Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

A BAYESIAN SIMULATION APPROACH FOR SUPPLY CHAIN SYNCHRONIZATION

Bianica Pires
Joshua Goldstein
Dave Higdon
Paul Sabin
Gizem Korkmaz
Stephanie Shipp
Sallie Keller

Social and Decision Analytics Laboratory
Biocomplexity Institute of Virginia Tech
Arlington, VA 22203, USA

Shan Ba
Ken Hamall
Art Koehler

The Procter & Gamble Company
Cincinnati, OH 45202, USA

Shane Reese

Brigham Young University
Provo, UT 84602, USA

ABSTRACT

While simulation has been used extensively to model supply chain processes, the use of a Bayesian approach has been limited. However, Bayesian modeling brings key advantages, especially in cases of uncertainty. In this paper, we develop a data informatics model that could be used to realize a digital synchronized supply chain. To realize this model, we take a hybrid approach that combines Bayesian modeling with discrete-event simulation and apply it to the supply chain process of a Procter & Gamble (P&G) manufacturing and distribution facility. We use approximately one year of transactional data to inform our model, including information on customer orders, production, raw materials, inventory, and shipments. A driving force for creating this model is to better understand and manage the balance between inventory, profit, and service.

1 INTRODUCTION

While simulation has been widely used to model supply chain processes (Jahangirian et al. 2010), informing these models through the use of transactional data has not been extensively explored. This could in part be due to the difficulties around stakeholder engagement since simulations often require extensive data gathering and processing (McNaught and Chan 2011). Recent advances in data collection and storage, however, facilitate the integration of these new data sources with simulation models. Bayesian statistical methods have proven to be a useful framework for combining observational data with simulation-based models while accounting for the various sources of uncertainty (Reese et al. 2004) – a key characteristic throughout many parts of the supply chain (McNaught and Chan 2011). Bayesian approaches have been applied to various areas of manufacturing, such as fault diagnosis (e.g., Jin et al. 2012), quality control (e.g., Correa et al. 2009), reliability (e.g., Li and Meeker 2014, Celik and Son 2010), and supply chain disruptions (e.g., Soberanis 2010). Moreover, hybrid approaches that combine two or more simulation techniques (e.g., discrete-event simulation, system dynamics) have seen a rise in popularity due to the increased trend in providing "enterprise wide solutions" that take into account the impact that different parts of an organization have on one another (Jahangirian et al. 2010). Other hybrid approaches have explored the use of these simulation techniques with Bayesian modeling (e.g., Xu and Son 2013). Chick

(2004) , in particular, stresses the benefits of combining Bayesian methods with discrete-event simulation, including for input modeling, response surface modeling, and uncertainty analysis. While such previous models have explored hybrid methods, we integrate Bayesian modeling and discrete-event simulation with disparate big data streams across a supply chain system (from customer orders to deliveries). This allows us to better understand and manage the uncertainty, variation, and dependencies in the supply chain and to make predictions of behaviors under new conditions.

In this paper, we describe the development of a data informatics model that could be used to realize a digital synchronized supply chain. To realize this model, we take a hybrid approach that combines Bayesian modeling with discrete-event simulation and apply it to the supply chain process of a Procter & Gamble (P&G) facility that carries out both manufacturing operations and distribution activities. We inform the model using approximately one year of transactional data. A driving force for creating this model is to better understand the balance between inventory, profit, and service. It should be noted that while we simulate the main processes between customer orders and customer shipments, a complete supply chain model would also need to include data on consumers and raw material suppliers. In the remainder of this paper, we discuss the details for developing the model and computing profit and service (Section 2), illustrate how the simulation can be used to explore optimal production planning strategies (Section 3), and summarize the paper (Section 4).

2 METHODOLOGY

The methodology for developing the data informatics model of the supply chain process of a P&G facility involved three high-level steps: (1) preparation of the data used to inform the model, (2) development of four simulators that together capture the supply chain process from customer orders to deliveries, and (3) calculation of actual and simulated profits to validate the model. This section discusses each of these steps in more detail. Note that due to the sensitivity of the data, we have removed information that might disclose the dates, location, units, and details on data structures such as table names and fields. Also, we use the terms product to refer to an item with a unique stock keeping unit (SKU) and product grouping to refer to one or more products which can be placed in the same product category.

2.1 Informing Models Through Data

The development of a data model is critical to understanding data flows and informing the development of the simulation models. We used one year of transactional data, including information on customer orders, production, raw materials, inventory, shipments, and deliveries. This resulted in the gathering of over 20 tables, some of which contained millions of data records. The data model consists of a set of clean, accurate, and reliable tables, where the relationships between tables are understood and tables have been stored in a centralized location for easy access for model building.

Development of the data model was a highly iterative process between subject-matter experts, data experts, and modelers. It consisted of the following high-level steps: (1) data source discovery and acquisition, which included identifying and obtaining internal data to use in modeling; (2) data profiling to understand and describe the data; (3) data cleaning and restructuring for purposes of analysis; (4) data linkages to understand the relationships between tables; and (4) data exploration, including visualization, generation of descriptive statistics, and analysis of patterns and inconsistencies. As is often the case with data-driven model building (Dasu and Johnson 2003), over half of the project's efforts went towards development of the data model. Figure 1 is a notional diagram of the transactional data across different silos of the data (e.g., orders, production). Due to the sensitivity of the data, the data model with table names, fields, and detailed linkages is not shown.

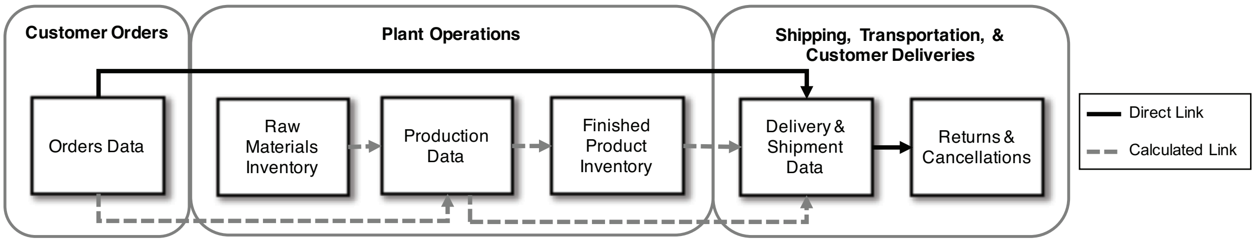

Figure 1: A notional diagram of the transactional data used to model the supply chain system. Direct links are established via unique identifiers. Calculated links are estimated through the data.

2.2 The Data-Informed Simulation Model

Using a combination of Bayesian modeling and discrete-event simulation, we developed a data informatics model to simulate the supply chain process of a P&G facility. To capture the processes from customer orders to deliveries, the model is broken out into four simulators: the Orders Simulator, the Production Simulator, the Production Planner, and the Shipment Simulator.

2.2.1 Orders Simulator

The Orders Simulator is constructed so that we can simulate how orders might change under different scenarios (e.g., altered prices, different order incentives, changes in product lineup, different customer make-up). It estimates volumes at aggregated levels (e.g., by customer, by product grouping) using Bayesian dynamic linear models (Harrison and West 1999). Given these aggregated forecasts, the number of orders, the product mix, and the order volume are derived for each customer from historical data that are altered to ensure aggregates match the forecasts. This alteration of orders is carried out using Bayesian updating.

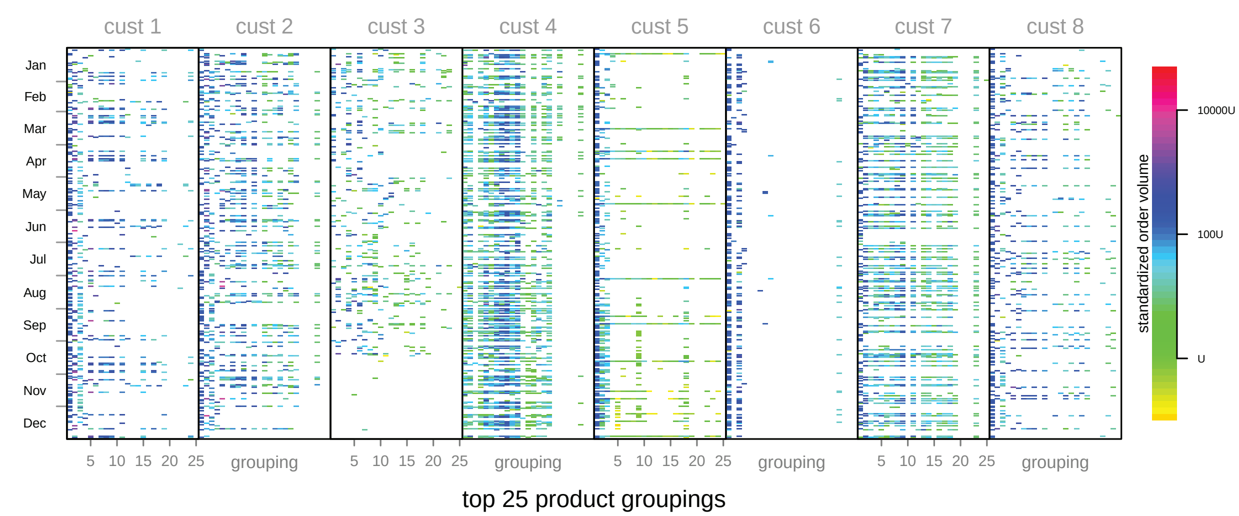

Figure 2: Individual orders from eight customers over the course of a year. Each row of the image shows a single order's volume for each of the top 25 product groupings. Each column of the image shows an order from the corresponding day of the year. Hence each image is 365×25.

Thousands of orders arrive daily for products that are produced at and/or shipped from this P&G facility. While individual orders differ depending on replenishment needs, prices, promotions, and incentives, different customers show unique patterns in the types and volumes of products they order. Figure 2 shows the pattern of orders by day over the course of a year for several high volume customers. Each pixel in the figure represents a product grouping and the color represents the volume ordered. The figure shows that the regularity and variety of products ordered varies by customer. While order content and volume can show substantial variation from one order to the next, the overall volume of orders has a fair amount of regularity when aggregated by product grouping and month (as shown in Figure 3). Hence the Order Simulator will need to capture this regularity at aggregated levels, while also reproducing the individual variation present in Figure 2.

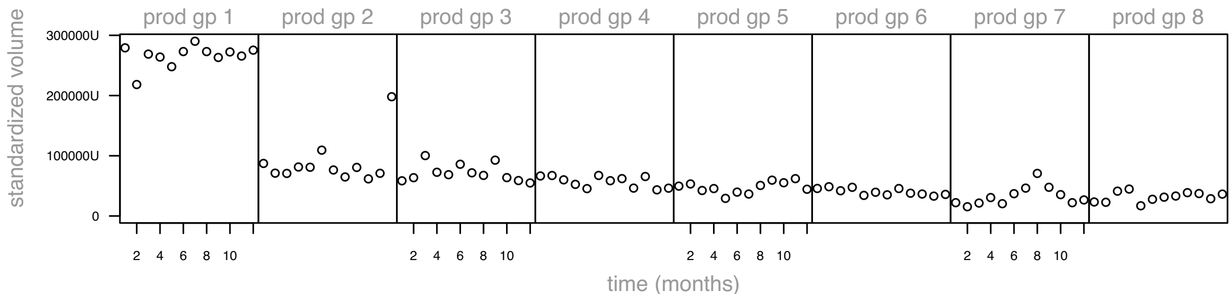

Figure 3: Monthly order volumes aggregated for each of 8 product groupings. Dynamic linear models are used to forecast the total monthly volume.

The orders are simulated using a combination of Bayesian forecasting at an aggregated level and empirically-based estimation at the detailed customer × product × order level. The approach for simulating a month of orders is outlined below.

1. A time-series model (Harrison and West 1999) is used to forecast the total volume of orders – aggregated over all customers and product groups – expected for the distribution center for the month.
2. This total volume is distributed over customers and product groupings based on the historical distribution from the previous month. This distribution could be estimated from an alternative time window or could account for known changes that will be present for the month being forecasted.
3. For each forecasted volume by customer, a set of orders that match previous customer order behavior are constructed that sums to the forecasted volume by customer and product grouping. Below are the details for this step:
   (a) For each customer, estimate the monthly order rate μ using historical data (typically the previous month). The simulated total order volume for the month is drawn from the resulting predictive distribution for the model.
   (b) Sample the estimated number of customer orders N from a Poisson distribution with mean μ.
   (c) Alter the ordered amounts in the sampled orders so that the volume aggregates appropriately over each of the product groups. This is done by treating the collection of orders as a vector random variable $\{v_{po}\}$ where $p = 1, \ldots, P$ indexes product, and $o = 1, \ldots, N$ indexes order. We construct a distribution for this vector by combining distributions that characterize different features of this collection of customer orders. A description of these distributions is shown in Table 1. These densities are combined using Bayes rule (Dawid et al. 1995), and then sampled using Markov chain Monte Carlo (Robert and Casella 1999).
   (d) Change the order dates of the sampled orders to reflect dates compatible with the month being simulated.

This algorithm results in a collection of orders that (1) reflects previous order specific product grouping behavior for each customer, (2) matches specified totals by product grouping and customer for that month, and (3) matches the forecasted grand total volume for the month. Figure 4 shows a number of simulated orders. Here the forecast is based on data up to the month being forecasted. The figure shows actual and simulated cumulative order volumes for seven different products over the course of one month. It is important that the simulated orders capture the monthly and daily variations seen in the actual orders so that the optimization of production planning and inventory levels will be robust to variation in the ordering.

Table 1: Distributions used to characterize the different features of customer orders.

density	description
$p(v_{po})$	Independent normal distributions for each v_{po} that is centered at the actual order amount, with standard deviation of 15%.
$f(\sum_{p\in G_k,o} v_{po})$	For each product group G_k, $f(\cdot)$ is a normal density with mean equal to the total volume determined in step 2, and a standard deviation of 10%.
$I[v_{po} \geq 0]$	Constraints ensuring that each order volume is either 0 or positive.

Using this holdout approach, we find that the simulated order timings and volume match to within about 5% for very regular products (products 1-4 in Figure 4). Less regular products can show substantial variation in order timing and volume over the course of a month (products 5-7 in Figure 4). For example, product 5 in Figure 4 sees more than half its monthly volume arrive on a single day. The fitted order simulator generally recognizes these high variation products, producing substantial variation in the simulated order histories.

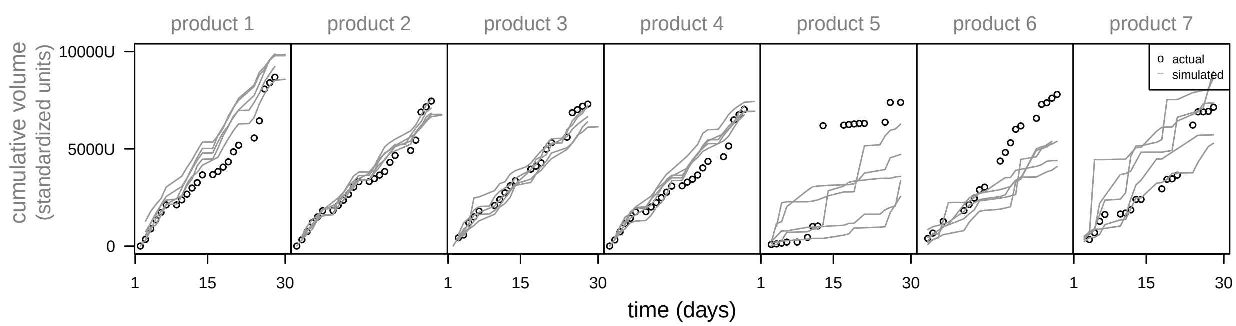

Figure 4: Cumulative order amounts for seven products over the course of a month. Simulations are shown by the gray lines; the actual orders are given by the black circles. These simulations are trained with data prior to the actual month shown.

2.2.2 Production Simulator

A substantial proportion of customer order volume are for products that are produced on-site. As orders arrive, a key question is how to allocate inventory, production time, and resources to fill these orders in a timely and cost-effective manner. The Production Simulator simulates the time required to produce a specified volume of any given product with variation in time and raw material usage that is consistent with historical data. This is a key tool in identifying production planning strategies to manage costs and product inventory in the face of variation in order volume and composition, as well as in production timing and efficiency.

Production occurs year round for 24 hours a day; it is affected by seasonality, time of the day, shift changes, which assembly lines are being used, and which product is being produced. Changeovers, where the line moves from producing one product to another can delay production. Some changeovers are relatively simple, such as when the items being produced remain the same but only the packaging or quantity changes, while other changeovers are more involved, requiring a different set of raw materials and line setup. Historical data allow us to estimate the changeover time distribution as well as production rates.

Production rates are estimated using a year of production data in which the start time, end time, and product are recorded for each production run on each assembly line. The production rate $\lambda_i(t)$ for each production run i, which starts at time t_{i0} and finishes at t_{i1}. This production rate depends on product k_i, assembly line ℓ_i, as well as temporal effects determined by the interval (t_{i0}, t_{i1}): season (day of year), day of the week, and time of day. The production rate $\lambda_i(t)$ is modeled as a linear combination of effects

$$\lambda_i(t) = \lambda_0 + \lambda_{k_i} + \lambda_{\ell_i} + \lambda_{\text{season}}(t) + \lambda_{\text{day}}(t) + \lambda_{\text{hour}}(t),$$

where λ_0, λ_{k_i}, and λ_{ℓ_i} are constants, accounting for the overall, product, and line effects on the production rate, $\lambda_{\text{season}}(t)$ is a term that varies smoothly over the course of the year (modeled as a spline (Higdon 2002)), $\lambda_{\text{day}}(t)$ is a constant for each day of the week, and $\lambda_{\text{hour}}(t)$ is a constant for each hour of the day. The amount produced y_i by run i is modeled as

$$y_i = \int_{t_{i0}}^{t_{i1}} \lambda_i(t)dt + \varepsilon_i,$$

where the error term ε_i accounts for additional variation between runs. These errors are modeled as independent and identically distributed (iid) $N(0, \sigma_\varepsilon^2)$ random variables.

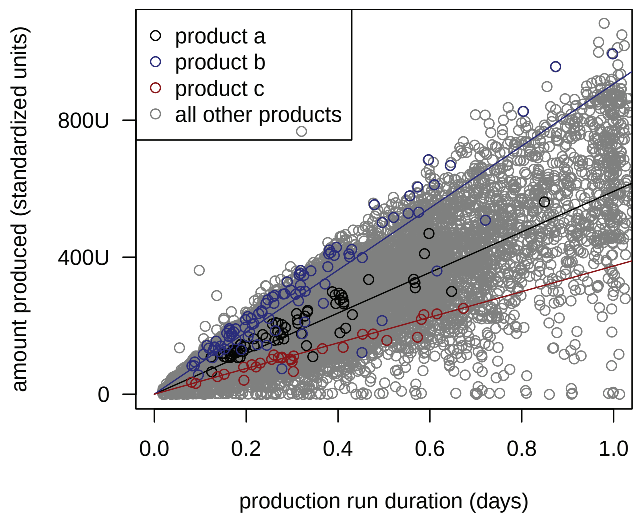

Figure 5: Production rate estimates (slope of the lines) from historical production data (circles). Each circle denotes a historical production run. The slope of the three lines show the estimated average production rate for three different products. The colored circles show results of production runs involving these three products; grey circles show all other production runs.

The unknown parameters – rate terms and σ_ε^2 – are estimated using a Bayesian hierarchical model (Gelman et al. 2004). In addition, we estimate uncertainty for each of these parameters. Figure 5 shows the estimated average production rates along with historical production runs for three different products.

The Production Simulator then uses these estimated rates, their uncertainties, and estimates of variation between similar production runs to estimate the outcome of any planned production run. The Production Simulator simulates the time required to produce a specified volume of any given product and/or the volume produced for a given start and stop time for a production run.

2.2.3 Production Planner

The Production Simulator interacts with the Orders Simulator through the Production Planner, which incorporates rules for scheduled production and adapts to the demand of incoming orders as shown in Figure 6.

The Production Planner has the flexibility to restrict the products produced on an assembly line, to speed up or slow down production on a line, and to remove lines. Inputs include the length of a production

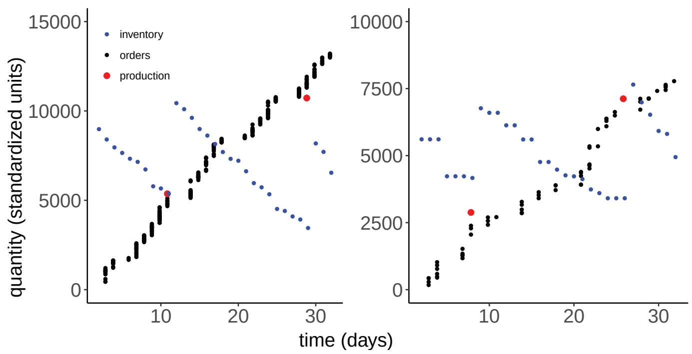

Figure 6: Simulated cumulative order volume (black circles), cumulative production (red circles), and daily inventory (blue circles) for two products over time.

run, the time between runs, the speed of each line, minimum level of acceptable safety stock in inventory, and parameters controlling the logic behind which product to produce next. Each time a line becomes available, the Production Planner chooses a product and desired amount to produce. It then calls the Production Simulator to retrieve a stochastic realization of the time the line runs.

The logic parameters weight different attributes of each product based on both simulated data and historical empirical data, allowing for a trade-off in prioritization as the simulation runs based on multiple factors. For instance, the demand weight w_d is larger for products with unfulfilled demand, the profitability weight w_p is larger for products with higher historical profitability, and the rarity weight w_r is larger for products that are rarely ordered. For a given setting of weights $w_i = (w_d, w_p, w_r)$, the Production Planner outputs the types and quantities of products produced on each line and in inventory.

2.2.4 Shipment Simulator

The Shipment Simulator is a discrete-event simulation that models the loading, shipment, and delivery of finished products to customers. There are two main inputs: (1) simulated orders, which is the output from the Orders Simulator, and (2) simulated inventory, which is an output from the Production Planner. The Shipment Simulator loops through each order and performs a series of steps to determine the truck for which to load the shipment. The main steps include (1) checking that there is available inventory to satisfy the quantity and product ordered, (2) checking for an available loading docks or space within a truck on the loading dock, (3) loading the order onto a truck, and (4) calculating the estimated delivery date and time of the order. A detailed process diagram of the Shipment Simulator is shown in Figure 7. In this initial model we make a set of simplifying assumptions, which are listed below:

- There is a maximum number of loading docks and a new truck cannot be loaded unless there is an available dock.
- Each order can fit in a truck (i.e., orders do not get split into multiple trucks).
- There is an infinite number of trucks. Thus, if there is an available dock, there is an available truck.
- There is a maximum number of pallets that can be loaded in a single truck.
- Shipments can have multiple legs and thus, trucks can ship multiple orders. An order will share the same truck if it meets the following criteria: (1) The customers in the orders must have shared a route in the historical data (2) the quantity of pallets of the orders must not exceed the maximum number of pallets that can fit in a truck.
- If there are multiple orders and multiple customers in a truck, information from existing routes is used to estimate travel time.
- Before leaving the loading dock, trucks that are not full will wait a fixed number of hours for additional orders.

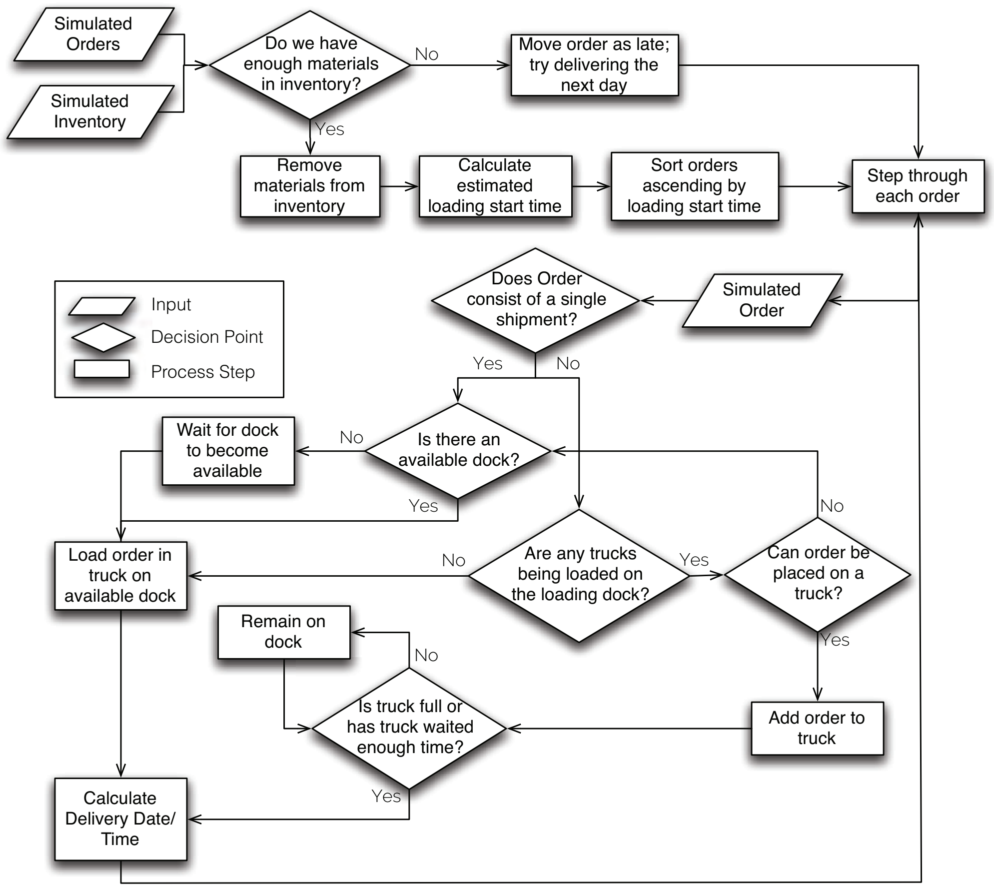

Figure 7: Process diagram of the Shipment Simulator.

The Shipment Simulator outputs one observation per shipment, including an order ID, shipment ID, truck ID, quantity shipped (in number of pallets), loading start time, loading end time, travel time, delivery date, and delivery time. Results were compared to actual observations to ensure the shipment simulator was running correctly and that the assumptions made were reasonable.

2.3 Calculating Costs and Profits

We estimated profits for both simulated and actual orders to validate simulation results (discussed in Section 3). To calculate the associated profit for each actual order in P&G's transactional data we developed an empirically-based simulation. In this simulation, we assume that orders are shipped using a First-In-First-Out (FIFO) process. For each order in the transactional data, we search for a production run in the data of the same product with a production date on or before the order's shipment date. If one more more production runs are found, we sort them in ascending order by production date. We loop through the production runs (starting with the oldest) until there is sufficient quantity to satisfy the order. These production runs are then linked to the order. We also track inventory flows by removing the quantity shipped from the inventory amounts that day. The detailed steps are shown in Figure 8.

We perform a similar process using results from our simulation model. Output from the Production Planner provides information on each production run, including the simulated production times and quantities produced on each assembly line. As with the empirically-based simulation, the simulated production is matched to incoming customer orders using a FIFO process.

Note that in our simulation model we are simulating a closed system with a single facility. Thus, every production run can be linked to an order and any orders that are unfulfilled when simulated production fails to meet demand are marked as such. We use this to calculate a simulated estimate of service rate, which is given by the percent volume of unfulfilled product. In reality, however, the supply chain is not a closed system with a single facility and we must also estimate production and shipment costs for orders coming from other facilities. This posed a challenge given that we had production runs, shipments, and excess inventory that we could not link to an order in our data. At this point, these associated costs were not substantial so we simply removed them from our analysis.

We now know the production information for each order (simulated and actual), including the quantity of raw materials used, the start and end times of production runs, and days spent in inventory for each unit ordered. We use this information to calculate each order's associated production, inventory, and overhead costs. Moreover, each simulated and actual order can be directly linked to its corresponding shipment information, which comes from the Shipment Simulator and the shipment data, respectively. This allows us to directly determine transportation costs for each order. We also have estimated revenue amounts for each simulated order from the Orders Simulator and actual revenue amounts from the orders data. Profits in both cases are then simply calculated by subtracting the total cost of the order from its revenue total. Simulated and actual profit, revenue, and cost results are shown in Figure 9(a).

3 A SIMULATION-BASED INVESTIGATION

At this point, we have informed statistical and discrete-event simulation models using transactional data to describe and simulate the supply chain system of a P&G facility. In this simulated system, we can efficiently and accurately compute key metrics such as costs, revenue, profit, and service (e.g., the proportion of on-time deliveries, proportion of filled orders, time from order to delivery). Similarly, we can compute these same metrics using empirical data. However, empirical data alone only informs us of the supply chain system as it is currently being managed, making it difficult to predict how the system will change under different production planning strategies.

We begin by first attempting to simulate the behavior of the supply chain system for a particular *test* month. We initialized the simulated system using two months of prior data in order to get inventory, production and shipment levels that were similar to the beginning of the test month's levels. This required that we calibrate the production planner to mimic the planning behavior of the facility for the test month. In order to do this, we encoded our knowledge about the production planning strategy used at the facility in a prior distribution. We then adjusted the parameters in the simulation model by matching key metrics from the prior months (e.g., profits, revenue, inventory, production costs, transportation costs) as shown in Figure 9(a). Due to computational demands of running the complete simulation model, which includes all four simulators described in the previous section, we incorporated ideas from Bayesian computer model calibration (Schonlau et al. 1998, Kennedy and O'Hagan 2001, Higdon et al. 2004).

Once the Production Planner is calibrated to reproduce the actual results of the test month, we explored the space of possible planner strategies that might produce better profit and service characteristics than are currently being realized. For each item delivered k, we compute $p_k \cdot I_k$ where p_k is the profit gained by delivering item k, and $I_k = 1$ if the item was delivered on or prior to its requested delivery date, $I_k = 0$ otherwise. We parameterized the Production Planner so that the relative weight of demand, profitability, and rarity could be adjusted, affecting how the Production Planner assigns production runs to lines given the customer orders (as discussed in Section 2.2.3). For a given setting of $w_i = (w_d, w_p, w_r)$, the simulation then produces an output y_i, which is the profit p_k for all items delivered I_k on time summed over the test month M. This is

$$y_i = \sum_{k \in M} p_k \cdot I_k.$$

We carried out a sequence of runs $w_1, \ldots, w_m$, drawn uniformly over the space of weights on demand, profitability, and rarity, and recorded $y_1, \ldots, y_m$ as shown in Figure 9(b). With these simulations, we fit a response surface using a Gaussian process emulator (Sacks et al. 1989) to the simulation output so we can seek out an optimal setting w^{opt} for the production planner (the $+$ in Figure 9(b)). Since this proposed optimal setting has not been used to simulate the supply chain process, we also run the complete simulation at this planning setting w^{opt} to produce a hypothetical result for the test month. Results show an increase in profit and on time delivery (service), while reducing inventory. Note that we cannot quantify these results due to the sensitivity of the data used in this study.

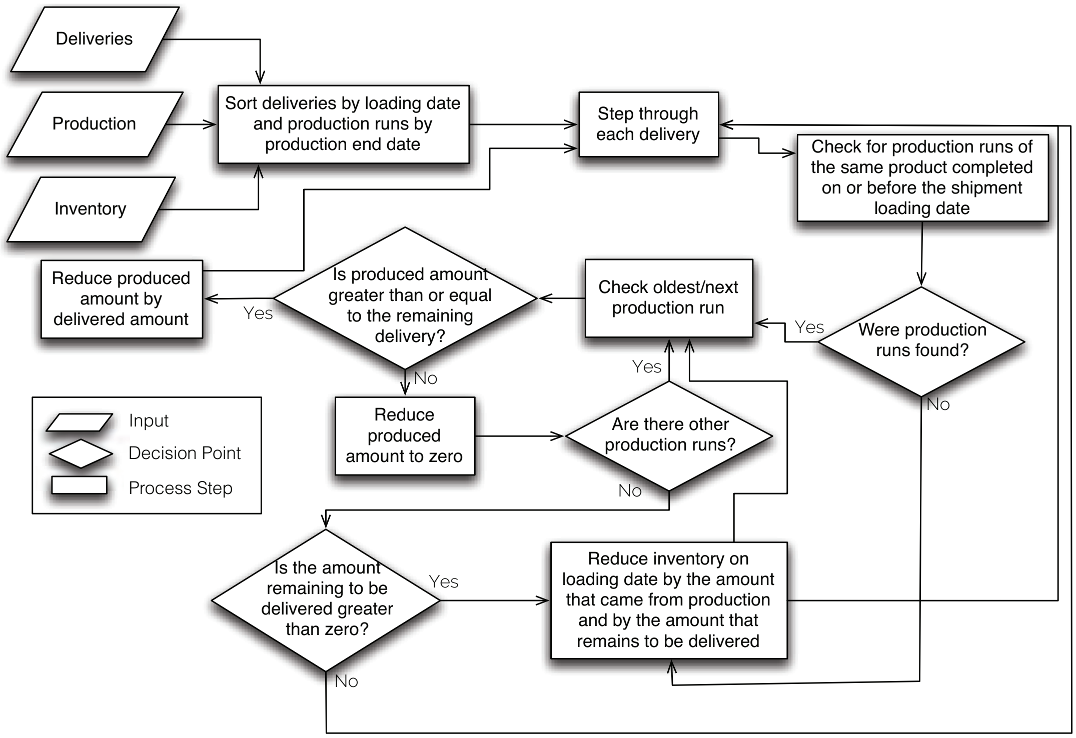

Figure 8: High-level process diagram of profits calculation.

4 CONCLUSION

An interconnected collection of disparate, big-data streams was used to inform the development of a series of statistical models and discrete-event simulations, which together captured supply chain processes from orders to deliveries. This allowed us to explore the behavior of the supply chain system within a P&G facility with its many dependencies and interactions. The above example focused on optimizing a family of production planning strategies. Such an approach would likely be beneficial in exploring a broader set of planning strategies. One could also explore the impact of altering the production capabilities of the production center, or the impact of incentives to alter customer ordering behavior.

From a broader perspective, this approach serves as one particular template for leveraging big data to gain understanding and insight about a complex system. This use of separate, but linked statistical and discrete-event models, allows the different sources of data to inform our knowledge about the supply chain system, and to predict its behavior under new conditions. This approach naturally captures variation in customer ordering behavior, production efficiency, and shipping times. It also captures uncertainty in statistical model parameters and forecasts. This results in simulation-based predictions that reflect these different sources of variation and uncertainty as well as their dependencies.

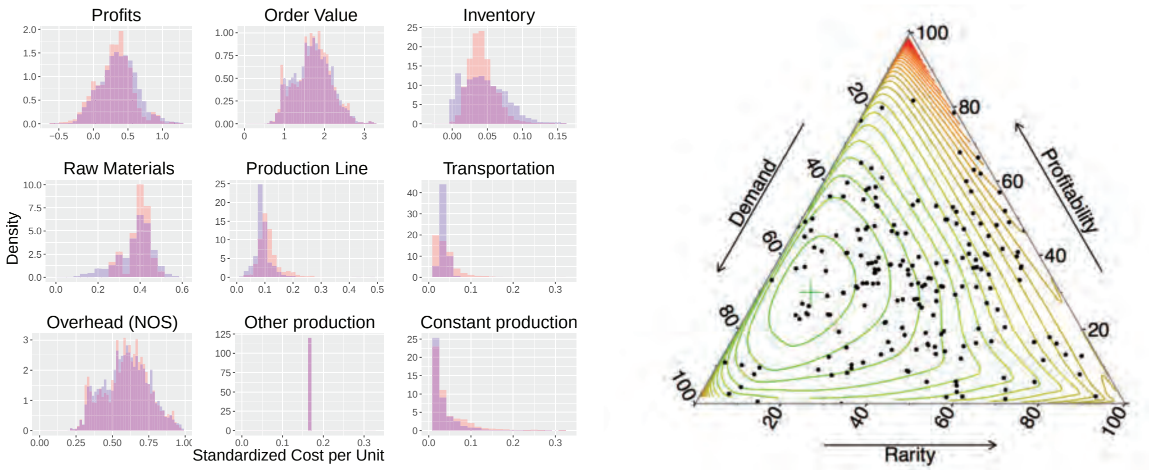

Figure 9: Left: Comparison of key outputs from the simulator (blue) and derived values from the test month (pink). Histograms show profits, order value, etc., for items delivered over the course of the test month. The overlap between the simulated and derived value distributions is shown in purple. Right: Simulated profit for all on-time deliveries over the course of the test month as parameters of the production planner are varied. A response surface was fit to the simulation output, suggesting an optimal planner setting (+).

REFERENCES

Celik, N., and Y.-J. Son. 2010. "State Estimation of a Supply Chain Using Improved Resampling Rules for Particle Filtering". In *Proceedings of the 2010 Winter Simulation Conference*, eds. B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Ycesan, 1998–2010. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Chick, S. E. 2004. "Bayesian Methods for Discrete Event Simulation". In *Proceedings of the 2004 Winter simulation Conference*, eds. R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 89–100. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Correa, M., C. Bielza, and J. Pamies-Teixeira. 2009. "Comparison of Bayesian Networks and Artificial Neural Networks for Quality Detection in a Machining Process". *Expert Systems with Applications* 36 (3): 7270–7279.

Dasu, T., and T. Johnson. 2003. *Exploratory Data Mining and Data Cleaning*, Volume 479. John Wiley & Sons.

Dawid, A., M. DeGroot, J. Mortera, R. Cooke, S. French, C. Genest, M. Schervish, D. Lindley, K. McConway, and R. Winkler. 1995. "Coherent Combination of Experts' Opinions". *Test* 4 (2): 263–313.

Gelman, A., J. Carlin, H. Stern, and D. Rubin. 2004. *Bayesian Data Analysis*. CRC press.

Harrison, J., and M. West. 1999. *Bayesian Forecasting & Dynamic Models*. Springer New York.

Higdon, D. 2002. "Space and Space-Time Modeling Using Process Convolutions". In *Quantitative Methods for Current Environmental Issues*, edited by C. Anderson, V. Barnett, P. C. Chatwin, and A. H. El-Shaarawi, 37–56. London: Springer Verlag.

Higdon, D., M. Kennedy, J. Cavendish, J. Cafeo, and R. D. Ryne. 2004. "Combining Field Observations and Simulations for Calibration and Prediction". *SIAM Journal of Scientific Computing* 26:448–466.

Jahangirian, M., T. Eldabi, A. Naseer, L. K. Stergioulas, and T. Young. 2010. "Simulation in Manufacturing and Business: A Review". *European Journal of Operational Research* 203:1–13.

Jin, S., Y. Liu, and Z. Lin. 2012. "A Bayesian Network Approach for Fixture Fault Diagnosis in Launch of the Assembly Process". *International Journal of Production Research* 50 (23): 6655–6666.

Kennedy, M., and A. O'Hagan. 2001. "Bayesian Calibration of Computer Models". *Journal of the Royal Statistical Society (Series B)* 68:425–464.

Li, M., and W. Q. Meeker. 2014. "Application of Bayesian Methods in Reliability Data Analyses". *Journal of Quality Technology* 46 (1): 1.

McNaught, K., and A. Chan. 2011. "Bayesian Networks in Manufacturing". *Journal of Manufacturing Technology Management* 22 (6): 734–747.

Reese, C. S., A. G. Wilson, M. Hamada, H. F. Martz, and K. J. Ryan. 2004. "Integrated Analysis of Computer and Physical Experiments". *Technometrics* 46 (2): 153–164.

Robert, C. P., and G. Casella. 1999. *Monte Carlo Statistical Methods*. Springer-Verlag Inc.

Sacks, J., W. J. Welch, T. J. Mitchell, and H. P. Wynn. 1989. "Design and Analysis of Computer Experiments". *Statistical science*:409–423.

Schonlau, M., W. J. Welch, and D. R. Jones. 1998. "Global Versus Local Search in Constrained Optimization of Computer Models". *Lecture Notes-Monograph Series*:11–25.

Soberanis, I. E. D. 2010. *An Extended Bayesian Network Approach for Analyzing Supply Chain Disruptions*. Ph. D. thesis, University of Iowa.

Xu, D., and Y.-J. Son. 2013. "An Integrated Simulation, Markov Decision Processes and Game Theoretic Framework for Analysis of Supply Chain Competitions". In *Proceedings of the 2013 Winter Simulation Conference*, eds. R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 3930–3931. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

BIANICA PIRES is a Research Scientist at the Social and Decision Analytics Laboratory (SDAL) within the Biocomplexity Institute of Virginia Tech (VT). Dr. Pires holds a Ph.D. in computational social science from George Mason University. Her email address is bpires@vt.edu.

JOSHUA GOLDSTEIN is a Postdoctoral Associate at SDAL. Dr. Goldstein is a statistician and received his Ph.D. at Penn State. His email address is joshg22@vt.edu.

DAVE HIGDON is a Professor at SDAL. Dr. Higdon holds a Ph.D. in statistics from University of Washington. His email address is dhigdon@vt.edu.

SHANE REESE is a Professor of Statistics at Brigham Young University. Dr. Reese holds a Ph.D. in statistics from Texas A&M University. His email address is reese@stat.byu.edu.

PAUL SABIN is an Analytics Specialist for ESPN and a Ph.D. student in statistics at SDAL. Mr. Sabin received his M.S. and B.S. in statistics from Brigham Young University. His email address is psabin88@vt.edu.

GIZEM KORKMAZ is a Research Assistant Professor at SDAL. Dr. Korkmaz holds a Ph.D. in Economics from the European University Institute. Her email address is gkorkmaz@vt.edu.

SHAN BA is a Senior Statistician at the Procter & Gamble (P&G) Company. Dr. Ba received his Ph.D. in Industrial Engineering (Statistics) from Georgia Tech. His email address is ba.s@pg.com.

KEN HAMALL is a Corporate Modeling and Simulation Discipline Leader at P&G. Mr. Hamall received his M.S. in Mechanical Engineering from Georgia Tech. His email address is hamall.km@pg.com.

ART KOEHLER, Associate Director of Baby, Fem and Family Sector, Modeling & Simulation at P&G, has thirty eight years of industry experience in product technical engineering, quality engineering, reliability engineering, and development of computer aided engineering methods and managing their applications for consumer product design and manufacturing process applications. His email address is koehler.aj@pg.com.

SALLIE KELLER is a Professor of Statistics and the Director of SDAL. Dr. Keller's areas of expertise are social and decision informatics, statistical underpinnings of data science, uncertainty quantification, and data access and confidentiality. Her email address is sallie41@vt.edu.

STEPHANIE SHIPP is the Deputy Director of SDAL. Dr. Shipp's research spans topics the use of data to inform policy, metropolitan analytics, and measurement of innovation. Her email address is steph19@vt.edu.