

# On the Learning Benefits of Resource Flexibility

Jiri Chod\*

Mihalis G. Markakis†

Nikolaos Trichakis‡

August 4, 2020

## Abstract

Resource flexibility, arguably among the most celebrated operational concepts, is known to provide firms facing demand uncertainty with such benefits as risk pooling, revenue-maximization optionality, and operational hedging. In this article we uncover a heretofore unknown benefit: we establish that resource flexibility facilitates learning the demand, when the latter is censored, which could, in turn, enable firms to make better-informed future operational decisions, thereby increasing profitability. Further, we quantify these learning benefits of flexibility and find that they could be of the same order of magnitude as the extensively studied risk pooling benefits of flexibility. This suggests that flexibility’s learning benefits could be a first-order consideration and that extant theories, which view flexibility only as the ability to act ex post, could be underestimating its true value when learning the demand is desirable, for example, when it enables managers to make better ex ante capacity, assortment, or pricing decisions in future periods.

Keywords: Flexibility, Learning, Risk Pooling, Censoring.

## 1. Introduction

Should a firm invest in resource flexibility or not? Analyzing the calculus associated with this dilemma has spanned almost four decades of research going back to the work of Eppen [1979] on stock centralization. On the one hand, flexibility tends to increase operating costs due to increased complexity and lack of specialization, among other reasons. On the other hand, flexibility also carries well-known and extensively studied benefits such as the reduction of mismatch between supply and demand through risk pooling [Jordan and Graves, 1995], the option to allocate resources to the most profitable use [Van Mieghem, 1998], and the ability to serve as an operational hedge

---

\*Carroll School of Management, Boston College; [chodj@bc.edu](mailto:chodj@bc.edu)

†IESE Business School, University of Navarra; [MMarkakis@iese.edu](mailto:MMarkakis@iese.edu)

‡MIT Operations Research Center and Sloan School of Management; [ntrichakis@mit.edu](mailto:ntrichakis@mit.edu)

[Van Mieghem, 2007]. It is because of all these benefits that flexibility is often defined as the “ability to react” to demand uncertainty realization [Golden and Powell, 2000].

Herein, we theorize that resource flexibility carries an important benefit that has been overlooked in the academic literature so far. Namely, we argue that resource flexibility can help firms to learn about the underlying demand when the latter is not fully observable due to censoring. Learning could, in turn, enable firms to make better-informed operational decisions, such as capacity choices, in subsequent periods, thereby increasing their profitability. Put differently, besides the ability to react *ex post* uncertainty realization, this paper propounds that flexibility may allow firms to also make better *ex ante* decisions in future periods. This further suggests that firms contemplating investing in flexibility could be underestimating its true value when accounting merely for its “reactive” abilities.

To study the learning benefits of flexibility, we consider a two-period model in which a firm faces two uncertain demand streams in each of the two periods. Demand is fulfilled using one of two types of resources: either non-flexible, so that each demand requires a dedicated resource, or flexible, so that each demand can be fulfilled using a single flexible resource. Resource capacities used in each period have to be chosen at the beginning of that period. To guide capacity choices, the firm has a prior distribution on demand. The firm can also “learn” and update its prior via realized demand samples that are subject to censoring if demand exceeds available capacity. We first consider a benchmark case in which demands across periods are independent, so that learning is inconsequential. Then, we consider a case in which demands are correlated across time, so that there is potential value to learning that could increase profits. Within this model, we seek to address the following questions: Do flexible resources lead to a higher profit increase when learning the realized demand is valuable, i.e., do flexible resources carry a learning benefit over non-flexible ones? If so, how large are the learning benefits of flexibility? Finally, what are the drivers of the learning benefits of flexibility and what do they depend on?

We derive two principal results that can be summarized as follows. First, we establish that, indeed, flexibility usually carries learning benefits. In particular, we show that when learning becomes valuable, profitability of the flexible system increases more than that of the non-flexible system, as long as profit margins are not too high. Second, we quantify the learning benefits of flexibility and find them to be *comparable* to flexibility’s risk pooling benefits in terms of profitability; see Table 1 for indicative results. Given the well-recognized significance of the latter, we conclude that flexibility’s learning benefits could be a first-order consideration for a firm.

Furthermore, our analysis sheds some light on what drives the learning benefits of flexibility.

Profit Margin	Learning/Pooling
5%	2.4×
10%	1.5×
20%	0.9×
40%	0.2×

**Table 1:** The ratio of flexibility’s learning benefits over its risk pooling benefits for different profit margins. The reported values correspond to a scenario in which the unit cost of flexible capacity is equal to that of dedicated capacity; as flexibility becomes more expensive, learning benefits become even more significant relative to pooling benefits. See Section 4.1 for a detailed discussion.

To this end, we first show that when demand is serially correlated, the flexible system usually affords lower variance for the second-period demand posterior. To understand what drives the improvement in the second-period forecasting accuracy, we then examine how flexibility affects censoring of the first-period demand. In a parsimonious one-period version of our model, we show that the flexible system tends to result in less censoring than the non-flexible system, providing an intuitive explanation of what ultimately drives flexibility’s learning benefits in the two-period model. But, going one step further, what makes the flexible system effective at mitigating censoring? As might be expected, when demand is serially correlated and averting censoring is valuable, the latter can be achieved by expanding first-period capacities, in both flexible and non-flexible systems. Interestingly, we show that the relative capacity increase is usually higher for the non-flexible system compared to the flexible one. The above results hint that the learning benefits of flexibility are achieved through effective censoring mitigation owing to a more efficient use of resources, rather than to excess capacity.

A common thread among our results is that the learning benefits of flexibility emerge only in settings in which profit margins are moderate or low. It should be noted that settings in which profit margins are high do not represent fertile ground to draw meaningful conclusions about learning benefits. This is because high margins justify ample capacity for both the flexible and the non-flexible systems, regardless of whether learning carries any value. Therefore, in this regime demand censoring becomes negligible and the learning benefits become minuscule for both systems.

Profit margin aside, we identify two more key factors affecting the learning benefits of flexibility: demand variability and serial correlation. As the coefficient of variation of first-period demand decreases, i.e., as there is less uncertainty regarding the first-period demand, the learning benefits of flexibility decrease as well, and preference for the flexible system wanes overall. Similarly, as demand correlation across periods weakens, i.e., as there is less information to be extracted from realized first-period demand, the learning benefits of flexibility decrease, and the flexible system

becomes less attractive overall.

Our theory is relevant in settings in which demand information is censored and, in particular, backorders are not allowed. The latter is often, but not always, the case in retail, as well as in certain service industries. The most salient example of demand censoring is the case of brick-and-mortar retail, whereby a customer who does not see their preferred product on the shelf, walks away or buys a substitute. Retail firms that are able to reduce lost sales by deploying flexibility strategies, such as delayed differentiation, can use the supplementary demand information to make better inventory, pricing, or assortment decisions thereafter. This can be particularly useful for fast-fashion retailers, like ZARA and H&M, which are alleged to learn fashion trends faster than their competitors by leveraging the frequent reordering and redesigning capabilities that their flexible operating systems offer. Within service industries, Liu et al. [2002] describe the demand censoring effect in the hotel and airline industries: “Even when reservation systems are designed to record data on lost opportunities, the data recorded are insufficient for inferring unconstrained demand.” Accordingly, airlines, hotels, or car rental companies, for example, that use flexible service capacity/upgrades [Netessine et al., 2002b] could profit from learning benefits. On the contrary, our theory would be less relevant in settings in which backorders are allowed, as is often the case in B2B settings, e.g., a car manufacturer that accepts backorders from dealerships.

As firms strive to redesign their supply chains from cost-efficient to responsive, resource flexibility is recognized as a key enabler of responding to fast-changing market conditions. In summary, our findings imply that resource flexibility can be also an important enabler of learning the changes in market conditions more efficiently, which could provide a valuable competitive advantage. That is, when learning matters, the value of flexibility could be higher than suggested by the existing literature, and so would be the cost threshold for which a firm should adopt a flexible system.

The remainder of the paper is organized as follows. After briefly reviewing the related literature on flexibility and learning, we present the models considered in this work in Section 2. Sections 3 and 4 are focused on establishing and quantifying the learning benefits of flexibility. Section 5 deals with shedding some light on drivers and factors of flexibility’s learning benefits. Section 6 concludes with a brief summary and directions for future research.

## **Related Literature**

Our paper contributes primarily to the flexibility literature. Various cost drivers of flexibility have already been documented [Van Mieghem and Allon, 2008], however, there is also recent research on the topic; either identifying new drivers, such as increased financing costs [Iancu et al., 2016],

or empirically documenting and quantifying their effects on productivity [Choudhary et al., 2018]. Focusing on benefits, the literature has thus far identified at least three distinct benefits of resource flexibility. The first one is the reduction of the mismatch between supply and demand through risk-pooling, and it can be captured through many different strategies including stock centralization [Eppen, 1979, Schwarz, 1989], investing in flexible production capacity [Chod and Rudi, 2005, Fine and Freund, 1990, Goyal and Netessine, 2011, Jordan and Graves, 1995, Netessine et al., 2002a, Van Mieghem, 1998], or postponement of product differentiation [Lee and Tang, 1997, Swaminathan and Tayur, 1998]. Second, resource flexibility often provides the real option to allocate the flexible resource to the most profitable use, e.g., the high-margin product [Van Mieghem, 1998], or in a global setting, the market with a favorable exchange rate [Ding et al., 2007, Huchzermeier and Cohen, 1996]. Finally, resource flexibility can be sometimes used as an operational hedge to mitigate profit variability [Chod et al., 2010, Ding et al., 2007, Van Mieghem, 2007]. We are not aware of any papers that would study the effect of resource flexibility on learning. Our contribution to the flexibility literature is to identify this effect, and to suggest that this literature undervalues true benefits of resource flexibility when learning is taken into consideration.

Our work is also novel within the literature on learning with censored demand data. There are two distinct, and complementary, strands in this literature: (i) the parametric approach, applied to both perishable [Besbes et al., 2017, Ding et al., 2002, Harpaz et al., 1982, Jain et al., 2015, Lariviere and Porteus, 1999, Lu et al., 2008, Mersereau, 2015] and nonperishable products [Bensoussan and Guo, 2015, Bisi et al., 2011, Chen, 2010, Chen and Plambeck, 2008]; (ii) the nonparametric approach, either establishing the asymptotic convergence to the optimal decision [Burnetas and Smith, 2000, Godfrey and Powell, 2001, Huh et al., 2011, Kunnumkal and Topaloglu, 2008, Maglaras and Eren, 2015, Powell et al., 2004, van Ryzin and McGill, 2000], or minimizing the regret [Besbes and Muharremoglu, 2013, Huh and Rusmevichientong, 2009, Lugosi et al., 2017, Shi et al., 2016]. A common feature of the above works is that the production/procurement capabilities of the firm are exogenous, so no connection is drawn between a firm’s ability to learn through censored data and its flexibility, or the lack thereof.

## 2. Model

To address our research questions, we focus on a two-product capacity decision problem similar to the ones studied in the literature on flexibility [Fine and Freund, 1990], extended over two time periods, both of which involve capacity decisions. Demand is serially correlated across these periods

and this introduces learning considerations, because demand observations in the first period carry potential benefits in the second. The end result is a dynamic problem with two interdependent capacity decision problems.

Apart from the above dynamic model though, we need a benchmark to compare our findings against. For that purpose, we consider the same model but in the absence of learning considerations; in particular, assuming that demands are uncorrelated across the two periods. In that case, the capacity decisions made by the firm at the beginning of the first period, have no bearing on what transpires in the second period. In that sense, they become identical to capacity decisions made within an equivalent static problem in which the firm's objective is to maximize its expected profits over a single period.

The analysis of the equivalent static model will also guide our efforts to focus on and distill the learning benefits of flexibility. In particular, in the static, one-period model, the advantage of flexible over non-flexible resources is solely attributed to their pooling benefits. If the per unit cost of flexible capacity is low enough, pooling benefits would outweigh capacity investment costs and flexibility would be preferable, but if it is too high, the latter would no longer be the case. There is then a break-even unit cost of flexible capacity, for which the pooling benefits are cancelled by capacity costs and the firm becomes indifferent between the flexible and the non-flexible systems. We use this break-even cost to distill the learning benefits of flexibility, because for this particular cost, any difference in expected profit, in the context of the dynamic model, would be solely attributed to learning.

## 2.1. Static Model

We consider a profit-maximizing firm that markets two products or services, possibly two versions of the same product customized for two different markets. We henceforth refer to them simply as products, and index them by  $i \in \{1, 2\}$ . Demand for product  $i$  is a continuous random variable, denoted by  $D_i$ , and  $F_i$  is its marginal c.d.f. These distributions can be thought of as part of a prior that encapsulates the firm's information/beliefs regarding future demand. The sale of product  $i$  generates unit revenue  $p_i$ , and  $p_1 \geq p_2$  without loss of generality.

Sales are constrained by the capacity of a critical resource, such as finished product inventory, WIP, machinery, or workforce, that has to be determined before demand is known. We consider two production systems.

**Non-flexible system.** In the non-flexible system, each product requires a product-specific resource, whose capacity, denoted as  $Q_i$  for product  $i$ , is built at a constant unit cost  $c$ . The profit margin

of product  $i$  is then  $p_i - c$ . Assuming equal unit capacity cost for both products simplifies the exposition without any loss of generality, as one can simply redefine one unit of either capacity.

The sales of product  $i$  are

$$S_{N,i} = \min \{D_i, Q_i\}, \quad (1)$$

and the firm chooses capacities of the two resources,  $\mathbf{Q}$ , so as to maximize expected profit

$$\pi_N = \mathbf{p}'\mathbb{E}[\mathbf{S}_N] - \mathbf{c}'\mathbf{Q}, \quad (2)$$

where boldface is used to denote vectors, whose elements are the respective quantities for the two products. Note that the expectation is taken with respect to the joint prior distribution of the demands for the two products—in essence, with respect to the beliefs of the firm regarding future demand—in line with the standard approach of stochastic optimization in the case of imperfect state information [Bertsekas, 2005]. Let

$$\bar{\mathbf{Q}}_N = \arg \max_{\mathbf{Q}} \pi_N \quad (3)$$

be the vector of optimal non-flexible capacity levels, which are standard newsvendor solutions; see Lemma 1 in Appendix A.

**Flexible system.** In the flexible system, the firm can satisfy both demands using a single flexible resource, which could be inventory of undifferentiated WIP, flexible machinery, or cross-trained workforce. The capacity of the flexible resource, denoted as  $Q_F$ , is built at a constant unit cost  $c_F$ .

We assume that demand is realized so as to enable revenue-maximizing fulfillment, as is standard in the flexibility literature [Fine and Freund, 1990, Jordan and Graves, 1995, Van Mieghem, 1998]. In our context, this means that demand for the more profitable product 1 is realized first. This could be the result of the firm's price skimming strategy: firms often introduce their products to high-valued markets first, delaying introduction to low-valued markets [Moorthy and Png, 1992]. Cellphones, for example, which require antenna customization depending on location, are often first introduced to the U.S. and certain European markets at high prices, before being made available more widely at lower prices. Any similar price-skimming strategies that also involve some product customization, such as watering down product features together with lowering price, would fit well within our model.

Under this assumption, it is optimal to make the entire flexible capacity available to satisfying

product 1's demand.<sup>1</sup> The sales of the two products are thus

$$S_{F,1} = \min \{D_1, Q_F\} \text{ and } S_{F,2} = \min \{D_2, Q_F - S_{F,1}\}, \quad (4)$$

and the expected profit is equal to

$$\pi_F(c_F) = \mathbf{p}'\mathbb{E}[\mathbf{S}_F] - c_F Q_F. \quad (5)$$

We note that for the flexible system, we usually highlight the dependence of various quantities on the unit cost of flexibility  $c_F$  explicitly, as, for example, in the equation above. This is because  $c_F$  will emerge as an important parameter in our analysis.

We denote the profit-maximizing flexible capacity level by

$$\bar{Q}_F(c_F) = \arg \max_{Q_F} \pi_F(c_F), \quad (6)$$

which satisfies a newsvendor-type optimality condition; see Lemma 2 in Appendix A.

**Risk pooling benefits of flexibility.** In this static model, any benefits that the flexible system may have over the non-flexible one have to be attributed to risk pooling, since no learning takes place. We define the *risk pooling benefits of flexibility* as the profitability difference between the flexible and the non-flexible systems within this static model:

$$\mathcal{P}(c_F) \equiv \max_{Q_F} \pi_F(c_F) - \max_{\mathbf{Q}} \pi_N. \quad (7)$$

As a side note, we use the term risk pooling benefits in a somewhat wider sense compared to the literature on flexibility, including also benefits from revenue-maximization optionality.

**Break-even unit cost of flexibility.** If the unit cost of flexible capacity  $c_F$  is relatively low compared to that of non-flexible  $c$ , for example if  $c_F = c$ , the profitability of the flexible system is higher, and then the risk pooling benefits of flexibility are positive. If  $c_F$  becomes too high, however, for example if  $c_F = 2c$ , then the profitability of the non-flexible system is higher, and then the risk pooling benefits of flexibility are negative. It can also be readily checked that the pooling benefits  $\mathcal{P}$  are monotonically decreasing in  $c_F$ . Given all this, there exists then a *break-even unit cost of*

---

<sup>1</sup>We revisit the demand fulfillment process in Appendix B.3.

flexibility  $\tilde{c}_F$  for which the pooling benefits vanish. Formally,  $\tilde{c}_F$  is defined as

$$\tilde{c}_F : \max_{Q_F} \pi_F(\tilde{c}_F) = \max_{\mathbf{Q}} \pi_N. \quad (8)$$

That is, for  $c_F = \tilde{c}_F$  the firm is indifferent between the two systems. For  $c_F < \tilde{c}_F$  (resp.  $c_F > \tilde{c}_F$ ) the flexible system (resp. non-flexible) becomes preferable.

## 2.2. Dynamic Model

We build the dynamic model that constitutes the focus of our analysis, on the static model introduced in Section 2.1. For brevity, we present only the differences between the two models.

We consider a two-period model, and we use superscript  $t \in \{1, 2\}$  to denote the time period. For example,  $D_i^t$  stands for the demand for product  $i$  in period  $t$ .

At the beginning of period  $t$ , the firm decides how much capacity to build in order to satisfy demand  $\mathbf{D}^t = (D_1^t, D_2^t)'$ , which is later realized and fulfilled in the same way as in Section 2.1. Unused first-period capacity becomes obsolete and cannot be used in the second period. As before, we consider two operating systems, which differ only in the type of capacity used to satisfy the first-period demand: non-flexible vs. flexible. In particular, in both systems, the firm uses non-flexible capacities to satisfy the second-period demand—a modeling choice that, as we explain later in more detail, allows us to disentangle the learning benefits of flexibility from its pooling benefits.

We assume that demand is serially correlated, i.e., the demand in the second period is equal to the first-period demand plus some noise:

$$D_i^2 = D_i^1 + \epsilon_i, \quad i \in \{1, 2\}. \quad (9)$$

Serially-correlated demand models enjoy excellent support in the academic literature, having been employed by foundational works, from Economics to Operations, and are supported by ample empirical evidence; see [Fisher and Raman, 2010, Kahn, 1987, Lee et al., 2004].

With serially-correlated demands, the accuracy of the second-period demand forecast depends on the first-period capacity decision, as well as the realization of the first-period demand. To see this, suppose that the firm has  $Q$  capacity units available to satisfy the first-period demand for the first product,  $D_1^1$ , which turns out to be  $d$ . If  $Q > d$ , then the firm observes that  $D_1^1 = d$ , and updates its forecast for second-period demand for the first product to be  $d + \epsilon_1$ . If  $Q \leq d$ , then the firm stocks out, and updates its forecast for second-period demand for the first product to be  $(D_1^1 \mid D_1^1 \geq Q) + \epsilon_1$ . Consequently, there are potential learning benefits associated with building

first-period capacity, be it non-flexible or flexible. What we wish to quantify are the added learning benefits that flexible resources carry vis-à-vis dedicated resources and, in particular, to compare them with the classical risk-pooling benefits of flexibility.

All decisions are made so as to maximize expected profits. The decision problem at hand is a Dynamic Program with imperfect state information. Let  $\mathbb{P}$  be the measure of the probability space. The state of the Dynamic Program corresponds to the most up-to-date belief about the demand distributions. At  $t = 1$ , the belief is simply given by  $\mathbb{P}$ . At  $t = 2$ , however, the belief is updated in a Bayesian fashion after observing the first-period sales  $\mathbf{S}^1$ , i.e., the belief is given by the associated conditional probability measure  $\mathbb{P}(\cdot | \mathbf{S}^1)$ .

Let  $V_X^t(\cdot)$  be the value function, or optimal expected profits-to-go, at the beginning of period  $t$ , where  $X = N$  for the non-flexible system and  $X = F$  for the flexible one.

For the non-flexible system, the Bellman equation at time  $t$ , for some up-to-date belief  $\mathbb{G}$ , can be written as

$$V_N^t(\mathbb{G}) = \max_{\mathbf{Q}^t \in \mathbb{R}_+^2} \mathbb{E}_{\mathbb{G}} \left[ \mathbf{p}' \mathbf{S}_N^t - \mathbf{c}' \mathbf{Q}^t + V_N^{t+1} \left( \mathbb{G} \left( \cdot | \mathbf{S}_N^t \right) \right) \right], \quad t = 1, 2, \quad (10)$$

with boundary condition  $V_N^3(\cdot) = 0$ .

For the flexible system, recall that flexible resources are used only at  $t = 1$ , and therefore we have

$$V_F^2(\mathbb{G}) = \max_{\mathbf{Q}^2 \in \mathbb{R}_+^2} \mathbb{E}_{\mathbb{G}} \left[ \mathbf{p}' \mathbf{S}_N^2 - \mathbf{c}' \mathbf{Q}^2 \right], \quad (11)$$

and

$$V_F^1(\mathbb{G}) = \max_{Q_F^1 \in \mathbb{R}_+} \mathbb{E}_{\mathbb{G}} \left[ \mathbf{p}' \mathbf{S}_F^1 - c_F Q_F^1 + V_F^2 \left( \mathbb{G} \left( \cdot | \mathbf{S}_F^1 \right) \right) \right]. \quad (12)$$

For notational convenience, we simply write the firm's optimal expected profits for this two-period model as  $V_N$  and  $V_F(c_F)$  for the two cases, i.e.,

$$V_N \equiv V_N^1(\mathbb{P}), \quad V_F(c_F) \equiv V_F^1(\mathbb{P}).$$

Let  $\bar{\mathbf{Q}}_N^1$  and  $\bar{Q}_F^1$  also denote the first-period optimal capacity choices for the non-flexible and the flexible systems, respectively.

Before we proceed with some useful definitions, we note that our model could reasonably approximate capacity decision problems facing firms in consumer electronics or fashion apparel industries. Specifically, these firms tend to release their products periodically, for example, on an annual basis. To launch their products in each such period, firms usually have to make stocking decisions before

demand is realized. Moreover, products in each period tend to be somewhat similar to products in the previous period, e.g., newer versions of cellphones are often similar to previous ones, albeit with some extra features. Consequently, there is some serial correlation of demands across periods, and therefore, observing demand in one period could better inform capacity decisions in future periods.

**Learning benefits of flexibility.** In the first period, the use of flexible capacity offers solely risk pooling benefits, since no learning has taken place yet. In the second period, however, dedicated capacities are used in both cases. So, we can only have learning benefits, which may arise due to a more accurate demand forecast, resultant from using flexible resources in the first period. Therefore, the incremental profit from using flexible capacity,  $V_F(c_F) - V_N$ , is attributed to both risk pooling and learning benefits. Consequently, we can measure the learning benefits of flexibility, denoted by  $\mathcal{L}(c_F)$ , as

$$\mathcal{L}(c_F) \equiv V_F(c_F) - V_N - \mathcal{P}(c_F), \quad (13)$$

where  $\mathcal{P}(c_F)$  is given by Eq. (7).

As a technical remark, note that the actual risk pooling benefits in the first period of this model would typically deviate from  $\mathcal{P}(c_F)$ , given that both systems adjust their capacity decisions in anticipation of learning. Notwithstanding this deviation, the fair way to measure the learning benefits is still as we write above, simply because any decrease in the actual risk pooling benefits in the first period represents a cost associated with learning, and should therefore be accounted for in the net learning benefits of flexibility.

**Independent demands across periods.** As discussed previously, we consider for benchmark purposes the case in which second-period demands are independent from first-period demands. In particular, we assume that  $D_i^2$  is the sum of a demand term and a noise term, similar to (9), but the demand term is now independent and identically distributed to first-period demand. In this case, the problem then decouples into two single-period problems. For system  $X \in \{N, F\}$ , we denote by  $V_X^\perp$  the associated optimal expected profits, and by  $\bar{Q}_N^\perp$  and  $\bar{Q}_F^\perp$  the associated optimal first-period capacity choice.

### 3. Establishing the Learning Benefits of Flexibility

We proceed to the analysis of the dynamic model of Section 2.2 and provide a mix of analytical results and numerical experiments that establish the learning benefits of flexibility. For this part of the analysis, we consider a “symmetric” case, in which products have equal margins, i.e.,  $(p - c)/p$ ,

and  $\{D_1^1, D_2^1\}$  are independent random variables, distributed identically to the uniform distribution in  $[x, x + U]$ ,  $i = 1, 2$ . We assume that the second-period noise terms are independent across products, and also uniformly distributed:  $\epsilon_i \sim \mathcal{U}[-\xi x, \xi x]$ , for some  $\xi \in [0, 1]$ . It should be noted that even for  $\xi = 0$ , the value functions involve, in general, non-convex optimization. Despite this technical challenge, we are still able to derive insightful analytical results.

We complement our analysis with numerical experiments, which we anchor at a base-case parameter setting throughout the paper.<sup>2</sup> To test the robustness of the derived insights, we complement the latter with further experiments, in which we unilaterally change different features of the model. In particular, we perform three robustness checks: (i) we change the stochastic variability of the first-period demand. This is captured by the *Coefficient of Variation (CoV)* of the prior distribution, which is controlled by parameter  $x$  in our model; (ii) we change the intensity of the second-period demand noise, which affects the information value that can be extracted from the first-period sales. Using Information Theory verbiage, this is captured by the *Signal-to-Noise Ratio (SNR)*, which is controlled by parameter  $\xi$  in our model; (iii) we change the demand distribution, considering Normal and Gamma.

More concretely, we compare the profits of the flexible system  $V_F(c_F)$  with that of the non-flexible system  $V_N$ . Any profitability advantage that the flexible system might have could be attributed to both pooling and learning benefits. Algebraically, this can be seen directly from (13) that yields  $V_F(c_F) - V_N = \mathcal{P}(c_F) + \mathcal{L}(c_F)$ .

To disentangle the effects of pooling and learning, we carry out the comparison for the break even cost of flexibility  $c_F = \tilde{c}_F$ , because under this cost the pooling benefits are nullified according to (7)-(8),  $\mathcal{P}(\tilde{c}_F) = 0$ , and learning benefits becomes the main ground of comparison, i.e.,

$$V_F(\tilde{c}_F) - V_N = \mathcal{L}(\tilde{c}_F). \quad (14)$$

In other words, the comparison of  $V_F(\tilde{c}_F)$  to  $V_N$  informs precisely which resource type becomes preferable exclusively due to learning considerations. Another way to see this is by noting that, when there is no demand correlation across time, the firm is indifferent between the two systems for  $c_F = \tilde{c}_F$ , i.e.,  $V_F^\perp(\tilde{c}_F) = V_N^\perp$ . By comparing  $V_F(\tilde{c}_F)$  to  $V_N$ , we study how learning affects preference for resource type. Therefore,  $V_F(\tilde{c}_F)$  being larger than  $V_N$  would establish the learning benefits of flexibility.

**Proposition 1.** *Suppose that  $\xi = 0$  and  $x = 0$ . Then, there exists a profit margin threshold, below*

---

<sup>2</sup>The parameters in the base-case setting take the values  $p = 1$ ,  $x = 50$ ,  $U = 100$ , and  $\xi = 0$ .

(resp. above) which  $V_F(\tilde{c}_F) \geq V_N$  (resp.  $V_F(\tilde{c}_F) \leq V_N$ ).

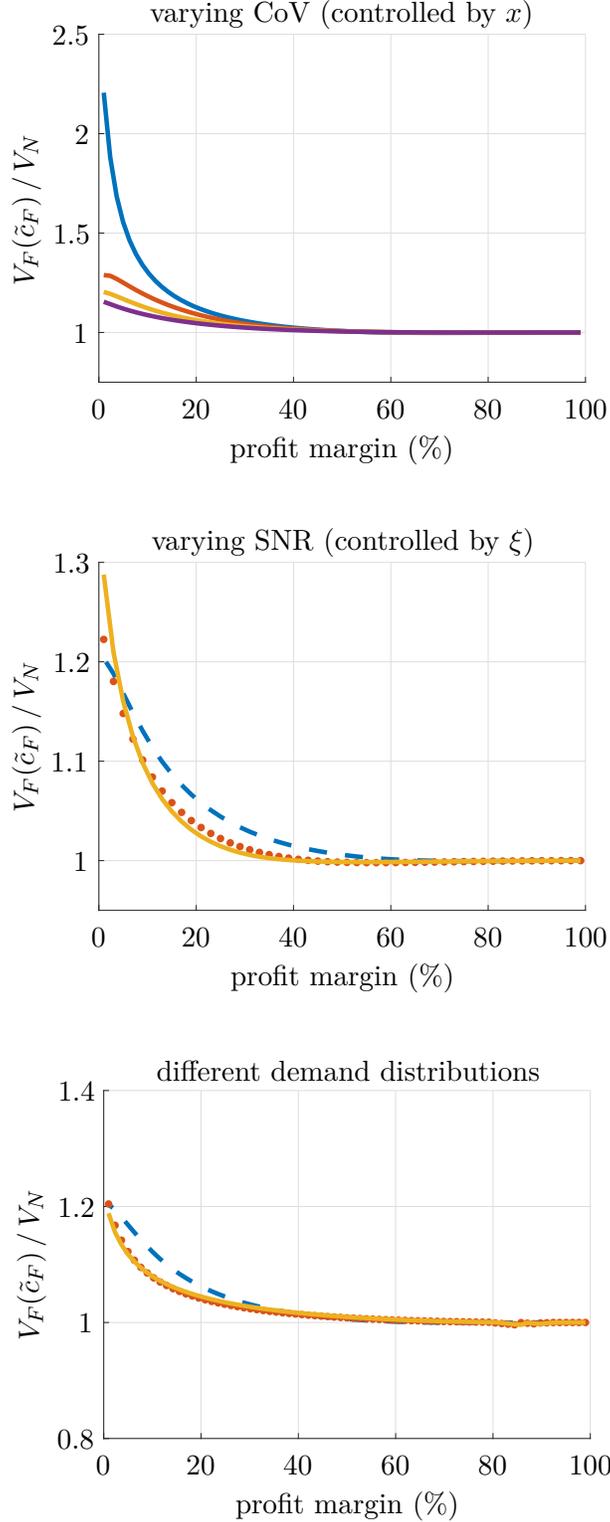
The result shows that if the demand is perfectly correlated in time and the margin is not high, the flexible system is more profitable than the non-flexible one. This implies that the relative increase in expected profit attributed to learning is higher in the flexible system compared to the non-flexible one. For  $\xi = 0$  and  $x = 0$ , we find the profit margin threshold  $((p - c)/p)$  to be approximately 59%. Thus, unless margins are higher than 59%, flexibility becomes more preferable in the presence of learning considerations. While the exact value of this threshold is sensitive to different parameters and modeling assumptions that we have made, the main insight is robust: the learning benefits of flexibility are established for moderate and low margins.

We provide further insight into the robustness of the above result through numerical experiments. In particular, in Figure 1 we study the ratio of the optimal expected profits of the flexible system over the non-flexible one,  $V_F(\tilde{c}_F)/V_N$ , as a function of the profit margin.

- In the top panel, we consider varying values of the first-period demand CoV (as controlled by parameter  $x$ ). A ratio larger (resp. smaller) than one reveals preference for flexible (resp. non-flexible) capacity. We note that all curves indeed exhibit a threshold-type structure as prescribed by Proposition 1. Moreover, for lower CoV (higher  $x$ ), the margin threshold below which flexibility carries learning benefits is decreasing, i.e., preference for flexibility due to learning wanes. The latter is because lower first-period demand CoV implies that there is less uncertainty to be resolved/learned in relative terms, and this reduces the learning benefits.
- In the center panel, we consider varying SNR values in the second period (as controlled by  $\xi$ ). As before, all curves exhibit a threshold-type structure. Also, for lower SNR (higher  $\xi$ ), the margin threshold decreases, because lower SNR again reduces the value of learning.
- In the bottom panel, we consider additional demand distributions, namely Normal and Gamma, the parameters of which are chosen so that they have the same mean and standard deviation as the Uniform distribution in our base case.

All findings are consistent with the main message of this section: what makes the flexible system attractive in this dynamic model is its learning capabilities; once they are limited, the attractiveness of the flexible system wanes accordingly.

Proposition 1 and the numerical analyses in Figure 1 also reveal the important dependence of the learning benefits of flexibility on profit margin. We observe that for low to moderate margins the profits of the flexible system become noticeably larger than the non-flexible one's. For high



**Figure 1:** The ratio of the optimal expected profits of the flexible system over the non-flexible one,  $V_F(\tilde{c}_F)/V_N$ , as a function of the profit margin for the base-case parameter setting. In the top panel, we consider varying CoV as controlled by  $x \rightarrow 0, 25, 50, 75$  (0 uppermost, 75 lowermost); in the center panel, we consider varying SNR as controlled by  $\xi \rightarrow 0, .5, 1$  (0 dashed, 1 solid); in the bottom panel, we consider different demand distributions, Uniform (dashed), Normal (dotted), and Gamma (solid).

margins, however, the profits of the two systems become almost indistinguishable. The intuition behind this is that high profit margins prompt ample capacity investments in the first period regardless of whether learning carries any value, to the extent that demand censoring becomes negligible and the learning benefits become minuscule. That is, for this regime we have  $\mathcal{L}(\tilde{c}_F) \approx 0$  and therefore  $V_F(\tilde{c}_F) \approx V_N$ . Therefore, high margins do not represent fertile ground to draw meaningful conclusions about learning benefits. The important takeaway is that for low to medium profit margins, learning considerations make flexible resources more preferable, whereas for high margins learning considerations become almost inconsequential.

## 4. Quantifying the Learning Benefits of Flexibility

In the next part of the analysis, we quantify the learning benefits of flexibility. We do so in two ways: first, we compare them against the risk pooling benefits of flexibility; second, we analyze uncaptured profits, i.e., how much is left on the table if the firm were to ignore the underlying demand correlation and overlook the learning benefits. We use the same setup as in Section 3.

### 4.1. Learning Benefits vs Pooling Benefits

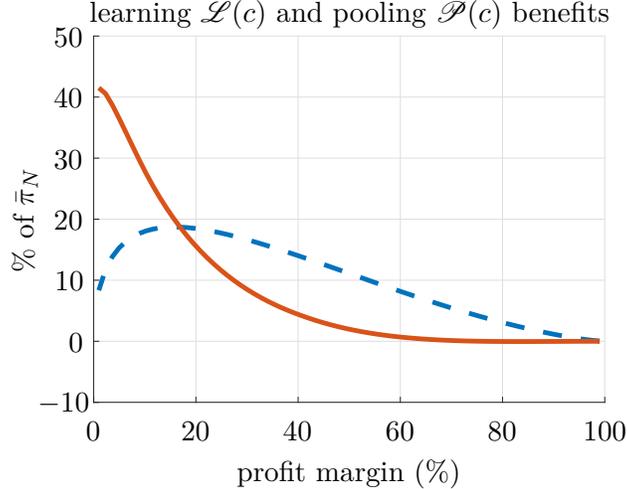
We compare the learning benefits of flexibility against its widely studied risk pooling benefits. If the former are of the same order of magnitude as the latter, albeit in our parsimonious model, this will imply that the learning benefits of flexibility could be a first-order consideration for a firm, given the profound impact that risk pooling concepts have had both in academia and in practice.

To conduct our comparison of learning vs pooling benefits, we focus on unit costs of flexible capacity  $c_F$  that are smaller than the break-even cost  $\tilde{c}_F$ , so that the pooling benefits are non-negative, i.e.,  $\mathcal{P}(c_F) \geq 0$ . Recall that for  $c_F > \tilde{c}_F$  the pooling benefits are negative, whereas for  $c_F = \tilde{c}_F$  the pooling benefits vanish, i.e.,  $\mathcal{P}(\tilde{c}_F) = 0$ .

We begin by setting the bar quite high: we ask whether the learning benefits could be in fact *larger* than flexibility’s pooling benefits. Against this backdrop, we find the following:

**Proposition 2.** *Suppose that  $\xi = 0$  and  $x = 0$ . Then, there exists a profit margin threshold below which  $\mathcal{L}(c_F) \geq \mathcal{P}(c_F)$ , for all  $c_F \leq \tilde{c}_F$ .*

Under the assumptions of Proposition 2, we find the threshold to be approximately 12%. In other words, if the profit margin is lower than  $\sim 12\%$ , flexibility’s learning benefits are indeed larger than its pooling benefits. Note that this threshold is “tight:” in the proof of the proposition, we construct an example in which the margin is slightly higher and  $\mathcal{L}(c_F) < \mathcal{P}(c_F)$ .



**Figure 2:** The normalized learning benefits of flexibility,  $\mathcal{L}(c)/\bar{\pi}_N$ , in solid line, and the normalized risk pooling benefits of flexibility,  $\mathcal{P}(c)/\bar{\pi}_N$ , in dashed line, for the base-case parameter setting.

It is worth highlighting that, given how celebrated the pooling benefits of flexibility have been in the field of Operations Management, requiring the learning benefits to be as large, or larger, is a rather tall order. By proving the very existence of a non-trivial parameter range over which this dominance holds, Proposition 2 provides compelling evidence for the significance of flexibility’s learning benefits over a broad range of relevant settings. Indeed, this generalization is corroborated by a series of numerical experiments that we conduct in order to capture additional settings, and which we discuss next.

In Figure 2 we plot the learning and risk pooling benefits of flexibility for  $c_F = c$  and for the base-case parameter setting, as functions of the profit margin. Both benefits are normalized by the corresponding optimal expected profits of the non-flexible system over one period, which we denote with  $\bar{\pi}_N \equiv \max_{\mathbf{Q}} \pi_N$ . We observe that for moderate margins, around 30%, or lower, the learning benefits of flexibility are of the same order of magnitude as the risk pooling benefits, and actually become higher as the profit margin decreases. In the case of high margins, between 40 – 60%, the risk pooling benefits of flexibility dominate. As we remark above, for high margins, the non-flexible system justifies enough capacity, and thus “learns quite well” already. Therefore, the added learning benefit of flexibility is marginal. Finally, for margins even higher than 60%, both learning and risk pooling benefits fade, because capacity becomes so ample as to render flexibility unnecessary altogether.

In Figure 3, we focus on the base case parameter setting and plot the ratio of the learning over the pooling benefits for  $c_F = c$  as a function of the profit margin:

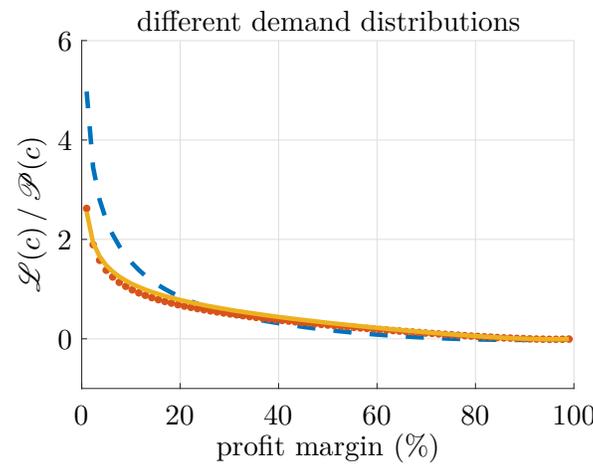
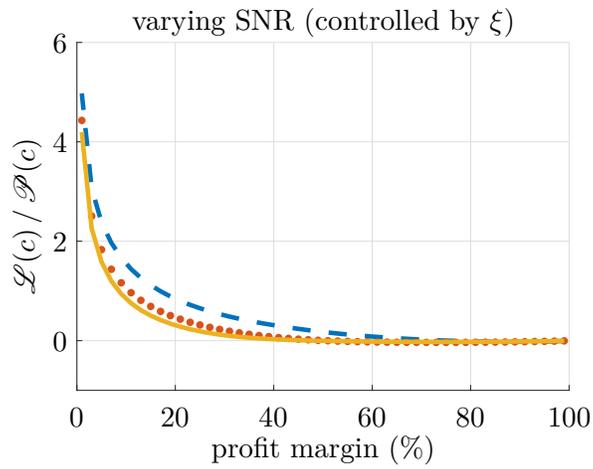
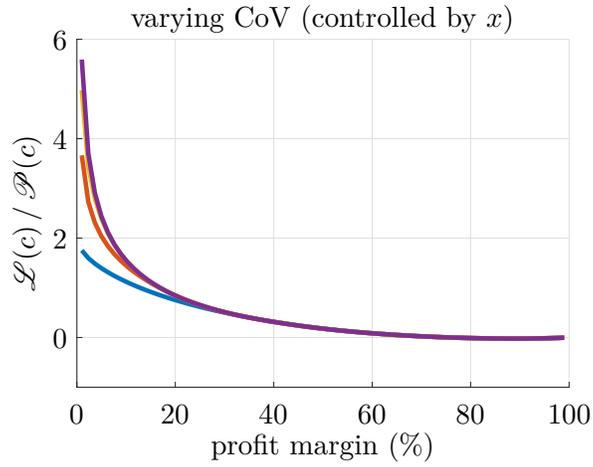
- In the top panel, we vary the first-period demand CoV, as controlled by  $x$ . We observe that as  $x$  increases and the CoV decreases, the ratio tends to increase. This suggests that as there is less first-period uncertainty, in relative terms, the pooling benefits tend to decline much faster than the learning benefits of flexibility. The curves also illustrate that there is a profit margin threshold below which learning dominates pooling benefits, and this threshold is increasing with  $x$ , i.e., with decreased first-period uncertainty.
- In the center panel, we vary the second-period SNR, as controlled by  $\xi$ . We observe that as  $\xi$  increases and the SNR decreases, the ratio  $\mathcal{L}(c)/\mathcal{P}(c)$  tends to decrease. This is because the pooling benefits are unaffected by the second period SNR, whereas the learning benefits fade with decreasing SNR, as learning becomes harder. The curves also illustrate that the profit margin threshold below which learning dominates pooling benefits is decreasing with  $\xi$ , i.e., with decreased second-period SNR.
- In the bottom panel, we consider additional demand distributions, namely Normal and Gamma, the parameters of which are chosen so that they have the same mean and standard deviation as the Uniform distribution in our base case.

Furthermore, in Figure 4, we consider a similar setup, but now assume that the unit cost of flexible capacity comes at a premium  $\alpha \geq 0$ , so that  $c_F = (1 + \alpha)c$ , and we vary parameter  $\alpha$ . We observe that the ratio  $\mathcal{L}(c_F)/\mathcal{P}(c_F)$  tends to increase with  $\alpha$ . Consequently, the threshold above which the learning benefits dominate the risk pooling benefits, appears to be increasing as flexible capacity becomes more expensive, i.e., as  $\alpha$  increases: when  $\alpha = 0$ , learning benefits are significant for margins less than 30%; when  $\alpha = 5\%$ , learning benefits are significant for margins less than 50%, approximately.

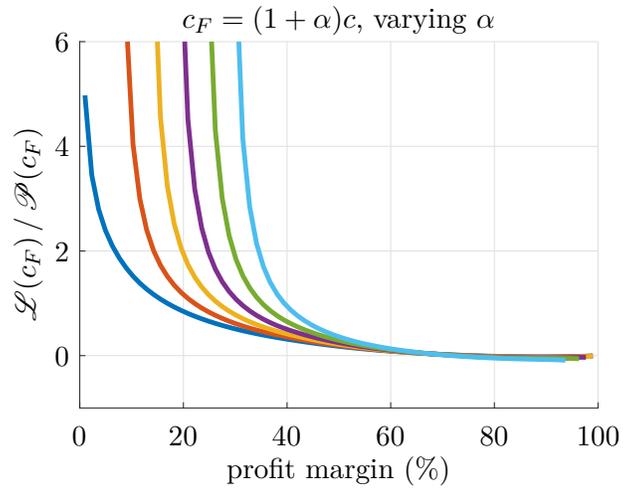
All these results reinforce our intuition about the dependence of flexibility’s learning benefits on profit margins. Specifically, in line with our earlier findings, the learning benefits of flexibility are significant and of the same order of magnitude as, or dominate the pooling benefits, for moderate and low margins.

## 4.2. Uncaptured Profits

To complement the results above, we provide additional quantification of the learning benefits of flexibility that focuses on uncaptured profit: we consider the case in which the firm is ignorant about demand correlation, acting myopically by optimizing one-period profits in both periods without updating its prior, and we measure how much money is “left on the table.”



**Figure 3:** The ratio of the learning to the pooling benefits of flexibility,  $\mathcal{L}(c)/\mathcal{P}(c)$ , as a function of profit margin for the base-case parameter setting. In the top panel, we consider varying CoV as controlled by  $x \rightarrow 0, 25, 50, 75$  (0 lowermost, 75 uppermost); in the center panel, we consider varying SNR as controlled by  $\xi \rightarrow 0, .5, 1$  (0 dashed, 1 solid); in the bottom panel, we consider different demand distributions, Uniform (dashed), Normal (dotted), and Gamma (solid).



**Figure 4:** The ratio of the learning to the pooling benefits of flexibility,  $\mathcal{L}(c_F) / \mathcal{P}(c_F)$ , as a function of profit margin for the base-case parameter setting, where  $c_F = (1 + \alpha)c$  for varying cost premium  $\alpha \rightarrow 1, 1.01, \dots, 1.05$  (1 leftmost, 1.05 rightmost).

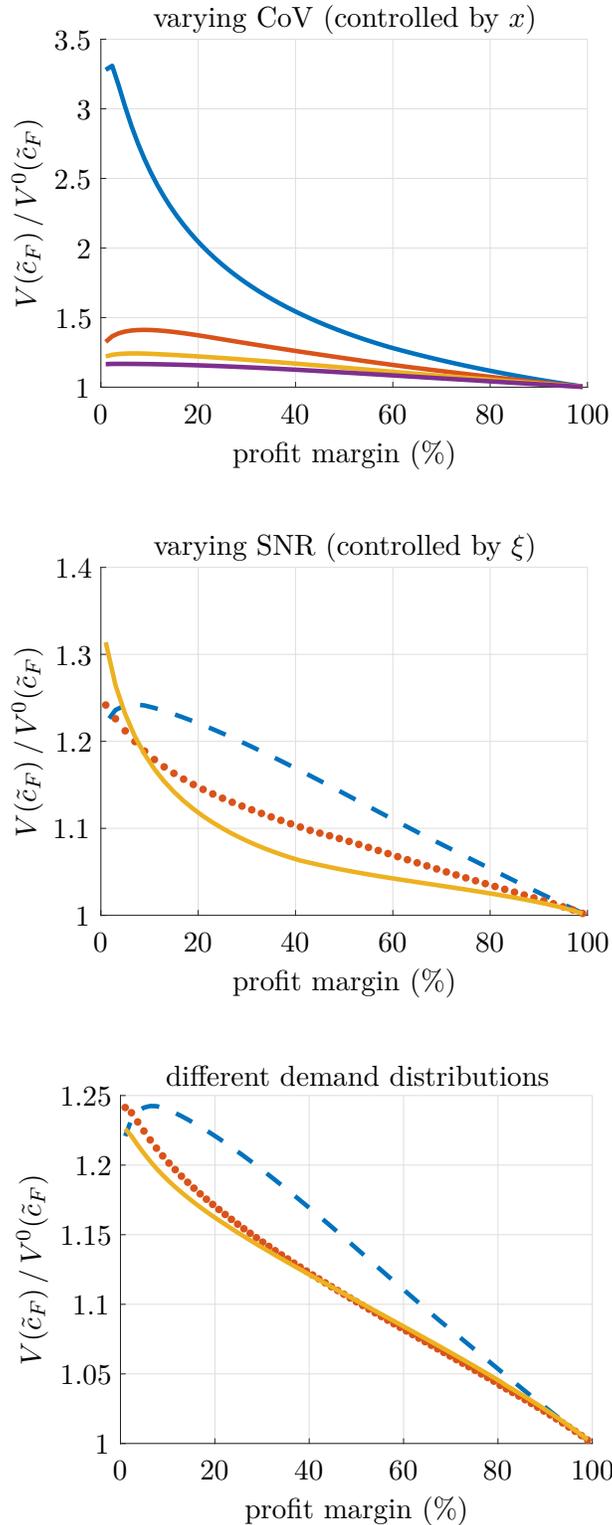
When accounting for learning, the firm generates profits  $V_X$ , when choosing system  $X \in \{N, F\}$ . Let  $V(c_F) \equiv \max\{V_N, V_F(c_F)\}$  then be the maximum profit it can generate when choosing the most profitable system. Suppose now that the firm was ignorant about the underlying demand correlation, but rather assumed independent demands, as discussed in the Section 2. Let  $V_X^0$  be the optimal profits it would generate under this assumption, when choosing system  $X \in \{N, F\}$ . Accordingly,  $V^0(c_F) \equiv \max\{V_N^0, V_F^0(c_F)\}$  would be the optimal profit among the two systems. The ratio  $V(c_F)/V^0(c_F)$  is a metric of how much the firm improves in terms of expected profit by accounting for demand correlation/learning. Equivalently, it relates directly to the money that the firm would leave on the table if it did not take learning into account. Again, we focus on the case of break-even unit cost of flexibility,  $c_F = \tilde{c}_F$ , so that the firm is indifferent between the flexible and the non-flexible system in the absence of learning opportunities.

In Figure 5, we present the results of our numerical analysis: In the top panel, we plot  $V(\tilde{c}_F)/V^0(\tilde{c}_F)$  as a function of the profit margin for varying CoV; in the center panel, we consider varying SNR; in the bottom panel, we consider additional demand distributions, calibrated to have the same mean and standard deviation as the Uniform distribution in the base case. Overall, these experiments reveal that uncaptured profits tend to decrease with profit margins, which is consistent with our finding that learning benefits are important for moderate to low margins.

## 5. Understanding the Learning Benefits of Flexibility

So far we have shown that, in the regime of moderate and low profit margins, flexibility possesses learning capabilities that are considerable relative to any learning capabilities that dedicated resources may have, as well as the risk pooling benefits of flexibility. In this section, our goal is to understand the learning benefits of flexibility better and derive additional insights.

Specifically, we start by establishing, perhaps unsurprisingly, that flexibility achieves a less variable second-period demand forecast, as measured by the variance of the posterior distribution, which, in turn, explains in part its (monetary) learning benefits. Then, we investigate whether the flexible system’s ability to generate a less variable forecast is due to excess capacity expansion relative to the non-flexible one. To this we provide a negative answer: the relative increase in capacity that demand learning justifies is lower in the flexible system compared to the non-flexible one. This leads to the logical conclusion that the flexible system forecasts “better”—and, thus, inherits the associated monetary learning benefits—because it makes more intelligent use of its available capacity. On that end, we prove that flexibility mitigates censoring more effectively



**Figure 5:** The ratio  $V(\tilde{c}_F)/V^0(\tilde{c}_F)$  as a function of profit margin for the base-case parameter setting. In the top panel, we consider varying CoV as controlled by  $x \rightarrow 0, 25, 50, 75$  (0 uppermost, 75 lowermost); in the center panel, we consider varying SNR as controlled by  $\xi \rightarrow 0, .5, 1$  (0 dashed, 1 solid); in the bottom panel, we consider different demand distributions, Uniform (dashed), Normal (dotted), and Gamma (solid).

than dedicated capacity under a variety of practical settings, demand distributions, and model parameters. This provides an intuitive explanation regarding the origins of the learning benefits of flexibility.

### 5.1. Variance of the Posterior

We begin by exploring which system affords a “better” second-period demand forecast. Since in our setting, demand forecasts are distributions, and not mere point estimates, we consider a posterior distribution with a smaller variance to be a better forecast.

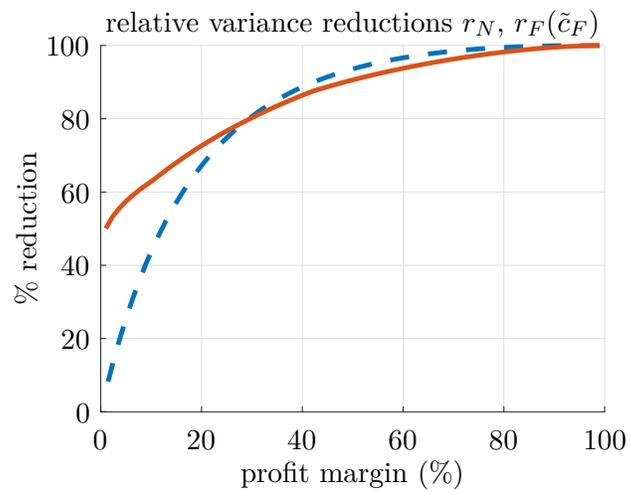
Let  $\text{VAR}_F(c_F)$  denote the expected (with respect to the prior distribution) sum of variances of the posterior for the two second-period demands for the flexible system;  $\text{VAR}_N$  denotes the same quantity for the non-flexible system; and  $\text{VAR}^0$  denotes the corresponding sum of variances of the prior distribution. We consider the relative *variance reduction* between the prior and posterior distribution:

$$r_X \equiv \frac{\text{VAR}^0 - \text{VAR}_X}{\text{VAR}^0}, \quad X \in \{N, F\}.$$

A larger value for  $r_X$  indicates a larger variance reduction and, as we argue above, suggests a better demand forecast. As in Section 3, we conduct our comparison of the two systems under the break-even cost of flexibility  $\tilde{c}_F$ , so as to isolate any effects from pooling.

**Proposition 3.** *Suppose that  $x = 0$  and  $\xi = 0$ . There exists a profit margin threshold, below (resp. above) which  $r_F(\tilde{c}_F) \geq r_N$  (resp.  $r_F(\tilde{c}_F) \leq r_N$ ).*

Under the assumptions in Proposition 3, we find the threshold below which the flexible system achieves a larger variance reduction to be approximately 27%. Figure 6 illustrates the variance reduction for the two systems under the base-case parameters, as function of the profit margin. As expected, increasing profit margins lead to higher installed capacities and, therefore, higher variance reductions for both systems. For profit margins less than 30% or so, we observe that flexible resources yield a much higher reduction than non-flexible resources. For larger margins, both systems manage to reduce demand variance substantially, owing to ample capacities, and almost by the same amount. This pattern is very much consistent with the one that emerged from our profitability comparison in Section 3: for moderate to low margins, flexibility has a considerable edge, whereas for high margins, the performance of the two systems becomes almost indistinguishable.



**Figure 6:** The variance reduction for the flexible system,  $r_F(\tilde{c}_F)$ , in solid line, and for the non-flexible system,  $r_N$ , in dashed line, as functions of the profit margin for the base-case parameter setting.

## 5.2. Capacity Expansion

It is well documented that, when a firm attempts to maximize expected profits while learning the demand, higher capacity levels are justified compared to the case in which there is no value to learning, so as to mitigate censoring; a phenomenon that the literature terms “information stalking,” e.g., see [Lariviere and Porteus, 1999]. It can be readily checked that in our model the first-period capacity is, accordingly, larger in the presence of learning considerations, rather than in their absence, that is  $\bar{Q}_X^1 \geq \bar{Q}_X^\perp$ , for both  $X \in \{N, F\}$ . We denote the resultant relative *capacity expansion* due to learning considerations by

$$\kappa_X \equiv \frac{\bar{Q}_X^1 - \bar{Q}_X^\perp}{\bar{Q}_X^\perp}, \quad X \in \{N, F\}.$$

As in previous sections, we conduct our comparison of the two systems under the break-even cost of flexibility  $\tilde{c}_F$  so as to isolate any effects from risk pooling.

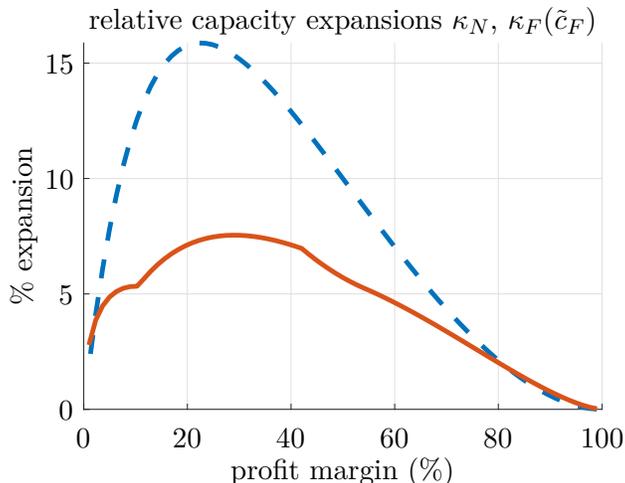
**Proposition 4.** *Suppose that  $x = 0$  and  $\xi = 0$ . Then, there exist  $\underline{\zeta}$  and  $\bar{\zeta}$  such that  $\kappa_N \geq \kappa_F(\tilde{c}_F)$  if and only if the profit margin satisfies  $\underline{\zeta} \leq \frac{p-c}{p} \leq \bar{\zeta}$ .*

The result suggests that, unless profit margins are too low or too high, learning prompts a higher capacity expansion for the non-flexible system compared with the flexible system. Under the assumptions in Proposition 4, we find the lower and upper profit thresholds between which the flexible system exhibits a lower capacity expansion to be  $\underline{\zeta} \approx 4\%$  and  $\bar{\zeta} \approx 76.5\%$ , respectively. Figure 7 depicts the relative capacity expansions for the two systems,  $\kappa_N$  and  $\kappa_F(\tilde{c}_F)$ , as functions of the profit margin for the base-case parameter setting. We observe that for profit margins outside the range established in Proposition 4, the two systems aim for rather diminishing and almost identical expansions.

The key takeaway is that the relative expansion in capacity installed at the beginning of the first period for the flexible system tends to be less than that for the non-flexible system. Hence, flexibility’s learning benefits—the monetary benefits established in Proposition 1, as well as the forecasting benefits established in Proposition 3—appear to be achieved precisely through the efficient use of resources that flexibility allows, not through excess capacity.

## 5.3. Censoring and Learning

We have seen so far that flexibility’s (monetary and informational) learning benefits are achieved through the efficient use of resources that it allows. What does “efficient use” really mean though?



**Figure 7:** The capacity expansions for the non-flexible system,  $\kappa_N$ , in dashed line, and for the flexible system,  $\kappa_F(\tilde{c}_F)$ , in solid line, as functions of the profit margin for the base-case parameter setting.

We show that flexibility mitigates censoring more effectively than dedicated capacity, unless the products have very different economic characteristics or their demand characteristics exhibit certain asymmetries. Interestingly, we find this to hold true regardless of whether the firm is myopic or takes learning into account. These findings provide an intuitive explanation regarding the origins of its learning benefits.

First, let us attempt to put things in a broader perspective. Apart from generating revenues, a firm’s sales produce exact or censored demand observations, or “demand signals.” One can envision a variety of situations where demand signals are subsequently used to update the firm’s prior, and this updating can generate value. For instance, there are many different ways in which the prior can be updated: demand could be auto-correlated across time, and the demand signals can be used to improve demand forecast for a future period, like in the model of Section 2; alternatively, the demand signals can be used to update the estimate of an uncertain demand distribution parameter, again to improve demand forecast for a future period. There are also many different ways in which the updated prior can generate value: the firm can use it, for example, to make better future capacity decisions, as in the model of Section 2, or assortment, or pricing decisions; or, it can use it to simply provide better service quality by knowing its customers better; or, it could even sell it to third-party firms that specialize in market research.

As the examples above illustrate, the prior updating process and the resultant value generation process are both context-specific, and their analyses would require very different models. However, a unifying feature of all such models is that exact demand observations lead to more accurate updating of the prior, and thus could generate more value to the firm, than censored demand

observations. Our intention in this part of the analysis is to abstract away from the specifics of both prior updating and value generation processes, while capturing in the most parsimonious fashion the core idea that these processes are more efficient under exact demand observations. Therefore, we focus on the static model of Section 2.1, and we investigate the ability of flexible resources to mitigate censoring and, thus, produce exact demand observations. Specifically, in each of the two system, flexible and non-flexible, we count the number of product demands that the firm observes censored, i.e.,

$$C_X \equiv \sum_{i=1}^2 \mathbf{1}(S_{X,i} < D_i), \quad X \in \{N, F\},$$

and define the *censoring mitigation capability* of each system as

$$\mathcal{L}_X \equiv 2 - \mathbb{E}[C_X], \quad X \in \{N, F\}. \quad (15)$$

We say that the flexible system mitigates censoring more effectively than the non-flexible one if  $\mathcal{L}_F \geq \mathcal{L}_N$ , i.e., if the firm expects (with respect to the prior distributions/beliefs) to observe more product demands uncensored by using the flexible system.

In what follows, we compare the censoring mitigation capabilities of the two systems for the case of the *break-even unit cost* of flexibility. Again, we first study the “symmetric” case in which the two products have i.i.d. demands and equal margins. Studying the symmetric case is already of practical interest, given that products that can be produced using the same flexible resources are often inherently similar to each other, e.g., cellphones with GSM or CDMA antennas. Furthermore, studying the symmetric case enables us to arrive at baseline results that prove useful in the second step of this analysis, in which we explore the effect of each product characteristic on the comparison of the censoring mitigation capabilities of the two systems.

**Proposition 5(a).** *Consider the static model of Section 2.1 with  $c_F = \tilde{c}_F$ . Assume that the demands are independent and distributed identically to a uniform distribution, and that the products have the same margin, i.e.,  $p_1 = p_2$ . Then, the flexible system mitigates censoring more effectively than the non-flexible one, i.e.,*

$$\mathcal{L}_F \geq \mathcal{L}_N.$$

In the presence of heterogeneity with respect to the demand/ or the economic characteristics of the two products, we establish the following result:

**Proposition 5(b).** *Consider the static model of Section 2.1 with  $c_F = \tilde{c}_F$ . Assume that the demands are independent and uniformly distributed with means  $\mu$  and standard deviations  $\sigma$ . The flexible*

system mitigates censoring more effectively than the non-flexible one, i.e.,  $\mathcal{L}_F \geq \mathcal{L}_N$ ,

(i) if  $\mu_1 \neq \mu_2$  and  $\sigma_1 = \sigma_2$ ,  $p_1 = p_2$ .

(ii) if  $\sigma_2 \leq \tilde{\lambda}\sigma_1$  and  $\mu_1 = \mu_2$ ,  $p_1 = p_2$ , for some  $\tilde{\lambda} > 1$ .

(iii) if  $p_1 \leq 3p_2$  and  $\mu_1 = \mu_2$ ,  $\sigma_1 = \sigma_2$ .

We conclude by showing that when underage is costlier than overage, the effect of different demand characteristics we that find in our previous result fades.

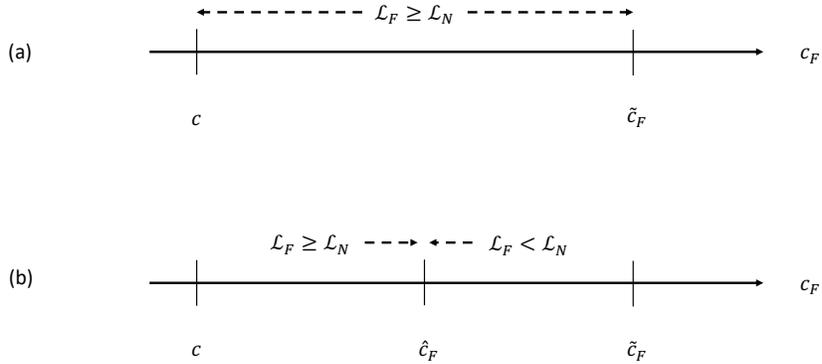
**Proposition 5(c).** *Consider the setting in Proposition 5(b) and further assume that  $p_i - c \geq c$ . Then, we can select  $\tilde{\lambda} \rightarrow \infty$  and the results continue to hold.*

Notably, the effect of different economic characteristics persists in this case, so that the learning capabilities of flexible resources dominate in this scenario only if margins do not differ by too much, specifically, if  $p_1 \leq 3p_2$ .

We also compare the censoring mitigation capabilities of the two systems, assuming normally distributed demands, in an extensive numerical study whose details are provided in Appendix B.1. In summary, we compare  $\mathcal{L}_F$  to  $\mathcal{L}_N$  while varying the fractile  $(p_i - c)/p_i$ , the demand variability  $\sigma_i/\mu_i$ , the demand correlation  $\rho$ , the asymmetry in mean demand  $\mu_2/\mu_1$ , the asymmetry in demand variability  $\sigma_2/\sigma_1$ , the asymmetry in unit revenues  $p_2/p_1$ , and the various combinations thereof. We find that flexibility mitigates censoring more effectively than dedicated capacity, as long as the economic characteristics of the two products are not vastly different. Namely, the only instances in which  $\mathcal{L}_F < \mathcal{L}_N$  are those in which  $p_1 \geq 4p_2$ .

In Appendix B.2 we analyze two more scenarios, that of *equal costs*, where flexible capacity has the same unit cost as dedicated capacity, and that of *equal capacities*, where the flexible system's resource capacity equals the sum of dedicated capacities of the non-flexible system. The former case simply provides a benchmark and, as expected, the flexible system dominates the non-flexible one in terms of censoring mitigation, irrespective of the demand distribution and the profit margin. The latter case could be relevant, for example, when pre-installed dedicated resources can be converted to flexible resources at a fixed cost (and negligible variable cost).

Although the technical analysis of each of the three scenario presents its own nuances, the following unifying finding emerges: *the flexible system mitigates censoring more effectively than the non-flexible one, unless the products have very different economic characteristics or their demand characteristics exhibit certain asymmetries.*



**Figure 8:** Range of the unit cost of flexible capacity for which the flexible system mitigates censoring more effectively than the non-flexible one. Figure 8(a) depicts a scenario in which the conditions of Proposition 5(b) are satisfied. In this case, flexibility facilitates learning in the entire range of the relevant unit cost of flexible capacity,  $[c, \tilde{c}_F]$ . Figure 8(b) depicts the situation otherwise. In that case, flexibility facilitates learning only in the interval  $[c, \hat{c}_F]$ , for some  $\hat{c}_F < \tilde{c}_F$ .

In Appendix B.3 we perform a robustness check wherein learning is accounted for in the firm’s objective, and we derive findings that further back the statement above.

Finally, we extrapolate from the three individual scenarios to the entire range of  $c_F$  that justifies investing in flexibility even in the absence of learning, namely, the interval  $[c, \tilde{c}_F]$ . Note that as  $c_F$  decreases, the optimal capacity, and thus the capability of the flexible system to mitigate censoring increases, or, in the worst case, stay the same. This implies that as  $c_F$  decreases, the conditions for flexibility to mitigate censoring more effectively than dedicated capacity become less stringent. This monotonicity, together with Propositions 6 and 5(b), give rise to the following results for the general case of arbitrary  $c_F$ .

**Corollary 1.** *Consider the static model of Section 2.1 with  $c_F \in [c, \tilde{c}_F]$ . Suppose that the assumptions of Proposition 5(b) are satisfied. If conditions (i)-(iii) are met, then*

$$\mathcal{L}_F \geq \mathcal{L}_N, \quad \forall c_F.$$

*Otherwise, there exists a threshold  $\hat{c}_F \in [c, \tilde{c}_F)$  such that*

$$\mathcal{L}_F \geq \mathcal{L}_N, \quad \forall c_F \leq \hat{c}_F, \quad \text{and} \quad \mathcal{L}_F < \mathcal{L}_N, \quad \forall c_F > \hat{c}_F.$$

Corollary 1 reinforces the insight that flexibility always mitigates censoring more effectively, unless demand distributions or margins are “too different” from each other—where “too different” is used in the sense of Proposition 5(b). Figure 8 provides a graphical illustration of these results.

## 6. Conclusions, Limitations and Future Research

In this paper we establish a heretofore unknown benefit of resource flexibility: it can facilitate learning the demand, when the latter is not fully observable or censored. In turn, this allows for more accurate future forecasting and better-informed operational decisions *ex ante* uncertainty resolution in future periods, ultimately increasing profitability and operations performance.

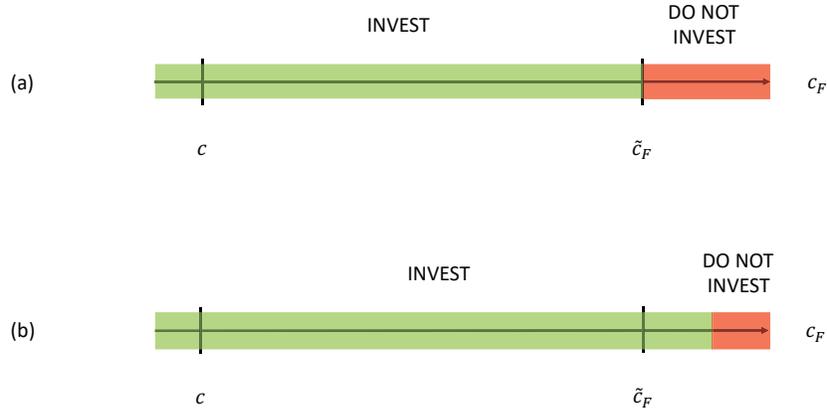
Our analysis shows that for low to moderate profit margins the learning benefits of flexibility could be substantial, and of the same order of magnitude as the celebrated risk-pooling benefits in terms of profitability. We also illustrate how these benefits derive from a more efficient use of resources—rather than excess capacity—that mitigates censoring. The effects tend to accentuate as prior uncertainty increases; or as correlation between prior and future demand increases, i.e., as demand observations have greater informational content in predicting the future.

For settings in which margins are high, ample capacities are already justified, and censoring becomes rare. Consequently, learning benefits fade. Similarly, our theory is less relevant to settings in which demand is observed in a credible way, as is the case, for example, in certain B2B interactions.

Our work has the potential to change the way operations managers think about flexibility in business environments in which demand learning is valuable: beyond the mere ability to hedge against uncertainty, flexibility can help reduce it, allowing firms to make better-informed decisions in future periods. As learning capabilities often represent a competitive priority in today’s rapidly changing business landscape, learning emerges as one of flexibility’s distinct benefits alongside risk pooling, revenue maximization, and operational hedging. This also suggests that the true value of flexibility could be greater than suggested by the existing models, which account merely for its reactive abilities. Therefore, a firm that cares about learning would be willing to pay more for flexibility than a firm that does not; see Figure 9 for a graphical illustration.

Our research could open up a whole new range of possible research avenues. To point out a few, it would be interesting to quantify the monetary value of the learning benefits of flexibility in additional settings. In this paper, we make the first step towards this goal, but one could explicitly model other contingent decisions that may benefit from superior demand forecasts, such as assortment, promotions, and pricing.

This line of research can also be thought as extending the seminal works of Fine and Freund [1990] and Van Mieghem [1998]: given the learning benefits of flexibility, how would the insights of these papers be affected in multi-period settings? Another interesting research direction would be to study learning benefits of process flexibility in the context of Jordan and Graves [1995]. Does limited



**Figure 9:** Range of the unit cost of flexible capacity  $c_F$ , for which it is optimal to invest in the flexible system over the non-flexible one. Figure 9(a) depicts the situation in the absence of any learning considerations. Figure 9(b) illustrates the case of a firm that takes learning into account. Then, the range over which it is optimal to invest in flexibility is larger.

flexibility, configured to chain products and resources together to the greatest extent possible, yield most of the *learning* benefits of total flexibility? Finally, we focus on resource flexibility, which subsumes many different forms such as employing a flexible manufacturing system, cross-training workforce, or delaying product differentiation. How about different types of flexibility, such as quick response or lowering setup costs?

## References

- Alain Bensoussan and Pengfei Guo. Technical note—managing nonperishable inventories with learning about demand arrival rate through stockout times. *Operations Research*, 63(3):602–609, 2015.
- Dimitri P Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA, 2005.
- Omar Besbes and Alp Muharremoglu. On implications of demand censoring in the newsvendor problem. *Management Science*, 59(6):1407–1424, 2013.
- Omar Besbes, Juan Chaneton, and Ciamac C. Moallemi. The exploration-exploitation trade-off in the newsvendor problem. *Columbia Business School Research Paper 14-61*, 2017.
- Kostas Bimpikis and Mihalis G. Markakis. Inventory pooling under heavy-tailed demand. *Management Science*, 62(6):1800–1813, 2016.

- Arnab Bisi, Maqbool Dada, and Surya Tokdar. A censored-data multiperiod inventory problem with newsvendor demand distributions. *Manufacturing & Service Operations Management*, 13(4):525–533, 2011.
- Apostolos N. Burnetas and Craig E. Smith. Adaptive ordering and pricing for perishable products. *Operations Research*, 48(3):436–443, 2000.
- Li Chen. Bounds and heuristics for optimal bayesian inventory control with unobserved lost sales. *Operations Research*, 58(2):396–413, 2010.
- Li Chen and Erica L. Plambeck. Dynamic inventory management with learning about the demand distribution and substitution probability. *Manufacturing & Service Operations Management*, 10(2):236–256, 2008.
- Jiri Chod and Nils Rudi. Resource flexibility with responsive pricing. *Operations Research*, 53(3):532–548, 2005.
- Jiri Chod, Nils Rudi, and Jan A. Van Mieghem. Operational flexibility and financial hedging: Complements or substitutes? *Management Science*, 56(6):1030–1045, 2010.
- Vivek Choudhary, Sameer Hasija, and Serguei Netessine. Do flexibility and chaining really help? an empirical analysis of automotive plant networks. Technical report, SSRN, Dec 2018. URL [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3301302](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3301302).
- Qing Ding, Lingxiu Dong, and Panos Kouvelis. On the integration of production and financial hedging decisions in global markets. *Operations Research*, 55(3):470–489, 2007.
- Xiaomei Ding, Martin L. Puterman, and Arnab Bisi. The censored newsvendor and the optimal acquisition of information. *Operations Research*, 50(3):517–527, 2002.
- Gary D. Eppen. Note—effects of centralization on expected costs in a multi-location newsboy problem. *Management Science*, 25(5):498–501, 1979.
- Charles H. Fine and Robert M. Freund. Optimal investment in product-flexible manufacturing capacity. *Management Science*, 36(4):449–466, 1990.
- Marshall Fisher and Ananth Raman. *The new science of retailing: how analytics are transforming the supply chain and improving performance*. Harvard Business Review Press, 2010.

- Gregory A. Godfrey and Warren B. Powell. An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Science*, 47(8):1101–1112, 2001.
- William Golden and Philip Powell. Towards a definition of flexibility: In search of the holy grail? *Omega*, 28(4):373 – 384, 2000.
- Manu Goyal and Serguei Netessine. Volume flexibility, product flexibility, or both: the role of demand correlation and product substitution. *Manufacturing & Service Operations Management*, 13(2):180–193, 2011.
- Giora Harpaz, Wayne Y. Lee, and Robert L. Winkler. Learning, experimentation, and the optimal output decisions of a competitive firm. *Management Science*, 28(6):589–603, 1982.
- Arnd Huchzermeier and Morris A. Cohen. Valuing operational flexibility under exchange rate risk. *Operations Research*, 44(1):100–113, 1996.
- Woonghee Tim Huh and Paat Rusmevichientong. A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research*, 34(1):103–123, 2009.
- Woonghee Tim Huh, Retsef Levi, Paat Rusmevichientong, and James B. Orlin. Adaptive data-driven inventory control with censored demand based on kaplan-meier estimator. *Operations Research*, 59(4):929–941, 2011.
- Dan A Iancu, Nikolaos Trichakis, and Gerry Tsoukalas. Is operating flexibility harmful under debt? *Management Science*, 63(6):1730–1761, 2016.
- Aditya Jain, Nils Rudi, and Tong Wang. Demand estimation and ordering under censoring: stock-out timing is (almost) all you need. *Operations Research*, 63(1):134–150, 2015.
- William C. Jordan and Stephen C. Graves. Principles on the benefits of manufacturing process flexibility. *Management Science*, 41(4):577–594, 1995.
- James A Kahn. Inventories and the volatility of production. *The American Economic Review*, pages 667–679, 1987.
- Sumit Kunnunkal and Huseyin Topaloglu. Using stochastic approximation methods to compute optimal base-stock levels in inventory control problems. *Operations Research*, 56(3):646–664, 2008.

- Martin A. Lariviere and Evan L. Porteus. Stalking information: bayesian inventory management with unobserved lost sales. *Management Science*, 45(3):346–363, 1999.
- Hau L. Lee and Christopher S. Tang. Modelling the costs and benefits of delayed product differentiation. *Management Science*, 43(1):40–53, 1997.
- Hau L Lee, Venkata Padmanabhan, and Seungjin Whang. Information distortion in a supply chain: the bullwhip effect. *Management science*, 50(12\_supplement):1875–1886, 2004.
- Patrick H Liu, Stuart Smith, Eric B Orkin, and George Carey. Estimating unconstrained hotel demand based on censored booking data. *Journal of Revenue and pricing Management*, 1(2):121–138, 2002.
- Xiangwen Lu, Jing-Sheng Song, and Kaijie Zhu. Technical note—analysis of perishable-inventory systems with censored demand data. *Operations Research*, 56(4):1034–1038, 2008.
- Gábor Lugosi, Mihalis G. Markakis, and Gergely Neu. On the hardness of inventory management with censored demand data. *arXiv Paper 1710.05739*, 2017.
- Costis Maglaras and Serkan Eren. A maximum entropy joint demand estimation and capacity control policy. *Production and Operations Management*, 24(3):438–450, 2015.
- Adam J. Mersereau. Demand estimation from censored observations with inventory record inaccuracy. *Manufacturing & Service Operations Management*, 17(3):335–349, 2015.
- K Sridhar Moorthy and Ivan PL Png. Market segmentation, cannibalization, and the timing of product introductions. *Management science*, 38(3):345–359, 1992.
- Serguei Netessine, Gregory Dobson, and Robert A. Shumsky. Flexible service capacity: Optimal investment and the impact of demand correlation. *Operations Research*, 50(2):375–388, 2002a.
- Serguei Netessine, Gregory Dobson, and Robert A Shumsky. Flexible service capacity: Optimal investment and the impact of demand correlation. *Operations Research*, 50(2):375–388, 2002b.
- Pablo A Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical programming*, 96(2):293–320, 2003.
- Warren Powell, Andrzej Ruszczyński, and Huseyin Topaloglu. Learning algorithms for separable approximations of discrete stochastic optimization problems. *Mathematics of Operations Research*, 29(4):814–836, 2004.

- Gennady Samorodnitsky and Murad S. Taqqu. *Stable non-gaussian random processes*. Chapman & Hall, London, UK, 1994.
- Leroy B. Schwarz. A model for assessing the value of warehouse risk-pooling: Risk-pooling over outside-supplier leadtimes. *Management Science*, 35(7):828–842, 1989.
- Cong Shi, Weidong Chen, and Izak Duenyas. Nonparametric data-driven algorithms for multi-product inventory systems with censored demand. *Operations Research*, 64(2):362–370, 2016.
- Jayashankar M. Swaminathan and Sridhar R. Tayur. Managing broader product lines through delayed differentiation using vanilla boxes. *Management Science*, 44(12-part-2):S161–S172, 1998.
- J Van Mieghem and G Allon. Operations strategy. *Dynamic Ideas: Belmont, MA, USA*, 2008.
- Jan A. Van Mieghem. Investment strategies for flexible resources. *Management Science*, 44(8):1071–1078, 1998.
- Jan A. Van Mieghem. Risk mitigation in newsvendor networks: Resource diversification, flexibility, sharing, and hedging. *Management Science*, 53(8):1269–1288, 2007.
- Garrett van Ryzin and Jeff McGill. Revenue management without forecasting or optimization: an adaptive algorithm for determining airline seat protection levels. *Management Science*, 46(6):760–775, 2000.