Characterizing the Distinguishability of Product Distributions through Multicalibration

Cassandra Marcussen, Aaron (Louie) Putterman, Salil Vadhan



Harvard University

Outline

- Problem setting
- Our results
- Previous work
- Multicalibration Theorem
- Proof overviews
- Conclusion

Outline

- **Problem setting**
- Our results
- Previous work
- Multicalibration Theorem
- Proof overviews
- Conclusion

Distinguishing between Distributions

- **"Hypothesis testing" problem:** Given samples $x_1, x_2, \ldots, x_k$ promised to be drawn from distribution X_0 or X_1, how to decide which distribution the samples are from?

Distinguishing between Distributions

- **"Hypothesis testing" problem:** Given samples $x_1, x_2, \ldots, x_k$ promised to be drawn from distribution X_0 or X_1, how to decide which distribution the samples are from?

Distinguishing between Distributions

- **"Hypothesis testing" problem:** Given samples $x_1, x_2, \ldots, x_k$ promised to be drawn from distribution X_0 or X_1, how to decide which distribution the samples are from?

- **Statistical (information-theoretic) limit**: captured by total variation distance between $X_0^{\otimes k}$ and $X_1^{\otimes k}$: $d_{TV}\left(X_0^{\otimes k}, X_1^{\otimes k}\right)^*$.

$^*\, d_{TV}(X, Y) = \frac{1}{2} \sum_{i \in \mathcal{X}} \left| \mathbb{P}[X = i] - \mathbb{P}[Y = i] \right|$

Distinguishing between Distributions

- **"Hypothesis testing" problem:** Given samples $x_1, x_2, \ldots, x_k$ promised to be drawn from distribution X_0 or X_1, how to decide which distribution the samples are from?

- **Statistical (information-theoretic) limit**: captured by total variation distance between $X_0^{\otimes k}$ and $X_1^{\otimes k}$: $d_{TV}\left(X_0^{\otimes k}, X_1^{\otimes k}\right)^*$.

- Our setting is **efficient distinguishers:** study distinguishers computable by small circuits

$$^* d_{TV}(X, Y) = \frac{1}{2} \sum_{i \in \mathcal{X}} \left| \mathbb{P}[X = i] - \mathbb{P}[Y = i] \right|$$

Terminology

- **Statistical distinguishability:** Captured by *total variation distance*.

Terminology

- **Statistical distinguishability:** Captured by *total variation distance*.

- **Computational distinguishability** of X, Y wrt circuits of size s ($\mathscr{C}_s$): Study

$$A(X, Y, C) = \left| \mathbb{P}[C(X) = 1] - \mathbb{P}[C(Y) = 1] \right|.$$

Terminology

- **Statistical distinguishability:** Captured by *total variation distance.*
- **Computational distinguishability** of X, Y wrt circuits of size s ($\mathscr{C}_s$): Study

$$A(X, Y, C) = \left| \mathbb{P}[C(X) = 1] - \mathbb{P}[C(Y) = 1] \right|.$$

  - If, for every $C \in \mathscr{C}_s$, $A(X, Y, C) \leq \varepsilon$, X, Y are (s, ε)-*indistinguishable* wrt circuits of size s. (Notation: $X \approx_\varepsilon Y$)
  - If there exists an $C \in \mathscr{C}_s$ with $A(X, Y, C) \geq \varepsilon$, then X, Y are (s, ε)-*distinguishable* wrt circuits of size s*.

***Note:** $A(X, Y, C) \leq d_{TV}(X, Y)$.

Statistical and computational distinguishability

How many samples $x_1, x_2, \ldots, x_k$ to distinguish X_0 from X_1 with constant advantage?

Statistical and computational distinguishability

How many samples $x_1, x_2, \ldots, x_k$ to distinguish X_0 from X_1 with constant advantage?

- Let $d_H(X_0, X_1)$ be the Hellinger distance* between X_0 and X_1.
- **Statistical distinguishability:** $k = \Theta(d_H^{-2}(X_0, X_1))$

* $d_H^2(X_0, X_1) = \frac{1}{2} \sum_{i \in \mathcal{X}} \left(\sqrt{\mathbb{P}[X_0 = i]} - \sqrt{\mathbb{P}[X_1 = i]} \right)^2$

Statistical and computational distinguishability

How many samples $x_1, x_2, \ldots, x_k$ to distinguish X_0 from X_1 with constant advantage?

- Let $d_H(X_0, X_1)$ be the Hellinger distance* between X_0 and X_1.
- **Statistical distinguishability:** $k = \Theta(d_H^{-2}(X_0, X_1))$
- **Computational distinguishability:** No instance-optimal characterization.

* $d_H^2(X_0, X_1) = \frac{1}{2} \sum_{i \in \mathcal{X}} \left(\sqrt{\mathbb{P}[X_0 = i]} - \sqrt{\mathbb{P}[X_1 = i]} \right)^2$

Outline

- Problem setting
- **Our results**
- Previous work
- Multicalibration Theorem
- Proof overviews
- Conclusion

Our work

- Reduce computational distinguishability of X_0, X_1 to statistical distinguishability of some $\widetilde{X}_0 \approx X_0$ and $\widetilde{X}_1 \approx X_1$

Our work

- Reduce computational distinguishability of X_0, X_1 to statistical distinguishability of some $\widetilde{X}_0 \approx X_0$ and $\widetilde{X}_1 \approx X_1$

- Prove tight, instance-optimal characterization of sample complexity of efficient distinguishers:

$$k = \Theta(d_H^{-2}(\widetilde{X}_0, \widetilde{X}_1))$$

Our work

- Reduce computational distinguishability of X_0, X_1 to statistical distinguishability of some $\widetilde{X}_0 \approx X_0$ and $\widetilde{X}_1 \approx X_1$
- Prove tight, instance-optimal characterization of sample complexity of efficient distinguishers:

$$k = \Theta(d_H^{-2}(\widetilde{X}_0, \widetilde{X}_1))$$

- Proof relies on Multicalibration Theorem from algorithmic fairness
  - Previous applications to graph theory and complexity theory [Dwork, Lee, Lin, Tankala 2023], [Casacuberta, Dwork, Vadhan 2024]

Roadmap

Multicalibration Theorem (for functions)

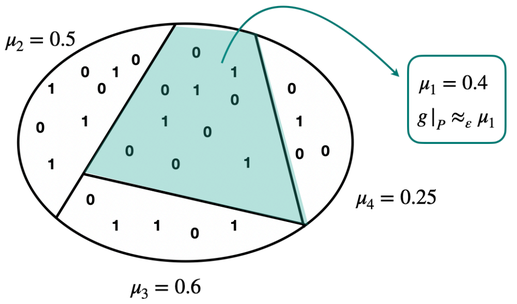

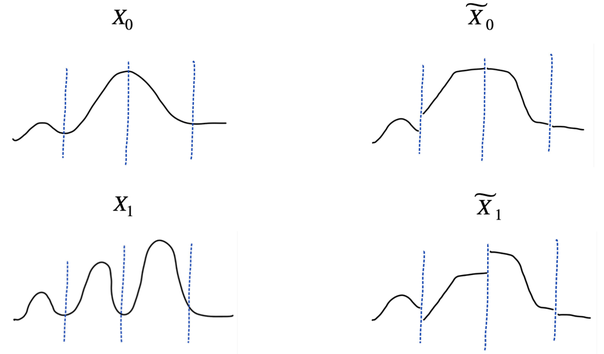

Multicalibration for Distinguishing

Main Theorem

Pseudo-Hellinger Characterization

Main theorem

- **Theorem [M, Putterman, Vadhan].** For every pair of random variables X_0, X_1, every integer s and $\varepsilon > 0$, $\exists$ random variables $\widetilde{X}_0, \widetilde{X}_1$ s.t. $\forall k > 0$,

  (1) $X_b \approx_\varepsilon \widetilde{X}_b$ by circuits of size s, for $b \in \{0,1\}$.

Main theorem

- **Theorem [M, Putterman, Vadhan].** For every pair of random variables X_0, X_1, every integer s and $\varepsilon > 0$, $\exists$ random variables $\widetilde{X}_0, \widetilde{X}_1$ s.t. $\forall k > 0$,

  (1) $X_b \approx_\varepsilon \widetilde{X}_b$ by circuits of size s, for $b \in \{0,1\}$.

  (2) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) + 2k\varepsilon)$-indistinguishable by circuits of size s.

Main theorem

- **Theorem [M, Putterman, Vadhan].** For every pair of random variables X_0, X_1, every integer s and $\varepsilon > 0$, $\exists$ random variables $\widetilde{X}_0, \widetilde{X}_1$ s.t. $\forall k > 0$,

  (1) $X_b \approx_\varepsilon \widetilde{X}_b$ by circuits of size s, for $b \in \{0,1\}$.

  (2) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) + 2k\varepsilon)$-indistinguishable by circuits of size s.

  (3) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) - 2k\varepsilon)$-distinguishable by circuits of size $s' = O(sk/\varepsilon^6) + \text{poly}(k/\varepsilon)^*$.

* Parameters are tightened by Dwork and Tankala in subsequent work.

Main theorem

- **Theorem [M, Putterman, Vadhan].** For every pair of random variables X_0, X_1, every integer s and $\varepsilon > 0$, $\exists$ random variables $\widetilde{X}_0, \widetilde{X}_1$ s.t. $\forall k > 0$,

  (1) $X_b \approx_\varepsilon \widetilde{X}_b$ by circuits of size s, for $b \in \{0,1\}$.

  (2) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) + 2k\varepsilon)$-indistinguishable by circuits of size s.

  (3) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) - 2k\varepsilon)$-distinguishable by circuits of size $s' = O(sk/\varepsilon^6) + \text{poly}(k/\varepsilon)^*$.

  (4) Above also holds with $\widetilde{X}_1 = X_1$, for $s' = O(sk/\varepsilon^{12}) + \text{poly}(k/\varepsilon)^*$.

* Parameters are tightened by Dwork and Tankala in subsequent work.

Roadmap

Multicalibration Theorem (for functions)

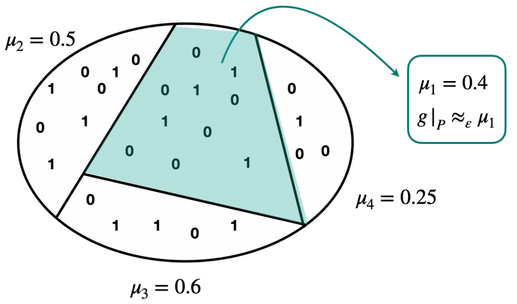

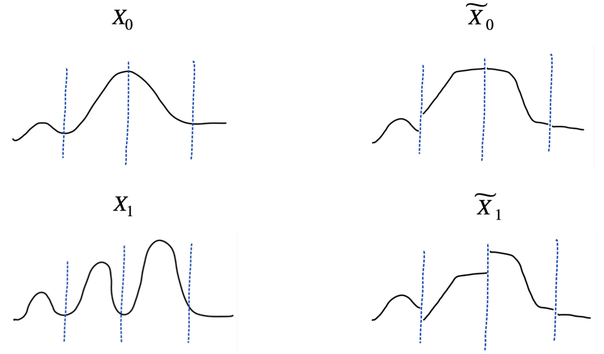

Multicalibration for Distinguishing

Main Theorem

Pseudo-Hellinger Characterization

Pseudo-Hellinger distance

- Main Theorem implies: distinguishing X_0, X_1 needs $k = \Theta(\delta^{-2})$ samples, for δ = "pseudo-Hellinger distance":

Pseudo-Hellinger distance

- Main Theorem implies: distinguishing X_0, X_1 needs $k = \Theta(\delta^{-2})$ samples, for δ = "pseudo-Hellinger distance":

- **Definition (Pseudo-Hellinger distance)** (s, ε)-pseudo-Hellinger distance between X_0, X_1 is *smallest* δ s.t. $\exists\, \widetilde{X}_0, \widetilde{X}_1$ satisfying:

  - (1) $X_0 \approx_\varepsilon \widetilde{X}_0$ for circuits of size s.

  - (2) $X_1 \approx_\varepsilon \widetilde{X}_1$ for circuits of size s.

  - (3) $d_H(\widetilde{X}_0, \widetilde{X}_1) \leq \delta$.

Pseudo-Hellinger distance

- Main Theorem implies: distinguishing X_0, X_1 needs $k = \Theta(\delta^{-2})$ samples, for δ = "pseudo-Hellinger distance":

- **Definition (Pseudo-Hellinger distance)** (s, ε)-pseudo-Hellinger distance between X_0, X_1 is *smallest* δ s.t. $\exists\, \widetilde{X}_0, \widetilde{X}_1$ satisfying:
  - (1) $X_0 \approx_\varepsilon \widetilde{X}_0$ for circuits of size s.
  - (2) $X_1 \approx_\varepsilon \widetilde{X}_1$ for circuits of size s.
  - (3) $d_H(\widetilde{X}_0, \widetilde{X}_1) \leq \delta$.

- $k = \Theta(\delta^{-2})$ is *computational analog* of statistical setting, where $k = \Theta(d_H^{-2}(X_0, X_1))$.

Outline

- Problem setting
- Our results
- **Previous work**
- Multicalibration Theorem
- Proof overviews
- Conclusion

Relationship to previous work

- **Computational distinguishability:** [Halevi-Rabin 2008, Geier 2022]
  If X_0, X_1 are δ-indistinguishable for size s circuits, then $X_0^{\otimes k}, X_1^{\otimes k}$ are $(1-(1-\delta)^k+\varepsilon)$-indistinguishable for circuits of size $s/\mathrm{poly}(k, 1/\varepsilon)$.

  - Parallels the statistical bound: $d_{TV}(X_0^{\otimes k}, X_1^{\otimes k}) \leq 1-(1-d_{TV}(X_0, X_1))^k$.

Relationship to previous work

- **Computational distinguishability:** [Halevi-Rabin 2008, Geier 2022]
  If X_0, X_1 are δ-indistinguishable for size s circuits, then $X_0^{\otimes k}, X_1^{\otimes k}$ are $(1-(1-\delta)^k+\varepsilon)$-indistinguishable for circuits of size $s/\mathrm{poly}(k, 1/\varepsilon)$.
  - Parallels the statistical bound: $d_{TV}(X_0^{\otimes k}, X_1^{\otimes k}) \leq 1-(1-d_{TV}(X_0, X_1))^k$.
  - This is *not instance optimal*.
  - A computational analog of $d_{TV}(X_0^{\otimes k}, X_1^{\otimes k})$ is missing.
  - How many samples to distinguish with constant advantage?

Outline

- Problem setting
- Our results
- Previous work
- **Multicalibration Theorem**
- Proof overviews
- Conclusion

Roadmap

Multicalibration Theorem (for functions)

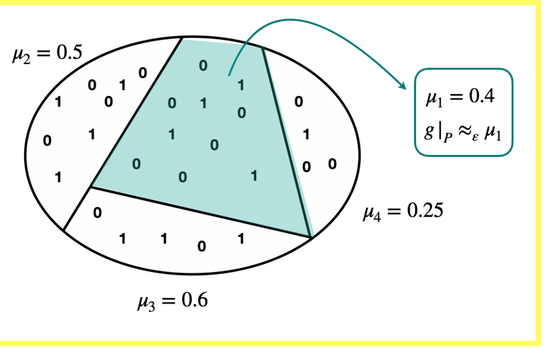

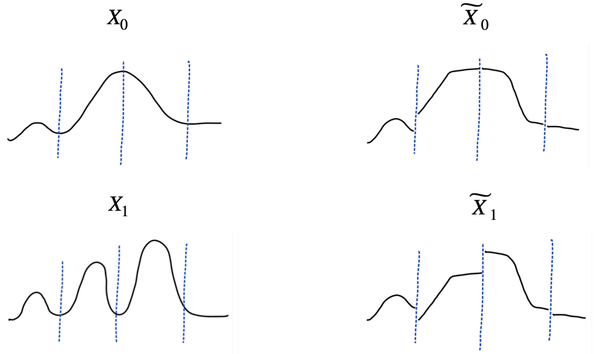

Multicalibration for Distinguishing

Main Theorem

Pseudo-Hellinger Characterization

Multicalibration Theorem

Key technical ingredient, from algorithmic fairness

- **Theorem:** [Hébert-Johnson, Kim, Reingold, Rothblum 2018*]
  For all randomized $g : \{0,1\}^n \to \{0,1\}$, distribution X on $\{0,1\}^n$, $s \in \mathbb{N}$, $\varepsilon > 0$, there exists partition $\{0,1\}^n = P_0 \cup P_1 \cup \ldots \cup P_k$, $k = O(1/\varepsilon)$ such that:

  - For every part P_i of partition except for P_0, g is ε-indistinguishable from its expectation (wrt X) over P_i by size s circuits.

* As formulated by Casacuberta, Dwork, and Vadhan, 2024.

Multicalibration Theorem

Key technical ingredient, from algorithmic fairness

- **Theorem:** [Hébert-Johnson, Kim, Reingold, Rothblum 2018*]
  For all randomized $g : \{0,1\}^n \to \{0,1\}$, distribution X on $\{0,1\}^n$, $s \in \mathbb{N}$, $\varepsilon > 0$, there exists partition $\{0,1\}^n = P_0 \cup P_1 \cup \ldots \cup P_k$, $k = O(1/\varepsilon)$ such that:

  - For every part P_i of partition except for P_0, g is ε-indistinguishable from its expectation (wrt X) over P_i by size s circuits.

  - Size of (possibly distinguishable) part P_0 is $\leq \varepsilon \cdot 2^n$.

* As formulated by Casacuberta, Dwork, and Vadhan, 2024.

Multicalibration Theorem

Key technical ingredient, from algorithmic fairness

- **Theorem:** [Hébert-Johnson, Kim, Reingold, Rothblum 2018*]
  For all randomized $g : \{0,1\}^n \to \{0,1\}$, distribution X on $\{0,1\}^n$, $s \in \mathbb{N}$, $\varepsilon > 0$, there exists partition $\{0,1\}^n = P_0 \cup P_1 \cup \ldots \cup P_k$, $k = O(1/\varepsilon)$ such that:

  - For every part P_i of partition except for P_0, g is ε-indistinguishable from its expectation (wrt X) over P_i by size s circuits.
  - Size of (possibly distinguishable) part P_0 is $\leq \varepsilon \cdot 2^n$.
  - Partition is computable by a circuit of size $s \cdot \text{poly}(1/\varepsilon)$.

* As formulated by Casacuberta, Dwork, and Vadhan, 2024.

Multicalibration Theorem

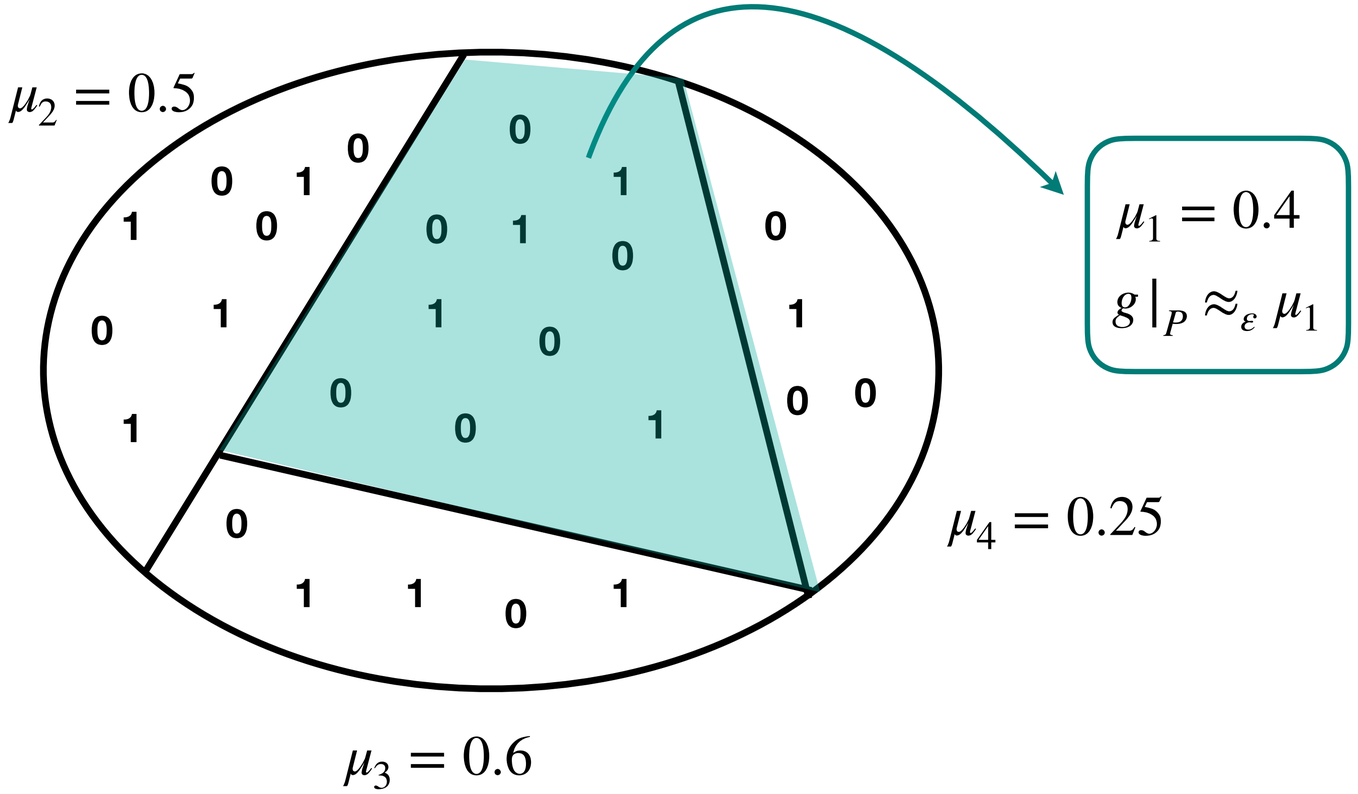

Multicalibration Theorem

Key technical ingredient, from algorithmic fairness

- **Theorem:** [Hébert-Johnson, Kim, Reingold, Rothblum 2018*] For all randomized $g : \{0,1\}^n \to \{0,1\}$, distribution X on $\{0,1\}^n$, $s \in \mathbb{N}$, $\varepsilon > 0$, exists a partition $\{0,1\}^n = P_0 \cup P_1 \cup \ldots \cup P_k$, $k = O(1/\varepsilon)$ such that:
  - For $i = 1,2,\ldots,k$, $(X, g(X))|_{X \in P_i}$ is ε-indistinguishable from $(X|_{X \in P_i}, \text{Bern}(\mu_i))$ by size s circuits, where $\mu_i = \mathbb{E}\left[g(X) \,|\, X \in P_i\right]$.
  - Let $\alpha_i = |P_i|/2^n$. Then $\alpha_0 \leq \varepsilon$.
  - There is a circuit $p : \{0,1\}^n \to \{0,1,\ldots,k\}$ of size $s \cdot \text{poly}(1/\varepsilon)$ such that $P_i = p^{-1}(i)$.

* As formulated by Casacuberta, Dwork, and Vadhan, 2024.

Roadmap

Multicalibration Theorem (for functions)

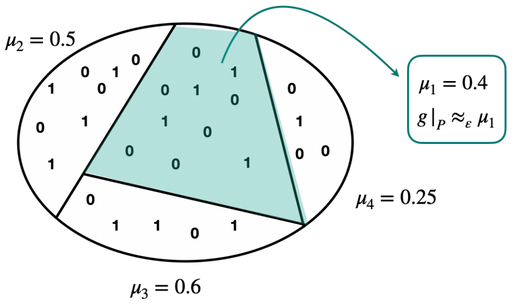

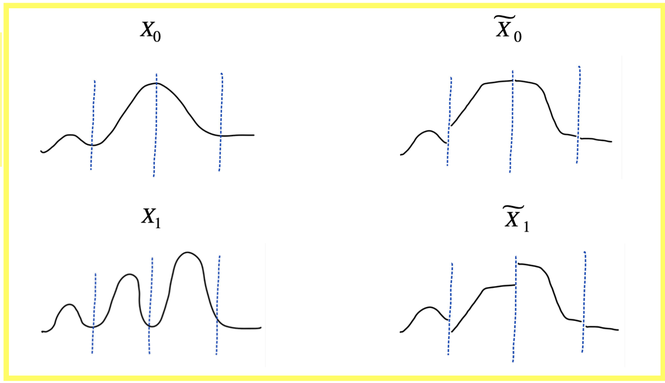

Multicalibration for Distinguishing

Main Theorem

Pseudo-Hellinger Characterization

Implications of Multicalibration

- **Lemma** [Casacuberta, Dwork, Vadhan 2024]
  If g is ε-indistinguishable from its expectation $\mu_i = \mathbb{E}\left[g(X) \mid X \in P_i\right]$ in P_i wrt size s circuits, then $X_{P_i}\big|_{g(x)=1}$ is ε'-indistinguishable from $X_{P_i}\big|_{g(x)=0}$ wrt size s' circuits, for slightly smaller ε', s'.

Implications of Multicalibration

- **Lemma** [Casacuberta, Dwork, Vadhan 2024]
  If g is ε-indistinguishable from its expectation $\mu_i = \mathbb{E}\left[g(X) \mid X \in P_i\right]$ in P_i wrt size s circuits, then $X_{P_i}|_{g(x)=1}$ is ε'-indistinguishable from $X_{P_i}|_{g(x)=0}$ wrt size s' circuits, for slightly smaller ε', s'.

- Multicalibration compares a *function* to its *expectation.* We want to compare *distributions*.

- Choose g, X such that $X_{P_i}|_{g(x)=1} = X_1|_{P_i}$ and $X_{P_i}|_{g(x)=0} = X_0|_{P_i}$, and define $\widetilde{X}_0$ by $\mathbb{P}[\widetilde{X}_0 = x] = \mathbb{P}[X_0 \in P_i] \cdot \mathbb{P}[X|_{P_i} = x]$ (and likewise for $\widetilde{X}_1$).

What we gain from multicalibration

It allows us to construct random variables $\widetilde{X}_0, \widetilde{X}_1$ for X_0, X_1 with the following properties.

- **"Multicalibration for Distinguishing" Theorem [M, Putterman, Vadhan]:** For every X_0, X_1, every positive s, and every $\varepsilon > 0$, there exists a "low-complexity" partition function $p : \mathcal{X} \to [m]$ for $m = O(1/\varepsilon)$, and random variables $\widetilde{X}_0, \widetilde{X}_1$ such that:

What we gain from multicalibration

It allows us to construct random variables $\widetilde{X}_0, \widetilde{X}_1$ for X_0, X_1 with the following properties.

- **"Multicalibration for Distinguishing" Theorem [M, Putterman, Vadhan]:** For every X_0, X_1, every positive s, and every $\varepsilon > 0$, there exists a "low-complexity" partition function $p : \mathcal{X} \to [m]$ for $m = O(1/\varepsilon)$, and random variables $\widetilde{X}_0, \widetilde{X}_1$ such that:
  - $\widetilde{X}_0 \approx_\varepsilon X_0$ and $\widetilde{X}_1 \approx_\varepsilon X_1$ for circuits of size s.

What we gain from multicalibration

It allows us to construct random variables $\widetilde{X}_0, \widetilde{X}_1$ for X_0, X_1 with the following properties.

- **"Multicalibration for Distinguishing" Theorem [M, Putterman, Vadhan]:** For every X_0, X_1, every positive s, and every $\varepsilon > 0$, there exists a "low-complexity" partition function $p : \mathcal{X} \to [m]$ for $m = O(1/\varepsilon)$, and random variables $\widetilde{X}_0, \widetilde{X}_1$ such that:
  - $\widetilde{X}_0 \approx_\varepsilon X_0$ and $\widetilde{X}_1 \approx_\varepsilon X_1$ for circuits of size s.
  - $p(X_0) \stackrel{d}{=} p(\widetilde{X}_0)$ and $p(X_1) \stackrel{d}{=} p(\widetilde{X}_1)$.

What we gain from multicalibration

It allows us to construct random variables $\widetilde{X}_0, \widetilde{X}_1$ for X_0, X_1 with the following properties.

- **"Multicalibration for Distinguishing" Theorem [M, Putterman, Vadhan]:** For every X_0, X_1, every positive s, and every $\varepsilon > 0$, there exists a "low-complexity" partition function $p : \mathcal{X} \to [m]$ for $m = O(1/\varepsilon)$, and random variables $\widetilde{X}_0, \widetilde{X}_1$ such that:
  - $\widetilde{X}_0 \approx_\varepsilon X_0$ and $\widetilde{X}_1 \approx_\varepsilon X_1$ for circuits of size s.
  - $p(X_0) \stackrel{d}{=} p(\widetilde{X}_0)$ and $p(X_1) \stackrel{d}{=} p(\widetilde{X}_1)$.
  - In every part of the partition, $\widetilde{X}_0, \widetilde{X}_1$ are identically distributed.

Multicalibration for Distinguishing

X_0

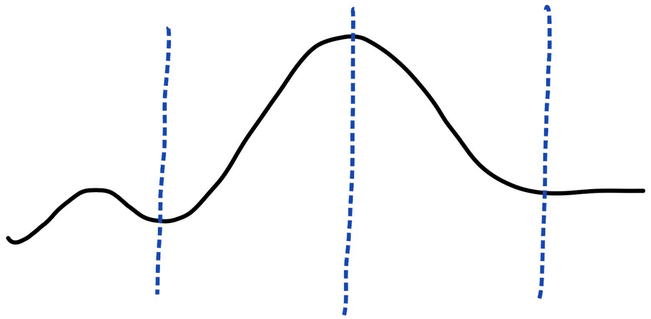

$\widetilde{X}_0$

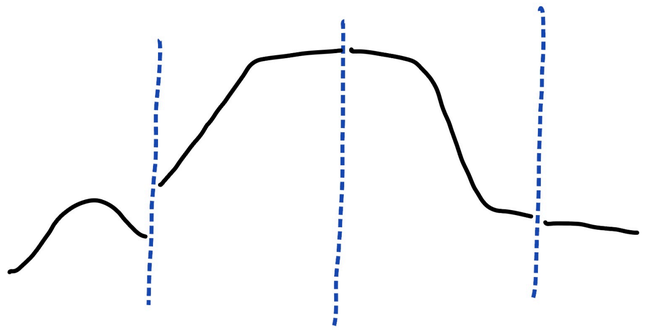

X_1

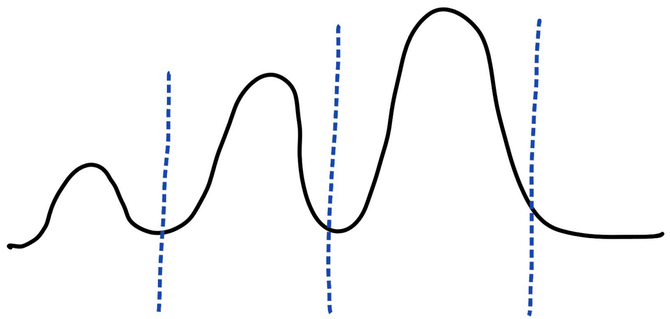

$\widetilde{X}_1$

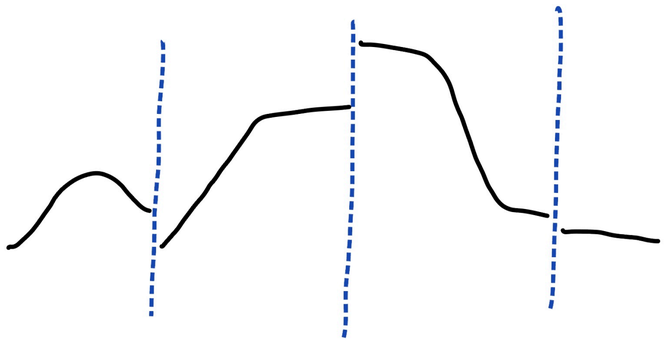

Multicalibration for Distinguishing

X_0 $\approx_\varepsilon$ $\widetilde{X}_0$

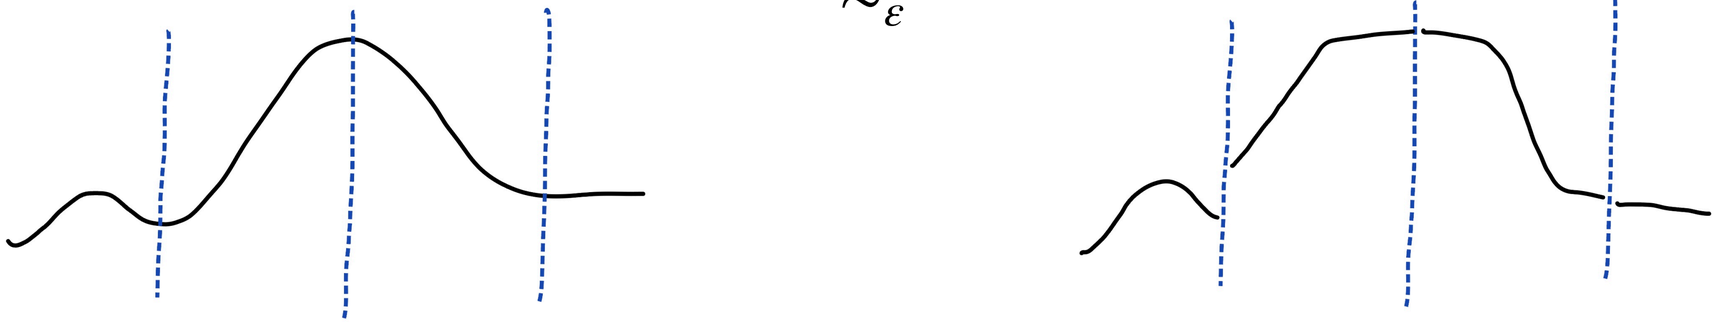

X_1 $\approx_\varepsilon$ $\widetilde{X}_1$

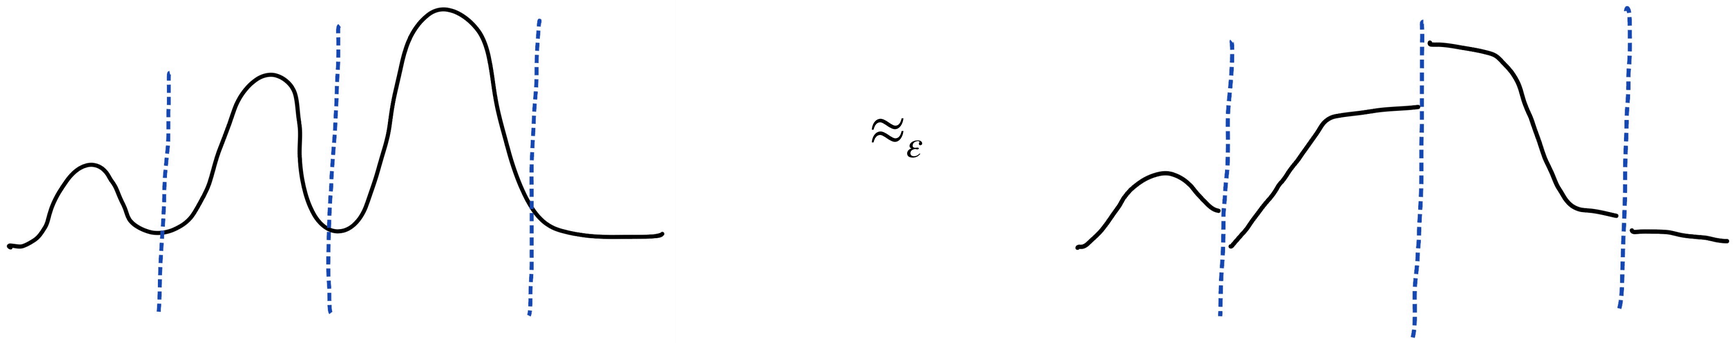

Recall: $\approx_\varepsilon$ means "ε-indistinguishable with respect to size-s circuits"

Multicalibration for Distinguishing

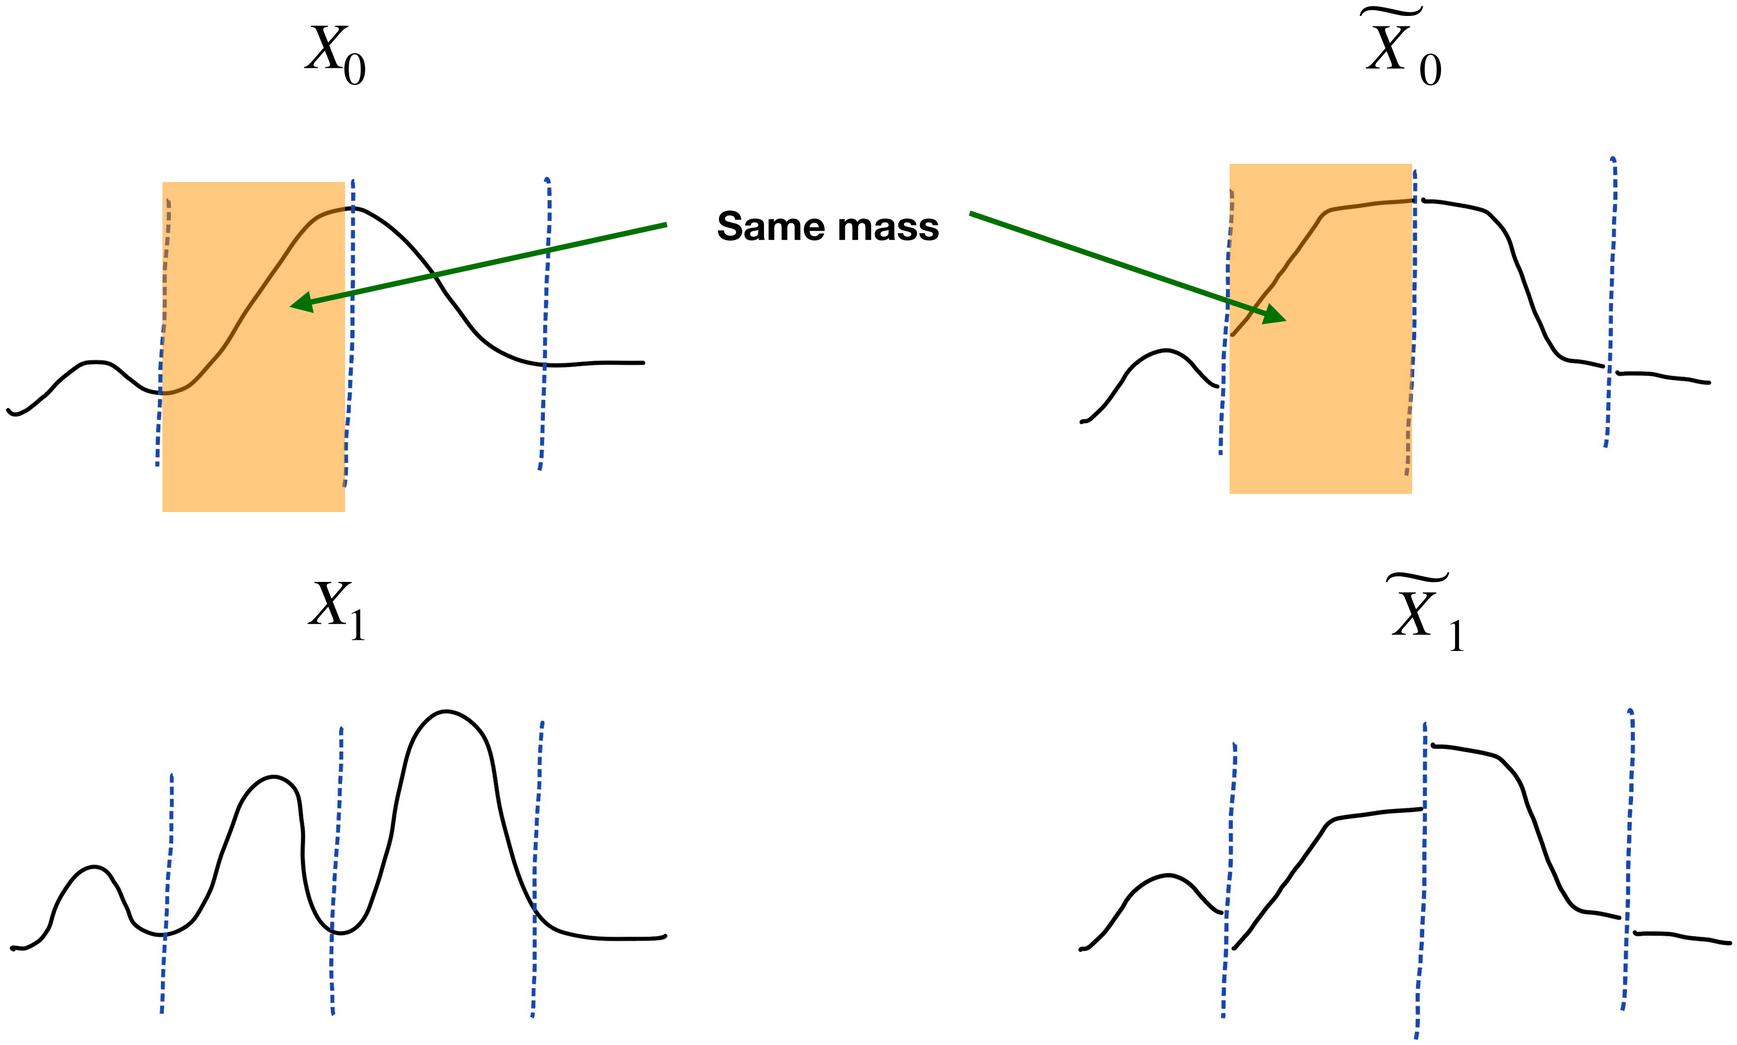

Multicalibration for Distinguishing

X_0

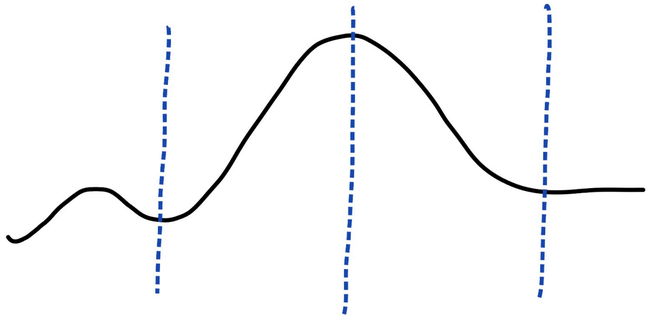

$\widetilde{X}_0$

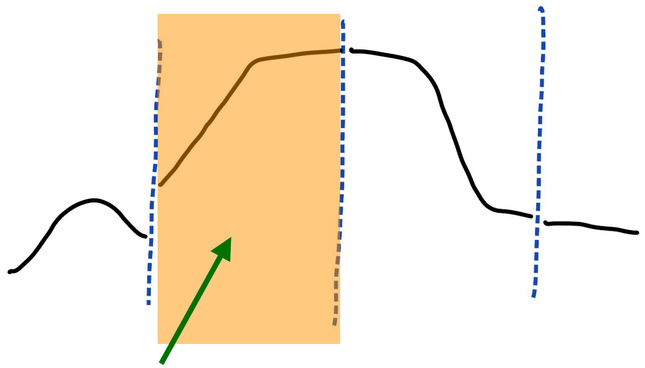

$\stackrel{d}{=}$

X_1

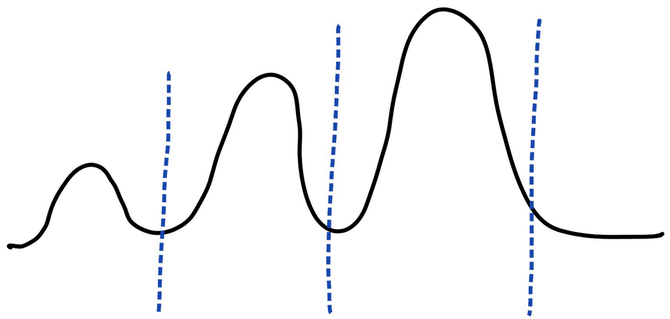

$\widetilde{X}_1$

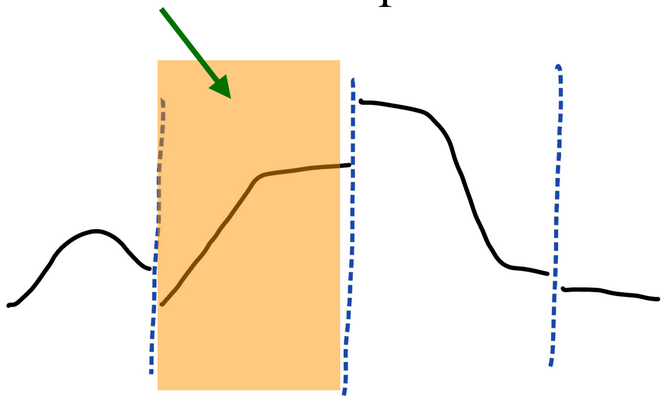

Outline

- Problem setting
- Our results
- Previous work
- Multicalibration Theorem
- **Proof overviews**
- Conclusion

Roadmap

Multicalibration Theorem (for functions)

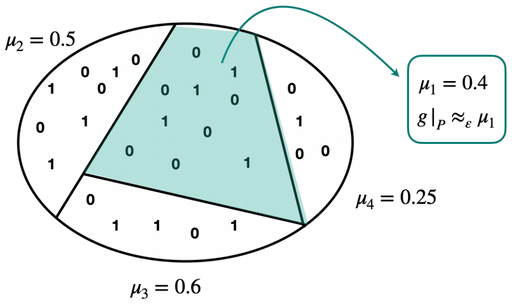

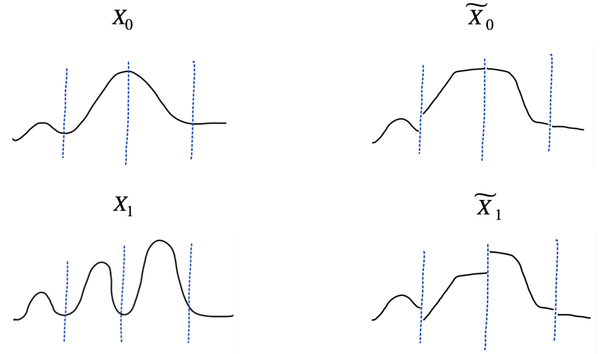

Multicalibration for Distinguishing

Main Theorem

Pseudo-Hellinger Characterization

Main theorem (reminder)

- **Theorem [M, Putterman, Vadhan].** For every pair of random variables X_0, X_1, every integer s and $\varepsilon > 0$, $\exists$ random variables $\widetilde{X}_0, \widetilde{X}_1$ s.t. $\forall k > 0$,

(1) $X_b \approx_\varepsilon \widetilde{X}_b$ by circuits of size s, for $b \in \{0,1\}$. ✓

(2) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) + 2k\varepsilon)$-indistinguishable by circuits of size s.

(3) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) - 2k\varepsilon)$-distinguishable by circuits of size $s' = O(sk/\varepsilon^6) + \text{poly}(k/\varepsilon)$.

Main theorem (reminder)

- **Theorem [M, Putterman, Vadhan].** For every pair of random variables X_0, X_1, every integer s and $\varepsilon > 0$, $\exists$ random variables $\widetilde{X}_0, \widetilde{X}_1$ s.t. $\forall k > 0$,

  (1) $X_b \approx_\varepsilon \widetilde{X}_b$ by circuits of size s, for $b \in \{0,1\}$.

  (2) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) + 2k\varepsilon)$-indistinguishable by circuits of size s.

  (3) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) - 2k\varepsilon)$-distinguishable by circuits of size $s' = O(sk/\varepsilon^6) + \text{poly}(k/\varepsilon)$.

Proof Outline of Indistinguishability

- MC for distinguishing Theorem gives: $\widetilde{X}_0 \approx_\varepsilon X_0$ and $\widetilde{X}_1 \approx_\varepsilon X_1$ for circuits of size s.
- By a hybrid argument: $\widetilde{X}_0^{\otimes k}, X_0^{\otimes k}$ and $\widetilde{X}_1^{\otimes k}, X_1^{\otimes k}$ are $k\varepsilon$-indistinguishable for circuits of size s.
- Thus, size s circuits cannot distinguish $X_0^{\otimes k}, X_1^{\otimes k}$ better than $d_{TV}\left(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}\right) + 2k\varepsilon$.

Main theorem (reminder)

- **Theorem [M, Putterman, Vadhan].** For every pair of random variables X_0, X_1, every integer s and $\varepsilon > 0$, $\exists$ random variables $\widetilde{X}_0, \widetilde{X}_1$ s.t. $\forall k > 0$,

  (1) $X_b \approx_\varepsilon \widetilde{X}_b$ by circuits of size s, for $b \in \{0,1\}$.

  (2) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) + 2k\varepsilon)$-indistinguishable by circuits of size s.

  (3) $X_0^{\otimes k}$ and $X_1^{\otimes k}$ are $(d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k}) - 2k\varepsilon)$-distinguishable by circuits of size $s' = O(sk/\varepsilon^6) + \text{poly}(k/\varepsilon)$.

Proof Outline of Distinguishability

- Start with distinguishing $\widetilde{X}_0, \widetilde{X}_1$.

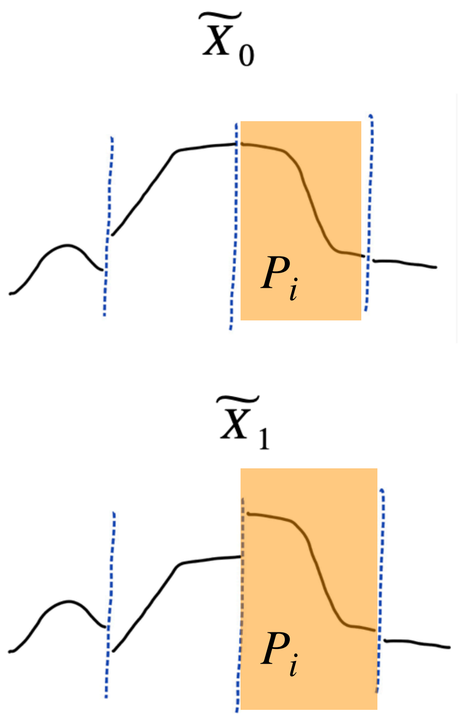

Proof Outline of Distinguishability

- Start with distinguishing $\widetilde{X}_0, \widetilde{X}_1$.
- Consider x sampled from either $\widetilde{X}_0, \widetilde{X}_1$. Let P_i be the part of partition x is in.

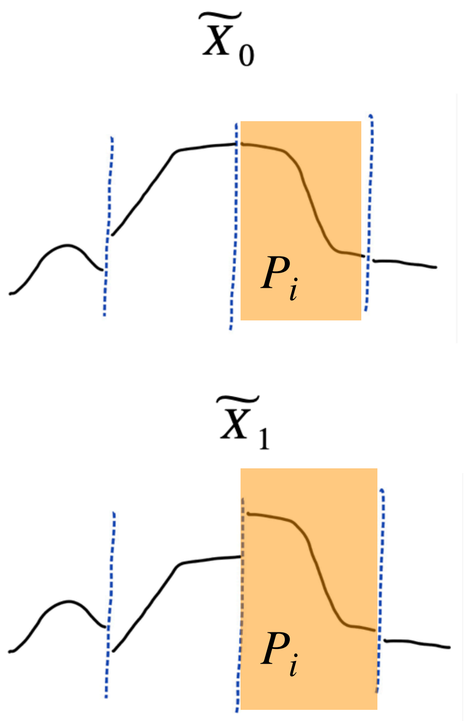

Proof Outline of Distinguishability

- Start with distinguishing $\widetilde{X}_0, \widetilde{X}_1$.
- Consider x sampled from either $\widetilde{X}_0, \widetilde{X}_1$. Let P_i be the part of partition x is in.
- $\widetilde{X}_{0|P_i} \overset{d}{=} \widetilde{X}_{1|P_i} \Rightarrow$ no information from sample x *besides* label i of part x is in.

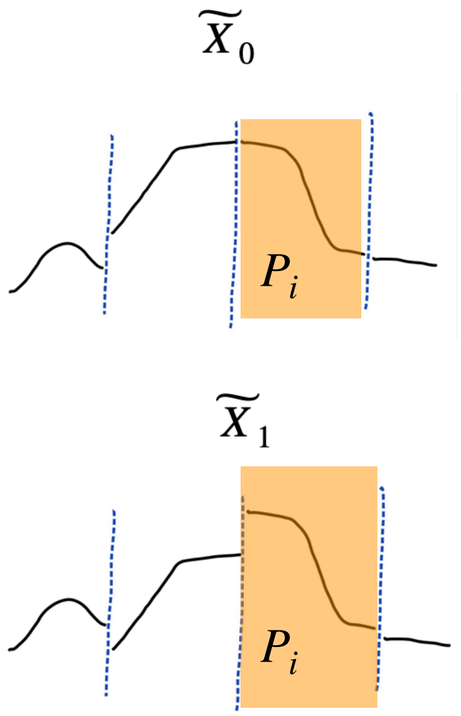

Proof Outline of Distinguishability

- Start with distinguishing $\widetilde{X}_0, \widetilde{X}_1$.
- Consider x sampled from either $\widetilde{X}_0, \widetilde{X}_1$. Let P_i be the part of partition x is in.
- $\widetilde{X}_{0|P_i} \stackrel{d}{=} \widetilde{X}_{1|P_i} \Rightarrow$ no information from sample x *besides* label i of part x is in.
- Recipe for *optimal* distinguisher between $\widetilde{X}_0, \widetilde{X}_1$:

  For each sample x seen, compute $P_i : x \in P_i$.

  See if $\mathbb{P}[\widetilde{X}_0 \in P_i] \geq \mathbb{P}[\widetilde{X}_1 \in P_i]$.

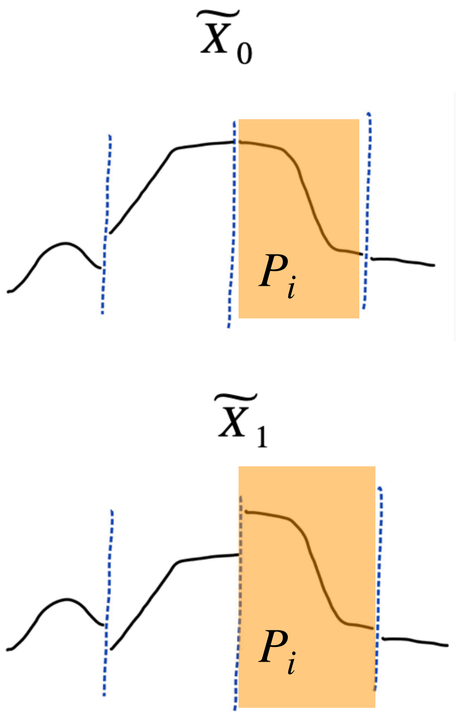

Proof Outline of Distinguishability

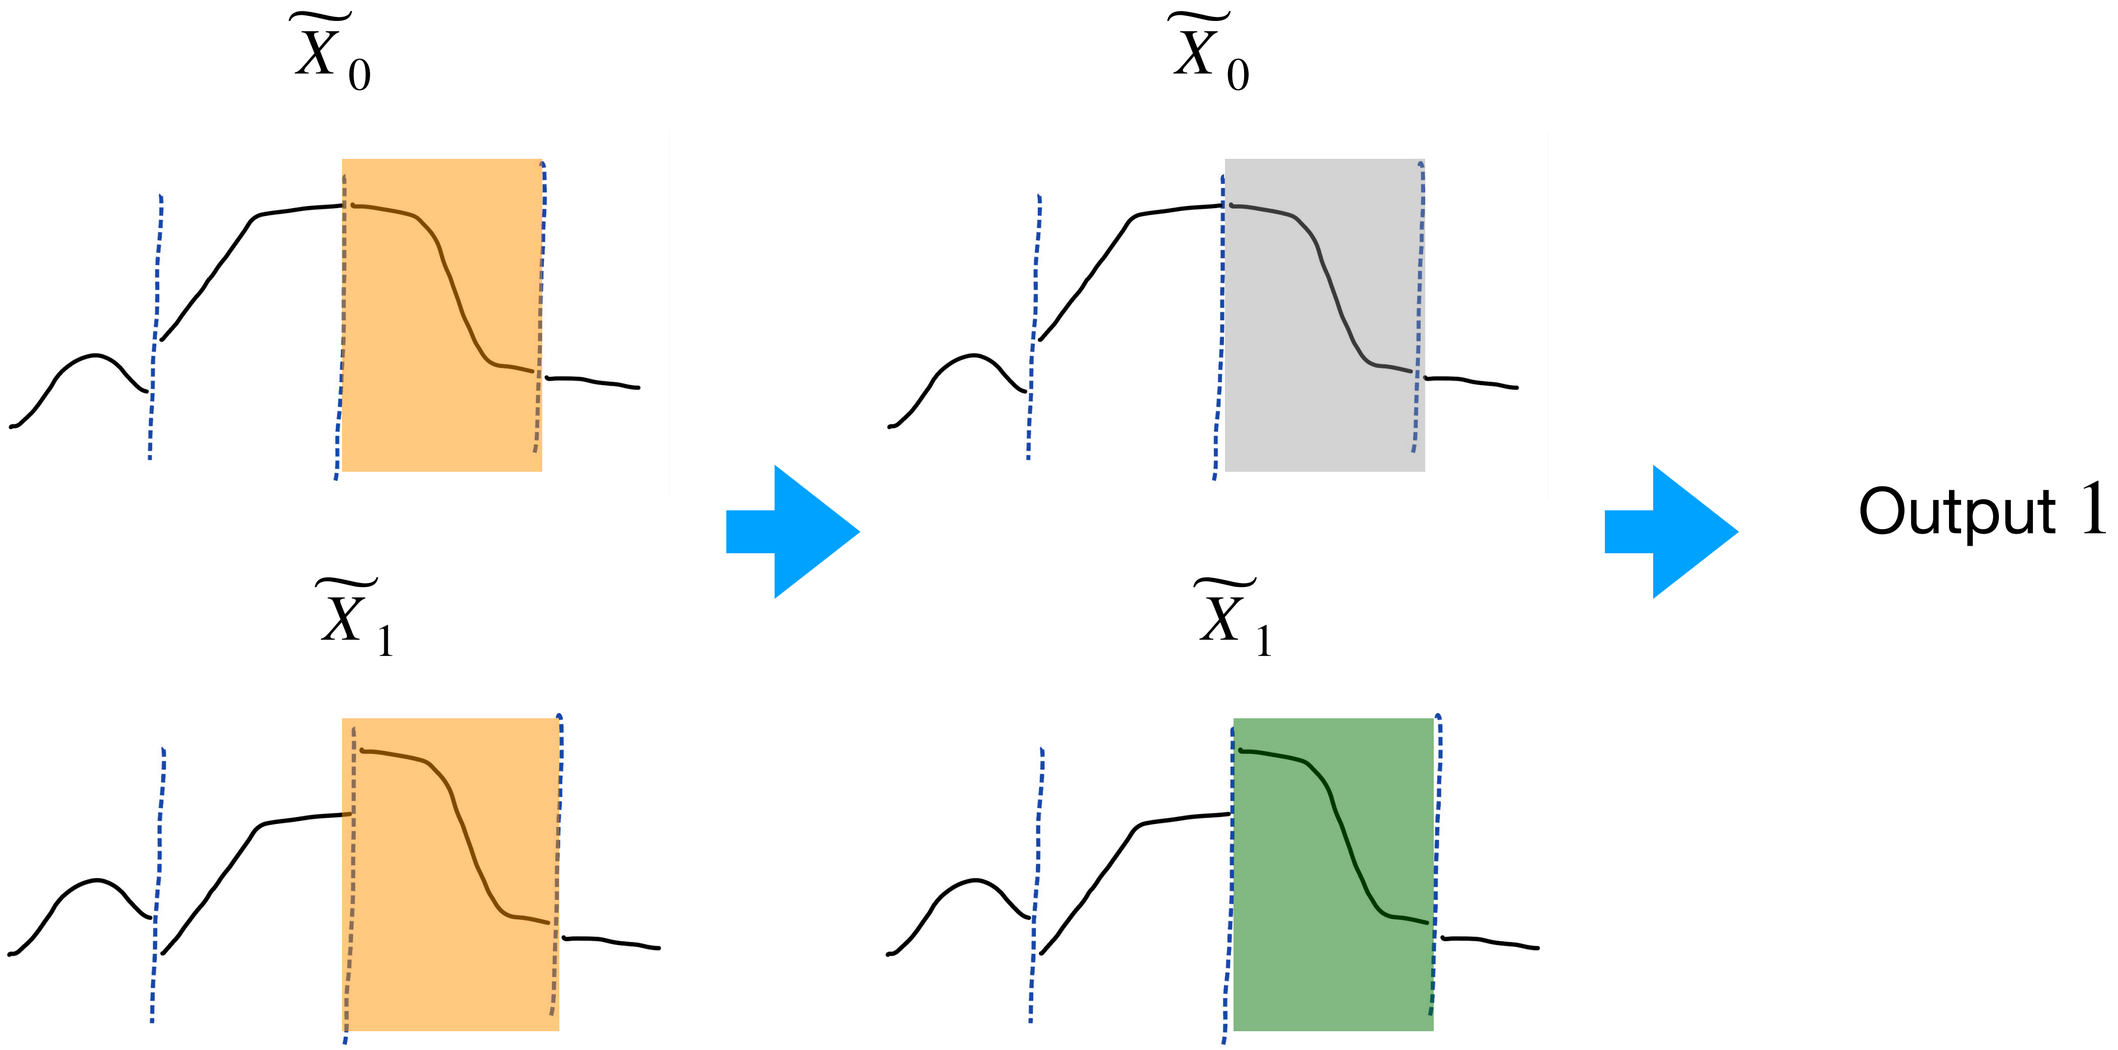

Proof Outline of Distinguishability

- Keep running product over samples $x_1, x_2, \ldots, x_k$:

$$L(x_1, x_2, \ldots, x_k) = \prod_{j=1}^{k} \frac{\mathbb{P}_{x \sim \widetilde{X}_1}\left[p(x) = p(x_j)\right]}{\mathbb{P}_{x \sim \widetilde{X}_0}\left[p(x) = p(x_j)\right]}.$$

If $L(x_1, x_2, \ldots, x_k) \geq 1$, output $\widetilde{X}_1$. Else output $\widetilde{X}_0$.

Proof Outline of Distinguishability

- Keep running product over samples $x_1, x_2, \ldots, x_k$:

$$L(x_1, x_2, \ldots, x_k) = \prod_{j=1}^{k} \frac{\mathbb{P}_{x \sim \widetilde{X}_1}\left[p(x) = p(x_j)\right]}{\mathbb{P}_{x \sim \widetilde{X}_0}\left[p(x) = p(x_j)\right]}.$$

  If $L(x_1, x_2, \ldots, x_k) \geq 1$, output $\widetilde{X}_1$. Else output $\widetilde{X}_0$.

- L is *likelihood ratio test: optimal distinguisher,* advantage $d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k})$.

  L can be encoded by small circuits because p is encoded by small circuits and has a small support.

Proof Outline of Distinguishability

- Keep running product over samples $x_1, x_2, \ldots, x_k$:

$$L(x_1, x_2, \ldots, x_k) = \prod_{j=1}^{k} \frac{\mathbb{P}_{x \sim \widetilde{X}_1}\left[p(x) = p(x_j)\right]}{\mathbb{P}_{x \sim \widetilde{X}_0}\left[p(x) = p(x_j)\right]}.$$

  If $L(x_1, x_2, \ldots, x_k) \geq 1$, output $\widetilde{X}_1$. Else output $\widetilde{X}_0$.

- L is *likelihood ratio test: optimal distinguisher,* advantage $d_{TV}(\widetilde{X}_0^{\otimes k}, \widetilde{X}_1^{\otimes k})$.

  L can be encoded by small circuits because p is encoded by small circuits and has a small support.

- $\widetilde{X}_0 \approx_\varepsilon X_0$ and $\widetilde{X}_1 \approx_\varepsilon X_1$ for small circuits $\Rightarrow$ L distinguishes $X_0^{\otimes k}, X_1^{\otimes k}$.

Outline

- Problem setting
- Our results
- Previous work
- Multicalibration Theorem
- Proof overviews
- **Conclusion**

Conclusion

- Prove a tight, instance-optimal characterization of the sample complexity of efficient distinguishers. For random variables $X_0 \approx \widetilde{X}_0$ and $X_1 \approx \widetilde{X}_1$,

$$k = \Theta(d_H^{-2}(\widetilde{X}_0, \widetilde{X}_1)).$$

- Key technical ingredient: Multicalibration Theorem from algorithmic fairness

Thank you!