



Brief paper

Mean–variance optimization of discrete time discounted Markov decision processes[☆]

Li Xia

CFINS, Department of Automation, TNLIS, Tsinghua University, Beijing 100084, China



ARTICLE INFO

Article history:

Received 13 September 2016
 Received in revised form 2 June 2017
 Accepted 28 August 2017

Keywords:

Markov decision process
 Mean–variance optimization
 Variance criterion
 Sensitivity-based optimization

ABSTRACT

In this paper, we study a mean–variance optimization problem in an infinite horizon discrete time discounted Markov decision process (MDP). The objective is to minimize the variance of system rewards with the constraint of mean performance. Different from most of works in the literature which require the mean performance already achieve optimum, we can let the discounted performance equal any constant. The difficulty of this problem is caused by the quadratic form of the variance function which makes the variance minimization problem not a standard MDP. By proving the decomposable structure of the feasible policy space, we transform this constrained variance minimization problem to an equivalent unconstrained MDP under a new discounted criterion and a new reward function. The difference of the variances of Markov chains under any two feasible policies is quantified by a difference formula. Based on the variance difference formula, a policy iteration algorithm is developed to find the optimal policy. We also prove the optimality of deterministic policy over the randomized policy generated in the mean-constrained policy space. Numerical experiments demonstrate the effectiveness of our approach.

© 2017 Elsevier Ltd. All rights reserved.

1. Introduction

The mean–variance optimization is an important problem in stochastic optimization and its origin can go back to the pioneering work by H. Markowitz, 1990 Nobel Laureate in Economics, on the modern portfolio management (Markowitz, 1952). In financial engineering, the mean indicates the return of assets and the variance indicates the risk of assets. The objective of the mean–variance optimization is to find an optimal policy such that the mean and the variance of system rewards are optimized in tradeoff and the efficient frontier (a curve comprised of *Pareto optima*) is obtained.

The mean–variance optimization is first proposed in a static optimization form in Markowitz's original paper (Markowitz, 1952), in which the decision variables are the investment percentage of securities and the securities returns are described as random variables with known means and variances. Then, the mean–variance optimization is further studied in a dynamic optimization form and Markov decision processes (MDPs) are widely used as an important analytical model. The difficulty of this problem mainly comes from the *non-additiveness* of the variance criterion which

makes the principle of *consistent choice* in dynamic programming invalid (Sobel, 1982; Xia, 2016). Such invalidness means that the optimal action selection during $[t + 1, \infty)$ may be not optimal for the action selection during $[t, \infty)$. In the literature, there are different ways to study the mean–variance optimization. Many works studied the variance minimization of MDPs after the mean performance is already maximized (Guo & Song, 2009; Hernandez-Lerma, Vega-Amaya, & Carrasco, 1999; Huang & Chen, 2012). For such problem, the variance minimization problem can be transformed to another standard MDP under an equivalent average or discounted criterion. There are also studies that use the policy gradient approach to study the mean–variance optimization when the policy is parameterized (Prashantha & Ghavamzadeh, 2013; Tamar, Castro, & Mannor, 2012).

Sobel (1982) gave an early study on the mean–variance optimization in a discrete time discounted Markov chain, but no optimization algorithm was presented in that paper. Chung (1994) and Sobel (1994) studied the variance minimization problem in a discrete time Markov chain with the constraint that the long-run average performance is larger than a given constant. This problem was transformed to a sequence of linear programming problems, which may have concerns of computation efficiency since the number of sequential problems may be large. Guo, Huang, and Zhang (2015) and Huo, Zou, and Guo (2017) studied the mean–variance optimization problem in a continuous time Markov chain with unbounded transition rates and state–action dependent discount factors, where the performance is accumulated until a

[☆] The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Bart De Schutter under the direction of Editor Christos G. Cassandras.

E-mail address: xial@tsinghua.edu.cn.

certain state is reached. There are certainly numerous other excellent works about the mean–variance optimization in the literature. However, most of the works in the literature either require a condition of optimal mean performance or reformulate the problem as variations of mathematical programming. Although linear programming may be used to study the mean–variance optimization in some cases, it does not utilize the structure of Markov systems and the efficiency is not satisfactory. Policy iteration is a classical approach in dynamic programming and it usually has a high convergence efficiency. There is little work to study the mean–variance optimization using policy iteration, at the condition that the mean performance equals a given value.

In this paper, we study a mean–variance optimization problem in an infinite horizon discrete time discounted Markov chain. The objective is to find the optimal policy with the minimal variance of rewards from the policy set in which the discounted performance equals a given constant. The motivation of this problem can be explained with a financial example. People may not always choose an asset portfolio with the maximal expected return, since a portfolio with big return usually has big risk (quantified by variance). People always like to seek a portfolio with minimal risk and acceptable return. The solution with minimal risk and fixed return is called *Pareto optimum*. All the Pareto optimal solutions compose a curve called *Pareto frontier*, or *efficient frontier* in financial engineering.

The difficulty of such mean–variance optimization problem mainly comes from two aspects. The first one is the difficulty caused by the non-additiveness of the variance criterion, which makes the mean–variance optimization not a standard MDP and policy iteration is not applicable directly. Another difficulty comes from the fact that the policy set with a fixed mean performance usually has no satisfactory structure, such as that described later in [Theorem 1](#). For example, the policy set whose long-run average performance equals a given constant may not be decomposable as that in [Theorem 1](#). Without such property, dynamic programming and policy iteration cannot be used for these problems.

In this paper, we use the sensitivity-based optimization theory to study this nonstandard MDP problem. For the policy set in which the discounted performance equals a given constant, we prove that this policy set is decomposable on the action space and the action can be chosen independently at every state. A difference formula is derived to quantify the variance difference under any two feasible policies. The original variance minimization problem with constraints is transformed to a standard unconstrained MDP under an equivalent discounted criterion with a new discount factor β^2 and a new reward function, where β is the discount factor of the original Markov chain. With this equivalent MDP, we prove the existence of the optimal policy for this mean–variance optimization problem. A policy iteration algorithm is developed to find the optimal policy with the minimal variance. The optimality of deterministic policy is also proved, compared with randomized policies generated in the mean-constrained policy space. Finally, we conduct a numerical experiment to demonstrate the effectiveness of our approach. The efficient frontier of this numerical example is also analyzed.

This paper is a continued work compared with our previous papers ([Xia, 2016, 2017](#)), which aim to minimize the variance of the long-run average performance of the Markov chain without considering the constraint of mean performance. The targeted models in these papers are different, so are the main results. To the best of our knowledge, this is the first paper that develops a policy iteration algorithm to minimize the variance of a discrete time discounted Markov chain at the condition of any given discounted performance.

2. Problem formulation

We consider a finite MDP in discrete time. X_t is denoted as the system state at time t , $t = 0, 1, \dots$. The state space is finite and denoted as $S = \{1, 2, \dots, S\}$, where S is the size of the state space. We only consider the deterministic and stationary policy d which is a mapping from the state space to the action space. If the current state is i , the policy d determines to choose an action a from a finite action space $\mathcal{A}(i)$ and a system reward $r(i, a)$ is obtained. The system will transit to a new state j with a transition probability $p(j|i, a)$ at the next time epoch, where $i, j \in S$ and $a \in \mathcal{A}(i)$. Obviously, we have $\sum_{j \in S} p(j|i, a) = 1$. Since d is a mapping in the state space, we have $a = d(i)$ and $d(i) \in \mathcal{A}(i)$ for all $i \in S$. We define the policy space \mathcal{D} as the family of all deterministic stationary policies. For each $d \in \mathcal{D}$, $\mathbf{P}(d)$ is denoted as a transition probability matrix and its (i, j) th element is $p(j|i, d(i))$, and $\mathbf{r}(d)$ is denoted as a column vector and its i th element is $r(i, d(i))$. We assume that the Markov chain is ergodic for any policy in \mathcal{D} . The discount factor of the MDP is β , $0 < \beta < 1$. For initial state i , the discounted performance of the MDP under policy d is defined as below.

$$J(d, i) := \mathbb{E}_i^d \left[\sum_{t=0}^{\infty} \beta^t r(X_t, d(X_t)) \right], \quad i \in S, \quad (1)$$

where \mathbb{E}_i^d is an expectation operator of the Markov chain at the condition that the initial state is i and the policy is d . $\mathbf{J}(d)$ is an S -dimensional column vector and its i th element is $J(d, i)$. The variance of the discounted Markov chain is defined as below.

$$\sigma^2(d, i) := \mathbb{E}_i^d \left[\left(\sum_{t=0}^{\infty} \beta^t r(X_t, d(X_t)) \right) - J(d, i) \right]^2, \quad i \in S. \quad (2)$$

We observe that $\sigma^2(d, i)$ quantifies the variance of the limiting random variable $\sum_{t=0}^{\infty} \beta^t r(X_t, d(X_t))$. $\sigma^2(d)$ is the variance vector of the discounted Markov chain and its i th element is $\sigma^2(d, i)$.

Denote λ as a given mean reward function on S . That is, λ is an S -dimensional column vector and its i th element is denoted as $\lambda(i)$, $i \in S$. The set of all feasible policies with which the discounted performance of the Markov chain equals λ is defined as below.

$$\mathcal{D}_\lambda := \{d \in \mathcal{D} | J(d, i) = \lambda(i), \text{ for all } i \in S\}. \quad (3)$$

Note that \mathcal{D} is a deterministic stationary policy set, so is \mathcal{D}_λ . In this paper, we do not consider randomized stationary policies. The optimality of deterministic policies will be studied in the next section, see [Theorem 5](#). It is easy to see that the policy set \mathcal{D}_λ may be empty if the value of λ is not chosen properly. In this paper, we assume that \mathcal{D}_λ is not empty, which is similar to the assumption in Markowitz's mean–variance portfolio problem ([Markowitz, 1952; Zhou & Yin, 2003](#)). For a given discounted performance vector λ , \mathcal{D}_λ may contain more than one policy. The objective is to find an optimal policy from \mathcal{D}_λ such that the variance of the Markov chain is minimized. The mathematical formulation is written as below.

$$\min_{d \in \mathcal{D}_\lambda} \{ \sigma^2(d, i) \}, \quad \text{for all } i \in S. \quad (4)$$

That is, we aim to find an optimal policy among all feasible policies whose discounted performance is equal to a given constant vector λ , such that the variance of discounted rewards is minimized. We denote such a *mean–variance optimal policy* as d_λ^* . The existence of the solution d_λ^* to the problem (4) is not guaranteed because the minimization in (4) is over every state $i \in S$, i.e., (4) can be viewed as a multi-objective optimization problem. For a general multi-objective optimization problem, it is possible that no solution can dominate all the other solutions on the value of every dimension of objective function. In the next section, we will discuss the existence of such optimal policy d_λ^* and develop an optimization algorithm to find it. Moreover, we have the following remarks.

Remark 1. In the literature on mean–variance optimization of MDPs, most of the works study the variance minimization at the constraint that the mean performance is already maximized. As a comparison in our problem (4), λ can be given as any feasible value.

Remark 2. If d_λ^* exists for a given λ , then d_λ^* is also called an *efficient policy* for the mean–variance optimization problem. $(\sigma^2(d_\lambda^*), \lambda)$ is called an *efficient point* and all the efficient points with different λ 's compose the *efficient frontier* of that problem.

3. Main results

We use the sensitivity-based optimization theory to study this mean–variance optimization problem. This theory was proposed by Dr. X.-R. Cao and the key idea is the difference formula that quantifies the performance difference of Markov systems under any two policies (Cao, 2007).

Consider a policy $d \in \mathcal{D}$. For simplicity, we omit d by default in the following notations \mathbf{J} , σ^2 , \mathbf{P} , \mathbf{r} . Consider another arbitrary policy $d' \in \mathcal{D}$, and the corresponding notations under policy d' are written as \mathbf{J}' , σ'^2 , \mathbf{P}' , \mathbf{r}' . For the discounted performance defined in (1), we can rewrite it in a recursive form as below.

$$\mathbf{J} = \mathbf{r} + \beta \mathbf{P} \mathbf{J}. \quad (5)$$

We further have

$$\mathbf{J} = (\mathbf{I} - \beta \mathbf{P})^{-1} \mathbf{r}, \quad (6)$$

where \mathbf{I} is an S -dimensional identity matrix. Note that the matrix $(\mathbf{I} - \beta \mathbf{P})$ is invertible and $(\mathbf{I} - \beta \mathbf{P})^{-1}$ is always nonnegative based on the following observations

$$(\mathbf{I} - \beta \mathbf{P})^{-1} = \sum_{n=0}^{\infty} \beta^n \mathbf{P}^n, \quad (7)$$

and \mathbf{P} is a nonnegative matrix. We further observe that if the Markov chain with \mathbf{P} is ergodic, all the elements of $(\mathbf{I} - \beta \mathbf{P})^{-1}$ are positive.

For a new policy $d' \in \mathcal{D}$, we similarly have

$$\mathbf{J}' = \mathbf{r}' + \beta \mathbf{P}' \mathbf{J}'. \quad (8)$$

Subtracting (5) from (8), we have

$$\mathbf{J}' - \mathbf{J} = \mathbf{r}' - \mathbf{r} + \beta(\mathbf{P}' - \mathbf{P})\mathbf{J} + \beta \mathbf{P}'(\mathbf{J}' - \mathbf{J}). \quad (9)$$

Therefore, we derive the following formula about the difference of the discounted performance of Markov chains

$$\mathbf{J}' - \mathbf{J} = (\mathbf{I} - \beta \mathbf{P}')^{-1} [\beta(\mathbf{P}' - \mathbf{P})\mathbf{J} + \mathbf{r}' - \mathbf{r}]. \quad (10)$$

In the above formula, \mathbf{J} in the right-hand side is also called the *performance potential* (or value function) in the sensitivity-based optimization theory (Cao, 2007). For a discounted Markov chain, the performance potential is the same as the discounted performance, but they are different in other general cases.

Remark 3. We can derive a policy iteration algorithm directly based on (10). With (7), we see that all the elements of $(\mathbf{I} - \beta \mathbf{P}')^{-1}$ are positive for ergodic Markov chains. If we choose $(\mathbf{P}', \mathbf{r}')$ to make the vector $\beta \mathbf{P}' \mathbf{J} + \mathbf{r}'$ as large as possible, then we see $\mathbf{J}' - \mathbf{J} \geq 0$ and the policy is improved.

Based on the difference formula (10), we study the structure of the policy set \mathcal{D}_λ in which the discounted performance equals a given constant λ . The following theorem describes the decomposable structure of \mathcal{D}_λ .

Theorem 1. For a given constant λ , the policy set \mathcal{D}_λ in which the discounted performance equals λ is the Cartesian product of feasible action space at every state, i.e.,

$$\mathcal{D}_\lambda = \mathcal{A}_\lambda(1) \times \mathcal{A}_\lambda(2) \times \cdots \times \mathcal{A}_\lambda(S), \quad (11)$$

where $\mathcal{A}_\lambda(i)$, $i \in S$, is defined as

$$\mathcal{A}_\lambda(i) := \{all\ a \in \mathcal{A}(i) | r(i, a) + \beta \sum_{j \in S} p(j|i, a) \lambda(j) = \lambda(i)\}. \quad (12)$$

Proof. Suppose $d \in \mathcal{D}_\lambda$ and the corresponding transition probability matrix and reward function are denoted as \mathbf{P} and \mathbf{r} , respectively. With (5), we have

$$\lambda = \mathbf{r} + \beta \mathbf{P} \lambda. \quad (13)$$

Consider an arbitrary policy $d' \in \mathcal{D}$ with \mathbf{P}' , \mathbf{r}' , and the corresponding discounted performance is denoted as \mathbf{J}' .

First, we want to prove that if $d' \in \mathcal{A}_\lambda(1) \times \mathcal{A}_\lambda(2) \times \cdots \times \mathcal{A}_\lambda(S)$, then $d' \in \mathcal{D}_\lambda$. Since $d' \in \mathcal{A}_\lambda(1) \times \mathcal{A}_\lambda(2) \times \cdots \times \mathcal{A}_\lambda(S)$, we have

$$r(i, d'(i)) + \beta \sum_{j \in S} p(j|i, d'(i)) \lambda(j) = \lambda(i), \quad i \in S. \quad (14)$$

That is, we have

$$\mathbf{r}' + \beta \mathbf{P}' \lambda = \lambda. \quad (15)$$

Substituting (13) into the above equation, we have

$$\mathbf{r}' + \beta \mathbf{P}' \lambda = \mathbf{r} + \beta \mathbf{P} \lambda. \quad (16)$$

With the difference formula (10), we directly have

$$\mathbf{J}' - \lambda = (\mathbf{I} - \beta \mathbf{P}')^{-1} [\beta(\mathbf{P}' - \mathbf{P})\lambda + \mathbf{r}' - \mathbf{r}]. \quad (17)$$

Substituting (16) into (17), we obtain

$$\mathbf{J}' - \lambda = (\mathbf{I} - \beta \mathbf{P}')^{-1} \mathbf{0} = \mathbf{0}. \quad (18)$$

Therefore, $\mathbf{J}' = \lambda$ and $d' \in \mathcal{D}_\lambda$.

Second, we want to prove that if $d' \in \mathcal{D}_\lambda$, then $d' \in \mathcal{A}_\lambda(1) \times \mathcal{A}_\lambda(2) \times \cdots \times \mathcal{A}_\lambda(S)$. Since $d' \in \mathcal{D}_\lambda$, we know $\mathbf{J}' = \lambda$. Therefore, the difference formula (17) becomes

$$\beta(\mathbf{P}' - \mathbf{P})\lambda + \mathbf{r}' - \mathbf{r} = (\mathbf{I} - \beta \mathbf{P}')(\mathbf{J}' - \lambda) = \mathbf{0}. \quad (19)$$

Substituting (13) into (19), we can directly derive (15) and its componentwise form (14). Therefore, we obtain $d' \in \mathcal{A}_\lambda(1) \times \mathcal{A}_\lambda(2) \times \cdots \times \mathcal{A}_\lambda(S)$.

Combining the above two results, we have $\mathcal{D}_\lambda = \mathcal{A}_\lambda(1) \times \mathcal{A}_\lambda(2) \times \cdots \times \mathcal{A}_\lambda(S)$ and the theorem is proved. \square

With Theorem 1, we know that \mathcal{D}_λ is decomposable as the product of all feasible action spaces $\mathcal{A}_\lambda(i)$, which indicates that the action selections at different states in the policy set \mathcal{D}_λ are independent. This property is important for us to develop policy iteration algorithms.

If the performance constraint is quantified under the long-run average criterion, the decomposable structure of \mathcal{D}_λ may not hold, which makes the policy iteration inapplicable to such problem. It is partly because the analog of (12) is not valid for the long-run average case. If the constraint is that the long-run average performance equals the maximum, \mathcal{D}_λ is decomposable and the policy iteration is applicable to the variance criterion (Guo & Song, 2009; Hernandez-Lerma et al., 1999).

We further study how to minimize the variance of the Markov chain. Consider a policy $d \in \mathcal{D}$, the variance is defined as (2). The

symbol d is also omitted in the following notations for simplicity. We rewrite (2) as follows:

$$\begin{aligned}\sigma^2(i) &= \mathbb{E}_i \left[\left(\sum_{t=0}^{\infty} \beta^t r(X_t) \right) - J(i) \right]^2 \\ &= \mathbb{E}_i \left[r(X_0) + \sum_{t=1}^{\infty} \beta^t r(X_t) \right]^2 - J^2(i) \\ &= r^2(i) + 2\beta r(i) \sum_{j \in \mathcal{S}} p(j|i) \mathbb{E}_j \left[\sum_{t=0}^{\infty} \beta^t r(X_t) \right] \\ &\quad + \beta^2 \sum_{j \in \mathcal{S}} p(j|i) \mathbb{E}_j \left[\sum_{t=0}^{\infty} \beta^t r(X_t) \right]^2 - J^2(i),\end{aligned}\quad (20)$$

where the last equality holds using the conditional expectation and the Markovian property. Recursively substituting (1) and (2) into (20), we obtain

$$\begin{aligned}\sigma^2(i) &= r^2(i) + 2\beta r(i) \sum_{j \in \mathcal{S}} p(j|i) J(j) + \beta^2 \sum_{j \in \mathcal{S}} p(j|i) J^2(j) \\ &\quad - J^2(i) + \beta^2 \sum_{j \in \mathcal{S}} p(j|i) \sigma^2(j).\end{aligned}\quad (21)$$

For the current policy d and state i , we define a new reward function $h(i)$ as below, $i \in \mathcal{S}$.

$$h(i) = r^2(i) + 2\beta r(i) \sum_{j \in \mathcal{S}} p(j|i) J(j) + \beta^2 \sum_{j \in \mathcal{S}} p(j|i) J^2(j) - J^2(i). \quad (22)$$

Substituting the above equation into (21), we obtain

$$\sigma^2(i) = h(i) + \beta^2 \sum_{j \in \mathcal{S}} p(j|i) \sigma^2(j), \quad i \in \mathcal{S}. \quad (23)$$

We observe that the above equation has a recursive form for $\sigma^2(i)$'s. Comparing (23) with (5), we can see that $\sigma^2(i)$ can be viewed as the discounted performance of the Markov chain with discount factor β^2 and new reward function $h(i)$, $i \in \mathcal{S}$. With (22), we further define this new reward function in a vector form as below.

$$\mathbf{h} := \mathbf{r}^2_{\odot} + 2\beta \mathbf{r} \odot (\mathbf{P}\mathbf{J}) + \beta^2 \mathbf{P}\mathbf{J}^2_{\odot} - \mathbf{J}^2_{\odot}, \quad (24)$$

where \odot indicates the Hadamard product (componentwisely), i.e.,

$$\mathbf{J}^2_{\odot} := \mathbf{J} \odot \mathbf{J} := (J^2(1), J^2(2), \dots, J^2(S))^T. \quad (25)$$

Based on the above analysis, we directly have the following theorem.

Theorem 2. *The variance of a β -discounted Markov chain is equivalent to the discounted performance of the same Markov chain with discount factor β^2 and reward function \mathbf{h} defined in (24). That is, we have*

$$\boldsymbol{\sigma}^2 = \mathbf{h} + \beta^2 \mathbf{P}\boldsymbol{\sigma}^2. \quad (26)$$

Note that the new reward function (22) can be further rewritten as below.

$$h(i) := \sum_{j \in \mathcal{S}} p(j|i) [r(i) + \beta J(j)]^2 - J^2(i). \quad (27)$$

We can further define a sample path version of the above definition as below.

$$h(X_t) \stackrel{\mathbb{E}}{=} [r(X_t) + \beta J(X_{t+1})]^2 - J^2(X_t), \quad (28)$$

where $\stackrel{\mathbb{E}}{=}$ means that the equality holds by expectation. With Theorem 2 and the above definitions, we can further rewrite the

variance of the Markov chain as below.

$$\sigma^2(i) = \mathbb{E}_i \left[\sum_{t=0}^{\infty} \beta^2 h(X_t) \right], \quad i \in \mathcal{S}. \quad (29)$$

Besides the above sample path version, the value of $\boldsymbol{\sigma}^2$ can also be computed in other ways that are described by the following theorem.

Theorem 3.

(a) With (24), we have

$$\boldsymbol{\sigma}^2 = (\mathbf{I} - \beta^2 \mathbf{P})^{-1} \mathbf{h}. \quad (30)$$

(b) With $\mathbf{f} := \mathbf{r}^2_{\odot} + 2\beta \mathbf{r} \odot (\mathbf{P}\mathbf{J})$, we have

$$\boldsymbol{\sigma}^2 = (\mathbf{I} - \beta^2 \mathbf{P})^{-1} \mathbf{f} - \mathbf{J}^2_{\odot}. \quad (31)$$

Proof. (a) Since the matrix $(\mathbf{I} - \beta^2 \mathbf{P})$ is invertible, (30) is directly derived from (26) in Theorem 2.

(b) Substituting $\mathbf{h} = \mathbf{f} + \beta^2 \mathbf{P}\mathbf{J}^2_{\odot} - \mathbf{J}^2_{\odot}$ into (30), we have

$$\begin{aligned}\boldsymbol{\sigma}^2 &= (\mathbf{I} - \beta^2 \mathbf{P})^{-1} (\mathbf{f} + \beta^2 \mathbf{P}\mathbf{J}^2_{\odot} - \mathbf{J}^2_{\odot}) \\ &= (\mathbf{I} - \beta^2 \mathbf{P})^{-1} \mathbf{f} + \sum_{n=0}^{\infty} (\beta^2 \mathbf{P})^n (\beta^2 \mathbf{P}\mathbf{J}^2_{\odot} - \mathbf{J}^2_{\odot}) \\ &= (\mathbf{I} - \beta^2 \mathbf{P})^{-1} \mathbf{f} - \mathbf{J}^2_{\odot}.\end{aligned}\quad (32)$$

The theorem is proved. \square

Comparing (31) with (6), we observe that $\boldsymbol{\sigma}^2 + \mathbf{J}^2_{\odot}$ equals the discounted performance of the same Markov chain with discount factor β^2 and reward function \mathbf{f} .

Below, we study the variance difference formula of the Markov chain under any two policies $d, d' \in \mathcal{D}$. With Theorem 2, we can view the variance as a special form of the discounted performance. Therefore, we directly apply the difference formula (10) and obtain

$$\boldsymbol{\sigma}^{2'} - \boldsymbol{\sigma}^2 = (\mathbf{I} - \beta^2 \mathbf{P}')^{-1} [\beta^2 (\mathbf{P}' - \mathbf{P}) \boldsymbol{\sigma}^2 + \mathbf{h}' - \mathbf{h}], \quad (33)$$

where \mathbf{P}' and \mathbf{h}' are the transition probability matrix and the equivalent reward function (24) of the Markov chain under the policy d' .

$$\begin{aligned}\boldsymbol{\sigma}^{2'} - \boldsymbol{\sigma}^2 &= (\mathbf{I} - \beta^2 \mathbf{P}')^{-1} [\beta^2 (\mathbf{P}' - \mathbf{P}) \boldsymbol{\sigma}^2 \\ &\quad + \mathbf{r}^2_{\odot} + 2\beta \mathbf{r}' \odot (\mathbf{P}'\boldsymbol{\lambda}) + \beta^2 \mathbf{P}'\boldsymbol{\lambda}^2_{\odot} \\ &\quad - \mathbf{r}^2_{\odot} - 2\beta \mathbf{r} \odot (\mathbf{P}\boldsymbol{\lambda}) - \beta^2 \mathbf{P}\boldsymbol{\lambda}^2_{\odot}] \\ &= (\mathbf{I} - \beta^2 \mathbf{P}')^{-1} [\beta^2 (\mathbf{P}' - \mathbf{P}) (\boldsymbol{\sigma}^2 + \boldsymbol{\lambda}^2_{\odot}) \\ &\quad + \mathbf{r}^2_{\odot} + 2\beta \mathbf{r}' \odot (\mathbf{P}'\boldsymbol{\lambda}) - \mathbf{r}^2_{\odot} - 2\beta \mathbf{r} \odot (\mathbf{P}\boldsymbol{\lambda})].\end{aligned}\quad (34)$$

Remark 4. Different from Remark 3, we cannot do the policy iteration based on (33) because the value of \mathbf{h}' is unknown. We need to pre-compute \mathbf{J}' before we compute the value of \mathbf{h}' based on (24). The computation of \mathbf{J}' under every possible policy $d' \in \mathcal{D}$ is a brute-force enumeration, which is unacceptable.

Fortunately, our original problem (4) is to find the optimal policy with the minimal variance from the policy set \mathcal{D}_{λ} . For any $d, d' \in \mathcal{D}_{\lambda}$, their discounted performances are equal to $\boldsymbol{\lambda}$. That is, $\mathbf{J} = \mathbf{J}' = \boldsymbol{\lambda}$, $\forall d, d' \in \mathcal{D}_{\lambda}$. Applying this fact and (24) to (33), we derive the variance difference formula (34) under any two policies $d, d' \in \mathcal{D}_{\lambda}$.

To obtain a concise form of (34), we further define the following column vectors \mathbf{g} and \mathbf{f} with elements

$$g(i) = \sigma^2(i) + \lambda^2(i), \quad i \in \mathcal{S}, \quad (35)$$

- For any given feasible constraint value λ , arbitrarily choose an initial policy $d^{(0)} \in \mathcal{D}_\lambda$ and set $k = 0$, where \mathcal{D}_λ is determined by $\mathcal{D}_\lambda = \mathcal{A}_\lambda(1) \times \mathcal{A}_\lambda(2) \times \cdots \times \mathcal{A}_\lambda(S)$ and $\mathcal{A}_\lambda(i)$ is determined by (12), $i \in \mathcal{S}$.
- For the current policy $d^{(k)}$, compute the value of $\mathbf{g}(d^{(k)})$ using (35) or (38).
- Generate a new policy $d^{(k+1)}$ using the following policy improvement

$$d^{(k+1)}(i) := \operatorname{argmin}_{a \in \mathcal{A}_\lambda(i)} \left\{ \beta^2 \sum_{j \in \mathcal{S}} p(j|i, a) \mathbf{g}(d^{(k)}, j) + r^2(i, a) + 2\beta r(i, a) \sum_{j \in \mathcal{S}} p(j|i, a) \lambda(j) \right\}, \quad i \in \mathcal{S}, \quad (40)$$
 where we choose $d^{(k+1)}(i) = d^{(k)}(i)$ if possible.
- If $d^{(k+1)} = d^{(k)}$, stop and output $d^{(k)}$ as the optimal policy; otherwise, let $k \leftarrow k + 1$ and repeat step 2.

Algorithm 1: A policy iteration algorithm to solve the constrained variance minimization problem (4).

$$f(i) = r^2(i) + 2\beta r(i) \sum_{j \in \mathcal{S}} p(j|i) \lambda(j), \quad i \in \mathcal{S}. \quad (36)$$

We see that \mathbf{g} can be viewed as the performance potential of the equivalent Markov chain with discount factor β^2 and cost function \mathbf{f} , and (36) can be viewed as a special case of \mathbf{f} defined in Theorem 3(b) with $\mathbf{J} = \lambda$.

Substituting (35) and (36) into (34), we obtain the variance difference formula of Markov chains under any two policies $d, d' \in \mathcal{D}_\lambda$ as follows:

$$\sigma^{2'} - \sigma^2 = (\mathbf{I} - \beta^2 \mathbf{P}')^{-1} [\beta^2 (\mathbf{P}' - \mathbf{P}) \mathbf{g} + \mathbf{f}' - \mathbf{f}]. \quad (37)$$

Putting the symbol d back to (35) and (36), we have

$$g(d, i) = \mathbb{E}_i^d \left[\sum_{t=0}^{\infty} \beta^t r(X_t, d(X_t)) \right]^2. \quad (38)$$

$$f(i, d(i)) = r^2(i, d(i)) + 2\beta r(i, d(i)) \sum_{j \in \mathcal{S}} p(j|i, d(i)) \lambda(j). \quad (39)$$

From the above equations, we can see that the value of \mathbf{f}' (with element $f(i, d'(i))$) under every possible $d' \in \mathcal{D}_\lambda$ is known since \mathbf{r}' (with element $r(i, d'(i))$) and \mathbf{P}' (with element $p(j|i, d'(i))$) are given values. The difficulty mentioned in Remark 4 is avoided. We can further develop a policy iteration algorithm to solve this constrained variance minimization problem (4). The details are described in Algorithm 1.

In the first step of Algorithm 1, we have to compute $\mathcal{A}_\lambda(i)$ by using (12). Since λ is given, we can enumerate every action $a \in \mathcal{A}(i)$ to see if the equation in (12) can hold. The total number of comparisons used in (12) is $\prod_{i \in \mathcal{S}} |\mathcal{A}(i)|$, which is affordable compared with the value iteration or policy iteration. The derivation of (40) can be intuitively understood as the policy improvement step in a standard policy iteration where we aim to equivalently minimize the discounted performance of the MDP with discount factor β^2 and cost function \mathbf{f} , according to the results of Theorem 2 and Theorem 3(b).

With the variance difference formula (34) or (37), we can derive the following theorem about the existence of the optimal policy and the convergence of Algorithm 1.

Theorem 4. The optimal policy d_λ^* for the problem (4) exists and Algorithm 1 can converge to d_λ^* .

The proof is omitted due to the capacity limited. A detailed version can be found in arXiv:1708.07095.

Based on Theorem 3, we can further derive the following optimality equation that the optimal σ^{2*} should satisfy.

$$\sigma^{2*}(i) = \min_{a \in \mathcal{A}_\lambda(i)} \left\{ h(i, a) + \beta^2 \sum_{j \in \mathcal{S}} p(j|i, a) \sigma^{2*}(j) \right\}, \quad (41)$$

for every state $i \in \mathcal{S}$. With this optimality equation, we can also develop a value iteration algorithm to solve the mean-constrained variance minimization problem. The value iteration algorithm will converge to the optimal value function σ^{2*} which is exactly the solution to the original problem (4). The algorithm is analogous to that in the classical MDP theory and we omit the details.

The main results of this paper have been obtained so far. We return to study a fundamental problem about the optimality of deterministic policy. All the above results are based on the problem formulation (4) in Section 2, where we limit our optimization in the deterministic policy space \mathcal{D}_λ . Below, we extend to study the reward variance of randomized policies.

Here we consider a special category of randomized policies that are generated from policy space \mathcal{D}_λ spanned by $\mathcal{A}_\lambda(i)$'s. That is, at each state i , we randomly choose actions a from $\mathcal{A}_\lambda(i)$ according to a probability distribution $\theta_{i,a}$, $i \in \mathcal{S}$ and $a \in \mathcal{A}_\lambda(i)$. Obviously, we have $0 \leq \theta_{i,a} \leq 1$ and $\sum_{a \in \mathcal{A}_\lambda(i)} \theta_{i,a} = 1$, $\forall i \in \mathcal{S}$. We denote θ as a vector composed of elements $\theta_{i,a}$'s. Different θ corresponds to different randomized policy and we denote the randomized policy as d^θ . Based on the sensitivity-based optimization theory, we obtain Theorem 5 as follows:

Theorem 5. For any randomized policy d^θ generated from \mathcal{D}_λ , we have $\mathbf{J}(d^\theta) = \lambda$ and $\sigma^2(d^\theta) \geq \sigma^2(d_\lambda^*)$. That is, we do not need to consider d^θ 's for problem (4).

Proof. The transition probability matrix and the reward function under the randomized policy d^θ are written as below, respectively.

$$\begin{aligned} p^\theta(j|i) &:= \sum_{a \in \mathcal{A}_\lambda(i)} \theta_{i,a} p(j|i, a), \\ r^\theta(i) &:= \sum_{a \in \mathcal{A}_\lambda(i)} \theta_{i,a} r(i, a). \end{aligned} \quad (42)$$

From definition (1), we can derive the following equation similar to (6).

$$\mathbf{J}(d^\theta) = (\mathbf{I} - \beta \mathbf{P}^\theta)^{-1} \mathbf{r}^\theta, \quad (43)$$

where \mathbf{P}^θ and \mathbf{r}^θ are the corresponding quantities under the randomized policy d^θ .

First, we compare the mean performance between the randomized policy d^θ and any deterministic policy $d \in \mathcal{D}_\lambda$. Similar to the difference formula (10), we can derive

$$\mathbf{J}(d^\theta) - \mathbf{J} = (\mathbf{I} - \beta \mathbf{P}^\theta)^{-1} [\beta (\mathbf{P}^\theta - \mathbf{P}) \mathbf{J} + \mathbf{r}^\theta - \mathbf{r}]. \quad (44)$$

In the above equation, $(\mathbf{I} - \beta \mathbf{P}^\theta)^{-1}$ is a positive matrix and we discuss the element in the square bracket. The square bracket in (44) is an S -dimensional column vector and its i th element can be written as below.

$$\beta \sum_{j \in \mathcal{S}} (p^\theta(j|i) - p(j|i)) \mathbf{J}(j) + r^\theta(i) - r(i).$$

$$\sum_{a \in \mathcal{A}_\lambda(i)} \theta_{i,a} \left\{ \sum_{j \in \mathcal{S}} [p(j|i, a) - p(j|i, d(i))] \sigma^2(j) + r^2(i, a) - r^2(i, d(i)) + 2\beta r(i, a) \sum_{j \in \mathcal{S}} p(j|i, a) \lambda(j) - 2\beta r(i, d(i)) \sum_{j \in \mathcal{S}} p(j|i, d(i)) \lambda(j) \right\}. \quad (48)$$

Box I.

Substituting (42) into the above equation, we obtain

$$\sum_{a \in \mathcal{A}_\lambda(i)} \theta_{i,a} \left[\sum_{j \in \mathcal{S}} \beta (p(j|i, a) - p(j|i, d(i))) J(j) + r(i, a) - r(i, d(i)) \right].$$

Since $d(i) \in \mathcal{A}_\lambda(i)$ and $J(i) = \lambda(i)$ for all $i \in \mathcal{S}$, with (12) we can see that the above equation equals 0. Therefore, we have $\mathbf{J}(d^\theta) - \mathbf{J} = 0$ and $\mathbf{J}(d^\theta) = \lambda$.

Then, we compare the reward variance between policy d^θ and $d \in \mathcal{D}_\lambda$. Similar to (21), we can also obtain the variance $\sigma^2(d^\theta, i)$ under policy d^θ as below.

$$\begin{aligned} \sigma^2(d^\theta, i) = & \sum_{a \in \mathcal{A}_\lambda(i)} \theta_{i,a} \left\{ r^2(i, a) + 2\beta r(i, a) \sum_{j \in \mathcal{S}} p(j|i, a) J(d^\theta, j) \right. \\ & \left. + \beta^2 \sum_{j \in \mathcal{S}} p(j|i, a) [\sigma^2(d^\theta, j) + J^2(d^\theta, j)] - J^2(d^\theta, i) \right\}. \end{aligned}$$

We denote \mathbf{h}^θ as a new reward function and its element $h^\theta(i), i \in \mathcal{S}$, is defined as below.

$$\begin{aligned} h^\theta(i) := & \sum_{a \in \mathcal{A}_\lambda(i)} \theta_{i,a} \left\{ r^2(i, a) + 2\beta r(i, a) \sum_{j \in \mathcal{S}} p(j|i, a) J(d^\theta, j) \right\} \\ & + \beta^2 \sum_{j \in \mathcal{S}} p^\theta(j|i) J^2(d^\theta, j) - J^2(d^\theta, i). \end{aligned} \quad (45)$$

Therefore, the result in Theorem 2 also holds for this randomized policy d^θ and we have

$$\sigma^2(d^\theta) = (\mathbf{I} - \beta^2 \mathbf{P}^\theta)^{-1} \mathbf{h}^\theta. \quad (46)$$

Similarly, we can derive the variance difference formula between policy d^θ and d as below.

$$\sigma^2(d^\theta) - \sigma^2 = (\mathbf{I} - \beta^2 \mathbf{P}^\theta)^{-1} [\beta^2 (\mathbf{P}^\theta - \mathbf{P}) \sigma^2 + \mathbf{h}^\theta - \mathbf{h}]. \quad (47)$$

Since $(\mathbf{I} - \beta^2 \mathbf{P}^\theta)^{-1}$ is a positive matrix, we study the value of the element of the square bracket in the above equation. Substituting (42) and (45) into (47), we can derive (48) in Box I to represent the i th element of the square bracket in (47), where we use the fact $\mathbf{J}(d^\theta) = \mathbf{J} = \lambda$. Since the value of the large bracket in (48) has no relation to d^θ , we can view it as a given value. With (47), it is easy to verify that the optimal θ^* with the minimal variance must satisfy a necessary condition: $\theta_{i,a}^* \in \{0, 1\}, \forall a \in \mathcal{A}_\lambda(i), i \in \mathcal{S}$. That is, the optimal policy is a deterministic one and we have $\sigma^2(d^\theta) \geq \sigma^2(d_\lambda^*)$. The theorem is proved. \square

4. Numerical example

Consider a discrete time Markov chain with state space $\mathcal{S} = \{1, 2\}$. The action space is $A(1) = \{1, 2, 3\}$ and $A(2) = \{1, 2, 3, 4\}$. The transition probabilities are $p(2|1, a) = a/4, p(1|1, a) = 1 - a/4$ for $a \in A(1)$, and $p(1|2, a) = a/4, p(2|2, a) = 1 - a/4$ for $a \in A(2)$. The rewards in state-action pairs are $r(1, 1) = 1, r(1, 2) = \frac{3}{4}, r(1, 3) = \frac{19}{32}; r(2, 1) = \frac{5}{2}, r(2, 2) = 2, r(2, 3) = 3, r(2, 4) = \frac{13}{4}$. The discount factor is $\beta = 0.5$. Obviously, the number of total policies in \mathcal{D} is 12 and all the possible policies are denoted as follows:

$$\begin{aligned} d_1 = (1, 1), d_2 = (1, 2), d_3 = (1, 3), d_4 = (1, 4), \\ d_5 = (2, 1), d_6 = (2, 2), d_7 = (2, 3), d_8 = (2, 4), \\ d_9 = (3, 1), d_{10} = (3, 2), d_{11} = (3, 3), d_{12} = (3, 4). \end{aligned}$$

Using (6) and (30), we can compute the value of \mathbf{J} and σ^2 of the Markov chain under every possible policy. The computation results are listed in Table 1.

From Table 1, we see that there exist some policies under which the discounted Markov chain has the same mean, even the same variance. We let $\lambda = (2.5, 4.5)^T$ and verify the main results derived in Section 3. From Table 1, we see that $\mathcal{D}_\lambda = \{d_1, d_3, d_4, d_5, d_7, d_8\}$ in which the discounted performance equals λ . It is easy to verify that $\mathcal{A}_\lambda(1) = \{1, 2\}$ and $\mathcal{A}_\lambda(2) = \{1, 3, 4\}$. Therefore, we have $\mathcal{D}_\lambda = \mathcal{A}_\lambda(1) \times \mathcal{A}_\lambda(2)$ and the result of Theorem 1 is verified in this example. If we let $\lambda = (2.125, 3.375)^T$, we can also verify that $\mathcal{D}_\lambda = \mathcal{A}_\lambda(1) \times \mathcal{A}_\lambda(2)$, where $\mathcal{D}_\lambda = \{d_6, d_{10}\}, \mathcal{A}_\lambda(1) = \{2, 3\}$, and $\mathcal{A}_\lambda(2) = \{2\}$.

Below, we verify the policy iteration algorithm to find the optimal policy for the problem (4) when $\lambda = (2.5, 4.5)^T$. We arbitrarily choose an initial policy from the policy set \mathcal{D}_λ , say we choose $d^{(0)} = d_5$. We compute $\mathbf{g}(d^{(0)})$ using (35) and obtain $\mathbf{g}(d^{(0)}) = (6.5722, 20.5056)^T$. Then we use the policy improvement (40) to generate a new policy as follows: for state $i = 1$, we have $d^{(1)}(1) = \operatorname{argmin}_{a \in \{1, 2\}} \{6.5139, 6.5722\} = 1$; for state $i = 2$,

we have $d^{(1)}(2) = \operatorname{argmin}_{a \in \{1, 3, 4\}} \{20.5056, 20.5139, 20.3306\} = 4$.

Therefore, $d^{(1)} = (1, 4) = d_4$. For $d^{(1)}$, we repeat the above process and compute $\mathbf{g}(d^{(1)}) = (6.4853, 20.3088)^T$. Then we again use (40) to generate the next policy as follows: for state $i = 1$, we have $d^{(2)}(1) = \operatorname{argmin}_{a \in \{1, 2\}} \{6.4853, 6.5368\} = 1$; for state $i =$

2, we have $d^{(2)}(2) = \operatorname{argmin}_{a \in \{1, 3, 4\}} \{20.4632, 20.4853, 20.3088\} = 4$.

Therefore, $d^{(2)} = (1, 4) = d_4$. We have $d^{(2)} = d^{(1)}$ and the stopping criterion is satisfied. Algorithm 1 stops and outputs d_4 as the optimal policy with the minimal variance among the policy set \mathcal{D}_λ . From Table 1, we find that d_4 is truly the optimal policy with the minimal variance $\sigma^2 = (0.2353, 0.0588)^T$ among the policy set $\mathcal{D}_\lambda = \{d_1, d_3, d_4, d_5, d_7, d_8\}$. Therefore, the convergence of Algorithm 1 is verified and Algorithm 1 converges to the optimal policy only through 1 iteration.

From Table 1, we observe another interesting fact that the policies d_1 and d_3 have the same mean and variance. To compare these two policies, we can further consider other high order performance metrics, such as the n th-order bias optimality (Cao & Zhang, 2008). Moreover, we can also investigate the Pareto optimality of policies from Table 1. For the objective of maximizing the mean and minimizing the variance, we can find: d_2 dominates d_{10} and d_{10} dominates d_6 , i.e., $\mathbf{J}(d_2) > \mathbf{J}(d_{10}) > \mathbf{J}(d_6)$ and $\sigma^2(d_2) < \sigma^2(d_{10}) < \sigma^2(d_6)$; d_4 dominates all the other policies in $\mathcal{D}_\lambda = \{d_1, d_3, d_4, d_5, d_7, d_8\}$; d_{12} dominates d_4, d_9 , and d_{11} . Therefore, the mean and variance of $\{d_2, d_{12}\}$ comprise the efficient frontier of this example.

5. Conclusion

In this paper, we study the variance minimization problem of a discrete time discounted MDP where the discounted performance is equal to a given constant. By transforming this constrained variance minimization problem to an unconstrained MDP with discount factor β^2 and new reward function \mathbf{h} , we develop a policy iteration algorithm to efficiently solve such category of

Table 1

The mean and variance of the discounted Markov chain under every possible policy.

	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8	d_9	d_{10}	d_{11}	d_{12}
J	$\begin{pmatrix} 2.5 \\ 4.5 \end{pmatrix}$	$\begin{pmatrix} 2.2857 \\ 3.4286 \end{pmatrix}$	$\begin{pmatrix} 2.5 \\ 4.5 \end{pmatrix}$	$\begin{pmatrix} 2.5 \\ 4.5 \end{pmatrix}$	$\begin{pmatrix} 2.5 \\ 4.5 \end{pmatrix}$	$\begin{pmatrix} 2.125 \\ 3.375 \end{pmatrix}$	$\begin{pmatrix} 2.5 \\ 4.5 \end{pmatrix}$	$\begin{pmatrix} 2.5 \\ 4.5 \end{pmatrix}$	$\begin{pmatrix} 2.6172 \\ 4.5234 \end{pmatrix}$	$\begin{pmatrix} 2.125 \\ 3.375 \end{pmatrix}$	$\begin{pmatrix} 2.6312 \\ 4.5562 \end{pmatrix}$	$\begin{pmatrix} 2.6364 \\ 4.5682 \end{pmatrix}$
σ^2	$\begin{pmatrix} 0.25 \\ 0.25 \end{pmatrix}$	$\begin{pmatrix} 0.0834 \\ 0.1052 \end{pmatrix}$	$\begin{pmatrix} 0.25 \\ 0.25 \end{pmatrix}$	$\begin{pmatrix} 0.2353 \\ 0.0588 \end{pmatrix}$	$\begin{pmatrix} 0.3222 \\ 0.2556 \end{pmatrix}$	$\begin{pmatrix} 0.1302 \\ 0.1302 \end{pmatrix}$	$\begin{pmatrix} 0.3235 \\ 0.2647 \end{pmatrix}$	$\begin{pmatrix} 0.2963 \\ 0.0741 \end{pmatrix}$	$\begin{pmatrix} 0.2271 \\ 0.2271 \end{pmatrix}$	$\begin{pmatrix} 0.1034 \\ 0.1264 \end{pmatrix}$	$\begin{pmatrix} 0.2316 \\ 0.2316 \end{pmatrix}$	$\begin{pmatrix} 0.1964 \\ 0.0491 \end{pmatrix}$

optimization problems. The success of this approach depends on the decomposable structure of the policy set \mathcal{D}_λ , as stated in Theorem 1. Such property may not hold if we consider the long-run average criterion instead of the discounted criterion. Therefore, the variance minimization problem of MDPs with a constraint on the long-run average performance is a future research topic. Another fundamental topic is to study the optimality of stationary policy for this constrained variance minimization problem, which may be generally different from the standard constrained MDPs in the literature. Moreover, how to extend our results to more general cases, such as continuous time continuous state Markov processes or finite horizon Markov chains, is another interesting topic deserving further investigation.

Acknowledgments

The author would like to thank the three anonymous referees and the associate editor for their valuable comments which help improve the paper. This work was supported in part by the National Key Research and Development Program of China (2016YFB0901900), the National Natural Science Foundation of China (61573206, 61203039, U1301254), the National 111 International Collaboration Project (B06002), and the special fund of SuZhou-Tsinghua Innovation Leading Action (2016SZ0202).

References

- Cao, X. R. (2007). *Stochastic learning and optimization—a sensitivity-based approach*. New York: Springer.
- Cao, X. R., & Zhang, J. (2008). The n th-order bias optimality for multi-chain Markov decision processes. *IEEE Transactions on Automatic Control*, 53, 496–508.
- Chung, K. J. (1994). Mean–variance tradeoffs in an undiscounted MDP: the unichain case. *Operations Research*, 42, 184–188.
- Guo, X., Huang, X., & Zhang, Y. (2015). On the first passage g -mean–variance optimality for discounted continuous-time Markov decision processes. *SIAM Journal on Control and Optimization*, 53, 1406–1424.
- Guo, X., & Song, X. Y. (2009). Mean–variance criteria for finite continuous-time Markov decision processes. *IEEE Transactions on Automatic Control*, 54, 2151–2157.
- Hernandez-Lerma, O., Vega-Amaya, O., & Carrasco, G. (1999). Sample-path optimality and variance-minimization of average cost Markov control processes. *SIAM Journal on Control and Optimization*, 38, 79–93.

- Huang, Y., & Chen, X. A sensitivity-based construction approach to sample-path variance minimization of Markov decision processes, In *Proceedings of the 2012 Australian control conference*, Sydney, Australia, (pp. 215–220).
- Huo, H., Zou, X., & Guo, X. (2017). The risk probability criterion for discounted continuous-time Markov decision processes. *Discrete Event Dynamic Systems: Theory and Applications*, 27, 675–699.
- Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7, 77–91.
- Prashantha, L. A., & Ghavamzadeh, M. (2013). Actor-critic algorithms for risk-sensitive MDPs, In *Advances in neural information processing systems (NIPS'13)*, (pp. 252–260).
- Sobel, M. J. (1982). The variance of discounted Markov decision. *Journal of Applied Probability*, 19, 794–802.
- Sobel, M. J. (1994). Mean–variance tradeoffs in an undiscounted MDP. *Operations Research*, 42, 175–183.
- Tamar, A., Castro, D. D., & Mannor, S. (2012). Policy gradients with variance related risk criteria, In *Proceedings of the 29th International conference on machine learning (ICML)*, Edinburgh, Scotland, (pp. 935–942).
- Xia, L. (2016). Optimization of Markov decision processes under the variance criterion. *Automatica*, 73, 269–278.
- Xia, L. (2017). Variance minimization of parameterized Markov decision processes. *Discrete Event Dynamic Systems: Theory and Applications*. <http://dx.doi.org/10.1007/s10626-017-0258-5>.
- Zhou, X. Y., & Yin, G. (2003). Markowitz's mean–variance portfolio selection with regime switching: A continuous-time model. *SIAM Journal on Control and Optimization*, 42, 1466–1482.



Li Xia is an associate professor at the Center for Intelligent and Networked Systems (CFINS), Department of Automation, Tsinghua University, Beijing China. He received the B.S. degree and the Ph.D. degree in Control Theory in 2002 and 2007 respectively, both from Tsinghua University. After graduation, he worked at IBM Research China as a research staff member (2007–2009) and at the King Abdullah University of Science and Technology (KAUST) Saudi Arabia as a postdoctoral research fellow (2009–2011). Then he returned to Tsinghua University in 2011. He was a visiting scholar Stanford University, the Hong Kong University of Science Technology, etc. He serves/served as an associate editor and program committee member of a number of international journals and conferences. His research interests include the methodology research in stochastic learning and optimization, queueing theory, Markov decision processes, reinforcement learning, and the application research in building energy, energy Internet, industrial Internet, Internet of things, etc. He is a senior member of IEEE.