

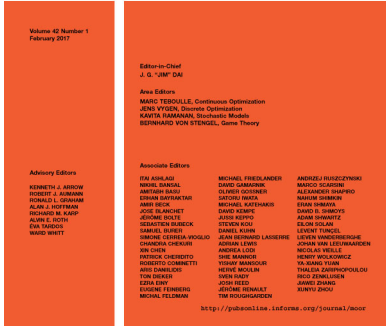


Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Risk-Averse Approximate Dynamic Programming with Quantile-Based Risk Measures

Daniel R. Jiang, Warren B. Powell



To cite this article:

Daniel R. Jiang, Warren B. Powell (2017) Risk-Averse Approximate Dynamic Programming with Quantile-Based Risk Measures. *Mathematics of Operations Research*

Published online in *Articles in Advance* 13 Nov 2017

<https://doi.org/10.1287/moor.2017.0872>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2017, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Risk-Averse Approximate Dynamic Programming with Quantile-Based Risk Measures

 Daniel R. Jiang,^a Warren B. Powell^b
^aDepartment of Industrial Engineering, University of Pittsburgh, Pittsburgh, Pennsylvania 15261; ^bDepartment of Operations Research and Financial Engineering, Princeton University, Princeton, New Jersey 08540

Contact: drjiang@pitt.edu,  <http://orcid.org/0000-0002-5388-8061> (DRJ); powell@princeton.edu,

 <http://orcid.org/0000-0002-4364-7555> (WBP)

Received: September 4, 2015

Revised: June 12, 2016; January 4, 2017; April 9, 2017

Accepted: May 4, 2017

Published Online in *Articles in Advance*: November 13, 2017

MSC2010 Subject Classification: Primary: 90-C39; secondary: 93-E35, 62-L20

OR/MS Subject Classification: decision analysis: risk; dynamic programming

<https://doi.org/10.1287/moor.2017.0872>
Copyright: © 2017 INFORMS

Abstract. In this paper, we consider a finite-horizon Markov decision process (MDP) for which the objective at each stage is to minimize a quantile-based risk measure (QBRM) of the sequence of future costs; we call the overall objective a *dynamic quantile-based risk measure* (DQBRM). In particular, we consider optimizing dynamic risk measures where the one-step risk measures are QBRMs, a class of risk measures that includes the popular *value at risk* (VaR) and the *conditional value at risk* (CVaR). Although there is considerable theoretical development of risk-averse MDPs in the literature, the computational challenges have not been explored as thoroughly. We propose data-driven and simulation-based approximate dynamic programming (ADP) algorithms to solve the risk-averse sequential decision problem. We address the issue of inefficient sampling for risk applications in simulated settings and present a procedure, based on importance sampling, to direct samples toward the “risky region” as the ADP algorithm progresses. Finally, we show numerical results of our algorithms in the context of an application involving risk-averse bidding for energy storage.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/moor.2017.0872>.

Keywords: approximate dynamic programming • dynamic risk measures • energy trading • reinforcement learning • Q-learning

1. Introduction

Sequential decision problems, in the form of Markov decision processes (MDPs), are most often formulated with the objective of minimizing an expected sum of costs or maximizing an expected sum of rewards (Puterman [51], Bertsekas and Tsitsiklis [9], Powell [49]). However, it is becoming more and more evident that solely considering the expectation is insufficient as risk-preferences can vary greatly from application to application. Broadly speaking, the expected value can fail to be useful in settings containing either heavy-tailed distributions or rare, but high-impact events. For example, heavy-tailed distributions arise frequently in finance (electricity prices are well-known to possess this feature; see Byström [14], Kim and Powell [37]). In this case, the mean of the distribution itself may not necessarily be a good representation of the randomness of the problem; instead, it is likely useful to introduce a measure of risk on the tail of the distribution as well. The rare event situation is, in a sense, the inverse case of the heavy-tail phenomenon, but it can also benefit from a risk measure other than the expectation. To illustrate, certain problems in operations research can be complicated by critical events that happen with small probability, such as guarding against stock-outs and large back-orders in inventory problems (see Glasserman and Liu [29]) or managing the risk of the failure of a high-value asset (see Enders et al. [25]). In these circumstances, the merit of a policy might be measured by the number of times that a *bad event* happens over some time horizon.

One way to introduce risk-aversion into sequential problems is to formulate the objective using *dynamic risk measures* (Ruszczyński [60]). A rough preview, without formal definitions, of our optimization problem is as follows: We wish to find a policy that minimizes risk, as assessed by a certain type of dynamic risk measure. The objective can be written as

$$\min_{\pi \in \Pi} \rho_0^\alpha (C_1^\pi + \rho_1^\alpha (C_2^\pi + \cdots + \rho_{T-1}^\alpha (C_T^\pi) \cdots)),$$

where Π is a set of policies, $\{C_t^\pi\}$ are costs under policy π , and $\{\rho_t^\alpha\}$ are one-step risk measures (i.e., components of the overall dynamic risk measure). Precise definitions are given in the subsequent sections. We focus on the case where the objective at each stage is to optimize a *quantile-based risk measure* (QBRM) of future costs; we call the overall objective a *dynamic quantile-based risk measure* (DQBRM).

This paper makes the following contributions. First, we describe a new data-driven or simulation-based ADP algorithm, called *Dynamic-QBRM ADP*, that is similar in spirit to established asynchronous algorithms like Q-learning (see Watkins and Dayan [73]) and lookup table approximate value iteration (see, e.g., Bertsekas and Tsitsiklis [9], Powell [49]), where one state is updated per iteration. The second contribution of the paper is a companion sampling procedure to Dynamic-QBRM ADP, which we call *risk-directed sampling* (RDS). As we describe above, when dealing with risk, there is a large class of problems in which we are inherently dealing with rare, but very costly events. Broadly speaking, the evaluation of a QBRM that is focused on the tail of the distribution (e.g., CVaR at, say, the 99% level) depends crucially on efficiently directing the algorithm toward sampling these “risky” regions. In this part of the paper, we consider the question: is there a way to learn, as the ADP algorithm progresses, the interesting values of the information process to sample?

The paper is organized as follows. We first provide a literature review in Section 2. In Section 3, we give our problem formulation, a brief introduction to dynamic risk measures, and the definition of a class of quantile-based risk measures. Next, we introduce the algorithm for solving risk-averse MDPs in Section 4 and give a theoretical analysis in Section 5. In Section 6, we discuss sampling issues and describe the companion sampling procedure. We show numerical results on an example energy trading application in Section 7 and conclude in Section 8.

2. Literature Review

The theory of *dynamic risk measures* and the notion of *time-consistency* (see e.g. Riedel [53], Artzner et al. [4], Cheridito et al. [17]) is extended to the setting of sequential optimization problems in Ruszczyński and Shapiro [61] and Ruszczyński [60], in which it is proved that any time-consistent dynamic risk measure can be written as compositions of *one-step conditional risk measures* (these are simply risk measures defined in a conditional setting, analogous to the conditional expectation for the traditional case). From this, a Bellman recursion is obtained, becoming a familiar way of characterizing optimal policies. Building on the theory of dynamic programming, versions of exact value iteration and policy iteration are also developed in Ruszczyński [60]. Later, in Çavus and Ruszczyński [15], these exact methods are analyzed in the more specific case of undiscounted transient models.

Under the assumption that we use one-step *coherent risk measures*, as axiomatized in Artzner et al. [3], the value functions of a risk-averse Markov decision process with a convex terminal value function can be easily shown to satisfy convexity using the fact that coherent risk measures are convex and monotone. Therefore, the traditional method of stochastic dual dynamic programming (SDDP) of Pereira and Pinto [44] for multistage, risk-neutral problems, which relies on the *convexity* of value functions, can be adapted to the risk-averse case. This idea is successfully explored in Philpott and de Matos [47], Shapiro et al. [69], and Philpott et al. [48], with applications to the large-scale problem of hydro-thermal scheduling using one-step mean-CVaR (convex combination of mean and CVaR) and one-step mean-upper semideviation risk measures. The main drawbacks of risk-averse SDDP are: (1) the cost function must be linear in the state; (2) some popular risk measures, such as *value at risk* (VaR), are excluded because they are not coherent and do not imply convex value functions; and (3) the full risk measure (which can be recast as an expectation in certain instances) has to be computed at every iteration. Since no convexity or linearity assumptions are made in this paper, we take an alternative approach from the SDDP methods and instead assume the setting of finite state and action spaces, as in Q-learning. At the same time, because the default implementation of our approach does not take advantage of structure, it is limited to smaller problems. Extensions to the methods proposed in this paper for exploiting structure can be made by following techniques such as those discussed in Powell et al. [50], Nascimento and Powell [42], and Jiang and Powell [32].

Recursive stochastic approximation methods have been applied to estimating quantiles in static settings (see Tierney [71], Bardou and Frikha [6], Kan [35]). Related to our work, a policy gradient method for optimizing MDPs (with a risk-neutral objective) under a CVaR constraint is given in Chow and Ghavamzadeh [18]. All of these methods are related to ours in the sense that the minimization formula Rockafellar and Uryasev [57, Theorem 10] for CVaR is optimized with gradient techniques. In our multistage setting with dynamic risk measures (which is also coupled with optimal control), there are some new interesting complexities, including the fact that every new observation (or data point) is generated from an *imperfect distribution* of future costs that is “bootstrapped” from the previous estimate of the value function. This means that not only are the observations inherently biased, but the errors compound over time—this was not the case for the static setting considered in earlier work. Under reasonable assumptions, we analyze both the almost sure convergence and convergence rates for our proposed algorithms.

Our risk-directed sampling procedure is inspired by adaptive importance sampling strategies from the literature, such as the celebrated *cross-entropy method* of Rubinstein [58]. See, e.g., Al-Qaq et al. [2], Bardou and Frikha [6], Egloff and Leippold [24], and Ryu and Boyd [63] for other similar approaches. The critical difference in our approach is that in an ADP setting, we have the added difficulty of not being able to assume perfect knowledge of the objective function; rather, our observations are noisy and biased. To our knowledge, this is the first time an adaptive sampling procedure has been combined with a value function approximation algorithm in the risk-averse MDP literature. The closest paper is by Kozmík and Morton [39], which considers an importance sampling approach for policy evaluation.

3. Problem Formulation

In this section, we establish the setting of the paper. In particular, we describe the risk-averse model, introduce the concept of dynamic risk measures, and define a class of quantile-based risk measures.

3.1. Model

We consider an MDP with a finite time-horizon, $t = 0, 1, 2, \dots, T$, where the last decision is made at time $t = T - 1$, so that the set of decision epochs is given by $\mathcal{T} = \{0, 1, 2, \dots, T - 1\}$. Given a probability space $(\Omega, \mathcal{F}, \mathbf{P})$, we define a discrete-time stochastic process $\{W_t\}_{t=0}^T$, with $W_t \in \mathcal{W}$ for all t , as the exogenous information process in the sequential decision problem, where W_t is adapted to a filtration $\{\mathcal{F}_t\}_{t=0}^T$, with $\{\emptyset, \Omega\} = \mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \dots \subseteq \mathcal{F}_T \subseteq \mathcal{F}$. We assume that all sources of randomness in the problem are encapsulated by the process $\{W_t\}$ and that it is independent across time. For computational tractability, we work in the setting of finite state and action spaces. Let the state space be denoted \mathcal{S} , and let the action space be \mathcal{A} , where $|\mathcal{S}| < \infty$ and $|\mathcal{A}| < \infty$. The set of feasible actions for each state $s \in \mathcal{S}$, written \mathcal{A}_s , is a subset of \mathcal{A} . The set $\mathcal{U} = \{(s, a) \in \mathcal{S} \times \mathcal{A} : a \in \mathcal{A}_s\}$ is the set of all feasible state-action pairs. The stochastic process describing the states of the system is $\{S_t\}_{t=0}^T$, where S_t is an \mathcal{F}_t -measurable random variable taking values in \mathcal{S} , and a_t is a feasible action determined by the decision maker using S_t . Furthermore, let \mathcal{X}_t denote the space of \mathcal{F}_t -measurable random variables and $\mathcal{X}_{t,T} = \mathcal{X}_t \times \dots \times \mathcal{X}_T$.

We model the system using a *transition function* or *system model* $S^M: \mathcal{S} \times \mathcal{A} \times \mathcal{W} \rightarrow \mathcal{S}$, which produces the next state S_{t+1} given a current state S_t , action a_t , and an outcome of the exogenous process W_{t+1} : $S_{t+1} = S^M(S_t, a_t, W_{t+1})$. The *cost* for time t is given by $c_t(S_t, a_t, W_{t+1})$, where $c_t: \mathcal{S} \times \mathcal{A} \times \mathcal{W} \rightarrow \mathbb{R}$ is the *cost function*. A policy is a sequence of decision functions $\{A_0^\pi, A_1^\pi, \dots, A_{T-1}^\pi\}$ indexed by $\pi \in \Pi$, where Π is the index set of all policies. Each decision function $A_t^\pi: \mathcal{S} \rightarrow \mathcal{A}$ is a mapping from a state to a feasible action, such that $A_t^\pi(s) \in \mathcal{A}_s$ for any state s . Let the sequence of costs under a policy π be represented by the process C_t^π for $t = 1, 2, \dots, T$, where

$$C_t^\pi = c_{t-1}(S_{t-1}^\pi, A_{t-1}^\pi(S_{t-1}^\pi), W_t) \in \mathcal{X}_t,$$

where $\{S_t^\pi\}$ are the states visited while following policy π . Note that C_t^π refers to the cost from time $t - 1$, but the index of t refers to its measurability: C_t^π depends on information only known at time t .

3.2. Review of Dynamic Risk Measures

In this subsection, we briefly introduce the notion of a dynamic risk measure; for a more detailed treatment, see, e.g., Frittelli and Gianin [27], Riedel [53], Pflug and Ruszczyński [46], Boda and Filar [10], Cheridito et al. [17], and Acciaio and Penner [1]. Our presentation closely follows that of Ruszczyński [60], which develops the theory of dynamic risk measures in the context of MDPs. First, a *conditional risk measure* is a mapping $\rho_{t,T}: \mathcal{X}_{t,T} \rightarrow \mathcal{X}_t$ that satisfies the following monotonicity requirement: for $X, Y \in \mathcal{X}_{t,T}$ and $X \leq Y$ (componentwise and almost surely), $\rho_{t,T}(X) \leq \rho_{t,T}(Y)$.

Given a sequence of future costs C_t, \dots, C_T , the intuitive meaning of $\rho_{t,T}(C_t, \dots, C_T)$ is a *certainty equivalent* cost (i.e., at time t , one is indifferent between incurring $\rho_{t,T}(C_t, \dots, C_T)$ and the alternative of being subjected to the stream of stochastic future costs). See Rudloff et al. [59] for an in-depth discussion regarding the certainty equivalent interpretation in the context of multistage stochastic models. A *dynamic risk measure* $\{\rho_{t,T}\}_{t=0}^T$ is a sequence of conditional risk measures $\rho_{t,T}: \mathcal{X}_{t,T} \rightarrow \mathcal{X}_t$, which allows us to evaluate the future risk at any time t using $\rho_{t,T}$. Of paramount importance to the theory of dynamic risk measures is the notion of *time-consistency*, which says that if from the perspective of some future time τ , one sequence of costs is riskier than another and the two sequences of costs are identical from the present until τ , then the first sequence is also riskier from the present perspective (see Ruszczyński [60] for the full technical definition).

Other definitions of time-consistency can be found in the literature, e.g., Boda and Filar [10], Cheridito and Stadje [16], and Shapiro [68]. Though they may differ technically, these definitions share the same intuitive spirit. Under the conditions:

$$\rho_{t,T}(0, \dots, 0) = 0 \quad \text{and} \quad \rho_{t,T}(C_t, C_{t+1}, \dots, C_T) = C_t + \rho_{t,T}(0, C_{t+1}, \dots, C_T),$$

it is proven in Ruszczyński [60] that for some *one-step conditional risk measures* $\rho_t: \mathcal{X}_{t+1} \rightarrow \mathcal{X}_t$, a time-consistent, dynamic risk measure $\{\rho_{t,T}\}_{t=0}^T$ can be expressed using the following nested representation:

$$\rho_{t,T}(C_t, \dots, C_T) = C_t + \rho_t(C_{t+1} + \rho_{t+1}(C_{t+2} + \dots + \rho_{T-1}(C_T) \dots)),$$

It is thus clear that we can take the reverse approach and define a time-consistent dynamic risk measure by simply specifying a set of one-step conditional risk measures $\{\rho_t\}_{t=0}^T$. This is a common method that has been used in the literature when applying the theory of dynamic risk measures in practice (see, e.g., Philpott and de Matos [47], Philpott et al. [48], Shapiro et al. [69], Kozmík and Morton [39], Rudloff et al. [59]).

3.3. Quantile-Based Risk Measures

In this paper, we focus on simulation techniques where the one-step conditional risk measure belongs to a specific class of risk measures called *quantile-based risk measures* (QBRM). Although the term *quantile-based risk measure* has been used in the literature to refer to risk measures that are similar in spirit to VaR and CVaR (see, e.g., Dowd and Blake [22], Neise [43], Sereda et al. [67]), it has not been formally defined. First, let us describe these two popular risk measures, which serve to motivate a more general definition for a QBRM.

Also known as the *quantile risk measure*, VaR is a staple of the financial industry (see, e.g., Duffie and Pan [23]). Given a real-valued random variable X (representing a loss) and a risk level $\alpha \in (0, 1)$, the *VaR* or *quantile* of X is defined to be

$$\text{VaR}^\alpha(X) = q^\alpha(X) = \inf_u \{ \mathbf{P}(X \leq u) \geq \alpha \}.$$

To simplify our notation, we use $q^\alpha(X)$ in the remainder of this paper. It is well known that VaR does not satisfy *coherency* (Artzner et al. [3]), specifically the axiom of *subadditivity*, an appealing property that encourages diversification. Despite this, several authors have given arguments in favor of VaR. For example, Danielsson et al. [20] concludes that in practical situations, VaR typically exhibits subadditivity. Dhaene et al. [21] and Ibragimov and Walden [30] give other points of view on why VaR should not be immediately dismissed as an effective measure of risk. A nested version of VaR for use in a multistage setting is proposed in Cheridito and Stadje [16], though practical implications have not been explored in the literature.

CVaR is a coherent alternative to VaR and has been both studied and applied extensively in the literature. Although the precise definitions may slightly differ, CVaR is also known by names such as *expected shortfall*, *average value at risk*, or *tail conditional expectation*. Given a general random variable X , the following characterization is given in Rockafellar and Uryasev [57]:

$$\text{CVaR}^\alpha(X) = \inf_u \left\{ u + \frac{1}{1-\alpha} \mathbf{E}[(X-u)^+] \right\} = q^\alpha(X) + \frac{1}{1-\alpha} \mathbf{E}[(X - q^\alpha(X))^+].$$

Applications of risk-averse MDPs using dynamic risk measures have largely focused on combining CVaR with expectation; once again, see Philpott and de Matos [47], Philpott et al. [48], Shapiro et al. [69], Kozmík and Morton [39], and Rudloff et al. [59].

For the purposes of this paper, we offer the following general definition of a QBRM that allows dependence on more than one quantile; the definition includes the above two examples as special cases.

Definition 1 (Quantile-Based Risk Measure (QBRM)). Let X be a real-valued random variable. A *quantile-based risk measure* ρ^α can be written as the expectation of a function of X and finitely many of its quantiles. More precisely, ρ^α takes the form

$$\rho^\alpha(X) = \mathbf{E}[\Phi(X, q^{\alpha_1}(X), q^{\alpha_2}(X), \dots, q^{\alpha_m}(X))], \quad (1)$$

where $\alpha \in \mathbb{R}^m$ is a vector of m risk levels, $\alpha_1, \alpha_2, \dots, \alpha_m$, and a function $\Phi: \mathbb{R}^{m+1} \rightarrow \mathbb{R}$, chosen so that ρ^α satisfies monotonicity, translation invariance, and positive homogeneity (see Artzner et al. [3] for the precise definitions and note that we interpret X as a random loss or a cost).

Our definition of QBRMs is largely motivated by practical considerations. First, the definition covers the two most widely used risk measures, VaR and CVaR, as special cases under a single framework; in addition, the flexibility allows for the specification of more sophisticated risk measures that may or may not be coherent. As previously mentioned, there are situations where nonconvex (and thus, not coherent) risk measures are appropriate (Dhaene et al. [21]). Another motivation for this definition of a QBRM is that it allows us to easily construct a risk measure such that $\text{VaR}^\alpha(X) \leq \rho^\alpha(X) \leq \text{CVaR}^\alpha(X)$, because, as Belles-Sampera et al. [7] points out, one issue with VaR is that it can underestimate large losses, but at the same time, some practitioners of the financial and insurance industries find CVaR to be too conservative.

We see that VaR is trivially a QBRM with $\Phi(X, q) = q$. CVaR can also be easily written as a QBRM, using the function $\Phi(X, q) = q + (1 - \alpha)^{-1}(X - q)^+$. Although our approach can be applied to any risk measure of the form (1), we use the CVaR risk measure in the empirical work of Section 7, due to its popularity in a variety of application areas.

3.4. Dynamic Quantile-Based Risk Measures

Notice that, so far, we have developed QBRMs in a “static” setting (the value of the risk measure is in \mathbb{R}) for simplicity. Given a random variable $X \in \mathcal{X}_{t+1}$ and a risk level $\alpha \in (0, 1)$, the conditional counterpart for the quantile is

$$q_t^\alpha(X) = \inf_{U \in \mathcal{X}_t} \{P(X \leq U \mid \mathcal{F}_t) \geq \alpha\} \in \mathcal{X}_t.$$

Using this new definition, we can similarly extend the definition of a QBRM to the conditional setting by replacing (1) with

$$\rho_t^\alpha(X) = \text{Ebig}[\Phi(X, q_t^{\alpha_1}(X), q_t^{\alpha_2}(X), \dots, q_t^{\alpha_m}(X)) \mid \mathcal{F}_t],$$

and replacing the required properties of monotonicity, translation invariance, and positive homogeneity in Definition 1 with their conditional forms given in Ruszczyński [60] (denoted therein by A2, A3, and A4). For the sake of notational simplicity, let us assume that all parameters, i.e., $m, \alpha_1, \dots, \alpha_m, \Phi$ are static over time, but we remark that an extension to time-dependent (and even state-dependent) versions of the one-step conditional risk measure is possible. Let $\tilde{\rho}_t^\alpha$ be a (conditional) QBRM that measures tail risk. In applications, a weighted combination of a tail risk measure with the traditional expectation ensures that the resulting policies are not driven completely by the tail behavior of the cost distribution; we may use QBRMs of the form $\rho_t^\alpha(X) = (1 - \lambda)\text{E}[X \mid \mathcal{F}_t] + \lambda \tilde{\rho}_t^\alpha(X)$, where $\lambda \in [0, 1]$.

Using one-step conditional risk measures as building blocks, we can define a dynamic risk measure to be $\{\rho_t^\alpha\}_{t=0}^T$, which we refer to as a *dynamic quantile-based risk measure* (DQBRM). The dynamic risk measures obtained when $\rho_t^\alpha = \text{VaR}_t^\alpha$ and $\rho_t^\alpha = \text{CVaR}_t^\alpha$ (the conditional forms of VaR and CVaR) are precisely the time-consistent risk measures suggested in Cheridito and Stadje [16] under the names *composed value at risk* and *composed conditional value at risk*.

3.5. Objective Function

We are interested in finding optimal risk-averse policies under objective functions specified using a DQBRM. The problem is

$$\min_{\pi \in \Pi} \rho_0^\alpha(C_1^\pi + \rho_1^\alpha(C_2^\pi + \dots + \rho_{T-1}^\alpha(C_T^\pi) \dots)). \quad (2)$$

The upcoming theorem, proven in Ruszczyński [60], gives the Bellman-like optimality equations for a risk-averse model. We state it under the assumption that the current period contribution is random, differing slightly from the original statement. A point of clarification: the original theorem is proved within the setting where the one-step risk measures satisfy conditional forms of the axioms of Artzner et al. [3] for coherent risk measures. In our setting, however, the QBRM ρ_t^α is only assumed to satisfy (conditional forms of) monotonicity, positive homogeneity, and translation invariance, but not necessarily convexity. The crucial step of the proof given in Ruszczyński [60] relies only on monotonicity and an associated *interchangeability* property (see Ruszczyński and Shapiro [62, Theorem 7.1], Ruszczyński [60, Theorem 2]). The assumption of convexity is therefore not necessary for the following theorem.

Theorem 1 (Bellman Recursion for Dynamic Risk Measures, Ruszczyński [60]). *The sequential decision problem (2) has optimal value functions given by*

$$V_t^*(s) = \min_{a_t \in \mathcal{A}_s} \rho_t^\alpha(c_t(s, a_t, W_{t+1}) + V_{t+1}^*(S_{t+1})), \quad \text{for all } s \in \mathcal{S}, t \in \mathcal{T},$$

$$V_T^*(s) = 0, \quad \text{for all } s \in \mathcal{S}.$$

The decision functions of an optimal policy π^* are given by

$$A_t^{\pi^*}(s) \in \arg \min_{a_t \in \mathcal{A}_s} \rho_t^\alpha(c_t(s, a_t, W_{t+1}) + V_{t+1}^*(S_{t+1})), \quad \text{for all } s \in \mathcal{S}, t \in \mathcal{T},$$

which map to a minimizing action of the optimality equation.

For computational purposes, we are interested in interchanging the minimization operator and the risk measure ρ_t^α and thus appeal to the *state-action value function* or *Q-factor* formulation of the Bellman equation. Define the state-action value function over the state-action pairs $(s, a) \in \mathcal{U}$ to be $Q_t^*(s, a) = \rho_t^\alpha(c_t(s, a, W_{t+1}) + V_{t+1}^*(S_{t+1}))$, for $t \in \mathcal{T}$ and let $Q_T^*(s, a) = 0$. Thus, the counterpart to the recursion in Theorem 1 is

$$Q_t^*(s, a) = \rho_t^\alpha\left(c_t(s, a, W_{t+1}) + \min_{a' \in \mathcal{A}_{S_{t+1}}} Q_{t+1}^*(S_{t+1}, a')\right), \quad (3)$$

with the minimization occurring inside of the risk measure.

3.6. Some Remarks on Notation

For simplicity, we henceforth refer to Q^* simply as the *optimal value function*. Let $d = |\mathcal{U}|$ and $D = |\mathcal{U}|(T + 1)$. We consider Q^* to be a vector in \mathbb{R}^D with components $Q_t^*(s, a)$. We also frequently use the notation $Q_t^* \in \mathbb{R}^d$ for some $t \leq T$, by which we mean Q^* restricted to the components $Q_t^*(s, a)$ for all $(s, a) \in \mathcal{U}$. We adopt this system for any vector in \mathbb{R}^D (e.g., \bar{Q}^n , $u^{i,*}$, and $\bar{u}^{i,n}$ to be defined later). The norms used in this paper are $\|\cdot\|_1$, $\|\cdot\|_2$, and $\|\cdot\|_\infty$, the l_1 -norm, the Euclidean norm, and the maximum norm, respectively. Example usages of the latter two are

$$\|Q_t^*\|_2 = \left(\sum_{(s,a) \in \mathcal{U}} Q_t^*(s, a)^2\right)^{1/2} \quad \text{and} \quad \|Q_t^*\|_\infty = \max_{(s,a) \in \mathcal{U}} |Q_t^*(s, a)|.$$

The following naming convention is used throughout the paper and appendix: stochastic processes denoted using ϵ , i.e., $\epsilon_{t+1}^{q,n}$, $\epsilon_{t+1}^{i,n}$, and $\epsilon_{t+1}^{h,n}$ are *conditionally unbiased* noise sequences and represent Monte Carlo sampling error. On the other hand, the processes denoted using ξ , i.e., $\xi_{t+1}^{q,n}$, $\xi_{t+1}^{i,n}$, and $\xi_{t+1}^{h,n}$, are *biased* noise and represent approximation error from using a value function approximation. For a vector v , $\text{diag}(v)$ is the diagonal matrix whose entries are the components of v . Lastly, for a nonnegative function $f: \mathbb{R} \rightarrow \mathbb{R}$, its *support* is represented by the notation $\text{supp}(f) = \{x \in \mathbb{R}: f(x) > 0\}$.

4. Algorithm

In this section, we introduce the risk-averse ADP algorithm for dynamic quantile-based risk measures, which aims to approximate the value function Q^* in order to produce near-optimal policies.

4.1. Overview of the Main Idea

Like most ADP and reinforcement learning algorithms, the algorithm that we develop in this paper to solve (2) is based on the recursive relationship of (3). The basic structure for the algorithm is a time-dependent version of Q-learning or approximate value iteration (see Powell [49, Chapter 10] for a discussion). Recall the form of the QBRM:

$$\rho_t^\alpha(X) = \mathbf{E}\left[\Phi(X, q_t^{\alpha_1}(X), q_t^{\alpha_2}(X), \dots, q_t^{\alpha_m}(X)) \mid \mathcal{F}_t\right].$$

The main idea of our approach is to approximate the quantiles $q_t^{\alpha_i}(X)$ and then combine the approximations to form an estimate of the risk measure. In essence, every observation of the exogenous information process (real or simulated data) can be utilized to give an updated approximation of each of the m quantiles. A second step then takes the observation and the quantile approximations to generate a refined approximation of the optimal value function Q^* . This type of logic is implemented using many concurrent *stochastic gradient* (Robbins and Monro [54], Kushner and Yin [40]) steps within a framework that walks through a single forward trajectory of states and actions on each iteration.

It turns out that there is a convenient characterization of the quantile through the so-called CVaR *minimization formula*. Given a real-valued, integrable random variable X , a risk level $\alpha_i \in (0, 1)$, and $u \in \mathbb{R}$, Rockafellar and Uryasev [56] proves that

$$q^{\alpha_i}(X) \in \arg \min_{u \in \mathbb{R}} \mathbf{E}\left[u + \frac{1}{1 - \alpha_i}(X - u)^+\right]. \quad (4)$$

Although the main result of Rockafellar and Uryasev [56] is that the optimal value of the optimization problem gives the CVaR $^{\alpha_i}(X)$, the characterization of the quantile as the minimizer is particularly useful for our purposes. It suggests the use of *stochastic approximation* or *stochastic gradient descent* algorithms (Robbins and Monro [54], Kushner and Yin [40]) to iteratively optimize (4).

With this intuition in mind, let us move back to the context of the MDP and define the auxiliary variables $u^{i,*} \in \mathbb{R}^D$, for each $i \in \{1, 2, \dots, m\}$, to be the α_i -quantiles of the future costs (recall that the quantiles α_i are defined as an argument to our QBRM in Definition 1). The component at time t and state (s, a) is

$$u_t^{i,*}(s, a) = q^{\alpha_i} \left(c_t(s, a, W_{t+1}) + \min_{a' \in \mathcal{A}_{S_{t+1}}} Q_{t+1}^*(S_{t+1}, a') \right), \quad (5)$$

for each $i \in \{1, 2, \dots, m\}$. Using (1), this allows us to take advantage of the equation

$$Q_t^*(s, a) = \mathbf{E} \left[\Phi(c_t(s, a, W_{t+1}) + \min_{a' \in \mathcal{A}_{S_{t+1}}} Q_{t+1}^*(S_{t+1}, a'), u_t^{1,*}(s, a), \dots, u_t^{m,*}(s, a)) \right]. \quad (6)$$

The relationship between Q^* and $u^{i,*}$ is fundamental to our algorithmic approach, which keeps track of mutually dependent approximations $\{\bar{u}^{i,n}\}$ and $\{\bar{Q}^n\}$ to the optimal values $u^{i,*}$ and Q^* , respectively.

4.2. The Dynamic-QBRM ADP Algorithm

Before discussing the details, we need some additional notation. Clearly, at each time t , the random quantity with which we are primarily concerned (and attempt to approximate) is the future cost given the optimal value function Q^* . Thus, we explicitly define its distribution function for every (s, a) :

$$F_t(x | s, a) = \mathbf{P} \left[c_t(s, a, W_{t+1}) + \min_{a' \in \mathcal{A}_{S_{t+1}}} Q_{t+1}^*(S_{t+1}, a') \leq x \right].$$

Recall that d is the cardinality of the state-action space. Next, suppose $\bar{u}_t^i \in \mathbb{R}^d$ is an approximation of $u_t^{i,*}$ and for each t and $i \in \{1, 2, \dots, m\}$, define the stochastic gradient mapping $\psi_t^i: \mathbb{R}^d \times \mathbb{R}^d \times \mathcal{W} \rightarrow \mathbb{R}^d$ to perform the stochastic gradient computation:

$$\psi_t^i(\bar{u}_t^i, \bar{Q}_{t+1}, W_{t+1})(s, a) = 1 - \frac{1}{1 - \alpha_i} \mathbf{1} \left\{ c_t(s, a, W_{t+1}) + \min_{a' \in \mathcal{A}_{S_{t+1}}} \bar{Q}_{t+1}(S_{t+1}, a') \geq \bar{u}_t^i(s, a) \right\}. \quad (7)$$

To avoid confusion, we note that this is the stochastic gradient associated with the minimization formula (4) for CVaR, but this step is necessary for any QBRM, even if we are not utilizing CVaR.

The second piece of notation we need is a specialized, stochastic version of the Bellman operator to the risk-averse case: for each t , we define the mapping $H_t: \mathbb{R}^d \times \dots \times \mathbb{R}^d \times \mathcal{W} \rightarrow \mathbb{R}^d$, with $m + 1$ arguments in \mathbb{R}^d , to represent an approximation of the term within the expectation of (6):

$$H_t(\bar{u}_t^1, \dots, \bar{u}_t^m, \bar{Q}_{t+1}, W_{t+1})(s, a) = \Phi \left(c_t(s, a, W_{t+1}) + \min_{a' \in \mathcal{A}_{S_{t+1}}} \bar{Q}_{t+1}(S_{t+1}, a'), \bar{u}_t^1(s, a), \dots, \bar{u}_t^m(s, a) \right).$$

Therefore, (6) can be rewritten using the stochastic Bellman operator H_t by replacing all approximate quantities with their true values:


$$Q_t^*(s, a) = \mathbf{E} \left[H_t(u_t^{1,*}, \dots, u_t^{m,*}, Q_{t+1}^*, W_{t+1})(s, a) \right]. \quad (8)$$

The Dynamic-QBRM ADP algorithm that we describe in the next section consists of both outer and inner iterations: for each outer iteration n , we step through the entire time horizon of the problem $t \in \mathcal{T}$. At time t , iteration n , the relevant quantities for our algorithms are a state-action pair $(S_t^n, a_t^n) \in \mathcal{U}$ and two samples $W_{t+1}^{u,n}, W_{t+1}^{q,n} \in \mathcal{W}$ from the distribution of W_{t+1} corresponding to the “two steps” of our algorithm, one for approximating the auxiliary variables $u^{i,*}$ and the second for approximating the value function Q^* . Figure 1 illustrates the main idea behind the algorithm: we merge the results of m adaptive minimizations of (4), corresponding to estimates of the m quantiles, into an estimate of the optimal value function, \bar{Q}_t^n . The estimate \bar{Q}_t^n is then used to produce estimates of the relevant quantities for the previous time period. Note that the m objective functions shown in the figure differ only in their risk levels α_i . The arrows on the curves indicate that the minimizations are achieved via gradient descent steps.

Now that we are in an algorithmic setting, we consider a new probability space $(\Omega, \mathcal{G}, \mathbf{P})$, where $\mathcal{G} = \sigma\{(S_t^n, a_t^n, W_{t+1}^{u,n}, W_{t+1}^{q,n}), n \geq 0, t \leq T\}$. In order to describe the history of the algorithms, we define:

$$\mathcal{G}_t^n = \sigma\{(S_\tau^k, a_\tau^k, W_\tau^{u,k}, W_\tau^{q,k}), k < n, \tau \leq T\} \cup \{(S_\tau^n, a_\tau^n, W_\tau^{u,n}, W_\tau^{q,n}), \tau \leq t\},$$

for $t \in \mathcal{T}$ and $n \geq 1$, with $\mathcal{G}_t^0 = \{\emptyset, \Omega\}$ for all $t \leq T$. We therefore have a filtration that obeys $\mathcal{G}_t^n \subseteq \mathcal{G}_{t+1}^n$ for $t \in \mathcal{T}$ and $\mathcal{G}_T^n \subseteq \mathcal{G}_0^{n+1}$, coinciding precisely with the progression of the algorithm. The random variables (S_t^n, a_t^n) are

Figure 1. (Color online) Main algorithm idea.

generated according to some sampling policy (to be discussed later) while $W_{t+1}^{u,n}$ and $W_{t+1}^{q,n}$ are generated from the distribution of the exogenous process W_{t+1} .

Crucial to many ADP algorithms is the *stepsize* (or *learning rate*). In our case, we use $\{\gamma_t^n\}$ and $\{\eta_t^n\}$ for smoothing new observations with previous estimates, where $\gamma_t^n, \eta_t^n \in \mathbb{R}^d$ for each t and n and are \mathcal{G}_t^n -measurable. The stepsize γ_t^n is used to update our approximation of $u_i^{1,*}$ while the stepsize η_t^n is used to update Q_t^* ; see Algorithm 1. We articulate the asynchronous nature of our algorithm by imposing the following condition on the stepsizes (included in Assumption 1 of Section 5):

$$(s, a) \neq (S_t^n, a_t^n) \implies \gamma_t^n(s, a) = \eta_t^n(s, a) = 0,$$

which causes updates to only happen for states that we actually visit.

Algorithm 1 (Dynamic-QBRM ADP)

Input: Initial estimates of the value function $\bar{Q}^0 \in \mathbb{R}^D$ and quantiles $\bar{u}^{i,0} \in \mathbb{R}^D$ for $i \in \{1, 2, \dots, m\}$.

Stepsize rules γ_t^n and η_t^n for all n, t .

Output: Approximations $\{\bar{Q}^n\}$ and $\{\bar{u}^{i,n}\}$.

- 1 Set $\bar{Q}_T^n = 0$ for all n .
- for** $n = 1, 2, \dots$ **do**
- 2 Choose an initial state (S_0^n, a_0^n) .
- for** $t = 0, 1, \dots, T-1$ **do**
- 3 Draw samples of the information process $W_{t+1}^{u,n}, W_{t+1}^{q,n} \in \mathcal{W}$.
- 4 Update auxiliary variable approximations for $i = 1, \dots, m$:

$$\bar{u}_t^{i,n} = \Pi_{\mathcal{X}_t^u} \{ \bar{u}_t^{i,n-1} - \text{diag}(\gamma_t^n) \psi_t^i(\bar{u}_t^{i,n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{u,n}) \}.$$
- 5 Compute an estimate of the future cost based on the current approximations:

$$\hat{q}_t^n = H_t(\bar{u}_t^{1,n-1}, \dots, \bar{u}_t^{m,n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{q,n}).$$
- 6 Update approximation of value function:

$$\bar{Q}_t^n = \Pi_{\mathcal{X}_t^q} \{ \bar{Q}_t^{n-1} - \text{diag}(\eta_t^n)(\bar{Q}_t^{n-1} - \hat{q}_t^n) \}.$$
- 7 Choose next state (S_{t+1}^n, a_{t+1}^n) .
- end**
- end.**

Stochastic approximation theory often requires a projection step onto a compact set (giving bounded iterates) to ensure convergence (Kushner and Yin [40]). Hence, for each t and (s, a) , let $\mathcal{X}_t^u(s, a) \subseteq \mathbb{R}$ and $\mathcal{X}_t^q(s, a) \subseteq \mathbb{R}$ be compact intervals and let

$$\mathcal{X}_t^u = \{x \in \mathbb{R}^d : x(s, a) \in \mathcal{X}_t^u(s, a)\} \quad \text{and} \quad \mathcal{X}_t^q = \{x \in \mathbb{R}^d : x(s, a) \in \mathcal{X}_t^q(s, a)\},$$

be our projection sets at time t . The Euclidean projection operator to a set \mathcal{X} is given by the usual definition:

$$\Pi_{\mathcal{X}}(y) = \arg \min_{x \in \mathcal{X}} \|y - x\|_2^2.$$

These sets may be chosen arbitrarily large in practice and our first theoretical result (almost sure convergence) will continue to hold. However, there is a tradeoff: If, in addition, we want our convergence *rate* results to hold, then these sets also cannot be too large (see Assumption 3).

The precise steps of Dynamic-QBRM ADP are given in Algorithm 1. A main characteristic of the algorithm is that sequences $\{\bar{u}^{i,n}\}$ and $\{\bar{Q}^n\}$ are intertwined (i.e., they depend on each other). Consequently, there are multiple levels of approximation being used throughout the steps of the algorithm. The theoretical results of the subsequent sections shed light onto these issues.

5. Analysis of Convergence

In this section, we state and prove convergence theorems for Algorithm 1. First, we give an overview of our analysis and the relationship to existing work.

5.1. A Preview of Results

The two main results of this section are: (1) the almost sure convergence of Dynamic-QBRM ADP, and (2) a convergence rate result under a particular sampling policy called ε -greedy. The proof of almost sure convergence uses techniques from the stochastic approximation literature (Kushner and Yin [40]), which were applied to the field of reinforcement learning and Q-learning in Tsitsiklis [72], Jaakkola et al. [31] and Bertsekas and Tsitsiklis [9]. However, our algorithm differs from risk-neutral Q-learning in that it tracks multiple quantities, $\bar{u}_t^{1,n}, \bar{u}_t^{2,n}, \dots, \bar{u}_t^{m,n}, \bar{Q}_t^n$, over a horizon $1, 2, \dots, T$. The intuition behind the proof is that multiple “stochastic approximation instances” are pasted together in order to obtain overall convergence of all relevant quantities. Accordingly, the interdependence of various approximations means that in several parts of the proof, we require careful analysis of *biased noise terms* (or *approximation error*) in addition to *unbiased statistical error*. See, e.g., Kearns and Singh [36], Even-Dar and Mansour [26] and Azar et al. [5], for convergence rate results for standard Q-learning. The proof technique used to analyze the high probability convergence rate of risk-neutral Q-learning in Even-Dar and Mansour [26] is based on the same types of stochastic approximation results that we utilize in this paper.

Let us now make a few remarks regarding some simplifying assumptions made in this paper. As proven in Rockafellar and Uryasev [57, Theorem 10], the set of minimizers $\arg \min_{u \in \mathbb{R}} \mathbb{E}[u + (1 - \alpha_i)^{-1}(X - u)^+]$ is a nonempty, closed, and bounded interval for a general X . We shall for ease of presentation, however, make assumptions (strictly increasing and continuous cdf, Assumption 2(iii)) to guarantee that $u_t^{i,*}(s, a)$ is the unique minimizer when X is the optimal future cost $c_t(s, a, W_{t+1}) + \min_{a' \in \mathcal{A}_{S_{t+1}}} Q_{t+1}^*(S_{t+1}, a')$ and that gradient computations to remain valid. This assumption is sufficient for almost sure convergence (Theorem 2). To further examine the convergence rate of the algorithm (Theorem 3), we must additionally have Assumption 3, which states that the density of the future cost exists and is positive within the constraint sets \mathcal{X}_t^n —this provides us the technical condition of strong convexity (discussed more in Section 5.3 below).

Since $\min_{a' \in \mathcal{A}_{S_{t+1}}} Q_{t+1}^*(S_{t+1}, a')$ has a discrete distribution, the assumptions hold only in certain situations: an obvious case is when the current stage cost has a density and is independent of S_{t+1} . For example, such a property holds when W_{t+1} can be written as two independent components (W_{t+1}^1, W_{t+1}^2) where the current stage cost depends on W_{t+1}^1 and the downstream state depends on W_{t+1}^2 . This model is relevant in a number of applications; notable examples include multi-armed bandits (Whittle [74]), shortest path problems with random edge costs (Ryzhov and Powell [64]), trade execution with temporary (and temporally independent) price impact (Bertimas and Lo [8]), and energy trading in two-settlement markets (Löhndorf et al. [41]). Small algorithmic extensions (requiring more complex notation) to handle the general case are possible, but the fundamental concepts would remain unchanged. Hence, we will assume the cleaner setting for the purposes of this paper.

5.2. Almost Sure Convergence

First, we discuss the necessary algorithmic assumptions, many of which are standard to the field of stochastic approximation.

Assumption 1. For all $(s, a) \in \mathcal{U}$ and $t \in \mathcal{T}$, suppose the following are satisfied:

- (i) $\gamma_t^n(s, a) = \tilde{\gamma}_t^{n-1} \mathbf{1}_{\{(s, a) = (S_t^n, a_t^n)\}}$, for some $\tilde{\gamma}_t^{n-1} \in \mathbb{R}$ that is \mathcal{G}_t^{n-1} -measurable,
- (ii) $\eta_t^n(s, a) = \tilde{\eta}_t^{n-1} \mathbf{1}_{\{(s, a) = (S_t^n, a_t^n)\}}$, for some $\tilde{\eta}_t^{n-1} \in \mathbb{R}$ that is \mathcal{G}_t^{n-1} -measurable,
- (iii) $\sum_{n=1}^{\infty} \gamma_t^n(s, a) = \infty$, $\sum_{n=1}^{\infty} \eta_t^n(s, a) = \infty$ a.s.,
- (iv) $\sum_{n=1}^{\infty} \gamma_t^n(s, a)^2 < \infty$, $\sum_{n=1}^{\infty} \eta_t^n(s, a)^2 < \infty$ a.s.,
- (v) $\exists \varepsilon > 0$, such that state sampling policy satisfies

$$\mathbf{P}((S_t^n, a_t^n) = (s, a) \mid \mathcal{G}_{t-1}^n) \geq \varepsilon \quad \text{and} \quad \mathbf{P}((S_0^n, a_0^n) = (s, a) \mid \mathcal{G}_T^{n-1}) \geq \varepsilon,$$

- (vi) the projection sets are chosen large enough so that $u_t^{i,*} \in \mathcal{X}_t^u$ for each $i \in \{1, 2, \dots, m\}$ and $Q_t^* \in \mathcal{X}_t^Q$.

Assumption 1(i) and 1(ii) represent the asynchronous nature of the algorithm, sending the stepsize to zero whenever a state is not visited, while 1(iii) and 1(iv) are standard conditions on the stepsize. Assumption 1(v) is an *exploration* requirement; by the Extended Borel-Cantelli Lemma (see Breiman [12]), sampling with this exploration requirement guarantees that we will visit every state infinitely often with probability one. In particular, for the case with an ε -greedy sampling policy (i.e., explore with probability ε , follow current policy otherwise), this assumption holds. We discuss this policy in greater detail in Section 5. Part (vi) is a technical assumption. The second group of assumptions that we present are related to the problem parameters.

Assumption 2. *The following hold:*

(i) *the risk-aversion function $\Phi: \mathbb{R}^{m+1} \rightarrow \mathbb{R}$ (from the QBRM within the one-step conditional risk measure ρ_t^α) is Lipschitz continuous with constant $L_\Phi > 0$, i.e., for all $v, w \in \mathbb{R}^{m+1}$, $|\Phi(v) - \Phi(w)| \leq L_\Phi \|v - w\|_1$,*

(ii) *$\exists C_{\max} > 0$ such that $\mathbf{E}[c_t(s, a, W_{t+1})^2] \leq C_{\max}$ for all $(s, a) \in \mathcal{U}$ and $t \in \mathcal{T}$,*

(iii) *the distribution function $F_t(x | s, a)$ is strictly increasing and Lipschitz continuous with constant $L_F > 0$, i.e.,*

$$|F_t(x | s, a) - F_t(y | s, a)| \leq L_F |x - y|,$$

for all $x, y \in \mathbb{R}$, $(s, a) \in \mathcal{U}$, and $t \in \mathcal{T}$.

As an example for Assumption 2(i), in the case of the QBRM being CVaR, where $\Phi(X, q) = q + (1 - \alpha)^{-1}(X - q)^+$, it is easy to see that we can take $L_\Phi = 1 + 1/(1 - \alpha)$. Assumption 2(ii) states that the second moment of the cost function is bounded. Assumptions 2(iii) and 1(vi) together imply that $u_t^{i,*}(s, a)$ is the unique $u \in \mathcal{X}_t^u(s, a)$ such that $F_t(u | s, a) = \alpha_i$.

Lemma 1. *Under Assumptions 1 and 2, if $\bar{Q}_{t+1}^n \rightarrow Q_{t+1}^*$ almost surely, then the sequences of iterates $\bar{u}_t^{i,n}$ generated in Step 4 of Algorithm 1 satisfy $\bar{u}_t^{i,n} \rightarrow u_t^{i,*}$ almost surely for each $i \in \{1, 2, \dots, m\}$.*

We distinguish between two types of noise sequences (in n) for each $t \in \mathcal{T}$, the *statistical error* and the *approximation error*, denoted by $\epsilon_{t+1}^{i,n} \in \mathbb{R}^d$ and $\xi_{t+1}^{i,n} \in \mathbb{R}^d$, respectively. The definitions are

$$\begin{aligned} \epsilon_{t+1}^{i,n} &= \psi_t^i(\bar{u}_t^{i,n-1}, Q_{t+1}^*, W_{t+1}^{u,n}) - \mathbf{E}[\psi_t^i(\bar{u}_t^{i,n-1}, Q_{t+1}^*, W_{t+1})], \\ \xi_{t+1}^{i,n} &= \psi_t^i(\bar{u}_t^{i,n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{u,n}) - \psi_t^i(\bar{u}_t^{i,n-1}, Q_{t+1}^*, W_{t+1}^{u,n}), \end{aligned}$$

and we see that the random variable $\epsilon_{t+1}^{i,n}$ represents the error that the sample gradient deviates from its mean, computed using the *true future cost* distribution (i.e., assuming we have Q_{t+1}^*). On the other hand, $\xi_{t+1}^{i,n}$ is the error between the two evaluations of ψ_t^i given the same sample $W_{t+1}^{u,n}$, due only to the difference between \bar{Q}_{t+1}^{n-1} and Q_{t+1}^* . Rearranging, we have

$$\psi_t^i(\bar{u}_t^{i,n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{u,n}) = \mathbf{E}[\psi_t^i(\bar{u}_t^{i,n-1}, Q_{t+1}^*, W_{t+1})] + \epsilon_{t+1}^{i,n} + \xi_{t+1}^{i,n},$$

which implies that the update given in Step 4 of Algorithm 1 can be rewritten as

$$\bar{u}_t^{i,n} = \Pi_{\mathcal{X}_t^u} \left\{ \bar{u}_t^{i,n-1} - \text{diag}(\gamma_t^n) \left[\mathbf{E}[\psi_t^i(\bar{u}_t^{i,n-1}, Q_{t+1}^*, W_{t+1})] + \epsilon_{t+1}^{i,n} + \xi_{t+1}^{i,n} \right] \right\}.$$

Note that the term in the square brackets is a *biased stochastic gradient* and observe that it is bounded (since ψ_t^i only takes two finite values). For the present inductive step at time t , let us fix a state (s, a) . It now becomes convenient for us to view $\bar{u}_t^{i,n}(s, a)$ as a stochastic process in n , adapted to the filtration $\{\mathcal{G}_{t+1}^n\}_{n \geq 0}$ (since $\mathcal{G}_{t+1}^n \subseteq \mathcal{G}_{t+1}^{n+1} \subseteq \mathcal{G}_{t+1}^{n+2} \dots$). It is clear that by the definition of $\epsilon_{t+1}^{i,n}(s, a)$:

$$\mathbf{E}[\epsilon_{t+1}^{i,n}(s, a) | \mathcal{G}_{t+1}^{n-1}] = 0 \quad a.s. \quad (9)$$

Therefore, $\epsilon_{t+1}^{i,n}(s, a)$ are unbiased increments that can be referred to as *martingale difference noise*. Before continuing, notice the following useful fact:

$$\min_{a \in \mathcal{A}_s} Q_{t+1}^*(s, a) - \min_{a \in \mathcal{A}_s} \bar{Q}_{t+1}^{n-1}(s, a) \leq \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_\infty. \quad (10)$$

The proof follows from $\min v = -\max(-v)$ and $\max v - \max w \leq \max |v - w|$, where the minimum and maximum are taken over the components of some vectors v and w . Now, let $S_{t+1}^n = S^M(s, a, W_{t+1}^{u,n})$. Expanding the definition of $\xi_{t+1}^{i,n}(s, a)$ and using (10), we obtain

$$\begin{aligned} \xi_{t+1}^{i,n}(s, a) &= \frac{1}{1 - \alpha_i} \left[\mathbf{1} \left\{ c_t(s, a, W_{t+1}^{u,n}) + \min_{a' \in \mathcal{A}_{S_{t+1}^n}} Q_{t+1}^*(S_{t+1}^n, a') \geq \bar{u}_t^{i,n-1}(s, a) \right\} \right. \\ &\quad \left. - \mathbf{1} \left\{ c_t(s, a, W_{t+1}^{u,n}) + \min_{a' \in \mathcal{A}_{S_{t+1}^n}} \bar{Q}_{t+1}^{n-1}(S_{t+1}^n, a') \geq \bar{u}_t^{i,n-1}(s, a) \right\} \right] \\ &\leq \frac{1}{1 - \alpha_i} \left[\mathbf{1} \left\{ c_t(s, a, W_{t+1}^{u,n}) + \min_{a' \in \mathcal{A}_{S_{t+1}^n}} Q_{t+1}^*(S_{t+1}^n, a') \geq \bar{u}_t^{i,n-1}(s, a) \right\} \right. \\ &\quad \left. - \mathbf{1} \left\{ c_t(s, a, W_{t+1}^{u,n}) + \min_{a' \in \mathcal{A}_{S_{t+1}^n}} Q_{t+1}^*(S_{t+1}^n, a') \geq \bar{u}_t^{i,n-1}(s, a) + \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_\infty \right\} \right]. \end{aligned}$$

Using the shorthand $F_t(\cdot) = F_t(\cdot | s, a)$ and taking the conditional expectation of both sides, we get (almost surely)

$$|\mathbf{E}[\xi_{t+1}^{i,n}(s, a) | \mathcal{G}_{t+1}^{n-1}]| \leq \frac{1}{1 - \alpha_i} |F_t(\bar{u}_t^{i,n-1}(s, a) + \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_\infty) - F_t(\bar{u}_t^{i,n-1}(s, a))| \leq \frac{L_F}{1 - \alpha_i} \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_\infty, \quad (11)$$

where the second inequality follows by Assumption 2(iii). Since we assumed in the statement of the lemma that $\bar{Q}_{t+1}^n \rightarrow Q_{t+1}^*$ almost surely, it must be the case that

$$\mathbf{E}[\xi_{t+1}^{i,n}(s, a) | \mathcal{G}_{t+1}^{n-1}] \rightarrow 0 \quad a.s., \quad (12)$$

and hence the noise “vanishes asymptotically” in expectation. Now, given the boundedness of ψ_i^i (and hence, finite second moment of ψ_i^i), the unbiasedness property (9), the vanishing noise property (12), the stepsize and sampling properties of Assumption 1, and the uniqueness of $u_i^*(s, a)$ from Assumption 2(iii), we can apply a classical theorem of stochastic approximation, Kushner and Yin [40, Theorem 2.4], to conclude that $\bar{u}_t^{i,n}(s, a) \rightarrow u_i^{i,*}(s, a)$ almost surely for each i . Because we chose an arbitrary (s, a) , this convergence holds for all $(s, a) \in \mathcal{U}$. \square

Theorem 2 (Almost Sure Convergence). Choose initial approximations $\bar{Q}^0 \in \mathbb{R}^D$ and $\bar{u}^{i,0} \in \mathbb{R}^D$ for each $i \in \{1, 2, \dots, m\}$ so that $\bar{Q}_t^0 \in \mathcal{X}_t^q$ and $\bar{u}_t^{i,0} \in \mathcal{X}_t^u$ for all $t \in \mathcal{T}$. Under Assumptions 1–2, Algorithm 1 generates a sequence of iterates \bar{Q}^n that converge almost surely to the optimal value function Q^* .

Sketch of Proof. The idea of the proof is to induct backwards on t and repeatedly apply Lemma 1. At each step, the induction hypothesis is that $\bar{Q}_{t+1}^n \rightarrow Q_{t+1}^*$ from which we obtain convergence of $\bar{u}_t^{i,n} \rightarrow u_i^{i,*}$ from Lemma 1. After making sure certain technical details are satisfied, stochastic approximation theory allows us to show that $\bar{Q}_t^n \rightarrow Q_t^*$. The full details are given in Appendix A. \square

5.3. Convergence Rate

In this section, we discuss the convergence rate (in terms of the expected deviation to Q^*) of the procedure described in Algorithm 1. Because we are working with an asynchronous algorithm where only one state per time period is visited every iteration, it is necessary for us to specify the assumed *state sampling policy* for visiting states before deriving convergence rate results. Due to Assumption 1(i)–1(ii), we naturally must consider the sampling policy and the stepsize sequences jointly, as is done in the upcoming proposition. We employ the ε -greedy policy, a popular choice that balances exploitation and exploration using a tunable parameter ε ; it is defined as follows. For any iteration $n > 0$, choose (S_0^n, a_0^n) independently and uniformly at random. At iteration n , time $t > 0$, let

$$s_{t+1}^n = S^M(S_t^n, a_t^n, W_{t+1}^n) \quad \text{and} \quad a_{t+1}^n = \arg \min_{a \in \mathcal{A}} \bar{Q}_{t+1}^{n-1}(s_{t+1}^n, a)$$

and define a Bernoulli random variable X_{t+1}^n with parameter $1 - \varepsilon d$ that is independent from \mathcal{G}_t^n (i.e., the rest of the system). The next state to visit is selected by the rule

$$(S_{t+1}^n, a_{t+1}^n) = \begin{cases} (s_{t+1}^n, a_{t+1}^n), & \text{if } X_{t+1}^n = 1, \\ \text{choose uniformly over } \mathcal{U}, & \text{otherwise.} \end{cases} \quad (13)$$

Note that Assumption 1(v) is clearly satisfied as each state is visited with probability at least ε . We also choose our stepsize sequences with $\tilde{\gamma}_t^{n-1}$ and $\tilde{\eta}_t^{n-1}$ as *deterministic harmonic sequences* (for ease of analysis) so that

$$\gamma_t^n(s, a) = \frac{\gamma_t^n}{n} \mathbf{1}_{\{(s,a)=(S_t^n, a_t^n)\}} \quad \text{and} \quad \eta_t^n(s, a) = \frac{\eta_t^n}{n} \mathbf{1}_{\{(s,a)=(S_t^n, a_t^n)\}}, \quad (14)$$

where $\gamma_t > 0$ and $\eta_t > 0$ are deterministic, time-dependent constants. Since neither $\tilde{\gamma}_t^{n-1} = \gamma_t^n/n$ nor $\tilde{\eta}_t^n = \eta_t^n/n$ depend on the history of visited states, they are known as *centralized learning rates* (see Szepesvari and Littman [70] for a discussion). The main difficulty for the centralized stepsizes of (14) is that when the frequency of visits to a state decays quickly enough, then Assumption 1(iii) and 11(iv) may not hold (for example, consider when a state is only visited on iteration numbers that are powers of two). In fact, it is often not immediately obvious when Assumption 1(iii) and 1(iv) are satisfied. The next proposition shows that under the ε -greedy sampling policy, the states are visited often enough that the assumption remains satisfied.

Proposition 1. *Under the ε -greedy sampling policy given in (13) and the deterministic harmonic stepsizes given in (14), Assumption 1(iii)–1(iv) is satisfied.*

Proof. See Appendix A. \square

For a technical reason needed to prove the convergence rate results, we need to constrain the iterates of the algorithm to a region *within* the support of the distribution of future costs (stated formally in Assumption 3 below). We shall see that when this assumption is satisfied, we get a notion of *strong convexity*. To be more precise, observe that $\partial^2 \mathbf{E}[u + (X - u)^+ / (1 - \alpha_i)] / \partial u^2 = f_X(u) / (1 - \alpha)$, where f_X is the density of X . If u is constrained to be within the support of the distribution of X (i.e., where $f_X > 0$), then we are able to lower bound the second derivative by some constant (that depends on the constraint set), thereby attaining strong convexity within the region. This is useful for deriving the convergence rate results (see Lemma 2); unfortunately, this condition is in general difficult to check in practice.

Assumption 3. *The density $f_t(x | s, a) = dF_t(x | s, a) / dx$ exists and the stochastic approximation projection set is within the support of the density: $\mathcal{X}_t^u(s, a) \subseteq \text{supp}(f_t(\cdot | s, a))$ for all $(s, a) \in \mathcal{U}$ and $t \in \mathcal{T}$. Let*

$$l_f = \min_{(s, a, t)} \min\{x \in \mathcal{X}_t^u(s, a): f_t(x | s, a)\}$$

be a positive real number that lower bounds the density function over all t and $(s, a) \in \mathcal{U}$.

We now provide a few lemmas that will be useful in establishing the final result. The first lemma relates the error of $\bar{u}_t^{i, n}$ to the error in the last iteration (i.e., of $\bar{u}_t^{i, n-1}$) and the error of the value function in the next time period (i.e., of \bar{Q}_{t+1}^{n-1}).

Lemma 2. *Under Assumptions 1(vi)–3, the ε -greedy sampling policy of (13), and the deterministic harmonic stepsizes given in (14), the sequence of approximations $\bar{u}_t^{i, n}$ generated by Algorithm 1 satisfies, for any $\kappa > 0$,*

$$\mathbf{E}[\|\bar{u}_t^{i, n} - u_t^{i, *}\|_2^2] \leq \left[1 - \frac{\gamma_t}{n}(2\varepsilon C_{l_f} - \kappa C_{L_f})\right] \mathbf{E}[\|\bar{u}_t^{i, n-1} - u_t^{i, *}\|_2^2] + \frac{\gamma_t}{n\kappa} \mathbf{E}[\|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_2^2] + C_{\alpha_i} \frac{\gamma_t^2}{n^2},$$

where $C_{l_f} = l_f / (1 - \alpha_i)$, $C_{L_f} = L_f^2 / (1 - \alpha_i)^2$, and $C_{\alpha_i} = [\max(1, \alpha_i / (1 - \alpha_i))]^2$.

Proof. Fix an i . Recall the update equation given in Step 5 of Algorithm 1:

$$\bar{u}_t^{i, n} = \Pi_{\mathcal{X}_t^u} \left\{ \bar{u}_t^{i, n-1} - \text{diag}(\gamma_t^n) \psi_t^i(\bar{u}_t^{i, n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{u, n}) \right\}.$$

Using the non-expansive property of the projection operator, we have

$$\begin{aligned} \|\bar{u}_t^{i, n} - u_t^{i, *}\|_2^2 &= \|\Pi_{\mathcal{X}_t^u} \left\{ \bar{u}_t^{i, n-1} - \text{diag}(\gamma_t^n) \psi_t^i(\bar{u}_t^{i, n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{u, n}) \right\} - \Pi_{\mathcal{X}_t^u} \{u_t^{i, *}\}\|_2^2 \\ &\leq \|\bar{u}_t^{i, n-1} - u_t^{i, *} - \text{diag}(\gamma_t^n) \psi_t^i(\bar{u}_t^{i, n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{u, n})\|_2^2 \\ &\leq \|\bar{u}_t^{i, n-1} - u_t^{i, *}\|_2^2 + C_{\alpha_i} \frac{\gamma_t^2}{n^2} - 2(\bar{u}_t^{i, n-1} - u_t^{i, *})^\top \text{diag}(\gamma_t^n) \psi_t^i(\bar{u}_t^{i, n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{u, n}). \end{aligned} \quad (15)$$

Recall from the proof of Lemma 1 that we can write

$$\psi_t^i(\bar{u}_t^{i, n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{u, n}) = \mathbf{E}[\psi_t^i(\bar{u}_t^{i, n-1}, Q_{t+1}^*, W_{t+1})] + \epsilon_{t+1}^{i, n} + \xi_{t+1}^{i, n},$$

from which we see that the cross-term $(\bar{u}_t^{i, n-1} - u_t^{i, *})^\top \text{diag}(\gamma_t^n) \psi_t^i(\bar{u}_t^{i, n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{u, n})$ in the chain of inequalities (15) can be expanded into three terms. These terms can now be analyzed separately. By (7) and Assumption 2(iv), we see that the derivative of $\mathbf{E}[\psi_t^i(u, Q_{t+1}^*, W_{t+1})]$ in $u(s, a)$ can be expressed as

$$\frac{\partial}{\partial u(s, a)} \mathbf{E}[\psi_t^i(u, Q_{t+1}^*, W_{t+1})](s, a) = \frac{f_t(u(s, a) | s, a)}{1 - \alpha_i} \geq \frac{l_f}{1 - \alpha_i}. \quad (16)$$

Since every state is visited with probability larger than ε , we know that $\mathbf{E}[\gamma_t^n(s, a) \mid \mathcal{G}_{t+1}^{n-1}] \geq \varepsilon \gamma_t/n$. Combining this with (16) and $\mathbf{E}[\psi_t^i(u_t^*, Q_{t+1}^*, W_{t+1})] = 0$, it follows that (almost surely)

$$\begin{aligned} & \mathbf{E}[(\bar{u}_t^{i,n-1} - u_t^{i,*})^\top \text{diag}(\gamma_t^n) \mathbf{E}[\psi_t^i(\bar{u}_t^{i,n-1}, Q_{t+1}^*, W_{t+1})] \mid \mathcal{G}_{t+1}^{n-1}] \\ & \geq \frac{\varepsilon \gamma_t}{n} (\bar{u}_t^{i,n-1} - u_t^{i,*})^\top [\mathbf{E}[\psi_t^i(\bar{u}_t^{i,n-1}, Q_{t+1}^*, W_{t+1})] - \mathbf{E}[\psi_t^i(u_t^*, Q_{t+1}^*, W_{t+1})]] \geq \frac{\varepsilon \gamma_t}{n} \frac{l_f}{1 - \alpha_i} \|\bar{u}_t^{i,n-1} - u_t^{i,*}\|_2^2. \end{aligned} \quad (17)$$

Note that $\gamma_t^n(s, a)$ depends on (S_t^n, a_t^n) and $\varepsilon_{t+1}^{i,n}(s, a)$ depends on $W_{t+1}^{u,n}$. By independence of (S_t^n, a_t^n) and $W_{t+1}^{u,n}$ and the unbiased property of $\varepsilon_{t+1}^{i,n}(s, a)$ of (9),

$$\mathbf{E}[(\bar{u}_t^{i,n-1} - u_t^{i,*})^\top \text{diag}(\gamma_t^n) \varepsilon_{t+1}^{i,n} \mid \mathcal{G}_{t+1}^{n-1}] = 0. \quad (18)$$

Moving on to the third term, using (11), the fact that γ_t^n has exactly one nonzero component, and the monotonicity of l_p norms, we can deduce

$$\begin{aligned} \mathbf{E}[-(\bar{u}_t^{i,n-1} - u_t^{i,*})^\top \text{diag}(\gamma_t^n) \xi_{t+1}^{i,n} \mid \mathcal{G}_{t+1}^{n-1}] & \leq \frac{\gamma_t}{n} \|\bar{u}_t^{i,n-1} - u_t^{i,*}\|_\infty \frac{L_F}{1 - \alpha_i} \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_\infty \\ & \leq \frac{\gamma_t}{n} \frac{L_F}{1 - \alpha_i} \|\bar{u}_t^{i,n-1} - u_t^{i,*}\|_2 \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_2. \end{aligned}$$

Using the inequality $2ab \leq a^2 \kappa + b^2/\kappa$ for $\kappa > 0$ on the above, we arrive at

$$\mathbf{E}[-2(\bar{u}_t^{i,n-1} - u_t^{i,*})^\top \text{diag}(\gamma_t^n) \xi_{t+1}^{i,n} \mid \mathcal{G}_{t+1}^{n-1}] \leq \frac{\gamma_t L_F^2 \kappa}{n(1 - \alpha_i)^2} \|\bar{u}_t^{i,n-1} - u_t^{i,*}\|_2^2 + \frac{\gamma_t}{n \kappa} \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_2^2. \quad (19)$$

Finally, the statement of the lemma follows by taking expectations of the inequalities (15), (17)–(19) and combining. \square

Similarly, the next lemma relates the error of the approximate value function \bar{Q}_t^n to the error in the last iteration (i.e., of \bar{Q}_t^{n-1}), the error of the value function in the next time period (i.e., of \bar{Q}_{t+1}^{n-1}), and the error of all of the approximate quantiles $\bar{u}_t^{i,n}$. Because the analysis is similar to that of Lemma 2, we relegate the proof to Appendix A.

Lemma 3. Under the same conditions as Lemma 2, the sequence of approximations \bar{Q}_t^n generated by Algorithm 1 satisfies, for any $\kappa_0, \kappa_1, \dots, \kappa_m > 0$,

$$\begin{aligned} \mathbf{E}[\|\bar{Q}_t^n - Q_t^*\|_2^2] & \leq \left[1 - \frac{\eta_t}{n} \left(2\varepsilon - L_\Phi \kappa_0 - L_\Phi \sum_{i=1}^m \kappa_i \right) \right] \mathbf{E}[\|\bar{Q}_t^{n-1} - Q_t^*\|_2^2] \\ & \quad + \frac{\eta_t L_\Phi}{n \kappa_0} \mathbf{E}[\|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_2^2] + \frac{\eta_t}{n} \sum_{i=1}^m \frac{L_\Phi}{\kappa_i} \mathbf{E}[\|\bar{u}_t^{i,n-1} - u_t^{i,*}\|_2^2] + C_H \frac{\eta_t^2}{n^2}, \end{aligned}$$

where C_H bounds the term

$$\mathbf{E}[\bar{Q}_t^{n-1}(s, a) - H_t(\bar{u}_t^{1,n-1}, \dots, \bar{u}_t^{m,n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1})(s, a)]^2 \leq C_H,$$


for all $(s, a) \in \mathcal{U}$.

Proof. See Appendix A. \square

With these preliminary results in mind, we move on to the theorem that states our $\mathcal{O}(1/n)$ convergence rate and provide a sketch of the proof.

Theorem 3 (Convergence Rate). Choose initial approximations $\bar{Q}^0 \in \mathbb{R}^D$ and $\bar{u}^{i,0} \in \mathbb{R}^D$ for each $i \in \{1, 2, \dots, m\}$ so that $\bar{Q}_t^0 \in \mathcal{X}_t^Q$ and $\bar{u}_t^{i,0} \in \mathcal{X}_t^u$ for all $t \in \mathcal{T}$. Under Assumptions 1(vi)–3, the ε -greedy sampling policy of (13), and the deterministic harmonic stepsizes given in (14), the sequences of iterates $\bar{u}^{i,n}$ for $i \in \{1, \dots, m\}$ and \bar{Q}^n generated by Algorithm 1 satisfy convergence rates of the form $\mathbf{E}[\|\bar{u}^{i,n} - u^{i,*}\|_2^2] \leq \mathcal{O}(1/n)$ for $i \in \{1, \dots, m\}$ and $\mathbf{E}[\|\bar{Q}^n - Q^*\|_2^2] \leq \mathcal{O}(1/n)$.

Sketch of Proof. The proof is also by backwards induction on t , where the induction hypothesis at each step is that \bar{Q}_{t+1}^n converges at a rate of $\mathcal{O}(1/n)$. Applying Lemma 2 along with some additional analysis, it is possible to show that $\bar{u}_t^{i,n}$ also converges at a rate of $\mathcal{O}(1/n)$. Lemma 3 then completes the proof by showing the desired rate for \bar{Q}_t^n . The details are given in Appendix A. \square

Figure 2. Sample paths of Dynamic-QBRM ADP.

Although both $\bar{u}^{i,n}$ and \bar{Q}^n converge at a rate of $\mathcal{O}(1/n)$, the slower sequence is \bar{Q}^n . To see why this step is slower, one can compare Lemmas 2 and 3. From Lemma 2, we see that the error of the quantity $\bar{u}_t^{i,n}$ depends only on the “last iteration error” (the error of $\bar{u}_t^{i,n-1}$) and the error of the next stage value function \bar{Q}_{t+1}^{n-1} . In contrast, Lemma 3 shows that the error of \bar{Q}_t^n also depends on the error of the current stage quantile $\bar{u}_t^{i,n-1}$, in addition to its own “last iteration error” and the error of the next stage value function. The interpretation of these bounds is that accuracy of $\bar{u}_t^{i,n-1}$ is needed before we obtain an accurate approximation of \bar{Q}_t^n , exactly as our intuition would suggest, given Figure 1. Indeed, this is observed in empirical experiments (see Section 7), motivating the second contribution of this paper, a procedure aimed toward speeding up the slow step of the ADP algorithm.

6. Efficient Sampling of the “Risky” Region

Arguably, the biggest practical issue with a Monte Carlo-based algorithm in the setting of risk-averse decision making is the question of *sampling*. To illustrate, suppose the one-step conditional risk measure is CVaR at a level of $\alpha = 0.99$. Because α is close to 1, the iterates of the Dynamic-QBRM ADP algorithm are volatile and exhibit poor empirical convergence rates, as shown in Figure 2. In this section, we discuss a method to address this issue.

6.1. Overview of the Main Idea

As we have mentioned, Dynamic-QBRM ADP can be applied in situations where the distribution of the information process $\{W_t\}$ is unknown, a common assumption for ADP algorithms. An example of when such a paradigm can be useful is when one has access to real data, but no stochastic model from which to simulate. Another example is a black-box simulator where the user has little to no control of its parameters. In these scenarios, a good remedy to any sampling issue is to implement an *adaptive stepsize rule*, similar to the likes of George and Powell [28], Schaul et al. [66], and Ryzhov et al. [65] (the third reference develops a stepsize rule specifically in the context of ADP), that can adjust based on previously observed data points. However, if the distribution of the stochastic process $\{W_t\}_{t=0}^T$ is known, then we can propose a new companion procedure to control the sampling process as our ADP algorithm progresses. The procedure takes advantage of the idea of *importance sampling* (see, e.g., Bucklew [13]) and is inspired by adaptive sampling approaches like the *cross-entropy method* of Rubinstein [58].


Recall from the results of the previous section that the convergence of \bar{Q}_t^n , the approximation of the value function, is expected to be slower than that of $\bar{u}_t^{i,n}$, the auxiliary variable. In this section, we propose a procedure called *risk-directed sampling* (RDS) to improve the sampling efficiency for the step of Dynamic-QBRM ADP where \bar{Q}_t^n is updated, i.e., Step 6 of Algorithm 1. The main idea, as illustrated in Figure 3, is to use the ADP approximations to drive the learning of the sampling distribution and simultaneously use the sampling distribution to generate the samples for updating the ADP approximations. As the figure shows, the hope is that our procedure updates the distribution in such a way that we dedicate samples to the regions of high risk (from where we may otherwise not see many samples).

6.2. Risk-Directed Sampling

Suppose the distribution of the exogenous information W_{t+1} has a density $p_t(w)$. Notice that by (6), we have for any (s, a) and t ,

$$Q_t^*(s, a) = \int H_t(u_t^{1,*}, u_t^{2,*}, \dots, u_t^{m,*}, Q_{t+1}^*(w))(s, a) p_t(w) dw. \quad (20)$$

Figure 3. RDS algorithm idea.



For convenience, we use the shorthand notation

$$H_t^*(w | s, a) = H_t(u_t^{1,*}, u_t^{2,*}, \dots, u_t^{m,*}, Q_{t+1}^*, w)(s, a),$$

to emphasize the variable of integration, w . From the principle of importance sampling (see, e.g., Bucklew [13]), it is known that to produce a low-variance estimate of $Q_t^*(s, a)$ using Monte Carlo sampling, one should sample from a distribution whose density is nearly proportional to the absolute value of the integrand of (20).

Our approach takes advantage of this proportionality property of the optimal density and directly constructs an approximation to the absolute value of the integrand of (20) by minimizing a *mean squared error*. The idea is to capture the “risky regions,” i.e., the parts of the outcome space where the integrand is large. When the parametric class of sampling distributions is chosen to be a mixture class (as we do), this introduces a simplification by allowing us to effectively remove the constraint $\|\theta\|_1 = 1$ from the optimization problem. In addition, we prove that our method converges without exact knowledge of the function $H_t^*(w | s, a)$. Instead, the algorithm works in conjunction with Dynamic-QBRM ADP by using approximations defined by

$$H_t^n(w | s, a) = H_t(\bar{u}_t^{1,n-1}, \bar{u}_t^{2,n-1}, \dots, \bar{u}_t^{m,n-1}, \bar{Q}_{t+1}^{n-1}, w)(s, a),$$

where $\bar{u}_t^{1,n-1}, \bar{u}_t^{2,n-1}, \dots, \bar{u}_t^{m,n-1}, \bar{Q}_{t+1}^{n-1}$ are outputs from Dynamic-QBRM ADP. Like the main ADP method, the procedure is fully adaptive and updates to the sampling distribution are made at every iteration. To our knowledge, an adaptive importance sampling approach for constructing risk-averse policies has not been considered in the literature. However, Kozmík and Morton [39] employs importance sampling from a different perspective: for the *evaluation* of risk-averse policies in stochastic programming.

Let $\{\phi_t^k\}_{k=1}^K$ be the set of densities for K prespecified basis distributions (we also refer to these as *basis functions*) for time t , from which we create a mixture distribution used for sampling. Note that it is often sufficient to specify one set of distributions for all t , as we do in Section 7. Our goal is to take the traditional regression approach and develop an approximation of the form $\sum_k \bar{\theta}_t^{k,n}(s, a) \phi_t^k(w) \approx |H_t^*(w | s, a)| p_t(w)$, where $\bar{\theta}_t^{k,n}(s, a)$ is the weight of the k -th basis density at iteration n and state (s, a) . For convenience, we also use the notation $\bar{\theta}_t^n(s, a)$ to be a vector in \mathbb{R}^K with components $\bar{\theta}_t^{k,n}(s, a)$ for $k = 1$ to K . Choose a (large) compact subset ${}^{\circlearrowleft}\mathcal{W} \subseteq {}^{\circlearrowright}\mathcal{W}$ and define a random variable W^u that is uniformly distributed over ${}^{\circlearrowleft}\mathcal{W}$; i.e., assume that its density is given by the function $p^u(w) = C_{{}^{\circlearrowleft}\mathcal{W}} \mathbf{1}_{\{w \in {}^{\circlearrowleft}\mathcal{W}\}}$ for some constant $C_{{}^{\circlearrowleft}\mathcal{W}}$. In what follows, we measure the mean squared error using this distribution, though of course, other reasonable choices may exist and the algorithm can be adapted accordingly. For a function $F: {}^{\circlearrowleft}\mathcal{W} \rightarrow \mathbb{R}$ and $\phi = (\phi^1, \phi^2, \dots, \phi^K)^\top$, define the projection operator

$$\Pi_\phi F = \arg \min_{\theta \geq 0} \mathbb{E}[\|\theta^\top \phi(W^u) - F(W^u)\|^2], \quad (21)$$

which maps F to a nonnegative coefficient vector θ corresponding to the best, i.e., minimum mean squared error, approximation under ϕ .

Proposition 2. *Under the condition that $\mathbb{E}[\phi(W^u)\phi(W^u)^\top]$ is positive definite, the optimization problem of (21) has a unique solution.*

Proof. If the positive definiteness condition holds, then the problem is a strictly convex quadratic program (see, e.g., Boyd and Vandenberghe [11]). \square

Exactly in the spirit of importance sampling, we define the optimal sampling density (within the mixture class) to be the probability density function corresponding to the best fit of the integrand using the given basis functions, i.e., $[\Pi_\phi[|H_t^*(\cdot | s, a)| p_t(\cdot)]]^\top \phi(w)$, which is computed by normalizing the weight vector so that it has unit norm. The sampling density we use at time t , iteration n , and state (s, a) and the corresponding likelihood ratio are given by

$$\bar{p}_t^{n-1}(w | s, a) \propto [\bar{\theta}_t^{n-1}(s, a)]^\top \phi(w) \quad \text{and} \quad L_{t+1}^n(s, a) = p_t(W_{t+1}^{q,n}) \bar{p}_t^{n-1}(W_{t+1}^{q,n} | s, a)^{-1},$$

where $L_{t+1}^n \in \mathbb{R}^d$ is defined for each t and n and $W_{t+1}^{q,n} \in \mathcal{W}$ is the random sample generated by from $\bar{p}_t^{n-1}(w | s, a)$. Assume that if $\bar{\theta}_t^{n-1}(s, a) = \mathbf{0}$ (the zero vector), then we set $\bar{p}_t^{n-1}(w | s, a) \propto \sum_k \phi^k(w)$. Finally, to update our coefficients $\bar{\theta}_t^n(s, a)$ from one iteration to the next, we need to introduce another \mathcal{G}_t^n -measurable stepsize sequence $\{\beta_t^n\}$ with $\beta_t^n \in \mathbb{R}^d$ for each t . The description of the new algorithm is given in Algorithm 2 below.

Besides the additional input parameters, there are three points of departure from Algorithm 1. In Step 3, the information process sample $W_{t+1}^{q,n}$ is drawn according to $\bar{p}_t^{n-1}(w | s, a)$ rather than the true distribution, $p_t(w)$. In Step 5, we include an additional factor corresponding to the likelihood ratio L_{t+1}^n in order to apply importance sampling. Step 7 is new and represents the updating step for the sampling coefficients $\bar{\theta}_t^n(s, a)$ using a stochastic approximation step. Before moving on to the convergence result, we need another stepsize assumption and an assumption on the basis functions. In addition, Lemma 5 verifies a technical condition that we need for the convergence result of Theorem 5.

Algorithm 2 (Dynamic-QBRM ADP with Risk-Directed Sampling)

Input: Initial estimates of the value function $\bar{Q}^0 \in \mathbb{R}^D$ and quantiles $\bar{u}^{i,0} \in \mathbb{R}^D$ for $i \in \{1, 2, \dots, m\}$.

Basis distributions ϕ^k and initial sampling coefficients $\bar{\theta}^0(s, a) \in \mathbb{R}^K$.

Stepsize rules γ_t^n , η_t^n , and β_t^n for all n, t .

Output: Approximations $\{\bar{Q}^n\}$ and $\{\bar{u}^{i,n}\}$.

```

1 Set  $\bar{Q}_T^n = 0$  for all  $n$ .
   for  $n = 1, 2, \dots$  do
2     Choose an initial state  $(S_0^n, a_0^n)$ .
     for  $t = 0, 1, \dots, T - 1$  do
3         Draw a sample of the information process  $W_{t+1}^{u,n}$  from the distribution of  $W_{t+1}$ .
         Draw an IS sample  $W_{t+1}^{q,n}$  so that  $W_{t+1}^{q,n}(s, a) \sim \bar{p}_t^{n-1}(w | s, a)$ .
4         Update auxiliary variable approximations for  $i = 1, \dots, m$ :
            $\bar{u}_t^{i,n} = \Pi_{\mathcal{X}_t^u} \{ \bar{u}_t^{i,n-1} - \text{diag}(\gamma_t^n) \psi_t^i(\bar{u}_t^{i,n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{u,n}) \}$ .
5         Compute an estimate of the future cost based on the current approximations:
            $\hat{q}_t^n = \text{diag}(L_{t+1}^n) [H_t(\bar{u}_t^{1,n-1}, \dots, \bar{u}_t^{m,n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{q,n})]$ .
6         Update approximation of value function:
            $\bar{Q}_t^n = \Pi_{\mathcal{Q}_t^q} \{ \bar{Q}_t^{n-1} - \text{diag}(\eta_t^n) (\bar{Q}_t^{n-1} - \hat{q}_t^n) \}$ .
7         Update the sampling coefficients for each state. Let  $w = W_{t+1}^{q,n}(s, a)$  and
            $\bar{\theta}_t^n(s, a) = [\bar{\theta}_t^{n-1}(s, a) - \beta_t^n(s, a) [(\bar{\theta}_t^{n-1}(s, a))^\top \phi(w) - |H_t^n(w | s, a)| p_t(w)] \phi(w) p^u(w) \bar{p}_t^{n-1}(w | s, a)^{-1}]^+$ ,
           where  $[\cdot]^+$  is taken componentwise.
8         Choose next state  $(S_{t+1}^n, a_{t+1}^n)$ .
     end
   end.
```

Assumption 4. For all $(s, a) \in \mathcal{U}$ and $t \in \mathcal{T}$, suppose β_t^n is \mathcal{G}_t^n -measurable and

- (i) $\beta_t^n(s, a) = \bar{\beta}_t^{n-1} \mathbf{1}_{\{(s, a) = (S_t^n, a_t^n)\}}$, for some $\bar{\beta}_t^{n-1} \in \mathbb{R}$ that is \mathcal{G}_t^{n-1} -measurable.
- (ii) $\sum_{n=1}^\infty \beta_t^n(s, a) = \infty$, $\sum_{n=1}^\infty \beta_t^n(s, a)^2 < \infty$ a.s.
- (iii) $\mathbf{E}[(\bar{\beta}_t^{n-1})^2] \leq \mathcal{O}(n^{-1-\epsilon})$ for some $\epsilon > 0$.

Assumption 5. With regard to the basis distributions ϕ^k , $k \in \{1, 2, \dots, K\}$, the following hold:

- (i) $\mathbf{E}[\phi(W^u)\phi(W^u)^\top]$ is a positive definite matrix (so that we may apply Proposition 2),
- (ii) Let $\hat{q}_t^n \in \mathbb{R}^d$ be defined as in Step 5 of Algorithm 2. There exists a constant C_ϕ such that the following holds for any t and (s, a) : given $\theta \geq 0$ and $\|\theta\|_1 = 1$, the tails of the distribution $\sum_k \theta^k \phi^k$ are “heavy” enough to guarantee that if $W_{t+1}^{q,n} \sim \sum_k \theta^k \phi^k$, then $\mathbf{E}[\hat{q}_t^n(s, a)^2] \leq C_\phi$.

We remark that due to the compactness of the sets \mathcal{X}_t^u and \mathcal{X}_t^q , many of the terms in the definition of \hat{q}_t^n are bounded; therefore, the crucial terms that affect Assumption 5(ii) are L_{t+1}^n and the cost function c_t . The condition that $\mathbf{E}[\beta_t^n(s, a)^2] \leq \mathcal{O}(n^{-1-\epsilon})$ is not particularly difficult to satisfy; for example, our deterministic harmonic stepsizes satisfy the condition with $\epsilon = 1$. In addition, polynomial rules of the form $\beta_t^n(s, a) = n^{-1/2-\epsilon/2} \mathbf{1}_{\{(s, a) = (S_t^n, a_t^n)\}}$ work as well. We state the convergence result for Algorithm 2 in Theorem 5. Since Step 7 of Algorithm 2 does not project to a compact set, we cannot make use of Kushner and Yin [40, Theorem 2.4] as we did before. Instead, our convergence result is derived from a theorem of Pflug [45] for stochastic approximation, which requires the result of Lemma 5. The rest of the proof is mostly standard and is thus deferred to Appendix A.

Because the importance sampled stochastic processes are corrected for in expectation using the likelihood ratio L_{t+1}^n , we should not expect the rate of convergence to change. Indeed, we have the following theorem.

Theorem 4 (Convergence Rate). Choose initial approximations $\bar{Q}^0 \in \mathbb{R}^D$ and $\bar{u}^{i,0} \in \mathbb{R}^D$ for each $i \in \{1, 2, \dots, m\}$ so that $\bar{Q}_t^0 \in \mathcal{X}_t^q$ and $\bar{u}_t^{i,0} \in \mathcal{X}_t^u$ for all $t \in \mathcal{T}$. Under Assumptions 1(vi)–5, the ε -greedy sampling policy of (13), and the deterministic harmonic stepsizes given in (14), the sequences of iterates $\bar{u}^{i,n}$ for $i \in \{1, \dots, m\}$ and \bar{Q}^n generated by Algorithm 1 satisfy convergence rates of the form $\mathbf{E}[\|\bar{u}^{i,n} - u^{i,*}\|_2^2] \leq \mathcal{O}(1/n)$ for $i \in \{1, \dots, m\}$ and $\mathbf{E}[\|\bar{Q}^n - Q^*\|_2^2] \leq \mathcal{O}(1/n)$.

Proof. The proof is analogous to that of Theorem 3. We need versions of Lemmas 2 and 3 for Algorithm 2. The difference is that we need to deal with the term $L_{t+1}^n(s, a)$, which has expectation equal to 1. \square

Theorem 5 (Convergence of Risk-Directed Sampling Procedure). Choose initial approximations $\bar{Q}^0 \in \mathbb{R}^D$ and $\bar{u}^{i,0} \in \mathbb{R}^D$ for each $i \in \{1, 2, \dots, m\}$ so that $\bar{Q}_t^0 \in \mathcal{X}_t^q$ and $\bar{u}_t^{i,0} \in \mathcal{X}_t^u$ for all $t \in \mathcal{T}$. Under Assumptions 1(vi)–5, the ε -greedy sampling policy of (13), and the deterministic harmonic stepsizes given in (14), Algorithm 2 generates a sequence of iterates \bar{Q}^n that converge almost surely to the optimal value function Q^* . Moreover, the sampling coefficients $\bar{\theta}_t^n(s, a)$ converge to the optimal sampling coefficients under ϕ :

$$\bar{\theta}_t^n(s, a) \longrightarrow \Pi_\phi[H_t^*(\cdot | s, a) | p_t(\cdot)] \quad a.s.$$

for each t and $(s, a) \in \mathcal{U}$.

Proof. See Appendix A. \square

Not surprisingly, the new algorithm retains the same theoretical properties of the standard Dynamic-QBRM ADP without RDS. In addition, we now have a sampling density that converges to the optimal sampling density as the algorithm progresses (Theorem 5). What remains for us to explore are the empirical convergence rates of the two approaches.

7. Numerical Results

The recent surge of interest in energy and sustainability has shown that when pertaining to the question of risk, the problem of optimal control of energy storage assets is an especially rich application domain. In this section, we illustrate our proposed ADP algorithm by way of a stylized energy trading and bidding problem in which both heavy tails and extreme events play a prominent role.

7.1. Model

We consider the problem of using energy storage to trade in the electricity market, i.e., *energy arbitrage*, with the caveat that there is the possibility of some financial penalty when the amount of stored energy is low. For example, since storage for backup purposes is rarely in use, one might consider using it to generate a stream of revenue by interacting with the market. However, this immediately introduces a source of risk: in the rare event when backup is needed, it is crucial that there is enough energy to cover demand. Failure to do can cause complications, so it is useful to consider a risk-averse policy to this problem. See Xi et al. [75] for a detailed model of the shared storage situation, but solved with a risk-neutral objective. When no penalty is assessed, we assume that there is a modest reward. Finally, we introduce a bidding aspect to the problem where one must place bids prior to the desired transaction time (see Jiang and Powell [33] for a detailed model of hour-ahead bidding).

For $t = 0$ to $t = T$, let $S_t \in \mathcal{S} = \{0, 1, \dots, S_{\max}\}$ be the amount of energy in storage and let $P_t \sim \log \mathcal{N}(\mu_P(t), \sigma_P(t)^2)$ be the (heavy-tailed) spot price of electricity. Also, let $U_t \sim \mathcal{N}(0, \sigma_U^2)$ be independent of P_t , and suppose our simple model of storage-based penalties is as follows. Given two constants $0 < a < b$, where a represents the “rate of reward” and b represents the “rate of penalty,” assume that the reward/penalty assessed at time $t + 1$ is given by

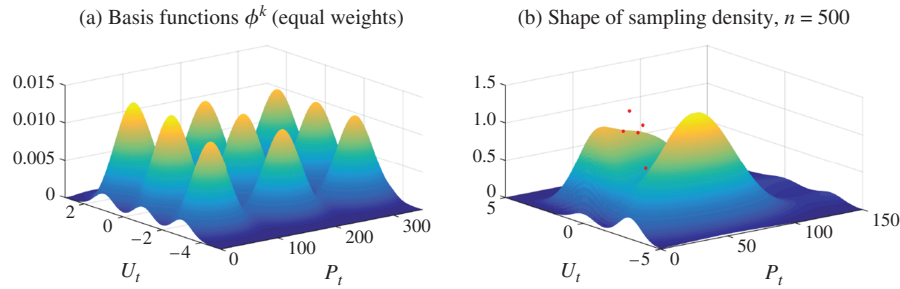
$$F_{t+1} = |\mu_S(S_t) + U_{t+1}| [b \cdot \mathbf{1}_{\{\mu_S(S_t) + U_{t+1} < 0\}} - a \cdot \mathbf{1}_{\{\mu_S(S_t) + U_{t+1} \geq 0\}}],$$

where $\mu_S: \mathcal{S} \rightarrow \mathbb{R}$ is a nondecreasing function, signifying that as S_t increases, the possibility of penalty decreases (rare event). Our two-dimensional action at each time t is given by

$$a_t = (b_t^-, b_t^+) \in \mathcal{A} \subseteq \{(b^-, b^+): 0 \leq b^- \leq b^+ \leq b_{\max}\},$$

where $|\mathcal{A}| < \infty$. There are no constraints, so $\mathcal{A}_s = \mathcal{A}$ for all $s \in \mathcal{S}$. We call b_t^- the *buy bid* and b_t^+ the *sell bid*: if P_t fluctuates below the buy bid, we are obligated to buy from the market, and if P_t rises above the sell bid, we are obligated to sell to the market. In addition, we are penalized the amount of the spot price if we are to sell but the storage device is empty, i.e., $S_t = 0$. In this problem, our information process W_t is given by the pair (P_t, U_t)

Figure 4. (Color online) Example illustration of risk-directed sampling ($\lambda = 0.5$).



and is independent of the past. We find it most natural to model this problem in the sense of maximizing revenues (or “contributions”) rather than minimizing costs; hence, the *contribution function* is

$$c_t(S_t, a_t, W_{t+1}) = -F_{t+1} + P_{t+1} [\mathbf{1}_{\{b_t^+ < P_{t+1}\}} - \mathbf{1}_{\{b_t^- > P_{t+1}\}} - \mathbf{1}_{\{S_t=0\}} \mathbf{1}_{\{b_t^+ < P_{t+1}\}}],$$

and the transition function is given by $S_{t+1} = [\min\{S_t + \mathbf{1}_{\{b_t^- > P_{t+1}\}} - \mathbf{1}_{\{b_t^+ < P_{t+1}\}}, S_{\max}\}]^+$. The objective is to optimize the risk-averse model given in (2) with the “min” operator replaced by a “max” and taking the QBRM to be mean-CVaR at $\alpha = 0.99$ (this refers to the lower 0.01 tail of the distribution since we are now in the setting of rewards, not costs). Various values of λ are considered.

7.2. Parameter Choices

We let $T = 12$ and the size of the storage device be $S_{\max} = 6$. Next, suppose the values of μ_S are chosen so that $\mathbf{P}(\mu_S(S_t) + U_{t+1} < 0)$ are 0.1, 0.05, 0.02, 0.01, 0.01, 0.001, and 0.001 for $S_t = 0, 1, \dots, S_{\max}$ (representing increasingly rare penalties) with $\sigma_U^2 = 1$. The penalty parameters are set to be $a = 5$ and $b = 500$, making the penalty events relatively severe. We denote the (seasonal) mean of the spot price process by $m(t) = \mathbf{E}(P_t) = 50 \sin(4\pi t/T) + 100$, with a constant variance $v = \text{Var}(P_t) = 3,000$. By the properties of the lognormal distribution, this leads to the parameters $\mu_p(t) = \log(m(t)(1 + v/m(t)^2)^{-1/2})$ and $\sigma_p(t) = \sqrt{\log(1 + v/m(t)^2)}$. The maximum bid is chosen to be $b_{\max} = 500$ and each dimension (the range from 0 to 500) is discretized into increments of 50, resulting in an MDP with approximately 7,000 states.

Finally, we discuss our choice of parameters for the risk-directed sampling (RDS) procedure. For ϕ^1 , we simply take the true distribution of W_{t+1} and for the remaining ϕ^k , we select a set of bivariate normal distributions placed in a grid, with mean of the first component taking values in $\{50, 175, 300\}$ and the mean of the second component taking values in $\{-3, -1, 1\}$. The standard deviations are chosen to be 750 and 0.25 for the respective components. This is a fairly general choice of basis functions that uses very little problem specific information and/or tuning. Figure 4(b) shows an example of the shape of the sampling density after 500 iterations of RDS for a fixed state of $(S_t, b_t^-, b_t^+) = (0, 150, 300)$ at $t = T - 1$ with $\lambda = 0.5$ (note: the z-axis is not normalized here and the range of the P_t axis has been decreased to focus on the nonzero areas). As we would expect, the algorithm has chosen to allocate a relatively large sampling effort toward small values of U_t , presumably due to the fact that these events are very costly with $b = 500$ and thus critical to the estimation of CVaR. The last few observations are shown as red points.

7.3. Evaluating Policies

Here, we describe the procedure for evaluating policies under the objective (2). Let $\{A_0^\pi, A_1^\pi, \dots, A_{T-1}^\pi\}$ be a policy; to compute (2), let $S_{t+1}^\pi = S^M(s, A_t^\pi(s), W_{t+1})$ and we solve the recursive equations given by

$$V_t^\pi(s) = \rho_t^\alpha (c_t(s, A_t^\pi(s), W_{t+1}) + V_{t+1}^\pi(S_{t+1}^\pi)), \quad \text{for all } s \in \mathcal{S}, t \in \mathcal{T}, \quad V_T^\pi(s) = 0, \quad \text{for all } s \in \mathcal{S}. \quad (22)$$

The value of the policy is then given by $V_0^\pi(S_0)$, but since W_t is continuously distributed, we cannot, in general, solve these equations exactly. For this reason, we select a large, finite sample $\hat{\Omega} \subseteq \Omega$, and apply the standard *sample average approximation* (SAA) technique of Kleywegt et al. [38] to the Bellman recursion (22) along with the standard linear programming method suggested in Rockafellar and Uryasev [56] for computing conditional value at risk. Next, in order to have an optimality benchmark against which ADP policies can be compared, we also apply the SAA technique to the Bellman recursion for the *optimal* policy specified in Theorem 1 (with $\rho_t = \rho_t^\alpha$ and the “min” replaced by “max”).

We use 50,000 realizations of $W_t = (P_t, U_t)$ for each $t = 1, 2, \dots, T$ in our simulations below. For the purposes of our numerical work and Figure 7, we refer to the SAA optimal policy (which has value function $V_{0,\beta}^*(S_0)$)

as our benchmark for “100% optimal” and a corresponding SAA *myopic policy*, denoted by π_m , (i.e., the policy which takes the action to maximize $\rho_t^\alpha(c_t(s, a, W_{t+1}))$ with zero continuation value to be “0% optimal.” More precisely, we have

$$\% \text{ optimality of policy } \pi \text{ (with respect to myopic policy } \pi_m) = \frac{V_0^\pi(S_0) - V_0^{\pi_m}(S_0)}{V_0^*(S_0) - V_0^{\pi_m}(S_0)}.$$

Recall that larger $V_0^\pi(S_0)$ is better, as we have switched to the maximization setting for the numerical application.

7.4. Results

Let us first show a few plots that illustrate the effectiveness of the RDS procedure. Once again, we fix the state to be $(S_t, b_t^-, b_t^+) = (0, 150, 300)$ at $t = T - 1$ and set $\lambda = 0.5$, but now we focus on the evolution of the approximations \bar{Q}_t^n and \bar{u}_t^n ; representative sample paths are shown in Figure 5 (the first plot is from Figure 2, repeated for comparison purposes).

To make the comparison as fair as possible, the random number generator seed is set so that the sequence of observations used to update \bar{u}_t^n is the same in both cases (i.e., the gray lines are identical). The true limit points are approximately $u_t^*(S_t, a_t) \approx -555$ and $Q_t^*(S_t, a_t) \approx -387$, and we notice that in the case of Figure 5(b), a decent approximation of Q_t^* is obtained around visit 2,500, while the approximation in Figure 5(a) does not settle until the end. This drastic difference in empirical convergence rate can have a significant impact on the performance of Dynamic-QBRM ADP as the large errors shown in Figure 5(a) are propagated backwards in time.

In fact, we can illustrate this impact by examining the value function approximations back at time $t = 0$, with and without RDS, compared to the SAA optimal value function. Figure 6 compares the two methods after $N = 5,000,000$ iterations by varying the bid dimensions of a fixed state $S_0 = 0$. Notice that in Figure 6(b), the approximation from the RDS procedure very closely resembles the optimal value function of Figure 6(c) even after errors are propagated through $T = 12$ steps. The same cannot be said of the approximation in Figure 6(a).

Lastly, we illustrate that in addition to the improved empirical convergence rates observed in Figures 5 and 6, running Dynamic-QBRM ADP with RDS has a noticeable effect on the resulting policies as well. Using the backward recursive SAA evaluation procedure described above, we plot optimality percentages of risk-averse policies produced by the two variants of the algorithm for $\lambda \in \{0.6, 0.55, 0.5, 0.45, 0.4\}$ in Figure 7. We see that RDS has the advantage in all cases, especially during the early iterations. Figure 7 uses a log-scale to display the early progress of the algorithms every 50,000 iterations from $N = 0$ to $N = 1,000,000$ and also the asymptotic progress every 1,000,000 iterations from $N = 1,000,000$ to $N = 5,000,000$. Our simulations also suggest that as λ decreases, the advantage of using RDS diminishes; this is not surprising because decreasing λ corresponds to deemphasizing the “risky” events that RDS seeks to sample. In fact, we find that in this particular example

Figure 5. Sample paths of approximations generated by Dynamic-QBRM ADP ($\lambda = 0.5$).




Figure 6. (Color online) Surface plots of value function approximations at $t = 0$ ($\lambda = 0.5$).


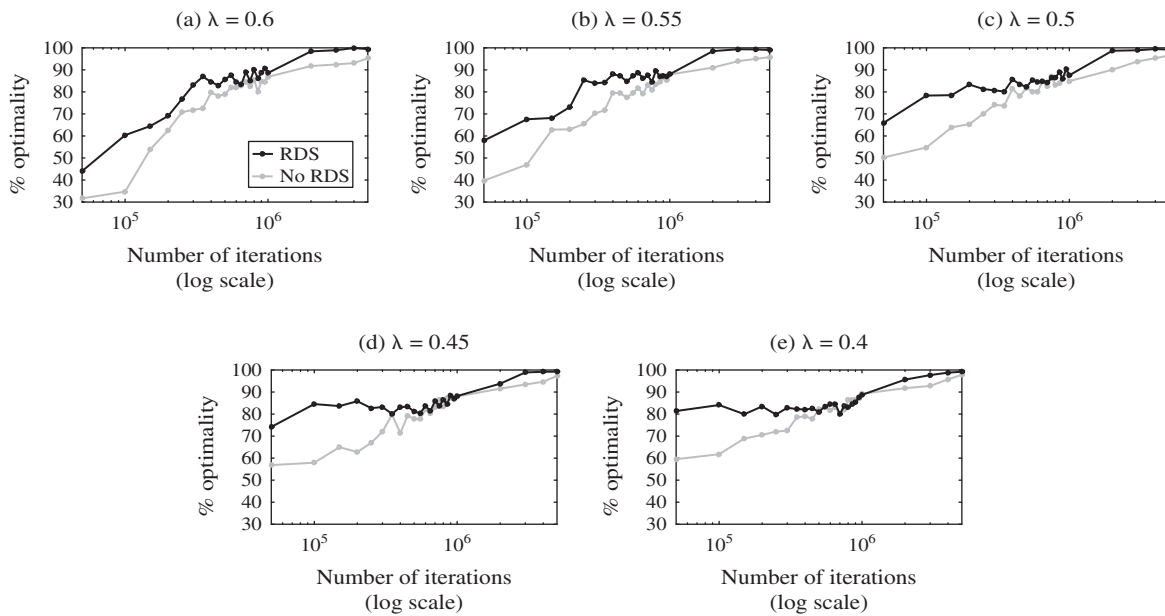


Figure 7. Comparison of Dynamic-QBRM ADP with and without RDS.**Table 1.** CPU times for $N = 5,000,000$ iterations of ADP and SAA optimal solution (minutes).

Method	$\lambda = 0.6$	$\lambda = 0.55$	$\lambda = 0.5$	$\lambda = 0.45$	$\lambda = 0.4$
Dynamic-QBRM ADP	810.24	819.40	816.40	830.69	811.45
Dynamic-QBRM ADP with RDS	926.25	983.33	989.87	926.85	973.50
SAA Backward Recursion	2,119.74	2,058.36	2,163.91	2,118.84	2,072.79

problem, the algorithms are practically indistinguishable for $\lambda \leq 0.3$. Computation times for the three methods are given in Table 1; notice that we are able to run the ADP methods for five million iterations in roughly half of the time it takes to compute the SAA optimal solution (the large number of scenarios presents a challenge for the optimization routine implemented in CPLEX). The online companion contains additional numerical results of evaluating our policies on more “practical” metrics of risk and reward, rather than the dynamic risk measure objective function. We observe that the policies generated by our algorithms behave in an intuitively appealing way.

8. Conclusion

In this paper, we propose a new ADP algorithm for solving risk-averse MDPs (Ruszczynski [60]) under a class of quantile-based risk measures. The algorithm uses a two-step updating procedure to approximate the optimal value function and we prove that it converges almost surely with a rate of $\mathcal{O}(1/n)$. We also show a companion sampling procedure to more efficiently sample the “risky” region of the outcome space and then prove that the sampling distribution converges to one that is, in a sense, the best within a parametric class. Next, we describe an energy storage and bidding application, on which we demonstrate that the RDS sampling approach provides significant benefits in terms of the empirical convergence rate. Moreover, we illustrate that the approximate policies produced by Dynamic-QBRM ADP behave in an intuitive and appealing way in terms of the tradeoff between risk and reward (see online companion), suggesting that it can be readily applied to a variety of problem settings.

Acknowledgments

The authors are grateful to the Adam Schwartz (the area editor), the associate editor, and the reviewers, whose detailed comments have helped to significantly improve earlier versions of this paper. The authors would also like to thank Andrzej Ruszczyński for valuable conversations and comments.

Appendix A. Proofs

Theorem 2 (Almost Sure Convergence). *Choose initial approximations $\bar{Q}^0 \in \mathbb{R}^D$ and $\bar{u}^{i,0} \in \mathbb{R}^D$ for each $i \in \{1, 2, \dots, m\}$ so that $\bar{Q}_t^0 \in \mathcal{X}_t^q$ and $\bar{u}_t^{i,0} \in \mathcal{X}_t^u$ for all $t \in \mathcal{T}$. Under Assumptions 1–2, Algorithm 1 generates a sequence of iterates \bar{Q}^n that converge almost surely to the optimal value function Q^* .*

Proof. We induct backwards on t , starting with the base case $t = T$. Since $Q_T^n(s, a) = Q_T^*(s, a) = 0$ for all n and all (s, a) , convergence trivially holds. The induction hypothesis is that $\bar{Q}_{t+1}^n(s, a) \rightarrow Q_{t+1}^*(s, a)$ almost surely, and we aim to show the same statement with t replacing $t + 1$. As we did before, we analyze the statistical error and approximation error, $\epsilon_{t+1}^{q,n}, \xi_{t+1}^{q,n} \in \mathbb{R}^d$, except they are now with respect to the stochastic Bellman operator H_t :

$$\begin{aligned} \epsilon_{t+1}^{q,n} &= Q_t^* - H_t(u_t^{1,*}, \dots, u_t^{m,*}, Q_{t+1}^*, W_{t+1}^{q,n}), \\ \xi_{t+1}^{q,n} &= H_t(u_t^{1,*}, \dots, u_t^{m,*}, Q_{t+1}^*, W_{t+1}^{q,n}) - H_t(\bar{u}_t^{1,n-1}, \dots, \bar{u}_t^{m,n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{q,n}). \end{aligned}$$

Therefore, the update of Step 6 of Algorithm 1 is equivalent to

$$\bar{Q}_t^n = \Pi_{\mathcal{X}_t^q} \left\{ \bar{Q}_t^{n-1} - \text{diag}(\eta_t^n) [\bar{Q}_t^{n-1} - Q_t^* + \epsilon_{t+1}^{q,n} + \xi_{t+1}^{q,n}] \right\}. \quad (\text{A.1})$$

It is once again clear that $\mathbb{E}[\epsilon_{t+1}^{q,n}(s, a) \mid \mathcal{G}_{t+1}^{n-1}] = 0$ almost surely. Furthermore, we can easily argue that for some positive constant C_H ,

$$\mathbb{E} \left[\left[\bar{Q}_t^{n-1}(s, a) - H_t(\bar{u}_t^{1,n-1}, \dots, \bar{u}_t^{m,n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1}^{q,n})(s, a) \right]^2 \right] \leq C_H, \quad (\text{A.2})$$

using Assumption 2(i), Assumption 2(ii), the fact that $\rho_t^a(0) = \Phi(0, \dots, 0) = 0$, and the boundedness of the iterates \bar{Q}_{t+1}^{n-1} and \bar{u}_t^n .

Now, fix an (s, a) and let $S_{t+1}^n = S^M(s, a, W_{t+1}^{q,n})$. Expanding and using Assumption 2(i) and the property (10) we see that

$$\begin{aligned} |\xi_{t+1}^{q,n}(s, a)| &\leq L_\Phi \left| \min_{a' \in \mathcal{A}_{S_{t+1}^n}} Q_{t+1}^*(S_{t+1}^n, a') - \min_{a' \in \mathcal{A}_{S_{t+1}^n}} \bar{Q}_{t+1}^{n-1}(S_{t+1}^n, a') \right| + L_\Phi \sum_i |\bar{u}_t^{i,n-1}(s, a) - u_t^{i,*}(s, a)| \\ &\leq L_\Phi \left[\|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_\infty + \sum_i \|\bar{u}_t^{i,n-1} - u_t^{i,*}\|_\infty \right]. \end{aligned} \quad (\text{A.3})$$

Taking conditional expectation, applying Lemma 1, and using the induction hypothesis (which tells us that $\bar{u}_t^{i,n} \rightarrow u_t^{i,*}$ and $\bar{Q}_{t+1}^n \rightarrow Q_{t+1}^*$ almost surely), we conclude

$$\mathbb{E}[\xi_{t+1}^{q,n}(s, a) \mid \mathcal{G}_{t+1}^{n-1}] \rightarrow 0 \quad \text{a.s.},$$

satisfying one of the conditions of Kushner and Yin [40, Theorem 2.4]. In addition, note the bounded second moment condition of (A.2), the conditional unbiasedness of $\epsilon_{t+1}^{q,n}(s, a)$, the stepsize, sampling, and truncation properties of Assumption 1. We once again have the ingredients to apply the stochastic approximation convergence theorem Kushner and Yin [40, Theorem 2.4] (the objective function for applying the theorem is $q \mapsto \|q - Q_t^*\|_2^2$) to the update equation of (A.1) in order to conclude

$$\bar{Q}_t^n(s, a) \rightarrow Q_t^*(s, a), \quad \text{a.s.},$$

for every $(s, a) \in \mathcal{U}$, completing both the inductive step and the proof. \square

Proposition 1. *Under the ε -greedy sampling policy given in (13) and the deterministic harmonic stepsizes given in (14), Assumption 1(iii)–1(iv) is satisfied.*

Proof. Let us consider, for a fixed t and (s, a) , the stepsize sequence

$$\gamma_t^n(s, a) = \frac{\gamma_t}{n} \mathbf{1}_{\{(s, a) = (S_t^n, a_t^n)\}}.$$

The argument for the case of $\eta_t^n(s, a)$ is, of course, exactly the same and is omitted. First, notice that the second part, i.e., $\sum_{n=1}^\infty \gamma_t^n(s, a)^2 < \infty$ almost surely, is trivial. We thus focus on proving that the second part of the assumption, that $\sum_{n=1}^\infty \gamma_t^n(s, a) = \infty$ almost surely, holds in this case. An iteration n is called a *visit* to state (s, a) if $(s, a) = (S_t^n, a_t^n)$ (so the stepsize is nonzero). Let $\{\Delta_t^k(s, a)\}_{k \geq 1}$ be a process that describes the *interarrival times* for visits to the state (s, a) ; in other words, $\Delta_t^k(s, a)$ is the number of iterations that pass between the $(k-1)$ -st and k -th visits to the state (s, a) . The *arrival time* sequence (the iterations for which (s, a) is visited) is thus given by the sum

$$N_t^k(s, a) = \sum_{k'=1}^k \Delta_t^{k'}(s, a), \quad \text{for } k > 0.$$

To simplify notation, we henceforth drop the dependence of these processes on (s, a) and use $\Delta_t^k = \Delta_t^k(s, a)$ and $N_t^k = N_t^k(s, a)$. By selecting only the iterations for which the stepsizes are nonzero, we see that

$$\sum_{k=1}^\infty \gamma_t^{N_t^k} = \sum_{n=1}^\infty \gamma_t^n(s, a).$$

Analogous to N_t^k process, define the deterministic sequence

$$n_t^k = \left\lceil \sum_{k'=1}^k \delta_t^{k'} \right\rceil \quad \text{and} \quad \delta_t^k = \frac{-2}{\log(1-\varepsilon)} \log k.$$

Observe that under the ε -greedy sampling policy, the event $\{(s, a) \neq (S_t^n, a_t^n)\}$ occurs with probability at most $1 - \varepsilon$. Hence, by independence, we can show that

$$\mathbf{P}(\Delta_t^k \geq \delta_t^k) \leq (1 - \varepsilon)^{\lceil \delta_t^k \rceil} \leq \frac{1}{k^2}.$$

By the Borel-Cantelli Lemma, with probability 1, the events $\{\Delta_t^k \geq \delta_t^k\}$ occur finitely often. Hence, there exists an almost surely finite K_t such that for all $k \geq K_t$, it is true that $\Delta_t^k \leq \delta_t^k$. Clearly, $n_t^k \leq \mathcal{O}(k \log k)$ and

$$\sum_{k=1}^{\infty} \gamma_t^{n_t^k} = \sum_{k=1}^{\infty} \frac{\gamma_t}{n_t^k} = \infty \quad (\text{A.4})$$

holds. Let $C_1 = \sum_{i=1}^{K_t-1} (\gamma_t / N_t^i)$ and $C_2 = \sum_{i=1}^{K_t-1} \Delta_t^i$ be almost surely finite random variables. Some simple manipulations yield (the following chain of inequalities may be analyzed ω -wise to obtain the *a.s.* qualification)

$$\sum_{k=1}^{\infty} \gamma_t^{N_t^k} = \sum_{k=1}^{\infty} \frac{\gamma_t}{N_t^k} = C_1 + \sum_{k=K_t}^{\infty} \frac{\gamma_t}{C_2 + \sum_{i=K_t}^k \Delta_t^i} \geq C_1 + \sum_{k=K_t}^{\infty} \frac{\gamma_t}{C_2 + \sum_{i=K_t}^k \delta_t^i} \geq C_1 + \sum_{k=K_t}^{\infty} \frac{\gamma_t}{C_2 - \sum_{i=0}^{K_t-1} \delta_t^i + n_t^k} = \infty \quad \text{a.s.},$$

where the final equality follows from (A.4). \square

Lemma 3. Under the same conditions as Lemma 2, the sequence of approximations \bar{Q}_t^n generated by Algorithm 1 satisfies, for any $\kappa_0, \kappa_1, \dots, \kappa_m > 0$,

$$\begin{aligned} \mathbf{E}[\|\bar{Q}_t^n - Q_t^*\|_2^2] &\leq \left[1 - \frac{\eta_t}{n} \left(2\varepsilon - L_\Phi \kappa_0 - L_\Phi \sum_{i=1}^m \kappa_i \right) \right] \mathbf{E}[\|\bar{Q}_t^{n-1} - Q_t^*\|_2^2] \\ &\quad + \frac{\eta_t L_\Phi}{n \kappa_0} \mathbf{E}[\|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_2^2] + \frac{\eta_t}{n} \sum_{i=1}^m \frac{L_\Phi}{\kappa_i} \mathbf{E}[\|\bar{u}_t^{i, n-1} - u_t^{i,*}\|_2^2] + C_H \frac{\eta_t^2}{n^2}, \end{aligned}$$

where C_H bounds the term

$$\mathbf{E}[\bar{Q}_t^{n-1}(s, a) - H_t(\bar{u}_t^{1, n-1}, \dots, \bar{u}_t^{m, n-1}, \bar{Q}_{t+1}^{n-1}, W_{t+1})(s, a)]^2 \leq C_H,$$

for all $(s, a) \in \mathcal{U}$.

Proof. Note that although we omit its algebraic form, the existence of C_g is guaranteed by Assumption 2(ii) and the boundedness of $\bar{Q}_t^{n-1}(s, a)$. Let us recall the update equation given in Step 7 of Algorithm 1 can be rewritten as

$$\bar{Q}_t^n = \Pi_{\mathcal{X}_t^n} \left\{ \bar{Q}_t^{n-1} - \text{diag}(\eta_t^n) [\bar{Q}_t^{n-1} - Q_t^* + \epsilon_{t+1}^{q,n} + \xi_{t+1}^{q,n}] \right\}.$$

Expanding, we have

$$\|\bar{Q}_t^n - Q_t^*\|_2^2 \leq \|\bar{Q}_t^{n-1} - Q_t^*\|_2^2 + \|\text{diag}(\eta_t^n) [\bar{Q}_t^{n-1} - Q_t^* + \epsilon_{t+1}^{q,n} + \xi_{t+1}^{q,n}]\|_2^2 - 2(\bar{Q}_t^{n-1} - Q_t^*)^\top \text{diag}(\eta_t^n) [\bar{Q}_t^{n-1} - Q_t^* + \epsilon_{t+1}^{q,n} + \xi_{t+1}^{q,n}]. \quad (\text{A.5})$$

We focus on the cross term. First, by the ε -greedy sampling policy, notice that

$$\mathbf{E}[(\bar{Q}_t^{n-1} - Q_t^*)^\top \text{diag}(\eta_t^n) (\bar{Q}_t^{n-1} - Q_t^*) \mid \mathcal{G}_{t+1}^{n-1}] \geq \frac{\varepsilon \eta_t}{n} \|\bar{Q}_t^{n-1} - Q_t^*\|_2^2, \quad (\text{A.6})$$

and by the definition of $\epsilon_{t+1}^{q,n}$,

$$\mathbf{E}[(\bar{Q}_t^{n-1} - Q_t^*)^\top \text{diag}(\eta_t^n) \epsilon_{t+1}^{q,n} \mid \mathcal{G}_{t+1}^{n-1}] = 0. \quad (\text{A.7})$$

Using the bound (A.3), the fact that η_t^n contains exactly one nonzero component, and the monotonicity of the l_p norms, we can see that

$$\begin{aligned} \mathbf{E}[-(\bar{Q}_t^{n-1} - Q_t^*)^\top \text{diag}(\eta_t^n) \xi_{t+1}^{q,n} \mid \mathcal{G}_{t+1}^{n-1}] &\leq \frac{\eta_t}{n} L_\Phi \left[\|\bar{Q}_t^{n-1} - Q_t^*\|_\infty \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_\infty + \|\bar{Q}_t^{n-1} - Q_t^*\|_\infty \sum_i \|\bar{u}_t^{i, n-1} - u_t^{i,*}\|_\infty \right] \\ &\leq \frac{\eta_t}{n} L_\Phi \left[\|\bar{Q}_t^{n-1} - Q_t^*\|_2 \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_2 + \|\bar{Q}_t^{n-1} - Q_t^*\|_2 \sum_i \|\bar{u}_t^{i, n-1} - u_t^{i,*}\|_2 \right]. \end{aligned}$$

Again applying $2ab \leq a^2 \kappa + b^2 / \kappa$, we see that for any constants $\kappa_k > 0$, $k = 0, 1, \dots, m$,

$$\mathbf{E}[-2(\bar{Q}_t^{n-1} - Q_t^*)^\top \text{diag}(\eta_t^n) \xi_{t+1}^{q,n} \mid \mathcal{G}_{t+1}^{n-1}] \leq \frac{\eta_t}{n} L_\Phi \left[\left(\kappa_0 + \sum_{i=1}^m \kappa_i \right) \|\bar{Q}_t^{n-1} - Q_t^*\|_2^2 + \frac{1}{\kappa_0} \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_2^2 + \sum_{i=1}^m \frac{1}{\kappa_i} \|\bar{u}_t^{i, n-1} - u_t^{i,*}\|_2^2 \right]. \quad (\text{A.8})$$

The statement of the lemma follows by taking expectations of (A.5)–(A.8), using the C_H bound, and combining. \square

Theorem 3 (Convergence Rate). Choose initial approximations $\bar{Q}^0 \in \mathbb{R}^D$ and $\bar{u}^{i,0} \in \mathbb{R}^D$ for each $i \in \{1, 2, \dots, m\}$ so that $\bar{Q}_t^0 \in \mathcal{X}_t^q$ and $\bar{u}_t^{i,0} \in \mathcal{X}_t^u$ for all $t \in \mathcal{T}$. Under Assumptions 1(vi)–3, the ε -greedy sampling policy of (13), and the deterministic harmonic stepsizes given in (14), the sequences of iterates $\bar{u}^{i,n}$ for $i \in \{1, \dots, m\}$ and \bar{Q}^n generated by Algorithm 1 satisfy convergence rates of the form $\mathbb{E}[\|\bar{u}^{i,n} - u^{i,*}\|_2^2] \leq \mathcal{O}(1/n)$ for $i \in \{1, \dots, m\}$ and $\mathbb{E}[\|\bar{Q}^n - Q^*\|_2^2] \leq \mathcal{O}(1/n)$.

Proof. First, we state a special case of Chung’s Lemma (see Chung [19]), which is useful for analyzing a specific type of sequence that arises often in recursive optimization algorithms (see, e.g., Juditsky et al. [34], Rakhlin et al. [52]).

Lemma 4 (Chung [19]). Consider a sequence $\{a^n\}$. Suppose that the following recursive inequality holds for some $b > 1$ and every $n \geq 1$:

$$a^n \leq \left(1 - \frac{b}{n}\right)a^{n-1} + \frac{c}{n^2}.$$

Then, if $k \geq \max\{c/(b-1), a^0\}$, it follows that $a^n \leq k/n$ for $n \geq 1$.

To prove the theorem, we use an induction argument for \bar{Q}^n . Let us first consider $t = T$ as the base case. Since $\mathbb{E}[\|\bar{Q}_T^n - Q_T^*\|_2^2] = 0$, it is trivially $\mathcal{O}(1/n)$, so let us take our induction hypothesis to be $\mathbb{E}[\|\bar{Q}_{t+1}^n - Q_{t+1}^*\|_2^2] \leq \mathcal{O}(1/n)$ for a particular $t < T$. Hence, there exists $k_{t+1}^q > 0$ independent of n such that $\mathbb{E}[\|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_2^2] \leq k_{t+1}^q/n$ holds for $n \geq 1$ (note the $n-1$ on the left hand side). We first apply Lemma 2 with $\gamma_t = 2/(\varepsilon C_{l_f})$ and $\kappa = \varepsilon C_{l_f}/C_{L_f}$, so by Lemma 4, we can take

$$k_t^i = 2 \max \left\{ \frac{2C_{L_f}}{\varepsilon^2 C_{l_f}^2} k_{t+1}^q + \frac{4C_{\alpha_i}}{\varepsilon^2 C_{l_f}^2}, \|\bar{u}_t^{i,0} - u_t^{i,*}\|_2^2 \right\}$$

to satisfy $\mathbb{E}[\|\bar{u}_t^n - u_t^{i,*}\|_2^2] \leq k_t^i/(2n)$ for $n \geq 1$. It follows that $\mathbb{E}[\|\bar{u}_t^{n-1} - u_t^{i,*}\|_2^2] \leq k_t^i/n$ for $n \geq 2$. Finally, we apply Lemma 3 with $\eta_t = 2/\varepsilon$, $\kappa_0 = \varepsilon/(2L_\Phi)$, and $\kappa_t = \varepsilon/(2mL_\Phi)$ for all i . Again by Lemma 4, we can choose

$$k_t^q = \max \left\{ \frac{4L_\Phi^2}{\varepsilon^2} k_{t+1}^q + \frac{4mL_\Phi^2}{\varepsilon^2} \sum_i k_t^i + \frac{4C_H}{\varepsilon^2}, \|\bar{Q}_t^0 - Q_t^*\|_2^2 \right\},$$

which gives $\mathbb{E}[\|\bar{Q}_t^n - Q_t^*\|_2^2] \leq k_t^q/n$ for $n \geq 2$, the $\mathcal{O}(1/n)$ rate needed to complete the inductive step. Therefore, we can conclude that

$$\mathbb{E}[\|\bar{Q}^n - Q^*\|_2^2] = \sum_{t=0}^T \mathbb{E}[\|\bar{Q}_t^n - Q_t^*\|_2^2] \leq \mathcal{O}(1/n),$$

as desired. As for the convergence of $\mathbb{E}[\|\bar{u}^{i,n} - u^{i,*}\|_2^2]$, we note that the values k_t^i obtained through the induction above can be used to deduce the $\mathcal{O}(1/n)$ convergence rate. \square

Similar to the sequences $\xi_{t+1}^{i,n}$ and $\xi_{t+1}^{q,n}$ used in the proofs of Lemma 1 and Theorem 2, let us introduce another useful process $\xi_{t+1}^{h,n}(s, a) \in \mathbb{R}^K$ for each t and (s, a) . As in Step 7 of Algorithm 2, let $w = W_{t+1}^{q,n}$. We define

$$\xi_{t+1}^{h,n}(s, a) = [H_t^*(w | s, a) - H_t^n(w | s, a)] \phi(w) p_t(w) p^u(w) \bar{p}_t^{n-1}(w | s, a)^{-1},$$

representing the error from using $H_t^n(w | s, a)$ to approximate $H_t^*(w | s, a)$.

Lemma 5. Under the ε -greedy sampling policy of (13) and Assumptions 1(vi)–5, it follows that

$$\sum_{n=1}^{\infty} \bar{\beta}_t^{n-1} \mathbb{E}[\|\xi_{t+1}^{h,n}(s, a)\|_2 | \mathcal{G}_{t+1}^{n-1}] < \infty \quad a.s.$$

for each t and (s, a) .

Proof. For simplicity, let $w = W_{t+1}^{q,n}$ and $P = \phi(w) p_t(w) p^u(w) \bar{p}_t^{n-1}(w | s, a)^{-1}$. We notice that

$$\|\xi_{t+1}^{h,n}(s, a)\|_2 \leq \|P\|_2 |H_t^*(w | s, a) - H_t^n(w | s, a)| \leq \|P\|_2 L_\Phi \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_2 + \|P\|_2 L_\Phi \sum_i \|\bar{u}_t^{i,n-1} - u_t^{i,*}\|_2,$$

where the second inequality follows from the same steps used in deriving (A.3) and then applying the monotonicity of l_p norms. Squaring, applying the inequality $2ab \leq a^2\kappa + b^2/\kappa$ for any $\kappa > 0$ to the cross terms, and using the fact that each component of P is bounded (due to p^u), we can write

$$\|\xi_{t+1}^{h,n}(s, a)\|_2^2 \leq C_Q \|\bar{Q}_{t+1}^{n-1} - Q_{t+1}^*\|_2^2 + \sum_i C_{u,i} \|\bar{u}_t^{i,n-1} - u_t^{i,*}\|_2^2, \tag{A.9}$$

for some constants $C_Q, C_{u,1}, \dots, C_{u,m} \geq 0$. Taking expectation of (A.9) and using the convergence rate result of Theorem 4, we see that $\mathbb{E}[\|\xi_{t+1}^{h,n}(s, a)\|_2^2] \leq \mathcal{O}(1/n)$. Now, by Cauchy-Schwarz and Assumption 4(iii),

$$\mathbb{E}[\bar{\beta}_t^{n-1} \mathbb{E}[\|\xi_{t+1}^{h,n}(s, a)\|_2 | \mathcal{G}_{t+1}^{n-1}]] = \mathbb{E}[\bar{\beta}_t^{n-1} \|\xi_{t+1}^{h,n}(s, a)\|_2] \leq \sqrt{\mathbb{E}[(\bar{\beta}_t^{n-1})^2] \mathbb{E}[\|\xi_{t+1}^{h,n}(s, a)\|_2^2]} \leq \mathcal{O}(n^{-1-\varepsilon/2}),$$

so it is clear that the terms $\mathbb{E}[\bar{\beta}_t^{n-1} \mathbb{E}[\|\xi_{t+1}^{h,n}(s, a)\|_2 | \mathcal{G}_{t+1}^{n-1}]]$ are summable. By the monotone convergence theorem, we conclude

$$\mathbb{E} \left[\sum_{n=1}^{\infty} \bar{\beta}_t^{n-1} \mathbb{E}[\|\xi_{t+1}^{h,n}(s, a)\|_2 | \mathcal{G}_{t+1}^{n-1}] \right] < \infty,$$

from which the statement of the lemma follows (notice that the term within the expectation must be finite almost surely). \square

Theorem 5 (Convergence of Risk-Directed Sampling Procedure). Choose initial approximations $\bar{Q}^0 \in \mathbb{R}^D$ and $\bar{u}^{i,0} \in \mathbb{R}^D$ for each $i \in \{1, 2, \dots, m\}$ so that $\bar{Q}_t^0 \in \mathcal{X}_t^q$ and $\bar{u}_t^{i,0} \in \mathcal{X}_t^u$ for all $t \in \mathcal{T}$. Under Assumptions 1(vi)–5, the ε -greedy sampling policy of (13), and the deterministic harmonic stepsizes given in (14), Algorithm 2 generates a sequence of iterates \bar{Q}^n that converge almost surely to the optimal value function Q^* . Moreover, the sampling coefficients $\bar{\theta}_t^n(s, a)$ converge to the optimal sampling coefficients under ϕ :

$$\bar{\theta}_t^n(s, a) \longrightarrow \Pi_\phi[|H_t^*(\cdot | s, a)|p_t(\cdot)] \quad a.s.$$

for each t and $(s, a) \in \mathcal{U}$.

Proof. First, notice that the convergence of \bar{Q}^n and $\bar{u}^{i,n}$ is (mostly) unaffected by the addition of the sampling procedure. By the principle of importance sampling, the factor of $\text{diag}(L_t^n)$ corrects, in expectation, for the fact that $W_{t+1}^{q,n}$ is sampled from the importance distribution $\bar{p}_t^{n-1}(w | s, a)$ rather than $p_t(w)$. In addition, Assumption 5 implies that the condition stating that the second moment of the gradient term is finite, which is needed for Kushner and Yin [40, Theorem 2.4], still holds and therefore the convergence follows.

We now focus on the last part of the theorem and analyze the convergence of the sampling coefficients. Let

$$\theta_t^*(s, a) = \Pi_\phi[|H_t^*(\cdot | s, a)|p_t(\cdot)],$$

and we aim to show $\bar{\theta}_t^n(s, a) \rightarrow \theta_t^*(s, a)$ almost surely. The proof technique is standard and uses a supermartingale convergence argument; see, e.g. Pflug [45, Theorem 5.3], but several aspects need to be adapted for technical reasons in our setting. Throughout this proof, fix a t and (s, a) . To simplify notation, we use the shorthand $\bar{p}_t^{n-1}(w) = \bar{p}_t^{n-1}(w | s, a)$, $w = W_{t+1}^{q,n} \sim \bar{p}_t^{n-1}(w)$, $H_t^n(w) = H_t^n(w | s, a)$, and $H_t^*(w) = H_t^*(w | s, a)$. First, we decompose the gradient term. Define a function $h_t: \mathbb{R}^K \rightarrow \mathbb{R}^K$, so that for $\theta \in \mathbb{R}^K$,

$$h_t(\theta) = \mathbf{E}[\theta^\top \phi(w) - |H_t^*(w)|p_t(w)]\phi(w)p^u(w)\bar{p}_t^{n-1}(w)^{-1}.$$

Note that h_t is affine in θ and is the gradient of the (strictly convex) quadratic objective function in (21), meaning that we can find positive constants $C_{h,1}$ and $C_{h,2}$ such that

$$\|h_t(\theta)\|_2^2 \leq C_{h,1}\|\theta - \theta_t^*(s, a)\|_2^2 + C_{h,2}. \quad (\text{A.10})$$

Next, the noise term $\epsilon_{t+1}^{h,n}(s, a)$ is given by

$$\epsilon_{t+1}^{h,n}(s, a) = [(\bar{\theta}_t^{n-1}(s, a))^\top \phi(w) - |H_t^*(w)|p_t(w)]\phi(w)p^u(w)\bar{p}_t^{n-1}(w)^{-1} - h_t(\bar{\theta}_t^{n-1}(s, a)),$$

and it follows that $\mathbf{E}[\epsilon_{t+1}^{h,n}(s, a) | \mathcal{G}_{t+1}^{n-1}] = 0$. Using the fact that $h_t(\bar{\theta}_t^{n-1}(s, a))$ is a deterministic affine function of $\bar{\theta}_t^{n-1}(s, a)$, we observe that $\epsilon_{t+1}^{h,n}(s, a)$ can be written in the form $X[\bar{\theta}_t^{n-1}(s, a) - \theta_t^*(s, a)] + Z$ where $X \in \mathbb{R}^{K \times K}$ and $Z \in \mathbb{R}^K$ are random. Since $p^u(w) = 0$ for w outside of a compact set, the entries of both X and Z are bounded. Let $\|\cdot\|_2$ denote the operator norm whenever its argument is a matrix and we have

$$\|\epsilon_{t+1}^{h,n}(s, a)\|_2^2 \leq [\|X\|_2\|\bar{\theta}_t^{n-1}(s, a) - \theta_t^*(s, a)\|_2 + \|Z\|_2]^2 \leq C_{\epsilon,1}\|\bar{\theta}_t^{n-1}(s, a) - \theta_t^*(s, a)\|_2^2 + C_{\epsilon,2}, \quad (\text{A.11})$$

for some constants $C_{\epsilon,1}$ and $C_{\epsilon,2}$. We also reproduce the definition of $\xi_{t+1}^{h,n}(s, a)$ given previously:

$$\xi_{t+1}^{h,n}(s, a) = [|H_t^*(w)| - |H_t^n(w)|]\phi(w)p_t(w)p^u(w)\bar{p}_t^{n-1}(w)^{-1}.$$

Step 7 of Algorithm 2 can therefore be written

$$\bar{\theta}_t^n(s, a) = [\bar{\theta}_t^{n-1}(s, a) - \beta_t^n(s, a)[h_t(\bar{\theta}_t^{n-1}(s, a)) + \epsilon_{t+1}^{h,n}(s, a) + \xi_{t+1}^{h,n}(s, a)]]^+.$$

For convenience, let us define

$$A_t^n = \bar{\theta}_t^n(s, a) - \theta_t^*(s, a) \quad \text{and} \quad \hat{h}_t^n = h_t(\bar{\theta}_t^{n-1}(s, a)) + \epsilon_{t+1}^{h,n}(s, a) + \xi_{t+1}^{h,n}(s, a).$$

Hence,

$$\|A_t^n\|_2^2 \leq \|[\bar{\theta}_t^{n-1}(s, a) - \beta_t^n(s, a)\hat{h}_t^n]^+ - [\theta_t^*(s, a)]^+\|_2^2 \leq \|A_t^{n-1}\|_2^2 - 2\beta_t^n(s, a)(A_t^{n-1})^\top \hat{h}_t^n + (\beta_t^n(s, a))^2 \|\hat{h}_t^n\|_2^2. \quad (\text{A.12})$$

Taking conditional expectation of the cross term, using Assumption 4, noting the inequality $(A_t^{n-1})^\top h_t(\bar{\theta}_t^{n-1}(s, a)) > 0$ (by strict convexity), and using Cauchy-Schwarz, we get

$$\mathbf{E}[-2\beta_t^n(s, a)(A_t^{n-1})^\top \hat{h}_t^n | \mathcal{G}_{t+1}^{n-1}] \leq -2\varepsilon\beta_t^{n-1}(A_t^{n-1})^\top h_t(\bar{\theta}_t^{n-1}(s, a)) + 2\beta_t^{n-1}\|A_t^{n-1}\|_2\mathbf{E}[\|\xi_{t+1}^{h,n}(s, a)\|_2 | \mathcal{G}_{t+1}^{n-1}] \quad (\text{A.13})$$

Moving onto the third term of (A.12), we have

$$\begin{aligned} \mathbf{E}[\|\hat{h}_t^n\|_2^2 | \mathcal{G}_{t+1}^{n-1}] &\leq [\|h_t(\bar{\theta}_t^{n-1}(s, a))\|_2^2 + \mathbf{E}[\|\epsilon_{t+1}^{h,n}(s, a)\|_2^2 + \|\xi_{t+1}^{h,n}(s, a)\|_2^2 + 2\|\epsilon_{t+1}^{h,n}(s, a)\|_2\|\xi_{t+1}^{h,n}(s, a)\|_2 \\ &\quad + 2\|h_t(\bar{\theta}_t^{n-1}(s, a))\|_2\|\epsilon_{t+1}^{h,n}(s, a)\|_2 + 2\|h_t(\bar{\theta}_t^{n-1}(s, a))\|_2\|\xi_{t+1}^{h,n}(s, a)\|_2 | \mathcal{G}_{t+1}^{n-1}]] \\ &\leq 3[\mathbf{E}[\|\epsilon_{t+1}^{h,n}(s, a)\|_2^2 + \|\xi_{t+1}^{h,n}(s, a)\|_2^2 | \mathcal{G}_{t+1}^{n-1}] + \|h_t(\bar{\theta}_t^{n-1}(s, a))\|_2^2] \\ &\leq [C_{\hat{h},1}\|A_t^{n-1}\|_2^2 + C_{\hat{h},2} + 3\mathbf{E}[\|\xi_{t+1}^{h,n}(s, a)\|_2^2 | \mathcal{G}_{t+1}^{n-1}]] \end{aligned} \quad (\text{A.14})$$

where the second inequality is due to the inequality $2ab \leq a^2 + b^2$ and the third inequality, with new constants $C_{\hat{h},1}$ and $C_{\hat{h},2}$, is due to the bounds (A.10) and (A.11). Combining (A.12)–(A.14),

$$\mathbf{E}[\|A_t^n\|_2^2 | \mathcal{G}_{t+1}^{n-1}] \leq \|A_t^{n-1}\|_2^2(1 + \zeta_t^{n-1}) + \mu_t^{n-1} - \nu_t^{n-1},$$

where

$$\begin{aligned} \zeta_t^{n-1} &= C_{\hat{h},1}(\tilde{\beta}_t^{n-1})^2 + 2\tilde{\beta}_t^{n-1}\mathbf{E}[\|\xi_{t+1}^{h,n}(s,a)\|_2 | \mathcal{G}_{t+1}^{n-1}], \\ \mu_t^{n-1} &= 3(\tilde{\beta}_t^{n-1})^2\mathbf{E}[\|\xi_{t+1}^{h,n}(s,a)\|_2^2 | \mathcal{G}_{t+1}^{n-1}] + C_{\hat{h},2}(\tilde{\beta}_t^{n-1})^2, \\ \nu_t^{n-1} &= 2\varepsilon\tilde{\beta}_t^{n-1}(A_t^{n-1})^\top h_t(\bar{\theta}_t^{n-1}(s,a)). \end{aligned}$$

Note that $\sum_{n=1}^\infty (\tilde{\beta}_t^{n-1})^2 < \infty$ almost surely by Assumption 4(iii) and the monotone convergence theorem (we can apply the same logic as Lemma 5). This, together with Lemma 5, allows us to conclude that $\sum_{n=1}^\infty \zeta_t^{n-1} < \infty$ almost surely. Moreover, since $\tilde{\beta}_t^{n-1}$ must converge to zero and (A.9) implies $\mathbf{E}[\|\xi_{t+1}^{h,n}(s,a)\|_2^2 | \mathcal{G}_{t+1}^{n-1}]$ converges to zero, we can see that $\sum_{n=1}^\infty \mu_t^{n-1} < \infty$ almost surely. The well-known supermartingale convergence lemma of Robbins and Siegmund [55] tells us that $\|A_t^n\|_2$ converges and $\sum_{n=1}^\infty \nu_t^{n-1} < \infty$ almost surely. Let $D_\delta = \{\lim_{n \rightarrow \infty} \|A_t^n\|_2 > \delta\}$. On the event D_δ , using strict convexity, we know that $(A_t^{n-1})^\top h_t(\bar{\theta}_t^{n-1}(s,a))$ is positive and bounded away from zero for any n . Along with $\sum_{n=1}^\infty \tilde{\beta}_t^{n-1} = \infty$, this implies $\sum_{n=1}^\infty \nu_t^{n-1} = \infty$, which in turn shows us that D_δ must occur with probability zero for any $\delta > 0$. \square

References

- [1] Acciaio B, Penner I (2011) Dynamic risk measures. Di Nunno, G, Øksendal, B, eds. *Advanced Mathematical Methods for Finance* (Springer, Berlin), 1–34.
- [2] Al-Qaq WA, Devetsikiotis M, Townsend JK (1995) Stochastic gradient optimization of importance sampling for the efficient simulation of digital communication systems. *IEEE Trans. Comm.* 43(12):2975–2985.
- [3] Artzner P, Delbaen F, Eber J, Heath D (1999) Coherent measures of risk. *Math. Finance* 9(3):203–228.
- [4] Artzner P, Delbaen F, Eber J, Heath D, Ku H (2006) Coherent multiperiod risk adjusted values and Bellman’s principle. *Ann. Oper. Res.* 152(1):5–22.
- [5] Azar MG, Ghavamzadeh M, Kappen HJ (2011) Reinforcement learning with a near optimal rate of convergence. Technical Report, Radboud University Nijmegen, Netherlands.
- [6] Bardou O, Frikha N (2009) Computing VaR and CVaR using stochastic approximation and adaptive unconstrained importance sampling. *Monte Carlo Methods Appl.* 15(3):173–210.
- [7] Belles-Sampera J, Guillén M, Santolino M (2014) Beyond value-at-risk: GlueVaR distortion risk measures. *Risk Anal.* 34(1):121–134.
- [8] Bertimas D, Lo AW (1998) Optimal control of execution costs. *J. Financial Markets* 1(1):1–50.
- [9] Bertsekas DP, Tsitsiklis JN (1996) *Neuro-Dynamic Programming* (Athena Scientific, Belmont, MA).
- [10] Boda K, Filar JA (2006) Time consistent dynamic risk measures. *Math. Methods Oper. Res.* 63(1):169–186.
- [11] Boyd S, Vandenberghe L (2004) *Convex Optimization* (Cambridge University Press, Cambridge, UK).
- [12] Breiman L (1992) *Probability* (SIAM, Philadelphia).
- [13] Bucklew J (2004) *Introduction to Rare Event Simulation* (Springer, New York).
- [14] Byström H (2005) Extreme value theory and extremely large electricity price changes. *Internat. Rev. Econom. Finance* 14(1):41–55.
- [15] Çavus O, Ruszczyński A (2014) Computational methods for risk-averse undiscounted transient Markov models. *Oper. Res.* 62(2):401–417.
- [16] Cheridito P, Stadje M (2009) Time-inconsistency of VaR and time-consistent alternatives. *Finance Res. Lett.* 6(1):40–46.
- [17] Cheridito P, Delbaen F, Kupper M (2006) Dynamic monetary risk measures for bounded discrete-time processes. *Electronic J. Probab.* 11(3):57–106.
- [18] Chow Y, Ghavamzadeh M (2014) Algorithms for CVaR optimization in MDPs. *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA), 3509–3517.
- [19] Chung KL (1954) On a stochastic approximation method. *Ann. Math. Statist.* 25(3):463–483.
- [20] Danielsson J, Jorgensen B, Samorodnitsky G, Sarma M, de Vries C (2005) Subadditivity reexamined: The case for value-at-risk. Discussion paper, 549. Financial Markets Group, London School of Economics and Political Science, London, UK.
- [21] Dhaene J, Laeven RJA, Vanduffel S, Darkiewicz G, Goovaerts MJ (2006) Can a coherent risk measure be too subadditive? *J. Risk Insurance* 75(2):365–386.
- [22] Dowd K, Blake D (2006) After VaR: The theory, estimation, and insurance applications of quantile-based risk measures. *J. Risk Insurance* 73(2):193–229.
- [23] Duffie D, Pan J (1997) An overview of value at risk. *J. Derivatives* 4(3):7–49.
- [24] Egloff D, Leippold M (2010) Quantile estimation with adaptive importance sampling. *Ann. Statist.* 38(2):1244–1278.
- [25] Enders J, Powell WB, Egan D (2010) A dynamic model for the failure replacement of aging high-voltage transformers. *Energy Systems* 1(1):31–59.
- [26] Even-Dar E, Mansour Y (2004) Learning rates for Q-learning. *J. Machine Learn. Res.* 5:1–25.
- [27] Frittelli M, Gianin ER (2004) Dynamic convex risk measures. Szegö, *Risk Measures for the 21st Century* (John Wiley & Sons, Chichester, UK), 227–248.
- [28] George AP, Powell WB (2006) Adaptive stepsizes for recursive estimation with applications in approximate dynamic programming. *Machine Learn.* 65(1):167–198.
- [29] Glasserman P, Liu TW (1996) Rare-event simulation for multistage production-inventory systems. *Management Sci.* 42(9):1292–1307.
- [30] Ibragimov R, Walden J (2007) The limits of diversification when losses may be large. *J. Banking and Finance* 31(8):2551–2569.
- [31] Jaakkola T, Jordan MI, Singh SP (1994) On the convergence of stochastic iterative dynamic programming algorithms. *Neural Computation* 6:1185–1201.

- [32] Jiang DR, Powell WB (2015) An approximate dynamic programming algorithm for monotone value functions. *Oper. Res.* 63(6):1489–1511.
- [33] Jiang DR, Powell WB (2015) Optimal hour-ahead bidding in the real-time electricity market with battery storage using approximate dynamic programming. *INFORMS J. Comput.* 27(3):525–543.
- [34] Juditsky A, Lan G, Nemirovski A, Shapiro A (2009) Robust stochastic approximation approach to stochastic programming. *SIAM J. Optim.* 19(4):1574–1609.
- [35] Kan YS (2011) On the convergence of a stochastic approximation procedure for estimating the quantile criterion in the case of a discontinuous distribution function. *Automation Remote Control* 72(2):283–288.
- [36] Kearns MJ, Singh SP (1999) Finite-sample convergence rates for Q-learning and indirect algorithms. *Advances in Neural Information Processing Systems*, 996–1002.
- [37] Kim JH, Powell WB (2011) An hour-ahead prediction model for heavy-tailed spot prices. *Energy Econom.* 33(6):1252–1266.
- [38] Kleywegt AJ, Shapiro A, Homem-de Mello T (2002) The sample average approximation method for stochastic discrete optimization. *SIAM J. Optimization* 12(2):479–502.
- [39] Kozmík V, Morton DP (2015) Evaluating policies in risk-averse multi-stage stochastic programming. *Math. Programming* 152(1):275–300.
- [40] Kushner HJ, Yin GG (2003) *Stochastic Approximation and Recursive Algorithms and Applications*, Vol. 35 (Springer, New York).
- [41] Löhndorf N, Wozabal D, Minner S (2013) Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Oper. Res.* 61(4):810–823.
- [42] Nascimento JM, Powell WB (2009) An optimal approximate dynamic programming algorithm for the lagged asset acquisition problem. *Math. Oper. Res.* 34(1):210–237.
- [43] Neise F (2008) *Risk Management in Stochastic Integer Programming* (Springer Fachmedien, Wiesbaden, Germany).
- [44] Pereira MVE, Pinto LMVG (1991) Multi-stage stochastic optimization applied to energy planning. *Math. Programming* 52(1):359–375.
- [45] Pflug GC (1996) *Optimization of Stochastic Models* (Kluwer Academic Publishers, New York).
- [46] Pflug GC, Ruszczyński A (2005) Measuring risk for income streams. *Computational Optim. Appl.* 32(1–2):161–178.
- [47] Philpott AB, de Matos VL (2012) Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *Eur. J. Oper. Res.* 218(2):470–483.
- [48] Philpott AB, de Matos VL, Finardi E (2013) On solving multistage stochastic programs with coherent risk measures. *Oper. Res.* 61(4):957–970.
- [49] Powell WB (2011) *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, 2nd ed. (John Wiley & Sons, New York).
- [50] Powell WB, Ruszczyński A, Topaloglu H (2004) Learning algorithms for separable approximations of discrete stochastic optimization problems. *Math. Oper. Res.* 29(4):814–836.
- [51] Puterman ML (2014) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (Wiley, New York).
- [52] Rakhlin A, Shamir O, Sridharan K (2012) Making gradient descent optimal for strongly convex stochastic optimization. *Proc. 29th Internat. Conf. Machine Learn. ICML '12* (Omnipress, Madison, WI), 449–456.
- [53] Riedel F (2004) Dynamic coherent risk measures. *Stochastic Processes and Their Appl.* 112(2):185–200.
- [54] Robbins H, Monro S (1951) A stochastic approximation method. *The Annals of Mathematical Statistics* 22(3):400–407.
- [55] Robbins H, Siegmund D (1971) A convergence theorem for nonnegative almost supermartingales and some applications. Rustagi JS, ed. *Optimizing Methods in Statistics* (Academic Press, New York), 233–257.
- [56] Rockafellar RT, Uryasev S (2000) Optimization of conditional value-at-risk. *J. Risk* 2:21–41.
- [57] Rockafellar RT, Uryasev S (2002) Conditional value-at-risk for general loss distributions. *J. Banking and Finance* 26(7):1443–1471.
- [58] Rubinstein RY (1999) The simulated entropy method for combinatorial and continuous optimization. *Methodology Comput. Appl. Probab.* 1(2):127–190.
- [59] Rudloff B, Street A, Valladão DM (2014) Time consistency and risk averse dynamic decision models: Definition, interpretation and practical consequences. *Eur. J. Oper. Res.* 234(3):743–750.
- [60] Ruszczyński A (2010) Risk-averse dynamic programming for Markov decision processes. *Math. Programming* 125(2):235–261.
- [61] Ruszczyński A, Shapiro A (2006) Conditional risk mappings. *Math. Oper. Res.* 31(3):544–561.
- [62] Ruszczyński A, Shapiro A (2006) Optimization of convex risk functions. *Math. Oper. Res.* 31(3):433–452.
- [63] Ryu EK, Boyd SP (2015) Adaptive importance sampling via stochastic convex programming. Preprint arXiv:1412.4845.
- [64] Ryzhov IO, Powell WB (2011) Information collection on a graph. *Oper. Res.* 59(1):188–201.
- [65] Ryzhov IO, Frazier PI, Powell WB (2015) A new optimal stepsize for approximate dynamic programming. *IEEE Trans. Automatic Control* 60(3):743–758.
- [66] Schaul T, Zhang S, LeCun Y (2013) No more pesky learning rates. *J. Machine Learn. Res.* 28(2):343–351.
- [67] Serrada EN, Bronshtein EM, Rachev ST, Fabozzi FJ, Sun EW, Stoyanov SV (2010) Distortion risk measures in portfolio optimization. *Handbook of Portfolio Construction* (Springer, New York), 649–673.
- [68] Shapiro A (2009) On a time consistency concept in risk averse multistage stochastic programming. *Oper. Res. Lett.* 37(3):143–147.
- [69] Shapiro A, Tekaya W, da Costa JP, Soares MP (2013) Risk neutral and risk averse stochastic dual dynamic programming method. *Eur. J. Oper. Res.* 224(2):375–391.
- [70] Szepesvari C, Littman ML (1996) Generalized Markov decision processes: Dynamic programming and reinforcement learning algorithms. Technical Report, Bolyai Institute of Mathematics, Hungary.
- [71] Tierney L (1983) A space-efficient recursive procedure for estimating a quantile of an unknown distribution. *SIAM J. Scientific Statist. Comput.* 4(4):706–711.
- [72] Tsitsiklis JN (1994) Asynchronous stochastic approximation and Q-learning. *Machine Learn.* 16(3):185–202.
- [73] Watkins CJ, Dayan P (1992) Q-learning. *Machine Learn.* 8(3–4):279–292.
- [74] Whittle P (1980) Multi-armed bandits and the Gittins Index. *J. Roy. Statist. Soc. Ser. B (Methodological)* 42(2):143–149.
- [75] Xi X, Sioshansi R, Marano V (2014) A stochastic dynamic programming model for co-optimization of distributed energy storage. *Energy Systems* 5(3):475–505.