A $d^{1/2+o(1)}$ Monotonicity Tester for Boolean Functions on d-Dimensional Hypergrids

FOCS 2023

Hadley Black
UCLA

Deeparnab Chakrabarty
Dartmouth

C. Seshadhri
UCSC

Monotonicity Testing

A central problem in property testing proposed by Goldreich-Goldwasser-Lehman-Ron-Samorodnitsky 99

- We consider $f\colon [n]^d \to \{0,1\}$
- f is **monotone** if $f(x) \leq f(y)$ whenever $x \prec y$
- Partial order: $x \preceq y$ iff $x_i \leq y_i, \forall i \in [d]$

Distance to monotonicity:

$$\varepsilon(f) = n^{-d} \cdot \min_{h \text{ monotone}} \# x : f(x) \neq h(x)$$

Given $f\colon [n]^d \to \{0,1\}$ and $\varepsilon > 0$...

1. if f monotone: **accept** w.p. $> 2/3$
2. if $\varepsilon(f) > \varepsilon$: **reject** w.p. $> 2/3$

* Non-adaptive queries

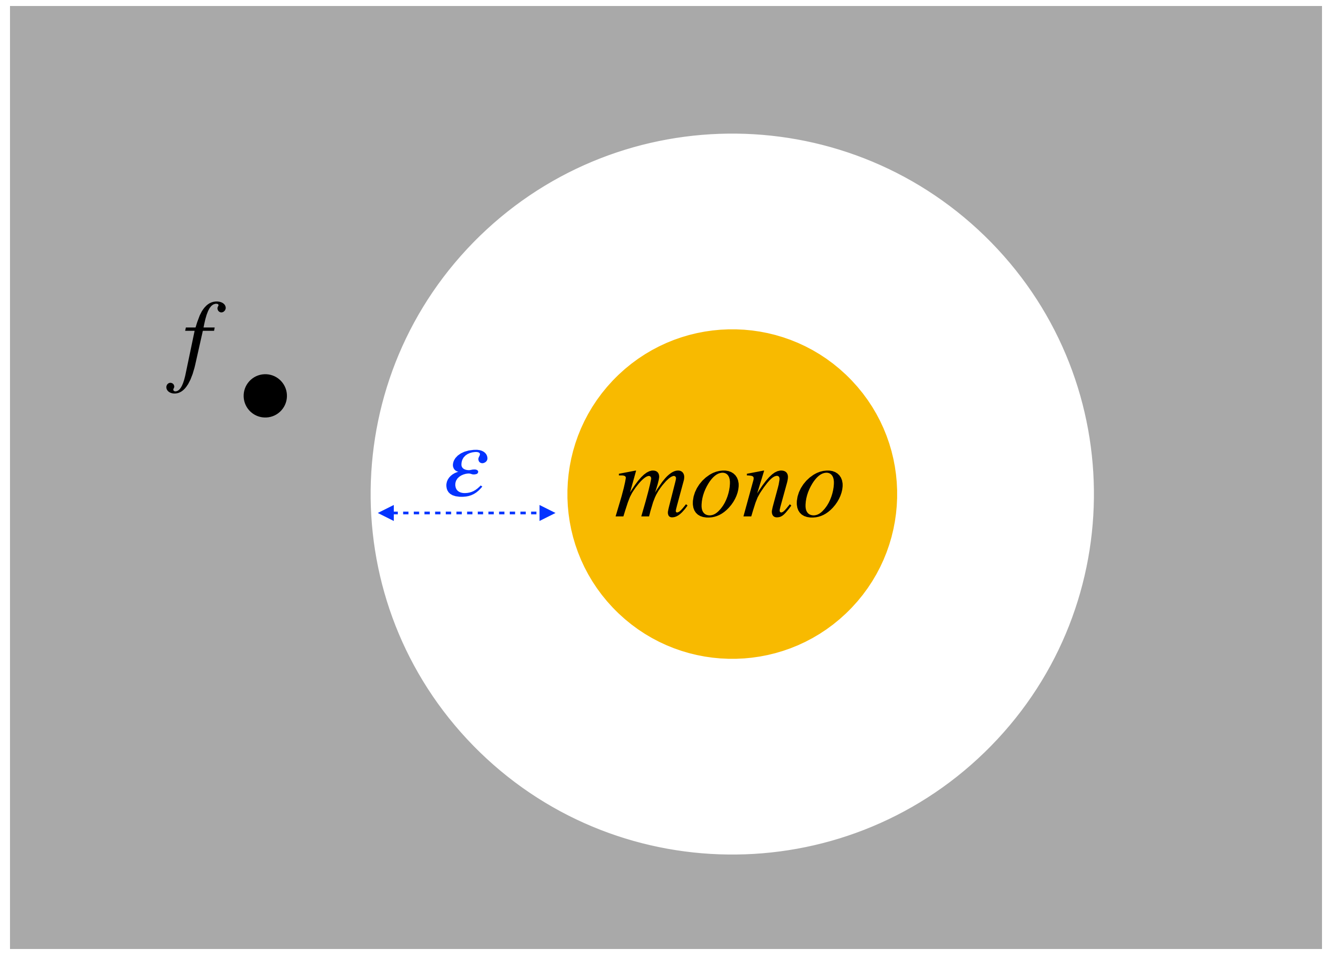

Abridged History of Non-adaptive Testing (for brevity let $\varepsilon = \Omega(1)$)

The Hypercube ($n = 2$)

Khot-Minzer-Safra FOCS 15 $\widetilde{O}(d^{1/2})$

Chen-Waingarten-Xie STOC 17 $\widetilde{\Omega}(d^{1/2})$

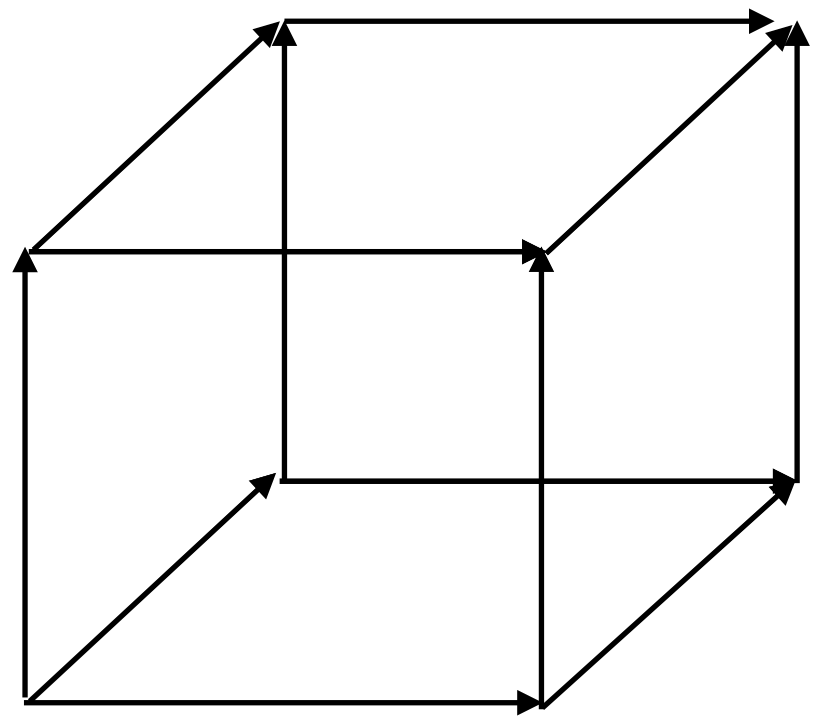

The **Hypergrid** ($n \geq 2$) $\approx$

Dodis-Goldreich-Lehman-Raskhodnikova-Ron-Samorodnitsky 99 and Berman-Raskhodnikova-Yaroslavtsev 14 $\widetilde{O}(d)$

Black-Chakrabarty-Seshadhri SODA 18, 20 $\widetilde{O}(d^{5/6})$

Braverman-Khot-Kindler-Minzer ITCS 23,
Black-Chakrabarty-Seshadhri STOC 23 $\widetilde{O}(\text{poly}(n) \cdot d^{1/2})$

Black-Chakrabarty-Seshadhri FOCS 23 $d^{1/2+o(1)}$

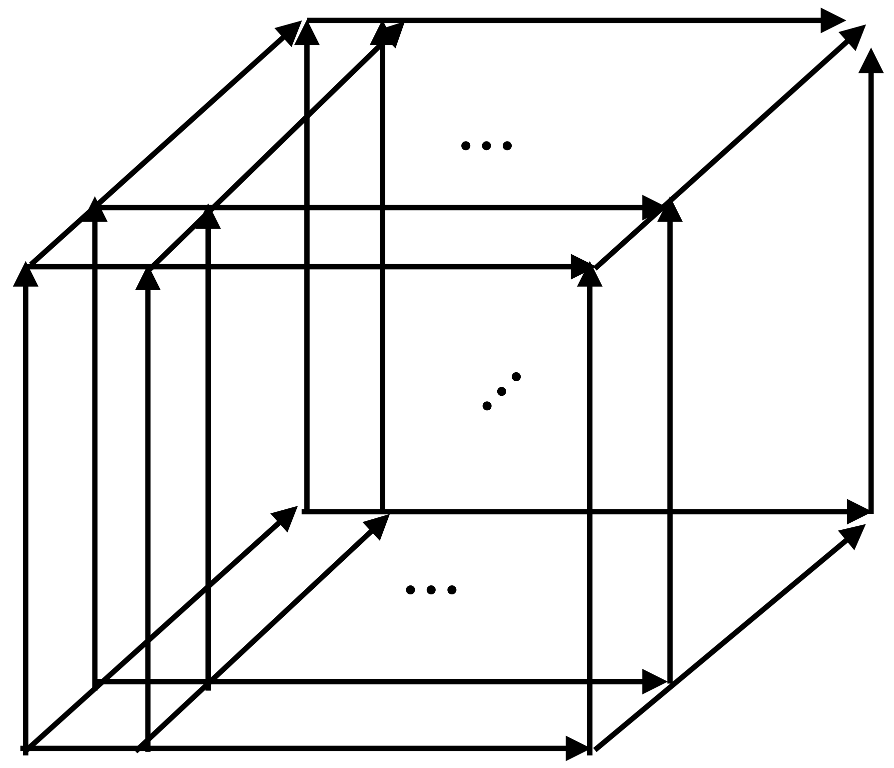

The Hypercube: a (very) brief history

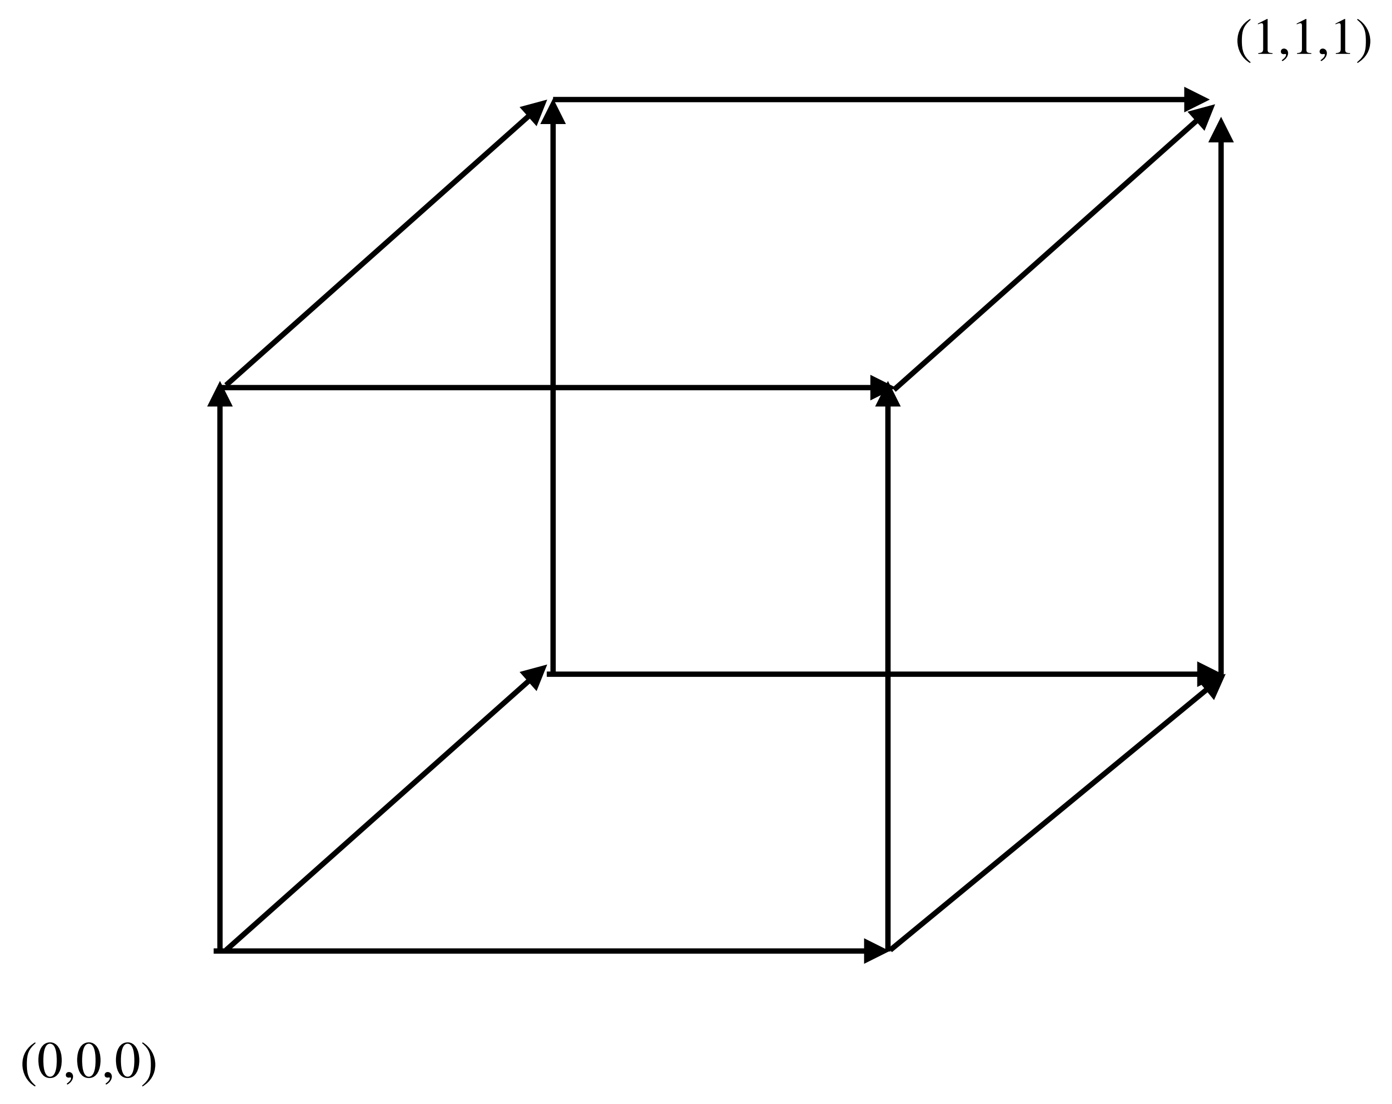

The Edge Tester

Edge tester (GGLRS [99])

- Sample an edge (x, y)
- Reject if $f(x) > f(y)$.

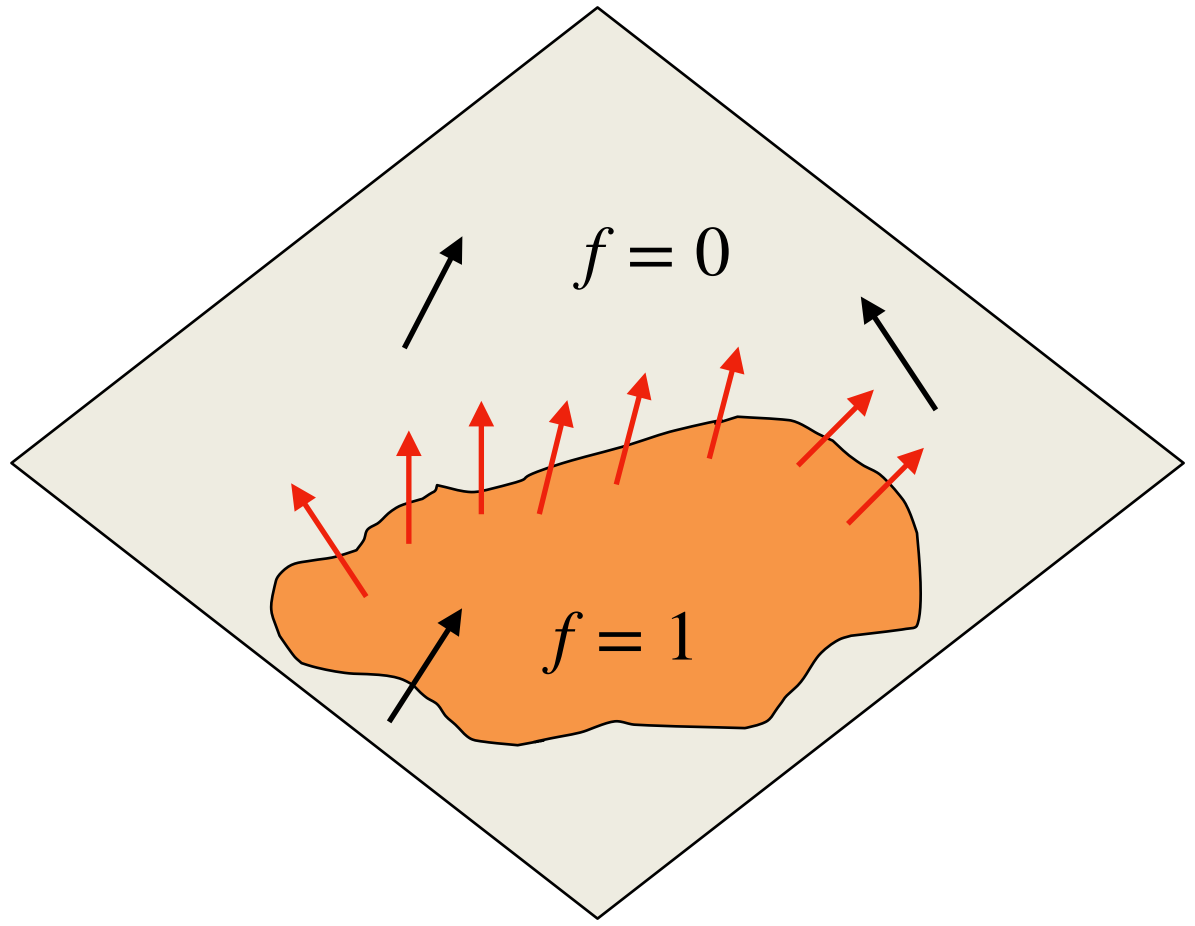

How many decreasing edges are there when $\varepsilon(f) > \varepsilon$?

Negative influence

$$I_f^- = \frac{\textbf{\# edges } (x, y) : f(x) > f(y)}{2^d}$$

Total influence

$$I_f = \frac{\textbf{\# edges } (x, y) : f(x) \neq f(y)}{2^d}$$

Theorem: (GGLRS [99]) $I_f^- \geq \Omega(\varepsilon(f))$

Theorem: (Poincaré) $I_f \geq \Omega(var(f))$

$\implies$ # decreasing edges $\geq \Omega(\varepsilon(f)) \cdot 2^d$

Total # edges is $d \cdot 2^{d-1}$

$\implies$ Edge test succeeds with probability $\Omega(\varepsilon/d)$

$\implies$ Repeat $O(d/\varepsilon)$ times!

Question:
Is this a tight analysis of the edge tester?

The Path Tester

Theorem: GGLRS [99] $I_f^- \geq \Omega(\varepsilon(f))$

- To beat $O(d)$ requires something other than the edge tester

Is this inequality tight? Yes.

anti-dictator function: $f(x) = 1 - x_1$

$\varepsilon(f) = 1/2$ and $I_f^- = 1/2$

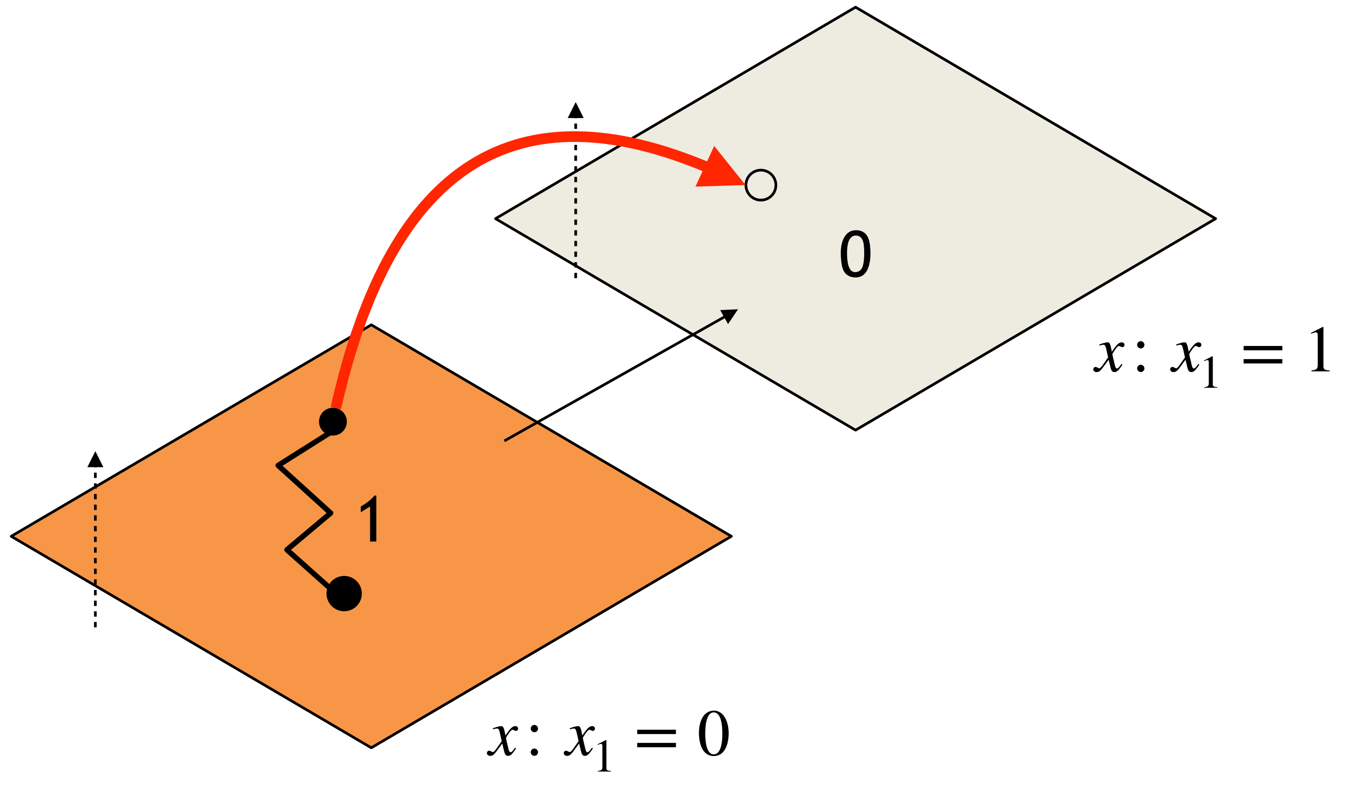

Path tester (CS [14])

- Sample $x \prec y$ which differ on $\approx d^{1/2}$ bits
- Reject if $f(x) > f(y)$

- We succeed with probability $\Omega(d^{-1/2})$ for the anti-dictator function

- Why? Decreasing edges are **spread** amongst the vertices

Question: Is there a more nuanced way to understand boundary?

Theorem: Talagrand [93] $\mathbb{E}_x[I_f(x)^{1/2}] = \Omega(var(f))$

Theorem: KMS [15] $\mathbb{E}_x[I_f^-(x)^{1/2}] = \widetilde{\Omega}(\varepsilon(f))$

$\Longrightarrow$ Can test with $\widetilde{O}(d^{1/2})$ queries by combining edge tester and path tester

The Hypergrid

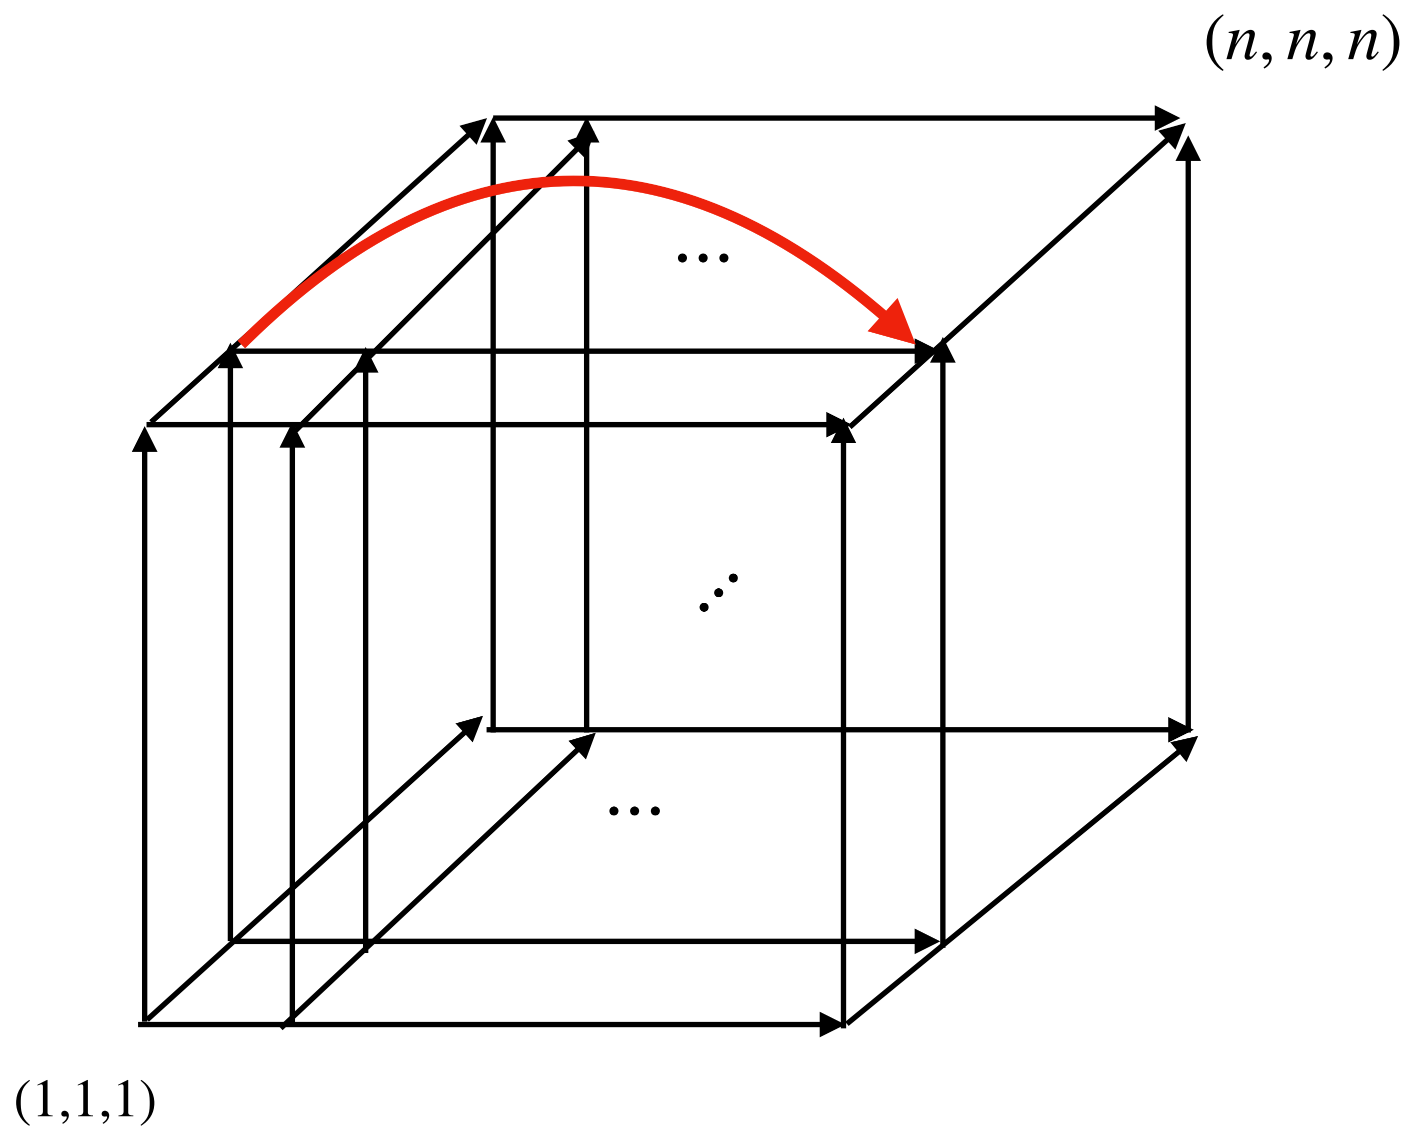

The (fully augmented) hypergrid:

DAG with…

- Vertex set: $[n]^d$
- Edges: (x, y) which differ on 1 coordinate **by any value**

The Directed Talagrand Inequality for Hypergrids (STOC 23)

Thresholded Influence:
Given $f\colon [n]^d \to \{0,1\}$ and $x \in [n]^d$

$\Phi_f(x) = \#i \in [d]:$ there exists a decreasing i-edge incident to x

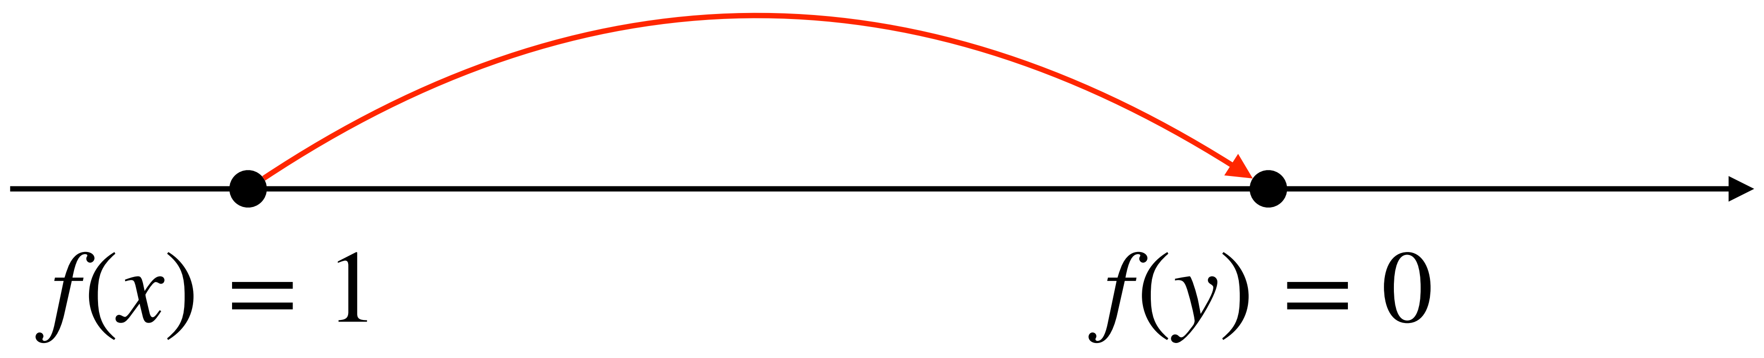

Theorem (BCS STOC 23): For any $f\colon [n]^d \to \{0,1\}$,

$$\mathbb{E}_x[\Phi_f(x)^{1/2}] = \Omega\left(\frac{\varepsilon(f)}{\log n}\right)$$

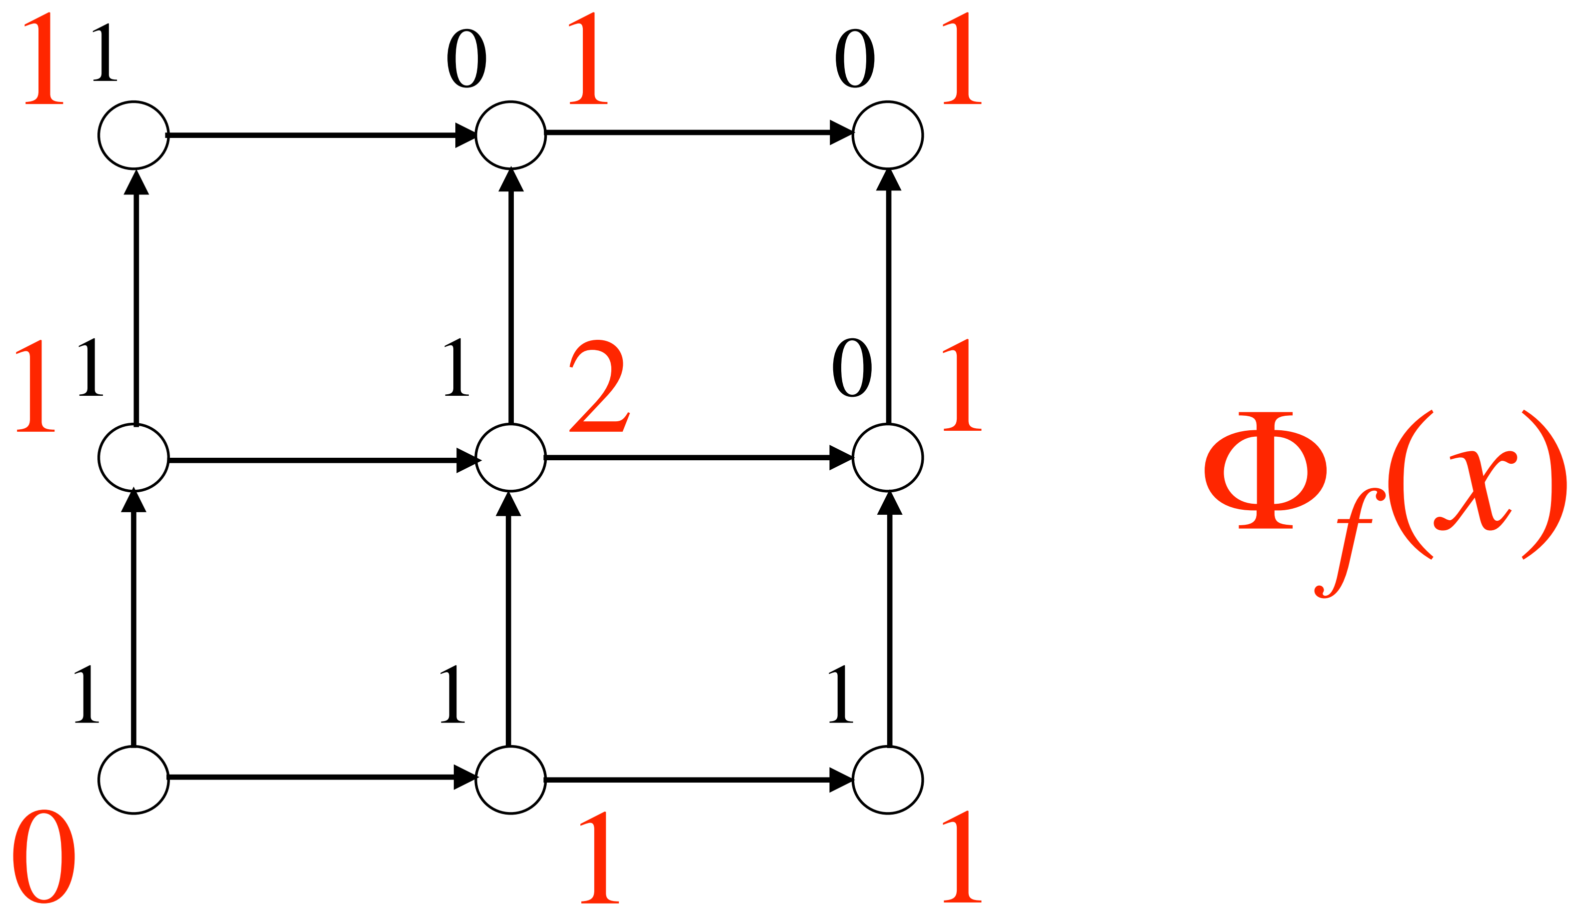

- When $n = 2,\ \Phi_f(x) = I_f^-(x)$

$\Longrightarrow$ Generalizes the directed Talagrand inequality by KMS

How does this inequality help us analyze monotonicity testers?

Good Subgraphs (let's assume $\varepsilon(f) = \Omega(1)$)

Theorem (BCS STOC 23): For any $f\colon [n]^d \to \{0,1\}$,

$$\mathbb{E}_x[\Phi_f(x)^{1/2}] = \Omega\left(\frac{\varepsilon(f)}{\log n}\right)$$

$\Longrightarrow$

Good subgraph lemma (KMS 15, informal):

For some Δ, there is bipartite subgraph of decreasing edges $G(U, V, E)$ with max degree Δ and $|E| \geq \widetilde{\Omega}(\sqrt{\Delta} \cdot n^d)$

$f = 1$ 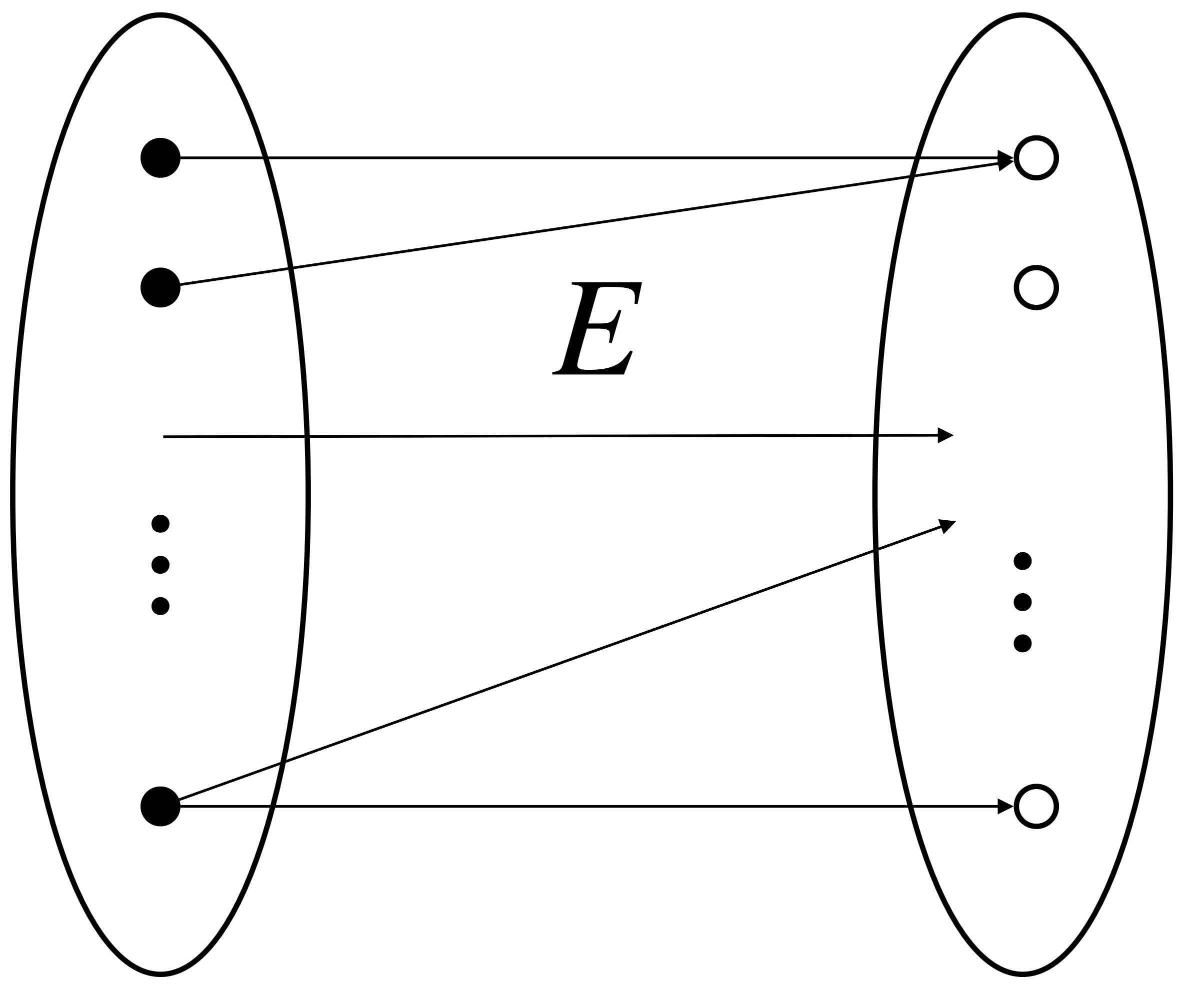 $f = 0$

$\Delta = d \quad \Longrightarrow \quad |E| \geq \Omega(\sqrt{d} \cdot n^d)$

$\Delta = 1 \quad \Longrightarrow \quad |E| \geq \Omega(n^d)$ and E is a **matching**

Plan for the Rest of the Talk

Matching assumption:
$f\colon [n]^d \to \{0,1\}$

There is a **matching** E of $\Omega(n^d)$ **decreasing** edges

$\Longrightarrow$ Tester analysis

1) An $O(n\sqrt{d})$ tester under matching assumption
   - These techniques are due to KMS 15

2) Sketch for $O(\log n\sqrt{d})$ tester under matching assumption
   - We may assume $n = \text{poly}(d)$ and so $O(\log n\sqrt{d}) = \widetilde{O}(\sqrt{d})$

(BCS SODA 20, Harms-Yoshida ICALP 22)

An $O(n\sqrt{d})$ Query Tester

Assumption: There is a **matching** E of $\Omega(n^d)$ decreasing edges

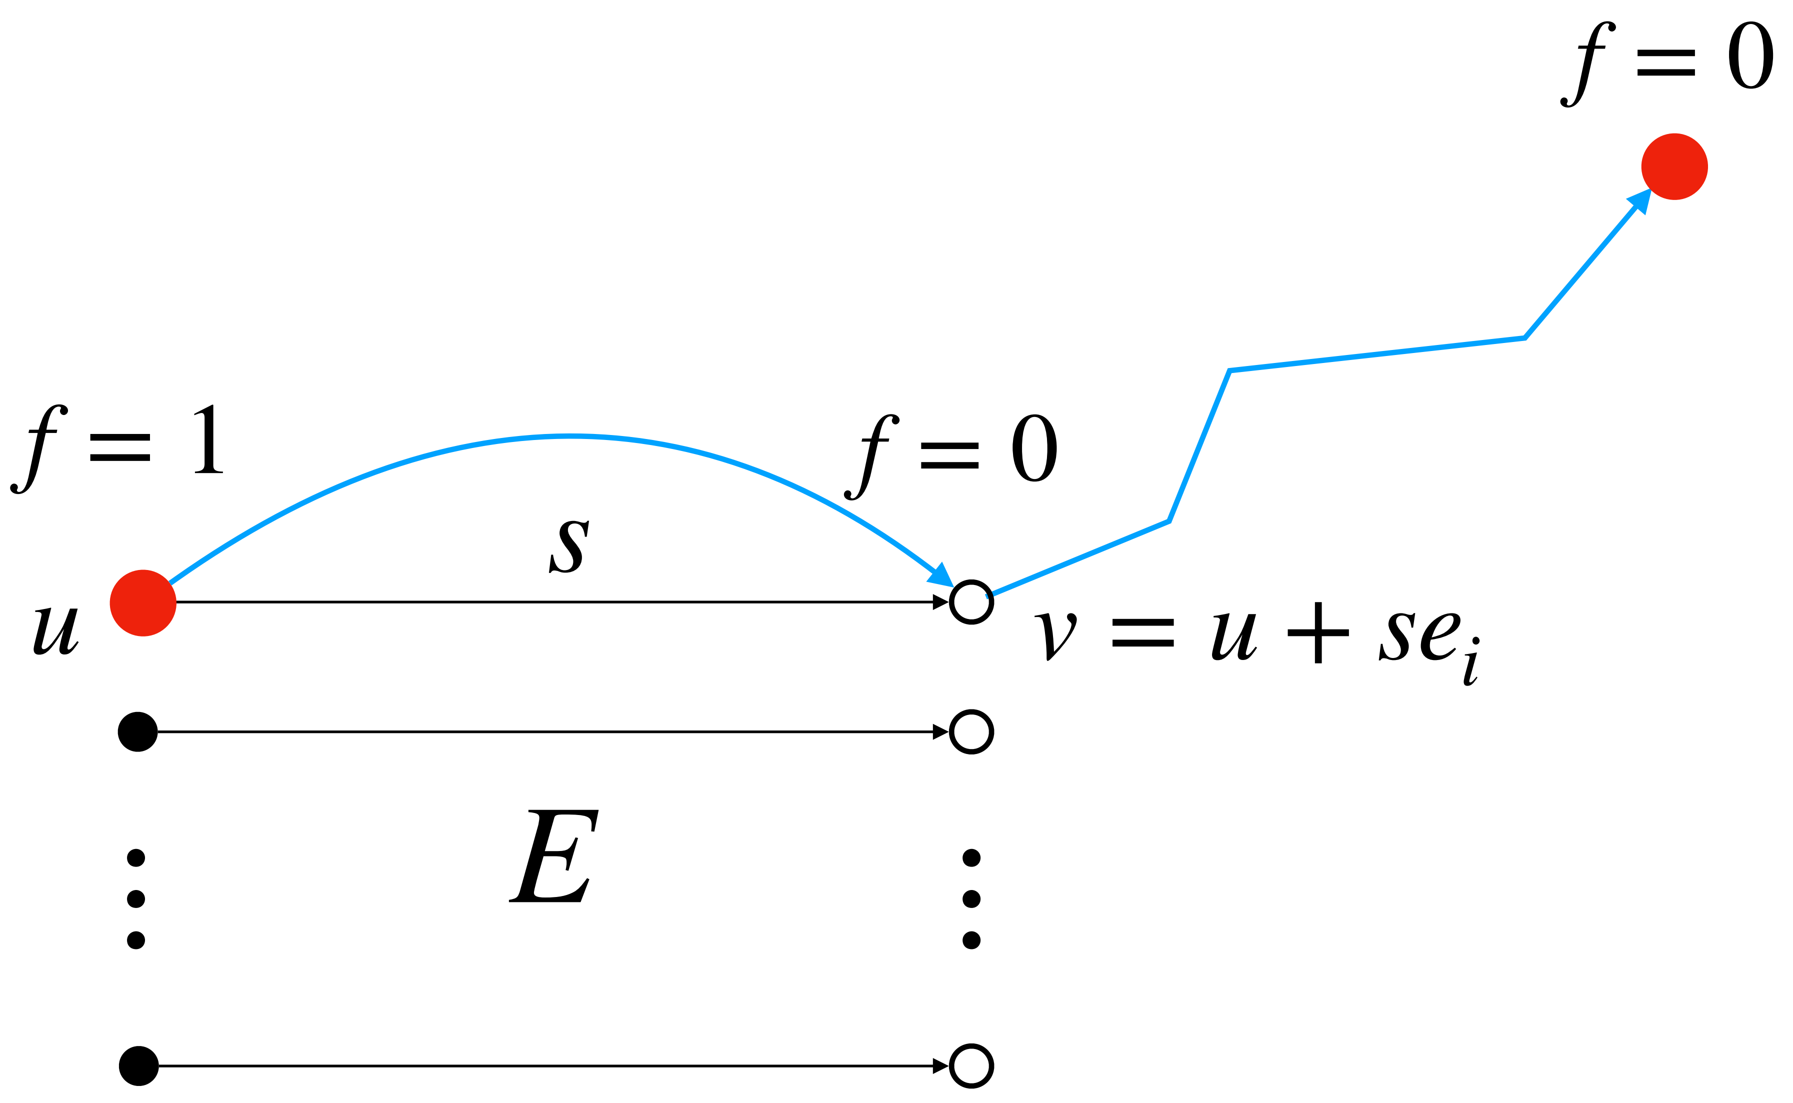

Walk distribution 1:

- Sample x u.a.r. and a set $T \subseteq [d]$ of $\tau \approx \sqrt{d}$ random coordinates
  - For $i \in T$: set $y_i \in [x_i, n]$ u.a.r.
  - For $i \notin T$: set $y_i = x_i$

Def: $y \in [n]^d$ is **persistent** if a random walk from y of length $\tau - 1$ leads to z with $f(z) = f(y)$ with probability ≥ 0.9

Lemma (KMS, informal):
If # decreasing edges $< n\sqrt{d} \cdot n^d$, then # non-persistent points is $o(n^d)$

$\implies$ All endpoints of E are **persistent**

$$\mathbb{P}[f(x) = 1 \wedge f(y) = 0]$$

$$\geq \sum_{(u,u+se_i)\in E} \underbrace{\mathbb{P}[x = u]}_{n^{-d}} \cdot \underbrace{\mathbb{P}[T \ni i]}_{\Omega(d^{-1/2})} \cdot \underbrace{\mathbb{P}[y_i - x_i = s]}_{\Omega(n^{-1})} \cdot \underbrace{\mathbb{P}[f(y) = 0]}_{0.9}$$

$$\geq \Omega(n^{-1}d^{-1/2})$$

End of the story for constant n…

What about **larger n?**

Internal Points

Assumption: There is a **matching** E of $\Omega(n^d)$ decreasing edges

- Let $I(u,v)$ = interval of points between u and v

… let's assume $I(u,v)$ is at least half 0's for all $(u,v) \in E$

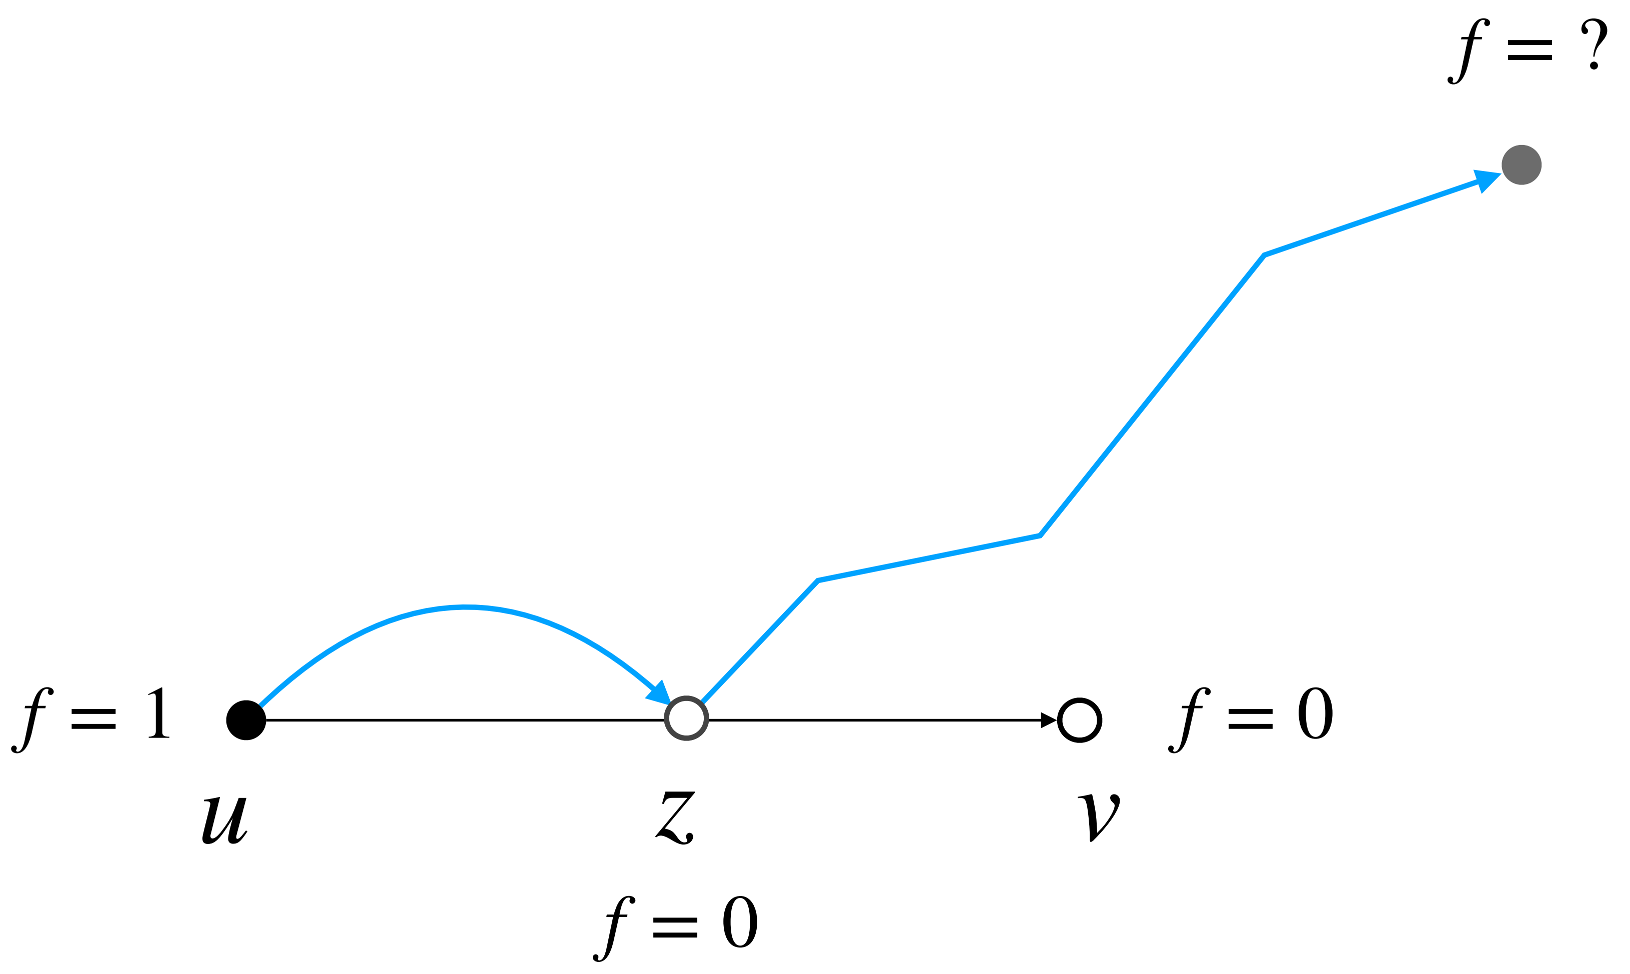

Walk distribution 2:

- Sample x u.a.r. and $T \subseteq [d]$ of $\tau \approx \sqrt{d}$ random coordinates
  - **For $i \in T$, choose $p \in [\log n]$ u.a.r**
    - **set $y_i \in [x_i, x_i + 2^p]$ u.a.r**

But can we argue a random point in $I(u,v)$ is **persistent?**

Short answer: **No.**

$$\left| \bigcup_{(u,v)\in E} I(u,v) \setminus \{u,v\} \right|$$ may be very small

Walk 2 analysis: condition on passing through u.a.r $z \in I(u,v)$

$$\approx \log^{-1} n$$

Red Edges and the Shifted Path Test

Assumption: There is a **matching** E of $\Omega(n^d)$ decreasing edges

Def: Call z *mostly-zero-below* (**mzb**) if a $(\tau - 1)$-length **downward** random walk from z ends at a 0 with prob. ≥ 0.9

Def: Call an edge (u, v) **red** if a $(\tau - 1)$-length **upward** random walk from a u.a.r. $w \in I(u, v)$ ends at z which is **mzb** with prob. ≥ 0.01

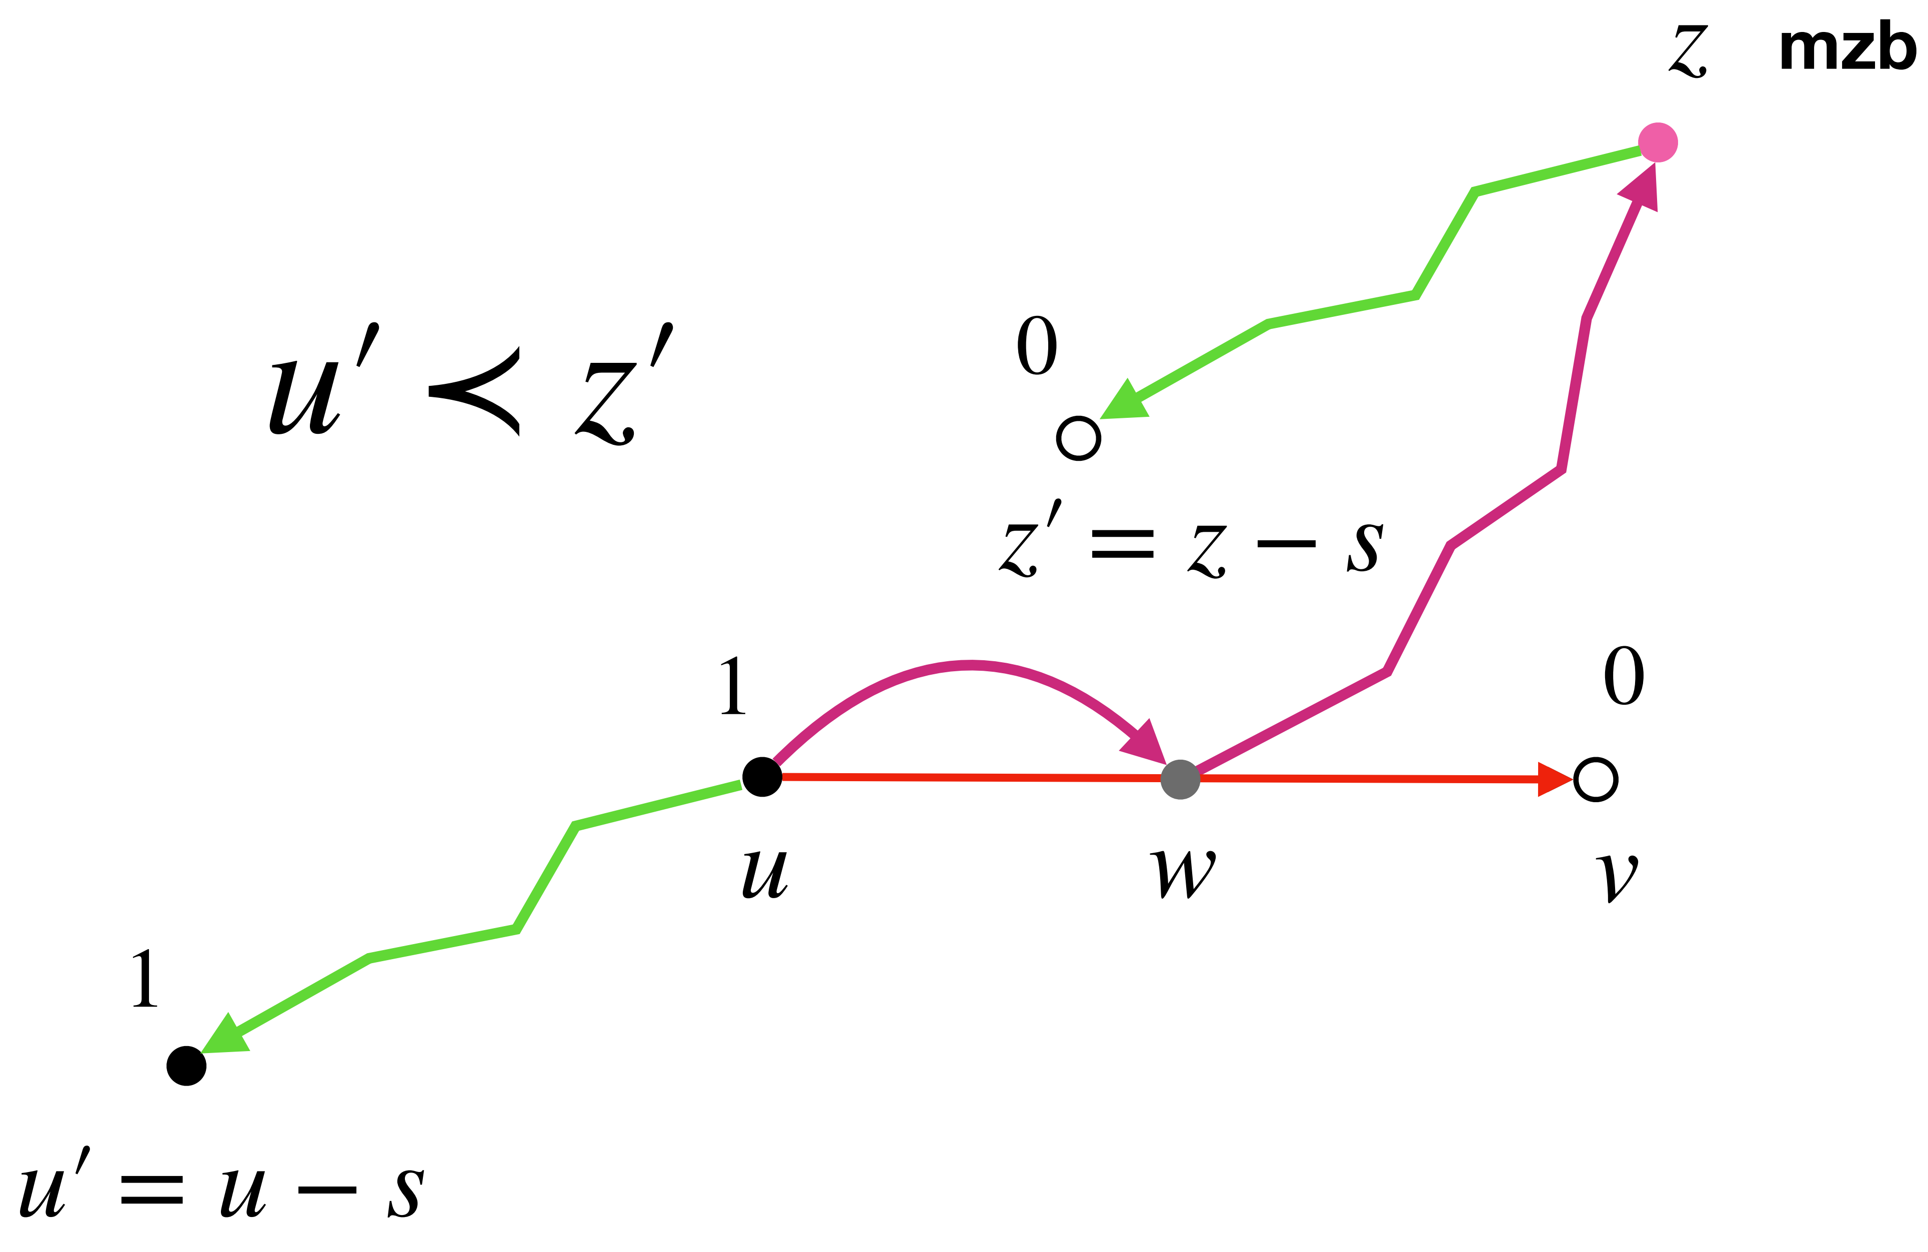

Recall: All endpoints of E are **persistent**

E is mostly **red** $\implies$ **Upward walk** + **downward shift** finds a violation with prob. $\Omega(d^{-1/2} \log^{-1} n)$

What if E is mostly **not red**?

Non-red Edges

Assumption: There is a **matching** E of $\Omega(n^d)$ decreasing edges

Def: Call an edge (u, v) **red** if a $(\tau - 1)$-length upward random walk from a u.a.r. $w \in I(u, v)$ ends at z which is **mzb** with prob. ≥ 0.01

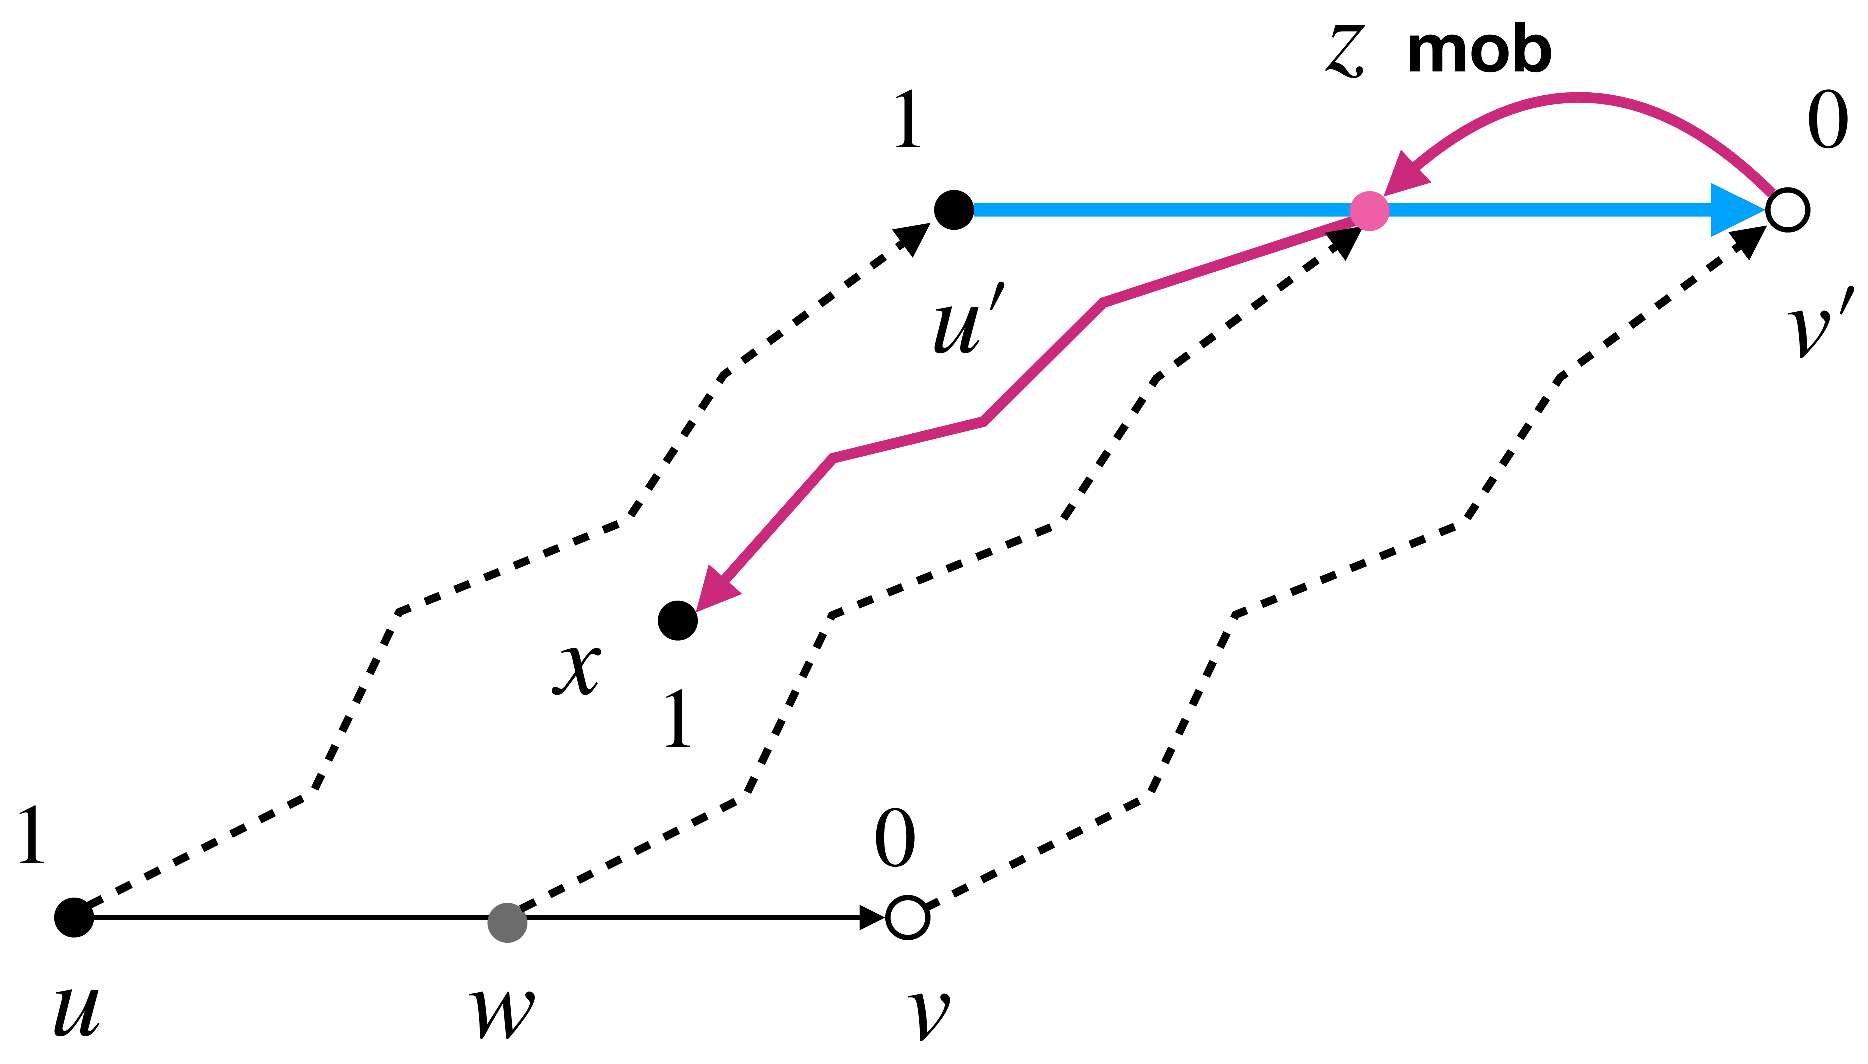

$\Longrightarrow$ (u, v) **non-red** $\Longrightarrow$ z mostly-**one**-below with prob. ≥ 0.99

Recall: u, v are **persistent**

Consider (u', v') a random translation of (u, v)

$\Longrightarrow$ With high prob. $f(u') = 1, f(v') = 0$ and most of $I(u', v')$ is **mostly-one-below**

$\Longrightarrow$ A **downward random walk** from v' discovers a violation with probability $\Omega(d^{-1/2} \log^{-1} n)$

Def: Call an edge (u', v') **blue** if a constant fraction of $I(u', v')$ is **mostly-one-below**

Red-blue Win-win Argument

Assumption: There is a **matching** E of $\Omega(n^d)$ decreasing edges

Def: Call an edge (u, v) **red** if a $(\tau - 1)$-length upward random walk from a u.a.r. $w \in I(u, v)$ ends at z which is **mzb** with prob. ≥ 0.01

Def: Call an edge (u', v') **blue** if a constant fraction of $I(u', v')$ is **mob**

Case 1: E is mostly **red**

$\Longrightarrow$ **Upward walk** + **downward shift** finds a violation with prob. $\Omega(d^{-1/2} \log^{-1} n)$

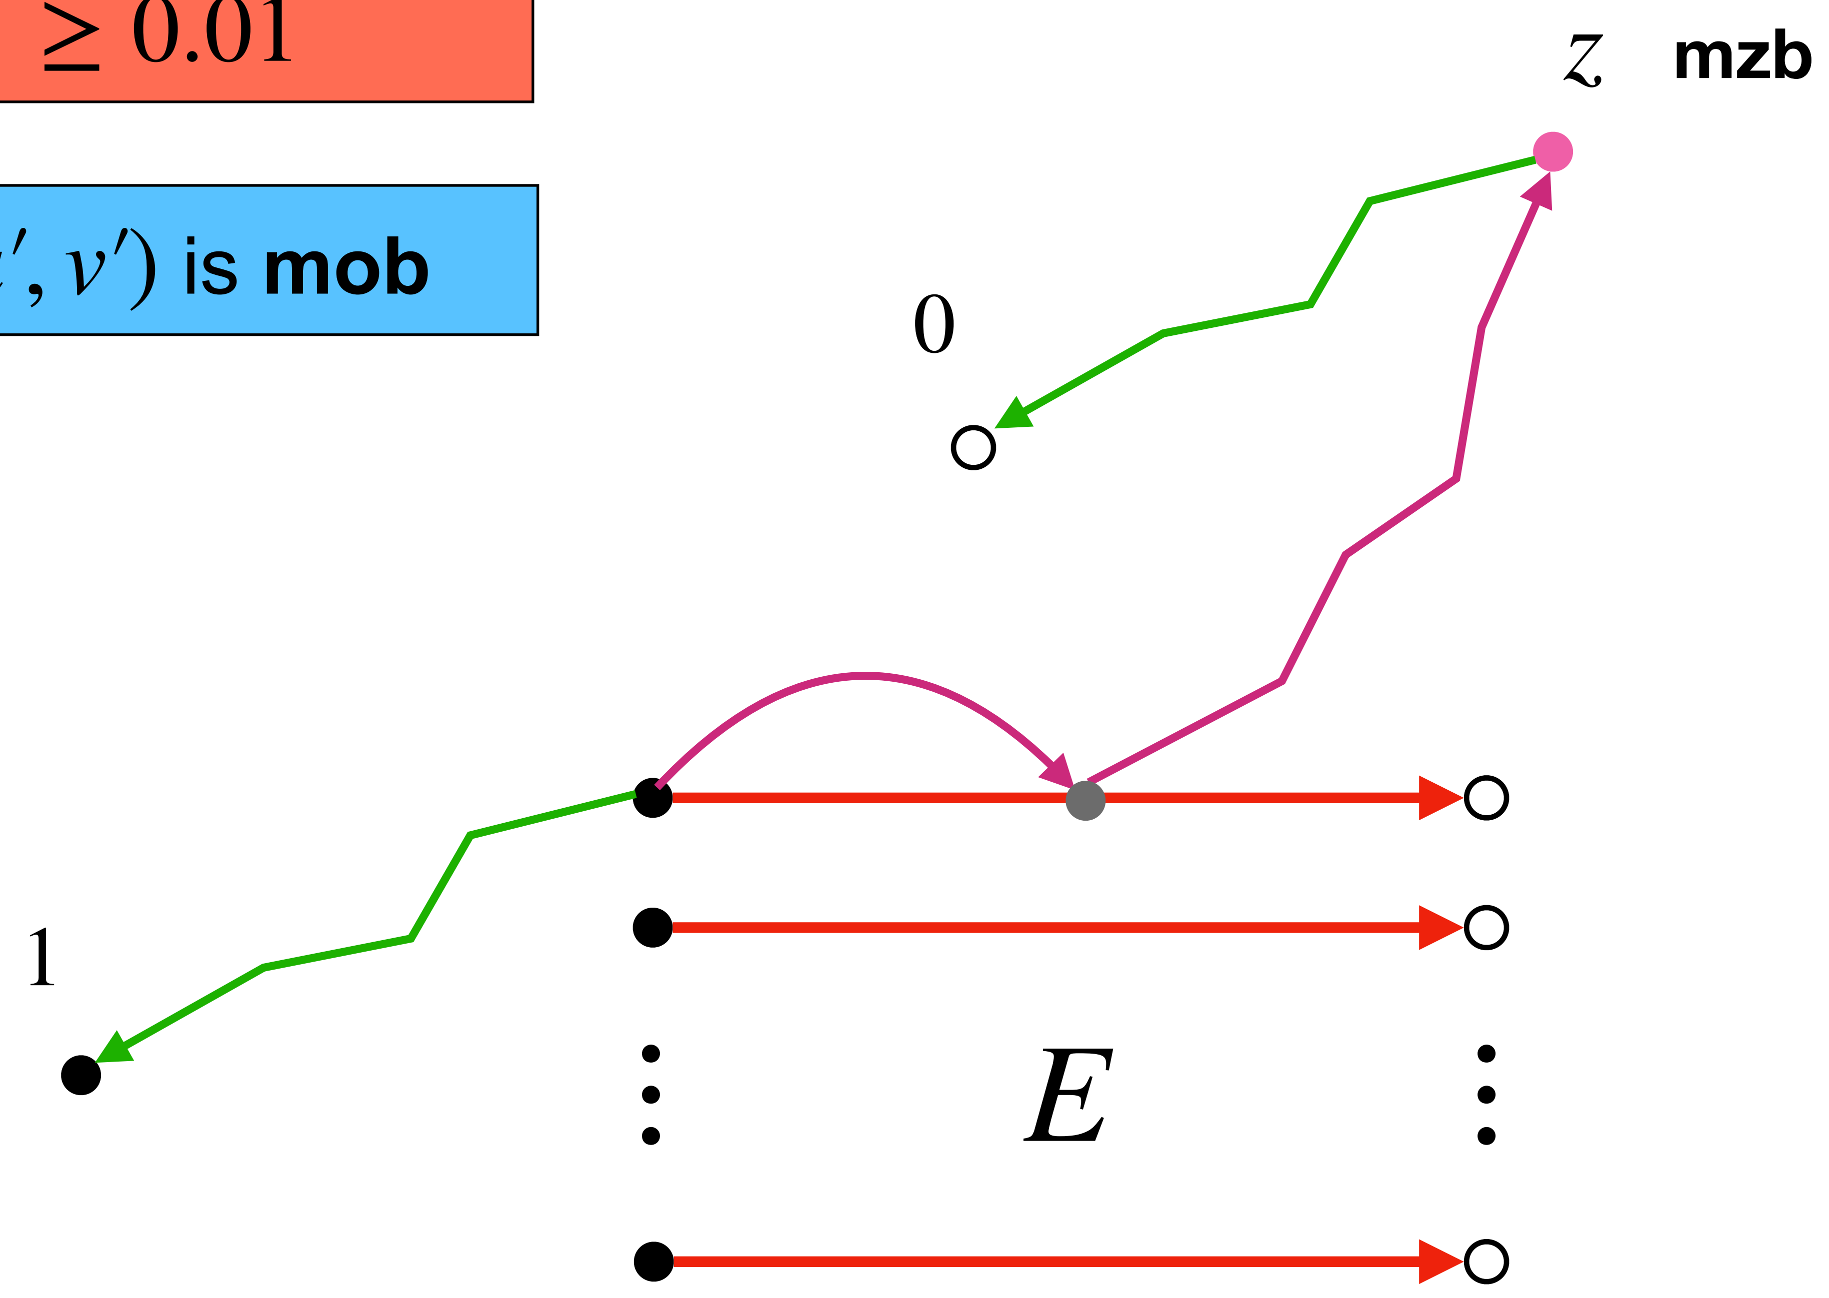

Red-blue Win-win Argument

Assumption: There is a **matching** E of $\Omega(n^d)$ decreasing edges

Def: Call an edge (u, v) **red** if a $(\tau - 1)$-length upward random walk from a u.a.r. $w \in I(u, v)$ ends at z which is **mzb** with prob. ≥ 0.01

Def: Call an edge (u', v') **blue** if a constant fraction of $I(u', v')$ is **mob**

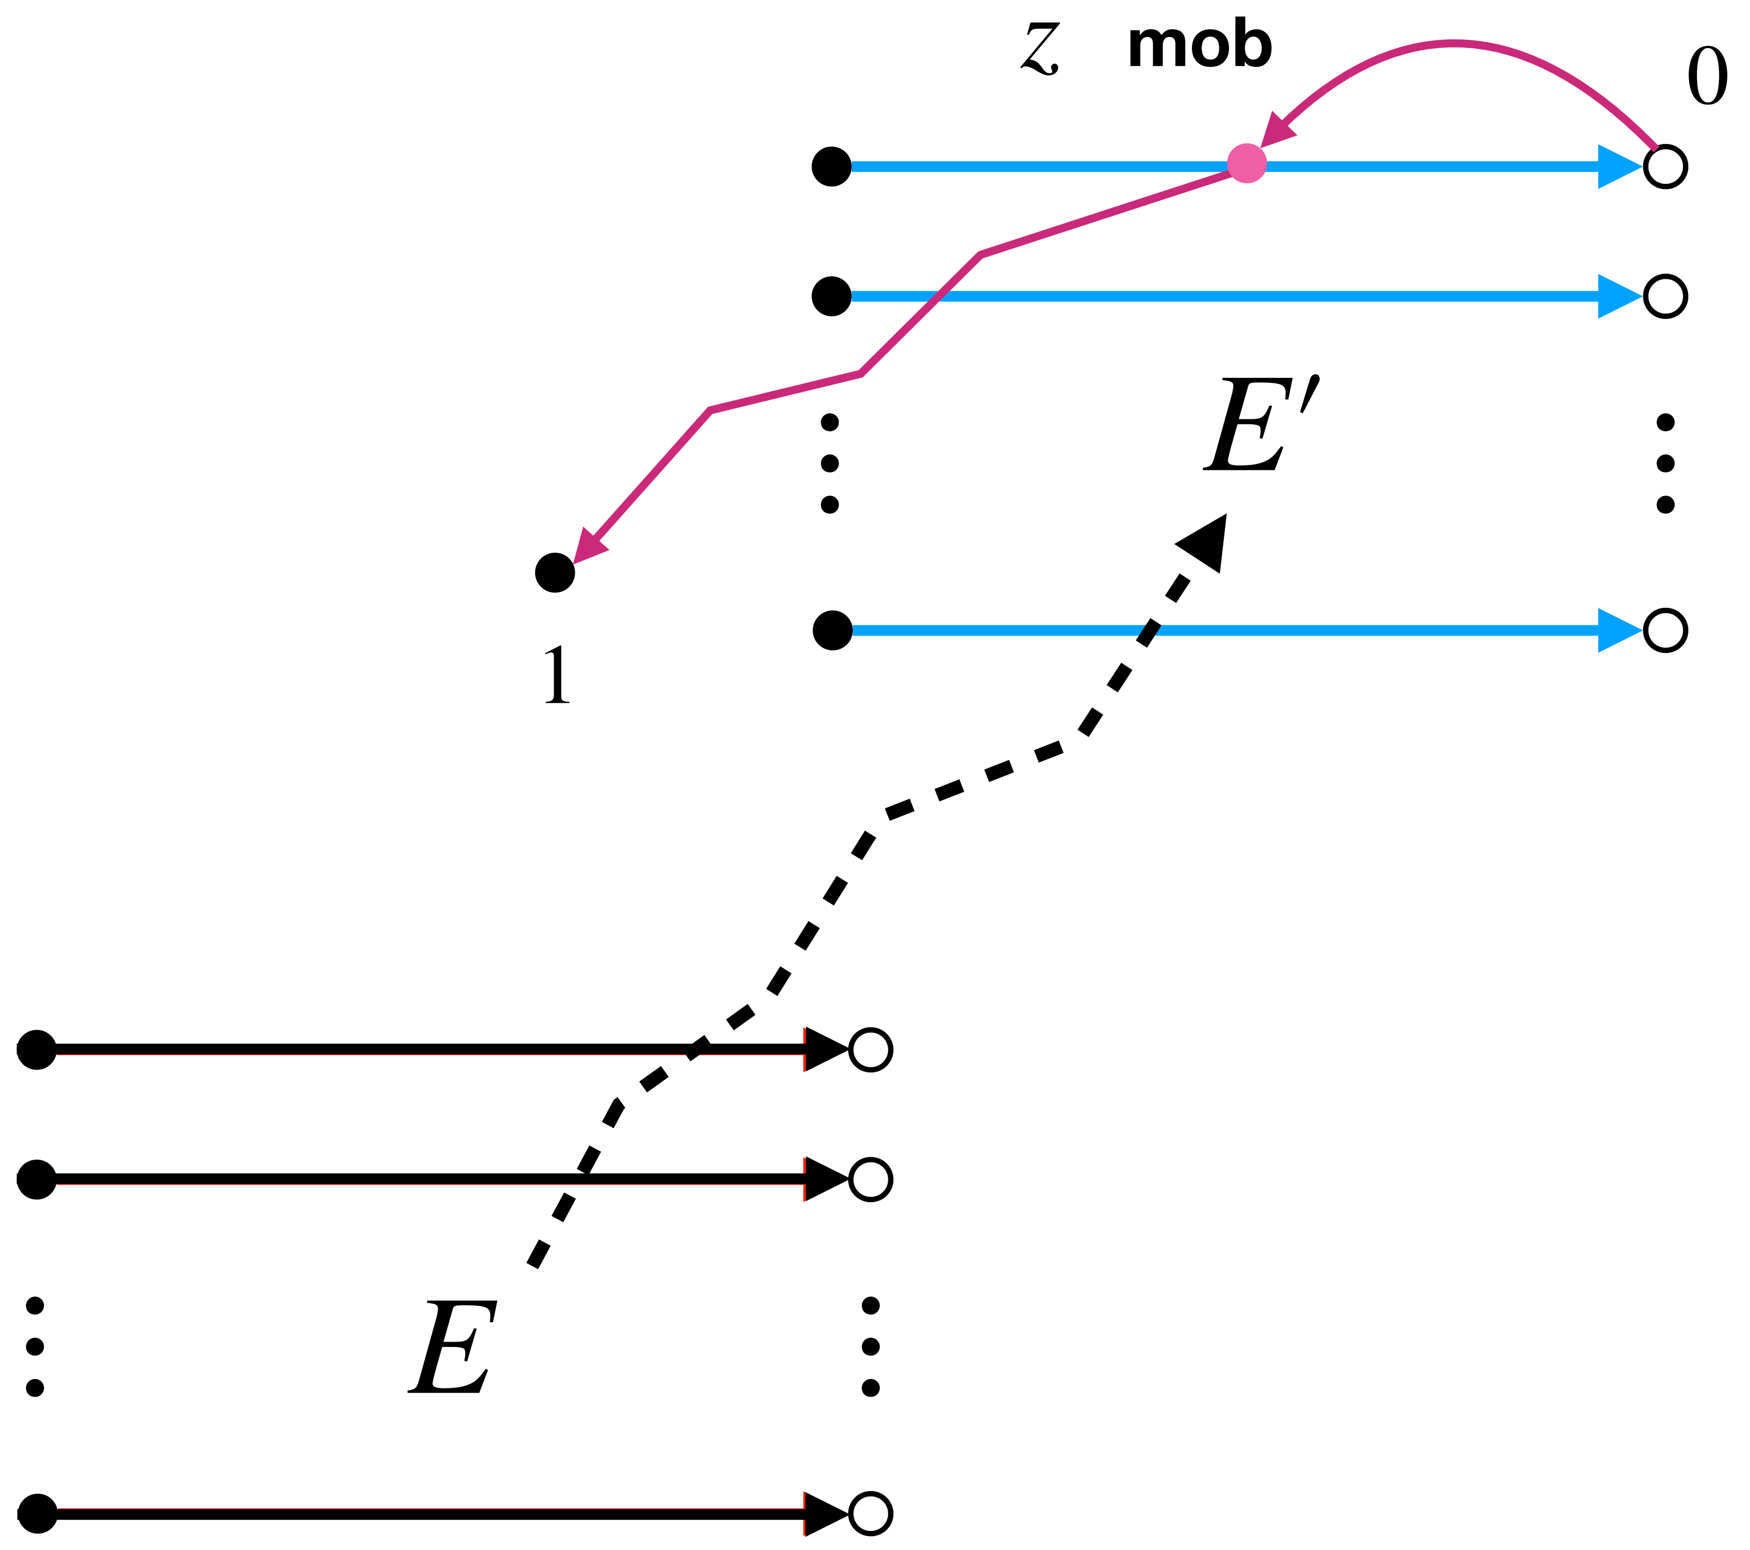

Case 1: E is mostly **red**

$\Longrightarrow$ **Upward walk** + **downward shift** finds a violation with prob. $\Omega(d^{-1/2} \log^{-1} n)$

Case 2: E is mostly **non-red**

$\Longrightarrow$ Flow argument: there exists another matching E' which is mostly **blue**

$\Longrightarrow$ **Downward walk** finds a violation with prob. $\Omega(d^{-1/2} \log^{-1} n)$ $\square$

Summary

BCS FOCS 23

There is a $\varepsilon^{-2} d^{1/2+o(1)}$ query monotonicity tester for functions $f\colon [n]^d \to \{0,1\}$

$\approx$ Resolves the nonadaptive monotonicity testing problem

BCS FOCS 23

Translating good subgraphs

+

BCS STOC 23

Directed Talagrand inequality for hypergrids

+

BCS SODA 20

Domain reduction for hypergrids

Open question:
Adaptive monotonicity testing?

Chen-Waingarten-Xie 17
$\widetilde{\Omega}(d^{1/3})$

?

Khot-Minzer-Safra 15
$\widetilde{O}(d^{1/2})$