University of Groningen

Flexible lot sizing in hybrid make-to-order/make-to-stock production planning

Beemsterboer, Bart; Land, Martin; Teunter, Ruud

Published in:
European Journal of Operational Research

DOI:
[10.1016/j.ejor.2017.01.015](https://doi.org/10.1016/j.ejor.2017.01.015)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2017

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Beemsterboer, B., Land, M., & Teunter, R. (2017). Flexible lot sizing in hybrid make-to-order/make-to-stock production planning. *European Journal of Operational Research*, 260(3), 1014-1023.
<https://doi.org/10.1016/j.ejor.2017.01.015>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Production, Manufacturing and Logistics

Flexible lot sizing in hybrid make-to-order/make-to-stock production planning

Bart Beemsterboer, Martin Land, Ruud Teunter*

University of Groningen, Dept. Operations, PO Box 800, 9700 AV Groningen, The Netherlands

ARTICLE INFO

Article history:

Received 23 June 2015

Accepted 9 January 2017

Available online 12 January 2017

Keywords:

Production

Production planning

Lot sizing

Hybrid production

Markov Decision Process

ABSTRACT

Hybrid make-to-order/make-to-stock production systems are difficult to control. Batch production of make-to-stock products allows for efficient capacity usage, but fixed batch sizes can make the system less responsive to make-to-order customers. Using Markov Decision Process modeling, we show that the optimal policy for such hybrid systems varies the lot size in response to make-to-order backlogs as well as stock levels. We evaluate the performance of this flexible lot sizing policy for a two-product hybrid system. We find that it leads to savings of up to 23% compared to policies that use either completely or partly fixed lot sizes. We also find that flexible lot sizing is especially beneficial for systems with a low load and where make-to-stock is important in terms of production capacity usage and cost.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

In many industries, customization has led to a situation where many production companies offer a high variety of products in order to remain competitive. Typically, some ‘standard’ fast moving products will exhibit relatively stable demand, whereas most are slow moving and have more volatile demand. It is well established in the literature that a mix of production approaches is needed for such companies. See Köber and Heinecke (2012) for a case in the agricultural machinery industry, and Romsdal, Arica, Strandhagen, and Dreyer (2013) and Soman, Van Donk, and Gaalman (2007) for applications in the food industry.

To minimize holding cost and avoid obsolescence, it is preferable to make products to order (MTO). However, doing so for all products would imply that volatile demand processes cannot be smoothed and lead to high peaks in the required production capacity, which in turn leads to long delivery times for periods with high demand and under-utilization of capacity in periods with low demand. To overcome these disadvantages, companies can make some of their fast moving products to stock (MTS) in periods with low demand. This kills two birds with one stone, as it better utilizes the capacity when demand is low and reduces the capacity requirement when demand is high by delivering the standard products from stock. Making some products to order and others to stock is referred to as a *hybrid* production strategy.

The difficulty of hybrid production planning is that make-to-order (MTO) and make-to-stock (MTS) production require different planning approaches. MTO production is triggered by demand, allowing for customization of products, and usually focuses on a timely delivery of the products. Standardized MTS products are preferably delivered of the shelf to guarantee high service levels and can be produced in lots to limit setup time and/or costs. Because of these different objectives, hybrid MTO-MTS systems are difficult to control.

When it comes to production planning of hybrid systems, many firms may be tempted to restrict to simple rules that are known to perform well in pure MTO or MTS systems, for instance by using an (r, Q) reordering policy for MTS items that triggers a replenishment batch of size Q when an item's inventory position drops to a level r . However, using fixed MTS lot sizes may be problematic for hybrid production systems. Batch production is generally applied for efficiency reasons, allowing a machine setup to be used for a series of products. In hybrid systems, however, working on large MTS batches may lead to conflicts as MTO orders have to be delivered within a restricted lead time. Put differently, a focus on a timely delivery of MTO orders may impose restrictions on the MTS lot size. Hence, intuitively it would be preferred to use a more flexible lot sizing approach that takes both the stock level and the MTO order book into account, and even during the production of the batch.

As will appear from our literature review in the next section, few authors have addressed the importance of lot size flexibility for hybrid MTO-MTS production systems. Furthermore, surprisingly, none of them have considered machine setup times although these obviously affect capacity usage and play a key role in

* Corresponding author.

E-mail addresses: b.j.beemsterboer@gmail.com (B. Beemsterboer), m.j.land@rug.nl (M. Land), r.h.teunter@rug.nl (R. Teunter).

determining lot sizes in the MTS literature. The main contribution of this paper is to study the importance of flexible lot sizing for hybrid MTO-MTS production systems, taking setup times into account. As we are the first to do so, our study is of an exploratory nature. We study a stylized setting with one product MTO and one product MTS. For this relatively simple setting, we are able to computationally determine the optimal production strategy using Markov Decision Programming (MDP). This approach cannot easily be extended to larger, realistic systems, but we are aiming at general insights on (flexible) lot sizing that may guide practitioners. Indeed, a key finding is that in order to achieve a good performance (in terms of holding and lateness costs), it is essential that the lot size can be adapted at any time in response to changes in demand. We call this full lot sizing flexibility and show numerically that it can lead to cost savings of around 20% compared to fixing the lot size when a batch is started.

The remainder of this paper is organized as follows. Relevant literature is reviewed in Section 2. We present our model and Markov Decision Process formulation in Section 3. In Section 4, we discuss the optimal policy for an illustrative example and examine the resulting MTS lot sizes. We conduct a numerical investigation in Section 5, in which we compare our flexible model with two reference models that use fixed batch sizes. Conclusions and directions for further research are given in Section 6.

2. Literature review

Lot sizing was first considered in the literature for pure MTS production systems. In early models, demand and production rates are often assumed non-random, so it is logical to use fixed lot sizes as in the traditional (r, Q) model. See, for instance, Elmaghraby (1978) for a review on the multi-item case, known as the Economic Lot Scheduling Problem (ELSP). Later models consider random demand, so that it becomes beneficial to let lot sizes depend on the state of the production system, giving rise to the Stochastic ELSP (SELSP). See Sox, Jackson, Bowman, and Muckstadt (1999) and Winands, Adan, and Van Houtum (2011) for review studies.

One of the first to consider hybrid production systems was Williams (1984), who addresses several questions, such as how to prioritize MTO and MTS items, assuming the traditional (r, Q) structure with a fixed lot size. Rajagopalan (2002) uses the same policy for MTS items, but also consider the decision which items to make to stock and which to make to order, besides selecting the MTS policy parameters. Federgruen and Katalan (1999) analyze the impact of adding an MTO item to an MTS production system. The production runs for the MTS items are assumed to last until a predefined order-up-to level is attained. Hence, the resulting batch size depends on the state of the system, but only through the inventory level of the product that is currently manufactured. Chang, Pai, Yuan, Wang, and Li (2003) and Wu, Jiang, and Chang (2008) propose special-purpose planning methods for hybrid production systems, based on a strict priority of MTO orders and on the rationale that MTS production is used to fill the remaining capacity. Hence, the lot sizes depend on the amount of workload in the system.

Soman, van Donk, and Gaalman (2006) compare four existing SELSP procedures in the context of a hybrid production environment: the method known as the Economic Manufacturing Quantity and those proposed by Vergin and Lee (1978), Leachman and Gascon (1988), and Fransoo, Sridharan, and Bertrand (1995). In a follow-up study, Soman et al. (2007) execute a case study in a food processing company, in which they propose an extensive planning approach. In both papers, Soman et al. (2007) allow the MTS lot size to depend on the state of all products through the run-out times of MTS units and remaining lead times of MTO orders. Sev-

eral other articles consider hybrid production planning without addressing MTS lot sizing (e.g. Hadj Youssef, Van Delft, & Dallery, 2004; Sox, Thomas, & McClain, 1997).

Contrary to the previous approaches, the following models implicitly use a more flexible way of determining lot sizes. Carr and Duenyas (2000) and Iravani, Liu, and Simchi-Levi (2012) consider contract manufacturers, who produce on an MTS basis for contracted customers and on an MTO basis for others. Beemsterboer, Land, and Teunter (2016) determine the potential benefits of a planning approach for hybrid production systems that fully considers the state of both MTS and MTO. All use simple, two-product Markov Decision Process models to analyze MTO order acceptance and/or planning decisions in hybrid production systems. Carr and Duenyas (2000) and Iravani et al. (2012) determine decisions based on the MTS inventory level and the amount of MTO products, whereas Beemsterboer et al. (2016) additionally take the remaining MTO lead times into account. In these models, the MTS lot size does not have to be determined at the start of the batch, but for every decision epoch, the production decision is reviewed so that the eventual lot size is the result of a series of production decisions that can be based on demands during the production of the batch. However, Carr and Duenyas (2000), Iravani et al. (2012) and Beemsterboer et al. (2016) do not include machine setup times in their models. As a consequence, the capacity advantage, which is the most important reason to produce standardized products in batches, is not considered in these models, and the benefits of flexible lot sizing cannot be determined in this respect.

Our paper aims at demonstrating the benefits of a flexible lot sizing approach for MTS products in hybrid production systems. It is flexible in that lot sizes can be state dependent but also in that the size does not have to be determined at the start of the batch. We analyze a Markov Decision Process model of a hybrid production system in order to determine optimal production decisions for each period, based on the most up-to-date state of the system. From a modeling perspective, we follow the approach of Carr and Duenyas (2000), Iravani et al. (2012) and Beemsterboer et al. (2016). However, we include a setup time in our model so that efficiency considerations of batch production come into play. Moreover, to account for the timing of MTO orders we include an MTO lead time allowance, whereas Carr and Duenyas (2000) and Iravani et al. (2012) instead use a proxy cost so that all MTO orders are modeled as being 'late' from the moment they arrive.

3. Model formulation

In this section, we model a hybrid production system. In Section 3.1, we provide the framework of the model, which we then formulate as a Markov Decision Process in Section 3.2. A list of notations is given in Table 1.

3.1. Production system

We consider a system that manufactures two products, one on an MTO basis and one on an MTS basis. The system can only manufacture one product at a time. For both products, the system requires a setup. We assume that MTO products are customized and that the system therefore requires a setup for each product. By contrast, MTS products are standardized products and only one setup is needed to produce a batch.

In order to formulate our problem as a Markov Decision Process in the next section, we make the following simplifying assumptions. We use a discrete time framework and we assume that the setup and unit processing time of both products are all equal to each other and given as one period. In order to ensure finiteness of the state space, we will use bounded demand distributions and we define a maximum MTS inventory level I and a maximum amount

Table 1
Notations.

Demand parameters	d_o	Mean MTO demand per period
	d_s	Mean MTS demand per period
	$d_o^{\max}$	Maximum MTO demand per period
	$d_s^{\max}$	Maximum MTS demand per period
Cost parameters	h	MTS holding costs (per unit, per period)
	g	MTO lateness costs (per unit, per period)
	b_s	MTS lost sales costs (per unit)
	b_o	MTO lost sales costs (per unit)
System parameters	K	Maximum # MTO orders in the system
	L	MTO lead time allowance
Demand probabilities	$p_o(j)$	Probability of j MTO demands in a period
	$p_s(j)$	Probability of j MTS demands in a period
State space	i	MTS inventory level
	k_l	# MTO orders in the system for l periods ($l = 0, \dots, L - 1$)
	k_L	# late MTO orders
	m	Machine setup status
Other	k	Total number of MTO orders

of MTO orders in the system K . We assume that excess MTO demand is lost. The maximum inventory level L is not treated as a system control parameter, but set large enough so as not to affect the numerical results.

Demands for MTO and MTS are random and follow discrete and bounded probability distributions, and we assume that these are independent. The probability of j demands in a period is denoted as $p_o(j)$ for MTO and $p_s(j)$ for MTS. We denote the mean demands per period as d_o for MTO and d_s for MTS and maximum demands as $d_o^{\max}$ (MTO) and $d_s^{\max}$ (MTS), so we have $d_o = \sum_{j=0}^{d_o^{\max}} j p_o(j)$ and $d_s = \sum_{j=0}^{d_s^{\max}} j p_s(j)$. MTS demands have to be satisfied from inventory on hand or are otherwise lost. MTO orders have a promised lead time of L periods (not including the period in which the order arrives). Late orders should still be delivered, but incur a penalty cost per time unit. In our system, MTO orders differ only in their remaining time until the due date, and therefore it is optimal to produce them in a first come, first served fashion to realize the lowest lateness costs.

In each period, the production system may either perform a setup (MTO, MTS) or manufacture a product (MTO, MTS). Logically, the system may only manufacture a product when it is set up for the concerning product type. Producing a unit of MTO brings the system back to a state where no setup is completed, while producing a unit of MTS maintains the setup status so that another MTS unit can be produced in the succeeding period. An MTO setup, which concerns customized products, can only be performed when there is at least one MTO order in the system. By contrast, an MTS setup (concerning the standardized product) can always be performed. We also allow the system to maintain the MTS setup status without producing a unit, which makes idling redundant. This is modeled by doing another setup. Finally, we assume that performing or maintaining a setup does not constrain the machine to produce the concerning product type in the succeeding period, i.e. production may be 'aborted' after a setup.

The order of events in each period is as follows. First, the system decides on an action. Second, if this action is to produce a unit of MTO or MTS, it becomes available. Third, demands occur and are either accepted/satisfied or lost.

The objective is to minimize the expected costs per period. These costs are composed of per period holding costs for MTS products, lateness costs for MTO products, and lost sales costs for MTO and MTS products. The MTS holding costs are given as h per unit per period, late MTO orders are penalized by a cost of g per period and lost sales costs are b_o per missed MTO demand and b_s per missed MTS demand. Note that we do not include an MTS setup cost, as setups are already 'penalized' by taking up potential production time. Both a positive setup time (as we consider) and

a positive setup cost provide an incentive for MTS batching. Our sole focus on setup time allows any batching in an optimal policy to directly be attributed to the setup time. It also reduces the state space of the Markov Decision Process formulation (presented in the next section) for finding an optimal policy.

As the production decision is determined period by period, and the MTS lot size is not decided in advance, our model follows the flexible approach described in Section 1.

3.2. Markov Decision Process

We model the problem described above as a discrete-time Markov Decision Process, which requires the definition of a state space $\mathcal{S}$, an action space $\mathcal{A}$, costs, and transition probabilities.

3.2.1. State space

The state of the system is described by the MTS inventory level i , the numbers of accepted MTO orders k_l , $l = 0, \dots, L - 1$, that have been in the system for l periods (not including the period of arrival) and so have $L - l$ periods remaining until the due date, the number of outstanding late orders k_L , and the machine setup status m , indicating whether the machine is not set up ($m = 1$) which happens after MTO production has finished, set up for MTO ($m = 2$), or set up for MTS ($m = 3$). Denoting the state space by $\mathcal{S}$, a state $s \in \mathcal{S}$ is described by the tuple $s = (i, k_0, \dots, k_L, m)$. For convenience, we refer to the MTO part of the state space $(k_0, \dots, k_L)$ as the *order state*. We further define k as the total number of MTO orders in the system, i.e. $k = \sum_{j=0}^L k_j$.

3.2.2. Action space

The action space is denoted by $\mathcal{A} = \{ou, op, su, sp\}$, where these actions correspond to, respectively, an MTO setup, MTO production, an MTS setup, and MTS production. Action *ou* (MTO setup) is admissible when the system is not set up for MTO yet, and only when MTO orders are at hand, i.e. in states $s = (i, k_0, \dots, k_L, m)$ with $m \neq 2$ and $k \geq 1$. Action *op* (MTO production) is admissible when the system is set up for MTO and again when MTO orders are at hand, i.e. in states s with $m = 2$ and $k \geq 1$. Action *su* (MTS setup) is admissible in any state and action *sp* (MTS production) is admissible when the system is set up for MTS and when the inventory level is below its maximum, i.e. in states s with $m = 3$ and $i = 0, \dots, L - 1$.

3.2.3. Costs

We now define the expected costs in state s if action a is taken, denoted as $c^a(s)$. Writing $[x]^+$ for $\max\{x, 0\}$, these are given as

$$c^{ou}(s) = c^{su}(s) = hi + gk_L + b_s \sum_{j=0}^{d_s^{\max}} (p_s(j)[j - i]^+) + b_o \sum_{j=0}^{d_o^{\max}} (p_o(j)[k + j - K]^+),$$

$$c^{op}(s) = hi + gk_L + b_s \sum_{j=0}^{d_s^{\max}} (p_s(j)[j - i]^+) + b_o \sum_{j=0}^{d_o^{\max}} (p_o(j)[k + j - K - 1]^+),$$

$$c^{sp}(s) = hi + gk_L + b_s \sum_{j=0}^{d_s^{\max}} (p_s(j)[j - i - 1]^+) + b_o \sum_{j=0}^{d_o^{\max}} (p_o(j)[k + j - K]^+).$$

The first and second terms represent the holding and lateness costs, respectively. The third and fourth terms are, respectively, the expected MTS and MTO lost sales costs. For action *op* (MTO production) and action *sp* (MTS production), the product that is manufactured becomes available before the demand occurs so that the expected excess demand of the concerning product type decreases by one.

3.2.4. Transition probabilities

We next discuss the transition probabilities $\pi^a(s, s')$, denoting the probability that the state at the end of a period is $s' = (i', k'_0, \dots, k'_L, m')$ given that the state at the beginning is s and action a is taken. For simplicity, we restrict the discussion to transitions for which lost sales can be ignored, i.e. transitions to states for which $i' > 0$ and $\sum_{j=0}^L k'_j < K$. For other transition probabilities, which we provide in the Appendix A, the possibility of lost sales leads to slightly different expressions, but the dynamics are essentially the same. We first consider the setup status transitions for each action, followed by the order state transitions, and then combine those into complete system state transitions with corresponding probabilities.

Following the definitions above, action *ou* (MTO setup) can be applied if $m \neq 2$ and sets the setup status to $m' = 2$. Action *op* (MTO production) can be applied if $m = 2$, resetting the setup status to $m' = 1$ afterward. Action *su* (MTS setup) can be applied regardless of the setup status but sets it to $m' = 3$ afterward, and, finally, action *sp* (MTS production) applies if $m = 3$ and maintains that status. Accordingly, we define setup status transition sets for each action:

$$\mathcal{M}_{ou} = \{(s, s') | m \neq 2, m' = 2\}, \quad (1)$$

$$\mathcal{M}_{op} = \{(s, s') | m = 2, m' = 1\}, \quad (2)$$

$$\mathcal{M}_{su} = \{(s, s') | m' = 3\}, \quad (3)$$

$$\mathcal{M}_{sp} = \{(s, s') | m = 3, m' = 3\}. \quad (4)$$

The order state transitions only depend on whether MTO production takes place (action *op*) or not (actions *ou*, *su*, *sp*). If it does not, then the order state transitions are as described by the following set.

$$\mathcal{O}_I = \{(s, s') | k'_l = k_{l-1} \text{ for } l = 1, \dots, L-1 \text{ and } k'_L = k_L + k_{L-1}\} \quad (5)$$

If production does take place, then an order with the smallest remaining lead time will be satisfied (recalling from Section 3.1 that first come, first served is optimal). Denoting 'the longest time that an order has been in the system' by $l^* = \max\{l | k_l \geq 1\}$, and defining $k'_l = k_l$ for $l \in \{0, \dots, L\} \setminus l^*$ and $k'_{l^*} = k_{l^*} - 1$, we obtain the second order transition set.

$$\mathcal{O}_{II} = \{(s, s') | k'_l = k'_{l-1} \text{ for } l = 1, \dots, L-1 \text{ and } k'_L = k'_L + k'_{L-1}\} \quad (6)$$

Combining the transition sets of setup status and order state, and using independence of MTO and MTS demands, we obtain the following complete state transition probabilities.

$$\pi^{ou}(s, s') = p_o(k'_0) p_s(i - i') \quad \text{for } (s, s') \in \mathcal{M}_{ou} \cap \mathcal{O}_I, \quad (7)$$

$$\pi^{op}(s, s') = p_o(k'_0) p_s(i - i') \quad \text{for } (s, s') \in \mathcal{M}_{op} \cap \mathcal{O}_{II}, \quad (8)$$

$$\pi^{su}(s, s') = p_o(k'_0) p_s(i - i') \quad \text{for } (s, s') \in \mathcal{M}_{su} \cap \mathcal{O}_I, \quad (9)$$

$$\begin{aligned} \pi^{sp}(s, s') &= p_o(k'_0) p_s(i - i' + 1) & \text{for } (s, s') \in \mathcal{M}_{sp} \cap \mathcal{O}_I, \\ \pi^a(s, s') &= 0 & \text{elsewhere.} \end{aligned} \quad (10)$$

4. Policy structure

Before performing an extensive numerical exploration using various parameterizations, we start by discussing the typical structure of the optimal policy for an illustrative example that has a relatively small state space, allowing all states to be represented in a table. We select the smallest possible (but positive) values for the maximum demands per period d_o^{max} and d_s^{max} , i.e. $d_o^{max} = d_s^{max} = 1$. This means that we define the demand per period of both product types as Bernoulli distributions. For the Bernoulli distribution, the mean equals the probability of a success, and so we set those probabilities to d_o and d_s . For this example their values are selected as 0.25 for both MTO and MTS, i.e. $d_o = d_s = 0.25$. Notice that producing an MTO unit (setup and production) takes a total of two periods, so an MTO demand of $d_o = 0.25$ leads to a 'machine load' contribution of 50%, i.e. 50% of the machine capacity is needed to satisfy all MTO demand. For MTS production, the required number of periods per unit lies between one, in the extreme case of producing an infinitely long batch, and two, in case we would perform a setup for each unit. So $d_s = 0.25$ implies a machine load contribution of MTS between 25% and 50% if no demands are lost and hence a total machine load of 75–100%.

The remaining parameters are as follows. We set the MTO lead time allowance to $L = 3$, which allows making a distinction between more urgent and less urgent MTO orders. The maximum amount of MTO orders in the system is set to $K = 5$. Finally, we normalize the MTS holding costs to $h = 1$ and we set $g = 8$ and $b_s = b_o = 250$, i.e. the unit MTO lateness costs is 8 times larger and the lost sales costs of both product types are 250 times larger than the MTS holding costs. These cost ratios guarantee that the resulting optimal policy avoids MTO lateness before lost sales come into view, and also sufficiently avoids MTS lost sales. This parameterization yields 36 different order states and we have a resulting optimal policy that requires a maximum inventory level of 5. As we also have three setup status values, there are 648 states altogether.

Table 2 shows the optimal policy for this example. The optimal actions are split into three columns, each corresponding to a machine setup status (not set up, set up for MTO, and set up for MTS). Each column contains six smaller columns corresponding to all (attainable) inventory levels. Each row corresponds to an order state, containing, from left to right, the numbers of MTO orders just received, received one period earlier, received two periods earlier, and received longer ago so that they are late. Note that, because performing an MTO setup is not allowed when there are no orders in the system, the order states (0, 0, 0, 0) and (1, 0, 0, 0) in the MTO setup status cannot be attained by definition.

The optimal policy shown in Table 2 reveals several insights. First of all, when the machine is not set up yet, the system performs an MTS setup either when there is no MTS inventory or when there are no MTO orders left in the system, and performs an MTO setup in all other cases. Second, the system always continues producing MTO when an MTO setup is completed. This is logical, as a switch to MTS would imply that the MTO setup that had just been performed is wasted, spoiling valuable production capacity. When set up for MTS, MTS production continues until a particular inventory level is reached, which depends on the order state. For the order state (0, 0, 0, 0), the system stops producing MTS from an inventory level of five, after which the system maintains its MTS setup status. In the order state (1, 0, 0, 0), MTS production continues until an inventory level of four is reached, after which the system switches to MTO. In all other order states, the system switches to MTO at an inventory level of three.

Although the differences in these 'switching' levels are small (they can be much larger for other examples), these reveal

Table 2
Optimal policy for illustrative example Actions: ou = MTO setup, op = MTO production, su = MTS setup, sp = MTS production Parameter settings: $d_o = d_s = 0.25$, $L = 3$, $K = 5$, $g = 8$, $b_s = b_o = 250$.

Setup status →	Not set up ($m = 1$)						For MTO ($m = 2$)						For MTS ($m = 3$)					
Inventory level →	0	1	2	3	4	5	0	1	2	3	4	5	0	1	2	3	4	5
Order state ↓																		
(0, 0, 0, 0)	su	su	su	su	su	su							sp	sp	sp	sp	sp	su
(1, 0, 0, 0)	su	ou	ou	ou	ou	ou							sp	sp	sp	sp	ou	ou
(0, 1, 0, 0)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 1, 0, 0)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 0, 1, 0)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 0, 1, 0)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 1, 1, 0)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 1, 1, 0)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 0, 0, 1)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 0, 0, 1)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 1, 0, 1)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 1, 0, 1)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 0, 1, 1)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 0, 1, 1)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 1, 1, 1)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 1, 1, 1)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 0, 0, 2)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 0, 0, 2)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 1, 0, 2)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 1, 0, 2)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 0, 1, 2)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 0, 1, 2)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 1, 1, 2)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 1, 1, 2)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 0, 0, 3)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 0, 0, 3)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 1, 0, 3)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 1, 0, 3)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 0, 1, 3)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 0, 1, 3)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 1, 1, 3)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 0, 0, 4)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 0, 0, 4)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 1, 0, 4)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(1, 1, 0, 4)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou
(0, 0, 0, 5)	su	ou	ou	ou	ou	ou	op	op	op	op	op	op	sp	sp	sp	sp	ou	ou

important benefits of a state-dependent switching level. In particular the difference in switching level between ‘no MTO orders’ (first row) and ‘MTO orders present’ (all other rows) is beneficial; if there are no MTO orders, the production capacity is used to restock up to high inventory levels, allowing the production of MTS to be less urgent when MTO orders arrive. In addition, we observe a different switching level for the second row and all lower rows. This is because in the order state (1, 0, 0, 0), the single MTO order is not urgent in the sense that it can still be completed on time if the production is postponed by one period, while in all other order states a lateness penalty will be incurred if the system does not switch to MTO immediately. Other considered examples confirmed these insights. Hence, we observe that an optimal policy may indeed benefit from the possibility to determine a production decision based on the state of both product types.

Fig. 1 shows a histogram of the distribution of the realized batch sizes for this example. Apparently, a high variation of the MTS batch size is needed to obtain minimum costs. The optimal fixed batch size in this example would have been 3, but the optimal flexible batch size is often much smaller with an average of 2.09 and a standard deviation of 1.35. It equals one in 46% and two in another 23% of the cases, and 15% of the batches become larger than size 3. This suggests that fixing the batch size may be far from optimal. A more extensive numerical study in the next section will indeed reveal that allowing the batch size to vary leads to cost savings of around 10% in a variety of settings.

Fig. 1. Distribution of the flexible MTS batch sizes in the optimal policy of the example.

Table 3

Parameters and results of the experiments.

#	Parameters								Machine load (%)	Costs			Savings relative	
	d_o	d_s	L	K	h	g	b_s	b_o		NF	PF	FF	to PF (%)	to NF (%)
1	0.25	0.25	7	8	1	8	250	250	75 – 100	5.0	4.8	4.5	6.0	9.5
2	0.20	0.25	7	8	1	8	250	250	65 – 90	3.6	3.4	3.0	10.9	14.6
3	0.30	0.25	7	8	1	8	250	250	85 – 110	7.8	7.6	7.4	2.7	5.5
4	0.25	0.20	7	8	1	8	250	250	70 – 90	3.8	3.6	3.4	6.2	10.2
5	0.25	0.30	7	8	1	8	250	250	80 – 110	6.5	6.3	6.0	4.5	7.9
6	0.25	0.25	6	8	1	8	250	250	75 – 100	5.3	5.1	4.9	5.4	8.5
7	0.25	0.25	8	8	1	8	250	250	75 – 100	4.6	4.5	4.2	6.6	10.3
8	0.25	0.25	7	6	1	8	250	250	75 – 100	5.0	4.8	4.5	6.0	9.6
9	0.25	0.25	7	10	1	8	250	250	75 – 100	5.0	4.8	4.5	6.0	9.4
10	0.25	0.25	7	8	0.5	8	250	250	75 – 100	3.2	3.1	2.9	5.1	8.6
11	0.25	0.25	7	8	2	8	250	250	75 – 100	7.6	7.4	6.9	6.8	10.1
12	0.25	0.25	7	8	1	4	250	250	75 – 100	4.3	4.1	3.8	8.2	11.7
13	0.25	0.25	7	8	1	16	250	250	75 – 100	5.4	5.3	5.0	4.7	7.2
14	0.25	0.25	7	8	1	8	125	250	75 – 100	4.1	4.0	3.8	6.3	8.1
15	0.25	0.25	7	8	1	8	500	250	75 – 100	5.4	5.2	4.9	6.4	10.1
16	0.25	0.25	7	8	1	8	250	125	75 – 100	5.0	4.8	4.5	6.1	9.7
17	0.25	0.25	7	8	1	8	250	500	75 – 100	5.0	4.8	4.5	5.9	9.4

5. Numerical investigation

This section provides a numerical study of the benefits of flexible MTS lot sizing compared to fixed batch sizes, by comparing the performance of the model defined in Section 3 with two reference models that use fixed MTS batch sizes. We introduce these reference models in Section 5.1. We then discuss the design of a first set of experiments in Section 5.2 and we present the results in Section 5.3. As these do not reveal the effect of the demand mix, we conducted additional experiments which we discuss in Section 5.4.

5.1. Reference models

Recall from Section 1 that besides the flexible approach, two other lot sizing approaches exist in the literature. These fix the MTS lot sizes either from the start of the batch or throughout the entire policy. We adopt these approaches into the following reference models. In the interest of consistency, the model developed in Section 3 will be referred to as *Fully Flexible*.

Partly flexible. When starting an MTS batch, its size is selected in advance, and switching during production is not possible. Switching after an MTO setup is also not possible for consistency. Hence, for each state, an optimal batch size is selected.

Not flexible. The MTS batch size is fixed throughout the entire policy, i.e. one optimal batch size is obtained across all states.

The optimization approach for the reference models is very similar to that for Fully Flexible. For Partly Flexible, we extend the Markov Decision Process defined in Section 3.2 so that each possible MTS batch size is associated with an action. For Not Flexible, we also use an adapted Markov Decision Process, but as the fixed MTS batch size has to be provided beforehand, we repeatedly solve instances of this model using a one-dimensional grid search, each with a different predefined MTS batch size, in order to obtain the overall optimal fixed batch size.

5.2. Experimental design

We first select and study a base case, followed by a sensitivity study in which we vary more parameters. The demand and cost parameters of the base case are selected equal to those of the illustrative example in Section 4: $d_o = d_s = 0.25$, $h = 1$, $g = 8$ and $b_s = b_o = 250$. Recall that $d_o = d_s = 0.25$ implies a machine load of 75–100%. We use Bernoulli demand distributions again, implying $d_o^{max} = d_s^{max} = 1$, which keeps the state space and calculation

times acceptable, also for the larger values of K and L that we will consider.

The maximum number of MTO orders K is mainly included in the model to bound the size of the state space and we prefer to have a value that is large enough so that it has little influence on the behavior of the optimal policies. Pretesting showed that a value of $K = 8$ is large enough for our purposes. To allow for completion of larger MTS batches we set our lead time allowance for MTO to $L = 7$.

In order to examine the dependence of the benefits of flexible lot sizing on the parameters, we conduct a sensitivity analysis in which we select a smaller and a larger value for each of the parameters d_o , d_s , L , K , h , g , b_s and b_o , as specified in Table 3. The parameters are varied one by one, i.e. all other parameters are fixed at their base value. The demand distribution type is not varied. This gives a total of 17 instances to be tested for each of the three models. In addition, we performed a full factorial design experiment where each of the eight parameters takes on either the low or the base value, giving $2^8 = 256$ instances in total. A full factorial design with three values per parameter would be too time-consuming. Moreover, ensuring that all settings are relevant (i.e. total demand does not exceed capacity) limits the range of values that can be considered per parameter, which is why we consider only the base and low values in the full factorial design. We center the main discussion around the 17 instances where we vary one parameter at a time. However, where relevant, we will refer to the results from the full factorial study that are presented in Appendix B.

5.3. Results

The right part of Table 3 overviews the results. The experimental setup allows us to distinguish two types of savings resulting from: (1) the flexibility to select various lot sizes, and (2) the flexibility to determine lot sizes once production has started. The latter type, represented by the savings of Fully Flexible relative to Partly Flexible, yields around 6% and varies considerably, ranging from 3% to 11% for these experiments. The flexibility to select various lot sizes adds another 3–4%, bringing the total savings of the fully flexible versus the not flexible model to around 10% for most of our experiments, ranging from 6% to 15%.

Before discussing the key sensitivity findings, we observe that varying the maximum amount of MTO orders (experiments 8 and 9) has a very limited effect on the costs and savings. This is what we aimed for as we introduced this parameter for technical reasons (see Section 3). We also remark that all main findings on what parameters affect the savings from flexible lot sizing most

Table 4
Effect of varying demand mix.

Total demand rate	Demand parameters (d_o , d_s)			Savings relative to NF (%)		
	Min. machine load			Min. machine load		
	80%	75%	70%	80%	75%	70%
0.4	(0.40, 0.0)	(0.35, 0.05)	(0.30, 0.1)	0.0	1.0	7.8
0.45	(0.35, 0.1)	(0.30, 0.15)	(0.25, 0.2)	4.4	7.1	10.2
0.5	(0.30, 0.2)	(0.25, 0.25)	(0.20, 0.3)	5.7	9.5	13.5
0.55	(0.25, 0.3)	(0.20, 0.35)	(0.15, 0.4)	7.9	11.0	16.8
0.6	(0.20, 0.4)	(0.15, 0.45)	(0.10, 0.5)	9.7	13.5	22.7

and in what direction are confirmed by the full factorial results presented in Appendix B.

Examining the results of the individual experiments, we distinguish the following two main effects. Considering experiments 2–5, in which we vary the demand rates, we observe that lower demand rates lead to higher savings. The effect is stronger for a varying MTO demand rate, but essentially applies to both. Production systems with a larger total demand rate have a higher machine load and, as a result, incur more costs. The savings in absolute terms, however, are not notably lower than systems with a lower machine load. However, as these systems incur overall more costs, the *relative* savings become lower.

Another effect is revealed by considering the experiments in which we vary the cost parameters. Experiments 11, 12, 15 and 16 lead to larger savings than the base case, while their counterparts (10, 13, 14 and 17) lead to smaller savings. Either when an MTO cost parameter is decreased, or when and MTS cost parameter is increased, the savings become larger. From these observations, we learn that whenever MTS becomes more costly, more is to be gained by varying lot sizes. This is intuitive: the models differ in the way they handle MTS production and hence when this type becomes more important from a cost perspective, more is to be gained. Considering experiments 6 and 7, where the MTO lead time allowance is varied, the same explanation applies. When MTO orders have more time to be produced, MTO lateness costs are incurred less often, and MTS becomes relatively more important from a cost perspective. In addition, when the MTO lead time allowance increases, there is less need for sophisticated MTS planning and therefore lower relative savings for the proposed advanced scheduling approaches.

The experiments so far do not allow us to draw clear conclusions with respect to the demand mix. Experiments 2 and 3 show that an increasing MTO demand rate, obviously leading to larger share of MTO, leads to smaller savings. By contrast, experiments 4 and 5 show that a decreasing MTS demand rate, also leading to a larger share of MTO, leads to increased savings. These seemingly contradictory results are related to the fact that the experiments could not isolate the demand mix (variations), because changing the demand rates not only affects the demand mix but also the machine load, for which we already have observed the effects. In order to find the effect of the demand mix, we conduct a second set of experiments.

5.4. Effect of demand mix

It is impossible to keep both the total demand rates and the machine load bounds constant when varying the product mix, because only MTS products can be batch produced. Therefore, in these additional experiments we apply a two-dimensional setting in which we vary demand rates as shown in Table 4. We increase the share of MTS in the demand mix in two ways: with a fixed total demand rate and a decreasing minimum machine load (from left to right), and with a fixed minimum machine load and an increasing total demand rate (from top to bottom). The two ways have opposite effects on the actual machine load. However, both

ways show increasing savings. Hence, we may conclude that, in general, an increasing share of MTS leads to higher savings.

The savings range from 0% (with $d_o = 0.4$ and no MTS demand) to 23% (with $d_o = 0.1$ and $d_s = 0.5$). Logically, when there is no MTS demand, no batching can be applied at all so there are no savings. By contrast, when MTS constitutes a large share of the total demand, the savings are substantial. Again, it is explained through the main difference of the models, which differ especially in the way they plan MTS production. Fully Flexible allows for better control of the inventory level than the other models, leading to both lower inventory holding costs and lower MTS lost sales costs. Thus, overall, we have shown that when MTS becomes more important either from a cost or from a demand perspective, the benefits of flexible lot sizing increase.

6. Conclusion

We studied the benefits of varying MTS lot sizes in hybrid production systems. We proposed a ‘fully flexible’ approach, where lot sizes are not fixed at any time but rather evolve in response to demands during the production of the batch.

Analysis of this approach with a Markov Decision Process has shown that the optimized fully flexible approach leads to considerable variation in lot sizes in order to reduce costs. For example, in a situation where an optimal fixed batch size of 3 would be applied, optimization under flexibility was shown to bring the average lot size down to 2.09 with a standard deviation of 1.35.

We compared our approach with two approaches that consider lot sizes that are fixed either from the start of the batch (partly flexible) or throughout the entire policy (not flexible) for a two-product hybrid system. The numerical experiments showed that this fully flexible approach may save up to 23% of the costs. Our design allowed us to attribute this percentage to two types of flexibility. The majority of the savings, up to 19%, is due to flexibility created by the possibility to review the production decision in each period. A smaller part of the benefits was attributed to the flexibility to differentiate lot sizes based on the state of the production system.

The savings are especially large in three situations. First, relatively high savings are obtained for low load situations. Second, a large share of MTS demand corresponds with higher savings, and third, when the MTS costs are relatively large, we obtain high savings. This confirms the intuition that savings will relate strongly to MTS characteristics, as lot sizing also relates to the MTS products.

Our results also bear relevance for planning hybrid MTO-MTS production in Enterprise Resource Planning (ERP). The main difference between MTO and MTS from an ERP perspective is that the former are based on sales orders (requiring linkage between sales order and production process modules), whereas the latter rely on sales forecasts. The ERP steps of converting those orders or forecasts into production plans are similar, and are therefore relatively easy to combine. However, our results show that to achieve a good system performance, ERP processes should build in flexibility to adapt lot sizes to changing conditions.

We conducted this study in a setting with a number of simplifying assumptions that allowed us to use an analytical optimization procedure and study the lot sizing decision in isolation. This helped us clarify the mechanisms that determine the advantage of flexible lot sizes. A next step in research could be to distill lot sizing rules based on our findings, and analyze them using queuing theory which may lead to further insights based on explicit formulas. Another promising avenue of further research is to test the performance of a fully flexible approach in more complex settings. Important extensions include: multiple products; variations in the production times, setup times or MTO lead times; and considering

a more detailed representation of a production facility, e.g. a job shop.

Appendix A. Transition probabilities of the Markov Decision Process

This appendix provides a complete overview of the transition probabilities of the Markov Decision Process introduced in Section 3.2. Recall from this section the state sets that describe the setup status and order state transitions. These are as provided by (1)–(6) as follows (we refer to Section 3.2 for explanation).

$$\mathcal{M}_{ou} = \{(s, s') | m \neq 2, m' = 2\},$$

$$\mathcal{M}_{op} = \{(s, s') | m = 2, m' = 1\},$$

$$\mathcal{M}_{su} = \{(s, s') | m' = 3\},$$

$$\mathcal{M}_{sp} = \{(s, s') | m = 3, m' = 3\},$$

$$\mathcal{O}_I = \{(s, s') | k'_l = k_{l-1} \text{ for } l = 1, \dots, L-1 \text{ and } k'_L = k_L + k_{L-1}\},$$

$$\mathcal{O}_{II} = \{(s, s') | k'_l = k''_{l-1} \text{ for } l = 1, \dots, L-1 \text{ and } k'_L = k''_L + k''_{L-1}\}.$$

We now provide the transition probabilities in four sets. The first refers to transitions to states with $i' > 0$ and $k' < K$, for which lost sales do not occur. These are given by (7)–(10) in Section 3.2 as

$$\pi^{ou}(s, s') = p_o(k'_0) p_s(i - i') \quad \text{for } (s, s') \in \mathcal{M}_{ou} \cap \mathcal{O}_I,$$

$$\pi^{op}(s, s') = p_o(k'_0) p_s(i - i') \quad \text{for } (s, s') \in \mathcal{M}_{op} \cap \mathcal{O}_{II},$$

$$\pi^{su}(s, s') = p_o(k'_0) p_s(i - i') \quad \text{for } (s, s') \in \mathcal{M}_{su} \cap \mathcal{O}_I,$$

$$\pi^{sp}(s, s') = p_o(k'_0) p_s(i - i' + 1) \quad \text{for } (s, s') \in \mathcal{M}_{sp} \cap \mathcal{O}_I.$$

The second set concerns transitions to states with $i' > 0$ and $k' = K$. These states are not only reached with exactly $K - k$ MTO demands ($K - k + 1$ for action *op*), but also with higher numbers of demands, because excess MTO demands are lost. Hence, these probabilities deviate from (7)–(10) by summing the MTO demand probabilities accordingly, and are thus given as

$$\pi^{ou}(s, s') = \left(\sum_{j=K-k}^{d_o^{\max}} p_o(j) \right) p_s(i - i') \quad \text{for } (s, s') \in \mathcal{M}_{ou} \cap \mathcal{O}_I,$$

$$\pi^{op}(s, s') = \left(\sum_{j=K-k+1}^{d_o^{\max}} p_o(j) \right) p_s(i - i') \quad \text{for } (s, s') \in \mathcal{M}_{op} \cap \mathcal{O}_{II},$$

$$\pi^{su}(s, s') = \left(\sum_{j=K-k}^{d_o^{\max}} p_o(j) \right) p_s(i - i') \quad \text{for } (s, s') \in \mathcal{M}_{su} \cap \mathcal{O}_I,$$

$$\pi^{sp}(s, s') = \left(\sum_{j=K-k}^{d_o^{\max}} p_o(j) \right) p_s(i - i' + 1) \quad \text{for } (s, s') \in \mathcal{M}_{sp} \cap \mathcal{O}_I.$$

Similarly, the third set applies to transitions to states with $i' = 0$ and $k' < K$, which are reached not only with i MTS demands (with $i + 1$ for action *sp*) but also with more demands, as excess MTS demands are lost. This leads to

$$\pi^{ou}(s, s') = p_o(k'_0) \left(\sum_{j=i}^{d_s^{\max}} p_s(j) \right) \quad \text{for } (s, s') \in \mathcal{M}_{ou} \cap \mathcal{O}_I,$$

$$\pi^{op}(s, s') = p_o(k'_0) \left(\sum_{j=i}^{d_s^{\max}} p_s(j) \right) \quad \text{for } (s, s') \in \mathcal{M}_{op} \cap \mathcal{O}_{II},$$

$$\pi^{su}(s, s') = p_o(k'_0) \left(\sum_{j=i}^{d_s^{\max}} p_s(j) \right) \quad \text{for } (s, s') \in \mathcal{M}_{su} \cap \mathcal{O}_I,$$

$$\pi^{sp}(s, s') = p_o(k'_0) \left(\sum_{j=i+1}^{d_s^{\max}} p_s(j) \right) \quad \text{for } (s, s') \in \mathcal{M}_{sp} \cap \mathcal{O}_I.$$

The fourth set combines the adjustments made in the previous two sets and gives the transition probabilities to states with $k' = K$ and $i' = 0$. These yield

$$\pi^{ou}(s, s') = \left(\sum_{j=K-k}^{d_o^{\max}} p_o(j) \right) \left(\sum_{j=i}^{d_s^{\max}} p_s(j) \right) \quad \text{for } (s, s') \in \mathcal{M}_{ou} \cap \mathcal{O}_I,$$

$$\pi^{op}(s, s') = \left(\sum_{j=K-k+1}^{d_o^{\max}} p_o(j) \right) \left(\sum_{j=i}^{d_s^{\max}} p_s(j) \right) \quad \text{for } (s, s') \in \mathcal{M}_{op} \cap \mathcal{O}_{II},$$

$$\pi^{su}(s, s') = \left(\sum_{j=K-k}^{d_o^{\max}} p_o(j) \right) \left(\sum_{j=i}^{d_s^{\max}} p_s(j) \right) \quad \text{for } (s, s') \in \mathcal{M}_{su} \cap \mathcal{O}_I,$$

$$\pi^{sp}(s, s') = \left(\sum_{j=K-k}^{d_o^{\max}} p_o(j) \right) \left(\sum_{j=i+1}^{d_s^{\max}} p_s(j) \right) \quad \text{for } (s, s') \in \mathcal{M}_{sp} \cap \mathcal{O}_I.$$

In all other cases, $\pi^a(s, s') = 0$.

Appendix B. Full factorial design of experiments

Besides varying each of the parameter values separately (low, base, high) as reported in the main text (see also Table 3), we also conduct a full factorial design of experiments. To limit the total computational time to about a week in total, we only consider two values (low, base) for each parameter in that experiment, giving a total of $2^8 = 256$ combinations. There are two main reasons for performing this full factorial design. First, it allows us to confirm (or deny) effects observed in the main text for the individual variation of each parameter. Second, we can check for interaction effects between the different parameters. Recall from the main text that we determine the cost savings resulting from full flexibility compared to both partial and no flexibility. It turns out that regression analyses on the full factorial results leads to similar results for both ways of defining savings, and so only the results for full compared to no flexibility are presented.

The first order effects of a regression on all eight parameters on the savings are reported in Table B.5. These confirm the results reported in the main text from the base case with individual parameter variations. The demand rates d_o and d_s , the lead time L , the MTS holding costs h and the MTO lateness costs q are the most influential drivers for the reported savings. Indeed, as the P -values show, these are the only significant effects at the significance level of 0.001. Moreover, the directions of these effects are confirmed by the signs of the coefficients.

To check for possible interaction effects, we perform another regression where we include, besides the 5 significant direct effects, all their $(5 \times 4)/2 = 10$ interaction effects. The results are presented in Table B.6. It appears that only the four interaction effects of MTO demand d_o with the other parameters are significant. These interaction effects are presented graphically in Fig. B.1. Note that for all four factors (MTS demand d_s , the holding cost h , the

Table B1
Full factorial design: regression on all parameters.

	Coefficient	Standard error	P-value
Intercept	0.327234	0.009386	0.000000
d_o	-1.051250	0.018386	0.000000
d_s	-0.260625	0.018386	0.000000
L	0.014094	0.000919	0.000000
K	0.000000	0.000460	1.000000
h	0.025500	0.001839	0.000000
q	-0.004461	0.000230	0.000000
b_o	-0.000003	0.000007	0.658935
b_s	0.000023	0.000007	0.001578

Fig. B1. Significant interaction effects.

Table B2

Full factorial design: regression including interaction effects on parameters with a significant direct effect.

	Coefficient	Standard error	P-value
Intercept	0.180469	0.076054	0.018438
d_o	-0.521875	0.248459	0.036733
d_s	-0.990625	0.248459	0.000089
L	0.062656	0.010190	0.000000
h	0.097562	0.027974	0.000580
q	-0.008367	0.003497	0.017488
$d_o \times d_s$	4.400000	0.605951	0.000000
$d_o \times L$	-0.215000	0.030298	0.000000
$d_o \times h$	-0.297500	0.060595	0.000002
$d_o \times q$	0.016875	0.007574	0.026815
$d_s \times L$	-0.021250	0.030298	0.483748
$d_s \times h$	-0.100000	0.060595	0.100190
$d_s \times q$	-0.007813	0.007574	0.303375
$L \times h$	0.003250	0.003030	0.284486
$L \times q$	0.000359	0.000379	0.343615
$h \times q$	-0.000625	0.000757	0.410107

lead time L , and the MTO lost sales cost q), the slope of the line is indeed slightly moderated by the value of MTO demand d_o , but the direction of the line is unchanged and in accordance with the discussions in the main text.

References

- Beemsterboer, B., Land, M., & Teunter, R. (2016). Hybrid mto-mts production planning: An explorative study. *European Journal of Operational Research*, 248, 453–461.
- Carr, S., & Duenyas, I. (2000). Optimal admission control and sequencing in a make-to-stock/make-to-order production system. *Operations research*, 48, 709–720.
- Chang, S.-H., Pai, P.-F., Yuan, K.-J., Wang, B.-C., & Li, R.-K. (2003). Heuristic pac model for hybrid mto and mts production environment. *International Journal of Production Economics*, 85, 347–358.
- Elmaghraby, S. E. (1978). The economic lot scheduling problem (elsp): Review and extensions. *Management Science*, 24, 587–598.
- Federgruen, A., & Katalan, Z. (1999). The impact of adding a make-to-order item to a make-to-stock production system. *Management Science*, 45, 980–994.
- Fransoo, J. C., Sridharan, V., & Bertrand, J. W. M. (1995). A hierarchical approach for capacity coordination in multiple products single-machine production systems with stationary stochastic demands. *European Journal of Operational Research*, 86, 57–72.
- Hadj Youssef, K., Van Delft, C., & Dallery, Y. (2004). Efficient scheduling rules in a combined make-to-stock and make-to-order manufacturing system. *Annals of Operations Research*, 126, 103–134.
- Iravani, S. M., Liu, T., & Simchi-Levi, D. (2012). Optimal production and admission policies in make-to-stock/make-to-order manufacturing systems. *Production and Operations Management*, 21, 224–235.
- Köber, J., & Heinecke, G. (2012). Hybrid production strategy between make-to-order and make-to-stock a case study at a manufacturer of agricultural machinery with volatile and seasonal demand. In *Proceedings of the 45th CIRP conference on manufacturing systems 2012* (pp. 453–458).
- Leachman, R. C., & Gascon, A. (1988). A heuristic scheduling policy for multi-item, single-machine production systems with time-varying, stochastic demands. *Management Science*, 34, 377–390.
- Rajagopalan, S. (2002). Make to order or make to stock: model and application. *Management Science*, 48, 241–256.
- Romsdal, A., Arica, E., Strandhagen, J. O., & Dreyer, H. C. (2013). Tactical and operational issues in a hybrid mto-mts production environment: The case of food production. In C. Emmanouilidis, M. Taisch, & D. Kiritsis (Eds.), *Advances in production management systems: Competitive manufacturing for innovative products and services* (pp. 614–621). Heidelberg: Springer.
- Soman, C. A., van Donk, D. P., & Gaalman, G. (2006). Comparison of dynamic scheduling policies for hybrid make-to-order and make-to-stock production systems with stochastic demand. *International Journal of Production Economics*, 104, 441–453.
- Soman, C. A., Van Donk, D. P., & Gaalman, G. J. (2007). Capacitated planning and scheduling for combined make-to-order and make-to-stock production in the food industry: An illustrative case study. *International Journal of Production Economics*, 108, 191–199.
- Sox, C. R., Jackson, P. L., Bowman, A., & Muckstadt, J. A. (1999). A review of the stochastic lot scheduling problem. *International Journal of Production Economics*, 62, 181–200.

- Sox, C. R., Thomas, L. J., & McClain, J. O. (1997). Coordinating production and inventory to improve service. *Management Science*, 43, 1189–1197.
- Vergin, R. C., & Lee, T. N. (1978). Scheduling rules for multiple product single machine system with stochastic demand. *Infor*, 16, 64–73.
- Williams, T. (1984). Special products and uncertainty in production/inventory systems. *European Journal of Operational Research*, 15, 46–54.
- Winands, E. M., Adan, I. J., & Van Houtum, G. (2011). The stochastic economic lot scheduling problem: A survey. *European Journal of Operational Research*, 210, 1–9.
- Wu, M.-C., Jiang, J.-H., & Chang, W.-J. (2008). Scheduling a hybrid mto/mts semiconductor fab with machine-dedication features. *International Journal of Production Economics*, 112, 416–426.