This is the accepted manuscript of the following article: Jinzhi Bu, David Simchi-Levi, Li Wang (2022) Offline Pricing and Demand Learning with Censored Data. Management Science 69(2):885-903, which has been published in final form at https://doi.org/10.1287/mnsc.2022.4382.

Offline Pricing and Demand Learning with Censored Data

Jinzhi Bu
Department of Logistics and Maritime Studies, Faculty of Business, The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong, jinzhi.bu@polyu.edu.hk

David Simchi-Levi
Institute for Data, Systems, and Society, Department of Civil and Environmental Engineering, and Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139, dslevi@mit.edu

Li Wang
Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139, li_w@mit.edu

We study a single product pricing problem with demand censoring in an offline data-driven setting. In this problem, a retailer has a finite amount of inventory, and faces a random demand that is price sensitive in a linear fashion with unknown price sensitivity and base demand distribution. Any unsatisfied demand that exceeds the inventory level is lost and unobservable. We assume that the retailer has access to an offline dataset consisting of triples of historical price, inventory level and potentially censored sales quantity. The retailer's objective is to use the offline dataset to find an optimal price, maximizing her expected revenue with finite inventories. Due to demand censoring in the offline data, we show that the existence of near-optimal algorithms in a data-driven problem – which we call *problem identifiability* – is not always guaranteed. We develop a necessary and sufficient condition for problem identifiability by comparing the solutions to two *distributionally robust optimization* problems. We propose a novel data-driven algorithm that hedges against the distributional uncertainty arising from censored data, with provable finite-sample performance guarantees regardless of problem identifiability and offline data quality. Specifically, we prove that, for identifiable problems, the proposed algorithm is near-optimal, and, for unidentifiable problems, its worst-case revenue loss approaches the best-achievable minimax revenue loss that any data-driven algorithm must incur. Numerical experiments demonstrate that our proposed algorithm is highly effective and significantly improves both the expected and worst-case revenues compared with three regression-based algorithms.

Key words: price optimization, demand censoring, data-driven algorithm, offline learning, finite-sample analysis

1. Introduction

For many businesses, the importance of a good pricing strategy cannot be overstated: the price of a product has a direct influence on customers' willingness to purchase, affecting the business's sales volumes, revenue figures, and the bottom line. Pricing is also an integral part of a business's operations, since the optimization of pricing decisions in practice is often conducted with many other operational factors in consideration. For example, the presence of a brand strategy or a government regulation may set a range for a product's price, and the availability of inventory may put a hard limit on the sales volumes. In this paper, we distill such fundamental pricing problems faced by businesses every day into a single-product pricing problem with finite inventory over a price range.

For any price optimization model, the relationship between price and demand is an essential input. In reality, such a relationship is hardly known exactly by the businesses. Nevertheless, with the rapid advances in information technology, many businesses have been collecting large amounts of historical pricing and sales data from past selling seasons, to help them learn the demand-price relationship and make *data-driven* pricing decisions (Foxman 2013). One crucial challenge related to historical sales data is that, when stock-out happens, customers whose demands are unsatisfied may simply leave the system without any purchase records, and in this case, the retailer is unable to observe the amount of lost-sales. This phenomenon is known as *demand censoring*, which could happen in both brick-and-mortar and e-commerce settings. Therefore, when sales quantities are capped by the levels of inventory available for sale, they become censored demand data. If demand censoring is not accounted for properly – for example, simply ignoring this effect and treating sales data as uncensored demand – it often leads to biased and inconsistent estimation on demand levels, deluding businesses to make suboptimal pricing decisions. In this paper, we study algorithms that are specially designed to handle historical censored demand data, and to make good data-driven pricing decisions in the presence of finite inventories.

1.1. Model Overview and Research Question

We consider a retailer who wants to find an optimal price over a given price range for a single product with a finite level of inventory. Any unsatisfied demand that exceeds the inventory level is lost and unobservable. We assume the demand is random and price sensitive in a linear fashion, which has been commonly adopted and studied in the literature (cf. den Boer 2015). Although the retailer does not know the exact price sensitivity or the distribution of the random base demand, she has access to a historical dataset collected from past selling seasons, called *offline data*, in the format of (price, inventory, sales) triples, where the sales volumes are independent samples from the random demands censored by the inventory levels. The retailer's objective is to devise an algorithm that uses offline data to help her make pricing decisions, maximizing her expected revenue with finite inventories.

Ideally, we want to build data-driven algorithms that achieve small expected revenue loss compared to the optimal expected revenue with high probability. We consider an algorithm near-optimal, if it can guarantee an expected revenue loss to any infinitesimal level with arbitrarily high probability, when given an offline dataset of sufficiently large size. However, due to different levels of demand censoring, not all data have the same "quality": a highly censored dataset contains little information about the true demand, making it extremely difficult and even intractable for data-driven algorithms to be near-optimal. In this paper, we aim to address the following two questions:

(1) How does demand censoring affect the *existence* of near-optimal algorithms in our data-driven pricing problems?

(2) How to design a data-driven algorithm that can be applied practically regardless of data qualities, with *provable* performance guarantees?

1.2. Key Challenges

In the presence of upper bounds on the observable sales quantities, some demand information that is crucial to the optimization of pricing may be inevitably lost in a censored demand dataset. To illustrate this, consider a *deterministic* linear demand model with given price sensitivity but unknown base demand quantity. Suppose a retailer observes her entire inventory was sold out on a certain day, i.e., the sales quantity equaled the inventory level. This suggests that her current price may be too low, and the retailer could possibly get a higher revenue by raising the price. However, such information from this censored demand observation is only directional, in that the retailer is unable to compute the optimal price adjustment, since the actual demand level was censored and not observable. For a *random* linear demand model, this means that demand censoring completely blocks the distributional information about demand quantities above the inventory level, inducing additional uncertainty that we need to pay extra attention to when designing data-driven algorithms.

Intuitively, a dataset that reveals more distributional information about the random demand is considered to have a better quality, and data-driven algorithms that take it as an input are more likely to have good performance guarantees, while a dataset with little to none demand information may disqualify all data-driven algorithms from being near-optimal. However, we must point out that data quality is in fact a relative measure. In a problem with a smaller inventory level or price range, the task of finding an optimal price may require less demand information, and hence a dataset previously believed having insufficient information may turn out to be adequate to empower near-optimal data-driven algorithms.

Overall, we believe that, not only is it important for us to study algorithms that perform well with censored data of good quality, but it is also necessary to understand what performance guarantees data-driven algorithms may and could obtain when given potentially low-quality datasets.

From the perspective of algorithmic design, there are three main challenges:

(1) Historical sales data are biased samples of the true demands due to demand censoring. Hence, a naive application of linear regression typically leads to a biased and inconsistent estimate on

the price sensitivity. Therefore, the first challenge is to design a procedure to form an unbiased estimate of the price sensitivity in a linear demand model using censored demand samples.

(2) In our pricing problem, where sales volumes are capped by the available inventory level, an optimal price that maximizes the expected revenue not only depends on the price sensitivity in the demand function, but is also decided by the *distribution* of the random base demand and the given inventory level. However, due to censored data, a portion of the distribution of random base demand is always blocked, bringing uncertainty in computing the expected revenues. Hence, the second challenge is to deal with the additional ambiguity of the base demand distribution.

(3) The third challenge is to build a data-driven algorithm that is "universal" and "robust": it is desired to be applied to all censored demand datasets and have provable performance guarantees regardless of data qualities. We also require the algorithm to be completely *agnostic* to any information that may be unrealistic to obtain in real business settings, for example, the support and distribution family of the random base demand.

1.3. Main Results and Contributions

We summarize our main results and contributions as follows.

Ambiguity set and worst-case revenue loss. In this paper, we take a novel approach to measure the performance of data-driven algorithms in the presence of demand censoring. We define an *ambiguity set* of distributions, a concept borrowed from *distributionally robust optimization* (DRO), to succinctly capture the distributional knowledge about the random base demand that can be exploited from the offline data, as well as the uncertainty created by the information loss in censored data. As no data-driven algorithms may further narrow down the ambiguity set without additional information, we measure an algorithm's performance by its *worst-case revenue loss*, which is defined as the maximum gap between the expected revenues generated by the optimal price and the algorithm's price among all candidate distributions in the ambiguity set. We also define a *minimax revenue loss* as the minimum worst-case revenue loss that any pricing policy must incur.

Necessary and sufficient condition for problem identifiability. We distinguish between two classes of data-driven problems: 1) identifiable problems, in which there exists a data-driven algorithm whose worst-case revenue loss converges to zero as the size of offline data increases; 2) unidentifiable problems, in which the worst-case revenue loss of *any* algorithm does not converge to zero. We provide a necessary and sufficient condition for the problem identifiability, which states that the data-driven problem is identifiable if and only if the solutions to two DRO problems are identical. This condition is further decomposed into three different cases depending on the location of the unconstrained maximizers of the two DRO problems. Moreover, we also show that there exists some threshold for a quality metric that we define for the offline censored data, referred to as the *observable boundary*, such that the problem is identifiable when the observable boundary is above that threshold.

Development of a data-driven algorithm. We design a data-driven algorithm that hedges against the distributional uncertainty caused by censored data, and is robust to any possible distributions in the ambiguity set. The algorithm takes three major steps. First, we estimate the price sensitivity based on the observation that the true price sensitivity is a solution to a quantile-based linear regression problem. By focusing on the empirical quantiles of the random demands whose left-hand parts are unlikely to be censored, we construct a consistent estimate for the price sensitivity with guaranteed finite-sample probability bound. Second, by leveraging the structural property of the so-called *optimistic* and *pessimistic* revenues, we compute their empirical counterparts using the sample average approximation (SAA) approach. Third, the algorithm suggests a price based on our exact characterization on the worst-case revenue loss through the optimistic and pessimistic revenues, and their empirical versions obtained from the second step.

Finite-sample probability bound. We provide theoretical performance guarantees of our algorithm for both identifiable and unidentifiable problems. When the underlying problem is identifiable, we establish a finite-sample high probability bound for the worst-case revenue loss of our algorithm being within any given small error, translating to the number of historical samples needed to guarantee a pre-specified error of ε with a confidence level of $1-\delta$ in the order of $O(\varepsilon^{-2}\log(\delta^{-1}))$. When the underlying problem is unidentifiable, we show that, although the worst-case revenue loss for any algorithm is always positive, our algorithm is guaranteed to be close to the minimax revenue loss within any small error, with the same high probability bound. Therefore, regardless of the data quality, our algorithm attains the best-achievable performance up to any accuracy and confidence levels. In the case of unidentifiable problems, we also prove that the algorithm can successfully detect unidentifiability with high probability by comparing the solutions to two empirical DRO problems.

Numerical study. The numerical results demonstrate that our proposed algorithm significantly improves three baseline algorithms based on simple linear regression, with smaller relative optimality gap evaluated under the distribution we use to generate the data, smaller worst-case revenue loss over the entire ambiguity set, and more accurate estimation on the price sensitivity. We also compare our proposed algorithm with its modified version that applies the so-called KM estimator, and find that the modified algorithm enjoys a slightly better empirical performance when the data size is small, while performs almost identically to the original algorithm when the data size becomes larger. Finally, we investigate the effects of the data quality and inventory level on the empirical performance of our algorithm.

1.4. Related Literature

The effect of demand censoring has been studied extensively for many pricing and inventory models in online demand learning, where sales data are generated by the decision maker's actions on the fly. Some papers study the repeated newsvendor problem with censored demand, e.g., Huh and Rusmevichientong (2009), Huh et al. (2011), Besbes and Muharremoglu (2013), Lugosi et al. (2017). Later studies analyze more complicated multi-period inventory and pricing management problems with censored demand, e.g., Huh et al. (2009), Zhang et al. (2018), Agrawal and Jia (2019), Zhang et al. (2020), Yuan et al. (2019), Chen and Chao (2020), Chen and Shi (2019), Chen (2019), Chen et al. (2020a), Chen et al. (2020b). Our paper differs from these works by considering an offline data-driven problem, where the historical data are given exogenously. Moreover, in contrast with the online setting where the demand distribution can be learned by adjusting pricing and inventory decisions, in the offline setting, partial information are always lost due to censoring effect, and it is not necessary that a near-optimal pricing policy can always be identified, leading to the dichotomy of identifiable problems and unidentifiable problems.

This paper is also related to data-driven algorithms for pricing and inventory models in an offline learning setting, where the entire dataset is available before the algorithm starts. Levi et al. (2007) study both single-period and multi-period inventory problems, and propose algorithms by approximating the true demand distribution with an empirical distribution based on the SAA approach. They develop bounds on the number of required samples to guarantee that the expected cost of their algorithm is arbitrarily close to the expected cost of the optimal policy. Levi et al. (2015) improve the bound in Levi et al. (2007) for the newsvendor model. Cheung and Simchi-Levi (2019) study the multi-period capacitated inventory system, and analyze the performance of the SAA method by comparing the empirical dynamic program constructed from SAA with the original dynamic program. Qin et al. (2019) study the data-driven joint inventory and pricing model, and develop an algorithm with guaranteed sample complexity results. A recent work by Ban and Rudin (2019) applies machine learning algorithms to the data-driven newsvendor with feature information. Ban et al. (2020) compare across three different approaches for solving data-driven newsvendor, and study the impact of model misspecification. However, all of the above papers assume that demand samples can be fully observed, and there is no issue of demand censoring in the historical data, while we consider the scenario that only the sales data are available to the retailer, which are potentially censored due to limited inventories.

To the best of our knowledge, Ban (2020) is the only paper that studies offline learning with censored demand in OR/MS literature. The author considers a multi-period inventory system with fixed ordering cost and unknown demand distribution, and develops a nonparametric estimation procedure for the (S, s) policy which is consistent and asympotically normal. In this paper, we consider a different context of a single-period pricing problem, and take a different angle to study this data-driven pricing problem by investigating the effect of demand censoring on the identifiability of a near-optimal algorithm.

In a recent paper by Feng and Shanthikumar (2018), the authors discuss how the relevant concepts emerged with Big Data can be incorporated into demand management and manufacturing. In particular, the authors suggest a new way of learning under demand censoring through the *data-integrated personalized demand model*. As the first step to approach the problem of how to conduct demand learning and pricing with censored data, in this paper, we do not consider the more complicated personalized demand model with additional covariates information. Instead, we focus on a simpler setting with a basic data structure consisting of the price, sales and inventory information. Although simple, our model still fits into the scenarios, especially in brick-and-mortar stores, where the retailer collects demand data in an aggregate way, and personalized pricing is restricted due to the use of common price tags, fairness concern, etc. Under this base model, this paper mainly addresses two theoretical questions proposed in Section 1.1, and also provides practical guidance for the retailer regarding how to apply a DRO framework to make pricing decisions with censored data. We leave the study for more complicated personalized demand models as future research. See Section 7 for further discussions.

Censored regression models have been studied extensively in econometrics and statistics literature. Since the pioneering work by Tobin (1958), a variety of estimation methods have been developed for both parametric and nonparametric models, see, e.g., Cox (1972), Miller (1976), Buckley and James (1979), Koul et al. (1981), Powell (1984), Fernandez (1986), Lewbel and Linton (2002). Under certain model assumptions, these works typically propose estimates for the regression coefficients, and prove asymptotic properties, e.g., consistency and asymptotic normality. Although we share a similar task of model parameter estimation as a part of our problem, the focus of this paper is very different. Our work is oriented towards a revenue maximization problem with limited inventory, for which the optimal price depends on both the price sensitivity and the distribution of the random base demand. Therefore, our objective is to design a pricing algorithm that not only estimates the unknown price sensitivity accurately, but also hedges against the extra distributional uncertainty of base demand arising from demand censoring by leveraging the specific structural property of the pricing problem. We also adapt one of the most commonly-used algorithms in statistics proposed by Buckley and James (1979) to estimate the demand parameters and revenue function in our setting, and numerically compare the performances of this algorithm with ours. The numerical result shows that our algorithms perform consistently better than the benchmark with more accurate estimate on the price sensitivity.

1.5. Structure and Notations

This paper is organized as follows. In Section 2, we introduce the basic model and describe the data-driven pricing problem with censored demand. In Section 3, we define the concept of problem identifiability, and provide a necessary and sufficient condition for problem identifiability. Then we develop a data-driven algorithm in Section 4 and present its theoretical performance in Section 5. A numerical study is provided in Section 6, and we conclude the paper in Section 7. Most of the technical proofs are presented in the online Appendix.

For any integer $n \in \mathbb{N}$, we use $[n]$ to denote the set $\{1, 2, \cdots, n\}$. For any event A, $\mathbb{1}_{\{A\}}$ is the indicator function of A, which equals one if A holds, and equals zero if A does not hold. For two non-negative functions $f(x)$ and $g(x)$ defined on $\mathbb{R}^+$, we use $f(x) = O(g(x))$ to represent that there exists some constant $C > 0$ such that $f(x) \le Cg(x)$ for any $x > 0$. The cumulative distribution function of a random variable is abbreviated as "c.d.f.". For any given discrete random variable with c.d.f. $G(\cdot)$, its breakpoints are defined as those $x_0 \in \mathbb{R}$ such that $\lim_{x \downarrow x_0} G(x) < G(x_0)$.

2. Model Formulation

We consider a retailer who sells a single product in the face of a finite amount of inventory. The amount of inventory available for sale, denoted by a nonnegative real number $y \in \mathbb{R}^+$, is given exogenously as a pre-determined quantity to the retailer, and she charges a unit price p from the feasible set $\in [\underline{p}, \overline{p}] \subset (0, \infty)$ for each product sold, up to y units. We assume that the market demand for the product with the price set to p is $D(p) := a - bp + \eta$, where a is the deterministic base demand, b is the price sensitivity belonging to the range $[\underline{b}, \overline{b}] \subset (0, \infty)$, and η is the random noise with mean zero. For convenience, we use ξ to denote the random base demand, i.e., $\xi := a + \eta$ and $D(p) = \xi - bp$, and let $F_\xi(\cdot)$ denote the c.d.f. of ξ. In this paper, we focus on the linear demand model since it is commonly used in practice. Moreover, the linear model is also shown to perform considerably well in the literature under different problem contexts (e.g., Dawes 1979 in clinical prediction, Besbes and Zeevi 2015 in dynamic pricing with online learning). We assume that $D(p)$ is nonnegative with probability one for any feasible price $p \in [\underline{p}, \overline{p}]$. Any extra demands over y units are lost and unobservable to the retailer, and the observed sales quantity is the minimum between the market demand $D(p)$ and the inventory level y. For any given inventory level y, the retailer's goal is to find an optimal price p^* that maximizes the expected revenue $R(p) := p \cdot \mathbb{E}[\min\{D(p), y\}]$, i.e.,

$$p^* := \max\Big\{\arg\max_{p \in [\underline{p}, \overline{p}]} R(p)\Big\}. \tag{1}$$

We also denote the unconstrained maximizer of the expected revenue by $p^\dagger$, i.e., $p^\dagger := \max\big\{\arg\max_{p \in \mathbb{R}} R(p; y)\big\}$. It is worth noting that $R(p)$, p^*, and $p^\dagger$ are all functions of the inventory level y, but for simplicity, we drop their dependence on y in the notation.

Suppose the retailer knows the price sensitivity b and the distribution of ξ, the optimal price can be computed efficiently by solving a concave optimization problem according to the following proposition. We call it the *full-information problem.*

PROPOSITION 1. *For any given inventory level $y \geq 0$, the expected revenue $R(p)$ is concave in $p \in [\underline{p}, \overline{p}]$, and the optimal price p^* is the projection of $p^\dagger$ to $[\underline{p}, \overline{p}]$, i.e., $p^* = Proj(p^\dagger, [\underline{p}, \overline{p}]) = \min\{\max\{p^\dagger, \underline{p}\}, \overline{p}\}$.*

In this paper, we do not assume the optimal price p^* to be an interior point of the price interval $[\underline{p}, \overline{p}]$; rather, it can be any price in $[\underline{p}, \overline{p}]$, including the boundary $\{\underline{p}, \overline{p}\}$. In many business settings, the price range is given to the pricing manager as a *hard constraint* and may not be adjusted arbitrarily, due to company policies, competitor's pricing, or even government regulations. Therefore, one cannot expand the range of feasible prices trying to make p^* an interior point. Moreover, in the presence of exogenous inventory constraints, it is possible that for some inventory level y, the global maximizer $p^\dagger$ does not belong to $[\underline{p}, \overline{p}]$, and thus p^* is obtained by projecting to the boundary of $[\underline{p}, \overline{p}]$. Therefore, we believe the absence of such an interior-point assumption for the optimal price is closer to the business practice.

Furthermore, we do not require the retailer to be able to distinguish between the case that the demand is exactly the same as the inventory level and the case that the demand is censored by the inventory level. This is because the sales quantities under both cases are the same, and without additional information, it is impossible to know whether $D(p) = y$ or $D(p) > y$. In other words, we only require the retailer to be able to fully observe demand $D(p)$ when it is strictly less than the inventory level y.

2.1. Data-Driven Problem with Censored Demand

In reality, the retailer does not know b or the distribution of ξ. As a result, she would not be able to compute the optimal price p^*. Fortunately, the retailer typically has access to some historical data, also called offline data, to help her learn the demand model and make pricing decisions by solving a *data-driven problem.*

2.1.1. Offline data with demand censoring. We assume that the retailer has experimented K different historical prices $p_i \in [\underline{p}, \overline{p}]$, where $K \geq 2$, each with a possibly different historical inventory level y_i. For each price-inventory pair (p_i, y_i), the retailer has N_i independent sales samples $S_i^1 \leq S_i^2 \leq \ldots \leq S_i^{N_i}$ sorted in ascending order without loss of generality. Each S_i^j is generated by the random demand at price p_i, subject to demand censoring by the inventory level y_i, i.e.,

$$S_i^j := \min\{D_i^j, y_i\} = \min\{\xi_i^j - p_i, y_i\},$$

where each ξ_i^j is an independent sample of the random base demand ξ. Therefore, we can describe the offline data using the set $\mathcal{S} = \{(p_i, y_i, S_i^j) : i \in [K], j \in [N_i]\}$, with each element consisting of three components: unit price, inventory level, and sales quantity.

A data-driven algorithm $\mathcal{A}$ takes the offline dataset $\mathcal{S}$ and the known intervals, $[\underline{p}, \overline{p}]$ and $[\underline{b}, \overline{b}]$, as inputs, and outputs a pricing policy in the form of a function $p^{\mathcal{A}}(\cdot) : \mathbb{R}^+ \to [\underline{p}, \overline{p}]$ that maps a given inventory level y to a feasible price $p^{\mathcal{A}}(y) \in [\underline{p}, \overline{p}]$. Again, for simplicity, we use the notation $p^{\mathcal{A}}$ instead of $p^{\mathcal{A}}(y)$, as there is no ambiguity caused by dropping the dependence on y. The optimality gap of the algorithm $\mathcal{A}$ is defined as the difference between the expected revenues generated by the optimal price and by the algorithm's price, i.e., $R(p^*) - R(p^{\mathcal{A}})$.

2.1.2. Data metrics. We introduce two important data metrics: one related to the "size" and another related to the "quality" of the offline data. The first metric is $N := \min_{i \in [K]} N_i$, which we use to denote the *size* of the dataset $\mathcal{S}$. The second metric is $\lambda := \max_{i \in [K]}\{y_i + bp_i\}$, which we call the *observable boundary* of the distribution of ξ. Since a sample ξ_i^j of the random base demand ξ is "theoretically observable" if and only if $\xi_i^j - bp_i < y_i$, the larger the value of $y_i + bp_i$ is, the more *true* samples of ξ we can observe. Therefore, the upper bound for the values of ξ that are theoretically observable is given by $\lambda := \max_{i \in [K]}\{y_i + bp_i\}$. Note that the second metric λ requires the knowledge of the price sensitivity b and cannot be directly computed from the historical data.

We also define a quantity that is closely related to λ and concerns the degree to which historical demands are censored by the associated inventory levels. For each $i \in [K]$, let γ_i denote the probability that the random variable $D(p_i)$ is strictly less than y_i, i.e., $\gamma_i = \mathbb{P}[\xi < y_i + bp_i]$. We define γ be the maximal value among all γ_i, i.e., $\gamma = \max_{i \in [K]} \gamma_i$, or equivalently, $\gamma = \mathbb{P}[\xi < \lambda]$. Again, we point out that, like λ, γ and γ_i's are not computable without knowing b and the underlying demand distribution. Nevertheless, the empirical version of γ_i, defined by $\hat{\gamma}_i := \frac{1}{N_i} \sum_{j=1}^{N_i} \mathbb{1}_{\{S_i^j < y_i\}}$, can be computed directly from the offline data. Note that $\hat{\gamma}_i$ is an unbiased sample of γ_i, i.e., $\mathbb{E}[\hat{\gamma}_i] = \gamma_i$. We assume that $\hat{\gamma}_i > 0$ for each $i \in [K]$, i.e., there exists at least one uncensored demand observation for each price-inventory pair in the dataset, which can be easily verified from the data. This also implicitly assumes that $\gamma_i > 0$ for each $i \in [K]$, and thus, $\gamma > 0$.

For the data-driven pricing problem defined above, we make the following assumption.

Assumption 1. *We assume that on the left-hand side of λ, ξ is either a discrete random variable with finite support $\{\beta_i : i \in [M]\}$ for some $M \in \mathbb{N}^+$ and $\beta_1 < \beta_2 < \ldots < \beta_M < \lambda$, or a continuous random variable with probability density distribution $f_\xi(\cdot)$. Besides, there exists some constant $\underline{g} > 0$ such that if ξ is discrete in $[0, \lambda)$, $\frac{F_\xi(\beta_{k+1}) - F_\xi(\beta_k)}{\beta_{k+1} - \beta_k} \geq \underline{g}$ for each $k \in [M-1]$, and if ξ is continuous in $[0, \lambda)$, $f_\xi(x) \geq \underline{g}$ for all $x \in [0, \lambda)$.*

If ξ is discrete, the existence of $\underline{g}$ in Assumption 1 is automatically guaranteed by letting $\underline{g} := \min_{k\in[M-1]} \frac{F_\xi(\beta_{k+1})-F_\xi(\beta_k)}{\beta_{k+1}-\beta_k}$. If ξ is continuous, Assumption 1 imposes a lower bound on its probability density function, which is satisfied by many commonly used continuous distributions, e.g., uniform distribution, truncated normal distribution, and is also widely assumed in the literature, e.g., Chen et al. (2019). Intuitively, $\underline{g}$ measures how flat the c.d.f. $F_\xi(\cdot)$ is, and the smaller $\underline{g}$ is, the flatter $F_\xi(\cdot)$ will be. It can be easily verified that under Assumption 1, if ξ is discrete, $F_\xi(\beta_j)-F_\xi(\beta_i) \geq \underline{g}(\beta_j-\beta_i)$ for any $1 \leq i \leq j \leq M$, and if ξ is continuous, $F_\xi(y)-F_\xi(x) \geq \underline{g}(y-x)$ for any $0 \leq x \leq y < \lambda$. This property enables learning the price sensitivity and will be used in later analysis. It is also worth noting that we do not impose any assumption on the distribution of ξ on the right-hand side of λ, which is allowed to have an arbitrary shape instead.

3. Identifiability of Near-Optimal Data-Driven Algorithms

In this section, we introduce the concept of *problem identifiability* and characterize an exact condition for problem identifiability. The necessity of introducing this concept is as follows. In our censored-demand setting, the retailer's ability to estimate the distribution of ξ heavily relies on the observable boundary λ of the offline data. By definition, λ is the maximum values of ξ samples that can be uncovered from the data, which implies that all samples of ξ that are greater than or equal to λ are unobservable. Therefore, even with infinite number of samples, no data-driven algorithms are able to learn the c.d.f. $F_\xi(\cdot)$ of ξ over the region $[\lambda,\infty)$. As a result, this negatively affects the retailer's ability to find near-optimal prices to any accuracy and confidence levels through data-driven algorithms, and encourages the discussion on the existence of near-optimal algorithms in the data-driven problems – which we call *problem identifiability*.

3.1. Problem Identifiability and Ambiguity Set of Distributions

As described above, demand censoring brings a crucial challenge to the data-driven problems regarding the problem identifiability. Due to demand censoring, the retailer is completely agnostic about the c.d.f. $F_\xi(\cdot)$ of ξ over the region $[\lambda,\infty)$, and she is not able to further reduce her uncertainty about $F_\xi(\cdot)$ than limiting all possible distributions to an *ambiguity set* of

$$\mathcal{F}(\lambda, F_\xi) = \Big\{ F \text{ is a c.d.f.: } F(x) = F_\xi(x), \forall x < \lambda \Big\},$$

which includes all c.d.f.s that coincide with the true unknown c.d.f. $F_\xi(\cdot)$ over the region $(-\infty,\lambda)$, but may take arbitrary shapes over the region $[\lambda,\infty)$. In some sense, the ambiguity set $\mathcal{F}(\lambda, F_\xi)$ represents all the knowledge about the distribution of the random base demand ξ that a data-driven algorithm can possibly learn from the censored offline dataset $\mathcal{S}$. For notational ease, we simply use $\mathcal{F}$ to denote $\mathcal{F}(\lambda, F_\xi)$.

We point out that the concept of ambiguity set was introduced in the literature on *distributionally robust optimization* (DRO), and the DRO framework is designed to accommodate all possible probability measures with the encoded available distributional information from historical data (see, e.g., Rahimian and Mehrotra 2019 for a recent review on the studies of DRO). In this paper, we borrow this terminology and construct a family of base demand distributions which are consistent with the prior knowledge from the offline data, as well as to capture the information loss caused by demand censoring.

3.1.1. Definition of problem identifiability. For any c.d.f. $F \in \mathcal{F}$, we define $R_F(p)$ as the expected revenue with the expectation taken with respect to $F(\cdot)$, and $p_F^\dagger$ and p_F^* as the optimal price to $R_F(p)$ in $\mathbb{R}$ and $[\underline{p}, \overline{p}]$ respectively. Similar to Proposition 1, $p_F^* = \mathrm{Proj}(p_F^\dagger, [\underline{p}, \overline{p}])$. We now formalize the definition of *problem identifiability.*

Definition 1. For any inventory level y, price range $[\underline{p}, \overline{p}]$, and base demand distribution F_ξ, the data-driven problem (defined in Section 2.1) with a dataset $\mathcal{S}$ of parameter λ is called *identifiable*, if there exists some algorithm $\mathcal{A}$ with the output price $p^\mathcal{A}$ such that, for any $\varepsilon > 0$,

$$\lim_{N\to\infty} \mathbb{P}\left[\max_{F\in\mathcal{F}} \left\{R_F(p_F^*) - R_F(p^\mathcal{A})\right\} < \varepsilon\right] = 1. \tag{2}$$

If no such algorithm exists, the data-driven problem in Section 2.1 is called *unidentifiable*.

In the above definition, we measure the performance of a data-driven algorithm by $\max_{F\in\mathcal{F}}\{R_F(p_F^*) - R_F(p^\mathcal{A})\}$, which is the maximum revenue loss among all possible c.d.f.s in the ambiguity set $\mathcal{F}$, and is referred to as the *worst-case revenue loss*. The reason why we choose this measure is as follows. For any data-driven algorithm, without extra information, each $F \in \mathcal{F}$ is equally likely to be the true c.d.f., and when faced with a potentially adversarial "nature" arbitrarily choosing a distribution $F \in \mathcal{F}$, no algorithms may offer revenue loss guarantees that are better than $\max_{F\in\mathcal{F}}\{R_F(p_F^*) - R_F(p^\mathcal{A})\}$. An algorithm is called *near-optimal* if it satisfies (2). It's worth noting that since the ambiguity set $\mathcal{F}$ and the price sensitivity b are unknown, whether the problem is identifiable is also unknown to the seller. Thus, the results developed in this section is theoretical rather than practical in determining whether the problem is identifiable or not.

The worst-case framework is also borrowed from the DRO literature, in which any feasible solution is evaluated under the worst-case expectation with respect to all possible distributions in the ambiguity set. Similarly, the performance of a data-driven algorithm in our problem setting is measured by the worst-case revenue loss. We also point out that since the algorithm $p^\mathcal{A}$ is data-dependent, when defining problem identifiability, the probability measure in (2) is taken over the distribution of all the offline sales data $\{S_i^j : i \in [K], j \in [N_i]\}$, under the assumption that $\{(p_i, y_i) : i \in [K]\}$ are fixed constants.

3.1.2. Optimistic and pessimistic revenues. With the ambiguity set $\mathcal{F}$, we can define the *optimistic revenue* $R_{\max}(p)$ and *pessimistic revenue* $R_{\min}(p)$, by maximizing and minimizing the expected revenue in the ambiguity set respectively. That is, for any $p \in \mathbb{R}$,

$$R_{\max}(p) = \max_{F \in \mathcal{F}} R_F(p), \quad R_{\min}(p) = \min_{F \in \mathcal{F}} R_F(p).$$

Furthermore, we define the *optimistic price* and *pessimistic price* as the maximizers of the optimistic and pessimistic revenues:

$$\begin{cases} p^*_{\max} = \max \left\{ \arg\max_{p \in [\underline{p}, \overline{p}]} R_{\max}(p) \right\} \\ p^*_{\min} = \max \left\{ \arg\max_{p \in [\underline{p}, \overline{p}]} R_{\min}(p) \right\} \end{cases}.$$

Thus, $p^*_{\max}$ and $p^*_{\min}$ are the solutions to two distributionally robust optimization problems: the *maximax* problem, and the *maximin* problem.

We also have the following closed-form expression for $R_{\max}(\cdot)$ and $R_{\min}(\cdot)$. Let $\tilde{p} := \frac{\lambda - y}{b}$.

PROPOSITION 2. *For any* $p \in [\underline{p}, \overline{p}]$,

$$R_{\max}(p) = \mathbb{1}_{\{p < \tilde{p}\}} R_{F_\xi}(p) + \mathbb{1}_{\{p \geq \tilde{p}\}} R_1(p),$$
$$R_{\min}(p) = \mathbb{1}_{\{p < \tilde{p}\}} R_{F_\xi}(p) + \mathbb{1}_{\{p \geq \tilde{p}\}} R_2(p),$$

where $R_1(p) := p(-\gamma bp + \mathbb{E}_{F_\xi}[\min\{\xi, \lambda\}] + (1-\gamma)(y - \lambda))$ *and* $R_2(p) := p(-bp + \mathbb{E}_{F_\xi}[\min\{\xi, \lambda\}])$.

It is worth noting that the construction of $R_1(\cdot)$ and $R_2(\cdot)$, as well as $R_{\max}(\cdot)$ and $R_{\min}(\cdot)$, does not require distributional knowledge about the random base demand over the region $[\lambda, \infty)$. This means that, with more data, it is possible for a data-driven algorithm to approximate the optimistic and pessimistic revenues with higher accuracy. We defer the proof of Proposition 2 to online Appendix C.1.

3.2. Necessary and Sufficient Condition

In this section, we present the first main result of this paper: a necessary and sufficient condition for problem identifiability, in Theorem 1.

THEOREM 1 **(Necessary and sufficient condition for identifiability)**. *For any inventory level* y*, price range* $[\underline{p}, \overline{p}]$*, and base demand distribution* F_ξ*, a data-driven problem (defined in Section 2.1) with a dataset* $\mathcal{S}$ *of parameter* λ *is identifiable if and only if*

$$p^*_{\max} = p^*_{\min}. \tag{3}$$

Theorem 1 states that, if the optimistic price $p^*_{\max}$ is equal to the pessimistic price $p^*_{\min}$, the data-driven problem is identifiable, i.e., there exists a data-driven algorithm that can guarantee near-optimal expected revenue to any accuracy and confidence levels, as the size metric N of the offline data $\mathcal{S}$ increases; if $p^*_{\max} \neq p^*_{\min}$, no such algorithms exist. To shed some light on the intuition of Theorem 1, we present the following proposition that provides upper and lower bounds on the optimal price p^*_F for any c.d.f. F in the ambiguity set.

PROPOSITION 3. *For any $F \in \mathcal{F}$, the following inequality holds: $p^*_{\min} \leq p^*_F \leq p^*_{\max}$.*

Proposition 3 asserts that the optimal price p^*_F according to any distribution $F \in \mathcal{F}$ is between the optimistic and pessimistic prices. This further implies that, although the true optimal price $p^*_{F_\xi}$ cannot be directly estimated from data due to the distributional information loss from demand censoring, we can estimate its upper and lower bounds, $p^*_{\max}$ and $p^*_{\min}$, using only information that is available in the offline data.

Therefore, when $p^*_{\max} = p^*_{\min}$, the optimal price p^*_F must equal to $p^*_{\max}$ and $p^*_{\min}$, regardless of which distribution $F \in \mathcal{F}$ the optimal price corresponds to. As a result, as shown in Theorem 1, the data-driven problem in this case is identifiable, i.e., it is possible for an algorithm to identify a near-optimal price, since the estimation of $p^*_{\max}$ and $p^*_{\min}$ does not require information beyond what the data may provide and it is expected to keep improving as data size grows. Nevertheless, when $p^*_{\max} \neq p^*_{\min}$, it is unfortunate that the uncertainty induced by demand censoring does not allow algorithmic revenue guarantees to any accuracy and confidence levels, as the optimal price p^*_F could take many possible values ranging from $p^*_{\min}$ to $p^*_{\max}$.

In the next proposition, we discuss in more details on when identifiability occurs (or equivalently, when $p^*_{\max} = p^*_{\min}$ by Theorem 1). Define $p^\dagger_{\max} := \max \arg\max_{p \in \mathbb{R}} R_{\max}(p)$ and $p^\dagger_{\min} := \max \arg\max_{p \in \mathbb{R}} R_{\min}(p)$.

PROPOSITION 4. *The condition $p^*_{\max} = p^*_{\min}$ holds if and only if one of the following three cases happens: (i) $p^\dagger_{\max} = p^\dagger_{\min} \leq \tilde{p}$; (ii) $p^\dagger_{\max} > p^\dagger_{\min} \geq \tilde{p}$ and $p^\dagger_{\max} \leq \underline{p}$; (iii) $p^\dagger_{\max} > p^\dagger_{\min} \geq \tilde{p}$ and $p^\dagger_{\min} \geq \overline{p}$.*

Proposition 4 describes the three cases for identifiability, which are also depicted in Figure 1. In case (i) when $p^\dagger_{\max} = p^\dagger_{\min} \leq \tilde{p}$, since $R_{\max}(p) = R(p) = R_{\min}(p)$ when $p \leq \tilde{p}$ from Proposition 2, it can be easily verified from concavity of $R(\cdot)$ that the three unconstrained maximizers $p^\dagger_{\max}$, $p^\dagger$ and $p^\dagger_{\min}$ are identical. Thus, the corresponding constrained optimal prices, $p^*_{\max}$, p^* and $p^*_{\min}$, obtained by projecting $p^\dagger_{\max}$, $p^\dagger$, and $p^\dagger_{\min}$ onto the feasible price range $[\underline{p}, \overline{p}]$ respectively, are also identical, leading to an identifiable problem. In case (ii) and case (iii) when $p^\dagger_{\max} > p^\dagger_{\min} \geq \tilde{p}$, although the unconstrained maximizers exceed $\tilde{p}$, which may result in the relationship $p^\dagger_{\min} < p^\dagger < p^\dagger_{\max}$, the data-driven problem can still be identifiable if $p^\dagger_{\max} \leq \underline{p}$ or $p^\dagger_{\min} \geq \overline{p}$. This is because after projecting $p^\dagger_{\max}$, $p^\dagger$ and $p^\dagger_{\min}$ onto $[\underline{p}, \overline{p}]$, the constrained optimal prices $p^*_{\max}$, p^* and $p^*_{\min}$ turn out to be identical again. In fact, $p^*_{\max} = p^* = p^*_{\min} = \underline{p}$ in case (ii), and $p^*_{\max} = p^* = p^*_{\min} = \overline{p}$ in case (iii). Note that in

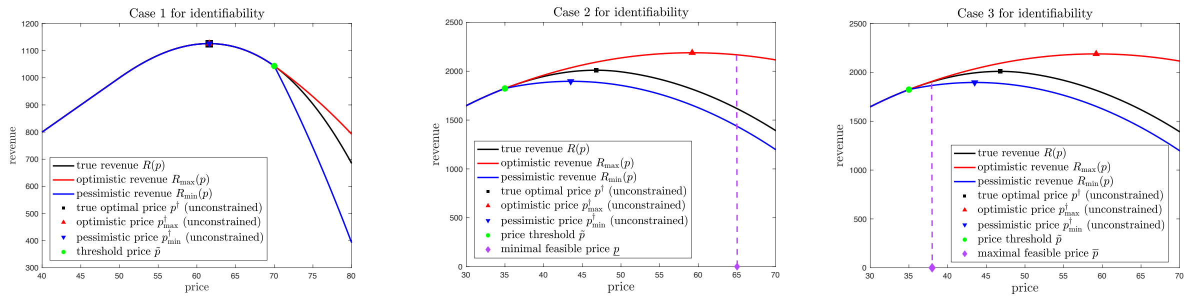

Figure 1 **Three cases for identifiability in Proposition 4.**

this paper, for the reason explained after Proposition 1, we do not assume the optimal price p^* is an interior point of $[\underline{p}, \overline{p}]$. The absence of this assumption makes cases (ii) and (iii) possible to occur, and the data-driven problem identifiable under all the above three cases.

3.2.1. Relations to data quality metric. As the data quality metric, the observable boundary λ governs the set of distributions to be included in the ambiguity set $\mathcal{F}$, which determines the values of the optimistic and pessimistic prices $p^*_{\max}$ and $p^*_{\min}$. Therefore, λ has a direct impact on the identifiability of data-driven problems, and we summarize it in Proposition 5.

PROPOSITION 5. *Consider two data-driven problems derived from the same underlying pricing problem, but with two different offline datasets $\mathcal{S}$ and $\mathcal{S}'$, whose observable boundaries λ and λ' satisfy $\lambda < \lambda'$. If the problem with dataset $\mathcal{S}$ is identifiable, then the problem with dataset $\mathcal{S}'$ is also identifiable.*

Proposition 5 indicates that the set of observable boundaries under which the data-driven problem is identifiable must be an interval. Therefore, there exists some threshold $\tilde{\lambda}$ such that if $\lambda > \tilde{\lambda}$, the data-driven problem is identifiable, and if $\lambda < \tilde{\lambda}$, the data-driven problem is unidentifiable. In other words, the identifiability of a data-driven problem corresponds to whether the observable boundary λ of the offline data $\mathcal{S}$ is sufficiently large. The intuition is that, as λ becomes larger, the retailer is able to observe a wider range of samples from the random base demand ξ, gaining more information about the underlying demand model, and making it more likely to identify a near-optimal price. We defer the proof of Proposition 5 to online Appendix C.5.

Proposition 5 also has the following implication for the retailer to set pricing and inventory levels in the offline data-collection process. Recall that $\lambda = \max_{i\in[K]}\{y_i + bp_i\}$, which combined with Proposition 5, implies that to increase the chance of obtaining an identifiable data-driven pricing problem, the retailer needs to set higher inventory levels as well as higher prices in the offline data-collection stage. This is quite intuitive since under the same price, demands are less likely to be censored with a higher inventory level, and under the same inventory level, the expected demand is lower with a higher price and thus is more likely to be fully observed.

3.3. Minimax Revenue Loss in Unidentifiable Problems

For the unidentifiable data-driven problems, it is certainly unfortunate to know that no near-optimal algorithms can be found, i.e., no data-driven algorithms can guarantee expected revenue to any accuracy and confidence levels. However, it is still important for the retailer to know what levels of revenue loss that she may possibly get by applying a data-driven algorithm, and whether "good" algorithms can be designed to limit such a revenue loss, albeit not to *any* accuracy levels.

We define the *minimax revenue loss* Δ as the minimum expected revenue loss that any data-driven algorithms must incur, when faced with an adversary deciding which distribution $F \in \mathcal{F}$ to choose, i.e.,

$$\Delta := \min_{p \in [\underline{p}, \overline{p}]} \max_{F \in \mathcal{F}} \{R_F(p_F^*) - R_F(p)\}.$$

In Proposition 6, we provide a lower bound on Δ when $p^*_{\max} \neq p^*_{\min}$.

Proposition 6. *Suppose $p^*_{\max} \neq p^*_{\min}$. Then the minimax revenue loss Δ is lower bounded by*

$$\Delta \geq \frac{\gamma b}{(\sqrt{\gamma}+1)^2} \left(p^*_{\max} - p^*_{\min}\right)^2. \tag{4}$$

Recall that we assume $\gamma > 0$ in Section 2. Then Proposition 6 shows that when the problem is unidentifiable, no data-driven algorithms may guarantee a lower expected revenue loss than a positive constant $\frac{\gamma b}{(\sqrt{\gamma}+1)^2}\left(p^*_{\max} - p^*_{\min}\right)^2$, thus disqualifying all algorithms from being near-optimal. We also remark that the lower bound in the RHS of (4) is tight, in that there exists some instance for the data-driven problem for which the minimax revenue loss Δ equals $\frac{\gamma b}{(\sqrt{\gamma}+1)^2}(p^*_{\max} - p^*_{\min})^2$. In fact, if $\tilde{p} < p^{\dagger}_{\min}$, and the global maximizers of $R_1(\cdot)$ and $R_2(\cdot)$ are within the price range $[\underline{p}, \overline{p}]$, we can prove that $p^*_{\max} = \arg\max_{p\in\mathbb{R}} R_1(p)$, $p^*_{\min} = \arg\max_{p\in\mathbb{R}} R_2(p)$, leading to $\Delta = \frac{\gamma b}{(\sqrt{\gamma}+1)^2}(p^*_{\max} - p^*_{\min})^2$. The proof of Proposition 6 is provided in online Appendix C.6.

Nonetheless, we believe that the nonexistence of near-optimal algorithms in unidentifiable problems should not be interpreted as a discouragement of using data-driven algorithms in pricing problems with demand censoring. Rather, it calls for more awareness and consideration of the unique challenges brought by demand censoring, when designing and applying data-driven algorithms in pricing problems.

4. Data-Driven Algorithm under Censored Demand

In this section, we present a Data-Driven Algorithm under Censored Demand (D2ACD) for the price optimization problem with finite inventories.

4.1. Algorithm Description

As discussed in Section 3, when demand observations are censored, some distributional information is inevitably lost in the offline data, and Proposition 6 shows that, due to the information loss, it is impossible to accurately estimate the value of the optimal price in some cases, resulting in a strictly positive minimax revenue loss. D2ACD addresses this challenge by avoiding direct estimations of the true revenue functions and the true optimal prices; instead, the algorithm estimates the optimistic and pessimistic revenues, and suggests prices by leveraging the structural property of the worst-case revenue loss in our pricing problem, and striking a balance between the two estimated revenues.

The description of D2ACD is presented in Algorithm 1, which consists of three major steps. In terms of the notation, we use $\texttt{SAA}(\mathcal{S})$ to denote the empirical c.d.f. constructed from sample average approximation with the given samples $\mathcal{S}$.

Although the steps in Algorithm 1 involve many details, the ideas are quite intuitive and can be largely considered as a series of two general tasks: estimation and optimization. In the first task (Step 1 and Step 2), D2ACD performs the estimation for price sensitivity as well as optimistic and pessimistic revenues; in the second task (Step 3), price optimization is performed based on the values of the estimated revenue functions. We next provide a detailed explanation on Steps 1-3 in the following three subsections respectively and discuss some important properties guaranteed by D2ACD.

4.2. Illustration for Step 1 in D2ACD

As mentioned earlier, a naive linear regression approach will lead to biased estimate for the price sensitivity. To overcome this challenge, we make the following observation. For any $u \in (0,1)$, let $F_i^{-1}(u)$ and $F_j^{-1}(u)$ be the u-quantile for random demands $\xi - bp_i$ and $\xi - bp_j$ respectively. Then the price sensitivity b is the solution to the following regression:

$$b = \arg\min_{b' \in [\underline{b}, \bar{b}]} \min_{a' \in \mathbb{R}} \sum_{i=1}^{K} \left(F_i^{-1}(u) - (a' - b'p_i)\right)^2. \tag{5}$$

To see why (5) holds, we notice that for any $i, j \in [K]$, $\xi - bp_i$ and $\xi - bp_j$ have two otherwise identical distributions that are shifted $b|p_i - p_j|$ from each other. Hence, $b(p_i - p_j) = -(F_i^{-1}(u) - F_j^{-1}(u))$. By simple algebra, we can verify the following equation satisfied by b:

$$b \sum_{i=1}^{K} (p_i - \frac{1}{K}\sum_{j=1}^{K} p_j)^2 = -\sum_{i=1}^{K} (p_i - \frac{1}{K}\sum_{j=1}^{K} p_j) \cdot (F_i^{-1}(u) - \frac{1}{K}\sum_{j=1}^{K} F_j^{-1}(u)),$$

or equivalently, b satisfies (5). With this observation, the key idea is to replace the true quantile $F_i^{-1}(u)$ in (5) with an empirical quantile below which the demand data are unlikely to be censored.

For each $i \in [K]$, we first construct an empirical c.d.f. $\hat{F}_i^{\mathrm{SAA}}$ based on the sales data under price p_i. Then we define a quantile function $\hat{F}_i^{-1}$ as the inverse of a piecewise linear approximation to $\hat{F}_i^{SAA}$.

Algorithm 1: Data-Driven Algorithm under Censored Demand (D2ACD)

Input: y, $[\underline{p},\overline{p}]$, $[\underline{b},\overline{b}]$, $\left\{(p_i,y_i,S_i^j): i\in[K], j\in[N_i]\right\}$, $\hat{\gamma}_{\min} := \min_{i\in[K]}\hat{\gamma}_i$

Step 1: estimating price sensitivity b

for $i\in[K]$ **do**

  $\hat{F}_i^{SAA}(\cdot) := \texttt{SAA}(\{S_i^j : j\in[N_i]\})$;

  Sort the breakpoints of $\hat{F}_i^{SAA}(\cdot)$ in an increasing order, and denote them by $\hat{\beta}_{i,1} < \hat{\beta}_{i,2} < \cdots < \hat{\beta}_{i,\hat{M}_i}$;

  $\hat{F}_i^{-1}(\cdot)$: modified quantile function of $\hat{F}_i^{SAA}(\cdot)$ defined as

$$\hat{F}_i^{-1}(x) := \begin{cases} \frac{\left(x-\hat{F}_i^{SAA}(\hat{\beta}_{i,k})\right)\cdot\left(\hat{\beta}_{i,k+1}-\hat{\beta}_{i,k}\right)}{\hat{F}_i^{SAA}(\hat{\beta}_{i,k+1})-\hat{F}_i^{SAA}(\hat{\beta}_{i,k})} + \hat{\beta}_{i,k}, & \text{if } x\in(\hat{F}_i^{SAA}(\hat{\beta}_{i,k}),\hat{F}_i^{SAA}(\hat{\beta}_{i,k+1})] \text{ for } k\in[\hat{M}_i-1]; \\ \hat{\beta}_{i,1}, & \text{if } x\le \hat{F}_i^{SAA}(\hat{\beta}_{i,1}); \end{cases}$$

end for

Solve the linear regression: $\hat{b} = \arg\min_{b'\in[\underline{b},\overline{b}]}\min_{a'\in\mathbb{R}}\sum_{i=1}^K\left(\hat{F}_i^{-1}(\hat{\gamma}_{\min}) - (a'-b'p_i)\right)^2$;

Step 2: estimating optimistic and pessimistic revenues $R_{\max}(\cdot)$ **and** $R_{\min}(\cdot)$

$\hat{i}^* := \arg\max_{i\in[K]}\{y_i+\hat{b}p_i\}$;

$\hat{\lambda} := y_{\hat{i}^*}+\hat{b}p_{\hat{i}^*}$, $\hat{\gamma} := \max_{i\in[K]}\hat{\gamma}_i$;

$\hat{R}(p) := p\cdot\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{S_{\hat{i}^*}^j+\hat{b}p_{\hat{i}^*}-\hat{b}p, y\}$;

$\hat{R}_1(p) := p\cdot\left(-\hat{\gamma}\hat{b}p+\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}(S_{\hat{i}^*}^j+\hat{b}p_{\hat{i}^*})+(1-\hat{\gamma})(y-\hat{\lambda})\right)$;

$\hat{R}_2(p) := p\cdot\left(-\hat{b}p+\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}(S_{\hat{i}^*}^j+\hat{b}p_{\hat{i}^*})\right)$;

$\hat{\tilde{p}} := (\hat{\lambda}-y)/\hat{b}$;

$\hat{R}_{\max}(p) := \mathbb{1}_{\{p<\hat{\tilde{p}}\}}\hat{R}(p)+\mathbb{1}_{\{p\ge\hat{\tilde{p}}\}}\hat{R}_1(p)$;

$\hat{R}_{\min}(p) := \mathbb{1}_{\{p<\hat{\tilde{p}}\}}\hat{R}(p)+\mathbb{1}_{\{p\ge\hat{\tilde{p}}\}}\hat{R}_2(p)$;

Step 3: computing algorithm's price p^{D2ACD}

$\hat{p}^*_{\max} := \arg\max_{p\in[\underline{p},\overline{p}]}\hat{R}_{\max}(p)$, $\hat{p}^*_{\min} := \arg\max_{p\in[\underline{p},\overline{p}]}\hat{R}_{\min}(p)$;

$\hat{W}_{\max}(p) := \hat{R}_{\max}(\hat{p}^*_{\max})-\hat{R}_{\max}(p)$;

$\hat{W}_{\min}(p) := \hat{R}_{\min}(\hat{p}^*_{\min})-\hat{R}_{\min}(p)$;

$p^{\mathsf{D2ACD}} = \arg\min_{p\in[\underline{p},\overline{p}]}\max\left\{\hat{W}_{\max}(p),\hat{W}_{\min}(p)\right\}$;

Output: p^{D2ACD}

This piecewise linear function is constructed by drawing straight lines joining any two consecutive breakpoints of $\hat{F}_i^{SAA}$. The reason for which we take the piecewise linear approximation as a bridge is to guarantee $\hat{F}_i^{-1}$ is uniquely defined on $[\hat{\beta}_{i,1},\hat{\beta}_{i,\hat{M}_i}]$, which will facilitate our analysis. The figures for $\hat{F}_i^{SAA}$ and $\hat{F}_i^{-1}$ are depicted in the blue step function and the red piecewise linear function in Figure 2 respectively. Since the demand data under p_i are uncensored below y_i, from the definition of $\hat{\gamma}_{\min}$, the demand is also uncensored below the empirical $\hat{\gamma}_{\min}$-quantile $\hat{F}_i^{-1}(\hat{\gamma}_{\min})$. Therefore, $\hat{F}_i^{-1}(\hat{\gamma}_{\min})$ intuitively approximates the true quantile (after piecewise linear approximation to the true c.d.f. of $\xi - bp_i$) with high accuracy.

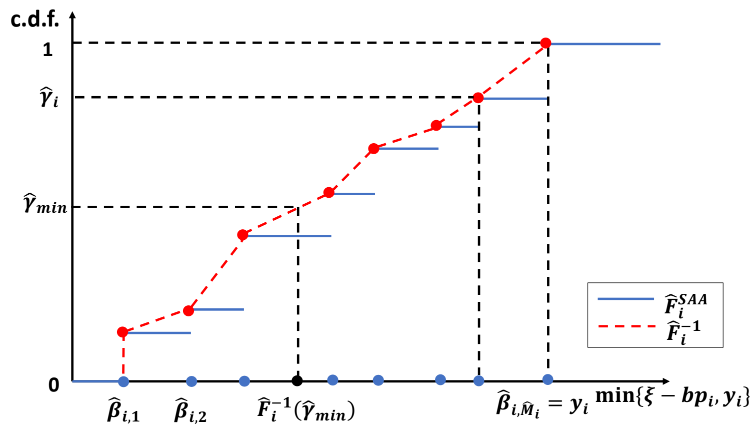

Figure 2 Computation of the modified empirical quantile under price p_i in Step 1 of D2ACD.

The following lemma establishes a finite-sample probability bound on the estimate $\hat{b}$.

LEMMA 1 **(Convergence of price sensitivity).** *For any $\alpha > 0$, $|\hat{b} - b| \le \alpha$ with probability at least* $1 - 2Ke^{-\frac{1}{9K}\sum_{i=1}^{K}(p_i - \frac{1}{K}\sum_{j=1}^{K}p_j)^2 \underline{g}^2 N\alpha^2}$.

Lemma 1 guarantees that regardless of whether the data-driven problem is identifiable or not, $\hat{b}$ is always consistent. This is ensured by taking advantage of the uncensored parts in the sales data when constructing $\hat{b}$. Besides, the probability bound in Lemma 1 increases in the price variance $\frac{1}{K}\sum_{i=1}^{K}(p_i - \frac{1}{K}\sum_{j=1}^{K}p_j)^2$. This is intuitive since as the offline prices become more dispersive, i.e., the variance of offline prices increases, the estimation for b is more accurate, and the probability bound is closer to one. In addition, the bound in Lemma 1 also increases in $\underline{g}$ (defined in Assumption 1), showing that as the true c.d.f. of random base demand ξ becomes steeper, the probability bound gets closer to one.

To prove Lemma 1, we need to carefully bound the distance between the empirical $\hat{\gamma}_{\min}$-quantile and the true $\hat{\gamma}_{\min}$-quantile of random demand $\xi - bp_i$ for each $i \in [K]$. Suppose to the contrary, these two quantiles are "far away" from each other. By leveraging the property guaranteed by $\underline{g}$, we can show that the empirical c.d.f. is also bounded away from the true c.d.f. at some uncensored point. This leads to a contradiction with the so-called Dvoretzk-Kiefer-Wolfowitz (DKW) inequality, which claims a uniform high-probability bound on the distance between the empirical c.d.f. and the true c.d.f. for the uncensored part. The formal statement of DKW inequality and the proof of Lemma 1 are provided in online Appendices A and D.1 respectively.

4.3. Illustration for Step 2 in D2ACD

In Step 2, we estimate the empirical optimistic revenue $\hat{R}_{\max}(\cdot)$ and pessimistic revenue $\hat{R}_{\min}(\cdot)$. According to Proposition 2, we need to compute the empirical revenue $\hat{R}(\cdot)$ and two empirical quadratic functions $\hat{R}_1(\cdot)$ and $\hat{R}_2(\cdot)$. To compute $\hat{R}(p)$, we replace the unknown expected sales $\mathbb{E}[\min\{\xi - bp, y\}]$ with its empirical counterpart $\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}} \min\{S_{\hat{i}^*}^j + \hat{b}p_{\hat{i}^*} - \hat{b}p, y\}$ by replacing b with $\hat{b}$ and adding back the estimated pricing effect $\hat{b}p_i$. To compute $\hat{R}_1(\cdot)$ and $\hat{R}_2(\cdot)$, we replace b, γ, λ and $\mathbb{E}[\min\{\xi, \lambda\}]$ in the expressions of $R_1(\cdot)$ and $R_2(\cdot)$ in Proposition 2 with their empirical version $\hat{b}$, $\hat{\gamma}$, $\hat{\lambda}$, and $\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}} \min\{S_{\hat{i}^*}^j + \hat{b}p_{\hat{i}^*}, \hat{\lambda}\} = \frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}} (S_{\hat{i}^*}^j + \hat{b}p_{\hat{i}^*})$, respectively.

We emphasize that since the price sensitivity is unknown and the sales data are censored, the empirical revenue function $\hat{R}(p)$ we obtain from the above procedure converges to the true revenue function $R(p)$ only when $p \le \min\{\tilde{p}, \hat{\tilde{p}}\}$. However, since the optimistic and pessimistic revenues depend on the distribution of ξ only through its left portion $[0, \lambda)$, which can be estimated from the historical data accurately when the data size becomes larger, we can then prove that $\hat{R}_{\max}(\cdot)$ and $\hat{R}_{\min}(\cdot)$ converge to $R_{\max}(\cdot)$ and $R_{\min}(\cdot)$ respectively on the entire price range $[\underline{p}, \overline{p}]$. This is formally stated in the following lemma, where we use the notations $\overline{\lambda} := \max_{i\in[K]}\{y_i + \overline{b}p_i\}$, $C_1 := 5\overline{p}^2 + \frac{(2y+3\overline{b}\overline{p}+2\overline{\lambda})^2}{\underline{b}^2}$, and $C_2 := \overline{p}(y + \overline{\lambda} + \overline{p}\overline{b})$.

Lemma 2 **(Convergence of empirical optimistic and pessimistic revenues)**. *For any* $\alpha > 0$, $\sup_{p\in[\underline{p},\overline{p}]}\left|\hat{R}_{\max}(p) - R_{\max}(p)\right| \le \alpha$ *and* $\sup_{p\in[\underline{p},\overline{p}]}\left|\hat{R}_{\min}(p) - R_{\min}(p)\right| \le \alpha$ *with probability at least* $1 - 2Ke^{-\frac{1}{36(y+\overline{b}\overline{p})^2\overline{p}^2}N\alpha^2} - 2Ke^{-\frac{\frac{1}{K}\sum_{i=1}^K (p_i - \frac{1}{K}\sum_{j=1}^K p_j)^2 \underline{g}^2}{36C_1^2}N\alpha^2} - 2Ke^{-\frac{1}{18C_2^2}N\alpha^2} - 2Ke^{-\frac{1}{18\overline{p}^2\overline{\lambda}^2}N\alpha^2}$.

Lemma 2 shows that with high probability, $\hat{R}_{\max}(\cdot)$ converges to $R_{\max}(\cdot)$ and $\hat{R}_{\min}(\cdot)$ converges to $R_{\min}(\cdot)$ *uniformly* in range $p \in [\underline{p}, \overline{p}]$. This property of uniform approximation will be crucial to the analysis of the algorithm's performance in Section 5. In particular, to prove this result, we invoke a new concentration inequality established in Lemma 14 of Qin et al. (2019), which bounds the probability $\mathbb{P}[\sup_{y\in\mathcal{Y}} |\frac{1}{n}\sum_{i=1}^n f(X_i, y) - \mathbb{E}[f(X_1, y)]| \le \alpha]$ for *i.i.d.* random variables $X_1, X_2, \ldots, X_n$, and bounded monotone functions $f(x, y)$ under any given accuracy level $\alpha > 0$ (see Lemma A.3 in online Appendix A). The proof of Lemma 2 is in online Appendix D.2.

Remark 1. When estimating $R_{\max}(\cdot)$ and $R_{\min}(\cdot)$, we only utilize the sales data from the single pair $(p_{\hat{i}^*}, y_{\hat{i}^*})$. The reason is that under the current SAA method, replacing $\hat{i}^*$ with other indices or pooling all the data together with potentially different censoring inventory levels leads to biased estimate for the distribution of ξ. To fully use the historical data while still maintaining consistency, one possible approach is to replace the SAA method with the so-called KM estimator introduced by Kaplan and Meier (1958) to estimate the distribution of ξ. In Section 6, we propose a new algorithm D2ACD-KM modified from D2ACD by using the KM estimator in Step 2, and compare its numerical performance with D2ACD. Since the theoretical analysis of D2ACD-KM in our problem context is much more complicated and difficult than D2ACD, we still focus on D2ACD in the presentation and analysis of the main results of this paper.

4.4. Illustration for Step 3 in D2ACD

In Step 3, we compute the algorithm's price p^{D2ACD}. The following proposition provides a crucial foundation for achieving the goal of designing a robust algorithm that hedges against the distributional uncertainty.

PROPOSITION 7. *For any price $p \in [\underline{p}, \overline{p}]$, the worst-case revenue loss satisfies*

$$\max_{F \in \mathcal{F}} \{R_F(p_F^*) - R_F(p)\} = \max\{R_{\max}(p_{\max}^*) - R_{\max}(p),\, R_{\min}(p_{\min}^*) - R_{\min}(p)\}\,. \tag{6}$$

Proposition 7 provides a closed-form expression for the worst-case revenue loss for any pricing decision $p \in [\underline{p}, \overline{p}]$ through the optimistic revenue and pessimistic revenue. This indicates that to evaluate the worst-case revenue loss within the ambiguity set, it suffices to consider the best-possible distribution and the worst-possible distribution. In this case, the minimax revenue loss is simplified to

$$\Delta = \min_{p \in [\underline{p}, \overline{p}]} \max\{R_{\max}(p_{\max}^*) - R_{\max}(p), R_{\min}(p_{\min}^*) - R_{\min}(p)\}. \tag{7}$$

Motivated from the above identity (7), we set the price p^{D2ACD} to minimize the empirical counterpart of the RHS in (7), i.e., $\max\{\hat{R}_{\max}(\hat{p}_{\max}^*) - \hat{R}_{\max}(p), \hat{R}_{\min}(\hat{p}_{\min}^*) - \hat{R}_{\min}(p)\}$. Since the empirical optimistic and pessimistic revenues $\hat{R}_{\max}(\cdot)$ and $\hat{R}_{\min}(\cdot)$ approximate the true revenues $R_{\max}(\cdot)$ and $R_{\min}(\cdot)$ with high accuracy as guaranteed in Lemma 2, intuitively, the worst-case revenue loss of the algorithm's price could also be close to the minimax revenue loss with high probability.

5. Finite-Sample Probability Bound of D2ACD Algorithm

In this section, we present the theoretical probability bound of D2ACD, and we consider both cases in terms of problem identifiability. Specifically, in Section 5.1, we show that when the optimistic and pessimistic prices are equal, i.e., $p_{\max}^* = p_{\min}^*$, D2ACD is near-optimal. Moreover, in Section 5.2, we study the case that $p_{\max}^* \neq p_{\min}^*$, and show D2ACD can approach the best-achievable performance within any accuracy and confidence levels.

5.1. Identifiable Problems

One of the main results of this paper is the development of an exact condition for problem identifiability of the data-driven pricing problems with demand censoring. In Section 3, we show that there exists a near-optimal algorithm if and only if $p_{\max}^* = p_{\min}^*$ (Theorem 1). As the "only if" direction is already proved in Proposition 6, we in this section show affirmatively that our proposed algorithm D2ACD is near-optimal when $p_{\max}^* = p_{\min}^*$, closing the "if" direction in the proof of Theorem 1.

Theorem 2 **(Sample complexity for identifiable problems)**. *Suppose* $p^*_{\max} = p^*_{\min}$. *For any* $\varepsilon > 0$ *and* $\delta \in (0,1)$, *if* $N \geq N(\varepsilon, \delta)$, *then with probability at least* $1-\delta$, *the worst-case revenue loss of* D2ACD *in Section 4 is no greater than* ε, *i.e.,*

$$\mathbb{P}\left[\max_{F\in\mathcal{F}}\left\{R_F(p^*_F) - R_F(p^{\mathsf{D2ACD}})\right\} \leq \varepsilon\right] \geq 1-\delta,$$

where $N(\varepsilon,\delta) = C\frac{1}{\varepsilon^2}\log\frac{8K}{\delta} = O(\frac{1}{\varepsilon^2}\log\frac{1}{\delta})$ *for some constant* C.

Theorem 2 shows that when the level of confidence probability is $1-\delta$, the number of samples needed to guarantee the worst-case revenue loss is within the error of ε, also referred to as the *sample complexity*, is in the order of $O(\frac{1}{\varepsilon^2}\log\frac{1}{\delta})$, which is quadratic in $1/\varepsilon$ and logarithmic in $1/\delta$. Such a dependency of the sample complexity on ε and δ is consistent with the existing results of Levi et al. (2015) in a different context of the data-driven newsvendor problem without censoring. The result in Theorem 2 can also be equivalently stated as follows: for any given size metric N and error level $\varepsilon > 0$, then $\mathbb{P}\left[\max_{F\in\mathcal{F}}\left\{R_F(p^*_F) - R_F(p^{\mathsf{D2ACD}})\right\} \leq \varepsilon\right] \geq 1 - 8K\exp(-\frac{N\varepsilon^2}{C})$, where the resulting probability bound has an exponential rate in N. The exact expression of constant C can be found in the proof of Theorem 2 in Appendix.

5.2. Unidentifiable Problems

When $p^*_{\max} \neq p^*_{\min}$, it has been shown in Proposition 6 that no data-driven algorithm can achieve a worst-case revenue loss less than the lower bound in inequality (4), and the minimax revenue loss Δ is strictly positive. Nevertheless, the revenue loss incurred by our algorithm is proven to converge to the best-achievable revenue loss Δ as the size metric N increases, with guaranteed finite-sample probability bound. The result is stated as follows.

Theorem 3 **(Sample complexity for unidentifiable problems)**. *Suppose* $p^*_{\max} \neq p^*_{\min}$. *For any* $\varepsilon > 0$ *and* $\delta \in (0,1)$, *if* $N \geq N(\varepsilon, \delta)$, *then with probability at least* $1-\delta$, *the worst-case revenue loss of* D2ACD *in Section 4 is within* ε *distance of the minimax revenue loss, i.e.,*

$$\mathbb{P}\left[\max_{F\in\mathcal{F}}\left\{R_F(p^*_F) - R_F(p^{\mathsf{D2ACD}})\right\} - \Delta \leq \varepsilon\right] \geq 1-\delta,$$

where $N(\varepsilon,\delta)$ *is defined in Theorem 2.*

Similar to the case of identifiable problems, for the unidentifiable problems, the sample complexity to guarantee the gap between the worst-case revenue loss under D2ACD and the minimax revenue loss Δ is within the error of ε is also in the order of $O\left(\frac{1}{\varepsilon^2}\log\frac{1}{\delta}\right)$. Correspondingly, the convergence rate of the probability bound for event $\max_{F\in\mathcal{F}}\left\{R_F(p^*_F) - R_F(p^{\mathsf{D2ACD}})\right\} - \Delta \leq \varepsilon$ is also exponentially fast in the size metric N. Since Δ is the best-possible revenue loss when one has full knowledge about the ambiguity set, Theorem 3 demonstrates that even if in the case of unidentifiable problem, D2ACD achieves the near-best performance. The proof of Theorem 3 is given in Appendix.

One related question the retailer may be faced with in practice is: when the underlying data-driven problem is unidentifiable, can we apply the algorithm to detect it? This is answered in the following proposition.

PROPOSITION 8. *Suppose* $p^*_{\max} \neq p^*_{\min}$. *Then*

$$\mathbb{P}\left[\hat{p}^*_{\max} \neq \hat{p}^*_{\min}\right] \geq 1 - 2Ke^{-\frac{1}{324(y+\overline{b}\overline{p})^2\overline{p}^2}N\alpha_0^2} - 2Ke^{-\frac{\frac{1}{K}\sum_{i=1}^K (p_i - \frac{1}{K}\sum_{j=1}^K p_j)^2 \underline{g}^2}{324C_1^2}N\alpha_0^2}$$
$$- 2Ke^{-\frac{1}{162C_2^2}N\alpha_0^2} - 2Ke^{-\frac{1}{162\overline{p}^2\overline{\lambda}^2}N\alpha_0^2},$$

where C_1 *and* C_2 *are constants defined in Lemma 2, and* $\alpha_0 = \frac{\gamma b}{(\sqrt{\gamma}+1)^2}\left(p^*_{\max} - p^*_{\min}\right)^2$.

Proposition 8 shows that when the true data-driven problem is unidentifiable, by checking whether $\hat{p}^*_{\max}$ and $\hat{p}^*_{\min}$ are different, the probability of the algorithm's failure to detect unidentifiability is very small as N increases. The intuition for this result is as follows. Recall that $p^*_{\max}$ and $p^*_{\min}$ are solved from $R_{\max}(\cdot)$ and $R_{\min}(\cdot)$ respectively. According to Lemma 2, as N grows, the empirical revenues $\hat{R}_{\max}(\cdot)$ and $\hat{R}_{\min}(\cdot)$ approximate $R_{\max}(\cdot)$ and $R_{\min}(\cdot)$ with higher accuracy respectively. Intuitively, their maximizers $\hat{p}^*_{\max}$ and $\hat{p}^*_{\min}$ will also approximate $p^*_{\max}$ and $p^*_{\min}$ with higher accuracy respectively. When the problem is unidentifiable, there is always a gap between $p^*_{\max}$ and $p^*_{\min}$, and thus, $\hat{p}^*_{\max}$ and $\hat{p}^*_{\min}$ are also bounded away from each other. The proof of Proposition 8 can be found in online Appendix E.

Proposition 8 also provides some guidance for the retailer in pricing practice. Proposition 3 shows that the interval $[p^*_{\min}, p^*_{\max}]$ characterizes the location of the true optimal price. The smaller this interval is, the closer the data-driven problem is to an identifiable problem. In practice, while the retailer cannot compute $p^*_{\max}$ or $p^*_{\min}$ since she does not know the distribution of ξ or the price sensitivity b, the empirical prices $\hat{p}^*_{\max}$ and $\hat{p}^*_{\min}$ can be computed based on Step 3 of Algorithm 1. Then the empirical interval $[\hat{p}^*_{\min}, \hat{p}^*_{\max}]$ can be seen as an approximation to the true interval $[p^*_{\min}, p^*_{\max}]$, which approximately captures the location of the true optimal price. Proposition 8 suggests that, if the data size is sufficiently large, but the gap between $\hat{p}^*_{\max}$ and $\hat{p}^*_{\min}$ is large, the problem is very likely to be unidentifiable, and the retailer should be very cautious when making the pricing decisions using censored data. In this case, modeling the uncertainty of demand distribution and solving the data-driven price optimization problem in a more robust way will be recommended.

6. Numerical Experiments

In this section, we conduct a numerical study on synthetic datasets to test the effectiveness of the proposed D2ACD. We first describe the numerical design and introduce three baseline algorithms and a modified algorithm from D2ACD, referred to as D2ACD-KM in Section 6.1. Then we compare D2ACD with the baseline algorithms and D2ACD-KM in Section 6.2. Finally, we discuss the effects of the data quality and inventory level on the performance of D2ACD in Section 6.3.

6.1. Description of Experiments

6.1.1. Three baseline algorithms and D2ACD-KM. We first introduce the three baseline algorithms that we compare D2ACD against in Section 6.2.

(1) The first algorithm LR-IncludeAll trains a linear regression model on all sales data, censored or not, to estimate the parameters (a,b), naively treating all sales data as true demands and completely ignoring the effect of demand censoring. After computing parameter estimates $(\hat{a},\hat{b})$, LR-IncludeAll constructs an empirical c.d.f. of η using the SAA approach based on the residual samples $\{S_i^j-(\hat{a}+\hat{b}p_i): i\in[K], j\in[N_i]\}$, and outputs a price that maxmizes the empirical revenue.

(2) The second algorithm LR-ExcludeCensored is aware of the fact that the offline data are subject to demand censoring when constructing a linear regression model, but it simply removes all the data points when the observed sales equal the inventory levels. The empirical revenue and algorithm's price are calculated using a similar approach to that in LR-IncludeAll, but with all censored residual samples excluded in the construction of the empirical c.d.f. of η.

(3) The third algorithm Buckley&James modified from Buckley and James (1979) is more intelligent than the first two baseline algorithms. It applies the KM estimator to a linear regression model by modifying the least squares normal equations when estimating the parameters (a,b). We refer interested readers to Buckley and James (1979) for a detailed description of this algorithm. The remaining steps for estimating the distribution of η, constructing the empirical revenue, and computing the algorithm suggested price, proceed in a similar manner as that of LR-IncludeAll.

We also introduce a modified algorithm from D2ACD, referred to as D2ACD-KM. Similarly, D2ACD-KM also proceeds in three major steps. The estimation for the price sensitivity in Step 1 and the procedure for computing the suggested price in Step 3 remain the same as D2ACD. The only difference lies in estimating the optimistic and pessimistic revenues in Step 2, where D2ACD-KM computes the empirical c.d.f. of ξ by applying a KM subroutine (see Algorithm 2 in online Appendix F for a detailed description) that leverages the KM estimator and utilizes all the historical data. This is in contrast to D2ACD that applies the SAA method and only uses the sales data associated with $(p_{\hat{i}^*}, y_{\hat{i}^*})$. We will also compare the two algorithms D2ACD and D2ACD-KM in Section 6.2.

6.1.2. Problem setup and performance measures. We set the deterministic base demand $a=100$, price sensitivity $b=1$, and random noise η as a centered Geometric random variable with parameter $1/30$. The inventory level and price range available to the retailer are $y=80$ and

$[\underline{p}, \overline{p}] = [30, 80]$, and the range of price sensitivity values is set to $[\underline{b}, \overline{b}] = [0.1, 3]$. Each set of the offline data is randomly sampled at $K = 2$ price-inventory pairs (p_1, y_1) and (p_2, y_2), which we set to $(40, 70)$ and $(60, 20)$ in Section 6.2 and to seven different combinations in Section 6.3. For any size metric $N \in \{20, \cdots, 200\}$ with increment of 1, we create a dataset $\mathcal{S}_N$ with N samples from (p_1, y_1) and N samples from (p_2, y_2).

For any algorithm $\mathcal{A}$, let $p^{\mathcal{A}}$ be the algorithm suggested price with $\mathcal{S}_N$, $[\underline{b}, \overline{b}]$, y, and $[\underline{p}, \overline{p}]$ given as input to $\mathcal{A}$. The algorithm performance on $\mathcal{S}_N$ is measured by two metrics:

(1) The *relative optimality gap* under the true c.d.f. $F_\xi(\cdot)$, defined as $\frac{R(p^*) - R(p^{\mathcal{A}})}{R(p^*)} \times 100\%$;

(2) The worst-case revenue loss over the ambiguity set $\mathcal{F}$, defined as $\max_{F \in \mathcal{F}} \{R_F(p_F^*) - R_F(p^{\mathcal{A}})\}$.

6.2. Algorithm Performance Comparison

In this subsection, we first compare the performances of D2ACD with those of LR-IncludeAll, LR-ExcludeCensored, and Buckley&James, averaged over 50 randomly generated problem instances of the same setup with $(p_1, y_1) = (40, 70)$ and $(p_2, y_2) = (60, 20)$, for each data size $N \in \{20, \cdots, 200\}$. In this setting, by computing the values of $p^*_{\max}$ and $p^*_{\min}$, the data-driven problem is verified to be unidentifiable.

In Figure 3, we plot the relative optimality gaps of the baseline and proposed algorithms, LR-IncludeAll, LR-ExcludeCensored, Buckley&James and D2ACD, as the percentages of revenue losses with respect to the optimal revenue $R(p^*)$. Figure 3 shows that D2ACD significantly improves over

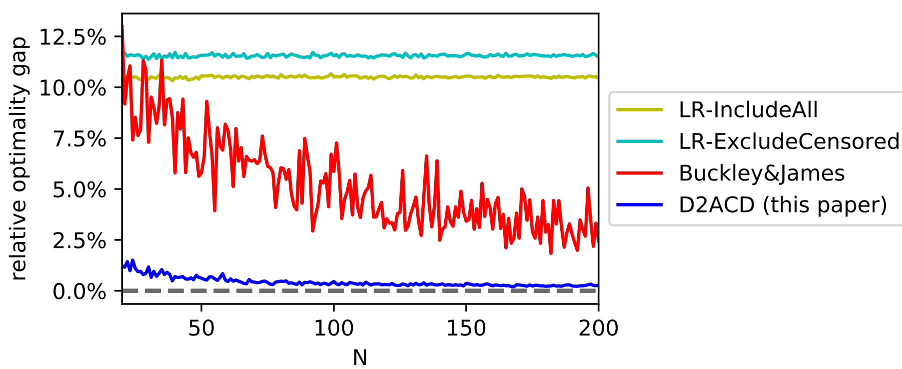

Figure 3 Relative optimality gaps of D2ACD **and baseline algorithms with different values of** N**.**

all three baseline algorithms, and, in this particular problem setup, it is able to keep the relative optimality gap from $< 2\%$ when $N = 20$ to $< 0.5\%$ when $N = 200$. As the sizes of the datasets get larger, D2ACD and Buckley&James enjoy smaller revenue losses, while the performances of LR-IncludeAll and LR-ExcludeCensored remain the same as their estimation biases do not get reduced with more data.

In Figure 4, we plot the worst-case revenue losses of the proposed and baseline algorithms. It is obvious that, among all four algorithms, D2ACD achieves the lowest worst-case revenue loss, which gets closer to the minimax revenue loss as the data size gets larger. We point out that the revenue

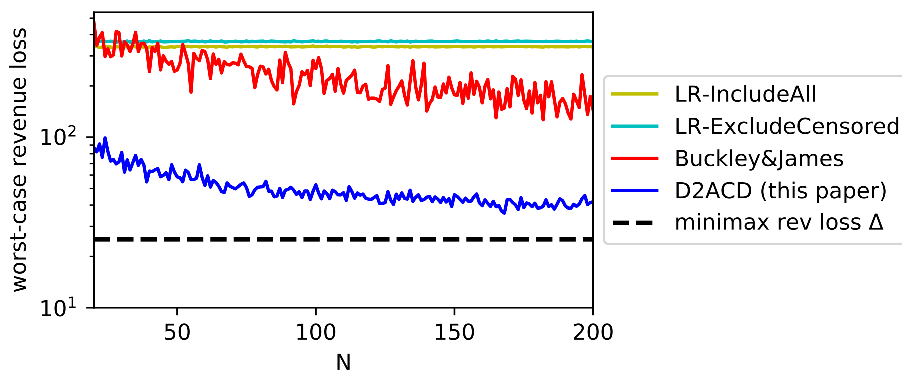

Figure 4 Worst-case revenue losses of D2ACD **and baseline algorithms with different values of** N**.**

losses are displayed in log-scale, and therefore the D2ACD's improvement over other algorithms is even more significant than the differences appearing on the chart. Figure 4 shows that the price given by D2ACD is not only effective in the particular demand distribution that we use to generate the offline datasets, but also highly robust to all possible distributions that may output the same datasets.

We next offer explanations on the superior performance of D2ACD by demonstrating the accuracy of D2ACD's price sensitivity estimates compared with those of the baseline algorithms and by illustrating D2ACD's pricing mechanism that is designed to be robust. Figure 5 plots the price sensitivity estimates $\hat{b}$ of the baseline and proposed algorithms on datasets of different sizes, where for each $20 \le N \le 200$, we plot the estimate $\hat{b}$ for each algorithm on one randomly generated dataset. It is clear that D2ACD attains the most accurate estimators of price sensitivity among all algorithms, and the results obtained from repeated experiments become more concentrated around the true price sensitivity as the data sizes grow, agreeing with our theoretical result in Lemma 1. The plots of the estimates of LR-IncludeAll and LR-ExcludeCensored show that ignoring the censoring effect leads to biased estimates and simply removing censored data does not solve the problem either. In particular, if dropping censored data, the selected samples for the linear regression are restricted to those below the inventory level, and thus, `LR-ExcludeCensored` will introduce additional bias to the regression model in an *endogenous* way. See more discussion on such endogeneity in online Appendix G. The third baseline algorithm Buckley&James slightly improves over the other two baseline algorithms, as its estimators sometimes do come very close

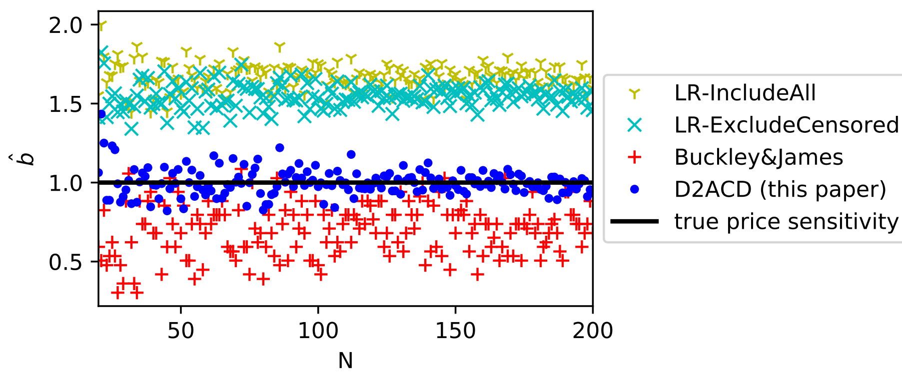

Figure 5 **Price sensitivity estimates $\hat{b}$ of D2ACD and baseline algorithms with different values of N.**

to the true price sensitivity. However, Buckley&James requires iteratively finding a fixed point of a discontinuous piecewise linear function, which may have zero, one, or multiple fixed points, and therefore its estimates are very susceptible to the choice of initial points. In the particular problem instances that we test, we see from Figure 5 that the price sensitivity estimates of Buckley&James tend to converge at fixed points that are lower than the true price sensitivity.

Figure 6 illustrates D2ACD's pricing mechanism: the algorithm first narrows down the location of the optimal price p^* to the interval between its empirical optimistic price $\hat{p}^*_{\max}$ and pessimistic price $\hat{p}^*_{\min}$, and then, within the interval, it finds a price p^{D2ACD} that minimizes the empirical worst-case revenue loss, based on $\hat{R}_{\max}(\cdot)$ and $\hat{R}_{\min}(\cdot)$, D2ACD's estimated upper and lower bounds of the true revenue function $R(\cdot)$. The intuition is that, when the data size is large, the empirical optimistic

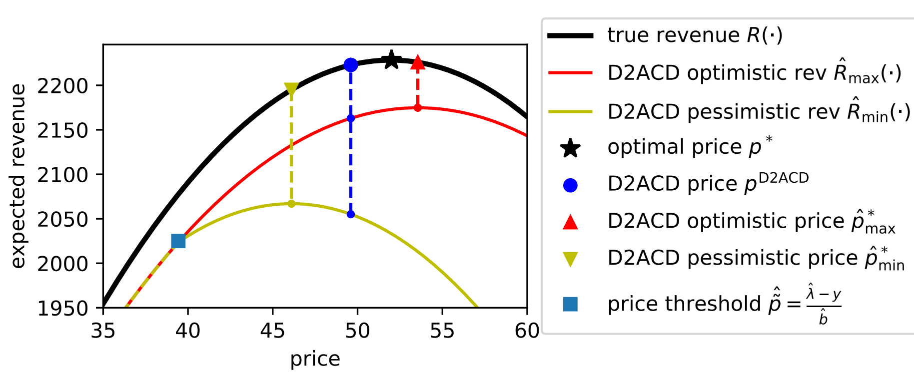

Figure 6 **Locations of the optimal price and D2ACD's prices in a problem instance with $N = 50$.**

and pessimistic prices are very close to the true optimistic and pessimistic prices respectively, which are the boundaries of the interval that confines the optimal price p^*, as suggested by Proposition 3. Therefore, with more data, the distance between p^{D2ACD} and p^* converges to a constant that is no greater than $|p^*_{\max} - p^*_{\min}|$ and leads to expected revenue loss of $< 0.5\%$ in the problem set

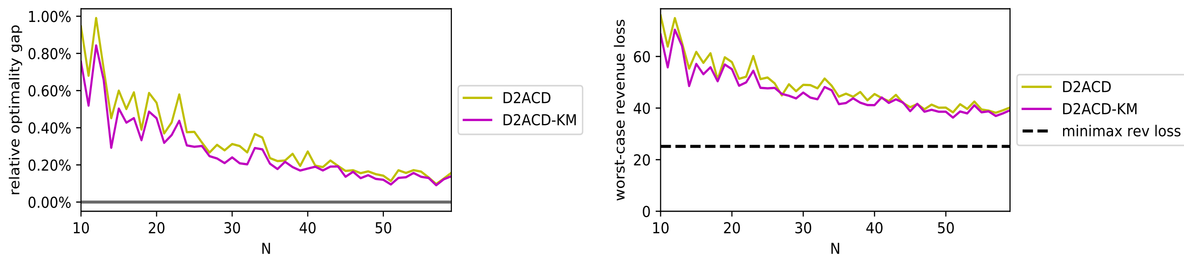

Figure 7 Comparison between D2ACD and D2ACD-KM with $K = 10$.

that we use to generate Figure 3. Similarly, when the data size is large, $\hat{R}_{\max}(\cdot)$ and $\hat{R}_{\min}(\cdot)$ well approximate $R_{\max}(\cdot)$ and $R_{\min}(\cdot)$ respectively, according to Proposition 2. Therefore, since p^{D2ACD} is set to minimize the empirical worst-case revenue loss based on $\hat{R}_{\max}(\cdot)$ and $\hat{R}_{\min}(\cdot)$, its true worst-case revenue loss calculated based on $R_{\max}(\cdot)$ and $R_{\min}(\cdot)$ is also close to the minimax revenue loss Δ, as shown on Figure 4.

To end this subsection, we compare the performances of D2ACD and D2ACD-KM. Through our experiments, we find that under the setting of two historical price-inventory pairs as in Figures 3 and 4, D2ACD achieves almost the same numerical performance as D2ACD-KM in terms of both the relative optimality gap and worst-case revenue loss. To better distinguish between the performances of the two algorithms, we consider the following ten pairs of historical pricing and inventory levels: $\{(35,45),(45,35),(55,25),(65,15),(75,5),(30,80),(40,70),(50,60),(60,50),(70,40)\}$, each of which is associated with N sales observations. In this case, D2ACD-KM utilizes much more historical information than D2ACD for estimating the distribution of ξ in Step 2. The other parameters remain the same as before. The relative optimality gap and the worst-case revenue loss incurred by D2ACD and D2ACD-KM are plotted in Figure 7. When the data size N is small, e.g., $N < 40$, D2ACD-KM performs slightly better than D2ACD under both measures of the relative optimality gap and the worst-case revenue loss, although the improvement is not significant. When N exceeds 40, the gap between D2ACD and D2ACD-KM is negligible, indicating the almost identical performance of D2ACD and D2ACD-KM with larger N. These are mainly because under our assumed data structure, the SAA method applied to the "best-quality" data $\{(p_{\hat{i}^*}, y_{\hat{i}^*}, S^j_{\hat{i}^*}) : j \in [N_{\hat{i}^*}]\}$ already leads to an accurate estimation for the distribution of ξ even with small to moderate data size N.

6.3. Effects of Data Quality and Inventory Level

In Section 6.3, we investigate the effects of different data qualities and inventory levels on the performance of the proposed D2ACD. As discussed in Section 2.1, the observable boundary λ is used as the metric of data quality. In this subsection, we also include the absolute difference between the two prices used to generate the offline dataset, denoted by $\rho := |p_1 - p_2|$, as another metric of data quality. In Table 1, we report the relative optimality gap of D2ACD under different combinations of

Data quality	Inventory level					Data quality	Inventory level				
(ρ,λ)	$y=20$	$y=40$	$y=60$	$y=80$	$y=100$	(ρ,λ)	$y=20$	$y=40$	$y=60$	$y=80$	$y=100$
$(10,105)$	<u>1.54%</u>	<u>2.02%</u>	1.56%	1.90%	3.96%	$(20,90)$	<u>0.38%</u>	0.57%	0.62%	2.59%	4.76%
$(15,105)$	<u>0.76%</u>	<u>0.66%</u>	1.01%	1.01%	2.13%	$(20,95)$	<u>0.65%</u>	0.19%	0.45%	1.75%	3.49%
$(20,105)$	<u>0.67%</u>	<u>0.41%</u>	0.53%	0.86%	1.64%	$(20,100)$	<u>0.50%</u>	<u>0.33%</u>	0.78%	1.45%	2.64%
$(30,105)$	<u>0.36%</u>	<u>0.32%</u>	0.48%	0.52%	1.51%	$(20,105)$	<u>0.50%</u>	<u>0.54%</u>	0.61%	0.87%	1.41%

Table 1 Relative optimality gap of D2ACD **under different** (ρ,λ) **and** y **with** $N=100$.

data quality metrics (ρ,λ) and inventory level y, averaged over 100 randomly generated datasets. When the problem is identifiable with the given quality metrics (ρ,λ) and inventory level y, we underline the relative optimality gap to differentiate it from the unidentifiable cases. Table 1 shows that the values of λ and y affect the identifiability of the underlying data-driven problem, and the values of ρ, λ and y also jointly influence the performance of D2ACD.

First, when λ and y are fixed, as ρ increases, the relative optimality gap of D2ACD decreases, leading to a better performance. For example, when $y=20$ and $\lambda=105$, as ρ increases from 10 to 30, the relative optimality gap decreases from 1.54% to 0.36%. This is intuitive since ρ measures how dispersive historical prices are, and as ρ becomes larger, the estimate $\hat{b}$ for the price sensitivity b becomes more accurate, and the algorithm is more likely to generate higher revenue.

Second, when y is fixed, as λ increases, an unidentifiable problem switches to an identifiable one. For example, for fixed $y=40$, the problem is unidentifiable when $\lambda\le 95$, and becomes identifiable when $\lambda\ge 100$. This is also consistent with our Proposition 5, which shows that when λ exceeds some threshold, the problem becomes identifiable. Moreover, for unidentifiable problems with a relatively high inventory level y, the relative optimality gap decreases as the observable boundary λ increases. For example, when $y=80$, as λ increases from 90 to 105, the relative optimality gap decreases from 2.59% to 0.87%. This is because a larger λ provides more information about the unknown distribution of ξ, and therefore leads to a better performance of D2ACD. For identifiable problems and unidentifiable problems with a relatively low inventory level y, the relative optimality gaps are consistently close to zero under different values of λ.

Third, when λ is fixed, as y increases, an identifiable problem switches to an unidentifiable one. For example, for $\lambda=105$, the problem is identifiable when $y\le 40$ and becomes unidentifiable when $y\ge 60$. Moreover, a higher inventory level y leads to a larger relative optimality gap and a worse performance of D2ACD. This is because to evaluate the expected revenue under a fixed price p, we need to know the distribution of ξ on the left-hand side of $y+bp$. If y is very small, it is more likely that $y+bp$ is within the observable boundary λ, and therefore, the problem is more likely to be identifiable. If y is extremely large, e.g., $y = \infty$, the optimal price is close to $a/(2b)$, which makes the problem unidentifiable since with censored offline data, no algorithm is able to uncover the exact value of a, let alone achieving an optimality gap close to zero.

7. Conclusion and Future Research

In this paper, we study a data-driven pricing problem with unknown demand model, for which demand censoring brings a crucial challenge to the existence of near-optimal data-driven algorithms. We define the notion of problem identifiability by constructing an ambiguity set of demand distributions, and measuring the performance of a data-driven algorithm in terms of the worst-case revenue loss. We provide an exact condition for a problem to be identifiable by comparing the optimal solutions to two distributionally robust counterparts of the full-information pricing problem. Moreover, we develop a data-driven algorithm that can be applied for both identifiable and unidentifiable problems, whose worst-case revenue loss converges to the best-achievable revenue loss for both cases. Numerical experiments are conducted to demonstrate the effectiveness of our proposed algorithms.

Demand censoring is widely observed in many business environments, especially in the retailing industry. This paper, by studying a data-driven pricing problem with a simple linear demand model, also aims to provide several important managerial insights. When businesses use offline data to optimize pricing decisions, the effect of demand censoring cannot be simply ignored. In a censored dataset, some partial distributional information is inevitably lost, creating additional demand uncertainty for businesses to manage. As shown in the numerical experiments (Section 6), naive treatments of censored data can lead to inaccurate demand estimation and suboptimal pricing decisions. A holistic approach to modeling the additional uncertainty in the demand learning process, for example, using an ambiguity set to capture all potential demand distributions as shown in Section 3.1, and the ambiguous stochastic optimization approach, for example, using the worst-case revenue loss to measure an algorithm's performance, are crucial to allow businesses to practice data-driven price optimization under uncertainty in a more robust manner.

This work also opens up several future research opportunities for offline demand learning and pricing in the presence of censored data. First, nowadays, retailers especially in the online retailing, have access to massive amount of covariates information about customers and demands, e.g., sales of substitutable products, customers' historical purchase behavior, etc. By leveraging such information, it is hopeful to make better inference about the lost-sales quantity, and build a less conservative data-driven pricing problem with a refined ambiguity set. Second, our model assumes the unknown demand curve belongs to the linear function class. It would be interesting to consider other classes of demand functions, in which case, it is important to study how the demand ambiguity set should be modified accordingly and to understand how demand censoring affects the problem identifiability under the new demand classes. Third, in this paper, we consider a data-driven pricing problem with given inventory capacity. In practice, firms may want to coordinate both pricing and inventory decisions to maximize the profit function $p\mathbb{E}[\min\{\xi - bp, y\}] - cy$. It will be interesting to investigate how to extend the current framework and results to the joint pricing and inventory management problem with censored data. One key challenge is that the profit function is generally not jointly concave in the decision variable (p, y). The technique developed by Feng et al. (2020) may be useful in dealing with such a non-concavity issue. Finally, it will also be important to extend the current single-period model to multi-period dynamic pricing problem. In such a model, we can still apply a similar worst-case framework with a properly defined ambiguity set to capture the information loss and measure the performance of a data-driven algorithm. In the analysis of a multi-period pricing model, the attention would be diverted to constructing a tractable empirical dynamic program as well as analyzing its relationship with the full-information dynamic program (see Cheung and Simchi-Levi 2019), which is more challenging than the current single-period model is.

Acknowledgments: The authors are grateful to the department editor J. George Shanthikumar, the associate editor and two referees for their constructive comments and suggestions which have helped to significantly improve both the content and exposition of this paper.

Appendix. Proof of Theorems 2 and 3

In this appendix, we prove Theorems 2 and 3 based on Proposition 7 and Lemma 2. For notation convenience, let $p^*_\Delta = \arg\min_{p\in[\underline{p},\overline{p}]} \max\{R_{\max}(p^*_{\max}) - R_{\max}(p), R_{\min}(p^*_{\min}) - R_{\min}(p)\}$ and for any $\alpha > 0$, let $\mathcal{B}(\alpha)$ be the following event:

$$\mathcal{B}(\alpha) = \left\{ |\hat{R}_{\max}(p) - R_{\max}(p)| \le \alpha, |\hat{R}_{\min}(p) - R_{\min}(p)| \le \alpha, \forall p \in [\underline{p}, \overline{p}] \right\}. \tag{8}$$

Note that from Proposition 7, when $p^*_{\max} = p^*_{\min}$, we have

$$\max_{F\in\mathcal{F}} \{R_F(p^*_F) - R_F(p^*_{\max})\} = \max\{R_{\max}(p^*_{\max}) - R_{\max}(p^*_{\max}), R_{\min}(p^*_{\min}) - R_{\min}(p^*_{\min})\} = 0,$$

which implies $\Delta = 0$ for this case. Thus, to prove both Theorems 2 and 3, it suffices to show that for any $\varepsilon > 0$ and $\delta \in (0,1)$, when $N \ge N(\varepsilon, \delta)$,

$$\mathbb{P}\left[\max_{F\in\mathcal{F}}\{R_F(p^*_F) - R_F(p^{\mathsf{D2ACD}})\} - \Delta \le \varepsilon\right] \ge 1 - \delta. \tag{9}$$

Recall that $p^{\mathsf{D2ACD}} = \arg\min_{p\in[\underline{p},\overline{p}]} \max\left\{\hat{W}_{\max}(p), \hat{W}_{\min}(p)\right\}$. Suppose $\mathcal{B}(\alpha)$ holds, then we obtain

$$\max_{F\in\mathcal{F}}\{R_F(p^*_F) - R_F(p^{\mathsf{D2ACD}})\}$$

$$
\begin{aligned}
&= \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p^{\mathsf{D2ACD}}),\, R_{\min}(p^*_{\min}) - R_{\min}(p^{\mathsf{D2ACD}})\right\} \\
&\le \max\left\{(\hat{R}_{\max}(p^*_{\max}) + \alpha) - (\hat{R}_{\max}(p^{\mathsf{D2ACD}}) - \alpha),\, (\hat{R}_{\min}(p^*_{\min}) + \alpha) - (\hat{R}_{\min}(p^{\mathsf{D2ACD}}) - \alpha)\right\} \\
&\le \max\left\{\hat{R}_{\max}(\hat{p}^*_{\max}) - \hat{R}_{\max}(p^{\mathsf{D2ACD}}),\, \hat{R}_{\min}(\hat{p}^*_{\min}) - \hat{R}_{\min}(p^{\mathsf{D2ACD}})\right\} + 2\alpha \\
&\le \max\left\{\hat{R}_{\max}(\hat{p}^*_{\max}) - \hat{R}_{\max}(p^*_{\Delta}),\, \hat{R}_{\min}(\hat{p}^*_{\min}) - \hat{R}_{\min}(p^*_{\Delta})\right\} + 2\alpha \\
&\le \max\left\{(R_{\max}(\hat{p}^*_{\max}) + \alpha) - (R_{\max}(p^*_{\Delta}) - \alpha),\, (R_{\min}(\hat{p}^*_{\min}) + \alpha) - (R_{\min}(p^*_{\Delta}) - \alpha)\right\} + 2\alpha \\
&\le \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p^*_{\Delta}),\, R_{\min}(p^*_{\min}) - R_{\min}(p^*_{\Delta})\right\} + 4\alpha \\
&= \Delta + 4\alpha, \qquad (10)
\end{aligned}
$$

where the first identity follows from Proposition 7, the first inequality follows from (8), the second inequality follows from the optimality of $\hat{p}^*_{\max}$ and $\hat{p}^*_{\min}$, and $p^*_{\max}, p^*_{\min} \in [\underline{p}, \overline{p}]$, the third inequality follows from the optimality of p^{D2ACD} and $p^*_{\Delta} \in [\underline{p}, \overline{p}]$, the fourth inequality follows from (8), the last inequality follows from the optimality of $p^*_{\max}$ and $p^*_{\min}$, and $\hat{p}^*_{\max}, \hat{p}^*_{\min} \in [\underline{p}, \overline{p}]$, and the last identity follows from the definition of p^*_{Δ}.

For any $\varepsilon > 0$, from (10) and by letting $\alpha = \frac{1}{4}\varepsilon$, we have

$$
\begin{aligned}
&\mathbb{P}\left[\max_{F\in\mathcal{F}}\{R_F(p^*_F) - R_F(p^{\mathsf{D2ACD}})\} - \Delta \le \varepsilon\right] \\
&\ge \mathbb{P}\left[\left|\hat{R}_{\max}(p) - R_{\max}(p)\right| \le \frac{1}{4}\varepsilon, \left|\hat{R}_{\min}(p) - R_{\min}(p)\right| \le \frac{1}{4}\varepsilon, \forall p \in [\underline{p}, \overline{p}]\right] \\
&\ge 1 - 2Ke^{-\frac{1}{576(y+\overline{b}\overline{p})^2\overline{p}^2}N\varepsilon^2} - 2Ke^{-\frac{\frac{1}{K}\sum_{i=1}^K(p_i - \frac{1}{K}\sum_{j=1}^K p_j)^2 \underline{g}^2}{576C_1^2}N\alpha^2} - 2Ke^{-\frac{1}{288C_2^2}N\alpha^2} - 2Ke^{-\frac{1}{288\overline{p}^2\overline{\lambda}^2}N\alpha^2},
\end{aligned}
$$

where the last inequality follows from Proposition 2. Thus, when $N \ge N(\varepsilon,\delta) = C\frac{1}{\varepsilon^2}\log\frac{8K}{\delta}$, where $C := \max\left\{576(y+\overline{b}\overline{p})^2\overline{p}^2, \frac{576C_1^2}{\frac{1}{K}\sum_{i=1}^K(p_i - \frac{1}{K}\sum_{j=1}^K p_j)^2\underline{g}^2}, 288C_2^2, 288\overline{p}^2\overline{\lambda}^2\right\}$, event $\max_{F\in\mathcal{F}}\{R_F(p^*_F) - R_F(p^{\mathsf{D2ACD}})\} - \Delta \le \varepsilon$ holds with probability lower bounded by $1-\delta$. □

References

Agrawal, Shipra, Randy Jia. 2019. Learning in structured mdps with convex cost functions: Improved regret bounds for inventory management. *Proceedings of the 2019 ACM Conference on Economics and Computation*. 743–744.

Ban, Gah-Yi. 2020. Confidence intervals for data-driven inventory policies with demand censoring. *Operations Research* **68**(2) 309–326.

Ban, Gah-Yi, Zhenyu Gao, Fabian Taigel. 2020. Model mis-specification in newsvendor decisions: A comparison of frequentist parametric, bayesian parametric and nonparametric approaches. *Available at SSRN 3495733* .

Ban, Gah-Yi, Cynthia Rudin. 2019. The big data newsvendor: Practical insights from machine learning. *Operations Research* **67**(1) 90–108.

Besbes, Omar, Alp Muharremoglu. 2013. On implications of demand censoring in the newsvendor problem. *Management Science* **59**(6) 1407–1424.

Besbes, Omar, Assaf Zeevi. 2015. On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Management Science* **61**(4) 723–739.

Buckley, Jonathan, Ian James. 1979. Linear regression with censored data. *Biometrika* **66**(3) 429–436.

Chen, Boxiao. 2019. Data-driven inventory control with shifting demand. *Available at SSRN 3503139* .

Chen, Boxiao, Xiuli Chao. 2020. Dynamic inventory control with stockout substitution and demand learning. *Management Science* .

Chen, Boxiao, Xiuli Chao, Hyun-Soo Ahn. 2019. Coordinating pricing and inventory replenishment with nonparametric demand learning. *Operations Research* **67**(4) 1035–1052.

Chen, Boxiao, Xiuli Chao, Cong Shi. 2020a. Nonparametric algorithms for joint pricing and inventory control with lost-sales and censored demand. *Available at SSRN 2836057* .

Chen, Boxiao, Xiuli Chao, Yining Wang. 2020b. Data based dynamic pricing and inventory control with censored demand and limited price changes. *Available at SSRN 2700747* .

Chen, Boxiao, Cong Shi. 2019. Tailored base-surge policies in dual-sourcing inventory systems with demand learning. *Available at SSRN 3456834* .

Cheung, Wang Chi, David Simchi-Levi. 2019. Sampling-based approximation schemes for capacitated stochastic inventory control models. *Mathematics of Operations Research* **44**(2) 668–692.

Cox, David R. 1972. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)* **34**(2) 187–202.

Dawes, Robyn M. 1979. The robust beauty of improper linear models in decision making. *American psychologist* **34**(7) 571.

den Boer, Arnoud V. 2015. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science* **20**(1) 1–18.

Feng, Qi, Sirong Luo, J George Shanthikumar. 2020. Integrating dynamic pricing with inventory decisions under lost sales. *Management Science* **66**(5) 2232–2247.

Feng, Qi, J George Shanthikumar. 2018. How research in production and operations management may evolve in the era of big data. *Production and Operations Management* **27**(9) 1670–1684.

Fernandez, Luis. 1986. Non-parametric maximum likelihood estimation of censored regression models. *Journal of Econometrics* **32**(1) 35–57.

Foxman, Simone. 2013. The us is home to one third of the world's data – here's who's storing it. https://qz.com/104868/the-us-is-home-to-one-third-of-the-worlds-data-heres-whos-storing-it. Quartz.

Hoeffding, Wassily. 1994. Probability inequalities for sums of bounded random variables. *The Collected Works of Wassily Hoeffding*. Springer, 409–426.

Huh, Woonghee Tim, Ganesh Janakiraman, John A Muckstadt, Paat Rusmevichientong. 2009. An adaptive algorithm for finding the optimal base-stock policy in lost sales inventory systems with censored demand. *Mathematics of Operations Research* **34**(2) 397–416.

Huh, Woonghee Tim, Retsef Levi, Paat Rusmevichientong, James B Orlin. 2011. Adaptive data-driven inventory control with censored demand based on kaplan-meier estimator. *Operations Research* **59**(4) 929–941.

Huh, Woonghee Tim, Paat Rusmevichientong. 2009. A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research* **34**(1) 103–123.

Kaplan, Edward L, Paul Meier. 1958. Nonparametric estimation from incomplete observations. *Journal of the American statistical association* **53**(282) 457–481.

Koul, H, V v Susarla, J Van Ryzin, et al. 1981. Regression analysis with randomly right-censored data. *The Annals of statistics* **9**(6) 1276–1288.

Levi, Retsef, Georgia Perakis, Joline Uichanco. 2015. The data-driven newsvendor problem: new bounds and insights. *Operations Research* **63**(6) 1294–1306.

Levi, Retsef, Robin O Roundy, David B Shmoys. 2007. Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research* **32**(4) 821–839.

Lewbel, Arthur, Oliver Linton. 2002. Nonparametric censored and truncated regression. *Econometrica* **70**(2) 765–779.

Lugosi, Gábor, Mihalis G Markakis, Gergely Neu. 2017. On the hardness of inventory management with censored demand data. *arXiv preprint arXiv:1710.05739* .

Massart, Pascal. 1990. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The annals of Probability* 1269–1283.

Miller, Rupert G. 1976. Least squares regression with censored data. *Biometrika* **63**(3) 449–464.

Powell, James L. 1984. Least absolute deviations estimation for the censored regression model. *Journal of Econometrics* **25**(3) 303–325.

Qin, Hanzhang, David Simchi-Levi, Li Wang. 2019. Data-driven approximation schemes for joint pricing and inventory control models. *Available at SSRN 3354358* .

Rahimian, Hamed, Sanjay Mehrotra. 2019. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659* .

Tobin, James. 1958. Estimation of relationships for limited dependent variables. *Econometrica: journal of the Econometric Society* 24–36.

Yuan, Hao, Qi Luo, Cong Shi. 2019. Marrying stochastic gradient descent with bandits: Learning algorithms for inventory systems with fixed costs. *Available at SSRN 3329611* .

Zhang, Huanan, Xiuli Chao, Cong Shi. 2018. Perishable inventory systems: Convexity results for base-stock policies and learning algorithms under censored demand. *Operations Research* **66**(5) 1276–1286.

Zhang, Huanan, Xiuli Chao, Cong Shi. 2020. Closing the gap: A learning algorithm for lost-sales inventory systems with lead times. *Management Science* **66**(5) 1962–1980.

Online Appendix for "Offline Pricing and Demand Learning with Censored Data"

Appendix A: Some Concentration Inequalities

Before providing the omitted proofs for the results in this paper, we first introduce several useful concentration inequalities.

The following lemma presents the classic Hoeffding inequality for bounded random variables.

LEMMA A.1 **(Hoeffding inequality, Hoeffding 1994)**. *Let $X_1,\dots,X_n$ denote n i.i.d. samples of a random variable with mean μ and bounded support $[\underline{x},\overline{x}]$. Then, for any $\varepsilon>0$,*

$$\mathbb{P}\left[\left|\frac{1}{n}\sum_{i=1}^{n}X_i-\mu\right|\le\varepsilon\right]\ge 1-2e^{-\frac{2n\varepsilon^2}{(\overline{x}-\underline{x})^2}}.$$

The next lemma quantifies how close an empirical distribution function constructed from *i.i.d.* samples is to the true distribution function, which is proved in Massart (1990).

LEMMA A.2 **(Dvoretzk-Kiefer-Wolfowitz inequality, Massart 1990)**. *Let $\hat{F}_n(\cdot)$ denote the empirical distribution function for n i.i.d. samples of a random variable with c.d.f. $F(\cdot)$. Then, for any $\varepsilon>0$,*

$$\mathbb{P}\left[\sup_{x\in\mathbb{R}}|\hat{F}_n(x)-F(x)|\le\varepsilon\right]\ge 1-2e^{-2n\varepsilon^2}.$$

The following lemma is a generalization of the Dvoretzk-Kiefer-Wolfowitz inequality when the random variable of interest is not necessarily an indicator function, but a general bounded and monotone function.

LEMMA A.3 **(Lemma 14 in Qin et al. 2019)**. *Let $X_1,X_2,\dots,X_n$ be i.i.d. random variables, and $\mathcal{X},\mathcal{Y}\subseteq\mathbb{R}$. Suppose $f(x,y):\mathcal{X}\times\mathcal{Y}\to[L,U]$ is a measurable function and monotone in $x\in\mathcal{X}$ for any given $y\in\mathcal{Y}$. Then, for any $\varepsilon>0$,*

$$\mathbb{P}\left[\sup_{y\in\mathcal{Y}}\left|\frac{1}{n}\sum_{i=1}^{n}f(X_i,y)-\mathbb{E}[f(X_1,y)]\right|\le\varepsilon\right]\ge 1-2e^{-\frac{2n\varepsilon^2}{(U-L)^2}}.$$

Appendix B: Omitted Proof in Section 2

B.1. Proof of Proposition 1

Proof. Note that $R(p;y)=\mathbb{E}[\min\{p(\xi-bp),py\}]$. For any realization of ξ, $p(\xi-bp)$ is concave in p as $b>0$, and in addition, py is linear and therefore is also concave in p. Thus, for any realization of ξ, $\min\{p(\xi-bp),py\}$ is also concave in p, which implies concavity of $R(p;y)$. From concavity, we have $p^*=\min\{\max\{p^\dagger,\underline{p}\},\overline{p}\}$, which completes the proof of Proposition 1. □

Appendix C: Omitted Proofs in Section 3

C.1. Proof of Proposition 2

Proof. For any $F\in\mathcal{F}$ and any $p\ge 0$, we have

$$\begin{aligned}R_F(p)&=p\cdot\mathbb{E}_F\left[(-bp+\xi)\cdot\mathbb{1}_{\{\xi<y+bp\}}+y\cdot\mathbb{1}_{\{\xi\ge y+bp\}}\right]\\&=p\cdot\big(-bp\mathbb{P}_F\left[\xi<y+bp\right]+\mathbb{E}_F\left[\xi\cdot\mathbb{1}_{\{\xi<y+bp\}}\right]+y\cdot(1-\mathbb{P}_F\left[\xi<y+bp\right])\big).\end{aligned}\qquad\text{(C.1)}$$

Also, by the definition of $\mathcal{F} = \Big\{F \text{ is a } c.d.f.\text{: } F(x) = F_\xi(x), \forall x < \lambda\Big\}$, we have, for any $F \in \mathcal{F}$, and any $x < \lambda$,

$$\mathbb{P}_F\left[\xi < x\right] = \mathbb{P}_{F_\xi}\left[\xi < x\right], \quad \mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\xi < x\}}\right] = \mathbb{E}_{F_\xi}\left[\xi \cdot \mathbb{1}_{\{\xi < x\}}\right]. \tag{C.2}$$

Recall that $\tilde{p} = \frac{\lambda - y}{b}$. We next consider two cases.

<u>**Case 1:** $p < \tilde{p}$.</u> In this case, $y + bp < \lambda$. By (C.1) and (C.2), we have

$$R_F(p) = p \cdot \big(- bp\mathbb{P}_{F_\xi}\left[\xi < y + bp\right] + \mathbb{E}_{F_\xi}\left[\xi \cdot \mathbb{1}_{\{\xi < y + bp\}}\right] + y \cdot (1 - \mathbb{P}_{F_\xi}\left[\xi < y + bp\right])\big) = R_{F_\xi}(p).$$

Thus, $R_{\max}(p) = R_{\min}(p) = R_{F_\xi}(p)$.

<u>**Case 2:** $p \geq \tilde{p}$.</u> In this case, we have $y + bp \geq \lambda$. By (C.2), we have

$$\begin{aligned}
&\mathbb{P}_F\left[\xi < y + bp\right] = \mathbb{P}_{F_\xi}\left[\xi < \lambda\right] + \mathbb{P}_F\left[\lambda \leq \xi < y + bp\right] = \gamma + \mathbb{P}_F\left[\lambda \leq \xi < y + bp\right],\\
&\mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\xi < y + bp\}}\right] = \mathbb{E}_{F_\xi}\left[\xi \cdot \mathbb{1}_{\{\xi < \lambda\}}\right] + \mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right] = \gamma \mathbb{E}_{F_\xi}[\xi | \xi < \lambda] + \mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right].
\end{aligned}$$

Then, by (C.1), we have

$$\begin{aligned}
R_F(p) &= p \cdot \Big(- bp\big(\gamma + \mathbb{P}_F\left[\lambda \leq \xi < y + bp\right]\big) + \gamma \mathbb{E}_{F_\xi}[\xi | \xi < \lambda] + \mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right] + y\big(1 - \gamma - \mathbb{P}_F\left[\lambda \leq \xi < y + bp\right]\big)\Big)\\
&= p \cdot \big(- \gamma bp + \gamma \mathbb{E}_{F_\xi}[\xi | \xi < \lambda] + (1 - \gamma)y\big) + p \cdot \mathbb{E}_F\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right].
\end{aligned}$$

For different $F \in \mathcal{F}$, $p \cdot \big(- \gamma bp + \gamma \mathbb{E}_{F_\xi}[\xi | \xi < \lambda] + (1 - \gamma)y\big)$ is invariant. Thus, to optimize $R_F(p)$ over $F \in \mathcal{F}$, it is equivalent to optimizing $p \cdot \mathbb{E}_F\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right]$. Depending on the sign of p, it can be further translated to optimizing $\mathbb{E}_F\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right]$.

Note that by the definition of $\mathcal{F}$, a distribution $F \in \mathcal{F}$ has fixed probability distribution over $(-\infty, \lambda)$, but has full freedom in assigning probability mass over $[\lambda, \infty)$. To solve $\max_{F \in \mathcal{F}} \mathbb{E}_F\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right]$, we observe that $\mathbb{E}_F\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right] \leq 0$. Hence, a distribution $\overline{F} \in \mathcal{F}$ maximizes $\mathbb{E}_F\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right]$ for $F \in \mathcal{F}$ by setting $\mathbb{P}_{\overline{F}}[\lambda \leq \xi < y + bp] = 0$. Thus,

$$\begin{aligned}
&\max_{F \in \mathcal{F}} \left\{\mathbb{E}_F\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right]\right\} = \mathbb{E}_{\overline{F}}\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right] = 0,\\
&R_{\overline{F}}(p) = p \cdot \big(- \gamma bp + \mathbb{E}_{F_\xi}[\min\{\xi, \lambda\}] + (1 - \gamma)(y - \lambda)\big)] = R_1(p).
\end{aligned}$$

To solve $\min_{F \in \mathcal{F}} \mathbb{E}_F\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right]$, we observe that $\mathbb{E}_F[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}] \geq (-bp + \lambda - y) \cdot (1 - \gamma)$. Hence, a distribution $\underline{F} \in \mathcal{F}$ minimizes $\mathbb{E}_F\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right]$ for $F \in \mathcal{F}$ by assigning all remaining probability mass apart from γ to a single point λ, i.e, $\mathbb{P}_{\underline{F}}[\xi = \lambda] = 1 - \gamma$. Thus,

$$\begin{aligned}
&\min_{F \in \mathcal{F}} \left\{\mathbb{E}_F\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right]\right\} = \mathbb{E}_{\underline{F}}\left[(-bp + \xi - y) \cdot \mathbb{1}_{\{\lambda \leq \xi < y + bp\}}\right] = (-bp + \lambda - y)(1 - \gamma),\\
&R_{\underline{F}}(p) = p \cdot \big(- bp + \mathbb{E}_{F_\xi}[\min\{\xi, \lambda\}]\big) = R_2(p).
\end{aligned}$$

If $p \geq 0$, then $R_{\max}(p) = R_{\overline{F}}(p) = R_1(p)$ and $R_{\min}(p) = R_{\underline{F}}(p) = R_2(p)$. If $p < 0$, then $R_{\max}(p) = R_{\underline{F}}(p) = R_2(p)$ and $R_{\min}(p) = R_{\overline{F}}(p) = R_1(p)$. Combining both cases of $p \geq 0$ and $p < 0$, we have $R_{\max}(p) = \max\{R_1(p), R_2(p)\}$ and $R_{\min}(p) = \min\{R_1(p), R_2(p)\}$ for any $p \geq \tilde{p}$.

Finally, combining both cases of $p < \tilde{p}$ and $p \geq \tilde{p}$, we have

$$\begin{aligned}
R_{\max}(p) &= \mathbb{1}_{\{p < \tilde{p}\}} R_{F_\xi}(p) + \mathbb{1}_{\{p \geq \tilde{p}\}} \max\big\{R_1(p), R_2(p)\big\},\\
R_{\min}(p) &= \mathbb{1}_{\{p < \tilde{p}\}} R_{F_\xi}(p) + \mathbb{1}_{\{p \geq \tilde{p}\}} \min\big\{R_1(p), R_2(p)\big\},
\end{aligned}$$

which implies Proposition 2 since when $p \geq \max\{0, \tilde{p}\}$, $R_1(p) \geq R_2(p)$. $\square$

C.2. Proof of Theorem 1

Proof. When $p^*_{\max} = p^*_{\min}$, from the result in Theorem 2, the algorithm D2ACD proposed in Section 4 satisfies the following property: for any $\varepsilon > 0$,

$$\lim_{N\to\infty} \mathbb{P}\left[\max_{F\in\mathcal{F}}\{R_F(p^*_F) - R_F(p^{\mathcal{A}})\} \le \varepsilon\right] = 0.$$

Thus, by definition, the data-driven problem is identifiable.

When $p^*_{\max} \neq p^*_{\min}$, it follows from the result in Proposition 6, for any $0 < \varepsilon < \frac{\gamma b}{(\sqrt{\gamma}+1)^2}(p^*_{\max} - p^*_{\min})^2$, and any data size N and data-driven algorithm $\mathcal{A}$,

$$\mathbb{P}\left[\max_{F\in\mathcal{F}}\{R_F(p^*_F) - R_F(p^{\mathcal{A}})\} > \varepsilon\right] = 1.$$

Therefore, there does not exist a data-driven algorithm satisfying identity (2). Thus, the problem is unidentifiable. □

C.3. Proof of Proposition 3

Proof. Recall that $p^\dagger_F = \max \arg\max_{p\in\mathbb{R}} R_F(p)$ for $F \in \mathcal{F}$, and $\tilde{p} = \frac{\lambda - y}{b}$. We next divide the proof into two cases, and we fix an arbitrary distribution $F \in \mathcal{F}$ in our analysis.

Case 1: $p^\dagger_{\min} < \tilde{p}$. We prove that $p^\dagger_{\max} = p^\dagger_F = p^\dagger_{\min}$.

Since $p^\dagger_{\min}$ is the maximal maximizer of $R_{\min}(\cdot)$, and $R_{\min}(\cdot)$ is concave from Proposition 1 and its definition, we know that $R_{\min}(\cdot)$ strictly decreases in $(p^\dagger_{\min}, \tilde{p})$. Since $R_F(\cdot)$ is concave and $R_F(p) = R_{\min}(p)$ for $p \le \tilde{p}$ from Proposition 1, we must have $p^\dagger_F = p^\dagger_{\min}$. Then we have the following inequality:

$$R_{\max}(p^\dagger_{\min}) \ge R_F(p^\dagger_{\min}) = R_F(p^\dagger_F) > R_F(p), \forall p > p^\dagger_F = p^\dagger_{\min}.$$

Since F can be arbitrary, the above inequality implies

$$R_{\max}(p^\dagger_{\min}) > R_{\max}(p), \forall p > p^\dagger_F = p^\dagger_{\min},$$

which implies $p^\dagger_{\min} \ge p^\dagger_{\max}$. However, since $R_{\max}(p^\dagger_{\min}) \ge R_F(p^\dagger_{\min}) = R_F(p^\dagger_F) \ge R_F(p)$ for any $p \in \mathbb{R}$ and $F \in \mathcal{F}$, then we also have $p^\dagger_{\max} \ge p^\dagger_{\min}$. Therefore, $p^\dagger_{\max} = p^\dagger_{\min}$.

Case 2: $p^\dagger_{\min} \ge \tilde{p}$. In this case, we divide our proof into four steps.

Step 1: We prove that $p^\dagger_F \ge \tilde{p}$ and $p^\dagger_{\max} \ge \tilde{p}$. Suppose to the contrary, $p^\dagger_F < \tilde{p}$, then $p^\dagger_F < p^\dagger_{\min}$, and

$$R_F(p^\dagger_F) = R_{\min}(p^\dagger_F) \le R_{\min}(p^\dagger_{\min}) \le R_F(p^\dagger_{\min}),$$

where the identity holds since $p^\dagger_F < \tilde{p}$ by assumption. This implies $p^\dagger_{\min} \in \arg\max_{p\in\mathbb{R}} R_F(p)$, which leads to contradiction with $p^\dagger_F < p^\dagger_{\min}$ and the fact that $p^\dagger_F$ is the maximal maximizer of $R_F(\cdot)$. Similarly, we can prove $p^\dagger_{\max} \ge \tilde{p}$.

Step 2: We prove that $p^\dagger_{\min} > 0$.

If $\tilde{p} > 0$, then obviously we have $p^\dagger_{\min} \ge \tilde{p} > 0$. If $\tilde{p} \le 0$, then $\lambda - y \le 0$. By Proposition 2, for any $p \ge \tilde{p}$, $R_{\min}(p) = \min\{R_1(p), R_2(p)\}$. Since $R_1(\cdot)$, $R_2(\cdot)$ are quadratic functions, we have $p^\dagger_{\min} = \arg\max_{p\in\mathbb{R}} R_1(p)$ or $p^\dagger_{\min} = \arg\max_{p\in\mathbb{R}} R_2(p)$. It is easy to calculate that

$$\arg\max_{p\in\mathbb{R}} R_1(p) = \frac{\mathbb{E}_{F_\xi}[\min\{\xi,\lambda\}] + (1-\gamma)(y-\lambda)}{2\gamma b} > 0,$$
$$\arg\max_{p\in\mathbb{R}} R_2(p) = \frac{\mathbb{E}_{F_\xi}[\min\{\xi,\lambda\}]}{2b} > 0.$$

Thus, $p^{\dagger}_{\min} > 0$.

Step 3: We prove that $p^{\dagger}_F \geq p^{\dagger}_{\min}$. If $p^{\dagger}_{\min} = \tilde{p}$, then the result follows directly from Step 1.

Then we only need to focus on $p^{\dagger}_{\min} > \tilde{p}$. In this case, since $R_{\min}(p) = R_2(p)$ for $p > \max\{\tilde{p}, 0\}$ from Proposition 2 and $R_2(p)$ is quadratic, then $R_{\min}(p) = R_2(p)$ strictly increases in $(\max\{\tilde{p}, 0\}, p^{\dagger}_{\min})$, and decreases in $(p^{\dagger}_{\min}, \infty)$. Then we must have $p^{\dagger}_{\min} = \arg\max_{p\in\mathbb{R}} R_2(p)$. In this case, for any $p \in [\tilde{p}, p^{\dagger}_{\min})$,

$$
\begin{aligned}
&R_F(p) - R_F(p^{\dagger}_{\min}) \\
&\quad= (-bp^2 + b(p^{\dagger}_{\min})^2) + (p - p^{\dagger}_{\min}) \cdot \mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\xi < y + b\tilde{p}\}}\right] + \left(p - p^{\dagger}_{\min}\right)\mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{y + b\tilde{p} \leq \xi \leq y + bp\}}\right] \\
&\qquad+ \left(p\mathbb{E}_F\left[(y + bp) \cdot \mathbb{1}_{\{y + bp < \xi \leq y + bp^{\dagger}_{\min}\}}\right] - p^{\dagger}_{\min} \cdot \mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{y + bp < \xi \leq y + bp^{\dagger}_{\min}\}}\right]\right) \\
&\qquad+ \left(p\mathbb{E}_F\left[(y + bp) \cdot \mathbb{1}_{\{\xi > y + bp^{\dagger}_{\min}\}}\right] - p^{\dagger}_{\min}\mathbb{E}_F\left[(y + bp^{\dagger}_{\min}) \cdot \mathbb{1}_{\{\xi > y + bp^{\dagger}_{\min}\}}\right]\right) \\
&\quad\leq (-bp^2 + b(p^{\dagger}_{\min})^2) + (p - p^{\dagger}_{\min}) \cdot \mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\xi < y + b\tilde{p}\}}\right] + \left(p - p^{\dagger}_{\min}\right) \cdot (y + b\tilde{p}) \cdot \left(F(y + bp) - \gamma\right) \\
&\qquad+ \left(p - p^{\dagger}_{\min}\right) \cdot (y + bp) \cdot \left(F(y + bp^{\dagger}_{\min}) - F(y + bp)\right) + \left(p - p^{\dagger}_{\min}\right) \cdot (y + bp + bp^{\dagger}_{\min}) \cdot (1 - F(y + bp^{\dagger}_{\min})) \\
&\quad\leq \left(-bp^2 + b(p^{\dagger}_{\min})^2\right) + \left(p - p^{\dagger}_{\min}\right)\mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\xi < y + b\tilde{p}\}}\right] + \left(p - p^{\dagger}_{\min}\right) \cdot (y + b\tilde{p}) \cdot (1 - \gamma) \\
&\quad= R_2(p) - R_2(p^{\dagger}_{\min}),
\end{aligned}
\tag{C.3}
$$

where the first inequality follows from that $p^{\dagger}_{\min} > 0$ in Step 2, and the second inequality follows from $y + bp + bp^{\dagger}_{\min} \geq y + bp \geq y + b\tilde{p} > 0$ and $p - p^{\dagger}_{\min} < 0$. Note that $R_2(p) - R_2(p^{\dagger}_{\min}) = R_2(p) - \max_{p'\in\mathbb{R}}\{R_2(p')\} < 0$. Thus, $R_F(p) < R_F(p^{\dagger}_{\min})$ for any $p \in [\tilde{p}, p^{\dagger}_{\min})$, which implies $p^{\dagger}_F \geq p^{\dagger}_{\min}$.

Step 4: We prove that $p^{\dagger}_F \leq p^{\dagger}_{\max}$. If $p^{\dagger}_{\max} = \tilde{p}$, suppose $p^{\dagger}_F > p^{\dagger}_{\max}$, then we have

$$
R_{\max}(p^{\dagger}_F) \geq R_F(p^{\dagger}_F) \geq R_F(p^{\dagger}_{\max}) = R_F(\tilde{p}) = R_{\max}(\tilde{p}) = R_{\max}(p^{\dagger}_{\max}),
$$

which contradicts with $p^{\dagger}_{\max} = \max\arg\max_{p\in\mathbb{R}} R_{\max}(p)$. Thus, we must have $p^{\dagger}_F \leq p^{\dagger}_{\max}$.

Now we focus on $p^{\dagger}_{\max} > \tilde{p}$. Note that it can be easily shown that $\exists F_{\max} \in \mathcal{F}$ such that $p^{\dagger}_{\max} = p^{\dagger}_{F_{\max}}$, and the result in Step 3 holds for any $F \in \mathcal{F}$, then we must have $p^{\dagger}_{\max} = p^{\dagger}_{F_{\max}} \geq p^{\dagger}_{\min} > 0$. In this case, since $R_{\max}(p) = R_1(p)$ for $p \geq \max\{\tilde{p}, 0\}$ from Proposition 2, and $p^{\dagger}_{\max} > \max\{\tilde{p}, 0\}$, it follows from concavity of $R_1(\cdot)$ that $p^{\dagger}_{\max} = \arg\max_{p\in\mathbb{R}} R_1(p)$. In this case, for any $p > p^{\dagger}_{\max}$, we have

$$
\begin{aligned}
&R_F(p) - R_F(p^{\dagger}_{\max}) \\
&\quad= \left(-bp^2 + b(p^{\dagger}_{\max})^2\right) + (p - p^{\dagger}_{\max}) \cdot \mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\xi < y + b\tilde{p}\}}\right] + (p - p^{\dagger}_{\max}) \cdot \mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{y + b\tilde{p} \leq \xi \leq y + bp^{\dagger}_{\max}\}}\right] \\
&\qquad+ \left(p\mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{y + bp^{\dagger}_{\max} < \xi \leq y + bp\}}\right] - p^{\dagger}_{\max}\mathbb{E}_F\left[(y + bp^{\dagger}_{\max}) \cdot \mathbb{1}_{\{y + bp^{\dagger}_{\max} < \xi \leq y + bp\}}\right]\right) \\
&\qquad+ \left(p\mathbb{E}_F\left[(y + bp) \cdot \mathbb{1}_{\{\xi > y + bp\}}\right] - p^{\dagger}_{\max}\mathbb{E}_F\left[(y + bp^{\dagger}_{\max}) \cdot \mathbb{1}_{\{\xi > y + bp\}}\right]\right) \\
&\quad\leq \left(-bp^2 + b(p^{\dagger}_{\max})^2\right) + \left(p - p^{\dagger}_{\max}\right)\mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\xi < y + b\tilde{p}\}}\right] + \left(p - p^{\dagger}_{\max}\right)\left(y + bp^{\dagger}_{\max}\right)\left(F(y + bp^{\dagger}_{\max}) - \gamma\right) \\
&\qquad+ \left(p - p^{\dagger}_{\max}\right)\left(y + bp + bp^{\dagger}_{\max}\right)\left(F(y + bp) - F(y + bp^{\dagger}_{\max})\right) + \left(p - p^{\dagger}_{\max}\right)\left(y + bp + bp^{\dagger}_{\max}\right)\left(1 - F(y + bp)\right) \\
&\quad\leq \left(-bp^2 + b(p^{\dagger}_{\max})^2\right) + \left(p - p^{\dagger}_{\max}\right) \cdot \mathbb{E}_F\left[\xi \cdot \mathbb{1}_{\{\xi < y + b\tilde{p}\}}\right] + \left(p - p^{\dagger}_{\max}\right) \cdot \left(y + bp + bp^{\dagger}_{\max}\right) \cdot (1 - \gamma) \\
&\quad= R_1(p) - R_1(p^{\dagger}_{\max}),
\end{aligned}
\tag{C.4}
$$

where the first inequality follows from that $p - p^{\dagger}_{\max} > 0$ and $p > p^{\dagger}_{\max} > 0$, the second inequality follows from that $y + bp + bp^{\dagger}_{\max} \geq y + bp$ and $p - p^{\dagger}_{\max} > 0$. Note that $R_1(p) - R_1(p^{\dagger}_{\max}) = R_1(p) - \max_{p'\in\mathbb{R}} R_1(p') < 0$. Thus, $R_F(p) < R_F(p^{\dagger}_{\max})$ for any $p > p^{\dagger}_{\max}$, which implies $p^{\dagger}_F \leq p^{\dagger}_{\max}$.

Combining Step 1 – Step 4, for Case 2 that $p^\dagger_{\min} \geq \tilde{p}$, we also have $p^\dagger_{\min} \leq p^\dagger_F \leq p^\dagger_{\max}$.

Since $R_{\min}(p)$ is concave, it then follows $p^*_{\min} = \text{Proj}(p^\dagger_{\min}, [\underline{p}, \overline{p}])$. In addition, in both Case 1 and Case 2, $R_{\max}(p)$ increases when $p \leq p^\dagger_{\max}$ and decreases when $p > p^\dagger_{\max}$. Thus, $p^*_{\max} = \text{Proj}(p^\dagger_{\max}, [\underline{p}, \overline{p}])$. Therefore, $p^*_{\min} \leq p^*_F \leq p^*_{\max}$ for any $F \in \mathcal{F}$. □

C.4. Proof of Proposition 4

Proof. We first show the "if" direction. The proof is divided into three cases.

(i) When $p^\dagger_{\max} = p^\dagger_{\min} \leq \tilde{p}$ in Case (i) happens, from the proof of Proposition 3, we have $p^\dagger_{\max} = p^\dagger_{\min} = p^\dagger_F$ for any $F \in \mathcal{F}$, and thus, $p^*_{\max} = p^*_{\min}$.

(ii) When $p^\dagger_{\max} > p^\dagger_{\min} \geq \tilde{p}$ and $p^\dagger_{\max} \leq \underline{p}$ in Case (ii) happen, it can be verified that $R_{\max}(\cdot)$ decreases in $[\underline{p}, \overline{p}]$, and therefore, from Proposition 3 and the definitions of $p^*_{\max}$ and $p^*_{\min}$, $p^*_{\max} = p^*_{\min} = \underline{p}$.

(iii) When $p^\dagger_{\max} > p^\dagger_{\min} \geq \tilde{p}$ and $p^\dagger_{\min} \geq \overline{p}$ in Case (iii) happen, it can be verified that $R_{\min}(\cdot)$ increases in $[\underline{p}, \overline{p}]$, and therefore, from Proposition 3 and the definitions of $p^*_{\max}$ and $p^*_{\min}$, $p^*_{\max} = p^*_{\min} = \overline{p}$.

We next show the "only if" direction. Suppose all the three cases described in Proposition 4 fail to hold, then it can be verified that we must have $p^\dagger_{\max} \neq p^\dagger_{\min}$, $p^\dagger_{\min} \geq \tilde{p}$, $p^\dagger_{\max} > \underline{p}$ and $p^\dagger_{\min} < \overline{p}$. From the proof of Proposition 3, we then have $p^\dagger_{\max} > p^\dagger_{\min}$, $p^*_{\max} = \text{Proj}(p^\dagger_{\max}, [\underline{p}, \overline{p}])$, and $p^*_{\min} = \text{Proj}(p^\dagger_{\min}, [\underline{p}, \overline{p}])$. By discussing the following four cases: $p^\dagger_{\min} \leq \underline{p} < p^\dagger_{\max} \leq \overline{p}$, $p^\dagger_{\min} \leq \underline{p} < \overline{p} < p^\dagger_{\max}$, $\underline{p} < p^\dagger_{\min} < p^\dagger_{\max} \leq \overline{p}$, and $\underline{p} < p^\dagger_{\min} < \overline{p} < p^\dagger_{\max}$, we can verify that $p^*_{\max} \neq p^*_{\min}$, which completes the proof of the "only if" direction. □

C.5. Proof of Proposition 5

Proof. For notation convenience, we define $p^\dagger_i = \arg\max_{p \in \mathbb{R}} R_i(p)$, $i = 1, 2$. In this proof, we add a superscript $'$ to the relevant quantities associated with the dataset $\mathcal{S}'$.

We first note that from Theorem 1 and the proof of Proposition 3, the following claim holds:

Claim 1. *Any data-driven problem is identifiable if and only if one of the following conditions holds: (i)* $p^\dagger_{\min} = p^\dagger_{\max} \in (\underline{p}, \overline{p})$*; (ii)* $p^\dagger_{\max} \leq \underline{p}$*; (iii)* $p^\dagger_{\min} \geq \overline{p}$.

We next consider the above three cases for the problem with dataset $\mathcal{S}$ of parameter λ.

Case 1: $p^\dagger_{\min} = p^\dagger_{\max} \in (\underline{p}, \overline{p})$. In this case, from the proof of Proposition 3, we have $p^\dagger_{\min} = p^\dagger = p^\dagger_{\max} \leq \tilde{p}$. Since $\tilde{p}' = (\lambda' - y)/b > (\lambda - y)/b = \tilde{p}$ and $p^\dagger = p^{\dagger'}$, we have $p^{\dagger'} < \tilde{p}'$. In this case, it can be verified from concavity that $p^{\dagger'}_{\min} = p^{\dagger'} < \tilde{p}'$. Then it follows from Case 1 in the proof of Proposition 3 that $p^{\dagger'}_{\min} = p^{\dagger'}_{\max}$ and $p^{*'}_{\max} = p^{*'}_{\min}$, which then implies from Theorem 1 that the problem with the dataset $\mathcal{S}'$ is identifiable.

Case 2: $p^\dagger_{\max} \leq \underline{p}$. In this case, we can assume $p^\dagger_{\min} \geq \tilde{p}$. Otherwise, if $p^\dagger_{\min} < \tilde{p}$, from Case 1 in the proof of Proposition 3, we have $p^\dagger_{\min} = p^\dagger_{\max} = p^\dagger < \tilde{p}$, and the result can be obtained from similar arguments in Case 1. In addition, it also suffices to consider $p^{\dagger'}_{\min} \geq \tilde{p}'$, since otherwise, from Case 1 in the proof of Proposition 3, $p^{\dagger'}_{\min} = p^{\dagger'}_{\max}$ and $p^{*'}_{\min} = p^{*'}_{\max}$, then the problem with dataset $\mathcal{S}'$ is already identifiable.

Since $p^\dagger_{\min} \geq \tilde{p}$, from Case 2 in the proof of Proposition 3, we have either $p^\dagger_{\min} = p^\dagger_{\max} = \tilde{p}$, or $p^\dagger_{\max} = p^\dagger_1$. If the first case happens, then the result follows from similar arguments in Case 1. Similarly, since $p^{\dagger'}_{\min} \geq \tilde{p}'$, we have either $p^{\dagger'}_{\min} = p^{\dagger'}_{\max} = \tilde{p}$, or $p^{\dagger'}_{\max} = p^{\dagger'}_1$. If the first case happens, the problem with dataset $\mathcal{S}'$ is already identifiable. Thus, we only need to consider the case when both $p^\dagger_{\max} = p^\dagger_1$ and $p^{\dagger'}_{\max} = p^{\dagger'}_1$ hold. In this case, if we can prove $p_1^{\dagger'} \le p_1^\dagger$, it then follows that $p_{\max}^{\dagger'} = p_1^{\dagger'} \le p_1^\dagger = p_{\max}^\dagger \le \underline{p}$, which implies from Claim 1 that the problem with $\mathcal{S}'$ is also identifiable. We now prove $p_1^{\dagger'} \le p_1^\dagger$ as follows:

$$\begin{aligned}
p_1^\dagger - p_1^{\dagger'} &= \frac{\mathbb{E}[\xi \cdot \mathbb{1}_{\{\xi<\lambda\}}] + (1-\gamma)y}{2\gamma b} - \frac{\mathbb{E}[\xi \cdot \mathbb{1}_{\{\xi<\lambda'\}}] + (1-\gamma')y}{2\gamma' b} \\
&= \frac{1}{2b\gamma\gamma'}\Big(\mathbb{E}[\xi \cdot \mathbb{1}_{\{\xi<\lambda\}}]\cdot \mathbb{E}[\mathbb{1}_{\{\xi<\lambda'\}}] - \mathbb{E}[\xi \cdot \mathbb{1}_{\{\xi<\lambda'\}}]\cdot \mathbb{E}[\mathbb{1}_{\{\xi<\lambda\}}] + \mathbb{E}[\mathbb{1}_{\{\lambda\le\xi<\lambda'\}}]y\Big) \\
&= \frac{1}{2b\gamma\gamma'}\Big(\mathbb{E}[\xi \cdot \mathbb{1}_{\{\xi<\lambda\}}]\cdot \mathbb{E}[\mathbb{1}_{\{\lambda\le\xi<\lambda'\}}] - \mathbb{E}[\xi \cdot \mathbb{1}_{\{\lambda\le\xi<\lambda'\}}]\cdot \mathbb{E}[\mathbb{1}_{\{\xi<\lambda\}}] + \mathbb{E}[\mathbb{1}_{\{\lambda\le\xi<\lambda'\}}]y\Big) \\
&\ge \frac{1}{2b\gamma\gamma'}\Big(\mathbb{E}[\xi \cdot \mathbb{1}_{\{\xi<\lambda\}}]\cdot \mathbb{E}[\mathbb{1}_{\{\lambda\le\xi<\lambda'\}}] - \mathbb{E}[\lambda' \cdot \mathbb{1}_{\{\lambda\le\xi<\lambda'\}}]\cdot \mathbb{E}[\mathbb{1}_{\{\xi<\lambda\}}] + \mathbb{E}[\mathbb{1}_{\{\lambda\le\xi<\lambda'\}}]y\Big) \\
&\ge \frac{1}{2b\gamma\gamma'}\Big(\mathbb{E}[\xi \cdot \mathbb{1}_{\{\xi<\lambda\}}]\cdot \mathbb{E}[\mathbb{1}_{\{\lambda\le\xi<\lambda'\}}] - \mathbb{E}[(y+bp^\dagger) \cdot \mathbb{1}_{\{\lambda\le\xi<\lambda'\}}]\cdot \mathbb{E}[\mathbb{1}_{\{\xi<\lambda\}}] + \mathbb{E}[\mathbb{1}_{\{\lambda\le\xi<\lambda'\}}]y\Big) \\
&= \frac{1}{2b\gamma\gamma'}\cdot \mathbb{E}[\mathbb{1}_{\{\lambda\le\xi<\lambda'\}}]\cdot\Big(y - \mathbb{E}\big[\big(y+bp^\dagger-\xi\big)\big)\cdot \mathbb{1}_{\{\xi<\lambda\}}\big]\Big) \\
&\ge \frac{1}{2b\gamma\gamma'}\cdot \mathbb{E}[\mathbb{1}_{\{\lambda\le\xi<\lambda'\}}]\cdot\Big(y - \mathbb{E}\big[\big(y+bp^\dagger-\xi\big)\cdot \mathbb{1}_{\{\xi<y+bp^\dagger\}}\big]\Big) \\
&= \frac{1}{2b\gamma\gamma'}\cdot \mathbb{E}[\mathbb{1}_{\{\lambda\le\xi<\lambda'\}}]\cdot \mathbb{E}\big[\min\{y, bp^\dagger-\xi\}\big] \\
&\ge 0,
\end{aligned}$$

where the second identity follows from $\gamma = \mathbb{E}[\mathbb{1}_{\{\xi<\lambda\}}]$ and $\gamma' = \mathbb{E}[\mathbb{1}_{\{\xi<\lambda'\}}]$, the second inequality holds because $p^\dagger \ge \tilde{p}'$ (from Case 2 – Step 1 in the proof of Proposition 3) and $\tilde{p}' = (\lambda' - y)/b$ imply $\lambda' \le y + bp^\dagger$, the third inequality holds since $\lambda \le y + bp^\dagger$, and $\xi < \lambda$ implies $y + bp^\dagger - \xi \ge \lambda - \xi > 0$, and the last inequality holds since $\mathbb{E}[\min\{y, bp^\dagger - \xi\}]$ is the expected sales at the maximizer $p^\dagger$ and is therefore greater than zero. This completes the proof of $p_1^{\dagger'} \le p_1^\dagger$.

Case 3: $p_{\min}^\dagger \ge \overline{p}$. In this case, similar to the reason in Case 2, it suffices to consider $p_{\min}^\dagger \ge \tilde{p}$, and $p_{\min}^{\dagger'} \ge \tilde{p}'$. Then from Case 2 in the proof of Proposition 3, we have four different subcases: (i) $p_{\min}^\dagger = \tilde{p}$ and $p_{\min}^{\dagger'} = \tilde{p}'$; (ii) $p_{\min}^\dagger = \tilde{p}$ and $p_{\min}^{\dagger'} = p_2^{\dagger'}$; (iii) $p_{\min}^\dagger = p_2^\dagger$, $p_{\min}^{\dagger'} = \tilde{p}'$; (iv) $p_{\min}^\dagger = p_2^\dagger$ and $p_{\min}^{\dagger'} = p_2^{\dagger'}$.

In subcase (i), it follows from $\lambda < \lambda'$ that $p_{\min}^{\dagger'} = \tilde{p}' > \tilde{p} = p_{\min}^\dagger \ge \overline{p}$. In subcase (ii), we have $p_{\min}^{\dagger'} \ge \tilde{p}' > \tilde{p} = p_{\min}^\dagger \ge \overline{p}$. In subcase (iii), since

$$p_2^\dagger = \frac{\mathbb{E}[\min\{\xi,\lambda\}]}{2b} \le \frac{\mathbb{E}[\min\{\xi,\lambda'\}]}{2b} = p_2^{\dagger'},$$

we have $\overline{p} \le p_{\min}^\dagger = p_2^\dagger \le p_2^{\dagger'} \le \tilde{p}' = p_{\min}^{\dagger'}$. In subcase (iv), as proved in case (iii), $p_2^{\dagger'} \ge p_2^\dagger$, then it follows that $p_{\min}^{\dagger'} = p_2^{\dagger'} \ge p_2^\dagger = p_{\min}^\dagger \ge \overline{p}$. Combining all the above four subcases, we have $p_{\min}^{\dagger'} \ge \overline{p}$, which implies from Claim 1 that the problem with $\mathcal{S}'$ is identifiable. □

C.6. Proof of Proposition 6

Proof. Similar to the proof of Proposition 5, for notation convenience, we define $p_i^\dagger = \arg\max_{p\in\mathbb{R}} R_i(p)$, $i = 1, 2$.

Note that from Proposition 7, it suffices to derive the lower bound for $\min_{p\in[\underline{p},\overline{p}]} \max\big\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\big\}$. We first prove that when $p^*_{\max} \ne p^*_{\min}$,

$$\begin{aligned}
&\min_{p\in[\underline{p},\overline{p}]} \max\big\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\big\} \\
&\quad= \min_{p\in[p^*_{\min}, p^*_{\max}]} \max\big\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\big\}. \qquad (C.5)
\end{aligned}$$

Since $p^*_{\max} \neq p^*_{\min}$, it can be verified from the proof of Proposition 3 that $p^\dagger_2 \le p^\dagger_{\min} \le p^*_{\min} < p^*_{\max} \le p^\dagger_{\max} = p^\dagger_1$. In this case, we have

$$\begin{aligned}
&\min_{p\in[\underline{p},\overline{p}],p\le p^*_{\min}} \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\right\} \\
&\quad\ge \min_{p\in[\underline{p},\overline{p}],p\le p^*_{\min}} \{R_{\max}(p^*_{\max}) - R_{\max}(p)\} = R_{\max}(p^*_{\max}) - R_{\max}(p^*_{\min}) \\
&\quad= \max\{R_{\max}(p^*_{\max}) - R_{\max}(p^*_{\min}), R_{\min}(p^*_{\min}) - R_{\min}(p^*_{\min}\},
\end{aligned} \tag{C.6}$$

where the identity holds since $p^*_{\min} < p^*_{\max} \le p^\dagger_{\max}$ and $R_{\max}(p)$ increases when $p \le p^\dagger_{\max}$. Similarly,

$$\begin{aligned}
&\min_{p\in[\underline{p},\overline{p}],p\ge p^*_{\max}} \max\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\} \\
&\quad\ge \min_{p\in[\underline{p},\overline{p}],p\ge p^*_{\max}} \{R_{\min}(p^*_{\min}) - R_{\min}(p)\} = R_{\min}(p^*_{\min}) - R_{\min}(p^*_{\max}) \\
&\quad= \max\{R_{\max}(p^*_{\max}) - R_{\max}(p^*_{\max}), R_{\min}(p^*_{\min}) - R_{\min}(p^*_{\max})\}.
\end{aligned} \tag{C.7}$$

Combining inequalities (C.6) and (C.7), we know that the minimum value of the function $\max\{R_{\max}(p^*_{\max}) - R_{\max}(p), R_{\min}(p^*_{\min}) - R_{\min}(p)\}$ must be taken at $p \in [p^*_{\min}, p^*_{\max}]$, which gives identity (C.5).

With identity (C.5), it suffices to prove

$$\min_{p\in[p^*_{\min},p^*_{\max}]} \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\right\} \ge \frac{\gamma b}{(\sqrt{\gamma}+1)^2}(p^*_{\max} - p^*_{\min})^2.$$

Since $R_{\max}(p) = R_1(p)$ and $R_{\min}(p) = R_2(p)$ for any $p \ge \max\{\tilde{p}, 0\}$, and $p^*_{\min} \ge \tilde{p}$ (otherwise, it can be easily verified that $p^*_{\max} = p^*_{\min}$, leading to contradiction with the assumption that $p^*_{\max} \neq p^*_{\min}$), we have $R_{\max}(p) = R_1(p)$ and $R_{\min}(p) = R_2(p)$ for all $p \in [p^*_{\min}, p^*_{\max}]$. Then for any $p \in [p^*_{\min}, p^*_{\max}]$, we have

$$R_{\max}(p^*_{\max}) - R_{\max}(p) = R_1(p^*_{\max}) - R_1(p) \ge R_1(p^\dagger_1) - R_1(p^\dagger_1 - (p^*_{\max} - p)) = \gamma b (p^*_{\max} - p)^2, \tag{C.8}$$

where the inequality follows from concavity of $R_1(\cdot)$ and $p \le p^*_{\max} \le p^\dagger_{\max} = p^\dagger_1$, and the identity follows from Taylor expansion of the quadratic function $R_1(p)$ at the minimizer $p^\dagger_1$. Similarly,

$$R_{\min}(p^*_{\min}) - R_{\min}(p) = R_2(p^*_{\min}) - R_2(p) \ge R_2(p^\dagger_2) - R_2(p^\dagger_2 + (p - p^*_{\min})) = b(p - p^*_{\min})^2. \tag{C.9}$$

Combining (C.8) and (C.9), we have

$$\begin{aligned}
&\min_{p\in[p^*_{\min},p^*_{\max}]} \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\right\} \\
&\quad\ge \min_{p\in[p^*_{\min},p^*_{\max}]} \max\left\{\gamma b(p^*_{\max} - p)^2,\, b(p - p^*_{\min})^2\right\} \\
&\quad= \frac{\gamma b}{(\sqrt{\gamma}+1)^2}(p^*_{\max} - p^*_{\min})^2,
\end{aligned}$$

where the last identity follows from simple calculation. □

Appendix D: Omitted Proof in Section 4

D.1. Proof of Lemma 1

We prove Lemma 1 for both cases of continuous and discrete base demand distributions in the range $[0, \lambda)$. Recall that in Step 1 of D2ACD, for each $i \in [K]$, the breakpoints of the empirical c.d.f. $\hat{F}^{\text{SAA}}_i(\cdot)$ are denoted by $\hat{\beta}_{i,1} < \hat{\beta}_{i,2} < \ldots < \hat{\beta}_{i,\hat{M}_i}$. Let $\mathcal{B} := \cup_{i\in[K]}\{\hat{\beta}_{i,j} + bp_i : j \in [\hat{M}_i - 1]\}$. Note that we exclude the last breakpoint

$\hat{\beta}_{i,\hat{M}_i}$ for each price p_i to facilitate our later discussion. Besides, let $\hat{M}$ be the cardinality of $\mathcal{B}$, and $\hat{\beta}_1 < \hat{\beta}_2 < \ldots < \hat{\beta}_{\hat{M}}$ be all the non-repetitive elements in $\mathcal{B}$. Now we focus on an arbitrary and fixed $i \in [K]$.

Consider a censored random demand at price p_i with the distribution identical to $\min\{\xi - bp_i, \lambda - bp_i\}$, and the c.d.f. is denoted by $F_i^{cens}(x)$. Since $y_i \le \lambda - bp_i$, then $F_i^{\text{cens}}(x)$ equals the true c.d.f. for $\min\{\xi - bp_i, y_i\}$ when $x < y_i$. Since $\hat{F}_i^{SAA}(\cdot)$ is the empirical c.d.f. for $\min\{\xi - bp_i, y_i\}$, from Lemma A.2, we get, for any $\alpha > 0$, the event

$$|\hat{F}_i^{SAA}(x) - F_i^{cens}(x)| \le \alpha \text{ for all } x < y_i \tag{D.1}$$

holds with probability at least $1 - 2e^{-2N_i\alpha^2}$.

The remaining proof consists of four major steps.

Step 1. Construct two modified quantile functions $\tilde{F}_i^{-1}$ and $\hat{F}_i^{-1}$. We first construct a piecewise linear function $\tilde{F}_i(x)$ defined on $x \in [\hat{\beta}_1 - bp_i, \hat{\beta}_{\hat{M}} - bp_i]$ based on the true c.d.f. $F_i^{cens}(x)$ and the set of breakpoints $\mathcal{B}$. Specifically, for any $x \in [\hat{\beta}_k - bp_i, \hat{\beta}_{k+1} - bp_i]$ and $k \in [\hat{M} - 1]$, let $\tilde{F}_i(x)$ be a linear function connecting $(\hat{\beta}_k, F_i^{cens}(\hat{\beta}_k - bp_i))$ and $(\hat{\beta}_{k+1}, F_i^{cens}(\hat{\beta}_{k+1} - bp_i))$, i.e.,

$$\tilde{F}_i(x) = \frac{F_i^{\text{cens}}(\hat{\beta}_{k+1} - bp_i) - F_i^{\text{cens}}(\hat{\beta}_k - bp_i)}{\hat{\beta}_{k+1} - \hat{\beta}_k}\big(x - (\hat{\beta}_k - bp_i)\big) + F_i^{\text{cens}}(\hat{\beta}_k - bp_i).$$

When ξ is continuous on $[0, \lambda)$, $F_i^{cens}(\cdot)$ is strictly increasing in its support, so $\tilde{F}_i(\cdot)$ is also strictly increasing on $[\hat{\beta}_1 - bp_i, \hat{\beta}_{\hat{M}} - bp_i]$. When ξ is discrete on $[0, \lambda)$, for each $j \in [K]$, the largest breakpoint $\hat{\beta}_{j,\hat{M}_j}$ under price p_j may appear due to the censoring effect but does not come from the true breakpoints of $\xi - bp_j$, i.e., $\hat{\beta}_{j,\hat{M}_j} = y_j$ but $y_j \notin \{\beta_1 - bp_j, \beta_2 - bp_j, \ldots, \beta_M - bp_j\}$ (recall from Assumption 1 that $\beta_1, \beta_2, \ldots, \beta_M$ are the breakpoints for the true c.d.f. of ξ on the left-hand side of λ). Therefore, in the definition of $\mathcal{B}$, we do not include $\{\hat{\beta}_{j,\hat{M}_j} : j \in [K]\}$. With this construction, we can easily verify that the set $\{\hat{\beta}_k - bp_i : k \in [\hat{M}]\}$ must belong to the whole set of breakpoints of the c.d.f. $F_i^{\text{cens}}(\cdot)$. Thus, when ξ is discrete on $[0, \lambda)$, $F_i^{\text{cens}}(\hat{\beta}_1 - bp_i) < F_i^{\text{cens}}(\hat{\beta}_2 - bp_i) < \ldots < F_i^{\text{cens}}(\hat{\beta}_{\hat{M}} - bp_i)$ and $\tilde{F}_i(\cdot)$ also strictly increases on $[\hat{\beta}_1 - bp_i, \hat{\beta}_{\hat{M}} - bp_i]$. Then we let $\tilde{F}_i^{-1}(\cdot)$ be the inverse function of $\tilde{F}_i(\cdot)$ on $[\hat{\beta}_1 - bp_i, \hat{\beta}_{\hat{M}} - bp_i]$. For $0 \le x < \tilde{F}_i(\hat{\beta}_1 - bp_i)$, we set $\tilde{F}_i^{-1}(x) = \hat{\beta}_1 - bp_i$, and for $\tilde{F}_i(\hat{\beta}_{\hat{M}} - bp_i) < x \le 1$, we set $\tilde{F}_i^{-1}(x) = \hat{\beta}_{\hat{M}} - bp_i$. In this way, $\tilde{F}_i^{-1}$ is well defined on $[0, 1]$. Besides, it can be verified that for any $x_1, x_2 \in [\hat{\beta}_1 - bp_i, \hat{\beta}_{\hat{M}} - bp_i]$ with $x_1 < x_2$,

$$\tilde{F}_i(x_2) \ge \tilde{F}_i(x_1) + \underline{g}(x_2 - x_1). \tag{D.2}$$

To see this, when ξ is discrete on $[0, \lambda)$, $\mathcal{B} \subseteq \{\beta_1, \beta_2, \ldots, \beta_M\}$, and it follows from the definition of $\underline{g}$ and construction of $\tilde{F}_i(\cdot)$ that the slope of each piece in $\tilde{F}_i(\cdot)$ is greater than $\underline{g}$, which then implies (D.2). When ξ is continuous on $[0, \lambda)$, Assumption 1 implies $F_i^{\text{cens}}(x_2) \ge F_i^{\text{cens}}(x_1) + \underline{g}(x_2 - x_1)$ for any $0 \le x_1 < x_2 < \lambda - bp_i$, and since $\tilde{F}_i(x) = F_i^{\text{cens}}(x)$ when $x + bp_i \in \mathcal{B}$, the slope of each piece in $\tilde{F}_i(\cdot)$ is also greater than $\underline{g}$, which also implies (D.2).

We then define a piecewise linear function $\hat{F}_i(\cdot)$ on $[\hat{\beta}_{i,1}, \hat{\beta}_{i,\hat{M}_i}]$ based on the empirical c.d.f. $\hat{F}_i^{\text{SAA}}(\cdot)$ and the set of breakpoints $\{\hat{\beta}_{i,1}, \hat{\beta}_{i,2}, \ldots, \hat{\beta}_{i,\hat{M}_i}\}$. For any $x \in [\hat{\beta}_{i,k}, \hat{\beta}_{i,k+1}]$ and $k \in [\hat{M}_i - 1]$, let $\hat{F}_i(x)$ be a linear function connecting the $(\hat{\beta}_{i,k}, \hat{F}_i^{\text{SAA}}(\hat{\beta}_{i,k}))$ and $(\hat{\beta}_{i,k+1}, \hat{F}_i^{\text{SAA}}(\hat{\beta}_{i,k+1}))$, i.e.,

$$\hat{F}_i(x) = \frac{\hat{F}_i^{\text{SAA}}(\hat{\beta}_{i,k+1}) - \hat{F}_i^{\text{SAA}}(\hat{\beta}_{i,k})}{\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k}}\big(x - \hat{\beta}_{i,k}\big) + \hat{F}_i^{\text{SAA}}(\hat{\beta}_{i,k}).$$

With this construction, $\hat{F}_i^{-1}(\cdot)$ defined in Step 1 of D2ACD is the inverse of $\hat{F}_i(x)$ when $x \in [\hat{\beta}_{i,1}, \hat{\beta}_{i,\hat{M}_i-1}]$, and equals $\hat{\beta}_{i,1}$ when $0 \le x < \hat{F}_i^{\mathrm{SAA}}(\hat{\beta}_{i,1})$.

For each $k \in [\hat{M}_i - 1]$, since $\hat{\beta}_{i,k} + bp_i \in \mathcal{B}$ and $\tilde{F}_i(\cdot)$ is identical to $F_i^{\mathrm{cens}}(\cdot)$ in the set $\{\hat{\beta}_1 - bp_i, \hat{\beta}_2 - bp_i, \ldots, \hat{\beta}_{\hat{M}} - bp_i\}$, it follows that $\tilde{F}_i(\hat{\beta}_{i,k}) = F_i^{cens}(\hat{\beta}_{i,k})$. This, combined with $\hat{F}_i(\hat{\beta}_{i,k}) = \hat{F}_i^{SAA}(\hat{\beta}_{i,k})$, $\hat{\beta}_{i,\hat{M}_i-1} < y_i$ and inequality (D.1), implies

$$|\hat{F}_i(\hat{\beta}_{i,k}) - \tilde{F}_i(\hat{\beta}_{i,k})| \le \alpha \text{ for all } k \in [\hat{M}_i - 1]. \tag{D.3}$$

Step 2. Bound the range of the true $\hat{\gamma}_{\min}$-quantile. Recall that $\hat{\gamma}_{\min} = \min_{i\in[K]} \hat{\gamma}_i$. Define $d_i := \tilde{F}_i^{-1}(\hat{\gamma}_{\min})$ and $\hat{d}_i := \hat{F}_i^{-1}(\hat{\gamma}_{\min})$. Since $\hat{\gamma}_{\min} \le \hat{\gamma}_i = \hat{F}_i^{\mathrm{SAA}}(\hat{\beta}_{i,\hat{M}_i-1})$, we then have $\hat{\beta}_{i,1} \le \hat{d}_i \le \hat{\beta}_{i,\hat{M}_i-1}$. Let $k \in [\hat{M}_i - 2]$ denote the index of breakpoint $\hat{\beta}_{i,k}$ such that $\hat{\beta}_{i,k} \le \hat{d}_i \le \hat{\beta}_{i,k+1}$.

Conditioning on event (D.1), we bound the range of d_i by considering two cases.

Case 1: $\tilde{F}_i(\hat{\beta}_{i,k}) < \hat{\gamma}_{\min}$. In this case, since $\hat{\beta}_{i,k}$ is also a breakpoint of $\tilde{F}_i$ and $d_i = \tilde{F}_i^{-1}(\hat{\gamma}_{\min})$, we must have $d_i > \hat{\beta}_{i,k}$. For any $x > \hat{\beta}_{i,k} + \big(\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + \alpha\big)/\underline{g}$ and $x \in [\hat{\beta}_1 - bp_i, \hat{\beta}_{\hat{M}} - bp_i]$, we have

$$\tilde{F}_i(x) \ge \tilde{F}_i(\hat{\beta}_{i,k}) + \underline{g}(x - \hat{\beta}_{i,k}) \ge \hat{F}_i(\hat{\beta}_{i,k}) - \alpha + \underline{g}(x - \hat{\beta}_{i,k}) > \hat{\gamma}_{\min},$$

where the first inequality follows from (D.2) and $\hat{\beta}_{i,k} \in [\hat{\beta}_1 - bp_i, \hat{\beta}_{\hat{M}} - bp_i]$, and the second inequality follows from (D.3). Hence, $\hat{\beta}_{i,k} < d_i \le \hat{\beta}_{i,k} + \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + \alpha}{\underline{g}}$.

Case 2: $\tilde{F}_i(\hat{\beta}_{i,k}) \ge \hat{\gamma}_{\min}$. In this case, if $d_i = \hat{\beta}_{i,1}$, since $\hat{\beta}_{i,k}$ is also a breakpoint of $\tilde{F}_i$, we have $d_i \le \hat{\beta}_{i,k}$. If $d_i > \hat{\beta}_{i,1}$, then $\tilde{F}_i(d_i) = \hat{\gamma}_{\min} \le \tilde{F}_i(\hat{\beta}_{i,k})$, then $d_i \le \hat{\beta}_{i,k}$. Again, due to inequalities (D.2) and (D.3), for any $x < \hat{\beta}_{i,k} - \big(\hat{F}_i(\hat{\beta}_{i,k}) + \alpha - \hat{\gamma}_{\min}\big)/\underline{g}$ and $x \in [\hat{\beta}_1 - bp_i, \hat{\beta}_{\hat{M}} - bp_i]$, we have

$$\tilde{F}_i(x) \le \tilde{F}_i(\hat{\beta}_{i,k}) - \underline{g}(\hat{\beta}_{i,k} - x) \le \hat{F}_i(\hat{\beta}_{i,k}) + \alpha - \underline{g}(\hat{\beta}_{i,k} - x) < \hat{\gamma}_{\min}.$$

Hence, $\hat{\beta}_{i,k} - \frac{\hat{F}_i(\hat{\beta}_{i,k}) + \alpha - \hat{\gamma}_{\min}}{\underline{g}} \le d_i \le \hat{\beta}_{i,k}$.

Step 3. Bound the range of the empirical $\hat{\gamma}_{\min}$-quantile. If $\hat{\gamma}_{\min} < \hat{F}_i(\hat{\beta}_{i,1})$, then $\hat{d}_i = \hat{\beta}_{i,1}$. Otherwise, we have $\hat{\beta}_{i,k} \le \hat{d}_i \le \hat{\beta}_{i,k+1}$ and $\hat{F}_i(\hat{\beta}_{i,k}) \le \hat{\gamma}_{\min} \le \hat{F}_i(\hat{\beta}_{i,k+1})$. For this case of $\hat{\gamma}_{\min} \ge \hat{F}_i(\hat{\beta}_{i,1})$, we consider the following two cases to bound the range of $\hat{d}_i$ conditioning on event (D.1).

Case 1: $\hat{\beta}_{i,k+1} \le \hat{\beta}_{i,k} + \big(\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha\big)/\underline{g}$. In this case, we have $\hat{\beta}_{i,k} \le \hat{d}_i \le \hat{\beta}_{i,k+1} \le \hat{\beta}_{i,k} + \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}}$. Therefore, $\hat{d}_i \in [\hat{\beta}_{i,k}, \hat{\beta}_{i,k} + \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}}]$.

Case 2: $\hat{\beta}_{i,k+1} > \hat{\beta}_{i,k} + \big(\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha\big)/\underline{g}$. In this case, we have, for any $x \in (\hat{\beta}_{i,k} + \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}}, \hat{\beta}_{i,k+1}] \subset (\hat{\beta}_{i,k}, \hat{\beta}_{i,k+1}]$,

$$\begin{aligned}\hat{F}_i(x) &= \hat{F}_i(\hat{\beta}_{i,k}) + \frac{\hat{F}_i(\hat{\beta}_{i,k+1}) - \hat{F}_i(\hat{\beta}_{i,k})}{\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k}} \cdot (x - \hat{\beta}_{i,k}) \\ &> \hat{F}_i(\hat{\beta}_{i,k}) + \frac{\hat{F}_i(\hat{\beta}_{i,k+1}) - \hat{F}_i(\hat{\beta}_{i,k})}{\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k}} \cdot \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}},\end{aligned} \tag{D.4}$$

where the identity holds since $\hat{F}_i$ is linear in $[\hat{\beta}_{i,k}, \hat{\beta}_{i,k+1}]$. Moreover,

$$\hat{F}_i(\hat{\beta}_{i,k+1}) \ge \tilde{F}_i(\hat{\beta}_{i,k+1}) - \alpha \ge \tilde{F}_i(\hat{\beta}_{i,k}) + \underline{g}(\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k}) - \alpha \ge \hat{F}_i(\hat{\beta}_{i,k}) + \underline{g}(\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k}) - 2\alpha, \tag{D.5}$$

where the first and last inequalities hold due to inequality (D.3) and $k \in [\hat{M}_i - 2]$, and the second inequality follows from inequality (D.2) and $k \in [\hat{M}_i - 2]$. Hence, combining inequalities (D.4) and (D.5), we have

$$\begin{aligned}\hat{F}_i(x) &> \hat{F}_i(\hat{\beta}_{i,k}) + \frac{\underline{g}(\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k}) - 2\alpha}{\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k}} \cdot \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}} \\ &= \hat{F}_i(\hat{\beta}_{i,k}) + \hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha\left(1 - \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}(\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k})}\right) \\ &= \hat{\gamma}_{\min} + 2\alpha\left(1 - \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}(\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k})}\right).\end{aligned}$$

Since $\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k} > \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}}$ by assumption, $1 - \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}(\hat{\beta}_{i,k+1} - \hat{\beta}_{i,k})} > 0$ and $\hat{F}_i(x) > \hat{\gamma}_{\min}$. Therefore, $\hat{d}_i \le \hat{\beta}_{i,k} + \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}}$.

Considering all the possible ranges of d_i in Step 2 and $\hat{d}_i$ in Step 3, we have

$$|d_i - \hat{d}_i| \le \left(\hat{\beta}_{i,k} + \frac{\hat{\gamma}_{\min} - \hat{F}_i(\hat{\beta}_{i,k}) + 2\alpha}{\underline{g}}\right) - \left(\hat{\beta}_{i,k} - \frac{\hat{F}_i(\hat{\beta}_{i,k}) + \alpha - \hat{\gamma}_{\min}}{\underline{g}}\right) = \frac{3\alpha}{\underline{g}}. \tag{D.6}$$

Step 4. Bound the estimation error $|b - \hat{b}|$. Since for any $1 \le i < j \le K$, $\tilde{F}_i(\cdot)$ and $\tilde{F}_j(\cdot)$ are identical except being shifted by a distance $b|p_i - p_j|$ horizontally, and $\tilde{F}_i^{-1}(\cdot)$ is the inverse of $\tilde{F}_i(\cdot)$, it follows that for any $x \in [0,1]$ and $i, j \in [K]$, we have $b(p_i - p_j) = -(\tilde{F}_i^{-1}(x) - \tilde{F}_j^{-1}(x))$. This implies

$$b = -\frac{\sum_{i=1}^K (p_i - \frac{1}{K}\sum_{j=1}^K p_j) \cdot \left(d_i - \frac{1}{K}\sum_{j=1}^K d_j\right)}{\sum_{i=1}^K (p_i - \frac{1}{K}\sum_{j=1}^K p_j)^2}. \tag{D.7}$$

In addition, since $b^\dagger$ is the least-square solution to the linear regression defined in Step 1, it can be easily verified from the optimality condition that

$$b^\dagger = -\frac{\sum_{i=1}^K (p_i - \frac{1}{K}\sum_{j=1}^K p_j) \cdot \left(\hat{d}_i - \frac{1}{K}\sum_{j=1}^K \hat{d}_j\right)}{\sum_{i=1}^K (p_i - \frac{1}{K}\sum_{j=1}^K p_j)^2}. \tag{D.8}$$

Therefore, when event (D.1) holds for any $i \in [K]$,

$$\begin{aligned}|b - \hat{b}| &\le \frac{1}{\sum_{i=1}^K \left(p_i - \frac{1}{K}\sum_{j=1}^K p_j\right)^2} \cdot \left|\sum_{i=1}^K \left(p_i - \frac{1}{K}\sum_{j=1}^K p_j\right) \cdot \left((d_i - \frac{1}{K}\sum_{j=1}^K d_j) - (\hat{d}_i - \frac{1}{K}\sum_{j=1}^K \hat{d}_j)\right)\right| \\ &= \frac{1}{\sum_{i=1}^K \left(p_i - \frac{1}{K}\sum_{j=1}^K p_j\right)^2} \cdot \left|\sum_{i=1}^K \left(p_i - \frac{1}{K}\sum_{j=1}^K p_j\right) \cdot (d_i - \hat{d}_i)\right| \\ &\le \frac{1}{\sum_{i=1}^K \left(p_i - \frac{1}{K}\sum_{j=1}^K p_j\right)^2} \sqrt{\sum_{i=1}^K \left(p_i - \frac{1}{K}\sum_{j=1}^K p_j\right)^2} \sqrt{\sum_{i=1}^K \left(d_i - \hat{d}_i\right)^2} \\ &\le \frac{1}{\sqrt{\sum_{i=1}^K \left(p_i - \frac{1}{K}\sum_{j=1}^K p_j\right)^2}} \sqrt{K\left(\frac{3\alpha}{\underline{g}}\right)^2} \\ &= \frac{1}{\sqrt{\frac{1}{K}\sum_{i=1}^K \left(p_i - \frac{1}{K}\sum_{j=1}^K p_j\right)^2}} \frac{3\alpha}{\underline{g}},\end{aligned}$$

where the first inequality follows from identities (D.7), (D.8) and the definition of $\hat{b}$, the second inequality follows from Cauchy-Schwarz inequality, the third inequality follows from (D.6).

Finally, by union bound, for any $\alpha > 0$, event (D.1) holds for each $i \in [K]$ with probability at least $1 - \sum_{i=1}^{K} 2e^{-2N_i\alpha^2}$. Since $N_i \geq N$ for each $i \in [K]$, by rearranging the terms, we get

$$\mathbb{P}\left[|b - \hat{b}| \leq \alpha\right] \geq 1 - 2Ke^{-\frac{\frac{1}{K}\sum_{i=1}^{K}(p_i - \frac{1}{K}\sum_{j=1}^{K}p_j)^2\underline{g}^2}{9}N\alpha^2}$$

for any $\alpha > 0$. $\square$

D.2. Proof of Lemma 2

Proof. To prove Lemma 2, we first need to establish the following two technical lemmas.

Lemma D.1 provides a probability bound on the difference between the empirical revenue $\hat{R}(p)$ and the true revenue $R(p)$ when p is in the range $[0, \max\{\min\{\hat{\tilde{p}}, \tilde{p}, \overline{p}\}, 0\}]$. Lemma D.2 establishes a probability bound on the difference between γ and $\hat{\gamma}$, referred to as the *true uncensored degree* and *empirical uncensored degree* respectively. The proofs of Lemma D.1 and Lemma D.2 can be found in Appendices D.4 and D.5 respectively.

Lemma D.1 **(Convergence of empirical revenue)**. *For any $\alpha > 0$, with probability at least $1 - 2Ke^{-\frac{N\alpha^2}{(y+\hat{b}\overline{p})^2\overline{p}^2}}$, $\left|\hat{R}(p) - R(p)\right| \leq 3\overline{p}^2|\hat{b} - b| + \alpha$ for all $0 \leq p \leq \max\{\min\{\hat{\tilde{p}}, \tilde{p}, \overline{p}\}, 0\}$.*

Lemma D.2 **(Convergence of empirical uncensored degree)**. *For any $\alpha > 0$, with probability at least $1 - 2Ke^{-2N\alpha^2}$, $|\hat{\gamma} - \gamma| \leq \alpha$.*

Now we start to prove Lemma 2. Recall that $\hat{i}^* = \arg\max_{i\in[K]}\{y_i + \hat{b}p_i\}$. For notation convenience, let $i^* := \arg\max_{i\in[K]}\{y_i + bp_i\}$. For any $p \in [\underline{p}, \overline{p}] \subset (0, \infty)$, from Proposition 2,

$$R_{\max}(p) = \begin{cases} R(p), & \text{if } p < \tilde{p}, \\ R_1(p), & \text{if } p \geq \tilde{p}; \end{cases} \quad \text{and} \quad R_{\min}(p) = \begin{cases} R(p), & \text{if } p < \tilde{p}, \\ R_2(p), & \text{if } p \geq \tilde{p}. \end{cases}$$

In addition, from the construction of D2ACD, for $p \in [\underline{p}, \overline{p}] \subset (0, \infty)$,

$$\hat{R}_{\max}(p) = \begin{cases} \hat{R}(p), & \text{if } p < \hat{\tilde{p}}, \\ \hat{R}_1(p), & \text{if } p \geq \hat{\tilde{p}}; \end{cases} \quad \text{and} \quad \hat{R}_{\min}(p) = \begin{cases} \hat{R}(p), & \text{if } p < \hat{\tilde{p}}, \\ \hat{R}_2(p), & \text{if } p \geq \hat{\tilde{p}}. \end{cases}$$

Our subsequent discussion is based on the assumption that the following four events hold:

(i) $\left|\hat{R}(p) - R(p)\right| \leq \alpha_0$ for all $0 \leq p \leq \max\{\min\{\hat{\tilde{p}}, \tilde{p}, \overline{p}\}, 0\}$,

(ii) $|\hat{b} - b| \leq \alpha_1$,

(iii) $|\hat{\gamma} - \gamma| \leq \alpha_2$,

(iv) $\left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}} \min\{\xi_{\hat{i}^*}^j, \lambda\} - \mathbb{E}[\min\{\xi, \lambda\}]\right| \leq \alpha_3$,

Without loss of generality, we can assume that $[\min\{\tilde{p}, \hat{\tilde{p}}\}, \max\{\tilde{p}, \hat{\tilde{p}}\}] \subseteq [\underline{p}, \overline{p}]$. Otherwise, we just focus on the interaction of the two intervals $[\min\{\tilde{p}, \hat{\tilde{p}}\}, \max\{\tilde{p}, \hat{\tilde{p}}\}]$ and $[\underline{p}, \overline{p}]$. For any $p \in [\underline{p}, \overline{p}]$, we discuss the following three cases: $p < \min\{\tilde{p}, \hat{\tilde{p}}\}$, $p > \max\{\tilde{p}, \hat{\tilde{p}}\}$ and $p \in [\min\{\tilde{p}, \hat{\tilde{p}}\}, \max\{\tilde{p}, \hat{\tilde{p}}\}]$.

Case 1: $p < \min\{\tilde{p}, \hat{\tilde{p}}\}$. In this case, we have

$$\hat{R}_{\max}(p) - R_{\max}(p) = \hat{R}(p) - R(p), \quad \text{and} \quad \hat{R}_{\min}(p) - R_{\min}(p) = \hat{R}(p) - R(p).$$

Since event (i) holds, then we have

$$\left|\hat{R}_{\max}(p) - R_{\max}(p)\right| \leq \alpha_0 \text{ and } \left|\hat{R}_{\min}(p) - R_{\min}(p)\right| \leq \alpha_0.$$

Case 2: $p > \max\{\tilde{p}, \hat{\tilde{p}}\}$**.** In this case, we have,

$$\hat{R}_{\max}(p) - R_{\max}(p) = \hat{R}_1(p) - R_1(p), \quad \text{and} \quad \hat{R}_{\min}(p) - R_{\min}(p) = \hat{R}_2(p) - R_2(p).$$

Since

$$R_1(p) = p \cdot \big(-\gamma b p + \mathbb{E}[\min\{\xi,\, \lambda\}] + (1-\gamma)(y-\lambda)\big), \quad \text{and} \quad R_2(p) = p \cdot \big(-bp + \mathbb{E}[\min\{\xi,\, \lambda\}]\big),$$

and it follows from $S^j_{\hat{i}^*} + \hat{b}p_{\hat{i}^*} \le y_{\hat{i}^*} + \hat{b}p_{\hat{i}^*} = \hat{\lambda}$ and the construction of $\hat{R}_1(\cdot)$ and $\hat{R}_2(\cdot)$ that

$$\hat{R}_1(p) = p \cdot \Big(-\hat{\gamma}\hat{b}p + \frac{1}{N_{\hat{i}^*}} \sum_{j=1}^{N_{\hat{i}^*}} \min\{S^j_{\hat{i}^*} + \hat{b}p_{\hat{i}^*}, \hat{\lambda}\} + (1-\hat{\gamma})(y - \hat{\lambda})\Big),$$

$$\hat{R}_2(p) = p \cdot \Big(-\hat{b}p + \frac{1}{N_{\hat{i}^*}} \sum_{j=1}^{N_{\hat{i}^*}} \min\{S^j_{\hat{i}^*} + \hat{b}p_{\hat{i}^*}, \hat{\lambda}\}\Big),$$

we have

$$\begin{aligned}
&|\hat{R}_1(p) - R_1(p)| \\
&= \Big| p\Big((\gamma b - \hat{\gamma}\hat{b})p + \big(\frac{1}{N_{\hat{i}^*}} \sum_{j=1}^{N_{\hat{i}^*}} \min\{S^j_{\hat{i}^*} + \hat{b}p_{\hat{i}^*}, \hat{\lambda}\} - \mathbb{E}[\min\{\xi,\, \lambda\}]\big) + \big((1-\hat{\gamma})(y-\hat{\lambda}) - (1-\gamma)(y-\lambda)\big)\Big)\Big| \\
&\le \overline{p}\Big(\overline{p}|\hat{\gamma}\hat{b} - \gamma b| + \Big| \frac{1}{N_{\hat{i}^*}} \sum_{j=1}^{N_{\hat{i}^*}} \min\{S^j_{\hat{i}^*} + \hat{b}p_{\hat{i}^*}, \hat{\lambda}\} - \mathbb{E}[\min\{\xi,\, \lambda\}]\Big| + |\hat{\lambda} - \lambda| + y|\hat{\gamma} - \gamma| + |\hat{\gamma}\hat{\lambda} - \gamma\lambda|\Big),
\end{aligned}$$

and

$$\begin{aligned}
|\hat{R}_2(p) - R_2(p)| &= \Big| p\Big((b - \hat{b})p + \big(\frac{1}{N_{\hat{i}^*}} \sum_{j=1}^{N_{\hat{i}^*}} \min\{S^j_{\hat{i}^*} + \hat{b}p_{\hat{i}^*}, \hat{\lambda}\} - \mathbb{E}[\min\{\xi,\, \lambda\}]\big)\Big)\Big| \\
&\le \overline{p}\Big(\overline{p}|\hat{b} - b| + \Big| \frac{1}{N_{\hat{i}^*}} \sum_{j=1}^{N_{\hat{i}^*}} \min\{S^j_{\hat{i}^*} + \hat{b}p_{\hat{i}^*}, \hat{\lambda}\} - \mathbb{E}[\min\{\xi,\, \lambda\}]\Big|\Big).
\end{aligned}$$

To further bound $|\hat{R}_1(p) - R_1(p)|$ and $|\hat{R}_2(p) - R_2(p)|$, we establish several inequalities as follows:

$$|\hat{\gamma}\hat{b} - \gamma b| \le |\hat{\gamma}\hat{b} - \hat{\gamma}b| + |\hat{\gamma}b - \gamma b| \le |\hat{b} - b| + \overline{b}|\hat{\gamma} - \gamma|, \tag{D.9}$$

$$|\hat{\lambda} - \lambda| \le \overline{p}|\hat{b} - b|, \tag{D.10}$$

$$|\hat{\gamma}\hat{\lambda} - \gamma\lambda| \le |\hat{\gamma}\hat{\lambda} - \hat{\gamma}\lambda| + |\hat{\gamma}\lambda - \gamma\lambda| \le |\hat{\lambda} - \lambda| + \overline{\lambda}|\hat{\gamma} - \gamma|, \tag{D.11}$$

where inequality (D.10) holds since

$$\begin{cases} \hat{\lambda} \ge y_{i^*} + \hat{b}p_{i^*} = y_{i^*} + bp_{i^*} - (b - \hat{b})p_{i^*} \ge \lambda - |\hat{b} - b|\overline{p}, \\ \hat{\lambda} = y_{\hat{i}^*} + bp_{\hat{i}^*} + (\hat{b} - b)p_{\hat{i}^*} \le \lambda + |\hat{b} - b|\overline{p}. \end{cases}$$

In addition, we also have

$$\begin{aligned}
&\Big| \frac{1}{N_{\hat{i}^*}} \sum_{j=1}^{N_{\hat{i}^*}} \min\{S^j_{\hat{i}^*} + \hat{b}p_{\hat{i}^*}, \hat{\lambda}\} - \mathbb{E}[\min\{\xi, \lambda\}]\Big| \\
&= \Big| \frac{1}{N_{\hat{i}^*}} \sum_{j=1}^{N_{\hat{i}^*}} \min\{\xi^j_{\hat{i}^*} + (\hat{b} - b)p_{\hat{i}^*}, \hat{\lambda}\} - \mathbb{E}[\min\{\xi, \lambda\}]\Big|
\end{aligned}$$

$$
\begin{aligned}
&\leq \Big|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi_{\hat{i}^*}^j+(\hat{b}-b)p_{\hat{i}^*},\hat{\lambda}\}-\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi_{\hat{i}^*}^j,\hat{\lambda}\}\Big|+\Big|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi_{\hat{i}^*}^j,\hat{\lambda}\}-\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi_{\hat{i}^*}^j,\lambda\}\Big| \\
&\quad+\Big|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi_{\hat{i}^*}^j,\lambda\}-\mathbb{E}[\min\{\xi,\lambda\}]\Big| \\
&\leq \overline{p}\cdot|\hat{b}-b|+|\hat{\lambda}-\lambda|+\Big|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi_{\hat{i}^*}^j,\lambda\}-\mathbb{E}[\min\{\xi,\lambda\}]\Big|. \qquad \text{(D.12)}
\end{aligned}
$$

Combining inequalities (D.9)-(D.12) and events (i)-(iv), we have for any $p>\max\{\tilde{p},\hat{\tilde{p}}\}$,

$$
\begin{aligned}
\Big|\hat{R}_{\max}(p)-R_{\max}(p)\Big| &= |\hat{R}_1(p)-R_1(p)| \\
&\leq 5\overline{p}^2|\hat{b}-b|+\overline{p}(y+\overline{\lambda}+\overline{p}\overline{b})|\hat{\gamma}-\gamma|+\overline{p}\Big|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi_{\hat{i}^*}^j,\lambda\}-\mathbb{E}[\min\{\xi,\lambda\}]\Big| \\
&\leq 5\overline{p}^2\alpha_1+\overline{p}(y+\overline{\lambda}+\overline{p}\overline{b})\alpha_2+\overline{p}\alpha_3, \qquad \text{(D.13)}
\end{aligned}
$$

and

$$
\Big|\hat{R}_{\min}(p)-R_{\min}(p)\Big| = |\hat{R}_2(p)-R_2(p)|\leq 3\overline{p}^2\alpha_1+\overline{p}\alpha_3. \qquad \text{(D.14)}
$$

Case 3: $p\in[\min\{\tilde{p},\hat{\tilde{p}}\},\max\{\tilde{p},\hat{\tilde{p}}\}]$**.** In this case, we first note that

$$
|\hat{\tilde{p}}-\tilde{p}| = \Big|\frac{(\lambda-y)(\hat{b}-b)+b(\lambda-\hat{\lambda})}{\hat{b}b}\Big|\leq\frac{(\overline{\lambda}+y)|\hat{b}-b|+\overline{b}|\hat{\lambda}-\lambda|}{\underline{b}^2}\leq\frac{\overline{\lambda}+y+\overline{b}\overline{p}}{\underline{b}^2}\alpha_1, \qquad \text{(D.15)}
$$

where the last inequality follows from event (ii) and inequality (D.10).

Suppose $\tilde{p}\leq\hat{\tilde{p}}$. Since $\tilde{p}\leq p\leq\hat{\tilde{p}}$, we have $\hat{R}_{\max}(p)=\hat{R}(p)$, $R_{\max}(p)=R_1(p)$, and thus,

$$
\begin{aligned}
|\hat{R}_{\max}(p)-R_{\max}(p)| = |\hat{R}(p)-R_1(p)| &\leq |\hat{R}(p)-\hat{R}(\hat{\tilde{p}})|+|\hat{R}(\hat{\tilde{p}})-R_1(\hat{\tilde{p}})|+|R_1(\hat{\tilde{p}})-R_1(p)| \\
&= |\hat{R}(p)-\hat{R}(\hat{\tilde{p}})|+|\hat{R}_1(\hat{\tilde{p}})-R_1(\hat{\tilde{p}})|+|R_1(\hat{\tilde{p}})-R_1(p)|,
\end{aligned}
$$

where the second identity follows from $\hat{R}(\hat{\tilde{p}})=\hat{R}_1(\hat{\tilde{p}})$. Since $\hat{R}(\cdot)$ and $R_1(\cdot)$ are both Lipschitz continuous in $[\underline{p},\overline{p}]$ with Lipschitz constant $y+\overline{b}\overline{p}$ and $y+2\overline{b}\overline{p}+2\overline{\lambda}$, respectively, with $\overline{\lambda}$ defined as $\max_{i\in[K]}\{y_i+\overline{b}p_i\}$, and $|p-\hat{\tilde{p}}|\leq|\hat{\tilde{p}}-\tilde{p}|$, we further have

$$
\begin{aligned}
|\hat{R}_{\max}(p)-R_{\max}(p)| &\leq (2y+3\overline{b}\overline{p}+2\overline{\lambda})|\hat{\tilde{p}}-\tilde{p}|+5\overline{p}^2\alpha_1+\overline{p}(y+\overline{\lambda}+\overline{p}\overline{b})\alpha_2+\overline{p}\alpha_3 \\
&\leq \Big(\frac{(2y+3\overline{b}\overline{p}+2\overline{\lambda})^2}{\overline{b}^2}+5\overline{p}^2\Big)\alpha_1+\overline{p}(y+\overline{\lambda}+\overline{p}\overline{b})\alpha_2+\overline{p}\alpha_3,
\end{aligned}
$$

where the second inequality follows from (D.15). Similarly, we have

$$
\begin{aligned}
|\hat{R}_{\min}(p)-R_{\min}(p)| = |\hat{R}(p)-R_2(p)| &\leq |\hat{R}(p)-\hat{R}(\hat{\tilde{p}})|+|\hat{R}_2(\hat{\tilde{p}})-R_2(\hat{\tilde{p}})|+|R_2(\hat{\tilde{p}})-R_2(p)| \\
&\leq (y+3\overline{b}\overline{p}+\overline{\lambda})|\hat{\tilde{p}}-\tilde{p}|+3\overline{p}^2\alpha_1+\overline{p}\alpha_3 \\
&\leq \Big(\frac{(y+3\overline{b}\overline{p}+\overline{\lambda})^2}{\overline{b}^2}+3\overline{p}^2\Big)\alpha_1+\overline{p}\alpha_3,
\end{aligned}
$$

where the second inequality holds since $\hat{R}(\cdot)$ and $R_2(\cdot)$ are Lipschitz continuous in $[\underline{p},\overline{p}]$ with Lipschitz constant $y+\overline{b}\overline{p}$ and $2\overline{b}\overline{p}+\overline{\lambda}$ respectively, and the last inequality follows from (D.15).

Suppose $\tilde{p} > \hat{\tilde{p}}$. For any $\hat{\tilde{p}} \le p \le \tilde{p}$, we have $\hat{R}_{\max}(p) = \hat{R}_1(p)$, $R_{\max}(p) = R(p)$, and thus,

$$\begin{aligned}
|\hat{R}_{\max}(p) - R_{\max}(p)| &= |\hat{R}_1(p) - R(p)| \le |\hat{R}_1(p) - \hat{R}_1(\hat{\tilde{p}})| + |\hat{R}_1(\hat{\tilde{p}}) - R(\hat{\tilde{p}})| + |R(\hat{\tilde{p}}) - R(p)| \\
&= |\hat{R}_1(p) - \hat{R}_1(\hat{\tilde{p}})| + |\hat{R}(\hat{\tilde{p}}) - R(\hat{\tilde{p}})| + |R(\hat{\tilde{p}}) - R(p)| \\
&\le \alpha_0 + (2y + 3\bar{b}\overline{p} + \overline{\lambda})|\hat{\tilde{p}} - \tilde{p}| \\
&\le \alpha_0 + \frac{(2y + 3\bar{b}\overline{p} + \overline{\lambda})^2}{\underline{b}^2}\alpha_1,
\end{aligned}$$

where the second inequality follows from event (i), and Lipschitz continuity of $\hat{R}_1(\cdot)$ and $R(\cdot)$ in $[\underline{p}, \overline{p}]$. Similarly, for any $\hat{\tilde{p}} \le p \le \tilde{p}$, we have $\hat{R}_{\min}(p) = \hat{R}_2(p)$, $R_{\min}(p) = R(p)$, and thus,

$$\begin{aligned}
|\hat{R}_{\min}(p) - R_{\min}(p)| &\le |\hat{R}_2(p) - \hat{R}_2(\hat{\tilde{p}})| + |\hat{R}(\hat{\tilde{p}}) - R(\hat{\tilde{p}})| + |R(\hat{\tilde{p}}) - R(p)| \\
&\le \alpha_0 + (y + 3\bar{b}\overline{p} + \overline{\lambda})|\hat{\tilde{p}} - \tilde{p}| \\
&\le \alpha_0 + \frac{(y + 3\bar{b}\overline{p} + \overline{\lambda})^2}{\underline{b}^2}\alpha_1,
\end{aligned}$$

where the second inequality follows from event (i) and Lipschitz continuity of $\hat{R}_2(\cdot)$ and $R(\cdot)$ in $[\underline{p}, \overline{p}]$, and the last inequality follows from (D.15).

Combining all the above three cases, conditioned on the four events (i)-(iv), we have

$$\begin{aligned}
|\hat{R}_{\max}(p) - R_{\max}(p)| &\le \alpha_0 + \left(5\overline{p}^2 + \frac{(2y + 3\bar{b}\overline{p} + 2\overline{\lambda})^2}{\underline{b}^2}\right)\alpha_1 + \overline{p}(y + \overline{\lambda} + \overline{p}\bar{b})\alpha_2 + \overline{p}\alpha_3, \\
|\hat{R}_{\min}(p) - R_{\min}(p)| &\le \alpha_0 + \left(5\overline{p}^2 + \frac{(2y + 3\bar{b}\overline{p} + 2\overline{\lambda})^2}{\underline{b}^2}\right)\alpha_1 + \overline{p}(y + \overline{\lambda} + \overline{p}\bar{b})\alpha_2 + \overline{p}\alpha_3.
\end{aligned}$$

Recall that $C_1 = 5\overline{p}^2 + \frac{(2y+3\bar{b}\overline{p}+2\overline{\lambda})^2}{\underline{b}^2}$ and $C_2 = \overline{p}(y + \overline{\lambda} + \overline{p}\bar{b})$. Then for any $\alpha > 0$, by letting $\alpha_0 = \alpha/6 + 3\overline{p}^2|\hat{b} - b|$, $\alpha_1 = \alpha/(2C_1)$, $\alpha_2 = \alpha/(6C_2)$ and $\alpha_3 = \alpha/(6\overline{p})$, we have

$$\begin{aligned}
&\mathbb{P}\left[\left|\hat{R}_{\max}(p) - R_{\max}(p)\right| \le \alpha, \left|\hat{R}_{\min}(p) - R_{\min}(p)\right| \le \alpha, \forall p \in [\underline{p}, \overline{p}]\right] \\
&\ge 1 - \mathbb{P}\left[\exists p \in [0, \max\{\min\{\hat{\tilde{p}}, \tilde{p}, \overline{p}\}\}], \left|\hat{R}(p) - R(p)\right| > \frac{\alpha}{6} + 3\overline{p}^2|\hat{b} - b|\right] - \mathbb{P}\left[|\hat{b} - b| > \frac{\alpha}{2C_1}\right] - \mathbb{P}\left[|\hat{\gamma} - \gamma| > \frac{\alpha}{6C_2}\right] \\
&\quad - \mathbb{P}\left[\left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}} \min\{\xi_{\hat{i}^*}^j, \lambda\} - \mathbb{E}[\min\{\xi, \lambda\}]\right| > \frac{\alpha}{6\overline{p}}\right] \\
&\ge 1 - 2Ke^{-\frac{1}{36(y+\bar{b}\overline{p})^2\overline{p}^2}N\alpha^2} - 2Ke^{-\frac{\frac{1}{K}\sum_{i=1}^K (p_i - \frac{1}{K}\sum_{j=1}^K p_j)^2 \underline{g}^2}{36C_1^2}N\alpha^2} - 2Ke^{-\frac{1}{18C_2^2}N\alpha^2} - 2Ke^{-\frac{1}{18\overline{p}^2\overline{\lambda}^2}N\alpha^2},
\end{aligned}$$

where the first inequality follows from the union bound, and the second inequality follows from Lemma 1, Lemma D.1, Lemma D.2, and the Hoeffding inequality in Lemma A.1. □

D.3. Proof of Proposition 7

Proof. Similar to the proofs of Proposition 5 and Proposition 6, for notation convenience, we define $p_{\max}^\dagger = \max\arg\max_{p\in\mathbb{R}} R_{\max}(p)$, $p_{\min}^\dagger = \max\arg\max_{p\in\mathbb{R}} R_{\min}(p)$.

We first note that for any $p \in [\underline{p}, \overline{p}]$ such that $p \le \tilde{p}$,

$$\begin{aligned}
\max_{F\in\mathcal{F}} \left\{R_F(p_F^*) - R_F(p)\right\} &= \max_{F\in\mathcal{F}}\{R_F(p_F^*)\} - R_{\max}(p) = R_{\max}(p_{\max}^*) - R_{\max}(p) \\
&= \max\left\{R_{\max}(p_{\max}^*) - R_{\max}(p),\, R_{\min}(p_{\min}^*) - R_{\min}(p)\right\},
\end{aligned}$$

where the first identity holds since when $p \le \tilde{p}$, $R_F(p) = R_{\max}(p)$ for any $F \in \mathcal{F}$, and the last identity holds since $R_{\min}(p^*_{\min}) \le R_{\max}(p^*_{\max})$, and when $p \le \tilde{p}$, $R_{\max}(p) = R_{\min}(p)$. Thus, identity (6) holds for this case.

Then it suffices to prove (6) for $p \in [\max\{\tilde{p}, \underline{p}\}, \overline{p}]$. Note that here we can assume $\tilde{p} < \overline{p}$ such that the interval $[\max\{\tilde{p}, \underline{p}\}, \overline{p}]$ is nonempty, since otherwise, any p in the set $[\underline{p}, \overline{p}]$ satisfies $p \le \tilde{p}$, and the result is already proved from the above discussion. We next consider the following two cases: $p^*_{\max} \ne p^*_{\min}$ and $p^*_{\max} = p^*_{\min}$.

Case 1: $p^*_{\max} \ne p^*_{\min}$. If $p^*_{\max} \ne p^*_{\min}$, it can be verified from the proof of Proposition 3 that for any $F \in \mathcal{F}$, $\tilde{p} \le p^\dagger_F$. Since $R_F(p)$ is concave in p, then $R_F(p) \le R_F(\tilde{p})$ for all $p < \tilde{p}$. Hence, we have

$$\max_{F\in\mathcal{F}} \left\{R_F(p^*_F) - R_F(p)\right\} = \max_{F\in\mathcal{F}} \max_{p'\in[\underline{p},\overline{p}]} \left\{R_F(p') - R_F(p)\right\} = \max_{F\in\mathcal{F}} \max_{p'\in[\max\{\underline{p},\tilde{p}\},\overline{p}]} \left\{R_F(p') - R_F(p)\right\}. \tag{D.16}$$

In what follows, we define $l(p) \triangleq \max_{F\in\mathcal{F}} \max_{p'\in[\max\{\underline{p},\tilde{p}\},\overline{p}]} \{R_F(p') - R_F(p)\}$. We next divide our proof into two steps.

Step 1. In this step, we prove the following result: for any $p \in [\max\{\underline{p}, \tilde{p}\}, \overline{p}]$,

$$l(p) = \max\left\{ \max_{p'\in[\max\{\underline{p},\tilde{p}\},\overline{p}],p'\ge p} \left\{R_{\max}(p') - R_{\max}(p)\right\}, \max_{p'\in[\max\{\underline{p},\tilde{p}\},\overline{p}],p'\le p} \left\{R_{\min}(p') - R_{\min}(p)\right\}\right\}, \tag{D.17}$$

To show (D.17), it suffices to prove the following two identities: for any $p_1, p_2 \in [\max\{\underline{p}, \tilde{p}\}, \overline{p}]$ with $p_1 \le p_2$,

$$\max_{F\in\mathcal{F}} \left\{R_F(p_2) - R_F(p_1)\right\} = R_{\max}(p_2) - R_{\max}(p_1), \tag{D.18}$$

$$\max_{F\in\mathcal{F}} \left\{R_F(p_1) - R_F(p_2)\right\} = R_{\min}(p_1) - R_{\min}(p_2). \tag{D.19}$$

To see (D.18), after replacing $p^\dagger_{\max}$ by p_1 and p by p_2 in inequality (C.4) in the proof of Proposition 3, we can prove that: for any $F \in \mathcal{F}$, and $\max\{\underline{p}, \tilde{p}\} \le p_1 \le p_2 \le \overline{p}$,

$$R_F(p_2) - R_F(p_1) \le R_1(p_2) - R_1(p_1) = R_{\max}(p_2) - R_{\max}(p_1).$$

Moreover, a distribution $\overline{F} \in \mathcal{F}$ defined as: $\overline{F}(x) = F_\xi(x)$ for any $x < \lambda$, $\overline{F}(x) = \gamma$ for all $\lambda \le x < y + bp_2$, and $\overline{F}(x) = F_\xi(x)$ for all $x \ge y + bp_2$, achieves the upper bound $R_{\max}(p_2) - R_{\max}(p_1)$ and therefore maximizes $R_F(p_2) - R_F(p_1)$ over all $F \in \mathcal{F}$. Thus, identity (D.18) holds. Similar arguments can be applied to prove identity (D.19) by using inequality (C.3) in the proof of Proposition 3, and thus we omit the details.

Step 2. In this step, we use (D.17) from Step 1 to prove for any $p \in [\max\{\underline{p}, \tilde{p}\}, \overline{p}]$,

$$l(p) = \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\right\}. \tag{D.20}$$

We first observe when $p^*_{\max} \ne p^*_{\min}$, we have $\tilde{p} \le p^\dagger_{\min} \le p^*_{\min} < p^*_{\max} \le p^\dagger_{\max}$. Then we consider three cases: $p \le p^*_{\min}$, $p > p^*_{\max}$ and $p^*_{\min} \le p \le p^*_{\max}$.

If $p < p^*_{\min}$, it can be verified that $p^*_{\min} = p^\dagger_{\min}$. Otherwise, suppose $p^*_{\min} > p^\dagger_{\min}$, since $R_{\min}(\cdot)$ is concave, we must have $p^*_{\min} = \underline{p} > p^\dagger_{\min}$, leading to contradiction with $\underline{p} \le p < p^*_{\min}$. In this case, $R_{\min}(\cdot)$ increases in $[\max\{\tilde{p}, \underline{p}\}, p]$. Therefore, $\max_{p'\in[\max\{\underline{p},\tilde{p}\},\overline{p}],p'\le p}\{R_{\min}(p') - R_{\min}(p)\} = 0$, and

$$\begin{aligned} l(p) &= \max_{p'\in[\max\{\underline{p},\tilde{p}\},\overline{p}],p'\ge p} \left\{R_{\max}(p') - R_{\max}(p)\right\} = R_{\max}(p^*_{\max}) - R_{\max}(p) \\ &= \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\right\}, \end{aligned} \tag{D.21}$$

where the second identity holds since $p^*_{\max} \in [\max\{\underline{p}, \tilde{p}\}, \overline{p}]$ and $p^*_{\max} \geq p$, and the last identity holds due to the following inequality:

$$R_{\max}(p^*_{\max}) - R_{\max}(p) \geq R_{\max}(p^*_{\min}) - R_{\max}(p) \geq R_{\min}(p^*_{\min}) - R_{\min}(p),$$

where the second inequality holds since $R_1(p) - R_2(p) = b(1-\gamma)p(p-\tilde{p})$ increases in $p \geq \max\{\tilde{p}, 0\}$, $R_{\max}(p) = R_1(p)$, $R_{\min}(p) = R_2(p)$ for $p \geq \max\{\tilde{p}, 0\}$, and p, $p^*_{\min} \geq \max\{\tilde{p}, 0\}$.

If $p > p^*_{\max}$, we can similarly prove that $p^*_{\max} = p^\dagger_{\max}$, $R_{\max}(\cdot)$ decreases in (p, ∞), $\max_{p' \in [\max\{\underline{p}, \tilde{p}\}, \overline{p}], p' \geq p} \{R_{\max}(p') - R_{\max}(p)\} = 0$ and

$$l(p) = R_{\min}(p^*_{\min}) - R_{\min}(p) = \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\right\}. \tag{D.22}$$

If $p^*_{\min} \leq p \leq p^*_{\max}$, it can be easily verified that $\max_{p' \in [\max\{\underline{p}, \tilde{p}\}, \overline{p}], p' \geq p} R_{\max}(p') = R_{\max}(p^*_{\max})$ and $\max_{p' \in [\max\{\underline{p}, \tilde{p}\}, \overline{p}], p' \leq p} R_{\min}(p') = R_{\min}(p^*_{\min})$. Thus, we obtain identity (D.20) from identity (D.17).

Case 2: $p^*_{\max} = p^*_{\min}$. If $p^*_{\max} = p^*_{\min}$, we will prove the result by considering two cases: $p^\dagger_{F_\xi} < \tilde{p}$ and $p^\dagger_{F_\xi} \geq \tilde{p}$. Recall that $F_\xi(\cdot)$ is the true c.d.f. for the distribution of ξ.

Subcase 1: $p^\dagger_{F_\xi} < \tilde{p}$. When $p^\dagger_{F_\xi} < \tilde{p}$, it is easy to see from the definitions of F and $R_F(\cdot)$ that $R_F(\cdot)$ decreases in $(\tilde{p}, \infty)$, and $p^\dagger_F = p^\dagger_{F_\xi} < \tilde{p}$ for all $F \in \mathcal{F}$. Under this circumstance, we prove the result by considering two cases: $\underline{p} < \tilde{p}$ and $\underline{p} \geq \tilde{p}$.

If $\underline{p} < \tilde{p}$, then we have $p^*_F = p^*_{\max} = p^*_{\min} = \max\{p^\dagger_{F_\xi}, \underline{p}\}$, and $R_F(p^*_F) = R_{\max}(p^*_{\max}) = R_{\min}(p^*_{\min})$. Then,

$$\begin{aligned}
\max_{F \in \mathcal{F}} \left\{R_F(p^*_F) - R_F(p)\right\} &= R_{\min}(p^*_{\min}) + \max_{F \in \mathcal{F}: p^\dagger_F < \tilde{p}} \{-R_F(p)\} \\
&= R_{\min}(p^*_{\min}) - R_{\min}(p) \\
&= \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\right\},
\end{aligned}$$

where the last identity holds since $R_{\max}(p^*_{\max}) - R_{\max}(p) = R_{\min}(p^*_{\min}) - R_{\max}(p) \leq R_{\min}(p^*_{\min}) - R_{\min}(p)$.

If $\underline{p} \geq \tilde{p}$, we have $p^*_F = p^*_{\max} = p^*_{\min} = \underline{p}$ for any $F \in \mathcal{F}$. In this case, $\tilde{p} \leq \underline{p} = p^*_F \leq p \leq \overline{p}$, and we have

$$\begin{aligned}
\max_{F \in \mathcal{F}} \left\{R_F(p^*_F) - R_F(p)\right\} &= R_{\min}(\underline{p}) - R_{\min}(p) \\
&= \max\{R_{\max}(\underline{p}) - R_{\max}(p), R_{\min}(\underline{p}) - R_{\min}(p)\} \\
&= \max\{R_{\max}(p^*_{\max}) - R_{\max}(p), R_{\min}(p^*_{\min}) - R_{\min}(p)\},
\end{aligned}$$

where the first identity follows from (D.19), the second identity holds due to a similar reason to the last identity in (D.21).

Subcase 2: $p^\dagger_{F_\xi} \geq \tilde{p}$. If $p^\dagger_{F_\xi} \geq \tilde{p}$, it is easily verified that $p^\dagger_F \geq \tilde{p}$ for any $F \in \mathcal{F}$. Moreover, if $p^*_{F_0} < \tilde{p}$ for some $F_0 \in \mathcal{F}$, we must have $p^*_{\min} = p^*_{\max} = p^*_F = \overline{p}$ and $R_{\max}(p^*_{\max}) = R_{\min}(p^*_{\min}) = R_F(p^*_F)$ for all $F \in \mathcal{F}$, which then implies

$$\max_{F \in \mathcal{F}} \left\{R_F(p^*_F) - R_F(p)\right\} = R^*_{\min}(p^*_{\min}) - R_{\min}(p) = \max\{R_{\max}(p^*_{\max}) - R_{\max}(p), R_{\min}(p^*_{\min}) - R_{\min}(p)\}.$$

Thus, the result holds. Then we only need to focus on $p^*_F \geq \tilde{p}$ for all $F \in \mathcal{F}$. Under this circumstance, we prove the result by considering two cases: $\tilde{p} \leq p^*_{\min} \leq p$, and $\tilde{p} \leq p < p^*_{\min}$.

When $\tilde{p} \le p^*_{\min} \le p$, we have

$$\begin{aligned}\max_{F\in\mathcal{F}}\{R_F(p^*_F) - R_F(p)\} &= \max_{F\in\mathcal{F}}\{R_F(p^*_{\min}) - R_F(p)\} = R_{\min}(p^*_{\min}) - R_{\min}(p)\\ &= \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\right\},\end{aligned}$$

where the first identity holds since $p^*_{\max} = p^*_{\min}$ implies $p^*_F = p^*_{\min}$ for any $F \in \mathcal{F}$, the second identity follows from similar arguments to (D.19), and the last identity holds due to a similar reason to the last identity in (D.21).

When $\tilde{p} \le p < p^*_{\min}$, we can similarly prove

$$\begin{aligned}\max_{F\in\mathcal{F}}\{R_F(p^*_F) - R_F(p)\} &= \max_{F\in\mathcal{F}}\{R_F(p^*_{\max}) - R_F(p)\} = R_{\max}(p^*_{\max}) - R_{\max}(p)\\ &= \max\left\{R_{\max}(p^*_{\max}) - R_{\max}(p),\, R_{\min}(p^*_{\min}) - R_{\min}(p)\right\}.\end{aligned}$$

Combining Case 1 and Case 2, the proof of Proposition 7 is complete. ☐

D.4. Proof of Lemma D.1

Proof. If $\min\{\hat{\tilde{p}}, \tilde{p}, \overline{p}\} \le 0$, the range $[0, \max\{\min\{\hat{\tilde{p}}, \tilde{p}, \overline{p}\}, 0\}]$ becomes $\{0\}$, and the result is trivial since $|\hat{R}(0) - R(0)| = 0$. We next assume $\min\{\hat{\tilde{p}}, \tilde{p}, \overline{p}\} > 0$.

For any $0 \le p \le \min\{\hat{\tilde{p}}, \tilde{p}, \overline{p}\}$, we define a random variable $X_p = \xi + (\hat{b} - b)p_{\hat{i}^*} - \hat{b}p$, where ξ follows the distribution of $F_\xi(\cdot)$ and is independent of $\hat{b}$. Then

$$\begin{aligned}&\left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{S^j_{\hat{i}^*} + \hat{b}(p_{\hat{i}^*} - p),\, y\} - \mathbb{E}_\xi\left[\min\{\xi - bp,\, y\}\right]\right|\\ &\le \left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{S^j_{\hat{i}^*} + \hat{b}(p_{\hat{i}^*} - p),\, y\} - \mathbb{E}_\xi\left[\min\{X_p,\, y\}\right]\right| + \left|\mathbb{E}_\xi\left[\min\{X_p,\, y\}\right] - \mathbb{E}_\xi\left[\min\{\xi - bp,\, y\}\right]\right|, \quad \text{(D.23)}\end{aligned}$$

where the subscript ξ is added to demonstrate that the expectation is taken with respect to the r.v. ξ with c.d.f. $F_\xi(\cdot)$.

We now bound the first term in (D.23). Note that when $0 \le p \le \min\{\hat{\tilde{p}}, \tilde{p}, \overline{p}\}$,

$$\begin{aligned}&\left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{S^j_{\hat{i}^*} + \hat{b}(p_{\hat{i}^*} - p),\, y\} - \mathbb{E}_\xi\left[\min\{X_p,\, y\}\right]\right|\\ &= \left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi^j_{\hat{i}^*} - bp_{\hat{i}^*},\, y - \hat{b}(p_{\hat{i}^*} - p)\} - \mathbb{E}_\xi\left[\min\{\xi - bp_i,\, y - \hat{b}(p_i - p)\}\right]\right|\\ &\le \left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi^j_{\hat{i}^*} - bp_{\hat{i}^*},\, y - \hat{b}(p_{\hat{i}^*} - p)\} - \frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi^j_{\hat{i}^*} - bp_{\hat{i}^*},\, y - b(p_{\hat{i}^*} - p)\}\right|\\ &\quad + \left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi^j_{\hat{i}^*} - bp_{\hat{i}^*},\, y - b(p_{\hat{i}^*} - p)\} - \mathbb{E}_\xi\left[\min\{\xi - bp_{\hat{i}^*},\, y - b(p_{\hat{i}^*} - p)\}\right]\right|\\ &\quad + \left|\mathbb{E}_\xi\left[\min\{\xi - bp_{\hat{i}^*},\, y - b(p_{\hat{i}^*} - p)\}\right] - \mathbb{E}_\xi\left[\min\{\xi - bp_{\hat{i}^*},\, y - \hat{b}(p_{\hat{i}^*} - p)\}\right]\right|\\ &\le |\hat{b} - b|\cdot|p_{\hat{i}^*} - p| + \left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi^j_{\hat{i}^*} - bp,\, y\} - \mathbb{E}_\xi\left[\min\{\xi - bp,\, y\}\right]\right| + |\hat{b} - b|\cdot|p_{\hat{i}^*} - p|\\ &\le 2\overline{p}|\hat{b} - b| + \left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi^j_{\hat{i}^*},\, y + bp\} - \mathbb{E}_\xi\left[\min\{\xi,\, y + bp\}\right]\right|, \quad \text{(D.24)}\end{aligned}$$

where the first identity holds since when $p \le \hat{\tilde{p}}$, $\min\{S^j_{\hat{i}^*} + \hat{b}(p_{\hat{i}^*} - p), y\} = \min\{\xi^j_{\hat{i}^*} - bp_{\hat{i}^*}, y_{\hat{i}^*}, y - \hat{b}(p_{\hat{i}^*} - p)\} + \hat{b}(p_{\hat{i}^*} - p) = \min\{\xi^j_{\hat{i}^*} - bp_{\hat{i}^*}, y - \hat{b}(p_{\hat{i}^*} - p)\} + \hat{b}(p_{\hat{i}^*} - p)$, and $\min\{X_p, y\} = \min\{\xi - bp_{\hat{i}^*}, y - \hat{b}(p_{\hat{i}^*} - p)\} + \hat{b}(p_{\hat{i}^*} - p)$, the second inequality holds since $|\min\{x,y\} - \min\{x,z\}| \le |y - z|$ for any $x, y, z \in \mathbb{R}$, and the last inequality follows from $p_i \le \overline{p}$ and $p \ge 0$. To further bound the second term in the right-hand side of (D.24), we apply Lemma A.3 by letting $f(x,p) = \min\{x, y + bp\}$ for fixed y and b, $\mathcal{X} = [0, y + b\overline{p}]$ and $\mathcal{Y} = [0, \overline{p}]$. It is easy to verify that $f(\cdot)$ is lower bounded by zero, upper bounded by $y + b\overline{p}$, and monotone in x for any given p. Thus, from the union bound, Lemma A.3 and $N_i \ge N$ for each $i \in [K]$, we have, for any given $\alpha > 0$,

$$\begin{aligned}
&\mathbb{P}\left[\sup_{p\in[0,\overline{p}]}\left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi^j_{\hat{i}^*},\, y+bp\} - \mathbb{E}_\xi\big[\min\{\xi,\, y+bp\}\big]\right| \ge \alpha\right] \qquad \text{(D.25)}\\
&= \mathbb{P}\left[\sup_{p\in[0,\overline{p}]}\left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}} f(\xi^j_{\hat{i}^*}, p) - \mathbb{E}_\xi[f(\xi,p)]\right| \ge \alpha\right]\\
&\le \sum_{i=1}^{K}\mathbb{P}\left[\sup_{p\in[0,\overline{p}]}\left|\frac{1}{N_i}\sum_{j=1}^{N_i} f(\xi^j_i, p) - \mathbb{E}_\xi[f(\xi,p)]\right| \ge \alpha\right]\\
&\le 2Ke^{-\frac{2N\alpha^2}{(y+b\overline{p})^2}}. \qquad \text{(D.26)}
\end{aligned}$$

We next bound the second term in inequality (D.23). Note that $X_p = \xi - bp + (\hat{b} - b)(p_{\hat{i}^*} - p)$, then, for $0 \le p \le \min\{\hat{\tilde{p}}, \tilde{p}, \overline{p}\}$, we have

$$\begin{aligned}
&\left|\mathbb{E}_\xi\big[\min\{X_p,\, y\}\big] - \mathbb{E}_\xi\big[\min\{\xi - bp,\, y\}\big]\right|\\
&= \left|\mathbb{E}_\xi\big[\min\{\xi - bp + (\hat{b} - b)(p_{\hat{i}^*} - p),\, y\}\big] - \mathbb{E}_\xi\big[\min\{\xi - bp,\, y\}\big]\right|\\
&\le |\hat{b} - b| \cdot |p_{\hat{i}^*} - p|\\
&\le \overline{p}|\hat{b} - b|. \qquad \text{(D.27)}
\end{aligned}$$

Therefore, combining inequalities (D.23), (D.24), (D.25) and (D.27), we conclude that

$$\begin{aligned}
&\mathbb{P}\left[\sup_{0\le p\le\max\{\min\{\hat{\tilde{p}},\tilde{p},\overline{p}\},0\}}\left|\hat{R}(p) - R(p)\right| \le 3\overline{p}^2|\hat{b} - b| + \alpha\right]\\
&\ge \mathbb{P}\left[\sup_{0\le p\le\max\{\min\{\hat{\tilde{p}},\tilde{p},\overline{p}\},0\}}\left|\frac{1}{N_{\hat{i}^*}}\sum_{j=1}^{N_{\hat{i}^*}}\min\{\xi^j_{\hat{i}^*},\, y+bp\} - \mathbb{E}_\xi\big[\min\{\xi,\, y+bp\}\big]\right| \le \frac{\alpha}{\overline{p}}\right]\\
&\ge 1 - 2Ke^{-\frac{N\alpha^2}{(y+b\overline{p})^2\overline{p}^2}},
\end{aligned}$$

where the first inequality follows from inequalities (D.23), (D.24) and (D.27), and the second inequality follows from inequality (D.25). □

D.5. Proof of Lemma D.2

Proof. By definition, $\hat{\gamma} = \max_{i\in[K]}\hat{\gamma}_i = \max_{i\in[K]}\frac{1}{N_i}\sum_{j=1}^{N_i}\mathbb{1}_{\{S^j_i < y_i\}} = \max_{i\in[K]}\frac{1}{N_i}\sum_{j=1}^{N_i}\mathbb{1}_{\{\xi^j_i < y_i + bp_i\}}$, and $\gamma = \max_{i\in[K]}\gamma_i = \max_{i\in[K]}\mathbb{P}[\xi < y_i + bp_i] = \max_{i\in[K]}\mathbb{E}[\mathbb{1}_{\{\xi < y_i + bp_i\}}]$.

For each $i \in [K]$, $\mathbb{1}_{\{\xi^j_i < y_i + bp_i\}}$ is an *i.i.d.* sample of the random variable $\mathbb{1}_{\{\xi < y_i + bp_i\}}$. Therefore, by Lemma A.1, we have $|\hat{\gamma}_i - \gamma_i| \le \alpha$, with probability at least $1 - 2e^{-2N_i\alpha^2}$, for each $i \in [K]$. By applying the union bound and using the fact that $N \le N_i$ for all $i \in [K]$, we have, with probability at least $1 - 2Ke^{-2N\alpha^2}$, that $|\hat{\gamma}_i - \gamma_i| \le \alpha$ holds for all $i \in [K]$. Given $|\hat{\gamma}_i - \gamma_i| \le \alpha$ for each $i \in [K]$, we have $\hat{\gamma} = \max_{i\in[K]} \hat{\gamma}_i \le \max_{i\in[K]} \gamma_i + \alpha = \gamma + \alpha$, and $\hat{\gamma} = \max_{i\in[K]} \hat{\gamma}_i \ge \max_{i\in[K]} \gamma_i - \alpha = \gamma - \alpha$.

Hence, we have $\mathbb{P}[|\hat{\gamma} - \gamma| \le \alpha] \ge 1 - 2Ke^{-2N\alpha^2}$. □

Appendix E: Proof of Proposition 8 in Section 5

Proof. Similar to the proofs of Theorems 2 and 3, let $\mathcal{B}(\alpha)$ be the following event for any $\alpha > 0$:

$$\mathcal{B}(\alpha) = \left\{ |\hat{R}_{\max}(p) - R_{\max}(p)| \le \alpha, |\hat{R}_{\min}(p) - R_{\min}(p)| \le \alpha, \forall p \in [\underline{p}, \overline{p}] \right\}.$$

Note that when $\hat{p}^*_{\max} = \hat{p}^*_{\min}$, we have $p^{\mathsf{D2ACD}} = \hat{p}^*_{\max} = \hat{p}^*_{\min}$, and from the second inequality of (10) in the proof of Theorems 2 and 3, we have

$$\mathbb{P}\left[\hat{p}^*_{\max} = \hat{p}^*_{\min}, \mathcal{B}(\alpha)\right] \le \mathbb{P}\left[\max_{F\in\mathcal{F}} \{R_F(p^*_F) - R_F(p^{\mathcal{A}})\} \le 2\alpha\right].$$

Thus, since when $p^*_{\max} \ne p^*_{\min}$ and $\Delta \ge \alpha_0$, we must have $\mathbb{P}\left[\hat{p}^*_{\max} = \hat{p}^*_{\min}, \mathcal{B}(\alpha)\right] = 0$ if $\alpha < \frac{1}{2}\alpha_0$. Let $\alpha = \frac{1}{3}\alpha_0$ and denote $\overline{\mathcal{B}}(\alpha)$ as the opposite event of $\mathcal{B}(\alpha)$, then we have

$$\begin{aligned}\mathbb{P}\left[\hat{p}^*_{\max} = \hat{p}^*_{\min}\right] &= \mathbb{P}\left[\hat{p}^*_{\max} = \hat{p}^*_{\min}, \mathcal{B}(\alpha)\right] + \mathbb{P}\left[\hat{p}^*_{\max} = \hat{p}^*_{\min}, \overline{\mathcal{B}}(\alpha)\right] = \mathbb{P}\left[\hat{p}^*_{\max} = \hat{p}^*_{\min}, \overline{\mathcal{B}}(\alpha)\right] \le \mathbb{P}\left[\overline{\mathcal{B}}(\alpha)\right] \\ &\le 2Ke^{-\frac{1}{324(y+\overline{b}\overline{p})^2\overline{p}^2}N\alpha_0^2} + 2Ke^{-\frac{\frac{1}{K}\sum_{i=1}^K (p_i - \frac{1}{K}\sum_{j=1}^K p_j)^2 \underline{g}^2}{324C_1^2}N\alpha_0^2} + 2Ke^{-\frac{1}{162C_2^2}N\alpha_0^2} + 2Ke^{-\frac{1}{162\overline{p}^2\overline{\lambda}^2}N\alpha_0^2},\end{aligned}$$

where the second inequality follows from Lemma 2. □

Appendix F: Description of KM Subroutine

In this appendix, we describe the KM subroutine used in $\mathsf{D2ACD\text{-}KM}$. See Algorithm 2. Specifically, the

---

Algorithm 2: The `KM` subroutine

Input: $\left\{(p_i, S_i^j, y_i) : i \in [K], j \in [N_i]\right\}$, $N_{[K]} = \sum_{i=1}^K N_i$, $\hat{b}$, $\hat{\lambda}$, $\hat{\gamma}$, p, y

Sort samples $\{(S_i^j + \hat{b}p_i, \mathbb{1}_{\{S_i^j < y_i\}}) : i \in [K], j \in [N_i]\}$ from smallest to largest according to $S_i^j + \hat{b}p_i$; ties are broken by putting censored observations with $\mathbb{1}_{\{S_i^j < y_i\}} = 0$ after uncensored observations with $\mathbb{1}_{\{S_i^j < y_i\}} = 1$; let the ordered observations be denoted as $\{(W_1, \delta_1), (W_2, \delta_2), \ldots, (W_{N_{[K]}}, \delta_{N_{[K]}})\}$;

Let $\hat{F}^{\mathrm{KM}}(x)$ be the empirical c.d.f. for ξ: $\hat{F}^{\mathrm{KM}}(x) = 1 - \prod_{1\le i\le N_{[K]}: W_i \le x} \left(\frac{N_{[K]}-i}{N_{[K]}-i+1}\right)^{\delta_i}$;

$\hat{R}(p) := p \cdot \mathbb{E}_{\xi\sim\hat{F}^{\mathrm{KM}}}[\min\{\xi - \hat{b}p, y\}]$;

$\hat{R}_1(p) := p \cdot \Big(-\hat{\gamma}\hat{b}p + \mathbb{E}_{\xi\sim\hat{F}^{\mathrm{KM}}}[\xi] + (1-\hat{\gamma})(y - \hat{\lambda})\Big)$;

$\hat{R}_2(p) := p \cdot \Big(-\hat{b}p + \mathbb{E}_{\xi\sim\hat{F}^{\mathrm{KM}}}[\xi]\Big)$;

$\hat{\hat{p}} := (\hat{\lambda} - y)/\hat{b}$;

$\hat{R}_{\max}(p) := \mathbb{1}_{\{p<\hat{\hat{p}}\}}\hat{R}(p) + \mathbb{1}_{\{p\ge\hat{\hat{p}}\}}\hat{R}_1(p)$;

$\hat{R}_{\min}(p) := \mathbb{1}_{\{p<\hat{\hat{p}}\}}\hat{R}(p) + \mathbb{1}_{\{p\ge\hat{\hat{p}}\}}\hat{R}_2(p)$;

Output: $\Big(\hat{R}_{\max}(p), \hat{R}_{\min}(p)\Big)$

---

inputs of the KM subroutine include the whole offline datset, the number of total historical samples $N_{[K]}$, the estimated price sensitivity $\hat{b}$ from Step 1 of Algorithm 1, the empirical quantities $\hat{\lambda}$ and $\hat{\gamma}$ defined in line 12 of Algorithm 1, any given price p, and the inventory level y in the data-driven problem (1). The empirical c.d.f. $\hat{F}^{KM}(\cdot)$ in line 3 of Algorithm 2 is defined through the product-limit formula that is typically used to compute the KM estimator. Note that the original product-limit formula uses the indicator function $\mathbb{1}_{\{D_i^j \leq y_i\}}$, which, however, may not be available in practice. Thus, we approximate this quantity through $\mathbb{1}_{\{S_i^j < y_i\}}$, which can always be computed based on the available sales and inventory information. The output of the KM subroutine is the estimated optimistic and pessimistic revenues at price p, which then serves as the input of D2ACD-KM's Step 3 for generating the suggested price.

Appendix G: Further Discussion on LR-ExcludeCensored

In this appendix, we further explain why linear regression on uncensored data, i.e., `LR-ExcludeCensored`, fails to generate a consistent estimate for the price sensitivity. We present a simple numerical example. Suppose the demand model is given by

$$D(p) = 100 - p + \eta, \quad \eta \sim \text{Uniform}(-10, 10).$$

We assume that the offline dataset is given by $\{(p_i, y, S_i) : i \in [N]\}$, where historical prices $p_1, p_2, \ldots, p_N$ are independently drawn from a uniform distribution in $[15, 40]$, and the sales quantity S_i given each price p_i is independently drawn from the above demand model after truncated by the constant censoring inventory level $y = 75$. `LR-ExcludeCensored` computes the estimates of base demand and price sensitivity by solving the following linear regression:

$$(\hat{a}, \hat{b}) = \arg \min_{a' \in \mathbb{R}, b' \in \mathbb{R}} \sum_{1 \leq i \leq N : S_i < y} \left(S_i - (a' - b' p_i)\right)^2 . \tag{G.1}$$

Figure 8 plots the result of LR-ExcludeCensored under a randomly generated dataset in the green line. To see how much bias `LR-ExcludeCensored` may introduce, we also plot the true demand function in red, and the fitted line of `LR-WithoutCensoring` in blue, which is obtained from linear regression on the full dataset without information loss, i.e., $\{(p_i, D_i) : i \in [N]\}$.

Undoubtedly, `LR-WithoutCensoring` approximates the true demand function very well, since there is no censoring issue in the data. By contrast, `LR-ExcludeCensored` leads to biased estimate for the price sensitivity and base demand. This is because if dropping censored data, the selected samples for the linear regression are restricted to the data points below the inventory level, i.e., those in orange, leading to the issue of *endogeneity*. In fact, a sample (p_i, S_i) is selected for the linear regression if and only if $100 - p_i + \eta_i < y$ (or equivalently, $\eta_i < p_i + y - 100$). Thus, the underlying distribution of the selected noise samples essentially follows $\text{Uniform}(-10, \min\{p + y - 100, 10\})$, which is correlated with the explanatory variable p in the regression model. Due to such endogeneity, the estimate of the regression coefficient in the above linear regression (G.1) is biased.

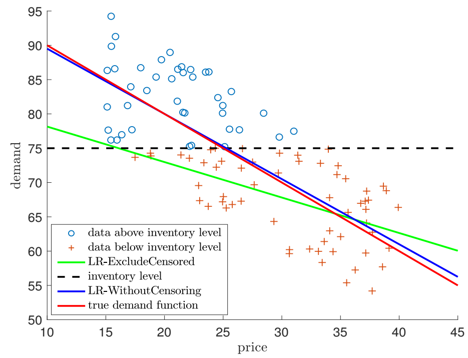

Figure 8 **Biased estimation from** LR-ExcludeCensored**.**