Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.

INVESTIGATING THE HIDDEN LOSSES CAUSED BY OUT-OF-SHELF EVENTS: A MULTI-AGENT-BASED SIMULATION

Priscilla Avegliano
Carlos Cardonha

IBM Research
Rua Tutóia, 1157
São Paulo, SP CEP 04007900, BRAZIL

ABSTRACT

Out-of-shelf events refer to periods of time in which items of a certain product are not available to customers. It is clear that incidents of this nature result in economic loss, but their side effects are much more profound: since there is no record of missed sales opportunities, the estimated demand curve tends to be inaccurate. As a result, order placement strategies employed by retailers are based on imprecise forecast models, so further out-of-shelf events are very likely to occur: a vicious cycle, hence, arises. In this work, we propose a multi-agent-based simulation to evaluate the impact of out-of-shelf events that considers the reactions of customers towards these incidents and retailers' ordering strategies. Our results show that these events have a significant effect on demand estimation and that multi-agent-based simulations may provide interesting insights and support for the development of more accurate forecast models in retail.

1 INTRODUCTION

Demand estimation is one of the most important challenges faced by the retail industry. Motivated by the unquestionable economical relevance of the problem, researchers and practitioners have been dedicating a considerable amount of effort to identify accurate forecast models. However, there is still a lack of satisfactory solutions for certain real-world scenarios (Papakiriakopoulos, Pramatari, and Doukidis 2009). Typically, statistical models are imprecise when they fail to consider relevant parameters of the problem, and in this context, the *impact of out-of-shelf events in fluctuations of product demand* has been one of these overlooked aspects. The term *out-of-shelf (OOS)* (or *on-shelf unavailability*) is used when a certain product is not available to customers on the shelf at a retail store.

Despite the substantial losses that it inflicts on retailers and manufacturers, the level of on-shelf unavailability remains steady at around 8% worldwide (a value that grows up to 12% for items on sale) since the 90's (Gruen, Corsten, and Bharadwaj 2002). OOS is difficult to prevent because it can be triggered by a multitude of root causes; obviously, when a retailer runs out of items in the back-room stock (which can be a consequence of poor demand forecasts or phantom stocks), the problem is unavoidable; moreover, on-shelf unavailability also takes place if replenishment plans are misaligned with sales velocity; finally, product misplacement is another serious related issue, caused both by planogram' nonconformities (*e.g.*, employees do not put the items on the correct shelf) and by customer interaction (*e.g.*, someone takes an item and abandon it in an incorrect shelf).

The impact of OOS events for retailers is probably more complex than a superficial analysis would allow one to see, since it has several direct and indirect consequences. It is estimated that on-shelf unavailability is the cause of 3.9% to 4.5% sales losses (Gruen, Corsten, and Bharadwaj 2002), so it is clear that the most immediate effect is economically significant. However, in addition to that, we should also take into account that customers facing OOS events may have different types of reactions; while some people postpone their purchase, others may search for the same item in a different store, and other will eventually buy a similar item from a different brand. We therefore claim that any serious analysis about the consequences of OOS events must take into account this heterogeneity in customer behaviour.

We present in this article a multi-agent-based simulation model that provides the environment for a deep analysis of the lasting effects of OOS events. We adopted this methodology because it enables the representation of (1) different customers' reactions when faced with an on-shelf unavailability and (2) the social processes that influences customer's reactions. The aggregation of different user behaviours, networks of social influence, and ordering strategies based on a product's availability determine the overall characteristics of the demand curve.

In our experiments, we validate the simulation model by checking (and actually observing) the occurrence of the so-called *bullwhip effect*, which says that demand variability tends to increase as we go beyond the customer-retailer relationship. Our results show that the effects of OOS incidents propagate over time and make traditional demand forecast algorithms inaccurate. As a consequence, they also tend to trigger other OOS incidents in the future, which results in a self-reinforced vicious cycle.

This article is structured as follows. Section 2 provides an overview of related work, which basically consists of articles about OOS events, the bullwhip effect, and multi-agent simulations of socio-economic scenarios. Section 3 is dedicated to the description of our simulation model, and in Section 4 we present and analyse our computational experiments. Finally, in Section 5, we present our conclusions and propose directions for future work.

2 RELATED WORK

On-shelf unavailability is a recurrent topic in the literature (specially in Marketing), since it has been identified as a common concern to retailers and consumer packaging good companies for a long time. More precisely, to the best of our knowledge, scientific studies dedicated to the measurement of OOS levels and to the identification of their root causes are being conducted since the 1960's (Grocer 1968a, Grocer 1968b).

Several articles address the problem of estimating the volume of sales that did not take place due to on-shelf unavailability. (Gruen, Corsten, and Bharadwaj 2002, Corsten and Gruen 2003) evaluate these rates by aggregating data collected from audits performed at retail stores and from points-of-sale (POS). (Fernie and Grant 2008) estimate that an out-of-shelf event takes place when a given product p is sold every t minutes in average and no sale is registered at the POS for a period $t' \geq t$. These three works estimate the economical impact of on-shelf unavailability as the product of OOS levels, percentage of customers that do not buy a product when faced with an OOS situation, and sales volume. This evaluation assumes that the consequences of on-shelf unavailability are restricted to the time intervals where OOS events took place.

(Anupindi, Dada, and Gupta 1998) present an approach that is closer to the one employed in our work. The authors introduce a method to estimate demand by taking into account *periods of OOS*, that is, the amount of time in which shelves remain empty. In essence, the method considers OOS periods as being moments of time for which data about sales is missing. In order to estimate the absent values, the authors employ interpolation techniques. This solution disregards possible customer reactions and, as a consequence, assume that OOS events do not affect demand.

(Papakiriakopoulos, Pramatari, and Doukidis 2009) propose the use of Decision Tree Learning techniques (Mitchell 1997) to identify patterns of user behaviour stimulated by on-shelf unavailability. This approach was adopted because the authors could not find a statistical distribution that fits satisfactorily the data collected from the points-of-sale and explains how demand behaves.

According to (Chen, Ryan, and Simchi-Levi 2000), *exponential smoothing* is currently the most popular method used to estimate demand; moreover, in the same article, the authors prove analytically that this strategy fosters the bullwhip effect. This is a phenomenon that takes place in supply chains and basically refers to the fact that demand variability tends to increase as we go beyond the customer-retailer relationship.

More precisely, orders placed by retailers to suppliers present a higher variance than customer demand observed in the stores (Forrester 1961). Thus, the small perturbations that can be generated by on-shelf unavailability are propagated and amplified through all the links of a supply chain. (Chen, Drezner, Ryan, and Simchi-Levi 2000) proved that *moving averages* also have the same problem. In a sense, both articles are related to ours, as they investigate the propagation of demand variation on the supply chain. However, we focus specifically on the evaluation of OOS events and on variations of demand perceived by the retailer - variations that, according to the bullwhip effect, will be amplified through the supply chain.

In this article, we employ a multi-agent model to investigate the impact of OOS events on demand forecast. The use of agent-based approaches is not new in the context of socio-economic scenarios. Since the seminal work of Schelling (Schelling 2006), which investigates how micro-decisions can impact dramatically the overall performance of a complex system, an extensive research has been conducted in order to identify and simulate misbehaviours that justify emergent phenomena in Economy. Just to name a few, we can mention (Kodia, Said, and Ghedira 2009), which proposes a multi-agent simulation of stock-markets, (Kephart, Hanson, Levine, Grosof, Sairamesh, Segal, and White 1998), which investigates and simulate open information-driven markets, and (Said, Drogoul, and Bouron 2001), which develops a multi-agent simulator modeling consumer behaviour.

It is possible to identify work in the literature proposing the use of agents to model, simulate, and/or optimize a supply chain (see *e.g.* (Swaminathan, Smith, and Sadeh 1998) and (Zarandi, Pourakbar, and Turksen 2008)). These articles explore the decentralized yet coupled relationship that occurs among producers, distributors, etc., that is, they focus on the internal dynamics of a supply chain. Differently, we focus on how the heterogeneous cognitive and social processes that takes place when customers face OOS events affect demand.

Another interesting approach is presented in (Parunak, Savit, and Riolo 1998), which uses supply network as a scenario for agent-based and equation-based simulations. The results obtained from the agent-based simulation are quite similar to the ones presented in this article regarding the bullwhip effect. However, both works diverge on their target simulation goals: while Parunak concentrates on the supply network relations' between suppliers of a single product, we focus on the modeling of consumer behaviours in an environment of multiple products.

3 MULTI-AGENT MODEL

In order to evaluate the impact of OOS events on demand fluctuations, it is crucial to consider that customers may have different reactions towards on-shelf unavailability. In order to create a simulation environment where this heterogeneity is considered, we designed a discrete-event multi-agent-based model to address the problem. The proposed formulation employs three classes of entities: *customers*, *retailers*, and *producers*. The characteristics of these agents are detailed below.

3.1 Customers

We defined the behaviour of customers based on the results of surveys published in the literature (see *e.g.*, (Campo, Gijsbrechts, and Nisol 2000), (Campo, Gijsbrechts, and Nisol 2004), (Sloot, Verhoef, and Franses 2002), and (Gruen, Corsten, and Bharadwaj 2002)). In these works, the authors measure brand loyalty and classify customers' reaction to OOS events using the following categories (together with their respective probability of occurrence):

- 31% try to buy the same product at another store;
- 26% substitute the product for another one from a different brand;
- 19% substitute the product for another one from the same brand;
- 15% delay the purchase;
- 9% do not purchase any product.

The percentages above vary if we restrict the scope either to first-need or to hedonistic products (see (Sloot, Verhoef, and Franses 2002)), but we do not make distinctions in this work.

We partitioned the set of customers in two profile groups: *loyal clients* and *switchers*. Customers who buy the product at another store, delay their purchase, or do not buy anything in OOS event are grouped in the loyal profile, whereas clients who buy either a similar product from the same brand or from some other brand are classified as switchers. Based on the percentages above, 55% of the agents in our simulations are loyal, while the remaining 45% are switchers.

Although loyal customers tend to be more "stable", it is reasonable to assume that they have a limited amount of patience with on-shelf unavailability. Therefore, we assume that these clients maintain a record of previous OOS incidents. Namely, if they experience OOS events once, they becomes less "attached", switching and becoming loyal to some other brand if they face on-shelf unavailability once again.

Figures 1 and 2 depicts the state transition diagram governing changes of preferences for loyal customers and switchers, respectively, in a scenario consisting of 2 products. In these figures, every transition (indicated by the arrows) is triggered by an OOS events.

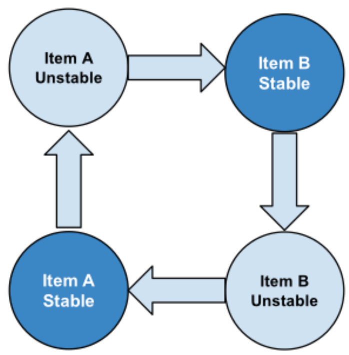

Figure 1: State diagram for loyal customers

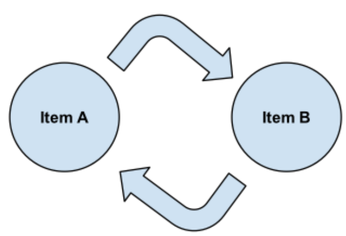

Figure 2: State diagram for switchers

We also assume that, for each customer, there are exactly two clients to which we assign a non-zero influence level (*e.g.*, some integer number between 1 and 10). After facing an OOS event and deciding not to purchase items from his preferred product P again, customer C checks if her most influential "neighbour" is currently purchasing a different product P'. If this is the case, C switches to P'. If not, C verifies the other neighbour's preference and repeat the procedure. Finally, if both neighbours are currently buying P, C will adopt a randomly chosen product $P' \neq P$.

3.2 Retailer and Producers

As the simulation was aimed at replicating the root causes of OOS events, it contemplates not only on-shelf unavailability, but also stock-outs. Hence, for each product, the retailer has a *back-room stock* and a *shelf* to dispose the items, both with limited capacity. Moreover, each product is associated to a producer, who is always able to deliver as many items as required and who always takes the same amount of time to make deliveries.

In our simulations, retailers employ exponential smoothing to estimate product demand and, consequently, to determine the number of items that should be ordered. The forecasted value at instant t is denoted by f_t and is described by the following equation:

$$f_t = \alpha * x_{t-1} + (1-\alpha) * f_{t-1} \tag{1}$$

$$f_0 = x_0 \tag{2}$$

The estimation of demand f_t for period t is given by a convex combination of the real demand registered in the previous period, denoted by x_{t-1}, and the estimation computed in the previous period, given by f_{t-1}. That is, $\alpha \in [0,1]$ reflects the behaviour one wishes to impose to f_t: values closer to 0 give a higher weight to historical data, while values closer to 1 privilege recent measurements and induce forecasts to react more quickly to recent changes in sales velocity.

Retailers also keep a threshold k for the number of items of each product in their back-room storage. If the stock goes below k at instant t, p products are automatically ordered, where p is the maximum between f_t and k. A new order can only be placed after the delivery of the previous one.

Let us suppose that a retailer makes a new order at time t, that the previous order was placed at time t', that i items were sold in this timespan, and that the producer takes the amount of time d to deliver the order. Since back-room space is limited, orders will ideally have the minimum number of products that will still be large enough to fulfil the upcoming sales requests. Therefore, the value of k is given by

$$k = d\left(\frac{i}{t-t'}\right).$$

The value of p is given by

$$p = d\left(\max(f_t, k) + \sigma_t\right),$$

where σ_t denotes the standard deviation of the previous demand forecasts (which are computed in every step t) and $\max(f_t, k)$ denotes the maximum between f_t and k. The intuition behind this formulation is the following: as the retailer knows that orders take d time steps to be delivered, he wants to have enough items in the store in order to avoid OOS events during the waiting time. Demand is estimated according to the average sales velocity observed between the previous and the new order, while σ_t accounts for safety margin.

Finally, for each step of the simulation, each product P has a market share $s_P \in [0,1]$, which indicates the percentage of customers having P as their favourite product. We assume that the products in our simulation compose the complete market of some specific item, so the sum of all market shares is always equal to one. The values s_P change over time according to the occurrence of OOS events and the profile of customers involved in the incidents. Market share rates are important to the simulation in the sense that a higher market share increases the probability of a product to be recommended to a dissatisfied customer by an influencer neighbour.

3.3 Description of the simulation

Our simulator considers one retailer, three competitor products (each one associated to exactly one producer), and n = 100 customers. The simulations take place in 365 steps and producers always take 7 steps to deliver the orders.

The market shares start with values 0.5, 0.3, and 0.2. Based on these values, we define the capacity of the back-room storages. Namely, if product P starts with market share s_P, its back-room will have space for $10ns_P$ items (which is equal to 500, 300, and 200 for the first, second, and third products, respectively). Although the values of s_P change over time, the back-room space of each P remains fixed. Similarly, the shelf of product P has space for ns_P items (50, 30, and 20 slots for the first, second, and third products, respectively). In the first step, shelves and back-rooms are completely full.

In the beginning of each step, shelves are replenished to their maximum capacities. If there are not enough products in the back-room, it is emptied and the shelf will not be completely filled. After the replenishment phase, the system checks the current status of the back-room storages. The retailer should place an order for a given product if the number of items is smaller than k and if an order has not been placed in the previous 6 steps, which means that the retailer is not waiting the deliver of a previous order. Values f_t are computed as indicated by Equation 1 with $\alpha = 0.9$, so demand forecasts react quickly to recent changes in sales velocity. We assume that these operations take place at night or very early in the morning, when clients do not have access to the store.

Once the "management" phase is finished, customers start to visit the store. In each step, all the clients try to buy products following a randomly selected ordering. Non-determinism is important here, since it minimizes the number of simulations in which certain customers face OOS events more frequently than others.

As we mentioned previously, customers may change their preferences after facing on-shelf unavailability. In our simulation, loyal clients that did experience an OOS event do not buy anything, but internally register the incident. The other clients have their preferences immediately changed (according to the interaction with neighbours, as described in Section 3.1) and try to purchase their new favourite product. If on-shelf unavailability is faced after this change, customers pick randomly some other available item, but without changing their preference for the next step.

In the first simulated step, each client had its demand attended, since the shelf space was equal to the number of consumers for each product. This first period of stability is important in order to provide accurate data for the forecast method. We modeled our retailers to present a bold attitude, trying to minimize the back-room space occupied by stocked products and, consequently, presenting a very low safety margin of products. To trigger the first OOS event, we deliberately initialized k with a low value. More precisely, for each product P, the first order submitted to the producers of P takes place when the back-room stock has less than $700s_P$, making the demand of the following 7 days unachievable. In the next steps, once the first on-shelf unavailability happens, an instability is triggered in the market.

4 COMPUTATIONAL RESULTS

We designed a computational experiment to address the following research question: *Does the effect of OOS events propagate over time?* This hypothesis naturally emerges from the fact that on-shelf unavailability leads to instability in demand and, as a consequence, to poor forecasts. All the results were generated after 100 executions of the simulation model described in Section 3.

Bullwhip Effect

In order to validate our simulation model, we checked if it could reproduce the bullwhip effect. We recall that, according to (Chen, Ryan, and Simchi-Levi 2000), this phenomenon always takes place in scenarios where retailers employ exponential smoothing to estimate demand, which is exactly the case in our simulation. Since we are considering scenarios where the retailer does not have any information about sales losses, we define the *bullwhip ratio* as the ratio between the variance of the size of orders placed by the retailer to the producer and the variance of the volume of sales. That is, the bullwhip effect is observed in situations where the bullwhip ratio is greater than 1.

We performed 100 simulations of 3500 steps each, and in Figure 3 we present box-plots showings the distributions of the lowest bullwhip ratios identified in each execution for each product. Since all the values are (well) above 1, we conclude that the simulation model is indeed reproducing the bullwhip effect. Figure 4 shows the evolution of the bullwhip ratio over time for each product in one of our simulations.

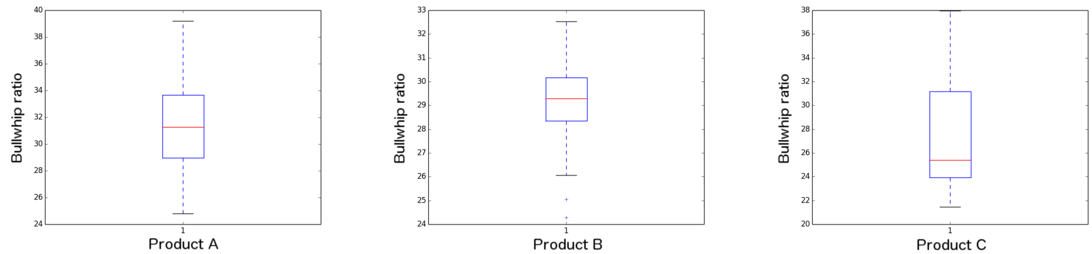

Figure 3: Boxplot of the lowest bullwhip ratios obtained in 100 simulations

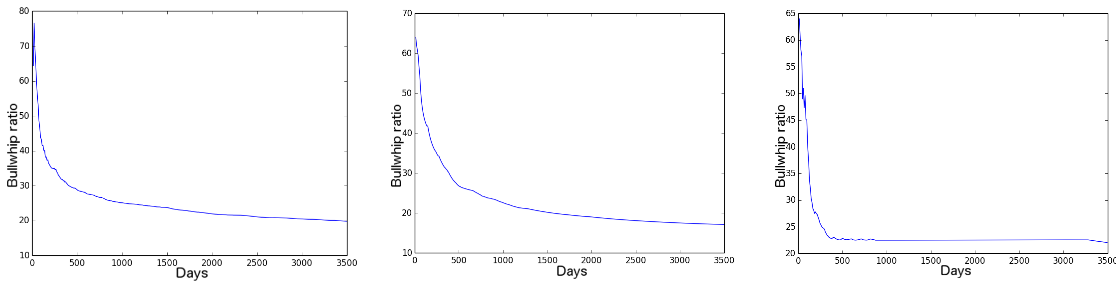

Figure 4: Evolution of bullwhip ratio over time for each product in one scenario

Propagation of OOS events

We investigated the propagation of OOS events over time by performing 100 simulations of 365 steps each, representing the sales dynamic for a period of one year. Figure 5 presents the number of items of type A sold in each step in one of these simulations. On-shelf unavailability can be identified in steps where sales drop abruptly and eventually go to zero.

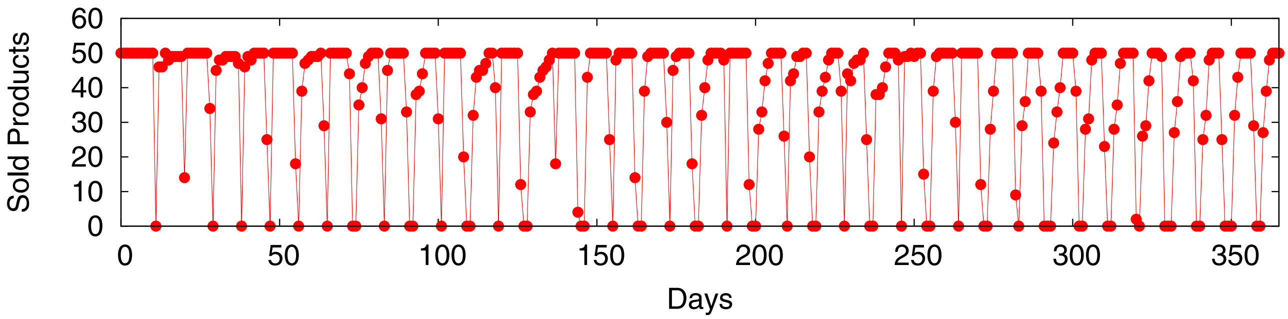

Figure 5: Evolution of sales over 365 days

Figure 6 is the composition of box-plots from the simulations and shows the moving-average of the number of sales opportunities that were not realized due to OOS events on a time-frame period of 30 simulation cycles for each step.

Discussions

Figure 5 shows that OOS events have a significant effect on sales. One can see that recovery from on-shelf unavailability is frequently not immediate and, in certain cases, rather slow. It is possible to identify several sequences where the shelves were already clearly full but sales where below their maximum capacity, as evidenced, for instance, on the 200^{th} step.

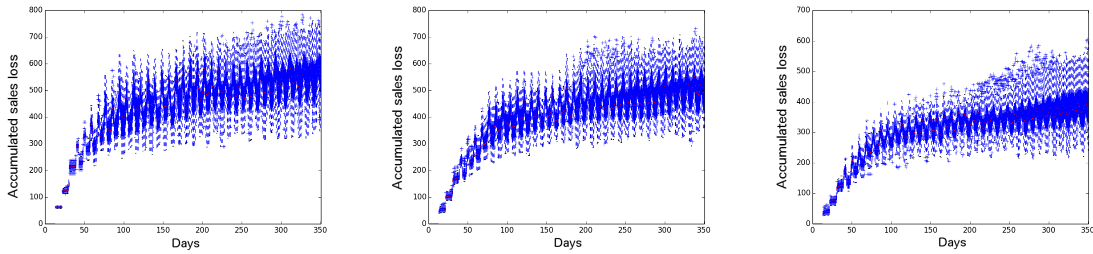

Figure 6: Accumulated number of OOS sales losses on a 30-day period

Figure 5 also reveals that this instability, in the long term, is very harmful for the retailer: as times passes and more OOS events takes place, sequences of steps where 50 products were sold become less frequent and smaller. Moreover, the system tends to stabilize in a lower mean value of daily sales, clearly leading to inaccurate forecasts that result in losses in the long term. Since this effect is distant in time from the initial OOS event that triggered the whole process, this association of cause-consequence can become unnoticed at a first glance.

It can be observed in Figure 6 that the incidence of OOS events increases over time and that this phenomenon is transversal to all the products. It also shows that, after a strong impact in the short term, the number of missed sales tends to stabilize (which is reasonable, since we considered a stable consumer market), but we remark that the economical impact on this plateau is significant: whereas the overall number of possible sales in a 30-day period is equal to 3000, accumulated sales loss can reach a level of more than 1500 in the same period.

These results allow us to conclude that the consequences of OOS incidents are not circumscribed to their duration, and the main reason for this fact is that they affect the preference of customers, leading to significant changes in demand. Therefore, a deeper investigation should be carried out in order to enhance current forecast methods by taking demand fluctuations caused by OOS events into account, avoiding, thus, a perpetual cycle of error propagation.

5 CONCLUSIONS

We presented in this article a multi-agent-based simulator used to evaluate the impact of out-of-shelf events on retail stores. Our computational experiments show that the effects of OOS incidents propagate over time. As a consequence, demand forecasting becomes inaccurate and further OOS incidents are more likely to occur, leading to a dynamics that creates a self-reinforced vicious cycle. Hence, more efficient forecast methods should be developed, since on-shelf unavailability cannot be seen anymore as an isolated event.

There are several possible extensions of this work. For instance, as we observed the occurrence of the bullwhip effect in our experiments, it would be interesting to investigate their relation with OOS events. From the point of view of multi-agent-based systems, it is possible to extend the proposed simulator in several different ways. For example, if we consider scenarios with non-competing products, it will be possible to investigate the correlation between different items and their associated OOS events. If many producers deliver similar products, we could incorporate variations on prices, number of delivered items, delivery time, etc.. These are just a few examples coming from a very large list, so there is clearly space for improvement in our simulator.

Multi-agent-based simulations can be considerably leveraged by the incorporation of optimization modules. For instance, if retailers want to use their back-room stocks in the best possible way by giving more space to items that sell more quickly, they will have to solve a constrained multi-dimensional knapsack problem, a classical combinatorial optimization problem that can be addressed by techniques such as Mixed Linear Integer Programming (Kellerer, Pferschy, and Pisinger 2004). Distribution of goods from producers to stores is also an example of problem where methods from Operations Research could not only be applied, but also have their use validated.

Finally, we conclude that the multi-agent-based approach used in the simulator allowed for the replication of an emergent phenomena, triggered by OOS events. This overlooked aspect may be one of the causes of the poor performance of current demand forecast models currently employed in the retail industry. Therefore, improved forecast methods that consider the bullwhip effect might not only reduce local damages in retail stores, but in the whole supply chain.

REFERENCES

Anupindi, R., M. Dada, and S. Gupta. 1998. "Estimation of consumer demand with stock-out based substitution: An application to vending machine products". *Marketing Science* 17 (4): 406–423.

Campo, K., E. Gijsbrechts, and P. Nisol. 2000. "Towards understanding consumer response to stock-outs". *Journal of Retailing* 76 (2): 219–242.

Campo, K., E. Gijsbrechts, and P. Nisol. 2004. "Dynamics in consumer response to product unavailability: do stock-out reactions signal response to permanent assortment reductions?". *Journal of Business Research* 57 (8): 834–843.

Chen, F., Z. Drezner, J. K. Ryan, and D. Simchi-Levi. 2000. "Quantifying the bullwhip effect in a simple supply chain: the impact of forecasting, lead times, and information". *Management science* 46 (3): 436–443.

Chen, F., J. K. Ryan, and D. Simchi-Levi. 2000. "The impact of exponential smoothing forecasts on the bullwhip effect". *Naval Research Logistics* 47 (4): 269–286.

Corsten, D., and T. Gruen. 2003. "Desperately seeking shelf availability: an examination of the extent, the causes, and the efforts to address retail out-of-stocks". *International Journal of Retail & Distribution Management* 31 (12): 605–617.

Fernie, J., and D. B. Grant. 2008. "On-shelf availability: the case of a UK grocery retailer". *International Journal of Logistics Management* 19 (3): 293–308.

Forrester, J. W. 1961. *Industrial dynamics*, Volume 2. MIT press Cambridge, MA.

Grocer, P. 1968a. "The out of stock study: Part I". *October. S1–S16*.

Grocer, P. 1968b. "The out of stock study: Part II". *November. S17–S32*.

Gruen, T. W., D. Corsten, and S. Bharadwaj. 2002. *Retail out-of-stocks: A worldwide examination of extent, causes and consumer responses*. Grocery Manufacturers of America Washington, DC.

Kellerer, H., U. Pferschy, and D. Pisinger. 2004. *Knapsack Problems*. 1 ed. Springer Verlag.

Kephart, J. O., J. E. Hanson, D. O. Levine, B. N. Grosof, J. Sairamesh, R. Segal, and S. R. White. 1998. "Emergent Behavior in Information Economies.". In *ICMAS*, edited by Y. Demazeau, 441–442: IEEE Computer Society.

Kodia, Z., L. B. Said, and K. Ghedira. 2009. "SiSMar: social multi-agent based simulation of stock market". In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, 1345–1346. International Foundation for Autonomous Agents and Multiagent Systems.

Mitchell, T. M. 1997. *Machine Learning*. 1st ed. New York, NY, USA: McGraw-Hill, Inc.

Papakiriakopoulos, D., K. Pramatari, and G. Doukidis. 2009. "A decision support system for detecting products missing from the shelf based on heuristic rules". *Decision Support Systems* 46 (3): 685–694.

Parunak, H. V. D., R. Savit, and R. L. Riolo. 1998. "Agent-based modeling vs. equation-based modeling: A case study and users guide". In *Multi-agent systems and agent-based simulation*, 10–25. Springer.

Said, L. B., A. Drogoul, and T. Bouron. 2001. "Multi-agent based simulation of consumer behaviour: Towards a new marketing approach". In *International Congress on Modelling and Simulation Proceedings*.

Schelling, T. C. 2006. *Micromotives and macrobehavior*. WW Norton & Company.

Sloot, L., P. C. Verhoef, and P. Franses. 2002. "The impact of brand and category characteristics on consumer stock-out reactions". Research Paper ERS-2002-106-MKT, Erasmus Research Institute of Management (ERIM).

Swaminathan, J. M., S. F. Smith, and N. M. Sadeh. 1998. "Modeling supply chain dynamics: A multiagent approach*". *Decision sciences* 29 (3): 607–632.

Zarandi, M., M. Pourakbar, and I. Turksen. 2008. "A Fuzzy agent-based model for reduction of bullwhip effect in supply chain systems". *Expert systems with applications* 34 (3): 1680–1691.

AUTHOR BIOGRAPHIES

PRISCILLA AVEGLIANO is a Software Engineer at IBM Research Brazil. She is graduated in Computer Engineering (2005) and holds a M.Sc. in Artificial Intelligence (2007) from University of São Paulo. Currently, her research is foccused on computational social simulation, multi-agent and complex systems.

CARLOS CARDONHA is a researcher in the Systems of Engagement and Insight group at IBM Research Brazil since January 2012. He is currently working on mathematical modelling, optimization, and machine learning to problems related to smarter cities, retail, industrial scheduling, etc. Carlos holds a B.Sc. (2005) and a M.Sc. (2006) degree in Computer Science from the Universidade de São Paulo and a Ph.D. (2011) in Mathematics from the Technische Universität Berlin.