QUANTILE OPTIMIZATION FOR HEAVY-TAILED DISTRIBUTIONS USING ASYMMETRIC SIGNUM FUNCTIONS

JAE HO KIM* AND WARREN B. POWELL†

Abstract. In this paper, we present a provably convergent algorithm for computing the quantile of a random variable that does not require storing all of the sample realizations. We then present an algorithm for optimizing the quantile of a random function which may be characterized by a heavy-tailed distribution where the expectation is not defined. The algorithm is illustrated in the context of electricity trading in the presence of storage, where electricity prices are known to be heavy-tailed with infinite variance.

Key words. quantile optimization, stochastic gradient, signum functions, energy systems

AMS subject classification. .

1. Introduction.

For the past several decades, academic research in stochastic control and optimization has been focused mainly on maximizing or minimizing the expectation of a random process ([12],[17],[19],[28],[29]). A separate line of research has focused on worst-case outcomes under the umbrella of robust optimization [3]. However, when we are dealing with a non-stationary and complex or heavy-tailed system, the expectation is difficult or impossible to compute and it is very often not meaningful to even try to approximate the expectation. In those cases, it is rational to optimize our policy based on quantiles, which are zeroth order statistics derived directly from the probability distribution. The cumulative distribution function (CDF) and the inverse CDF, also known as the quantile function, can almost always be computed in a practical manner. Optimizing expectations is especially difficult in various financial markets including equities, commodities, and electricity markets, many of which are notorious for being non-stationary and extremely volatile and heavy-tailed ([2], [4], [5], [8]). For example, daily volatilities of 20-30% are common in electricity markets while the equity market also exhibits occasional heavy-tailed crashes as happened recently during the 1987 meltdown, the LTCM crisis in 1998, the dot-com bubble burst in 2001, and of course, during the 2008 financial meltdown.

A heavy-tailed distribution is defined by its structure of the decline in probabilities for large deviations. In a heavy-tailed environment, the usual statistical tools at our disposal can be tricked into producing erroneous results from observations of data in a finite sample and jumping to wrong conclusions. For example, suppose we want to compute the variance from a finite number of i.i.d. samples whose underlying

*email: jack@princeton.edu

†email: powell@princeton.edu / Author is supported in part by a grant from SAP and AFOSR contract FA9550-08-1-0195.

distribution is heavy-tailed such that the second and higher-order moments are not finite. If we use the standard empirical formula for computing the variance, we will get “a number,” thus giving the illusion of finiteness, but they typically show a high level of instability and grows steadily as we add more and more samples [35]. In this paper, we present a framework for optimizing the median or any other quantile of a stochastic process and show how it can be applied to optimizing a trading policy in the electricity markets. Since the quantile function always exists and is relatively stable even in a highly volatile and non-stationary environment, optimizing the quantile is a robust procedure, and almost always a reasonable thing to do in a complex, volatile, and heavy-tailed environment or in a situation where the worst-case scenario is unknown or too extreme.

For the past decade, many countries have deregulated their electricity markets. Electricity prices in deregulated markets are known to be very volatile and as a result participants in those markets face large risks. Recently, there has been a growing interest among electricity market participants about the possibility of using storage devices to mitigate the effect of the volatilities in electricity prices. However, storage devices have capacity limits and significant conversion losses, and it is unclear how one should use a storage device to maximize its value. Fortunately, unlike stock prices, it is well-known that electricity prices in deregulated markets exhibit reversion to the long-term “average,” allowing us to make directional bets ([6], [7], [22], [27], [32]). The value of a storage device is determined by how good we are at making such bets - buying electricity when we believe the price will go up and selling electricity when we believe the price will go down. When using a storage device, we automatically lose 10-30% of the energy through conversion losses. The conversion loss can be seen as a transaction cost, and hence the storage is valuable only if the price process is volatile enough so that we can sell at a price that is sufficiently higher than the price we buy to make up for the loss due to conversion.

The statistics community has long recognized that in the context of linear regression, minimizing the mean-squared or absolute-value error is an arbitrary choice. For certain applications, minimizing some quantile of the error terms instead can be a viable alternative. This approach is generally referred to as quantile regression, which minimizes the empirical expectation of a loss function on a given finite sample set via deterministic optimization. A comprehensive review of quantile regression and the relevant literature can be found in [18]. Meanwhile, the signal processing community has recognized that when the data is non-stationary and highly volatile, applying median filters [1] using order statistics yields better results than applying traditional filters based on minimizing empirical expectations.

The method described in this paper is different from the ones described in [18] or [1]. Instead of applying order statistics and deterministic algorithms on a given data set, we use an adaptation of the recursive algorithm pioneered by Robbins and Monro [31]. In addition, while algorithms for estimating quantiles require storing the history of observations, our adaptation performs quantile optimization without this requirement, making our implementation much more compact. The main contribution of this paper is the development of an algorithm which avoids the need to store the history of our process, as well as the requirement that an expectation exists, by replacing the stochastic gradient with the asymmetric signum function. Unlike the stochastic approximation algorithms that are derived from Robbins and Monro [31], we do not assume that the expectation of the random variable or random function exists. Our algorithm is applicable to heavy-tailed random variables and random functions whose mean and variance are not necessarily well-defined.

This paper is organized as follows. In §2, we present a provably convergent algorithm that allows us to find any quantile of a random variable. In §3, we present a framework for optimizing the quantile of a random function. In §4, we apply our quantile optimization method to trading electricity in the spot market in the presence of storage. In §5, we compare the trading strategy for maximum profit based on quantiles to the standard trading policy based on the hour of the day. In §6, we summarize our conclusions.

2. Computing the Quantile of a Random Variable. In this section, we provide a simple, provably convergent algorithm that allows us to compute the quantile of a random variable without storing the history of observations. Assume $X_1, X_2, ...$ are i.i.d. and continuous random variables where X_n is $\mathcal{F}_n$-measurable,

$$-\infty < X_n < \infty,$$

and

$$q_\alpha := \inf \left\{ b \in \mathbb{R} : \mathbb{P}\left[X_i \leq b\right] \geq \alpha \right\}, \ \forall n \geq 1,$$

for some $\alpha \in (0,1)$. We do not assume that $\mathbb{E}\left[X_n\right]$ exists. For example, $(X_n)_{n>1}$ can be i.i.d. Cauchy random variables. Given $\mathcal{F}_n$, the most common and straightforward method for approximating q_α is to sort $(X_i)_{1\leq i\leq n}$ in increasing order and pick the $n\alpha^{th}$ smallest number. However, in order to find the exact quantile using this method, we need an infinite amount of memory to store all $(X_i)_{1\leq i\leq n}$ as $n \to \infty$, which is not practical. We present an algorithm that only requires us to store one estimate of the quantile at any given time and show that it converges to the true quantile as $n \to \infty$.

Theorem 2.1. *Let $Y_0 \in \mathbb{R}$ be some finite number and*

$$Y_n = Y_{n-1} - \gamma_{n-1} sgn_\alpha \left(Y_{n-1} - X_n\right),$$

where

$$sgn_\alpha(u) = \begin{cases} 1-\alpha, & \textit{if} \quad u \geq 0, \\ -\alpha, & \textit{if} \quad u < 0, \end{cases}$$

is the asymmetric signum function and the stochastic stepsize $\gamma_n \geq 0$ *satisfies*

$$\sum_{n=0}^{\infty} \gamma_n = \infty \;\; a.s.$$

and

$$\sum_{n=0}^{\infty} (\gamma_n)^2 < \infty \;\; a.s.$$

Then,

$$\lim_{n\to\infty} Y_n \overset{a.s.}{=} q_\alpha.$$

Proof: We know that Y_n is $\mathcal{F}_n$-measurable and

$$(Y_n - q_\alpha)^2 = (Y_{n-1} - q_\alpha)^2 - 2\gamma_{n-1}(Y_{n-1} - q_\alpha) \cdot sgn(Y_{n-1} - X_n) + (\gamma_{n-1})^2 \lambda_n,$$

where

$$\lambda_n := sgn_\alpha^2 (Y_{n-1} - X_n)$$

and hence $0 \leq \lambda_n \leq 1$. Since

$$\begin{aligned} \mathbb{E}[sgn_\alpha(Y_{n-1} - X_n) \mid \mathcal{F}_{n-1}] &= (1-\alpha)\mathbb{P}[X_n \leq Y_{n-1} \mid \mathcal{F}_{n-1}] - \alpha\mathbb{P}[X_n > Y_{n-1} \mid \mathcal{F}_{n-1}] \\ &= (1-\alpha)\mathbb{P}[X_n \leq Y_{n-1} \mid \mathcal{F}_{n-1}] - \alpha(1 - \mathbb{P}[X_n \leq Y_{n-1} \mid \mathcal{F}_{n-1}]) \\ &= \mathbb{P}[X_n \leq Y_{n-1} \mid \mathcal{F}_{n-1}] - \alpha \\ &\leq 0 \text{ if and only if } Y_{n-1} \leq q_\alpha, \end{aligned}$$

we know that

$$\mathbb{E}[(Y_{n-1} - q_\alpha)\, sgn_\alpha(Y_{n-1} - q_\alpha) \mid \mathcal{F}_{n-1}] \geq 0. \tag{2.1}$$

Therefore,

$$\mathbb{E}\left[(Y_n - q_\alpha)^2 \mid \mathcal{F}_{n-1}\right] \leq (Y_{n-1} - q_\alpha)^2 + (\gamma_{n-1})^2, \;\forall n \geq 1. \tag{2.2}$$

Next, define

$$W_n := (Y_n - q_\alpha)^2 + \sum_{m=n}^{\infty} (\gamma_m)^2, \;\forall n \in \mathbb{N}_+. \tag{2.3}$$

Then, from (2.2),

$$\mathbb{E}\left[W_{n+1} \mid \mathcal{F}_n\right] \le W_n,\ \forall n \in \mathbb{N}_+,$$

showing that $(W_n)_{n\ge 1}$ is a positive supermartingale. Then, by the martingale convergence theorem,

$$\lim_{n\to\infty} (Y_n - q_\alpha)^2 = \lim_{n\to\infty} W_n \overset{a.s.}{=} W_\infty, \tag{2.4}$$

where

$$0 \le W_\infty < \infty.$$

Moreover, since

$$0 \le \mathbb{E}\left[W_n \mid \mathcal{F}_0\right] \le W_0 < \infty,\ \forall n \in \mathbb{N}_+,$$

by the dominated convergence theorem,

$$\mathbb{E}\left[W_\infty \mid \mathcal{F}_0\right] = \mathbb{E}\left[\lim_{N\to\infty} W_N \mid \mathcal{F}_0\right] = \lim_{N\to\infty} \mathbb{E}\left[W_N \mid \mathcal{F}_0\right]. \tag{2.5}$$

Define

$$\beta_{n-1,n} := \mathbb{E}\left[(Y_{n-1} - q_\alpha) \cdot sgn_\alpha\left(Y_{n-1} - X_n\right) \mid \mathcal{F}_{n-1}\right].$$

From (2.1), $\beta_{n-1,n} \ge 0$, and hence

$$\begin{aligned}
\beta_{0,n} &= \mathbb{E}\left[(Y_{n-1} - q_\alpha) \cdot sgn_\alpha\left(Y_{n-1} - X_n\right) \mid \mathcal{F}_0\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[(Y_{n-1} - q_\alpha) \cdot sgn_\alpha\left(Y_{n-1} - X_n\right) \mid \mathcal{F}_{n-1}\right] \mid \mathcal{F}_0\right] \\
&= \mathbb{E}\left[\beta_{n-1,n} \mid \mathcal{F}_0\right] \ge 0.
\end{aligned}$$

Next, from (2.3),

$$\begin{aligned}
W_N - (Y_0 - q_\alpha)^2 &= (Y_N - q_\alpha)^2 - (Y_0 - q_\alpha)^2 + \sum_{m=N}^{\infty} (\gamma_m)^2 \\
&= \sum_{n=1}^{N} \left\{(Y_n - q_\alpha)^2 - (Y_{n-1} - q_\alpha)^2\right\} + \sum_{m=N}^{\infty} (\gamma_m)^2 \\
&= -2 \sum_{n=1}^{N} \gamma_{n-1} \left(Y_{n-1} - q_\alpha\right) \cdot sgn_\alpha\left(Y_{n-1} - X_n\right) + \sum_{n=1}^{N} \left(\gamma_{n-1}\right)^2 \lambda_n + \sum_{m=N}^{\infty} (\gamma_m)^2.
\end{aligned}$$

Therefore, from (2.5),

$$\begin{aligned}
\mathbb{E}\left[W_\infty \mid \mathcal{F}_0\right] - \left(Y_0 - q_\alpha\right)^2 &= \lim_{N\to\infty} \mathbb{E}\left[W_N - \left(Y_0 - q_\alpha\right)^2 \mid \mathcal{F}_0\right] \\
&= \sum_{n=1}^{N} \left(\gamma_{n-1}\right)^2 \lambda_n + \sum_{m=N}^{\infty} \left(\gamma_m\right)^2 \\
&\quad -2 \lim_{N\to\infty} \sum_{n=1}^{N} \gamma_{n-1} \mathbb{E}\left[\left(Y_{n-1} - q_\alpha\right) \cdot sgn_\alpha\left(Y_{n-1} - X_n\right) \mid \mathcal{F}_0\right] \\
&= \sum_{n=1}^{N} \left(\gamma_{n-1}\right)^2 \lambda_n + \sum_{m=N}^{\infty} \left(\gamma_m\right)^2 - 2 \sum_{n=1}^{\infty} \gamma_{n-1} \beta_{0,n}.
\end{aligned}$$

Since the left-hand side is finite and

$$\sum_{n=1}^{N} \left(\gamma_{n-1}\right)^2 \lambda_n + \sum_{m=N}^{\infty} \left(\gamma_m\right)^2 \leq \sum_{n=1}^{\infty} \left(\gamma_m\right)^2 < \infty,$$

we must have

$$\sum_{n=1}^{\infty} \gamma_{n-1} \beta_{0,n} < \infty.$$

Since $\beta_{0,n} \geq 0$ and

$$\sum_{n=1}^{\infty} \gamma_{n-1} > \infty,$$

we must have

$$\lim_{n\to\infty} \beta_{0,n} = \lim_{n\to\infty} \mathbb{E}\left[\beta_{n-1,n} \mid \mathcal{F}_0\right] = 0.$$

However, since $\beta_{n-1,n}$ is a non-negative random variable, the expectation of it is zero if and only if $\beta_{n-1,n}$ goes to 0 with probability 1 as $n \to \infty$. In other words,

$$\lim_{n\to\infty} \beta_{n-1,n} = \lim_{n\to\infty} \left(Y_{n-1} - q_\alpha\right) \mathbb{E}\left[sgn_\alpha\left(Y_{n-1} - X_n\right) \mid \mathcal{F}_{n-1}\right] \overset{a.s.}{=} 0.$$

Since

$$\mathbb{E}\left[sgn_\alpha\left(Y_{n-1} - X_n\right) \mid \mathcal{F}_{n-1}\right] = 0$$

if and only if $Y_{n-1} = q_\alpha$,

$$\lim_{n\to\infty} Y_n \overset{a.s.}{=} q_\alpha.$$

$\square$.

3. Optimizing the Quantile of a Random Function. In this section, we show how to optimize the quantile of a sub-differentiable function. We show that a monotonicity condition in a random function and its derivative with respect to a sample realization allows us to optimize the quantile of the random function using a simple, provably convergent algorithm. We also show that the well-known newsvendor problem has the aforementioned monotonicity.

3.1. Monotonicity Assumptions. Let Ω be the set of all possible outcomes and let $\Theta \subseteq \mathbb{R}$ be a compact control space. We start by defining our control variable $\theta \in \Theta$ as a scalar for a clear and concise presentation of our concepts and the proof of convergence for our algorithm. Then, we generalize our process to the vector-valued case. Let $F : \Theta \times \Omega \mapsto \mathbb{R}$ be a function such that $F(\theta, \omega)$ is convex and continuous with respect to $\theta \in \Theta$ for all fixed $\omega \in \Omega$. We also assume that $F(\theta, \omega)$ is differentiable with respect to $\theta \in \Theta$ almost everywhere. For example, $F(\theta, \omega)$ may be piecewise-linear with respect to $\theta \in \Theta$ with a finite number of break-points for all fixed $\omega \in \Omega$. Let $F(\theta)$ denote the random variable whose specific sample realization is $F(\theta, \omega)$. Also, let Q_α denote the operator such that for any random variable X,

$$Q_\alpha X = \inf \{ b \in \mathbb{R} : \mathbb{P}[X \leq b] \geq \alpha \}$$

gives the α-quantile of the random variable X. The following propositions show the monotonicity assumptions that allow us to interchange the subgradient $\frac{\partial}{\partial\theta}$ and the quantile operator Q_α.

Proposition 3.1. *Assume*

$$F(\theta, \omega_1) \leq F(\theta, \omega_2) \text{ if and only if } \frac{\partial}{\partial\theta} F(\theta, \omega_1) \leq \frac{\partial}{\partial\theta} F(\theta, \omega_2), \tag{3.1}$$

$\forall \omega_1, \omega_2 \in \Omega$ *and* $\forall \theta \in \Theta$. *Then,*

$$Q_\alpha \frac{\partial}{\partial\theta} F(\theta) = \frac{\partial}{\partial\theta} Q_\alpha F(\theta), \ \forall \theta \in \Theta.$$

Proposition 3.2. *Assume*

$$F(\theta, \omega_1) \leq F(\theta, \omega_2) \text{ if and only if } \frac{\partial}{\partial\theta} F(\theta, \omega_1) \geq \frac{\partial}{\partial\theta} F(\theta, \omega_2), \tag{3.2}$$

$\forall \omega_1, \omega_2 \in \Omega$ *and* $\forall \theta \in \Theta$. *Then,*

$$Q_{1-\alpha} \frac{\partial}{\partial\theta} F(\theta) = \frac{\partial}{\partial\theta} Q_\alpha F(\theta), \ \forall \theta \in \Theta.$$

Proof: For some fixed $\theta \in \Theta$, let $\omega^* \in \Omega$ be the sample realization such that

$$Q_\alpha F(\theta) = F(\theta, \omega^*),$$

and let

$$\mathcal{A} := \{ \omega \in \Omega \mid F(\theta, \omega) \leq F(\theta, \omega^*) \}$$

and

$$\mathcal{B} := \{ \omega \in \Omega \mid F(\theta, \omega) > F(\theta, \omega^*) \}.$$

Then, if (3.1) holds,

$$\frac{\partial}{\partial\theta}F(\theta,\omega_1) \leq \frac{\partial}{\partial\theta}F(\theta,\omega^*) < \frac{\partial}{\partial\theta}F(\theta,\omega_2),\ \forall\omega_1 \in \mathcal{A}, \omega_2 \in \mathcal{B},$$

and hence

$$Q_\alpha \frac{\partial}{\partial\theta}F(\theta) = \frac{\partial}{\partial\theta}F(\theta,\omega^*) = \frac{\partial}{\partial\theta}Q_\alpha F(\theta).$$

On the other hand, if (3.2) holds,

$$\frac{\partial}{\partial\theta}F(\theta,\omega_1) \geq \frac{\partial}{\partial\theta}F(\theta,\omega^*) > \frac{\partial}{\partial\theta}F(\theta,\omega_2),\ \forall\omega_1 \in \mathcal{A}, \omega_2 \in \mathcal{B},$$

and hence

$$Q_{1-\alpha} \frac{\partial}{\partial\theta}F(\theta) = \frac{\partial}{\partial\theta}F(\theta,\omega^*) = \frac{\partial}{\partial\theta}Q_\alpha F(\theta).$$

□

We can illustrate **Proposition 3.2** through the following example. Suppose $\Theta = [0,1]$ and

$$F(\theta) = (1-\theta)U,\ \forall\theta \in \Theta,$$

where $U \sim \mathcal{U}(0,1)$ is a uniform random variable. Then, $\forall\omega_1, \omega_2 \in \Omega$ and $\forall\theta \in \Theta$,

$$F(\theta,\omega_1) = (1-\theta)U(\omega_1) \leq (1-\theta)U(\omega_2) = F(\theta,\omega_2)$$

if and only if $U(\omega_1) \leq U(\omega_2)$, and thus

$$\frac{\partial}{\partial\theta}F(\theta,\omega_1) = -U(\omega_1) \geq -U(\omega_2) = \frac{\partial}{\partial\theta}F(\theta,\omega_2),$$

satisfying (3.2). Next, $Q_\alpha F(\theta) = \alpha(1-\theta)$ and hence

$$\frac{\partial}{\partial\theta}Q_\alpha F(\theta) = -\alpha,\ \forall 0 < \alpha < 1.$$

On the other hand,

$$\frac{\partial}{\partial\theta}F(\theta) = -U \text{ and hence } Q_{1-\alpha}\frac{\partial}{\partial\theta}F(\theta) = -\alpha,$$

satisfying **Proposition 3.2.**

Note that the classical stochastic gradient algorithm that allows us to optimize the expectation of a function works because the expectation operator is interchangeable with the differential operator as long as the expectation is finite. We do not need to satisfy this condition when optimizing quantiles. Our goal is to develop an algorithm that is analogous to the classical stochastic gradient algorithm that allows us to optimize the quantile of a function.

3.2. Quantile Optimization. We begin by stating the algorithm for quantile optimization and proving its convergence.

THEOREM 3.3. *Define*

$$\theta^{\alpha} := \underset{\theta\in\Theta}{\operatorname{argmin}}\ Q_{\alpha}\ F(\theta).$$

For some $n \in \mathbb{N}_+$ and $\theta_{n-1} \in \Theta$, let $F_n(\theta_{n-1})$ be the $\mathcal{F}_n$-measurable random variable whose distribution is that of $F(\theta_{n-1})$. Denote

$$F_n'(\theta_{n-1}) := \frac{\partial}{\partial\theta} F_n(\theta,\omega)\,|_{\theta=\theta_{n-1}}.$$

Let $\theta_0 \in \Theta$ and

$$\theta_n = \begin{cases} \theta_{n-1} - \gamma_{n-1} sgn_{\alpha}\left(F_n'(\theta_{n-1})\right), & \text{if} \quad (3.1) \text{ holds for } F(\theta) \\ \theta_{n-1} - \gamma_{n-1} sgn_{1-\alpha}\left(F_n'(\theta_{n-1})\right), & \text{if} \quad (3.2) \text{ holds for } F(\theta) \end{cases}, \ \forall n \in \mathbb{N}_+,$$

where the stochastic stepsize $\gamma_n \geq 0$ satisfies

$$\sum_{n=0}^{\infty} \gamma_n = \infty \ a.s. \tag{3.3}$$

and

$$\sum_{n=0}^{\infty} (\gamma_n)^2 < \infty \ a.s. \tag{3.4}$$

Then,

$$\lim_{n\to\infty} \theta_n \stackrel{a.s.}{=} \theta^{\alpha}.$$

Proof: Assume (3.1) holds and let $\theta_0 \in \Theta$ and

$$\theta_n = \theta_{n-1} - \gamma_{n-1} sgn_{\alpha}\left(F_n'(\theta_{n-1})\right).$$

Then,

$$(\theta_n - \theta^{\alpha})^2 = (\theta_{n-1} - \theta^{\alpha})^2 - 2\gamma_{n-1}(\theta_{n-1} - \theta^{\alpha}) \cdot sgn_{\alpha}\left(F_n'(\theta_{n-1})\right) + \left(\gamma_{n-1}\right)^2 \lambda_n,$$

where

$$\lambda_n^{\alpha} := sgn_{\alpha}^2\left(F_n'(\theta_{n-1})\right).$$

From assumption (3.1), if $\theta_{n-1} \leq \theta^{\alpha}$,

$$Q_{\alpha} F_n'(\theta_{n-1}) = \frac{\partial}{\partial\theta} Q_{\alpha} F(\theta)\,|_{\theta=\theta_{n-1}} \leq 0$$

and if $\theta_{n-1} > \theta^{\alpha}$,

$$Q_{\alpha} F_n' (\theta_{n-1}) = \frac{\partial}{\partial \theta} Q_{\alpha} F(\theta) |_{\theta=\theta_{n-1}} > 0.$$

Therefore,

$$\begin{aligned}
\mathbb{E}\left[sgn_{\alpha}\left(F_n'(\theta_{n-1})\right) \mid \mathcal{F}_{n-1}\right] &= (1-\alpha)\,\mathbb{P}\left[F_n'(\theta_{n-1}) \geq 0 \mid \mathcal{F}_{n-1}\right] - \alpha \mathbb{P}\left[F_n'(\theta_{n-1}) < 0 \mid \mathcal{F}_{n-1}\right] \\
&= (1-\alpha)\,\mathbb{P}\left[F_n'(\theta_{n-1}) \geq 0 \mid \mathcal{F}_{n-1}\right] - \alpha\left(1 - \mathbb{P}\left[F_n'(\theta_{n-1}) \geq 0 \mid \mathcal{F}_{n-1}\right]\right) \\
&= \mathbb{P}\left[F_n'(\theta_{n-1}) \geq 0 \mid \mathcal{F}_{n-1}\right] - \alpha \\
&\leq 0 \text{ if and only if } \theta_{n-1} \leq \theta^{\alpha},
\end{aligned}$$

and hence

$$\mathbb{E}\left[(\theta_{n-1} - \theta^{\alpha}) \cdot sgn_{\alpha}\left(F_n'(\theta_{n-1})\right) \mid \mathcal{F}_{n-1}\right] \geq 0.$$

Then,

$$\mathbb{E}\left[(\theta_n - \theta^{\alpha})^2 \mid \mathcal{F}_{n-1}\right] \leq (\theta_{n-1} - \theta^{\alpha})^2 + \left(\gamma_{n-1}\right)^2, \ \forall n \geq 1.$$

Using the same proof technique used in section §2, we can show that

$$\lim_{n \to \infty} \theta_n \overset{a.s.}{=} \theta^{\alpha}.$$

Similarly, if assumption (3.2) holds instead, we can let $\theta_0 \in \Theta$ and

$$\theta_n = \theta_{n-1} - \gamma_{n-1} sgn_{1-\alpha}\left(F_n'(\theta_{n-1})\right).$$

Then,

$$(\theta_n - \theta^{\alpha})^2 = (\theta_{n-1} - \theta^{\alpha})^2 - 2\gamma_{n-1}(\theta_{n-1} - \theta^{\alpha}) \cdot sgn_{1-\alpha}\left(F_n'(\theta_{n-1})\right) + \left(\gamma_{n-1}\right)^2 \lambda_n,$$

where

$$\lambda_n^{1-\alpha} := sgn_{1-\alpha}^2\left(F_n'(\theta_{n-1})\right).$$

Under assumption (3.2), if $\theta_{n-1} \leq \theta^{\alpha}$,

$$Q_{1-\alpha} F_n'(\theta_{n-1}) = \frac{\partial}{\partial \theta} Q_{\alpha} F(\theta) |_{\theta=\theta_{n-1}} \leq 0$$

and if $\theta_{n-1} > \theta^{\alpha}$,

$$Q_{1-\alpha} F_n'(\theta_{n-1}) = \frac{\partial}{\partial \theta} Q_{\alpha} F(\theta) |_{\theta=\theta_{n-1}} > 0.$$

Therefore,

$$\begin{aligned}
\mathbb{E}\left[sgn_{1-\alpha}\left(F_n'\left(\theta_{n-1}\right)\right) \mid \mathcal{F}_{n-1}\right] &= \alpha\mathbb{P}\left[F_n'\left(\theta_{n-1}\right) \geq 0 \mid \mathcal{F}_{n-1}\right] - (1-\alpha)\,\mathbb{P}\left[F_n'\left(\theta_{n-1}\right) < 0 \mid \mathcal{F}_{n-1}\right] \\
&= \alpha\mathbb{P}\left[F_n'\left(\theta_{n-1}\right) \geq 0 \mid \mathcal{F}_{n-1}\right] - (1-\alpha)\left(1 - \mathbb{P}\left[F_n'\left(\theta_{n-1}\right) \geq 0 \mid \mathcal{F}_{n-1}\right]\right) \\
&= \mathbb{P}\left[F_n'\left(\theta_{n-1}\right) \geq 0 \mid \mathcal{F}_{n-1}\right] - (1-\alpha) \\
&\leq 0 \text{ if and only if } \theta_{n-1} \leq \theta^\alpha,
\end{aligned}$$

and hence

$$\mathbb{E}\left[\left(\theta_{n-1} - \theta^\alpha\right) \cdot sgn_{1-\alpha}\left(F_n'\left(\theta_{n-1}\right)\right) \mid \mathcal{F}_{n-1}\right] \geq 0.$$

Again,

$$\mathbb{E}\left[\left(\theta_n - \theta^\alpha\right)^2 \mid \mathcal{F}_{n-1}\right] \leq \left(\theta_{n-1} - \theta^\alpha\right)^2 + \left(\gamma_{n-1}\right)^2, \ \forall n \geq 1,$$

and we can show that

$$\lim_{n\to\infty} \theta_n \overset{a.s.}{=} \theta^\alpha.$$

□

While we have shown the proof for scalar θ, the algorithm can be readily extended to the multivariable case as long as the monotonicity assumption holds in each of the variables and corresponding partial derivatives. For a simple illustration, suppose $F\left(\overrightarrow{\theta}, \omega\right)$ is convex in $\overrightarrow{\theta} = \left(\theta^1, \theta^2, ..., \theta^m\right) \in \Theta$ for some m-dimensional compact space $\Theta \in \mathbb{R}^m$ and fixed $\omega \in \Omega$.

Theorem 3.4. *For each $1 \leq i \leq m$, assume that either*

$$\frac{\partial}{\partial\theta^i}F\left(\overrightarrow{\theta}, \omega_1\right) \leq \frac{\partial}{\partial\theta^i}F\left(\overrightarrow{\theta}, \omega_2\right) \text{ if and only if } F\left(\overrightarrow{\theta}, \omega_1\right) \leq F\left(\overrightarrow{\theta}, \omega_2\right) \tag{3.5}$$

holds true or

$$\frac{\partial}{\partial\theta^i}F\left(\overrightarrow{\theta}, \omega_1\right) \geq \frac{\partial}{\partial\theta^i}F\left(\overrightarrow{\theta}, \omega_2\right) \text{ if and only if } F\left(\overrightarrow{\theta}, \omega_1\right) \leq F\left(\overrightarrow{\theta}, \omega_2\right) \tag{3.6}$$

holds true, $\forall\omega_1, \omega_2 \in \Omega$ and $\forall\overrightarrow{\theta} \in \Theta$. For some $n \in \mathbb{N}_+$ and $\theta_{n-1} \in \Theta$, let $F_n\left(\overrightarrow{\theta}_{n-1}\right)$ be the $\mathcal{F}_n$-measurable random variable whose distribution is that of $F\left(\overrightarrow{\theta}_{n-1}\right)$. Then, if we let

$$\overrightarrow{\theta}_0 = \left(\theta_0^1, \theta_0^2, ..., \theta_0^m\right) \text{ for some } \overrightarrow{\theta}_0 \in \Theta$$

and

$$\theta_n^i = \begin{cases} \theta_{n-1}^i - \gamma_{n-1}^i sgn_\alpha\left(\frac{\partial}{\partial\theta^i}F_n\left(\overrightarrow{\theta}\right)|_{\overrightarrow{\theta} = \overrightarrow{\theta}_{n-1}}\right), & \text{if (3.5) holds true,} \\ \theta_{n-1}^i - \gamma_{n-1}^i sgn_{1-\alpha}\left(\frac{\partial}{\partial\theta^i}F_n\left(\overrightarrow{\theta}\right)|_{\overrightarrow{\theta} = \overrightarrow{\theta}_{n-1}}\right), & \text{if (3.6) holds true,} \end{cases}$$

while the stochastic stepsizes $\gamma_n^i \geq 0$ *satisfy*

$$\sum_{n=0}^{\infty} \gamma_n^i = \infty \ a.s.$$

and

$$\sum_{n=0}^{\infty} \left(\gamma_n^i\right)^2 < \infty \ a.s,$$

$\forall 1 \leq i \leq m$*, then*

$$\lim_{n\to\infty} \overrightarrow{\theta}_n \overset{a.s.}{=} \overrightarrow{\theta}^{\alpha}.$$

where

$$\overrightarrow{\theta}^{\alpha} = \underset{\overrightarrow{\theta} \in \Theta}{\operatorname{argmin}} \ Q_\alpha \ F\left(\overrightarrow{\theta}\right).$$

3.3. Initialization and the Scaling of Stepsizes. We have shown that the algorithm for finding some quantile of a random variable and optimizing some quantile of a random function converges as long as the stepsize $\gamma_n \geq 0$ satisfies (3.3) and (3.4). Extensive study on various stepsizes for optimizing the expectation of a random function has been done. A comprehensive review of the literature can be found in [9] and [28]. The typical approach for finding the optimal stepsizes in the case of optimizing expectations is as follows. Suppose $(X_i)_{1\leq i\leq n}$ are i.i.d random variables with finite variance and we want to find the best stepsizes $\left(\gamma_{i-1}\right)_{1\leq i\leq n}$ that gives the best estimate for the $\mathbb{E}\left[X_i\right]$. Then, the mean-squared error

$$\min_{\gamma_0,\ldots,\gamma_n} \mathbb{E}\left(\sum_{i=1}^{n} \gamma_{i-1} X_i - \mathbb{E}\left[X_i\right]\right)^2$$

is minimized when $\gamma_{n-1} = \frac{1}{n}$, $\forall n$. When $(X_i)_{1\leq i\leq n}$ are not i.i.d, different stepsizes can give better results.

However, when computing quantiles using asymmetric signum functions, it is difficult to determine what the optimal stepsizes should be even for the i.i.d case because it is unclear what we should try to minimize. We cannot minimize the expectation or variance in a heavy-tailed environment. Thus, instead of trying to identify stepsizes that are optimal in some sense, we present an ad-hoc method for practical use. Let $(X_i)_{1\leq i\leq n}$ be i.i.d random variables and assume we want to compute some α-quantile. However, suppose the α-quantile of X_i is 10^6, for example. Then, if we use the following algorithm

$$Y_n = Y_{n-1} - \gamma_{n-1} sgn_\alpha\left(Y_{n-1} - X_n\right),$$

with $Y_0 = 0$ and $\gamma_{n-1} = \frac{1}{n}$, it will take prohibitively long for the algorithm to give us a reasonable estimate. Thus, finding a reasonable initialization point Y_0 and stepsizes γ_{n-1} is critical. We propose initially taking sample realizations $(X_i)_{1\le i\le k}$ for some small k and sorting it in increasing order and pick the $k\alpha^{th}$ smallest number and let it be Y_0. Next, we let

$$\gamma_{n-1} := \frac{\sigma}{n} \tag{3.7}$$

where

$$\sigma := \frac{\left(.75\text{-quantile of the set } (X_i)_{1\le i\le k}\right) - \left(.25\text{-quantile of the set } (X_i)_{1\le i\le k}\right)}{2}$$

is a scaling factor.

For the quantile optimization problem, first pick $(\theta_m)_{1\le m\le M}$ evenly dividing the compact space Θ for some small M. For each θ_m, we can generate $(F_i(\theta_m))_{1\le i\le k}$ to find an approximation to the α-quantile of $F(\theta_m)$. Then, we let our initialization point θ_0 be the one that gives the minimum α-quantile. That is, initialization requires generating $M\times k$ samples. Next, we let the stepsize to be given by equation (3.7) where we replace X_i with $F_i(\theta_0)$. When optimizing expectations using the standard Robbins-Monro algorithm, the scaling factor has to strike a balance between the units of $\theta \in \Theta$ and the scale of the stochastic gradient, which can be highly volatile, ranging between large negative values and large positive values. However, when optimizing the quantile using our algorithm, the stepsizes have to be scaled based on θ alone. If we have a rough idea of where the optimal θ may fall, we can scale the stepsize accordingly.

3.4. Newsvendor Problem. To develop an understanding of the properties of the algorithm, we begin by illustrating it in the context of a simple newsvendor problem ([26]). Given a random demand $D \in \mathbb{R}_+$, we usually want to compute the optimal stocking quantity $\theta \in \mathbb{R}_+$ that maximizes the expected profit

$$G(\theta) := \mathbb{E}\left[p\min\{\theta, D\}\right] - c\theta$$

where $p > 0$ is the unit price, $c > 0$ is the unit cost, and $p > c$. For this simple problem, the optimal solution

$$\theta^* = Q_{\frac{p-c}{p}} D$$

is the $\frac{p-c}{p}$-quantile of the random variable D, which can be computed using the algorithm shown in section §2. However, instead of maximizing the expected profit, suppose a manager wants to maximize some α-quantile of the profit

$$q_\alpha := Q_\alpha\left[p\min\{\theta, D\} - c\theta\right].$$

This is equivalent to minimizing

$$q_{1-\alpha} := Q_{1-\alpha}\left[c\theta - p\min\{\theta, D\}\right].$$

Let Ω be the set of all possible outcomes. For each $\omega \in \Omega$, let $D(\omega)$ be the sample realization of the demand and denote

$$\begin{aligned} F(\theta, \omega) &:= c\theta - p \min\{\theta, D(\omega)\} \\ &= \begin{cases} \theta(c-p), & \text{if} \quad \theta \leq D(\omega), \\ c\theta - pD(\omega), & \text{if} \quad \theta > D(\omega). \end{cases} \end{aligned}$$

Then,

$$\frac{\partial}{\partial \theta} F(\theta, \omega) = \begin{cases} c - p, & \text{if} \quad \theta \leq D(\omega), \\ c, & \text{if} \quad \theta > D(\omega). \end{cases}$$

Next, $\forall \omega_1, \omega_2 \in \Omega$ and $\forall \theta \in \mathbb{R}_+$, we know that $F(\theta, \omega_1) \leq F(\theta, \omega_2)$ if and only if $D(\omega_1) \geq D(\omega_2)$, which again is true if and only if

$$\frac{\partial}{\partial \theta} F(\theta, \omega_1) \leq \frac{\partial}{\partial \theta} F(\theta, \omega_2).$$

Let $(D_n)_{n \geq 1}$ be i.i.d random variables. If we let $\theta_0 = 0$ and

$$\begin{aligned} \theta_n &= \theta_{n-1} - \frac{1}{n} sgn_\alpha \left(F_n'(\theta_{n-1})\right) \\ &= \begin{cases} \theta_{n-1} + \frac{\alpha}{n}, & \text{if} \quad \theta_{n-1} \leq D_n \\ \theta_{n-1} - \frac{1-\alpha}{n}, & \text{if} \quad \theta_{n-1} > D_n \end{cases} \end{aligned} \tag{3.8}$$

where D_n is $\mathcal{F}_n$-measurable, then

$$\lim_{n \to \infty} \theta_n \overset{a.s.}{=} \underset{\theta \in \mathbb{R}_+}{\operatorname{argmax}} Q_\alpha \left[p \min\{\theta, D\} - c\theta\right],$$

from **Theorem 3.3**. This implies

$$\underset{\theta \in \mathbb{R}_+}{\operatorname{argmax}} Q_\alpha \left[p \min\{\theta, D\} - c\theta\right] = Q_\alpha D.$$

For illustration, suppose $p = 10$, $c = 1$, and $(D_n)_{n \geq 1}$ is i.i.d with pareto distribution whose tail index is 1/2:

$$\mathbb{P}[D_n \leq y] = \begin{cases} 1 - \left(\frac{1}{y}\right)^{1/2}, & \text{if} \quad y \geq 1 \\ 0, & \text{if} \quad y < 1 \end{cases}$$

for all $n \geq 1$. For given $\theta \geq 1$, the expected profit is

$$\begin{aligned} \mathbb{E}[p \min\{\theta, D_n\}] - c\theta &= \mathbb{E}[10 \min\{\theta, D_n\}] - \theta \\ &= 5 \int_{y=1}^{\theta} y^{-1/2} dy + 10\theta \mathbb{P}[D_n > \theta] - \theta \\ &= 10 y^{1/2} \big|_{y=1}^{\theta} + 10\theta^{1/2} - \theta \\ &= 20\theta^{1/2} - 10 - \theta. \end{aligned}$$

Then, we obtain a maximum expected profit of 90 if

$$\theta^* = Q_{\frac{p-c}{p}} D_n = Q_{\frac{9}{10}} D_n = 100.$$

However, while $\theta^* = 100$ maximizes the expected profit, the risk that we lose a lot of money is very high. For our newsvendor example, the probability that we lose at least \$50 while following our optimal policy is given by

$$\mathbb{P}\left[p \min\left\{\theta^*, D_n\right\} - c\theta^* \leq -50\right] = \mathbb{P}\left[10 \min\left\{100, D_n\right\} \leq 50\right] = \mathbb{P}\left[D \leq 5\right] = 1 - \left(\frac{1}{5}\right)^{1/2} \approx .55.$$

A policy that has a 55% chance of losing 50% or more of the initial investment of 100 is a disastrous policy even if the expected return on investment is 90%. In a heavy-tailed environment, focusing on the expected profit leads one to take on imprudent risks. The above policy has only a 30% chance of making a profit of 10% return or more. Instead of trying to maximize the expectation, suppose we want to maximize some quantile. Since

$$\underset{\theta \in \mathbb{R}_+}{\operatorname{argmax}} Q_\alpha \left[p \min\left\{\theta, D_n\right\} - c\theta\right] = Q_\alpha D_n = \frac{1}{(1-\alpha)^2},$$

we maximize the median profit if $\theta^* = 4$. Then, the median profit is 36 while the expected profit is 26 and the minimum amount of profit we can make is 6. That is, we have a 100% chance of making more than 150% profit on the initial investment and our expected return on investment is 650%. The median return on our investment is 900%. Again, when $\theta^* = 4$, we have a 50% chance of making a profit of 36 or more, while we only have a 30% chance of making a profit of 10 or more when $\theta^* = 100$. In fact, $\theta^* = 100$ maximizes the 0.9-quantile of the profit. This example illustrates that it is not always reasonable to try to maximize the expectation even when we can compute the expectation. The merit of the more conservative policy derived from quantile optimization become even more conspicuous if we conduct independent experiments multiple times and observe what happens to the cumulative profit. Let

$$U_i = 10 \min\left\{100, D_i\right\} - 100$$

be the profit of a business that follows the expectation-maximizing policy ($\theta^* = 100$) at the i^{th} iteration. Let

$$V_n = \sum_{i=1}^{n} U_i$$

be the cumulative profit of a business after the n^{th} iteration. Left plot on the **Figure 3.1** shows ten different sample paths of V_n as well as their average.

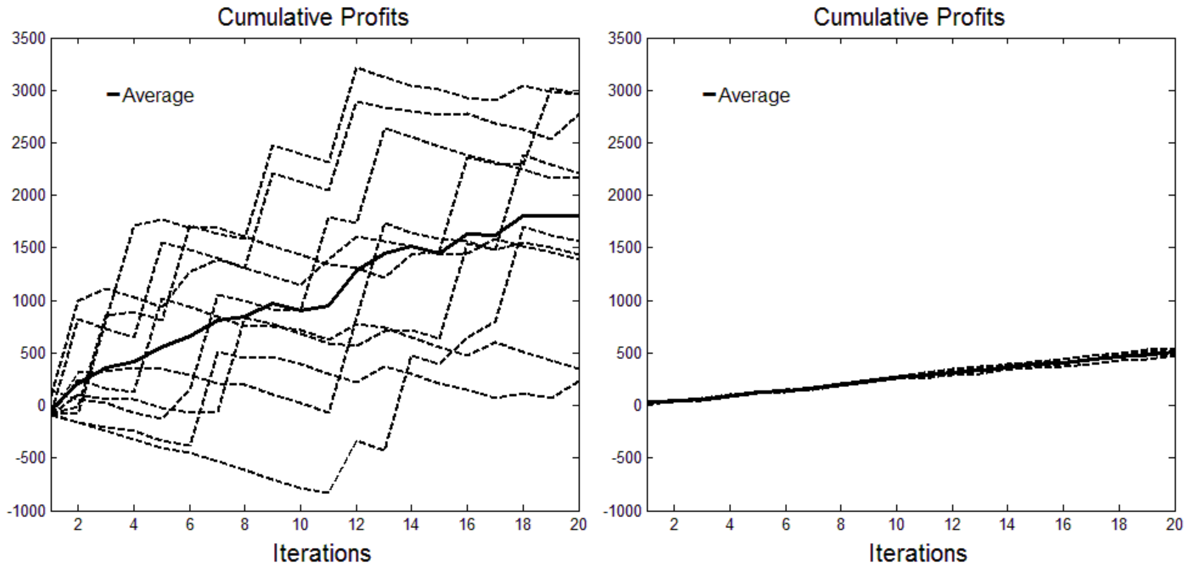

FIG. 3.1. *Cumulative profit obtained by following the expectation-maximizing policy (left) and by following the median-maximizing policy (right), for ten sample paths*

Next, let

$$X_i = 10 \min \{4, D_i\} - 4$$

be the profit of a business that follows the median-maximizing policy ($\theta^* = 4$) at the i^{th} iteration. Let

$$Y_n = \sum_{i=1}^{n} X_i$$

be the cumulative profit. Right plot on the **Figure 3.1** shows several different sample paths following the median-maximization policy. By comparing the two plots, we can see that while V_n grows a lot more faster than Y_n and thus much more profitable on average, it is much more volatile than Y_n. We can observe how risky the expectation-maximizing policy can be; one sample path of V_n reaches below -\$700 after the 11^{th} iteration, implying that the business would have gone bankrupt if it did not have initial capital in excess of \$600.

For the above example, we happened to know the closed-form expression for the quantiles, but there are many applications where we do not have a formula in a closed-form. Often, we do not even know which distribution we are sampling from. In such cases, we must rely on the algorithm (3.8). **Figure 3.2** shows the convergence of the algorithm (3.8) for $\alpha = 0.5, 0.684$, and 0.8, where $\theta_n \to 4, 10$, and 25, respectively. The horizontal axis is shown in the logarithmic scale. **Table 3.1** shows the expected profit and the probability of loss for different quantiles.

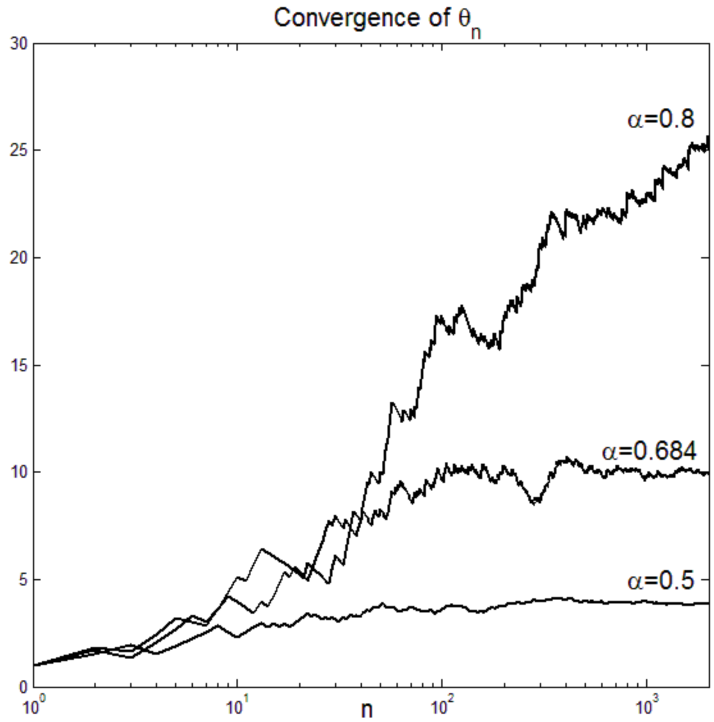

FIG. 3.2. *Convergence of* θ_n *when* $\alpha = 0.5, 0.684,$ *and* 0.8

	θ^*	Expected Profit	Probability of Loss
$\alpha = 0.5$	4	36	0
$\alpha = 0.684$	10	43.2	0
$\alpha = 0.8$	25	65	0.37
$\alpha = 0.9$	100	90	0.68

TABLE 3.1

Optimizing the quantile of the profit for the newsvendor problem

4. Trading Electricity in the Presence of Storage. Trading on the electricity spot market introduces the challenge of dealing with spot prices which are well-known to be heavy tailed ([2],[4],[8]). In this section, we show how quantile optimization can be used to derive a trading strategy in the presence of storage. At each time t, we sign a contract and determine whether to buy or sell electricity during the time interval $[t, t+1)$. We propose a trading strategy and a performance metric that can be assumed to be approximately i.i.d, thus allowing us to apply our quantile optimization algorithm. We use electricity spot market price data from Ercot and PJM West Hub. Ercot is a utility that covers most of Texas. PJM West covers the Chicago metropolitan area. Ercot and PJM represent competitive electricity markets where the price is settled by the market based on the supply and the demand in real time. The electricity price in Ercot is particularly volatile because of the significant role wind energy plays in the market. While demand for electricity is always random, reliance on intermittent energy such as wind or solar makes the supply side of the equation random as well, thus making a large mismatch between supply and demand more likely, resulting in violent price swings.

4.1. Model. Throughout this section, we assume the amount of electricity we can deliver to the market by discharging our storage for a one hour period is one megawatt-hour. We assume that the round-trip efficiency of our storage is $0 < \rho < 1$. ρ is between .65 and .80 for most of the commercially available storage devices [33]. We must import $1/\rho$ megawatt-hours of electricity when charging our storage device for every megawatt-hour of energy that we want to release. Throughout this section, we assume it takes 8 hours to fully charge or discharge our storage device. Thus, our storage capacity is 8 megawatt-hours. Moreover, we assume that our storage capacity is sufficiently small compared to the overall market so that we can always discharge our storage and sell one-megawatt-hour of electricity to the market and we can always buy one-megawatt-hour of electricity and charge our storage if we choose to do so.

State Variables

Let $t \in \mathbb{N}_+$ be a discrete time index corresponding to the decision epoch.

R_t = discretized storage level at time t. $R_t \in \{0, 1, 2, ..., 8\}, \forall t$. One unit of storage corresponds to one megawatt-hour of energy.

p_t = price of electricity per megawatt-hour traded during the time interval $[t, t+1), \forall t$.

When we discharge our storage one unit to sell one megawatt-hour of electricity, we earn p_t. When we buy electricity to charge our storage, we are buying $1/\rho$ megawatt-hour of electricity, thus we have to pay p_t/ρ.

h_t = the electricity price history over the last 99 hours = $(p_{t'})_{t-99 \le t' \le t}$. The history h_t is needed to compute the quantile of the price process.

$S_t = (R_t, h_t)$ = state of the system at time t.

Decision (Action) Variable

x_t = decision that tells us to discharge (sell), hold, or charge (buy) our storage at time t. $x_t \in \mathcal{X}(R_t)$ where

$$\mathcal{X}(R_t) = \begin{cases} \{0, 1\}, & \text{if} \quad R_t = 0, \\ \{-1, 0, 1\}, & \text{if} \quad 0 < R_t < 8, \\ \{-1, 0\}, & \text{if} \quad R_t = 8, \end{cases}$$

is the feasible set for the decision variable. $x_t = -1$ indicates discharge, $x_t = 0$ indicates hold, and $x_t = 1$ indicates charge.

Exogenous Process

$\widehat{p}_t = p_t - p_{t-1} =$ random variable that captures the evolution of the price process $(p_t)_{t \geq 0}$.

Let Ω be the set of all possible outcomes and let $\mathcal{F}$ be a σ-algebra on the set, with filtrations $\mathcal{F}_t$ generated by the information given up to time t :

$$\mathcal{F}_t = \sigma\left(S_0, x_0, \widehat{p}_1, S_1, x_1, \widehat{p}_2, S_2, x_2, ..., \widehat{p}_t, S_t, x_t\right).$$

$\mathbb{P}$ is the probability measure on the measure space $(\Omega, \mathcal{F})$. We have defined the state of our system at time t as all variables that are $\mathcal{F}_t$-measurable and needed to compute our decision at time t.

Storage Transition Function

$$R_{t+1} = R_t + x_t. \tag{4.1}$$

If $x_t = 1$, we import electricity from the market and store it with some conversion factor ρ. That is, we purchase $1/\rho$ megawatt-hours of electricity from the market, but it only fills up one megawatt-hour worth of electricity in our storage due to the conversion loss. If $x_t = -1$, the potential energy in the storage is converted into electricity with a conversion factor of 1, and sold to the market.

4.2. Electricity Price Behavior. At time t, let $p_t^{(i)}$ denote the order statistics of the past hundred hours of data $(p_{t'})_{t-99 \leq t' \leq t}$ sorted in increasing order:

$$p_t^{(1)} \leq p_t^{(2)} \leq ... \leq p_t^{(100)}.$$

Then, let

$$q_t := \min\left\{b \in \{1, 2, ..., 100\} : p_t^{(b)} \geq p_t\right\} \tag{4.2}$$

be the index of p_t in the order statistics. For the hourly electricity spot market price from Texas Ercot and PJM West Hub,

$$\mathbb{P}\left[q_t \leq b\right] \approx \frac{b}{100},\ \forall b \in \{1, 2, ..., 100\}. \tag{4.3}$$

Next, **Table 4.1** shows the behavior of the electricity price for Texas Ercot. One can see that for lower values of q_t, the probability that the price will increase grows, and vice versa. The results are similar for prices for the PJM West hub, as shown in **Table 4.2**. This is because we may assume that the electricity

$b =$	10	20	30	40	50	60	70	80	90
$P\left[p_{t+1} \geq \frac{1}{0.7}p_t \mid q_t \leq b\right] =$	.41	.30	.24	.20	.17	.15	.14	.13	.12
$P\left[p_{t+1} \geq \frac{1}{0.75}p_t \mid q_t \leq b\right] =$	.44	.33	.27	.23	.20	.18	.16	.15	.14
$P\left[p_{t+1} \geq \frac{1}{0.8}p_t \mid q_t \leq b\right] =$	.47	.37	.32	.28	.25	.22	.21	.19	.18
$P\left[p_{t+1} \geq \frac{1}{0.9}p_t \mid q_t \leq b\right] =$	.55	.47	.44	.40	.38	.35	.33	.32	.30
$P\left[p_{t+1} \leq p_t \mid q_t > b\right] =$	.52	.53	.54	.56	.57	.59	.62	.66	.74

TABLE 4.1

Dependence of electricity price behavior on q_t for Texas Ercot

$b =$	10	20	30	40	50	60	70	80	90
$P\left[p_{t+1} \geq \frac{1}{0.7}p_t \mid q_t \leq b\right] =$	.173	.132	.121	.119	.115	.110	.106	.102	.098
$P\left[p_{t+1} \geq \frac{1}{0.75}p_t \mid q_t \leq b\right] =$	.211	.174	.153	.151	.148	.144	.139	.134	.129
$P\left[p_{t+1} \geq \frac{1}{0.8}p_t \mid q_t \leq b\right] =$	.249	.205	.193	.191	.190	.186	.181	.176	.171
$P\left[p_{t+1} \geq \frac{1}{0.9}p_t \mid q_t \leq b\right] =$	.393	.339	.323	.320	.317	.310	.304	.297	.289
$P\left[p_{t+1} \leq p_t \mid q_t > b\right] =$	.52	.54	.55	.56	.58	.60	.65	.72	.83

TABLE 4.2

Dependence of electricity price behavior on q_t for PJM West Hub

price is median-reverting, as shown in [16]. According to [16], the electricity price is heavy-tailed and non-stationary so that the empirical mean cannot be used to characterize the behavior of the price. Instead of using standard Gaussian jump-diffusion processes with mean-reversion, [16] shows that the price behavior is more appropriately modelled as being median-reverting with an underlying heavy-tailed process. Thus, if the current price level is at the higher end of the prices realized in the past one hundred hours, we know there is a high probability that the price will go down due to the median-reversion, and vice versa.

4.3. A Robust Trading Policy. We believe that a robust policy for trading electricity is a policy whose performance is consistent across different sample paths of the price process. The policy optimized based on past data should remain optimal going into the future. In other words, if we have a parameterized policy, the parameters that were optimal for year 2008 and the parameters that were optimal for year 2009 should be almost identical. This implies that by the end of the year 2008, we would know what policy we will use to trade in year 2009 and we would know that the policy will perform well. Our goal is to devise a robust trading policy that utilizes the fact that electricity price is median-reverting, as shown in the previous section.

In prior research, we have shown that the electricity price is heavy-tailed [16], which implies that our trading policy should only depend on zeroth-order statistics, i.e., the quantile function. Since we do not want to have negative exposure to heavy-tailed price movements, we must always buy low and sell high, which is possible because the price process is median-reverting.

In the previous section, we have shown that we can ascertain how high or low the current price level is by comparing it to the percentile constructed from the order statistics. Given a set of potential quantiles $1 \leq \theta^L \leq 50$ and $51 \leq \theta^H \leq 100$, our policy is

$$X_t^{\pi}(S_t) = \begin{cases} 1, & \text{if} \quad R_t < 8 \text{ and } q_t \leq \theta^L \\ -1, & \text{if} \quad R_t > 0 \text{ and } q_t \geq \theta^H \\ 0, & \text{otherwise} \end{cases}. \tag{4.4}$$

where

$$\pi := \left(\theta^L, \theta^H\right).$$

θ^L determines when to buy electricity. Of course, the lower the threshold θ^L, the higher the probability that the price will increase by more than a factor of $1/\rho$, as shown in **Table 4.1** and **Table 4.2**. On the other hand, for smaller values of θ^L, we will buy electricity less frequently and thus we will have less opportunity to make money, as implied from (4.3). Similarly, the higher the threshold θ^H, the higher the probability that the price will decrease, but we will be able to sell electricity less frequently, reducing the opportunities to make money.

Proposition 4.1. *Given $\mathcal{F}_t$, define*

$$F_{t+1}\left(\theta^L, \theta^H, \omega\right) = \begin{cases} \theta^L \cdot \left(p_{t+1} - \frac{1}{\rho} p_t\right)(\omega), & \text{if} \quad q_t \leq \theta_t^L \\ 0, & \text{if} \quad \theta^L < q_t < \theta^H \\ \left(100 - \theta^H\right) \cdot \left(p_t - p_{t+1}\right)(\omega), & \text{if} \quad \theta^H \leq q_t \end{cases} \tag{4.5}$$

to be our objective function. Note that

$$\mathbb{P}\left[q_t \leq \theta^L\right] \approx \frac{\theta^L}{100} \text{ and } \mathbb{P}\left[q_t \geq \theta^H\right] \approx \frac{100 - \theta^H}{100},$$

as shown in (4.3). Then, θ^L satisfies the monotonicity condition (3.1) while θ^H satisfies (3.2).

Proof: Let $F_{t+1}\left(\theta^L, \theta^H\right)$ denote the random variable of which specific sample realization is $F_{t+1}\left(\theta^L, \theta^H, \omega\right)$. Then, for $\forall \omega \in \Omega$, $F_{t+1}\left(\theta^L, \theta^H, \omega\right)$ is a concave function of $\left(\theta^L, \theta^H\right)$. When $\theta^L \geq q_t$, we buy electricity and hence we would make a profit (in a mark-to-market sense) if $p_{t+1} - \frac{1}{\rho} p_t \geq 0$ and have a loss if $p_{t+1} - \frac{1}{\rho} p_t < 0$. When $\theta^H \leq q_t$, we would have correctly sold electricity if $p_{t+1} - p_t \leq 0$, and incur a loss in opportunity cost if $p_{t+1} - p_t \geq 0$.

Next, we know that

$$\frac{\partial}{\partial\theta^L}F_{t+1}\left(\theta^L,\theta^H,\omega\right)=\begin{cases}\left(p_{t+1}-\frac{1}{\rho}p_t\right)(\omega), & \text{if} \quad q_t\le\theta^L_t,\\ 0, & \text{else.}\end{cases}$$

and

$$\frac{\partial}{\partial\theta^H}F_{t+1}\left(\theta^L,\theta^H,\omega\right)=\begin{cases}\left(p_{t+1}-p_t\right)(\omega), & \text{if} \quad \theta^H\le q_t,\\ 0, & \text{else.}\end{cases}$$

For some $\omega^1,\omega^2\in\Omega$,

$$F_{t+1}\left(\theta^L,\theta^H,\omega^1\right)\le F_{t+1}\left(\theta^L,\theta^H,\omega^2\right)$$

if and only if one of the following three conditions hold:

1. $q_t\le\theta^L$ and $\left(p_{t+1}-\frac{1}{\rho}p_t\right)\left(\omega^1\right)\le\left(p_{t+1}-\frac{1}{\rho}p_t\right)\left(\omega^2\right)$
2. $\theta^L<q_t<\theta^H$
3. $\theta^H\le q_t$ and $\left(p_{t+1}-p_t\right)\left(\omega^1\right)\le\left(p_{t+1}-p_t\right)\left(\omega^2\right)$,

These conditions hold true if and only if

$$\frac{\partial}{\partial\theta^L}F_{t+1}\left(\theta^L,\theta^H,\omega^1\right)\le\frac{\partial}{\partial\theta^L}F_{t+1}\left(\theta^L,\theta^H,\omega^2\right)$$

and

$$\frac{\partial}{\partial\theta^H}F_{t+1}\left(\theta^L,\theta^H,\omega^1\right)\ge\frac{\partial}{\partial\theta^H}F_{t+1}\left(\theta^L,\theta^H,\omega^2\right).$$

Therefore, θ^L satisfies (3.1) while θ^H satisfies (3.2). □

Then, our goal is to compute $\left(\theta^{L*},\theta^{H*}\right)$ such that

$$\left(\theta^{L*},\theta^{H*}\right):=\underset{1\le\theta^L\le 50,50<\theta^H<100}{\operatorname{argmax}}\ Q_\alpha\ F_{t+1}\left(\theta^L,\theta^H\right),$$

which can be achieved through the following algorithm

$$\theta^L_t=\begin{cases}\theta^L_{t-1}-\frac{1}{t}(1-\alpha), & \text{if} \quad q_{t-1}\le\theta^L_{t-1} \text{ and } p_t<\frac{1}{\rho}p_{t-1}\\ \theta^L_{t-1}+\frac{1}{t}\alpha, & \text{else}\end{cases} \tag{4.6}$$

and

$$\theta^H_t=\begin{cases}\theta^H_{t-1}-\frac{1}{t}\alpha, & \text{if} \quad \theta^H_{t-1}\le q_{t-1} \text{ and } p_t<p_{t-1}\\ \theta^H_{t-1}+\frac{1}{t}(1-\alpha), & \text{else}\end{cases}. \tag{4.7}$$

While we have proved that the above algorithm will converge if $\left(F^L_t\right)_{t\ge1}$ are i.i.d, $\left(F^L_t\right)_{t\ge1}$ computed from real price data is not likely to be i.i.d. Real data is almost never i.i.d, and we must numerically test if the above algorithm shows a reasonable path of convergence towards a reasonable answer.

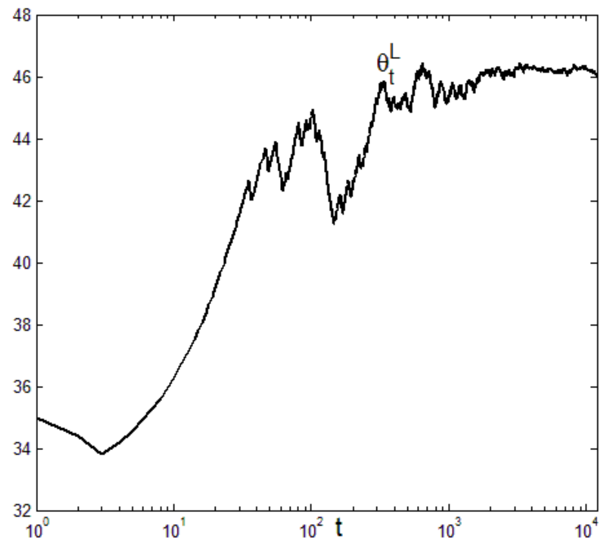

FIG. 4.1. θ_t^L *computed from (4.6) as* $t \to \infty$, *for* $\alpha = 0.5$ *and* $\rho = 0.7$

	$\alpha = 0.3$	$\alpha = 0.4$	$\alpha = 0.5$
$\rho = 0.7$	36.2	43.6	49.4
$\rho = 0.75$	37.1	44.6	50.5
$\rho = 0.8$	39.0	46.3	51.6

TABLE 4.3

θ^L *computed from PJM West Hub Data*

Figure 4.1 shows the rate of convergence for the algorithm (4.6) for the Texas Ercot data. The horizontal axis uses the logarithmic scale. We can see that the algorithm converges in about one thousand iterations. **Table 4.3** shows θ^L computed from the above algorithm using PJM West hub price data for different quantile parameters α and round-trip conversion factor ρ. The lower the value of ρ, the greater the conversion loss, and hence our decision to buy must be more conservative, leading to a lower θ^L. And of course, a larger value of α leads to more risk taking and hence higher θ^L. We show results for $\alpha \leq 0.5$ because for larger α, we end up with policies that take too much risk and buy electricity too readily and end up not making a reasonable amount of profit. Similarly, **Table 4.4** shows θ^L computed from the above algorithm using Texas Ercot price data. Next, **Table 4.5** shows θ^H computed from PJM West and Texas Ercot data for some $\alpha \leq 0.6$. For $\alpha > 0.6$, we end up with policies that take too much risk and sell our electricity too easily and end up not making a reasonable amount of profit.

	$\alpha = 0.3$	$\alpha = 0.4$	$\alpha = 0.5$
$\rho = 0.7$	32.6	39.6	46.0
$\rho = 0.75$	33.7	40.8	47.2
$\rho = 0.8$	35.5	42.2	49.0

Table 4.4

θ^L computed from Texas Ercot Data

	$\alpha = 0.3$	$\alpha = 0.4$	$\alpha = 0.5$	$\alpha = 0.6$
PJM West	87.9	75.2	62.4	50.1
Texas Ercot	85.4	72.1	59.7	45.5

Table 4.5

θ^H computed from Texas Ercot and PJM West Data

5. Maximum Profit Trading Policy. In the previous section, we have shown how to compute θ^L and θ^H based on our risk-appetite determined by α. If we want to find out exactly how much risk we would have had to accept in the past to maximize our profit, we can first find θ^L and θ^H that maximizes the cumulative profit

$$C\left(\theta^L, \theta^H\right) := -\sum_{t=0}^{T} p_t x_t,$$

through deterministic search. Then, α that is closest to producing the above to θ^L and θ^H tells us what the necessary risk-appetite should have been for us to have maximized our profit. Assuming $\rho = .75$, for year 2006, when the median price of the year was \$47.1 in Texas Ercot, $\theta^L = 39$ and $\theta^H = 71$ would have given us the highest profit of \$7.6 per hour, and these correspond to $\alpha \approx 0.36$. In other words, we should have been fairly conservative in controlling the probability of losses in order to maximize our profit. This is because the heavy-tailed spikes in the electricity prices have a disproportionate impact on profits and losses. Instead of trading often and trying to make a profit as often as possible, we must be patient and wait long enough for large price deviations to occur. These large deviations actually occur often enough so that we can eventually get a good price for buying and selling electricity at a good profit. For year 2007, when the median price of the year was \$48.3, $\theta^L = 36$ and $\theta^H = 68$ would have been optimal, giving us a profit of \$7.2 per hour. This again corresponds to about $\alpha \approx 0.36$. Thus, even though the electricity price is highly volatile and non-stationary with wildly differing price distributions over time, the risk-appetite that would have been optimal for year 2006 was still optimal for year 2007. Results are similar for year 2008 and 2009 and for the data from PJM West hub. The trading policy based on optimizing the quantiles is truly robust

year	median price	optimal (θ_L, θ_H)	optimal profit	last year's (θ_L, θ_H)	profit
2007	\$48.3	$(36, 68)$	\$7.2/hour	$(39, 71)$	\$7.1/hour
2008	\$48.7	$(39, 68)$	\$17.0/hour	$(36, 68)$	\$16.8/hour
2009	\$23.5	$(39, 69)$	\$6.3/hour	$(39, 68)$	\$6.2/hour

Table 5.1

Performance of quantile based trading policy in Ercot

- it is not sensitive to the particular sample paths of data. This is because the policy captures the core characteristics of the data. The hourly electricity spot price is median-reverting and heavy-tailed.

Table 5.1 shows the performance of the above trading policy from year 2007 to 2009 in Ercot. We can see that the difference between the profit we can make using the optimal parameters $\left(\theta^L, \theta^H\right)$ and the profit we can make using the parameters computed from the previous year's data, is negligible, implying that the quantile-based trading policy is robust.

6. Conclusion. In this paper, we have presented a simple, provably convergent algorithm for optimizing the quantile of a random function. The algorithm follows the spirit of Robbins and Monro [31], and does not require us to store and sort the data. In the real world, we must often deal with non-stationary and heavy-tailed data. In that case, a policy based on higher-order statistics such as the expectation and the variance can be unstable at best or enormously risky, as we have shown with our example of newsvendor problem. We also showed that in the electricity spot market, a simple trading policy based on quantile optimization is robust and generate good profit.

REFERENCES

[1] Arce, G.R, Paredes, J.L. 2000. Recursive weighted median filters admitting negative weights and their optimization. *IEEE Trans. Signal Processing.* **48**(3) 768-779.

[2] Bouchaud, J.P. 2001. Power laws in economics and finance: some ideas from physics. *Quantitative Finance.* **1**(1) 105-112.

[3] Bertsimas, D, Thiele, A. 2006. A robust optimization approach to inventory theory. *Operations Research.* **54**(1) 150-168

[4] Breymann, W. Dias, A., Embrecths, P. 2003. Dependence structure for multivariate high-frequency data in finance. *Quantitative Finance.* **3**(1) 1-14.

[5] Bystrom, H.N.E. 2002. Extreme value theory and extremely large electricity price changes. *International Review of Economics & Finance.* **14**(1) 41-55.

[6] Cartea, A., Figueroa, M.G. 2005. Pricing in electricity markets: a mean-reverting jump-diffusion model with seasonality. *Applied Mathematical Finance.* **12**(4) 313-335.

[7] Eydeland, A., Wolyniec, K. 2003. *Energy and Power Risk Management.* John Wiley & Sons, New Jersey.

[8] Farmer, J.D., Lillio, F. 2004. On the origin of power law tails in price fluctuations. *Quantitative Finance.* **4**(1) C7-C11.

[9] George, A.P., Powell, W.B. 2006. Adaptive stepsizes for recursive estimation with applications in approximate dynamic programming. *Machine Learning.* **65**(1) 167-198.

[10] Hill, B.M. 1975. A simple general approach to inference about the tail of a distribution. *Annals of Statistics.* **3**(5) 1163-1174.

[11] Heidelberger, P., Lewis, P.A.W. 1984. Quantile Estimation in Dependent Sequences. *Operations Research.* **32**(1) 185-209.

[12] Heyman, D.P., Sobel, M.J. 2003. *Stochastic Models in Operations Research.* Dover Publications, New York.

[13] Hsing, T. 1991. On tail index estimation using dependent data. *Annals of Statistics.* **19**(3) 1547-1569.

[14] Huisman, R., Maiheu, R. 2003. Regime jumps in electricity prices. *Energy Economics.* **25**(5) 425-434.

[15] Kibzun, A.I., Kurbakovskiy, V.Y. 1991. Guaranteeing approach to solving quantile optimization problems. *Annals of Operations Research.* **30**(1) 81-93.

[16] Kim, J.H., Powell, W.B. 2011. An Hour-Ahead Prediction Model for Heavy-Tailed Spot Prices. Submitted to *Energy Economics.*

[17] Kleywegt, A.J., Shapiro, A., and Homem-de-Mello, T. 2002. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization.* **12**(2) 479-502.

[18] Koenker, R. 2005. *Quantile Regression.* Cambridge University Press, New York.

[19] Kushner, H.J, Yin, G. 2003. *Stochastic approximation and recursive algorithms and applications.* Sringer-Verlag, New York.

[20] Lai, T.Z. 2003. Stochastic Approximations. *Annals of Statistics.* 31(2) 391-406.

[21] LeBaron, B. 2001. Stochastic volatiliy as a simple generator of apparent financial power laws and long memory. *Quantitative Finance.* **1**(6) 621-631.

[22] Lucia, J.J., Schwartz, E.S. 2002. Electricity prices and power derivatives: Evidence from the Nordic power exchange. *Review of Derivatives Research.* **5**(1) 5-20.

[23] Peng, L. 2001. Estimating the mean of a heavy-tailed distribution. *Statistics and Probability Letters.* **52**(3) 255-264.

[24] Peng, L. 2004. Empirical-likelihood based confidence interval for the mean with a heavy-tailed distribution. *Annals of Statistics.* **32**(3) 1192-1214.

[25] Peng, L., Qi, Y. 2006. Confidence regions for high quantiles of a heavy-tailed distribution. *Annals of Statistics.* **34**(4) 1964-1986.

[26] Petruzzi, N.C., Dada, M. 1999. Pricing and the newsvendor problem: A review with extensions. *Operations Research.* **47**(2) 183-194.

[27] Pilipovic, D. 2007. *Energy Risk: Valuing and Managing Energy Derivatives.* Mcgraw-Hill, New York.

[28] Powell, W.B. 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality.* John Wiley and Sons, New York.

[29] Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming.* John Wiley and Sons, New York.

[30] Resnick, S., Starica, C. 1998. Tail index estimation for dependent data. *Annals of Applied Probability.* **8**(4) 1156-1183.

[31] Robbins, H., Monro, S. 1951. A stochastic Approximation Method. *Annals of Mathematical Statistics.* **22**(3) 400-407.

[32] Schwartz, E.S. 1997. Stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance.* **52**(3) 923-974.

[33] Sioshansi, R., Denholm, P., Jenkin, T., Weiss, J. 2009. Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Economics.* **31**(2) 269-277.

[34] Taleb, N.N. 2007. Black swans and the domains of statistics. *The American Statistician.* **61**(3) 198-200.

[35] Taleb, N. N. 2009. Errors, robustness, and the fourth quadrant. *International Journal of Forecasting.* **25**(4) 744-759.

[36] Yin, L., Yang, R., Gabbouj, M., Neuvo, Y. 1996. Weighted Median Filters: A Tutorial. *IEEE Transactions on Circuits and Systems.* **43**(3) 157-192.