# Monte Carlo Integration II: Multiple Importance Sampling

CS-E5520 Spring 2023
Jaakko Lehtinen
with many slides from Frédo Durand

# Recap: Reflectance Equation

$$L(x \rightarrow \mathbf{v}) =$$

**outgoing radiance**

$$\int_{\Omega} f_r(x, \mathbf{l} \rightarrow \mathbf{v}) L(x \leftarrow \mathbf{l}) \cos\theta \, \mathrm{dl}$$

**integral over hemisphere**

**BRDF**

**incoming radiance**

**cosine of incident angle**

**L_in * cos = incident differential irradiance**

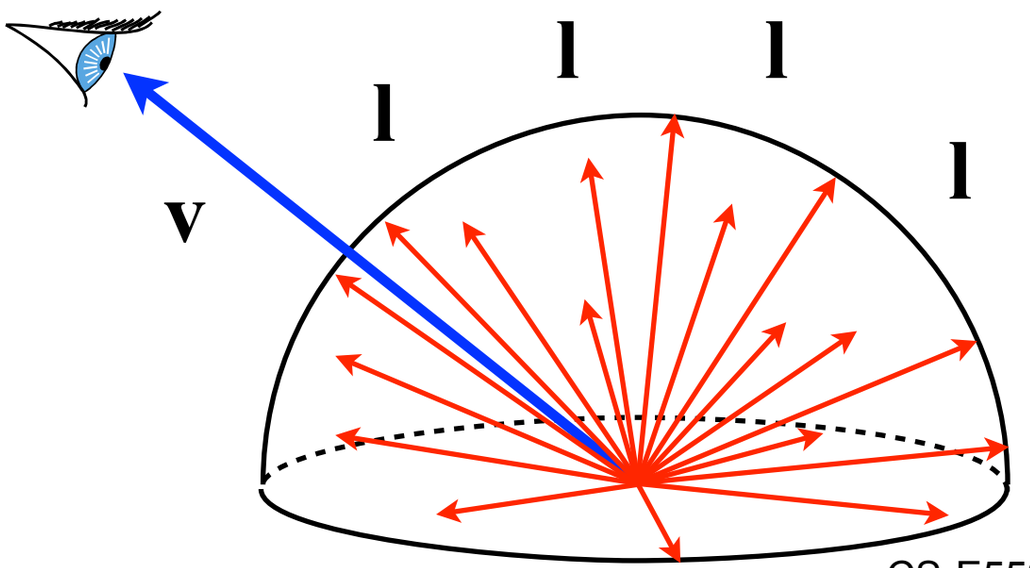

# Recap: Reflectance Equation

$$L(x \rightarrow \mathbf{v}) =$$

$$\int_{\Omega} f_r(x, \mathbf{l} \rightarrow \mathbf{v})\, L(x \leftarrow \mathbf{l})\, \cos\theta \, \mathrm{d}\mathbf{l}$$

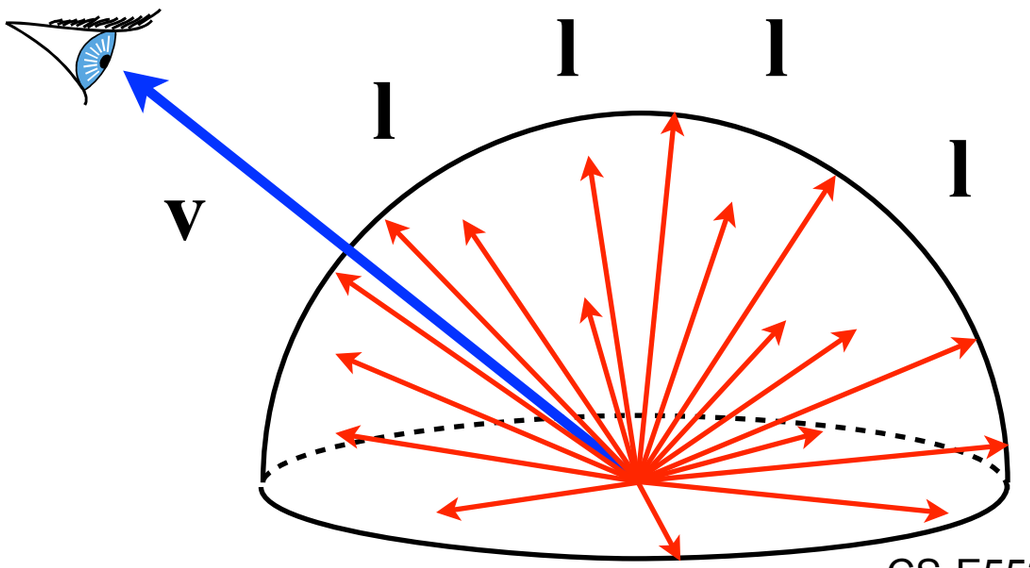

# Imp. Sampling According to BRDF

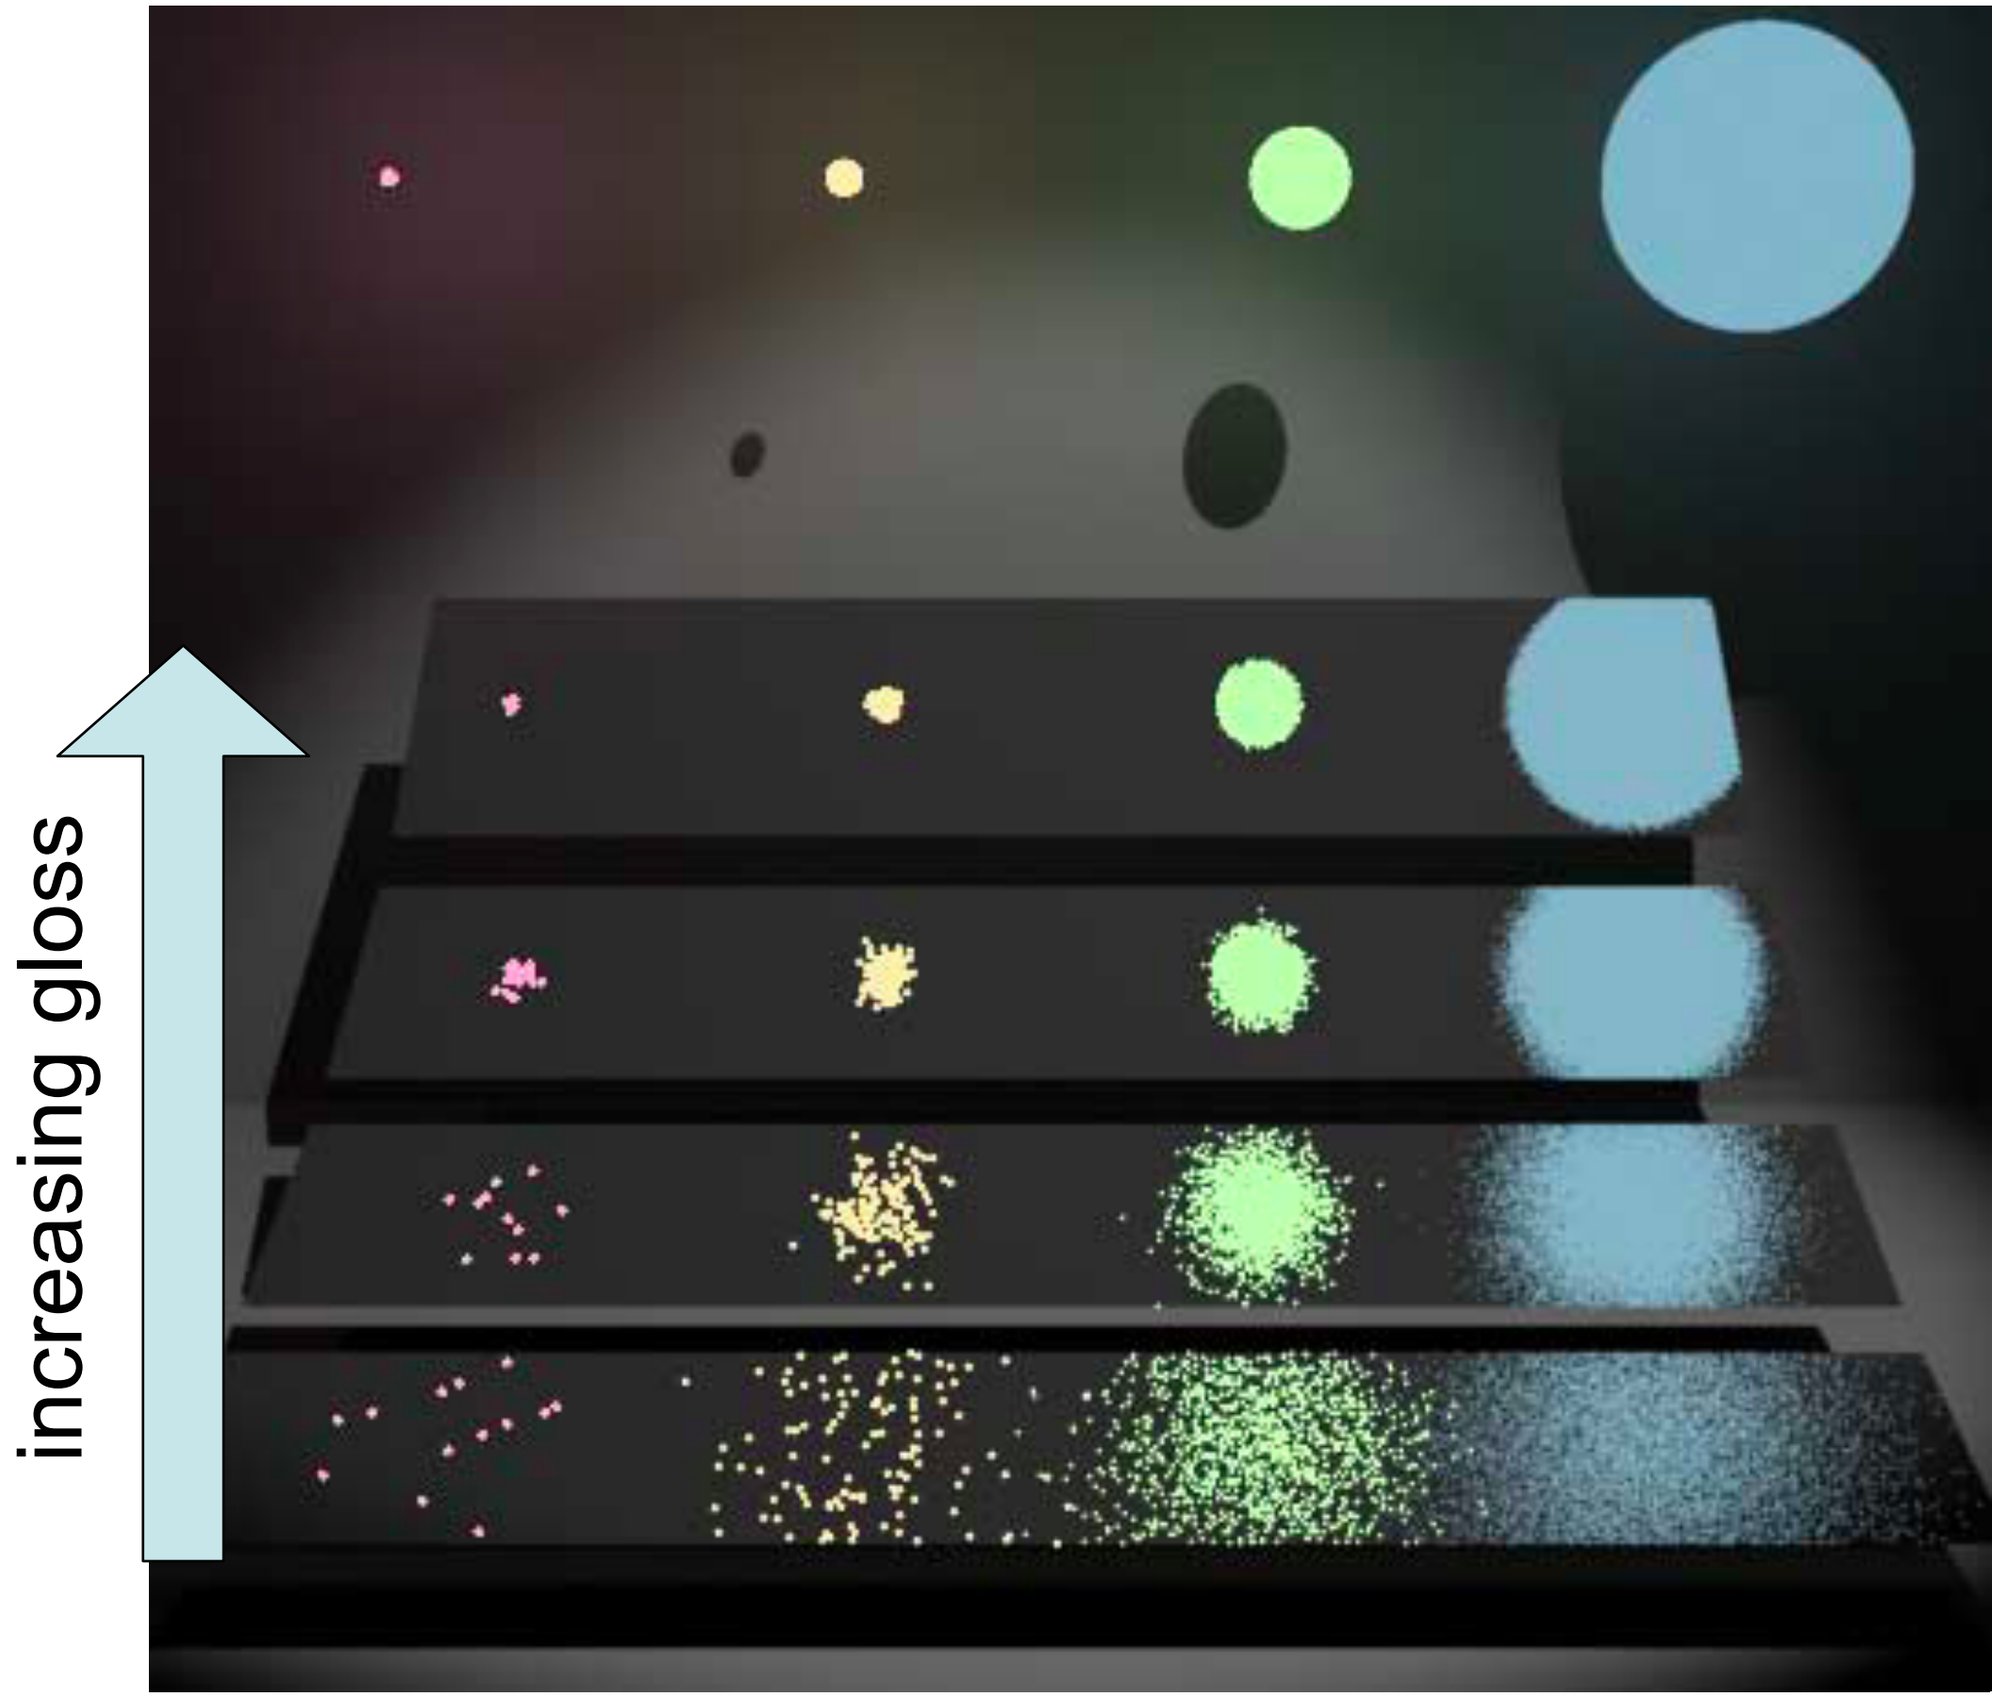

# Imp. Sampling According to Light

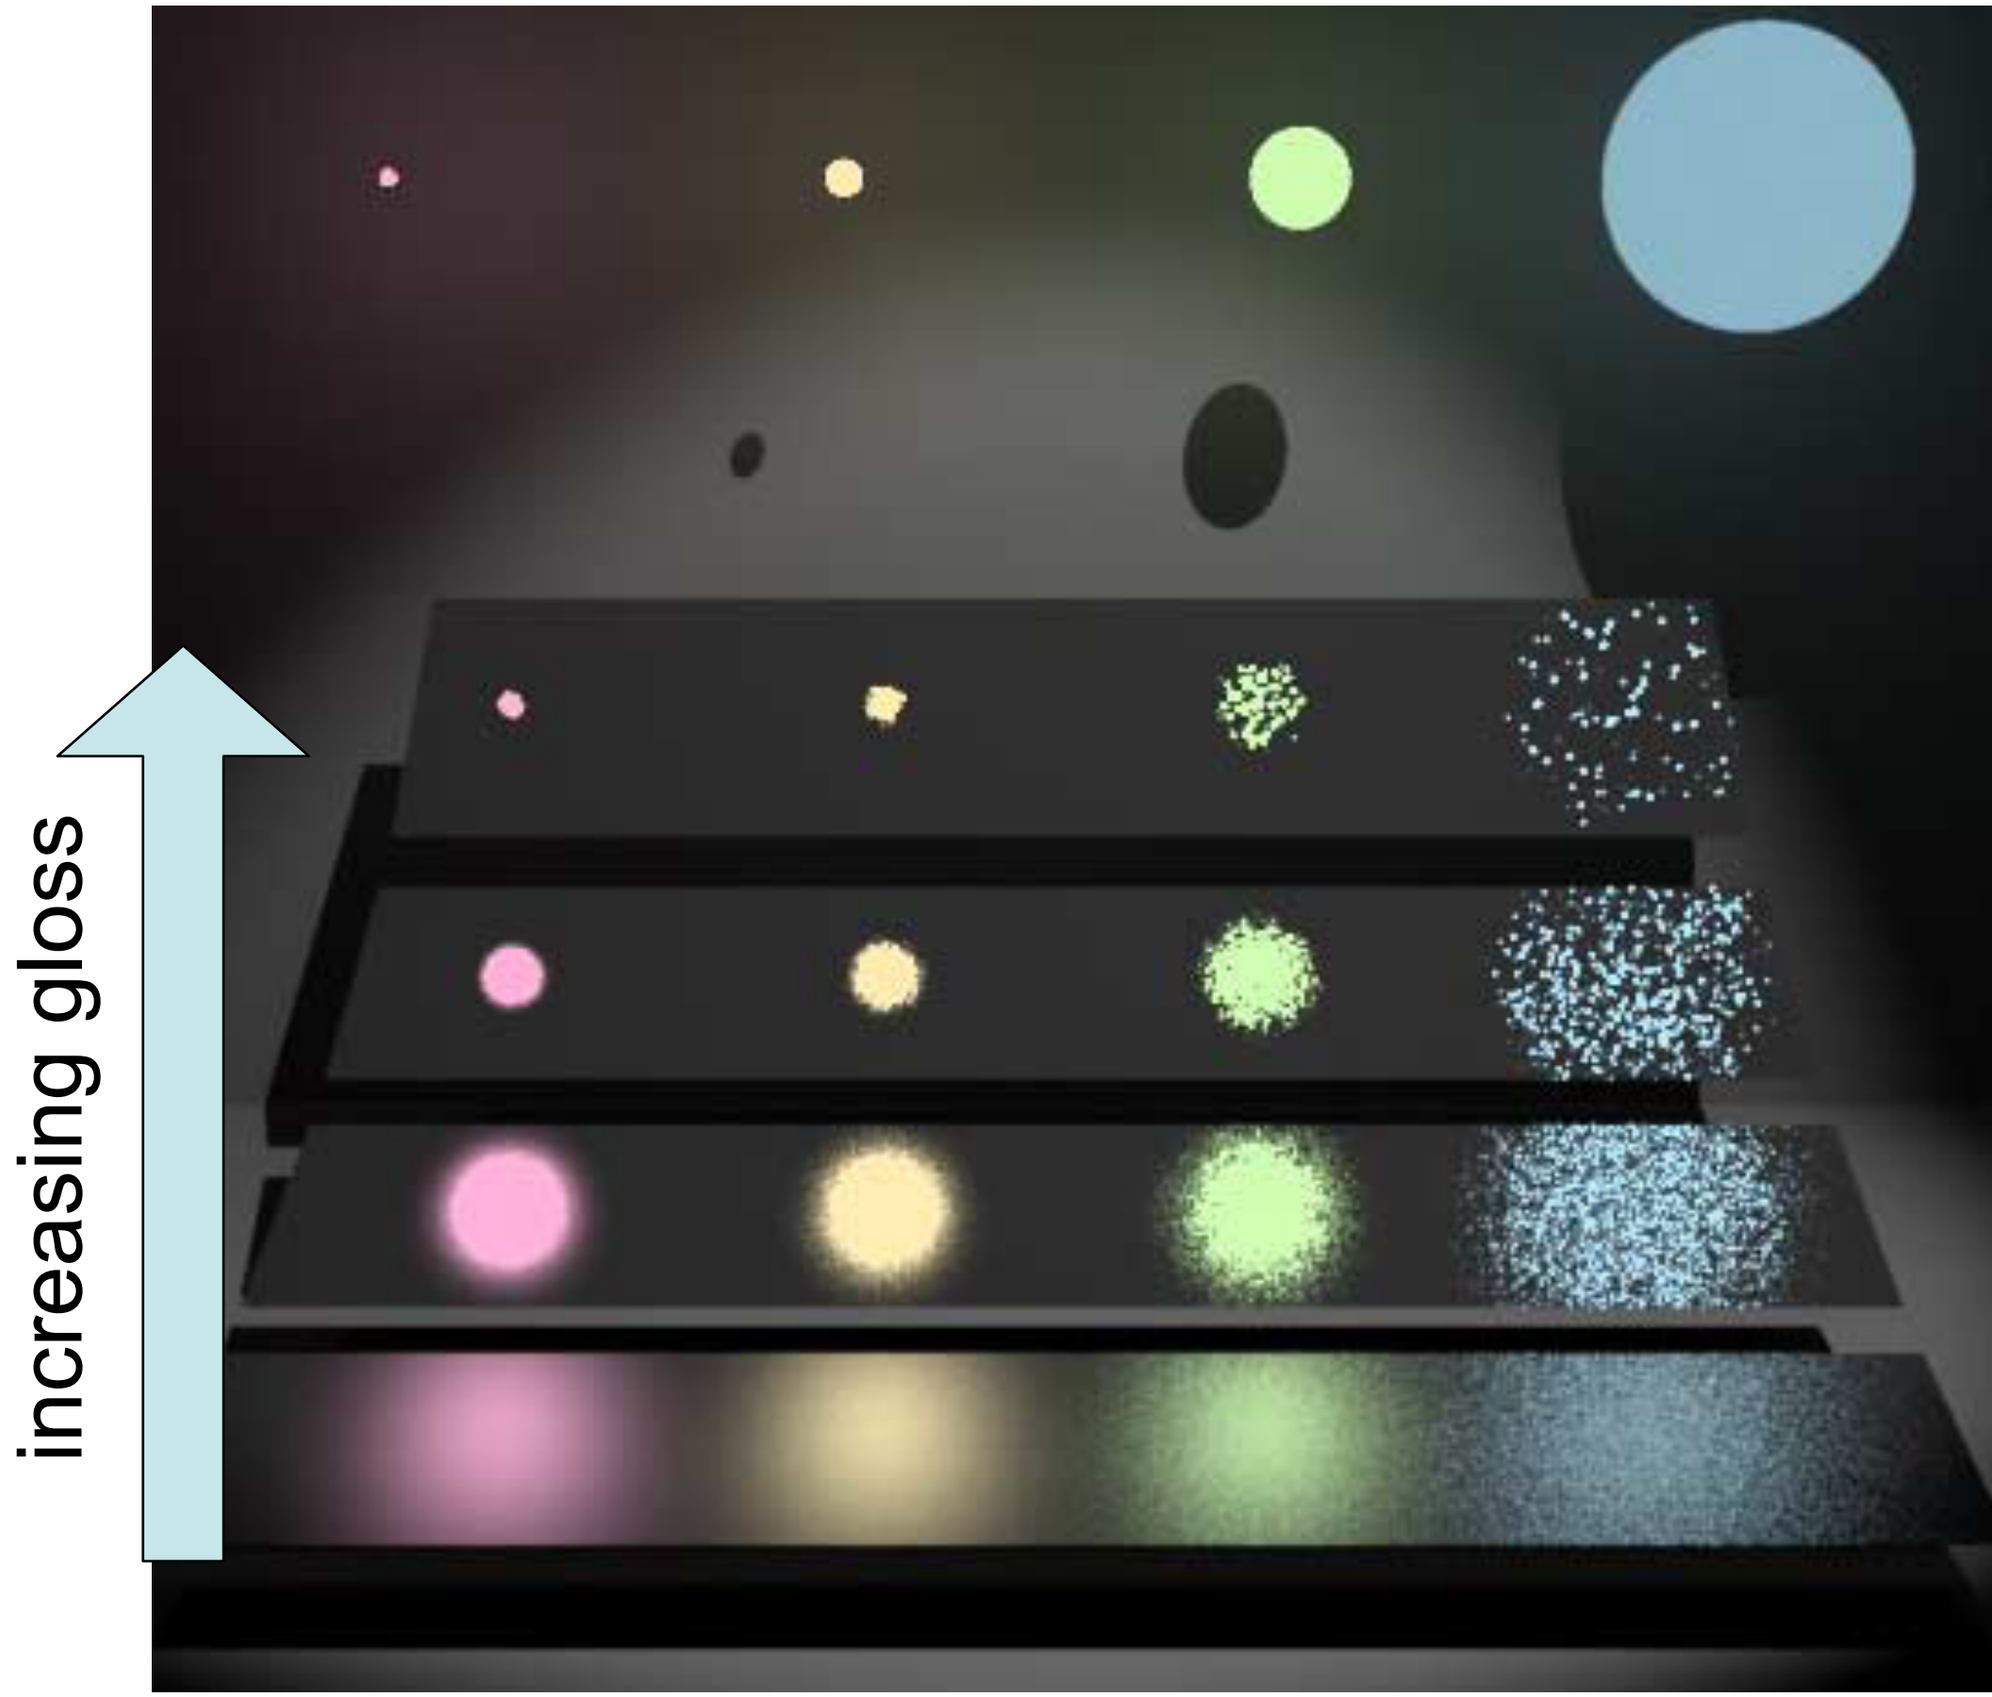

# Multiple Importance Sampling (MIS)

- If integrand *f* has a complex shape that consists of distinct features that are easy to sample from individually, we can use multiple PDFs and combine them in a nice way so that we got lower variance
  - See Veach and Guibas 1995

# What's Going on Here?

- Dull gloss/diffuse surface, importance sample BRDF

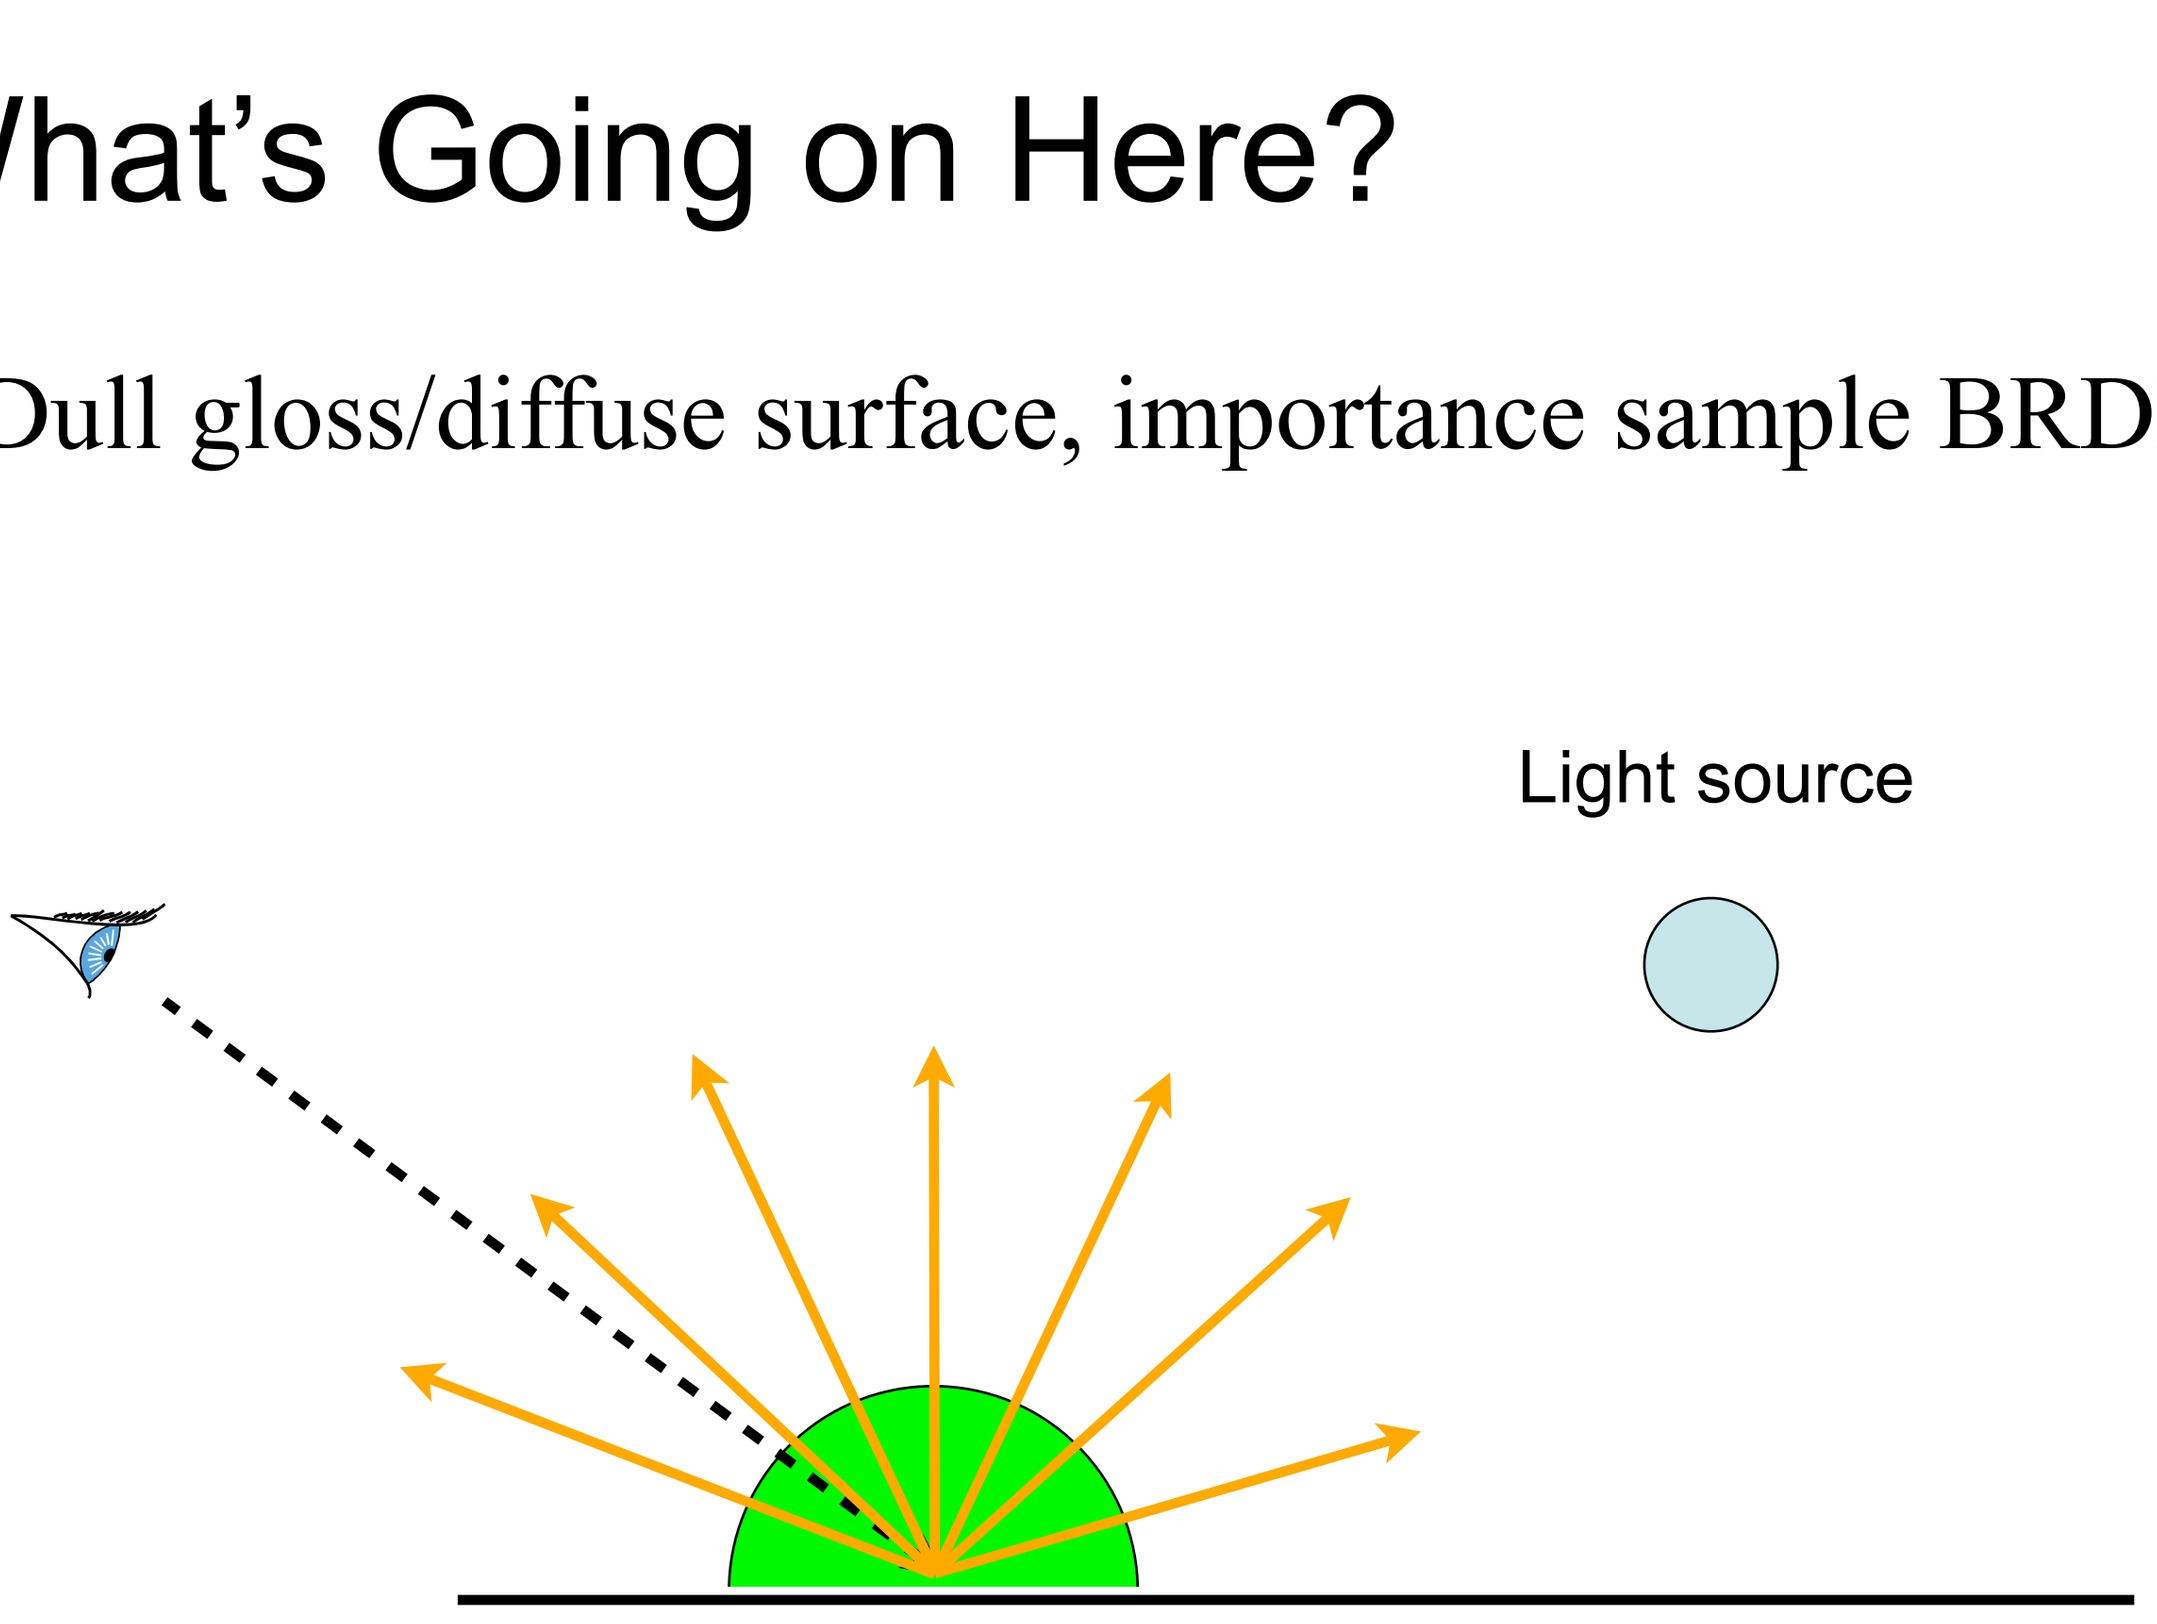

# What's Going on Here?

- Dull gloss/diffuse surface, importance sample BRDF

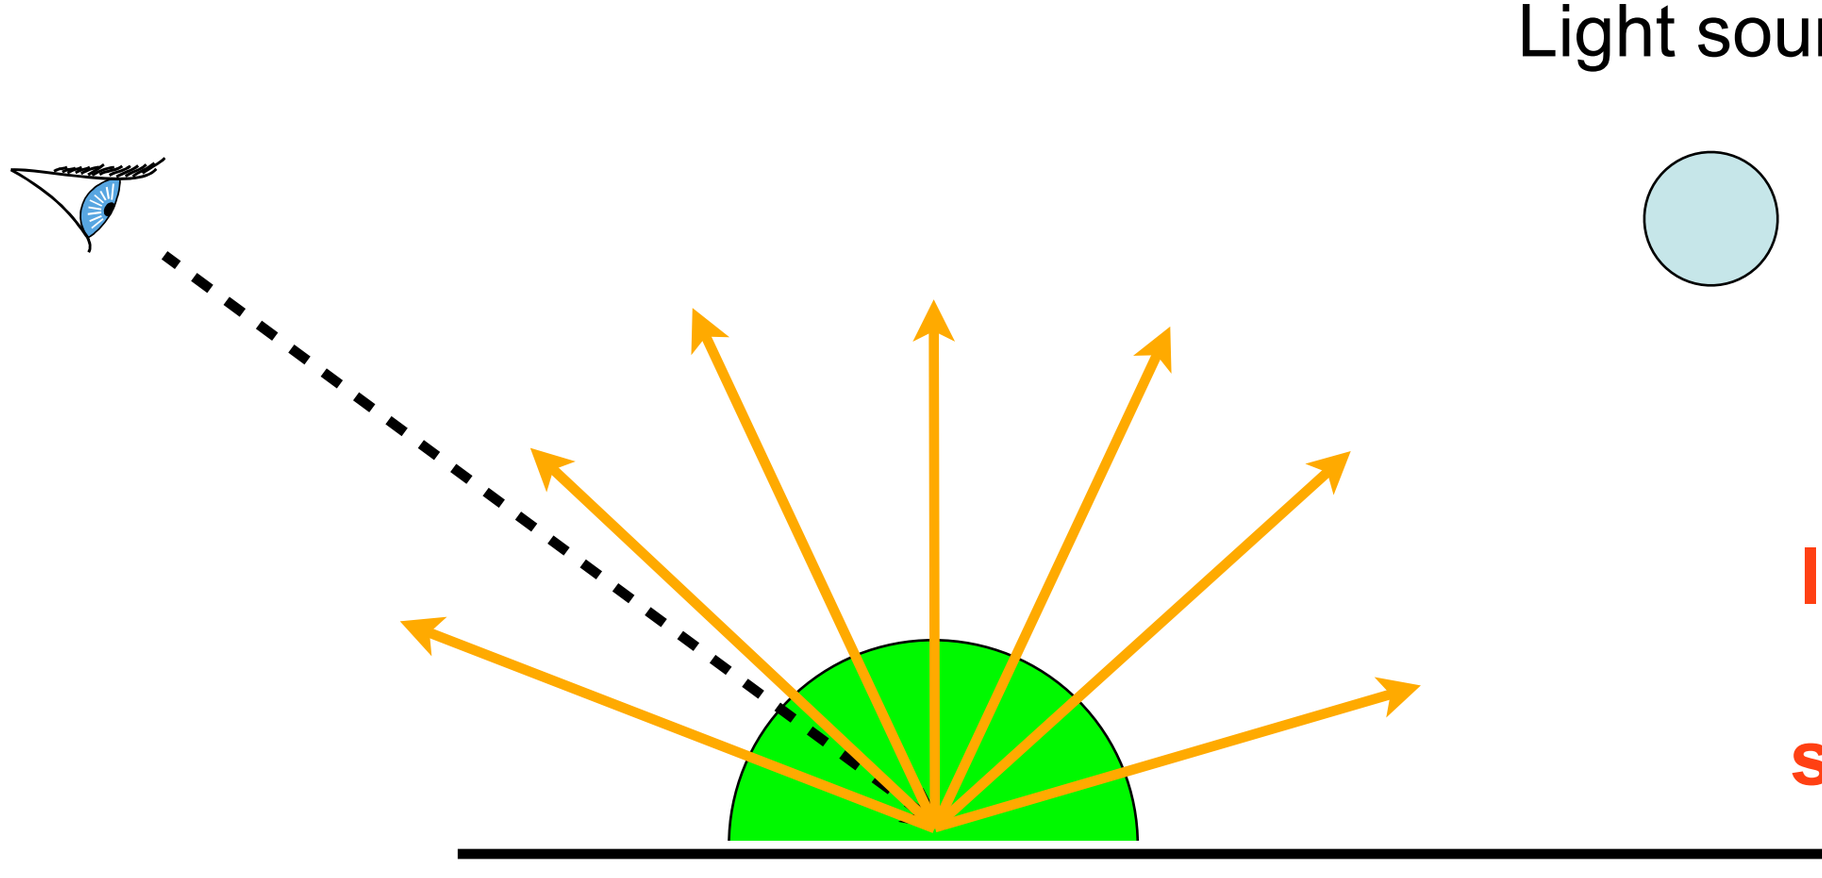


**Only few directions actually carry light, so we are using our samples poorly**

# Here Makes Sense to Sample Light

- Dull gloss/diffuse surface, importance sample light

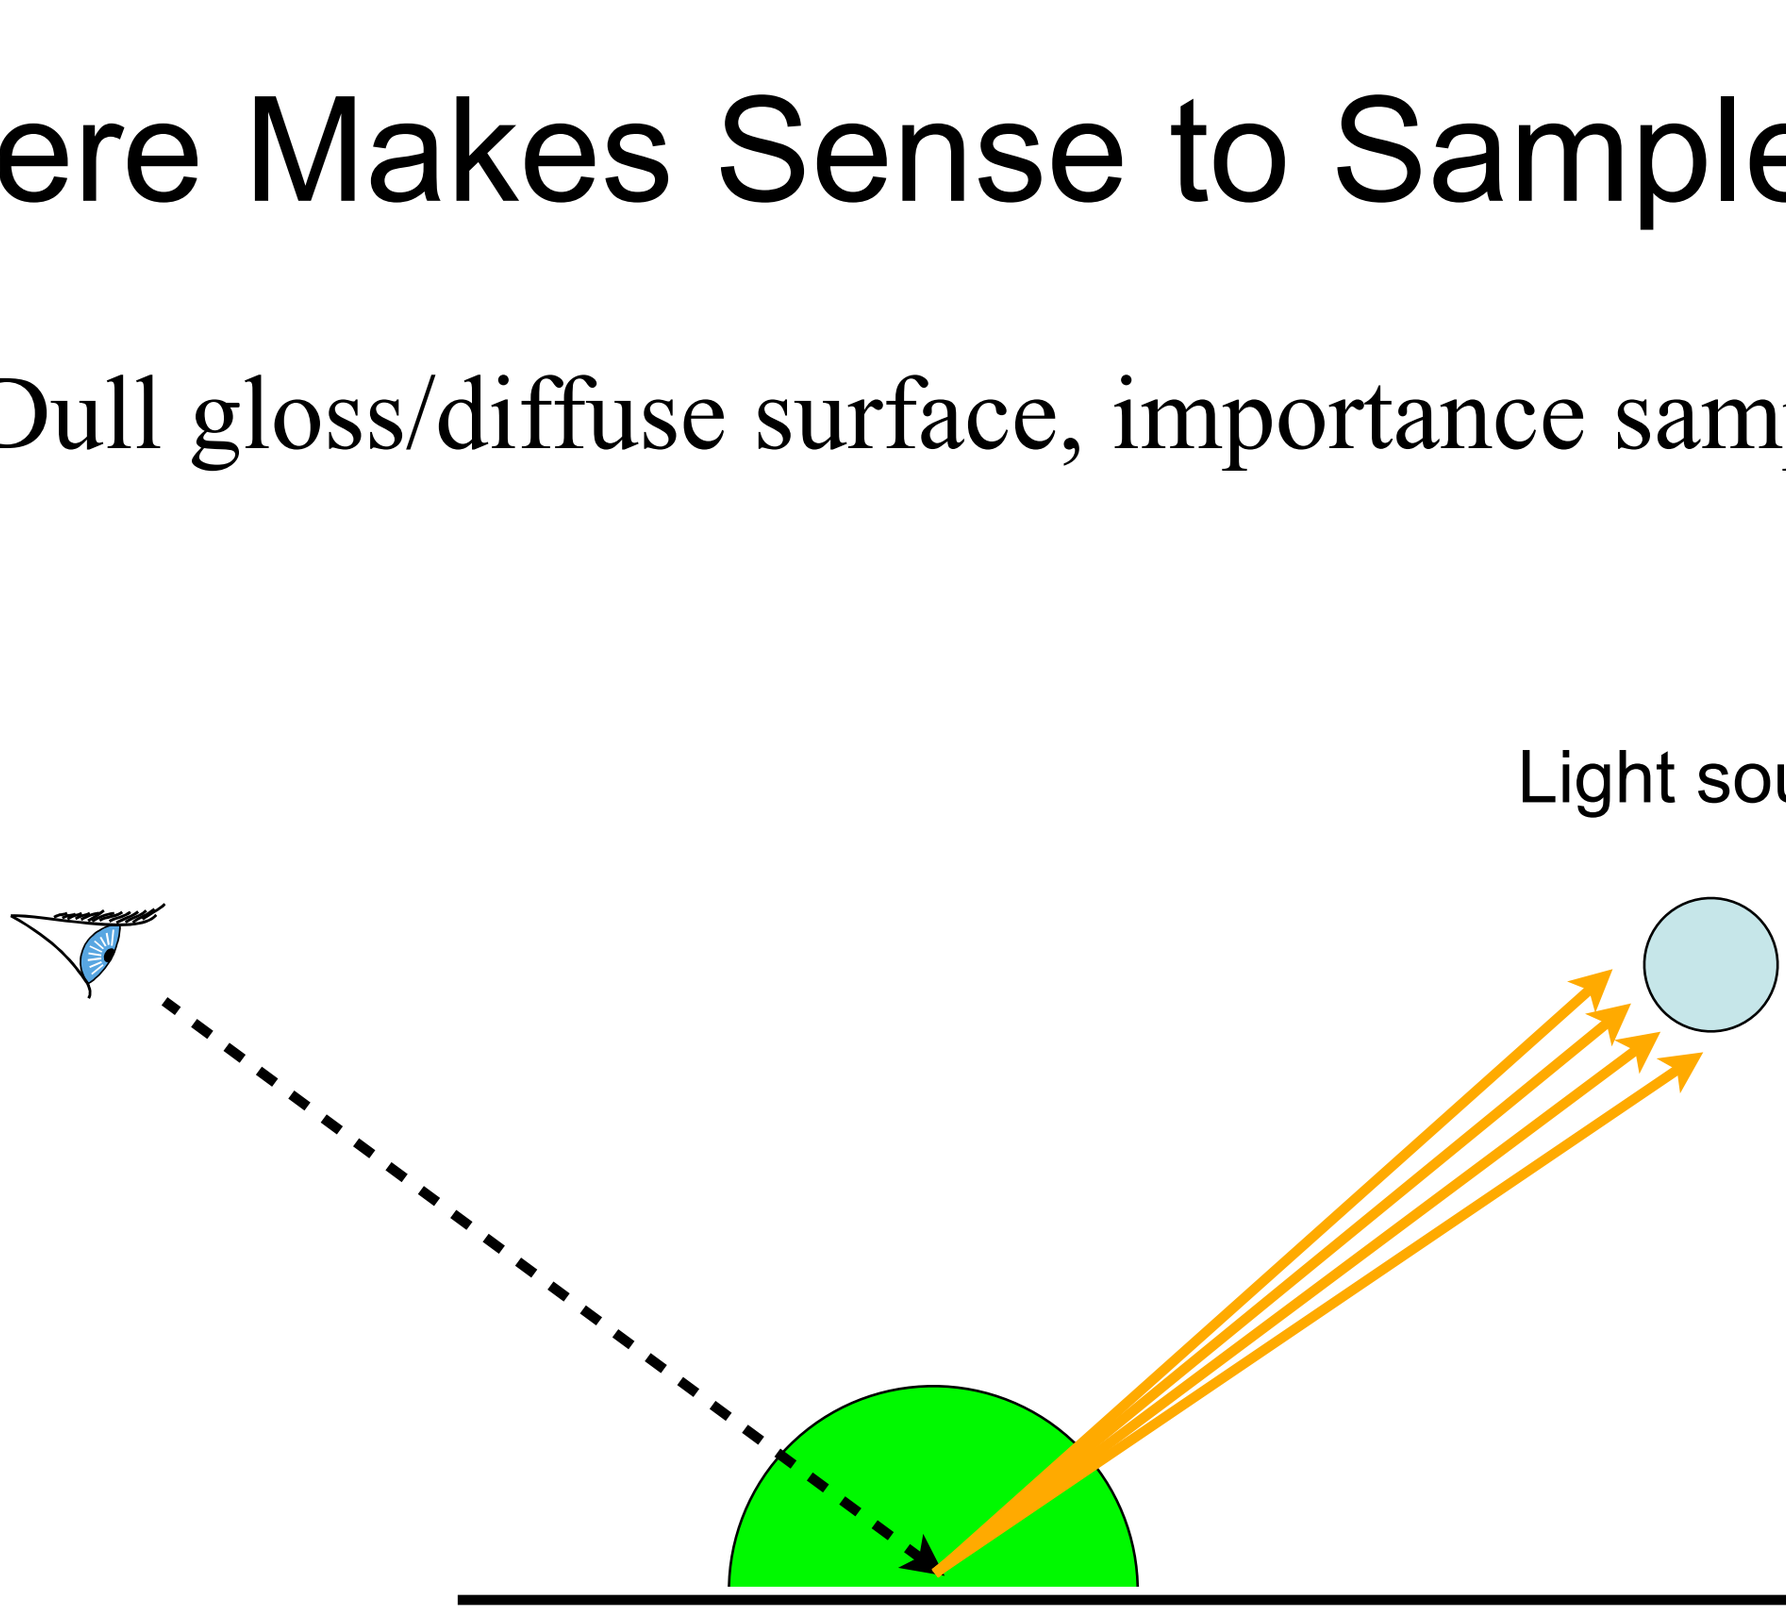


**Sampling the light uses sample budget better**

# What's Going on Here?

- Highly glossy surface, narrow lobe, large light source, importance sample light

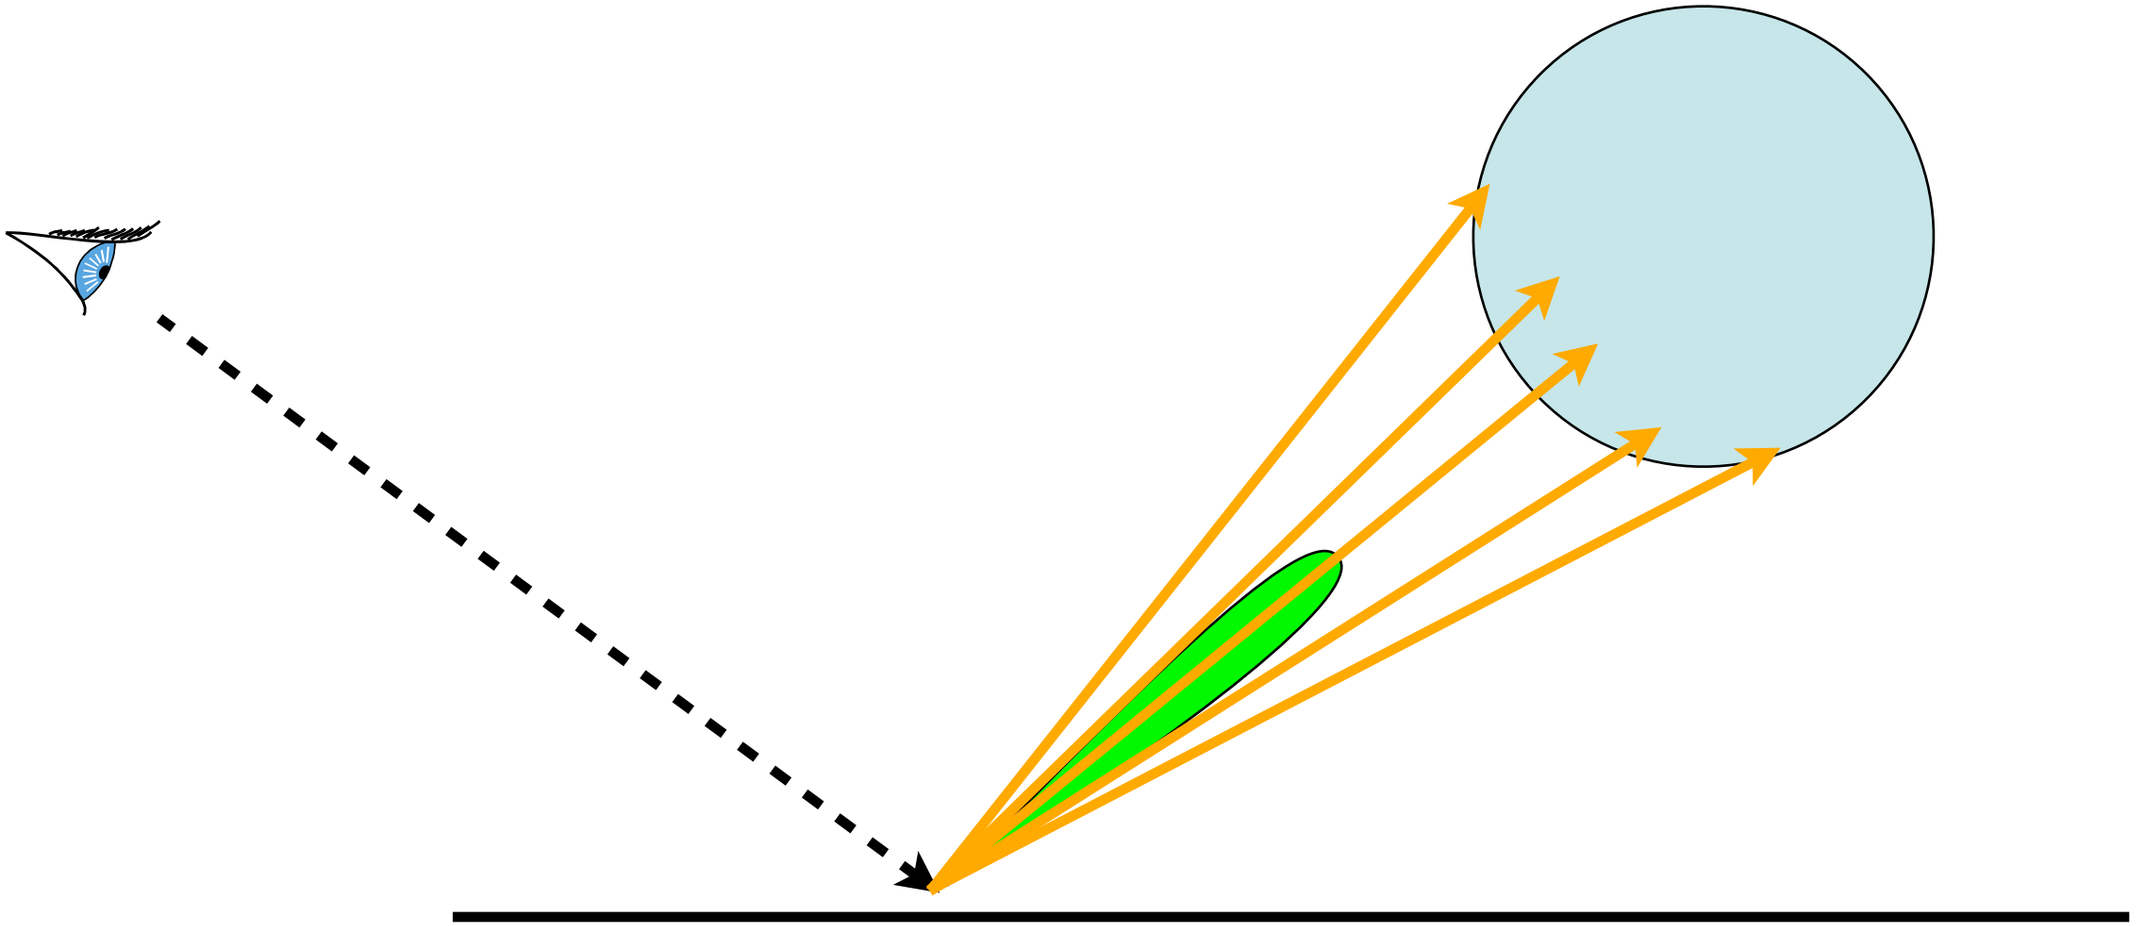

# What's Going on Here?

- Highly glossy surface, narrow lobe, large light source, importance sample light

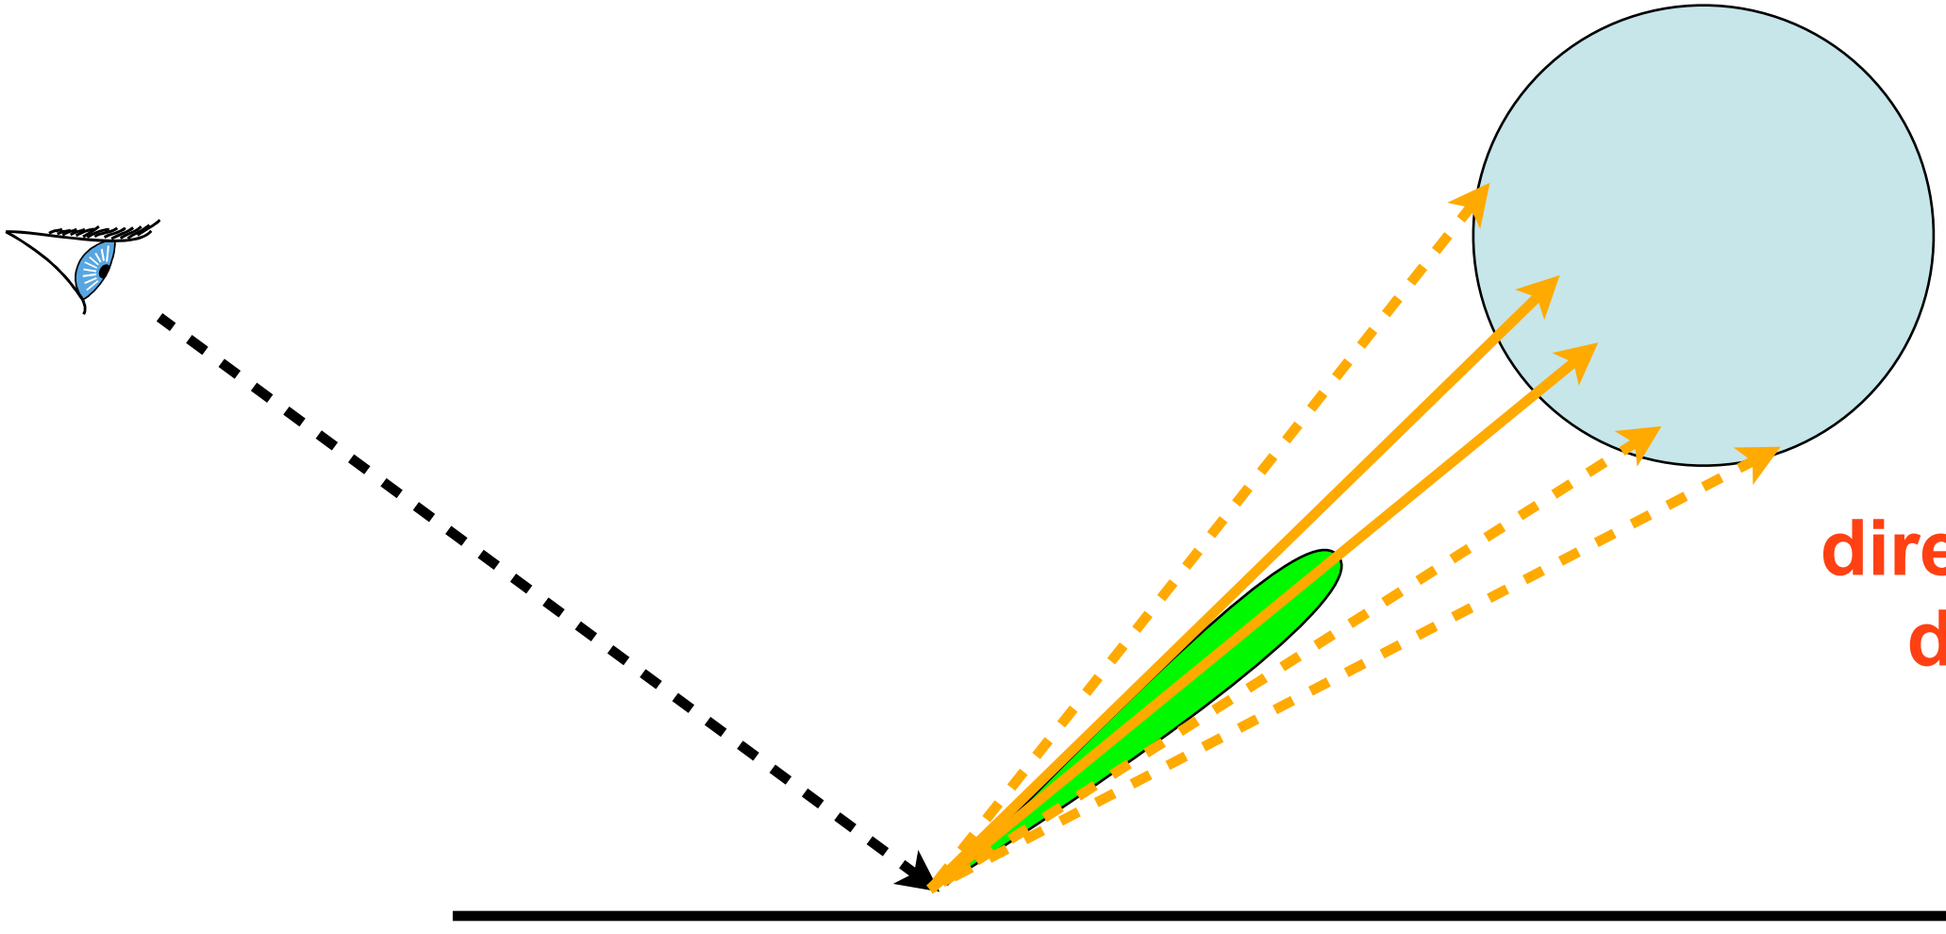

**Again we sample directions that do not carry irradiance**

# Here, Better to Sample BRDF

- Highly glossy surface, narrow lobe, large light source, importance sample light

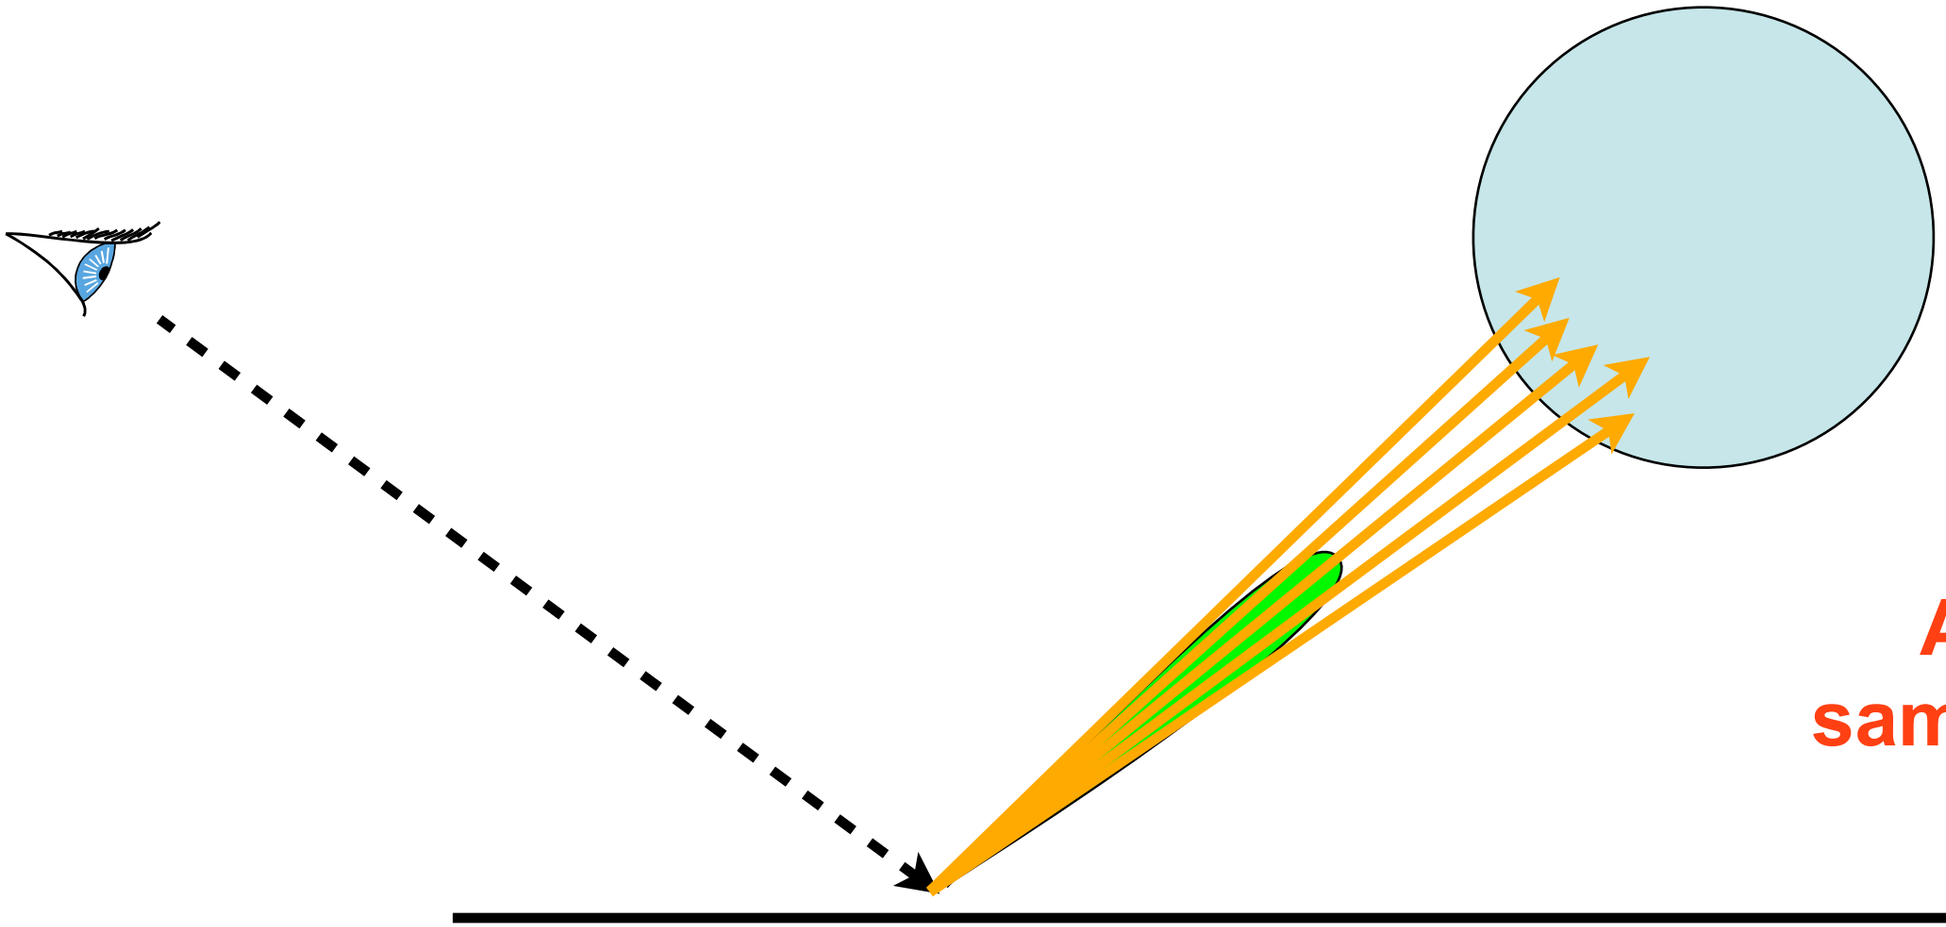

**Again, each sample counts more**

# Multiple Importance Sampling ✔

**MIS = Sample both ways and optimally combine the samples**

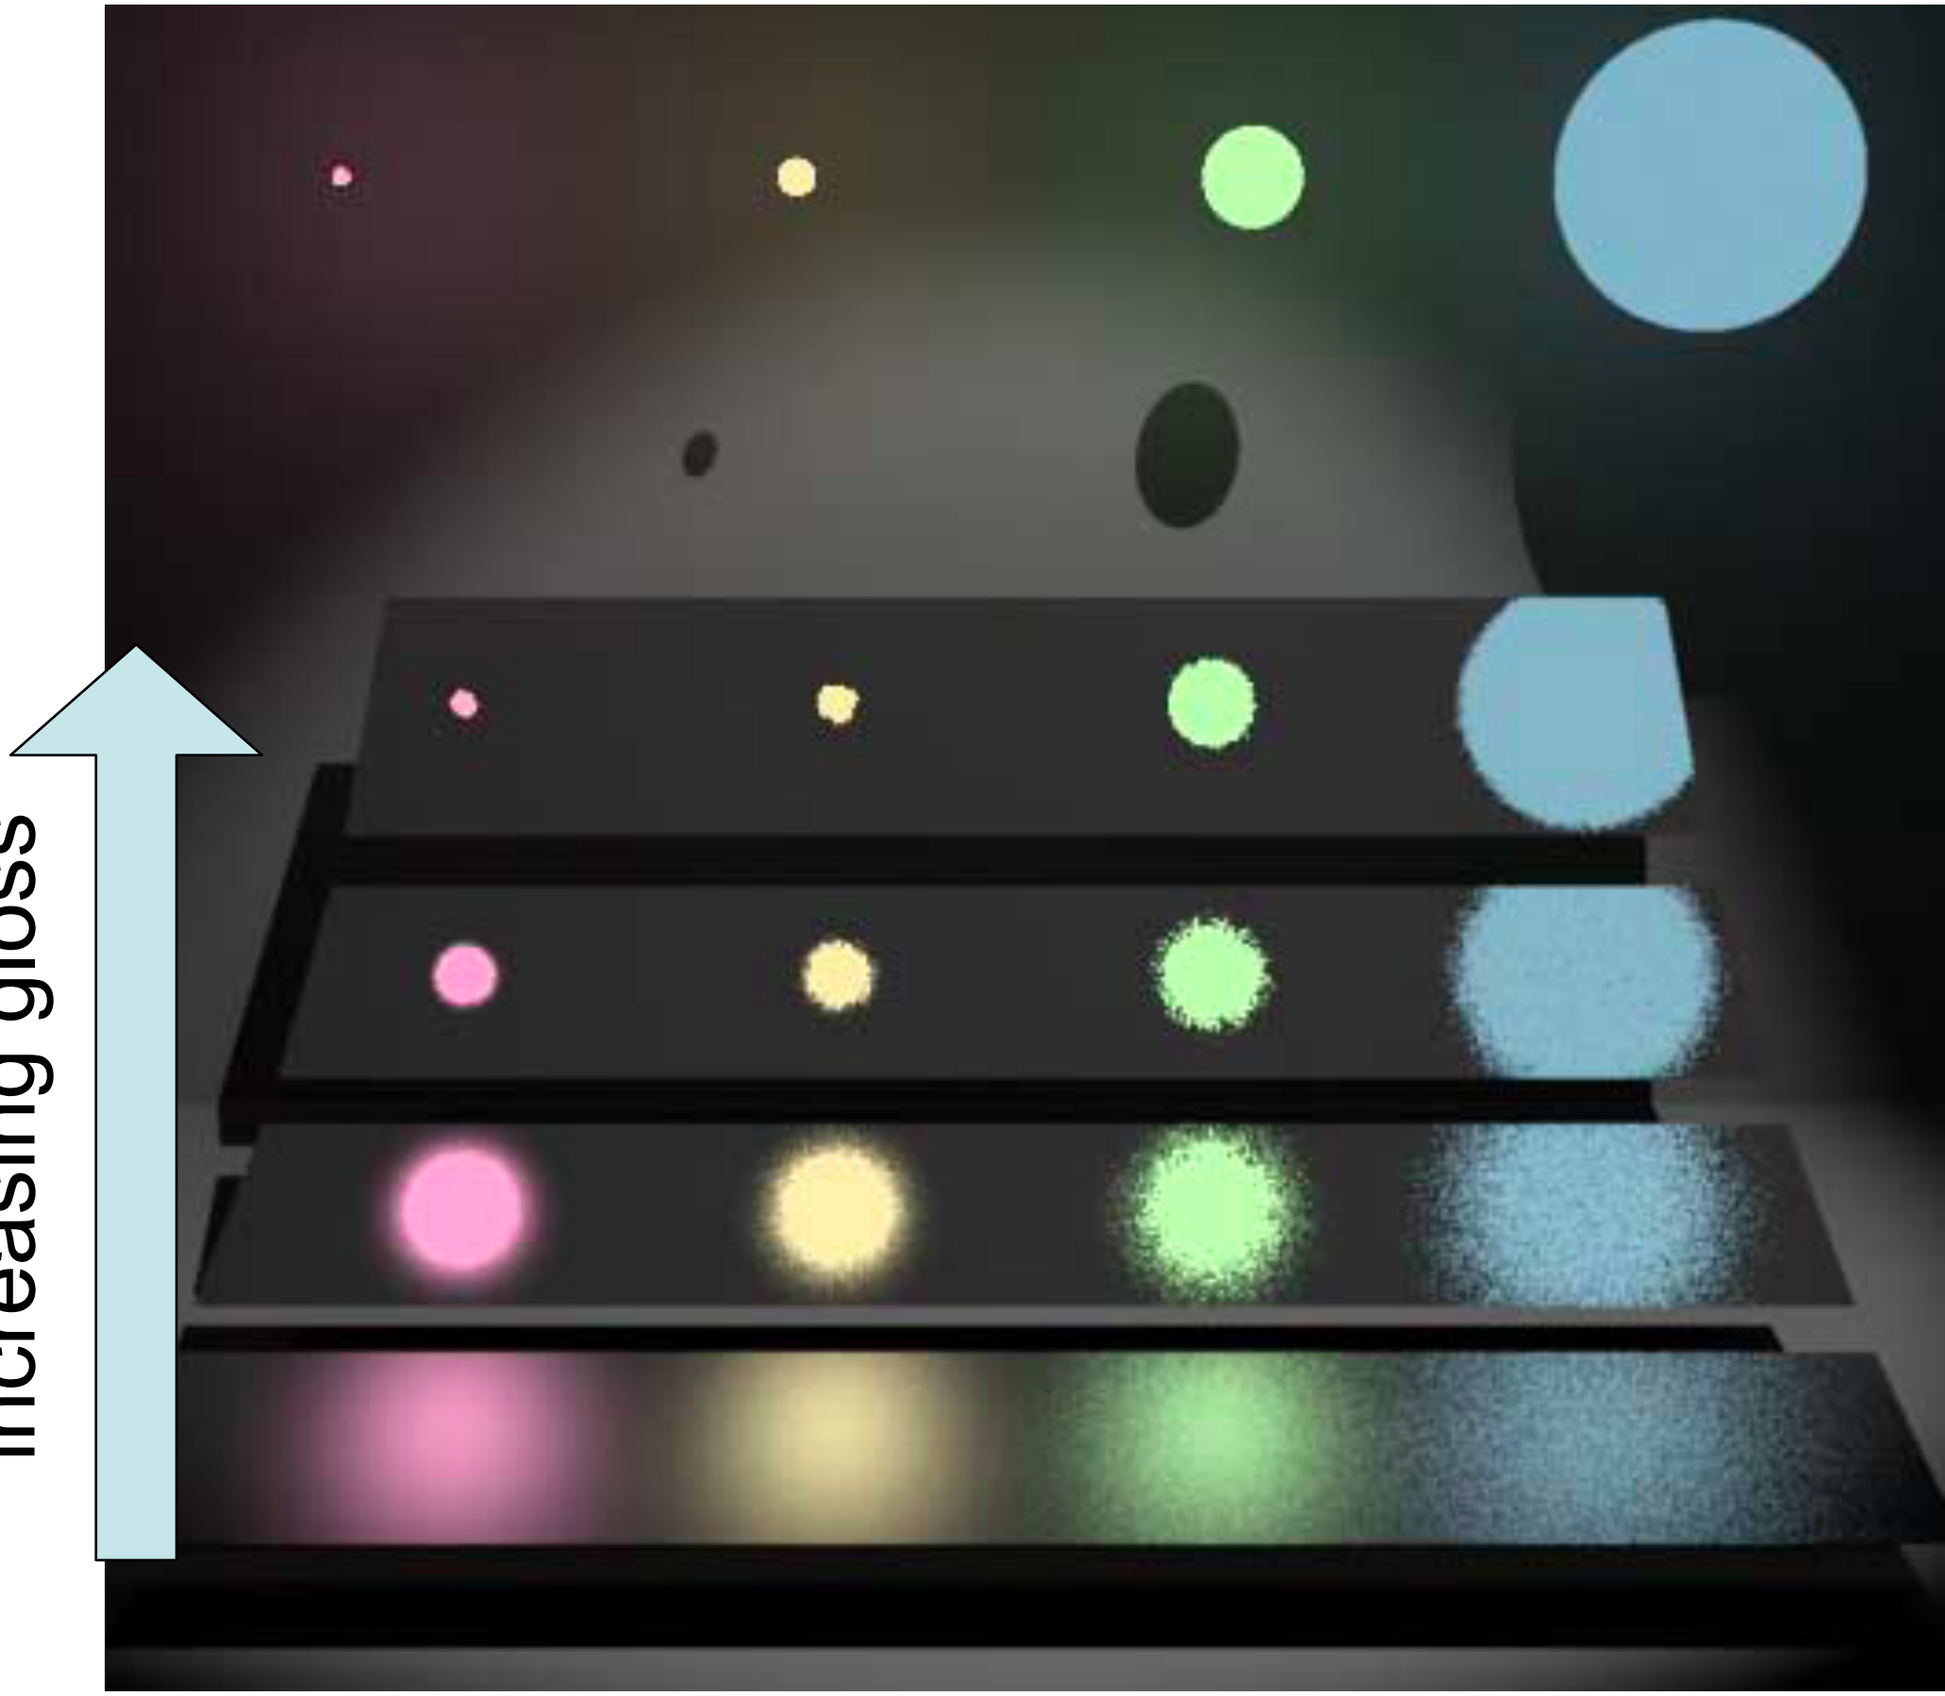

# Ok, how do you do it?

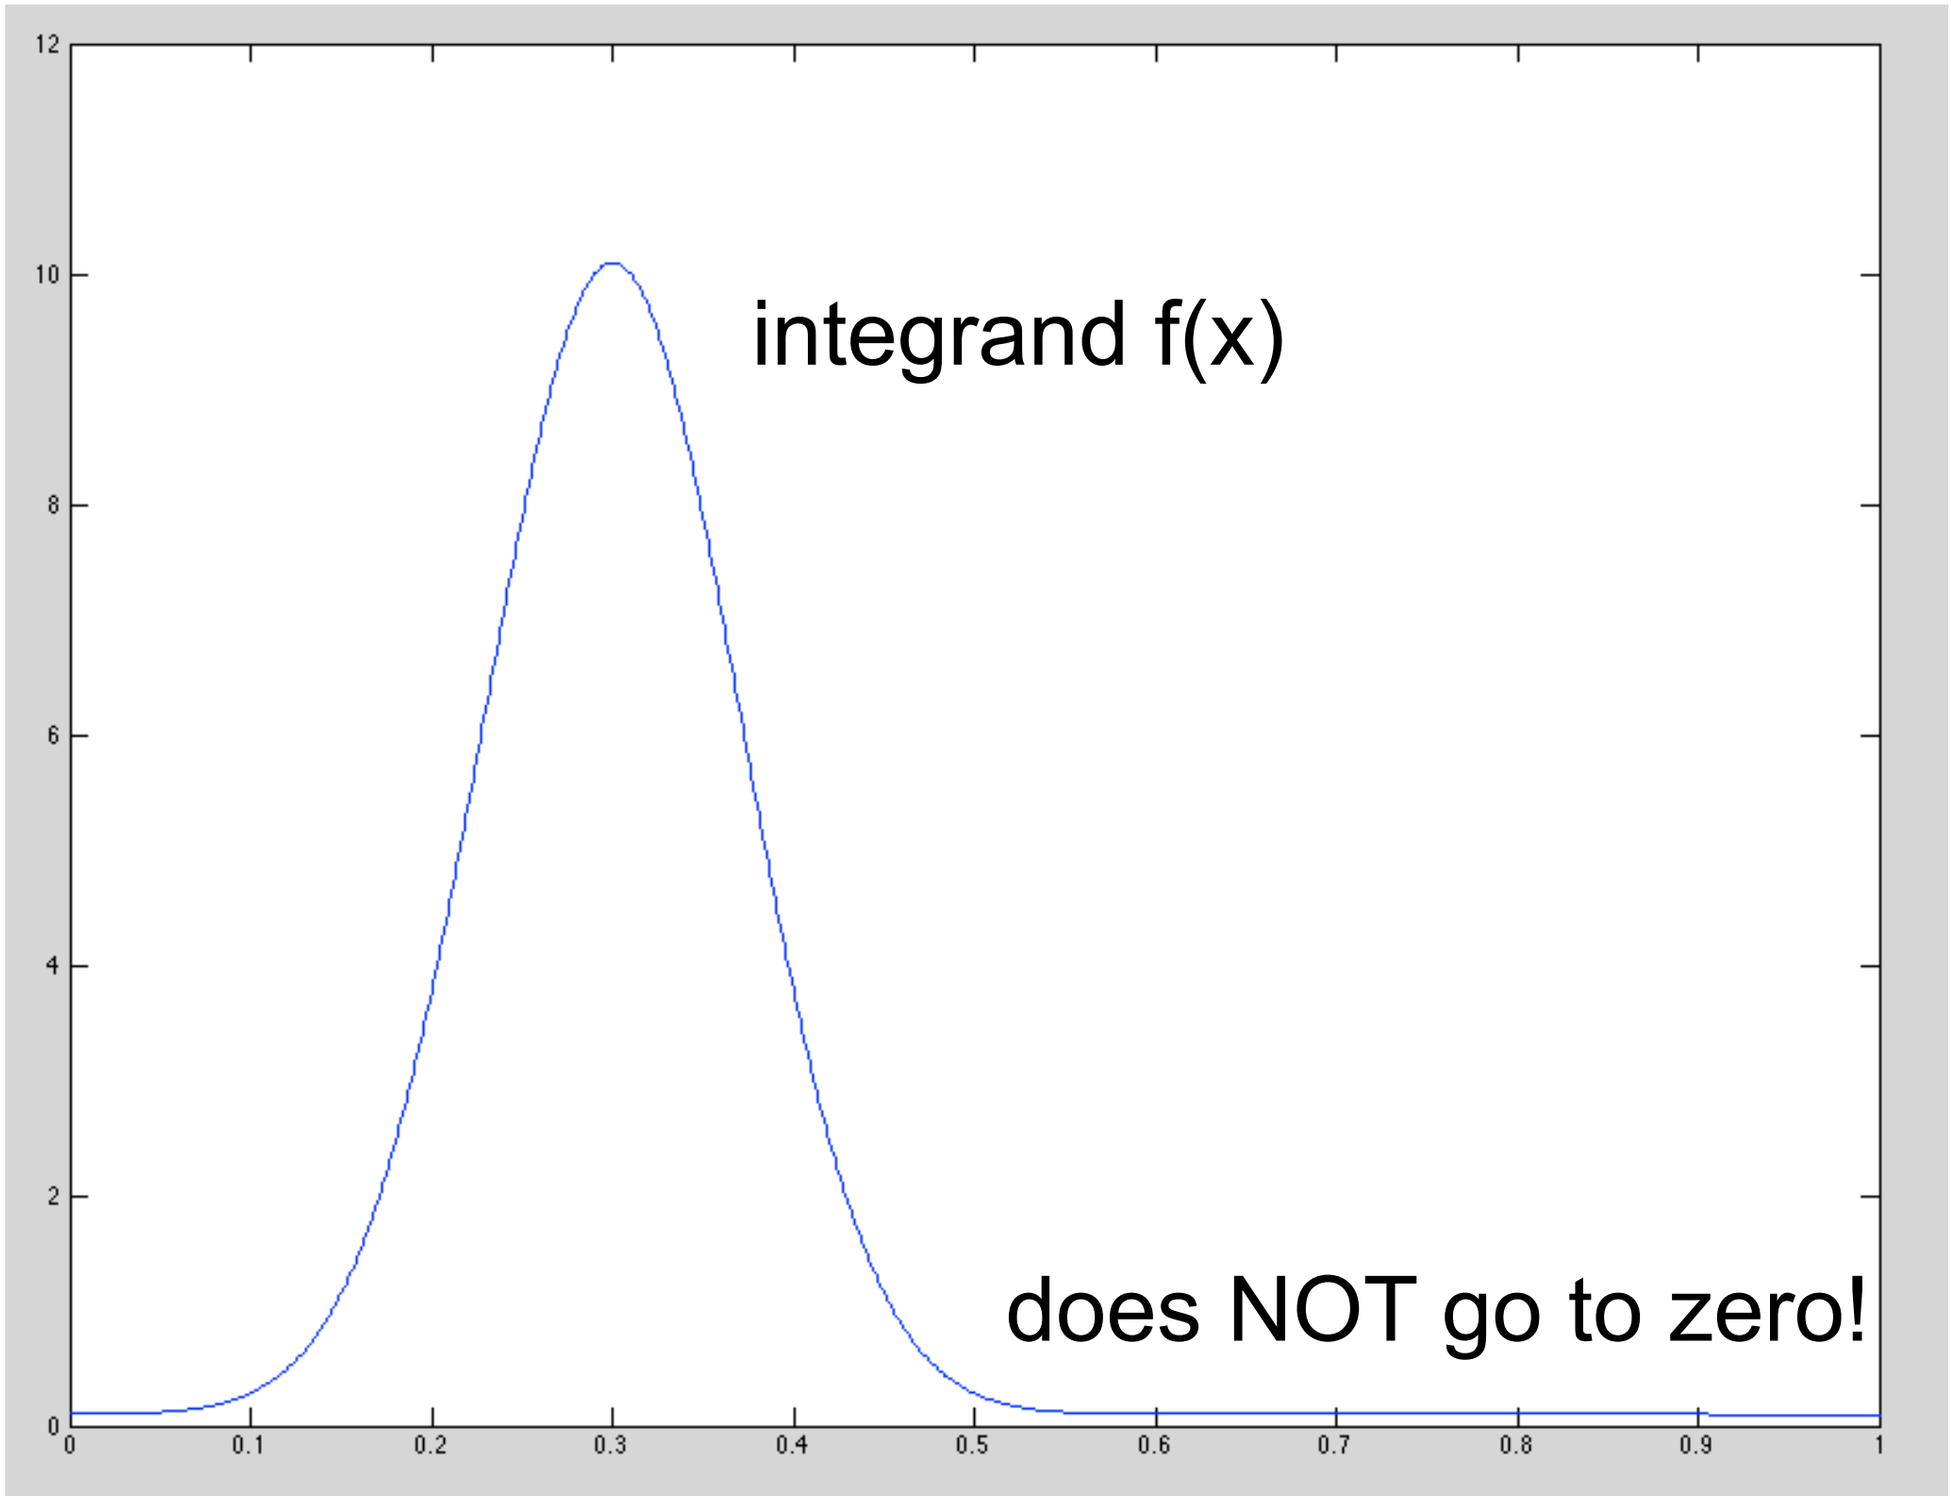

# Why is the Red Gaussian bad for IS?

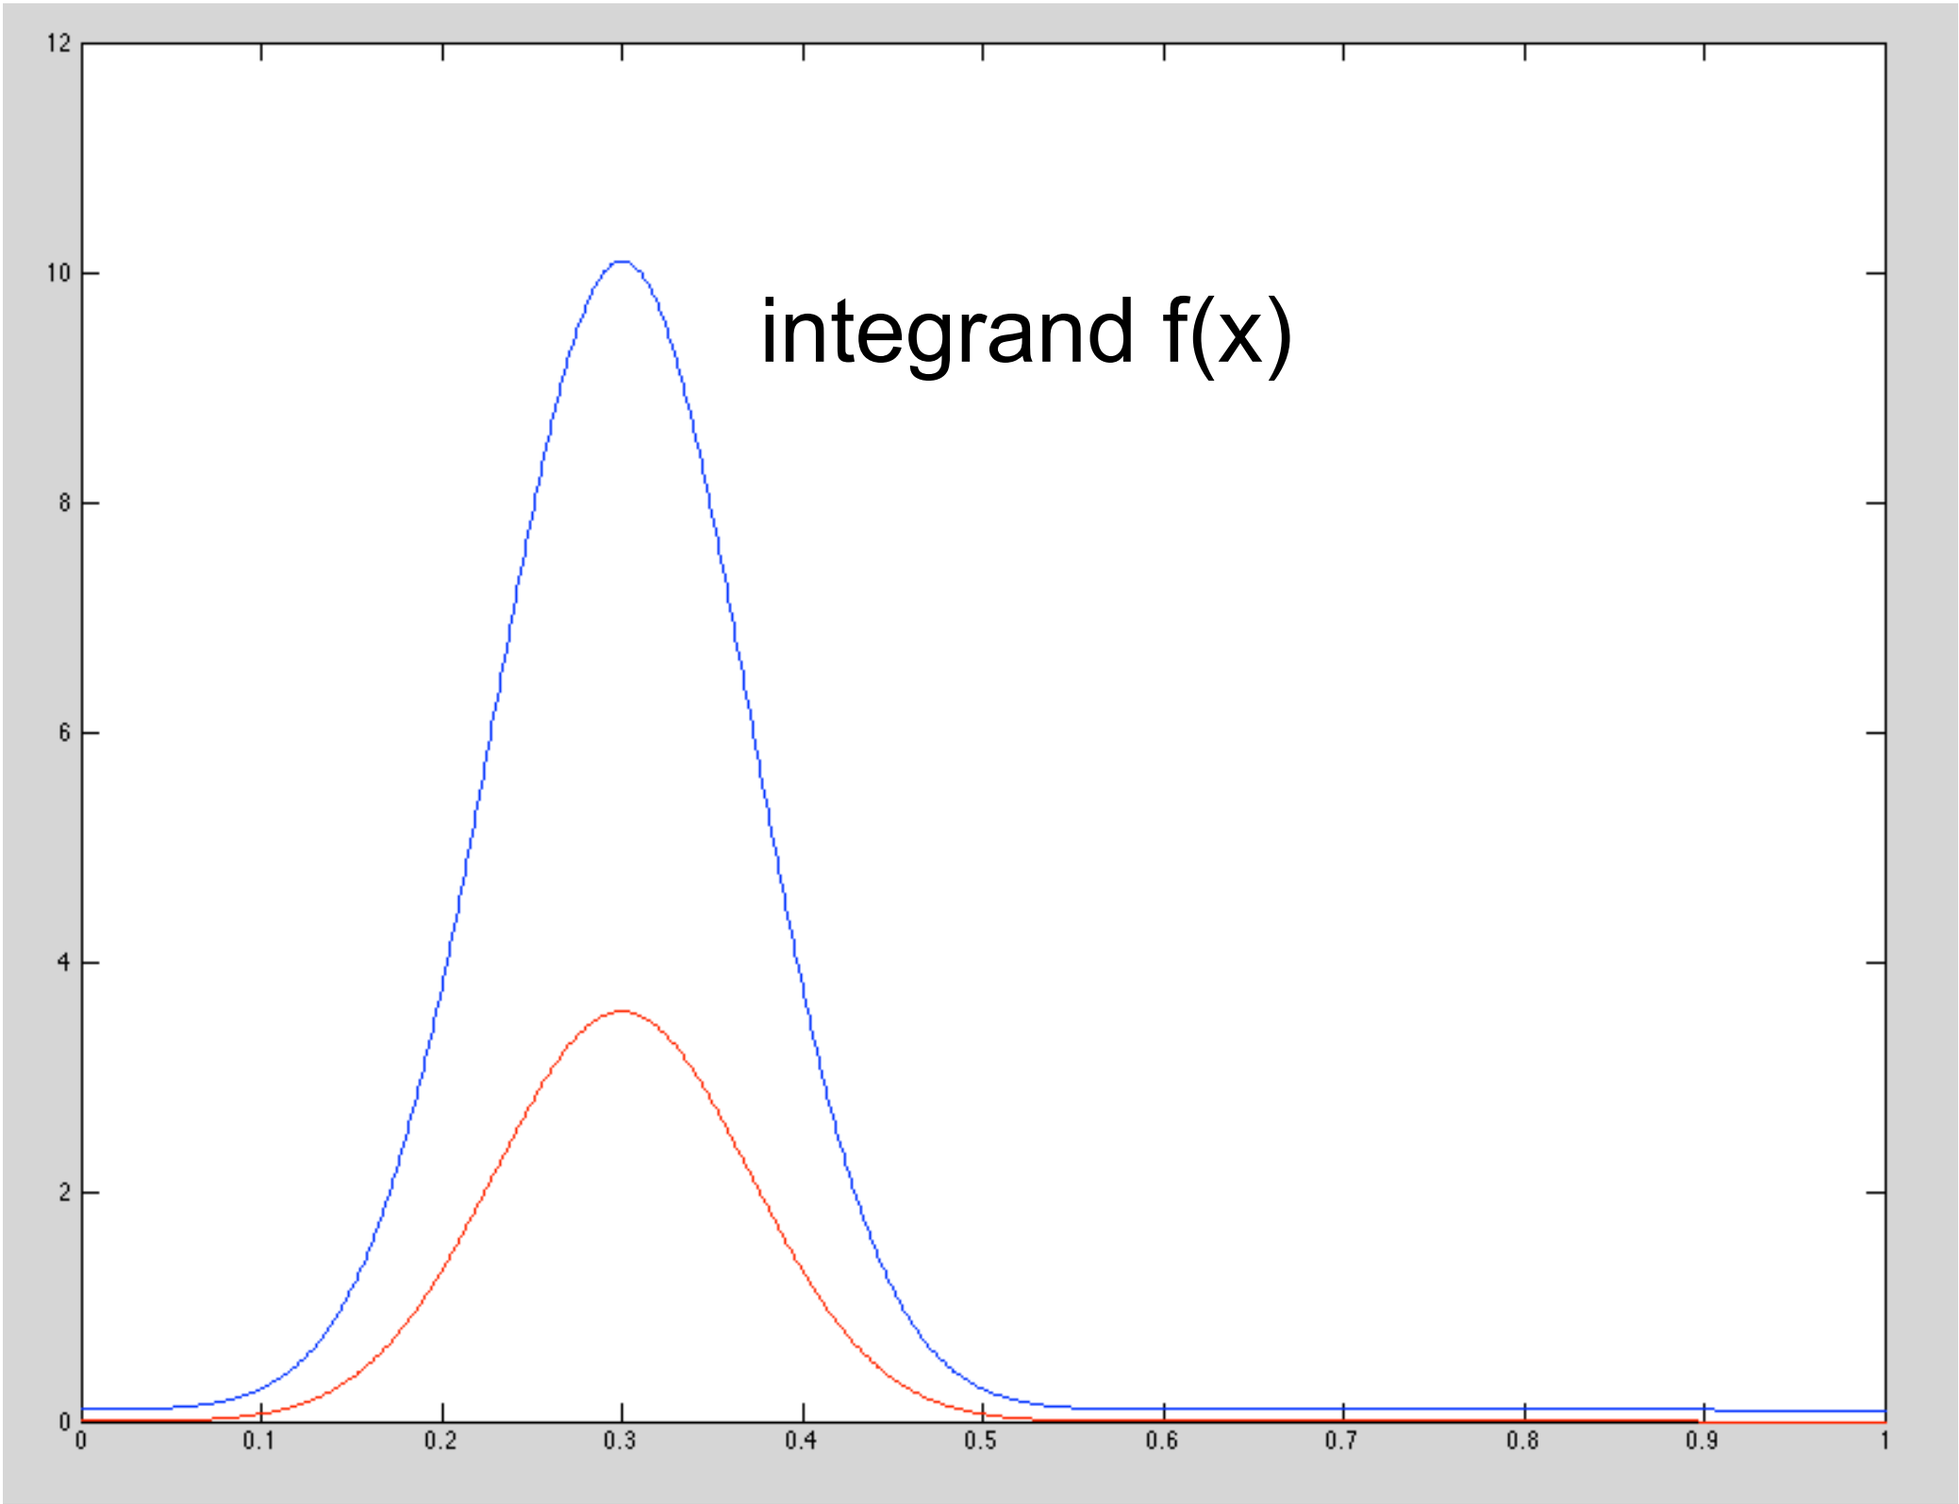

# Why the Red Gaussian is *bad* for IS

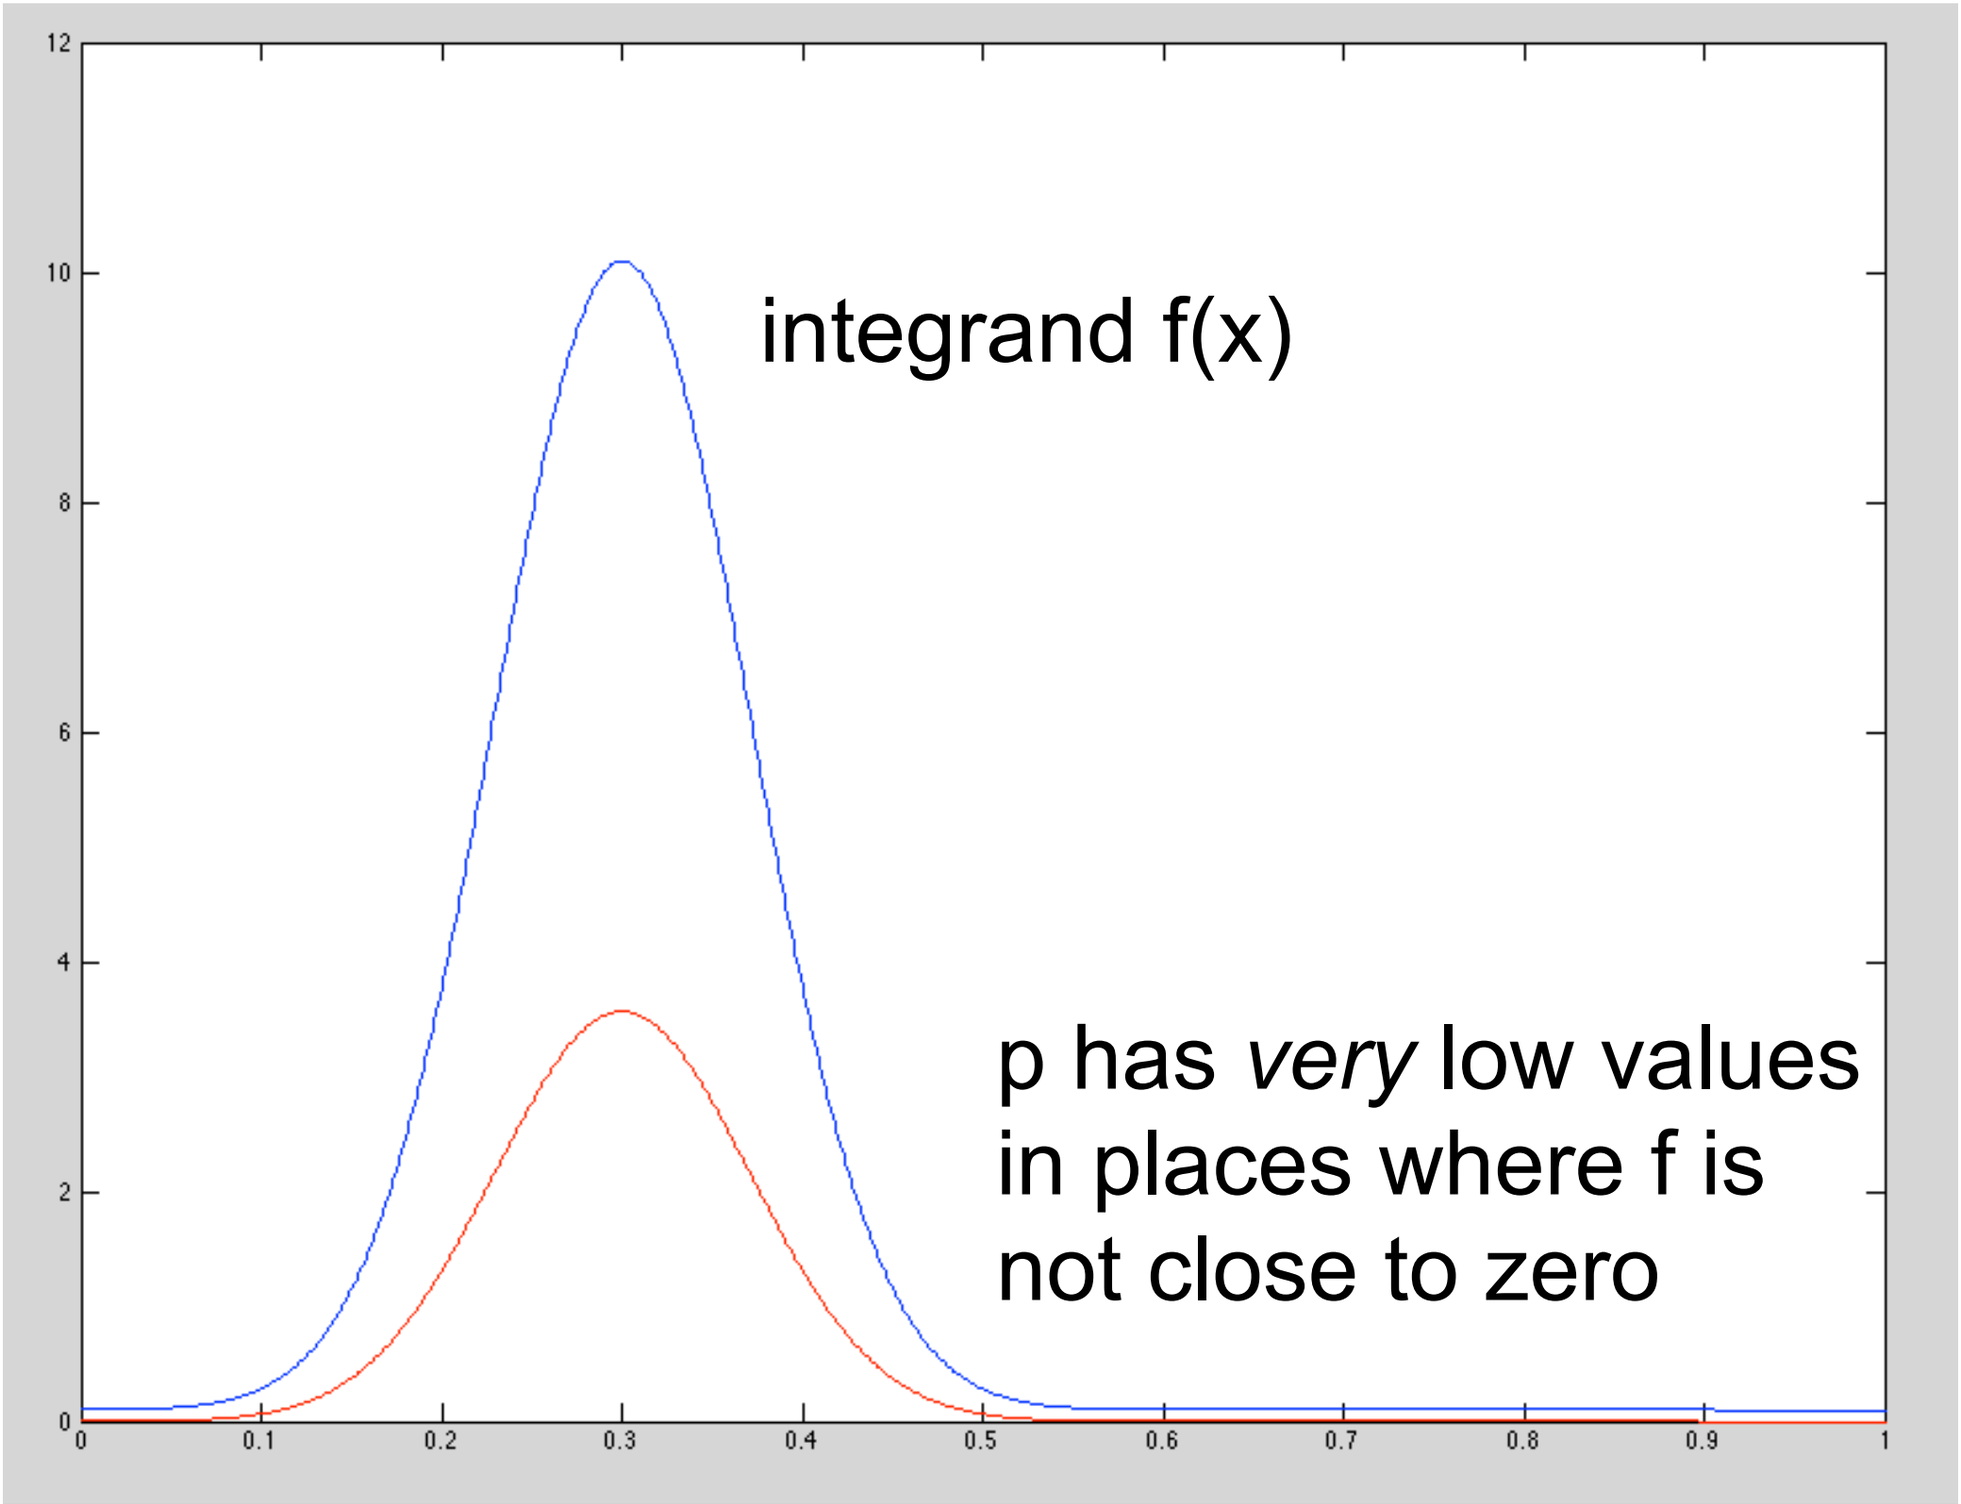

# Why This Matters

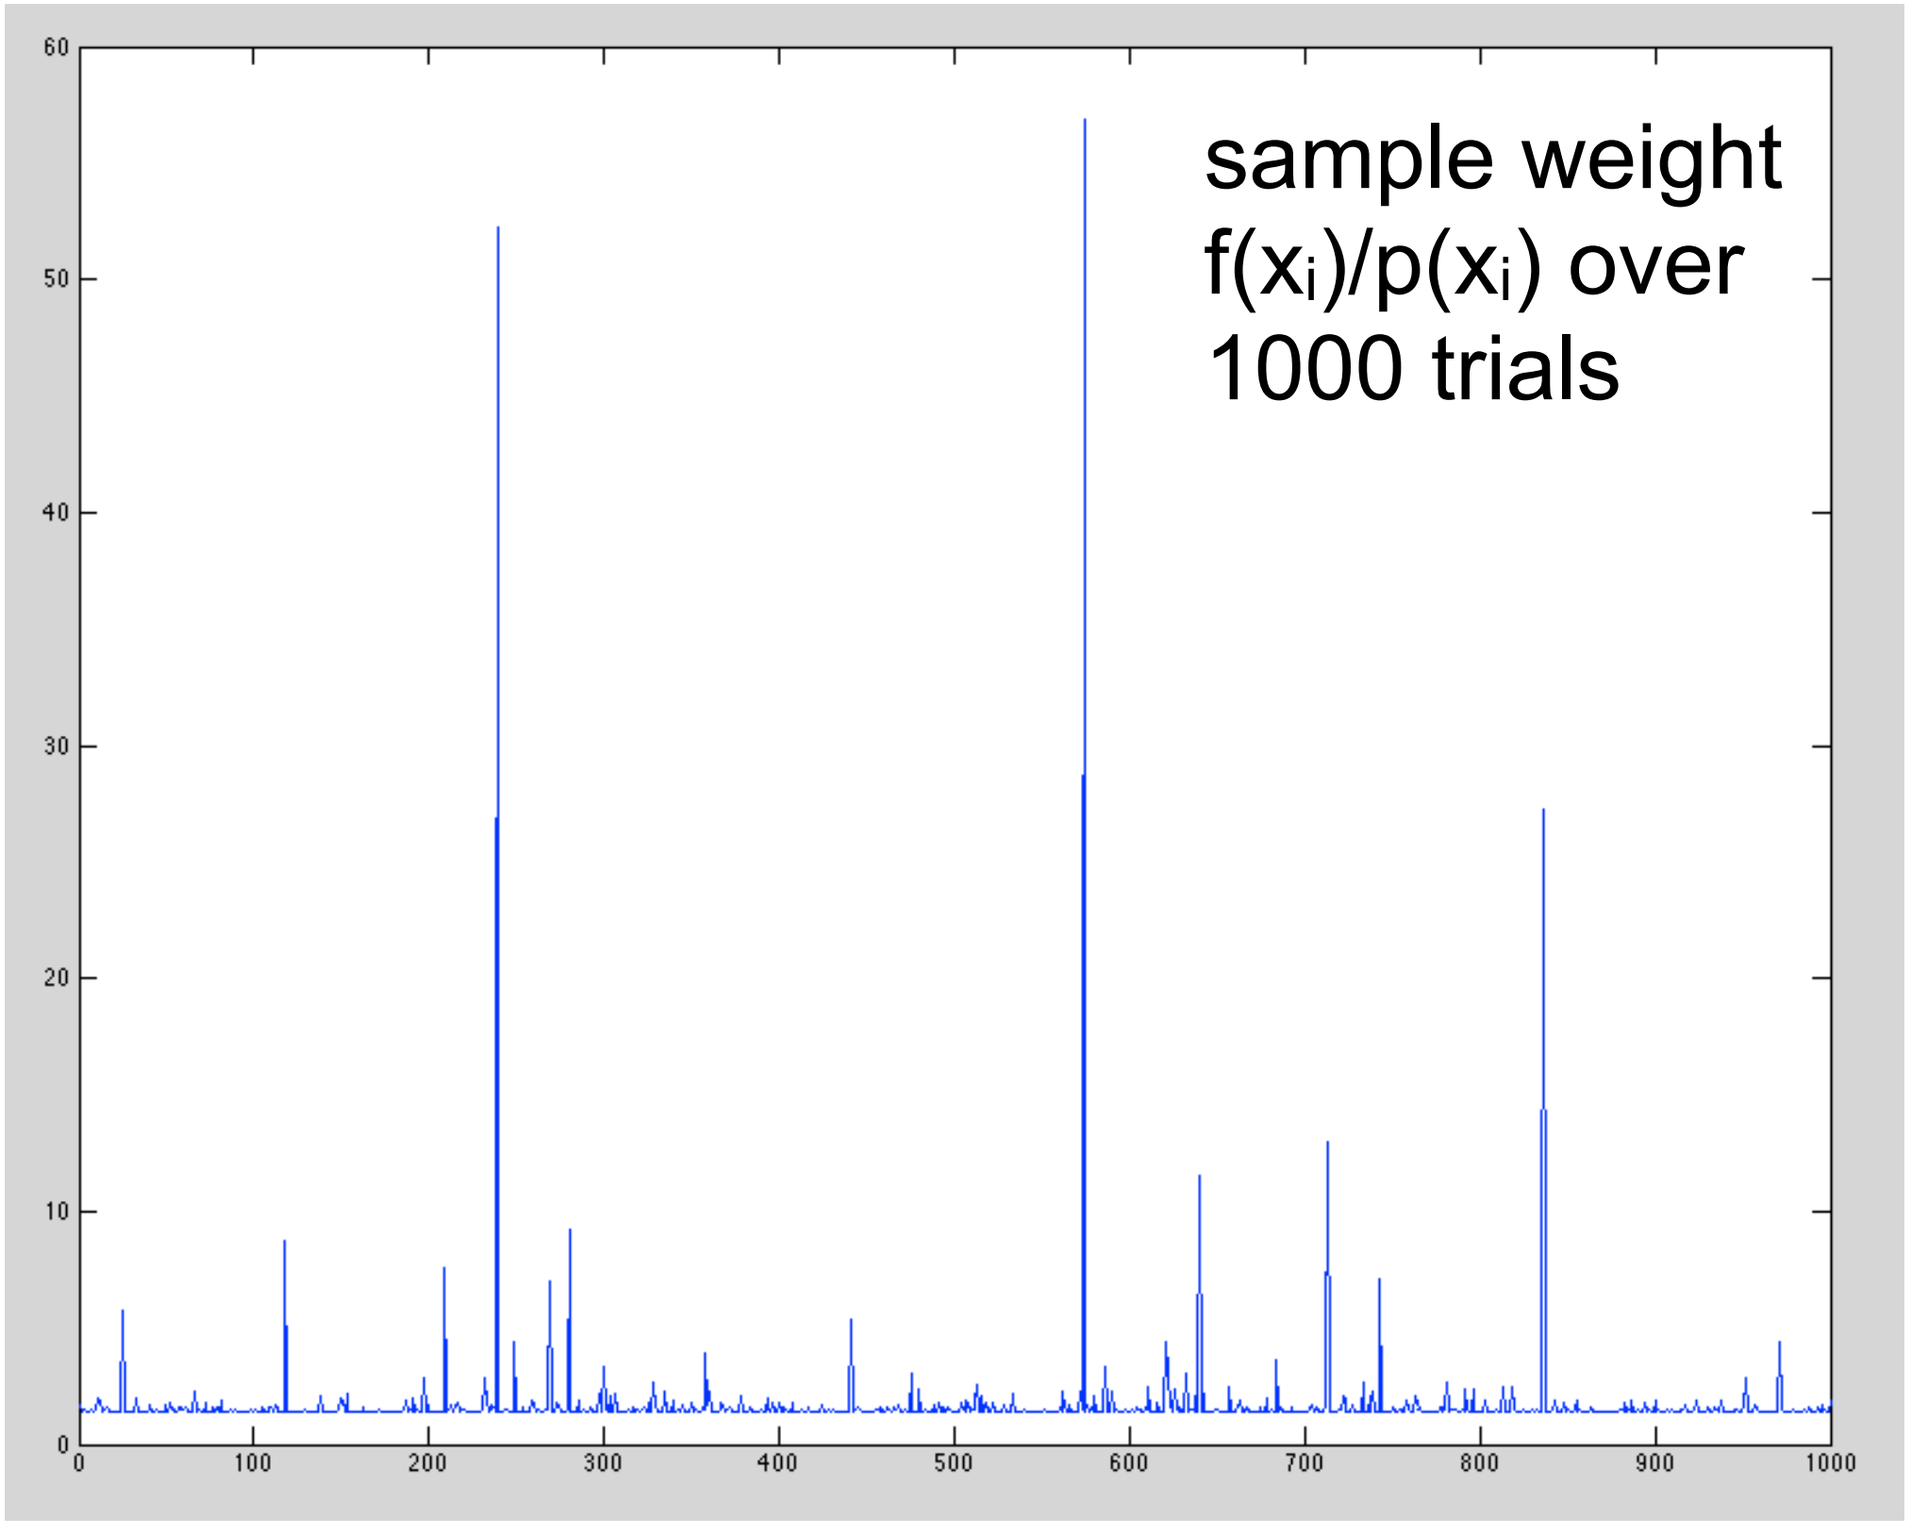

# Why This Matters

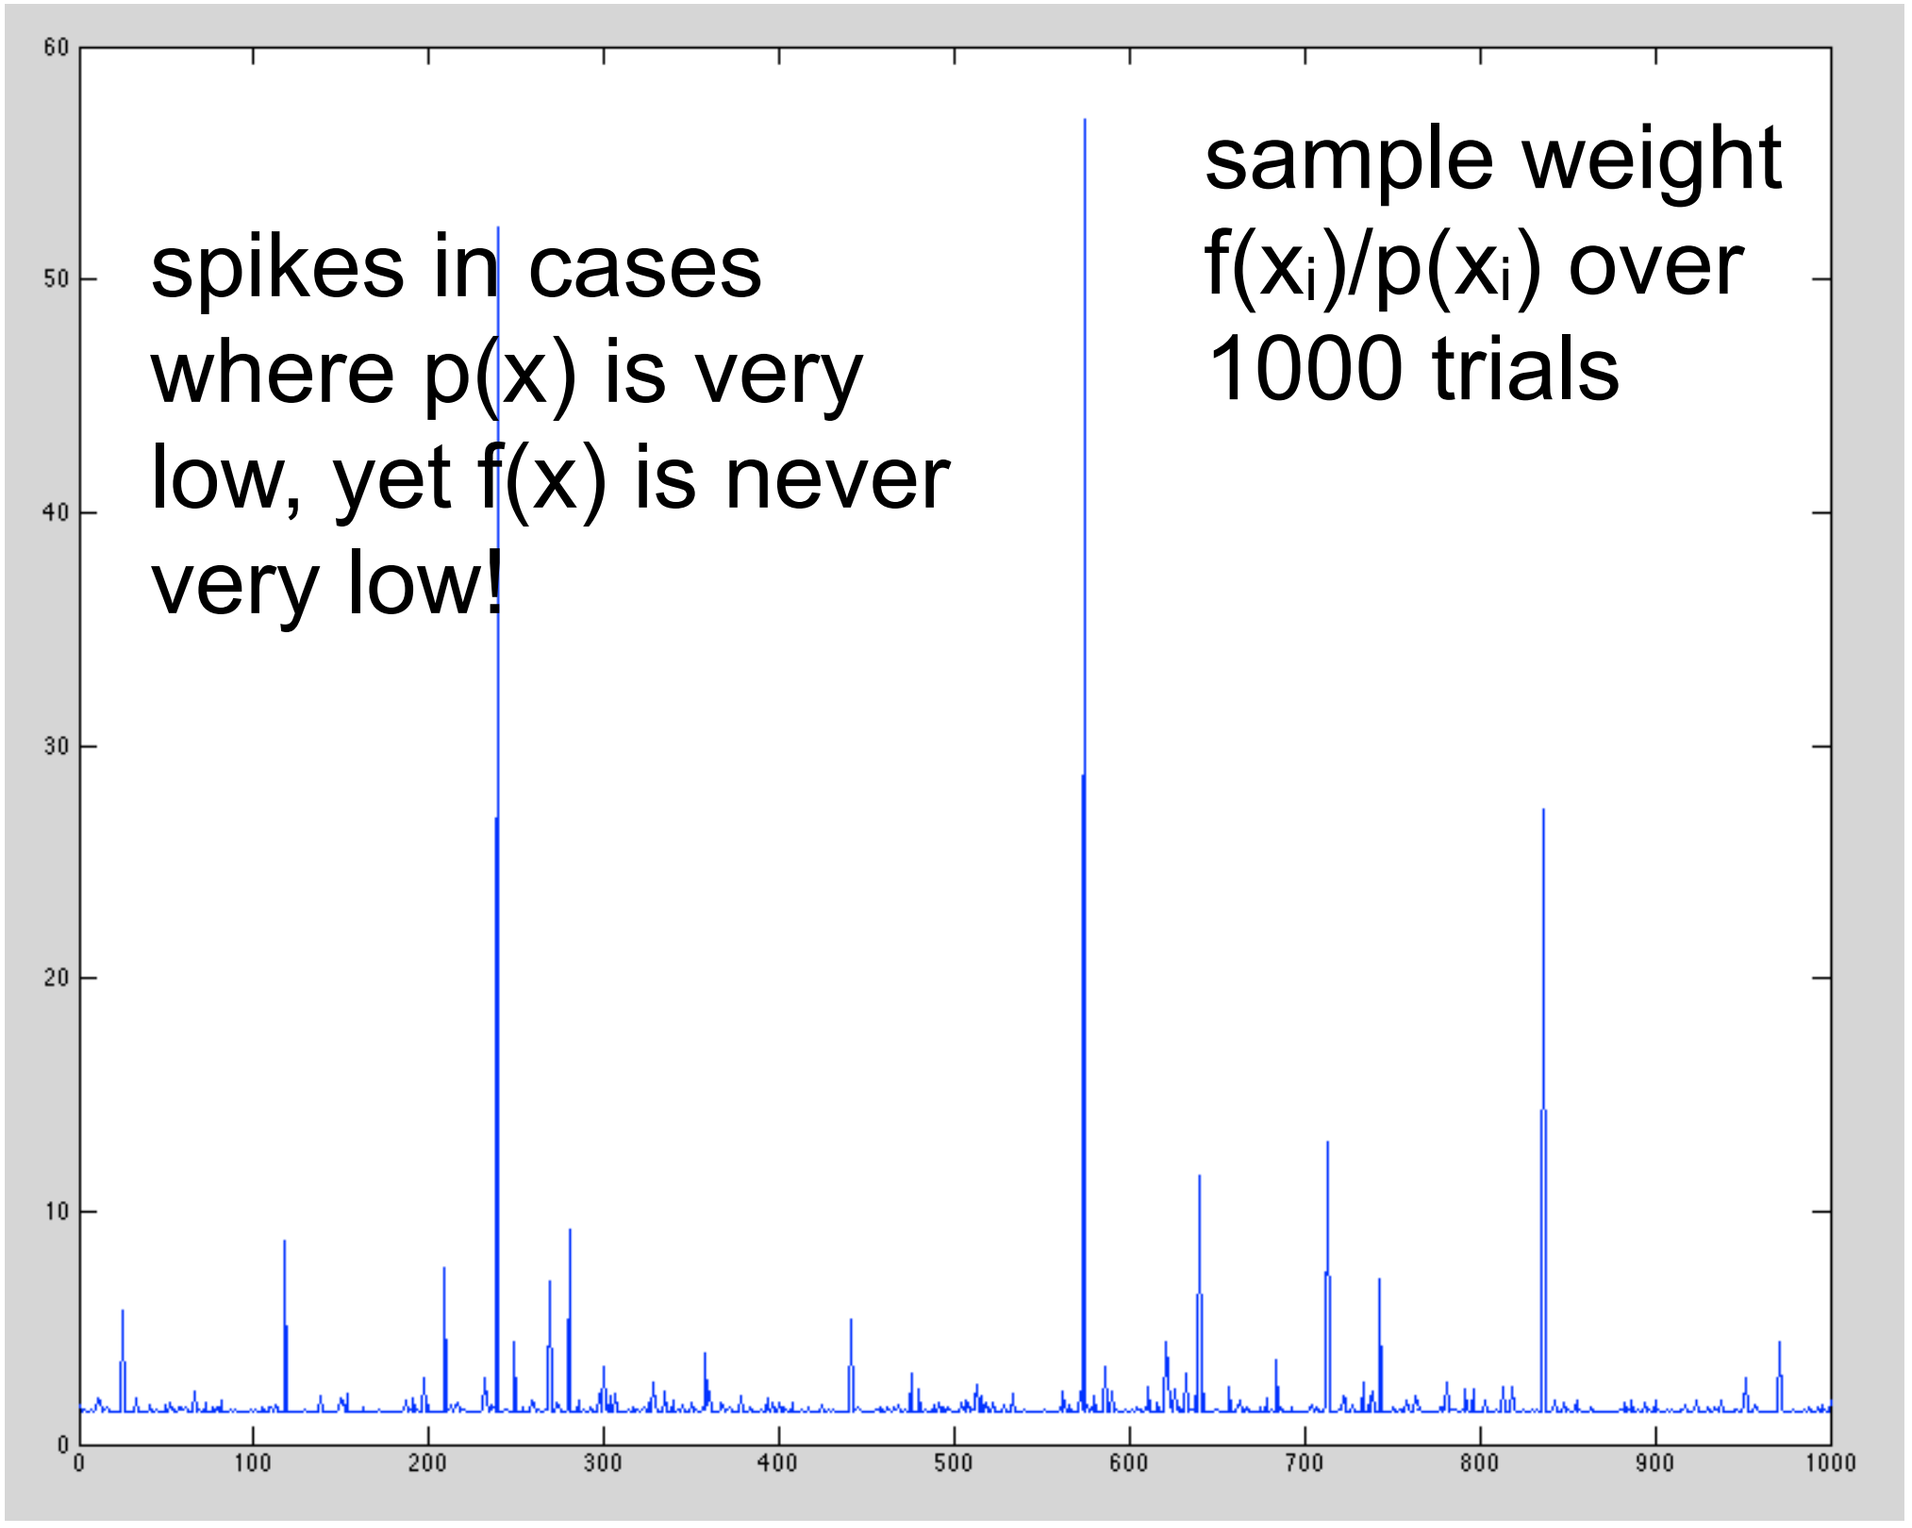

# Why This Matters

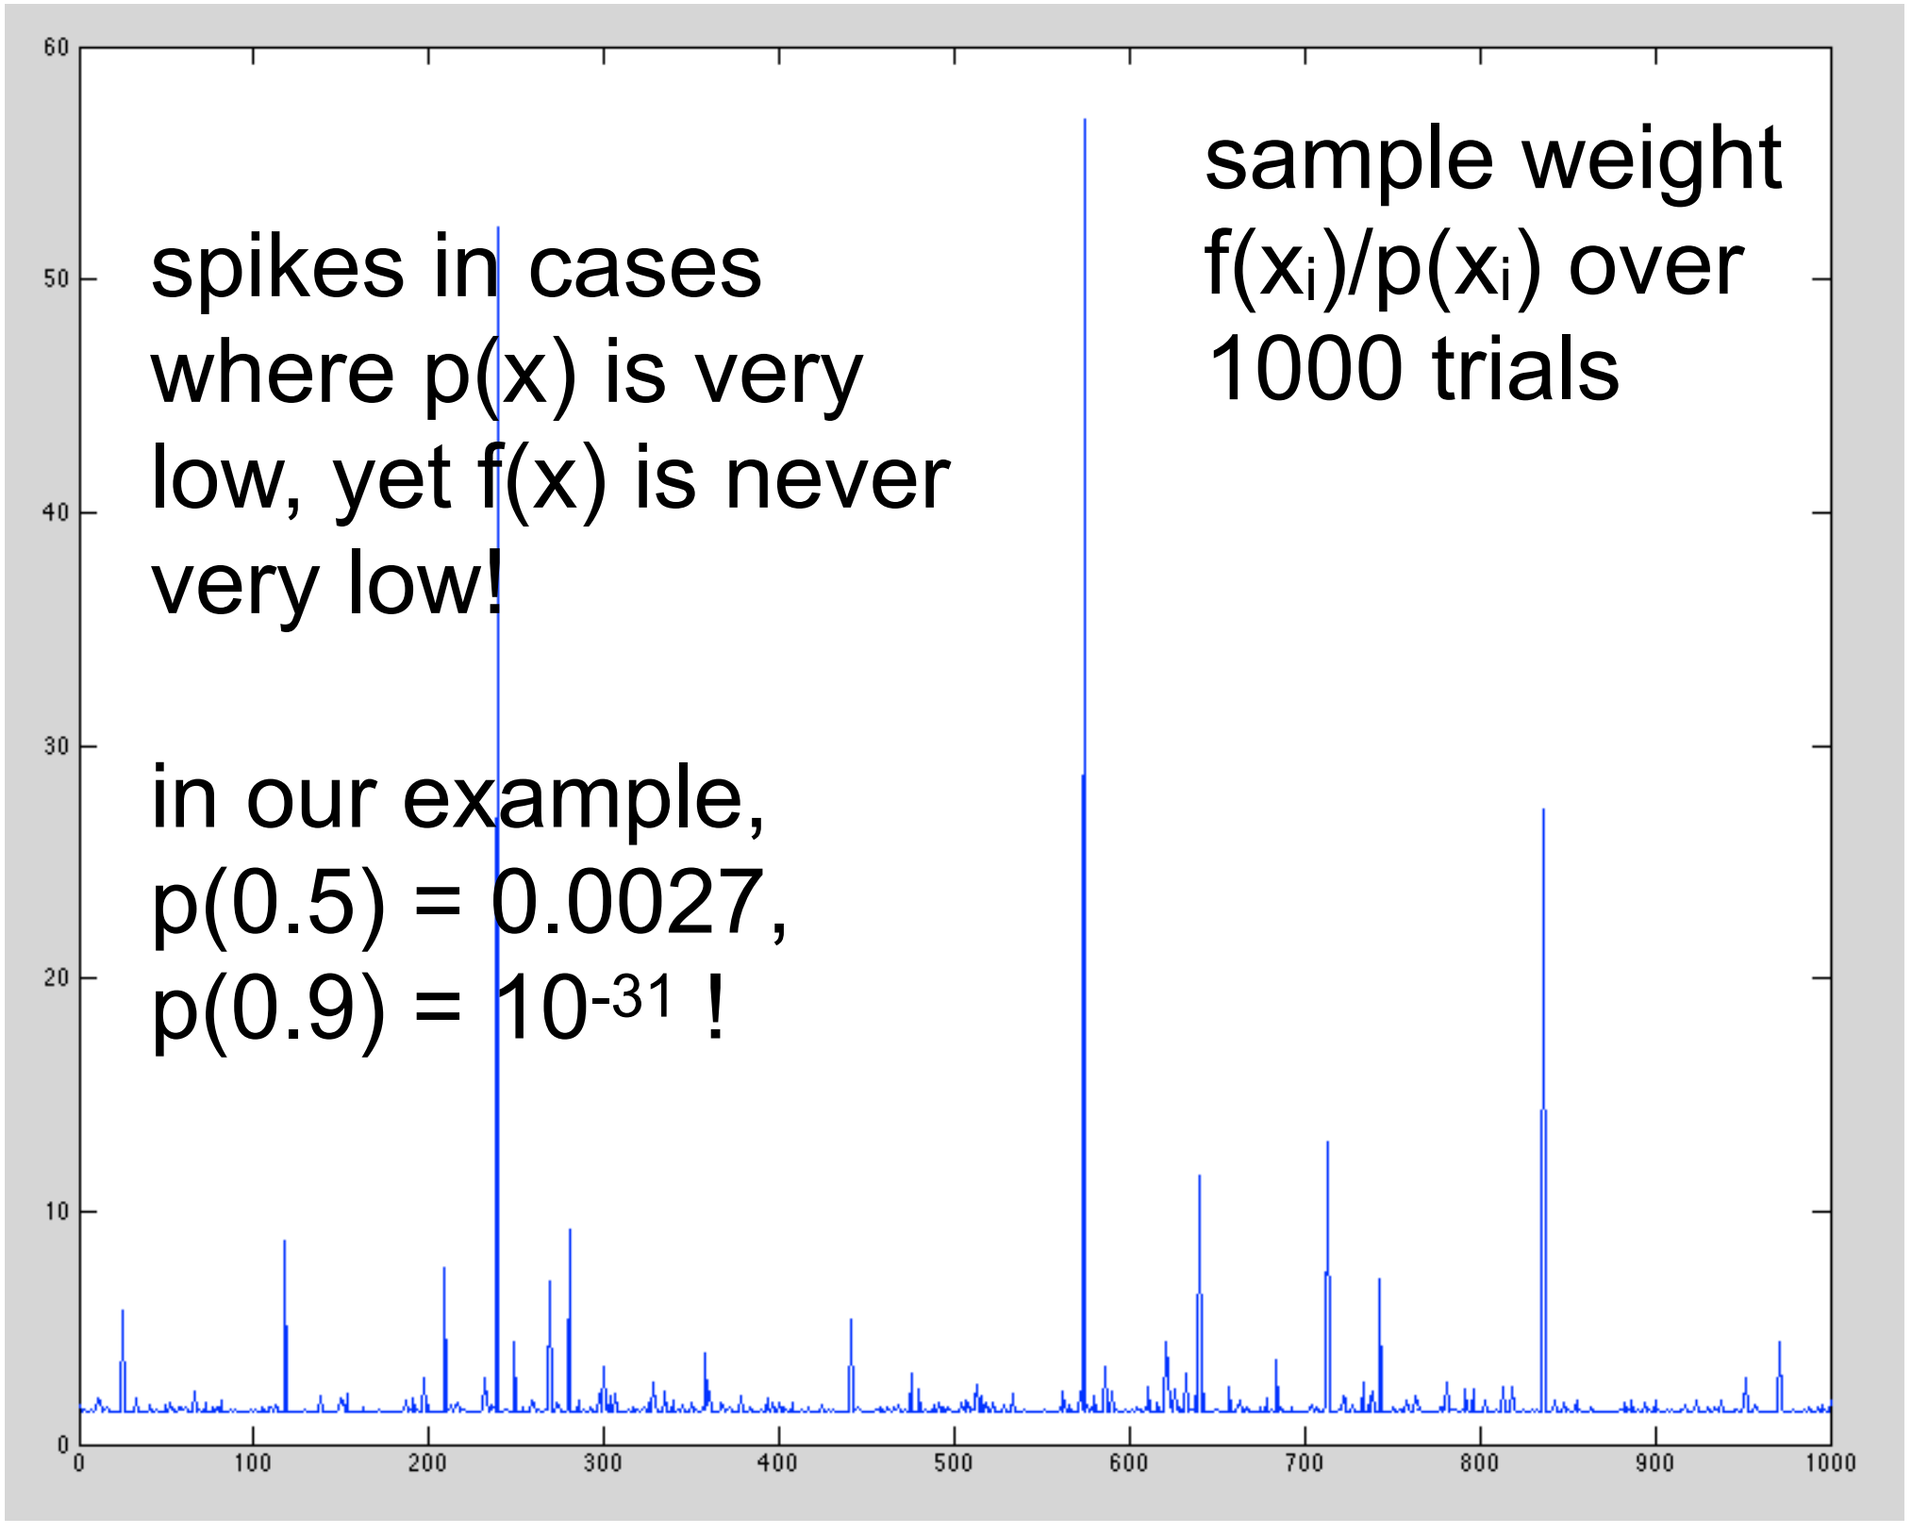

# Spikes get worse with higher N

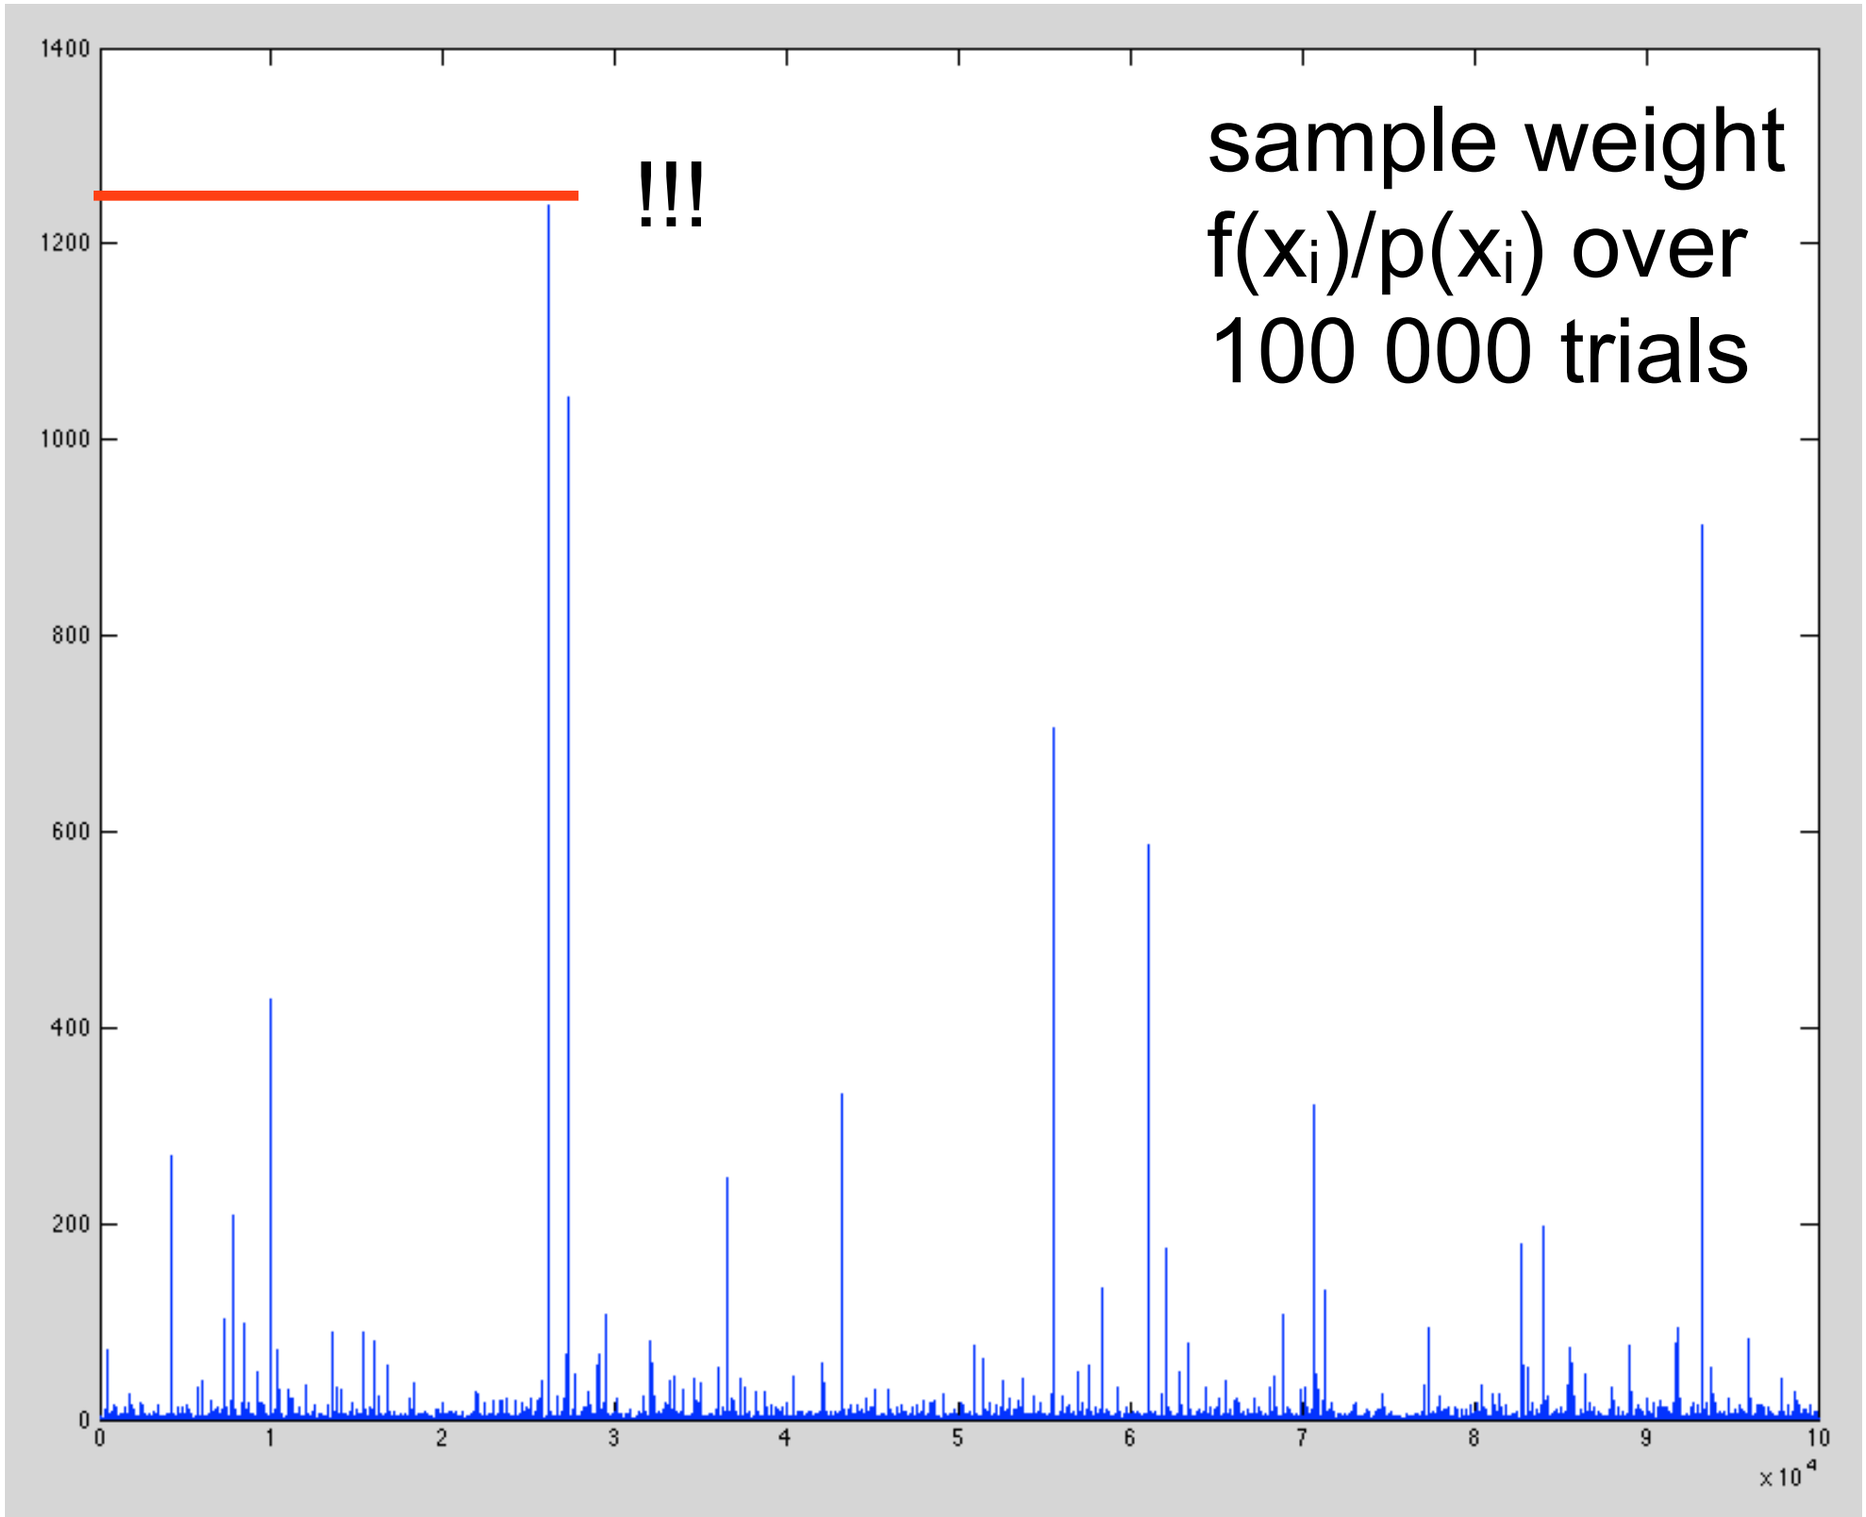

# Effect of Spikes on Integral Estimate

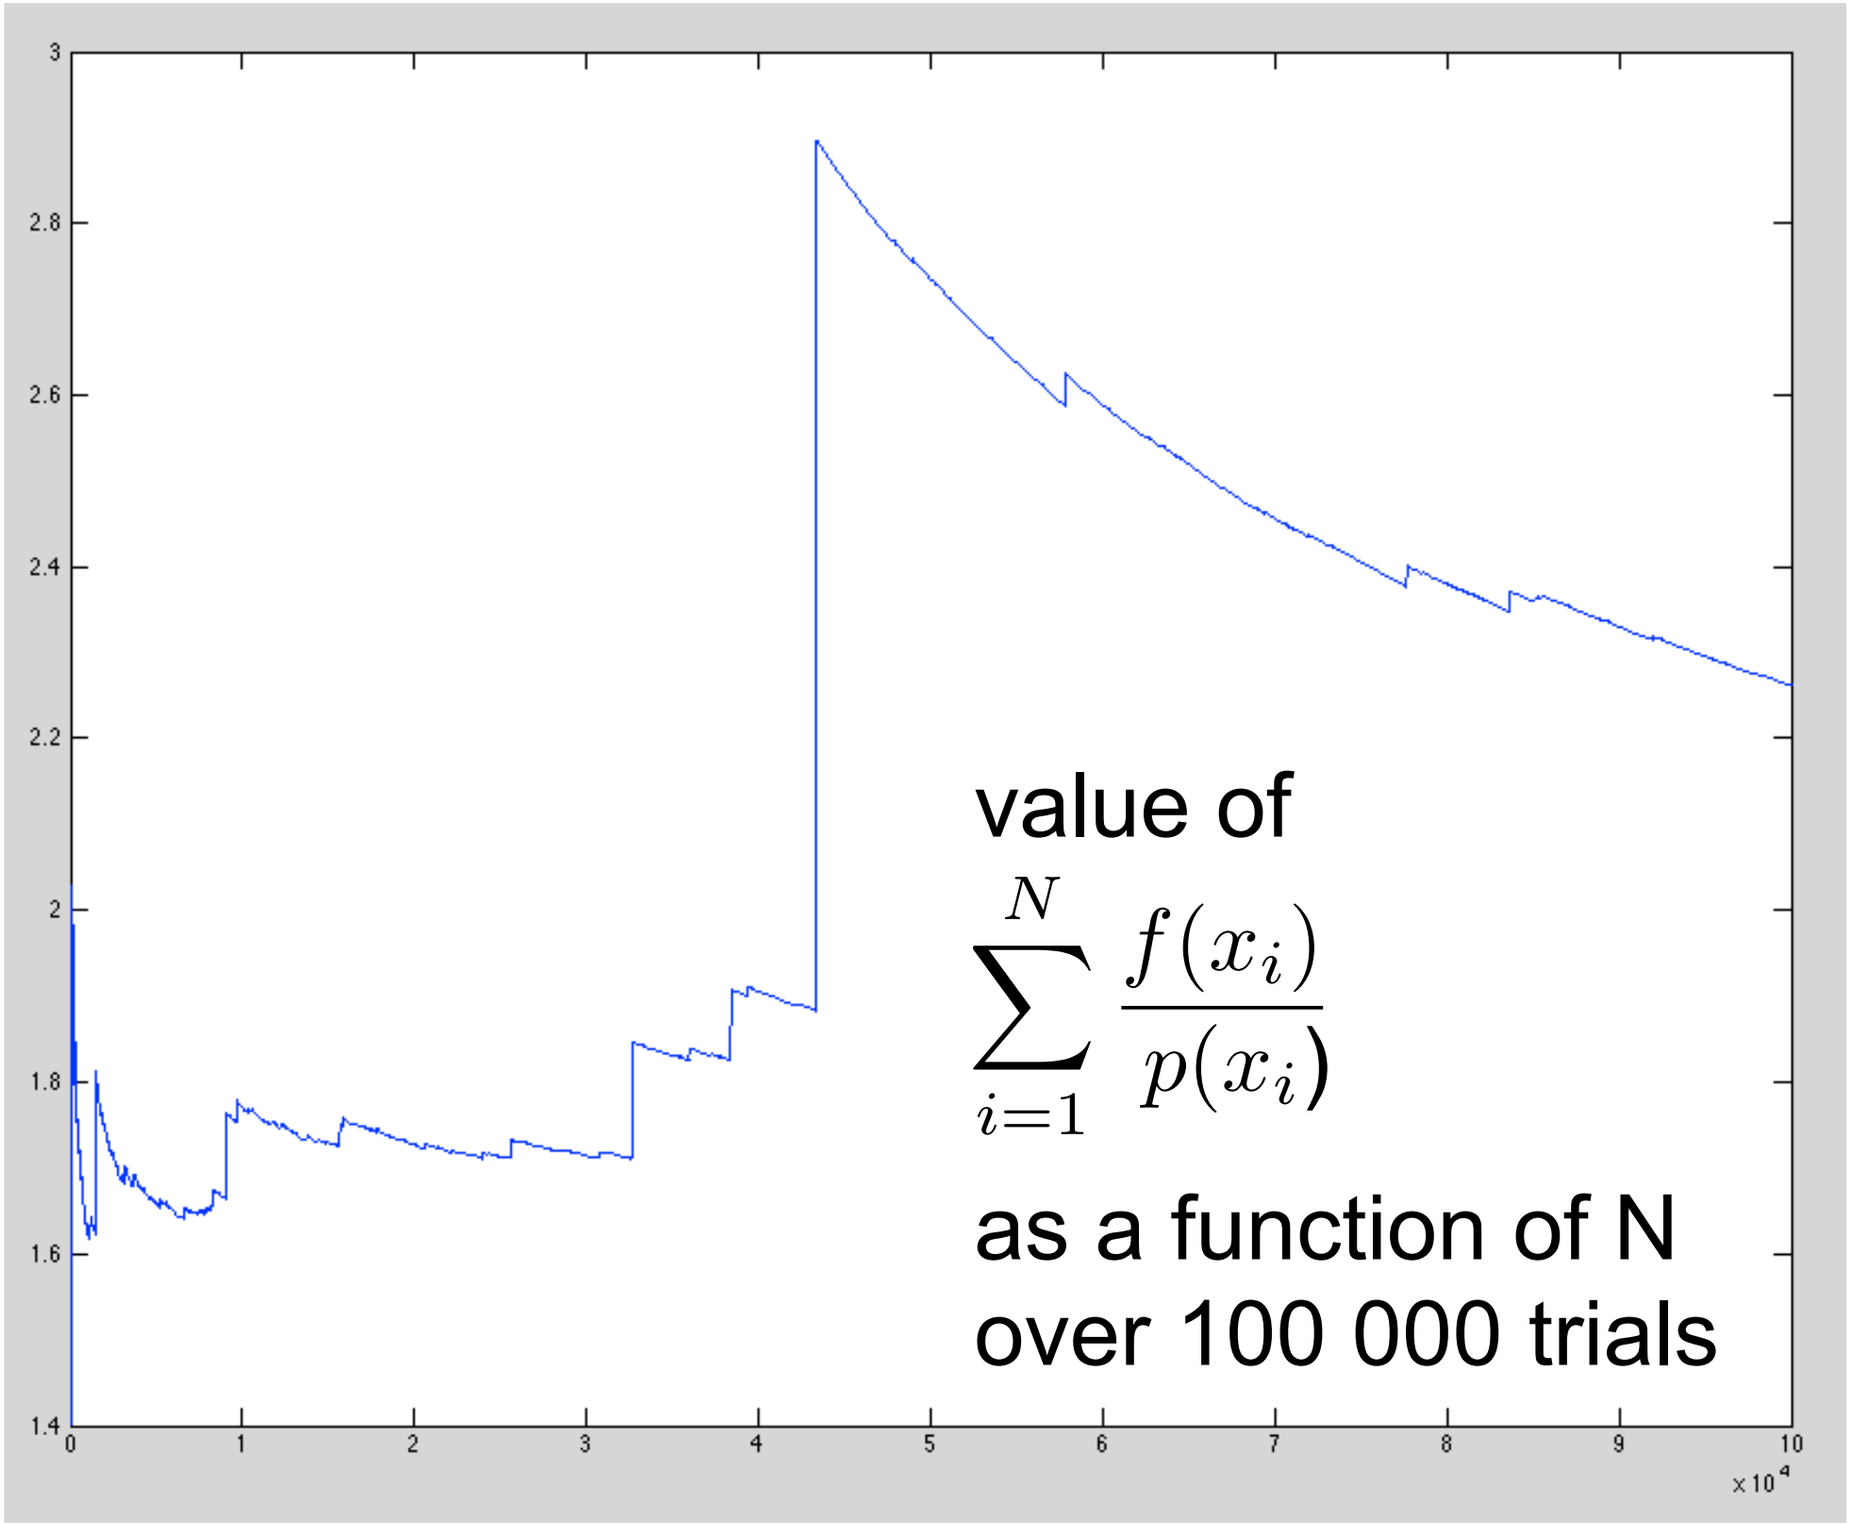

# Effect of Spikes on Integral Estimate

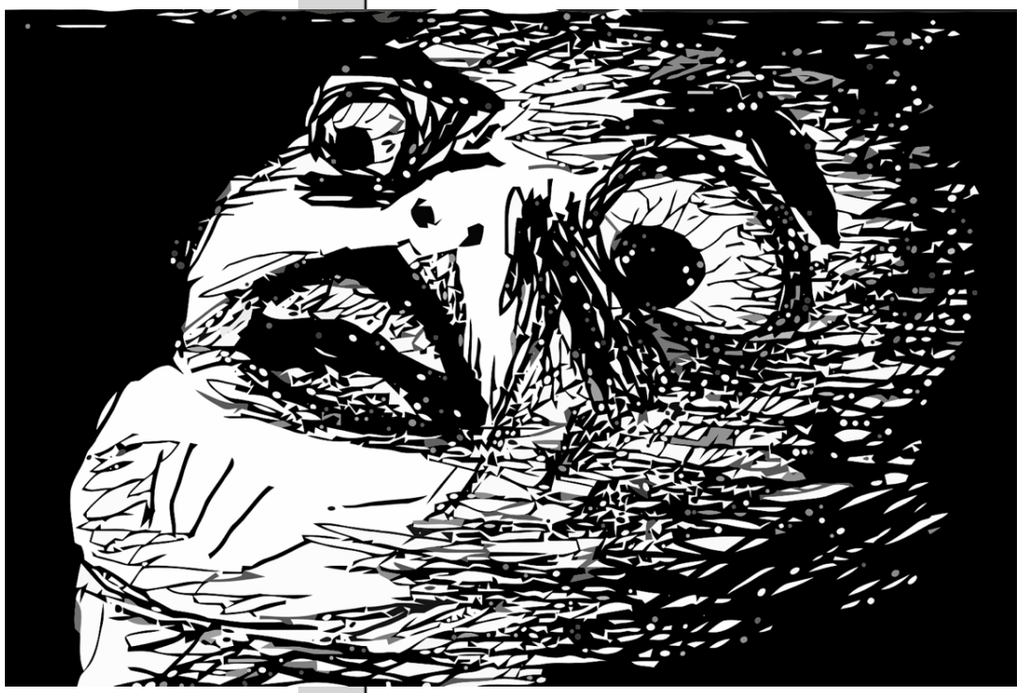




value of

$$\sum_{i=1}^{N} \frac{f(x_i)}{p(x_i}$$

as a function of N
over 100 000 trials

# Graph of f/p (note log scale in y!)

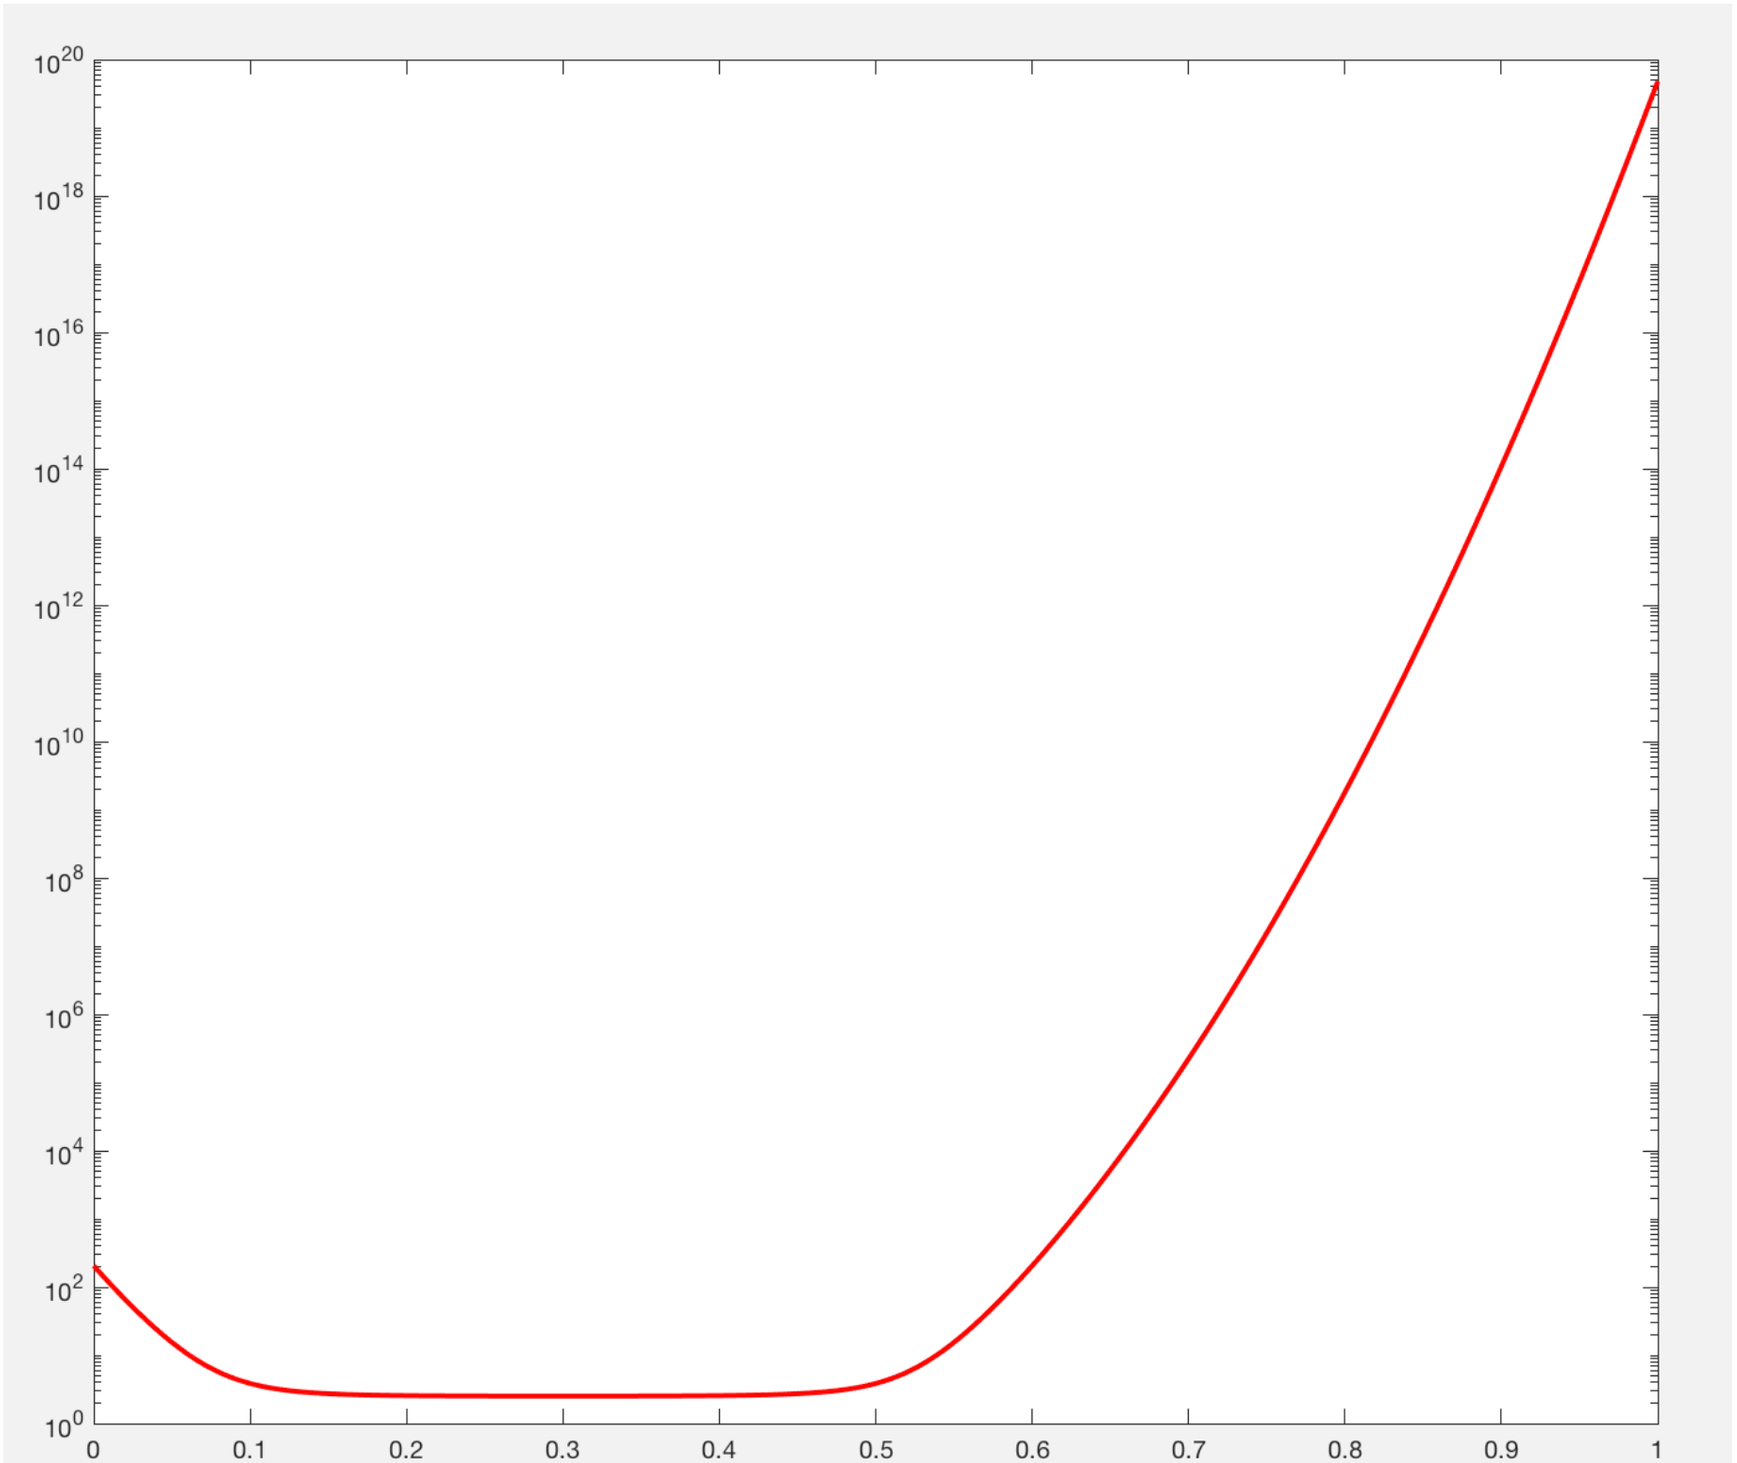

# Better: Let's mix in a constant PDF

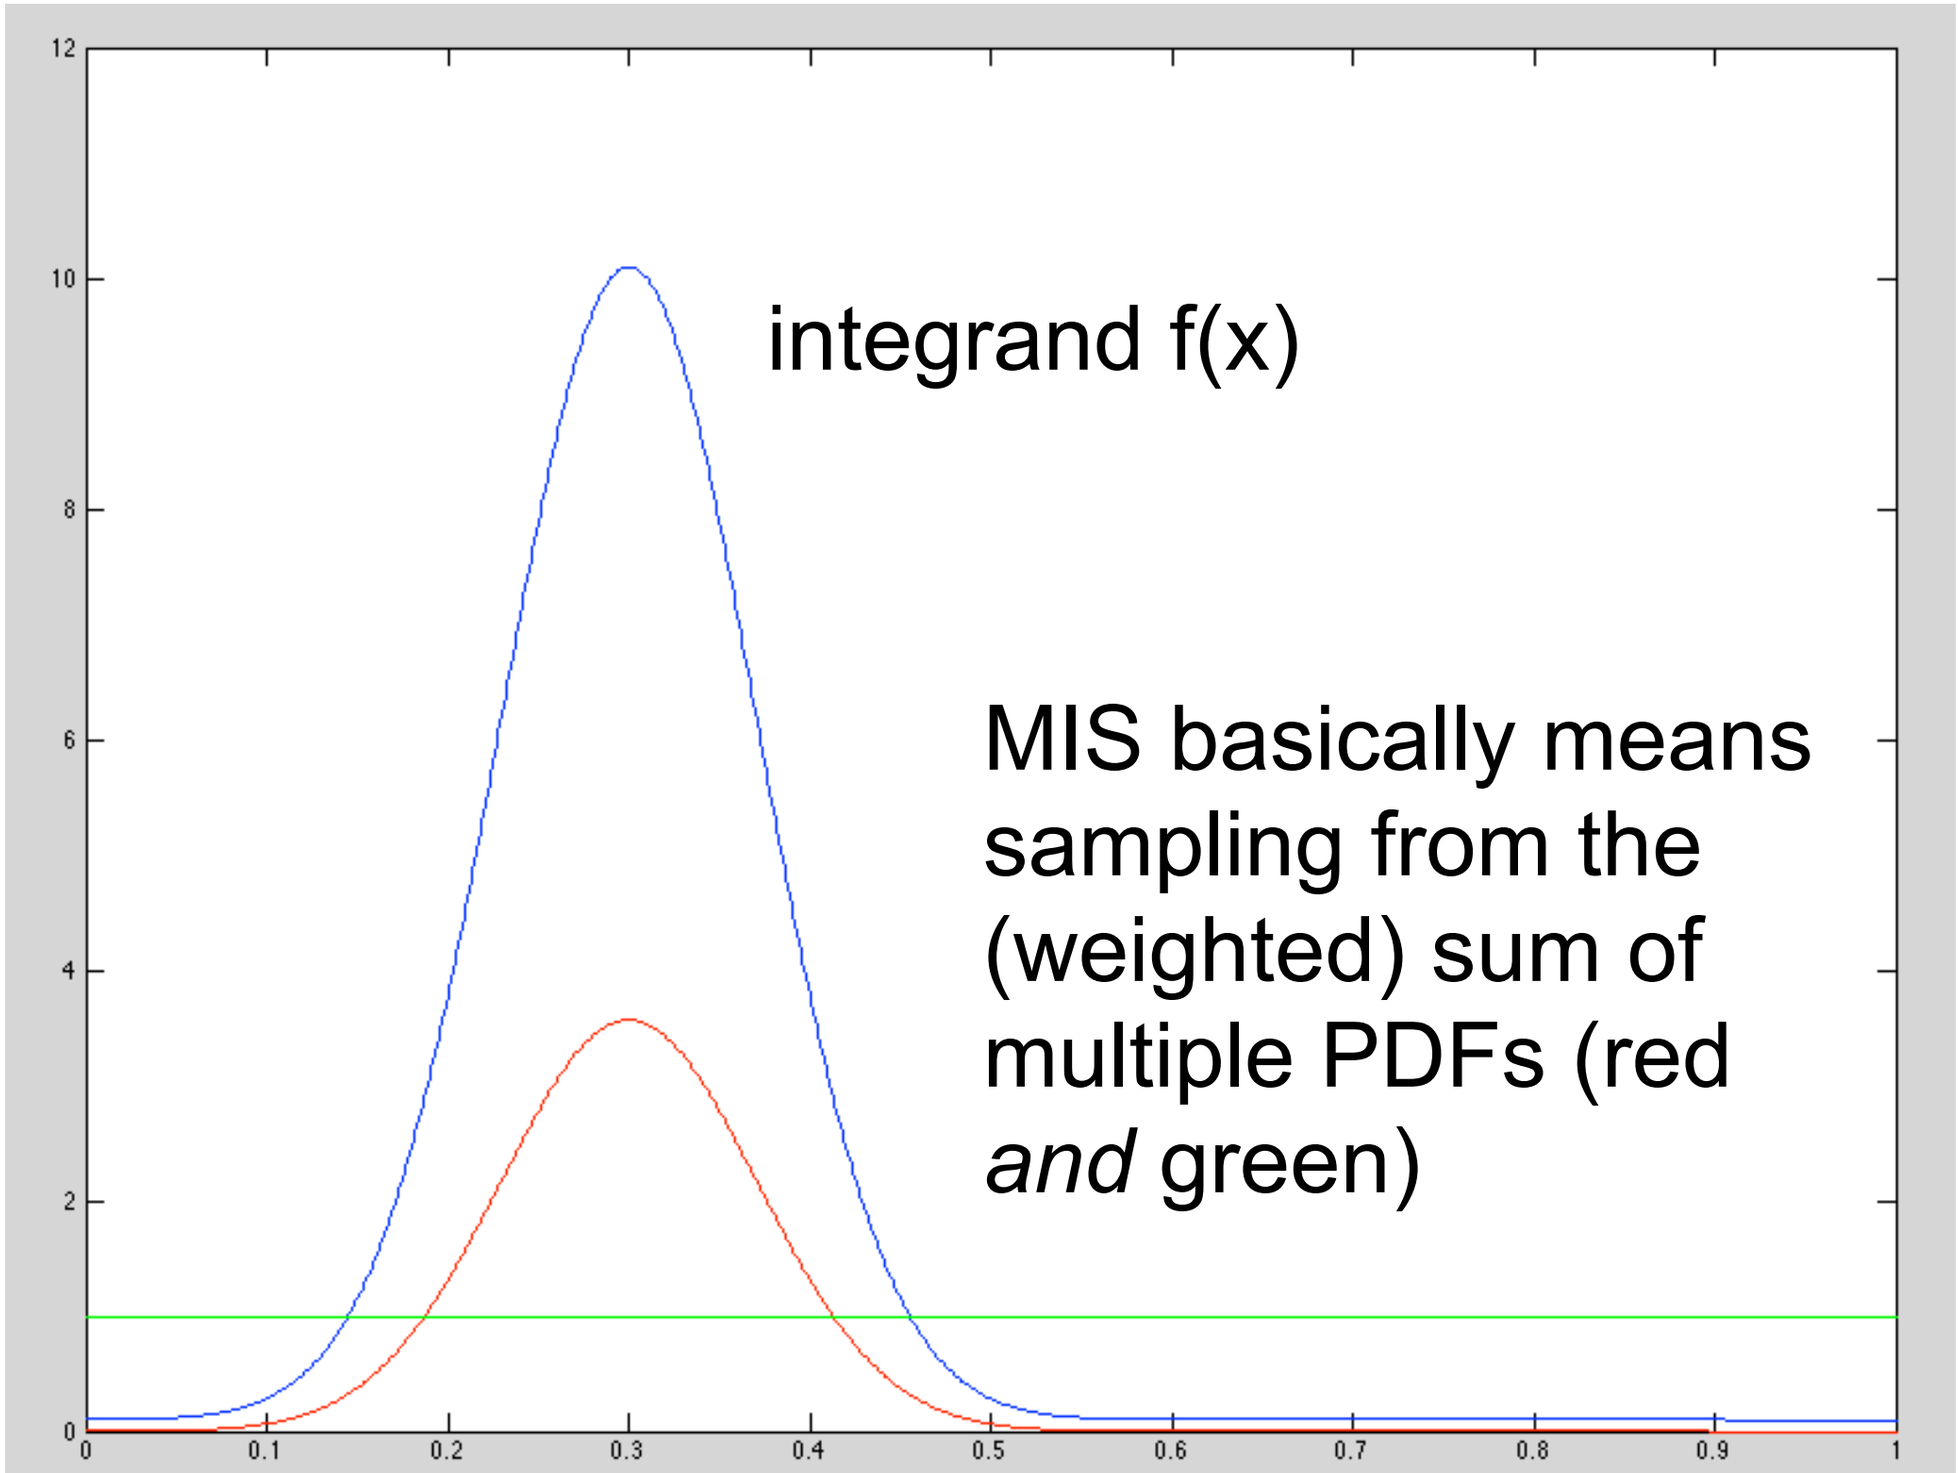

# Basic MIS Recipe

- You have M sampling distributions.
- For each sample *i*
  - Pick one distribution at random, let's say it's the *j*th one
    - You can't do much better than equal chances, i.e. using probability p(*j*) = 1/M for all *j* (Veach 1995, Sec. 5.2) (I assume this below.)
  - Draw a sample $x_i$ from the *j*th distribution
  - Compute

$$W_i = \frac{f(x_i)}{\sum_{j=1}^{M} p(j) p_j(x_i)}$$

  - Take the average of the $W_i$
  - Done!

# What's Going On?

- The above process generates samples with the joint distribution

$$\bar{p}(x) = \sum_{j=1}^{M} p(j) p_j(x)$$

- Hence, we're just computing *f/p* with this new PDF!
  - Note that the p(j)'s are a discrete distribution, their sum must be 1

- *This is an unbiased estimate, just like regular MC.*

# Ha!

– sample weight $f(x_i)/p(x_i)$ over 1000 trials

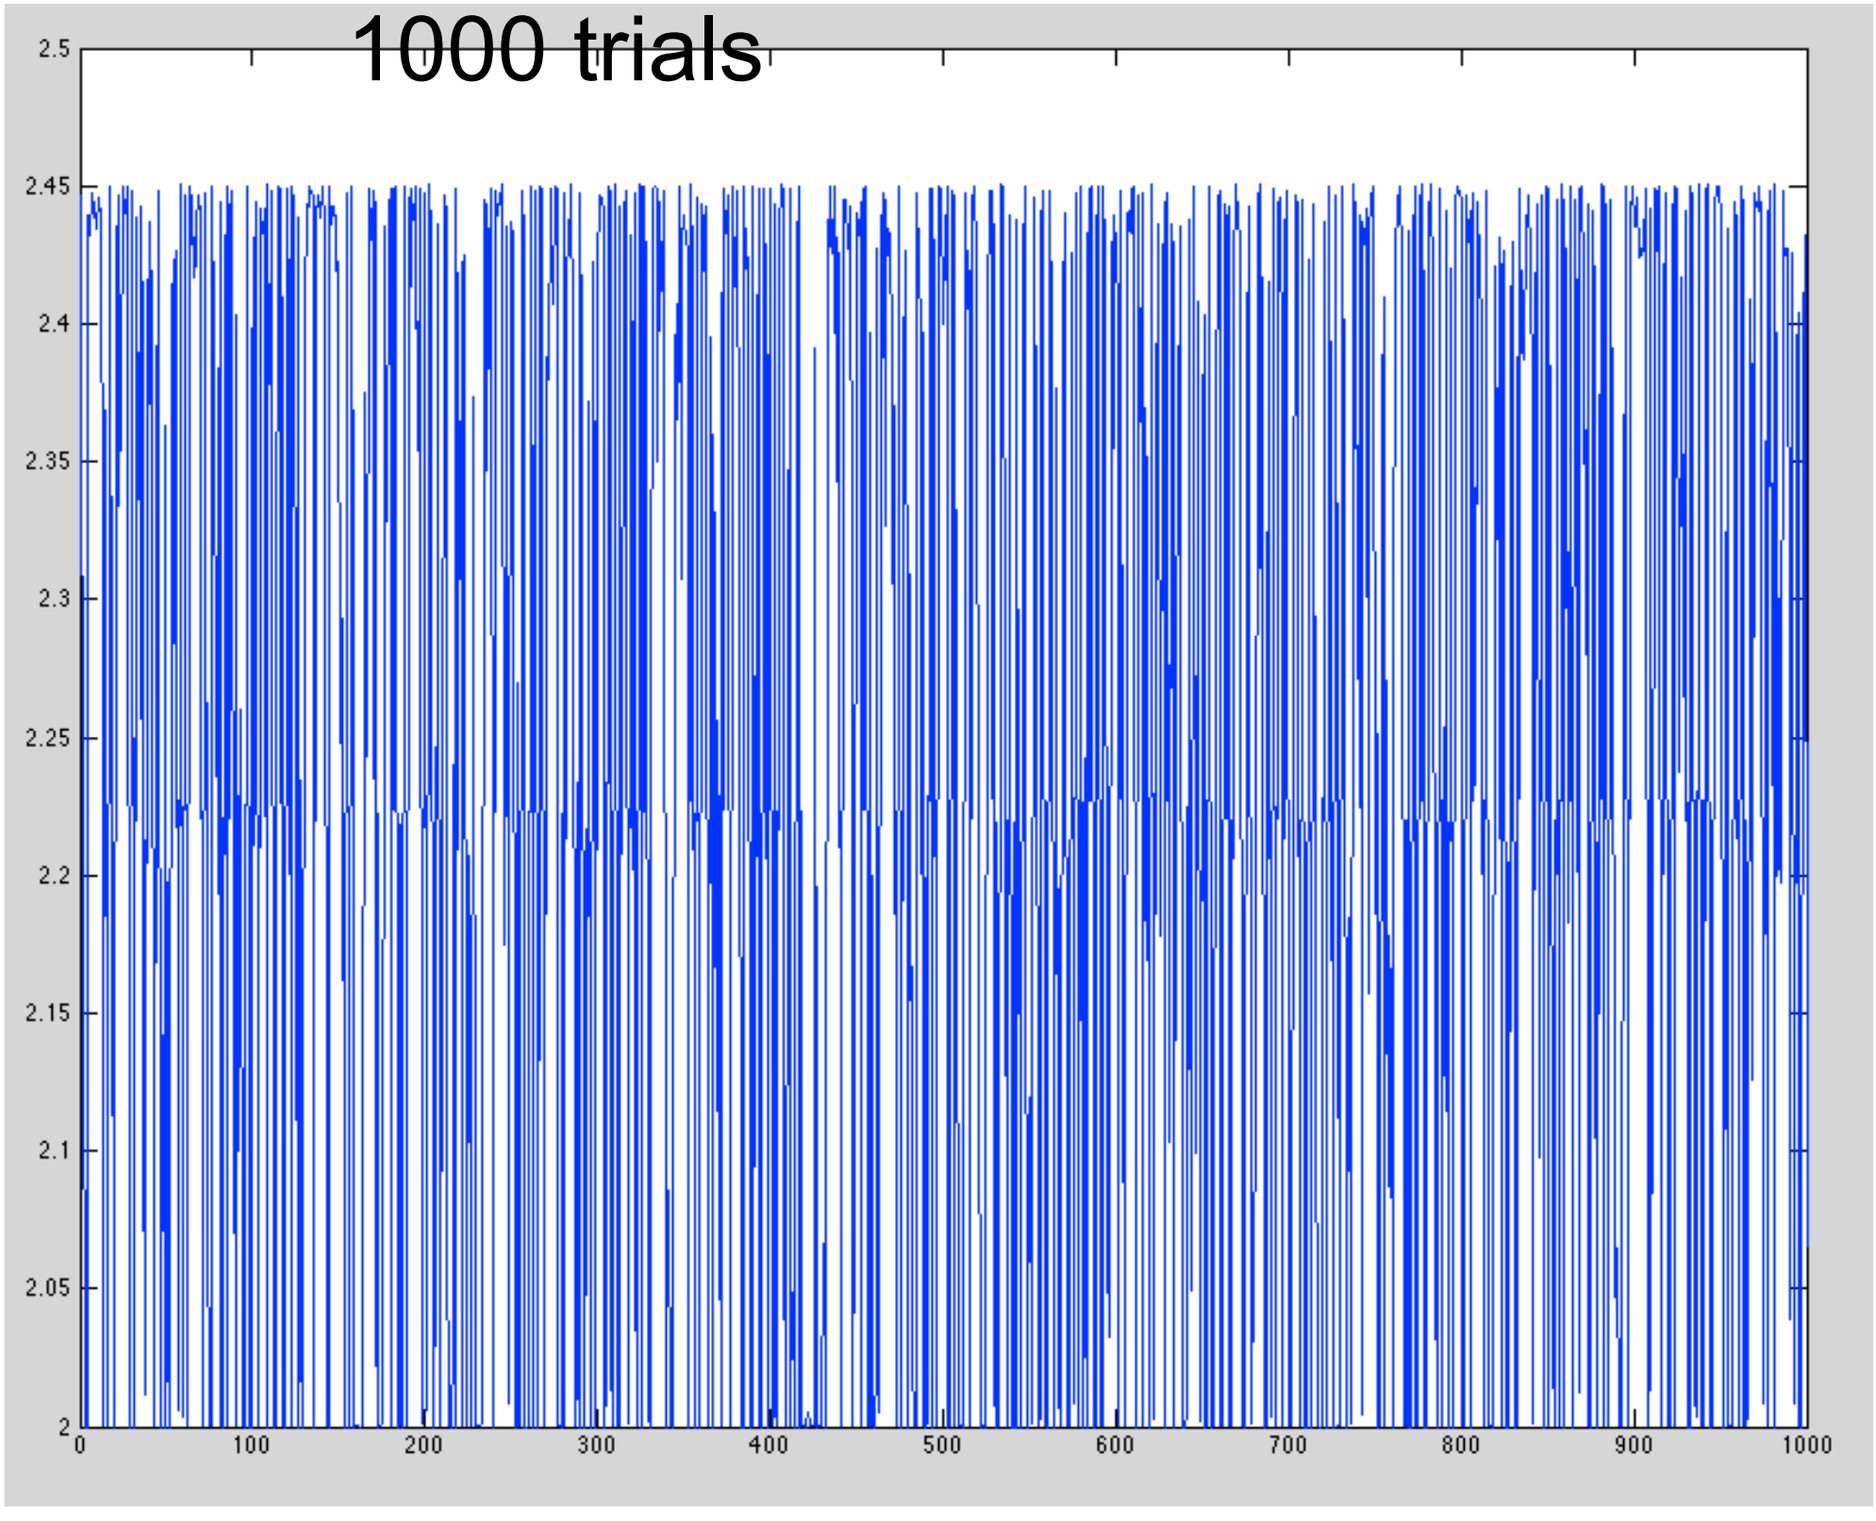

# Integral Estimate, No MIS, 100k samples

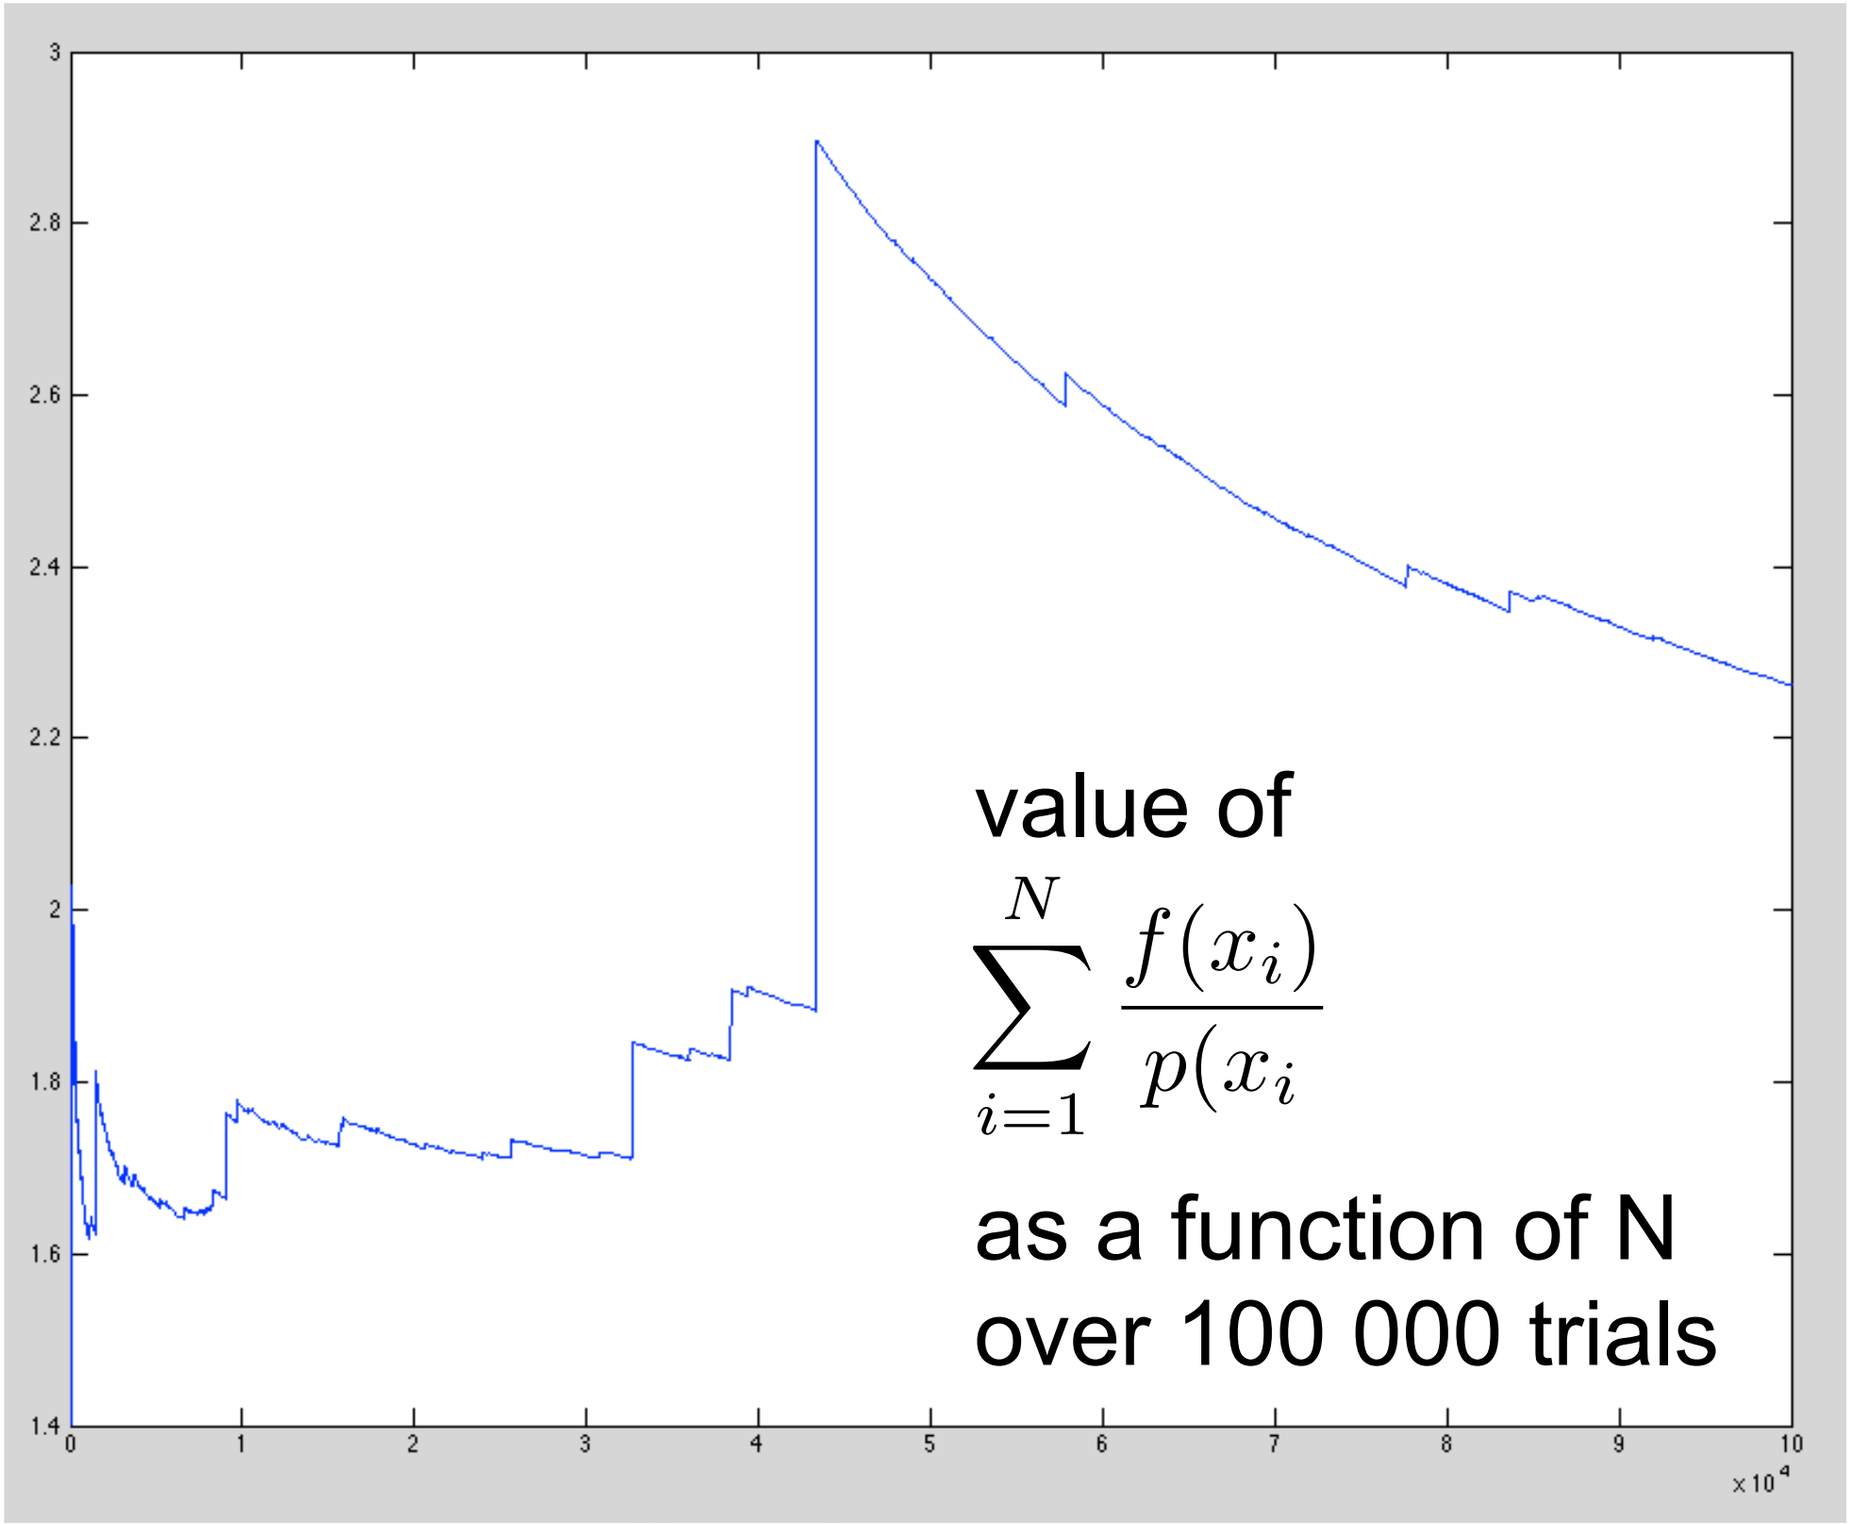

# Integral Estimate, MIS, 1k samples

(100x fewer than previous terrible non-MIS result)

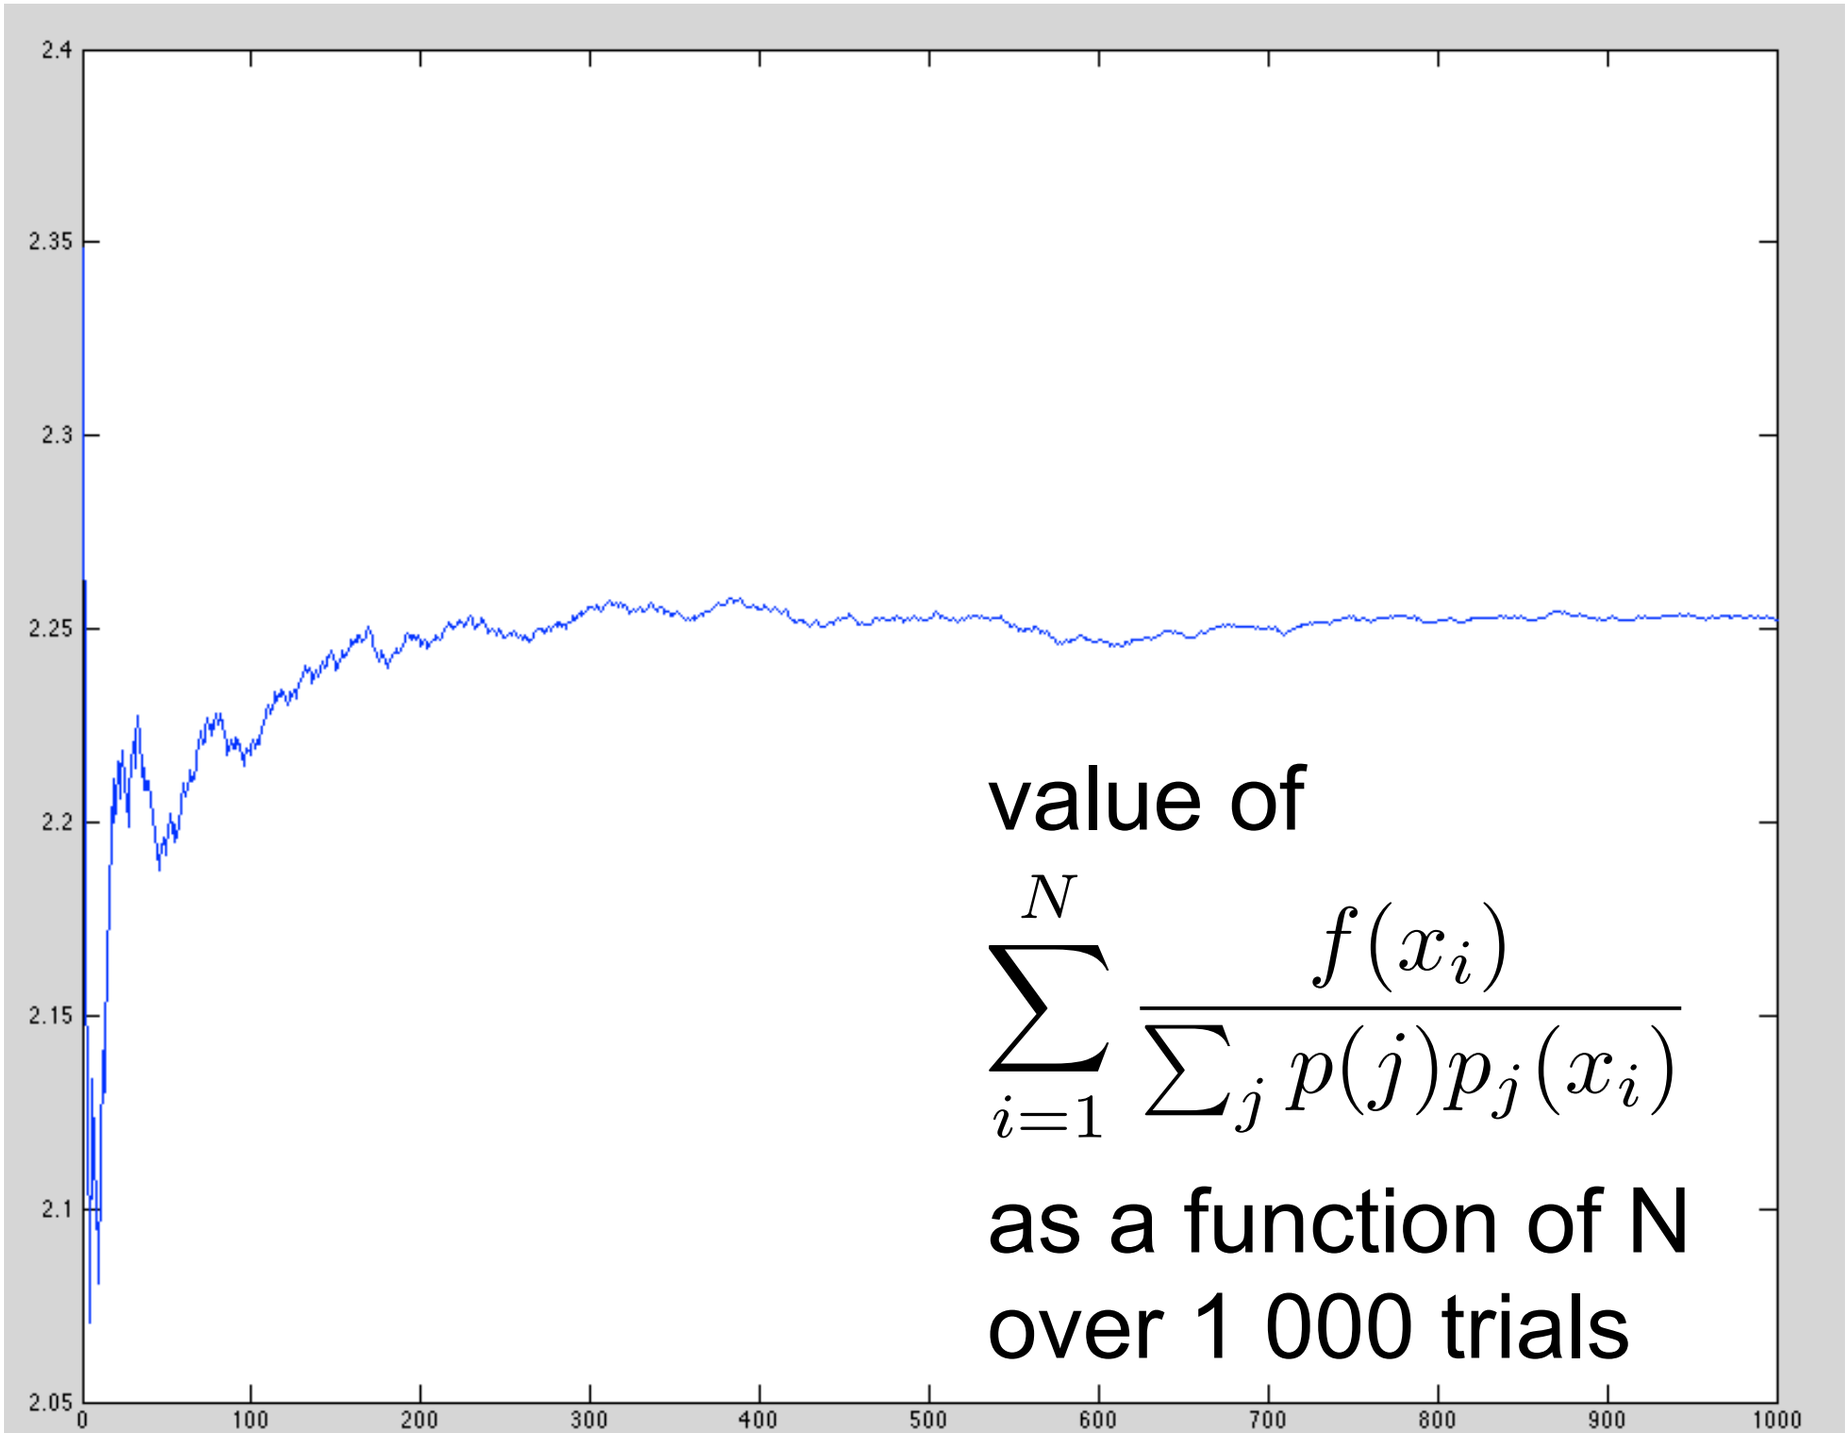

# Integral Estimate, MIS, 1k samples

(100x fewer than previous terrible non-MIS result)

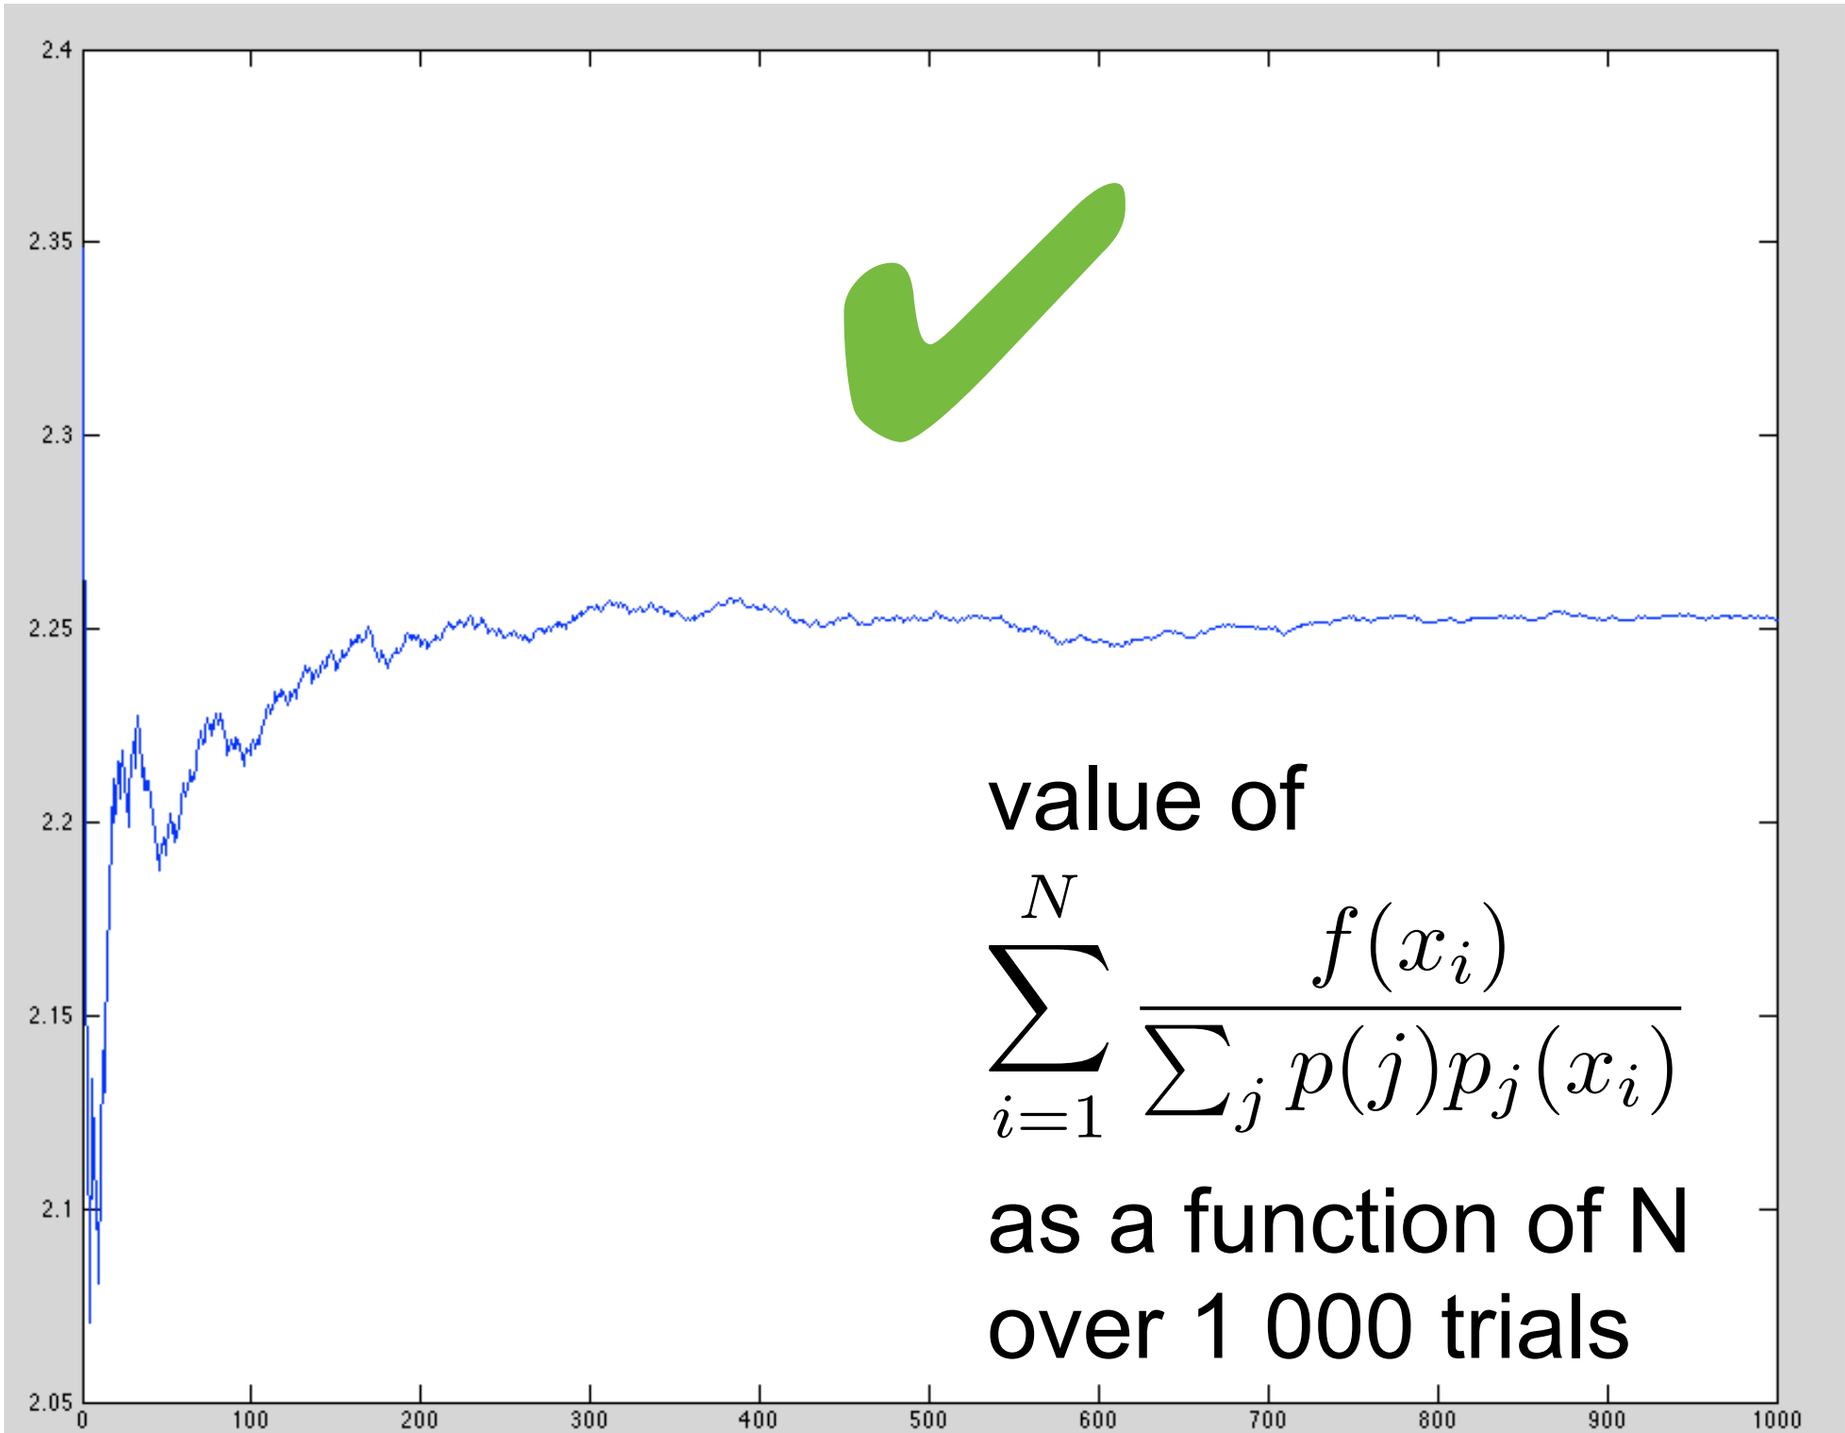

# Bells And Whistles

- This is the basic intuition and approach.

- Veach's 1995 paper contains a long treatment on how to choose the relative weighting between the PDFs and more general ways of constructing $\bar{p}(x)$ based on the individual distributions.

- However, we won't go into this. This process is really general and applies wherever MC can be applied.

# Example: Use in a Path Tracer

- Apart from the direct eye ray, our basic path tracer only accounts for light through shadow rays
  - If the extension ray, which is sampled from the BRDF, hits a light source, we set its contribution to zero.
  - Is this the best we can do?

- Indeed, we can repurpose the extension ray for another purpose: we'll try to make the light connection by both light sampling and BRDF sampling.
  - However we deterministically use both samplers, no random picking.

# Multiple Importance Sampling ✔

**MIS = Sample both ways and optimally combine the samples**

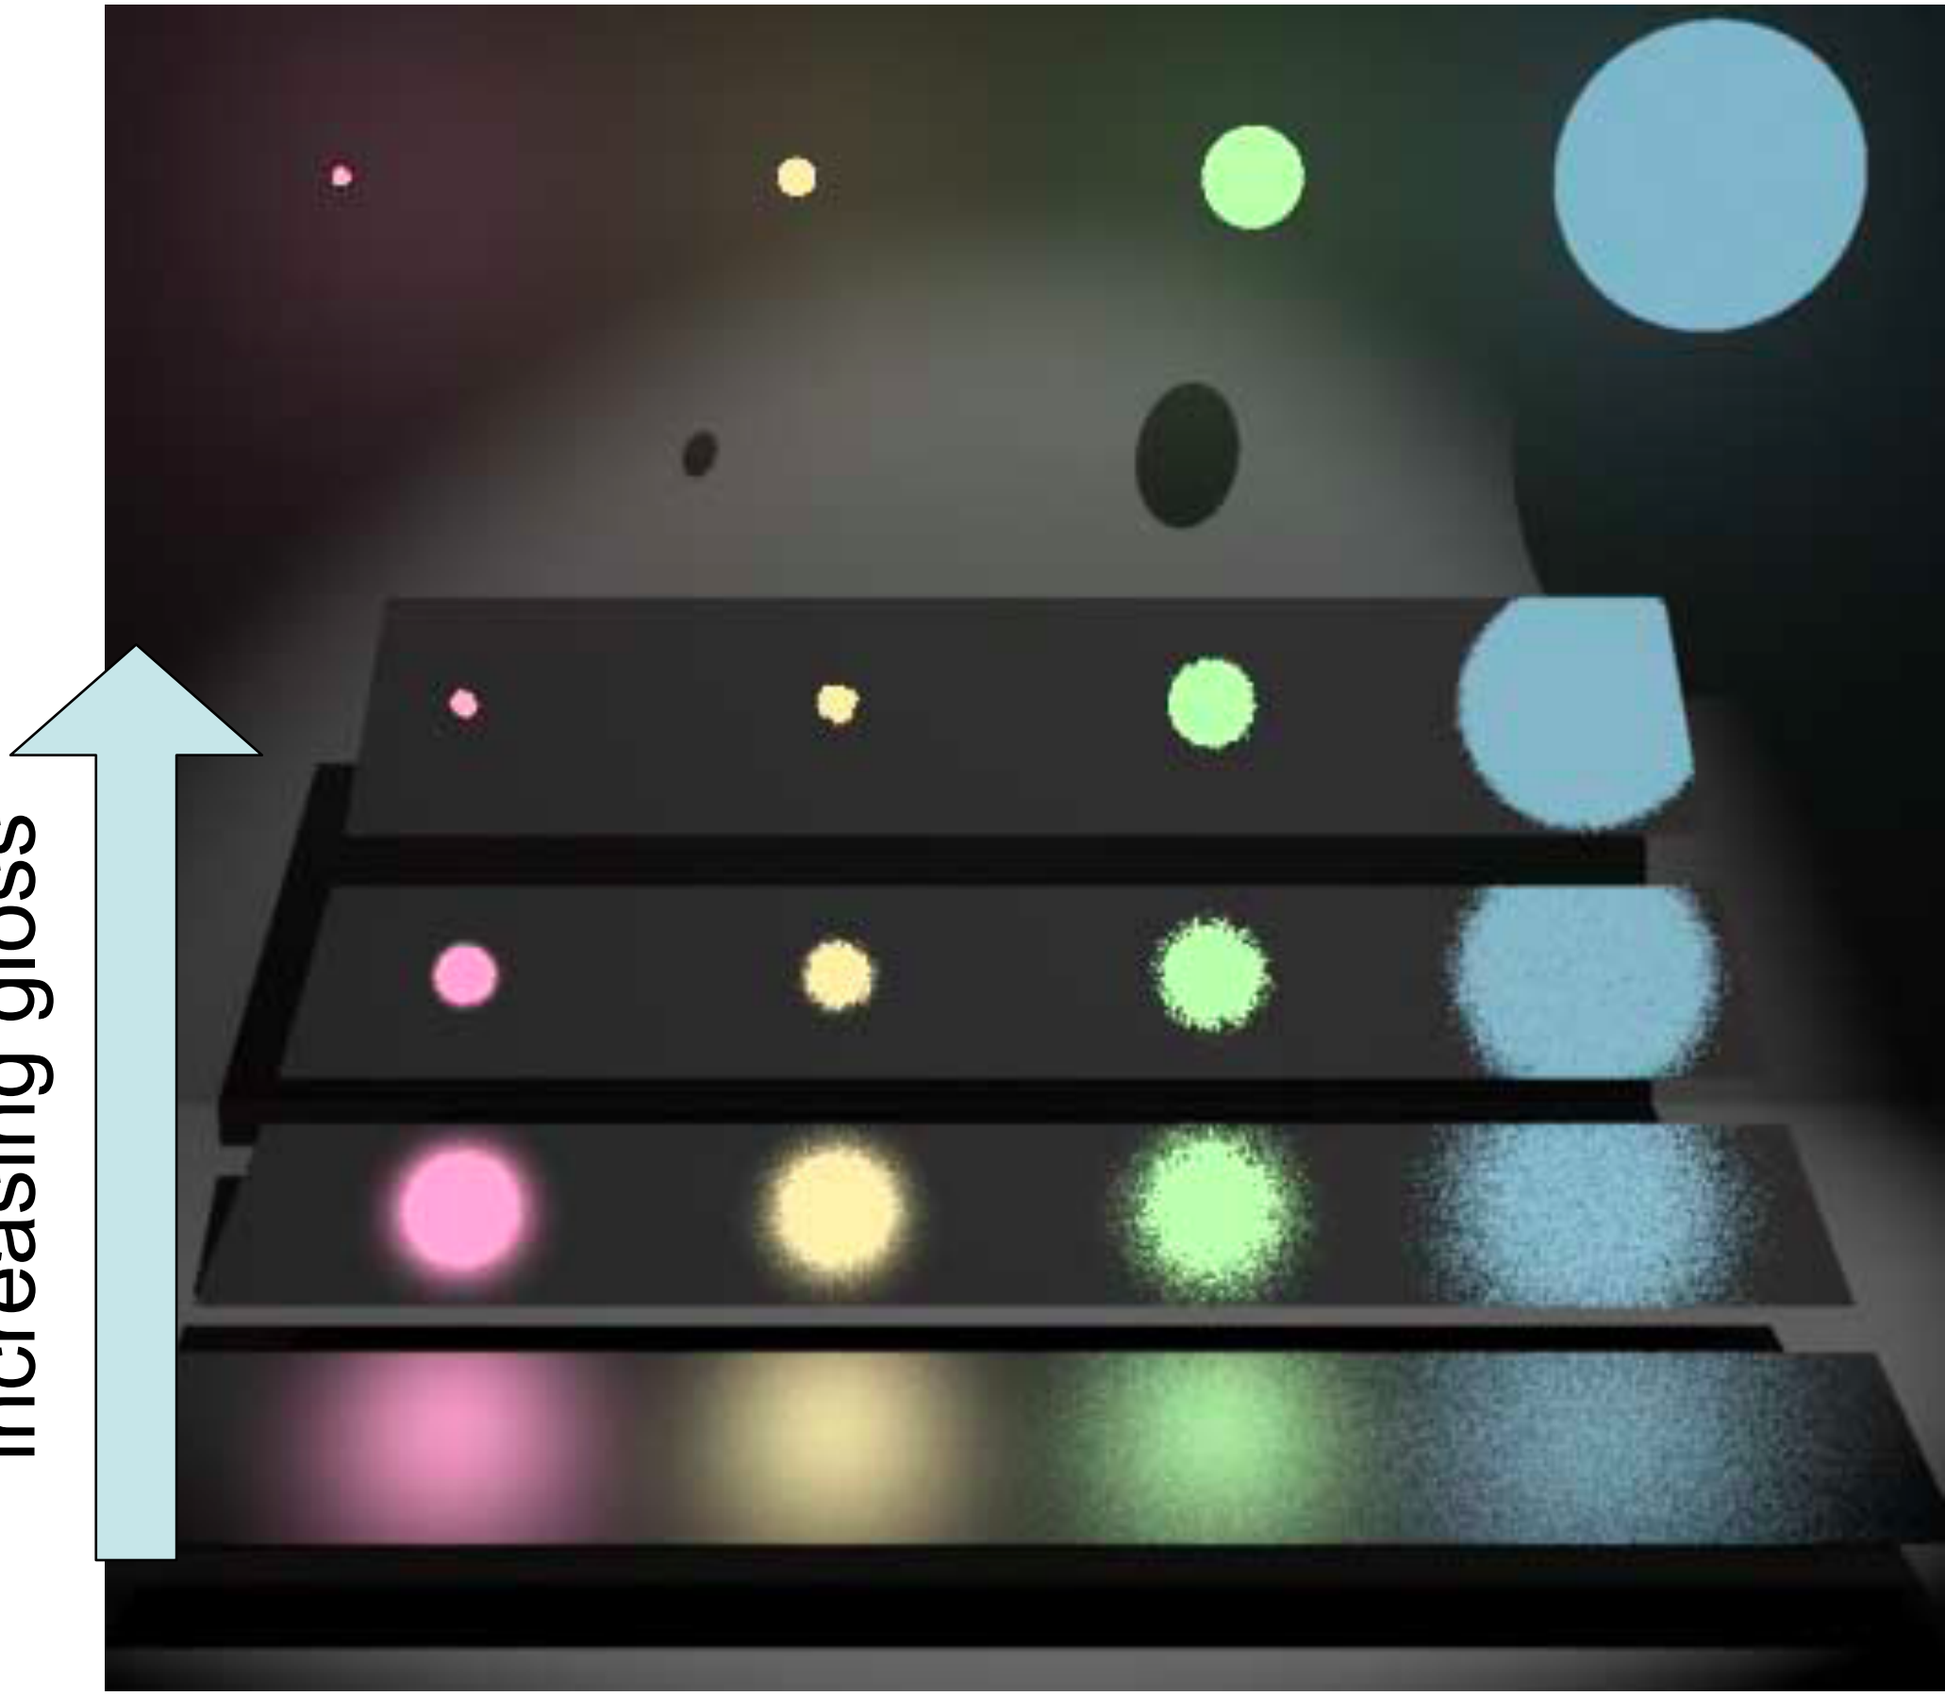

# Questions?

estudibasic, Rendered using Maxwell