

Mitigating the cost of anarchy in supply chain systems ^{*}

H. Edwin Romeijn[†] Wilco van den Heuvel[‡] Joseph Geunes[§]

ECONOMETRIC INSTITUTE REPORT 2012-03

Abstract

In a decentralized two-stage supply chain where a supplier serves a retailer who, in turn, serves end customers, operations decisions based on local incentives often lead to suboptimal system performance. Operating decisions based on local incentives may in such cases lead to a degree of system disorder or anarchy, wherein one party's decisions put the other party and/or the system at a disadvantage. While models and mechanisms for such problem classes have been considered in the literature, little work to date has considered such problems under nonstationary demands and fixed replenishment order costs. This paper models such two-stage problems as a class of Stackelberg games where the supplier announces a set of time-phased ordering costs to the retailer over a discrete time horizon of finite length, and the retailer then creates an order plan, which then serves as the supplier's demand. We provide metrics for characterizing the degree of efficiency (and anarchy) associated with a solution, and provide a set of easily understood and implemented mechanisms that can increase this efficiency and reduce the negative impacts of anarchic decisions.

1 Introduction

1.1 Background and motivation

This paper considers a two-stage supply chain system containing a supplier who serves a retailer with a single product, where costs as well as external demand parameters are nonstationary. Such systems are often not centrally managed by a single decision maker. Rather, the supplier and retailer are separate parties with conflicting objectives. To avoid unnecessary costs (or loss of

^{*}This work was supported in part by the National Science Foundation under grants no. CMMI-0926508 and CMMI-0927930 and by the Netherlands Organization for Scientific Research (NWO).

[†]Department of Industrial and Operations Engineering, The University of Michigan, 1205 Beal Avenue, Ann Arbor, Michigan 48109-2117; e-mail: romeijn@umich.edu.

[‡]Econometric Institute, Erasmus University Rotterdam, Rotterdam, The Netherlands; e-mail: wvandenheuvel@few.eur.nl.

[§]Department of Industrial and Systems Engineering, University of Florida, P.O. Box 116595, Gainesville, Florida 32611-6595; e-mail: geunes@ise.ufl.edu.

potential profits) from the perspective of the entire two-stage system, the decision makers could, in principle, explicitly agree *a priori* on the application of a globally optimal planning policy with an associated agreement on revenue sharing. However, such explicit coordination is often undesirable or impractical (if not impossible) since it is deemed to be too costly, too risky, or too cumbersome. Therefore, such explicit coordination is often absent in practice, and the individual decision makers tend to make decisions that they perceive to be in their best interest, solely based on a local perspective and in response to local incentives. In particular, the supplier will generally quote ordering costs (or prices) to the retailer and make production output decisions, while a retailer must make ordering decisions.

The past fifteen years has witnessed a high volume of activity in supply chain coordination research, with the overwhelming majority of the high-impact results applying to systems with either time-invariant costs and demand parameters, or situations where the time dimension is ignored (for example, similar assumptions to those made for newsvendor or economic order quantity problems are employed). Analysis of supply chain coordination in such systems has led to relatively simple coordination mechanisms and policies that are especially attractive and valuable from a practical point of view. However, due to the large number of degrees of freedom in the nonstationary systems that we will study, coordination mechanisms can quickly become complex and too cumbersome to implement in practice. Therefore, we will focus on the development and analysis of mechanisms that can be easily explained, understood, and implemented. For example, by manipulating per-unit wholesale price, a supplier directly affects the local incentives of retailers. This naturally leads to questions as to how a supplier can use wholesale prices to improve both its own performance and that of the supply chain as a whole. Wholesale price manipulation provides a supplier just one source of flexibility in affecting a retailer's local incentives and, therefore, global system performance. The supply chain coordination literature contains a number of additional contractible parameters that a supply chain member can use to affect its partners incentives and operations policies (e.g., side payments, rebates, holding cost subsidies), each providing a source of flexibility in influencing supply chain performance. We are particularly interested in mechanisms for manipulating anarchy in the supply chain that draw on as few of these sources of flexibility as possible, in order to ensure ease of implementation, while at the same time achieving desirable supply chain performance.

We will consider a setting in which (i) the retailer faces an economic lot-sizing problem to find a

sequence of order quantities to minimize the total cost required to satisfy all external demands; and (ii) the supplier faces an economic lot-sizing problem to find a sequence of procurement quantities that minimizes the total cost required to fulfill the retailer's orders. Assuming that the supplier has some wholesale pricing flexibility, the supplier can essentially set some of the cost parameters that impact the ordering decisions of the retailer; these decisions then, in turn, affect the procurement decisions by the supplier. The supplier should therefore consider this dynamic when taking its initial action. In this case, a Stackelberg game arises where the supplier is the Stackelberg leader and the retailer is the follower. A solution to such a game is characterized by the well-known concept of a Nash Equilibrium, wherein no agent can improve its position by choosing to unilaterally deviate from the equilibrium solution. Assuming that the supplier has complete information on the retailer's decision problem, the supplier attempts to solve the game using backward induction. That is, the supplier may characterize the retailer's optimal response to its initial action and then optimize this initial action. The resulting problem falls in the general class of so-called *bilevel programming problems*. Due to the presence of an economic lot-sizing cost structure at both the supplier and the retailer level, we in fact obtain a mixed-integer bilevel optimization problem. Since bilevel linear optimization has been shown to be NP-Hard (Ben-Ayed and Blair [2]), mixed-integer bilevel programming is, in general, NP-Hard as well. However, due to the attractiveness of efficiently implementable simple coordination mechanisms, we will consider practical classes of so-called *leader action-space restrictions*, i.e., subsets of the leader's action space. We will show that many of these mechanisms will make identifying a Nash Equilibrium solution for the Stackelberg leader-follower game tractable, with the problem often even solvable in polynomial time.

1.2 Literature review

Although the concepts of double marginalization and competition for channel profits are quite old (Cournot [9]), the work of Monahan [19] and Lal and Staelin [16] appear to be some of the early works to consider differing supplier-buyer incentives in a supply chain operations context, with Monahan [19] focusing on the buyer's perspective, and Lal and Staelin [16] justifying the use of quantity discounts by considering both buyer and supplier costs. Lee and Rosenblatt [18] subsequently generalized Monahan's model to violate the lot-for-lot ordering assumption, while Weng [27] extended the analysis to multiple homogenous buyers and price-dependent demand. Corbett

and de Groote [8] then considered this problem class under information asymmetry, while Klastorin, Moinzadeh, and Son [15] used price discounts to provide incentives for customers to coordinate order timing with the supplier. This latter approach is similar in spirit to the model of Viswanathan and Piplani [24], who applied price incentives to entice retailers to order in common replenishment periods. Chen, Federgruen, and Zheng [6, 7] considered multiple non-identical retailers with price-dependent demands, and showed how the system-optimal solution can be achieved under decentralized management using a discount scheme based on annual volume, order quantity, and order frequency. This stream of literature focuses on the tension between supplier and buyer preferences under the standard EOQ cost and demand assumptions. Beyond the multi-echelon work that employs standard EOQ assumptions, little work exists considering supply chain coordination under dynamic demands and costs, which serves as one of the focus areas of our work. An exception is the recent work of Sahin and Robinson [22], who considered the impacts of varying degrees of explicit cooperation and information sharing among partners in a make-to-order supply chain. Our work is fundamentally different from this, as we consider settings in which explicit coordination is infeasible, and where mechanisms are required to influence anarchy in the system.

While we have discussed supply chain coordination models under deterministic demand assumptions, a substantial segment of the literature in this area considers uncertain demand where, in addition to marginalization issues, supply chain partners also face risk allocation issues. Just as decentralized models under deterministic demand are hampered by failing to consider overall supply chain margin, decentralized models under stochastic demand do not consider the overall supply chain risk picture (see Cachon [4] for a discussion of risk allocation in supply chains). A great deal of the work in this area uses the one-period newsvendor setting to demonstrate the need for and benefits of coordination mechanisms under uncertain demand, and proposes a number of contract mechanisms to achieve coordination. Cachon [3] provides an in-depth literature survey on the use of contract mechanisms to coordinate multiple stages in a newsvendor setting, discussing past work on buy-back contracts, quantity flexibility contracts, and revenue sharing contracts, among others. Toptal and Çetinkaya [23] extend this analysis to explicitly incorporate the impacts of realistic transportation cost structures. While a number of past works (e.g., Lariviere and Porteus [17], Cachon and Lariviere [5]) have illustrated the inability of price-only contracts to lead to coordinated supply chain operations in decentralized systems, Perakis and Roels [20]

recently quantified the loss of efficiency, or price of anarchy, of price-only contracts in decentralized supply chains in a single-period setting. The metrics we will define for measuring supply chain efficiency (and the price of anarchy) are similar in spirit to theirs. However, our focus will expand beyond the newsvendor context, and we will introduce coordination mechanisms that can be used to manipulate the anarchy in complex systems.

2 Anarchy and efficiency

Following recent work in transportation systems, we will refer to the cost resulting from individual decision makers following locally optimal decisions in decentralized systems as the *cost of anarchy* in the supply chain (for discussions on the cost, or price, of anarchy in transportation and supply chain systems, please see Roughgarden [21] and Perakis and Roels [20], respectively). We will denote the *centralized system optimal cost* in the supply chain by C^* and the *total cost under anarchy* incurred when individual decision makers make locally optimal decisions by C^A . In general, we expect to have $C^* < C^A$, so that anarchy is costly in terms of the entire supply chain, and we define the cost of anarchy as follows:

Definition 2.1. *The cost of anarchy in a supply chain system is defined as the difference between the optimal cost under anarchy and the system optimal cost:*

$$\text{cost of anarchy} = C^A - C^*.$$

When the cost of anarchy is large, supply chain partners may be able to provide local incentives to improve the efficiency of the supply chain. We will refer to the application of such incentives in order to change local decision making as *manipulated anarchy*. Such levers can be viewed as somewhat analogous to using toll pricing to manipulate anarchy in transportation systems. Denoting the *total cost under manipulated anarchy* by C^M , we should then have $C^* \leq C^M \leq C^A$. We will define the *efficiency* of a coordination mechanism as the fraction of the cost of anarchy that it avoids:

Definition 2.2. *The efficiency of a supply chain system under a particular coordination mechanism*

is defined to be the fraction by which cost of anarchy can be reduced:

$$\text{coordination mechanism efficiency} = \frac{C^A - C^M}{C^A - C^*}.$$

Finally, we will define the *potential* of a coordination mechanism as the fraction of the cost of anarchy that is avoided if a central decision maker were to optimize the sum of the local objectives of both parties (using the mechanism), leading to a supply chain cost of C^P satisfying

$$C^* \leq C^P \leq C^M \leq C^A.$$

Definition 2.3. *The efficiency potential of a supply chain system under a particular coordination mechanism is defined as:*

$$\text{coordination mechanism potential} = \frac{C^A - C^P}{C^A - C^*}.$$

Of course achieving a cost of C^P would require some limited coordination which, perhaps even to a greater extent than full coordination, is unlikely to be encountered in practical situations. However, the coordination mechanism potential does provide insight into the extent to which a coordination mechanism is able to reach its full potential without actually requiring some form of coordination.

In the remainder of this paper, we are particularly interested in mechanisms for manipulating anarchy in the supply chain that are easy to implement while at the same time achieving a desirable supply chain performance.

3 Two-stage supply chain system

Recall that the two-stage serial supply chain system that we consider in this paper consists of a supplier and a retailer, where the retailer faces a sequence of deterministic and dynamic demands for a single product over a finite planning horizon. Each player in the system wishes to minimize their cost over the planning horizon. We assume for convenience that both the supplier and the retailer have unlimited procurement capacities in all order periods.

3.1 Preliminaries

3.1.1 Notation and solution structure

We denote the sequence of external demands by d_t ($t = 1, \dots, T$), where T is the length of the planning horizon. We assume that the structure of procurement and inventory holding costs for both parties are such that, in both centralized and decentralized systems, we can restrict ourselves to solutions that satisfy the zero-inventory ordering (ZIO) property at both the supplier and retailer level. In other words, it suffices to consider binary decision variables y_t^s and y_t^r ($t = 1, \dots, T$), where these variables take the value 1 in a setup or order period and 0 in the other periods. We will denote the corresponding binary vectors by y^s and y^r . Since we will focus on coordination mechanisms between the two parties, we will pay particular attention to the costs at the retailer level. For convenience, we will assume that the cost functions at the retailer can be characterized by a finite parameter vector v , the elements of which can be (to some extent to be characterized later) influenced or set by the supplier.

Although some of our results hold in the very general setting described above, most of our concrete algorithmic contributions apply to the following cost structure at the retailer level. In each period t ($t = 1, \dots, T$), the retailer incurs fixed ordering costs F_t^r if it places an order with the supplier in period t , and variable net production and inventory costs $g_{t\tau}^r$ per unit produced in period t and used to satisfy external demand in period τ . We use the following general form to represent the variable net production and inventory costs:

$$g_{t\tau}^r = p_t^s + a_t^r + \sum_{j=t}^{\tau-1} h_j^r + (p_t^s + a_t^r) \sum_{j=t}^{\tau-1} \rho_j^r, \quad \tau = t + 1, \dots, T; t = 1, \dots, T,$$

where p_t^s denotes the unit price charged by the supplier, a_t^r denotes a unit transportation and handling cost, h_t^r denotes the holding cost per unit in inventory at the end of the period, and ρ_t^r denotes the cost rate of capital in period t ($t = 1, \dots, T$). If $\rho_t^r = 0$ ($t = 1, \dots, T$), then the cost structure is equal to the standard form often used in economic lot-sizing models, which we refer to as Cost Structure I. Alternatively, if $h_t^r = 0$ ($t = 1, \dots, T$), then holding costs are viewed as an opportunity cost of the capital invested in inventory, and we refer to this structure as Cost Structure II. Whenever convenient, we use $\rho_{t\tau}^r$ as a shorthand notation for $\sum_{j=t}^{\tau-1} \rho_j^r$. Furthermore, in case the cost rate of capital is modeled in a compounded way, we should replace $\sum_{j=t}^{\tau-1} \rho_j^r$ by

$\prod_{j=t}^{\tau-1} (1 + \rho_j) - 1$. In any case, the vector v consists of elements F_t^r and $g_{t\tau}^r$ (for $\tau = t + 1, \dots, T$ and $t = 1, \dots, T$).

3.1.2 Centralized decision making

Under centralized decision making, the decision maker should seek the joint supplier-retailer plan that minimizes total system costs. Zangwill [28] presents a dynamic programming approach to solve this two-echelon lot-sizing problem, which was shown by Van Hoesel et al. [14] to have a running time of $\mathcal{O}(T^3)$. We will denote the centralized cost function by $C(y^s, y^r)$, noting that the optimal solution value of this problem equals C^* , the centralized system optimal cost.

3.1.3 Decentralized decision making

Under decentralized decision making (i.e., anarchy), the retailer starts by optimizing its procurement plan by solving a standard single-level dynamic production planning problem. Subsequently, the supplier solves a similar problem, where the supplier's observed demand pattern is given by the retailer order quantities (or, equivalently, by the retailer order periods). The supplier and retailer cost functions can be written as $C^s(y^s|y^r)$ and $C^r(y^r)$. Note that the total system costs satisfy

$$C(y^s, y^r) = C^s(y^s|y^r) + C^r(y^r).$$

It is well-known that each supply chain partner's optimization problem can be solved in $\mathcal{O}(T^2)$ time (see, e.g., Wagner and Whitin [26]). Under additional assumptions on the cost structures these problems can even be solved in $\mathcal{O}(T \log T)$ or $\mathcal{O}(T)$ time (see Aggarwal and Park [1], Federgruen and Tzur [11], and Wagelmans et al. [25]). The sum of the optimal solution values of these problems is C^A , the total cost under anarchy.

3.2 Stackelberg game to improve supplier profit

In order to mitigate the cost of anarchy, the supplier can choose to manipulate the retailer's cost coefficients in an attempt to reduce the total system cost by reducing (or, at least, not increasing) either party's cost. In particular, the supplier may offer any combination of wholesale price discounts, fixed and variable cost rebates, and holding cost subsidies (in the remainder, we will often refer to any such offering from the supplier to the retailer as a *promotion* or *incentive*). This

then leads to a Stackelberg game in which the supplier plays the role of Stackelberg leader and the retailer becomes the Stackelberg follower. Clearly, the supplier must ensure that its actions do not make the retailer worse off than under the current prices, since otherwise the retailer may choose another supply source (thus we assume some prevailing market wholesale prices exist that serve as the planned p_t^s values for $t = 1, \dots, T$). The objective is then to determine the incentives the supplier should offer in order to achieve minimum cost, while ensuring that the retailer is at least as well off as under the current prevailing market wholesale prices.

More formally, let θ denote a parameter (or parameter vector) characterizing the supplier's manipulation of the cost coefficients v (incentives), and Θ the set of feasible incentives. We will discuss this set in more detail later, but for now it suffices to say that this set will be characterized by (i) concavity of the retailer's cost function; (ii) the retailer being no worse off under the incentive; and (iii) convenience or intuitive nature. The retailer will now incur costs characterized by $v(\theta)$. For convenience, we will let $\theta = \mathbf{0}$ correspond to the case of unchanged parameter values, i.e., $v(\mathbf{0}) = v$, where $\mathbf{0}$ is a vector of all zeroes of the same dimension as θ .

This framework provides for a potentially very large degree of flexibility in the supplier's action space, depending on how the cost functions are parameterized. However, we will pay particular attention to far more realistic settings where, for clarity, convenience, and ease of communication, the supplier's flexibility is reduced. For example, it seems reasonable to expect the supplier to be able to provide rebates to the retailer in order periods to help defray the retailer's fixed order costs, or to be able to offer a discounted price in some or all periods. In many cases, it would seem reasonable to restrict attention to schemes in which the supplier focuses on cost parameters of a single type, and/or on highly structured rebates, such as *stationary* absolute or relative rebates of, for example, fixed ordering costs or the purchase price. Such structured rebates would be characterized by a low dimension of the set Θ .

We next discuss the follower and leader problems in the Stackelberg game, as well as the set of feasible incentives and mechanism efficiency in more detail.

3.2.1 Retailer problem

As in the case of independent decision making under the prevailing market (wholesale) prices, the retailer independently determines an optimal order schedule, but now based on the new cost

parameters. We will denote the corresponding cost function by $C^r(y^r|\theta)$, and the optimal value function to the retailer problem as a function of θ by

$$\hat{C}^r(\theta) \equiv \min_{y^r \in \{0,1\}^T} C^r(y^r|\theta).$$

Let $\Theta(y^r)$ denote the set of incentives that lead the retailer to select plan y^r , i.e.,

$$\hat{C}^r(\theta) = C^r(y^r|\theta) \quad \text{if } \theta \in \Theta(y^r). \quad (1)$$

Of course each set of incentives yields a retailer order plan so that, by definition,

$$\bigcup_{y^r \in \{0,1\}^T} \Theta(y^r) = \Theta.$$

However, property (1) alone does not guarantee that the sets $\Theta(y^r)$ form a partition of Θ , i.e., that $\Theta(y^r) \cap \Theta(\bar{y}^r) = \emptyset$ whenever $y^r \neq \bar{y}^r$, since the retailer problem may have alternate optimal solutions. If we assume that the retailer makes the choice that is *least* costly for the supplier we can simply redefine $\Theta(y^r)$ to denote the set of incentives for which y^r is an optimal retailer order plan, i.e., through equation (1). Otherwise, we should assume that we have a partition that implicitly characterizes which order plan the retailer selects in case of multiple optima. A conservative choice would assume that the retailer makes the choice that is *most* costly for the supplier. We will consider both of these alternatives in more detail later.

3.2.2 Supplier problem

The supplier problem is substantially more challenging to solve, as is often the case for the leader's problem in such games. Given a vector of incentives $\theta \in \Theta$ and the resulting retailer order plan y^r , the supplier wishes to optimize its total costs $C^s(y^s|y^r)$. However, the supplier is also responsible for covering the difference between the *true* retailer cost and the cost charged to the retailer under the incentives. This means that the optimal supplier cost as a function of $\theta \in \Theta$ is given by

$$\begin{aligned} \hat{C}^s(\theta) &\equiv \min_{y^s \in \{0,1\}^T} C^s(y^s|y^r) + [C^r(y^r) - \hat{C}^r(\theta)] && \theta \in \Theta(y^r), y^r \in \{0,1\}^T \\ &= \min_{y^s \in \{0,1\}^T} C(y^s, y^r) - \hat{C}^r(\theta) && \theta \in \Theta(y^r), y^r \in \{0,1\}^T. \end{aligned}$$

The supplier then chooses the incentives that minimize its cost:

$$\inf_{\theta \in \Theta} \hat{C}^s(\theta). \quad (2)$$

(Note that, at this point, we do not know whether this problem has an optimal solution, hence the use of “inf” instead of “min”.) This problem can be reformulated as

$$\begin{aligned} \inf_{\theta \in \Theta} \hat{C}^s(\theta) &= \inf_{\theta \in \Theta} \left[\min_{y^s \in \{0,1\}^T} C(y^s, y^r) - \hat{C}^r(\theta) \right] = \min_{y^r \in \{0,1\}^T} \inf_{\theta \in \Theta(y^r)} \left[\min_{y^s \in \{0,1\}^T} C(y^s, y^r) - \hat{C}^r(\theta) \right] \\ &= \min_{y^r \in \{0,1\}^T} \left[\min_{y^s \in \{0,1\}^T} C(y^s, y^r) - \sup_{\theta \in \Theta(y^r)} \hat{C}^r(\theta) \right] \\ &= \min_{y^r \in \{0,1\}^T} \left[\min_{y^s \in \{0,1\}^T} C(y^s, y^r) - \sup_{\theta \in \Theta(y^r)} C^r(y^r | \theta) \right]. \end{aligned} \quad (3)$$

In principle, we could solve this problem by enumerating all 2^T potential retailer order plans $y^r \in \{0, 1\}^T$ and solve the following problem for each of these:

$$\min_{y^s \in \{0,1\}^T} C(y^s, y^r) - \sup_{\theta \in \Theta(y^r)} C^r(y^r | \theta). \quad (4)$$

The first term in equation (4) can be evaluated relatively easily for any given y^r by solving a single-level lot-sizing problem. The second term is, in general, much more challenging. In Section 3.3 we will study this problem in more detail for practical special cases.

3.2.3 Feasible incentives

Let Θ_0 denote the set of incentives that ensures that the retailer’s cost function is concave and Θ_1 the set of incentives that ensures that the retailer is no worse off under the incentives than under the market prices. Furthermore, let Θ_2 be some set of incentives that is convenient, intuitive, and easy to justify to the retailer; for example absolute cost discounts that are nonnegative, or relative rebates that are between 0 and 1. Then $\Theta = \Theta_0 \cap \Theta_1 \cap \Theta_2$ and clearly we have that

$$\Theta_1 = \left\{ \theta : \hat{C}^r(\theta) \leq \hat{C}^r(\mathbf{0}) \right\}. \quad (5)$$

Observe that the supplier will only make such an offer to the retailer (and choose to initiate the Stackelberg game) if the supplier is better off as a result of the supplier's optimal offer, which is guaranteed if $\mathbf{0} \in \Theta$. It is easy to see that $\mathbf{0} \in \Theta_0 \cap \Theta_1$, and it seems reasonable to assume that $\mathbf{0} \in \Theta_2$ as well, which means that any feasible offer by the supplier will improve (or, at least, not deteriorate) the system performance.

3.2.4 Mechanism efficiency

For a given incentive vector $\theta \in \Theta$, we can define the corresponding total system costs as

$$\hat{C}(\theta) = \hat{C}^s(\theta) + \hat{C}^r(\theta). \quad (6)$$

The system costs under optimal manipulated anarchy costs are then given by

$$C^M(\Theta) = \inf_{\theta \in \arg \min_{\theta' \in \Theta} \hat{C}^s(\theta')} \hat{C}(\theta).$$

Moreover, we can express the optimal system costs under (non-manipulated) anarchy as

$$C^A = C^M(\{\mathbf{0}\}) = \hat{C}(\mathbf{0}).$$

Finally, the mechanism cost under limited coordination, which is used to evaluate mechanism potential, is equal to

$$C^P = \inf_{\theta \in \Theta} \hat{C}(\theta).$$

3.3 Lot-sizing costs and linear incentive mechanisms

Considering the specific cost structures from Section 3.1.1, we note that many convenient incentive mechanisms, such as absolute or relative cost rebates or subsidies, can be represented linearly as $v(\theta) = v + V\theta$, where $V \neq 0$ is a matrix of appropriate dimensions. (Note that, as desired, we have $v(\mathbf{0}) = v$.) For convenience, let $x_{t\tau}^r$ denote the quantity ordered by the retailer in period t for satisfying demand in period τ using order plan y^r , and let x^r be the vector of these elements (note that y^r uniquely determines x^r). It is then easy to see that

$$C^r(y^r|\theta) = (y^r, x^r)^\top v(\theta) = (y^r, x^r)^\top (v + V\theta) = C^r(y^r) + (y^r, x^r)^\top V\theta, \quad \theta \in \Theta(y^r),$$

i.e., $C^r(y^r|\theta)$ is linear in $\theta \in \Theta(y^r)$, and thus $\hat{C}^r(\theta)$ is piecewise-linear in $\theta \in \Theta$. In fact, it now follows that

$$\hat{C}^r(\theta) = \min_{y^r \in \{0,1\}^T} \left(C^r(y^r) + (y^r, x^r)^\top V \theta \right), \quad \theta \in \Theta, \quad (7)$$

so that \hat{C}^r is piecewise-linear and concave in θ . In fact, under the assumption that, in case of multiple optima, the retailer selects the order plan that is least costly to the supplier, the sets $\Theta(y^r)$ are closed polyhedra (where the values of θ for which two such polyhedra intersect correspond to cases where the retailer solution is nonunique). This has the attractive computational side-effect that the supplier's optimum is attained, so that we can replace “inf” and “sup” by “min” and “max” in equations (2), (3), and (4).

Now a potentially tractable solution approach to the supplier problem can be developed if we can efficiently identify the closed polyhedra representing \hat{C}^r . Recall that the application of action-space restrictions can often be of significant practical value. Therefore, we will in the remainder of this section study action spaces of dimension one and two. In these cases, a significant reduction in the number of retailer order plans that need to be considered can be achieved, sometimes even leading to polynomial-time solvability of the supplier problem. Section 4 will then discuss and study examples of specific incentive mechanisms.

3.4 Single-parameter case

We first consider the case where θ is a scalar parameter (so that both v and V are vectors of the same dimension as (y^r, x^r)). By equation (3), the supplier's problem is to minimize a piecewise-linear function that is obtained by subtracting a piecewise-linear concave function (\hat{C}^r) from a step-function whose value only depends on the appropriate retailer order plan y^r . The piecewise-linearity of \hat{C}^r means that we can limit consideration to a finite set of incentives, say $\hat{\Theta}$, to be determined later. The heart of our approach to the supplier's problem will therefore be to (i) obtain an explicit characterization of the function \hat{C}^r , (ii) identify the appropriate retailer order plan for each candidate incentive $\theta \in \hat{\Theta}$, and (iii) identify the optimal incentive. We will use the well-known method usually attributed to Eisner and Severance [10] to achieve (i), a slight modification of a standard dynamic programming approach to the economic lot-sizing problem for (ii), and simple enumeration for (iii).

3.4.1 Characterizing the function \hat{C}^r

We will focus on the (common) situation where $\Theta_0 \cap \Theta_2 = [\theta^L, \theta^U]$ (where $\theta^L \leq 0 \leq \theta^U$). The following method, due to Eisner and Severance [10], can then be used to construct the function \hat{C}^r on this set as follows:

0. Find an optimal retailer order plan corresponding to both θ^L and θ^U ; each of these yields a line, and the lower envelope of these yields an upper bound on \hat{C}^r on Θ .
1. Let $\bar{\theta}$ be a breakpoint of the current upper bound on \hat{C}^r that has not been investigated yet. If such a breakpoint does not exist: stop; the current upper bound is \hat{C}^r .
2. Find an optimal retailer order plan under incentives $\bar{\theta}$ and update the upper bound on \hat{C}^r . Return to Step 1.

Note that if $\Theta_0 \cap \Theta_2 = [\theta^L, \infty)$ then in Step 0 the optimal retailer plan corresponding to θ^U is replaced by the one that optimizes the retailer problem with cost vector V^+ defined as

$$V^+ = \lim_{\theta \rightarrow \infty} \frac{v(\theta)}{\theta}.$$

Similarly, if $\Theta_0 \cap \Theta_2 = (-\infty, \theta^U]$ then in Step 0 the optimal retailer plan corresponding to θ^L is replaced by the one that optimizes the retailer problem with cost vector V^- defined as

$$V^- = \lim_{\theta \rightarrow -\infty} \frac{v(\theta)}{\theta}.$$

(And of course both in the uncommon situation that $\Theta_0 = \Theta_2 = \mathbb{R}$.) Eisner and Severance [10] show that the running time of this algorithm is $\mathcal{O}(BR)$ time, where B is the number of segments of \hat{C}^r on $\Theta_0 \cap \Theta_2$ and R is the running time required to solve the retailer problem for a fixed value of θ . Finally, let $\hat{\Theta}$ be the set of breakpoints of \hat{C}^r on Θ (including θ^L and θ^U provided these are in Θ_1), supplemented by $\theta = 0$ and the nonzero value of θ satisfying $\hat{C}^r(\theta) = \hat{C}^r(0)$ (provided such a value exists). (Note that if there are more than two solutions to this equation, then $\Theta_1 = \mathbb{R}$ and no supplemental breakpoints are needed.)

3.4.2 Identifying the optimal incentive

The supplier's cost function \hat{C}^s is a piecewise linear function on Θ in which discontinuities and/or slope changes occur only at values of $\theta \in \hat{\Theta}$. If we can assume that, in case of alternative optima, the retailer chooses the order plan that is least costly for the supplier, we also have that \hat{C}^s is lower semi-continuous, which implies that the minimum of \hat{C}^s is attained for some $\theta \in \hat{\Theta}$. Alternatively, if the supplier does not know which alternative optimal solution the retailer will choose and hence takes a worst-case approach by assuming that the retailer will choose the order plan that is most costly for the supplier, then \hat{C}^s is upper semi-continuous, which implies that the minimum of \hat{C}^s is either attained for some $\theta \in \hat{\Theta}$ or it is not attained and the supplier will choose a value arbitrarily close to an element of $\hat{\Theta}$. Let us therefore limit our attention to this finite set of values. As the retailer's problem for any fixed θ is an economic lot-sizing problem (which can be solved by dynamic programming, i.e., as a shortest path problem in an acyclic network), we can easily identify *all* retailer order plans that are optimal for any $\theta \in \hat{\Theta}$. Let $Y^r(\theta)$ denote the set of optimal retailer order plans for $\theta \in \Theta$. For a given $\theta \in \hat{\Theta}$ and an optimal retailer order plan $\hat{y}^r \in Y^r(\theta)$, the supplier then solves an economic lot-sizing problem as well with demands implied by the retailer order plan. Again assuming that the retailer chooses the order plan that is best for the supplier we have

$$\hat{C}_1^s(\theta) = \min_{\hat{y}^r \in Y^r(\theta)} \left\{ \min_{y^s \in \{0,1\}^T} C(y^s, \hat{y}^r) - \hat{C}^r(\theta) \right\} \quad \theta \in \Theta, \quad (8)$$

so that an optimal incentive is $\theta^* \in \arg \min_{\theta \in \hat{\Theta}} \hat{C}_1^s(\theta)$ and $\min_{\theta \in \Theta} \hat{C}_1^s(\theta) = \hat{C}_1^s(\theta^*)$. If, on the other hand, the supplier does not know which alternative optimal solution the retailer will choose and takes a worst-case approach we have

$$\hat{C}_2^s(\theta) = \max_{\hat{y}^r \in Y^r(\theta)} \left\{ \min_{y^s \in \{0,1\}^T} C(y^s, \hat{y}^r) - \hat{C}^r(\theta) \right\} \quad \theta \in \Theta,$$

so that

$$\inf_{\theta \in \Theta} \hat{C}_2^s(\theta) = \min_{\theta \in \hat{\Theta}} \min \left\{ \lim_{\theta' \uparrow \theta} \hat{C}_2^s(\theta'), \lim_{\theta' \downarrow \theta} \hat{C}_2^s(\theta') \right\}. \quad (9)$$

Then the supplier will wish to choose a value of θ arbitrarily close to an element $\theta^* \in \hat{\Theta}$ for which the outer minimum in the right-hand-side of equation (9) is attained, in the direction indicated by

which of the two expressions attains the inner minimum. For convenience, albeit with a slight abuse of notation in the latter case, let $\hat{C}^s(\theta^*) = \hat{C}_1^s(\theta^*)$ or $\hat{C}^s(\theta^*) = \min \left\{ \lim_{\theta \uparrow \theta^*} \hat{C}_2^s(\theta), \lim_{\theta \downarrow \theta^*} \hat{C}_2^s(\theta) \right\}$, depending on the assumptions about the retailer behavior.

Finally, this typically means that the system costs under optimal manipulated anarchy are equal (or very close) to

$$C^M(\Theta) = \hat{C}^s(\theta^*) + \hat{C}^r(\theta^*).$$

In case θ^* is not unique we can refine this by assuming that the supplier uses total system cost (and hence retailer costs) as its secondary objective.

3.4.3 Running time

It follows immediately from the discussion in Section 3.4.1 that the running time required to solve the supplier problem depends on the number of breakpoints of the continuous concave piecewise linear function \hat{C}^r . This function is obtained by solving a (1-dimensional) parametric economic lot-sizing problem where the cost parameters are linear functions of the parameter θ . It is well-known that a T -period economic lot-sizing problem can be solved as a shortest path problem on a complete acyclic directed graph with $T+1$ nodes. Therefore, it suffices to analyze a parametric shortest path problem in order to bound the number of line pieces and hence the number of breakpoints of \hat{C}^r .

Theorem 3.1. *The optimal value function of a one-dimensional parametric shortest path problem on a complete acyclic directed graph with $n \in \mathbb{N}$ nodes is no more than $\frac{1}{2}(2n-1)^{\log_2(2n-1)}$.*

Proof. Our analysis is based on Gusfield [13], who analyzes a parametric shortest path problem on general graphs. We start the analysis by considering a one-dimensional parametric shortest path problem on a complete acyclic directed graph with $n \in \mathbb{N}$ nodes. Let ℓ_n be the maximum number of segments of the optimal value function of such a problem with n nodes. First, we only consider values of n that are powers of 2, i.e., $n = 2^k$ nodes for some $k \in \mathbb{N}$. We will derive an upper bound on ℓ_{2n} by considering the concave piecewise linear function corresponding to all paths using arc (i, j) with $1 \leq i \leq n$ and $n+1 \leq j \leq 2n$ (see Figure 1). Since the sum of 2 continuous piecewise linear functions with p and q segments is a continuous piecewise linear function with at most $p+q-1$ segments, this function consists of at most $\ell_i + \ell_{2n-j+1} - 1$ segments. Moreover,

since every path uses some arc (i, j) , it follows that

$$\ell_{2n} \leq \sum_{i=1}^n \sum_{j=n+1}^{2n} (\ell_i + \ell_{2n-j+1} - 1) \leq \sum_{i=1}^n \sum_{j=n+1}^{2n} 2\ell_n = 2n^2\ell_n.$$

Furthermore, taking $n = 2^k$ (so $k = \log_2 n$), using $\ell_2 = 1$, and applying induction on k we obtain

$$\ell_n = \ell_{2^k} \leq \prod_{i=1}^{k-1} 2(2^i)^2 = \prod_{i=1}^{k-1} 2^{2i+1} = 2^{\sum_{i=1}^{k-1} (2i+1)} = 2^{k^2-1} = \frac{1}{2}2^{k^2} = \frac{1}{2}2^{\log_2^2 n} = \frac{1}{2}n^{\log_2 n}.$$




Figure 1: Analysis of the parametric shortest path problem

Finally, consider a graph with an arbitrary number of nodes n . In that case, let $k = \lceil \log_2 n \rceil$ and $m = 2^k$, so that $m \leq 2n - 1$. Clearly, a complete acyclic directed graph with n nodes is embedded in a complete acyclic directed graph with m nodes and hence the latter graph can be used to determine an upper bound on the number of segments. Using the previous results, it follows that the number of line segments corresponding to the parametric shortest path problem on a graph with n nodes is bounded by $\frac{1}{2}(2n - 1)^{\log_2(2n-1)}$. \square

(Note that this bound is a slight improvement on the one derived by Gusfield [13], which is equal to $\frac{1}{2}(2n + 1)^{1+\log_2(n)}$.) Applying the result of Theorem 3.1 to the parametric lot-sizing problem yields the following corollary:

Corollary 3.2. *The function \hat{C}^r consists of at most $\frac{1}{2}(2T + 1)^{\log_2(2T+1)}$ linear segments.*

Note that the bound from Corollary 3.2 grows subexponentially in T , that is, while it grows faster than any polynomial function it grows more slowly than any exponential function in T . Since the number of breakpoints and line pieces are of the same order, the supplier problem can be solved in subexponential time as well. This is summarized in the following corollary.

Corollary 3.3. *The running time required to solve the supplier problem under a linear single-parameter incentive mechanism is $\mathcal{O}((2T + 1)^{\log_2(2T+1)}T^2)$.*

The term T^2 represents the running time required solve a lot-sizing problem with the general cost structure presented in Section 3.1.1. In Section 4 we will show that the supplier problem can be solved in polynomial time under some practical coordination mechanisms.

3.5 Two-parameter case

We next consider the case where θ is a two-dimensional parameter vector, i.e., $\theta = (\theta_1, \theta_2)$. The supplier problem can be solved in a similar way as outlined for the single-parameter case in Section 3.4. The main difference is in the characterization of the function \hat{C}^r . Note that, in the two-dimensional case, \hat{C}^r is the lower envelope of a finite number of planes. Hence, \hat{C}^r is a piecewise linear and concave function in two dimensions and so it is represented by faces, edges (the intersection of two planes), and vertices (the intersection of three or more planes). Similar to the single-parameter case, it is not difficult to see that the optimal solution to the supplier problem is found at or near one of the vertices of \hat{C}^r . Steps (ii) and (iii) remain unchanged except that the term ‘breakpoint’ needs to be replaced by ‘vertex’. Therefore, we will only focus on determining the function \hat{C}^r in (i).

Let us assume that, as will often be the case, $\Theta_0 \cap \Theta_2 = \{(\theta_1, \theta_2) : \theta_1^L \leq \theta_1 \leq \theta_1^U, \theta_2^L \leq \theta_2 \leq \theta_2^U\}$. The method of Fernández-Baca and Srinivasan [12] can be used to find the vertices of \hat{C}^r . In fact, they show how to partition $\Theta_0 \cap \Theta_2$ into convex polygonal regions such that every region corresponds to a face of \hat{C}^r . The main steps of the procedure are as follows:

0. Find an optimal retailer order plan corresponding to (θ_1^L, θ_2^L) , (θ_1^L, θ_2^U) , (θ_1^U, θ_2^L) and (θ_1^U, θ_2^U) ; each of these yields a plane, and the minimum of these yields an upper bound on \hat{C}^r .
1. Let $(\bar{\theta}_1, \bar{\theta}_2)$ be a vertex of the current upper bound on \hat{C}^r that has not been investigated yet. If such a vertex does not exist: stop, and the current upper bound is in fact \hat{C}^r .
2. Find an optimal retailer order plan under $(\bar{\theta}_1, \bar{\theta}_2)$ and update the upper bound on \hat{C}^r , i.e., find the intersection of the new plane and the current upper bound. Return to Step 1.

As shown by Fernández-Baca and Srinivasan [12], the running time of this algorithm is $\mathcal{O}(FR + F^2)$ with F the number of faces of \hat{C}^r and R the running time to solve the lot-sizing problem for a

fixed value of θ . It remains to intersect $\Theta_0 \cap \Theta_2$ with Θ_1 which, through the application of the same algorithm, leads to both the addition and removal of a set of breakpoints. Then let $\hat{\Theta}$ be the set of incentives corresponding to the vertices of \hat{C}^r on Θ (including (θ_1^L, θ_2^L) , (θ_1^L, θ_2^U) , (θ_1^U, θ_2^L) and (θ_1^U, θ_2^U) provided these are in Θ_1). In order to determine the running time to solve the supplier problem, we need an upper bound on the number of vertices. By using Euler's formula $F + V = E + 1$ with V and E the number of vertices and edges of \hat{C}^r , respectively, Fernández-Baca and Srinivasan [12] argue that $E = \mathcal{O}(F)$. In turn, this implies that $V = \mathcal{O}(F)$. This leads to the following theorem:

Theorem 3.4. *The running time required to solve the supplier's problem under a linear 2-parameter mechanism is $\mathcal{O}(FT^2 + F^2)$ with F the number of faces of \hat{C}^r .*

The final step of the algorithm is to determine the optimal incentive. This is straightforward when we can assume that, in case of multiple retailer optima, the retailer chooses the order plan that is best for the supplier. As in the case of a single parameter, if the supplier must assume the worst-case behavior by the retailer, the optimal incentive is not necessarily attained, but a value with cost that is close to optimal can be found analogously to the procedure outlined in Section 3.4.2.

4 Practical coordination mechanisms

This section describes a set of potential single-parameter mechanisms that may be easily implemented in practice, and which allow characterizing optimal mechanism parameter values. These mechanisms apply discounts or rebates to various cost types, for the two cost structures discussed in Section 3.1.1. We first consider mechanisms that apply for all periods in a planning horizon and later discuss the implications of limiting these mechanisms to particular time windows, or subsets of the planning horizon.

4.1 Constant discounts and rebates over the full planning horizon

It will be convenient to define the aggregate demand in a sequence of consecutive periods: $d_{t\tau} = \sum_{s=t}^{\tau} d_s$. Moreover, throughout this section we will represent a retailer procurement plan by a sequence of periods $t_1 < \dots < t_n < t_{n+1} \equiv T + 1$ in which an order is placed. In other words, such

a sequence corresponds to the retailer order plan $y_t^r = 1$ if $t \in \{t_1, \dots, t_n\}$ and $y_t^r = 0$ otherwise. The cost of a retailer order plan (t_1, \dots, t_n) as a function of the incentive parameter θ is then given by

$$\sum_{i=1}^n \left(F_{t_i}^r(\theta) + \sum_{j=t_i}^{t_{i+1}-1} g_{t_i,j}^r(\theta) d_j \right). \quad (10)$$

We will discuss practical coordination mechanisms of the following form:

$$F_t^r(\theta) = F_t^r - \phi_t \theta \quad t = 1, \dots, T \quad (11)$$

$$g_{t\tau}^r(\theta) = g_{t\tau}^r - \gamma_{t\tau} \theta \quad \tau = t + 1, \dots, T; t = 1, \dots, T \quad (12)$$

with $\phi_t, \gamma_{t\tau} \geq 0$ ($\tau = t + 1, \dots, T; t = 1, \dots, T$). Although such combined single-parameter mechanisms are possible, we will limit ourselves to *fixed-cost-based* mechanisms of the form (11) (i.e., with $\gamma_{t\tau} = 0$ for $\tau = t + 1, \dots, T, t = 1, \dots, T$) and *variable-cost-based* mechanisms of the form (12) (i.e., with $\phi_t = 0$ for $t = 1, \dots, T$). For mechanisms of the form (11)–(12) the cost of a particular retailer order plan can be written as a linear function of the incentive parameter θ :

$$C^r(y^r|\theta) = C_0(t_1, \dots, t_n) - \theta C_1(t_1, \dots, t_n) \quad (13)$$

where the intercept is equal to

$$C_0(t_1, \dots, t_n) = \sum_{i=1}^n \left(F_{t_i}^r + \sum_{j=t_i}^{t_{i+1}-1} g_{t_i,j}^r d_j \right)$$

and the slope function for the two cases is equal to

$$C_1(t_1, \dots, t_n) = \sum_{i=1}^n \phi_{t_i} + \sum_{i=1}^n \sum_{j=t_i}^{t_{i+1}-1} \gamma_{t_i,j} d_j.$$

Moreover, it is easy to see that, in these cases, we have $\Theta_0 = [0, \theta_0^U]$ where

$$\theta_0^U = \min \left\{ \min_{t=1, \dots, T} \frac{F_t^r}{\phi_t}, \min_{\tau=t+1, \dots, T; t=1, \dots, T} \frac{g_{t\tau}^r}{\gamma_{t\tau}} \right\}$$

(where we interpret the ratio to be ∞ if the denominator is 0). For absolute discounts and rebates, a choice of $\Theta_2 = \mathbb{R}^+$ seems reasonable. Furthermore, we will parameterize relative discounts in

such a way that it is intuitive to let $\Theta_2 = [0, 1]$. In all of the cases that we will study it turns out that $\theta_0^U \geq 1$, so that in those cases we could then simply set $\theta_0^U = 1$ without loss of generality.

In Table 1, we introduce $g_{t\tau}^r(\theta)$ and $F_t^r(\theta)$ associated with several relevant mechanisms, and identify the corresponding parameters ϕ_t ($t = 1, \dots, T$) and $\gamma_{t\tau}$ ($\tau = t + 1, \dots, T$; $t = 1, \dots, T$). It is interesting to note that, for some incentive schemes, the slope function takes on a very special form that allows us to analyze these in more detail. Firstly, for absolute transportation cost rebates as well as for absolute price discounts in Cost Structure I the slope function reduces to $C_1(t_1, \dots, t_n) = d_{1T}$, so that these incentives will not change the optimal retailer order plan since the cost of *any* retailer order plan will decrease by the same amount. In other words, these incentives are not useful for the supplier. Secondly, for absolute order cost discounts we have $C_1(t_1, \dots, t_n) = n$. This means that, among all plans with exactly n order periods, the ones that have a minimum value for the intercept $C_0(t_1, \dots, t_n)$ dominate the others. In other words, the retailer's cost as a function of θ has just $\mathcal{O}(T)$ segments, which ensures that the running time of our algorithm is polynomial in the planning horizon T .

Coordination mechanism	$g_{t\tau}^r(\theta)$	$\gamma_{t\tau}$
Absolute transportation cost rebate	$\sum_{j=t}^{\tau-1} h_j^r + ((a_t^r - \theta) + p_t^s)(1 + \rho_{t\tau}^r)$	$1 + \rho_{t\tau}^r$
Absolute price discount	$\sum_{j=t}^{\tau-1} h_j^r + (a_t^r + (p_t^s - \theta))(1 + \rho_{t\tau}^r)$	$1 + \rho_{t\tau}^r$
Absolute holding cost rebate	$(a_t^r + p_t^s) + \sum_{j=t}^{\tau-1} (h_j^r + (a_t^r + p_t^s)\rho_j^r - \theta)$	$\tau - t$
Relative transportation cost rebate	$\sum_{j=t}^{\tau-1} h_j^r + ((1 - \theta)a_t^r + p_t^s)(1 + \rho_{t\tau}^r)$	$a_t^r(1 + \rho_{t\tau}^r)$
Relative price discount	$\sum_{j=t}^{\tau-1} h_j^r + (a_t^r + (1 - \theta)p_t^s)(1 + \rho_{t\tau}^r)$	$p_t^s(1 + \rho_{t\tau}^r)$
Relative holding cost rebate	$(a_t^r + p_t^s) + \sum_{j=t}^{\tau-1} (1 - \theta)(h_j^r + (a_t^r + p_t^s)\rho_j^r)$	$h_{t\tau}^r + (a_t^r + p_t^s)\rho_{t\tau}^r$
Coordination mechanism	$F_t^r(\theta)$	ϕ_t
Absolute order cost discount	$F_t^r - \theta$	1
Relative order cost discount	$(1 - \theta)F_t^r$	F_t^r

Table 1: Overview of relevant mechanisms

4.2 Constant price discounts and rebates in a time window

The incentive mechanisms we have discussed thus far have somewhat limited flexibility, because they require offering the proposed discount or rebate in each period. In this section we will assume that a particular incentive is only offered in periods $\mathcal{T} \subseteq \{1, \dots, T\}$ (which, although the elements of \mathcal{T} need not be consecutive, we will refer to as a *time window*) and consider practical fixed-cost and variable-cost-based coordination mechanisms of the form in equation (11) (where $\phi_t = 0$ for

$t \notin \mathcal{T}$) or the following slight generalization of that in equation (12):

$$g_{t\tau}^r(\theta) = g_{t\tau}^r - \gamma_{t\tau}^{\mathcal{T}}\theta \quad \tau = t + 1, \dots, T; t = 1, \dots, T; \mathcal{T} \subseteq \{1, \dots, T\}$$

where we are now able to set certain coefficients $\gamma_{t\tau}^{\mathcal{T}} = 0$ depending on the time window. For example, we could have a base set of values $\gamma_{t\tau}$ as in Section 4.1 and then define

$$\gamma_{t\tau}^{\mathcal{T}} = \gamma_{t\tau} 1_{\{t \in \mathcal{T}\}} \quad \tau = t + 1, \dots, T; t = 1, \dots, T; \mathcal{T} \subseteq \{1, \dots, T\}$$

where 1_A denotes an indicator function having the value 1 if event A occurs and 0 otherwise. This would correspond to a case where a time window based incentive is applied whenever an item is procured in a period in the time window. We then again obtain equation (13) with

$$C_0(t_1, \dots, t_n) = \sum_{i=1}^n \left(F_{t_i}^r + \sum_{j=t_i}^{t_{i+1}-1} g_{t_i, j}^r d_j \right)$$

and

$$C_1(t_1, \dots, t_n) = \sum_{i=1}^n 1_{\{t_i \in \mathcal{T}\}} \phi_{t_i} + \sum_{i=1}^n \sum_{j=t_i}^{t_{i+1}-1} \gamma_{t_i, j}^{\mathcal{T}} d_j. \quad (14)$$

Clearly there are $\mathcal{O}(2^T)$ possible time window choices. In order to ensure that the incentive scheme remains easily implemented and communicated to the retailer and that the supplier can identify optimal incentive parameters with reasonable computational effort, the following classes of time windows could be relevant in practice:

- (i) *fixed period time window*: time window of the form $\mathcal{T} = \{t\}$ for some $1 \leq t \leq T$;
- (ii) *initial periods time window*: time window of the form $\mathcal{T} = \{1, \dots, t\}$ for some $1 \leq t \leq T$;
- (iii) *remaining periods time window*: time window of the form $\mathcal{T} = \{t, \dots, T\}$ for some $1 \leq t \leq T$;
- (iv) *general time window*: time window of the form $\{t, \dots, \tau\}$ for some $1 \leq t \leq \tau \leq T$;
- (v) *periodic time window*: time windows of the form $\cup_{k=0}^{m-1} \{t + kT', \dots, \tau + kT'\} \cap \{1, \dots, T\}$ for $1 \leq t \leq \tau \leq T'$, where $m = \lfloor (T - 1)/T' \rfloor + 1$.

It is easy to see that there are $\mathcal{O}(T)$ time windows in classes (i)–(iii), $\mathcal{O}(T^2)$ time windows in class (iv), and $\mathcal{O}(T^3)$ time windows in class (v). We can apply the procedure developed in Section 3.4 for each potential time window in order to identify the supplier’s best combination of time window and incentive. However, for some incentive mechanisms applied with a fixed-period time window (i), as well as for absolute order cost discounts, price discounts, and transportation cost rebates, it can be shown that our algorithm runs in polynomial time in the planning horizon T (see also the discussion at the end of Section 4.1). We will discuss these in the remainder of this section.

4.2.1 Fixed period time windows

Under fixed period time windows we assume the incentive is offered only in a single period of the time horizon, i.e., $\mathcal{T} = \{\hat{t}\}$.

Fixed-cost–based incentives Under an order cost discount in a single period we obtain that

$$C_1(t_1, \dots, t_n) = \sum_{i=1}^n 1_{\{t_i \in \mathcal{T}\}} \phi_{t_i} = \begin{cases} \phi_{\hat{t}} & \text{if } \hat{t} \in \{t_1, \dots, t_n\} \\ 0 & \text{otherwise} \end{cases}$$

so that, for a given fixed period time window, there are only 2 slope values associated with production plan cost as a function of θ , which in turn implies that \hat{C}^r contains at most two line segments.

Variable-cost–based incentives When the time window based incentive is applied whenever an item is procured in period \hat{t} , i.e., the values $\gamma_{t\tau}^{\mathcal{T}}$ are of the form $\gamma_{t\tau}^{\mathcal{T}} = \gamma_{t\tau} 1_{\{t=\hat{t}\}}$ with the values of $\gamma_{t\tau}$ as in Section 4.1, we have that

$$C_1(t_1, \dots, t_n) = \sum_{i=1}^n \sum_{j=t_i}^{t_{i+1}-1} \gamma_{t_{ij}}^{\mathcal{T}} d_j = \begin{cases} \sum_{j=\hat{t}}^{t_{i+1}-1} \gamma_{t_{ij}} d_j & \text{if } \hat{t} = t_i \\ 0 & \text{otherwise} \end{cases}$$

so that, for a given fixed period time window there are no more than T slope values associated with production plan cost as a function of θ , which in turn implies that \hat{C}^r contains at most $\mathcal{O}(T)$ line segments.

Similarly, when a holding cost rebate is applied for all items in inventory in period \hat{t} under Cost

Structure I we have that $\gamma_{t\tau}^{\mathcal{T}} = h_{\hat{t}}^r 1_{\{t \leq \hat{t} < \tau\}}$, which yields

$$C_1(t_1, \dots, t_n) = \sum_{i=1}^n \sum_{j=t_i}^{t_{i+1}-1} \gamma_{t_i j}^{\mathcal{T}} d_j = h_{\hat{t}}^r \sum_{j=\hat{t}+1}^{t_{i+1}-1} d_j \quad \text{if } t_i \leq \hat{t} < t_{i+1}$$

so that, once again, for a given fixed period time window there are no more than T slope values associated with production plan cost as a function of θ , which in turn implies that \hat{C}^r contains at most $\mathcal{O}(T)$ line segments.

4.2.2 Absolute order cost and price discounts

In this section we consider an arbitrary but fixed time window \mathcal{T} .

Absolute order cost discount As noted above, an absolute order cost discount is characterized by $\phi_t = 1$ for $t \in \mathcal{T}$. Combining this with equation (14) we obtain that $C_1(t_1, \dots, t_n) = |\{t_1, \dots, t_n\} \cap \mathcal{T}|$. Analogously to the discussion at the end of Section 4.1 this means that, among all plans with exactly n' order periods in time window \mathcal{T} , the ones that have a minimum value for the intercept $C_0(t_1, \dots, t_n)$ dominate the others. In other words, the retailer's cost is a piecewise linear and concave function with at most $|\mathcal{T}| = \mathcal{O}(T)$ segments.

Absolute price discount (Cost Structure I) or transportation cost rebate Let \mathcal{T} be a general time window of the form $\{t, \dots, \tau\}$. As mentioned previously, an absolute price discount (under Cost Structure I) or transportation cost rebate is characterized by $\gamma_{t\tau} = 1$ for $t \in \mathcal{T}$. Combining this with equation (14) we obtain that

$$C_1(t_1, \dots, t_n) = \sum_{i=1}^n 1_{\{t_i \in \mathcal{T}\}} d_{t_i, t_{i+1}-1}$$

which implies that the retailer's cost is a piecewise linear and concave function with $\mathcal{O}(T^2)$ segments.

5 Computational results

5.1 Experimental design

In order to investigate the effectiveness of different coordination mechanisms on the system efficiency, we have implemented and tested the algorithms of Section 3 on five sets of problem instances with all parameters drawn from an integer uniform distribution over some interval $[a, b]$. For every problem set, demands are generated from the interval $[1, 20]$. The distributions of the other parameters can be found in Table 2. A single value means that the parameter is time-invariant, and p_t^w is the price charged by a wholesaler to the supplier. Note that the first two sets are equivalent since the parameters are time-invariant. The retailer faces Cost Structure I in sets A–C and Cost Structure II in sets A, D, and E. We assume that the supplier faces Cost Structure I for all sets. Note that the mean values of the parameters of sets B and D are equal to the values of the parameters of set A. For comparative purposes, the mean values of sets C and E are set 2 times higher than the mean values of Sets B and D, respectively.

Set	F_t^s	p_t^w	h_t^s	F_t^r	p_t^s	a_t^r	h_t^r	ρ_t
A	50	2	2	50	3	3	3	0
A	50	2	2	50	3	3	0	$\frac{1}{2}$
B	$5 \times [5, 15]$	$[1, 3]$	$[1, 3]$	$5 \times [5, 15]$	$[1, 5]$	$[1, 5]$	$[1, 5]$	0
C	$10 \times [5, 15]$	$[1, 7]$	$[1, 7]$	$10 \times [5, 15]$	$[1, 11]$	$[1, 11]$	$[1, 11]$	0
D	$5 \times [5, 15]$	$[1, 3]$	$[1, 3]$	$5 \times [5, 15]$	3	3	0	$\frac{1}{6}[1, 5]$
E	$10 \times [5, 15]$	$[1, 7]$	$[1, 7]$	$10 \times [5, 15]$	3	3	0	$\frac{1}{6}[1, 11]$

Table 2: Characteristics of the problem instances

For every problem set, we generated 100 instances with planning horizons $T \in \{5, 10, 15\}$ for which the optimal system cost under full coordination is strictly smaller than the system cost under anarchy, i.e., $C^* < C^A$ (so that system efficiency was well-defined). Overall, for 77% of all generated instances this was the case, with the corresponding number rapidly increasing in the planning horizon: the corresponding fractions for $T = 5, 10, 15$ were 60%, 85%, and 95%, respectively.

We consider 6 different single-parameter coordination mechanisms including fixed rebates, price discounts and holding cost subsidies, both in absolute and in relative terms; these are denoted by $F_a, p_a, h_a, F_r, p_r,$ and $h_r,$ respectively. Since p_t^s and a_t^r are generated from the same distribution in

our test bed, price discounts and transportation cost promotions will have the same effect on the system efficiency. Therefore, we only use price discounts as a coordination mechanism and we do not consider transportation cost discounts.

The lower bounds on the coordination parameter θ are set to zero, which means that the retailer is not worse off under the coordination scheme. Moreover, the upper bounds on θ are chosen such that the adjusted cost parameters are non-negative. If this does not hold, the retailer may make an infinite profit and hence the supplier an infinite loss. For instance, this happens in case $p_t^s - \theta < 0$, in which case the retailer will order an infinite amount in period t .

Besides single-parameter coordination mechanisms, we also test two-parameter mechanisms. The two-parameter mechanisms consist of all feasible combinations of two single-parameter mechanisms as presented before, and hence are still relatively easy to implement in practice. Note that not all combinations lead to feasible mechanisms. For example, the combination of an absolute and relative price discount is not feasible in case of time-invariant supplier prices p_t^s . All feasible combinations lead to 12 mechanisms in total.

We assume that the supplier does not know which optimal solution the retailer will choose in case of multiple alternatives and hence takes a worst-case approach by assuming that the retailer will choose the order plan that is most costly for the supplier (see equation (9)). Finally, we compare the single and two-parameter mechanisms when they are either valid over the whole horizon or in a periodic time window (see Section 4.2). In the latter case, the supplier solves problem (3) for each time window, and chooses the time window with the lowest cost.

5.2 The results

Table 3 shows the mechanism efficiency (ME) and mechanism potential (MP) when coordination occurs over the whole horizon (WH) or in a time window (TW) in the case of single-parameter coordination. We see that holding cost subsidies lead to the largest mechanism efficiency in general. As already observed in Section 4.1, absolute price discounts over the whole horizon do not lead to a positive efficiency in the case of Cost Structure I (sets A–C). Under whole horizon coordination, fixed rebates hardly lead to any efficiency improvements either: no more than 5% on average.

The efficiency improves significantly when coordination occurs in a time window instead of over the whole horizon. Time window coordination leads to an average efficiency between 56% and 85%

for at least one of the 6 coordination mechanisms. We observe that the holding cost subsidies lead to the largest mechanism efficiency in case of Cost Structure II, while this maximum is attained for either price discounts or holding cost subsidies in case of Cost Structure I. The increase in efficiency can be explained by the increased flexibility of the supplier; that is, the supplier can set the discounts in the periods where needed. This leads to higher supplier profits and a higher mechanism efficiency in general. We also observe that the efficiency under time window coordination is less sensitive to the length of the horizon than under whole horizon coordination. While the mechanism efficiency decreases over time under whole horizon coordination, it remains rather stable under time window coordination.

When comparing the mechanism efficiency to the mechanism potential, we observe small gaps. Especially when the mechanism potential is relatively low (under 30%), the mechanism efficiency differs by no more than 2 percentage points. This means that not much efficiency is lost due to the profit maximizing supplier, but the system as a whole benefits from the coordination mechanism. We may conclude from Table 3 that the cost of anarchy can be reduced considerably by employing single-parameter mechanisms.

Set	$T = 5$						$T = 10$						$T = 15$					
	F_a	p_a	h_a	F_r	p_r	h_r	F_a	p_a	h_a	F_r	p_r	h_r	F_a	p_a	h_a	F_r	p_r	h_r
A WH ME	0.00	0.00	0.74	0.00	0.00	0.74	0.00	0.00	0.63	0.00	0.00	0.63	0.00	0.00	0.56	0.00	0.00	0.56
A WH MP	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00
A TW ME	0.04	0.63	0.81	0.04	0.63	0.81	0.09	0.63	0.85	0.09	0.63	0.85	0.13	0.60	0.85	0.13	0.60	0.85
A TW MP	0.04	0.96	1.00	0.04	0.96	1.00	0.10	0.94	1.00	0.10	0.94	1.00	0.13	0.87	1.00	0.13	0.87	1.00
B WH ME	0.03	0.00	0.29	0.05	0.15	0.32	0.02	0.00	0.16	0.03	0.17	0.18	0.02	0.00	0.14	0.03	0.15	0.15
B WH MP	0.06	0.00	0.36	0.09	0.34	0.55	0.07	0.00	0.28	0.09	0.45	0.34	0.05	0.00	0.26	0.07	0.53	0.31
B TW ME	0.24	0.59	0.62	0.24	0.59	0.64	0.30	0.58	0.53	0.30	0.58	0.55	0.27	0.59	0.55	0.26	0.62	0.58
B TW MP	0.24	0.81	0.69	0.24	0.81	0.74	0.31	0.77	0.58	0.32	0.82	0.63	0.27	0.76	0.62	0.28	0.85	0.67
C WH ME	0.04	0.00	0.13	0.04	0.19	0.22	0.02	0.00	0.18	0.02	0.13	0.18	0.02	0.00	0.10	0.02	0.11	0.10
C WH MP	0.08	0.00	0.18	0.10	0.44	0.35	0.05	0.00	0.27	0.07	0.48	0.30	0.05	0.00	0.15	0.07	0.53	0.23
C TW ME	0.25	0.56	0.46	0.25	0.56	0.48	0.27	0.62	0.49	0.26	0.64	0.52	0.24	0.63	0.46	0.25	0.66	0.50
C TW MP	0.25	0.82	0.51	0.25	0.83	0.53	0.28	0.76	0.52	0.28	0.85	0.56	0.25	0.79	0.50	0.27	0.91	0.58
D WH ME	0.01	0.11	0.45	0.02	0.11	0.51	0.01	0.11	0.40	0.01	0.11	0.39	0.00	0.07	0.30	0.01	0.07	0.29
D WH MP	0.02	0.59	0.57	0.03	0.59	0.69	0.02	0.62	0.62	0.05	0.62	0.68	0.01	0.53	0.52	0.02	0.53	0.56
D TW ME	0.22	0.63	0.81	0.22	0.63	0.83	0.21	0.70	0.80	0.21	0.70	0.85	0.22	0.64	0.79	0.23	0.64	0.81
D TW MP	0.22	0.99	0.91	0.22	0.99	0.95	0.22	1.00	0.91	0.22	1.00	0.96	0.23	0.99	0.90	0.23	0.99	0.97
E WH ME	0.03	0.14	0.27	0.04	0.14	0.39	0.01	0.14	0.30	0.01	0.14	0.38	0.01	0.07	0.23	0.02	0.07	0.25
E WH MP	0.05	0.48	0.38	0.07	0.48	0.61	0.04	0.57	0.42	0.06	0.57	0.66	0.03	0.45	0.33	0.05	0.45	0.53
E TW ME	0.27	0.70	0.77	0.28	0.70	0.79	0.24	0.73	0.80	0.24	0.73	0.82	0.28	0.70	0.77	0.28	0.70	0.80
E TW MP	0.27	0.98	0.86	0.28	0.98	0.90	0.25	0.99	0.86	0.25	0.99	0.95	0.29	0.98	0.86	0.31	0.98	0.93

Table 3: Mechanism efficiency and potential for single-parameter coordination

Table 4 shows the mechanism efficiency and potential for two-parameter coordination mecha-

nisms. For space reasons, we have selected four combinations out of the twelve combinations that perform well in terms of efficiency. It turned out that an incentive combination consisting of either holding cost subsidies and price discounts, or holding cost subsidies and fixed rebates performed best on all data sets. Note that there are no results for the combination consisting of relative holding cost subsidies and relative price discounts under Cost Structure II (the same holds for the combination with relative holding cost subsidies and absolute price discounts). The reason is that the function \hat{C}^r is bilinear in the coordination parameters instead of linear, and hence the methodology of Section 3.5 can no longer be applied.

As in the single-parameter case, we observe that time window coordination leads to much better results than whole horizon coordination: in some cases the mechanism efficiency is even twice as high. Although there is no efficiency improvement under whole horizon coordination for set A when compared to single-parameter coordination, there is a reasonable efficiency improvement in the other cases. Time window coordination leads to an average efficiency between 63% and 88% for at least one of the four selected coordination mechanisms. Furthermore, we see that the mechanism potential is close to or equal to 100% for at least one coordination mechanism in all cases. This means that a relatively simple two-parameter coordination mechanism can potentially result in an (almost) fully coordinated supply chain in terms of system cost.

So far, Tables 3 and 4 have shown that simple coordination mechanisms may lead to considerable efficiency improvements. However, it is not clear that both the supplier and retailer benefit from this. Therefore, we report the average retailer share (ARS), maximum retailer share (MRS) and maximum supplier share (MSS) for a subset of the instances in Table 5. Given a problem instance and a coordination mechanism, the retailer (supplier) share is defined as the relative part of the difference between the system cost under anarchy and the optimal system cost under the mechanism, that comes as a benefit to the retailer (supplier). The average (maximum) retailer share is the average (maximum) over all instances. A similar definition holds for the average/maximum supplier share. Note that the average retailer and supplier share add up to 1 and hence we only report the average retailer shares. Furthermore, Table 5 also shows the number of instances for which an efficiency improvement (NEI) occurs, i.e., the number of times that a coordination mechanism leads to a reduction of the cost of anarchy.

It follows from Table 5 that the average retailer share is higher under whole horizon coordination

	$T = 5$				$T = 10$				$T = 15$			
	F_a, h_r	h_a, p_r	F_r, h_r	p_r, h_r	F_a, h_r	h_a, p_r	F_r, h_r	p_r, h_r	F_a, h_r	h_a, p_r	F_r, h_r	p_r, h_r
A WH ME	0.74	0.74	0.74	0.74	0.63	0.63	0.63	0.63	0.56	0.56	0.56	0.56
A WH MP	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
A TW ME	0.81	0.81	0.81	0.81	0.85	0.85	0.85	0.85	0.85	0.85	0.85	0.85
A TW MP	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
B WH ME	0.35	0.42	0.37	0.44	0.20	0.27	0.21	0.31	0.17	0.25	0.17	0.25
B WH MP	0.65	0.71	0.64	0.89	0.44	0.70	0.46	0.79	0.35	0.78	0.36	0.84
B TW ME	0.79	0.84	0.79	0.84	0.72	0.77	0.72	0.79	0.69	0.75	0.69	0.77
B TW MP	0.90	0.97	0.90	0.99	0.83	0.97	0.85	0.99	0.83	0.98	0.84	0.99
C WH ME	0.27	0.31	0.27	0.40	0.21	0.27	0.20	0.31	0.12	0.20	0.12	0.22
C WH MP	0.46	0.67	0.45	0.80	0.35	0.74	0.36	0.82	0.29	0.70	0.31	0.82
C TW ME	0.68	0.82	0.68	0.84	0.69	0.77	0.68	0.82	0.63	0.76	0.65	0.79
C TW MP	0.73	0.96	0.74	0.97	0.79	0.98	0.81	0.99	0.76	0.97	0.80	0.99
D WH ME	0.53	0.45	0.53	-	0.41	0.41	0.39	-	0.29	0.29	0.30	-
D WH MP	0.77	0.74	0.71	-	0.71	0.77	0.69	-	0.62	0.64	0.57	-
D TW ME	0.88	0.85	0.88	-	0.88	0.85	0.88	-	0.83	0.82	0.83	-
D TW MP	0.99	0.99	1.00	-	1.00	1.00	1.00	-	1.00	1.00	1.00	-
E WH ME	0.41	0.29	0.42	-	0.39	0.32	0.39	-	0.26	0.26	0.27	-
E WH MP	0.71	0.61	0.66	-	0.71	0.68	0.68	-	0.60	0.59	0.55	-
E TW ME	0.87	0.86	0.88	-	0.87	0.86	0.87	-	0.83	0.80	0.84	-
E TW MP	1.00	0.98	1.00	-	1.00	0.99	1.00	-	0.98	0.99	0.99	-

Table 4: Mechanism efficiency and potential for two-parameter coordination

than under time window coordination. The average retailer share is not more than 22% under time window coordination, while this quantity is often around 40% under whole horizon coordination. Apparently, the supplier benefits from the fact that he/she can choose the time window and he/she gets a bigger share of the total savings. However, the fact that both the maximum retailer and supplier shares are close to or equal to 1 shows that there is quite some variability in these statistics. Finally, we see that time window coordination often leads to an efficiency improvement for at least one coordination mechanism: more than 81 times in the case of single-parameter coordination and more than 93 times in the case of two-parameter coordination. Again, this suggests that the cost of anarchy can be reduced considerably by applying the simple mechanisms.

Finally, we report some statistics on the observed running times of the algorithms. For a given problem instance and a coordination mechanism, the supplier problem is solved in a fraction of a second. Therefore, we report the number of line segments of the retailer's optimal value function \hat{C}^r in the case of single-parameter coordination and the number of faces of \hat{C}^r in the case of two-parameter coordination. In particular, we counted the (possibly non-unique) production plans corresponding to line segments under single-parameter coordination, or equivalently, the number of retailer plans that are optimal for at least two different values of θ . In the case of two-parameter

			F_a	p_a	h_a	F_r	p_r	h_r	F_a, h_r	h_a, p_r	F_r, h_r	p_r, h_r
A	WH	ARS	-	-	0.42	-	-	0.42	0.42	0.42	0.42	0.42
A	WH	MRS	-	-	1.00	-	-	1.00	1.00	1.00	1.00	1.00
A	WH	MSS	-	-	0.97	-	-	0.97	0.97	0.97	0.97	0.97
A	WH	NEI	0	0	70	0	0	70	70	70	70	70
A	TW	ARS	0.01	0.22	0.11	0.01	0.22	0.11	0.11	0.11	0.11	0.11
A	TW	MRS	0.27	0.90	0.59	0.16	0.90	0.59	0.59	0.59	0.59	0.59
A	TW	MSS	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
A	TW	NEI	37	86	97	37	86	97	97	97	97	97
C	WH	ARS	0.32	-	0.30	0.17	0.50	0.44	0.42	0.42	0.40	0.48
C	WH	MRS	0.86	-	0.75	0.38	0.94	0.96	0.96	0.93	0.96	0.96
C	WH	MSS	1.00	-	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
C	WH	NEI	4	0	17	4	19	18	22	31	21	34
C	TW	ARS	0.01	0.15	0.06	0.02	0.16	0.10	0.06	0.10	0.08	0.10
C	TW	MRS	0.45	0.90	0.56	0.26	0.86	0.88	0.88	0.51	0.88	0.74
C	TW	MSS	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
C	TW	NEI	60	92	81	60	92	83	93	96	94	96
E	WH	ARS	0.64	0.27	0.43	0.52	0.27	0.41	0.40	0.45	0.42	-
E	WH	MRS	0.78	0.98	0.99	0.84	0.98	1.00	1.00	0.99	1.00	-
E	WH	MSS	0.50	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	-
E	WH	NEI	2	14	34	5	14	38	38	38	40	-
E	TW	ARS	0.01	0.17	0.09	0.02	0.17	0.12	0.12	0.10	0.12	-
E	TW	MRS	0.21	0.80	0.77	0.65	0.80	0.82	0.82	0.77	0.82	-
E	TW	MSS	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	-
E	TW	NEI	66	92	97	66	92	98	99	97	99	-

Table 5: Retailer/supplier shares and number of efficiency improvements for $T = 15$

coordination, we counted the number of (possibly non-unique) production plans corresponding to faces, or equivalently, the number of retailer plans that are optimal for at least three linearly independent points $\theta = (\theta_1, \theta_2)$.

Table 6 shows the number of line segments and faces of a subset of the mechanisms for Set E. In the time windows case, the averages and maxima are taken over all possible time windows. Furthermore, the lines denoted by ‘All’ give the maxima over all problem sets. It follows from the table that the average number of line segments is somewhat lower than T , while the maximum just exceeds T . The figures show a linear behavior as opposed to the subexponential behavior that was derived in Corollary 3.2. Clearly, since we have only tested a limited number of instances, no general conclusions can be drawn. The number of faces seem to increase somewhat faster than a linear function in T , although the numbers are still relatively small. For example, the possible number of retailer plans for a 15 periods instance is $2^{14} = 16,384$, while the maximum number of faces encountered equals 84. We conclude that running time is not an issue when applying the methodology to the instances of our test bed.

			$T = 5$			$T = 10$			$T = 15$		
mechanism			F_r	p_r	h_r	F_r	p_r	h_r	F_r	p_r	h_r
E	WH	Avg	3.2	2	3.2	6	2.9	6.3	8.8	3.2	9.1
E	WH	Max	5	4	6	8	6	11	12	7	14
E	TW	Avg	2.1	2.1	2.2	3.4	2.7	3.5	4.8	3.2	4.6
E	TW	Max	6	5	6	16	9	11	22	15	19
All	TW	Max	7	6	6	16	11	11	22	18	20
mechanism			F_a, h_r	h_a, p_r	F_r, h_r	F_a, h_r	h_a, p_r	F_r, h_r	F_a, h_r	h_a, p_r	F_r, h_r
E	WH	Avg	5.2	3.1	6	13.1	5.2	12.1	21.5	6.8	17.8
E	WH	Max	9	6	8	27	13	17	42	18	24
E	TW	Avg	3.5	2.9	4.1	7.1	5	9.7	10.9	6.9	15.9
E	TW	Max	11	10	13	34	23	46	54	36	76
All	TW	Max	11	10	13	36	24	46	62	50	84

Table 6: Number of line segments/faces

6 Concluding remarks

While a great deal of attention has justifiably been given to two-stage supply chain coordination problems in the operations literature, to our knowledge, this literature does not account for the effects of dynamic demands or economies of scale in inventory replenishment. This paper provides a modeling framework that addresses these complexities, as well as mechanisms that permit a supplier to improve operations efficiency in a two-stage supply chain when explicit coordination is not possible. We model this problem as a Stackelberg game in which the supplier sets coordination parameter values that guarantee that neither party is worse off (when compared to prevailing wholesale prices). As our computational tests showed, relatively simple and easily implemented mechanisms involving one or two parameters can indeed lead to non-trivial efficiency improvements without the need for explicit coordination.

References

- [1] A. Aggarwal and J.K. Park. Improved algorithms for economic lot-size problems. *Operations Research*, 41(3):549–571, 1993.
- [2] O. Ben-Ayed and C.E. Blair. Computational difficulties of bilevel linear programming. *Operations Research*, 38(3):556–560, 1990.
- [3] G.P. Cachon. Supply chain coordination with contracts. In S.C. Graves and T. de Kok, editors, *Handbook of Operations Research and Management Science: Supply Chain Management*. Kluwer, Amsterdam, The Netherlands, 2003.

- [4] G.P. Cachon. The allocation of inventory risk in a supply chain: push, pull, and advance-purchase discounts. *Management Science*, 50(2):222–238, 2004.
- [5] G.P. Cachon and M. Lariviere. Supply chain coordination with revenue sharing: strengths and limitations. *Management Science*, 51(1):30–44, 2005.
- [6] F. Chen, A. Federgruen, and Y.-S. Zheng. Coordination mechanisms for decentralized distribution systems with one supplier and multiple retailers. *Management Science*, 47(5):693–708, 2001.
- [7] F. Chen, A. Federgruen, and Y.-S. Zheng. Near-optimal pricing and replenishment strategies for a retail/distribution system. *Operations Research*, 49(6):839–853, 2001.
- [8] C. Corbett and X. de Groot. A suppliers optimal quantity discount policy under asymmetric information. *Management Science*, 3:444–450, 2000.
- [9] A. Cournot. *Researches into the Mathematical Principles of the Theory of Wealth*. Macmillan, New York, 2001. Original title: Recherches sur les principes mathématiques de la théorie des richesses (1838), translated by N.T. Bacon.
- [10] M.J. Eisner and D.G. Severance. Mathematical techniques for efficient record segmentation in large shared databases. *Journal of the Association of Computing Machinery*, 23:619–635, June 1976.
- [11] A. Federgruen and M. Tzur. A simple forward algorithm to solve general dynamic lot sizing models with n periods in $O(n \log n)$ or $O(n)$. *Management Science*, 37:909–925, 1991.
- [12] D. Fernández-Baca and S. Srinivasan. Constructing the minimization diagram of a two-parameter problem. *Operations Research Letters*, 10:87–93, 1991.
- [13] D. M. Gusfield. *Sensitivity analysis for combinatorial optimization*. PhD thesis, University of California, Berkeley, 1980. Memorandum UCB/ERL M80/22D - Electronics Research Laboratory.
- [14] S. van Hoesel, H.E. Romeijn, D. Romero Morales, and A.P.M. Wagelmans. Integrated lot-sizing in serial supply chains with production capacities. *Management Science*, 51(11):1706–1719, 2005.
- [15] T.D. Klastorin, K. Moinzadeh, and J. Son. Coordinating orders in supply chains through price discounts. *IIE Transactions*, 34:679–689, 2002.
- [16] R. Lal and R. Staelin. An approach for developing an optimal discount pricing policy. *Management Science*, 30:1524–1539, 1984.
- [17] M. Lariviere and E. Porteus. Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing & Service Operations Management*, 3(4):293–305, 2001.
- [18] H.L. Lee and J. Rosenblatt. A generalized quantity discount pricing model to increase suppliers profits. *Management Science*, 33:1167–1185, 1986.

- [19] J.P. Monahan. A quantity discount pricing model to increase vendor profits. *Management Science*, 30:720–726, 1984.
- [20] G. Perakis and G. Roels. The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Management Science*, 53(8):1249–1268, 2007.
- [21] T. Roughgarden. Selfish routing and the price of anarchy. *Optima*, 74:1–14, May 2007.
- [22] F. Sahin and E.P. Robinson. Information sharing and coordination in make-to-order supply chains. *Journal of Operations Management*, 23:579–598, 2005.
- [23] A. Topal and S. Çetinkaya. Contractual agreements for coordination and vendor-managed delivery under explicit transportation considerations. *Naval Research Logistics*, 53:397–417, 2006.
- [24] S. Viswanathan and R. Piplani. Coordinating supply chain inventories through common replenishment epochs. *European Journal of Operational Research*, 129:277–286, 2001.
- [25] A. Wagelmans, S. van Hoesel, and A. Kolen. Economic lot sizing: an $O(n \log n)$ algorithm that runs in linear time in the Wagner-Whitin case. *Operations Research*, 40-S1:S145–S156, 1992.
- [26] H.M. Wagner and T.M. Whitin. Dynamic version of the economic lot size model. *Management Science*, 5:89–96, 1958.
- [27] Z.K. Weng. Channel coordination and quantity discounts. *Management Science*, 41:1509–1522, 1995.
- [28] W.I. Zangwill. A backlogging model and a multi-echelon model of a dynamic economic lot size production system – a network approach. *Management Science*, 15(9):506–527, 1969.