Composition Theorems for Differential Privacy

Composition Theorems for Differential Privacy

We will define a composition of mechanisms $\mathcal{M}_1, \mathcal{M}_2, ..., \mathcal{M}_k$ as $\mathcal{M}(x)$,
Where $\mathcal{M}(x) = \langle \mathcal{M}_1(x), \mathcal{M}_2(x), ..., \mathcal{M}_k(x) \rangle$

Composition Theorems for Differential Privacy

We will define a composition of mechanisms $\mathcal{M}_1, \mathcal{M}_2, ..., \mathcal{M}_k$ as $\mathcal{M}(x)$,
Where $\mathcal{M}(x) = \langle \mathcal{M}_1(x), \mathcal{M}_2(x), ..., \mathcal{M}_k(x) \rangle$

Basic Composition
If $\mathcal{M}_1, ...\mathcal{M}_k$ are each (ϵ, δ) *differentially private*, then:

$$\mathcal{M}\ is\ (k\epsilon, k\delta)\ \textit{differentially private}$$

If we are willing to tolerate an increase in the δ term, the privacy parameter ϵ only needs to degrade proportionally to $\sqrt{k}$:

Advanced Composition
If $\mathcal{M}_1, ...\mathcal{M}_k$ are each (ϵ, δ) *differentially private* then for all $\delta' > 0$,

$$\mathcal{M}\ is\ \left(O\left(\sqrt{k \log\left(1/\delta'\right)} \cdot \epsilon + k\epsilon\left(e^{\epsilon} - 1\right)\right), k\delta + \delta'\right)\ \textit{differentially private}.$$

Definition (*differentially private*) For $\epsilon \geq 0,\ \delta \in [0,1]$, we say that randomized mechanism $\mathcal{M} : X^n \longrightarrow R$ is (ϵ, δ) *differentially private* if for every two neighboring DBs $x \sim x' \in X^n$ (DBs that differ on one row), the output distribution of mechanism $\mathcal{M}$ on x should be "similar" to that of $\mathcal{M}$ on x' with $1-\delta$ "confidence":

$$\forall S \subseteq R, Pr\left[\mathcal{M}\left(x\right) \in S\right] \leq e^{\epsilon} \cdot Pr\left[\mathcal{M}\left(x'\right) \in S\right] + \delta$$

Definition (*differentially private*) For $\epsilon \geq 0,\ \delta \in [0,1]$, we say that randomized mechanism $\mathcal{M} : X^n \longrightarrow R$ is (ϵ, δ) *differentially private* if for every two neighboring DBs $x \sim x' \in X^n$ (DBs that differ on one row), the output distribution of mechanism $\mathcal{M}$ on x should be "similar" to that of $\mathcal{M}$ on x' with $1-\delta$ "confidence":

$$\forall S \subseteq R, Pr\left[\mathcal{M}\left(x\right) \in S\right] \leq e^{\epsilon} \cdot Pr\left[\mathcal{M}\left(x'\right) \in S\right] + \delta$$

Definition ((ϵ, δ)*-indistinguishable*) We call two random variables Y and Y' taking values in R (ϵ, δ)*-indistinguishable* if:

$$\forall S \subseteq R, Pr\left[Y \in S\right] \leq e^{\epsilon} \cdot Pr\left[Y' \in S\right] + \delta,\ and$$
$$Pr\left[Y' \in S\right] \leq e^{\epsilon} \cdot Pr\left[Y \in S\right] + \delta$$

Another interpretation for *differentially private* mechanism $\mathcal{M}$ is that for every two neighboring DBs $x \sim x' \in X^n$, The output distribution of mechanism $\mathcal{M}$ on x and x' are (ϵ, δ)*-indistinguishable* variables.

Lemma Two *random variables* Y *and* Y' *are* (ϵ, δ) *indistinguishable if and only if there are two events* $E = E(Y)$ *and* $E' = E'(Y')$ such that:

- $Pr[E], Pr[E'] \geq 1 - \delta,\ and$
- $Y|_E$ and $Y'|_{E'}$ are $(\epsilon, 0) - indistinguishable$

Lemma Two *random variables* Y *and* Y' *are* (ϵ, δ) *indistinguishable if and only if there are two events* $E = E(Y)$ *and* $E' = E'(Y')$ such that:

- $Pr[E], Pr[E'] \geq 1 - \delta,$ *and*
- $Y|_E$ and $Y'|_{E'}$ are $(\epsilon, 0) - indistinguishable$

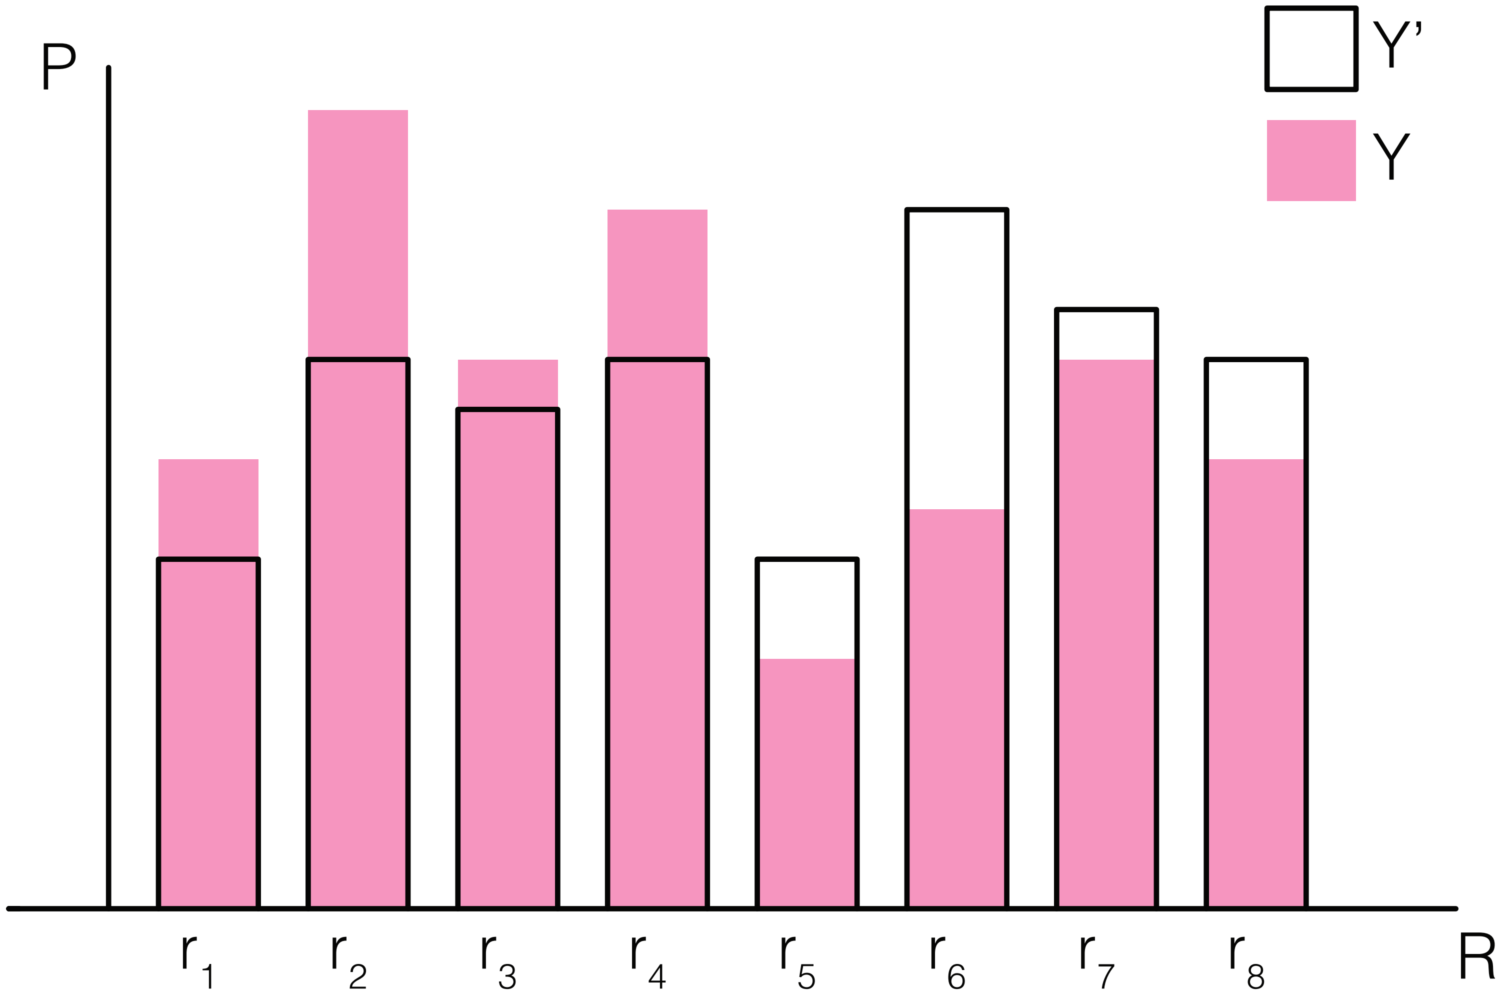

We will mark the bad group as:

$$Bad = \{r_i : e^\epsilon P_{Y'}(r_i) \leq P_Y(r_i)\}$$

since Y and Y' are (ϵ, δ) indistinguishable, it holds that:

$$P_Y(Bad) \leq e^\epsilon P_{Y'}(Bad) + \delta.$$

Which means that:

$$\gamma = \sum_{r_i \in Bad} P_Y(r_i) - e^\epsilon P_{Y'}(r_i) \leq \delta$$

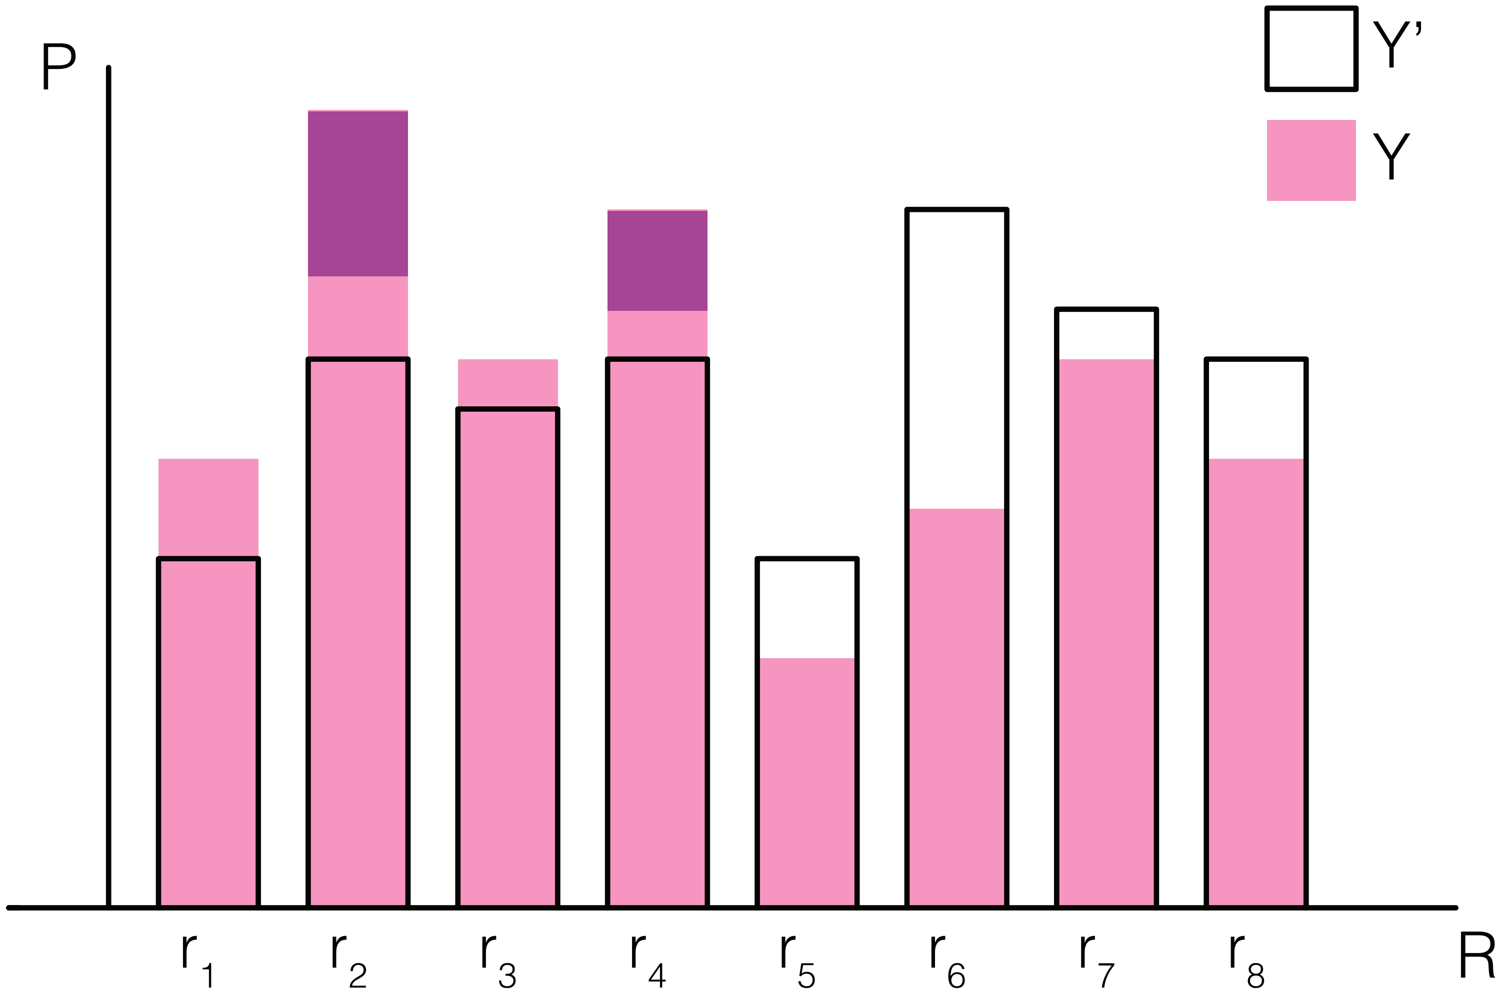

We will define the event $\bar{E}$ as follows:

$$\forall r_i \in Bad.\ if\ Y = r_i\ than\ \bar{E}\ happens\ with\ probability\ \frac{P_Y(r_i) - e^{\epsilon} P_{Y'}(r_i)}{P_Y(r_i)}.$$

We get that

$$P(\bar{E}) = \sum_{r_i \in Bad} P_Y(r_i) \cdot \frac{P_Y(r_i) - e^{\epsilon} P_{Y'}(r_i)}{P_Y(r_i)} = \gamma \leq \delta$$

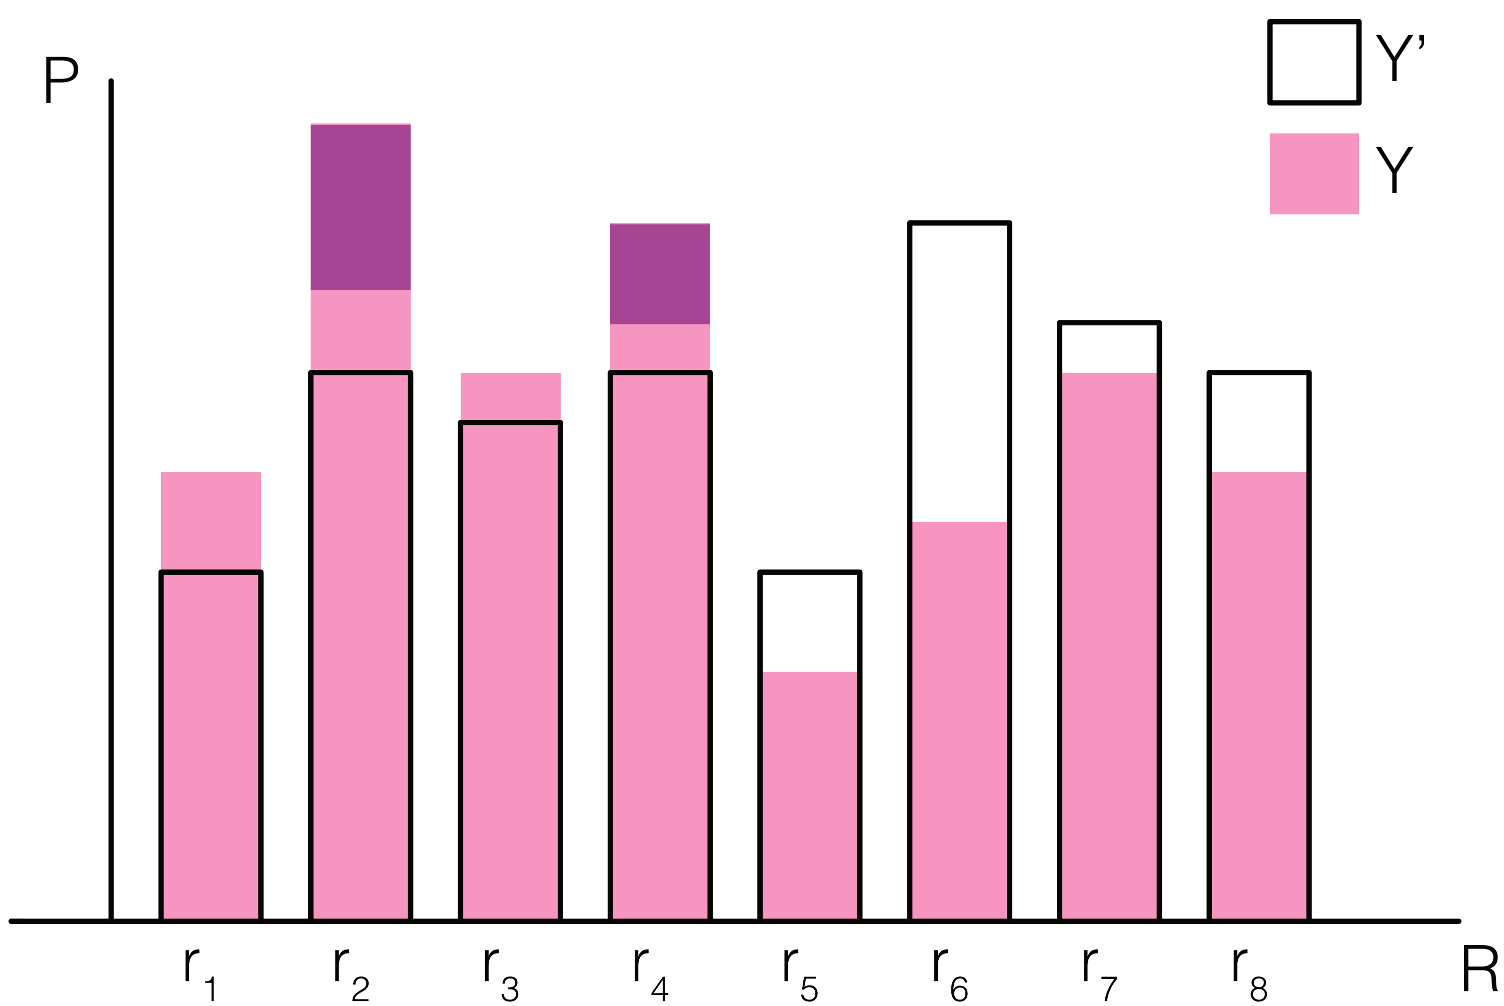

We have *fixed* the bad cases when $e^{\epsilon} P(Y' = r) \leq P(Y = r)$ by looking at

$$P(Y = r|E) = \frac{P(Y = r)}{P(E = r)},$$

But, while doing so, we also scale the cases where $P(Y = r) \leq P(Y' = r)$

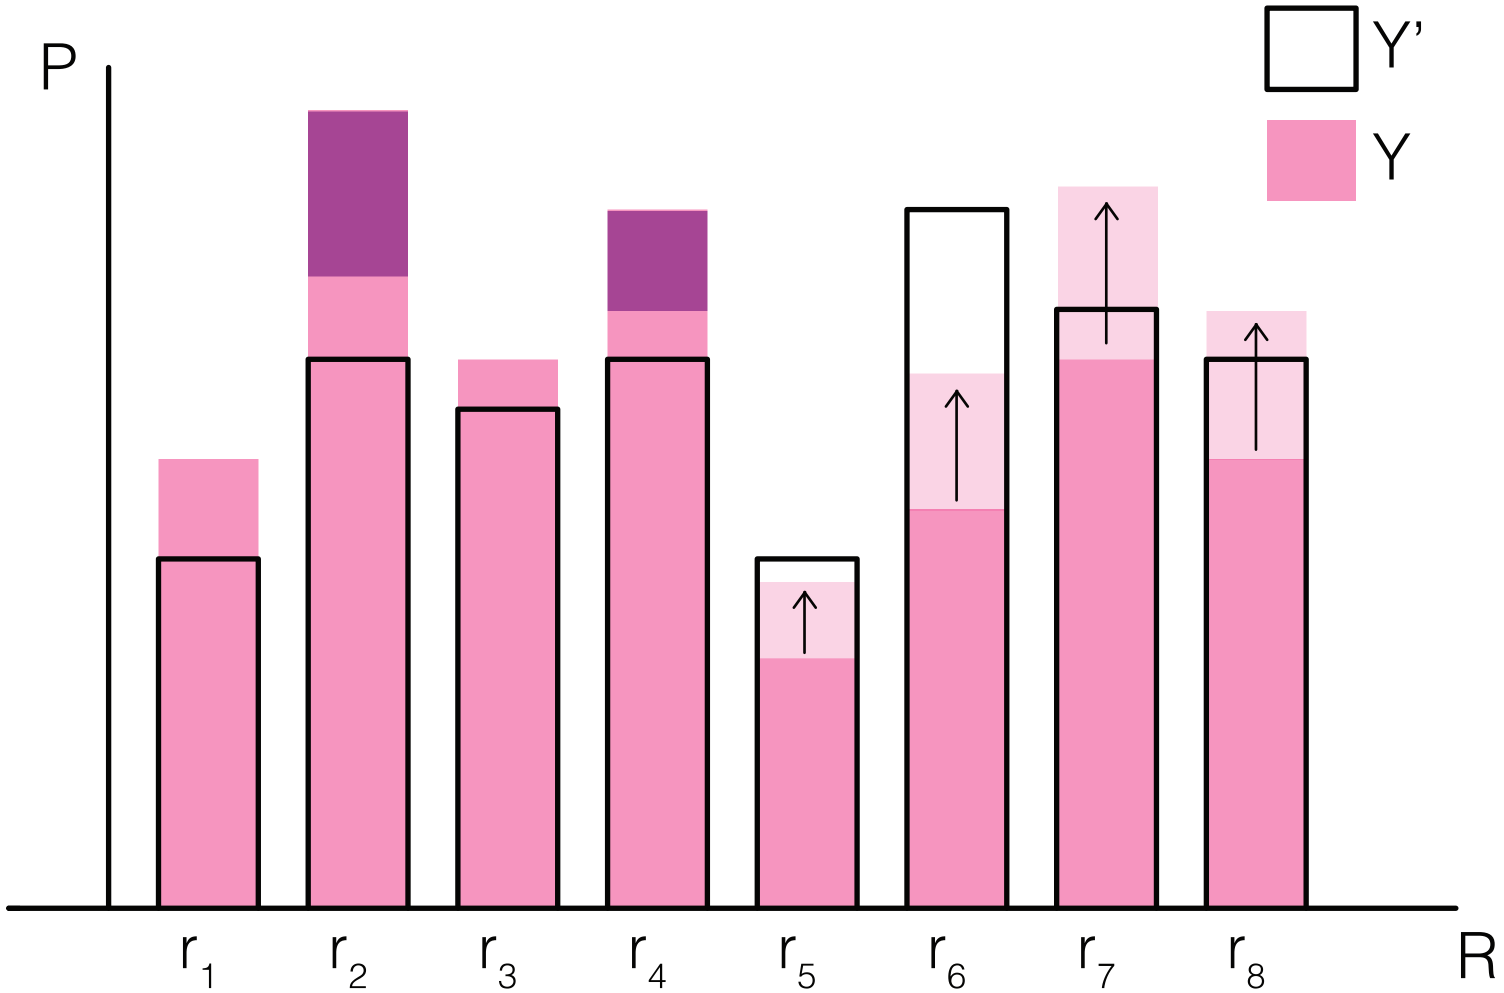

We will correct it by reduce the same γ from $P(Y')$. We will mark group S as:

$$s = \{r_i : (P_Y(r_i) \leq P_{Y'}(r_i)\}$$

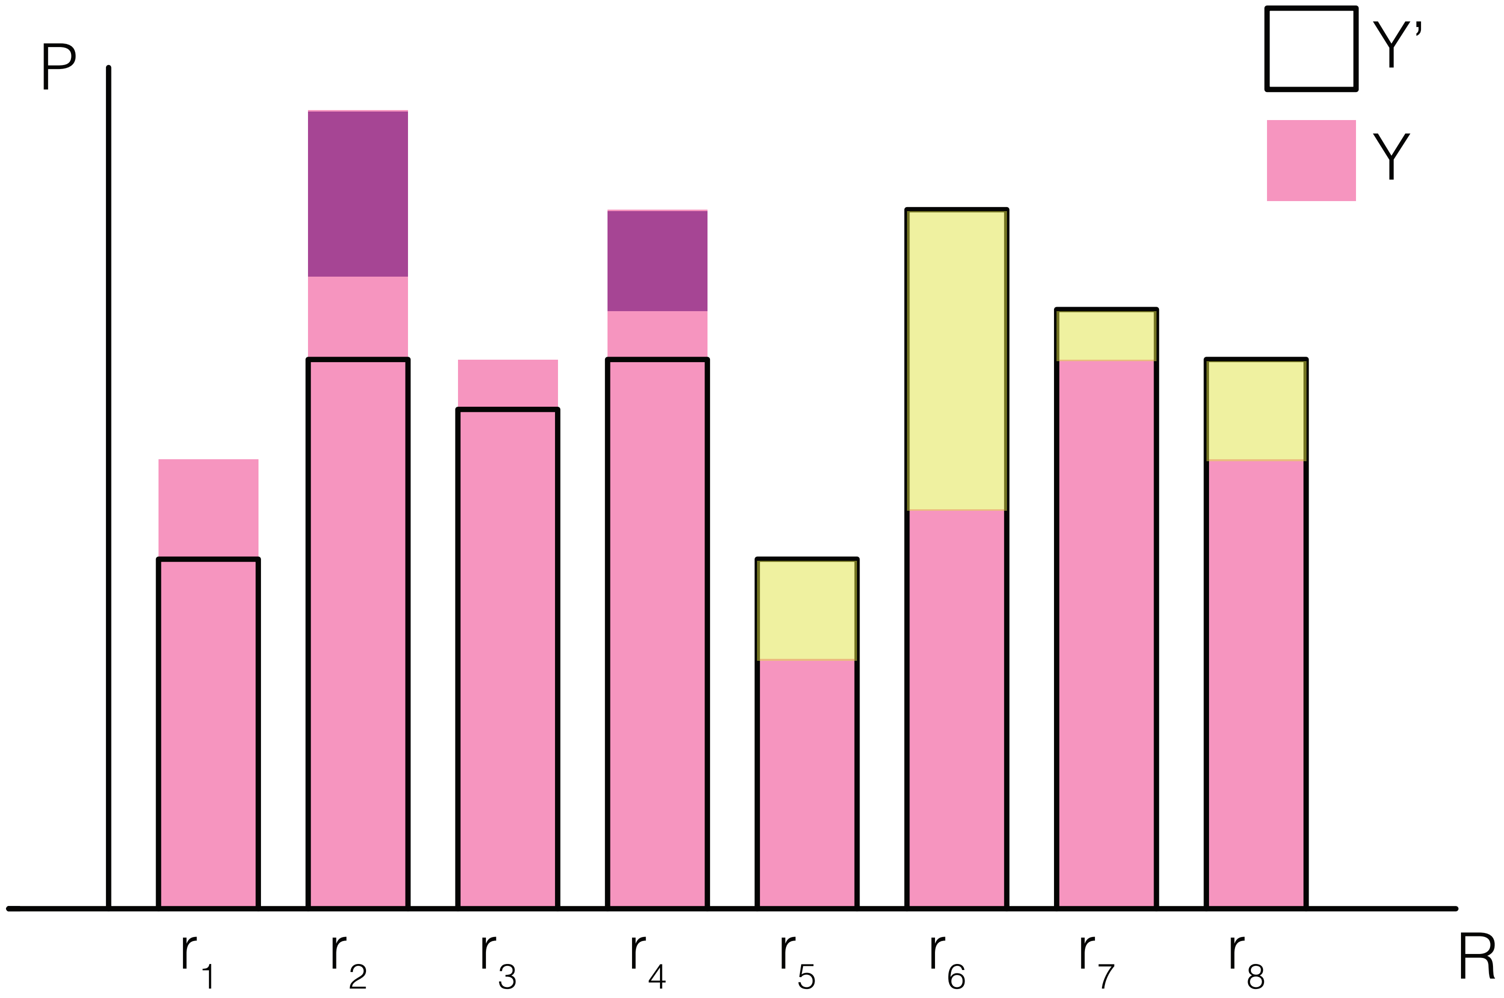

We will correct it by reduce the same γ from $P(Y')$. We will mark group S as:

$$s = \{r_i : (P_Y(r_i) \leq P_{Y'}(r_i)\}$$

and define event $\bar{E}'$ to happened with probability γ by *reducing* the gap between $P(Y)$ and $P(Y')$ in S.

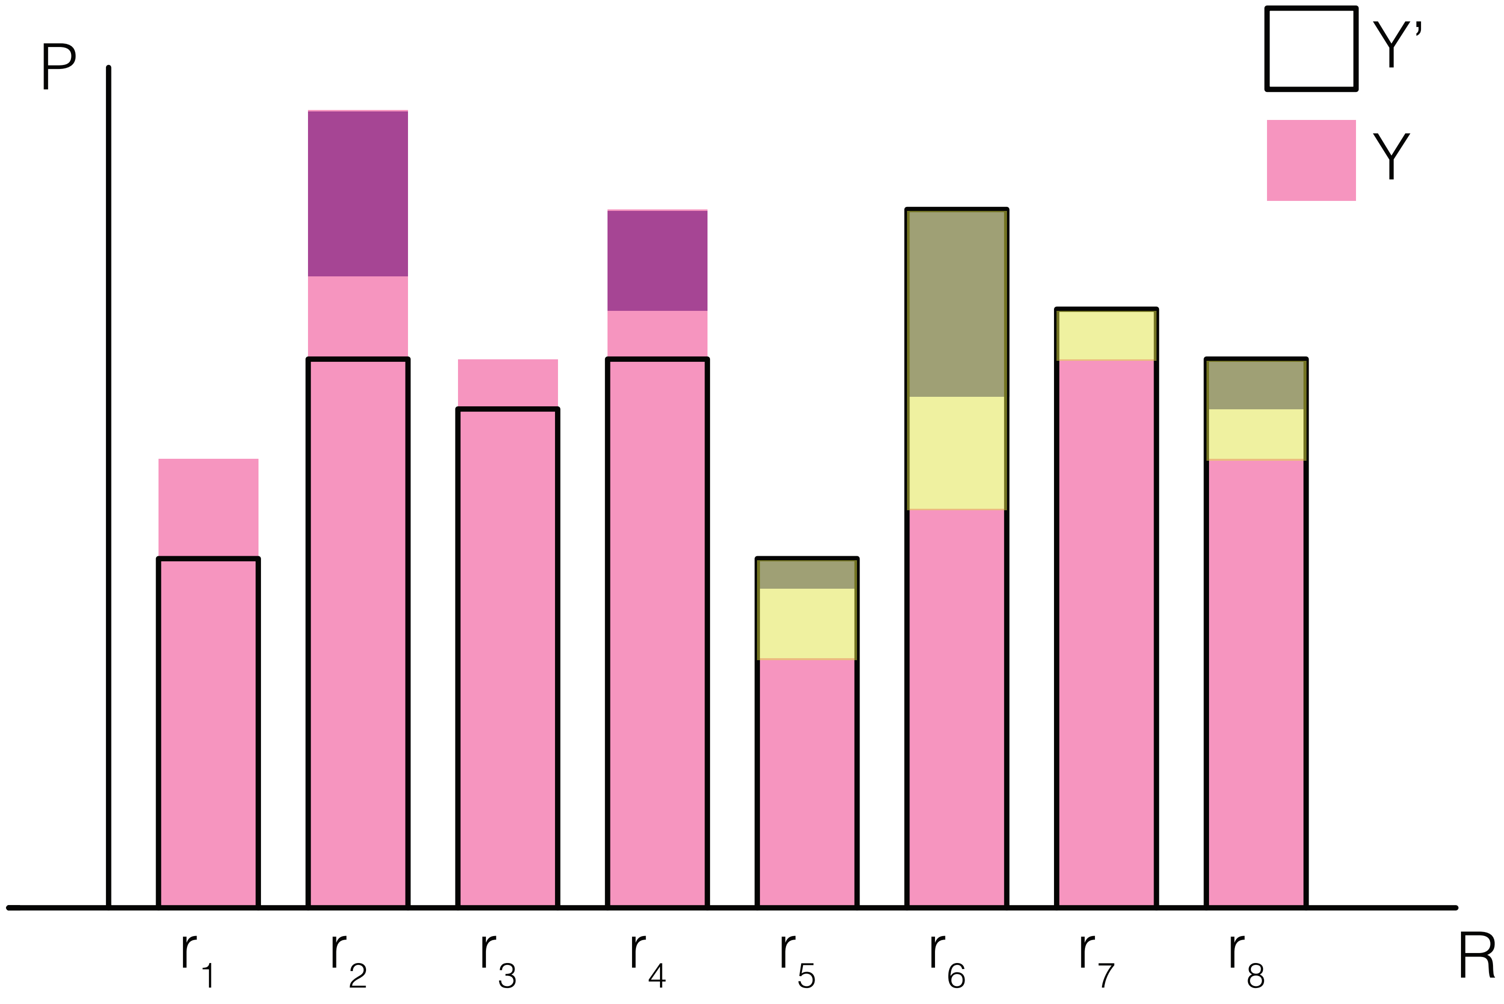

Overall:

- $P(\bar{E}), P(\bar{E}') \leq \delta \longrightarrow P(E), P(E') > 1 - \delta$
- $P(Y|E) \leq e^{\epsilon} P(Y'|E) \longrightarrow Y|_E$ and $Y'|_{E'}$ are $(\epsilon, 0) - indistinguishable$

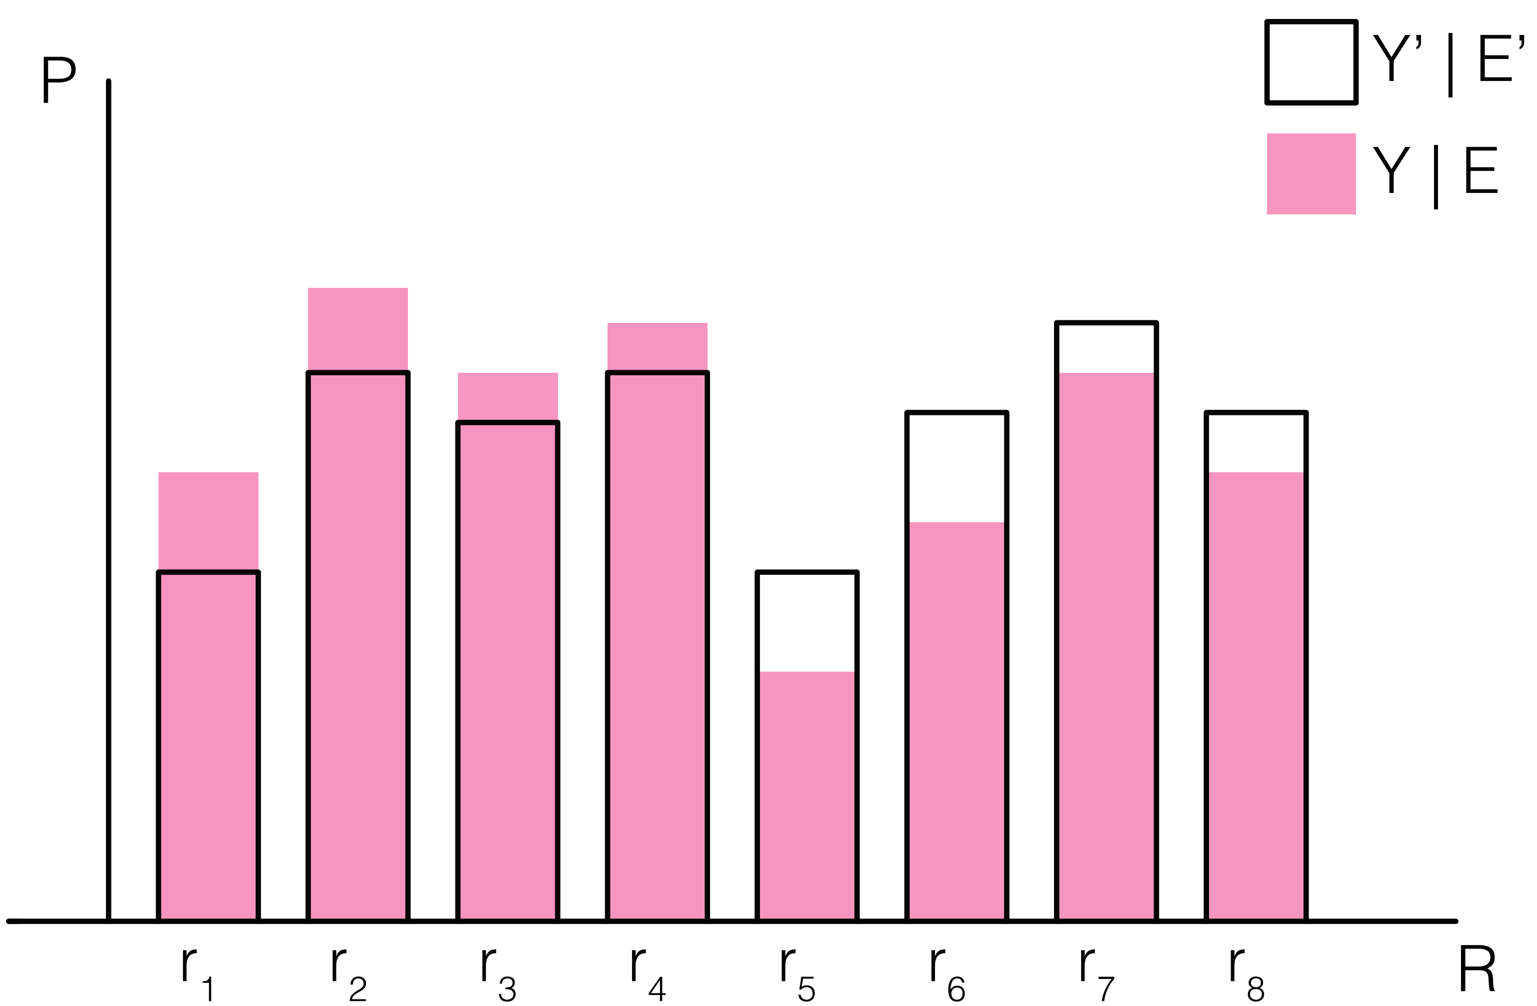

Basic Composition

If $\mathcal{M}_1, ...\mathcal{M}_k$ are each (ϵ, δ) *differentially private*, then:

$$\mathcal{M}\ is\ (k\epsilon, k\delta)\ \textit{differentially private}$$

Advanced Composition
If $\mathcal{M}_1, ...\mathcal{M}_k$ are each (ϵ, δ) *differentially private* then for all $\delta' > 0$,

$$\mathcal{M}\ is\ \left(O\left(\sqrt{k\log\left(1/\delta'\right)}\cdot\epsilon + k\epsilon\left(e^{\epsilon}-1\right)\right), k\delta+\delta'\right)\ \textit{differentially private}.$$

To simplify the proof, we will assume that:

- $\delta = 0$
- $\epsilon \leq 1$ s.t. $\epsilon\left(e^{\epsilon}-1\right) \approx \epsilon^2$
- $k < 1/\epsilon^2$

The tuple $\left(O\left(\sqrt{k\log\left(1/\delta'\right)}\cdot\epsilon + k\epsilon\left(e^{\epsilon}-1\right)\right), k\delta+\delta'\right)$ become $\left(O\left(\sqrt{k\log\left(1/\delta'\right)}\cdot\epsilon\right), \delta'\right)$

Definition (*privacy loss*)

$$L_{\mathcal{M}}^{x \to x'}(r) = \ln\left(\frac{Pr\left[\mathcal{M}(x) = r\right]}{Pr\left[\mathcal{M}(x') = r\right]}\right) = -L_{\mathcal{M}}^{x' \to x}(r)$$

Definition (*privacy loss*)

$$L_{\mathcal{M}}^{x \to x'}(r) = \ln\left(\frac{Pr\left[\mathcal{M}(x) = r\right]}{Pr\left[\mathcal{M}(x') = r\right]}\right) = -L_{\mathcal{M}}^{x' \to x}(r)$$

Definition (*KL-Divergence*). The Kullback—Leibler divergence between two random variables Y and Z taking values from the same domain is defined to be:

$$D(Y \| Z) = \mathbb{E}_{y \sim Y}\left[\ln \frac{Pr[Y = y]}{Pr[Z = y]}\right]$$

Definition (*privacy loss*)

$$L_{\mathcal{M}}^{x \to x'}(r) = \ln\left(\frac{Pr\left[\mathcal{M}(x) = r\right]}{Pr\left[\mathcal{M}(x') = r\right]}\right) = -L_{\mathcal{M}}^{x' \to x}(r)$$

Definition (*KL-Divergence*). The Kullback—Leibler divergence between two random variables Y and Z taking values from the same domain is defined to be:

$$D(Y \| Z) = \mathbb{E}_{y \sim Y}\left[\ln \frac{Pr[Y = y]}{Pr[Z = y]}\right]$$

Notice that $\mathbb{E}_{r \sim R}\left[L_{\mathcal{M}}^{x \to x'}(r)\right] = D\left(\mathcal{M}_i(x) \| \mathcal{M}_i(x')\right)$

The *Max Divergence* between two random variables Y and Z is defined by:

$$D_\infty(Y \| Z) = \max_{S \subseteq Supp(Y)}\left[\ln \frac{Pr[Y \in S]}{Pr[Z \in S]}\right].$$

And finally, the $\delta-$Approximate Max Divergence between Y and Z is:

$$D_\infty^\delta(Y \| Z) = \max_{S \subseteq Supp(Y) : Pr[Y \in S] \geq \delta}\left[\ln \frac{Pr[Y \in S] - \delta}{Pr[Z \in S]}\right].$$

Definition (*privacy loss*)

$$L_{\mathcal{M}}^{x \to x'}(r) = \ln\left(\frac{Pr\left[\mathcal{M}(x) = r\right]}{Pr\left[\mathcal{M}(x') = r\right]}\right) = -L_{\mathcal{M}}^{x' \to x}(r)$$

Lemma If $\mathcal{M}_i$ is ϵ *differentially private*, where $\epsilon \leq 1$, than

$$E_{r \in R}\left[L_{\mathcal{M}_i}^{x \to x'}(r)\right] = D\left[\mathcal{M}_i(x) \| \mathcal{M}_i(x')\right] \leq 2\epsilon^2$$

Advanced Composition
If $\mathcal{M}_1, ...\mathcal{M}_k$ are each (ϵ, δ) *differentially private* then for all $\delta' > 0$,

$$\mathcal{M}\ is \quad \left(O\left(\sqrt{k\log\left(1/\delta'\right)}\cdot\epsilon\right), \delta'\right) \quad \textit{differentially private}.$$

Lemma (Hoeffding's Inequality). Let $X_1,\ \ldots, X_k$ be independent real-valued random variables such that for every i, X_i is bounded by $[a_i, b_i]$, than:

$$Pr\left(S_k \geq E\left[S_k\right] + t\right) \leq \exp\left(\frac{-2t^2}{\sum_{i=1}^{k}\left(b_i - a_i\right)^2}\right),$$

$$where\ S_k = \sum_{i=1}^{k} X_i$$

Advanced Composition
If $\mathcal{M}_1, ...\mathcal{M}_k$ are each (ϵ, δ) *differentially private* then for all $\delta' > 0$,

$$\mathcal{M}\ is\ \ \left(O\left(\sqrt{k\log\left(1/\delta'\right)}\cdot\epsilon\right), \delta'\right)\ \ \textit{differentially private}.$$

Lemma (Azuma's Inequality). Let $C_1,\ ...\ , C_k$ be real-valued random variables such that for every $i \in [k],\ Pr[|C_i| \leq \alpha] = 1$ and for every $c_1,\ ...\ , c_{i-1}$, we have

$$E[C_i|C_1 = c_1,\ ...\ C_{i-1} = c_{i-1}] \leq \beta$$

Than, for every $z > 0$, we have

$$Pr\left[\sum_{i=1}^{k} C_i > k\beta + z\sqrt{k}\cdot\alpha\right] \leq e^{-z^2/2}$$