ROBUST REVENUE MAXIMIZATION UNDER MINIMAL STATISTICAL INFORMATION

Alexandros Tsigonias-Dimitriadis

Operations Research Group & RTG AdONE
TU Munich

(Joint work with Yiannis Giannakopoulos and Diogo Poças)

WINE 2020 - Virtual Talk

Technische
Universität
München

Myerson's optimal auction

Myerson's optimal auction

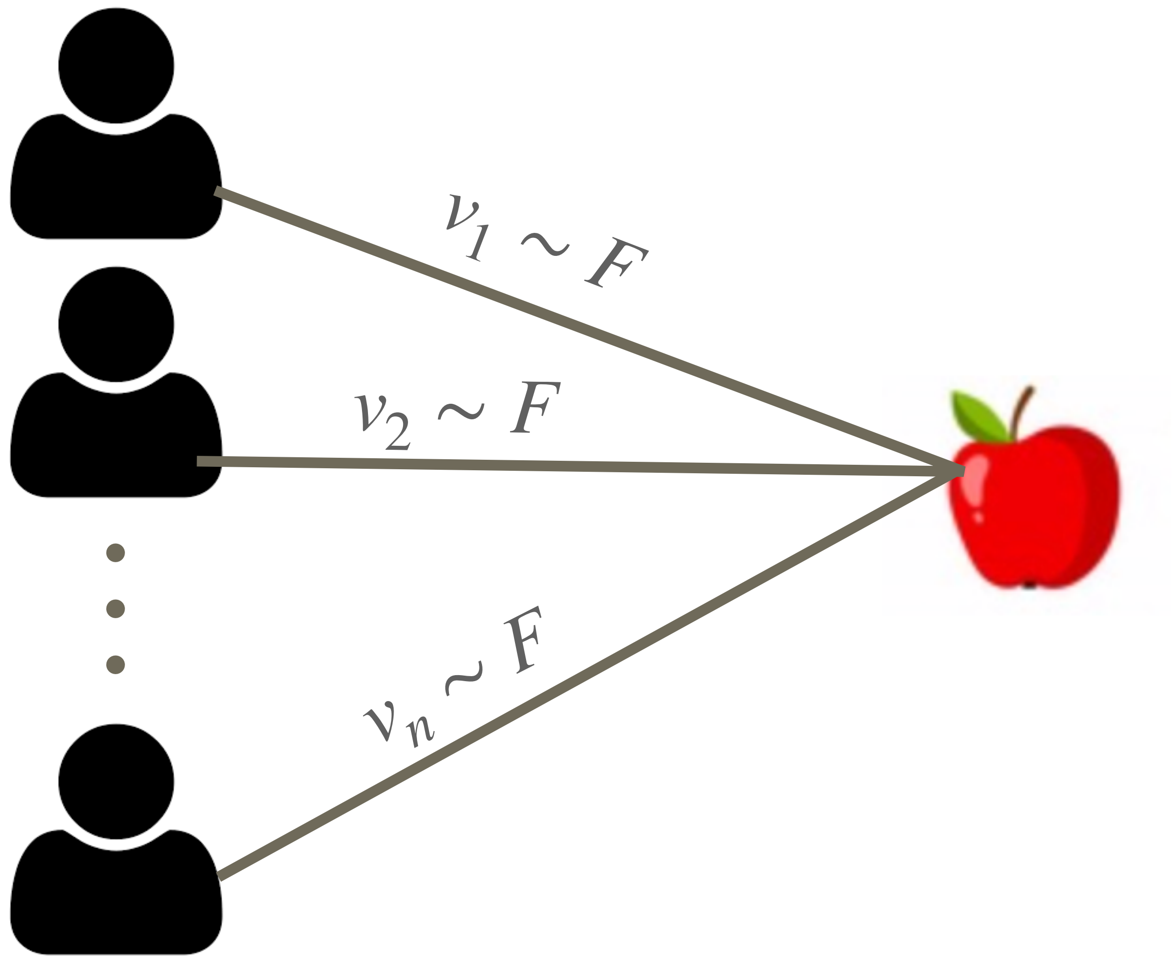

Myerson's optimal auction

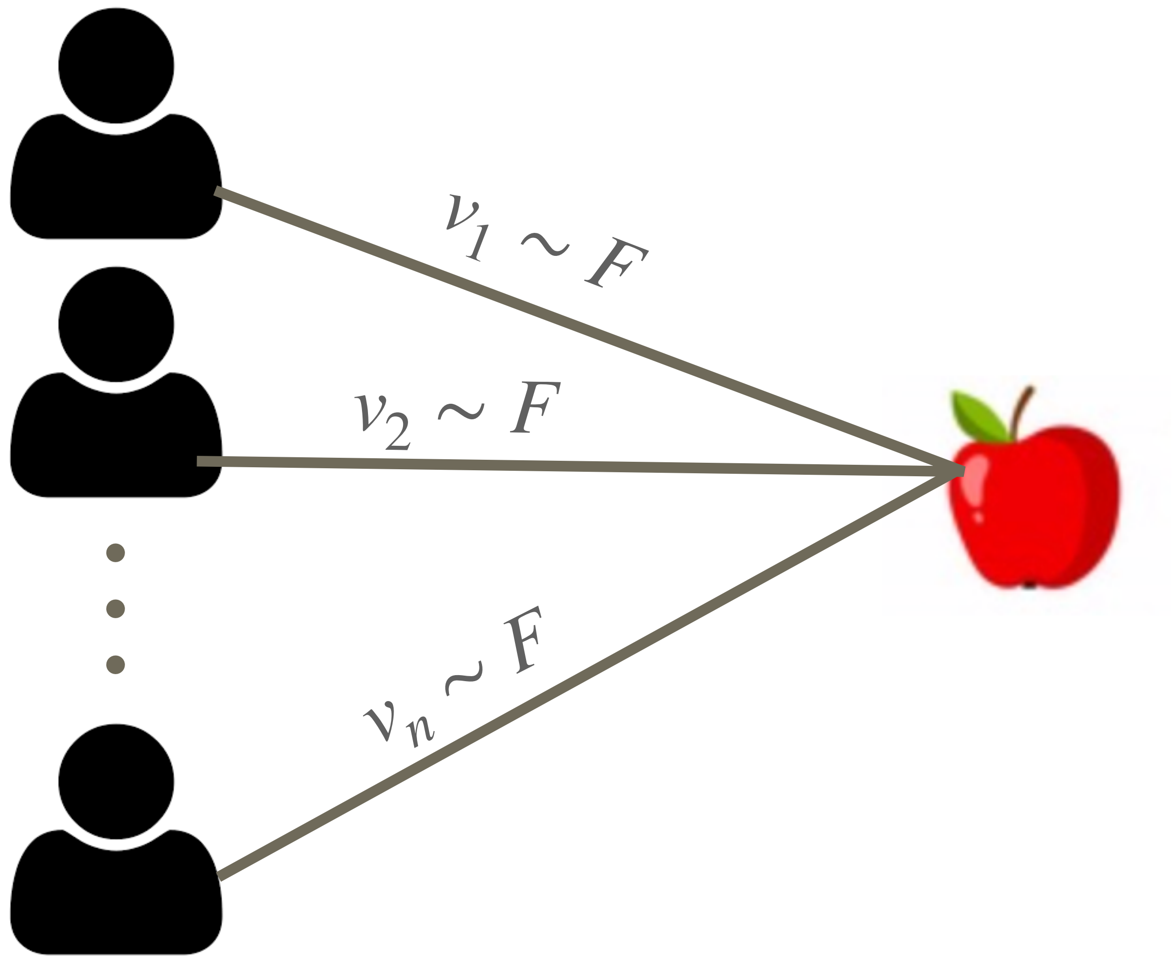

Q: What is the revenue-maximizing auction?

Myerson's optimal auction

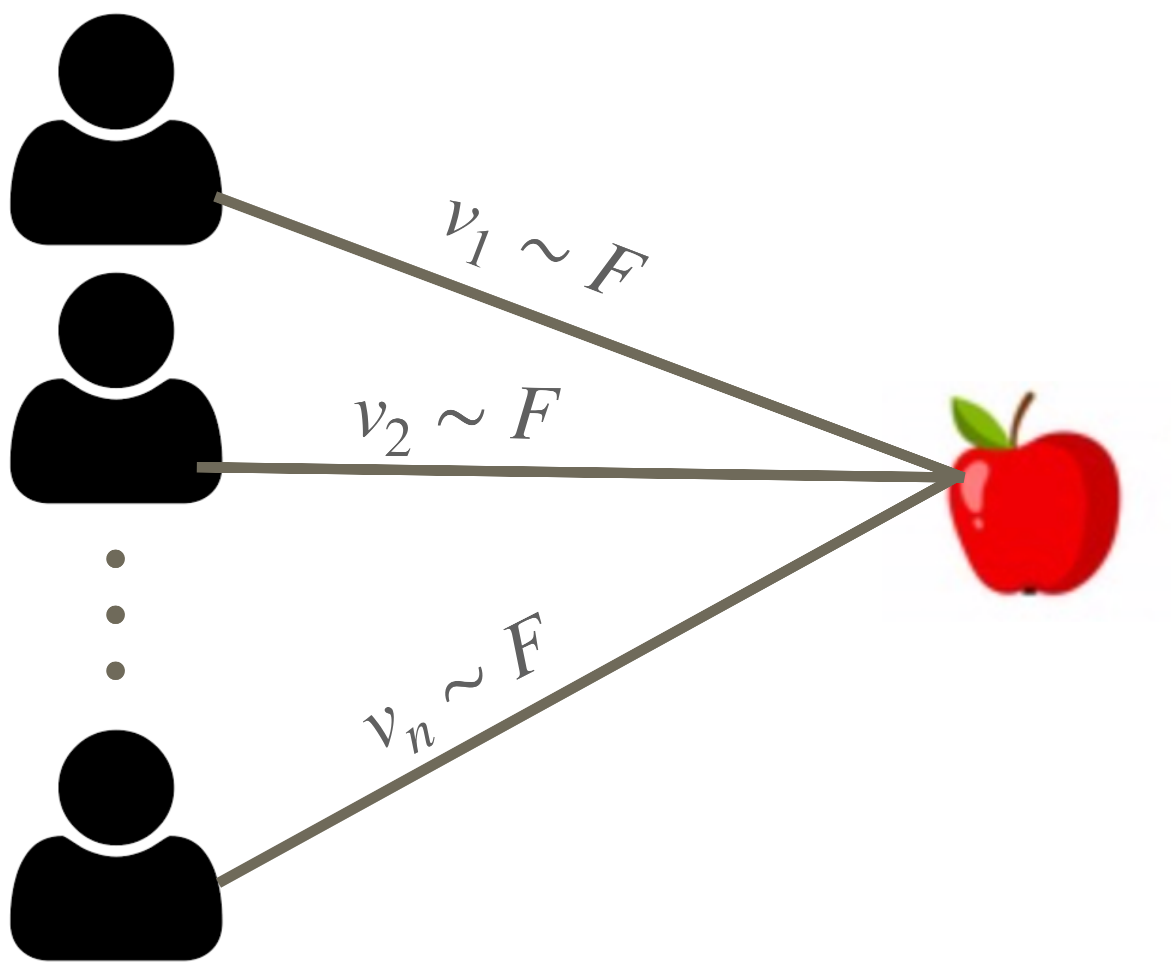

Q: What is the revenue-maximizing auction?

A: Second-price auction with a reserve price!

Myerson's optimal auction

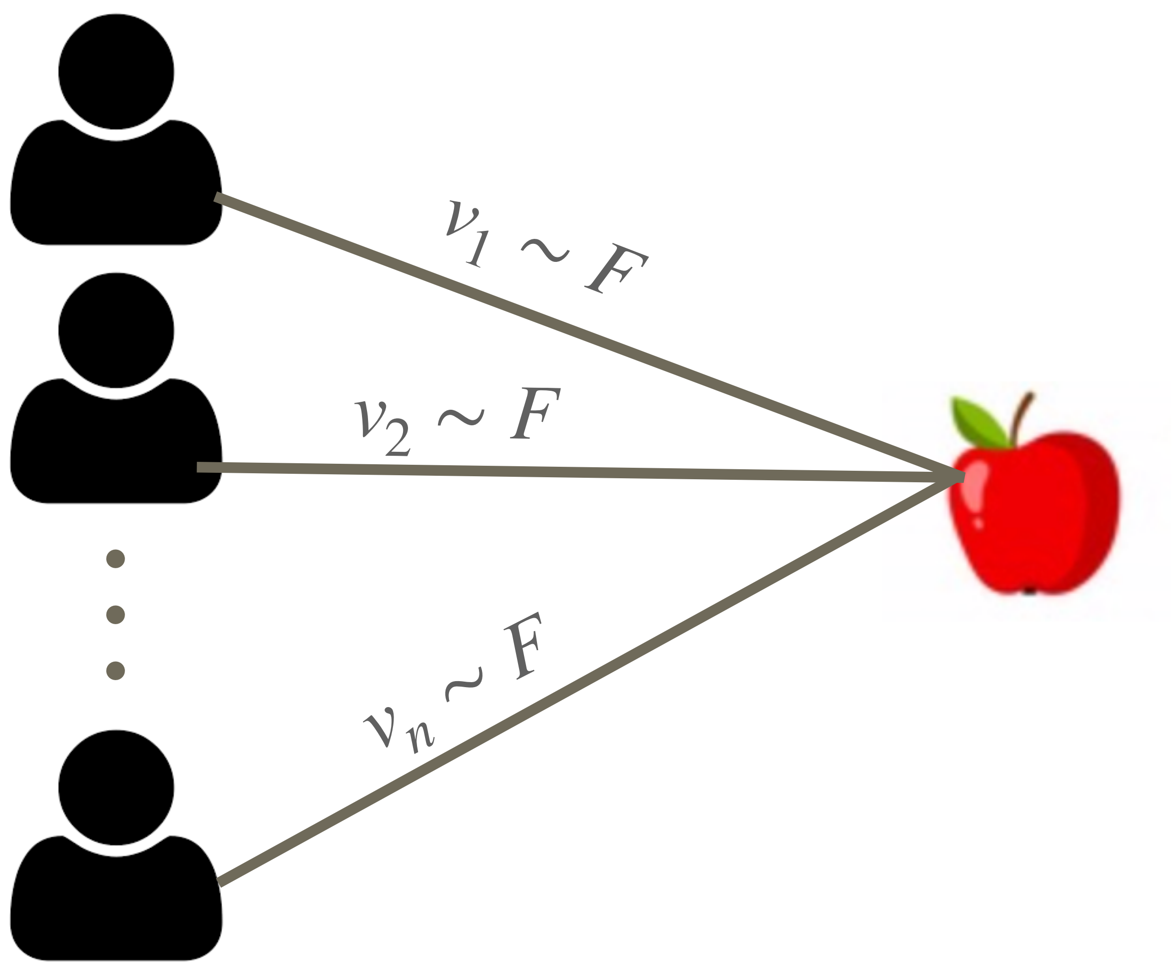

Q: What is the revenue-maximizing auction?

A: Second-price auction with a reserve price!

Full knowledge of distribution!

Motivation of our work

Motivation of our work

Exact knowledge of the distributions is rare in practice.

Motivation of our work

Exact knowledge of the distributions is rare in practice.

Need for detail-free mechanisms (Wilson's doctrine): Relax strong assumptions.

Motivation of our work

Exact knowledge of the distributions is rare in practice.

Need for detail-free mechanisms (Wilson's doctrine): Relax strong assumptions.

Settings with no access to the underlying distribution (e.g. data privacy), but statistics available.

Motivation of our work

Exact knowledge of the distributions is rare in practice.

Need for detail-free mechanisms (Wilson's doctrine): Relax strong assumptions.

Settings with no access to the underlying distribution (e.g. data privacy), but statistics available.

Sample access vs. knowledge of moments.
[Cole and Roughgarden STOC '14, Gonczarowski and Weinberg FOCS '18, Huang et al. SICOMP '18, …]

Robust Auction Design - Our model

Setting: Single additive buyer, m items.

Assumptions: Seller knows only μ_j, σ_j of each item's j marginal distribution.

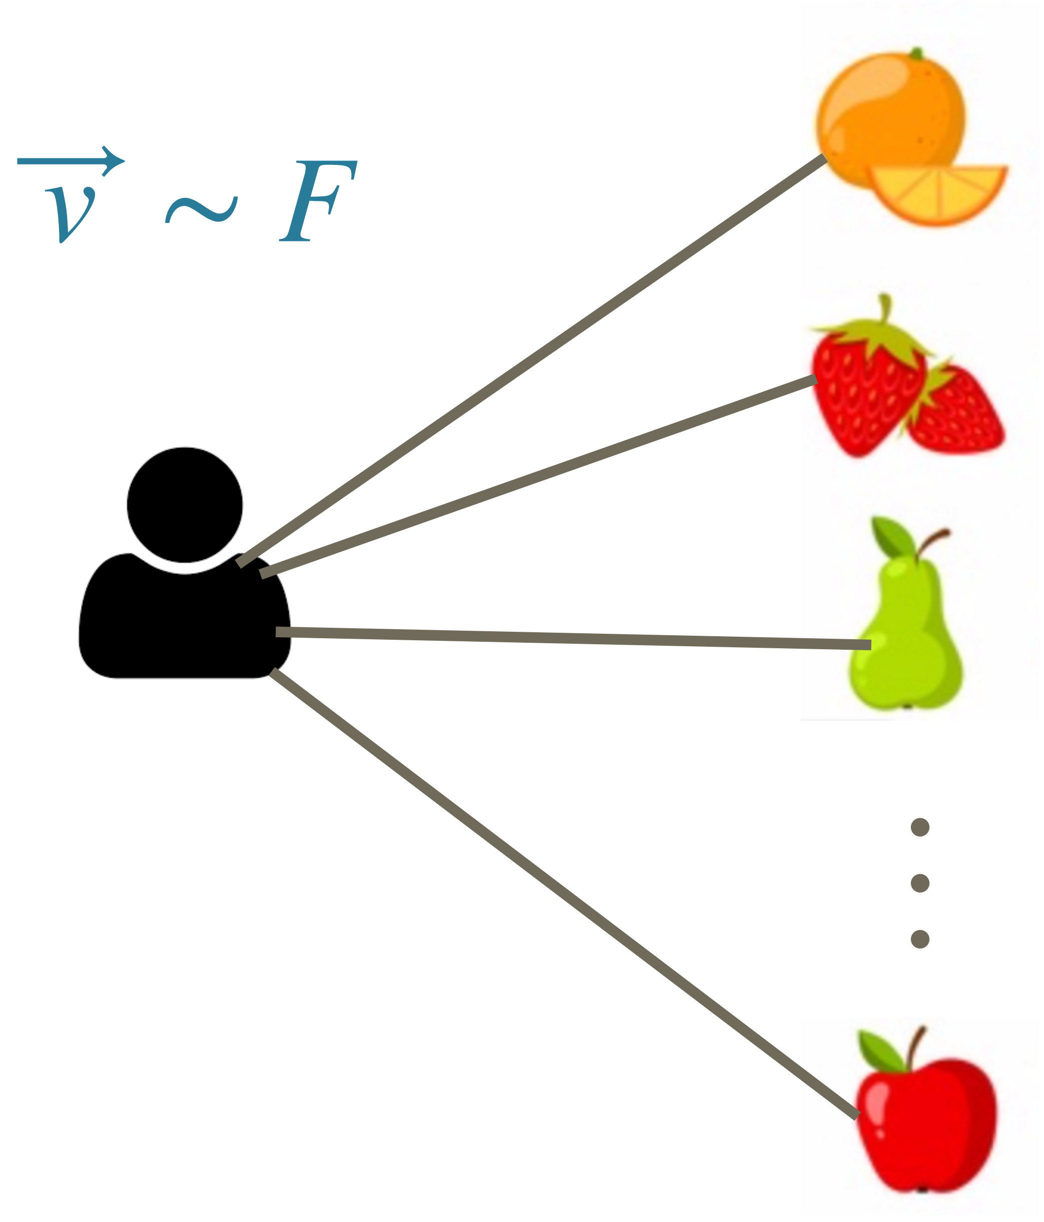

Robust Auction Design - Our model

Setting: Single additive buyer, m items.

Assumptions: Seller knows only μ_j, σ_j of each item's j marginal distribution.

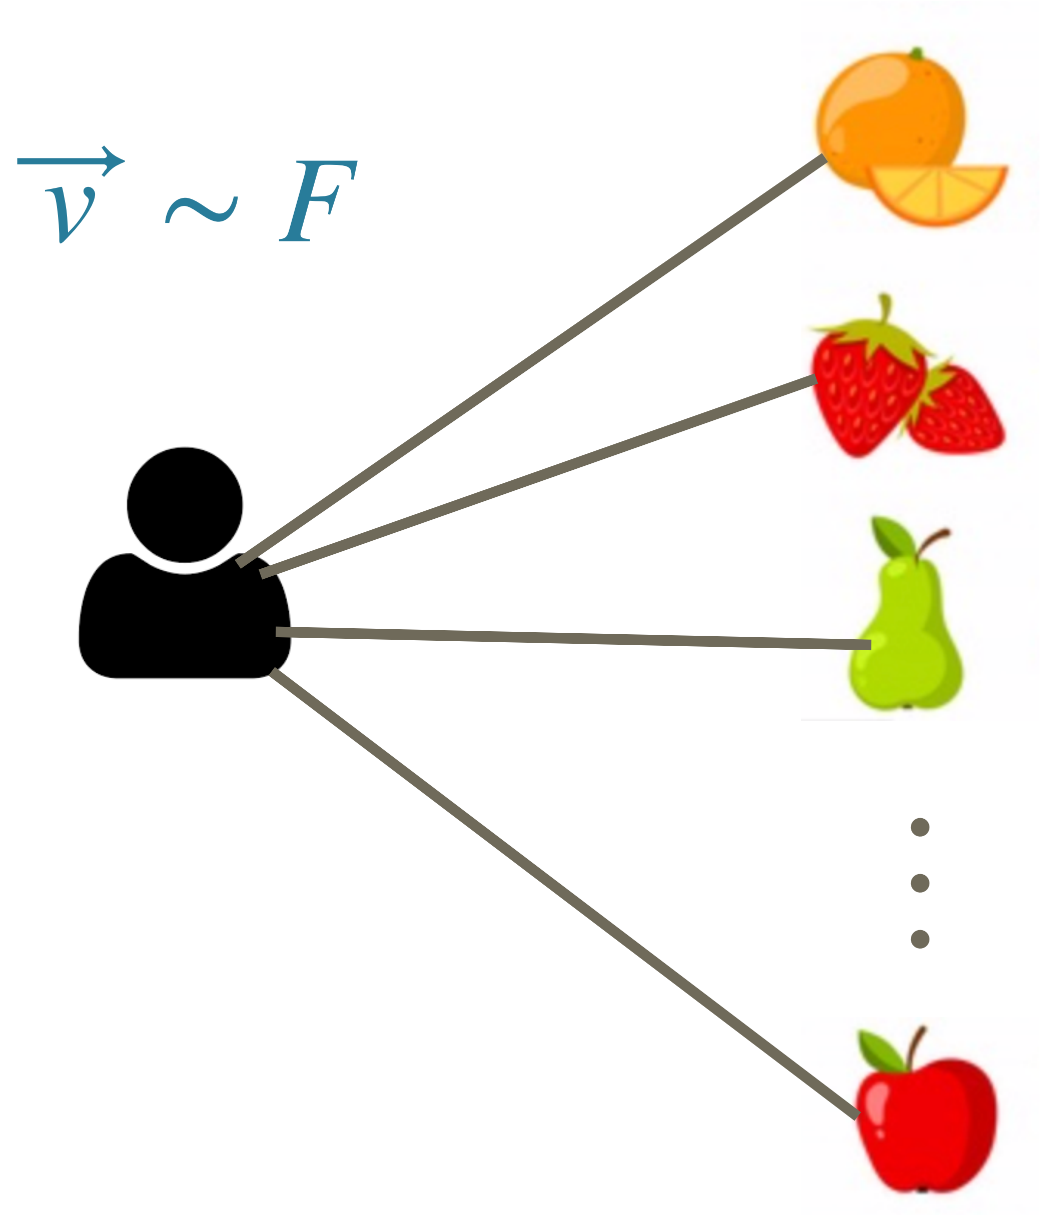

Robust Auction Design - Our model

<u>Setting</u>: Single additive buyer, m items.

<u>Assumptions</u>: Seller knows only μ_j, σ_j of each item's j marginal distribution.

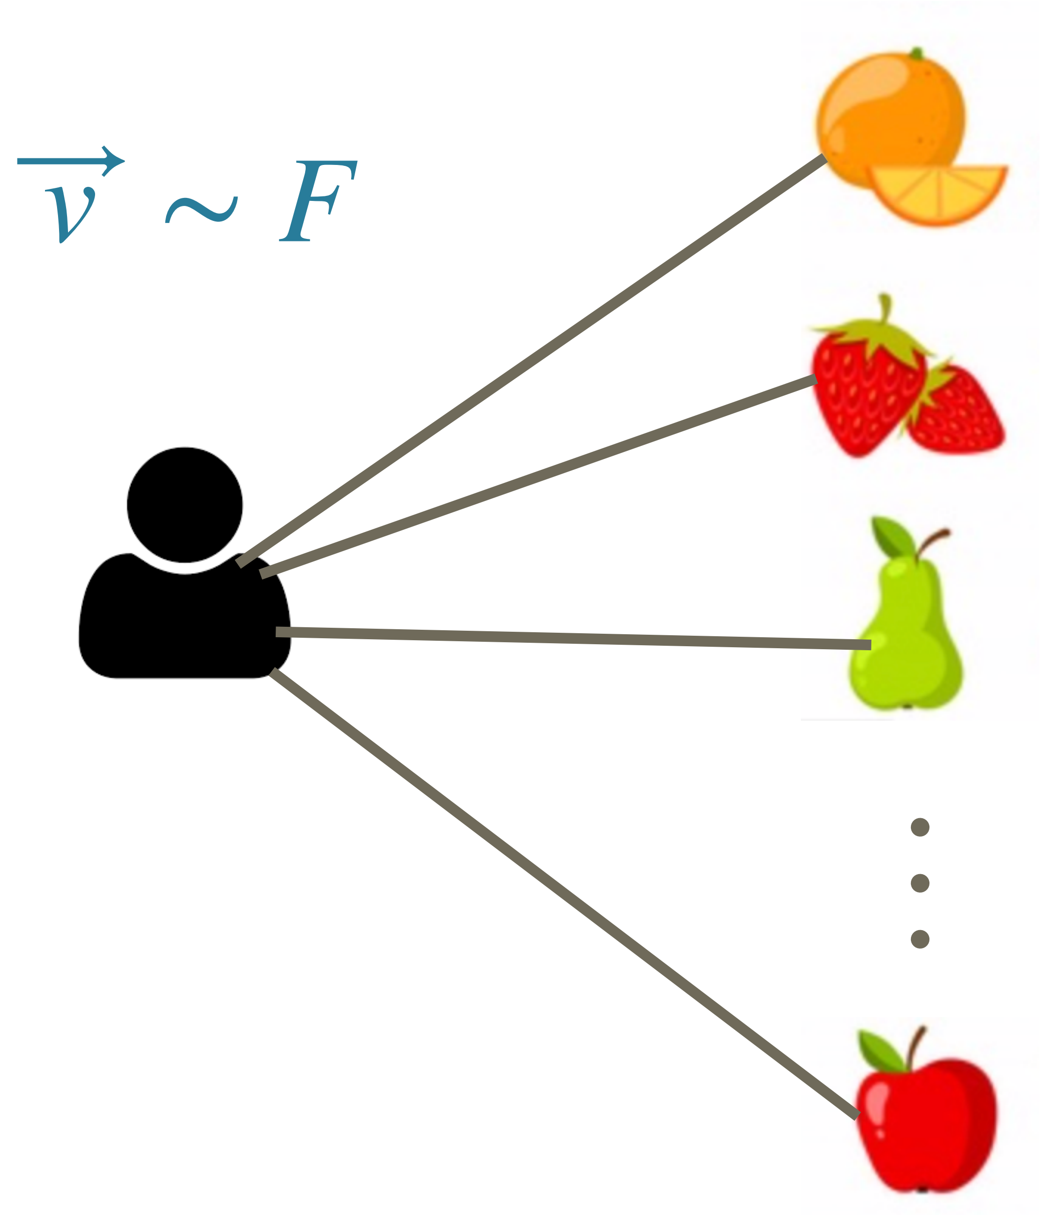

Robust Auction Design - Our model

Setting: Single additive buyer, m items.

Assumptions: Seller knows only μ_j, σ_j of each item's j marginal distribution.

$\vec{v} \sim F$

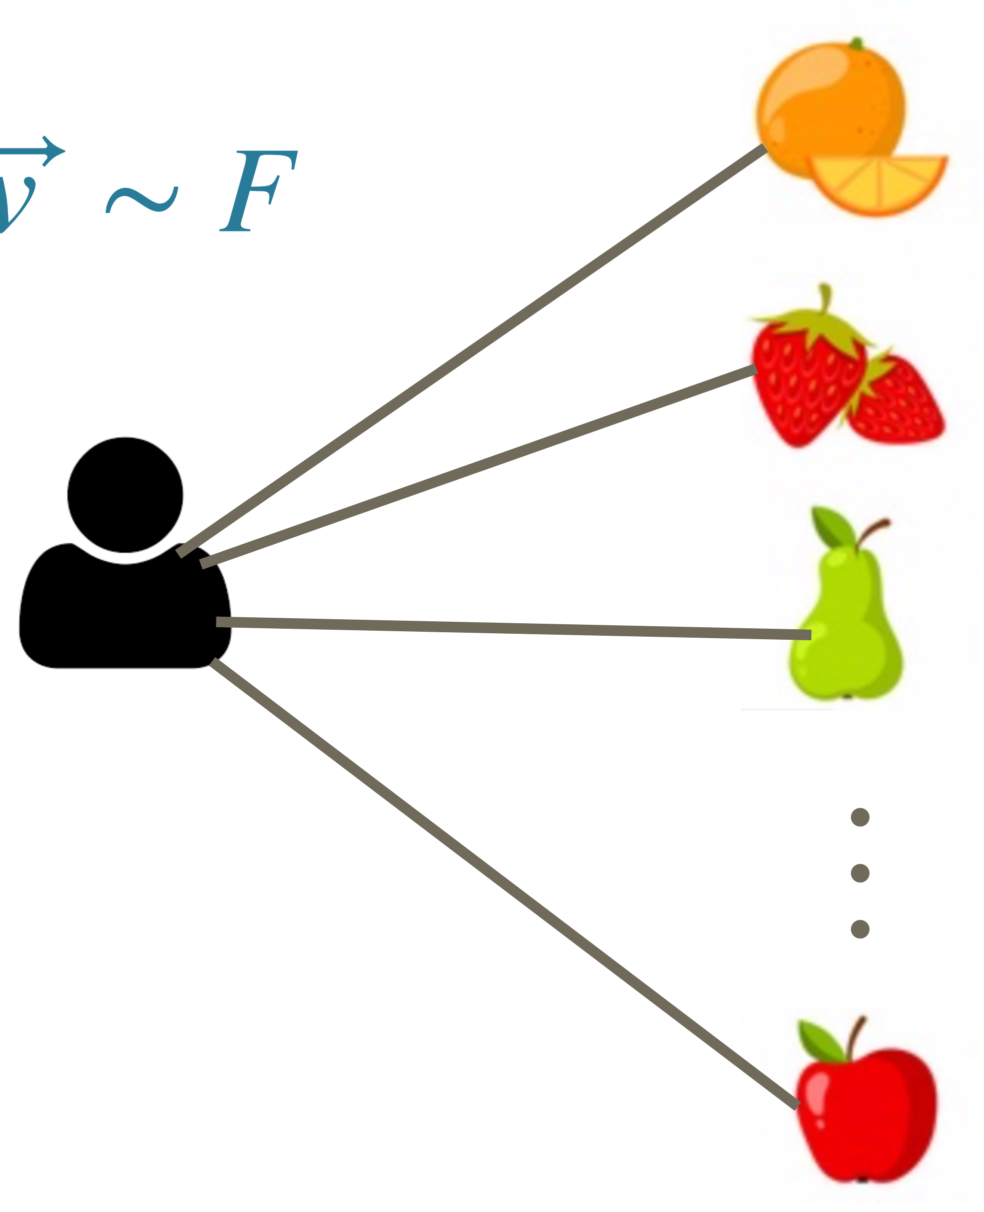

Robust Auction Design - Our model

<u>Setting</u>: Single additive buyer, m items.

<u>Assumptions</u>: Seller knows only μ_j, σ_j of each item's j marginal distribution.

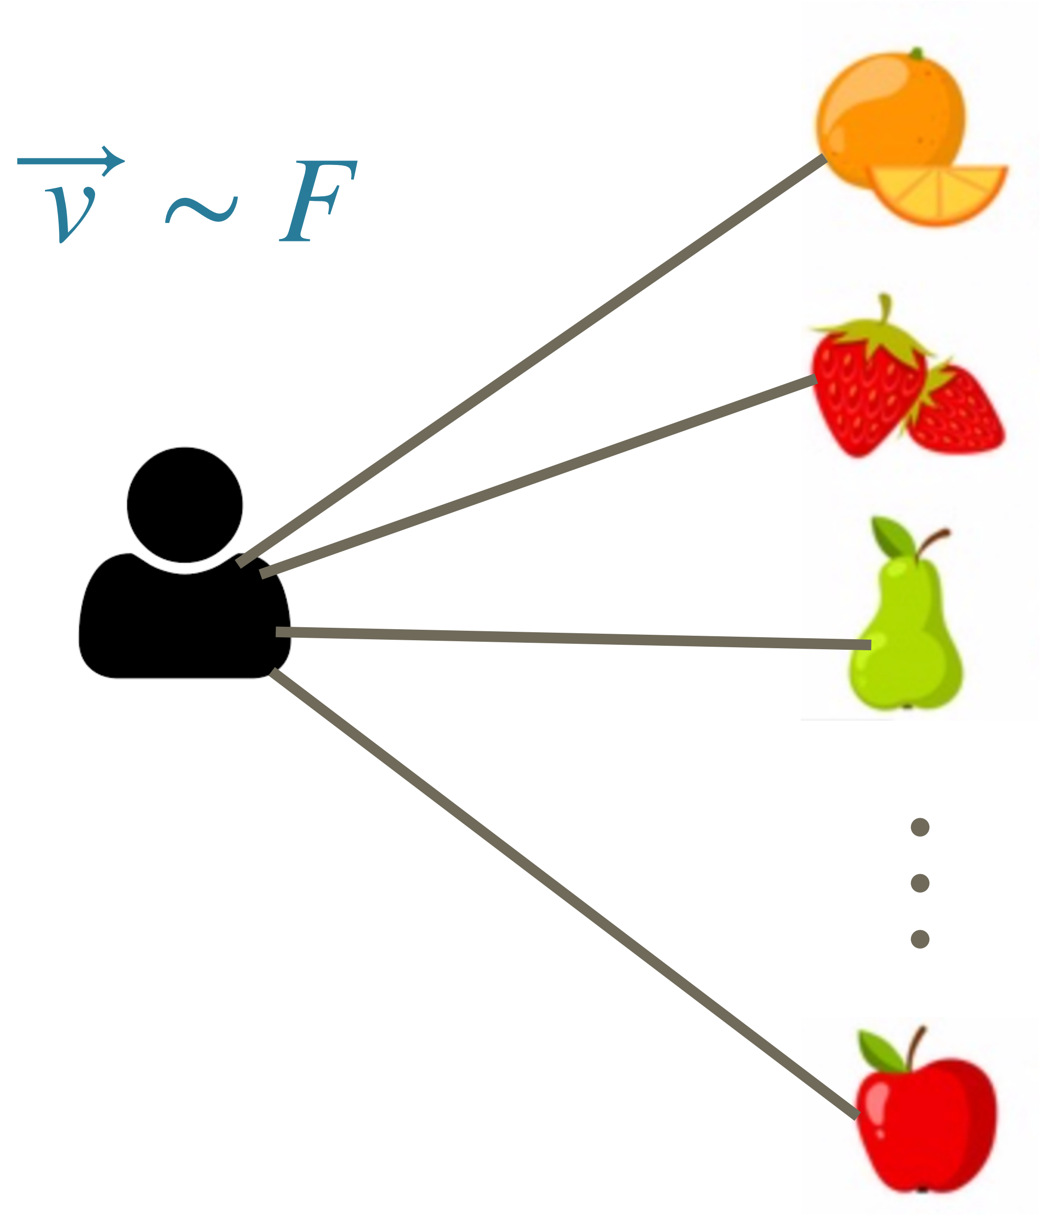

Robust Auction Design - Our model

Setting: Single additive buyer, m items.

Assumptions: Seller knows only μ_j, σ_j of each item's j marginal distribution.

$\vec{v} \sim F$

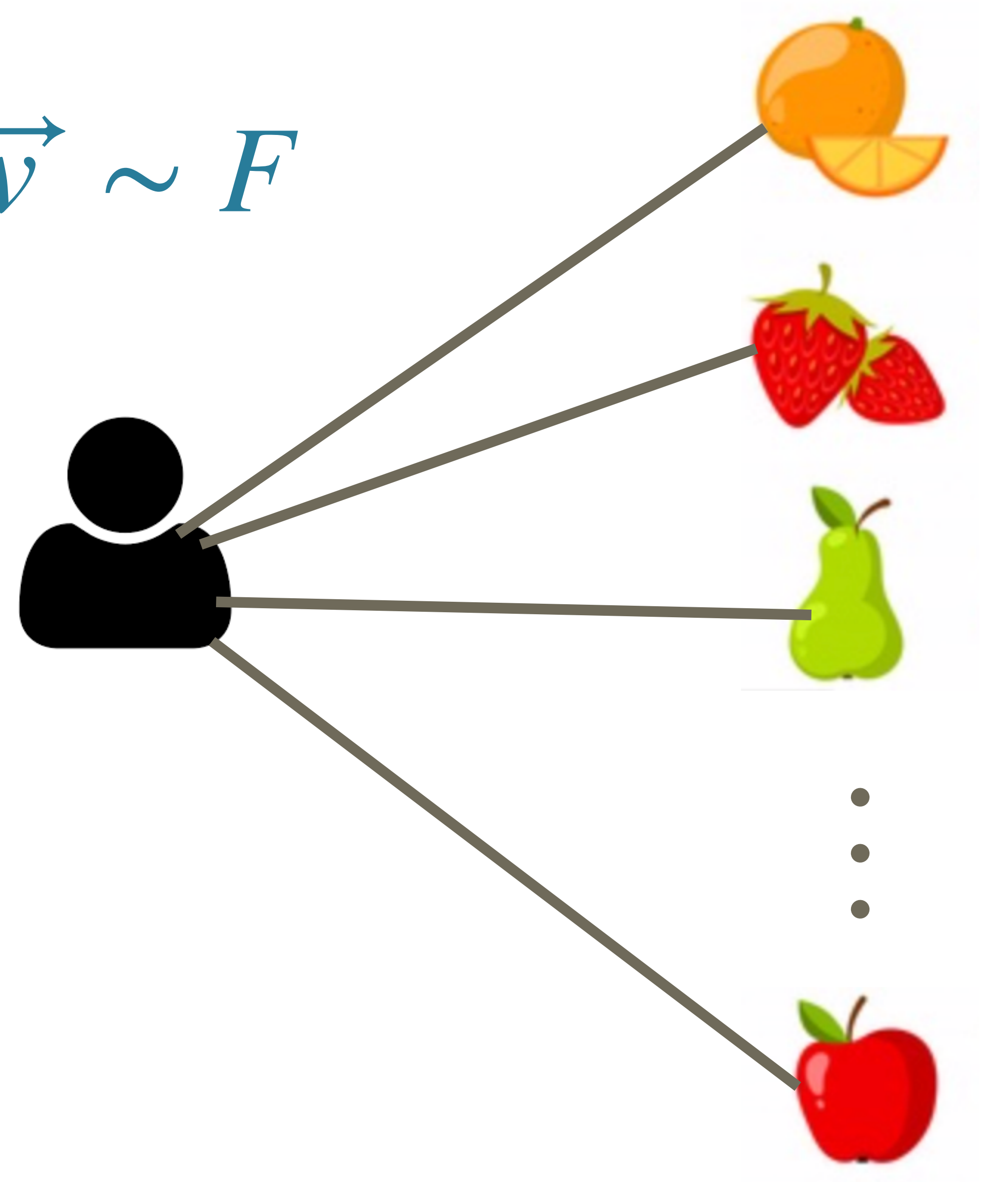

Question: Can we design mechanisms that provide good approximation guarantees?

A minimal example

A minimal example

A minimal example

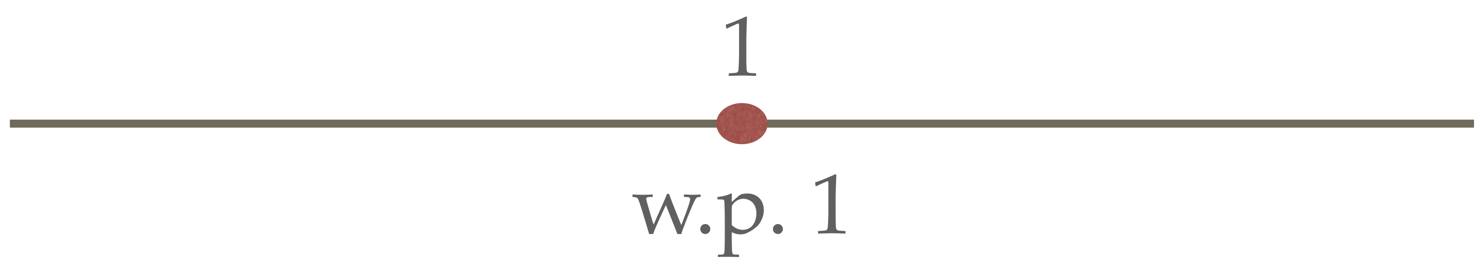

Dist. 1

A minimal example

A minimal example

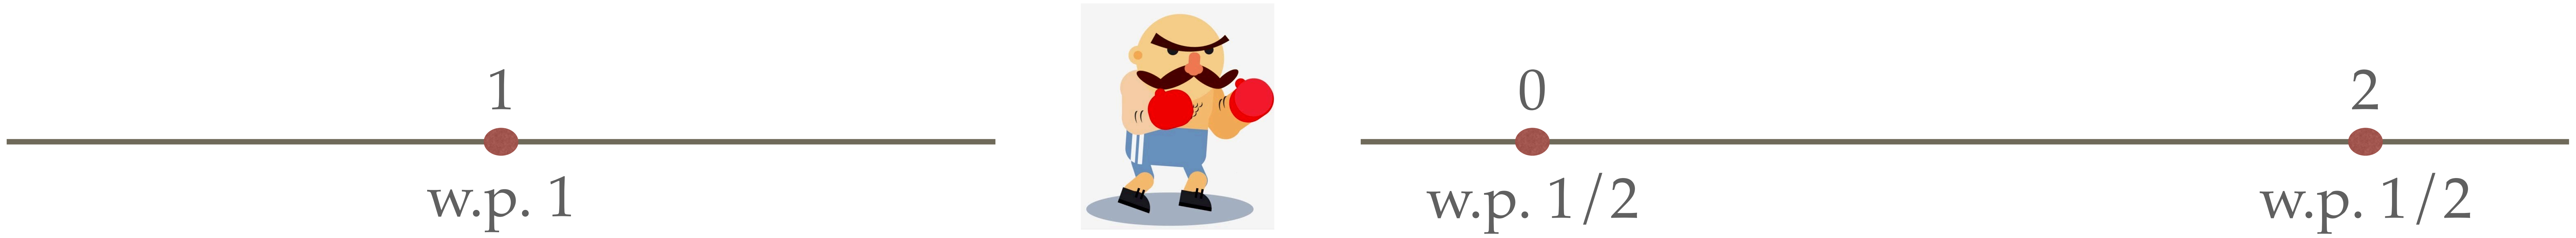

Dist. 1 Dist. 2

Best deterministic pricing ?

A minimal example

1

w.p. 1

Dist. 1

0 — w.p. 1/2

2 — w.p. 1/2

Dist. 2

Best deterministic pricing ?

Price at 1 → The adversary will play Dist. 2

$E[Rev] = 1 \cdot 1/2 = 1/2$

A minimal example

1

w.p. 1

Dist. 1

0   2

w.p. 1/2   w.p. 1/2

Dist. 2

Best deterministic pricing ?   Price at 1 → The adversary will play Dist. 2

$E[Rev] = 1 \cdot 1/2 = 1/2$

Simple randomized pricing ?

A minimal example

1

w.p. 1

Dist. 1

0 w.p. 1/2 — 2 w.p. 1/2

Dist. 2

Best deterministic pricing ? Price at 1 → The adversary will play Dist. 2

$E[Rev] = 1 \cdot 1/2 = 1/2$

Simple randomized pricing ? Price = $\begin{cases} 1, \text{ w.p. } 2/3 \\ 2, \text{ w.p. } 1/3 \end{cases}$

A minimal example

1
w.p. 1

Dist. 1

0   2
w.p. 1/2   w.p. 1/2

Dist. 2

Best deterministic pricing ? Price at 1 → The adversary will play Dist. 2
E[Rev] = $1 \cdot 1/2 = 1/2$

Simple randomized pricing ? Price = $\begin{cases} 1, \text{ w.p. } 2/3 \\ 2, \text{ w.p. } 1/3 \end{cases}$

Adversary plays Dist.1 → E[Rev] = $1 \cdot 2/3 = 2/3$

A minimal example

1

w.p. 1

Dist. 1

0 w.p. 1/2 — 2 w.p. 1/2

Dist. 2

Best deterministic pricing ? Price at 1 $\longrightarrow$ The adversary will play Dist. 2

$E[Rev] = 1 \cdot 1/2 = 1/2$

Simple randomized pricing ? Price = $\begin{cases} 1, \text{ w.p. } 2/3 \\ 2, \text{ w.p. } 1/3 \end{cases}$

Adversary plays Dist.1 $\longrightarrow$ $E[Rev] = 1 \cdot 2/3 = 2/3$

Adversary plays Dist.2 $\longrightarrow$ $E[Rev] = 1 \cdot 2/3 \cdot 1/2 + 2 \cdot 1/3 \cdot 1/2 = 2/3$

A minimal example

1 (w.p. 1) — Dist. 1

0 (w.p. 1/2), 2 (w.p. 1/2) — Dist. 2

Best deterministic pricing ? Price at 1 $\longrightarrow$ The adversary will play Dist. 2

$E[\text{Rev}] = 1 \cdot 1/2 = 1/2$

Simple randomized pricing ? $\text{Price} = \begin{cases} 1, \text{ w.p. } 2/3 \\ 2, \text{ w.p. } 1/3 \end{cases}$

Adversary plays Dist.1 $\longrightarrow$ $E[\text{Rev}] = 1 \cdot 2/3 = 2/3$

Adversary plays Dist.2 $\longrightarrow$ $E[\text{Rev}] = 1 \cdot 2/3 \cdot 1/2 + 2 \cdot 1/3 \cdot 1/2 = 2/3$

Previous results

Previous results

[Azar and Micali '12 & ITCS '13]: Deterministic, single-item, single-bidder.

Exact solution to the *maximin expected revenue*.

Previous results

[Azar and Micali '12 & ITCS '13]: Deterministic, single-item, single-bidder.

Exact solution to the *maximin expected revenue*.
(Non-tight) upper & lower bounds for the *ratio* that grow **quadratically** in $r = \sigma/\mu$.

Previous results

[Azar and Micali ’12 & ITCS ’13]: Deterministic, single-item, single-bidder.

Exact solution to the *maximin expected revenue*.
(Non-tight) upper & lower bounds for the *ratio* that grow **quadratically** in $r = \sigma/\mu$.

[Azar, Daskalakis, Micali, Weinberg SODA ’13]: Generalizes some results to n bidders.

Previous results

[Azar and Micali '12 & ITCS '13]: Deterministic, single-item, single-bidder.

Exact solution to the *maximin expected revenue*.
(Non-tight) upper & lower bounds for the *ratio* that grow **quadratically** in $r = \sigma/\mu$.

[Azar, Daskalakis, Micali, Weinberg SODA '13]: Generalizes some results to n bidders.

Seller knows some parameters of the distributions (ex. medians).
Revenue & social welfare approximation under regularity or MHR assumptions.

Previous results

[Azar and Micali '12 & ITCS '13]: Deterministic, single-item, single-bidder.

Exact solution to the *maximin expected revenue*.
(Non-tight) upper & lower bounds for the *ratio* that grow **quadratically** in $r = \sigma/\mu$.

[Azar, Daskalakis, Micali, Weinberg SODA '13]: Generalizes some results to n bidders.

Seller knows some parameters of the distributions (ex. medians).
Revenue & social welfare approximation under regularity or MHR assumptions.

[Carrasco et al. JET '18]: Maximin opt for single-item, single-bidder and *randomized* mechanisms.
[Che '19]: Generalization to n bidders.

Previous results

[Azar and Micali '12 & ITCS '13]: Deterministic, single-item, single-bidder.

Exact solution to the *maximin expected revenue*.
(Non-tight) upper & lower bounds for the *ratio* that grow **quadratically** in $r = \sigma/\mu$.

[Azar, Daskalakis, Micali, Weinberg SODA '13]: Generalizes some results to n bidders.

Seller knows some parameters of the distributions (ex. medians).
Revenue & social welfare approximation under regularity or MHR assumptions.

[Carrasco et al. JET '18]: Maximin opt for single-item, single-bidder and *randomized* mechanisms.
[Che '19]: Generalization to n bidders.

[Suzdaltsev '20]: Maximin opt. revenue, n bidders & single item, seller knows means & UB on support.

Previous results

[Azar and Micali '12 & ITCS '13]: Deterministic, single-item, single-bidder.

Exact solution to the *maximin expected revenue*.
(Non-tight) upper & lower bounds for the *ratio* that grow **quadratically** in $r = \sigma/\mu$.

[Azar, Daskalakis, Micali, Weinberg SODA '13]: Generalizes some results to n bidders.

Seller knows some parameters of the distributions (ex. medians).
Revenue & social welfare approximation under regularity or MHR assumptions.

[Carrasco et al. JET '18]: Maximin opt for single-item, single-bidder and *randomized* mechanisms.
[Che '19]: Generalization to n bidders.

[Suzdaltsev '20]: Maximin opt. revenue, n bidders & single item, seller knows means & UB on support.

None of the above tailored to the ratio benchmark! (+multi-item)

Main contributions of our paper

Our quantity of interest is the robust approximation ratio:

$r = \frac{\sigma}{\mu}$ is the CV.

$$\mathrm{APX}(\vec{\mu}, \vec{\sigma}) = \inf_{\textit{mechs}} \sup_{\textit{distribs}} \frac{\mathrm{OPT}(F)}{\mathrm{REV}(A; F)}$$

Main contributions of our paper

Our quantity of interest is the robust approximation ratio:

$r = \frac{\sigma}{\mu}$ is the CV.

$$\mathrm{APX}(\vec{\mu}, \vec{\sigma}) = \inf_{\textit{mechs}} \sup_{\textit{distribs}} \frac{\mathrm{OPT}(F)}{\mathrm{REV}(A; F)}$$

Thm1: The deterministic $\mathrm{APX}(\mu, \sigma)$ of selling a single (μ, σ)-distributed item is $\rho_D(r) \approx 1 + 4 \cdot r^2$.
This is achieved by offering a take-it-or-leave-it price of $p = \frac{\rho_D(r)}{2\rho_D(r) - 1} \cdot \mu$.

Main contributions of our paper

Our quantity of interest is the robust approximation ratio:

$r = \frac{\sigma}{\mu}$ is the CV.

$$\mathrm{APX}(\vec{\mu}, \vec{\sigma}) = \inf_{\textit{mechs}} \sup_{\textit{distribs}} \frac{\mathrm{OPT}(F)}{\mathrm{REV}(A; F)}$$

Thm1: The deterministic $\mathrm{APX}(\mu, \sigma)$ of selling a single (μ, σ)-distributed item is $\rho_D(r) \approx 1 + 4 \cdot r^2$. This is achieved by offering a take-it-or-leave-it price of $p = \frac{\rho_D(r)}{2\rho_D(r) - 1} \cdot \mu$.

Thm2: The randomized $\mathrm{APX}(\mu, \sigma)$ for single items is $\rho(r) \approx 1 + \ln(1 + r^2)$. It is *asymptotically tight* and is achieved by randomization over posted prices.

Main contributions of our paper

Our quantity of interest is the robust approximation ratio:

$r = \frac{\sigma}{\mu}$ is the CV.

$$\mathrm{APX}(\vec{\mu}, \vec{\sigma}) = \inf_{mechs} \sup_{distribs} \frac{\mathrm{OPT}(F)}{\mathrm{REV}(A; F)}$$

Thm1: The deterministic $\mathrm{APX}(\mu, \sigma)$ of selling a single (μ, σ)-distributed item is $\rho_D(r) \approx 1 + 4 \cdot r^2$. This is achieved by offering a take-it-or-leave-it price of $p = \frac{\rho_D(r)}{2\rho_D(r) - 1} \cdot \mu$.

Thm2: The randomized $\mathrm{APX}(\mu, \sigma)$ for single items is $\rho(r) \approx 1 + \ln(1 + r^2)$. It is *asymptotically tight* and is achieved by randomization over posted prices.

Thm3: When selling m (possibly correlated) $(\vec{\mu}, \vec{\sigma})$-distributed items then $\mathrm{APX}(\vec{\mu}, \vec{\sigma})$ is $\mathcal{O}(\log r_{\max})$. Mechanism: Sell each item separately with the lottery of Thm2.

The bounds for small values of r

Deterministic

$\mathcal{O}(r^2)$

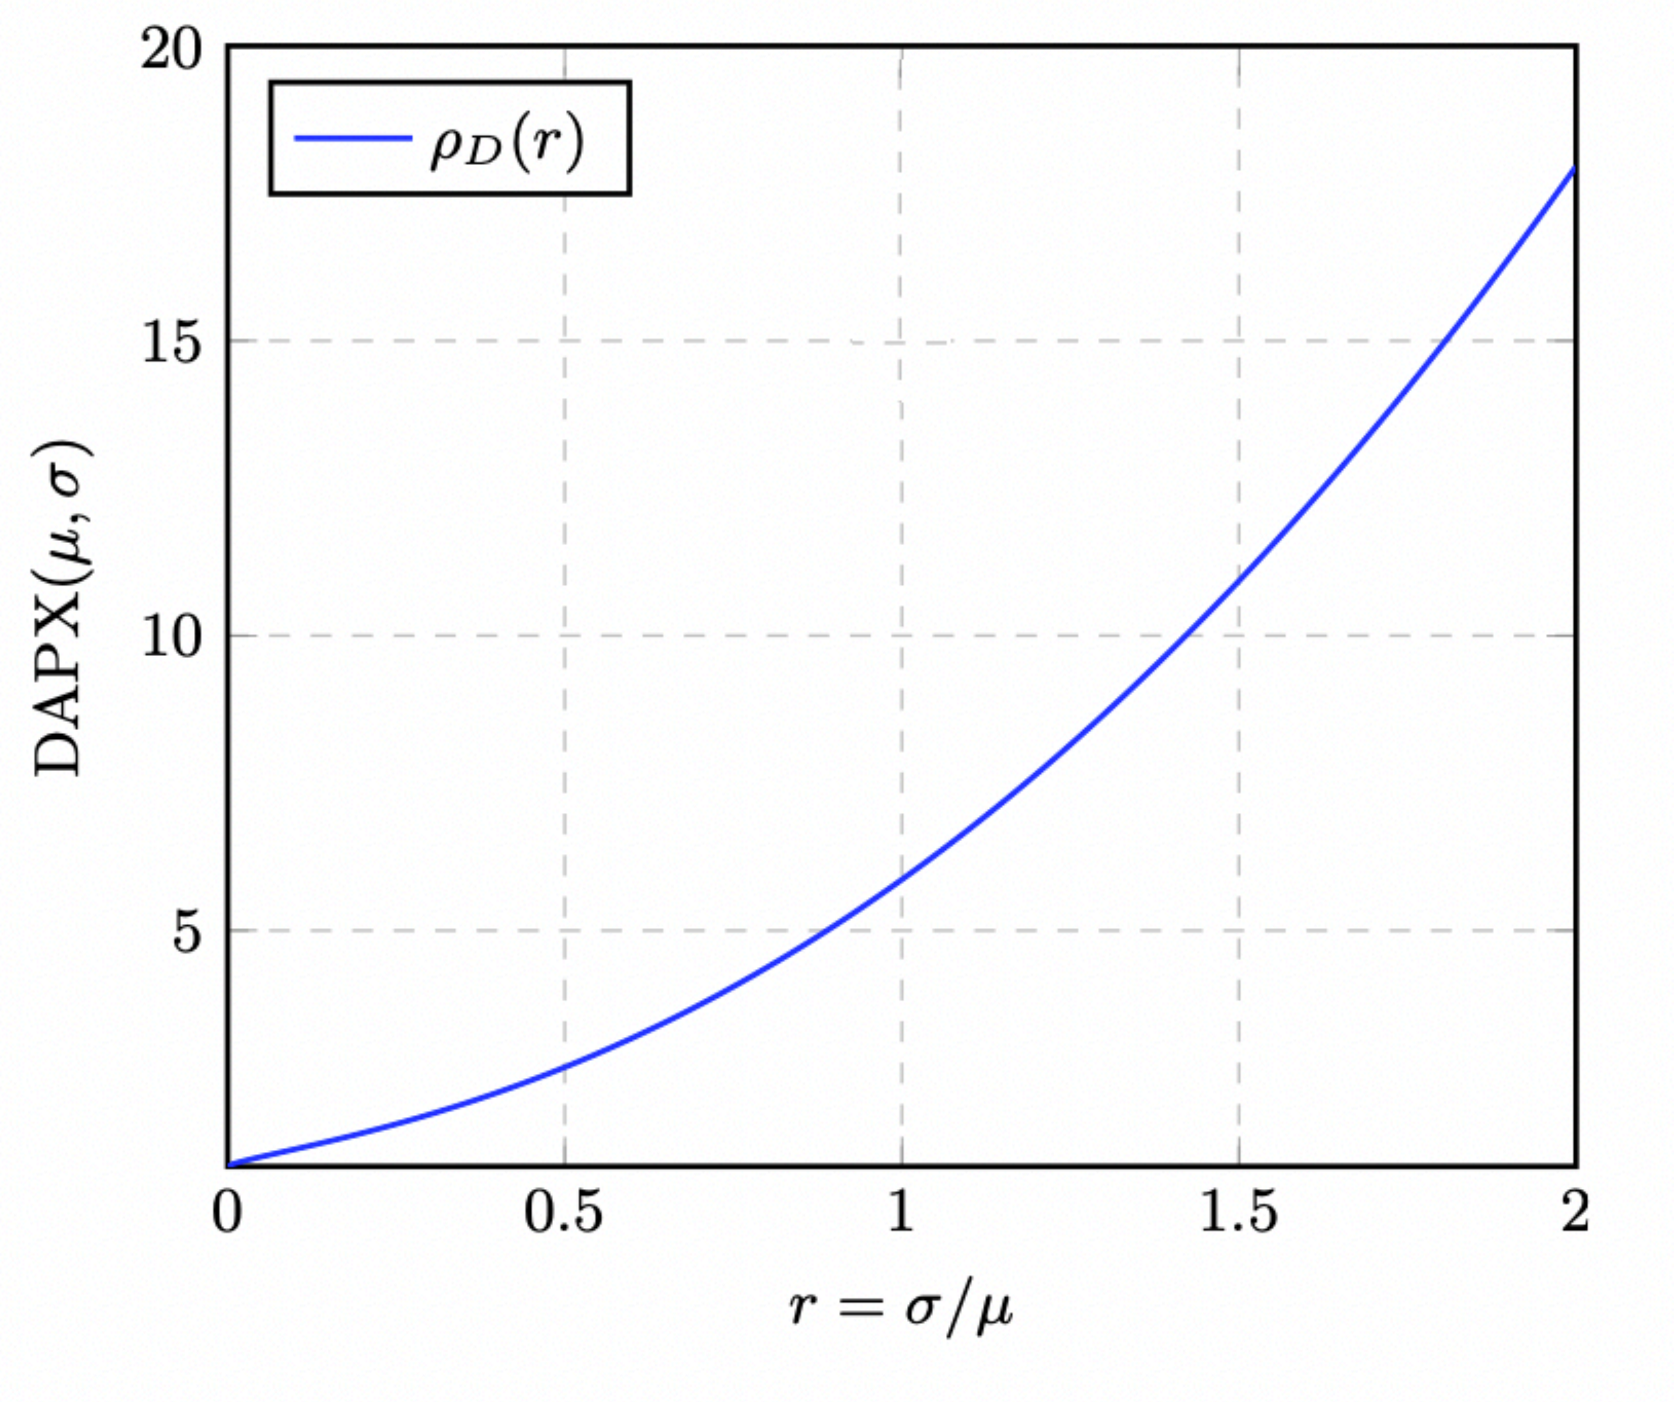

Randomized

$\mathcal{O}(\log r)$

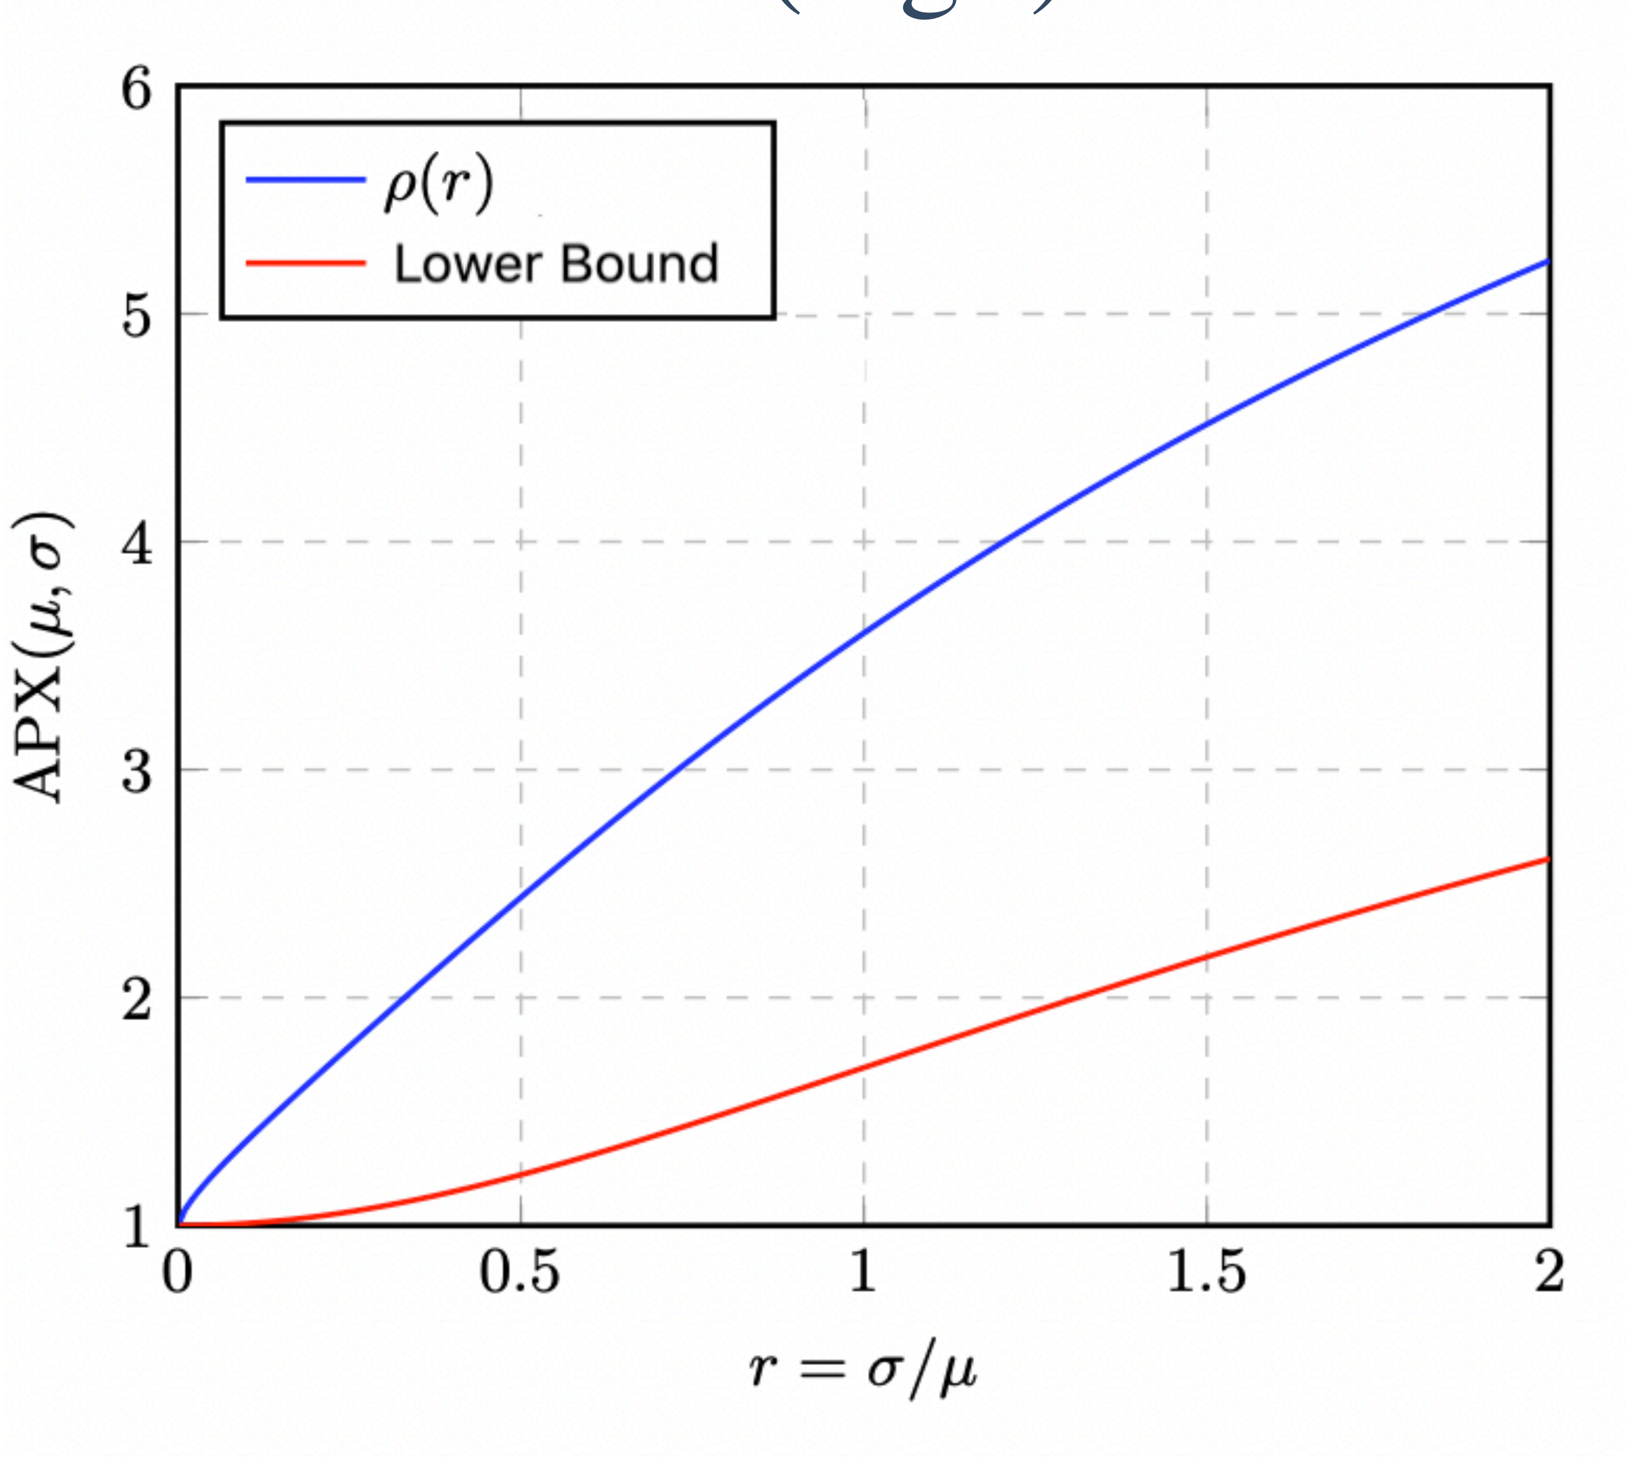

The bounds for small values of r

Deterministic

$\mathcal{O}(r^2)$

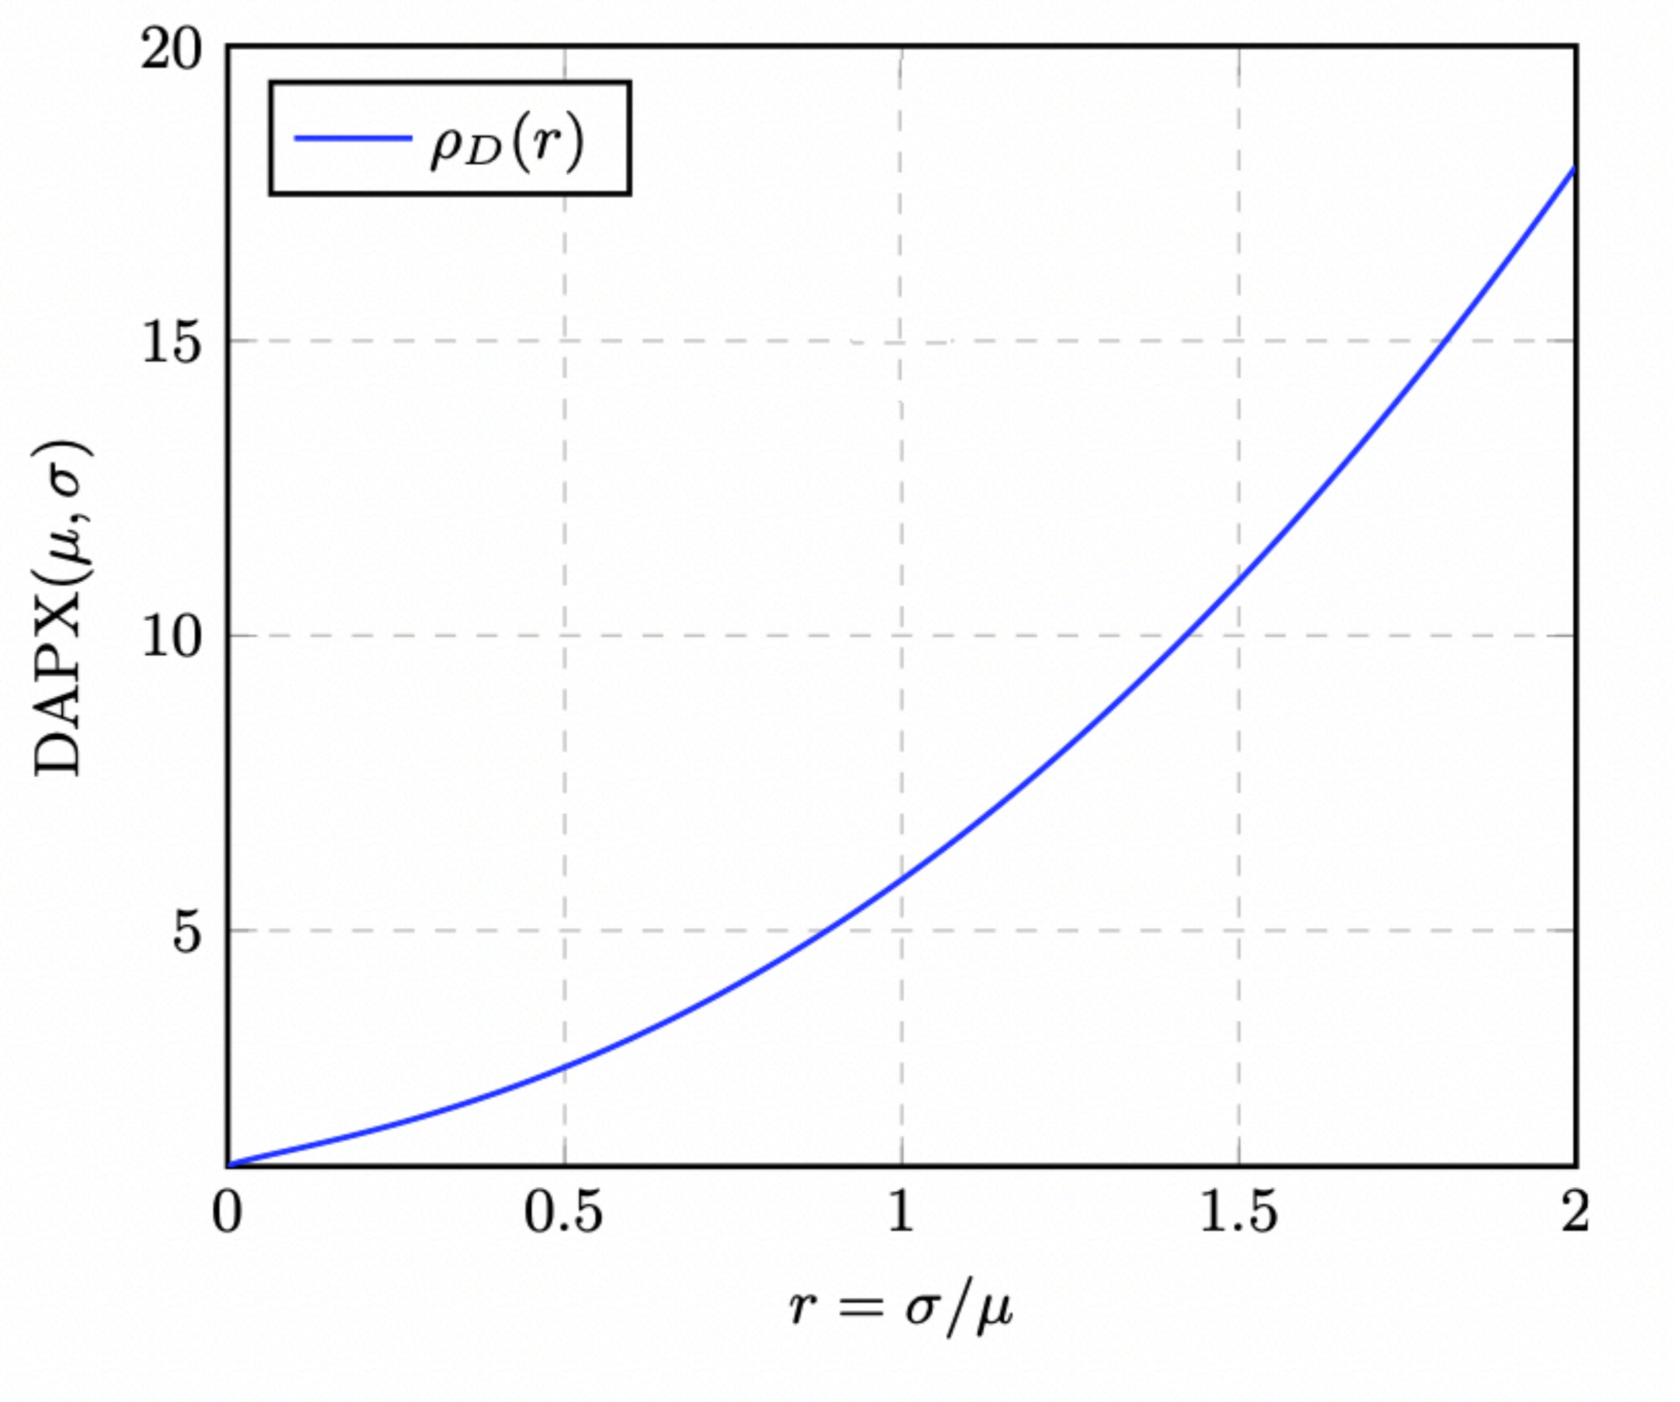

Randomized

$\mathcal{O}(\log r)$

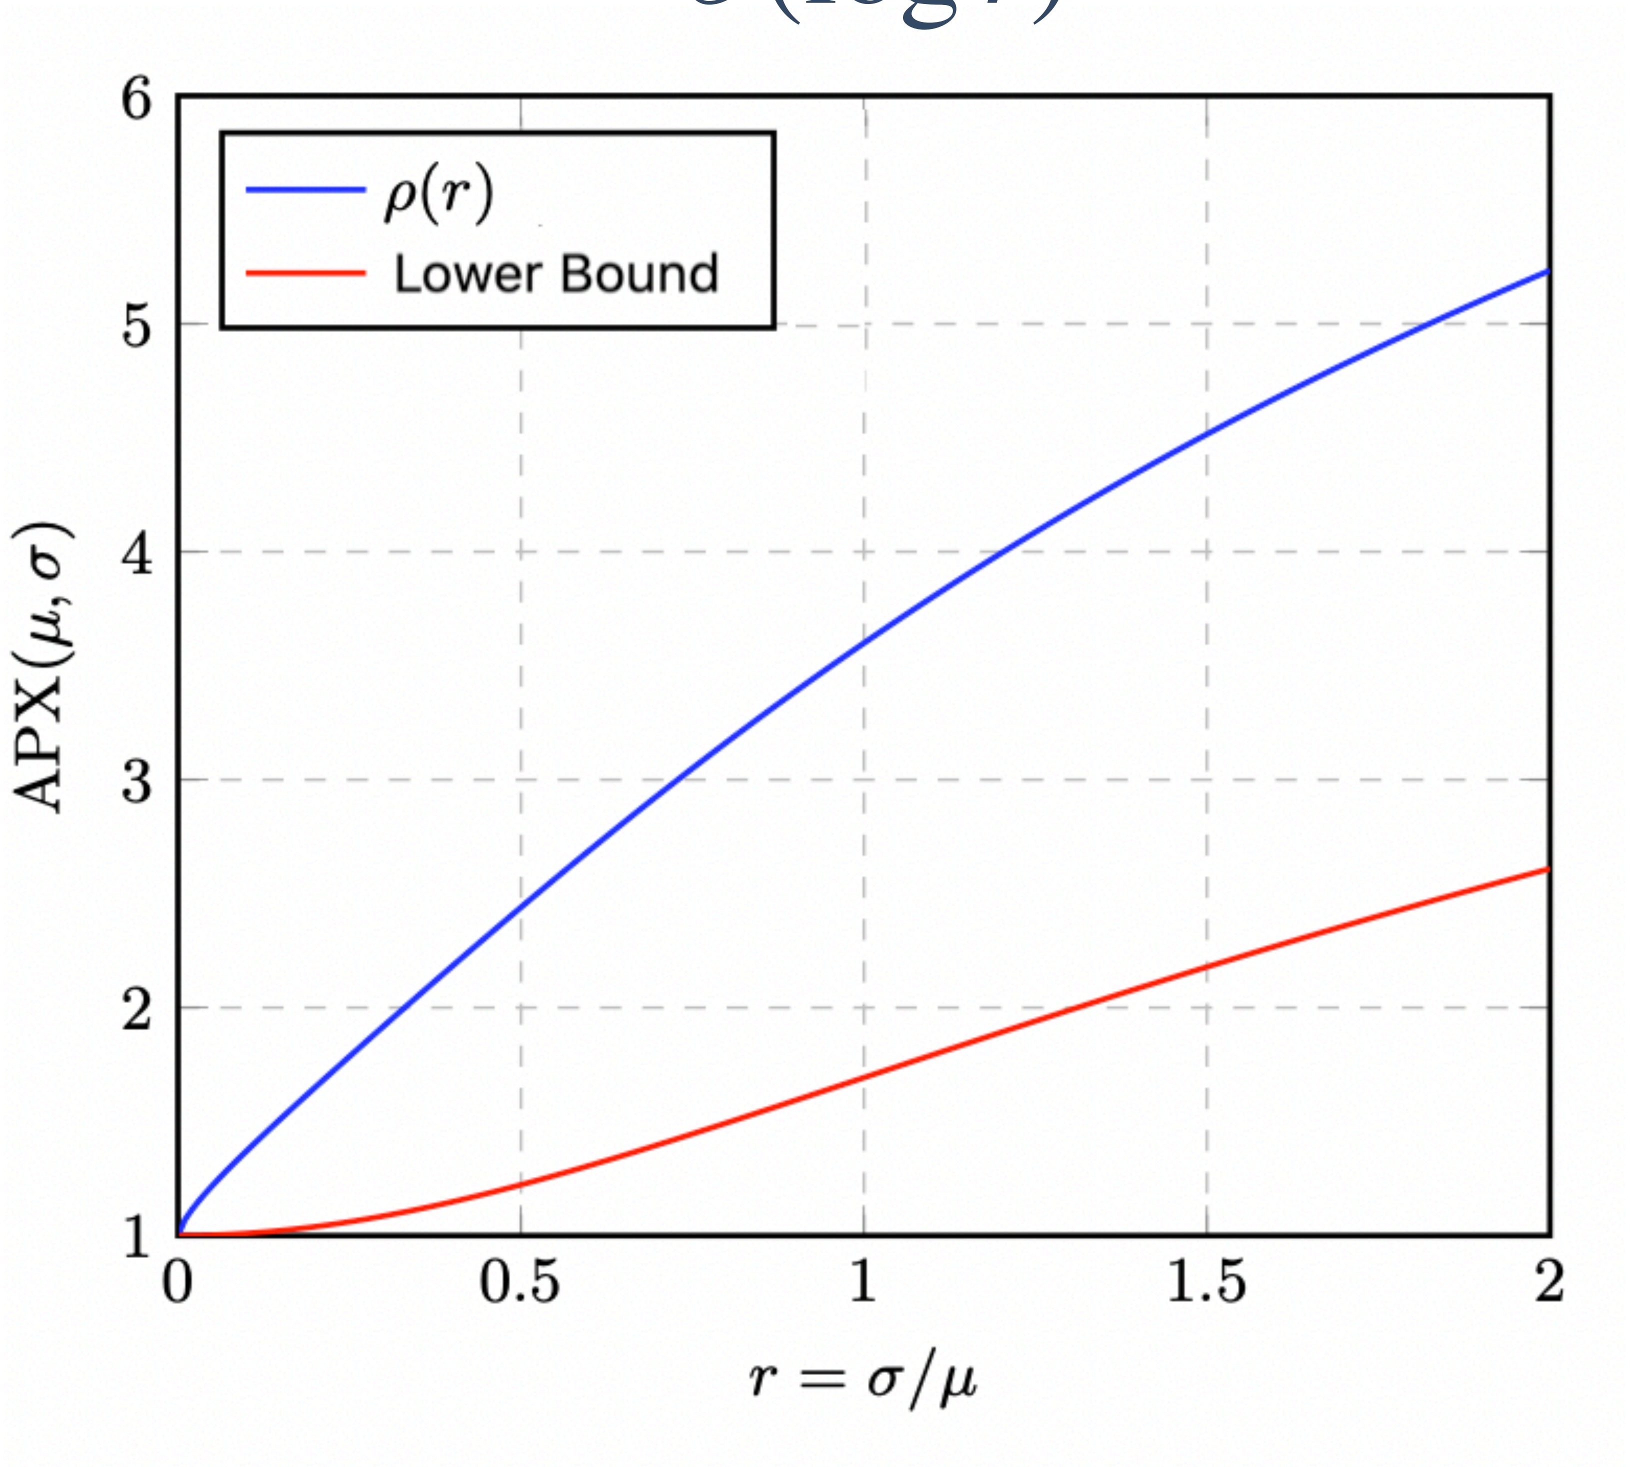

CV is small for known classes of distributions (e.g. MHR).

If bounded by universal constant, then APX$(\vec{\mu}, \vec{\sigma})$ constant!

Lower bound for randomized algorithms

Lower bound for randomized algorithms

Proof via generalized Yao's principle

Lower bound for randomized algorithms

Proof via generalized Yao's principle

$$\inf_{A \in \mathbb{A}_1} \sup_{F \in \mathbb{F}_{\mu,\sigma}} \frac{\mathrm{OPT}(F)}{\mathrm{REV}(A; F)} \geq \sup_{(B,F) \in \Delta_{\mu,\sigma}} \inf_{p \geq 0} \frac{\mathbb{E}_{\varepsilon \sim B}[\mathrm{OPT}(F_\varepsilon)]}{\mathbb{E}_{\varepsilon \sim B}[\mathrm{REV}(p; F_\varepsilon)]}$$

Lower bound for randomized algorithms

Proof via generalized Yao's principle

$$\inf_{A \in \mathbb{A}_1} \sup_{F \in \mathbb{F}_{\mu,\sigma}} \frac{\mathrm{OPT}(F)}{\mathrm{REV}(A;F)} \geq \underbrace{\sup_{(B,F) \in \Delta_{\mu,\sigma}}}_{\text{"Distributions over Distributions"}} \inf_{p \geq 0} \frac{\mathbb{E}_{\varepsilon \sim B}[\mathrm{OPT}(F_\varepsilon)]}{\mathbb{E}_{\varepsilon \sim B}[\mathrm{REV}(p;F_\varepsilon)]}$$

Lower bound for randomized algorithms

Proof via generalized Yao's principle

$$\inf_{A \in \mathbb{A}_1} \sup_{F \in \mathbb{F}_{\mu,\sigma}} \frac{\mathrm{OPT}(F)}{\mathrm{REV}(A; F)} \geq \underbrace{\sup_{(B,F) \in \Delta_{\mu,\sigma}}}_{\text{“Distributions over Distributions”}} \inf_{p \geq 0} \boxed{\frac{\mathbb{E}_{\varepsilon \sim B}[\mathrm{OPT}(F_\varepsilon)]}{\mathbb{E}_{\varepsilon \sim B}[\mathrm{REV}(p; F_\varepsilon)]}}$$

Lower bound for randomized algorithms

Proof via generalized Yao's principle

$$\inf_{A \in \mathbb{A}_1} \sup_{F \in \mathbb{F}_{\mu,\sigma}} \frac{\mathrm{OPT}(F)}{\mathrm{REV}(A; F)} \geq \underbrace{\sup_{(B,F) \in \Delta_{\mu,\sigma}}} \inf_{p \geq 0} \frac{\mathbb{E}_{\varepsilon \sim B}[\mathrm{OPT}(F_\varepsilon)]}{\mathbb{E}_{\varepsilon \sim B}[\mathrm{REV}(p; F_\varepsilon)]}$$

"Distributions over Distributions"

Posterior distribution $\mathbb{E}_{\varepsilon \sim B}[F_\varepsilon](z) = \int F(z; \varepsilon) dB(\varepsilon)$

Lower bound - Construction of the mixture

Lower bound - Construction of the mixture

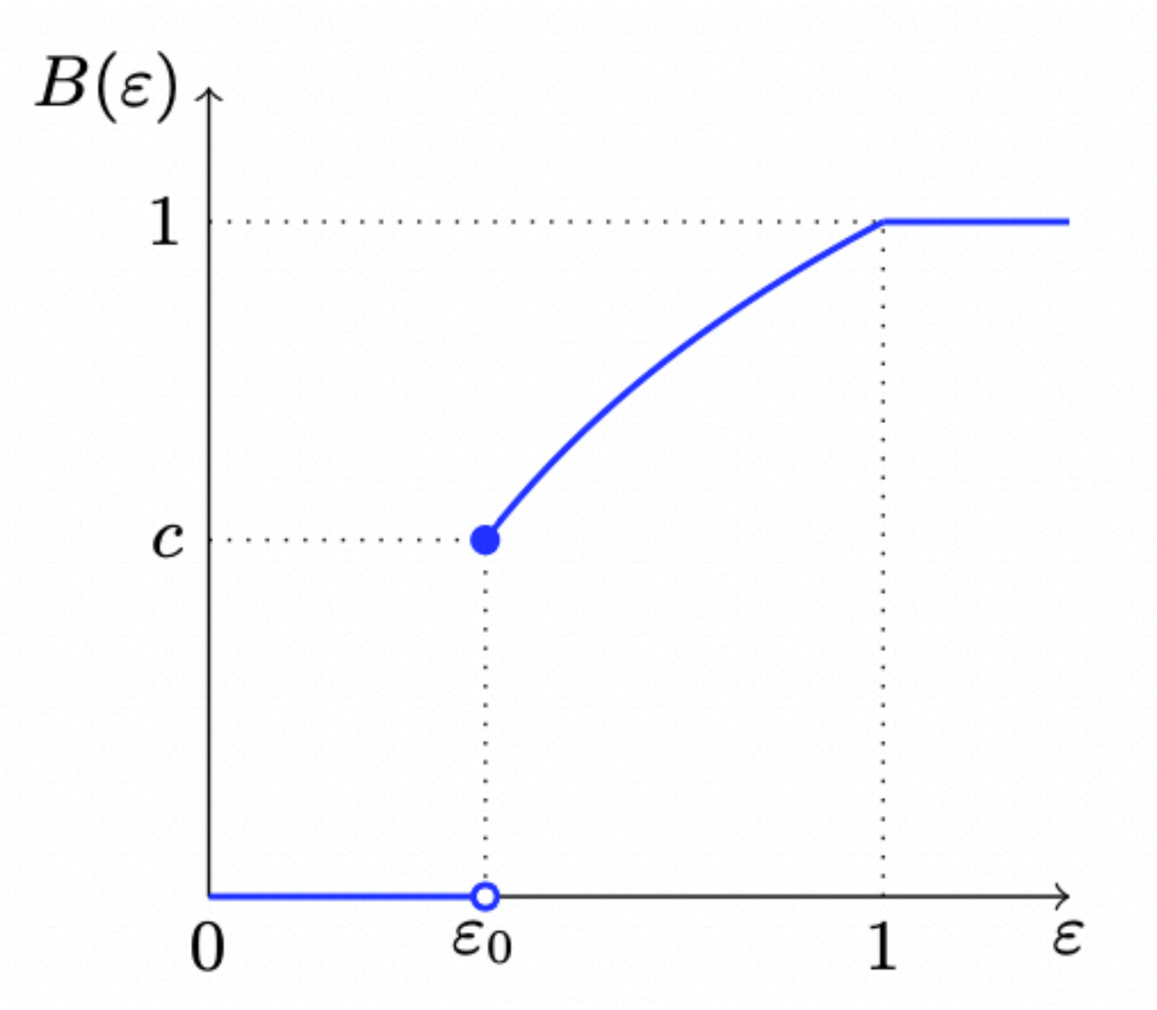

Lower bound - Construction of the mixture

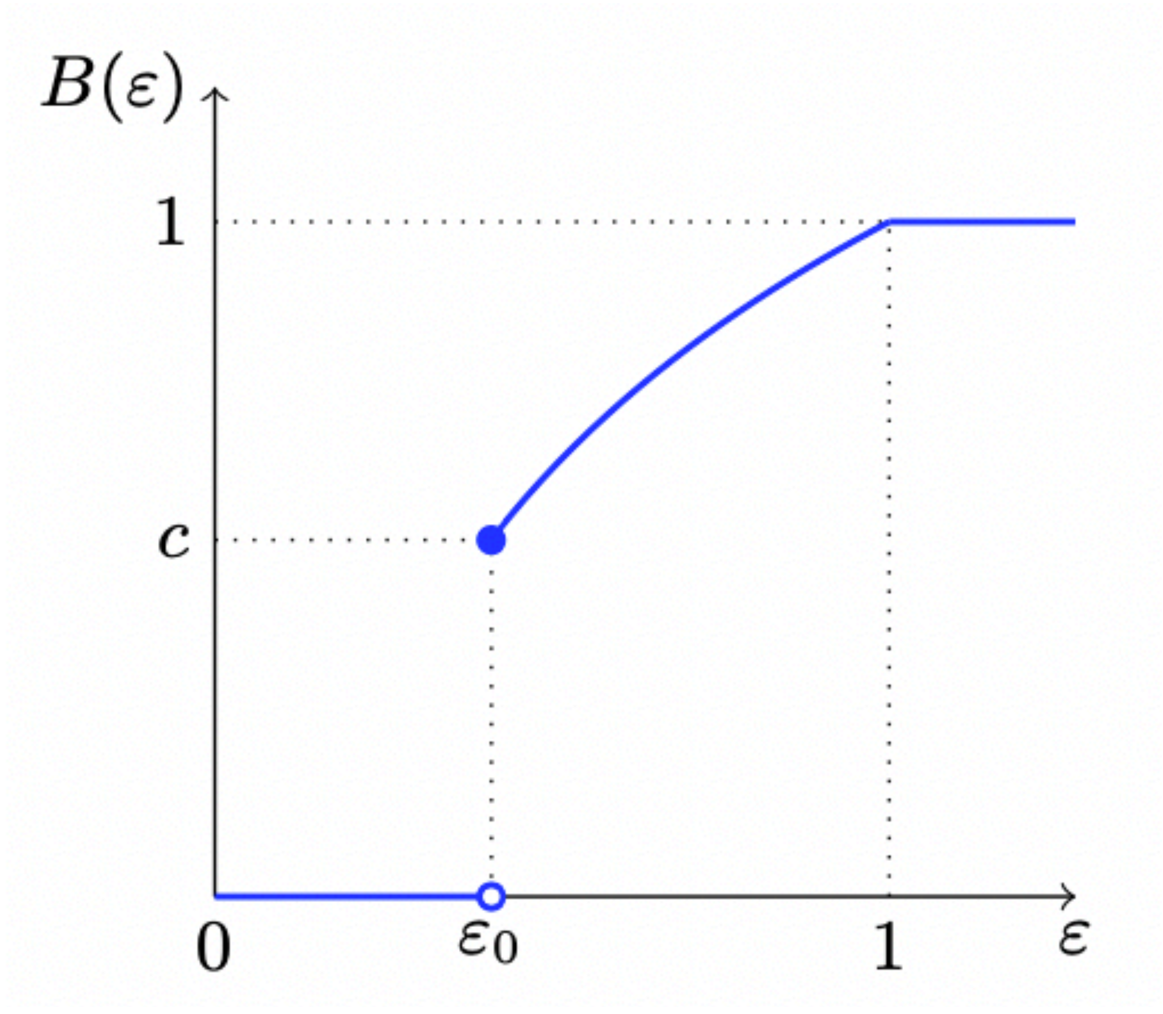

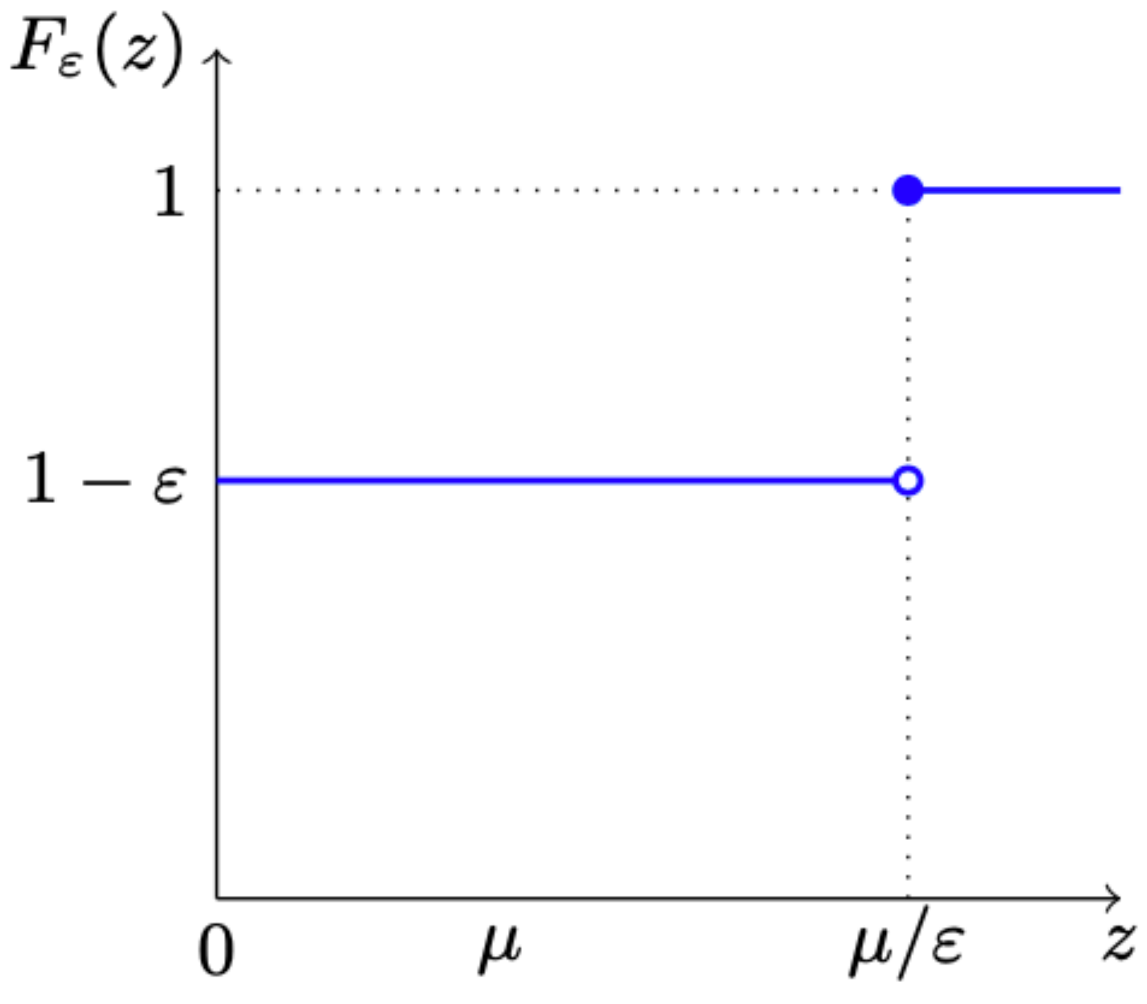

Lower bound - Construction of the mixture

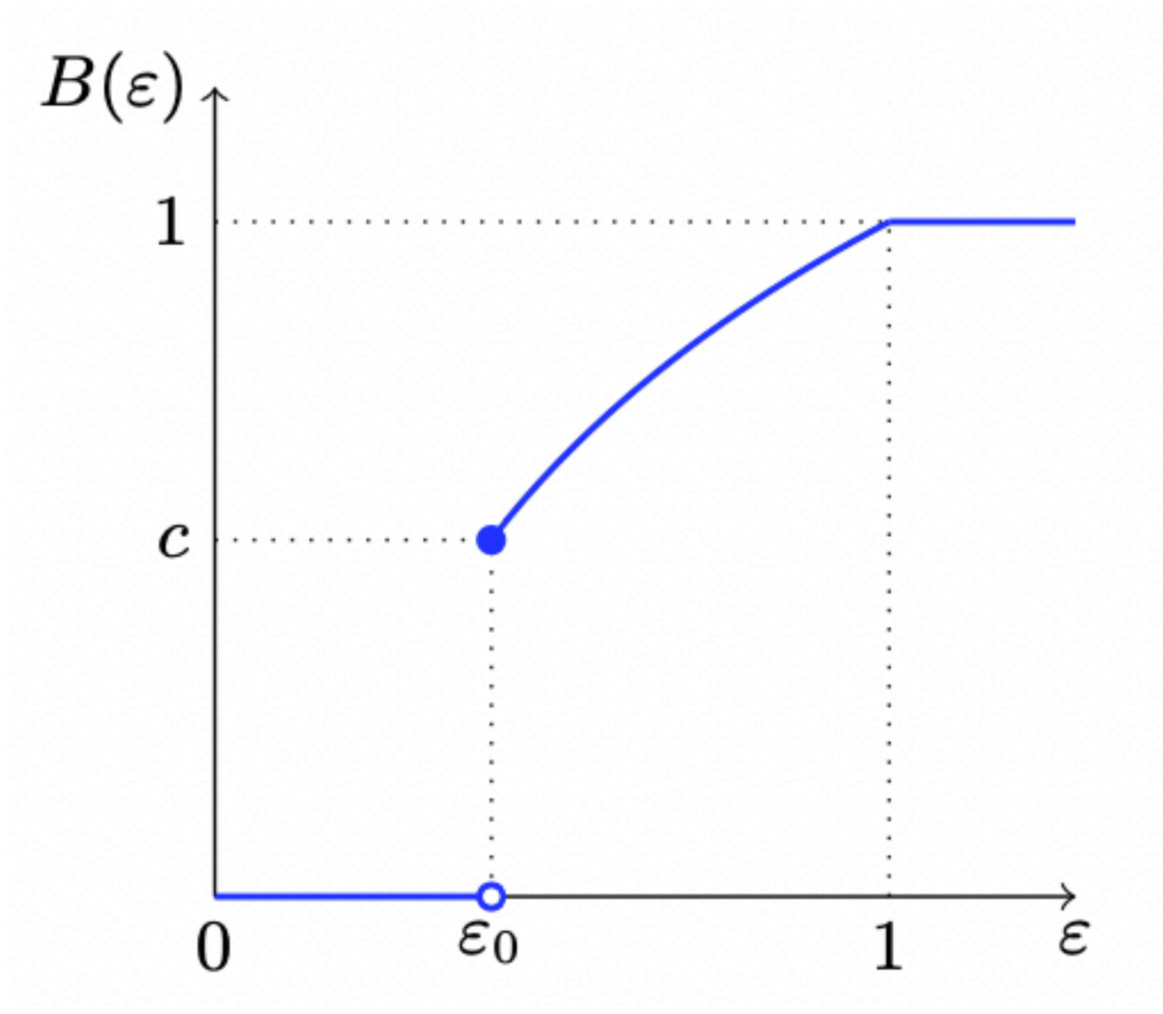

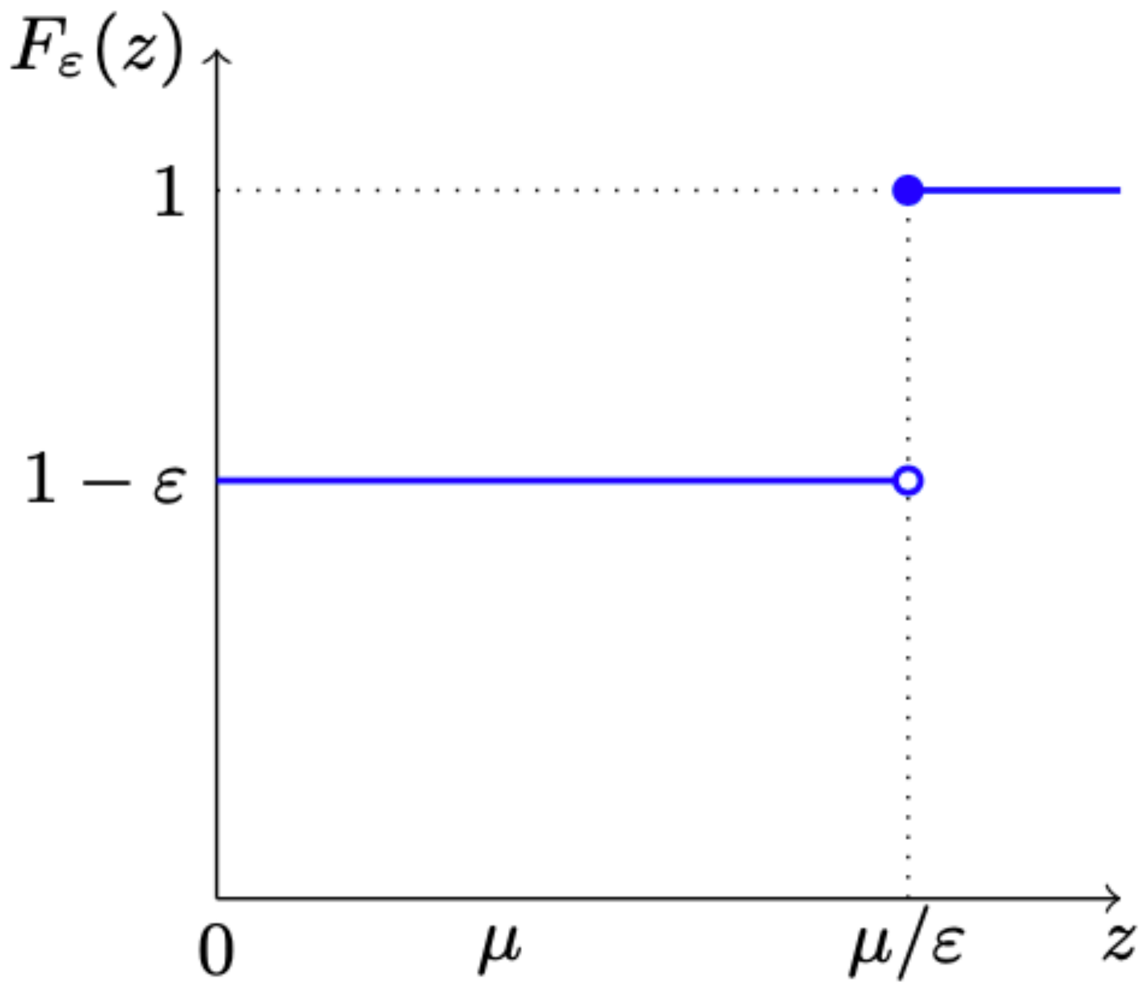

$$G(z) = \mathbb{E}_{\varepsilon \sim B}[F_\epsilon](z)$$

Lower bound - Construction of the mixture

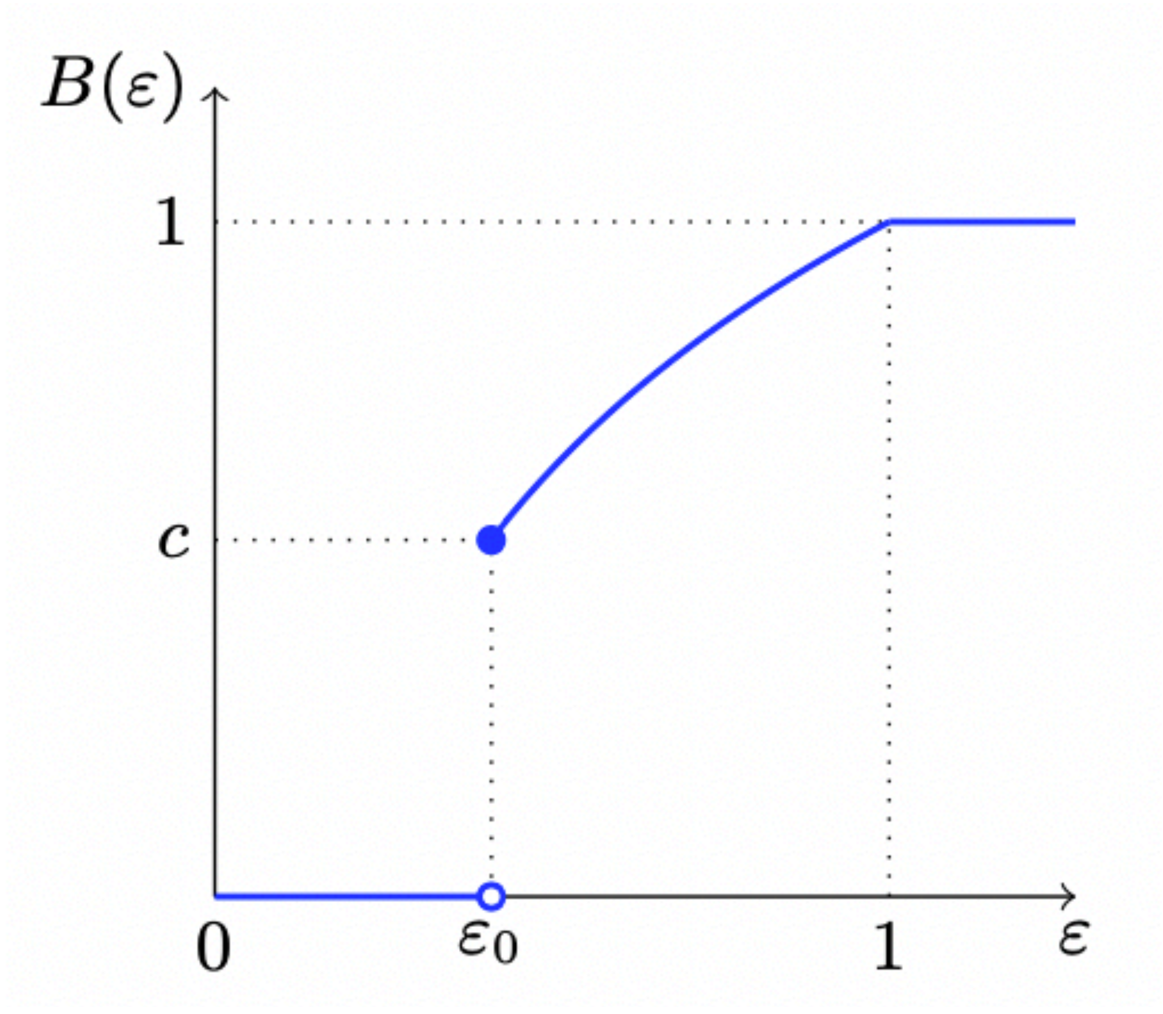

*

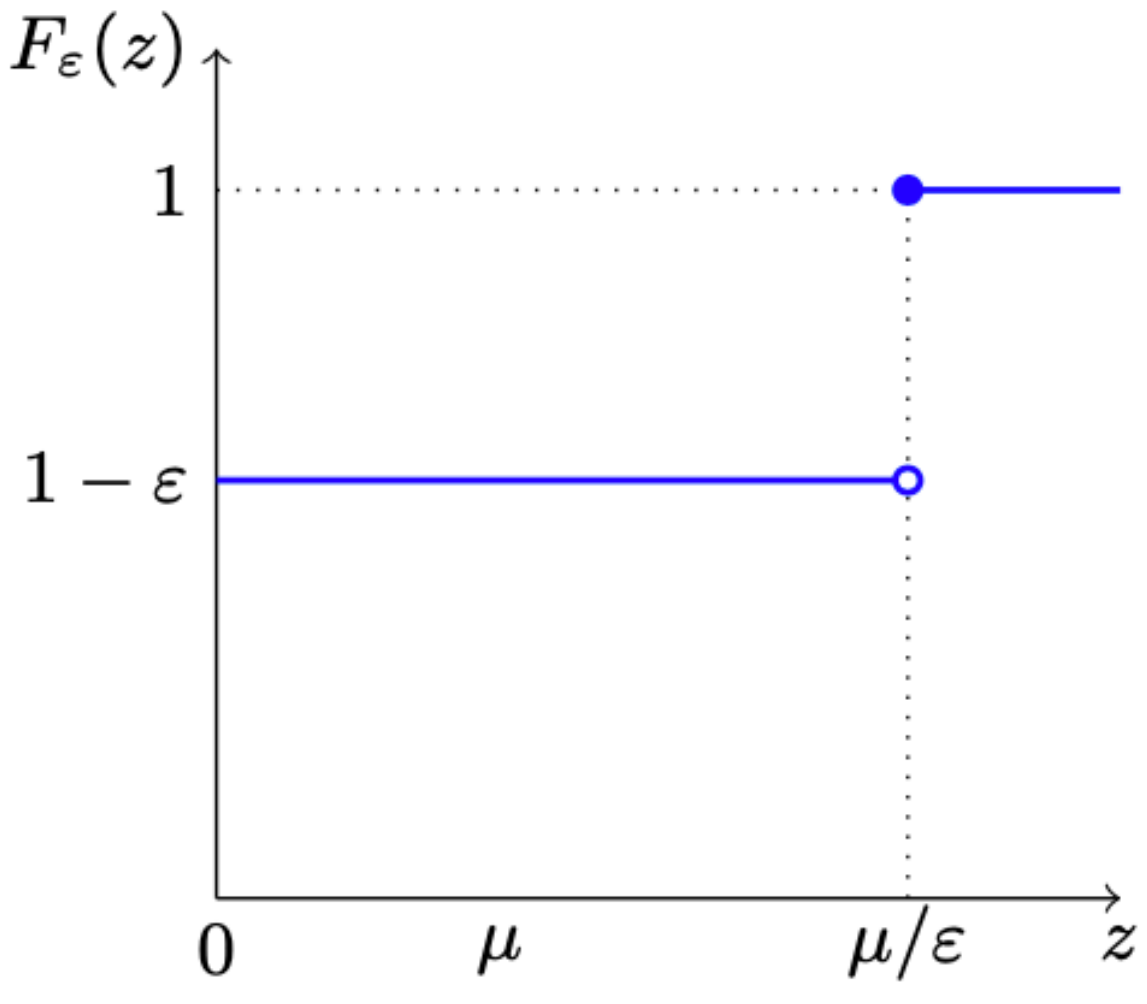

=

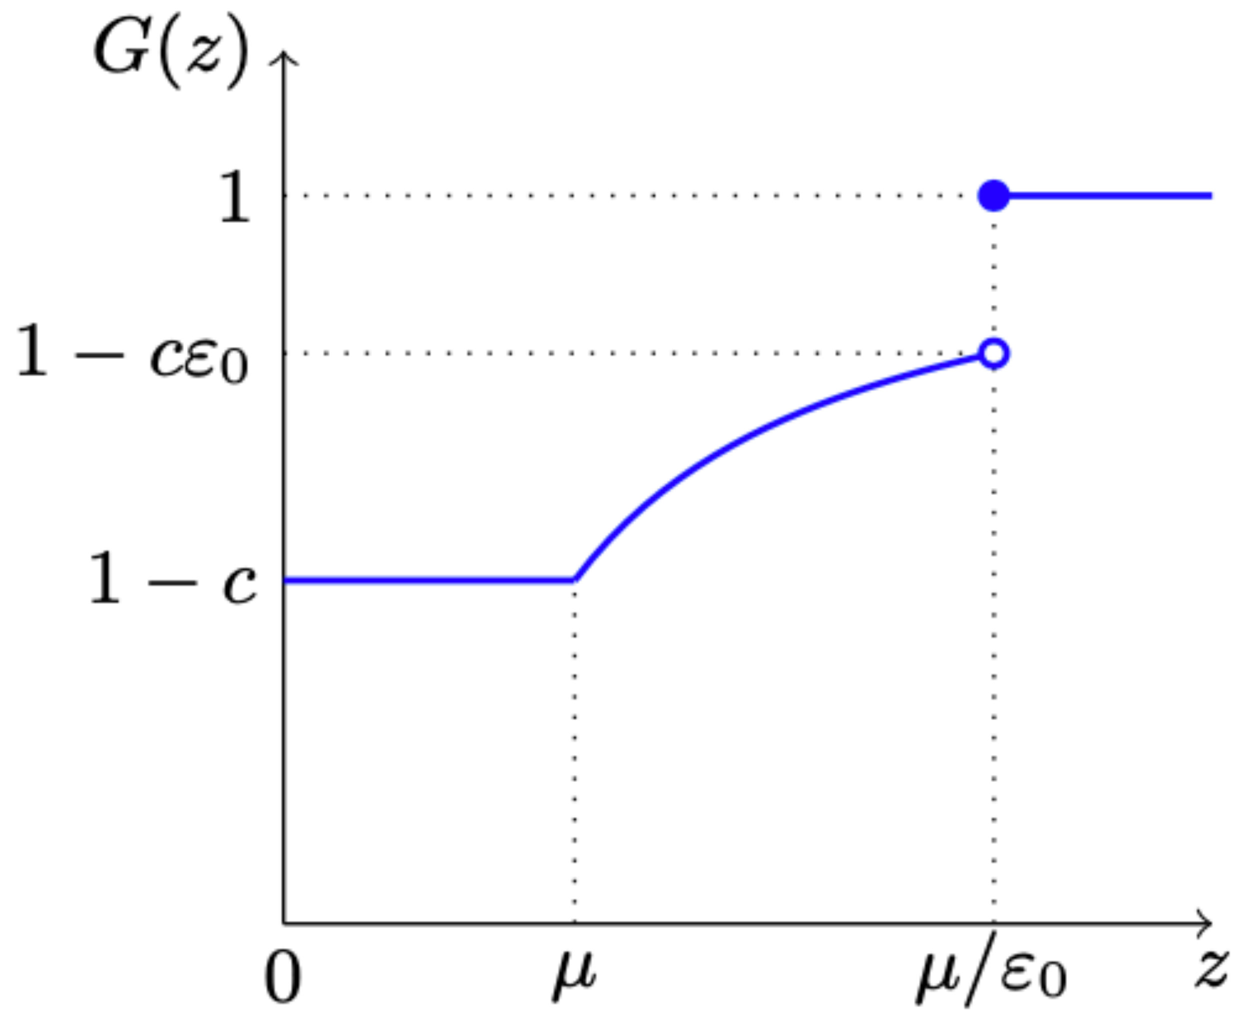

$$G(z) = \mathbb{E}_{\varepsilon \sim B}[F_\epsilon](z)$$

Lower bound - Construction of the mixture

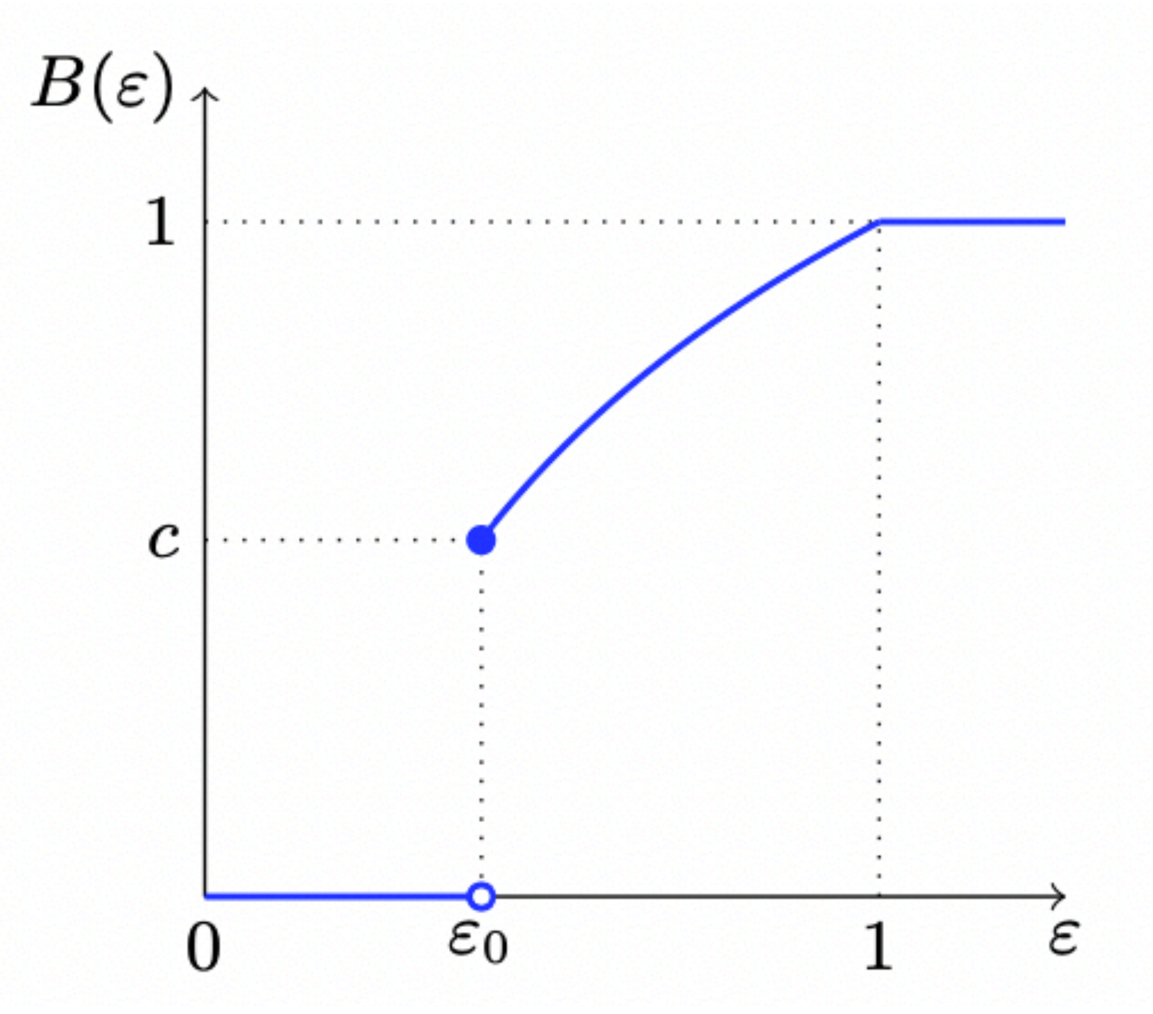

*

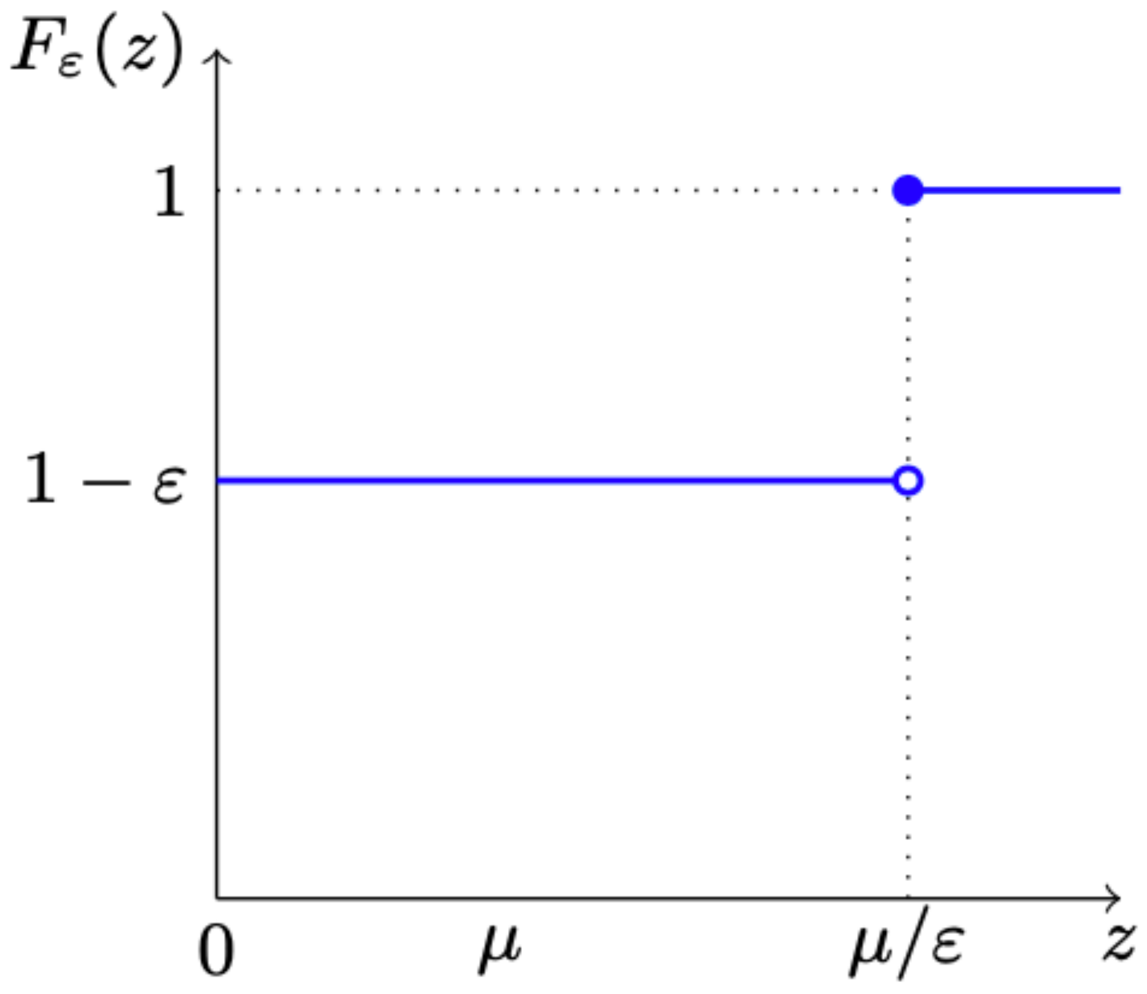

=

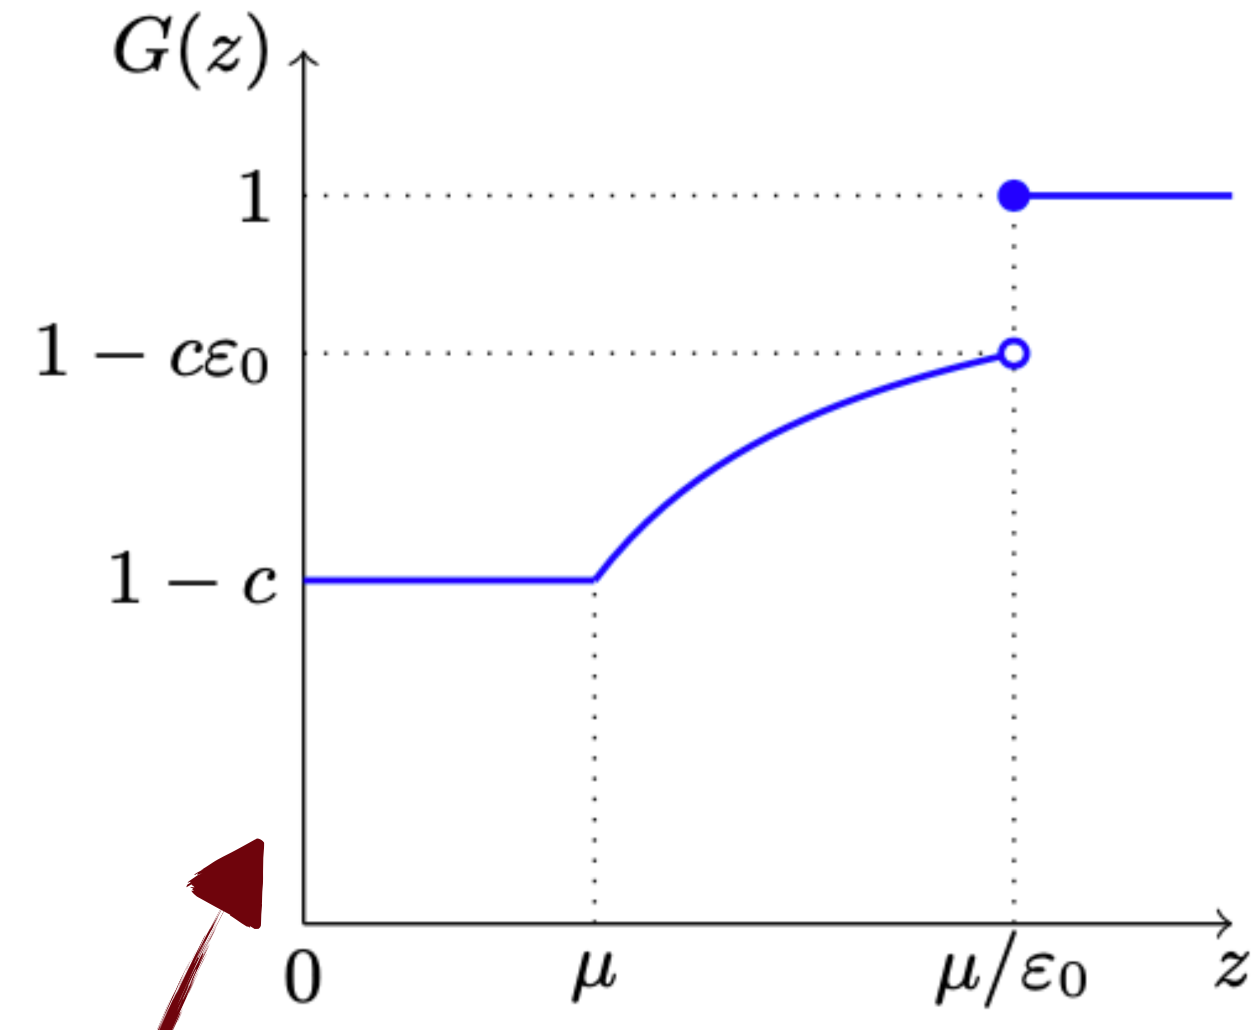

$G(z) = \mathbb{E}_{\varepsilon\sim B}[F_\epsilon](z)$

Truncated equal revenue

Open questions

Open questions

- Multiple bidders - multiple items (generalize ours & [Azar et al. SODA ’13])

Open questions

- Multiple bidders - multiple items (generalize ours & [Azar et al. SODA ’13])
- Intervals of confidence

Open questions

- Multiple bidders - multiple items (generalize ours & [Azar et al. SODA '13])
- Intervals of confidence
- Broader classes of valuations

Open questions

- Multiple bidders - multiple items (generalize ours & [Azar et al. SODA ’13])
- Intervals of confidence
- Broader classes of valuations
- Higher-order moments - the “moment complexity”