Statement of Research Contributions

Cong Shi

Generally speaking, my research thus far has focused on prescribing *efficient and effective algorithms with theoretical performance guarantees* for several core and important models arising in the context of *inventory management*, *supply chain management*, *revenue management*, and *service operations.* In what follows, I shall categorize my work into three broad areas, and summarize my major contributions in each area.

I: Learning Algorithms for Inventory and Supply Chain Systems

The focus of this research thrust is to design efficient learning algorithms based on observable demand data to make provably-good adaptive operational decisions in stochastic inventory systems. Traditionally, these models are considered under the assumption that the demand distribution is fully known as part of the input. However, in practical scenarios, this assumption is often unrealistic, and even correctly specifying the parametric form of the demand distribution could be difficult. My research takes a *nonparametric* approach that directly works with past observable demand data collected over time to make adaptive decisions [7].

It is worthwhile highlighting several unique challenges in which the general learning algorithms in the OR/MS and Computer Science and Statistics literature cannot be readily employed in our settings. First, in most retail settings, customers typically walk away in the face of stock-out, and therefore the system is unable to keep track of these lost-sales. Thus, the observable demand data is, in fact, the sales data, which is also known as the *censored demand* data. Second, the inventory (and pricing) decisions may impact the cost/profit function over an extended period, due to the *multi-dimensional* state variable in the underlying stochastic inventory system. Third, the stochastic inventory system has *hard physical constraints*, e.g., positive inventory carry-over, warehouse capacity constraint, ordering/production capacity constraint, and these constraints limit the search space in a dynamic way. Our performance measure is the natural notion of *regret*, which is the difference between the cost of our proposed learning algorithm and the clairvoyant optimal cost (had all the distributions of randomness been specified *a priori*). We have tackled the following core and important inventory/pricing problems over the past five years.

Perishable Inventory Systems under Censored Demand [23]. Perishable products are prevalent in our daily lives, e.g., meat, fruit, vegetable, dairy products, and frozen foods in the supermarket, pharmaceuticals like drugs and vitamins, and the blood products in the blood bank. On the other hand, perishable inventory systems (with a general positive product lifetime m) are notoriously hard to solve, due to the need to keep track of the age information of the on-hand inventory which leads to an m-dimensional state vector. We first show that the total cost of any base-stock system is convex in the base-stock level on every sample path. Next, restricting to the class of base-stock policies, we develop the first nonparametric learning algorithm in which the average regret scales as $O(1/\sqrt{T})$ for any planning horizon T, which is shown to be theoretically the best scaling (for any learning algorithms). Our algorithm and analysis require a novel method for computing a valid cycle subgradient and the construction of a bridging problem (that hypothetically renews a subset of older inventory). Our numerical results also confirm the efficacy of the proposed learning algorithm.

Lost-Sales Inventory Systems with Lead Times under Censored Demand [22]. The periodic-review inventory control problem with lost-sales and a positive lead time L is an *iconic* problem in the theory of inventory management. While prevalent in applications, the problem is challenging due to the need to keep track of the pipeline inventory vector which again leads to a L-dimensional state vector. Restricting to the class of base-stock policies, we develop a new nonparametric learning algorithm in which the average regret scales as $O(1/\sqrt{T})$, which improves upon the existing result $O(1/\sqrt[3]{T})$ by [8]. More importantly, our result closes the gap between upper and lower regret bounds, and attains the optimal scaling. The algorithm uses a random cycle-updating rule based on an auxiliary simulated system, and also involves two new concepts, namely, the withheld on-hand inventory and the double-phase cycle gradient estimation. The techniques developed are effective for learning systems with complex systems dynamics and lasting impact on decisions.

Joint Inventory and Pricing Systems under Censored Demand [3]. The periodic-review joint inventory control and pricing problem has received much attention in the literature, and is a core model in operations management. The lost-sales version of the problem is well-known to be challenging because the profit function fails to be jointly concave in inventory level and price. The sampled analog of the profit function could even be *multimodal* in price. We give the first nonparametric learning algorithm in which the average regret scales as $O(T^{-1/5}(\log T)^{1/4})$. Our algorithm integrates exploration and exploitation through carefully designed cycles, and searches the decision space through a new sparse discretization scheme.

Capacitated Inventory Systems under Censored Demand [16]. We have also considered stochastic inventory systems under capacity constraints, which again are core models in operations management. The first system involves multiple products and a warehouse capacity constraint, which is very common in practice. The second system involves deterministic or random production/ordering capacity. For each system, we characterize the structure of clairvoyant optimal policies and then give the first nonparametric learning algorithm in which the average regret scales as $O(1/\sqrt{T})$, which is theoretically the optimal scaling.

We believe this line of research is well aligned with the important opportunity that now exists to advance *data-driven algorithmic decision-making under uncertainty*. Moreover, it adds an important dimension to the general theory of online learning and reinforcement learning, since the firms often face a realistic stochastic supply chain system where system dynamics are complex, constraints are abundant, and information about uncertainty in the system is typically censored. It is, therefore, important to analyze the structure of the underlying system more closely, and devise an efficient and effective learning algorithm that can generate better data, which is then feedback to the algorithm to make better decisions. This forms a virtuous cycle.

II: Approximation Algorithms for Inventory and Supply Chain Systems

Many core problems in inventory and supply chain management are naturally cast as multistage stochastic optimization problems. They are typically formulated as dynamic programs. However, the resulting dynamic programs are usually extremely hard to solve to optimality (or even approximately), both in theory and practice. The computational obstacle is the enormous state-space that gives little hope to any straightforward attempt to solve these dynamic programs, which is usually referred to as the *curse of dimensionality*. Motivated by the approximation algorithm for the classical multi-period inventory problem [12], we set out to study several more realistic stochastic inventory problems. Each problem requires new algorithmic ideas and techniques, as detailed below.

Approximation Algorithms for Perishable Inventory Systems [2, 24, 1]. Perishable products such as meat, fruit, vegetable, dairy products, and frozen foods constitute the majority of supermarket sales of more than $660 billion yearly. Perishable products also include virtually all pharmaceuticals (e.g., drugs and vitamins), which represent another multi-billion dollar industry. Developing effective inventory control policies not only increases firms revenue but also significantly reduces wastage (due to spoilage and outdating) and makes the supply chain "green and environmentally friendly. Despite its importance, these perishable inventory control problems are notoriously hard in inventory theory, due to the need to keep track a *multi-dimensional* state variable (recording the age information of inventory) in the underlying dynamic programming formulation. It is also well-known in the literature that the structure of optimal policies is extremely complicated, and indeed only a partial characterization (monotonicity and sensitivity) is available. There is a lack of efficient algorithms to compute a good inventory control policy − let alone establishing any theoretical performance guarantees. We develop the first set of approximation algorithms for perishable inventory systems, perishable inventory systems with fixed cost, perishable inventory systems with capacity and positive lead times in [2, 24, 1], respectively. The computational performance is consistently excellent. To establish their worst-case approximation ratios, we introduce the notion of *trimmed on-hand inventory level*, a *transient unit-matching rule* to dynamically match the supply and demand units, and the notion of *associated demand processes* provides the right future demand information to establish the desired results.

Approximation Algorithms for Stochastic Lot-Sizing Problems [14, 18]. The stochastic inventory problem with fixed cost is perhaps the most fundamental problem in inventory management, and it is well-known that state-dependent (s, S) policies are optimal. If ordering capacity is further considered, the structure of optimal policies can only be partially characterized. In both settings, computing the exact optimal policy is intractable, under generally nonstationary and correlated demand structures that allow for dynamic forecast updates. We develop two approximation algorithms (with worst-case performance guarantees) for stochastic inventory control models with fixed costs and/or ordering capacities.

Joint-Replenishment Problem (JRP) and Inventory Routing Problem (IRP) [15]. We consider two fundamental deterministic multiproduct inventory optimization problems over a planning horizon T with non-stationary demands and general submodular cost functions. Note that [6] has considered the same problem but with three special classes of submodular functions. We present a unified approach that yields $O(\log T/\log\log T)$-approximation algorithms for both problems with polynomial holding costs. This is the first sub-logarithmic approximation ratio for either problem.

Approximation Algorithms for Service-Level Constrained Inventory Systems [9]. In practice, service levels are typically much easier to quantify than costs. We give two approximation algorithms to α-service-level (or ready rate) constrained inventory systems. The core concept developed is called the delayed forced holding and production cost, which is proven effective in dealing with such inventory systems.

III: Other Important Classes of Problems in OR/MS

Besides the above two major lines of research, I am also broadly interested in several other important topics in OR/MS, such as revenue management [4, 25], contract theory [13], scheduling and resource allocation [20, 11, 19], and transportation [21]. In the interest of space, we will only pick two what-we-believe very exciting classes of problems to highlight our main results and contributions [17, 5].

Process Flexibility for Multi-Period Production Systems [17]. Process flexibility is one of the key operational strategies that firms adopt to cope with the variabilities of product demands and reduce the operational costs. In the seminal paper [10], the authors observed that with the sparse chaining flexibility structure, one often obtains almost the same benefit as the fully flexible system. To this date, much of the theoretical analysis on this topic focuses on a single-period model. However, there is little theoretical result on the multi-period backlogging model, due to the curse of dimensionality brought by any non-trivial flexibility structure. This is a significant gap in the literature (as most practical systems run in multi-period environments). We, therefore, develop a theory for the design of process flexibility in a multi-period make-to-order production system. We propose and formalize a notion of "effective chaining" termed the *Generalized Chaining Gap* (GCG for short), which can be viewed as a natural extension of classical chaining structure from the process flexibility literature. Using the GCG, we prove that in a general system with high capacity utilization, one only needs a sparse flexibility structure with $m + n$ arcs to achieve similar performance as full flexibility, where m and n are equal to the number of plants and products in the system, respectively. We also provide a simple and efficient algorithm for finding such sparse structures. Moreover, we show that the requirement of $m + n$ arcs is tight, by explicitly constructing systems in which even the best flexibility structure with $m+n-1$ arcs cannot achieve the same asymptotic performance as full flexibility. As a result, our analysis not only echoes "a little bit of flexibility goes a long way", a recurrent theme in the process flexibility literature, but also quantifies *exactly how much* flexibility is needed in highly utilized multi-period make-to-order environment.

Inventory and Revenue Management with Strategic Customers [5]. We consider a joint inventory control and pricing problem in [5] wherein customers are *forward-looking*, i.e., they strategize their purchasing times (anticipating price promotions in the future) but incur delay disutility from postponing the purchasing decisions. A customer's arrival time and product valuation are her private information. We allow for heterogeneity in valuation as well as a positive correlation between valuation and delay disutility. The seller seeks to maximize her long-run average expected revenue less inventory costs. Through a tractable upper bound constructed by solving a dynamic mechanism design problem, we completely characterize the optimal joint pricing, delivery, and inventory policy, which turns out to be a simple cyclic policy. This result is significant for the following two reasons. First, there is strong empirical evidence that customers exhibit forward-looking behaviors. Second, finding the optimal policy is not straightforward given the complex dynamics introduced by the game between the seller and the customers.

References

[1] Chao, X., X. Gong, C. Shi, C. Yang, H. Zhang, S. X. Zhou. 2017. Approximation algorithms for capacitated perishable inventory systems with positive lead times. *Management Science*, to appear.

[2] Chao, X., X. Gong, C. Shi, H. Zhang. 2015. Approximation algorithms for perishable inventory systems. *Operations Research* **63**(3) 585–601.

[3] Chen, B., X. Chao, C. Shi. 2015. Nonparametric algorithms for joint pricing and inventory control with lost-sales and censored demand. *Mathematics of Operations Research*, Major Revision.

[4] Chen, Y., R. Levi, C. Shi. 2017. Revenue management of reusable resources with advanced reservations. *Production and Operations Management* **26**(5) 836–859.

[5] Chen, Y., C. Shi. 2018. Joint pricing and inventory management with strategic customers. *Operations Research*, to appear (pending a minor revision).

[6] Cheung, M., A. N. Elmachtoub, R. Levi, D. B. Shmoys. 2016. The submodular joint replenishment problem. *Mathematical Programming* **158**(1-2) 207–233.

[7] Huh, W. H., P. Rusmevichientong. 2009. A non-parametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research* **34**(1) 103–123.

[8] Huh, W. T., G. Janakiraman, J. A. Muckstadt, P. Rusmevichientong. 2009. An adaptive algorithm for finding the optimal base-stock policy in lost sales inventory systems with censored demand. *Mathematics of Operations Research* **34**(2) 397–416.

[9] Jiang, Y., C. Shi, S. Shen. 2018. Service-level constrained inventory systems. *Production and Operations Management*, Major Revision.

[10] Jordan, W. C., S. C. Graves. 1995. Principles on the benefits of manufacturing process flexibility. *Management Science* **41**(4) 577–594.

[11] Keyvanshokooh, E., C. Shi, M. P. Van Oyen. 2018. Advance online scheduling with overtime: A primal-dual approach. *Manufacturing & Service Operations Management*, Major Revision.

[12] Levi, R., M. Pál, R. O. Roundy, D. B. Shmoys. 2007. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* **32**(4) 821–838.

[13] Levi, R., G. Perakis, C. Shi, W. Sun. 2018. Strategic capacity planning problems in revenue sharing joint ventures. *Production and Operations Management*, Major Revision.

[14] Levi, R., C. Shi. 2013. Approximation algorithms for the stochastic lot-sizing problem with order lead times. *Operations Research* **61**(3) 593–602.

[15] Nagarajan, V., C. Shi. 2016. Approximation algorithms for inventory problems with submodular or routing costs. *Mathematical Programming* **160**(1-2) 225–244.

[16] Shi, C., W. Chen, I. Duenyas. 2016. Nonparametric data-driven algorithms for multiproduct inventory systems with censored demand. *Operations Research* **64**(2) 362–370.

[17] Shi, C., Y. Wei, Y. Zhong. 2018. Process flexibility for multi-period production systems. *Operations Research*, to appear.

[18] Shi, C., H. Zhang, X. Chao, R. Levi. 2014. Approximation algorithms for capacitated stochastic inventory systems with setup costs. *Naval Research Logistics* **61**(4) 304–319.

[19] Shi, C., H. Zhang, C. Qin. 2015. A faster algorithm for the resource allocation problem with convex cost functions. *Journal of Discrete Algorithms* **34** 137–146.

[20] Xu, Y., C. Shi, I. Duenyas. 2016. Priority rules for multi-task due-date scheduling under varying processing costs. *Production and Operations Management* **25**(12) 2086–2102.

[21] Yu, M., Y. Ding, R. Lindsey, C. Shi. 2016. A data-driven approach to manpower planning at us–canada border crossings. *Transportation Research Part A: Policy and Practice* **91** 34–47.

[22] Zhang, H., X. Chao, C. Shi. 2017. Closing the gap: A learning algorithm for the lost-sales inventory system with lead times. *Management Science*, Major Revision.

[23] Zhang, H., X. Chao, C. Shi. 2018. Perishable inventory systems: Convexity results for base-stock policies and learning algorithms under censored demand. *Operations Research*, to appear.

[24] Zhang, H., C. Shi, X. Chao. 2016. Approximation algorithms for perishable inventory systems with setup costs. *Operations Research* **64**(2) 432–440.

[25] Zhang, H., C. Shi, C. Qin, C. Hua. 2016. Stochastic regret minimization for revenue management problems with nonstationary demands. *Naval Research Logistics* **63**(6) 433–448.