Monte Carlo Simulation Model for the Economic Impact of an Inventory-Dependent Business using s,S Inventory Policy

Ezequiel L. Castaño, Gastón Amengual

Universidad Tecnológica Nacional Facultad Regional Rosario

ecastano@frro.utn.edu.ar, gamengual@frro.utn.edu.ar

Abstract

Risks are present in all companies and are an essential factor for investors, and not including them in the decision-making process can lead to inefficient decisions. With the latest computational tools, simulations can be run to forecast possible outcomes and yield new insights. One of the most important criteria for profit-making in stock dependent businesses is proper and timely inventory management. This work introduces a discrete-event simulation model capable of integrating both logistic elements, such as inventory control, and also economic variables, such as capital flow. A comparative analysis of how different configurations of the s,S Inventory Policy can lead to contrasting economic scenarios is presented along with relevant statistical summaries of the net profit after simulating six months.

Keywords

s,S Policy; Inventory System; Discrete Event Simulation; Monte Carlo Methods

Introduction

Risk assessment is a key element to evaluate any investment. Entrepreneurs and managers must decide what to do with the available data, which is usually incomplete and noisy, and the unavoidable uncertainty, produced by the lack of information or otherwise, leads often to prejudicial decisions.

The advances in computational tools and electronic components made simulations once considered implausible, possible, and feasible. Running thousands of simulations and performing statistical analysis with the results allows more informed and timely decision-making.

When simulations are combined with modern optimization techniques, it is possible to obtain information regarding which the optimum initial conditions would be (*e.g.* initial inventory level, minimum inventory, and inventory policies parameters), so that key performance indicators (KPIs) are maximized (*e.g.* capital level after 6 months).

In the context of business risk assessment, it is now reasonable to run simulations that can include multiple perspectives at the same time, incorporating those essential for most businesses, such as logistic and economic.

In this paper, a simulation model is presented using a Monte Carlo approach. The simulation integrates both logistic and economic perspectives, as it involves inventory policy parameters, particularly, the s,S policy [1], and analyzes the impact on the evolution of the inventory level and capital flow over time. It also introduces behavioral economic elements, such as varying the arrival and sale rates based on whether the customer expectations (price and inventory availability) are satisfied. The KPI of interest is the capital level at the end of the simulation.

The triggering questions that motivate this study are: (1) How do different s,S values affect the capital flow? (2) Is there an underlying pattern on the s,S values that leads to different economic scenarios? (3) How much time would it take to recover the initial investment? (4) Is the expected capital flow heteroscedastic? (5) Does the simulation ever enter a steady state?

The next section describes previous works done in this direction, then the methodology and model description are explained, the following section details the experiments performed, and finally, the results are summarized along with the conclusions.

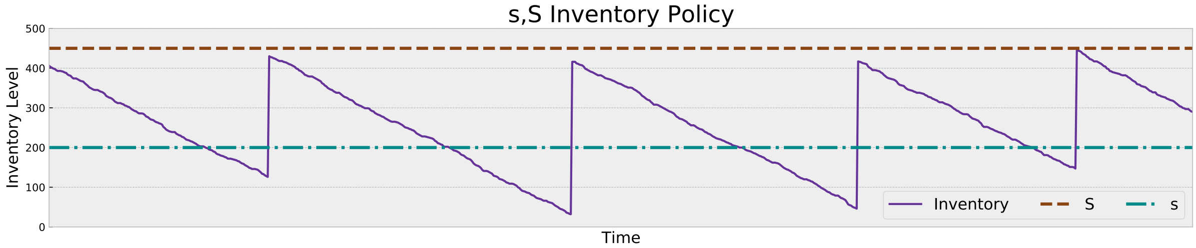

Figure 1: S,s Policy example with $s = 200$ and $S = 450$

State of the Art

The use of the s,S policy was proposed in 1950 by [1]. Other works in the inventory control field, like [2], compare only a few runs of a simulation, in this case using (Q,R) policy, instead of focusing on other KPIs such as the ones described in the previous section.

Simulations involving Monte Carlo Methods (MC) take advantage of modern computational resources and repeat a given experiment hundreds or thousands of times, avoiding the huge uncertainty present when doing inference with small sample size. As a system grows more complex, analytical and numerical approaches are extremely difficult to solve, and a simulation (either stochastic or deterministic) can be useful to replicate those systems, which generally include many variables interacting with each other.

There are three well-known simulation paradigms. First, System Dynamics [3] uses differential equations to represent populations, and relationships between the variables to reflect interactions in the actual phenomena. Second, Agent-Based Simulation [4], focuses on the elements rather than on the interactions. Finally, Discrete Event Simulation [5] relies on modeling each interaction in the actual system as an event occurring on a specific time, and a global clock which determines which the next event is.

In this paper, a Discrete Event Simulation is used, as it is the paradigm that best represents the system [6]. In the following section, details about the implementation and the assumptions are described.

Methodology

In this section, several elements needed to replicate the study are presented. First, the general assumptions, then the input and output variables, and finally how the simulation model combines the mentioned variables.

Policy Description

The s,S policy is an inventory policy that states that when the present inventory level falls below a certain threshold s, an order must be scheduled, the size of the order equals a given inventory level S minus the current inventory level. In Figure 1, an example of this policy is presented.

Mathematically, let $I(t)$ be the inventory level at a time t, then the amount Z to order is calculated as follows:

$$Z = \begin{cases} S - I(t) & \text{if } I(t) < s \\ 0 & \text{otherwise} \end{cases}$$

Assumptions

To reduce the complexity to a manageable level, the following assumptions were made:

Economic Assumptions

- The economic landscape is stable (negligible inflation).
- The niche market is kept constant (no major changes in the demand).
- The firm sells only one product.
- The customer will always want to buy if the price is less or equal than their expectations.

Logistic Assumptions

- The product has a single supplier.
- All necessary products are ordered and delivered at once.

- It is not possible to cancel or modify an existing order.
- The payment of the order is done at the moment of its scheduling.
- The inventory level is never negative.
- When an order arrives, if the amount ordered is greater than the maximum inventory, the exceeding inventory is lost.
- The stock only contain non-perishable products

Behavioural Economics Assumptions

- If the customer expectation is satisfied, the arrival rate of customers increases, if not, it decreases.

Variables

In the present work, the following convention is used: the input quantities of the simulation are called *Parameters*, the results of the simulation are called *Output Variables*, and the values the Parameters take at the beginning of the simulation are called *Initial Conditions* (Figure 2). The distinction between Parameters and Initial Conditions is made because some of the parameters are modified by a feedback loop during the simulation. To distinguish the parameters, those with feedback loop are expressed with a time subindex (*e.g.* λ_t) and those without it don't have a time subindex (*e.g.* Z).

It is expected that the *Output Variables* exhibit random behavior due to the stochastic elements in the simulation. Furthermore, each of the *Output Variables* are assumed i.i.d (Independent and Identically Distributed) for a given set of initial conditions and a given time t. Therefore, thanks to the Law of Large Numbers, Confidence Intervals can be generated for the mean (or median) of each Output Variable.
For this simulation, if a steady state for an *Output Variable* is reached, it is expected that for a time t_i (when that state begins) onwards the variable will be i.i.d. $\forall\, t_j \,/\, t_j \geq t_i$.

Parameters

The parameters of the model, along with their description, are detailed below, grouped in their respective domain.

Customer arrival

Each customer arrives according to a Poisson Process. The distribution of the number of daily arrivals is not constant throughout the simulation, instead, the *Average Daily Arrivals* is adjusted based on whether the demand was satisfied or not (following rules explained in the previous section). The change of this parameter follows a logistic growth with the minimum and maximum values given as additional parameters.

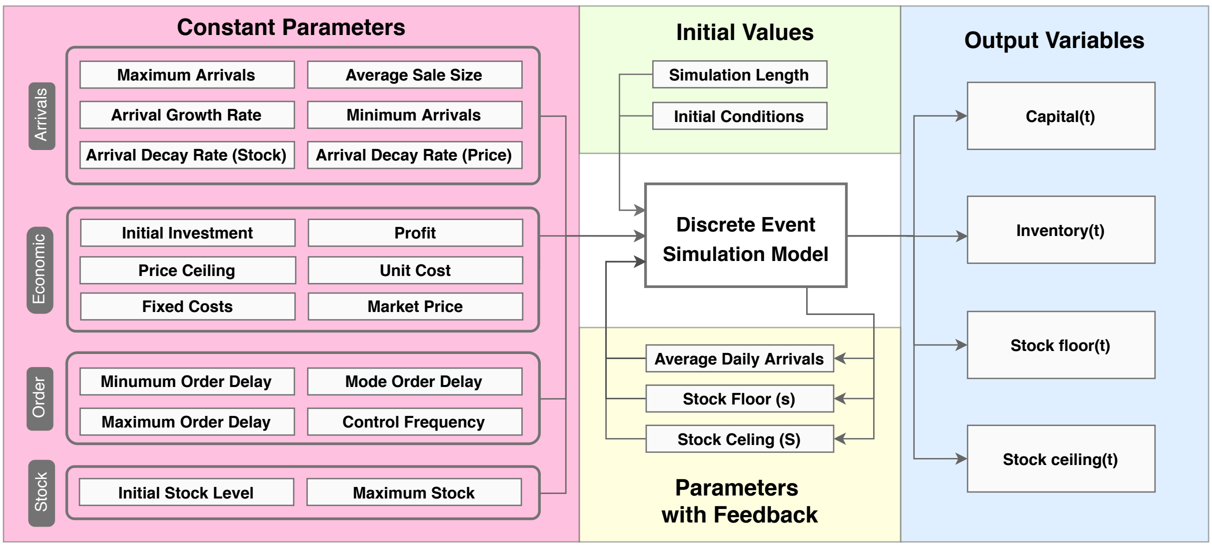

Figure 2: Diagram of the different components for the simulation. Including Parameters with and without Feedback loop, Initial Values and Output Variables

$$Daily\ Arrivals \sim Poisson(\lambda_t)$$

where λ_t is the *Average Daily Arrivals* for a time t. The adjustment for λ_t is done as follows:

$$change(P,r,K) = P\frac{r(K-P)}{K}$$

$$\lambda_{t+1} = \begin{cases} max(m,\ \lambda_t + change(\lambda_t, G, M)) \text{ If A} \\ max(m,\ \lambda_t - change(\lambda_t, D_p, M)) \text{ If B} \\ max(m,\ \lambda_t - change(\lambda_t, D_i, M)) \text{ If C} \end{cases}$$

Where A means that the customer expectations were satisfied, B that they were not due to price and C were not due to inventory. Moreover, G is the logistic growth rate (*Arrival Growth Rate*), D_p is the decay rate due to price (*Arrival Decay Rate (Price)*), D_i is the decay rate due to missing inventory (*Arrival Decay Rate (Inventory)*), M is the theoretical maximum for number of daily arrivals (*Maximum Arrivals*) and m is the theoretical minimum for the number of daily arrivals (*Minimum Arrivals*). In this implementation, the function is a truncated logistic function.

Each customer purchases, on average, a certain amount of products *Average Sale Size*, which is assumed to follow a Poisson distribution. Namely:

$$Average\ Sale\ Size \sim Poisson(A)$$

where A is the *Average Sale Size*.

Business and economic

At the beginning of the simulation, an *Initial Investment* is made, which constitutes the initial capital level to be updated with each sell and each order.

The *Profit* is the financial gain expected over the total costs. The total costs are determined by the cost of each unit *Unit Cost*, the fixed costs, the latter are the overall costs of the different services, taxes and employees *Fixed Costs*. For simplicity, the fixed costs are proportionally distributed during the whole month to smoothen the capital flow curve.

The product price, *Price*, increases or decreases according to the fluctuations in the *Average Daily Arrivals*. Furthermore, it is assumed that there is a *Price Ceiling* at which the product can be sold and is thus the maximum selling price, and a minimum price, calculated so that the firm breaks even, *i.e.*, it covers the total costs.

To calculate the product Price, it is necessary to estimate the total cost. The process to calculate the Price is described in the following steps:

1. Estimate the number of units sold in a month, U_{month}.
   $U_{month} = \lambda_t \cdot A \cdot \frac{365}{12}$
2. Estimate the monthly cost of those units C_{month}.
   $C_{month} = U_{month} * C_{unit} + Fixed$
3. Estimate the minimum price to break even Min_{price}.
   $Min_{price} = C_{month}/U_{month}$
4. Calculate Price *Price*.
   $Price = C_{month} * (1 + Profit)/U_{month}$
5. Truncate to limits if necessary.
   $Price = min(max(Price, Min_{price}), Price\ Ceiling)$

It is assumed that the customer price expectation is at least the *Market Price* with a positive percentual variation of up to 20%, mathematically:

$$\text{P} \sim MP * (1 + \text{Triag}(0, 0.1, 0.2))$$

where P is the customer price expectations, MP is the *Market Price* and Triag is the triangular distribution with parameters $(min, mode, max)$.

In summary, the price is calculated as follows:

$$\text{Price} = min(PC,\ C_{month} * (1 + profit))$$

where PC is the *Price Ceiling*.

Inventory and Orders

At the beginning of the simulation, the *Initial Inventory Level* is set and is updated during the simulation with each sell and each order. Also, the initial *Stock Floor* (s) and *Stock Ceiling* (S) are established, and are both modified during the simulation. The inventory is checked with a given *Control Frequency*, where the inventory is compared to

the *Stock Floor*, and, in case it is lower than it, a new order is scheduled according to the s,S Policy. In case there is not enough capital available to pay the order, the order size is reduced to the maximum size that can be paid.
The *Maximum Inventory* establishes the maximum amount of products that can be hold as inventory. When the order arrives, if the total inventory level is greater than the *Maximum Inventory*, it is round down to this value. Namely:

$$I(t+1) = min(\text{Maximum Inventory}, I(t)$$

where $I(t)$ represents the inventory level for a given time t. This rounding is always done after the arrival of a new order.
The mechanism to replenish the inventory is to order new products and wait until the supplier delivers them. The order delay *OD*, the time that takes to the order to arrive, follows a triangular distribution as follows:

$$\text{OD} \sim \text{Triag}(OD_a, OD_b, OD_c)$$

where OD_a is the *Minimum Order Delay*, OD_b is the *Maximum Order Delay*, and OD_c is the *Mode Order Delay*.

Simulation Length
The *Simulation Length* is the number of consecutive days to run the simulation. Ideally, it should be at the end of a month to better resemble reality.

Output Variables

The Output Variables are functions of the time t that consider the specific value for t. All the Output Variables are listed below:

- $Capital(t)$: capital level at time t
- $Inventory(t)$: inventory level at time t
- $Stockfloor(t)$: stock floor at time t
- $Stockceiling(t)$: stock ceiling at time t

A desirable characteristic in the capital variable is heteroscedasticity, this implies that the variability of the $Capital(t)$ will widen as time progresses. Namely:

$$\sigma^2(Capital(t_i)) < \sigma^2(Capital(t_j)),\ t_i < t_j$$

This property should be present in the simulation, as it represents more truthfully the reality, because, as time progresses, the precision of the model should decrease, since the more further the point in time, the greater the uncertainty. For instance, it is reasonable to assume that the uncertainty for 6 months is much less than for 36 months.

Model

To combine all the aforementioned variables, a discrete-event paradigm is used, because it was proven useful in the past for this scenario [6]. This type of simulation has three main components, (1) the state of the system, where the value of all the variables is stored, (2) a series of events that are run every time they are scheduled and that change the state of the system, and, (3) a global clock which determines which is the next event to be executed and always advances to the most imminent of the future events, where both the state of the system and the knowledge of the times of future events are updated.

Events

There are 3 main events present in the simulation, (1) Control, (2) Customer Arrival, and (3) Order Arrival. Each is described in detail below:
Control: this event is scheduled periodically with a fixed frequency *Control Frequency* during the simulation. It checks whether the *Inventory(t)* is less than the *Stock Floor(t)*, and, if so, the order size is set to the difference of *Stock Ceiling(t)* and *Inventory level*. If the *Capital(t)* is greater or equal than the ordering cost (*Unit Cost* times order size), an Order event is scheduled according to the supplier delay explained in the previous section. If *Capital(t)* is not big enough, a new order amount is calculated so that its ordering cost is less or equal than the capital level. If the order amount is greater than 1, an Order event is scheduled. Finally, a new control event is scheduled following the given frequency period.
Customer Arrival: this event is scheduled using the distribution associated with the

Common Parameters

Average Daily Arrivals	3	Maximum Arrivals	500	Minimum Arrivals	1	Average Sale Size	2
Growth Rate	0.005	Decay Rate (Inventory)	0.1	Decay Rate (Price)	0.001	Initial Investment	5000
Profit	2	Unit Cost	3	Fixed Costs	5000	Minimum Order Delay	0.5
Mode Order Delay	1	Maximum Order Delay	6	Market Price	7	Price Ceiling	14
Maximum Inventory	3000						

Table 1: Common Economic Parameters

Daily Arrivals. When a customer arrives, the sale size is calculated and, if that amount of stock is available and the price fits the customer expectations, then the sale is performed which means (1) the capital level increases by the amount sold times the price and (2) the inventory level decreases by the amount sold. Finally, the *Price* and the *Average Daily Arrivals* are updated and a new arrival event is scheduled.

Order Arrival: when an order arrives, the *Inventory(t)* is increased by the order size, and then round to the *Maximum Stock* in case it is greater than it. The time of the next order is not specified, as it is assigned in the Control event.

Experiments

The main focus of this paper is to evaluate the economic impact of the inventory-related parameters, therefore, different experiments were performed to show the flexibility of the developed model. In particular, four relevant economic scenarios are considered, which the simulation model should be able to replicate, being *NP* the net profit (*i.e.* Total Income - Total Expenses):

1. Profit without Losses (PNL): the company presents profits without intermediate losses, *i.e.* *NP* is positive and the capital is never negative during the simulation.
2. Profit with Losses (PL): the company presents profits with intermediate losses, *i.e.* *NP* is positive and the capital could have been negative during the simulation.
3. Breaks Even (E): the firm breaks even, *i.e.* *NP* can be positive or negative but it is approximately zero.
4. Loss (L): the company presents losses, *i.e.* *NP* is negative.

To calculate *NP*, the *Initial Investment* is subtracted from the *Capital(t)* at the end of the simulation. The simulation length is set to 6 months.

The parameters are divided into two groups, (1) those who are kept fixed in all scenarios and (2) those who are scenario-dependent. The only parameters that differ in each scenario are those related to the inventory control, namely *Control Frequency*, *Initial Inventory Level*, *Stock Floor*, *Stock Ceiling*. In Table 1 all the economic (fixed) parameters are listed.

Scenario-Specific Parameters

To test whether the simulation model is capable of reproducing the mentioned scenarios, an optimization was done to get the values. The optimization method used was Bayesian Optimization through the Python Package Optuna [7], and the objective functions were:

1. PNL: maximize *NP* but return 0 if at any time the *Capital(t)* is less than or equal to 0.
2. PL: maximize *NP*.
3. E: minimize $abs(NP)$.
4. L: minimize *NP*.

The values of the parameters were the best among 20 independent optimizations, of 1000 iterations each, each consisting of 5 replicates of the whole simulation. The values are shown in Table 2.

Results

For each scenario, 100 runs were performed. The model was implemented using the Python programming language and its scientific toolbox, including Numpy [8],

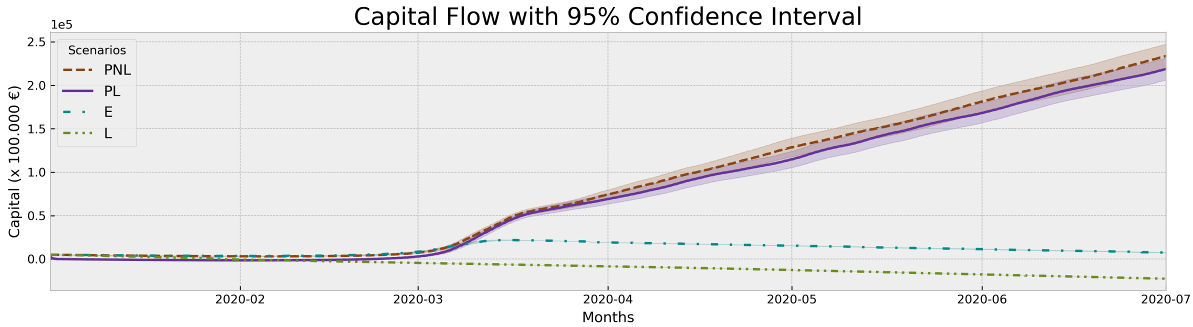

Figure 3: Confidence Interval Chart

Scenario Dependent Parameters

	PNL	PL	E	L
Control Frequency	1	1	22	27
Inventory Level	2754	441	2349	0
Stock Floor	2807	2827	3	0
Stock Ceiling	2799	2889	65	137

Table 2: Scenario-Specific Parameters

Pandas [9], and Matplotlib [10].
In Table 3, following the principle behind [11], the Confidence Interval (CI) ($\alpha = 0.05$) for mean of the Capital Level at the end of the simulation (6 months) is shown for each scenario. In Figure 3, a Confidence Interval Chart for the mean of the Capital Level in each scenario is shown. Furthermore, the change in width of the CI over time is on average positive for the whole time domain, indicating an always growing variance, as illustrated in Figure 4. This was also proved analytically by a rejected Student's t-test (mean$=0$, $\alpha = 0.05$) and a rejected Levene test ($\alpha = 0.05$). Also, all 100 runs for the PNL scenario can be seen in Figure 5.

Capital Confidence Interval (95%)

PNL	233774 ± 13869	PL	218853 ± 12658
E	7559 ± 337	L	-22557 ± 25

Table 3: Confidence Interval (95%) for the Capital Level after 6 Months

Discussion

The work uses economic parameters that were arbitrarily chosen, however, a fit to real data and a posterior calibration is needed before implementing the proposed model in a real-world scenario. Moreover, some of the assumptions may be too restrictive, such as the complete avoidance of inflation or the fact that the market price is fixed.

Conclusions

First, clear and distinct patterns emerged for each scenario: (1) the lower the control fre-

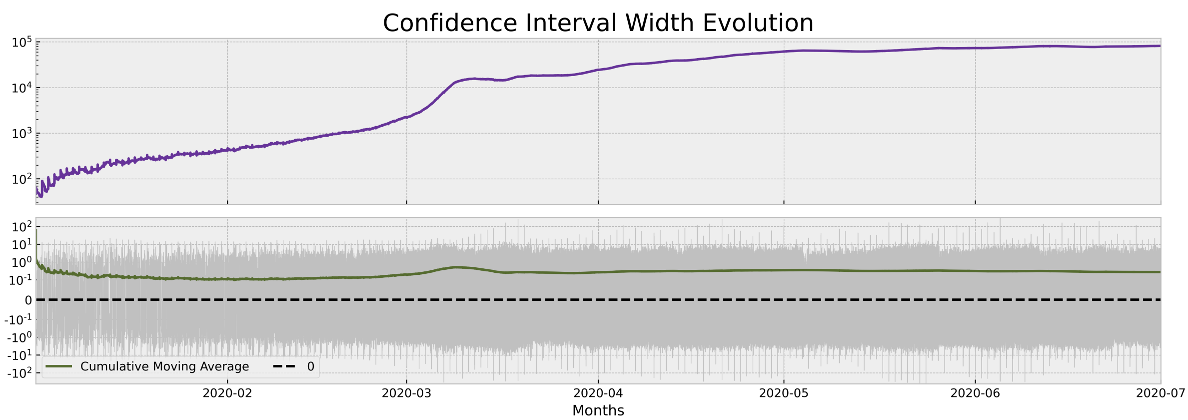

Figure 4: Width of Capital Flow CI. Width over Time (above), width instant change with cumulative moving average (below)

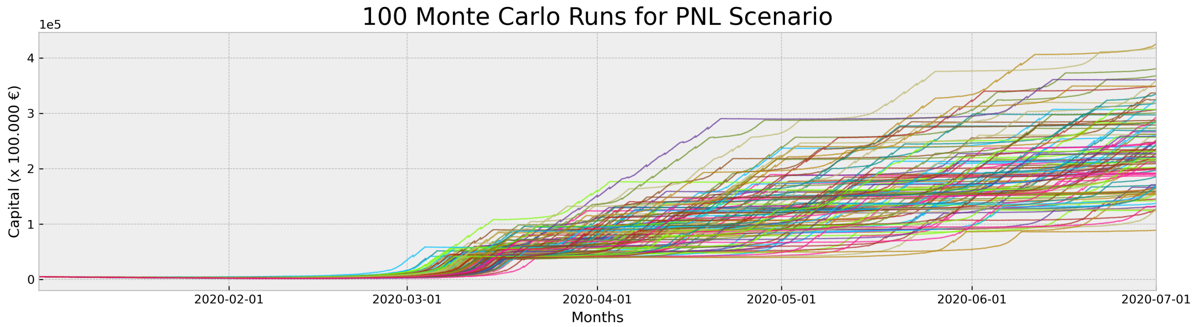

Figure 5: Capital Flow - 100 Simulations

quency, the higher the profit; (2) to maximize the profit, the initial stock should be as high as possible, although low initial stock does not imply breaking even or losing; (3) in both profitable scenarios, stock floor, and stock ceiling were close to the maximum stock possible (greater than the 80th percentile of possible values); (4) the most critical factor to break even or losing is to have either a low stock ceiling or a high control frequency.

Second, the return of the investment in both profitable scenarios was between the third and fourth month. Additionally, sustaining losses throughout the simulation did not generate higher profits than applying a no-loss policy.

Furthermore, as explained in the Results Section, the variance increases continuously over time, implying heteroscedasticity and therefore reflecting the basic intuition that long-term forecasts are less precise than short-term predictions.

Finally, in none of the scenarios tested the simulation reaches a steady state. If desired, a more complex control system should be incorporated.

References

[1] K. J. Arrow, T. Harris, and J. Marschak, "Optimal inventory policy," *Econometrica: Journal of the Econometric Society*, pp. 250–272, 1951.

[2] G. Bonilla-Enriquez and S.-O. Caballero-Morales, "Simulation model for assessment of non deterministic inventory control techniques," *Asian Journal of Research in Computer Science*, pp. 63–70, 2020.

[3] J. W. Forrester, "System dynamics, systems thinking, and soft or," *System dynamics review*, vol. 10, no. 2-3, pp. 245–256, 1994.

[4] E. Bonabeau, "Agent-based modeling: Methods and techniques for simulating human systems," *Proceedings of the national academy of sciences*, vol. 99, no. suppl 3, pp. 7280–7287, 2002.

[5] G. S. Fishman, "Principles of discrete event simulation," 1978.

[6] A. M. Law, W. D. Kelton, and W. D. Kelton, *Simulation modeling and analysis*, vol. 3. McGraw-Hill New York, 2000.

[7] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A next-generation hyperparameter optimization framework," in *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.

[8] S. van der Walt, S. C. Colbert, and G. Varoquaux, "The numpy array: A structure for efficient numerical computation," *Computing in Science Engineering*, vol. 13, no. 2, pp. 22–30, 2011.

[9] Wes McKinney, "Data Structures for Statistical Computing in Python," in *Proceedings of the 9th Python in Science Conference* (Stéfan van der Walt and Jarrod Millman, eds.), pp. 56 – 61, 2010.

[10] J. D. Hunter, "Matplotlib: A 2d graphics environment," *Computing in Science Engineering*, vol. 9, no. 3, pp. 90–95, 2007.

[11] M. Lin, H. C. Lucas Jr, and G. Shmueli, "Research commentary—too big to fail: large samples and the p-value problem," *Information Systems Research*, vol. 24, no. 4, pp. 906–917, 2013.