

ADAPTIVE METHODS ARE PREFERABLE IN HIGH PRIVACY SETTINGS: AN SDE PERSPECTIVE

005 **Anonymous authors**

006 Paper under double-blind review

ABSTRACT

011 Differential Privacy (DP) is becoming central to large-scale training as privacy
 012 regulations tighten. We revisit how DP noise interacts with *adaptivity* in optimiza-
 013 tion through the lens of *stochastic differential equations*, providing the first SDE-
 014 based analysis of private optimizers. Focusing on DP-SGD and DP-SignSGD
 015 under per-example clipping, we show a sharp contrast under fixed hyperpara-
 016 meters: DP-SGD converges at a privacy-utility trade-off $\mathcal{O}(1/\varepsilon^2)$ with speed inde-
 017 pendent of ε , while DP-SignSGD converges at a speed *linear in ε* with a $\mathcal{O}(1/\varepsilon)$
 018 trade-off, dominating in high-privacy or high-noise regimes. Under optimal learn-
 019 ing rates, both methods reach comparable theoretical asymptotic performance;
 020 however, the optimal learning rate of DP-SGD scales linearly with ε , while that
 021 of DP-SignSGD is essentially ε -independent. This makes adaptive methods far
 022 more practical, as their hyperparameters transfer across privacy levels with little or
 023 no re-tuning. Empirical results confirm our theory across training and test metrics,
 024 and extend from DP-SignSGD to DP-Adam.

1 INTRODUCTION

027 The rapid deployment of large-scale machine learning systems has intensified the demand for rig-
 028 orous privacy guarantees. In sensitive domains such as healthcare or conversational agents, even
 029 the disclosure of a single training example can have serious consequences. Legislation and policy
 030 initiatives show that AI regulation is tightening rapidly. In the United States, the *Executive Order*
 031 of October 30, 2023 mandates developers of advanced AI systems to share safety test results and
 032 promotes privacy-preserving techniques such as differential privacy (House, 2023). Complementing
 033 this, the National Institute of Standards and Technology (NIST), a U.S. federal agency, released draft
 034 guidance (SP 800-226) on privacy guarantees in AI (NITS, 2023a) and included “privacy-enhanced”
 035 as a key dimension in its AI Risk Management Framework (RMF 1.0) (NITS, 2023b). In Europe, the
 036 *EU AI Act* sets binding obligations for high-risk systems (EU, 2023), while ENISA recommends in-
 037 tegrating data protection into AI development (EU, 2024). In this context, Differential Privacy (DP)
 038 (Dwork et al., 2006) is therefore emerging as the de facto standard for ensuring user-level confi-
 039 dentiality in stochastic optimization. By injecting carefully calibrated noise into the training process, DP
 040 optimizers protect individual data points while inevitably trading off some population-level utility.

041 A central open question is how differential privacy noise influences optimization dynamics, and
 042 in particular, how it interacts with adaptivity and batch noise. In this work, we revisit this prob-
 043 lem through the lens of *stochastic differential equations* (SDEs), which, over the last decade, have
 044 proven to be a powerful tool for analyzing optimization algorithms (Li et al., 2017; Mandt et al.,
 045 2017; Compagnoni et al., 2023). While SDEs have not yet been applied to DP methods, here we
 046 use them to uncover a key and previously overlooked phenomenon: *DP noise affects adaptive and*
 047 *non-adaptive methods in structurally different ways*. We focus on two fundamental DP optimizers:
 048 DP-SGD (Abadi et al., 2016) and DP-SignSGD. The former serves as the baseline for DP
 049 optimization; Although the latter is not widely used in practice, it is substantially simpler to ana-
 050 lyze than the popular DP optimizer DP-Adam (Gylberth et al., 2017; Zhou et al., 2020b; Li et al.,
 051 2021a; McKenna et al., 2025). Relying on SignSGD as a proxy for Adam is standard in prior work
 052 (Compagnoni et al., 2025c; Balles & Hennig, 2018; Zou et al., 2021; Peng et al., 2025; Li et al.,
 053 2025), and this motivates our focus on DP-SignSGD for the theoretical development. Importantly,
 setting $\beta_1 = \beta_2 = 0$ reduces DP-Adam to DP-SignSGD. We leave the study of more advanced
 DP optimizers to future work, as each would require a separate technical treatment. Under standard

assumptions and with per-example clipping, our analysis isolates how the privacy budget ε , which governs the overall privacy level, influences the dynamics.

In practice, private training is usually performed across a range of privacy budgets ε , and for each value one searches for the best-performing hyperparameters. A change in ε can therefore arise either from this exploratory sweep or from stricter regulatory requirements. To capture these situations, we study two complementary protocols. **Protocol A (fixed hyperparameters):** To examine the situation when re-tuning is not feasible, e.g., low budget, we **we first fix a privacy budget ε and find the optimal configuration (η, C, B, \dots) via grid search**. Then, we analyze how performance changes if training were repeated under different ε , without adjusting hyperparameters, therefore isolating the impact of ε on the performance. **Protocol B (best-tuned per ε):** When re-tuning is allowed, we **search the optimal** hyperparameters (i.e., (η, C, B, \dots)) for each ε , thereby isolating the *intrinsic scaling of the optimal learning rates with respect to ε* .

Contributions. Our work makes the following contributions:

1. We provide the first SDE-based analysis of differentially private optimizers, using this framework to expose how DP noise interacts with adaptivity and batch noise;
2. **Protocol A:** We show that DP-SGD converges at a speed *independent* of ε , with a privacy-utility trade-off that scales as $\mathcal{O}(1/\varepsilon^2)$ (consistent with prior work);
3. **Protocol A:** We prove a novel result for DP-SignSGD: its convergence speed scales linearly in ε , while its privacy-utility trade-off scales as $\mathcal{O}(1/\varepsilon)$;
4. **Protocol A:** When batch noise is sufficiently large, DP-SignSGD always dominates. When batch noise is small, the outcome depends on the privacy budget: for strict privacy ($\varepsilon < \varepsilon^*$), DP-SignSGD is preferable, while for looser privacy ($\varepsilon > \varepsilon^*$), DP-SGD has better performance;
5. **Protocol B:** We theoretically derive that the optimal learning rate of DP-SGD scales as $\eta^* \propto \varepsilon$, while the optimal learning rate of DP-SignSGD is ε -independent. This tuning allows the two methods to reach theoretically *comparable* asymptotic performance, including at very small ε ;
6. We empirically validate all our theoretical insights on real-world tasks, and show that the qualitative insights extend from training to *test* loss and from DP-SignSGD to DP-Adam.

In summary, our results refine the privacy-utility landscape, **which, to our knowledge, has not yet provided a definitive answer as to which of DP-SGD or DP-Adam/DP-SignSGD performs best, and under which conditions**. Under Protocol A, adaptivity is preferable in stricter privacy regimes: DP-SignSGD converges more slowly but achieves better utility when ε is small or batch noise is large, whereas DP-SGD converges faster but suffers sharper degradation. Under Protocol B, both methods achieve comparable asymptotic performance; however, adaptive methods are far more practical, as their optimal learning rate is essentially ε -independent, allowing it to transfer across privacy levels with little or no re-tuning. This matters not only for computational cost but also for privacy, since each hyperparameter search consumes additional budget (Papernot & Steinke, 2021). In contrast, DP-SGD requires an ε -dependent learning rate tuned *ad hoc*, making it brittle if the sweep grid misses the “right” value. Intuitively, adaptive methods inherently adjust to the scale of DP noise, whereas non-adaptive methods require explicit tuning of the learning rate to counter the effect of privacy noise.

2 RELATED WORK

SDE approximations. SDEs have long been used to analyze discrete-time optimization algorithms (Helmke & Moore, 1994; Kushner & Yin, 2003). Beyond their foundational role, these approximations have been applied to practical tasks such as learning-rate tuning (Li et al., 2017; 2019) and batch-size selection (Zhao et al., 2022). Other works have focused on deriving convergence bounds (Compagnoni et al., 2023; 2024; 2025c), uncovering scaling laws that govern optimization dynamics (Jastrzebski et al., 2018; Compagnoni et al., 2025c;a), and revealing implicit effects such as regularization (Smith et al., 2021; Compagnoni et al., 2023) and preconditioning (Xiao et al., 2025; Marshall et al., 2025). In particular, SDE-based techniques have been used to study a broad class of modern adaptive optimizers, including RMSProp, Adam, AdamW, and SignSGD, as well as minimax and distributed variants (Compagnoni et al., 2024; 2025c;a; Xiao et al., 2025). Most analyses rely on weak approximations, as rigorously formalized by Li et al. (2017), although some works have also considered heavy-tailed batch noise via Lévy-driven SDEs to capture non-Gaussianity (Simsekli et al., 2019; Zhou et al., 2020a). Despite this progress, prior work has exclusively focused on non-private optimization. To our knowledge, ours is the first to extend the

108 SDE lens to differentially private optimizers, including explicit convergence rates and stationary
 109 distributions as functions of the privacy budget.
 110

111 **Differential privacy in optimization.** Differentially private training is most commonly implemented via DP-SGD (Abadi et al., 2016), which clips per-example gradients to a fixed norm bound
 112 to control sensitivity and injects calibrated Gaussian noise into the averaged update. Advanced
 113 accounting methods such as the moments accountant (Abadi et al., 2016) and Rényi differential
 114 privacy (Mironov, 2017; Wang et al., 2019), combined with privacy amplification by subsampling
 115 (Balle et al., 2018; 2020), allow practitioners to track the cumulative privacy cost tightly over
 116 many updates and have made large-scale private training feasible. A central challenge is that clipping,
 117 while essential for privacy, also alters the optimization dynamics: overly aggressive thresholds
 118 bias gradients and can stall convergence (Chen et al., 2020), prompting extensive work on how
 119 to set or adapt the clipping norm. Approaches include rule-based or data-driven thresholds, such
 120 as AdaClip (Pichapati et al., 2019) and quantile-based adaptive clipping (Andrew et al., 2021),
 121 as well as recent analyses that characterize precisely how the clipping constant influences convergence
 122 (Koloskova et al., 2023). Together, these contributions have positioned DP-SGD and its variants
 123 as the standard backbone for differentially private optimization.
 124

125 **Adaptive DP optimizers.** Adaptive methods such as AdaGrad (Duchi et al., 2011; McMahan &
 126 Streeter, 2010), RMSProp (Tieleman & Hinton, 2012), and Adam (Kingma & Ba, 2015) generally
 127 outperform non-adaptive SGD in non-private training. However, this performance *gap* under DP
 128 constraints; *i*) narrows considerably (Zhou et al., 2020b; Li et al., 2022); *ii*) essentially vanishes
 129 when both optimizers are carefully tuned, as observed for large-scale LLM fine-tuning in Li et al.
 130 (2021a, App. S). **Consistently with non-DP training, non-adaptive methods are sometimes still pre-ferred in vision tasks (De et al., 2022). Therefore, which of DP-SGD and DP-Adam is preferable**
 131 **remains an open question.** Under assumptions that include bounded/convex domain, bounded gradient
 132 norm, bounded gradient noise, convexity of the loss, and possibly without performing clipping of
 133 the per-sample gradients, several strategies have been theoretically and empirically explored to miti-gate
 134 the drop in performance of adaptive methods in DP. These include bias-corrected DP-Adam
 135 variants (Tang & Lécuyer, 2023; Tang et al., 2023), the use of non-sensitive auxiliary data (Asi
 136 et al., 2021), and scale-then-privatize techniques that exploit adaptivity before noise injection (Li
 137 et al., 2023; Ganesh et al., 2025). A most recent related work by (Jin & Dai, 2025) studies Noisy
 138 SignSGD: Conceptually, they investigate how the sign compressor amplifies privacy, and argue that
 139 the sign operator itself provides privacy amplification beyond the Gaussian mechanism. Their analysis
 140 establishes convergence guarantees in the distributed learning setting while relying on *bounded*
 141 *gradient norms and bounded variance* assumptions, thereby avoiding the need for clipping and ex-plicitly
 142 leaving its study to future work.
 143

144 We view these contributions as providing valuable theoretical and empirical advances in the design
 145 of adaptive private optimizers, clarifying many important aspects of their behavior as well as trying
 146 to restore the aforementioned performance *gap*. Yet, the fundamental question of *which privacy*
 147 *regimes are most favorable to adaptivity* remains largely unanswered, and addressing it could explain
 148 at least one aspect of the nature of this *gap*. Our work addresses this *open question* by analyzing *why*
 149 and *when adaptivity matters* under DP noise, identifying the regimes where adaptive methods domi-nate
 150 and where they match non-adaptive ones. Crucially, we incorporate *per-example clipping*, a
 151 central element of DP-SGD, and a heavy-tailed batch noise model that captures unbounded variance.
 152

3 PRELIMINARIES

153 **General Setup and Noise Assumptions.** We model the loss function with a differentiable function
 154 $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with global minimum $f^* = 0$: This is not restrictive, as one can always consider
 155 the suboptimality $f(x) - f^*$ and rename it as f . Regarding noise assumptions, recent literature com-monly
 156 assumes that the stochastic gradient of the loss function on a minibatch γ can be decomposed
 157 as $\nabla f_\gamma(x) = \nabla f(x) + Z_\gamma$ where batch noise Z_γ is modeled with a Gaussian (Ahn et al., 2012; Chen
 158 et al., 2014; Mandt et al., 2016; Stephan et al., 2017; Zhu et al., 2019; Jastrzebski et al., 2018; Wu
 159 et al., 2020; Xie et al., 2021), often with constant covariance matrix (Li et al., 2017; Mertikopoulos
 160 & Staudigl, 2018; Raginsky & Bouvrie, 2012; Zhu et al., 2019; Mandt et al., 2016; Ahn et al.,
 161 2012; Jastrzebski et al., 2018). **In this work, we refine the standard noise assumption to distinguish**
 162 **the two regimes induced by per-example clipping in DP training.** Since clipping is applied at the
 163 datapoint level, each mini-batch contains a mix of *clipped* and *unclipped* gradients. For unclipped
 164 datapoints, we follow the usual literature and model the batch-averaged noise as Gaussian. For

clipped datapoints, which do not benefit from batch averaging, we model the per-example noise as multivariate Student- t , $Z_\gamma \sim \sigma_\gamma t_\nu(0, I_d)$, capturing potentially heavy-tailed behaviour and recovering the Gaussian case as $\nu \rightarrow \infty$. See Assumption B.2 and Remark B.2 for more details. Finally, we use the following approximation, formally derived in Lemma A.2: $\mathbb{E} \left[\frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|} \right] \approx \frac{\nabla f(x)}{\sigma_\gamma \sqrt{d}}$. The approximation is valid under two assumptions: *i*) The parameter dimension d is sufficiently large ($d = \Omega(10^4)$), consistent with modern deep learning models that often reach billions of trainable parameters; *ii*) The signal-to-noise ratio satisfies $\frac{\|\nabla f(x)\|_2^2}{2\sigma_\gamma^2} \ll d$: This condition has been thoroughly empirically studied by Malladi et al. (2022) (Appendix G), who observed that across multiple tasks and architectures the ratio $\frac{\|\nabla f(x)\|_2^2}{2\sigma_\gamma^2}$ never exceeds $\mathcal{O}(10^2)$, well below typical values of d . Therefore, this signal-to-noise ratio need not be small: We simply require it to be smaller than d — See Remark A.1 for more details, including experimental validations. We highlight that our experiments confirm that the insights derived from our theoretical results carry over to real-world tasks. Importantly, while our theory is developed for DP-SignSGD, we further validate that the same insights hold empirically for DP-Adam, showing that our insights extend directly to this widely used private optimizer, as well as also transfer from training to test loss. This highlights both the mildness of the assumptions and the robustness of the analysis.

SDE approximation. The following definition formalizes in which sense a continuous-time model, such as a solution to an SDE, can accurately describe the dynamics of a discrete-time process, such as an optimizer. Drawn from the field of numerical analysis of SDEs (see Mil'shtein (1986)), it quantifies the disparity between the discrete and the continuous processes. Simply put, the approximation is meant in a *weak sense*, meaning in distribution rather than path-wise: We require their expectations to be close over a class of test functions with polynomial growth, meaning that all the moments of the processes become closer at a rate of η^α and thus their distributions.

Definition 3.1 Let $0 < \eta < 1$ be the learning rate, $\tau > 0$ and $T = \lfloor \frac{\tau}{\eta} \rfloor$. We say that a continuous time process X_t over $[0, \tau]$, is an order- α weak approximation of a discrete process x_k , if for any polynomial growth function g , $\exists M > 0$, independent of the learning rate η , such that for all $k = 0, 1, \dots, T$, $|\mathbb{E}g(X_{k\eta}) - \mathbb{E}g(x_k)| \leq M\eta^\alpha$.

Remark 3.1 (Validity of the SDE approximation) To guarantee that the SDE model is a first-order weak approximation of the optimizer dynamics in the sense of Definition 3.1, one shows that the first two moments of the one-step increments of the optimizer and of the SDE match up to $O(\eta^2)$, while all higher-order terms in the Taylor expansion are collected in an $O(\eta^2)$ remainder (see Appendix B). This implies that the discrepancy between the two processes scales as $O(\eta)$ for any test function of polynomial growth and any finite time horizon. The neglected $O(\eta^2)$ terms could, in principle, be retained by deriving higher-order SDEs. However, to the best of our knowledge, such second-order models have been derived (Li et al., 2017), but have not yet led to additional practical insight in the analysis of optimisation algorithms. Finally, Figure C.1 empirically compares the discrete algorithms with their SDE counterparts on quadratic and quartic objectives, confirming that for the step sizes used in our experiments, the first-order SDEs closely track the discrete dynamics.

While we refer the reader to Appendix B for technical details, we illustrate with a basic example. The SGD iterates follow $x_{k+1} = x_k - \eta \nabla f_{\gamma_k}(x_k)$, and, as shown in Li et al. (2017), it can be approximated in continuous time by the first-order SDE

$$dX_t = -\nabla f(X_t)dt + \sqrt{\eta} \sqrt{\Sigma(X_t)}dW_t, \quad (1)$$

where $\Sigma(x) = \frac{1}{n} \sum_{i=1}^n (\nabla f(x) - \nabla f_i(x))(\nabla f(x) - \nabla f_i(x))^\top$ is the gradient noise covariance. Intuitively, the iterates drift along the gradient while the stochasticity scales with this covariance.

Differential Privacy. Here, we outline the relevant background of foundational prior work in DP optimization. We adopt the standard (ε, δ) -DP framework (Dwork et al., 2006).

Definition 3.2 A random mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ is said to be (ε, δ) -differentially private if for any two adjacent datasets $d, d' \in \mathcal{D}$ (i.e., they differ in 1 sample) and for any subset of outputs $S \subseteq \mathcal{R}$ it holds that $\mathbb{P}[\mathcal{M}(d) \in S] \leq e^\varepsilon \mathbb{P}[\mathcal{M}(d') \in S] + \delta$.

In this work, we consider example-level differential privacy applied by a central trusted aggregator. We implement this using the sub-sampled Gaussian mechanism (Dwork & Roth, 2014; Mironov

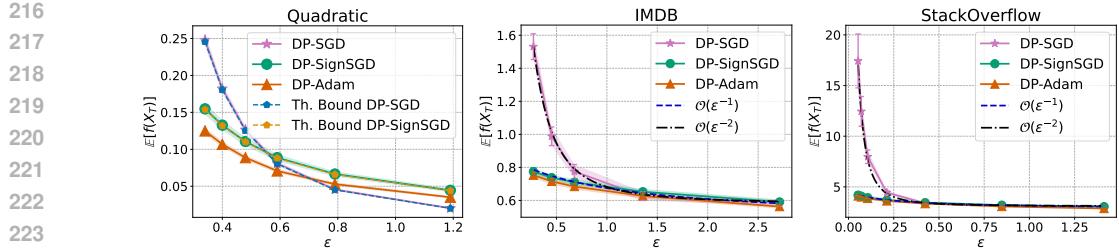


Figure 1: Empirical validation of the privacy-utility trade-off predicted by Thm. 4.1 and Thm. 4.3, comparing DP-SGD, DP-SignSGD, and DP-Adam: Our focus is on verifying the functional dependence of the asymptotic loss levels in terms of ϵ . **Left:** On a quadratic convex function $f(x) = \frac{1}{2}x^\top Hx$, the observed empirical loss values perfectly match the theoretical predictions (Eq. 7, Eq. 10). **Center and Right:** Logistic regressions on the IMDB dataset (center) and the StackOverflow dataset (right), confirm the same pattern: the utility of DP-SGD scales as $\frac{1}{\epsilon^2}$, while the utility of DP-SignSGD scales linearly as $\frac{1}{\epsilon}$. Across all settings, we observe that the insights obtained for DP-SignSGD extend to DP-Adam as well as to the test loss (see Figure C.4). For experimental details see Appendix C.2.

et al., 2019) to perturb the SGD updates: At each iteration, a random mini-batch is drawn, per-example gradients are clipped to a fixed bound to limit sensitivity, and Gaussian noise is added to the averaged clipped gradients. The following definition formalizes these mechanisms and provides the update rules for DP-SGD and DP-SignSGD.

Definition 3.3 For $k \geq 0$, learning rate η , variance σ_{DP}^2 , and batches γ_k of size B modeled as i.i.d. uniform random variables taking values in $\{1, \dots, n\}$. Let g_k be the private gradient, defined as

$$g_k := \frac{1}{B} \sum_{i \in \gamma_k} \mathcal{C}(\nabla f_i(x_k)) + \frac{1}{B} \mathcal{N}(0, C^2 \sigma_{DP}^2 I) \quad (2)$$

and $\mathcal{C}[\cdot]$ be the clipping function

$$\mathcal{C}(x) = \begin{cases} C \frac{x}{\|x\|_2} & \text{if } \|x\|_2 \geq C \\ x & \text{if } \|x\|_2 < C \end{cases} \quad (3)$$

The iterates of DP-SGD are defined as

$$x_{k+1} = x_k - \eta g_k, \quad (4)$$

while those of DP-SignSGD are defined as

$$x_{k+1} = x_k - \eta \text{sign}[g_k], \quad (5)$$

where $\text{sign}[\cdot]$ is applied component-wise. Finally, those of DP-Adam are defined in Eq. 208.

We say that an optimizer is in Phase 1 if the argument of \mathcal{C} is larger than C and Phase 2 otherwise.

The following theorem from (Abadi et al., 2016) gives the conditions under which DP-SGD, and thus also DP-SignSGD, is a differentially-private algorithm.

Theorem 3.1 For $q = \frac{B}{n}$ where B is the batch size, n is the number of training points, and number of iterations T , $\exists c_1, c_2$ s.t. $\forall \epsilon < c_1 q^2 T$, if the noise multiplier σ_{DP} satisfies $\sigma_{DP} \geq c_2 \frac{q \sqrt{T \log(1/\delta)}}{\epsilon}$, DP-SGD is (ϵ, δ) -differentially private for any $\delta > 0$. In the following, we will often use $\sigma_{DP} = \frac{\sqrt{T}\Phi}{\epsilon}$, where $\Phi := q \sqrt{\log(1/\delta)}$ to indicate the DP noise multiplier.

4 THEORETICAL RESULTS

In this section, we investigate how the privacy budget ϵ influences convergence speed and shapes the privacy-utility trade-offs in both the loss and the gradient norm. To do so, we leverage SDE models for DP-SGD and DP-SignSGD, which can be found in Theorem B.5 and Theorem B.10, respectively, and are experimentally validated in Figure C.1. In addition, we provide the first stationary distributions for these optimizers, presented in Theorem B.9 and Theorem B.15 in the Appendix. This section is organized as follows:

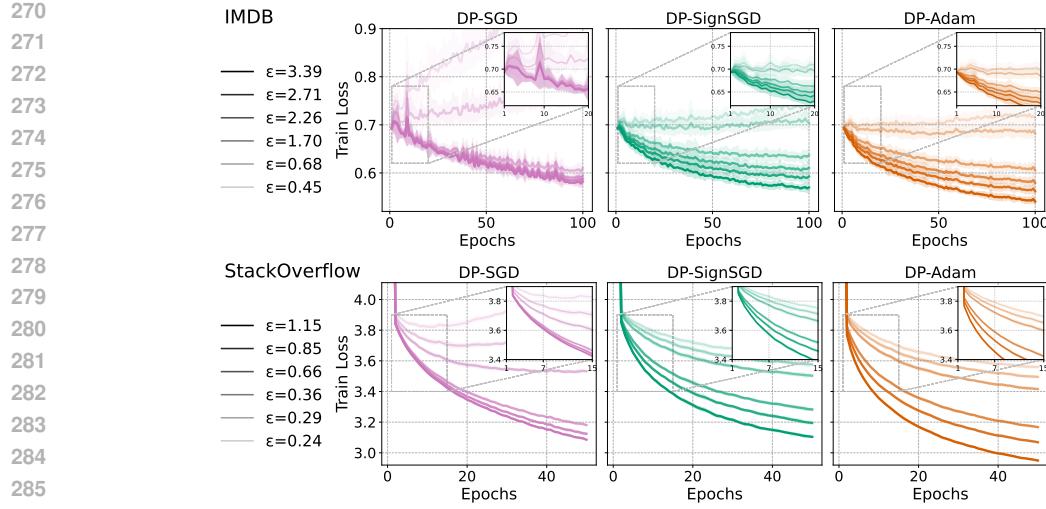


Figure 2: Empirical validation of the convergence speeds predicted by Thm. 4.1 and Thm. 4.3. We compare DP-SGD, DP-SignSGD, and DP-Adam as we train a logistic regression on the IMDB dataset (**Top Row**) and on the StackOverflow dataset (**Bottom Row**). In both tasks, we verify that when DP-SGD converges, its speed is unaffected by ϵ . As expected, it diverges when ϵ is too small. Regarding DP-SignSGD and DP-Adam, they are faster when ϵ is large and never diverge even when this is small. Crucially, Figure C.5 shows that these insights are also verified on the test loss. For experimental details see Appendix C.3.

1. Protocol A (Section 4.1). Section 4.1.1 analyzes DP-SGD, yielding bounds for the loss (Thm. 4.1) and the gradient norm (Thm. 4.2) in the μ -PL and L -smooth cases, respectively: We observe that the convergence speed is *independent* of ϵ , while the privacy-utility trade-off scales as $\mathcal{O}(1/\epsilon^2)$. Section 4.1.2 analyzes DP-SignSGD, and Thm. 4.3 and Thm. 4.4) show a qualitatively different behavior: Convergence speed scales linearly with ϵ , while the privacy-utility terms scale as $\mathcal{O}(1/\epsilon)$, making adaptivity preferable if the privacy budget is small enough. Finally, Theorem 4.5 in Section 4.1.3 shows that when batch noise is large enough, DP-SignSGD always dominates. When batch noise is small, the outcome depends on the privacy budget: There exists ϵ^* such that for strict privacy ($\epsilon < \epsilon^*$), DP-SignSGD is preferable, while for looser privacy ($\epsilon > \epsilon^*$), DP-SGD is better.

2. Protocol B (Section 4.2). In this section, we derive the optimal learning rates of DP-SGD and DP-SignSGD: That of DP-SGD scales linearly in ϵ , while that of DP-SignSGD is independent of it. Under these parameter choices, they achieve the same asymptotic neighbourhoods.

We empirically validate our theoretical insights on real datasets¹. Crucially, the same insights derived from DP-SignSGD *empirically* extend to DP-Adam as well as to test metrics: This underscores the mildness of our assumptions and the depth of our analysis.

Notation. In the following, we use the symbol \lesssim to suppress absolute numerical constants (e.g., 2, 4, etc.), and never problem-dependent quantities such as d , μ , L , or ϵ : This convention lightens the presentation. Finally, observe that $\Phi := q\sqrt{\log(1/\delta)} = \frac{B}{n}\sqrt{\log(1/\delta)} \Rightarrow \frac{\Phi}{B} = \frac{1}{n}\sqrt{\log(1/\delta)}$. We will often use ϵ to highlight the privacy budget in relevant formulas.

4.1 PROTOCOL A: FIXED HYPERPARAMETERS

Following the tuning routine of Li et al. (2023), we conduct extensive grid search to select a configuration (η, C, B, \dots) for one ϵ and keep them unchanged as we vary ϵ . In particular, η does *not* depend on ϵ or on other hyperparameters. This absolute comparison exposes structural differences in how DP noise interacts with adaptive vs non-adaptive updates.

¹For all our experiments, we use the official GitHub repository <https://github.com/kenziyuliu/DP2> released with the Google paper Li et al. (2023).

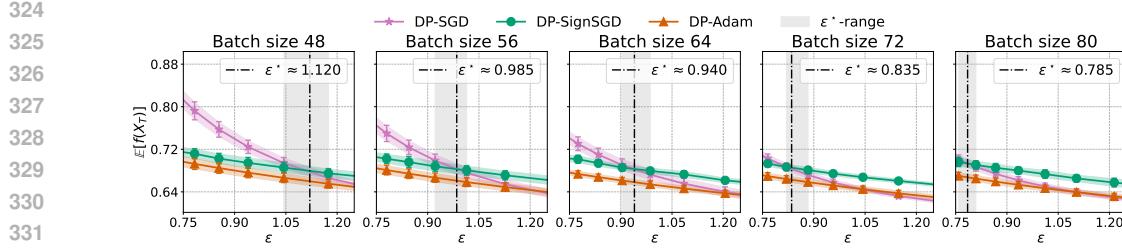


Figure 3: Logistic regression on IMDB Dataset: From left to right, we decrease the batch noise, i.e., increase the batch size, taking values $B \in \{48, 56, 64, 72, 80\}$: As per Theorem 4.5, the privacy threshold ε^* that determines when DP-SignSGD is more advantageous than DP-SGD shifts to the left. This confirms that if there is more noise due to the batch size, less privacy noise is needed for DP-SignSGD to be preferable over DP-SGD. For experimental details see Appendix C.4.

4.1.1 DP-SGD: THE PRIVACY-UTILITY TRADE-OFF IS $\mathcal{O}(1/\varepsilon^2)$

By definition, DP-SGD might alternate between a *clipped* and an *unclipped* phase. We first take a didactic perspective to analyze each phase separately to isolate the role of ε on the dynamics, while Theorem B.8 covers the case where these phases are mixed.

Theorem 4.1 *Let f be μ -PL and L -smooth, then we have that during*

- *Phase 1, i.e., when the gradient is clipped, the loss satisfies:*

$$\mathbb{E}[f(X_t)] \lesssim f(X_0) \underbrace{e^{-\frac{\mu C}{\sigma_\gamma \sqrt{d}} t}}_{\text{Decay}} + \left(1 - e^{-\frac{\mu C}{\sigma_\gamma \sqrt{d}} t}\right) \underbrace{\frac{T \eta d^{\frac{3}{2}} L C \sigma_\gamma}{\mu} \left(\frac{\varepsilon^2}{dT} + \frac{\Phi^2}{B^2}\right)}_{\text{Privacy-Utility Trade-off}} \frac{1}{\varepsilon^2}; \quad (6)$$

- *Phase 2, i.e., when the gradient is not clipped, the loss satisfies:*

$$\mathbb{E}[f(X_t)] \lesssim f(X_0) \underbrace{e^{-\mu t}}_{\text{Decay}} + (1 - e^{-\mu t}) \underbrace{\frac{T \eta d L}{\mu} \left(\frac{\varepsilon^2 \sigma_\gamma^2}{B T} + C^2 \frac{\Phi^2}{B^2}\right)}_{\text{Privacy-Utility Trade-off}} \frac{1}{\varepsilon^2}. \quad (7)$$

The decay rates are independent of ε : in Phase 2 they depend only on μ , while in Phase 1 normalization spreads the signal over the sphere of radius C (Vershynin, 2018, Ch. 3), giving a rate proportional to $C/(\sigma_\gamma \sqrt{d})$. In both phases, the privacy-utility term scales as $1/\varepsilon^2$.

We now turn to analyzing SDE dynamics assuming only L -smoothness of f . The following theorem presents a bound on the expected gradient norm *across* both phases **together**: We observe that the expected gradient norm admits the same $\mathcal{O}(1/\varepsilon^2)$ scaling.

Theorem 4.2 *Let f be L -smooth, $K_1 := \max\{1, \frac{\sigma_\gamma \sqrt{d}}{C}\}$, and $K_2 := \max\{\frac{\sigma_\gamma^2}{B}, \frac{C^2}{d}\}$. Then,*

$$\mathbb{E}[\|\nabla f(X_{\tilde{t}})\|_2^2] \lesssim K_1 \left(\frac{f(X_0)}{\eta T} + \eta d L \left(K_2 + \frac{C^2 (\frac{a}{B})^2 T \log(1/\delta)}{\varepsilon^2} \right) \right), \quad (8)$$

where \tilde{t} is a random time with uniform distribution over $[0, \tau]$.

Takeaway. Theorem 4.1 separates two effects: the *decay* terms, which determine the convergence speed, and the *privacy-utility* terms, which determine the asymptotic neighbourhood under DP. Our results show that the convergence speed of DP-SGD is unaffected by the privacy budget ε : Figure 2 confirms empirically that, whenever DP-SGD does not diverge, its convergence speed is independent of ε . Additionally, the privacy-utility trade-off scales as $\mathcal{O}(1/\varepsilon^2)$: This insight is validated in Figure 1: on a quadratic function (left panel) the observed loss matches the theoretical values from Theorem 4.1, and the same scaling is reproduced when training logistic regression on IMDB and StackOverflow (center and right panels). The behavior also persists on the test loss (Figure C.4).

4.1.2 DP-SIGNSGD: THE PRIVACY-UTILITY TRADE-OFF IS $\mathcal{O}(1/\varepsilon)$

As for DP-SGD, we isolate the effect of ε on the dynamics of DP-SignSGD and study the loss in each phase **separately**, while Theorem B.14 covers the case where these phases are mixed.

378 **Theorem 4.3** Let f be μ -PL and L -smooth. Then, we have that during
 379
 380 • Phase 1, i.e., when the gradient is clipped, the loss satisfies:

$$\mathbb{E}[f(X_t)] \lesssim f(X_0) \underbrace{e^{\frac{-\mu B}{\sigma_\gamma \sqrt{dT}} \frac{\epsilon}{\Phi} t}}_{\text{Decay}} + \left(1 - e^{\frac{-\mu B}{\sigma_\gamma \sqrt{dT}} \frac{\epsilon}{\Phi} t}\right) \underbrace{\frac{\sqrt{T} \eta L d^{\frac{3}{2}} \sigma_\gamma \Phi}{\mu B} \frac{\epsilon}{\Phi}}_{\text{Privacy-Utility Trade-off}}; \quad (9)$$

385 • Phase 2, i.e., when the gradient is not clipped, the loss satisfies:

$$\mathbb{E}[f(X_t)] \lesssim f(X_0) \underbrace{e^{\sqrt{\frac{\epsilon^2}{B} + \frac{C^2 \Phi^2}{B^2}} T}}_{\text{Decay}} + \left(1 - e^{\sqrt{\frac{\epsilon^2}{B} + \frac{C^2 \Phi^2}{B^2}} T}\right) \underbrace{\frac{\sqrt{T} \eta L d}{\mu} \sqrt{\frac{\epsilon^2 \sigma_\gamma^2}{BT} + \frac{C^2 \Phi^2}{B^2}} \frac{1}{\epsilon}}_{\text{Privacy-Utility Trade-off}}. \quad (10)$$

390 The decay rate scales proportionally with ϵ in both phases (Eq. 9 and Eq. 10), unlike DP-SGD,
 391 where it is independent of ϵ (Eq. 6 and Eq. 7). At the same time, the privacy-utility term in both
 392 phases scales as $\mathcal{O}(1/\epsilon)$, which *might* be more favorable than the $\mathcal{O}(1/\epsilon^2)$ scaling of DP-SGD in
 393 high-privacy regimes, e.g., if ϵ is sufficiently small.

394 Assuming only L -smoothness of f , the following theorem presents a bound on the expected gradient
 395 norm *across* both phases **together**. As the bound scales as $\mathcal{O}(1/\epsilon)$, it suggests that adaptivity
 396 *might* mitigate the effect of large privacy noise on performance. **Intuitively**, the $\text{sign}[\cdot]$ effectively
 397 **clips the privatized gradient signal, capping the update magnitude and reducing sensitivity to noise**
 398 **corruption**.

399 **Theorem 4.4** Let f be L -smooth and $K_3 := \max \left\{ \sqrt{\frac{\sigma_\gamma^2 \epsilon^2}{BT} + \frac{C^2 \Phi^2}{B^2}}, \frac{\sigma_\gamma \Phi}{B} \sqrt{d} \right\}$. Then,

$$\mathbb{E} [\|\nabla f(X_{\tilde{t}})\|_2^2] \lesssim K_3 \left(\frac{f(X_0)}{\eta \sqrt{T}} + \eta d L \sqrt{T} \right) \frac{1}{\epsilon}, \quad (11)$$

403 where \tilde{t} is a random time with uniform distribution over $[0, \tau]$.

405 **Takeaway:** Theorem 4.3 suggests that the privacy noise directly enters the convergence dynamics
 406 of DP-SignSGD, making its behavior qualitatively different from DP-SGD: The center column of
 407 Figure 2 confirms that DP-SignSGD converges faster for larger ϵ . Additionally, it also shows that
 408 it never diverges as drastically as DP-SGD for small ϵ . This is better shown in Figure 1, where we
 409 validate that the asymptotic loss scales with $\frac{1}{\epsilon}$, while that of DP-SGD scales with $\frac{1}{\epsilon^2}$. Therefore,
 410 adaptive methods are preferable in high-privacy settings, and all these insights are verified also for
 411 DP-Adam and generalize to the test loss (Figure C.4).

4.1.3 WHEN ADAPTIVITY REALLY MATTERS UNDER FIXED HYPERPARAMETERS.

413 In this subsection, we quantify when an adaptive method such as DP-SignSGD achieves better
 414 utility than DP-SGD. To this end, we compare *Privacy-Utility* terms of Phase 2 for both methods
 415 and derive conditions on the two sources of noise that govern the dynamics: the batch noise size σ_γ
 416 and the privacy budget ϵ .

417 **Theorem 4.5** If $\sigma_\gamma^2 \geq B$, then DP-SignSGD always achieves a better privacy-utility trade-off
 418 than DP-SGD. If $\sigma_\gamma^2 < B$, there exists a critical privacy level $\epsilon^* = \sqrt{\frac{C^2 TB}{n^2(B - \sigma_\gamma^2)} \log\left(\frac{1}{\delta}\right)}$ such that
 419 DP-SignSGD outperforms DP-SGD in utility whenever $\epsilon < \epsilon^*$.

421 **Takeaway:** This result makes the comparison explicit: *i*) Under **large batch noise** ($\sigma_\gamma^2 \geq B$),
 422 DP-SignSGD achieves a better utility than DP-SGD; *ii*) Under **small batch noise** ($\sigma_\gamma^2 < B$), the
 423 best optimizer depends on the privacy budget. For strict privacy ($\epsilon < \epsilon^*$), DP-SignSGD has better
 424 utility, while for looser privacy ($\epsilon > \epsilon^*$), DP-SGD achieves better overall performance. Thus, ϵ^*
 425 marks the threshold at which the advantage shifts from adaptive to non-adaptive methods when batch
 426 noise is small. By contrast, when batch noise is large, adaptive methods are already known to be
 427 more effective (Compagnoni et al., 2025b;a), and the effect of DP noise is only marginal relative
 428 to the intrinsic stochasticity of the gradients. We verify this result empirically in Figure 3: As we
 429 increase the batch size B , ϵ^* decreases, in accordance with our theoretical prediction.

430 **Practical Implication.** If hyperparameter re-tuning is infeasible and the target regime involves
 431 stronger privacy constraints, e.g., lower privacy budget ϵ , or high stochasticity from small batches,
 adaptive methods are preferable. Otherwise, DP-SGD is the method of choice.

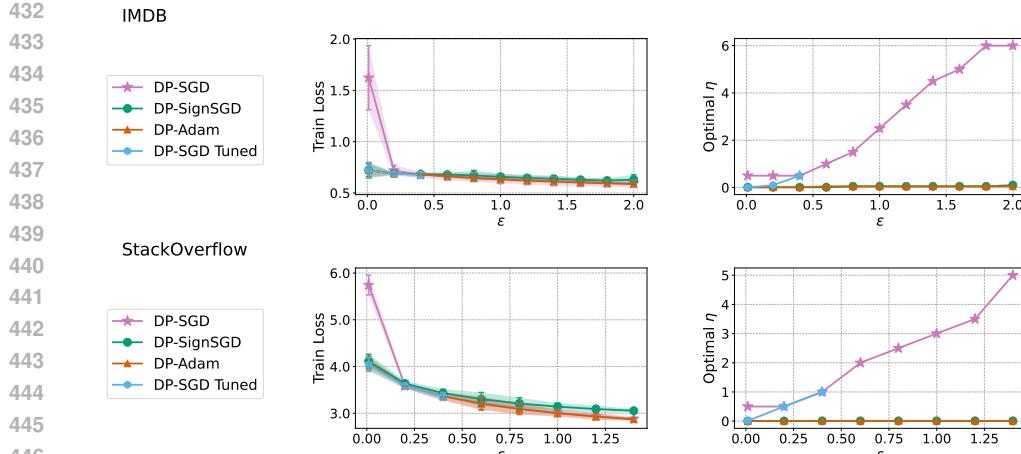


Figure 4: Empirical verification of Thm. 4.6 and Thm. 4.7 under Protocol B on the **IMDB dataset (Top Row)** and on the **StackOverflow dataset (Bottom Row)**. We tune (η, C) of each optimizer for each ϵ and confirm that: *i*) all methods achieve comparable performance across privacy budgets; *ii*) the optimal η of DP-SGD scales linearly with ϵ , while that of adaptive methods is essentially ϵ -independent; *iii*) failing to sweep over the “best” range of learning rates causes DP-SGD to severely underperform, whereas adaptive methods are resilient. On the **left**, DP-SGD degrades sharply for small ϵ . Indeed, the **right** panels shows that the selected optimal η flattens out, while the theoretical one would have linearly decayed more: The “best” η was simply missing from the grid. *A posteriori*, re-running the sweep with a larger grid (DP-SGD Tuned) recovers the scaling law and matches the performance of adaptive methods. For experimental details see Appendix C.5.

4.2 PROTOCOL B: BEST-TUNED HYPERPARAMETERS

We now mirror standard practice by allowing (η, C) to be *tuned* over an extensive grid search for each target privacy budget ϵ . In contrast to Protocol A, this leads us to derive the theoretical optimal learning rates, which, just as in empirical tuning, are allowed to depend on ϵ explicitly.

To select the optimal learning rate η^* for DP-SGD, we minimize the bound in Thm. 4.2 and consequently derive the implied optimal privacy-utility trade-off for DP-SGD in the L -smooth case.

Theorem 4.6 (DP-SGD) *Let $\eta^* = \min \left\{ \sqrt{\frac{f(X_0)}{dLT\sigma_\gamma^2}}, \sqrt{\frac{f(X_0)}{dL}} \frac{\epsilon n}{CT} \right\}$, then the expected gradient norm bound of DP-SGD is $\tilde{O} \left(\frac{C \sqrt{dL f(X_0)}}{\epsilon n} \right)$, as we ignore logarithmic terms and those decaying in T .*

This result aligns with the best-known privacy-utility trade-off obtained in prior works in these settings (Koloskova et al., 2023; Bassily et al., 2014). Importantly, we notice that the optimal learning rate of DP-SGD scales linearly in ϵ and that the resulting asymptotic performance scales like $\frac{1}{\epsilon}$.

To derive the optimal learning rate η^* of DP-SignSGD, we minimize the bound in Theorem 4.4, and derive a privacy-utility trade-off in the L -smooth case.

Theorem 4.7 (DP-SignSGD) *Let $\eta^* = \sqrt{\frac{f(X_0)}{dLT}}$. The expected asymptotic gradient norm bound of DP-SignSGD is $\tilde{O} \left(\frac{C \sqrt{dL f(X_0)}}{\epsilon n} \right)$, as we ignore logarithmic terms and those decaying in T .*

Importantly, we observe that the asymptotic neighborhood of DP-SignSGD matches that of DP-SGD, while the optimal learning rate is independent of ϵ . This suggests that adaptivity automatically handles the privacy noise injection: This facilitates the transferability of optimal parameters to setups that require higher privacy. In contrast, DP-SGD needs retuning of the hyperparameters.

Takeaway: Our theory shows that while optimal learning rate scalings differ, the induced neighborhoods match. As shown in Figure 4, our experiments verify that: *i*) DP-SGD, DP-SignSGD, and DP-Adam exhibit similar asymptotic performance across a broad range of ϵ , including very small values; *ii*) the optimal learning rate of DP-SGD is linear in ϵ , while those of adaptive methods are seemingly independent of it.

486 **Practical implication.** Hyperparameter searches are not free under DP: each evaluation consumes
 487 a portion of the privacy budget (Papernot & Steinke, 2021), making fine learning-rate grids costly.
 488 This asymmetrically impacts the two optimizers. For DP-SGD, the optimal step size scales linearly
 489 with ε (Thm. 4.6), so the “right” η^* moves as privacy tightens. If a fixed sweep grid misses a value
 490 close to η^* , the performance of DP-SGD can degrade sharply. This is illustrated in our experiments
 491 (Fig. 4): in the left panel, the performance of DP-SGD collapses because the selected “optimal”
 492 η plateaus instead of decaying linearly as predicted (right panel) — the true η^* was simply absent
 493 from the grid. By contrast, the optimal step size of DP-SignSGD (and empirically DP-Adam)
 494 is essentially ε -invariant (Thm. 4.7), so a single well-chosen η transfers across privacy levels with
 495 little or no re-tuning. This mechanism also helps explain prior empirical reports that non-adaptive
 496 methods deteriorate more severely under stricter privacy (Zhou et al., 2020b, Fig. 1), (Li et al.,
 497 2023, Fig. 5), (Asi et al., 2021, Fig. 2): a plausible cause is that their fixed grids did not track the
 498 ε -dependent η^* for DP-SGD. Importantly, when both optimizers are carefully tuned, DP-SGD and
 499 DP-Adam achieve matching performance in large-scale LLM fine-tuning (Li et al., 2021a, App. S).

500 5 CONCLUSION

501 We studied how differential privacy noise interacts with adaptive compared to non-adaptive optimi-
 502 zation through the lens of SDEs: To our knowledge, this is the first SDE-based analysis of DP
 503 optimizers. Our results include explicit upper bounds on the expected loss and gradient norm, opti-
 504 mal learning rates, as well as the first characterization of stationary distributions for DP optimizers.

505 Under a *fixed-hyperparameter* scenario (Protocol A), the analysis reveals a sharp contrast: *i*)
 506 DP-SGD converges at a speed independent of the privacy budget ε while incurring a $\mathcal{O}(1/\varepsilon^2)$
 507 privacy-utility trade-off; *ii*) DP-SignSGD converges at a speed proportional to ε while exhib-
 508 iting a $\mathcal{O}(1/\varepsilon)$ privacy-utility trade-off. Additionally, when batch noise is large, adaptive methods
 509 dominate in terms of utility, as the effect of DP noise is marginal compared to the intrinsic stochas-
 510 ticity of the gradients, confirming known insights from non-private optimization. When batch noise
 511 is small, the preferable method depends on the privacy budget: for strict privacy, DP-SignSGD
 512 yields better utility, while for looser privacy, DP-SGD achieves better overall performance.

513 Under a *best-tuned* scenario (Protocol B), the picture changes: theory and experiments agree that
 514 the optimal learning rate of DP-SGD scales linearly with ε , whereas the optimal learning rate of
 515 DP-SignSGD (and empirically DP-Adam) is approximately ε -independent. With this tuning, the
 516 induced privacy-utility trade-offs match in order and the methods achieve comparable asymptotic
 517 performance, including at very small ε . A practical implication is that adaptive methods require less
 518 re-tuning if regulations mandate tighter privacy budgets.

519 We validated these theoretical insights on both synthetic and real datasets. Importantly, we also
 520 demonstrated that the qualitative behavior observed for DP-SignSGD extends empirically to
 521 DP-Adam and to test metrics, underscoring the strength and generality of our framework.

522 **Practitioner guidance.** Under higher privacy requirements, e.g., regulations mandate a smaller
 523 ε , if per- ε re-tuning of the hyperparameters is impractical because retraining/tuning is expected to
 524 be costly (Protocol A), prefer an *adaptive* private optimizer such as DP-SignSGD (or DP-Adam):
 525 their performance scales more favorably as ε decreases compared to DP-SGD.

526 When re-tuning is feasible (Protocol B): Both DP-SGD and adaptive methods can reach comparable
 527 asymptotic performance. However, hyperparameter searches are not free under DP: each sweep
 528 consumes additional privacy budget (Papernot & Steinke, 2021), making fine grids expensive. This
 529 creates an asymmetric risk: DP-SGD requires an ε -dependent learning rate ($\eta^* \propto \varepsilon$), so if the sweep
 530 grid does not track this scaling, its performance can degrade sharply. In contrast, adaptive methods
 531 retain a portable, ε -independent learning rate, making them more robust and less costly to tune
 532 across privacy levels.

533
 534
 535
 536
 537
 538
 539

540 REFERENCES
541

542 Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and
543 Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC*
544 *Conference on Computer and Communications Security*, CCS’16. ACM, October 2016. doi:
545 10.1145/2976749.2978318. URL <http://dx.doi.org/10.1145/2976749.2978318>.

546 Sungjin Ahn, Anoop Korattikara, and Max Welling. Bayesian posterior sampling via stochastic
547 gradient fisher scoring. *arXiv preprint arXiv:1206.6380*, 2012.

548 Galen Andrew, Om Thakkar, Brendan McMahan, and Swaroop Ramaswamy. Differentially private
549 learning with adaptive clipping. *Advances in Neural Information Processing Systems*, 34:17455–
550 17466, 2021.

551 Hilal Asi, John Duchi, Alireza Fallah, Omid Javidi, and Kunal Talwar. Private adaptive gradient
552 methods for convex optimization, 2021. URL <https://arxiv.org/abs/2106.13756>.

553 Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight
554 analyses via couplings and divergences. *Advances in neural information processing systems*, 31,
555 2018.

556 Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy profiles and amplification by subsampling.
557 *Journal of Privacy and Confidentiality*, 10(1), 2020.

558 Lukas Balles and Philipp Hennig. Dissecting adam: The sign, magnitude and variance of stochastic
559 gradients. In *International Conference on Machine Learning*, pp. 404–413. PMLR, 2018.

560 Raef Bassily, Adam Smith, and Abhradeep Thakurta. Differentially private empirical risk minimization:
561 Efficient algorithms and tight error bounds. *arXiv preprint arXiv:1405.7085*, 2014.

562 Tianqi Chen, Emily Fox, and Carlos Guestrin. Stochastic gradient hamiltonian monte carlo. In
563 *International conference on machine learning*, pp. 1683–1691. PMLR, 2014.

564 Xiangyi Chen, Steven Z Wu, and Mingyi Hong. Understanding gradient clipping in private sgd: A
565 geometric perspective. *Advances in Neural Information Processing Systems*, 33:13773–13782,
566 2020.

567 Enea Monzio Compagnoni, Luca Biggio, Antonio Orvieto, Frank Norbert Proske, Hans Kersting,
568 and Aurelien Lucchi. An sde for modeling sam: Theory and insights. In *International Conference
569 on Machine Learning*, pp. 25209–25253. PMLR, 2023.

570 Enea Monzio Compagnoni, Antonio Orvieto, Hans Kersting, Frank Proske, and Aurelien Lucchi.
571 Sdes for minimax optimization. In *International Conference on Artificial Intelligence and Statistics*,
572 pp. 4834–4842. PMLR, 2024.

573 Enea Monzio Compagnoni, Rustem Islamov, Frank Norbert Proske, and Aurelien Lucchi. Unbiased
574 and sign compression in distributed learning: Comparing noise resilience via SDEs. In *The 28th
575 International Conference on Artificial Intelligence and Statistics*, 2025a.

576 Enea Monzio Compagnoni, Tianlin Liu, Rustem Islamov, Frank Norbert Proske, Antonio Orvieto,
577 and Aurelien Lucchi. Adaptive methods through the lens of sdes: Theoretical insights on the role
578 of noise, 2025b. URL <https://arxiv.org/abs/2411.15958>.

579 Enea Monzio Compagnoni, Tianlin Liu, Rustem Islamov, Frank Norbert Proske, Antonio Orvieto,
580 and Aurelien Lucchi. Adaptive methods through the lens of SDEs: Theoretical insights on the
581 role of noise. In *The Thirteenth International Conference on Learning Representations*, 2025c.

582 Soham De, Leonard Berrada, Jamie Hayes, Samuel L. Smith, and Borja Balle. Unlocking high-
583 accuracy differentially private image classification through scale, 2022. URL [https://arxiv.org/ab-s/2204.13650](https://arxiv.org/ab-
584 s/2204.13650).

585 John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and
586 stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 07 2011.

594 Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends*
 595 *Theor. Comput. Sci.*, 9(3–4):211–407, August 2014. ISSN 1551-305X. doi: 10.1561/0400000042.
 596 URL <https://doi.org/10.1561/0400000042>.

597

598 Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data,
 599 ourselves: Privacy via distributed noise generation. In Serge Vaudenay (ed.), *Advances in Cryptology - EUROCRYPT 2006*, pp. 486–503, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
 600 ISBN 978-3-540-34547-3.

601

602 EU. Eu ai act: First regulation on artificial intelligence. <https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>, 2023. Accessed: 2025-09-24.

603

604

605 EU. Artificial intelligence and next generation technologies. <https://www.enisa.europa.eu/topics/artificial-intelligence-and-next-gen-technologies>, 2024. Accessed: 2025-09-24.

606

607

608 Arun Ganesh, Brendan McMahan, and Abhradeep Thakurta. On design principles for private adaptive optimizers, 2025. URL <https://arxiv.org/abs/2507.01129>.

609

610 Roan Gylberth, Risman Adnan, Setiadi Yazid, and T. Basaruddin. Differentially private optimization
 611 algorithms for deep neural networks. In *2017 International Conference on Advanced Computer
 612 Science and Information Systems (ICACSIS)*, pp. 387–394, 2017. doi: 10.1109/ICACSIS.2017.8
 613 355063.

614

615 Uwe Helmke and John B Moore. *Optimization and Dynamical Systems*. Springer London, 1st
 616 edition, 1994.

617

618 White House. Fact sheet: President biden issues executive order on safe, secure, and trustworthy
 619 artificial intelligence. <https://bidenwhitehouse.archives.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/>, 2023. Accessed: 2025-09-24.

620

621 Stanisław Jastrzebski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Ben-
 622 gio, and Amos Storkey. Three factors influencing minima in sgd. *ICANN 2018*, 2018.

623

624 Richeng Jin and Huaiyu Dai. Noisy SIGNSGD is more differentially private than you (might) think.
 625 In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=thCqMz1ZXw>.

626

627 Kaggle. Stack overflow data on kaggle. <https://www.kaggle.com/datasets/stackoverflow/stackoverflow>, 2022.

628

629

630 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

631

632 Anastasia Koloskova, Hadrien Hendrikx, and Sebastian U. Stich. Revisiting Gradient Clipping:
 633 Stochastic bias and tight convergence guarantees. In *ICML 2023 - 40th International Conference on Machine Learning*, pp. 1–19, Honolulu, Hawaii, United States, July 2023. URL <https://openreview.net/pdf?id=C3DXiFTrve>.

634

635

636 Harold Kushner and G George Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35. Springer Science & Business Media, 2003.

637

638

639 Bingrui Li, Wei Huang, Andi Han, Zhanpeng Zhou, Taiji Suzuki, Jun Zhu, and Jianfei Chen. On the
 640 optimization and generalization of two-layer transformers with sign gradient descent, 2025. URL
 641 <https://arxiv.org/abs/2410.04870>.

642

643 Qianxiao Li, Cheng Tai, and E Weinan. Stochastic modified equations and adaptive stochastic
 644 gradient algorithms. In *International Conference on Machine Learning*, pp. 2101–2110. PMLR,
 2017.

645

646 Qianxiao Li, Cheng Tai, and E Weinan. Stochastic modified equations and dynamics of stochastic
 647 gradient algorithms i: Mathematical foundations. *The Journal of Machine Learning Research*, 20(1):1474–1520, 2019.

648 Tian Li, Manzil Zaheer, Sashank J. Reddi, and Virginia Smith. Private adaptive optimization with
 649 side information, 2022. URL <https://arxiv.org/abs/2202.05963>.

650

651 Tian Li, Manzil Zaheer, Ken Ziyu Liu, Sashank J. Reddi, H. Brendan McMahan, and Virginia
 652 Smith. Differentially private adaptive optimization with delayed preconditioners, 2023. URL
 653 <https://arxiv.org/abs/2212.00309>.

654 Xuechen Li, Florian Tramer, Percy Liang, and Tatsunori Hashimoto. Large language models can be
 655 strong differentially private learners. *arXiv preprint arXiv:2110.05679*, 2021a.

656

657 Zhiyuan Li, Sadhika Malladi, and Sanjeev Arora. On the validity of modeling SGD with stochastic
 658 differential equations (SDEs). In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan
 659 (eds.), *Advances in Neural Information Processing Systems*, 2021b.

660 Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher
 661 Potts. Learning word vectors for sentiment analysis. In Dekang Lin, Yuji Matsumoto, and Rada
 662 Mihalcea (eds.), *Proceedings of the 49th Annual Meeting of the Association for Computational
 663 Linguistics: Human Language Technologies*, pp. 142–150, Portland, Oregon, USA, June 2011.
 664 Association for Computational Linguistics. URL <https://aclanthology.org/P11-1015/>.

665 Sadhika Malladi, Kaifeng Lyu, Abhishek Panigrahi, and Sanjeev Arora. On the SDEs and scaling
 666 rules for adaptive gradient algorithms. In *Advances in Neural Information Processing Systems*,
 667 2022.

668

669 Stephan Mandt, Matthew Hoffman, and David Blei. A variational analysis of stochastic gradient
 670 algorithms. In *International conference on machine learning*, pp. 354–363. PMLR, 2016.

671 Stephan Mandt, Matthew D Hoffman, and David M Blei. Stochastic gradient descent as approximate
 672 bayesian inference. *JMLR* 2017, 2017.

673

674 Noah Marshall, Ke Liang Xiao, Atish Agarwala, and Elliot Paquette. To clip or not to clip: the
 675 dynamics of SGD with gradient clipping in high-dimensions. In *The Thirteenth International
 676 Conference on Learning Representations*, 2025.

677 Ryan McKenna, Yangsibo Huang, Amer Sinha, Borja Balle, Zachary Charles, Christopher A.
 678 Choquette-Choo, Badih Ghazi, George Kaassis, Ravi Kumar, Ruibo Liu, Da Yu, and Chiyuan
 679 Zhang. Scaling laws for differentially private language models, 2025. URL <https://arxiv.org/abs/2501.18914>.

680

681 H. Brendan McMahan and Matthew Streeter. Adaptive bound optimization for online convex opti-
 682 mization, 2010. URL <https://arxiv.org/abs/1002.4908>.

683

684 Panayotis Mertikopoulos and Mathias Staudigl. On the convergence of gradient-like flows with
 685 noisy gradient input. *SIAM Journal on Optimization*, 28(1):163–197, 2018.

686

687 GN Mil'shtein. Weak approximation of solutions of systems of stochastic differential equations.
 688 *Theory of Probability & Its Applications*, 30(4):750–766, 1986.

689

690 Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th computer security foundations sym-
 691 posium (CSF)*, pp. 263–275. IEEE, 2017.

692

693 Ilya Mironov, Kunal Talwar, and Li Zhang. Rényi differential privacy of the sampled gaussian
 694 mechanism, 08 2019.

695

696 NITS. Nist offers draft guidance for evaluating privacy protection technique in the ai era. [https://www.nist.gov/news-events/news/2023/12/nist-offers-draft-guidance-evaluating-privacy-
 697 protection-technique-ai-era](https://www.nist.gov/news-events/news/2023/12/nist-offers-draft-guidance-evaluating-privacy-protection-technique-ai-era), 2023a. Accessed: 2025-09-24.

698

699 NITS. Artificial intelligence risk management framework (ai rmf 1.0). Technical report, National
 700 Institute of Standards and Technology, 2023b. URL <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

701

Nicolas Papernot and Thomas Steinke. Hyperparameter tuning with renyi differential privacy. *arXiv
 702 preprint arXiv:2110.03620*, 2021.

702 Courtney Paquette, Kiwon Lee, Fabian Pedregosa, and Elliot Paquette. Sgd in the large: Average-
 703 case analysis, asymptotics, and stepsize criticality. In *Conference on Learning Theory*, pp. 3548–
 704 3626. PMLR, 2021.

705

706 Hanyang Peng, Shuang Qin, Yue Yu, Fangqing Jiang, Hui Wang, and Zhouchen Lin. Simple con-
 707 vergence proof of adam from a sign-like descent perspective, 2025. URL <https://arxiv.org/abs/2507.05966>.

708

709 Venkatadheeraj Pichapati, Ananda Theertha Suresh, Felix X Yu, Sashank J Reddi, and Sanjiv Kumar.
 710 Adaclip: Adaptive clipping for private sgd. *arXiv preprint arXiv:1908.07643*, 2019.

711

712 Maxim Raginsky and Jake Bouvrie. Continuous-time stochastic mirror descent on a network: Vari-
 713 ance reduction, consensus, convergence. In *2012 IEEE 51st IEEE Conference on Decision and
 714 Control (CDC)*, pp. 6793–6800. IEEE, 2012.

715

716 Umut Simsekli, Levent Sagun, and Mert Gurbuzbalaban. A tail-index analysis of stochastic gradient
 717 noise in deep neural networks. In *International Conference on Machine Learning*, 2019.

718

719 Samuel L. Smith, Benoit Dherin, David G. T. Barrett, and Soham De. On the origin of implicit
 720 regularization in stochastic gradient descent. *ArXiv*, abs/2101.12176, 2021.

721

722 Mandt Stephan, Matthew D Hoffman, David M Blei, et al. Stochastic gradient descent as approxi-
 723 mate bayesian inference. *Journal of Machine Learning Research*, 18(134):1–35, 2017.

724

725 Qiaoyue Tang and Mathias Lécuyer. Dp-adam: Correcting dp bias in adam’s second moment esti-
 726 mation, 2023. URL <https://arxiv.org/abs/2304.11208>.

727

728 Qiaoyue Tang, Frederick Shpilevskiy, and Mathias Lécuyer. Dp-adambc: Your dp-adam is actually
 729 dp-sgd (unless you apply bias correction), 2023. URL <https://arxiv.org/abs/2312.14334>.

730

731 TensorFlow Federated. Tensorflow federated stack overflow dataset. https://www.tensorflow.org/federated/api_docs/python/tff/simulation/datasets/stackoverflow/load_data, 2022.

732

733 Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running
 734 average of its recent magnitude., 2012. URL https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf.

735

736 Roman Vershynin. *High-dimensional probability : an introduction with applications in data science*.
 737 Cambridge series in statistical and probabilistic mathematics ; 47. Cambridge University Press,
 Cambridge, United Kingdom ;, 2018. ISBN 9781108415194.

738

739 Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. Subsampled renyi differential pri-
 740 vacy and analytical moments accountant. In Kamalika Chaudhuri and Masashi Sugiyama (eds.),
 741 *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*,
 742 volume 89 of *Proceedings of Machine Learning Research*, pp. 1226–1235. PMLR, 16–18 Apr 2019.

743

744 Jingfeng Wu, Wenqing Hu, Haoyi Xiong, Jun Huan, Vladimir Braverman, and Zhanxing Zhu. On the
 745 noisy gradient descent that generalizes as sgd. In *International Conference on Machine Learning*,
 746 pp. 10367–10376. PMLR, 2020.

747

748 Ke Liang Xiao, Noah Marshall, Atish Agarwala, and Elliot Paquette. Exact risk curves of signSGD
 749 in high-dimensions: quantifying preconditioning and noise-compression effects. In *Forty-second
 750 International Conference on Machine Learning*, 2025.

751

752 Zeke Xie, Li Yuan, Zhanxing Zhu, and Masashi Sugiyama. Positive-negative momentum: Ma-
 753 nipulating stochastic gradient noise to improve generalization. In *Proceedings of the 38th Inter-
 754 national Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning
 755 Research*, pp. 11448–11458. PMLR, 18–24 Jul 2021.

756

757 Jim Zhao, Aurelien Lucchi, Frank Norbert Proske, Antonio Orvieto, and Hans Kersting. Batch size
 758 selection by stochastic optimal control. In *Has it Trained Yet? NeurIPS 2022 Workshop*, 2022.

756 Pan Zhou, Jia Shi Feng, Chao Ma, Caiming Xiong, Steven Chu Hong Hoi, et al. Towards theoreti-
757 cally understanding why sgd generalizes better than adam in deep learning. *Advances in Neural*
758 *Information Processing Systems*, 33:21285–21296, 2020a.

759

760 Yingxue Zhou, Xiangyi Chen, Mingyi Hong, Zhiwei Steven Wu, and Arindam Banerjee. Private
761 stochastic non-convex optimization: Adaptive algorithms and tighter generalization bounds,
762 2020b. URL <https://arxiv.org/abs/2006.13501>.

763

764 Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. The anisotropic noise in stochastic
765 gradient descent: Its behavior of escaping from sharp minima and regularization effects. *ICML*,
766 2019.

767

768 Difan Zou, Yuan Cao, Yuanzhi Li, and Quanquan Gu. Understanding the generalization of adam in
769 learning neural networks with proper regularization, 2021. URL <https://arxiv.org/abs/2108.11371>.

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 811 812 813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836 837 838 839 840 841 842 843 844 845 846 847 848 849 850 851 852 853 854 855 856 857 858 859 860 861 862 863 Appendix

CONTENTS

A Technical Results	16
B Theoretical Framework	18
B.1 DP-SGD	20
B.2 DP-SignSGD	26
C Experimental Details and Additional Results	35
C.1 DP-SGD and DP-SignSGD: SDE Validation (Figure C.1).	37
C.2 Asymptotic Loss Bound (Figures 1 and C.4)	37
C.3 Convergence Speed Analysis (Figure 2)	38
C.4 When Adaptivity Really Matters (Figure 3)	38
C.5 Best-Tuned Hyperparameters (Figures 4)	39
C.6 Stationary Distributions	39
C.7 Additional Results — Test Loss	40
D Limitations	41

837 838 839 840 841 842 843 844 845 846 847 848 849 850 851 852 853 854 855 856 857 858 859 860 861 862 863 A TECHNICAL RESULTS

In this section, we introduce some technical results used in the derivation of the SDEs.

Lemma A.1 *Let $X \sim \mathcal{N}(\mu, \sigma^2 I_d)$ and fix a tolerance $\epsilon > 0$. If $\frac{\|\mu\|_2^2}{2\sigma^2(d+2)} < \epsilon$, for $d \rightarrow \infty$, we have that $\mathbb{E}\left(\frac{X}{\|X\|_2}\right) = \sqrt{\frac{1}{d}\frac{\mu}{\sigma}} + \mathcal{O}\left(\frac{1}{d^{3/2}}\right)$.*

Proof: Let us remember that if $X \sim \mathcal{N}(\mu, \sigma^2 I_d)$,

$$\mathbb{E}\left(\frac{X}{\|X\|_2^k}\right) = \frac{\Gamma\left(\frac{d}{2} + 1 - \frac{k}{2}\right)}{(2\sigma^2)^{k/2} \Gamma\left(\frac{d}{2} + 1\right)} {}_1F_1\left(\frac{k}{2}; \frac{d+2}{2}; -\frac{\|\mu\|_2^2}{2\sigma^2}\right) \mu, \quad (12)$$

where ${}_1F_1(a; b; z)$ is Kummer's confluent hypergeometric function. We know that

$$\lim_{d \rightarrow \infty} \frac{\Gamma\left(\frac{d}{2} + 1 - \frac{1}{2}\right)}{\Gamma\left(\frac{d}{2} + 1\right)} \underset{k=1}{\sim} \sqrt{\frac{2}{d}} + \mathcal{O}\left(\frac{1}{d^{3/2}}\right). \quad (13)$$

Let $z = \frac{\|\mu\|_2^2}{2\sigma^2}$. If $d > z$, by expanding the series, we have

$${}_1F_1\left(\frac{1}{2}; \frac{d}{2} + 1; -z\right) = \sum_{n \geq 0} \frac{a^{(n)}(-z)^n}{b^{(n)} n!} = 1 - \frac{z}{d+2} + \mathcal{O}\left(\frac{z}{d}\right)^2 < 1 - \epsilon. \quad (14)$$

Combining everything together, we obtain $\mathbb{E}\left(\frac{X}{\|X\|_2}\right) = \sqrt{\frac{1}{d}\frac{\mu}{\sigma}} + \mathcal{O}\left(\frac{1}{d^{3/2}}\right)$. \square

Lemma A.2 *Let $K(\nu) = \sqrt{\frac{2}{\nu}} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})}$ and $X \sim t_\nu(\mu, \sigma^2 I_d)$, for $\nu \geq 1$. Fix a tolerance $\epsilon > 0$: If $\frac{\|\mu\|_2^2}{2\sigma^2(d+2)} < \epsilon$, for $d \rightarrow \infty$, we have that $\mathbb{E}\left(\frac{X}{\|X\|_2}\right) = K(\nu) \sqrt{\frac{1}{d}\frac{\mu}{\sigma}} + \mathcal{O}\left(\frac{1}{d^{3/2}}\right)$.*

864 **Proof:** One can write $X = \mu + \frac{\sigma Z}{\sqrt{S/\nu}}$, where $Z \sim \mathcal{N}(0, I_d)$ and $S \sim \chi_\nu^2$ are independent. Define
 865 $\tau = \frac{\sigma}{\sqrt{S/\nu}}$, then, conditioning on S and applying Lemma A.1, we have
 866
 867

$$868 \quad 869 \quad \mathbb{E} \left[\frac{X}{\|X\|_2} \middle| S \right] = \sqrt{\frac{1}{d} \frac{\mu}{\tau} + \epsilon \mathcal{O}(d^{-3/2})}. \quad (15)$$

871 Remembering that $\mathbb{E}[\sqrt{S}] = \sqrt{2} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})}$, we have
 872
 873

$$874 \quad 875 \quad \mathbb{E} \left[\frac{X}{\|X\|_2} \right] = \mathbb{E}_S \left[\left(\sqrt{\frac{1}{d} \frac{\mu}{\tau}} + \epsilon \mathcal{O}(d^{-3/2}) \right) \right] \quad (16)$$

$$876 \quad 877 \quad = \left(\sqrt{\frac{1}{d} \frac{\mu}{\sigma \sqrt{\nu}}} \mathbb{E}_S[\sqrt{S}] + \epsilon \mathcal{O}(d^{-3/2}) \right) \quad (17)$$

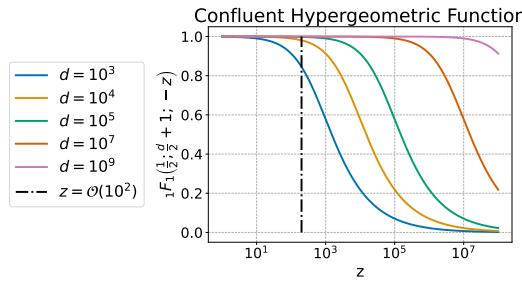
$$878 \quad 879 \quad = K(\nu) \sqrt{\frac{1}{d} \frac{\mu}{\sigma}} + \epsilon \mathcal{O}\left(\frac{1}{d^{3/2}}\right). \quad (18)$$

880
 881
 882
 883
 884

885 **Remark A.1** As discussed in the main paper, our analysis is based on the following two assumptions:
 886

887 i) The number of trainable parameters is large, specifically $d = \Omega(10^4)$;
 888
 889 ii) The signal-to-noise ratio satisfies $\frac{\|\nabla f(x)\|_2^2}{2\sigma_\gamma^2} \ll d$.

890 First, they ensure that the approximation of the confluent hypergeometric function in Equation 14
 891 is highly accurate. Second, neither condition is restrictive for modern deep learning models.
 892 The dimensionality assumption is trivially satisfied by contemporary architectures, which routinely
 893 have millions of parameters. Regarding the second assumption, Malladi et al. (2022) empirically
 894 measured the signal-to-noise ratio $\frac{\|\nabla f(x)\|_2^2}{2\sigma_\gamma^2}$ across a wide range of large-scale architectures and
 895 datasets, and consistently found values of at most $O(10^2)$. Thus, the regime in which our approxi-
 896 mation is valid closely matches the regime observed in practice. This is further supported by our
 897 experimental results, which confirm our theoretical predictions across all models and tasks consid-
 898 ered in this paper. In Figure A.1, we numerically evaluate the confluent hypergeometric function
 899 for varying values of d , and show that, for sufficiently large parameter counts, the approximation
 900 remains tight throughout the realistic signal-to-noise ratio range reported in (Malladi et al., 2022).
 901
 902



914 Figure A.1: Numerical validation of the approximation used in Equation 14. For several values of
 915 d , we plot the confluent hypergeometric function as a function of the signal-to-noise ratio z . In
 916 the realistic range observed in (Malladi et al., 2022), approximating this function by 1 is extremely
 917 accurate.

918 **B THEORETICAL FRAMEWORK**
919

920 In this section, we introduce the theoretical framework, assumptions, and notations used to formally
921 derive the SDE models used in this paper. We briefly recall the definition of L -smoothness and μ -PL
922 functions. Then we introduce the set of functions of polynomial growth G .
923

924 **Definition B.1** A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is L -smooth if it is differentiable and its gradient is L -
925 Lipschitz continuous, namely
926

927
$$\|\nabla f(x) - \nabla f(y)\|_2 \leq L\|x - y\|_2 \quad \forall x, y. \quad (19)$$
928

929 **Definition B.2** A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ admitting a global minima x^* satisfies the Polyak-
930 Lojasiewicz inequality if, for some $\mu > 0$ and for all $x \in \mathbb{R}^d$, it holds
931

932
$$f(x) - f^* \leq \frac{1}{2\mu} \|\nabla f(x)\|_2^2. \quad (20)$$
933

934 In this case, we say that the function f is μ -PL.
935

936 **Definition B.3** Let G denote the set of continuous functions $g : \mathbb{R}^d \rightarrow \mathbb{R}$ of at most polynomial
937 growth, namely such that there exist positive integers $k_1, k_2 > 0$ such that $|g(x)| < k_1(1 + \|x\|_2^2)^{k_2}$,
938 for all $x \in \mathbb{R}^d$.
939

940 To simplify the notation, we will write
941

942
$$b(x + \eta) = b_0(x) + \eta b_1(x) + \mathcal{O}(\eta^2),$$
943

944 whenever there exists $g \in G$, independent of η , such that
945

946
$$|b(x + \eta) - b_0(x) - \eta b_1(x)| \leq g(x)\eta^2.$$
947

948 We now introduce the definition of weak approximation, which formalizes in which sense the solution
949 to an SDE, which is a continuous-time random process, models a discrete-time optimizer.
950

951 **Definition B.4** Let $0 < \eta < 1$, $\tau > 0$ and $T = \lfloor \frac{\tau}{\eta} \rfloor$. We say that a continuous time process X_t over
952 $[0, \tau]$, is an order α weak approximation of a discrete process x_k , for $k = 0, \dots, N$, if for every
953 $g \in G$, there exists M , independent of η , such that for all $k = 0, 1, \dots, N$
954

955
$$|\mathbb{E}g(X_{k\eta}) - \mathbb{E}g(x_k)| \leq M\eta^\alpha.$$
956

957 This framework focuses on approximation in a *weak sense*, meaning in distribution rather than path-
958 wise. Since G contains all polynomials, all the moments of both processes become closer at a rate
959 of η^α and thus their distributions. Thus, while the processes exhibit similar average behavior, their
960 sample paths may differ significantly, justifying the term weak approximation.
961

962 **Remark B.1** Our continuous-time models are derived using the standard order-1 weak-
963 approximation framework (see Section 2). In particular, Definition B.4 together with Theorems B.5
964 and B.10 show that the drift and covariance of the discrete updates match those of the SDE up
965 to $O(\eta^2)$, so that the weak error over finite horizons is $O(\eta)$. In line with the existing literature,
966 we therefore restrict attention to first-order SDEs; higher-order SDEs have been derived in special
967 cases but, to the best of our knowledge, have not led to additional practical insights. Finally, Figure
968 C.1 empirically confirms that our SDEs closely track their corresponding algorithms on simple
969 landscapes, following the standard validation practice in the field (Compagnoni et al., 2025c).
970

971 The key ingredient for deriving the SDE is given by the following result (see Theorem 1, (Li et al.,
972 2017)), which provides sufficient conditions to get a weak approximation in terms of the single step
973 increments of both X_t and x_k . Before stating the theorem, we list the regularity assumption under
974 which we are working.
975

976 **Assumption B.1** Assume that the following conditions are satisfied:
977

978

- $f, f_i \in \mathcal{C}_b^8(\mathbb{R}^d, \mathbb{R})$;
979

980

- f, f_i and its partial derivatives up to order 7 belong to G ;
- $\nabla f, \nabla f_i$ satisfy the following Lipschitz condition: there exists $L > 0$ such that

$$\|\nabla f(u) - \nabla f(v)\|_2 + \sum_{i=1}^d \|\nabla f_i(u) - \nabla f_i(v)\|_2 \leq L\|u - v\|_2;$$

- $\nabla f, \nabla f_i$ satisfy the following growth condition: there exists $M > 0$ such that

$$\|\nabla f(x)\|_2 + \sum_{i=1}^n \|\nabla f_i(x)\|_2 \leq M(1 + \|x\|_2).$$

Assumption B.2 Assume that the stochastic gradient can be written as $\nabla f_\gamma = \nabla f + Z_\gamma$. In Phase 1 (clipping regime), the batch noise Z_γ is modeled as heavy-tailed, e.g., a Student-t distribution with ν degrees of freedom and scale σ_γ : for $\nu = \infty$ we recover the Gaussian case, while if $\nu < 2$ the variance is unbounded and if $\nu = 1$ the distribution becomes a Cauchy, therefore the expectation is unbounded as well. In Phase 2 (non-clipping regime), the batch noise is modeled as a Gaussian of variance $\frac{\sigma_\gamma^2}{B}$, reflecting the averaging effect of i.i.d., per-sample gradients.

Remark B.2 The distinction between the two phases stems from the effect of per-example clipping on the noise distribution. In Phase 1, clipping is applied before the batch average, so the noise of each individual stochastic gradient is not smoothed by averaging and can remain strongly heavy-tailed; in this regime, a Gaussian model is no longer appropriate. We therefore model the Phase 1 noise as a multivariate Student-t, which both captures this heavy-tailed behaviour and admits tractable expressions for our SDE analysis, while recovering the Gaussian model used in Phase 2 in the limit of large degrees of freedom.

Lemma B.3 Let $0 < \eta < 1$. Consider a stochastic process $X_t, t \geq 0$ satisfying the SDE

$$dX_t = b(X_t)dt + \sqrt{\eta}\sigma(X_t)dW_t, \quad X_0 = x \quad (21)$$

where b, σ together with their derivatives belong to G . Define the one-step difference $\Delta = X_\eta - x$, and indicate the i -th component of Δ with Δ_i . Then we have

1. $\mathbb{E}\Delta_i = b_i\eta + \frac{1}{2} \left[\sum_{j=1}^d b_j \partial_j b_i \right] \eta^2 + \mathcal{O}(\eta^3) \quad \forall i = 1, \dots, d;$
2. $\mathbb{E}\Delta_i \Delta_j = [b_i b_j + \sigma \sigma_{ij}^\top] \eta^2 + \mathcal{O}(\eta^3) \quad \forall i, j = 1, \dots, d;$
3. $\mathbb{E} \prod_{j=1}^s \Delta_{i_j} = \mathcal{O}(\eta^3) \quad \forall s \geq 3, i_j = 1, \dots, d.$

All functions above are evaluated at x .

Theorem B.4 Let $0 < \eta < 1$, $\tau > 0$ and set $T = \lfloor \tau/\eta \rfloor$. Let Assumption B.1 hold and let X_t be a stochastic process as in Lemma B.3. Define $\bar{\Delta} = x_1 - x$ to be the increment of the discrete-time algorithm, and indicate the i -th component of $\bar{\Delta}$ with $\bar{\Delta}_i$. If in addition there exist $K_1, K_2, K_3, K_4 \in G$ so that

1. $|\mathbb{E}\Delta_i - \mathbb{E}\bar{\Delta}_i| \leq K_1(x)\eta^2, \quad \forall i = 1, \dots, d;$
2. $|\mathbb{E}\Delta_i \Delta_j - \mathbb{E}\bar{\Delta}_i \bar{\Delta}_j| \leq K_2(x)\eta^2, \quad \forall i, j = 1, \dots, d;$
3. $|\mathbb{E} \prod_{j=1}^s \Delta_{i_j} - \mathbb{E} \prod_{j=1}^s \bar{\Delta}_{i_j}| \leq K_3(x)\eta^2, \quad \forall s \geq 3, \forall i_j = 1, \dots, d;$
4. $\mathbb{E} \prod_{j=1}^s |\bar{\Delta}_{i_j}| \leq K_4(x)\eta^2, \quad \forall i_j = 1, \dots, d.$

Then, there exists a constant C so that for all $k = 0, 1, \dots, N$ we have

$$|\mathbb{E}g(X_{k\eta}) - \mathbb{E}g(x_k)| \leq C\eta. \quad (22)$$

We say Eq. 21 is an order 1 weak approximation of the update step of x_k .

1026 B.1 DP-SGD
1027

1028 This subsection provides the formal derivation of the SDE model for DP-SGD and formal statements
1029 of Theorem 4.1, Theorem 4.2, and Theorem B.9. Since, by construction, the dynamic of the method
1030 shifts stochastically between two phases, we first model and study each phase separately.
1031

1032 **Theorem B.5** *Let $0 < \eta < 1$, $\tau > 0$ and set $T = \lfloor \tau/\eta \rfloor$ and $K(\nu) = \sqrt{\frac{2}{\nu} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})}}$. Let $x_k \in \mathbb{R}^d$
1033 denote a sequence of DP-SGD iterations defined in Eq. 4. Assume Assumption B.1 and Assumption
1034 B.2. Let X_t be the solution of the following SDEs with initial condition $X_0 = x_0$:*
1035

1036 • Phase 1:

$$1037 dX_t = -\frac{CK(\nu)}{\sigma_\gamma \sqrt{d}} \nabla f(X_t) dt + \sqrt{\eta} \sqrt{\bar{\Sigma}(X_t)} dW_t, \quad (23)$$

1039 where $\bar{\Sigma}(x) = C^2 \left(\mathbb{E} \left[\frac{\nabla f_\gamma(x) \nabla f_\gamma(x)^\top}{\|\nabla f_\gamma(x)\|_2^2} \right] - \frac{K(\nu)^2}{\sigma_\gamma \sqrt{d}} \nabla f(x) \nabla f(x)^\top + \frac{\sigma_{DP}^2}{B^2} I_d \right)$.
1040

1041 • Phase 2:

$$1042 dX_t = -\nabla f(X_t) dt + \sqrt{\eta} \sqrt{\bar{\Sigma}(X_t)} dW_t, \quad (24)$$

1043 where $\bar{\Sigma}(x) = \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2} \right) I_d$.
1044

1045 Then, Eq. 23 and Eq. 24 are an order 1 approximation of the discrete update of Phase 1 and Phase 2
1046 of DP-SGD, respectively.
1047

1048 **Proof:** • Phase 1: Let $Z_{DP} \sim \mathcal{N} \left(0, \frac{C^2 \sigma_{DP}^2}{B^2} I_d \right)$ be the differentially-private noise injected via
1049 Gaussian Mechanism and denote with $\bar{\Delta} = x_1 - x$ the one-step increment for Phase 1. Applying
1050 Lemma A.2 with tolerance $\epsilon = \eta$ and by definition we have
1051

$$1053 \mathbb{E}[\bar{\Delta}] = -\eta \mathbb{E}_{\gamma, DP} \left[C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} \right] = -\eta \frac{CK(\nu)}{\sigma_\gamma \sqrt{d}} \nabla f(x) + \mathcal{O}(\eta^2). \quad (25)$$

1055 Then, the second moment becomes
1056

$$1058 \text{Cov}(\bar{\Delta}) = \mathbb{E} \bar{\Delta} \bar{\Delta}^\top - \mathbb{E} [\bar{\Delta}] \mathbb{E} [\bar{\Delta}^\top] \quad (26)$$

$$1060 = \eta^2 \mathbb{E}_{\gamma, DP} \left[\left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} - \frac{CK(\nu)}{\sigma_\gamma \sqrt{d}} \nabla f(x) + \mathcal{O}(\eta^2) \right) \right. \quad (27)$$

$$1063 \left. \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} - \frac{CK(\nu)}{\sigma_\gamma \sqrt{d}} \nabla f(x) + \mathcal{O}(\eta^2) \right)^\top \right] \quad (28)$$

$$1066 = \eta^2 \left(C^2 \mathbb{E}_{\gamma, DP} \left[\frac{\nabla f_\gamma(x) \nabla f_\gamma(x)^\top}{\|\nabla f_\gamma(x)\|_2^2} \right] + Z_{DP} Z_{DP}^\top \right. \quad (29)$$

$$1068 \left. - \frac{C^2 K(\nu)^2}{\sigma_\gamma^2 d} \nabla f(x) \nabla f(x)^\top \right) + \mathcal{O}(\eta^4) \quad (30)$$

$$1071 = \eta^2 \left(C^2 \mathbb{E} \left[\frac{\nabla f_\gamma(x) \nabla f_\gamma(x)^\top}{\|\nabla f_\gamma(x)\|_2^2} \right] - \frac{C^2 K(\nu)^2}{\sigma_\gamma \sqrt{d}} \nabla f(x) \nabla f(x)^\top + \frac{C^2 \sigma_{DP}^2}{B^2} I_d \right) + \mathcal{O}(\eta^4). \quad (31)$$

1074 Define now
1075

$$1076 b(x) := -\frac{CK(\nu)}{\sigma_\gamma \sqrt{d}} \nabla f(x) \quad (32)$$

$$1078 \bar{\Sigma}(x) := C^2 \mathbb{E} \left[\frac{\nabla f_\gamma(x) \nabla f_\gamma(x)^\top}{\|\nabla f_\gamma(x)\|_2^2} \right] - \frac{C^2 K(\nu)^2}{\sigma_\gamma \sqrt{d}} \nabla f(x) \nabla f(x)^\top + \frac{C^2 \sigma_{DP}^2}{B^2} I_d. \quad (33)$$

1080 Then, from Lem B.3 and Thm. B.4 the claim follows.
1081

1082 • Phase 2: Following the same steps as above, one obtains:
1083

$$\mathbb{E}[\bar{\Delta}] = -\eta \mathbb{E}_{\gamma, DP} [\nabla f_\gamma(x) + Z_{DP}] = -\eta \nabla f(x), \quad (34)$$

1084 and
1085

$$\text{Cov}(\bar{\Delta}) = \eta^2 \mathbb{E} [(\nabla f_\gamma(x) + Z_{DP} - \nabla f(x)) (\nabla f_\gamma(x) + Z_{DP} - \nabla f(x))^\top] \quad (35)$$

$$= \eta^2 \mathbb{E} [(\nabla f_\gamma(x) - \nabla f(x)) (\nabla f_\gamma(x) - \nabla f(x))^\top] + \eta^2 \frac{C^2 \sigma_{DP}^2}{B^2} I_d \quad (36)$$

$$= \eta^2 \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2} \right) I_d \quad (37)$$

1093 Define
1094

$$b(x) := -\nabla f(x). \quad (38)$$

$$\bar{\Sigma}(x) := \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2} \right) I_d. \quad (39)$$

1099 Finally, from Lem B.3 and Thm. B.4 the claim follows. \square
1100

1101 **Theorem B.6** Let f be L -smooth and μ -PL. Then, for $t \in [0, \tau]$, we have that
1102

1103 • Phase 1, i.e., when the gradient is clipped, the loss satisfies:
1104

$$\mathbb{E}[f(X_t)] \lesssim f(X_0) e^{-\frac{\mu C}{\sigma_\gamma \sqrt{d}} t} + \left(1 - e^{-\frac{\mu C}{\sigma_\gamma \sqrt{d}} t}\right) \frac{T \eta d^{\frac{3}{2}} L C \sigma_\gamma}{\mu} \left(\frac{\varepsilon^2}{dT} + \frac{\Phi^2}{B^2}\right) \frac{1}{\varepsilon^2}; \quad (40)$$

1107 • Phase 2, i.e., when the gradient is not clipped, the loss satisfies:
1108

$$\mathbb{E}[f(X_t)] \lesssim f(X_0) e^{-\mu t} + (1 - e^{-\mu t}) \frac{T \eta d L}{\mu} \left(\frac{\varepsilon^2 \sigma_\gamma^2}{B T} + C^2 \frac{\Phi^2}{B^2}\right) \frac{1}{\varepsilon^2}. \quad (41)$$

1111 **Proof:** • Phase 1: By construction we have
1112

$$\text{Tr}(\bar{\Sigma}(x)) \leq C^2 + d \frac{C^2 \sigma_{DP}^2}{B^2}. \quad (42)$$

1116 Since f is μ -PL and L -smooth it follows that $2\mu f(x) \leq \|\nabla f(x)\|_2^2$ and $\nabla^2 f(x) \preceq L I_d$. Hence, by
1117 applying the Itô formula we have
1118

$$df(X_t) = -\frac{CK(\nu)}{\sigma_\gamma \sqrt{d}} \|\nabla f(X_t)\|_2^2 dt + \frac{\eta}{2} \text{Tr}(\nabla^2 f(X_t) \bar{\Sigma}(X_t)) dt + \mathcal{O}(\text{Noise}) \quad (43)$$

$$\leq -2\mu \frac{CK(\nu)}{\sigma_\gamma \sqrt{d}} f(X_t) dt + \frac{\eta d L}{2} \left(\frac{C^2}{d} + \frac{C^2 \sigma_{DP}^2}{B^2}\right) dt + \mathcal{O}(\text{Noise}). \quad (44)$$

1123 Therefore,
1124

$$\mathbb{E}[f(X_t)] \leq f(X_0) e^{-2\mu \frac{CK(\nu)C}{\sigma_\gamma \sqrt{d}} t} + \left(1 - e^{-2\mu \frac{CK(\nu)C}{\sigma_\gamma \sqrt{d}} t}\right) \frac{\eta d^{\frac{3}{2}} L C \sigma_\gamma}{4\mu CK(\nu)} \left(\frac{C^2}{d} + \frac{C^2 \sigma_{DP}^2}{B^2}\right). \quad (45)$$

1128 Let us now remind that
1129

$$\sigma_{DP} = \frac{q \sqrt{T \log(1/\delta)}}{\varepsilon}, \quad (46)$$

1131 then
1132

$$\mathbb{E}[f(X_t)] \leq f(X_0) e^{-2\mu \frac{CK(\nu)C}{\sigma_\gamma \sqrt{d}} t} + \left(1 - e^{-2\mu \frac{CK(\nu)C}{\sigma_\gamma \sqrt{d}} t}\right) \frac{\eta d^{\frac{3}{2}} L C \sigma_\gamma}{4\mu CK(\nu)} \left(\frac{1}{d} + \frac{T q^2 \log(1/\delta)}{B^2 \varepsilon^2}\right). \quad (47)$$

1134 • Phase 2: Similarly to Phase 1, we have
 1135

$$\text{Tr}(\bar{\Sigma}(x)) = d \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2} \right). \quad (48)$$

1138 Again using the fact that f is μ -PL and L -smooth and by applying the Itô formula, one obtains
 1139

$$df(X_t) \leq -\|\nabla f(X_t)\|_2^2 dt + \frac{\eta dL}{2} \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2} \right) + \mathcal{O}(\text{Noise}) \quad (49)$$

1142 from which we have
 1143

$$\mathbb{E}[f(X_t)] \leq f(X_0) e^{-2\mu t} + (1 - e^{-2\mu t}) \frac{\eta dL}{4\mu} \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2} \right). \quad (50)$$

1146 Hence, by expanding σ_{DP}
 1147

$$\mathbb{E}[f(X_t)] \leq f(X_0) e^{-2\mu t} + (1 - e^{-2\mu t}) \frac{\eta dL}{4\mu} \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 q^2 T \log(1/\delta)}{B^2 \varepsilon^2} \right). \quad (51)$$

1150 Finally, let $\Phi = q\sqrt{\log(1/\delta)}$ and suppress all problem-independent constants, such as $2, \pi, K(\nu)$,
 1151 to obtain the claim.
 1152

□

1155 **Theorem B.7** Let f be L -smooth and define

$$K_1 := \max \left\{ 1, \frac{\sigma_\gamma \sqrt{d}}{CK(\nu)} \right\} \quad K_2 := \max \left\{ \frac{C^2}{d}, \frac{\sigma_\gamma^2}{B} \right\}. \quad (52)$$

1159 then

$$\mathbb{E}[\|\nabla f(X_{\tilde{t}})\|_2^2] \lesssim K_1 \left(\frac{f(X_0)}{\eta T} + \frac{\eta dL}{2} \left(K_2 + \frac{C^2 \left(\frac{q}{B} \right)^2 T \log(1/\delta)}{\varepsilon^2} \right) \right), \quad (53)$$

1162 where $\tilde{t} \sim \text{Unif}(0, \tau)$.
 1163

1164 **Proof:** Since f is L -smooth and by applying the Itô formula to Phase 1 we have:
 1165

$$df(X_t) \leq -\frac{CK(\nu)}{\sigma_\gamma \sqrt{d}} \|\nabla f(X_t)\|_2^2 dt + \frac{\eta}{2} \text{Tr}(L\bar{\Sigma}(X_t)) dt + \mathcal{O}(\text{Noise}) \quad (54)$$

$$\leq -\frac{CK(\nu)}{\sigma_\gamma \sqrt{d}} \|\nabla f(X_t)\|_2^2 dt + \frac{\eta dL}{2} \left(\frac{C^2}{d} + \frac{C^2 \sigma_{DP}^2}{B^2} \right) dt + \mathcal{O}(\text{Noise}). \quad (55)$$

1170 Similarly, in Phase 2 we obtain
 1171

$$df(X_t) \leq -\|\nabla f(X_t)\|_2^2 dt + \frac{\eta}{2} \text{Tr}(L\bar{\Sigma}(X_t)) dt + \mathcal{O}(\text{Noise}) \quad (56)$$

$$\leq -\|\nabla f(X_t)\|_2^2 dt + \frac{\eta dL}{2} \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2} \right) dt + \mathcal{O}(\text{Noise}). \quad (57)$$

1177 Let K_1 and K_2 as in Eq. 52. Then, by integrating and taking the expectation, we have
 1178

$$\mathbb{E} \int_0^\tau \|\nabla f(X_t)\|_2^2 dt \leq K_1 \left(f(X_0) - f(X_\tau) + \frac{\tau \eta dL}{2} \left(K_2 + \frac{C^2 \sigma_{DP}^2}{B^2} \right) \right) \quad (58)$$

$$\Rightarrow \mathbb{E} \int_0^\tau \frac{1}{\tau} \|\nabla f(X_t)\|_2^2 dt \leq K_1 \left(\frac{f(X_0) - f(X_\tau)}{\tau} + \frac{\eta dL}{2} \left(K_2 + \frac{C^2 \sigma_{DP}^2}{B^2} \right) \right) \quad (59)$$

$$\Rightarrow \mathbb{E} [\|\nabla f(X_{\tilde{t}})\|_2^2] \leq K_1 \left(\frac{2\varepsilon^2 (f(X_0) - f(X_\tau)) + \eta^2 dL T K_2}{2\eta T} + \frac{\eta dL T C^2 q^2 \log(1/\delta)}{B^2} \right) \frac{1}{\varepsilon^2}$$

1185 where the last step follows from the Law of the Unconscious Statistician and $\tilde{t} \sim \text{Unif}(0, \tau)$. Finally,
 1186 by suppressing problem-independent constants, $2, \pi$, we obtain the claim.
 1187

□

1188 B.1.1 MIXED-PHASE GRADIENT BOUND
1189

1190 In this section, we extend the two-phase SDE derivation to a single mixed setting. This is important
1191 because, at any point during training, some per-example gradients may exceed the clipping threshold
1192 while others remain below it. We show that in this scenario the same bound holds as in Theorem B.7,
1193 where it was previously derived under a worst-case approach.

1194 **Theorem B.8** *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be L -smooth. Then, we can write the SDE of DP-SGD as*

$$1196 dX_t = b_{\text{mix}}(X_t) dt + \sqrt{\eta} \Sigma_{\text{mix}}(X_t)^{1/2} dW_t, \quad (60)$$

1197 where the drift and covariance satisfy, for all x ,

$$1199 \langle \nabla f(x), b_{\text{mix}}(x) \rangle \leq -\frac{1}{K_1} \|\nabla f(x)\|^2, \quad (61)$$

$$1201 \text{Tr } \Sigma_{\text{mix}}(x) \leq d \left(K_2 + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} \right), \quad (62)$$

1204 where K_1, K_2 are defined in Equation 52. Therefore, for $\tilde{t} \sim \text{Unif}(0, \tau)$,

$$1205 1206 \mathbb{E} \|f(X_{\tilde{t}})\|_2^2 \leq K_1 \left(\frac{f(X_0)}{\eta T} + \frac{\eta d L}{2} \left(K_2 + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} \right) \right), \quad (63)$$

1208 i.e., the same L -smooth convergence bound as in Theorem B.7 holds for any mixture of clipped and
1209 unclipped samples in each mini-batch.

1211 **Proof:** We proceed in three steps: i) drift under mixed clipping, ii) covariance under mixed clip-
1212 ping using an explicit decomposition of G_k into clipped and unclipped parts, and iii) Itô's formula
1213 and the final bound.

1214 • Step 1: Drift of the mixed batch. The DP-SGD update can be written as

$$1216 1217 x_{k+1} = x_k - \eta \left(G_k + \frac{1}{B} Z_{\text{DP}} \right), \quad \text{where } G_k := \frac{1}{B} \sum_{i \in \gamma_k} \mathcal{C}(\nabla f_i(x_k)). \quad (64)$$

1218 Let

$$1219 1220 S_{1,k} := \{i \in \gamma_k : \|\nabla f_i(x_k)\| \geq C\}, \quad S_{2,k} := \{i \in \gamma_k : \|\nabla f_i(x_k)\| < C\}, \quad (65)$$

$$1221 1222 B_k := |S_{1,k}|, \quad p_k := \frac{B_k}{B} \in [0, 1]. \quad (66)$$

1223 Intuitively, p_k represents the probability of a sample being in Phase 1. Define the per-sample contribu-
1224 tions

$$1225 Y_i := \mathcal{C}(\nabla f_i(x_k)), \quad i \in S_{1,k}, \quad X_i := \nabla f_i(x_k), \quad i \in S_{2,k}, \quad (67)$$

1226 and the corresponding batch averages

$$1228 1229 g_k^{(1)} := \frac{1}{B} \sum_{i \in S_{1,k}} Y_i, \quad g_k^{(2)} := \frac{1}{B} \sum_{i \in S_{2,k}} X_i, \quad (68)$$

1230 so that

$$1231 1232 G_k = g_k^{(1)} + g_k^{(2)}. \quad (69)$$

1233 In the same way as in the proof of Theorem B.5, we have that

$$1234 1235 \mathbb{E}[Y_i] = a_1 \nabla f(x_k), \quad a_1 := \frac{C K(\nu)}{\sigma_{\gamma} \sqrt{d}}, \quad (70)$$

1236 and

$$1237 1238 \mathbb{E}[X_i] = \nabla f(x_k). \quad (71)$$

1239 Conditioned on the sets $S_{1,k}, S_{2,k}$, the Y_i are i.i.d. over $S_{1,k}$ and the X_i are i.i.d. over $S_{2,k}$, so

$$1240 1241 \mathbb{E}[g_k^{(1)} | S_{1,k}] = \frac{1}{B} \sum_{i \in S_{1,k}} \mathbb{E}[Y_i] = \frac{B_k}{B} \mathbb{E}[Y_i] = p_k a_1 \nabla f(x_k), \quad (72)$$

$$\mathbb{E}[g_k^{(2)} \mid S_{2,k}] = \frac{1}{B} \sum_{i \in S_{2,k}} \mathbb{E}[X_i] = \frac{B - B_k}{B} \mathbb{E}[X_i] = (1 - p_k) \nabla f(x_k). \quad (73)$$

Thus

$$\mathbb{E}[G_k \mid S_{1,k}, S_{2,k}] = p_k a_1 \nabla f(x_k) + (1 - p_k) \nabla f(x_k). \quad (74)$$

Recall the definition of K_1 from Equation 52

$$K_1 := \max \left\{ 1, \frac{\sigma_\gamma \sqrt{d}}{CK(\nu)} \right\}. \quad (75)$$

Then, it holds $a_1 = CK(\nu)/(\sigma_\gamma \sqrt{d}) \geq 1/K_1$ and $a_2 := 1 \geq 1/K_1$. Since $a_{\text{mix}}(p_k)$ is a convex combination of a_1 and a_2 ,

$$a_{\text{mix}}(p_k) = p_k a_1 + (1 - p_k) a_2 \geq \min\{a_1, a_2\} \geq \frac{1}{K_1} \quad \forall p_k \in [0, 1]. \quad (76)$$

Therefore, the drift in the SDE limit satisfies

$$b_{\text{mix}}(x) = -a_{\text{mix}}(p(x)) \nabla f(x), \quad \langle \nabla f(x), b_{\text{mix}}(x) \rangle \leq -\frac{1}{K_1} \|\nabla f(x)\|^2. \quad (77)$$

• Step 2: Covariance of the mixed batch. We now compute the gradient noise covariance and show it is a convex combination of the pure-phase covariances, which are already derived in Theorem B.5. Define the centered contributions

$$U_i := Y_i - \mathbb{E}[Y_i], \quad i \in S_{1,k}, \quad V_j := X_j - \mathbb{E}[X_j], \quad j \in S_{2,k}. \quad (78)$$

Then

$$g_k^{(1)} - \mathbb{E}[g_k^{(1)} \mid S_{1,k}] = \frac{1}{B} \sum_{i \in S_{1,k}} U_i, \quad g_k^{(2)} - \mathbb{E}[g_k^{(2)} \mid S_{2,k}] = \frac{1}{B} \sum_{j \in S_{2,k}} V_j. \quad (79)$$

Hence

$$G_k - \mathbb{E}[G_k \mid S_{1,k}, S_{2,k}] = \frac{1}{B} \sum_{i \in S_{1,k}} U_i + \frac{1}{B} \sum_{j \in S_{2,k}} V_j. \quad (80)$$

Let

$$\Sigma_1^{\text{single}}(x_k) := \text{Cov}(Y_i) = \text{Cov}(U_i), \quad \Sigma_2^{\text{single}}(x_k) := \text{Cov}(X_j) = \text{Cov}(V_j). \quad (81)$$

be the covariances of a single data-point. Conditioned on the sets $S_{1,k}, S_{2,k}$, the random vectors $\{U_i : i \in S_{1,k}\}$ and $\{V_j : j \in S_{2,k}\}$ are independent and zero-mean. Thus

$$\text{Cov}(G_k \mid S_{1,k}, S_{2,k}) = \text{Cov} \left(\frac{1}{B} \sum_{i \in S_{1,k}} U_i + \frac{1}{B} \sum_{j \in S_{2,k}} V_j \mid S_{1,k}, S_{2,k} \right) \quad (82)$$

$$= \frac{1}{B^2} \text{Cov} \left(\sum_{i \in S_{1,k}} U_i \right) + \frac{1}{B^2} \text{Cov} \left(\sum_{j \in S_{2,k}} V_j \right), \quad (83)$$

where cross terms vanish by independence. Using i.i.d. within each group, we have

$$\text{Cov} \left(\sum_{i \in S_{1,k}} U_i \right) = B_k \Sigma_1^{\text{single}}(x_k), \quad \text{Cov} \left(\sum_{j \in S_{2,k}} V_j \right) = (B - B_k) \Sigma_2^{\text{single}}(x_k), \quad (84)$$

therefore

$$\Sigma_{\text{grad}}(x_k; S_{1,k}, S_{2,k}) := \text{Cov}(G_k \mid x_k, S_{1,k}, S_{2,k}) = \frac{B_k}{B^2} \Sigma_1^{\text{single}}(x_k) + \frac{B - B_k}{B^2} \Sigma_2^{\text{single}}(x_k). \quad (85)$$

Since $p_k = B_k/B$ and $1 - p_k = (B - B_k)/B$, we obtain

$$\Sigma_{\text{grad}}(x_k; S_{1,k}, S_{2,k}) = \frac{p_k}{B} \Sigma_1^{\text{single}}(x_k) + \frac{1 - p_k}{B} \Sigma_2^{\text{single}}(x_k). \quad (86)$$

In the pure-phase SDEs of Theorem B.5, the *batch-level* (gradient and DP) covariances are given by

$$\Sigma_1(\bar{x}) = \frac{1}{B} \Sigma_1^{\text{single}}(x) + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} I_d, \quad \Sigma_2(\bar{x}) = \frac{1}{B} \Sigma_2^{\text{single}}(x) + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} I_d. \quad (87)$$

1296 From equation 86, the *gradient* part of the mixed-phase covariance is
 1297

$$1298 \Sigma_{\text{grad}}(x_k; S_{1,k}, S_{2,k}) = p_k \left(\Sigma_1(x_k) - \frac{C^2 \sigma_{\text{DP}}^2}{B^2} I_d \right) + (1 - p_k) \left(\Sigma_2(x_k) - \frac{C^2 \sigma_{\text{DP}}^2}{B^2} I_d \right). \quad (88)$$

1300 Adding the DP noise term $\frac{C^2 \sigma_{\text{DP}}^2}{B^2} I_d$ back in, the *full* covariance of the DP-SGD increment in the
 1301 mixed batch is
 1302

$$1303 \Sigma_{\text{mix}}(x_k; S_{1,k}, S_{2,k}) = \Sigma_{\text{grad}}(x_k; S_{1,k}, S_{2,k}) + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} I_d \quad (89)$$

$$1304 = p_k \Sigma_1(x_k) + (1 - p_k) \Sigma_2(x_k). \quad (90)$$

1306 Thus, at the SDE level, the mixed-phase covariance is exactly a convex combination of the pure-
 1307 phase covariances Σ_1 and Σ_2 . From Theorem B.6, we have the trace bounds
 1308

$$1309 \text{Tr } \Sigma_1(\bar{x}) \leq C^2 + d \frac{C^2 \sigma_{\text{DP}}^2}{B^2}, \quad (91)$$

$$1311 \text{Tr } \Sigma_2(\bar{x}) = d \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} \right). \quad (92)$$

1313 Let K_2 as in Equation 52 we can write, for $r \in \{1, 2\}$,

$$1315 \text{Tr } \Sigma_r(\bar{x}) \leq d \left(K_2 + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} \right). \quad (93)$$

1318 Using equation 90, for any $p_k \in [0, 1]$,

$$1319 \text{Tr } \Sigma_{\text{mix}}(x_k) = p_k \text{Tr } \Sigma_1(x_k) + (1 - p_k) \text{Tr } \Sigma_2(x_k) \leq d \left(K_2 + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} \right). \quad (94)$$

1322 Hence, the mixed-phase covariance satisfies exactly the same worst-case trace bound as the pure-
 1323 phase covariances.

1324 • Step 3: Itô bound and convergence. Finally, we can rewrite the SDE of DP-SGD as follows:

$$1325 dX_t = b_{\text{mix}}(X_t) dt + \sqrt{\eta} \Sigma_{\text{mix}}(X_t)^{1/2} dW_t, \quad (95)$$

1327 where, for all x ,

$$1328 \langle \nabla f(x), b_{\text{mix}}(x) \rangle \leq -\frac{1}{K_1} \|\nabla f(x)\|^2, \quad (96)$$

$$1330 \text{Tr } \Sigma_{\text{mix}}(x) \leq d \left(K_2 + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} \right). \quad (97)$$

1332 Since f is L -smooth, $\nabla^2 f(x) \preceq L I_d$. By Itô's formula,

$$1334 df(X_t) = \langle \nabla f(X_t), b_{\text{mix}}(X_t) \rangle dt + \frac{\eta}{2} \text{Tr}(\nabla^2 f(X_t) \Sigma_{\text{mix}}(X_t)) dt + \mathcal{O}(\text{Noise}). \quad (98)$$

1336 Using the drift and covariance bounds,

$$1337 df(X_t) \leq -\frac{1}{K_1} \|\nabla f(X_t)\|_2^2 dt + \frac{\eta dL}{2} \left(K_2 + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} \right) dt + \mathcal{O}(\text{Noise}). \quad (99)$$

1340 Integrating from 0 to $\tau := \eta T$,

$$1341 f(X_\tau) - f(X_0) \leq -\frac{1}{K_1} \int_0^\tau \|\nabla f(X_t)\|_2^2 dt + \frac{\eta dL}{2} \left(K_2 + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} \right) \tau + \mathcal{O}(\text{Noise}). \quad (100)$$

1343 Rearranging,

$$1345 \frac{1}{K_1} \int_0^\tau \|\nabla f(X_t)\|_2^2 dt \leq f(X_0) - f(X_\tau) + \frac{\eta dL}{2} \left(K_2 + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} \right) \tau + \mathcal{O}(\text{Noise}). \quad (101)$$

1347 Taking expectations,

$$1349 \frac{1}{K_1} \mathbb{E} \int_0^\tau \|\nabla f(X_t)\|_2^2 dt \leq \mathbb{E}[f(X_0) - f(X_\tau)] + \frac{\eta dL}{2} \left(K_2 + \frac{C^2 \sigma_{\text{DP}}^2}{B^2} \right) \tau. \quad (102)$$

1350 Let $\tilde{t} \sim \text{Unif}(0, \tau)$. Then, by the Law on Unconscious Statistician,

$$1352 \mathbb{E}\|\nabla f(X_{\tilde{t}})\|^2 = \frac{1}{\tau} \mathbb{E} \int_0^\tau \|\nabla f(X_t)\|_2^2 dt \leq K_1 \left(\frac{f(X_0)}{\tau} + \frac{\eta d L}{2} \left(K_2 + \frac{C^2 \sigma_{DP}^2}{B^2} \right) \right). \quad (103)$$

1354 Since $\tau = \eta T$, this is exactly the gradient-norm bound as in Theorem B.7, with the same constants
1355 K_1, K_2 , now rigorously shown to hold under arbitrary mixtures of clipped and unclipped samples
1356 at each iteration.

□

1359 We now derive the stationary distribution of DP-SGD at convergence: We empirically validate this
1360 result in Figure C.3.

1361 **Theorem B.9** *Let $f(x) = \frac{1}{2}x^\top Hx$ where $H = \text{diag}(\lambda_1, \dots, \lambda_d)$. The stationary distribution at
1362 convergence of DP-SGD is*

$$1364 \mathbb{E}[X_\tau], \text{Cov}(X_\tau)) = \left(X_0 e^{-H\tau}, \frac{T\eta}{2\varepsilon^2} \left(\frac{\varepsilon^2 \sigma_\gamma^2}{BT} + \frac{C^2 q^2 \log(1/\delta)}{B^2} \right) (1 - e^{-2H\tau}) H^{-1} \right). \quad (104)$$

1367 **Proof:** Since H is diagonal, we can isolate each component. Furthermore, since $f(\cdot)$ is quadratic
1368 we can rewrite the SDE as:

$$1370 dX_{t,i} = -\lambda_i X_{t,i} + \sqrt{\eta} \sqrt{\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2}} dW_{t,i}. \quad (105)$$

1372 We have immediately that

$$1374 \mathbb{E}[X_{t,i}] = X_{0,i} e^{-\lambda_i t}. \quad (106)$$

1375 Applying the Itô isometry, we obtain:

$$\begin{aligned} 1376 \mathbb{E}[(X_{t,i} - \mathbb{E}[X_{t,i}])^2] \\ 1377 &= \eta \mathbb{E} \left[\int_0^t \left(e^{-\lambda_i(t-s)} \sqrt{\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2}} dW_s \right)^\top \left(e^{-\lambda_i(t-s)} \sqrt{\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2}} dW_s \right) \right] \\ 1378 &= \eta \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 \sigma_{DP}^2}{B^2} \right) \int_0^t e^{-2\lambda_i(t-s)} ds \\ 1381 &= \frac{\eta}{2\lambda_i} \left(\frac{\sigma_\gamma^2}{B} + \frac{C^2 q^2 T \log(1/\delta)}{B^2 \varepsilon^2} \right) (1 - e^{-2\lambda_i t}) \\ 1384 &= \frac{T\eta}{2\varepsilon^2 \lambda_i} \left(\frac{\varepsilon^2 \sigma_\gamma^2}{BT} + \frac{C^2 q^2 \log(1/\delta)}{B^2} \right) (1 - e^{-2\lambda_i t}). \end{aligned}$$

□

1392 B.2 DP-SIGNSGD

1393 This subsection provides the formal derivation of the SDE model for DP-SignSGD and formal
1394 statements of Theorem 4.3, Theorem 4.4, and Theorem B.15. Similarly to DP-SGD, the dynamics
1395 of the method shifts again between two phases; we first model and study each phase separately.

1397 **Theorem B.10** *Let $0 < \eta < 1$, $\tau > 0$ and set $T = \lfloor \tau/\eta \rfloor$ and $K(\nu) = \sqrt{\frac{2}{\nu} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})}}$, $\nu \geq 1$. Let
1398 $x_k \in \mathbb{R}^d$ denote a sequence of DP-SignSGD iterations defined in Eq. 5. Assume Assumption B.1
1399 and Assumption B.2. Let X_t be the solution of the following SDEs with initial condition $X_0 = x_0$:*

1400 **• Phase 1:**

$$1403 dX_t = -\mathbb{E}_\gamma \left[\text{Erf} \left(\frac{B}{\sigma_{DP} \sqrt{2}} \frac{\nabla f_\gamma(X_t)}{\|\nabla f_\gamma(X_t)\|_2} \right) \right] dt + \sqrt{\eta} \sqrt{\bar{\Sigma}(X_t)} dW_t, \quad (107)$$

1404 where $\bar{\Sigma}(x) = I_d - \mathbb{E}_\gamma \left[\text{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \right]^2$.
 1405

1406 • Phase 2:

1407
$$dX_t = -\text{Erf} \left(\frac{\nabla f(X_t)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right) dt + \sqrt{\eta} \sqrt{\bar{\Sigma}(X_t)} dW_t, \quad (108)$$

 1408
 1409
 1410
 1411

1412 where $\bar{\Sigma}(x) = I_d - \text{Erf} \left(\frac{\nabla f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right)^2$ and $\text{Erf}(\cdot)$ is applied component-wise.
 1413
 1414
 1415

1416 Then, Eq. 107 and Eq. 108 are an order 1 approximation of the discrete update of Phase 1 and
 1417 Phase 2 of DP-SignSGD, respectively.

1418 **Proof:** The proof is virtually identical to that of Theorem B.5. Hence, we highlight only the
 1419 necessary details for each phase. Let $\bar{\Delta} = x_1 - x$ be the one-step increment.
 1420

1421 • Phase 1: We begin by computing the first moment:

1422
$$\mathbb{E}[\bar{\Delta}] = -\eta \mathbb{E}_{\gamma, DP} \left[\text{sign} \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} \right) \right]. \quad (109)$$

 1423
 1424

1425 Remember that, for any random variable Y , we have

1426
$$\mathbb{E}[\text{sign}(Y)] = 1 - 2\mathbb{P}(Y < 0), \quad (110)$$

 1427

1428 and that if furthermore $Y \sim \mathcal{N}(0, 1)$, then

1429
$$\Phi(y) = \frac{1}{2} \left(1 + \text{Erf} \left(\frac{y}{\sqrt{2}} \right) \right). \quad (111)$$

 1430
 1431

1432 Since $Z_{DP} \sim \mathcal{N} \left(0, \frac{C^2 \sigma_{DP}^2}{B^2} \right)$, we have that
 1433

1434
$$1 - 2\mathbb{P} \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} < 0 \right) = 1 - 2\Phi \left(-\frac{B}{C\sigma_{DP}} C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \quad (112)$$

 1435
 1436

1437
$$= 1 - \left(1 + \text{Erf} \left(-\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \right) \quad (113)$$

 1438
 1439

1440
$$= \text{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right). \quad (114)$$

 1441

1442 Thus

1443
$$\mathbb{E}[\bar{\Delta}] = -\eta \mathbb{E}_\gamma \left[\text{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \right]. \quad (115)$$

 1444
 1445

1446 The second moment is instead

1447
$$\text{Cov}(\bar{\Delta})_{ij} = \eta^2 \mathbb{E}_{\gamma, DP} \left[\left(\text{sign} \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} \right) - \mathbb{E}_\gamma \left[\text{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \right] \right)_i \right. \\ \left. \left(\text{sign} \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} \right) - \mathbb{E}_\gamma \left[\text{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \right] \right)_j \right] \quad (116)$$

 1448
 1449
 1450
 1451
 1452

1453
$$= \eta^2 \mathbb{E}_{\gamma, DP} \left[\text{sign} \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} \right)_i \text{sign} \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} \right)_j \right] \quad (117)$$

 1454
 1455

1456
$$- \eta^2 \mathbb{E}_\gamma \left[\text{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right)_i \right] \mathbb{E}_\gamma \left[\text{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right)_j \right]. \quad (118)$$

 1457

1458 If $i = j$, we have
 1459

1460
 1461
$$\bar{\Delta}_{ii} = \eta^2 - \eta^2 \mathbb{E}_\gamma \left[\operatorname{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \right]^2. \quad (119)$$

 1462
 1463

1464 Otherwise, we have
 1465

1466
 1467
$$\mathbb{E}_{\gamma, DP} \left[\operatorname{sign} \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} \right)_i \operatorname{sign} \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} \right)_j \right] \quad (120)$$

 1468

1469
 1470
$$= \mathbb{E}_\gamma \left[\mathbb{E}_{DP} \left[\operatorname{sign} \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} \right)_i \right] \mathbb{E}_{DP} \left[\operatorname{sign} \left(C \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} + Z_{DP} \right)_j \right] \right]$$

 1471

1472
 1473
$$= \mathbb{E}_\gamma \left[\operatorname{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right)_i \operatorname{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right)_j \right] \quad (121)$$

 1474

1475
 1476
$$= \mathbb{E}_\gamma \left[\operatorname{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right)_i \right] \mathbb{E}_\gamma \left[\operatorname{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right)_j \right]. \quad (122)$$

 1477

1478
 1479 Where we used the independence of the i -th and j -th components. Hence
 1480

1481
 1482
$$\operatorname{Cov}(\bar{\Delta})_{ij} = 0. \quad (123)$$

 1483

1484
 1485 Finally, we have
 1486

1487
 1488
$$\operatorname{Cov}(\bar{\Delta}) = \eta^2 I_d - \eta^2 \mathbb{E}_\gamma \left[\operatorname{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \right]^2. \quad (124)$$

 1489

1490
 1491 Define now
 1492

1493
 1494
$$b(x) = -\mathbb{E}_\gamma \left[\operatorname{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \right] \quad (125)$$

 1495

1496
 1497
$$\bar{\Sigma}(x) = I_d - \mathbb{E}_\gamma \left[\operatorname{Erf} \left(\frac{B}{\sigma_{DP}\sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \right]^2. \quad (126)$$

 1498

1499
 1500 Then, from Lem B.3 and Thm. B.4 the claim follows.
 1501

1502 • Phase 2: Remember that, from Assumption B.2, $\nabla f_\gamma = \nabla f + Z_\gamma$, where $Z_\gamma \sim \mathcal{N} \left(0, \frac{\sigma_\gamma^2}{B} \right)$. We
 1503 calculate the expected increment
 1504

1505
 1506
$$\mathbb{E}[\bar{\Delta}] = -\eta \mathbb{E}[\operatorname{sign}(\nabla f_\gamma(x) + Z_{DP})] \quad (127)$$

 1507

1508
 1509
$$= -\eta \mathbb{E}[\operatorname{sign}(\nabla f(x) + Z_\gamma + Z_{DP})] \quad (128)$$

 1510

1511
 1512
$$= -\eta \operatorname{Erf} \left(\frac{\nabla f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right). \quad (129)$$

 1513

1512 Instead, the covariance becomes
 1513

$$1514 \text{Cov}(\bar{\Delta})_{ij} = \eta^2 \mathbb{E}_{\gamma, DP} \left[\left(\text{sign}(\nabla f_\gamma + Z_{DP}) - \text{Erf} \left(\frac{\nabla f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right) \right)_i \right. \\ 1515 \left. \left(\text{sign}(\nabla f_\gamma + Z_{DP}) - \text{Erf} \left(\frac{\nabla f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right) \right)_j \right] \quad (130)$$

$$1516 \\ 1517 \\ 1518 \\ 1519 \\ 1520 \\ 1521 \\ 1522 \\ 1523 \\ 1524 \\ 1525 \\ 1526 \\ 1527 \\ 1528$$

$$= \eta^2 \mathbb{E}_{\gamma, DP} [\text{sign}(\nabla f_\gamma + Z_{DP})_i \text{sign}(\nabla f_\gamma + Z_{DP})_j] \quad (132)$$

$$- \eta^2 \text{Erf} \left(\frac{\partial_i f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right) \text{Erf} \left(\frac{\partial_j f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right). \quad (133)$$

1529 If $i = j$, we have
 1530

$$1531 \\ 1532 \\ 1533 \\ 1534 \\ 1535$$

$$\text{Cov}(\bar{\Delta})_{ii} = \eta^2 \left(1 - \text{Erf} \left(\frac{\partial_i f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right)^2 \right); \quad (134)$$

1536 while if $i \neq j$
 1537

$$1538 \text{Cov}(\bar{\Delta})_{ij} = \eta^2 \text{Erf} \left(\frac{\partial_i f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right) \text{Erf} \left(\frac{\partial_j f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right) \quad (135)$$

$$1542 \\ 1543 \\ 1544 \\ 1545 \\ 1546$$

$$- \eta^2 \text{Erf} \left(\frac{\partial_i f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right) \text{Erf} \left(\frac{\partial_j f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right) = 0, \quad (136)$$

1547 Define now

$$1548 \\ 1549 \\ 1550 \\ 1551 \\ 1552 \\ 1553 \\ 1554 \\ 1555 \\ 1556$$

$$b(x) = - \text{Erf} \left(\frac{\nabla f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right) \quad (137)$$

$$\bar{\Sigma}(x) = I_d - \text{Erf} \left(\frac{\nabla f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right)^2. \quad (138)$$

1557 Then, from Lem B.3 and Thm. B.4 the claim follows.
 1558

□

1560 **Corollary B.11** *Under our assumptions, the SDEs (Eq. 107 and Eq. 108) modelling the two phases
 1561 of DP-SignSGD as follows:*

1562 **• Phase 1:**

$$1563 \\ 1564 \\ 1565$$

$$dX_t = - \sqrt{\frac{2}{d\pi} \frac{BK(\nu)}{\sigma_{DP} \sigma_\gamma}} \nabla f(X_t) dt + \sqrt{\eta} \sqrt{I_d - \frac{2}{d\pi} \frac{B^2 K(\nu)^2}{\sigma_{DP}^2 \sigma_\gamma^2} \text{diag}(\nabla f(X_t))^2} dW_t; \quad (139)$$

1566 • Phase 2:

$$1568 dX_t = -\sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B}}} \nabla f(X_t) + \eta \sqrt{I_d - \frac{2}{\pi} \frac{1}{\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B}}} \text{diag}(\nabla f(X_t))^2 dW_t. \quad (140)$$

1571 **Proof:** Let us w.l.o.g. assume that $\left| \frac{\partial_i f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right| \ll 1$ when $\|\nabla f_\gamma(x)\|_2 \geq C$: This is not
1572 restrictive when the number of trainable parameters d is large as it is under our assumptions. Additionally,
1573 we recall that under our assumptions, $\frac{|\partial_i f(x)|}{\sqrt{2(\sigma_{DP}^2 + \sigma_\gamma^2/B)}} \ll 1$. Then one can write:
1574

1576 • Phase 1: Since $\left| \frac{\partial_i f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right| \ll 1$, one can approximate the error function in a neighborhood of 0
1577 as follows: $\text{Erf}(x) \sim \frac{2}{\sqrt{\pi}} x$. Thanks to Lemma A.1, we have
1578

$$1579 \mathbb{E}_\gamma \left[\text{Erf} \left(\frac{B}{\sigma_{DP} \sqrt{2}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right) \right] = \mathbb{E}_\gamma \left[\sqrt{\frac{2}{\pi}} \frac{B}{\sigma_{DP}} \frac{\nabla f_\gamma(x)}{\|\nabla f_\gamma(x)\|_2} \right] = \sqrt{\frac{2}{d\pi}} \frac{BK(\nu)}{\sigma_{DP} \sigma_\gamma} \nabla f(x) \quad (141)$$

1582 Therefore, Eq. 107 becomes

$$1584 dX_t = -\sqrt{\frac{2}{d\pi}} \frac{BK(\nu)}{\sigma_{DP} \sigma_\gamma} \nabla f(X_t) dt + \sqrt{\eta} \sqrt{I_d - \frac{2}{d\pi} \frac{B^2 K(\nu)^2}{\sigma_{DP}^2 \sigma_\gamma^2}} \text{diag}(\nabla f(X_t))^2 dW_t. \quad (142)$$

1587 • Phase 2: Since $\left| \frac{\partial_i f(x)}{\sqrt{2(\sigma_{DP}^2 + \sigma_\gamma^2/B)}} \right| \ll 1$ for $i = 1, \dots, d$, one can use the same argument as before
1588 to use a linear approximation of the error function. In detail, one has
1589

$$1590 \text{Erf} \left(\frac{\nabla f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \right) = \frac{2}{\sqrt{\pi}} \frac{\nabla f(x)}{\sqrt{2 \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} = \sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B}}} \nabla f(x). \quad (143)$$

1595 Therefore, Eq. 108 becomes

$$1596 dX_t = -\sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B}}} \nabla f(X_t) + \eta \sqrt{I_d - \frac{2}{\pi} \frac{1}{\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B}}} \text{diag}(\nabla f(X_t))^2 dW_t. \quad (144)$$

1599 \square

1600 **Theorem B.12** Let f be L -smooth and μ -PL. Then, for $t \in [0, \tau]$, we have that

1602 • Phase 1, i.e., when the gradient is clipped, the loss satisfies:

$$1604 \mathbb{E}[f(X_t)] \lesssim f(X_0) e^{\frac{-\mu B}{\sigma_\gamma \sqrt{dT}} \frac{\varepsilon}{\Phi} t} + \left(1 - e^{\frac{-\mu B}{\sigma_\gamma \sqrt{dT}} \frac{\varepsilon}{\Phi} t} \right) \frac{\sqrt{T} \eta L d^{\frac{3}{2}} \sigma_\gamma \Phi}{\mu B} \frac{\varepsilon}{\varepsilon}; \quad (145)$$

1606 • Phase 2, i.e., when the gradient is not clipped, the loss satisfies:

$$1608 \mathbb{E}[f(X_t)] \lesssim f(X_0) e^{\sqrt{\varepsilon^2 \frac{\sigma_\gamma^2}{B} + \frac{C^2 \Phi^2}{B^2} T}} + \left(1 - e^{\sqrt{\varepsilon^2 \frac{\sigma_\gamma^2}{B} + \frac{C^2 \Phi^2}{B^2} T}} \right) \frac{\sqrt{T} \eta L d}{\mu} \sqrt{\frac{\varepsilon^2 \sigma_\gamma^2}{B T} + \frac{C^2 \Phi^2}{B^2} \frac{1}{\varepsilon}}. \quad (146)$$

1612 **Proof:** First of all, observe that, in both phases, it holds that $\bar{\Sigma}(x) \preceq I_d$.

1613 • Phase 1: Since f is μ -PL and L -smooth it follows that $2\mu f(x) \leq \|\nabla f(x)\|^2$ and $\nabla^2 f(x) \preceq L I_d$.
1614 Then, By applying the Itô formula, we have

$$1616 df(X_t) \leq -\sqrt{\frac{2}{d\pi T}} \frac{K(\nu)}{\sigma_\gamma} \frac{B\varepsilon}{q \log(1/\delta)} \|\nabla f(X_t)\|_2^2 dt + \frac{\eta}{2} \text{Tr}(\nabla^2 f(X_t) I_d) dt + \mathcal{O}(\text{Noise}) \quad (147)$$

$$1618 \leq -2\mu \sqrt{\frac{2}{d\pi T}} \frac{K(\nu)}{\sigma_\gamma} \frac{B\varepsilon}{q \log(1/\delta)} f(X_t) dt + \frac{\eta d L}{2} dt + \mathcal{O}(\text{Noise}). \quad (148)$$

1620 Therefore,

$$1622 \mathbb{E}[f(X_t)] \leq f(X_0) e^{-2\mu \left(\sqrt{\frac{2}{d\pi T}} \frac{K(\nu)}{\sigma_\gamma} \frac{B\varepsilon}{q \log(1/\delta)} \right) t} \quad (149)$$

$$1624 + \left(1 - e^{-2\mu \left(\sqrt{\frac{2}{d\pi T}} \frac{K(\nu)}{\sigma_\gamma} \frac{B\varepsilon}{q \log(1/\delta)} \right) t} \right) \sqrt{\frac{\pi T}{2}} \frac{\eta d^{\frac{3}{2}} L \sigma_\gamma}{4\mu K(\nu)} \frac{q \log(1/\delta)}{B\varepsilon}. \quad (150)$$

1626 • Phase 2: As for Phase 1, by applying the Itô formula one has

$$1628 df(X_t) \leq -\sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\varepsilon^2 B^{-1} \sigma_\gamma^2 + B^{-2} C^2 q^2 \log(1/\delta) T}} \varepsilon \|\nabla f(X_t)\|_2^2 dt \quad (151)$$

$$1631 + \frac{\eta}{2} \text{Tr}(\nabla^2 f(X_t) I_d) dt + \mathcal{O}(\text{Noise}) \quad (152)$$

$$1633 \leq -2\mu \sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\varepsilon^2 B^{-1} \sigma_\gamma^2 + B^{-2} C^2 q^2 \log(1/\delta) T}} \varepsilon f(X_t) dt + \frac{\eta d L}{2} dt + \mathcal{O}(\text{Noise}). \quad (153)$$

1636 Therefore

$$1638 \mathbb{E}[f(X_t)] \leq f(X_0) e^{-\sqrt{\frac{2}{\pi}} \frac{2\mu}{\sqrt{\varepsilon^2 B^{-1} \sigma_\gamma^2 + B^{-2} C^2 q^2 \log(1/\delta) T}} \varepsilon t} \quad (154)$$

$$1640 + \left(1 - e^{-\sqrt{\frac{2}{\pi}} \frac{2\mu}{\sqrt{\varepsilon^2 B^{-1} \sigma_\gamma^2 + B^{-2} C^2 q^2 \log(1/\delta) T}} \varepsilon t} \right) \sqrt{\frac{\pi T}{2}} \frac{\eta d L}{4\mu} \sqrt{\frac{\varepsilon^2 \sigma_\gamma^2}{B T} + \frac{C^2 q^2 \log(1/\delta)}{B^2}} \frac{1}{\varepsilon}.$$

1643 Finally, by suppressing problem-independent constants, such as $2, \pi, K(\nu)$, the thesis follows.

1644 \square

1646 **Theorem B.13** Let f be an L -smooth function. Define

$$1648 K_3 = \max \left\{ \sqrt{\frac{d\pi}{2}} \frac{\sigma_\gamma q \sqrt{\log(1/\delta)}}{B K(\nu)}, \sqrt{\frac{\pi}{2}} \sqrt{\frac{\varepsilon^2 \sigma_\gamma^2}{B T} + \frac{C^2 q^2 \log(1/\delta)}{B^2}} \right\}. \quad (155)$$

1651 Then

$$1652 \mathbb{E} [\|\nabla f(X_{\tilde{t}})\|_2^2] \lesssim K_3 \left(\frac{f(X_0)}{\eta \sqrt{T}} + \eta d L \sqrt{T} \right) \frac{1}{\varepsilon}, \quad (156)$$

1654 where $\tilde{t} \sim \text{Unif}(0, \tau)$.

1655

1656 **Proof:** Since in both phases the diffusion coefficient $\bar{\Sigma}(x) \preceq I_d$, the drift is the only term worth
1657 comparing for a worst-case analysis. Let then K_3 as in Eq. 155. Applying the Itô formula to the
1658 worst-case SDE we have

$$1659 df(X_t) \leq -\varepsilon (\sqrt{T} K_3)^{-1} \|\nabla f(X_t)\|_2^2 dt + \frac{\eta}{2} \text{Tr}(\nabla^2 f(X_t) I_d) dt + \mathcal{O}(\text{Noise}) \quad (157)$$

$$1661 \leq -\varepsilon (\sqrt{T} K_3)^{-1} \|\nabla f(X_t)\|_2^2 dt + \frac{\eta d L}{2} dt + \mathcal{O}(\text{Noise}). \quad (158)$$

1663 Then, by integrating and taking the expectation

$$1665 \mathbb{E} \int_0^\tau \|\nabla f(X_t)\|_2^2 dt \leq K_3 \sqrt{T} \left(f(X_0) + \frac{\eta d L \tau}{2} \right) \varepsilon^{-1} \quad (159)$$

$$1667 \implies \mathbb{E} \int_0^\tau \frac{1}{\tau} \|\nabla f(X_t)\|_2^2 dt \leq \frac{K_3}{\eta \sqrt{T}} \left(f(X_0) + \frac{\eta d L \tau}{2} \right) \varepsilon^{-1} \quad (160)$$

$$1670 \implies \mathbb{E} [\|\nabla f(X_{\tilde{t}})\|_2^2] \leq K_3 \left(\frac{f(X_0)}{\eta \sqrt{T}} + \frac{\eta d L \sqrt{T}}{2} \right) \frac{1}{\varepsilon} \quad (161)$$

1672 where in the last step we used the Law of the Unconscious Statistician and $\tilde{t} \sim \text{Unif}(0, \tau)$. Finally,
1673 by suppressing problem-independent constants, we get the thesis. \square

1674 B.2.1 MIXED-PHASE GRADIENT BOUND
1675

1676 Analogously to Section B.1.1, we extend the two-phase SDE derivation to a single mixed setting.
1677 Recall that, at any point during training, some per-example gradients may lie above the clipping
1678 threshold while others remain below it. The next result shows that, even in this more realistic mixed
1679 regime, we recover the same upper bound on the gradient norm as in Theorem B.13.

1680
1681
1682 **Theorem B.14** *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be L -smooth. Then, we can write the SDE of DP-SignSGD as*

$$1683 \quad dX_t = b_{\text{mix}}(X_t) dt + \sqrt{\eta} \Sigma_{\text{mix}}(X_t)^{1/2} dW_t, \quad (162)$$

1684 *where*

$$1685 \quad b_{\text{mix}}(x) = -\mathbb{E} \left[\text{Erf} \left(\frac{B}{C\sigma_{DP}\sqrt{2}} G(x) \right) \right], \quad (163)$$

$$1686 \quad \bar{\Sigma}_{\text{mix}}(x) = I_d - \mathbb{E} \left[\text{Erf} \left(\frac{B}{C\sigma_{DP}\sqrt{2}} G(x) \right) \right]^2, \quad (164)$$

1687 *and*

$$1688 \quad G(x) = \frac{1}{B} \sum_{i \in \gamma_k} \mathcal{C}(\nabla f_i(x)) \quad (165)$$

1689 *Define*

$$1690 \quad K_4 = \max \left\{ \sqrt{\frac{\pi d}{2}} \frac{\sigma_\gamma q \sqrt{\log(1/\delta)}}{BK(\nu)}, \sqrt{\frac{\pi}{2}} \frac{Cq \sqrt{\log(1/\delta)}}{B} \right\}. \quad (166)$$

1691 *Then*

$$1692 \quad \mathbb{E} \left[\|\nabla f(X_{\tilde{t}})\|_2^2 \right] \leq K_4 \left(\frac{f(X_0)}{\eta\sqrt{T}} + \frac{\eta d L \sqrt{T}}{2} \right) \frac{1}{\varepsilon}, \quad (167)$$

1693 *where $\tilde{t} \sim \text{Unif}(0, \tau)$.*

1694
1695 **Remark B.3** *By construction we have $K_4 \leq K_3$, so Theorem B.14 provides a formally tighter
1696 upper bound than Theorem B.13. However, note that the first term in the definitions of K_3 and
1697 K_4 (Equations 155 and 166, respectively) scales as \sqrt{d} . Since d is assumed to be large, this term
1698 typically dominates the maximum in both constants. As a consequence, in the high-dimensional
1699 regime of interest we effectively have $K_3 = K_4$, and the improvement from the mixed-phase analysis
1700 is negligible in practice.*

1701 **Proof:** We divide the proof into two steps: i) SDE derivation, ii) gradient bound.

1702 • **Step 1: SDE derivation.** Using the same notation as in the proof of Theorem B.8, we write the
1703 update of DP-SignSGD as

$$1704 \quad x_{k+1} = x_k - \eta \left(G_k + \frac{1}{B} Z_{DP} \right), \quad (168)$$

1705 *where*

$$1706 \quad G_k := p_k \frac{1}{B_k} \sum_{i=1}^{B_k} \mathcal{C}(\nabla f_i(x_k)) + (1-p_k) \frac{1}{B-B_k} \sum_{i=1}^{B-B_k} \nabla f_i(x_k). \quad (169)$$

1707 Since $Z_{DP} \sim \mathcal{N}$ we have

$$1708 \quad \mathbb{E}[x_{k+1} - x_k] = -\eta \mathbb{E} \left[\text{sign} \left(G_k + \frac{1}{B} Z_{DP} \right) \right] \quad (170)$$

$$1709 \quad = -\eta \mathbb{E} \left[\text{Erf} \left(\frac{B}{C\sigma_{DP}\sqrt{2}} G_k \right) \right] \quad (171)$$

1728

1729 and

1730

1731
$$\text{Cov}(x_{k+1} - x_k) = \eta^2 \mathbb{E} \left[\left(\text{sign} \left(G_k + \frac{1}{B} Z_{DP} \right) - \mathbb{E} \left[\text{Erf} \left(\frac{B}{C\sigma_{DP}\sqrt{2}} G_k \right) \right] \right) \right] \quad (172)$$

1732

1733

1734
$$\left(\text{sign} \left(G_k + \frac{1}{B} Z_{DP} \right) - \mathbb{E} \left[\text{Erf} \left(\frac{B}{C\sigma_{DP}\sqrt{2}} G_k \right) \right] \right)^\top \quad (173)$$

1735

1736

1737
$$= \eta^2 \left(I_d - \mathbb{E} \left[\text{Erf} \left(\frac{B}{C\sigma_{DP}\sqrt{2}} G_k \right) \right]^2 \right). \quad (174)$$

1738

If we define

1739

1740
$$b_{\text{mix}}(x) = -\mathbb{E} \left[\text{Erf} \left(\frac{B}{C\sigma_{DP}\sqrt{2}} G(x) \right) \right], \quad (175)$$

1741

1742

1743
$$\bar{\Sigma}_{\text{mix}}(x) = I_d - \mathbb{E} \left[\text{Erf} \left(\frac{B}{C\sigma_{DP}\sqrt{2}} G(x) \right) \right]^2, \quad (176)$$

1744

1745

Then, the following SDE is an order-1 approximation of the update step of DP-SignSGD

1746

1747
$$dX_t = b_{\text{mix}}(X_t) + \sqrt{\eta} \sqrt{\bar{\Sigma}_{\text{mix}}(X_t)} dW_t. \quad (177)$$

1748

1749

• Step 2: Gradient bound. As argued in the proof of Corollary B.11, we assume $\left| \frac{B}{C\sigma_{DP}\sqrt{2}} G(x) \right| \ll 1$ without loss of generality. Therefore, applying the Itô formula, we have

1750

1751

1752
$$df_t \leq -\sqrt{\frac{2}{\pi}} \frac{B}{C\sigma_{DP}} \nabla f(X_t)^\top \mathbb{E}[G_k] dt + \frac{\eta}{2} \text{Tr}(\nabla^2 f(X_t) \bar{\Sigma}(X_t)) dt + \mathcal{O}(\text{Noise}) \quad (178)$$

1753

1754

1755
$$\leq -(p_k a_1 + (1-p_k) a_2) \nabla f(X_t) dt + \frac{\eta dL}{2} dt + \mathcal{O}(\text{Noise}), \quad (179)$$

1756

1757

where $a_1 = \sqrt{\frac{2}{\pi}} \frac{B}{\sigma_{DP}} \frac{K(\nu)}{\sigma_\gamma \sqrt{d}}$ and $a_2 = \sqrt{\frac{2}{\pi}} \frac{B}{C\sigma_{DP}}$. Expand σ_{DP} and define

1758

1759

1760

1761

1762
$$K_4 = \max \left\{ \sqrt{\frac{\pi d}{2}} \frac{\sigma_\gamma q \sqrt{\log(1/\delta)}}{BK(\nu)}, \sqrt{\frac{\pi}{2}} \frac{Cq \sqrt{\log(1/\delta)}}{B} \right\}. \quad (180)$$

1763

Then

1764

1765
$$\varepsilon \sqrt{T^{-1}} K_4^{-1} \leq p_k a_1 + (1-p_k) a_2, \quad \forall p_k \in [0, 1] \quad (181)$$

1766

Then

1767

1768
$$df_t \leq -\varepsilon \sqrt{T^{-1}} K_4^{-1} \nabla f(X_t) dt + \frac{\eta dL}{2} dt + \mathcal{O}(\text{Noise}). \quad (182)$$

1769

Then, by taking the expectation and integrating over $[0, \tau]$

1770

1771

1772

1773

1774

1775

1776

1777

1778
$$\mathbb{E} \int_0^\tau \|\nabla f(X_t)\|_2^2 dt \leq \varepsilon^{-1} \sqrt{T} K_4 \left(f(X_0) + \frac{\eta dL\tau}{2} \right) \quad (183)$$

1779

1780

1781

1782
$$\implies \mathbb{E} \int_0^\tau \frac{1}{\tau} \|\nabla f(X_t)\|_2^2 dt \leq \frac{K_4}{\eta \sqrt{T}} \left(f(X_0) + \frac{\eta dL\tau}{2} \right) \varepsilon^{-1} \quad (184)$$

1783

1784

1785

1786
$$\implies \mathbb{E} \left[\|\nabla f(X_t)\|_2^2 \right] \leq K_4 \left(\frac{f(X_0)}{\eta \sqrt{T}} + \frac{\eta dL\sqrt{T}}{2} \right) \frac{1}{\varepsilon}, \quad (185)$$

1787

where in the last step we used the Law of the Unconscious Statistician and $\tilde{t} \sim \text{Unif}(0, \tau)$.

□

Finally, we derive the stationary distribution of DP-SignSGD: We empirically validate it in Fig. C.3.

33

1782 **Theorem B.15** Let $f(x) = \frac{1}{2}x^\top Hx$ where $H = \text{diag}(\lambda_1, \dots, \lambda_d)$. The stationary distribution of
 1783 Phase 2 is

$$1785 \mathbb{E}[X_T] = X_0 e^{-KH\tau}; \quad (186)$$

$$1786 \text{Cov}(X_T) = X_0^2 e^{-2KH\tau} \left(e^{-\eta K^2 H\tau} - 1 \right) \quad (187)$$

$$1788 + \eta (2KH + \eta H^2 K^2)^{-1} \left(1 - e^{-(2KH + \eta K^2 H^2)\tau} \right) \quad (188)$$

1790 where $K = \sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\varepsilon^2 B^{-1} \sigma_\gamma^2 + B^{-2} C^2 q^2 \log(1/\delta) T}} \varepsilon$.

1793 **Proof:** Since H is diagonal, we can work component-wise. Let us remember the SDE:

$$1795 dX_{t,i} = -\sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B}}} \lambda_i X_{t,i} + \sqrt{\eta} \sqrt{1 - \frac{2}{\pi \left(\frac{C^2 \sigma_{DP}^2}{B^2} + \frac{\sigma_\gamma^2}{B} \right)}} \lambda_i^2 X_{t,i}^2 dW_t. \quad (189)$$

1798 To ease the notation, we write $K = \sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\varepsilon^2 B^{-1} \sigma_\gamma^2 + B^{-2} C^2 q^2 \log(1/\delta) T}} \varepsilon$. Hence, we can write $X_{t,i}$
 1799 in closed form as

$$1802 X_{t,i} = x_{0,i} e^{-K\lambda_i t} + \sqrt{\eta} \int_0^t e^{-K\lambda_i(t-s)} \sqrt{1 - K^2 \lambda_i^2 X_{t,i}^2} dW_s. \quad (190)$$

1804 Due to the properties of the stochastic integral, we immediately have

$$1806 \mathbb{E}[X_{t,i}] = X_{0,i} e^{-\sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{\varepsilon^2 B^{-1} \sigma_\gamma^2 + B^{-2} C^2 q^2 \log(1/\delta) T}} \varepsilon \lambda_i t}. \quad (191)$$

1809 Using the Itô formula on $g(x) = x^2$, we have

$$1811 d(X_{t,i}^2) = -2K\lambda_i X_{t,i}^2 dt + \frac{\eta}{2} 2dt - \frac{\eta}{2} 2K^2 \lambda_i^2 X_{t,i}^2 dt + \mathcal{O}(\text{Noise}) \quad (192)$$

$$1813 \Rightarrow \mathbb{E}[X_{t,i}^2] = X_{0,i}^2 e^{-(2K\lambda_i + \eta K^2 \lambda_i^2)t} + \frac{\eta}{2K\lambda_i + \eta \lambda_i^2 K^2} \left(1 - e^{-(2K\lambda_i + \eta K^2 \lambda_i^2)t} \right), \quad (193)$$

1815 therefore

$$1817 \text{Cov}(X_{t,i}) = \mathbb{E}[X_{t,i}^2] - \mathbb{E}[X_{t,i}]^2 \quad (194)$$

$$1819 = X_{0,i}^2 e^{-(2K\lambda_i + \eta K^2 \lambda_i^2)t} + \frac{\eta}{2K\lambda_i + \eta \lambda_i^2 K^2} \left(1 - e^{-(2K\lambda_i + \eta K^2 \lambda_i^2)t} \right) - X_{0,i}^2 e^{-2K\lambda_i t}$$

$$1821 = X_{0,i}^2 e^{-2K\lambda_i t} \left(e^{-\eta K^2 \lambda_i^2 t} - 1 \right) + \frac{\eta}{2K\lambda_i + \eta \lambda_i^2 K^2} \left(1 - e^{-(2K\lambda_i + \eta K^2 \lambda_i^2)t} \right). \quad (195)$$

1823 \square

1825 Finally, we present a result that allows us to determine which of DP-SignSGD and DP-SignSGD
 1826 is more advantageous depending on the training setting.

1828 **Corollary B.16** If $\frac{\sigma_\gamma^2}{B} \geq 1$, then DP-SignSGD always achieves a better privacy-utility trade-off
 1829 than DP-SGD, though its convergence is slower. If $\frac{\sigma_\gamma^2}{B} < 1$, there exists a critical privacy level

$$1832 \varepsilon^* = \sqrt{\frac{C^2 TB}{n^2 (B - \sigma_\gamma^2)} \log\left(\frac{1}{\delta}\right)}, \quad (196)$$

1834 such that DP-SignSGD outperforms DP-SGD in utility whenever $\varepsilon < \varepsilon^*$, but still converges more
 1835 slowly than DP-SGD.

1836 **Proof:** The Phase 2 asymptotic terms at $t = T$ are
 1837

$$1838 A_{\text{SGD}} = \frac{T\eta dL}{\mu} \left(\frac{\varepsilon^2 \sigma_\gamma^2}{TB} + C^2 \frac{\Phi^2}{B^2} \right) \frac{1}{\varepsilon^2}, \quad A_{\text{Sign}} = \frac{\sqrt{T}\eta dL}{\mu} \sqrt{\frac{\varepsilon^2 \sigma_\gamma^2}{TB} + C^2 \frac{\Phi^2}{B^2}} \frac{1}{\varepsilon}. \quad (197)$$

1840 We compare $A_{\text{Sign}} < A_{\text{SGD}}$. Cancelling the common factor $\frac{\eta dL}{\mu}$ gives
 1841

$$1842 \frac{\sqrt{T}}{\varepsilon} \sqrt{\frac{\varepsilon^2 \sigma_\gamma^2}{TB} + C^2 \frac{\Phi^2}{B^2}} < \frac{T}{\varepsilon^2} \left(\frac{\varepsilon^2 \sigma_\gamma^2}{TB} + C^2 \frac{\Phi^2}{B^2} \right). \quad (198)$$

1845 Multiplying by ε^2 and dividing by the positive square root yields
 1846

$$1847 \varepsilon \sqrt{T} < T \sqrt{\frac{\varepsilon^2 \sigma_\gamma^2}{TB} + C^2 \frac{\Phi^2}{B^2}}. \quad (199)$$

1848 All quantities are non-negative, so squaring preserves the inequality:
 1849

$$1850 \varepsilon^2 T < T^2 \left(\frac{\varepsilon^2 \sigma_\gamma^2}{TB} + C^2 \frac{\Phi^2}{B^2} \right) \iff \left(1 - \frac{\sigma_\gamma^2}{B} \right) \varepsilon^2 < C^2 \frac{\Phi^2}{B^2} T. \quad (200)$$

1852 Using $\frac{\Phi}{B} = \frac{1}{n} \sqrt{\log(1/\delta)}$ gives
 1853

$$1854 \left(1 - \frac{\sigma_\gamma^2}{B} \right) \varepsilon^2 < \frac{C^2}{n^2} T \log\left(\frac{1}{\delta}\right). \quad (201)$$

1856 If $\frac{\sigma_\gamma^2}{B} \geq 1$, the left coefficient is non-positive and the inequality holds for all $\varepsilon > 0$. If $\frac{\sigma_\gamma^2}{B} < 1$,
 1857 solving for ε yields
 1858

$$1859 \varepsilon < \sqrt{\frac{C^2 TB}{n^2 (B - \sigma_\gamma^2)} \log\left(\frac{1}{\delta}\right)} = \varepsilon^*, \quad (202)$$

1861 which proves the claim. \square
 1862

1863 Interestingly, by keeping η and C depend on the optimizer, we get
 1864

$$1865 \sqrt{T} \eta_{\text{sign}} \sqrt{\frac{\sigma_\gamma^2}{BT} + \frac{C_{\text{sign}}^2 \Phi^2}{B^2 \varepsilon^2}} < T \eta_{\text{sgd}} \left(\frac{\sigma_\gamma^2}{BT} + \frac{C_{\text{sgd}}^2 \Phi^2}{B^2 \varepsilon^2} \right). \quad (203)$$

1868 We observe that if $\sigma_\gamma \rightarrow \infty$, DP-SignSGD is always better than DP-SGD, while if $\sigma_\gamma \rightarrow 0$, there
 1869 is always a threshold ε^* . Since the algebraic expressions are complex, we believe this is enough to
 1870 show that our insight is much more general than the case derived here and presented in the main
 1871 paper.

1873 C EXPERIMENTAL DETAILS AND ADDITIONAL RESULTS

1875 Our empirical analysis is based on the official GitHub repository <https://github.com/kenziyuliu/DP2> released with the Google paper (Li et al., 2023). In particular we consider the two following
 1876 classification problems:
 1877

1879 **IMDB** (Maas et al., 2011) is a sentiment analysis dataset for movie reviews, posed as a binary
 1880 classification task. It contains 25,000 training samples and 25,000 test samples, with each review
 1881 represented using a vocabulary of 10,000 words. We train a logistic regression model with 10,001
 1882 parameters.

1883 **StackOverflow** (Kaggle, 2022), (TensorFlow Federated, 2022) is a large-scale text dataset derived
 1884 from Stack Overflow questions and answers. Following the setup in (TensorFlow Federated, 2022),
 1885 we consider the task of predicting the tag(s) associated with a given sentence, but we restrict our
 1886 experiments to the standard centralized training setting rather than the federated one. We randomly
 1887 select 246,092 sentences for training and 61,719 for testing, each represented with 10,000 features.
 1888 The task is cast as a 500-class classification problem, yielding a model with approximately 5 million
 1889 parameters.

1890 **Optimizers.** We train both classification problems using DP-SGD, DP-SignSGD and DP-Adam.
 1891 For $k \geq 0$, learning rate η , variance σ_{DP}^2 , and batches γ_k of size B modeled as i.i.d. uniform random
 1892 variables taking values in $\{1, \dots, n\}$. Let g_k be the private gradient, defined as
 1893

$$1894 \quad g_k := \frac{1}{B} \sum_{i \in \gamma_k} \mathcal{C}(\nabla f_i(x_k)) + \frac{1}{B} \mathcal{N}(0, C^2 \sigma_{\text{DP}}^2 I_d) \quad (204)$$

1895 and $\mathcal{C}[\cdot]$ be the clipping function
 1896

$$1897 \quad \mathcal{C}(x) = \begin{cases} C \frac{x}{\|x\|_2} & \text{if } \|x\|_2 \geq C \\ x & \text{if } \|x\|_2 < C \end{cases} \quad (205)$$

1900 The iterates of DP-SGD are defined as
 1901

$$1902 \quad x_{k+1} = x_k - \eta g_k, \quad (206)$$

1903 while those of DP-SignSGD are defined as
 1904

$$1905 \quad x_{k+1} = x_k - \eta \text{sign}[g_k], \quad (207)$$

1906 where $\text{sign}[\cdot]$ is applied component-wise. The update rule of DP-Adam is defined as follows:
 1907

$$1908 \quad m_{k+1} = \beta_1 m_k + (1 - \beta_1) g_k, \quad \hat{m}_{k+1} = \frac{m_{k+1}}{1 - \beta_1^{k+1}},$$

$$1909 \quad v_{k+1} = \beta_2 v_k + (1 - \beta_2) g_k^2, \quad \hat{v}_{k+1} = \frac{v_{k+1}}{1 - \beta_2^{k+1}}, \quad (208)$$

$$1910 \quad x_{k+1} = x_k - \eta \frac{\hat{m}_{k+1}}{\sqrt{\hat{v}_{k+1}} + \epsilon},$$

1911 where g_k is the privatized stochastic gradient and is defined in Equation 204.
 1912

1913 **Hyper-parameters.** Unless stated otherwise, we fix the following hyperparameters in our experiments:
 1914 for IMDB and StackOverflow respectively, we train for 100, 50 epochs with batch size
 1915 $B = 64$. The choice of batch size follows the setting in (Li et al., 2023). We also aimed to avoid
 1916 introducing unnecessary variability, keeping the focus on the direction suggested by our theoretical
 1917 results. Finally, we set $\delta = 10^{-5}, 10^{-6}$, corresponding to the rule $\delta = 10^{-k}$, where k is the smallest
 1918 integer such that $10^{-k} \leq 1/n$ for the training dataset size n .
 1919

1920 **Protocol A.** we perform a grid search on *learning rate* $\eta = \{0.001, 0.01, 0.1, 1, 3, 5, 10\}$ and
 1921 *clipping threshold* $C = \{0.1, 0.25, 0.5, 1, 5\}$ for DP-SGD, DP-SignSGD and DP-Adam on both
 1922 datasets, using $\sigma_{\text{DP}} = 1$: this gives $\epsilon = 2.712$ and $\epsilon = 0.424$ for IMDB and StackOverflow
 1923 respectively. We summarize the best set of hyperparameters for each method on both datasets in
 1924 Table C.1.
 1925

1926 Dataset	1927 DP-SGD	1928 DP-SignSGD	1929 DP-Adam
1930 IMDB	(5, 0.5)	(0.1, 0.5)	(0.1, 0.5)
1931 StackOverflow	(3, 0.25)	(0.01, 0.5)	(0.01, 0.5)

1932 Table C.1: Tuned hyperparameters for different methods across the two datasets. The values refer to
 1933 (learning rate, clipping parameter); For DP-Adam we also used $\beta_1 = 0.9, \beta_2 = 0.999$ and adaptivity
 1934 $\epsilon = 10^{-8}$ in both cases.
 1935

1936 **Protocol B.** For each noise multiplier, we tune a new pair of learning rate and clipping pa-
 1937 rameter by performing a grid search. **IMDB:** For DP-SignSGD and DP-Adam, we consider
 1938 the following learning rates $\eta = \{0.01, 0.05, 0.10, 0.15, 0.22, 0.27, 0.33, 0.38, 0.44, 0.50\}$ and
 1939 clipping thresholds $C = \{0.05, 0.1, 0.25, 0.5\}$, while for DP-SGD we consider a different
 1940 range of learning rates $\eta = \{0.5, 0.7, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0, 5.5, 6.0\}$ and
 1941 $C = \{0.1, 0.25, 0.5\}$. This tuning is designed to identify the best hyperparameters across a
 1942 broad range of privacy budgets $\epsilon = \{0.01, 0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.4, 1.6, 1.8, 2.0\}$, which
 1943

1944 correspond to the following noise multipliers: $\{271.23, 13.56, 6.78, 4.52, 3.39, 2.71, 2.26, 1.94, 1.70, 1.51, 1.36\}$. **StackOverflow**: For DP-Adam we consider the following learning rates $\{0.001, 0.003, 0.005, 0.01, 0.03, 0.05, 0.1, 0.5\}$, for DP-SignSGD we add $\{0.008, 0.015, 0.02, 0.03, 0.04\}$ to the list, while for DP-SGD we consider a different range of learning rates $\eta = \{0.1, 0.5, 1.0, 2.0, 2.5, 3.0, 3.5, 4.0, 5.0\}$. For the clipping thresholds we consider $C = \{0.05, 0.1, 0.25, 0.35, 0.5, 1.0\}$ for every method. This tuning is designed to identify the best hyperparameters across a broad range of privacy budgets $\varepsilon = \{0.01, 0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.4\}$, which correspond to the following noise multipliers: $\{42.384, 2.119, 1.060, 0.706, 0.530, 0.424, 0.353, 0.303\}$.

1953

1954 C.1 DP-SGD AND DP-SIGNSGD: SDE VALIDATION (FIGURE C.1).

1955

1956 In this section, we describe how we validated the SDE models derived in Theorem B.5 and Theorem B.10 (Figure C.1). In line with works in the literature Compagnoni et al. (2025c;a), we optimize a quadratic and a quartic function. We run both DP-SGD and DP-SignSGD, calculating the full gradient and injecting noise as described in Assumption B.2. Similarly, we integrate our SDEs using the Euler-Maruyama algorithm (See, e.g., (Compagnoni et al., 2025c), Algorithm 1) with $\Delta t = \eta$. Results are averaged over 200 repetitions. For each of the two functions, the details are presented in the following paragraphs.

1962

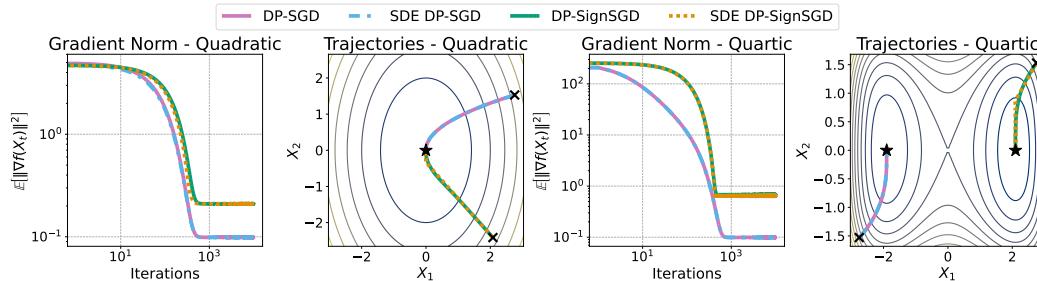
1963 **Quadratic function:** We consider the quadratic function $f(x) = \frac{1}{2}x^\top Hx$, with $H = 0.1 \text{ diag}(2, 1, \dots, 1)$, in dimension $d = 1024$. The clipping parameter is set to $C = 5$, and each 1964 algorithm is run for $T = 10000$ iterations. The gradient noise scale is $\sigma_\gamma = 1/\sqrt{d}$. The learning 1965 rate is $\eta = 0.1$ for DP-SGD and $\eta = 0.01$ for DP-SignSGD. The differential privacy parameters 1966 are $(\varepsilon, \delta, q) = (5, 10^{-4}, 10^{-4})$, corresponding to a noise multiplier of $\sigma_{DP} = 0.03$. The initial point 1967 is sampled as $x_0 = \frac{50}{\sqrt{d}} \mathcal{N}(0, I_d)$, using an independent seed for each method.

1968

1969 **Quartic function:** We also test on the quartic function $f(x) = \frac{1}{2} \sum_{i=0}^{d-1} H_{ii} x_i^2 + \frac{\lambda}{4} \sum_{i=0}^{d-1} x_i^4 - \xi \sum_{i=0}^{d-1} x_i^3$, where $H = \text{diag}(-2, 1, \dots, 1)$, $\lambda = 0.5$, and $\xi = 0.1$. The problem dimension, 1970 clipping, and number of iterations are the same: $d = 1024$, $C = 5$, $T = 10000$, with gradient noise 1971 $\sigma_\gamma = 1/\sqrt{d}$. Both methods use a learning rate of $\eta = 0.01$. The differential privacy parameters 1972 are $(\varepsilon, \delta, q) = (5, 10^{-4}, 10^{-4})$ for DP-SGD and $(5, 10^{-4}, 2 \times 10^{-4})$ for DP-SignSGD, corresponding 1973 to noise multipliers $\sigma_{DP} = 0.03$ and $\sigma_{DP} = 0.06$, respectively. Initialization is $x_0 = \frac{50}{\sqrt{d}} \mathcal{N}(0, I_d)$ 1974 for DP-SGD and $y_0 = -x_0$ for DP-SignSGD.

1975

1976



1977

1978

1979

1980

1981

1982

1983

1984

1985

1986

1987

1988

1989

1990

1991

1992

1993

1994

1995

1996

1997

Figure C.1: Consistent with Theorem B.5 and Theorem B.10, we empirically validate that the SDEs of DP-SGD and DP-SignSGD model their respective optimizers. For a convex quadratic function (**left two panels**) and a nonconvex quartic function (**right two panels**), the SDEs accurately track both the trajectories and the gradient norm of the corresponding algorithms, averaged over 200 runs.

C.2 ASYMPTOTIC LOSS BOUND (FIGURES 1 AND C.4)

This section refers to Figure 1 and Figure C.4. We consider three different scenarios: A quadratic function, IMDB, and StackOverflow. Each setup is optimized using DP-SGD, DP-SignSGD, and DP-Adam, and we plot the final averaged training loss across a range of privacy levels. In the left panel, we include the exact bounds from Theorem 4.1 and Theorem 4.3 to show agreement with

1998 theory; in the central and right panels, we compare the final losses with the trends in ε predicted by
 1999 the same theorems. Experimental details are as follows.

2000 **Quadratic:** $f(x) = \frac{1}{2}x^\top Hx$, $H = 10I_d$; $d = 1024$, $C = 5$, $T = 50000$, $\sigma_\gamma = 0.01$; learning rate
 2001 $\eta = 0.01 \cdot \eta_t$ with $\eta_t = (1 + \eta t)^{-0.6}$; Adam parameters: $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$. We
 2002 used 8 noise multipliers, linearly spaced from 0 to 2, which with $q = 10^{-4}$, $\delta = 10^{-4}$ correspond to
 2003 $\varepsilon \in \{\infty, 6.78, 2.38, 1.19, 0.79, 0.59, 0.48, 0.40, 0.34\}$.

2004 **IMDB:** Hyperparameters are given in Table C.1. We performed 10 runs for each
 2005 noise multiplier $\{0.5, 1.0, 2.0, 4.0, 6.0, 8.0, 10.0, 12.0\}$, yielding the following values for ε
 2006 $\{5.425, 2.712, 1.356, 0.678, 0.452, 0.339, 0.271, 0.226\}$, respectively. We report the average
 2007 training and test loss of the final epoch with confidence bounds (Figure 1 and Figure C.4).

2008 **StackOverflow:** Hyperparameters are given in Table C.1. We performed 3 runs using
 2009 for each multiplier $\{0.1, 0.3, 0.5, 1.0, 2.0, 4.0, 6.0, 8.0\}$, yielding the following values for ε
 2010 $\{4.238, 1.413, 0.848, 0.424, 0.212, 0.106, 0.071, 0.053\}$, respectively. We report the average
 2011 training and test loss of the final epoch with confidence bounds (Figure 1 and Figure C.4).

2012 C.3 CONVERGENCE SPEED ANALYSIS (FIGURE 2)

2013 This section refers to Figure 2. We consider two different scenarios: IMDB and StackOverflow.
 2014 Each setup is optimized using DP-SGD, DP-SignSGD, and DP-Adam and six different privacy
 2015 levels: We plot the average trajectories of the training losses and observe that, when it converges,
 2016 the convergence speed of DP-SGD does not depend on the level of privacy, while the two adaptive
 2017 method are more resilient to the demands of high levels of privacy, but their convergence speed
 2018 changes for every ε , as predicted in Theorem 4.3.

2019 **IMDB:** Hyperparameters are given in Table C.1. We performed 10 runs for
 2020 each noise multiplier $\{0.8, 1.0, 1.2, 1.6, 4.0, 6.0\}$ and corresponding epsilons
 2021 $\{3.390, 2.712, 2.260, 1.695, 0.678, 0.452\}$. We report the average trajectories of the training
 2022 loss with confidence bounds (Figure 2).

2023 **StackOverflow:** Hyperparameters are given in Table C.1. We performed 3 runs
 2024 for each noise multiplier $\{0.37, 0.5, 0.64, 1.19, 1.46, 1.73\}$ and corresponding epsilons
 2025 $\{1.146, 0.848, 0.662, 0.356, 0.290, 0.245\}$. We report the average trajectories of the training loss
 2026 with confidence bounds (Figure 2).

2027 C.4 WHEN ADAPTIVITY REALLY MATTERS (FIGURE 3)

2028 This section refers to Figure 3 and Figure C.2. Each setup is optimized using DP-SGD,
 2029 DP-SignSGD, and DP-Adam. We consider different batch sizes and for each we plot the final
 2030 loss values for different privacy levels, similarly to Section C.2. We highlight the possible range of
 2031 ε^* and a dash-dotted line to mark its approximate value, suggested by each graph. As predicted by
 2032 Theorem 4.5, the empirical value of ε^* shifts left as we increase the batch size. Experimental details
 2033 are as follows.

2034 **IMDB:** Hyperparameters are given in Table C.1. We select a wide range of noise multipliers:
 2035 $\{0.5, 1.0, 1.2, 1.5, 1.8, 2.0, 2.2, 2.5, 2.8, 3.0, 3.2, 3.5, 3.8, 4.0, 4.5, 5.0, 6.0, 8.0, 10.0, 12.0\}$
 2036 and increasing batch sizes $B = \{48, 56, 64, 72, 80\}$. The corresponding epsilons are

2037 $B = 48: \{4.698, 2.349, 1.879, 1.566, 1.342, 1.174, 1.044, 0.940, 0.854, 0.783, 0.723, 0.671,$
 2038 $0.626, 0.587, 0.522, 0.470, 0.391, 0.294, 0.235, 0.196\}$;

2039 $B = 56: \{5.070, 2.535, 2.028, 1.690, 1.449, 1.268, 1.127, 1.014, 0.922, 0.845, 0.780, 0.724,$
 2040 $0.676, 0.634, 0.563, 0.507, 0.423, 0.317, 0.254, 0.211\}$;

2041 $B = 64: \{5.425, 2.712, 2.170, 1.808, 1.550, 1.356, 1.205, 1.085, 0.986, 0.904, 0.835, 0.775,$
 2042 $0.723, 0.678, 0.603, 0.542, 0.452, 0.339, 0.271, 0.226\}$;

2043 $B = 72: \{5.740, 2.870, 2.296, 1.913, 1.640, 1.435, 1.276, 1.148, 1.044, 0.957, 0.883, 0.820,$
 2044 $0.765, 0.717, 0.638, 0.574, 0.478, 0.359, 0.287, 0.239\}$;

2052 $B = 80$: $\{6.070, 3.035, 2.428, 2.023, 1.734, 1.517, 1.349, 1.214, 1.104, 1.012, 0.934, 0.867,$
 2053 $0.809, 0.759, 0.674, 0.607, 0.506, 0.379, 0.303, 0.253\}$.

2054 For each batch size, we performed 10 runs and plotted the average final value of the Train Loss and
 2055 the empirical ε^* : these observed values follow the direction indicated in Thm. 4.5. For visualization
 2056 purposes, we show only a smaller window of ε values satisfying $0.75 \leq \varepsilon \leq 1.25$.

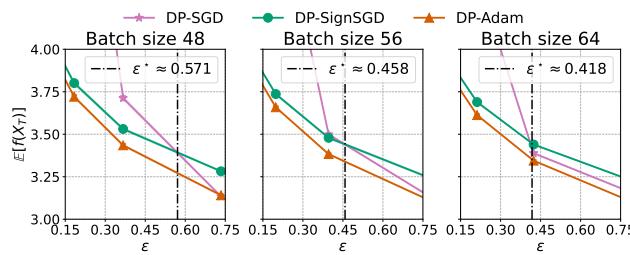
2057 **StackOverflow:** Due to the higher computational cost required, with our limited resources we
 2058 managed to select only a restricted range of noise multipliers: $\{0.1, 0.3, 0.5, 1.0, 2.0, 4.0, 6.0, 8.0\}$
 2059 and batch sizes: $\{48, 56, 64\}$. The corresponding epsilons are

2061 $B = 48$: $\{1.223, 0.734, 0.367, 0.184, 0.092, 0.061, 0.046\}$;

2062 $B = 56$: $\{1.322, 0.793, 0.396, 0.198, 0.099, 0.066, 0.050\}$;

2063 $B = 64$: $\{1.413, 0.848, 0.424, 0.212, 0.106, 0.071, 0.053\}$.

2064 For each batch size, we performed 3 runs and plotted the average final value of the Train Loss and
 2065 the empirical ε^* : these observed values follow the direction indicated in Thm. 4.5. For visualization
 2066 purposes, we show only a smaller window of ε values satisfying $0.08 \leq \varepsilon \leq 1.1$.



2069
 2070 Figure C.2: StackOverflow: From left to right, we decrease the batch noise, i.e., increase the batch
 2071 size, taking values $B = 48, 56, 64$: As per Theorem 4.5, the privacy threshold ε^* that determines
 2072 when DP-SignSGD is more advantageous than DP-SGD shifts to the left. This confirms that if
 2073 there is more noise due to the batch size, less privacy noise is needed for DP-SignSGD to be
 2074 preferable over DP-SGD.

2075

2076 C.5 BEST-TUNED HYPERPARAMETERS (FIGURES 4)

2077 This section refers to Figure 4. On top of the hyperparameter sweep performed described in Sec-
 2078 tion C, we additionally tune DP-SGD for the smaller values of ε . As predicted by Theorem 4.6, the
 2079 optimal learning rate for DP-SGD scales with ε , while those of the adaptive methods are almost con-
 2080 stant. Furthermore, we observe that once we reach the limits of the hyperparameter grid, DP-SGD
 2081 loses performance drastically.

2082 **IMDB:** We additionally tune DP-SGD using $\eta = \{0.001, 0.005, 0.01, 0.05, 0.1, 0.5\}$ and $C =$
 2083 $\{0.1, 0.25, 0.5\}$ and add the corresponding values using the cyan line. On the left, we plot the
 2084 average of the final 5 train loss values and confidence bound for each method against the privacy
 2085 budget ε ; On the right, we focus on the scaling of the optimal learning rate with respect to ε .

2086 **StackOverflow:** We additionally tune DP-SGD using $\eta = \{0.001, 0.01, 0.05\}$ and add the corre-
 2087 sponding values using the cyan line. As above, on the left, we plot the average of the final 5 training
 2088 loss values and confidence bounds for each method against the privacy budget ε ; on the right, we
 2089 focus on the scaling of the optimal learning rate with respect to ε .

2090

2091 C.6 STATIONARY DISTRIBUTIONS

2092 In this paragraph, we describe how we validated the convergence behavior predicted in Theorem B.9
 2093 and Theorem B.15. To produce Figure C.3, we run both DP-SGD and DP-SignSGD on $f(x) =$
 2094 $\frac{1}{2}x^\top Hx$, where $H = \text{diag}(2, 1)$, $x_0 = (0.01, 0.005)$, $\eta = 0.001$, $\sigma_\gamma = \sigma_{DP} = 0.1$, $C = 5$.
 2095 We average over 20000 runs and plot the evolution of the moments compared to the theoretical
 2096 prediction provided in Theorem B.9 and Theorem B.15.

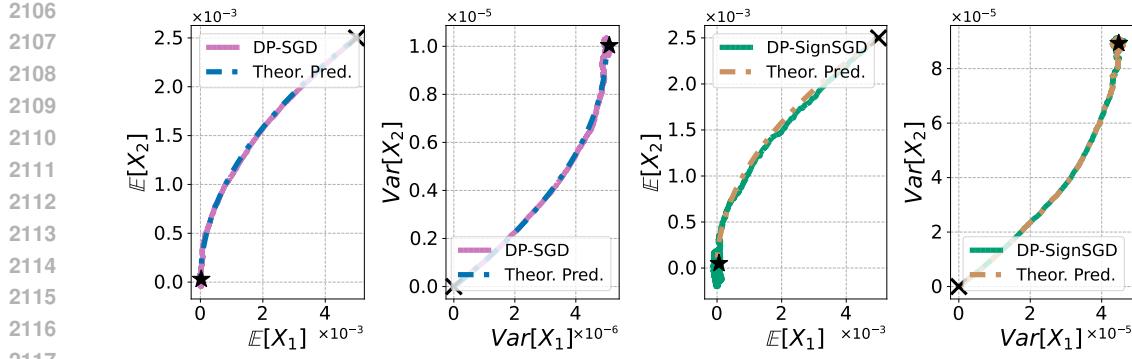


Figure C.3: The empirical dynamics of the first and second moments of the iterates X_t of DP-SGD (left two panels) and of DP-SignSGD (right two panels) match that prescribed in Theorem B.9 and Theorem B.15, respectively.

C.7 ADDITIONAL RESULTS — TEST LOSS

Interestingly, the insights provided in Theorem 4.1 and Theorem 4.3 regarding both the asymptotic bound and the convergence speed extend, in practice, also to the test loss. In the same set-up of Section C.2, we plot the asymptotic values of the Test Loss and interpolate with $\mathcal{O}(1/\varepsilon)$ and $\mathcal{O}(1/\varepsilon^2)$ to show that they match the predicted scaling.

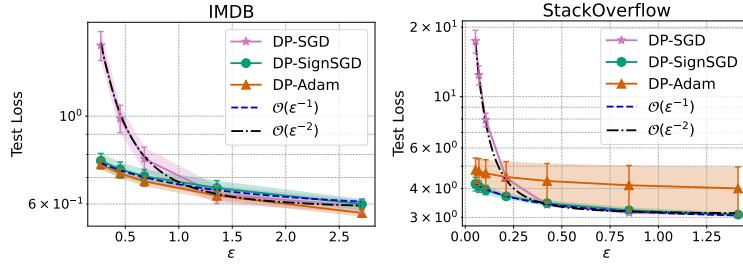


Figure C.4: Privacy-utility trade-off on the *test loss*, comparing DP-SGD, DP-SignSGD, and DP-Adam. **Left:** Logistic regression on the IMDB dataset. **Right:** Logistic regression on the StackOverflow dataset. In both cases, the empirical scalings predicted by Thm. 4.1 and Thm. 4.3 carry over from training to test: DP-SGD follows the $\frac{1}{\varepsilon^2}$ trend, while adaptive methods follow the $\frac{1}{\varepsilon}$ trend. This demonstrates that not only do our theoretical insights generalize to the widely used DP-Adam, but also extend from *training* to *test loss*.

Similarly, in the same set-up as Section C.3, we plot the trajectories of the Test Loss (Fig. C.5): we observe that once again the convergence speed of DP-SGD is not affected by the choice of ε , while adaptive methods clearly present different ε -dependent rates.

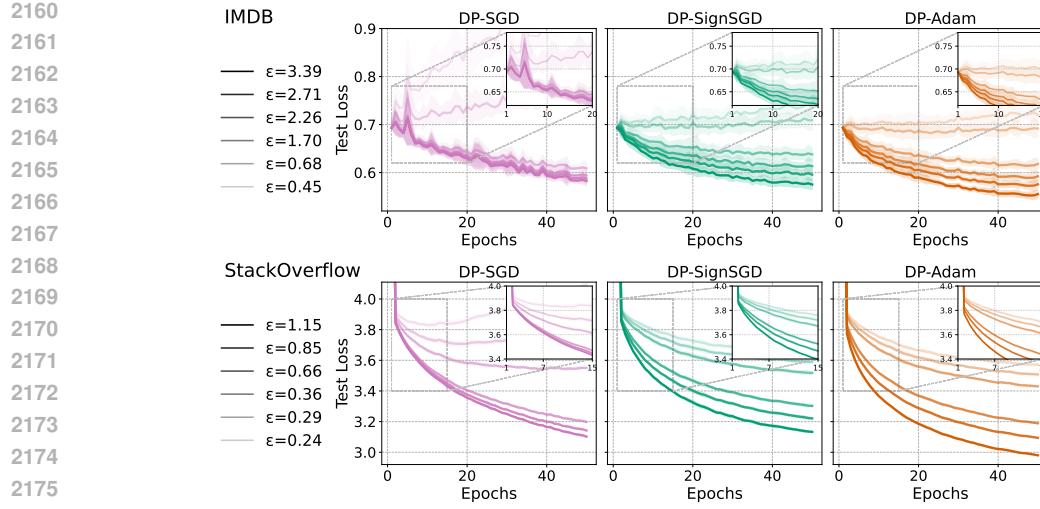


Figure C.5: We compare the Test Loss of DP-SGD, DP-SignSGD, and DP-Adam as we train a logistic regression on the IMDB dataset (**Top Row**) and on the StackOverflow dataset (**Bottom Row**).

D LIMITATIONS

As highlighted by Li et al. (2021b), the approximation capability of SDEs can break down when the learning rate η is large or when certain regularity assumptions on ∇f and the noise covariance matrix are not fulfilled. Although such limitations can, in principle, be alleviated by employing higher-order weak approximations, our position is that the essential function of SDEs is to provide a simplified yet faithful description of the discrete dynamics that offers practical insight. We do not anticipate that raising the approximation order beyond what is required to capture curvature-dependent effects would deliver substantial additional benefits.

We stress that our SDE formulations have been thoroughly validated empirically: the derived SDEs closely track their corresponding optimizers across a wide range of architectures, including MLPs, CNNs, ResNets, and ViTs (Paquette et al., 2021; Malladi et al., 2022; Compagnoni et al., 2024; 2025c;a; Xiao et al., 2025; Marshall et al., 2025).

Acknowledgments. We acknowledge the use of OpenAI’s ChatGPT as a writing assistant to help us rephrase and refine parts of the manuscript. All technical content, derivations, and scientific contributions remain the sole responsibility of the authors.