Gaussian Sampling over the Integers: Efficient, Generic, Constant-Time

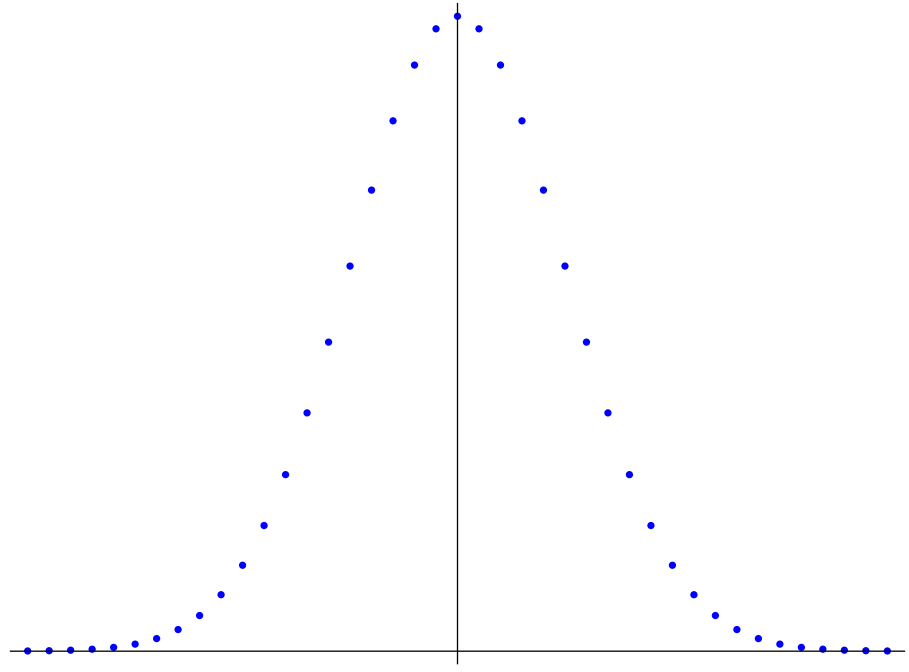

Daniele Micciancio

UCSD

Michael Walter

UCSD

eprint.iacr.org/2017/259

Definitions

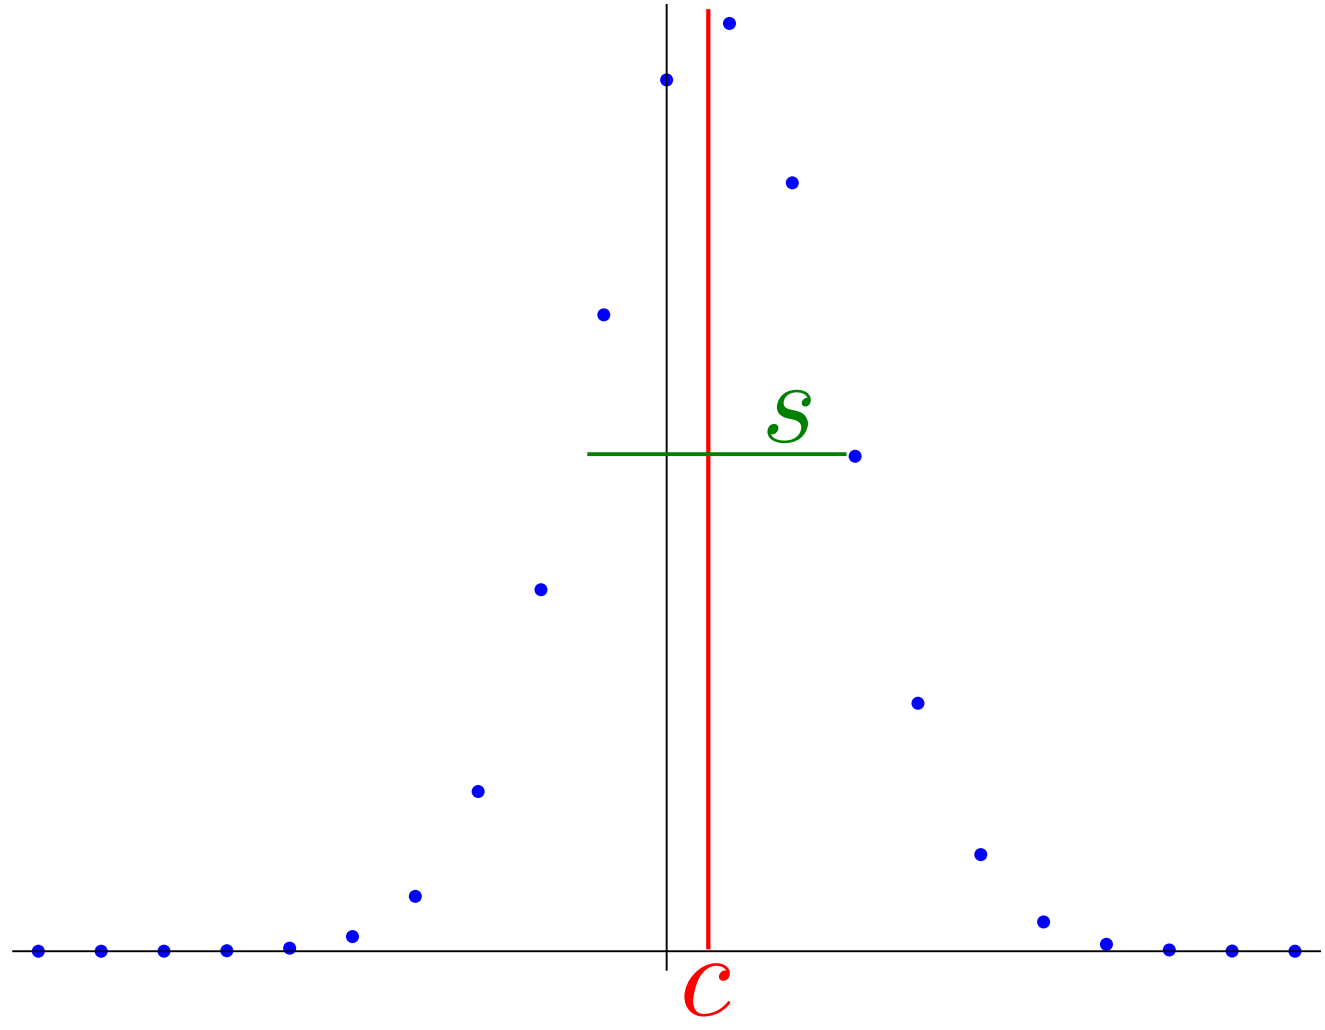

Definitions

$$\rho(x) = e^{-\pi x^2}$$

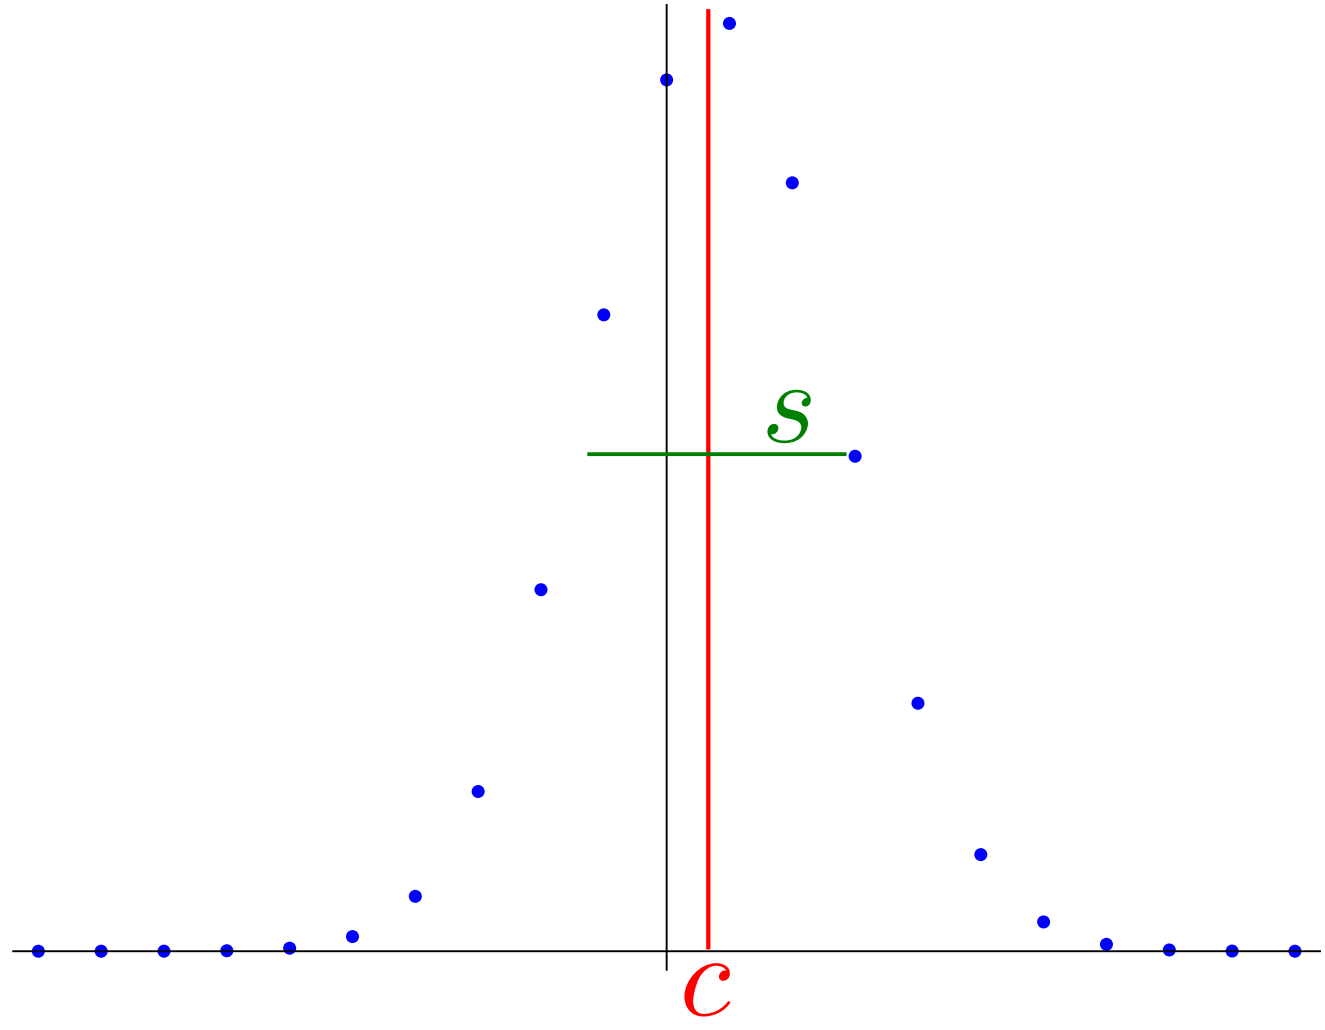

Definitions

$$\rho(x) = e^{-\pi x^2}$$

$$\rho_{c,s}(x) = \rho((x-c)/s)$$

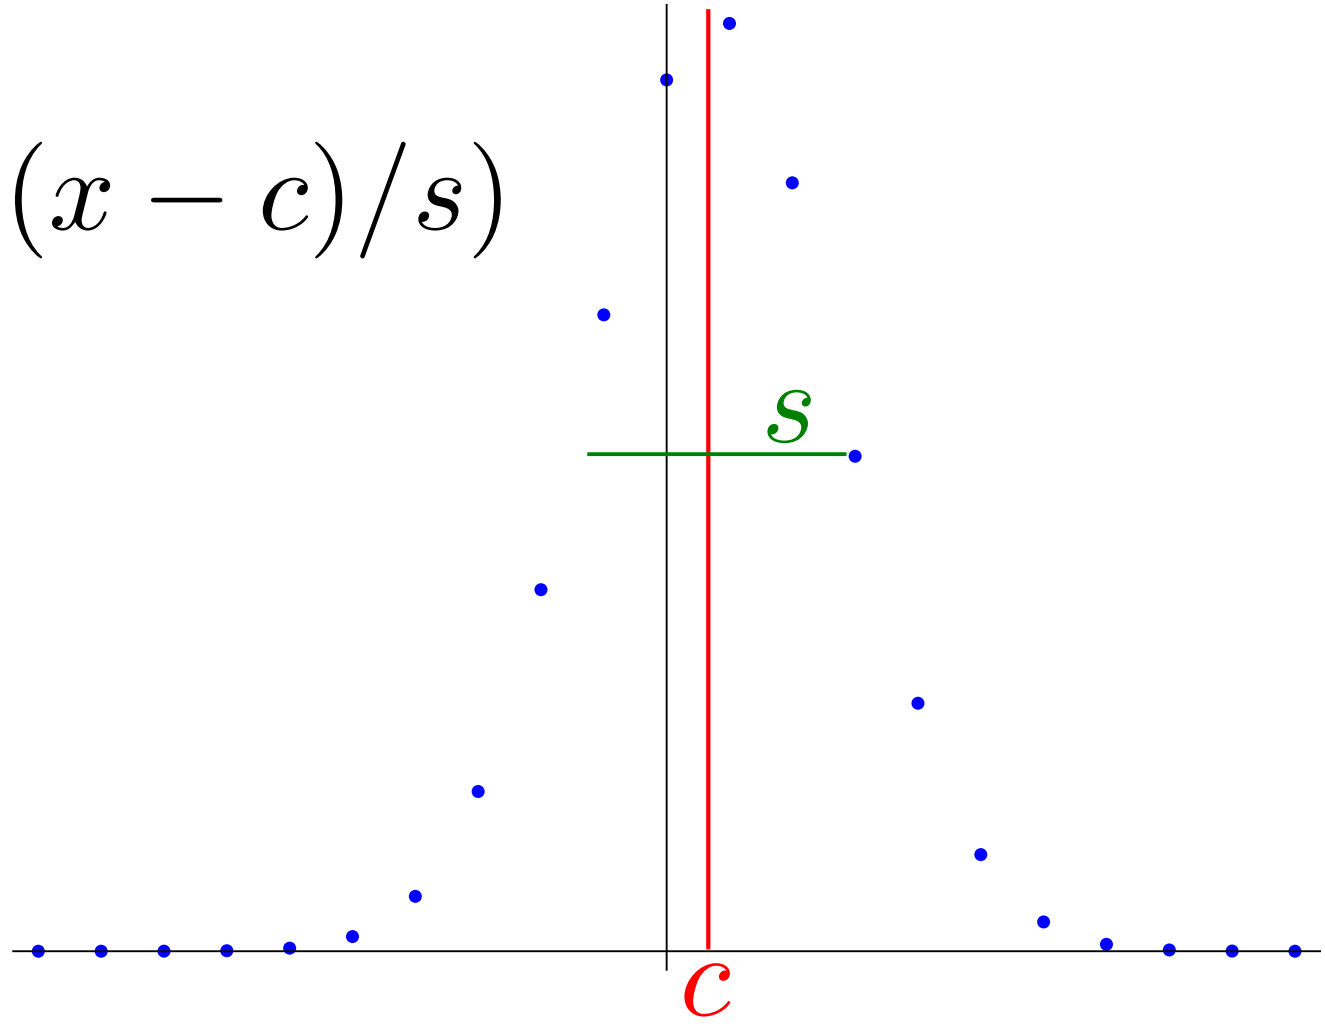

Definitions

$$\rho(x) = e^{-\pi x^2}$$

$$\rho_{c,s}(x) = \rho((x-c)/s)$$

$$\mathcal{D}_{\mathbb{Z},c,s}(x) \sim \rho_{c,s}(x)$$

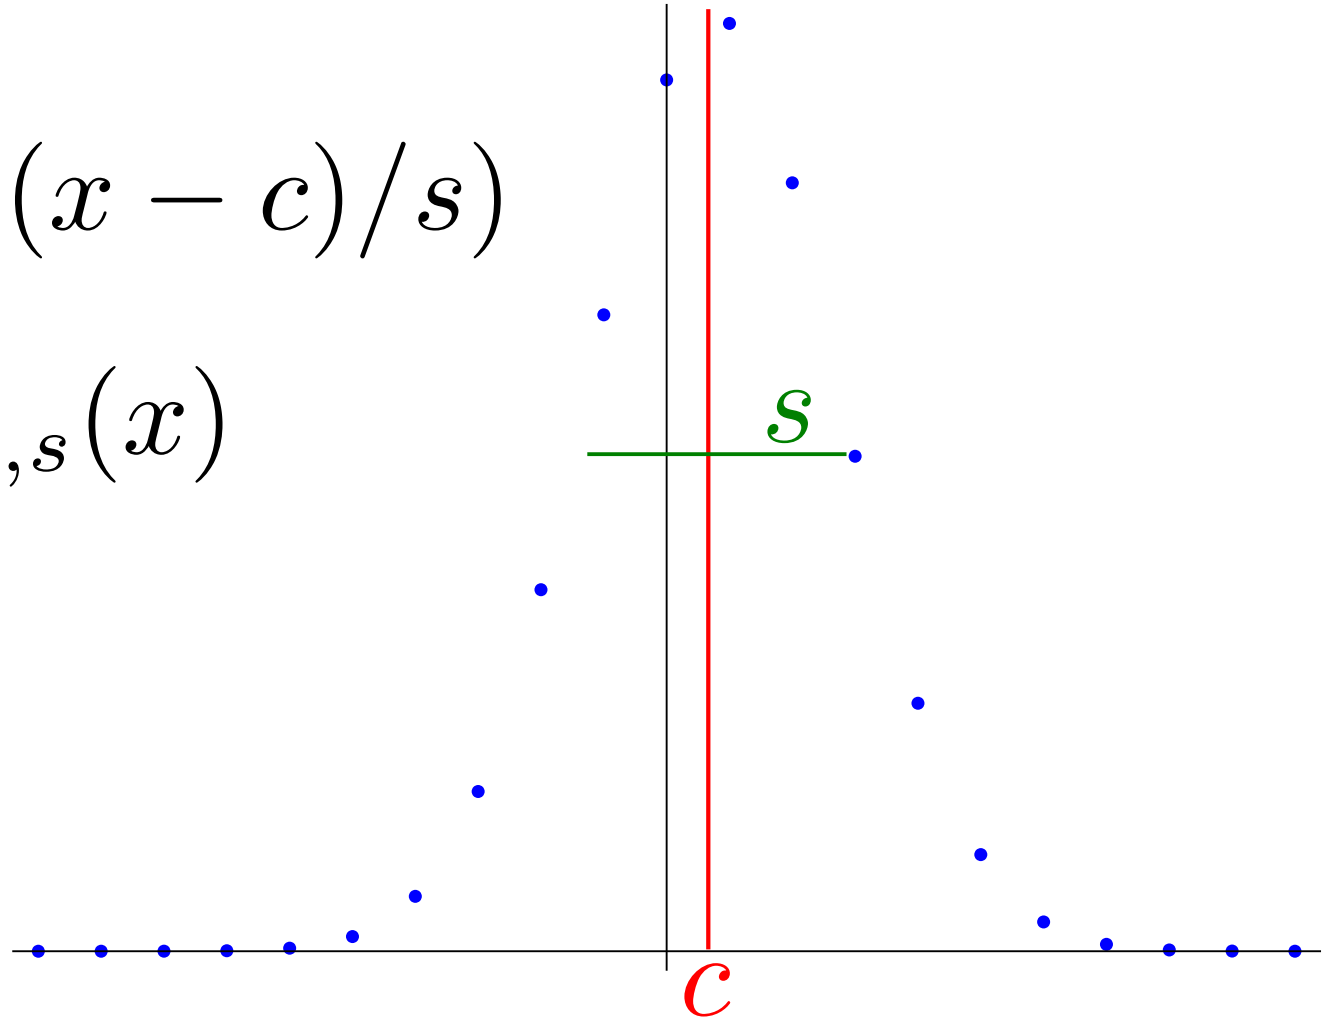

Applications

Sample from $\mathcal{D}_{c,s}$ where

c, s

Application:

Algorithms:

Applications

Sample from $\mathcal{D}_{c,s}$ where

c, s fixed

Application:

Algorithms:

Applications

Sample from $\mathcal{D}_{c,s}$ where

c, s fixed

Application: LWE based schemes $(c = 0)$

Algorithms:

Applications

Sample from $\mathcal{D}_{c,s}$ where

c, s	fixed	variable

Application: LWE based schemes $(c = 0)$

Algorithms:

Applications

Sample from $\mathcal{D}_{c,s}$ where

c, s	fixed	variable
Application:	LWE based schemes $(c = 0)$	Lattice Trapdoors
Algorithms:		

Applications

Sample from $\mathcal{D}_{c,s}$ where

c, s	fixed	variable
Application:	LWE based schemes $(c = 0)$	Lattice Trapdoors
Algorithms:	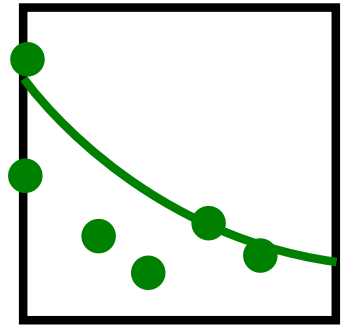	

Applications

Sample from $\mathcal{D}_{c,s}$ where

c, s	fixed	variable
Application:	LWE based schemes $(c = 0)$	Lattice Trapdoors
Algorithms:	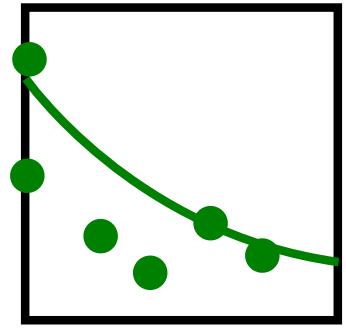	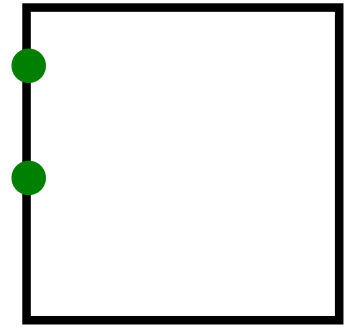

Applications

Sample from $\mathcal{D}_{c,s}$ where

c, s	fixed	variable
Application:	LWE based schemes $(c = 0)$	Lattice Trapdoors
Algorithms:	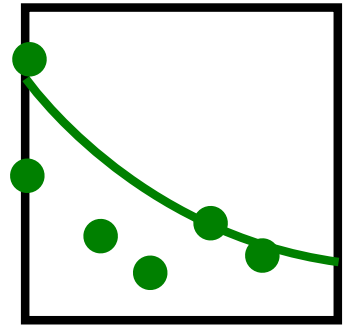	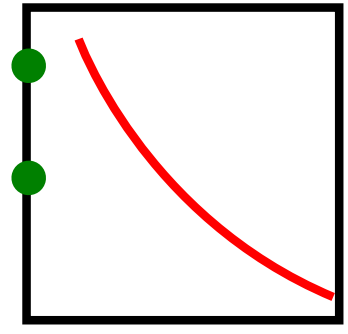

DGS: Rejection Sampling

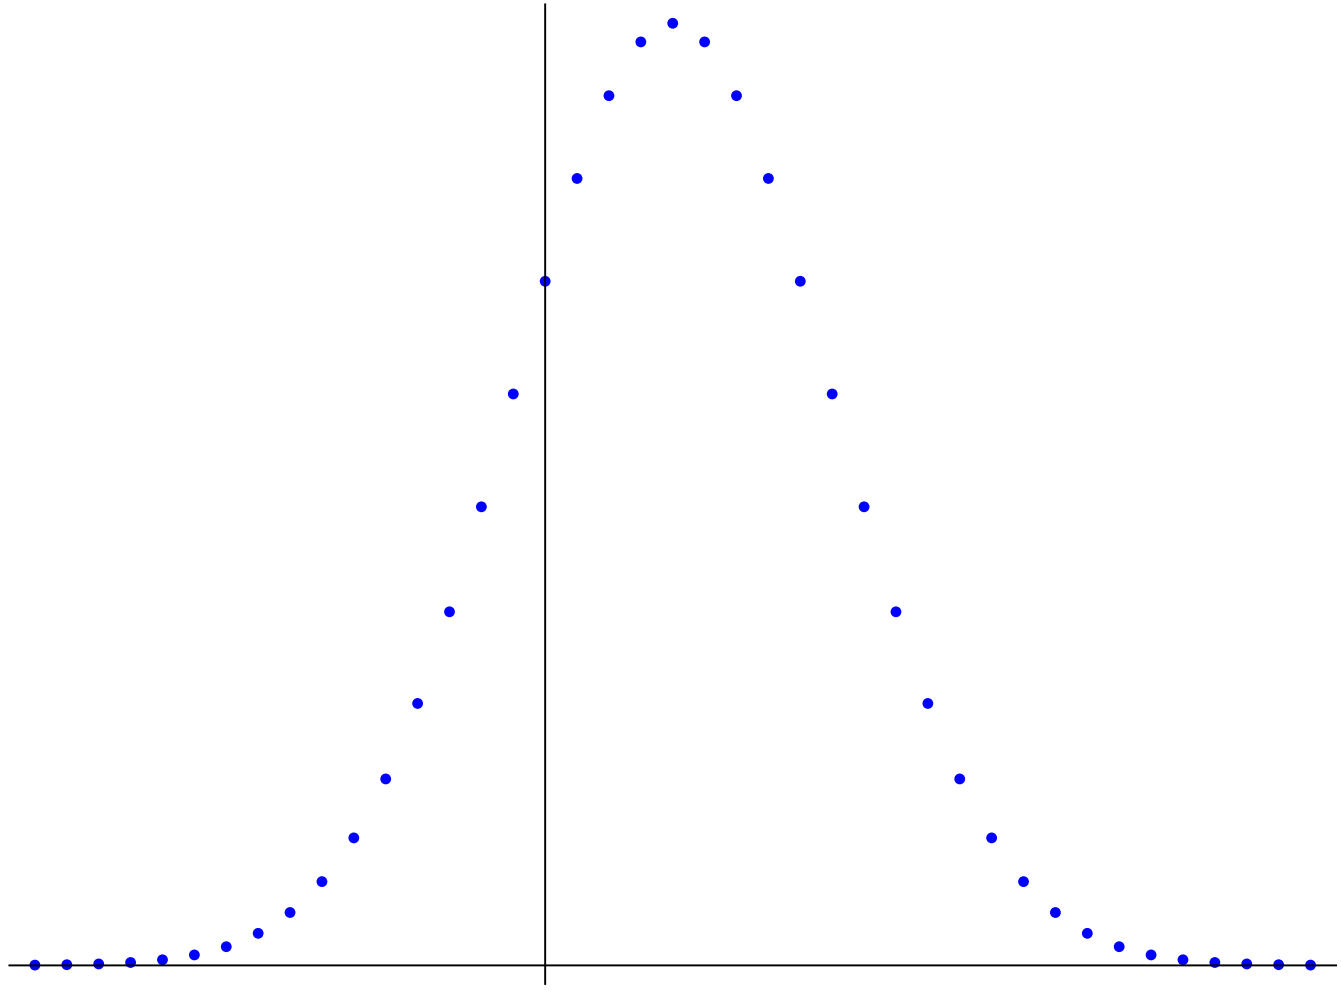

DGS: Rejection Sampling

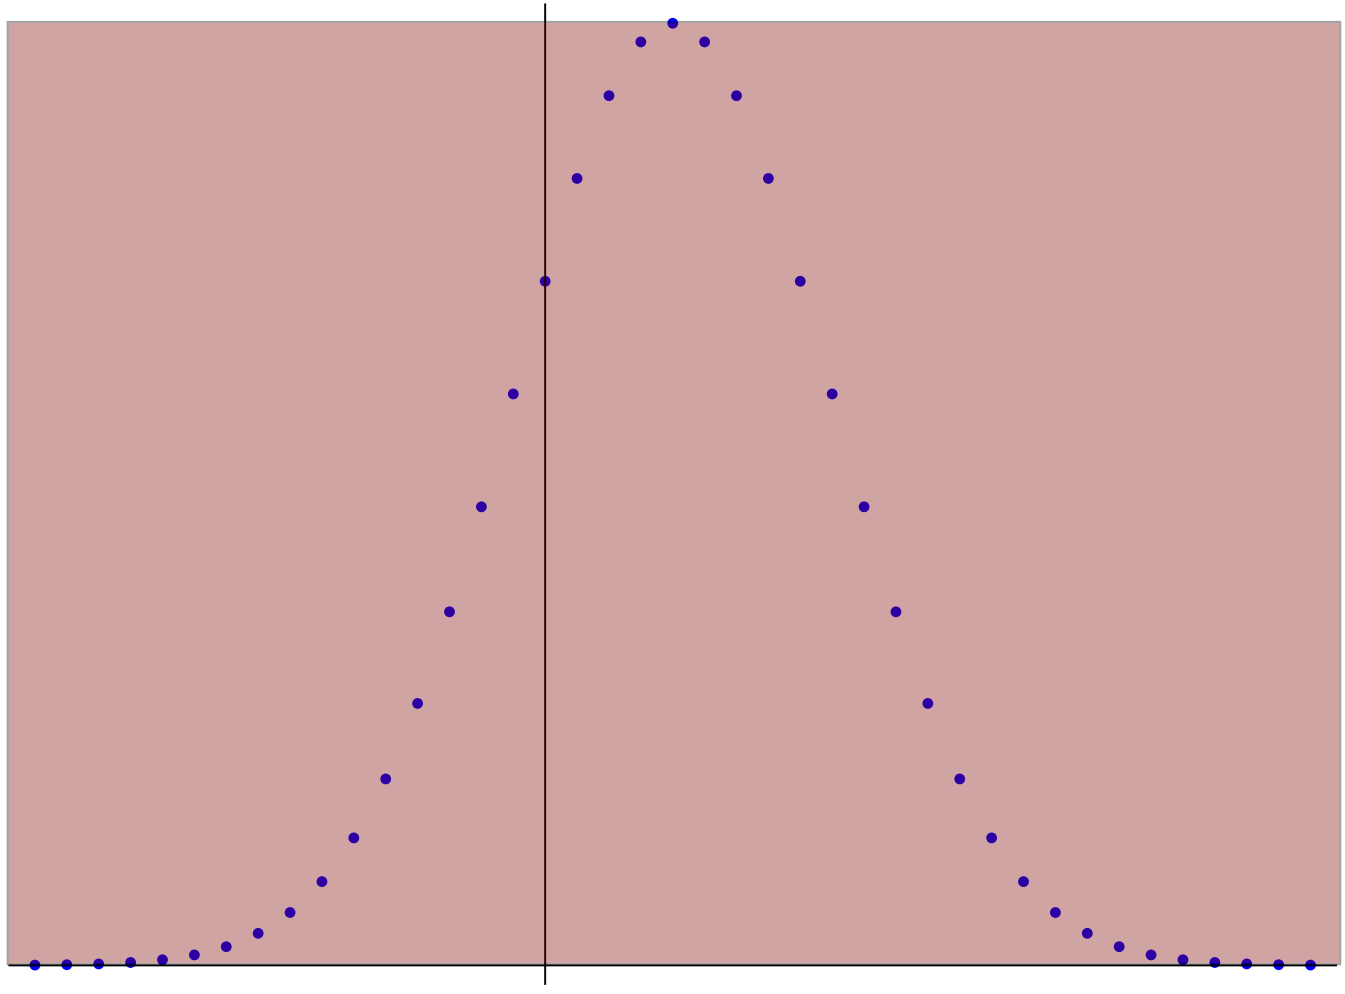

DGS: Rejection Sampling

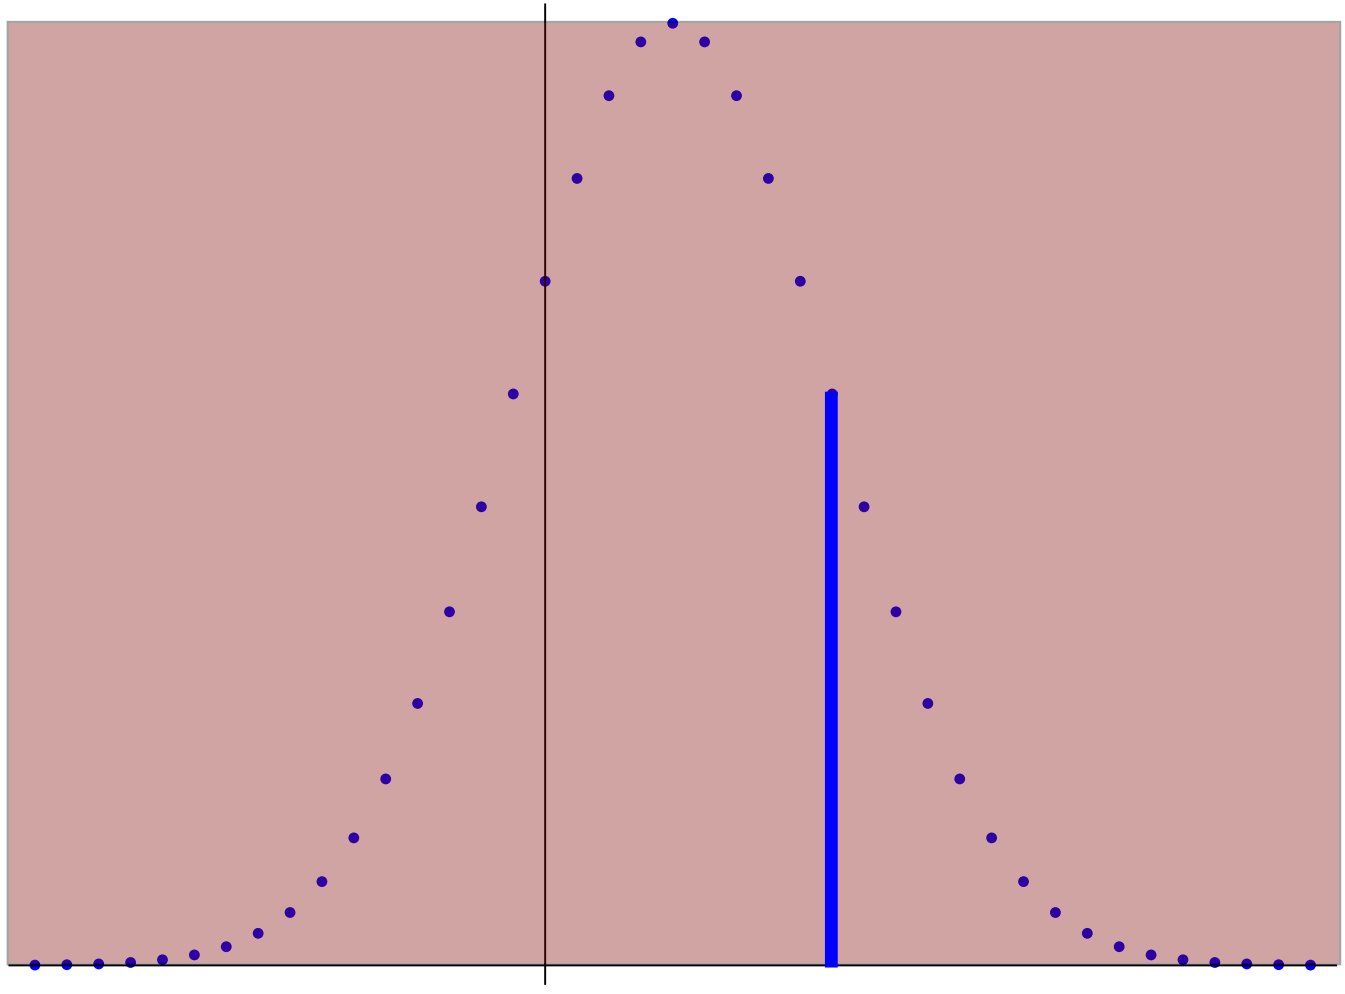

DGS: Rejection Sampling

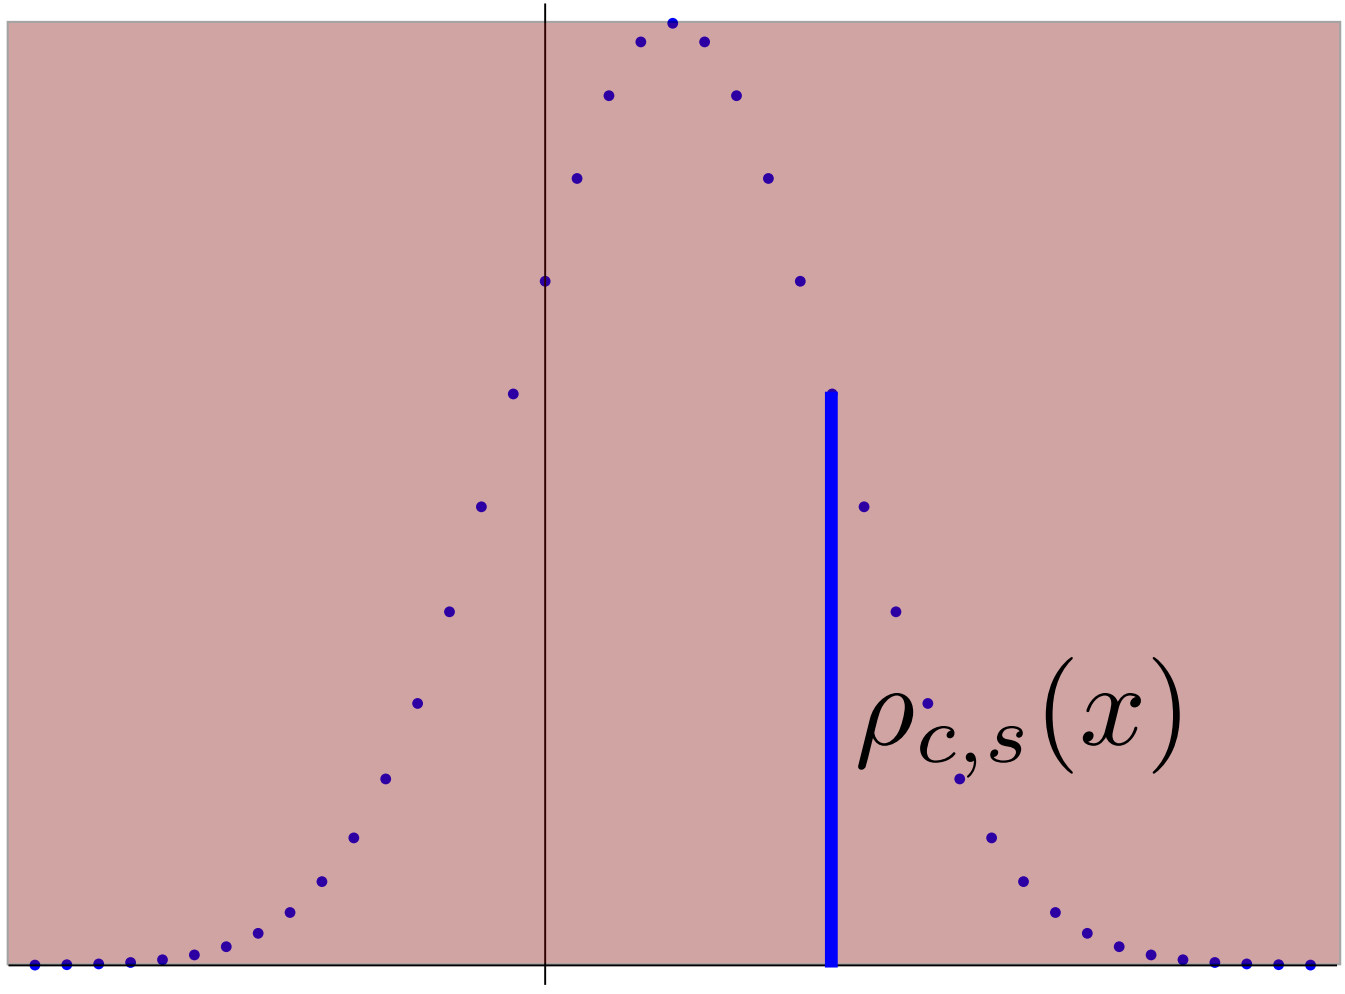

DGS: Inversion Sampling

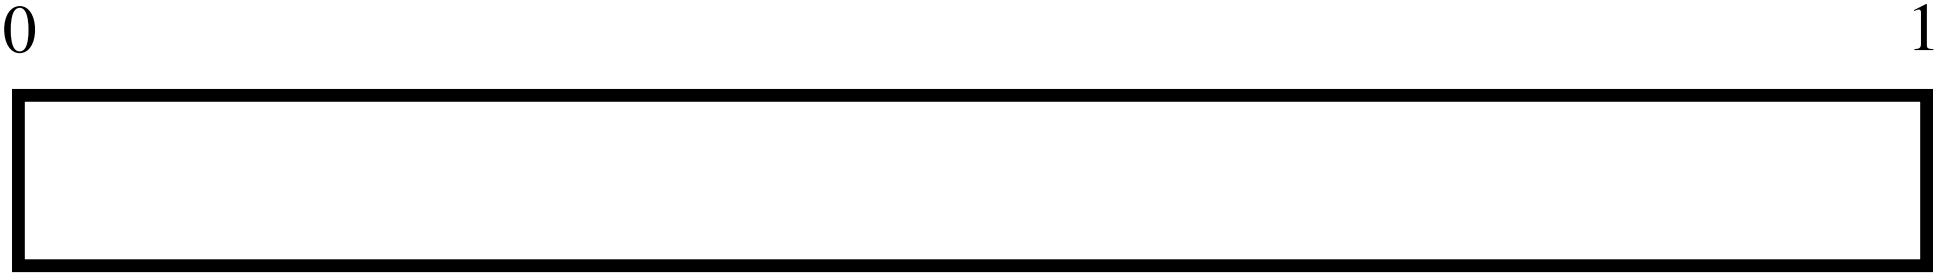

DGS: Inversion Sampling

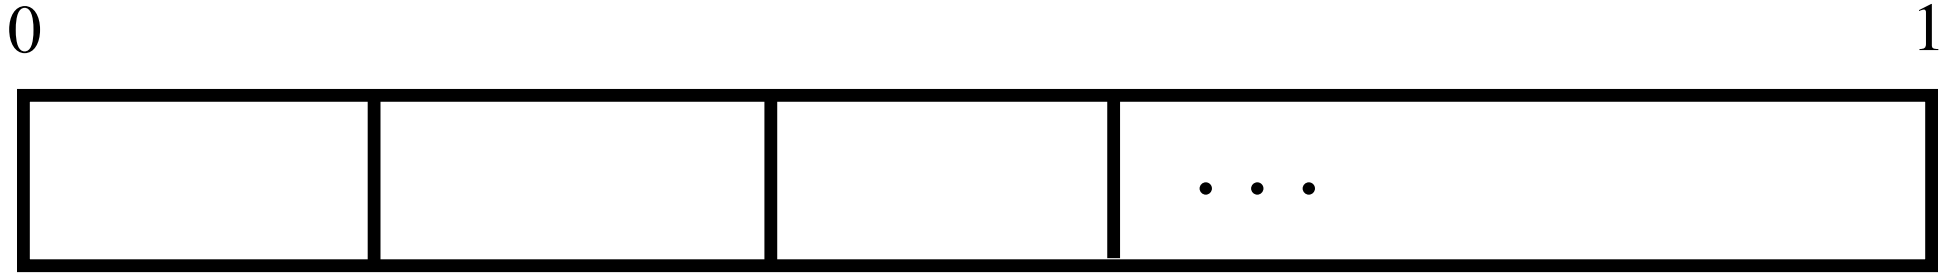

DGS: Inversion Sampling

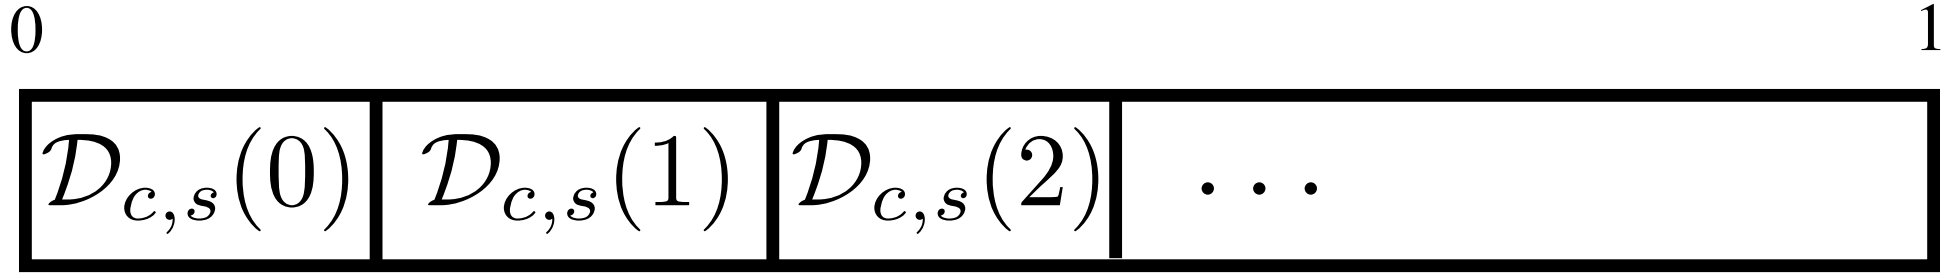

DGS: Inversion Sampling

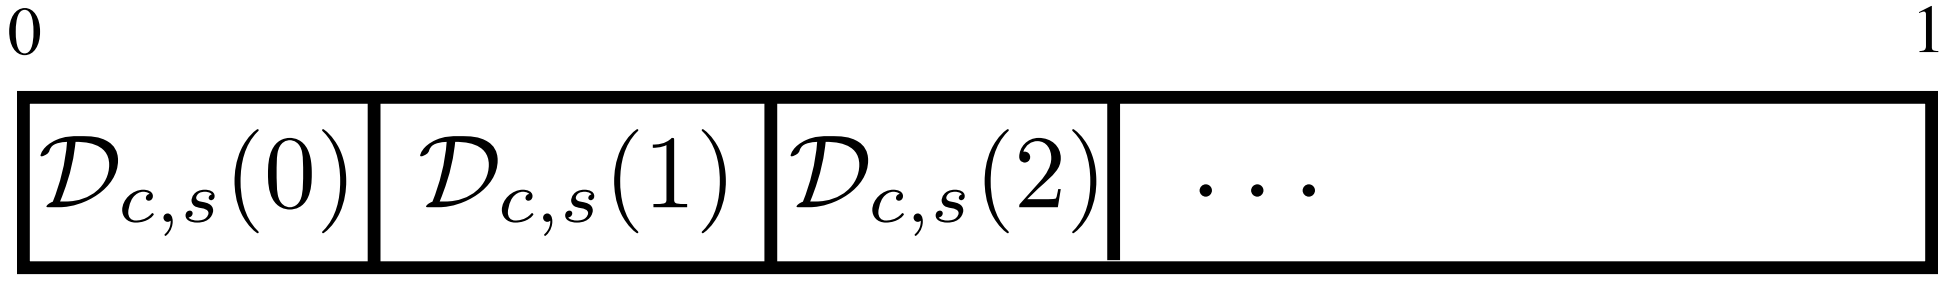

$$r \xleftarrow{\$} [0, 1)$$

DGS: Inversion Sampling

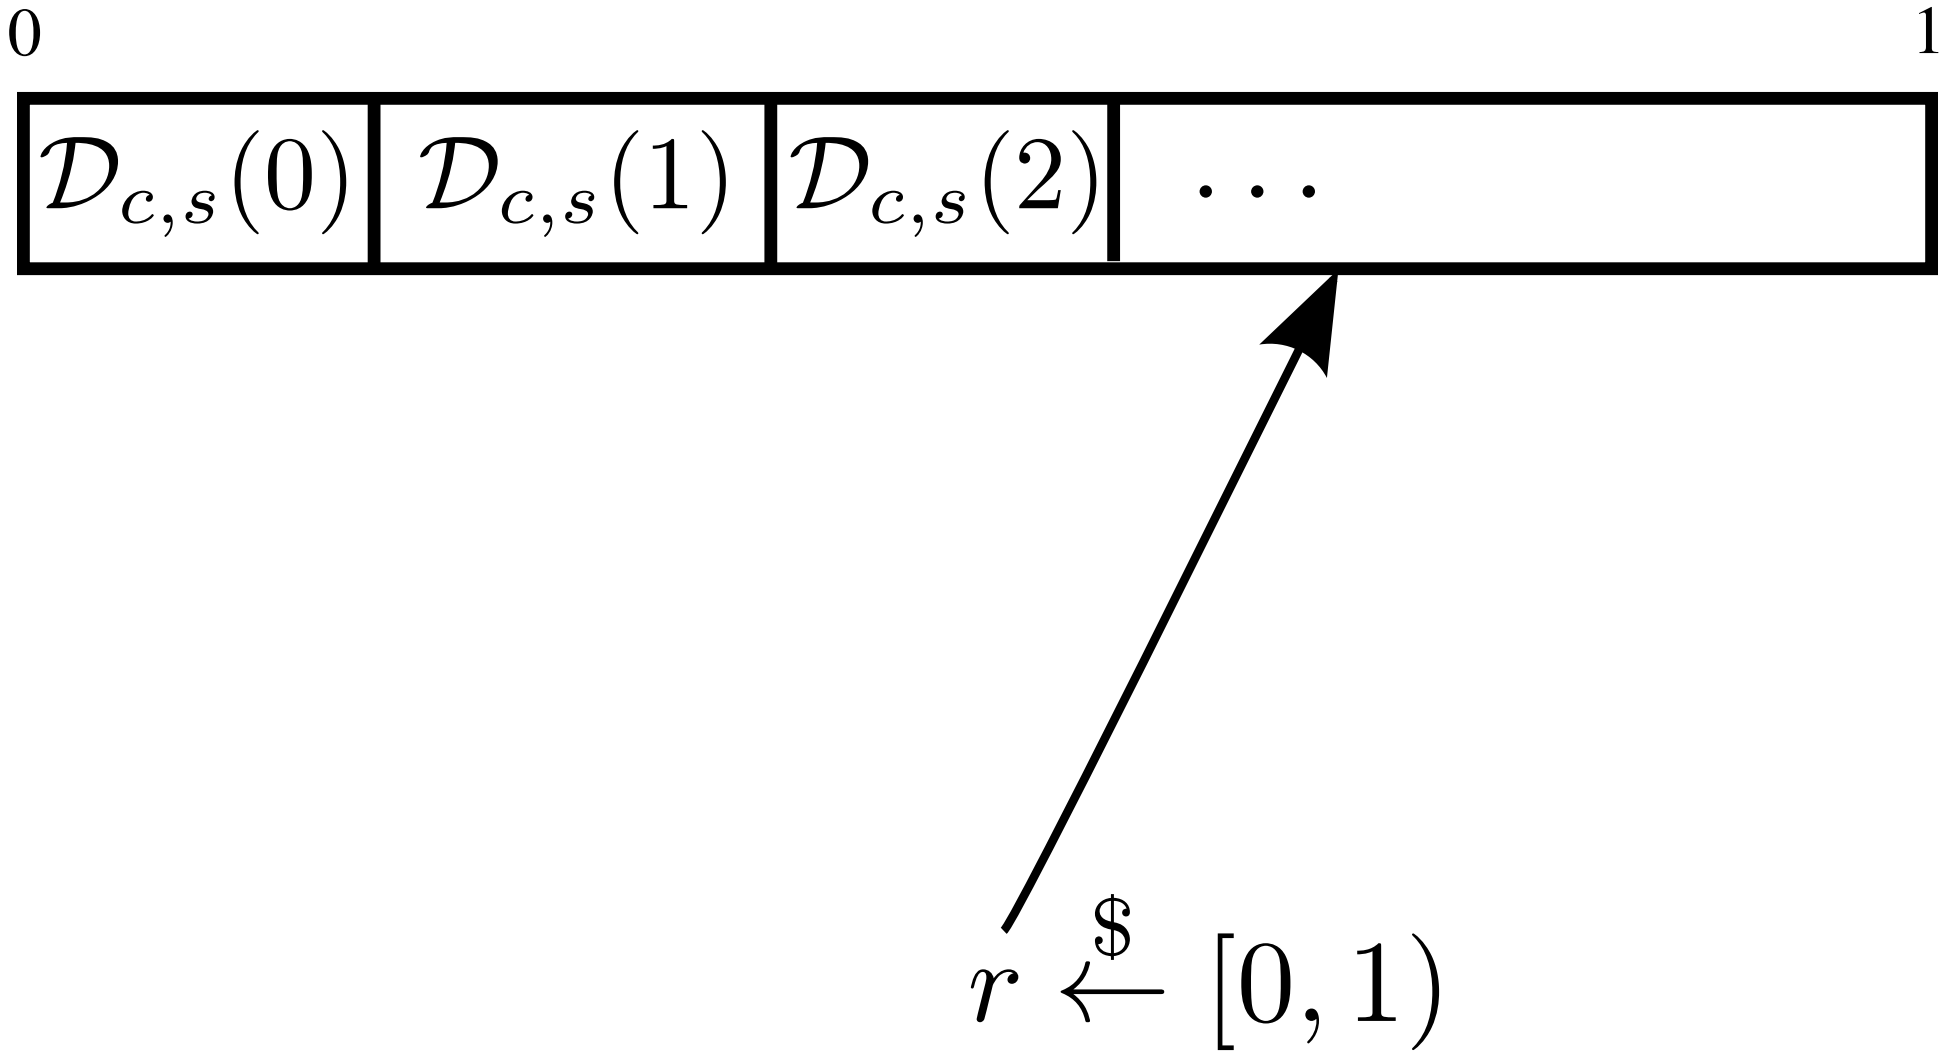

High Level Overview

High Level Overview

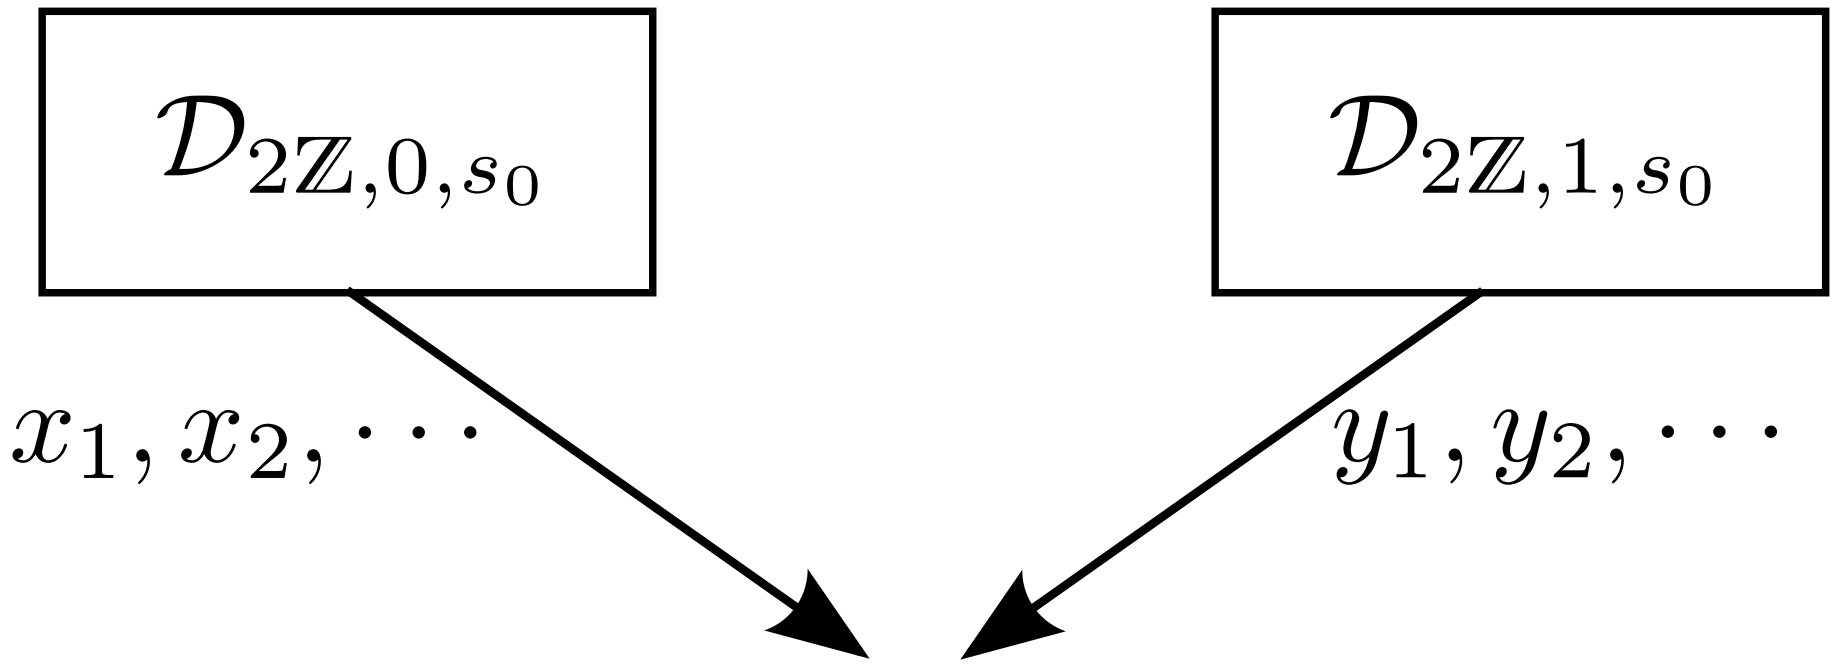

High Level Overview

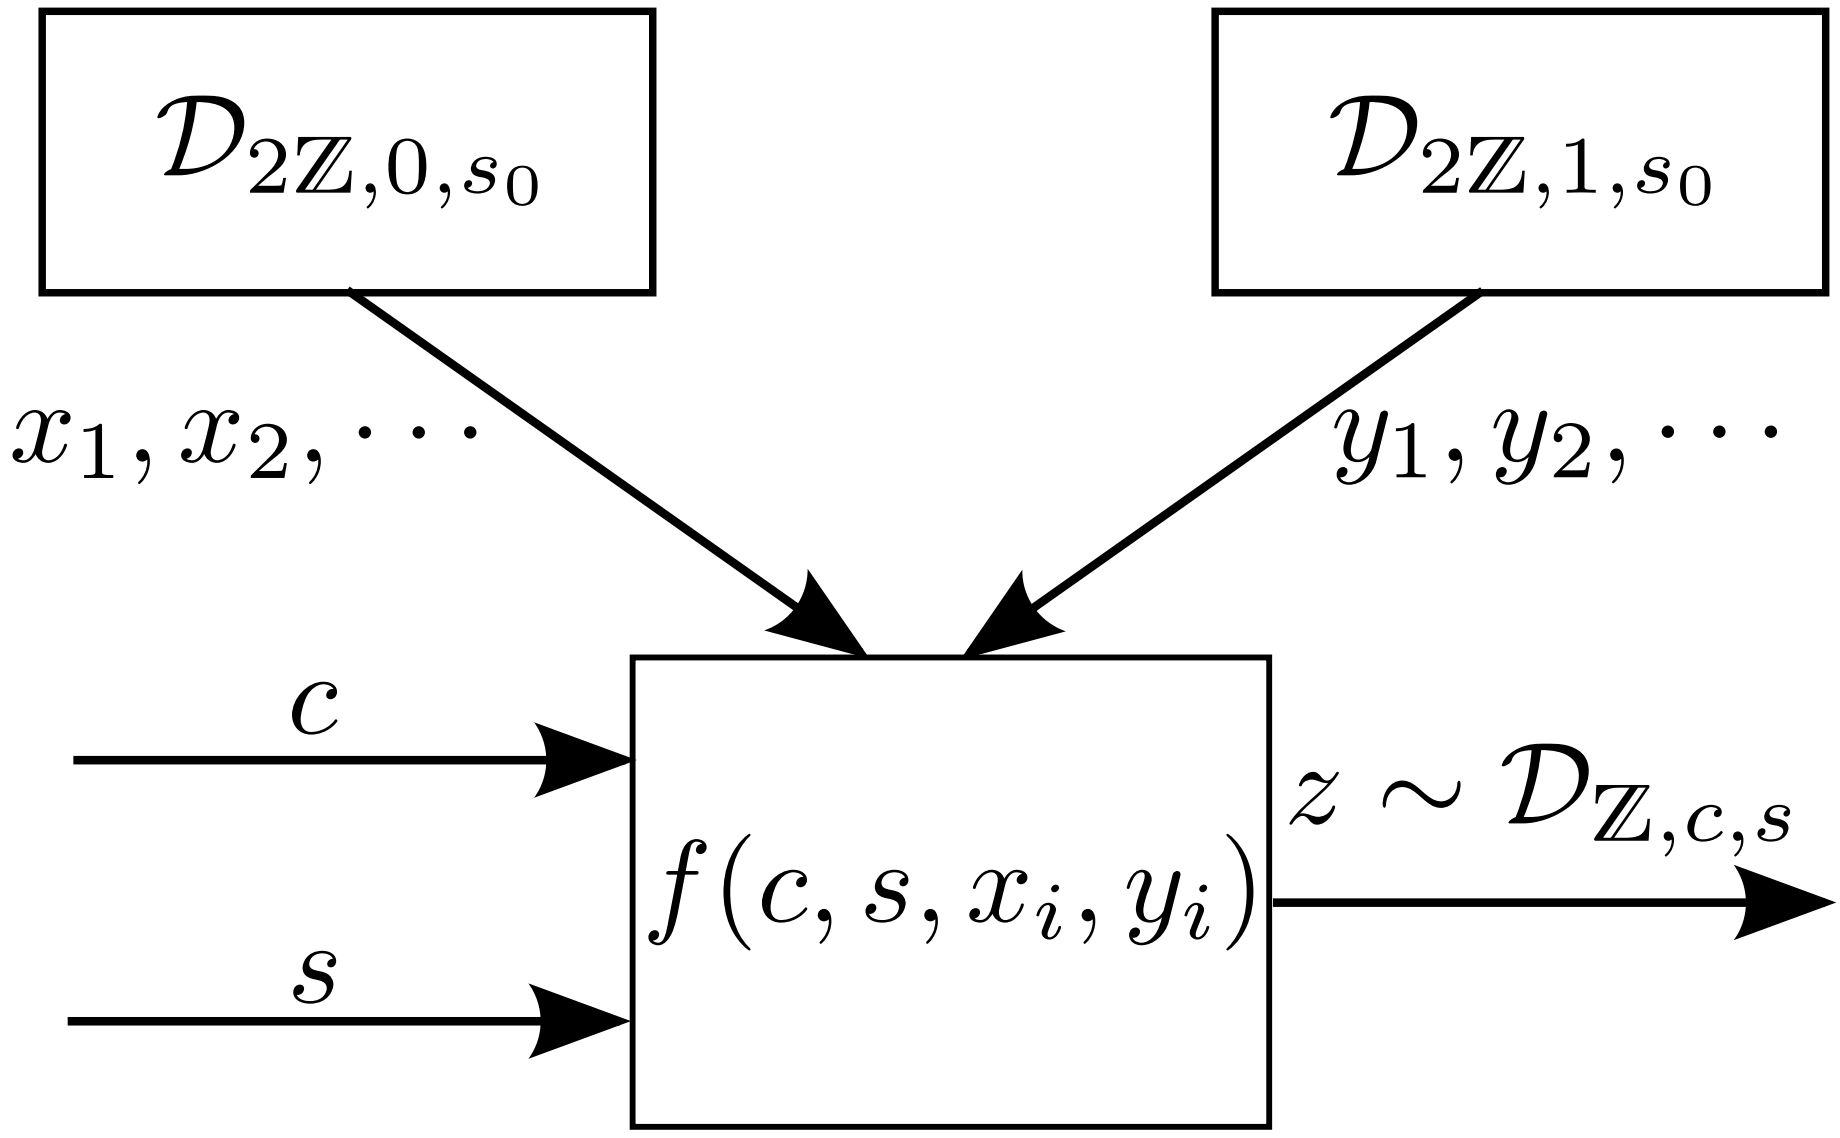

Main Tool: Convolution

Main Tool: Convolution

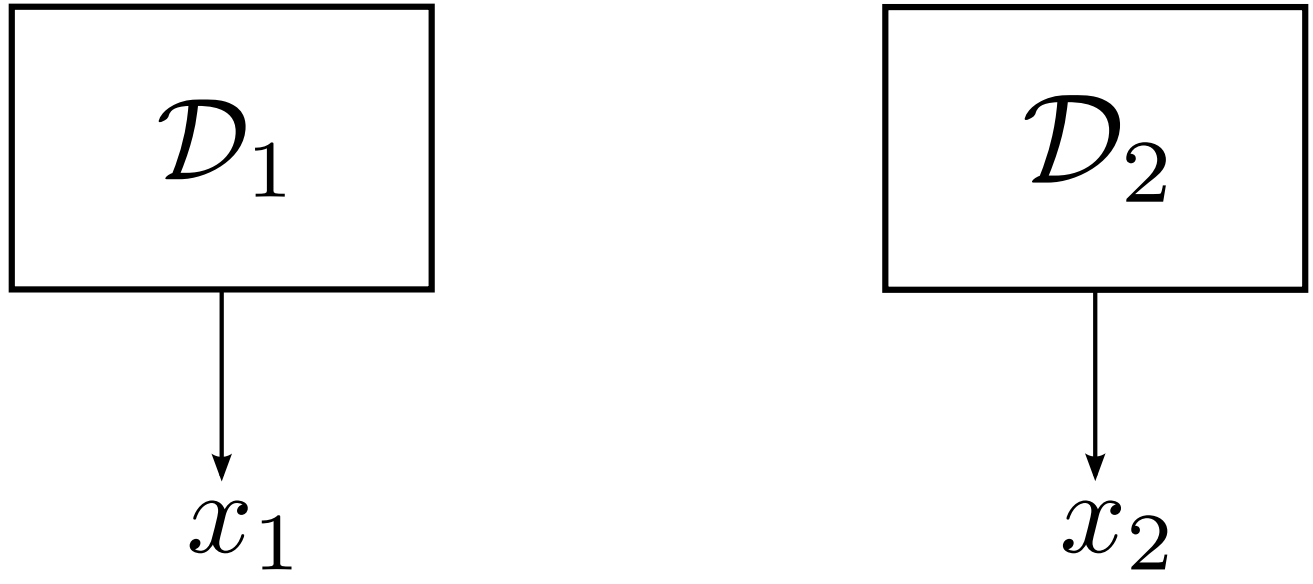

Main Tool: Convolution

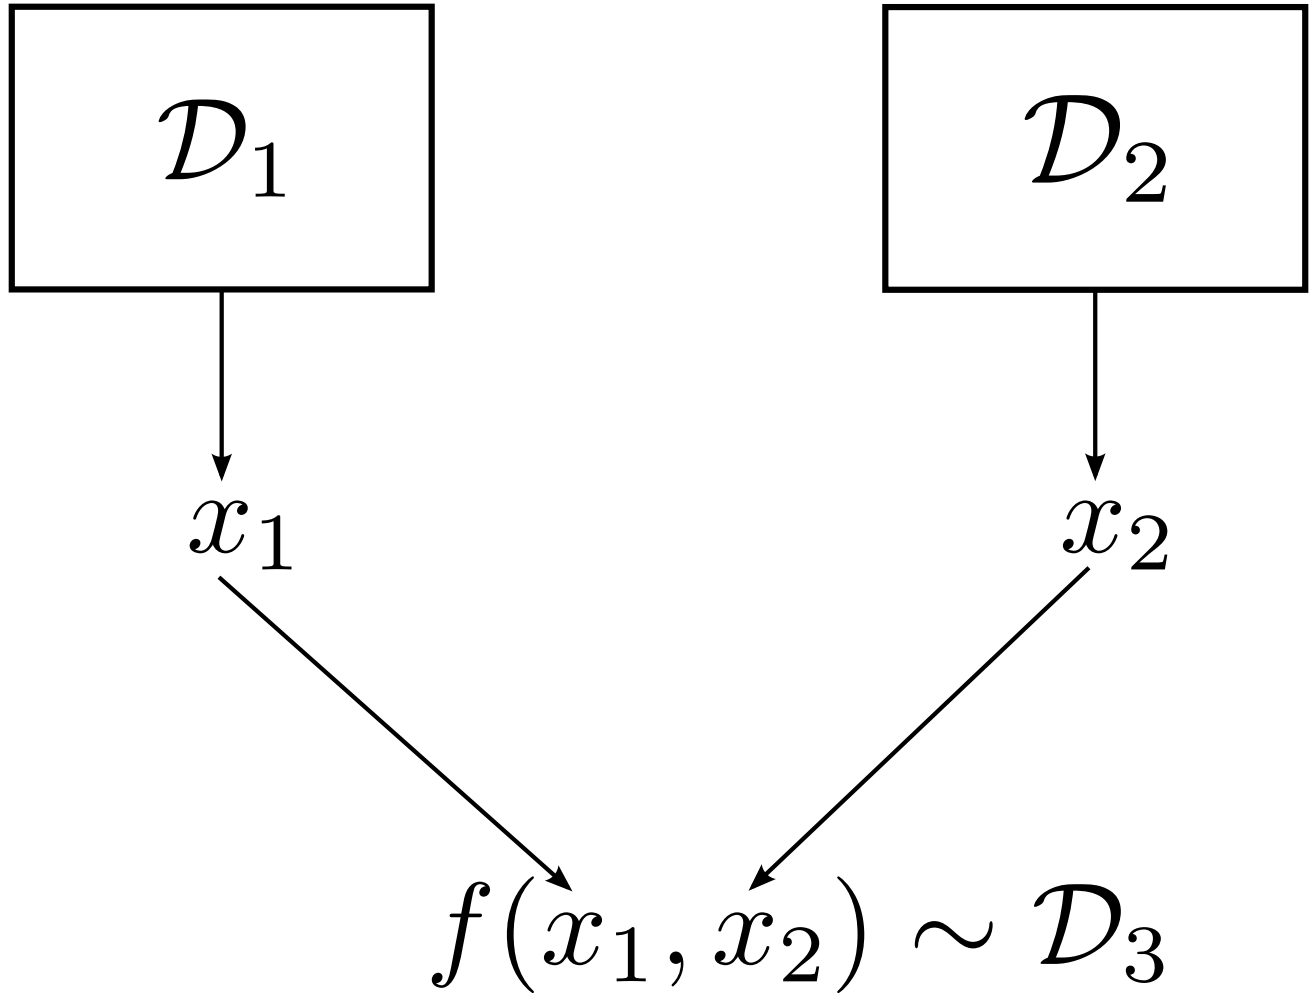

Key Idea

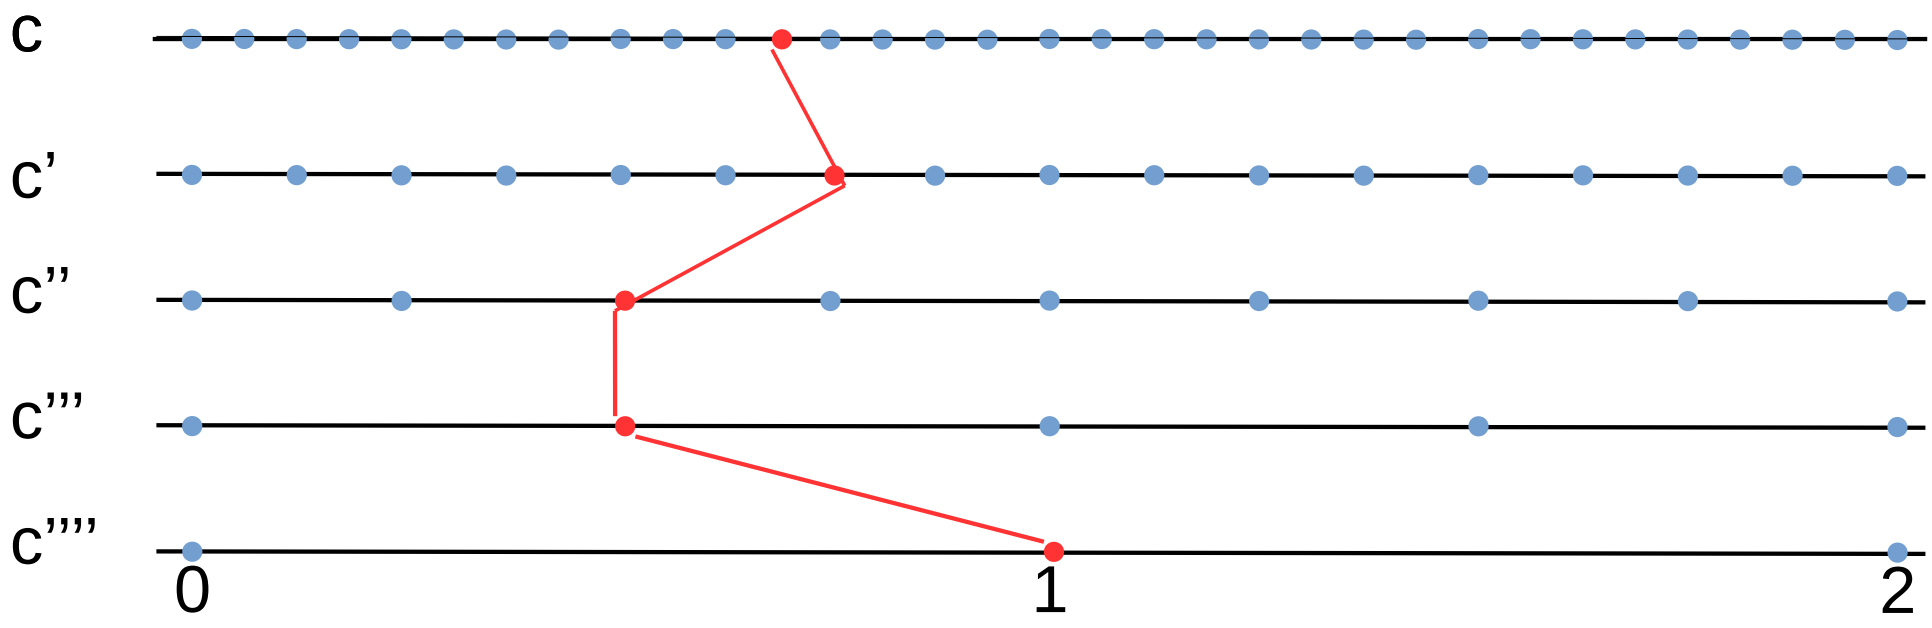

Key Idea

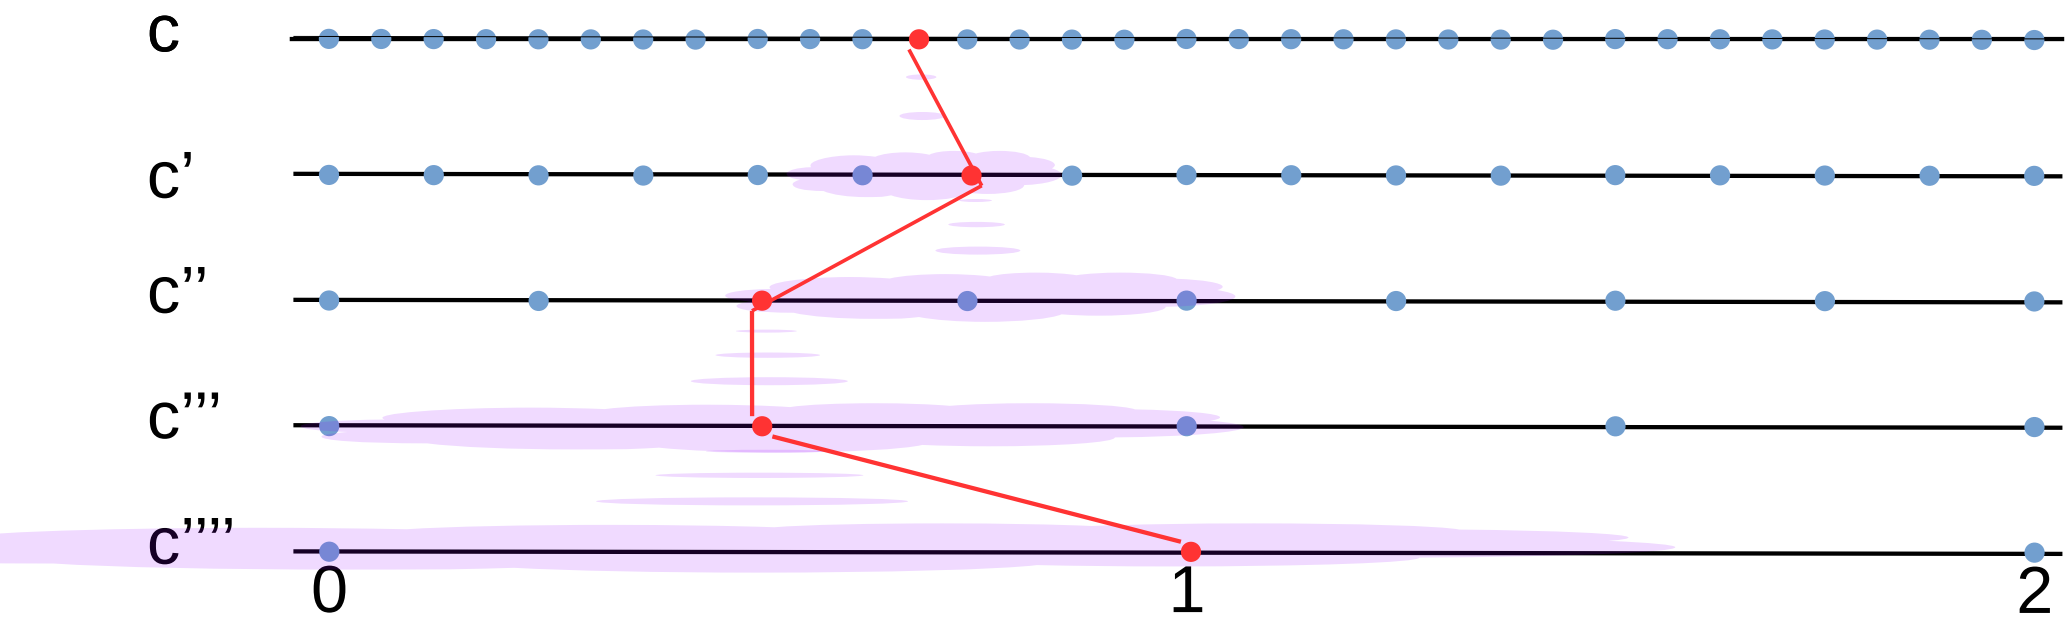

Key Idea

$$c = 0.0011011b_k \in \mathbb{Z}/2^k$$

Key Idea

$$c = 0.0011011b_k \in \mathbb{Z}/2^k$$

$$c = 0.0011001 \in \mathbb{Z}/2^{k-1}$$

Key Idea

$$c = 0.0011011b_k \in \mathbb{Z}/2^k$$

$$x \leftarrow \mathcal{D}_{2\mathbb{Z},b_k,s}$$

$$c = 0.0011001 \in \mathbb{Z}/2^{k-1}$$

Key Idea

$$c = 0.0011011b_k \in \mathbb{Z}/2^k$$

$$\downarrow \quad x \leftarrow \mathcal{D}_{2\mathbb{Z},b_k,s}$$

$$\downarrow \quad x \leftarrow (x + b_k)/2^k$$

$$c = 0.0011001 \in \mathbb{Z}/2^{k-1}$$

Key Idea

$$c = 0.0011011b_k \in \mathbb{Z}/2^k$$

$$\downarrow$$

$$x \leftarrow \mathcal{D}_{2\mathbb{Z},b_k,s}$$
$$x \leftarrow (x + b_k)/2^k$$
$$c \leftarrow c - x$$

$$c = 0.0011001 \in \mathbb{Z}/2^{k-1}$$

Arbitrary Center

$$c = 0.00110101100011111100101\ldots$$

Arbitrary Center

$$c = 0.00110101100011111100101\ldots$$

Arbitrary Center

$$c = 0.00110101100011111100101\ldots$$

Arbitrary Center

$$c = 0.\underbrace{0011010110}_{k}00111100101\ldots$$

Arbitrary Center

$$c = 0.\underbrace{00110101100}_{k}0011110010 1\ldots$$

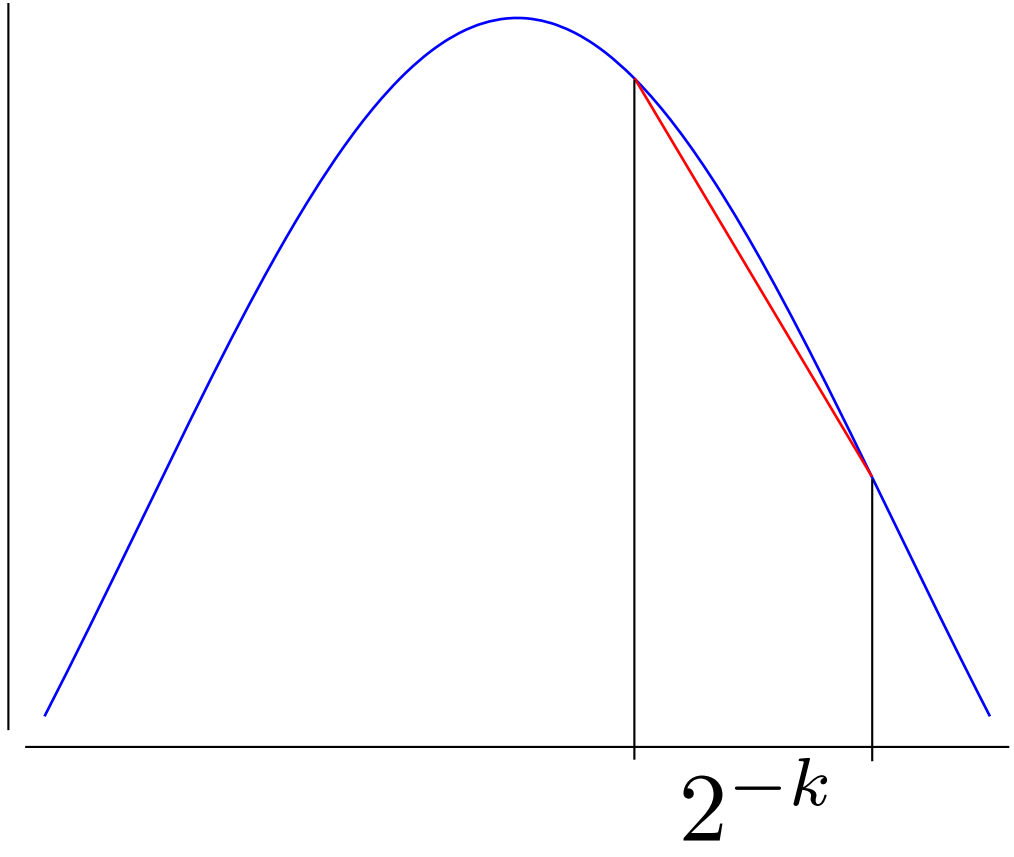

Approximate Sampler

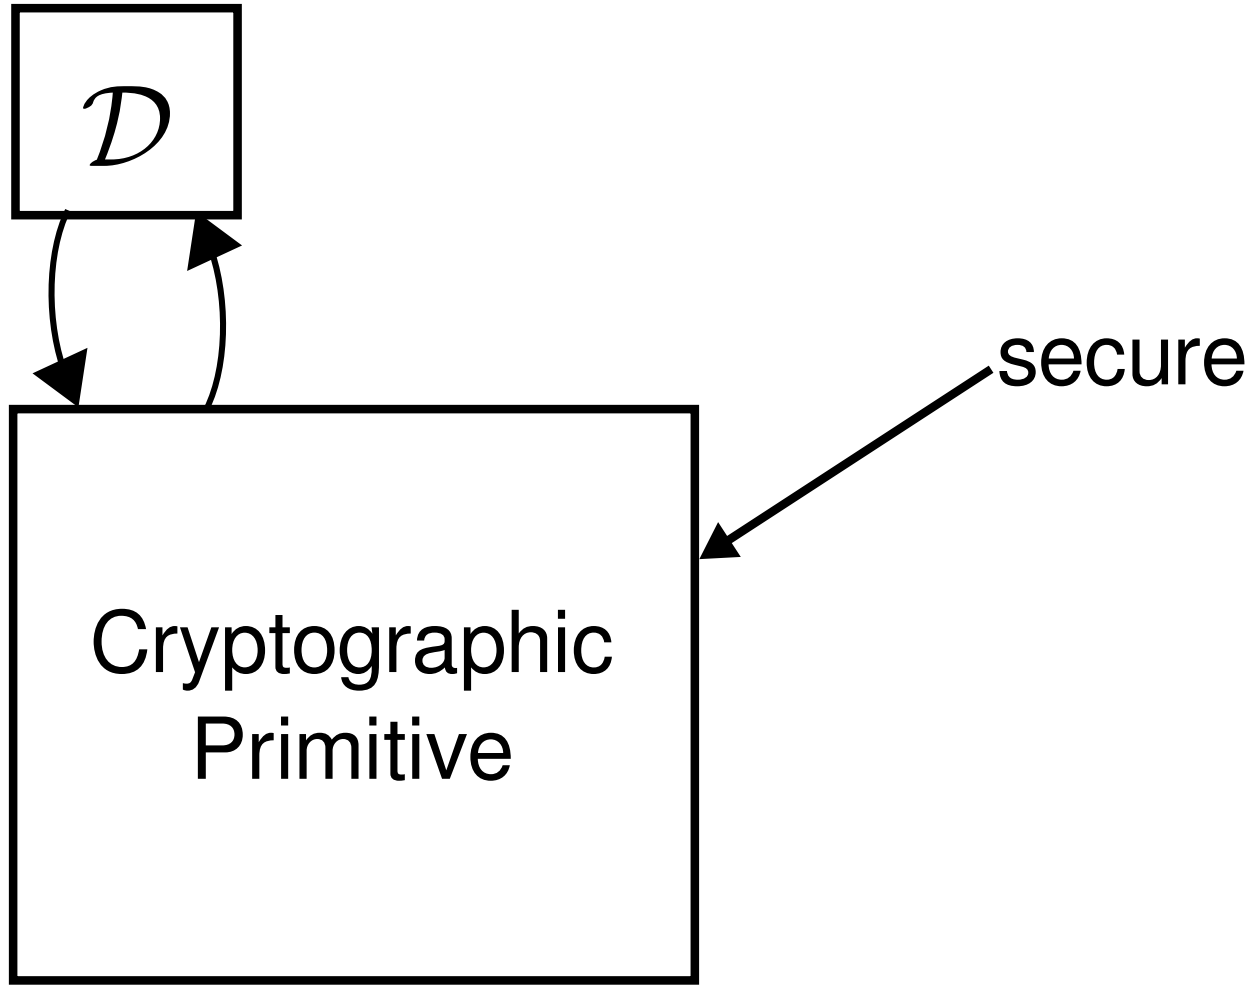

Approximate Sampler

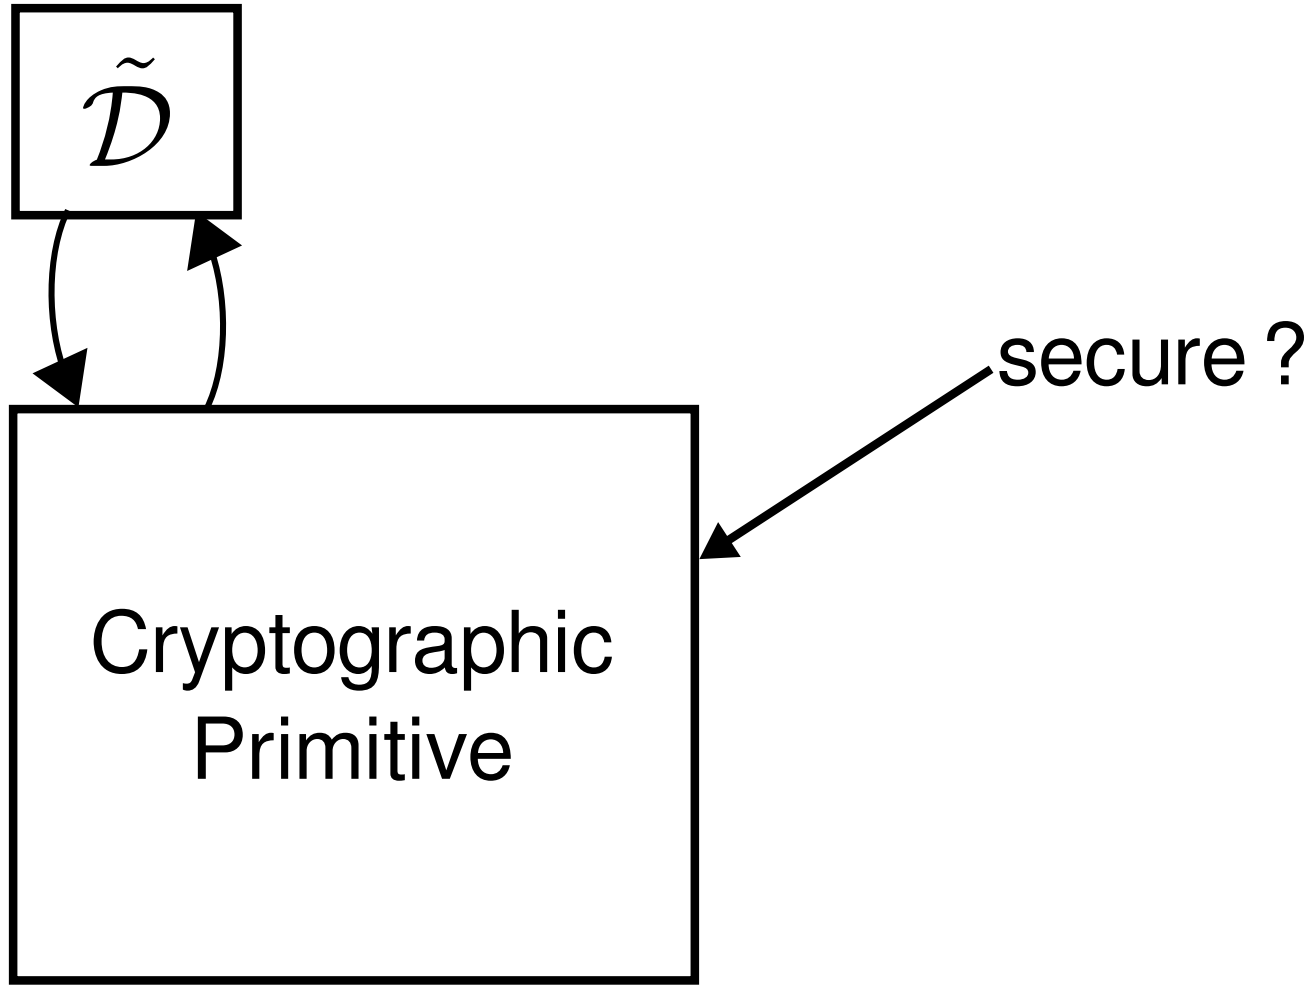

Bit Security

$$\frac{\text{Resources}}{\text{Advantage}}$$

Bit Security

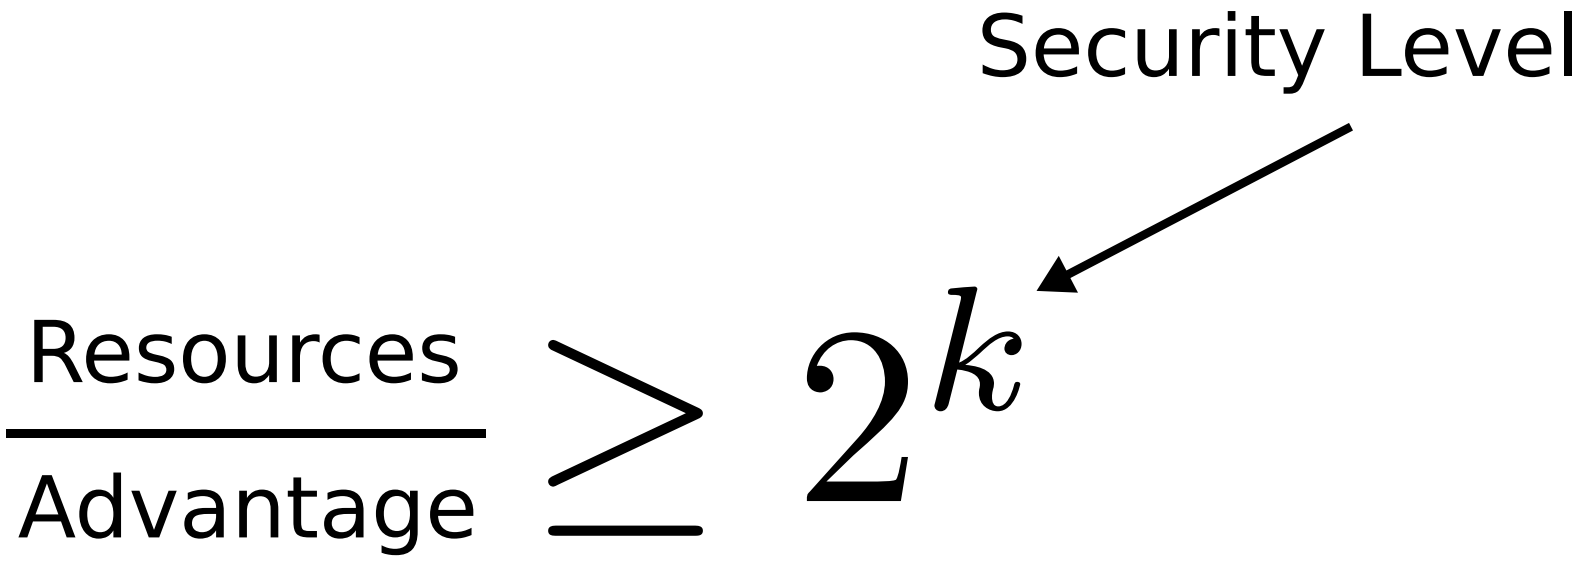

Security Proof

Security Proof

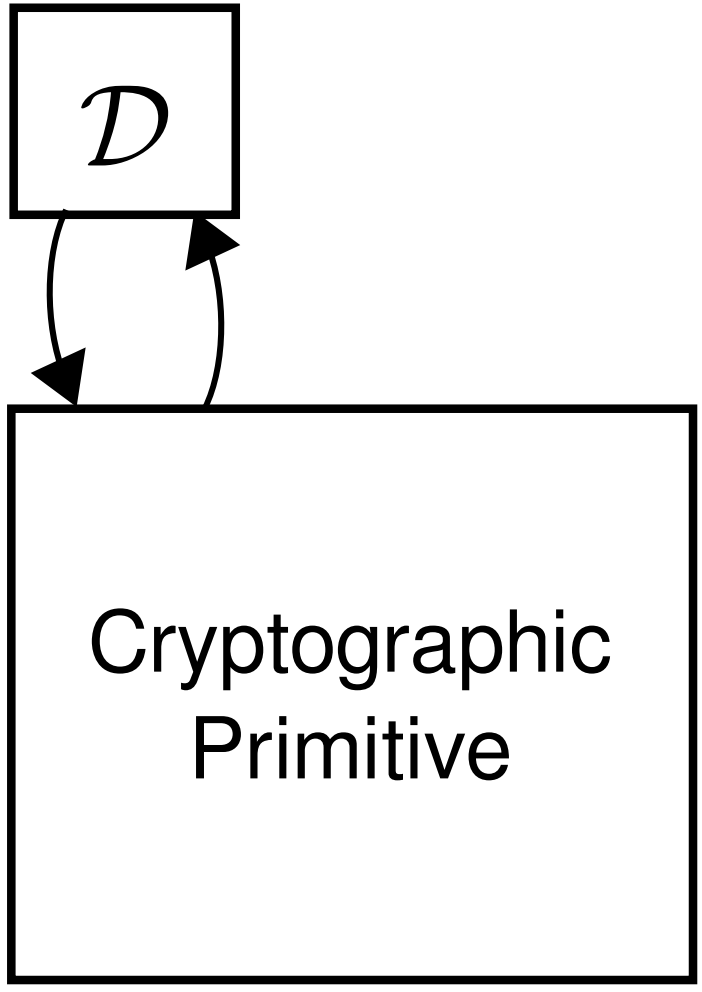

Security Proof

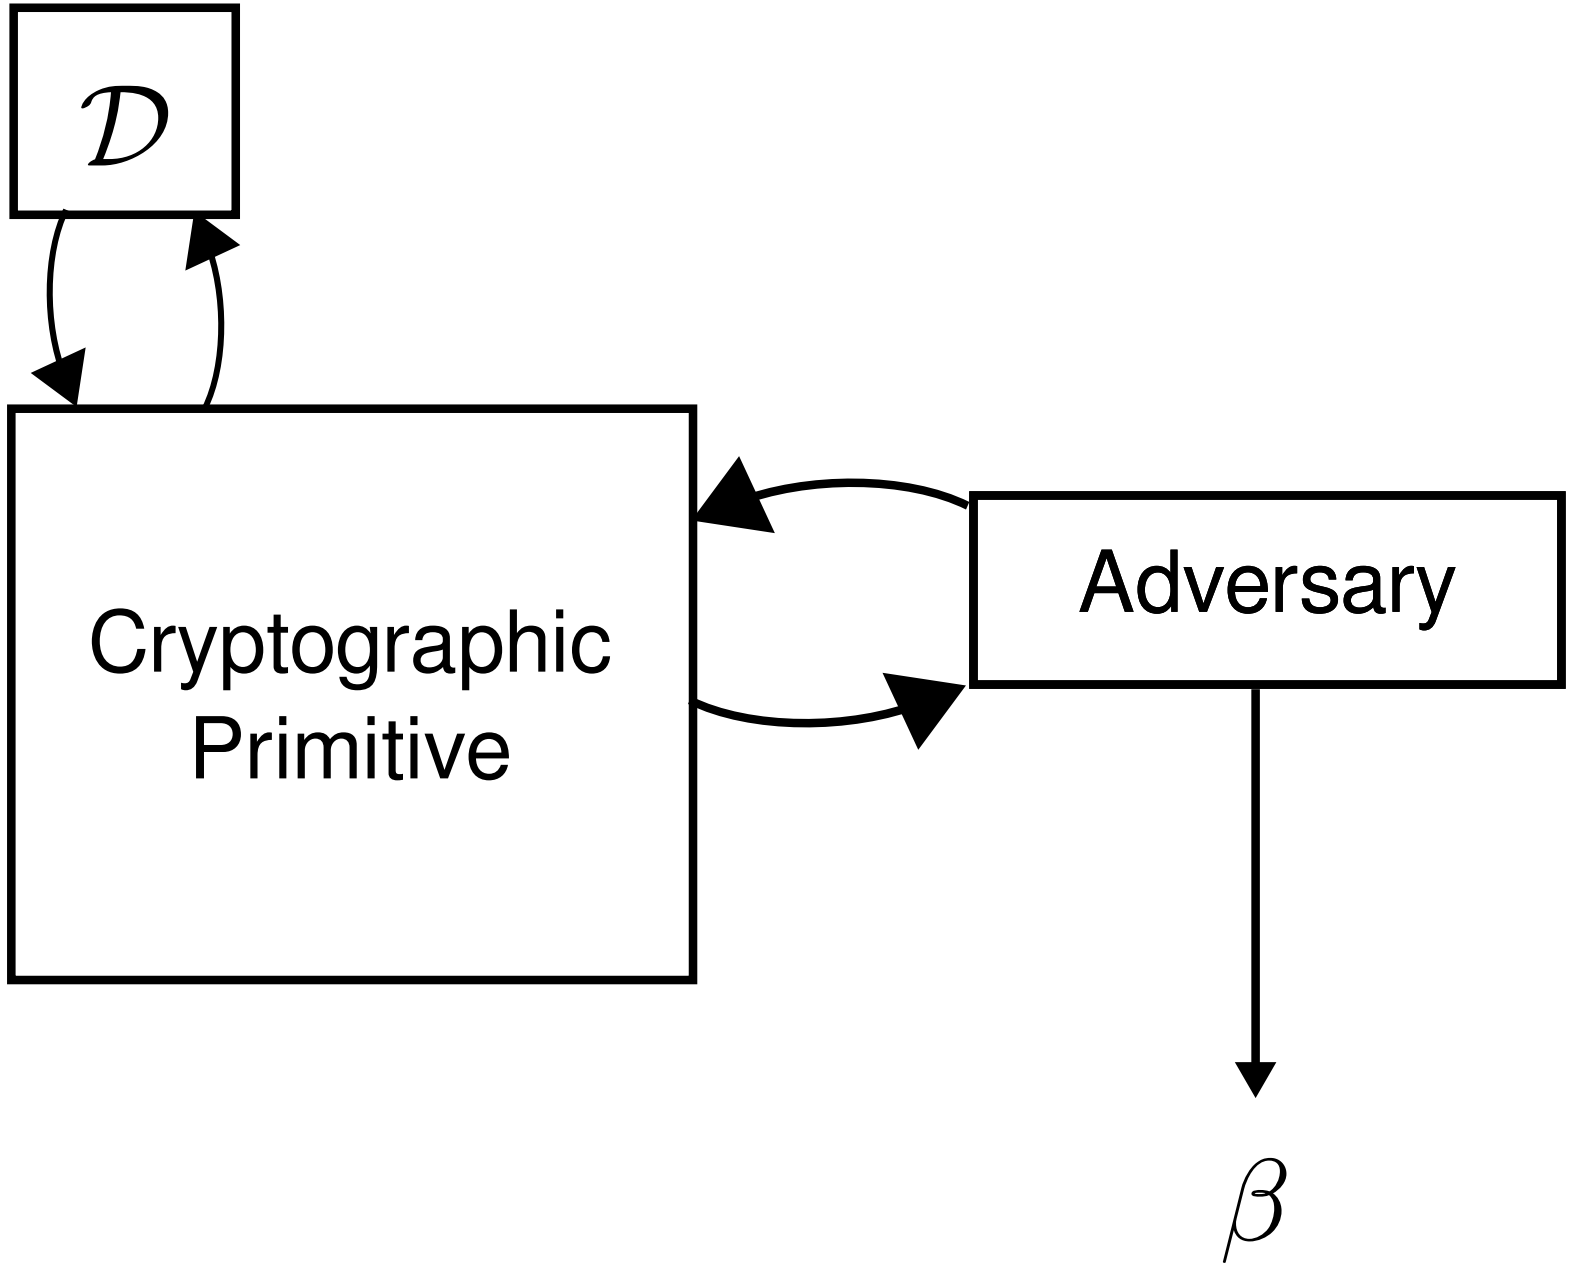

Security Proof

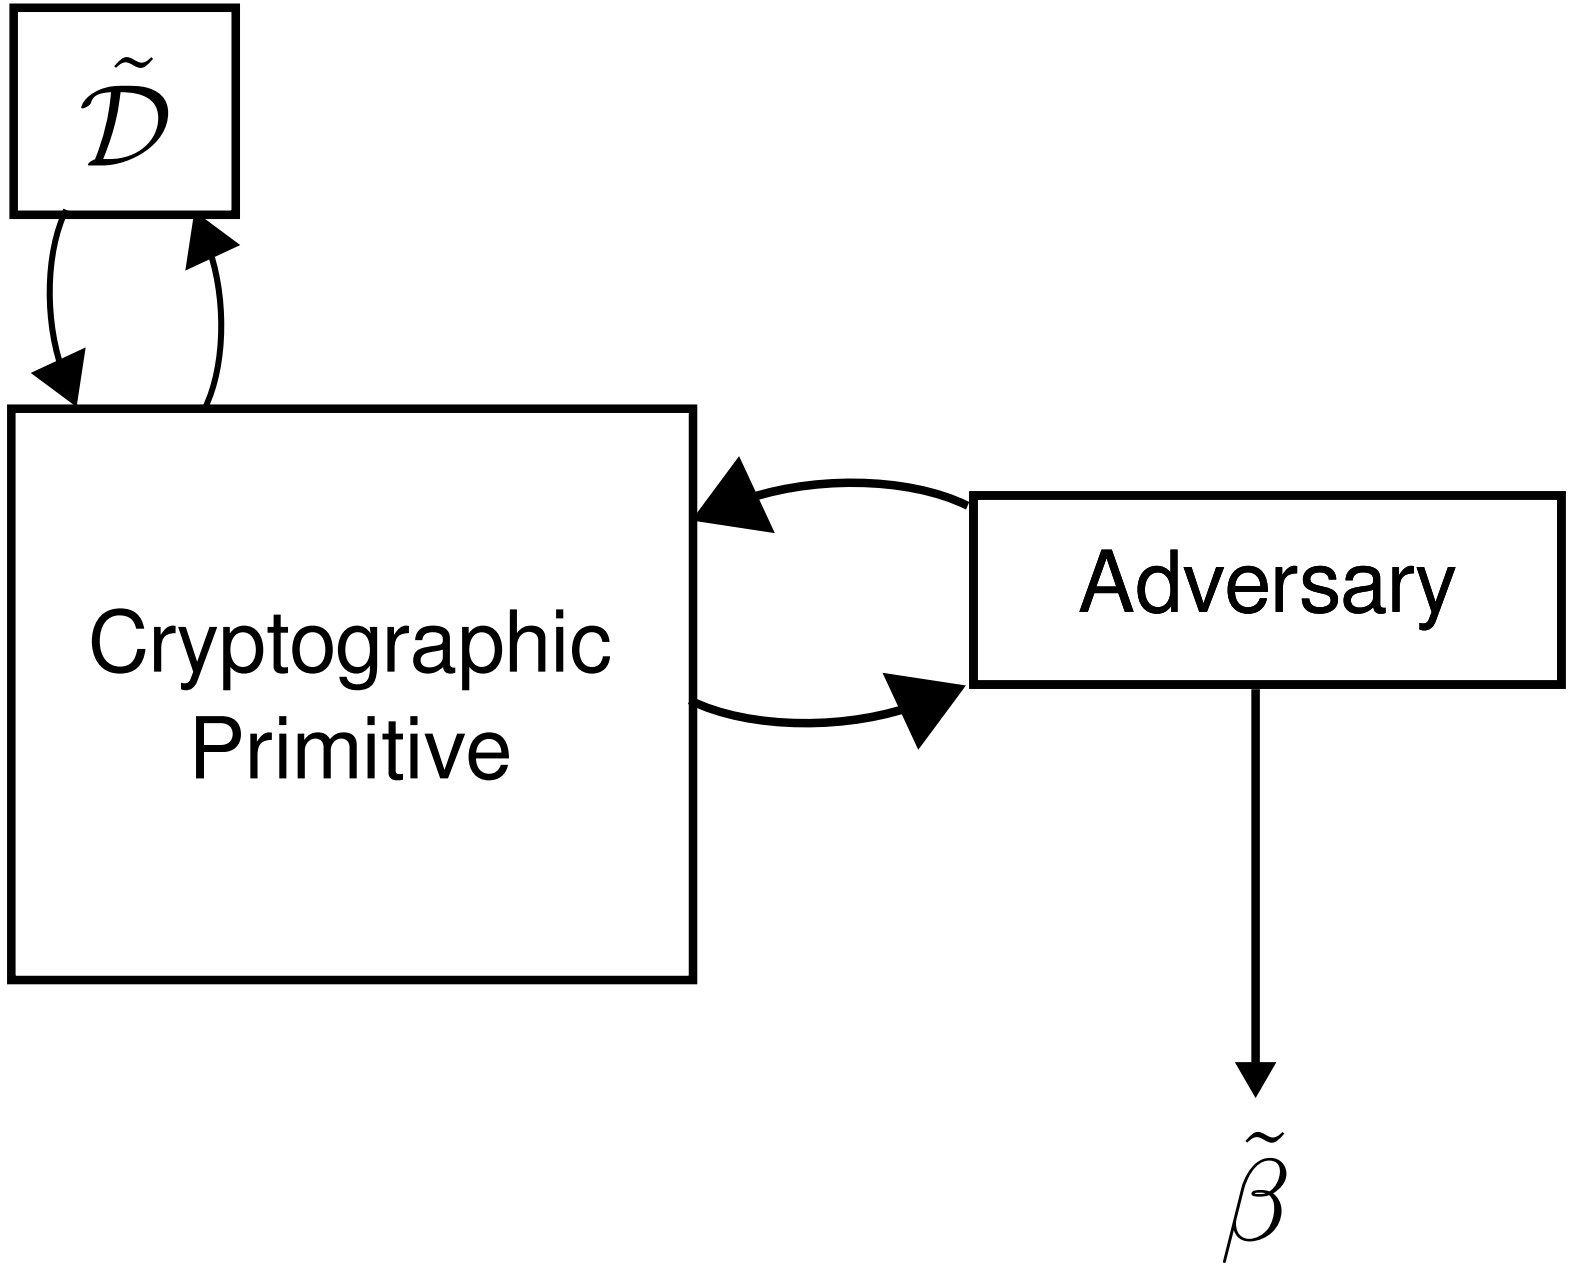

Security Proof

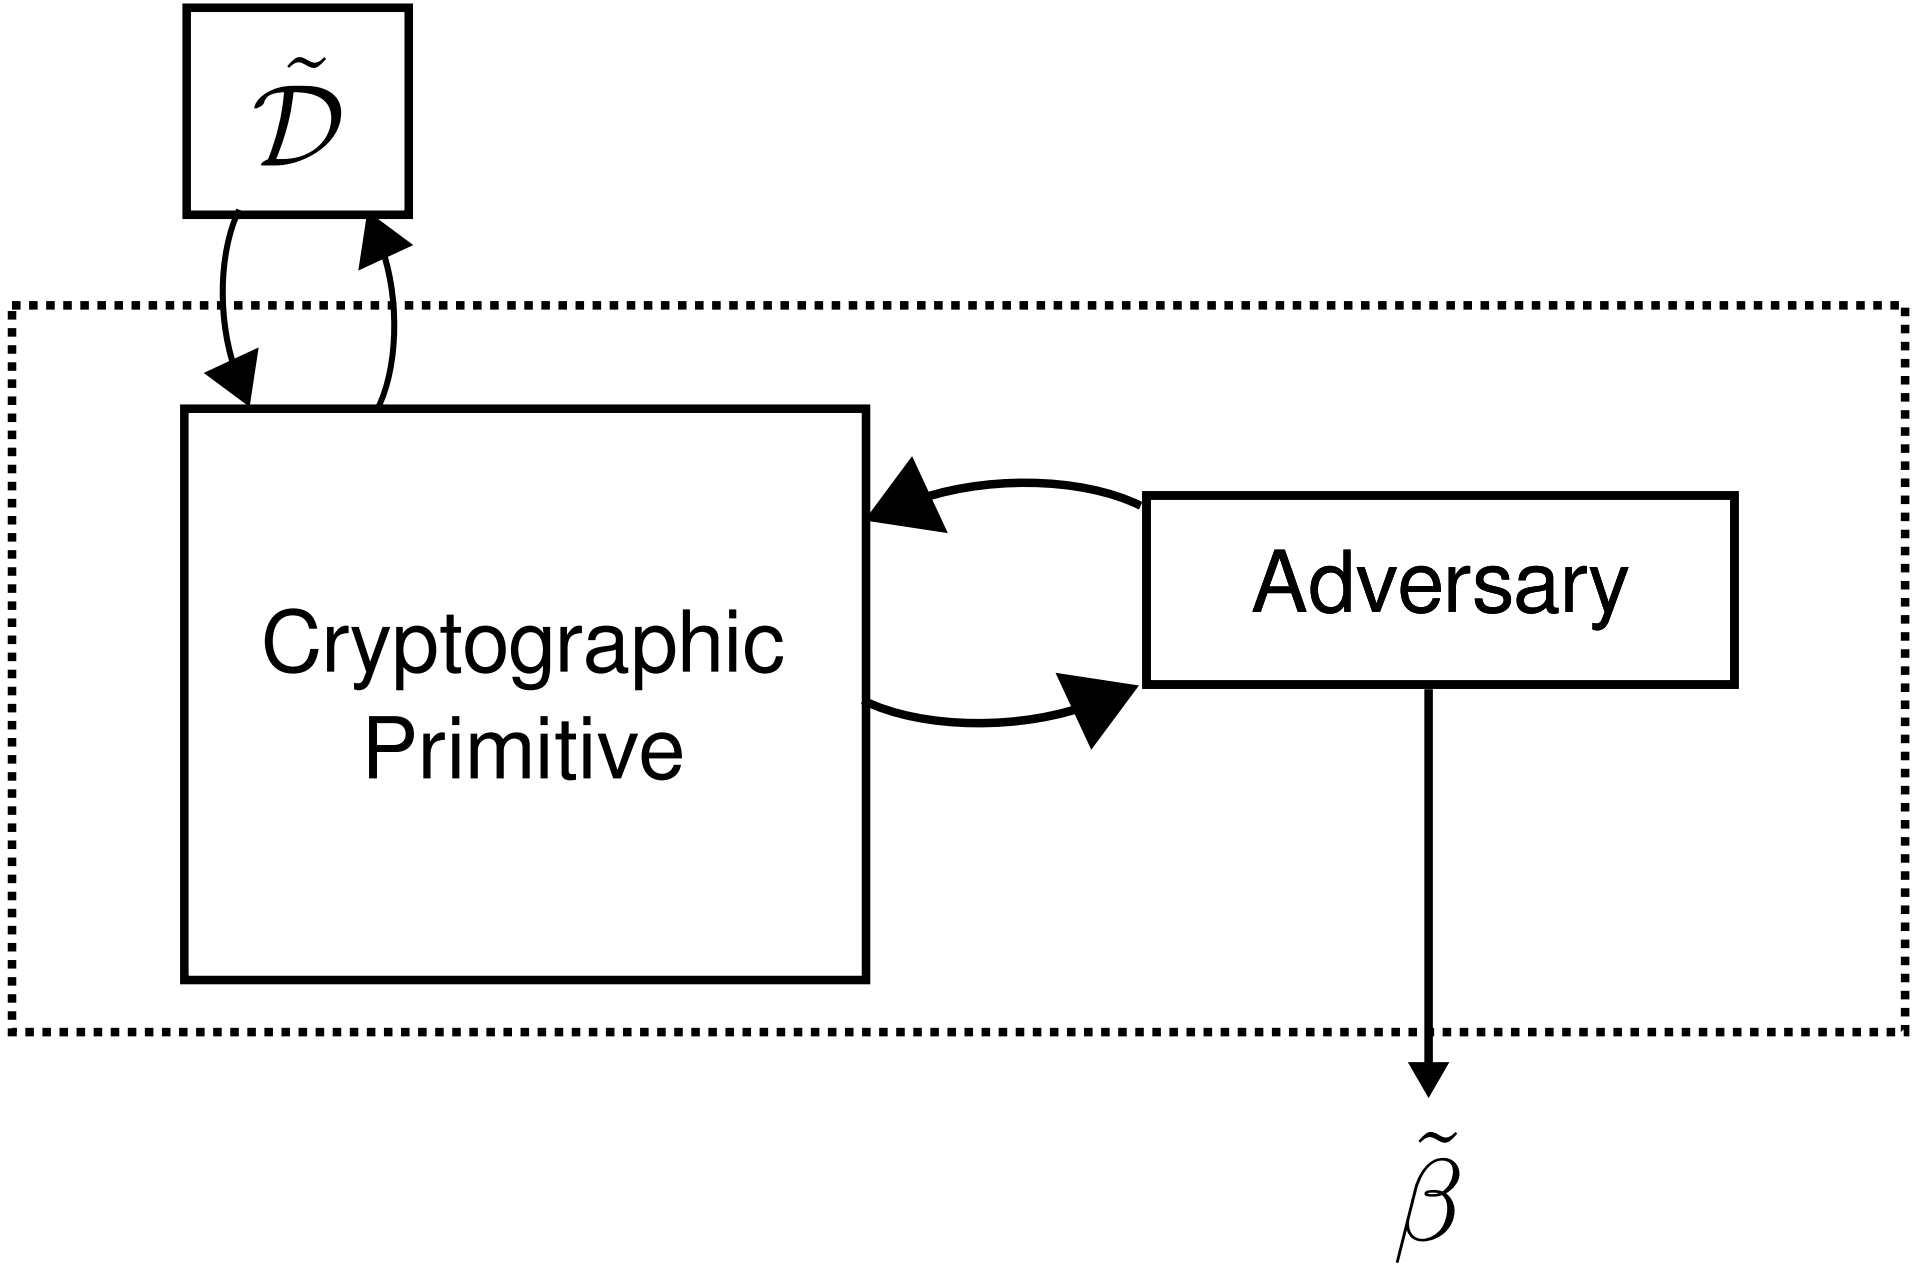

Security Proof

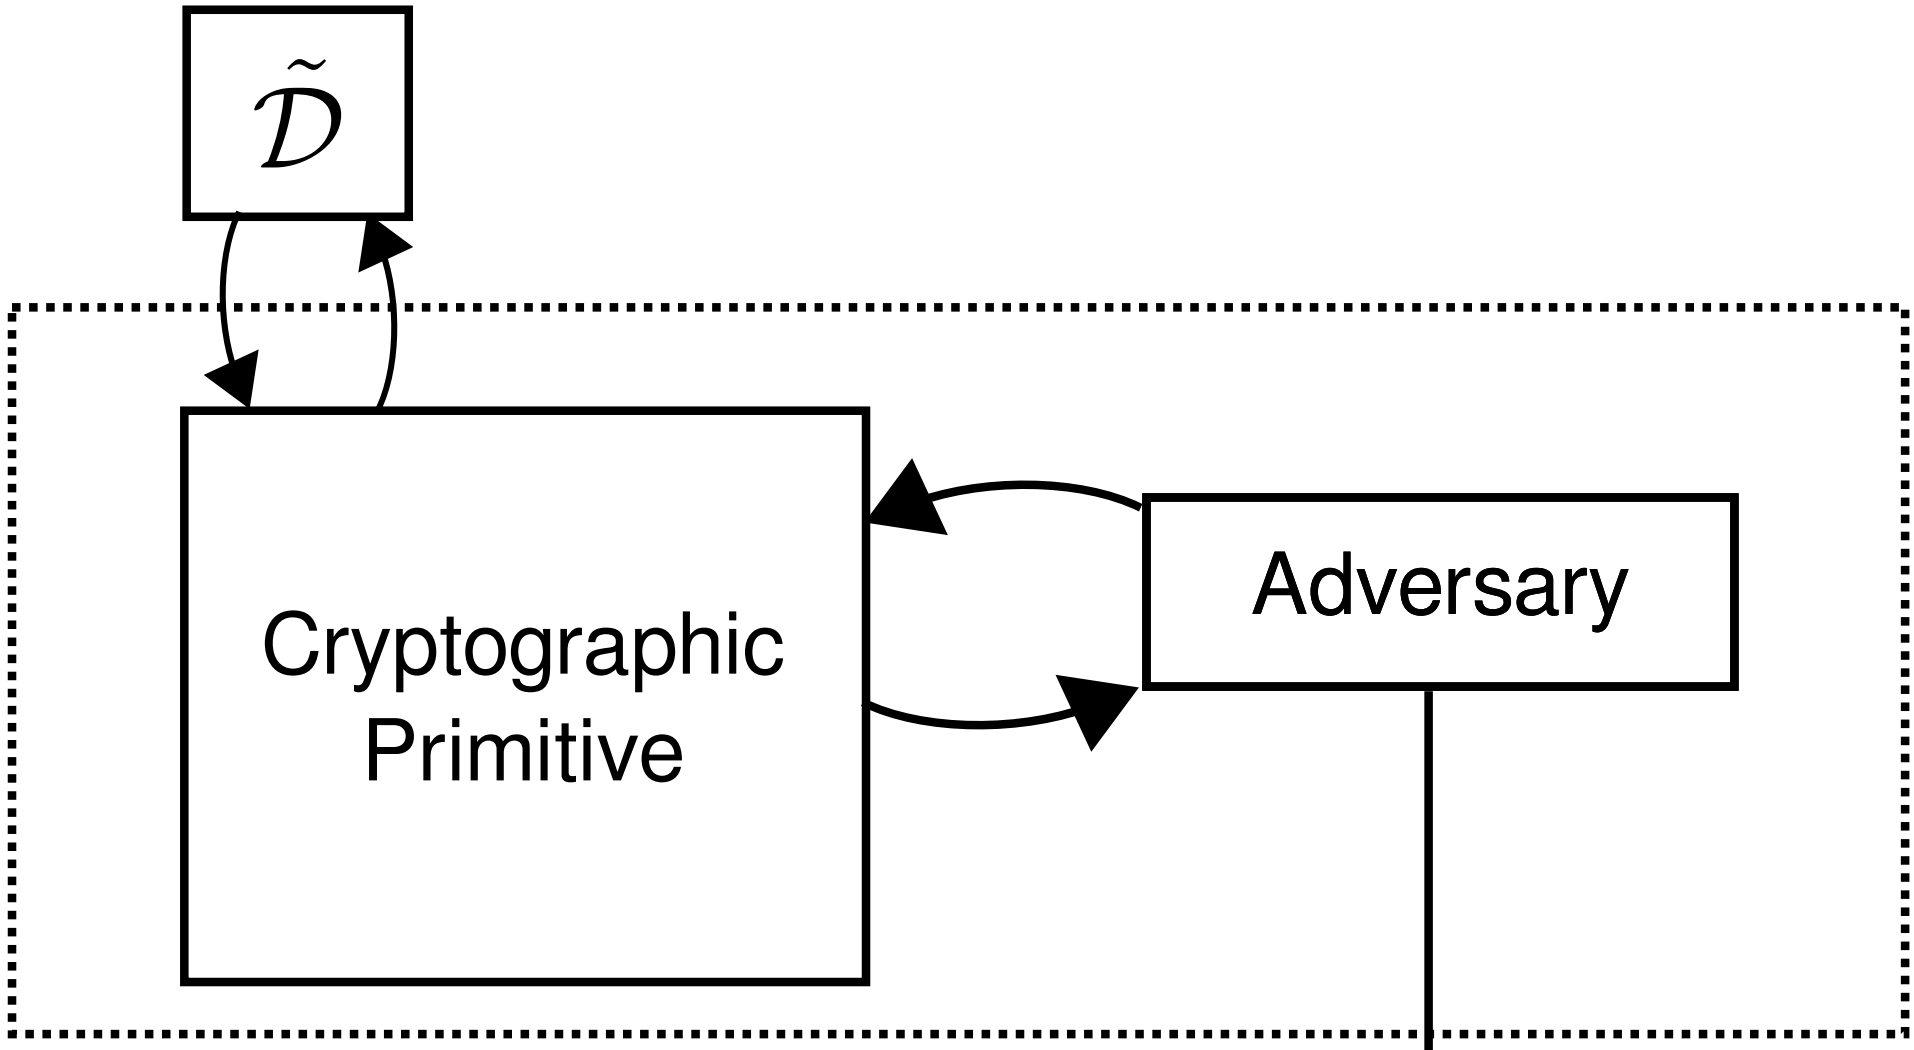

Works for any $\mathcal{D}, \tilde{\mathcal{D}}$

Security of Approximate Samplers

$$\Delta_{SD}(\mathcal{P}, \mathcal{Q}) = \max_{E \in \Omega} (\mathcal{P}(E) - \mathcal{Q}(E))$$

Security of Approximate Samplers

$$\Delta_{SD}(\mathcal{P}, \mathcal{Q}) = \max_{E \in \Omega} \left(\mathcal{P}(E) - \mathcal{Q}(E)\right)$$

$$\tilde{\beta} \leq \beta + \Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q)$$

Security of Approximate Samplers

$$\Delta_{SD}(\mathcal{P}, \mathcal{Q}) = \max_{E \in \Omega} \left(\mathcal{P}(E) - \mathcal{Q}(E)\right)$$

Advantage against

approximate scheme

ideal scheme

$$\tilde{\beta} \leq \beta + \Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q)$$

Security of Approximate Samplers

$$\Delta_{SD}(\mathcal{P}, \mathcal{Q}) = \max_{E \in \Omega} \left(\mathcal{P}(E) - \mathcal{Q}(E)\right)$$

Advantage against

approximate scheme

ideal scheme

number of queries

$$\tilde{\beta} \leq \beta + \Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q)$$

Security of Approximate Samplers

$$\Delta_{SD}(\mathcal{P}, \mathcal{Q}) = \max_{E \in \Omega} \left(\mathcal{P}(E) - \mathcal{Q}(E)\right)$$

Advantage against

approximate scheme

ideal scheme

number of queries

$$\tilde{\beta} \leq \beta + \Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q)$$

$$\beta \geq \tilde{\beta} - \Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q)$$

Security of Approximate Samplers

$$\Delta_{SD}(\mathcal{P}, \mathcal{Q}) = \max_{E \in \Omega} \left(\mathcal{P}(E) - \mathcal{Q}(E)\right)$$

Advantage against

approximate scheme

ideal scheme

number of queries

$$\tilde{\beta} \leq \beta + \Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q)$$

$$\beta \geq \tilde{\beta} - \Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q)$$

$$\Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q) \leq \mathbf{?}$$

Security of Approximate Samplers

Classical Approach:

Security of Approximate Samplers

Classical Approach:

$$\Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q) \leq q\Delta_{SD}(\mathcal{D}, \tilde{\mathcal{D}})$$

Security of Approximate Samplers

Classical Approach:

$$\Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q) \leq q\Delta_{SD}(\mathcal{D}, \tilde{\mathcal{D}})$$

$$\Longrightarrow$$

$$\Delta_{SD}(\mathcal{D}, \tilde{\mathcal{D}}) \lesssim \frac{1}{q}$$

Security of Approximate Samplers

Modern Approach:

Security of Approximate Samplers

Modern Approach:

$$\Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q) \leq \sqrt{KL(\mathcal{D}^q || \tilde{\mathcal{D}}^q)}$$

Security of Approximate Samplers

Modern Approach:

$$\Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q) \leq \sqrt{KL(\mathcal{D}^q || \tilde{\mathcal{D}}^q)}$$

$$\leq \sqrt{qKL(\mathcal{D} || \tilde{\mathcal{D}})}$$

Security of Approximate Samplers

Modern Approach:

$$\Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q) \leq \sqrt{KL(\mathcal{D}^q || \tilde{\mathcal{D}}^q)}$$

$$\leq \sqrt{qKL(\mathcal{D}||\tilde{\mathcal{D}})}$$

$$\Longrightarrow$$

$$\sqrt{KL(\mathcal{D}||\tilde{\mathcal{D}})} \lesssim \frac{1}{\sqrt{q}}$$

Recall: Convolution

Recall: Convolution

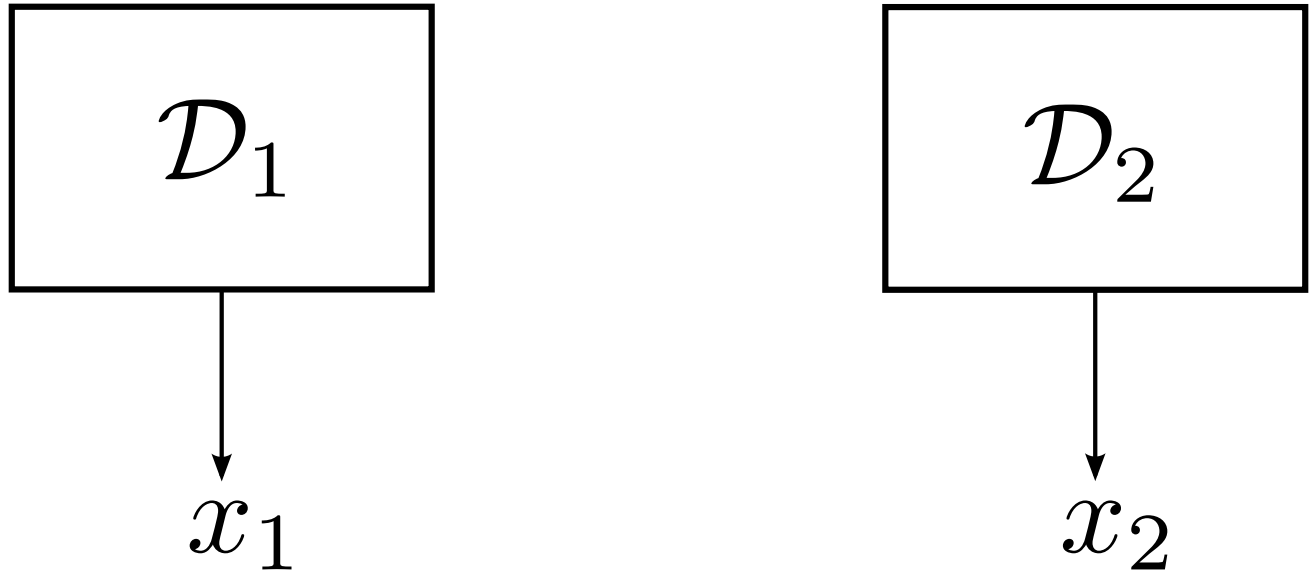

Recall: Convolution

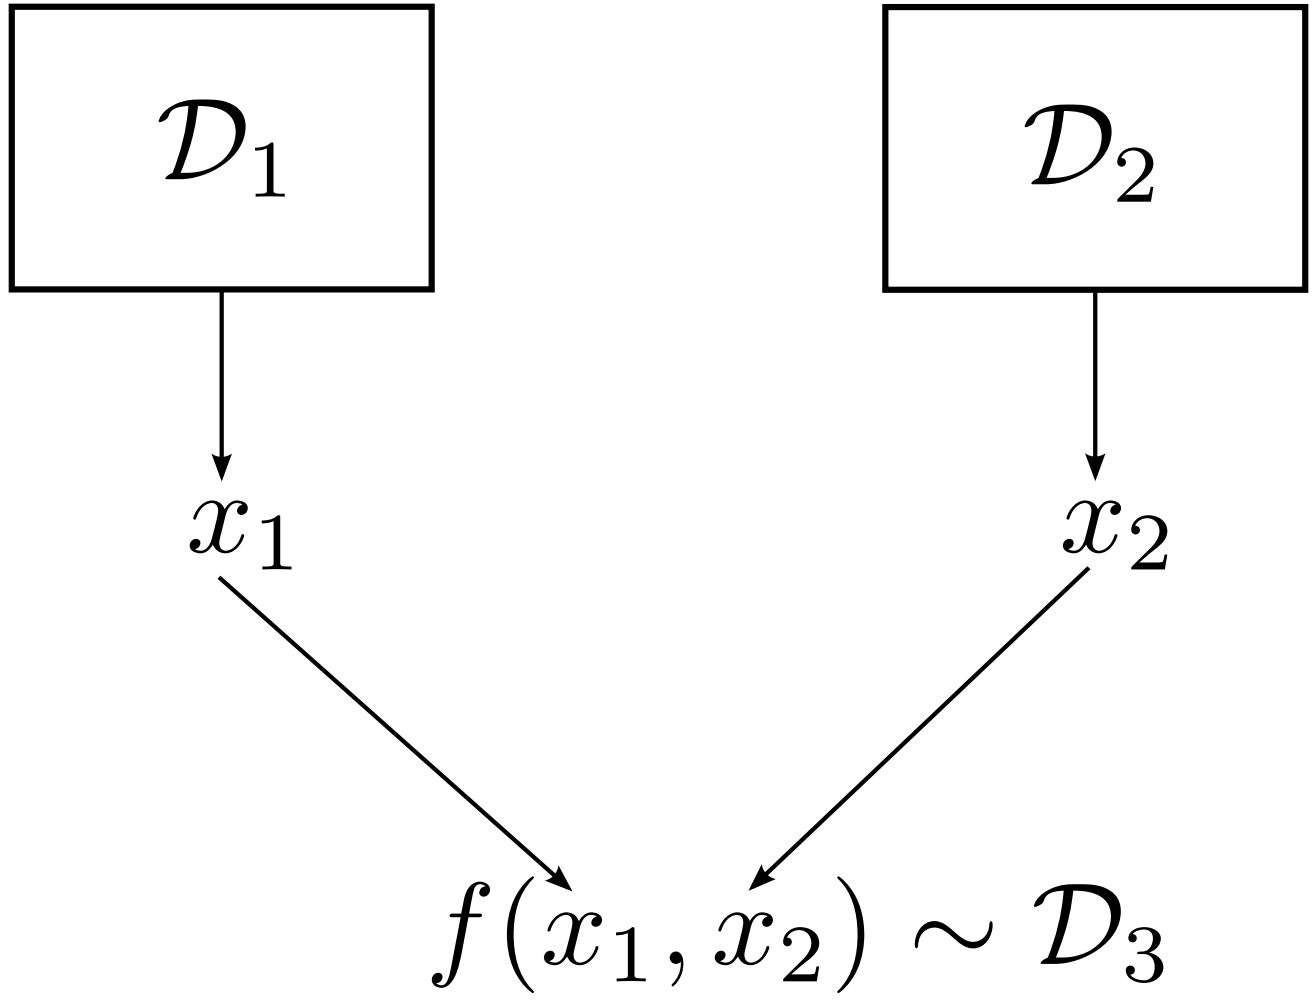

Recall: Convolution

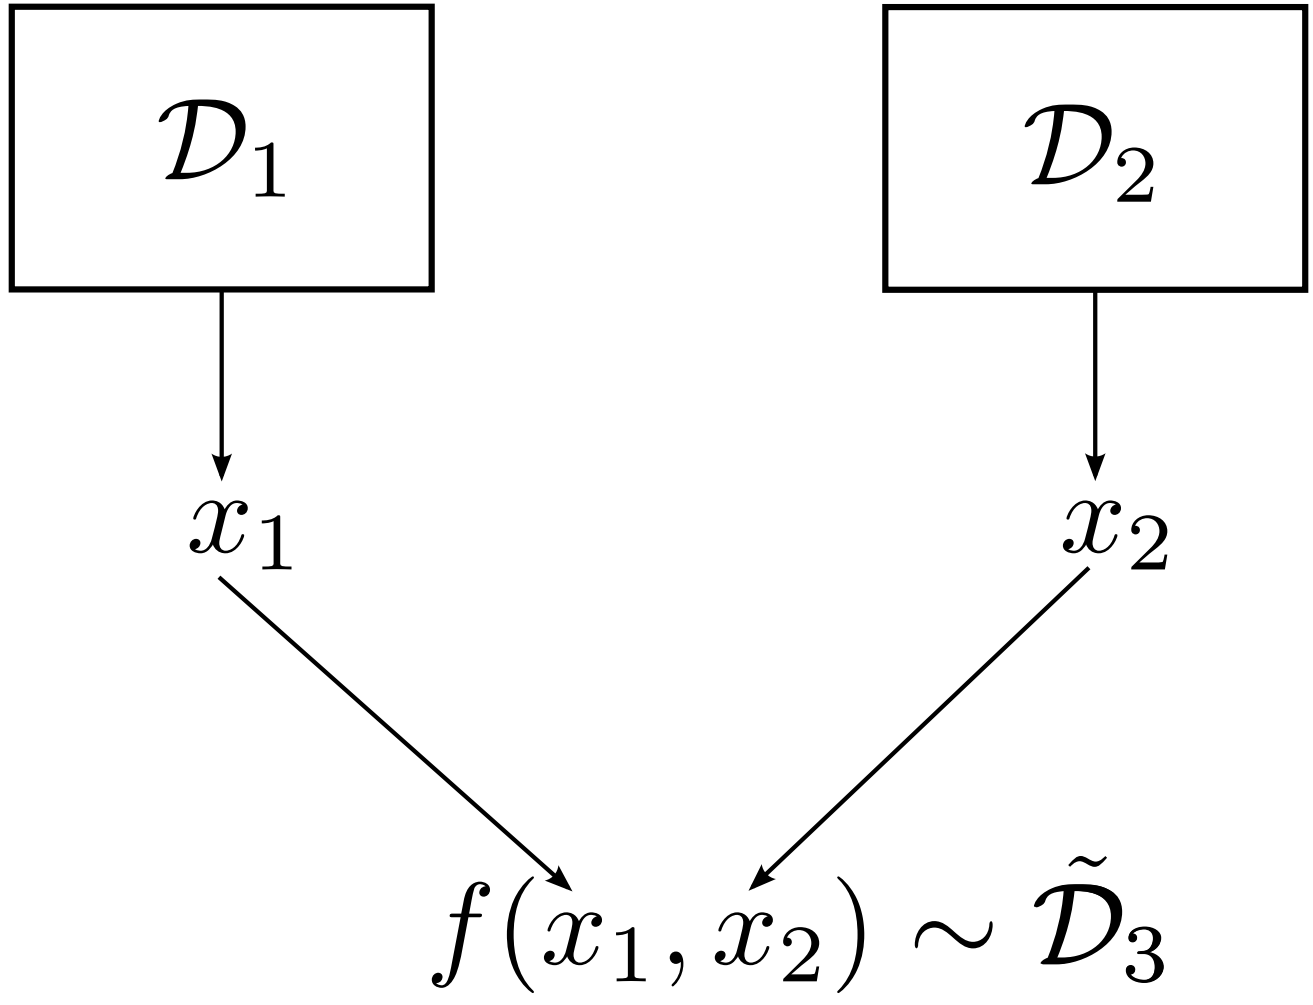

Recall: Convolution

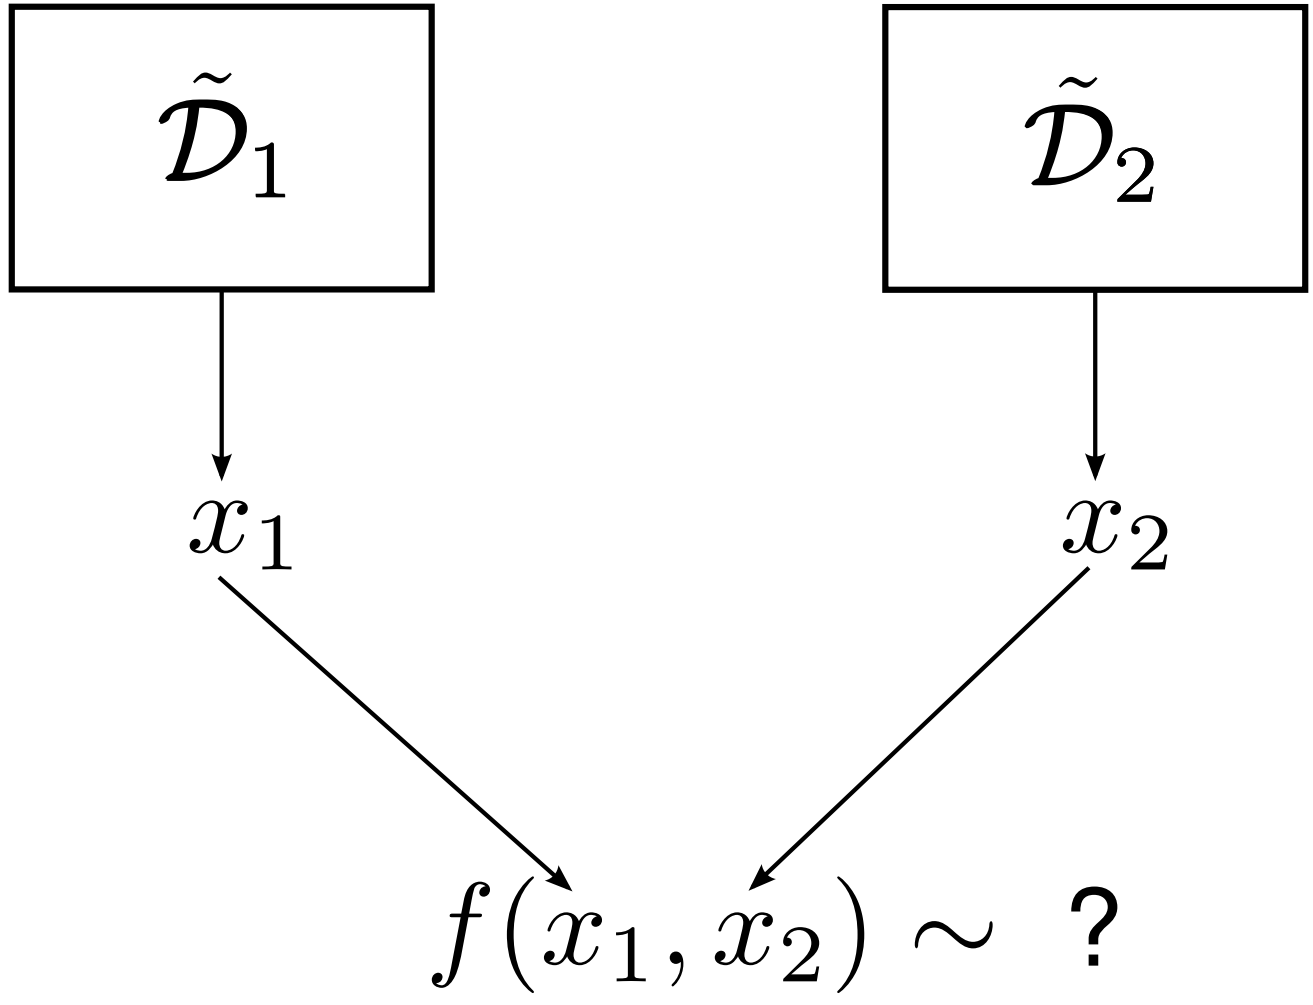

Security of Approximate Samplers

Our Approach:

Security of Approximate Samplers

Our Approach:

$$\Delta_{ML}(\mathcal{P}, \mathcal{Q}) = \max_x |\ln \mathcal{P}(x) - \ln \mathcal{Q}(x)|$$

Security of Approximate Samplers

Our Approach:

$$\Delta_{ML}(\mathcal{P}, \mathcal{Q}) = \max_x |\ln \mathcal{P}(x) - \ln \mathcal{Q}(x)|$$

$$\Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q) \leq \sqrt{q}\Delta_{ML}(\mathcal{D}, \tilde{\mathcal{D}})$$

Security of Approximate Samplers

Our Approach:

$$\Delta_{ML}(\mathcal{P}, \mathcal{Q}) = \max_x |\ln \mathcal{P}(x) - \ln \mathcal{Q}(x)|$$

$$\Delta_{SD}(\mathcal{D}^q, \tilde{\mathcal{D}}^q) \leq \sqrt{q}\Delta_{ML}(\mathcal{D}, \tilde{\mathcal{D}})$$

$$\Longrightarrow$$

$$\Delta_{ML}(\mathcal{D}, \tilde{\mathcal{D}}) \lesssim \frac{1}{\sqrt{q}}$$

Security of Approximate Samplers

$$\Delta_{SD} \qquad KL$$

Metric

Strong Security

Security of Approximate Samplers

	Δ_{SD}	KL
Metric		
Strong Security		

Security of Approximate Samplers

	Δ_{SD}	KL
Metric		
Strong Security		

Security of Approximate Samplers

	Δ_{SD}	KL
Metric		
Strong Security		

Security of Approximate Samplers

	Δ_{SD}	KL
Metric		
Strong Security		

Security of Approximate Samplers

	Δ_{SD}	KL	Δ_{ML}
Metric			
Strong Security			

Security of Approximate Samplers

	Δ_{SD}	KL	Δ_{ML}
Metric			
Strong Security			

Security of Approximate Samplers

	Δ_{SD}	KL	Δ_{ML}
Metric			
Strong Security