Fully-Adaptive Composition in Differential Privacy

Justin Whitehouse [1] **Aaditya Ramdas** [1] **Ryan Rogers** [2] **Zhiwei Steven Wu** [1]

Abstract

Composition is a key feature of differential privacy. Well-known advanced composition theorems allow one to query a private database quadratically more times than basic privacy composition would permit. However, these results require that the privacy parameters of all algorithms be fixed *before* interacting with the data. To address this, Rogers et al. (2016) introduced *fully adaptive* composition, wherein both algorithms and their privacy parameters can be selected adaptively. They defined two probabilistic objects to measure privacy in adaptive composition: *privacy filters*, which provide differential privacy guarantees for composed interactions, and *privacy odometers*, time-uniform bounds on privacy loss. There are substantial gaps between advanced composition and existing filters and odometers. First, existing filters place stronger assumptions on the algorithms being composed. Second, these odometers and filters suffer from large constants, making them impractical. We construct filters that match the rates of advanced composition, including constants, despite allowing for adaptively chosen privacy parameters. En route we also derive a privacy filter for approximate zCDP. We also construct several general families of odometers. These odometers match the tightness of advanced composition at an arbitrary, preselected point in time, or at all points in time simultaneously, up to a doubly-logarithmic factor. We obtain our results by leveraging advances in martingale concentration. In sum, we show that fully adaptive privacy is obtainable at almost no loss.

[1]Carnegie Mellon University [2]LinkedIn. Correspondence to: Justin Whitehouse <jwhiteho@andrew.cmu.edu>.

Proceedings of the 40th International Conference on Machine Learning, Honolulu, Hawaii, USA. PMLR 202, 2023. Copyright 2023 by the author(s).

1. Introduction

Differential privacy (Dwork et al., 2006b) is an algorithmic criterion that provides meaningful guarantees of individual privacy for conducting analysis on sensitive data. Intuitively, an algorithm is differentially private if similar inputs induce similar distributions on outputs. More formally, an algorithm $A : \mathcal{X} \to \mathcal{Y}$ is differentially private if, for any set of outcomes $G \subset \mathcal{Y}$ and any *neighboring* inputs $x, x' \in \mathcal{X}$,

$$\mathbb{P}(A(x) \in G) \leq e^{\epsilon}\mathbb{P}(A(x') \in G) + \delta, \tag{1}$$

where ϵ and δ are the privacy parameters of the algorithm.

A key property of differential privacy is graceful composition. Suppose $A_1, \ldots, A_n$ are algorithms such that each A_m is (ϵ_m, δ_m)-differentially private. Advanced composition (Dwork et al., 2010; Kairouz et al., 2015) states that, for any $\delta' > 0$, the *composed* sequence of algorithms is (ϵ, δ)-differentially private, where $\delta = \delta' + \sum_{m \leq n} \delta_m$, and

$$\epsilon = \sqrt{2\log\left(\frac{1}{\delta'}\right)\sum_{m\leq n}\epsilon_m^2} + \sum_{m \leq n}\epsilon_m\left(\frac{e^{\epsilon_m}-1}{e^{\epsilon_m}+1}\right). \tag{2}$$

When all privacy parameters are the same and small, we roughly have $\epsilon = O(\sqrt{n}\epsilon_m)$. Hence, analysts can make use of sensitive datasets with a slow degradation of privacy.

However, there is a major disconnect between most existing results on privacy composition and modern data analysis. As analysts view the outputs of algorithms, the future manner in which they interact with the data changes. Advanced composition allows analysts to adaptively select algorithms, but not privacy parameters. In many cases, analysts may wish to choose the subsequent privacy parameters based on the outcomes of the previous private algorithms. For example, if an analyst learns, from past computations, that they only need to run one more computation, they should be able to use the remainder of their privacy budget in the final round. Likewise, if an analyst is having a hard time deriving conclusions, they should be allowed to adjust privacy parameters to extend the allowable number of computations.

This desideratum has motivated the study of *fully adaptive* composition, wherein one is allowed to adaptively select the privacy parameters of the algorithms. Rogers et al. (2016) define two probabilistic objects which can be used to ensure

privacy guarantees in fully adaptive composition. The first, called a *privacy filter*, is an adaptive stopping condition that ensures an entire interaction between an analyst and a dataset retains a pre-specified target privacy level, even when the privacy parameters are chosen adaptively. The second, called a *privacy odometer*, provides a sequence of high-probability upper bounds on how much privacy has been lost up to any point in time. While this work took the first steps towards fully adaptive composition, their filters and odometers suffered from large constants and the latter suffered from sub-optimal asymptotic rates.

We show that, as long as a target privacy level is pre-specified, one can obtain the same rate as advanced composition, including constants. We also construct families of privacy odometers that are not only tighter than the originals, but can be optimized for various target levels of privacy. Overall, we show that full adaptivity is not a cost—but rather a feature—of differential privacy.

1.1. Related Work

Privacy Composition: There is a long line of work on privacy composition. The "basic composition" theorem states that, when composing private algorithms, the privacy parameters (both ϵ and δ) add up linearly (Dwork et al., 2006b;a; Dwork and Lei, 2009). The "advanced composition" theorem allows the total ϵ to grow sublinearly with a small degradation on δ (Dwork et al., 2010). Later work (Kairouz et al., 2015; Murtagh and Vadhan, 2016) studies "optimal" composition, a computationally intractable formula that tightly characterizes the overall privacy of composed mechanisms.

More recently, several variants of privacy have been studied including (zero)-concentrated differential privacy (zCDP) (Bun and Steinke, 2016; Dwork and Rothblum, 2016), Renyi differential privacy (RDP) (Mironov, 2017), and f-differential privacy (f-DP) (Dong et al., 2021). These all exhibit tighter composition results than differential privacy, but for restricted classes of mechanisms. These results do not allow adaptive choices of privacy parameters.

Privacy Filters and Odometers: Rogers et al. (2016) originally introduced privacy filters and odometers, which allow privacy composition with adaptively selected privacy parameters. While their contributions provide a decent approximation of advanced composition, their bounds suffer from large constants, which prevents practical usage. Our work directly improves over these initial results. First, we construct privacy filters essentially matching advanced composition. We also provide flexible families of privacy odometers that outperform those of Rogers et al. (2016).

Feldman and Zrnic (2021) leverage RDP to construct Rényi filters, where they require individual mechanisms to satisfy RDP. Since our proof establishes a new privacy filter for approximate zCDP (Bun and Steinke, 2016), our results also extend to approximate RDP (Papernot and Steinke, 2022), which directly generalizes their Rényi filter. Even though it is also possible to obtain a privacy filter for (ϵ, δ)-DP through Rényi filters (Feldman and Zrnic, 2021), this result requires a stronger assumption that algorithms being composed satisfy *probabilistic* (i.e. point-wise) differential privacy (Kasiviswanathan and Smith, 2014). Since converting from differential privacy to probabilistic differential privacy can be costly (see Lemma 2), our filters demonstrate an improvement by avoiding the conversion cost.

More recently, Koskela et al. (2022) and Smith and Thakurta (2022) provide privacy filters for Gaussian DP (GDP) (Dong et al., 2021). However, their results do not hold for more general mechanisms under f-DP and therefore cannot handle algorithms with rare "catastrophic" privacy failure events, in which the privacy loss goes to infinity. Both of our (ϵ, δ)-filter and approximate zCDP filters can handle such events.

Feldman and Zrnic (2021) and Lécuyer (2021) construct RDP odometers. The former work sequentially composes Rényi filters and the latter work simultaneously runs multiple Rényi filters and takes a union bound. Neither odometer provides high probability, time-uniform bounds on privacy loss, making these results incomparable to our own. We believe our notion of odometers, which aligns with that of Rogers et al. (2016), is more natural.

To prove our results, we leverage time-uniform concentration results for martingales (Howard et al., 2020; 2021). The bounds in these papers directly improve over related self-normalized concentration results (de la Pena et al., 2004; Chen et al., 2014). These latter bounds were leveraged in Rogers et al. (2016) to construct filters and odometers.

1.2. Summary of Contributions

In this work, we provide two primary contributions. We present these results in full rigor following a brief discussion of privacy basics and martingale theory in Section 2.

Privacy Filters: In Theorem 2 of Section 3, we construct *privacy filters* that match the rate of advanced composition. Our filters significantly improve over those of Rogers et al. (2016). In fact, our proof first derives a privacy filter for approximate zCDP (Bun and Steinke, 2016) (and also approximate RDP (Papernot and Steinke, 2022)), which implies a filter for (ϵ, δ)-differential privacy using known conversion results. Our results then extend the existing filters for pure RDP in Feldman and Zrnic (2021). This extension allows us to capture a broader class of algorithms and avoids the conversion loss when translating bounds between pure RDP and (ϵ, δ)-differential privacy. We state our result in Informal

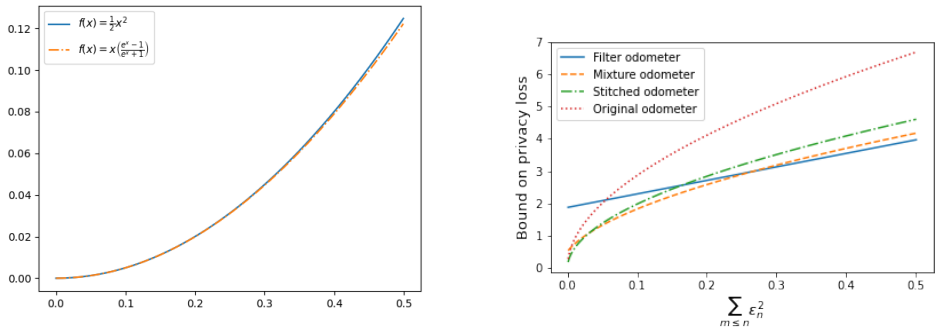

(a) Comparing lower order terms

(b) Comparing privacy odometers

Figure 1: Figure 1a compares the lower order terms of advanced composition and our privacy filter. Figure 1b compares the original odometer of Rogers et al. (2016) with our odometers (filter, mixture, and stitched).

Theorem 1 below.[1]

Privacy Odometers: In Theorem 3 of Section 4, we construct improved *privacy odometers* — that is, sequences of upper bounds on privacy loss which are all simultaneously valid with high probability. Our three families of odometers theoretically and empirically outperform those of Rogers et al. (2016). See Figure 1b for a comparison.

Informal Theorem 1 (Improved Privacy Filter)**.** *Fix target privacy parameters $\epsilon > 0$ and $\delta > 0$, and suppose $(A_n)_{n\ge1}$ is an adaptively selected sequence of algorithms. Assume that A_n is (ϵ_n, δ_n)-DP conditioned on the outputs of the first $n-1$ algorithms, where ϵ_n and δ_n may depend on outputs of $A_1, \ldots, A_{n-1}$. If a data analyst stops interacting with the data before $\sqrt{2\log\left(\frac{1}{\delta}\right)\sum_{m\le n+1}\epsilon_m^2} + \frac{1}{2}\sum_{m\le n+1}\epsilon_m^2 > \epsilon$, then the entire interaction is (ϵ, δ)-DP.*

Informal Theorem 1 almost recovers advanced composition when all parameters ϵ_n and δ_n are fixed prior to interacting with the dataset. The only difference is a slight gap in the lower order term, as $\epsilon\left(\frac{e^\epsilon-1}{e^\epsilon+1}\right) \le \frac{1}{2}\epsilon^2$. (In fact the difference between the left and right hand sides is $O(\epsilon^4)$, as can be checked with a Taylor expansion.) Figure 1a demonstrates that this gap is negligible for small values of ϵ, which is the natural setting for differential privacy.

Our second major contribution is the construction of several families of privacy odometers. These odometers give a running bound on privacy loss in settings where a target level of overall privacy is not known. Our constructed odometers are significantly tighter than the originals (Rogers et al., 2016), as can be seen in Figure 1b.

Our key insight is to view adaptive privacy composition as depending not on the number of algorithms being composed, but rather on the sums of squares of privacy parameters, $\sum_{m\le n}\epsilon_m^2$. This shift to looking at "intrinsic time" allows us to apply recent advances in time-uniform concentration (Howard et al., 2020; 2021) to privacy loss martingales. Overall, our results show that their is essentially no cost for fully adaptive private data analysis.

2. Background on Differential Privacy

Throughout, we assume all algorithms map from a space of datasets $\mathcal{X}$ to outputs in a measurable space, typically either denoted $(\mathcal{Y}, \mathcal{G})$ or $(\mathcal{Z}, \mathcal{H})$. For a sequence of algorithms $(A_n)_{n\ge1}$, we often consider the composed algorithm $A_{1:n} := (A_1, \ldots, A_n)$. For more background on measure-theoretic matters, as well as on the notion of neighboring datasets, see Appendix A.

We start by formalizing a generalization of differential privacy in which the privacy parameters of an algorithm A_n can be functions of the outputs of $A_1, \ldots, A_{n-1}$. In particular, we replace the probabilities in Equation (1) with conditional probabilities given relevant random variables.

Definition 1 (Conditional Differential Privacy)**.** *Suppose A and B are algorithms mapping from a space $\mathcal{X}$ to measurable spaces $(\mathcal{Y}, \mathcal{G})$ and $(\mathcal{Z}, \mathcal{H})$ respectively. Suppose $\epsilon, \delta : \mathcal{Z} \to \mathbb{R}_{\ge0}$ are measurable functions. We say the algorithm A is (ϵ, δ)-differentially private conditioned on B if, for any neighbors $x, x' \in \mathcal{X}$ and for all measurable sets $G \in \mathcal{G}$, we have*

$$\mathbb{P}\left(A(x) \in G \mid B(x)\right) \le e^{\epsilon(B(x))}\mathbb{P}\left(A(x') \in G \mid B(x)\right) + \delta(B(x)).$$

For conciseness, we will write either ϵ or $\epsilon(x)$ for $\epsilon(B(x))$

[1] In Appendix D, we provide an alternative proof for our privacy filter result through reductions to generalized randomized response. While it gives the exact same rates, we believe it could be of independent interest. For example, it may be useful for obtaining filters with rates like the optimal composition (Murtagh and Vadhan, 2016; Kairouz et al., 2015), which used a similar reduction to randomized response in their analysis.

and likewise δ or $\delta(x)$ for $\delta(B(x))$.

In the nth round of adaptive composition, we will set $A := A_n$ and $B := A_{1:n-1}$. In this setting, the analyst has functions $\epsilon_n, \delta_n : \mathcal{Y}^{n-1} \to \mathbb{R}_{\geq 0}$ and takes the nth round privacy parameters to be $\epsilon_n(A_{1:n-1}(x))$ and $\delta_n(A_{1:n-1}(x))$. In other words, the analyst uses the outcome of the first $n-1$ algorithms to decide the level of privacy for the nth algorithm, ensuring that A_n is (ϵ_n, δ_n)-differentially private conditioned on $A_{1:n-1}$.

We will also leverage the notion of *zero-concentrated differential privacy (zCDP)* (Bun and Steinke, 2016), which often provides a cleaner analysis for privacy composition. First, we will recall the definition of Rényi divergence.

Definition 2. *The Rényi divergence from P to Q of order $\lambda \geq 1$ is defined as*

$$D_\lambda(P\|Q) := \frac{1}{\lambda - 1} \log \left(\mathbb{E}_{Y \sim P} \left[\left(\frac{P(Y)}{Q(Y)} \right)^{\lambda - 1} \right] \right).$$

The notion of zCDP bounds the Rényi divergence from $A(x)$ to $A(x')$ for any neighbors x and x'. We will focus on a more general definition called approximate zCDP (Bun and Steinke, 2016; Papernot and Steinke, 2022) that permits a small probability of unbounded Rényi divergence. For the purpose of adaptive composition, we will state the conditional counterpart of this definition.[2]

Definition 3 (Conditional approximate zCDP)**.** *Supppose A and B are algorithms with inputs in space $\mathcal{X}$ and outputs in measurable spaces $(\mathcal{Y}, \mathcal{G})$ and $(\mathcal{Z}, \mathcal{H})$. Suppose $\delta, \rho : \mathcal{Z} \to \mathbb{R}_{\geq 0}$ are measurable. We say the algorithm A is δ-approximate ρ-zCDP conditioned on B if, for any neighboring datasets x, x', there exist distributions P', P'', Q', Q'' such that the conditional outputs are distributed according to the following mixture distributions:*

$$A(x) \mid B(x) \sim (1 - \delta(B(x)))P' + \delta(B(x))P''$$
$$A(x') \mid B(x) \sim (1 - \delta(B(x))Q' + \delta(B(x))Q'',$$

where for all $\lambda \geq 1$, $D_\lambda(P'\|Q') \leq \rho(B(x))\lambda$ and $D_\lambda(Q'\|P') \leq \rho(B(x))\lambda$. For succinctness, we will write $\rho(x)$ for $\rho(B(x))$ and $\delta(x)$ for $\delta(B(x))$.

We will also use the notions of filtration and martingales.

Filtration and Martingales: A process $(X_n)_{n\in\mathbb{N}}$ is said to be a martingale with respect to a filtration $(\mathcal{F}_n)_{n\in\mathbb{N}}$ if, for all $n \in \mathbb{N}$, (a) X_n is $\mathcal{F}_n$-measurable, (b) $\mathbb{E}|X_n| < \infty$, and (c) $\mathbb{E}(X_n \mid \mathcal{F}_{n-1}) = X_{n-1}$. Correspondingly, $(X_n)_{n\in\mathbb{N}}$ is a supermartingale if $\mathbb{E}(X_n \mid \mathcal{F}_{n-1}) \leq X_{n-1}$. In our context, we will consider the natural filtration $(\mathcal{F}_n(x))_{n\in\mathbb{N}}$ generated by $(A_n(x))_{n\geq 1}$. In our proofs, we construct the appropriate (super)martingales so that we can leverage the optional stopping theorem and time-uniform concentration to obtain privacy filters and odometers (Ville, 1939; Howard et al., 2020; 2021). We present a full exposition of the mathematical tools in Appendix A and B.

3. Privacy Filters

We now provide our main results on privacy filter. In general, a privacy filter is a function N that takes the privacy parameters of a sequence of private algorithms as input and decides to stop at some point so that the composition of these algorithms satisfies a pre-specified level of privacy. We will first present a privacy filter for approximate zCDP (Theorem 1), which will immediately imply the privacy filter result for (ϵ, δ)-DP (Theorem 2). Since approximate zCDP bounds Rényi divergence of all orders λ, our proof for Theorem 1 also directly implies a privacy fiter for approximate RDP (Papernot and Steinke, 2022), which generalizes the RDP filter by Feldman and Zrnic (2021).

Our (ϵ, δ)-DP filter improves on the rate of the original filter presented in Rogers et al. (2016) and matches the rate of advanced composition that requires pre-fixed choices of privacy parameters. Even though it is also possible to obtain an (ϵ, δ)-DP filter through the result of Feldman and Zrnic (2021), our privacy filters avoid their conversion costs and provide a tighter bound.[3]

We can now state our general privacy filter in terms of approximate zCDP.

Theorem 1 (Approximate zCDP filter)**.** *Let $(A_n)_{n\geq 1}$ be an adaptive sequence of algorithms, and, for any x, let $\mathcal{F} \equiv (\mathcal{F}_n(x))_{n\in\mathbb{N}}$ be the natural filtration generated by $(A_n(x))_{n\in\mathbb{N}}$. Assume that δ_n, ρ_n are predictable with respect to $\mathcal{F}$, meaning that they are $\mathcal{F}_{n-1}(x)$-measurable. For any $n \geq 1$, assume that A_n is δ_n-approximate ρ_n-zCDP conditioned on $\mathcal{F}_{n-1}(x)$. Consider the stopping function $N \colon \mathbb{R}^\infty_{\geq 0} \times \mathbb{R}^\infty_{\geq 0} \to \mathbb{N}$ given by*

$$N((\rho_n)_{n\geq 1}, (\delta_n)_{n\geq 1}) := \inf \left\{ n \colon \rho < \sum_{m \leq n+1} \rho_m \quad \text{or} \quad \delta < \sum_{m \leq n+1} \delta_m \right\}$$

Then $A_{1:N(\cdot)}(\cdot) \colon \mathcal{X} \to \mathcal{Y}$ is δ-approximate ρ-zCDP.

We note that the above theorem immediately implies a privacy filter for approximate RDP, and thus Theorem 1 can be

[2] The approximate zCDP definition we state uses the convex mixture formulation adapted from Papernot and Steinke (2022), since it is more convenient for our proof. In Appendix C.1, we will show that this definition is equivalent to the original definition in Bun and Steinke (2016).

[3] Feldman and Zrnic (2021, Section 4.3) apply Rényi filters to algorithms which satisfy (conditional) probabilistic differential privacy (pDP). In general, a lossy conversion from (ϵ, δ)-DP to (ϵ, δ)-pDP is required to apply their filter.

viewed as a strict generalization of the work of Feldman and Zrnic (2021). Further, Theorem 1 implies a privacy filter under (ϵ, δ)-differential privacy. To show this implication, we will use the following conversion results.

Lemma 1 ((Bun and Steinke, 2016))**.** *If A satisfies (ϵ, δ)-DP, then A satisfies δ-approximate $\frac{1}{2}\epsilon^2$-zCDP. If A satisfies δ-approximate ρ-zCDP, then A satisfies $(\rho + 2\sqrt{\rho \ln(1/\delta')}, \delta + (1-\delta)\delta')$-DP.*

We can now obtain our (ϵ, δ)-privacy filter by a conversion of individual approximate differential privacy parameters to approximate zCDP ones, application of the approximate zCDP filter, and the conversion of approximate zCDP back to approximate differential privacy.

Theorem 2 ((ϵ, δ)-DP filter)**.** *Suppose $(A_n)_{n\geq 1}$ is a sequence of algorithms such that, for any $n \geq 1$, A_n is (ϵ_n, δ_n)-differentially private conditioned on $A_{1:n-1}$. Let $\epsilon > 0$ and $\delta = \delta' + \delta''$ be target privacy parameters such that $\delta' > 0, \delta'' \geq 0$. Let $N : \mathbb{R}^\infty_{\geq 0} \times \mathbb{R}^\infty_{\geq 0} \to \mathbb{N}$ be given by*

$$N((\epsilon_n)_{n\geq 1}, (\delta_n)_{n\geq 1}) := \inf\left\{ n : \epsilon < \sqrt{2\log\left(\frac{1}{\delta'}\right) \sum_{m\leq n+1} \epsilon_m^2} + \frac{1}{2}\sum_{m\leq n+1} \epsilon_m^2 \text{ or } \delta'' < \sum_{m\leq n+1} \delta_m \right\}.$$

Then, the algorithm $A_{1:N(\cdot)}(\cdot) : \mathcal{X} \to \mathcal{Y}^\infty$ is (ϵ, δ)-DP, where $N(x) := N((\epsilon_n(x))_{n\geq 1}, (\delta_n(x))_{n\geq 1})$.

Now we provide a proof for Theorem 1. Recall that the output under a privacy filter is a random vector $A_{1:N(x)}(x) = (A_1(x), \ldots, A_{N(x)}(x))$. In our proof, we will also consider the unstopped process $A(x) = (A_1(x), A_2(x), \ldots)$.

Proof of Theorem 1. Let x, x' be neighbors, and P, Q denote the likelihoods of the observed output $A(x)$ when the inputs are x, x' respectively. The likelihoods of observing the stopped process $A_{1:N(x)}(x)$ under x and x' are:

$$P(A_{1:N(x)}(x)) = \prod_{n=1}^{N(x)} P(A_n(x) \mid \mathcal{F}_{n-1}(x)), \quad (3)$$

$$Q(A_{1:N(x)}(x)) = \prod_{n=1}^{N(x)} Q(A_n(x) \mid \mathcal{F}_{n-1}(x)). \quad (4)$$

It suffices to show that the two likelihoods can be decomposed as weighted mixtures of P' and P'', and Q' and Q'' respectively such that the mixture weights on P' and Q' are at least $(1-\delta)$ and for all $\lambda \geq 1$,

$$\max\left\{ D_\lambda\left(P'(A_{1:N(x)}(x)) \| Q'(A_{1:N(x)}(x))\right), \; D_\lambda\left(Q'(A_{1:N(x)}(x)) \| P'(A_{1:N(x)}(x))\right) \right\} \leq \rho\lambda. \quad (5)$$

By our assumption of conditional approximate zCDP at each step n, we can write $P(A_n(x) \mid \mathcal{F}_{n-1}(x))$ and $Q(A_n(x) \mid \mathcal{F}_{n-1}(x))$ as the following convex combinations:

$$P(A_n(x) \mid \mathcal{F}_{n-1}(x)) = (1-\delta_n(x))P'_n(A_n(x) \mid \mathcal{F}_{n-1}(x)) + \delta_n(x)P''_n(A_n(x) \mid \mathcal{F}_{n-1}(x)),$$

$$Q(A_n(x) \mid \mathcal{F}_{n-1}(x)) = (1-\delta_n(x))Q'_n(A_n(x) \mid \mathcal{F}_{n-1}(x)) + \delta_n(x)Q''_n(A_n(x) \mid \mathcal{F}_{n-1}(x)),$$

such that for all $\lambda \geq 1$, we have both

$$D_\lambda\left(P'_n(A_n(x) \mid \mathcal{F}_{n-1}(x)) \,\|\, Q'_n(A_n(x) \mid \mathcal{F}_{n-1}(x))\right) \leq \rho_n(x)\lambda, \quad (6)$$

$$D_\lambda\left(Q'_n(A_n(x) \mid \mathcal{F}_{n-1}(x)) \,\|\, P'_n(A_n(x) \mid \mathcal{F}_{n-1}(x))\right) \leq \rho_n(x)\lambda. \quad (7)$$

Now consider the product measures P' and Q' such that for any $n \geq 1$,

$$P'(A_{1:n}(x)) = \prod_{m=1}^{n} P'_m(A_m(x) \mid \mathcal{F}_{m-1}(x)) \text{ and}$$

$$Q'(A_{1:n}(x)) = \prod_{m=1}^{n} Q'_m(A_m(x) \mid \mathcal{F}_{m-1}(x)). \quad (8)$$

We will establish inequality (5). For any fixed $\lambda \geq 1$, consider the following processes:

$$M_n := \sum_{m\leq n} \left\{ \log\left(\frac{P'_m(A_m(x) \mid \mathcal{F}_{m-1}(x))}{Q'_m(A_m(x) \mid \mathcal{F}_{m-1}(x))}\right) - \lambda\rho_m(x) \right\}, \quad (9)$$

$$X_n := \exp\left((\lambda - 1)M_n\right). \quad (10)$$

By Lemma 6, X_n is a nonnegative P'-supermartingale with respect to $(\mathcal{F}_n(x))_{n\in\mathbb{N}}$. By the optional stopping theorem for nonnegative supermartingales (Lemma 5), we have

$$\mathbb{E}_{P'}[X_{N(x)}] \leq \mathbb{E}_{P'}[X_0] = 1. \quad (11)$$

By plugging in the definition of X_n and the stopping criterion of N, we can bound the Rényi divergence $D_\lambda\left(P'(A_{1:N(x)}(x)) \| Q'(A_{1:N(x)}(x))\right) \leq \rho\lambda$ (see Lemma 7), and so inequality (5) holds by symmetry.

Finally, by Lemma 8, we can rewrite both P and Q as weighted mixtures containing P' and Q', with weights at least $1-\delta$. This completes the proof. □

4. Privacy Odometers

Previously, we constructed privacy filters that matched the rate of advanced composition while allowing *both* algorithms and privacy parameters to be chosen adaptively. While privacy filters require the total level of privacy to be fixed in advance, it is desirable to track the privacy loss at all steps without a pre-fixed budget (Ligett et al., 2017). We now study privacy odometers which provide sequences of upper bounds on accumulated privacy loss that are valid at all points in time simultaneously with high probability.

4.1. Background on Privacy Loss and Odometers

To formally introduce privacy odometers, we will first revisit the notion of *privacy loss*, which measures how much information is revealed about the underlying input dataset. For neighbors $x, x' \in \mathcal{X}$, let p^x and $p^{x'}$ be the densities of $A(x)$ and $A(x')$ respectively. The privacy loss between $A(x)$ and $A(x')$ is defined as

$$\mathcal{L}(x, x') := \log\left(\frac{p^x(A(x))}{p^{x'}(A(x))}\right). \tag{12}$$

By Equation (12), a negative privacy loss suggests that the input is more likely to be x', and likewise a positive privacy loss suggests that the input is more likely to be x. We now generalize privacy loss to its conditional counterpart.

Definition 4 (Conditional Privacy Loss)**.** *Suppose A and B are as in Definition 1. Suppose $x, x' \in \mathcal{X}$ are neighbors. Let $p^x(\cdot|\cdot), p^{x'}(\cdot|\cdot) : \mathcal{Y} \times \mathcal{Z} \to \mathbb{R}_{\geq 0}$ be conditional densities for $A(x)$ and $A(x')$ respectively given $B(x)$.*[4] *The privacy loss between $A(x)$ and $A(x')$ conditioned on B is given by*

$$\mathcal{L}_B(x, x') := \log\left(\frac{p^x(A(x)|B(x))}{p^{x'}(A(x)|B(x))}\right).$$

Suppose A_n is the nth algorithm being run and we have already observed $A_{1:n-1}(x)$ for some unknown input $x \in \mathcal{X}$. If we are trying to guess whether x or a neighbor x' produced the data, we would consider the privacy loss between $A_n(x)$ and $A_n(x')$ conditioned on $A_{1:n-1}(x)$. It is straightforward to characterize the privacy loss of a composed algorithm $A_{1:n}$ in terms of the privacy loss of each constituent algorithm $A_1, \cdots, A_n$. Namely, from Bayes rule,

$$\mathcal{L}_{1:n}(x, x') = \sum_{m \leq n} \mathcal{L}_m(x, x'), \tag{13}$$

where $\mathcal{L}_m(x, x')$ is shorthand for the conditional privacy loss between $A_m(x)$ and $A_m(x')$ given $A_{1:m-1}(x)$, per Definition 4. Equation (13) also holds at arbitrary random times $N(x)$ that only depend on the dataset $x \in \mathcal{X}$ through observed algorithm outputs.

The simple decomposition of privacy loss noted above motivates the study of an "alternative", probabilistic definition of differential privacy. Intuitively, an algorithm should be differentially private if, with high probability, the privacy loss is small. More formally, an algorithm $A : \mathcal{X} \to \mathcal{Y}$ is said to be (ϵ, δ)-*probabilistically differentially private*, or (ϵ, δ)-pDP for short, if, for all neighboring inputs $x, x' \in \mathcal{X}$, we have $\mathbb{P}(|\mathcal{L}(x, x')| > \epsilon) \leq \delta$. In the previous line (as well as in the remainder of the section), the randomness in $\mathcal{L}(x, x')$ comes from the randomized algorithm A.

Unfortunately, as noted by Kasiviswanathan and Smith (2014) (in which pDP is called *point-wise indistinguishability*), pDP is a strictly stronger notion than DP. In particular, if an algorithm is (ϵ, δ)-pDP, it is also (ϵ, δ)-DP. The converse in general requires a costly conversion.

Lemma 2 (Conversions between DP and pDP (Kasiviswanathan and Smith, 2014))**.** *If A is (ϵ, δ)-pDP, then A is also (ϵ, δ)-DP. Conversely, if A is (ϵ, δ)-DP, then A is $(2\epsilon, \frac{2\delta}{\epsilon e^{\epsilon}})$-pDP.*

We note that that Guingona et al. (2023) have recently shown that other possible conversion rates from probabilistic differential privacy to approximate differential privacy are possible. However, we note that these conversions require trading off tightness in the approximation parameter ϵ and the approximation parameter δ. In particular, a fully tight conversion from probabilistic differenial privacy to approximate differential privacy is not possible. We will work with the conditional counterpart of probabilistic differential privacy (pDP).

Definition 5 (Conditional Probabilistic Differential Privacy)**.** *Suppose $A : \mathcal{X} \to \mathcal{Y}$ and $B : \mathcal{X} \to \mathcal{Z}$ are algorithms, and $\epsilon, \delta : \mathcal{Z} \to \mathbb{R}_{\geq 0}$ are measurable. Then, A is said to be (ϵ, δ)-probabilistically differentially private conditioned on B if, for any neighbors $x, x' \in \mathcal{X}$, we have*

$$\mathbb{P}(|\mathcal{L}_B(x, x')| > \epsilon(B(x))|B(x)) \leq \delta(B(x)).$$

While in Theorem 2 we assumed that the algorithms being composed were *conditionally differentially private*, here, we need to assume *conditional probabilistic privacy*. This is because our goal is not differential privacy, but rather tight control over privacy loss. We conjecture that a version of our privacy odometer (in Theorem 3) that replaces pDP by DP and leaves all else identical does not hold. Our intuition for this conjecture is that there exist simple examples of algorithms satisfying (ϵ, δ)-DP that don't satisfy (ϵ, δ)-pDP (see Appendix F, for instance). We believe that, by sequentially composing such algorithms and using anti-concentration results, one can show that some odometers fail to be valid. We leave this as potential future work. In sequential composition, we would assume the nth algorithm A_n is (ϵ_n, δ_n)-pDP conditioned on $A_{1:n-1}$. The privacy parameters would be given as functions of $A_{1:n-1}(x)$. Now we state the definition of privacy odometer, which provides bounds on privacy loss under arbitrary stopping conditions (e.g. conditions based on model accuracy).

Definition 6 (Privacy Odometer (Rogers et al., 2016))**.** *Let $(A_n)_{n \geq 1}$ be an adaptive sequence of algorithms such that, for all $n \geq 1$, A_n is (ϵ_n, δ_n)-pDP conditioned on*

[4] To ensure the existence of conditional densities, it suffices to assume that $\mathcal{Y}$ and $\mathcal{Z}$ are *Polish spaces* under some metrics $d_{\mathcal{Y}}$ and $d_{\mathcal{Z}}$, and that $\mathcal{G}$ and $\mathcal{H}$ are the corresponding Borel σ-algebras associated with $d_{\mathcal{Y}}$ and $d_{\mathcal{Z}}$ (Durrett, 2019). These measurability assumptions are not restrictive, as Euclidean spaces, countable spaces, and Cartesian products of the two satisfy these assumption.

$A_{1:n-1}$. *Let* $(u_n)_{n\geq 1}$ *be a sequence of functions where* $u_n : \mathbb{R}_{\geq 0}^{n-1} \times \mathbb{R}_{\geq 0}^{n-1} \to \mathbb{R}_{\geq 0}$. *Let* $\delta \in (0,1)$ *be a target confidence parameter. For* $x \in \mathcal{X}, n \geq 1$, *define* $U_n(x) := u_n(\epsilon_{1:n-1}(x), \delta_{1:n-1}(x))$. *Then,* $(u_n)_{n\geq 1}$ *is called a* δ*-privacy odometer if, for all* $x, x' \in \mathcal{X}$ *neighbors, we have*

$$\mathbb{P}\left(\exists n \geq 1 : \mathcal{L}_{1:n}(x, x') > U_n(x)\right) \leq \delta.$$

4.2. Improved Privacy Odometers

We construct our privacy odometers in Theorem 3. Our technical centerpiece is time-uniform concentration inequalities for martingales (Ville, 1939; Howard et al., 2020; 2021). For a martingale $(M_n)_{n\in\mathbb{N}}$ and confidence level $\delta > 0$, time-uniform concentration inequalities provides bounds $(U_n)_{n\in\mathbb{N}}$ satisfying $\mathbb{P}(\exists n \in \mathbb{N} : M_n > U_n) \leq \delta$. Thus, if we can create a martingale from privacy loss, we can use time-uniform concentration to construct odometers. Our proof first considers the case where each A_n is $(\epsilon_n, 0)$-pDP and the *privacy loss martingale* $(M_n)_{n\in\mathbb{N}}$ (Dwork et al., 2010) is given by $M_0 = 0$ and:

$$M_n := M_n(x, x') := \mathcal{L}_{1:n}(x, x') - \sum_{m\leq n} \mathbb{E}\left(\mathcal{L}_m(x, x') | \mathcal{F}_{n-1}(x)\right) \tag{14}$$

We then extend to the case of $\delta_n \geq 0$ via conditioning.

To construct their filters and odometers, Rogers et al. (2016) use self-normalized concentration inequalities (de la Pena et al., 2004; Chen et al., 2014). We instead use advances in time-uniform martingale concentration (Howard et al., 2020; 2021), which yields tighter results.

Theorem 3. *Suppose* $(A_n)_{n\geq 1}$ *is a sequence of algorithms such that, for any* $n \geq 1$, A_n *is* (ϵ_n, δ_n)*-pDP conditioned on* $A_{1:n-1}$. *Let* $\delta = \delta' + \delta''$ *be a target approximation parameter such that* $\delta' > 0, \delta'' \geq 0$. *Define* $N := N((\delta_n)_{n\geq 1}) := \inf\left\{n \in \mathbb{N} : \delta'' < \sum_{m\leq n+1} \delta_m\right\}$ *and* $V_n := \sum_{m\leq n} \epsilon_m^2$. *Define the following:*

1. ***Filter odometer.*** *For any* $\epsilon > 0$, *let* $y^* := \left(-\sqrt{2\log\left(\frac{1}{\delta'}\right)} + \sqrt{2\log\left(\frac{1}{\delta'}\right) + \epsilon}\right)^2$. *Define the functions* $(u_n^F)_{n\geq 1}$ *by*

$$u_n^F(\epsilon_{1:n}, \delta_{1:n}) := \begin{cases} \infty & n > N \\ \frac{\sqrt{2y^*\log\left(\frac{1}{\delta'}\right)}}{2} + \frac{\sqrt{2\log\left(\frac{1}{\delta'}\right)}}{2\sqrt{y^*}} V_n + \frac{1}{2}V_n & \text{otherwise.} \end{cases}$$

2. ***Mixture odometer.*** *For any* $\gamma > 0$, *define the sequence of functions* $(u_n^M)_{n\geq 1}$ *by*

$$u_n^M(\epsilon_{1:n}, \delta_{1:n}) := \begin{cases} \infty & n > N \\ \sqrt{2\log\left(\frac{1}{\delta'}\sqrt{\frac{V_n+\gamma}{\gamma}}\right)(\gamma + V_n)} + \frac{1}{2}V_n & \text{otherwise.} \end{cases}$$

3. ***Stitched odometer.*** *For any* $v_0 > 0$, *define the sequence of functions* $(u_n^S)_{n\geq 1}$ *by*

$$u_n^S(\epsilon_{1:n}, \delta_{1:n}) := \begin{cases} \infty & n > N \text{ or } V_n < v_0 \\ 1.7\sqrt{V_n\left(\log\log\left(\frac{2V_n}{v_0}\right) + 0.72\log\left(\frac{5.2}{\delta'}\right)\right)} + \frac{1}{2}V_n & \text{else.} \end{cases}$$

Then, any of the sequences $(u_n^F)_{n\geq 1}$, $(u_n^M)_{n\geq 1}$, *or* $(u_n^S)_{n\geq 1}$ *is a* δ*-privacy odometer.*

The proof of Theorem 3 can be found in Appendix E. We now provide intuition for our odometers, which are plotted in Figure 3. Our insight is to view odometers not as functions of the number of algorithms being composed, but rather as functions of the intrinsic time $\sum_{m\leq n} \epsilon_m^2$. This reframing allows us to leverage the various time-uniform concentration inequalities discussed in Appendix B. The filter odometer is the tightest odometer when the value $\sum_{m\leq n} \epsilon_m^2$ is close to a "fixed time" y^*, but the tightness drops off precipitously when $\sum_{m\leq n} \epsilon_m^2$ is far from y^*. The mixture odometer, which is named after the *the method of mixtures* (Robbins, 1970; de la Peña et al., 2007; Howard et al., 2021), sacrifices tightness at any fixed point in time to obtain overall tighter bounds on privacy loss. This odometer can be numerically optimized, in terms of ρ, for tightness at a predetermined value $\sum_{m\leq n} \epsilon_m^2$. The stitched odometer, whose name derives from Theorem 6, is similarly tight across time. This odometer requires that $\sum_{m\leq n} \epsilon_m^2$ exceed some pre-selected "variance" v_0 before becoming nontrivial (i.e. finite). Larger values of v_0 will yield tighter odometers, albeit at the cost of losing bound validity when accumulated variance is small. With this intuition, we can compare our odometers to the original presented in Rogers et al. (2016).

Lemma 3 (Theorem 6.5 in Rogers et al. (2016))**.** *Assume the same setup as Theorem 3, and fix* $\delta = \delta' + \delta''$, *where* $\frac{1}{e} \geq \delta' > 0$ *and* $\delta'' \geq 0$. *Define the sequence of functions* $(u_n^R)_{n\geq 1}$ *by*

$$u_n^R(\epsilon_{1:n}, \delta_{1:n}) := \begin{cases} \infty, & n > N \\ \sqrt{2V_n\left(\log(110e) + 2\log\left(\frac{\log(|x|)}{\delta'}\right)\right)} + \frac{1}{2}V_n & n \leq N, V_n \in \left[\frac{1}{|x|^2}, 1\right] \\ \sqrt{2\left(\frac{1}{|x|^2} + V_n\right)\left(1 + \frac{1}{2}\log\left(1 + |x|^2 V_n\right)\right)\log\log\left(\frac{4}{\delta'}\log_2(|x|)\right)} + \frac{1}{2}V_n, & \text{otherwise} \end{cases},$$

where $|x|$ *denotes the number of elements in dataset* x. *Then,* $(u_n^R)_{n\geq 1}$ *is a* δ*-privacy odometer.*

Our new odometers improve over the one presented in Lemma 3. First, the above odometer has an explicit dependence on dataset size. In learning settings, datasets are large, degrading the quality of the odometer. Secondly, the tightness of the odometer drops off outside of the interval $\left[\frac{1}{|x|^2}, 1\right]$. If *any* privacy parameter of an algorithm

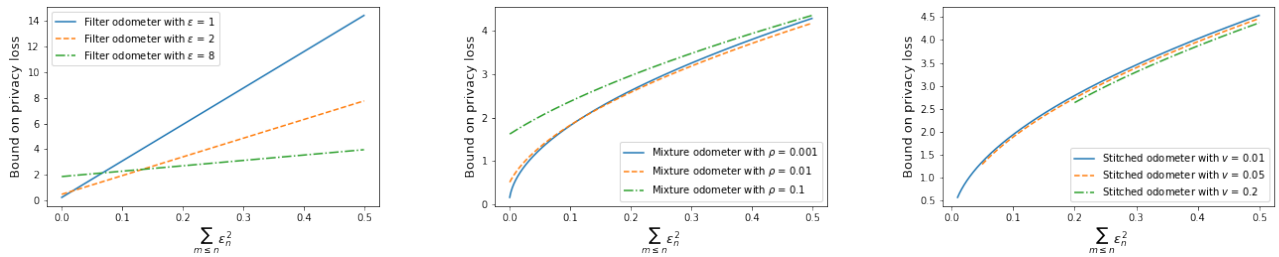

(a) Comparing filter odometers **(b)** Comparing mixture odometers **(c)** Comparing stitched odometers

Figure 2: Comparison of filter, mixture, and stitched odometers plotted as functions of $\sum_{m \le n} \epsilon_m^2$. We set $\delta' = 10^{-6}$ and assume all algorithms being composed are purely differentially private for simplicity.

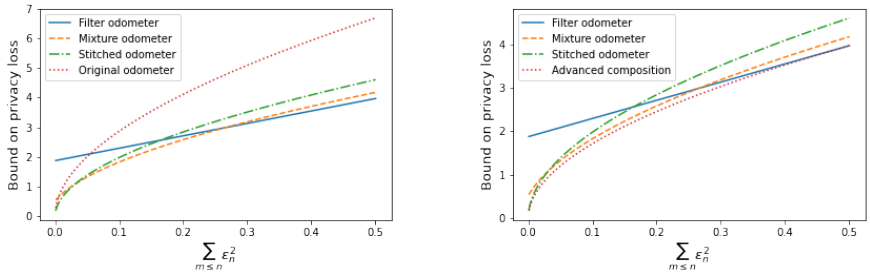

(a) New odometers vs. original **(b)** New odometers vs. pointwise advanced composition

Figure 3: Figure 3a compares our odometers to the original. Figure 3b compares them with advanced composition optimized point-wise. The curve plotted for advanced composition is valid at any fixed time, but not uniformly over time. Our odometers nevertheless provide a close approximation.

being composed exceeds 1, the bound becomes significantly looser. Lastly, and perhaps most simply, the form of the odometer is complicated. Our odometers all have relatively straightforward dependence on the intrinsic time $\sum_{m \le n} \epsilon_m^2$.

We now examine the rates of all odometers. For simplicity, let $v := \sum_{m \le n} \epsilon_m^2$. The stitched odometer has a rate of $O(\sqrt{v \log \log(v)})$ in its leading term, asymptotically matching the law of the iterated logarithm (Robbins, 1970) up to constants. Both the original privacy odometer and the mixture odometer have a rate of $O(\sqrt{v \log(v)})$, demonstrating worse asymptotic performance. The filter odometer has the worst asymptotic performance, growing linearly as $O(v)$. This does not mean the stitched odometer is the best odometer, since target levels of privacy are often kept small.

To empirically compare odometers, it suffices to consider the setting of *pure* differential privacy, as the odometers identically depend on $(\delta_n)_{n \ge 1}$. Each presented odometer can be viewed as a function of v, allowing us to compare odometers by plotting their values for a continuum of v. Figure 3a shows that there is no clearly tightest odometer. All odometers, barring the original, dominate for some window of values of v. While the stitched odometer is asymptotically best, the mixture odometer is tighter for small values of v. Likewise, if one knows an approximate target privacy level, the filter odometer is tightest. This behavior is expected from our understanding of martingale concentration (Howard et al., 2020; 2021): there is no uniformly tightest boundary containing (with probability $1 - \delta$) the entire path of a martingale; boundaries that are tight early must be looser later, and vice versa. In fact, we conjecture that our bounds are essentially unimprovable in general — this conjecture stems from the fact that the time-uniform martingale boundaries employed have error probability *essentially* equal to δ, which in turn stems from the deep fact that for continuous-path (and thus continuous-time) martingales, Ville's inequality (Fact 4)—that underlies the derivation of these boundaries—holds with exact equality. Since we operate in discrete-time, the only looseness in Ville's inequality stems from lower-order terms that reflect the possibility that at the stopping time, the value of the stopped martingale may not be *exactly* the value at the boundary.

In Figure 3b, we compare our odometers with advanced composition optimized in a point-wise sense for all values of v simultaneously. This boundary *is not a valid odometer*, as advanced composition only holds at a prespecified point in intrinsic time v. Our odometers are almost tight with advanced composition for the values of v plotted. Our filter

odometer lies tangent to the advanced composition curve, as expected from Section 5.2 of Howard et al. (2020).

5. Future Directions

There are many open problems related to fully adaptive composition. For example, even though privacy filters has been studied under the notion of Gaussian DP (Smith and Thakurta, 2022; Koskela et al., 2022), privacy filters and odometers have not been studied for general f-DP (Dong et al., 2021). It also has not been investigated whether adaptivity in privacy parameter selection improves the performance of iterative algorithms such as private SGD. Intuitively, it should be beneficial to let the iterates of an algorithm guide future choices of privacy parameters. Optimal composition results (Kairouz et al., 2015; Murtagh and Vadhan, 2016; Zhu et al., 2022) have yet to be considered in a setting where privacy parameters are adaptively selected. In Appendix D, we provide another proof of Theorem 2, which leverages a reduction of private algorithms to generalized randomized response. Since such a reduction was used in the proofs of Kairouz et al. (2015) and Murtagh and Vadhan (2016), we believe this proof can be useful for optimal composition with adaptively chosen privacy parameters.

6. Acknowledgements

AR acknowledges support from NSF DMS 1916320 and an ARL IoBT CRA grant. Research reported in this paper was sponsored in part by the DEVCOM Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196 (ARL IoBT CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

ZSW and JW were supported in part by the NSF CNS2120667, NSF Award #2120667, a CyLab 2021 grant, a Google Faculty Research Award, and a Mozilla Research Grant.

JW acknowledges support from NSF GRFP grants DGE1745016 and DGE2140739.

References

David Blackwell. Equivalent comparisons of experiments. *The annals of mathematical statistics*, pages 265–272, 1953.

Mark Bun and Thomas Steinke. Concentrated differential privacy: simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016.

Shanshan Chen, Zhenping Wang, Wenfei Xu, and Yu Miao. Exponential inequalities for self-normalized martingales. *Journal of Inequalities and Applications*, 2014(289):1–12, 2014.

Victor H de la Pena, Michael J Klass, and Tze Leung Lai. Self-normalized processes: exponential inequalities, moment bounds, and iterated logarithm laws. *Annals of Probability*, pages 1902–1933, 2004.

Victor H. de la Peña, Michael J. Klass, and Tze Leung Lai. Pseudo-maximization and self-normalized processes. *Probability Surveys*, 4:172 – 192, 2007. doi: 10.1214/07-PS119.

Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. In *Journal of the Royal Statistical Society: Series B*, pages 1–35, 2021.

Richard Durrett. *Probability: theory and examples*. Duxbury Press, Belmont, CA, second edition, 1996. ISBN 0-534-24318-5.

Rick Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.

Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, pages 371–380, 2009.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.

Cynthia Dwork and Guy N. Rothblum. Concentrated differential privacy. *CoRR*, abs/1603.01887, 2016.

Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–503. Springer, 2006a.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography*, pages 265–284, Berlin, Heidelberg, 2006b. Springer Berlin Heidelberg.

Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 51–60. IEEE, 2010.

Vitaly Feldman and Tijana Zrnic. Individual privacy accounting via a Rényi filter. *Advances in Neural Information Processing Systems*, 2021.

Vincent Guingona, Alexei Kolesnikov, Julianne Nierwinski, and Avery Schweitzer. Comparing approximate and probabilistic differential privacy parameters. *Information Processing Letters*, page 106380, 2023.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform Chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257 – 317, 2020. doi: 10.1214/18-PS321.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49 (2):1055 – 1080, 2021. doi: 10.1214/20-AOS1991.

Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In *International Conference on Machine Learning*, pages 1376–1385. PMLR, 2015.

Shiva P Kasiviswanathan and Adam Smith. On the semantics of differential privacy: A Bayesian formulation. *Journal of Privacy and Confidentiality*, 6(1), 2014.

Emilie Kaufmann and Wouter M Koolen. Mixture martingales revisited with applications to sequential tests and confidence intervals. *Journal of Machine Learning Research*, 22(246):1–44, 2021.

Antti Koskela, Marlon Tobaben, and Antti Honkela. Individual privacy accounting with gaussian differential privacy. *CoRR*, abs/2209.15596, 2022. doi: 10.48550/arXiv.2209.15596. URL https://doi.org/10.48550/arXiv.2209.15596.

Mathias Lécuyer. Practical privacy filters and odometers with Rényi differential privacy and applications to differentially private deep learning. *arXiv Preprint arXiv:2103.01379*, 2021.

Katrina Ligett, Seth Neel, Aaron Roth, Bo Waggoner, and Steven Z Wu. Accuracy first: Selecting a differential privacy level for accuracy constrained erm. *Advances in Neural Information Processing Systems*, 30, 2017.

Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.

Jack Murtagh and Salil Vadhan. The complexity of computing the optimal composition of differential privacy. In *Theory of Cryptography Conference*, pages 157–175. Springer, 2016.

Nicolas Papernot and Thomas Steinke. Hyperparameter tuning with renyi differential privacy. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL https://openreview.net/forum?id=-70L8lpp9DF.

Herbert Robbins. Statistical methods related to the law of the iterated logarithm. *The Annals of Mathematical Statistics*, 41(5):1397–1409, 1970.

Ryan M Rogers, Aaron Roth, Jonathan Ullman, and Salil Vadhan. Privacy odometers and filters: pay-as-you-go composition. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.

Adam D. Smith and Abhradeep Thakurta. Fully adaptive composition for gaussian differential privacy. *CoRR*, abs/2210.17520, 2022. doi: 10.48550/arXiv.2210.17520. URL https://doi.org/10.48550/arXiv.2210.17520.

Jean Ville. Etude critique de la notion de collectif. *Bull. Amer. Math. Soc*, 45(11):824, 1939.

Yuqing Zhu, Jinshuo Dong, and Yu-Xiang Wang. Optimal accounting of differential privacy via characteristic function. In *International Conference on Artificial Intelligence and Statistics*, pages 4782–4817. PMLR, 2022.

A. Measure-Theoretic Formalism

Below, we provide some measure-theoretic formalisms and details regarding datasets and neighboring relations.

Neighboring Datasets: Roughly speaking, an algorithm is differentially private if it difficult to distinguish between output distributions when the algorithm is run on similar inputs. In general, this notion of similarity amongst inputs is defined as a *neighboring relation* $\sim$ between elements on the input space $\mathcal{X}$. In particular, if two inputs (also referred to as datasets or databases) $x, x' \in \mathcal{X}$ satisfy the neighboring relation $x \sim x'$, the we say x and x' are *neighbors*.

There are several canonical examples of neighboring relations on the space of inputs $\mathcal{X}$. One example is where $\mathcal{X} = \mathbb{X}^n$ for some data domain $\mathbb{X}$. The data domain can be viewed as the set of all possible individual entries for a dataset, and the space $\mathbb{X}^n$ correspondingly contains all possible n element datasets. In this setting, databases $x, x' \in \mathcal{X}$ may be considered neighbors if x and x' differ in exactly one entry. Another slightly more general setting is when $\mathcal{X} = \mathbb{X}^*$, i.e., all possible datasets of finite size. In this situation, the earlier notion of neighboring still makes sense. However, in addition, we may say input datasets x and x' are neighbors if x can be obtained from x' by either adding

or deleting an element. This is a very natural notion of neighboring, as under such a relation an algorithm would be differentially private if it were difficult to determine the presence or absence of an individual. Our work is agnostic to the precise choice of neighboring relation. As such, we choose to leave the notion as general as possible.

Algorithms and Random Variables: We will consider algorithms as randomized mappings $A : \mathcal{X} \to \mathcal{Y}$ taking inputs from $\mathcal{X}$ to some output space $\mathcal{Y}$. To be fully formal, we consider the output space $\mathcal{Y}$ as a *measurable space* $(\mathcal{Y}, \mathcal{G})$, where $\mathcal{G}$ is some σ-algebra denoting possible events. Recall that a σ-algebra $\mathcal{S}$ for a set S is simply a subset of 2^S containing S and $\emptyset$ that is closed under countable union, intersection, and complements. When we say A is an algorithm having inputs in some space $\mathcal{X}$, we really mean $A(x)$ is a $\mathcal{Y}$-valued random variable for any $x \in \mathcal{X}$. The space $\mathcal{X}$ need not have an associated σ-algebra, as algorithm inputs are essentially just indexing devices. Given a sequence of algorithms $(A_n)_{n\geq 1}$, $(A_n(x))_{n\geq 1}$ is a sequence of $\mathcal{Y}$-valued random variables, for any $x \in \mathcal{X}$.[5]

Since we are dealing with the composition of algorithms, we write $A_{1:n}(x)$ as shorthand for the random vector of the first n algorithm outputs, i.e. $A_{1:n}(x) = (A_1(x), \ldots, A_n(x))$. Formally, the random vector $A_{1:n}(x)$ takes output values in the product measurable space $(\mathcal{Y}^n, \mathcal{G}^{\otimes n})$ where $\mathcal{G}^{\otimes n}$ denotes the n-fold product σ-algebra of $\mathcal{G}$ with itself. Likewise, since the number of algorithm outputs one views in fully-adaptive composition may be random, if N is a random time (i.e. a $\mathbb{N}$-valued random variable), we will often consider the random vector $A_{1:N}(x) = (A_1(x), \ldots, A_N(x))$.

Filtrations and Stopping Times: Since privacy composition involves sequences of random outputs, we will use the measure-theoretic notion of a *filtration*. If we have fixed an input $x \in \mathcal{X}$, we can assume the random sequence $(A_n(x))_{n\geq 1}$ is defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Given such a probability space, a filtration $(\mathcal{F}_n)_{n\in\mathbb{N}}$ of $\mathcal{F}$ is a sequence of σ-algebras satisfying: (i) $\mathcal{F}_n \subset \mathcal{F}_{n+1}$ for all $n \in \mathbb{N}$, and (ii) $\mathcal{F}_n \subset \mathcal{F}$ for all $n \in \mathbb{N}$. Given an arbitrary $\mathcal{Y}$-valued discrete-time stochastic process $(X_n)_{n\geq 1}$, it is often useful to consider the *natural filtration* $(\mathcal{F}_n)_{n\in\mathbb{N}}$ given by $\mathcal{F}_n := \sigma(X_m : m \leq n)$ and $\mathcal{F}_0 = \{\emptyset, \Omega\}$. Intuitively, a filtration formalizes the notion of accumulating information over time. In particular, in the context of the natural filtration generated by a stochastic process, the nth σ-algebra in the filtration $\mathcal{F}_n$ essentially represents the entirety of information contained in the first n random variables. In other words, if one is given $\mathcal{F}_n$, they would know all possible events/outcomes that could have occurred up to and including timestep n.

Lastly, we briefly mention the notion of a *stopping time*, as this measure-theoretic object is necessary to define privacy filters. Given a filtration $(\mathcal{F}_n)_{n\in\mathbb{N}}$, a random time N is said to be a stopping time with respect to $(\mathcal{F}_n)_{n\in\mathbb{N}}$ if, for any n, the event $\{N \leq n\} \in \mathcal{F}_n$. In words, a random time N is a stopping time if given the information in $\mathcal{F}_n$ we can determine whether or not we should have stopped by time n. Stopping times are essential to the study of fully-adaptive composition, as a practitioner of privacy will need to use the adaptively selected privacy parameters to determine whether or not to stop interacting with the underlying sensitive database.

[5] Even if algorithms have different types of outputs (maybe some algorithms have categorical outputs while others output real-valued vectors), $\mathcal{Y}$ can still be made appropriately large to contain all possible outcomes.

B. Martingale Inequalities

In this appendix, we provide a thorough exposition into the concentration inequalities leveraged in this paper. First, at the heart of supermartingale concentration is Ville's inequality (Ville, 1939), which can be viewed as a time-uniform version of Markov's inequality.

Lemma 4 (Ville's Inequality (Ville, 1939))**.** *Let $(X_n)_{n\in\mathbb{N}}$ be a nonnegative supermartingale with respect to some filtration $(\mathcal{F}_n)_{n\in\mathbb{N}}$. Then, for any confidence parameter $\delta \in (0,1)$, we have $\mathbb{P}\left(\exists n \in \mathbb{N} : X_n \geq \frac{\mathbb{E}X_0}{\delta}\right) \leq \delta$.*

We do not directly leverage Ville's inequality in this work, but all inequalities we use can be directly proven from Lemma 4 (Howard et al., 2020; 2021). In short, each inequality in this supplement is proved by carefully massaging a martingale of interest into a non-negative supermartingale.

Another useful tool we will leverage is Doob's optional stopping theorem.

Lemma 5 (Optional stopping theorem (Durrett, 1996))**.** *Let $(X_n)_{n\in\mathbb{N}}$ be a nonnegative supermartingale with respect to some filtration $(\mathcal{F}_n)_{n\in\mathbb{N}}$. Then $\mathbb{E}[X_\tau] \leq \mathbb{E}[X_0]$ for all stopping times τ that are potentially infinite.*

For our alternative proof of the privacy filter (in Section D), we leverage the following special case of a recent advance in time-uniform martingale concentration (Howard et al., 2020). The following Theorem 4 is just a special case of the main result in Howard et al. (2020), and we include the proof for completeness. When we say a random variable X is σ^2-subGaussian conditioned on some sigma-algebra $\mathcal{G}$, we mean that, for all $\lambda \geq 0$,

$$\mathbb{E}\left(e^{\lambda X} \mid \mathcal{G}\right) \leq e^{\lambda^2\sigma^2/2}.$$

In particular, if X is σ^2-subGaussian as above, this does not imply that $-X$ is σ-subGaussian (because the condition is only assumed for $\lambda \geq 0$). In general, X can have different behaviors in its left and right tail, see for example the

discussion of the differing tails of the empirical variance of Gaussians in Howard et al. (2021).

Theorem 4. *Let $(M_n)_{n\in\mathbb{N}}$ be a martingale with respect to some filtration $(\mathcal{F}_n)_{n\in\mathbb{N}}$ such that $M_0 = 0$ almost surely. Moreover, let $(\sigma_n)_{n\geq 1}$ be a $(\mathcal{F}_n)_{n\in\mathbb{N}}$-predictable sequence of random variables such that, conditioned on $\mathcal{F}_{n-1}$, $\Delta M_n := M_n - M_{n-1}$ is σ_n^2-subGaussian. Define $V_n := \sum_{m\leq n}\sigma_m^2$. Then, we have, for all $a, b > 0$,*

$$\mathbb{P}\left(\exists n \in \mathbb{N} : M_n \geq \frac{b}{2} + \frac{b}{2a}V_n\right) \leq \exp\left(\frac{-b^2}{2a}\right).$$

Proof of Theorem 4. Let $(M_n)_{n\in\mathbb{N}}$ be the martingale listed in the theorem statement. Observe that, for any $a, b > 0$, the process $(X_n)_{n\in\mathbb{N}}$ given by

$$X_n := \exp\left(\frac{b}{a}M_n - \frac{b^2}{2a^2}\sum_{m\leq n}\sigma_m^2\right)$$

is a non-negative supermartingale. As such, applying Ville's inequality (Lemma 4) yields

$$\mathbb{P}\left(\exists n \in \mathbb{N} : X_n > \exp\left(\frac{b^2}{2a}\right)\right) \leq \exp\left(-\frac{b^2}{2a}\right).$$

Now, on such event, taking logs and rearranging yields

$$\frac{b}{a}M_n \leq \frac{b^2}{2a} + \frac{b^2}{2a^2}\sum_{m\leq n}\sigma_m^2.$$

Multiplying both sides by $\frac{a}{b}$ finishes the proof. □

The predictable process $(V_n)_{n\in\mathbb{N}}$ is a proxy for the accumulated variance of $(M_n)_{n\in\mathbb{N}}$ up to any fixed point in time. In particular, the process $(V_n)_{n\in\mathbb{N}}$ can be thought of as yielding the "intrinsic time" of the process. The free parameters a and b thus allow us to optimize the tightness of the boundary for some intrinsic moment in time. This is ideal for us, as, for the sake of composition, the target privacy parameter ϵ can guide us in finding a point in intrinsic time (that is, in terms of the process $(V_n)_{n\in\mathbb{N}}$) to optimize for. We discuss how to apply this inequality to prove privacy composition results both in this supplement and in Section 3.

We also leverage the following martingale inequalities from Howard et al. (2021) in Section 4, where we construct various families of time-uniform bounds on privacy loss in fully-adaptive composition. These inequalities take on a more complicated form than Theorem 4, but we explain the intuition behind them in the sequel. The first bound we present relies on the method of mixtures for martingale concentration, which stems back to Robbins' work in the 1970s (Robbins, 1970). There are many good resources providing an introduction to the method of mixtures (de la Peña et al., 2007; Kaufmann and Koolen, 2021; Howard et al., 2021).

Theorem 5. *Let $(M_n)_{n\in\mathbb{N}}$ be a martingale with respect to some filtration $(\mathcal{F}_n)_{n\in\mathbb{N}}$ such that $M_0 = 0$ almost surely. Moreover, let $(\sigma_n)_{n\geq 1}$ be a $(\mathcal{F}_n)_{n\in\mathbb{N}}$-predictable sequence of random variables such that, conditioned on $\mathcal{F}_{n-1}$, $\Delta M_n := M_n - M_{n-1}$ is σ_n^2-subGaussian. Define $V_n := \sum_{m\leq n}\sigma_m^2$ and choose a tuning parameter $\gamma > 0$. Then, for any $\delta > 0$, we have*

$$\mathbb{P}\left(\exists n \in \mathbb{N} : M_n \geq \sqrt{2(V_n+\gamma)\log\left(\frac{1}{\delta}\sqrt{\frac{V_n+\gamma}{\gamma}}\right)}\right) \leq \delta.$$

The next inequality relies on the recent technique of boundary stitching, first presented in Howard et al. (2021). Intuitively, the technique works by breaking intrinsic time — that is, time according to the accumulated variance process $(V_n)_{n\in\mathbb{N}}$ — into roughly geometrically spaced pieces. Then, one optimizes a tight-boundary in each region and takes a union bound. The actual details are more technical, but are not needed in this work.

Theorem 6. *Let $(M_n)_{n\in\mathbb{N}}$ be a martingale with respect to $(\mathcal{F}_n)_{n\in\mathbb{N}}$ such that $M_0 = 0$ almost surely. Moreover, let $(\sigma_n)_{n\geq 1}$ be a $(\mathcal{F}_n)_{n\in\mathbb{N}}$-predictable sequence of random variables such that, conditioned on $\mathcal{F}_{n-1}$, both $\Delta M_n := M_n - M_{n-1}$ and $-\Delta M_n$ are σ_n^2-subGaussian. Define $V_n := \sum_{m\leq n}\sigma_m^2$ and choose a starting intrinsic time $v_0 > 0$. Then, for any $\delta \in (0,1)$, we have*

$$\mathbb{P}\left(\exists n \in \mathbb{N} : M_n \geq 1.7\sqrt{V_n\left(\log\log\left(\frac{2V_n}{v_0}\right) + .72\log\left(\frac{5.2}{\delta}\right)\right)} \text{ and } V_n \geq v_0\right) \leq \delta.$$

Note that the original version of Theorem 6 as found in Howard et al. (2021) has more free parameters to optimize over, but we have already simplified the expression to make the result more readable. The free parameter $v_0 > 0$ in the above boundary gives the intrinsic time at which the boundary becomes non-trivial (i.e., the tightest available upper bound before $V_n \geq v_0$ is ∞).

We qualitatively compare these bounds in Section 4, wherein we construct various time-uniform bounds on privacy loss processes. For now, Theorem 4 can be thought of as providing a tight upper bound on a martingale at a single point in intrinsic time, providing loose guarantees elsewhere. On the other hand, Theorems 5 and 6 provide decently tight control over a martingale at all points in intrinsic time simultaneously, although at the cost of sacrificing tightness at any given fixed point.

C. Details in Proof of Approx-zCDP Filter

C.1. Equivalence of Approximate zCDP Definitions

We will show that our definition of approximate zCDP is equivalent to the original definition of approximate zCDP due to Bun and Steinke (2016). Let us first restate their definition as a condition on a private algorithm A.

Condition 1 (Original definition of Bun and Steinke (2016)). *For any neighboring datasets x, x', there exist events E and E' such that for all $\lambda \geq 1$,*

$$\begin{aligned}
&D_\lambda(A(x) \mid E \| A(x') \mid E') \leq \rho\lambda, \\
&D_\lambda(A(x') \mid E' \| A(x) \mid E) \leq \rho\lambda, \\
&\mathbb{P}(A(x) \in E) \geq 1 - \delta, \text{ and} \\
&\mathbb{P}(A(x') \in E') \geq 1 - \delta.
\end{aligned}$$

Our definition is adapted from the approximate Rényi differential privacy definition due to Papernot and Steinke (2022). We restate the (unconditional) definition below.

Condition 2 (Adapted from Papernot and Steinke (2022)). *For any neighboring datasets x, x', there exist distributions P', P'', Q', Q'' such that the outputs are distributed according to the following mixture distributions:*

$$A(x) \sim (1-\delta)P' + \delta P'', \qquad A(x') \sim (1-\delta)Q' + \delta Q''$$

with for all $\lambda \geq 1$, $D_\lambda(P' \| Q') \leq \rho\lambda$ and $D_\lambda(P' \| Q') \leq \rho\lambda$.

Theorem 7. *Conditions 1 and 2 are equivalent.*

Proof of Theorem 7. Fix any neighbors x, x'. Suppose an algorithm A satisfies Condition 1 for some events E, E'. Then we could let P' and Q' be the conditional distributions $\mathbb{P}(A(x) \in \cdot \mid A(x) \in E)$ and $\mathbb{P}(A(x') \in \cdot \mid A(x') \in E')$ respectively. Then let

$$\begin{aligned}
P''(\cdot) &= \frac{1}{\delta}\Big(\mathbb{P}(A(x) \in \cdot \mid A(x) \in E^c)\mathbb{P}(A(x) \in E^c) \\
&\quad + P'(\cdot)\left(\mathbb{P}(A(x) \in E) - (1-\delta)\right)\Big), \\
Q''(\cdot) &= \frac{1}{\delta}\Big(\mathbb{P}(A(x') \in \cdot \mid A(x') \in E'^c)\mathbb{P}(A(x') \in E'^c) \\
&\quad + Q'(\cdot)\left(\mathbb{P}(A(x') \in E') - (1-\delta)\right)\Big).
\end{aligned}$$

Then $A(x)$ is distributed according to the mixture $(1-\delta)P' + \delta P''$, and $A(x')$ is distributed according to the mixture $(1-\delta)Q' + \delta Q''$. Thus, A also satisfies condition 2 given that $D_\lambda(P' \| Q') \leq \lambda\rho$ and $D_\lambda(Q' \| P') \leq \lambda\rho$ by our assumption of Condition 1.

Now suppose A satisfies Condition 2 for some pairs of distributions (P', P'') and (Q', Q''). Then we can view the output distribution of $A(x)$ as generating a Bernoulli random variable C such that with probability $(1-\delta)$, $C = 1$ and $A(x)$ draws an outcome from P' and with probability $C = 0$ and $A(x)$ draws an outcome from P''. Similarly, we can view $A(x')$ as flipping a coin C' such that $A(x')$ draws an outcome from Q' when $C' = 1$. Then letting the events E be all the randomness of $A(x)$ such that $C = 1$ and E' be all the randomness of $A(x')$ such that $C' = 1$ satisfies condition 1. □

C.2. Missing Proofs

Lemma 6. *The process $\{X_n\}_{n\geq 1}$ defined in* (10) *is a P'-nonnegative supermartingale with respect to $(\mathcal{F}_n(x))_{n\in\mathbb{N}}$.*

Proof of Lemma 6. For any $t \geq 0$,

$$\begin{aligned}
&\mathbb{E}_{P'}[X_{t+1} \mid \mathcal{F}_t(x)] \\
&= \mathbb{E}_{P'}\Bigg[X_t \exp\Bigg((\lambda-1)\log\left(\frac{P'_{t+1}(A_{t+1}(x) \mid \mathcal{F}_t(x))}{Q'_{t+1}(A_{t+1}(x) \mid \mathcal{F}_t(x))}\right) \\
&\qquad - \lambda(\lambda-1)\rho_{t+1}(x)\Bigg) \mid \mathcal{F}_t(x)\Bigg] \\
&= X_t\, \mathbb{E}_{P'}\left[\left(\frac{P'_{t+1}(A_{t+1}(x) \mid \mathcal{F}_t(x))}{Q'_{t+1}(A_{t+1}(x) \mid \mathcal{F}_t(x))}\right)^{(\lambda-1)} \mid \mathcal{F}_t(x)\right] \\
&\qquad \cdot \exp(-\lambda(\lambda-1)\rho_{t+1}(x)) \\
&\leq X_t \exp(\lambda(\lambda-1)\rho_{t+1}(x)) \exp(-\lambda(\lambda-1)\rho_{t+1}(x)) \\
&= X_t,
\end{aligned}$$

where the last inequality follows from the Renyi divergence bound due to approximate zCDP. □

Lemma 7. *Consider measures P' and Q' defined in* (8). *Their Rényi divergence satisfies*

$$D_\lambda\left(P'(A_{1:N(x)}(x)) \| Q'(A_{1:N(x)}(x))\right) \leq \rho\lambda.$$

Proof of Lemma 7. By the definition of X_n and that $\mathbb{E}_{P'}[X_{N(x)}] \leq \mathbb{E}_{P'}[X_0] = 1$, we have

$$\begin{aligned}
&\mathbb{E}_{A_{1:N(x)}(x)\sim P'}\left[\exp\left((\lambda-1)M_{N(x)}\right)\right] \leq 1 \iff \\
&\mathbb{E}_{P'}\Bigg[\exp\Bigg((\lambda-1)\sum_{m\leq N(x)}\Bigg\{\log\left(\frac{P'_m(A_m(x) \mid \mathcal{F}_{m-1}(x))}{Q'_m(A_m(x) \mid \mathcal{F}_{m-1}(x))}\right) \\
&\qquad - \lambda\rho_m(x)\Bigg\}\Bigg)\Bigg] \leq 1 \iff \\
&\mathbb{E}_{P'}\Bigg[\left(\frac{P'(A_{1:N(x)})}{Q'(A_{1:N(x)})}\right)^{\lambda-1} \\
&\qquad \cdot \exp\left(-(\lambda-1)\lambda\sum_{m\leq N(x)}\rho_m(x)\right)\Bigg] \leq 1.
\end{aligned}$$

By the definition of stopping time N, we have $\sum_{m\leq N(x)}\rho_m(x) \leq \rho$, which implies the stated Renyi divergence bound. □

Lemma 8. *Let likelihood functions P, Q, P', Q' be defined in* (3)*,* (4)*, and* (8)*. Then there exists likelihood functions P'' and Q'' such that*

$$P = (1-\delta)P' + \delta P'',$$
$$Q = (1-\delta)Q' + \delta Q''.$$

Proof of Lemma 8. We will show the decomposition for P, and the proof follows identically for the decomposition of Q. First, we can express the likelihood $P(A_{1:N(x)}(x))$ as follows:

$$\begin{aligned} P(A_{1:N(x)}(x)) &= \prod_{n=1}^{N(x)} P(A_n(x) \mid \mathcal{F}_{n-1}(x)) \\ &= \prod_{n=1}^{N(x)} \big[(1-\delta_n(x))P'_n(A_n(x) \mid \mathcal{F}_{n-1}(x)) \\ &\qquad + \delta_n(x)P''_n(A_n(x) \mid \mathcal{F}_{n-1}(x))\big] \\ &= \sum_{S \subseteq \{1,\ldots,N(x)\}} w_S \cdot f_S(A_{1:N(x)}(x)) \end{aligned}$$

where

$$f_S(A_{1:N(x)}(x)) := \prod_{n\in S} P''_n(A_n(x) \mid \mathcal{F}_{n-1}(x)) \prod_{n \le N(x), n\notin S} P'_n(A_n(x) \mid \mathcal{F}_{n-1}(x))$$

and $w_S = \left(\prod_{n\in S}\delta_n(x)\prod_{n\in\mathbb{N}\setminus S}(1-\delta_n(x))\right)$. Note that each f_S is a likelihood of the stopped process $A_{1:N(x)}(x)$ under input data set x, and $f_\emptyset = P'(A_{1:N(x)}(x))$. Thus, it suffices to show that $w_\emptyset \ge 1-\delta$ almost surely. To see this, we have

$$w_\emptyset = \prod_{n\le N(x)}(1-\delta_n(x)) \ge 1 - \sum_{n\le N(x)}\delta_n(x) \ge 1-\delta.$$

□

D. An Alternative Proof for Theorem 2

As a first step in our alternative proof of Theorem 2, it is easier to consider the case where each algorithm A_n satisfies conditional (ϵ_n, δ_n)-pDP, as this condition provides a high-probability bound on the privacy loss. This allows us to use the martingale machinery in Appendix B to prove tight composition results.

Lemma 9. *Theorem 2 holds under the stronger assumption that, for any $n \ge 1$, A_n is (ϵ_n, δ_n)-pDP conditioned on $A_{1:n-1}$.*

Before we can prove Lemma 9, we need to following bound on the conditional expectation of privacy loss, which can be immediately obtained from the bound on expected privacy loss presented in Bun and Steinke (2016).

Lemma 10 (Proposition 3.3 in Bun and Steinke (2016))**.** *Suppose A and B are algorithms such that A is ϵ-differentially private conditioned on B. Then, for any input dataset $x \in \mathcal{X}$ and neighboring dataset $x' \sim x$, we have that*

$$\mathbb{E}\left(\mathcal{L}(x,x')|B(x)\right) \le \frac{1}{2}\left(\epsilon(B(x))\right)^2.$$

Now, we prove Lemma 9.

Proof of Lemma 9. To begin, we assume that the algorithms $(A_n)_{n\ge1}$ satisfy $(\epsilon_n, 0)$-pDP conditioned on $A_{1:n-1}$. We will show how to alleviate this assumption on the approximation parameter in the second half of the proof. Fix an input database $x \in \mathcal{X}$. For convenience, we denote by $(\mathcal{F}_n(x))_{n\in\mathbb{N}}$ the natural filtration generated by $(A_n(x))_{n\ge1}$. Since we have fixed $x \in \mathcal{X}$, for notational simplicity, we write ϵ_n for the random variable $\epsilon_n(A_{1:n-1}(x))$ and define δ_n similarly. Additionally, by N we mean the stopping time $N((\epsilon_n)_{n\in\mathbb{N}}, (\delta_n)_{n\in\mathbb{N}})$. Recall that we have already argued that, for any neighboring dataset $x' \sim x$, the process

$$M_n := M_n(x,x') = \mathcal{L}_{1:n}(x,x') - \sum_{m\le n}\mathbb{E}\left(\mathcal{L}_m(x,x')|\mathcal{F}_{m-1}(x)\right)$$

is a martingale with respect to $(\mathcal{F}_n(x))_{n\in\mathbb{N}}$. Further observe that its increments $\Delta M_n := \mathcal{L}_n(x,x') - \mathbb{E}\left(\mathcal{L}_n(x,x')|\mathcal{F}_{n-1}(x)\right)$ are ϵ_n^2-subGaussian conditioned on $\mathcal{F}_{n-1}(x)$.

Thus, by Theorem 4, we know that, for any $b, a > 0$, we have

$$\mathbb{P}\left(\exists n \in \mathbb{N} : M_n \ge \frac{b}{2} + \frac{b}{2a}V_n\right) \le \exp\left(\frac{-b^2}{2a}\right),$$

where the process $(V_n)_{n\in\mathbb{N}}$ given by $V_n := \sum_{m\le n}\epsilon_m^2$ is the accumulated variance up to and including time n. Thus, it suffices to optimize the free parameters a and b to prove the result.

To do this, consider the following function $f : \mathbb{R}_{\ge0} \to \mathbb{R}_{\ge0}$ given by

$$f(y) = \sqrt{2\log\left(\frac{1}{\delta'}\right)y} + \frac{1}{2}y.$$

Clearly, f is a quadratic polynomial in $\sqrt{y}$ which is strictly increasing. In particular, one can readily check that

$$y^* := \left(-\sqrt{2\log\left(\frac{1}{\delta'}\right)} + \sqrt{2\log\left(\frac{1}{\delta'}\right) + \epsilon}\right)^2 \tag{15}$$

solves the equation $f(y) = \epsilon$, where $\epsilon > 0$ is the target privacy parameter.

As such, setting $a := y^*$ and $b := \sqrt{2\log\left(\frac{1}{\delta'}\right)y^*}$ yields

$$\exp\left(\frac{-b^2}{a}\right) = \exp\left(\frac{-2y^*\log\left(\frac{1}{\delta'}\right)}{y^*}\right) = \delta'.$$

Furthermore, expanding the definition of $(M_n)_{n\in\mathbb{N}}$, we see that for the selected parameters the parameters yield, with probability at least $1-\delta'$, for all $n \le N$ we have:

$$\begin{aligned}
&\mathcal{L}_{1:n}(x,x') \le \frac{b}{2} + \frac{b}{2a}V_n + \sum_{m\le n} \mathbb{E}\left(\mathcal{L}_m(x,x') \mid \mathcal{F}_{m-1}\right) \\
&\le \frac{b}{2} + \frac{b}{2a}\sum_{m\le n}\epsilon_m^2 + \frac{1}{2}\sum_{m\le n}\epsilon_m^2 \\
&= \frac{1}{2}\sqrt{2\log\left(\frac{1}{\delta'}\right)y^*} + \frac{1}{2}\frac{\sqrt{2\log\left(\frac{1}{\delta'}\right)y^*}}{y^*}\sum_{m\le n}\epsilon_m^2 + \frac{1}{2}\sum_{m\le n}\epsilon_m^2 \\
&\le \frac{1}{2}\sqrt{2\log\left(\frac{1}{\delta'}\right)y^*} + \frac{1}{2}\sqrt{2\log\left(\frac{1}{\delta'}\right)y^*} + \frac{1}{2}\sum_{m\le n}\epsilon_m^2 \\
&= \sqrt{2\log\left(\frac{1}{\delta'}\right)y^*} + \frac{1}{2}\sum_{m\le n}\epsilon_m^2 \le \sqrt{2\log\left(\frac{1}{\delta'}\right)y^*} + \frac{1}{2}y^* = \epsilon.
\end{aligned}$$

Thus, we have proven the desired result in the case where all algorithms have $\delta_n = 0$.

Now, we show how to generalize our result to the case where the approximation parameters δ_n are not identically zero. Define the events

$$\begin{aligned}
A &:= \{\exists n \le N : \mathcal{L}_{1:n}(x,x') > \epsilon\}, \text{ and} \\
B &:= \{\exists n \le N : \mathcal{L}_n(x,x') > \epsilon_n\}.
\end{aligned}$$

Our goal is to show that, with N defined as in the statement of Theorem 2, that $\mathbb{P}(A) \le \delta$. Simply using Bayes rule, we have that

$$\mathbb{P}(A) = \mathbb{P}(A\cap B^c) + \mathbb{P}(A\cap B) \le \mathbb{P}(A|B^c) + \mathbb{P}(B) \le \delta' + \mathbb{P}(B),$$

where the second inequality follows from our already-completed analysis in the case that $\delta_n = 0$. Now, we show that $\mathbb{P}(B) \le \delta''$, which suffices to prove the result as we have, by assumption, $\delta = \delta' + \delta''$.

Define the modified privacy loss random variables $(\widetilde{\mathcal{L}}_n(x,x'))_{n\in\mathbb{N}}$ by

$$\widetilde{\mathcal{L}}_n(x,x') := \begin{cases} \mathcal{L}_n(x,x') & n \le N \\ 0 & \text{otherwise} \end{cases}.$$

Likewise, define the modified privacy parameter random variables $\widetilde{\epsilon}_n$ and $\widetilde{\delta}_n$ in an identical manner. Then, we can bound $\mathbb{P}(B)$ in the following manner:

$$\begin{aligned}
&\mathbb{P}(\exists n \le N : \mathcal{L}_n(x,x') > \epsilon_n) = \mathbb{P}\left(\exists n \in \mathbb{N} : \widetilde{\mathcal{L}}_n(x,x') > \widetilde{\epsilon}_n\right) \\
&\le \sum_{n=1}^{\infty}\mathbb{P}\left(\widetilde{\mathcal{L}}_n(x,x') > \widetilde{\epsilon}_n\right) = \sum_{n=1}^{\infty}\mathbb{E}\mathbb{P}\left(\widetilde{\mathcal{L}}_n(x,x') > \widetilde{\epsilon}_n | \mathcal{F}_{n-1}\right) \\
&\le \sum_{n=1}^{\infty}\mathbb{E}\widetilde{\delta}_n = \mathbb{E}\left[\sum_{n=1}^{\infty}\widetilde{\delta}_n\right] = \mathbb{E}\left[\sum_{n\le N}\delta_n\right] \le \delta''.
\end{aligned}$$

Thus, we have have proven the desired result in the general case. □

Our key insight above is to view filters as functions of the "intrinsic time" determined by privacy parameters, $\sum_{m\le n}\epsilon_m^2$. Lemma 9 can also be obtained leveraging the analysis for Rényi filters (Feldman and Zrnic, 2021). However, our approach to proving Theorem 2 has the advantage that it does not require reductions between different modes of privacy. While Lemma 10, which bounds expected privacy loss, does require some complicated analysis, we only ever need to apply Lemma 9 to instances of randomized response, in which case computing the privacy loss bound is trivial.

We now use Lemma 9 to prove Theorem 2. Recall that Lemma 2 shows that algorithms that satisfy pDP also satisfy DP, but the converse is not true and may require a conversion cost. To avoid this cost, we define following generalization of randomized response.

Definition 7 (Conditional Randomized Response)**.** *Let $\mathcal{R} := \{0, 1, \top, \bot\}$ and $2^{\mathcal{R}}$ be the corresponding power set of $\mathcal{R}$. Then, R taking inputs in $\{0,1\}$ to outputs in the measurable space $(\mathcal{R}, 2^{\mathcal{R}})$ is an instance of (ϵ,δ)-randomized response if, for $b \in \{0,1\}$, $R(b)$ outputs the following:*

$$R(b) = \begin{cases} b & \text{with probability } (1-\delta)\frac{e^\epsilon}{1+e^\epsilon} \\ 1-b & \text{with probability } (1-\delta)\frac{1}{1+e^\epsilon} \\ \top & \text{with probability } \delta \text{ if } b = 1 \\ \bot & \text{with probability } \delta \text{ if } b = 0. \end{cases}$$

More generally, suppose $B : \{0,1\} \to \mathcal{Z}$ is a randomized algorithm. For functions $\epsilon, \delta : \mathcal{Z} \to \mathbb{R}_{\ge 0}$, we say R is an instance of (ϵ,δ)-randomized response conditioned on B if, for any true input $b' \in \{0,1\}$ and hypothesized alternative $b \in \{0,1\}$, the conditional probability $\mathbb{P}(R(b) \in \cdot | B(b') = z)$ is the same as the law of $(\epsilon(z), \delta(z))$-randomized response with input bit b.

Conditional (ϵ,δ)-randomized response satisfies both conditional (ϵ,δ)-DP and conditional (ϵ,δ)-pDP. We will leverage the fact that it satisfies both privacy definitions with the same parameters. A surprising result in the nonadaptive setting is that *any* (ϵ,δ)-DP algorithm can be viewed as a randomized post-processing of (ϵ,δ)-randomized response (Kairouz et al., 2015). We generalize this result to the adaptive conditional setting below. In the language of Blackwell's comparison of experiments (Blackwell, 1953), instances of randomized response are "sufficient" for instances of arbitrary DP algorithms, and we prove that the same is true for conditional randomized response and conditionally DP algorithms. In what follows, by a transition kernel ν, we mean that for any $b \in \mathcal{Z}$ and $r \in \mathcal{R}$, $\nu(\cdot, r \mid b)$ is a probability measure on $(\mathcal{Y}, \mathcal{G})$.

Lemma 11 (Reduction to Conditional Randomized Response)**.** *Let A and B map from $\mathcal{X}$ to measurable spaces $(\mathcal{Y},\mathcal{G})$ and $(\mathcal{Z},\mathcal{H})$, respectively. Suppose A is (ϵ,δ)-differentially private conditioned on B. Fix neighbors*

$x_0, x_1 \in \mathcal{X}$, and let R be an instance of (ϵ, δ)-randomized response conditioned on B', where $B' : \{0,1\} \to \mathcal{Z}$ is the restricted algorithm satisfying $B'(b) = B(x_b)$. Then, there is a transition kernel $\nu : \mathcal{G} \times \mathcal{R} \times \mathcal{Z} \to [0,1]$ such that, for all $b, b' \in \{0,1\}$, $\mathbb{P}\left(A(x_b) \in \cdot \mid B'(b')\right) = \nu_{b,b'}$, where $\nu_{b,b'} = \mathbb{E}\left(\nu(\cdot, R(b) \mid B'(b')) \mid B'(b')\right)$.[6]

Lemma 11 tells us that the conditional distribution obtained by averaging the kernel $\nu(\cdot, R(b) \mid B'(b'))$ over the randomness in $R(b)$ matches the conditional distribution of $A(x_b)$. To prove Lemma 11, first recall the important fact that *any* differentially private algorithm can be viewed as a post-processing of randomized response (Kairouz et al., 2015), as stated in Lemma 12 below.

Lemma 12 (Reduction to Randomized Response (Kairouz et al., 2015))**.** *Let algorithm $A : \mathcal{X} \to \mathcal{Y}$ be (ϵ, δ)-DP. Let R be an instance of (ϵ, δ)-randomized response. Then, for any neighbors $x_0, x_1 \in \mathcal{X}$, there is a transition kernel $\nu : \mathcal{G} \times \mathcal{R} \to [0,1]$ such that for $b \in \{0,1\}$, we have $\mathbb{P}(A(x_b) \in \cdot) = \nu_b$, where*[7] *$\nu_b = \mathbb{E}\nu(\cdot, R(b))$.*

In Lemma 11 of Section 3, we generalized Lemma 12 to the case of conditional differential privacy. To do this, we introduced *conditional randomized response* in Definition 7. In conditional randomized response, on the event $\{B = z\}$, the conditional laws of $R(0)$ and $R(1)$ just become that of regular randomized response with some known privacy parameters $\epsilon(z)$ and $\delta(z)$. We now prove Lemma 11.

Proof of Lemma 11. Let $b, b' \in \{0,1\}$ be arbitrary. For any outcome $\{B'(b') = z\}$, let $\mathbb{P}_z(A(x_b) \in \cdot)$ be the probability measure $\mathbb{P}(A(x_b) \in \cdot | B'(b') = z)$. In particular, this measure does not depend on the input bit b'. By the assumptions of conditional differential privacy (Definition 1), it follows that under the probability measure $\mathbb{P}_z$, $A(x_b)$ is $(\epsilon(z), \delta(z))$-differentially private. Moreover, it also follows that R is an instance of $(\epsilon(z), \delta(z))$-randomized response under $\mathbb{P}_z$. Consequently, Lemma 12 yields the existence of a kernel ν_z such that $\mathbb{P}_z(A(x_b) \in \cdot) = \mathbb{E}_z \nu_z(\cdot, R(b))$, where the averaged measure is as defined in Footnote 7. Setting $\nu(\cdot, R(b)|z) := \nu_z(\cdot, R(b))$, we see that

$$\mathbb{P}(A(x_b) \in \cdot \mid B'(b') = z) = \mathbb{E}\left(\nu(\cdot, R(b) \mid z) \mid B'(b') = z\right),$$

which thus yields

$$\mathbb{P}(A(x_b) \in \cdot \mid B'(b')) = \mathbb{E}\left(\nu(\cdot, R(b) \mid B'(b')) \mid B'(b')\right),$$

where the conditionally averaged measure is as described in Footnote 6 in the main body of the paper. This proves the desired result. □

Lastly, before proving Theorem 2, we need the following lemma. This lemma essentially tells us that if A is (ϵ, δ)-pDP conditioned on B, and A' is a randomized post-processing algorithm, then releasing the vector (A, A') is also (ϵ, δ)-pDP conditioned on B. Note that this is *not* in contradiction with the converse direction of Lemma 2, as releasing the output of A' alone may not satisfy conditional (ϵ, δ)-pDP. But once we observe A, since A' is a post-processing, we can gleam no more information about the true underlying dataset.

Lemma 13. *Suppose A, B are algorithms with inputs in $\mathcal{X}$ and outputs in measurable spaces $(\mathcal{Y}, \mathcal{G})$ and $(\mathcal{Z}, \mathcal{H})$ respectively. Assume A is (ϵ, δ)-pDP conditioned on B. Let $(S, \mathcal{S})$ be a measurable space and suppose $\mu : \mathcal{S} \times \mathcal{Y} \times \mathcal{Z} \to [0,1]$ is a conditional transition kernel. Suppose $A' : \mathcal{X} \to S$ is an algorithm satisfying*

$$\mathbb{P}\left(A'(x) \in \cdot | A(x') = y, B(x') = z\right) = \mu(\cdot, y \mid z), \quad (16)$$

for all $y \in \mathcal{Y}, z \in \mathcal{Z}$, and $x, x' \in \mathcal{X}$. Then, the joint algorithm $(A, A') : \mathcal{X} \to \mathcal{Y} \times S$ is also (ϵ, δ)-pDP conditioned on B.

Proof of Lemma 13. Let $x, x' \in \mathcal{X}$ be arbitrary neighboring datasets. Let $q_B^x, q_B^{x'}$ be the corresponding conditional joint densities of $(A(x), A'(x))$ and $(A(x'), A'(x'))$ given $B(x)$ respectively. Likewise, let $p_B^x, p_B^{x'}$ be the corresponding conditional densities of $A(x)$ and $A(x')$ respectively conditioned on $B(x)$, and $q_{B,A}^x, q_{B,A}^{x'}$ the conditional densities of $A'(x)$ and $A'(x')$ given $A(x)$ and $B(x)$. Let $\mathcal{L}_B^{(A,A')}(x, x')$ denote the joint privacy loss between $(A(x), A'(x))$ and $(A(x'), A'(x'))$ given $B(x)$, while $\mathcal{L}_B^A(x, x')$ denotes the privacy loss between $A(x)$ and $A(x')$ given $B(x)$. We have, using Bayes rule,

$$\begin{aligned}
\mathcal{L}_B^{(A,A')}(x, x') &= \log\left(\frac{q_B^x(A(x), A'(x) \mid B(x))}{q_B^{x'}(A(x), A'(x) \mid B(x))}\right) \\
&= \log\left(\frac{p_B^x(A(x) \mid B(x))}{p_B^{x'}(A(x) \mid B(x))} \cdot \frac{q_{B,A}^x(A'(x) \mid B(x), A(x))}{q_{B,A}^{x'}(A'(x) \mid B(x), A(x))}\right) \\
&= \log\left(\frac{p_B^x(A(x) \mid B(x))}{p_B^{x'}(A(x) \mid B(x))}\right) = \mathcal{L}_B^{(A)}(x, x'),
\end{aligned}$$

[6] By $\nu_{b,b'}(\cdot) := \mathbb{E}\left(\nu(\cdot, R(b) \mid B'(b')) \mid B'(b')\right)$, we mean that $\nu_{b,b'}$ is the (random) averaged probability measure:

$$\begin{aligned}
\nu_{b,b'}(\cdot) = &\ \mathbb{P}(R(b) = 1 \mid B'(b'))\nu(\cdot, 1 \mid B'(b')) \\
&+ \mathbb{P}(R(b) = 0 \mid B'(b'))\nu(\cdot, 0 \mid B'(b')) \\
&+ \mathbb{P}(R(b) = \perp \mid B'(b'))\nu(\cdot, \perp \mid B'(b')) \\
&+ \mathbb{P}(R(b) = \top \mid B'(b'))\nu(\cdot, \top \mid B'(b')).
\end{aligned}$$

[7] By $\nu_b(\cdot) := \mathbb{E}\nu(\cdot, R(b))$, we mean ν_b is the averaged probability measure given by

$$\begin{aligned}
\nu_b(\cdot) = &\ \mathbb{P}(R(b) = 1)\nu(\cdot, 1) + \mathbb{P}(R(b) = 0)\nu(\cdot, 0) \\
&+ \mathbb{P}(R(b) = \perp)\nu(\cdot, \perp) + \mathbb{P}(R(b) = \top)\nu(\cdot, \top).
\end{aligned}$$

The first equality on the second line follows from the assumption outlined in Equation (16). More specifically, since we have

$$\mathbb{P}\left(A'(x) \in \cdot | A(x), B(x)\right) = \mu(\cdot, A(x) \mid B(x)) = \mathbb{P}\left(A'(x') \in \cdot | A(x), B(x)\right),$$

it follows that the conditional densities $q^{x}_{B,A}$ and $q^{x'}_{B,A}$ are equal almost surely. Since A is (ϵ, δ)-pDP conditioned on B, the result now follows. □

We now can prove Theorem 2 using these tools.

Proof of Theorem 2. Fix arbitrary neighbors $x_0, x_1 \in \mathcal{X}$. Let $(R_n)_{n\geq 1}$ be a sequence of algorithms such that R_n is an instance of (ϵ_n, δ_n)-randomized response conditioned on $A'_{1:n-1} : \{0,1\} \to \mathcal{Y}^{n-1}$, where $A'_m : \{0,1\} \to \mathcal{Y}$ is the restricted algorithm given by $A'_m(b) := A_m(x_b)$, for all $m \geq 1$. Lemma 11 guarantees the existence of a sequence of transition kernels $(\nu_n)_{n\geq 1}$, $\nu_n : \mathcal{G} \times \mathcal{R} \times \mathcal{Y}^{n-1} \to [0,1]$ such that, for all $n \geq 1$ and $b, b' \in \{0,1\}$, we have $\mathbb{P}(A'_n(b) \in \cdot \mid A'_{1:n-1}(b')) = \nu^{(n)}_{b,b'}$ almost surely. Here, $\nu^{(n)}_{b,b'}$ is the averaged conditional probability, as defined in terms of ν_n in Lemma 11 and Footnote 6. This equality means we can find an underlying probability space (i.e. a coupling) such that the random post-processing draws from the kernel $\nu_n(\cdot, R_n(b) \mid A'_{1:n-1}(b'))$ equal $A'_n(b)$ almost surely, for all $n \geq 1$.

Now, for any $n \geq 1$, since R_n is an instance of (ϵ_n, δ_n)-randomized response conditioned on $A'_{1:n-1}$, it follows that R_n is in fact (ϵ_n, δ_n)-pDP conditioned on $A'_{1:n-1}$. Moreover, this also implies that R_n is (ϵ_n, δ_n)-pDP conditioned on $(A'_{1:n-1}, R_{1:n-1})$, since, by definition, ϵ_n and δ_n only depend on the realizations of $R_{1:n-1}$ through the outputs of $A'_{1:n-1}$. By Lemma 13, it follows that for all $n \geq 1$, the algorithm (R_n, A'_n) is (ϵ_n, δ_n)-pDP conditioned on $(R_{1:n-1}, A'_{1:n-1})$. Thus, by Lemma 9, it follows that the composed algorithm $(R_{1:N'(\cdot)}(\cdot), A'_{1:N'(\cdot)}(\cdot))$ is (ϵ, δ)-DP, where $N'(b) := N(x_b)$ and ϵ, δ and N, are as outlined in the statement of Theorem 2.

Lastly, since differential privacy is closed under arbitrary post-processing (Dwork and Roth, 2014), it follows that $A'_{1:N'(\cdot)}(\cdot)$ is (ϵ, δ)-differentially private. Since x_0 and x_1 were arbitrary neighboring inputs, the result follows, i.e. $A_{1:N(\cdot)}(\cdot) : \mathcal{X} \to \mathcal{Y}^{\infty}$ is (ϵ, δ)-differentially private. □

E. Proof for Privacy Odometers in Theorem 3

We now show the formal proof for our privacy odometers presented in Theorem 3 in Section 4.

Theorem 3. As in the proof of Lemma 9, we first consider the case where $\delta_n = 0$ for all $n \geq 1$. In this case, fix an input dataset $x \in \mathcal{X}$ and a neighboring dataset $x' \in \mathcal{X}$. Let $(M_n)_{n\in\mathbb{N}}$ be the corresponding privacy loss martingale as outlined in Equation (14), where we implicitly hide the dependence on x, x', which are fixed. Let $(u_n)_{n\geq 1}$ be one of the sequences outlined in the theorem statement, and define $U_n := u_n(\epsilon_{1:n}, \delta_{1:n})$ for all $n \geq 1$, where once again we write ϵ_n and δ_n for $\epsilon_n(A_{1:n-1}(x))$ and $\delta_n(A_{1:n-1}(x))$ respectively. It follows from Theorems 4, 5, and 6 that

$$\mathbb{P}\left(\exists n \in \mathbb{N} : M_n > B_n\right) \leq \delta,$$

for $B_n = U_n - \frac{1}{2}\sum_{m\leq n} \epsilon_m^2$. Recalling that $M_n = \sum_{m\leq n}\{\mathcal{L}_m(x,x') - \mathbb{E}(\mathcal{L}_m(x,x')|\mathcal{F}_{n-1}(x))\}$ and that $\mathbb{E}(\mathcal{L}_n(x,x')|\mathcal{F}_{n-1}(x)) \leq \frac{1}{2}\epsilon_n^2$ for all $n \in \mathbb{N}$, it thus follows that

$$\mathbb{P}\left(\exists n \in \mathbb{N} : \mathcal{L}_{1:n}(x,x') > U_n\right) \leq \delta,$$

where $(\mathcal{F}_n(x))_{n\geq 1}$ is again the natural filtration generated by $(A_n(x))_{n\geq 1}$. Thus, since $x \sim x'$ were arbitrary, we have shown that $(u_n)_{n\geq 1}$ is a δ-privacy odometer in the case $\delta_n = 0$ for all $n \geq 1$.

To generalize to the case where δ_n may be nonzero, we can apply precisely the same argument used in the second part of the proof of Lemma 9, thus proving the general result. □

F. An Algorithm Satisfying (ϵ, δ)-DP but not (ϵ, δ)-pDP

In this appendix, we construct a simple algorithm taking binary inputs that satisfies (ϵ, δ)-DP but not (ϵ, δ)-pDP. In particular, this provides intuition as to why we conjecture our odometers constructed in Section 4 would not hold under the assumption that the algorithms being composed satisfy (ϵ, δ)-DP in general.

To this end, fix a privacy parameter $\epsilon > 0$ and an approximation parameter $\delta \in (0,1)$. Let $A : \{0,1\} \to \{0,1,\top,\bot\}$ be an instance of (ϵ, δ)-randomized response, and let $B : \{0,1\} \to \{0,1\}$ be defined by

$$B(b) := \begin{cases} 1 & \text{if } A(b) \in \{1, \top\}, \\ 0 & \text{otherwise.} \end{cases}$$

Since differential privacy is closed under arbitrary post-processing, it follows that the constructed algorithm B is (ϵ, δ)-differentially private. On the other hand, setting $x = 1$,

$x' = 0$, we note that on the event $\{B(1) = 1\}$,

$$\begin{aligned}
\mathcal{L}_B(1,0) &= \log\left(\frac{\mathbb{P}(B(1)=1)}{\mathbb{P}(B(0)=1)}\right) \\
&= \log\left(\frac{\mathbb{P}(A(1)=1)+\mathbb{P}(A(1)=\top)}{\mathbb{P}(A(0)=1)+\mathbb{P}(A(0)=\top)}\right) \\
&= \log\left(\frac{\delta+(1-\delta)\frac{e^\epsilon}{1+e^\epsilon}}{(1-\delta)\frac{1}{1+e^\epsilon}}\right) \\
&= \log\left(\frac{\delta+e^\epsilon}{1-\delta}\right) > \epsilon.
\end{aligned}$$

Since straightforward calculation yields

$$\mathbb{P}(B(1)=1) = (1-\delta)\frac{e^\epsilon}{1+e^\epsilon} + \delta > \delta,$$

we see that B does not satisfy (ϵ, δ)-pDP.