Transfer Learning, Cross Learning and Co-Learning with Operational Data Analytics (ODA)

Qi Feng†, Lei Li‡, and J. George Shanthikumar†

†Mitch Daniels School of Business, Purdue University, West Lafayette, IN 47907

annabellefeng@purdue.edu shanthikumar@purdue.edu

‡Department of Logistics and Maritime Studies, Faculty of Business, The Hong Kong Polytechnic University, Hung Hom, Hong Kong, China

leihk.li@polyu.edu.hk

March 15, 2026

[Accepted for Publication at Management Science]

Abstract: Making decisions with limited data and incomplete statistical characterization is challenging. The typical statistical-machine-learning approaches would call for migrating the experience of a related system with ample data through *transfer learning* or leveraging the similarity of multiple systems with limited data through *data pooling*. We, instead, develop new solution concepts to learn across related systems by adapting the parametric Operational Data Analytics (ODA) framework, which is known to produce uniformly optimal data-integrated decisions in the corresponding parametric settings, for non-parametric decision-making. We demonstrate, through the application of newsvendor systems, that transfer learning can, indeed, improve decision performance in the focal system by utilizing a model pre-trained with ample data in a related system. However, through the lens of the ODA framework, the best transfer-learning decision falls in a subclass of operational statistics, limiting the ultimate optimality. In contrast, the ODA cross-learning approach utilizes the ample data from the related system to mimic the stochastic environment of the focal system. When the data from the old system are sufficiently large, the cross-learning solutions derived outperform *any* transfer-learning solution, and they are shown to asymptotically approach the parametric ODA solutions. When there are multiple related systems with limited data, we aggregate the data from different systems to create a generic stochastic environment for the decision-making problem, which facilitates the implementation of the parametric ODA solutions. We show that the derived *co-learning* solutions are asymptotically optimal for the aggregate system and for each sub-system. This approach outperforms the existing data-pooling techniques in the sense that the latter focuses only on the aggregated performance, and the chosen solution may be (asymptotically) suboptimal for individual sub-systems. Our results underscore the roles of domain knowledge and the structural relationships between the data and the decision in designing efficient learning solutions with limited data. Though we demonstrate our development through the application of newsvendor systems, the solutions developed in this study apply to a much wider class of operational decision-making problems that exhibit certain homogeneous properties.

Keywords: Operational Data Analytics; Transfer Learning; Cross Learning; Co-Learning; Small Sample

1 Introduction

Data are playing an increasingly important role in today's business. In many cases, a lot of data are needed to train a model to gain the predictive or prescriptive power from the model. However, the accessible training data are often limited to fully realize the value of the model. Data-driven modeling is particularly challenging when there is a lack of sufficient statistical characterization of the uncertainty involved in the system.

Decisions on sourcing, production and distribution for new products are often made when only limited sales data are available. Especially in the fashion industry, the product life cycle can be as short as twelve weeks or as long as thirty weeks (Berg et al. 2018), and it is not possible to collect enough data to train a decision model in such situations. Even for products with long life cycles, market conditions evolve with the macroeconomic and competitive environment, significantly altering the patterns of product consumption. For example, COVID-19 has significantly changed consumer behaviors and shattered the demand forecasts for many products (Brea et al. 2020). The recent increase in online shopping and new environmental regulations are reshaping the demand for paper products, imposing challenges on pulp and paper companies' planning for shifting market trends (Feber et al. 2022). In these situations, the past data may become irrelevant to understand the changed demand patterns. Decisions must be made based on limited recent data.

Though the data for the product, service or market of interest are limited, it is often the case that firms have been operating similar processes either in the past or in parallel. For example, a newly launched product may be an upgraded version of an old product that has been offered for quite some time, and thus some demand patterns of the old product may be expected to carry over to the new product (Hu et al. 2019, Baardman et al. 2018). It is also common that a retailer operates distribution centers and stores across geographic regions (Garvin and Levesque 2008). Product demands across geographic regions may share some similarity (Hitsch et al. 2021, Xu and Bastani 2025). To improve decision efficiency, one may either leverage the experience from some past systems or combine the knowledge across parallel systems. In this paper, we develop learning solutions for such scenarios and demonstrate the results through newsvendor systems.

The premise for learning across systems is the statistical similarity among the data sets collected from different systems. In the context of newsvendor systems, we identify, from the data published by JD.com, that the demands of a product across different distribution centers can be presented by scaling some common random variable, though both the scale parameters and the distribution family are unknown. This observation plays a pivotal role in developing data-driven decisions for the distribution centers.

When significant experience with a related system is available, the idea of *transfer learning* can be adopted, which applies a pre-trained model with established performance for that system to analyze the focal system with limited training data. As the model is already pre-trained, fine-tuning it for the focal system requires much less data and time than training a completely new model. In the context of newsvendor systems, the focal system may be a newly opened distribution center, while a related system can be some center with similar characteristics that has been in operation for a long time. Alternatively, the focal system may be a newly launched product, while a related system can be some substitutable product, targeting the same customer segments, that has been offered for a long time. In such situations, transfer learning would migrate a well-trained data-driven solution from the related system to a decision in the focal system by directly utilizing the fact that the demand data from the two systems differ only in their scale parameters. Our numerical experiments confirm that, indeed, the transfer-learning solutions significantly improve the decision efficiency compared with training some oracle solution for the focal system without utilizing the knowledge from the related system. The shortcoming of transfer learning, however, lies in the fact that the chosen solution is always within a special set of data-integration models (in particular, the *scaled class of operational statistics*; see Section 3.1 for an explanation). As a result, regardless of the sample size, the performance is always capped by the theoretically optimal solution within this class, leading to an optimality gap in general.

To reduce the optimality gap, we note that any data-integrated solution cannot be superior to the parametric ODA (Operational Data Analytics) solution, which is uniformly optimal when the distribution family is known (Feng and Shanthikumar 2023). Therefore, the design of an efficient learning solution should appropriately balance *data integration* and *solution validation*, the two pillars of the ODA framework, based on the structural property of the parametric solution and the data availability. This is the main philosophy behind *cross learning*. Specifically, given the relationship between the data sets, we utilize the ample data from the related system to generate many problem instances that mimic the stochastic environment of the focal system. This approach effectively enhances the power of validating the candidate data-integration models (i.e., the operational statistics). We demonstrate that cross learning can significantly outperform transfer learning. Moreover, as the data of the related system gets ample, the cross-learning performance approaches the best possible data-integrated solution, i.e., that under the corresponding optimal parametric

ODA solution.

In many situations, there may be multiple related systems, for each of which a decision needs to be made with limited data. A typical example is a new-product release in multiple geographic regions. The stocking level in each region must be determined when there are only limited demand data during the initial period. Though different regional markets may exhibit different demand patterns, there can be common features of consumer preference across the regions. Traditionally, *data pooling* is often applied to such a situation, which combines the data from different related systems to utilize the common statistical property of the data sets. Such an approach, though leading to performance improvement over the combined systems, does not guarantee the efficiency of individual systems. We, instead, propose *co-learning*. Exploring the common structure in the data-to-decision mapping of the parametric ODA solutions across systems, we transform the data from different systems to create a common generic environment for the decision-making problem. In an essential contrast to the data pooling approaches, our way of combining data across systems is based on the properties of the operational statistics (i.e., the direct data-to-decision mapping). This approach allows us to efficiently approximate the parametric ODA solutions to achieve the asymptotic optimality for individual systems within the corresponding classes of operational statistics, as the number of involved systems gets large.

The remainder of the paper is organized as follows. The next section presents the related literature and articulates our contribution. Section 3 lays out the problem and provides a brief review of the parametric ODA solution. In Section 4, we discuss several existing solution approaches, and demonstrate the application of transfer learning. In Section 5, we develop several cross-learning solutions based on the ODA framework, which exhibit significant improvement over the transfer-learning solutions. Section 6 concerns multiple parallel systems, for which the ODA-based co-learning solutions are developed. We conclude the study in Section 7.

2 Literature Review

Our study, focusing on decision-making with data supplementing the lack of statistical knowledge, intersects with three streams of literature: data-integrated decision-making in operations, transfer learning, and data pooling.

2.1 Data-Integrated Operational Decision-Making

There is a rapidly growing literature in operations management on decision-making based on data. Our work falls into the stream of static learning, in which a decision needs to be made based on a set of historical data. In the parametric setting where some parameters are unknown for a known distribution family, Hayes (1969) points out that a reasonable data-integrated newsvendor solution would be naturally biased, and proposes some potential solutions. Liyanage and Shanthikumar (2005) define the notion of operational statistics, and derive a uniformly optimal solution when the demand faced by the newsvendor is exponentially distributed. Chu et al. (2008) generalize the result to a general demand distribution with an unknown scale parameter, and Ramamurthy et al. (2012) analyze the situation where there are three unknown parameters of the demand distribution. Akcay et al. (2011) propose to bias the safety-stock factor in the estimation-and-then-optimization solution when the demand comes from the Johnson transformation system (which can be obtained by specific monotone transformations of a standard log-normal distribution). In a similar spirit, Janssen et al. (2009) set the service-level target by correcting the safety-stock factor based on the moment estimation of normal demands. A recent paper by Chu et al. (2025) characterizes the uniformly optimal solution for the price-setting newsvendor problem. The learning approaches developed in our paper, building on the solution derived by Chu et al. (2008) and Chu et al. (2025), do not assume any known distribution family.

In the non-parametric setting, sample average approximations (e.g., Levi et al. 2007a, Huh et al. 2009, Homem-de Mello and Bayraksan 2014, Qin et al. 2022), quantile regressions (e.g., Amrani and Khmelnitsky 2017, Harsha et al. 2021), and order statistics (e.g., Besbes and Mouchtaki 2023) have been analyzed to derive data-integrated inventory decisions. Implementation of these approaches often requires a significant amount of data (Gupta and Rusmevichientong 2021). To account for the potential issue of overfitting, empirical risk minimization (e.g., Ban and R. 2019) and robust optimization (e.g., Lim et al. 2006, Ben-Tal et al. 2013) have been proposed. Gotoh et al. (2018, 2021) suggest that robust optimization with specific deviation measures, including the Kullback-Leibler divergence, χ^2-divergence, and Hellinger distance, produces solutions close to that derived with variance as the regularizer. Feng and Shanthikumar (2023) demonstrate how these solutions can be unified with a non-parametric ODA framework through adaptive boosting. While we also analyze adaptive boosting of an existing solution, the boosting approach in our analysis must leverage the data across different systems.

2.2 Transfer Learning

In the machine learning literature, transfer learning is an important topic. By transferring the knowledge from different but related systems with abundant data, an improvement can be achieved for classification, regression, or clustering performed on the focal system. The reader is referred to Pan and Yang (2009), Weiss et al. (2016), Zhuang et al. (2020) for comprehensive surveys of various methods and applications of transfer learning. Transfer learning is popular in natural language processing (Alyafeai et al. 2020), virtual categorization (Shao et al. 2014), human activity classification (Cook et al. 2013), and software defect classification (Ma et al. 2012).

In the operations literature, transfer learning has been used to enhance the performance of predictive models with limited data. For example, Hao et al. (2021) develop the prediction of COVID-19 transmission over a spatial network by a long short-term memory model and multi-task learning. They utilize the past influenza data in the same geographic regions to estimate the model parameters. Hu et al. (2019) and Baardman et al. (2018) develop new product sales forecasts by matching the new product with some past products. The old product data are used to form clusters and the new product is classified into one of the clusters. Transfer learning is used to reduce training time and avoid overfitting. Oroojlooyjadid et al. (2022) analyze the beer game using deep reinforcement learning. They directly transfer the parameters trained from some layers of a related system to the focal system. Qin et al. (2020) apply transfer learning in the deep reinforcement learning of a ride-sharing platform. Bastani (2021) transfers estimated regression parameters from a proxy task to a focal task with limited data. They propose a two-step estimation with LASSO regularization and show that the estimator can achieve the same level of accuracy as that of the popular heuristics with up to exponentially less data of the focal task.

As a major distinction from the existing studies of transfer learning which focus on predictive modeling, we develop decision-making approaches that utilize data from related systems to supplement the limited data in the focal system. We derive the transfer-learning solution based on the structural property of the decision-making problem. More importantly, we develop the cross-learning solutions that capture the key statistical and structural characteristics of the parametric ODA solution, and show that the proposed cross-learning solutions outperform the transfer-learning solutions in both the large and small sample performance.

2.3 Data Pooling

When there are multiple related systems with limited data, pooling the data across the systems can improve the overall estimation accuracy. The famous James-Stein estimator is a biased (shrunken) mean estimator of a correlated Gaussian vector with the same variance. There are many applications and extensions of the James-Stein approach. For example, Duan and Wang (2023) analyze data pooling in a multi-tasking system with shrinkage. Xu and Bastani (2025) discuss the same problem using LASSO to address data sparsity. In the operations literature, several recent studies have analyzed data pooling in various contexts. Bastani et al. (2022) apply the empirical Bayesian method to study the problem of pricing for multiple products with a common unknown prior distribution of the demand parameters. They show that the proposed meta dynamic pricing algorithm generates a sublinear regret in the number of products. Gupta and Kallus (2022) consider a large number of problems, each with limited data, that need to be solved simultaneously. The data for each problem are i.i.d. draws from some distribution with a finite support.

The domain knowledge is critical for designing data pooling strategies. As a general theme, data pooling is viable when there are statistical similarities among the systems from which the data are pooled (e.g., the common standard deviation in James-Stein's analysis, and a common prior in Gupta and Kallus's study). As an essential departure from the existing data pooling work, which utilizes the statistics development for estimation, our co-learning solution explores the structural properties of the parametric ODA solution that directly links the data to the decision. The data are pooled accordingly to address the lack of the distributional knowledge in the non-parametric environment. More importantly, our ODA co-learning solutions aim toward optimizing not only the efficiency of the overall systems, but also the efficiency of the individual systems.

3 The Problem

In many operating environments, there are multiple related systems involving similar decision-making processes. In each system, a decision $y \in \mathcal{Y} \subset \mathbb{R}_+$ needs to be made and the outcome of system performance is affected by a nonnegative random parameter X with distribution F_X. In most practical situations, the set $\mathcal{Y}$ of feasible decisions is bounded (by, e.g., machine hours, staffing level, storage space, supply contract, technology capability, etc.). The profit generated from the system, when the realization of X is x, is $\psi(y,x)$. The objective is to maximize the expected profit

$$\phi[y, F_X] = \mathsf{E}[\psi(y,X)] = \int_{x\in\mathbb{R}_+} \psi(y,x)\mathrm{d}F_X(x),\ y \in \mathcal{Y}. \tag{1}$$

If we know the distribution F_X, the theoretically optimal decision is

$$y^*[F_X] \in \arg\max\Big\{\phi[y, F_X],\ y \in \mathcal{Y}\Big\}.$$

Such a decision-making problem can arise in many applications. We use newsvendor systems as the running example throughout the paper to demonstrate our development, while we should point out that our analysis directly applies to many other contexts (see the applications described by Feng and Shanthikumar 2023). For the newsvendor application, a system can be a distribution center that carries the product of interest. The decision y corresponds to the inventory level of the product (which is likely to be limited by the shelf space or production capacity), and the random parameter X corresponds to the demand of that product (which is nonnegative). Suppose that the procurement cost is c and the selling price is p with $c < p$, then $\psi(y,x) = p\min\{y,x\} - cy$, and the newsvendor solution is $y^*[F_X] = \bar{F}_X^{\mathrm{inv}}(c/p)$, where $\bar{F}_X^{\mathrm{inv}}$ is the inverse of the survival function $\bar{F}(x) = 1 - F_X(x)$. We note that the newsvendor profit function exhibits the following property.

Assumption 1 (Profit Function)

1. $\psi(\cdot,\cdot)$ is Lipschitz continuous almost everywhere.

2. For some fixed ι and κ, the profit function $\psi(y,x)$ satisfies

$$\psi(\alpha^\iota y, \alpha x) = \alpha^\kappa \psi(y,x), y \in \mathcal{Y}, x \in \mathbb{R}_+,\ \forall \alpha \in \mathbb{R}_+.$$

3. $\psi(y,x), y \in \mathcal{Y}, x \in \mathbb{R}_+$ is bounded.

It is easy to verify that the newsvendor profit $\psi(y,x) = p\min\{y,x\} - cy$ satisfies the second condition with $\iota = \kappa = 1$. This assumption suggests that when the demand and order quantity are both scaled by some constant α, the profit is also scaled by the same constant. Moreover, in most practical applications, including the newsvendor model analyzed here, the feasible decision is generally bounded, and the profit function is bounded over the feasible set.

In reality, we may not know the exact demand distribution. Instead, data are collected to supplement the demand information for decision-making. The development of data-integrated decisions depends on our (partial) knowledge of demand and the kind of data available. Though the data of the decision-making problem at hand may be limited, the firm may have sufficient understanding of its operating environment and enough experiences of managing similar situations. For example, when we examine the data set of JD.com, a Chinese e-commerce platform, we find an interesting feature of the demand data—The coefficient of variation (i.e., the ratio of the standard deviation to the average) of the demand across different distribution centers stays constant; See a demonstration in Figure 1 and additional details in Online Appendix C. In Figure 1, each data point specifies the average and the standard deviation of the daily demand for a product in one distribution center over 31 days. Though the daily demands are different in different distribution centers, the ratio of their sample standard deviation to their sample average stays almost constant. Furthermore, a pair-wise Kolmogorov–Smirnov test suggests that the normalized demands come from the same distribution. These observations suggest that the demand X in every distribution center is of the form $X = \theta Z$ for some common random variable Z and a center-specific parameter θ, so that the coefficient of variation $\mathsf{Cv}[X] = \sqrt{\mathsf{Var}[X]}/\mathsf{E}[X] = \mathsf{Cv}[Z]$ stays the same across distribution centers. Thus, we make the following assumption.

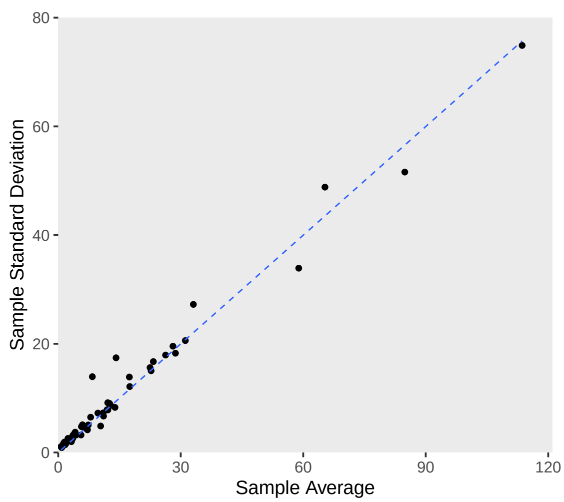

Note. Each point specifies the sample average and sample standard deviation of the daily demands for SKU 068f4481b3 in one of the 60 distribution centers of JD.com over 31 days. The data set is available at `https://connect.informs.org/msom/events/datadriven2020`.

Figure 1: The daily product demands at different distribution centers of JD.com

Assumption 2 (Random Parameter)

1. *The distribution function of X satisfies $F_X(x|\theta) = F_Z(x/\theta), x \in \mathbb{R}_+$, where $\theta \in \mathbb{R}_+$ is the scale parameter and F_Z is the distribution function of some random variable Z with* $\mathsf{E}[Z] = 1$*,*

$\mathsf{E}[Z^2] < \infty$ *and* $\mathsf{E}[Z^\iota] < \infty$.

2. *The random variable* Z *has a density* $f_Z(z) > 0$ *for* $z \in \mathbb{R}_+$.

In practice, most product demands are integers and thus the demand distribution is not continuous. Our assumption of a continuous distribution is a reasonable approximation in the retail application, typically with high demand granularity, in which a distribution center serves a significant customer base so that the expected demand θ cannot be too small.

When it is not viable to compute the theoretically optimal solution $y^*[F_X]$, the approach for data-integrated decision-making depends critically on our knowledge of the operating environment, which defines the domain of validation. In the next subsection, we discuss the scenario where the random variable X is known up to the scale parameter (i.e., with unknown θ and known F_Z). This scenario serves as the benchmark for the evaluation of learning solutions of unknown distribution families.

3.1 Benchmark: The Parametric ODA Solutions

In some situations, firms may have extensive experience with the business, which allows for a confident characterization of the statistical nature of the system. For example, in the context of product selling, Hu et al. (2019) argue that most ready-to-launch products are unlikely to be entirely new. Knowledge of past products is often relevant to infer the characteristics of the new product and make operational decisions (e.g., Kahn 2006, Lin et al. 2022). In such situations, parametric approaches are developed under the premise that the random variables involved in the system come from some known distribution family.

When the random variable described in Assumption 2 is known up to its scale, i.e., F_Z is known and θ is unknown, the uniformly optimal solution, known as the parametric Operational Data Analytics (ODA) solution, can be explicitly computed (Feng and Shanthikumar 2023). In this subsection, we briefly describe this solution, which constitutes a stepping stone to developing the ODA cross-learning and co-learning solutions.

The inputs to the ODA framework are the *data-generation model* and the *domain of validation*. The data $\boldsymbol{X} = (X_1, X_2, \ldots, X_n)$ are generated as i.i.d. draws from some distribution F_X of a focal system. Given that we know F_X up to its scale (i.e., known F_Z and unknown θ), the statistical domain of validation is

$$\mathcal{D}(F_Z) = \{F_X : F_X(x|\theta) = F_Z(x/\theta),\ \forall x \in \mathbb{R}_+, \theta \in \mathbb{R}_+\}.$$

The two pillars of the ODA framework are the *data-integration model* and the *validation model*. To formulate the data-integration model, we note that the ultimate solution y is a statistic, called the *operational statistic*, as it is an implementable decision, i.e., $y : \mathbb{R}_+^n \to \mathbb{R}_+$. In view of the property of the profit function $\psi(y, x)$ in Assumption 1, if the realization x is scaled by a factor α, the optimal decision y must be scaled by a factor of α^ι. Thus, it is natural to focus on the following data-integration model:

$$\mathcal{H}_+^n(\iota, \pi) = \big\{y : \mathbb{R}_+^n \to \mathbb{R}_+; y(\boldsymbol{x}) = y(\boldsymbol{x}_\pi); y(\alpha \boldsymbol{x}) = \alpha^\iota y(\boldsymbol{x}), \alpha \geq 0\big\}, \tag{2}$$

where $\boldsymbol{x}_\pi$ denotes any permutation of $\boldsymbol{x}$, and π denotes the permutation-invariant property. The class $\mathcal{H}_+^n(\iota, \pi)$ contains the family of order-ι positively homogeneous permutation-invariant functions. Feng and Shanthikumar (2023) prove that, for problems satisfying Assumptions 1 and 2, there does *not* exist any solution that is uniformly better than the optimal operational statistic within $\mathcal{H}_+^n(\iota, \pi)$.

When implementing an operational statistic $y : \mathbb{R}_+^n \to \mathbb{R}_+$, we obtain an expected profit of

$$\mathsf{E}\big[\phi[y(\boldsymbol{X}), F_X(\cdot|\theta)]\big]. \tag{3}$$

Note that the expectation is taken over the random sample $\boldsymbol{X}$. The optimal operational statistic within the class of order-ι positively homogeneous permutation-invariant functions (referred to as the *homogeneous class* thereafter) is

$$y^*_{\mathrm{H}\iota}(\cdot) \quad = \quad \arg\max \big\{\mathsf{E}[\phi\big[y(\boldsymbol{X}), F_X(\cdot|\theta)]\big] : y \in \mathcal{H}_+^n(\iota, \pi)\big\}. \tag{4}$$

Thus, the *validation model* optimizes the decision as a statistic (i.e., a function of the data). The solution $y^*_{\mathrm{H}\iota}(\cdot)$ is uniformly optimal when the profit function satisfies Assumption 1. Chu et al. (2008, 2025) explicitly derive the solution for the newsvendor problem (with $\iota = 1$).

To gain some intuition of the ODA framework, we observe that for any $y \in \mathcal{H}_+^n(\iota, \pi)$, the expected profit satisfies

$$\mathsf{E}\big[\phi[y(\boldsymbol{X}), F_X(\cdot|\theta)]\big] = \theta^\kappa \mathsf{E}[\psi(y(\theta^{-1}\boldsymbol{X}), \theta^{-1}X)] = \theta^\kappa \mathsf{E}[\psi(y(\boldsymbol{Z}), Z)]. \tag{5}$$

In other words, the profit obtained from implementing an operational statistic from the data-integration model $\mathcal{H}_+^n(\iota, \pi)$ depends on the unknown parameter θ only through a scale factor, θ^κ. Let $\hat{\mu}_{\boldsymbol{z}}$ denote the average of vector $\boldsymbol{z}$, i.e., $\hat{\mu}_{\boldsymbol{z}} = \frac{1}{n}\sum_{i=1}^n \boldsymbol{z}$. We further define the *base set* as

$$\mathcal{B}^n = \{\boldsymbol{z}_B \in \mathbb{R}_+^n : \hat{\mu}_{\boldsymbol{z}_B} = 1\},$$

which contains all base points (i.e., samples with average one), and define the set of samples that can be obtained by scaling a given base point $\boldsymbol{z}_B \in \mathcal{B}^n$ as $\mathcal{S}(\boldsymbol{z}_B) = \{\boldsymbol{x} : \boldsymbol{x} = \alpha \boldsymbol{z}_B, \alpha > 0\}$. Then the collection of sets $\{\mathcal{S}(\boldsymbol{z}_B), \boldsymbol{z}_B \in \mathcal{B}^n\}$ forms a partition of the sample space $\mathbb{R}_+^n$. We can derive (see the details in Feng and Shanthikumar 2023)

$$\mathsf{E}[\phi[y(\boldsymbol{X}), F_X(\cdot|\theta)]] \quad = \quad \theta^\kappa \int_{\boldsymbol{z}_B \in \mathcal{B}^n} \phi_B[y(\boldsymbol{z}_B), \boldsymbol{z}_B, F_Z] \mathrm{d}\boldsymbol{z}_B, \tag{6}$$

where

$$\phi_B[y, \boldsymbol{z}_B, F_Z] = \int_{\boldsymbol{x} \in S(\boldsymbol{z}_B)} \phi[(\hat{\mu}_{\boldsymbol{x}})^\iota y, F_Z] \prod_{i=1}^n f_Z(x_i) \mathrm{d}\boldsymbol{x}. \tag{7}$$

Thus, finding the optimal operational statistic $y^*_{\mathrm{H}\iota}$ in (4) boils down to computing the base operational statistic

$$y_B^{\mathrm{OS}}(\boldsymbol{z}_B) = \arg\max \left\{\phi_B[y, \boldsymbol{z}_B, F_Z] : y \in \mathbb{R}_+\right\}, \; \boldsymbol{z}_B \in \mathcal{B}^n, \tag{8}$$

and the optimal operational statistic can be mapped from the base operational statistic through appropriate scaling:

$$y^*_{\mathrm{H}\iota}(\boldsymbol{x}) = (\hat{\mu}_{\boldsymbol{x}})^\iota y_B^{\mathrm{OS}}(\boldsymbol{z}_B), \boldsymbol{x} \in \mathcal{S}(\boldsymbol{z}_B). \tag{9}$$

It is important to recognize that the optimal base operational statistic y_B^{OS} is independent of the unknown scale parameter θ. In other words, systems involving a common F_Z but different scale parameters should make the same decisions for any sample falling in the base set $\mathcal{B}^n$. This observation would be essential to design the learning solutions later.

For the purpose of comparison, Feng and Shanthikumar (2023) also propose the *scaled class* of operational statistics as an alternative data-integration model:

$$\mathcal{SC}_+^n(\iota, \pi) = \left\{y : \mathbb{R}_+^n \to \mathbb{R}_+; y(\boldsymbol{x}) = (\hat{\mu}_{\boldsymbol{x}})^\iota \gamma, \gamma \geq 0\right\}. \tag{10}$$

It is easy to see that $\mathcal{SC}_+^n(\iota, \pi) \subset \mathcal{H}_+^n(\iota, \pi)$. In other words, the optimal scaled operational statistic is in general inferior to that of the homogeneous class. The validation of the scaled family leads to the optimal scaled operational statistic:

$$y^*_{\mathrm{SC}\iota}(\cdot) \quad = \quad \arg\max \left\{\mathsf{E}\left[\phi[y(\boldsymbol{X}), F_X(\cdot|\theta)]\right] : y \in \mathcal{SC}_+^n(\iota, \pi)\right\}. \tag{11}$$

Though the true value of θ is unknown, Feng and Shanthikumar (2023) show that the uniformly optimal operational statistics for both the scaled class and the homogeneous class can be derived explicitly as described in the next theorem.

Theorem 1 (The Parametric ODA Solutions) *Suppose F_Z is known.*

i) The operational statistic $y^*_{\mathrm{H}\iota}(\boldsymbol{x})$ maximizes

$$\phi_{\mathrm{H}\iota}(y, \boldsymbol{x}) = \int_{\hat{x}=0}^{\infty} \psi(y, \hat{x}) \int_{\eta=0}^{\infty} f_X(\hat{x}|\eta) \frac{\eta^{-\iota-1} \prod_{i=1}^{n} f_X(x_i|\eta)}{\int_{\alpha=0}^{\infty} \alpha^{-\iota-1} \prod_{i=1}^{n} f_X(x_i|\alpha) \mathrm{d}\alpha} \mathrm{d}\eta \mathrm{d}\hat{x}$$

over $y \in \mathcal{Y}$.

ii) The operational statistic $y^*_{\mathrm{SC}\iota}(\boldsymbol{x})$ satisfies $y^*_{\mathrm{SC}\iota}(\boldsymbol{x}) = (\hat{\mu}_{\boldsymbol{x}})^{\iota} y_{\mathrm{SC}\iota}[F_Z]$, where

$$y_{\mathrm{SC}\iota}[F_Z] = \arg\max \left\{ \int_{\boldsymbol{z} \in \mathbb{R}_+^n} \int_{\hat{z} \in \mathbb{R}_+} \psi\big((\hat{\mu}_{\boldsymbol{z}})^{\iota}\gamma, \hat{z}\big) f_Z(\hat{z}) \prod_{i=1}^{n} f_Z(z_i) \mathrm{d}\hat{z} \mathrm{d}\boldsymbol{z} : \gamma \geq 0 \right\}. \tag{12}$$

In Online Appendix B, we provide explicit derivations of the optimal operational statistics for several known distribution families of F_X. It is interesting to note that $y^*_{\mathrm{H}\iota}(\boldsymbol{x}) = y^*_{\mathrm{SC}\iota}(\boldsymbol{x})$ when F_Z follows a gamma distribution.

It is important to remark that the above optimal operational statistics can be derived by directly solving the validation models in (4) and (11) for their respective data-integration models because of the knowledge of F_Z. Feng and Shanthikumar (2023) suggest that, in view of the Bayesian interpretation of the parametric ODA framework, the solution can be efficiently computed with a simulation algorithm. When F_Z is unknown, direct optimization of (4) or (11) is not possible. To derive an appropriate solution, one needs to modify the data-integration model based on the available knowledge and formulate the corresponding validating model to approximate the validation model. The solution performance must be bounded from above by that of $y^*_{\mathrm{H}\iota}$ or $y^*_{\mathrm{SC}\iota}$ within the respective classes. Thus, $y^*_{\mathrm{H}\iota}(\boldsymbol{x})$ and $y^*_{\mathrm{SC}\iota}(\boldsymbol{x})$ are natural benchmarks to evaluate the quality of the learning solutions in the non-parametric setting.

3.2 Learning Among Related Systems

When we do not have much knowledge of the focal system, the statistical domain of validation becomes

$$\mathcal{D}(\mathcal{F}_Z) = \{F_X : F_X(x|\theta) = F_Z(x/\theta), \forall x \in \mathbb{R}_+, \theta \in \mathbb{R}_+, F_Z \in \mathcal{F}_Z\}, \tag{13}$$

where $\mathcal{F}_Z$ is the set of all distribution functions with mean one. Certainly, $\mathcal{D}(F_Z) \subset \mathcal{D}(\mathcal{F}_Z)$, and the lack of knowledge of F_Z imposes significant challenges to decision-making. Inferring the distributional characteristics of F_Z requires a large sample, which is often not available.

In reality, the firm may be operating similar systems in the past or at the same time. Each system has its own statistical environment, while all the systems are related because of the common structure $X = \theta Z$. When the data collected from one system are limited, we can leverage the data from other systems to improve the decision performance. We consider three learning approaches that are applicable to different situations depending on the knowledge, the experience, and the data availability from relevant systems.

Transfer Learning. Transfer learning is a well-studied approach in the machine learning literature. It is used to improve the performance of one system by transferring a well-trained model of a similar (old) system from which significant experience has already been gained. In particular, for our problem, we take an oracle solution of the old system from which we have collected ample data. We recognize that the old system with θ^o and the focal system with θ both satisfy Assumptions 1 and 2. Thus, the theoretically optimal solutions for θ^o and θ are related through a ratio of $(\theta^o/\theta)^\iota$. Consequently, we can transfer the oracle solution from the old system to the focal system through appropriate scaling, as described in Section 4.

Cross Learning. Unlike transfer learning, which directly applies a pre-trained model of the old system to the focal system, cross learning attempts to create the stochastic environment of the focal system using the data collected from the old system. Specifically, we generate, using the data from the old system, many instances of problems that mimic the decision-making environment of the focal system. These problem instances reflect the statistical nature of the focal system, and help enrich the knowledge of the focal system in decision-making. With the simulated problem instances, we can expand the data-integration model beyond the scaled class, resulting in improved learning performance from transfer learning.

Co-Learning. When there are multiple systems operating under similar environments, data collected from parallel systems may be pooled to assist decision-making and improve the overall performance. The premise of co-learning is the similarity among the systems, which suggests a common structure of data-to-decision mapping among different systems. As we discuss in Section 6, with appropriate aggregation of the data from different systems, it is possible to validate the data-integration model for the generic system that captures the stochastic environment of the decision-making problems in individual systems. The co-learning solution enhances the performance of not only the aggregate system but also the individual systems.

Next, we discuss in detail each of the learning solutions.

4 Transfer Learning

In this section, we demonstrate how to implement the well-studied transfer-learning approach for our decision-making problem. When we do not know F_Z and θ to statistically characterize the random parameter $X =^d \theta Z$, a sufficiently large sample is needed to train an efficient solution for the focal system. When the data $\boldsymbol{X} = (X_1, X_2, \ldots, X_n)$ collected are limited (i.e., n is small), however, it is unlikely to develop a prevailing solution with only the available information of the focal system.

Instead, we may have some experience in operating a similar system, which we label as the "old" system. The old system is similar to the focal one in the sense that the random parameter involved takes the form $X^o = \theta^o Z$. For the old system, historical data $\boldsymbol{X}^o = \{X_1^o, X_2^o, \ldots, X_{n^o}^o\}$ is available with a sufficiently large n^o, i.e., $n^o >> n$. Moreover, a well-trained oracle solution $y^o(\boldsymbol{x}^o)$ has been implemented.

Assumption 3 *The oracle solution $y^o \in \mathcal{H}_+^{n^o}(\iota, \pi)$ is consistent and bounded, i.e.,*

$$\lim_{n^o \to \infty} y^o(\boldsymbol{X}^o) = y^*[F_{X^o}] \text{ a.s. and } y^o(\boldsymbol{X}^o) < \infty \text{ a.s.}$$

The task here is to transfer the experience of the old system to develop a solution for the focal system with limited data of the latter.

4.1 Potential Oracle Solutions

There are several data-integrated approaches that have been proposed for the newsvendor problem, which may be used as oracle solutions. In this subsection, we briefly describe them. The solutions presented below are derived for the problem defined in (1) with the validation domain defined in (13) and i.i.d. observations $\boldsymbol{X} = \{X_1, X_2, \ldots, X_n\}$.

- **Sample average approximation:** The expected profit $\phi[y, F_X]$ is approximated by the average profit under each observation, i.e., $\varphi(y, \boldsymbol{X}) = \frac{1}{n}\sum_{i=1}^{n} \psi(y, X_i)$. The sample average approximation solution is

$$y_{\text{SAA}}(\boldsymbol{X}) = \arg\max \left\{ \varphi(y, \boldsymbol{X}), y \in \mathbb{R}_+ \right\}.$$

This solution also coincides with the retrospectively optimal solution or the optimal solution under empirical distribution under Assumptions 1 and 2; See Feng and Shanthikumar (2023).

- **Regularized sample average:** Accounting for the potential risk of overfitting, we may choose to regularize the sample average profit by its variability, measured by either the variance or standard deviation (see, e.g., Levi et al. 2007b, Huh and Janakiraman 2008, Homem-de Mello and Bayraksan 2014, Qin et al. 2022). Specifically, the variance of the sample average profit is $\mathsf{Var}[\varphi(y,\boldsymbol{X})] = \frac{1}{n}\sum_{i=1}^{n}\big(\psi(y,X_i)\big)^2 - \big(\varphi(y,\boldsymbol{X})\big)^2$. With a penalty $\beta > 0$, the regularized solutions are

$$
\begin{aligned}
y_{\text{Re-Var}}(\boldsymbol{X},\beta) &= \arg\max\big\{\varphi(y,\boldsymbol{X}) - \beta\mathsf{Var}[\varphi(y,\boldsymbol{X})] : y \geq 0\big\},\\
y_{\text{Re-Std}}(\boldsymbol{X},\beta) &= \arg\max\big\{\varphi(y,\boldsymbol{X}) - \beta\sqrt{\mathsf{Var}[\varphi(y,\boldsymbol{X})]} : y \geq 0\big\}.
\end{aligned}
$$

  The parameters $\beta_{\text{Re-Var}}(\boldsymbol{X})$ and $\beta_{\text{Re-Std}}(\boldsymbol{X})$ are obtained through cross validation or bootstrapping.

- **Robust optimization:** An alternative to regularizing the objective function is to specify an uncertainty set of distribution functions based on the observed data (e.g., Gilboa and Schmeidler 1989, Lim et al. 2006, Zhao et al. 2022). Specifically, an h-divergence measure is defined as $d_h(f,\hat{f}_{X|\boldsymbol{X}}) = \int_x h(\frac{f(x)}{\hat{f}_{X|\boldsymbol{X}}(x)})\hat{f}_{X|\boldsymbol{X}}(x)\mathrm{d}x$, where $\hat{f}_{X|\boldsymbol{X}}(x) = \frac{1}{n}\sum_{i=1}^{n}\mathbb{I}_{x=X_i}$ is the empirical distribution and h is the divergence function (e.g., Kullback-Leibler divergence, χ^2-divergence, and Hellinger distance). The uncertainty set should contain distributions close enough to the empirical distribution, i.e., $\mathcal{F}_{\boldsymbol{X},h}(d_{\max}) = \{f : d_h(f,\hat{f}_{X|\boldsymbol{X}}) \leq d_{\max}\}$ for some threshold $d_{\max} > 0$. One may solve for the constrained solution or penalized solution:

$$
\begin{aligned}
y_{\text{Ro-C}}(\boldsymbol{X},\beta) &= \arg\max\Big\{\min\Big\{\sum_{i=1}^{n}\psi(y,x_i)f(x_i) : f \in \mathcal{F}_{\boldsymbol{X},h}(\beta)\Big\}, y \geq 0\Big\},\\
y_{\text{Ro-P}}(\boldsymbol{X},\beta) &= \arg\max\Big\{\min\Big\{\sum_{i=1}^{n}\psi(y,x_i)f(x_i) - \beta d_h(f,\hat{f}_{X|\boldsymbol{X}}) : f \in \mathcal{F}_{\boldsymbol{X},h}(d_{\max})\Big\}, y \geq 0\Big\}.
\end{aligned}
$$

  The parameters $\beta_{\text{Ro-C}}(\boldsymbol{X})$ and $\beta_{\text{Ro-P}}(\boldsymbol{X})$ are obtained through cross validation or bootstrapping. Empirically, we find the robust solutions very similar to the regularized solutions. Sim et al. (2025) and Long et al. (2023) develop a closely related approach, called robust satisficing, for data-integrated decision-making. Feng and Shanthikumar (2023) show that these two approaches are equivalent under mild conditions through the lens of the ODA framework.

- **Order statistics:** Let $\boldsymbol{X}_{[]}$ be the ascending sequence of $\boldsymbol{X}$ (i.e., $X_{[1]} \leq X_{[2]} \leq \ldots \leq X_{[n]}$). Recognizing that the SAA solution corresponds to the $\lfloor nc/p \rfloor$th or $\lceil nc/p \rceil$th order statistic, depending on which one gives a higher sample average profit, one can look for an operational statistic by taking a vector of weights $\boldsymbol{w}$ (with sum one) over the order statistics $\boldsymbol{X}_{[]}$, i.e., $y(\boldsymbol{w}, \boldsymbol{X}) = \sum_{i=1}^{n} w_i X_{[i]}$ (see Liyanage and Shanthikumar 2005). Besbes and Mouchtaki (2023) choose the weight vector $\boldsymbol{w}$ through robust optimization, which generates the eventual solution $y_{\mathrm{OrS}}(\boldsymbol{x})$. Alternatively, one may cross validate the weights.

- **ODA based on empirical distribution:** We normalize the sample data by its average to obtain $\boldsymbol{Z} = \boldsymbol{X}/\hat{\mu}_{\boldsymbol{X}}$. The distribution of Z is estimated using the smoothed empirical density of $\boldsymbol{Z}$:

$$\tilde{f}_{Z|\boldsymbol{Z}}(\xi) = \frac{1}{n\lambda} \sum_{i=1}^{n} \kappa_\lambda(\xi, Z_i), \tag{14}$$

  where $\lambda > 0$ is the width and κ_λ is the Kernel function. For example, $\kappa_\lambda(\xi, z) = \frac{1}{\sqrt{2\pi}} e^{\frac{(\xi - z)^2}{2\lambda^2}}$, with $\lim_{n\to\infty} \lambda = 0$, is the widely used Gaussian kernel. Using $\tilde{f}_{Z|\boldsymbol{Z}}$ to replace f_Z in Theorem 1, we can derive the ODA solutions $y_{\mathrm{H}\iota\text{-Emp}}$ and $y_{\mathrm{SC}\iota\text{-Emp}}$.

It is important to note that all of the above solutions belong to the homogeneous class $\mathcal{H}_+^n(\iota, \pi)$ (see Remark 2 in Feng and Shanthikumar 2023), provided that the h-divergence measure employed in the robust optimization formulation is itself homogeneous (e.g., the χ^2-divergence or the Kullback–Leibler divergence). In fact, a nonhomogeneous solution may be questionable in practice. In our newsvendor example (with $\iota = 1$), the optimal ordering decision should remain unchanged when demand data are recorded in different measurement units (e.g., by item or by carton). This invariance implies that any reasonable order quantity must be homogeneous of degree one.

4.2 Transfer Learning from the Oracle Solution of the Old System

We would like to transfer a well-trained, consistent oracle solution $y^o(\boldsymbol{X}^o)$ from the old system to the focal system with limited data $\boldsymbol{X}$. To do so, we recognize that the data from both systems are connected through the structure described in Assumption 2. We first normalize the data from the old system as $\boldsymbol{Z}^o = \boldsymbol{X}^o/\hat{\mu}_{\boldsymbol{X}^o}$. When the sample size n^o is sufficiently large, Z_i^o converges to the underlying random variable Z, as suggested in the next lemma.

Lemma 1 *Suppose $X^o =^d \theta^o Z$ with $\mathsf{E}[Z] = 1$ and $0 < \theta^o < \infty$. Then (superscript d stands for "in distribution")*

$$\lim_{n^o \to \infty} \frac{X_i^o}{\hat{\mu}_{\boldsymbol{X}^o}} =^d Z.$$

Because the oracle solution $y^o(\boldsymbol{X}^o)$ belongs to $\mathcal{H}_+^n(\iota,\pi)$, and both $\boldsymbol{X}$ and $\boldsymbol{X}^o$ satisfy Assumption 2, we can apply (9) to transfer $y^o(\boldsymbol{X}^o)$ to the decision in the focal system as

$$y_{\text{TrL:o}}(\boldsymbol{X}) = \frac{(\hat{\mu}_{\boldsymbol{X}})^\iota}{(\hat{\mu}_{\boldsymbol{X}^o})^\iota} y^o(\boldsymbol{X}^o) \in \mathcal{SC}_+^n(\iota,\pi). \tag{15}$$

Thus, the transfer-learning solution is dominated by the optimal scaled operational statistic defined in (11) with a known F_Z, i.e.,

$$\mathsf{E}[\psi(y_{\text{TrL:o}}(\boldsymbol{X}), X)] \le \mathsf{E}[\psi(y^*_{SC\iota}(\boldsymbol{X}), X)]. \tag{16}$$

Unfortunately, $y_{\text{TrL:o}}$ does not converge to the best scaled operational statistic $y^*_{\text{SC}\iota}(\boldsymbol{X}) = (\hat{\mu}_{\boldsymbol{X}})^\iota y_{\text{SC}}[F_Z]$ characterized in Theorem 1(ii).

Theorem 2 (Consistency of Transfer-Learning Solution) *Given a consistent solution y^o of the old system (i.e., $\lim_{n^o\to\infty} y^o(\boldsymbol{X}^o) \to y^*[F_{X^o}]$, a.s.), we have*

$$\lim_{n^o\to\infty} y_{\text{TrL:o}}(\boldsymbol{X}) =^d (\hat{\mu}_{\boldsymbol{X}})^\iota y^*[F_Z] \quad \textit{and} \quad \lim_{n^o,n\to\infty} y_{\text{TrL:o}}(\boldsymbol{X}) =^d y^*[F_X]. \tag{17}$$

In addition, if $\Pr\left\{\left|y^o(\boldsymbol{X}^o) - y^*[F_{X^o}]\right| \ge \epsilon\right\} \le \mathcal{O}\big(g^o(n^o,\epsilon)\big)$ for some function $g^o(n^o,\epsilon)$ (e.g., $g^o(n^o,\epsilon) = 1/(n^o\epsilon^2)$ for y^{SAA}) and for any $\epsilon > 0$, then

$$\Pr\left\{\left|y_{\text{TrL:o}}(\boldsymbol{X}) - (\hat{\mu}_{\boldsymbol{X}})^\iota y^*[F_Z]\right| \ge \epsilon\right\} \le \mathcal{O}\Big(g^o(n^o,\epsilon)\Big) + \mathcal{O}\Big(\frac{1}{n^o\epsilon^2}\Big), \ \textit{and} \tag{18}$$

$$\Pr\left\{\left|y_{\text{TrL:o}}(\boldsymbol{X}) - y^*[F_X]\right| \ge \epsilon\right\} \le \mathcal{O}\Big(g^o(n^o,\epsilon)\Big) + \mathcal{O}\Big(\frac{1}{n^o\epsilon^2}\Big) + \mathcal{O}\Big(\frac{1}{n\epsilon^2}\Big). \tag{19}$$

Table 1 illustrates the performance of the transfer-learning solutions against training the corresponding oracle solutions in the focal system, where Ψ_k denotes the simulated expected profit and the subscript k denotes the solution method. Here we choose Gamma distribution because, by Theorem 1, the optimal homogeneous operational statistic coincides with the optimal scaled operational statistic (i.e., $\Psi_{\text{H1}} = \Psi_{\text{SC1}}$) in this case, which allows for clear insights with a single benchmark. Certainly, any solution without the knowledge of F_Z would lead to a profit less than the benchmark profit, Ψ_{H1}. Because there is a significant variability in the demand reflected by a high coefficient of variation ($\mathsf{Cv}[Z] = 3.16$), the oracle solution performs poorly with only five observations from the focal system. Transfer learning, migrating the oracle solution from the old system, can greatly improve the profit, and the improvement becomes more significant when the sample size n^o in the old system gets larger. Similar observations are obtained when Z follows a beta distribution as shown in Online Appendix D.

5 Cross Learning

When we transfer the oracle solution in the previous section, the performance is capped by that under the optimal scaled operational statistic $y^*_{\mathrm{SC}\iota}$; recall Equation (16). Unless in special cases (i.e., F_Z follows a gamma distribution), the optimal homogeneous operational statistic $y^*_{\mathrm{H}\iota}$ is superior to $y^*_{\mathrm{SC}\iota}$. This suggests room for improvement over the transfer-learning solutions.

We note that the transfer-learning solution, though leveraging some well-trained solution from the old system, does not fully explore the information contained in the historical data. As a result, a significant risk remains in transferring the oracle solution due to the limited data of the focal system. To address this shortcoming, we, instead, utilize the ample data from the old system to create the stochastic environment for the decision-making problem in the focal system. Once such environment is created, we can utilize the structure of data-to-decision mapping exhibited in the parametric ODA solution to derive an efficient solution of the focal system.

Specifically, based on our discussion in Section 3.1, the problem of the focal system with data $\boldsymbol{X}$ boils down to deriving a statistic of the normalized data $\boldsymbol{Z}_B = \boldsymbol{X}/\hat{\mu}_{\boldsymbol{X}}$ (which is a base point). In the parametric setting (where F_Z is known), the objective to derive the base operational statistic is $\phi_B[y, \boldsymbol{z}_B, F_Z]$ defined in Equation (7). It is important to note that this profit function is independent of the unknown parameter θ, and thus the optimal base operational statistic does not depend on θ. In other words, the optimal base operational statistic $y^*_{\mathrm{H}\iota}(\boldsymbol{z}_B)$ is the same for the old system and the focal system. Therefore, we can seek an efficient solution of the base point $\boldsymbol{z}_B$ in a generic system with random parameter Z, instead of X.

We use the data from the old system to create the generic system. When the sample size gets large, the normalized data $\boldsymbol{Z}^o = \boldsymbol{X}^o/\hat{\mu}_{\boldsymbol{X}^o}$ reflect the distribution of Z very well; Recall Lemma 1. From Equations (5) and (6), we have, for any decision $y \in \mathcal{H}^n_+(\iota, \pi)$,

$$\mathsf{E}[\phi[y(\boldsymbol{Z}), F_Z]] = \int_{\boldsymbol{z}_B \in \mathcal{B}^n} \phi_B[y(\boldsymbol{z}_B), \boldsymbol{z}_B, F_Z] \mathrm{d}\boldsymbol{z}_B,$$

so that $y(\boldsymbol{z}) = (\hat{\mu}_{\boldsymbol{z}})^\iota y(\boldsymbol{z}_B)$ for $\boldsymbol{z} = \mathcal{S}(\boldsymbol{z}_B)$ and $\boldsymbol{z}_B \in \mathcal{B}^n$. We need to estimate the profit $\phi_B[y(\boldsymbol{z}_B), \boldsymbol{z}_B, F_Z]$ defined in Equation (7) using the simulated data.

To fully utilize the data collected from the old system, we create many problem instances from $\boldsymbol{Z}^o$ that closely mimic the decision-making problem in the focal system, instead of directly replacing F_Z by $\boldsymbol{Z}^o$ to estimate the corresponding expected profit. Specifically, we randomly generate samples of size n based on the empirical density $\tilde{f}_{Z|\boldsymbol{Z}^o}$ in (14):

$$\boldsymbol{Z}^{(j)} = (Z_1^{(j)}, Z_2^{(j)}, \ldots, Z_n^{(j)}),\ j = 1, 2, \ldots, m,$$

with a large m. Denote $\mathbf{Z}^{()} = (\boldsymbol{Z}^{(j)} : j = 1, 2, \ldots, m)$, $\boldsymbol{Z}_B^{(j)} = \boldsymbol{Z}^{(j)}/\hat{\mu}_{\boldsymbol{Z}^{(j)}}$ and $\mathbf{Z}_B^{()} = (\boldsymbol{Z}_B^{(j)}, j = 1, 2, \ldots, m)$. Thus, each simulated sample $\boldsymbol{Z}^{(j)}$ with n i.i.d. observations gives an instance of the generic system whose random parameter follows the density $\tilde{f}_{Z|\boldsymbol{Z}^o}$. The problem of the generic system is to make a decision with a sample size of n, the situation faced in the focal system. If we can derive an efficient base operational statistic for the generic system, we obtain the operational statistic of the focal system in view of Equation (9). We will discuss several ways of cross learning using the simulated generic systems.

5.1 Cross Learning for Optimal Scaled Operational Statistics

We first develop a cross-learning solution optimized over the scaled class $\mathcal{SC}_+^n(\iota, \pi)$. This solution directly contrasts the transfer-learning solution, a particular element in $\mathcal{SC}_+^n(\iota, \pi)$. With our construction of generic systems to capture the stochastic environment for the decision-making, we can derive the optimal solution for the validating model in (20) within $\mathcal{SC}_+^n(\iota, \pi)$, instead of scaling a given oracle solution. Thus, the cross-learned scaled operational statistics always dominate *any* transfer-learning solution with sufficiently large n^o.

Note from the parametric ODA solution $y^*_{\mathrm{SC}\iota}(\boldsymbol{X}) = (\hat{\mu}_{\boldsymbol{X}})^\iota y_{\mathrm{SC}\iota}[F_Z]$, $y_{\mathrm{SC}\iota}[F_Z]$ is independent of the data $\boldsymbol{X}$ and it corresponds to the operational statistic for *any* base point. Thus, for the focal system, we look for an operational statistic of the form $y(\boldsymbol{X}) = (\hat{\mu}_{\boldsymbol{X}})^\iota \gamma$, where γ is the surrogate for $y_{\mathrm{SC}\iota}[F_Z]$ and is the decision for all base points (including $\boldsymbol{z}_B$). Because the solution is the same for samples with the same average, we can approximate the expected profit $\mathsf{E}[\phi(y(\boldsymbol{Z}), F_Z)]$ by taking the sample average over $\mathbf{Z}^{()}$, and formulate the validating model as

$$\hat{\phi}_{\mathrm{CrL:SC}\iota}(\gamma, \mathbf{Z}^{()}) = \frac{1}{m(m-1)n} \sum_{j=1}^{m} \sum_{\substack{i \in \{1,2,\ldots,n\} \\ \ell \in \{1,2,\ldots,m\} \setminus j}} \psi((\hat{\mu}_{\boldsymbol{Z}^{(j)}})^\iota \gamma, Z_i^{(\ell)}). \tag{20}$$

In computing $\hat{\phi}_{\mathrm{CrL:SC}\iota}$, we use each sample $\boldsymbol{Z}^{(j)}$ in the first argument of ψ to mimic a scenario of implementing the scaled operational statistic for a sample of size n. All other simulated data in $\mathbf{Z}^{()}$ are used to represent the random parameter Z in the second argument of ψ. Based on this estimated profit, we can derive the cross-learning solution within the class of scaled operational statistics as

$$y_{\mathrm{CrL:SC}\iota}(\boldsymbol{X}) = (\hat{\mu}_{\boldsymbol{X}})^\iota \gamma_{\mathrm{CrL:SC}\iota},$$

where

$$\gamma_{\mathrm{CrL:SC}\iota} = \arg\max\{\hat{\phi}_{\mathrm{CrL:SC}\iota}(\gamma, \mathbf{Z}^{()}) : \gamma \in \mathbb{R}^+\}.$$

It is important to note that $\gamma_{\text{CrL:SC}\iota}$ depends on the sample size n. Thus, the cross-learning solution accounts for the variability across random samples of the same size n.

We may want to interpret the scaled operational statistic $y_{\text{CrL:SC}\iota}$ as a boosted solution, corresponding to a constant-boosting parameter $\gamma(\boldsymbol{x}) = \gamma, \forall \boldsymbol{x} \in \mathbb{R}^n_+$ (see the discussion in Section 5.2). Specifically, consider a naive candidate solution that sets $y^c(\boldsymbol{x}) = \hat{\mu}_{\boldsymbol{x}}$. By applying a boosting constant γ, we obtain the scaled operational statistic. However, this candidate solution is generally inconsistent, and boosting an inconsistent estimator does not necessarily improve performance, even with large sample sizes. In particular, in the newsvendor model, the optimal order quantity is not the mean demand in general. Interestingly, the scaled solution $y_{\text{CrL:SC}\iota}$ is consistent, as established in the following theorem. By Assumption 1, there exist a, b such that $a \le \psi(y, x) \le b, y \in \mathcal{Y}, x \in \mathcal{X}$.

Theorem 3 (Consistency of Cross-Learned Scaled Statistics) *Under Assumptions 1 and 2, if $\Gamma \subset \mathbb{R}_+$ is compact and $m \propto n^o$, then for any $\epsilon > 0$ and a sufficiently large n^o,*

$$\Pr\left\{ \sup_{\gamma \in \Gamma} \left| \hat{\phi}_{\text{CrL:SC}\iota}(\gamma, \mathbf{Z}^{()}) - \mathsf{E}[\psi((\hat{\mu}_{\boldsymbol{Z}})^{\iota}\gamma, Z)] \right| \ge \epsilon \right\} \le \mathcal{O}\Big(\exp\Big(- \frac{2n^o\epsilon^2 - 4C\sqrt{n^o}\epsilon}{(b-a)^2} \Big) \Big), \tag{21}$$

where $C > 0$. In addition, if $y_{\text{SC}\iota}[F_Z]$ defined in Equation (12) *is unique on Γ, then*

$$y_{\text{CrL:SC}\iota}(\boldsymbol{x}) \xrightarrow{p} y^*_{\text{SC}\iota}(\boldsymbol{x}),\ \forall \boldsymbol{x} \in \mathbb{R}^n_+,\ \textit{as } n^o \to \infty. \tag{22}$$

Furthermore, if $y_{\text{SC}\iota}[F_Z]$ is an unique and interior maximizer on Γ, $\psi(\cdot, x)$ is continuously differentiable with Lipschitz continuity and bounded $\nabla\psi(\cdot, x)$ over Γ a.s., and $\mathsf{E}[\psi(\cdot, X)]$ is strongly concave, then for any $\epsilon > 0$ and a sufficiently large n^o,

$$\Pr\left\{ \left| y_{\text{CrL:SC}\iota}(\boldsymbol{x}) - y^*_{\text{SC}\iota}(\boldsymbol{x}) \right| \ge \epsilon \right\} \le \mathcal{O}\Big(\exp\Big(- \frac{2n^o\epsilon^2 - 4C\sqrt{n^o}\epsilon}{(b-a)^2} \Big) \Big). \tag{23}$$

Theorem 3 suggests that the empirical profit $\hat{\phi}_{\text{CrL:SC}\iota}$ converges uniformly in probability to the actual expected profit obtained from implementing any scaled operational statistic, as the sample size of the old data gets large. We shall note that the convergence is in probability as the empirical profit is computed based on data $\mathbf{Z}^{()}$, which is simulated from the empirical density $\tilde{f}_{Z|\boldsymbol{Z}^o}$. Given that we are restricted to the scaled class here, the best possible solution is the parametric scaled operational statistic $y^*_{\text{SC}\iota}$, which is uniformly optimal within the scaled class. The cross-learning solution approaches this optimal solution asymptotically, which makes an essential contrast to the transfer-learning solutions; recall (17).

Recall from Theorem 2 that the transfer-learning solutions converge to some element within the scaled class, but not the optimal scaled operational statistics. Thus, the cross-learned scaled operational statistic dominates any transfer-learning solution with a sufficiently large n^o. Empirically,

Table 1 demonstrates that the profit $\Psi_{\text{CrL:SC1}}$ obtained by implementing the scaled cross-learning solution for the newsvendor problem is significantly higher than $\Psi_{\text{TrL:}j}$ obtained from any transfer-learning solution, even though the tested n^o is not very large. Moreover, the gap between $\Psi_{\text{CrL:SC1}}$ and the theoretical upper bound Ψ_{H1} is small even with a relatively small sample ($n^o = 200$) of the old system, suggesting the power of cross learning.

5.2 Cross Learning Through Boosting of a Candidate Solution

The solution derived in the previous subsection is constrained within the scaled operational statistics. We may seek potential improvements by expanding the solution space. One way is to boost from some candidate policy, y^c, that is known to be consistent. Examples of y^c include those oracle solutions described in Section 4.1. Boosting is a widely adopted approach in statistics to improve predictive power with limited data (see, e.g., Freund and Schapire 1997, Friedman et al. 2000, Hastie et al. 2009). In practice, a candidate solution is often chosen as one that is known to produce reasonable performance historically. We shall also assume that y^c satisfies Assumption 3 to ensure the consistency of the boosted solution; see Theorem 4 below.

To identify a subclass of operational statistics that goes beyond the scaled class using the candidate class, we recognize that a data-integrated model satisfies $y(\alpha \boldsymbol{X}) = \alpha^\iota y(\boldsymbol{X})$ and $\psi(\alpha^\iota y(\boldsymbol{X}), \alpha x) = \alpha^\kappa \psi(y(\boldsymbol{X}), x)$. Therefore, we can focus on the following class of operational statistics obtained by *targeted boosting* of the candidate solution y^c:

$$\mathcal{Y}^n_{\text{TB-OS}}(y^c) = \left\{ y : \mathbb{R}^n_+ \to \mathbb{R}_+, y(\boldsymbol{x}) = \gamma(\boldsymbol{x}) y^c(\boldsymbol{x}), \gamma \in \mathcal{H}^n_+(0, \pi) \right\} \ \subset\ \mathcal{H}^n_+(\iota, \pi). \tag{24}$$

That is, we look for a boosting statistic $\gamma(\boldsymbol{x})$ to scale the candidate solution. The boosting statistic γ is homogeneous of degree zero to ensure that the boosted solution stays within the homogeneous class. It is important to note that the boosted class $\mathcal{Y}^n_{\text{TB-OS}}(y^c)$ is only a subset of the homogeneous class $\mathcal{H}^n_+(\iota, \pi)$ for finite samples. Identifying the operational statistic for the base point $\boldsymbol{z}_B$ boils down to finding the boosting value $\gamma(\boldsymbol{z}_B)$ for $\boldsymbol{z}_B$. As $\gamma(\boldsymbol{x})$ is order-zero homogeneous, any sample within the partition $\mathcal{S}(\boldsymbol{z}_B)$ should have the same boosting value; Recall Equation (7).

To compute the objective for the base operational statistic, we observe from Equation (7) that we need to aggregate all the samples in the partition $\mathcal{S}(\boldsymbol{z}_B)$ for the base point $\boldsymbol{z}_B$. To avoid overfitting, we cluster samples with base points that are close enough to $\boldsymbol{z}_B$, and treat these samples as if they belong to $\mathcal{S}(\boldsymbol{z}_B)$ in computing the objective of the decision for $\boldsymbol{z}_B$. Specifically, let $d(\cdot, \cdot)$ be a distance measure and $\eta > 0$ be a closeness threshold. For our analysis, define $\boldsymbol{z}_{[]}$ as the ascending sequence of $\boldsymbol{z}$ such that $z_{[1]} \leq z_{[2]} \leq \cdots \leq z_{[n]}$. Then, we use the following Euclidean distance to determine the similarity of two samples $\boldsymbol{z}$ and $\hat{\boldsymbol{z}}$:

$$d(\boldsymbol{z}, \hat{\boldsymbol{z}}) = \sqrt{\frac{1}{n} \sum_{i=1}^{n} \left(z_{[i]} - \hat{z}_{[i]}\right)^2}. \tag{25}$$

We cluster samples that are close enough to $\boldsymbol{Z}_B$, i.e., $\{j : d(\boldsymbol{Z}_B, \boldsymbol{Z}_B^{(j)}) \leq \eta\}$. An immediate observation is that for $\eta > 0$, any simulated sample $\boldsymbol{Z}^{(j)}$ has a strictly positive probability to be within this set, as established in the next lemma.

Lemma 2 *Let $\hat{\boldsymbol{Z}}$ be an i.i.d sample of size n drawn from $\tilde{f}_{Z|\boldsymbol{z}^o}$ defined in (14) for given $\boldsymbol{Z}^o = \boldsymbol{z}^o$. Given $\eta > 0$ and $\boldsymbol{z}_B \in \mathcal{B}^n$, for sufficiently large n^o, there exists $\underline{C}_n > 0$ such that*

$$\Pr\{d(\hat{\boldsymbol{Z}}/\hat{\mu}_{\hat{\boldsymbol{Z}}}, \boldsymbol{z}_B) \leq \eta\} \geq \underline{C}_n \frac{\pi^{(n-1)/2}}{\Gamma(\frac{n-1}{2}+1)} \eta^{n-1}.$$

We further note that the average of the simulated sample within the cluster converges in distribution to that in partition $\mathcal{S}(\boldsymbol{z}_B)$ as the sample size in the old system gets large and the distance threshold gets close to zero.

Lemma 3 *Consider a continuous density f_Z and a sequence of continuous densities $\{f_{n^o}, n^o \in \mathbb{N}\}$ over support $\mathbb{R}_+$ such that $\lim_{n^o \to \infty} f_{n^o} \to f_Z$ almost everywhere. Let $\boldsymbol{Z}$ and $\boldsymbol{Z}_{n^o}$ be i.i.d. draws of f_Z and f_{n^o}, respectively, both of size n. Then, for a given $\boldsymbol{z}_B \in \{\boldsymbol{z} \in \mathbb{R}_+^n : \hat{\mu}_{\boldsymbol{z}} = 1\}$,*

$$\left[\hat{\mu}_{\boldsymbol{Z}_{n^o}} \middle| d(\boldsymbol{Z}_{n^o}/\hat{\mu}_{\boldsymbol{Z}_{n^o}}, \boldsymbol{z}_B) \leq \eta\right] \xrightarrow{d} \left[\hat{\mu}_{\boldsymbol{Z}_{n^o}} \middle| \boldsymbol{Z}_{n^o}/\hat{\mu}_{\boldsymbol{Z}_{n^o}} = \boldsymbol{z}_B\right], \text{ as } \eta \to 0, \quad \text{and}$$
$$\left[\hat{\mu}_{\boldsymbol{Z}_{n^o}} \middle| \boldsymbol{Z}_{n^o}/\hat{\mu}_{\boldsymbol{Z}_{n^o}} = \boldsymbol{z}_B\right] \xrightarrow{d} \left[\hat{\mu}_{\boldsymbol{Z}} \middle| \boldsymbol{Z}/\hat{\mu}_{\boldsymbol{Z}} = \boldsymbol{z}_B\right], \text{ as } n^o \to \infty.$$

We would like to have η decreasing in m so that $\lim_{m\to\infty} \eta = 0$. Also, we would like to have the number of samples satisfying $\{d(\boldsymbol{Z}^{(j)}/\hat{\mu}_{\boldsymbol{Z}^{(j)}}, \boldsymbol{z}_B) \leq \eta\}$ gets large as m gets large. Then, by Lemma 3, we can use the samples which, after normalizing, are close enough to $\boldsymbol{z}_B$ to compute the integration over $\mathcal{S}(\boldsymbol{z}_B)$ in Equation (7). Specifically, we estimate the cross-learning objective for the base point $\boldsymbol{z}_B$ as

$$\hat{\phi}_{\text{CrL:c}}(\gamma, y^c(\cdot), \boldsymbol{z}_B; \eta, \mathbf{Z}^{()}) = \frac{\sum_{j=1}^{m} \mathbb{I}_{\{d(\boldsymbol{z}_B, \boldsymbol{Z}_B^{(j)}) \leq \eta\}} \sum_{\substack{i \in \{1,2,\ldots,n\} \\ \ell \in \{1,2,\ldots,m\} \setminus j}} \psi(\gamma y^c(\boldsymbol{Z}^{(j)}), Z_i^{(\ell)})}{(m-1)n \sum_{j=1}^{m} \mathbb{I}_{\{d(\boldsymbol{z}_B, \boldsymbol{Z}_B^{(j)}) \leq \eta\}}}. \tag{26}$$

The validating model finds the boosting statistic for the base point as

$$\gamma_{\text{CrL:c}}(\boldsymbol{z}_B) = \arg\max \left\{\hat{\phi}_{\text{CrL:c}}(\gamma, y^c(\cdot), \boldsymbol{z}_B; \eta, \mathbf{Z}^{()}) : \gamma \in \mathbb{R}_+\right\},$$

and the cross-learning solution for the base point is

$$y_{\mathrm{CrL:c}}(\boldsymbol{z}_B) = \gamma_{\mathrm{CrL:c}}(\boldsymbol{z}_B) y^c(\boldsymbol{z}_B).$$

Once we compute the cross-learning solution for the base point, we can project it to the observed sample to obtain $y_{\mathrm{CrL:c}}(\boldsymbol{x}) = (\hat{\mu}_{\boldsymbol{x}})^{\iota} y_{\mathrm{CrL:c}}(\boldsymbol{x}/\hat{\mu}_{\boldsymbol{x}})$ as suggested by Equation (9).

We should underscore the fundamental difference between transfer learning from the oracle solution in Section 4.2 and cross learning through targeted boosting. The transfer-learning solution directly applies the pre-trained oracle solution of the old system to the focal system through the common data structure (i.e., $X^o = \theta^o Z$ and $X = \theta Z$). In our problem, solution transfer boils down to scaling the oracle solution by the ratio of sample averages. In an essential contrast, cross learning ignores any solution y^o trained for the old system, but fully utilizes the data $\boldsymbol{X}^o$. The candidate solution y^c is trained using the data from the focal system, and the data from the old system are used to create the generic systems. Then, the candidate solution is fine-tuned in the generic systems.

From the example in Table 1, we observe that boosting a candidate solution trained in the focal system reveals superior performance compared to transferring the same solution pre-trained from the old system. More importantly, the performance is more robust with the former approach in the sense that the variability (measured by standard deviation) is much lower and the worst case of the resulting profit is much higher. Another experiment with a beta distribution is reported in Online Appendix D, from which the observations are consistent.

Although the boosted class $\mathcal{Y}^n_{\text{TB-OS}}(y^c)$ may not contain the best homogeneous solution $y^*_{H\iota}$ for a finite n^o, the performance of the cross-learning solution approaches that of the parametric ODA solution $y^*_{H\iota}$ as n^o gets large.

Theorem 4 (Consistency of Cross-Learned Boosting Statistics) *Under Assumptions 1 and 2, if $\Gamma \subset \mathbb{R}_+$ is compact, $y^c \in \mathcal{H}^{n^o}_+(\iota, \pi)$ is bounded and satisfies Assumption 3, $m \propto (n^o)^n$, and $\eta \propto m^{-\frac{1}{n}}$, then for any $\epsilon > 0$ and a sufficiently large n^o,*

$$\Pr\left\{\sup_{\gamma\in\Gamma}\left|\hat{\phi}_{\mathrm{CrL:c}}(\gamma, y^c(\cdot), \boldsymbol{z}_B; \eta, \mathbf{Z}^{()}) - \phi_B\big[\gamma y^c(\boldsymbol{z}_B), \boldsymbol{z}_B, F_Z\big]\right| \geq \epsilon\right\} \leq \mathcal{O}\bigg(\exp\Big(-\frac{2n^o\epsilon^2 - 4C\sqrt{n^o}\epsilon}{(b-a)^2}\Big)\bigg), \quad (27)$$

where $C > 0$. In addition, if the maximizer of $\phi_B\big[\gamma y^c(\boldsymbol{z}_B), \boldsymbol{z}_B, F_Z\big]$ is unique, then

$$y_{\mathrm{CrL:c}}(\boldsymbol{z}_B) \xrightarrow{p} y^*_{\mathrm{H}\iota}(\boldsymbol{z}_B), \text{ as } n^o \to \infty. \quad (28)$$

Furthermore, if the maximizer of $\phi_B\big[\gamma y^c(\boldsymbol{z}_B), \boldsymbol{z}_B, F_Z\big]$ is unique and interior on Γ, $\psi(\cdot, x)$ is continuously differentiable with Lipschitz continuity and bounded $\nabla\psi(\cdot, x)$ over Γ a.s., and $\mathsf{E}[\psi(\cdot, X)]$

is strongly concave, then for any $\epsilon > 0$ and a sufficiently large n^o,

$$\Pr\left\{\left|y_{\mathrm{CrL:c}}(\boldsymbol{z}_B) - y^*_{\mathrm{H}\iota}(\boldsymbol{z}_B)\right| \geq \epsilon\right\} \leq \mathcal{O}\Bigg(\exp\Big(-\frac{2n^o\epsilon^2 - 4C\sqrt{n^o}\epsilon}{(b-a)^2}\Big)\Bigg). \tag{29}$$

Similar to the scaled operational statistic, the cross-learned boosting solution is obtained by optimizing a subclass of the homogeneous operational statistics extended from some candidate solution. The difference lies in the fact that the scaling parameter is independent of the individual sample, while the boosting parameter depends on the observed sample. Thus, the variability associated with the scaled solution is generally lower when the sample size n^o in the old system is small (e.g., $n^o = 200$). This is evident from Table 1.

Compared to the transfer-learning solutions, the cross-learned boosting solutions exhibit two notable features, as shown in Table 1. First, the performance of a transfer-learned solution is directly tied to that of its corresponding directly learned candidate solution–when the candidate solution yields low profit, the transfer-learned solution also generates a relatively low profit. In contrast, the performance of the cross-learned boosting solution does not exhibit this relationship. Second, the profit variation among different boosting solution is much smaller than that observed among transfer-learning solutions or directly trained solutions. This suggests that the performance of the boosted solution is relatively insensitive to the choice of the candidate solution. As suggested by Theorem 4, when n^o becomes sufficiently large, all boosted learning solutions get close to the uniformly optimal solution within the homogeneous class, $y^*_{\mathrm{H}\iota}$.

5.3 Cross Learning for Optimal Homogeneous Operational Statistics

In the previous subsection, we have looked for a solution within the boosted class $\mathcal{Y}^n_{\text{TB-OS}}(y^c)$ derived based on a candidate solution y^c. Because the boosted class is a subset of the homogeneous class, i.e., $\mathcal{Y}^n_{\text{TB-OS}}(y^c) \subset \mathcal{H}^n_+(\iota, \pi)$, the best possible data-integration solution, $y^*_{\mathrm{H}\iota}$, may be outside of $\mathcal{Y}^n_{\text{TB-OS}}(y^c)$ with a finite n^o. In this subsection, we look for optimizing the solution over the entire homogeneous class $\mathcal{H}^n_+(\iota, \pi)$.

Different from the philosophy of boosting in statistics, the optimal homogeneous solution is constructed based on the parametric ODA framework. Specifically, the key is to compute the optimal operational statistic for the base point $\boldsymbol{z}_B$ corresponding to the data $\boldsymbol{X}$ observed from the focal system. Based on our earlier discussion, this operational statistic can be derived from the generic system with random variable Z. For that, we need to estimate the objective in (7) of the parametric ODA validation with a solution $y(\boldsymbol{x}) = (\hat{\mu}_{\boldsymbol{x}})^\iota \gamma$, where $\gamma = y(\boldsymbol{z}_B)$. Thus, the validating model in (26) now becomes

$$\hat{\phi}_{\mathrm{CrL:H}\iota}(\gamma, \boldsymbol{z}_B; \eta, \mathbf{Z}^{()}) = \frac{\sum_{j=1}^{m} \mathbb{I}_{\{d(\boldsymbol{z}_B, \boldsymbol{Z}_B^{(j)}) \leq \eta\}} \sum_{\substack{i \in \{1,2,\ldots,n\} \\ \ell \in \{1,2,\ldots,m\} \setminus j}} \psi((\hat{\mu}_{\boldsymbol{Z}^{(j)}})^{\iota}\gamma, Z_i^{(\ell)})}{(m-1)n \sum_{j=1}^{m} \mathbb{I}_{\{d(\boldsymbol{z}_B, \boldsymbol{Z}_B^{(j)}) \leq \eta\}}}. \tag{30}$$

The optimal base operational statistic is

$$y_{\mathrm{CrL:H}\iota}(\boldsymbol{z}_B) = \arg\max \big\{ \hat{\phi}_{\mathrm{CrL:H}\iota}(\gamma, \boldsymbol{z}_B; \eta, \mathbf{Z}^{()}), \gamma > 0 \big\},$$

and the solution for the focal system is $y_{\mathrm{CrL:H}\iota}(\boldsymbol{x}) = (\hat{\mu}_{\boldsymbol{x}})^{\iota} y_{\mathrm{CrL:H}\iota}(\boldsymbol{x}/\hat{\mu}_{\boldsymbol{x}})$.

Comparing (26) and (30), it may appear that $y_{\mathrm{CrL:H}\iota}$ is obtained by boosting a candidate solution $y^c = (\hat{\mu}_{\boldsymbol{x}})^{\iota}$. However, this candidate solution is in general not consistent. In particular, ordering mean demand is in general suboptimal for newsvendor model. Thus, Theorem 4 does not apply to (30). Moreover, one may treat (20) as a special case of (30) with $\eta \to \infty$ in the latter. However, these two models converge to different limits with sufficiently large samples because of the different data-integration models.

Theorem 5 (Consistency of Cross-Learned Homogeneous Statistics) *Under Assumptions 1 and 2, if $\Gamma \subset \mathbb{R}_+$ is compact, $m \propto (n^o)^n$, and $\eta \propto m^{-\frac{1}{n}}$, then for any $\epsilon > 0$ and a sufficiently large n^o,*

$$\Pr\left\{ \sup_{\gamma \in \Gamma} \left| \hat{\phi}_{\mathrm{CrL:H}\iota}(\gamma; \boldsymbol{z}_B; \eta, \boldsymbol{Z}^{()}) - \phi_B[\gamma, \boldsymbol{z}_B, F_Z] \right| \geq \epsilon \right\} \leq \mathcal{O}\Bigg(\exp\Big(- \frac{2n^o\epsilon^2 - 4C\sqrt{n^o}\epsilon}{(b-a)^2} \Big) \Bigg), \tag{31}$$

where $C > 0$. In addition, if $y^*_{\mathrm{H}\iota}(\boldsymbol{x})/(\hat{\mu}_{\boldsymbol{x}})^{\iota}$ is the unique maximizer on Γ, then

$$y_{\mathrm{CrL:H}\iota}(\boldsymbol{z}_B) \xrightarrow{p} y^*_{\mathrm{H}\iota}(\boldsymbol{z}_B), \ \text{as } n^o \to \infty. \tag{32}$$

Furthermore, if $y^*_{\mathrm{H}\iota}(\boldsymbol{x})/(\hat{\mu}_{\boldsymbol{x}})^{\iota}$ is a unique and interior maximizer on Γ, $\psi(\cdot, x)$ is continuously differentiable with Lipschitz continuous and bounded $\nabla\psi(\cdot, x)$ over Γ a.s., and $\mathsf{E}[\psi(\cdot, X)]$ is strongly concave, then for any $\epsilon > 0$ and a sufficiently large n^o,

$$\Pr\left\{ \left| y_{\mathrm{CrL:H}\iota}(\boldsymbol{z}_B) - y^*_{\mathrm{H}\iota}(\boldsymbol{z}_B) \right| \geq \epsilon \right\} \leq \mathcal{O}\Bigg(\exp\Big(- \frac{2n^o\epsilon^2 - 4C\sqrt{n^o}\epsilon}{(b-a)^2} \Big) \Bigg). \tag{33}$$

Certainly, when the data from the old system offer rich insights into the stochastic environment of the associated decision-making problem (i.e., when n^o is sufficiently large), the optimal homogeneous operational statistic is the most desired solution, as it approaches the best possible data-integrated decision, $y^*_{\mathrm{H}\iota}$. When the data from the old system is not too large, the simulated samples falling into the cluster may not fully reflect the actual statistical nature of any random sample corresponding to the base point. In this case, the cross-learned homogeneous solution may not exhibit superiority against other cross-learning solutions. This is because, with limited data, a homogeneous solution may overfit, while a solution with fewer unknown parameters produces smaller variability.

As a general observation from Table 1, the cross-learning solutions are robust in the sense that they not only improve the transfer-learning solutions, but also produce profits with significantly reduced variability measured by the standard deviation. As the sample size of focal system, n^o, gets large, the performance of every cross-learning solution converges to the benchmark Ψ_{H1}. Among all these solutions, the cross-learned scaled operational statistic always outperforms others because the optimal solution within the homogeneous class is a scaled operational statistic when Z follows a gamma distribution. Unlike boosted or homogeneous operational statistics, the optimal scaled operational statistics can be obtained by optimizing a single parameter, $\gamma_{\mathrm{Crl:SC1}}$, while fully utilizing all data from generic systems without clustering. This approach can yield higher solution precision compared to boosted or homogeneous statistics with limited samples. Therefore, despite being a special case of the homogeneous class, the scaled class is interesting to study in its own right. Certainly, with a different distribution (see the experiment on beta distribution in Online Appendix D), other solutions can outperform the scaled operational statistics.

We remark that the key idea of cross learning is to simulate a large number m of generic systems to effectively capture the distributional information of F_Z contained in the data from the old system. From a theoretical standpoint, ensuring the efficiency of the learning solution requires $m \to \infty$. Empirically, however, our numerical experiments indicate that a much smaller m (e.g., on the order of 10,000) is sufficient to achieve substantial performance improvements.

6 Co-Learning

When there are multiple focal systems and limited data are available from each system, we may explore the statistical similarity among the data sets. Based on such similarity, one may pool the data among different systems to improve the quality of prediction and efficiency of decision-making. The widely applied data pooling approach, originated from James and Stein (1961), combines the data from different systems and aims at improving the overall performance of all systems.

The co-learning idea we propose, instead, explores the similarity of the data-to-decision structure among all systems and designs solutions that improve not only the aggregate performance

Table 1: Performance of Transfer/Cross-Learning Solutions Against Directly Trained Solutions

		Directly Trained Solutions				Transfer-Learning Solutions				Cross-Learning Solutions			
			Ave	Stdev	Min		Ave	Stdev	Min		Ave	Stdev	Min
$n=5$	Benchmark	Ψ_{H1}	0.076	0.07	-2.6								
	$n^o=200$	Ψ_{SAA}	-0.598	2.59	-99.1	$\Psi_{\text{TrL:SAA}}$	0.011	0.38	-16.2	$\Psi_{\text{CrL:SAA}}$	0.063	0.15	-10.1
		$\Psi_{\text{Re-Var}}$	-0.159	1.82	-105.0	$\Psi_{\text{TrL:Re-Var}}$	0.034	0.30	-16.4	$\Psi_{\text{CrL:Re-Var}}$	0.063	0.15	-10.4
		$\Psi_{\text{Re-Std}}$	-0.273	2.14	-105.0	$\Psi_{\text{TrL:Re-Std}}$	0.027	0.33	-16.4	$\Psi_{\text{CrL:Re-Std}}$	0.057	0.14	-10.1
		Ψ_{OrS}	-1.439	5.19	-149.9	$\Psi_{\text{TrL:OrS}}$	0.005	0.41	-16.3	$\Psi_{\text{CrL:OrS}}$	0.063	0.15	-10.0
		$\Psi_{\text{SC1-Emp}}$	-0.675	2.77	-96.0	$\Psi_{\text{TrL:SC1-Emp}}$	0.002	0.41	-16.5	$\Psi_{\text{CrL:SC1-Emp}}$	0.062	0.15	-9.5
		$\Psi_{\text{H1-Emp}}$	-0.839	3.25	-101.0	$\Psi_{\text{TrL:H1-Emp}}$	-0.001	0.43	-17.9	$\Psi_{\text{CrL:H1-Emp}}$	0.064	0.14	-9.3
										$\Psi_{\text{CrL:SC1}}$	0.069	0.11	-4.4
										$\Psi_{\text{CrL:H1}}$	0.063	0.15	-10.0
	$n^o=1000$	Ψ_{SAA}	-0.598	2.59	-99.1	$\Psi_{\text{TrL:SAA}}$	0.033	0.27	-8.1	$\Psi_{\text{CrL:SAA}}$	0.073	0.09	-3.5
		$\Psi_{\text{Re-Var}}$	-0.159	1.82	-105.0	$\Psi_{\text{TrL:Re-Var}}$	0.042	0.24	-8.1	$\Psi_{\text{CrL:Re-Var}}$	0.074	0.09	-3.7
		$\Psi_{\text{Re-Std}}$	-0.273	2.14	-105.0	$\Psi_{\text{TrL:Re-Std}}$	0.041	0.24	-8.1	$\Psi_{\text{CrL:Re-Std}}$	0.065	0.08	-3.5
		Ψ_{OrS}	-1.439	5.19	-149.9	$\Psi_{\text{TrL:OrS}}$	0.032	0.27	-8.1	$\Psi_{\text{CrL:OrS}}$	0.074	0.09	-3.4
		$\Psi_{\text{SC1-Emp}}$	-0.675	2.77	-96.0	$\Psi_{\text{TrL:SC1-Emp}}$	0.030	0.28	-8.0	$\Psi_{\text{CrL:SC1-Emp}}$	0.072	0.08	-3.3
		$\Psi_{\text{H1-Emp}}$	-0.838	3.24	-101.1	$\Psi_{\text{TrL:H1-Emp}}$	0.030	0.28	-8.2	$\Psi_{\text{CrL:H1-Emp}}$	0.072	0.08	-3.2
										$\Psi_{\text{CrL:SC1}}$	0.074	0.08	-2.6
										$\Psi_{\text{CrL:H1}}$	0.074	0.09	-3.5
$n=10$	Benchmark	Ψ_{H1}	0.105	0.06	-1.4								
	$n^o=200$	Ψ_{SAA}	-0.217	1.31	-40.5	$\Psi_{\text{TrL:SAA}}$	0.068	0.22	-8.0	$\Psi_{\text{CrL:SAA}}$	0.085	0.14	-11.9
		$\Psi_{\text{Re-Var}}$	-0.096	1.11	-51.2	$\Psi_{\text{TrL:Re-Var}}$	0.080	0.17	-8.4	$\Psi_{\text{CrL:Re-Var}}$	0.082	0.16	-14.8
		$\Psi_{\text{Re-Std}}$	-0.253	1.50	-51.2	$\Psi_{\text{TrL:Re-Std}}$	0.076	0.19	-8.4	$\Psi_{\text{CrL:Re-Std}}$	0.081	0.16	-11.9
		Ψ_{OrS}	-0.578	2.06	-51.4	$\Psi_{\text{TrL:OrS}}$	0.064	0.23	-8.4	$\Psi_{\text{CrL:OrS}}$	0.086	0.15	-11.7
		$\Psi_{\text{SC1-Emp}}$	-0.251	1.32	-42.4	$\Psi_{\text{TrL:SC1-Emp}}$	0.063	0.23	-8.2	$\Psi_{\text{CrL:SC1-Emp}}$	0.085	0.14	-12.5
		$\Psi_{\text{H1-Emp}}$	-0.276	1.46	-41.8	$\Psi_{\text{TrL:H1-Emp}}$	0.061	0.24	-8.7	$\Psi_{\text{CrL:H1-Emp}}$	0.089	0.12	-10.6
										$\Psi_{\text{CrL:SC1}}$	0.095	0.09	-3.5
										$\Psi_{\text{CrL:H1}}$	0.086	0.15	-11.0
	$n^o=1000$	Ψ_{SAA}	-0.217	1.31	-40.5	$\Psi_{\text{TrL:SAA}}$	0.087	0.14	-3.9	$\Psi_{\text{CrL:SAA}}$	0.097	0.07	-1.9
		$\Psi_{\text{Re-Var}}$	-0.096	1.11	-51.2	$\Psi_{\text{TrL:Re-Var}}$	0.093	0.12	-3.9	$\Psi_{\text{CrL:Re-Var}}$	0.097	0.08	-2.5
		$\Psi_{\text{Re-Std}}$	-0.253	1.50	-51.2	$\Psi_{\text{TrL:Re-Std}}$	0.092	0.12	-3.9	$\Psi_{\text{CrL:Re-Std}}$	0.095	0.08	-3.1
		Ψ_{OrS}	-0.578	2.06	-51.4	$\Psi_{\text{TrL:OrS}}$	0.087	0.14	-3.9	$\Psi_{\text{CrL:OrS}}$	0.102	0.07	-2.5
		$\Psi_{\text{SC1-Emp}}$	-0.251	1.32	-42.4	$\Psi_{\text{TrL:SC1-Emp}}$	0.086	0.14	-4.0	$\Psi_{\text{CrL:SC1-Emp}}$	0.099	0.07	-2.6
		$\Psi_{\text{H1-Emp}}$	-0.276	1.46	-41.5	$\Psi_{\text{TrL:H1-Emp}}$	0.086	0.14	-4.0	$\Psi_{\text{CrL:H1-Emp}}$	0.098	0.06	-1.5
										$\Psi_{\text{CrL:SC1}}$	0.103	0.07	-1.9
										$\Psi_{\text{CrL:H1}}$	0.102	0.07	-2.6

Notes. $Z \sim Gamma(0.1, 10)$ (implying $\mathsf{Cv}[Z] = 3.16$), $\mathsf{E}[X] = 1.5$, $\mathsf{E}[X^o] = 15$, and $n = \{5, 10\}$. The profit function of the focal system is $\psi(y, x) = 10 \min\{y, x\} - 3y$, and the expected profit under the parametric homogeneous solution is $\Psi_{\text{H1}} = 0.076$. The subscripts TrL:j and CrL:j stand for, respectively, transfer-learning and cross-learning solutions, where $j = \text{SAA}, \text{Re-Var}, \text{Re-Std}, \text{OrS}, \text{SC1-Emp}$, or H1-Emp stands for the sample-average approximation, regularization with variance, regularization with standard deviation, order statistics, scaled operational statistic with empirical distribution, or homogeneous operational statistic with empirical distribution, respectively. The result is generated by computing the average, standard deviation, and minimum of the actual profit Ψ_j out of $100,000$ randomly generated instances.

but also the individual performance. To achieve that, we recognize that for systems described by Assumptions 1-2, the best data-integrated solution (i.e., the parametric ODA solution of the homogeneous class, y^*_{H1}) does not depend on the different scale parameters across the systems, but only on the distribution of the common random factor Z, as discussed in Section 3.1. In particular, our discussion of the parametric ODA solution (recall Equation (6)) suggests that the optimal base operational statistic, $\{y_B^{\mathrm{OS}}(\boldsymbol{z}_B), \boldsymbol{z}_B \in \mathcal{B}^{n-1}\}$, is independent of the unknown scale parameter θ and the observed sample (when we know F_Z). This property is essential for combining the data across the systems to derive the co-learning solution.

The random variable involved in a system is of the form $X = \Theta Z$. We have data from k systems denoted as $\mathbf{X} = \{\boldsymbol{X}_j : j = 1, 2, \ldots, k\}$, where $\boldsymbol{X}_j = (X_{j:1}, X_{j:2}, \ldots, X_{j:n})$ consists of n i.i.d. observations from system j. Because we do not know the Θ value of any system, each $X_{j:i}$ is a realization of X. However, we know that $(X_{j:i} : i \in \{1, 2, \ldots, n\})$ share the same Θ value. To distinguish, we denote this value by Θ_j, a random copy of Θ, and denote $\boldsymbol{Z}_j = \boldsymbol{X}_j/\Theta_j$. Then, $\boldsymbol{Z}_j$ consists of n i.i.d. copies of F_Z.

For any n-element vector $\boldsymbol{a}$, let $\boldsymbol{a}_{\text{-}\ell} = (a_i : i = 1, 2, \ldots, n; i \neq \ell)$ denote the subvector excluding the ℓth element. Note that $\boldsymbol{X}_{j:\text{-}\ell}$ and $X_{j:\ell}$ are independent. Thus, we can use the subvector $\boldsymbol{X}_{j:\text{-}\ell}$ to define the statistics and $X_{j:\ell}$ to validate. Specifically, consider implementing an operational statistic $y \in \mathcal{H}_+^{n-1}(\iota, \pi)$. Then by Equation (3), the expected profit collected from system j is

$$\mathsf{E}\big[\phi[y(\boldsymbol{X}_{j:\text{-}\ell}), F_X]|\Theta = \Theta_j\big] = \mathsf{E}\big[\psi(y(\boldsymbol{X}_{j:\text{-}\ell}), X_{j:\ell})|\Theta = \Theta_j\big] = \Theta_j^\kappa \mathsf{E}[\psi(y(\boldsymbol{Z}_{j:\text{-}\ell}), Z_{j:\ell})]. \tag{34}$$

Though the value of Θ is different across different systems, by Equation (6), the optimal base operational statistic $\{y_B^{\mathrm{OS}}(\boldsymbol{z}_B), \boldsymbol{z}_B \in \mathcal{B}^{n-1}\}$, is the same regardless of the value of Θ. Thus, if y is an efficient solution for some system, then the base operational statistic is $y_B(\boldsymbol{X}_{\text{-}1}/\hat{\mu}_{\boldsymbol{X}_{\text{-}1}}) = y(\boldsymbol{X}_{\text{-}1})/(\hat{\mu}_{\boldsymbol{X}_{\text{-}1}})^\iota$ for any observed sample $\boldsymbol{X}_{\text{-}1}$ corresponding to any Θ. Thus, when data $\boldsymbol{X}_{j:\text{-}1}$ are observed in system j, one can simply implement $(\hat{\mu}_{\boldsymbol{X}_{j:\text{-}1}})^\iota y_B(\boldsymbol{X}_{j:\text{-}1}/\hat{\mu}_{\boldsymbol{X}_{j:\text{-}1}}) = y(\boldsymbol{X}_{j:\text{-}1})$. In other words, given the common data structure $X = \Theta Z$, the functional form of the decisions for all systems should be the same.

When choosing a $y \in \mathcal{H}_+^{n-1}(\iota, \pi)$, we can empirically approximate the average expected profit obtained from k systems as

$$\hat{\phi}(y(\cdot), \mathbf{X}) = \frac{1}{kn} \sum_{j=1}^{k} \sum_{\ell=1}^{n} \psi(y(\boldsymbol{X}_{j:\text{-}\ell}), X_{j:\ell}). \tag{35}$$

In the above calculation, we randomly pick one element from the data $\boldsymbol{X}_j$ from each system to validate an $(n-1)$-dimensional operational statistic. Our goal is to identify an appropriate operational statistic y to maximize the overall empirical profit calculated in (35). Given the limited data in individual systems, implementing a candidate policy y^c (discussed in Section 4.1) for each system can be inefficient. We look to expand the candidate policy to a class of operational statistics, and formulate an appropriate validating model that utilizes all the available data. Two solution approaches are discussed in the next subsections.

We note that while we assume the same sample size across all systems, our analysis can be easily extended to scenarios where the sample size n_j varies across systems. When $(\max_{1\le j\le k} n_j - \min_{1\le j\le k} n_j)$ is small, we choose $n_0 = \min\{n_j : j = 1, 2, \ldots, k\} - 1$. Define $\mathcal{S}_j = \{(i_1, \ldots, i_{n_j - n_0}) : \cup_{\ell=1}^{n_j - n_0}\{i_\ell\} \subset \cup_{i=1}^{n_j}\{i\}, \cap_{\ell=1}^{n_j-n_0}\{i_\ell\} = \emptyset\}$ as the collection of all $(n_j - n_0)$-element subsets of indices $\{1, 2, \ldots, n_j\}$, and $\boldsymbol{X}_{j:\text{-}S} = (X_{j:i} : i \in \cup_{i_0=1}^{n_j}\{i_0\}\backslash S)$ as the subset of $\boldsymbol{X}_j$ excluding the elements in $S \in \mathcal{S}_j$. Then the empirical approximation of the expected profit from system j can be computed as

$$\hat{\phi}_j(y(\cdot), \boldsymbol{X}_j) = \frac{1}{(n_j - n_0)\binom{n_j}{n_0}} \sum_{S\in\mathcal{S}_j} \sum_{\ell\in S} \psi(y(\boldsymbol{X}_{j:\text{-}S}, X_{j:\ell})).$$

Then the average profit across all systems becomes $\hat{\phi}(y(\cdot), \mathbf{X}) = \frac{1}{k}\sum_{j=1}^{k} \hat{\phi}_j(y(\cdot), \boldsymbol{X}_j)$.

When $(\max_{1\le j\le k} n_j - \min_{1\le j\le k} n_j)$ is large, we may choose a threshold n_0 and apply co-learning only to systems with sample size above n_0, while applying cross learning to those with small sample size.

Unlike in the situation of cross learning where one can generate as many generic systems as needed to capture the stochastic nature of the decision-making, the number of available systems k for co-learning is fixed and limited in reality. It is thus reasonable to resort to boosting a candidate solution.

6.1 Co-Learning with Constant Boosting from a Candidate Solution

Suppose that we are given a candidate solution $y^c \in \mathcal{H}_+^{n-1}(\iota, \pi)$ (e.g., any oracle solution described in Section 4.1). We would like to expand from this solution to a class of operational statistics, while keeping the homogeneous property. An immediate way is *constant boosting*:

$$\mathcal{Y}_{\text{CB-OS}}^{n-1}(y^c) = \left\{y : \mathbb{R}_+^{n-1} \to \mathbb{R}_+, y(\boldsymbol{x}) = \gamma y^c(\boldsymbol{x}), \gamma \in \mathbb{R}_+\right\} \subseteq \mathcal{H}_+^{n-1}(\iota, \pi). \tag{36}$$

Because y^c belongs to the homogeneous class, so does any element in $\mathcal{Y}_{\text{CB-OS}}^{n-1}(y^c)$.

Based on our discussion of the parametric ODA solution in Section 3.1, if we know F_Z, then we can solve the validation model for the constant-boosting class:

$$\max\left\{\mathsf{E}[\psi(\gamma y^c(\boldsymbol{X}_{\text{-}1}), X_1)] : \gamma \in \mathbb{R}_+\right\},$$

where $\boldsymbol{X}$ is a vector of n i.i.d. draws of $X = \Theta Z$ for some specific draw of $\Theta = \Theta_j$. Because the form of the optimal operational statistic is independent of individual Θ_j, we use (35) to approximate the expected profit of any system. Then, the boosting constant can be computed as

$$\gamma_{\text{CoL-CB:c}} = \arg\max\left\{\hat{\phi}(\gamma y^c(\cdot), \mathbf{X}) : \gamma > 0\right\}.$$

In this case, the co-learning operational statistic is

$$y_{\text{CoL-CB:c}}(\boldsymbol{x}) = \gamma_{\text{CoL-CB:c}} y^c(\boldsymbol{x}),\ \boldsymbol{x} \in \mathbb{R}_+^{n-1}.$$

Theorem 6 (Co-Learning Solution with Constant Boosting) *Under Assumptions 1, 2 and 3, for any $\epsilon > 0$,*

$$\Pr\left\{\sup_{\gamma\in\Gamma}\left|\hat{\phi}(\gamma y^c(\cdot), \mathbf{X}) - \mathsf{E}[\Theta^\kappa]\mathsf{E}[\psi(\gamma y^c(\boldsymbol{Z}_{\text{-}1}), Z_1)]\right| \geq \epsilon\right\} \leq \mathcal{O}\left(\exp\left(-\frac{2k\epsilon^2 - 4C\sqrt{k}\epsilon}{(b-a)^2}\right)\right).$$

Moreover, if $\mathsf{E}[\psi(\gamma y^c(\boldsymbol{Z}_{\text{-}1}), Z_1)]$ has a unique maximizer on $\gamma \in \Gamma$, denoted by $\gamma^*_{\text{CoL-CB:c}} := \text{argmax}_{\gamma\geq 0}$ $\{\mathsf{E}[\psi(\gamma y^c(\boldsymbol{Z}_{\text{-}1}), Z_1)]\}$, then

$$y_{\text{CoL-CB:c}}(\boldsymbol{x}) \xrightarrow{p} \gamma^*_{\text{CoL-CB:c}} y^c(\boldsymbol{x}),\ \forall \boldsymbol{x} \in \mathbb{R}_+^{n-1},\ \textit{as } k \to \infty.$$

Furthermore, if the maximizer of $\mathsf{E}[\psi(\gamma y^c(\boldsymbol{Z}_{\text{-}1}), Z_1)]$ is unique and interior on Γ, $\psi(\cdot, x)$ is continuously differentiable with Lipschitz continuous and bounded $\nabla\psi(\cdot, x)$ over Γ a.s., and $\mathsf{E}[\psi(\cdot, X)]$ is strongly concave, then for any $\epsilon > 0$,

$$\Pr\left\{\left|y_{\text{CoL-CB:c}}(\boldsymbol{x}) - \gamma^*_{\text{CoL-CB:c}} y^c(\boldsymbol{x})\right| \geq \epsilon\right\} \leq \mathcal{O}\left(\exp\left(-\frac{2k\epsilon^2 - 4C\sqrt{k}\epsilon}{(b-a)^2}\right)\right).$$

Theorem 6 suggests that the co-learning solution with constant boosting is asymptotically optimal within the class $\mathcal{Y}^{n-1}_{\text{CB-OS}}(y^c)$ provided that the candidate solution y^c comes from the homogeneous class. We underscore that this solution achieves optimality for both the combined systems and individual systems. This is because the way we pool the data to construct the operational statistics accounts for the fact that parametric ODA solutions are related to the same base point, and the base operational statistic is the same across all systems. The constant-boosting solution utilizes the data from different systems to mimic the uncertain environment faced by any individual system, which determines the appropriate choice of the operational statistic.

A candidate solution that is not consistent is the sample average, i.e., $y^c(\boldsymbol{x}) = \frac{1}{n-1}\sum_{i=1}^{n-1} x_i$. In this case, the co-learning solution with constant boosting converges to the optimal scaled operational statistic, $y^*_{\text{SC1}}(\boldsymbol{x})$.

We should remark that, though the observed data contain n samples, the co-learning solution constructed is an $(n-1)$-dimensional function. For system j with sample $\boldsymbol{x}_j \in \mathbb{R}_+^n$, we may implement the above solution by randomly splitting the vector $\boldsymbol{x}$ into $\boldsymbol{x}_{j:\text{-}\ell} \cup x_{j:\ell}$, and making a decision $y_{\text{CoL-CB:c}}(\boldsymbol{x}_{j:\text{-}\ell})$. One drawback of this approach is that when n is small, the boosting parameter $\gamma_{\text{CoL-CB:c}}$ depends heavily on n as the solution may be sensitive to how the vector is split. Alternatively, we may compute the decision y_j for system j through the equation

$$\frac{1}{n}\sum_{n=1}^{n} \psi(y_j, \boldsymbol{X}_j) = \frac{1}{n}\sum_{n=1}^{n} \psi(\gamma_{\text{CoL-CB:c}} y^c(\boldsymbol{X}_{j:\text{-}\ell}), X_{j:\ell}).$$

In other words, we look for a decision that leads to a sample-average profit that is equal to the average profit derived from implementing the constant-boosting solution for the $(n-1)$-element subvectors. It is easy to see that this solution is asymptotically optimal based on Theorem 6.

The constant-boosting solution adjusts the candidate solution uniformly across all systems through a common boosting parameter $\gamma_{\text{CoL-CB:c}}$, without exploring the differences across the systems that may be reflected in the data $\mathbf{X}$. To address such possibility, we introduce co-learning with targeted boosting.

6.2 Co-Learning with Targeted Boosting from a Candidate Solution

The given candidate solution $y^c \in \mathcal{H}_+^{n-1}(\iota, \pi)$ may not be sensitive to the distinctions among the involved systems. In this case, forming a class of operational statistics that allows the flexibility to distinguish the systems based on the observed data can adjust the candidate solution for better performance. For that purpose, we can apply the *targeted boosting* to adjust the candidate solution based on the sample $\boldsymbol{x}_j$ observed for the individual focal system. Specifically, the boosting parameter should be an order-zero homogeneous function $\gamma(\boldsymbol{x})$ so that the targeted-boosting solution $\gamma(\boldsymbol{x}_j)y^c(\boldsymbol{x}_j)$ is order-ι homogeneous; recall the targeted-boosting class $\mathcal{Y}^{n-1}_{\text{TB-OS:c}}(y^c)$ defined in (24). This means that every system has its own boosting parameter γ and thus a total of k parameters need to be determined. As a result, the best γ value is chosen for the observed sample of each individual system, and there is no value of pooling the data. To avoid such overfitting, we cluster the samples using the distance measure defined in (25) with a distance threshold $\eta > 0$. The estimated pooled profit for a base point $\boldsymbol{z}_B \in \mathbb{R}_+^{n-1}$ becomes

$$\hat{\phi}_{\text{CoL-TB}}(\gamma, y^c(\cdot), \boldsymbol{z}_B, \eta, \mathbf{X}) = \frac{\sum_{j=1}^{k}\sum_{\ell=1}^{n} \mathbb{I}_{\{d(\boldsymbol{z}_B, \boldsymbol{X}_{j:\text{-}\ell}/\hat{\mu}_{\boldsymbol{X}_{j:\text{-}\ell}}) \le \eta\}} \psi(\gamma y^c(\boldsymbol{X}_{j:\text{-}\ell}), X_{j:\ell})}{\sum_{j=1}^{k}\sum_{\ell=1}^{n} \mathbb{I}_{\{d(\boldsymbol{z}_B, \boldsymbol{X}_{j:\text{-}\ell}/\hat{\mu}_{\boldsymbol{X}_{j:\text{-}\ell}}) \le \eta\}}},$$

provided that there exists some $j \in \{1, 2, \ldots, k\}$ and $\ell \in \{1, 2, \ldots, n\}$ such that $d(\boldsymbol{z}_B, \boldsymbol{X}_{j:\text{-}\ell}/\hat{\mu}_{\boldsymbol{X}_{j:\text{-}\ell}}) \le \eta$. The validating model is

$$\max\left\{\hat{\phi}_{\text{CoL-TB}}(\gamma, y^c(\cdot), \boldsymbol{z}_B, \eta, \mathbf{X}), \gamma \in \mathbb{R}_+\right\}.$$

The co-learning operational statistic is

$$y_{\text{CoL-TB:c}}(\boldsymbol{x}) = \gamma_{\text{CoL-TB:c}}(\boldsymbol{x}) y^c(\boldsymbol{x}),\ \boldsymbol{x} \in \mathbb{R}_+^{n-1},$$

where $\gamma_{\text{CoL-TB:c}}(\boldsymbol{x}) = \arg\max\{\hat{\phi}_{\text{CoL-TB}}(\gamma, y^c(\cdot), \boldsymbol{x}/\hat{\mu}_{\boldsymbol{x}}, \eta, \mathbf{X}), \gamma \in \mathbb{R}_+\}$. As the operational statistic derived is an $(n-1)$-dimensional function, we can adopt the same treatment as that for constant boosting to derive the decision for an observed sample of each system. The empirical implementation of clustering and boosting is discussed in detail in Online Appendix D.

Theorem 7 (Co-Learning Solution with Targeted Boosting) *Under Assumptions 1, 2 and 3, if $\eta \to 0$ and $k\eta^{n-2} \to \infty$ as $k \to \infty$ (e.g., $\eta \propto k^{-\frac{1}{n-1}}$), given any base point $\boldsymbol{z}_B \in \mathcal{B}^{n-1}$, then for any $\epsilon > 0$ and a sufficiently large k,*

$$\begin{aligned} &\Pr\left\{\sup_{\gamma\in\Gamma}\left|\hat{\phi}_{\text{CoL-TB:c}}(\gamma, y^c(\cdot), \boldsymbol{z}_B, \eta, \mathbf{X}) - \mathsf{E}[\Theta^\kappa]\phi_B\left[\gamma y^c(\boldsymbol{z}_B), \boldsymbol{z}_B, F_Z\right]\right| \ge \epsilon\right\} \\ \le\ & \mathcal{O}\left(\exp\left(-\frac{2k\eta^{n-2}\epsilon^2 - 4C\sqrt{k\eta^{n-2}}\epsilon}{(b-a)^2}\right)\right). \end{aligned}$$

Moreover, if $\phi_B[\gamma y^c(\boldsymbol{z}_B), \boldsymbol{z}_B, F_Z]$ has a unique maximizer on $\gamma \in \Gamma$, denoted by $\gamma^*_{\text{CoL-TB:c}} :=$ $\text{argmax}_{\gamma\ge0}\{\phi_B\left[\gamma y^c(\boldsymbol{z}_B), \boldsymbol{z}_B, F_Z\right]\}$, then

$$y_{\text{CoL-TB:c}}(\boldsymbol{x}) \xrightarrow{p} \gamma^*_{\text{CoL-TB:c}} y^c(\boldsymbol{x}),\ \forall \boldsymbol{x} \in \mathbb{R}_+^{n-1},\ \text{as } k \to \infty.$$

Furthermore, if the maximizer of $\phi_B[\gamma y^c(\boldsymbol{z}_B), \boldsymbol{z}_B, F_Z]$ is unique and interior on Γ, $\psi(\cdot, x)$ is continuously differentiable with Lipschitz continuity and bounded $\nabla\psi(\cdot, x)$ over Γ a.s., and $\mathsf{E}[\psi(\cdot, X)]$ is strongly concave, then for any $\epsilon > 0$ and a sufficiently large k,

$$\Pr\left\{\left|y_{\text{CoL-TB:c}}(\boldsymbol{x}) - \gamma^*_{\text{CoL-TB:c}} y^c(\boldsymbol{x})\right| \ge \epsilon\right\} \le \mathcal{O}\left(\exp\left(-\frac{2k\eta^{n-2}\epsilon^2 - 4C\sqrt{k\eta^{n-2}}\epsilon}{(b-a)^2}\right)\right),\ \forall \boldsymbol{x} \in \mathbb{R}_+^{n-1}.$$

To ensure the asymptotic optimality of the co-learning solution with targeted boosting, we need to appropriately choose the clustering threshold. One way is to set $\eta = 1/k^{\frac{1}{n-1}}$, with which $\eta \to 0$ and $k\eta^{n-2} = k^{\frac{1}{n-1}} \to \infty$. With this η value, the targeted-boosting statistic $y_{\text{CoL-TB:c}}$ achieves asymptotic optimality for the combined systems and the individual systems.

We demonstrate the empirical performance of the co-learning solutions in Table 2. It is clear that co-learning can significantly improve the profit compared with directly training individual

Table 2: Performance of Co-Learning Solutions Against Directly Trained Solutions

			Directly Trained Solutions			Constant-Boosting Solutions				Targeted-Boosting Solutions			
			Ave	Stdev	Min		Ave	Stdev	Min		Ave	Stdev	Min
$n = 5$	Benchmark	Ψ_{H1}	0.076	0.03	-0.3								
	$k = 20$	Ψ_{SAA}	-0.608	0.80	-15.8	$\Psi_{\text{CoL-CB:SAA}}$	0.012	0.15	-7.3	$\Psi_{\text{CoL-TB:SAA}}$	0.010	0.14	-7.3
		$\Psi_{\text{Re-Var}}$	-0.181	0.61	-15.5	$\Psi_{\text{CoL-CB:Re-Var}}$	-0.012	0.20	-8.0	$\Psi_{\text{CoL-TB:Re-Var}}$	-0.006	0.19	-8.0
		$\Psi_{\text{Re-Std}}$	-0.284	0.66	-15.2	$\Psi_{\text{CoL-CB:Re-Std}}$	0.001	0.13	-7.1	$\Psi_{\text{CoL-TB:Re-Std}}$	0.004	0.12	-7.1
		Ψ_{OrS}	-1.461	1.62	-29.9	$\Psi_{\text{CoL-CB:OrS}}$	0.036	0.07	-3.3	$\Psi_{\text{CoL-TB:OrS}}$	0.031	0.07	-3.3
		$\Psi_{\text{SC1-Emp}}$	-0.689	0.88	-15.4	$\Psi_{\text{CoL-CB:SC1-Emp}}$	0.018	0.11	-5.4	$\Psi_{\text{CoL-TB:SC1-Emp}}$	0.016	0.11	-5.4
		$\Psi_{\text{H1-Emp}}$	-0.824	1.01	-21.0	$\Psi_{\text{CoL-CB:H1-Emp}}$	0.016	0.13	-6.6	$\Psi_{\text{CoL-TB:H1-Emp}}$	0.013	0.12	-5.8
						$\Psi_{\text{CoL-CB:SC1}}$	0.057	0.06	-1.6	$\Psi_{\text{CoL-TB:H1}}$	0.051	0.06	-1.6
	Benchmark	Ψ_{H1}	0.076	0.01	0.1								
	$k = 500$	Ψ_{SAA}	-0.616	0.17	-1.8	$\Psi_{\text{CoL-CB:SAA}}$	0.045	0.01	0.0	$\Psi_{\text{CoL-TB:SAA}}$	0.045	0.01	0.0
		$\Psi_{\text{Re-Var}}$	-0.184	0.12	-1.4	$\Psi_{\text{CoL-CB:Re-Var}}$	0.022	0.01	-0.1	$\Psi_{\text{CoL-TB:Re-Var}}$	0.026	0.01	-0.1
		$\Psi_{\text{Re-Std}}$	-0.290	0.14	-1.6	$\Psi_{\text{CoL-CB:Re-Std}}$	0.020	0.00	0.0	$\Psi_{\text{CoL-TB:Re-Std}}$	0.024	0.00	0.0
		Ψ_{OrS}	-1.473	0.33	-3.7	$\Psi_{\text{CoL-CB:OrS}}$	0.049	0.01	0.0	$\Psi_{\text{CoL-TB:OrS}}$	0.050	0.01	0.0
		$\Psi_{\text{SC1-Emp}}$	-0.698	0.21	-2.0	$\Psi_{\text{CoL-CB:SC1-Emp}}$	0.045	0.01	0.0	$\Psi_{\text{CoL-TB:SC1-Emp}}$	0.047	0.01	0.0
		$\Psi_{\text{H1-Emp}}$	-0.834	0.21	-2.2	$\Psi_{\text{CoL-CB:H1-Emp}}$	0.045	0.01	0.0	$\Psi_{\text{CoL-TB:H1-Emp}}$	0.047	0.01	0.0
						$\Psi_{\text{CoL-CB:SC1}}$	0.075	0.01	0.0	$\Psi_{\text{CoL-TB:H1}}$	0.075	0.01	0.0
$n = 10$	Benchmark	Ψ_{H1}	0.105	0.03	-0.1								
	$k = 20$	Ψ_{SAA}	-0.216	0.41	-8.3	$\Psi_{\text{CoL-CB:SAA}}$	0.037	0.08	-3.8	$\Psi_{\text{CoL-TB:SAA}}$	0.040	0.07	-3.8
		$\Psi_{\text{Re-Var}}$	-0.109	0.36	-8.4	$\Psi_{\text{CoL-CB:Re-Var}}$	0.026	0.10	-4.0	$\Psi_{\text{CoL-TB:Re-Var}}$	0.031	0.09	-3.2
		$\Psi_{\text{Re-Std}}$	-0.158	0.38	-8.4	$\Psi_{\text{CoL-CB:Re-Std}}$	0.015	0.10	-4.9	$\Psi_{\text{CoL-TB:Re-Std}}$	0.018	0.09	-4.9
		Ψ_{OrS}	-0.582	0.66	-10.4	$\Psi_{\text{CoL-CB:OrS}}$	0.049	0.08	-2.6	$\Psi_{\text{CoL-TB:OrS}}$	0.052	0.07	-2.6
		$\Psi_{\text{SC1-Emp}}$	-0.258	0.42	-7.5	$\Psi_{\text{CoL-CB:SC1-Emp}}$	0.043	0.07	-2.2	$\Psi_{\text{CoL-TB:SC1-Emp}}$	0.046	0.07	-2.2
		$\Psi_{\text{H1-Emp}}$	-0.275	0.46	-8.5	$\Psi_{\text{CoL-CB:H1-Emp}}$	0.040	0.07	-3.2	$\Psi_{\text{CoL-TB:H1-Emp}}$	0.043	0.07	-3.2
						$\Psi_{\text{CoL-CB:SC1}}$	0.093	0.04	-0.7	$\Psi_{\text{CoL-TB:H1}}$	0.089	0.04	-0.7
	Benchmark	Ψ_{H1}	0.105	0.01	0.1								
	$k = 500$	Ψ_{SAA}	-0.216	0.08	-0.9	$\Psi_{\text{CoL-CB:SAA}}$	0.052	0.01	0.0	$\Psi_{\text{CoL-TB:SAA}}$	0.055	0.01	0.0
		$\Psi_{\text{Re-Var}}$	-0.108	0.07	-0.7	$\Psi_{\text{CoL-CB:Re-Var}}$	0.044	0.01	0.0	$\Psi_{\text{CoL-TB:Re-Var}}$	0.048	0.01	0.0
		$\Psi_{\text{Re-Std}}$	-0.158	0.08	-0.8	$\Psi_{\text{CoL-CB:Re-Std}}$	0.036	0.01	0.0	$\Psi_{\text{CoL-TB:Re-Std}}$	0.038	0.01	0.0
		Ψ_{OrS}	-0.582	0.13	-1.4	$\Psi_{\text{CoL-CB:OrS}}$	0.068	0.01	0.0	$\Psi_{\text{CoL-TB:OrS}}$	0.070	0.01	0.0
		$\Psi_{\text{SC1-Emp}}$	-0.258	0.10	-0.9	$\Psi_{\text{CoL-CB:SC1-Emp}}$	0.059	0.01	0.0	$\Psi_{\text{CoL-TB:SC1-Emp}}$	0.061	0.01	0.0
		$\Psi_{\text{H1-Emp}}$	-0.274	0.09	-1.0	$\Psi_{\text{CoL-CB:H1-Emp}}$	0.056	0.01	0.0	$\Psi_{\text{CoL-TB:H1-Emp}}$	0.059	0.01	0.0
						$\Psi_{\text{CoL-CB:SC1}}$	0.105	0.01	0.1	$\Psi_{\text{CoL-TB:H1}}$	0.104	0.01	0.1

Notes. $Z \sim Gamma(0.1, 10)$ (implying $\mathsf{Cv}[Z] = 3.16$), $\Theta \sim 1.5 \times Exp(1)$, and $n = \{5, 10\}$. The profit function is $\psi(y, x) = 10 \min\{y, x\} - 3y$. The average profits over k systems $\Psi_{\text{CoL-TB:}j}$ and $\Psi_{\text{CoL-CB:}j}$ are derived from, respectively, constant-boosting and targeted-boosting candidate solution j, where $j = \text{SAA}, \text{Re-Var}, \text{Re-Std}, \text{OrS}, \text{SC1-Emp}$, or H1-Emp stands for the sample-average approximation, regularization with variance, regularization with standard deviation, order statistics, scaled operational statistic with empirical distribution, or homogeneous operational statistic with empirical distribution, respectively. The result is generated by computing the average, standard deviation, and minimum of Ψ_{j} out of $30,000$ randomly generated problem instances.

systems with their own data, even when the number of systems is small (i.e., $k = 20$). Moreover, with the same candidate solution, constant boosting may outperform targeted boosting with a small k. When pooling a large number of systems (i.e., $k = 500$), however, targeted boosting may improve the solution quality over constant boosting. This is because only a single parameter

$\gamma_{\text{CoL-CB:c}}$ is optimized in constant boosting, while the boosting parameter $\gamma_{\text{CoL-TB:c}}(\cdot)$ is a function of the base point $\boldsymbol{Z}_B$ in targeted boosting. Constant boosting avoids overfitting with a small k, while targeted boosting offers additional flexibility with a large k. This observation underscores the balance between data integration (specifying the class of operational statistics) and solution validation (using the available data) in designing the ODA framework. Similar observations are obtained in systems with a beta distribution as shown in Online Appendix D. Overall, to maximize average profit, co-learning constant-boosting solution with sample average could outperform other solutions when k is small, while co-learning targeted boosting solution with sample average could outperform others when k is large except for gamma distributions.

7 Concluding Remarks

We design solutions to facilitate learning across systems under the Operational Data Analytics (ODA) framework. As an essential departure from the existing transfer learning and data pooling approaches, the ODA learning solutions developed in this paper utilize the structural relationship between the data and the theoretically optimal solution for the focal system, and reduces the optimality gap by appropriately using the data from other related systems. Specifically, when a well-trained model is available from a relevant system, the traditional transfer learning would directly fine-tune the pre-trained solution to fit the focal system with the limited data utilizing the similarity of the data from the systems. The ODA cross-learning approach we propose, in contrast, would mimic the environment of the focal system using the data from the related system, so that we can better understand how the focal system should react to the uncertainties faced in deriving a solution. The traditional data pooling approach would focus on exploring the common statistical features of different systems to improve the aggregate system performance when combining the data sets across the systems. The ODA co-learning approach we propose, in distinction, would identify the common structure of the solution-to-data mapping across the systems to maximize not only the aggregate system performance but also the individual system performance. All in all, the development of ODA learning solutions rely heavily on our understanding of the parametric ODA solutions, and aims toward supplementing the lack of statistical characterization with the data from related systems to approach the parametric ODA solutions. We establish asymptotic optimality of the developed solutions and demonstrate their superior performance when the sample size of the focal system is extremely small.

In this study, we explore systems whose uncertainties are represented through scaling a common random parameter. One may deploy the development here to explore other common structures among related systems. For example, Feng et al. (2025) derive the parametric ODA solution for distribution families with unknown shift. Our analysis can be directly adapted to develop the cross-learning and co-learning solutions for the corresponding nonparametric setting when the demands across systems vary by their shifts. Further developments are needed to extend the learning solutions to discontinuous demand functions (den Boer and Keskin 2020). In addition, the design of solutions for dynamic decision-making requires careful evaluation to avoid the risks associated with incomplete learning (Keskin and Zeevi 2018). Recognizing the connection of the homogeneous property to Euler's equation, one may extend the analysis for systems involving multiple decisions (Feng and Shanthikumar 2023).

Our analysis establishes the theoretical foundation for cross learning and co-learning within a common distribution family, thereby providing a stepping stone toward extending such learning to similar but non-identical distribution families. However, quantifying "similarity" between distributions is nontrivial. The appropriate metric and clustering technique (e.g., Keskin et al. 2024) should be selected in accordance with the specific learning approach applied to the pooled data. Further theoretical development is needed to characterize the trade-off between the benefits of data pooling and the potential noise or bias introduced by combining data from heterogeneous distribution families.

8 Acknowledgement

Lei Li's research is partly supported by the Research Grants Council of Hong Kong [Grant 15515324] and The Hong Kong Polytechnic University under grant 1-BEAT.

References

Akcay A, Biller B, Tayur S (2011) Improved inventory targets in the presence of limited historical demand data. *Manufacturing & Service Operations Management* 13(3):297–309.

Alyafeai Z, AlShaibani MS, Ahmad I (2020) A survey on transfer learning in natural language processing. *arXiv preprint arXiv:2007.04239* .

Amrani H, Khmelnitsky E (2017) Estimation of quantiles of non-stationary demand distributions. *IISE Transactions* 49(4):381–394.

Baardman L, Levin I, Perakis G, Singhvi D (2018) Leveraging comparables for new product sales forecasting. *Production and Operations Management* 27(12):2340–2343.

Ban GY, R C (2019) The big data newsvendor: Practical insights from machine learning. *Operations Research* 67(1):90–108.

Bastani H (2021) Predicting with proxies: Transfer learning in high dimension. *Management Science* 67(5):2964–2984.

Bastani H, Simchi-Levi D, Zhu R (2022) Meta dynamic pricing: Transfer learning across experiments. *Management Science* 68(3):1865–1881.

Ben-Tal A, Hertog DD, Waegenaere AD, Melenberg B, Rennen G (2013) Robust solutions of optimization problems affected by uncertain probabilities. *Management Science* 59(2):341–357.

Berg A, Lobis M, Rölkens F, Simon P (2018) Faster fashion: How to shorten the apparel calendar. *Mckinsey & Company* Https://www.mckinsey.com/industries/retail/our-insights/faster-fashion-how-to-shorten-the-apparel-calendar.

Besbes O, Mouchtaki O (2023) How big should your data really be? Data-driven newsvendor and the transient of learning. *Management Science* 69(10):5695–6415.

Brea C, Bicanic S, Li YN, Bhardwaj S (2020) Predicting consumer demand in an unpredictable world. *Harvard Business Review* 1–7.

Chu LY, Feng Q, Shanthikumar JG, Shen ZJM, Wu J (2025) Solving the price-setting newsvendor problem with parametric operational data analytics (oda). *Management Science* 71(8):6627–6646.

Chu LY, Shanthikumar JG, Shen ZJM (2008) Solving operational statistics via a bayesian analysis. *Operations Research Letters* 36:110–116.

Cook D, Feuz KD, Krishnan NC (2013) Transfer learning for activity recognition: A survey. *Knowledge and information systems* 36(3):537–556.

den Boer AV, Keskin NB (2020) Discontinuous demand functions: Estimation and pricing. *Management Science* 66(10):4516–4534.

Duan Y, Wang K (2023) Adaptive and robust multi-task learning. *The Annals of Statistics* 51(5):2015–2039.

Feber D, Gruenewald F, Lingqvist O, Nordigaarden D, Vainberg G (2022) Perspectives on paper and forest products in 2022: How can CEOs navigate today's era of transformational change? *Mckinsey & Company* Https://mckinsey.com/industries/packaging-and-paper/our-insights/perspectives-on-paper-and-forest-products-in-2022-how-can-ceos-navigate-todays-era-of-transformational-change.

Feng Q, Shanthikumar JG (2023) The framework of parametric and nonparametric operational data analytics. *Production and Operations Management* 32(9):2685–2703.

Feng Q, Shanthikumar JG, Wu J (2025) Contextual data-integrated newsvendor solution with operational data analytics (oda). *Management Science* 71(11):9384–9403.

Freund Y, Schapire RE (1997) A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences* 55(1):119–139.

Friedman J, Hastie T, Tibshirani R (2000) Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *The annals of statistics* 28(2):337–407.

Garvin DA, Levesque LC (2008) The multiunit enterprise. *Harvard Business Review* 86(6):106–142.

Gilboa I, Schmeidler D (1989) Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18:141–153.

Gotoh J, Kim MJ, Lim AEB (2018) Robust empirical optimization is almost the same as mean–variance optimization. *Operations research letters* 46(4):448–452.

Gotoh J, Kim MJ, Lim AEB (2021) Calibration of distributionally robust empirical optimization models. *Operations Research* 69(5):1630–1650.

Gupta V, Kallus N (2022) Data pooling in stochastic optimization. *Management Science* 68(3):1595–1615.

Gupta V, Rusmevichientong P (2021) Small-data, large-scale linear optimization with uncertain objectives. *Management Science* 67(1):220–241.

Hao S, Xu Y, Mukherjee UK, Seshadri S, Souyris S, Ivanov A, Ahsen M, Sridhar P (2021) Hotspots for emerging epidemics: Multi-task and transfer learning over mobility networks. *Available at SSRN 3858274* .

Harsha P, Natarajan R, Subramanian D (2021) A prescriptive machine-learning framework to the price-setting newsvendor problem. *Informs Journal on Optimization* 3(3):227–253.

Hastie T, Tibshirani R, Friedman JH, Friedman JH (2009) *The elements of statistical learning: data mining, inference, and prediction*, volume 2 (Springer).

Hayes RH (1969) Statistical estimation problems in inventory control. *Management Science* 15(11):686–701.

Hitsch G, Hortacsu A, Lin X (2021) Prices and promotions in US retail markets. *Quantitative Marketing and Economics* 1–80.

Homem-de Mello T, Bayraksan G (2014) Monte carlo sampling-based methods for stochastic optimization. *Surveys in Operations Research and Management Science* 19(1):56–85.

Hu K, Acimovic J, Erize F, Thomas DJ, Mieghem JAV (2019) Forecasting new product life cycle curves: Practical approach and empirical analysis. *Manufacturing & Service Operations Management* 21(1):66–85.

Huh W, Janakiraman G (2008) (s, S) optimality in joint inventory-pricing control: an alternate approach. *Operations Research* 56(3):783–790.

Huh WT, Janakiraman G, Muckstadt HA, Rusmevichientong P (2009) Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Science* 55(3):404–420.

James W, Stein C (1961) Estimation of the mean of the multivariate normal distribution. *The Annals of Statistics* 9(6):1135–1151.

Janssen E, Strijbosch L, Brekelmans R (2009) Assessing the effects of using demand parameters estimates in inventory control and improving the performance using a correction function. *International Journal of Production Economics* 118(1):34–42.

Kahn KB (2006) *New product forecasting: an applied approach* (Routledge).

Keskin NB, Li Y, Sunar N (2024) Data-driven clustering and feature-based retail electricity pricing with smart meters. *Operations Research* .

Keskin NB, Zeevi A (2018) On incomplete learning and certainty-equivalence control. *Operations Research* 66(4):1136–1167.

Levi R, Pal M, Roundy RO, Shmoys DB (2007a) Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* 32(2):284–302.

Levi R, Roundy RO, Shmoys DB (2007b) Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research* 32(4):821–839.

Lim AEB, Shanthikumar JG, Shen ZJM (2006) Model uncertainty, robust optimization and learning. *Informs Tutorial* 2006:66–94.

Lin S, Chen F, Li Y, Shen ZJM (2022) Data-driven newsvendor rgularized by a profit risk constraint. *Production and Operations Management* 31(4):1630–1644.

Liyanage L, Shanthikumar JG (2005) A practical inventory control policy using operational statistics. *Operations research letters* 33:341–348.

Long DZ, Sim M, Zhou M (2023) Robust satisficing. *Operations Research* 71(1):61–82.

Ma Y, Luo G, Zeng X, Chen A (2012) Transfer learning for cross-company software defect prediction. *Information and Software Technology* 54(3):248–256.

Oroojlooyjadid A, Nazari M, Snyder LV, Takáč M (2022) A deep *q*-network for the beer game: Deep reinforcement learning for inventory optimization. *Manufacturing & Service Operations Management* 24(1):285–304.

Pan SJ, Yang Q (2009) A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* 22(10):1345–1359.

Qin H, Simchi-Levi D, Wang L (2022) Data-driven approximation schemes for joint pricing and inventory control models. *Management Science* 68(9):6355–7064.

Qin Z, Tang X, Jiao Y, Zhang F, Xu Z, Zhu H, Ye J (2020) Ride-hailing order dispatching at didi via reinforcement learning. *INFORMS Journal on Applied Analytics* 50(5):272–286.

Ramamurthy V, G Shanthikumar J, Shen ZJM (2012) Inventory policy with parametric demand: Operational statistics, linear correction, and regression. *Production and Operations Management* 21(2):291–308.

Shao L, Zhu F, Li X (2014) Transfer learning for visual categorization: A survey. *IEEE transactions on neural networks and learning systems* 26(5):1019–1034.

Sim M, Tang Q, Zhou M, Zhu T (2025) The analytics of robust satisficing: predict, optimize, satisfice, then fortify. *Operations Research* 73(5):2708–2728.

Weiss K, Khoshgoftaar TM, Wang D (2016) A survey of transfer learning. *Journal of Big data* 3(1):1–40.

Xu K, Bastani H (2025) Multitask learning and bandits via robust statistics. *Management Science* 71(9):7752–7773.

Zhao L, Chakrabarti D, Muthuraman K (2022) Unified classical and robust optimization for least squares. *Available at SSRN 3182422* .

Zhuang F, Qi Z, Duan K, Xi D, Zhu Y, Zhu H, Xiong H, He Q (2020) A comprehensive survey on transfer learning. *Proceedings of the IEEE* 109(1):43–76.