

AGGREGATED SIMULATION MODELING TO ASSESS PRODUCT-SPECIFIC SAFETY STOCK TARGETS DURING MARKET UP- AND DOWNSWINGS: A CASE STUDY

Cas Rosman¹, Eric Weijers¹, Kai Schelthoff¹, Willem van Jaarsveld², Alp Akcay², and Ivo Adan²

¹Dept. of Supply Chain Innovation, NXP Semiconductors N.V., Eindhoven, NETHERLANDS

²Dept. of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, Eindhoven, NETHERLANDS

ABSTRACT

In this study, we propose an aggregated simulation model of the back-end supply chain for manufacturing semiconductors. The simulation model is applied to real-world data from NXP Semiconductors N.V. to assess the need for safety stock at the die bank during market up- and downswings, respectively. To model demand uncertainty, we use future forecasts and adjust them by sampling from discrete distributions of historical forecast errors. For modeling supply, we propose an aggregated simulation model of the back-end supply chain and assume a front-end process that produces to forecast (Make To Stock) with the addition of safety stock. We conclude from experiments that during market up- or downswings, the impact of safety stock target levels on supply chain performance differs significantly. The proposed method allows supply chain managers to assess the impact of safety stock target levels on key performance indicators.

1 INTRODUCTION

The semiconductor supply chain consists of a front-end and a back-end process. In the front-end process, individual components (transistors, capacitors, resistors, etc.) are patterned in the raw wafers. In the back-end process, the patterned wafers are cut into single dies, which are then assembled and tested. Both processes are depicted in Figure 1.

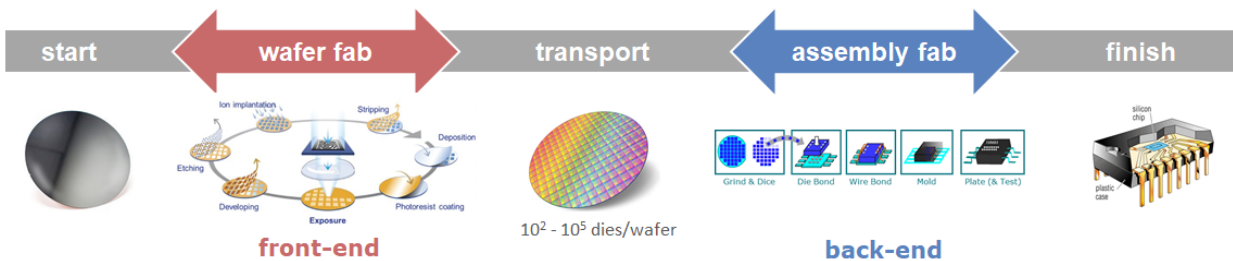


Figure 1: Overview of the semiconductor manufacturing process (NXP Semiconductors N.V. 2024).

In semiconductor supply chains the front-end process typically has a Lead Time between 10-15 weeks (Mönch, Uzsoy, and Fowler 2018b), whereas the back-end process typically has a Lead Time between 1-2 weeks. Because of this difference in Lead Times, different manufacturing paradigms are used in both processes. To prevent long Lead Times to customers, in front-end manufacturing, production orders start based on forecasted demand (Make To Stock). In the back-end manufacturing process, production orders start based on incoming customer orders (Make To Order). Both processes are decoupled by a die bank, an inventory point where wafers coming from the front-end can be stored until they are pulled by the back-end process. The die bank is usually selected as the Customer Order Decoupling Point (CODP) in semiconductor supply chains.

In semiconductor supply chains, managing inventory is crucial for balancing On-Time Delivery (OTD) performance with inventory costs. Forecasting demand is complex due to volatile market demand and short product life cycles, making traditional techniques prone to time lag (Wang and Chen 2019). The semiconductor industry's cyclic demand, influenced by economic factors, increases forecasting challenges and potentially leads to biased forecasts (Rapp et al. 2022). Biased forecasts towards over or under forecasting pose risks of excess inventory or material shortages, impacting key performance indicators. Safety stock can mitigate the risk of material shortages, however, holding inventory entails risks of costs due to obsolescence by quality degradation or market demand. Despite the necessity of safety stock for meeting OTD targets, their implementation requires careful consideration due to the high risks involved with keeping too much inventory. Thus, effectively managing inventory in semiconductor supply chains involves navigating the trade-offs between OTD performance and inventory.

To decide on the desired trade-off, practitioners need to assess many demand/supply scenarios in a short period to prepare for and adapt to unforeseen changes in the electronics market. Current methods often require high data maintenance effort, are computationally expensive, or require high effort from practitioners. In this study, we propose an aggregated simulation model of the back-end manufacturing process to assess the impact of safety stock target levels at the die bank on OTD performance in both market up- or downswings. The proposed aggregated simulation model is based on the existing master data structure of the company (no additional data maintenance needed), only models key bottleneck resources (reduces computational expense), and allows practitioners to quickly set up and analyze different scenarios. We demonstrate the use of the aggregated simulation model by assessing different safety stock target levels during market up- or downswings. To the best of our knowledge, we are the first to assess the impact of market up- or downswings on the need for safety stock in semiconductor supply chains using simulation.

The paper is structured as follows, after discussing relevant literature in section 2, we propose our method in section 3. The method is applied to a real-world use case in section 4. The paper is concluded with a discussion in section 5, managerial insights in section 6, and a conclusion in section 7.

2 LITERATURE REVIEW

In their survey on semiconductor supply chain models, Uzsoy et al. (2018), identify a literature stream that focuses on the location and sizing of safety stocks in semiconductor supply chains. The identified literature stream is highly relevant to this study. We further explore the literature stream by summarizing the literature on the application of simulation models and the impact of forecast inaccuracy/bias for inventory control in semiconductor supply chain management.

Forstner and Mönch (2015) utilize a genetic algorithm (GA) and detailed simulation to assess delivery performance under optimal safety stock levels. However, their model's detailed nature makes it computationally intensive, limiting scalability to larger networks. Similarly, Persson et al. (2017) employ simulation to determine the necessary safety stock level amidst intermittent product demand. In contrast, Afridi et al. (2020) train a Deep Reinforcement Learning (DRL) agent to identify the optimal replenishment policy for Vendor Managed Inventory (VMI) settings, which does not directly apply to our context. Kim et al. (2022) investigate the necessity of a die bank in a large semiconductor manufacturer's supply chain using simulation but do not extensively address the impact of forecast accuracy in their models. Forecast accuracy is a critical factor influencing inventory control decisions, as highlighted by seminal studies such as Heath and Jackson (1994) and more recently by Nataraja et al. (2023). In a series of publications Manary and Willems (2008), Manary et al. (2009) and Manary et al. (2019) develop models to set safety stock targets when forecast bias is known, leading to significant increases in gross profits for Intel. Nakashima et al. (2014) focus on assessing inventory targets under various bias profiles, considering stochastic inventory control problems within uncertain supply chains. Furthermore, Norouzi and Uzsoy (2014) study the effect of rolling forecast uncertainty, while Albey et al. (2015) model the evolution of demand information using a Martingale Model of Forecast Evolution. Ziarnetzky et al. (2019) investigate production planning models for front-end wafer fabs with stochastic demand, highlighting the importance of incorporating safety stock

and production decisions under chance constraints. In a vendor-managed inventory (VMI) setting, Diaz et al. (2022) simulate the semiconductor crisis caused by customer behavior, emphasizing the role of demand reduction severity in the replenishment process and the bullwhip effect it creates. Similarly, Ehm et al. (2023) simulate disruptive scenarios on inventory models and the bullwhip effect in a vendor-managed setting.

We conclude that research on assessing inventory targets in semiconductor supply chains has been a subject of great interest in the literature. Although such problems have been studied in the semiconductor industry, few studies have addressed the need for practitioners to model many scenarios with minimal data management effort, computational expense, and effort to set up scenarios. Therefore, this study, which establishes an aggregated simulation model of the back-end processes, is original, practical and applicable to actual production planning. Additionally, the need for the industry to respond effectively to market up- and downswings using safety stock has not yet been investigated. To the best of our knowledge, this paper is the first with a case study that verifies the importance of inventory management considering a die bank through simulation during market up- and downswings.

3 METHOD

We propose an aggregated simulation model of the back-end part of semiconductor supply chains. The model is applied to real-world scenarios in section 4 for assessing safety stock target levels at the die bank, taking into account demand uncertainty and capacity constraints. This section consists of two parts. First, we propose to model demand uncertainty using future forecasts and discrete distributions of historical forecast errors. Second, we propose an aggregated simulation model of the back-end part of the semiconductor supply chain to simulate the supply of products to the end customers.

3.1 Demand Modeling

In this research, we include one stochastic element, demand uncertainty. To model future demand uncertainty we first compute the discrete distribution of historical forecast errors. The required input data is illustrated in Table 1. For each **Item** $p \in \{A, B, \dots, P\}$, we collect historical forecast errors. The **Requested Demand** for a certain **Demand Week** represents the actual quantities requested by the customer (not actual sales) to be delivered in that given week. Similarly, the **Forecast Quantity** reflects the capacity-unconstrained forecast for actual customer demand in that same Demand week. The **Look-ahead Period** indicates how far ahead in time the forecast for a particular Demand Week is made. So a forecast made for Demand Week 30 with a Look-ahead Period of 12 represents the forecast made in Week 18 for Demand Week 30. Forecast errors typically increase for forecasts made further out in time due to a growing number of uncertainties, therefore, it is important to consider the Look-ahead Period when analyzing forecast errors. Finally, the **Forecast Error** is computed by subtracting the requested demand from the forecasted demand value for each Demand Week. A positive forecast error indicates over forecasting, meaning the customer did not require as many products as anticipated. Conversely, a negative forecast error suggests that actual demand exceeds expectations, potentially leading to material shortages.

Table 1: Historical forecast errors.

Item	Demand Week	Requested Demand	Forecast Quantity	Forecast Error	Look-ahead Period
A	1	100	110	10	2
B	2	120	120	0	5
...
P	10	100	90	-10	3

As proposed by Vandepuut (2020), instead of looking at the forecast error per week (per Item), we look at the forecast error over multiple consecutive weeks (Look-ahead Periods). In other words, for each Item, we consider what could happen during the replenishment Lead Time (time it takes to order and manufacture additional patterned wafers) at the front-end. For example, if a product has a 12-week front-end Lead Time with weekly orders, we will look at the forecast error distribution over thirteen consecutive weeks. We remove exceptionally high forecast errors from the data by keeping only the forecast errors between the 5th and 95th percentiles of all computed forecast errors. Extreme forecast errors can be caused by multiple reasons (e.g. human error). Therefore, we remove extreme cases of over- and under forecasting from the data.

In this study we analyze two scenarios (market up- and downswings), the historical forecast errors are therefore split into parts where the market was in an up- or downswing. To determine the moments that the market was in an up- or downswing we apply the definition proposed by Aubry and Renou-Maissant (2013) and split all the datasets accordingly.

For each simulation run, we randomly sample from the historical forecast error distributions and add the sampled forecast errors to the future forecast. See Figure 2 for an example, the blue confidence intervals are based on the discrete values of historical forecast errors.

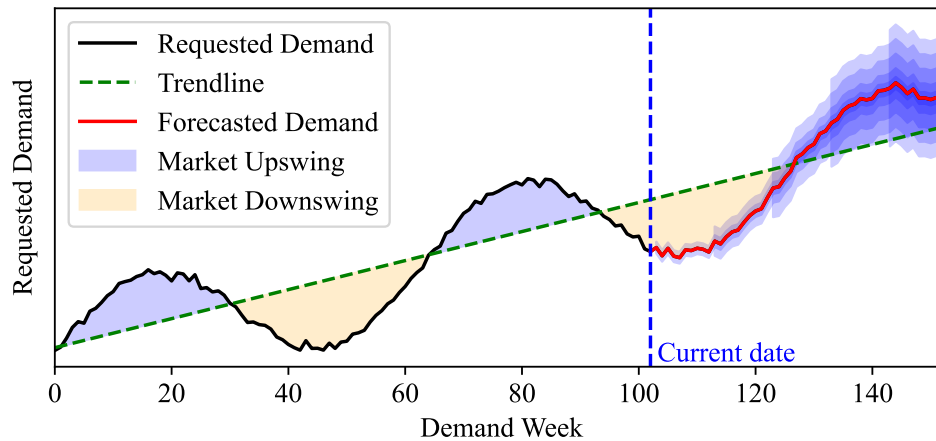


Figure 2: Example of future demand uncertainty.

3.2 Supply Modeling

In the semiconductor industry, the Customer Order Decoupling Point (CODP) is typically placed at the die bank (Kim and Kim 2012). By doing so, the CODP decouples the push production system in the front-end and the pull production system in the back-end. This positioning introduces a capacity constraint between customer demand and the CODP posed by the back-end, as limited capacity is available at the back-end. This constraint implies that a higher safety stock level does not guarantee a higher service level, as limited capacity will limit the amount of End Items that can be produced from stock. To deal with the capacity constraint, we propose an approach in which the evaluation of the effect of safety stock target levels on related KPIs is done by simulating the back-end supply chain and thus the capacity constraints. To reflect the required level of detail in the simulation model, we propose to use a Discrete Event Simulation (DES).

Detailed supply chain simulation models require large amounts of data and suffer from long computation times to produce statistically valid results (Mönch et al. 2018a). In this research, we propose to use an aggregated back-end supply chain simulation model based on existing master data structures. Since the goal of this study is to assess safety stock **targets** at the die bank, there is no need to model the front-end. We assume a "perfect" front-end process that produces the forecasted demand including safety stock. Using a

"higher" aggregation level reduces the required data maintenance effort and limits computational expense. A "higher" aggregation level does imply a lower accuracy of results compared to a real-life setting. We argue that this trade-off is acceptable because the use of the existing master data structures makes the model fast to set up and easier to adopt within the semiconductor industry.

Following the master data structure, we aggregate resources into capacity groups. Resources can only be grouped when they perform the same operation and are at the same production location. We further simplify the model by assuming deterministic capacity group parameters. As such, per consumed Item (patterned wafer) - End Item (chip) combination, each capacity group has a fixed yield, usage rate, cycle time, and weekly capacity. We assume that back-end cycle times are guaranteed under the given capacity restrictions. The different combinations, referred to as supply chains, are prioritized based on cycle time, where the shortest cycle time is assigned the highest priority. All supply chain parameters per End Item are summarized in Table 2.

Table 2: Supply chain input parameters.

Supply chain input parameters	
End Item	Supply chains
Supply chain	Location consumed Item
	Production location
	Cycle time
	Priority
	Capacity groups
Capacity group	Weekly capacity
	Yield
	Usage rate

In Figure 3, we present an example of an aggregated back-end supply chain. A supply chain is defined as a connection between two stock points, in this case, the die bank and an Internal Warehouse (IWH) to store End Items before shipment to customers. All supply chains include the back-end operations: Logistics Group (LG), Wire-Bonding (WB), Molding (MD), and Final Testing (FT). Each of these operations is linked to the corresponding capacity group (see hexagons) with an available capacity denoted in hours per week. Furthermore, each supply chain has a yield, which we assume to be fixed over time. The CODP for all Items is located at the die bank.




Figure 3: Back-end supply chain model.

Other than the supply chain parameters discussed above and the demand modeling discussed in subsection 3.1, further input parameters to the simulation model are the safety stock target levels per End Item at the CODP. Safety stock target levels are expressed in coverage weeks of demand. The safety stock target level can be calculated as $SS \text{ target level} = (\text{sum of all demand over the risk horizon} / \text{number of weeks in the risk horizon}) * \text{coverage weeks}$. The safety stock target levels need to be set by the user and are the only variable in the simulation model.

As forecast, demand, and capacity have a weekly granularity, the simulation model uses weekly time buckets. Each week, consumed Item inventory at the CODP is replenished and End Item demand occurs. The

quantity of inventory replenishment, or pipeline inventory, is determined using the reorder policy defined as $reorder\ quantity = SS\ target\ level + forecast\ over\ replenishment\ lead\ time - pipeline\ inv - inv\ on\ hand$.

We assume that the quantity of consumed Items is determined using the reorder policy and is completely delivered by the front-end with a constant Lead Time. Additionally, we assume the following;

- An End Item can only be produced when enough capacity and required patterned wafers are available.
- When an End Item can be scheduled in one week of demand occurrence, it is considered to be on time.
- End Items that cannot be produced on time will be delayed and prioritized over new incoming demand.
- Within the forecast data, we distinguish between market up- and downswings. We assume that, within each dataset, the quality of forecast data does not change over time.
- The simulation model is a steady-state model. A warm-up interval with the length of the consumed Item replenishment Lead Time is used.
- End Items are scheduled according to the first-come, first-served (FCFS) queuing principle.

4 APPLICATION

In this section, we apply the proposed methodology from section 3 to high-runner products from NXP Semiconductors N.V. We study two real-world scenarios, one where the market was in an upswing and one where the market was in a downswing. We test a different set of safety stock target levels for each scenario to study the impact on OTD and inventory position performance indicators.

4.1 Demand Uncertainty

For all End Items in the problem set, we compute the discrete distribution of historical forecast errors (using real-world data) in both market up- and downswing scenarios, see Figure 4. We analyze data from all up- and downswings that occurred between the beginning of 2017 and the end of 2023. The relative forecast error (%) on the y-axis indicates the percentage difference between the forecast (from the past) and the actual requested demand. A positive forecast error means that we are over forecasting, while a negative relative forecast error means we are under forecasting. Over forecasting can lead to inventory building at the die bank and increases the risk of obsolete stock, while under forecasting can result in material shortages and a decrease in On-Time Delivery performance. In Figure 4, the boxes represent data points between the 25th and 75th percentile while the whiskers represent data points between the 5th and 95th percentiles. Any data points outside the range of the whiskers are considered outliers and are represented as dots in the figure. The outliers will be removed from the data.

The figure indicates that the risk of over forecasting is higher during market downswings. As mentioned in section 1, this can be explained by the delay in adjusting the forecasting models when demand is going down. During market upswings, the forecast bias reverses and there is a higher probability of under forecasting, possibly leading to stock outs and late deliveries.

We conclude that the market phase (up- or downswing) needs to be taken into account when modeling demand uncertainty using forecast errors. Overfitting of historical forecast errors can occur when sampling from the entire set of forecast errors while the market is in a certain phase. Therefore, we model two scenarios, market up- or downswing in subsection 4.2 reducing the risk of overfitting past forecast errors.

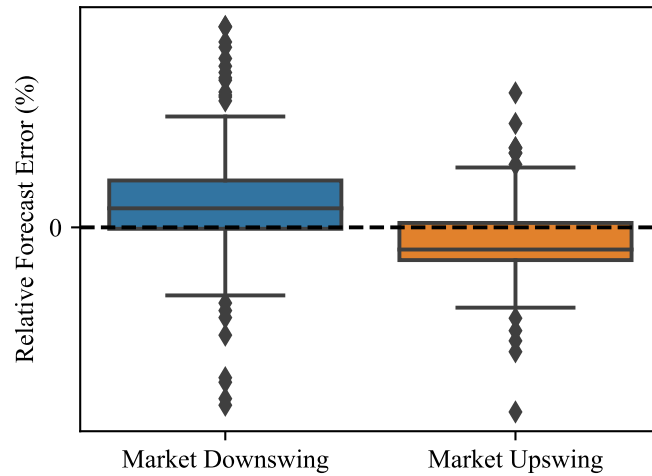


Figure 4: Historical forecast errors per market phase.

4.2 Results

We apply our method to assess the effect of various product-specific safety stock target levels on the performance indicators OTD and inventory position at the die bank during market up- and downswings. We define OTD as fulfilled demand / total demand. Safety stock targets, expressed as coverage weeks, typically range from 0 to 2 weeks for high-runner End Items and can be set per End Item. Each scenario (market up- or downswing) tests all combinations of three safety stock target levels (0, 1, and 2 coverage weeks) per End Item, each run 100 times. The trade-off between OTD and inventory position is shown in Figure 5 for both scenarios. Crosses represent mean values with coefficient of variance, normalized between 0 and 1 for both OTD and inventory position. Thicker crosses indicate Pareto optimal solutions, representing the optimal set. More inventory generally improves OTD until capacity constraints limit further improvement, with greater inventory leading to increased OTD certainty. Comparing market up- and downswings reveals significantly shifted optimal fronts, with less inventory needed due to lower demand during downswings and biased forecasts toward underestimation during upswings necessitating more inventory. Based on these insights, practitioners can select safety stock targets to analyze further. For example, if the strategic target is a minimum OTD performance of 0.8, practitioners can identify the closest case on the Pareto front. Typically, the two closest points are considered for further analysis to understand the desired trade-off.




Figure 5: Pareto set showing optimal safety stock level settings.

In Figure 6 we show detailed simulation results where inventory position and OTD performance over time are compared for both up- and downswing market conditions. For each market condition, the two closest points to 0.8 OTD are selected. For each point, OTD and inventory position mean and 95% CI over time are shown for one End Item. For both market conditions, the mean of the point closest to 0.8 OTD (scenario 1 in Figure 6) is above the threshold. When analyzing the coefficient of variance, due to uncertainty, a value below 0.8 is possible. Therefore, a second point from the optimal set (scenario 2 in Figure 6) is analyzed. Although having a higher inventory position, the second point has a lower variance for both performance indicators creating more certainty for decision makers. This result is reflected in the detailed simulation results. For both performance indicators, we see that scenario 2 has a smaller CI and thus provides a higher certainty to the decision maker that the desired performance will be achieved. Furthermore, when up- and downswing results in the detailed plots are compared, the plots show that in a market downswing, both inventory position and OTD are more stable over time. Due to lower demand during a market downswing, less inventory is needed and the capacity constraint is less significant, leading to a more stable delivery performance than in a market upswing. This is desirable because a stable OTD guarantees delivery performance to customers, whereas the unstable upswing results show that OTD is below the threshold for multiple periods. In a market upswing, managers can therefore improve delivery performance by keeping inventory position higher compared to a market downswing. A higher inventory position, achieved by higher safety stock level settings, can mitigate biased forecasts toward under forecasting. Furthermore, capacity should be increased to meet higher demand. As the detailed plots show, ample capacity during downswing results in more stable performance indicators. It is expected that the same holds for market upswing conditions.

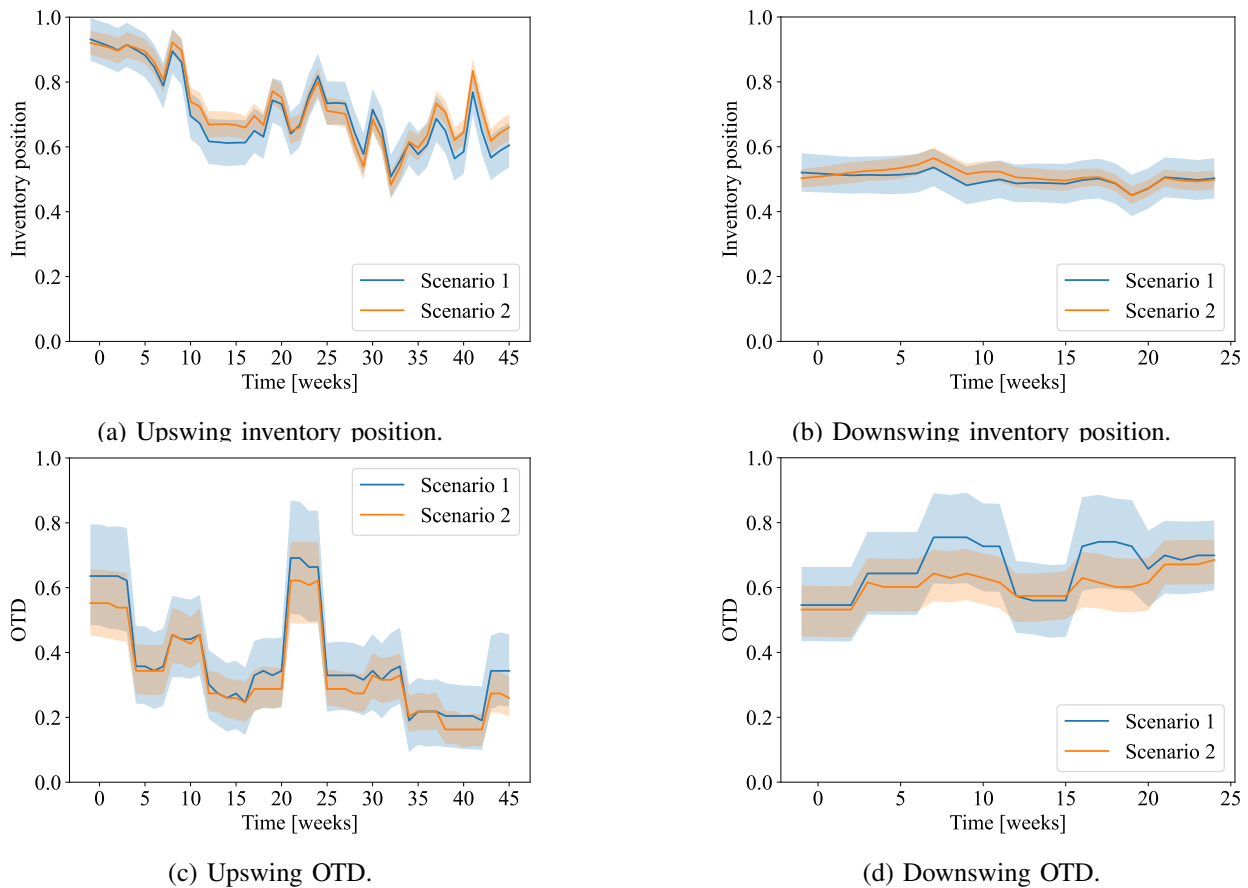


Figure 6: OTD and inventory position over time for the two scenarios closest to 0.8 OTD.

5 DISCUSSION

Originality. We propose an aggregated simulation model of the back-end semiconductor supply chain. The aggregation level is based on the existing master data structure of NXP Semiconductors N.V., eliminating the need for additional master data management effort to maintain the simulation model. We model demand uncertainty and demand by using discrete distributions of historical forecast errors and future forecasts. The use of historical forecast errors to model future demand uncertainty is well known in the inventory control literature, however, combining this approach with an aggregated simulation model of the back-end manufacturing process has not yet been studied in the literature to our knowledge.

Methodology. We model the back-end part of the semiconductor supply chain on an aggregated level (based on existing master data) to obtain the following benefits: eliminate the need for additional (master) data maintenance, reduce computational expense, and minimize effort to set up/analyze different supply/demand scenarios. To simplify the simulation model further, we made several assumptions as listed in subsection 3.2. The potential disadvantage of the proposed aggregation level and simplifying assumptions is that important details can be aggregated away, possibly leading to less accurate results. In this study, we do not compare our results with more detailed simulation approaches. We validated the initial simulation results based on expert feedback. The validation of input data is still pending.

We model demand uncertainty by using historical forecast errors because it is a proven concept in the literature and the required data is widely available in the industry. However, when applying this approach there is a risk of overfitting past forecast errors on future demand. We assume demand forecasting will not improve and therefore historical forecast error distributions are still relevant for the future.

Application. We apply the proposed simulation model to assess the need for safety stock of high-runner Items during a market up- or downswings to demonstrate the impact of market phases on forecast bias and subsequently on the need for safety stock. The proposed simulation model can be applied for assessing safety stock targets for other scenarios and Items as well. For example, analyzing the need for safety stock during product ramp-up or ramp-down is an interesting case to study. We do recommend practitioners to carefully analyze the historical forecast error distributions when setting up new scenarios to prevent overfitting (e.g. market up- or downswings). The conclusions we draw based on the application are based on data from NXP Semiconductors N.V., therefore, when replicating the experiments at other companies the conclusions need to be reevaluated.

6 MANAGERIAL IMPLICATIONS

The proposed aggregated simulation approach allows practitioners to test the impact of a range of safety stock target levels on On-Time Delivery (OTD) performance in different supply/demand scenarios. Due to the stochastic nature of the simulation model practitioners can analyze the certainty of reaching their performance targets as well as the mean expected performance. This allows practitioners to balance between "optimal" and "certain" supply chain performance. Additionally, practitioners can study the impact of forecast bias on the operational performance of the supply chain. The above-mentioned abilities help practitioners to prepare for and adapt to supply/demand scenarios, in turn improving their potential for making qualitative decisions on safety stock target levels that are in line with the objectives of the company. Improving safety stock targets can greatly impact the supply chain performance of the company, balancing the costs of holding inventory with On-Time Delivery performance.

By only modeling key bottleneck resources that are already aligned with the current master data structure of the company, we establish a simulation approach that requires no additional data maintenance efforts from the company to maintain. Additionally, the aggregated approach reduces computational expense compared to more detailed modeling approaches.

In this study, we focus on safety stock setting, however, there are other interesting applications (e.g. high-level capacity decisions) of our proposed simulation model. For practitioners, it would be of great interest to further explore the potential applications of the proposed aggregated simulation model.

7 CONCLUSIONS

In this study, we propose an aggregated simulation model of the back-end semiconductor supply chain. To model future demand uncertainty we use historical distributions of forecast errors and future forecasts. We demonstrate our proposed method by analyzing safety stock requirements during market up- or downswings, taking into account potential forecast bias and inaccuracy.

After analyzing all forecast errors made between 2017 and 2023, we conclude that the forecast in our problem subset is biased and that the forecast bias is correlated to the market phase (up- or downswings). During market upswings, the forecast is biased towards over forecasting, while during market downswings the forecast is biased towards under forecasting. Based on our simulation results, we conclude that the trade-off between safety stock targets and On-Time Delivery (OTD) performance behaves significantly different during a market up- or downswing. The results show that during a market downswing (over forecasting) there is less need for safety stock. Over forecasting triggers more wafers to be started in the front-end, hence, decreasing the need for high safety stock targets to deal with demand variation. During a market upswing (under forecasting) the opposite is true, resulting in higher required safety targets to achieve the desired OTD performance.

In future research, it would be of great interest to expand the study by including Items in ramp-up or ramp-down life cycle phases in the application and applying optimization algorithms (e.g. Genetic Algorithms) to the safety stock target level setting. Furthermore, on the current aggregation level improvements can be made to the simulation model by relaxing assumptions on fixed cycle time and fixed yields. To further validate the performance (e.g. runtime and accuracy) of our proposed simulation model, it is of great future interest to compare our results to that of more detailed modeling approaches. This study's conclusions rely on data from one company. Future research should explore applying the method to data from other companies to validate the results across the semiconductor industry.

ACKNOWLEDGMENTS

The project AIMS5.0 is supported by the Chips Joint Undertaking and its members, including the top-up funding by National Funding Authorities from involved countries under grant agreement no. 101112089.

REFERENCES

- Afridi, M. T., S. Nieto-Isaza, H. Ehm, T. Ponsignon and A. Hamed. 2020. "A Deep Reinforcement Learning Approach for Optimal Replenishment Policy in A Vendor Managed Inventory Setting For Semiconductors". In *2020 Winter Simulation Conference (WSC)*, 1753–1764 <https://doi.org/10.1109/WSC48552.2020.9384048>.
- Albey, E., A. Norouzi, K. G. Kempf, and R. Uzsoy. 2015. "Demand Modeling with Forecast Evolution: an Application to Production Planning". *IEEE Transactions on Semiconductor Manufacturing* 28(3):374–384.
- Aubry, M. and P. Renou-Maissant. 2013. "Investigating the Semiconductor Industry Cycles". *Applied Economics* 45(21):3058–3067.
- Diaz, M. F. L., H. Ehm, and A. Ismail. 2022. "Simulated-Based Analysis of Recovery Actions Under Vendor-Managed Inventory Amid Black Swan Disruptions in the Semiconductor Industry: A Case Study from Infineon Technologies AG". In *2022 Winter Simulation Conference (WSC)*, 3513–3524 <https://doi.org/10.1109/WSC57314.2022.10015305>.
- Ehm, H., C. H. Chung, S. K. Chowdhury, M. Ratusny and A. Ismail. 2023. "The Bullwhip Effect in End-To-End Supply Chains: The Impact of Reach-Based Replenishment Policies with a Long Cycle Time Supplier". In *2023 Winter Simulation Conference (WSC)*, 2194–2205 <https://doi.org/10.1109/WSC60868.2023.10407267>.
- Forstner, L. and L. Mönch. 2015. "Using simulation-based optimization to determine production strategies and safety stock levels in semiconductor supply chains". In *2015 IEEE International Conference on Automation Science and Engineering (CASE)*, 655–656.
- Heath, D. C. and P. L. Jackson. 1994. "Modeling the Evolution of Demand Forecasts ITH Application to Safety Stock Analysis in Production/Distribution Systems". *IIE transactions* 26(3):17–30.
- Kim, D., Y. S. Park, H. Kim, K. S. Park and I. Moon. 2022. "Inventory Policy for Postponement Strategy in the Semiconductor Industry with a Die Bank". *Simulation Modelling Practice and Theory* 117:102498.
- Kim, J. and S. Kim. 2012. "Positioning a Decoupling Point in a Semiconductor Supply Chain under Demand and Lead Time Uncertainty". *International Journal of Advanced Logistics* 1(2):31–45.

- Manary, M. P., B. Wieland, S. P. Willems, and K. G. Kempf. 2019. "Analytics Makes Inventory Planning a Lights-Out Activity at Intel Corporation". *INFORMS Journal on Applied Analytics* 49(1):52–63.
- Manary, M. P. and S. P. Willems. 2008. "Setting Safety-Stock Targets at Intel in the Presence of Forecast Bias". *Interfaces* 38(2):112–122.
- Manary, M. P., S. P. Willems, and A. F. Shihata. 2009. "Correcting Heterogeneous and Biased Forecast Error at Intel for Supply Chain Optimization". *Interfaces* 39(5):415–427.
- Mönch, L., R. Uzsoy, and J. W. Fowler. 2018a. "A Survey of Semiconductor Supply Chain Models Part I: Semiconductor Supply Chains, Strategic Network Design, and Supply Chain Simulation". *International Journal of Production Research* 56(13):4524–4545.
- Mönch, L., R. Uzsoy, and J. W. Fowler. 2018b. "A Survey of Semiconductor Supply Chain Models Part III: Master Planning, Production Planning, and Demand Fulfilment". *International Journal of Production Research* 56(13):4565–4584.
- Nakashima, K., T. Sornmanpong, H. Ehm, and G. Yachi. 2014. "Comparison between Two Types of Inventory Targets under Variability of a Semiconductor Supply Chain". *International Journal of Business and Economics Research* 13(1):74–80.
- Nataraja, B. M., N. Singh, I. Adan, T. Aggarwal and K. Herps. 2023. "Using Simulation to Assess the Reliability of Forecasts in High-Tech Industry". In *2023 Winter Simulation Conference (WSC)*, 817–828. IEEE.
- Norouzi, A. and R. Uzsoy. 2014. "Modeling the Evolution of Dependency between Demands, with Application to Inventory Planning". *IIE Transactions* 46(1):55–66.
- NXP Semiconductors N.V. 2024. "Semiconductor manufacturing process". Accessed 9th April 2024.
- Persson, F., M. Axelsson, F. Edlund, C. Lanshed, A. Lindstrom and F. Persson. 2017. "Using simulation to determine the safety stock level for intermittent demand". In *2017 Winter Simulation Conference (WSC)*, 3768–3779 <https://doi.org/10.1109/WSC.2017.8248089>.
- Rapp, H. P., J. Möbert, S. Schneider, and D. B. AG. 2022. "Extraordinary Semiconductor Cycle Triggered by One-Time Events, Cyclical and Geopolitical Effects". *Germany Monitor: Deutsche Bank Research*.
- Uzsoy, R., J. W. Fowler, and L. Mönch. 2018. "A Survey of Semiconductor Supply Chain Models Part II: Demand Planning, Inventory Management, and Capacity Planning". *International Journal of Production Research* 56(13):4546–4564.
- Vandeput, N. 2020. *Inventory optimization: Models and simulations*. Walter de Gruyter GmbH & Co KG.
- Wang, C.-H. and J.-Y. Chen. 2019. "Demand Forecasting and Financial Estimation Considering the Interactive Dynamics of Semiconductor Supply-Chain Companies". *Computers & Industrial Engineering* 138:106104.
- Ziarnetzky, T., L. Mönch, and R. Uzsoy. 2019. "Simulation-Based Performance Assessment of Production Planning Models with Safety Stock and Forecast Evolution in Semiconductor Wafer Fabrication". *IEEE Transactions on Semiconductor Manufacturing* 33(1):1–12.

AUTHOR BIOGRAPHIES

CAS ROSMAN is a doctoral candidate in the Department of Planning IT at NXP Semiconductors N.V. and in the Department of Industrial Engineering and Innovation Sciences at the Eindhoven University of Technology. He obtained his master's degree in Operations Management & Logistics at the Eindhoven University of Technology. His primary research interest is in the area of simulation modeling and decentralized decision-making in supply chain management. His email address is cas.rosman@nxp.com.

ERIC WEIJERS is a doctoral candidate in the Department of Supply Chain Operations at NXP Semiconductors N.V. and in the Department of Industrial Engineering and Innovation Sciences at the Eindhoven University of Technology. He obtained his master's degree in Operations Management & Logistics at the Eindhoven University of Technology. His primary research interest is in the area of simulation modeling and optimization in supply chain management. His email address is eric.weijers@nxp.com.

KAI SCHELTHOFF is the head of the Supply Chain Innovation team at NXP Semiconductors N.V. He obtained his PhD at the Karlsruhe Institute of Technology about data-driven cycle time estimation in semiconductor wafer fabrication using a concatenated machine learning approach, in collaboration with Robert Bosch GmbH. He is an expert in machine learning applications for estimation, classification, and optimization in Supply Chain Operations and Manufacturing. His email address is kai.schelthoff@nxp.com.

WILLEM VAN JAARSVELD is an Associate Professor in Stochastic Optimization and Machine Learning at Eindhoven University of Technology (TU/e). His main research interest is stochastic optimization, using a diverse set of methodologies including Deep Reinforcement Learning, Stochastic Programming and Dynamic Programming. Applications areas include data-driven inventory control, production planning, supply chain management, and maintenance logistics. His email address is w.l.v.jaarsveld@tue.nl.

ALP AKCAY is an Associate Professor in the Operations, Planning, Accounting, and Control (OPAC) group at Eindhoven University of Technology. His research interests include statistical decision-making under uncertainty, simulation design and analysis, and approximate dynamic programming with applications in manufacturing, maintenance, and supply chain management. His e-mail address is a.e.akcay@tue.nl.

IVO ADAN is a Full Professor in the section Operations, Planning, Accounting and Control (department of Industrial Engineering & Innovation Sciences) at Eindhoven University of Technology (TU/e) and holds the Manufacturing Networks chair. His expertise and tuition areas include probability theory / statistics, operations research, manufacturing networks, stochastic operations research, and queueing models. His e-mail address is i.adan@tue.nl.