A prescriptive machine learning framework to the price-setting newsvendor problem

Pavithra Harsha
IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, pharsha@us.ibm.com

Ramesh Natarajan
IBM T. J. Watson Research Center (past affiliation when this research was done)
Amazon, 515 Westlake Avenue N, Seattle, WA 98109 , ramesn@amazon.com

Dharmashankar Subramanian
IBM T. J. Watson Research Center, Yorktown Heights, NY 10598, dharmash@us.ibm.com

We develop a practical framework for modeling the price-setting newsvendor problem, which includes statistical estimation and price optimization methods for estimating the optimal solutions and associated confidence intervals. We present a novel and exact reformulation of the problem that leads to the framework which requires as input the estimates of only three distinct aspects of the demand distribution: the mean, quantile and superquantile (in contrast to the full-demand distribution), and it provides asymptotically optimal solutions under mild conditions, if these estimates are consistent. To estimate these quantities in a data-driven, distribution-free fashion potentially with multi-dimensional observational datasets, we investigate statistical estimators based on generalized linear regression (GLR), mixed-quantile regression (MQR), and superquantile regression (SQR). We propose a novel and exact large-scale decomposition method that is computationally efficient for SQR, and extend the MQR estimation method by relaxing its implicit assumptions of homoskedasticity (these two extensions are of independent interest). Our computational experiments, first, indicate the importance of flexible statistical estimation methods that inherently account for heteroskedasticity, and second, suggest that quantile-based methods such as MQR and SQR provide better solutions for a wide range of demand distributions, although for certain location-scale demand distributions similar to the Normal distribution, GLR may be preferable.

Key words: Pricing, newsvendor, statistics: estimation, decomposition algorithm, heteroskedastic least squares, quantile regression, conditional value-at-risk (CVaR), superquantile regression

1. Introduction

The classical price-setting newsvendor problem occupies a central and prominent role in pricing and inventory decision theory. In its simplest form, a firm must simultaneously and jointly determine the optimal price and optimal order quantity for a product with a *known* stochastic price-dependent demand, in order to maximize the expected profit during a single inventory period. The newsvendor profit function accounts for underage and overage costs which are respectively associated with scenarios when the demand is higher or lower than the order quantity, as well as the revenue that is obtained from the observed demand.

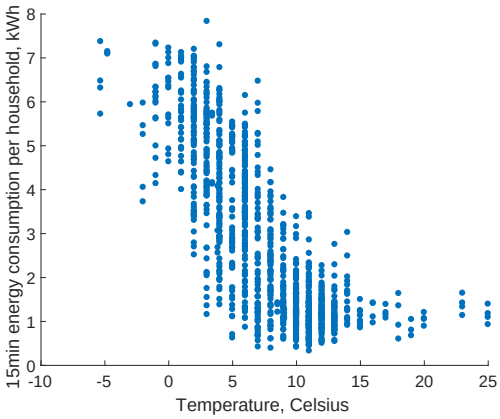

Figure 1 **Average energy consumption in kWh per household recorded every 15 mins between 7-9am from April 1, 2006 to March 31, 2007, when prices are in the [40,60] cents range, in a dynamic pricing experiment in the Pacific Northwest GridWise Testbed Demonstration Project (Hammerstrom et al. 2007).**

In many of these existing and emerging application areas, however, the relevant demand function is a random variable, whose probability distribution and dependence on price is not explicitly known, and must be modeled and estimated from historical sales data. Furthermore, many other drivers besides price must be included in the demand response model for statistical accuracy, and conditional heteroskedasticity effects in the demand distribution must also be taken into account. The objective of this paper is to address this issue using data-driven, distribution-free, machine learning methods to characterize the stochastic demand response, while the primary novelty is that the statistical modeling methods are carried out in close conjunction with the requirements of the optimization problem, here, the classical price-setting newsvendor.

One motivation for our work is the potential application of the price-setting newsvendor model and related schemes to the emerging electricity smart grid for demand response planning. Here, the electric utility may simultaneously decide on both the scheduled generation and certain demand-shaping price incentives, so as to minimize the impact on the expected operational profits. In this application, other factors besides price, such as weather and time-of-day effects, will influence the demand response and demand variability, and must be taken into consideration in the statistical modeling. For example, consider the energy consumption data in Fig. 1 recorded during the morning peak in a dynamic price experiment at the Olympic Peninsula. Notice the significant impact of temperature on demand and the presence of temperature-dependent heteroskedasticity. The consumption patterns can also display significant daily and weekly dependencies (not shown). One approach here is to estimate a separate demand model from historical data for each combination of the levels of the relevant external factors; however, this is clearly impractical when the number of such combinations is very large. Besides, many potential combinations are rarely observed in

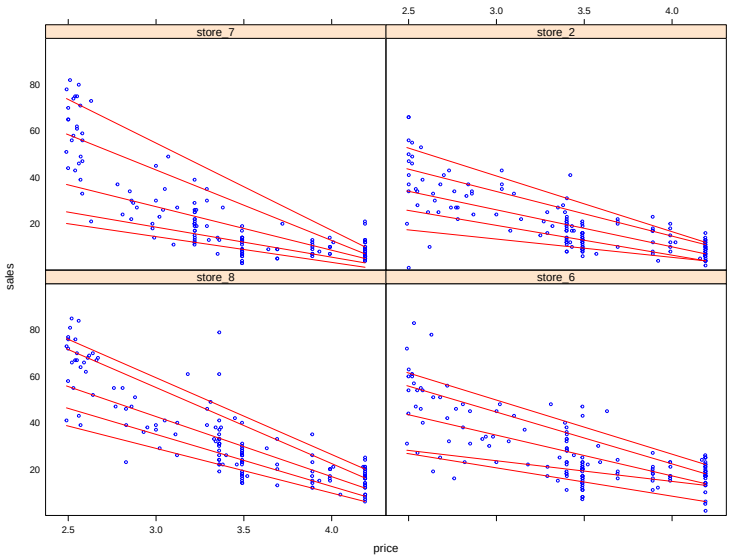

Figure 2 **Retail sales data for a non-seasonal product as a function of price across several stores (data anonymized to protect confidentiality). The red lines are the quantiles at levels 0.1,0.25,0.5,0.75 and 0.9.**

the historical data. Therefore, it is preferable to capture these high-dimensional relationships using flexible, multivariate regression methods (e.g., see, Hastie et al. 2001), to obtain robust demand response models for operational planning.

The same issues arise for applications of the price-setting newsvendor model in the retail, manufacturing and services sectors. For example, the data in Fig. 2 shows a price-dependent heteroskedasticity in individual stores sales of a product, as well as distributional variations across the individual store locations (perhaps due to their distinct store characteristics and shopper demographics). We also observe a monotonic price-dependent variance function with the lowest demand variability occurring at the largest price. A non-monotonic price-dependent variance function is also quite likely to occur in practical applications. For example, the lowest demand variability can come in the middle range of prices where one has a good understanding of the market (Raz and Porteus 2006). The methods described in this paper are relevant in all these settings.

1.1. Contributions

This paper focuses on the classical price-setting newsvendor problem and provides results that address the two common variants of this problem: the *lost sales* setting where the excess demand is entirely lost and the *emergency order* setting where the excess demand is also met but at a high cost that is exogenous to the price. Our contributions are as follows:

1. In this paper we develop a **practical prescriptive machine learning framework for the price-setting newsvendor problem**, which includes statistical estimation and price optimization methods for estimating the optimal solutions and associated confidence intervals. The

framework utilizes a novel and exact reformulation of the price-setting newsvendor problem that we present in the paper and requires as input the estimates of only three distinct aspects of the demand distribution, namely the mean, quantile and superquantile (also known as conditional value-at-risk, CVaR). In contrast to the current state of the art, the proposed framework does not require the complete price-dependent demand distribution for optimization and hence tailored data-driven machine learning methods can be utilized in the proposed framework in a plug and play fashion. Under some mild conditions, we show that if these data-driven estimates are consistent then the proposed prescriptive framework is asymptotically optimal.

2. We **investigate three different distribution-free machine learning methods**, which are based on generalized linear regression (GLR), mixed-quantile regression (MQR), and superquantile regression (SQR). The choice of these methods is motivated by their ability to obtain consistent estimators for the mean, quantile and superquantile of the corresponding unknown demand distribution. Any combination of these techniques, or other equivalent techniques, may be used to estimate the three distinct quantities of interest, leading to a profusion of ways for implementing the desired optimization computations. For each method, we investigate the inherent tradeoffs between the flexibility of modeling, the scalability to large data sets and the ability to capture conditional heteroskedastic effects[1].
3. We provide **extensions to current CVaR estimation methods**. We develop an efficient and exact large-scale decomposition method to solve large instances of the SQR which currently does not scale beyond a few hundred sample points. The proposed method is a novel cutting plane algorithm that is shown to be empirically far more tractable than the original SQR formulation. We also extend the MQR method by relaxing its implicit assumption of a homoskedastic data set. These extensions have wider applicability and are of independent interest besides the price-setting newsvendor application discussed in this paper (e.g., in financial applications).
4. We carry out **computational analysis and comparisons and derive insights** across different regression methods for stochastic demand models with different functional and noise characteristics as applied to the price-setting newsvendor problem by varying the sample size and the number of covariates. The regression methods that model heteroskedasticity implicitly or explicitly are observed to be more effective in the presence of heteroskedasticity, over methods that ignore it, and are competitive in the absence of heteroskedasticity. Our computational experiments also suggest that quantile based methods such as MQR and SQR provide better solutions for a wider range of demand models, except for the case when the noise terms have a location-scale form that is similar to the *Normal* distribution (e.g., symmetric, unimodal, and

[1] Conditional heteroskedasticity occurs when any chosen subset of the covariates may affect the conditional (on covariates) variance of the corresponding linear regression error.

not heavy-tailed) when GLR methods may be preferred. Not surprisingly, regularization techniques within the framework are particularly useful for big-data settings for stable out-of-sample performance of the overall decision problem. Finally, we present results of the proposed methods for estimating the optimal solutions and associated confidence intervals (using bootstrapping) for a practical data environment using the energy consumption data (see Fig. 1).

1.2. Background and Relevant work

The evolving literature on the coordination of pricing and inventory decisions has been reviewed by Elmaghraby and Keskinocak (2003), Chan et al. (2004), Yano and Gilbert (2005), and more recently by Chen and Simchi-Levi (2010), while Özer and Phillips (2012) have discussed the importance of empirical and statistical studies of pricing.

A synthesis of the literature on the price-setting newsvendor problem related to the existence and uniqueness conditions for the lost sales formulation is provided in Petruzzi and Dada (1999). They cover the case when the stochastic price-demand relationship is specified in a certain form, e.g., the additive model with a linear demand function (Mills 1959), the multiplicative model with a iso-elastic demand function (Karlin and Carr 1962), and the mixed additive-multiplicative model (Young 1978). In all these cases, the mean demand is specified as a monotonic decreasing function of the price, and the variance is specified as a non-increasing function of price (further, typically, a constant variance is assumed in the additive case, and monotonic decreasing variance is assumed in the multiplicative and mixed additive-multiplicative case). Yao et al. (2006) have extended these results to a more general class of price-demand functions for the additive and multiplicative models. Kocabıyıkoğlu and Popescu (2011) provide further generalizations, in particular, including the case when the demand variance may be a non-monotonic function of the price for mixed additive-multiplicative models. Such a non-monotonic variance function is quite likely in practice (see references and discussion below).

Arikan and Jammernegg (2009) have reviewed a number of approaches for modeling the stochastic price-demand relationship in the literature, and they characterize the following two approaches as being distinctive. The first, due to Lau and Lau (1988), models the first four lower-order moments of the demand distribution as a function of price. The second, due to Raz and Porteus (2006), models the individual quantiles of the demand distribution as piecewise-linear functions of the price. In these two "distinctive" yet complementary approaches of specifying the price-dependent heteroskedasticity effects in the demand model, the model parameters are specified based on subjective assessment of experts, rather than directly estimated from historical sales data. For instance, in Lau and Lau (1988), the first four moments are subjectively matched to the corresponding moments of a four-parameter beta distribution. Raz and Porteus (2006) recommend obtaining the subjective

estimates for the demand variability at a few selected price points, which are then interpolated and extended throughout the price range of interest using piecewise-linear functions.

A fully distribution-free statistical estimation approach, by contrast, is not constrained by the need to obtain the subjective estimates in some convenient manner. So, for example, multiple demand drivers can be directly incorporated and estimated in the demand model, whereas the subjective assessment of these multivariate effects would be difficult at best. While the subjective approach is not the focus of this paper, it is nevertheless a useful alternative approach, particularly when there is no historical data available (e.g., for new products that have little or no sales history).

For the *standard newsvendor* problem, the only decision variable is the order quantity (and price is not a decision variable). Beutel and Minner (2012) describe an approach where the demand model can comprise of multiple drivers that include the effects of price, price changes and weather. They observe that the inclusion of these additional drivers in the demand modeling substantially improves the accuracy of the demand forecasts, leading to a corresponding reduction in the safety stock requirements. They describe two multiple-covariate dependent modeling approaches for estimating the optimal order quantity for the standard newsvendor problem: a linear regression method with a first-order heteroskedasticity correction; a linear programming formulation equivalent to quantile regression. Ban and Rudin (2018) additionally have proposed using regularization and kernel-based methods for high-dimensional problems. The extension of these approaches to the price-setting newsvendor problem is not straightforward, since this also requires the estimation of the CVAR.

Other related distribution-free perspectives on the standard newsvendor problem include, for example, the use of bootstrap confidence intervals for newsvendor quantile estimates (Bookbinder and Lordahl 1989), robust optimization (Scarf et al. 1958, Gallego and Moon 1993, Perakis and Roels 2008), non-parametric approaches in censored data environments (Godfrey and Powell 2001, Huh et al. 2011, Huh and Rusmevichientong 2009, Besbes and Muharremoglu 2013), operational statistics (Liyanage and Shanthikumar 2005, Chu et al. 2008) and sampling-based bounds (Levi et al. 2007, 2015, Ban and Rudin 2018). However, none of these papers, have considered price as a decision variable, and barring Ban and Rudin (2018), neither have they considered the case where multiple demand drivers are present.

Some of the other variants of the price-setting newsvendor problem in the literature include: (a) the coordination of pricing and inventory for an assortment of products where the demand of any item depends on the prices of all the items (Aydin and Porteus 2008); (b) use of an alternate objective such as a risk-averse profit objective as opposed to a traditional risk-neutral one (Agrawal and Seshadri 2000, Chen et al. 2009); (c) multi-period models with backordered inventory (e.g., Federgruen and Heching 1999). The focus in these papers is on the existence and uniqueness of the optimal decisions and related structural results (in similar spirit to Petruzzi and Dada 1999).

Our work is related to the data-driven pricing literature where prices are optimized using forecasted demand obtained from historical observations and its drivers, and inventory is not a decision variable (Caro and Gallien 2012, Ferreira et al. 2016, Harsha et al. 2019, Ettl et al. 2019). Recent variants include adaptive learning with (1) distribution-free perspectives (Bertsimas and Vayanos 2017) and (2) regret minimization (Qiang and Bayati 2016, Cohen et al. 2016).

Our framework requires a statistical estimated CVaR of the demand distribution, for which we have considered the following two recent approaches to CVaR or superquantile regression. The *mixed quantile regression* method proposed by Chun et al. (2012) estimates the CVaR using a linear programming formulation similar to quantile regression. Their method is however restricted to homoskedastic or constant variance distributions only. The *superquantile regression* method proposed by Rockafellar et al. (2014) also estimates CVaR using a linear programming (LP) formulation, which is however derived based on the risk quadrangle (see Rockafellar and Uryasev 2013). Although this method makes no assumptions about homoskedasticity, as indicated by the authors in their paper, it does not scale beyond a few hundred sample points because of size explosion in the LP (both in the number of constraints and variables).

2. Price-Setting Newsvendor Problem

Consider a single-product, profit-maximizing firm, which at the beginning of an inventory period has to set a unit product price $p \in \mathcal{P}$, where $\mathcal{P}$ is a closed and continuous set of feasible prices. Simultaneously, the firm also has to set the order quantity x for stocking the product at the unit procurement cost c. The stochastic price-dependent demand is assumed to be fully observed, and denoted by $D(p, \mathbf{z})$, where $\mathbf{z}$ denotes the external drivers that influence demand, as elucidated further below. Any unsold stock units at the end of the inventory period are redeemed at the unit salvage price s. Note that $p > c > s$ is required in order to have a meaningful and non-trivial newsvendor problem.

The external drivers $\mathbf{z}$ in the stochastic demand $D(p, \mathbf{z})$ may include the effects of time-of-day, day-of-week, season, weather, holidays, special events, promotional incentives, and so on. Furthermore, when dynamic time-series effects are considered, $\mathbf{z}$ may also include the lagged effects of demand, price and other relevant drivers. We assume that all the influences on demand, besides price p, which are included in $\mathbf{z}$ have their values precisely known at the beginning of the inventory period. Therefore, any demand drivers that are unknown, unmeasured or uncertain are not included in $\mathbf{z}$, and their effects may be considered to be part of the stochastic component of $D(p, \mathbf{z})$. Finally, in order to focus the discussion, we assume that the observational data for estimating the demand

function $D(p, \mathbf{z})$ is collected in an uncensored setting, where the price and other relevant drivers could be varied independently.[2]

We consider two common variants of the price-setting newsvendor problem that arise in different applications.

Lost Sales Formulation: For a product retailer, we are primarily concerned with this variant of the price-setting newsvendor problem. Here, if the observed demand $D(p, \mathbf{z})$ exceeds the order quantity x, then the resulting underage is associated with a unit cost $p - c + v$ where v may be the monetary equivalent of the loss of consumer goodwill that is incurred due to the out-of-stock situation. On the other hand, if the observed demand $D(p, \mathbf{z})$ is lower than the order quantity x, then the resulting overage is associated with a unit cost $c - s$. The product retailer aims to maximize the expected profit by jointly optimizing the two decision variables, viz., the unit price p and the order quantity x. As is well known, this optimization problem can be formulated as:

$$\Pi_{ls}(\mathbf{z}): \quad \max_{p \in \mathcal{P}, x} \quad (p-c)E[D(p,\mathbf{z})] - (p-c+v)E\left[D(p,\mathbf{z}) - x\right]^+ - (c-s)E\left[x - D(p,\mathbf{z})\right]^+. \quad (2.1)$$

Emergency Order Formulation: For the electricity provider, by contrast, we are primarily concerned with this variant of the formulation. Here, the order quantity x represents the pre-scheduled electricity generation. If the observed demand $D(p, \mathbf{z})$ exceeds x, then the resulting shortfall is immediately procured from the spot market or from spinning reserve, but at the premium fixed unit procurement cost $m > c$. The resulting underage unit cost $m - c$ is also often referred to as the unanticipated stock-replenishment costs. The unit salvage price $s < c$ is associated with the excess in the pre-scheduled generation; for example, s may represent the contracted per unit sell price with a bulk storage farm. The electricity provider also aims to maximize the expected profitability by jointly optimizing the two decision variables, viz., the unit price p and the order quantity x. Similarly, this optimization problem can be formulated as:

$$\Pi_{eo}(\mathbf{z}): \quad \max_{p \in \mathcal{P}, x} \quad (p-c)E[D(p,\mathbf{z})] - (m-c)E\left[D(p,\mathbf{z}) - x\right]^+ - (c-s)E\left[x - D(p,\mathbf{z})\right]^+. \quad (2.2)$$

We now describe a reformulation of (2.1) and (2.2) that is suitable for the implementation of the required optimization procedures.

2.1. Optimization Formulation

The properties of the newsvendor objective function have been widely studied (Zipkin 2000). For the lost sales (2.1) and emergency order (2.2) formulations, the objective function is concave in x for given p. Therefore, in both cases, there is a unique optimal solution for x given p, denoted $\mathrm{VaR}_\alpha[D(p, \mathbf{z})]$, and given by

[2] These assumptions are made to be able to focus on the key contributions of the paper. There is vast literature on statistical modeling techniques that can be adapted to relax these assumptions (for the censored setting, see e.g. Huh et al. (2011), and for the price-endogenous setting, see e.g. Phillips et al. (2012), Bertsimas and Kallus (2016)). These can be used in our framework to obtain consistent estimators of the quantities of interest.

$$\text{VaR}_\alpha[D(p,\mathbf{z})] = \inf\left\{x \geq 0 : F_{D(p,\mathbf{z})}(x) \geq \alpha\right\}, \tag{2.3}$$

where $F_{D(p,\mathbf{z})}(.)$ is the cumulative distribution function (c.d.f.) of the random variable $D(p,\mathbf{z})$. The *critical quantile* or the *newsvendor quantile* $\alpha \in [0,1]$ is denoted by α_{ls} for the lost sales formulation, and by α_{eo} for the emergency order formulation respectively, with

$$\alpha_{ls} = \frac{p-c+v}{p-s+v}, \quad \alpha_{eo} = \frac{m-c}{m-s}. \tag{2.4}$$

The quantity $\text{VaR}_\alpha[D(p,\mathbf{z})]$ in Eq. (2.3) is the α-level value-at-risk of $D(p,\mathbf{z})$, or equivalently, the α-level quantile function of $D(p,\mathbf{z})$.

In Eq. (2.4), the value α_{eo} depends only on the known problem parameters m, c and s. The value α_{ls}, however, depends on the decision variable p, in addition to the specified parameters c, s and v. While this distinction between α_{eo} and α_{ls} is important for the respective optimization procedures, as described later below, for notational brevity, and to emphasize the common aspects of the optimization formulation, we suppress the dependence of α_{ls} on p, and also omit the subscripts on α_{ls} and α_{eo} below (except where this distinction is explicitly required).

Then, substituting the conditional optimal value for x from Eq. (2.3) into either (2.1) or (2.2) results in the following reduced objective function which only involves the decision variable p:

$$\Pi_{ls}(\mathbf{z}) \text{ or } \Pi_{eo}(\mathbf{z}): \quad \max_{p \in P} \quad (p-s)E[D(p,\mathbf{z})] - (c-s)\text{CVaR}_\alpha[D(p,\mathbf{z})]. \tag{2.5}$$

Here $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ denotes the α-level conditional value-at-risk of $D(p,\mathbf{z})$, which is defined in Rockafellar and Uryasev (2000) as

$$\text{CVaR}_\alpha[D(p,\mathbf{z})] = \min_x \left[x + \frac{1}{(1-\alpha)} E\left[D(p,\mathbf{z}) - x\right]^+\right]. \tag{2.6}$$

For a continuous random variable $D(p,\mathbf{z})$, this is identical to the conditional expected value in the upper α tail, given by

$$\begin{aligned} \text{CVaR}_\alpha[D(p,\mathbf{z})] &= E[D(p,\mathbf{z}) | D(p,\mathbf{z}) \geq \text{VaR}_\alpha[D(p,\mathbf{z})]] \\ &\equiv \frac{1}{1-\alpha}\int_\alpha^1 \text{VaR}_\tau[D(p,\mathbf{z})]d\tau. \end{aligned} \tag{2.7}$$

For discrete or mixed discrete-continuous $D(p,\mathbf{z})$, there is an equivalent definition to Eq. (2.7) which is given later below.

The conditional value-at-risk $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ in Eq. (2.6) arises in diverse disciplines, although the terminology may vary depending on the interpretation of the random variable $D(p,\mathbf{z})$. For example, in the electricity distribution industry, if $D(p,\mathbf{z})$ denotes the stochastic electricity demand, and if α denotes the quantile level of $D(p,\mathbf{z})$ corresponding to the pre-scheduled electricity generation, then $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ is essentially equivalent to the well-known reliability metric *LOLE - Loss of Load in Expectation* (Billinton and Allan 1996, Harsha et al. 2013). In more recent literature, the conditional value-at-risk is also referred to as the *superquantile* (Rockafellar et al. 2014).

2.1.1. Discussion and Perspective The optimization formulation in Eq. (2.5) involves a specific linear combination of the two quantities $E[D(p,\mathbf{z})]$ and $\text{CVaR}_\alpha[D(p,\mathbf{z})]$. To our knowledge, the price-setting newsvendor problem has not previously been posed in this form, which is reminiscent of the mean-CVaR objective function used in certain risk optimization settings, e.g., Rockafellar and Uryasev (2000). A third quantity $\text{VaR}_\alpha[D(p,\mathbf{z})]$ is required to obtain the optimal order quantity in Eq. (2.3).

In summary, the specification of the three quantities $E[D(p,\mathbf{z})]$, $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ and $\text{VaR}_\alpha[D(p,\mathbf{z})]$ is sufficient to obtain the desired optimal solutions to the price-setting newsvendor problem from Eq. (2.3) and Eq. (2.5). If the stochastic demand function $D(p,\mathbf{z})$ is available in some standard, explicit form, these three quantities can be directly evaluated (e.g., using closed-form expressions available for many standard distributions; see, Andreev et al. 2005, Nadarajah et al. 2014). From the perspective of this paper, however, the optimization formulation in Eq. (2.3) and Eq. (2.5) suggests that it may be fruitful to directly estimate these quantities from the historical sales data, without the intermediate step of explicitly ascertaining $D(p,\mathbf{z})$. A variety of statistical estimation techniques can be used for this purpose, which are capable of flexibly modeling the respective functional dependencies on the demand drivers, with minimal assumptions on the form and distribution of $D(p,\mathbf{z})$, as described further below.

2.2. Optimization Approaches

We outline one possible approach for obtaining the optimal price and optimal order quantity for the optimization formulation described in Section 2.1.

Procedure 1 Optimization Procedure for Price-Setting Newsvendor Problem (Emergency Order)

Input: Given m, c, s and $\tilde{\mathbf{z}}$.

1: Estimate $E[D(p,\mathbf{z})]$ as a function of $(p,\mathbf{z})$.

2: Obtain α from Eq. (2.4).

3: Estimate $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ as a function of $(p,\mathbf{z})$.

4: Solve for the optimal price $p^*(\tilde{\mathbf{z}}) \in \mathcal{P}$ from Eq. (2.5) using the estimates $E[D(p,\mathbf{z})]$ and $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ obtained in steps 1 and 3 above at the given $\tilde{\mathbf{z}}$.

5: Estimate optimal order quantity $x^* = \text{VaR}_\alpha[D(p,\mathbf{z})]|_{p^*,\tilde{\mathbf{z}}}$ from Eq. (2.3).

Output: optimal price p^*, and optimal order quantity x^*.

Procedure 1 is given for the emergency order setting. The corresponding procedure for the lost sales setting implements Steps 2 and 3 as sub-routines of Steps 4 to obtain the optimal price estimate, p^*. Finally, the corresponding optimal order quantity estimate is obtained from Step 5. A pictorial representation of our distribution-free prescriptive machine learning framework is provided

in Fig. 3, which captures the high-level implementation details of the procedures for the emergency order and lost sales settings. We make some remarks about Procedure 1.

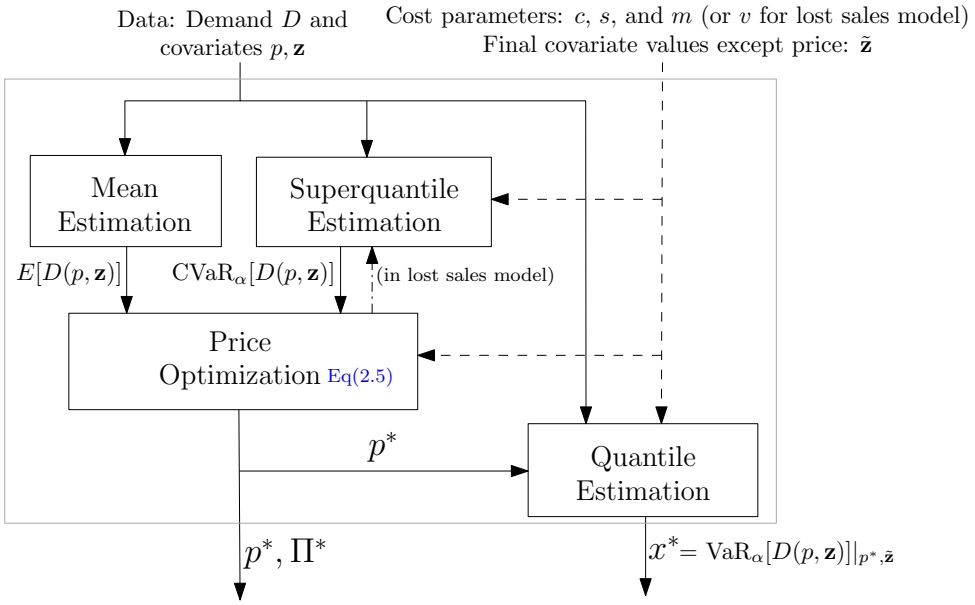

Figure 3 **A distribution-free prescriptive machine learning framework to the price-setting newsvendor problem.**

First, this procedure can not only be used when the explicit form of the stochastic demand function $D(p,\mathbf{z})$ is known, but also when the three quantities $E[D(p,\mathbf{z})]$, $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ and $\text{VaR}_\alpha[D(p,\mathbf{z})]$ are directly estimated from the data. In fact, under some mild conditions, if the data-driven estimates are consistent then the proposed method is also asymptotically optimal.

Claim 1. *We assume that the data-driven estimators* $E[D(p,\mathbf{z})]$, $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ *and* $\text{VaR}_\alpha[D(p,\mathbf{z})]$ *are (pointwise) consistent estimators that are also continuous in* $p \in P$ *We also assume that the true objective of* (2.5) *has a unique global maximum. Then the proposed data-driven method for the emergency order setting is asymptotically optimal (consistent).*

Additionally, if we assume that the estimator for $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ *uniformly converges in probability in* $\alpha \in [\alpha_{\min}, \alpha_{\max}]$ *where* $\alpha_{\min} = \frac{p_{\min}-c+v}{p_{\min}-s+v}$ *and* $\alpha_{\max} = \frac{p_{\max}-c+v}{p_{\min}-s+v}$ *and that the estimator is also jointly continuous in* (p,α) *then the proposed method for the lost-sales setting is also asymptotically optimal (consistent).*

Second, the objective of problem (2.5), which is a univariate function of price p, need not be concave in general.[3] Sufficient conditions for concavity of the emergency order setting are easy to deduce, for example $(p-c)E[D(p,\mathbf{z})]$ and $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ should be concave and convex in p

[3] Note that Claim 1 only requires a unique global optimal and therefore true even if there are multiple local maxima. The claim does not comment about the ease of solving the optimization problem (2.5), which we address here.

respectively. In this setting, the choice of the price dependence can help induce such a relationship, for example if both $E[D(p,\mathbf{z})]$ and $\text{CVaR}_\alpha[D(p,\mathbf{z})]$ are linear in p, then (2.5) is a quadratic function. For the lost sales setting, the concavity conditions depend more intricately on the demand distribution, and in particular, the increasing elasticity of *lost-sales rate*, defined as probability that the demand is greater than the order quantity. We refer the reader to Petruzzi and Dada (1999), Yao et al. (2006), Kocabıyıkoğlu and Popescu (2011) whose primary focus has been to derive the exact conditions for this setting. As our paper focuses on data-driven procedures without any distributional assumptions about the data generating process, in general it is not possible to ensure that these concavity conditions are satisfied by the observational data (except in artificial settings, e.g., with simulated data sets). Also, imposing any restrictions on the price dependence may also reduce the efficacy of capturing the true demand-price relationship in the data. Consequently, even though the corresponding optimal price always exists (i.e., for a continuous objective function (2.5) in a closed set $\mathcal{P}$), there may be multiple local maxima if the concavity conditions are violated in either of these optimization settings. Therefore, a standard univariate, derivative-free, non-linear optimization procedure (Brent 1973) can be used to directly obtain a solution for (2.5). When there is a unique maximum, this solution is the global optimum. Otherwise, solutions from multiple initial starting points using grid-search over the domain, can be used to examine and improve solution quality.

3. Some Relevant Statistical Estimation Methods

As noted in Section 2.1, the optimization formulation in Eq. (2.3) and Eq. (2.5) is completely specified in terms of the three quantities $E[D(p,\mathbf{z})]$, $\text{VaR}_\alpha[D(p,\mathbf{z})]$, and $\text{CVaR}_\alpha[D(p,\mathbf{z})]$. Each of these quantities can be directly estimated by a variety of different techniques and any appropriate combination of these techniques can be used in the optimization formulation, leading to a profusion of ways for implementing the computations. In practice, there is no prior knowledge of the corresponding demand distribution and it has to be estimated from historical sales data. Moreover, the estimations often involve large-scale data sets (either in the number of observations or covariates or both) that may exhibit a heteroskedastic relationship[4] between demand and multiple other covariates (also referred to as features/attributes/ explanatory variables) like the examples discussed in Section 1. This motivates us to consider (ideally) flexible, scalable data-driven, distribution-free, consistent statistical estimation approaches that can capture heteroskedasticity in the data.

[4] In a heteroskedastic data set, the variance of the linear regression error (or residual, in practice), conditional on the covariates, is not constant; it is sensitive to the covariates.

We consider three specific distribution-free approaches, which are based on generalized linear regression (GLR), mixed quantile regression (MQR), and superquantile regression (SQR) respectively. GLR is a popular and well-studied consistent mean estimation method where heteroskedasticity is explicitly modeled, using a location-scale model. The by-product of GLR are residuals which provide an estimate of the quantiles and superquantiles at specified $\alpha-$levels. To focus on the main contributions of the paper and for completeness, we defer the description of the GLR methodology adopted to estimate the three quantities of interest to Appendix 3.2 and refer the readers to this section for details. Similar to GLR, quantile regression (QR) is a popular and well-studied consistent quantile estimation method, where heteroskedasticity is implicitly modeled. With foundations built up on QR, SQR and MQR are two recent and different methods described in Rockafellar et al. (2014) and Chun et al. (2012) respectively that estimate the superquantiles directly from the data. The MQR method is restricted to homoskedastic distributions only, while the SQR method implicitly models heteroskedasticity in a similar spirit to QR but not scalable to large data sets. We provide appropriate extensions for adaptation to these techniques to the price-setting newsvendor problem in the following sections especially around scalability and heteroskedastic modeling for SQR and MQR respectively, when they fall short along these dimensions.

The possible combinations for implementing the computations for the price-setting newsvendor problem, just based on the different approaches discussed here, are shown in Table 1. The subscripts 0 or α represent the input quantile levels associated with the estimation procedures. Because $E[D(p,\mathbf{z})] = \text{CVaR}_0[D(p,\mathbf{z})]$, a module that computes CVaR can also be used to compute the mean.

$E[D(p,\mathbf{z})]$	$\text{VaR}_\alpha[D(p,\mathbf{z})]$	$\text{CVaR}_\alpha[D(p,\mathbf{z})]$
GLR	GLR+Residuals_α	GLR + Residuals_α
MQR_0	QR_α	MQR_α
SQR_0		SQR_α

Table 1 Possible statistical estimation methods to estimate the different quantities of interest.

We now describe the uniform notation we use across the different estimation methods. The response variable is denoted by Y, which is typically the demand $D(p,\mathbf{z})$ itself.[5] The set of regression covariates are denoted by $\mathbf{X}$. This set includes the constant or intercept term (except where explicitly indicated), along with other terms that involve linear and/or non-linear functional transformations of $(p,\mathbf{z})$, possibly with interactions terms. We assume that $\mathbf{X}$ and Y are random variables

[5] Quantile estimates are invariant under monotonic transformations (i.e., Y can be monotonic transformation of $D(p,\mathbf{z})$). On the other hand, mean and the superquantile estimates are not invariant (e.g., $E[\log Y] \neq \log E[Y]$) and therefore, here we fix $Y = D(p,\mathbf{z})$.

(denoting them by capitals) and that the decision maker has N independent and identically distributed (i.i.d) observations of $\langle \mathbf{X}, Y \rangle$ (from their underlying stochastic models), which we denote (in lower cases) by $\langle \mathbf{x}_i, y_i \rangle$ for $i = 1, \cdots, N$.[6] All expectations in the paper are conditional expectations (in the covariates), and this includes VaR and CVaR. This means, for example, expected value of Y is $E[Y|\mathbf{X}]$.

The regression models typically involve the estimation of $E[Y|\mathbf{X}]$, $\text{VaR}_\alpha[Y|\mathbf{X}]$ or $\text{CVaR}_\alpha[Y|\mathbf{X}]$. We assume that the regression functions involve linear combinations of the regression covariates with the coefficients being the parameters to be estimated from historical data, based on the appropriate regression formulation. We note that the regression functions are typically linear in the parameters, the ability to include nonlinear and interaction terms in the covariate effects, as well as nonlinear response transformations, provides sufficient generality for modeling a wide range of functional forms required in practical applications (Hastie et al. 2001).

3.1. Motivation for quantile-based methods

The fundamental building block of quantile-based methods is quantile regression, which is a method to estimate $\text{VaR}_\alpha[Y|\mathbf{X}]$ given α. One motivating reason for these methods is that it enables a broader class of stochastic demand functions to be modeled as it implicitly models heteroskedasticity. For instance, the specific example below describes a response variable for which different covariates are significant at different quantile levels. This characteristic cannot be modeled using the parameterization that is used for location-scale GLR models described in Appendix B.

EXAMPLE 1. This example describes a stochastic demand function with different demand drivers at different quantile levels. Consider the random variable ϵ, and let

$$Y = \boldsymbol{\beta}^T \mathbf{X} + \boldsymbol{\gamma}_1^T \mathbf{Z_1} \min\{\epsilon, \lambda\} + \boldsymbol{\gamma}_2^T \mathbf{Z_2} \max\{\lambda, \epsilon\}, \tag{3.1}$$

where the constant $\lambda = \text{VaR}_\zeta[\epsilon]$ for some value $\zeta \in (0,1)$ and $\mathbf{Z_1}, \mathbf{Z_2}$ are select non-intersecting covariates from $\mathbf{X}$. Note that

$$\text{VaR}_\alpha[\min\{\epsilon, \lambda\}] = \begin{cases} \lambda, & \alpha \geq \zeta, \\ \text{VaR}_\alpha[\epsilon], & \alpha < \zeta. \end{cases}$$
$$\text{VaR}_\alpha[\max\{\epsilon, \lambda\}] = \begin{cases} \text{VaR}_\alpha[\epsilon], & \alpha \geq \zeta, \\ \lambda, & \alpha < \zeta, \end{cases}$$

and similarly since

$$\text{CVaR}_\alpha[\min\{\epsilon, \lambda\}] = \begin{cases} \lambda, & \alpha \geq \zeta, \\ \text{CVaR}_\alpha[\epsilon] + (\lambda - \tau)\frac{1-\zeta}{1-\alpha}, & \alpha < \zeta. \end{cases}$$
$$\text{CVaR}_\alpha[\max\{\epsilon, \lambda\}] = \begin{cases} \text{CVaR}_\alpha[\epsilon], & \alpha \geq \zeta, \\ \frac{\zeta-\alpha}{1-\alpha}\lambda + \frac{1-\zeta}{1-\alpha}\tau, & \alpha < \zeta. \end{cases}$$

where $\tau = \text{CVaR}_\zeta[\epsilon]$, we have

[6] The approach can be extended to the non i.i.d settings using specialized methods in the regression literature for addressing serial and spatial correlation effects to improve the efficiency of the statistical estimators.

$$\text{VaR}_\alpha[Y|\mathbf{X},\mathbf{Z_1},\mathbf{Z_2}] = \boldsymbol{\beta}^T\mathbf{X} + \text{VaR}_\alpha[\min\{\epsilon,\lambda\}]\boldsymbol{\gamma}_1^T\mathbf{Z_1} + \text{VaR}_\alpha[\max\{\epsilon,\lambda\}]\boldsymbol{\gamma}_2^T\mathbf{Z_2}.$$

along with a similar expression for $\text{CVaR}_\alpha[Y|\mathbf{X},\mathbf{Z_1},\mathbf{Z_2}]$.

A situation highlighted in this specific example arises frequently in practice. For example, the uppermost quantiles of the stochastic electricity demand are likely to be quite sensitive to price and weather covariates, whereas the lowermost quantiles are quite insensitive to these covariates.

The estimate for $\text{VaR}_\alpha[Y]$ is obtained using quantile regression (Koenker and Bassett 1978), given a response Y, covariates X and quantile level α, in the form:

$$\text{VaR}_\alpha[Y|\mathbf{X}] = \boldsymbol{\beta}_v^T\mathbf{X}. \tag{3.2}$$

Quantile regression involves solving the following optimization problem to estimate $\boldsymbol{\beta}_v$:

$$\mathbf{QR}: \quad \hat{\boldsymbol{\beta}}_v = \arg\min_{\boldsymbol{\beta}} \quad \frac{1}{N}\sum_{i=1}^{N}\psi_\alpha(y_i - \boldsymbol{\beta}^T\mathbf{x}_i) \tag{3.3}$$

where $\psi_\theta(t) = \theta[t]^+ + (1-\theta)[-t]^+$ and $\theta = [0,1]$. QR results in a consistent estimator as long as the true quantile is linear in the covariate space(Koenker and Bassett Jr 1982), as in a (linear) location- (linear) scale model for example (or the above example). In particular, for a fixed data sample, the estimates of these quantile regression coefficients will depend on α, and one can use significance tests to ascertain if the differences in these coefficients across the different α values are indicative of heteroskedasticity. QR can be re-written as a simple linear program and standard efficient subroutines are available in most commercial statistical software packages.

3.2. Superquantile Regression (SQR)

Recently, Rockafellar et al. (2014) proposed *superquantile regression* for estimating $\text{CVaR}_\alpha[Y|\mathbf{X}]$ conditional on a set of covariates $\mathbf{X}$. In this method, the authors provide a suitable modification of the error measure used in quantile regression leading to the superquantile regression estimates for $\text{CVaR}_\alpha[Y|\mathbf{X}]$. This modified error measure is based on an auxiliary response variable whose quantiles, by construction, are equivalent to the desired superquantiles $\text{CVaR}_\alpha[Y|\mathbf{X}]$. This approach has strong theoretical foundations in the fundamental theory of risk measure and the risk quadrangle (Rockafellar and Uryasev 2013). While we are not aware of a theoretical proof of the consistency of this estimation method, we believe it has properties similar to quantile regression because it applies the essential idea of quantile regression to the above stated auxiliary response variable. This method is more appealing than GLR and MQR, as it inherently models heteroskedasticity with no additional specifications (e.g., model or covariates to model heteroskedasticity). Also, unlike MQR that we discuss in the following section, there are no numerical discretizations required. Below we describe the method, its drawbacks and our extensions.

The estimate for $\text{CVaR}_\alpha[Y]$ is obtained using superquantile regression (Rockafellar et al. 2014), given a response Y, covariates X and quantile level α, in the form:

$$\text{CVaR}_\alpha[Y|\mathbf{X}] = \beta_{c,0} + \boldsymbol{\beta}_c^T\mathbf{X} \tag{3.4}$$

where the constant term is explicit and $\mathbf{X}$ consists of columns for all the covariates except the column of ones unlike the standard notation we use in the rest of the paper. Then the corresponding minimization of the error measure in superquantile regression is formulated as

$$\boldsymbol{\beta}_c = \arg\min_{\boldsymbol{\beta}} \frac{1}{1-\alpha}\int_\alpha^1 \text{CVaR}_\tau[Y-\boldsymbol{\beta}^T\mathbf{X}]d\tau - E[Y-\boldsymbol{\beta}^T\mathbf{X}], \tag{3.5}$$

$$\beta_{c,0} = \text{CVaR}_\alpha[Y-\boldsymbol{\beta}_c^T\mathbf{X}]. \tag{3.6}$$

With data samples $\{\mathbf{x}_i, y_i\}$, the residual random variable, $Y-\boldsymbol{\beta}_c^T\mathbf{X}$, has a discrete support, thereby leading to a cumulative distribution function which has a piecewise constant structure. This structure, along with the formulation of the CVaR estimation as a minimization problem in Rockafellar and Uryasev (2000), allows Problem (3.5) to be expressed as the following nonlinear mathematical program,

$$\begin{aligned}\hat{\boldsymbol{\beta}}_c = \arg\min_{\boldsymbol{\beta},\mathbf{U}} \;& \frac{1}{1-\alpha}\sum_{k=N_\alpha}^{N-1}(\kappa_k-\kappa_{k-1})U_k + \frac{1}{N(1-\alpha)}\sum_{k=N_\alpha}^{N-1}\sum_{i=1}^{N} a_k \max\{y_i-\boldsymbol{\beta}^T\mathbf{x_i}-U_k,0\} \\ &+ \frac{1}{N(1-\alpha)}\max_{i=1\ldots N}(y_i-\boldsymbol{\beta}^T\mathbf{x}_i) - \frac{1}{N}\sum_{i=1}^{N}(y_i-\boldsymbol{\beta}^T\mathbf{x}_i)\end{aligned} \tag{3.7}$$

In the above model, $N_\alpha = \lceil N\alpha \rceil$, and the decision variables include $\boldsymbol{\beta}_c \in \mathbb{R}^n$, $\mathbf{U} \in \mathbb{R}^{N-N_\alpha}$ where n is the number of coefficients to be estimated except the intercept term. Further $\kappa_{N_\alpha-1} = \alpha$, and $\kappa_k = \frac{k}{N}$, which capture the various piecewise constant levels within the limits of the integration in Eq. (3.5), and $a_k = \ln(1-\kappa_{k-1}) - \ln(1-\kappa_k)$.

Linearization using additional decision variables, $\mathbf{V} \in \mathbb{R}^{N(N-N_\alpha)}$, and $W \in \mathbb{R}$ yields the following linear program as developed in Rockafellar et al. (2014).

$$\begin{aligned}\textbf{SQR}: \; \hat{\boldsymbol{\beta}}_c = \arg\min_{\boldsymbol{\beta},\mathbf{U},\mathbf{V},W} \;& \frac{1}{1-\alpha}\sum_{k=N_\alpha}^{N-1}(\kappa_k-\kappa_{k-1})U_k + \frac{1}{N(1-\alpha)}\sum_{k=N_\alpha}^{N-1}\sum_{i=1}^{N} a_k V_{ki} \\ &+ \frac{1}{N(1-\alpha)}W - \frac{1}{N}\sum_{i=1}^{N}(y_i-\boldsymbol{\beta}^T\mathbf{x}_i),\end{aligned} \tag{3.8}$$

$$\text{s.t.}, \quad V_{ki} \geq y_i - \boldsymbol{\beta}^T\mathbf{x}_i - U_k, \quad \forall\; k = N_\alpha,\ldots,N-1,\; i=1,\ldots,N, \tag{3.9}$$

$$V_{ki} \geq 0, \qquad \forall\; k = N_\alpha,\ldots,N-1,\; i=1,\ldots,N, \tag{3.10}$$

$$W \geq y_i - \boldsymbol{\beta}^T\mathbf{x}_i, \qquad \forall\; i=1,\ldots,N. \tag{3.11}$$

The estimation approach can then be summarized as Procedure 2. The methodology for the price optimization is then similar to that described in Section 3.3, except that the CVaR terms can now estimated using superquantile regression.

Procedure 2 Superquantile regression for CVaR estimation

Input: Data $\{\mathbf{x}_i, y_i\}$ for $i = 1, ..., N$, and the level α. Note that $\mathbf{x}_i$ does not include the constant 1 covariate corresponding to the intercept term.

1: Set up and solve the linear program **SQR**. The solution produces an estimate for $\hat{\boldsymbol{\beta}}_c$.

2: Obtain constant term, using Eq. (3.6), by computing the empirical CVaR of the residual corresponding to $\hat{\boldsymbol{\beta}}_c$. This is computable as,

$$\hat{\beta}_{c,0} = \frac{1}{N\alpha} \sum_{i=1}^{\lfloor N\alpha \rfloor} R^{(i)} + \Big(\frac{N\alpha - \lfloor N\alpha \rfloor}{N\alpha} \Big) R^{(\lceil N\alpha \rceil)},$$

where $R = (Y - \hat{\boldsymbol{\beta}}_c^T \mathbf{X})$, and $R^{(i)}$ represent the decreasing order statistics of R over the empirical sample $\{\mathbf{x}_i, y_i\}$, $i = 1, \ldots, N$, i.e. $R^{(1)} \geq \cdots \geq R^{(N)}$.

Output: Estimates $[\hat{\beta}_{c,0}; \hat{\boldsymbol{\beta}}_c]$ for superquantile regression, i.e., $\text{CVaR}_\alpha[Y|\mathbf{X}] = \hat{\beta}_{c,0} + \hat{\boldsymbol{\beta}}_c^T \mathbf{X}$.

3.2.1. An Efficient Algorithm For Superquantile Regression

Although the above method makes no assumptions about homoskedasticity, which has also been noted by the authors in their paper, it does not scale beyond a few hundred sample points. This is because the above formulation **SQR** involves $\mathcal{O}(N^2)$ number of variables as well as constraints. When N is large, this quadratic complexity makes the above formulation impractical in terms of the computational time needed to solve the linear program. We instead present an alternative linearization of Eq. (3.7), **RSQR** (reformulated SQR), which enables the derivation of an efficient cutting plane algorithm. We begin with a technical observation which is then used to arrive at the alternative linear reformulation.

Claim 2. *Let $\mathcal{N} = \{1 \ldots N\}$. Then, for any fixed index k, the following equality holds, where $\mathcal{P}(\mathcal{N})$ denotes the power set, i.e. the set of all subsets of $\mathcal{N}$.*

$$\sum_{i=1}^{N} \max\{y_i - \boldsymbol{\beta}^T \mathbf{x}_i - U_k, 0\} = \max_{J_k \in \mathcal{P}(\mathcal{N})} \sum_{i \in J_k} (y_i - \boldsymbol{\beta}^T \mathbf{x}_i - U_k) \tag{3.12}$$

Expressing the expectation in Eq. (3.7) as a finite summation and applying Eq. (3.12) leads to the following formulation.

$$\begin{aligned} \hat{\boldsymbol{\beta}}_c = \underset{\boldsymbol{\beta}, \mathbf{U}, W}{\arg\min} \quad & \frac{1}{1-\alpha} \sum_{k=N_\alpha}^{N-1} (\kappa_k - \kappa_{k-1}) U_k + \frac{1}{N(1-\alpha)} \sum_{k=N_\alpha}^{N-1} a_k \max_{J_k \in \mathcal{P}(\mathcal{N})} \sum_{i \in J_k} (y_i - \boldsymbol{\beta}^T \mathbf{x}_i - U_k) \\ & + \frac{1}{N(1-\alpha)} W - \frac{1}{N} \sum_{i=1}^{N} (y_i - \boldsymbol{\beta}^T \mathbf{x}_i), \end{aligned} \tag{3.13}$$

$$\text{s.t.,} \quad W \geq y_i - \boldsymbol{\beta}^T \mathbf{x}_i, \quad \forall\, i = 1, \ldots, N. \tag{3.14}$$

We may then linearize the above formulation using an exponential number of constraints, with only $\mathcal{O}(N)$ additional variables, T_k, as follows.

$$\textbf{RSQR}: \ \hat{\boldsymbol{\beta}}_c = \underset{\boldsymbol{\beta},\mathbf{U},W,\mathbf{T}}{\arg\min} \quad \frac{1}{1-\alpha}\sum_{k=N_\alpha}^{N-1}(\kappa_k-\kappa_{k-1})U_k + \frac{1}{N(1-\alpha)}\sum_{k=N_\alpha}^{N-1} a_k T_k \tag{3.15}$$

$$+\frac{1}{N(1-\alpha)}W - \frac{1}{N}\sum_{i=1}^{N}(y_i-\boldsymbol{\beta}^T\mathbf{x}_i),$$

$$\text{s.t.,} \quad T_k \geq \sum_{i\in J_k} y_i - \boldsymbol{\beta}^T\mathbf{x}_i - U_k, \quad \forall\ J_k \in \mathcal{P}(\mathcal{N}),\ k=N_\alpha,\ldots,N-1, \tag{3.16}$$

$$T_k \geq 0, \qquad \forall\ k=N_\alpha,\ldots,N-1, \tag{3.17}$$

$$W \geq y_i - \boldsymbol{\beta}^T\mathbf{x}_i, \qquad \forall\ i=1,\ldots,N. \tag{3.18}$$

The proposed reformulation allows us to derive an efficient decomposition algorithm. We first present the following claim that lets us successfully seed the following algorithm in its very first iteration. Let $\textbf{RSQR}-\textbf{RELAX}$ denote a relaxation of **RSQR**, where we replace $\mathcal{P}(\mathcal{N})$ in constraint Eq. (3.15) with respective subsets $\mathcal{J}_k \subseteq \mathcal{P}(\mathcal{N})$, for each index k.

Claim 3. *Let $\mathcal{J}_k = \{\mathcal{N}\}, \forall\ k$. Then, the corresponding relaxation,* $\textbf{RSQR}-\textbf{RELAX}$ *is a bounded linear program.*

The decomposition algorithm may then be presented as in Procedure 3.

Theorem 1. *Procedure 3 converges in finite time and solves problem* **RSQR**, *equivalently* **SQR**, *upon convergence.*

In practice, convergence is realized in far fewer iterations than the cardinality of the power set. The above decomposition code written in MATLAB using CPLEX as the MIP solver is publicly available with sample synthetic datasets [7] for downloading. We present below an empirical investigation of the computational performance of the algorithm.

3.2.2. Computational Results We consider the heteroskedastic generating model **G1** described in Section 4.1 below with a unit Normal error distribution. Fig. 4 shows the average computational time needed to solve formulation **SQR** using Procedure 2, as well as the proposed reformulation using Procedure 3, as a function of N, namely, the number of sample points at two quantile levels $\alpha = 0.7, 0.85$. For each value of N and α, the plot shows the mean value and error bars estimated over 200 independently generated data sets, each of size N. All computations were carried out using Matlab/CPLEX on 64-bit Macbook Pro, Intel®Core™ i7 @ 2.5 GHz, 16 GB RAM. While Procedure 2 fails to acceptably scale beyond a five hundred sample points, the proposed decomposition algorithm performs well even for really large sets. In our experiments, we observe that it is often more stable to implement $\textbf{RSQR}-\textbf{RELAX}$ in terms of the scaled covariate $\hat{\mathbf{X}} = (\mathbf{X}-\mu_{\mathbf{X}})/\sigma_{\mathbf{X}}$ without any loss of generality.

[7] http://researcher.watson.ibm.com/researcher/files/us-dharmash/SuperquantileDecomposition.zip

Procedure 3 Decomposition Algorithm for Superquantile regression used for CVaR estimation

Input: Data $\{\mathbf{x}_i, y_i\}$ for $i = 1, ..., N$, and the level α. Note that $\mathbf{x}_i$ does not include the constant covariate corresponding to the intercept term.

1: Initialize $\mathcal{J}_k = \{\mathcal{N}\}$, $\forall\ k = N_\alpha, \ldots, N-1$.

2: Solve the relaxed linear program $\mathbf{RSQR-RELAX}$ implied by the current value of $\mathcal{J}_k$. Let $\hat{\boldsymbol{\beta}}$ and $\hat{U}_k$ be the optimal solution values for these variables.

3: For each k, identify the most violated constraint, relative to the full set of constraints in Eq. (3.15). This is computable in $\mathcal{O}(N)$ effort for each k as

$$J_k^* = \{i \in \mathcal{N} \mid y_i - \hat{\boldsymbol{\beta}}^T \mathbf{x}_i - \hat{U}_k > 0\},$$

and do $\mathcal{J}_k = \mathcal{J}_k \cup J_k^*$. This results in adding a constraint for each k.

4: Repeat steps 2-3, until convergence of the LP solution in step 2. Let the converged solution for variable $\boldsymbol{\beta}$ be denoted as $\hat{\boldsymbol{\beta}}_c$.

5: Obtain constant term, using Eq. (3.6), by computing the empirical CVaR of the residual corresponding to $\hat{\boldsymbol{\beta}}_c$. This is computable as,

$$\hat{\beta}_{c,0} = \frac{1}{N\alpha} \sum_{i=1}^{\lfloor N\alpha \rfloor} R^{(i)} + \left(\frac{N\alpha - \lfloor N\alpha \rfloor}{N\alpha}\right) R^{(\lceil N\alpha \rceil)},$$

where $R = (Y - \hat{\boldsymbol{\beta}}_c^T \mathbf{X})$, and $R^{(i)}$ represent the decreasing order statistics of R over the empirical sample $\{\mathbf{x}_i, y_i\}$, $i = 1, \ldots, N$, i.e. $R^{(1)} \geq \cdots \geq R^{(N)}$.

Output: Estimates $[\hat{\beta}_{c,0}; \hat{\boldsymbol{\beta}}_c]$ for superquantile regression, i.e., $\text{CVaR}_\alpha[Y|\mathbf{X}] = \hat{\beta}_{c,0} + \hat{\boldsymbol{\beta}}_c^T \mathbf{X}$.

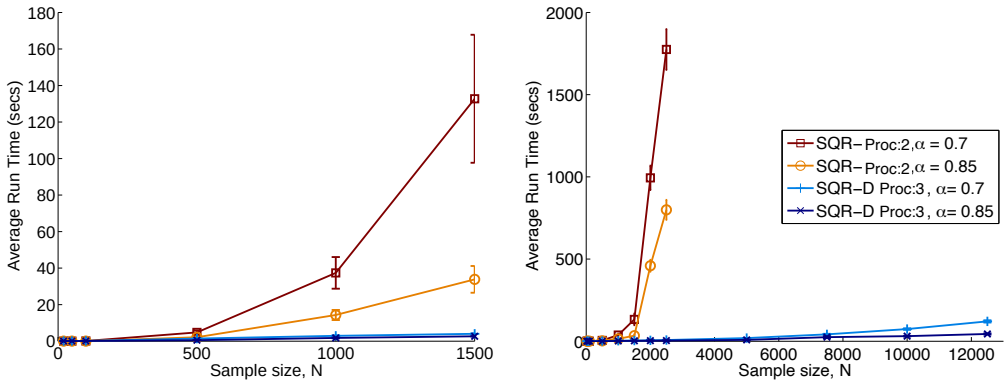

Figure 4 **Computational run times for Procedure 2 and Procedure 3 as a function of N for $\alpha = 0.7, 0.85$. The left plot is zoomed version of the right plot for the range [0, 1500].**

3.3. Mixed Quantile Regression (MQR)

An alternative method for the evaluation of $\text{CVaR}_\alpha[Y|\mathbf{X}]$ is motivated from the module for estimating $\text{VaR}_\alpha[Y|\mathbf{X}]$ given α exploiting Eq. (2.7). More specifically, it is based on the numerical quadrature of the integral in Eq. (2.7). Hence $\text{CVaR}_\alpha[Y|\mathbf{X}]$ will be a linear combination, with appropriate quadrature weights, of $\text{VaR}_{\alpha'}[Y|\mathbf{X}]$ evaluated at certain quadrature nodes α', where

$\alpha < \alpha' < 1$ (we typically use quadrature rules that only involve nodes that are in the interior of the interval to avoid the estimation of the extremal quantiles $\alpha' = 0, 1$). From Eq. (3.2), the corresponding estimate for $\text{CVaR}_\alpha[Y|\mathbf{X}]$ also has the form (Peracchi and Tanase 2008, Leorato et al. 2012):

$$\text{CVaR}_\alpha[Y|\mathbf{X}] = \boldsymbol{\beta}_c^T \mathbf{X}. \tag{3.19}$$

A related method was first described in Rockafellar et al. (2008), and further explored by Chun et al. (2012) where it is aptly referred to as a mixed quantile regression (MQR). However, as presented there, MQR is restricted (by design as explained below) to homoskedastic data sets, where the variance of the linear regression error (or residual, in practice), conditional on the covariates, is constant. The authors show that the estimator is consistent under a homoskedastic setting with (additive) i.i.d errors. Below we extend this method to incorporate heteroskedasticity, i.e. where the variance of the linear regression error is sensitive to covariates. After we describe our extension, we show how we recover the formulation by Chun et al. (2012) as a special case.

In order to describe our formulation, we let $\sum_{j=1}^{M} w_j \text{VaR}_{\alpha_j}[Y|\mathbf{X}]$ be the discretization of Eq. (2.7), where α_j and w_j denote the quadrature nodes and weights respectively. Then, we have

$$\textbf{MQR}: \quad \hat{\boldsymbol{\beta}}_c = \underset{\boldsymbol{\beta}, \boldsymbol{\tau}_j}{\arg\min} \quad \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{M} w_j \psi_{\alpha_j} \left(y_i - (\boldsymbol{\tau}_j + \boldsymbol{\beta})^T \mathbf{x}_i \right), \tag{3.20}$$

$$\text{s.t.}, \quad \sum_{j=1}^{M} w_j \boldsymbol{\tau}_j = \mathbf{0}, \tag{3.21}$$

$$\bar{\mathbf{e}}.\boldsymbol{\tau}_j = \mathbf{0} \quad \forall\ j = 1, \ldots, M. \tag{3.22}$$

where $\psi_\theta(t)$ is the loss function for quantile regression with $\sum_{j=1}^{M} w_j = 1$ and $\bar{\mathbf{e}}$ is a vector such that $\bar{e}_p = 1$ if the p^{th} covariate (besides intercept) does not affect the variance of the linear regression error and 0 otherwise. For example, if a simple mid-point quadrature rule is used for the discretization of Eq. (2.7), then with $\Delta := M^{-1}(1-\alpha)$, we have $w_j = (1-\alpha)^{-1}\Delta$ and $\alpha_j := \alpha + (j - 0.5)\Delta,\ j = 1, .., M$. Similarly, if a Gauss-Legendre quadrature rule is used, then $w_j = 0.5(1-\alpha)\delta_j$ and $\alpha_j = 0.5\left[(1-\alpha)\xi_j + (1+\alpha)\right]$ where δ_j and $\xi_j \in (-1, 1)$ respectively denote the weights and nodes of the corresponding M-point quadrature rule.

The objective function in **MQR** is a weighted sum of individual loss functions for each quantile level α_j, with the corresponding regression function $(\boldsymbol{\beta} + \boldsymbol{\tau}_j)^T \mathbf{X}$. Constraint (3.21) ensures that $\hat{\boldsymbol{\beta}}_c^T \mathbf{X}$ is the desired estimate of $\text{CVaR}_\alpha[Y|\mathbf{X}]$. Constraint (3.22) imposes the condition that $\boldsymbol{\tau}_j$ is set to zero for any covariate (besides intercept) that does not affect the variance of the linear regression error. To understand the impact of this constraint, let us notionally partition the space of covariates into two sets, namely, those that show heteroskedasticity in their relationship with demand, i.e. those that affect the variance of the corresponding linear regression error (say, set A),

and those that do not (say set B). For covariates in set B, constraint (3.22) restricts all quantiles to have an identical coefficient (the elements of the $\boldsymbol{\tau}_j$ are set to zero) and thus ensuring parallelism in that covariate dimension. On the other hand, for covariates in set A, without constraint (3.22), the elements in $\boldsymbol{\tau}_j$ have a degree of freedom, that they can vary across quantiles, and thus the estimation method can choose to have quantiles that are not parallel to obtain a better fit, similar to Fig. 2. Thus the above **MQR** formulation by design allows for conditional homoskedasticity and conditional heteroskedasticity over the space of covariates.

In the case of a fully homoskedastic data set, i.e. where all covariates are notionally in set B above, all $\boldsymbol{\tau}_j$'s are set to zero except for the intercept term, we recover the MQR formulation in Chun et al. (2012) as a special case. In the fully heteroskedastic case, where all $\boldsymbol{\tau}_j$'s can be non-zero, the problem can be decoupled (with a change of variables $\boldsymbol{\beta}_j := \boldsymbol{\beta} + \boldsymbol{\tau}_j$) and evaluated with independent and individual quantile regressions for each α_j, which can be then aggregated to obtain $\boldsymbol{\beta}$ $(= \sum_{j=1}^{M} w_j \boldsymbol{\beta}_j)$. This can be implemented with the widely-available quantile regression modules. On the other hand, the extended MQR method described here is more general, i.e. it can limit the subset of covariates over which heteroskedasticity is manifested and requires implementing a specialized linear program. It requires this partition of covariates as an input, and some preliminary analysis using t-tests, should be used to ascertain if the differences in the coefficients obtained from quantile regression across the different α_j values are significant, and therefore indicative of heteroskedasticity, and to accordingly partition the covariates.

In summary, given $\mathbf{X}$ and the homoskedastic covariates, the three quantities needed to solve the price-setting newsvendor problem (see Fig. 3 and Section 2.2) are as follows:

$$\widehat{E}[Y|\mathbf{X}] = \hat{\tilde{\boldsymbol{\beta}}}_c^T \mathbf{X}, \tag{3.23}$$

$$\widehat{\mathrm{VaR}}_\alpha[Y|\mathbf{X}] = \hat{\boldsymbol{\beta}}_v^T \mathbf{X}, \tag{3.24}$$

$$\widehat{\mathrm{CVaR}}_\alpha[Y|\mathbf{X}] = \hat{\boldsymbol{\beta}}_c^T \mathbf{X}, \tag{3.25}$$

where $\hat{\tilde{\boldsymbol{\beta}}}_c, \hat{\boldsymbol{\beta}}_v$ and $\hat{\boldsymbol{\beta}}_c$ are outputs of subroutines that solve optimization formulations **MQR**, **QR**, and **MQR** with input quantiles $0, \alpha$ and α respectively.

4. Monte-Carlo Simulation Study

In this section, the proposed statistical estimation methodologies for the price-setting newsvendor problem described in Section 3, are evaluated through a Monte Carlo simulation study. To fix ideas, we focus on the lost-sales formulation of the price-setting newsvendor problem, although the general conclusions carry over to the emergency order formulation as well.

In the Monte Carlo simulation, we use an explicitly known stochastic price-demand function and obtain the exact optimal solutions to the price-setting newsvendor problem using the methods

in Section 2.2. We also generate simulated data sets from these known stochastic price-demand functions, and estimate the optimal solutions for each simulated data set using the methods in Section 3. These results are used to evaluate the statistical properties of the estimated optimal solutions with respect to the true solution, as well as the coverage and length of their estimated bootstrap confidence intervals, described further below.

For a pricing problem, in general, in the experiments with real-data, the true demand response for a price unobserved in the history is unknown. This makes the evaluation of even a method, let alone across multiple methods (such as the alternatives considered in Table 1), difficult or subjective at best and dependent of the quality of fit of a predictive model. Therefore, we focus on Monte Carlo simulation results in this section of the paper.

4.1. Simulated Data Sets

Denoting the stochastic price-demand functions by Y, we consider two such explicit functions which are motivated from Eqs. (B.1) and (3.1) respectively (for simplicity of exposition, price is the only demand driver that is included; extension are discussed in Section 4.5):

$$\textbf{G1.} \qquad Y = \beta_0 + \beta_1 p + \left(\gamma_0 + \gamma_1 p + \gamma_2 p^2\right)\, \epsilon, \tag{4.1}$$

where ϵ is a random variable with mean 0 that is specified further below, and $\beta_0 = 200.0,\ \beta_1 = -35.0,\ \gamma_0 = 36.0,\ \gamma_1 = -12.0,\ \gamma_2 = 2.1$.

$$\textbf{G2.} \qquad Y = \beta_0 + \beta_1 p + \beta_2 p^2 + (\gamma_0 + \gamma_1 p)\, \epsilon^- + \gamma_2 p^2\, \epsilon^+, \tag{4.2}$$

where $\epsilon^- = \min\{\epsilon, 0\},\ \epsilon^+ = \max\{0, \epsilon\}$ and ϵ is $N(0,1)$ i.e., a Normal distribution with mean 0 and standard deviation 1. Here, $\beta_0 = 215.0,\ \beta_1 = -37.0,\ \beta_2 = -1.5\ \text{CVaR}_\epsilon(0.5) = -1.1968,\ \gamma_0 = 36.0, \gamma_1 = -4.0, \gamma_2 = 3$.

The demand functions in the generating models **G1** and **G2** have means that are decreasing linear functions of price, and variances that are non-monotonic quadratic functions of price as motivated by Raz and Porteus (2006). Note that the mean for the generating model **G2** is obtained by evaluating $\text{CVaR}_0[Y|p]$. As discussed in Section 1.2, price-demand functions with these characteristics are of practical importance.

The parameters for the price-setting newsvendor problem are taken to be $c = 1.0,\ s = 0.5$, and $v = 1.0$. The unit price p, which is the decision variable, is constrained to the interval $(1.5, 4.0)$. These parameter values are inspired by an example in Lau and Lau (1988), although that paper only considered a homoskedastic demand models with normal errors.

For the random variable ϵ in the generating model **G1**, we consider the following distributions:

1. Normal(0,1): Normal distribution with mean 0 and standard deviation 1.

Generating Model	Distribution for ϵ	Optimal Solutions		
		Price, p^*	Order Quantity, x^*	Profit Π^*
G1	Normal	3.32	105.57	178.74
	Gamma	3.28	114.77	167.76
	Lognormal	3.22	113.60	155.85
	Student T	3.28	111.5	169.58
	Mixture	3.34	134.18	184.41
G2	Normal	3.16	119.05	169.04

Table 2 True optima for the lost sales price-setting newsvendor problem with the stochastic demand model.

2. Gamma(2,1): Gamma distribution with shape 2 and rate 1 (equivalently with mean 2 and standard deviation $\sqrt{2}$), recentered to have mean 0.
3. Log-normal(0,1): Log-normal distribution with mean 0 and standard deviation 1 on the variable's log-scale, recentered to have mean 0.
4. Student's t(3): Student's t-distribution with 3 degrees of freedom, with mean 0 and standard deviation $\sqrt{3}$.
5. Mixture(-2,2): Mixture of two normal distributions, $N(2,1)$ and $N(-2,1)$, with equal weight and standard deviations 1 each.

The Gamma(2) and Log-normal(0,1) distributions which are re-centered to have mean 0, are asymmetric distributions. The Student's t(3) distribution is symmetric but is heavy-tailed. The Mixture(2,-2) distribution is also symmetric but is bi-modal unlike the rest.

The exact optimal solutions to the price-setting newsvendor problem for the price-demand functions in generating model **G1** and **G2** are given in Table 2. The corresponding sample estimators for the quantities in Table 2, are respectively denoted by $\hat{p}^*$ for optimal price, $\hat{x}^*$ for optimal order quantity, and $\hat{\Pi}^*$ for optimal profit.

The number of data points in each simulated data set is denoted by N and our results are obtained for values ranging from $N = 50$ to $N = 1500$. The covariate values for p for the individual cases in each simulated data set are obtained by independent uniform sampling from the allowed range in the interval $(1.5, 4.0)$. The number of simulated data sets used in the Monte Carlo evaluation is denoted by N_{mc} and is chosen to be 200.

4.2. Estimation methods used in experiments

Table 1 described the profusion of ways for implementing the computations to solve the price-setting newsvendor problem. To fix ideas, we implement and compare the following techniques:

- For the generating model, **G1**, we use the GLR method to evaluate $E[D(p, \mathbf{z})]$ and compare three methods that differ in the way they compute $\text{CVaR}_\alpha[D(p, \mathbf{z})]$ and $\text{VaR}_\alpha[D(p, \mathbf{z})]$ respectively. The first method uses the residuals of the GLR for the superquantile and quantile estimations,

the second method uses MQR and QR and the last method uses SQR and QR respectively. We refer to these three methods as GLR, MQR and SQR respectively.

The reason we choose the same mean estimator across the three different methods is to do with many reasons: (1) Eq. (4.1) has the same form as Eq. (B.1) and Eq. (B.2) (i.e., set $\mu = \beta_0 + \beta_1 p$, $\sqrt{\phi} = \gamma_0 + \gamma_1 p + \gamma_2 p^2$, $V(\mu) = 1$, with $g(\mu)$ being the identity link function, and $h(\phi)$ being the square-root link function). (2) The mean coefficient retrieval is expected to be good if the weights are a good approximation in the weighted least squares subroutine (step 2 when $V(\mu) = 1$) in GLR (3) A common mean estimator across methods also enables one to focus on the quality of the superquantile and the quantile estimation methods.

- For the generating model, **G2**, we use two different mean estimation methods: GLR and MQR_0 (referred to as GLR mean and MQR mean respectively) and compare these against the three different ways described above to estimate the superquantile and quantile, i.e., with GLR residuals, MQR-QR and SQR-QR (referred to as GLR, MQR and SQR respectively).

For both the generating models, **G1** and **G2**, we additionally include as a baseline the standard ordinary least squares (OLS) approach that (incorrectly) assumes homoskedasticity, with the goal of highlighting the importance of modeling heteroskedasticity.

The GLR estimation was based on Procedure 4 and the corresponding residuals, as also the case with OLS, are computed using Eqs. (B.3–B.4). In our implementations, we define the GLR procedure to have successful convergence if the number of iterations in Procedure 4 is less than 50. The QR and MQR methods refer to the solutions of **QR** and **MQR** respectively while SQR refers to the decomposition algorithm presented in Procedure 3. For the MQR method we use a simple uniform discretization where $\Delta = 0.01$.

4.3. Performance Metrics - Perfect Hindsight

In this simulation study, our goal is to understand the statistical performance of the estimation techniques in retrieving the true optimal price and the true optimal order quantity. To avoid information overload we only focus on the effect of the estimated price and order quantity on the true realized objective that we denote by 'Realized $\text{Profit}_{i,m}$' for an finite sample instance i using the method m. We compare this realized objective against the maximum realizable objective, denoted by 'Profit*' had we offered the true optimal price, p^* and stocked the true optimal order quantity, x^*. We use (absolute) error as a measure of performance between the two objectives i.e.,

$$\text{AE}_{i,m} = \frac{\text{Profit}^* - \text{Realized Profit}_{i,m}}{\text{Profit}^*} \tag{4.3}$$

We estimate this measure of error for every finite sample Monte Carlo data set i using a variety of methods m described in Section 4.2. We present the mean and standard error of the $\text{AE}_{i,m}$

over all the instances i for each method m. We increase the number of observations to study the statistical unbiasedness (mean decreasing to 0), consistency (standard error going to zero) and rate-of-convergence across methods.

4.4. Results and discussion

We summarize our results about the mean and standard error of $\text{AE}_{i,m}$ in Tables 3 and 4. Table 3 focuses on generating model **G1** and presents the results for various distributions of error discussed in Section 4.1 and four different estimation procedures discussed in Section 4.2. Table 4 focuses on generating model **G2** and presents the results for seven different estimation methods. In the tables, N refers to the number of data points in each instance of the data set 'Instances' denotes the number of instances, amongst the $N_{mc} = 200$ randomly generated instances, where the GLR procedure successfully converged. The mean and standard error results are presented only on the instances where the GLR procedure successfully converged.

For both the generating models **G1** and **G2**, with an exception of sparse instances in the Normal and Gamma distributions for **G1**, the OLS method underperforms compared to all the other methods, and the performance significantly degrades with the size of N. Specifically, when $N = 1500$, the best method relative to OLS is as high as $\sim$92% better for **G2** and greater than $\sim$50% for **G1**. Our observation of the poor performance of OLS seems different from Chun et al. (2012) who compared the fit of OLS and MQR methods for superquantile estimation in a homoskedastic setting (GLR is OLS in a homoskedastic setting) and conclude that even though theoretically both estimators are consistent, Monte Carlo simulations indicate that OLS estimators perform better than their MQR counterparts. The key in a homoskedastic setting to capture the mean accurately while the quantiles and superquantiles capture the effect of the empirical noise distribution from the residuals. Unlike this, in a heteroskedastic model (e.g., models **G1** and **G2**) it is not only important to capture the mean, but also the other quantities accurately (and estimating one depends on estimating the other for the GLR estimator). Here OLS is unable to estimate the mean well (and hence even its residuals), while its counterparts GLR that explicitly models heteroskedasticity and MQR and SQR that implicitly model heteroskedasticity are effective.

It can be observed from Table 3 that for generating model **G1** the GLR method tends to outperform the MQR and SQR for the Normal and Gamma distributions ($\sim$ 17% and 8% respectively relative to quantile methods when $N = 1500$) especially as N becomes larger. The reverse is true, i.e., GLR tends to underperform compared to MQR and SQR for the other distributions such as Student t, Lognormal and Mixture for larger N ($\sim$ 50%, 38% and 8% respectively relative to GLR when $N = 1500$). Recall that in this experiment, we focus only the variation in the quantile and superquantile estimation procedures, while the mean for GLR, MQR and SQR methods are derived from the GLR estimator.

N	Normal					Gamma				
	Instances	OLS	GLR	SQR	MQR	Instances	OLS	GLR	SQR	MQR
50	198	**0.907** (0.092)	0.916 (0.101)	1.099 (0.112)	1.010 (0.089)	192	2.209 (0.230)	**2.077** (0.208)	2.4 (0.220)	2.435 (0.232)
100	200	0.469 (0.051)	**0.410** (0.029)	0.478 (0.036)	0.477 (0.036)	200	0.898 (0.106)	**0.752** (0.064)	0.875 (0.074)	0.847 (0.075)
250	200	0.189 (0.013)	**0.146** (0.011)	0.171 (0.013)	0.171 (0.013)	200	0.331 (0.026)	**0.266** (0.020)	0.312 (0.022)	0.303 (0.021)
500	200	0.131 (0.010)	**0.081** (0.006)	0.095 (0.007)	0.095 (0.007)	200	0.193 (0.013)	**0.137** (0.011)	0.165 (0.013)	0.161 (0.013)
1000	200	0.094 (0.005)	**0.036** (0.003)	0.041 (0.003)	0.041 (0.003)	200	0.146 (0.010)	**0.069** (0.005)	0.082 (0.007)	0.082 (0.007)
1500	200	0.079 (0.004)	**0.024** (0.002)	0.029 (0.002)	0.029 (0.002)	200	0.111 (0.006)	**0.043** (0.003)	0.047 (0.003)	0.046 (0.003)

N	Student t					Lognormal				
	Instances	OLS	GLR	SQR	MQR	Instances	OLS	GLR	SQR	MQR
50	177	2.080 (0.234)	2.117 (0.216)	1.889 (0.177)	**1.816** (0.182)	157	3.616 (0.414)	3.160 (0.472)	3.282 (0.356)	**2.989** (0.339)
100	191	1.079 (0.131)	1.123 (0.118)	0.938 (0.091)	**0.926** (0.094)	185	2.424 (0.295)	1.561 (0.154)	1.805 (0.207)	**1.373** (0.147)
250	198	0.474 (0.039)	0.430 (0.033)	**0.334** (0.023)	0.336 (0.024)	186	0.863 (0.154)	0.560 (0.046)	0.565 (0.089)	**0.446** (0.035)
500	197	0.274 (0.018)	0.282 (0.024)	0.181 (0.014)	**0.177** (0.014)	194	0.423 (0.086)	0.342 (0.031)	0.248 (0.023)	**0.245** (0.023)
1000	198	0.151 (0.01)	0.120 (0.009)	**0.072** (0.005)	**0.072** (0.005)	197	0.211 (0.016)	0.159 (0.013)	0.123 (0.011)	**0.122** (0.011)
1500	197	0.123 (0.007)	0.106 (0.014)	**0.052** (0.003)	0.053 (0.004)	198	0.177 (0.016)	0.141 (0.013)	**0.087** (0.007)	**0.087** (0.007)

N	Mixture				
	Instances	OLS	GLR	SQR	MQR
50	198	6.485 (0.321)	5.606 (0.301)	5.064 (0.301)	**4.967** (0.301)
100	200	4.822 (0.250)	3.636 (0.233)	3.464 (0.236)	**3.450** (0.233)
250	200	4.058 (0.185)	2.295 (0.128)	2.212 (0.129)	**2.210** (0.128)
500	200	3.244 (0.123)	1.378 (0.067)	1.273 (0.067)	**1.271** (0.067)
1000	200	3.298 (0.094)	1.036 (0.048)	**0.970** (0.048)	0.972 (0.048)
1500	200	3.279 (0.071)	0.748 (0.041)	0.687 (0.041)	**0.685** (0.041)

Table 3 Mean and standard error (in brackets) for the AE of the profit using the generating model G1. We boldface the method that has the best performance at each N.

N	Instances	OLS	GLR mean			MQR mean		
			GLR	SQR	MQR	GLR	SQR	MQR
50	189	2.967	2.059	2.139	2.264	2.062	2.082	**2.014**
		(0.329)	(0.238)	(0.241)	(0.271)	(0.240)	(0.253)	(0.235)
100	199	1.277	0.681	0.704	0.698	**0.648**	0.664	0.678
		(0.156)	(0.062)	(0.056)	(0.058)	(0.054)	(0.051)	(0.053)
250	200	0.738	0.300	0.235	0.231	0.301	**0.222**	0.225
		(0.065)	(0.020)	(0.017)	(0.017)	(0.021)	(0.016)	(0.016)
500	200	0.655	0.235	0.141	0.138	0.237	**0.136**	0.140
		(0.033)	(0.015)	(0.012)	(0.012)	(0.015)	(0.012)	(0.012)
1000	200	0.497	0.144	0.059	0.058	0.143	**0.056**	0.058
		(0.017)	(0.007)	(0.004)	(0.004)	(0.007)	(0.004)	(0.004)
1500	200	0.481	0.131	0.037	0.037	0.131	**0.036**	0.038
		(0.013)	(0.005)	(0.003)	(0.003)	(0.005)	(0.003)	(0.003)

Table 4 **Mean and standard error (in brackets) for the AE of the profit using the generating model G2. We boldface the method that has the best performance at each N.**

In Table 4 for generating model **G2**, the SQR and MQR methods with either mean estimation method tends to dominate over the GLR based method ($\sim 70\%$ relative to GLR when $N = 1500$). This is surprising because generating model **G2** only employs a Normal distribution of error and for generating model **G1** under a Normal distribution GLR method dominated over the MQR and SQR methods. It is possible here that the MQR mean estimation method here performs slightly better than the GLR mean estimation method as heteroskedasticity is implicitly model instead of an explicit location-scale model assumed by a GLR which in this case does not approximate the generating model accurately. We believe this is where quantile based methods like QR, MQR and SQR are powerful in comparison to GLR. For smaller data sets, no method statistically dominates another method although some methods seem to have smaller mean for the same standard errors.

We also note that across all the results that it is harder to distinguish the MQR and the SQR methods as their mean performances are very similar with near identical standard error levels. We do see that MQR tends to slightly outperform SQR for smaller sized data sets in generating model **G1** and SQR tends to slightly outperform in generating model **G2**.

To understand the impact of the degree of heteroskedasticity on the above results, we consider variation of the generating model **G1** at different levels of heteroskedasticity by introducing a scalar multiplier H to parameters γ_1 and γ_2 respectively. If $H = 0$, the setting is homoskedastic. The results above were presented when $H = 1$. Finally, if $H = 2$, we increase the heteroskedasticity even more; and to differentiate the settings between $H = 1$ and $H = 2$, we refer to them as low and high heteroskedasticty respectively. Table 5 presents the results of three methods in these settings when $N = 250,\ 500$ and 1000: the baseline OLS approach, the complete GLR method (GLR mean and GLR residuals for superquantile and quantile estimates) and the SQR method (MQR_0 for mean,

Distribution	N	Homoskedastic			Low heteroskedasticity			High heteroskedasticity		
		OLS	GLR	SQR	OLS	GLR	SQR	OLS	GLR	SQR
Normal	250	**0.289**	0.322	0.357	0.204	**0.149**	0.174	0.622	**0.019**	0.031
		(0.022)	(0.022)	(0.029)	(0.013)	(0.010)	(0.011)	(0.018)	(0.001)	(0.003)
	500	**0.135**	0.152	0.180	0.122	**0.077**	0.089	0.593	**0.009**	0.016
		(0.010)	(0.011)	(0.012)	(0.008)	(0.006)	(0.006)	(0.013)	(0.001)	(0.001)
	1000	**0.079**	0.089	0.105	0.089	**0.039**	0.049	0.598	**0.005**	0.007
		(0.006)	(0.007)	(0.008)	(0.005)	(0.003)	(0.003)	(0.009)	(0.000)	(0.000)
Gamma	250	0.807	**0.722**	0.818	0.365	**0.276**	0.337	0.587	**0.030**	0.042
		(0.090)	(0.059)	(0.062)	(0.063)	(0.022)	(0.029)	(0.021)	(0.002)	(0.003)
	500	**0.285**	0.299	0.323	0.213	**0.138**	0.173	0.560	**0.016**	0.024
		(0.023)	(0.023)	(0.024)	(0.016)	(0.010)	(0.013)	(0.015)	(0.001)	(0.002)
	1000	**0.145**	0.148	0.181	0.143	**0.070**	0.084	0.526	**0.008**	0.014
		(0.011)	(0.011)	(0.013)	(0.009)	(0.005)	(0.006)	(0.009)	(0.001)	(0.001)
Student t	250	**0.733**	1.073	0.834	0.431	0.381	**0.315**	0.765	0.049	**0.048**
		(0.076)	(0.095)	(0.087)	(0.042)	(0.025)	(0.022)	(0.025)	(0.004)	(0.003)
	500	0.460	0.624	**0.407**	0.223	0.229	**0.169**	0.704	0.027	**0.024**
		(0.044)	(0.053)	(0.032)	(0.015)	(0.016)	(0.012)	(0.019)	(0.002)	(0.002)
	1000	**0.195**	0.370	0.200	0.156	0.116	**0.074**	0.649	0.016	**0.012**
		(0.017)	(0.040)	(0.015)	(0.008)	(0.009)	(0.005)	(0.012)	(0.001)	(0.001)
Lognormal	250	2.260	1.385	**1.118**	0.887	0.598	**0.464**	0.479	0.078	**0.070**
		(0.278)	(0.104)	(0.095)	(0.131)	(0.051)	(0.040)	(0.023)	(0.007)	(0.006)
	500	0.891	0.930	**0.574**	0.482	0.318	**0.238**	0.416	**0.033**	0.037
		(0.126)	(0.080)	(0.043)	(0.089)	(0.028)	(0.021)	(0.015)	(0.003)	(0.003)
	1000	0.377	0.430	**0.283**	0.226	0.204	**0.126**	0.408	0.027	**0.022**
		(0.035)	(0.033)	(0.023)	(0.019)	(0.020)	(0.012)	(0.011)	(0.003)	(0.002)
Mixture	250	**4.027**	4.499	4.131	3.602	1.949	**1.801**	4.097	0.294	**0.291**
		(0.255)	(0.267)	(0.258)	(0.162)	(0.101)	(0.099)	(0.267)	(0.016)	(0.015)
	500	**2.677**	2.732	2.742	3.207	1.501	**1.305**	3.835	0.209	**0.200**
		(0.153)	(0.164)	(0.158)	(0.131)	(0.079)	(0.076)	(0.030)	(0.011)	(0.011)
	1000	1.690	1.860	**1.682**	3.370	1.083	**0.979**	3.862	0.148	**0.136**
		(0.091)	(0.104)	(0.099)	(0.090)	(0.051)	(0.051)	(0.021)	(0.008)	(0.008)

Table 5 Mean and standard error (in brackets) of the AE in profit using a modified generating model G1 at different heteroskedasticity levels. We boldface the method that has the best performance for each distribution under each heteroskedasticity level.

SQR for superquantile and QR for quantile). When the data has heteroskedasticity (H=1 or 2), our observation, similar to that noted earlier, is that the methods that model heteroskedasticity are particularly effective. Their performance relative to OLS, which ignores heteroskedasticity, shows significant gains ($\sim$ 50% and 90% when $H = 1$ and 2 respectively when $N = 1000$). On the other hand, in the case of homoskedasticity, the OLS method tends to outperform in most distributions (not surprisingly as it coincides with the assumptions of the generating model and also similar to that observed by Chun et al. (2012) for the superquantile estimation problem), but the best

heteroskedastic method is competitive relative to OLS ($<\sim 3\%$ degradation when $N = 1000$ except for the Normal distribution where it is $\sim 13\%$). Here, in the Lognormal distribution SQR has a notable improvement ($\sim 25\%$ relative to OLS when $N = 1000$).

We now summarize our observations and lessons learnt from the above experiments. Our results in the context of the price-setting newsvendor suggest that ignoring heteroskedasticity can lead to significant errors, and that amongst the methods that model heteroskedasticity the SQR and the MQR methods result in better solutions for a wide range of generating models over GLR. GLR performs best in Normal and Gamma distributions but if the error distributions are highly asymmetric or heavy tailed or bi-modal (even though symmetric) or possess heteroskedastic effects that cannot fully be explained by a variance predictor (i.e., quantiles that depend differently on the different covariates aside from the effects of the noise), a mean-variance model captured by GLR may not have the best performance. Sometimes, GLR can even fail to converge, more so in some distributions over others.

For quantile or superquantile estimations, unless one expects the error distributions to have a unimodal, symmetric, non-heavy-tailed or homoskedastic behavior (i.e., similar to the *Normal* distribution), our results suggest that quantile-based methods such as QR, MQR or SQR may be better. For the mean estimation, quantile-based methods such as MQR may be preferred to GLR whenever the quantiles have different dependence on the covariates (aside from the effects of the noise distribution). However, when this is not the case, as in generating model **G1**, methods like GLR can outperform MQR and we believe this stems from the BLUE property of weighted least squares. Between MQR and SQR, SQR has some very interesting theoretical properties in terms of risk measure and being part of the risk quadrangle but statistically we are unable to distinguish MQR and SQR. Based on the above discussion, we gather that it is important the user of these techniques understands the data and uses the insights from the application area together with some of our conclusions to gauge the best technique that suits the data.

4.5. Non-price covariates, out-of-sample performance and regularization

We now consider an extension where the generating model **G1** has additional non-price covariates as described below, and empirically test the out-of-sample performance.

$$\textbf{G3.} \qquad Y = \beta_0 + \beta_1 p + \beta_3 T + \sum_{k=1}^{K} X_k + \left(\gamma_0 + \gamma_1 p + \gamma_2 p^2 + \gamma_3 T + \sum_{k=1}^{K} X_k\right) \epsilon, \tag{4.4}$$

where ϵ is a Normal (0,1) random variable, $\beta_0 = 200.0,\ \beta_1 = -35.0,\ \beta_3 = 10.0,\ \gamma_0 = 36.0,\ \gamma_1 = -12.0,\ \gamma_2 = 2.1$ and $\gamma_3 = 2.0$. The covariates T and $X_k, k = 1, \cdots, K$ are other independently and identically drawn covariates from their respective distributions besides price p. Specifically, we assume T is uniform between [-1,1] and half of the $X_k, k = 1, \cdots, K$ are Normal (0,1) and the other

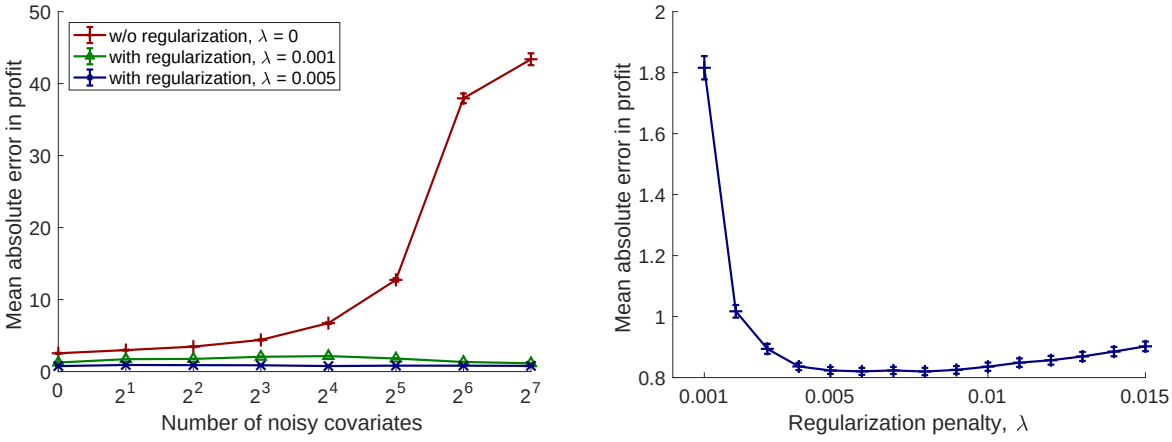

(a) (b)

Figure 5 **Mean and standard error (error bars) of the out-of-sample absolute error as a function of the noisy covariates K. Here, $N_{traning} = 100$, $N_{out-of-sample} = 20$ and number of instances $N_{mc} = 200$**

half are Bernoulli with 0 or 1 outcomes with probability 0.5. We increase the dimensionality of the model by varying K and measure the out-of-sample performance using the mean and standard deviation of the absolute error in profit as described in Eq. (4.3). Note here that the true profit and the realized profit depend on the out-of-sample covariates. The red line in Fig. 5(a) plots the mean and error bars using the mixed-quantile (or quantile) estimation method for the mean and CVaR (or VaR) over 200 instances, each with 20 out of sample observations when $N_{training} = 100$. We observe smaller error and standard deviation when the number of non-price covariates is small, and it is not surprising that it scales very quickly with the number of covariates due to overfitting. To manage overfitting at higher dimensionality, we use Lasso (l_1 norm) regularization, a standard feature selection technique, for each individual estimation method within our framework (other methods such as l_2 norm ridge penalty or elastic-net can also be used; see Ban and Rudin (2018) who motivate regularization in the context of data-driven estimations for inventory problems). The blue and green lines in Fig. 5(a) plot the out-of-sample performance for different penalty levels (we use the same penalty for mean, CVaR and VaR estimations and the results will only improve if we fine tune the penalty for each method). As we observe in the plot, regularization techniques are particularly useful in the big data settings with high-dimensional covariates, since they lead to stable regression estimates with better predictive power. Fig. 5(b) shows the variation in the error with the l_1 norm penalty when $K = 32$ and it indicates the potential opportunity in decrease in error with penalty tuning.

4.6. Run time comparisons

For the comparisons, we consider the heteroskedastic generating model **G1** with a unit Normal error distribution. The price optimization with the SQR method with 1000 observations converges

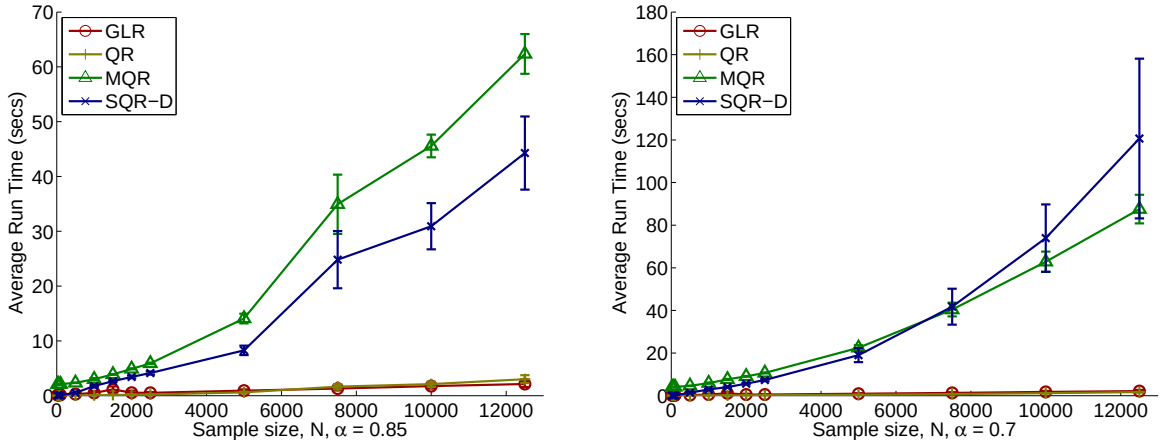

Figure 6 **Computational run times for GLR, QR, MQR and SQR-Decomposition as a function of N.**

on average in 15 iterations (correspondingly the CVaR evaluations at different α) and 20 seconds in total across all iterations. As the price optimization is a non-linear optimization problem, the number of iterations and in the lost sales settings, the number of CVaR estimations as well the α at which the CVaR is evaluated, varies across instances as well as choice of the CVaR estimator. For ease of comparison, we focus on the run-times of the individual procedures of GLR, QR, MQR and SQR at different α values as opposed to the overall time for solving the price-setting newsvendor problem. Fig. 6 shows the average computational time needed to solve the different estimation methods as a function of N for two different α values (i.e., 0.7 and 0.85). For each value of N and α, the plot shows the mean value of the run times and the corresponding error (standard deviation) bars estimated over 200 independently generated data sets, each of size N.

GLR as a practical technique for larger data sets is much faster and grows at a negligible rate compared to MQR and SQR at any quantile level, and even more so at quantiles closer to 0. Observe also that the standard deviation of this method is very small. The time taken for MQR in comparison to SQR is in the same order and a bit higher in many cases except for $\alpha = 0.7$ and large N. The run time of MQR can be tuned up or down by decreasing or increasing Δ which is currently set to 0.01. This reduces the number of discretizations over which the quantiles are estimated. The decomposition method of SQR on the other hand uses the finest discretization that can be generated with the residual data set. Observe also that the standard deviations of MQR and SQR to increase for larger N (and smaller α for SQR in particular). To estimate a single quantile QR is a fast practical routine with similar performance guarantees like GLR.

4.7. Bootstrap confidence intervals

An important aspect of a practical method for the price-setting newsvendor problem is obtaining confidence intervals for the estimated optimal solutions and the resulting estimated optimal profit

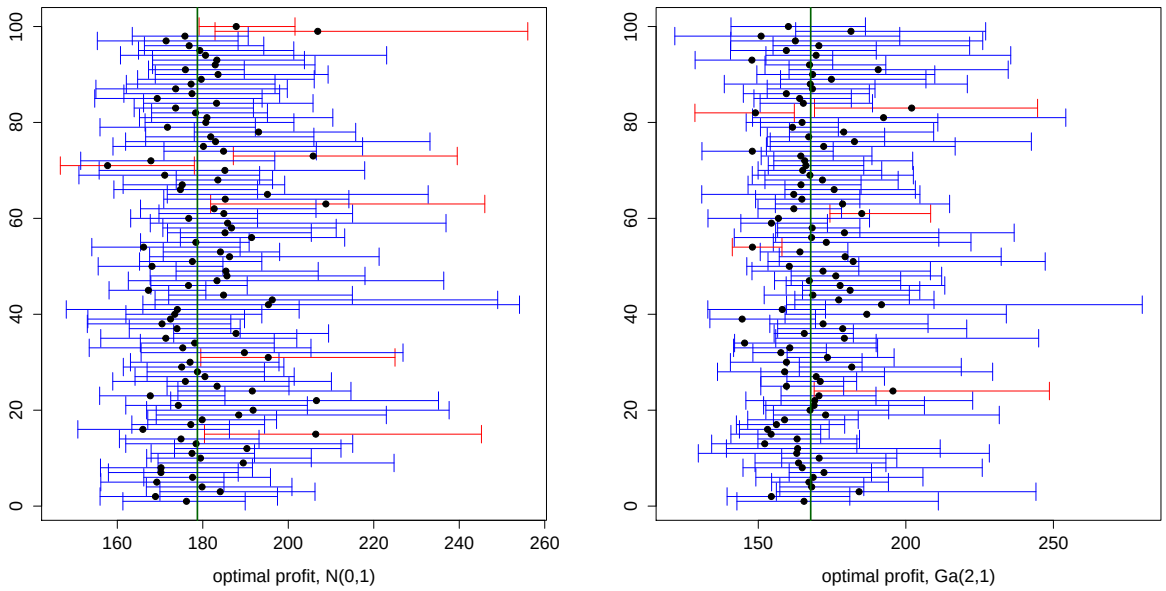

Figure 7 **Coverage results for the 95% bootstrap confidence intervals for the estimated optimal profit obtained using heteroskedastic regression for the sample size $N = 50$. The confidence intervals are shown along with the exact optimal profit in each case for $N_{\text{mc}} = 100$ Monte Carlo simulations (the intervals shown marked in red do not contain the exact optimal profit).**

$\hat{\Pi}^*$. These confidence intervals can clarify the impact of the model fit and estimation errors on the optimal solution. We propose to use the non-parametric bootstrap (Efron and Tibshirani 1994) to obtain the desired confidence intervals. In particular, we use the "paired bootstrap" whereby the bootstrap data sets are generated by sampling entire individual cases with replacement from the original data set. We implement our estimation procedure on each bootstrap data set to recover a bootstrap confidence interval for the quantity of interest for each original data set.

For an explicitly-known stochastic price-demand function, the coverage and accuracy of the resulting bootstrap confidence intervals can be evaluated through a Monte Carlo simulation study. Fig. 7 shows the results of the Monte Carlo evaluation of the 95% bootstrap confidence intervals of the estimated optimal profit $\hat{\Pi}^*$ over $N_{\text{mc}} = 100$ instances (original data sets) of model **G1** using GLR based estimators for mean, quantile and superquantile respectively. For each of the Monte Carlo instance, the 95% bootstrap confidence intervals are presented. The coverage of the bootstrap method is evaluated as the fraction of simulated Monte Carlo data sets for which the exact optimal profit is within the 95% bootstrap confidence intervals for the corresponding estimated optimal profit. Fig. 7 shows that these bootstrap confidence intervals are quite adequate.

5. An application of the framework with real data

We use the energy consumption data, a portion of which was depicted in Fig. 1, as an example to motivate that the approaches proposed in the paper are well suited to realistic data sets. The data consists of 2,313 observations in total and each observation corresponds to the average energy consumption in kWh per household recorded every 15 mins between 7-9am on weekdays from April 1, 2006 to March 31, 2007 along with a price and temperature measurement. We derive two variables related to the temperature covariate which we refer to as the cooling degree day ($CDD = [T-18]^+$) and the heating degree day ($HDD = [15-T]^+$), where T is the temperature. These are well understood temperature boundaries where energy consumption patterns change: in the former case households have air conditioners turned on, while in the latter case heaters are turned on. Using price, CDD and HDD as covariates of the consumed energy, we identify the optimal price, optimal order quantity and optimal predicted profit as a function of temperature for the emergency order formulation of the price-setting newsvendor problem with $\alpha_{eo} = 0.85$ and $c = 4$. Fig. 8 shows the bootstrap confidence intervals of the optimal price, the optimal order quantity and the optimal predicted profit as a function of temperature using two approaches that differ in the CVaR estimation procedure (SQR and MQR respectively). Both approaches use MQR and QR to get estimates of the mean and quantile respectively. Interestingly, both approaches result in fairly close values of optimal decisions and predicted profits, with negligible differences at temperatures higher than 5^oC. As opposed to constructing a multi-dimensional joint distribution of the demand as a function of the covariates to identify the optimal decision quantities, the key benefit of the proposed method is to perform the optimization in conjunction with the specific estimation techniques that characterize certain statistics of the demand response as a function of the covariates in a data-driven distribution-free manner for the prescriptive analytics.

Appendix A: Proof of Claim 1

The main part of our proof is to show that the estimate of the objective function in (2.5) uniformly converges in probability in $p \in P$. We begin with the estimators and then construct the case for the objective.

Suppose $Q\big(\mathbf{X}(p,\mathbf{z}),\boldsymbol{\beta}\big)$ is one of the statistical estimators of interest for a given $\alpha \in [0,1)$, where $\mathbf{X}(p,\mathbf{z})$ and $\boldsymbol{\beta}$ represent the covariates and the parameters of the data-driven estimator. (Note that the mean estimator does not depend on α, while CVaR and VaR do.) We denote the estimate of $Q\big(\mathbf{X}(p,\mathbf{z}),\boldsymbol{\beta}\big)$ with sample data of size N by the random variable $\hat{Q}\big(\mathbf{X}(p,\mathbf{z}),\boldsymbol{\beta}_N\big)$. As $Q\big(\mathbf{X}(p,\mathbf{z}),\boldsymbol{\beta}\big)$ is consistent, we know that given any $\epsilon > 0, \delta > 0$, $\exists\ N_o(\epsilon,\delta,\boldsymbol{\beta})$ such that $P\Big(|\hat{Q}\big(\mathbf{X}(p,\mathbf{z}),\boldsymbol{\beta}_N\big) - Q\big(\mathbf{X}(p,\mathbf{z}),\boldsymbol{\beta}\big)| > \epsilon\Big) < \delta,\ \forall\ N > N_o(\epsilon,\delta,\boldsymbol{\beta})$. Observe that $N_o(\epsilon,\delta,\boldsymbol{\beta})$ is independent of p because $\boldsymbol{\beta}$ is independent of the covariates $\mathbf{X}(p,\mathbf{z})$ and hence p as well. Therefore $Q\big(\mathbf{X}(p,\mathbf{z}),\boldsymbol{\beta}\big)$ converges uniformly in probability over $p \in P$.

We first consider the emergency order setting and prove its uniform convergence. We denote the objective function and its estimate by $\Pi(p)$ and $\hat{\Pi}(p)$ respectively. In this setting, α is a fixed constant, say α_o. We

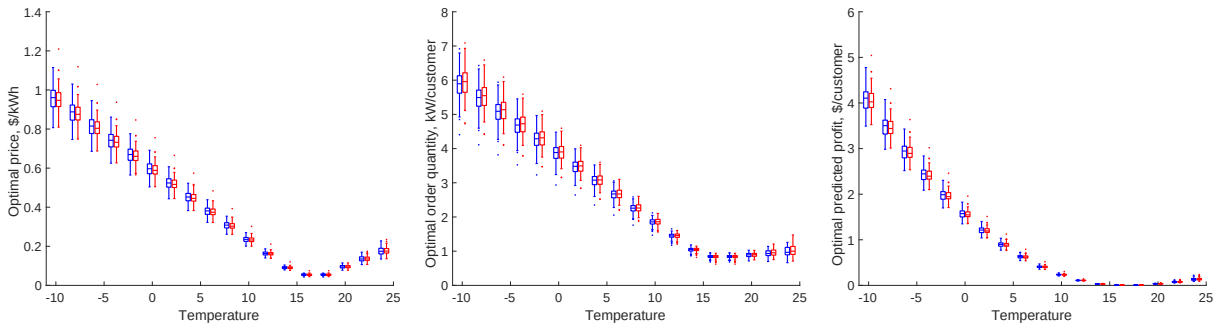

Figure 8 **Bootstrap confidence intervals for the energy consumption data as a function of temperature using two methods to estimate CVaR: SQR (blue, left) and MQR (red, right) respectively and the MQR method to estimate the mean. The box plots show median value (solid horizontal line) and 25-75th percentile values (boundary of boxes), with the whiskers extending to the extreme points excluding outliers marked with a dot.**

have suppressed the dependence on $\mathbf{z}$ for notational brevity. Denote $E(p), \hat{E}(p)$ and $C(p), \hat{C}(p)$ the true and estimated quantities of $E[D(p,\mathbf{z})]$ and $\text{CVaR}_{\alpha_o}[D(p,\mathbf{z})]$ respectively.

For a given ϵ, δ, we construct $N_o(\epsilon,\delta)$ for the convergence of $\Pi(p)$ argument as follows: $N_o(\epsilon,\delta) = \max\{N_o^E, N_o^C\}$ where N_o^E, N_o^C are the corresponding N_o's (see above) for the mean and the CVaR estimators with input parameters as follows: (1) $N_o^E = N_o\left(\epsilon_1, \delta, \beta^E\right)$ for $E[D(p,\mathbf{z})]$ estimator where $\epsilon_1 = \frac{\epsilon}{p_{\max}-c}$; and (2) $N_o^C = N_o\left(\epsilon_2, \delta, \beta^{C,\alpha_o}\right)$ for $\text{CVaR}_{\alpha_o}[D(p,\mathbf{z})]$ estimator where $\epsilon_2 = \frac{\epsilon}{c-s}$. Therefore, for any $p \in P$, we have

$$P\left(|\hat{\Pi}(p) - \Pi(p)| > \epsilon\right) = P\left(\left|(p-s)[\hat{E}(p) - E(p)] - (c-s)[\hat{C}(p) - C(p)]\right| > \epsilon\right) \tag{A.1}$$

$$\leq P\left(\left|(p-s)[\hat{E}(p) - E(p)]\right| > \epsilon\right) + P\left(\left|(c-s)[\hat{C}(p) - C(p)]\right| > \epsilon\right) \tag{A.2}$$

$$\leq P\left(\left|\hat{E}(p) - E(p)\right| > \epsilon_1\right) + P\left(\left|\hat{C}(p) - C(p)\right| > \epsilon_2\right) \tag{A.3}$$

$$< \delta \qquad \forall\ N > N_o(\epsilon,\delta) \tag{A.4}$$

As $N_o(\epsilon,\delta)$ is independent of p by by construction, we have the uniform convergence in probability of $\hat{\Pi}(p)$ to $\Pi(p)$ for all $p \in P$ in the emergency order setting.

We now consider the lost sales setting. Here α is a function of p denoted by $\alpha(p)$. Barring the new notations we need for the CVaR estimator that we describe here, we use the same notation as above. Denote $C(p), \hat{C}(p)$ the true and estimated quantity of $\text{CVaR}_{\alpha(p)}[D(p,\mathbf{z})]$. Define $N_o^C = N_o\left(\epsilon_2, \delta, \beta^{C,\alpha(p)}\right)$ for $\text{CVaR}_{\alpha(p)}[D(p,\mathbf{z})]$ estimator where $\epsilon_2 = \frac{\epsilon}{c-s}$. Because of our additional restrictive assumption in the lost-sales setting, we know that $N_o\left(\epsilon_2, \delta, \beta^{C,\alpha(p)}\right)$ is independent of p for all $p \in P$ and hence N_o^C as well. And with the same proof above, we have the uniform convergence in probability of $\Pi(p)$ for $p \in P$ even in the lost sales setting.

As the estimators for $E[D(p,\mathbf{z})]$ and $\text{CVaR}_{\alpha}[D(p,\mathbf{z})]$ are jointly continuous in p (and jointly in (p,α) in the lost-sales setting), it is easy to conclude that $\hat{\Pi}(p)$ is continuous in p.

Together with the condition on the unique global maximum for $\Pi(p)$, we meet all the conditions of consistency theorem of extremum estimators, Theorem 4.1.1 in Amemiya 1985, which implies that the estimated optimal solution p_N^* from (2.5) is consistent to the true optimal price p^{**}.

Next we show the consistency of the estimated x_N^* using (2.3) to the true optimal x^{**}. We want to show that given an $\epsilon > 0, \delta > 0$, $\exists\ N_o(\epsilon, \delta, \boldsymbol{\beta}, p^{**})$ such that $\forall\ N > N_o(\epsilon, \delta, \boldsymbol{\beta}, p^{**})$, we have $P\Big(|\hat{Q}\big(\mathbf{X}(p_N^*, \mathbf{z}), \boldsymbol{\beta}_N\big) - Q\big(\mathbf{X}(p^{**}, \mathbf{z}), \boldsymbol{\beta}\big)| > \epsilon\Big) < \delta$ where $Q(.), \hat{Q}(.)$ represents the the true VaR and its estimator. Now because $\text{VaR}_\alpha[D(p, \mathbf{z})]$ is a consistent estimator, we have for the same ϵ and δ, $\exists\ N_o(\epsilon, \frac{\delta}{2}, \boldsymbol{\beta})$ such that $\forall\ N > N_o(\epsilon, \frac{\delta}{2}, \boldsymbol{\beta})$, we have $P\Big(|\hat{Q}\big(\mathbf{X}(p^{**}, \mathbf{z}), \boldsymbol{\beta}_N\big) - Q\big(\mathbf{X}(p^{**}, \mathbf{z}), \boldsymbol{\beta}\big)| > \epsilon\Big) < \frac{\delta}{2}$. We also know that every estimate of VaR for a given set of observations is continuous in p. Therefore for the same ϵ, $\exists \eta > 0$ such that if $|p_N^* - p^{**}| < \eta$ implies $|\hat{Q}\big(\mathbf{X}(p_N^*, \mathbf{z}), \boldsymbol{\beta}_N\big) - \hat{Q}\big(\mathbf{X}(p^{**}, \mathbf{z}), \boldsymbol{\beta}_N\big)| < \epsilon$. But from consistency of price p_N^* to the true p^{**}, we know that for the same ϵ, η above, $\exists\ N_o(\eta, \frac{\delta}{2}, p^{**})$ such that $\forall\ N > N_o(\epsilon, \frac{\delta}{2}, p^{**})$, $P(|p_N^* - p^{**}| > \eta) < \frac{\delta}{2}$. This means $P(|\hat{Q}\big(\mathbf{X}(p_N^*, \mathbf{z}), \boldsymbol{\beta}_N\big) - \hat{Q}\big(\mathbf{X}(p^{**}, \mathbf{z}), \boldsymbol{\beta}_N\big)|) > \epsilon) < \frac{\delta}{2}$. Therefore we know that for $N > N_o(\epsilon, \delta, \boldsymbol{\beta}, p^{**}) = \max\big\{N_o(\epsilon, \frac{\delta}{2}, \boldsymbol{\beta}), N_o(\epsilon, \frac{\delta}{2}, p^{**})\big\}$, we have $P(|\hat{Q}\big(\mathbf{X}(p^*, \mathbf{z}), \boldsymbol{\beta}_N\big) - Q\big(\mathbf{X}(p^{**}, \mathbf{z}), \boldsymbol{\beta}\big)| > \epsilon) \leq P(|\hat{Q}\big(\mathbf{X}(p_N^*, \mathbf{z}), \boldsymbol{\beta}_N\big) - \hat{Q}\big(\mathbf{X}(p^{**}, \mathbf{z}), \boldsymbol{\beta}_N\big)|) > \epsilon) + P\Big(|\hat{Q}\big(\mathbf{X}(p^{**}, \mathbf{z}), \boldsymbol{\beta}_N\big) - Q\big(\mathbf{X}(p^{**}, \mathbf{z}), \boldsymbol{\beta}\big)| > \epsilon\Big) < \delta$. This proves the consistency of x_N^* to the true optimal x^{**}.

With the consistency results on p_N^* and x_N^* we conclude that the proposed method is asymptotically optimal.

□

Appendix B: Generalized Linear Regression (GLR)

The first approach is based on heteroskedastic regression using generalized linear models (GLM). To describe this approach in full generality, we need a specification of the mean-variance relationship and denoting the mean by μ, we assume the variance is $\phi V(\mu)$, where $V(\mu)$ is the variance function, and ϕ is the dispersion parameter. For example, a common specification for the variance function is $V(\mu) = \mu^\theta$ for some fixed θ. Regression models for μ and ϕ can be specified in the form of generalized linear models, and the model parameters can be estimated by a two stage iterative procedure where the inner stage uses GLM maximum quasi-likelihood optimization (i.e., iterative re-weighted least squares) to separately estimate the updates to the mean and the dispersion. For more information on this method we refer the reader to Nelder and Pregibon (1987) and Davidian and Carroll (1987, 1988) respectively. We note that method maximizes quasi likelihood retains the desirable properties of consistency, efficiency and asymptotic normality of the maximum likelihood estimates, even if the distribution of Y is not explicitly known (so long as this distribution has a finite second moment) (Wedderburn 1974).

Model Specification and Estimation. To fix ideas for our context, we consider the following generating model for the stochastic demand function:

$$Y = \mu + \sqrt{\phi V(\mu)}\epsilon, \tag{B.1}$$

where ϵ is a random variable whose distribution is independent of $\mathbf{X}$ with $E[\epsilon] = 0$ and $E[\epsilon^2] = 1$, with

$$g(\mu) = \boldsymbol{\beta}^T\mathbf{X}, \quad h(\phi) = \boldsymbol{\gamma}^T\mathbf{Z}, \tag{B.2}$$

where $\boldsymbol{\beta}, \boldsymbol{\gamma}$ are the respective regression parameters, and $g : \mathbb{R} \to \text{Range}(Y)$, and $h : \mathbb{R} \to \mathbb{R}^+$ are the respective *link functions* (that which provides the relationship between the linear predictor and the mean of the response variable) for μ and ϕ, and $\mathbf{X}$ and $\mathbf{Z}$ denote the respective set of covariates in the mean and dispersion models (Nelder and Lee 1992, Smyth et al. 2001).

The model in Eq. (B.1) and Eq. (B.2) is equivalent to the additive-multiplicative demand model in the inventory literature, and the covariates $\mathbf{X}$ and $\mathbf{Z}$ represent the demand drivers including price in this model. We outline the steps for estimating μ and ϕ in Procedure 4. This can be easily implemented using existing software for fitting GLM models, that require the specification of the response, covariates, variance function, dispersion parameter and link function.

Procedure 4 Heteroskedastic regression

Input: Data $\{\mathbf{x}_i, \mathbf{z}_i, y_i\}$ for $i = 1, ..., N$, the variance function $V(\mu)$, and the link functions $g(.)$ and $h(.)$ respectively.

1: Set the initial values for $\hat{\phi}_i$ for $i = 1, ..., N$.

2: Use GLM fitting to obtain the mean regression parameters $\hat{\boldsymbol{\beta}}$, using response y_i, covariates $\mathbf{x}_i$, variance function $V(\mu)$, dispersion $\hat{\phi}_i$, and link function $g(.)$. Set $\hat{\mu}_i = g^{-1}(\hat{\boldsymbol{\beta}}^T \mathbf{x}_i)$ and obtain the Pearson residuals $\hat{d}_i = \frac{(y_i - \hat{\mu}_i)^2}{V(\hat{\mu}_i)}$.

3: Use GLM fitting to obtain the dispersion regression parameters $\hat{\boldsymbol{\gamma}}$, using response $\hat{d}_i$, covariates $\mathbf{z_i}$, variance function ϕ^2, dispersion 2, and link function $h(.)$. Set $\phi_i = h^{-1}(\hat{\boldsymbol{\gamma}}^T \mathbf{z}_i)$.

4: Repeat from step 2 till $\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}$ converge.

Output: Estimates $\hat{\boldsymbol{\beta}}$ and $\hat{\boldsymbol{\gamma}}$ for heteroskedastic regression.

An estimate of the VaR and CVaR of ϵ can now be obtained from the empirical distribution of the adjusted residuals $\hat{\epsilon}$ where $\hat{\epsilon}_i = \frac{y_i - \hat{\mu}_i}{\sqrt{\hat{\phi}_i V(\hat{\mu}_i)}}$. We denote the empirical cdf $F_{\hat{\epsilon}}(u) = \frac{1}{N} \sum_{i=1}^{N} \mathrm{I}_{[\hat{\epsilon}_i]}(u)$, where $\mathrm{I}_{[\hat{\epsilon}_i]}(u)$ is the indicator function which takes the value 1 if $(u - \hat{\epsilon}_i) \geq 0$, and 0 otherwise. Then,

$$\widehat{\mathrm{VaR}}_\alpha[\hat{\epsilon}] = \inf\{u : F_{\hat{\epsilon}}(u) \geq \alpha\}, \tag{B.3}$$

$$\widehat{\mathrm{CVaR}}_\alpha[\hat{\epsilon}] = \lambda_\alpha(\hat{\epsilon}) \widehat{\mathrm{VaR}}_\alpha[\hat{\epsilon}] + (1 - \lambda_\alpha(\hat{\epsilon}))\, E\left[\hat{\epsilon} | \hat{\epsilon} > \widehat{\mathrm{VaR}}_\alpha[\hat{\epsilon}]\right], \tag{B.4}$$

$$\text{where } \lambda_\alpha(\hat{\epsilon}) = \frac{F_{\hat{\epsilon}}\left(\widehat{\mathrm{VaR}}_\alpha[\hat{\epsilon}]\right) - \alpha}{1 - \alpha}.$$

This description of CVaR for discrete distributions is given by Rockafellar and Uryasev (2002). The hat on VaR and CVaR are to denote that they are estimates.

In summary, the three quantities of interest for the price-setting newsvendor problem (see Fig. 3 and Section 2.2) are as follows:

$$\widehat{E}[Y|\mathbf{X}, \mathbf{Z}] = \hat{\mu}, \tag{B.5}$$

$$\widehat{\mathrm{VaR}}_\alpha[Y|\mathbf{X}, \mathbf{Z}] = \hat{\mu} + \sqrt{\hat{\phi} V(\hat{\mu})}\, \widehat{\mathrm{VaR}}_\alpha[\hat{\epsilon}], \tag{B.6}$$

$$\widehat{\mathrm{CVaR}}_\alpha[Y|\mathbf{X}, \mathbf{Z}] = \hat{\mu} + \sqrt{\hat{\phi} V(\hat{\mu})}\, \widehat{\mathrm{CVaR}}_\alpha[\hat{\epsilon}], \tag{B.7}$$

where $\hat{\mu} = g^{-1}(\hat{\boldsymbol{\beta}}^T \mathbf{X})$ and $\hat{\phi} = h^{-1}(\hat{\boldsymbol{\gamma}}^T \mathbf{Z})$. Here, $\hat{\boldsymbol{\beta}}$ and $\hat{\boldsymbol{\gamma}}$ are the outputs of estimation Procedure 4.

Appendix C: Proof of Claim 2

Consider the subset $J_k^* = \{i \in \mathcal{N} \mid y_i - \boldsymbol{\beta}^T \mathbf{x}_i - U_k > 0\}$. Then, we have

$$\sum_{i=1}^{N} \max\{y_i - \boldsymbol{\beta}^T \mathbf{x}_i - U_k, 0\} = \sum_{i \in J_k^*} \max\{y_i - \boldsymbol{\beta}^T \mathbf{x}_i - U_k, 0\} \tag{C.1}$$

since any index $i \notin J_k^*$ contributes zero to the summation on the left hand side. Similarly,

$$J_k^* \in \arg\max_{J_k \in \mathcal{P}(\mathcal{N})} \sum_{i \in J_k} (y_i - \boldsymbol{\beta}^T \mathbf{x}_i - U_k) \tag{C.2}$$

This is because any subset, say, $J_k \subset J_k^*$ can be augmented with elements from $J_k^* \setminus J_k$ to strictly increase the objective function, while no superset $J_k \supset J_k^*$ can possibly increase the objective function relative to J_k^* due to its definition. Taken together, Eqs. (C.1–C.2) lead to Eq. (3.12). □

Appendix D: Proof of Claim 3

We firstly note that constructing a finite, feasible solution for the corresponding dual LP is sufficient to establish boundedness of the above LP, due to weak duality. The corresponding dual LP is:

$$\max_{p,q} \quad \sum_{i=1}^{N} p_i y_i + \sum_{k=N_\alpha}^{N-1} \sum_{i=1}^{N} y_i q_k \tag{D.1}$$

$$\text{s.t.,} \quad q_k \leq \frac{a_k}{N(1-\alpha)}, \qquad \forall\ k = N_\alpha, \ldots, N-1, \tag{D.2}$$

$$q_k = \frac{\kappa_k - \kappa_{k-1}}{N(1-\alpha)}, \qquad \forall\ k = N_\alpha, \ldots, N-1, \tag{D.3}$$

$$\sum_{i=1}^{N} p_i = \frac{1}{N(1-\alpha)}, \tag{D.4}$$

$$\sum_{i=1}^{N} p_i x_{i,l} + \sum_{k=N_\alpha}^{N-1} \sum_{i=1}^{N} x_{i,l} q_k = \frac{1}{N} \sum_{i=1}^{N} x_{i,l}, \qquad \forall\ l = 1, \ldots, n, \tag{D.5}$$

$$p_i, q_k \geq 0, \qquad \forall\ i = 1, \ldots, N, \quad k = N_\alpha, \ldots, N-1. \tag{D.6}$$

Consider the candidate solution that evidently satisfies the non-negativity constraints, as well as constraints Eq. (D.3) and Eq. (D.4).

$$\tilde{p}_i = \frac{1}{N^2(1-\alpha)}, \quad \tilde{q}_k = \frac{\kappa_k - \kappa_{k-1}}{N(1-\alpha)} \tag{D.7}$$

Constraint Eq. (D.2) is satisfied because, using a series expansion for the natural logarithm (where each $|\kappa_k| < 1$), we have

$$a_k = \ln(1-\kappa_{k-1}) - \ln(1-\kappa_k) = \sum_{j=1}^{\infty} \frac{\kappa_k^j - \kappa_{k-1}^i}{j} > \kappa_k - \kappa_{k-1} \tag{D.8}$$

The final constraint Eq. (D.5) is also satisfied as verifiable via substitution, where $\forall\ l = 1, \ldots, n$,

$$\text{LHS} = \textstyle\sum_{i=1}^{N} x_{i,l} \left(\frac{1}{N^2(1-\alpha)} + \sum_{k=N_\alpha}^{N-1} \frac{\kappa_k - \kappa_{k-1}}{N(1-\alpha)} \right) = \sum_{i=1}^{N} x_{i,l} \left(\frac{1}{N^2(1-\alpha)} + \frac{1}{N(1-\alpha)} \left(\frac{N-1}{N} - \alpha \right) \right) = \text{RHS}.$$

We also note that feasibility of the above relaxed primal, namely $\mathbf{RSQR-RELAX}$ is self-evident. Taken together, these imply a finite, non-empty optimal solution for $\mathbf{RSQR-RELAX}$. □

Appendix E: Proof of Theorem 1

Step 2 is guaranteed to result in a finite, non-empty solution in the very first iteration, due to Claim 3 and thereby successfully seeds the delayed constraint-generation procedure. Finite convergence is guaranteed due to the finiteness of the power set, $\mathcal{P}(\mathcal{N})$. Convergence in the LP solution is achieved in step 3, when no new constraint is identifiable for each index k, i.e. J_k^* is already present in $\mathcal{J}_k$. Upon such convergence, it can be seen that the (final) converging linear program $\mathbf{RSQR-RELAX}$ is a relaxation of **RSQR** with respect to the representation of Eq. (3.15), but it also satisfies all the unrepresented constraints from **RSQR**. Thereby, its solution is also optimal for **RSQR**, and equivalently **SQR**. □

References

Agrawal V, Seshadri S (2000) Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing & Service Operations Management* 2(4):410–423.

Amemiya T (1985) *Advanced econometrics* (Harvard university press).

Andreev A, Kanto A, Malo P (2005) On closed-form calculation of CVaR. Technical report, Helsinki School of Economics.

Arikan E, Jammernegg W (2009) The newsvendor problem with a general price dependent demand distribution. *Available at SSRN 1523303* .

Aydin G, Porteus EL (2008) Joint inventory and pricing decisions for an assortment. *Operations Research* 56(5):1247–1255.

Ban GY, Rudin C (2018) The big data newsvendor: Practical insights from machine learning. *Operations Research* 67(1):90–108.

Bertsimas D, Kallus N (2016) The power and limits of predictive approaches to observational-data-driven optimization. *arXiv preprint arXiv:1605.02347* .

Bertsimas D, Vayanos P (2017) Data-driven learning in dynamic pricing using adaptive optimization. *Optimization Online* .

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

Beutel A, Minner S (2012) Safety stock planning under causal demand forecasting. *International Journal of Production Economics* 140(2):637–645.

Billinton R, Allan RN (1996) *Reliability evaluation of power systems* (Plenum Press).

Bookbinder JH, Lordahl AE (1989) Estimation of inventory re-order levels using the bootstrap statistical procedure. *IIE transactions* 21(4):302–312.

Brent RP (1973) *Algorithms for minimization without derivatives* (Courier Dover Publications).

Caro F, Gallien J (2012) Clearance pricing optimization for a fast-fashion retailer. *Operations Research* 60(6):1404–1422.

Chan LMA, Shen ZJM, Simchi-Levi D, Swann JL (2004) Coordination of pricing and inventory decisions: A survey and classification. *Series in Operations Research and Management Science* 335–392.

Chen X, Simchi-Levi D (2010) Pricing and inventory management. *The Handbook of Pricing Management, Oxford University Press, Oxford, UK* .

Chen Y, Xu M, Zhang ZG (2009) Technical notea risk-averse newsvendor model under the cvar criterion. *Operations Research* 57(4):1040–1044.

Chu LY, Shanthikumar JG, Shen ZM (2008) Solving operational statistics via a bayesian analysis. *Operations Research Letters* 36(1):110–116.

Chun SY, Shapiro A, Uryasev S (2012) Conditional value-at-risk and average value-at-risk: Estimation and asymptotics. *Operations Research* 60(4):739–756.

Cohen M, Lobel I, Paes Leme R (2016) Feature-based dynamic pricing. *Available at SSRN 2737045* .

Davidian M, Carroll RJ (1987) Variance function estimation. *Journal of the American Statistical Association* 82(400):1079–1091.

Davidian M, Carroll RJ (1988) A note on extended quasi-likelihood. *Journal of the Royal Statistical Society. Series B (Methodological)* 74–82.

Efron B, Tibshirani RJ (1994) *An introduction to the bootstrap*, volume 57 (CRC press).

Elmaghraby W, Keskinocak P (2003) Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science* 49(10):1287–1309.

Ettl M, Harsha P, Papush A, Perakis G (2019) A data-driven approach to personalized bundle pricing and recommendation. *Manufacturing & Service Operations Management* .

Federgruen A, Heching A (1999) Combined pricing and inventory control under uncertainty. *Operations Research* 47(3):454–475.

Ferreira KJ, Lee BHA, Simchi-Levi D (2016) Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* 18(1):69–88.

Gallego G, Moon I (1993) The distribution free newsboy problem: review and extensions. *Journal of the Operational Research Society* 825–834.

Godfrey GA, Powell WB (2001) An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Science* 47(8):1101–1112.

Hammerstrom DJ, et al. (2007) Pacific northwest gridwise testbed demonstration projects. *Part I. Olympic Peninsula Project* .

Harsha P, Sharma M, Natarajan R, Ghosh S (2013) An inventory theory framework for the analysis of probabilistic demand response schemes. *IEEE Transactions on Smart Grid* 4(4):2274–2284.

Harsha P, Subramanian S, Ettl M (2019) A practical price optimization approach for omnichannel retailing. *Informs Journal on Optimization* 1(3):241–264.

Hastie T, Tibshirani R, Friedman J (2001) The elements of statistical learning theory.

Huh WT, Levi R, Rusmevichientong P, Orlin JB (2011) Adaptive data-driven inventory control with censored demand based on kaplan-meier estimator. *Operations Research* 59(4):929–941.

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

Karlin S, Carr CR (1962) Prices and optimal inventory policy. *Studies in Applied Probability and Management Science* 159–172.

Kocabıyıkoğlu A, Popescu I (2011) An elasticity approach to the newsvendor with price-sensitive demand. *Operations Research* 59(2):301–312.

Koenker R, Bassett G (1978) Regression quantiles. *Econometrica* 46(1):33–50.

Koenker R, Bassett Jr G (1982) Robust tests for heteroscedasticity based on regression quantiles. *Econometrica: Journal of the Econometric Society* 43–61.

Lau AHL, Lau HS (1988) The newsboy problem with price-dependent demand distribution. *IIE transactions* 20(2):168–175.

Leorato S, Peracchi F, Tanase AV (2012) Asymptotically efficient estimation of the conditional expected shortfall. *Computational Statistics & Data Analysis* 56(4):768–784.

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

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

Liyanage LH, Shanthikumar JG (2005) A practical inventory control policy using operational statistics. *Operations Research Letters* 341–348.

Mills ES (1959) Uncertainty and price theory. *The Quarterly Journal of Economics* 73(1):116–130.

Nadarajah S, Zhang B, Chan S (2014) Estimation methods for expected shortfall. *Quantitative Finance* 2(14):271–291.

Nelder JA, Lee Y (1992) Likelihood, quasi-likelihood and pseudolikelihood: some comparisons. *Journal of the Royal Statistical Society. Series B (Methodological)* 273–284.

Nelder JA, Pregibon D (1987) An extended quasi-likelihood function. *Biometrika* 74(2):221–232.

Özer Ö, Phillips R (2012) *The Oxford handbook of pricing management* (Oxford University Press).

Peracchi F, Tanase AV (2008) On estimating the conditional expected shortfall. *Applied Stochastic Models in Business and Industry* 24(5):471–493.

Perakis G, Roels G (2008) Regret in the newsvendor model with partial information. *Operations Research* 56(1):188–203.

Petruzzi NC, Dada M (1999) Pricing and the newsvendor problem: A review with extensions. *Operations Research* 47(2):183–194.

Phillips R, Simsek AS, VanRyzin G (2012) Endogeneity and price sensitivity in customized pricing. *Columbia University Center for Pricing and Revenue Management Working Paper* 4.

Qiang S, Bayati M (2016) Dynamic pricing with demand covariates. *Available at SSRN 2765257* .

Raz G, Porteus EL (2006) A fractiles perspective to the joint price/quantity newsvendor model. *Management science* 52(11):1764–1777.

Rockafellar RT, Royset JO, Miranda SI (2014) Superquantile regression with applications to buffered reliability, uncertainty quantification, and conditional value-at-risk. *European Journal of Operational Research* 234(1):140–154.

Rockafellar RT, Uryasev S (2000) Optimization of conditional value-at-risk. *Journal of Risk* 2:21–41.

Rockafellar RT, Uryasev S (2002) Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance* 26(7):1443–1471.

Rockafellar RT, Uryasev S (2013) The fundamental risk quadrangle in risk management, optimization and statistical estimation. *Surveys in Operations Research and Management Science* 18(1):33–53.

Rockafellar RT, Uryasev S, Zabarankin M (2008) Risk tuning with generalized linear regression. *Mathematics of Operations Research* 33:712–729.

Scarf H, Arrow KJ, Karlin S (1958) A min-max solution of an inventory problem. *Studies in the mathematical theory of inventory and production* 10:201–209.

Smyth GK, Huele AF, Verbyla AP (2001) Exact and approximate reml for heteroskedastic regression. *Statistical modelling* 1(3):161–175.

Wedderburn RWM (1974) Quasi-likelihood functions, generalized linear models, and the gauss-newton method. *Biometrika* 61(3):439–447.

Yano CA, Gilbert SM (2005) Coordinated pricing and production/procurement decisions: A review. *Managing Business Interfaces* 65–103.

Yao L, Chen Y, Yan H (2006) The newsvendor problem with pricing: Extensions. *International Journal of Management Science and Engineering Management* 1(1):3–16.

Young L (1978) Price, inventory and the structure of uncertain demand. *New Zealand Oper. Res* 6(1):157–177.

Zipkin PH (2000) *Foundations of inventory management*, volume 2 (McGraw-Hill New York).