# ROBUST STOCHASTIC APPROXIMATION APPROACH TO STOCHASTIC PROGRAMMING* 

A. NEMIROVSKI ${ }^{\dagger}$, A. JUDITSKY ${ }^{\ddagger}$, G. LAN $^{\dagger}$, AND A. SHAPIRO ${ }^{\dagger}$


#### Abstract

In this paper we consider optimization problems where the objective function is given in a form of the expectation. A basic difficulty of solving such stochastic optimization problems is that the involved multidimensional integrals (expectations) cannot be computed with high accuracy. The aim of this paper is to compare two computational approaches based on Monte Carlo sampling techniques, namely, the stochastic approximation (SA) and the sample average approximation (SAA) methods. Both approaches, the SA and SAA methods, have a long history. Current opinion is that the SAA method can efficiently use a specific (say, linear) structure of the considered problem, while the SA approach is a crude subgradient method, which often performs poorly in practice. We intend to demonstrate that a properly modified SA approach can be competitive and even significantly outperform the SAA method for a certain class of convex stochastic problems. We extend the analysis to the case of convex-concave stochastic saddle point problems and present (in our opinion highly encouraging) results of numerical experiments.


Key words. stochastic approximation, sample average approximation method, stochastic programming, Monte Carlo sampling, complexity, saddle point, minimax problems, mirror descent algorithm

AMS subject classifications. 90C15, 90C25
DOI. 10.1137/070704277

1. Introduction. In this paper we first consider the following stochastic optimization problem:

$$
\begin{equation*}
\min _{x \in X}\{f(x)=\mathbb{E}[F(x, \xi)]\}, \tag{1.1}
\end{equation*}
$$

and then we deal with an extension of the analysis to stochastic saddle point problems. Here $X \subset \mathbb{R}^{n}$ is a nonempty bounded closed convex set, $\xi$ is a random vector whose probability distribution $P$ is supported on set $\Xi \subset \mathbb{R}^{d}$ and $F: X \times \Xi \rightarrow \mathbb{R}$. We assume that the expectation

$$
\begin{equation*}
\mathbb{E}[F(x, \xi)]=\int_{\Xi} F(x, \xi) d P(\xi) \tag{1.2}
\end{equation*}
$$

is well defined and finite valued for every $x \in X$. Moreover, we assume that the expected value function $f(\cdot)$ is continuous and convex on $X$. Of course, if for every $\xi \in \Xi$ the function $F(\cdot, \xi)$ is convex on $X$, then it follows that $f(\cdot)$ is convex. With these assumptions, (1.1) becomes a convex programming problem.

A basic difficulty of solving stochastic optimization problem (1.1) is that the multidimensional integral (expectation) (1.2) cannot be computed with a high accuracy for dimension $d$, say, greater than five. The aim of this paper is to compare two

[^0]computational approaches based on Monte Carlo sampling techniques, namely, the stochastic approximation (SA) and the sample average approximation (SAA) methods. To this end we make the following assumptions.
(A1) It is possible to generate an independent identically distributed (iid) sample $\xi_{1}, \xi_{2}, \ldots$, of realizations of random vector $\xi$.
(A2) There is a mechanism (an oracle), which, for a given input point $(x, \xi) \in$ $X \times \Xi$ returns stochastic subgradient-a vector $\mathrm{G}(x, \xi)$ such that $\mathrm{g}(x):=\mathbb{E}[\mathrm{G}(x, \xi)]$ is well defined and is a subgradient of $f(\cdot)$ at $x$, i.e., $\mathrm{g}(x) \in \partial f(x)$.

Recall that if $F(\cdot, \xi), \xi \in \Xi$, is convex and $f(\cdot)$ is finite valued in a neighborhood of a point $x$, then (cf. Strassen [28])

$$
\begin{equation*}
\partial f(x)=\mathbb{E}\left[\partial_{x} F(x, \xi)\right] . \tag{1.3}
\end{equation*}
$$

In that case we can employ a measurable selection $\mathrm{G}(x, \xi) \in \partial_{x} F(x, \xi)$ as a stochastic subgradient. At this stage, however, this is not important, we shall see later other relevant ways for constructing stochastic subgradients.

Both approaches, the SA and SAA methods, have a long history. The SA method is going back to the pioneering paper by Robbins and Monro [21]. Since then SA algorithms became widely used in stochastic optimization (see, e.g., $[3,6,7,20,22]$ and references therein) and, due to especially low demand for computer memory, in signal processing. In the classical analysis of the SA algorithm (it apparently goes back to the works [5] and [23]) it is assumed that $f(\cdot)$ is twice continuously differentiable and strongly convex and in the case when the minimizer of $f$ belongs to the interior of $X$, exhibits asymptotically optimal rate ${ }^{1}$ of convergence $\mathbb{E}\left[f\left(x_{t}\right)-f_{*}\right]=O\left(t^{-1}\right)$ (here $x_{t}$ is $t$ th iterate and $f_{*}$ is the minimal value of $f(x)$ over $\left.x \in X\right)$. This algorithm, however, is very sensitive to a choice of the respective stepsizes. Since "asymptotically optimal" stepsize policy can be very bad in the beginning, the algorithm often performs poorly in practice (e.g., [27, section 4.5.3.]).

An important improvement of the SA method was developed by Polyak [18] and Polyak and Juditsky [19], where longer stepsizes were suggested with consequent averaging of the obtained iterates. Under the outlined "classical" assumptions, the resulting algorithm exhibits the same optimal $O\left(t^{-1}\right)$ asymptotical convergence rate, while using an easy to implement and "robust" stepsize policy. It should be mentioned that the main ingredients of Polyak's scheme - long steps and averaging - were, in a different form, proposed already in Nemirovski and Yudin [15] for the case of problems (1.1) with general-type Lipschitz continuous convex objectives and for convex-concave saddle point problems. The algorithms from [15] exhibit, in a nonasymptotical fashion, the $O\left(t^{-1 / 2}\right)$ rate of convergence. It is possible to show that in the general convex case (without assuming smoothness and strong convexity of the objective function), this rate of $O\left(t^{-1 / 2}\right)$ is unimprovable. For a summary of early results in this direction, see Nemirovski and Yudin [16].

The SAA approach was used by many authors in various contexts under different names. Its basic idea is rather simple: generate a (random) sample $\xi_{1}, \ldots, \xi_{N}$, of size $N$, and approximate the "true" problem (1.1) by the sample average problem

$$
\begin{equation*}
\min _{x \in X}\left\{\hat{f}_{N}(x)=N^{-1} \sum_{j=1}^{N} F\left(x, \xi_{j}\right)\right\} . \tag{1.4}
\end{equation*}
$$

[^1]Note that the SAA method is not an algorithm; the obtained SAA problem (1.4) still has to be solved by an appropriate numerical procedure. Recent theoretical studies (cf. $[11,25,26]$ ) and numerical experiments (see, e.g., $[12,13,29]$ ) show that the SAA method coupled with a good (deterministic) algorithm could be reasonably efficient for solving certain classes of two-stage stochastic programming problems. On the other hand, classical SA-type numerical procedures typically performed poorly for such problems.

We intend to demonstrate in this paper that a properly modified SA approach can be competitive and even significantly outperform the SAA method for a certain class of stochastic problems. The mirror descent SA method we propose here is a direct descendent of the stochastic mirror descent method of Nemirovski and Yudin [16]. However, the method developed in this paper is more flexible than its "ancestor": the iteration of the method is exactly the prox-step for a chosen prox-function, and the choice of prox-type function is not limited to the norm-type distance-generating functions. Close techniques, based on subgradient averaging, have been proposed in Nesterov [17] and used in [10] to solve the stochastic optimization problem (1.1). Moreover, the results on large deviations of solutions and applications of the mirror descent SA to saddle point problems, to the best of our knowledge, are new.

The rest of this paper is organized as follows. In section 2 we focus on theory of the SA method applied to (1.1). We start with outlining the relevant-to-our-goals part of the classical " $O\left(t^{-1}\right)$ " SA theory (section 2.1), along with its " $O\left(t^{-1 / 2}\right)$ " modifications (section 2.2). Well-known and simple results presented in these sections pave the road to our main developments carried out in section 2.3 . In section 3 we extend the constructions and results of section 2.3 to the case of the convexconcave stochastic saddle point problem. In concluding section 4 we present results (in our opinion, highly encouraging) of numerical experiments with the SA algorithm (sections 2.3 and 3) applied to large-scale stochastic convex minimization and saddle point problems. Section 5 gives a short conclusion for the presented results. Finally, some technical proofs are given in the appendix.

Throughout the paper, we use the following notation. By $\|x\|_{p}$, we denote the $\ell_{p}$ norm of vector $x \in \mathbb{R}^{n}$, in particular, $\|x\|_{2}=\sqrt{x^{T} x}$ denotes the Euclidean norm, and $\|x\|_{\infty}=\max \left\{\left|x_{1}\right|, \ldots,\left|x_{n}\right|\right\}$. By $\Pi_{X}$, we denote the metric projection operator onto the set $X$, that is, $\Pi_{X}(x)=\arg \min _{x^{\prime} \in X}\left\|x-x^{\prime}\right\|_{2}$. Note that $\Pi_{X}$ is a nonexpanding operator, i.e.,

$$
\begin{equation*}
\left\|\Pi_{X}\left(x^{\prime}\right)-\Pi_{X}(x)\right\|_{2} \leq\left\|x^{\prime}-x\right\|_{2} \forall x^{\prime}, x \in \mathbb{R}^{n} \tag{1.5}
\end{equation*}
$$

By $O(1)$, we denote positive absolute constants. The notation $\lfloor a\rfloor$ stands for the largest integer less than or equal to $a \in \mathbb{R}$ and $\lceil a\rceil$ for the smallest integer greater than or equal to $a \in \mathbb{R}$. By $\xi_{[t]}=\left(\xi_{1}, \ldots, \xi_{t}\right)$, we denote the history of the process $\xi_{1}, \ldots$, up to time $t$. Unless stated otherwise, all relations between random variables are supposed to hold almost surely.
2. Stochastic approximation, basic theory. In this section we discuss theory and implementations of the SA approach to the minimization problem (1.1).
2.1. Classical SA algorithm. The classical SA algorithm solves (1.1) by mimicking the simplest subgradient descent method. That is, for chosen $x_{1} \in X$ and a sequence $\gamma_{j}>0, j=1, \ldots$, of stepsizes, it generates the iterates by the formula

$$
\begin{equation*}
x_{j+1}=\Pi_{X}\left(x_{j}-\gamma_{j} \mathrm{G}\left(x_{j}, \xi_{j}\right)\right) \tag{2.1}
\end{equation*}
$$

Of course, the crucial question of that approach is how to choose the stepsizes $\gamma_{j}$. Let $x_{*}$ be an optimal solution of (1.1). Note that since the set $X$ is compact and $f(x)$ is continuous, (1.1) has an optimal solution. Note also that the iterate $x_{j}=x_{j}\left(\xi_{[j-1]}\right)$ is a function of the history $\xi_{[j-1]}=\left(\xi_{1}, \ldots, \xi_{j-1}\right)$ of the generated random process and hence is random.

Denote

$$
\begin{equation*}
A_{j}=\frac{1}{2}\left\|x_{j}-x_{*}\right\|_{2}^{2} \quad \text { and } a_{j}=\mathbb{E}\left[A_{j}\right]=\frac{1}{2} \mathbb{E}\left[\left\|x_{j}-x_{*}\right\|_{2}^{2}\right] \tag{2.2}
\end{equation*}
$$

By using (1.5) and since $x_{*} \in X$ and hence $\Pi_{X}\left(x_{*}\right)=x_{*}$, we can write

$$
\begin{align*}
A_{j+1} & =\frac{1}{2}\left\|\Pi_{X}\left(x_{j}-\gamma_{j} \mathrm{G}\left(x_{j}, \xi_{j}\right)\right)-x_{*}\right\|_{2}^{2} \\
& =\frac{1}{2}\left\|\Pi_{X}\left(x_{j}-\gamma_{j} \mathrm{G}\left(x_{j}, \xi_{j}\right)\right)-\Pi_{X}\left(x_{*}\right)\right\|_{2}^{2}  \tag{2.3}\\
& \leq \frac{1}{2}\left\|x_{j}-\gamma_{j} \mathrm{G}\left(x_{j}, \xi_{j}\right)-x_{*}\right\|_{2}^{2} \\
& =A_{j}+\frac{1}{2} \gamma_{j}^{2}\left\|\mathrm{G}\left(x_{j}, \xi_{j}\right)\right\|_{2}^{2}-\gamma_{j}\left(x_{j}-x_{*}\right)^{T} \mathrm{G}\left(x_{j}, \xi_{j}\right)
\end{align*}
$$

Since $x_{j}=x_{j}\left(\xi_{[j-1]}\right)$ is independent of $\xi_{j}$, we have

$$
\begin{align*}
\mathbb{E}\left[\left(x_{j}-x_{*}\right)^{T} \mathrm{G}\left(x_{j}, \xi_{j}\right)\right] & =\mathbb{E}\left\{\mathbb{E}\left[\left(x_{j}-x_{*}\right)^{T} \mathrm{G}\left(x_{j}, \xi_{j}\right) \mid \xi_{[j-1]}\right]\right\} \\
& =\mathbb{E}\left\{\left(x_{j}-x_{*}\right)^{T} \mathbb{E}\left[\mathrm{G}\left(x_{j}, \xi_{j}\right) \mid \xi_{[j-1]}\right]\right\}  \tag{2.4}\\
& =\mathbb{E}\left[\left(x_{j}-x_{*}\right)^{T} \mathrm{~g}\left(x_{j}\right)\right] .
\end{align*}
$$

Assume now that there is a positive number $M$ such that

$$
\begin{equation*}
\mathbb{E}\left[\|\mathrm{G}(x, \xi)\|_{2}^{2}\right] \leq M^{2} \quad \forall x \in X \tag{2.5}
\end{equation*}
$$

Then, by taking expectation of both sides of (2.3) and using (2.4), we obtain

$$
\begin{equation*}
a_{j+1} \leq a_{j}-\gamma_{j} \mathbb{E}\left[\left(x_{j}-x_{*}\right)^{T} \mathrm{~g}\left(x_{j}\right)\right]+\frac{1}{2} \gamma_{j}^{2} M^{2} \tag{2.6}
\end{equation*}
$$

Suppose further that the expectation function $f(x)$ is differentiable and strongly convex on $X$, i.e., there is constant $c>0$ such that

$$
f\left(x^{\prime}\right) \geq f(x)+\left(x^{\prime}-x\right)^{T} \nabla f(x)+\frac{1}{2} c\left\|x^{\prime}-x\right\|_{2}^{2}, \quad \forall x^{\prime}, x \in X
$$

or equivalently that

$$
\begin{equation*}
\left(x^{\prime}-x\right)^{T}\left(\nabla f\left(x^{\prime}\right)-\nabla f(x)\right) \geq c\left\|x^{\prime}-x\right\|_{2}^{2} \forall x^{\prime}, x \in X \tag{2.7}
\end{equation*}
$$

Note that strong convexity of $f(x)$ implies that the minimizer $x_{*}$ is unique. By optimality of $x_{*}$, we have that

$$
\left(x-x_{*}\right)^{T} \nabla f\left(x_{*}\right) \geq 0 \quad \forall x \in X
$$

which together with (2.7) implies that $\left(x-x_{*}\right)^{T} \nabla f(x) \geq c\left\|x-x_{*}\right\|_{2}^{2}$. In turn, it follows that $\left(x-x_{*}\right)^{T} g \geq c\left\|x-x_{*}\right\|_{2}^{2}$ for all $x \in X$ and $g \in \partial f(x)$, and hence

$$
\mathbb{E}\left[\left(x_{j}-x_{*}\right)^{T} \mathrm{~g}\left(x_{j}\right)\right] \geq c \mathbb{E}\left[\left\|x_{j}-x_{*}\right\|_{2}^{2}\right]=2 c a_{j}
$$

Therefore, it follows from (2.6) that

$$
\begin{equation*}
a_{j+1} \leq\left(1-2 c \gamma_{j}\right) a_{j}+\frac{1}{2} \gamma_{j}^{2} M^{2} \tag{2.8}
\end{equation*}
$$

Let us take stepsizes $\gamma_{j}=\theta / j$ for some constant $\theta>1 /(2 c)$. Then, by (2.8), we have

$$
a_{j+1} \leq(1-2 c \theta / j) a_{j}+\frac{1}{2} \theta^{2} M^{2} / j^{2}
$$

It follows by induction that

$$
\begin{equation*}
\mathbb{E}\left[\left\|x_{j}-x_{*}\right\|_{2}^{2}\right]=2 a_{j} \leq Q(\theta) / j \tag{2.9}
\end{equation*}
$$

where

$$
\begin{equation*}
Q(\theta)=\max \left\{\theta^{2} M^{2}(2 c \theta-1)^{-1},\left\|x_{1}-x_{*}\right\|_{2}^{2}\right\} \tag{2.10}
\end{equation*}
$$

Suppose further that $x_{*}$ is an interior point of $X$ and $\nabla f(x)$ is Lipschitz continuous, i.e., there is constant $L>0$ such that

$$
\begin{equation*}
\left\|\nabla f\left(x^{\prime}\right)-\nabla f(x)\right\|_{2} \leq L\left\|x^{\prime}-x\right\|_{2} \forall x^{\prime}, x \in X \tag{2.11}
\end{equation*}
$$

Then

$$
\begin{equation*}
f(x) \leq f\left(x_{*}\right)+\frac{1}{2} L\left\|x-x_{*}\right\|_{2}^{2}, \quad \forall x \in X \tag{2.12}
\end{equation*}
$$

and hence

$$
\begin{equation*}
\mathbb{E}\left[f\left(x_{j}\right)-f\left(x_{*}\right)\right] \leq \frac{1}{2} L \mathbb{E}\left[\left\|x_{j}-x_{*}\right\|_{2}^{2}\right] \leq \frac{1}{2} L Q(\theta) / j \tag{2.13}
\end{equation*}
$$

where $Q(\theta)$ is defined in (2.10).
Under the specified assumptions, it follows from (2.9) and (2.13), respectively, that after $t$ iterations, the expected error of the current solution in terms of the distance to $x_{*}$ is of order $O\left(t^{-1 / 2}\right)$, and the expected error in terms of the objective value is of order $O\left(t^{-1}\right)$, provided that $\theta>1 /(2 c)$. The simple example of $X=\left\{x:\|x\|_{2} \leq 1\right\}$, $f(x)=\frac{1}{2} c x^{T} x$, and $\mathrm{G}(x, \xi)=\nabla f(x)+\xi$, with $\xi$ having standard normal distribution $\mathcal{N}\left(0, I_{n}\right)$, demonstrates that the outlined upper bounds on the expected errors are tight within factors independent of $t$.

We have arrived at the $O\left(t^{-1}\right)$ rate of convergence in terms of the expected value of the objective mentioned in the Introduction. Note, however, that the result is highly sensitive to a priori information on $c$. What would happen if the parameter $c$ of strong convexity is overestimated? As a simple example, consider $f(x)=x^{2} / 10$, $X=[-1,1] \subset \mathbb{R}$, and assume that there is no noise, i.e., $\mathrm{G}(x, \xi) \equiv \nabla f(x)$. Suppose, further that we take $\theta=1$ (i.e., $\gamma_{j}=1 / j$ ), which will be the optimal choice for $c=1$, while actually here $c=0.2$. Then the iteration process becomes

$$
x_{j+1}=x_{j}-f^{\prime}\left(x_{j}\right) / j=\left(1-\frac{1}{5 j}\right) x_{j}
$$

and hence starting with $x_{1}=1$,

$$
\begin{aligned}
x_{j}= & \prod_{s=1}^{j-1}\left(1-\frac{1}{5 s}\right)=\exp \left\{-\sum_{s=1}^{j-1} \ln \left(1+\frac{1}{5 s-1}\right)\right\}>\exp \left\{-\sum_{s=1}^{j-1} \frac{1}{5 s-1}\right\} \\
& >\exp \left\{-\left(0.25+\int_{1}^{j-1} \frac{1}{5 t-1} d t\right)\right\}>\exp \left\{-0.25+0.2 \ln 1.25-\frac{1}{5} \ln j\right\} \\
& >0.8 j^{-1 / 5}
\end{aligned}
$$

That is, the convergence is extremely slow. For example, for $j=10^{9}$, the error of the iterated solution is greater than 0.015 . On the other hand, for the optimal stepsize factor of $\theta=1 / c=5$, the optimal solution $x_{*}=0$ is found in one iteration.

It could be added that the stepsizes $\gamma_{j}=\theta / j$ may become completely unacceptable when $f$ loses strong convexity. For example, when $f(x)=x^{4}, X=[-1,1]$, and there is no noise, these stepsizes result in a disastrously slow convergence: $\left|x_{j}\right| \geq$ $O\left([\ln (j+1)]^{-1 / 2}\right)$. The precise statement here is that with $\gamma_{j}=\theta / j$ and $0<x_{1} \leq \frac{1}{6 \sqrt{\theta}}$, we have that $x_{j} \geq \frac{x_{1}}{\sqrt{1+32 \theta x_{1}^{2}[1+\ln (j+1)]}}$ for $j=1,2, \ldots$.

We see that in order to make the SA "robust"-applicable to general convex objectives rather than to strongly convex ones - one should replace the classical stepsizes $\gamma_{j}=O\left(j^{-1}\right)$, which can be too small to ensure a reasonable rate of convergence even in the "no noise" case, with "much larger" stepsizes. At the same time, a detailed analysis shows that "large" stepsizes poorly suppress noise. As early as in [15] it was realized that in order to resolve the arising difficulty, it makes sense to separate collecting information on the objective from generating approximate solutions. Specifically, we can use large stepsizes, say, $\gamma_{j}=O\left(j^{-1 / 2}\right)$ in (2.1), thus avoiding too slow motion at the cost of making the trajectory "more noisy." In order to suppress, to some extent, this noisiness, we take, as approximate solutions, appropriate averages of the search points $x_{j}$ rather than these points themselves.
2.2. Robust SA approach. Results of this section go back to Nemirovski and Yudin [15, 16]. Let us look again at the basic relations (2.2), (2.5), and (2.6). By convexity of $f(x)$, we have that $f(x) \geq f\left(x_{t}\right)+\left(x-x_{t}\right)^{T} \mathrm{~g}\left(x_{t}\right)$ for any $x \in X$, and hence

$$
\mathbb{E}\left[\left(x_{t}-x_{*}\right)^{T} \mathrm{~g}\left(x_{t}\right)\right] \geq \mathbb{E}\left[f\left(x_{t}\right)-f\left(x_{*}\right)\right] .
$$

Together with (2.6), this implies (recall that $a_{t}=\mathbb{E}\left[\frac{1}{2}\left\|x_{t}-x_{*}\right\|_{2}^{2}\right]$ )

$$
\gamma_{t} \mathbb{E}\left[f\left(x_{t}\right)-f\left(x_{*}\right)\right] \leq a_{t}-a_{t+1}+\frac{1}{2} \gamma_{t}^{2} M^{2} .
$$

It follows that whenever $1 \leq i \leq j$, we have

$$
\begin{equation*}
\sum_{t=i}^{j} \gamma_{t} \mathbb{E}\left[f\left(x_{t}\right)-f\left(x_{*}\right)\right] \leq \sum_{t=i}^{j}\left[a_{t}-a_{t+1}\right]+\frac{1}{2} M^{2} \sum_{t=i}^{j} \gamma_{t}^{2} \leq a_{i}+\frac{1}{2} M^{2} \sum_{t=i}^{j} \gamma_{t}^{2}, \tag{2.14}
\end{equation*}
$$

and hence, setting $\nu_{t}=\frac{\gamma_{t}}{\sum_{\tau=i}^{\gamma} \gamma_{\tau}}$,

$$
\begin{equation*}
\mathbb{E}\left[\sum_{t=i}^{j} \nu_{t} f\left(x_{t}\right)-f\left(x_{*}\right)\right] \leq \frac{a_{i}+\frac{1}{2} M^{2} \sum_{t=i}^{j} \gamma_{t}^{2}}{\sum_{t=i}^{j} \gamma_{t}} \tag{2.15}
\end{equation*}
$$

Note that $\nu_{t} \geq 0$ and $\sum_{t=i}^{j} \nu_{t}=1$. Consider the points

$$
\begin{equation*}
\tilde{x}_{i}^{j}=\sum_{t=i}^{j} \nu_{t} x_{t}, \tag{2.16}
\end{equation*}
$$

and let

$$
\begin{equation*}
D_{X}=\max _{x \in X}\left\|x-x_{1}\right\|_{2} \tag{2.17}
\end{equation*}
$$

By convexity of $X$, we have $\tilde{x}_{i}^{j} \in X$, and, by convexity of $f$, we have $f\left(\tilde{x}_{i}^{j}\right) \leq$ $\sum_{t=i}^{j} \nu_{t} f\left(x_{t}\right)$. Thus, by (2.15) and in view of $a_{1} \leq D_{X}^{2}$ and $a_{i} \leq 4 D_{X}^{2}, i>1$, we get

$$
\begin{align*}
& \text { (a) } \mathbb{E}\left[f\left(\tilde{x}_{1}^{j}\right)-f\left(x_{*}\right)\right] \leq \frac{D_{X}^{2}+M^{2} \sum_{t=1}^{j} \gamma_{t}^{2}}{2 \sum_{t=1}^{j} \gamma_{t}} \quad \text { for } \quad 1 \leq j, \\
& \text { (b) } \mathbb{E}\left[f\left(\tilde{x}_{i}^{j}\right)-f\left(x_{*}\right)\right] \leq \frac{4 D_{X}^{2}+M^{2} \sum_{t=i}^{j} \gamma_{t}^{2}}{2 \sum_{t=i}^{j} \gamma_{t}} \quad \text { for } \quad 1<i \leq j . \tag{2.18}
\end{align*}
$$

Based on the resulting bounds on the expected inaccuracy of approximate solutions $\tilde{x}_{i}^{j}$, we can now develop "reasonable" stepsize policies along with the associated efficiency estimates.

Constant stepsizes and basic efficiency estimate. Assume that the number $N$ of iterations of the method is fixed in advance and that $\gamma_{t}=\gamma, t=1, \ldots, N$. Then it follows by (2.18(a)) that

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{1}^{N}\right)-f\left(x_{*}\right)\right] \leq \frac{D_{X}^{2}+M^{2} N \gamma^{2}}{2 N \gamma} \tag{2.19}
\end{equation*}
$$

Minimizing the right-hand side of (2.19) over $\gamma>0$, we arrive at the constant stepsize policy

$$
\begin{equation*}
\gamma_{t}=\frac{D_{X}}{M \sqrt{N}}, \quad t=1, \ldots, N \tag{2.20}
\end{equation*}
$$

along with the associated efficiency estimate

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{1}^{N}\right)-f\left(x_{*}\right)\right] \leq \frac{D_{X} M}{\sqrt{N}} \tag{2.21}
\end{equation*}
$$

With the constant stepsize policy (2.20), we also have, for $1 \leq K \leq N$,

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{K}^{N}\right)-f\left(x_{*}\right)\right] \leq \frac{D_{X} M}{\sqrt{N}}\left[\frac{2 N}{N-K+1}+\frac{1}{2}\right] \tag{2.22}
\end{equation*}
$$

When $K / N \leq 1 / 2$, the right-hand side of (2.22) coincides, within an absolute constant factor, with the right-hand side of (2.21). Finally, for a constant $\theta>0$, passing from the stepsizes $(2.20)$ to the stepsizes

$$
\begin{equation*}
\gamma_{t}=\frac{\theta D_{X}}{M \sqrt{N}}, \quad t=1, \ldots, N \tag{2.23}
\end{equation*}
$$

the efficiency estimate becomes

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{K}^{N}\right)-f\left(x_{*}\right)\right] \leq \max \left\{\theta, \theta^{-1}\right\} \frac{D_{X} M}{\sqrt{N}}\left[\frac{2 N}{N-K+1}+\frac{1}{2}\right], 1 \leq K \leq N \tag{2.24}
\end{equation*}
$$

Discussion. We conclude that the expected error in terms of the objective of Robust $S A$ algorithm (2.1), (2.16), with constant stepsize policy (2.20), after $N$ iterations is of order $O\left(N^{-1 / 2}\right)$ in our setting. Of course, this is worse than the rate $O\left(N^{-1}\right)$ for the classical SA algorithm as applied to a smooth strongly convex function attaining minimum at a point from the interior of the set $X$. However, the error bounds (2.21)
and (2.22) are guaranteed independently of any smoothness and/or strong convexity assumptions on $f$. All that matters is the convexity of $f$ on the convex compact set $X$ and the validity of (2.5). Moreover, scaling the stepsizes by positive constant $\theta$ affects the error bound (2.24) linearly in $\max \left\{\theta, \theta^{-1}\right\}$. This can be compared with a possibly disastrous effect of such scaling in the classical SA algorithm discussed in section 2.1. These observations, in particular the fact that there is no necessity in "fine tuning" the stepsizes to the objective function $f$, explain the adjective "robust" in the name of the method. Finally, it can be shown that without additional, as compared to convexity and (2.5), assumptions on $f$, the accuracy bound (2.21) within an absolute constant factor is the best one allowed by statistics (cf. [16]).

Varying stepsizes. When the number of steps is not fixed in advance, it makes sense to replace constant stepsizes with the stepsizes

$$
\begin{equation*}
\gamma_{t}=\frac{\theta D_{X}}{M \sqrt{t}}, \quad t=1,2, \ldots \tag{2.25}
\end{equation*}
$$

From (2.18(b)) it follows that with this stepsize policy, one has, for $1 \leq K \leq N$,

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{K}^{N}\right)-f\left(x_{*}\right)\right] \leq \frac{D_{X} M}{\sqrt{N}}\left[\frac{2}{\theta}\left(\frac{N}{N-K+1}\right)+\frac{\theta}{2} \sqrt{\frac{N}{K}}\right] \tag{2.26}
\end{equation*}
$$

Choosing $K$ as a fixed fraction of $N$, i.e., setting $K=\lceil r N\rceil$, with a fixed $r \in(0,1)$, we get the efficiency estimate

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{K}^{N}\right)-f\left(x_{*}\right)\right] \leq C(r) \max \left\{\theta, \theta^{-1}\right\} \frac{D_{X} M}{\sqrt{N}}, \quad N=1,2, \ldots \tag{2.27}
\end{equation*}
$$

with an easily computable factor $C(r)$ depending solely on $r$. This bound, up to a factor depending solely on $r$ and $\theta$, coincides with the bound (2.21), with the advantage that our new stepsize policy should not be adjusted to a fixed-in-advance number of steps $N$.
2.3. Mirror descent SA method. On a close inspection, the robust SA algorithm from section 2.2 is intrinsically linked to the Euclidean structure of $\mathbb{R}^{n}$. This structure plays the central role in the very construction of the method (see (2.1)), the same as in the associated efficiency estimates, like (2.21) (since the quantities $D_{X}, M$ participating in the estimates are defined in terms of the Euclidean norm, see (2.17) and (2.5)). By these reasons, from now on, we refer to the algorithm from section 2.2 as the (robust) Euclidean $S A$ (E-SA). In this section we develop a substantial generalization of the E-SA approach allowing us to adjust, to some extent, the method to the geometry, not necessary Euclidean, of the problem in question. We shall see in the meantime that we can gain a lot, both theoretically and numerically, from such an adjustment. A rudimentary form of the generalization to follow can be found in Nemirovski and Yudin [16], from where the name "mirror descent" originates.

Let $\|\cdot\|$ be a (general) norm on $\mathbb{R}^{n}$ and $\|x\|_{*}=\sup _{\|y\| \leq 1} y^{T} x$ be its dual norm. We say that a function $\omega: X \rightarrow \mathbb{R}$ is a distance-generating function modulus $\alpha>0$ with respect to $\|\cdot\|$, if $\omega$ is convex and continuous on $X$, the set

$$
\begin{equation*}
X^{o}=\{x \in X: \partial \omega(x) \neq \emptyset\} \tag{2.28}
\end{equation*}
$$

is convex (note that $X^{o}$ always contains the relative interior of $X$ ) and restricted to $X^{o}, \omega$ is continuously differentiable and strongly convex with parameter $\alpha$ with
respect to $\|\cdot\|$, i.e.,

$$
\begin{equation*}
\left(x^{\prime}-x\right)^{T}\left(\nabla \omega\left(x^{\prime}\right)-\nabla \omega(x)\right) \geq \alpha\left\|x^{\prime}-x\right\|^{2} \forall x^{\prime}, x \in X^{o} . \tag{2.29}
\end{equation*}
$$

A simple example of a distance-generating function is $\omega(x)=\frac{1}{2}\|x\|_{2}^{2}$ (modulus 1 with respect to $\left.\|\cdot\|_{2}, X^{o}=X\right)$.

Let us define function $V: X^{o} \times X \rightarrow \mathbb{R}_{+}$as follows:

$$
\begin{equation*}
V(x, z)=\omega(z)-\left[\omega(x)+\nabla \omega(x)^{T}(z-x)\right] . \tag{2.30}
\end{equation*}
$$

In what follows we shall refer to $V(\cdot, \cdot)$ as prox-function associated with distancegenerating function $\omega(x)$ (it is also called Bregman distance [4]). Note that $V(x, \cdot)$ is nonnegative and is a strongly convex modulus $\alpha$ with respect to the norm $\|\cdot\|$. Let us define prox-mapping $P_{x}: \mathbb{R}^{n} \rightarrow X^{o}$, associated with $\omega$ and a point $x \in X^{o}$, viewed as a parameter, as follows:

$$
\begin{equation*}
P_{x}(y)=\arg \min _{z \in X}\left\{y^{T}(z-x)+V(x, z)\right\} \tag{2.31}
\end{equation*}
$$

Observe that the minimum in the right-hand side of (2.31) is attained since $\omega$ is continuous on $X$ and $X$ is compact, and all the minimizers belong to $X^{o}$, whence the minimizer is unique, since $V(x, \cdot)$ is strongly convex on $X^{o}$. Thus, the prox-mapping is well defined.

For $\omega(x)=\frac{1}{2}\|x\|_{2}^{2}$, we have $P_{x}(y)=\Pi_{X}(x-y)$ so that (2.1) is the recurrence

$$
\begin{equation*}
x_{j+1}=P_{x_{j}}\left(\gamma_{j} \mathrm{G}\left(x_{j}, \xi_{j}\right)\right), \quad x_{1} \in X^{o} \tag{2.32}
\end{equation*}
$$

Our goal is to demonstrate that the main properties of the recurrence (2.1) (which from now on we call the $E$-SA recurrence) are inherited by (2.32), whatever be the underlying distance-generating function $\omega(x)$.

The statement of the following lemma is a simple consequence of the optimality conditions of the right-hand side of (2.31) (proof of this lemma is given in the appendix).

Lemma 2.1. For every $u \in X, x \in X^{o}$, and $y \in \mathbb{R}^{n}$, one has

$$
\begin{equation*}
V\left(P_{x}(y), u\right) \leq V(x, u)+y^{T}(u-x)+\frac{\|y\|_{*}^{2}}{2 \alpha} \tag{2.33}
\end{equation*}
$$

Using (2.33) with $x=x_{j}, y=\gamma_{j} \mathrm{G}\left(x_{j}, \xi_{j}\right)$, and $u=x_{*}$, we get

$$
\begin{equation*}
\gamma_{j}\left(x_{j}-x_{*}\right)^{T} \mathrm{G}\left(x_{j}, \xi_{j}\right) \leq V\left(x_{j}, x_{*}\right)-V\left(x_{j+1}, x_{*}\right)+\frac{\gamma_{j}^{2}}{2 \alpha}\left\|\mathrm{G}\left(x_{j}, \xi_{j}\right)\right\|_{*}^{2} \tag{2.34}
\end{equation*}
$$

Note that with $\omega(x)=\frac{1}{2}\|x\|_{2}^{2}$, one has $V(x, z)=\frac{1}{2}\|x-z\|_{2}^{2}, \alpha=1,\|\cdot\|_{*}=\|\cdot\|_{2}$. That is, (2.34) becomes nothing but the relation (2.6), which played a crucial role in all the developments related to the E-SA method. We are about to process, in a completely similar fashion, the relation (2.34) in the case of a general distancegenerating function, thus arriving at the mirror descent SA. Specifically, setting

$$
\begin{equation*}
\Delta_{j}=\mathrm{G}\left(x_{j}, \xi_{j}\right)-\mathrm{g}\left(x_{j}\right) \tag{2.35}
\end{equation*}
$$

we can rewrite (2.34), with $j$ replaced by $t$, as

$$
\begin{equation*}
\gamma_{t}\left(x_{t}-x_{*}\right)^{T} \mathrm{~g}\left(x_{t}\right) \leq V\left(x_{t}, x_{*}\right)-V\left(x_{t+1}, x_{*}\right)-\gamma_{t} \Delta_{t}^{T}\left(x_{t}-x_{*}\right)+\frac{\gamma_{t}^{2}}{2 \alpha}\left\|\mathrm{G}\left(x_{t}, \xi_{t}\right)\right\|_{*}^{2} \tag{2.36}
\end{equation*}
$$

Summing up over $t=1, \ldots, j$, and taking into account that $V\left(x_{j+1}, u\right) \geq 0, u \in X$, we get

$$
\begin{equation*}
\sum_{t=1}^{j} \gamma_{t}\left(x_{t}-x_{*}\right)^{T} \mathrm{~g}\left(x_{t}\right) \leq V\left(x_{1}, x_{*}\right)+\sum_{t=1}^{j} \frac{\gamma_{t}^{2}}{2 \alpha}\left\|\mathrm{G}\left(x_{t}, \xi_{t}\right)\right\|_{*}^{2}-\sum_{t=1}^{j} \gamma_{t} \Delta_{t}^{T}\left(x_{t}-x_{*}\right) \tag{2.37}
\end{equation*}
$$

Setting $\nu_{t}=\frac{\gamma_{t}}{\sum_{i=1}^{j} \gamma_{i}}, t=1, \ldots, j$, and

$$
\begin{equation*}
\tilde{x}_{1}^{j}=\sum_{t=1}^{j} \nu_{t} x_{t} \tag{2.38}
\end{equation*}
$$

and invoking convexity of $f(\cdot)$, we have

$$
\begin{aligned}
\sum_{t=1}^{j} \gamma_{t}\left(x_{t}-x_{*}\right)^{T} \mathrm{~g}\left(x_{t}\right) & \geq \sum_{t=1}^{j} \gamma_{t}\left[f\left(x_{t}\right)-f\left(x_{*}\right)\right] \\
& =\left(\sum_{t=1}^{j} \gamma_{t}\right)\left[\sum_{t=1}^{j} \nu_{t} f\left(x_{t}\right)-f\left(x_{*}\right)\right] \\
& \geq\left(\sum_{t=1}^{j} \gamma_{t}\right)\left[f\left(\tilde{x}_{j}\right)-f\left(x_{*}\right)\right]
\end{aligned}
$$

which combines with (2.37) to imply that

$$
\begin{equation*}
f\left(\tilde{x}_{1}^{j}\right)-f\left(x_{*}\right) \leq \frac{V\left(x_{1}, x_{*}\right)+\sum_{t=1}^{j} \frac{\gamma_{t}^{2}}{2 \alpha}\left\|\mathrm{G}\left(x_{t}, \xi_{t}\right)\right\|_{*}^{2}-\sum_{t=1}^{j} \gamma_{t} \Delta_{t}^{T}\left(x_{t}-x_{*}\right)}{\sum_{t=1}^{j} \gamma_{t}} \tag{2.39}
\end{equation*}
$$

Let us suppose, as in the previous section (cf. (2.5)), that we are given a positive number $M_{*}$ such that

$$
\begin{equation*}
\mathbb{E}\left[\|\mathrm{G}(x, \xi)\|_{*}^{2}\right] \leq M_{*}^{2} \forall x \in X \tag{2.40}
\end{equation*}
$$

Taking expectations of both sides of (2.39) and noting that (i) $x_{t}$ is a deterministic function of $\xi_{[t-1]}=\left(\xi_{1}, \ldots, \xi_{t-1}\right)$, (ii) conditional on $\xi_{[t-1]}$, the expectation of $\Delta_{t}$ is 0 , and (iii) the expectation of $\left\|\mathrm{G}\left(x_{t}, \xi_{t}\right)\right\|_{*}^{2}$ does not exceed $M_{*}^{2}$, we obtain

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{1}^{j}\right)-f\left(x_{*}\right)\right] \leq \frac{\max _{u \in X} V\left(x_{1}, u\right)+(2 \alpha)^{-1} M_{*}^{2} \sum_{t=1}^{j} \gamma_{t}^{2}}{\sum_{t=1}^{j} \gamma_{t}} \tag{2.41}
\end{equation*}
$$

Assume from now on that the method starts with the minimizer of $\omega$ :

$$
x_{1}=\operatorname{argmin}_{X} \omega(x) .
$$

Then, from (2.30), it follows that

$$
\begin{equation*}
\max _{z \in X} V\left(x_{1}, z\right) \leq D_{\omega, X}^{2} \tag{2.42}
\end{equation*}
$$

where

$$
\begin{equation*}
D_{\omega, X}:=\left[\max _{z \in X} \omega(z)-\min _{z \in X} \omega(z)\right]^{1 / 2} \tag{2.43}
\end{equation*}
$$

Consequently, (2.41) implies that

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{1}^{j}\right)-f\left(x_{*}\right)\right] \leq \frac{D_{\omega, X}^{2}+\frac{1}{2 \alpha} M_{*}^{2} \sum_{t=1}^{j} \gamma_{t}^{2}}{\sum_{t=1}^{j} \gamma_{t}} \tag{2.44}
\end{equation*}
$$

Constant stepsize policy. Assuming that the total number of steps $N$ is given in advance and $\gamma_{t}=\gamma, t=1, \ldots, N$, optimizing the right-hand side of (2.44) over $\gamma>0$ we arrive at the constant stepsize policy

$$
\begin{equation*}
\gamma_{t}=\frac{\sqrt{2 \alpha} D_{\omega, X}}{M_{*} \sqrt{N}}, \quad t=1, \ldots, N \tag{2.45}
\end{equation*}
$$

and the associated efficiency estimate

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{1}^{N}\right)-f\left(x_{*}\right)\right] \leq D_{\omega, X} M_{*} \sqrt{\frac{2}{\alpha N}} \tag{2.46}
\end{equation*}
$$

(cf. (2.20), (2.21)). For a constant $\theta>0$, passing from the stepsizes (2.45) to the stepsizes

$$
\begin{equation*}
\gamma_{t}=\frac{\theta \sqrt{2 \alpha} D_{\omega, X}}{M_{*} \sqrt{N}}, t=1, \ldots, N \tag{2.47}
\end{equation*}
$$

the efficiency estimate becomes

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{1}^{N}\right)-f\left(x_{*}\right)\right] \leq \max \left\{\theta, \theta^{-1}\right\} D_{\omega, X} M_{*} \sqrt{\frac{2}{\alpha N}} \tag{2.48}
\end{equation*}
$$

We refer to the method (2.32), (2.38), and (2.47) as the (robust) mirror descent $S A$ algorithm with constant stepsize policy.

Probabilities of large deviations. So far, all our efficiency estimates were upper bounds on the expected nonoptimality, in terms of the objective, of approximate solutions generated by the algorithms. Here we complement these results with bounds on probabilities of large deviations. Observe that by Markov inequality, (2.48) implies that

$$
\begin{equation*}
\operatorname{Prob}\left\{f\left(\tilde{x}_{1}^{N}\right)-f\left(x_{*}\right)>\varepsilon\right\} \leq \frac{\sqrt{2} \max \left\{\theta, \theta^{-1}\right\} D_{\omega, X} M_{*}}{\varepsilon \sqrt{\alpha N}} \forall \varepsilon>0 \tag{2.49}
\end{equation*}
$$

It is possible, however, to obtain much finer bounds on deviation probabilities when imposing more restrictive assumptions on the distribution of $\mathrm{G}(x, \xi)$. Specifically, assume that

$$
\begin{equation*}
\mathbb{E}\left[\exp \left\{\|\mathrm{G}(x, \xi)\|_{*}^{2} / M_{*}^{2}\right\}\right] \leq \exp \{1\} \quad \forall x \in X \tag{2.50}
\end{equation*}
$$

Note that condition (2.50) is stronger than (2.40). Indeed, if a random variable $Y$ satisfies $\mathbb{E}[\exp \{Y / a\}] \leq \exp \{1\}$ for some $a>0$, then by Jensen inequality, $\exp \{\mathbb{E}[Y / a]\} \leq$ $\mathbb{E}[\exp \{Y / a\}] \leq \exp \{1\}$, and therefore, $\mathbb{E}[Y] \leq a$. Of course, condition (2.50) holds if $\|\mathrm{G}(x, \xi)\|_{*} \leq M_{*}$ for all $(x, \xi) \in X \times \Xi$.

PROPOSITION 2.2. In the case of (2.50) and for the constant stepsizes (2.47), the following holds for any $\Omega \geq 1$ :

$$
\begin{equation*}
\operatorname{Prob}\left\{f\left(\tilde{x}_{1}^{N}\right)-f\left(x_{*}\right)>\frac{\sqrt{2} \max \left\{\theta, \theta^{-1}\right\} M_{*} D_{\omega, X}(12+2 \Omega)}{\sqrt{\alpha N}}\right\} \leq 2 \exp \{-\Omega\} \tag{2.51}
\end{equation*}
$$

Proof of this proposition is given in the appendix.

Varying stepsizes. Same as in the case of E-SA, we can modify the mirror descent SA algorithm to allow for time-varying stepsizes and "sliding averages" of the search points $x_{t}$ in the role of approximate solutions, thus getting rid of the necessity to fix in advance the number of steps. Specifically, consider

$$
\begin{align*}
\bar{D}_{\omega, X} & :=\sqrt{2} \sup _{x \in X^{o}, z \in X}\left[\omega(z)-\omega(x)-(z-x)^{T} \nabla \omega(x)\right]^{1 / 2}  \tag{2.52}\\
& =\sup _{x \in X^{o}, z \in X} \sqrt{2 V(x, z)}
\end{align*}
$$

and assume that $\bar{D}_{\omega, X}$ is finite. This is definitely so when $\omega$ is continuously differentiable on the entire $X$. Note that for the E-SA, that is, with $\omega(x)=\frac{1}{2}\|x\|_{2}^{2}, \bar{D}_{\omega, X}$ is the Euclidean diameter of $X$.

In the case of (2.52), setting

$$
\begin{equation*}
\tilde{x}_{i}^{j}=\frac{\sum_{t=i}^{j} \gamma_{t} x_{t}}{\sum_{t=i}^{j} \gamma_{t}} \tag{2.53}
\end{equation*}
$$

summing up inequalities (2.34) over $K \leq t \leq N$, and acting exactly as when deriving (2.39), we get for $1 \leq K \leq N$,

$$
f\left(\tilde{x}_{K}^{N}\right)-f\left(x_{*}\right) \leq \frac{V\left(x_{K}, x_{*}\right)+\sum_{t=K}^{N} \frac{\gamma_{t}^{2}}{2 \alpha}\left\|\mathrm{G}\left(x_{t}, \xi_{t}\right)\right\|_{*}^{2}-\sum_{t=K}^{N} \gamma_{t} \Delta_{t}^{T}\left(x_{t}-x_{*}\right)}{\sum_{t=K}^{N} \gamma_{t}} .
$$

Noting that $V\left(x_{K}, x_{*}\right) \leq \frac{1}{2} \bar{D}_{\omega, X}^{2}$ and taking expectations, we arrive at

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{K}^{N}\right)-f\left(x_{*}\right)\right] \leq \frac{\frac{1}{2} \bar{D}_{\omega, X}^{2}+\frac{1}{2 \alpha} M_{*}^{2} \sum_{t=K}^{N} \gamma_{t}^{2}}{\sum_{t=K}^{N} \gamma_{t}} \tag{2.54}
\end{equation*}
$$

(cf. (2.44)). It follows that with a decreasing stepsize policy

$$
\begin{equation*}
\gamma_{t}=\frac{\theta \bar{D}_{\omega, X} \sqrt{\alpha}}{M_{*} \sqrt{t}}, \quad t=1,2, \ldots \tag{2.55}
\end{equation*}
$$

one has for $1 \leq K \leq N$,

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{K}^{N}\right)-f\left(x_{*}\right)\right] \leq \frac{\bar{D}_{\omega, X} M_{*}}{\sqrt{\alpha} \sqrt{N}}\left[\frac{2}{\theta} \frac{N}{N-K+1}+\frac{\theta}{2} \sqrt{\frac{N}{K}}\right] \tag{2.56}
\end{equation*}
$$

(cf. (2.26)). In particular, with $K=\lceil r N\rceil$ for a fixed $r \in(0,1)$, we get an efficiency estimate

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{K}^{N}\right)-f\left(x_{*}\right)\right] \leq C(r) \max \left\{\theta, \theta^{-1}\right\} \frac{\bar{D}_{\omega, X} M_{*}}{\sqrt{\alpha} \sqrt{N}} \tag{2.57}
\end{equation*}
$$

completely similar to the estimate (2.27) for the E-SA.
Discussion. Comparing (2.21) to (2.46) and (2.27) to (2.57), we see that for both the Euclidean and the mirror descent robust SA, the expected inaccuracy, in terms of the objective, of the approximate solution built in course of $N$ steps is $O\left(N^{-1 / 2}\right)$. A benefit of the mirror descent over the Euclidean algorithm is in the
potential possibility to reduce the constant factor hidden in $O(\cdot)$ by adjusting the norm $\|\cdot\|$ and the distance-generating function $\omega(\cdot)$ to the geometry of the problem.

Example. Let $X=\left\{x \in \mathbb{R}^{n}: \sum_{i=1}^{n} x_{i}=1, x \geq 0\right\}$ be a standard simplex. Consider two setups for the mirror descent SA:

- Euclidean setup, where $\|\cdot\|=\|\cdot\|_{2}$ and $\omega(x)=\frac{1}{2}\|x\|_{2}^{2}$, and
- $\ell_{1}$-setup, where $\|\cdot\|=\|\cdot\|_{1}$, with $\|\cdot\|_{*}=\|\cdot\|_{\infty}$ and $\omega$ is the entropy function

$$
\begin{equation*}
\omega(x)=\sum_{i=1}^{n} x_{i} \ln x_{i} \tag{2.58}
\end{equation*}
$$

The Euclidean setup leads to the Euclidean robust SA, which is easily implementable (computing the prox-mapping requires $O(n \ln n)$ operations) and guarantees that

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{1}^{N}\right)-f\left(x_{*}\right)\right] \leq O(1) \max \left\{\theta, \theta^{-1}\right\} M N^{-1 / 2} \tag{2.59}
\end{equation*}
$$

with $M^{2}=\sup _{x \in X} \mathbb{E}\left[\|\mathrm{G}(x, \xi)\|_{2}^{2}\right]$, provided that the constant $M$ is known and the stepsizes (2.23) are used (see (2.24), (2.17), and note that the Euclidean diameter of $X$ is $\sqrt{2}$ ).

The $\ell_{1}$-setup corresponds to $X^{o}=\{x \in X: x>0\}, D_{\omega, X}=\sqrt{\ln n}, \alpha=1$, and $x_{1}=\operatorname{argmin}_{X} \omega=n^{-1}(1, \ldots, 1)^{T}$ (see appendix). The associated mirror descent SA is easily implementable: the prox-function here is

$$
V(x, z)=\sum_{i=1}^{n} z_{i} \ln \frac{z_{i}}{x_{i}}
$$

and the prox-mapping $P_{x}(y)=\operatorname{argmin}_{z \in X}\left[y^{T}(z-x)+V(x, z)\right]$ can be computed in $O(n)$ operations according to the explicit formula

$$
\left[P_{x}(y)\right]_{i}=\frac{x_{i} e^{-y_{i}}}{\sum_{k=1}^{n} x_{k} e^{-y_{k}}}, \quad i=1, \ldots, n
$$

The efficiency estimate guaranteed with the $\ell_{1}$-setup is

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{1}^{N}\right)-f\left(x_{*}\right)\right] \leq O(1) \max \left\{\theta, \theta^{-1}\right\} \sqrt{\ln n} M_{*} N^{-1 / 2}, \tag{2.60}
\end{equation*}
$$

with

$$
M_{*}^{2}=\sup _{x \in X} \mathbb{E}\left[\|\mathrm{G}(x, \xi)\|_{\infty}^{2}\right]
$$

provided that the constant $M_{*}$ is known and the constant stepsizes (2.47) are used (see (2.48) and (2.40)). To compare (2.60) and (2.59), observe that $M_{*} \leq M$, and the ratio $M_{*} / M$ can be as small as $n^{-1 / 2}$. Thus, the efficiency estimate for the $\ell_{1}$-setup never is much worse than the estimate for the Euclidean setup, and for large $n$, can be far better than the latter estimate:

$$
\sqrt{\frac{1}{\ln n}} \leq \frac{M}{\sqrt{\ln n} M_{*}} \leq \sqrt{\frac{n}{\ln n}}, N=1,2, \ldots
$$

both the upper and the lower bounds being achievable. Thus, when $X$ is a standard simplex of large dimension, we have strong reasons to prefer the $\ell_{1}$-setup to the usual Euclidean one.

Note that $\|\cdot\|_{1}$-norm can be coupled with "good" distance-generating functions different from the entropy one, e.g., with the function

$$
\begin{equation*}
\omega(x)=(\ln n) \sum_{i=1}^{n}\left|x_{i}\right|^{1+\frac{1}{\ln n}}, \quad n \geq 3 \tag{2.61}
\end{equation*}
$$

Whenever $0 \in X$ and $\operatorname{Diam}_{\|\cdot\|_{1}}(X) \equiv \max _{x, y \in X}\|x-y\|_{1}$ equal to 1 (these conditions can always be ensured by scaling and shifting $X$ ), for the just-outlined setup, one has $\bar{D}_{\omega, X}=O(1) \sqrt{\ln n}, \alpha=O(1)$, so that the associated mirror descent robust SA guarantees that with $M_{*}^{2}=\sup _{x \in X} \mathbb{E}\left[\|\mathrm{G}(x, \xi)\|_{\infty}^{2}\right]$ and $N \geq 1$,

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{\lceil r N\rceil}^{N}\right)-f\left(x_{*}\right)\right] \leq C(r) \frac{M_{*} \sqrt{\ln n}}{\sqrt{N}} \tag{2.62}
\end{equation*}
$$

(see (2.57)), while the efficiency estimate for the Euclidean robust SA is

$$
\begin{equation*}
\mathbb{E}\left[f\left(\tilde{x}_{\lceil r N\rceil}^{N}\right)-f\left(x_{*}\right)\right] \leq C(r) \frac{M \operatorname{Diam}_{\|\cdot\|_{2}}(X)}{\sqrt{N}} \tag{2.63}
\end{equation*}
$$

with

$$
M^{2}=\sup _{x \in X} \mathbb{E}\left[\|\mathrm{G}(x, \xi)\|_{2}^{2}\right] \text { and } \operatorname{Diam}_{\|\cdot\|_{2}}(X)=\max _{x, y \in X}\|x-y\|_{2}
$$

Ignoring logarithmic in $n$ factors, the second estimate (2.63) can be much better than the first estimate (2.62) only when $\operatorname{Diam}_{\|\cdot\|_{2}}(X) \ll 1=\operatorname{Diam}_{\|\cdot\|_{1}}(X)$, as it is the case, e.g., when $X$ is an Euclidean ball. On the other hand, when $X$ is an $\|\cdot\|_{1}$-ball or its nonnegative part (which is the simplex), so that the $\|\cdot\|_{1^{-}}$and $\|\cdot\|_{2}$-diameters of $X$ are of the same order, the first estimate (2.62) is much more attractive than the estimate (2.63) due to potentially much smaller constant $M_{*}$.

Comparison with the SAA approach. We compare now theoretical complexity estimates for the robust mirror descent SA and the SAA methods. Consider the case when (i) $X \subset \mathbb{R}^{n}$ is contained in the $\|\cdot\|_{p}$-ball of radius $R, p=1,2$, and the SA in question is either the E-SA $(p=2)$, or the SA associated with $\|\cdot\|_{1}$ and the distancegenerating function ${ }^{2}$ (2.61), (ii) in SA, the constant stepsize rule (2.45) is used, and (iii) the "light tail" assumption (2.50) takes place.

Given $\varepsilon>0, \delta \in(0,1 / 2)$, let us compare the number of steps $N=N_{\text {SA }}$ of SA, which, with probability $\geq 1-\delta$, results in an approximate solution $\tilde{x}_{1}^{N}$ such that $f\left(\tilde{x}_{1}^{N}\right)-f\left(x_{*}\right) \leq \varepsilon$, with the sample size $N=N_{\text {SAA }}$ for the SAA resulting in the same accuracy guarantees. According to Proposition 2.2 we have that $\operatorname{Prob}\left[f\left(\tilde{x}_{1}^{N}\right)-f\left(x_{*}\right)>\varepsilon\right] \leq \delta$ for

$$
\begin{equation*}
N_{\mathrm{SA}}=O(1) \varepsilon^{-2} D_{\omega, X}^{2} M_{*}^{2} \ln ^{2}(1 / \delta) \tag{2.64}
\end{equation*}
$$

where $M_{*}$ is the constant from (2.50) and $D_{\omega, X}$ is defined in (2.43). Note that the constant $M_{*}$ depends on the chosen norm, $D_{\omega, X}^{2}=O(1) R^{2}$ for $p=2$, and $D_{\omega, X}^{2}=$ $O(1) \ln (n) R^{2}$ for $p=1$.

This can be compared with the estimate of the sample size (cf. [25, 26])

$$
\begin{equation*}
N_{\mathrm{SAA}}=O(1) \varepsilon^{-2} R^{2} M_{*}^{2}\left[\ln (1 / \delta)+n \ln \left(R M_{*} / \varepsilon\right)\right] \tag{2.65}
\end{equation*}
$$

[^2]We see that both SA and SAA methods have logarithmic in $\delta$ and quadratic (or nearly so) in $1 / \varepsilon$ complexity in terms of the corresponding sample sizes. It should be noted, however, that the SAA method requires solution of the corresponding (deterministic) problem, while the SA approach is based on simple calculations as long as stochastic subgradients could be easily computed.
3. Stochastic saddle point problem. We show in this section how the mirror descent SA algorithm can be modified to solve a convex-concave stochastic saddle point problem. Consider the following minimax (saddle point) problem:

$$
\begin{equation*}
\min _{x \in X} \max _{y \in Y}\{\phi(x, y)=\mathbb{E}[\Phi(x, y, \xi)]\} \tag{3.1}
\end{equation*}
$$

Here $X \subset \mathbb{R}^{n}$ and $Y \subset \mathbb{R}^{m}$ are nonempty bounded closed convex sets, $\xi$ is a random vector whose probability distribution $P$ is supported on set $\Xi \subset \mathbb{R}^{d}$, and $\Phi: X \times Y \times$ $\Xi \rightarrow \mathbb{R}$. We assume that for every $(x, y) \in X \times Y$, the expectation

$$
\mathbb{E}[\Phi(x, y, \xi)]=\int_{\Xi} \Phi(x, y, \xi) d P(\xi)
$$

is well defined and finite valued and that the expected value function $\phi(x, y)$ is convex in $x \in X$ and concave in $y \in Y$. It follows that (3.1) is a convex-concave saddle point problem. In addition, we assume that $\phi(\cdot, \cdot)$ is Lipschitz continuous on $X \times Y$. It is well known that, in the above setting, (3.1) is solvable, i.e., the corresponding "primal" and "dual" optimization problems

$$
\min _{x \in X}\left[\max _{y \in Y} \phi(x, y)\right] \quad \text { and } \quad \max _{y \in Y}\left[\min _{x \in X} \phi(x, y)\right]
$$

respectively, are solvable with equal optimal values, denoted $\phi^{*}$, and pairs $\left(x^{*}, y^{*}\right)$ of optimal solutions to the respective problems form the set of saddle points of $\phi(x, y)$ on $X \times Y$.

As in the case of the minimization problem (1.1), we assume that neither the function $\phi(x, y)$ nor its sub/supergradients in $x$ and $y$ are available explicitly. However, we make the following assumption.
( $\left.\mathbf{A}^{\prime} \mathbf{2}\right)$ We have at our disposal an oracle which, given an input of point $(x, y, \xi) \in$ $X \times Y \times \Xi$, returns a stochastic subgradient, that is, $(n+m)$-dimensional vector $\mathrm{G}(x, y, \xi)=\left[\begin{array}{c}\mathrm{G}_{x}(x, y, \xi) \\ -\mathrm{G}_{y}(x, y, \xi)\end{array}\right]$ such that vector

$$
\mathrm{g}(x, y)=\left[\begin{array}{c}
\mathrm{g}_{x}(x, y) \\
-\mathrm{g}_{y}(x, y)
\end{array}\right]=\left[\begin{array}{c}
\mathbb{E}\left[\mathrm{G}_{x}(x, y, \xi)\right] \\
-\mathbb{E}\left[\mathrm{G}_{y}(x, y, \xi)\right]
\end{array}\right]
$$

is well defined, $\mathrm{g}_{x}(x, y) \in \partial_{x} \phi(x, y)$, and $-\mathrm{g}_{y}(x, y) \in \partial_{y}(-\phi(x, y))$.
For example, if for every $\xi \in \Xi$ the function $\Phi(\cdot, \cdot, \xi)$ is convex-concave and the respective subdifferential and integral operators are interchangeable, we ensure ( $\mathrm{A}^{\prime} 2$ ) by setting

$$
\mathrm{G}(x, y, \xi)=\left[\begin{array}{c}
\mathrm{G}_{x}(x, y, \xi) \\
-\mathrm{G}_{y}(x, y, \xi)
\end{array}\right] \in\left[\begin{array}{c}
\partial_{x} \Phi(x, y, \xi) \\
\partial_{y}(-\Phi(x, y, \xi))
\end{array}\right] .
$$

Let $\|\cdot\|_{x}$ be a norm on $\mathbb{R}^{n}$ and $\|\cdot\|_{y}$ be a norm on $\mathbb{R}^{m}$, and let $\|\cdot\|_{*, x}$ and $\|\cdot\|_{*, y}$ stand for the corresponding dual norms. As in section 2.1, the basic assumption we make about the stochastic oracle (aside from its unbiasedness, which we have already postulated) is that we know positive constants $M_{*, x}^{2}$ and $M_{*, y}^{2}$ such that

$$
\begin{equation*}
\mathbb{E}\left[\left\|\mathrm{G}_{x}(u, v, \xi)\right\|_{*, x}^{2}\right] \leq M_{*, x}^{2} \text { and } \mathbb{E}\left[\left\|\mathrm{G}_{y}(u, v, \xi)\right\|_{*, y}^{2}\right] \leq M_{*, y}^{2} \forall(u, v) \in X \times Y \tag{3.2}
\end{equation*}
$$

3.1. Mirror SA algorithm for saddle point problems. We equip $X$ and $Y$ with distance-generating functions $\omega_{x}: X \rightarrow \mathbb{R}$ modulus $\alpha_{x}$ with respect to $\|\cdot\|_{x}$, and $\omega_{y}: Y \rightarrow \mathbb{R}$ modulus $\alpha_{y}$ with respect to $\|\cdot\|_{y}$. Let $D_{\omega_{x}, X}$ and $D_{\omega_{y}, Y}$ be the respective constants (see definition (2.42)). We equip $\mathbb{R}^{n} \times \mathbb{R}^{m}$ with the norm

$$
\begin{equation*}
\|(x, y)\|=\sqrt{\frac{\alpha_{x}}{2 D_{\omega_{x}, X}^{2}}\|x\|_{x}^{2}+\frac{\alpha_{y}}{2 D_{\omega_{y}, Y}^{2}}\|y\|_{y}^{2}} \tag{3.3}
\end{equation*}
$$

so that the dual norm is

$$
\begin{equation*}
\|(\zeta, \eta)\|_{*}=\sqrt{\frac{2 D_{\omega_{x}, X}^{2}}{\alpha_{x}}\|\zeta\|_{*, x}^{2}+\frac{2 D_{\omega_{y}, Y}^{2}}{\alpha_{y}}\|\eta\|_{*, y}^{2}} \tag{3.4}
\end{equation*}
$$

and set

$$
\begin{equation*}
M_{*}^{2}=\frac{2 D_{\omega_{x}, X}^{2}}{\alpha_{x}} M_{*, x}^{2}+\frac{2 D_{\omega_{y}, Y}^{2}}{\alpha_{y}} M_{*, y}^{2} \tag{3.5}
\end{equation*}
$$

It follows by (3.2) that

$$
\begin{equation*}
\mathbb{E}\left[\|\mathrm{G}(x, y, \xi)\|_{*}^{2}\right] \leq M_{*}^{2} . \tag{3.6}
\end{equation*}
$$

We use the notation $z=(x, y)$ and equip the set $Z=X \times Y$ with the distancegenerating function

$$
\omega(z)=\frac{\omega_{x}(x)}{2 D_{\omega_{x}, X}^{2}}+\frac{\omega_{y}(y)}{2 D_{\omega_{y}, Y}^{2}}
$$

It is immediately seen that $\omega$ indeed is a distance-generating function for $Z$ modulus $\alpha=1$ with respect to the norm $\|\cdot\|$ and that $Z^{o}=X^{o} \times Y^{o}$ and $D_{\omega, Z}=1$. In what follows, $V(z, u): Z^{o} \times Z \rightarrow \mathbb{R}$ and $P_{z}(\zeta): \mathbb{R}^{n+m} \rightarrow Z^{o}$ are the prox-function and prox-mapping associated with $\omega$ and $Z$ (see (2.30), (2.31)).

We are ready now to present the mirror SA algorithm for saddle point problems. This is the iterative procedure (compare with (2.32))

$$
\begin{equation*}
z_{j+1}=P_{z_{j}}\left(\gamma_{j} \mathrm{G}\left(z_{j}, \xi_{j}\right)\right) \tag{3.7}
\end{equation*}
$$

where the initial point $z_{1} \in Z$ is chosen to be the minimizer of $\omega(z)$ on $Z$. As before (compare with (2.38)), we define approximate solution $\tilde{z}_{1}^{j}=\left(\tilde{x}_{1}^{j}, \tilde{y}_{1}^{j}\right)$ of (3.1) after $j$ iterations as

$$
\begin{equation*}
\tilde{z}_{1}^{j}=\frac{\sum_{t=1}^{j} \gamma_{t} z_{t}}{\sum_{t=1}^{j} \gamma_{t}} \tag{3.8}
\end{equation*}
$$

We refer to the procedure (3.7), (3.8) as the saddle point mirror $S A$ algorithm.
Let us analyze convergence properties of the algorithm. We measure quality of an approximate solution $\tilde{z}=(\tilde{x}, \tilde{y})$ by the error

$$
\epsilon_{\phi}(\tilde{z}):=\left[\max _{y \in Y} \phi(\tilde{x}, y)-\phi_{*}\right]+\left[\phi_{*}-\min _{x \in X} \phi(x, \tilde{y})\right]=\max _{y \in Y} \phi(\tilde{x}, y)-\min _{x \in X} \phi(x, \tilde{y})
$$

By convexity of $\phi(\cdot, y)$, we have

$$
\phi\left(x_{t}, y_{t}\right)-\phi\left(x, y_{t}\right) \leq\left(x_{t}-x\right)^{T} \mathrm{~g}_{x}\left(x_{t}, y_{t}\right) \quad \forall x \in X
$$

and by concavity of $\phi(x, \cdot)$,

$$
\phi\left(x_{t}, y\right)-\phi\left(x_{t}, y_{t}\right) \leq\left(y-y_{t}\right)^{T} \mathrm{~g}_{y}\left(x_{t}, y_{t}\right) \quad \forall y \in Y
$$

so that for all $z=(x, y) \in Z$,

$$
\phi\left(x_{t}, y\right)-\phi\left(x, y_{t}\right) \leq\left(x_{t}-x\right)^{T} \mathrm{~g}_{x}\left(x_{t}, y_{t}\right)+\left(y-y_{t}\right)^{T} \mathrm{~g}_{y}\left(x_{t}, y_{t}\right)=\left(z_{t}-z\right)^{T} \mathrm{~g}\left(z_{t}\right)
$$

Using once again the convexity-concavity of $\phi$, we write

$$
\begin{align*}
\epsilon_{\phi}\left(\tilde{z}_{1}^{j}\right) & =\max _{y \in Y} \phi\left(\tilde{x}_{1}^{j}, y\right)-\min _{x \in X} \phi\left(x, \tilde{y}_{1}^{j}\right) \\
& \leq\left[\sum_{t=1}^{j} \gamma_{t}\right]^{-1}\left[\max _{y \in Y} \sum_{t=1}^{j} \gamma_{t} \phi\left(x_{t}, y\right)-\min _{x \in X} \sum_{t=1}^{j} \gamma_{t} \phi\left(x, y_{t}\right)\right]  \tag{3.9}\\
& \leq\left[\sum_{t=1}^{j} \gamma_{t}\right]^{-1} \max _{z \in Z} \sum_{t=1}^{j} \gamma_{t}\left(z_{t}-z\right)^{T} \mathrm{~g}\left(z_{t}\right)
\end{align*}
$$

To bound the right-hand side of (3.9), we use the result of the following lemma (its proof is given in the appendix).

Lemma 3.1. In the above setting, for any $j \geq 1$, the following inequality holds:

$$
\begin{equation*}
\mathbb{E}\left[\max _{z \in Z} \sum_{t=1}^{j} \gamma_{t}\left(z_{t}-z\right)^{T} \mathrm{~g}\left(z_{t}\right)\right] \leq 2+\frac{5}{2} M_{*}^{2} \sum_{t=1}^{j} \gamma_{t}^{2} \tag{3.10}
\end{equation*}
$$

Now to get an error bound for the solution $\tilde{z}_{1}^{j}$, it suffices to substitute inequality (3.10) into (3.9) to obtain

$$
\begin{equation*}
\mathbb{E}\left[\epsilon_{\phi}\left(\tilde{z}_{1}^{j}\right)\right] \leq\left[\sum_{t=1}^{j} \gamma_{t}\right]^{-1}\left[2+\frac{5}{2} M_{*}^{2} \sum_{t=1}^{j} \gamma_{t}^{2}\right] \tag{3.11}
\end{equation*}
$$

Constant stepsizes and basic efficiency estimates. For a fixed number of steps $N$, with the constant stepsize policy

$$
\begin{equation*}
\gamma_{t}=\frac{2 \theta}{M_{*} \sqrt{5 N}}, \quad t=1, \ldots, N \tag{3.12}
\end{equation*}
$$

condition (3.6) and estimate (3.11) imply that

$$
\begin{align*}
\epsilon_{\phi}\left(\tilde{z}_{1}^{N}\right) & \leq 2 \max \left\{\theta, \theta^{-1}\right\} M_{*} \sqrt{\frac{5}{N}} \\
& =2 \max \left\{\theta, \theta^{-1}\right\} \sqrt{\frac{10\left[\alpha_{y} D_{\omega_{x}, X}^{2} M_{*, x}^{2}+\alpha_{x} D_{\omega_{y}, Y}^{2} M_{*, y}^{2}\right]}{\alpha_{x} \alpha_{y} N}} . \tag{3.13}
\end{align*}
$$

Variable stepsizes. Same as in the minimization case, assuming that

$$
\begin{align*}
\bar{D}_{\omega, Z} & :=\sqrt{2} \sup _{z \in Z^{o}, w \in Z}\left[\omega(w)-\omega(z)-(w-z)^{T} \nabla \omega(z)\right]^{1 / 2} \\
& =\sqrt{2}\left[\sup _{z \in Z^{o}, w \in Z} V(z, w)\right]^{1 / 2} \tag{3.14}
\end{align*}
$$

is finite, we can pass from constant stepsizes on a fixed "time horizon" to decreasing stepsize policy

$$
\gamma_{t}=\frac{\theta \bar{D}_{\omega, Z}}{M_{*} \sqrt{t}}, \quad t=1,2, \ldots
$$

(compare with (2.55) and take into account that we are in the situation of $\alpha=1$ ), and from the averaging of all iterates to the "sliding averaging"

$$
\tilde{z}_{i}^{j}=\frac{\sum_{t=i}^{j} \gamma_{t} z_{t}}{\sum_{t=i}^{j} \gamma_{t}}
$$

arriving at the efficiency estimates (compare with (2.56) and (2.57))

$$
\begin{array}{cll}
\epsilon\left(\tilde{z}_{K}^{N}\right) \leq \frac{\bar{D}_{\omega, Z} M_{*}}{\sqrt{N}}\left[\frac{2}{\theta} \frac{N}{N-K+1}+\frac{5 \theta}{2} \sqrt{\frac{N}{K}}\right], & 1 \leq K \leq N,  \tag{3.15}\\
\epsilon\left(\tilde{z}_{\lceil r N\rceil}^{N}\right) \leq C(r) \max \left\{\theta, \theta^{-1}\right\} \frac{\bar{D}_{\omega, Z} M_{*}}{\sqrt{N}}, & r \in(0,1) .
\end{array}
$$

Probabilities of large deviations. Assume that instead of (3.2), the following stronger assumption holds:

$$
\begin{align*}
& \mathbb{E}\left[\exp \left\{\left\|\mathrm{G}_{x}(u, v, \xi)\right\|_{*, x}^{2} / M_{*, x}^{2}\right\}\right] \leq \exp \{1\}  \tag{3.16}\\
& \mathbb{E}\left[\exp \left\{\left\|\mathrm{G}_{y}(x, y, \xi)\right\|_{*, y}^{2} / M_{*, y}^{2}\right\}\right] \leq \exp \{1\}
\end{align*}
$$

Proposition 3.2. In the case of (3.16), with the stepsizes given by (3.12) and (3.6), one has, for any $\Omega>1$,

$$
\begin{equation*}
\operatorname{Prob}\left\{\epsilon_{\phi}\left(\tilde{z}_{1}^{N}\right)>\frac{(8+2 \Omega) \max \left\{\theta, \theta^{-1}\right\} \sqrt{5} M_{*}}{\sqrt{N}}\right\} \leq 2 \exp \{-\Omega\} \tag{3.17}
\end{equation*}
$$

Proof of this proposition is given in the appendix.
3.2. Application to minimax stochastic problems. Consider the following minimax stochastic problem:

$$
\begin{equation*}
\min _{x \in X} \max _{1 \leq i \leq m}\left\{f_{i}(x)=\mathbb{E}\left[F_{i}(x, \xi)\right]\right\} \tag{3.18}
\end{equation*}
$$

where $X \subset \mathbb{R}^{n}$ is a nonempty bounded closed convex set, $\xi$ is a random vector whose probability distribution $P$ is supported on set $\Xi \subset \mathbb{R}^{d}$, and $F_{i}: X \times \Xi \rightarrow \mathbb{R}$, $i=1, \ldots, m$. We assume that the expected value functions $f_{i}(\cdot), i=1, \ldots, m$, are well defined, finite valued, convex, and Lipschitz continuous on $X$. Then the minimax problem (3.18) can be formulated as the following saddle point problem:

$$
\begin{equation*}
\min _{x \in X} \max _{y \in Y}\left\{\phi(x, y)=\sum_{i=1}^{m} y_{i} f_{i}(x)\right\} \tag{3.19}
\end{equation*}
$$

where $Y=\left\{y \in \mathbb{R}^{m}: \sum_{i=1}^{m} y_{i}=1, y \geq 0\right\}$.
Assume that we are able to generate independent realizations $\xi_{1}, \ldots$, of random vector $\xi$, and, for given $x \in X$ and $\xi \in \Xi$, we can compute $F_{i}(x, \xi)$ and its stochastic subgradient $\mathrm{G}_{i}(x, \xi)$ such that $\mathrm{g}_{i}(x)=\mathbb{E}\left[\mathrm{G}_{i}(x, \xi)\right]$ is well defined and $\mathrm{g}_{i}(x) \in \partial f_{i}(x)$,
$x \in X, i=1, \ldots, m$. In other words, we have a stochastic oracle for the problem (3.19) such that assumption $\left(\mathrm{A}^{\prime} 2\right)$ holds, with

$$
\mathrm{G}(x, y, \xi)=\left[\begin{array}{c}
\sum_{i=1}^{m} y_{i} \mathrm{G}_{i}(x, \xi)  \tag{3.20}\\
-\left(F_{1}(x, \xi), \ldots, F_{m}(x, \xi)\right)
\end{array}\right]
$$

and

$$
\mathrm{g}(x, y)=\mathbb{E}[\mathrm{G}(x, y, \xi)]=\left[\begin{array}{c}
\sum_{i=1}^{m} y_{i} \mathrm{~g}_{i}(x)  \tag{3.21}\\
-\left(f_{1}(x), \ldots, f_{m}(x)\right)
\end{array}\right] \in\left[\begin{array}{c}
\partial_{x} \phi(x, y) \\
-\partial_{y} \phi(x, y)
\end{array}\right]
$$

Suppose that the set $X$ is equipped with norm $\|\cdot\|_{x}$, whose dual norm is $\|\cdot\|_{*, x}$, and a distance-generating function $\omega$ modulus $\alpha_{x}$ with respect to $\|\cdot\|_{x}$, and let $R_{x}^{2}=\frac{D_{\omega_{x}, x}^{2}}{\alpha_{x}}$. We equip the set $Y$ with the norm $\|\cdot\|_{y}=\|\cdot\|_{1}$, so that $\|\cdot\|_{*, y}=\|\cdot\|_{\infty}$, and with the distance-generating function

$$
\omega_{y}(y)=\sum_{i=1}^{m} y_{i} \ln y_{i}
$$

and set $R_{y}^{2}=\frac{D_{\omega_{y}, Y}^{2}}{\alpha_{y}}=\ln m$. Next, following (3.3), we set

$$
\|(x, y)\|=\sqrt{\frac{\|x\|_{x}^{2}}{2 R_{x}^{2}}+\frac{\|y\|_{1}^{2}}{2 R_{y}^{2}}}
$$

and hence

$$
\|(\zeta, \eta)\|_{*}=\sqrt{2 R_{x}^{2}\|\zeta\|_{*, x}^{2}+2 R_{y}^{2}\|\eta\|_{\infty}^{2}}
$$

Let us assume uniform bounds:

$$
\max _{1 \leq i \leq m} \mathbb{E}\left[\left\|\mathrm{G}_{i}(x, \xi)\right\|_{*, x}^{2}\right] \leq M_{*, x}^{2}, \quad \mathbb{E}\left[\max _{1 \leq i \leq m}\left|F_{i}(x, \xi)\right|^{2}\right] \leq M_{*, y}^{2}, \quad i=1, \ldots, m
$$

Note that

$$
\mathbb{E}\left[\|\mathrm{G}(x, y, \xi)\|_{*}^{2}\right]=2 R_{x}^{2} \mathbb{E}\left[\left\|\sum_{i=1}^{m} y_{i} \mathrm{G}_{i}(x, \xi)\right\|_{*, x}^{2}\right]+2 R_{y}^{2} \mathbb{E}\left[\|F(x, \xi)\|_{\infty}^{2}\right]
$$

and since $y \in Y$,

$$
\left\|\sum_{i=1}^{m} y_{i} \mathrm{G}_{i}(x, \xi)\right\|_{*, x}^{2} \leq\left(\sum_{i=1}^{m} y_{i}\left\|\mathrm{G}_{i}(x, \xi)\right\|_{*, x}\right)^{2} \leq \sum_{i=1}^{m} y_{i}\left\|\mathrm{G}_{i}(x, \xi)\right\|_{*, x}^{2}
$$

It follows that

$$
\begin{equation*}
\mathbb{E}\left[\|\mathrm{G}(x, y, \xi)\|_{*}^{2}\right] \leq M_{*}^{2}, \tag{3.22}
\end{equation*}
$$

where

$$
M_{*}^{2}=2 R_{x}^{2} M_{*, x}^{2}+2 R_{y}^{2} M_{*, y}^{2}=2 R_{x}^{2} M_{*, x}^{2}+2 M_{*, y}^{2} \ln m
$$

Let us now use the saddle point mirror SA algorithm (3.7)-(3.8) with the constant stepsize policy

$$
\gamma_{t}=\frac{2}{M_{*} \sqrt{5 N}}, \quad t=1,2, \ldots, N
$$

When substituting the value of $M_{*}$, we obtain the following from (3.13):

$$
\begin{align*}
\mathbb{E}\left[\epsilon_{\phi}\left(\tilde{z}_{1}^{N}\right)\right] & =\mathbb{E}\left[\max _{y \in Y} \phi\left(\hat{x}_{1}^{N}, y\right)-\min _{x \in X} \phi\left(x, \hat{y}_{1}^{N}\right)\right]  \tag{3.23}\\
& \leq 2 M_{*} \sqrt{\frac{5}{N}} \leq 2 \sqrt{\frac{10\left[R_{x}^{2} M_{*, x}^{2}+M_{*, x}^{2} \ln m\right]}{N}}
\end{align*}
$$

Discussion. Looking at the bound (3.23), one can make the following important observation. The error of the saddle point mirror SA algorithm in this case is "almost independent" of the number $m$ of constraints (it grows as $O(\sqrt{\ln m})$ as $m$ increases). The interested reader can easily verify that if an E-SA algorithm were used in the same setting (i.e., the algorithm tuned to the norm $\|\cdot\|_{y}=\|\cdot\|_{2}$ ), the corresponding bound would grow with $m$ much faster (in fact, our error bound would be $O(\sqrt{m})$ in that case).

Note that properties of the saddle point mirror SA can be used to reduce significantly the arithmetic cost of the algorithm implementation. To this end let us look at the definition (3.20) of the stochastic oracle: In order to obtain a realization $\mathrm{G}(x, y, \xi)$, one has to compute $m$ random subgradients $\mathrm{G}_{i}(x, \xi), i=1, \ldots, m$, and then their convex combination $\sum_{i=1}^{m} y_{i} \mathrm{G}_{i}(x, \xi)$. Now let $\eta$ be an independent of $\xi$ and uniformly distributed on $[0,1]$ random variable, and let $\imath(\eta, y):[0,1] \times Y \rightarrow\{1, \ldots, m\}$ equal to $i$ when $\sum_{s=1}^{i-1} y_{s}<\eta \leq \sum_{s=1}^{i} y_{s}$. That is, random variable $\hat{\imath}=\imath(\eta, y)$ takes values $1, \ldots, m$ with probabilities $y_{1}, \ldots, y_{m}$. Consider random vector

$$
\mathrm{G}(x, y,(\xi, \eta))=\left[\begin{array}{c}
\mathrm{G}_{\imath(\eta, y)}(x, \xi)  \tag{3.24}\\
-\left(F_{1}(x, \xi), \ldots, F_{m}(x, \xi)\right)
\end{array}\right]
$$

We refer to $\mathrm{G}(x, y,(\xi, \eta))$ as a randomized oracle for problem (3.19), the corresponding random parameter being $(\xi, \eta)$. By construction, we still have $\mathbb{E}[\mathrm{G}(x, y,(\xi, \eta))]=$ $\mathrm{g}(x, y)$, where g is defined in (3.21), and, moreover, the same bound (3.22) holds for $\mathbb{E}\left[\|\mathrm{G}(x, y,(\xi, \eta))\|_{*}^{2}\right]$. We conclude that the accuracy bound (3.23) holds for the error of the saddle point mirror SA algorithm with randomized oracle. On the other hand, in the latter procedure only one randomized subgradient $\mathrm{G}_{\hat{\imath}}(x, \xi)$ per iteration is computed. This simple idea is further developed in another interesting application of the saddle point mirror SA algorithm to bilinear matrix games, which we discuss next.
3.3. Application to bilinear matrix games. Consider the standard matrix game problem, that is, problem (3.1) with

$$
\phi(x, y)=y^{T} A x+b^{T} x+c^{T} y
$$

where $A \in \mathbb{R}^{m \times n}$, and $X$ and $Y$ are the standard simplexes:

$$
X=\left\{x \in \mathbb{R}^{n}: x \geq 0, \sum_{j=1}^{n} x_{j}=1\right\}, Y=\left\{y \in \mathbb{R}^{m}: y \geq 0, \sum_{i=1}^{m} y_{i}=1\right\}
$$

In the case in question it is natural to equip $X$ (respectively, $Y$ ) with the $\|\cdot\|_{1}$-norm on $\mathbb{R}^{n}$ (respectively, $\mathbb{R}^{m}$ ). We choose entropies as the corresponding distance-generating functions:

$$
\omega_{x}(x)=\sum_{i=1}^{n} x_{i} \ln x_{i}, \quad \omega_{y}(x)=\sum_{i=1}^{m} y_{i} \ln y_{i} \quad\left[\Rightarrow \frac{D_{\omega_{x}, X}^{2}}{\alpha_{x}}=\ln n, \frac{D_{\omega_{y}, Y}^{2}}{\alpha_{y}}=\ln m\right]
$$

According to (3.3), we set

$$
\begin{equation*}
\|(x, y)\|=\sqrt{\frac{\|x\|_{1}^{2}}{2 \ln n}+\frac{\|y\|_{1}^{2}}{2 \ln m}} \Rightarrow\|(\zeta, \eta)\|_{*}=\sqrt{2\|\zeta\|_{\infty}^{2} \ln n+2\|\eta\|_{\infty}^{2} \ln m} \tag{3.25}
\end{equation*}
$$

In order to compute the estimates $\mathrm{G}(x, y, \xi)$ of $\mathrm{g}(x, y)=\left(b+A^{T} y,-c-A x\right)$, to be used in the saddle point mirror SA iterations (3.7), we use the randomized oracle

$$
\mathrm{G}(x, y, \xi)=\left[\begin{array}{c}
c+A^{\imath\left(\xi_{1}, y\right)}  \tag{3.26}\\
-b-A_{\imath\left(\xi_{2}, x\right)}
\end{array}\right]
$$

where $\xi_{1}$ and $\xi_{2}$ are independent uniformly distributed on $[0,1]$ random variables and $\hat{j}=\imath\left(\xi_{1}, y\right), \hat{i}=\imath\left(\xi_{2}, x\right)$ are defined as in (3.24) (i.e., $\hat{j}$ can take values $1, \ldots, m$, with probabilities $y_{1}, \ldots, y_{m}$ and $\hat{i}$ can take values $1, \ldots, n$, with probabilities $x_{1}, \ldots, x_{n}$ ), and $A_{j},\left[A^{i}\right]^{T}$ are $j$ th column and $i$ th row in $A$, respectively.

Note that

$$
\mathrm{g}(x, y) \equiv \mathbb{E}[\mathrm{G}(x, y,(\hat{j}, \hat{i}))] \in\left[\begin{array}{c}
\partial_{x} \phi(x, y)  \tag{3.27}\\
\partial_{y}(-\phi(x, y))
\end{array}\right]
$$

Besides this,

$$
\begin{aligned}
\left|\mathrm{G}(x, y, \xi)_{i}\right| \leq \max _{1 \leq j \leq m}\left\|A^{j}+b\right\|_{\infty}, & 1 \leq i \leq n \\
\left|\mathrm{G}(x, y, \xi)_{i}\right| \leq \max _{1 \leq j \leq n}\left\|A_{j}+c\right\|_{\infty}, & n+1 \leq i \leq n+m,
\end{aligned}
$$

whence, invoking (3.25), for any $x \in X, y \in Y$, and $\xi$,

$$
\begin{equation*}
\|\mathrm{G}(x, y, \xi)\|_{*}^{2} \leq M_{*}^{2}=2 \ln n \max _{1 \leq j \leq m}\left\|A^{j}+b\right\|_{\infty}^{2}+2 \ln m \max _{1 \leq j \leq n}\left\|A_{j}+c\right\|_{\infty}^{2} \tag{3.28}
\end{equation*}
$$

The bottom line is that our stochastic gradients along with the just-defined $M_{*}$ satisfy both ( $\mathrm{A}^{\prime} 2$ ) and (3.16), and therefore with the constant stepsize policy (3.12), we have

$$
\begin{equation*}
\mathbb{E}\left[\epsilon_{\phi}\left(\tilde{z}_{1}^{N}\right)\right]=\mathbb{E}\left[\max _{y \in Y} \phi\left(\tilde{x}_{1}^{N}, y\right)-\min _{x \in X} \phi\left(x, \tilde{y}_{1}^{N}\right)\right] \leq 2 M_{*} \sqrt{\frac{5}{N}} \tag{3.29}
\end{equation*}
$$

(cf. (3.13)). In our present situation, Proposition 3.2 in a slightly refined form (for proof, see the appendix) reads as follows.

Proposition 3.3. With the constant stepsize policy (3.12), for the just-defined algorithm, one has for any $\Omega \geq 1$, that

$$
\begin{equation*}
\operatorname{Prob}\left\{\epsilon_{\phi}\left(\tilde{z}_{1}^{N}\right)>2 M_{*} \sqrt{\frac{5}{N}}+\frac{4 \bar{M}}{\sqrt{N}} \Omega\right\} \leq \exp \left\{-\Omega^{2} / 2\right\} \tag{3.30}
\end{equation*}
$$

where

$$
\begin{equation*}
\bar{M}=\max _{1 \leq j \leq m}\left\|A^{j}+b\right\|_{\infty}+\max _{1 \leq j \leq n}\left\|A_{j}+c\right\|_{\infty} \tag{3.31}
\end{equation*}
$$

Discussion. Consider a bilinear matrix game with $n \geq m, \ln (m)=O(1) \ln (n)$, and $b=c=0$ (so that $M_{*}=O(1) \sqrt{\ln n} \bar{M}$ and $\bar{M}=\max _{i, j}\left|A_{i j}\right|$; see (3.28), (3.31)). Suppose that we are interested to solve it within a fixed relative accuracy $\rho$, that is, to ensure that the (perhaps random) approximate solution $\tilde{z}_{1}^{N}$, which we get after $N$ iterations, satisfies the error bound

$$
\epsilon_{\phi}\left(\tilde{z}_{N}\right) \leq \rho \max _{1 \leq i, j \leq n}\left|A_{i j}\right|
$$

with probability at least $1-\delta$. According to (3.30), to this end, one can use the randomized saddle point mirror SA algorithm (3.7), (3.8), (3.26) with stepsizes (3.12), (3.28) and with

$$
\begin{equation*}
N=O(1) \frac{\ln n+\ln (1 / \delta)}{\rho^{2}} \tag{3.32}
\end{equation*}
$$

The computational cost of building $\tilde{z}_{1}^{N}$ with this approach is

$$
\mathcal{C}(\rho)=O(1) \frac{[\ln n+\ln (1 / \delta)][\mathcal{R}+n]}{\rho^{2}}
$$

arithmetic operations, where $\mathcal{R}$ is the arithmetic cost of extracting a column/row from $A$ given the index of this column/row. The total number of rows and columns visited by the algorithm does not exceed the number of steps $N$ as given in (3.32) so that the total number of entries in $A$ used in the course of the entire computation does not exceed

$$
M=O(1) \frac{n(\ln n+\ln (1 / \delta))}{\rho^{2}}
$$

When $\rho$ is fixed, $m=O(1) n$ and $n$ is large, $M$ is incomparably less than the total number $m n$ of entries in $A$. Thus, our algorithm exhibits sublinear-time behavior: it produces reliable solutions of prescribed quality to large-scale matrix games by inspecting a negligible, as $n \rightarrow \infty$, part of randomly selected data. Note that randomization here is critical. ${ }^{3}$ It can be seen that a deterministic algorithm, which is capable to find a solution with (deterministic) relative accuracy $\rho \leq 0.1$, has to "see" in the worst case at least $O(1) n$ rows/columns of $A$.
4. Numerical results. In this section, we report the results of our computational experiments where we compare the performance of the robust mirror descent SA method and the SAA method applied to three stochastic programming problems, namely: a stochastic utility problem, a stochastic max-flow problem, and a network planning problem with random demand. We also present a small simulation study of the performance of randomized mirror SA algorithm for bilinear matrix games.

The algorithms we were testing are the two variants of the robust mirror descent SA. The first variant, the E-SA, is as described in section 2.2; in terms of section 2.3, this is nothing but mirror descent robust SA with Euclidean setup; see the example in section 2.3. The second variant, referred to as the non-Euclidean SA (N-SA), is the mirror descent robust SA with $\ell_{1}$-setup; see, the example in section 2.3.

[^3]Table 4.1
Selecting stepsize policy.

| [method: N-SA, N:2,000, K:10,000, instance: L1] |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: |
|  | $\theta$ |  |  |  |
|  | 0.1 | 1 | 5 | 10 |
|  | -7.4733 | -7.8865 | -7.8789 | -7.8547 |
|  | -6.9371 | -7.8637 | -7.9037 | -7.8971 |

These two variants of the SA method are compared with the SAA approach in the following way: fixing an iid. sample (of size $N$ ) for the random variable $\xi$, we apply the three aforementioned methods to obtain approximate solutions for the test problem under consideration, and then the quality of the solutions yielded by these algorithms is evaluated using another iid. sample of size $K \gg N$. It should be noted that SAA itself is not an algorithm, and in our experiments, it was coupled with the non-Euclidean restricted memory level (NERML) [2]-a powerful deterministic algorithm for solving the sample average problem (1.4).

### 4.1. Preliminaries.

Algorithmic schemes. Both E-SA and N-SA were implemented according to the description in section 2.3, the number of steps $N$ being the parameter of a particular experiment. In such an experiment, we generated $\approx \log _{2} N$ candidate solutions $\tilde{x}_{i}^{N}$, with $N-i+1=\min \left[2^{k}, N\right], k=0,1, \ldots,\left\lceil\log _{2} N\right\rceil$. We then used an additional sample to estimate the objective at these candidate solutions in order to choose the best of these candidates, specifically, as follows: we used a relatively short sample to choose the two "most promising" of the candidate solutions, and then a large sample (of size $K \gg N)$ to identify the best of these two candidates, thus getting the "final" solution. The computational effort required by this simple postprocessing is not reflected in the tables to follow.

The stepsizes. At the "pilot stage" of our experimentation, we made a decision on which stepsize policy-(2.47) or (2.55) - to choose and how to identify the underlying parameters $M_{*}$ and $\theta$. In all our experiments, $M_{*}$ was estimated by taking the maxima of $\|\mathrm{G}(\cdot, \cdot)\|_{*}$ over a small (just 100) calls to the stochastic oracle at randomly generated feasible solutions. As about the value of $\theta$ and type of the stepsize policy $((2.47)$ or $(2.55))$, our choice was based on the results of experimentation with a single test problem (instance L1 of the utility problem, see below); some results of this experimentation are presented in Table 4.1. We have found that the constant stepsize policy (2.47) with $\theta=0.1$ for the E-SA and $\theta=5$ for the N-SA slightly outperforms other variants we have considered. This particular policy, combined with the aforementioned scheme for estimating $M_{*}$, was used in all subsequent experiments.

Format of test problems. All our test problems are of the form $\min _{x \in X} f(x)$, $f(x)=\mathbb{E}[F(x, \xi)]$, where the domain $X$ either is a standard simplex $\left\{x \in \mathbb{R}^{n}: x \geq\right.$ $\left.0, \sum_{i} x_{i}=1\right\}$ or can be converted into such a simplex by scaling of the original variables.

Notation in the tables. Below,

- $n$ is the design dimension of an instance,
- $N$ is the sample size (i.e., the number of steps in SA, and the size of the sample used to build the stochastic average in SAA),
- Obj is the empirical mean of the random variable $F(x, \xi), x$ being the approximate solution generated by the algorithm in question. The empirical means are taken over a large ( $K=10^{4}$ elements) dedicated sample,
- CPU is the $C P U$ time in seconds.

TABLE 4.2
$S A$ versus $S A A$ on the stochastic utility problem.

| - |  | $\mathrm{L} 1: n=500$ |  | $\mathrm{~L} 2: n=1,000$ |  | $\mathrm{~L} 3: n=2,000$ |  | $\mathrm{~L} 4: n=5,000$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ALG. | $N$ | Obj | CPU | Obj | CPU | Obj | CPU | Obj | CPU |
| N-SA | 100 | -7.7599 | 0 | -5.8340 | 0 | -7.1419 | 1 | -5.4688 | 3 |
|  | 1,000 | -7.8781 | 2 | -5.9152 | 2 | -7.2312 | 6 | -5.5716 | 13 |
|  | 2,000 | -7.8987 | 2 | -5.9243 | 5 | -7.2513 | 10 | -5.5847 | 25 |
|  | 4,000 | -7.9075 | 5 | -5.9365 | 12 | -7.2595 | 20 | -5.5935 | 49 |
| E-SA | 100 | -7.6895 | 0 | -5.7988 | 1 | -7.0165 | 1 | -4.9364 | 4 |
|  | 1,000 | -7.8559 | 2 | -5.8919 | 4 | -7.2029 | 7 | -5.3895 | 20 |
|  | 2,000 | -7.8737 | 3 | -5.9067 | 7 | -7.2306 | 15 | -5.4870 | 39 |
|  | 4,000 | -7.8948 | 7 | -5.9193 | 13 | -7.2441 | 29 | -5.5354 | 77 |
| SAA | 100 | -7.6571 | 7 | -5.6346 | 8 | -6.9748 | 19 | -5.3360 | 44 |
|  | 1,000 | -7.8821 | 31 | -5.9221 | 68 | -7.2393 | 134 | -5.5656 | 337 |
|  | 2,000 | -7.9100 | 72 | -5.9313 | 128 | -7.2583 | 261 | -5.5878 | 656 |
|  | 4,000 | -7.9087 | 113 | -5.9384 | 253 | -7.2664 | 515 | -5.5967 | 1,283 |

Table 4.3
The variability for the stochastic utility problem.

| - | N-SA |  |  | E-SA |  |  | SAA |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | Obj |  | CPU | Obj |  | CPU | Obj |  |  | CPU |
| Inst | $N$ | Mean | Dev | (Avg.) | Mean | Dev | (Avg.) | Mean | Dev |  |  |
| (Avg.) |  |  |  |  |  |  |  |  |  |  |  |
| L2 | 1,000 | -5.9159 | 0.0025 | 2.63 | -5.8925 | 0.0024 | 4.99 | -5.9219 | 0.0047 |  |  |
| L2 | 2,000 | -5.9258 | 0.0022 | 5.03 | -5.9063 | 0.0019 | 7.09 | -5.9328 | 0.0028 |  |  |
| 131.25 |  |  |  |  |  |  |  |  |  |  |  |

4.2. A stochastic utility problem. Our first experiment was carried out with the utility model

$$
\begin{equation*}
\min _{x \in X}\left\{f(x)=\mathbb{E}\left[\phi\left(\sum_{i=1}^{n}\left(i / n+\xi_{i}\right) x_{i}\right)\right]\right\} \tag{4.1}
\end{equation*}
$$

where $X=\left\{x \in \mathbb{R}^{n}: x \geq 0, \sum_{i=1}^{n} x_{i}=1\right\}, \xi_{i} \sim \mathcal{N}(0,1)$ are independent and $\phi(\cdot)$ is a piecewise linear convex function given by $\phi(t)=\max \left\{v_{1}+s_{1} t, \ldots, v_{m}+s_{m} t\right\}$, where $v_{k}$ and $s_{k}$ are certain constants. In our experiment, we used $m=10$ breakpoints, all located on $[0,1]$. The four instances L1, L2, L3, L4 we dealt with were of dimension varying from 500 to 2,000 , each instance - with its own randomly generated function $\phi$. All the algorithms were coded in ANSI C, and the experiments were conducted on an Intel PIV 1.6GHz machine with Microsoft windows XP professional.

We run each of the three aforementioned methods with various sample sizes on every one of the instances. The results are reported in Table 4.2.

In order to evaluate stability of the algorithms, we run each of them 100 times; the resulting statistics are shown in Table 4.3. In this relatively time-consuming experiment, we restrict ourselves with a single instance (L2) and just two sample sizes $(N=1,000$ and 2,000). In Table 4.3, "Mean" and "Dev" are, respectively, the mean and the deviation, over 100 runs, of the objective value Obj at the resulting approximate solution.

The experiments demonstrate that as far as the quality of approximate solutions is concerned, N-SA outperforms E-SA and is almost as good as SAA. At the same time, the solution time for $\mathrm{N}-\mathrm{SA}$ is significantly smaller than the one for SAA.
4.3. Stochastic max-flow problem. In the second experiment, we consider simple two-stage stochastic linear programming, namely, a stochastic max-flow problem. The problem is to optimize the capacity expansion of a stochastic network. Let

TABLE 4.4
$S A$ versus $S A A$ on the stochastic max-flow problem.

| - <br> $(m, n)$ |  | F1 <br> $(50,500)$ |  | F2 <br> $(100,1,000)$ |  | F3 <br> $(100,2,000)$ |  | F4 <br> $(250,5,000)$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ALG. | $N$ | Obj | CPU | Obj | CPU | Obj | CPU | Obj | CPU |
| N-SA | 100 | 0.1140 | 0 | 0.0637 | 0 | 0.1296 | 1 | 0.1278 | 3 |
|  | 1,000 | 0.1254 | 1 | 0.0686 | 3 | 0.1305 | 6 | 0.1329 | 15 |
|  | 2,000 | 0.1249 | 3 | 0.0697 | 6 | 0.1318 | 11 | 0.1338 | 29 |
|  | 4,000 | 0.1246 | 5 | 0.0698 | 11 | 0.1331 | 21 | 0.1334 | 56 |
| E-SA | 100 | 0.0840 | 0 | 0.0618 | 1 | 0.1277 | 2 | 0.1153 | 7 |
|  | 1,000 | 0.1253 | 3 | 0.0670 | 6 | 0.1281 | 16 | 0.1312 | 39 |
|  | 2,000 | 0.1246 | 5 | 0.0695 | 13 | 0.1287 | 28 | 0.1312 | 72 |
|  | 4,000 | 0.1247 | 9 | 0.0696 | 24 | 0.1303 | 53 | 0.1310 | 127 |
| SAA | 100 | 0.1212 | 5 | 0.0653 | 12 | 0.1310 | 20 | 0.1253 | 60 |
|  | 1,000 | 0.1223 | 35 | 0.0694 | 84 | 0.1294 | 157 | 0.1291 | 466 |
|  | 2,000 | 0.1223 | 70 | 0.0693 | 170 | 0.1304 | 311 | 0.1284 | 986 |
|  | 4,000 | 0.1221 | 140 | 0.0693 | 323 | 0.1301 | 636 | 0.1293 | 1,885 |

$G=(N, A)$ be a diagraph with a source node $s$ and a sink node $t$. Each $\operatorname{arc}(i, j) \in A$ has an existing capacity $p_{i j} \geq 0$ and a random implementing/operating level $\xi_{i j}$. Moreover, there is a common random degrading factor $\eta$ for all arcs in $A$. The goal is to determine how much capacity to add to the arcs, subject to a budget constraint, in order to maximize the expected maximum flow from $s$ to $t$. Denoting by $x_{i j}$ the capacity to be added to arc $(i, j)$, the problem reads

$$
\begin{equation*}
\max _{x}\left\{f(x)=\mathbb{E}[F(x ; \xi, \eta)]: \sum_{(i, j) \in A} c_{i j} x_{i j} \leq b, x_{i j} \geq 0, \quad \forall(i, j) \in A\right\} \tag{4.2}
\end{equation*}
$$

where $c_{i j}$ is the per unit cost for the capacity to be added, $b$ is the total available budget, and $F(x ; \xi, \eta)$ denotes the maximum $s-t$ flow in the network when the capacity of an arc $(i, j)$ is $\eta \xi_{i j}\left(p_{i j}+x_{i j}\right)$. Note that the above is a maximization rather than a minimization problem.

We assume that the random variables $\xi_{i j}, \theta$ are independent and uniformly distributed on $[0,1]$ and $[0.5,1]$, respectively, and consider the case of $p_{i j}=0, c_{i j}=1$ for all $(i, j) \in E$, and $b=1$. We randomly generated 4 network instances (referred to as F1, F2, F3, and F4) using the network generator GRIDGEN available on DIMACS challenge. The push-relabel algorithm [8] was used to solve the second stage max-flow problem.

In the first test, each algorithm (N-SA, E-SA, SAA) was run once at each test instance; the results are reported in Table 4.4, where $m, n$ stand for the number of nodes, respectively, arcs in $G$. Similar to the stochastic utility problem, we investigate the stability of the methods by running each of them 100 times. The resulting statistics is presented in Table 4.5, whose columns have exactly the same meaning as in Table 4.3.

This experiment fully supports the conclusions on the methods suggested by the experiments with the utility problem.
4.4. A network planning problem with random demand. In the last experiment, we consider the so-called SSN problem of Sen, Doverspike, and Cosares [24]. This problem arises in telecommunications network design where the owner of the network sells private-line services between pairs of nodes in the network, and the demands are treated as random variables based on the historical demand patterns.

TABLE 4.5
The variability for the stochastic max-flow problem.

| - |  | N-SA |  |  | E-SA |  |  | SAA |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | Obj |  | Avg. | Obj |  |  | Avg. | Obj |  | Avg. |
| Inst | $N$ | Mean | Dev | CPU | Mean | Dev | CPU | Mean | Dev | CPU |  |
| F2 | 1,000 | 0.0691 | 0.0004 | 3.11 | 0.0688 | 0.0006 | 4.62 | 0.0694 | 0.0003 | 90.15 |  |
| F2 | 2,000 | 0.0694 | 0.0003 | 6.07 | 0.0692 | 0.0002 | 6.91 | 0.0695 | 0.0003 | 170.45 |  |

The optimization problem is to decide where to add capacity to the network to minimize the expected rate of unsatisfied demands. Since this problem has been studied by several authors (see, e.g., $[12,24]$ ), it could be interesting to compare the results. Another purpose of this experiment is to investigate the behavior of the SA method when the Latin hyperplane sampling (LHS) variance reduction technique (introduced in [14]) is applied.

The problem has been formulated as a two-stage stochastic linear programming as follows:

$$
\begin{equation*}
\min _{x}\left\{f(x)=\mathbb{E}[F(x, \xi)]: x \geq 0, \sum_{i} x_{i}=b\right\} \tag{4.3}
\end{equation*}
$$

where $x$ is the vector of capacities to be added to the arcs of the network, $b$ (the budget) is the total amount of capacity to be added, $\xi$ denotes the random demand, and $F(x, \xi)$ represents the number of unserved requests, specifically,

$$
F(x, \xi)=\min _{s, f}\left\{\begin{array}{ll} 
& \sum_{i} \sum_{r \in R(i)} A_{r} f_{i r} \leq x+c  \tag{4.4}\\
\sum_{i} s_{i}: & \sum_{r \in R(i)} f_{i r}+s_{i}=\xi^{i}, \quad \forall i \\
f_{i r} \geq 0, s_{i} \geq 0, \quad \forall i, r \in R(i)
\end{array}\right\}
$$

Here,

- $R(i)$ is the set of routes used for traffic $i$ (traffic between the source-sink pair of nodes \# i),
- $\xi^{i}$ is the (random) demand for traffic $i$,
- $A_{r}$ are the route-arc incidence vectors (so that $j$ th component of $A_{r}$ is 1 or 0 depending on whether arc $j$ belongs to the route $r$ ),
- $c$ is the vector of current capacities, $f_{i r}$ is the fraction of traffic $i$ transferred via route $r$, and $s$ is the vector of unsatisfied demands.

In the SSN instance, there are $\operatorname{dim} x=89 \operatorname{arcs}$ and $\operatorname{dim} \xi=86$ source-sink pairs, and components of $\xi$ are independent random variables with known discrete distributions (from 3 to 7 possible values per component), which result in $\approx 10^{70}$ possible demand scenarios.

In the first test with the SSN instance, each of our 3 algorithms was run once without and once with the LHS technique; the results are reported in Table 4.6. We then tested the stability of algorithms by running each of them 100 times; see statistics in Table 4.7. Note that experiments with the SSN problem were conducted on a more powerful computer: Intel Xeon 1.86 GHz with Red Hat Enterprize Linux.

As far as comparison of our three algorithms is concerned, the conclusions are in full agreement with those for the utility and the max-flow problem. We also see that for our particular example, the LHS does not yield much of an improvement, especially when a larger sample size is applied. This result seems to be consistent with the observation in [12].

Table 4.6
$S A$ versus $S A A$ on the $S S N$ problem.

| - |  | Without LHS |  | With LHS |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Alg. | $N$ | Obj | CPU | Obj | CPU |
| N-SA | 100 | 11.0984 | 1 | 10.1024 | 1 |
|  | 1,000 | 10.0821 | 6 | 10.0313 | 7 |
|  | 2,000 | 9.9812 | 12 | 9.9936 | 12 |
|  | 4,000 | 9.9151 | 23 | 9.9428 | 22 |
| E-SA | 100 | 10.9027 | 1 | 10.3860 | 1 |
|  | 1,000 | 10.1268 | 6 | 10.0984 | 6 |
|  | 2,000 | 10.0304 | 12 | 10.0552 | 12 |
|  | 4,000 | 9.9662 | 23 | 9.9862 | 23 |
| SAA | 100 | 11.8915 | 24 | 11.0561 | 23 |
|  | 1,000 | 10.0939 | 215 | 10.0488 | 216 |
|  | 2,000 | 9.9769 | 431 | 9.9872 | 426 |
|  | 4,000 | 9.8773 | 849 | 9.9051 | 853 |

TABLE 4.7
The variability for the SSN problem.

| - | N-SA |  |  | E-SA |  |  | SAA |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | Obj |  | Avg. | Obj |  | Avg. | Obj |  | Avg. |
| $N$ | LHS | Mean | Dev | CPU | Mean | Dev | CPU | Mean | Dev | CPU |
| 1,000 | no | 10.0624 | 0.1867 | 6.03 | 10.1730 | 0.1826 | 6.12 | 10.1460 | 0.2825 | 215.06 |
| 1,000 | yes | 10.0573 | 0.1830 | 6.16 | 10.1237 | 0.1867 | 6.14 | 10.0135 | 0.2579 | 216.10 |
| 2,000 | no | 9.9965 | 0.2058 | 11.61 | 10.0853 | 0.1887 | 11.68 | 9.9943 | 0.2038 | 432.93 |
| 2,000 | yes | 9.9978 | 0.2579 | 11.71 | 10.0486 | 0.2066 | 11.74 | 9.9830 | 0.1872 | 436.94 |

4.5. N-SA versus E-SA. The data in Tables 4.3, 4.4, and 4.6 demonstrate that with the same sample size $N$, the N-SA somehow outperforms the E-SA in terms of both the quality of approximate solutions and the running time. ${ }^{4}$ The difference in solutions' quality, at the first glance, seems slim, and one could think that adjusting the SA algorithm to the "geometry" of the problem in question (in our case, to minimization over a standard simplex) is of minor importance. We, however, do believe that such a conclusion would be wrong. In order to get a better insight, let us come back to the stochastic utility problem. This test problem has an important advantage - we can easily compute the value of the objective $f(x)$ at a given candidate solution $x$ analytically. ${ }^{5}$ Moreover, it is easy to minimize $f(x)$ over the simplex-on a closest inspection, this problem reduces to minimizing an easy-to-compute univariate convex function so that we can approximate the true optimal value $f_{*}$ to high accuracy by bisection. Thus, in the case in question, we can compare solutions $x$ generated by various algorithms in terms of their "true inaccuracy" $f(x)-f_{*}$, and this is the rationale behind our "Gaussian setup." We can now exploit this advantage of the stochastic utility problem for comparing properly N-SA and E-SA. In Table 4.8, we present the true values of the objective $f(\bar{x})$ at the approximate solutions $\bar{x}$ generated by N-SA and E-SA as applied to the instances L1 and L4 of the utility problem (cf. Table 4.3) along with the inaccuracies $f(\bar{x})-f_{*}$ and the Monte Carlo estimates $\widehat{f}(\bar{x})$ of $f(\bar{x})$ obtained via 50,000 -element samples. We see that the difference in

[^4]TABLE 4.8
$N-S A$ versus $E-S A$.

| Method | Problem | $\widehat{f}(\bar{x}), f(\bar{x})$ | $f(\bar{x})-f_{*}$ | Time |
| :--- | :---: | :---: | :---: | :---: |
| N-SA, $N=2,000$ | L2: $n=1,000$ | $-5.9232 /-5.9326$ | 0.0113 | 5.00 |
| E-SA, $N=2,000$ | L2 | $-5.8796 /-5.8864$ | 0.0575 | 6.60 |
| E-SA, $N=10,000$ | L2 | $-5.9059 /-5.9058$ | 0.0381 | 39.80 |
| E-SA, $N=20,000$ | L2 | $-5.9151 /-5.9158$ | 0.0281 | 74.50 |
| N-SA, $N=2,000$ | L4: $n=5,000$ | $-5.5855 /-5.5867$ | 0.0199 | 25.00 |
| E-SA, $N=2,000$ | L4 | $-5.5467 /-5.5469$ | 0.0597 | 44.60 |
| E-SA, $N=10,000$ | L4 | $-5.5810 /-5.5812$ | 0.0254 | 165.10 |
| E-SA, $N=20,000$ | L4 | $-5.5901 /-5.5902$ | 0.0164 | 382.00 |

the inaccuracy $f(\bar{x})-f_{*}$ of the solutions produced by the algorithms is much more significant than is suggested by the data in Table 4.3 (where the actual inaccuracy is "obscured" by the estimation error and summation with $f_{*}$ ). Specifically, at the common for both algorithm sample sizes $N=2,000$, the inaccuracy yielded by N-SA is $3-5$ times less than the one for E-SA and in order to compensate for this difference, one should increase the sample size for E-SA (and hence the running time) by factor $5-10$. It should be added that in light of theoretical complexity analysis carried out in Example 2.3, the outlined significant difference in performances of N-SA and E-SA is not surprising; the surprising fact is that E-SA works at all.
4.6. Bilinear matrix game. We consider here a bilinear matrix game

$$
\min _{x \in X} \max _{y \in Y} y^{T} A x
$$

where both feasible sets are the standard simplexes in $\mathbb{R}^{n}: Y=X=\left\{x \in \mathbb{R}^{n}\right.$ : $\left.\sum_{i=1}^{n} x_{i}=1, x \geq 0\right\}$. We consider two versions of the randomized mirror SA algorithm (3.7), (3.8) for the saddle point problem. The first algorithm, the E-SA, uses $\frac{1}{2}\|x\|_{2}^{2}$ as $\omega_{x}, \omega_{y}$ and $\|\cdot\|_{2}$ as $\|\cdot\|_{x},\|\cdot\|_{y}$. The second algorithm, the N-SA, uses the entropy function (2.58) as $\omega_{x}, \omega_{y}$ and the norm $\|\cdot\|_{1}$ as $\|\cdot\|_{x},\|\cdot\|_{y}$. To compare the two procedures, we compute the corresponding approximate solutions $\tilde{z}_{1}^{N}$ and compute the exact values of the error:

$$
\epsilon\left(\tilde{z}_{1}^{N}\right)=\max _{y \in Y} y^{T} A \tilde{x}_{1}^{N}-\min _{x \in X}\left[\tilde{y}_{1}^{N}\right]^{T} A x, \quad i=1,2 .
$$

In our experiments we consider symmetric matrices $A$ of two kinds. The matrices of the first family, parameterized by $\alpha>0$, are given by

$$
A_{i j}=\left(\frac{i+j-1}{2 n-1}\right)^{\alpha}, \quad 1 \leq i, j \leq n
$$

The second family of matrices, which are also parameterized by $\alpha>0$, is given by

$$
A_{i j}=\left(\frac{|i-j|+1}{2 n-1}\right)^{\alpha}, \quad 1 \leq i, j \leq n
$$

We use the notations $E_{1}(\alpha)$ and $E_{2}(\alpha)$ to refer to the experiments with the matrices of the first and second kind with parameter $\alpha$. We present in Table 4.9 the results of experiments conducted for the matrices $A$ of size $10^{4} \times 10^{4}$. We made 100 simulation runs in each experiment and present the average error (column Mean), standard

Table 4.9
SA for bilinear matrix games.

|  | $E_{2}(2), \epsilon\left(\tilde{z}_{1}\right)=0.500$ |  |  | $E_{2}(1), \quad \epsilon\left(\tilde{z}_{1}\right)=0.500$ |  |  | $E_{2}(0.5), \quad \epsilon\left(\tilde{z}_{1}\right)=0.390$ |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{gathered} \mathrm{N}-\mathrm{SA} \\ N \end{gathered}$ | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | $\begin{aligned} & \text { CPU } \\ & \text { CPU } \end{aligned}$ | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | $\begin{aligned} & \hline \text { CPU } \\ & \text { CPU } \end{aligned}$ | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | CPU |
|  | Mean | Dev |  | Mean | Dev |  | Mean | Dev | CPU |
| 100 | 0.0121 | $3.9 \mathrm{e}-4$ | 0.58 | 0.0127 | $1.9 \mathrm{e}-4$ | 0.69 | 0.0122 | $4.3 \mathrm{e}-4$ | 0.81 |
| 1,000 | 0.00228 | $3.7 \mathrm{e}-5$ | 5.8 | 0.00257 | $2.2 \mathrm{e}-5$ | 7.3 | 0.00271 | $4.5 \mathrm{e}-5$ | 8.5 |
| 2,000 | 0.00145 | $2.1 \mathrm{e}-5$ | 11.6 | 0.00166 | $1.0 \mathrm{e}-5$ | 13.8 | 0.00179 | $2.7 \mathrm{e}-5$ | 16.4 |
| E-SA | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | $\begin{gathered} \text { CPU } \\ \text { (Avg.) } \end{gathered}$ | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | $\begin{gathered} \text { CPU } \\ \text { (Avg.) } \end{gathered}$ | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | CPU |
| $N$ | Mean | DEV |  | Mean | Dev |  | Mean | Dev | (Avg.) |
| 100 | 0.00952 | $1.0 \mathrm{e}-4$ | 1.27 | 0.0102 | $5.1 \mathrm{e}-5$ | 1.77 | 0.00891 | $1.1 \mathrm{e}-4$ | 1.94 |
| 1,000 | 0.00274 | $1.3 \mathrm{e}-5$ | 11.3 | 0.00328 | 7.8 e-6 | 17.6 | 0.00309 | $1.6 \mathrm{e}-5$ | 20.9 |
| 2,000 | 0.00210 | $7.4 \mathrm{e}-6$ | 39.7 | 0.00256 | $4.6 \mathrm{e}-6$ | 36.7 | 0.00245 | $7.8 \mathrm{e}-6$ | 39.2 |
|  | $E_{1}(2), \epsilon\left(\tilde{z}_{1}\right)=0.0625$ |  |  | $E_{1}(1), \epsilon\left(\tilde{z}_{1}\right)=0.125$ |  |  | $E_{1}(0.5), \epsilon\left(\tilde{z}_{1}\right)=0.138$ |  |  |
| N-SA | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | $\begin{gathered} \text { CPU } \\ (\text { Avg. }) \end{gathered}$ | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | $\begin{gathered} \text { CPU } \\ (\text { Avg. }) \end{gathered}$ | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | CPU |
| $N$ | Mean | Dev |  | Mean | Dev |  | Mean | Dev | (Avg.) |
| 100 | 0.00817 | 0.0016 | 0.58 | 0.0368 | 0.0068 | 0.66 | 0.0529 | 0.0091 | 0.78 |
| 1,000 | 0.00130 | $2.7 \mathrm{e}-4$ | 6.2 | 0.0115 | 0.0024 | 6.5 | 0.0191 | 0.0033 | 7.6 |
| 2,000 | 0.00076 | $1.6 \mathrm{e}-4$ | 11.4 | 0.00840 | 0.0014 | 11.7 | 0.0136 | 0.0018 | 13.8 |
| E-SA | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | $\begin{gathered} \text { CPU } \\ \text { (Avg.) } \end{gathered}$ | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | CPU | $\epsilon\left(\tilde{z}_{1}^{N}\right)$ |  | CPU |
| $N$ | Mean | Dev |  | Mean | Dev | (Avg.) | Mean | Dev | (Avg.) |
| 100 | 0.00768 | 0.0012 | 1.75 | 0.0377 | 0.0062 | 2.05 | 0.0546 | 0.0064 | 2.74 |
| 1,000 | 0.00127 | $2.2 \mathrm{e}-4$ | 17.2 | 0.0125 | 0.0022 | 19.9 | 0.0207 | 0.0020 | 18.4 |
| 2,000 | 0.00079 | $1.6 \mathrm{e}-4$ | 35.0 | 0.00885 | 0.0015 | 36.3 | 0.0149 | 0.0020 | 36.7 |

deviation (column Dev) and the average running time (with excluded time to compute the error of the resulting solution). For comparison, we also present the error of the initial solution $\tilde{z}_{1}=\left(x_{1}, y_{1}\right)$.

Our basic observation is as follows: Both N-SA and E-SA succeed to reduce the solution error reasonably fast. The N-SA implementation is preferable as it is more efficient in terms of running time. For comparison, it takes Matlab from 10 (for the simplest problem) to 35 seconds (for the hardest one) to compute just one answer $\mathrm{g}(x, y)=\left[\begin{array}{c}A^{T} y \\ -A x\end{array}\right]$ of the deterministic oracle.
5. Conclusions. It is shown that for a certain class of convex stochastic optimization and saddle point problems, robust versions of the SA approach have similar theoretical estimates of computational complexity, in terms of the required sample size, to the SAA method. Numerical experiments, reported in section 4, confirm this conclusion. These results demonstrate that for considered problems, a properly implemented mirror descent SA algorithm produces solutions of comparable accuracy to the SAA method for the same sample size of generated random points. On the other hand, the implementation (computational) time of the SA method is significantly smaller with a factor of up to $30-40$ for considered problems. Thus, both theoretical and numerical results suggest that the robust mirror descent SA is a viable alternative to the SAA approach, an alternative which at least deserves testing in particular applications. It is also shown that the robust mirror SA approach can be applied as a randomization algorithm to large-scale deterministic saddle point problems (in particular, to minimax optimization problems and bilinear matrix games) with encouraging results.
6. Appendix. Proof of Lemma 2.1. Let $x \in X^{o}$ and $v=P_{x}(y)$. Note that $v \in$ $\operatorname{argmin}_{z \in X}\left[\omega(z)+p^{T} z\right]$, where $p=\nabla \omega(x)-y$. Thus, $\omega$ is differentiable at $v$ and $v \in$
$X^{o}$. As $\nabla_{v} V(x, v)=\nabla \omega(v)-\nabla \omega(x)$, the optimality conditions for (2.31) imply that

$$
\begin{equation*}
(\nabla \omega(v)-\nabla \omega(x)+y)^{T}(v-u) \leq 0 \quad \forall u \in X \tag{6.1}
\end{equation*}
$$

For $u \in X$, we therefore have

$$
\begin{aligned}
V(v, u)-V(x, u)= & {\left[\omega(u)-\nabla \omega(v)^{T}(u-v)-\omega(v)\right] } \\
& -\left[\omega(u)-\nabla \omega(x)^{T}(u-x)-\omega(x)\right] \\
= & \nabla \omega(v)-\nabla \omega(x)+y)^{T}(v-u)+y^{T}(u-v) \\
\leq & y^{T}(u-v)-V(x, v)
\end{aligned}
$$

where the last inequality is due to (6.1). By Young's inequality, ${ }^{6}$ we have

$$
y^{T}(x-v) \leq \frac{\|y\|_{*}^{2}}{2 \alpha}+\frac{\alpha}{2}\|x-v\|^{2}
$$

while $V(x, v) \geq \frac{\alpha}{2}\|x-v\|^{2}$, due to the strong convexity of $V(x, \cdot)$. We get

$$
\begin{aligned}
V(v, u)-V(x, u) & \leq y^{T}(u-v)-V(x, v)=y^{T}(u-x)+y^{T}(x-v)-V(x, v) \\
& \leq y^{T}(u-x)+\frac{\|y\|_{*}^{2}}{2 \alpha}
\end{aligned}
$$

as required in (2.33).
Entropy as a distance-generating function on the standard simplex. The only property which is not immediately evident is that the entropy $w(x)=\sum_{i=1}^{n} x_{i} \ln x_{i}$ is strongly convex, modulus 1 with respect to $\|\cdot\|_{1}$-norm, on the standard simplex $X=\left\{x \in \mathbb{R}^{n}: x \geq 0, \sum_{i=1}^{n} x_{i}\right\}$. We are in the situation where $X^{o}=\{x \in X:$ $x>0\}$ and in order to establish the property in question, it suffices to verify that $h^{T} \nabla^{2} \omega(x) h \geq\|h\|_{1}^{2}$ for every $x \in X^{o}$. Here is the computation:
$\left[\sum_{i}\left|h_{i}\right|\right]^{2}=\left[\sum_{i}\left(x_{i}^{-1 / 2}\left|h_{i}\right|\right) x_{i}^{1 / 2}\right]^{2} \leq\left[\sum_{i} h_{i}^{2} x_{i}^{-1}\right]\left[\sum_{i} x_{i}\right]=\sum_{i} h_{i}^{2} x_{i}^{-1}=h^{T} \nabla^{2} \omega(x) h$,
where the inequality follows by Cauchy inequality.
Proof of Lemma 3.1. By (2.33), we have, for any $u \in Z$, that

$$
\begin{equation*}
\gamma_{t}\left(z_{t}-u\right)^{T} \mathrm{G}\left(z_{t}, \xi_{t}\right) \leq V\left(z_{t}, u\right)-V\left(z_{t+1}, u\right)+\frac{\gamma_{t}^{2}}{2}\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2} \tag{6.2}
\end{equation*}
$$

(recall that we are in the situation of $\alpha=1$ ). This relation implies that for every $u \in Z$, one has

$$
\begin{equation*}
\gamma_{t}\left(z_{t}-u\right)^{T} \mathrm{~g}\left(z_{t}\right) \leq V\left(z_{t}, u\right)-V\left(z_{t+1}, u\right)+\frac{\gamma_{t}^{2}}{2}\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2}-\gamma_{t}\left(z_{t}-u\right)^{T} \Delta_{t} \tag{6.3}
\end{equation*}
$$

where $\Delta_{t}=\mathrm{G}\left(z_{t}, \xi_{t}\right)-\mathrm{g}\left(z_{t}\right)$. Summing up these inequalities over $t=1, \ldots, j$, we get $\sum_{t=1}^{j} \gamma_{t}\left(z_{t}-u\right)^{T} \mathrm{~g}\left(z_{t}\right) \leq V\left(z_{1}, u\right)-V\left(z_{t+1}, u\right)+\sum_{t=1}^{j} \frac{\gamma_{t}^{2}}{2}\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2}-\sum_{t=1}^{j} \gamma_{t}\left(z_{t}-u\right)^{T} \Delta_{t}$.
Now we need the following simple lemma.

[^5]LEMmA 6.1. Let $\zeta_{1}, \ldots, \zeta_{j}$ be a sequence of elements of $\mathbb{R}^{n+m}$. Define the sequence $v_{t}, t=1,2, \ldots$ in $Z^{o}$ as follows: $v_{1} \in Z^{o}$ and

$$
v_{t+1}=P_{v_{t}}\left(\zeta_{t}\right), 1 \leq t \leq j
$$

Then, for any $u \in Z$, the following holds:

$$
\begin{equation*}
\sum_{t=1}^{j} \zeta_{t}^{T}\left(v_{t}-u\right) \leq V\left(v_{1}, u\right)+\frac{1}{2} \sum_{t=1}^{j}\left\|\zeta_{t}\right\|_{*}^{2} \tag{6.4}
\end{equation*}
$$

Proof. Using the bound (2.33) of Lemma 2.1 with $y=\zeta_{t}$ and $x=v_{t}$ (so that $\left.v_{t+1}=P_{v_{t}}\left(\zeta_{t}\right)\right)$ and recalling that we are in the situation of $\alpha=1$, we obtain the following for any $u \in Z$ :

$$
V\left(v_{t+1}, u\right) \leq V\left(v_{t}, u\right)+\zeta_{t}^{T}\left(u-v_{t}\right)+\frac{1}{2}\left\|\zeta_{t}\right\|_{*}^{2}
$$

Summing up from $t=1$ to $t=j$, we conclude that

$$
V\left(v_{j+1}, u\right) \leq V\left(v_{1}, u\right)+\sum_{t=1}^{j} \zeta_{t}^{T}\left(u-v_{t}\right)+\frac{1}{2} \sum_{t=1}^{j}\left\|\zeta_{t}\right\|_{*}^{2},
$$

which implies (6.4) due to $V(v, u) \geq 0$ for any $v \in Z^{o}, u \in Z$.
Applying Lemma 6.1 with $v_{1}=z_{1}, \zeta_{t}=-\gamma_{t} \Delta_{t}$, we get

$$
\begin{equation*}
\sum_{t=1}^{j} \gamma_{t} \Delta_{t}^{T}\left(u-v_{t}\right) \leq V\left(z_{1}, u\right)+\frac{1}{2} \sum_{t=1}^{j} \gamma_{t}^{2}\left\|\Delta_{t}\right\|_{*}^{2} \forall u \in Z \tag{6.5}
\end{equation*}
$$

Observe that

$$
\mathbb{E}\left\|\Delta_{t}\right\|_{*}^{2} \leq 4 \mathbb{E}\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2} \leq 4\left(\frac{2 D_{\omega_{x}, X}^{2}}{\alpha_{x}} M_{*, x}^{2}+\frac{2 D_{\omega_{y}, Y}^{2}}{\alpha_{y}} M_{*, y}^{2}\right)=4 M_{*}^{2}
$$

so that when taking the expectation of both sides of (6.5), we get

$$
\begin{equation*}
\mathbb{E}\left[\sup _{u \in Z}\left\{\sum_{t=1}^{j} \gamma_{t} \Delta_{t}^{T}\left(u-v_{t}\right)\right\}\right] \leq 1+2 M_{*}^{2} \sum_{t=1}^{j} \gamma_{t}^{2} \tag{6.6}
\end{equation*}
$$

(recall that $V\left(z_{1}, \cdot\right)$ is bounded by 1 on $Z$ ). Now we proceed exactly as in section 2.2: we sum up (6.3) from $t=1$ to $j$ to obtain

$$
\begin{align*}
& \sum_{t=1}^{j} \gamma_{t}\left(z_{t}-u\right)^{T} \mathrm{~g}\left(z_{t}\right) \leq V\left(z_{1}, u\right)+\sum_{t=1}^{j} \frac{\gamma_{t}^{2}}{2}\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2}-\sum_{t=1}^{j} \gamma_{t}\left(z_{t}-u\right)^{T} \Delta_{t}  \tag{6.7}\\
& =V\left(z_{1}, u\right)+\sum_{t=1}^{j} \frac{\gamma_{t}^{2}}{2}\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2}-\sum_{t=1}^{j} \gamma_{t}\left(z_{t}-v_{t}\right)^{T} \Delta_{t}+\sum_{t=1}^{j} \gamma_{t}\left(u-v_{t}\right)^{T} \Delta_{t}
\end{align*}
$$

When taking into account that $z_{t}$ and $v_{t}$ are deterministic functions of $\xi_{[t-1]}=$ $\left(\xi_{1}, \ldots, \xi_{t-1}\right)$ and that the conditional expectation of $\Delta_{t}, \xi_{[t-1]}$ being given, vanishes, we conclude that $\mathbb{E}\left[\left(z_{t}-v_{t}\right)^{T} \Delta_{t}\right]=0$. We take now suprema in $u \in Z$ and then
expectations on both sides of (6.7):

$$
\begin{aligned}
\mathbb{E}\left[\sup _{u \in Z} \sum_{t=1}^{j} \gamma_{t}\left(z_{t}-u\right)^{T} \mathrm{~g}\left(z_{t}\right)\right] \leq & \sup _{u \in Z} V\left(z_{1}, u\right)+\sum_{t=1}^{j} \frac{\gamma_{t}^{2}}{2} \mathbb{E}\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2} \\
& +\sup _{u \in Z} \sum_{t=1}^{j} \gamma_{t}\left(u-v_{t}\right)^{T} \Delta_{t} \\
(\operatorname{by}(6.6)) \leq & 1+\frac{M_{*}^{2}}{2} \sum_{t=1}^{j} \gamma_{t}^{2}+\left[1+2 M_{*}^{2} \sum_{t=1}^{j} \gamma_{t}^{2}\right] \\
= & 2+\frac{5}{2} M_{*}^{2} \sum_{t=1}^{j} \gamma_{t}^{2}
\end{aligned}
$$

and we arrive at (3.10).
Proof of Propositions 2.2 and 3.2. We provide here the proof of Proposition 3.2 only. The proof of Proposition 2.2 follows the same lines and can be easily reconstructed using the bound (2.39) instead of the relations (6.5) and (6.7) in the proof below.

First of all, with $M_{*}$ given by (3.6), one has

$$
\begin{equation*}
\forall(z \in Z): \mathbb{E}\left[\exp \left\{\|\mathrm{G}(z, \xi)\|_{*}^{2} / M_{*}^{2}\right\}\right] \leq \exp \{1\} \tag{6.8}
\end{equation*}
$$

Indeed, setting $p_{x}=\frac{2 D_{\omega_{x}, X}^{2} M_{*, x}^{2}}{\alpha_{x} M_{*}^{2}}, p_{y}=\frac{2 D_{\omega_{y}, Y}^{2} M_{*, y}^{2}}{\alpha_{y} M_{*}^{2}}$, we have $p_{x}+p_{y}=1$, whence, invoking (3.4),

$$
\mathbb{E}\left[\exp \left\{\|\mathrm{G}(z, \xi)\|_{*}^{2} / M_{*}^{2}\right\}\right]=\mathbb{E}\left[\exp \left\{p_{x}\left\|\mathrm{G}_{x}(z, \xi)\right\|_{*, x}^{2} / M_{*, x}^{2}+p_{y}\left\|\mathrm{G}_{y}(z, \xi)\right\|_{*, y}^{2} / M_{*, y}^{2}\right\}\right]
$$

and (6.8) follows from (3.16) by the Hölder inequality.
Setting $\Gamma_{N}=\sum_{t=1}^{N} \gamma_{t}$ and using the notation from the proof of Lemma 3.1, relations (3.9), (6.5), and (6.7) combined with the fact that $V\left(z_{1}, u\right) \leq 1$ for $u \in Z$, imply that

$$
\begin{equation*}
\Gamma_{N} \epsilon_{\phi}\left(\tilde{z}_{N}\right) \leq 2+\underbrace{\frac{1}{2} \sum_{t=1}^{N} \gamma_{t}^{2}\left[\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2}+\left\|\Delta_{t}\right\|_{*}^{2}\right]}_{\alpha_{N}}+\underbrace{\sum_{t=1}^{N} \gamma_{t}\left(v_{t}-z_{t}\right)^{T} \Delta_{t}}_{\beta_{N}} \tag{6.9}
\end{equation*}
$$

Now, from (6.8), it follows straightforwardly that

$$
\begin{equation*}
\mathbb{E}\left[\exp \left\{\left\|\Delta_{t}\right\|_{*}^{2} /\left(2 M_{*}\right)^{2}\right\}\right] \leq \exp \{1\}, \quad \mathbb{E}\left[\exp \left\{\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2} / M_{*}^{2}\right\}\right] \leq \exp \{1\} \tag{6.10}
\end{equation*}
$$

which, in turn, implies that

$$
\begin{equation*}
\mathbb{E}\left[\exp \left\{\alpha_{N} / \sigma_{\alpha}\right\}\right] \leq \exp \{1\}, \quad \sigma_{\alpha}=\frac{5}{2} M_{*}^{2} \sum_{t=1}^{N} \gamma_{t}^{2} \tag{6.11}
\end{equation*}
$$

and therefore, by Markov inequality, for any $\Omega>0$,

$$
\begin{equation*}
\operatorname{Prob}\left\{\alpha_{N} \geq(1+\Omega) \sigma_{\alpha}\right\} \leq \exp \{-\Omega\} \tag{6.12}
\end{equation*}
$$

Indeed, we have by (6.8)

$$
\left\|\mathrm{g}\left(z_{t}\right)\right\|_{*}=\left\|\mathbb{E}\left[\mathrm{G}\left(z_{t}, \xi_{t}\right) \mid \xi_{[t-1]}\right]\right\|_{*} \leq \sqrt{\mathbb{E}\left(\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2} \mid \xi_{[t-1]}\right)} \leq M_{*}
$$

and

$$
\left\|\Delta_{t}\right\|_{*}^{2}=\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)-\mathrm{g}\left(z_{t}\right)\right\|_{*}^{2} \leq\left(\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}+\left\|\mathrm{g}\left(z_{t}\right)\right\|_{*}\right)^{2} \leq 2\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2}+2 M_{*}^{2}
$$

which implies that

$$
\alpha_{N} \leq \sum_{t=1}^{N} \frac{\gamma_{t}^{2}}{2}\left[3\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2}+2 M_{*}^{2}\right]
$$

Further, by the Hölder inequality, we have the following from (6.8):

$$
\mathbb{E}\left[\exp \left\{\frac{\gamma_{t}^{2}\left[\frac{3}{2}\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2}+M_{*}^{2}\right]}{\frac{5}{2} \gamma_{t}^{2} M_{*}^{2}}\right\}\right] \leq \exp (1)
$$

Observe that if $r_{1}, \ldots, r_{i}$ are nonnegative random variables such that $\mathbb{E}\left[\exp \left\{r_{t} / \sigma_{t}\right\}\right] \leq$ $\exp \{1\}$ for some deterministic $\sigma_{t}>0$, then, by convexity of the exponent, $w(s)=$ $\exp \{s\}$ and

$$
\begin{equation*}
\mathbb{E}\left[\exp \left\{\frac{\sum_{t \leq i} r_{t}}{\sum_{t \leq i} \sigma_{t}}\right\}\right] \leq \mathbb{E}\left[\sum_{t \leq i} \frac{\sigma_{t}}{\sum_{\tau \leq i} \sigma_{\tau}} \exp \left\{r_{t} / \sigma_{t}\right\}\right] \leq \exp \{1\} \tag{6.13}
\end{equation*}
$$

Now applying (6.13) with $r_{t}=\gamma_{t}^{2}\left[\frac{3}{2}\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2}+M_{*}^{2}\right]$ and $\sigma_{t}=\frac{5}{2} \gamma_{t}^{2} M_{*}^{2}$, we obtain (6.11).

Now let $\zeta_{t}=\gamma_{t}\left(v_{t}-z_{t}\right)^{T} \Delta_{t}$. Observing that $v_{t}, z_{t}$ are deterministic functions of $\xi_{[t-1]}$, while $\mathbb{E}\left[\Delta_{t} \mid \xi_{[t-1]}\right]=0$, we see that the sequence $\left\{\zeta_{t}\right\}_{t=1}^{N}$ of random real variables forms a martingale difference. Besides this, by strong convexity of $\omega$ with modulus 1 w.r.t. $\|\cdot\|$ and due to $D_{\omega, Z} \leq 1$, we have

$$
u \in Z \Rightarrow 1 \geq V\left(z_{1}, u\right) \geq \frac{1}{2}\left\|u-z_{1}\right\|^{2}
$$

whence the $\|\cdot\|$-diameter of $Z$ does not exceed $2 \sqrt{2}$ so that $\left|\zeta_{t}\right| \leq 2 \sqrt{2} \gamma_{t}\left\|\Delta_{t}\right\|_{*}$, and therefore

$$
\mathbb{E}\left[\exp \left\{\left|\zeta_{t}\right|^{2} /\left(32 \gamma_{t}^{2} M_{*}^{2}\right)\right\} \mid \xi_{[t-1]}\right] \leq \exp \{1\}
$$

by (6.10). Applying Cramer's deviation bound, we obtain, for any $\Omega>0$,

$$
\begin{equation*}
\operatorname{Prob}\left\{\beta_{N}>4 \Omega M_{*} \sqrt{\sum_{t=1}^{N} \gamma_{t}^{2}}\right\} \leq \exp \left\{-\Omega^{2} / 4\right\} \tag{6.14}
\end{equation*}
$$

Indeed, for $0 \leq \gamma$, setting $\sigma_{t}=4 \sqrt{2} \gamma_{t} M_{*}$ and taking into account that $\zeta_{t}$ is a deterministic function of $\xi_{[t]}$, with $\mathbb{E}\left[\zeta_{t} \mid \xi_{[t-1]}\right]=0$ and $\mathbb{E}\left[\exp \left\{\zeta_{t}^{2} / \sigma_{t}^{2}\right\} \mid \xi_{[t-1]}\right] \leq \exp \{1\}$, we have

$$
\begin{aligned}
0<\gamma \sigma_{t} \leq 1 \Rightarrow & \left(\text { as } e^{x} \leq x+e^{x^{2}}\right) \\
\mathbb{E}\left[\exp \left\{\gamma \zeta_{t}\right\} \mid \xi_{[t-1]}\right] & \leq \mathbb{E}\left[\exp \left\{\gamma^{2} \zeta_{t}^{2}\right\} \mid \xi_{[t-1]}\right] \\
& \leq \mathbb{E}\left[\left(\exp \left\{\zeta_{t}^{2} / \sigma_{t}^{2}\right\}\right)^{\gamma^{2} \sigma_{t}^{2}} \mid \xi_{[t-1]}\right] \leq \exp \left\{\gamma^{2} \sigma_{t}^{2}\right\} ; \\
\gamma \sigma_{t}>1 \Rightarrow & \\
\mathbb{E}\left[\exp \left\{\gamma \zeta_{t}\right\} \mid \xi_{[t-1]}\right] & \leq \mathbb{E}\left[\exp \left\{\left.\left[\frac{1}{2} \gamma^{2} \sigma_{t}^{2}+\frac{1}{2} \zeta_{t}^{2} / \sigma_{t}^{2}\right\} \right\rvert\, \xi_{[t-1]}\right]\right. \\
& \leq \exp \left\{\frac{1}{2} \gamma^{2} \sigma_{t}^{2}+\frac{1}{2}\right\} \leq \exp \left\{\gamma^{2} \sigma_{t}^{2}\right\},
\end{aligned}
$$

that is, in both cases, $\mathbb{E}\left[\exp \left\{\gamma \zeta_{t}\right\} \mid \xi_{[t-1]}\right] \leq \exp \left\{\gamma^{2} \sigma_{t}^{2}\right\}$. Therefore,

$$
\mathbb{E}\left[\exp \left\{\gamma \beta_{i}\right\}\right]=\mathbb{E}\left[\exp \left\{\gamma \beta_{i-1}\right\} \mathbb{E}\left[\exp \left\{\gamma \zeta_{i}\right\} \mid \xi_{[i-1]}\right]\right] \leq \exp \left\{\gamma^{2} \sigma_{i}^{2}\right\} \mathbb{E}\left[\exp \left\{\gamma \beta_{i-1}\right\}\right],
$$

whence $\mathbb{E}\left[\exp \left\{\gamma \beta_{N}\right\}\right] \leq \exp \left\{\gamma^{2} \sum_{t=1}^{N} \sigma_{t}^{2}\right\}$, and thus, by Markov inequality for every $\Omega>0$, it holds

$$
\operatorname{Prob}\left\{\beta_{N}>\Omega \sqrt{\sum_{t=1}^{N} \sigma_{t}^{2}}\right\} \leq \exp \left\{\gamma^{2} \sum_{t=1}^{N} \sigma_{t}^{2}\right\} \exp \left\{-\gamma \Omega \sqrt{\sum_{t=1}^{N} \sigma_{t}^{2}}\right\} .
$$

When choosing $\gamma=\frac{1}{2} \Omega\left(\sum_{t=1}^{N} \sigma_{t}^{2}\right)^{-1 / 2}$, we arrive at (6.14).
Combining (6.9), (6.10), and (6.14), we get the following for any positive $\Omega$ and $\Theta$ :

$$
\begin{aligned}
\operatorname{Prob}\left\{\Gamma_{N} \epsilon_{\phi}\left(\tilde{z}_{t}\right)\right. & \left.>2+\frac{5}{2}(1+\Omega) M_{*}^{2} \sum_{t=1}^{N} \gamma_{t}^{2}+4 \sqrt{2} \Theta M_{*} \sqrt{\sum_{t=1}^{N} \gamma_{t}^{2}}\right\} \\
& \leq \exp \{-\Omega\}+\exp \left\{-\frac{1}{4} \Theta^{2}\right\} .
\end{aligned}
$$

When setting $\Theta=2 \sqrt{\Omega}$ and substituting (3.12), we obtain (3.17).
Proof of Proposition 3.3. As in the proof of Proposition 3.2, when setting $\Gamma_{N}=$ $\sum_{t=1}^{N} \gamma_{t}$ and using the relations (3.9), (6.5), and (6.7), combined with the fact that $\left\|\mathrm{G}\left(z, \xi_{y}\right)\right\|_{*} \leq M_{*}$, we obtain

$$
\begin{array}{r}
\Gamma_{N} \epsilon_{\phi}\left(\tilde{z}_{N}\right) \leq 2+\sum_{t=1}^{N} \frac{\gamma_{t}^{2}}{2}\left[\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|_{*}^{2}+\left\|\Delta_{t}\right\|_{*}^{2}\right]+\sum_{t=1}^{N} \gamma_{t}\left(v_{t}-z_{t}\right)^{T} \Delta_{t} \\
\leq 2+\frac{5}{2} M_{*}^{2} \sum_{t=1}^{N} \gamma_{t}^{2}+\underbrace{\sum_{t=1}^{N} \gamma_{t}\left(v_{t}-z_{t}\right)^{T} \Delta_{t}}_{\alpha_{N}} \tag{6.15}
\end{array}
$$

Recall that by definition of $\Delta_{t},\left\|\Delta_{t}\right\|_{*}=\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)-\mathrm{g}\left(z_{t}\right)\right\|_{*} \leq\left\|\mathrm{G}\left(z_{t}, \xi_{t}\right)\right\|+\left\|\mathrm{g}\left(z_{t}\right)\right\|_{*} \leq$ $2 M_{*}$.

Note that $\zeta_{t}=\gamma_{t}\left(v_{t}-z_{t}\right)^{T} \Delta_{t}$ is a bounded martingale difference, i.e., $\mathbb{E}\left(\zeta_{t} \mid \xi_{[t-1]}\right)=$ 0 and $\left|\zeta_{t}\right| \leq 4 \gamma_{t} \bar{M}$ (here $\bar{M}$ is defined in (3.31)). Then, by Azuma-Hoeffding's inequality [1] for any $\Omega \geq 0$,

$$
\begin{equation*}
\operatorname{Prob}\left(\alpha_{N}>4 \Omega \bar{M} \sqrt{\sum_{t=1}^{N} \gamma_{t}^{2}}\right) \leq e^{-\Omega^{2} / 2} . \tag{6.16}
\end{equation*}
$$

Indeed, let us denote $v_{t}=\left(v_{t}^{(x)}, v_{t}^{(y)}\right)$ and $\Delta_{t}=\left(\Delta_{t}^{(x)}, \Delta_{t}^{(y)}\right)$. When taking into account that $\left\|v_{t}^{(x)}\right\|_{1} \leq 1,\left\|v_{t}^{(y)}\right\|_{1} \leq 1$ and $\left\|x_{t}\right\|_{1} \leq 1,\left\|y_{t}\right\|_{1} \leq 1$, we conclude that

$$
\begin{aligned}
\left|\left(v_{t}-z_{t}\right)^{T} \Delta_{t}\right| & \leq\left|\left(v_{t}^{(x)}-x_{t}\right)^{T} \Delta_{t}^{(x)}\right|+\left|\left(v_{t}^{(y)}-y_{t}\right)^{T} \Delta_{t}^{(y)}\right| \\
& \leq 2\left\|\Delta_{t}^{(x)}\right\|_{\infty}+2\left\|\Delta_{t}^{(y)}\right\|_{\infty} \leq 4 \max _{1 \leq j \leq m}\left\|A^{j}+b\right\|_{\infty}+4 \max _{1 \leq j \leq n}\left\|A_{j}+c\right\|_{\infty} \\
& =4 \bar{M}
\end{aligned}
$$

We conclude from (6.15) and (6.16) that

$$
\operatorname{Prob}\left(\Gamma_{N} \epsilon_{\phi}\left(\tilde{z}_{N}\right)>2+\frac{5}{2} M_{*}^{2} \sum_{t=1}^{N} \gamma_{t}^{2}+4 \Omega \bar{M} \sqrt{\sum_{t=1}^{N} \gamma_{t}^{2}}\right) \leq e^{-\Omega^{2} / 2}
$$

and the bound (3.30) of the proposition can be easily obtained by substituting the constant stepsizes $\gamma_{t}$ as defined in (3.12).

## REFERENCES

[1] K. Azuma, Weighted sums of certain dependent random variables, Tökuku Math. J., 19 (1967), pp. 357-367.
[2] A. Ben-Tal and A. Nemirovski, Non-Euclidean restricted memory level method for large-scale convex optimization, Math. Program., 102 (2005), pp. 407-456.
[3] A. Benveniste, M. Métivier, and P. Priouret, Algorithmes Adaptatifs et Approximations Stochastiques, Masson, Paris, 1987 (in French). Adaptive Algorithms and Stochastic Approximations, Springer, New York, 1993 (in English).
[4] L.M. Bregman, The relaxation method of finding the common points of convex sets and its application to the solution of problems in convex programming, Comput. Math. Math. Phys., 7 (1967), pp. 200-217.
[5] K.L. Chung, On a stochastic approximation method, Ann. Math. Statist., 25 (1954), pp. 463483.
[6] Y. Ermoliev, Stochastic quasigradient methods and their application to system optimization, Stochastics, 9 (1983), pp. 1-36.
[7] A.A. Gaivoronski, Nonstationary stochastic programming problems, Kybernetika, 4 (1978), pp. 89-92.
[8] A.V. Goldberg and R.E. Tarjan, A new approach to the maximum flow problem, J. ACM, 35 (1988), pp. 921-940.
[9] M.D. Grigoriadis and L.G. Khachiyan, A sublinear-time randomized approximation algorithm for matrix games, Oper. Res. Lett., 18 (1995), pp. 53-58.
[10] A. Juditsky, A. Nazin, A. Tsybakov, and N. Vayatis, Recursive aggregation of estimators by the mirror descent algorithm with averaging, Probl. Inf. Transm., 41 (2005), pp. 368-384.
[11] A.J. Kleywegt, A. Shapiro, and T. Homem-de-Mello, The sample average approximation method for stochastic discrete optimization, SIAM J. Optim., 12 (2002), pp. 479-502.
[12] J. Linderoth, A. Shapiro, and S. Wright, The empirical behavior of sampling methods for stochastic programming, Ann. Oper. Res., 142 (2006), pp. 215-241.
[13] W.K. Mak, D.P. Morton, and R.K. Wood, Monte Carlo bounding techniques for determining solution quality in stochastic programs, Oper. Res. Lett., 24 (1999), pp. 47-56.
[14] M.D. McKay, R.J. Beckman, and W.J. Conover, A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, Technometrics, 21 (1979), pp. 239-245.
[15] A. Nemirovski and D. Yudin, On Cezari's convergence of the steepest descent method for approximating saddle point of convex-concave functions, Dokl. Akad. Nauk SSSR, 239 (1978) (in Russian). Soviet Math. Dokl., 19 (1978) (in English).
[16] A. Nemirovski and D. Yudin, Problem Complexity and Method Efficiency in Optimization, Wiley-Intersci. Ser. Discrete Math. 15, John Wiley, New York, 1983.
[17] Y. Nesterov, Primal-dual subgradient methods for convex problems, Math. Program., Ser. B, http://www.springerlink.com content-b441795t5254m533.
[18] B.T. Polyak, New stochastic approximation type procedures, Automat. i Telemekh., 7 (1990), pp. 98-107.
[19] B.T. Polyak and A.B. Juditsky, Acceleration of stochastic approximation by averaging, SIAM J. Control Optim., 30 (1992), pp. 838-855.
[20] G.C. Pflug, Optimization of Stochastic Models, The Interface Between Simulation and Optimization, Kluwer, Boston, 1996.
[21] H. Robbins and S. Monro, A stochastic approximation method, Ann. Math. Stat., 22 (1951), pp. 400-407.
[22] A. Ruszczyński and W. Syski, A method of aggregate stochastic subgradients with on-line stepsize rules for convex stochastic programming problems, Math. Prog. Stud., 28 (1986), pp. 113-131.
[23] J. Sacks, Asymptotic distribution of stochastic approximation, Ann. Math. Stat., 29 (1958), pp. 373-409.
[24] S. Sen, R.D. Doverspike, and S. Cosares, Network planning with random demand, Telecomm. Syst., 3 (1994), pp. 11-30.
[25] A. Shapiro, Monte Carlo sampling methods, in Stochastic Programming, Handbook in OR \& MS, Vol. 10, A. Ruszczyński and A. Shapiro, eds., North-Holland, Amsterdam, 2003.
[26] A. Shapiro and A. Nemirovski, On complexity of stochastic programming problems, in Continuous Optimization: Current Trends and Applications, V. Jeyakumar and A.M. Rubinov, eds., Springer, New York, 2005, pp. 111-144.
[27] J.C. SpalL, Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control, John Wiley, Hoboken, NJ, 2003.
[28] V. Strassen, The existence of probability measures with given marginals, Ann. Math. Statist., 38 (1965), pp. 423-439.
[29] B. Verweit, S. Ahmed, A.J. Kleywegt, G. Nemhauser, and A. Shapiro, The sample average approximation method applied to stochastic routing problems: A computational study, Comput. Optim. Appl., 24 (2003), pp. 289-333.


[^0]:    *Received by the editors October 1, 2007; accepted for publication (in revised form) August 26, 2008; published electronically January 21, 2009.
    http://www.siam.org/journals/siopt/19-4/70427.html
    ${ }^{\dagger}$ Georgia Institute of Technology, Atlanta, Georgia 30332 (nemirovs@isye.gatech.edu, glan@isye. gatech.edu, ashapiro@isye.gatech.edu). Research of the first author was partly supported by NSF award DMI-0619977. Research of the third author was partially supported by NSF award CCF0430644 and ONR award N00014-05-1-0183. Research of the fourth author was partly supported by NSF awards DMS-0510324 and DMI-0619977.
    ${ }^{\ddagger}$ Université J. Fourier, B.P. 53, 38041 Grenoble Cedex 9, France (Anatoli.Juditsky@imag.fr).

[^1]:    ${ }^{1}$ Throughout the paper, we speak about convergence in terms of the objective value.

[^2]:    ${ }^{2}$ In the second case, we apply the SA after the variables are scaled to make $X$ the unit $\|\cdot\|_{1}$-ball.

[^3]:    ${ }^{3}$ The possibility to solve matrix games in a sublinear-time fashion by a randomized algorithm was discovered by Grigoriadis and Khachiyan [9]. Their "ad hoc" algorithm is similar, although not completely identical to ours, and possesses the same complexity bounds.

[^4]:    ${ }^{4}$ The difference in running times can be easily explained: with $X$ being a simplex, the proxmapping for E-SA takes $O(n \ln n)$ operations versus $O(n)$ operations for N-SA.
    ${ }^{5}$ Indeed, $\left(\xi_{1}, \ldots, \xi_{n}\right) \sim \mathcal{N}\left(0, I_{n}\right)$, so that the random variable $\xi_{x}=\sum_{i}\left(a_{i}+\xi_{i}\right) x_{i}$ is normal with easily computable mean and variance, and since $\phi$ is piecewise linear, the expectation $f(x)=\mathbb{E}\left[\phi\left(\xi_{x}\right)\right]$ can be immediately expressed via the error function.

[^5]:    ${ }^{6}$ For any $u, v \in \mathbb{R}^{n}$, we have, by the definition of the dual norm, that $\|u\|_{*}\|v\| \geq u^{T} v$, and hence $\left(\|u\|_{*}^{2} / \alpha+\alpha\|v\|^{2}\right) / 2 \geq\|u\|_{*}\|v\| \geq u^{T} v$.

