Supply Network Coordination by Vendor Managed Inventory – A Mechanism Design Approach

Péter Egri[1], **József Váncza**[1,2]
[1]Fraunhofer Project Center for Production Management and Informatics
Computer and Automation Research Institute, Hungarian Academy of Sciences
Kende u. 13-17, 1111 Budapest, Hungary
[2]Department of Manufacturing Science and Technology, Budapest University of Technology and Economics
H-1111 Budapest, Egry József u. 1., Hungary

Abstract

The paper studies a generic coordination problem in supply networks with some retailers and a single supplier agent. The parties posses private information—the retailers on uncertain demand forecasts, the supplier on production costs—and seek to maximise their own utilities. So as to design a coordination mechanism that warrants the satisfaction of all market demand at maximal social welfare, we model the problem as a non-cooperative game. By relying on the conceptual apparatus of mechanism design theory, we provide an analytical explanation for the widely used vendor managed inventory (VMI) where the responsibility of planning and all the risks of over- and underproduction are at the supplier. After proving that under reasonable assumptions a fair sharing of these risks is not possible, we present a family of coordination mechanisms that are efficient and can be implemented without an individually existing mechanism. Beyond new managerial insights—like modelling VMI as a service where payment should also depend on the accuracy of information communicated—a novel approach is provided to handling supply networks consisting of autonomous agents.

1 Introduction

Supply networks are large and complex systems, characterised by the existence of numerous competitive agents, dynamic structures, uncertain knowledge and difficult planning and decision making problems. The uncoordinated actions in such a system lead to e.g., suboptimal performance, exemplified in a simple case by the well-known *prisoners' dilemma*. In supply networks the appearance of this phenomenon is called *double marginalisation*: since every enterprise concern their own profit when making decisions, the aggregate benefit is in general lower than if the enterprises were vertically integrated. This suboptimality also results in waste of materials, labour, energy, environmental resources and eventually causes significant financial losses for the enterprises.

In a vertically integrated supply network with multiple retailers and a supplier, centralising the replenishment and inventory management decisions at the supplier side is advantageous compared to the situation where each retailer has to decide individually. This centralisation approach is called *risk pooling*, and it is proved to result both in lower overall safety stocks and in lower average inventory levels [Simchi-Levi *et al.*, 2000].

In order to use the idea of risk pooling in vertically non-integrated networks, the *vendor managed inventory* (VMI) business model is applied frequently. In VMI the supplier takes all risks and full responsibility for managing a one-point inventory, while it has to fulfil the entire demand of the retailers, even if this requires additional costs due to extra capacity usage, overtime, outsourcing, or rush production orders [Simchi-Levi *et al.*, 2000]. This situation is clearly disadvantageous for the supplier, in fact, the main practical reason underlying VMI is the market power of the retailers, and not the mutual interest of the partners. Furthermore, since the retailers are not faced with the consequences of an imprecise forecast directly, they are not inspired to increase their efforts in accurate forecasting.

Sometimes the retailers have incentives even for distorting the forecasts. If the performance of the retailers are measured by the eventual shortage, then they tend to overplan demand and forward too optimistic plans towards the supplier. On the other hand, if the retailers are rewarded for overperforming the plans, then they tend to underestimate the demand. In both cases, the selfish distortion of information will introduce additional uncertainty into the demand forecasts, and lead to higher operational costs.

2 Related Work

Agents provide a natural metaphor for manufacturing in supply networks, where the knowledge is both incomplete (distributed) and imprecise (uncertain) [Egri and Váncza, 2007]. Unfortunately, most of the existing multi-agent models assume *benevolence*—i.e., the agents must implicitly share a common goal—, which holds in some situations, but certainly not in a supply chain of autonomous enterprises. Lack of benevolence can result not only in suboptimal behaviour, but also in the collapse of the production process, as an earlier study in decentralised production scheduling showed [Váncza and Márkus, 2000].

Therefore, the number of deployed multi-agent systems that are already running in real industrial environments is unsurprisingly small. An other important reason for this is that in the behaviour of a multi-agent system there is always an element of *emergence* which can be a serious barrier to the practical acceptance of agent-based solutions. Industry needs safeguards against unpredictable behaviour and guarantees regarding reliability, safety and operational performance [Monostori *et al.*, 2006].

Such guarantees can be given by applying the results of game theory and *mechanism design* [Rosenschein and Zlotkin, 1994; Shoham and Leyton-Brown, 2008; van der Krogt *et al.*, 2008]. For example, the auction theory has already demonstrated its value with several applications, even in the electronic markets. This success is mainly due to the fact that these markets are *well-structured* with clear, exact regulations. Where the conditions of control are given by formal rules, introducing the concepts and apparatus of game theory is a really promising approach. That is why it has recently become popular for analysing the behaviour of automated agents operating on the Internet [Dash *et al.*, 2003; Nisan *et al.*, 2007].

Considering enterprises with own objectives also resulted in various game theoretic formalisations, both cooperative and non-cooperative ones. The former approach is taken for studying coalition formation, stability analysis, bargaining or profit allocation [Nagarajan and Sosic, 2008]. The non-cooperative models on the other hand, usually apply sequential games; especially the use of the *principal – agent* model of contracting theory [Laffont, 2001; Salanié, 2005] is common. In the operational research literature this approach is called *supply chain coordination* [Arshinder *et al.*, 2008]. These researches are usually related to inventory management in distributed production planning problems. Most of the works study the distributed version of the one-period *newsvendor lot-sizing problem* due to its simple structure; for a review of newsvedor games we refer to [Cachon and Netessine, 2004].

Two main problems with the majority of the current studies in the literature are that they (i) consider rather special production problems lacking generality, and (ii) usually take only simple concepts from game theory (like e.g., the Stackelberg games [Hennet and Arda, 2008]). Both the supply chain management and the game theory literature have some practically more relevant results which should be combined and further studied, such as rolling horizon planning, hierarchical planning systems, repeated games, equilibrium learning, Vickrey – Clarke – Groves mechanisms, to name a few. In the mechanism design theory there are also recent achievements considering algorithmic issues, such as verification ([Nisan and Ronen, 2001]), distributed mechanisms ([Shneidman and Parkes, 2004]) and stochastic problems ([Ieong *et al.*, 2007; Papakonstantinou *et al.*, 2011]). This paper intends to be a further step on this way.

The remainder of the paper is organised as follows. In Section 3, we model distributed decision making in the supply network as a mechanism design problem and inspect some of its properties. We introduce some assumptions into the model in Section 4, in order to develop practically applicable efficient coordination mechanisms. In Section 5, we illustrate the performance of the VMI compared to the traditional order-based purchase on a numerical example. Finally, we conclude the results of this paper and suggest some future research directions.

3 A Mechanism Design Analysis

In this section we formalise the supply network model with n retailer agents and a supplier agent. For the sake of simplicity, we assume that the *retailers* are homogeneous, although this assumption can be relaxed and the results still remain valid. Retailer i has some private belief (forecast) about the future market *demand*, which is denoted by $\theta_i \in \Theta$. Demand is satisfied by production done at the *supplier* who has, in turn, private information about the cost factors. Since the exact demands $\xi_i \in D$ realise at some later time, only the forecasts can be considered when creating a *production plan* denoted by $x \in \mathcal{K}$. If the actual demand does not match the forecast, then the production will deviate from what was planned. If the demand was underestimated, new and costly production is necessary, while overestimation leads to extra, sometimes even to obsolete inventories. In both cases, the actual costs incurred are higher than planned. Therefore the production cost at the supplier is a function of the original production plan as well as of the realized demands: $c \in C = \{c : \mathcal{K} \times D^n \to \mathbb{R}\}$, which is a private information of the supplier. Note that we do not assume that an a priori distribution about the private information is known by the other agents, i.e., we regard a situation with *strict incomplete information*.

According to the classic mechanism design theory, an independent mediator, the *mechanism* is required for observing the agents' actions, making the decision and after realisation, transferring the payments among the agents. Some recent developments aim at omitting the mediator, which possibility we also will study in the next section. For the moment, let us define the mechanism as $\mathcal{M} = (f, t_1, \ldots, t_n, t_s)$, where $f : \Theta^n \times C \to \mathcal{K}$ is the choice function determining the production plan based on the forecasts and the cost function[1], $t_i : \Theta^n \times C \times D^n \to \mathbb{R}$ are the payment functions of the retailers ($i = 1, \ldots, n$), and $t_s : \Theta^n \times C \times D^n \to \mathbb{R}$ is the payment for the supplier.

Note two assumptions of this formulation. Firstly, this is a *direct-revelation* mechanism, i.e., the strategy of the agents is to share their private information (not necessarily truthfully) with the mediator. Secondly, we consider that the realised demands are commonly observable.

After the demands realise, an income arises at each retailer i from the sales: $v_i : D \to \mathbb{R}$. Now we can define the utility—the income minus the cost—for each retailer and the supplier:

$$u_i(\theta_i, \hat{\theta}, \hat{c}, \xi) = v_i(\xi_i) - t_i(\hat{\theta}, \hat{c}, \xi) \tag{1}$$

and

$$u_s(c, \hat{\theta}, \hat{c}, \xi) = t_s(\hat{\theta}, \hat{c}, \xi) - c(f(\hat{\theta}, \hat{c}), \xi) \tag{2}$$

respectively, if their private information is θ_i and c, but they claim $\hat{\theta}_i$ and $\hat{c}$ instead, with $\theta = (\theta_1, \ldots, \theta_n)$ and $\xi = (\xi_1, \ldots, \xi_n)$ denoting the forecast and demand vectors. (The

[1] Although the production planning problem is complex in general, in this paper we disregard computational issues, and assume that an optimal plan can be found for every possible forecast.

first parameters of the utility functions are the real private information, the second and third are the communicated parameters, and the last one is the realised demand.)

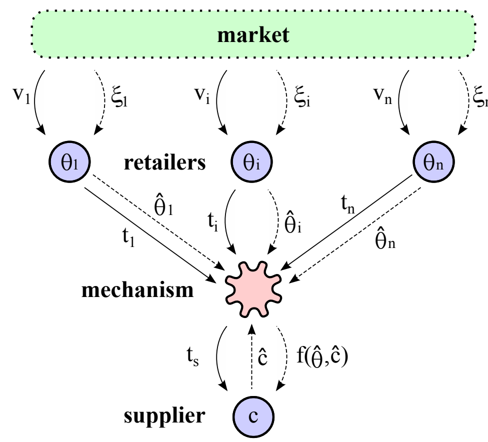

Figure 1: Mechanism design setting.

Figure 1 illustrates this mechanism design model, where the dashed arrows denote information flow, while the solid ones represent the monetary payments. The variables and functions inside the agents are private knowledge of the given agent and are unobservable for the others.

We are seeking such a mechanism, wherewith the performance of the production network as a whole is optimal. This can be guaranteed, if all the agents disclose their private information truthfully, and the mechanism uses an optimal planning choice function. Let us define these properties formally.

Definition 1 *A mechanism $\mathcal{M}$ is* (weakly) strategy-proof*, if truth telling is a* dominant strategy *for every agent, i.e., it maximizes their expected utility:* $\forall i, \forall \theta_i \in \Theta, \forall \hat{\theta} \in \Theta^n, \forall \hat{c} \in C:$

$$\mathbb{E}_{\tilde{\theta}}[u_i(\theta_i, \tilde{\theta}, \hat{c}, \xi)] \geq \mathbb{E}_{\tilde{\theta}}[u_i(\theta_i, \hat{\theta}, \hat{c}, \xi)], \tag{3}$$

where $\tilde{\theta} = (\hat{\theta}_1, \ldots, \hat{\theta}_{i-1}, \theta_i, \hat{\theta}_{i+1}, \ldots, \hat{\theta}_n)$*, and* $\forall c, \hat{c} \in C, \forall \hat{\theta} \in \Theta^n:$

$$\mathbb{E}_{\hat{\theta}}[u_s(c, \hat{\theta}, c, \xi)] \geq \mathbb{E}_{\hat{\theta}}[u_s(c, \hat{\theta}, \hat{c}, \xi)]. \tag{4}$$

Definition 2 *The choice function f is* efficient*, if it maximises* social welfare *(the sum of the utilities without the payments), i.e.,* $\forall \theta \in \Theta^n, \forall c \in C:$

$$f(\theta, c) \in \operatorname*{argmax}_{x \in \mathcal{K}} \mathbb{E}_\theta[\sum v_i(\xi_i) - c(x, \xi)] = \operatorname*{argmin}_{x \in \mathcal{K}} \mathbb{E}_\theta[c(x, \xi)]. \tag{5}$$

Firstly, we show that if a strategy-proof mechanism gives the same output for different cost functions, then it is expected to give the same payment to the supplier.

Proposition 1 *If $\mathcal{M}$ is a strategy-proof mechanism, $c, \hat{c} \in C, \theta \in \Theta^n$ such that $f(\theta, c) = f(\theta, \hat{c})$, then $\mathbb{E}_\theta[t_s(\theta, c, \xi)] = \mathbb{E}_\theta[t_s(\theta, \hat{c}, \xi)]$.*

Proof Let us assume that $\mathbb{E}_\theta[t_s(\theta, c, \xi)] < \mathbb{E}_\theta[t_s(\theta, \hat{c}, \xi)]$ (the other direction is analogous). But then

$$\mathbb{E}_\theta[t_s(\theta, c, \xi)] - \mathbb{E}_\theta[c(f(\theta, c), \xi)] < \mathbb{E}_\theta[t_s(\theta, \hat{c}, \xi)] - \mathbb{E}_\theta[c(f(\theta, \hat{c}), \xi)], \tag{6}$$

i.e., the mechanism is not strategy-proof, since the supplier with cost function c would reveal $\hat{c}$ instead. □

We now prove that if we are looking for an efficient strategy-proof mechanism, then t_s should be (in the sense of expected value) independent from the cost function of the supplier, therefore it excludes the possibility of cost sharing among the agents.

Theorem 2 *Let $\mathcal{M}$ be an efficient, strategy-proof mechanism. Then*

$$\forall \theta \in \Theta^n, \forall c, \hat{c} \in C: \mathbb{E}_\theta[t_s(\theta, c, \xi)] = \mathbb{E}_\theta[t_s(\theta, \hat{c}, \xi)]. \tag{7}$$

Proof The proof is similar to the proof of the uniqueness of *Groves mechanism* among efficient and strategy-proof mechanisms proved by [Green and Laffont, 1977], thus we exploit that the cost function can be arbitrary.

Let us consider a fixed θ, and indirectly assume that the statement of the theorem is false: $\exists c, \hat{c} \in C: \mathbb{E}_\theta[t_s(\theta, c, \xi)] > \mathbb{E}_\theta[t_s(\theta, \hat{c}, \xi)]$. Furthermore, let us define

$$\varepsilon = \mathbb{E}_\theta[t_s(\theta, c, \xi)] - \mathbb{E}_\theta[t_s(\theta, \hat{c}, \xi)] > 0. \tag{8}$$

Due to the *modus tollens* of Proposition 1, $f(\theta, c) \neq f(\theta, \hat{c})$. Because the cost function can be arbitrary, $\exists \tilde{c} \in C, \exists k \in \mathbb{R}:$

$$\mathbb{E}_\theta[\tilde{c}(f(\theta, \hat{c}), \xi)] = k \tag{9}$$

$$\mathbb{E}_\theta[\tilde{c}(x, \xi)] > k \qquad \forall x \neq f(\theta, \hat{c}) \tag{10}$$

$$\mathbb{E}_\theta[\tilde{c}(f(\theta, c), \xi)] < k + \varepsilon. \tag{11}$$

From the efficiency of the mechanism follows that $f(\theta, \tilde{c}) = f(\theta, \hat{c})$, and then from Proposition 1, we have $\mathbb{E}_\theta[t_s(\theta, \tilde{c}, \xi)] = \mathbb{E}_\theta[t_s(\theta, \hat{c}, \xi)]$. But then

$$\begin{aligned} \mathbb{E}_\theta[u_s(\tilde{c}, \theta, \tilde{c}, \xi)] &= \mathbb{E}_\theta[t_s(\theta, \tilde{c}, \xi)] - \mathbb{E}_\theta[\tilde{c}(f(\theta, \tilde{c}), \xi)] \\ &= \mathbb{E}_\theta[t_s(\theta, \hat{c}, \xi)] - k, \end{aligned} \tag{12}$$

and

$$\begin{aligned} \mathbb{E}_\theta[u_s(\tilde{c}, \theta, c, \xi)] &= \mathbb{E}_\theta[t_s(\theta, c, \xi)] - \mathbb{E}_\theta[\tilde{c}(f(\theta, c), \xi)] \\ &> \mathbb{E}_\theta[t_s(\theta, c, \xi)] - k - \varepsilon \\ &= \mathbb{E}_\theta[t_s(\theta, \hat{c}, \xi)] - k, \end{aligned} \tag{13}$$

thus the mechanism is not strategy-proof. □

The theorem proves the reasonable conjecture that the supplier can claim higher costs in such a way, that the optimal production plan does not change. Thus, the supplier may try to obtain more payment without increasing its costs.

4 Coordination Mechanisms for Supply Networks

Although Theorem 2 is rather negative if we are aimed at fair cost sharing, it has some positive consequences as well that we analyse in this section. Our main goal is to omit the necessity of an independent mediator which is unrealistic in a supply network, but at the same time, to preserve the favourable properties of the system. In what follows, we dissolve the two reasons for the existence of an independent mechanism: balancing the difference between the agents' payments and providing efficiency.

From now on, we assume that the payment of the supplier is independent from its revealed cost function—explained by the conclusions of Section 3—, and the next definition necessitates that the payments of the retailers are also independent from $\hat{c}$.

Definition 3 *A mechanism $\mathcal{M}$ is* budget-balanced*, if*

$$\forall \hat{\theta} \in \Theta^n, \forall \xi \in D^n : t_s(\hat{\theta}, \xi) = \sum t_i(\hat{\theta}, \xi), \tag{14}$$

i.e., there is no surplus or deficit for the mechanism, the total payment is distributed among the agents.

Note that requiring budget-balance excludes the application of the Vickrey – Clarke – Groves (VCG) mechanisms which is one of the main positive results of the classic mechanism design theory, and it is also frequently applied for solving algorithmic problems [Nisan *et al.*, 2007].

A direct corollary of the independence of t_s from $\hat{c}$ is that the supplier agent can maximise its utility by minimising its cost. This means that the efficiency of the mechanism corresponds with the supplier's interest, therefore the mechanism can be implemented by the supplier without requiring an independent mediator. There is no need to disclose its private information about the costs, and furthermore, even the specific planning algorithm and the resulted plan can be kept secret. In fact, this property is the essence of VMI.

In order to provide strategy-proofness, truth telling should be the dominant strategy for each retailer, independently from the decision and realised demand of the other retailers. Since the income is independent from the disclosed information, the utility is maximal when the payment is minimal.

Firstly, let us consider a trivial example for illustration, when the forecast is simply the expected value of the demand. In this case it is easy to see that for example the payment function in the form

$$t_i(\hat{\theta}, \xi) = \alpha_i |\hat{\theta}_i - \xi_i| + \beta_i(\hat{\theta}_{-i}, \xi), \tag{15}$$

where $\alpha_i > 0$ is a constant, β_i is an arbitrary function and $\hat{\theta}_{-i} = (\hat{\theta}_1, \ldots, \hat{\theta}_{i-1}, \hat{\theta}_{i+1}, \ldots, \hat{\theta}_n)$, is appropriate, since the first term is expected to be minimal when $\hat{\theta}_i = \mathbb{E}[\xi_i]$, and the second term is independent from $\hat{\theta}_i$.

If $\beta_i(\hat{\theta}_{-i}, \xi)$ depends on $\hat{\theta}_j$ and ξ_j $(j \neq i)$, this allows some profit sharing between the retailers, but requires cooperation between them. Otherwise $\beta_i(\hat{\theta}_{-i}, \xi) = \gamma_i(\xi_i)$ with some arbitrary γ_i function, and the retailers' profits are independent from each other. In this case, $\gamma_i(\xi_i)$ practically can be considered as the payment for the supplied products, while α_i defines the price of the flexible VMI service.

When the forecast becomes more complex, it is not straightforward to guarantee strategy-proofness. If for example we refine the previous model assuming the forecast is given by the *expected value* and the *standard deviation*, the task becomes more interesting. The difference between the expected and realised demand can be easily measured, but how can we estimate the accuracy of the standard deviation based only on one observation? In the following, we answer this question by presenting a strongly strategy-proof payment, which is a generalisation of the result published in [Egri, 2008], without assuming any particular distribution of the demand. Due to the lack of space, we omit the proof which is analogous to the one presented in the previously mentioned thesis.

Theorem 3 *Let us consider a one-period supply coordination network problem, where the forecasts are given by the expected values and the standard deviations, i.e., $\theta_i = (m_i, \sigma_i)$. Then the payment function in the form*

$$t_i(\hat{m}, \hat{\sigma}, \xi) = \alpha_i \left(\frac{(\hat{m}_i - \xi_i)^2}{\hat{\sigma}_i} + \hat{\sigma}_i \right) + \beta_i(\hat{m}_{-i}, \hat{\sigma}_{-i}, \xi), \tag{16}$$

where $\alpha_i > 0$ is a constant and β_i is an arbitrary function, is strongly strategy-proof.

One can notice the similarity between this payment and the payment defined by Eq. (15). Furthermore, there is a simple intuition behind the term $(\hat{m}_i - \xi_i)^2/\hat{\sigma}_i + \hat{\sigma}_i$: if a retailer states that the forecast is fairly precise (i.e., σ_i is small), it is ready to pay larger compensation for the difference between the expected and the realised demand. This could be avoided by stating higher uncertainty, but then this increases the second part of the term.

Such one-period problems are widely studied in the supply chain coordination literature due to their simple structure, however, in several practical cases they cannot be properly applied. In industrial problems involving longer horizons, usually some medium-term, discrete forecasts are used; in addition, the forecasts are often updated from time to time, on a rolling horizon. These more realistic cases can be approached in a similar way as the one-period problem: besides the payment for the products (the "β-part" of the payment), the *imprecision of the forecast* should be measured and used as the basis for the payment of the VMI service ("α-part"). For example, in [Váncza *et al.*, 2008] we present a strongly strategy-proof payment scheme for the *multi-period*, *rolling horizon* case with uncertain length of product life-cycle.

All in all, with VMI the supplier not only offers products, but also flexibility as a service. Accordingly, a composite payment function should be constructed: the retailers must pay not only (i) for the quantity delivered, but also (ii) for the deviation from the forecast, as well as (iii) for the uncertainty of the forecast. This payment compensates the supplier for the eventual obsolete inventory or the cost of extra production exceeding its original production plan.

5 Computational Study

In this section we illustrate on a simple example how an efficient strategy-proof mechanism can improve the performance of a supply network. We consider n retailers, and we assume that the ξ_i demands are independent and normally distributed with expected values m_i and standard deviations σ_i.

Firstly, we examine a suboptimal solution, where the retailers make firm orders $\hat{m}_i$ and pay a w_1 wholesale price for the supplied goods. When the ξ_i demand realises, either some surplus remain at retailer i, or it has to order again, but due to the urgency, on a higher w_2 price. The supplier works in *make-to-order* mode, i.e., it produces the normal orders with c_1 piecewise cost, and the urgent orders on a higher c_2 cost. Formally, this can be expressed as follows:

$$t_i(\hat{m}, \hat{\sigma}, \xi) = w_1 \hat{m}_i + w_2 \max(\xi_i - \hat{m}_i, 0), \tag{17}$$

and the emerging cost at the supplier becomes

$$c(\xi) = c_1 \sum \hat{m}_i + c_2 \sum \max(\xi_i - \hat{m}_i, 0). \tag{18}$$

One can derive that in this newsvendor-like case the optimal order quantity is

$$\hat{m}_i = F_i^{-1}(1 - \frac{w_1}{w_2}), \tag{19}$$

where F_i is the cumulative density function (CDF) of ξ_i, thus the optimal order quantity is not even equal with the expected demand.

However, if one applies a mechanism with a payment defined by Eq. (16), then the retailers truthfully reveal their private information about the expected values and standard deviations of the demand forecasts. Now, the expected value of the total demand will be the sum of the expected values, and since the demands are considered to be independent, the standard deviation of the total demand becomes $\sqrt{\sum \sigma_i^2}$. Furthermore, if the demand at the retailers are normally distributed, the distribution of the total demand will also be normal, therefore the optimal production quantity and cost are

$$x = F^{-1}(1 - \frac{c_1}{c_2}) \tag{20}$$

and

$$c(x, \xi) = c_1 x + c_2 \max(\sum \xi_i - x, 0), \tag{21}$$

where F is the CDF of the total demand.

Figure 2 illustrates the difference between the costs of the two approaches, depending on the number of retailers. We set the price parameters as $c_1 = 50$, $c_2 = 80$, $w_1 = 135$, and $w_2 = 165$. The expected value and standard deviation of the total demand was set to $m = 800$ and $\sigma = 100$, and we considered retailers with identical distributions, therefore their parameters were $m_i = m/n$ and $\sigma_i = \sigma/\sqrt{n}$. Each cost value indicated on the figure is an average made on 5000 simulation runs.

As it can be seen, the coordinated VMI is more efficient than the order-based supply even in the one retailer case due to the elimination of the double marginalisation. However, when the number of the retailers increases, risk pooling keeps the optimality in the network, while the uncoordinated approach quickly deviates from cost efficient performance.

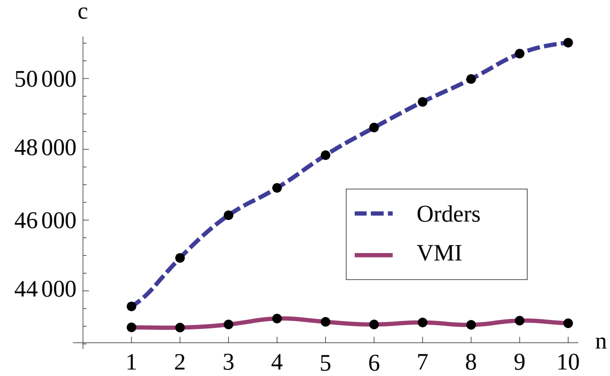

Figure 2: The cost in function of the number of retailer agents.

6 Conclusions and Further Work

In this paper we studied networks of autonomous retailers and a supplier, where the utilities depend on a stochastic market demand whose distribution is not known by the decision maker. The basic assumptions of our model are rooted in supply chain coordination models, specifically, we departed from the vendor managed inventory where demand anticipated and communicated by the retailer drives production planning at the supplier. The analysis was aimed at elaborating a direct revelation mechanism that maximizes overall utility, even though the partners are rational, expected utility maximising decision makers. The coordination problem called also for a budget-balancing solution. Furthermore, we were interested in finding an answer to a crucial question in supply chain coordination: the sharing of costs and profits. Due to the private information of the players, the traditional distribution methods of cooperative game theory were not applicable.

Firstly, we formulated a mechanism design model with the inherent incomplete information about the forecasted demand and production costs. We proved that in any mechanism that is efficient and where truth-telling is a dominant strategy for each player, the payment for the supplier's production efforts should be independent from its actual cost. Hence, this result excludes in general the possibility of sharing costs in a fair way among the agents.

Next, we presented a specific coordination mechanism that was efficient, budget-balanced and strategy-proof. The managerial insight behind these results is that VMI should be interpreted as a *service* provided by the supplier. In turn, the retailer's payment for this service should depend on the accuracy of the information it shares with the supplier. For the practical applicability, the mechanism requires (i) no new mediator party, and (ii) no new information exchange channels, however, it implies that (i) new innovative business models, (ii) efficient local planning at the supplier, and (iii) improved forecasting of market demands are essential.

Negative consequences of Theorem 2 are not entirely discouraging, though. By relaxing assumptions of our model in various ways, a number of open questions emerge, together with new opportunities for the application of mechanism design in supply chain management. In this paper we assumed that the entire demand should be fulfilled. One can study the situation when *lost sales* are allowed, in which case Theorem 2 holds no more. Furthermore, one may also study such mechanisms that are only approximately efficient.

Of course, not only the VMI supply scheme is important; the traditional order-based procurement can also be analysed by means of the mechanism design theory. In such cases negotiation processes—possibly in an automated way between enterprise information systems—could be developed for supporting decentralised decision making. This approach, however, cannot skip considering computational issues any more [Nisan and Ronen, 2001]. We made a first step in this direction in [Egri *et al.*, 2011], where we presented a simple protocol aimed at improving benefits both for a manufacturer and its supplier, without guaranteeing strict optimum.

Finally, a further goal is extending the two-echelon games and handling more complete supply networks. In such cases, achieving global optima through coordination is out of question due to the complexity of the interconnections, the local planning problems to be solved, as well as the dynamic behaviour of the network. Nevertheless, the analytic approach can help estimate the theoretic bounds of the system, and measure the performance gap between results of approximate optimisations and the theoretical global optimum.

Acknowledgements

This work has been supported by the Hungarian Scientific Research Fund OTKA No. T73376 and the National Office for Research and Technology OMFB No. 01638/2009 grants.

References

[Arshinder *et al.*, 2008] Arshinder, A. Kanda, and S.G. Deshmukh. Supply chain coordination: Perspectives, empirical studies and research directions. *International Journal of Production Economics*, 115(2):316–335, 2008.

[Cachon and Netessine, 2004] G.P. Cachon and S. Netessine. Game theory in supply chain analysis. In D. Simchi-Levi, S.D. Wu, and Z-J. Shen, editors, *Handbook of Quantitative Supply Chain Analysis: Modeling in the eBusiness Era*, pages 13–66. Kluwer, 2004.

[Dash *et al.*, 2003] R.K. Dash, N.R. Jennings, and D.C. Parkes. Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, 18:40–47, 2003.

[Egri and Váncza, 2007] P. Egri and J. Váncza. Cooperative production networks – multiagent modeling and planning. *Acta Cybernetica*, 18(2):223–238, 2007.

[Egri *et al.*, 2011] P. Egri, A. Döring, T. Timm, and J. Váncza. Collaborative planning with benefit balancing in Dynamic Supply Loops. *CIRP Journal of Manufacturing Science and Technology*, 2011. In print, doi: 10.1016/j.cirpj.2011.05.002.

[Egri, 2008] P. Egri. *Coordination in Production Networks*. PhD thesis, Eötvös Loránd University, Budapest, 2008.

[Green and Laffont, 1977] J. Green and J-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45(2):427–438, 1977.

[Hennet and Arda, 2008] J-C. Hennet and Y. Arda. Supply chain coordination: A game-theory approach. *Engineering Applications of Artificial Intelligence*, 21:399–405, 2008.

[Ieong *et al.*, 2007] S. Ieong, A. Man-Cho So, and M. Sundararajan. Stochastic mechanism design. In *Proceedings of the 3rd International Conference on Internet and Network Economics*, pages 269–280, 2007.

[Laffont, 2001] D. Laffont, J-J.and Martimort. *The Theory of Incentives : The Principal-Agent Model*. Princeton University Press, 2001.

[Monostori *et al.*, 2006] L. Monostori, J. Váncza, and S.R.T. Kumara. Agent-based systems for manufacturing. *CIRP Annals–Manufacturing Technology*, 55(2):697–720, 2006.

[Nagarajan and Sosic, 2008] M. Nagarajan and G. Sosic. Game-theoretic analysis of cooperation among supply chain agents: Review and extensions. *European Journal of Operational Research*, 187(3):719–745, 2008.

[Nisan and Ronen, 2001] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1-2):166–196, 2001.

[Nisan *et al.*, 2007] N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[Papakonstantinou *et al.*, 2011] A. Papakonstantinou, A. Rogers, E.H. Gerding, and N.R. Jennings. Mechanism design for the truthful elicitation of costly probabilistic estimates in distributed information systems. *Artificial Intelligence*, 175(2):648–672, 2011.

[Rosenschein and Zlotkin, 1994] J.S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, Cambridge, MA, USA, 1994.

[Salanié, 2005] B. Salanié. *The Economics of Contracts: A Primer, 2nd Edition*. The MIT Press, 2005.

[Shneidman and Parkes, 2004] J. Shneidman and D.C. Parkes. Specification faithfulness in networks with rational nodes. In *Proceedings of the 23rd Annual ACM Symposium on Principles of Distributed Computing*, pages 88–97, 2004.

[Shoham and Leyton-Brown, 2008] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2008.

[Simchi-Levi *et al.*, 2000] D. Simchi-Levi, P. Kaminsky, and E. Simchi-Levi. *Designing and Managing the Supply Chain: Concepts, Strategies, and Cases*. McGraw – Hill, New York, 2000.

[van der Krogt *et al.*, 2008] R. van der Krogt, M. de Weerdt, and Y. Zhang. Of mechanism design and multiagent planning. In *Proceeding of the 18th European Conference on Artificial Intelligence (ECAI 2008)*, pages 423–427, 2008.

[Váncza and Márkus, 2000] J. Váncza and A. Márkus. An agent model for incentive-based production scheduling. *Computers in Industry*, 43:173–187, 2000.

[Váncza *et al.*, 2008] J. Váncza, P. Egri, and L. Monostori. A coordination mechanism for rolling horizon planning in supply networks. *CIRP Annals – Manufacturing Technology*, 57(1):455–458, 2008.