Optimal Differential Privacy Composition for Exponential Mechanisms and the Cost of Adaptivity

Jinshuo Dong*[1], David Durfee[2], and Ryan Rogers[2]

[1]Applied Mathematics and Computational Sciences, University of Pennsylvania
[2]Applied Research, LinkedIn

June 26, 2020

Abstract

Composition is one of the most important properties of differential privacy (DP), as it allows algorithm designers to build complex private algorithms from DP primitives. We consider precise composition bounds of the overall privacy loss for exponential mechanisms, one of the most fundamental class of mechanisms in DP. We give explicit formulations of the optimal privacy loss for both the adaptive and nonadaptive settings. For the nonadaptive setting in which each mechanism has the same privacy parameter, we give an efficiently computable formulation of the optimal privacy loss. Furthermore, we show that there is a difference in the privacy loss when the exponential mechanism is chosen adaptively versus nonadaptively. To our knowledge, it was previously unknown whether such a gap existed for any DP mechanisms with fixed privacy parameters, and we demonstrate the gap for a widely used class of mechanism in a natural setting. We then improve upon the best previously known upper bounds for adaptive composition of exponential mechanism with efficiently computable formulations and show the improvement.

*Work done while at LinkedIn

Contents

1 Introduction

Differential privacy (DP) has emerged as the leading privacy benchmark in machine learning as well as data analytics on sensitive data sets. One of the most useful properties of DP is that it composes, with slight degradation in the overall privacy loss parameters. This allows algorithm designers to build complicated algorithms whose privacy analysis follows from the fact that each subroutine satisfies DP. Further, composition allows us to bound the amount of privacy loss, quantified by the (ε, δ) parameters in DP, consumed by an (adaptive) sequence of DP algorithms evaluated on the same dataset. Hence, there have been several works in DP that help bound the privacy loss in composition, starting with basic composition from Dwork et al. [8] and advanced composition from Dwork et al. [9]. More recently, there have been works that give the exact, optimal privacy loss bound when all that is known is that the individual algorithms are each DP: Kairouz et al. [11] give the optimal privacy loss bound in the homogeneous case, where all the privacy parameters for each algorithm are the same, and Murtagh and Vadhan [14] give the more general optimal privacy loss bound in the heterogeneous case, where all the privacy parameters can be different at each round.

Although these *black box* composition theorems give the best possible bound on the privacy loss over multiple rounds of general DP algorithms, one should be able to improve on this bound when considering specific subclasses of DP algorithms. One example of such a composition theorem that takes into account the particular algorithm being used at each round is in *moments accounting composition* from Abadi et al. [1]. For their setting, they use noisy stochastic gradient descent and account for the subsampling and Gaussian noise that is added to the gradients at each round in their overall privacy loss bound. In particular, they are able to save a factor of $O(\sqrt{\ln(k/\delta)})$ in the overall privacy parameter, where k is the number of gradient descent steps taken. Another example of *white box* composition is from Durfee and Rogers [6] who introduce *bounded range* (BR) as a property for DP algorithms, which leads to improved composition bounds compared to the general case optimal bound.

Arguably, the fundamental DP primitives are the following: randomized response [19], Laplace mechanism [8] or its discretized variant (geometric mechanism), Gaussian mechanism [7], and the exponential mechanism [13]. The optimal DP composition bounds [11, 14] follow by showing that each DP algorithm, once the neighboring datasets are fixed, can be written as randomized response composed with a post-processing function that is independent of the data. Hence, the optimal DP composition is essentially tailored to composing randomized response mechanisms. The geometric mechanism was shown to also achieve the optimal composition bound [11]. Optimal DP composition bound for Gaussian mechanisms is obtained as a special case of the general composition theorem in Dong et al. [4].

Hence, it is only natural to then ask: **what is the optimal DP composition bound over the class of exponential mechanisms?** This question is the primary focus of this work. As was shown in Durfee and Rogers [6], the exponential mechanism satisfies the BR property and hence enjoys their improved composition bound. The exponential mechanism provides a very general way to design DP algorithms over an arbitrary outcome space where a *quality score* measures the value of each possible outcome given the input datatset. In practice, the exponential mechanism is most often deployed when a maximum or minimum operation is needed in a DP algorithm.

Surprisingly, the answer to this question depends on whether the choice of exponential mechanism is adaptively chosen at each round or not. For the existing DP composition bounds, adaptivity in the choice of DP algorithm did not affect the overall privacy bound, even in the optimal privacy loss bounds. Rogers et al. [17] show that there is an asymptotic gap in the privacy loss bound when

the privacy parameters $\{\varepsilon_i\}_{i=1}^k$ are fixed in advance versus when an analyst can adaptively select the privacy parameters ε_i at each round i based on previous outcomes before i. However, we focus on the traditional view of DP that fixes all the privacy parameters up front.

In the local setting of differential privacy, interactivity and adaptivity have been shown to be significant in learning algorithms and estimation tasks, see [12, 18, 10], although for some estimation tasks in more restricted interactivity models, there is no distinction [5]. However, such interactivity models are not relevant to the central model since mechanisms are designed to take the full dataset as input rather than designing algorithms on each datum as in the local model. Our result is in a similar vein to these results that consider the possible impact to the privacy loss from giving the adversary additional power.

We find the gap here particularly interesting because this is such a natural setting and has practical interest in the deployment of top-k algorithms [6]. For such data queries, it would be ideal to allow the analyst to adaptively interact with a DP system, rather than having the analyst select all the mechanisms up front and produce results as a batch. For example, consider the exponential mechanism as simply reporting the (noisy) maximum index for some metric of interest, but only for a certain subgroup and the analyst specifies the classifications for this subgroup, such as company, job title, geographic location, etc. Even if we fix the privacy parameter, our privacy loss will increase if we allow the analyst to adaptively select these classifications in subsequent queries.

Both the nonadaptive and adaptive setting will have practical importance and the distinction will be important in efficiently computing the tightest possible bounds on the privacy loss. In particular, our nonadaptive and efficiently computable composition formulation can be applied in a *dashboard setting*, where the set of queries that are privately output for a dataset is predetermined, and could include top-k queries for all the metrics of interest. Further, we know that our bound cannot be improved in this setting. Alternatively, our efficiently computable improved bounds for the adaptive setting can be applied in an *API setting* mentioned above, where the analyst adaptively interacts with the DP system.

While the improvements we give here in bounding the overall privacy loss are not asymptotically significant, if we consider the amount of privacy loss to be fixed, then increasing the number of allowable queries by a constant factor can still have a substantial impact on practical deployments. From our results in Figure 1, our nonadaptive composition bound allows for about four times more queries than the optimal composition for general DP mechanisms given a fixed privacy loss budget. Furthermore, this optimal composition allows for about two times more queries than the improved bounds given in [6]. Additionally, in some settings our improvement for the adaptive composition bound of exponential mechanisms allows for about three times more queries than both the optimal composition for DP mechanisms and the improved bounds in [6].

1.1 Our contributions

We informally summarize our main contributions here and will give the formal statements in Section 3 once we have set up the requisite notation.

1. We show that there is indeed a gap between the optimal composition bound when an adversary can adaptively select different exponential mechanisms at each round as opposed to an adversary who selects all the exponential mechanisms in advance. This is in contrast

to traditional DP composition bounds, which showed no difference between these different adversaries in terms of the privacy loss.

2. For the nonadaptive adversary, we provide an explicitly computable formula for the optimal composition bound that can be computed in $O(k^2)$ time, where k is the total number of exponential mechanisms that are executed.

3. For the adaptive adversary, we provide an explicit formulation for the optimal composition, but in a recursive formulation that is intractable to compute for even reasonably sized k. We then improve upon the previous upper bound on the privacy loss by giving an improved KL divergence bound, and further provide a numerical scheme based on the moment generating function of the privacy loss to obtain an even better upper bound on the optimal composition.

Although we have presented the exponential mechanism as a specific DP mechanism, it is also important to discuss its generality. In particular, there is the folklore result that states that any (pure) DP mechanism can be written in terms of an exponential mechanism with a particular quality score, i.e. the log-density of the mechanism [13]. Hence, it might seem that the optimal k-fold adaptive composition bound over the class of exponential mechanisms, or BR mechanisms, is also the optimal k-fold adaptive composition bound over the class of all DP mechanisms. However, sometimes taking general DP mechanisms, such as randomized response or the Laplace mechanism, to the generic exponential mechanism form could result in a different overall privacy parameter. Hence, a general ε-DP mechanism can be written in terms of an exponential mechanism with parameter ε', which can be up to a factor 2 larger than ε. See Section 3.5 for more discussion.

2 Preliminaries

In this section, we set up the necessary notation and definitions for our results. It will be necessary in our analysis to use a generalized version of randomized response that corresponds to BR mechanisms. Similar to the work in the optimal composition bounds for DP mechanisms, our goal will be to reduce composition to adaptive calls of this more generalized randomized response than the one used in the optimal DP composition analysis [11, 14]. For this reduction, we give a more fine-grained definition of adaptive composition, that is equivalent to previous versions, but includes details that were not necessary for standard DP composition. In particular, the class of algorithms that we want to give a DP composition bound for is not closed under convex combinations. Thus, an adversary can randomize over different algorithms in the same class and the resulting algorithm is no longer in that class. Finally, we give the definition of optimal composition and an alternative formulation that will be easier to work with.

We first cover the standard differential privacy definition from [8, 7], where we will say that two datasets $x, x' \in \mathcal{X}$ are neighbors if they differ in the addition or deletion of one individual's data, sometimes denoted as $x \sim x'$.

Definition 2.1. *A mechanism $M : \mathcal{X} \to \mathcal{Y}$ is (ε, δ)-differentially-private (DP) if the following holds for any neighboring dataset x, x' and $S \subseteq \mathcal{Y}$:*

$$\Pr[M(x) \in S] \leqslant \mathrm{e}^{\varepsilon} \Pr[M(x') \in S] + \delta.$$

Also if $\delta = 0$, we simply write ε-DP.

We now present the definition of bounded range from Durfee and Rogers [6], which was useful in improving the composition bounds for their algorithms.

Definition 2.2. *A mechanism $M : \mathcal{X} \to \mathcal{Y}$ is ε-bounded-range (BR) if the following holds for any neighboring dataset x, x' and $y_1, y_2 \in \mathcal{Y}$:*

$$\frac{\Pr[M(x) = y_1]}{\Pr[M(x') = y_1]} \leqslant \mathrm{e}^{\varepsilon} \frac{\Pr[M(x) = y_2]}{\Pr[M(x') = y_2]}.$$

Note that for continuous outcome spaces, we can use the probability density function instead. We then have the following equivalent formulation of BR mechanisms, which will be easier to use in our analysis.

Corollary 2.1. *A mechanism $M : \mathcal{X} \to \mathcal{Y}$ is ε-BR if and only if for any neighboring databases x, x' there exists some $t \in [0, \varepsilon]$ such that for any outcome $y \in \mathcal{Y}$ we have*

$$t - \varepsilon \leq \ln\left(\frac{\Pr[M(x) = y]}{\Pr[M(x') = y]}\right) \leq t$$

We also have the following connection between BR and (pure) DP.

Lemma 2.1 (Corollary 4.2 in [6])**.** *If M is ε-BR then it is ε-DP. Furthermore, if M is ε-DP, then it is also 2ε-BR.*

We will now define the exponential mechanism in terms of a quality score $u : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ and its sensitivity $\Delta u := \max_{y \in \mathcal{Y}} \max_{x \sim x'} |u(x, y) - u(x', y)|$.

Definition 2.3 (Exponential Mechanism [13])**.** *The exponential mechanism $M_u : \mathcal{X} \to \mathcal{Y}$ samples an outcome $y \in \mathcal{Y}$ with probability proportional to* $\exp\left(\frac{\varepsilon u(x,y)}{2\Delta u}\right)$.

The factor of two accounts for the possibility that the normalization term can also change with a neighboring dataset and for some quality scores, i.e. monotonic, the factor of 2 is not necessary. We then have the following result.

Theorem 1. *The exponential mechanism is ε-DP [13]. Further, the exponential mechanism is ε-BR [6].*

Throughout the rest of this work, we will use a generalized version of randomized response, which our analysis will primarily focus on and we define below.

Definition 2.4 (Generalized Random Response)**.** *For any $\varepsilon \geq 0$ and $t \in [0, \varepsilon]$, let $\mathtt{RR}_{\varepsilon,t} : \{0, 1\} \to \{0, 1\}$ be a randomized mechanism in terms of probabilities $q_{\varepsilon,t}$ and $p_{\varepsilon,t}$ such that*

$$\mathtt{RR}_{\varepsilon,t}(0) = 0 \text{ w.p. } \frac{1 - e^{t-\varepsilon}}{1 - e^{-\varepsilon}} =: q_{\varepsilon,t} \quad \text{and} \quad \mathtt{RR}_{\varepsilon,t}(0) = 1 \text{ w.p. } \frac{e^{t-\varepsilon} - e^{-\varepsilon}}{1 - e^{-\varepsilon}} =: 1 - q_{\varepsilon,t}$$
$$\mathtt{RR}_{\varepsilon,t}(1) = 0 \text{ w.p. } \frac{e^{-t} - e^{-\varepsilon}}{1 - e^{-\varepsilon}} =: p_{\varepsilon,t} \quad \text{and} \quad \mathtt{RR}_{\varepsilon,t}(1) = 1 \text{ w.p. } \frac{1 - e^{-t}}{1 - e^{-\varepsilon}} =: 1 - p_{\varepsilon,t}.$$

Note the $\mathtt{RR}_{\varepsilon,\varepsilon/2}$ is simply the standard randomized response with privacy parameter $\varepsilon/2$ [19]. We will typically drop the dependence of ε in $\mathtt{RR}_{\varepsilon,t} \equiv \mathtt{RR}_t$ when it is clear from context. It will be useful to also define what we mean by optimal privacy parameters, which we will write as a function $\delta_{\mathtt{OPT}}$ of a mechanism and a global DP parameters ε_g.

Definition 2.5 (Optimal Privacy Parameters)**.** *Given a mechanism $M : \mathcal{X} \to \mathcal{Y}$ and any $\varepsilon \in \mathbb{R}$, we define the optimal δ to be*

$$\delta_{\texttt{OPT}}(M, \varepsilon) := \inf \left\{ \delta : \ M \text{ is } (\varepsilon, \delta)\text{-}DP \right\}$$

Further, if $\mathcal{M}$ is a class of mechanisms $M : \mathcal{X} \to \mathcal{Y}$, then for any $\varepsilon \in \mathbb{R}$, we define

$$\delta_{\texttt{OPT}}(\mathcal{M}, \varepsilon) := \sup_{M \in \mathcal{M}} \delta_{\texttt{OPT}}(M, \varepsilon)$$

Fact 1. *For any mechanism $M : \mathcal{X} \to \mathcal{Y}$ and $\varepsilon \in \mathbb{R}$*

$$\delta_{\texttt{OPT}}(M, \varepsilon) = \sup_{x \sim x'} \int_{y \in \mathcal{Y}} \max\{\Pr[M(x) = y] - e^{\varepsilon} \Pr[M(x') = y], 0\} dy \tag{1}$$

Proof. It follows immediately from definition that M is (ε, δ)-DP if and only if

$$\sup_{x \sim x'} \sup_{S \subseteq \mathcal{Y}} \left\{ \Pr[M(x) \in S] - \mathrm{e}^{\varepsilon} \Pr[M(x') \in S] \right\} \leq \delta$$

This immediately implies

$$\sup_{x \sim x'} \sup_{S \subseteq \mathcal{Y}} \left\{ \Pr[M(x) \in S] - \mathrm{e}^{\varepsilon} \Pr[M(x') \in S] \right\} = \delta_{\texttt{OPT}}(M, \varepsilon)$$

Furthermore, it is straightforward to see that for any neighbors x, x'

$$\sup_{S \subseteq \mathcal{Y}} \left\{ \Pr[M(x) \in S] - \mathrm{e}^{\varepsilon} \Pr[M(x') \in S] \right\} = \int_{y \in \mathcal{Y}} \max\{\Pr[M(x) = y] - e^{\varepsilon} \Pr[M(x') = y], 0\} dy$$

□

2.1 Improved semantics for the exponential mechanism

Here we present a slight modification to the traditional exponential mechanism presented in Definition 2.3. In particular, rather than presenting the probability of selecting different outcomes in terms of the quality score's *sensitivity*, we define it in terms of what we call the *range* of the quality score. This leads to a simpler formulation of the exponential mechanism that does not have to consider whether a quality score is monotonic or not, i.e. whether to include a factor of two or not in the probability sampling rate, and for this reason we only view our modification as a semantic improvement. Additionally, we present the following example, to show that defining the exponential mechanism in terms of the max sensitivity leads to unwanted properties, which suggests that sensitivity is not a canonical parameter that should appear in the exponential mechanism.

Example 1. *Let $u : \mathcal{X} \times [m] \to \mathbb{R}$ be an arbitrary quality score with sensitivity Δu. Consider an arbitrary function $f : \mathcal{X} \to \mathbb{R}$ on the data domain. We define the alternate quality score $u'(x, i) := u(x, i) + f(x)$. It is easy to see that*

$$\frac{e^{\varepsilon u(x,i)}}{\sum_i e^{\varepsilon u(x,i)}} = \frac{e^{\varepsilon u'(x,i)}}{\sum_i e^{\varepsilon u'(x,i)}}.$$

That is, the privacy properties of the exponential mechanism with quality score u and u' are equivalent. However, it is very common that $\Delta u \neq \Delta u'$. For example let $X = Y = \{0, 1\}$ and $u(x, y) = x + y$, $f(x) = 10x$ and hence $u'(x, y) = 11x + y$. Clearly, $\Delta u = 1$ and $\Delta u' = 11$.

Note that this example is carefully constructed to show that using sensitivity has unwanted properties and we found no examples of such utility functions used in the literature. However, it would be nice to have a definition that also optimally handles such utility functions, in addition to encapsulating the monotonic case in the definition.

Given a quality score $u : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, consider defining the exponential mechanism in terms of some function of the quality score $\phi(u)$, e.g. $\phi(u) \equiv 2\Delta u$ would give us the traditional exponential mechanism. Instead, let's consider the property that $\phi(u)$ needs to satisfy to ensure a mechanism $M : \mathcal{X} \to \mathcal{Y}$ is ε-BR, and hence ε-DP. Let $x, x' \in \mathcal{X}$ be neighbors and fix outcomes $y, y' \in \mathcal{Y}$. To ensure ε-BR, we require the following condition on $\phi(u)$ (note that the normalization factors cancel)

$$\frac{\exp\left(\frac{\varepsilon u(x,y)}{\phi(u)}\right)}{\exp\left(\frac{\varepsilon u(x',y)}{\phi(u)}\right)} \leq \mathrm{e}^{\varepsilon} \cdot \frac{\exp\left(\frac{\varepsilon u(x,y')}{\phi(u)}\right)}{\exp\left(\frac{\varepsilon u(x',y')}{\phi(u)}\right)} \qquad \Leftrightarrow \qquad u(x, y) - u(x', y) - u(x, y') + u(x', y') \leq \phi(u).$$

With this observation, we define the *range* $\widetilde{\Delta} u$ of a function u as the following

$$\begin{aligned}\widetilde{\Delta} u &:= \sup_{x \sim x', y, y' \in \mathcal{Y}} u(x, y) - u(x', y) - u(x, y') + u(x', y') \\ &= \sup_{x \sim x'} \left\{ \max_{y} \left\{ u(x, y) - u(x', y) \right\} - \min_{y'} \left\{ u(x, y') - u(x', y') \right\} \right\}\end{aligned}$$

We then have the following useful properties of the range.

Proposition 1. *The range $\widetilde{\Delta} u$ of a function $u : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ has the following properties*

- $\widetilde{\Delta} u = \widetilde{\Delta} u'$ *when* $u'(x, y) = u(x, y) + f(x)$
- $\widetilde{\Delta} u \leqslant 2 \cdot \Delta u$.
- $\widetilde{\Delta} u = \Delta u$ *if u is monotone.*

We then have the immediate result, which presents a variant of the exponential mechanism in terms of the range, rather than the sensitivity, of the quality score.

Proposition 2. *The mechanism $M : \mathcal{X} \to \mathcal{Y}$ that samples $y \in \mathcal{Y}$ with probability proportional to $\exp\left(\frac{\varepsilon u(x, y)}{\widetilde{\Delta} u}\right)$ is ε-BR, and hence ε-DP.*

2.2 Formally defining composition

We now present the definition of adaptive composition for DP algorithms in the setting introduced by Dwork et al. [9]. Our definition will be slightly more explicit in how we formulate the adversary. Specifically, the adaptive composition game in [9] does not explicitly allow the analyst to use its own personal randomness in picking the mechanism at each round. Defining the analyst in this

way is fine if the analyst selects a DP mechanism at each round, since we know that any convex combination of ε-DP mechanisms is still ε-DP. For example, if M' and M'' are arbitrary ε-DP mechanisms, then if we define the mechanism M to run M' with probability p and run M'' with probability $(1-p)$, the mechanism M is ε-DP. Therefore, any randomness used by the adversary in their choice of ε-DP mechanism can simply be simulated by another ε-DP mechanism and can be ignored in the definition.

In full generality, the class of mechanisms that we allow the analyst to select from at each round may not necessarily be closed under convex combinations. In particular, we will be considering the setting in which the class of mechanisms the adversary can choose from is restricted to ε-BR mechanisms, which is not closed under convex combinations, see Section 4.1. Hence, in the adaptive composition game `AdComp` presented in Algorithm 1, we decompose the adversary into a randomized and deterministic component. The adversary will then use personal randomness $\mathcal{R}$ at each round and based on this will then use a deterministic function $\mathcal{D}$ to select a mechanism M_i from the class of algorithms $\mathcal{E}_i$ at round i.[1] As one would expect and we will show, the adversary cannot add their own independent randomness that is data-independent and further degrade privacy. The output of the adaptive composition game will be a sequence of random coins the adversary uses and the outcomes from applying the mechanism for the corresponding databases (given the bit b), which we write as $R_0, A_1^b, \cdots, R_{k-1}, A_k^b, R_k$.

Algorithm 1 `AdComp`$(\mathcal{A} = (\mathcal{R}, \mathcal{D}), (\mathcal{E}_1, \cdots, \mathcal{E}_k), b)$, where $\mathcal{D}$ is a deterministic algorithm, $\mathcal{R}$ is a randomized algorithm, $\mathcal{E}_1, \cdots, \mathcal{E}_k$ are classes of randomized algorithms, and $b \in \{0,1\}$.

$r_0 \sim \mathcal{R}(\emptyset)$.
for $i = 1, \cdots, k$ **do**
  $\mathcal{D}(r_0, A_1^b, \cdots, r_{i-1})$ selects neighboring datasets $x^{i,0}, x^{i,1}$, and $M_i \in \mathcal{E}_i$
  $\mathcal{A}$ receives $A_i^b = M_i(x^{i,b})$
  $r_i \sim \mathcal{R}(r_0, A_1^b, \cdots, r_{i-1}, A_i^b)$
return view $V^b = (r_0, A_1^b, r_1, \cdots, r_{k-1}, A_k^b, r_k)$

Definition 2.6 (k-fold Adaptive Composition)**.** *Given classes of randomized algorithms $\vec{\mathcal{E}} = (\mathcal{E}_1, \cdots \mathcal{E}_k)$, we say $\vec{\mathcal{E}}$ is $(\varepsilon_g, \delta_g)$-DP under k-fold adaptive composition if for any adversary $\mathcal{A}$ and $b \in \{0,1\}$, along with any set S that is a subset of outputs of $\mathtt{AdComp}(\mathcal{A}, \vec{\mathcal{E}}, \cdot)$, we have*

$$\Pr[\mathtt{AdComp}(\mathcal{A}, \vec{\mathcal{E}}, b) \in S] \leq e^{\varepsilon_g} \Pr[\mathtt{AdComp}(\mathcal{A}, \vec{\mathcal{E}}, 1-b) \in S] + \delta_g.$$

It will be important to distinguish adaptive and nonadaptive adversaries in our composition bounds. The nonadaptive adversary selects all the mechanisms and neighboring datasets to be used at each round prior to any computation on the dataset. For this case, we can simply study the privacy guarantees of a mechanism $M = M_1 \times M_2 \times \cdots \times M_k$ where each M_i is ε_i-BR and $M(x) = (M_1(x), M_2(x), \cdots, M_k(x))$.

[1] Similarly, Rogers et al. [17] defined a *simulated* game which explicitly decomposed the adversary into a deterministic post-processing function of randomized response at each round and then used randomness at the beginning of all the interactions to simulate the adaptive randomness at each round. They showed that such a simulated game is equivalent to the adaptive parameter composition game, which allowed them to simply consider randomized response mechanisms at each round with a deterministic adversary.

3 Overview of results and techniques

Given the necessary notation and setup, we present formal statements of our main results along with the intuition and techniques used to prove these results. We detail the formal proofs in the sequel.

3.1 Reduction to generalized randomized response

Similar to [11, 14], we first want to identity the "worst-case" mechanism for the class of BR mechanisms, which is to say that any BR mechanism can be simulated through post-processing of this worst-case mechanism. It then follows that composition over the class of BR mechanisms can be reduced to simply considering composition of this worst-case mechanism, allowing for explicit computation of the optimal composition. For the class of ε-DP mechanisms, the worst-case mechanism was shown to be randomized response through both the hypothesis testing interpretation [11], and explicitly constructing the post-processing function [14]. Rather than consider the class of exponential mechanisms in our analysis, we will instead focus on the more general class of BR mechanisms due to the fact that the BR property in Corollary 2.1 closely matches the definition of (pure) DP. We also show in Section 4.2 that this definition is essentially equivalent to the standard use of the exponential mechanism, which is to say that the privacy loss is identical for the worst-case BR mechanism and the exponential mechanism. We then categorize the worst-case BR mechanisms similarly to analysis done in [11, 14]. More specifically, we know from Corollary 2.1 that if a mechanism $M : \mathcal{X} \to \mathcal{Y}$ is ε-BR, then for any neighboring x, x' there exists some $t \in [0, \varepsilon]$ such that for any $y \in \mathcal{Y}$,

$$t - \varepsilon \leq \ln\left(\frac{\Pr[M(x) = y]}{\Pr[M(x') = y]}\right) \leq t.$$

Note that if for each neighboring x, x' we have that $t = \varepsilon/2$, then M is also $\frac{\varepsilon}{2}$-DP. It then follows from [11, 14] that when $t = \varepsilon/2$ the worst-case mechanism is simply randomized response with parameter $\frac{\varepsilon}{2}$. Intuitively, this is the mechanism M such that for any $y \in \mathcal{Y}$ one of the bounds is tight, in other words

$$\ln\left(\frac{\Pr[M(x) = y]}{\Pr[M(x') = y]}\right) \in \left\{-\frac{\varepsilon}{2}, \frac{\varepsilon}{2}\right\}.$$

For our setting, this same intuition must hold for $t = \varepsilon/2$, and we then generalize this to any $t \in [0, \varepsilon]$ where the worst-case mechanism M is such that

$$\ln\left(\frac{\Pr[M(x) = y]}{\Pr[M(x') = y]}\right) \in \{t - \varepsilon, t\}.$$

This generalization is exactly our Definition 2.4, and using a similar interpretation to hypothesis testing, we show that for any given $t \in [0, \varepsilon]$ this is the worst-case mechanism that satisfies the BR property. While this result is largely unsurprising, in Section 4 we give a thorough treatment towards proving that both nonadaptive and adaptive composition can be reduced to this generalized random response at each step where some $t \in [0, \varepsilon]$ is chosen either nonadaptively or adaptively.

Note that for composition over ε-DP mechanisms, the worst-case mechanism is simply randomized response, hence there is no difference between the nonadaptive and adaptive setting because the worst-case is always randomized response regardless of previous outcomes. However, in our setting

the same conclusion is not necessarily true because the adversary now has the power of adaptively choosing $t \in [0, \varepsilon]$ at each round. We then first restrict our consideration to the nonadaptive setting and consider the optimal composition of this setting.

3.2 Nonadaptive optimal composition

As with the previous work on advanced and optimal composition for (ε, δ)-DP mechanisms, it does not suffice to simply consider one pair $(\varepsilon_g, \delta_g)$, but instead we consider a parameter as a function of the other parameter to get a full curve of privacy loss parameters. Note that throughout this work, we will use similar conventions to [14] in that (ε, δ) will denote the privacy parameters of a single mechanism and $(\varepsilon_g, \delta_g)$ will denote the global privacy parameters that are for the composition of these mechanisms. While the previous optimal composition bounds considered fixing $\delta_g \in [0, 1]$ and computing the optimal ε_g as a function of the δ_g, it will be easier for us to write δ_g as an explicit formula of ε_g given Fact 1, which is also seen in [11, 14]. These formulations are equivalent, so for simplicity we will instead consider fixing ε_g and computing the optimal δ_g.

Having restricted our consideration to the nonadaptive setting and reducing to the worst-case mechanisms being our generalized random response, it is then straightforward to obtain the optimal composition formula for the heterogeneous setting of ε_i-BR mechanisms. We define the following class $\mathcal{M}_{\texttt{BR}}^{1:k}$ of nonadaptive heterogeneous BR mechanisms and $\mathcal{M}_{\texttt{BR}}^{k}$ of nonadaptive homogeneous BR mechanisms as

$$\mathcal{M}_{\texttt{BR}}^{1:k} := \{M_1 \times \cdots \times M_k : M_i \text{ is } \varepsilon_i\text{-BR}\} \qquad \mathcal{M}_{\texttt{BR}}^{k} := \{M_1 \times \cdots \times M_k : M_i \text{ is } \varepsilon\text{-BR}\}. \tag{2}$$

Lemma 3.1. *Recall from Definition 2.4 we have $p_{\varepsilon_i,t_i}, q_{\varepsilon_i,t_i}$. We then have*

$$\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{1:k}, \varepsilon_g) = \sup_{\mathbf{t} \in \prod_{i\in[k]}[0,\varepsilon_i]} \sum_{S \subseteq \{1,\dots,k\}} \max\left\{ \prod_{i \notin S} q_{\varepsilon_i,t_i} \prod_{i \in S}(1 - q_{\varepsilon_i,t_i}) - e^{\varepsilon_g} \prod_{i \notin S} p_{\varepsilon_i,t_i} \prod_{i \in S}(1 - p_{\varepsilon_i,t_i}), 0 \right\}.$$

Note this formulation can in some ways be seen as a generalization of the following result from Murtagh and Vadhan [14], although we only state it in the nonadaptive setting (as well as fix $\delta_i = 0$), it does also hold in the adaptive setting.

Theorem 2 (Theorem 1.5 from Murtagh and Vadhan [14])**.** *Let $\mathcal{M}_{\texttt{DP}}^{1:k}$ be the class of nonadaptive composed mechanism $M = M_1 \times \cdots \times M_k$ where each M_i is ε_i-DP, then we have*

$$\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{DP}}^{1:k}, \varepsilon_g) = \frac{1}{\prod_{j=1}^{k}(1 + e^{\varepsilon_j})} \cdot \sum_{S \subseteq \{1,\dots,k\}} \max\left\{ \exp\left(\sum_{i \in S} \varepsilon_i\right) - e^{\varepsilon_g} \exp\left(\sum_{i \notin S} \varepsilon_i\right), 0 \right\}.$$

In particular, if we set $t_i = \frac{\varepsilon_i}{2}$ for all i instead of taking the sup, then this is equal to the LHS of the equation (1) in Theorem 1.5 for [14] where we replace ε_i with $\frac{\varepsilon_i}{2}$. Equivalently, by setting $t_i = \frac{\varepsilon_i}{2}$ for all i, this is the optimal composition of $\frac{\varepsilon_1}{2}$-DP, $\dots$, $\frac{\varepsilon_k}{2}$-DP mechanisms.

Similar to the result in Kairouz et al. [11] on optimal composition of DP mechanisms, we will restrict our consideration to the homogeneous setting where $\varepsilon_1, \cdots, \varepsilon_k = \varepsilon$ in an attempt to obtain a formulation that is efficiently computable. However, this formulation will be far more difficult to simplify than the optimal composition of DP mechanisms because of the supremum over $t_i \in [0, \varepsilon]$. Our simplification will require significant technical work that will ultimately be done in two key

steps: 1) we show that the supremum is achieved when all $t_i = t_j$ for $i \neq j$, and 2) we show that the supremum is achieved by a certain value $t_i = t^* \in [0, \varepsilon]$ contained in a set of at most k possible values. This will then yield an explicit and efficiently computable formulation of the optimal nonadaptive composition of BR mechanisms.

Theorem 3. *Consider the homogeneous case where $\varepsilon_i = \varepsilon$ for each $i \in [k]$, then we have for $p_{t_i} = p_{\varepsilon,t_i}$ given in Definition 2.4 and setting $t^*_\ell = \frac{\varepsilon_g + (\ell+1)\varepsilon}{k+1}$ where if $t^*_\ell \notin [0, \varepsilon]$, then we round it to the closest point in $[0, \varepsilon]$*

$$\delta_{\mathtt{OPT}}(\mathcal{M}^k_{BR}, \varepsilon_g) = \max_{0 \leqslant \ell \leqslant k} \sum_{i=0}^{k} \binom{k}{i} p_{t^*_\ell}^{k-i} (1 - p_{t^*_\ell})^i \max\left\{ \left(e^{kt^*_\ell - i\varepsilon} - e^{\varepsilon_g} \right), 0 \right\}.$$

Furthermore, this can be computed in $O(k^2)$ time.

Once again, we note that by instead setting $t^*_\ell = \frac{\varepsilon}{2}$, then this formulation is equivalent to the LHS of Theorem 1.4 in [14], which is a rephrasing of the original optimal composition formulation in [11], where we replace ε with $\varepsilon/2$.[2] We also plot the DP optimal composition bound where $\varepsilon/2$ is used as the DP privacy parameter for each individual mechanism in Figure 1. The improvement in this formulation over the optimal composition of ε-DP mechanisms is more than a factor of 2, and we empirically compare the bound for ε_g in Figure 1 as a function of k. In the figure, we label "DP OptComp" as the optimal composition bound for DP mechanisms from [14], "DR19" as the composition bound for ε-BR mechanisms from [6], and "BR OptComp" as the composition bound in Theorem 3.

Unfortunately, our proofs of this optimal composition formulation cannot be applied to the adaptive setting, pointing to the natural question of whether there is in fact further privacy loss when the adversary is given power to choose the mechanism based upon previous responses.

3.3 Additional power of adaptivity

In order to better explain the intuition behind optimal composition in both nonadaptive and adaptive settings, we rely upon the random walk interpretation of composition similar to analysis in [9, 17]. In particular, for composition of ε-DP mechanisms, we can instead consider a random walk on the real line beginning at 0, where with probability $\frac{e^\varepsilon}{e^\varepsilon+1}$ a step of ε is taken and with probability $\frac{1}{e^\varepsilon+1}$ a step of $-\varepsilon$ is taken. Given some ε_g, the goal of the adversary is to maximize the probability that the walk exceeds ε_g after k steps and the amount in which it exceeds ε_g. For achieving an upper bound on the composition as in [9], we can ignore the amount the walk exceeds ε_g and apply concentration bounds on the probability that the walk exceeds ε_g after k steps. The optimal composition from [11, 14] instead requires computing the resulting binomial distribution over the length of the walk to explicitly obtain both the probability and amount that each walk exceeds ε_g. In the nonadaptive setting, the reason we could also achieve an efficient formulation was because we proved that we can equivalently restrict all t_i to be equal and further we can restrict the possible t_i to a smaller set, so our computation once again became equivalent to examining each respective binomial distribution.

[2] Interestingly, this then implies that for any ε_g where this maximum is achieved with $t_\ell = \frac{\varepsilon}{2}$, we then have that the optimal composition of ε-BR mechanisms is equivalent to the optimal composition of $\frac{\varepsilon}{2}$-DP mechanisms for that specific ε_g. We have in fact tested this and found cases in which this is true, but could not find any discernible pattern for the specific values of ε_g when the optimal composition is equivalent.

For the composition of ε-DP mechanisms, the worst-case mechanism does not require an adversarial choice, however in our setting the adversary does have the power to choose each $t_i \in [0, \varepsilon]$ in the generalized random response mechanism. This choice of t_i will then exactly determine the length of the step in each direction, where either a step of t_i is taken or a step of $t_i - \varepsilon$ is taken. It might seem like the adversary would then always choose the maximum t_i, but the probability of taking that step is inversely related to the magnitude of the step. More specifically, the larger the adversary sets t_i, the smaller the probability that the step is taken in the positive direction, presenting a natural tradeoff. Following this random walk interpretation, we can then give an explicit optimal composition in the adaptive setting as a recursive formulation that incorporates the natural maximization problem.

We begin by simplifying our notation for adaptive composition and focusing on the homogeneous case where $\varepsilon_i = \varepsilon$ for $i \in [k]$ and will address the heterogenous case in Section 6. Given some fixed $\varepsilon > 0$, let $\overrightarrow{\mathcal{M}}^k_{\texttt{BR}} := (\mathcal{M}_{\texttt{BR}}, \ldots, \mathcal{M}_{\texttt{BR}})$ be such that $\mathcal{M}_{\texttt{BR}}$ is the class of ε-BR mechanisms. We will denote the family of adaptive composition games over all adversaries as the following

$$\mathcal{A}^k_{\texttt{BR}} := \{\texttt{AdComp}(\mathcal{A}, \overrightarrow{\mathcal{M}}^k_{\texttt{BR}}, \cdot) : \text{ adversary } \mathcal{A}\}. \tag{3}$$

We then have the following result. Note that we also consider the heterogenous case for $\varepsilon_1, \cdots, \varepsilon_k$ in Lemma 6.1

Lemma 3.2. *Given global parameter ε_g and q_{ε,t_i} from Definition 2.4, we have the following optimal privacy parameter where use set $\delta_{OPT}(\mathcal{A}^0_{BR}, \varepsilon_g) = \max\{1 - e^{\varepsilon_g}, 0\}$,*

$$\delta_{OPT}(\mathcal{A}^k_{BR}, \varepsilon_g) = \sup_{t_1 \in [0,\varepsilon]} \left\{ q_{\varepsilon,t_1} \delta_{OPT}(\mathcal{A}^{k-1}_{BR}, \varepsilon_g - t_1) + (1 - q_{\varepsilon,t_1}) \delta_{OPT}(\mathcal{A}^{k-1}_{BR}, \varepsilon_g + \varepsilon - t_1) \right\}.$$

Note that this formulation does not necessarily hold for the nonadaptive setting because the choice of t_2 cannot change based upon the result of the first mechanism, and the supremum for all possible t_i gets pulled to the beginning of the expression. It is exactly this difference that will give the adversary additional power in the adaptive setting because, relying upon our random walk intuition, the natural tradeoff between the magnitude of the step t_i and the probability of that step is actually dependent upon the current position of the random walk. For example, consider a walk that begins by taking several steps in the negative direction. In order to make up this increased distance and exceed ε_g it may then become necessary to increase the subsequent values of t_i despite this decreasing probability of these steps occurring. Similarly, if the walk begins by taking several steps in the positive direction, it may become favorable to choose more conservative values of t_i and increase the probability of taking these positive steps.

We rigorously confirm this intuition that will heavily rely upon having obtained an efficient formulation of the nonadaptive optimal composition. Furthermore, we confirm that this difference in privacy loss exists for all possible values of k in our composition, and almost all choices of ε_g. As would be expected, if $\varepsilon_g = k\varepsilon$ and basic composition can be applied, then there is no difference between optimal composition in the adaptive and nonadaptive setting. We further show that this slightly extends beyond just basic composition in which the adaptive and nonadaptive setting are equal, almost completely giving a full picture of when the adversary has additional power from adaptivity.

Theorem 4. *Recall the nonadaptive family of homogeneous ε-BR mechanisms $\mathcal{M}^k_{BR}$ from (2) and $\mathcal{A}^k_{BR}$ given in (3). For any $\varepsilon_g \in [0, (k-3)\varepsilon]$ we have,*

$$\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) > \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g).$$

Further, for any $\varepsilon_g \geq (k-1)\varepsilon$, we have

$$\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) = \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g).$$

Note that under these conditions the gap only exists for $k \geq 4$. We also show that the gap exists for $k = 2, 3$ in Section 6, but the conditions do not extend as nicely and we leave them out of the theorem statement here for simplicity.

We believe that the gap is quite small for all values of ε_g and k, however we believe that proving a strong upper bound on the gap would require significant technical work and leave it to future work. We can confirm this numerically for reasonable k, but due to the nature of the recursive formulation for the adaptive setting it is intractable to check this for larger values of k. Furthermore, these numerical methods become even more computationally difficult for the case of heterogenous privacy parameters and the gap for this setting may be much larger.

3.4 Improved and efficiently computable bounds for adaptive composition

While we gave an explicit formulation of the optimal composition for the adaptive setting of BR mechanisms, the computation is not tractable, and we suspect that it has similar hardness results to [14], which we leave to future work. Accordingly, we further improve the known efficiently computable upper bounds on the adaptive composition of ε-BR mechanisms from [6]. The previous work on ε-BR composition followed a similar approach to [9] applying both an Azuma-Hoeffding bound (on the variance) and a KL divergence bound (on the bias) to achieve a reasonably simple upper bound on the optimal composition. However, the previous work only considered using the BR property to improve the bound from Azuma-Hoeffding and did not consider improving the KL divergence bound. While these bounds are quite complex to generally compute, we note that for our generalized random response it will actually be quite simple to compute the explicit KL divergence. Using our reduction to this worst-case class of mechanisms and taking the supremum over all choices of t we can give a much improved bound on the KL divergence.

Corollary 3.1. *Let $\overrightarrow{\mathcal{M}} := (\mathcal{M}_1, \mathcal{M}_2, \cdots, \mathcal{M}_k)$ where each $\mathcal{M}_i$ is the class of ε_i-BR mechanisms. We then have that $\overrightarrow{\mathcal{M}}$ is $(\varepsilon_g(\delta_g), \delta_g)$-DP under k-fold adaptive composition for any $\delta_g \geq 0$ where*

$$\varepsilon_g(\delta) = \min\left\{\sum_{i=1}^k \varepsilon_i, \sum_{i=1}^k \left(\frac{\varepsilon_i}{1-e^{-\varepsilon_i}} - 1 - \ln\left(\frac{\varepsilon_i}{1-e^{-\varepsilon_i}}\right)\right) + \sqrt{\frac{1}{2}\sum_{i=1}^k \varepsilon_i^2 \ln(1/\delta)}\right\}.$$

This gives substantial improvements over the previous bound in some settings (and we will provide plots in Section 7), but we will further improve this bound. In particular, the bound given above considers the KL divergence and Azuma-Hoeffding separately, which is to say that the worst-case $t_i \in [0, \varepsilon_i]$ is chosen separately for these two bounds instead of choosing this supremum with respect to both. In order to improve this, we backtrack a step in this method and use the same techniques from the proof of Azuma-Hoeffding but apply our more exact characterization.

Theorem 5. *Let* $\overrightarrow{\mathcal{M}} := (\mathcal{M}_1, \mathcal{M}_2, \cdots, \mathcal{M}_k)$ *each* $\mathcal{M}_i$ *is the class of* ε_i*-BR mechanisms. We then have that* $\overrightarrow{\mathcal{M}}$ *is* $(\varepsilon_g, \delta_g(\varepsilon_g))$*-DP under k-fold adaptive composition for any* $\varepsilon_g \geq 0$ *where we define* $h_\varepsilon(\lambda) := \sup_{t \in [0,\varepsilon]} \lambda(\varepsilon - t) + \ln\left(1 + p_{\varepsilon,t}(e^{-\lambda\varepsilon} - 1)\right)$ *with* $p_{\varepsilon,t} = \frac{e^{-t} - e^{-\varepsilon}}{1 - e^{-\varepsilon}}$ *and*

$$\delta_g(\varepsilon_g) = \inf_{\lambda > 0} \mathrm{e}^{-\lambda\varepsilon_g + \sum_i h_{\varepsilon_i}(\lambda)}.$$

We present plots of our results in Figure 1 for the homogeneous case, plotting ε_g as a function of k. As stated earlier, we label "ε-DP OptComp" as the optimal composition bound for DP mechanisms from [14], "DR19" as the composition bound for ε-BR mechanisms from [6], and "BR OptComp" as the composition bound in Theorem 3, which only applies in the nonadaptive setting. Furthermore, we label "OptKL" as the bound from Corollary 3.1 and "MGF" as the bound in Theorem 5. To compare our bounds with simply using the optimal DP composition bound with a half the actual privacy parameter, we also plot the DP optimal composition bound with $\varepsilon/2$ with label "$\varepsilon/2$-DP OptComp". This last curve highlights the fact that ε-BR is almost the same as $\varepsilon/2$-DP when applying composition.

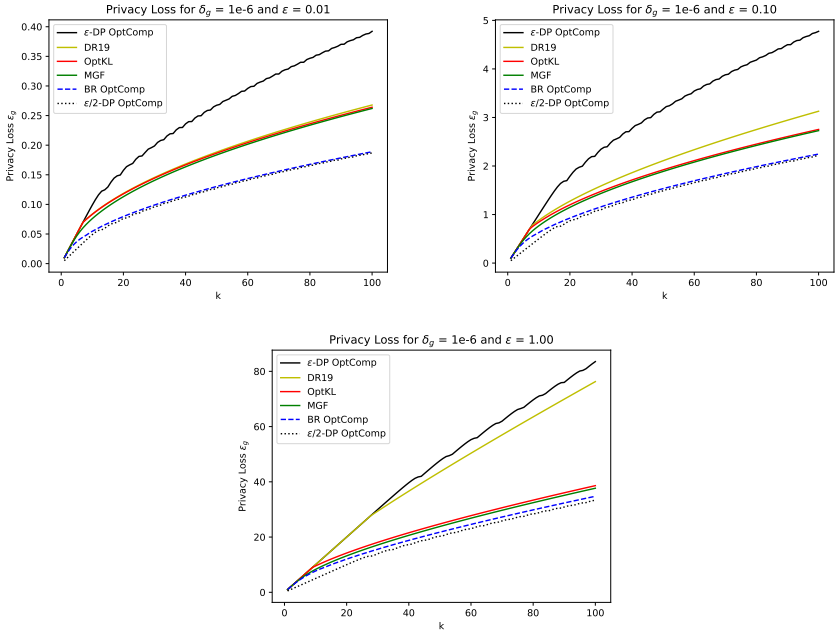

Figure 1. Comparison of optimal DP composition with the BR composition bounds in this work and in Durfee and Rogers [6]. The dashed curve only applies in the nonadaptive composition setting and the dotted curve uses the existing DP optimal composition bound with half the actual privacy parameter. We present results for $\delta_g = 10^{-6}$ and $\varepsilon \in \{0.01, 0.1, 1\}$.

3.5 Discussion of optimal DP composition bounds

Although ε-BR implies ε-DP, and the converse holds up to a factor of 2 in the privacy parameter, it is important to point out that our optimal composition analysis of BR mechanisms does not supersede the optimal composition of DP mechanisms. More specifically, consider the Laplace mechanism [8], which adds Laplace noise to a bounded sensitivity statistic. This mechanism is ε-DP, but it is also

2ε-BR yet it has a fixed value $t = \varepsilon$ for any neighboring datasets. As we will discuss more rigorously in our analysis, our optimal composition bounds for BR mechanisms follows from maximizing the bound over all sequences of t values. Hence, utilizing the optimal composition bound over BR mechanisms will result in a larger than necessary bound when considering Laplace mechanisms, and thus the optimal DP composition bounds from [11, 14] should be used. Alternatively, if we are composing exponential mechanisms that we know are ε-BR, then our composition bounds improves on the optimal composition of ε-DP mechanisms.

Consider the following example with randomized response. In this case $M_{\texttt{RR}} : \{0,1\} \to \{0,1\}$ and $M_{\texttt{RR}}(b;\varepsilon) = b$ with probability $\frac{e^\varepsilon}{e^\varepsilon+1}$. To fit this into the generic exponential mechanism, we require a quality score $u(b,b')$ and we need to calculate its sensitivity, or as we discussed in Proposition 2, its range. In this case $u(b,b') = \mathbb{1}\{b = b'\}$, which has sensitivity$\Delta u = 1$ and also has range $\tilde{\Delta} u = 2$. Whether we use the range or the sensitivity of the quality score, the generic exponential mechanism is then written as $M_u(b;\varepsilon) = \frac{e^{\varepsilon q(b,b)/2}}{e^{\varepsilon q(b,b)/2}+e^{\varepsilon q(b,1-b)/2}} = \frac{e^{\varepsilon/2}}{e^{\varepsilon/2}+1}$. Hence, we have $M_u(\cdot;2\varepsilon) = M_{\texttt{RR}}(\cdot;\varepsilon)$. The fact that randomized response can be written as $M_u(\cdot;2\varepsilon)$ implies that it is 2ε-BR, but we further note that there are only two neighboring databases for randomized response. This then allows for only one value $t \in [0, 2\varepsilon]$ from Corollary 2.1, where we see that $t = \varepsilon$ implies that this randomized response is also ε-DP. Accordingly, if we only knew the generic exponential form with parameter ε then our composition bounds would improve over the general optimal DP composition bounds from [14, 11]. However, if it is also known that each individual mechanism is also $\varepsilon/2$-DP, as is the case for randomized response with parameter $\varepsilon/2$, then the bounds from [14, 11] cannot be improved.

4 Bounded range and generalized random response

In this section, we show that k-fold adaptive composition over the class of BR mechanisms can be reduced to only considering adversaries that select a generalized random response mechanism at each step. First, we show that we can post-process the generalized random response to simulate any BR mechanism on neighboring inputs. For this proof, we will utilize the hypothesis testing interpretation of DP that was similarly used in [11] and then extended in [4]. We defer the analysis to Appendix A.

Lemma 4.1. *Let mechanism $M : \mathcal{X} \to \mathcal{Y}$ be ε-BR. For any neighboring databases $x^0, x^1 \in \mathcal{X}$, there exists some $t = t(M, x^0, x^1) \in [0, \varepsilon]$ and randomized function $\phi : \{0,1\} \to \mathcal{Y}$ that depends on M, x^0, x^1 such that for any $y \in Y$ and $b \in \{0,1\}$ we have the following equivalence in terms of the generalized randomized response mechanism from Definition 2.4.*

$$\Pr[M(x^b) = y] = \Pr[\phi(\texttt{RR}_{\varepsilon,t}(b)) = y]$$

We next show that k-fold adaptive composition over the class of BR mechanisms is equivalent to considering the class of generalized randomized response mechanisms instead.

Lemma 4.2. *Fix parameters $\varepsilon_1, \cdots, \varepsilon_k$. Let $\overrightarrow{\mathcal{M}} = (\mathcal{M}_1, \cdots, \mathcal{M}_k)$ be such that $\mathcal{M}_i$ is the class of ε_i-BR mechanisms, and let $\overrightarrow{\mathcal{RR}} = (\mathcal{RR}_1, \cdots, \mathcal{RR}_k)$ be the class such that $\mathcal{RR}_i := \{\texttt{RR}_{\varepsilon_i,t_i} : t_i \in [0, \varepsilon_i]\}$. We then have that $\overrightarrow{\mathcal{M}}$ is $(\varepsilon_g, \delta_g)$-DP under k-fold adaptive composition if and only if $\overrightarrow{\mathcal{RR}}$ is $(\varepsilon_g, \delta_g)$-DP under k-fold adaptive composition.*

Proof. Take any $\mathcal{A} = (\mathcal{D}, \mathcal{R})$ that selects mechanisms from $\mathcal{M}_i$ at round i and we will construct $\mathcal{A}' = (\mathcal{D}', \mathcal{R}')$ that selects mechanisms in $\mathcal{RR}_i$ in the following way. Replace the deterministic component $\mathcal{D}(r_0, A_1^b, \cdots r_{i-1})$ that selects neighbors x_i^0, x_i^1 and $M_i \in \mathcal{M}_i$ at each round i with $\mathcal{D}'(r_0, B_1^b, A_1^b, \cdots r_{i-1})$ that selects neighbors x_i^0, x_i^1 and $t_i(M_i, x_i^0, x_i^1)$ where $B_\ell^b = \mathtt{RR}_{t(\mathcal{M}_\ell, x_\ell^0, x_\ell^1)}(b)$ and $\ell < i$.

The new analyst $\mathcal{A}'$ receives $B_i^b = \mathtt{RR}_{t(\mathcal{M}_i, x_i^0, x_i^1)}(b)$ whereas $\mathcal{A}$ receives $A_i^b = M_i(x_i^b)$. We then construct the randomized component of $\mathcal{A}'$ in the following way. Rather than sample $r_i \sim \mathcal{R}(r_1, A_1^b, \cdots r_{i-1}, A_{i-1}^b)$, we sample $r_i' = (A_i^b, r_i) \sim \mathcal{R}'(r_1, B_1^b, A_1^b, \cdots, r_{i-1}, B_i^b)$ where first $A_i^b = \phi_i(B_i^b)$ and ϕ_i is the post-processing function described in Lemma 4.1 that depends on $\mathcal{M}_i, x_i^0, x_i^1$, then $r_i \sim \mathcal{R}(r_1, A_1^b, \cdots r_{i-1}, A_{i-1}^b)$, as before.

Given any outcome $(r_0, A_1^b, \cdots r_{k-1}, A_k^b, r_k)$, we know that there exists a post-processing function ψ such that for $b \in \{0, 1\}$

$$\Pr\left[\mathtt{AdComp}(\mathcal{A}, \overrightarrow{\mathcal{M}}, b) = \left(r_0, A_1^b, \cdots r_{k-1}, A_k^b, r_k\right)\right] = \Pr\left[\psi\left(\mathtt{AdComp}(\mathcal{A}', \overrightarrow{\mathcal{RR}}, b)\right) = \left(r_0, A_1^b, \cdots r_{k-1}, A_k^b, r_k\right)\right]$$

□

4.1 Handling convexity for BR composition

In this section, we discuss a technicality for adaptive composition of BR mechanisms. As discussed earlier, BR mechanisms are not closed under convex combinations, and this can be easily seen by simply considering a mechanism that has four possible outputs from randomizing over $\mathtt{RR}_{\varepsilon,t_1}$ and $\mathtt{RR}_{\varepsilon,t_2}$ where $t_1 \neq t_2$. This allows adversaries potentially additional power when they can randomize between different BR mechanisms at each round, which is not the case for classes of mechanisms that are closed under convex combinations, such as DP.

Despite this technicality, we will show that allowing the analyst this adaptive randomness at each step does not increase the privacy loss. Consider the same adaptive game in Algorithm 1, but now we take away the adversary's ability to add their own data-independent randomness at each round, which we will denote as $\mathcal{A} = (\emptyset, \mathcal{D})$. We will show that this has the same level of privacy regardless of the class of randomized algorithms used.

Definition 4.1 (Adaptive Composition without Adversarial Randomness)**.** *Given classes of randomized algorithms $\vec{\mathcal{E}} = (\mathcal{E}_1, \cdots \mathcal{E}_k)$, we say $\vec{\mathcal{E}}$ is $(\varepsilon_g, \delta_g)$ differentially private under k-fold adaptive composition without adversarial randomness if for any adversary $\mathcal{A} = (\emptyset, \mathcal{D})$ that does not have any randomness of its own and $b \in \{0, 1\}$, along with any set S that is a subset of outputs of $\mathtt{AdComp}((\emptyset, \mathcal{D}), \mathcal{E}, \cdot)$*

$$\Pr[\mathtt{AdComp}((\emptyset, \mathcal{D}), \vec{\mathcal{E}}, b) \in S] \leq e^{\varepsilon_g} \Pr[\mathtt{AdComp}((\emptyset, \mathcal{D}), \vec{\mathcal{E}}, 1-b) \in S] + \delta_g$$

Lemma 4.3. *Given any class of randomized algorithms $\vec{\mathcal{E}} = (\mathcal{E}_1, \cdots, \mathcal{E}_k)$, $\vec{\mathcal{E}}$ is $(\varepsilon_g, \delta_g)$-DP under k-fold adaptive composition without adversarial randomness if and only if $\vec{\mathcal{E}}$ is $(\varepsilon_g, \delta_g)$-DP under k-fold adaptive composition.*

Proof. We largely follow Lemma 3.4 in Rogers et al. [16] which shows the point-wise equivalence between an adversary that has access to internal randomness and with a deterministic adversary

who can then post-processes the final result. This is done by including *simulated* randomness for the deterministic adversary that can be fixed prior to any interaction with the dataset. One technical difference between our setting and theirs is that for them an adversary can select a DP algorithm, which is then a post-processing function of randomized response, at each round. This means that even if an adversary could additionally randomize between different DP algorithms at each round, the result is still DP. In our case, there is a difference between a deterministic adversary and an adversary that can randomize between BR mechanisms at each round, because the resulting mechanism may no longer be BR. However, we can just include this internal randomness of the adversary at each round in the simulated randomness from the analysis in Lemma 3.4 of [16]. Hence, we can analyze the DP guarantees for each realized value of simulated randomness. Lastly, the DP guarantee does not change under convex combinations of the realized simulated randomness, which shows that it suffices to only consider deterministic adversaries. □

Using Lemmas 4.2 and 4.3, we have the immediate result which shows that without loss of generality, we can consider deterministic adversaries that can select generalized randomized response mechanisms at each round.

Corollary 4.1. *Fix parameters $\varepsilon_1, \cdots, \varepsilon_k$. Let $\overrightarrow{\mathcal{M}} = (\mathcal{M}_1, \cdots, \mathcal{M}_k)$ be such that $\mathcal{M}_i$ is the class of ε_i-BR mechanisms, and let $\overrightarrow{\mathcal{RR}} = (\mathcal{RR}_1, \cdots, \mathcal{RR}_k)$ be the class such that $\mathcal{RR}_i := \{\texttt{RR}_{\varepsilon_i, t_i} : t_i \in [0, \varepsilon_i]\}$. We then have that $\overrightarrow{\mathcal{M}}$ is $(\varepsilon_g, \delta_g)$-DP under k-fold adaptive composition if and only if $\overrightarrow{\mathcal{RR}}$ is $(\varepsilon_g, \delta_g)$-DP under k-fold adaptive composition without adversarial randomness.*

4.2 Exponential mechanism equivalence to generalized random response

It was shown in Kairouz et al. [11] that the discretized version of the Laplace mechanism, i.e. the geometric mechanism, has the largest privacy degradation under composition. Similarly, we show that for certain quality scores the exponential mechanism is equal in distribution, up to a data independent post processing function, as the generalized randomized response mechanism. Among this class of quality scores is the commonly used score for counting queries. More specifically, if we run an exponential mechanism, then by post-processing we can achieve the same distribution as $\texttt{RR}_{\varepsilon,t}$ for some t, and likewise if we run $\texttt{RR}_{\varepsilon,t}$ with the same t, then by post-processing we can achieve the same distribution as the exponential mechanism. We first define the exponential mechanism that we will be considering. This mechanism is one of the most common uses of the exponential mechanism where each individual's data is a bit string over some domain, and the mechanism wants to output the maximum count for all individuals over this domain.

Definition 4.2. *Let $\mathcal{X} \equiv \{0,1\}^{n \times d}$ and $\mathbf{x} = (x_{i,j} : i \in [n], j \in [d]) \in \mathcal{X}$ for some $n \in \mathbb{N}$, and define $M_{CQ} : \mathcal{X} \to [d]$ to be the ε-DP exponential mechanism from Definition 2.3 with quality score $u(x, j) = \sum_{i=1}^{n} x_{i,j}$. Neighboring databases will result from the addition or subtraction of a bit string $x_i = \{0,1\}^d$. Note here that $\Delta u = 1$ and that u is also monotonic.*

Similar to the generalized random response mechanism, we then show that for any neighboring databases the log-ratio of the probability mass for any outcome $j \in [d]$ is only at the end points of the range.

Lemma 4.4. *For any neighboring databases $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ there exists some $t \in [0, \varepsilon]$ such that for any outcome $j \in [d]$*

$$\ln\left(\frac{\Pr[M_{CQ}(\mathbf{x}) = j]}{\Pr[M_{CQ}(\mathbf{x}') = j]}\right) \in \{t - \varepsilon, t\}$$

Proof. We first assume that $\mathbf{x}' = \mathbf{x} + x_i$ where $x_i \in \{0,1\}^d$. We first set

$$t = \ln\left(\frac{\sum_{j\in[d]} e^{\varepsilon u(\mathbf{x}',j)}}{\sum_{j\in[d]} e^{\varepsilon u(\mathbf{x},j)}}\right)$$

Note that we must have $t \in [0,\varepsilon]$ because $u(\mathbf{x},j)+1 \geq u(\mathbf{x}',j) \geq u(\mathbf{x},j)$ for all $j \in [d]$. We can then reduce our probability log-ratio to

$$\ln\left(\frac{\Pr[M_{CQ}(\mathbf{x}) = j]}{\Pr[M_{CQ}(\mathbf{x}') = j]}\right) = t + \ln\left(\frac{e^{\varepsilon u(\mathbf{x},j)}}{e^{\varepsilon u(\mathbf{x}',j)}}\right)$$

Applying our assumption that $\mathbf{x}' = \mathbf{x}+x_i$, by the definition of u we have $u(\mathbf{x}',j) = u(\mathbf{x},j)+x_{i,j}$, which reduces our expression to

$$\ln\left(\frac{\Pr[M_{CQ}(\mathbf{x}) = j]}{\Pr[M_{CQ}(\mathbf{x}') = j]}\right) = t - \varepsilon x_{i,j}$$

and this implies our desired result because $x_{i,j} \in \{0,1\}$. We assumed $\mathbf{x}' = \mathbf{x} + x_i$ and considering the other case is equivalent to flipping the fraction, where it follows from natural log properties that

$$\ln\left(\frac{\Pr[M_{CQ}(\mathbf{x}') = j]}{\Pr[M_{CQ}(\mathbf{x}) = j]}\right) = \varepsilon x_{i,j} - t$$

which also implies our desired result because $\varepsilon - t \in [0,\varepsilon]$.

□

This result is exactly why we consider the relation between this mechanism and generalized random response to be analogous to the relation between geometric noise and randomized response. For any outcome in the geometric mechanism, the magnitude of the log-ratio is always ε, but unlike randomized response there are many more than two possible outcomes. Essentially, we can consider geometric noise and this counting query mechanism to split the outcomes of their respective randomized response into many outcomes, which will be the post-processing function.

Corollary 4.2. *For any neighboring databases $\mathbf{x}^0, \mathbf{x}^1$ then there must exist some $t \in [0,\varepsilon]$ and post-processing functions ϕ and ϕ' such that $M_{CQ}(\mathbf{x}^b) \equiv \phi(\texttt{RR}_{\varepsilon,t}(b))$ and $\phi'(M_{CQ}(\mathbf{x}^b)) \equiv \texttt{RR}_{\varepsilon,t}(b)$*

Proof. Applying Lemma 4.4, we split the outcome indices in the following way with $b' \in \{0,1\}$

$$\mathcal{I}_{b'} = \left\{j \in [d] : \ln\left(\frac{\Pr[M_{CQ}(\mathbf{x}^0) = j]}{\Pr[M_{CQ}(\mathbf{x}^1) = j]}\right) = t - \varepsilon b'\right\}.$$

It is straightforward to see from Definition 2.4 that we also have

$$\ln\left(\frac{\Pr[\texttt{RR}_{\varepsilon,t}(0) = b']}{\Pr[\texttt{RR}_{\varepsilon,t}(1) = b']}\right) = t - \varepsilon b'.$$

Therefore, we must have for any $b \in \{0,1\}$ and $b' \in \{0,1\}$ that

$$\Pr[\mathtt{RR}_{\varepsilon,t}(b) = b'] = \sum_{j \in \mathcal{I}_{b'}} \Pr[M_{CQ}(\mathbf{x}^b) = j]$$

and our claim follows easily. □

From Corollary 4.1, we know that the adaptive composition of BR mechanisms can be reduced to the class of generalized random responses and that this class is parameterized over all $t \in [0, \varepsilon]$. In our proof of Lemma 4.4 we showed that the value t came from the log-ratio of the sum of exponential functions. For our definition of $\mathcal{X}$, the number of neighboring databases is countably infinite, so it is technically impossible for there to always exist some neighboring databases with a corresponding t over the uncountably infinite interval $[0, \varepsilon]$. However, we can find neighboring databases that give a log-ratio arbitrarily close to any given $t \in [0, \varepsilon]$, i.e. the set of possible t values from neighboring databases is dense in $[0, \varepsilon]$, and for all practical purposes we can consider them equivalent. Therefore, the adaptive composition game with this simple instantiation of the exponential mechanism is equivalent to an adversary being restricted to the class of generalized randomized response mechanisms at each round. This is comparable to the result in Kairouz et al. [11] that shows that the geometric mechanism achieves the worst case privacy composition bound since it also achieves the same privacy region as the standard randomized response once the neighboring datasets are fixed at each round.

5 Nonadaptive optimal composition

In this section, we first give the explicit formulation for the optimal composition of nonadaptive BR mechanisms originally stated in Lemma 3.1. The majority of the section will then be devoted to reducing this formulation to a simpler formula that can be computed in $O(k^2)$ time for the homogeneous composition case, i.e. all privacy parameters are the same at each round. This will then culminate in a proof of Theorem 3.

We will denote $\mathbf{t} = (t_1, \cdots, t_k) \in \prod[0, \varepsilon_i]$ where $\prod[0, \varepsilon_i] := [0, \varepsilon_1] \times \cdots \times [0, \varepsilon_k]$ and if all $\varepsilon_i = \varepsilon$ we will simply write $[0, \varepsilon]^k$. Recall from (2), we will denote the family of nonadaptive BR mechanisms as $\mathcal{M}_{\mathtt{BR}}^k$ for the homogeneous case and $\mathcal{M}_{\mathtt{BR}}^{1:k}$ for the heterogeneous case. Recall that we defined the optimal privacy parameters by fixing a global ε_g and giving a formula for $\delta_{\mathtt{OPT}}$ in terms of ε_g as in (1). Our first formulation follows immediately from Lemma 4.1.

Lemma 5.1.

$$\delta_{\mathit{OPT}}(\mathcal{M}_{\mathit{BR}}^{1:k}, \varepsilon_g) = \sup_{\mathbf{t} \in \prod[0,\varepsilon_i]} \max_{\mathbf{b} \in \{0,1\}^k} \sum_{\mathbf{y} \in \{0,1\}^k} \max\left\{ \prod_{i=1}^{k} \Pr[\mathit{RR}_{\varepsilon_i,t_i}(b_i) = y_i] - e^{\varepsilon_g} \prod_{i=1}^{k} \Pr[\mathit{RR}_{\varepsilon_i,t_i}(1-b_i) = y_i], 0 \right\}.$$

Proof. We know that DP is closed under post-processing, so from Lemma 4.1 we can restrict our consideration to $\mathtt{RR}_{\varepsilon_i,t_i}$ for $t_i \in [0, \varepsilon_i]$, along with $b_i \in \{0,1\}$. The formulation then follows from Definition 2.5 and Fact 1. □

We have the following symmetry result for the generalized randomized response mechanism, which will be useful in our analysis.

Claim 5.1. *For any $b \in \{0,1\}$ along with $\varepsilon \geq 0$ and $t \in [0,\varepsilon]$ we have*

$$\Pr[\mathit{RR}_{\varepsilon,t}(b) = b] = \Pr[\mathit{RR}_{\varepsilon,\varepsilon-t}(1-b) = 1-b].$$

We then use this symmetry property to show that the choice of b_i is irrelevant.

Corollary 5.1. *For any $\mathbf{t} \in \prod[0,\varepsilon_i]$ and $\mathbf{b} \in \{0,1\}^k$, and some fixed $b \in \{0,1\}$, there exists $\mathbf{t}' \in \prod[0,\varepsilon_i]$ such that*

$$\sum_{\mathbf{y}\in\{0,1\}^k} \max\left\{\prod_{i=1}^k \Pr[\mathit{RR}_{\varepsilon_i,t_i}(b_i) = y_i] - e^{\varepsilon_g}\prod_{i=1}^k \Pr[\mathit{RR}_{\varepsilon_i,t_i}(1-b_i) = y_i], 0\right\}$$
$$= \sum_{\mathbf{y}\in\{0,1\}^k} \max\left\{\prod_{i=1}^k \Pr[\mathit{RR}_{\varepsilon_i,t'_i}(b) = y_i] - e^{\varepsilon_g}\prod_{i=1}^k \Pr[\mathit{RR}_{\varepsilon_i,t'_i}(1-b) = y_i], 0\right\}$$

Proof. If $b_i = b$, then we can simply set $t'_i = t_i$. If $b_i \neq b$, then from Claim 5.1 we can set $t'_i = \varepsilon_i - t_i$ and the value of the summation will not change. □

It then follows that we can fix $b \in \{0,1\}$ to give a simpler expression, and this expression is also a generalization of the optimal composition bound in Theorem 2, where instead of the sup term over $\mathbf{t} \in \prod[0,\varepsilon_i]$, we can set each $t_i = \varepsilon_i/2$, and this becomes the optimal composition of $\frac{\varepsilon_i}{2}$-DP mechanisms.

Lemma 3.1. *Recall from Definition 2.4 we have $p_{\varepsilon_i,t_i}, q_{\varepsilon_i,t_i}$. We then have*

$$\delta_{\mathit{OPT}}(\mathcal{M}_{\mathit{BR}}^{1:k}, \varepsilon_g) = \sup_{\mathbf{t}\in\prod_{i\in[k]}[0,\varepsilon_i]} \sum_{S\subseteq\{1,\ldots,k\}} \max\left\{\prod_{i\notin S} q_{\varepsilon_i,t_i}\prod_{i\in S}(1-q_{\varepsilon_i,t_i}) - e^{\varepsilon_g}\prod_{i\notin S} p_{\varepsilon_i,t_i}\prod_{i\in S}(1-p_{\varepsilon_i,t_i}), 0\right\}.$$

Proof. Follows immediately from applying Corollary 5.1 with $b = 0$ to Lemma 5.1. □

5.1 Simplifying the optimal composition bound for the homogeneous case

Although we have a formula for the optimal composition bound over BR mechanisms, it is intractable to compute for even modest values of k. To help simplify things, we will now restrict our consideration to the homogeneous case, where all $\varepsilon_i = \varepsilon \geq 0$, and we will drop the ε from our notation, e.g. $p_{\varepsilon,t_i} \equiv p_{t_i}$. We conjecture that the heterogeneous case has a similar hardness result to compute as the result in Murtagh and Vadhan [14], but we leave that as an open problem.

Since we have shown that $\delta_{\mathtt{OPT}}(\mathcal{M}_{\mathtt{BR}}^k, \varepsilon_g)$ can be written as a sup over $\mathbf{t} \in [0,\varepsilon]^k$, we will define the function $\delta : [0,\varepsilon]^k \times \mathbb{R} \to [0,1]$ as the following

$$\delta(\mathbf{t}, \varepsilon_g) := \sum_{S\subseteq\{1,\ldots,k\}} \max\left\{\prod_{i\notin S} q_{t_i}\prod_{i\in S}(1-q_{t_i}) - e^{\varepsilon_g}\prod_{i\notin S} p_{t_i}\prod_{i\in S}(1-p_{t_i}), 0\right\}. \tag{4}$$

Written in this way, we have $\delta_{\mathtt{OPT}}(\mathcal{M}_{\mathtt{BR}}^k, \varepsilon_g) = \sup_{\mathbf{t}\in[0,\varepsilon]^k} \delta(\mathbf{t}, \varepsilon_g)$. We first show that when $\varepsilon_g \notin (-k\varepsilon, k\varepsilon)$, then the choice of $\delta(\mathbf{t}, \varepsilon_g)$ does not depend on $\mathbf{t} \in [0,\varepsilon]^k$. However, this region for ε_g is not typically interesting in most DP applications, since $\varepsilon_g = k\varepsilon$ is simply applying *basic composition* from Dwork et al. [8].

Lemma 5.2. *For any* $\mathbf{t} \in [0, \varepsilon]^k$*, if* $\varepsilon_g \leq -k\varepsilon$ *then* $\delta(\mathbf{t}, \varepsilon_g) = 1 - e^{\varepsilon_g}$*, and if* $\varepsilon_g \geq k\varepsilon$ *then* $\delta(\mathbf{t}, \varepsilon_g) = 0$*.*

Proof. Using the fact that $q_t = e^t p_t$ and $(1 - q_t) = e^{t-\varepsilon}(1 - p_t)$, we equivalently have

$$\delta(\mathbf{t}, \varepsilon_g) = \sum_{S \subseteq \{1,\ldots,k\}} \prod_{i \notin S} p_{t_i} \prod_{i \in S} (1 - p_{t_i}) \max\left\{e^{\sum t_i - |S|\varepsilon} - e^{\varepsilon_g}, 0\right\}$$

If $\varepsilon_g \geq k\varepsilon$ then $\max\{e^{\sum t_i - |S|\varepsilon} - e^{\varepsilon_g}, 0\} = 0$ for any $S \subseteq \{1, \ldots, k\}$. Similarly, if $\varepsilon_g \leq -k\varepsilon$ then $\max\{e^{\sum t_i - |S|\varepsilon} - e^{\varepsilon_g}, 0\} = e^{\sum t_i - |S|\varepsilon} - e^{\varepsilon_g}$ for any $S \subseteq \{1, \ldots, k\}$ and we get

$$\delta(\mathbf{t}, \varepsilon_g) = \sum_{S \subseteq \{1,\ldots,k\}} \left(\prod_{i \notin S} q_{t_i} \prod_{i \in S} (1 - q_{t_i}) - e^{\varepsilon_g} \prod_{i \notin S} p_{t_i} \prod_{i \in S} (1 - p_{t_i}) \right) = 1 - e^{\varepsilon_g}$$

□

For the remainder of our analysis, we will focus on the interesting setting where $\varepsilon_g \in (-k\varepsilon, k\varepsilon)$. Despite the large domain $[0, \varepsilon]^k$ of values to choose from in the $\sup_{\mathbf{t}}$ for $\delta_{\texttt{OPT}}$, we show that it suffices to consider the much smaller domain where each $t_i = t^*$ for some t^* for each $i \in [k]$. This result is crucial in determining a formula that can be computed efficiently for $\delta_{\texttt{OPT}}$. We first give an easy condition on what the t_i must satisfy to optimize the δ parameter which will be important for proving a strict inequality in the subsequent claim.

Lemma 5.3. *If* $\varepsilon_g \in (-k\varepsilon, k\varepsilon)$ *then for any* $\mathbf{t} \in [0, \varepsilon]^k$ *such that* $\delta(\mathbf{t}, \varepsilon_g) = \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)$*, we must have*

$$\varepsilon_g < \sum_{i=1}^{k} t_i < \varepsilon_g + k\varepsilon$$

Proof. Using the fact that $q_t = e^t p_t$ and $(1 - q_t) = e^{t-\varepsilon}(1 - p_t)$, we equivalently have

$$\delta(\mathbf{t}, \varepsilon_g) = \sum_{S \subseteq \{1,\ldots,k\}} \prod_{i \notin S} p_{t_i} \prod_{i \in S} (1 - p_{t_i}) \max\left\{e^{\sum t_i - |S|\varepsilon} - e^{\varepsilon_g}, 0\right\}$$

It then follows that if $\sum t_i \leq \varepsilon_g$ we must have

$$\max\left\{e^{\sum t_i - |S|\varepsilon} - e^{\varepsilon_g}, 0\right\} = 0$$

for any S and so $\delta(\mathbf{t}, \varepsilon_g) = 0$. However, if $\varepsilon_g < k\varepsilon$, then there must exist $\mathbf{t}$ such that $t_i < \varepsilon$ for each i and $\sum t_i > \varepsilon_g$. Setting $S = \emptyset$ we must have $p_{t_i} > 0$ for all i and $\max\{e^{\sum t_i} - e^{\varepsilon_g}, 0\} > 0$. Therefore, $\delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g) > 0$ and if $\sum t_i \leq \varepsilon_g$ we must have $\delta(\mathbf{t}, \varepsilon_g) < \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)$.

Similarly, if $\sum t_i \geq \varepsilon_g + k\varepsilon$ we must have the following for any subset S

$$\max\left\{e^{\sum t_i - |S|\varepsilon} - e^{\varepsilon_g}, 0\right\} = e^{\sum t_i - |S|\varepsilon} - e^{\varepsilon_g}$$

We then have the following,

$$\begin{aligned}\delta(\mathbf{t}, \varepsilon_g) &= \sum_{S \subseteq \{1,\ldots,k\}} \prod_{i \notin S} p_{t_i} \prod_{i \in S} (1 - p_{t_i}) \left(e^{\sum t_i - |S|\varepsilon} - e^{\varepsilon_g}\right) \\ &= \sum_{S \subseteq \{1,\ldots,k\}} \prod_{i \notin S} q_{t_i} \prod_{i \in S} (1 - q_{t_i}) - e^{\varepsilon_g} \prod_{i \notin S} p_{t_i} \prod_{i \in S} (1 - p_{t_i}) = 1 - e^{\varepsilon_g}\end{aligned}$$

By the same reasoning, we have $\delta(\mathbf{t}, \varepsilon_g) > 1 - e^{\varepsilon_g}$ if $e^{\sum t_i - |S|\varepsilon} - e^{\varepsilon_g} < 0$ for some $S \subseteq \{1, \cdots, k\}$ and all $t_i \in (0, \varepsilon)$, which implies $p_{t_i} \in (0, 1)$ for all i. Accordingly, we have $\delta(\mathbf{t}, \varepsilon_g) > 1 - e^{\varepsilon_g}$ if $\sum t_i < \varepsilon_g + k\varepsilon$, and if $\varepsilon_g > -k\varepsilon$, there must exist positive t_i such that $\sum t_i < \varepsilon_g + k\varepsilon$. Therefore if $\sum t_i \geq \varepsilon_g + k\varepsilon$, we must have $\delta(\mathbf{t}, \varepsilon_g) < \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)$. □

The next lemma shows that taking the average of some t_i, t_j can only increase the value of $\delta(\mathbf{t}, \varepsilon_g)$. Further, this will strictly increase the δ when the t_i satisfy the condition of the lemma above. We will be able to easily conclude from this that δ cannot be optimal if $t_i \neq t_j$ for some i, j

Lemma 5.4. *For any $\varepsilon_g \in \mathbb{R}$ and $\mathbf{t} \in [0, \varepsilon]^k$,*

$$\delta(\mathbf{t}, \varepsilon_g) \leq \delta\left(\left(\frac{t_1 + t_2}{2}, \frac{t_1 + t_2}{2}, t_3, ..., t_k\right), \varepsilon_g\right)$$

Further, the inequality is strict whenever $\varepsilon_g < \sum t_i < \varepsilon_g + k\varepsilon$ and $t_1 \neq t_2$.

The proof of this lemma will require quite a bit of technical detail which we relegate to Appendix B. We then have the immediate corollary.

Corollary 5.2. *For any $\varepsilon_g \in (-k\varepsilon, k\varepsilon)$ we must have the following for any $\mathbf{t} \in [0, \varepsilon]^k$ such that there exists some $t_i \neq t_j$*

$$\delta(\mathbf{t}, \varepsilon_g) < \delta_{\mathit{OPT}}(\mathcal{M}^k_{\mathit{BR}}, \varepsilon_g).$$

Proof. We will prove by contradiction and suppose $\delta(\mathbf{t}, \varepsilon_g) = \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)$ and $t_i \neq t_j$ for some pair of indices. Note that $\delta(\mathbf{t}, \varepsilon_g)$ is equal under permutation of the indices in $\mathbf{t}$, so without loss of generality, we let $t_1 \neq t_2$. From Lemma 5.3, we must have $\varepsilon_g < \sum t_i < \varepsilon_g + k\varepsilon$. We then apply Lemma 5.4 to get our contradiction

$$\delta(\mathbf{t}, \varepsilon_g) < \delta\left(\frac{t_1 + t_2}{2}, \frac{t_1 + t_2}{2}, t_3, ..., t_k\right) \leq \delta_{\texttt{OPT}}(\mathcal{M}^k_{\text{BR}}, \varepsilon_g)$$

□

We now prove the simplified formula for the optimal privacy parameters for the family $\mathcal{M}^k_{\texttt{BR}}$ of ε-BR mechanisms, although in the next subsection, we show that we can restrict the range $[0, \varepsilon]$ that the sup is over a smaller set.

Lemma 5.5. *For any $\varepsilon_g \in \mathbb{R}$ and $\varepsilon \geq 0$*

$$\delta_{\mathit{OPT}}(\mathcal{M}^k_{\mathit{BR}}, \varepsilon_g) = \sup_{t \in [0, \varepsilon]} \sum_{i=0}^{k} \binom{k}{i} p_t^{k-i} (1 - p_t)^i \max\left\{\left(e^{kt - i\varepsilon} - e^{\varepsilon_g}\right), 0\right\} \tag{5}$$

Proof. By Lemma 3.1 and our definition for $\delta(\mathbf{t}, \varepsilon_g)$ given in (4), $\delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g) = \sup_{\mathbf{t} \in [0, \varepsilon]^k} \delta(\mathbf{t}, \varepsilon_g)$. From Corollary 5.2 we know that for $\varepsilon_g \in (-k\varepsilon, k\varepsilon)$,

$$\delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g) = \sup_{t \in [0, \varepsilon]} \delta(t, \ldots, t, \varepsilon_g).$$

Furthermore, we know if $\varepsilon_g \geq k\varepsilon$ then $\delta(\mathbf{t}, \varepsilon_g) = 0$ for any $\mathbf{t} \in [0, \varepsilon]^k$, and also if $\varepsilon_g \leq -k\varepsilon$ then $\delta(\mathbf{t}, \varepsilon_g) = 1 - e^{\varepsilon_g}$ for any $\mathbf{t} \in [0, \varepsilon]^k$. Therefore,

$$\begin{aligned}\delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g) &= \sup_{t \in [0,\varepsilon]} \sum_{S \subseteq \{1,\dots,k\}} \prod_{i \notin S} p_t \prod_{i \in S} (1 - p_t) \max\left\{e^{kt - |S|\varepsilon} - e^{\varepsilon_g}, 0\right\} \\ &= \sup_{t \in [0,\varepsilon]} \sum_{S \subseteq \{1,\dots,k\}} p_t^{k-|S|} (1 - p_t)^{|S|} \max\left\{e^{kt - |S|\varepsilon} - e^{\varepsilon_g}, 0\right\}\end{aligned}$$

For each $i \in \{0, 1, \cdots, k\}$ there are $\binom{k}{i}$ subsets $S \subseteq \{1, \dots, k\}$ such that $|S| = i$, and grouping these together gives our desired equality. □

5.2 Efficiently computing the optimal composition bound

Now that we have a much simpler formulation of the optimal composition for BR mechanisms in (5), we will solve for the $t \in [0, \varepsilon]$ that maximizes this expression. Ultimately, we will show that there are only k different candidate values of t that maximizes $\delta((t, t, \cdots, t), \varepsilon_g)$, and give explicit expressions for these candidate values of t. These explicit expressions will also be necessary in later sections when we show that there is a difference between the adaptive and nonadaptive setting.

Since we no longer need to consider any $\mathbf{t} \in [0, \varepsilon]^k$ where $\mathbf{t}$ is not a scalar times the all ones vector, we will simplify our notation to be

$$\delta^k(t, \varepsilon_g) := \sum_{i=0}^{k} \binom{k}{i} p_t^{k-i} (1 - p_t)^i \max\left\{\left(e^{kt - i\varepsilon} - e^{\varepsilon_g}\right), 0\right\}. \tag{6}$$

Given that we want to find the t which maximizes this expression, our goal will be to take the partial derivative of this function with respect to t. The maximization within the expression will make this more difficult, however, because the maximization is over a variable term and zero, we will always be able to write $\delta_{\texttt{OPT}}$ in terms of the following function F_ℓ for some $\ell \in \{0, \cdots, k\}$ that will depend on t.

$$F_\ell(t, \varepsilon_g) := \sum_{i=0}^{\ell} \binom{k}{i} p_t^{k-i} (1 - p_t)^i \left(e^{kt - i\varepsilon} - e^{\varepsilon_g}\right). \tag{7}$$

This function is differentiable and we show its relation to $\delta^k(t, \varepsilon_g)$.

Lemma 5.6. *For any $\varepsilon_g \in \mathbb{R}$, $\varepsilon \geq 0$, and $t \in [0, \varepsilon]$, there must exist some $\ell \in [k]$ such that*

$$\delta^k(t, \varepsilon_g) = F_\ell(t, \varepsilon_g).$$

Proof. Note that $e^{kt - i\varepsilon} - e^{\varepsilon_g}$ decreases as i increases, which implies that for any $t \in [0, \varepsilon]$ there must exist some ℓ such that $\max\{e^{kt - i\varepsilon} - e^{\varepsilon_g}, 0\} = e^{kt - i\varepsilon} - e^{\varepsilon_g}$ for all $i \leq \ell$ and $\max\{e^{kt - i\varepsilon} - e^{\varepsilon_g}, 0\} = 0$ for all $i > \ell$. Therefore, because p_t and $(1 - p_t)$ are non-negative we have

$$\delta^k(t, \varepsilon_g) = F_\ell(t, \varepsilon_g).$$

□

It then follows that optimizing over $t \in [0, \varepsilon]$ for $\delta^k(t, \varepsilon_g)$ can be reduced to optimizing over $t \in [0, \varepsilon]$ for each $F_\ell(t, \varepsilon_g)$.

Corollary 5.3. *For any $\varepsilon_g \in \mathbb{R}$ and $\varepsilon \geq 0$,*

$$\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g) = \max_{0 \leq \ell \leq k} \{ \sup_{t \in [0,\varepsilon]} F_\ell(t, \varepsilon_g) \}.$$

Proof. Follows immediately from Lemma 5.6 and because for any ε_g and $t \in [0, \varepsilon]$, by definition $F_\ell(t, \varepsilon_g) \geq \delta^k(t, \varepsilon_g)$ for all ℓ. □

We will now individually solve each $\sup_{t \in [0,\varepsilon]} F_\ell(t, \varepsilon_g)$, which does not contain a maximization term and is differentiable. Our ultimate goal will be to solve $\frac{\partial F_\ell(t, \varepsilon_g)}{\partial t} = 0$, and we want explicit expressions for t, which will require a simple formulation of the partial derivate with respect to t. These explicit expressions will also be necessary for proving that there is a gap between the non-adaptive and adaptive settings. The proof for this will become quite involved with some surprisingly nice cancellation, and we relegate the details to Appendix B.

Lemma 5.7. *For $\varepsilon_g \in \mathbb{R}$, $\varepsilon \geq 0$, and $0 \leq \ell \leq k$*

$$\frac{\partial F_\ell(t, \varepsilon_g)}{\partial t} = (k - \ell) \binom{k}{\ell} p_t^{k-1-\ell} (1 - p_t)^\ell \frac{1}{1 - e^{-\varepsilon}} \left(e^{\varepsilon_g - t} - e^{kt - (\ell+1)\varepsilon} \right).$$

In order to prove that there is a gap between composition of adaptive and nonadaptive BR mechanisms, we will further utilize this exact characterization of the partial derivative to give a strict interpretation of the set of t that can achieve a maximization of our full expression. However, for giving an efficiently computable expression for optimal composition, the following simple corollary will suffice.

Corollary 5.4. *For $\varepsilon_g \in \mathbb{R}$, $\varepsilon \geq 0$, and $0 \leq \ell \leq k$*

$$\arg \sup_{t \in [0,\varepsilon]} F_\ell(t, \varepsilon_g) \in \left\{ 0, \varepsilon, \frac{\varepsilon_g + (\ell + 1)\varepsilon}{k + 1} \right\}.$$

Proof. Note that $p_t = 1$ when $t = 0$ and $p_t = 0$ when $t = \varepsilon$. Therefore $\frac{\partial F_\ell(t, \varepsilon_g)}{\partial t} = 0$ when $t \in \{0, \varepsilon\}$ or when $\varepsilon_g - t = kt - (\ell + 1)\varepsilon$ which evaluates to $t = \frac{\varepsilon_g + (\ell+1)\varepsilon}{k+1}$. □

We can now prove our main theorem for this section that gives an efficient computation of optimal composition in the non-adaptive setting, which we restate here.

Theorem 3. *Consider the homogeneous case where $\varepsilon_i = \varepsilon$ for each $i \in [k]$, then we have for $p_{t_i} = p_{\varepsilon, t_i}$ given in Definition 2.4 and setting $t_\ell^* = \frac{\varepsilon_g + (\ell+1)\varepsilon}{k+1}$ where if $t_\ell^* \notin [0, \varepsilon]$, then we round it to the closest point in $[0, \varepsilon]$*

$$\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g) = \max_{0 \leqslant \ell \leqslant k} \sum_{i=0}^{k} \binom{k}{i} p_{t_\ell^*}^{k-i} (1 - p_{t_\ell^*})^i \max \left\{ \left(e^{kt_\ell^* - i\varepsilon} - e^{\varepsilon_g} \right), 0 \right\}.$$

Furthermore, this can be computed in $O(k^2)$ time.

Proof. From Lemma 5.5 we have

$$\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g) = \sup_{t\in[0,\varepsilon]} \delta^k(t, \varepsilon_g)$$

From Lemma 5.6 and Corollary 5.3 we can restrict our consideration to values of $t \in [0, \varepsilon]$ that maximize $F_\ell(t, \varepsilon_g)$ for some $\ell \in [k]$. Applying Corollary 5.4 we can then restrict our consideration to t_ℓ for all $\ell \in [k]$, along with 0 and ε. Note that $p_t = 1$ when $t = 0$ and $p_t = 0$ when $t = \varepsilon$, so it is straightforward to verify that $\delta^k(0, \varepsilon_g) = \delta^k(\varepsilon, \varepsilon_g) = \max\{1 - e^{\varepsilon_g}, 0\}$ for any ε_g. In the proof of Lemma 5.3, we showed that $\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g) > 0$ and $\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g) > 1 - e^{\varepsilon_g}$ when $\varepsilon_g \in (-k\varepsilon, k\varepsilon)$, so it is irrelevant whether we include $0, \varepsilon$ in this setting. Finally, if $\varepsilon_g \notin (-k\varepsilon, k\varepsilon)$, then from Lemma 5.2 we have $\delta^k(t, \varepsilon_g) = \delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g) = \max\{1 - e^{\varepsilon_g}, 0\}$ for any t.

For the running time, first note that for any t we can compute $p_t^k(e^{kt} - e^{\varepsilon_g})$ in $O(k)$ time. Further, for any t, if we are given the values $\binom{k}{i} p_t^{k-i}(1 - p_t)^i$ and $e^{kt - i\varepsilon}$, then we can compute $\binom{k}{i+1} p_t^{k-(i+1)}(1 - p_t)^{i+1}$ and $e^{kt-(i+1)\varepsilon}$ in $O(1)$ time. Our running time of $O(k^2)$ then immediately follows. □

6 Adaptive optimal composition

In this section, we give the formulation for the optimal composition of BR mechanisms that can be chosen adaptively, which will be recursively defined and intractable even for reasonable k. We see no way to simplify this formulation and believe that exact computation (or even approximate) is hard, but we leave that for future work. We further show that there is in fact a gap between the optimal composition bound in the adaptive and the nonadaptive cases for all $k \geq 2$, and that this gap exists for almost all non-trivial ε_g.

We will set up some notation that is similar to what we presented in Section 3.3, although we extend it here to the heterogeneous case, where $\varepsilon_1, \cdots, \varepsilon_k$ need not be the same. Given some fixed $\varepsilon_1, \ldots, \varepsilon_k$, and mechanisms $(\mathcal{M}_1, \ldots, \mathcal{M}_k)$ be such that $\mathcal{M}_i$ is the class of ε_i-BR mechanisms. We then define the following family of mechanisms, which generalizes the homogeneous case $\mathcal{A}_{\texttt{BR}}^k$ given in (3),

$$\mathcal{A}_{\texttt{BR}}^{1:k} := \{\texttt{AdComp}(\mathcal{A}, (\mathcal{M}_1, \cdots, \mathcal{M}_k), \cdot) : \text{ adversary } \mathcal{A}\}. \tag{8}$$

The formulations and proofs in this section will rely upon recursive definitions, and it then becomes necessary to define the adaptive composition for different families of mechanisms, i.e. $\mathcal{A}_{\texttt{BR}}^{\ell:k} := \{\texttt{AdComp}(\mathcal{A}, (\mathcal{M}_\ell, \ldots, \mathcal{M}_k), \cdot) : \text{ adversary } \mathcal{A}\}$ for $\ell \in [k]$.

These definitions will then allow us to give an explicit recursive formulation of the optimal composition bounds for the k-fold adaptive composition of BR mechanisms. This formulation will follow from Corollary 4.1 which allows us to restrict our consideration to deterministically choosing t_i for our generalized random response, where this choice is conditional upon the previous outcomes. The proof will be straightforward, but notationally heavy.

Lemma 6.1. *Let $\mathcal{A}_{BR}^{1:k}$ be the class of adaptive k-fold composition of ε_i-BR mechanisms given in (8), then for any $\varepsilon_g \in \mathbb{R}$ and setting $\delta_{OPT}(\mathcal{A}_{BR}^{k+1:k}, \varepsilon_g) = \max\{1 - e^{\varepsilon_g}, 0\}$ we have,*

$$\delta_{OPT}(\mathcal{A}_{BR}^{1:k}, \varepsilon_g) = \sup_{t_1\in[0,\varepsilon_1]} \left\{ q_{\varepsilon_1,t_1} \delta_{OPT}(\mathcal{A}_{BR}^{2:k}, \varepsilon_g - t_1) + (1 - q_{\varepsilon_1,t_1}) \delta_{OPT}(\mathcal{A}_{BR}^{2:k}, \varepsilon_g + \varepsilon_1 - t_1) \right\}$$

Proof. We will prove this claim by induction, where the key will be to apply Corollary 4.1 which gives that we can equivalently restrict our consideration to adversaries without their own randomness and only consider mechanisms in the generalized randomized response class.

For our base case of $k = 1$, we have from Fact 1, Corollary 4.1, and using $\mathcal{A}^{1:1}_{\texttt{BR}}$ with privacy parameter ε_1,

$$\delta_{\texttt{OPT}}(\mathcal{A}^{1:1}_{\texttt{BR}}, \varepsilon_g) = \sup_{t_1 \in [0,\varepsilon_1]} \sup_{b_1 \in \{0,1\}} \sum_{y_1 \in \{0,1\}} \max\left\{\Pr[\texttt{RR}_{\varepsilon_1,t_1}(b_1) = y_1] - e^{\varepsilon_g} \Pr[\texttt{RR}_{\varepsilon_1,t_1}(1-b_1) = y_1], 0\right\}.$$

The symmetry of generalized random response from Claim 5.1 implies that we can fix $b_1 = 0$, and this reduces to

$$\delta_{\texttt{OPT}}(\mathcal{A}^{1:1}_{\texttt{BR}}, \varepsilon_g) = \sup_{t_1 \in [0,\varepsilon_1]} \left\{ \max\left\{q_{\varepsilon_1,t_1} - e^{\varepsilon_g} p_{\varepsilon_1,t_1}, 0\right\} + \max\left\{(1 - q_{\varepsilon_1,t_1}) - e^{\varepsilon_g}(1 - p_{\varepsilon_1,t_1}), 0\right\} \right\}.$$

Using the fact that $q_{\varepsilon_1,t_1} = e^{t_1} p_{\varepsilon_1,t_1}$ and $(1 - q_{\varepsilon_1,t_1}) = e^{t_1-\varepsilon_1}(1 - p_{\varepsilon_1,t_1})$, this reduces to our desired equality. We then assume for $k-1$, and again applying Fact 1 and Corollary 4.1 we have the following for the deterministic adversary $\mathcal{A} = (\mathcal{D}, \emptyset)$ without its own source of randomness and letting $\overrightarrow{\mathcal{RR}} = (\mathcal{RR}_1, \ldots, \mathcal{RR}_k)$ be the class such that $\mathcal{RR}_i := \{\texttt{RR}_{\varepsilon_i,t_i} : t_i \in [0, \varepsilon_i]\}$,

$$\delta_{\texttt{OPT}}(\mathcal{A}^{1:k}_{\texttt{BR}}, \varepsilon_g) = \sup_{\mathcal{A}=(\mathcal{D},\emptyset)} \sum_{\mathbf{y} \in \{0,1\}^k} \max\left\{\Pr[\texttt{AdComp}(\mathcal{A}, \overrightarrow{\mathcal{RR}}, b) = \mathbf{y}] - e^{\varepsilon_g} \Pr[\texttt{AdComp}(\mathcal{A}, \overrightarrow{\mathcal{RR}}, 1-b) = \mathbf{y}], 0\right\}.$$

We will expand this term by considering the first round where some $t_1 \in [0, \varepsilon_1]$ is chosen deterministically. Once again, we use the symmetry of generalized random response from Claim 5.1 to simply set $b_1 = 0$. The next choices are then dependent on this outcome, so the full expression becomes

$$\delta_{\texttt{OPT}}(\mathcal{A}^{1:k}_{\texttt{BR}}, \varepsilon_g) = \sup_{t_1 \in [0,\varepsilon_1]} \sum_{y_1 \in \{0,1\}} \sup_{\mathcal{A}=(\mathcal{D},\emptyset)} \left\{ \sum_{\mathbf{y} \in \{0,1\}^{k-1}} \max\left\{ \Pr[\texttt{RR}_{\varepsilon_1,t_1}(0) = y_1] \Pr[\texttt{AdComp}(\mathcal{A}, (\mathcal{RR}_2, \cdots, \mathcal{RR}_k), b) = \mathbf{y}] \right.\right.$$
$$\left.\left. - e^{\varepsilon_g} \Pr[\texttt{RR}_{\varepsilon_1,t_1}(1) = y_1] \Pr[\texttt{AdComp}(\mathcal{A}, (\mathcal{RR}_2, \cdots, \mathcal{RR}_k), 1-b) = \mathbf{y}], 0\right\}\right\}.$$

Again, we use the fact that $q_{\varepsilon_1,t_1} = e^{t_1} p_{\varepsilon_1,t_1}$ and $(1 - q_{\varepsilon_1,t_1}) = e^{t_1-\varepsilon_1}(1 - p_{\varepsilon_1,t_1})$ to pull them outside of the maximum in the expression, so that for $y_1 = 0$ the inner term then reduces to

$$q_{\varepsilon_1,t_1} \left(\sup_{\mathcal{A}=(\mathcal{D},\emptyset)} \left\{ \sum_{\mathbf{y} \in \{0,1\}^{k-1}} \max\left\{ \Pr[\texttt{AdComp}(\mathcal{A}, (\mathcal{RR}_2, \cdots, \mathcal{RR}_k), b) = \mathbf{y}] \right.\right.\right.$$
$$\left.\left.\left. - e^{\varepsilon_g - t_1} \Pr[\texttt{AdComp}(\mathcal{A}, (\mathcal{RR}_2, \cdots, \mathcal{RR}_k), 1-b) = \mathbf{y}], 0\right\}\right\}\right)$$
$$= q_{\varepsilon_1,t_1} \cdot \delta_{\texttt{OPT}}(\mathcal{A}^{2:k}_{\texttt{BR}}, \varepsilon_g - t_1).$$

This similarly follows for $y_1 = 1$, and we have our desired claim. □

Unfortunately, straightforward computation of this formulation is intractable, and we conjecture that it has a similar hardness result as in Murtagh and Vadhan [14], even in the homogenous setting. In later sections, we give improved bounds on adaptive composition for BR mechanisms, but for this section we instead focus on proving that there is indeed a gap between this optimal formulation and our formulation for the nonadaptive setting given in Theorem 3. Further, we show that this gap exists in the homogenous setting for all k and almost all choices of ε_g.

We now state the main result of this section, where we prove each claim in Lemmas 6.6 and 6.10, respectively.

Theorem 4. *Recall the nonadaptive family of homogeneous ε-BR mechanisms $\mathcal{M}^k_{BR}$ from (2) and $\mathcal{A}^k_{BR}$ given in (3). For any $\varepsilon_g \in [0, (k-3)\varepsilon]$ we have,*

$$\delta_{OPT}(\mathcal{A}^k_{BR}, \varepsilon_g) > \delta_{OPT}(\mathcal{M}^k_{BR}, \varepsilon_g).$$

Further, for any $\varepsilon_g \geq (k-1)\varepsilon$, we have

$$\delta_{OPT}(\mathcal{A}^k_{BR}, \varepsilon_g) = \delta_{OPT}(\mathcal{M}^k_{BR}, \varepsilon_g).$$

6.1 Gap between adaptive and nonadaptive optimal composition

In this section we show that there is a gap in the privacy loss between the adaptive and nonadaptive setting for BR mechanisms. Furthermore, we want to prove that this gap exists for all $k \geq 2$ and most ε_g. In fact, the only values of ε_g in which the privacy loss is equivalent is when ε_g is almost the bound from basic composition.

The general idea for proving the gap will be to also give the recursive definition for the nonadaptive optimal composition that must fix t for each recursive call. The goal will then be to show that at some point within this recursion the summation will strictly increase if the value for t is changed. This will require that we first fully characterize the possible values of t for the nonadaptive optimal composition. Fortunately, most of the heavy lifting in this regard was done in the previous section. With this characterization, we show that there is a gap when $k = 2$, and then further show that we can apply this gap for $k \geq 2$.

We will restrict our consideration to the simpler homogenous setting in which $\varepsilon_i = \varepsilon$ for all i, and use $\mathcal{A}^k_{\mathtt{BR}}$ as defined in (3) and $\mathcal{M}^k_{\mathtt{BR}}$ is the class of nonadaptive composed ε-BR mechanisms as in (2). We know that we can instead just restrict our consideration to the class of generalized random response, and the key to our the proof will be that we will be able to specify exactly which values of $t_1, \ldots, t_k$ maximize the privacy loss for the nonadaptive setting. We define this set as in terms of $\delta(\mathbf{t}, \varepsilon_g)$ from (4),

$$t_{\mathtt{OPT}}(\mathcal{M}^k_{\mathtt{BR}}, \varepsilon_g) := \{\mathbf{t} \in [0, \varepsilon]^k : \delta(\mathbf{t}, \varepsilon_g) = \delta_{\mathtt{OPT}}(\mathcal{M}^k_{\mathtt{BR}}, \varepsilon_g)\}.$$

From Corollary 5.2, we know that this set cannot contain any $\mathbf{t} \in [0, \varepsilon]^k$ such that $t_i \neq t_j$ in the interesting setting where $\varepsilon_g \in (-k\varepsilon, k\varepsilon)$. For the remainder of this section, we instead consider the definition to equivalently be

$$t_{\mathtt{OPT}}(\mathcal{M}^k_{\mathtt{BR}}, \varepsilon_g) = \{t \in [0, \varepsilon] : \delta^k(t, \varepsilon_g) = \delta_{\mathtt{OPT}}(\mathcal{M}^k_{\mathtt{BR}}, \varepsilon_g)\}$$

because when $\varepsilon_g \notin (-k\varepsilon, k\varepsilon)$ then there is not a gap between adaptivity and non-adaptivity, so we ignore this setting. We will further utilize our proofs from the previous section to show that we can further restrict this set.

Lemma 6.2. *Let $\varepsilon \geq 0$. If $\varepsilon_g \in (-k\varepsilon, k\varepsilon)$, then*

$$t_{OPT}(\mathcal{M}_{BR}^k, \varepsilon_g) \subseteq \left\{ \frac{\varepsilon_g + (\ell+1)\varepsilon}{k+1} : \ell \in \{0, \cdots, k-1\} \right\} \cap (0, \varepsilon).$$

Proof. From Lemma 5.6 and Corollary 5.3 we can restrict our consideration to values of $t \in [0, \varepsilon]$ that maximize $F_\ell(t, \varepsilon_g)$ for some $\ell \in [k]$. Furthermore, $F_\ell(t, \varepsilon_g)$ can only maximized at the endpoints of the interval or whenever $\frac{\partial F_\ell(t,\varepsilon_g)}{\partial t} = 0$. Thus, from Corollary 5.4 we have

$$t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g) \subseteq \left\{ t_\ell^* = \frac{\varepsilon_g + (\ell+1)\varepsilon}{k+1} : \ell \in \{0, k\} \right\} \cup \{0, \varepsilon\}.$$

By definition, we can remove all values outside of $[0, \varepsilon]$, so it then suffices to show that we can also remove $\{0, \varepsilon, t_k^*\}$. Note that $p_t = 1$ when $t = 0$ and $p_t = 0$ when $t = \varepsilon$ and recall $\delta^k(t, \varepsilon_g)$ from (6), so it is straightforward to verify that $\delta^k(0, \varepsilon_g) = \delta^k(\varepsilon, \varepsilon_g) = \max\{1 - e^{\varepsilon_g}, 0\}$ for any ε_g. In the proof of Lemma 5.3, we showed that $\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g) > 0$ and $\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g) > 1 - e^{\varepsilon_g}$ when $\varepsilon_g \in (-k\varepsilon, k\varepsilon)$, which implies $0, \varepsilon \notin t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g)$.

It then suffices to show $t_k^* \notin t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g)$. If $\varepsilon_g > 0$, then $t_k^* > \varepsilon$, so we only need to consider $\varepsilon_g \leq 0$. Note that $kt_k^* = k(\frac{\varepsilon_g}{k+1} + \varepsilon)$, so for any $i \leq k$ we have $kt_k^* - i\varepsilon \geq \frac{k}{k+1}\varepsilon_g$ which implies

$$\max\left\{ e^{kt_k^* - i\varepsilon} - e^{\varepsilon_g}, 0 \right\} = e^{kt_k^* - i\varepsilon} - e^{\varepsilon_g}.$$

Therefore, $\delta^k(t_k^*, \varepsilon_g) = 1 - e^{\varepsilon_g}$ and from above we know $\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g) > 1 - e^{\varepsilon_g}$ when $\varepsilon_g \in (-k\varepsilon, k\varepsilon)$, which implies $t_k^* \notin t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g)$ as desired. □

We now want to show that we can write the optimal nonadaptive composition in a similar form as the adaptive composition. This recursive formulation will then fix a value t throughout the recursion and $\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g)$ is then just the maximum value of this recursion over all $t \in [0, \varepsilon]$.

Corollary 6.1. *For $k \geq 1$ and for $\delta^k(t, \varepsilon_g)$ from (6), we have $\delta^0(t, \varepsilon_g) = \max\{1 - e^{\varepsilon_g}, 0\}$ and*

$$\delta^k(t, \varepsilon_g) = q_t \delta^{k-1}(t, \varepsilon_g - t) + (1 - q_t)\delta^{k-1}(t, \varepsilon_g + \varepsilon - t).$$

We relegate the proof of this corollary to Appendix C. Now that the formulations are similar, we show the intuitive fact that if at any point in the recursion either it is the case that either 1) switching the value of t, or 2) switching to the adaptive setting, will strictly increase that $\delta_{\texttt{OPT}}$ then there must be a gap between the nonadaptive and adaptive setting.

Lemma 6.3. *Fix the individual privacy parameter $\varepsilon > 0$, some global privacy parameter $\varepsilon_g \in \mathbb{R}$ and $k \geq 2$, along with some $t \in t_{OPT}(\mathcal{M}_{BR}^k, \varepsilon_g)$, if there exist $0 \leq \ell' \leq \ell < k$ such that either $\delta_{OPT}(\mathcal{M}_{BR}^{k-\ell}, \varepsilon_g - \ell t + \ell'\varepsilon) < \delta_{OPT}(\mathcal{A}_{BR}^{k-\ell}, \varepsilon_g - \ell t + \ell'\varepsilon)$ or $t \notin t_{OPT}(\mathcal{M}_{BR}^{k-\ell}, \varepsilon_g - \ell t + \ell'\varepsilon)$, then we must have*

$$\delta_{OPT}(\mathcal{M}_{BR}^k, \varepsilon_g) < \delta_{OPT}(\mathcal{A}_{BR}^k, \varepsilon_g).$$

This lemma will actually require quite a bit of technical detail, so we instead give a proof in Appendix C. With this property and our characterization of $t_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^k, \varepsilon_g)$, we now show that there is a gap for the base case of $k = 2$.

Lemma 6.4. *For any $\varepsilon_g \in (-\varepsilon/2, \varepsilon/2)$ we have*

$$\delta_{OPT}(\mathcal{M}_{BR}^2, \varepsilon_g) < \delta_{OPT}(\mathcal{A}_{BR}^2, \varepsilon_g)$$

Proof. From Lemma 6.2, we know that there exists and $\ell \in \{0,1\}$ such that $t_\ell = \frac{\varepsilon_g+(\ell+1)\varepsilon}{3} \in t_{\texttt{OPT}}(\mathcal{M}^2_{\texttt{BR}}, \varepsilon_g)$. Furthermore, if both $\varepsilon_g - t_\ell$ and $\varepsilon_g - t_\ell + \varepsilon$ are in $(-\varepsilon, \varepsilon)$, then we also must have $t_{\texttt{OPT}}(\mathcal{M}^1_{\texttt{BR}}, \varepsilon_g - t_\ell) = \frac{\varepsilon_g - t_\ell + \varepsilon}{2}$ and $t_{\texttt{OPT}}(\mathcal{M}^1_{\texttt{BR}}, \varepsilon_g - t_\ell + \varepsilon) = \frac{\varepsilon_g - t_\ell + 2\varepsilon}{2}$ which implies $t_{\texttt{OPT}}(\mathcal{M}^1_{\texttt{BR}}, \varepsilon_g - t_\ell) \neq t_{\texttt{OPT}}(\mathcal{M}^1_{\texttt{BR}}, \varepsilon_g - t_\ell + \varepsilon)$.

Therefore, by Lemma 6.3 it suffices to show that both $\varepsilon_g - t_\ell$ and $\varepsilon_g - t_\ell + \varepsilon$ are in $(-\varepsilon, \varepsilon)$, which is equivalent to showing $\varepsilon_g - t_\ell \in (-\varepsilon, 0)$. Plugging in for t_ℓ we then have

$$\varepsilon_g - \frac{\varepsilon_g + (\ell+1)\varepsilon}{3} \in (-\varepsilon, 0) \qquad \Leftrightarrow \qquad \varepsilon_g \in \left(\frac{(\ell-2)\varepsilon}{2}, \frac{(\ell+1)\varepsilon}{2}\right)$$

which holds for $\ell \in \{0,1\}$ by our assumption that $\varepsilon_g \in (-\varepsilon/2, \varepsilon/2)$. □

We will then apply this base case to the more general case for certain conditions by applying Lemma 6.3.

Lemma 6.5. *Given some $\varepsilon_g \in (-k\varepsilon, k\varepsilon)$ and $t \in t_{OPT}(\mathcal{M}^k_{BR}, \varepsilon_g)$. For $k \geq 4$, if $\varepsilon_g - (k-2)t < \varepsilon/2$ and $\varepsilon_g - (k-2)t + (k-2)\varepsilon > -\varepsilon/2$, then*

$$\delta_{OPT}(\mathcal{M}^k_{BR}, \varepsilon_g) < \delta_{OPT}(\mathcal{A}^k_{BR}, \varepsilon_g).$$

We relegate the proof of this lemma to Appendix C and will use this to show our desired result.

Lemma 6.6. *For any $\varepsilon_g \in [-(k-3)\varepsilon, (k-3)\varepsilon]$ and $k \geq 4$ we have*

$$\delta_{OPT}(\mathcal{M}^k_{BR}, \varepsilon_g) < \delta_{OPT}(\mathcal{A}^k_{BR}, \varepsilon_g).$$

Proof. We will prove for $\varepsilon_g \in [0, (k-3)\varepsilon]$ and the case of $\varepsilon_g \in [-(k-3)\varepsilon, 0]$ follows symmetrically. From Lemma 6.2 we know that for any $t \in t_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)$ we must have $t = \frac{\varepsilon_g+(\ell+1)\varepsilon}{k+1}$ for some $0 \leq \ell \leq k-1$. The general idea will then be to show that for any $t_\ell = \frac{\varepsilon_g+(\ell+1)\varepsilon}{k+1}$, if $t_\ell \in t_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)$, then $\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) > \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)$. We will split this into three cases.

Case I: ($\ell \geq 2$) For this setting, we want to show that we can apply Lemma 6.5 where we know $\varepsilon_g + (k-2)(\varepsilon - t) \geq -\varepsilon/2$ for any t because we are assuming $\varepsilon_g \geq 0$. It then suffices to show that $\varepsilon_g - (k-2)t_\ell < \varepsilon/2$. Plugging in for t_ℓ we have

$$\varepsilon_g - (k-2)t_\ell < \varepsilon/2 \qquad \Leftrightarrow \qquad 6\varepsilon_g < (2(k-2)(\ell+1) + (k+1))\,\varepsilon.$$

By assumption, we know $\varepsilon_g \leq (k-3)\varepsilon$, so for $\ell \geq 2$, we have

$$6\varepsilon_g \leq 6(k-3)\varepsilon < (7k-11)\varepsilon \leq (2(k-2)(\ell+1) + (k+1))\,\varepsilon.$$

and therefore $\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) > \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)$ by Lemma 6.5.

Case II: ($\ell = 0$) For this setting we have $t_0 = \frac{\varepsilon_g+\varepsilon}{k+1}$. By our assumption that $\varepsilon_g \in [-(k-3)\varepsilon, (k-3)\varepsilon]$, we must have $\varepsilon_g + \varepsilon - t_0 \in (-(k-1)\varepsilon, (k-1)\varepsilon)$. From Lemma 6.2 we then know $t_{\texttt{OPT}}(\mathcal{M}^{k-1}_{\texttt{BR}}, \varepsilon_g + \varepsilon - t_0) \subseteq \left\{\frac{\varepsilon_g+\varepsilon-t_0+(\ell'+1)\varepsilon}{k} : \ell' \in \{0, k-2\}\right\}$. We further see that for any $\ell' \geq 0$,

$$\frac{\varepsilon_g + \varepsilon}{k+1} < \frac{\varepsilon_g + \varepsilon}{k} \leq \frac{\varepsilon_g + \varepsilon - t_0 + (\ell'+1)\varepsilon}{k}.$$

This implies $t_0 \notin t_{\texttt{OPT}}(\mathcal{M}^{k-1}_{\texttt{BR}}, \varepsilon_g + \varepsilon - t_0)$ and so $\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) > \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)$ by Lemma 6.3.

Case III: $(\ell = 1)$ This will follow from the same argument as the previous case. For this setting we have $t_1 = \frac{\varepsilon_g + 2\varepsilon}{k+1}$. Once again, we use our more restrictive assumption that $\varepsilon_g \in [0, (k-3)\varepsilon]$, and therefore $\varepsilon_g + 2(\varepsilon - t_1) \geq -(k-2)\varepsilon$. Furthermore, we have

$$\varepsilon_g + 2\left(\varepsilon - \frac{\varepsilon_g + 2\varepsilon}{k+1}\right) = \frac{k-1}{k+1}(\varepsilon_g + 2\varepsilon) \leq \frac{(k-1)^2}{k+1}\varepsilon < (k-2)\varepsilon$$

where the last step follows because $(k-1)^2 < (k+1)(k-2)$ for $k > 1$. Thus $\varepsilon_g + 2(\varepsilon - t_1) \in (-(k-2)\varepsilon, (k-2)\varepsilon)$ and by Lemma 6.2, $t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-2}, \varepsilon_g + 2(\varepsilon - t_1)) \subseteq \left\{\frac{\varepsilon_g + 2(\varepsilon - t_1) + (\ell''+1)\varepsilon}{k-1} : \ell'' \in \{0, k-3\}\right\}$. It then follows that

$$\frac{\varepsilon_g + 2\varepsilon}{k+1} = \frac{\varepsilon_g + 2(\varepsilon - t_1)}{k-1} < \frac{\varepsilon_g + 2(\varepsilon - t_1) + (\ell''+1)\varepsilon}{k-1}$$

for any $\ell'' \geq 0$. This implies $t_1 \notin t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-2}, \varepsilon_g + 2(\varepsilon - t_1))$ and so $\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^k, \varepsilon_g) > \delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^k, \varepsilon_g)$ by Lemma 6.3.

□

6.2 Settings for equivalent adaptive and nonadaptive optimal composition

We also want to show that there is not a gap between adaptive and nonadaptive composition even in the non-trivial setting. More specifically, we will show that there is a gap not only when $\varepsilon_g \notin (-k\varepsilon, k\varepsilon)$ and basic composition can be applied. We first show that there is no gap for the trivial setting and then will extend this a bit.

Lemma 6.7. *For any $\varepsilon > 0$ and $\varepsilon_g \geq k\varepsilon$, we have $\delta_{OPT}(\mathcal{A}_{BR}^k, \varepsilon_g) = 0$, and for any $\varepsilon_g \leq -k\varepsilon$, we have $\delta_{OPT}(\mathcal{A}_{BR}^k, \varepsilon_g) = 1 - e^{\varepsilon_g}$.*

We leave the proof of the trivial setting to Appendix C. The basic idea for extending this interval a bit further will simply be to consider the case in which only one outcome can produce a positive probability, or equivalently, all but one outcome can produce a positive probability.

Lemma 6.8. *For any $\varepsilon > 0$ and $\varepsilon_g \geq (k-1)\varepsilon$ for $k \geq 1$, we have*

$$\delta_{OPT}(\mathcal{A}_{BR}^k, \varepsilon_g) = \sup_{\mathbf{t} \in [0,\varepsilon]^k} \left\{ \left(\prod_{i=1}^k q_{t_i} \right) \max\{1 - e^{\varepsilon_g - \sum t_i}, 0\} \right\}$$

Proof. We show this inductively. For $k = 1$, if $\varepsilon_g \geq 0$, then for any $t \in [0, \varepsilon]$, we must have $\varepsilon_g + \varepsilon - t \geq 0$ and $\max\{1 - e^{\varepsilon_g + \varepsilon - t}, 0\} = 0$. This then implies

$$\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^1, \varepsilon_g) = \sup_{t \in [0,\varepsilon]} q_t \max\{1 - e^{\varepsilon_g - t}, 0\}.$$

The inductive step for $k \geq 2$ follows equivalently, where for any $t \in [0, \varepsilon]$, we must have $\varepsilon_g + \varepsilon - t \geq (k-1)\varepsilon$, so from Lemma 6.7, we have

$$\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^k, \varepsilon_g) = \sup_{t \in [0,\varepsilon]} q_t \delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{k-1}, \varepsilon_g - t)$$

and we can then apply our inductive hypothesis because $k - 1 \geq 1$ and $\varepsilon_g - t \geq (k-2)\varepsilon$, which then gives the desired claim. □

For completeness, we also consider the symmetric case where ε_g can be negative, but leave the proof to Appendix C.

Lemma 6.9. *For any $\varepsilon > 0$ and $\varepsilon_g \leq -(k-1)\varepsilon$ with $k \geq 1$, we have*

$$\delta_{\texttt{OPT}}(\mathcal{A}^k_{BR}, \varepsilon_g) = 1 - e^{\varepsilon_g} + \sup_{\mathbf{t}\in[0,\varepsilon]^k} \left\{ \left(\prod_{i=1}^{k} (1 - q_{t_i}) \right) \left(e^{\varepsilon_g + k\varepsilon - \sum t_i} - 1 \right) \right\}.$$

We then have the following result that together with Lemma 6.6 covers almost all choices of $\varepsilon_g \in \mathbb{R}$.

Lemma 6.10. *For any $\varepsilon_g \geq (k-1)\varepsilon$ or $\varepsilon_g \leq -(k-1)\varepsilon$ we have*

$$\delta_{\texttt{OPT}}(\mathcal{A}^k_{BR}, \varepsilon_g) = \delta_{\texttt{OPT}}(\mathcal{M}^k_{BR}, \varepsilon_g)$$

Proof. Each case can be proven directly following similar reasoning as in Lemmas 6.8 and 6.9. But, we can more easily point out that in both cases, Lemmas 6.8 and 6.9 imply that the choices of $t_1, \ldots, t_k$ are not adaptively made, so we must have $\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) = \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)$. □

7 Improved and efficient adaptive composition bounds

Although we have presented the optimal composition bound in (6.1), directly computing it is intractable. We then aim to bound the privacy loss with computationally efficient bounds that improve on previous work. Given our formulation of the privacy loss in terms of a summation of individual generalized randomized response privacy loss variables, we then follow a similar analysis to concentration inequalities, e.g. Azuma-Hoeffding bounds, by bounding the moments of the privacy loss. We now present the main result of this section.

Theorem 5. *Let $\overrightarrow{\mathcal{M}} := (\mathcal{M}_1, \mathcal{M}_2, \cdots, \mathcal{M}_k)$ each $\mathcal{M}_i$ is the class of ε_i-BR mechanisms. We then have that $\overrightarrow{\mathcal{M}}$ is $(\varepsilon_g, \delta_g(\varepsilon_g))$-DP under k-fold adaptive composition for any $\varepsilon_g \geq 0$ where we define $h_\varepsilon(\lambda) := \sup_{t\in[0,\varepsilon]} \lambda(\varepsilon - t) + \ln\left(1 + p_{\varepsilon,t}(e^{-\lambda\varepsilon} - 1)\right)$ with $p_{\varepsilon,t} = \frac{e^{-t} - e^{-\varepsilon}}{1 - e^{-\varepsilon}}$ and*

$$\delta_g(\varepsilon_g) = \inf_{\lambda > 0} \mathrm{e}^{-\lambda\varepsilon_g + \sum_i h_{\varepsilon_i}(\lambda)}.$$

By Corollary 4.1 we can assume all the component mechanisms are generalized randomized response. Let L_i be log likelihood ratio in the i-th term in the total privacy loss and we will write $t_i = t_i(y_1, \cdots, y_{i-1})$

$$\begin{aligned} L_i(y_1, \ldots, y_i) &= \ln \frac{\Pr[\texttt{RR}_{\varepsilon_i, t_i}(1) = y_i | y_1, \cdots, y_{i-1}]}{\Pr[\texttt{RR}_{\varepsilon_i, t_i}(0) = y_i | y_1, \cdots, y_{i-1}]} \\ &= \begin{cases} t_i, & \text{if } y_i = 1, \\ t_i - \varepsilon_i, & \text{if } y_i = 0. \end{cases} \end{aligned}$$

Recall from Corollary 4.1 that we need only consider deterministic adversaries $\mathcal{A} = (\emptyset, \mathcal{D})$ and the class of mechanisms $\overrightarrow{\mathcal{M}}$ to be the class of generalized randomized response mechanisms $\overrightarrow{\mathcal{RR}}$. For simplicity, let P be the output distribution for $\texttt{AdComp}(\mathcal{A}, \overrightarrow{\mathcal{M}}, b)$ and Q be the output distribution

for $\texttt{AdComp}(\mathcal{A}, \overrightarrow{\mathcal{M}}, 1-b)$ on $\{0,1\}^k$ and $L = \sum_{i=1}^k L_i = \ln \frac{Q}{P}$ be the log likelihood ratio of the composed mechanism. Then we have

$$\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{1:k}, \varepsilon_g) = Q[L > \varepsilon_g] - e^{\varepsilon_g} P[L > \varepsilon_g].$$

The classical method introduced in Dwork et al. [9] make two approximations: (1) ignore the negative term $-e^{\varepsilon_g} P[L > \varepsilon_g]$ and (2) use moment generating function of L to bound the tail probability $Q[L > \varepsilon_g]$. We follow (1) and improve on (2) with the help of the reduction in Corollary 4.1. We present the proof of Theorem 5 to point out the stages in the analysis where other approaches used weaker bounds. The initial steps remain consistent:

$$\begin{aligned}
\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{1:k}, \varepsilon_g) &= Q[L > \varepsilon_g] - e^{\varepsilon_g} P[L > \varepsilon_g] \\
&\leqslant Q[L > \varepsilon_g] \\
&= \Pr\Big[\sum L_i > \varepsilon_g\Big] \\
&\leqslant \inf_{\lambda > 0} \Pr\left[e^{\lambda \sum L_i} > e^{\lambda \varepsilon_g}\right] \\
&\leqslant \inf_{\lambda > 0} e^{-\lambda \varepsilon_g} \cdot \mathbb{E}\left[e^{\lambda \sum L_i}\right].
\end{aligned}$$

With a standard conditional probability argument we have the following result.

Lemma 7.1. *If there is a function $U_i : (0, +\infty) \to \mathbb{R}$ such that for each $i = 1, 2, \ldots, k$ the following holds for any arbitrary outcomes $y_1, \ldots, y_{i-1}$ of the previous generalized randomized response mechanisms,*

$$\mathbb{E}_Q[e^{\lambda L_i} \mid y_1, \ldots, y_{i-1}] \leqslant e^{U_i(\lambda)},$$

then the following holds for any $\lambda > 0$,

$$\delta_{\mathit{OPT}}(\mathcal{A}_{\mathit{BR}}^{1:k}, \varepsilon_g) \leqslant e^{-(\lambda \varepsilon_g - \sum_i U_i(\lambda))}.$$

Different bounds correspond to different choices of $U_i(\lambda)$ in Lemma 7.1, which result in different bounds on $\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{1:k}, \varepsilon_g)$. For example, both Dwork et al. [9] and Durfee and Rogers [6] utilize the following lemma:

Lemma 7.2 (Hoeffding's lemma)**.** *If a random variable $X \in [a, b]$ then $\ln \mathbb{E}[e^{\lambda X}] \leqslant \frac{1}{8}(b-a)^2\lambda^2 + \lambda \mathbb{E}X$.*

We now walk through the following comparisons with previous work to highlight our improvement. Dwork et al. [9] only uses the fact that $L_i \in [-\varepsilon_i, \varepsilon_i]$ (which is weaker than ε-BR). It implies

(a) $\ln \mathbb{E}_Q[e^{\lambda L_i} \mid y_1, \ldots, y_{i-1}] \leqslant \frac{1}{2}\varepsilon_i^2\lambda^2 + \lambda \mathbb{E}_Q[L_i \mid y_1, \ldots, y_{i-1}]$

(b) $\mathbb{E}_Q[L_i \mid y_1, \ldots, y_{i-1}] \leqslant \varepsilon_i \tanh \frac{\varepsilon_i}{2} \leqslant \frac{1}{2}\varepsilon_i^2$.

For part (b), Dwork et al. [9] used a much rougher estimate. The $\frac{1}{2}\varepsilon_i^2$ upper bound appears in [3]. For the most refined bound in terms of hyperbolic tangent function, readers can refer to Lemma D.8 in [4].

Combining both (a) and (b), we have $U_i(\lambda) = \frac{1}{2}\varepsilon_i^2(\lambda^2 + \lambda)$, which we refer to as "Improved DRV10" in Figure 2.

Using the bounded range property from Durfee and Rogers [6], we know that for ε-BR there is a $t_i \in [0, \varepsilon_i]$ such that $a = t_i - \varepsilon_i, b = t_i$ in Hoeffding's lemma. A similar argument yields $U_i(\lambda) = \frac{1}{2}\varepsilon_i^2(\frac{1}{4}\lambda^2 + \lambda)$, which we label as "DR19" in Figure 2.

A straightforward improvement could come from a finer treatment of (b). By definition of L_i,

$$\begin{aligned}
\mathbb{E}_Q[L_i \mid y_1, \ldots, y_{i-1}] &= \mathrm{KL}\big(\mathrm{Bern}(q_{\varepsilon_i,t_i}) \| \mathrm{Bern}(p_{\varepsilon_i,t_i})\big) \\
&= q_{\varepsilon_i,t_i} \cdot \ln \frac{q_{\varepsilon_i,t_i}}{p_{\varepsilon_i,t_i}} + (1 - q_{\varepsilon_i,t_i}) \cdot \ln \frac{1 - q_{\varepsilon_i,t_i}}{1 - p_{\varepsilon_i,t_i}} \\
&= t_i q_{\varepsilon_i,t_i} + (t_i - \varepsilon_i)(1 - q_{\varepsilon_i,t_i}) \\
&= t_i - \frac{\varepsilon_i}{e^{\varepsilon_i} - 1}(e^{t_i} - 1).
\end{aligned}$$

A bit of calculus shows the above expression is maximized at $t_i = \ln \frac{e^{\varepsilon_i} - 1}{\varepsilon_i}$, and the value is

$$\mathrm{maxkl}(\varepsilon) := \frac{\varepsilon}{e^{\varepsilon} - 1} - 1 - \ln \frac{\varepsilon}{e^{\varepsilon} - 1}.$$

That is, we have replaced (b) with

(b′) $\mathbb{E}_Q[L_i \mid y_1, \ldots, y_{i-1}] \leqslant \mathrm{maxkl}(\varepsilon_i)$.

Combining (a) and (b′), we can use $U_i(\lambda) = \frac{1}{8}\varepsilon_i^2\lambda^2 + \lambda \cdot \mathrm{maxkl}(\varepsilon_i)$, which we label as "KL-improved DR19" in Figure 2. This observation on the expectation together with the Durfee and Rogers [6] bound that uses Azuma-Hoeffding, but with a weaker bound on the expectation term, we have the following result.

Corollary 3.1. *Let $\overrightarrow{\mathcal{M}} := (\mathcal{M}_1, \mathcal{M}_2, \cdots, \mathcal{M}_k)$ where each $\mathcal{M}_i$ is the class of ε_i-BR mechanisms. We then have that $\overrightarrow{\mathcal{M}}$ is $(\varepsilon_g(\delta_g), \delta_g)$-DP under k-fold adaptive composition for any $\delta_g \geq 0$ where*

$$\varepsilon_g(\delta) = \min\left\{ \sum_{i=1}^k \varepsilon_i, \sum_{i=1}^k \left(\frac{\varepsilon_i}{1 - e^{-\varepsilon_i}} - 1 - \ln\left(\frac{\varepsilon_i}{1 - e^{-\varepsilon_i}} \right) \right) + \sqrt{\frac{1}{2} \sum_{i=1}^k \varepsilon_i^2 \ln(1/\delta)} \right\}.$$

Instead of trying to come up with analytic, closed form upper bounds, we directly compute $\mathbb{E}_Q[e^{\lambda L_i} \mid y_1, \ldots, y_{i-1}]$, resorting to numerical tools when necessary. Recall that $p_{\varepsilon,t} = \frac{e^{-t} - e^{-\varepsilon}}{1 - e^{-\varepsilon}}$ and $q_{\varepsilon,t} = e^t p_{\varepsilon,t} = \frac{1 - e^{t-\varepsilon}}{1 - e^{-\varepsilon}}$, we then have the following result.

Lemma 7.3.

$$\mathbb{E}_Q[e^{\lambda L_i} \mid y_1, \ldots, y_{i-1}] = p_{\varepsilon_i,\varepsilon_i - t_i}^{\lambda+1} q_{\varepsilon_i,\varepsilon_i - t_i}^{-\lambda} + (1 - p_{\varepsilon_i,\varepsilon_i - t_i})^{\lambda+1} (1 - q_{\varepsilon_i,\varepsilon_i - t_i})^{-\lambda}$$

where $t_i = t_i(y_1, \cdots, y_{i-1})$.

Proof. Let P_i be the distribution for $\mathrm{Bern}(p_{\varepsilon_i,t_i})$ and Q_i be the distribution for $\mathrm{Bern}(q_{\varepsilon_i,t_i})$. Then

$$\begin{aligned}
\mathbb{E}_Q[e^{\lambda L_i} \mid y_1, \ldots, y_{i-1}] &= \int \left(\frac{Q_i}{P_i}\right)^{\lambda} \cdot Q_i \\
&= q_{\varepsilon_i,t_i}^{\lambda+1} p_{\varepsilon_i,t_i}^{-\lambda} + (1 - q_{\varepsilon_i,t_i})^{\lambda+1} (1 - p_{\varepsilon_i,t_i})^{-\lambda}
\end{aligned}$$

It is easy to verify that $q_{\varepsilon,\varepsilon-t} = 1 - p_{\varepsilon,t}$ and $p_{\varepsilon,\varepsilon-t} = 1 - q_{\varepsilon,t}$. Plugging these into the above expression yields the desired result. □

We now simplify the expression in Lemma 7.3 with the following function,

$$\begin{aligned} h_\varepsilon(\lambda) &= \sup_{t\in[0,\varepsilon]} \ln\left(p_{\varepsilon,t}^{\lambda+1} q_{\varepsilon,t}^{-\lambda} + (1-p_{\varepsilon,t})^{\lambda+1}(1-q_{\varepsilon,t})^{-\lambda}\right) \\ &= \sup_{t\in[0,\varepsilon]} \ln\left(p_{\varepsilon,t} e^{-\lambda t} + (1-p_{\varepsilon,t}) e^{-\lambda(t-\varepsilon)}\right) \\ &= \sup_{t\in[0,\varepsilon]} \lambda(\varepsilon - t) + \ln\left(1 + p_{\varepsilon,t}(e^{-\lambda\varepsilon} - 1)\right). \end{aligned}$$

The second line above makes use of the fact that

$$q_{\varepsilon,t} = e^t p_{\varepsilon,t} \text{ and } 1 - q_{\varepsilon,t} = e^{t-\varepsilon}(1 - p_{\varepsilon,t}).$$

Now it is easy to see that $U_i(\lambda)$ can be taken as $h_{\varepsilon_i}(\lambda)$, which we label as "General MGF" in Figure 2. Combining this with Lemma 7.1, we have Theorem 5.

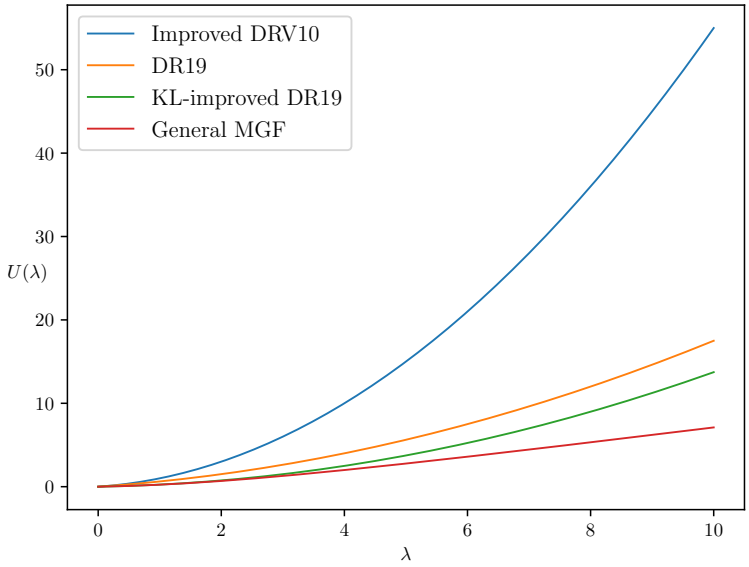

Figure 2. A unified view and comparison of composition theorems involving concentration inequalities. The figure shows graphs of different U functions (see Lemma 7.1) used in different results, such as from Dwork et al. [9] (labeled "Improved DRV10") and Durfee and Rogers [6] (labeled "DR19"). According to Lemma 7.1, smaller function U yields tighter privacy result. Theorem 5 uses the smallest U (labeled "General MGF") among all and is hence the tightest. All curves use $\varepsilon = 1$.

Numerical Issue We now point out a potential numeric issue in computing the function $h_\varepsilon(\lambda)$. Note that it can be simplified differently as

$$h_\varepsilon(\lambda) = \sup_{t\in[0,\varepsilon]} -\lambda t + \ln\left(p_{\varepsilon,t} + e^{\varepsilon\lambda}(1 - p_{\varepsilon,t})\right).$$

For comparision, the expression we use in Theorem 5 is

$$h_\varepsilon(\lambda) = \sup_{t\in[0,\varepsilon]} \lambda(\varepsilon - t) + \ln\left(1 + p_{\varepsilon,t}(e^{-\lambda\varepsilon} - 1)\right).$$

At first glance it may appear that the above two expressions are equal. However, the one used in the theorem is far more robust numerically, as in the optimization step, $\varepsilon\lambda$ can be large, which could make $e^{\varepsilon\lambda}$ beyond the range of floating point numbers.

8 Conclusion and future directions

In this work, we studied the privacy loss when composing multiple exponential mechanisms, which is a fundamental class of DP algorithms. We considered the privacy loss bounds when the exponential mechanisms can be adaptively selected at each round or when they are all selected in advance, as well as differentiated the homogeneous (all privacy parameters are the same) and the heterogeneous (privacy parameters can be different) case. We then made the connection between exponential mechanisms and the generalized randomized response mechanism to help simplify our privacy loss expressions. Although we provided formulas for each case, we only provided an efficient calculation for computing the optimal composition bound in the nonadaptive and homogeneous case. We conjecture that computing the optimal composition bound in the nonadaptive and heterogeneous case has similar hardness results as shown in Murtagh and Vadhan [14] and we leave the problem open for future work.

We then showed for the optimal homogenous composition bound that there is a separation between in the adaptive and nonadaptive case, which to our knowledge is a first of its kind result. We then provided improved and computationally efficient composition bounds for the adaptive and homogeneous case by tailoring concentration bounds for our particular setting. In order to better understand the adaptive composition bound, one potential direction for future work is to understand the asymptotics of the privacy loss bound, as $k \to \infty$. We conjecture that the asymptotic gap collapses between the optimal composition bound for the adaptive and nonadaptive cases, and leave that as future work to study. Furthermore, in the non-asymptotic setting we believe that the gap between adaptive and non-adaptive is quite small, and also leave proving a strong upper bound on this gap to future work.

Lastly, it is interesting to study composition bounds that account for different types of DP mechanisms at each round. General DP composition bounds can be used in cases where Laplace and exponential mechanisms are used, but perhaps those bounds can be improved with composition that accounts for exponential mechanisms and Laplace mechanisms separately. We leave this as an interesting direction of future work.

9 Acknowledgements

We thank our colleagues Reza Hosseini, Krishnaram Kenthapadi, Sean Peng, and Subbu Subramaniam for their helpful feedback and comments.

References

[1] M. Abadi, A. Chu, I. Goodfellow, B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *23rd ACM Conference on Computer and Communications Security (ACM CCS)*, pages 308–318, 2016. URL https://arxiv.org/abs/1607.00133.

[2] D. Blackwell. Comparison of experiments. Technical report, HOWARD UNIVERSITY Washington United States, 1950.

[3] M. Bun and T. Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference (TCC)*, pages 635–658, 2016.

[4] J. Dong, A. Roth, and W. J. Su. Gaussian differential privacy. *CoRR*, abs/1905.02383, 2019. URL http://arxiv.org/abs/1905.02383.

[5] J. Duchi and R. Rogers. Lower bounds for locally private estimation via communication complexity. In A. Beygelzimer and D. Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 1161–1191, Phoenix, USA, 25–28 Jun 2019. PMLR. URL http://proceedings.mlr.press/v99/duchi19a.html.

[6] D. Durfee and R. Rogers. Practical differentially private top-k selection with pay-what-you-get composition. *CoRR*, abs/1905.04273, 2019. URL http://arxiv.org/abs/1905.04273.

[7] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology (EUROCRYPT 2006)*, 2006.

[8] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Theory of Cryptography Conference*, pages 265–284, 2006.

[9] C. Dwork, G. N. Rothblum, and S. P. Vadhan. Boosting and differential privacy. In *51st Annual Symposium on Foundations of Computer Science*, pages 51–60, 2010.

[10] M. Joseph, J. Mao, S. Neel, and A. Roth. The role of interactivity in local differential privacy. *CoRR*, abs/1904.03564, 2019. URL http://arxiv.org/abs/1904.03564.

[11] P. Kairouz, S. Oh, and P. Viswanath. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 63(6):4037–4049, June 2017. ISSN 0018-9448. doi: 10.1109/TIT.2017.2685505.

[12] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

[13] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *48th Annual Symposium on Foundations of Computer Science*, 2007.

[14] J. Murtagh and S. Vadhan. The complexity of computing the optimal composition of differential privacy. In *Proceedings, Part I, of the 13th International Conference on Theory of Cryptography - Volume 9562*, TCC 2016-A, pages 157–175, Berlin, Heidelberg, 2016.

Springer-Verlag. ISBN 978-3-662-49095-2. doi: 10.1007/978-3-662-49096-9_7. URL https://doi.org/10.1007/978-3-662-49096-9_7.

[15] S. Oh and P. Viswanath. The composition theorem for differential privacy. *arXiv:1311.0776 [cs.DS]*, 2013.

[16] R. M. Rogers, A. Roth, J. Ullman, and S. P. Vadhan. Privacy odometers and filters: Pay-as-you-go composition. *CoRR*, abs/1605.08294, 2016. URL http://arxiv.org/abs/1605.08294.

[17] R. M. Rogers, A. Roth, J. Ullman, and S. P. Vadhan. Privacy odometers and filters: Pay-as-you-go composition. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1921–1929, 2016. URL http://papers.nips.cc/paper/6170-privacy-odometers-and-filters-pay-as-you-go-composition.

[18] A. Smith, A. Thakurta, and J. Upadhyay. Is interaction necessary for distributed private learning? In *IEEE Symposium on Security and Privacy*, 2017.

[19] S. Warner. Randomized response: a survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

A Proof of Lemma 4.1

In order to use this interpretation, we will need to first establish some notation. For a pair of probability distributions P and Q on a common probability space Ω, its trade-off function [4] describes the hardness of the hypothesis testing problem $H_0 : P$ vs $H_1 : Q$. Let $E \subseteq \Omega$ be an arbitrary rejection region and

$$\begin{aligned}\alpha_E &= P[E] \\ \beta_E &= 1 - Q[E]\end{aligned}$$

be the type I and type II errors of the test E respectively. Fix a level α_0 and let E run over all test with type I error at most α_0, the minimal type II error is

$$\inf\{\beta_E : E \text{ is a rejection region s.t. } \alpha_E \leqslant \alpha_0\}.$$

This correspondence of α_0 to the minimal type II error defines a function from $[0,1]$ to $[0,1]$. We will call this function $T(P,Q)$. Formally,

$$\begin{aligned}T(P,Q) : [0,1] &\to [0,1] \\ \alpha_0 &\mapsto \inf\{\beta_E : \alpha_E \leqslant \alpha_0\}\end{aligned}$$

For our proof, we will use this function T and apply Blackwell's theorem ([2], Theorem 10). The following form is taken from [4].

Theorem 6. *Let P, Q be probability distributions on Y and P', Q' be probability distributions on Z. The following two statements are equivalent:*

(a) $T(P,Q) \leqslant T(P',Q')$.

(b) There exists a randomized algorithm $\mathrm{Proc} : Y \to Z$ *such that* $\mathrm{Proc}(P) = P', \mathrm{Proc}(Q) = Q'$.

We now prove that we can post-process the generalized random response to simulate any BR mechanism on neighboring inputs.

Proof of Lemma 4.1. Let P be the outcome distribution of $M(x^0)$ and Q be the outcome distribution of $M(x^1)$. By Corollary 2.1, we know there exists some $t \in [0, \varepsilon]$ such that

$$t - \varepsilon \leqslant \ln \frac{Q(y)}{P(y)} \leqslant t.$$

Equivalently, for any event $E \subseteq \mathcal{Y}$,

$$\mathrm{e}^{t-\varepsilon} P[E] \leqslant Q[E] \leqslant \mathrm{e}^t P[E]. \tag{9}$$

Applying the same rule for the complement event E^c, we have

$$\mathrm{e}^{t-\varepsilon} P[E^c] \leqslant Q[E^c] \leqslant \mathrm{e}^t P[E^c]. \tag{10}$$

The second inequality of (9) and the first inequality of (10) imply

$$1 - \beta_E \leqslant \mathrm{e}^t \alpha_E, \quad \mathrm{e}^{t-\varepsilon}(1 - \alpha_E) \leqslant \beta_E. \tag{11}$$

Let the piece-wise linear function $l_{t,\varepsilon} : [0,1] \to [0,1]$ be defined as

$$l_{t,\varepsilon}(x) = \max\{1 - \mathrm{e}^t x, \mathrm{e}^{t-\varepsilon}(1-x)\}.$$

It's easy to see that (11) implies $T(P,Q) \geqslant l_{t,\varepsilon}$ pointwise in $[0,1]$.

Furthermore, it is straightforward to verify that $l_{t,\varepsilon} \equiv T(\mathtt{RR}_{\varepsilon,t}(0), \mathtt{RR}_{\varepsilon,t}(1))$ because the respective inequalities in (11) are tight for $E = \{0\}$ and $E = \{1\}$, respectively. Therefore, there must be a $t = t(M, x^0, x^1)$ such that

$$T\big(M(x^0), M(x^1)\big) \geqslant T\big(\mathtt{RR}_{\varepsilon,t}(0), \mathtt{RR}_{\varepsilon,t}(1)\big).$$

Applying Theorem 6 then gives our desired claim. □

B Omitted Proofs from Section 5

We provide here the proofs from Section 5 that were omitted.

B.1 Proof of Lemma 5.4

This lemma will be proven in two main sublemmas. First, we show that it holds for $k = 2$, then we show how we can reduce the general case to $k = 2$ by conditioning outcomes other than the first and second terms.

Lemma B.1. *For any $\varepsilon_g \in \mathbb{R}$ and $t_1, t_2 \in [0, \varepsilon]$*

$$\delta((t_1,t_2),\varepsilon_g) \leq \delta\left(\left(\frac{t_1+t_2}{2}, \frac{t_1+t_2}{2}\right), \varepsilon_g\right)$$

Further, the inequality is strict whenever $\varepsilon_g < t_1 + t_2 < \varepsilon_g + 2\varepsilon$ and $t_1 \neq t_2$.

Proof. Using the fact that $q_t = e^t p_t$ and $(1-q_t) = e^{t-\varepsilon}(1-p_t)$, we rewrite

$$\delta((t_1,t_2),\varepsilon_g) = \sum_{S\subseteq\{1,2\}} \prod_{i\notin S} p_{t_i} \prod_{i\in S}(1-p_{t_i}) \max\left\{e^{t_1+t_2-|S|\varepsilon} - e^{\varepsilon_g}, 0\right\}$$

We will then prove our desired inequality by considering four cases.

Case I ($t_1 + t_2 \leq \varepsilon_g$): This implies that $\max\{e^{t_1+t_2-|S|\varepsilon} - e^{\varepsilon_g}, 0\} = 0$ for any subset S and

$$\delta((t_1,t_2),\varepsilon_g) = \delta\left(\left(\frac{t_1+t_2}{2}, \frac{t_1+t_2}{2}\right), \varepsilon_g\right) = 0.$$

Case II ($t_1 + t_2 \geq \varepsilon_g + 2\varepsilon$): This implies $\max\{e^{t_1+t_2-|S|\varepsilon} - e^{\varepsilon_g}, 0\} = e^{t_1+t_2-|S|\varepsilon} - e^{\varepsilon_g}$ for any S, which gives

$$\delta((t_1,t_2),\varepsilon_g) = \sum_{S\subseteq\{1,2\}} \left(\prod_{i\notin S} q_{t_i} \prod_{i\in S}(1-q_{t_i}) - e^{\varepsilon_g} \prod_{i\notin S} p_{t_i} \prod_{i\in S}(1-p_{t_i})\right) = 1 - e^{\varepsilon_g}$$

and equivalently holds for $\delta(\left(\frac{t_1+t_2}{2}, \frac{t_1+t_2}{2}\right), \varepsilon_g)$.

Case III ($\varepsilon_g < t_1 + t_2 \leq \varepsilon_g + \varepsilon$)**:** This implies that $\max\{e^{t_1+t_2-|S|\varepsilon} - e^{\varepsilon_g}, 0\} = 0$ for any S such that $|S| > 0$. Therefore,

$$\delta((t_1, t_2), \varepsilon_g) = p_{t_1} p_{t_2} \left(e^{t_1+t_2} - e^{\varepsilon_g}\right)$$

Equivalently, we have

$$\delta\left(\left(\frac{t_1+t_2}{2}, \frac{t_1+t_2}{2}\right), \varepsilon_g\right) = p^2_{\frac{t_1+t_2}{2}} \left(e^{t_1+t_2} - e^{\varepsilon_g}\right)$$

We want strict inequality for this case, so it suffices to show $p_{t_1} p_{t_2} < p^2_{\frac{t_1+t_2}{2}}$. Plugging in the explicit formula for each p_t and performing some simple algebraic manipulations gives that this is equivalent to

$$2e^{-\frac{t_1+t_2}{2}} < e^{-t_1} + e^{-t_2}$$

which holds due to the strict-convexity of the exponential function.

Case IV ($\varepsilon_g + \varepsilon \leq t_1 + t_2 < \varepsilon_g + 2\varepsilon$)**:** This implies that $\max\{e^{t_1+t_2-|S|\varepsilon} - e^{\varepsilon_g}, 0\} = 0$ when $|S| = 2$. Therefore,

$$\delta((t_1, t_2), \varepsilon_g) = p_{t_1} p_{t_2} \left(e^{t_1+t_2} - e^{\varepsilon_g}\right) + (p_{t_1}(1 - p_{t_2}) + p_{t_2}(1 - p_{t_1})) \left(e^{t_1+t_2-\varepsilon} - e^{\varepsilon_g}\right)$$

From Case II, we know

$$\sum_{S \subseteq \{1,2\}} \prod_{i \notin S} p_{t_i} \prod_{i \in S} (1 - p_{t_i}) \left(e^{t_1+t_2-|S|\varepsilon} - e^{\varepsilon_g}\right) = 1 - e^{\varepsilon_g}$$

which yields

$$\delta((t_1, t_2), \varepsilon_g) = 1 - e^{\varepsilon_g} - (1 - p_{t_1})(1 - p_{t_2}) \left(e^{t_1+t_2-2\varepsilon} - e^{\varepsilon_g}\right).$$

This equivalently holds for $\delta(\left(\frac{t_1+t_2}{2}, \frac{t_1+t_2}{2}\right), \varepsilon_g)$ and because $e^{t_1+t_2-2\varepsilon} - e^{\varepsilon_g} < 0$, we have

$$\delta((t_1, t_2), \varepsilon_g) < \delta\left(\left(\frac{t_1+t_2}{2}, \frac{t_1+t_2}{2}\right), \varepsilon_g\right) \qquad \Leftrightarrow \qquad (1 - p_{t_1})(1 - p_{t_2}) < \left(1 - p_{\frac{t_1+t_2}{2}}\right)^2.$$

Once again, we plug in the explicit formula for each p_t and perform some simple algebraic manipulations to see that this is also equivalent to

$$2e^{-\frac{t_1+t_2}{2}} < e^{-t_1} + e^{-t_2}$$

and this again holds due to the strict-convexity of the exponential function.

□

We now want to extend this to $k > 2$, which will be done by fixing an arbitrary subset of $\{3, \cdots, k\}$ and show that the inequality holds when we restrict the summation to subsets of $\{1, \cdots, k\}$ that must contain that subset of $\{3, \cdots, k\}$. This will allow for easy cancellation. We will denote $\delta_U(\mathbf{t}, \varepsilon_g, S)$ for a set $U \subseteq [k]$ and $S \subseteq U$ as

$$\delta_U(\mathbf{t}, \varepsilon_g, S) := \prod_{i \in U \setminus S} p_{t_i} \prod_{i \in S} (1 - p_{t_i}) \cdot \sum_{S' \subseteq [k] \setminus U} \max\left\{e^{\sum_{j \in U} t_j - |S|\varepsilon} \prod_{i \notin U \cup S'} q_{t_i} \prod_{i \in S'} (1 - q_{t_i}) - e^{\varepsilon_g} \prod_{i \notin U \cup S'} p_{t_i} \prod_{i \in S'} (1 - p_{t_i}), 0\right\}.$$

Claim B.1. *Let $\varepsilon_g \in \mathbb{R}$. Then for any $\mathbf{t} \in [0,\varepsilon]^k$, we have for $U = \{3, \cdots, k\}$*

$$\delta(\mathbf{t}, \varepsilon_g) = \sum_{S \subseteq U} \delta_U(\mathbf{t}, \varepsilon_g, S)$$

Proof. We fix a set $S \subseteq \{3, \cdots, k\} = U$. Using the fact that $q_t = e^t p_t$ and $(1 - q_t) = e^{t-\varepsilon}(1 - p_t)$, we have

$$\prod_{i \in U \setminus S} q_{t_i} \prod_{i \in S} (1 - q_{t_i}) = e^{t_3 + \cdots + t_k - |S|\varepsilon} \prod_{i \in U \setminus S} p_{t_i} \prod_{i \in S} (1 - p_{t_i})$$

Therefore, we also have

$$\delta_U(\mathbf{t}, \varepsilon_g, S) = \sum_{S' \subseteq \{1,2\}} \max \left\{ \prod_{i \notin S' \cup S} q_{t_i} \prod_{i \in S' \cup S} (1 - q_{t_i}) - e^{\varepsilon_g} \prod_{i \notin S' \cup S} p_{t_i} \prod_{i \in S' \cup S} (1 - p_{t_i}), 0 \right\}$$

Summing over all S we can simply rewrite this summation over all subsets of $\{1, \cdots, k\}$, giving our desired equality. □

Lemma B.2. *For any $S \subseteq \{3, ..., k\} = U$, we have the following inequality*

$$\delta_U(\mathbf{t}, \varepsilon_g, S) \leq \delta_U \left(\left(\frac{t_1 + t_2}{2}, \frac{t_1 + t_2}{2}, t_3, ..., t_k \right), \varepsilon_g, S \right)$$

Further, the inequality is strict if $\varepsilon_g < \sum_{i=1}^k t_i - |S|\varepsilon < \varepsilon_g + 2\varepsilon$ and $t_1 \neq t_2$.

Proof. We fix $S \subseteq \{3, \cdots, k\}$. Let $\varepsilon_g' = \varepsilon_g + |S| - t_3 - \cdots - t_k$, and then by cancelling non-negative like terms it suffices to show

$$\sum_{S' \subseteq \{1,2\}} \max \left\{ \prod_{i \notin S'} q_{t_i} \prod_{i \in S'} (1 - q_{t_i}) - e^{\varepsilon_g'} \prod_{i \notin S'} p_{t_i} \prod_{i \in S'} (1 - p_{t_i}), 0 \right\}$$
$$\leq \sum_{S' \subseteq \{1,2\}} \max \left\{ \prod_{i \notin S'} q_{t'} \prod_{i \in S'} (1 - q_{t'}) - e^{\varepsilon_g'} \prod_{i \notin S'} p_{t'} \prod_{i \in S'} (1 - p_{t'}), 0 \right\}$$

where $t' = \frac{t_1 + t_2}{2}$. By definition, this is then equivalent to showing

$$\delta((t_1, t_2), \varepsilon_g') \leq \delta \left(\left(\frac{t_1 + t_2}{2}, \frac{t_1 + t_2}{2} \right), \varepsilon_g' \right)$$

which follows from Lemma B.1, and the strictness follows from the fact that $\varepsilon_g' = \varepsilon_g + |S| - \sum_{j>2} t_j$. □

With these we can now prove our main convexity lemma.

Proof of Lemma 5.4. It immediately follows from Claim B.1 and Lemma B.2 that for any $\mathbf{t} \in [0,\varepsilon]^k$

$$\delta(\mathbf{t}, \varepsilon_g) \leq \delta \left(\left(\frac{t_1 + t_2}{2}, \frac{t_1 + t_2}{2}, t_3, ..., t_k \right), \varepsilon_g \right)$$

Additionally, if we assume that $t_1 \neq t_2$ and $\varepsilon_g < \sum t_i < \varepsilon_g + k\varepsilon$, then there must exist some $\ell \in [0, k-2]$ such that $\varepsilon_g + \ell\varepsilon < \sum t_i < \varepsilon_g + (\ell + 2)\varepsilon$, which implies that $\varepsilon_g < \sum t_i - \ell\varepsilon < \varepsilon_g + 2\varepsilon$. Further, we know that for any $\ell \in [0, k-2]$ there exists $S \subseteq \{3, \cdots, k\}$ such that $|S| = \ell$. Therefore, for one of these subsets the inequality is strict and the sum must be a strict inequality as well. □

B.2 Proof of Lemma 5.7

Recall that we had the following definition, for which we wanted to compute the partial derivate with respect to t.

$$F_\ell(t, \varepsilon_g) := \sum_{i=0}^{\ell} \binom{k}{i} p_t^{k-i} (1 - p_t)^i \left(e^{kt - i\varepsilon} - e^{\varepsilon_g} \right) \tag{12}$$

We further split each $F_\ell(t, \varepsilon_g)$ into the individual terms to more easily differentiate the full summation with respect to t.

$$f_\ell(t, \varepsilon_g) := \binom{k}{\ell} p_t^{k-\ell} (1 - p_t)^\ell \left(e^{kt - \ell\varepsilon} - e^{\varepsilon_g} \right)$$

In particular, giving a much simpler formulation for the partial derivative will rely upon an inductive proof, so this definition will allow an even easier comparison between $F_\ell(t, \varepsilon_g)$ and $F_{\ell+1}(t, \varepsilon_g)$ that follows immediately from the definition.

Corollary B.1. *For any $\ell \in [1, k]$*

$$F_\ell(t, \varepsilon_g) = F_{\ell-1}(t, \varepsilon_g) + f_\ell(t, \varepsilon_g)$$

We first differentiate the simplest of these expressions $F_0(t, \varepsilon_g)$, and then we will ultimately use this as the base case for proving a simplified formulation of derivative for the general case.

Lemma B.3.

$$\frac{\partial F_0(t, \varepsilon_g)}{\partial t} = k p_t^{k-1} \frac{1}{1 - e^{-\varepsilon}} \left(e^{\varepsilon_g - t} - e^{kt - \varepsilon} \right)$$

Proof. By definition

$$F_0(t, \varepsilon_g) = p_t^k \left(e^{kt} - e^{\varepsilon_g} \right) = \left(\frac{e^{-t} - e^{-\varepsilon}}{1 - e^{-\varepsilon}} \right)^k \left(e^{kt} - e^{\varepsilon_g} \right)$$

Therefore, by basic differentiation rules

$$\begin{aligned} \frac{\partial F_0(t, \varepsilon_g)}{\partial t} &= \left(-k \frac{e^{-t}}{1 - e^{-\varepsilon}} \left(\frac{e^{-t} - e^{-\varepsilon}}{1 - e^{-\varepsilon}} \right)^{k-1} \left(e^{kt} - e^{\varepsilon_g} \right) \right) + \left(\frac{e^{-t} - e^{-\varepsilon}}{1 - e^{-\varepsilon}} \right)^k k e^{kt} \\ &= k p_t^{k-1} \frac{1}{1 - e^{-\varepsilon}} \left(-e^{-t} \left(e^{kt} - e^{\varepsilon_g} \right) + \left(e^{-t} - e^{\varepsilon} \right) e^{kt} \right) \end{aligned}$$

which easily reduces to our desired term.

□

To apply an inductive claim to the general case, we will also need to evaluate the partial derivative of the last term for each sum.

Lemma B.4. *For* $1 \leq \ell \leq k$

$$\frac{\partial f_\ell(t,\varepsilon_g)}{\partial t} = \binom{k}{\ell} p_t^{k-1-\ell}(1-p_t)^{\ell-1}\left(\frac{1}{1-e^{-\varepsilon}}\right)^2\Big((k-\ell)\left(e^{\varepsilon_g - t} + e^{(k-1)t-(\ell+1)\varepsilon}\right) + \ell\left(e^{(k-1)t-\ell\varepsilon} + e^{\varepsilon_g-\varepsilon-t}\right) - k\left(e^{\varepsilon_g-2t} + e^{kt-(\ell+1)\varepsilon}\right)\Big)$$

Proof. By definition

$$f_\ell(t,\varepsilon_g) = \binom{k}{\ell} p_t^{k-\ell}(1-p_t)^\ell\left(e^{kt-\ell\varepsilon} - e^{\varepsilon_g}\right)$$

We can consider this then to instead be $f_\ell(t,\varepsilon_g) = \binom{k}{\ell} f(t)\cdot g(t)\cdot h(t)$ with $f(t) = p_t^{k-\ell}$, $g(t) = (1-p_t)^\ell$, and $h(t) = e^{kt-\ell\varepsilon} - e^{\varepsilon_g}$. Applying basic differentiation rules and using the fact that $p_t = \frac{e^{-t}-e^{-\varepsilon}}{1-e^{-\varepsilon}}$, we obtain

$$\frac{\partial f_\ell(t,\varepsilon_g)}{\partial t} = \binom{k}{\ell}(k-\ell)\left(\frac{-e^{-t}}{1-e^{-\varepsilon}}\right)p_t^{k-1-\ell}(1-p_t)^\ell\left(e^{kt-\ell\varepsilon} - e^{\varepsilon_g}\right) + \binom{k}{\ell}\ell\left(\frac{e^{-t}}{1-e^{-\varepsilon}}\right)p_t^{k-\ell}(1-p_t)^{\ell-1}\left(e^{kt-\ell\varepsilon} - e^{\varepsilon_g}\right) + \binom{k}{\ell}ke^{kt-\ell\varepsilon}p_t^{k-\ell}(1-p_t)^\ell$$

We can pull out similar terms from each expression to achieve

$$\frac{\partial f_\ell(t,\varepsilon_g)}{\partial t} = \binom{k}{\ell}p_t^{k-1-\ell}(1-p_t)^{\ell-1}\left(\frac{1}{1-e^{-\varepsilon}}\right)^2\Big(-(k-\ell)e^{-t}(1-e^{-t})\left(e^{kt-\ell\varepsilon} - e^{\varepsilon_g}\right) + \ell e^{-t}(e^{-t}-e^{-\varepsilon})\left(e^{kt-\ell\varepsilon} - e^{\varepsilon_g}\right) + ke^{kt-\ell\varepsilon}\left(e^{-t}-e^{-\varepsilon}\right)\left(1-e^{-t}\right)\Big)$$

Further examination of the inner term by expanding each expression and cancelling like terms gives

$$\begin{aligned}-(k-\ell)e^{-t}(1-e^{-t})\left(e^{kt-\ell\varepsilon} - e^{\varepsilon_g}\right) + \ell e^{-t}(e^{-t}-e^{-\varepsilon})\left(e^{kt-\ell\varepsilon} - e^{\varepsilon_g}\right) + ke^{kt-\ell\varepsilon}\left(e^{-t}-e^{-\varepsilon}\right)\left(1-e^{-t}\right)\\ = (k-\ell)\left(e^{\varepsilon_g - t} + e^{(k-1)t-(\ell+1)\varepsilon}\right) + \ell\left(e^{(k-1)t-\ell\varepsilon} + e^{\varepsilon_g-\varepsilon-t}\right) - k\left(e^{\varepsilon_g-2t} + e^{kt-(\ell+1)\varepsilon}\right)\end{aligned}$$

This then implies our desired expression.

□

We now have the pieces to give a simpler evaluation of the partial derivative for the general case using an inductive argument. Surprisingly, with a bit of combinatorial and algebraic massaging, the full partial derivative will reduce to a rather simple expression.

Proof of Lemma 5.7. The base case of $\ell = 0$ is true from Lemma B.3. We then assume the claim for $\ell - 1$, and by Corollary B.1 we know $F_\ell(t, \varepsilon_g) = F_{\ell-1}(t, \varepsilon_g) + f_\ell(t, \varepsilon_g)$, which implies

$$\frac{\partial F_\ell(t, \varepsilon_g)}{\partial t} = \frac{\partial F_{\ell-1}(t, \varepsilon_g)}{\partial t} + \frac{\partial f_\ell(t, \varepsilon_g)}{\partial t}$$

Applying our inductive claim and Lemma B.4 we then have

$$\begin{aligned}\frac{\partial F_\ell(t, \varepsilon_g)}{\partial t} = (k - (\ell-1))\binom{k}{\ell-1} p_t^{k-1-(\ell-1)}(1-p_t)^{\ell-1}\frac{1}{1-e^{-\varepsilon}}\left(e^{\varepsilon_g - t} - e^{kt-\ell\varepsilon}\right) + \\ \binom{k}{\ell} p_t^{k-1-\ell}(1-p_t)^{\ell-1}\left(\frac{1}{1-e^{-\varepsilon}}\right)^2\Big((k-\ell)\left(e^{\varepsilon_g - t} + e^{(k-1)t-(\ell+1)\varepsilon}\right) \\ + \ell\left(e^{(k-1)t-\ell\varepsilon} + e^{\varepsilon_g - \varepsilon - t}\right) - k\left(e^{\varepsilon_g - 2t} + e^{kt-(\ell+1)\varepsilon}\right)\Big)\end{aligned}$$

We use the fact that $(k - (\ell-1))\binom{k}{\ell-1} = \ell\binom{k}{\ell}$ and this reduces to

$$\begin{aligned}\frac{\partial F_\ell(t, \varepsilon_g)}{\partial t} = \binom{k}{\ell} p_t^{k-1-\ell}(1-p_t)^{\ell-1}\left(\frac{1}{1-e^{-\varepsilon}}\right)^2\Big(\ell\left(e^{-t} - e^{-\varepsilon}\right)\left(e^{\varepsilon_g - t} - e^{kt-\ell\varepsilon}\right) + \\ (k-\ell)\left(e^{\varepsilon_g - t} + e^{(k-1)t-(\ell+1)\varepsilon}\right) + \ell\left(e^{(k-1)t-\ell\varepsilon} + e^{\varepsilon_g - \varepsilon - t}\right) - k\left(e^{\varepsilon_g - 2t} + e^{kt-(\ell+1)\varepsilon}\right)\Big)\end{aligned}$$

Further examination of the inner term by expanding each expression and cancelling like terms gives

$$\begin{aligned}&\ell\left(e^{-t} - e^{-\varepsilon}\right)\left(e^{\varepsilon_g - t} - e^{kt-\ell\varepsilon}\right) + (k-\ell)\left(e^{\varepsilon_g - t} + e^{(k-1)t-(\ell+1)\varepsilon}\right) \\ &\quad + \ell\left(e^{(k-1)t-\ell\varepsilon} + e^{\varepsilon_g - \varepsilon - t}\right) - k\left(e^{\varepsilon_g - 2t} + e^{kt-(\ell+1)\varepsilon}\right) \\ &\quad = (k-\ell)\left(e^{\varepsilon_g - t} - e^{\varepsilon_g - 2t} + e^{(k-1)t-(\ell+1)\varepsilon} - e^{kt-(\ell+1)\varepsilon}\right) \\ &\quad = (k-\ell)(1-e^{-t})\left(e^{\varepsilon_g - t} - e^{kt-(\ell+1)\varepsilon}\right)\end{aligned}$$

Substituting for this simplified expression and using the fact that $1 - p_t = \frac{1-e^{-t}}{1-e^{-\varepsilon}}$ then gives our desired result. □

C Omitted Proofs from Section 6

We provide here the proofs from Section 6 that were omitted.

C.1 Proofs from Section 6.1

Proof of Corollary 6.1. Note that by our definition, $q_t = e^t p_t$ and $1 - q_t = e^{t-\varepsilon}(1 - p_t)$, so we can equivalently write

$$\delta^k(t, \varepsilon_g) = \sum_{i=0}^{k} \binom{k}{i} q_t^{k-i} (1-q_t)^i \max\left\{\left(1 - e^{\varepsilon_g - kt + i\varepsilon}\right), 0\right\}.$$

We then prove by induction. For $k = 1$, the base case,

$$\delta^1(t, \varepsilon_g) = q_t \max\{1 - e^{\varepsilon_g - t}, 0\} + (1-q_t)\max\{1 - e^{\varepsilon_g - t + \varepsilon}, 0\},$$

and the claim follows by definition of $\delta^0(t, \varepsilon_g)$. We can then apply our inductive hypothesis to get both

$$\begin{aligned}
q_t \cdot \delta^{k-1}(\varepsilon_g - t) &= \sum_{i=0}^{k-1} \binom{k-1}{i} q_t^{k-i} (1-q_t)^i \max\left\{\left(1 - e^{\varepsilon_g - kt + i\varepsilon}\right), 0\right\}, \\
(1-q_t) \cdot \delta^{k-1}(\varepsilon_g - t + \varepsilon) &= \sum_{i=0}^{k-1} \binom{k-1}{i} q_t^{k-1-i} (1-q_t)^{i+1} \max\left\{\left(1 - e^{\varepsilon_g - kt + (i+1)\varepsilon}\right), 0\right\} \\
&= \sum_{i=1}^{k} \binom{k-1}{i-1} q_t^{k-i} (1-q_t)^i \max\left\{\left(1 - e^{\varepsilon_g - kt + i\varepsilon}\right), 0\right\}.
\end{aligned}$$

Our claim then follows from the fact that for any $i \in [1, k-1]$, we must have $\binom{k-1}{i-1} + \binom{k-1}{i} = \binom{k}{i}$. □

Proof of Lemma 6.3. We prove this inductively. For the base case $k = 2$, from Corollary 6.1 and our definition of $t_{\texttt{OPT}}(\mathcal{M}^2_{\texttt{BR}}, \varepsilon_g)$ we have

$$\delta_{\texttt{OPT}}(\mathcal{M}^2_{\texttt{BR}}, \varepsilon_g) = q_t \delta^1(t, \varepsilon_g - t) + (1-q_t)\delta^1(t, \varepsilon_g + \varepsilon - t)$$

If $t \notin t_{\texttt{OPT}}(\mathcal{M}^1_{\texttt{BR}}, \varepsilon_g - t + \ell'\varepsilon)$ for some $\ell' \in \{0, 1\}$, then $\delta_{\texttt{OPT}}(\mathcal{M}^1_{\texttt{BR}}, \varepsilon_g - t + \ell'\varepsilon) > \delta^1(t, \varepsilon_g - t + \ell'\varepsilon)$. Applying Lemma 6.1 for the homogeneous case,

$$\begin{aligned}
\delta_{\texttt{OPT}}(\mathcal{A}^2_{\texttt{BR}}, \varepsilon_g) &\geq q_t \delta_{\texttt{OPT}}(\mathcal{A}^1_{\texttt{BR}}, \varepsilon_g - t) + (1-q_t)\delta_{\texttt{OPT}}(\mathcal{A}^1_{\texttt{BR}}, \varepsilon_g - t + \varepsilon) \\
&> q_t \delta^1(t, \varepsilon_g - t) + (1-q_t)\delta^1(t, \varepsilon_g + \varepsilon - t) = \delta_{\texttt{OPT}}(\mathcal{M}^2_{\texttt{BR}}, \varepsilon_g)
\end{aligned}$$

This equivalently follows if $\delta_{\texttt{OPT}}(\mathcal{M}^1_{\texttt{BR}}, \varepsilon_g - t + \ell'\varepsilon) < \delta_{\texttt{OPT}}(\mathcal{A}^1_{\texttt{BR}}, \varepsilon_g - t + \ell'\varepsilon)$ for either $\ell' \in \{0, 1\}$.

The inductive step will then follow equivalently. Once again, we have

$$\delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g) = q_t \delta^{k-1}(t, \varepsilon_g - t) + (1-q_t)\delta^{k-1}(t, \varepsilon_g + \varepsilon - t)$$

which similarly implies

$$\begin{aligned}
\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) &\geq q_t \delta_{\texttt{OPT}}(\mathcal{A}^{k-1}_{\texttt{BR}}, \varepsilon_q - t) + (1-q_t)\delta_{\texttt{OPT}}(\mathcal{A}^{k-1}_{\texttt{BR}}, \varepsilon_g - t + \varepsilon) \\
&\geq q_t \delta_{\texttt{OPT}}(\mathcal{M}^{k-1}_{\texttt{BR}}, \varepsilon_q - t) + (1-q_t)\delta_{\texttt{OPT}}(\mathcal{M}^{k-1}_{\texttt{BR}}, \varepsilon_g - t + \varepsilon) \\
&\geq q_t \delta^{k-1}(t, \varepsilon_g - t) + (1-q_t)\delta^{k-1}(t, \varepsilon_g + \varepsilon - t) = \delta_{\texttt{OPT}}(\mathcal{M}^k_{\texttt{BR}}, \varepsilon_g)
\end{aligned}$$

The goal will then be to show that this inequality becomes strict if one of the conditions in the statement holds. First, suppose $t \notin t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1}, \varepsilon_g - t + \ell'\varepsilon)$ for either $\ell' \in \{0,1\}$, then $\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1}, \varepsilon_g - t + \ell'\varepsilon) > \delta_g^{k-1}(t, \varepsilon_g - t + \ell'\varepsilon)$ and the inequality must be strict. On the other hand, if $t \in t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1}, \varepsilon_g - t + \ell'\varepsilon)$ for both $\ell' \in \{0,1\}$, then this fits the condition of our inductive hypothesis, and we will then use this to prove our claim for the remaining cases.

Let $0 \le \ell' \le \ell < k$ be such that $\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{k-\ell}, \varepsilon_g - \ell t + \ell'\varepsilon) > \delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-\ell}, \varepsilon_g - \ell t + \ell'\varepsilon)$, and if $\ell = 0$, then the inequality holds trivially. If $\ell \ge 1$, then rewriting the inequality, we equivalently have both of the following inequalities,

$$\begin{aligned}
&\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{k-1-(\ell-1)}, \varepsilon_g - t - (\ell-1)t + \ell'\varepsilon) > \delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1-(\ell-1)}, \varepsilon_g - t - (\ell-1)t + \ell'\varepsilon),\\
\text{and}\quad &\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{k-1-(\ell-1)}, \varepsilon_g - t + \varepsilon - (\ell-1)t + (\ell'-1)\varepsilon)\\
&\qquad\qquad > \delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1-(\ell-1)}, \varepsilon_g - t + \varepsilon - (\ell-1)t + (\ell'-1)\varepsilon).
\end{aligned}$$

If $\ell \ge 1$, then we must have either $0 \le \ell' \le (\ell-1) < k-1$ or $0 \le (\ell'-1) \le (\ell-1) < k-1$. We can then apply our inductive hypothesis to achieve $\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{k-1}, \varepsilon_g - t) > \delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1}, \varepsilon_g - t)$, or $\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{k-1}, \varepsilon_g - t + \varepsilon) > \delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1}, \varepsilon_g - t + \varepsilon)$, respectively, which implies that our inequality is strict.

Similarly, let $0 \le \ell' \le \ell < k$ be such that $t \notin t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-\ell}, \varepsilon_g - \ell t + \ell'\varepsilon)$ By definition we cannot have $\ell = 0$, and we previously considered $\ell = 1$, so we assume $\ell > 1$ in order to apply our inductive claim. Rewriting the set $t_{\texttt{OPT}}$, we must then have both hold

$$\begin{aligned}
&t \notin t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1-(\ell-1)}, \varepsilon_g - t - (\ell-1)t + \ell'\varepsilon)\\
\text{and}\qquad &t \notin t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1-(\ell-1)}, \varepsilon_g - t + \varepsilon - (\ell-1)t + (\ell'-1)\varepsilon).
\end{aligned}$$

If $\ell > 1$, then $\ell - 1 > 0$ and either $0 \le \ell' \le (\ell-1) < k-1$ or $0 \le (\ell'-1) \le (\ell-1) < k-1$. Applying our inductive hypothesis, we have either case hold, respectively

$$\begin{aligned}
&\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{k-1}, \varepsilon_g - t) > \delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1}, \varepsilon_g - t),\\
\text{or } &\delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^{k-1}, \varepsilon_g - t + \varepsilon) > \delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^{k-1}, \varepsilon_g - t + \varepsilon).
\end{aligned}$$

This implies our inequality is strict. □

In order to prove Lemma 6.5, we will also need the following edge case.

Lemma C.1. *$t_{OPT}(\mathcal{M}_{BR}^2, -3\varepsilon/2) \cap t_{OPT}(\mathcal{M}_{BR}^2, \varepsilon/2) = \emptyset$ and $t_{OPT}(\mathcal{M}_{BR}^2, -\varepsilon/2) \cap t_{OPT}(\mathcal{M}_{BR}^2, \varepsilon/2) = \emptyset$*

Proof. For any $\varepsilon_g \in \{-3\varepsilon/2, -\varepsilon/2, \varepsilon/2, 3\varepsilon/2\}$, from Lemma 6.2 that $t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^2, \varepsilon_g) \subseteq \{\frac{\varepsilon_g + (\ell+1)\varepsilon}{3}\} \cap (0, \varepsilon)$ for $\ell \in \{0,1\}$. This then implies that $t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^2, -3\varepsilon/2) = \varepsilon/6$, $t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^2, -\varepsilon/2) \subseteq \{\varepsilon/6, \varepsilon/2\}$, $t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^2, \varepsilon/2) \subseteq \{\varepsilon/2, 5\varepsilon/6\}$, and $t_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^2, 3\varepsilon/2) = 5\varepsilon/6$. The claim then follows immediately. □

Proof of Lemma 6.5. By our assumptions, it immediately follows that either there exists $0 \le j \le k-2$ such that $\varepsilon_g - (k-2)t + j\varepsilon \in (-\varepsilon/2, \varepsilon/2)$, or we are in the edge case where there exists $0 \le j < k-2$ such that $\varepsilon_g - (k-2)t + j\varepsilon = -\varepsilon/2$. In first case, we know that $\delta_{\texttt{OPT}}(\mathcal{M}_{\texttt{BR}}^2, \varepsilon_g - (k-2)t_\ell + j\varepsilon) < \delta_{\texttt{OPT}}(\mathcal{A}_{\texttt{BR}}^2, \varepsilon_g - (k-2)t + j\varepsilon)$ from Lemma 6.4. In the second case (the edge case), if $j = 0$ then we know $j + 2 \le k - 2$ because $k \ge 4$, and from Lemma C.1 we must either

have $t \notin t_{\texttt{OPT}}(\mathcal{M}^2_{\texttt{BR}}, \varepsilon_g - (k-2)t_\ell + j\varepsilon)$ or $t \notin t_{\texttt{OPT}}(\mathcal{M}^2_{\texttt{BR}}, \varepsilon_g - (k-2)t_\ell + (j+2)\varepsilon)$. Otherwise, if $j > 0$, then we again have from Lemma C.1 that either $t \notin t_{\texttt{OPT}}(\mathcal{M}^2_{\texttt{BR}}, \varepsilon_g - (k-2)t_\ell + (j-1)\varepsilon)$ or $t \notin t_{\texttt{OPT}}(\mathcal{M}^2_{\texttt{BR}}, \varepsilon_g - (k-2)t_\ell + (j+1)\varepsilon)$.

In either case, we can immediately apply Lemma 6.3 to achieve our desired inequality. □

C.2 Proofs from Section 6.2

Proof of Lemma 6.7. We will prove both statements by induction, where the base case $\delta_{\texttt{OPT}}(\mathcal{A}^0_{\texttt{BR}}, \varepsilon_g) = \max\{1 - e^{\varepsilon_g}, 0\} = 0$ for $\varepsilon_g \geq 0$ and $\delta_{\texttt{OPT}}(\mathcal{A}^0_{\texttt{BR}}, \varepsilon_g) = \max\{1 - e^{\varepsilon_g}, 0\} = 1 - e^{\varepsilon_g}$ for $\varepsilon_g \leq 0$. For any $t \in [0, \varepsilon]$, if $\varepsilon_g \geq k\varepsilon$ we must have $\varepsilon_g - t \geq (k-1)\varepsilon$ and $\varepsilon_g - t + \varepsilon \geq (k-1)\varepsilon$. Similarly, if $\varepsilon_g \leq -k\varepsilon$ we must have $\varepsilon_g - t \leq -(k-1)\varepsilon$ and $\varepsilon_g - t + \varepsilon \leq -(k-1)\varepsilon$. Using Lemma 6.1 for the homogeneous case, we know

$$\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) = \sup_{t \in [0,\varepsilon]} \left\{ q_t \delta_{\texttt{OPT}}(\mathcal{A}^{k-1}_{\texttt{BR}}, \varepsilon_g - t) + (1 - q_t)\delta_{\texttt{OPT}}(\mathcal{A}^{k-1}_{\texttt{BR}}, \varepsilon_g + \varepsilon - t) \right\}$$

and applying our inductive hypothesis easily gives $\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) = 0$ for any $\varepsilon_g \geq k\varepsilon$. Applying our inductive hypothesis for $\varepsilon_g \leq -k\varepsilon$, we have for any t that

$$\begin{aligned} &q_t \delta_{\texttt{OPT}}(\mathcal{A}^{k-1}_{\texttt{BR}}, \varepsilon_g - t) + (1 - q_t)\delta_{\texttt{OPT}}(\mathcal{A}^{k-1}_{\texttt{BR}}, \varepsilon_g + \varepsilon - t) \\ &\qquad = \frac{1 - e^{t-\varepsilon}}{1 - e^{-\varepsilon}} \left(1 - e^{\varepsilon_g - t}\right) + \frac{e^{t-\varepsilon} - e^{-\varepsilon}}{1 - e^{-\varepsilon}} \left(1 - e^{\varepsilon_g - t + \varepsilon}\right) = 1 - e^{\varepsilon_g}. \end{aligned}$$

This then implies $\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) = 1 - e^{\varepsilon_g}$ for any $\varepsilon_g \leq -k\varepsilon$.

□

Proof of Lemma 6.9. We show this inductively. For $k = 1$, if $\varepsilon_g \leq 0$, then for any $t \in [0, \varepsilon]$, we must have $\varepsilon_g - t \leq 0$ and $\max\{1 - e^{\varepsilon_g - t}, 0\} = 1 - e^{\varepsilon_g - t}$. This then implies

$$\delta_{\texttt{OPT}}(\mathcal{A}^1_{\texttt{BR}}, \varepsilon_g) = \sup_{t \in [0,\varepsilon]} \left\{ q_t(1 - e^{\varepsilon_g - t}) + (1 - q_t) \max\{1 - e^{\varepsilon_g + \varepsilon - t}, 0\} \right\}.$$

If $\varepsilon_g \leq -\varepsilon$ then $\max\{1 - e^{\varepsilon_g + \varepsilon - t}, 0\} = 1 - e^{\varepsilon_g + \varepsilon - t}$ for any $t \in [0, \varepsilon]$ and $\delta_{\texttt{OPT}}(\mathcal{A}^1_{\texttt{BR}}, \varepsilon_g) = 1 - e^{\varepsilon_g}$ because from the proof of Lemma 6.7 we have $q_t(1 - e^{\varepsilon_g - t}) + (1 - q_t)(1 - e^{\varepsilon_g + \varepsilon - t}) = 1 - e^{\varepsilon_g}$ for any t. Furthermore, $e^{\varepsilon_g + \varepsilon - t} - 1 \leq 0$ for any $t \in [0, \varepsilon]$, so $\sup_{t \in [0,\varepsilon]} \left\{ (1 - q_t)(e^{\varepsilon_g + \varepsilon - t} - 1) \right\} = 0$ by setting $t = 0$, and we have our desired equality.

If $\varepsilon < \varepsilon_g \leq 0$, then there must exist some $t \in [0, \varepsilon]$ such that $\varepsilon_g + \varepsilon - t > 0$. Once again, we know $q_t(1 - e^{\varepsilon_g - t}) + (1 - q_t)(1 - e^{\varepsilon_g + \varepsilon - t}) = 1 - e^{\varepsilon_g}$ for any t. Consequently, the supremum must be achieved for some $t \in [0, \varepsilon_g + \varepsilon) \subset [0, \varepsilon]$ such that $1 - e^{\varepsilon_g + \varepsilon - t} < 0$. Thus,

$$\begin{aligned} \delta_{\texttt{OPT}}(\mathcal{A}^1_{\texttt{BR}}, \varepsilon_g) &= \sup_{t \in [0, \varepsilon_g + \varepsilon)} q_t(1 - e^{\varepsilon_g - t}) \\ &= \sup_{t \in [0, \varepsilon_g + \varepsilon)} \left\{ q_t(1 - e^{\varepsilon_g - t}) + (1 - q_t)(1 - e^{\varepsilon_g + \varepsilon - t}) + (1 - q_t)(e^{\varepsilon_g + \varepsilon - t} - 1) \right\} \\ &= 1 - e^{\varepsilon_g} + \sup_{t \in [0, \varepsilon_g + \varepsilon)} (1 - q_t)(e^{\varepsilon_g + \varepsilon - t} - 1) \\ &= 1 - e^{\varepsilon_g} + \sup_{t \in [0, \varepsilon]} (1 - q_t)(e^{\varepsilon_g + \varepsilon - t} - 1). \end{aligned}$$

The inductive step for $k \geq 2$ follows more easily, where for any $t \in [0, \varepsilon]$, we must have $\varepsilon_g - t \leq -(k-1)\varepsilon$, so from Lemma 6.7, we have

$$\delta_{\texttt{OPT}}(\mathcal{A}^k_{\texttt{BR}}, \varepsilon_g) = \sup_{t \in [0,\varepsilon]} \left\{ q_t(1 - e^{\varepsilon_g - t}) + (1 - q_t)\delta_{\texttt{OPT}}(\mathcal{A}^{2:k}_{\texttt{BR}}, \varepsilon_g + \varepsilon - t) \right\}.$$

We can then apply our inductive hypothesis because $k - 1 \geq 1$ and $\varepsilon_g + \varepsilon - t \leq -(k-2)\varepsilon$, and therefore

$$\delta_{\texttt{OPT}}(\mathcal{A}^{k-1}_{\texttt{BR}}, \varepsilon_g + \varepsilon - t) = 1 - e^{\varepsilon_g + \varepsilon - t} + \sup_{t_i \in [0,\varepsilon]} \prod_{i=1}^{k-1} (1 - q_{t_i}) \left(e^{\varepsilon_g + \varepsilon - t + (k-1)\varepsilon - \sum_{i=1}^{k-1} t_i} - 1 \right)$$

Plugging in this term and once again using the fact that $q_t(1 - e^{\varepsilon_g - t}) + (1 - q_t)(1 - e^{\varepsilon_g + \varepsilon - t}) = 1 - e^{\varepsilon_g}$ for any t, gives our desired equality. □