Isochronous Gaussian Sampling: From Inception to Implementation

With Applications to the Falcon Signature Scheme

James Howe - Thomas Prest - Thomas Ricosset - Mélissa Rossi

PQCrypto 2020

21-23 September

Falcon

(P-A Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Prest, T. Ricosset, G. Seiler, W. Whyte, Z. Zhang)

Falcon

(P-A Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Prest, T. Ricosset, G. Seiler, W. Whyte, Z. Zhang)

Based on the GPV framework

Gentry, Peikert and Vaikuntanathan STOC 2008

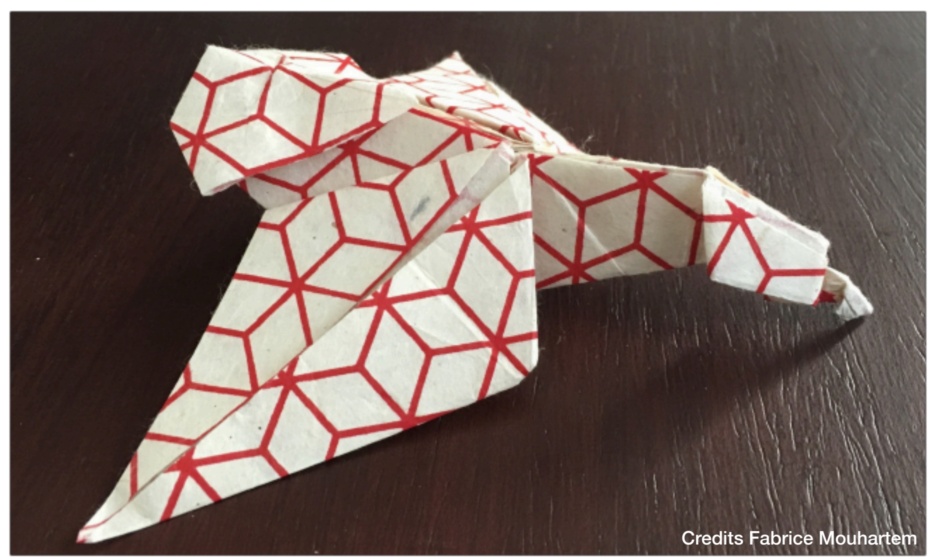

Falcon

(P-A Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Prest, T. Ricosset, G. Seiler, W. Whyte, Z. Zhang)

Based on the GPV framework

Gentry, Peikert and Vaikuntanathan STOC 2008

Relying on NTRU lattices

Hoffstein et al. ANTS 1998, CT-RSA 2003

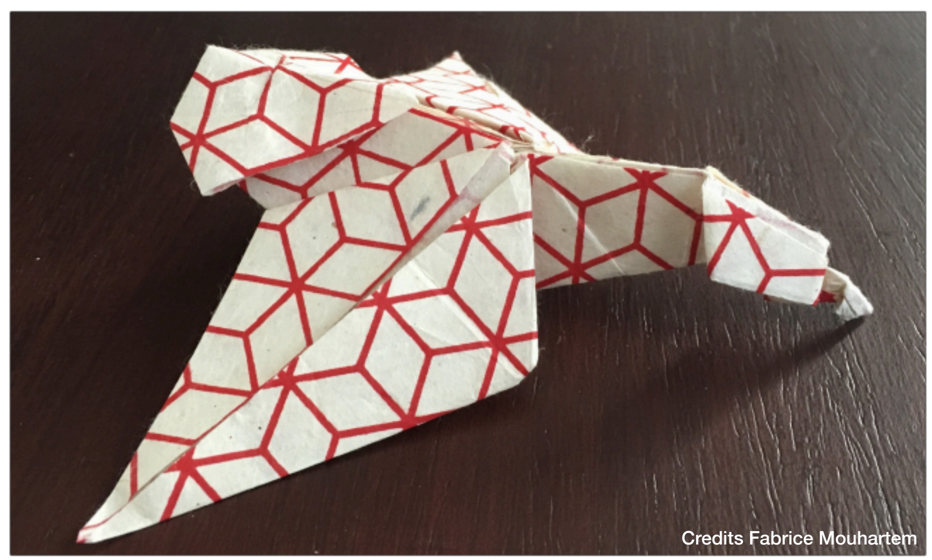

Falcon

(P-A Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Prest, T. Ricosset, G. Seiler, W. Whyte, Z. Zhang)

Based on the GPV framework

Gentry, Peikert and Vaikuntanathan STOC 2008

Relying on NTRU lattices

Hoffstein et al. ANTS 1998, CT-RSA 2003

Using Fast Fourier Orthogonalization

Ducas-Prest, ISSAC 2016

Falcon

(P-A Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Prest, T. Ricosset, G. Seiler, W. Whyte, Z. Zhang)

Based on the GPV framework

Gentry, Peikert and Vaikuntanathan STOC 2008

Relying on NTRU lattices

Hoffstein et al. ANTS 1998, CT-RSA 2003

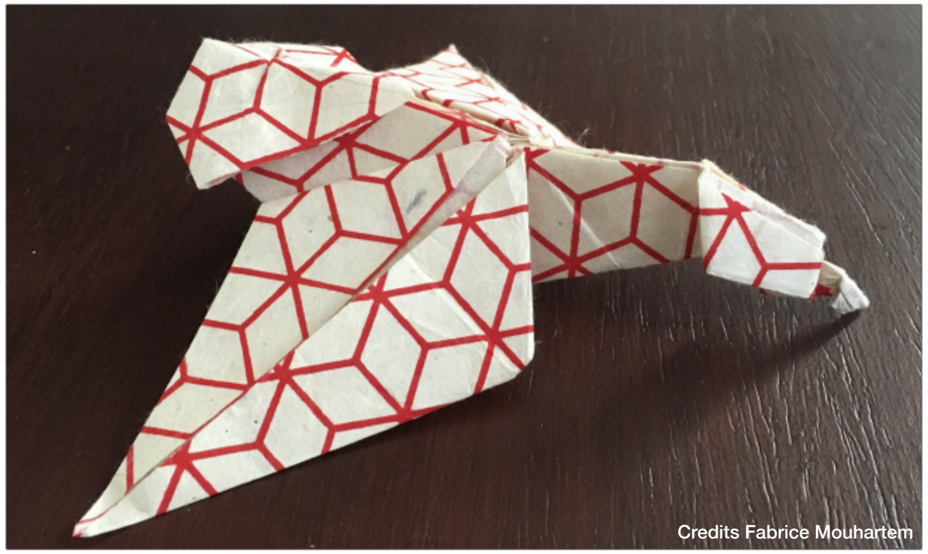

Using Fast Fourier Orthogonalization

Ducas-Prest, ISSAC 2016

Compact signatures

|s| + |pk| minimized

Falcon in a nutshell

$$\mathcal{R} = \frac{\mathbb{Z}_q[x]}{x^n + 1}$$

KeyGen()

- Generate matrices $\mathbf{A}, \mathbf{B}$ with coefficients in $\mathcal{R}$

  such that $\begin{cases} \mathbf{BA} = \mathbf{0} \\ \mathbf{B} \text{ has small coefficients} \end{cases}$
- $pk \leftarrow \mathbf{A}$
- $sk \leftarrow \mathbf{B}$

Sign(m,sk)

- Compute $\mathbf{c}$ such that $\mathbf{cA} = H(m)$
- $\mathbf{v} \leftarrow$ a vector in $\Lambda(\mathbf{B})$ close to $\mathbf{c}$
- $\mathbf{s} \leftarrow \mathbf{c} - \mathbf{v}$

Verify(m,pk,s)

Accept iff:

$\begin{cases} \mathbf{s} \text{ is short} \\ \mathbf{sA} = H(m) \end{cases}$

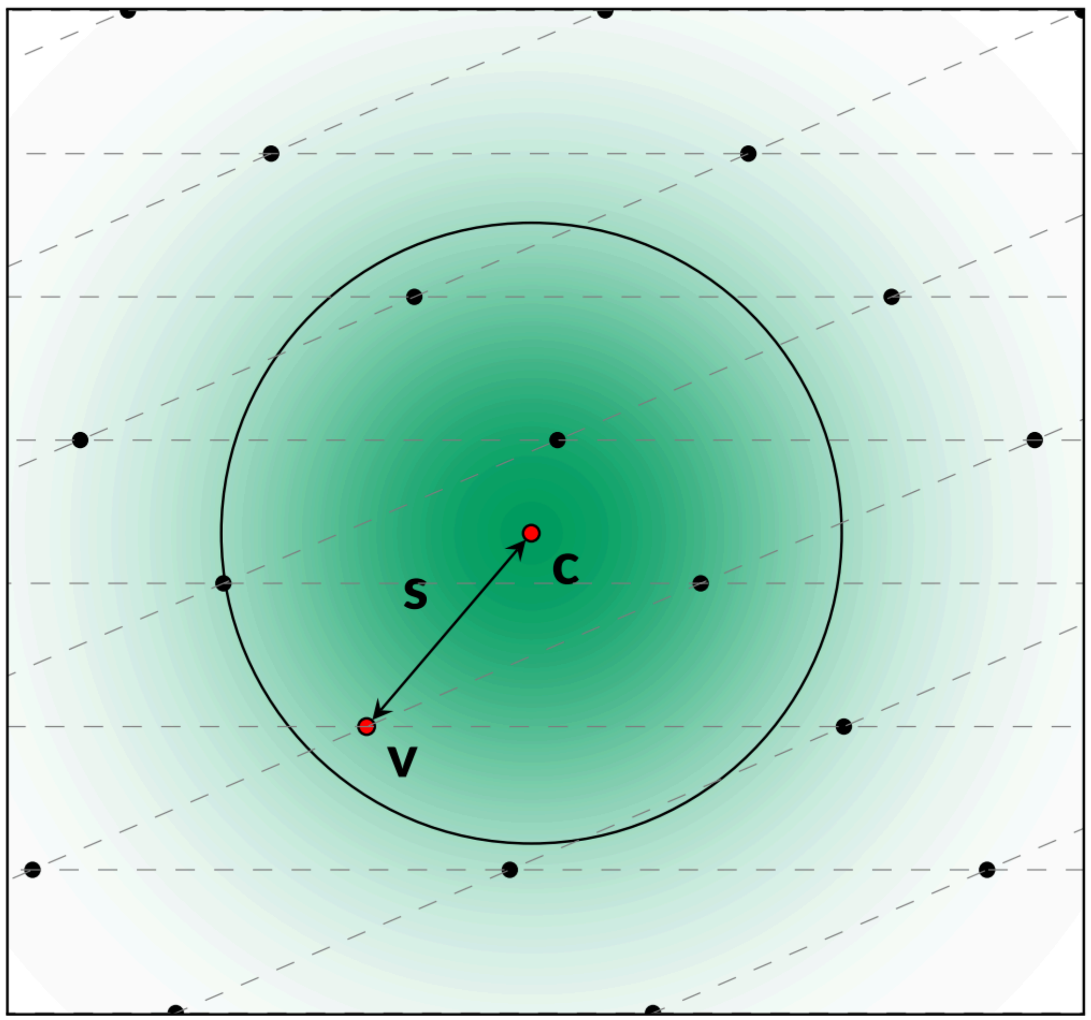

Falcon round I

Advantages

- ☑ Compact
- ☑ Fast
- ☑ GPV framework proved secure in the ROM and QROM (Boneh et al. ASIACRYPT 2011)

Falcon round I

Advantages

- ☑ Compact
- ☑ Fast
- ☑ GPV framework proved secure in the ROM and QROM (Boneh et al. ASIACRYPT 2011)

→ Selected to round II and later round III

Falcon and implementation attacks

Limitations

?
- ☐ Non Trivial to understand and implement
- ☐ Floating point arithmetic
- ☐ Side channel resistance not very studied

Side channel attacks targeting Gaussians

- Espitau et al. SAC'2016
- Fouque et al EUROCRYPT'2020

Implementation issues

Portability issues:
- Floating point arithmetics
- Many subtleties for implementing the Gaussian sampler

Falcon and implementation attacks

Limitations

?
- ☐ Non Trivial to understand and implement
- ☐ Floating point arithmetic
- ☐ Side channel resistance not very studied

Side channel attacks targeting Gaussians

- Espitau et al. SAC'2016
- Fouque et al EUROCRYPT'2020

Implementation issues

Portability issues:
- Floating point arithmetics
- Many subtleties for implementing the Gaussian sampler

Need for timing protection

« Constant time » is a confusing term

Constant time does not mean constant execution time

« Constant time » is a confusing term

Constant time does not mean constant execution time

« Constant time »

The execution time **does not depend on the private key.**

➡ Not necessarily constant !

« Constant time » is a confusing term

Constant time does not mean constant execution time

« Constant time »

The execution time **does not depend on the private key.**

➡ Not necessarily constant !

Better say **isochronous** ?

« Constant time » is a confusing term

Constant time does not mean constant execution time

« Constant time »

The execution time **does not depend on the private key.**

➡ Not necessarily constant !

Better say **isochronous** ?

Assumption

$+, -, \times, /$ Constant time on integers

Contributions of this work

- ☑ Integer arithmetic for the Gaussian sampling for Falcon
- ☑ Theoretically studied isochrony

- ☑ Test suite : Statistically Acceptable Gaussians (SAGA)
- ☑ Implementations

What is not isochronous in Falcon?

Sign(m,sk)

- Compute $\mathbf{c}$ such that $\mathbf{cA} = H(m)$
- $\mathbf{v} \leftarrow$ a vector in $\Lambda(\mathbf{B})$ close to $\mathbf{c}$
- $\mathbf{s} \leftarrow \mathbf{c} - \mathbf{v}$

What is not isochronous in Falcon?

Sign(m,sk)

- Compute $\mathbf{c}$ such that $\mathbf{cA} = H(m)$
- $\mathbf{v} \leftarrow$ a vector in $\Lambda(\mathbf{B})$ close to $\mathbf{c}$
- $\mathbf{s} \leftarrow \mathbf{c} - \mathbf{v}$

What is not isochronous in Falcon?

Sign(m,sk)

- Compute $\mathbf{c}$ such that $\mathbf{cA} = H(m)$
- $\mathbf{v} \leftarrow$ a vector in $\Lambda(\mathbf{B})$ close to $\mathbf{c}$ → ffsampling
- $\mathbf{s} \leftarrow \mathbf{c} - \mathbf{v}$

What is not isochronous in Falcon?

Sign(m,sk)

- Compute $\mathbf{c}$ such that $\mathbf{cA} = H(m)$
- $\mathbf{v} \leftarrow$ a vector in $\Lambda(\mathbf{B})$ close to $\mathbf{c}$
- $\mathbf{s} \leftarrow \mathbf{c} - \mathbf{v}$

→ ffsampling → Gaussian Sampling over $\mathbb{Z}$

What is not isochronous in Falcon?

Sign(m,sk)

- Compute $\mathbf{c}$ such that $\mathbf{cA} = H(m)$
- $\mathbf{v} \leftarrow$ a vector in $\Lambda(\mathbf{B})$ close to $\mathbf{c}$
- $\mathbf{s} \leftarrow \mathbf{c} - \mathbf{v}$

→ ffsampling → Gaussian Sampling over $\mathbb{Z}$

Except Gaussian sampling, other operations do not use conditional branching

Isochronous Gaussian sampling

Some literature on Gaussian Samplers:

- Sinha Roy, Vercauteren and Verbauwhede SAC'13
- Hulsing, Lange and Smeets PKC'18
- Micciancio and Walter CRYPTO'17
- Karmakar et al. DAC IEEE'19

Isochronous Gaussian sampling

Some literature on Gaussian Samplers:

- Sinha Roy, Vercauteren and Verbauwhede SAC'13
- Hulsing, Lange and Smeets PKC'18
- Micciancio and Walter CRYPTO'17
- Karmakar et al. DAC IEEE'19

Here we present a simple alternative dedicated to Falcon

Isochronous Gaussian sampling

Some literature on Gaussian Samplers:

- Sinha Roy, Vercauteren and Verbauwhede SAC'13
- Hulsing, Lange and Smeets PKC'18
- Micciancio and Walter CRYPTO'17
- Karmakar et al. DAC IEEE'19

Here we present a simple alternative dedicated to Falcon

Idea

Construct a distribution that looks somewhat like a Gaussian but is not statistically close, and use rejection sampling to correct the discrepancy.

The sampling distribution

$1.31 = \sigma_{min} \leq \sigma \leq \sigma_0 = 1.82$ $\mu \in [0,1)$

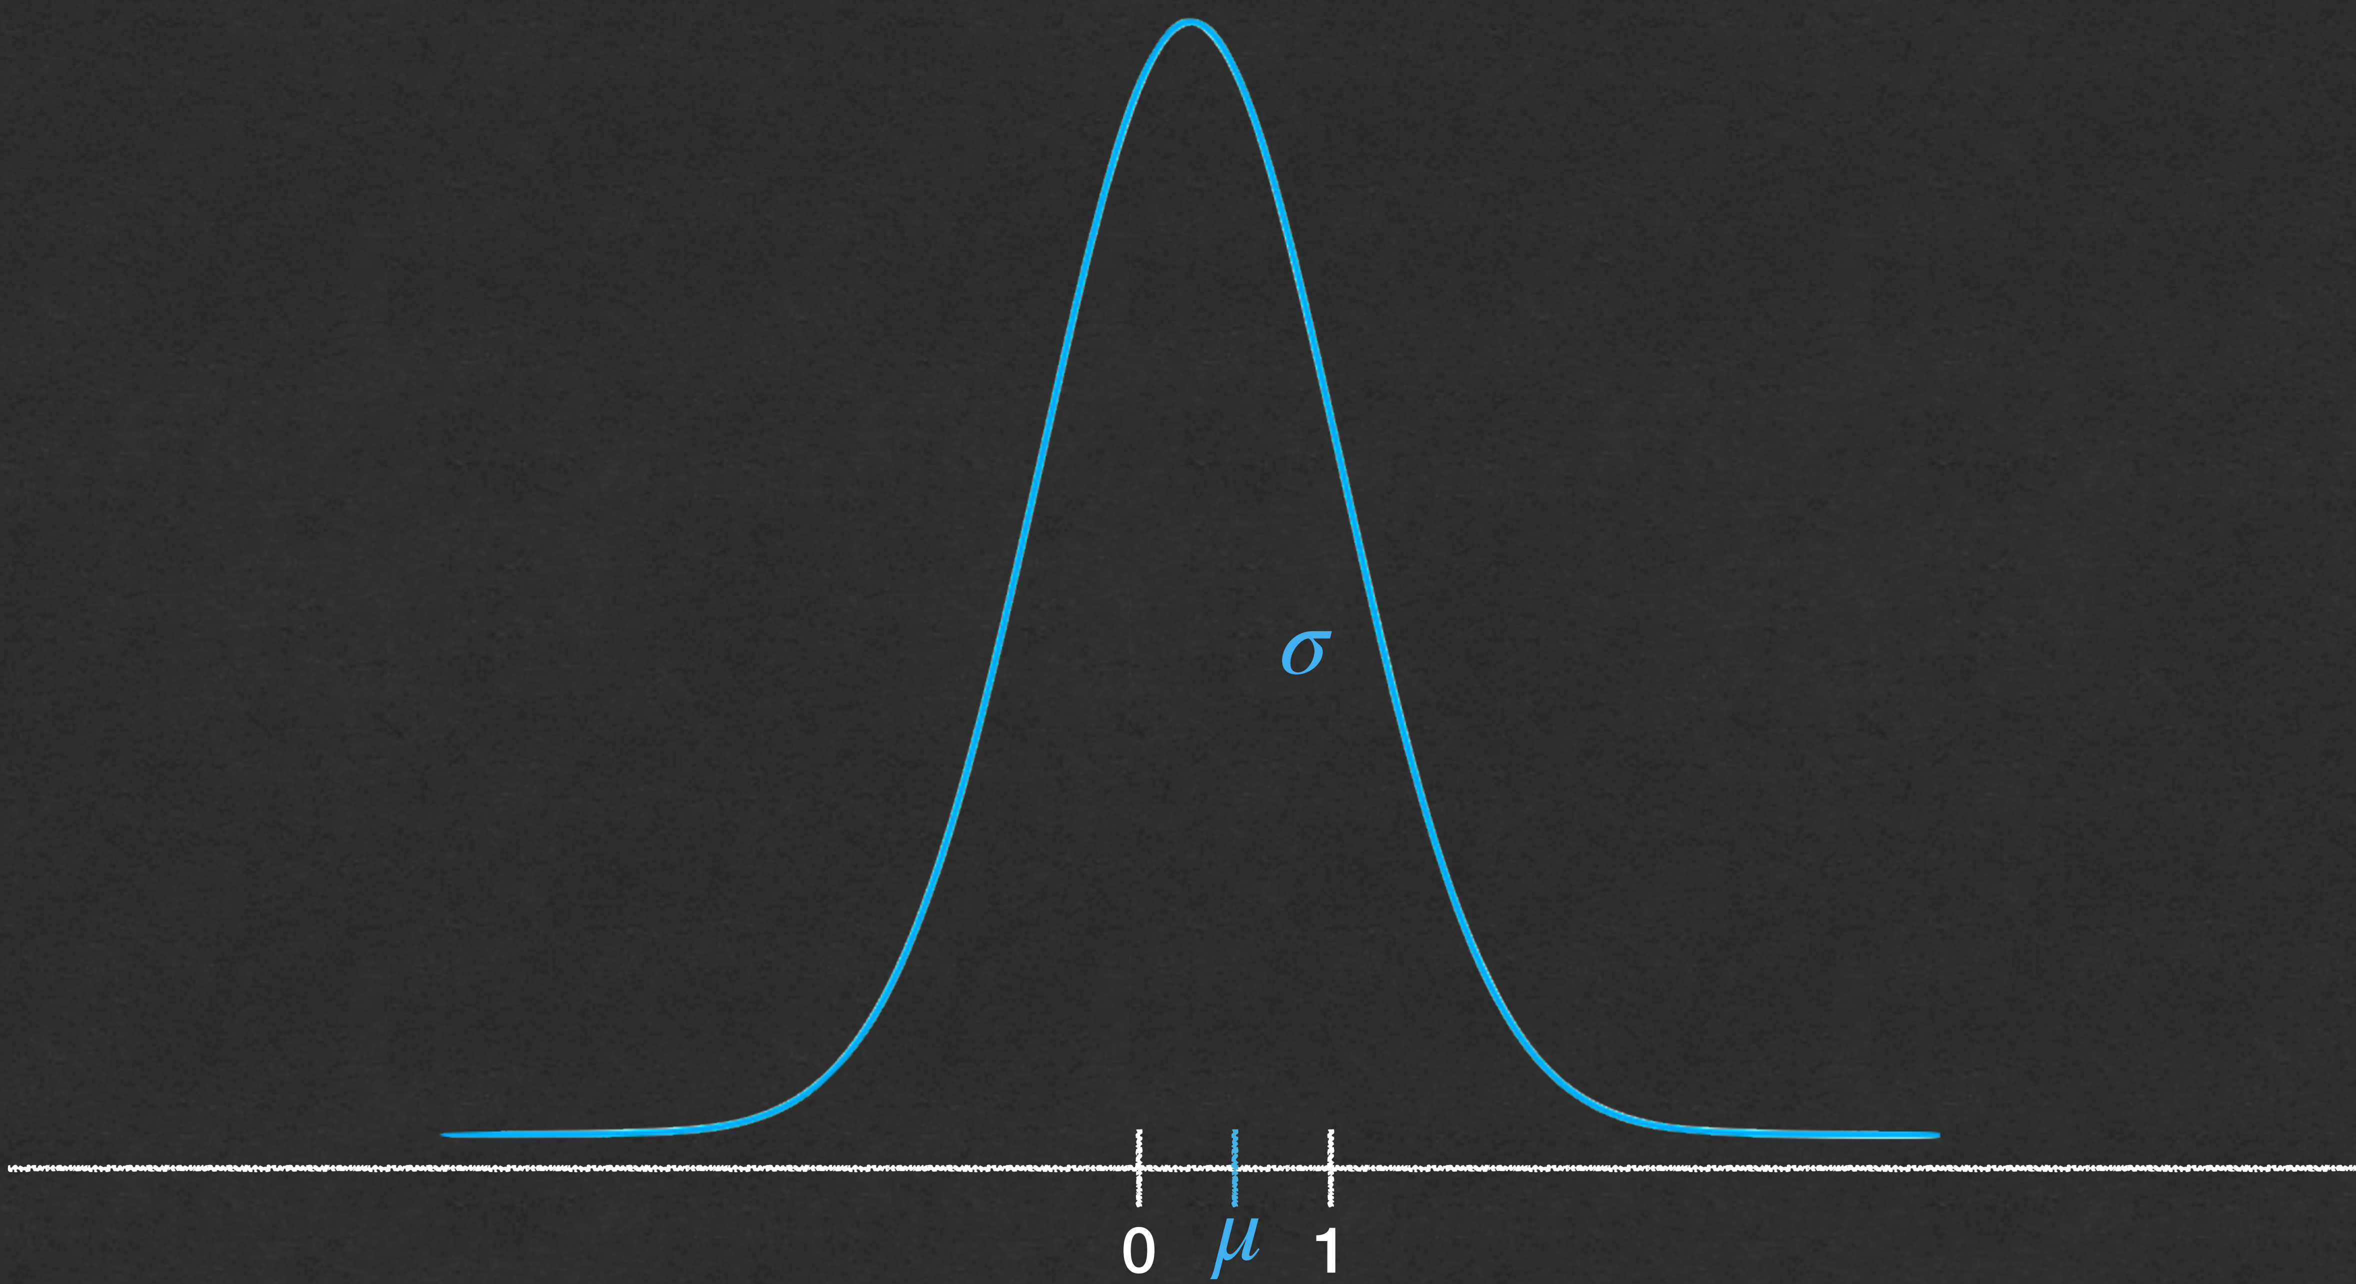

The sampling distribution

$1.31 = \sigma_{min} \leq \sigma \leq \sigma_0 = 1.82$

$\mu \in [0,1)$

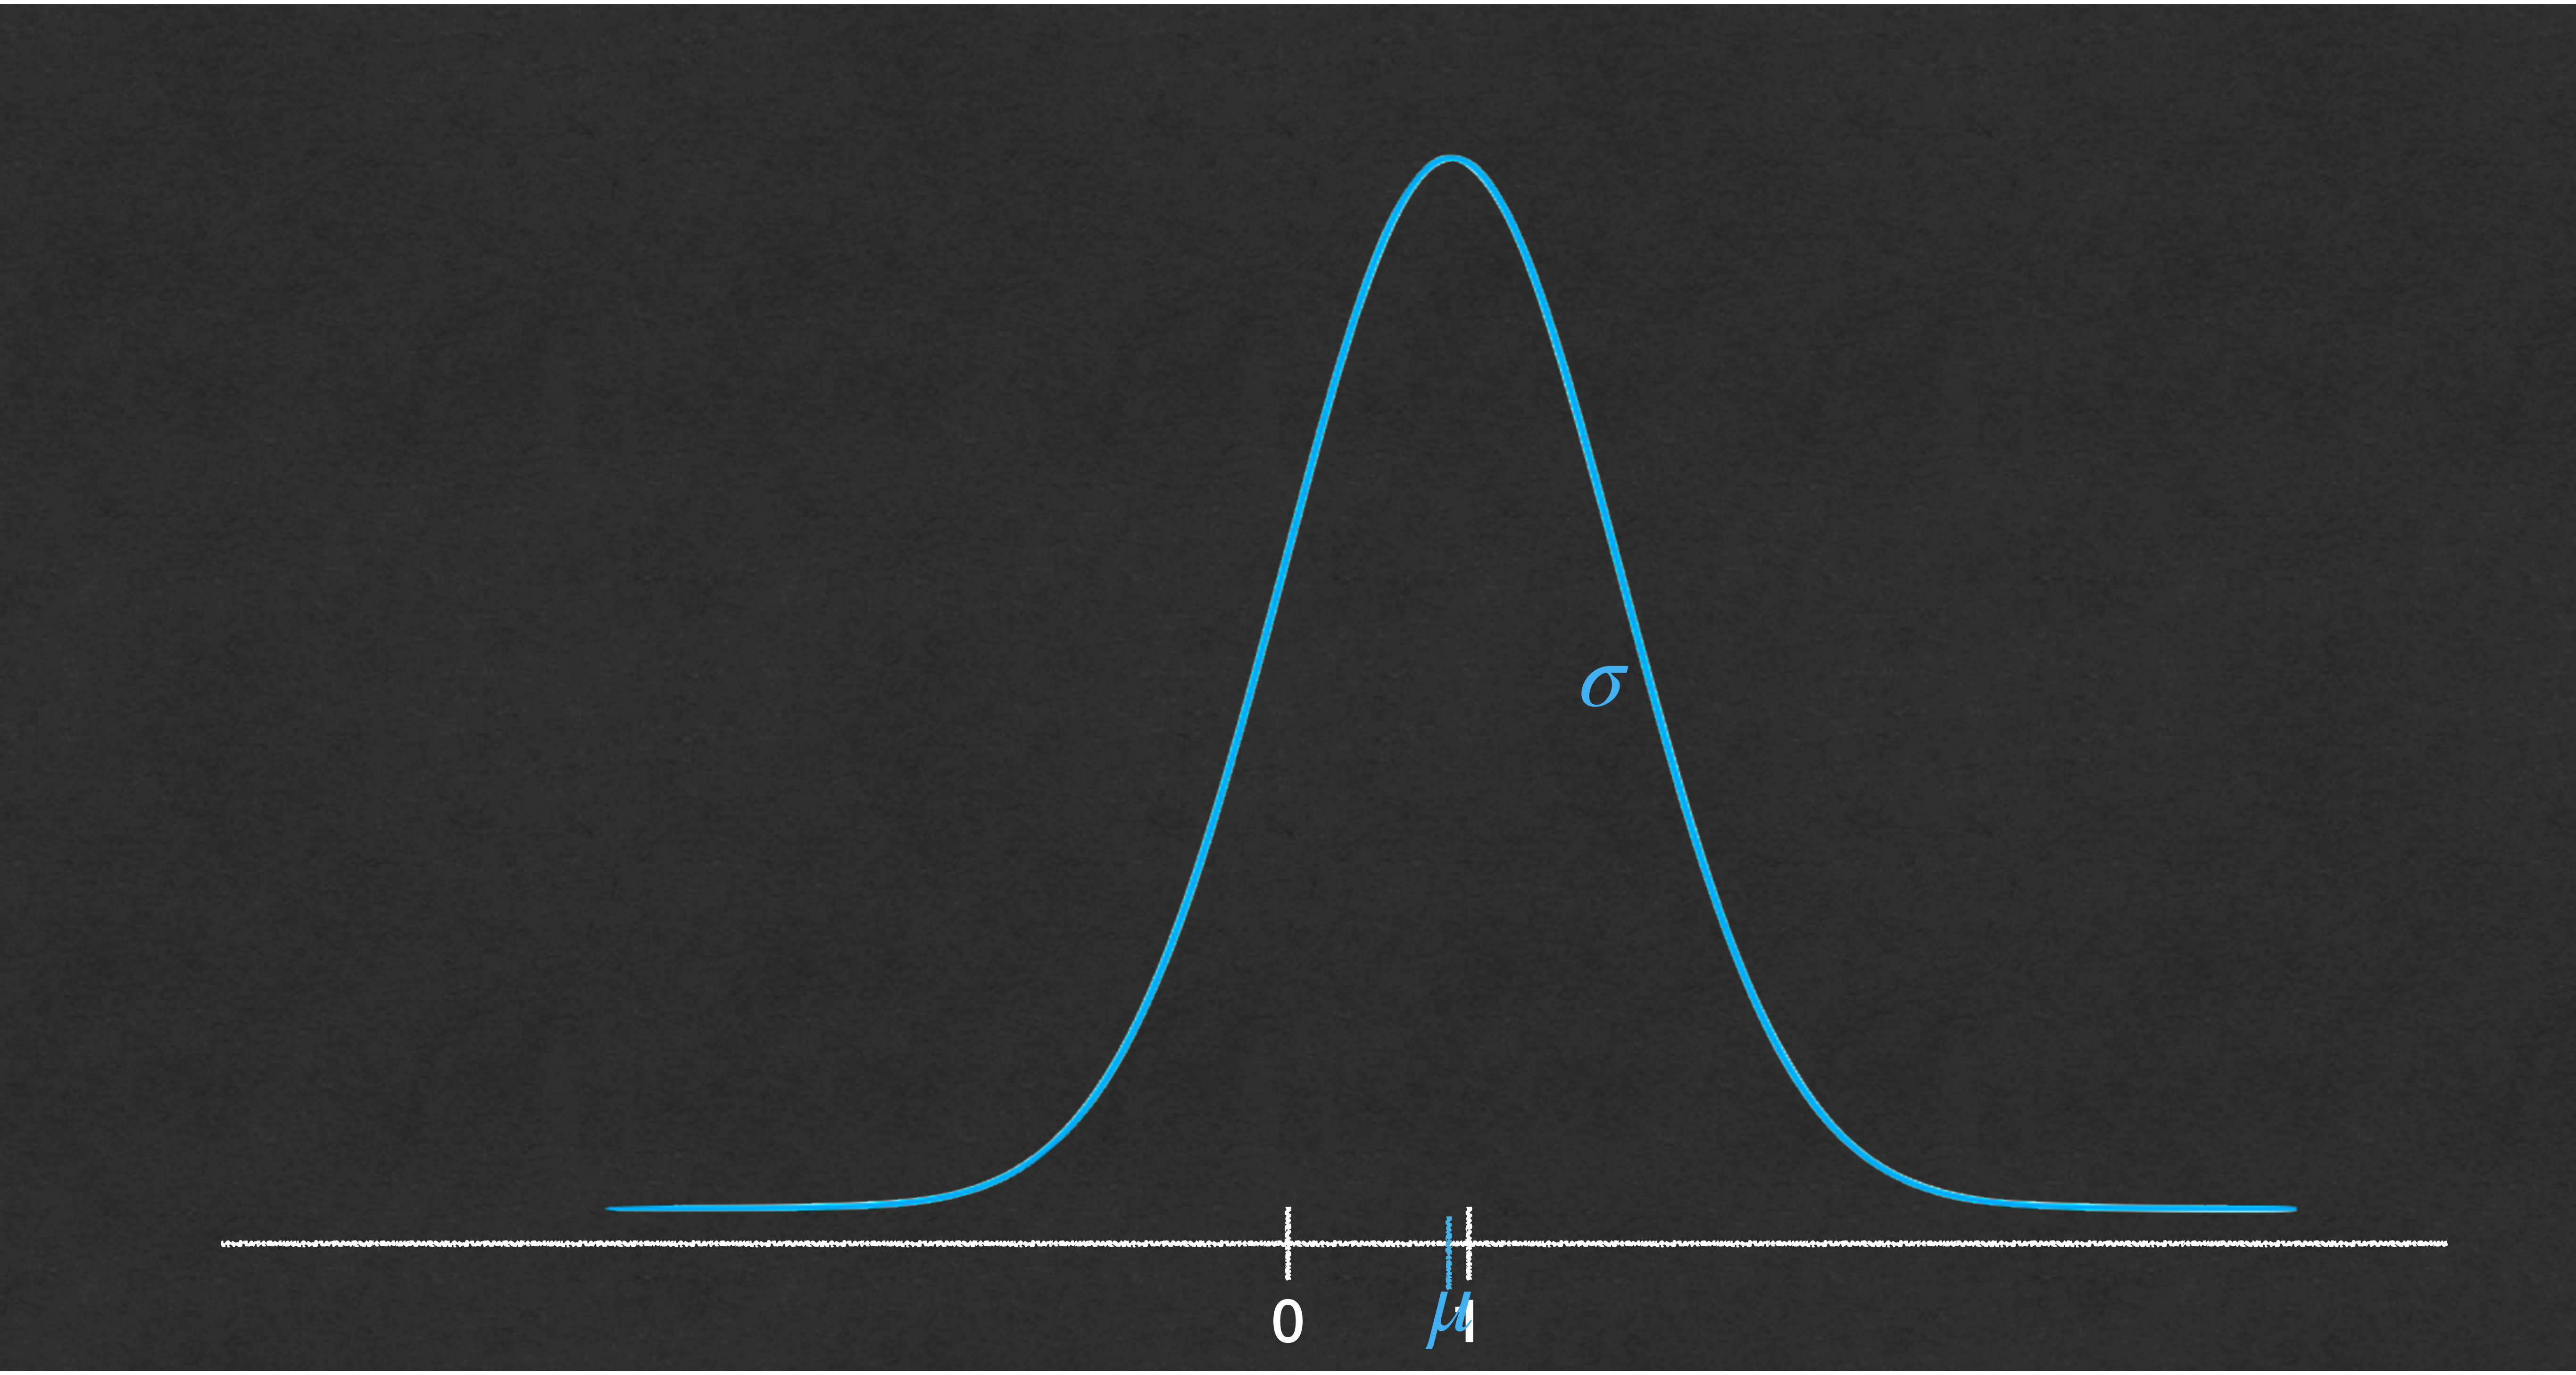

The sampling distribution

$1.31 = \sigma_{min} \leq \sigma \leq \sigma_0 = 1.82$

$\mu \in [0,1)$

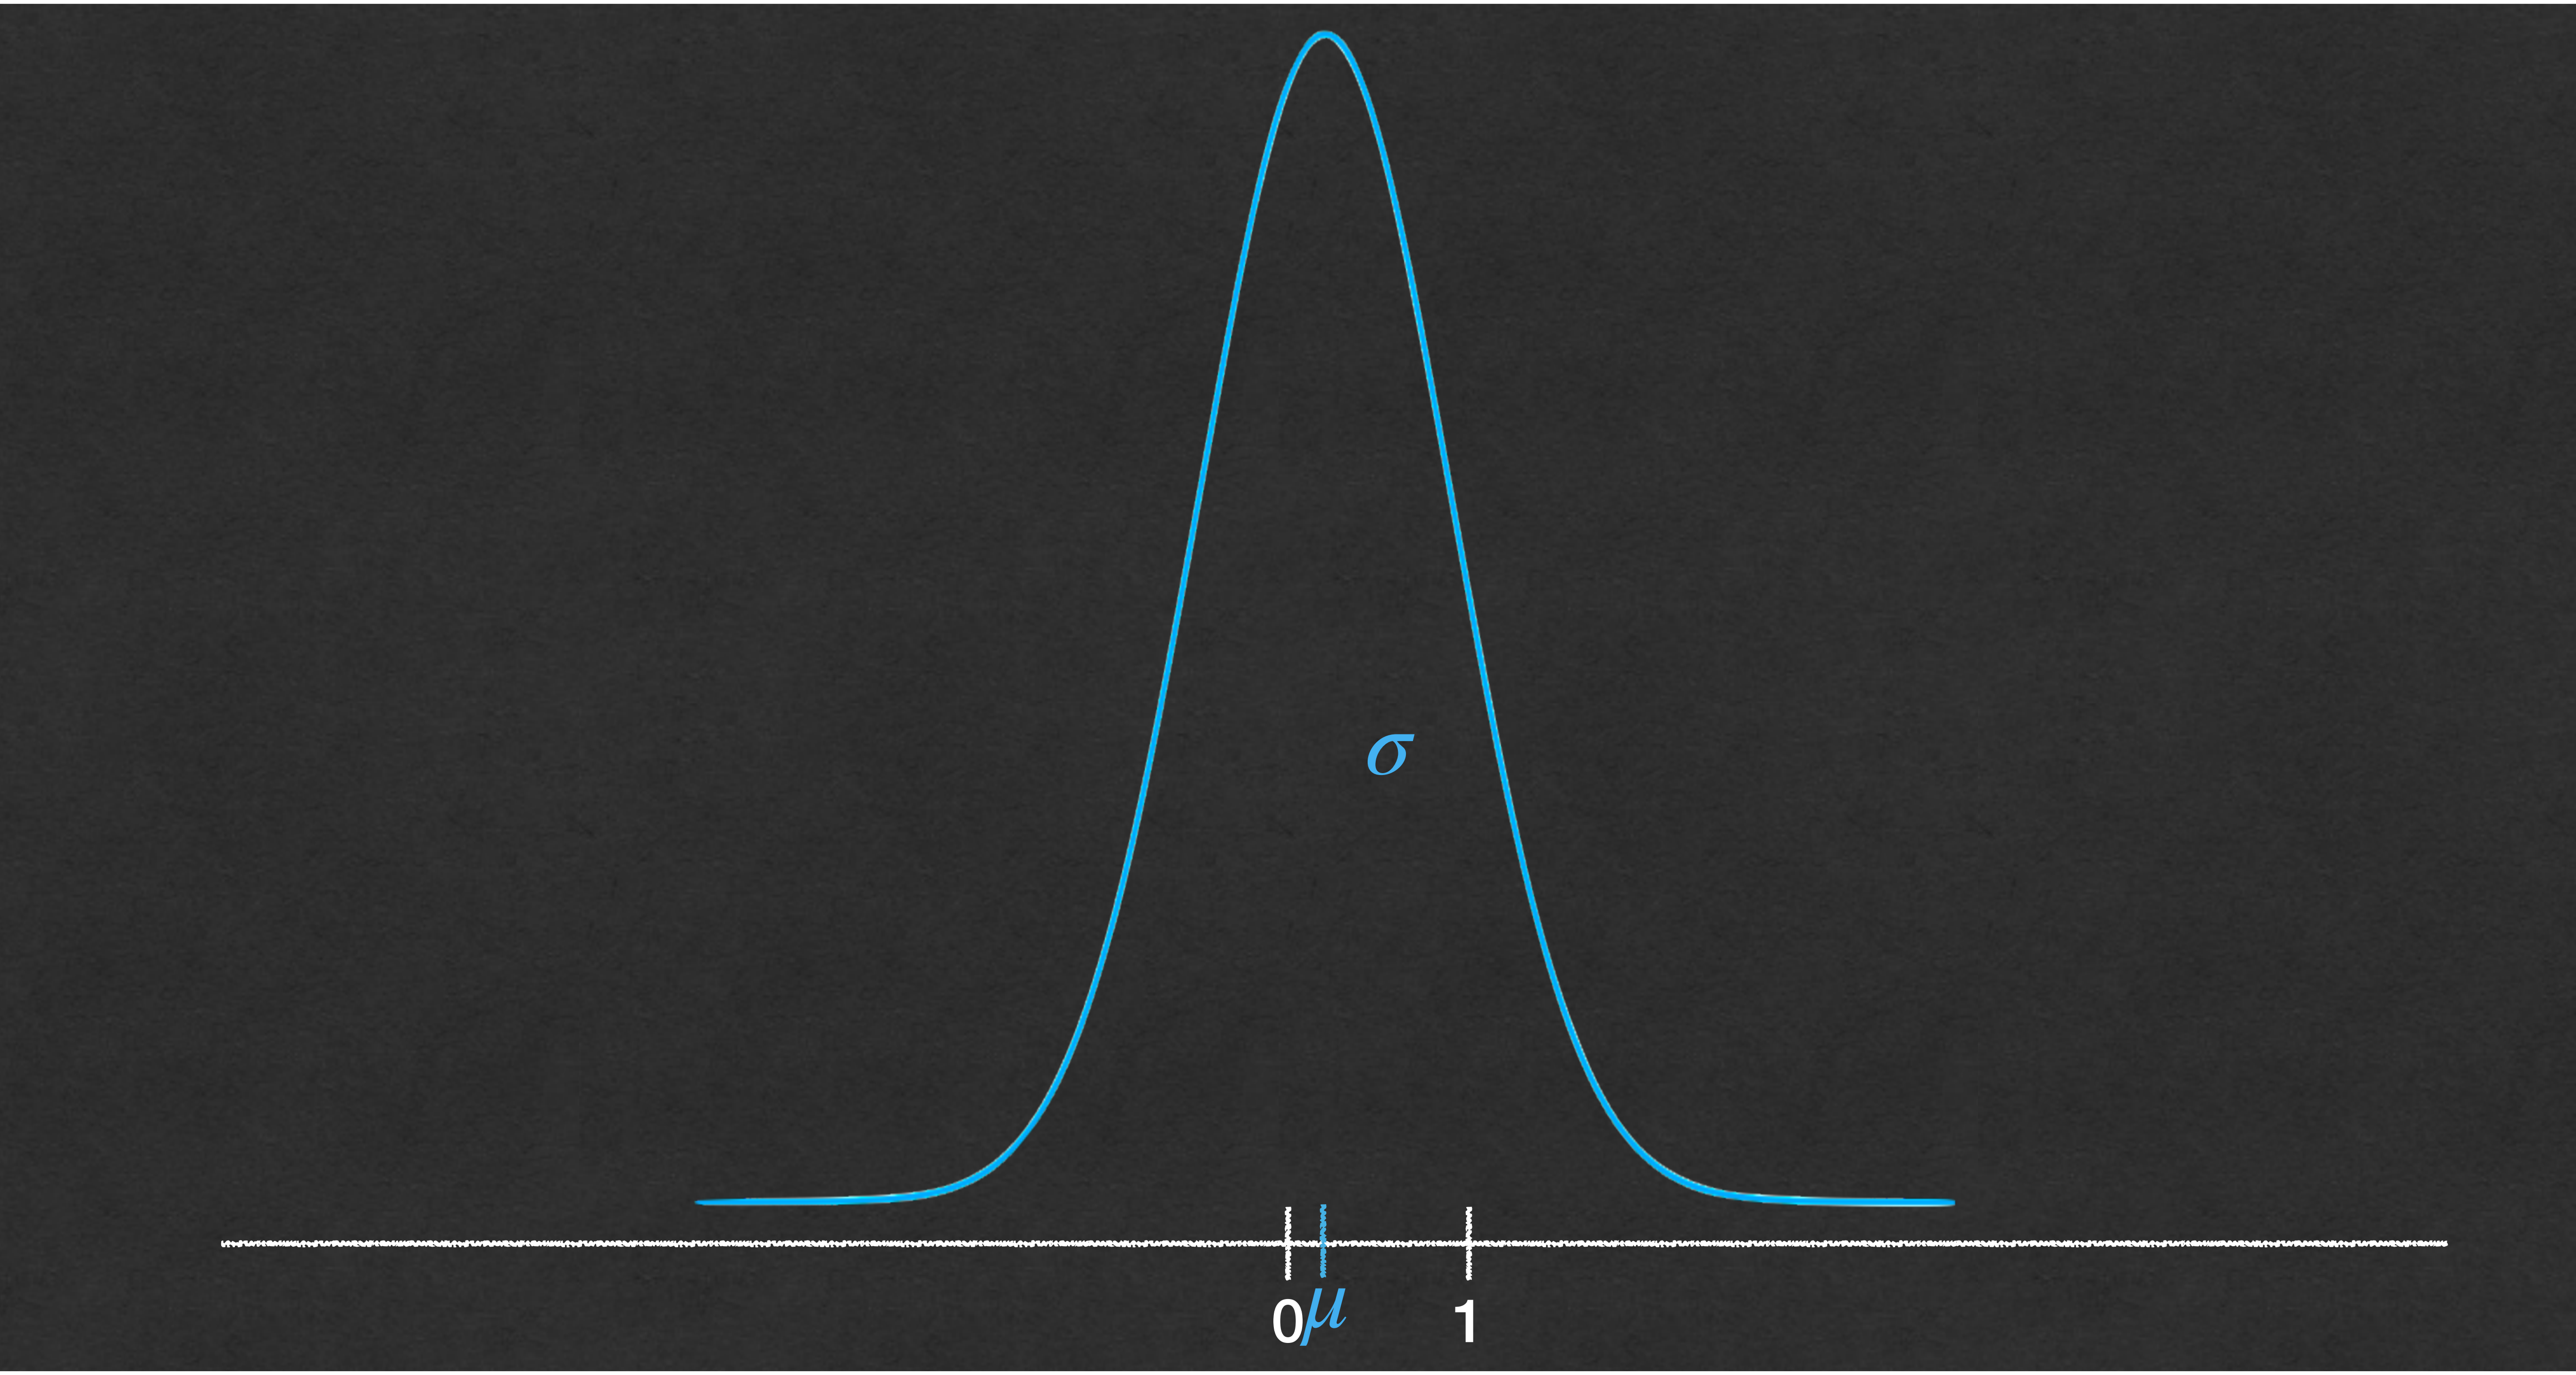

The technique

1 Draw an element z_0 from a centered half Gaussian of standard deviation σ_0

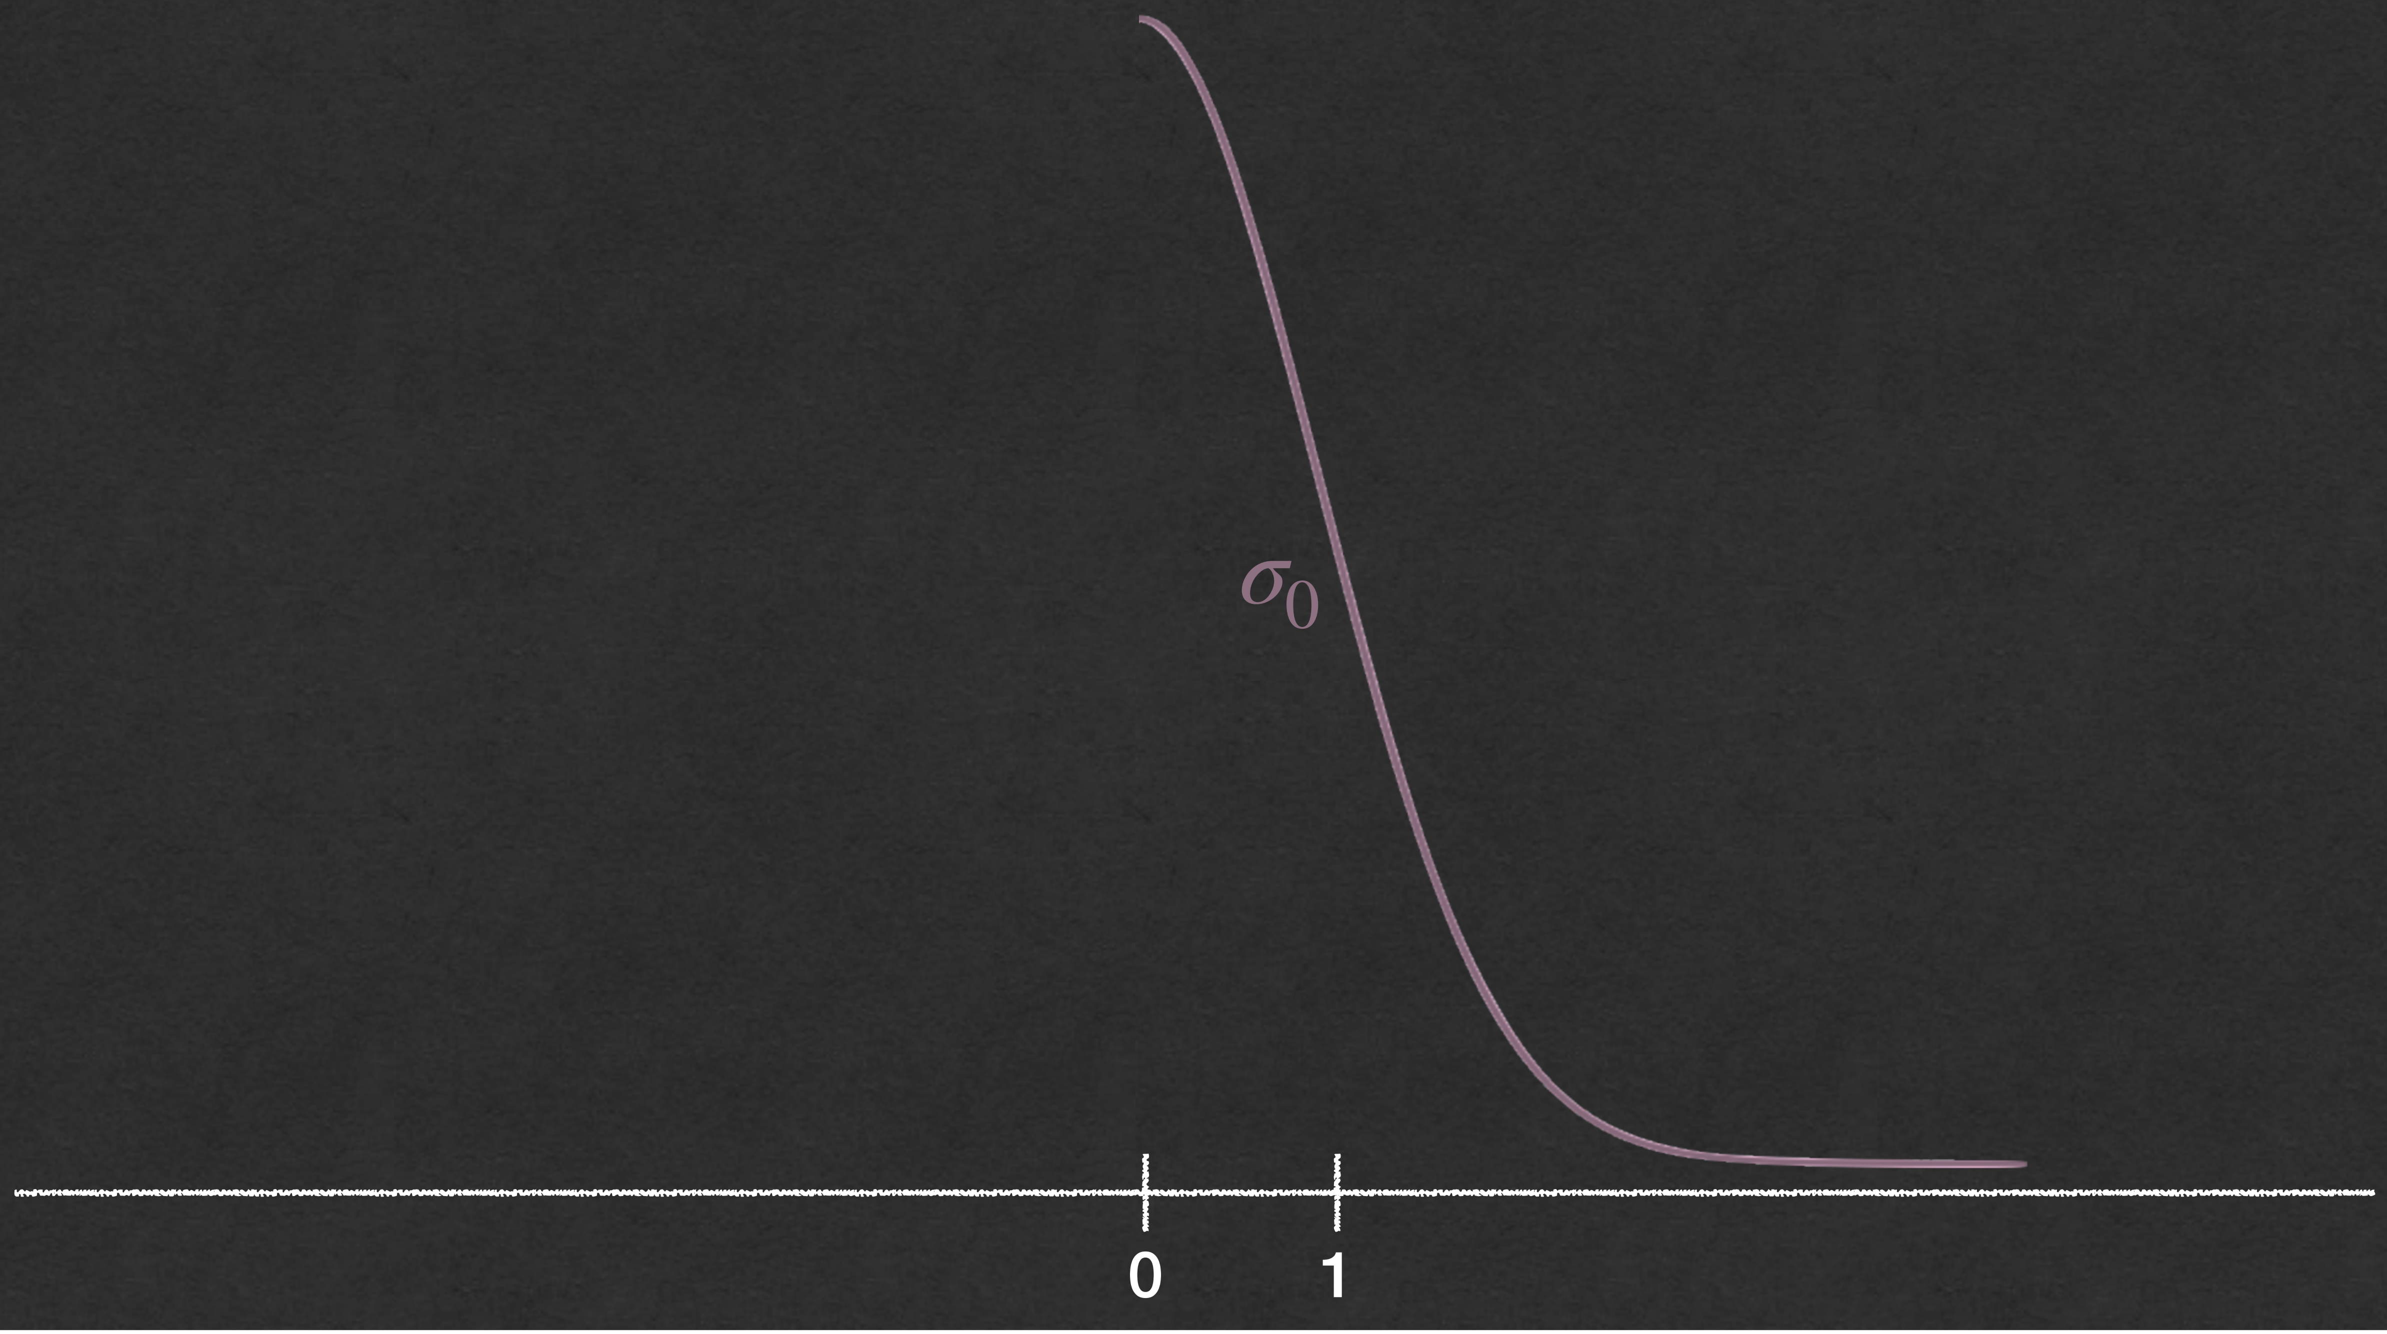

The technique

2 Draw b uniformly at random in {0,1} and compute $z \leftarrow (2b-1) \cdot z_0 + b$

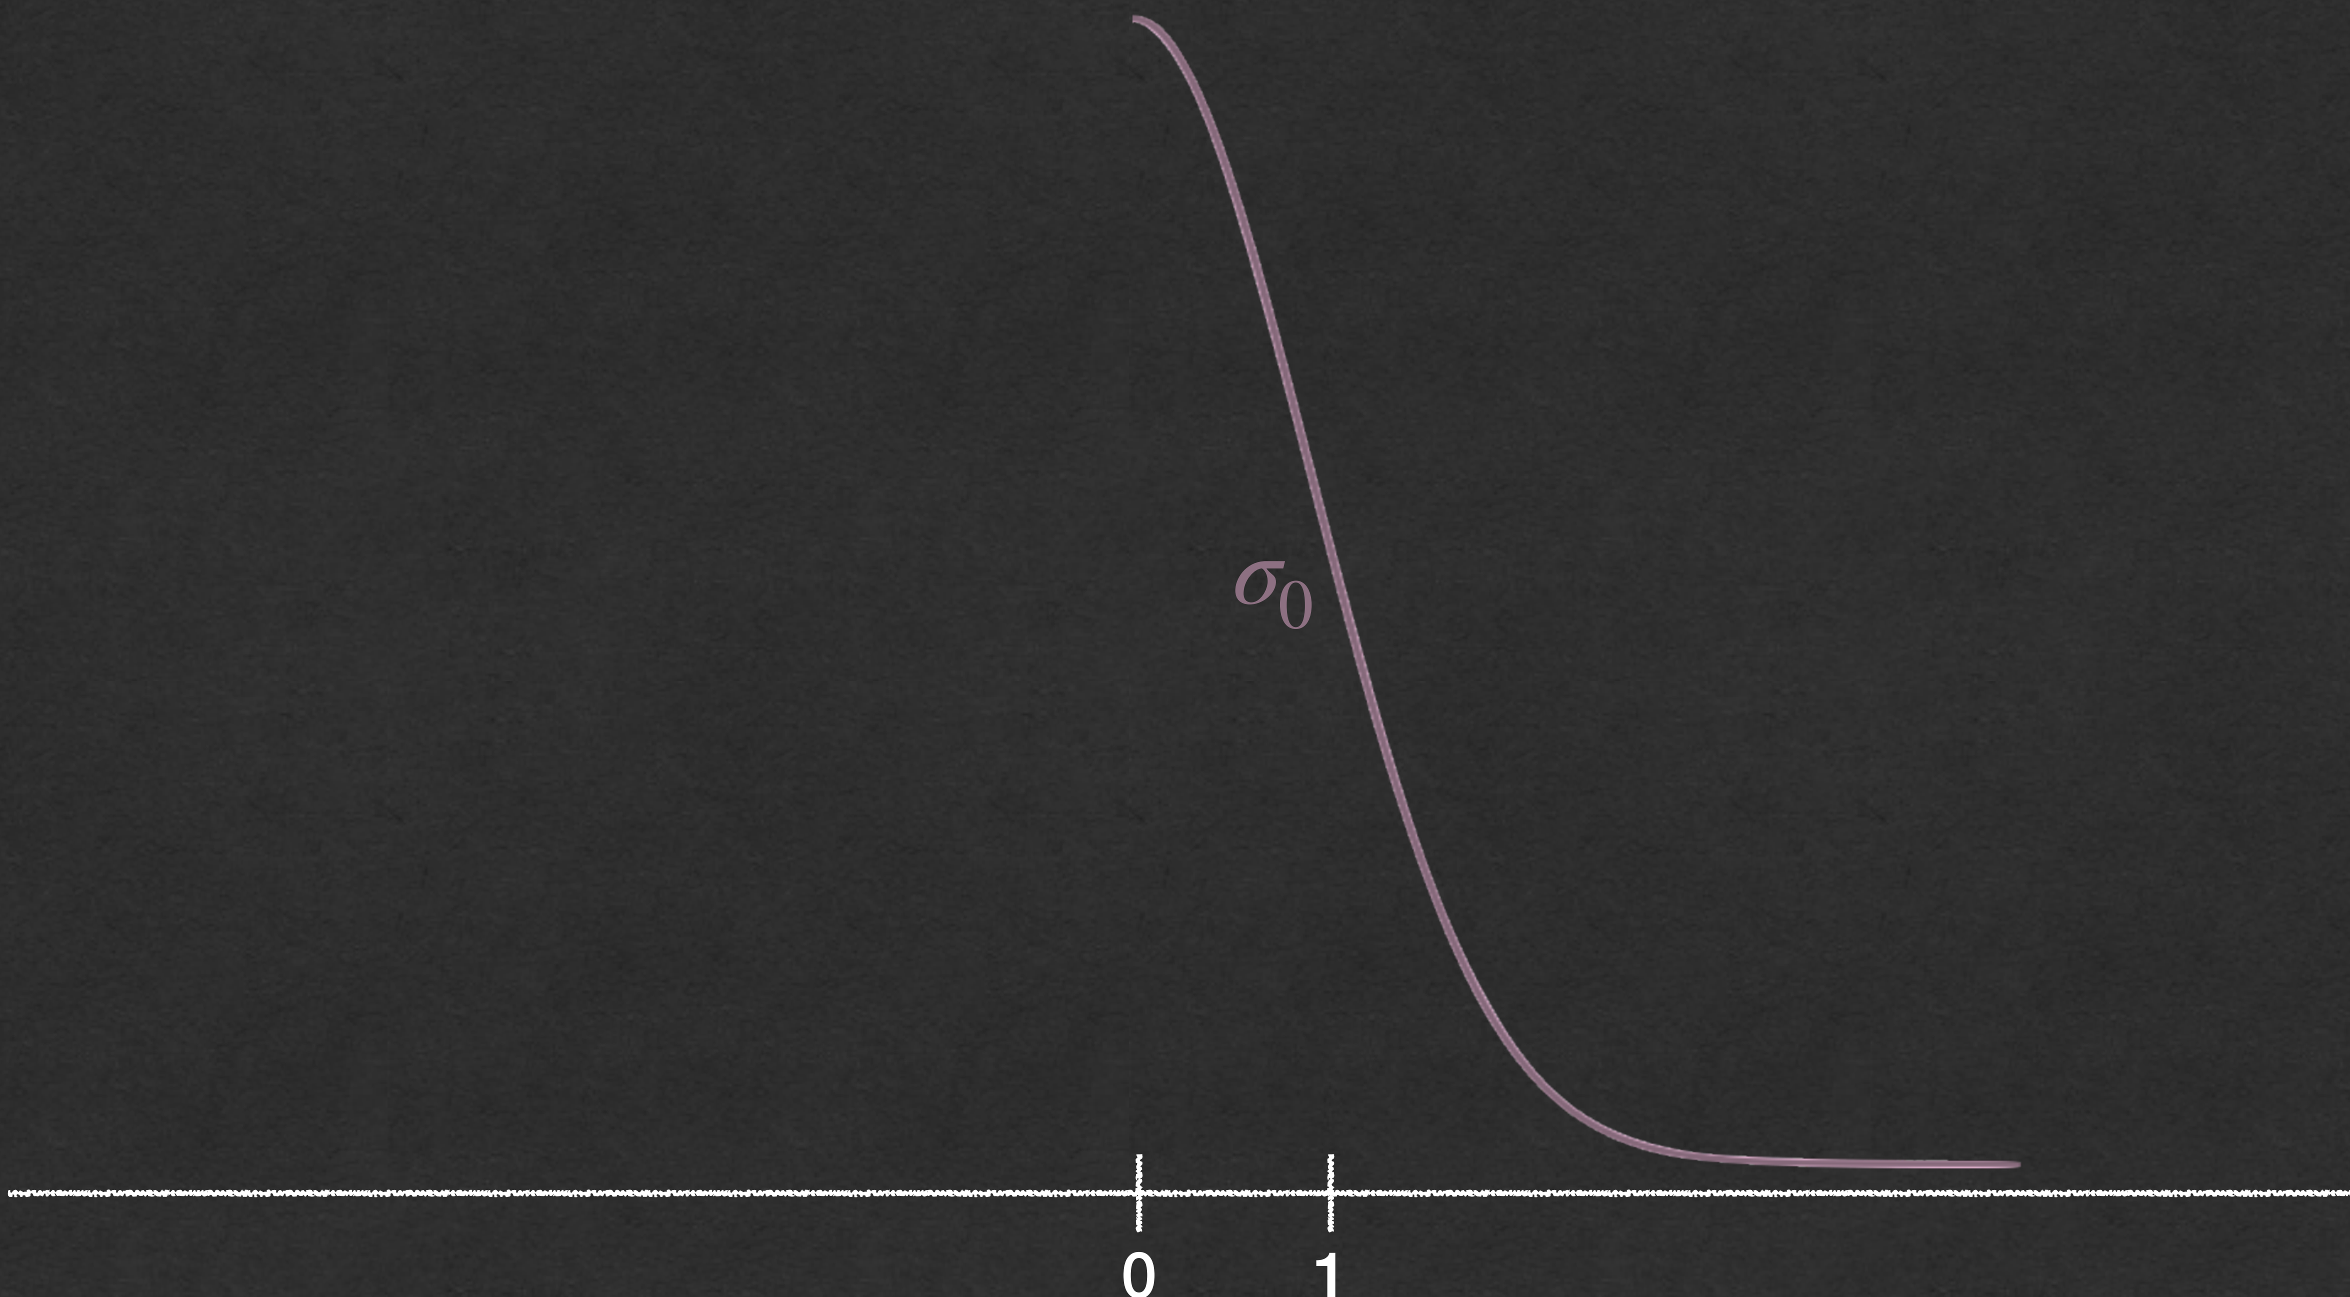

The technique

2 Draw b uniformly at random in {0,1} and compute $z \leftarrow (2b-1) \cdot z_0 + b$

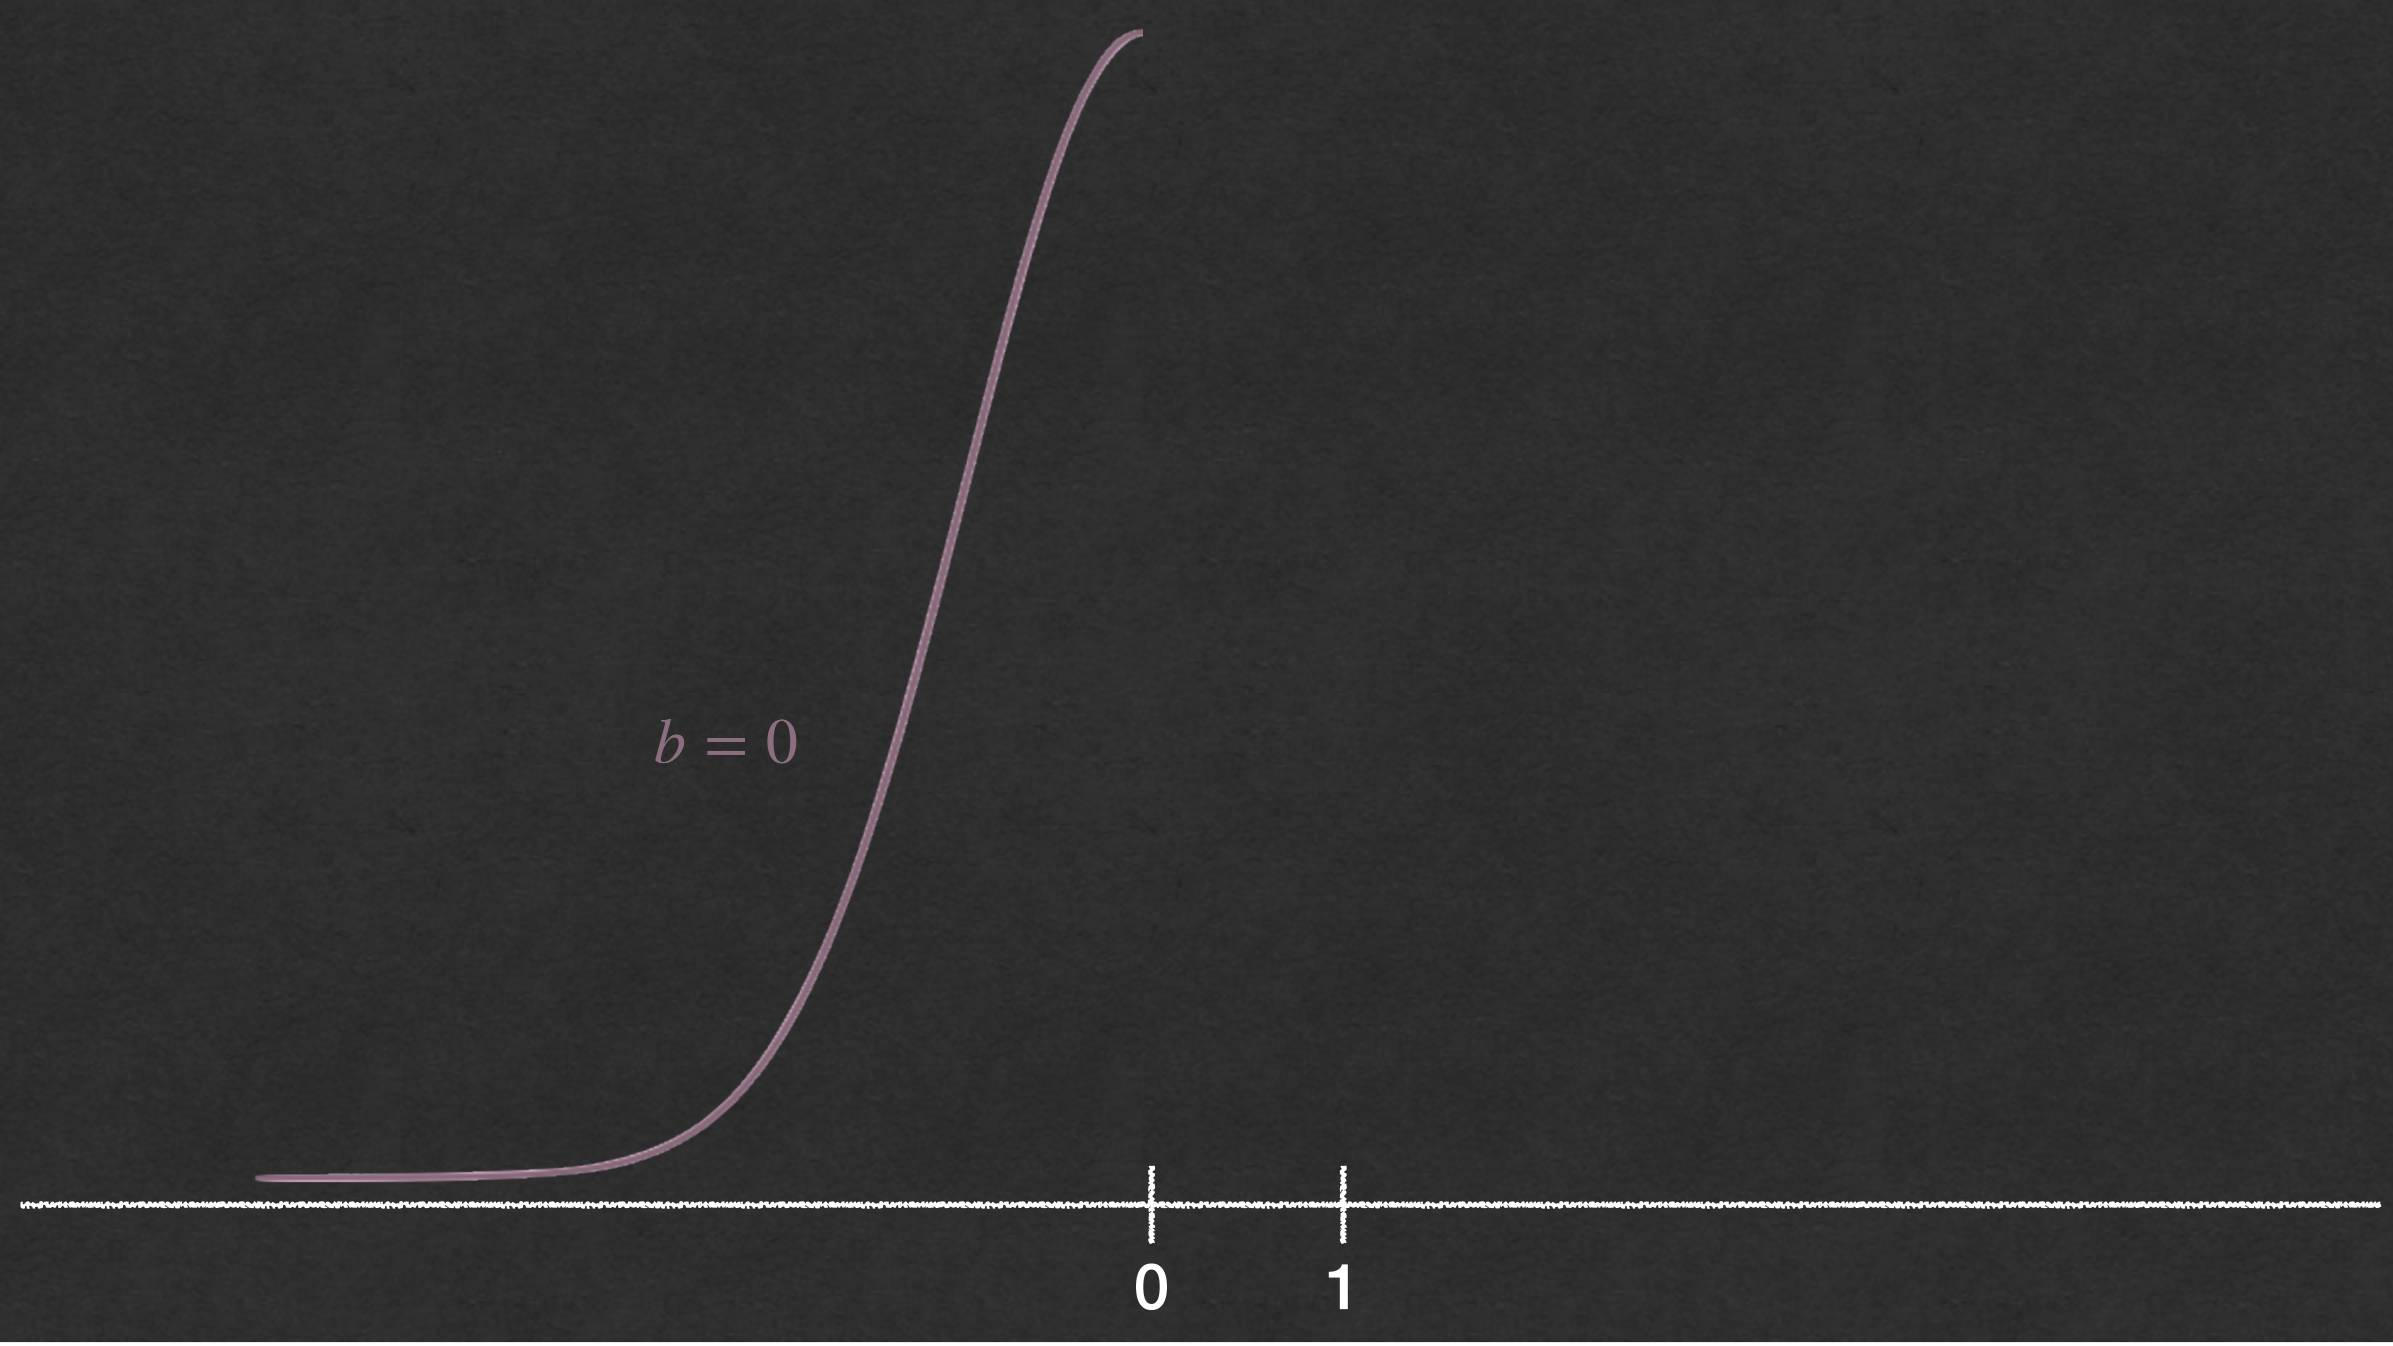

The technique

2 Draw b uniformly at random in {0,1} and compute $z \leftarrow (2b-1) \cdot z_0 + b$

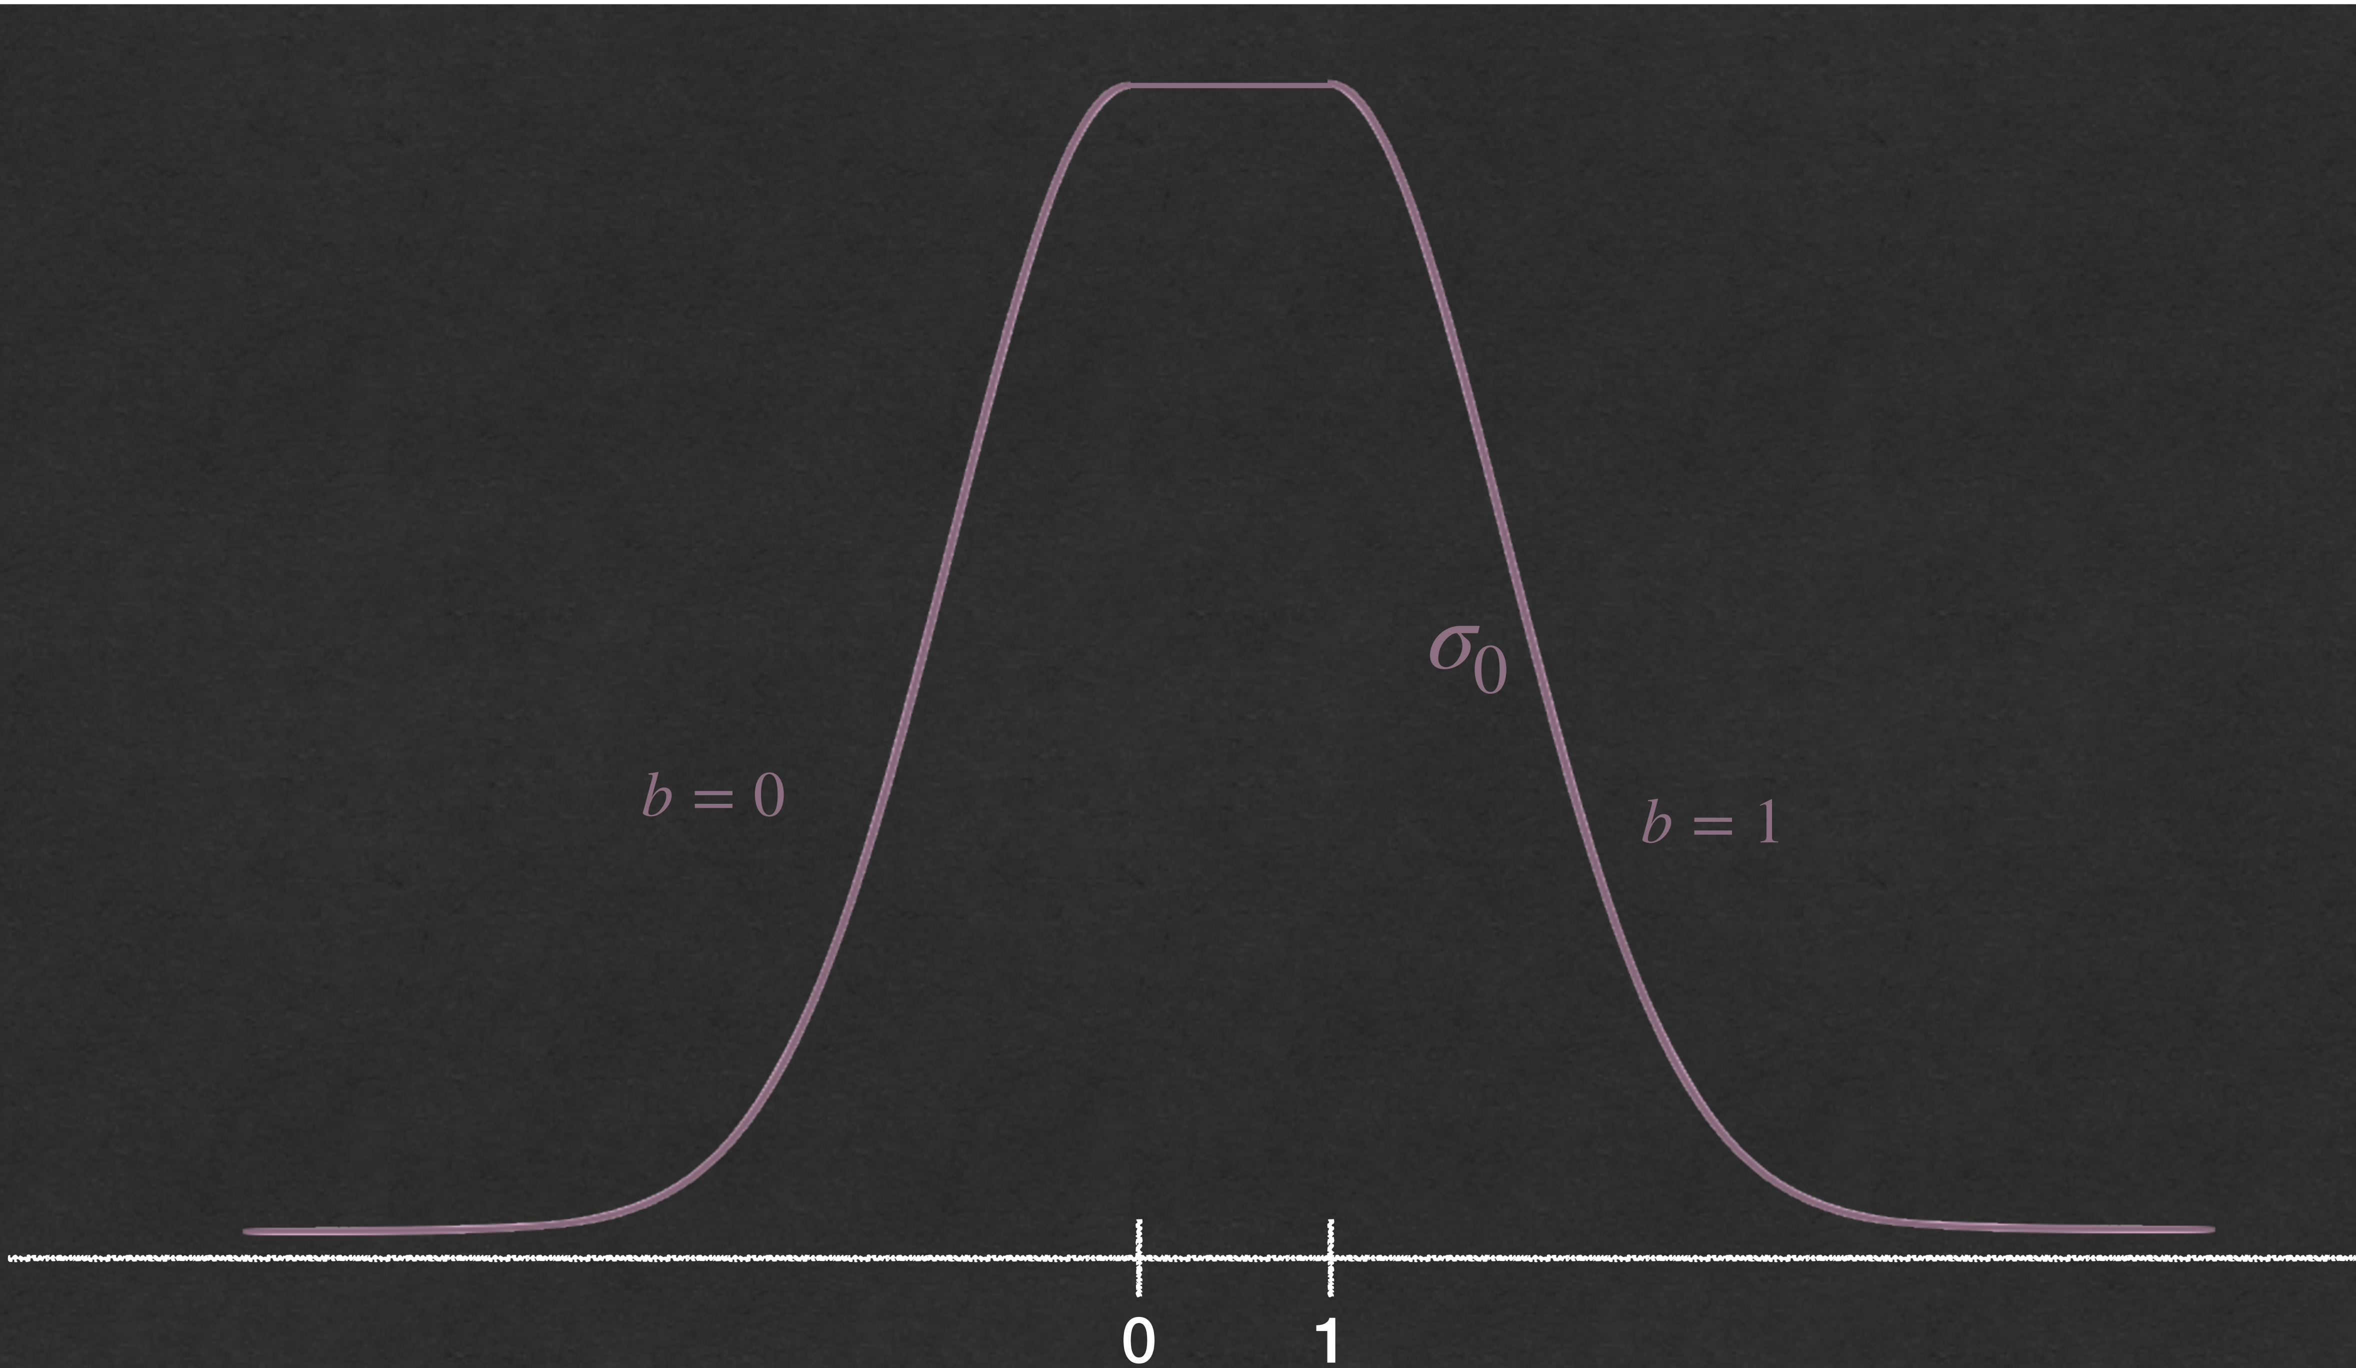

The technique

3 Rejection Sampling (Lyubashevsky EC 2012) Accept with probability $P_{\text{accept}} \propto \dfrac{D_{\sigma,\mu}(z)}{G_{\mathbb{Z},\sigma_0}(z)}$

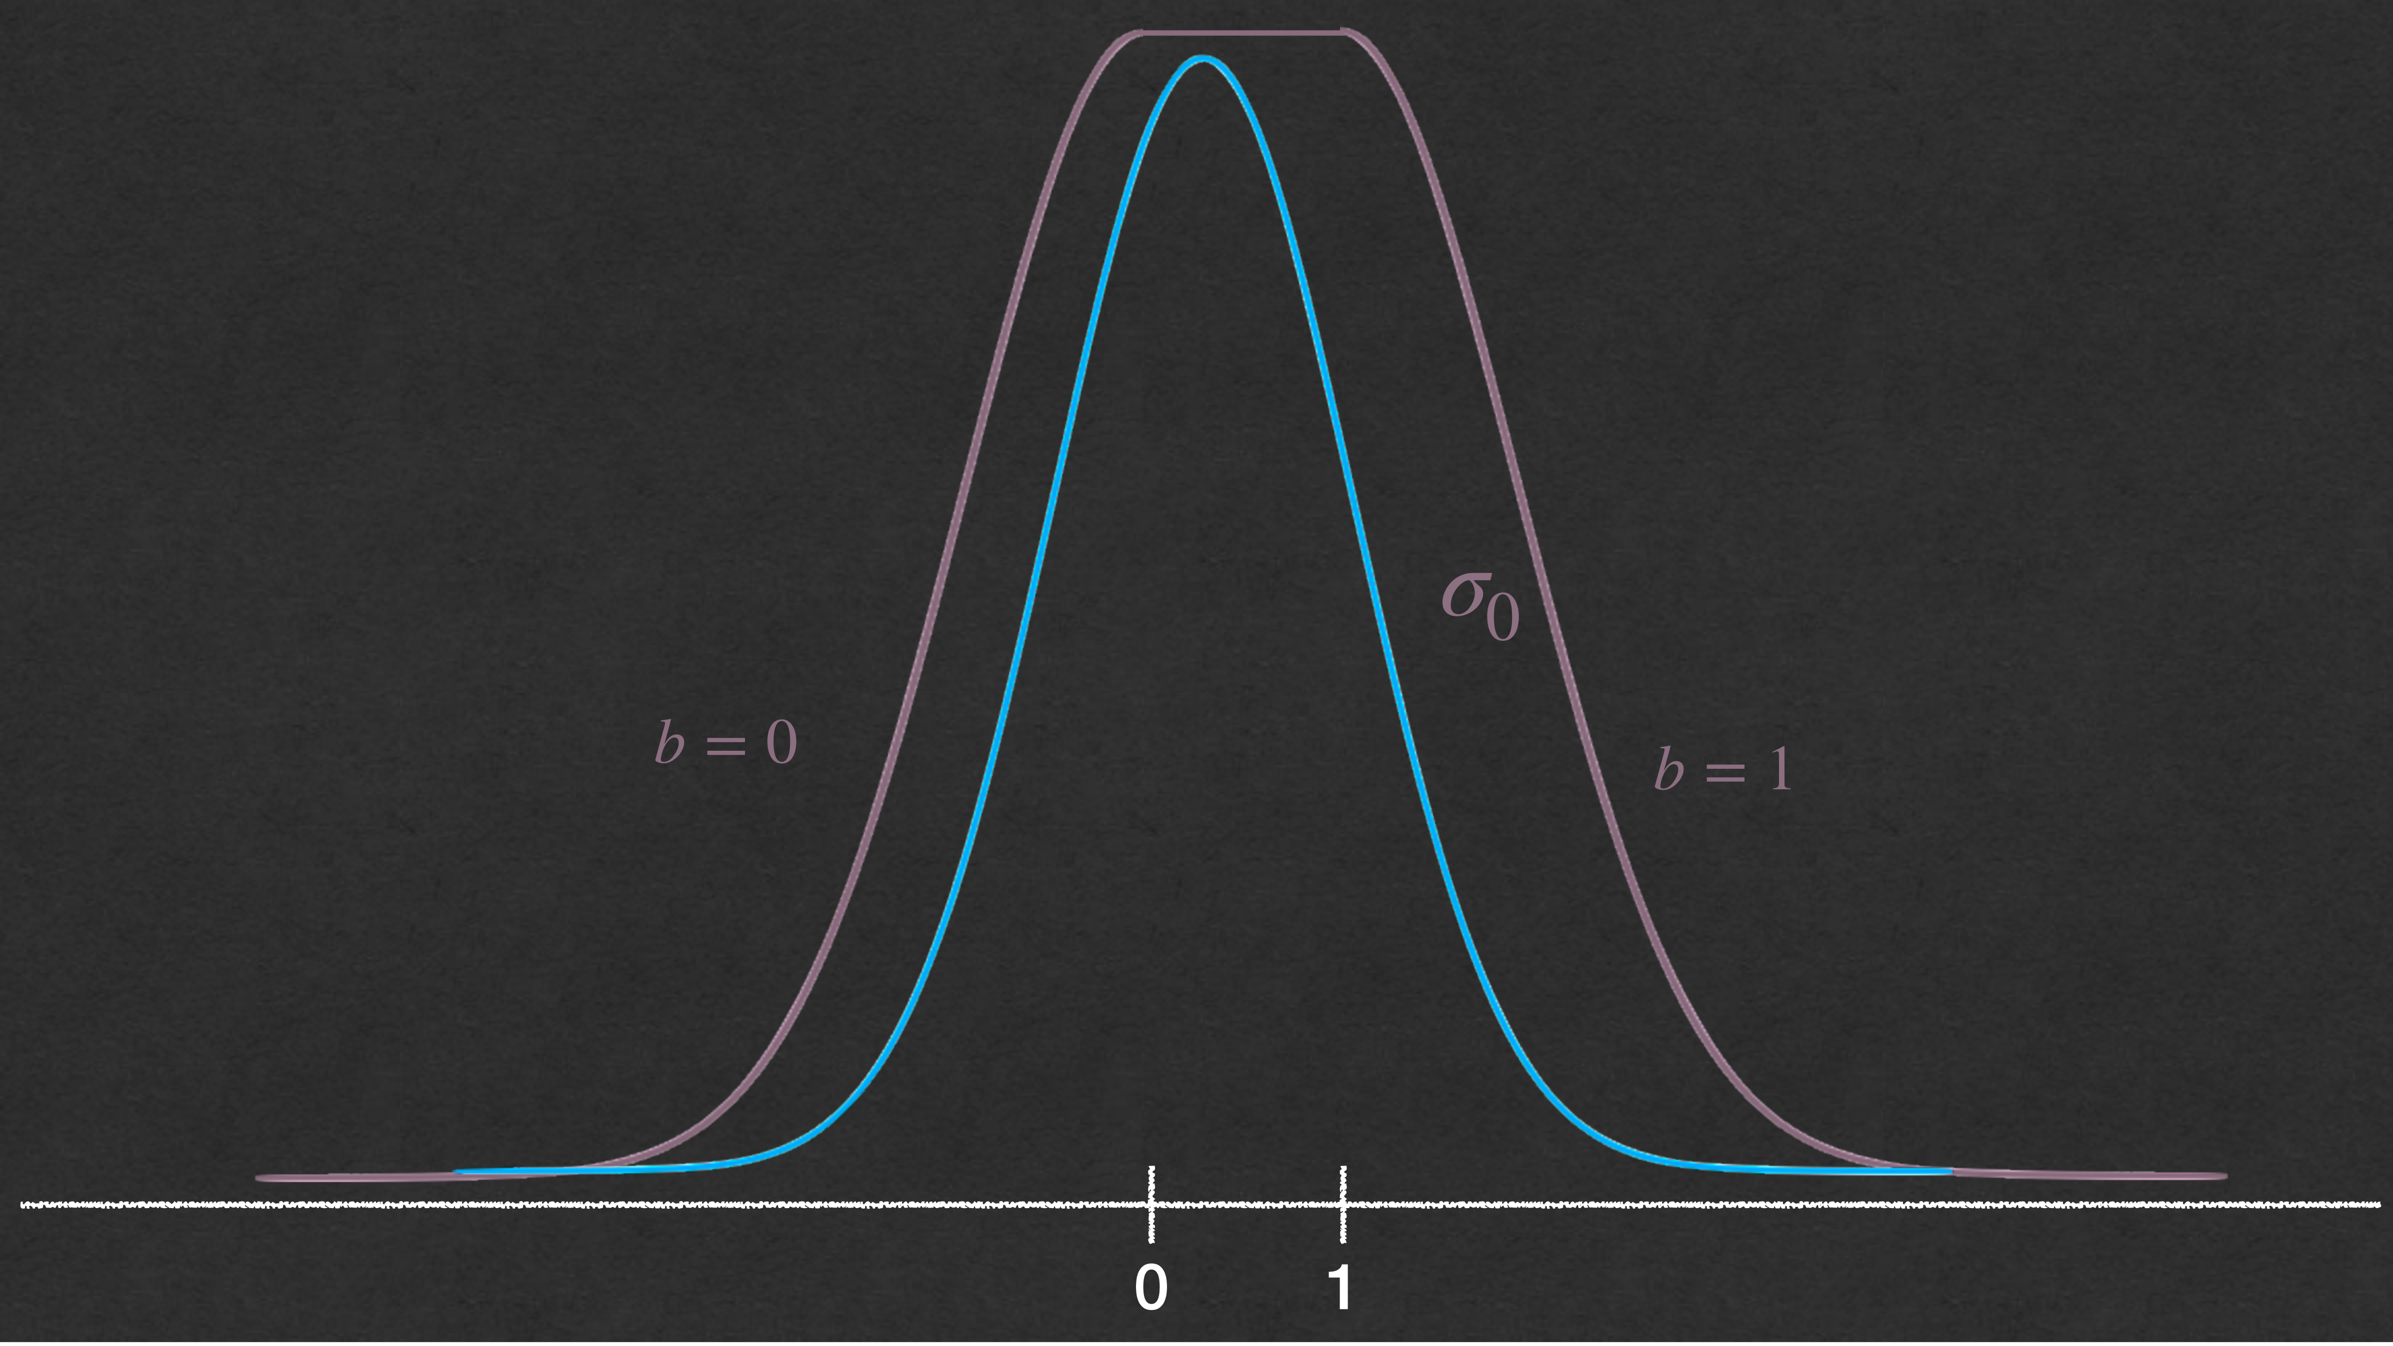

Falcon's Gaussian sampler

Algorithm SampleZ(σ, μ)

Require: $\mu \in [0,1), \sigma \leq \sigma_0$

Ensure: $z \sim D_{\mathbb{Z},\sigma,\mu}$

1. $z_0 \leftarrow$ Basesampler()
2. $b \leftarrow \{0,1\}$ uniformly
3. $z \leftarrow (2b-1) \cdot z_0 + b$
4. $x \leftarrow -\dfrac{(z-\mu)^2}{2\sigma^2} + \dfrac{z_0^2}{2\sigma_0^2}$
5. Accept with probability $\exp(x)$
   Restart to 1. otherwise

Falcon's Gaussian sampler

Algorithm SampleZ(σ, μ)

Require: $\mu \in [0,1), \sigma \leq \sigma_0$

Ensure: $z \sim D_{\mathbb{Z},\sigma,\mu}$

1. $z_0 \leftarrow$ Basesampler()
2. $b \leftarrow \{0,1\}$ uniformly
3. $z \leftarrow (2b-1) \cdot z_0 + b$
4. $x \leftarrow -\dfrac{(z-\mu)^2}{2\sigma^2} + \dfrac{z_0^2}{2\sigma_0^2}$
5. Accept with probability $\exp(x)$
   Restart to 1. otherwise

1.

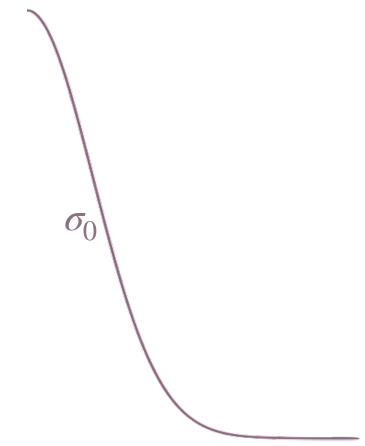

Falcon's Gaussian sampler

Algorithm SampleZ(σ, μ)

Require: $\mu \in [0,1), \sigma \leq \sigma_0$

Ensure: $z \sim D_{\mathbb{Z},\sigma,\mu}$

1. $z_0 \leftarrow$ Basesampler()
2. $b \leftarrow \{0,1\}$ uniformly
3. $z \leftarrow (2b-1) \cdot z_0 + b$
4. $x \leftarrow -\dfrac{(z-\mu)^2}{2\sigma^2} + \dfrac{z_0^2}{2\sigma_0^2}$
5. Accept with probability $\exp(x)$
   Restart to 1. otherwise

1.

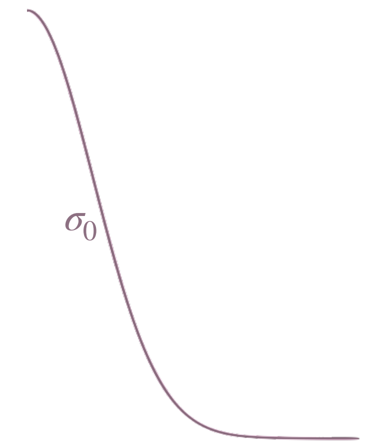

3.

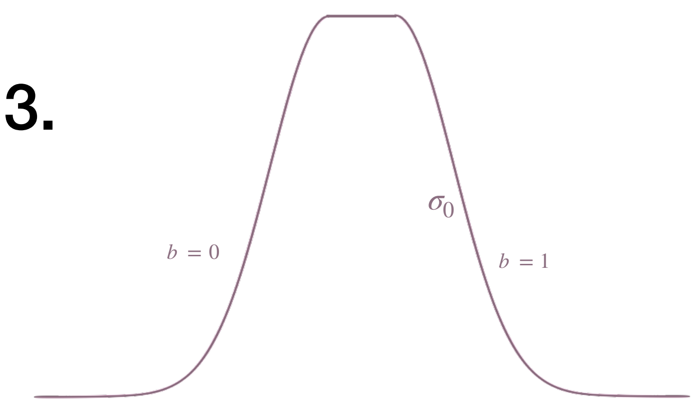

Falcon's Gaussian sampler

Algorithm SampleZ(σ, μ)

Require: $\mu \in [0,1), \sigma \leq \sigma_0$

Ensure: $z \sim D_{\mathbb{Z},\sigma,\mu}$

1. $z_0 \leftarrow$ Basesampler()
2. $b \leftarrow \{0,1\}$ uniformly
3. $z \leftarrow (2b-1)\cdot z_0 + b$
4. $x \leftarrow -\dfrac{(z-\mu)^2}{2\sigma^2} + \dfrac{z_0^2}{2\sigma_0^2}$
5. Accept with probability $\exp(x)$
   Restart to 1. otherwise

$$P_{\text{accept}} = \frac{\exp\left(-\frac{(z-\mu)^2}{2\sigma^2}\right)}{\exp\left(-\frac{z_0^2}{2\sigma_0^2}\right)}$$

1.

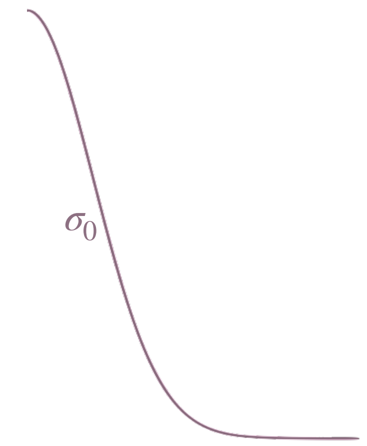

3.

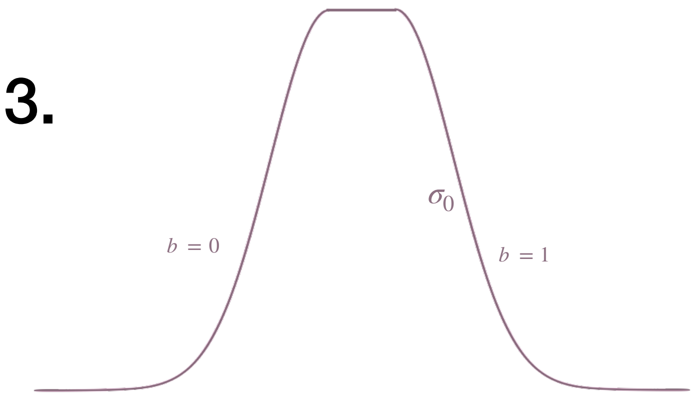

5.

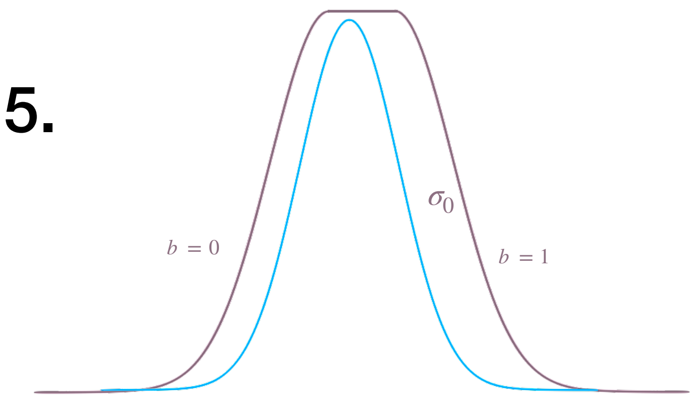

Isochronous Falcon Gaussian sampler

Algorithm SampleZ(σ, μ)

Require: $\mu \in [0,1), \sigma \leq \sigma_0$

Ensure: $z \sim D_{\mathbb{Z},\sigma,\mu}$

1. $z_0 \leftarrow$ Basesampler()
2. $b \leftarrow \{0,1\}$ uniformly
3. $z \leftarrow (2b-1) \cdot z_0 + b$
4. $x \leftarrow -\dfrac{(z-\mu)^2}{2\sigma^2} + \dfrac{z_0^2}{2\sigma_0^2}$
5. Accept with probability $\exp(x)$
   Restart to 1. otherwise

Isochronous Falcon Gaussian sampler

Algorithm SampleZ(σ, μ)

Require: $\mu \in [0,1), \sigma \leq \sigma_0$

Ensure: $z \sim D_{\mathbb{Z},\sigma,\mu}$

1. $z_0 \leftarrow$ Basesampler()
2. $b \leftarrow \{0,1\}$ uniformly
3. $z \leftarrow (2b-1) \cdot z_0 + b$
4. $x \leftarrow -\dfrac{(z-\mu)^2}{2\sigma^2} + \dfrac{z_0^2}{2\sigma_0^2}$
5. Accept with probability $\exp(x)$
   Restart to 1. otherwise

Isochrony details

1) Basesampler with a table
2) Polynomial approximation for exp
3) Make the number of iterations independent from the secret

Rényi divergence and security

Security analysis

1. Our sampler is isochronous with respect to the **standard deviation σ, the center μ and the sampled value z.**

2. Using our sampler on a λ-bit secure signature scheme provides $\lambda - 2$ bits of security.

See our paper for the proof

Rényi divergence and security

Security analysis

1. Our sampler is isochronous with respect to the **standard deviation σ, the center μ and the sampled value z.**
2. Using our sampler on a λ-bit secure signature scheme provides $\lambda - 2$ bits of security.

See our paper for the proof

Rényi divergence tool

Take two cryptographic schemes
- One with distribution $\mathscr{P}$
- One with an approximate distribution $\mathscr{Q}$ with the same support

Suppose that :

1. $\mathscr{P}$ and $\mathscr{Q}$ are close enough : $\left\| 1 - \frac{\mathscr{Q}}{\mathscr{P}} \right\|_\infty \leq 2^{-K}$
2. the number of sample queries is bounded

Then, the **bit security will remain almost the same.**

- T. Prest
  ASIACRYPT'17
- S. Bai, A. Langlois, T. Lepoint, D. Stehle, and R. Steinfeld.
  ASIACRYPT'15

A technical independence issue in isochrony

- T. Prest
  ASIACRYPT'17
- S. Bai, A. Langlois, T. Lepoint, D. Stehle, and R. Steinfeld.
  ASIACRYPT'15

Take two cryptographic schemes

- One with distribution $\mathscr{P}$
- One with an approximate distribution $\mathscr{Q}$ with the same support

Suppose that :

1. $\mathscr{P}$ and $\mathscr{Q}$ are close enough : $\left\| 1 - \frac{\mathscr{Q}}{\mathscr{P}} \right\|_\infty \leq 2^{-K}$
2. the number of sample queries is bounded

Then, the **bit security will remain almost the same.**

A technical independence issue in isochrony

- T. Prest
  ASIACRYPT'17
- S. Bai, A. Langlois, T. Lepoint, D. Stehle, and R. Steinfeld.
  ASIACRYPT'15

Take two cryptographic schemes
- One with distribution $\mathscr{P}$
- One with an approximate distribution $\mathscr{Q}$ with the same support

Suppose that :

1. $\mathscr{P}$ and $\mathscr{Q}$ are close enough : $R_a(\mathscr{Q}, \mathscr{P}) \leq \sqrt{2}$
2. the number of sample queries is bounded

Then, the **bit security will remain almost the same.**

A technical independence issue in isochrony

- T. Prest
  ASIACRYPT'17
- S. Bai, A. Langlois, T. Lepoint, D. Stehle, and R. Steinfeld.
  ASIACRYPT'15

Take two cryptographic schemes
- One with distribution $\mathscr{P}$
- One with an approximate distribution $\mathscr{Q}$ with the same support

Suppose that :

1. $\mathscr{P}$ and $\mathscr{Q}$ are close enough : $R_a(\mathscr{Q}, \mathscr{P}) \leq \sqrt{2}$
2. the number of sample queries is bounded

Then, the **bit security will remain almost the same.**

Let $\mathscr{P}$ and $\mathscr{Q}$ denote two distributions of a N-uple of variables (x_i).

Multiplicativity

If the random variables (x_i) are independent,

$$R_a(\mathscr{Q}, \mathscr{P}) = \Pi_i \; R_a(\mathscr{Q}_i, \mathscr{P}_i)$$

A technical independence issue in isochrony

- T. Prest
  ASIACRYPT'17
- S. Bai, A. Langlois, T. Lepoint, D. Stehle, and R. Steinfeld.
  ASIACRYPT'15

Take two cryptographic schemes
- One with distribution $\mathscr{P}$
- One with an approximate distribution $\mathscr{Q}$ with the same support

Suppose that :

1. $\mathscr{P}$ and $\mathscr{Q}$ are close enough : $R_a(\mathscr{Q}, \mathscr{P}) \leq \sqrt{2}$
2. the number of sample queries is bounded

Then, the **bit security will remain almost the same.**

Let $\mathscr{P}$ and $\mathscr{Q}$ denote two distributions of a N-uple of variables (x_i).

Multiplicativity

If the random variables (x_i) are independent,

$$R_a(\mathscr{Q}, \mathscr{P}) = \Pi_i \; R_a(\mathscr{Q}_i, \mathscr{P}_i)$$

$$R_a(\mathscr{Q}, \mathscr{P}) \leq \left(1 + \frac{1}{4N}\right)^N \leq \exp(1/4) \leq \sqrt{2}$$

A technical independence issue in isochrony

- T. Prest
  ASIACRYPT'17
- S. Bai, A. Langlois, T. Lepoint, D. Stehle, and R. Steinfeld.
  ASIACRYPT'15

Take two cryptographic schemes
- One with distribution $\mathscr{P}$
- One with an approximate distribution $\mathscr{Q}$ with the same support

Suppose that :

1. $\mathscr{P}$ and $\mathscr{Q}$ are close enough : $R_a(\mathscr{Q}, \mathscr{P}) \leq \sqrt{2}$
2. the number of sample queries is bounded

Then, the **bit security will remain almost the same.**

Let $\mathscr{P}$ and $\mathscr{Q}$ denote two distributions of a N-uple of variables (x_i).

Multiplicativity

If the random variables (x_i) are independent,

$$R_a(\mathscr{Q}, \mathscr{P}) = \Pi_i \; R_a(\mathscr{Q}_i, \mathscr{P}_i)$$

$$R_a(\mathscr{Q}, \mathscr{P}) \leq \left(1 + \frac{1}{4N}\right)^N \leq \exp(1/4) \leq \sqrt{2}$$

One issue

The multiplicativity result can only be applied if the distributions $\mathscr{P}_i$ are independent.

OK for Fiat-Shamir with aborts signatures.

Not ok for Falcon where σ and μ are dependent.

A technical independence issue in isochrony

- T. Prest
  ASIACRYPT'17
- S. Bai, A. Langlois, T. Lepoint, D. Stehle, and R. Steinfeld.
  ASIACRYPT'15

Take two cryptographic schemes
- One with distribution $\mathscr{P}$
- One with an approximate distribution $\mathscr{Q}$ with the same support

Suppose that :

1. $\mathscr{P}$ and $\mathscr{Q}$ are close enough : $R_a(\mathscr{Q}, \mathscr{P}) \leq \sqrt{2}$
2. the number of sample queries is bounded

Then, the **bit security will remain almost the same.**

Let $\mathscr{P}$ and $\mathscr{Q}$ denote two distributions of a N-uple of variables (x_i).

Multiplicativity

If the random variables (x_i) are independent,

$$R_a(\mathscr{Q}, \mathscr{P}) = \Pi_i \; R_a(\mathscr{Q}_i, \mathscr{P}_i)$$

$$R_a(\mathscr{Q}, \mathscr{P}) \leq \left(1 + \frac{1}{4N}\right)^N \leq \exp(1/4) \leq \sqrt{2}$$

One issue

The multiplicativity result can only be applied if the distributions $\mathscr{P}_i$ are independent.

OK for Fiat-Shamir with aborts signatures.

Not ok for Falcon where σ and μ are dependent.

Multiplicativity without independence assumption

If for every preceding drawn values of $x_{<i} = (x_0, \ldots, x_{i-1})$,

$$R_a(\mathscr{Q}_{i|x_{<i}}, \mathscr{P}_{i|x_{<i}}) \leq r_{a,i}\,.$$

Then, the Renyi divergence of $\mathscr{P}$ and $\mathscr{Q}$ is also bounded

$$R_a(\mathscr{Q}, \mathscr{P}) \leq \Pi_i \; r_{a,i}\,.$$

A technical independence issue in isochrony

- T. Prest
  ASIACRYPT'17
- S. Bai, A. Langlois, T. Lepoint, D. Stehle, and R. Steinfeld.
  ASIACRYPT'15

Take two cryptographic schemes

- One with distribution $\mathscr{P}$
- One with an approximate distribution $\mathscr{Q}$ with the same support

Suppose that :

1. $\mathscr{P}$ and $\mathscr{Q}$ are close enough : $R_a(\mathscr{Q}, \mathscr{P}) \leq \sqrt{2}$
2. the number of sample queries is bounded

Then, the **bit security will remain almost the same.**

Let $\mathscr{P}$ and $\mathscr{Q}$ denote two distributions of a N-uple of variables (x_i).

Multiplicativity

If the random variables (x_i) are independent,

$$R_a(\mathscr{Q}, \mathscr{P}) = \Pi_i \ R_a(\mathscr{Q}_i, \mathscr{P}_i)$$

One issue

The multiplicativity result can only be applied if the distributions $\mathscr{P}_i$ are independent.

OK for Fiat-Shamir with aborts signatures.

Not ok for Falcon where σ and μ are dependent.

Multiplicativity without independence assumption

If for every preceding drawn values of $x_{<i} = (x_0, \ldots, x_{i-1})$,

$$R_a(\mathscr{Q}_{i|x_{<i}}, \mathscr{P}_{i|x_{<i}}) \leq r_{a,i}\,.$$

Then, the Renyi divergence of $\mathscr{P}$ and $\mathscr{Q}$ is also bounded

$$R_a(\mathscr{Q}, \mathscr{P}) \leq \Pi_i \ r_{a,i}\,.$$

$$R_a(\mathscr{Q}, \mathscr{P}) \leq \left(1 + \frac{1}{4N}\right)^N \leq \exp(1/4) \leq \sqrt{2}$$

The isochronous sampler

- ☐ Basesampler with a table
- ☐ Polynomial approximation for exp
- ☐ Make the number of iterations independent from the secret

1) Sampling with a table

BaseSampler() close to $D_{\mathbb{Z}^+,\sigma_0}$

Cumulative Distribution Table (CDT) with w elements of θ bits

CDT sampling can be done in constant time if the algorithm reads the entire table each time and carry out each comparison

1) Sampling with a table

BaseSampler() close to $D_{\mathbb{Z}^+,\sigma_0}$

We provide a script that generates w and the CDT table for a given target precision $\epsilon = 2^{-80}$ and θ

CDT sampling can be done in constant time if the algorithm reads the entire table each time and carry out each comparison

1) Sampling with a table

BaseSampler() close to $D_{\mathbb{Z}^+,\sigma_0}$

We provide a script that generates w and the CDT table for a given target precision $\epsilon = 2^{-80}$ and θ

Algorithm Renyification(σ, ϵ, θ)

Require: $\sigma, \epsilon \leq 0, \theta$

Ensure: w, the CDT table

1. $w \leftarrow$ Smallest tailcut such that $R_a\left(D_{[w],\sigma_0}, D_{\mathbb{Z}^+,\sigma_0}\right) \leq 1 + \epsilon$
2. Compute the table values with a « clever » rounding
   1. For $z \geq 1$, $CDT(z) \leftarrow 2^{-\theta}\left\lfloor 2^{\theta} \cdot D_{[w],\sigma_0}(z) \right\rfloor$
   2. $CDT(0) \leftarrow 1 - \sum_{z \geq 1} CDT(z)$
3. Recompute Rényi divergence and return the new precision, w and CDT

1) CDT Sampling

$$R_\infty\left(\textbf{BaseSampler()}, D_{\mathbb{Z}^+,\sigma_0}\right) \leq 1 + 2^{-80}$$

For $\sigma_0 = 1.8205$, our script gave

$w = 19$ elements

$\theta = 72$ bits

$\epsilon = 80$

$\texttt{CDT}(0) = 2^{-72} \times 1697680241746640300030$

$\texttt{CDT}(1) = 2^{-72} \times 1459943456642912959616$

$\texttt{CDT}(2) = 2^{-72} \times 928488355018011056515$

$\texttt{CDT}(3) = 2^{-72} \times 436693944817054414619$

$\texttt{CDT}(4) = 2^{-72} \times 151893140790369201013$

$\texttt{CDT}(5) = 2^{-72} \times 39071441848292237840$

$\texttt{CDT}(6) = 2^{-72} \times 7432604049020375675$

$\texttt{CDT}(7) = 2^{-72} \times 1045641569992574730$

$\texttt{CDT}(8) = 2^{-72} \times 108788995549429682$

$\texttt{CDT}(9) = 2^{-72} \times 8370422445201343$

$\texttt{CDT}(10) = 2^{-72} \times 476288472308334$

$\texttt{CDT}(11) = 2^{-72} \times 20042553305308$

$\texttt{CDT}(12) = 2^{-72} \times 623729532807$

$\texttt{CDT}(13) = 2^{-72} \times 4354889437$

$\texttt{CDT}(14) = 2^{-72} \times 244322621$

$\texttt{CDT}(15) = 2^{-72} \times 3075302$

$\texttt{CDT}(16) = 2^{-72} \times 28626$

$\texttt{CDT}(17) = 2^{-72} \times 197$

$\texttt{CDT}(18) = 2^{-72} \times 1$

The isochronous sampler

- ☑ Basesampler with a table
- ☐ Polynomial approximation for exp
- ☐ Make the number of iterations independent from the secret

2) Polynomial approximation

Find P such that $$\left|\frac{P(x)-\exp(x)}{\exp(x)}\right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$$

and $$\left|\frac{P(x)-\exp(x)}{1-\exp(x)}\right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$$

Polynomial approximation tools

- Floating points option: FACCT by Zhao, Steinfeld and Sakzad 2018/1234

2) Polynomial approximation

Find P such that $$\left|\frac{P(x)-\exp(x)}{\exp(x)}\right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$$

and $$\left|\frac{P(x)-\exp(x)}{1-\exp(x)}\right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$$

Polynomial approximation tools

- Floating points option: FACCT by Zhao, Steinfeld and Sakzad 2018/1234
- Integer option: GALACTICS by Barthe et al. 2019/511

$\Longrightarrow$ 32-bit coefficients
degree 10

2) Polynomial approximation

Find P such that $\left| \dfrac{P(x) - \exp(x)}{\exp(x)} \right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$

and $\left| \dfrac{P(x) - \exp(x)}{1 - \exp(x)} \right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$

Polynomial approximation tools

- Floating points option: FACCT by Zhao, Steinfeld and Sakzad 2018/1234
- Integer option: GALACTICS by Barthe et al. 2019/511

$\Longrightarrow$ 32-bit coefficients, degree 10

Depending on the architecture, several tradeoffs

2) Polynomial approximation

Find P such that $$\left|\frac{P(x)-\exp(x)}{\exp(x)}\right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$$

and $$\left|\frac{P(x)-\exp(x)}{1-\exp(x)}\right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$$

Polynomial approximation tools

- Floating points option: FACCT by Zhao, Steinfeld and Sakzad 2018/1234
- Integer option: GALACTICS by Barthe et al. 2019/511

$\Longrightarrow$ 32-bit coefficients
degree 10

Depending on the architecture, several tradeoffs

Degree

2) Polynomial approximation

Find P such that $\left|\frac{P(x) - \exp(x)}{\exp(x)}\right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$

and $\left|\frac{P(x) - \exp(x)}{1 - \exp(x)}\right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$

Polynomial approximation tools

- Floating points option: FACCT by Zhao, Steinfeld and Sakzad 2018/1234
- Integer option: GALACTICS by Barthe et al. 2019/511

$\Longrightarrow$ 32-bit coefficients, degree 10

Depending on the architecture, several tradeoffs

Degree

Size

2) Polynomial approximation

Find P such that $\left|\frac{P(x) - \exp(x)}{\exp(x)}\right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$

and $\left|\frac{P(x) - \exp(x)}{1 - \exp(x)}\right| \leq 2^{-44} \quad \forall x \in [0, \ln(2)]$

Polynomial approximation tools

- Floating points option: FACCT by Zhao, Steinfeld and Sakzad 2018/1234
- Integer option: GALACTICS by Barthe et al. 2019/511

$\Longrightarrow$ 32-bit coefficients
degree 10

Depending on the architecture, several tradeoffs

Degree | Size | Depth

The isochronous sampler

- ☑ Basesampler with a table
- ☑ Polynomial approximation for exp
- ☐ Make the number of iterations independent from the secret

3) Number of iterations of the while loop

?

Zhao, Steinfeld and Sakzad (2018/1234)
Karmakar et al (2019/267)

- Could the number of iterations leak the secret?

3) Number of iterations of the while loop

? Zhao, Steinfeld and Sakzad (2018/1234)
Karmakar et al (2019/267)

- Could the number of iterations leak the secret?

The number of iterations follows a geometric distribution of average $\dfrac{2 \cdot \rho_{\sigma_0}(\mathbb{Z}^+)}{\rho_{\sigma,\mu}(\mathbb{Z})}$

3) Number of iterations of the while loop

?

Zhao, Steinfeld and Sakzad (2018/1234)
Karmakar et al (2019/267)

- Could the number of iterations leak the secret?

The average number of iterations is $\dfrac{2 \cdot \rho_{\sigma_0}(\mathbb{Z}^+)}{\frac{\sigma_{min}}{\sigma}\rho_{\sigma,\mu}(\mathbb{Z})}$

Tweak for Falcon's sampler

The acceptance probability P_{accept} is scaled by a factor $\dfrac{\sigma_{min}}{\sigma} \leq \dfrac{\sigma_{min}}{\sigma_{max}} \approx 0.73$

3) Number of iterations of the while loop

? Zhao, Steinfeld and Sakzad (2018/1234)
Karmakar et al (2019/267)

- Could the number of iterations leak the secret?

The average number of iterations is $\dfrac{2\cdot\rho_{\sigma_0}(\mathbb{Z}^+)}{\frac{\sigma_{min}}{\sigma}\rho_{\sigma,\mu}(\mathbb{Z})}$

Tweak for Falcon's sampler

The acceptance probability P_{accept} is scaled by a factor $\dfrac{\sigma_{min}}{\sigma} \leq \dfrac{\sigma_{min}}{\sigma_{max}} \approx 0.73$

Indeed, with a Poisson summation (under a Rényi divergence argument),

$$\rho_{\sigma,\mu}(\mathbb{Z}) \approx \sigma\sqrt{2\pi}$$

So, $$\frac{2\cdot\rho_{\sigma_0}(\mathbb{Z}^+)}{\frac{\sigma_{min}}{\sigma}\rho_{\sigma,\mu}(\mathbb{Z})} \approx \frac{2\cdot\rho_{\sigma_0}(\mathbb{Z}^+)}{\frac{\sigma_{min}}{\sigma}\sigma\sqrt{2\pi}} = \frac{2\cdot\rho_{\sigma_0}(\mathbb{Z}^+)}{\sigma_{min}\sqrt{2\pi}}$$

3) Number of iterations of the while loop

Zhao, Steinfeld and Sakzad (2018/1234)
Karmakar et al (2019/267)

- Could the number of iterations leak the secret?

The average number of iterations is $$\frac{2 \cdot \rho_{\sigma_0}(\mathbb{Z}^+)}{\frac{\sigma_{min}}{\sigma}\rho_{\sigma,\mu}(\mathbb{Z})}$$

Tweak for Falcon's sampler

The acceptance probability P_{accept} is scaled by a factor $\frac{\sigma_{min}}{\sigma} \leq \frac{\sigma_{min}}{\sigma_{max}} \approx 0.73$

Indeed, with a Poisson summation (under a Rényi divergence argument),

$$\rho_{\sigma,\mu}(\mathbb{Z}) \approx \sigma\sqrt{2\pi}$$

So, $$\frac{2 \cdot \rho_{\sigma_0}(\mathbb{Z}^+)}{\frac{\sigma_{min}}{\sigma}\rho_{\sigma,\mu}(\mathbb{Z})} \approx \frac{2 \cdot \rho_{\sigma_0}(\mathbb{Z}^+)}{\frac{\sigma_{min}}{\sigma}\sigma\sqrt{2\pi}} = \frac{2 \cdot \rho_{\sigma_0}(\mathbb{Z}^+)}{\sigma_{min}\sqrt{2\pi}}$$

- ✓ Independent from μ
- ✓ Independent from σ
- ✓ Independent from z

3) Number of iterations of the while loop

?

Zhao, Steinfeld and Sakzad (2018/1234)
Karmakar et al (2019/267)

- Could the number of iterations leak the secret?

The average number of iterations is $\dfrac{2 \cdot \rho_{\sigma_0}(\mathbb{Z}^+)}{\frac{\sigma_{min}}{\sigma}\rho_{\sigma,\mu}(\mathbb{Z})}$

Tweak for Falcon's sampler

The acceptance probability P_{accept} is scaled by a factor $\dfrac{\sigma_{min}}{\sigma} \leq \dfrac{\sigma_{min}}{\sigma_{max}} \approx 0.73$

Indeed, with a Poisson summation (under a Rényi divergence argument),

$$\rho_{\sigma,\mu}(\mathbb{Z}) \approx \sigma\sqrt{2\pi}$$

So, $\dfrac{2 \cdot \rho_{\sigma_0}(\mathbb{Z}^+)}{\frac{\sigma_{min}}{\sigma}\rho_{\sigma,\mu}(\mathbb{Z})} \approx \dfrac{2 \cdot \rho_{\sigma_0}(\mathbb{Z}^+)}{\frac{\sigma_{min}}{\sigma}\sigma\sqrt{2\pi}} = \dfrac{2 \cdot \rho_{\sigma_0}(\mathbb{Z}^+)}{\sigma_{min}\sqrt{2\pi}}$

✓ Independent from μ

✓ Independent from σ

✓ Independent from z

The whole algorithm
is constant time

Statistically Acceptable Gaussians

Our second contribution is SAGA, a statistical test suite.

Statistically Acceptable Gaussians

Our second contribution is SAGA, a statistical test suite.
We propose this because:

☑ Implementation failures are possible, e.g. inaccuracy or incorrectness in CDT table values.

☑ Implementation failures can also be found if the base Gaussian sampler is validated, but the outputs are not.

☑ Randomness / entropy levels not being sufficient.

☑ SAGA only works on outputs, thus it is completely agnostic to the sampling method or scheme used.

Statistically Acceptable Gaussians

Our second contribution is SAGA, a statistical test suite.
More specifically SAGA can validate:

☑ Univariate Gaussian samples for base Gaussian samplers useful for samplers in FrodoKEM, DLP-IBE, FHE, etc.

☑ Multivariate Gaussian samples for outputs of schemes useful for Falcon, DLP-IBE, LATTE, etc.

☑ Supplementary, graphical, and sanity check tests for things like rejection rates, uni-, and multi-variate normality.

SAGA Tests on Univariate Samples

- ☑ First we compare the Expected vs Empirical observations for mean, variance, skewness, and kurtosis.

- ☑ Secondly we perform a chi-squared normality test.

```
Testing a Gaussian sampler with center = -0.920619 and sigma = 1.711864
Number of samples: 100

Moments   |   Expected     Empiric
---------+----------------------
Mean:     |   -0.92062      -0.92000
St. dev. |   1.71186       1.51446
Skewness |   0.00000       -0.25650
Kurtosis |   0.00000       -0.26704

Chi-2 statistic:   4.033416341364921
Chi-2 p-value:     0.4015023295495953   (should be > 0.001)

How many outliers? 0

Is the sample valid? True
```

An example output for testing univariate samples from a (base) Gaussian sampler.

SAGA Tests on Univariate Samples

- ☑ First we compare the Expected vs Empirical observations for mean, variance, skewness, and kurtosis.
- ☑ Secondly we perform a chi-squared normality test.

```
Testing a Gaussian sampler with center = -0.920619 and sigma = 1.711864
Number of samples: 100

Moments   |   Expected      Empiric
----------+----------------------
Mean:     |   -0.92062       -0.92000
St. dev.  |   1.71186       1.51446
Skewness  |   0.00000       -0.25650
Kurtosis  |   0.00000       -0.26704

Chi-2 statistic:   4.033416341364921
Chi-2 p-value:     0.4015023295495953   (should be > 0.001)

How many outliers? 0

Is the sample valid? True
```

An example output for testing univariate samples from a (base) Gaussian sampler.

SAGA Tests on Univariate Samples

- ☑ First we compare the Expected vs Empirical observations for mean, variance, skewness, and kurtosis.

- ☑ Secondly we perform a chi-squared normality test.

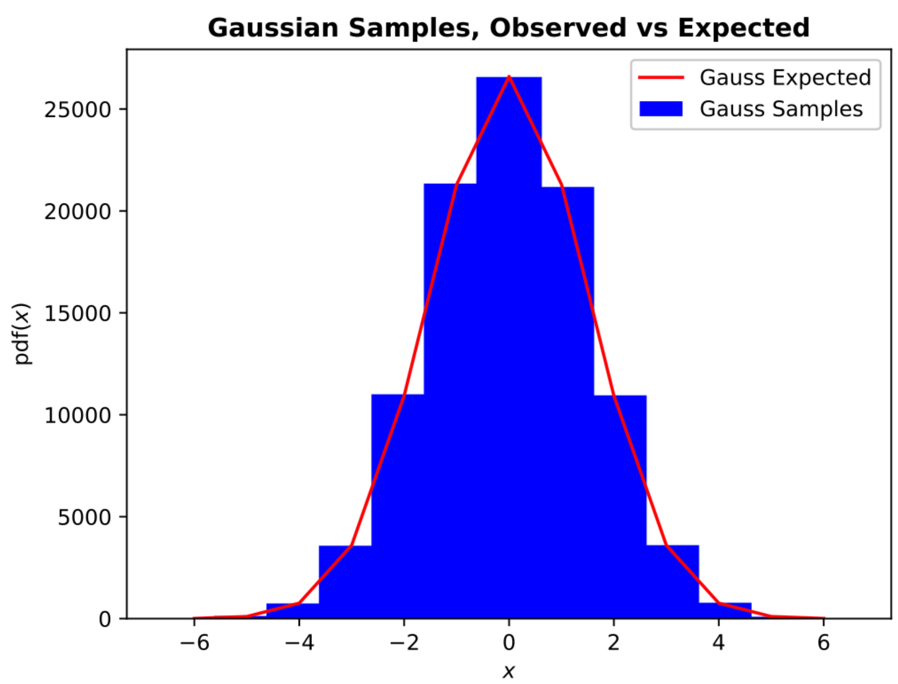

(a) Observed vs Expected Gaussian PDF

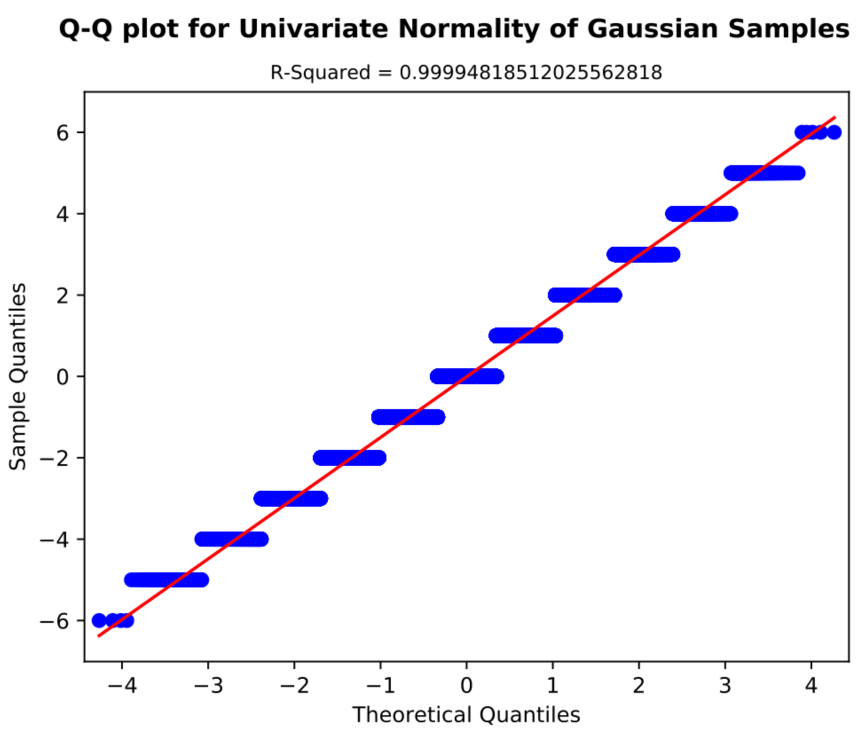

(b) QQ-plot of Observed vs Expected Quantiles

A visual representation of checking normality for the univariate Gaussian samples.

SAGA Tests on Univariate Samples

- ☑ First we compare the Expected vs Empirical observations for mean, variance, skewness, and kurtosis.
- ☑ Secondly we perform a chi-squared normality test.

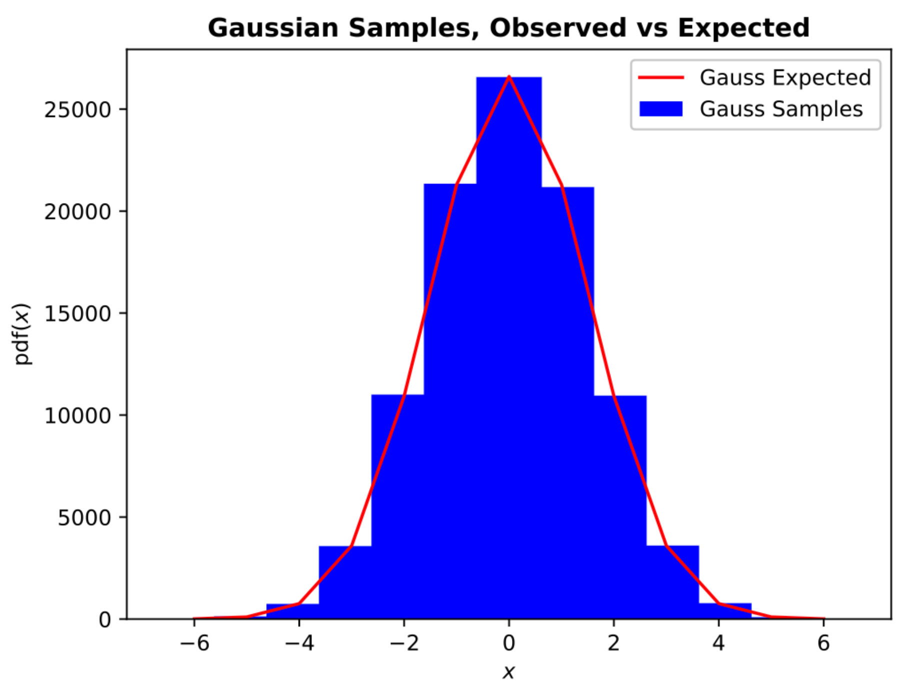

(a) Observed vs Expected Gaussian PDF

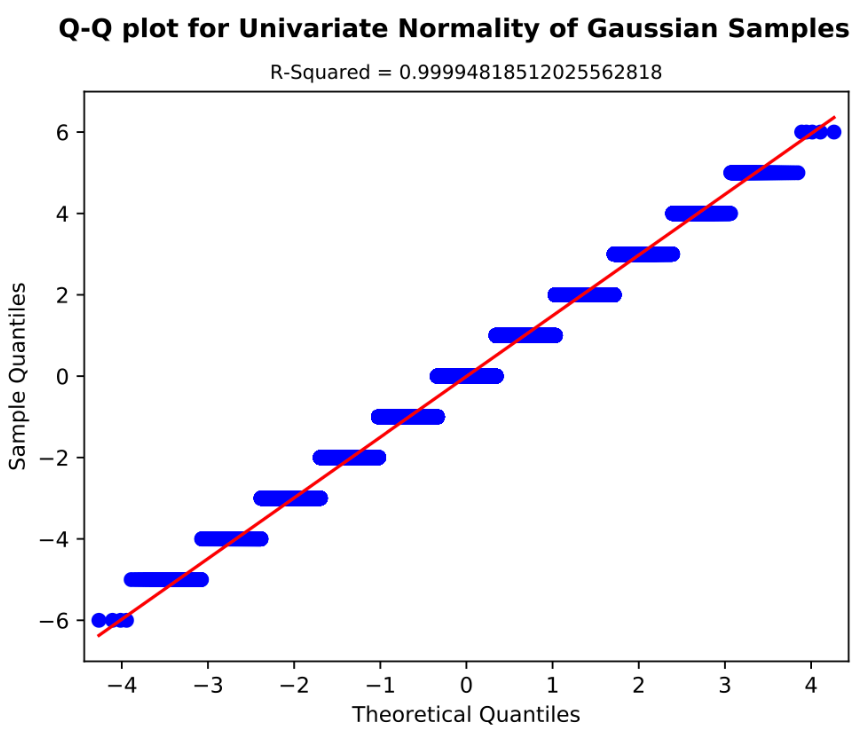

(b) QQ-plot of Observed vs Expected Quantiles

A visual representation of checking normality for the univariate Gaussian samples.

SAGA Tests on Multivariate Samples

1. Can we find errors if the base sampler is designed well?

- ☑ Incorrect tree designs in Falcon will affect its covariance.
- ☑ We thus posit that covariance in (block-)sub-diagonals:
  - ☑ grow in $O(\sqrt{n})$ for correct implementations and
  - ☑ grow in $O(n)$ for incorrect implementations.
- ☑ Test 3 uses this (p-value) in a chi-squared test.

```
1 - Covariance matrix (128 x 128):
[[ 0.997   -0.0021   0.0065 ...   0.0014   0.0012  -0.0039]
 [-0.0021   1.0001  -0.0014 ...   0.0032   0.0005  -0.0048]
 [ 0.0065  -0.0014   1.0028 ...  -0.0006   0.0074   0.0065]
 ...
 [ 0.0014   0.0032  -0.0006 ...   1.0063  -0.0022  -0.0005]
 [ 0.0012   0.0005   0.0074 ...  -0.0022   0.993   -0.0008]
 [-0.0039  -0.0048   0.0065 ...  -0.0005  -0.0008   1.0081]]

2 - P-value of Doornik-Hansen test:                 0.2453

3 - P-value of covariance diagonals test:           0.3244

4 - Gaussian coordinates (w/ st. dev. = sigma)?     128 out of 128
```

Example output for a correct implementation of Falcon.

SAGA Tests on Multivariate Samples

1. Can we find errors if the base sampler is designed well?

- ☑ Incorrect tree designs in Falcon will affect its covariance.
- ☑ We thus posit that covariance in (block-)sub-diagonals:
  - ☑ grow in $O(\sqrt{n})$ for correct implementations and
  - ☑ grow in $O(n)$ for incorrect implementations.
- ☑ Test 3 uses this (p-value) in a chi-squared test.

```
1 - Covariance matrix (128 x 128):
[[ 0.997   -0.0021   0.0065 ...   0.0014   0.0012  -0.0039]
 [-0.0021   1.0001  -0.0014 ...   0.0032   0.0005  -0.0048]
 [ 0.0065  -0.0014   1.0028 ...  -0.0006   0.0074   0.0065]
 ...
 [ 0.0014   0.0032  -0.0006 ...   1.0063  -0.0022  -0.0005]
 [ 0.0012   0.0005   0.0074 ...  -0.0022   0.993   -0.0008]
 [-0.0039  -0.0048   0.0065 ...  -0.0005  -0.0008   1.0081]]

2 - P-value of Doornik-Hansen test:                  0.2453

3 - P-value of covariance diagonals test:            0.3244

4 - Gaussian coordinates (w/ st. dev. = sigma)?      128 out of 128
```

Example output for a correct implementation of Falcon.

SAGA Tests on Multivariate Samples

1. Can we find errors if the base sampler is designed well?

- ☑ Incorrect tree designs in Falcon will affect its covariance.
- ☑ We thus posit that covariance in (block-)sub-diagonals:
  - ☑ grow in O(√n) for correct implementations and
  - ☑ grow in O(n) for incorrect implementations.
- ☑ Test 3 uses this (p-value) in a chi-squared test.

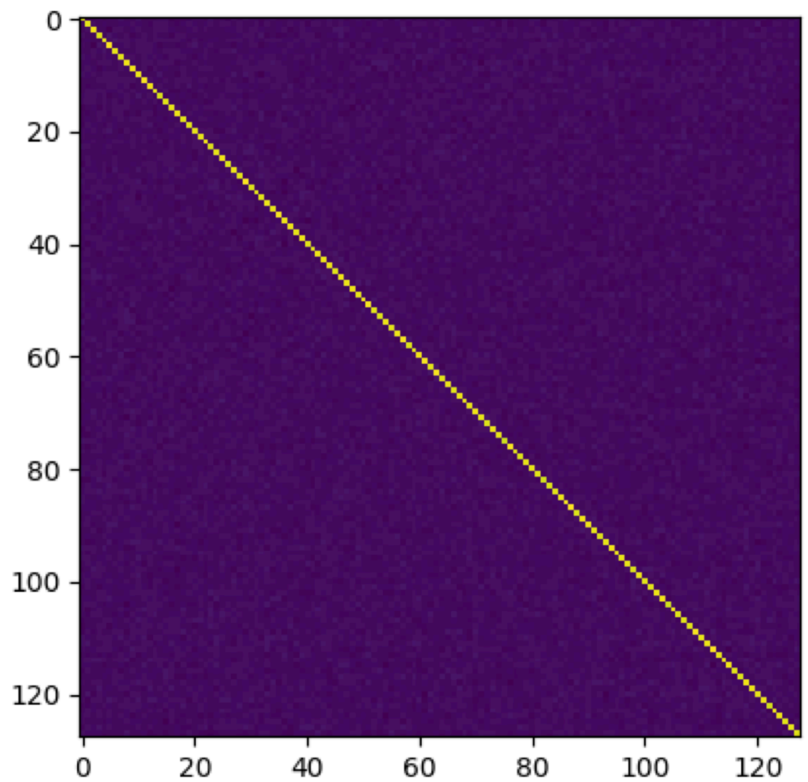

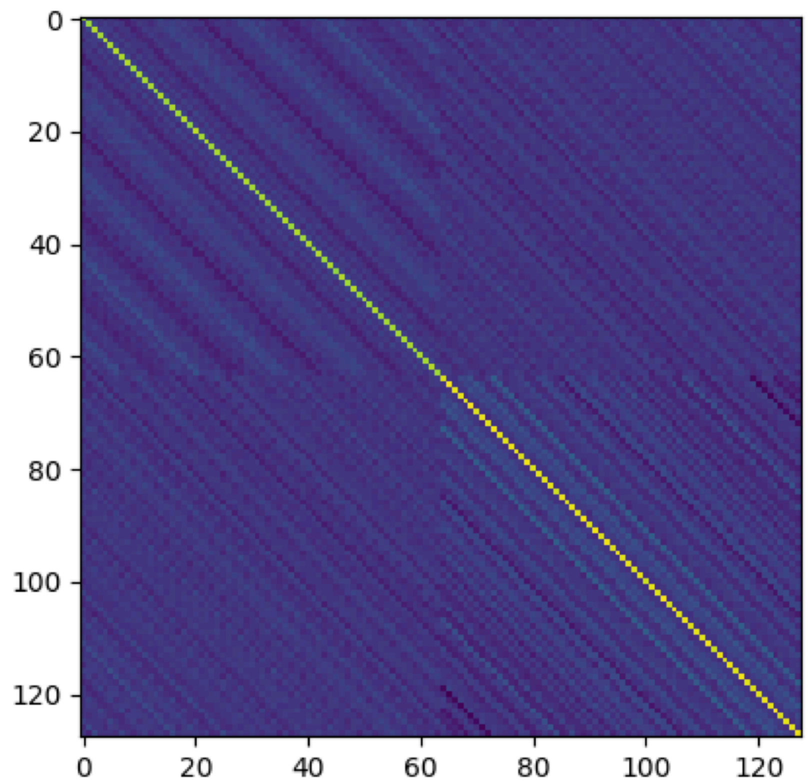

A properly functioning Falcon implementation VS an implementation with a mistake when constructing the Falcon tree.

SAGA Tests on Multivariate Samples

1. Can we find errors if the base sampler is designed well?

- ☑ Incorrect tree designs in Falcon will affect its covariance.
- ☑ We thus posit that covariance in (block-)sub-diagonals:
  - ☑ grow in O(√n) for correct implementations and
  - ☑ grow in O(n) for incorrect implementations.
- ☑ Test 3 uses this (p-value) in a chi-squared test.

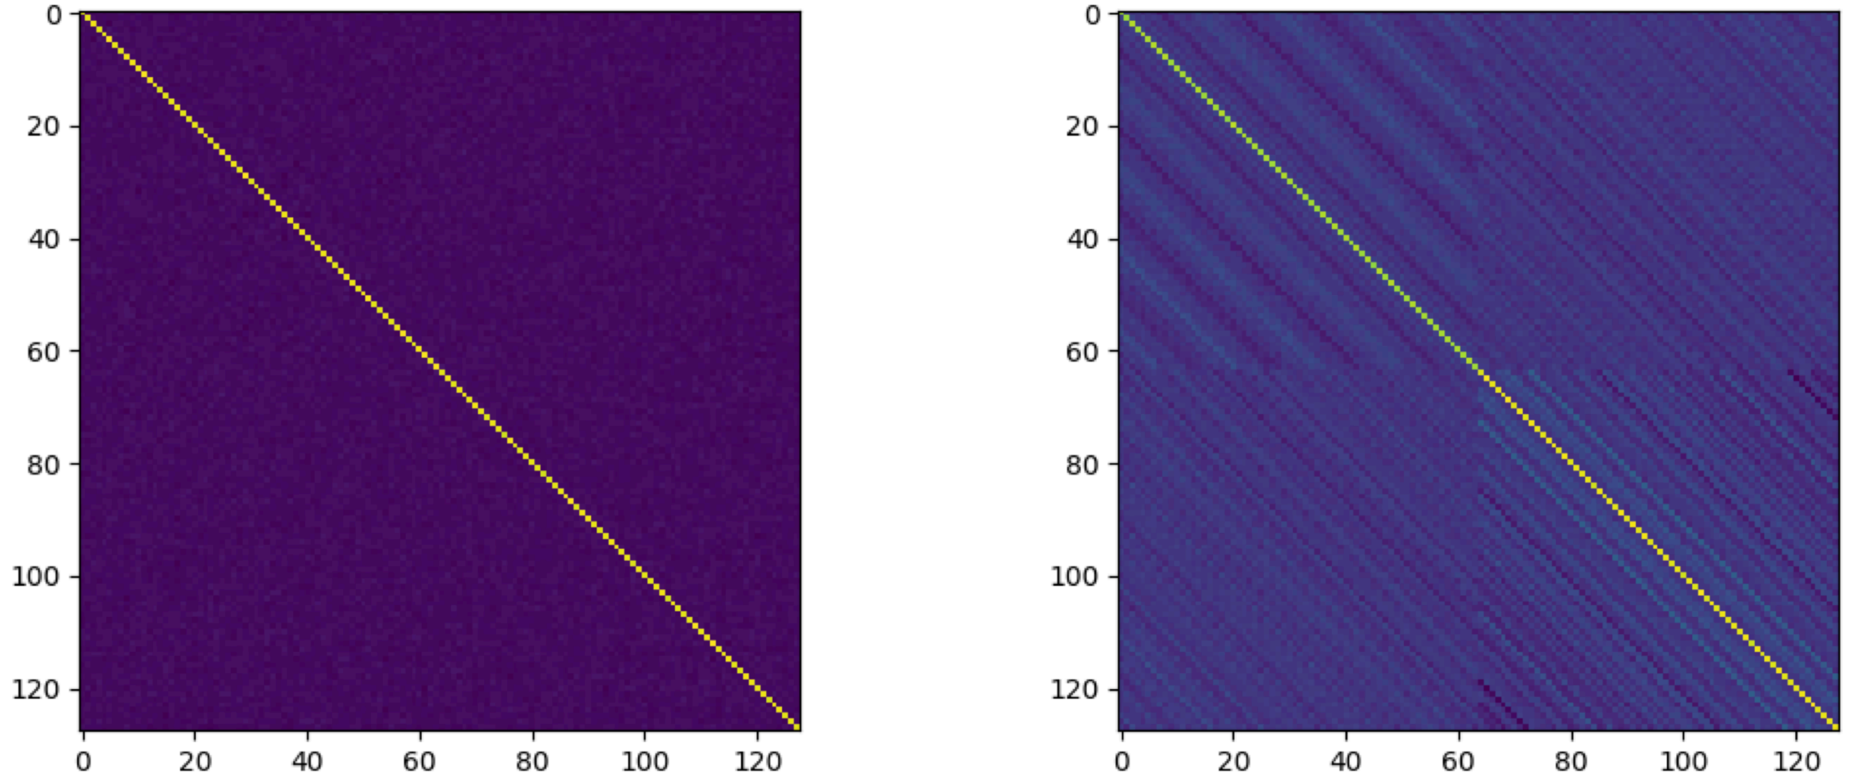

A properly functioning Falcon implementation VS an implementation with a mistake when constructing the Falcon tree.

SAGA Tests on Multivariate Samples

2. Performs a multivariate normality test.

- ☑ We implement the Doornik-Hansen test.
- ☑ Using skewness and kurtosis of the multivariate data.
- ☑ Other equivalent tests suffer with poor power.

```
1 - Covariance matrix (128 x 128):
[[ 0.997   -0.0021  0.0065 ...  0.0014  0.0012 -0.0039]
 [-0.0021  1.0001 -0.0014 ...  0.0032  0.0005 -0.0048]
 [ 0.0065 -0.0014  1.0028 ... -0.0006  0.0074  0.0065]
 ...
 [ 0.0014  0.0032 -0.0006 ...  1.0063 -0.0022 -0.0005]
 [ 0.0012  0.0005  0.0074 ... -0.0022  0.993  -0.0008]
 [-0.0039 -0.0048  0.0065 ... -0.0005 -0.0008  1.0081]]

2 - P-value of Doornik-Hansen test:                  0.2453

3 - P-value of covariance diagonals test:            0.3244

4 - Gaussian coordinates (w/ st. dev. = sigma)?     128 out of 128
```

Example output for a correct implementation of Falcon.

SAGA Tests on Multivariate Samples

2. Performs a multivariate normality test.

- ☑ We implement the Doornik-Hansen test.
- ☑ Using skewness and kurtosis of the multivariate data.
- ☑ Other equivalent tests suffer with poor power.

```
1 - Covariance matrix (128 x 128):
[[ 0.997   -0.0021  0.0065 ...  0.0014  0.0012 -0.0039]
 [-0.0021   1.0001 -0.0014 ...  0.0032  0.0005 -0.0048]
 [ 0.0065 -0.0014   1.0028 ... -0.0006  0.0074  0.0065]
 ...
 [ 0.0014  0.0032 -0.0006 ...  1.0063 -0.0022 -0.0005]
 [ 0.0012  0.0005  0.0074 ... -0.0022  0.993  -0.0008]
 [-0.0039 -0.0048  0.0065 ... -0.0005 -0.0008  1.0081]]

2 - P-value of Doornik-Hansen test:                  0.2453

3 - P-value of covariance diagonals test:            0.3244

4 - Gaussian coordinates (w/ st. dev. = sigma)?     128 out of 128
```

Example output for a correct implementation of Falcon.

SAGA Supplementary Tests

- ☑ Mahalanobis distance visualises multivariate normality.
- ☑ The distance measures std. devs. of each point from distribution.
- ☑ Empirical vs Expected should follow a chi-square distribution.

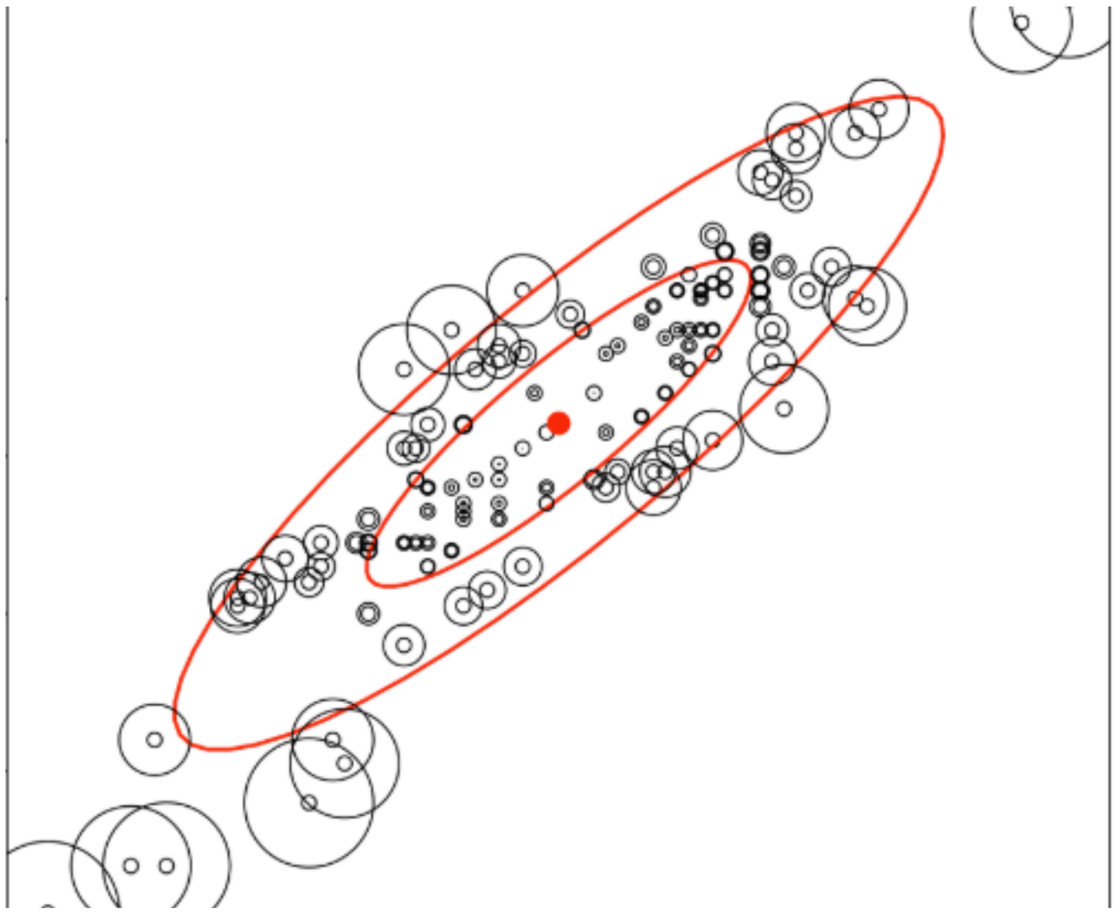

A visual representation of the Mahalanobis distance.

SAGA Supplementary Tests

- ☑ Mahalanobis distance visualises multivariate normality.
- ☑ The distance measures std. devs. of each point from distribution.
- ☑ Empirical vs Expected should follow a chi-square distribution.

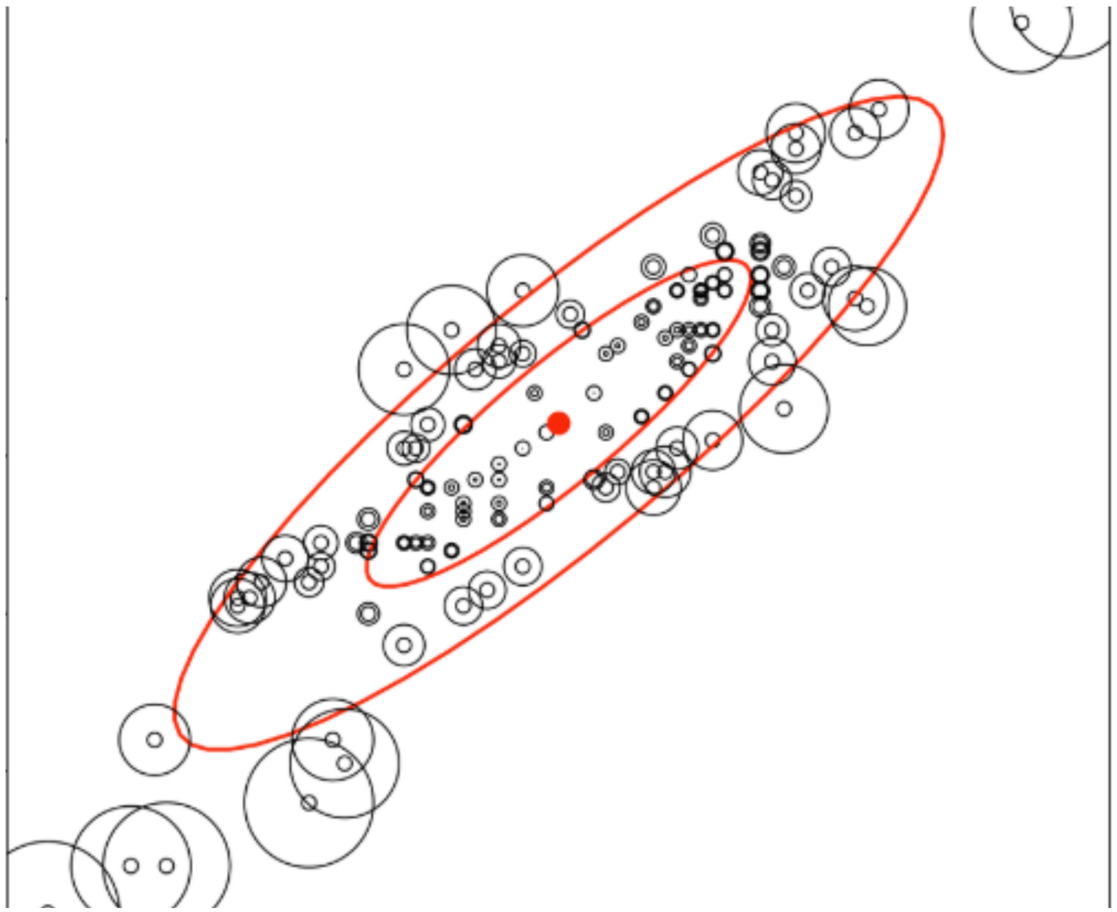

A visual representation of the Mahalanobis distance.

SAGA Supplementary Tests

- ☑ Mahalanobis distance visualises multivariate normality.
- ☑ The distance measures std. devs. of each point from distribution.
- ☑ Empirical vs Expected should follow a chi-square distribution.
- ☑ QQ-plot (+R^2 value) visualise this.

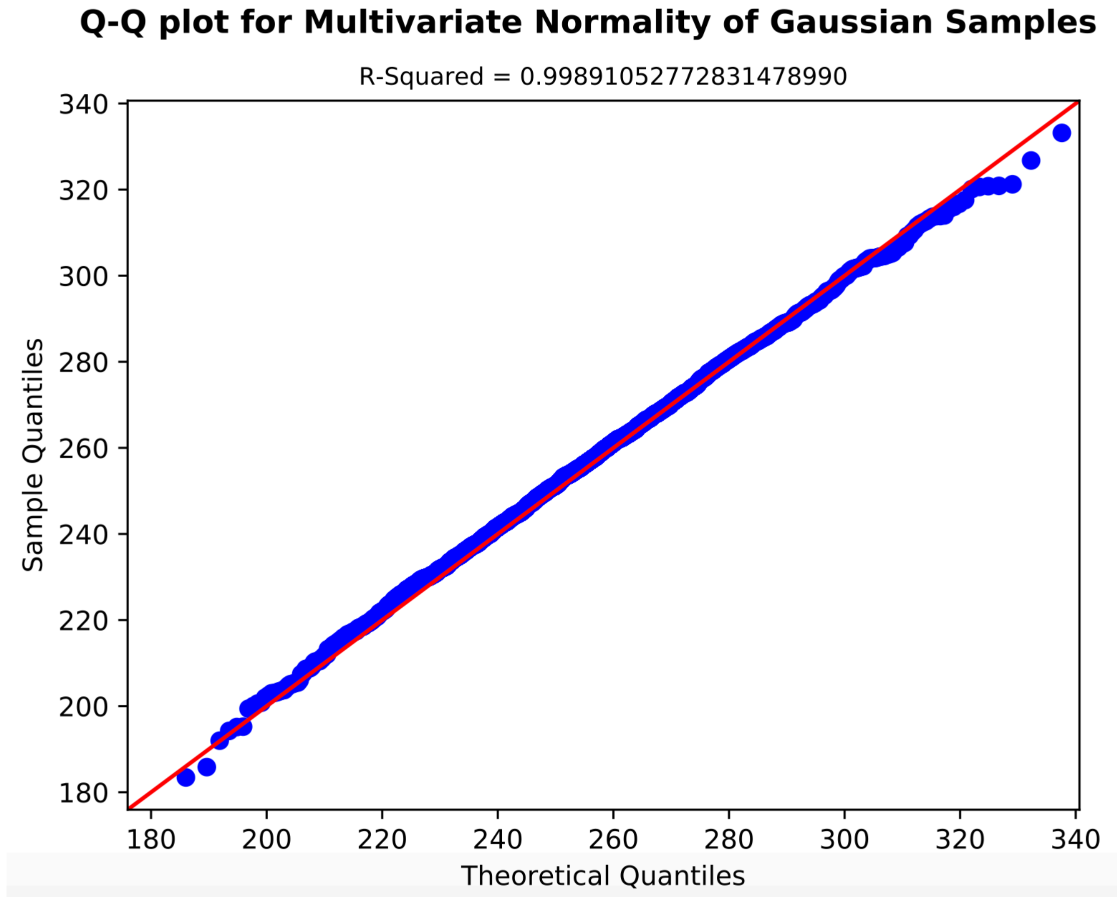

Visual multivariate normality tests.

SAGA Supplementary Tests

- ☑ Mahalanobis distance visualises multivariate normality.
- ☑ The distance measures std. devs. of each point from distribution.
- ☑ Empirical vs Expected should follow a chi-square distribution.
- ☑ QQ-plot (+R^2 value) visualise this.

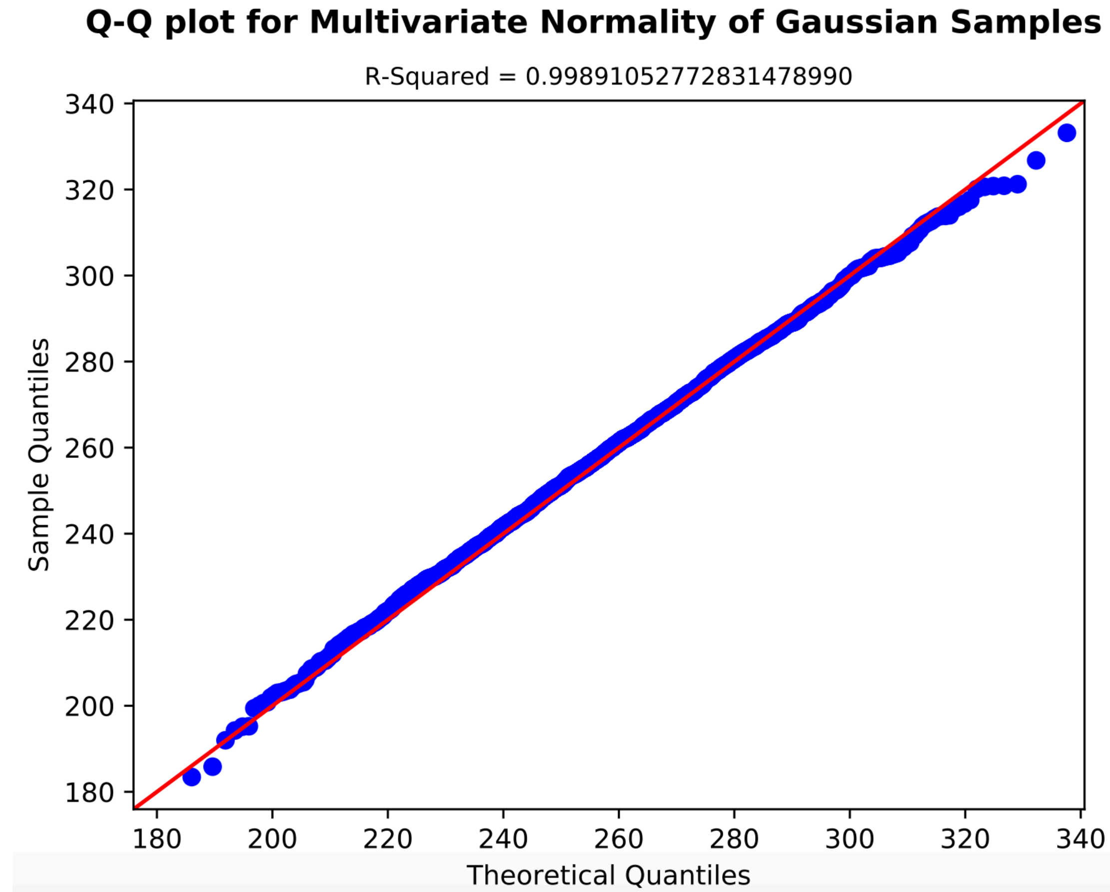

Visual multivariate normality tests.

SAGA Supplementary Tests

- ☑ Mahalanobis distance visualises multivariate normality.
- ☑ The distance measures std. devs. of each point from distribution.
- ☑ Empirical vs Expected should follow a chi-square distribution.
- ☑ QQ-plot (+R^2 value) visualise this.

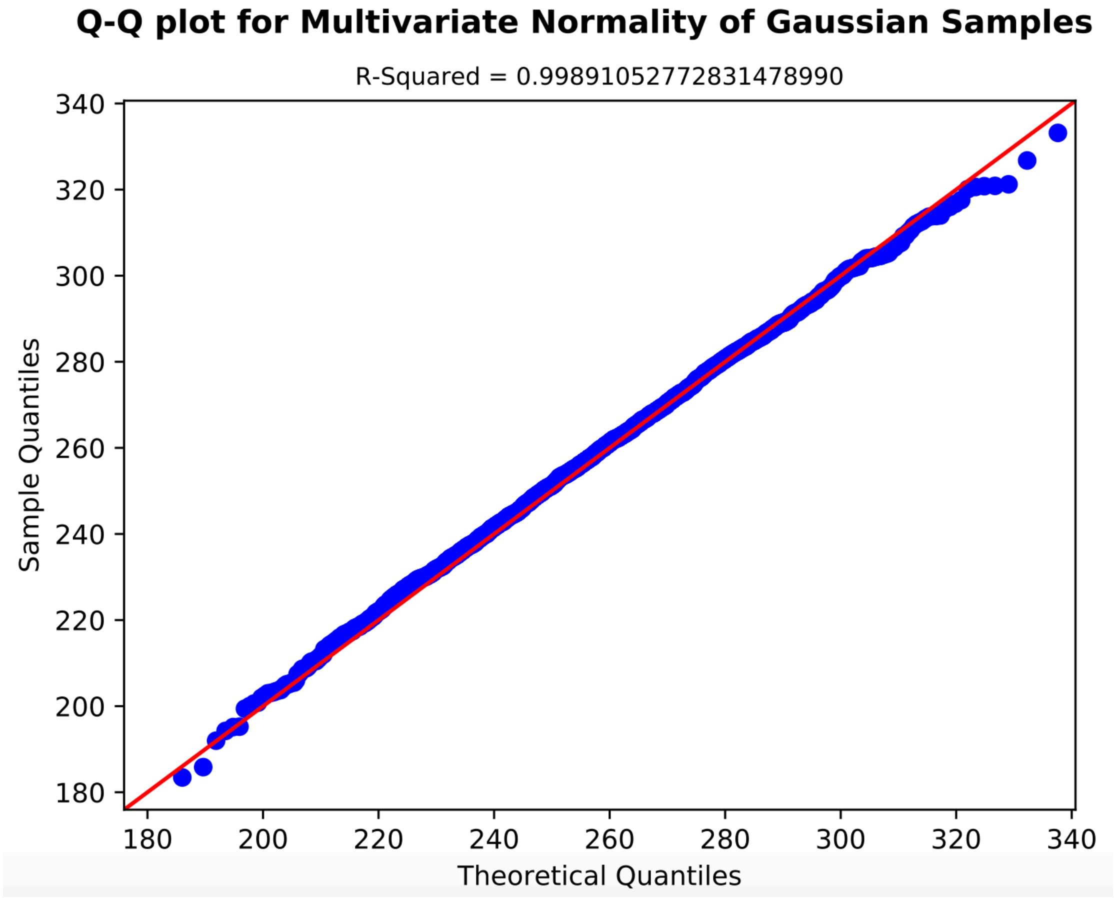

- ☑ Rejections modelled to observe the geometric decrease.

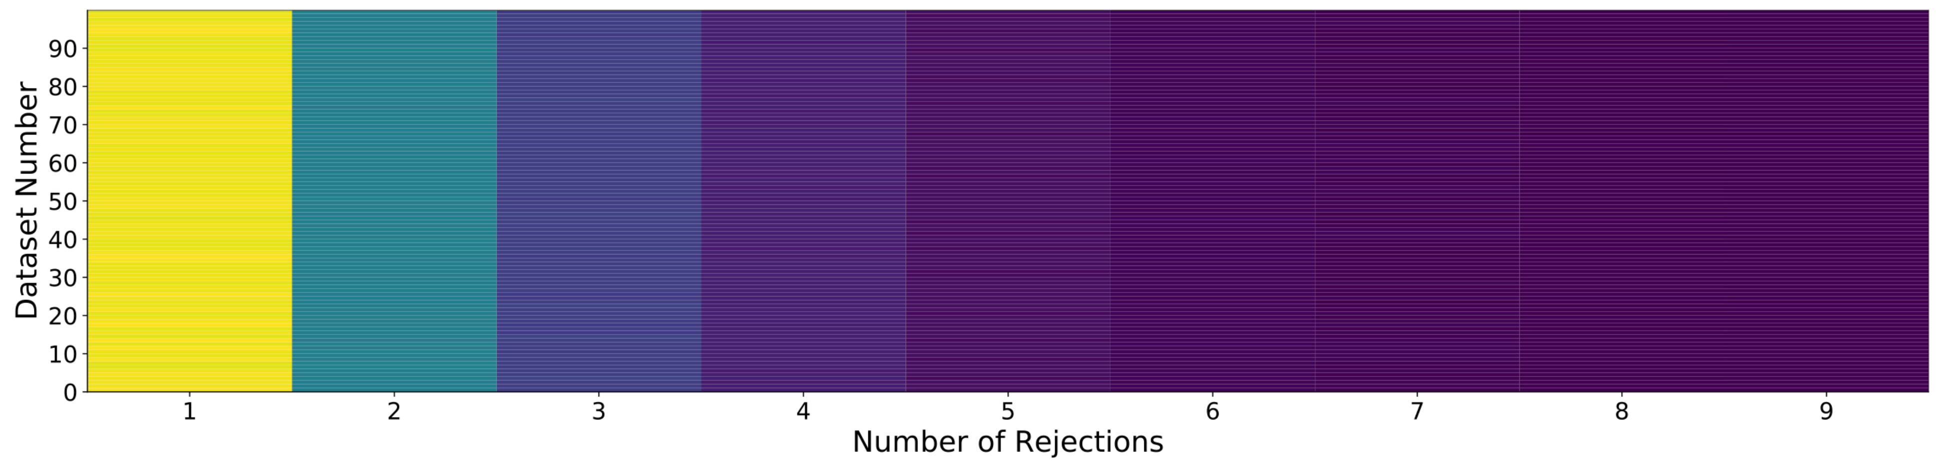

SAGA Supplementary Tests

- ☑ Mahalanobis distance visualises multivariate normality.
- ☑ The distance measures std. devs. of each point from distribution.
- ☑ Empirical vs Expected should follow a chi-square distribution.
- ☑ QQ-plot (+R^2 value) visualise this.

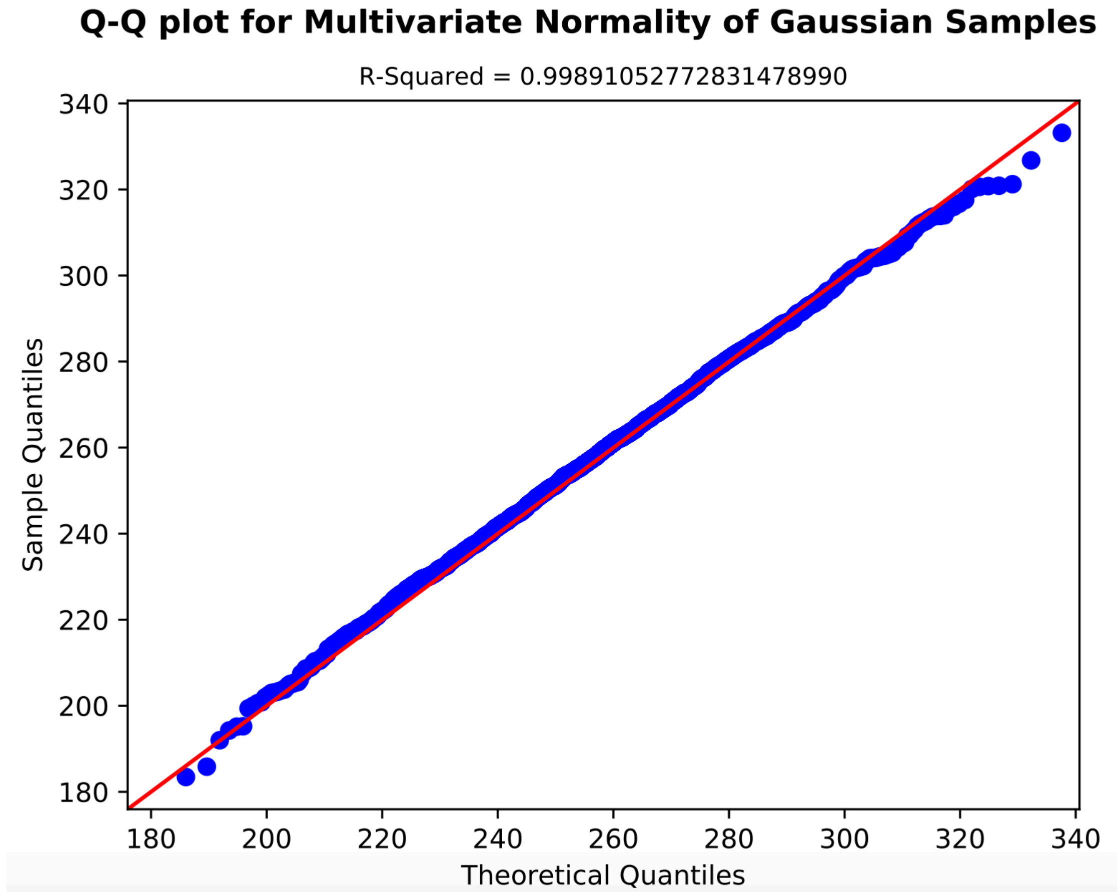

- ☑ Rejections modelled to observe the geometric decrease.

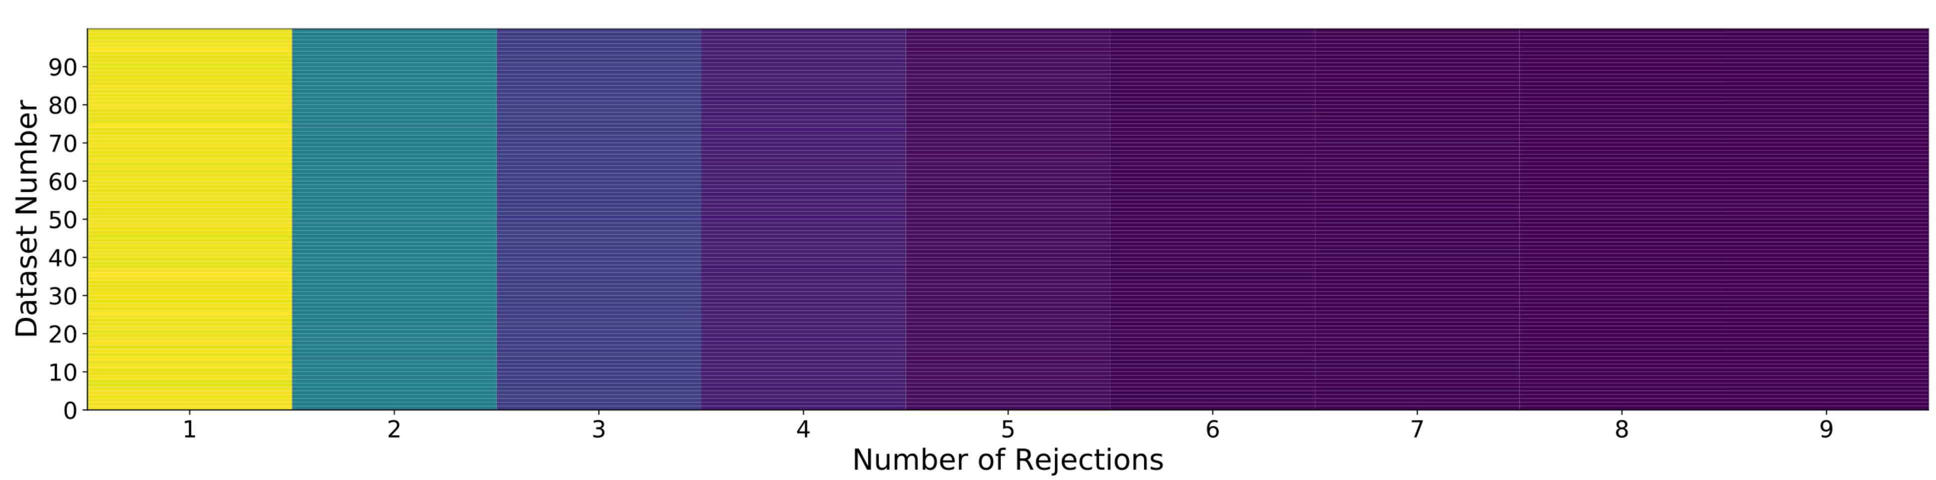

Implementations

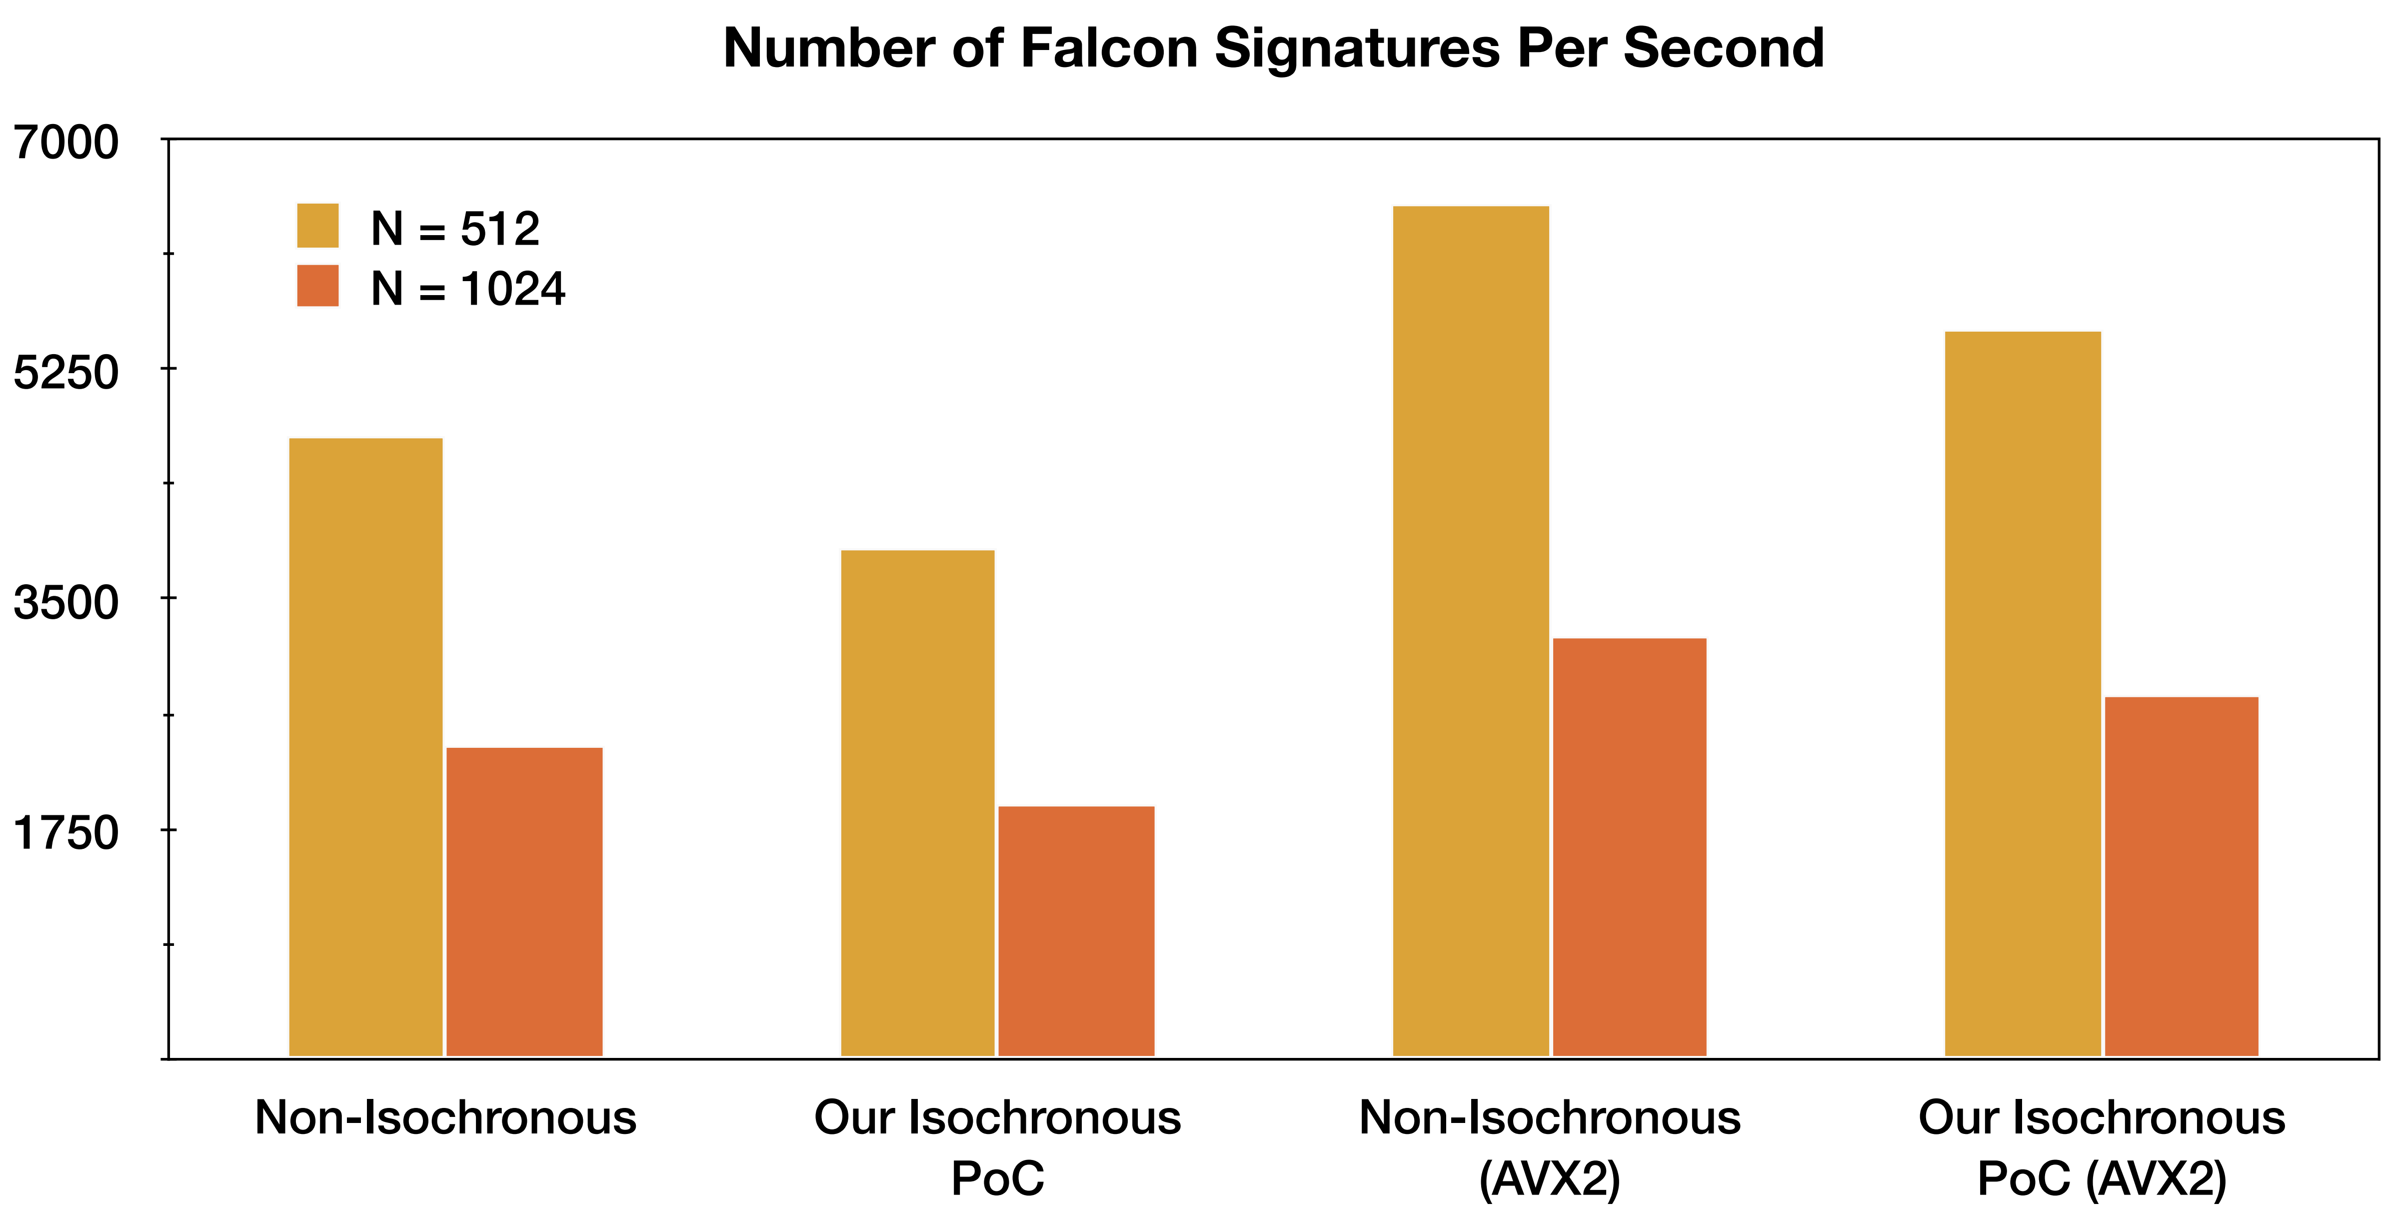

- ☑ Our sampler in Falcon on one Intel Core i7-6500U CPU @2.5GHz.
- ☑ The performance loss for isochrony is minimal (13% - 18%).

Implementations

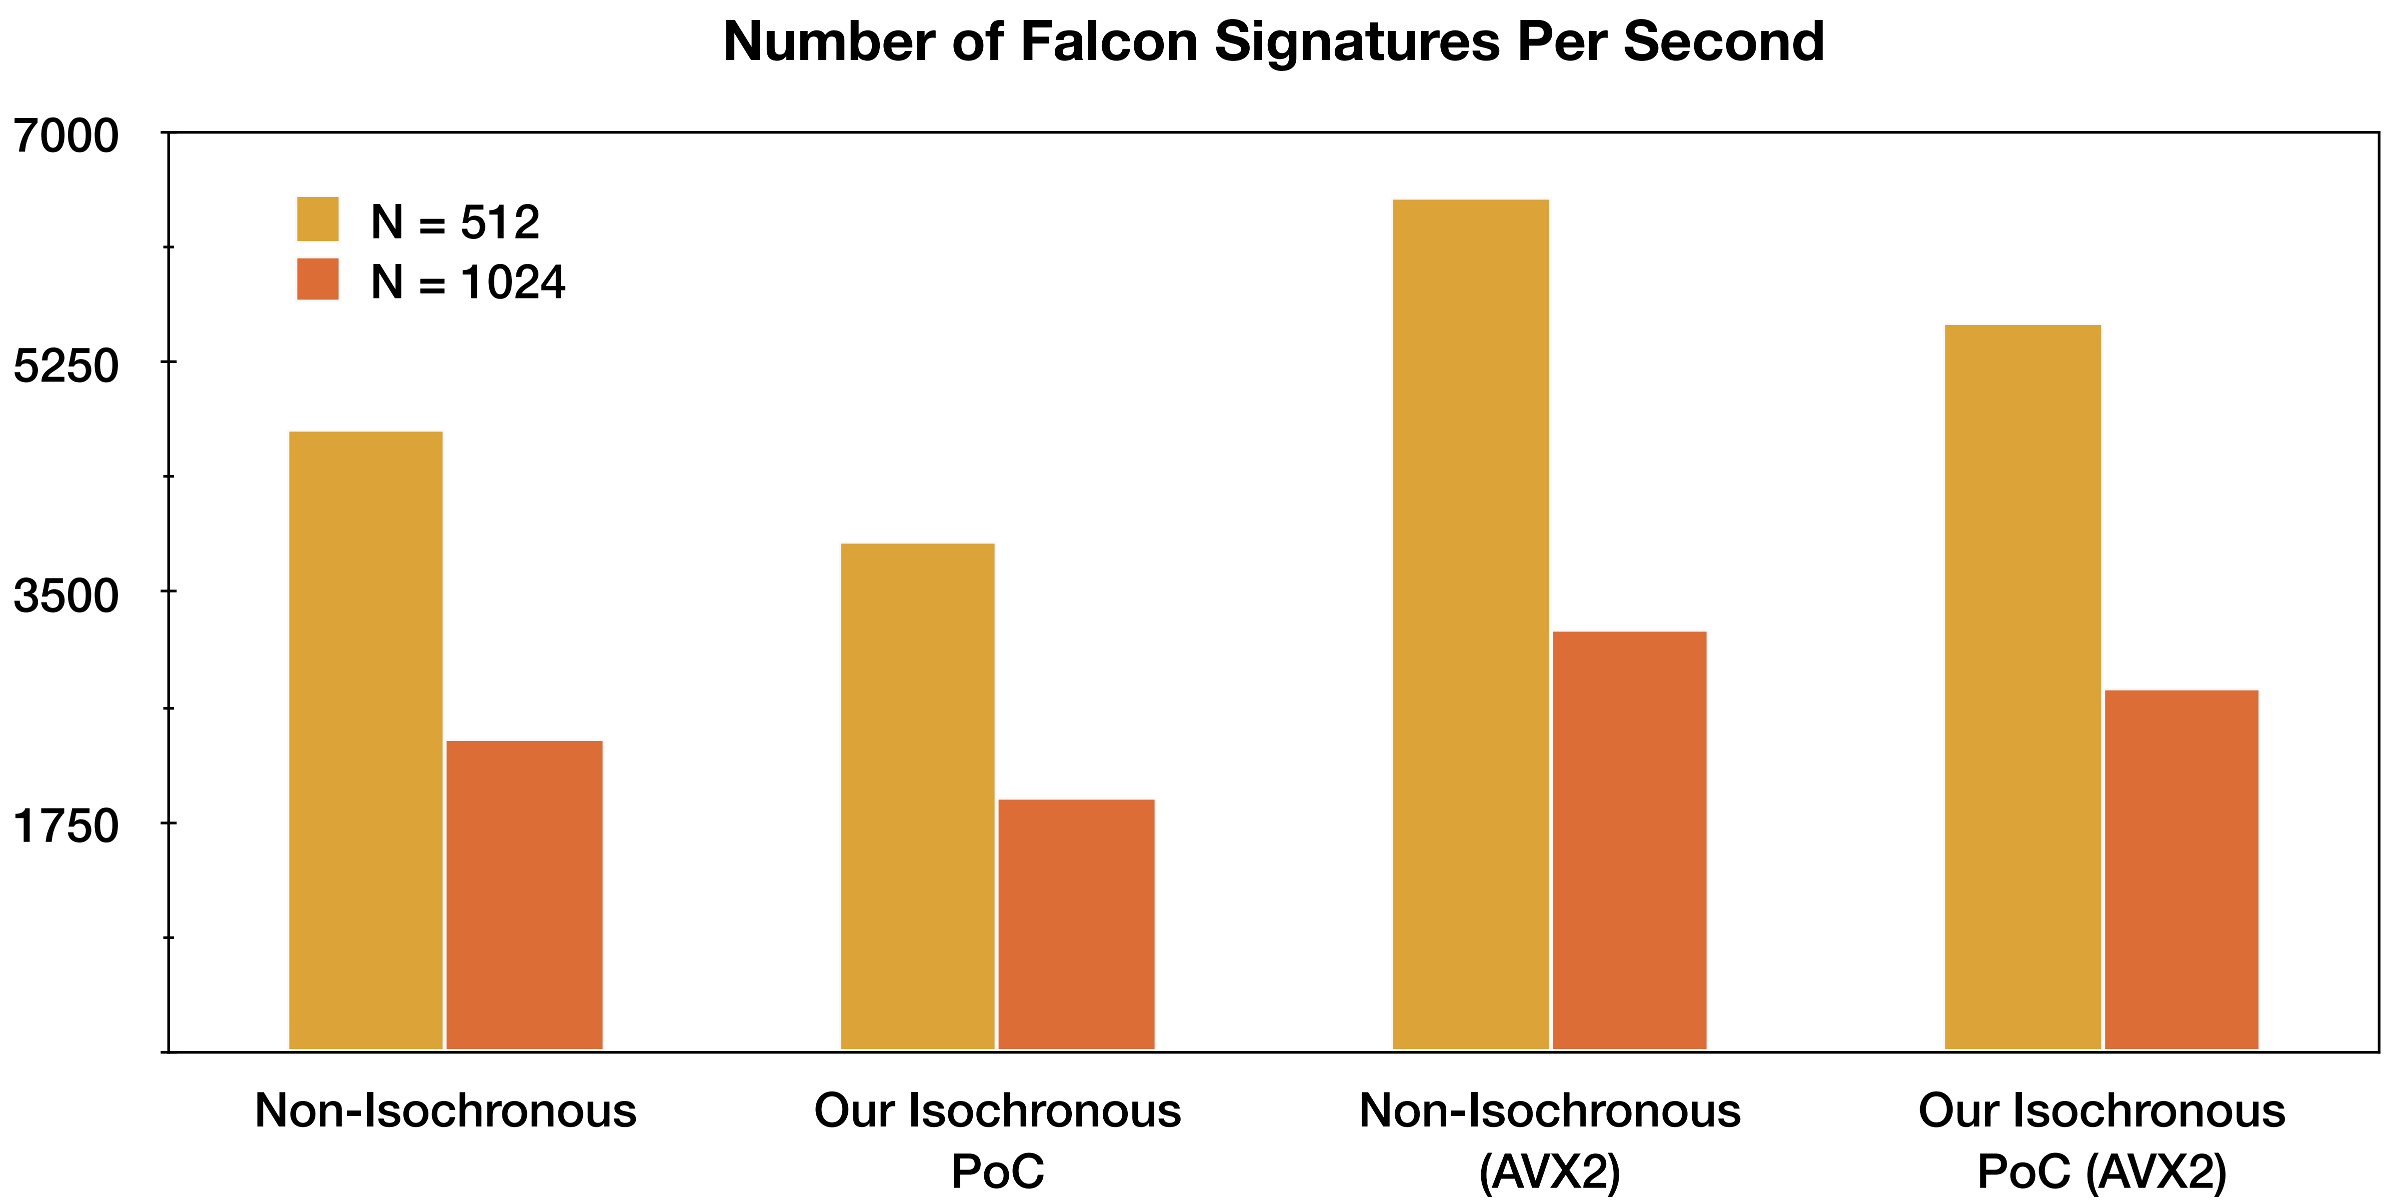

- ☑ Our sampler in Falcon on one Intel Core i7-6500U CPU @2.5GHz.
- ☑ The performance loss for isochrony is minimal (13% - 18%).

Thanks for Listening