

Two-Echelon Lot-Sizing with Asymmetric Information and Continuous Type Space

R.B.O. Kerckamp & W. van den Heuvel & A.P.M. Wagelmans

Econometric Institute Report EI2018-17
14th May 2018

Abstract

We analyse a two-echelon discrete lot-sizing problem with a supplier and a retailer under information asymmetry. We assume that all cost parameters are time independent and that the retailer has single-dimensional continuous private information, namely either his setup cost or his holding cost. The supplier uses mechanism design to determine a menu of contracts that minimises his expected costs, where each contract specifies the retailer's procurement plan and a side payment to the retailer. There is no restriction on the number of contracts in the menu.

To optimally solve this contracting problem we present a two-stage approach, based on a theoretical analysis. The first stage generates a list of procurement plans that is sufficient to solve the contracting problem to optimality. The second stage optimally assigns these plans to the retailer types and determines all side payments. The result is an optimal menu with finitely many contracts that pools retailer types. We identify cases for which the contracting problem can be solved in polynomial time and provide the corresponding algorithms. Furthermore, our analysis reveals that information asymmetry leads to atypical structures in the plans of the optimal menu, e.g., plans violating the zero-inventory property. Our solution approach and several results are directly applicable to more general problems as well.

1 Introduction

We consider a two-echelon supply chain consisting of a supplier and a retailer under a discrete lot-sizing setting with asymmetric information. The supply chain must come to an agreement for a joint procurement plan to satisfy market demand for a single indivisible product for a given time horizon. We assume that the market demand can be modelled as demand in discrete time periods and is known up front, leading to a discrete lot-sizing problem for the supply chain. That is, the joint procurement plan specifies the following for each time period up to the planning horizon. For the supplier, the options are to produce new products, to keep products in inventory for later time periods, and to transfer products to the retailer. For the retailer, these are to receive products from the supplier, to keep products in inventory, and to satisfy market demand. We assume all market demand must be satisfied and back-ordering is not allowed.

If all information is shared among the two parties in the supply chain, we have the traditional joint lot-sizing problem, which is well known and analysed thoroughly in the literature, e.g., in Zangwill (1969). We consider the case where there is only partial cooperation between the supplier and the retailer. Namely, the retailer has private information on his cost structure that he does not share with the supplier. Furthermore, we assume that the supplier and the retailer both act individually rationally and want to minimise their own costs. This partial cooperation typically leads to inefficiencies for the supply chain, see for example Inderfurth et al. (2013) and Perakis and Roels (2007). However, we consider the problem from the supplier's point of view, who wants to minimise his own costs, and thus perfect supply chain coordination is not a goal.

We assume that the supply chain uses a pull ordering strategy, i.e., the retailer has the initiative and the market power to place orders at the supplier. The supplier must satisfy these orders. Hence, by default the retailer will order according to his own individually optimal procurement plan, which is typically suboptimal for the supplier. The supplier has a single opportunity to offer the retailer a menu of contracts to persuade him to change his procurement plan. A single contract specifies the retailer's orders at the supplier and a side payment from the supplier to the retailer. By using a large enough side payment the supplier can convince the retailer to accept a different procurement plan. The menu

can contain any number of contracts. However, since the retailer can reject any offered contract and has private information on his cost structure, it is not trivial to design a menu of contracts that minimises the supplier's costs.

We consider the case where all cost parameters are time independent. Consequently, as we will show in Section 2, the only relevant costs are the supplier's setup cost of production, the retailer's setup cost for an order, and the holding costs for the inventory of both parties. We assume that the retailer's private information is either his setup cost or his holding cost, and lies in a certain interval. Thus, the private information is single dimensional, bounded, and continuous. We also assume that the supplier has a probability distribution for the retailer's private information.

The supplier uses mechanism design (see Laffont and Martimort (2002)) to construct a menu of contracts that minimises his expected costs, which requires solving a specific optimisation problem. We call this optimisation problem the *contracting problem*, which can be formulated as a mixed integer linear program with infinitely many variables and constraints. All details of the setting and the model will be given in Section 2.

Our goals are to analyse the contracting problem to obtain a tractable formulation and to determine efficiently solvable cases. Of particular interest is whether the information asymmetry changes the complexity class of the underlying optimisation problem. That is, if all information is shared the corresponding contracting problem turns out to be a traditional joint lot-sizing problem, which is solvable in polynomial time (see Zangwill (1969)). With information asymmetry the contracting problem is non-trivial, but can it still be solved in polynomial time? Before we state our main results for these questions, we first discuss the related literature to position our contribution.

1.1 Related literature

At its core, the described lot-sizing contracting problem is strongly related to the two-echelon lot-sizing problem. Also notice that the retailer's default plan, being individually optimal, follows from solving a single-level lot-sizing problem. Both these traditional lot-sizing problems have been analysed in detail in the literature. We refer to Wagner and Whitin (1958) and Zangwill (1969) for solution methods. As we will show, we need to solve several joint lot-sizing subproblems where either the number of retailer setups or the amount of retailer inventory is fixed. In a way, this relates to a parametric analysis (see Van Hoesel and Wagelmans (2000)) and stability regions of solutions (see Richter and Vörös (1989)). However, certain properties used in the previous references do not hold in general for our subproblems. For example, in specific cases the optimal menu of contracts contains procurement plans that do not satisfy the so-called zero-inventory property, implying that the retailer has unnecessary inventory when considered in isolation. More details are given in Section 3.

The lot-sizing contracting problem fits in the broader research field of the application of mechanism design to traditional optimisation problems. We focus on literature that considers related supply chain procurement problems with asymmetric information. See also Laffont and Martimort (2002) and Leng and Parlar (2005) for more general references on this topic.

Perhaps one of the most fundamental researched problems is the economic order quantity (EOQ) problem under information asymmetry. Compared to our lot-sizing setting, the EOQ problem considers a constant demand rate over time, an infinite time horizon, and divisible products. Several variations have been researched, such as the private information being continuous or discrete, and single- or two-dimensional (see for example Corbett and de Groote (2000), Inderfurth et al. (2013), Kerkkamp et al. (2018), and Pishchulov and Richter (2016)). Another setting is the newsvendor problem under information asymmetry, which considers a single period but with uncertain demand. This problem has been analysed in Burnetas et al. (2007), Cachon (2003), and Cakanyildirim et al. (2012) among others.

In these models a (single) order quantity describes the entire procurement plan and the total costs of each party have closed-form expressions in terms of this order quantity. In contrast, for the lot-sizing problem the total costs of each party cannot be expressed as (manageable) closed-form formulas. Instead, the costs follow from solving a combinatorial optimisation problem. This requires a different solution approach.

In Albrecht (2017) a coordination problem based on lot-sizing between a supplier and retailer is considered. Both parties only communicate the desired or supplied order quantities, no other information is shared. The focus lies on a heuristic coordination scheme which might lead to an optimal procurement plan for the entire supply chain. Certain conditions are identified for which this is indeed the case. In the proposed scheme, the retailer determines a list of individually optimal retailer plans, where each plan has a fixed number of retailer setups. The list is then offered to the supplier, who determines his

optimal response (a supplier plan) for each retailer plan. Finally, both parties jointly decide which of the resulting joint procurement plans is executed. These final negotiations should also include a way to divide the resulting profit gained from the coordination, for which several strategies are suggested but not analysed. Similar coordination and negotiation settings are analysed in for example Buer et al. (2013), Dudek and Stadtler (2005), and Dudek and Stadtler (2007).

The setting in Albrecht (2017) differs significantly from ours: our goal is to minimise the supplier's costs, not to achieve perfect supply chain coordination, and more information is available to the supplier. However, the coordination scheme has the following similarity to our case. As we will show in Section 3, with private setup cost it is sufficient for optimality to design a list of T plans, namely one plan for each possible number of retailer setups. In contrast, these plans follow from joint lot-sizing problems and are not individually optimal plans.

To our knowledge, only the works of Mobini et al. (2014) and Phouratsamay (2017) consider similar discrete lot-sizing problems under information asymmetry. In the setting of Mobini et al. (2014) the costs are time dependent and the retailer's private information is discrete and multi-dimensional. Several conditions are identified under which the retailer's behaviour, regarding the selection of contracts, is more structured. Furthermore, the case with private demand information is analysed. Phouratsamay (2017) also considers the lot-sizing contracting problem with time-dependent costs and discrete private information. Three contract variations are analysed: contracts without side payments, contracts where the side payments can only compensate the retailer's holding costs, and contracts with unrestricted side payments. If all information is shared among the supplier and retailer, the variant with restricted side payments is NP-hard and the other two are solvable in polynomial time. For the private information case, the variant without side payments is polynomially solvable, but the complexity for the others remain open. For all these cases a numerical study is performed, showing that using restricted side payments performs only slightly worse than using unrestricted side payments. We complement the work of Mobini et al. (2014) and Phouratsamay (2017) by considering continuous private information, which requires a different solution approach.

1.2 Contribution

We present and analyse a two-echelon discrete lot-sizing problem where the retailer has single-dimensional continuous private information. In this principal-agent contracting problem either the retailer's setup cost or his holding cost is private. To our knowledge, this type of problem has not been researched in the literature, and we are the first to analyse a principal-agent contracting problem with an underlying combinatorial structure and continuous private information. Based on a theoretical analysis, we propose a two-stage solution approach consisting of a plan-generation stage and a plan-assignment stage. We identify cases where these stages can be solved in polynomial time and give the corresponding algorithms. This provides further insights into the complexity of lot-sizing models with asymmetric information.

Moreover, we observe structural differences compared to traditional lot-sizing problems due to the information asymmetry, such as optimal menus with plans that violate the zero-inventory property. Furthermore, the contracting problem and several of our results have an intuitive graphical interpretation, which is also applicable to other problem settings. Therefore, we also describe a more general setting for which the (conceptual) model, the solution approach, and certain results are applicable as well.

The remainder is organised as follows. In Section 2 we formally introduce the setting of our problem and the associated optimisation model. In Section 3 we analyse this model, derive a solution approach, and prove complexity results. The generalisability of our results is the central topic of the discussion in Section 4, in which we also conclude our results.

2 The contracting problem

In this section we formalise the considered two-echelon lot-sizing contracting problem described in the introduction. In Section 2.1 we specify the lot-sizing setting, the two players involved in the problem, and their possible actions. The corresponding optimisation model is given in Section 2.2.

2.1 The setting

Our setting considers a discrete lot-sizing problem between a supplier and a retailer for a finite planning horizon $T \in \mathbb{N}_{\geq 1}$. The retailer needs to satisfy market demand $d_t \in \mathbb{N}_{>0}$ in each time period $t \in \mathcal{T} = \{1, \dots, T\}$ in the planning horizon. We assume that the products are indivisible, leading to discrete

demand, and that this demand is strictly positive and deterministic in each period. The strict positivity of the demand streamlines certain results and proofs, and will be discussed in Section 4.1. The market demand can be satisfied either from the retailer's inventory, i.e., surplus available from the previous time period, or directly from a retailer's order at the supplier. In turn, the supplier satisfies the retailer's orders either from available inventory or by setting up a new production.

In the entire supply chain lead-times are zero, all demand or orders must be met, and back-ordering is not allowed. Furthermore, in the first time period the starting inventory of the supplier and retailer are assumed to be zero and the retailer must end with zero inventory after the final time period T . Thus, to satisfy the market demand a procurement plan for the supply chain must be made. This plan specifies for each period $t \in \mathcal{T}$ the supplier's production quantity $x_t^S \in \mathbb{N}$ and the retailer's order quantity $x_t^R \in \mathbb{N}$ at the supplier. Since the demand is deterministic and there are no back-orders, these order quantities completely determine the flow of products in the supply chain. Hence, a procurement plan prescribes the setups, the order quantities, and the resulting inventory for the entire planning horizon.

In our setting all costs and revenues are time independent. Consequently, we can assume without loss of generality that the variable procurement costs and the revenue from sold products are zero in the supply chain. We will elaborate on this after giving the optimisation model. Therefore, there are two relevant types of costs involved for the supplier and retailer, namely setup cost and holding cost. If the retailer places an order he incurs a setup cost of $f \in \mathbb{R}_{>0}$ and keeping a unit of products in inventory costs $h \in \mathbb{R}_{>0}$ per time period for the retailer. Similarly, for the supplier we have setup cost $F \in \mathbb{R}_{>0}$ and holding cost $H \in \mathbb{R}_{>0}$.

As mentioned in the introduction, the retailer has single-dimensional private information, i.e, either his setup cost f or his holding cost h is private. To handle both cases, we use θ for his private cost and call θ the retailer's *type*. To be precise, if the setup cost f is private we define $f(\theta) = \theta$ and $h(\theta) = h$. In the other case, with private holding cost h , we define $f(\theta) = f$ and $h(\theta) = \theta$. We assume that the supplier has estimated the retailer's private information θ to follow a strictly positive continuous distribution $\omega : \Theta \rightarrow \mathbb{R}_{>0}$ on a closed interval $\Theta = [\underline{\theta}, \bar{\theta}] \subset \mathbb{R}_{>0}$.

By assumption, the retailer has the market power to enforce any retailer's procurement plan onto the supplier. Consequently, by default the retailer orders according to his individually optimal plan, which depends on his type. The corresponding retailer's default costs are denoted by $\phi^*(\theta)$ for type $\theta \in \Theta$ and follow from solving a traditional single-level lot-sizing problem. We refer to ϕ^* as the retailer's *default option*, also known as his reservation level. The supplier uses mechanism-design techniques by offering the retailer a menu of contracts to incentivise the retailer to alter his procurement plan. The menu effectively assigns a contract to each type $\theta \in \Theta$, where a contract prescribes the retailer's order quantities $x_t^R(\theta)$, $t \in \mathcal{T}$, and a side payment $z(\theta) \in \mathbb{R}$ from the supplier to the retailer. However, the retailer has the power to choose any of the offered contracts or his default option, whichever minimises his own costs. Therefore, this menu of contracts has to be specifically designed by the supplier, as will be made clear in the next section when discussing the optimisation model.

The overall goal of the supplier is to design a menu of contracts that minimises the supplier's expected net costs whilst ensuring that the retailer can satisfy the market demand. There is no restriction on the number of contracts in the menu, but an optimal menu with fewer contracts is preferred. Finally, the menu can be offered only once and there are no renegotiations.

2.2 The contracting model

We formulate an optimisation model to determine an optimal menu of contracts that minimises the supplier's expected net costs, as described in the previous section. To this end, let $y_t^R \in \mathbb{B}$ denote whether the retailer has a setup (places an order) at time $t \in \mathcal{T}$ and let $I_t^R \in \mathbb{N}$ be the retailer's ending inventory at time $t \in \mathcal{T}$. Similarly, we have the setup indicator $y_t^S \in \mathbb{B}$ and the ending inventory $I_t^S \in \mathbb{N}$ for the supplier. Recall that $x_t^S, x_t^R \in \mathbb{N}$ are the order quantities and $z \in \mathbb{R}$ is the side payment. The *contracting model* is defined as follows:

$$\min \int_{\underline{\theta}}^{\bar{\theta}} \omega(\theta) \left(F \sum_{t \in \mathcal{T}} y_t^S(\theta) + H \sum_{t \in \mathcal{T}} I_t^S(\theta) + z(\theta) \right) d\theta \quad (1)$$

subject to

$$I_0^S(\theta) = 0, \quad \forall \theta \in \Theta, \quad (2)$$

$$I_{t-1}^S(\theta) + x_t^S(\theta) = I_t^S(\theta) + x_t^R(\theta), \quad \forall \theta \in \Theta, t \in \mathcal{T}, \quad (3)$$

$$x_t^S(\theta) \leq M y_t^S(\theta), \quad \forall \theta \in \Theta, t \in \mathcal{T}, \quad (4)$$

$$I_0^R(\theta) = I_T^R(\theta) = 0, \quad \forall \theta \in \Theta, \quad (5)$$

$$I_{t-1}^R(\theta) + x_t^R(\theta) = I_t^R(\theta) + d_t, \quad \forall \theta \in \Theta, t \in \mathcal{T}, \quad (6)$$

$$y_t^R(\theta) \leq x_t^R(\theta) \leq M y_t^R(\theta), \quad \forall \theta \in \Theta, t \in \mathcal{T}, \quad (7)$$

$$y_t^S(\theta), y_t^R(\theta) \in \mathbb{B}, \quad \forall \theta \in \Theta, t \in \mathcal{T}, \quad (8)$$

$$x_t^S(\theta), x_t^R(\theta), I_t^S(\theta), I_t^R(\theta) \in \mathbb{N}, \quad \forall \theta \in \Theta, t \in \mathcal{T}, \quad (9)$$

$$f(\theta) \sum_{t \in \mathcal{T}} y_t^R(\hat{\theta}) + h(\theta) \sum_{t \in \mathcal{T}} I_t^R(\hat{\theta}) \equiv \phi(\mathbf{x}(\hat{\theta})|\theta), \quad \forall \theta, \hat{\theta} \in \Theta, \quad (10)$$

$$\phi(\mathbf{x}(\theta)|\theta) - z(\theta) \leq \phi^*(\theta), \quad \forall \theta \in \Theta, \quad (11)$$

$$\phi(\mathbf{x}(\theta)|\theta) - z(\theta) \leq \phi(\mathbf{x}(\hat{\theta})|\theta) - z(\hat{\theta}), \quad \forall \theta, \hat{\theta} \in \Theta. \quad (12)$$

Here, the objective (1) is to minimise the supplier's expected net costs, which consists of setup and holding costs and the side payment paid to the retailer. Constraints (2)-(9) are the lot-sizing constraints for the procurement plan of each contract, constraints (10) are for notational convenience, and constraints (11)-(12) are the mechanism-design constraints.

In particular, constraints (2) make sure that the supplier's inventory at the start of the planning horizon is zero. Constraints (3) model the supplier's inventory balance, i.e., the flow of products on the supplier's level. Next, constraints (4) enforce that a setup takes place if at least one unit of products is produced. Here, M is a suitably large number, e.g., $M = \sum_{t \in \mathcal{T}} d_t$.

Constraints (5)-(7) are similar and correspond to the retailer. Note that by assumption the supplier can only prescribe the retailer's order quantities $x_t^R(\theta)$, so he cannot force the retailer to incur the setup cost f by using a dummy order of zero products. This is reflected in the model by $y_t^R(\theta) \leq x_t^R(\theta)$ in (7), which is explicitly needed for correctness. We also enforce our assumption that $I_T^R(\theta) = 0$ in (5). This is in contrast to traditional lot-sizing models without information asymmetry.

Moreover, for given order quantities (x_t^R , $t \in \mathcal{T}$) the rest of the retailer's procurement plan (y_t^R , I_t^R , $t \in \mathcal{T}$) is fixed. In other words, there is a bijection between the retailer's order quantities and his procurement plan. Therefore, we denote a contract by $(\mathbf{x}(\theta), z(\theta))$, where $\mathbf{x}(\theta)$ encodes the retailer's procurement plan. We omit the superscript for the retailer in $\mathbf{x}(\theta)$ to simplify our notation.

In constraints (10) we define $\phi(\mathbf{x}(\hat{\theta})|\theta)$ as the retailer's lot-sizing costs when using plan $\mathbf{x}(\hat{\theta})$ and being type $\theta \in \Theta$. Next, constraints (11) are the Individual Rationality (IR) constraints, which imply that for retailer type θ the contract $(\mathbf{x}(\theta), z(\theta))$ leads to net costs that do not exceed his default costs $\phi^*(\theta)$. Constraints (12) are the Incentive Compatibility (IC) constraints and require for retailer type θ that contract $(\mathbf{x}(\theta), z(\theta))$ has the lowest net costs of all contracts. Thus, (11) and (12) ensure that the retailer of type θ will accept his intended contract $(\mathbf{x}(\theta), z(\theta))$.

We conclude this section with several remarks on the model related to information asymmetry. First, notice that the supplier in fact faces a bi-level optimisation problem. In our case, the retailer's response to a contract can be easily incorporated by (5)-(7), leading to a single-level optimisation model. However, care has to be taken to enforce the proper behaviour, i.e., no dummy setups and no excess supply of products, as explained above. Second, as $\phi^*(\theta) \leq \phi(\mathbf{x}|\theta)$ by definition for any feasible plan \mathbf{x} , any feasible contract has a non-negative side payment by (11). Finally, in the previous section we claimed that any time-independent variable cost or revenue can be assumed to be zero. It is trivial to verify that $\sum_{t \in \mathcal{T}} x_t^S(\theta) = \sum_{t \in \mathcal{T}} x_t^R(\theta) = \sum_{t \in \mathcal{T}} d_t$ for all $\theta \in \Theta$ in any optimal solution. Therefore, a non-zero time-independent variable cost/revenue either leads to a constant term in the objective or cancels out in (11) and (12). This is also the case if that cost/revenue is private information. Hence, we only need to include setup and holding costs.

Clearly, complicating factors in solving the contracting model are the infinitely many variables and constraints. In the next section, we describe a solution approach which leads to polynomial-time algorithms in certain cases.

3 Solution approach

To solve the contracting model introduced in Section 2.2 we propose a two-stage approach. In the first stage, a list of procurement plans for the supply chain is constructed such that the list is sufficient for solving the contracting problem in the second stage. Next, in the second stage, the plans are assigned to retailer types and appropriate side payments are determined. To justify this approach, we start by analysing the contracting model in Section 3.1. The plan assignment is discussed in Section 3.2 and the plan generation in Section 3.3. All corresponding proofs are given in Appendix A.


3.1 Analysis

First of all, let us state some well-known properties of the retailer's default option ϕ^* , i.e., it being the lower envelope of at most T linear functions in $\theta \in \Theta$, see Lemma 1. The stated *zero-inventory property* means that a setup only occurs if there is no inventory from the previous time period. In this case, it refers to only the retailer's level, i.e., $y_t^R I_{t-1}^R = 0$ for all $t \in \mathcal{T}$.

Lemma 1 (Van Hoesel and Wagelmans (2000)). *The retailer's default option $\phi^*(\theta)$ is piecewise linear, non-decreasing, concave, and continuous in the retailer type $\theta \in \Theta$. It consists of at most T linear segments and the corresponding retailer's default plans satisfy the zero-inventory property. A complete specification of ϕ^* can be determined in $\mathcal{O}(T^2)$ time.*

We have a graphical interpretation of Lemma 1 which is also useful for the contracting model. For any retailer plan we can plot its costs as a function of θ , see Figure 1a for a conceptual example. In this example we assume that only the 5 shown retailer plans exist to keep the figure legible. In a real example, there could be excessively, but still finitely, many feasible plans. The horizontal axis is the type space which contains Θ . The vertical axis is the retailer's total costs $\phi(\mathbf{x}|\theta)$. Each line is a retailer plan \mathbf{x} with the following properties. In case of private setup cost, the slope is equal to the number of retailer setups $\sum_{t \in \mathcal{T}} y_t^R$ and the intersection with the vertical axis is the retailer's total holding costs $h \sum_{t \in \mathcal{T}} I_t^R$. In case of private holding cost, the slope is the total amount of retailer inventory $\sum_{t \in \mathcal{T}} I_t^R$ and the intersection with the vertical axis is the retailer's total setup costs $f \sum_{t \in \mathcal{T}} y_t^R$. In either case, the slope implies the retailer's *private* costs and the intersection with the vertical axis is equal to the retailer's *public* costs.

In Figure 1a plans I, II, and IV form the retailer's default option ϕ^* (shown in red). Plans III and V are never optimal for the retailer. For the contracting model the optimal supplier plans are determined for these retailer plans. If, for example, plans III and V result in very low costs for the supplier, he can use side payments to incentivise the retailer to accept these plans instead of I, II, and IV, as shown in blue in Figure 1b. Side payments shift the lines vertically, leading to a new lower envelope, which must lie under ϕ^* for $\theta \in \Theta$ by constraints (11).



(a) Without side payments the retailer accepts contracts II or IV.

(b) Using side payments to incentivise the retailer to accept contracts III or V.

Figure 1: Conceptual graphical interpretation of the contracting model.

From the graphical interpretation it follows intuitively that the optimal menu leads to a piecewise linear, non-decreasing, concave, and continuous function (a lower envelope) in terms of $\theta \in \Theta$, which lies below ϕ^* in Θ . Hence, the slopes of the segments must be non-increasing. Furthermore, if multiple segments have the same slope, only one with the lowest supplier's net costs is required. This implies a

strong ordering in the slopes, i.e., either the number of retailer setups (private setup cost) or the retailer inventory (private holding cost) is strictly decreasing. This result is formalised in Lemma 2.

Lemma 2. *Without loss of optimality, any two distinct contracts $(\mathbf{x}(\theta), z(\theta))$ and $(\mathbf{x}(\hat{\theta}), z(\hat{\theta}))$ for some $\theta < \hat{\theta} \in \Theta$ in an optimal menu satisfy*

$$\begin{cases} \sum_{t \in \mathcal{T}} y_t^R(\theta) > \sum_{t \in \mathcal{T}} y_t^R(\hat{\theta}) & \text{if setup cost } f \text{ is private} \\ \sum_{t \in \mathcal{T}} I_t^R(\theta) > \sum_{t \in \mathcal{T}} I_t^R(\hat{\theta}) & \text{if holding cost } h \text{ is private} \end{cases} \quad (13)$$

A direct consequence of the strict ordering in Lemma 2 and the discrete nature of the involved quantities is that offering only a limited number of contracts is sufficient for optimality for the contracting problem. By doing so, multiple retailer types will be assigned the same contract, which is called *pooling*. Moreover, it follows that this pooling occurs in a structured way: the interval $[\underline{\theta}, \bar{\theta}]$ is partitioned into subintervals and each subinterval is assigned a unique contract. This effect is again intuitively clear from the graphical interpretation in Figure 1b. More details are provided in Corollary 3.

Corollary 3. *Without loss of optimality, an optimal menu partitions (pools) the retailer types into subintervals and consists of at most*

$$\begin{cases} T & \text{contracts if setup cost } f \text{ is private} \\ 1 + \sum_{t \in \mathcal{T}} (t-1)d_t & \text{contracts if holding cost } h \text{ is private} \end{cases}.$$

In particular, for such an optimal menu consisting of $K \in \mathbb{N}_{\geq 1}$ distinct contracts the types $[\underline{\theta}, \bar{\theta}]$ are partitioned into K closed subintervals $[\underline{\theta}_k, \bar{\theta}_k]$, $k \in \{1, \dots, K\}$, where the k -th contract is the most preferred contract for all types in the k -th subinterval $[\underline{\theta}_k, \bar{\theta}_k]$.

The maximum number of contracts stated in Corollary 3 is the number of feasible slopes that can be achieved. By Lemma 2 it is sufficient for optimality to design a single plan for each feasible slope, i.e., for each feasible number of retailer setups $\sum_{t \in \mathcal{T}} y_t^R$ (private setup cost) or retailer inventory $\sum_{t \in \mathcal{T}} I_t^R$ (private holding cost). It turns out that the procurement plans in an optimal menu can be determined independently from each other, each following from a modified joint lot-sizing problem, see Theorem 4.

Theorem 4. *Without loss of optimality, the lot-sizing variables of a contract $(\mathbf{x}(\theta), z(\theta))$ in an optimal menu satisfying*

$$\begin{cases} \sum_{t \in \mathcal{T}} y_t^R(\theta) = n & \text{if setup cost } f \text{ is private} \\ \sum_{t \in \mathcal{T}} I_t^R(\theta) = n & \text{if holding cost } h \text{ is private} \end{cases} \quad (14)$$

are determined by solving a corresponding joint lot-sizing problem, namely minimising

$$\begin{cases} \sum_{t \in \mathcal{T}} \left(Fy_t^S(\theta) + HI_t^S(\theta) + hI_t^R(\theta) \right) & \text{if setup cost } f \text{ is private} \\ \sum_{t \in \mathcal{T}} \left(Fy_t^S(\theta) + HI_t^S(\theta) + fy_t^R(\theta) \right) & \text{if holding cost } h \text{ is private} \end{cases} \quad (15)$$

under the constraints (2)-(9) and (14).

We call the joint lot-sizing problem of Theorem 4 the *n-plan generation problem*. Notice that the *n-plan generation problem* only includes the supplier's setup and holding costs and the retailer's *public* costs. The retailer's *private* costs are fixed by (14). We can obtain this result in the graphical interpretation as well. Consider the situation in Figure 1a. We can 'normalise' all plans by shifting them downwards so they intersect with the origin, by setting the side payment equal to the retailer's public costs of the plan. All plans with the same slope (see (14)) are now essentially equivalent and it is optimal to only use the plan with the lowest supplier's 'normalised' costs. That is, the plan for which (15) is minimal, as the normalisation incorporates the retailer's public costs into the supplier's costs.

From the above theoretical results, we conclude that it is sufficient for optimality to solve the *n-plan generation problem* for each feasible slope n and use these plans to design a menu of contracts. The described plan generation is the first stage of our solution approach. We postpone the analysis of the *n-plan generation problem* to Section 3.3. First, we continue in Section 3.2 with the second stage of the solution approach: the *plan assignment problem*, where we need to assign the plans to the retailer types by using side payments, leading to a menu of contracts.

3.2 Plan assignment

From Section 3.1 we can assume without loss of optimality that we have a finite list of procurement plans, obtained from the plan-generation stage. The next step is the *plan assignment* stage where we need to decide which plans of the list will be incorporated into contracts and how these plans/contracts are assigned to the retailer types. Before we state the plan assignment model in Section 3.2.2, we derive two properties in Section 3.2.1 that will simplify the model.

3.2.1 Properties

We first introduce additional notation. Let $K \in \mathbb{N}_{\geq 1}$ be the number of plans in the considered list. For now, assume that each plan is included into a contract. We can index and sort the contracts by decreasing slope of the retailer plan, resulting in (\mathbf{x}_k, z_k) for $k \in \mathcal{K} = \{1, \dots, K\}$. By Lemma 2 and Corollary 3 an optimal menu will partition $[\underline{\theta}, \bar{\theta}]$ into K subinterval $[\underline{\theta}_k, \bar{\theta}_k]$, where the k -th contract will be assigned to types in $[\underline{\theta}_k, \bar{\theta}_k]$.

For the plan assignment, we need to take the IR and IC constraints (11)-(12) into account. These infinitely many constraints can be made tractable by using the partition structure described above. See Lemma 5 for the result.

Lemma 5. *To determine an optimal menu with K distinct contracts (\mathbf{x}_k, z_k) , $k \in \mathcal{K} = \{1, \dots, K\}$, and corresponding partition subintervals $[\underline{\theta}_k, \bar{\theta}_k]$, the IR and IC constraints (11)-(12) are equivalent to:*

$$\begin{aligned} \phi(\mathbf{x}_1|\underline{\theta}) - z_1 &\leq \phi^*(\underline{\theta}), \\ \phi(\mathbf{x}_k|\underline{\theta}_k) - z_k &\leq \phi^*(\underline{\theta}_k), & \forall k \in \mathcal{K} \setminus \{1\}, \\ \phi(\mathbf{x}_K|\bar{\theta}) - z_K &\leq \phi^*(\bar{\theta}), \\ \phi(\mathbf{x}_k|\underline{\theta}_k) - z_k &= \phi(\mathbf{x}_{k-1}|\underline{\theta}_k) - z_{k-1}, & \forall k \in \mathcal{K} \setminus \{1\}. \end{aligned} \quad (16)$$

Consider Lemma 5 in the graphical interpretation. The plan assignment problem essentially consists of shifting the lines in Figure 1a vertically to construct an optimal lower envelope for domain Θ (seen in blue in Figure 1b). From the piecewise linearity, concavity, and continuity of ϕ^* and the new lower envelope, it follows immediately that we only need to consider the IR constraints at the breakpoints. Furthermore, (16) relates to the continuity of the constructed lower envelope.

The second property relates to redundant plans included in the list. In an optimal menu it might be the case that not all provided plans are assigned to retailer types. Ideally, having these redundant plans included in the list should not interfere with the optimisation process. Lemma 6 shows that this is indeed the case: redundant plans can safely be added without affecting the optimum.

Lemma 6. *Having redundant plans/contracts does not affect the plan assignment problem.*

Graphically, the lines of redundant plans can/are placed tangent to the constructed lower envelope. This does not affect the lower envelope (the optimum), but ensures feasibility according to the equivalent IR and IC constraints stated in Lemma 5. From this point onwards, given a menu of contracts, a plan k in the menu is called *assigned* if $\bar{\theta}_k > \underline{\theta}_k$ and *redundant* if $\bar{\theta}_k = \underline{\theta}_k$.

3.2.2 The plan assignment model

We can now formulate the plan assignment model. To do so in a unified way for both cases of private information, we introduce new notation for the supplier's lot-sizing costs, the retailer's public lot-sizing costs, and the slope corresponding to the private information. For a procurement plan \mathbf{x} the corresponding supplier's costs is

$$C = \sum_{t \in \mathcal{T}} (Fy_t^S + HI_t^S).$$

If setup cost f is private, we define

$$c^{\text{pub}} = h \sum_{t \in \mathcal{T}} I_t^R, \quad n = \sum_{t \in \mathcal{T}} y_t^R.$$

Otherwise, if holding cost h is private, we have

$$c^{\text{pub}} = f \sum_{t \in \mathcal{T}} y_t^R, \quad n = \sum_{t \in \mathcal{T}} I_t^R.$$

Hence, by definition, we have $\phi(\mathbf{x}_k|\theta) = c_k^{\text{pub}} + n_k\theta$ for $k \in \mathcal{K}$ and $\theta \in \Theta$. Note that in the plan-assignment stage C_k , c_k^{pub} , and n_k are all known parameters and follow from the provided list of plans. It is essential that the plans are sorted such that $n_k > n_{k+1}$ for all $k \in \mathcal{K}$ to ensure that the model is in line with Lemma 2.

The *plan assignment model* is given by:

$$\min \sum_{k \in \mathcal{K}} \left(\int_{\underline{\theta}_k}^{\bar{\theta}_k} \omega(\theta) d\theta \right) (C_k + z_k) \quad (17)$$

subject to

$$c_1^{\text{pub}} + n_1\theta - z_1 \leq \phi^*(\theta), \quad (18)$$

$$c_k^{\text{pub}} + n_k\underline{\theta}_k - z_k \leq \phi^*(\underline{\theta}_k), \quad \forall k \in \mathcal{K} \setminus \{1\}, \quad (19)$$

$$c_K^{\text{pub}} + n_K\bar{\theta} - z_K \leq \phi^*(\bar{\theta}), \quad (20)$$

$$c_{k-1}^{\text{pub}} - c_k^{\text{pub}} + (n_{k-1} - n_k)\underline{\theta}_k = z_{k-1} - z_k, \quad \forall k \in \mathcal{K} \setminus \{1\} \quad (21)$$

$$\theta_1 = \theta, \quad (22)$$

$$\underline{\theta}_k \leq \bar{\theta}_k, \quad \forall k \in \mathcal{K}, \quad (23)$$

$$\bar{\theta}_k = \underline{\theta}_{k+1}, \quad \forall k \in \mathcal{K} \setminus \{K\}, \quad (24)$$

$$\bar{\theta}_K = \bar{\theta}. \quad (25)$$

Here, (18)-(20) are the IR constraints and (21) the IC constraints as described in Lemma 5. The constraints (22)-(25) model the partition of $[\theta, \bar{\theta}]$ and the corresponding assignment of contracts to subintervals as stated in Corollary 3. Consequently, the integral in the objective (17) is the probability that the retailer accepts contract (\mathbf{x}_k, z_k) .

We emphasize again that the model is only correct if $n_k > n_{k+1}$ for all $k \in \mathcal{K}$ (by Lemma 2). Also, by Lemma 6 redundant plans/contracts can be added without affecting the optimum, provided that the ordering in n_k is maintained. Moreover, note that $\phi^*(\underline{\theta}_k)$ in (19) can be modelled with at most T linear constraints for each $k \in \mathcal{K}$ (see Lemma 1). Namely, replace (19) by

$$c_k^{\text{pub}} + n_k\underline{\theta}_k - z_k \leq c_l^* + n_l^*\underline{\theta}_k, \quad \forall l \in \mathcal{L}, k \in \mathcal{K} \setminus \{1\},$$

where the $\theta \mapsto c_l^* + n_l^*\theta$, $l \in \mathcal{L}$ ($|\mathcal{L}| \leq T$), correspond to the retailer's default plans and whose lower envelope is ϕ^* . Finally, by combining (21) the IC constraints imply for $k \in \mathcal{K}$ that

$$z_k = z_1 + c_k^{\text{pub}} - c_1^{\text{pub}} - \sum_{i=2}^k (n_{i-1} - n_i)\underline{\theta}_i.$$

Substituting this expression in the objective function, results in a separable non-linear objective:

$$\begin{aligned} & z_1 - c_1^{\text{pub}} + \sum_{k \in \mathcal{K}} \left(\int_{\underline{\theta}_k}^{\bar{\theta}_k} \omega(\theta) d\theta \right) \left(C_k + c_k^{\text{pub}} - \sum_{i=2}^k (n_{i-1} - n_i)\underline{\theta}_i \right) \\ &= z_1 - c_1^{\text{pub}} + \sum_{k \in \mathcal{K}} \left(\int_{\underline{\theta}_k}^{\bar{\theta}_k} \omega(\theta) d\theta \right) \left(C_k + c_k^{\text{pub}} \right) - \sum_{k=2}^K \left(\int_{\underline{\theta}_k}^{\bar{\theta}_k} \omega(\theta) d\theta \right) (n_{k-1} - n_k)\underline{\theta}_k. \end{aligned} \quad (26)$$

Thus, the plan assignment model has a formulation with linear constraints and a non-linear separable objective function. In general, such optimisation models are difficult to solve to optimality, but several (heuristic) solution approaches have been designed (see for example Bradley et al. (1977), Byrd et al. (2003), and Kolda et al. (2007)).

If the retailer type distribution ω is uniform, the plan assignment model has a hidden convexity. The standard formulation is still non-convex, but by using the reformulated objective function (26) we obtain a linearly-constrained convex-quadratic model. It is well known that these models can be solved efficiently (see Ye and Tse (1989)). This result is captured in Theorem 7 and its proof contains the details of the convex formulation.

Theorem 7. *If ω is a uniform distribution, then the plan assignment model can be formulated as a linearly-constrained convex-quadratic model. It can be solved in polynomial time in the number of contracts K by interior-point methods.*

3.3 Plan generation

In the plan-generation stage we need to solve several joint lot-sizing problems as described in Theorem 4. In Section 3.3.1 we first give properties of this problem that are common for the two private information cases. Then we focus on each case separately: private setup cost in Section 3.3.2 and private holding cost in Section 3.3.3.

3.3.1 Common properties

In a standard joint lot-sizing problem, i.e., without constraint (14), it is well known that there exists an optimal solution that satisfies the zero-inventory property. Such an optimal solution can be found in polynomial time using dynamic programming by its decomposition into independent subplans. In contrast, for the n -plan generation problem the optimal solution might not satisfy the zero-inventory property, as we will show later. However, the zero-inventory property always holds for the supplier's lot-sizing plan, see Lemma 8.

Lemma 8. *For an optimal solution for the n -plan generation problem, the supplier's lot-sizing plan must satisfy the zero-inventory property.*

Another property in certain joint lot-sizing problems is that the joint plan is *nested*. This means that a supplier setup implies a retailer setup in the same time period: $y_t^S = 1$ implies $y_t^R = 1$ for $t \in \mathcal{T}$. This property holds for the n -plan generation problem, as shown in Lemma 9.

Lemma 9. *For an optimal solution for the n -plan generation problem, the joint lot-sizing plan must be nested.*

The properties in Lemmas 8 and 9 imply that the main focus of the remaining analysis is the retailer's plan. In particular, how does the constraint on either the retailer setups or the retailer inventory affect the solution structure? We continue with analysing the n -plan generation problem separately for the two private information cases.

3.3.2 Private setup cost

In this section we prove that for private setup cost the plan generation problem can be solved in polynomial time by a dynamic-programming algorithm. An essential part of this algorithm is that we can decompose an optimal solution of the n -plan generation into independent subplans. An *independent subplan*, denoted by $(\underline{t}, \bar{t}, n)$, only considers the subproblem with time periods $\{\underline{t}, \dots, \bar{t}\} \subseteq \mathcal{T}$. It has a single supplier setup, in the initial time period \underline{t} , from which exactly all demand $\sum_{t=\underline{t}}^{\bar{t}} d_t$ is satisfied. Also, there is no inventory transferred to/from time periods not belonging to the subproblem. Finally, the subplan must have exactly n retailer setups. The decomposable structure into independent subplans is proven in Lemma 10.

Lemma 10. *Any optimal solution of the n -plan generation problem can be decomposed into independent subplans.*

The result of Lemma 10 implies that the optimal solution of the n -plan generation problem can be found by solving several appropriately chosen subproblems independently. In order to solve such a subproblem we need to determine the structure of its optimal solutions. The next result, Lemma 11, shows that the structure depends on whether $H \leq h$ or $H > h$.

Lemma 11. *Consider an optimal independent subplan prescribing exactly n retailer setups. If $H \leq h$ then this subplan satisfies the zero-inventory property (without loss of optimality if $H = h$). If $H > h$ then this subplan is unique: the retailer has setups only in the first n periods, where the post-initial orders are 1 unit of supply. In this case, the retailer's plan might not satisfy the zero-inventory property.*

Lemma 11 states that if $H \leq h$ the optimal solution satisfies the zero-inventory property. Hence, this case is similar to traditional joint lot-sizing problems and can be solved by dynamic programming. However, if $H > h$ there is a unique and straightforward optimal solution, which might violate the zero-inventory property. See Figure 2 for an example with $T = 5$ and $n = 3$. This figure is a *network flow graph*, where the arrows indicate strictly positive flow of products through the supply chain. That is, a vertical arrow is a setup and a horizontal arrow implies having inventory at that time period. The upper layer is the supplier's lot-sizing plan and the lower layer the retailer's plan. At the bottom the time

periods are displayed. In Appendix B we give an example where the (unique) optimal menu contains such a contract that violates the zero-inventory property.




Figure 2: The unique optimal subplan in the case of private setup cost, $H > h$, $T = 5$, and $n = 3$.

From Lemmas 10 and 11 the dynamic-programming approach should be clear. First, solve all related independent subproblems and then use these optimal subplans to construct an optimal solution for the n -plan generation problem by dynamic programming. Since we need to solve all n -plan generation problems ($n \in \{1, \dots, T\}$) we can reuse many computations. The approach is similar to the dynamic-programming algorithm in Zangwill (1969), but we need to fix the number of retailer setups and take the two cases $H \leq h$ and $H > h$ into account. Theorem 12 concludes these insights and its proof contains the specification of the dynamic-programming algorithm.

Theorem 12. *Solving all n -plan generation problems can be done in $\mathcal{O}(T^4)$ time by dynamic programming.*

To conclude the private setup cost case, we can use the dynamic-programming algorithm stated in the proof of Theorem 12 to construct a list of procurement plans that is sufficient for optimality for the contracting problem. This list can then be used in the plan-assignment stage to determine the optimal allocation of contracts to the retailer types and solve the contracting problem. In particular, if ω is a uniform distribution, the entire contracting problem can be solved to optimality in polynomial time by Theorems 7 and 12. We state this result in the next theorem.

Theorem 13. *If ω is a uniform distribution, then the contracting problem can be solved in polynomial time.*

3.3.3 Private holding cost

The case that the holding cost is private information appears to be more complicated than having private setup cost. In particular, a similar result as Lemma 10 does not hold for the n -plan generation problem in general. For example, certain amounts of retailer inventory (n values) cannot be achieved with plans that satisfy the zero-inventory property. Furthermore, if the supplier’s setup cost F is appropriately chosen, then it would be optimal to have several supplier setups in such a (sub)plan, disproving that the decomposition structure holds in general. The smallest example is $T = 2$, $d_1 = 1$, $d_2 = 2$, $F < H$, and the optimal plan for $n = 1$, see Figure 3c.




Figure 3: The optimal solutions for the n -plan generation problems in the case of private holding cost with $T = 2$, $d_1 = 1$, and $d_2 = 2$.

Moreover, we potentially need to solve pseudo-polynomially many n -plan generation problems by Corollary 3. We have not been able to determine an efficient combinatorial algorithm to solve all n -plan generation problems. However, there seems to be a redundancy in the complete list of plans. For example, the plans in Figures 3c and 3e lead to the same supplier’s lot-sizing costs, but the retailer’s lot-sizing costs are lower for the plan for $n = 0$ (Figure 3e). In other words, the plan of Figure 3c has a form of inefficiency. Unfortunately, it is not directly clear whether we can omit this plan, since the side payments need to be taken into account.

If we assume a uniform distribution for the retailer's type, then the following lemma provides further indications that certain plans are redundant. Lemma 14 states a necessary condition for assigning a plan with a slope that does not occur in ϕ^* .

Lemma 14. *Assume that ω is a uniform distribution. Consider $k \in \{2, \dots, K-1\}$ such that ϕ^* has no slopes n^* with $n_{k-1} > n^* > n_{k+1}$. If plan k is assigned, i.e., $\theta_k > \underline{\theta}_k$, then the following must hold:*

$$\frac{(C_k + c_k^{\text{pub}}) - (C_{k-1} + c_{k-1}^{\text{pub}})}{n_{k-1} - n_k} + \frac{(C_k + c_k^{\text{pub}}) - (C_{k+1} + c_{k+1}^{\text{pub}})}{n_k - n_{k+1}} < 0. \quad (27)$$

In particular, if we consider *all* possible plans ($n_{k+1} = n_k - 1$), Lemma 14 implies that an assigned plan $k \in \mathcal{K}$ such that n_k is not a slope of ϕ^* must satisfy

$$2(C_k + c_k^{\text{pub}}) - (C_{k-1} + c_{k-1}^{\text{pub}}) - (C_{k+1} + c_{k+1}^{\text{pub}}) < 0.$$

We can apply this to the example in Figure 3. Realise that $n = 1$ is not a slope of ϕ^* . The condition of Lemma 14 for $n = 1$ is

$$\begin{aligned} 0 &> 2(F + \min\{F, H\} + 2f) - (F + f) - (F + \min\{F, 2H\} + 2f) \\ &= f + 2 \min\{F, H\} - \min\{F, 2H\} \geq f > 0. \end{aligned}$$

This is a contradiction. Hence, the plan for $n = 1$ is never assigned.

Under the additional assumption that the supplier's setup cost F is high enough to prevent any additional supplier setups, we can solve the plan generation problem in polynomial time. In this case, all optimal plans have exactly a single supplier setup. The idea is to use Lemma 14 to exclude many plans and show that the remaining plans can be determined efficiently. Of particular interest are so-called extreme plans. We call a plan *m-extreme* if it has minimal retailer inventory with m retailer setups, where $m \in \{1, \dots, T\}$. It is trivial that these extreme plans must satisfy the zero-inventory property. Note that all default plans of ϕ^* are extreme plans. Lemma 15 shows that under the mentioned assumptions T extreme plans are sufficient for optimality for the contracting problem.

Lemma 15. *Assume that ω is a uniform distribution and $F > H \max_{\tau \in \mathcal{T}} \{(\tau - 1) \sum_{t=\tau}^T d_t\}$. A list consisting of an *m-extreme* plan for each $m \in \{1, \dots, T\}$ is sufficient for optimality for the contracting problem.*

Under the assumptions of Lemma 15, all *m-extreme* plans for fixed m have the same supplier's costs and retailer's public costs. Consequently, it is sufficient to determine any *m-extreme* plan. In this case, we can solve the (entire) plan generation problem by determining these T extreme plans, which can be done by dynamic programming. The result is a polynomial-time algorithm for the plan generation under the specified assumptions, see Lemma 16.

Lemma 16. *Assume that ω is a uniform distribution and $F > H \max_{\tau \in \mathcal{T}} \{(\tau - 1) \sum_{t=\tau}^T d_t\}$. Generating plans sufficient for optimality for the contracting problem can be done in $\mathcal{O}(T^3)$ time by dynamic programming.*

By combining Theorem 7 and Lemmas 15 and 16, we conclude that under the stated conditions the contracting problem can be solved in polynomial time, see Corollary 17.

Corollary 17. *If ω is a uniform distribution and $F > H \max_{\tau \in \mathcal{T}} \{(\tau - 1) \sum_{t=\tau}^T d_t\}$, then the contracting problem can be solved in polynomial time.*

From numerical experiments we have indications that similar results hold without the condition on the supplier's setup cost F . Furthermore, we have the following property. Consider a list containing the plans for *all* slopes of ϕ^* . Now keep all side payments fixed and focus on the plans with slopes different from ϕ^* . These other plans share a special property: they can always be removed from a feasible menu to obtain a new feasible menu (without changing the side payments). This property is obvious from the graphical interpretation and can potentially be used to exclude plans from consideration. To conclude, we conjecture that plans with the same slopes as ϕ^* are essential for optimality for the contracting problem. However, more research needs to be done for a formal proof and for other distributions for ω .

4 Discussion and conclusion

The modelling concept and solution approach is applicable to a broader range of problems. In this section, we discuss the generalisability of our results, propose research directions, and conclude our findings.

4.1 Demand assumption

One of our assumptions is that the demand in each period is strictly positive. This is not without loss of generality, especially due to the time-independent holding costs. We will discuss the consequences if demand can be zero.

First, we often use that the number of retailer setups lies between 1 and T . Instead, there is a maximum feasible number of retailer setups $1 \leq K \leq T$. This has no significant impact on the results. Second, the dynamic-programming algorithms for both plan generation problems need to be adjusted slightly to prevent dummy retailer setups. Consequently, fewer options need to be considered during the algorithm, so the complexity results still hold.

Only our results for the plan generation for private setup cost are significantly affected. In the proofs of Lemmas 10 and 11 we have explicitly mentioned where we use that demand is strictly positive. There is a very specific case for which the two stated proofs do not hold if demand can be zero: there needs to be a substructure that violates the zero-inventory property where this retailer inventory cannot be decreased without invalidating a retailer setup. All details are provided in Appendix C, which we summarise here.

A common assumption for lot-sizing problems is that value is added to the product as it moves downstream in the supply chain, increasing the holding cost. In other words, $H \leq h$ holds. Another interpretation is that the supplier benefits from economies of scale to have less holding costs. In Appendix C.1 we prove for the case $H \leq h$ that without loss of optimality a plan is assigned in an optimal menu only if it satisfies the zero-inventory property. We conclude that, when demand is non-negative and $H \leq h$, the plan generation problem is solvable in $\mathcal{O}(T^4)$ time by dynamic programming.

The other case, $H > h$, can be analysed using techniques similar to those in our proofs. The (unique) optimal n -plan can be non-decomposable, as shown in Appendix C.2. However, we show that an optimal plan consists of substructures similar to Figure 2. That is, the solution is fixed when we know the supplier setups and how many retailer setups occur in between supplier setups. This allows for a dedicated dynamic-programming algorithm with $\mathcal{O}(T^5)$ running time, which can potentially be improved.

We conclude that our results are still valid when demand is non-negative, albeit that some modifications are needed.

4.2 Generalisability

For several of our results we did not use any property of the lot-sizing problem, implying that these results are also valid for other problems. Here, we discuss the generalisability of our approach. We still refer to the two involved parties by the supplier and the retailer.

The more general setting is as follows. Given the decision variables (the plan) of the supplier, the retailer needs to solve a mixed integer linear programming (MILP) problem to determine his optimal plan, and vice versa. Both the supplier and the retailer want to minimise costs as their objective. We assume that for any retailer plan there exists a feasible supplier plan, in order to have a well-defined default option.

The retailer has single-dimensional private information $\theta \in [\underline{\theta}, \bar{\theta}] = \Theta \subseteq \mathbb{R}$, which must only affect his objective value (his costs). Let \mathbf{x} denote all decision variables from the supplier and the retailer. The retailer's costs for type $\theta \in \Theta$ are

$$\phi(\mathbf{x}|\theta) = (a^\top \mathbf{x} + a_0)\theta + (b^\top \mathbf{x} + b_0),$$

for given vectors a and b , and scalars a_0 and b_0 . Hence, we have public costs $c^{\text{pub}} = b^\top \mathbf{x} + b_0$ and slope $n = a^\top \mathbf{x} + a_0$. We assume that this slope $a^\top \mathbf{x} + a_0$ only takes on finitely many values over all feasible plans \mathbf{x} , which is essential. Consequently, the retailer's default option ϕ^* is the lower envelope of finitely many linear functions.

Under the described setting and assumptions, we can apply the same two-stage solution approach consisting of a plan-generation stage and a plan-assignment stage. In particular, all results related to the graphical interpretation and the plan assignment are valid, since these are independent of the lot-sizing setting. The main difference in the difficulty in solving the contracting problem lies in the plan generation. If the bi-level optimisation problem with the additional constraint that $(a^\top \mathbf{x} + a_0) = n$ for some slope n can be solved by a single-level MILP problem, then general MILP solvers can be used. Whether polynomial-time algorithms exist highly depends on the underlying optimisation problem. Also, if the number of plans to generate is high, it might be useful to use heuristics instead, which we will discuss below.

To conclude the generalisability, we provide the details on which results are still valid. First, since the default option ϕ^* is the lower envelope of finitely many linear functions, equivalent properties as in Lemma 1 hold (except for the complexity result). By trivially modifying the cost functions in our proofs, we obtain the following similar results. As in Lemma 2, the slopes of the assigned plans must be strictly decreasing. By the finitely many possible slopes, we get a bound on the number of contracts and properties similar to Corollary 3. In the equivalent of Theorem 4, the slope $a^\top \mathbf{x} + a_0$ must be fixed to n and the joint optimisation problem minimises the sum of the supplier’s costs and the retailer’s public costs $b^\top \mathbf{x} + b_0$. Regarding the plan assignment, Lemmas 5, 6, and 14, and Theorem 7 are valid, since they do not use any lot-sizing properties. The other results concern the plan generation and are specific to the lot-sizing problem.

4.3 Heuristics

Our analysis provides several research directions for heuristics. First of all, the plan generation can be restricted to (potentially suboptimal) plans that can be constructed efficiently, e.g., lot-sizing plans satisfying the zero-inventory property. Also, we can generate plans only for slopes with intuitive interpretations. For example, only plans with slopes that appear in ϕ^* or in the optima of the unrestricted traditional joint lot-sizing problem, when a sensitivity analysis is performed on the private cost parameter.

Second, by rescaling the side payment z_k as $\tilde{z}_k = z_k - c_k^{\text{pub}}$, the constraints of the plan assignment problem only depend on the slopes of the included plans. This leads to the following idea. First, add a place-holder plan for each possible slope. Second, determine lower bounds on the joint costs $C_k + c_k^{\text{pub}}$ for each $k \in \mathcal{K}$. These costs appear in the objective after rescaling the side payment. This results in a relaxation for the contracting problem. Then, we determine the exact joint costs for (a subset of) the place-holder plans assigned in the optimum of the relaxation. If we calculate the joint costs of all assigned plans, then we obtain an upper bound. By repeating this process, we get better lower and upper bounds, and a solution approach.

Third, we can design an iterative heuristic as follows. If we fix the number of contracts and their assignment to retailer types, then the resulting optimisation model is a mixed integer linear program. This model is given in Appendix D. For a given partition, we solve this model to obtain procurement plans. Then, solve the plan assignment model for these plans, resulting in a new partition. Switching between these models leads to an iterative heuristic.

4.4 Concluding remarks

If all information is shared in the supply chain, then it is trivial to show that the supplier’s contracting problem can be solved by a single joint lot-sizing problem. Namely, compared to the model with information asymmetry in Section 2.2, there is only a single type, a single IR constraint, and no IC constraints. In any optimal solution the IR constraint binds, i.e., $z = \phi(\mathbf{x}) - \phi^*$, where we omit the single type. Substituting this in the objective function results in a traditional joint lot-sizing problem that minimises the sum of the supplier’s costs and *all* retailer’s costs. As discussed before, this problem can be solved in polynomial time.

Our analysis and obtained results show that information asymmetry does not necessarily change the complexity class of the underlying optimisation problem. That is, we have identified cases for which we can solve the contracting problem under information asymmetry in polynomial time. However, clearly it is more complicated to determine the optimal solution. Not only do we need to solve multiple joint lot-sizing problems to generate a sufficient list of plans, we also need to solve the assignment model. Furthermore, the interdependence between the contracts through the side payments results in offering atypical procurement plans. For example, plans that do not satisfy the zero-inventory property or that are not decomposable into independent subproblems.

Although the plan generation for private setup cost can be solved in polynomial time for any instance, further research might narrow down which n -plans are sufficient for optimality for the contracting problem. For the case of private holding cost more research is needed to prove polynomial-time complexity under less restrictive assumptions (if valid at all). Finally, the above described generalisability provides research directions to other, more general, problem settings.

References

- Albrecht, M. (2017). “Coordination of dynamic lot-sizing in supply chains”. *Journal of the Operational Research Society* 68.3, pp. 322–330.
- Bradley, S. P., A. C. Hax, and T. L. Magnanti (1977). *Applied mathematical programming*. Addison-Wesley Publishing Company.
- Buer, T., J. Homberger, and H. Gehring (2013). “A collaborative ant colony metaheuristic for distributed multi-level uncapacitated lot-sizing”. *International Journal of Production Research* 51.17, pp. 5253–5270.
- Burnetas, A., S. M. Gilbert, and C. E. Smith (2007). “Quantity discounts in single-period supply contracts with asymmetric demand information”. *IIE Transactions* 39.5, pp. 465–479.
- Byrd, R. H., N. I. Gould, J. Nocedal, and R. A. Waltz (2003). “An algorithm for nonlinear optimization using linear programming and equality constrained subproblems”. *Mathematical Programming* 100.1, pp. 27–48.
- Cachon, G. P. (2003). “Supply chain coordination with contracts”. *Handbooks In Operations Research and Management Science* 11, pp. 227–339.
- Cakanyildirim, M., Q. Feng, X. Gan, and S. P. Sethi (2012). “Contracting and coordination under asymmetric production cost information”. *Production and Operations Management* 21.2, pp. 345–360.
- Corbett, C. J. and X. de Groote (2000). “A supplier’s optimal quantity discount policy under asymmetric information”. *Management Science* 46.3, pp. 444–450.
- Dudek, G. and H. Stadtler (2005). “Negotiation-based collaborative planning between supply chains partners”. *European Journal of Operational Research* 163.3, pp. 668–687.
- Dudek, G. and H. Stadtler (2007). “Negotiation-based collaborative planning in divergent two-tier supply chains”. *International Journal of Production Research* 45.2, pp. 465–484.
- Inderfurth, K., A. Sadrieh, and G. Voigt (2013). “The impact of information sharing on supply chain performance under asymmetric information”. *Production and Operations Management* 22.2, pp. 410–425.
- Karush, W. (1939). “Minima of functions of several variables with inequalities as side constraints”. Master thesis. Chicago, Illinois: Department of Mathematics, University of Chicago.
- Kerckamp, R. B. O., W. van den Heuvel, and A. P. M. Wagelmans (2018). “Two-echelon supply chain coordination under information asymmetry with multiple types”. *Omega* 76, pp. 137–159.
- Kolda, T. G., R. M. Lewis, and V. Torczon (2007). “Stationarity results for generating set search for linearly constrained optimization”. *SIAM Journal on Optimization* 17.4, pp. 943–968.
- Kuhn, H. W. and A. W. Tucker (1951). “Nonlinear programming”. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley: University of California Press, pp. 481–492.
- Laffont, J.-J. and D. Martimort (2002). *The theory of incentives: the principal-agent model*. Princeton University Press.
- Leng, M. and M. Parlar (2005). “Game theoretic applications in supply chain management: a review”. *INFOR: Information Systems and Operational Research* 43.3, pp. 187–220.
- Mobini, Z., W. van den Heuvel, and A. P. M. Wagelmans (2014). “Designing multi-period supply contracts in a two-echelon supply chain with asymmetric information”. *Econometric Institute Research Papers* (No. EI2014-28).
- Perakis, G. and G. Roels (2007). “The price of anarchy in supply chains: quantifying the efficiency of price-only contracts”. *Management Science* 53.8, pp. 1249–1268.
- Phouratsamay, S.-L. (2017). “Coordination des décisions de planification dans une chaîne logistique”. PhD thesis. Paris, France: University Pierre et Marie Curie.
- Pishchulov, G. and K. Richter (2016). “Optimal contract design in the joint economic lot size problem with multi-dimensional asymmetric information”. *European Journal of Operational Research* 253.3, pp. 711–733.
- Richter, K. and J. Vörös (1989). “On the stability region for multi-level inventory problems”. *European Journal of Operational Research* 41.2, pp. 169–173.
- Van Hoesel, C. P. M. and A. P. M. Wagelmans (2000). “Parametric analysis of setup cost in the economic lot-sizing model without speculative motives”. *International Journal of Production Economics* 66.1, pp. 13–22.
- Wagner, H. M. and T. M. Whitin (1958). “Dynamic version of the economic lot size model”. *Management Science* 5.1, pp. 89–96.

- Ye, Y. and E. Tse (1989). “An extension of Karmarkar’s projective algorithm for convex quadratic programming”. *Mathematical Programming* 44.1, pp. 157–179.
- Zangwill, W. I. (1969). “A backloging model and a multi-echelon model of a dynamic economic lot size production”. *Management Science* 15.9, pp. 506–527.

A Proofs of Section 3

This appendix contains all proofs of Section 3. Note that many results also have an intuitive graphical interpretation for which we refer to the main text.

A.1 Proofs of Section 3.1

Proof of Lemma 1. The default option ϕ^* follows from the single-level lot-sizing problem on the retailer's level. It is trivial that there are finitely many feasible procurement plans and each feasible plan is linear and non-decreasing in θ . By definition, ϕ^* is the point-wise minimum (the lower envelope) of these finitely many linear functions. Only the plans that for a given number of setups (ranging from 1 to T) minimise the retailer's holding costs can be minimisers and form the lower envelope. Consequently, these plans must satisfy the zero-inventory property. This proves the stated properties of ϕ^* . Finally, ϕ^* can be determined efficiently in $\mathcal{O}(T^2)$ time by using the method described in Van Hoesel and Wagelmans (2000). \square

Proof of Lemma 2. First, realise that by definition we have

$$\begin{aligned} \phi(\mathbf{x}(\theta)|\theta) - \phi(\mathbf{x}(\hat{\theta})|\theta) + \phi(\mathbf{x}(\hat{\theta})|\hat{\theta}) - \phi(\mathbf{x}(\theta)|\hat{\theta}) &= (f(\theta) - f(\hat{\theta})) \left(\sum_{t \in \mathcal{T}} y_t^R(\theta) - \sum_{t \in \mathcal{T}} y_t^R(\hat{\theta}) \right) \\ &\quad + (h(\theta) - h(\hat{\theta})) \left(\sum_{t \in \mathcal{T}} I_t^R(\theta) - \sum_{t \in \mathcal{T}} I_t^R(\hat{\theta}) \right). \end{aligned} \quad (28)$$

Hence, if

$$\begin{cases} \sum_{t \in \mathcal{T}} y_t^R(\theta) = \sum_{t \in \mathcal{T}} y_t^R(\hat{\theta}) & \text{if setup cost } f \text{ is private} \\ \sum_{t \in \mathcal{T}} I_t^R(\theta) = \sum_{t \in \mathcal{T}} I_t^R(\hat{\theta}) & \text{if holding cost } h \text{ is private} \end{cases} \quad (29)$$

then the right-hand side of (28) is equal to zero and

$$\phi(\mathbf{x}(\theta)|\theta) - \phi(\mathbf{x}(\hat{\theta})|\theta) = \phi(\mathbf{x}(\theta)|\hat{\theta}) - \phi(\mathbf{x}(\hat{\theta})|\hat{\theta}).$$

Second, the IC conditions state

$$\phi(\mathbf{x}(\theta)|\theta) - z(\theta) \leq \phi(\mathbf{x}(\hat{\theta})|\theta) - z(\hat{\theta}), \quad \phi(\mathbf{x}(\hat{\theta})|\hat{\theta}) - z(\hat{\theta}) \leq \phi(\mathbf{x}(\theta)|\hat{\theta}) - z(\theta),$$

implying that

$$\phi(\mathbf{x}(\theta)|\theta) - \phi(\mathbf{x}(\hat{\theta})|\theta) \leq z(\theta) - z(\hat{\theta}) \leq \phi(\mathbf{x}(\theta)|\hat{\theta}) - \phi(\mathbf{x}(\hat{\theta})|\hat{\theta}). \quad (30)$$

So if (29) is true, then (30) holds with equalities, resulting in

$$\phi(\mathbf{x}(\theta)|\theta) - z(\theta) = \phi(\mathbf{x}(\hat{\theta})|\theta) - z(\hat{\theta}), \quad \phi(\mathbf{x}(\hat{\theta})|\hat{\theta}) - z(\hat{\theta}) = \phi(\mathbf{x}(\theta)|\hat{\theta}) - z(\theta).$$

In other words, both types θ and $\hat{\theta}$ are indifferent to each other's contracts and these contracts can be interchanged without affecting feasibility.

Now consider an optimal menu of contracts. If (29) holds for types θ and $\hat{\theta}$, then assigning both types either contract $(\mathbf{x}(\theta), z(\theta))$ or $(\mathbf{x}(\hat{\theta}), z(\hat{\theta}))$ is feasible as shown above. Assigning the contract that leads to the lowest supplier's net costs cannot result in a worse objective value, i.e., the new menu must be optimal as well. By repeating this argument, we conclude that without loss of optimality distinct contracts do not satisfy (29).

Finally, from (30) it follows that (28) must be non-positive. For types $\theta < \hat{\theta}$ with distinct contracts we have either $f(\theta) < f(\hat{\theta})$ or $h(\theta) < h(\hat{\theta})$, depending on which cost parameter is private. Furthermore, with the above insight (28) must be strictly negative (without loss of optimality). Thus, (13) holds without loss of optimality. \square

Proof of Corollary 3. By Lemma 2 we can bound the number of distinct contracts. First, the total number of retailer setups lies between 1 and T . Since $d_t > 0$ for all $t \in \mathcal{T}$, all numbers $1, \dots, T$ of retailer setups are feasible. Second, the total retailer inventory lies between 0 (use the maximum number

of setups) and $\sum_{t \in \mathcal{T}} (t-1)d_t$ (use one setup). All discrete intermediate values are also feasible by appropriately delaying parts of the orders (starting with a single setup). Here, we use our assumption that the products are indivisible, i.e., the retailer order quantities must be discrete. Consequently, by the finite bounds given above and the discrete nature, Lemma 2 implies that there are only finitely many contracts in an optimal menu (without loss of optimality). Hence, retailer types must be pooled, i.e., some are offered the same contract.

The partitioning of the retailer types follows trivially from the ordering implied by (13). Only the technicality that we can use closed subintervals remains to be shown. Consider the case that the k -th contract (\mathbf{x}_k, z_k) is the most preferred contract for all types $(\underline{\theta}_k, \bar{\theta}_k]$, but not for type $\underline{\theta}_k$. Instead, type $\underline{\theta}_k$ strictly prefers the l -th contract (\mathbf{x}_l, z_l) :

$$\phi(\mathbf{x}_l|\underline{\theta}_k) - z_l < \phi(\mathbf{x}_k|\underline{\theta}_k) - z_k.$$

However, we also have

$$\begin{aligned} & \phi(\mathbf{x}_k|\theta_k) - z_k \leq \phi(\mathbf{x}_l|\theta_k) - z_l & \forall \theta_k \in (\underline{\theta}_k, \bar{\theta}_k] \\ \implies & \phi(\mathbf{x}_k|\underline{\theta}_k) - \phi(\mathbf{x}_l|\underline{\theta}_k) = \lim_{\theta_k \rightarrow \underline{\theta}_k} \{\phi(\mathbf{x}_k|\theta_k) - \phi(\mathbf{x}_l|\theta_k)\} \leq z_k - z_l. \end{aligned}$$

Here, the continuity of the limit follows from the fact that $\phi(\mathbf{x}_k|\theta_k) - \phi(\mathbf{x}_l|\theta_k)$ is continuous in θ_k . The result contradicts that $\underline{\theta}_k$ strictly prefers the l -th contract. Similar arguments hold for the other cases. Thus, each subinterval is closed. \square

Proof of Theorem 4. Consider a feasible menu and any of its contracts $(\mathbf{x}(\theta), z(\theta))$. Modify the lot-sizing variables of contract $(\mathbf{x}(\theta), z(\theta))$, resulting in $\bar{\mathbf{x}}(\theta)$, in any way such that it is a feasible lot-sizing plan satisfying (2)-(9) and (14). Adjust the side payment to $\bar{z}(\theta)$ to compensate for the change in costs for type θ :

$$\bar{z}(\theta) = z(\theta) + \phi(\bar{\mathbf{x}}(\theta)|\theta) - \phi(\mathbf{x}(\theta)|\theta).$$

By construction, we have

$$\begin{aligned} \phi(\bar{\mathbf{x}}(\theta)|\theta) - \bar{z}(\theta) &= \phi(\mathbf{x}(\theta)|\theta) - z(\theta) \leq \phi^*(\theta), \\ \phi(\bar{\mathbf{x}}(\theta)|\theta) - \bar{z}(\theta) &= \phi(\mathbf{x}(\hat{\theta})|\theta) - z(\theta) \leq \phi(\mathbf{x}(\hat{\theta})|\theta) - z(\hat{\theta}), \quad \forall \hat{\theta} \in \Theta. \end{aligned}$$

Furthermore, for $\hat{\theta} \in \Theta$ we get

$$\begin{aligned} \phi(\bar{\mathbf{x}}(\theta)|\hat{\theta}) - \bar{z}(\theta) &= \phi(\bar{\mathbf{x}}(\theta)|\hat{\theta}) - \phi(\bar{\mathbf{x}}(\theta)|\theta) + \phi(\mathbf{x}(\theta)|\theta) - \phi(\mathbf{x}(\theta)|\hat{\theta}) + \phi(\mathbf{x}(\theta)|\hat{\theta}) - z(\theta) \\ &= (f(\hat{\theta}) - f(\theta)) \left(\sum_{t \in \mathcal{T}} \bar{y}_t^R(\theta) - \sum_{t \in \mathcal{T}} y_t^R(\theta) \right) \\ &\quad + (h(\hat{\theta}) - h(\theta)) \left(\sum_{t \in \mathcal{T}} \bar{I}_t^R(\theta) - \sum_{t \in \mathcal{T}} I_t^R(\theta) \right) \\ &\quad + \phi(\mathbf{x}(\theta)|\hat{\theta}) - z(\theta) \\ &= \phi(\mathbf{x}(\theta)|\hat{\theta}) - z(\theta) \geq \phi(\mathbf{x}(\hat{\theta})|\hat{\theta}) - z(\hat{\theta}). \end{aligned}$$

Here, the last equality holds since both plans satisfy (14) and the inequality follows from the IC constraints. To conclude, the menu with the modified contract is feasible.

Finally, consider an optimal menu. We can modify each contract sequentially as described above, where the corresponding term in the objective is

$$\sum_{t \in \mathcal{T}} \left(F\bar{y}_t^S(\theta) + H\bar{I}_t^S(\theta) \right) + \bar{z}(\theta) = \sum_{t \in \mathcal{T}} \left(F\bar{y}_t^S(\theta) + H\bar{I}_t^S(\theta) \right) + z(\theta) + \phi(\bar{\mathbf{x}}(\theta)|\theta) - \phi(\mathbf{x}(\theta)|\theta),$$

where $z(\theta)$ and $\phi(\mathbf{x}(\theta)|\theta)$ are now constants. Furthermore, in

$$\phi(\bar{\mathbf{x}}(\theta)|\theta) = f(\theta) \sum_{t \in \mathcal{T}} \bar{y}_t^R(\theta) + h(\theta) \sum_{t \in \mathcal{T}} \bar{I}_t^R(\theta)$$

one of these terms is constant (equal to n times the retailer's type) by (14) and the other term does not depend on the retailer's type. The result now follows. \square

A.2 Proofs of Section 3.2

Proof of Lemma 5. For $k \in \mathcal{K}$ the IR and IC constraints must hold for all types $\theta_k \in [\underline{\theta}_k, \bar{\theta}_k]$ with respect to contract (\mathbf{x}_k, z_k) . First, we consider the IC constraints. For $k \in \mathcal{K} \setminus \{1\}$ two (adjacent) IC constraints state

$$\phi(\mathbf{x}_k|\underline{\theta}_k) - \phi(\mathbf{x}_{k-1}|\underline{\theta}_k) \leq z_k - z_{k-1} \leq \phi(\mathbf{x}_k|\bar{\theta}_{k-1}) - \phi(\mathbf{x}_{k-1}|\bar{\theta}_{k-1}).$$

Since $\bar{\theta}_{k-1} = \underline{\theta}_k$, equality holds throughout and we obtain

$$z_k - z_{k-1} = \phi(\mathbf{x}_k|\underline{\theta}_k) - \phi(\mathbf{x}_{k-1}|\underline{\theta}_k). \quad (31)$$

This shows necessity of (31) and we continue with proving its sufficiency. We denote the lot-sizing variables of retailer plan \mathbf{x}_i by $y_t^{R(i)}$ and $I_t^{R(i)}$. For $l < k$ we have

$$\begin{aligned} z_k - z_l &= \sum_{i=l+1}^k (z_i - z_{i-1}) \stackrel{(31)}{=} \sum_{i=l+1}^k (\phi(\mathbf{x}_i|\underline{\theta}_i) - \phi(\mathbf{x}_{i-1}|\underline{\theta}_i)) \\ &= \sum_{i=l+1}^k \left(f(\underline{\theta}_i) \left(\sum_{t \in \mathcal{T}} y_t^{R(i)} - \sum_{t \in \mathcal{T}} y_t^{R(i-1)} \right) + h(\underline{\theta}_i) \left(\sum_{t \in \mathcal{T}} I_t^{R(i)} - \sum_{t \in \mathcal{T}} I_t^{R(i-1)} \right) \right) \\ &\stackrel{(13)}{\leq} \sum_{i=l+1}^k \left(f(\theta_l) \left(\sum_{t \in \mathcal{T}} y_t^{R(i)} - \sum_{t \in \mathcal{T}} y_t^{R(i-1)} \right) + h(\theta_l) \left(\sum_{t \in \mathcal{T}} I_t^{R(i)} - \sum_{t \in \mathcal{T}} I_t^{R(i-1)} \right) \right) \\ &= \phi(\mathbf{x}_k|\theta_l) - \phi(\mathbf{x}_l|\theta_l), \end{aligned}$$

for any $\theta_l \in [\underline{\theta}_l, \bar{\theta}_l]$. Likewise, for $l > k$ we get

$$\begin{aligned} z_k - z_l &= - \sum_{i=k+1}^l (z_i - z_{i-1}) \stackrel{(31)}{=} - \sum_{i=k+1}^l (\phi(\mathbf{x}_i|\underline{\theta}_i) - \phi(\mathbf{x}_{i-1}|\underline{\theta}_i)) \\ &= - \sum_{i=k+1}^l \left(f(\underline{\theta}_i) \left(\sum_{t \in \mathcal{T}} y_t^{R(i)} - \sum_{t \in \mathcal{T}} y_t^{R(i-1)} \right) + h(\underline{\theta}_i) \left(\sum_{t \in \mathcal{T}} I_t^{R(i)} - \sum_{t \in \mathcal{T}} I_t^{R(i-1)} \right) \right) \\ &\stackrel{(13)}{\leq} - \sum_{i=k+1}^l \left(f(\theta_l) \left(\sum_{t \in \mathcal{T}} y_t^{R(i)} - \sum_{t \in \mathcal{T}} y_t^{R(i-1)} \right) + h(\theta_l) \left(\sum_{t \in \mathcal{T}} I_t^{R(i)} - \sum_{t \in \mathcal{T}} I_t^{R(i-1)} \right) \right) \\ &= \phi(\mathbf{x}_k|\theta_l) - \phi(\mathbf{x}_l|\theta_l), \end{aligned}$$

for any $\theta_l \in [\underline{\theta}_l, \bar{\theta}_l]$. Thus, all IC constraints are implied by (31).

Second, the IR constraints for $k \in \mathcal{K}$ are

$$\begin{aligned} &\phi(\mathbf{x}_k|\theta_k) - z_k \leq \phi^*(\theta_k) \quad \forall \theta_k \in [\underline{\theta}_k, \bar{\theta}_k] \\ \iff &\sup\{\phi(\mathbf{x}_k|\theta_k) - \phi^*(\theta_k) : \theta_k \in [\underline{\theta}_k, \bar{\theta}_k]\} \leq z_k. \end{aligned}$$

Since $\phi^*(\theta_k)$ is concave by Lemma 1 and $\phi(\mathbf{x}_k|\theta_k)$ is linear in θ_k , the difference $\phi(\mathbf{x}_k|\theta_k) - \phi^*(\theta_k)$ is a convex function in θ_k . The stated supremum is therefore attained at one of the border points $\underline{\theta}_k$ or $\bar{\theta}_k$ and the other IR constraints are redundant. In fact, more IR constraints are redundant, provided that the IC constraints hold. For $k \in \mathcal{K} \setminus \{K\}$ we have

$$\phi(\mathbf{x}_{k+1}|\underline{\theta}_{k+1}) - z_{k+1} \stackrel{(31)}{=} \phi(\mathbf{x}_k|\underline{\theta}_{k+1}) - z_k = \phi(\mathbf{x}_k|\bar{\theta}_k) - z_k.$$

This implies that we only need one of the IR constraints corresponding to types $\bar{\theta}_k$ and $\underline{\theta}_{k+1}$. \square

Proof of Lemma 6. Suppose $\underline{\theta}_k = \bar{\theta}_k$, i.e., the k -th contract is effectively not assigned and is redundant. First, we consider the IC constraints. From (16) we have

$$\begin{aligned} z_{k+1} - z_k &= \phi(\mathbf{x}_{k+1}|\underline{\theta}_{k+1}) - \phi(\mathbf{x}_k|\underline{\theta}_{k+1}), \\ z_k - z_{k-1} &= \phi(\mathbf{x}_k|\underline{\theta}_k) - \phi(\mathbf{x}_{k-1}|\underline{\theta}_k). \end{aligned}$$

Adding both equalities and using $\theta_k = \bar{\theta}_k = \underline{\theta}_{k+1}$ results in

$$\begin{aligned} z_{k+1} - z_{k-1} &= \phi(\mathbf{x}_{k+1}|\underline{\theta}_{k+1}) - \phi(\mathbf{x}_k|\underline{\theta}_{k+1}) + \phi(\mathbf{x}_k|\underline{\theta}_k) - \phi(\mathbf{x}_{k-1}|\underline{\theta}_k) \\ &= \phi(\mathbf{x}_{k+1}|\underline{\theta}_{k+1}) - \phi(\mathbf{x}_k|\underline{\theta}_{k+1}) + \phi(\mathbf{x}_k|\underline{\theta}_{k+1}) - \phi(\mathbf{x}_{k-1}|\underline{\theta}_{k+1}) \\ &= \phi(\mathbf{x}_{k+1}|\underline{\theta}_{k+1}) - \phi(\mathbf{x}_{k-1}|\underline{\theta}_{k+1}). \end{aligned}$$

This is the IC constraint if we would omit the redundant k -th contract from the model.

Second, consider the IR constraints. If $k \in \mathcal{K} \setminus \{K\}$ we have

$$\phi(\mathbf{x}_k|\underline{\theta}_k) - z_k = \phi(\mathbf{x}_k|\bar{\theta}_k) - z_k = \phi(\mathbf{x}_k|\underline{\theta}_{k+1}) - z_k \stackrel{(16)}{=} \phi(\mathbf{x}_{k+1}|\underline{\theta}_{k+1}) - z_{k+1}.$$

Thus, the corresponding IR constraints are implied by others and have no effect. Similarly, if $k = K$ it holds that

$$\phi(\mathbf{x}_K|\bar{\theta}_K) - z_K = \phi(\mathbf{x}_K|\underline{\theta}_K) - z_K = \phi(\mathbf{x}_K|\bar{\theta}_{K-1}) - z_K \stackrel{(16)}{=} \phi(\mathbf{x}_{K-1}|\bar{\theta}_{K-1}) - z_{K-1}.$$

Notice that this leads exactly to the second IR constraint required for type $K-1$ if we would omit the K -th contract from the model. We conclude that the redundant contract has no effect on the feasible region.

Finally, the probability that the k -th contract is selected by the retailer is $\int_{\underline{\theta}_k}^{\bar{\theta}_k} \omega(\theta) d\theta$. If $\underline{\theta}_k = \bar{\theta}_k$ then this probability is zero. Furthermore, other contracts are unaffected as argued above. Therefore, such redundant contracts do not affect the optimisation problem. \square

Proof of Theorem 7. First, define $\bar{\theta}_0 = \underline{\theta}$ for notational convenience. Since ω is a uniform distribution, the objective function in (17) becomes

$$\sum_{k \in \mathcal{K}} \frac{\bar{\theta}_k - \bar{\theta}_{k-1}}{\bar{\theta} - \underline{\theta}} (C_k + z_k) = \sum_{k=1}^{K-1} \frac{(C_k - C_{k+1} + z_k - z_{k+1})\bar{\theta}_k}{\bar{\theta} - \underline{\theta}} + \frac{(C_K + z_K)\bar{\theta} - (C_1 + z_1)\underline{\theta}}{\bar{\theta} - \underline{\theta}}.$$

We can use (21) to get

$$\begin{aligned} (C_k - C_{k+1} + z_k - z_{k+1})\bar{\theta}_k &= (C_k - C_{k+1} + c_k^{\text{pub}} - c_{k+1}^{\text{pub}} + (n_k - n_{k+1})\underline{\theta}_{k+1})\bar{\theta}_k \\ &= (n_k - n_{k+1})\bar{\theta}_k^2 + (C_k + c_k^{\text{pub}} - C_{k+1} - c_{k+1}^{\text{pub}})\bar{\theta}_k. \end{aligned}$$

Furthermore, we can repeatedly use (21) to obtain

$$\begin{aligned} z_k &= z_1 + \sum_{i=2}^k (c_i^{\text{pub}} - c_{i-1}^{\text{pub}} - (n_{i-1} - n_i)\underline{\theta}_i) = z_1 + c_k^{\text{pub}} - c_1^{\text{pub}} - \sum_{i=2}^k (n_{i-1} - n_i)\underline{\theta}_i \\ &= z_1 + c_k^{\text{pub}} - c_1^{\text{pub}} - \sum_{i=1}^{k-1} (n_i - n_{i+1})\bar{\theta}_i. \end{aligned} \tag{32}$$

Combining all results and substituting z_K , leads to the objective function

$$\begin{aligned} z_1 - c_1^{\text{pub}} &+ \frac{(C_K + c_K^{\text{pub}})\bar{\theta} - (C_1 + c_1^{\text{pub}})\underline{\theta}}{\bar{\theta} - \underline{\theta}} \\ &+ \sum_{k=1}^{K-1} \frac{(n_k - n_{k+1})\bar{\theta}_k^2 + (C_k + c_k^{\text{pub}} - C_{k+1} - c_{k+1}^{\text{pub}} - (n_k - n_{k+1})\bar{\theta})\bar{\theta}_k}{\bar{\theta} - \underline{\theta}}. \end{aligned} \tag{33}$$

This objective is convex quadratic, as $n_k > n_{k+1}$ holds.

After eliminating unnecessary variables and using (32), the quadratic formulation is to minimise (33)

subject to

$$\begin{aligned}
n_1 \underline{\theta} + c_1^{\text{pub}} - z_1 &\leq \phi^*(\underline{\theta}), & [\lambda_0] \\
n_k \bar{\theta}_k + \sum_{i=1}^{k-1} (n_i - n_{i+1}) \bar{\theta}_i + c_1^{\text{pub}} - z_1 &\leq \phi^*(\bar{\theta}_k), & \forall k \in \{1, \dots, K-1\}, \quad [\lambda_k] \\
n_K \bar{\theta} + \sum_{i=1}^{K-1} (n_i - n_{i+1}) \bar{\theta}_i + c_1^{\text{pub}} - z_1 &\leq \phi^*(\bar{\theta}), & [\lambda_K] \\
\underline{\theta} - \bar{\theta}_1 &\leq 0, & [\mu_{0,1}] \\
\bar{\theta}_k - \bar{\theta}_{k+1} &\leq 0, & \forall k \in \{1, \dots, K-2\}, \quad [\mu_{k,k+1}] \\
\bar{\theta}_{K-1} - \bar{\theta} &\leq 0. & [\mu_{K-1,K}]
\end{aligned}$$

Note that we do not need the IC constraints in the model any more. Also, the IR constraints have been rewritten by using

$$n_k \underline{\theta}_k + \sum_{i=1}^{k-1} (n_i - n_{i+1}) \bar{\theta}_i = n_{k-1} \bar{\theta}_{k-1} + \sum_{i=1}^{k-2} (n_i - n_{i+1}) \bar{\theta}_i.$$

Finally, the default values $\phi^*(\bar{\theta}_k)$ can each be modelled using at most T linear constraints, as shown in Section 3.2.2.

Thus, if ω is a uniform distribution, then the plan assignment model can be formulated as a linearly-constrained convex-quadratic model with $(K-1) + 1$ continuous variables and $2 + (K-1)T + K$ linear constraints. Furthermore, the coefficients of the model are polynomial in the input size of the contracting problem as they relate to total costs in lot-sizing plans. This convex-quadratic model can be solved in polynomial time in the number of contracts K by interior-point methods, see for example Ye and Tse (1989). \square

A.3 Proofs of Section 3.3

Proof of Lemma 8. Suppose there exists an optimal solution where the supplier's lot-sizing plan does not satisfy the zero-inventory property. Consequently, there exists a time $\tau \in \mathcal{T}$ such that $I_{\tau-1}^S > 0$ and $y_\tau^S = 1$. We can strictly improve the contract by shifting $I_{\tau-1}^S$ units from the origin order(s) to time τ , leading to a reduction in the total costs by at least $H > 0$. This contradicts the optimality of the original lot-sizing plan. Hence, the supplier's plan must satisfy the zero-inventory property in the optimal solution. \square

Proof of Lemma 9. Suppose there is an optimal solution where $y_\tau^S = 1$ and $y_\tau^R = 0$ for some $\tau \in \mathcal{T}$. The supplier order must transfer to the retailer at some time, so it must hold that $\tau < T$. Furthermore, as $y_\tau^R = 0$ it is feasible to delay the considered supplier order. By shifting the supplier order from time τ to time $\tau + 1$ the total costs reduce by $Hx_\tau^S > 0$, contradicting optimality. \square

Proof of Lemma 10. Consider an optimal solution and suppose there exist two supplier setups part of the same subplan. Let $t_1 < t_2 \in \mathcal{T}$ be the time periods of the first of such setups: $y_{t_1}^S = y_{t_2}^S = 1$ and $y_t^S = 0$ for all $t_1 < t < t_2$. By Lemmas 8 and 9 we conclude that $I_{t_1-1}^S = I_{t_2-1}^S = 0$ and $y_{t_1}^R = y_{t_2}^R = 1$. By assumption, both setups are part of the same subplan, so it must hold that $I_{t_2-1}^R > 0$. Furthermore, $I_{t_1-1}^R = 0$ and the origin of these $I_{t_2-1}^R$ products is the supplier order at time t_1 , since times t_1 and t_2 correspond the first occurrence of the described supplier setups.

Now, realise that we can feasibly shift one unit of supply from the retailer's inventory $I_{t_2-1}^R$ to the supplier and retailer setups at time t_2 without invalidating any retailer setups in time periods t_1, \dots, t_2-1 (for details, see below). This results in a total cost reduction of at least $h > 0$, contradicting optimality. Hence, no two supplier setups can be part of the same subplan, implying that an optimal solution results in the stated decomposition into independent subplans. Note that each subplan must prescribe the number of retailer setups to ensure that the retailer uses exactly n setups in total.

The feasibility of the described shift is guaranteed by the assumption that $d_t > 0$ for all $t \in \mathcal{T}$. The details are as follows. Arbitrarily follow one unit of supply from the supplier order at time t_1 to the retailer's inventory at time $t_2 - 1$, which prevents time t_2 from being the start of a new independent subplan. Let τ ($t_1 \leq \tau < t_2$) be the time this supply is transferred to the retailer's level, so $x_\tau^R > 0$.

If $x_\tau^R = 1$ we cannot use this path to remove the considered unit of supply as that would remove the retailer setup at time τ , causing issues with the fact that the total number of retailer setups is fixed. However, by assumption we have $I_{t-1}^R + x_t^R = I_t^R + d_t > I_t^R$ for all $t \in \mathcal{T}$. In particular, $I_{\tau-1}^R + 1 = I_{\tau-1}^R + x_\tau^R = I_\tau^R + d_\tau > I_\tau^R \geq 1$, implying $I_{\tau-1}^R > 0$. Thus, there exists an alternative supply path. By repeating this argument and using $I_{t_1-1}^R = 0$, we conclude that there must exist a supply path for which $x_\tau^R > 1$. This path can be used to feasibly shift supply. \square

Proof of Lemma 11. The idea of the proof is as follows. Consider such an optimal independent subplan, which by definition must have exactly one supplier setup (in the initial period). Hence, all post-initial demand is supplied from inventory (on the supplier's or retailer's level).

If $H < h$ it is strictly optimal to shift as much inventory as possible to the supplier's level, keeping the prescribed number of retailer setups in mind. Therefore, an optimal subplan must satisfy the zero-inventory property. If $H = h$ the described shifts are weakly optimal, so the zero-inventory property holds without loss of optimality.

If $H > h$ then it is strictly optimal to shift as much inventory as possible to the retailer's level. This implies placing all prescribed retailer setups as early as possible, i.e., in the first n periods, and reducing the post-initial retailer orders to the minimum of 1 unit of products. This unique optimal subplan satisfies the zero-inventory property only under very specific conditions.

We continue with a formal proof of the above argument. To simplify notation, we assume without loss of generality (by independence) that the considered optimal independent subplan spans all time periods $1, \dots, T$. Thus, by definition of the subplan: $y_1^S = 1, y_t^S = 0$ for $1 < t \leq T, I_0^S = I_T^S = 0, y_1^R = 1, \sum_{t \in \mathcal{T}} y_t^R = n, I_0^R = I_T^R = 0$.

Case I: suppose $H \leq h$.

Suppose the optimal subplan does not satisfy the zero-inventory property. Let $t_2 \in \mathcal{T}$ be the first time that the zero-inventory property is violated, which must be in the retailer's lot-sizing plan by Lemma 8. Let $1 \leq t_1 < t_2$ be the time of the preceding retailer setup, which must exist. Hence, we have $y_{t_1}^R = y_{t_2}^R = 1, y_t^R = 0$ for all $t_1 < t < t_2, I_{t_1-1}^R = 0$, and $I_{t_2-1}^R > 0$. We emphasize that by assumption $d_{t_1} > 0$ so $x_{t_1}^R > I_{t_2-1}^R$ must hold, which is essential.

Now, shift $I_{t_2-1}^R$ units of inventory to the supplier instead of the retailer as follows. We reduce the retailer order at time t_1 by $I_{t_2-1}^R$ and keep these units in the supplier's inventory, resulting in a change in total costs of $(H - h)(t_2 - t_1)I_{t_2-1}^R \leq 0$. If $H < h$ then this is a strict inequality, which contradicts optimality. If $H = h$ then this shift does not affect the total costs, but removes this violation of the zero-inventory property. By repeating this argument, an optimal subplan can be constructed which satisfies the zero-inventory property.

Case II: suppose $H > h$.

First, we show that all post-initial retailer orders must be the minimum of 1 unit of products. Suppose $1 < \tau \leq T$ exists such that $x_\tau^R > 1$. Shift a unit of supply from the retailer order at time τ to time 1, resulting in change in costs of $(\tau - 1)(h - H) < 0$. This contradicts optimality, so it must hold that $x_t^R \in \mathbb{B}$ for all $1 < t \leq T$.

Second, we prove that all n retailer setups must be in the first n periods. Suppose there exist $t_1, t_2 \in \mathcal{T}$ with $t_2 > t_1 + 1, x_{t_1}^R = x_{t_2}^R = 1$, and $x_t^R = 0$ for all $t_1 < t < t_2$. Shifting the retailer order from time t_2 to $t_1 + 1$ results in a change in costs of $(h - H)(t_2 - t_1 - 1) < 0$, which contradicts optimality. We conclude that all n retailer setups are in the first n periods.

To conclude this case, if $H > h$ there is a unique optimal solution given as follows (see Figure 2). The retailer setups are all in the first n periods: $y_t^R = 1$ if and only if $1 \leq t \leq n$. All post-initial retailer orders have size 1: $x_t^R = 1$ for all $1 < t \leq n$. All remaining demand is supplied from the retailer's inventory, by the order in the initial period: $x_1^R = \sum_{t \in \mathcal{T}} d_t - (n - 1)$. The total costs for the subplan are

$$F + nf + (H - h) \sum_{t=1}^n (t - 1) + h \sum_{t=1}^T (t - 1)d_t.$$

This subplan, the unique optimum, satisfies the zero-inventory property if and only if $n = 1$, or $n = T$ and $d_t = 1$ for all $1 < t \leq T$. Note that if the subplan spans a subset of the time periods, these expressions need to be trivially adjusted accordingly. \square

Proof of Theorem 12. The dynamic-programming (DP) algorithm that provides optimal solutions for all n -plan generation problems is based on Zangwill (1969). From Lemma 3 and Theorem 4 we know that we need to solve several joint lot-sizing problems where the number of retailer setups is fixed to $1, \dots, T$. The DP relies on the fact that any optimal solution is decomposable into independent subplans, see

Lemma 10. By considering all possible decompositions, we can determine an optimal solution. As such, we need to solve the corresponding subproblems, for which we use the insights of Lemma 11.

We now present the DP algorithm. The DP states are (\underline{t}, n) which corresponds to the joint lot-sizing problem with time periods $\{\underline{t}, \dots, T\} \subseteq \mathcal{T}$ and prescribes having exactly $1 \leq n \leq T$ retailer setups. Let $v(\underline{t}, n)$ be the corresponding optimal objective value. Thus, our list of optimal plans follows from the states $(1, n)$ with $n = 1, \dots, T$.

We also need states related to the subproblems. For the subproblems, we have DP states $(\underline{t}, \bar{t}, n)$ which corresponds to the joint lot-sizing problem with time periods $\{\underline{t}, \dots, \bar{t}\} \subseteq \mathcal{T}$ and where the number of retailer setups is fixed to $1 \leq n \leq T$. Let $w(\underline{t}, \bar{t}, n)$ be the optimal objective value *minus* supplier setup cost F of the corresponding optimal independent subplan.

The DP initialisation is given by

$$v(\underline{t}, 1) = F + w(\underline{t}, T, 1),$$

where we consider all feasible states: $1 \leq \underline{t} \leq T$. That is, if only a single retailer setup is allowed, it must be an independent subplan. The DP recursion is:

$$v(\underline{t}, n) = F + \min \left\{ w(\underline{t}, T, n), \min_{\underline{t} \leq \tau < T} \left\{ \min_{\substack{1 \leq m \leq 1 + \tau - \underline{t} \\ n + \tau - T \leq m \leq n - 1}} \{w(\underline{t}, \tau, m) + v(\tau + 1, n - m)\} \right\} \right\},$$

where we consider all feasible states ($n = 1$ is the initialisation): $1 \leq \underline{t} \leq T$ and $1 < n \leq (1 + T - \underline{t})$. Essentially, the DP recursion compares the non-decomposable subplan (\underline{t}, T, n) to all other feasibly decomposable subplans.

The solutions to the subproblems can be determined as follows. First, if only a single retailer setup is allowed ($n = 1$), there is a single feasible solution. We have for all feasible states $1 \leq \underline{t} \leq \bar{t} \leq T$:

$$w(\underline{t}, \bar{t}, 1) = f + h \sum_{t=\underline{t}}^{\bar{t}} (t - \underline{t}) d_t.$$

Next, we use Lemma 11 and need to consider two cases. If $H > h$ then we directly know the unique optimal solution and obtain

$$w(\underline{t}, \bar{t}, n) = nf + (H - h) \sum_{t=1}^n (t - 1) + h \sum_{t=\underline{t}}^{\bar{t}} (t - \underline{t}) d_t,$$

where we consider all feasible states (except $n = 1$): $1 \leq \underline{t} < \bar{t} \leq T$ and $1 < n \leq (1 + \bar{t} - \underline{t})$.

Otherwise, if $H \leq h$ there exists an optimal solution that satisfies the zero-inventory property and we can use a straightforward modification of the standard joint lot-sizing DP:

$$w(\underline{t}, \bar{t}, n) = \min_{\underline{t} \leq \tau \leq \bar{t} + 1 - n} \left\{ \left(f + h \sum_{t=\underline{t}}^{\tau} (t - \underline{t}) d_t + H(\tau + 1 - \underline{t}) \sum_{t=\tau+1}^{\bar{t}} d_t \right) + w(\tau + 1, \bar{t}, n - 1) \right\},$$

where we consider all feasible states (except $n = 1$): $1 \leq \underline{t} < \bar{t} \leq T$ and $1 < n \leq (1 + \bar{t} - \underline{t})$.

The optimal plans are constructed by keeping track of the optimal choices made during the DP. It remains to determine the complexity of this DP. Precomputing the summations of demand takes $\mathcal{O}(T^2)$ time. The calculation of one $w(\underline{t}, \bar{t}, n)$ value then takes at most $\mathcal{O}(T)$ time, leading to $\mathcal{O}(T^4)$ time to determine w . Next, one $v(\underline{t}, n)$ value needs $\mathcal{O}(T^2)$ time, resulting in $\mathcal{O}(T^4)$ time for a complete specification of v . Thus, the total complexity is polynomial, namely $\mathcal{O}(T^4)$ time. \square

Proof of Theorem 13. First, realise that solving all n -plan generation problems results in a list of T plans, so the number of contracts K is equal to T and is polynomial in the input size. The theorem now follows immediately by combining Theorems 7 and 12. \square

Proof of Lemma 14. If ω is a uniform distribution the plan assignment model can be reformulated into a convex model. We refer to the proof of Theorem 7 for the model and its dual variables. We will use the Karush-Kuhn-Tucker (KKT) conditions (see Karush (1939) and Kuhn and Tucker (1951)) to prove the lemma.

For dual feasibility, all dual variables λ_k ($k \in \{0, \dots, K\}$) and $\mu_{k,k+1}$ ($k \in \{0, \dots, K-1\}$) must be non-negative. The KKT stationarity condition for z_1 is

$$1 - \sum_{k=0}^K \lambda_k = 0. \quad (34)$$

The stationarity conditions for $\bar{\theta}_k$, $k \in \{1, \dots, K-1\}$, are

$$0 \in \frac{1}{\bar{\theta} - \theta} \left(2(n_k - n_{k+1})\bar{\theta}_k + ((C_k + c_k^{\text{pub}}) - (C_{k+1} + c_{k+1}^{\text{pub}}) - (n_k - n_{k+1})\bar{\theta}) \right) \\ + (n_k - \partial\phi^*(\bar{\theta}_k))\lambda_k + (n_k - n_{k+1}) \sum_{i=k+1}^K \lambda_i - \mu_{k-1,k} + \mu_{k,k+1}. \quad (35)$$

Here, we need the subdifferential $\partial\phi^*(\bar{\theta}_k)$, since ϕ^* is not differentiable in all points. In our case, $\partial\phi^*(\bar{\theta}_k)$ is the closed interval

$$\partial\phi^*(\bar{\theta}_k) = \left[\lim_{\theta_k \downarrow \bar{\theta}_k} \frac{d}{d\theta_k} \phi^*(\theta_k), \lim_{\theta_k \uparrow \bar{\theta}_k} \frac{d}{d\theta_k} \phi^*(\theta_k) \right].$$

Hence, in the optimal solution there exist subgradients, for simplicity denoted by $\frac{d}{d\theta_k} \phi^*(\bar{\theta}_k) \in \partial\phi^*(\bar{\theta}_k)$ for $k \in \mathcal{K}$, such that (35) is satisfied. In other words, for these subgradients it must hold that

$$\bar{\theta}_k = \frac{(C_{k+1} + c_{k+1}^{\text{pub}}) - (C_k + c_k^{\text{pub}})}{2(n_k - n_{k+1})} + \frac{1}{2} \left(\bar{\theta} - (\bar{\theta} - \theta) \sum_{i=k+1}^K \lambda_i \right) \\ - \frac{(\bar{\theta} - \theta)}{2(n_k - n_{k+1})} \left(\left(n_k - \frac{d}{d\theta_k} \phi^*(\bar{\theta}_k) \right) \lambda_k - \mu_{k-1,k} + \mu_{k,k+1} \right). \quad (36)$$

Now consider $k \in \{2, \dots, K-1\}$ such that ϕ^* has no slopes n^* with $n_{k-1} > n^* > n_{k+1}$. If plan k is assigned, we have $\bar{\theta}_k > \bar{\theta}_{k-1}$ and $\mu_{k-1,k} = 0$. Substituting (36) in $\bar{\theta}_k > \bar{\theta}_{k-1}$ results in the condition

$$\frac{(C_k + c_k^{\text{pub}}) - (C_{k-1} + c_{k-1}^{\text{pub}})}{2(n_{k-1} - n_k)} + \frac{(C_k + c_k^{\text{pub}}) - (C_{k+1} + c_{k+1}^{\text{pub}})}{2(n_k - n_{k+1})} \\ + \frac{(\bar{\theta} - \theta)}{2(n_k - n_{k+1})} \left(\left(n_{k+1} - \frac{d}{d\theta_k} \phi^*(\bar{\theta}_k) \right) \lambda_k + \mu_{k,k+1} \right) \\ + \frac{(\bar{\theta} - \theta)}{2(n_{k-1} - n_k)} \left(- \left(n_{k-1} - \frac{d}{d\theta_{k-1}} \phi^*(\bar{\theta}_{k-1}) \right) \lambda_{k-1} + \mu_{k-2,k-1} \right) < 0. \quad (37)$$

Note that the summations over λ_i almost cancel out: only λ_k remains which explains why the term $n_{k+1}\lambda_k$ is present. We claim that the terms in (37) containing λ_{k-1} and λ_k are non-negative.

First, if $\lambda_{k-1} > 0$ then the IR constraint binds at $\bar{\theta}_{k-1}$. This is only possible if $n_{k-1} \geq \frac{d}{d\bar{\theta}_{k-1}} \phi^*(\bar{\theta}_{k-1})$. If $n_{k-1} = \frac{d}{d\bar{\theta}_{k-1}} \phi^*(\bar{\theta}_{k-1})$, then the term with λ_{k-1} is zero. Otherwise, if $n_{k-1} > \frac{d}{d\bar{\theta}_{k-1}} \phi^*(\bar{\theta}_{k-1})$, we get a contradiction with our assumptions. Namely, in this case it must hold that $\lim_{\theta_{k-1} \downarrow \bar{\theta}_{k-1}} \frac{d}{d\theta_{k-1}} \phi^*(\theta_{k-1}) \leq n_{k+1} < n_k$ as there are no larger eligible slopes of ϕ^* by assumption. Assigning plan k will violate the IR constraints, which leads to a contradiction. We conclude that $(n_{k-1} - \frac{d}{d\bar{\theta}_{k-1}} \phi^*(\bar{\theta}_{k-1}))\lambda_{k-1} = 0$.

Second, if $\lambda_k > 0$ then $n_k > \frac{d}{d\bar{\theta}_k} \phi^*(\bar{\theta}_k)$, where the inequality is strict since n_k is not a slope of ϕ^* . In particular, it must hold that $\frac{d}{d\bar{\theta}_k} \phi^*(\bar{\theta}_k) \leq n_{k+1}$. This proves our claim.

Thus, all terms with dual variables in (37) are non-negative. Therefore, the other terms must be strictly negative in total, i.e., (27) must hold. \square

Proof of Lemma 15. First of all, realise that $\max_{\tau \in \mathcal{T}} \{(\tau - 1) \sum_{t=\tau}^T d_t\}$ is the maximum inventory that can be rerouted to a new supplier setup. So the assumption on the supplier's setup cost implies that it is never optimal to have more than a single supplier setup in a lot-sizing plan. Consequently, any optimal solution for the n -plan generation problem always has total costs

$$F + \left(\sum_{t \in \mathcal{T}} (t-1)d_t - n \right) H \quad (38)$$

for the supplier. Only the number of retailer setups can be minimised. Let $m_k \in \mathbb{N}_{\geq 1}$ be the number of retailer setups of the optimal n_k -plan.

Suppose we have determined a minimal list of plans sufficient for optimality for the contracting problem. Hence, each of these plans is assigned to retailer types. Add to this list the following T plans if not yet present: for $m \in \{1, \dots, T\}$ add any m -extreme plan, i.e., a plan that has minimal retailer inventory with m retailer setups. Denote the indices of the combined list by \mathcal{K} and the resulting indices of the extreme plans by $\mathcal{L} \subseteq \mathcal{K}$. It can be verified from the properties of ϕ^* that the slopes of ϕ^* are contained in $\{n_l : l \in \mathcal{L}\}$.

By definition of the extreme plans, we have $n_1 = \sum_{t \in \mathcal{T}} (t-1)d_t$, $n_K = 0$, and $1, K \in \mathcal{L}$. By our assumptions, and consequently by (38), all m -extreme plans have the same supplier's costs and retailer's public costs for fixed m . Therefore, it does not matter which m -extreme plan is added. Also, any two extreme plans in \mathcal{L} have different amounts of retailer inventory: the plan with the most retailer setups has the lowest retailer inventory. This trivially follows from adding retailer setups to the plan with less retailer setups, which must decrease the retailer inventory.

Consider a non-extreme plan $k \in \mathcal{K} \setminus \mathcal{L}$, so $1 < k < K$ holds. Realise that we can apply Lemma 14 to this plan. If this plan would be assigned in the optimal menu, then Lemma 14 states that the following condition must be met:

$$\begin{aligned} & \frac{(C_k + c_k^{\text{pub}}) - (C_{k-1} + c_{k-1}^{\text{pub}})}{n_{k-1} - n_k} + \frac{(C_k + c_k^{\text{pub}}) - (C_{k+1} + c_{k+1}^{\text{pub}})}{n_k - n_{k+1}} < 0 \\ \iff & \frac{(n_{k-1} - n_k)H + (m_k - m_{k-1})f}{n_{k-1} - n_k} + \frac{(n_{k+1} - n_k)H + (m_k - m_{k+1})f}{n_k - n_{k+1}} < 0 \\ \iff & \frac{m_k - m_{k-1}}{n_{k-1} - n_k} + \frac{m_k - m_{k+1}}{n_k - n_{k+1}} < 0. \end{aligned} \quad (39)$$

Define $\underline{k}, \bar{k} \in \mathcal{L}$ such that $n_{\underline{k}} > n_k > n_{\bar{k}}$ and $m_{\bar{k}} = m_k + 1$, so $\underline{k} < k < \bar{k}$. Note that these must exist. By definition, (39) must hold for each $\underline{k} < k < \bar{k}$. Adding these inequalities results in

$$\frac{m_{k+1} - m_{\underline{k}}}{n_{\underline{k}} - n_{k+1}} + \frac{m_{\bar{k}-1} - m_{\bar{k}}}{n_{\bar{k}-1} - n_{\bar{k}}} < 0,$$

as all interior terms cancel out. We claim that this is a contradiction. By definition, n_k is the minimum amount of retailer inventory when using m_k retailer setups, and likewise for $n_{\bar{k}}$ and $m_{\bar{k}}$. Since $n_k < n_{\bar{k}}$ for all $k > \bar{k}$ it follows that $m_k > m_{\bar{k}}$ for $k > \bar{k}$. Also, by definition we have $m_{\bar{k}} = m_{\bar{k}} + 1$ and, as a consequence, $m_k \geq m_{\bar{k}}$ for $k > \bar{k}$. By combining these insights, we obtain the following contradiction:

$$0 > \frac{m_{k+1} - m_{\underline{k}}}{n_{\underline{k}} - n_{k+1}} + \frac{m_{\bar{k}-1} - m_{\bar{k}}}{n_{\bar{k}-1} - n_{\bar{k}}} \geq \frac{1}{n_k - n_{k+1}} > 0.$$

We conclude that we only need to include plans which have minimal retailer inventory for each number of retailer setups. \square

Proof of Lemma 16. Recall that under the assumptions all m -extreme plan have the same supplier's costs and retailer's public costs for fixed m . Therefore, we only need to determine a plan that minimises the retailer's inventory and uses exactly m retailer setups. The dynamic-programming (DP) algorithm is as follows. The DP states are (\underline{t}, m) , which corresponds to the joint lot-sizing problem with time periods $\{\underline{t}, \dots, T\} \subseteq \mathcal{T}$ and prescribes having exactly $1 \leq m \leq T$ retailer setups. The corresponding minimal retailer inventory is denoted by $v(\underline{t}, m)$. Hence, our list of extreme plans follows from the states $(1, m)$ with $m = 1, \dots, T$.

The DP initialisation for $1 \leq \underline{t} \leq T$ is

$$v(\underline{t}, 1) = \sum_{t=\underline{t}}^T (t - \underline{t})d_t.$$

The DP recursion is given by

$$v(\underline{t}, m) = \min_{\underline{t} \leq \tau \leq T+1-m} \left\{ \sum_{t=\underline{t}}^{\tau} (t - \underline{t})d_t + v(\tau + 1, m - 1) \right\},$$

where we consider all feasible states (except $m = 1$): $1 \leq \underline{t} < T$ and $1 < m \leq (1 + T - \underline{t})$.

The plans are constructed by keeping track of the optimal choices made during the DP algorithm. Precomputing the summations of demand takes $\mathcal{O}(T^2)$ time. The calculation of each DP state then requires $\mathcal{O}(T)$ time, leading to $\mathcal{O}(T^3)$ time in total for the algorithm. \square

Proof of Corollary 17. By Lemma 15 it is sufficient for optimality to use a list consisting of an m -extreme plan for each $m \in \{1, \dots, T\}$. This list consists of T plans and it can be constructed in polynomial time by Lemma 16. Thus, the number of contracts K is equal to T and polynomial in the input size. The result now follows from Theorem 7. \square

B Example for private setup cost

In this appendix we give an example for the private setup cost case where the optimal menu assigns a contract that violates the zero-inventory property. We consider the smallest example by setting $T = 2$, $d_1 = 1$, and $d_2 = 2$, see Figure 4 for the optimal solution for each n -plan generation problem. Note that the plan in Figure 4a does not satisfy the zero-inventory property.




Figure 4: The optimal solutions for the n -plan generation problems in the case of private setup cost with $T = 2$, $d_1 = 1$, and $d_2 = 2$.

The cost parameters are as follows: $\theta = 1$, $\bar{\theta} = 5$, $F = 4$, $H = 2$, and $h = 1$. It is trivial to verify that the default option ϕ^* is given by

$$\phi^*(\theta) = \begin{cases} 2\theta & \text{if } \theta \in [1, 2] \\ \theta + 2 & \text{if } \theta \in (2, 5] \end{cases}.$$

Since $H > h$ and $F > H + h$ the plan in Figure 4a is the only optimum for $n = 2$. Furthermore, we let ω be a uniform distribution. For the first contract we have $n_1 = 2$ retailer setups, $C_1 = F + H = 6$, and $c_1^{\text{pub}} = h = 1$. For the second contract these parameters are $n_2 = 1$, $C_2 = F = 4$, and $c_2^{\text{pub}} = 2h = 2$.

If the optimal $\bar{\theta}_1$ lies in $[1, 2)$, then the first segment of ϕ^* binds. By (19) and (21) the side payments must be

$$\begin{aligned} z_2 &= c_2^{\text{pub}} + n_2 \theta_2 - \phi^*(\theta_2) = 2 + \bar{\theta}_1 - 2\bar{\theta}_1 = 2 - \bar{\theta}_1, \\ z_1 &= z_2 + c_1^{\text{pub}} - c_2^{\text{pub}} + (n_1 - n_2)\theta_2 = (2 - \bar{\theta}_1) + 1 - 2 + \bar{\theta}_1 = 1. \end{aligned}$$

Consequently, the objective function is

$$\frac{1}{\bar{\theta} - \theta} \left((\bar{\theta}_1 - \theta)(C_1 + z_1) + (\bar{\theta} - \bar{\theta}_1)(C_2 + z_2) \right) = \frac{1}{4} (\bar{\theta}_1^2 - 4\bar{\theta}_1 + 23).$$

Since the derivative is $\frac{1}{4}(2\bar{\theta}_1 - 4) < 0$ on $[1, 2)$, it is optimal set $\bar{\theta}_1$ as large as possible in $[1, 2)$.

If the optimal $\bar{\theta}_1$ lies in $(2, 5]$, then the second segment of ϕ^* binds. Similar to the previous case, this leads to side payments $z_2 = 2 + \bar{\theta}_1 - (\bar{\theta}_1 + 2) = 0$ and $z_1 = \bar{\theta}_1 - 1$. The resulting objective function is $\frac{1}{4}(\bar{\theta}_1^2 + 15)$, which has strictly positive derivative on $(2, 5]$. Hence, it is optimal to set $\bar{\theta}_1$ as small as possible in $(2, 5]$.

Finally, it can be similarly verified that if $\bar{\theta}_1 = 2$, then $z_1 = 1$, $z_2 = 0$, and the objective value is $\frac{19}{4}$. This corresponds to the limits of the above two cases. We conclude that this is the unique optimal solution, which assigns the plan in Figure 4a to types $[1, 2]$ and the plan in Figure 4d to types $[2, 5]$. Thus, the optimal menu assigns a plan that violates the zero-inventory property.

C Relaxing the demand assumption

If demand can be zero certain results are no longer valid or need to be adjusted. This appendix describes the required changes. In Assumption 1 we state the properties of the considered demand.

Assumption 1. The (integer) demand is strictly positive in the first and last period, and non-negative otherwise, i.e., $d_1, d_T \in \mathbb{N}_{>0}$ and $d_t \in \mathbb{N}$ for $1 < t < T$.

Since we have time-independent costs, the assumption $d_1, d_T > 0$ is without loss of generality. If either is zero, that period can be removed from the problem. Under Assumption 1 only the results for the plan generation for private setup cost are significantly affected, i.e., Lemmas 10 and 11, and Theorem 12. To address the issues we condition on whether $H \leq h$ (Section C.1) or $H > h$ (Section C.2). Before we do so, we mention the minor adjustments needed for the other results.

Throughout our results and proofs we use that the number of retailer setups lies between 1 and T . If some demand is zero, there is a maximum feasible number of retailer setups $1 \leq K \leq T$, which should be used instead. This has no significant impact on the results. Furthermore, dummy setups (with zero order quantity) need to be prevented in any dynamic-programming algorithm. This slight adjustment does not affect the complexity results.

For the remainder of this appendix we consider the plan generation problem for private setup cost f under Assumption 1. For the proofs we often refer to or reuse parts of the proofs for Lemmas 10 and 11, and Theorem 12.

C.1 Case $H \leq h$

This case is more regular and essentially all results still hold. In Lemma 18 we will show that any plan assigned in the optimal menu must satisfy the zero-inventory property (and thus be decomposable). This lemma replaces Lemmas 10 and 11 for this case, and leads to the same overall conclusion. Consequently, Theorem 12 is unaffected and this case is solvable in $\mathcal{O}(T^4)$ time by dynamic programming. We continue with the lemma and its proof.

Lemma 18. *Under Assumption 1, private setup cost f , and $H \leq h$, it suffices for optimality for the contracting problem to restrict the plan generation to plans satisfying the zero-inventory property.*

Proof. The idea is as follows. First, we show that without loss of optimality an optimal n -plan violates the zero-inventory property only if n is strictly larger than the maximum slope of ϕ^* . This implies that this plan has unnecessary many retailer setups. Next, we show that such plans are never assigned in an optimal menu. The details are given below.

Part 1: violating the zero-inventory property implies a larger slope than ϕ^* .

We first introduce some notation. For $t \in \mathcal{T}$ with $I_{t-1}^R > 0$ we can backtrack the flow of a unit of products from the inventory I_{t-1}^R to its origin at the supplier. At some time period $\tau \in \{1, \dots, t-1\}$ this unit of products is transferred from the supplier to the retailer by a retailer order, which we call the transfer time. Since there can be multiple options when backtracking, we define the indices $\mathcal{L}(t)$ such that $\{\tau_l : l \in \mathcal{L}(t)\} \subseteq \{1, \dots, t-1\}$ are *all possible* transfer times for a unit of products of the inventory I_{t-1}^R . That is, any path in the network flow graph from the supplier to the retailer's inventory I_{t-1}^R transfers to the retailer's layer at some time τ_l , $l \in \mathcal{L}(t)$. By definition, we have $x_{\tau_l}^R \geq 1$ for all $l \in \mathcal{L}(t)$. We denote the first and last transfer times by $\tau_1(t)$ and $\tau_L(t)$, respectively.

Suppose an optimal solution to the n -plan generation problem exists that violates the zero-inventory property, say at time $t_1 \in \mathcal{T}$ (and potentially elsewhere). By Lemma 8 this implies that $1 < t_1 \leq T$, $x_{t_1}^R > 0$, and $I_{t_1-1}^R > 0$.

If $l \in \mathcal{L}(t_1)$ with $x_{\tau_l}^R > 1$ exists, then we can reduce the inventory $I_{t_1-1}^R$ by shifting a unit of products to the supplier's inventory at time τ_l instead of the retailer's inventory (see Figure 5a). The costs change by $(t_1 - \tau_l)(H - h) \leq 0$. Since $\tau_l < t_1$ this contradicts optimality if $H < h$, otherwise (if $H = h$) we can exclude this solution without loss of optimality. Hence, for any time $t \in \mathcal{T}$ where the plan violates the zero-inventory property ($x_t^R > 0$ and $I_{t-1}^R > 0$) it must hold that $x_{\tau_l}^R = 1$ for all $l \in \mathcal{L}(t)$.

Suppose there exists a time $t \in \mathcal{T}$ with $y_t^R = 0$, $I_{t-1}^R > 0$, and $x_{\tau_l(t)}^R > 1$ for some $l \in \mathcal{L}(t)$. See also Figure 5b. Decrease I_{t-1}^R by shifting one unit of the flow through the latest time $\tau_L(t)$ to the supplier's inventory instead. This removes the retailer setup at time $\tau_L(t)$, but still leads to a feasible plan with a change in costs of $(H - h)(t - \tau_L(t)) \leq 0$. Reinsert the removed retailer setup at time t and shift a single unit of products from time $\tau_l(t) < t$ (which has $x_{\tau_l(t)}^R > 1$ by assumption) to this new setup. The result is a feasible n -plan with a change in costs of $(H - h)(t - \tau_l(t)) + (H - h)(t_1 - \tau_L(t)) \leq 0$,

compared to the original n -plan. This contradicts the optimality of the original plan if $H < h$ or we can exclude it without loss of optimality ($H = h$).

Thus, by combining the above, we have for all $t \in \mathcal{T}$ with $I_{t-1}^R > 0$ that $x_{\tau_l}^R = 1$ for all $l \in \mathcal{L}(t)$. Suppose there exists $t \in \mathcal{T}$ with $d_t > 0$ and $y_t^R = 0$, which implies that $I_{t-1}^R > 0$. Hence, $x_{\tau_l}^R = 1$ for all $l \in \mathcal{L}(t)$ must hold. Realise that this is only possible if $d_{\tau_1(t)} = 0$. We can reposition this retailer setup at time $\tau_1(t)$ to time t resulting in a change in costs of $(H - h)(t - \tau_1(t)) \leq 0$. Again, this results in a contradiction to optimality ($H < h$) or can be excluded without loss of optimality ($H = h$).

We conclude that $y_t^R = 1$ for all $t \in \mathcal{T}$ with $d_t > 0$. Furthermore, it must hold that $d_{\tau_1(t_1)} = 0$ as $x_{\tau_1(t_1)}^R = 1$. Therefore, the plan must have strictly more retailer setups than the number of time periods with strictly positive demand. In other words, the plan has enough retailer setups to directly satisfy all demand without any retailer inventory. However, there are redundant retailer setups (such as at $\tau_1(t_1)$) that only increase the number of retailer setups and lead to higher costs for both the supplier and the retailer. Consequently, the number of retailer setups n of this optimal plan exceeds the maximum slope of ϕ^* , which is essential.




Figure 5: Shifting flow in procurement plans. Reduced flows are red, increased flows are blue, purple setups are removed, and green setups are new.

Part 2: plans violating the zero-inventory property are not assigned in an optimal menu.

Suppose the considered plan, which violates the zero-inventory property, is assigned to retailer types in an optimal solution to the contracting problem. Let it be the k -th contract in the optimal menu and let n_k denote its number of retailer setups. By Lemma 6 we assume without loss of generality that the $(k + 1)$ -th contract has a plan with one less retailer setup, $n_{k+1} = n_k - 1$, but is potentially not assigned to retailer types. Consider the k -th plan and remove a unit of inventory from $I_{t_1-1}^R$, remove the retailer setup at time $\tau_L(t_1)$, and shift that supply to the supplier's inventory. We obtain a feasible $(n_k - 1)$ -plan with a change in costs of $(H - h)(t_1 - \tau_L(t_1)) \leq 0$. Hence, it must hold that

$$C_{k+1} + c_{k+1}^{\text{pub}} \leq C_k + c_k^{\text{pub}} + (H - h)(t_1 - \tau_L(t_1)) \leq C_k + c_k^{\text{pub}}.$$

By also using (21), we have

$$\begin{aligned} C_{k+1} + z_{k+1} &= C_{k+1} + c_{k+1}^{\text{pub}} - c_{k+1}^{\text{pub}} + z_{k+1} \leq C_k + c_k^{\text{pub}} - c_{k+1}^{\text{pub}} + z_{k+1} \\ &= C_k + z_k - (n_k - n_{k+1})\underline{\theta}_{k+1} < C_k + z_k. \end{aligned} \quad (40)$$

Thus, keeping the side payments z_{k+1}, \dots, z_K constant and replacing the k -th contract by the $(k + 1)$ -th leads to a strictly better objective value when considered in isolation. In terms of the partition, we only change $\underline{\theta}_{k+1}$ to $\underline{\theta}_k$, so the k -th contract is no longer assigned. However, the side payments of contracts $1, \dots, k$ also need to be adjusted for feasibility. This can be done by decreasing z_1, \dots, z_k by $(n_k - n_{k+1})(\underline{\theta}_{k+1} - \underline{\theta}_k) > 0$ according to (21). Namely, consider (21) for $k + 1$. Since $\underline{\theta}_{k+1}$ is changed to $\underline{\theta}_k$, z_k needs to decrease by $(n_k - n_{k+1})(\underline{\theta}_{k+1} - \underline{\theta}_k)$. Now consider (21) for k . As $\underline{\theta}_k$ remains unchanged, we need to decrease z_{k-1} by $(n_k - n_{k+1})(\underline{\theta}_{k+1} - \underline{\theta}_k)$. The previous argument also holds for all contracts $1, \dots, k - 1$. Graphically, we are shifting all lines $1, \dots, k$ vertically upwards such that the k -th line is no longer essential for the lower envelope. By (40) and the decrease in side payments, these modifications improve the objective value. All these modifications to the menu are feasible with respect to the IR constraints, since the plans of *all* contracts $1, \dots, k + 1$ have a slope of at least the maximum slope of ϕ^* . The latter property is essential for this argument.

We conclude that we have constructed a strictly better feasible menu where the considered k -th contract is not assigned to retailer types. This contradiction implies that it suffices for optimality for the contracting problem to restrict the plan generation to plans that satisfy the zero-inventory property. \square

C.2 Case $H > h$

In this case, there are instances for which any optimal n -plan is non-decomposable into independent subplans for certain values of n . For example, by appropriately choosing $H > h$, M , and F , the plan in Figure 6a is the unique optimum. In particular, for any $H > h$ this example can be extended such that the optimum is non-decomposable. Thus, Lemma 10 does not hold in general. The details are as follows.




Figure 6: Two possible solutions for the 5-plan generation problem in the case of private setup cost with $T = 6$, $d = (M, 0, M, M, 1, 1)$.

For given H and h with $H > h$, let $T \geq 6$ be such that

$$H > \frac{T-3}{T-5}h,$$

and set the demand to $d = (M, 0, M, M, 1, \dots, 1)$ for some $M \in \mathbb{N}_{>0}$ large enough. The parameters M and F are chosen such that it is optimal to have a supplier setup only at the time periods with demand M , i.e., time periods 1, 3, and 4.

Now consider the $(T-1)$ -plan generation problem. Any plan that does not have a supplier and retailer setup at times 1, 3, and 4 is suboptimal due to the choice of M and the resulting large holding costs. Suppose an optimal $(T-1)$ -plan does not have a retailer setup at time 2, then it must have retailer setups in all other periods (similar to Figure 6b). Perform the following shift to obtain a feasible plan similar to Figure 6a. Shift the supply provided by x_2^R to x_4^R , reposition the retailer setup at time T to time 2, and set $x_2^R = 1$ by shifting a unit of supply from the orders at time 3 to time 1. This leads to a change in costs of $(h-H)(T-4) + (H+h) = h(T-3) - H(T-5) < 0$ by choice of T . From this contradiction to optimality, we conclude that any optimal $(T-1)$ -plan is non-decomposable.

There are also instances where this non-decomposable plan is necessary for optimality for the contracting problem, i.e., any solution can be improved by adding it to the menu. For example, let ω be a uniform distribution and take $T = 6$, $d = (20, 0, 20, 20, 1, 1)$, $F = 20$, $H = 7$, $h = 2$, and $f \in [1, 20] = \Theta$. These parameters satisfy the conditions given above, so contract 2 with $n_2 = T-1 = 5$ is non-decomposable. Furthermore, an optimal solution is to assign contracts 2, 3, and 4 (with $n_2 = 5$, $n_3 = 4$, and $n_4 = 3$) to retailer types $[1, 2]$, $[2, 4]$, and $[4, 20]$, respectively. The side payments are $z = (3, 6, 2, 0, 24, 172)$ and the optimal objective value is 62. Removing the non-decomposable plan from the menu, leads to an optimal objective value of approximately 62.1, showing the necessity to include this plan.

With these insights, we provide a dynamic-programming algorithm to solve all n -plan generation problems. The running time is polynomial, namely $\mathcal{O}(T^5)$. The presented algorithm suffices for our goal to show that this case is also efficiently solvable. The result is stated in Lemma 19.

Lemma 19. *Under Assumption 1, private setup cost f , and $H > h$, solving all n -plan generation problems can be done in $\mathcal{O}(T^5)$ time by dynamic programming.*

Proof. For $H > h$ and some demand being zero, the optimal n -plan solution might be non-decomposable into independent subproblems (recall the example in Figure 6a). Therefore, a new approach is needed. The idea is as follows.

First, we show that an optimal plan consists of substructures as illustrated in Figure 7, which is similar to Figure 2. In particular, the optimal solution is fixed when we are given the supplier setups and how many retailer setups occur in between the supplier setups. Here, it is essential that $H > h$, so we know that as much inventory as possible is placed at the retailer (without invalidating retailer setups).

Next, we describe a dynamic-programming (DP) algorithm similar to that in the proof of Theorem 12. Since we might have a non-decomposable solution, we need to add the available inventory to the DP states. In an optimal solution these inventory states are non-zero only if removing that inventory would lead to dummy retailer setups. Hence, the inventory states are bounded by T , which is essential to obtain a polynomial-time algorithm. Below we give all the details.




Figure 7: The unique optimal substructure in the case of private setup cost and $H > h$. There are sequential supplier setups at time periods t_i and t_{i+1} , and n_i retailer setups in time periods $t_i, \dots, t_{i+1} - 1$.

Part 1: properties of the optimal n -plan.

Consider an optimal solution of the n -plan generation problem. Assume that $t_1 < t_2 \in \mathcal{T}$ exist such that $y_{t_1}^S = y_{t_2}^S = 1$ and $y_t^S = 0$ for $t_1 < t < t_2$. By Lemmas 8 and 9 we know that $I_{t_1-1}^S = I_{t_2-1}^S = 0$ and $y_{t_1}^R = y_{t_2}^R = 1$. Let $n_1 \in \{1, \dots, t_2 - t_1\}$ be the number of retailer setups in $t_1, \dots, t_2 - 1$. Note that if such t_1 and t_2 do not exist, then the optimum is an independent subplan and we can apply Lemma 11 to obtain similar properties as derived below.

First, we show that all retailer orders in periods $t_1 + 1, \dots, t_2 - 1$ must be the minimum of 1 unit of products. Suppose $t_1 < \tau < t_2$ exists such that $x_\tau^R > 1$. These products are satisfied from the supplier's inventory. Shift a unit of supply from the retailer order at time τ to time t_1 , resulting in change in costs of $(\tau - t_1)(h - H) < 0$. This contradicts optimality, so it must hold that $x_t^R \in \mathbb{B}$ for all $t_1 < t < t_2$.

Second, we prove that all n_1 retailer setups must be in the initial periods. Suppose there exist $t_1 \leq \tau_1 < \tau_2 < t_2$ with $\tau_2 > \tau_1 + 1$, $x_{\tau_1}^R = x_{\tau_2}^R = 1$, and $x_t^R = 0$ for all $\tau_1 < t < \tau_2$. Again, these products are satisfied from the supplier's inventory. Shifting the retailer order from time τ_2 to $\tau_1 + 1$ results in a change in costs of $(h - H)(\tau_2 - \tau_1 - 1) < 0$, which contradicts optimality. We conclude that all n_1 retailer setups are in the first n_1 periods of the considered time periods.

For the considered optimal solution we have for some $N \in \{1, \dots, T\}$ the supplier setups $t_1, \dots, t_N \in \mathcal{T}$ and the number of retailer setups $n_1, \dots, n_N \in \{1, \dots, T\}$ such that n_i retailer setups take place in time periods $t_i, \dots, t_{i+1} - 1$ for $i \in \{1, \dots, N - 1\}$ and n_N in t_N, \dots, T . From the above, we must have $x_{t_i}^S > 0$, $x_{t_i}^R > 0$, $x_{t_i+1}^R = \dots = x_{t_i+n_i-1}^R = 1$ for $i \in \{1, \dots, N\}$ and zero otherwise. See also Figure 7.

Now realise that the retailer order $x_{t_i}^R$ needs to be minimal in the optimal solution. Otherwise there is unnecessary retailer inventory at time $t_{i+1} - 1$ which can be transferred to the supplier and retailer setup at time t_{i+1} . This strictly reduces the total costs, contradicting optimality. The minimal order $x_{t_i}^R$ is at least 1 (to keep it a valid setup) and is such that all demand in $t_i, \dots, t_{i+1} - 1$ is satisfied, taking into account the supply from any available inventory $I_{t_i-1}^R$ and any additional retailer setups $x_{t_i+1}^R = \dots = x_{t_i+n_i-1}^R = 1$.

In particular, if $I_{t_i-1}^R > \sum_{j=1}^{i-1} n_j$ for some $i \in \{1, \dots, N\}$, then the retailer inventory $I_{t_i-1}^R$ can be reduced, which contradicts optimality. Namely, consider the artificial situation that $d_t = 0$ for all $1 \leq t < t_i$. The $\sum_{j=1}^{i-1} n_j$ retailer setups need to provide at least 1 supply each, which is transferred by retailer inventory $I_{t_i-1}^R$ to time t_i . If more inventory is supplied, then a unit of products can be shifted from a previous order to the order at time t_i , without invalidating a retailer setup. This strictly reduces the costs, proving the claim. Note that a better bound exists, but this suffices to obtain polynomial running time.

Part 2: the dynamic-programming algorithm.

We can now formulate a DP algorithm with a polynomial running time that solves all n -plan generation problems. The DP states are $(\underline{t}, n, m_{\text{in}})$ which corresponds to the joint lot-sizing subproblem with time periods $\{\underline{t}, \dots, T\} \subseteq \mathcal{T}$ and prescribes having exactly $1 \leq n \leq T$ retailer setups. Furthermore, the supplier inventory satisfies $I_{\underline{t}-1}^S = I_T^S = 0$ and for the retailer inventory we have $I_{\underline{t}-1}^R = m_{\text{in}} \in \{0, \dots, \underline{t} - 1\}$ and $I_T^R = 0$. Note that the bound on m_{in} follows from the above arguments on minimal remaining inventory. Let $v(\underline{t}, n, m_{\text{in}})$ be the corresponding optimal objective value. Our list of optimal plans follows from the states $(1, n, 0)$ with $n = 1, \dots, T$ (or up to the maximum feasible number of retailer setups).

We also need states $(\underline{t}, \bar{t}, n, m_{\text{in}})$ for $\underline{t} \leq \bar{t} \in \mathcal{T}$, $n \in \{1, \dots, 1 + \bar{t} - \underline{t}\}$, and $m_{\text{in}} \in \{0, \dots, \underline{t} - 1\}$ as follows. They correspond to subproblems spanning time periods $\underline{t}, \dots, \bar{t}$ with a supplier setup only at time \underline{t} : $y_{\underline{t}}^S = 1$ and $y_t^S = 0$ for $\underline{t} < t \leq \bar{t}$. Furthermore, there are exactly n retailer setups: $\sum_{t=\underline{t}}^{\bar{t}} y_t^R = n$. Finally, the supplier inventory satisfies $I_{\underline{t}-1}^S = I_{\bar{t}}^S = 0$ and for the retailer inventory we have $I_{\underline{t}-1}^R = m_{\text{in}}$.

By the above analysis there exists a unique optimal subplan for each feasible state $(\underline{t}, \bar{t}, n, m_{\text{in}})$. Let $w(\underline{t}, \bar{t}, n, m_{\text{in}})$ denote the corresponding costs. Furthermore, we have the corresponding (minimal) remaining inventory $I_{\bar{t}}^R$ denoted by $m_{\text{out}}(\underline{t}, \bar{t}, n, m_{\text{in}})$. Any infeasible states are assigned the value infinity,

i.e., they are omitted when determining the optimal plans.

The DP initialisation for feasible states is given by

$$v(\underline{t}, 1, m_{\text{in}}) = w(\underline{t}, T, 1, m_{\text{in}}).$$

The DP recursion for feasible states is:

$$v(\underline{t}, n, m_{\text{in}}) = \min \left\{ w(\underline{t}, T, n, m_{\text{in}}), \right. \\ \left. \min_{\underline{t} \leq \tau < T} \left\{ \min_{1 \leq \kappa \leq n-1} \{ w(\underline{t}, \tau, \kappa, m_{\text{in}}) + v(\tau + 1, n - \kappa, m_{\text{out}}(\underline{t}, \tau, \kappa, m_{\text{in}})) \} \right\} \right\}.$$

Certain options in the shown ranges might lead to infeasible states and should be omitted. There are $\mathcal{O}(T^3)$ many DP states and each takes $\mathcal{O}(T^2)$ time to compute. Thus, to determine v we need $\mathcal{O}(T^5)$ time.

It remains to solve the subproblems related to w and m_{out} . For a feasible state $(\underline{t}, \bar{t}, n, 0)$ the supplier order must be $x_{\underline{t}}^S = \max\{n, \sum_{t=\underline{t}}^{\bar{t}} d_t\}$ in order to supply all demand in $\underline{t}, \dots, \bar{t}$ and to have no dummy retailer setups. We can construct the corresponding subplan in $\mathcal{O}(T)$ time, from which we obtain I_t^R , $t \in \{\underline{t}, \dots, \bar{t}\}$, and in particular the remaining inventory $m_{\text{out}}(\underline{t}, \bar{t}, n, 0)$. The corresponding costs are

$$w(\underline{t}, \bar{t}, n, 0) = F + nf + H \sum_{t=1}^n (t-1) + h \sum_{t=\underline{t}}^{\bar{t}} I_t^R.$$

Thus, calculating $w(\underline{t}, \bar{t}, n, 0)$ and $m_{\text{out}}(\underline{t}, \bar{t}, n, 0)$ for all \underline{t}, \bar{t} , and n takes $\mathcal{O}(T^4)$ time. The other feasible states follow from

$$w(\underline{t}, \bar{t}, n, m+1) = \begin{cases} w(\underline{t}, \bar{t}, n, m) & \text{if } x_{\underline{t}}^R > 1 \text{ in state } (\underline{t}, \bar{t}, n, m), \\ w(\underline{t}, \bar{t}, n, m) + h(1 + \bar{t} - \underline{t}) & \text{otherwise} \end{cases}, \\ m_{\text{out}}(\underline{t}, \bar{t}, n, m+1) = \begin{cases} m_{\text{out}}(\underline{t}, \bar{t}, n, m) & \text{if } x_{\underline{t}}^R > 1 \text{ in state } (\underline{t}, \bar{t}, n, m), \\ m_{\text{out}}(\underline{t}, \bar{t}, n, m) + 1 & \text{otherwise} \end{cases}.$$

That is, if $x_{\underline{t}}^R > 1$ and if m_{in} increases by 1, then this additional initial inventory is used to satisfy demand in $\underline{t}, \dots, \bar{t}$. Furthermore, we decrease $x_{\underline{t}}^S$ and $x_{\underline{t}}^R$ by 1. The inventory $I_{\underline{t}}^R$ remains unchanged for $t \in \{\underline{t}, \dots, \bar{t}\}$. If $x_{\underline{t}}^R = 1$, the previously described changes would lead to a dummy retailer setup at time \underline{t} . Hence, the additional initial inventory is kept in inventory throughout the subplan.

In total, determining w and m_{out} takes $\mathcal{O}(T^4)$ time. Overall, the DP takes $\mathcal{O}(T^5)$ time, which can potentially be reduced since many m_{in} values are infeasible or not used. \square

D Fixed partition model

In this appendix we formulate the *fixed partition model*. In this model the number of contracts and their assignment to the retailer types is fixed, but the lot-sizing plans and side payments of these contracts need to be determined.

For $K \in \mathbb{N}_{\geq 1}$ we are given a K -partition of Θ , denoted by $[\underline{\theta}_k, \bar{\theta}_k]$ with $k \in \mathcal{K} = \{1, \dots, K\}$. It is allowed to have $\underline{\theta}_k = \bar{\theta}_k$, since (the proof of) Lemma 6 also holds for the fixed partition model. The fixed partition model is given by

$$\min \sum_{k \in \mathcal{K}} \left(\int_{\underline{\theta}_k}^{\bar{\theta}_k} \omega(\theta) d\theta \right) \left(F \sum_{t \in \mathcal{T}} y_t^{S(k)} + H \sum_{t \in \mathcal{T}} I_t^{S(k)} + z_k \right)$$

subject to

$$\begin{aligned}
I_0^{S(k)} &= 0, & \forall k \in \mathcal{K}, \\
I_{t-1}^{S(k)} + x_t^{S(k)} &= I_t^{S(k)} + x_t^{R(k)}, & \forall k \in \mathcal{K}, t \in \mathcal{T}, \\
x_t^{S(k)} &\leq M y_t^{S(k)}, & \forall k \in \mathcal{K}, t \in \mathcal{T}, \\
I_0^{R(k)} = I_T^{R(k)} &= 0, & \forall k \in \mathcal{K}, \\
I_{t-1}^{R(k)} + x_t^{R(k)} &= I_t^{R(k)} + d_t, & \forall k \in \mathcal{K}, t \in \mathcal{T}, \\
y_t^{R(k)} \leq x_t^{R(k)} &\leq M y_t^{R(k)}, & \forall k \in \mathcal{K}, t \in \mathcal{T}, \\
y_t^{S(k)}, y_t^{R(k)} &\in \mathbb{B}, & \forall k \in \mathcal{K}, t \in \mathcal{T}, \\
x_t^{S(k)}, x_t^{R(k)}, I_t^{S(k)}, I_t^{R(k)} &\in \mathbb{N}, & \forall k \in \mathcal{K}, t \in \mathcal{T}, \\
f(\theta_k) \sum_{t \in \mathcal{T}} y_t^{R(l)} + h(\theta_k) \sum_{t \in \mathcal{T}} I_t^{R(l)} &= \phi(\mathbf{x}_l | \theta_k), & \forall \theta_k \in [\underline{\theta}_k, \bar{\theta}_k], k, l \in \mathcal{K}, \\
\phi(\mathbf{x}_k | \theta_k) - z_k &\leq \phi^*(\theta_k), & \forall \theta_k \in [\underline{\theta}_k, \bar{\theta}_k], k \in \mathcal{K}, \\
\phi(\mathbf{x}_k | \theta_k) - z_k &\leq \phi(\mathbf{x}_l | \theta_k) - z_l, & \forall \theta_k \in [\underline{\theta}_k, \bar{\theta}_k], k, l \in \mathcal{K}.
\end{aligned}$$

By realising that Lemma 5 is also valid for this model, the IR and IC constraints can be replaced by the following finitely many constraints:

$$\begin{aligned}
\phi(\mathbf{x}_1 | \underline{\theta}) - z_1 &\leq \phi^*(\underline{\theta}), \\
\phi(\mathbf{x}_k | \underline{\theta}_k) - z_k &\leq \phi^*(\underline{\theta}_k), & \forall k \in \mathcal{K} \setminus \{1\}, \\
\phi(\mathbf{x}_K | \bar{\theta}) - z_K &\leq \phi^*(\bar{\theta}), \\
\phi(\mathbf{x}_k | \underline{\theta}_k) - z_k &= \phi(\mathbf{x}_{k-1} | \underline{\theta}_k) - z_{k-1}, & \forall k \in \mathcal{K} \setminus \{1\}.
\end{aligned}$$

The resulting model is a standard mixed integer linear program, provided that we use the linear modelling of ϕ^* as shown in Section 3.2.2.