An Overview of Quantified Derandomization

Roei Tell, Weizmann Institute of Science

Complexity Theory @ Oxford, July 2018

מכון ויצמן למדע
WEIZMANN INSTITUTE OF SCIENCE

Classical derandomization (CAPP)

› the standard derandomization problem

Given a circuit **$C \in \mathcal{C}$** over **n bits**, deterministically distinguish between the cases:

- **C accepts all but at most $2^n/3$ of its inputs**
- **C rejects all but at most $2^n/3$ of its inputs**

Classical derandomization (CAPP)

› lower bounds ⇒ derandomization

- When **$\mathcal{C}$=P/poly** equivalent to **prBPP=prP**
- Implied by **average-case lower bounds** for $\mathcal{C}$
  - hardness-randomness [Yao'82, BM'84, NW'94]
  - hardness amplification (e.g., [IW'99])
  - gives blackbox derandomization (i.e., a PRG)

Classical derandomization (CAPP)

- state of the art

- **P/poly:** ?
- **TC^0, NC^1:** ?
- **ACC^0:** sat in time 2^{n-n^ε} [Wil'11]
- **AC^0:** quasipoly time [AW'85, Bra'11, TX'12, Tal'17]
- **CNFs:** time $n^{\tilde{O}(\log\log n)}$ [LV'96, Baz'07, DETT'10, GMR'12]

Classical derandomization (CAPP)

- derandomization ⇒ lower bounds

- **Blackbox derand implies lower bounds**
  - output-set of PRG/HSG is “hard” function
- **Whitebox derand implies** (weaker) **lower bounds**
  - indirect arguments [IW’98, IKW’02, KI’04, Wil’11, BV’14, MW’18]
  - “hard” function in E^{NP}, NEXP, NQP, NTIME[$n^{\log^*(n)}$]
- **Faster derand ⇒ better lower bounds**
  - circuit size, explicitness of “hard” function

Quantified derandomization

- a relaxed derandomization problem [GW'14]

Given a circuit **C ∈ $\mathcal{C}$** over **n bits**, deterministically distinguish between the cases:

- **C accepts all but at most B(n) of its inputs**
- **C rejects all but at most B(n) of its inputs**

⇒ in the classical problem $B(n)=2^n/3$; we think of $B(n) = o(2^n)$

Quantified derandomization

› conflicting intuitions

- **In “complexity 101” they said that ⅓ is arbitrary!**
  - error-reduction: just how low can it take us?
- **For $B(n)=0$, I know how to solve the problem!**
  - detecting extremely small bias is easy
- **So is it easy or hard to detect extremely small bias?**

Quantified derandomization

- for a fixed circuit class $\mathcal{C}$

"Easy" vs "hard" values for B(n)

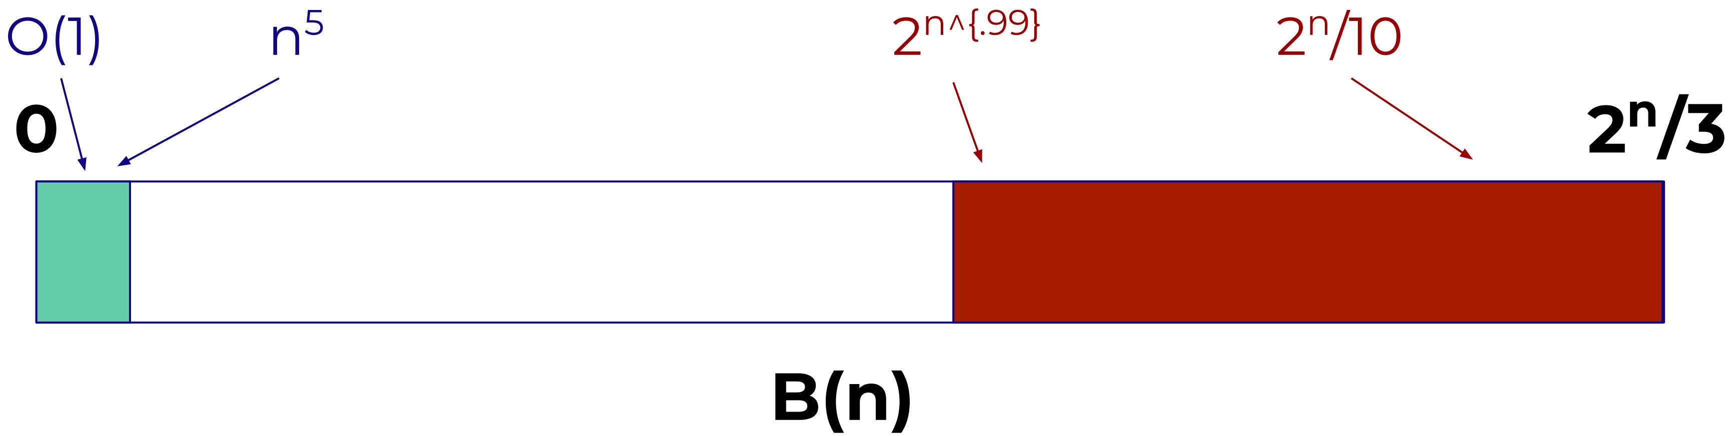

Quantified derandomization

- for a fixed circuit class $\mathcal{C}$

<u>Goal 1:</u> Understand! Get **tight results**

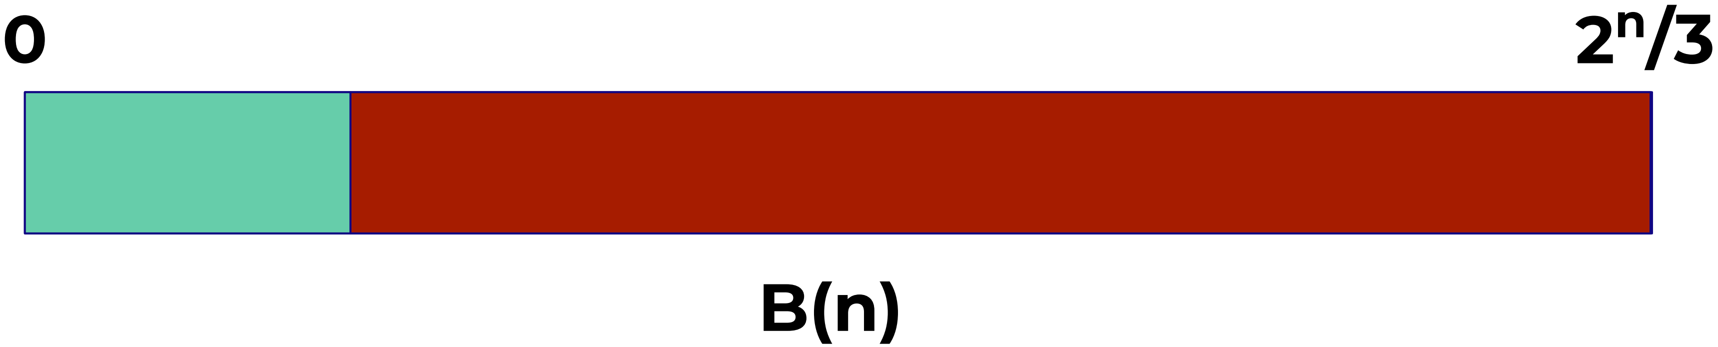

Quantified derandomization

- for a fixed circuit class $\mathcal{C}$

<u>Goal 1:</u> Understand! Get **tight results**

<u>Goal 2:</u> Make green and red cross ⇒ **standard derand**

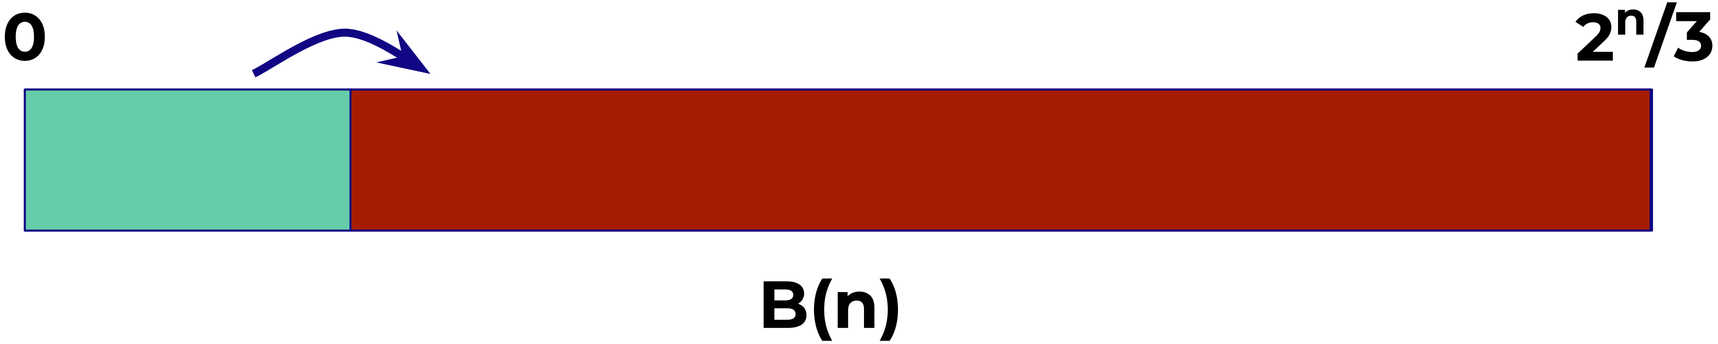

Quantified derandomization

› derandomization ⇒ lower bounds

- **Blackbox derand implies lower bounds** [1]
  - output-set of PRG/HSG still a “hard” function
- **Whitebox derand doesn’t** (necessary) **imply LBs**
  - implies LBs indirectly, via standard derandomization
- **No** (known) **speed vs. size trade-off**

1 assuming non-triviality: #exceptional inputs ≥ #outputs of HSG/PRG

Polynomials that vanish rarely

- Consider **degree-d polys $F^n \to F$** for finite field $F=F_q$
- Hitting-set for **all polys** has size **≥ (n+d choose d)**
- Is there a hitting-set for polys that **vanish on at most b(n)** of inputs of **size o((n+d choose d))**?

1 question interesting even for non-explicit hitting-sets!

Some known results

research directions that have been active

Overview of known results

- **Constant-depth circuits:**
  - AC^0 [GW'14, GVW'15, CL'16, T'17]
  - $AC^0[\oplus]$ [GW'14, T'17]
  - TC^0, LTF/PTF ckts [T'18, KL'18]
- **Polys that vanish rarely** [GW'14, T'17, in progress]
- **Proof systems** [GW'14]

AC^0: touching the threshold

› circuits of constant depth d

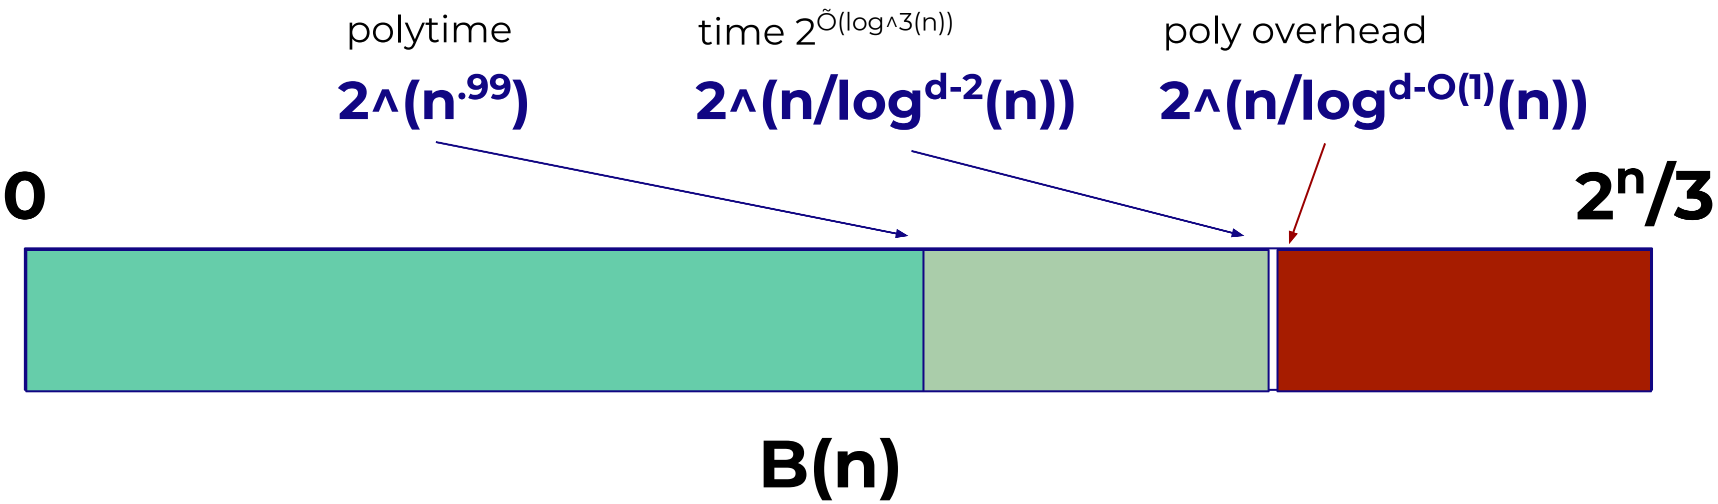

1 see [GW'14, GVW'15, CL'16, T'17]

TC^0, LTF and PTF circuits

- circuits of constant depth d

#wires	**lower bounds**	**quant derand with $B(n) \approx 2^{n^{.99}}$**
poly(n)	⇧	⇧
$n^{1+O(1/d)}$	bounds against specific funcs can be "magnified" [AK'10]	quant derand would imply standard derand of all TC^0 [T'18]
$n^{1+\exp(-d)}$	unconditional bds: parity, gen Andreev [IPS'97, CSS'16]	unconditional quant derand for LTF, PTF ckts [T'18,KL'18]

1 see [T'18, KL'18]

Polys that vanish rarely

› polys $F^n \to F$ of any degree d=d(n)

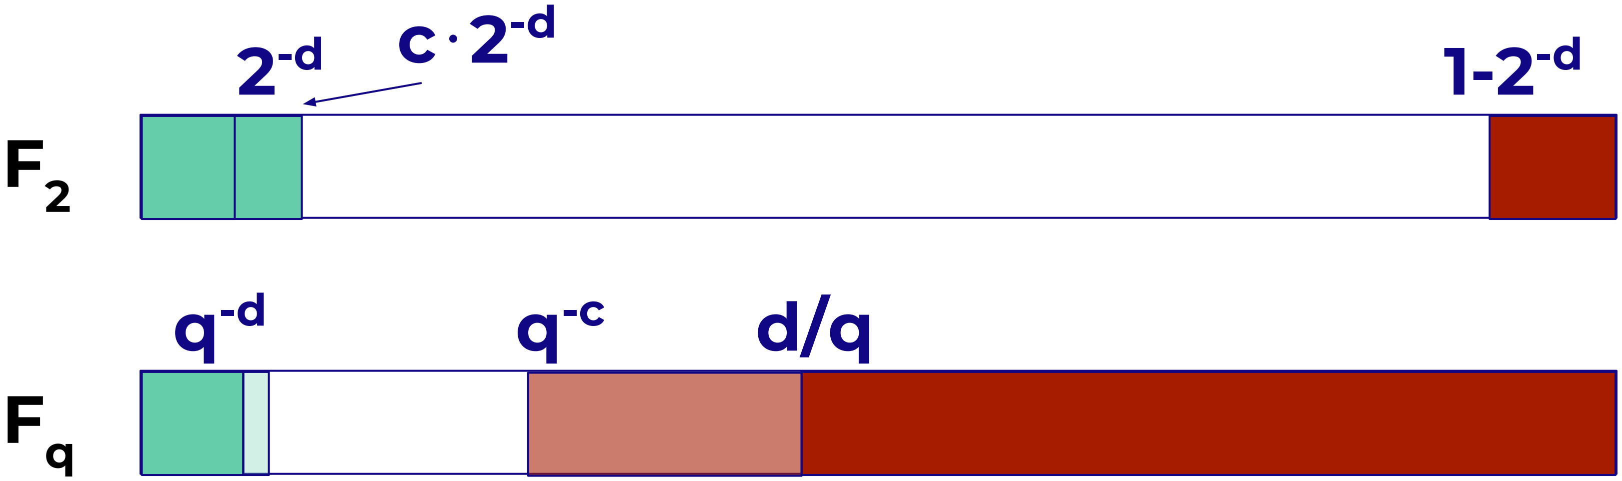

1 see [GW'14, T'17]; work in progress

Known techniques and their limitations

Deterministic restrictions

› high-level strategy suggested by [GW'14]

Idea: Given $C:\{0,1\}^n \to \{0,1\}$, find **simple function** that **approximates C** in large subset $S \subseteq \{0,1\}^n$, $|S| \gg B(n)$

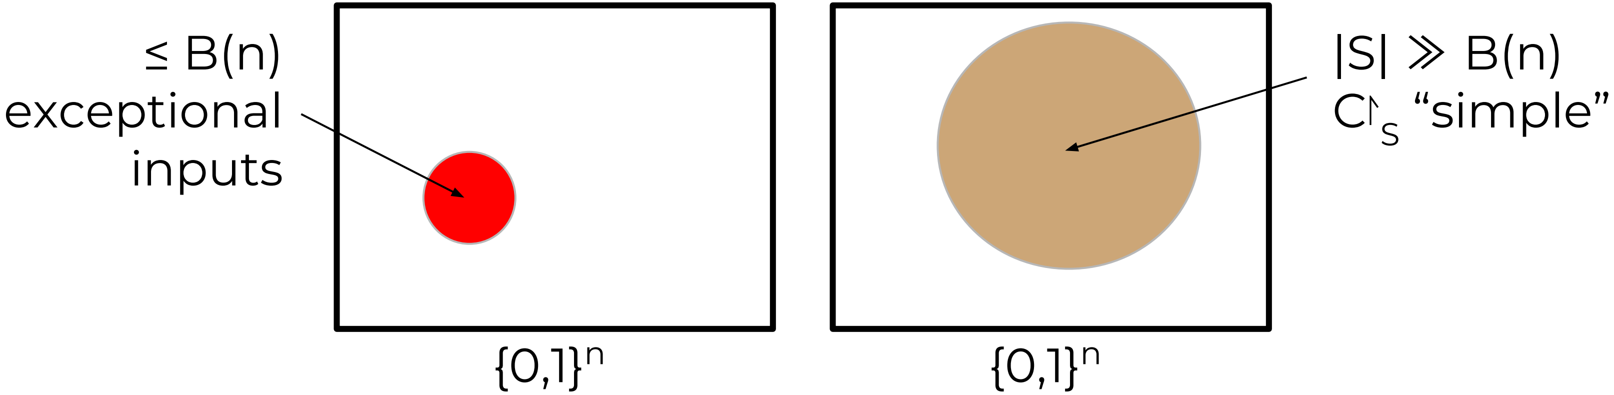

Deterministic restrictions

› comments

- **Obs: Method is "complete"**
- Subset **S not necessarily a subcube**
  - but we need to approx the bias of the simple func in S
- Can use **whitebox access** to circuit
- "Full derandomization" of restriction procedures
  - previous applications required only partial derand [AW'85]

Polys that vanish rarely

- several ad-hoc techniques
- **Structural results:**
  - biased polys approximated by **low-degree polys**
  - biased polys **constant on almost all large subspaces**
- **Biased ckts** have probabilistic representation

  as **biased polys** ⇒ approx by **low-degree polys**

Error-reduction

Input

$C:\{0,1\}^m \to \{0,1\}$

- depth d, size s
- at most $2^m/3$ bad inputs

Output

$C':\{0,1\}^n \to \{0,1\}$

- blow-up in d, s, n=n(m)
- preserves majority output
- at most **B(n) bad inputs**

Error-reduction

- using a seeded extractor / averaging sampler

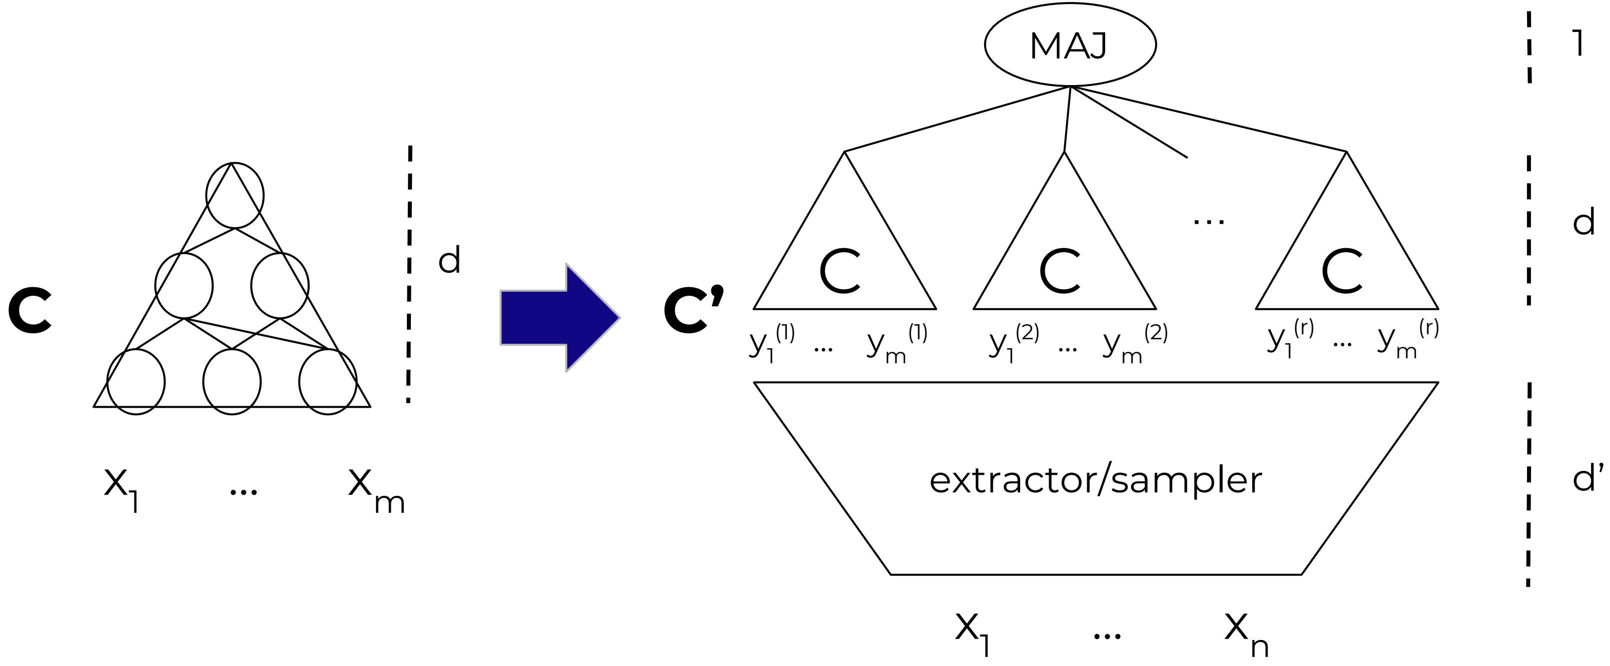

Error-reduction

› comments

› Extractors in "weak models" **barely studied before**

  › this led to fruitful study of extractors in AC^0, TC^0, polys

› Extractors are **an "overkill"**

  › we only need to sample one event, induced by circuit $C \in \mathcal{C}$

  › weaker notions: extractor for $\mathcal{C}$-events, whitebox extractor

1 AC^0-extractors for AC^0-tests cannot be significantly more efficient than AC^0-extractors for all tests

Limitation of blackbox techniques

Limitation of blackbox techniques

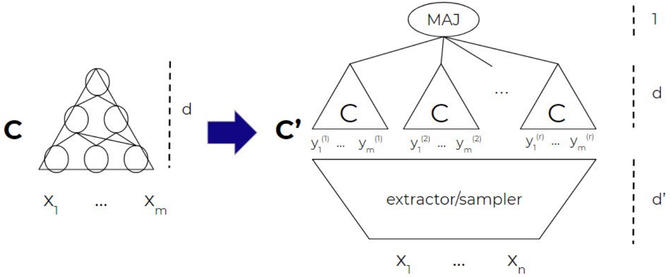

Step 1: Error-reduction

› extractor for $\mathcal{C}$-events

› doesn't depend on specific C

Idea: Given $C:\{0,1\}^n \to \{0,1\}$, find **simple function** that **approximates C** in large subset $S \subseteq \{0,1\}^n$, $|S| \gg B(n)$

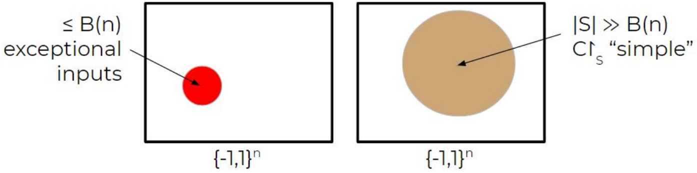

Step 2: Restrictions

› distribution over restrictions

› doesn't depend on specific C

Limitation of blackbox techniques

- **Thm:** For any class $\mathcal{C} \supseteq$ {polysize DNFs}, if there are
  1. $\mathcal{C}$-computable extractor with **B'(n) bad inputs** for error $\Omega(1)$
  2. distribution over sets of **size B(n)** that **simplifies every C ∈ $\mathcal{C}$** to a constant, wp > ½

  Then, necessarily **B(n) < B'(n).**

⇒ Naive comb of the two techs **cannot suffice for standard derand**

1 restriction procedures for "small $AC^0[\oplus]$", LTF ckts, PTF ckts already whitebox

Open problems are everywhere

here’s a carefully-trimmed list

Where next?

- few suggested directions

- **Non-deterministic algorithm** for quantified derand
  - suffice for "derand ⇒ lower bounds" [Wil'11]
  - can use collapse hypothesis & some advice [FS'16,MW'17]
- **Whitebox samplers** (sampler for specific circuit)
- HSGs for **polys $\mathbf{F}^n_q \to \mathbf{F}_q$ that vanish rarely**

Thank you!

⇒ relaxed circuit-analysis task
⇒ limitations on blackbox techniques
⇒ “interesting problem! perhaps relevant to stuff I like?”