



**HAL**  
open science

# Investigating Variance Definitions for Stochastic Mirror Descent with Relative Smoothness

Hadrien Hendrikx

► **To cite this version:**

Hadrien Hendrikx. Investigating Variance Definitions for Stochastic Mirror Descent with Relative Smoothness. 2024. hal-04830814

**HAL Id: hal-04830814**

**<https://hal.science/hal-04830814v1>**

Preprint submitted on 11 Dec 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

---

# Investigating Variance Definitions for Stochastic Mirror Descent with Relative Smoothness

---

Hadrien Hendrikx<sup>1</sup>

## Abstract

Mirror Descent is a popular algorithm, that extends Gradients Descent (GD) beyond the Euclidean geometry. One of its benefits is to enable strong convergence guarantees through smooth-like analyses, even for objectives with exploding or vanishing curvature. This is achieved through the introduction of the notion of *relative smoothness*, which holds in many of the common use-cases of Mirror descent. While basic deterministic results extend well to the relative setting, most existing stochastic analyses require additional assumptions on the mirror, such as strong convexity (in the usual sense), to ensure bounded variance. In this work, we revisit Stochastic Mirror Descent (SMD) proofs in the (relatively-strongly-) convex and relatively-smooth setting, and introduce a new (less restrictive) definition of variance which can generally be bounded (globally) under mild regularity assumptions. We then investigate this notion in more details, and show that it naturally leads to strong convergence guarantees for stochastic mirror descent. Finally, we leverage this new analysis to obtain convergence guarantees for the Maximum Likelihood Estimator of a Gaussian with unknown mean and variance.

## 1. Introduction

The central problem of this paper is to solve optimization problems of the following form:

$$\min_{x \in C} f(x), \text{ where } f(x) = \mathbb{E}[f_\xi(x)], \quad (1)$$

where  $C$  is a closed convex subset of  $\mathbb{R}^d$ , and  $f_\xi$  are differentiable convex functions (stochasticity is on the variable  $\xi$ ). The problems that we will consider typically arise from machine-learning use-cases, meaning that the dimension  $d$

can be very large. Therefore, first-order methods are very popular for solving these problems, since they usually scale very well with the dimension.

In standard machine learning setups, computing a gradient of  $f$  is very costly (or even impossible), since it requires computing gradients for all individual examples in the dataset. Yet, gradients of  $f_\xi$  are relatively cheap, and arbitrarily high precisions are generally not required. This makes Stochastic Gradient Descent (SGD) the method of choice (Bottou, 2010). The SGD updates can be written as:

$$x_{\text{SGD}}^+ = \arg \min_{u \in C} \left\{ \eta \nabla f_\xi(x)^\top u + \frac{1}{2} \|u - x\|^2 \right\}. \quad (2)$$

While the standard Euclidean geometry leading to Gradient Descent (GD) fits many use-cases quite well, several applications are better solved with *Mirror Descent* (MD), a generalization of GD which allows to better capture the geometry of the problem. For instance, the Kullback-Leibler divergence might be better suited to discriminating between probability distributions than the (squared) Euclidean norm, and this is something that one can leverage using MD with entropy as a mirror. As a matter of fact, many standard algorithms can be interpreted as MD, *i.e.*, as generalized first-order methods. This is for instance the case in statistics, where Expectation Minimization and Maximum A Posteriori estimators can be interpreted as running MD with specific mirror and step-sizes (Kunstner et al., 2021; Le Priol et al., 2021). Mirror descent can also be used to solve Poisson inverse problems, which have many applications in astronomy and medicine (Bertero et al., 2009), to reduce the communication cost of distributed algorithms (Shamir et al., 2014; Hendrikx et al., 2020b), or to solve convex quartic problems (Dragomir & Nesterov, 2023). In the online learning community as well, many standard algorithms such as Exponential Weight Updates or Follow-The-Regularized-Leader can be interpreted as running mirror descent (McMahan, 2011; Hoeffner et al., 2018). There are still many open questions regarding the convergence guarantees for most of the algorithms mentioned above. Therefore, progress on the understanding of MD can lead to a plethora of results on these applications, and more generally to a more consistent theory for Majorization-Minimization algorithms. This paper is

---

<sup>1</sup> Centre Inria de l'Univ. Grenoble Alpes CNRS, LJK, Grenoble, France. Correspondence to: Hadrien Hendrikx <hadrien.hendrikx@inria.fr>.

a stepping stone in this direction.

Let us now introduce the *mirror map*, or *potential function*  $h$ , together with the *Bregman divergence* with respect to  $h$ , which is defined for  $x, y \in \text{dom } h$  as:

$$D_h(x, y) = h(x) - h(y) - \nabla h(y)^\top (x - y). \quad (3)$$

We now introduce the Stochastic Mirror Descent (SMD) update, which can be found in its deterministic form in, e.g., Nemirovskij & Yudin (1983). SMD consists in replacing the squared Euclidean norm from the SGD update (2) by the Bregman divergence with respect to the mirror map  $h$ :

$$x^+(\eta, \xi) = \arg \min_{u \in C} \{ \eta \nabla f_\xi(x)^\top u + D_h(u, x) \}. \quad (4)$$

Note that since  $D_{\|\cdot\|^2}(x, y) = \|x - y\|^2$ , one can recover SGD by taking  $h = \frac{1}{2}\|\cdot\|^2$ . In this sense, Stochastic Mirror Descent can be viewed as standard SGD, but changing the way distances are computed, and so the geometry of the problem. Yet, this change significantly complicates the convergence analysis of the method, since the Bregman divergence, *in general*:

1. Does not satisfy the triangular inequality,
2. Is not symmetric,
3. Is not translation-invariant,
4. Is not convex in its second argument.

This means that analyzing mirror descent methods requires quite some care, and that many standard (S)GD results do not extend to the mirror setting. For instance, one can prove that mirror descent cannot be *accelerated* in general (Dragomir et al., 2021b). Similarly, applying techniques such as variance-reduction requires additional assumptions (Dragomir et al., 2021a).

To ensure that  $x^+(\eta, \xi)$  exists and is unique, we first make the following blanket assumption throughout the paper:

**Assumption 1.** *Function  $h : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{\infty\}$  is twice continuously differentiable and strictly convex on  $C$ . For every  $y \in \mathbb{R}^d$ , the problem  $\min_{x \in C} h(x) - x^\top y$  has a unique solution, which lies in  $\text{int } C$ , and all  $f_\xi$  are convex.*

Note that the regularity assumption on  $h$  could be relaxed, as discussed in Section 3, but we choose a rather strong one to make sure all the objects we will manipulate are well-defined. Interestingly, while mirror descent changes the way distances are computed to move away from the Euclidean geometry, standard analyses of mirror descent methods, and in particular in the online learning community, still require strong convexity and Lipschitz continuity

with respect to norms (Bubeck et al., 2015, Chapter 4). It is only recently that a *relative smoothness* assumption was introduced to study mirror descent (Bauschke et al., 2017; Lu et al., 2018), together with the corresponding relative strong convexity.

**Definition 1.** The function  $f$  is said to be  $L$ -relatively smooth and  $\mu$ -relatively strongly convex with respect to  $h$  if for all  $x, y \in C$ :

$$\mu D_h(x, y) \leq D_f(x, y) \leq L D_h(x, y). \quad (5)$$

Definition 1 extends the standard smooth and strongly convex assumptions that correspond to the case  $h = \frac{1}{2}\|\cdot\|^2$ , so that for all  $x \in C$ ,  $\nabla^2 h(x) = I$  the identity matrix. These assumptions allow MD analyses to generalize standard GD analyses, and in particular to obtain similar linear and sub-linear rates, with constant step-size and conditions adapted to the *relative* assumptions.

While the basic deterministic setting is now well-understood under relative assumptions, a good understanding of the stochastic setting remains elusive. In particular, as we will see in more details in the related work section, *all existing proofs somehow require the mirror  $h$  to be globally strongly convex with respect to a norm*, or have non-vanishing variance. The only case that can be analyzed tightly is under *interpolation* (there exists a point that minimizes all stochastic functions), or when using Coordinate Descent instead of SMD (Hanzely & Richtárik, 2021; Hendrikx et al., 2020a). This is a major weakness, as the goal of relative smoothness is precisely to avoid comparisons to norms. Indeed, even when these “absolute” regularity assumptions hold, the smoothness and strong convexity constants are typically very loose, and the theory is not representative of the observed behaviour of the algorithms.

However, as hinted at earlier, this was expected: acceleration is notoriously hard to achieve for mirror descent (and even impossible in general (Dragomir et al., 2021b)), and variance reduction typically encounters the same problems (Dragomir et al., 2021a). For stochastic updates, this comes from the fact that it is impossible to disentangle the stochastic gradient from the effect of the curvature of  $h$  at the point at which it is applied.

**Contribution and outline.** The main contribution of this paper is to introduce a new analysis for mirror descent, with a variance notion which is provably bounded under mild regularity assumptions: typically, the same as those required for the deterministic case. We introduce our new variance notion, and compare it with standard ones from the literature in Section 2. This new analysis is both simpler and tighter than existing ones, as shown in Section 3. Finally, we use our results to analyse the convergence of the Maximum Likelihood and Maximum A Posteriori estimators for a Gaussian with unknown mean and variance

in Section 4, and show that it is the first generic stochastic mirror descent analysis that obtains meaningful finite-time convergence guarantees in this case.

## 2. Variance Assumptions

We now focus on the various variance assumptions under which Stochastic Mirror Descent is analyzed. Some manipulations require technical lemmas, such as the duality property of the Bregman divergence or the Bregman co-coercivity lemma, which can be found in Appendix A.

We start by introducing our variance definition, prove a few good properties for it, and then compare it with the existing ones to highlight their shortcomings. The two key properties we would like to ensure (and which are not satisfied by other definitions) are: (i) boundedness without strong convexity of  $h$  or restricting the SMD iterates, and (ii) finiteness for  $\eta \rightarrow 0$  (with the appropriate scaling).

### 2.1. New variance definition

Let  $\eta > 0$ , and recall that  $x^+(\eta, \xi)$  is the result of a stochastic mirror descent step from  $x$  using function  $f_\xi$  with step-size  $\eta$  (Equation (4)). From now on, when clear from the context, we will simply denote this point  $x^+$ . Yet, although the dependence is now implicit, do keep in mind that  $x^+$  is a stochastic quantity that is not independent from  $\xi$  nor  $\eta$ , as this is critical in most results. Under Assumption 1,  $x^+$  writes:

$$\nabla h(x^+) = \nabla h(x) - \eta \nabla f_\xi(x). \quad (6)$$

Similarly, we denote by  $\overline{x^+}$  the deterministic Mirror Descent update, which is such that  $\nabla h(\overline{x^+}) = \nabla h(x) - \eta \nabla f(x)$ . We also introduce  $h^* : y \mapsto \arg \max_{x \in C} x^\top y - h(x)$  the convex conjugate of  $h$ , which verifies  $\nabla h^*(\nabla h(x)) = x$ . Let us now define the function

$$f_\eta(x) = f(x) - \frac{1}{\eta} \mathbb{E} [D_h(x, x^+)], \quad (7)$$

which is at the heart of our variance definition:

**Definition 2.** We define the variance of the stochastic mirror descent iterates given by (4) as:

$$\sigma_{\star, \eta}^2 = \frac{1}{\eta} \sup_{x \in C} (f(x_\star) - f_\eta(x)) = \frac{f^\star - f_\eta^\star}{\eta}, \quad (8)$$

where  $f^\star$  and  $f_\eta^\star$  are respectively the infima of  $f$  and  $f_\eta$ .

In order to simplify the presentation for this section, we make the following blanket assumption. However, note that most results hold without it, since derivations are performed for all  $x$ , and then specialized to  $x = x_\eta$ . In case  $x_\eta$  is not achieved (if it is on the boundary of  $C$  for instance), we can replace all results by a limit for  $x \rightarrow x_\eta$ .

**Assumption 2.** If  $\sigma_{\star, \eta}^2 < \infty$ , the supremum is realized at point  $x_\eta$ , so that  $\sigma_{\star, \eta}^2 = \frac{1}{\eta} (f(x_\star) - f_\eta(x_\eta))$ .

We now state various bounds on  $\sigma_{\star, \eta}^2$ , to help understand its behaviour. We start by positivity, which is an essential property that justifies the square in the definition.

**Proposition 2.1** (Positivity). *For all  $\eta > 0$ ,  $\sigma_{\star, \eta} \geq 0$ .*

This result is direct from the fact that  $f_\eta(x) \leq f(x)$  since  $D_h(x, x^+) \geq 0$  for all  $x \in C$  by convexity of  $h$ .

**Stochastic functions after a step.** We first provide a result that upper bounds  $\sigma_{\star, \eta}^2$  directly in terms of  $f_\xi$ .

**Proposition 2.2.** *If  $f_\xi$  is  $L$ -relatively-smooth with respect to  $h$  and  $\eta \leq 1/L$ , then:*

$$\sigma_{\star, \eta}^2 \leq \frac{1}{\eta} \left( f(x_\star) - \min_{x \in C} \mathbb{E} [f_\xi(x^+)] \right). \quad (9)$$

*Proof.* We write:

$$\begin{aligned} D_h(x, x^+) &= \langle \nabla h(x^+) - \nabla h(x), x^+ - x \rangle - D_h(x^+, x) \\ &= -\eta \nabla f_\xi(x)^\top (x^+ - x) - D_h(x^+, x) \\ &= \eta (D_{f_\xi}(x^+, x) - f_\xi(x^+) + f_\xi(x)) - D_h(x^+, x). \end{aligned}$$

The relative smoothness of  $f_\xi$  and the step-size condition imply that  $\eta D_{f_\xi}(x^+, x) \leq D_h(x^+, x)$ , leading to  $\frac{1}{\eta} D_h(x, x^+) \leq f_\xi(x) - f_\xi(x^+)$ , and the result follows.  $\square$

This bound offers a new point of view on the variance, which can be bounded as the difference between the optimum of  $f$ , and the optimum of a related function, in which we make one mirror descent step before evaluating each  $f_\xi$ . Note that we can also allow this gradient step on the  $f$  part, but it would not change anything since  $x_\star^+ = x_\star$ . In this sense, if we define the operator  $\mathcal{T}(g) = g(\nabla h^*(\nabla h(x) - \eta \nabla g(x)))$ , then the bound from proposition (2.2) becomes:

$$\sigma_{\star, \eta}^2 \leq \min_{x \in C} \mathcal{T}(\mathbb{E}[f_\xi])(x) - \min_{x \in C} \mathbb{E}[\mathcal{T}(f_\xi)](x). \quad (10)$$

We recognize the structure of a variance, as the difference between an operator applied to the expectation of a random variable, and the expectation of the operator applied to the random variable. Yet, compared to standard (Euclidean) analyses of SGD, it does not simply corresponds to the variance of the stochastic gradients (at optimum), and bears a more complex form. We will see that  $\sigma_{\star, \eta}^2$  can actually directly be expressed in such a form, by using the  $c$ -transform.

An interesting observation is that  $\mathbb{E}[f_\xi(x^+)]$  is reminiscent of meta-learning objectives (Finn et al., 2017). This connection highlights that the variance is bounded by the

difference between the standard objective and the MAML one.

**Finiteness.** Proposition 2.2 implies the following:

**Corollary 2.3.** *If  $f_\xi$  is  $L$ -relatively-smooth w.r.t.  $h$  and admits a minimum  $x_\star^\xi \in \text{int } C$  a.s., then for all  $\eta \leq 1/L$ ,  $\sigma_{\star,\eta}^2 \leq \frac{f(x_\star) - \mathbb{E}[f_\xi(x_\star^\xi)]}{\eta}$ . In particular,  $\sigma_{\star,\eta}^2$  is finite.*

This result directly comes from the fact that  $\min_{x \in C} \mathbb{E}[f_\xi(x^+)] \geq \mathbb{E}[\min_{x \in C} f_\xi(x^+)] \geq \mathbb{E}[f_\xi(x_\star^\xi)]$ . It shows that the standard regularity assumptions for the convergence of stochastic mirror descent guarantee that *the variance as introduced in Definition 2 remains bounded*. This is a strong result, that justifies the supremum in the variance definition. Indeed, **most other variance definitions require additional assumptions for the variance to remain bounded after the supremum**. Instead, we *globalize* the variance definition, by taking the supremum over the right quantity to ensure that it remains bounded over the whole domain without having to explicitly assume it.

Note that the bound from Corollary 2.3 has already been investigating in other settings for stochastic optimization (Loizou et al., 2021), as discussed in Section 2.2. While useful to show boundedness, this bound has a major drawback, which is that it explodes when the step-size  $\eta$  vanishes. This does not reflect what happens in practice, which is why we investigate finer bounds on  $\sigma_{\star,\eta}^2$ .

**Gradient norm at optimum.** A usual way of formulating variance is to express it as the norm of the difference between stochastic gradients and the deterministic gradients. While the previous bounds highlight dependencies on the gradient steps (through evaluations at  $x^+$ ), none of them really corresponds to “the size of the stochastic gradients at optimum”. The key subtlety is that when using mirror descent, it is important to also specify the point at which these gradients are applied, and the following proposition gives a bound of this flavor on  $\sigma_{\star,\eta}^2$ .

**Proposition 2.4.** *If  $f$  is  $L$ -relatively-smooth w.r.t.  $h$ ,  $\eta \leq 1/L$  and  $x_\star \in \text{int } C$  we have:*

$$\sigma_{\star,\eta}^2 \leq \frac{1}{\eta^2} \mathbb{E} \left[ D_h \left( \overline{x_\star^+}, x_\star^+ \right) \right]. \quad (11)$$

This can safely be considered as the Mirror Descent equivalent of  $\mathbb{E}[\|\nabla f_\xi(x_\star)\|^2]$ . Yet, a key difference is that stochastic gradients are evaluated at point  $x_\eta$  instead of  $x_\star$ , and  $\nabla f(x_\eta) \neq 0$  in general.

*Proof.* For all  $x$ , we have:

$$\begin{aligned} \mathbb{E} [D_h(x, x^+)] &= \mathbb{E} [D_{h^\star}(\nabla h(x^+), \nabla h(x))] \\ &= \mathbb{E} [D_{h^\star}(\nabla h(x) - \eta \nabla f_\xi(x), \nabla h(x))] \\ &= \mathbb{E} [D_{h^\star}(\nabla h(x) - \eta \nabla f(x), \nabla h(x))] \\ &+ \mathbb{E} [D_{h^\star}(\nabla h(x) - \eta \nabla f_\xi(x), \nabla h(x) - \eta \nabla f(x))] \\ &= \mathbb{E} [D_{h^\star}(\nabla h(x) - \eta [\nabla f(x) - \nabla f(x_\star)], \nabla h(x))] \\ &+ \mathbb{E} \left[ D_h \left( \overline{x^+}, x^+ \right) \right] \\ &\leq \eta D_f(x, x_\star) + \mathbb{E} \left[ D_h \left( \overline{x^+}, x^+ \right) \right]. \end{aligned}$$

Here, the first equality comes from the duality property of the Bregman divergence, the third one from the Bregman bias-variance decomposition Lemma (Pfau, 2013), and the last inequality from the Bregman cocoercivity Lemma (Dragomir et al., 2021a). All these technical results can be found in Appendix A. Therefore,

$$f_\eta(x) \geq f(x_\star) - \frac{1}{\eta} \mathbb{E} \left[ D_h(\overline{x^+}, x^+) \right], \quad (12)$$

and this is in particular true for  $x = x_\eta$ .  $\square$

**Limit behaviour.** A first observation is that both the  $D_h(x, x^+)$  term in the definition of  $f_\eta$  and our variance definition are scaled by  $\eta^{-1}$ , potentially indicating that they blow up when  $\eta \rightarrow 0$ . While it is clear that it is not the case in the Euclidean setting, this property holds more generally, as shown in the two following results.

**Proposition 2.5.** *Let  $x \in C$  and  $\eta_0 > 0$  be such that  $\mathbb{E} D_h(x, x^+(\eta_0, \xi)) < \infty$ . Then,  $f_\eta(x) \rightarrow f(x)$  for  $\eta \rightarrow 0$ .*

Note that uniform convergence of  $f_\eta$  to  $f$  would require that there exists  $\eta > 0$  such that  $\sup_{x \in C} D_h(x, x^+)$  is finite, which we cannot guarantee in general (it does not hold for  $f = g = \frac{1}{2} \|\cdot\|^2$  defined on  $\mathbb{R}^d$  for instance). Denote  $\|x\|_A^2 = x^\top A x$ , then:

**Proposition 2.6** (Small step-sizes limit). *If  $f_\xi$  are  $L$ -relatively-smooth and  $f$  has a unique minimizer  $x_\star$  then:*

$$\lim_{\eta \rightarrow 0} \sigma_{\star,\eta}^2 = \lim_{\eta \rightarrow 0} \frac{1}{\eta^2} \mathbb{E} [D_h(x_\star^+, x_\star)] \quad (13)$$

$$= \frac{1}{2} \mathbb{E} \left[ \|\nabla f_\xi(x_\star)\|_{\nabla^2 h(x_\star)^{-1}}^2 \right]. \quad (14)$$

This variance is actually the best we can hope for in the Bregman setting, which indicates the relevance of Definition 2. Indeed, this term exactly correspond to the variance one would obtain when making infinitesimal SMD steps from  $x_\star$ , i.e., the norm of the stochastic gradients at optimum in the geometry given by  $\nabla^2 h(x_\star)^{-1}$ .

## 2.2. Standard Assumptions

We now compare Definition 2 with several variance assumptions from the literature. Note that they typically “only” require the bounds to hold for all iterates over the trajectory. However, in the absence of proof that the iterates stay in certain regions of the space, suprema over the whole domain are required for all variance definitions.

**Euclidean case.** Let us now take a step back and look at Euclidean case,  $h = \frac{1}{2}\|\cdot\|^2$ , and assume that  $f$  is  $L$ -smooth. Writing Equation (11) with this specific  $h$  and replacing  $x_\eta$  by a supremum, we obtain:

$$\sigma_{\star,\eta}^2 \leq \sup_{x \in C} \mathbb{E} \left[ \frac{1}{2} \|\nabla f(x) - \nabla f_\xi(x)\|^2 \right], \quad (15)$$

which is a common though debatable variance assumption. Indeed, it involves a maximum over the domain, and is in particular not bounded in general even for simple examples like Linear Regression. Yet, we can recover another standard variance assumption by assuming the smoothness of all  $f_\xi$  (e.g., Gower et al. (2019)), which writes:

$$\sigma_{\star,\eta}^2 \leq \mathbb{E} \left[ \|\nabla f_\xi(x_\star)\|^2 \right]. \quad (16)$$

This result is obtained by writing that  $\|\nabla f_\xi(x)\|^2 \leq 2\|\nabla f_\xi(x) - \nabla f_\xi(x_\star)\|^2 + 2\|\nabla f_\xi(x_\star)\|^2$ , and bounding the first term using smoothness. In particular, we see that standard Euclidean variance definitions are natural bounds of  $\sigma_{\star,\eta}^2$ . Detailed derivations can be found in Appendix B.

**Divergence between stochastic and deterministic gradients.** An early variance definition for stochastic mirror descent in the relative setting comes from Hanzely & Richtárik (2021), who define  $\sigma_{\text{sym}}^2$  as:

$$\sigma_{\text{sym}}^2 = \frac{1}{\eta} \sup_{x \in C} \left\langle \nabla f(x) - \nabla f_\xi(x), x^+ - \overline{x^+} \right\rangle \quad (17)$$

$$= \frac{1}{\eta^2} \sup_{x \in C} \left[ D_h \left( x^+, \overline{x^+} \right) + D_h \left( \overline{x^+}, x^+ \right) \right], \quad (18)$$

where we recall that  $\overline{x^+}$  is such that  $\nabla h \left( \overline{x^+} \right) = \nabla h(x) - \eta \nabla f(x)$ . This quantity is sometimes called the symmetrized Bregman Divergence.

We remark two main things when comparing  $\sigma_{\text{sym}}^2$  with Proposition 2.4: (i)  $\sigma_{\star,\eta}^2$  is not symmetrized, and contains only one of the two terms, and (ii) the bound only needs to hold at  $x_\eta$  instead of for all  $x \in C$ . As a result, we directly obtain that  $\sigma_{\star,\eta}^2 \leq \sigma_{\text{sym}}^2$ , and  $\sigma_{\text{sym}}^2$  is actually infinite in most cases, whereas  $\sigma_{\star,\eta}^2$  is usually finite, as seen above.

**Stochastic gradients at optimum.** Dragomir et al. (2021a)

define the variance as:

$$\sigma_{DEH}^2 = \sup_{x \in C} \frac{1}{2\eta^2} \mathbb{E} \left[ D_{h^*}(\nabla h(x) - 2\eta \nabla f_\xi(x_\star), \nabla h(x)) \right] \quad (19)$$

$$= \sup_{x \in C} \mathbb{E} \left[ \|\nabla f_\xi(x_\star)\|_{\nabla^2 h^*(z(x))}^2 \right], \quad (20)$$

where  $z(x) \in [\nabla h(x), \nabla h(x) - \eta \nabla f_\xi(x_\star)]$ . The main interest of this definition is that stochastic gradients are only taken at  $x_\star$ . In particular, this variance is 0 in case there is interpolation (all stochastic functions share a common minimum). However, this quantity can blow up if  $h$  is not strongly convex, since in this case  $\nabla^2 h^*$  is not upper bounded (indeed, smoothness of the conjugate is ensured by strong convexity of the primal function (Kakade et al., 2009)). Following similar derivations, but *after the supremum has been taken*, we arrive at:

**Proposition 2.7.** *If  $f$  is  $L$ -relatively-smooth w.r.t.  $h$ , then for  $\eta < 1/(2L)$ , the variance can be bounded as:*

$$\sigma_{\star,\eta}^2 \leq \mathbb{E} \left[ \|\nabla f_\xi(x_\star)\|_{\nabla^2 h^*(z_\eta)}^2 \right], \quad (21)$$

where  $z_\eta \in [\nabla h(x_\eta), \nabla h(x_\eta) - \eta \nabla f_\xi(x_\star)]$ .

In particular, we obtain a finite bound without having to restrict the space.

**Functions variance.** Another variance definition that appears in the SGD literature is of the form  $f(x_\star) - \mathbb{E} \left[ f_\xi(x_\star^\xi) \right]$ , using the optima of the stochastic functions (Loizou et al., 2021). Unfortunately, the results derived with this definition do not obtain a vanishing variance term when  $\eta \rightarrow 0$ , unlike most other variance definitions, and contrary to what is observed in practice, that smaller step-sizes reduce the variance. The vanishing variance term can be obtained by rescaling by  $1/\eta$  (so considering  $(f(x_\star) - \mathbb{E} \left[ f_\xi(x_\star^\xi) \right]) / \eta$  instead), but the problem is that this variance definition would explode for  $\eta \rightarrow 0$ . This is because using such a definition would come down to performing the supremum step within the expectation in (9), using that  $f_\xi(x^+) \geq f_\xi(x_\star^\xi)$ , which is a very crude bound. Instead, Corolary 2.3 directly shows that our variance definition is tighter than this one, and in particular (i) it is bounded for all  $\eta > 0$ , (ii) it remains finite as  $\eta \rightarrow 0$  even with the proper rescaling (Proposition 2.6).

**Relation to  $c$ -transform.** Léger & Aubin-Frankowski (2023) revisited the analysis of gradient descent, viewing it as an alternate minimization method on transforms of  $f$ . This point of view subsumes many methods, including the Newton Method or Mirror Descent. Central to their analysis is the notion of  $c$ -transform  $f^c(y) = \sup_{x \in C} f(x) - c(x, y)$ , a standard quantity from optimal

transport (Villani et al., 2009). It turns out that for  $\eta \leq 1/L$ ,  $f_\eta$  is actually linked to the  $c$ -transform as:

$$f_\eta(x) = \mathbb{E} [f_\xi^c(x^+)], \quad (22)$$

where we use the cost  $c(x, y) = \frac{1}{\eta} D_h(x, y)$ . Since  $f(x_\star) = f^c(x_\star) = \arg \min_{x \in C} f(\overline{x^+})$ , denoting  $\mathcal{T}_c(g) = g^c(\nabla h^*(\nabla h(x) - \eta \nabla g(x)))$ , we have that:

$$\sigma_{\star, \eta}^2 = \frac{1}{\eta} \left( \min_{x \in C} \mathcal{T}_c(\mathbb{E} [f_\xi])(x) - \min_{x \in C} \mathbb{E} [\mathcal{T}_c(f_\xi)](x) \right). \quad (23)$$

This alternative way of looking at  $\sigma_{\star, \eta}^2$  suggests that it can also be expressed as a ‘‘variance’’, similarly to (10), without resorting to the bound from Proposition 2.2, and incorporating the notion of  $c$ -transform instead.

In this section, we have highlighted the connections with other definitions, and argued that  $f_\eta$  (and its minimum) is a relevant quantity. In particular, Definition 2 is the only definition that allows boundedness of the variance notion both after a supremum step over the iterates (and without strong convexity of  $h$ ) and in the  $\eta \rightarrow 0$  limit with the proper rescaling.

### 3. Convergence Analysis

Now that we have (extensively) investigated  $\sigma_{\star, \eta}^2$ , and the various interpretations that come from different bounds, we are ready to state the convergence results. Some proofs in this section are just sketched, but complete derivations can be found in Appendix C.

#### 3.1. Relatively Strongly Convex setting.

Recall that  $f_\eta^\star = \inf_{x \in C} f_\eta(x)$ . Starting from an arbitrary  $x^{(0)}$ , the sequence  $(x^{(k)})_{k \geq 0}$  is built as  $x^{(k+1)} = (x^{(k)})^+$  for  $k \in \{0, T\}$  for some  $T > 0$

**Theorem 3.1.** *If  $f$  is  $\mu$ -relatively-strongly-convex with respect to  $h$ , under a constant step-size  $\eta$ , the iterates obtained by SMD (Equation (6)) verify*

$$\eta \left[ \mathbb{E} [f_\eta(x^{(T)})] - f_\eta^\star \right] + \mathbb{E} [D_h(x_\star, x^{(T+1)})] \leq (1 - \eta\mu)^{T+1} D_h(x_\star, x^{(0)}) + \frac{\eta \sigma_{\star, \eta}^2}{\mu}. \quad (24)$$

Note that the (relatively) strongly-convex theorem has a standard form, and recovers usual MD results if we remove the variance, and standard SGD results if we take  $h = \frac{1}{2} \|\cdot\|^2$ . Let us now proceed to the proof of Theorem 3.1.

*Proof.* We start from a standard equality, which is a varia-

tion of *e.g.*, Dragomir et al. (2021a, Lemma 4):

$$\mathbb{E} [D_h(x_\star, x^+)] - D_h(x_\star, x) + \eta D_f(x_\star, x) \quad (25)$$

$$= -\eta [f(x) - f(x_\star)] + \mathbb{E} [D_h(x, x^+)] \quad (26)$$

$$= \eta \left[ f(x_\star) - \left( f(x) - \frac{1}{\eta} \mathbb{E} [D_h(x, x^+)] \right) \right] \quad (27)$$

$$= \eta [f(x_\star) - f_\eta(x)] \quad (28)$$

$$= -\eta [f_\eta(x) - f_\eta^\star] + \eta [f(x_\star) - f_\eta^\star]. \quad (29)$$

Using that  $D_f(x_\star, x) \geq \mu D_h(x_\star, x)$  (relative strong convexity), and remarking that  $f(x_\star) - f_\eta^\star = \eta \sigma_{\star, \eta}^2$ , we obtain:

$$\eta [f_\eta(x) - f_\eta^\star] + \mathbb{E} [D_h(x_\star, x^+)] \leq (1 - \eta\mu) D_h(x_\star, x) + \eta^2 \sigma_{\star, \eta}^2. \quad (30)$$

At this point, we can neglect the  $\eta [f_\eta(x) - f_\eta^\star] \geq 0$  terms and chain the inequalities for  $x = x^{(t)}$  for  $t$  from 0 to  $T$  to obtain the result.  $\square$

This proof is quite simple, and naturally follows from Lemma C.1. One can also note that *relative smoothness of  $f$  is not required to obtain Theorem 3.1*, which has no condition on the step-size. This is not a typo, but reflects the fact that *step-size conditions are needed to obtain a bounded variance*. Indeed, the variance as defined here entangles aspects tied with the error due to discretization (which is usually dealt with using smoothness), and the error due to stochasticity. This is natural, as the stochastic noise vanishes in the continuous limit ( $\eta \rightarrow 0$ ). Besides, the magnitude of the updates depends both on where the stochastic gradient is applied and on the step-size. Yet, the simplicity of the proof is partly due to this entanglement, meaning that we have deferred some of the complexity to the bounding of the variance term.

Also note that Theorem 3.1 uses constant step-sizes, but Equation (30) can be used with time-varying step-sizes, as is done for instance in the proof of Theorem 4.3.

A variant of Theorem 3.1 in which the discretization error is partly removed from the notion of variance writes:

**Corollary 3.2.** *If  $f$  is  $\mu$ -strongly-convex and  $L$ -relatively-smooth with respect to  $h$ , if  $\eta \leq 1/L$ , the iterates obtained by SMD (Equation (6)) with constant step-size  $\eta$  verify*

$$\eta \left[ \mathbb{E} [f_\xi((x^{(T)})^+)] - f_+^\star \right] + \mathbb{E} [D_h(x_\star, x^{(T+1)})] \leq (1 - \eta\mu)^{T+1} D_h(x_\star, x^{(0)}) + \frac{\eta}{\mu} \left[ \frac{f(x_\star) - f_+^\star}{\eta} \right], \quad (31)$$

where  $f_+^\star = \inf_{x \in C} \mathbb{E} [f_\xi(x^+)]$ .

This alternate version is obtained by the key result in the proof of Proposition 2.2, *i.e.*,  $f_\eta(x) \geq \mathbb{E} [f_\xi(x^+)]$ , in line (28). Note that in the deterministic case,  $f_+^\star = f(x_\star)$ , and we recover standard results.

### 3.2. Convex setting.

Let us now consider the convex case, meaning that  $\mu = 0$ .

**Theorem 3.3.** *If  $f$  is convex, the iterates obtained by SMD using a constant step-size  $\eta > 0$  verify*

$$\begin{aligned} \frac{1}{T+1} \sum_{k=0}^T \mathbb{E} \left[ f_\eta(x^{(k)}) - f_\eta^* + D_f(x_*, x^{(k)}) \right] \\ \leq \frac{D_h(x_*, x^{(0)})}{\eta(T+1)} + \eta \sigma_{*,\eta}^2. \end{aligned}$$

This theorem is obtained by summing Equation (29) for  $x = x^{(k)}$  for all  $k \in \{1, \dots, T\}$  and rearranging the terms. Note that we choose a constant step-size for simplicity, but varying step-size results can be obtained in the same way.

This case differs from standard convex analyses, in that we obtain a control on  $f_\eta(x^{(k)}) - f_\eta^* + D_f(x_*, x^{(k)})$  instead of the usual  $f(x^{(k)}) - f(x_*)$ . One of the main consequences is that we cannot get a control on the average iterate since Bregman divergences are in general not convex in their second argument, and  $f_\eta$  is not necessarily convex. This non-standard result is a direct consequence of our choice of variance definition, but it is actually a quantity that naturally arises in the analysis. Note that a variant involving  $f_\eta^*$  can be obtained in the same lines as Corollary 3.2.

**Controlling  $f_\eta$ .** The results in this section do not directly control the function gap  $f(x) - f^*$ , but rather the transformed one  $f_\eta(x) - f_\eta^*$ . Yet, the continuity result (in  $\eta$ ) from Proposition 2.5 shows that the bounds we provide can still be interpreted as relevant function values for small  $\eta$ .

**Controlling  $D_f(x_*, x^{(k)})$ .** An interesting property of  $D_f(x_*, x^{(k)})$  is that it can be linked with the size of the gradients of  $f$ , as shown by the following result.

**Proposition 3.4.** *If  $\nabla f(x_*) = 0$  then for all  $x \neq x_*$ ,*

$$D_f(x_*, x) \geq LD_{h^*} \left( \nabla h(x_*) + \frac{\nabla f(x)}{L}, \nabla h(x_*) \right) > 0.$$

This is a Bregman equivalent of controlling the gradient squared norm, with the additional benefit that the reference point at which we apply the gradient is the optimum  $x_*$ . Besides, Proposition 3.4 shows that  $D_f(x_*, x) > 0$  for  $x \neq x_*$  without requiring  $f$  to be strictly convex (only  $h$ ).

**Minimal assumptions on  $h$ .** Note that the theorems in this section do not actually require  $h$  to satisfy Assumption 1, but only that iterations can be written in the form of Equation 6 (which is guaranteed by Assumption 1). While Assumption 1 allows for instance to use the Bregman co-coercivity lemma with any points, or ensures that  $\nabla^2 h$  is well-defined, which we leverage extensively in Section 2, our theorems are much more general than this, and include

applications such as proximal gradient mirror descent (next remark) or the MAP for Gaussian Parameters Estimation (next section).

**Stochastic Mirror Descent with a Proximal term.** We briefly discuss a generalization of the mirror descent iterations, which writes:

$$x^+ = \arg \min_{u \in C} \eta [\nabla f_\xi(x)^\top u + g(u)] + D_h(u, x), \quad (32)$$

where  $g$  is a proper lower semi-continuous convex function. This is a ‘‘proximal’’ version, which for instance corresponds to projected stochastic mirror descent if  $g$  is the indicator of a convex set. Let  $\omega \in \partial g(x^+)$  be such that  $\nabla h(x^+) = \nabla h(x) - \eta [\nabla f_\xi(x) + \omega]$ , then this iteration can be rewritten as  $\nabla h(x^+) + \eta \omega = \nabla h(x) + \eta \omega_x - \eta [\nabla f_\xi(x) + \omega_x]$  for any  $\omega_x \in \partial g(x)$ . In particular, (32) can be interpreted as a Stochastic mirror descent step with objective  $f_\xi + g$  and mirror  $h + \eta g$ . While the mirror does not satisfy Assumption 1 (and in particular twice differentiability in case  $g$  is the indicator of a set), the iterations can still be written in the form of Equation (6). In particular, the theorems from Section 3 still apply, with the adapted variance definition involving function  $f + g$  and mirror  $h + \eta g$ . Similarly,  $f + g$  is  $1/\eta$  relatively-smooth with respect to  $h + \eta g$  as long as  $f$  is  $L$ -relatively-smooth with respect to  $h$  and  $\eta \leq 1/L$ . More details can be found in Appendix E.

## 4. MAP For Gaussian Parameters Estimation.

So far, we have proposed new variance definitions for the analysis of stochastic mirror descent, and we have shown that they compare favorably to existing ones, while leading to simple convergence proofs. In this section, we investigate the open problem formulated by Le Priol et al. (2021), which is to find non-asymptotic convergence guarantees for the KL-divergence of the Maximum A Posteriori (MAP) estimator. In particular, this example will highlight the relevance of the infimum step on  $f_\eta$ , since it gives the first generic analysis that obtains meaningful finite time convergence rates.

### 4.1. MAP and MLE of exponential families.

In this section, we rapidly review the formalism of exponential family. More details can be found in Le Priol et al. (2021), and Wainwright et al. (2008, Chapter 3). Let  $X$  be a random variable, and  $T$  a deterministic function, then the density of an exponential family for a sample  $x$  writes:

$$p_\theta(x) = p(x|\theta) = \exp(\langle \theta, T(x) \rangle - A(\theta)), \quad (33)$$

where  $A$  is often referred to as the log-partition function. In this case,  $\theta$  is called the natural parameter, and  $T$  is the sufficient statistic. Function  $A$  is convex, and we can thus establish a form of duality through convex conjugacy. The



entropy thus writes:

$$A^*(\mu) = \max_{\theta' \in \Theta} \langle \mu, \theta' \rangle - A(\theta'). \quad (34)$$

Parameter  $\mu$  is called the *mean* parameter, and the standard MAP estimator can be derived for  $n_0 \in \mathbb{N}$ ,  $\mu_0 \in \mathbb{R}$  as:

$$\mu_{\text{MAP}}^{(n)} = \frac{n_0 \mu^{(0)} + \sum_{i=1}^n T(X_i)}{n_0 + n} \quad (35)$$

The Maximum Likelihood Estimator (MLE) corresponds to taking  $n_0 = 0$ . An interesting observation is that  $\mu_{\text{MAP}}^{(n)}$  can be obtained recursively for  $n > 0$ , as

$$\mu_{\text{MAP}}^{(0)} = \mu^{(0)}, \quad \eta_n = (n + n_0)^{-1}, \quad (36)$$

$$\mu_{\text{MAP}}^{(n+1)} = \mu_{\text{MAP}}^{(n)} - \eta_n \nabla g_{X_n}(\mu_{\text{MAP}}^{(n)}), \quad (37)$$

with  $\nabla g_{X_n}(\mu) = \mu - T(X_n)$ . In terms of primal variable  $\theta^{(n)} = \nabla A^*(\mu_{\text{MAP}}^{(n)})$ , (37) writes:

$$\nabla A(\theta^{(n+1)}) = \nabla A(\theta^{(n)}) - \eta \nabla f_{X_n}(\theta^{(n)}), \quad (38)$$

where  $f_{X_n}(\theta) = A(\theta) - \langle \theta, T(X_n) \rangle$ , so that  $f(\theta) = A(\theta) - \langle \theta, \mu_* \rangle$ . We recognize stochastic mirror descent iterations, with mirror  $A$  and stochastic gradients  $\nabla f_X$ . Similar results on the MLE can be obtained by taking  $n_0 = 0$ . This key observation implies that **convergence guarantees on the MAP and the MLE can be deduced from stochastic mirror descent convergence guarantees**.

While this seems to be a very appealing way to obtain convergence guarantees for the MAP estimator, Le Priol et al. (2021) observe that none of the existing convergence rates for stochastic mirror descent allow to obtain meaningful rates for the convergence of the MAP for general exponential families. In particular, none of them recover the  $O(1/n)$  asymptotic convergence rate for estimating a Gaussian with unknown mean and covariance.

This is due to the variance definitions used in the existing analyses, that all have issues (not uniformly bounded over the domain, not decreasing with the step-size...) as discussed in Section 2. Our analysis fixes this problem, and thus yields finite-time guarantees for the MAP and MLE estimators for the estimation of a Gaussian with unknown mean and covariance. This shows the relevance of Assumption 2.

## 4.2. Full Gaussian (unknown mean and covariance)

The main problem studied in Le Priol et al. (2021) is that of the one-dimensional full-Gaussian case, where the goal is to estimate the mean and covariance of a Gaussian from i.i.d. samples  $X_1, \dots, X_n \sim \mathcal{N}(m_*, \Sigma_*)$ , with  $\Sigma_* > 0$ . Note that although notation  $\Sigma$  is usually reserved for the covariance matrix of a multivariate Gaussian, we use it for

a scalar value here to highlight the distinction with  $\sigma_{*,\eta}^2$ , the variance from stochastic mirror descent. In this case, the sufficient statistics write  $T(X) = (X, X^2)$ , and the log-partition and entropy functions are, up to constants:

$$A(\theta) = \frac{\theta_1^2}{-4\theta_2} - \frac{1}{2} \log(-\theta_2), \quad A^*(\mu) = -\frac{1}{2} \log(\mu_2 - \mu_1^2),$$

for  $\theta \in \Theta = \mathbb{R} \times \mathbb{R}_*^*$  and  $\mu \in \{(u, v), u^2 < v\}$ . The goal is to estimate  $D_A(\theta, \theta_*)$ , for which Le Priol et al. (2021) show that only partial solutions exist: results are either asymptotic, or rely on the objective being (approximately) quadratic. Note that there is a relationship between natural parameters, mean parameters, and  $(m, \Sigma^2)$ , the mean and covariance of the Gaussian we would like to estimate. In the following, we will often abuse notations, and write for instance  $D_A(\tilde{\theta}, \theta)$  in terms of  $(m, \Sigma^2)$  and  $(\tilde{m}, \tilde{\Sigma}^2)$  rather than  $\theta$  and  $\tilde{\theta}$ . More specifically:

$$D_A(\tilde{\theta}, \theta) = -\frac{1}{2} \log \left( \frac{\Sigma^2}{\tilde{\Sigma}^2} \right) - \frac{\tilde{\Sigma}^2 - \Sigma^2}{2\tilde{\Sigma}^2} + \frac{(\tilde{m} - m)^2}{2\tilde{\Sigma}^2}.$$

The update formulas for the parameters are given by:

$$m^+ = (1 - \eta)m + \eta X, \quad (39)$$

$$(\Sigma^2)^+ = (1 - \eta) [\Sigma^2 + \eta(m - X)^2]. \quad (40)$$

Note that in our case,  $\nabla^2 A = \nabla^2 f_\xi$  for all  $\xi$ , and so in particular  $D_A = D_{f_\xi} = D_f$ . Therefore, we obtain that:

$$\begin{aligned} f_\eta(\theta) - f(\theta_*) &= \frac{1}{2\eta} \mathbb{E} \left[ \log \left( (1 - \eta) \left( 1 + \eta \frac{(m - X)^2}{\Sigma^2} \right) \right) \right] \\ &\quad - \frac{1}{2} \log \left( \frac{\Sigma_*^2}{\Sigma^2} \right). \end{aligned} \quad (41)$$

At this point, the main question is: *What can we lower bound  $f_\eta$  by?* We start by showing the following proposition:

**Proposition 4.1.** *The iterations (38) are well-defined for  $\eta < 1$  in the sense that if  $\theta^{(n)} \in \Theta = \mathbb{R} \times \mathbb{R}_*^*$ , then  $\nabla A(\theta^{(n)}) - \eta \nabla f_{X_n}(\theta^{(n)}) \in \text{Range}(\nabla A)$  almost surely, so that  $\theta^{(n+1)} \in \Theta$  is well-defined almost surely. Besides,  $f_\xi$  is 1-relatively-smooth and 1-relatively-strongly-convex with respect to  $A$ .*

This result is a direct consequence of the fact that  $D_{f_\xi} = D_A$  for all  $\xi$ , and the fact that  $\nabla A(\theta) - \eta \nabla f_{X_n}(\theta) = (1 - \eta) \nabla A(\theta) + \eta T(X_n) \in \{(u, v), u^2 < v\}$  if  $\nabla A(\theta) \in \{(u, v), u^2 < v\}$ . One can then deduce from this result that the variance of the mirror descent iterations  $\sigma_{*,\eta}^2$  is bounded for  $\eta < 1$ , thanks to Corolary 2.3. We now investigate the properties of the minimizer of  $f_\eta$  in this case, and first prove the following lemma:

**Lemma 4.2.** *Let  $(m_\eta, \Sigma_\eta^2)$  be the minimizer of  $f_\eta$ . Then, for  $\eta < 1/3$ :*

$$m_\eta = m_*, \quad \Sigma_*^2 \geq \Sigma_\eta^2 \geq (1 - 3\eta) \Sigma_*^2. \quad (42)$$

In particular, the variance  $\sigma_{\star,\eta}^2$  verifies:

$$\sigma_{\star,\eta}^2 \leq -\frac{1}{2\eta} \log(1 - 3\eta). \quad (43)$$

For  $1/3 < \eta \leq 1 - \varepsilon$ ,  $\sigma_{\star,\eta}^2 \leq c_\varepsilon$ , where  $c_\varepsilon$  is a numerical constant that only depends on  $\varepsilon$ .

Note that we show in this example that  $\Sigma_\eta^2$  is arbitrarily close to  $\Sigma_\star^2$  as  $\eta \rightarrow 0$ , which is expected.

**Theorem 4.3.** *The MAP estimator satisfies:*

$$D_A(\theta_\star, \theta^{(n)}) \leq \frac{n_0 D_A(\theta_\star, \theta^{(0)}) + \frac{3}{2} \log(1 + \frac{n+1}{n_0}) + \Gamma}{n + n_0},$$

where  $\Gamma \geq 0$  is a numerical constant and  $\Gamma = 0$  if  $n_0 > 3$ .

Note that the numerical constants are not optimized. The main benefit is that we obtain an anytime result on the convergence of the MAP estimator for all  $n \geq 0, n_0 \geq 1$  directly from the general Stochastic Mirror Descent convergence theorem. Yet, the open problem from [Le Priol et al. \(2021\)](#) is not completely solved still. We discuss these aspects below.

**Reverse KL bound.** We obtain a bound on  $D_A(\theta_\star, \theta^{(n)})$ , instead of  $D_A(\theta^{(n)}, \theta_\star) = f(\theta) - f(\theta_\star)$ . Note that the second term can be controlled asymptotically thanks to the bound on  $f_\eta(\theta^{(n)}) - f_\eta(\theta_\eta)$ , and  $f_\eta \rightarrow f$  when  $\eta = 1/n \rightarrow 0$ , but we might also be able to exploit this control over the course of the iterations.

**Asymptotic convergence.** Theorem 4.3 leads to a  $O(\log n/n)$  asymptotic convergence rate instead of the expected  $O(1/n)$  ([Le Priol et al., 2021](#)). This probably indicates that the globalization step (replacing  $\theta$  by  $\theta_\eta$ ), although leading to a finite variance, is too crude in this case. Indeed,  $\theta_n$  actually has a lot of structure in this example, since  $\nabla A(\theta_n) = \frac{1}{n} \sum_{k=1}^n T(X_k)$ . The SMD analysis is oblivious to this structure, hence the gap. Note that we can get rid of the  $\log n$  factor and recover the right  $O(1/n)$  rate from the same analysis by using a slightly different estimator than the MAP (or MLE). This is done by setting the step-size as  $\eta_n = \frac{2}{n+1}$  for  $n > 1$ , and the analysis of this variant follows [Lacoste-Julien et al. \(2012\)](#), as detailed in [Appendix F.3](#).

**The special case of the MLE.** The MLE corresponds to  $n_0 = 0$ , which is not handled in our analysis since the first step corresponds to  $\eta = 1$ , which necessarily results in  $\theta_2^{(1)} = -\infty$  (which corresponds to  $\Sigma^2 = 0$ , as can be seen from (40)). If we consider that mirror descent is run from  $\theta^{(1)}$ , then we obtain  $\mathbb{E}[D_A(\theta_\star, \theta^{(1)})] = \infty$  in general, where the expectation is over the value of the first sample drawn. Therefore, we need to start the SMD analysis at  $\theta^{(2)}$  to fit the MLE into this framework, and in particular we need to be able to evaluate  $\mathbb{E}[D_A(\theta_\star, \theta^{(2)})]$ . This is further discussed in [Appendix F.4](#).

## 5. Conclusion

This paper introduces a new notion of variance for the analysis of stochastic mirror descent. This notion, based on the fact that a certain function  $f_\eta$  admits a minimum, is less restrictive than existing ones, has the right asymptotic scaling with the step-size and is bounded regardless of the trajectory of the iterates without further assumptions.

We strongly believe that our analysis of SMD opens up new perspectives. As an example, we use our SMD results to show convergence of the MAP for estimating a Gaussian with unknown mean and covariance. As evidenced in [Le Priol et al. \(2021\)](#), all existing generic analyses of stochastic mirror descent failed to obtain such results.

Beyond SMD, this new way of looking at variance could prove useful for the analysis of many majorization minimization algorithms, including for instance (Euclidean) proximal SGD. While we have focused on stochastic gradients of the objective, considering stochastic mirror maps is an interesting problem that would widely extend the range of applications, and likely require even more general variance notions.

## 6. Acknowledgements

I would like to start by thanking Rémi Le Priol and Frederik Kunstner for introducing me to MAP problem for Gaussian Parameters Estimation, and the interesting discussions that followed.

Many thanks go to Radu Dragomir, Baptiste Goujaud and Adrien Taylor for their many comments and suggestions, which immensely helped me write this paper.

## References

- Bauschke, H. H., Borwein, J. M., et al. Legendre functions and the method of random bregman projections. *Journal of convex analysis*, 4(1):27–67, 1997.
- Bauschke, H. H., Bolte, J., and Teboulle, M. A descent lemma beyond lipschitz gradient continuity: first-order methods revisited and applications. *Mathematics of Operations Research*, 42(2):330–348, 2017.
- Bertero, M., Boccacci, P., Desiderà, G., and Vicidomini, G. Image deblurring with poisson data: from cells to galaxies. *Inverse Problems*, 25(12):123006, 2009.
- Bottou, L. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT’2010: 19th International Conference on Computational Statistics Paris France, August 22-27, 2010 Keynote, Invited and Contributed Papers*, pp. 177–186. Springer, 2010.
- Bubeck, S. et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Dragomir, R.-A. and Nesterov, Y. Convex quartic problems: homogenized gradient method and preconditioning. *arXiv preprint arXiv:2306.17683*, 2023.
- Dragomir, R. A., Even, M., and Hendriks, H. Fast stochastic bregman gradient methods: Sharp analysis and variance reduction. In *International Conference on Machine Learning*, pp. 2815–2825. PMLR, 2021a.
- Dragomir, R.-A., Taylor, A. B., d’Aspremont, A., and Bolte, J. Optimal complexity and certification of bregman first-order methods. *Mathematical Programming*, pp. 1–43, 2021b.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.
- Gower, R. M., Loizou, N., Qian, X., Sailanbayev, A., Shulgin, E., and Richtárik, P. Sgd: General analysis and improved rates. In *International conference on machine learning*, pp. 5200–5209. PMLR, 2019.
- Hanzely, F. and Richtárik, P. Fastest rates for stochastic mirror descent methods. *Computational Optimization and Applications*, 79:717–766, 2021.
- Hendriks, H., Bach, F., and Massoulié, L. Dual-free stochastic decentralized optimization with variance reduction. *Advances in neural information processing systems*, 33:19455–19466, 2020a.
- Hendriks, H., Xiao, L., Bubeck, S., Bach, F., and Massoulié, L. Statistically preconditioned accelerated gradient method for distributed optimization. In *International conference on machine learning*, pp. 4203–4227. PMLR, 2020b.
- Hoeven, D., Erven, T., and Kotłowski, W. The many faces of exponential weights in online learning. In *Conference On Learning Theory*, pp. 2067–2092. PMLR, 2018.
- Kakade, S., Shalev-Shwartz, S., Tewari, A., et al. On the duality of strong convexity and strong smoothness: Learning applications and matrix regularization. *Unpublished Manuscript*, <http://ttic.uchicago.edu/shai/papers/KakadeShalevTewari09.pdf>, 2(1):35, 2009.
- Kunstner, F., Kumar, R., and Schmidt, M. Homeomorphic-invariance of em: Non-asymptotic convergence in kl divergence for exponential families via mirror descent. In *International Conference on Artificial Intelligence and Statistics*, pp. 3295–3303. PMLR, 2021.
- Lacoste-Julien, S., Schmidt, M., and Bach, F. A simpler approach to obtaining an  $\mathcal{O}(1/t)$  convergence rate for the projected stochastic subgradient method. *arXiv preprint arXiv:1212.2002*, 2012.
- Le Priol, R., Kunstner, F., Scieur, D., and Lacoste-Julien, S. Convergence rates for the map of an exponential family and stochastic mirror descent—an open problem. *arXiv preprint arXiv:2111.06826*, 2021.
- Léger, F. and Aubin-Frankowski, P.-C. Gradient descent with a general cost. *arXiv preprint arXiv:2305.04917*, 2023.
- Loizou, N., Vaswani, S., Laradji, I. H., and Lacoste-Julien, S. Stochastic polyak step-size for sgd: An adaptive learning rate for fast convergence. In *International Conference on Artificial Intelligence and Statistics*, pp. 1306–1314. PMLR, 2021.
- Lu, H., Freund, R. M., and Nesterov, Y. Relatively smooth convex optimization by first-order methods, and applications. *SIAM Journal on Optimization*, 28(1):333–354, 2018.
- McMahan, B. Follow-the-regularized-leader and mirror descent: Equivalence theorems and  $\ell_1$  regularization. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 525–533. JMLR Workshop and Conference Proceedings, 2011.
- Nemirovskij, A. S. and Yudin, D. B. Problem complexity and method efficiency in optimization. 1983.

- Pfau, D. A generalized bias-variance decomposition for bregman divergences. *Unpublished Manuscript*, 2013.
- Shamir, O., Srebro, N., and Zhang, T. Communication-efficient distributed optimization using an approximate newton-type method. In *International conference on machine learning*, pp. 1000–1008. PMLR, 2014.
- Villani, C. et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Wainwright, M. J., Jordan, M. I., et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.

## A. Technical results on Bregman divergences

As for the rest of this paper, Assumption 1 is assumed throughout this section. However, some of these results hold even with less regularity, and in particular do not require second order continuous differentiability.

**Lemma A.1** (Duality). *For all  $x, y \in C$ , it holds that:*

$$D_h(x, y) = D_{h^*}(\nabla h(y), \nabla h(x)) \quad (44)$$

See, e.g. Bauschke et al. (1997, Theorem 3.7) for the proof.

**Lemma A.2** (Symmetrized Bregman). *For all  $x, y \in C$ , it holds that:*

$$D_h(x, y) + D_h(y, x) = \langle \nabla h(x) - \nabla h(y), x - y \rangle \quad (45)$$

The proof immediately follows from the definition of the Bregman divergence. The following result corresponds to Dragomir et al. (2021a, Lemma 3).

**Lemma A.3** (Bregman cocoercivity). *If a convex function  $f$  is  $L$ -relatively-smooth with respect to  $h$ , then for all  $\eta \leq 1/L$ ,*

$$D_{h^*}(\nabla h(x) - \eta[\nabla f(x) - \nabla f(y)], \nabla h(x)) \leq \eta D_f(x, y). \quad (46)$$

Denoting  $x^{+y} = \nabla h^*(\nabla h(x) - \eta[\nabla f(x) - \nabla f(y)])$ , a tighter result actually writes:

$$D_h(x, x^{+y}) + \eta D_f(x^{+y}, y) \leq \eta D_f(x, y). \quad (47)$$

The proof of the tighter version is simply obtained by not using that  $D_f(x^{+y}, y) \geq 0$  in the original proof. While we don't directly use it in this paper, it is sometimes useful. We now introduce the generalized bias-variance decomposition Lemma (Pfau, 2013, Theorem 0.1).

**Lemma A.4.** *If  $X$  is a random variable, then for all  $u \in C$ ,*

$$\mathbb{E}[D_{h^*}(X, u)] = D_{h^*}(\mathbb{E}[X], u) + D_{h^*}(X, \mathbb{E}[X]). \quad (48)$$

## B. Missing results on the variances

We start this section by proving the following lemma, which in particular ensures that  $D_h(x, x^+)/\eta$  increases with  $\eta$  (and so decreases as  $\eta \rightarrow 0$ ).

**Lemma B.1.** *Let  $\phi_\xi : \eta \mapsto \frac{1}{\eta} D_h(x, x^+(\eta, \xi))$ . Then,  $\nabla \phi_\xi(\eta) = \frac{1}{\eta^2} D_h(x^+(\eta, \xi), x) \geq 0$ .*

*Proof.* First remark that since  $\nabla h(x^+) = \nabla h(x) - \eta \nabla f_\xi(x)$ , we can write

$$\begin{aligned} \nabla_\eta [D_h(x, x^+)] &= \nabla_\eta [h(x) - h(x^+) - \nabla h(x^+)^\top (x - x^+)] \\ &= -\nabla h(x^+)^\top \nabla_\eta x^+ + \nabla f_\xi(x)^\top (x - x^+) + \nabla h(x^+)^\top \nabla_\eta x^+ \\ &= \nabla f_\xi(x)^\top (x - x^+) \\ &= \frac{1}{\eta} (\nabla h(x) - \nabla h(x^+))^\top (x - x^+) = \frac{D_h(x, x^+) + D_h(x^+, x)}{\eta}. \end{aligned}$$

Then, the expression follows from

$$\nabla \phi_\xi(\eta) = \nabla_\eta \left[ \frac{1}{\eta} D_h(x, x^+) \right] = \frac{1}{\eta} \nabla_\eta [D_h(x, x^+)] - \frac{1}{\eta^2} D_h(x, x^+) = \frac{1}{\eta^2} D_h(x^+, x). \quad (49)$$

□

*Proof of Proposition 2.5.* We now prove that  $f_\eta \rightarrow f$  when  $\eta \rightarrow 0$ . To show this, we note that for any fixed  $x \in \text{int } C$ :

- For any fixed  $\xi$ ,  $\frac{1}{\eta}D_h(x, x^+) = \frac{\eta}{2}\|\nabla f_\xi(x)\|_{\nabla^2 h^*(z)}^2$  for  $z \in [\nabla h(x), \nabla h(x) - \eta\nabla f_\xi(x)]$ . Therefore,  $\frac{1}{\eta}D_h(x, x^+) \rightarrow 0$  for  $\eta \rightarrow 0$  since  $\nabla^2 h^*(\nabla h(x)) = (\nabla^2 h(x))^{-1} < \infty$  by strict convexity of  $h$ .
- Let  $\eta \leq \eta_0$ . Then, for all  $\xi$ ,  $\frac{1}{\eta}D_h(x, x^+(\eta, \xi)) \leq \frac{1}{\eta_0}D_h(x, x^+(\eta_0, \xi))$  since the function  $\eta \mapsto \frac{1}{\eta}D_h(x, x^+(\eta, \xi))$  is an increasing function (positive gradient using Lemma B.1).
- $\frac{1}{\eta_0}\mathbb{E}[D_h(x, x^+(\eta_0, \xi))]$  is finite.

Then, using the dominated convergence theorem, we obtain that we can invert the integral (expectation) and the limit, so that  $\lim_{\eta \rightarrow 0} \mathbb{E} \frac{1}{\eta}D_h(x, x^+) = \mathbb{E} \lim_{\eta \rightarrow 0} \frac{1}{\eta}D_h(x, x^+) = 0$ .  $\square$

*Proof of Proposition 2.6.* We prove this result by successively upper bounding and lower bounding  $\sigma_{\star, \eta}^2$ , and making  $\eta \rightarrow 0$ .

1 - *Upper bound on  $\sigma_{\star, \eta}^2$ .* One side is direct, by writing that  $f(x_\eta) \geq f(x_\star)$ :

$$\sigma_{\star, \eta}^2 = \frac{1}{\eta} \left( f(x_\star) - f(x_\eta) + \frac{1}{\eta} \mathbb{E} [D_h(x_\eta, x_\eta^+)] \right) \leq \frac{1}{\eta^2} \mathbb{E} [D_h(x_\eta, x_\eta^+)]. \quad (50)$$

From the proof of Proposition 2.5 we have pointwise convergence of  $f_\eta$  to  $f$ . Since  $f$  is convex and has a unique minimizer  $x_\star$ , then  $x_\eta \rightarrow x_\star$  for  $\eta \rightarrow 0$ , which leads to the result.

2 - *Lower bound on  $\sigma_{\star, \eta}^2$ .* By definition of  $x_\eta$  as the minimizer of  $f_\eta$ , we have  $f_\eta(x_\eta) \leq f_\eta(x_\star)$ , and so:

$$\sigma_{\star, \eta}^2 = \frac{f(x_\star) - f_\eta(x_\eta)}{\eta} \geq \frac{f(x_\star) - f_\eta(x_\star)}{\eta} = \frac{1}{\eta^2} \mathbb{E} [D_h(x_\star, x_\star^+)]. \quad (51)$$

$\square$

Let us now prove the following proposition, which follows the proof from Dragomir et al. (2021a).

*Proof of Proposition 2.7.* Let us prove that  $\sigma_{\star, \eta}^2 \leq \mathbb{E} [\|\nabla f_\xi(x_\star)\|_{\nabla^2 h^*(z_\eta)}^2]$ . We start by

$$D_h(x, x^+) = D_{h^*}(\nabla h(x) - \eta\nabla f_\xi(x), \nabla h(x)) \quad (52)$$

$$= D_{h^*}(\nabla h(x) - \eta[\nabla f_\xi(x) - \nabla f_\xi(x_\star)] - \eta\nabla f_\xi(x_\star), \nabla h(x)) \quad (53)$$

$$= D_{h^*}\left(\frac{(\nabla h(x) - 2\eta[\nabla f_\xi(x) - \nabla f_\xi(x_\star)]) + (\nabla h(x) - 2\eta\nabla f_\xi(x_\star))}{2}, \nabla h(x)\right). \quad (54)$$

Using the convexity of  $D_{h^*}$  in its first argument and then the Bregman cocoercivity lemma, we obtain for  $\eta \leq 1/2L$ :

$$D_h(x, x^+) \leq \frac{1}{2}D_{h^*}(\nabla h(x) - 2\eta[\nabla f_\xi(x) - \nabla f_\xi(x_\star)], \nabla h(x)) + \frac{1}{2}D_{h^*}(\nabla h(x) - 2\eta\nabla f_\xi(x_\star), \nabla h(x)) \quad (55)$$

$$\leq \eta D_{f_\xi}(x, x_\star) + \frac{1}{2}D_{h^*}(\nabla h(x) - 2\eta\nabla f_\xi(x_\star), \nabla h(x)). \quad (56)$$

Using that  $\mathbb{E} [D_{f_\xi}(x, x_\star)] = D_f(x, x_\star)$  and applying this to  $x = x_\eta$ , we obtain

$$\begin{aligned} \sigma_{\star, \eta}^2 &= \frac{f(x_\star) - f_\eta(x_\eta)}{\eta} \\ &= \frac{\mathbb{E} [D_h(x_\eta, x_\eta^+)] + \eta f(x_\star) - \eta f(x_\eta)}{\eta^2} \\ &\leq \frac{1}{2\eta^2} \mathbb{E} [D_{h^*}(\nabla h(x_\eta) - 2\eta\nabla f_\xi(x_\star), \nabla h(x_\eta))] + \frac{D_f(x_\eta, x_\star) + f(x_\star) - f(x_\eta)}{\eta} \\ &\leq \frac{1}{2\eta^2} \mathbb{E} [D_{h^*}(\nabla h(x_\eta) - 2\eta\nabla f_\xi(x_\star), \nabla h(x_\eta))] = \frac{1}{2\eta^2} \times \mathbb{E} \left[ \frac{1}{2} \|\nabla f_\xi(x_\star)\|_{\nabla^2 h^*(z_\eta)}^2 \right], \end{aligned}$$

and the result follows. The last inequality comes from the fact that if  $x_\star = \arg \min_{x \in C} f(x)$ , then  $-\nabla f(x_\star)$  is normal to  $C$  so  $-\nabla f(x_\star)^\top (x_\eta - x_\star) \leq 0$ .  $\square$

### C. Convergence results.

In this section, we detail the proofs of the various convergence theorems that were only sketched in the main text. We start by proving the first identity, which is a variation of *e.g.*, Dragomir et al. (2021a, Lemma 4), which we detail here for the sake of completeness.

**Lemma C.1.** *Let  $x^+ \in C$  be such that  $\nabla h(x^+) = \nabla h(x) - \eta \nabla f_\xi(x)$ , with  $f_\xi$  a random differentiable function such that  $\mathbb{E}[f_\xi] = f$ . Then, for all points  $y \in C$ ,*

$$\mathbb{E}[D_h(y, x^+)] - D_h(y, x) + \eta D_f(y, x) = -\eta[f(x) - f(y)] + \mathbb{E}[D_h(x, x^+)] \quad (57)$$

In particular, we can apply the result to  $y = x_*$ .

*Proof.* We give a slightly different proof than Dragomir et al. (2021a), and in particular this version of the identity is slightly more direct (though maybe less insightful) and does not require  $\nabla f(y) = 0$ . We write:

$$\begin{aligned} \mathbb{E}[D_h(y, x^+)] &= \mathbb{E}[h(y) - h(x^+) - \nabla h(x^+)^\top (y - x^+)] \\ &= \mathbb{E}[h(y) - h(x^+) - \nabla h(x^+)^\top (y - x) - \nabla h(x^+)^\top (x - x^+)] \\ &= \mathbb{E}[h(y) - h(x) - \nabla h(x)^\top (y - x) + \eta \nabla f_\xi(x)^\top (y - x) - \nabla h(x^+)^\top (x - x^+) + h(x) - h(x^+)] \\ &= D_h(y, x) + \eta \nabla f(x)^\top (y - x) + \mathbb{E}[D_h(x, x^+)] \\ &= D_h(y, x) - \eta D_f(y, x) + \eta[f(y) - f(x)] + \mathbb{E}[D_h(x, x^+)]. \end{aligned}$$

□

*Proof of Corollary 3.2.* We start back from Equation (28), and write, using that  $f_\eta(x) \geq \mathbb{E}[f_\xi(x^+)]$  (proof of Proposition 2.2):

$$\mathbb{E}[D_h(x_*, x^+)] - D_h(x_*, x) + \eta D_f(x_*, x) = \eta[f(x_*) - f_\eta(x)] \quad (58)$$

$$\leq \eta[f(x_*) - \mathbb{E}[f_\xi(x^+)]] \quad (59)$$

$$\leq -\eta[\mathbb{E}[f_\xi(x^+)] - f_*^+] + \eta^2 \left( \frac{f(x_*) - f_*^+}{\eta} \right). \quad (60)$$

The result follows naturally from using the relative strong convexity of  $f$ , leading to:

$$\eta[\mathbb{E}[f_\xi(x^+)] - f_*^+] + D_h(x_*, x^+) \leq (1 - \eta\mu)D_h(x_*, x) + \eta^2 \left[ \frac{f(x_*) - f_*^+}{\eta} \right]. \quad (61)$$

Then, we chain iterations as done for Theorem 3.1

□

*Proof of Theorem 3.3.* We also start from the same result as above, and write it for  $x = x^{(k)}$ , so that  $x^+ = x^{(k+1)}$ :

$$\mathbb{E}[D_h(x_*, x^{(k+1)})] - D_h(x_*, x^{(k)}) + \eta D_f(x_*, x^{(k)}) = \eta[f(x_*) - f_\eta(x^{(k)})] \leq -\eta[f_\eta(x^{(k)}) - f_\eta(x_\eta)] + \eta^2 \sigma_{*,\eta}^2. \quad (62)$$

Moving the  $f_\eta$  terms to the left, and summing this for  $k = 0$  to  $T$  leads to:

$$\eta \sum_{k=0}^T [f_\eta(x^{(k)}) - f_\eta(x_\eta) + D_f(x_*, x^{(k)})] \leq D_h(x_*, x^{(0)}) - D_h(x_*, x^{(k+1)}) + T\eta^2 \sigma_{*,\eta}^2. \quad (63)$$

The final result is obtained by dividing by  $\eta T$ , and the fact that  $D_h(x_*, x^{(k+1)}) \geq 0$ .

□

*Proof of Proposition 3.4.* We use Bregman cocoercivity (Lemma A.3) with  $\eta = \frac{1}{L}$  between  $x_*$  and  $x$  (instead of  $x$  and  $x_*$  as it had been done previously), which directly leads to:

$$D_{h^*} \left( \nabla h^*(x_*) - \frac{1}{L} [\nabla f(x_*) - \nabla f(x)] \right) \leq \frac{1}{L} D_f(x_*, x). \quad (64)$$

The first part of the proposition follows from the fact that  $\nabla f(x_\star) = 0$ . For the rest proof, we start with Inequality (32), which gives:

$$\begin{aligned} 0 &= D_{h^\star}(\nabla h(x_\star) - \frac{1}{L}\nabla f(x), \nabla h(x_\star)) \\ &= D_h\left(x_\star, \nabla h^\star\left(\nabla h(x_\star) - \frac{1}{L}\nabla f(x)\right)\right). \end{aligned}$$

At this point, strict convexity of  $h$  leads to  $\nabla h^\star(\nabla h(x_\star) - \frac{1}{L}\nabla f(x)) = x_\star$ , so that  $\nabla f(x) = 0$  by applying  $\nabla h$  on both sides.  $\square$

## D. Variation on the convex case

In this section, we quickly illustrate that the result we obtain is tightly linked to the notion of variance that we define. As an example, a variation of Theorem 3.3 can be obtained with a control on  $f(x) - f(x_\star)$ , but this requires a different notion of variance:

**Theorem D.1.** *If  $f$  is convex, the iterates obtained by SMD using a constant step-sizes  $\eta > 0$  verify*

$$f\left(\frac{1}{T}\sum_{k=0}^T x^{(k)}\right) - f(x_\star) \leq \frac{D_h(x_\star, x^{(0)})}{\eta T} + \eta \tilde{\sigma}_{\star, \eta}^2, \quad (65)$$

where

$$\tilde{\sigma}_{\star, \eta}^2 = \frac{1}{\eta} \max_{x \in C} \left\{ \frac{1}{\eta} \mathbb{E} [D_h(x, x^+)] - D_f(x_\star, x) \right\}. \quad (66)$$

Note that this alternative variance definition can be unbounded even when  $\sigma_{\star, \eta}^2$  is bounded, as is the case for instance in the Gaussian MAP example. Besides, it does not inherit from most of the good properties of  $\sigma_{\star, \eta}^2$  presented in Section 2, and cannot be compared to the other standard variance notions. The main case in which this alternative definition makes sense is the Euclidean case, in which  $\tilde{\sigma}_{\star, \eta}^2$  can be bounded using cocoercivity.

*Proof of Theorem D.1.* This proof directly starts from Lemma C.1:

$$\mathbb{E} [D_h(x_\star, x^{(k+1)})] = D_h(x_\star, x^{(k)}) - \eta D_f(x_\star, x^{(k)}) - \eta [f(x^{(k)}) - f(x_\star)] + \mathbb{E} [D_h(x^{(k)}, (x^{(k)})^+)] \quad (67)$$

$$= D_h(x_\star, x^{(k)}) - \eta [f(x^{(k)}) - f(x_\star)] + \eta \left[ \frac{1}{\eta} \mathbb{E} [D_h(x^{(k)}, (x^{(k)})^+)] - D_f(x_\star, x^{(k)}) \right] \quad (68)$$

$$\leq D_h(x_\star, x^{(k)}) - \eta [f(x^{(k)}) - f(x_\star)] + \eta^2 \tilde{\sigma}_{\star, \eta}^2. \quad (69)$$

Summing this for  $k = 0$  to  $T$ , and dividing by  $\eta T$  we obtain:

$$\frac{1}{T} \sum_{k=0}^T f(x^{(k)}) - f(x_\star) \leq \frac{D_h(x_\star, x^{(0)})}{\eta T} + \eta \tilde{\sigma}_{\star, \eta}^2 \quad (70)$$

The result on the average iterate then follows from convexity of  $f$ .  $\square$

## E. Stochastic Mirror Descent with a Proximal term

We are interested in this section in a variation of the original problem, where we would like to solve the following problem:

$$\min_{x \in C} f(x) + g(x), \quad (71)$$

where  $g$  is a convex proper lower semi-continuous function (but not necessarily differentiable). This problem can be solved using the following stochastic proximal mirror descent algorithm:

$$x^+ = \arg \min_{u \in C} g(u) + \nabla f_\xi(x)^\top u + \frac{1}{\eta} D_h(u, x). \quad (72)$$



In this case, the iterations write:

$$\nabla h(x^+) = \nabla h(x) - \eta [\nabla f_\xi(x) + \omega] \quad (73)$$

where  $\omega \in \partial g(x^+)$ . We now prove an equivalent for Lemma C.1.

**Lemma E.1.** *Let  $x^+ \in C$  be such that  $\nabla h(x^+) = \nabla h(x) - \eta [\nabla f_\xi(x) + \omega]$ , with  $f_\xi$  a random differentiable function such that  $\mathbb{E}[f_\xi] = f$  and  $\omega \in \partial g(x^+)$  where  $g$  is a convex proper lower semi-continuous function. Then, for all  $y \in C \cap \text{dom}g$ ,*

$$\mathbb{E}[D_h(y, x^+)] = D_h(y, x) - \eta D_f(y, x) - \eta [f(x) - f(y)] + \mathbb{E}[D_h(x, x^{+f}) - D_h(x^+, x^{+f})] + \eta \omega^\top (y - x^+), \quad (74)$$

where  $x^{+f}$  is the point such that  $\nabla h(x^{+f}) = \nabla h(x) - \eta \nabla f_\xi(x)$ .

*Proof.* We write:

$$\begin{aligned} D_h(y, x^+) &= h(y) - h(x^+) - \nabla h(x^+)^\top (y - x^+) \\ &= h(y) - h(x^{+f}) - \nabla h(x^{+f})^\top (y - x^+) + \eta \omega^\top (y - x^+) - h(x^+) + h(x^{+f}) \\ &= D_h(y, x^{+f}) - \nabla h(x^{+f})^\top (x^{+f} - x^+) + \eta \omega^\top (y - x^+) - h(x^+) + h(x^{+f}) \\ &= D_h(y, x^{+f}) - D_h(x^+, x^{+f}) + \eta \omega^\top (y - x^+) \end{aligned}$$

The result follows from applying Lemma C.1 to  $D_h(y, x^{+f})$ .  $\square$

Note that by abuse of notation, if we denote  $D_g(y, x^+) = g(y) - g(x^+) - \omega_x^\top (y - x^+)$ , and  $D_g(y, x) = g(y) - g(x) - \omega_x^\top (y - x)$  for any  $\omega_x \in \partial g(x)$ , then with a few lines of computations, and noting in particular that  $D_h(x, x^{+f}) - D_h(x^+, x^{+f}) = D_h(x, x^+) - [\nabla h(x^{+f}) - \nabla h(x^+)]^\top (x - x^+)$  we recover exactly the result of Lemma C.1 applied to the iterations in which we take (sub)-gradients of  $f + g$  with mirror  $h + \eta g$ . This is not surprising and justifies the remark about Stochastic Mirror Descent with a proximal term written out in the main text. In particular, Theorems 3.1 and 3.3 can be transposed directly to the composite  $(f + g)$  setting by simply defining generalized Bregman divergences where the gradient parts are replaced by the subgradients picked in the actual SMD steps.

While  $h + \eta g$  does not necessarily satisfy Assumption 1, the key point is that iterations can be written in the form of Equation (73), which is the case for instance if  $g$  is the indicator of a convex set.

Note that Corollary 3.2 also holds in the same way, since relative smoothness is only needed to obtain that  $\eta D_{f_\xi+g}(x, x^+) \leq D_{h+\eta g}(x, x^+)$ , which is equivalent to  $\eta D_{f_\xi}(x, x^+) \leq D_h(x, x^+)$ , which holds by  $L$ -relative smoothness of  $f$  with respect to  $h$  for  $\eta \leq 1/L$ .

## F. Gaussian case with unknown covariance.

In this section, we prove the various results for Gaussian estimation with unknown mean and covariance. For the sake of brevity, we only prove the propositions, and refer the interested reader to, *e.g.*, Le Priol et al. (2021) for standard results about the setting.

### F.1. Instanciation in the Stochastic mirror descent setting

We first write what the various divergences are in our setting, together with the mirror updates and finally the form of  $f_\eta$ . Following Le Priol et al. (2021, Section 4.2), we write that:

$$\theta_1 = \frac{m}{\Sigma^2}, \quad \theta_2 = -\frac{1}{2\Sigma^2}. \quad (75)$$

This allows us to express  $A(\theta)$  in terms of  $(m, \Sigma^2)$ :

$$A(\theta) = -\frac{1}{2} \log(-\theta_2) - \frac{\theta_1^2}{4\theta_2} = \frac{1}{2} \log(2\Sigma^2) + \frac{1}{2} \frac{m^2}{\Sigma^2} \quad (76)$$

**Proposition F.1.** *The Bregman divergence with respect to  $\tilde{\theta}, \theta$  writes:*

$$D_A(\tilde{\theta}, \theta) = -\frac{1}{2} \log\left(\frac{\Sigma^2}{\tilde{\Sigma}^2}\right) - \frac{\tilde{\Sigma}^2 - \Sigma^2}{2\tilde{\Sigma}^2} + \frac{(\tilde{m} - m)^2}{2\tilde{\Sigma}^2}. \quad (77)$$

*Proof.* We know that  $\nabla A(\theta) = \mu = (m, m^2 + \Sigma^2)$ . Therefore,

$$\nabla A(\theta)^\top (\tilde{\theta} - \theta) = m \left( \frac{\tilde{m}}{\tilde{\Sigma}^2} - \frac{m}{\Sigma^2} \right) - \frac{1}{2} (m^2 + \Sigma^2) \left( \frac{1}{\tilde{\Sigma}^2} - \frac{1}{\Sigma^2} \right) \quad (78)$$

$$= \frac{m\tilde{m}}{\tilde{\Sigma}^2} - \frac{m^2}{2\Sigma^2} - \frac{m^2}{2\tilde{\Sigma}^2} - \frac{1}{2} \left( \frac{\Sigma^2}{\tilde{\Sigma}^2} - 1 \right) \quad (79)$$

$$= -\frac{(m - \tilde{m})^2}{2\tilde{\Sigma}^2} + \frac{\tilde{m}^2}{2\tilde{\Sigma}^2} - \frac{m^2}{2\Sigma^2} - \frac{\Sigma^2 - \tilde{\Sigma}^2}{2\tilde{\Sigma}^2}. \quad (80)$$

Using Equation (76), we obtain:

$$\begin{aligned} D_A(\tilde{\theta}, \theta) &= A(\tilde{\theta}) - A(\theta) - \nabla A(\theta)^\top (\tilde{\theta} - \theta) \\ &= \frac{1}{2} \log(2\tilde{\Sigma}^2) - \frac{1}{2} \log(2\Sigma^2) + \frac{\Sigma^2 - \tilde{\Sigma}^2}{2\tilde{\Sigma}^2} + \frac{(m - \tilde{m})^2}{2\tilde{\Sigma}^2}, \end{aligned}$$

which finishes the proof.  $\square$

In the Gaussian with unknown covariance, the sufficient statistics are:

$$T(X) = (X, X^2), \quad (81)$$

where  $x \in \mathbb{R}$  is an observation drawn from  $\mathcal{N}(m_\star, \Sigma_\star)$ .

Let us now prove the proposition on the updates, which corresponds to Equations (39) and (40):

**Proposition F.2.** *In  $(m, \Sigma^2)$  parameters, the updates write:*

$$m^+ = (1 - \eta)m + \eta X, \quad (82)$$

$$(\Sigma^2)^+ = (1 - \eta) [\Sigma^2 + \eta(m - X)^2]. \quad (83)$$

*Proof.* Since the (stochastic) gradients write  $g(\mu) = \mu - T(X)$ , the iterations are defined by:

$$\mu_1^+ = (1 - \eta)\mu_1 + \eta X \quad (84)$$

$$\mu_2^+ = (1 - \eta)\mu_2 + \eta X^2. \quad (85)$$

Since  $(\mu_1, \mu_2) = (m, m^2 + \Sigma^2)$ , the update on  $m$  is immediate. For the update on  $\Sigma^2$ , we write:

$$\begin{aligned} (\Sigma^2)^+ &= \mu_2^+ - (m^+)^2 \\ &= (1 - \eta)\mu_2 + \eta X^2 - ((1 - \eta)m + \eta X)^2 \\ &= (1 - \eta)\Sigma^2 + (1 - \eta)m^2 + \eta X^2 - (1 - \eta)^2 m^2 - 2\eta(1 - \eta)Xm - \eta^2 X^2 \\ &= (1 - \eta)\Sigma^2 + \eta(1 - \eta)(m - X)^2. \end{aligned}$$

$\square$

We can now proceed to proving the form of  $f_\eta$ . We first start by writing that:

$$f(\theta) = A(\theta) - \theta^\top (m_\star, m_\star^2 + \Sigma_\star^2) \quad (86)$$

$$= \frac{1}{2} \log(2\Sigma^2) + \frac{1}{2} \frac{m^2}{\Sigma^2} - \frac{mm_\star}{\Sigma^2} + \frac{m_\star^2 + \Sigma_\star^2}{2\Sigma^2}. \quad (87)$$

Therefore,

$$f(\theta) = \frac{1}{2} \log(2\Sigma^2) + \frac{\Sigma_\star^2}{2\Sigma^2} + \frac{(m - m_\star)^2}{2\Sigma^2} \quad (88)$$

In particular,

$$f(\theta) - f(\theta_\star) = \frac{1}{2} \log \left( \frac{\Sigma^2}{\Sigma_\star^2} \right) + \frac{\Sigma_\star^2 - \Sigma^2}{2\Sigma^2} + \frac{(m - m_\star)^2}{2\Sigma^2} \quad (89)$$

Note that, as expected, this corresponds to  $D_A(\theta, \theta_*)$ , that we can also compute through Proposition F.1. We now write:

$$D_A(\theta, \theta^+) = -\frac{1}{2} \log \left( \frac{(\Sigma^2)^+}{\Sigma^2} \right) - \frac{\Sigma^2 - (\Sigma^2)^+}{2\Sigma^2} + \frac{(m - m^+)^2}{2\Sigma^2} \quad (90)$$

$$= -\frac{1}{2} \log \left( (1 - \eta) \left[ 1 + \eta \frac{(m - X)^2}{\Sigma^2} \right] \right) + \frac{(1 - \eta)(\Sigma^2 + \eta(m - X)^2) - \Sigma^2}{2\Sigma^2} + \frac{\eta^2(m - X)^2}{2\Sigma^2} \quad (91)$$

$$= -\frac{1}{2} \log \left( (1 - \eta) \left[ 1 + \eta \frac{(m - X)^2}{\Sigma^2} \right] \right) - \frac{\eta}{2} + \eta(1 - \eta) \frac{(m - X)^2}{2\Sigma^2} + \frac{\eta^2(m - X)^2}{2\Sigma^2} \quad (92)$$

$$= -\frac{1}{2} \log \left( (1 - \eta) \left[ 1 + \eta \frac{(m - X)^2}{\Sigma^2} \right] \right) - \frac{\eta}{2} + \eta \frac{(m - X)^2}{2\Sigma^2}. \quad (93)$$

Therefore,

$$f(\theta) - \frac{D_A(\theta, \theta^+)}{\eta} - f(\theta_*) \quad (94)$$

$$= \frac{1}{2} \log \left( \frac{\Sigma^2}{\Sigma_*^2} \right) + \frac{\Sigma_*^2 - \Sigma^2}{2\Sigma^2} + \frac{(m - m_*)^2}{2\Sigma^2} + \frac{1}{2\eta} \log \left( (1 - \eta) \left[ 1 + \eta \frac{(m - X)^2}{\Sigma^2} \right] \right) + \frac{1}{2} - \frac{(m - X)^2}{2\Sigma^2} \quad (95)$$

$$= \frac{1}{2} \log \left( \frac{\Sigma^2}{\Sigma_*^2} \right) + \frac{1}{2\eta} \log \left( (1 - \eta) \left[ 1 + \eta \frac{(m - X)^2}{\Sigma^2} \right] \right) + \frac{\Sigma_*^2}{2\Sigma^2} + \frac{(m - m_*)^2}{2\Sigma^2} - \frac{(m - X)^2}{2\Sigma^2}. \quad (96)$$

Finally,  $\mathbb{E}[(m - X)^2] = (m - m_*)^2 + \Sigma_*^2$ , and so:

$$f_\eta(\theta) - f(\theta_*) = \frac{1}{2} \log \left( \frac{\Sigma^2}{\Sigma_*^2} \right) + \frac{1}{2\eta} \mathbb{E} \left[ \log \left( (1 - \eta) \left[ 1 + \eta \frac{(m - X)^2}{\Sigma^2} \right] \right) \right], \quad (97)$$

which precisely corresponds to Equation (41). We now proceed to proving bounds on  $\theta_\eta$  for  $\eta < 1$ .

## F.2. Bounding the stochastic mirror descent variance $\sigma_{*,\eta}^2$ .

Now that we have an explicit form for  $f_\eta$ , we can characterize its minimizer  $\theta_\eta$ , and use this to prove results on  $f_\eta(\theta_\eta)$ , which will in turn lead to bounds on  $\sigma_{*,\eta}^2$ . This is the core of Lemma 4.2.

*Proof.* Proof of Lemma 4.2. The proof will proceed in three different stages:

- Differentiating  $f_\eta$  with respect to  $m$  and  $\Sigma^2$ .
- Using these expressions to obtain bounds on the  $(m_\eta, \Sigma_\eta^2)$  for which  $\nabla f_\eta$  is 0.
- Plugging these bounds into the expression of  $f_\eta$  to bound  $\Sigma_\eta^2$ .

**1 - Differentiating  $f_\eta$ .** Before differentiating, we rewrite:

$$f_\eta(\theta) - f(\theta_*) = \frac{1}{2} \log(\Sigma^2) + \frac{1}{2\eta} \mathbb{E} \left[ \log \left( 1 + \eta \frac{(m - X)^2}{\Sigma^2} \right) \right] - \frac{1}{2} \log(\Sigma_*^2) + \frac{1}{2\eta} \log(1 - \eta) \quad (98)$$

$$= -\frac{1 - \eta}{2\eta} \log(\Sigma^2) + \frac{1}{2\eta} \mathbb{E} \left[ \log(\Sigma^2 + \eta(m - X)^2) \right] - \frac{1}{2} \log(\Sigma_*^2) + \frac{1}{2\eta} \log(1 - \eta). \quad (99)$$

Indeed, the two terms on the right are constant and so do not matter. If we differentiate in  $m$ , we obtain:

$$\nabla_m f_\eta(\theta) = \mathbb{E} \left[ \frac{1}{2\eta} 2\eta \frac{m - X}{\Sigma^2} \frac{1}{\Sigma^2 + \eta(m - X)^2} \right] = \mathbb{E} \left[ \frac{m - X}{\Sigma^2 + \eta(m - X)^2} \right]. \quad (100)$$

Now, differentiating in  $\Sigma^2$  yields:

$$\nabla_{\Sigma^2} f_\eta(\theta) = -\frac{1 - \eta}{2\eta \Sigma^2} + \frac{1}{2\eta} \mathbb{E} \left[ \frac{1}{\Sigma^2 + \eta(m - X)^2} \right] = \frac{1}{2\Sigma^2} - \mathbb{E} \left[ \frac{(m - X)^2}{2\Sigma^2(\Sigma^2 + \eta(m - X)^2)} \right]. \quad (101)$$

**2 - Obtaining bounds on  $(m_\eta, \Sigma_\eta^2)$ .** The solution to  $\nabla_m f_\eta(\theta) = 0$  is  $m = m_\star$ . Indeed, it is direct to verify that in this case,  $\mathbb{E} \left[ \frac{\tilde{X}}{\Sigma^2 + \eta \tilde{X}^2} \right] = 0$  since  $\tilde{X} = m_\star - X$  is symmetric (with respect to 0). For  $m > m_\star$ ,  $\mathbb{E} \left[ \frac{\tilde{X}}{\Sigma^2 + \eta \tilde{X}^2} \right] > 0$  since we integrate the same values as the previous case, but now more mass is put on the positive values (and similarly for  $m < m_\star$ ). Note that this is the case regardless of  $\Sigma_\eta^2$ .

We are now interested in  $\Sigma_\eta^2$ . Note that we will not get such a clean expression as for  $m_\eta$ , but only bounds. From its expression, we deduce that  $\nabla_{\Sigma^2} f_\eta(\theta_\eta) = 0$  can be reformulated as:

$$\mathbb{E} \left[ \frac{(m_\eta - X)^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] = 1 \quad (102)$$

For the upper bound, we simply write that:

$$1 = \mathbb{E} \left[ \frac{(m_\eta - X)^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] \leq \mathbb{E} \left[ \frac{(m_\eta - X)^2}{\Sigma_\eta^2} \right] = \frac{\Sigma_\star^2}{\Sigma_\eta^2}, \quad (103)$$

from which we deduce that  $\Sigma_\eta^2 \leq \Sigma_\star^2$ . Let us now introduce some  $\alpha > 0$ . We have that:

$$\mathbb{E} \left[ \frac{(m_\eta - X)^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] = \mathbb{E} \left[ \frac{(m_\eta - X)^2}{\alpha - \alpha + \Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] = \mathbb{E} \left[ \frac{(m_\eta - X)^2}{\alpha} \frac{1}{1 - 1 + \frac{\Sigma_\eta^2 + \eta(m_\eta - X)^2}{\alpha}} \right] \quad (104)$$

We now use that for  $u \geq -1$ ,  $\frac{1}{1+u} \geq 1 - u$ , and so:

$$\mathbb{E} \left[ \frac{(m_\eta - X)^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] \geq \mathbb{E} \left[ \frac{(m_\eta - X)^2}{\alpha} \left( 1 - \left[ -1 + \frac{\Sigma_\eta^2 + \eta(m_\eta - X)^2}{\alpha} \right] \right) \right] \quad (105)$$

$$= \mathbb{E} \left[ \frac{(m_\eta - X)^2}{\alpha} \left( 2 - \frac{\Sigma_\eta^2}{\alpha} \right) - \eta \frac{(m_\eta - X)^4}{\alpha^2} \right]. \quad (106)$$

Now, recall that  $m_\eta = m_\star$ , so  $X - m_\eta \sim \mathcal{N}(0, \Sigma_\star)$ , leading to

$$1 = \mathbb{E} \left[ \frac{(m_\eta - X)^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] \geq \frac{\Sigma_\star^2}{\alpha} \left( 2 - \frac{\Sigma_\eta^2}{\alpha} \right) - \eta \frac{3\Sigma_\star^4}{\alpha^2}. \quad (107)$$

Rearranging terms, we obtain:

$$\frac{\alpha^2}{\Sigma_\star^2} - 2\alpha \geq -\Sigma_\eta^2 - 3\Sigma_\star^2, \text{ so } \Sigma_\eta^2 \geq \frac{2\alpha\Sigma_\star^2 - \alpha^2}{\Sigma_\star^2} - 3\eta\Sigma_\star^2. \quad (108)$$

We see that  $\alpha = \Sigma_\star^2$  maximizes the right term, and we obtain the desired result, *i.e.*:

$$\Sigma_\eta^2 \geq (1 - 3\eta)\Sigma_\star^2. \quad (109)$$

Unfortunately, we see that this bound is only informative for  $3\eta < 1$ . For the rest of the cases, we will use the Markov inequality instead, which writes for all  $a > 0$ :

$$\mathbb{P} \left( \frac{(m_\eta - X)^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \geq a \right) \leq \frac{1}{a} \mathbb{E} \left[ \frac{(m_\eta - X)^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] = \frac{1}{a}. \quad (110)$$

Yet,

$$\mathbb{P} \left( \frac{(m_\eta - X)^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \geq a \right) = \mathbb{P} \left( \frac{(m_\eta - X)^2}{\Sigma_\star^2} \geq \frac{a}{1 - \eta a} \frac{\Sigma_\eta^2}{\Sigma_\star^2} \right) = 2\mathbb{P} \left( \frac{X - m_\star}{\Sigma_\star} \geq \sqrt{\frac{a}{1 - \eta a}} \frac{\Sigma_\eta}{\Sigma_\star} \right). \quad (111)$$

Therefore, denoting  $\Phi$  the cumulative distribution function of the standard Gaussian, we have:

$$2 \left( 1 - \Phi \left( \sqrt{\frac{a}{1 - \eta a}} \frac{\Sigma_\eta}{\Sigma_\star} \right) \right) \leq \frac{1}{a}, \quad (112)$$

and since  $\Phi^{-1}$  is an increasing function, this leads to:

$$\sqrt{\frac{a}{1 - \eta a}} \frac{\Sigma_\eta}{\Sigma_\star} \geq \Phi^{-1} \left( 1 - \frac{1}{2a} \right), \quad (113)$$

so that:

$$\Sigma_\eta \geq \sqrt{1 - \eta a} \frac{\Phi^{-1} \left( 1 - \frac{1}{2a} \right)}{\sqrt{a}} \Sigma_\star \quad (114)$$

One can check that  $\Phi^{-1} \left( 1 - \frac{1}{2a} \right) / \sqrt{a} < 1$  for all  $a$ , which is consistent with the fact that  $\Sigma_\eta^2 \leq \Sigma_\star^2$ . Also note that for  $\eta = 1$ , a non-trivial bound would require  $a < 1$ , but then  $\Phi^{-1} \left( 1 - \frac{1}{2a} \right) \leq 0$  so (as expected), we cannot get better than  $\Sigma_\eta^2 \geq 0$ . However, the previous bounding (Equation (109)) is more precise for small  $\eta$  since  $\Phi^{-1} \left( 1 - \frac{1}{2a} \right) / \sqrt{a} < 1 - c$  with  $c > 0$  a constant regardless of  $a$ . In particular, for any  $\varepsilon$ , by using any  $1 < a < 1/(1 - \varepsilon)$ , we obtain that  $\Sigma_\eta^2 \geq \alpha_\varepsilon \Sigma_\star^2$  for some constant  $\alpha_\varepsilon$  that only depends on the  $a$  that we choose. In particular, we can handle the cases  $\eta = 1/2$  and  $\eta = 1/3$  that gave trivial results  $\Sigma_\eta^2 \geq 0$  with the previous bounds.

The last part consists in proving that  $f_\eta(\theta_\eta) - f(\theta_\star) \geq \frac{1}{2} \log \left( \frac{\Sigma_\eta^2}{\Sigma_\star^2} \right)$ . To do so, we start back from

$$f_\eta(\theta) - f(\theta_\star) = \frac{1}{2} \log \left( \frac{\Sigma^2}{\Sigma_\star^2} \right) + \frac{1}{2\eta} \mathbb{E} \left[ \log \left( (1 - \eta) \left[ 1 + \eta \frac{(m - X)^2}{\Sigma^2} \right] \right) \right],$$

and show that  $\mathbb{E} \left[ \log \left( (1 - \eta) \left[ 1 + \eta \frac{(m_\eta - X)^2}{\Sigma_\eta^2} \right] \right) \right] \geq 0$ . We start by the inequality  $\log(1 + x) \geq \frac{x}{1+x}$ , leading to:

$$\mathbb{E} \left[ \log \left( (1 - \eta) \left[ 1 + \eta \frac{(m_\eta - X)^2}{\Sigma_\eta^2} \right] \right) \right] \geq \mathbb{E} \left[ \frac{(1 - \eta) \left[ 1 + \eta \frac{(m_\eta - X)^2}{\Sigma_\eta^2} \right] - 1}{(1 - \eta) \left[ 1 + \eta \frac{(m_\eta - X)^2}{\Sigma_\eta^2} \right]} \right] \quad (115)$$

$$= \mathbb{E} \left[ \frac{\eta(1 - \eta) \frac{(m_\eta - X)^2}{\Sigma_\eta^2} - \eta}{(1 - \eta) \left[ 1 + \eta \frac{(m_\eta - X)^2}{\Sigma_\eta^2} \right]} \right] \quad (116)$$

$$= \eta \mathbb{E} \left[ \frac{(1 - \eta)(m_\eta - X)^2 - \Sigma_\eta^2}{(1 - \eta) [\Sigma_\eta^2 + \eta(m_\eta - X)^2]} \right] \quad (117)$$

Recall that the optimality conditions for  $(m_\eta, \Sigma_\eta^2)$  write:

$$1 = \mathbb{E} \left[ \frac{(m_\eta - X)^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] = \frac{1}{\eta} \mathbb{E} \left[ 1 - \frac{\Sigma_\eta^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right], \quad (118)$$

so that

$$\mathbb{E} \left[ \frac{\Sigma_\eta^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] = 1 - \eta. \quad (119)$$

Combining these, we obtain that

$$\begin{aligned} \mathbb{E} \left[ \log \left( (1 - \eta) \left[ 1 + \eta \frac{(m_\eta - X)^2}{\Sigma_\eta^2} \right] \right) \right] &\geq \eta \mathbb{E} \left[ \frac{(1 - \eta)(m_\eta - X)^2 - \Sigma_\eta^2}{(1 - \eta) [\Sigma_\eta^2 + \eta(m_\eta - X)^2]} \right] \\ &= \eta \left( \mathbb{E} \left[ \frac{(m_\eta - X)^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] - \frac{1}{1 - \eta} \mathbb{E} \left[ \frac{\Sigma_\eta^2}{\Sigma_\eta^2 + \eta(m_\eta - X)^2} \right] \right) = 0, \end{aligned}$$

which is the desired result.

The final result is obtained by plugging the lower bounds for  $\Sigma_\eta^2$  into this bound, leading to either  $\sigma_{\star,\eta}^2 \leq -\frac{1}{2\eta} \log(1 - 3\eta)$  for  $\eta < 1/3$  or  $\sigma_{\star,\eta}^2 \leq -\frac{1}{2\eta} \log \alpha_\varepsilon$  for  $\eta < 1 - \varepsilon$ .

□

### F.3. Unrolling the recursions to derive actual convergence results.

#### F.3.1. PROOF OF THEOREM 4.3

Now that we have bounded the stochastic mirror descent variance  $\sigma_{\star,\eta}^2$  in this setting, we can plug it into Theorem 3.1 to obtain finite-time convergence guarantees on the MAP and MLE estimators.

*Proof of Theorem 4.3.* Starting from Theorem 3.1, we obtain:

$$D_A(\theta_\star, \theta^{(k+1)}) \leq (1 - \eta)D_A(\theta_\star, \theta^{(k)}) - \frac{\eta}{2} \log(1 - 3\eta) \leq (1 - \eta)D_A(\theta_\star, \theta^{(k)}) + \frac{3\eta^2}{2}, \quad (120)$$

where the right term is replaced by  $c_\varepsilon$  (where  $c_\varepsilon = -\frac{1}{2} \log \alpha_\varepsilon$ ) for  $k \leq 3$ . Taking  $\eta = 1/k$  for  $k > 1$  and multiplying by  $k$  leads for  $k > 3$  to:

$$kD_A(\theta_\star, \theta^{(k+1)}) \leq (k - 1)D_A(\theta_\star, \theta^{(k)}) + \frac{3}{2k}. \quad (121)$$

Therefore, a telescopic sum leads to, for  $n_0 > 0$ :

$$(n + n_0)D_A(\theta_\star, \theta^{(n)}) \leq n_0D_A(\theta_\star, \theta^{(0)}) + \frac{3}{2} \sum_{k=n_0}^{n+n_0} \frac{1}{k} + 2c_{1/2}, \quad (122)$$

and so, since  $\sum_{k=n_0}^n \frac{1}{k} \leq \log(n + n_0 + 1) - \log(n_0)$ :

$$D_A(\theta_\star, \theta^{(n)}) \leq \frac{n_0D_A(\theta_\star, \theta^{(0)}) + (3/2) \log(1 + (n + 1)/n_0) + \Gamma}{n + n_0}, \quad (123)$$

where  $\Gamma = 2c_{1/2}$  and we actually have  $\Gamma = 0$  for  $n_0 > 3$ .

□

#### F.3.2. $O(1/n)$ CONVERGENCE RESULT.

We now consider a different estimator (from the MAP and the MLE), which we construct in the following way:

- Choose  $n_0 \geq 6$  and initial parameter  $\tilde{\theta}^{(n_0)}$ .
- Obtain  $\tilde{\theta}^{(n)}$  by performing  $n - n_0$  stochastic mirror descent steps from  $\tilde{\theta}^{(n_0)}$  with step-sizes  $\eta_k = 2/(k + 1)$  for  $k \in \{n_0, \dots, n\}$ .

This estimator is a modified version of the MAP, where  $n_0$  controls how much weight we would like to put on the prior, and  $\tilde{\theta}^{(n_0)}$  would typically be the same starting parameter as for the MAP estimator. This estimator is built so that we can use the convergence analysis from [Lacoste-Julien et al. \(2012\)](#) and obtain a  $O(1/n)$  convergence rate. Note that we make the  $n_0 \geq 6$  restriction for simplicity to ensure that  $\sigma_{\star,\eta}^2 \leq 3/2$ , but the result can be easily adapted to  $n_0 \geq 2$ .

**Proposition F.3.** *After  $n - n_0$  steps, this modified estimator  $\theta^{(n)}$  verifies:*

$$\mathbb{E}D_h(\theta_\star, \tilde{\theta}^{(n)}) \leq \frac{2n_0(n_0 - 1)}{n(n - 1)} D_h(\theta_\star, \tilde{\theta}^{(n_0)}) + \frac{6}{n}. \quad (124)$$

*Proof.* Let us note  $D_k = \mathbb{E} \left[ D_h(\theta_\star, \tilde{\theta}^{(k)}) \right]$ . In this case, using that  $\sigma_{\star,\eta}^2 \leq 3/2$ , Theorem 3.1 writes (since  $\mu = 1$ ):

$$D_{k+1} \leq (1 - \eta_k)D_k + \frac{3\eta_k^2}{2}. \quad (125)$$

At this point, we can multiply by  $k(k+1)$  on both sides, and take  $\eta_k = \frac{2}{k+1}$  for  $k \geq n_0$ . Remarking that  $1 - \eta_k = 1 - \frac{2}{k+1} = \frac{k-1}{k+1}$ , we obtain that:

$$(k+1)kD_{k+1} \leq k(k-1)D_k + \frac{6k}{k+1} \leq k(k-1)D_k + 6. \quad (126)$$

Unrolling this recursion from  $k = n_0$  to  $k = n-1$  (since  $(k+1)kD_{k+1} = L_{k+1}$ , where  $L_k = k(k-1)D_k$ ), we obtain:

$$n(n-1)D_n \leq n_0(n_0-1)D_{n_0} + \sum_{k=n_0}^{n-1} 6, \quad (127)$$

and the result follows by dividing by  $n(n-1)$ , and using that  $(n-n_0)/(n-1) \leq 1$ .  $\square$

#### F.4. The case of the MLE

For the MLE estimator, directly applying the mirror descent approach would require using  $\eta_0 = 1$ , starting from an arbitrary  $\theta^{(0)}$  (that would not affect the results anyway). The problem in this case is that  $D_h(\theta_*, \theta^{(1)})$  is infinite since  $\Sigma^{(2)} = 0$ . This also means that we cannot start the stochastic mirror descent algorithm from  $\theta^{(1)}$ , since the recursion would still involve the infinite  $D_h(\theta_*, \theta^{(1)})$ . Therefore, in the case of the MLE, considering that the first two samples are  $X^{(1)}$  and  $X^{(2)}$ , then the first two points are:

$$m^{(1)} = X^{(1)}, \Sigma^{(1)} = 0 \text{ and } m^{(2)} = \frac{X^{(1)} + X^{(2)}}{2}, (\Sigma^{(2)})^2 = \frac{(X^{(1)} - X^{(2)})^2}{4}. \quad (128)$$

More generally, a direct recursion for the MLE leads to:

$$m^{(n)} = \frac{1}{n} \sum_{k=1}^n X^{(k)}, \quad (\Sigma^{(n)})^2 = \frac{1}{n} \sum_{k=1}^n (X^{(k)} - m^{(n)})^2. \quad (129)$$

From this, we derive that:

$$\mathbb{E} [(\Sigma^{(n)})^2] = \mathbb{E} [(X^{(n)} - m^{(n)})^2] \quad (130)$$

$$= \mathbb{E} \left[ \left( \left(1 - \frac{1}{n}\right) X^{(n)} - \frac{n-1}{n} m^{(n-1)} \right)^2 \right] \quad (131)$$

$$= \left(\frac{n-1}{n}\right)^2 \mathbb{E} \left[ \left( X^{(n)} - m_* - (m^{(n-1)} - m_*) \right)^2 \right] \quad (132)$$

$$= \left(\frac{n-1}{n}\right)^2 \mathbb{E} \left[ \left( X^{(n)} - m_* \right)^2 + \left( m^{(n-1)} - m_* \right)^2 \right] \quad (133)$$

$$= \left(\frac{n-1}{n}\right)^2 \left( \Sigma_*^2 + \frac{1}{n-1} \Sigma_*^2 \right) = \left(1 - \frac{1}{n}\right) \Sigma_*^2, \quad (134)$$

where (133) comes from the fact that  $X^{(n)}$  and  $m^{(n-1)}$  are independent with mean  $m_*$ . Plugging this into the expression of  $D_h(\theta_*, \theta)$  for the MLE after  $n$  steps, we obtain:

$$D_h(\theta_*, \theta^{(n)}) = -\frac{1}{2} \mathbb{E} \left[ \log \frac{(\Sigma^{(n)})^2}{\Sigma_*^2} \right]. \quad (135)$$

Unfortunately, there is no closed-form for this expression for arbitrary  $n$ , hence the need for a more involved analysis, for instance through the mirror descent framework. For the case  $n = 2$  however (which is the one we are interested in), we obtain that

$$D_h(\theta_*, \theta^{(2)}) = -\frac{1}{2} \mathbb{E} \left[ \log \left( \frac{X^{(1)} - X^{(2)}}{2\Sigma_*} \right)^2 \right] = -\frac{1}{2} \mathbb{E} \left[ \log \frac{Y^2}{2} \right], \quad (136)$$

where  $Y = \frac{X^{(1)} - X^{(2)}}{\sqrt{2}\Sigma_*} \sim \mathcal{N}(0, 1)$ . Therefore, this can simply be treated as a constant that we can precisely evaluate numerically (for instance remarking that  $Y^2$  is gamma distributed and using results on logarithmic expectations of gamma distributions).

For  $n > 2$ , it is tempting to use the convexity of  $-\log$  to use a similar reasoning, but this only leads to a constant bound on  $D_h(\theta_*, \theta^{(n)})$ .