

The principal-agent problem in supply chain management – the simulation based framework*

by

Kinga Siuta and Daniel Kaszyński

SGH Warsaw School of Economics, Warsaw, Poland

ksiuta@sgh.waw.pl, dkaszy@sgh.waw.pl

Abstract: This research deals with a phenomenon well-known in socio-economic studies and referred to as the Agency Theory: the principal-agent problem. The agent is designated to act on behalf of the company owner, i.e., the principal, in the domain of supply chain management in the face of a supplier's default uncertainty. Each of the players is gain-oriented, but their objective functions and the degree of risk-bearing may significantly differ, leading to an agent's decisions not being optimal in terms of the principal's outcome. This article proposes a multi-period supply chain simulation model that can be used to perform a supply chain optimization and comparison between the agent and the principal. The proposed framework allows for differentiating the model parametrization depending on the industry, in particular the consequences of the inability to deliver the end product, costs of keeping safety stocks, or the uncertainty regarding the suppliers' delivery failure. As players' objective functions, we consider expected profit-based indicators and measures taking the inter-period variance into account. We find that the agent's actions may diverge from the principal's optimum if the agent's incentive system is not selected correctly. We also propose a solution that unifies the goals of the players. The recent COVID-19 pandemic amplifies the importance of such research. Many companies had to limit production capacities due to global lockdowns and, per the JIT production strategy, the prior safety stock levels were low.

Keywords: supply chain, quantitative models for decision-making, simulation framework, agency problem, principal-agent problem, ESOP

1. Introduction and motivation

In this paper, we investigate the well-known principal-agent problem (also referred to as the agency problem) where the optimal, in terms of the individual

*Submitted: December 2020; Accepted: March 2021.

objective function, agent's action may not align with the maximum value of the principal's objective function. The principal is a company owner, who delegates running the firm's activities – supply chain management – to the agent (or a manager). This framework is also common in the literature – usually the principal acts more statically by delegating the agent's decision-making process (see Jensen and Meckling, 1976, or Ross, 1963). Instead, it is the agent making all of the decisions and managing the supply chain in an uncertain environment. Namely, the company's suppliers may not deliver any products in some periods, which is the leading risk this work deals with. In this paper we will be focused on investigating two problems:

- 1) the difference in optimal strategies between the principal and the agent under different objective (i.e., profit) functions,
- 2) selection of a contract with the agent that will guarantee a satisfactory profit for the principal.

The frequently used approach to supply chain modeling generally assumes that 1) there is one particular objective function, 2) the objective function is cost-oriented, 3) the model is deterministic (Pourhejazy and Kwon, 2016). The framework presented in this article is, in contrast, focused on finding a profit-oriented objective function for a manager that would align with one of the principals, especially in uncertain (non-deterministic) conditions. The reason for introducing such an approach is that we believe that a principal-agent problem in supply chain management can be observed in reality, where disturbances may occur. It is related to the differences as to the incentives, e.g., 1) the agent does not bear the risk of failure of the enterprise to the same degree the principal does, or 2) there are potentially different time horizons, over which the objective functions are evaluated for both actors. Therefore, this work is our attempt to propose a model for supporting supply chain decision-making in uncertain conditions in the face of an agency problem.

Disturbances in our analysis will stem from a sudden inability to obtain ordered products from the suppliers, a problem experienced by many enterprises due to the COVID-19 pandemic. The agent will try to optimize their profit, while the principal is also concerned with profit volatility and protecting the company from high risks. The agent can control two factors: 1) safety stock levels, 2) the number of suppliers, while the principal constructs agent's contract. We believe that in a non-deterministic supply chain, an apparent principal-agent problem can be observed. This is why, in the analysis, we will firstly confirm whether the agency problem is, in fact, visible in such a system. Furthermore, we will investigate under what objective functions an agent would not perform the way the principal would anticipate and propose actions of the principal that can mitigate the potential gap between the two parties. We expect that in order to achieve a sufficient similarity of the actors' goals, introducing disturbance related risk-sharing is required.

The paper is structured as follows. In the next section, Section 2, we describe problems related to the proposed analysis. We formally define our simulation model and present its parametrization in Section 3. Section 4 is devoted mainly to the results we managed to obtain using the model, while in Section 5, we summarize the findings.

2. Related problems

The following problems are directly related to this paper: 1) group decision-making processes, 2) the agency problem, and 3) the supply chain management. Here, we will briefly discuss them in this exact order before describing the framework developed based on their combination.

Our analysis is rooted in group decision-making (GDM), i.e., a process, through which individuals collectively choose from different alternatives (Kersten, 1985). The main issue GDM deals with is how to integrate individual preferences into a group preference (Li, Kou and Peng, 2016). Agents can have wildly diverging opinions, which is why a solution with the maximum degree of agreement among decision-makers is searched for in GDM problems (Kersten, 1997). This article deals with a similar issue, the difference being that there is just one decision-maker that acts on someone else's behalf. Thus, the situation also requires some sort of a consensus between the parties as the principal employs the agent.

The problem we are referring to may be classified as coming from an even narrower group of Economic Theory of Agency. As formulated by Stephen A. Ross (1963), an agency relationship arises between multiple parties when there is one designated agent who acts for (on behalf of) or as a representative of the other(s), the principal, in a particular domain of decision problems, and this is the relationship we describe in this work. A scheme of this arrangement is presented in Fig. 1.

In practice, there are plenty of scenarios, where the agency problem arises between two parties. The most notable examples are 1) lender and borrower, 2) shareholder and corporate managers, 3) citizens and elected officials (Janda, 2006). In this paper, the agent will be responsible only for a part of the company's activities, namely supply chain management. Just as in Grossman and Hart (1983), the owner is assumed not to be able to monitor the manager's actions and can only see the financial outcome of their work.

Although different objective functions guide the principal, who owns the production resources, the agent employed to manage them is oriented at his individual gains, this leading, potentially, to a non-optimal solution in terms of the principal's objective function. Notably, very often it is not just the objective function that differentiates the manager and the firm owner. There are many cases in the literature where information asymmetry occurs, i.e., the agents have the knowledge the principal has no access to (Myerson, 1982). This can




Figure 1. The relation between principal and agent
Source: own elaboration

result from many years of manager's experience in this particular role and their in-depth knowledge of the industry. In contrast, the principal lacks expertise in how other entities are run and how they operate behind the scenes. Our primary focus in this aspect, just as in the case of other principal-agent problem-related publications, is the contract as the unit of analysis (Kuwornu, Kuiper and Pennings, 2009). The principal values the manager's performance and pays the compensation as agreed in their contract. Depending on this arrangement, the agent's work results can be closer or further from the firm owner's goals. One may even encounter a type of contract, according to which the agent would make hazardous decisions counting on big profits and not bearing any costs in case of a failure. According to this discussion, the agent contract, which reduces the agency problem should address the following issues:

1. Information asymmetry regarding the market (e.g., insolvency probability of suppliers and the agent's efforts).
2. Goal conflict resulting from various time-horizons of the principal and the agent.
3. Risk-adjusted performance measure allowing for risk-sharing between the principal and the agent.

Having this relation between parties, one may also expect various time-horizons related to the objective functions. Namely, the principal is more broadly and deeply associated with the company, and the agent, who acts as a manager, is much shallowly embedded into the organization. The relations that the parties have with the company should also be captured by the model, which is why

the objective function of the manager can, in particular, be related to a shorter time-horizon (i.e., the length of the contract), whereas the principal's objective function (e.g., profit) is coupled with the time horizon of company's lifetime. The discussion, presented in the paper, regarding the time-horizon of the decision making (i.e., whether the short-term outcomes should be favoured over the long-term results) is currently present in the scientific literature (Barton, Manyika and Williamson, 2017).

In terms of the agency theory, our simulation-based framework will help to assess the gap between the two objectives and identify a type of contract between principal and agent that can narrow this gap and reduce agency costs. Moreover, our simulation concerns supply chain management under uncertainty and the related risk factors (e.g., penalty imposed on the company, in case of not covering the demand). Tang (2006) pointed out that reducing the effects of operational and disruption-related risks, especially in the area of participation in the supply chain, can be achieved by improving the efficiency of operations in four areas:

1. Supply management.
2. Demand management.
3. Product management.
4. Information management.

This article is focused on the first aspect, i.e., managing supply. Works in this field concern, among other things, the supplier selection criteria (e.g., De Boer, Labro and Morlacchi, 2001), the durability of relationships with contractors (e.g., Tang 1999), the division of orders among suppliers (e.g., Minner, 2003), and coordination between individual producers, allowing for the improvement of efficiency in the network as a whole (e.g., Tsay, Nahmias and Agrawal, 1999). In managing the supply chain, an enterprise can control the following factors:

1. The number of suppliers and distributors.
2. Speed and cost of acquiring new business partners.
3. Type of the delivered product and the value of flows between the company and contractors from point 1.
4. Timeliness of deliveries.
5. Timeliness of payments.

In terms of the techniques used in this paper, according to the article by C. Elock Son (2018), the dominant methods for modelling supply chains are, among others, quantitative methods with an optimization approach, multivariate analysis, stochastic programming, panel research (questionnaires) or simulation (used in this article).

In supply chain simulation literature, authors deal with mitigating risks and consequences of operating in an uncertain environment. As it was noticed in *Modelling and simulation of a supply chain in an uncertain environment* (Petrovic, Roy and Petrovic, 1998), external supplier uncertainty is a

serious threat to successful supply chain operations. The inventory stock levels (i.e. safety buffers) can be increased to compensate for such adverse effects. Even though the industry has been moving towards JIT (Just-in-Time) production for the past decades, recent COVID-19 pandemic disturbances showed that safety stocks should be kept in order to ensure that the enterprise is functioning smoothly when facing a disruption in supply (Hadwick, 2020). Media even announced the end of the JIT era approaching, saying that JIT turned out to be too lean to deal with the sudden surge regarding raw materials and increased demand for Consumer-Packaged Goods (CPG). *Financial Times* claimed that the pandemic has shown that many businesses neglected vital safety margins, leading to outcomes as severe as closing their companies in an article of a very self-explanatory title “Companies should shift from ‘just in time’ to ‘just in case’ ” (*Financial Times*, 2020).

Inventory related decisions are, therefore, the key for building a reliable supply chain. The introduced strategy has to serve two purposes (Tang, 2006). Firstly, it should allow the firm to reduce costs and/or improve consumer satisfaction under stable conditions. Secondly, the company should be able to sustain its normal operations during a period of disruption. Only a solution satisfying both requirements can be called robust in optimizing a supply chain under uncertain conditions. It has been established that sourcing the same material or product from multiple suppliers can improve the chain’s durability. For example, after the supply disruptions observed at Philips, Nokia reconfigured the design of their basic phones to accept also slightly different chips from Philips’s other plants and even other suppliers (Tang, 2006).

There is significant amount of simulation-based work on the principal-agent problem (Páez-Pérez and Sánchez-Silva, 2019), as well as similar approach to modeling supply chains (Persson and Olhager, 2002), but we are not aware of a simulation combining these areas into the problem we model in this paper. Also, while the problem itself has been tackled with quantitative methods before (Wang, Chen, Liu, Guo and Chen, 2018), there is not much research exploring the area of simulation as a solution of the agent-principal dilemma in supply chain management, and we hope to contribute to this field.

Given the above discussion, in this paper we focus on managing in the multi-period framework, the company’s inventory, i.e., impacted both by the covering demand and replenishment purchases. The literature discusses the agency problem broadly, which is also considered in the here presented model – the agent’s objective function is different from the principal’s. Moreover, as in the real-world conditions, the discussed model allows for modelling of the uncertainty regarding: a) demand, which is normally distributed and b) suppliers’ deliveries, which may be impacted by random events (further referred also as a default event).

3. Problem formulation

The model is a multi-period simulation mirroring a part of the company's operations. It includes selling products and replenishing safety stock to a level determined by the manager (the agent). The enterprise increases its margins by obtaining quantity discounts from suppliers and, on the other hand, bears the costs of keeping inventory. Moreover, a penalty must be paid each time the firm does not fulfil its product demand. Therefore, a company loses potential profit if: 1) no quantity discount is granted, 2) there is large inventory to support, and also 3) a penalty must be paid to customers/distributors. Notably, while demand is not constant, there are no significant disruptions in this area. The leading cause of uncertainty in the system is the possibility of a delivery failure on the side of (at least) one of the suppliers. The principal bears all the risks of the company's operations, while the agent, as an employee, does not.

The details of the model's components are described in the next paragraphs. The simulation will be answering the following questions: 1) what are the differences in optimal strategies for both of these players under various objective functions, and 2) what contract with the agent will guarantee satisfactory profit for the principal.

Most often, uncertainty in parameters in inventory control problems has been modelled by probability distributions (Petrovic, Roy and Petrovic, 1998) that could be estimated based on historical data, if available. This paper will formulate the general simulation framework based on the theoretical and expert-based assumptions related to supply chain modelling. We take an approach not calibrated to any actual data to show qualitative relationships in the presented problem. From the practical point of view, the simulation could be supplemented with parameters estimated on empirical data – that is one of the potential future advancements regarding the model.

In this work, by paraphrasing Craighead et al. (2007), we define supply chain disruptions as unplanned and unanticipated events disturbing the normal flow of goods within a supply chain and, as a result, expose the modelled company to operational and financial risks. We will not be looking into the chain as a whole but a local fragment of it – suppliers concentrated around one enterprise. Moreover, we will not be investigating the geographical nature of the processes even though we are aware that, especially in terms of COVID-19 pandemic and country-wide lockdowns, spatial diversification of suppliers is a topic worth exploring in further work. We will also not be looking into reasons for the disruptions and suppliers' defaults – whether they originated at our suppliers or were transmitted to them through their supply chain connections. Our interest is whether the supplier defaulted resulting in the consequential inability to deliver any products that the modelled company has ordered. Hence, the supply chain management should protect the firm against the risk of suppliers' default while keeping the profit satisfying.

In the previous section, the five factors an enterprise can control in its supply chain were listed. One can imagine a grand simulation taking all of the mentioned aspects into account. However, to keep the analysis transparent, we have to assume that some of these factors are constant and let the agents make decisions only in the area of:

1. The number of suppliers.
2. Safety stock levels.

In this study, the distributors and demand in each timestamp are randomly distributed (from the normal distribution with $\mu = 1000$ and $\sigma = 250$). The speed and cost of acquiring business partners are both equal to zero; here we recognize the literature related to transaction costs, for instance: Williamson and Masten (1999). The products delivered and bought are homogenous and so are the suppliers – each of them is capable of producing the same number of products if the default risk has not occurred (in the given timestamp). Payments, both up and down the chain, are timely. In the case of a supplier’s default event, the company cannot substitute and source the goods anywhere last minute. It can only use its inventory (or safety stock – synonymous terms in our work) and if it is not sufficient – a penalty for the pre-agreed quantities of the product has to be paid. Hence, at the stage of the actual purchase of the stock replenishment at iteration t , the final order size SR_t is equal:

$$SR_t = \sum_{k=1}^K \min\{SP_{kt}, D_{kt}\},$$

where D_{kt} is the number of items the modelled company would like to purchase from supplier k at iteration t , and SP_{kt} is the k -th supplier’s production at iteration t . The latter depends on whether the given supplier bankrupted or not. We do not allow for fluctuations of suppliers’ production – it is either big enough to satisfy the firm’s demand (in our case a large number of 100 000 was applied), or zero, i.e.

$$SP_{kt} = 100000 \cdot I \begin{cases} 0 & \text{if the supplier defaults} \\ 1 & \text{otherwise} \end{cases} .$$

Therefore, the firm encounters three types of costs: 1) cost of products, which account for the quantity discounts, i.e., the more company purchases from the supplier, the lower the potentially negotiated price, 2) storage costs, related to keeping safety stocks, and 3) penalty costs, imposed if demand exceeds company’s inventory. The agents are motivated by quantity discounts, obtained when making larger orders in terms of potential profits and savings, just like the lack of penalties and inventory-related costs.

The values of the costs taken into the simulation are not calibrated to empirical data and are strictly exemplary. However, they may be adjusted to the type of business, product, and business environment, the modelled company

operates in. Namely, the penalty connected with delivery failure varies, for instance depending on client retention in the given business. Suppose it is easy to find another customer for our product. In that case, our failure's consequences will be minimal or even zero but the situation is profoundly different if finding another customer or distributor takes a lot of time/effort and is difficult in general. Also, if it is easy to find another supplier for the product we deliver, we become replaceable. Hence, to retain this client after a delivery failure, it may be hard to negotiate beneficial conditions for future deliveries. Notably, frequent failures in delivering the pre-agreed amounts of the final product to distributors can have other negative consequences over time when at some point, our business partners lose their trust in our professionalism and ability to fulfil the contract terms. We do not account for this factor in the analysis; instead, we use a constant penalty for each undelivered unit. Still, our simulation allows for easy implementation of an extension of an increase in penalty or other forms of punishment.

Another cost driver, whose calibration can differ depending on the case is the product storage cost. Some products have to be stored in particular conditions (e.g., they require low temperatures or some sort of maintenance, i.e., labour or other resources), or their depreciation rate is large (e.g., there is a hidden cost of product depreciation connected with an inability to sell the product at a high price) – in such a case, total costs related to storing the product would be more considerable. In contrast, some products are very easy to store, e.g., they do not take up much space, hence their storage cost is close to zero. Storage cost, therefore, varies from industry to industry (Rajhans, Patil and Kulkarni, 2015). Very often companies even calculate the exact number, especially if they outsource the order fulfilment system, so the calibration in this area should not be problematic.

Although our model does not explicitly include transportation aspects of supply chain management, an aspect, analogous to storage costs can be present in the area of transportation, especially of perishable goods (Anholcer, Hinc and Kawa, 2019). In fact, transportation losses can be partly accounted for in our model by making appropriate adjustments to storage costs and making the suppliers heterogenous. Namely, depending on transportation costs from a given supplier, different prices would apply. Last but not least, prices of the supplied product are not constant. In reality, suppliers offer quantity discounts connected with negotiating a better price if the purchase is significant. In our model, the prices depend on the number of products delivered as well. The basic price is modified based on each supplier's level, so if we refill the stock levels with products of only one company, we can obtain larger discounts. In contrast, one gets little or no discount in the case of getting the same number of products but divided among 5 or 10 suppliers. We assume that the highest discount possible is 10% (which is also modifiable depending on business reality) and, as shown in Figs. 2 and 3, there are two possible discount schemes: 1) following a sigmoid function, 2) non-linear quantity discount thresholds. For convenience




Figure 2. The sigmoid quantity discount function
Source: own elaboration

and comparability reasons, in both functions maximum discount is achieved at the level of demand for our products in one period.

The mechanism of booking the supply costs based on the discounted prices is also worth mentioning. Namely, after buying supply, the products are mixed with the already accumulated stock, and the new average price is calculated. The price is an average weighted with the number of units bought at each price. We classify the purchase as a cost when selling it, i.e., not at the stage of the actual purchase at the supplier. From an accounting point of view, unsold items do not constitute a cost because, as they are part of the inventory, they are not included in calculating the financial result in a given period. The supply costs are each time accounted for at the price of the current inventory mix.

After a short description of the model components, we will present the model dynamics using a simplified diagram of Fig. 4*. In the here presented simulation, both types of agents can manipulate the same parameters, namely: 1) the number of suppliers, 2) inventory goal that will be achieved after each stock replenishment. The other values that appear in the simulation stem from factors such as market situation (e.g., number of sold products), or firm's sector and type of business (e.g., penalty for delivery failure). In the model, we first initi-

*The full model Python code can be found in a GitHub repository <https://github.com/kingasiuta/SimulFrameworkSC>




Figure 3. Non-linear quantity discount thresholds
Source: own elaboration

ate some values, e.g., stock levels are set to *InventoryGoal* equal to maximum stock levels the company will want to keep. We assume that no discount was granted when making this purchase, and the enterprise paid the full price of 0.8 monetary units to its suppliers, whereas the final sales price for the company is one monetary unit. After the state of the inventory and the amount paid for it is established, the actual simulation may begin. Below, we present one iteration of the model, i.e., one period, e.g., a day, week or month.

After a given number of periods (e.g., indicating one year or five years), total profit is calculated; however, the principal's and agent's profits are calculated differently. For the principal, all profit values are added – it does not matter whether the company noted any gain. It means that the principal bears the risk of adverse financial results, and in some periods, a loss is noted. In contrast, the agents cannot take such a threat as they are employees – they have their basic salary (which is constant, hence is negligible in this analysis and set to zero) and get bonuses depending on how big the profits are in this period. As the actual percentage of the premium does not matter if it is constant over time, we assume it is equal to one. Therefore, the only values that bear an impact on the agent's salary are the positive profits – the agent does not bear any cost connected with negative financial outcomes in some periods. It is possible that from the agent's standpoint it is more economically practical not to get a modest bonus every period but to get a larger payoff in between mostly non-

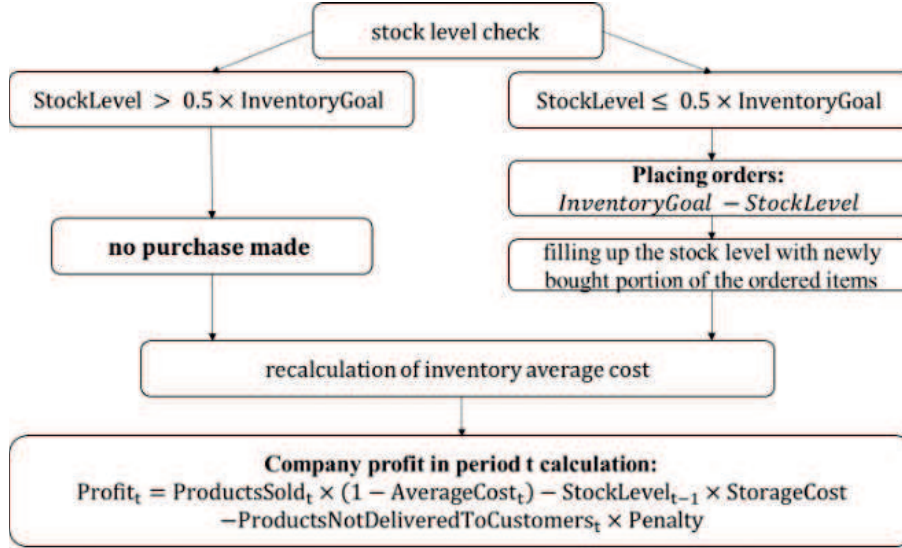


Figure 4. The scheme of oneiteration (one timestamp) of the simulation
Source: own elaboration

profitable period instead. Hence, the agent's salary variance is lower (because profit value is floored at zero) than the principal's, and their profits are equal or bigger than the firm owner's. We will also show different variants of calculating agent's bonuses, but the general rule of not punishing the agent with negative bonuses is kept to.

As a consequence, the company profit in a time frame t can be formed into the following equation:

$$Profit_t = PS_t \cdot (1 - AC_t) - SL_{t-1} \cdot SC - PND_t \cdot PT,$$

where PS_t – products sold at iteration t , AC_t – average cost of the inventory at iteration t , SL_{t-1} – stock level at the end of iteration $t - 1$, SC – storage cost, PND_t – products not delivered to customers at iteration t , PT – penalty. Note that the average cost AC_t is different at each step of the simulation as it is the result of the previously negotiated prices for the inventory and currently negotiated price for the ordered items. Given that the said cost is calculated after the stocks replenishment and before selling the products at iteration t , and that NP_t are the prices negotiated with the suppliers for the number of items SR_t bought in this step of the simulation as stock replenishment, the average cost at iteration t can be written down as

$$AC_t = \frac{SL_{t-1} \cdot AC_{t-1} + SR_t \cdot NP_t}{SL_t}.$$

Note that if no order was made $SR_t \cdot NP_t = 0$, and the average cost will derive strictly from the average prices at which the inventory was purchased.

For n iterations of the simulation, the profits for the principal and the agent are equal, respectively:

$$PrincipalProfit = \sum_{t=1}^n Profit_t$$

$$AgentProfit = \sum_{t=1}^n \max\{0, Profit_t\}.$$

In the experiments, we perform 10,000 iterations of the simulation for each problem parametrization. Being guided by the results and charts stability, we found that it was the sufficient number of iterations to avoid drawing conclusions from the noise instead of the actual patterns. Table 1 shows what parameters were manipulated and what were the ranges of changes in the study.

Table 1. Domain ranges of the simulation parameters

Parameter	Range
storage cost	0.05 – 0.1
penalty	0.3 – 0.6
default rate	0% – 20%
inventory goal	0 – 4000
number of suppliers	1 – 20

Source: own elaboration

4. Results of experiments

In this section, we will discuss the simulation results obtained from the model. First, we assume a no-disruptions scenario. In such a case, there are no supplier defaults, and the financial results of the agent and the principal are dependent only on how well their stock levels are adjusted to the demand they have to satisfy and the cost related to the inventory.

Figure 5 presents the curves of the profits for both the agent and the principal, assuming a different number of suppliers used (orders equally distributed across the suppliers). Each point on the curves presents the expected profit in each of the 10,000 iterations (agent's profits are floored at 0, the principal's are not). It allows us to interpret the plots as the expected profit generated in each of the iterations, assuming different strategies (i.e., inventory goal and the number of suppliers).

The first matter we would like to address is the difference in average profit between the principal and the agent. As mentioned in Section 3, the agent does not bear any costs related to noted losses. Hence, for the lowest inventory goals, when the firm cannot deliver all of the demanded products, the agent's profit is zero, whereas the principal bears all of the costs, ignored by the agent.




Figure 5. Profit profiles for storage cost=0.05, penalty=0.3, default rate=0%

Source: own elaboration

The second mechanism this chart shows is that a higher number of suppliers (absolute differences are, of course, more significant for a small number of suppliers) translates into lower profits – because of the difference in negotiated price, i.e., a gap in a quantity discount. Notably, when the number of suppliers is exceptionally high, no or a little discount is granted, so there is not much difference between 10 and 20 suppliers in terms of profit.

The third characteristic that the curves present is that after reaching the average profit peak at some inventory goal, a decline stemming from additional storage costs begins. The bigger the storage cost, the steeper the drop. In the example illustrated here, it equals 5% of the final product's price, i.e., 0.05.

The fourth feature is related to the curves' unimodality – in the single supplier case, after reaching the peak, the higher the inventory goal (i.e., the amount the company aims to store), the lower the profit. However, this mechanism is no longer valid for more suppliers. In the case of two suppliers, for the argument that maximized the profit for one supplier, the agent's payoff (but also the principal's) has a plateau – the optimum is shifted into higher (around 2500 units) values. This feature is related to a) lower quantity discount (i.e., more suppliers

result in lower quantity orders in each of them) but also b) higher inventory goal, and thus more significant (but less frequent) orders made. In this scenario, the company is making less frequent, yet significant orders. In the absence of risk (i.e., no default rates), this strategy is sub-optimal – the company is losing the discount (as the profit curve for one supplier is higher than in the case of more suppliers) for the diversification purposes, but yet there is no risk involved.

We will now shift into *real-world* scenarios, where suppliers, faced with external shocks, may stop operating or even default. Those scenarios, especially with significant global disruptions, are common and often encountered in supply chains. For instance, in terms of the recent COVID-19 pandemic, the GDP of the US has contracted by 9.5% in January-June 2020 (Greenwood, Iverson and Thesmar, 2020); for more on the effect of the recent pandemic on corporate bankruptcy rates we refer to Wang et al. (2020).

After introducing the risk into the model, the curves in Fig. 6 remained similar to the previous outcomes, with a few significant differences. In the agent's case, the optimal profit in the setup with two suppliers is almost identical as with one supplier – yet it is obtained with a much larger inventory goal (namely, a shift from 1100 to 2500). This means that the agent may take two distinct strategies: a) relying on one supplier (and hence earning quantity discount), or b) having two suppliers (and thus a lower probability of not covering the demand) and making less frequent, but larger, orders (and bearing the storage costs). On the other hand, the principal's profit is maximized when two suppliers are engaged (the latter strategy of the agent).




Figure 6. Profit profiles for storage cost=0.05, penalty=0.3, default rate=10%
Source: own elaboration




Figure 7. Profit profiles for storage cost=0.05, penalty=0.6, default rate=10%
Source: own elaboration

In the case of the two costs (that, as mentioned earlier, can be determined for each economy sector/business branch independently), the penalty cost regulates level and shape of the curve on the left-hand side from the extrema, while the storage costs balance the right hand side.

Figure 7 presents the scenario of having the penalty cost higher (0.6) than in previous examples. The depth of the loss generated for the principal, in case of not covering the demand, is bigger, resulting in even -500 units per iteration on average (no inventory goal area). Moreover, the higher the penalty, the more dominant is the two-suppliers strategy for the principal.

On the other hand, in a sector with higher storage costs, the second scenario (i.e., ordering less frequent but larger quantities from two suppliers) is less profitable (on average) than previously. Figure 8 presents the simulation results with doubled storage costs, *ceteris paribus*.

The previous figures showed the results for the simulation, meant to investigate the sensitivity analysis of the profit curve. As in the deterministic scenario, the left-hand side (i.e., to the left of the maximum) of the curve represents the cases where the company often did not cover the demand. In contrast, the right-hand side of the curve represents the pace at which the company loses money due to storage price. The model shows that the profit decreases along with the rise in the number of suppliers (above two suppliers) for both the agent and the principal. This effect is strictly related to the quantity discounts earned when larger quantities are ordered from a single company (as mentioned earlier – a single supplier with more frequent orders or two suppliers with less periodic replenish-




Figure 8. Profit profiles for storage cost=0.1, penalty=0.3, default rate=10%. Source: own elaboration

ment). However, the analysis above is related to the expected values of profits – the measure representing the *average* scenario (e.g., no tail-observations considered, such as heavy disturbances and a small number of co-occurring defaults). For a broader discussion, in Fig. 9 we also include the analysis presenting the standard deviations of profits (in order to account for the uncertainty over the expectations).

What is worth mentioning, regarding the curves above, is the fact that subsequent differences between various numbers of suppliers (e.g., 1 vs. 2, 2 vs. 3, ..., 10 vs. 11, ...) are diminishing, i.e., the marginal loss (due to smaller quantity discount) is diminishing with regard to the number of suppliers. On the other hand, the uncertainty of the profit (measured as the standard deviation of the profit) is narrowing, which results in a more predictable business outcome. From the shareholder (principal) perspective, that may not be negligible, especially if the current economic situation involves frequent disturbances. To calculate an integrated measure (i.e., expectations adjusted with the uncertainty), we utilized the Sharpe measure (Sharpe, 1994), penalizing the outcomes with severe uncertainty. We claim that, especially in risk-sensitive situations, the stakeholders may select a relatively smaller payoff conditional on smaller volatility. The impact of the uncertainty regarding the results may be neglected in a risk-neutral world, where the company could buy the insurance for potential losses. As this is usually not the case, the stakeholders may desire to tradeoff




Figure 9. Profit and standard deviation profiles for storage cost=0.05, penalty=0.3, default rate=10%. Source: own elaboration

the profit for a more certain output:

$$S_p = \frac{E[R_p]}{\sigma_p},$$

where R_p is the profit of the party p and σ_p is the standard deviation of the profit. The results are presented in Fig. 10.




Figure 10. Sharpe ratio for storage cost= 0.05, penalty= 0.3, default rate= 10%
Source: own elaboration

Notably, as the profit does not change significantly over the next suppliers but the standard deviation narrows down dramatically, one may observe that the most balanced strategy (i.e., the one accounting for the risk) is to strongly diversify the portfolio at low inventory levels. The larger the costs of not covering the demand (i.e., a penalty imposed), the more significant incentives the owner has to a) diversify the suppliers (the larger number of suppliers, the higher Sharpe ratio) and b) hold bigger safety buffers. In Fig. 11 we present the simulation results with heavy penalty costs (twice as big as in previous examples). What is worth noting is that in the case of an agent, the scenarios do not differ much – only the biggest numbers of suppliers (i.e., 10, 20) outperform significantly (the rest of the curves exhibit similar behaviour) in terms of the Sharpe ratios the other strategies/numbers of suppliers. However, it is the principal who can experience severe effects of no-diversification – for the principal, the one-supplier scenario is strongly distinguishable from the rest of the scenarios. From the equity point of view, as the profit function for the principal is related to the sum of profits over the whole time-horizon (which may be perceived as a valuation of the company – excluding the value of the assets and any liabilities), the best strategy is to hold almost 2,200 units of products in inventory.

Moreover, this effect is even more severe if the demand is more restrictive (i.e., the penalty for not covering the orders is higher). Figure 11 presents the Sharpe ratio, assuming the higher penalty imposed. In terms of the Sharpe ratio from the agent's perspective, the strategies of two suppliers (less frequent orders) and plenty of suppliers (10 or 20 suppliers) are almost identical.




Figure 11. Sharpe ratio for storage cost=0.05, penalty=0.6, default rate=10%
Source: own elaboration

Furthermore, as shown in Fig. 12, the higher cost of storage shifts this indifference into the predominance of a high number of suppliers (as the storage cost is high, the less recurring but larger orders will also be not preferable). In this setup, with large storage costs, the best strategy (assuming the Sharpe ratio index) is to have plenty of suppliers and no significant safety buffer (and hence also smaller quantity discounts).

A similar analysis, meant to highlight the discussion above, was performed for the higher default rate regime – 20%. Figure 13 presents the simulation results under the assumption that the suppliers default more frequently (hence the company is more exposed to not covering the demand). However, in this regime, the diversification advantages are more apparent – the average profits for both players are higher for the two-suppliers strategy.

The dominance of this strategy is clearer from the Sharpe ratio perspective – i.e., average profit adjusted with the uncertainty – Fig. 14 shows the Sharpe ratio assuming a higher default rate.

The agent's profit function may (and in practice does) incentivize them to select the principal's sub-optimal supply-chain strategies – as in the first example, the agent could choose one supplier with a narrower safety buffer (especially if




Figure 12. Sharpe ratio for storage cost=0.1, penalty=0.3, default rate=10%
Source: own elaboration

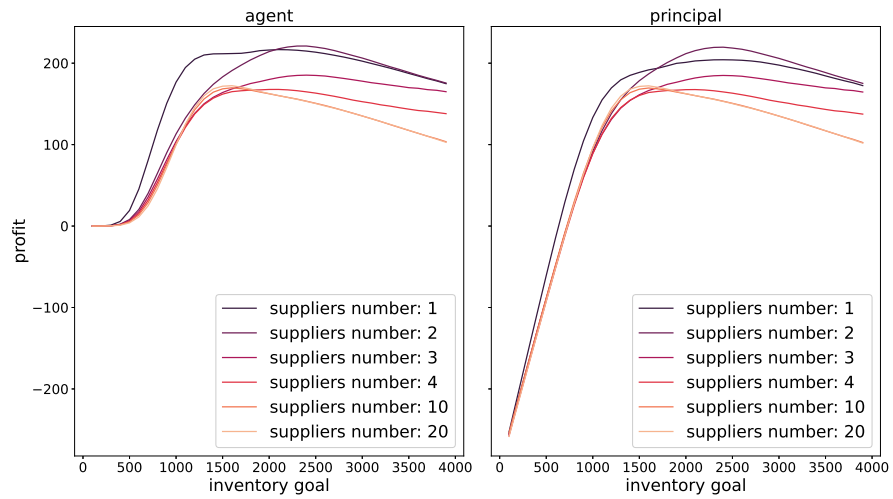


Figure 13. Profit profiles for storage cost=0.05, penalty=0.3, default rate=20%
Source: own elaboration




Figure 14. Sharpe ratio for storage cost=0.05, penalty=0.3, default rate=20%
Source: own elaboration

the storage costs are high). Two actions can be taken to incentivize the agent to act in favour of the principal. First, the agent's profit function might be linked with the one of the principal – e.g., via the stock price. In publicly traded companies, the managers (agents) are incentivized to protect the stock prices (held by the shareholders – principals). However, as discussed in the article *It's Time to Replace the Public Corporation* (Martin, 2021), this may lead to focusing on the short-term goals (e.g., smaller R&D investments).

In the respective analysis, we have applied a multiple-period bonus system, where for the principal, a simple 5-period mean is calculated. In contrast, in the agent's case, a significant change in the profit evaluation was made. Namely, negative profits are not ignored anymore. Instead, they are taken into account in the bonus calculation process (we account for both profits and losses). The premium is paid to the manager if the outcome is positive – the assumption of not punishing the agent is still in power. Thus, in Fig. 15 we depict the simple average bonus (without uncertainty), while in Fig. 16 we show the Sharpe ratio indicator. Notably, the profiles observed for both measures are now remarkably compatible. The agent participates in the principal's risk, which is why the shapes of the curves and the optimal solutions for both players overlap. Hence, the manager's decisions should be more aligned with the principal's interests.

The second action that may be taken to change the agent's perspective consists in imposing the medium- or long-term profits. This practice is commonly used in management contracts, usually implemented in employee stock ownership plans (ESOP). Embedding some of the agent's profit into the company




Figure 15. Profit profiles for storage cost=0.05, penalty=0.3, default rate=10%, 5-period value

Source: own elaboration




Figure 16. Sharpe ratio for storage cost=0.05, penalty=0.3, default rate=10%, 5-period value

Source: own elaboration

value (i.e., in our analysis, long-time profits) will allow the principal to shift the agent's incentives and hold more robust and safety strategies. As presented in Fig. 7, the agent's profit was almost identical in both strategies (i.e., frequent orders from one supplier or less periodic replenishments from two suppliers). However, upon adding the principal's perspective (accounting for the risk), the agent would be incentivized to hold the scenario more profitable for the principal – the one with lower uncertainty (especially when higher penalties for not meeting the demand are imposed).

5. Conclusions

This paper has proposed a simple simulation-based framework, allowing for analyzing the processes observed in supply chains. The simulation enables generalization, i.e., regarding different quantity discount functions or a bigger number of supplier layers, that is – a multi-layered supply chain. Its leading purpose was to allow for a straightforward analysis of the optimum decision-making in the supply chains, assuming some parametrization (e.g., the storage costs or penalty accounted for every time the company cannot cover the demand).

In terms of the supply chain, we have identified several parameters that bear an impact on the replenishment processes. As of the identified strategies, allowing to earn the quantity discounts, we found that both: a) small inventory goals with recurring orders from a single supplier, and b) large inventory goals and less frequent orders from multiple suppliers are the sweet spots in the simulation. Based on the sensitivity analysis, we have pointed out that indeed, if the storage costs are low, instead of frequent restocking, the quantity discount may be obtained by infrequent but large orders from several suppliers, reducing the risk related to suppliers' defaults. The former strategy is more cost-efficient in the regimes with high storage costs and low penalties related to not covering the demand. However, the latter strategy is more robust, as it utilizes multiple sources of products (and hence is less prone to not covering the demand).

As a concluding note, we emphasize the need of identifying the specification of the market that the company operates in – the final outcome is dependent on the level of risk encountered on the market (i.e., probability of default of suppliers), storage costs (which for different markets may vary heavily) and penalty imposed on not covering the demand (which is related to market/product/location-specific factors). Moreover, we have included the Sharpe ratio analysis to investigate the risk-adjusted expected profit of players – the superior strategies assumed a larger number of suppliers. In terms of risk-adjusted expected returns, we have pointed out that almost always (except for the deterministic case with no risk) a more diversified suppliers' portfolio is preferred.

Using the proposed model, we could assess the gap between the two objective functions (of the principal and the agent) in terms of the agency problem. Given

various incentives and time-horizons of the agent and the principal (i.e., the agent is in practice more short-medium term oriented), we pointed out that agents may be prone to selecting the decision parameters sub-optimal from the perspective of the principal. This effect is mainly embedded in the risk exposure of particular parties. An employee (i.e., agent) is not exposed to the negative returns (however, the principal may attempt to transfer some past losses on the agent) in practice. Our analysis also included the agent's prolonged time-horizons; if the profit accumulation time horizon is longer for an agent (even 5-period), the profit profile of the agent converges to the one of the principal. This finding provides the evidence that the ESOP or other forms of prolonging the agent's incentives (e.g., stock options) may reduce the agency problem – in practice, these solutions put agents into the principals' shoes. The debate that we have included is present in current literature and relates to various (short/long) time horizons of various stakeholders' perspectives.

As pointed out earlier, we assume that each of the economy sectors/business branches has a unique specification – namely, the cost parameters for various sectors will be different. The model generalizes the business process (i.e., replenishment) and allows for reparametrization to a specific sector/segment. One of our work's further extensions is the calibration of the model's parameters to empirical data. As discussed earlier, various business segments will exhibit different parametrization of the model. For example, some of the branches may easily attract new demand – i.e., penalty cost is lower – others will impose more attention to keeping the current customers. Moreover, the natural advancements to the model may include:

1. Other quantity discount functions than the sigmoidal and step functions we have used.
2. Heterogenous suppliers, i.e., in this paper, we assumed that the suppliers are homogenous.
3. Different profit time-discounting functions, i.e., we assumed no time-money value discounting.
4. A broader range of products, i.e., introducing complementary and substitute products.
5. An analytical model of the problem, supplemented with numerical experiments for testing the derived relationships.

Based on those advancements, the optimization techniques may be utilized to reveal the agent and principal's optimum strategies.

Bibliography

- ANHOLCER, M., HINC, T. AND KAWA, A. (2019) Losses in Transportation - Importance and Methods of Handling. In: A. Kawa and A. Maryniak, *SMART Supply Network*. Springer International Publishing, Berlin Heidelberg, 111-128.

- BARTON, D., MANYIKA, J. AND WILLIAMSON, S. K. (2017) Finally, evidence that managing for the long term pays off. *Harvard Business Review*, 9. <https://hbr.org/2017/02finally-proof-that-managing-for-the-long-term-pays-off>
- CRAIGHEAD, C. W., BLACKHURST, J., RUNGTUSANATHAM, M. AND HANDFIELDS, R. B. (2007) The Severity of Supply Chain Disruptions: Design Characteristics and Mitigation Capabilities. *Decision Sciences*, **38**(1), 131-156.
- DE BOER, L., LABRO, E. AND MORLACCHI, P. (2001) A review of methods supporting supplier selection. *European Journal of Purchasing and Supply Management*, **7**(2), 75–89.
- ELOCK SON, C. (2018) Supply Chain Risk Management: A Review of Thirteen Years of Research. *American Journal of Industrial and Business Management*, **8**(12), 2294.
- Financial Times* (2020, April 22) Companies should shift from ‘just in time’ to ‘just in case’. *Financial Times*. Retrieved from *Financial Times*: <https://www.ft.com/content/606d1460-83c6-11ea-b555-37a289098206>
- GREENWOOD, R. M., IVERSON, B. C. AND THESMAR, D. (2020) Sizing up corporate restructuring in the covid crisis. *NBER Working Paper*. no. 28104, November.
- GROSSMAN, S. AND HART, O. (1983) An Analysis of the Principal-Agent Problem. *Econometrica* **51**, 302-340.
- HADWICK, A. (2020, July 3) The end of just-in-time? *Reuters Events*.
- JANDA, K. (2006) *Lender and borrower as principal and agent*. Charles University Prague, Faculty of Social Sciences, Institute of Economic Studies.
- JENSEN, M. C. AND MECKLING, W. H. (1976) Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, **3**(4), 305-360.
- KERSTEN, G. E. (1985) NEGOT—group decision support system. *Information & Management*, **8**(5), 237-246.
- KERSTEN, G. E. (1997) Support for group decisions and negotiations an overview. In: J. Climavco, ed., *Multicriteria Analysis*, Springer Verlag, 332–346.
- KUWORNUN, J., KUIPER, W. AND PENNING, J. (2009) Agency problem and hedging in agri-food chains: Model and application. *Journal of Marketing Channels*, **16**(3), 265-289.
- LI, G., KOU, G. AND PENG, Y. (2016) A group decision making model for integrating heterogeneous information. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, **48**(6), 982-992.
- MARTIN, R. L. (2021) It’s Time to Replace the Public Corporation. *Harvard Business Review* **99**(1), 34–42.
- MINNER, S. (2003) Multiple-supplier inventory models in supply chain management: A review. *International Journal of Production Economics*, **81-82**, 265–279.

- MYERSON, R. (1982) Optimal coordination mechanisms in generalized principal-agent problems. *Journal of Mathematical Economics*, **10**(1), 67-81.
- PERSSON, F. AND OLHAGER, J. (2002) Performance simulation of supply chain designs. *International Journal of Production Economics*, **77**(3), 231-245.
- PETROVIC, D., ROY, R. AND PETROVIC, R. (1998) Modelling and simulation of a supply chain in an uncertain environment. *European Journal of Operational Research*, **109**(2), 299-309.
- POURHEJAZY, P. AND KWON, O. K. (2016) The New Generation of Operations Research Methods in Supply Chain Optimization: A Review. *Sustainability*, **8**(10), 1033. <https://doi.org/10.3390/su8101033>
- PÁEZ-PÉREZ, D. AND SÁNCHEZ-SILVA, M. (2019) A dynamic principal-agent framework for modeling the performance of infrastructure. *European Journal of Operational Research*, **254**(2), 576-592.
- RAJHANS, N., PATIL, S. AND KULKARNI, P. (2015) Mathematical Model for Calculating Inventory Carrying Cost. *IIIE Journal*, **8**(1), 12-14.
- ROSS, S. (1963) The economic theory of agency: The principal's problem. *The American Economic Review*, **63**(2), 134-139.
- SHARPE, W. F. (1994) The Sharpe Ratio. *Journal of Portfolio Management* **21**(1), 49-58.
- TANG, C. S. (2006) Perspectives in supply chain risk management. *International Journal of Production Economics*, **103**(2).
- TANG, C. S. (2006) Robust strategies for mitigating supply chain disruptions. *International Journal of Logistics*, **9**(1), 33-45.
- TSAY, A., NAHMIA, S. AND AGRAWAL, N. (1999) Modeling supply chain contracts: A review. In: *Quantitative Models for Supply Chain Management*. Springer, Boston.
- WANG, J., YANG, J., IVERSON, B. C. AND KLUENDER, R. (2020) Bankruptcy and the COVID-19 Crisis. *Harvard Business School*. Working Paper, SSRN 3690398.
- WANG, S., CHEN, K., LIU, Z., GUO, R.-Y. AND CHEN, S. (2018) An ontology-based approach for supply-chain quality control: From a principal-agent perspective. *Journal of Information Science*, **45**(3), 283-303.
- WILLIAMSON, O. AND MASTEN, S. (1999) *The Economics of Transaction Costs*. Edward Elgar Publishing.