, 24, 169-185.

Kaut, M., Wallace, S.W., 2007. Evaluation of scenario-generation methods for stochastic programming. Pacific Journal of Optimization, 3(2), 257--271.

Khouja, M., Mehrez, A., Rabinowitz, G., 1996. A two-item newsboy problem with substitutability. International Journal of Production Economics, 44, 267-275.

Kök, A.G., Fisher, M., Vaidyanathan, R., 2006, forthcoming. Assortment Planning: Review of Literature and Industry Practice. In Agrawal, N. and Smith, S.A. (Eds.). Retail Supply chain management. Kluwer Publishers.

Kök, A.G., Fisher, M.L., 2007. Demand Estimation and Assortment Optimization under Substitution: Methodology and Application. Operations research, 55(6), 1001-1021.

Mahajan, S., van Ryzin, G., 2001. Stocking retail assortments under dynamic consumer substitution. Operations Research, 49, 334–351.

Netessine, S., Rudi, N., 2003. Centralized and Competitive Inventory Models with Demand Substitution. Operations Research, 51(2), 329-341.

Noonan, P.S., 1995. When consumers choose: a multi-product, multi-location newsboy model with substitution. Working paper. Emory University, Atlanta, Georgia.

Parlar, M., Goyal, S.K., 1984. Optimal ordering policies for two substitutable products with stochastic demand. Opsearch 21(1), 1–15.

Pasternack, B., Drezner, Z., 1991. Optimal inventory policies for substitutable commodities with stochastic demand. Naval research Logistics, 38, 221-240.

Rajaram, K., Tang, C.S., 2001. The impact of product substitution on retail merchandising. European Journal of Operations Research, 135, 582–601.

Raman, A., Fisher, M.L., McClelland, A., 2001. Supply Chain Management at World Co. Ltd. Harvard Case Collection, HBC Publishing, 21.

van Ryzin, G., Mahajan, S., 1999. On the relationship between inventory costs and variety benefits in retail assortments. Management Science, 45, 1496–1509

Vaagen, H., Wallace S.W., 2008. Product variety arising from hedging in the fashion supply chains. International Journal of Production Economics 114, 431-455.

Vaagen, H., Wallace, S.W., Kaut, M., 2008. The value of numerical models – substitutable newsvendors. Working paper / Molde University College, no. 2008:16. Molde, Norway: Molde University College.