CSE 517 — MACHINE LEARNING

QUAN NGUYEN

BAYESIAN OPTIMIZATION

BLACK-BOX OPTIMIZATION

BLACK-BOX OPTIMIZATION

Want to computationally solve for

$$x^* = \arg\max_D f(x)$$

BLACK-BOX OPTIMIZATION

Want to computationally solve for

$$x^* = \arg\max_D f(x)$$

Challenges: The objective function f

BLACK-BOX OPTIMIZATION

Want to computationally solve for

$$x^* = \arg\max_D f(x)$$

Challenges: The objective function f

- is **expensive** to evaluate (money, time, safety conditions, etc.)

BLACK-BOX OPTIMIZATION

Want to computationally solve for

$$x^* = \arg\max_D f(x)$$

Challenges: The objective function f

- is **expensive** to evaluate (money, time, safety conditions, etc.)
- has **no analytical form** (e.g., $f(x) \neq x^2 + x - 1$)

BLACK-BOX OPTIMIZATION

Want to computationally solve for

$$x^* = \arg\max_D f(x)$$

Challenges: The objective function f

- is **expensive** to evaluate (money, time, safety conditions, etc.)
- has **no analytical form** (e.g., $f(x) \neq x^2 + x - 1$)
- has **no gradient** information (cannot run gradient descent, L-BFGS, etc.)

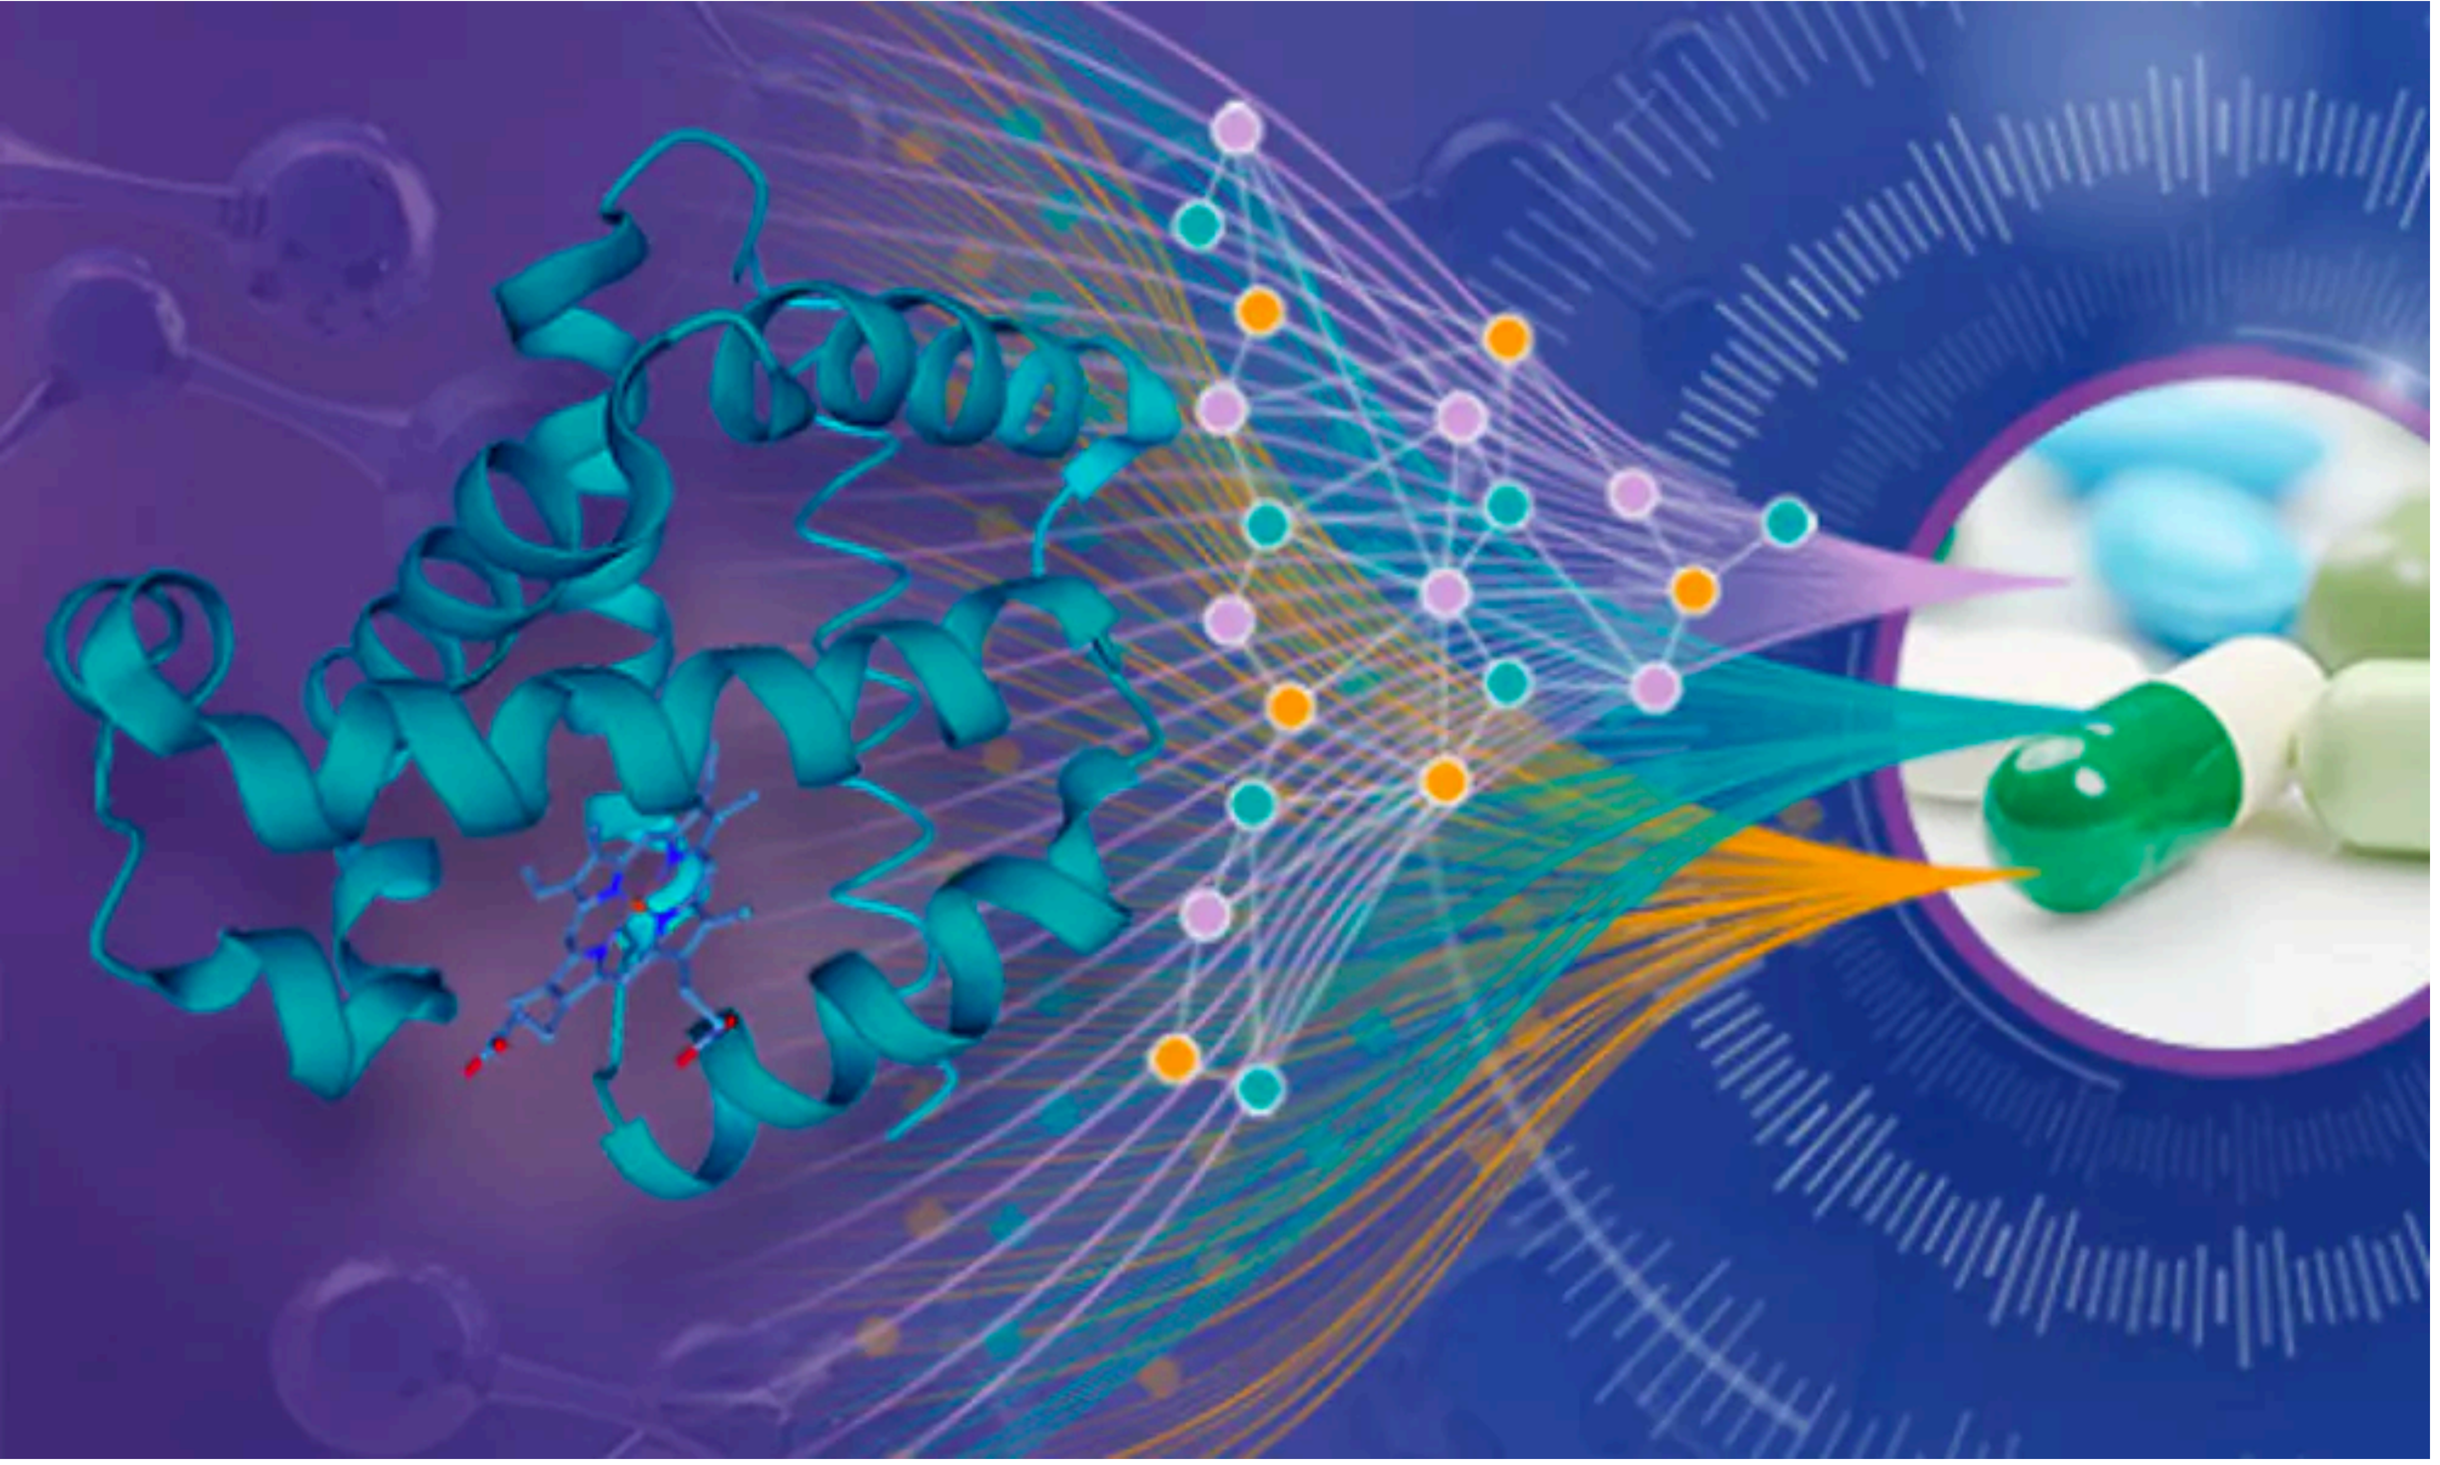

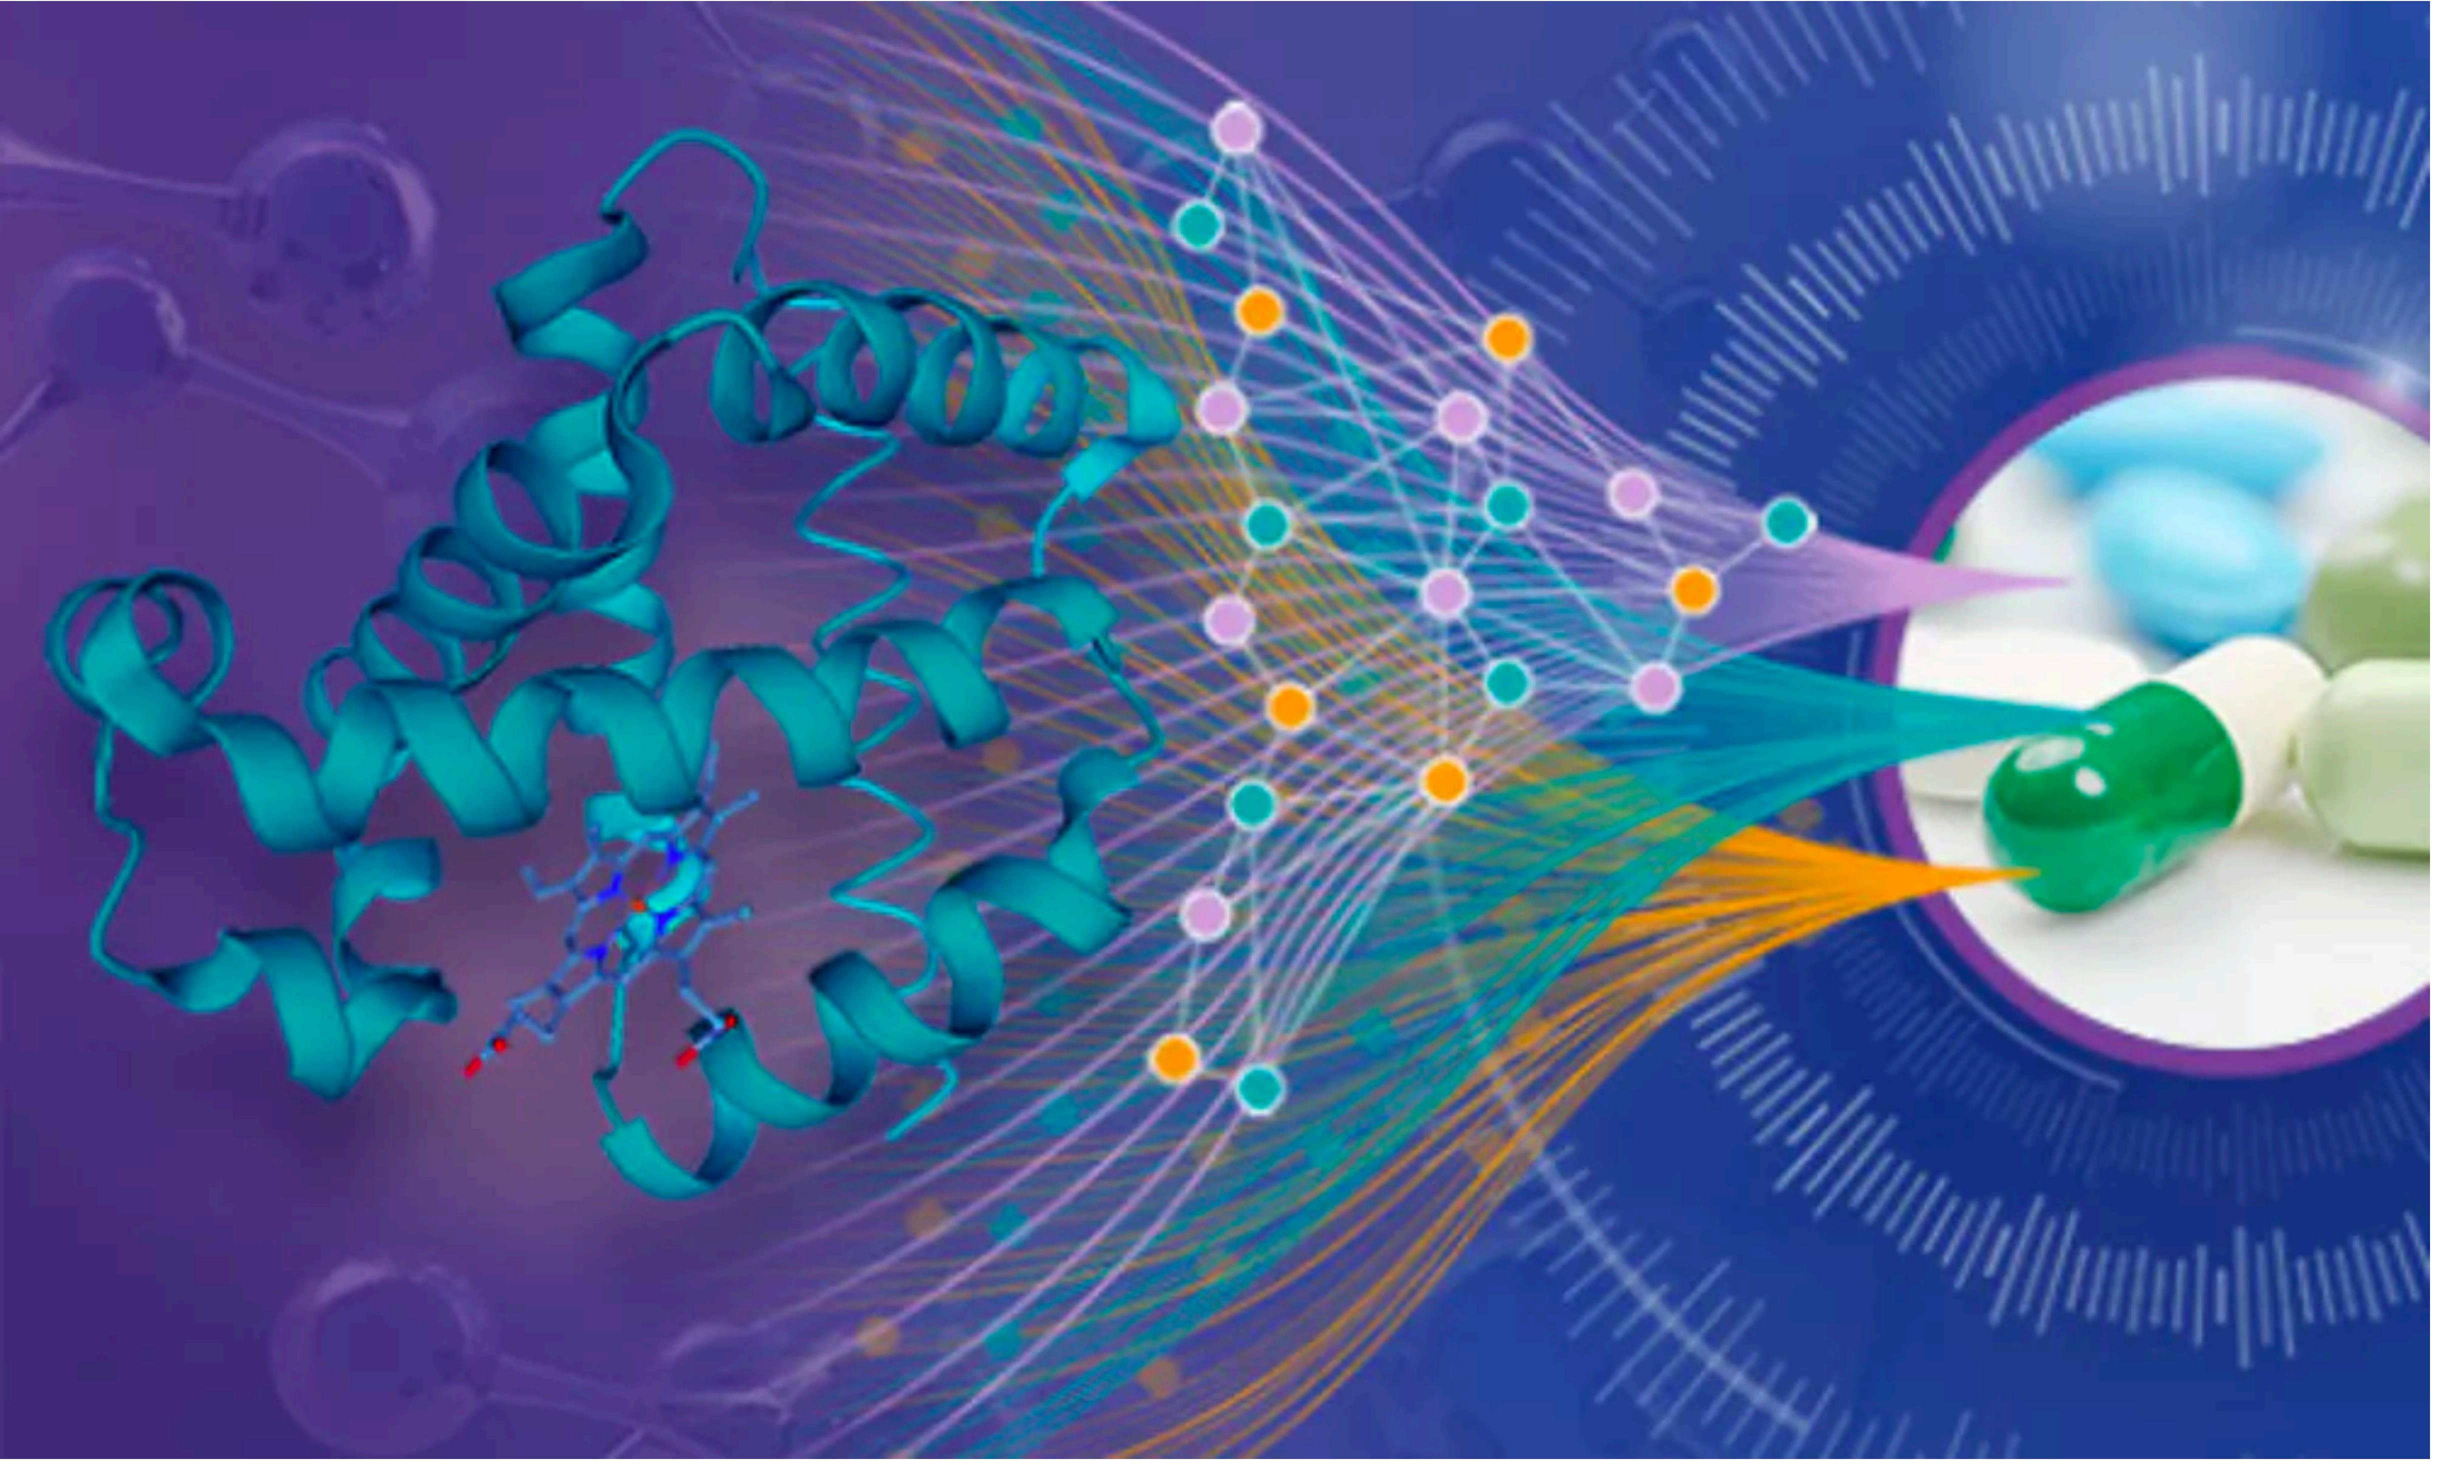

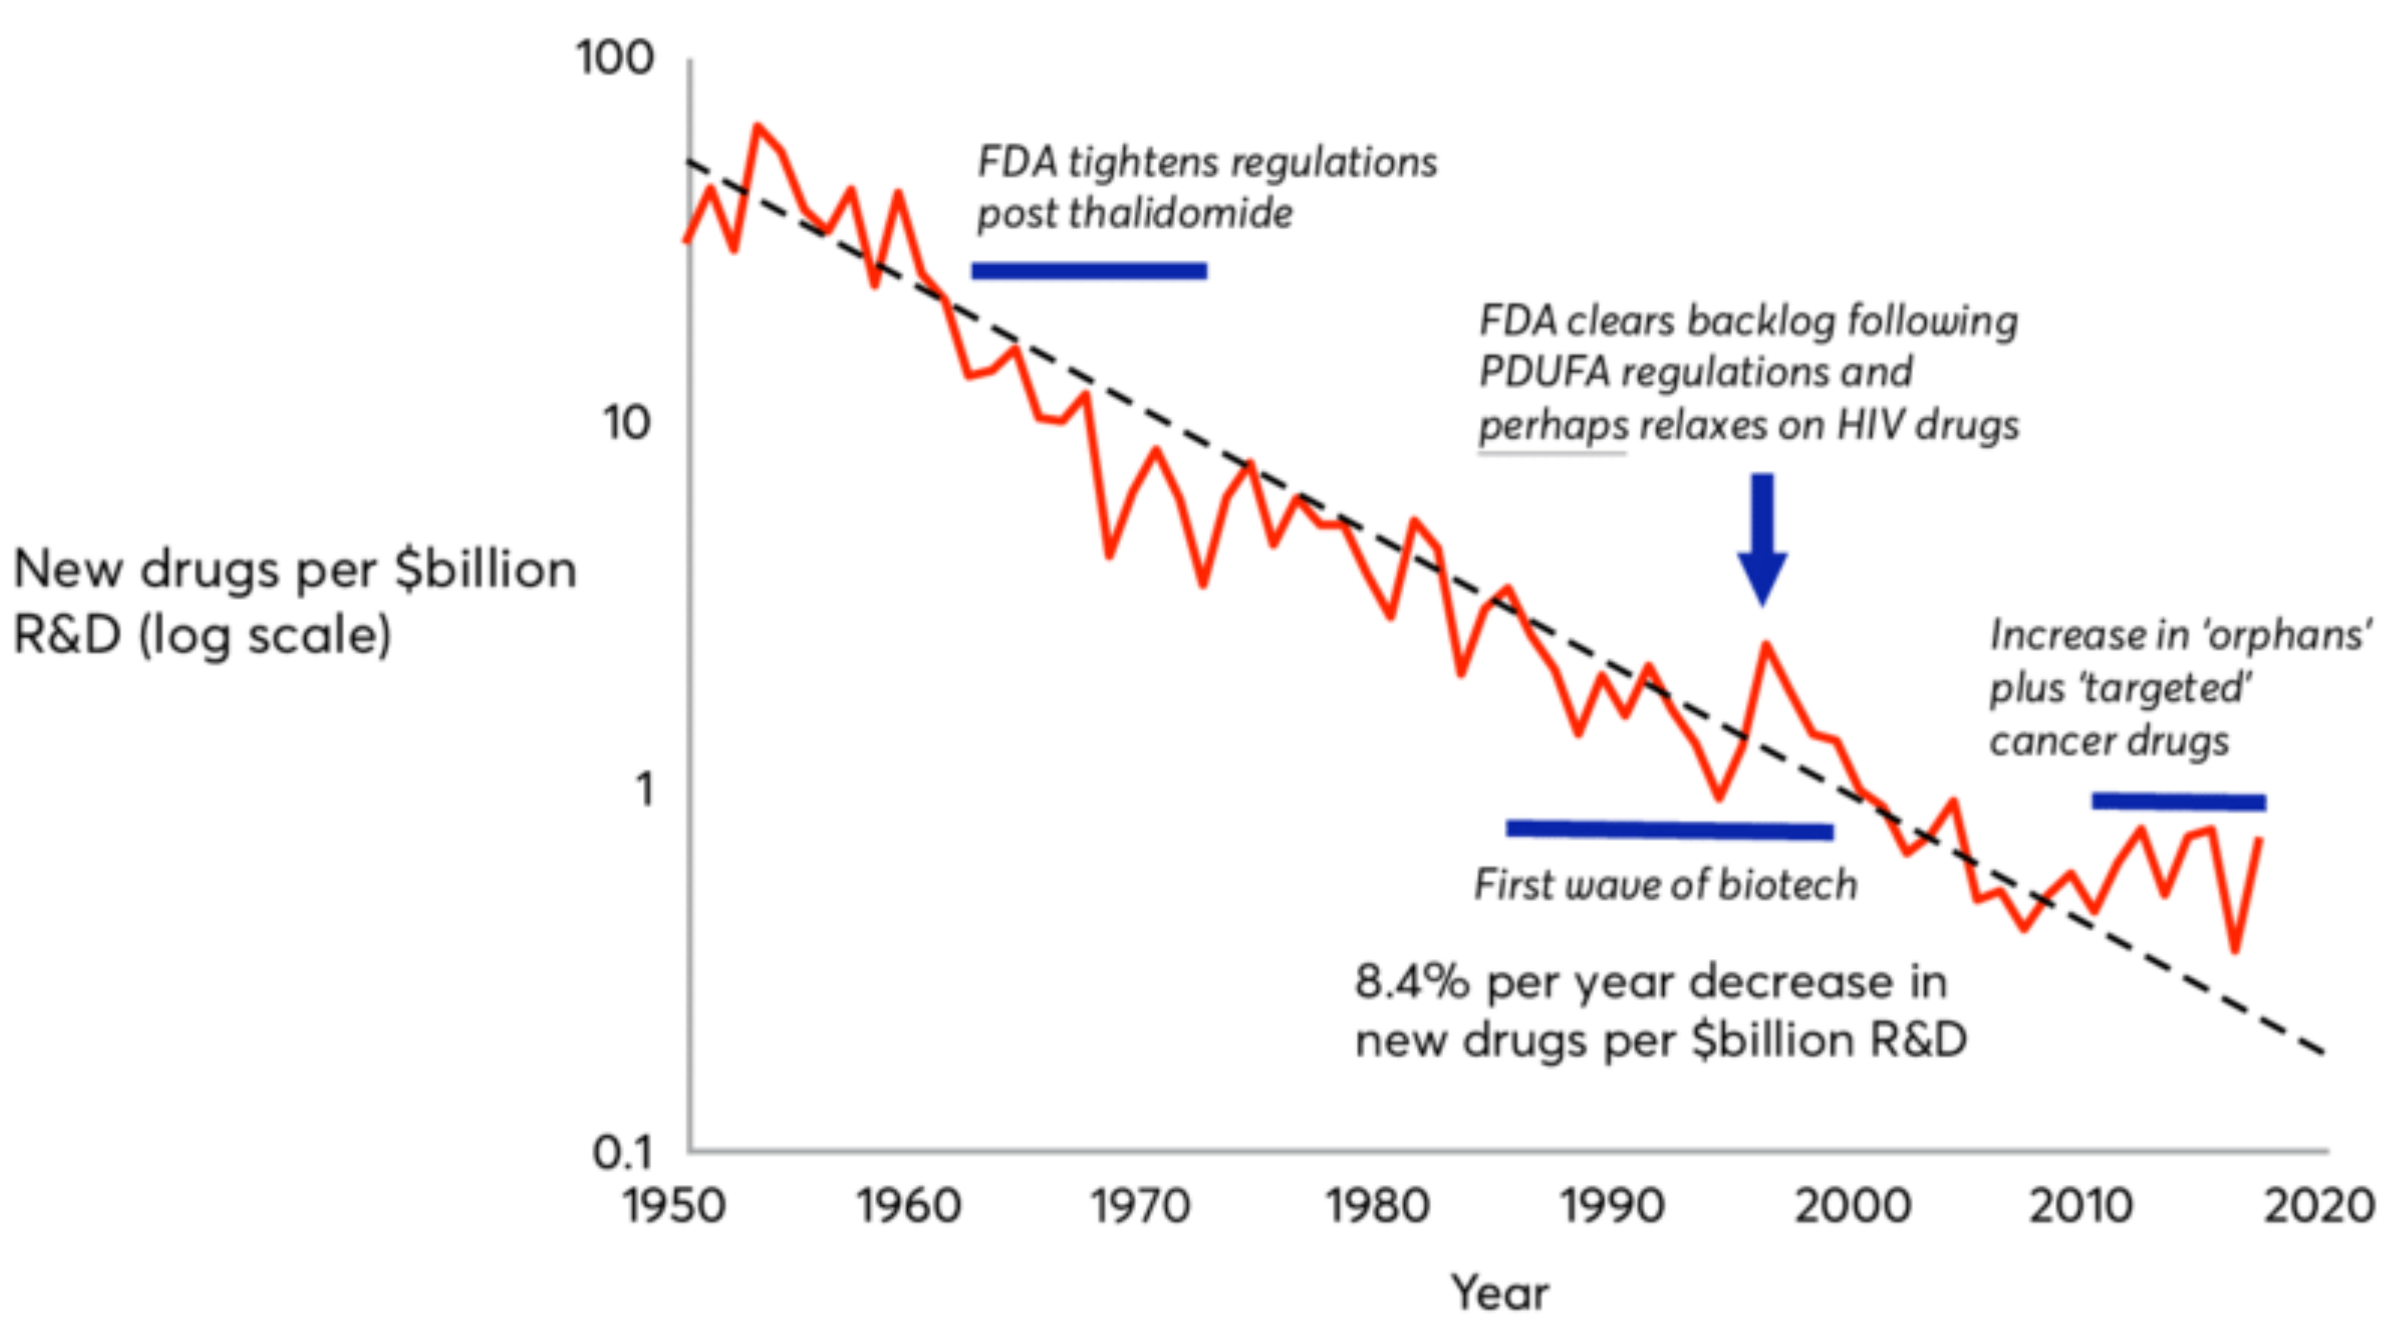

New drugs per $billion R&D (log scale)
100
10
1
0.1
FDA tightens regulations post thalidomide
FDA clears backlog following PDUFA regulations and perhaps relaxes on HIV drugs
Increase in 'orphans' plus 'targeted' cancer drugs
First wave of biotech
8.4% per year decrease in new drugs per $billion R&D
1950
1960
1970
1980
1990
2000
2010
2020
Year

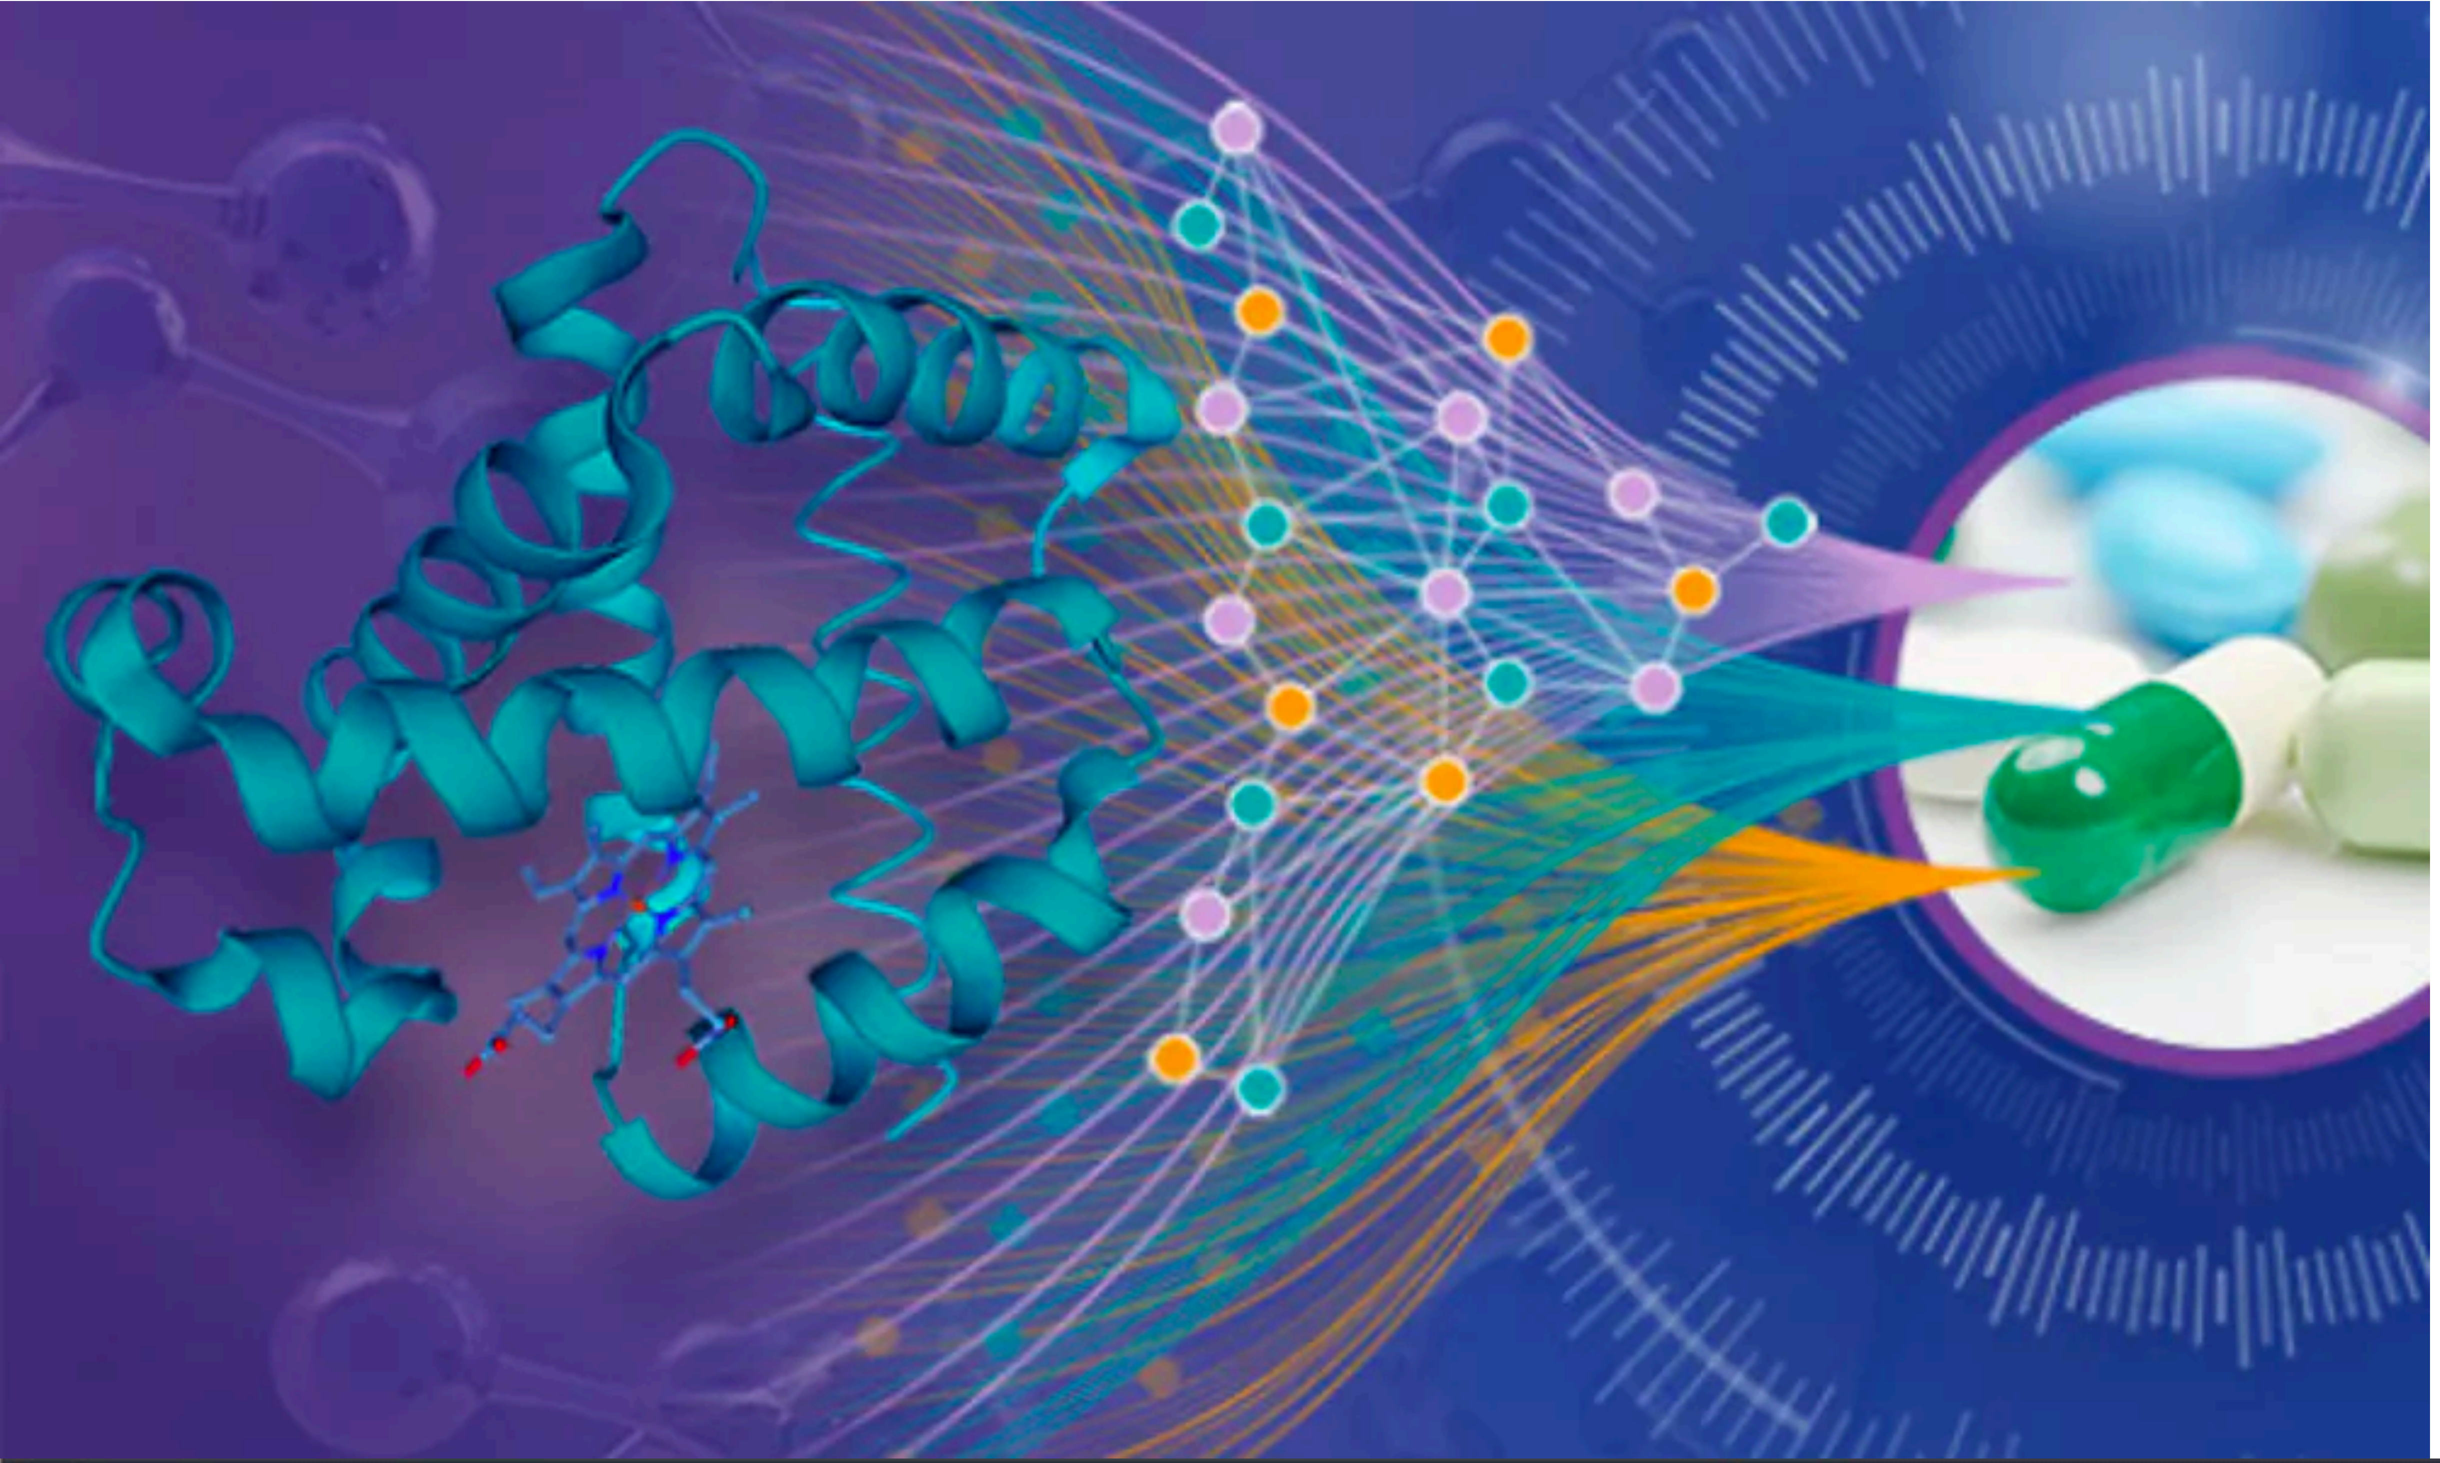

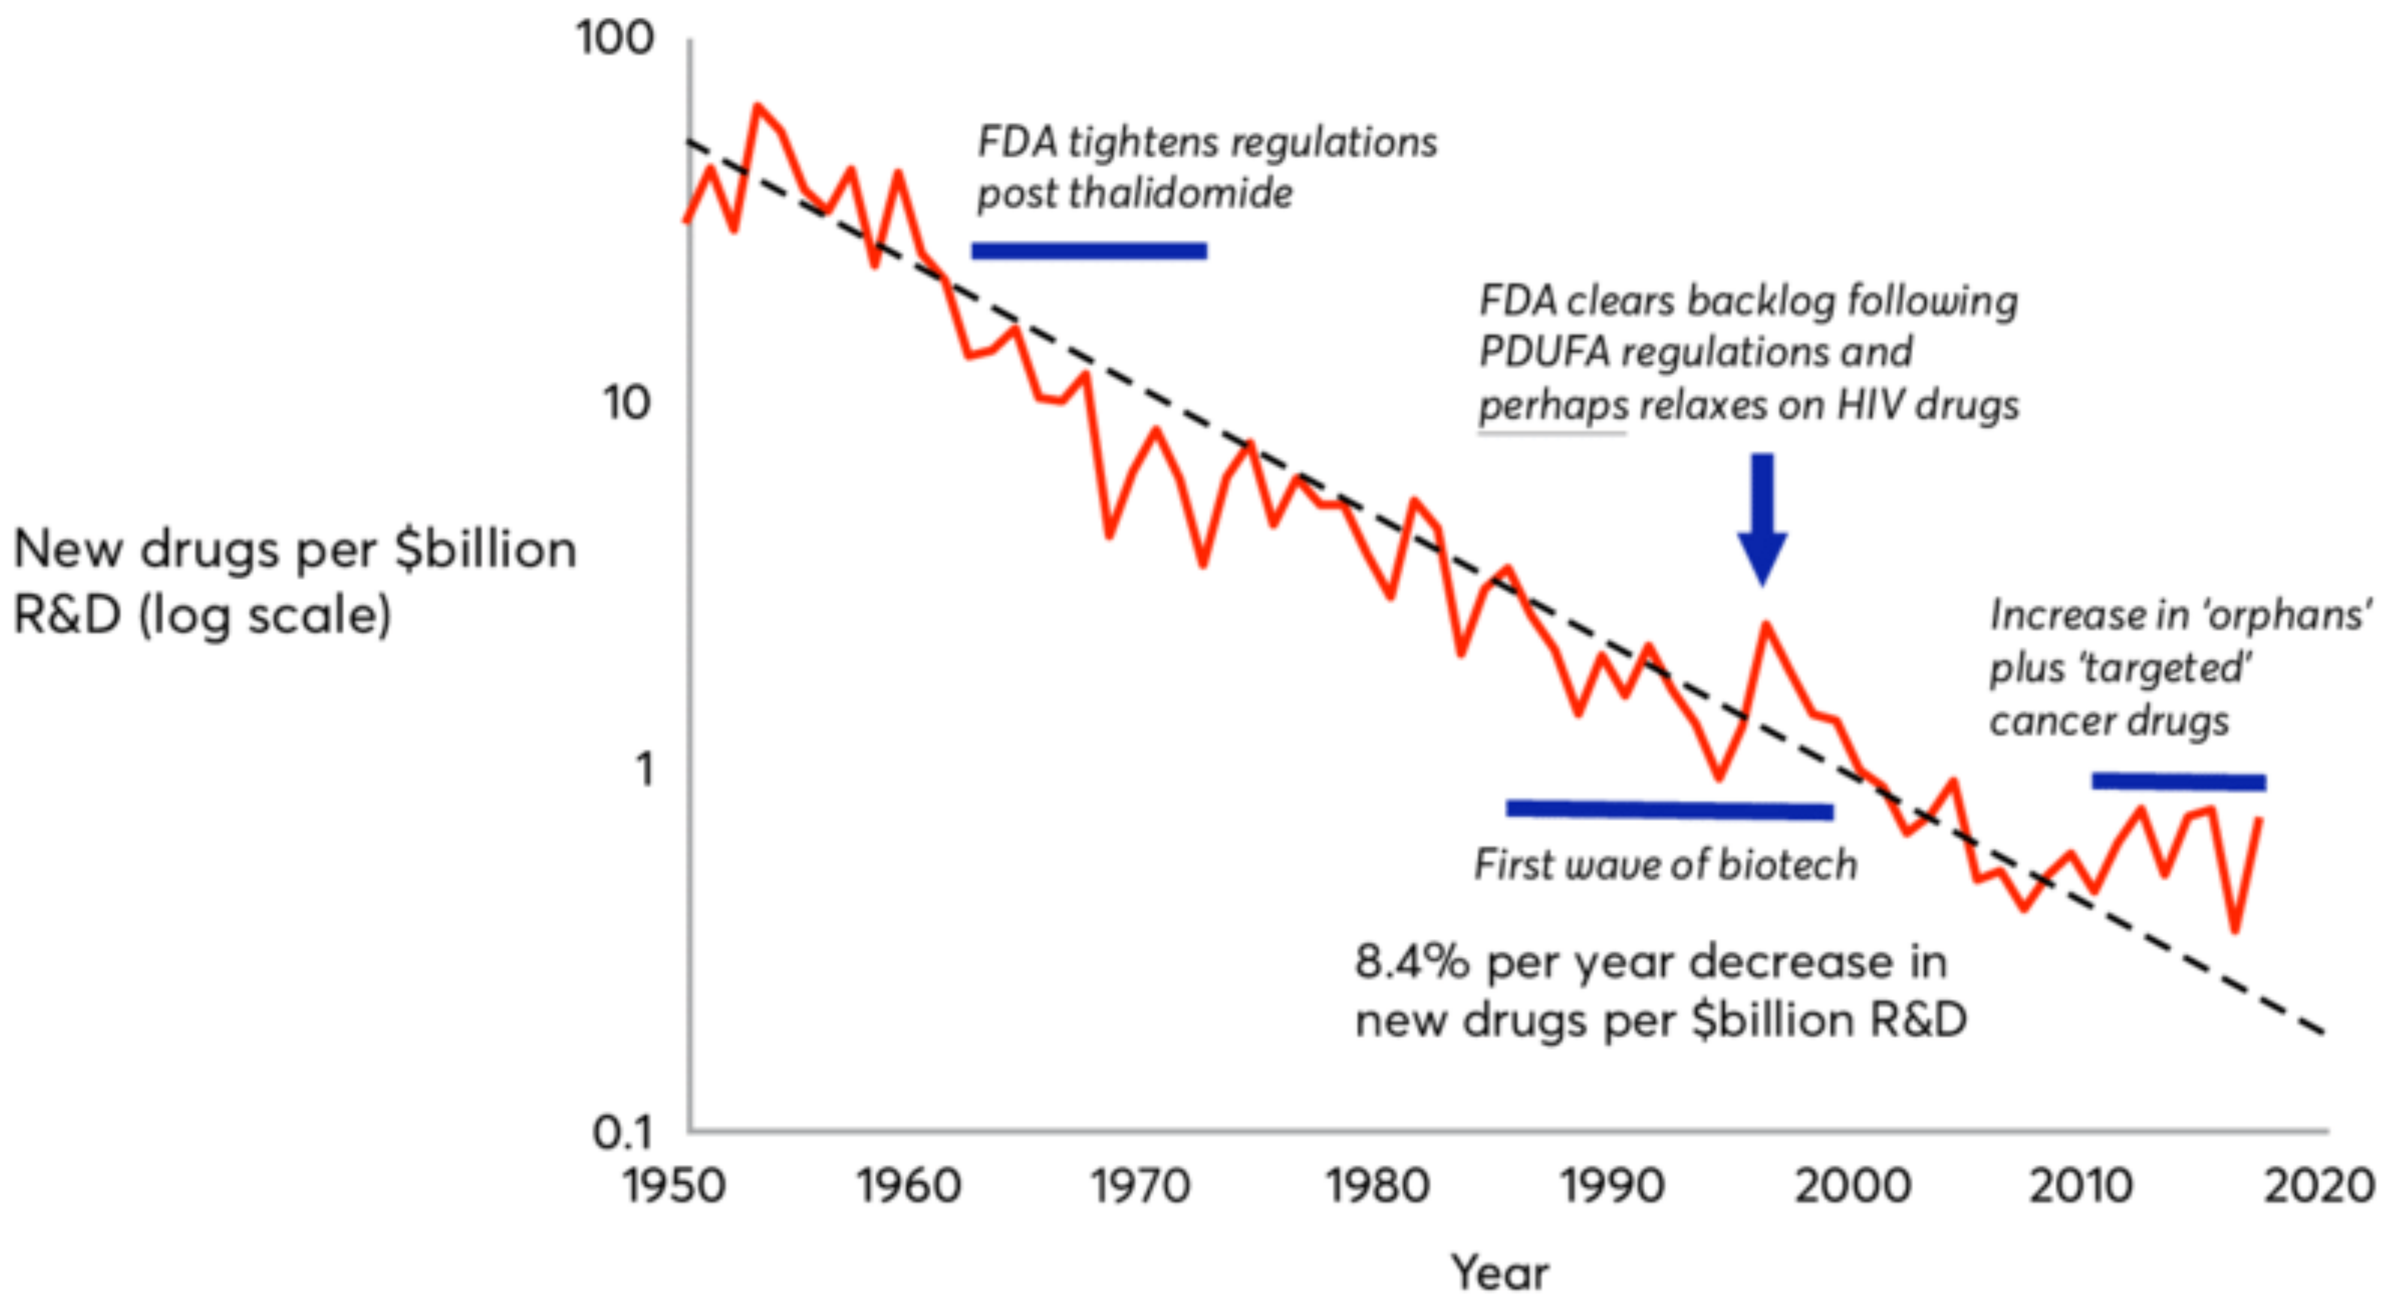

100
10
1
0.1
New drugs per $billion R&D (log scale)
FDA tightens regulations post thalidomide
FDA clears backlog following PDUFA regulations and perhaps relaxes on HIV drugs
Increase in 'orphans' plus 'targeted' cancer drugs
First wave of biotech
8.4% per year decrease in new drugs per $billion R&D
1950
1960
1970
1980
1990
2000
2010
2020
Year

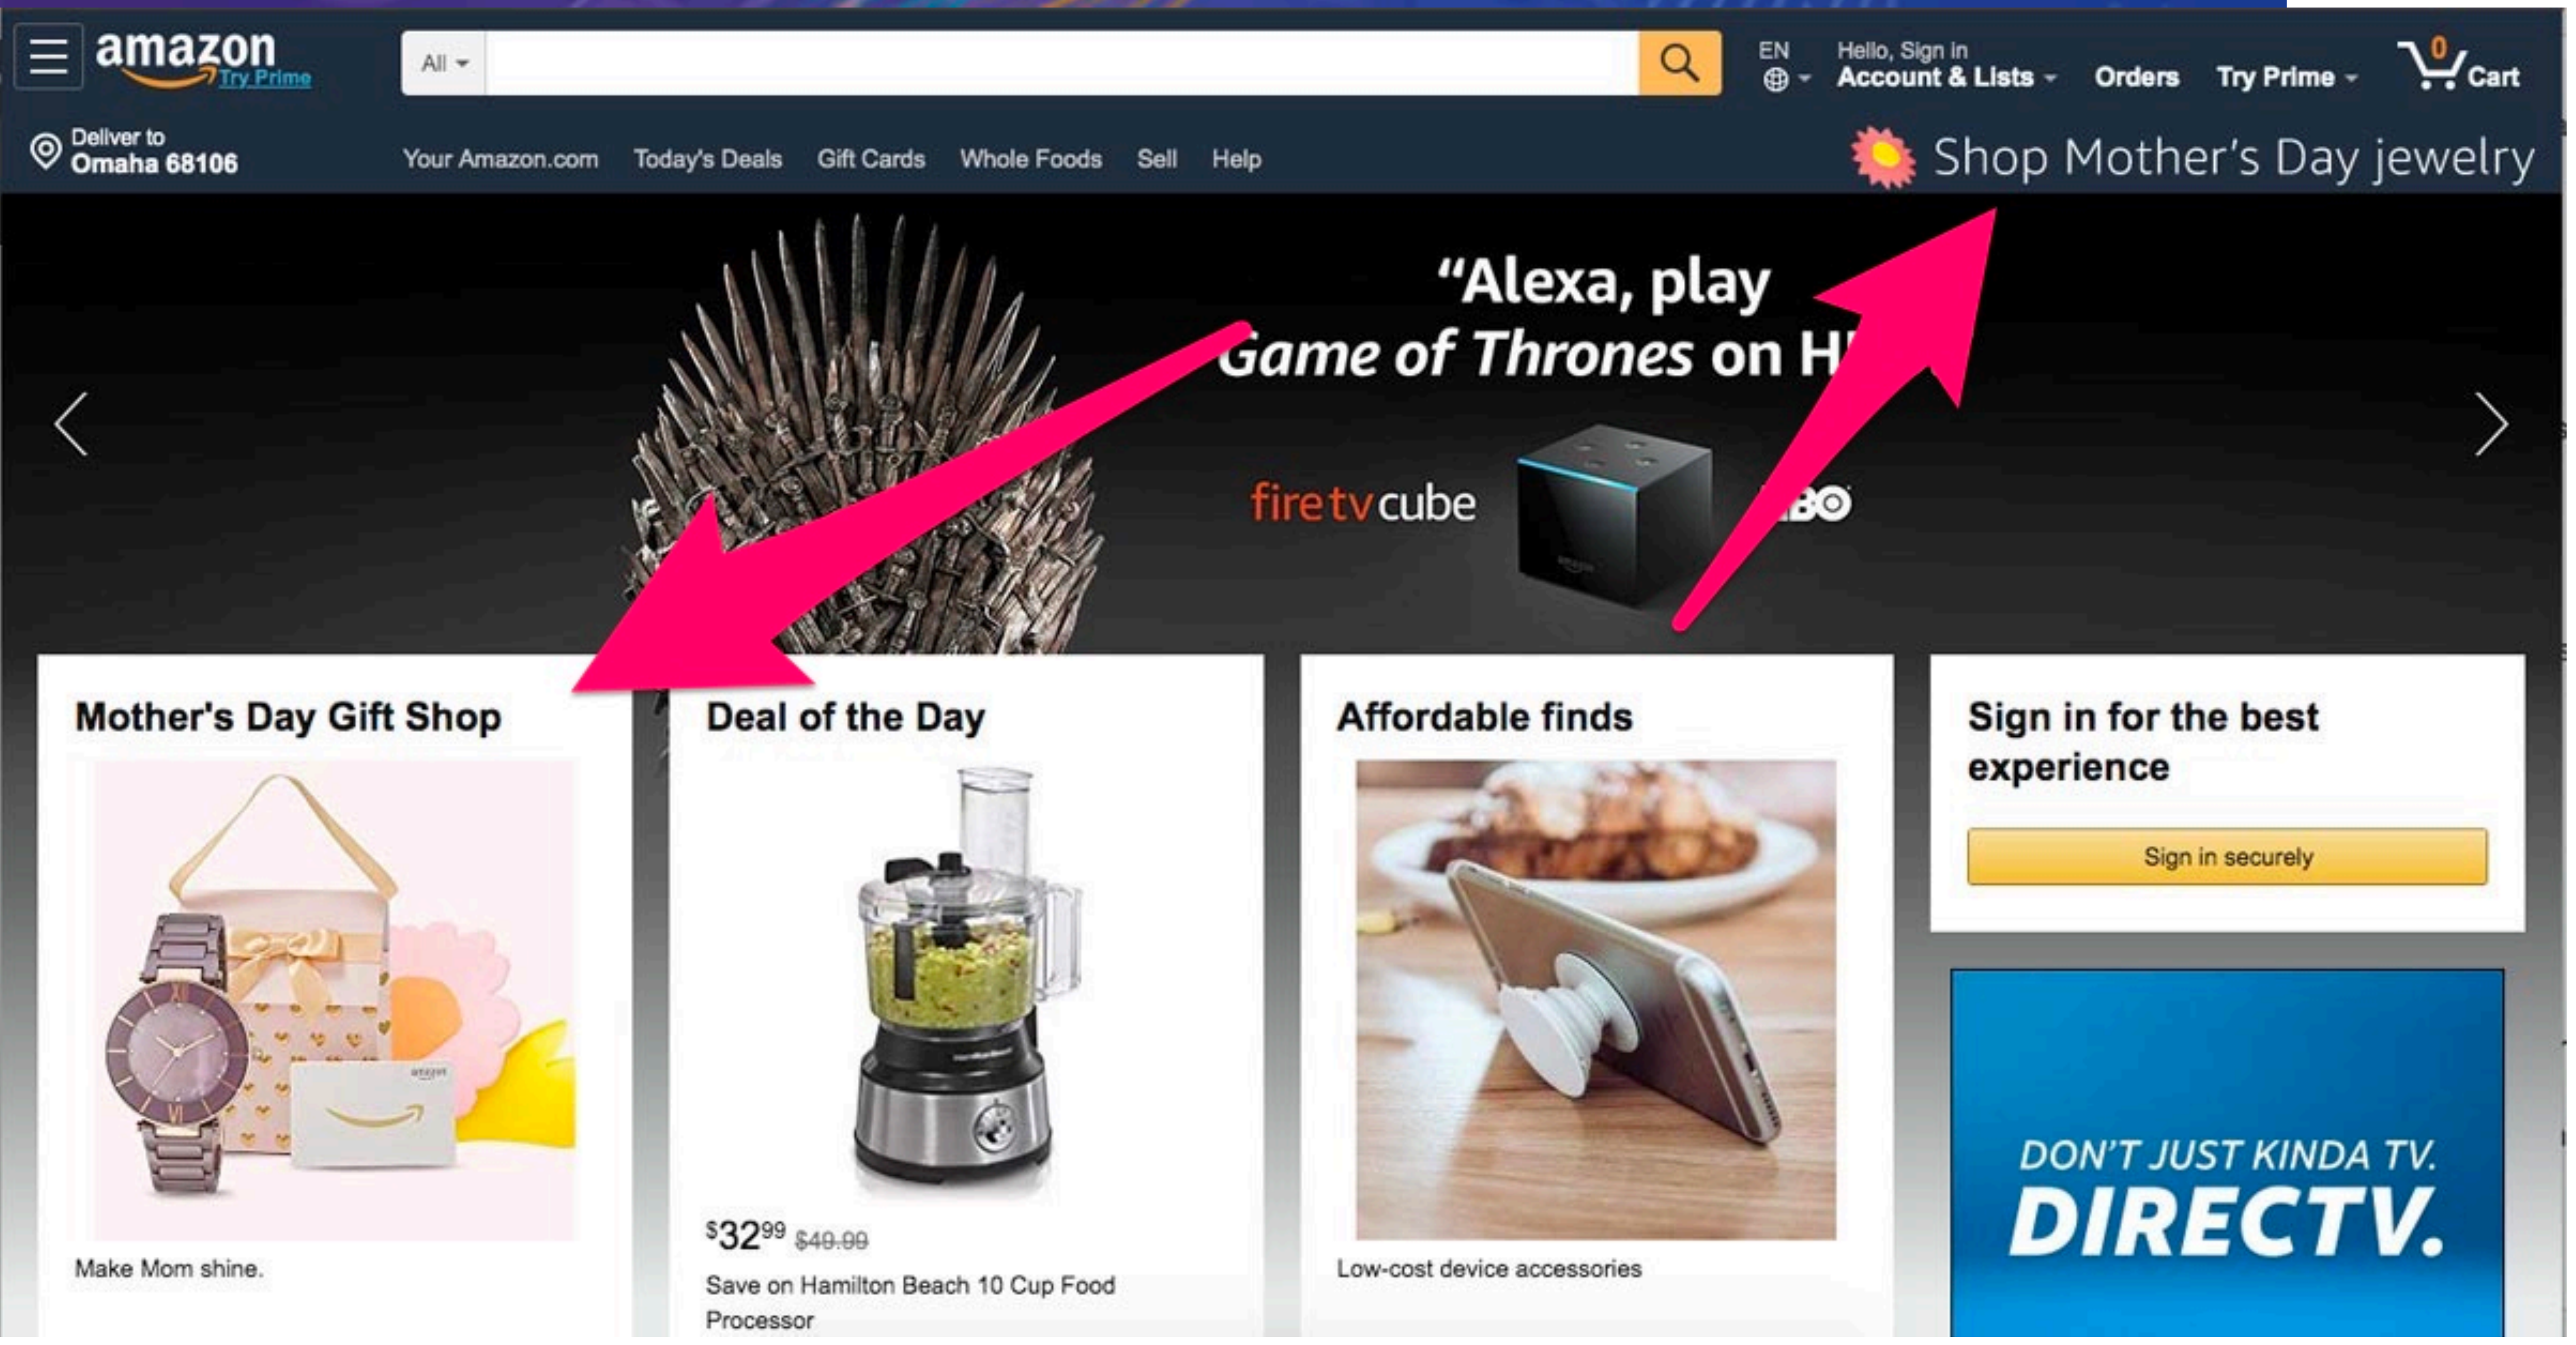

amazon
Try Prime
All
EN
Hello, Sign in
Account & Lists
Orders
Try Prime
0
Cart
Deliver to
Omaha 68106
Your Amazon.com
Today's Deals
Gift Cards
Whole Foods
Sell
Help
Shop Mother's Day jewelry
"Alexa, play
Game of Thrones on H
fire tv cube
Mother's Day Gift Shop
Make Mom shine.
Deal of the Day
$32.99 $49.99
Save on Hamilton Beach 10 Cup Food Processor
Affordable finds
Low-cost device accessories
Sign in for the best experience
Sign in securely
DON'T JUST KINDA TV.
DIRECTV.

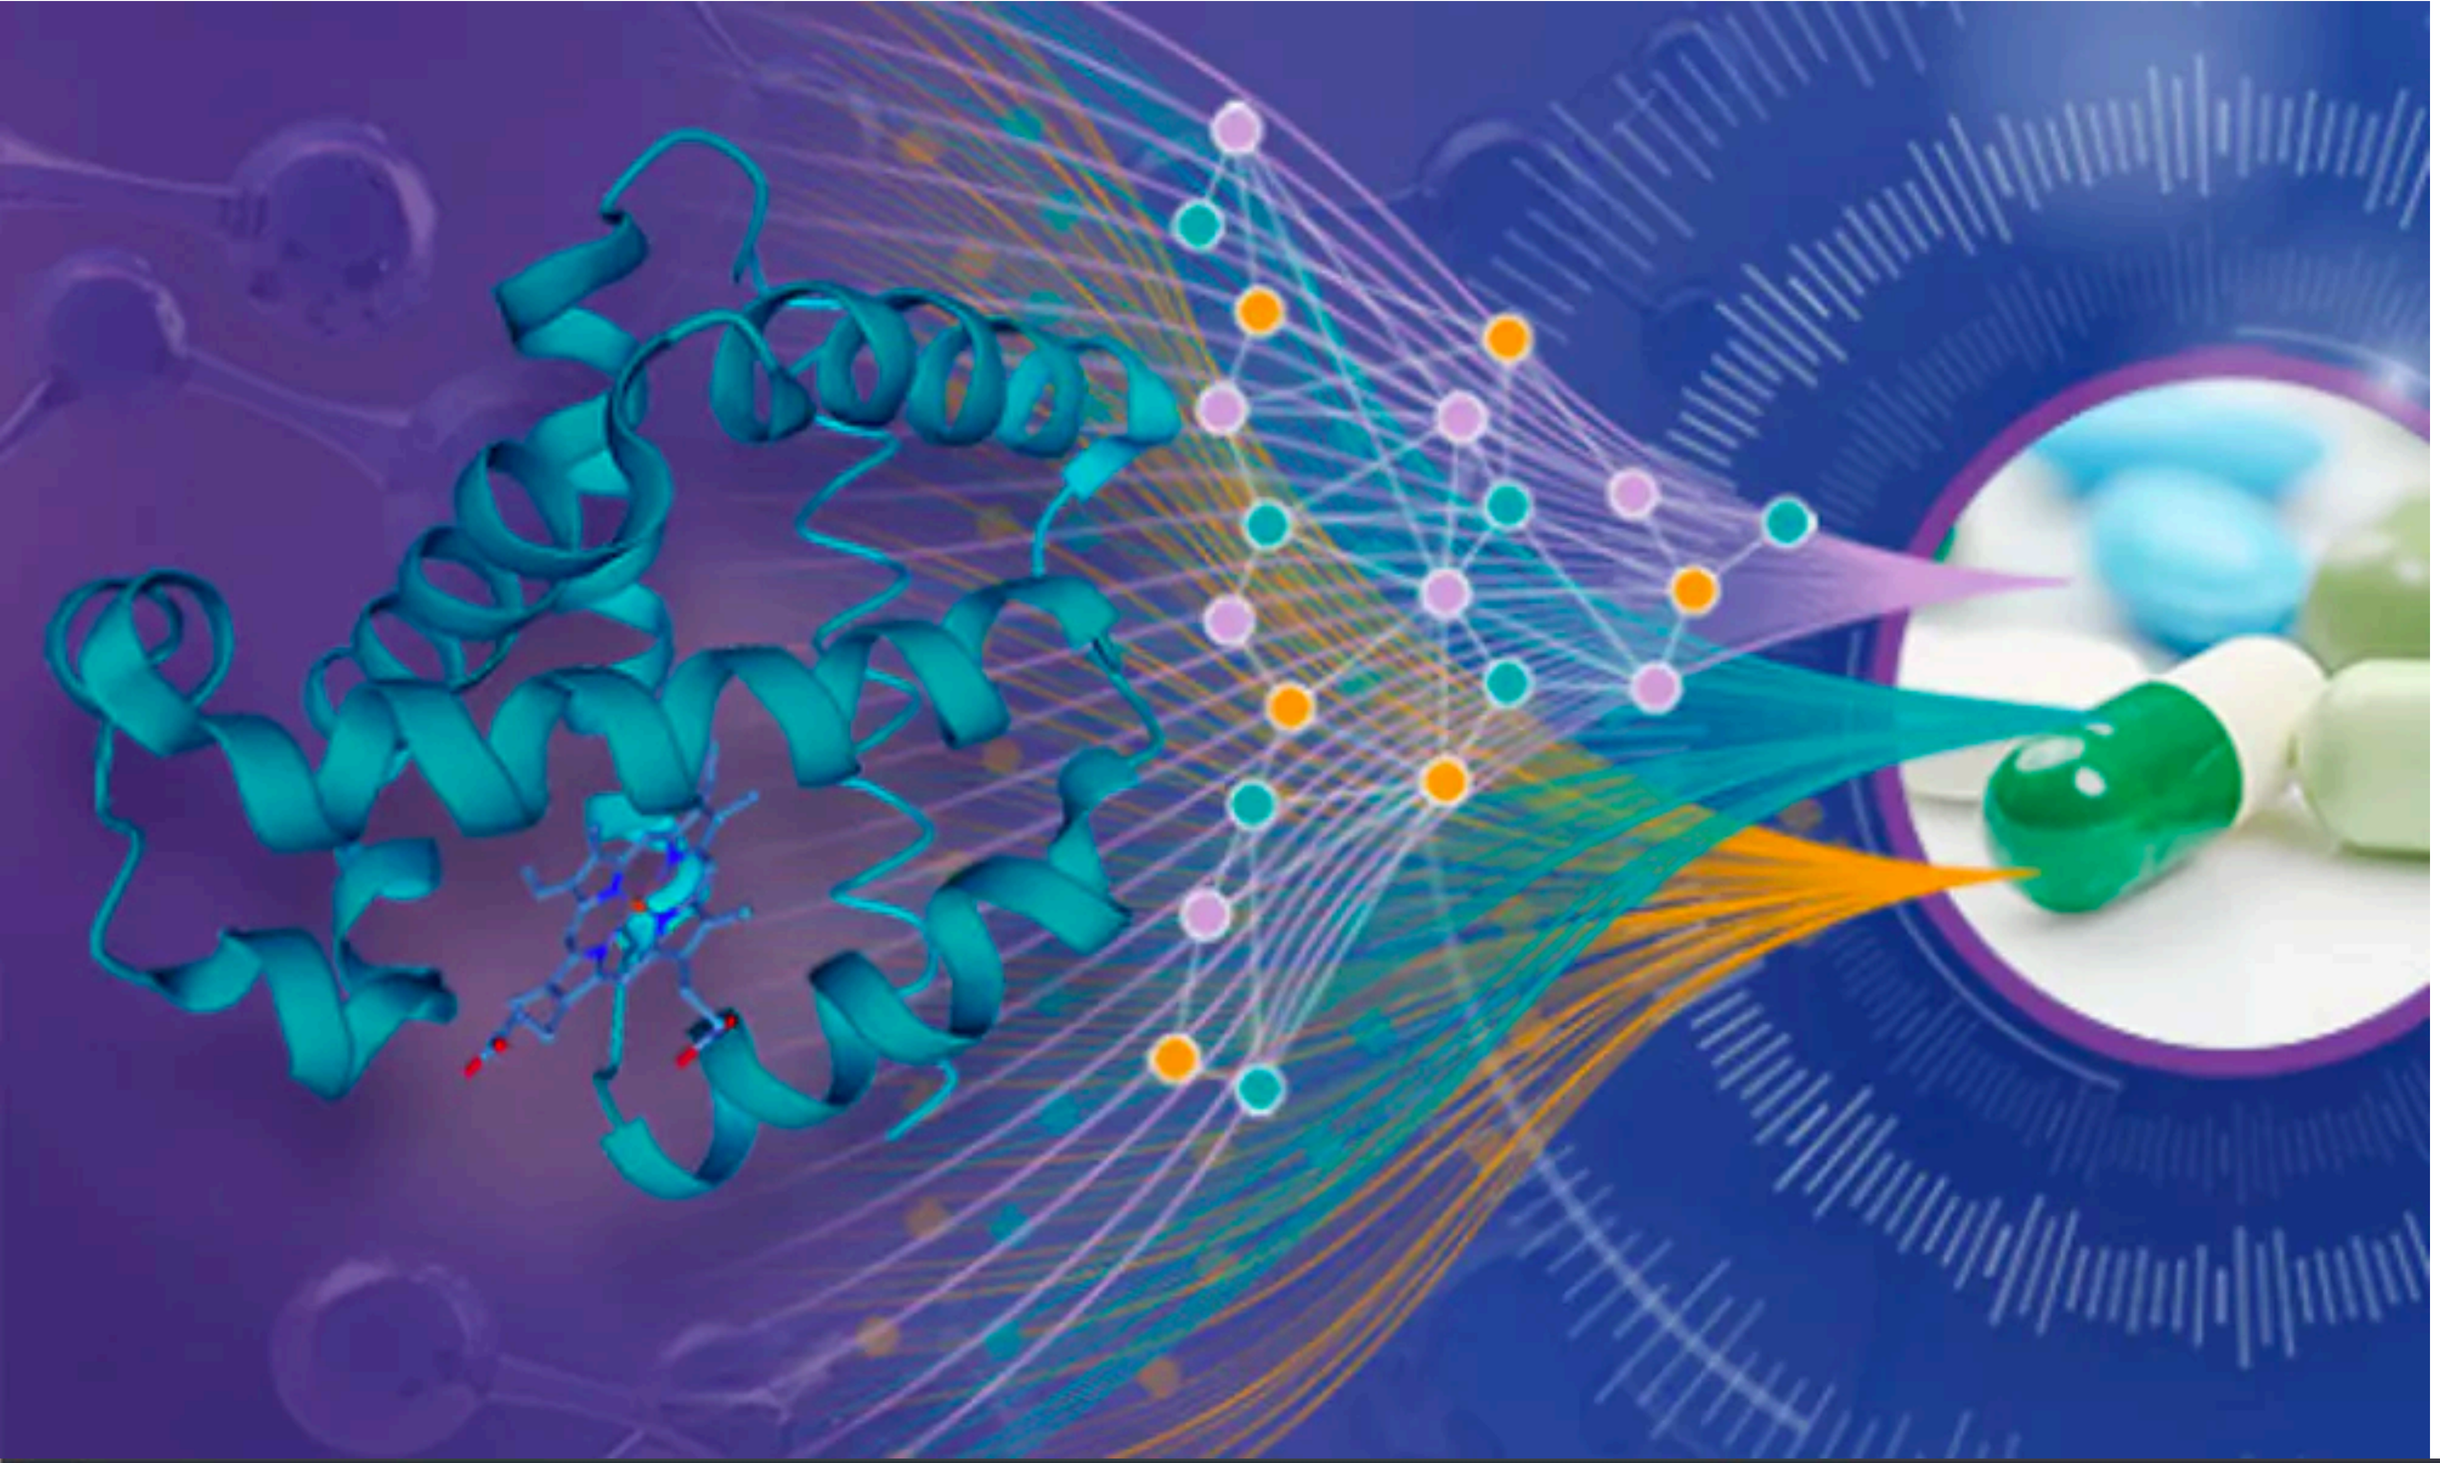

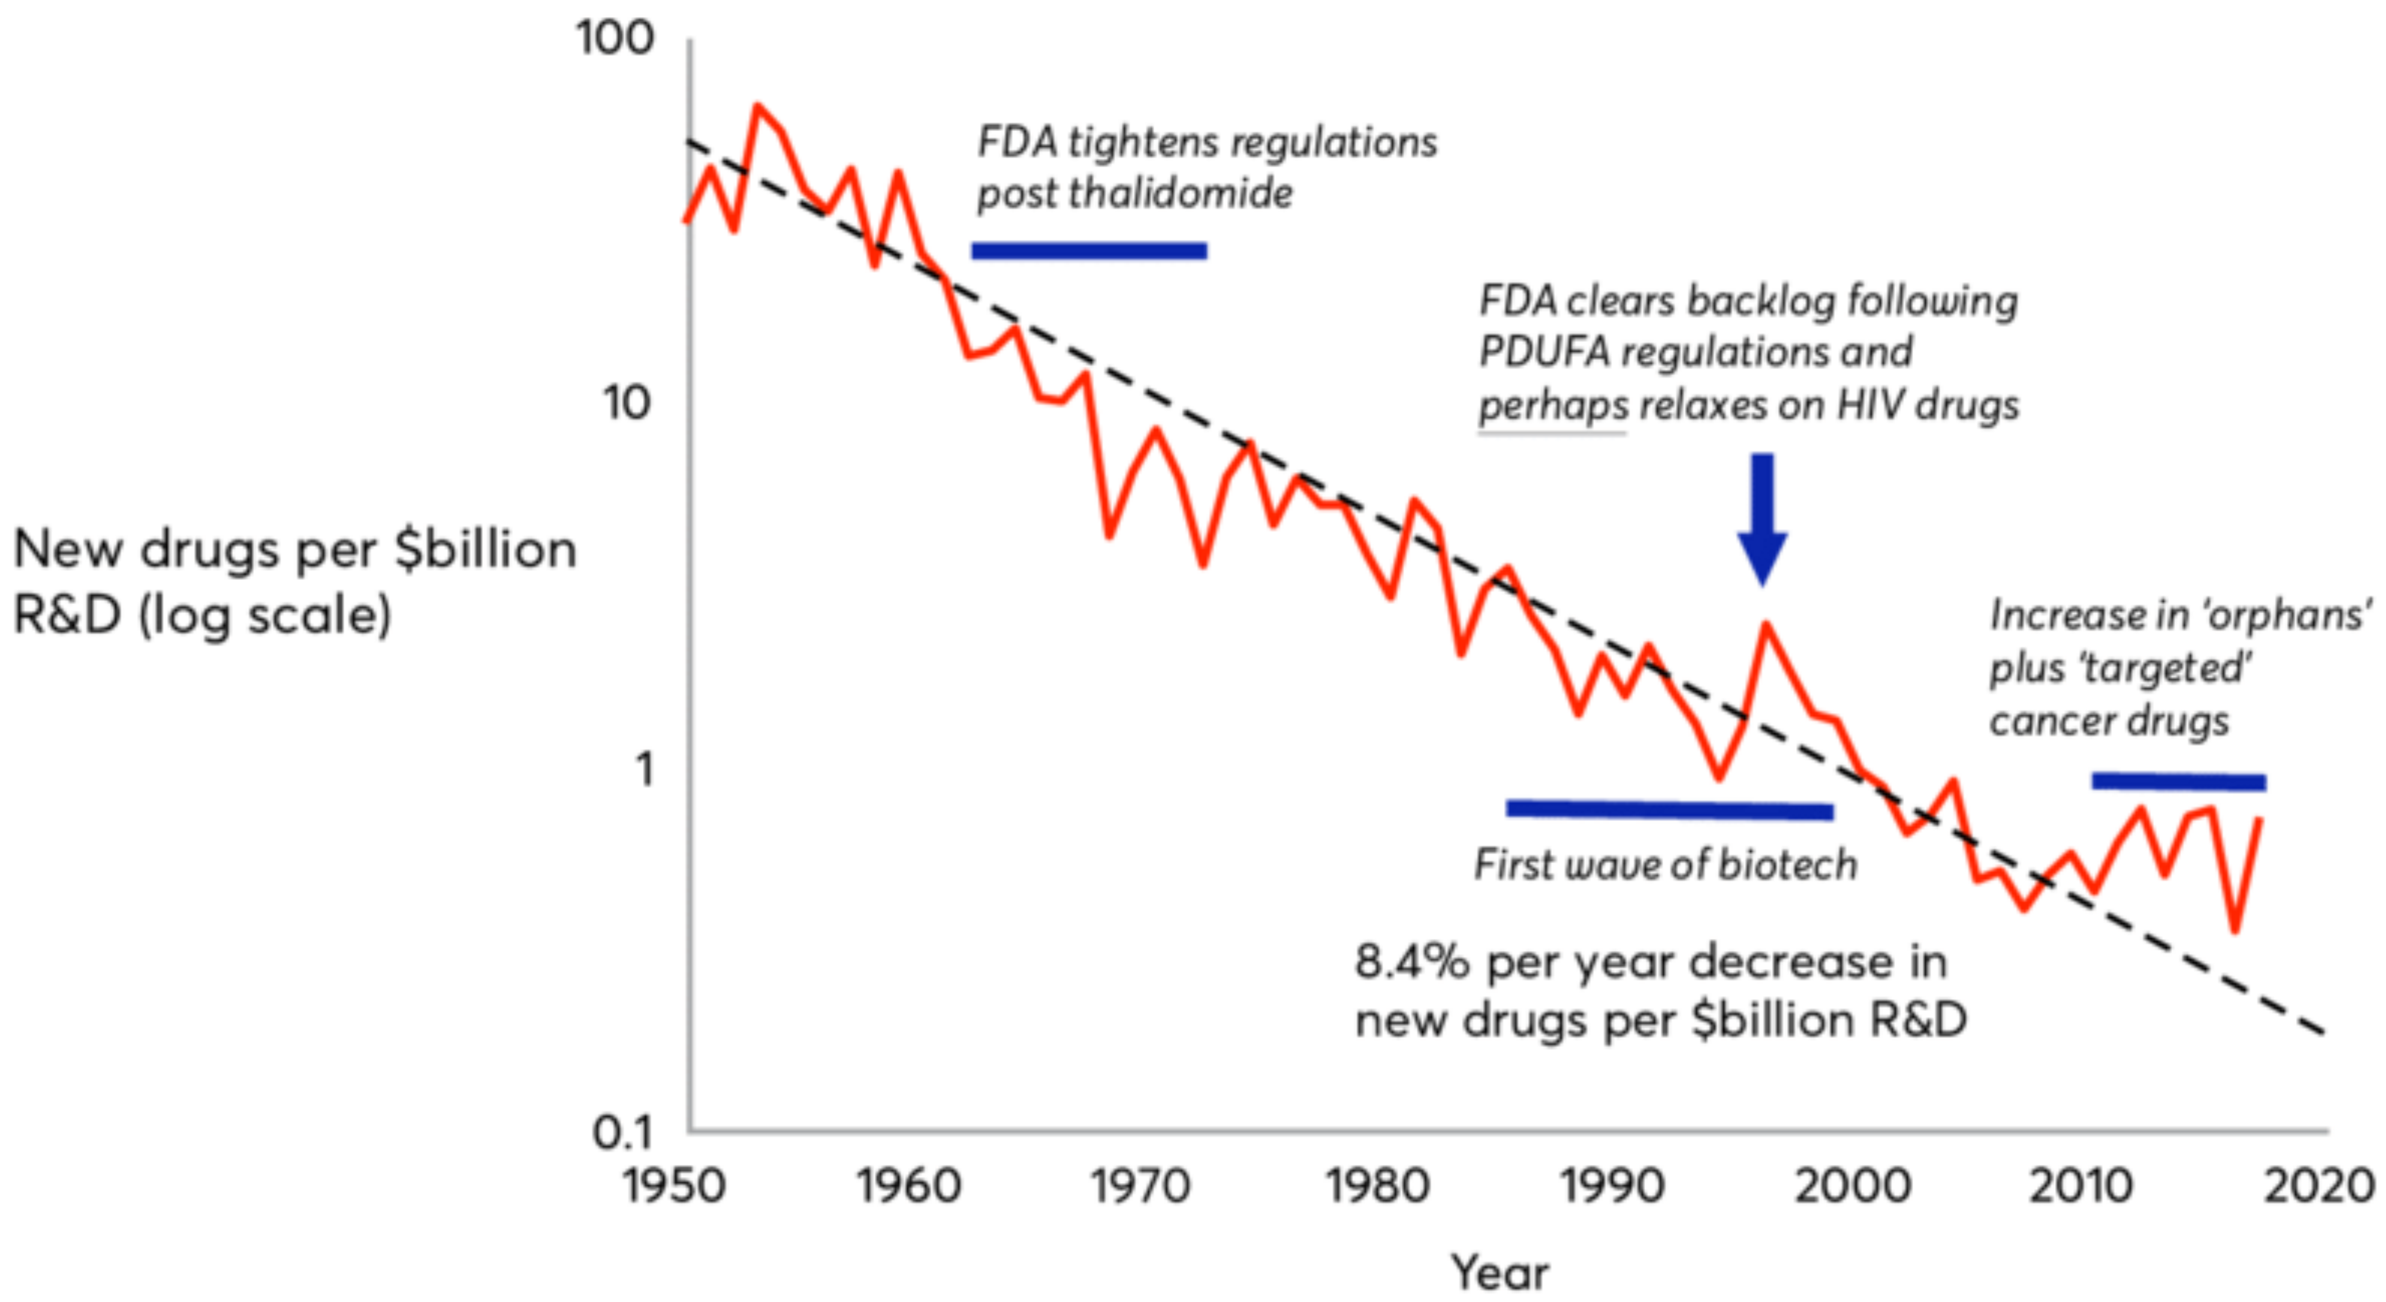

100
10
1
0.1
New drugs per $billion R&D (log scale)
FDA tightens regulations post thalidomide
FDA clears backlog following PDUFA regulations and perhaps relaxes on HIV drugs
Increase in 'orphans' plus 'targeted' cancer drugs
First wave of biotech
8.4% per year decrease in new drugs per $billion R&D
1950
1960
1970
1980
1990
2000
2010
2020
Year

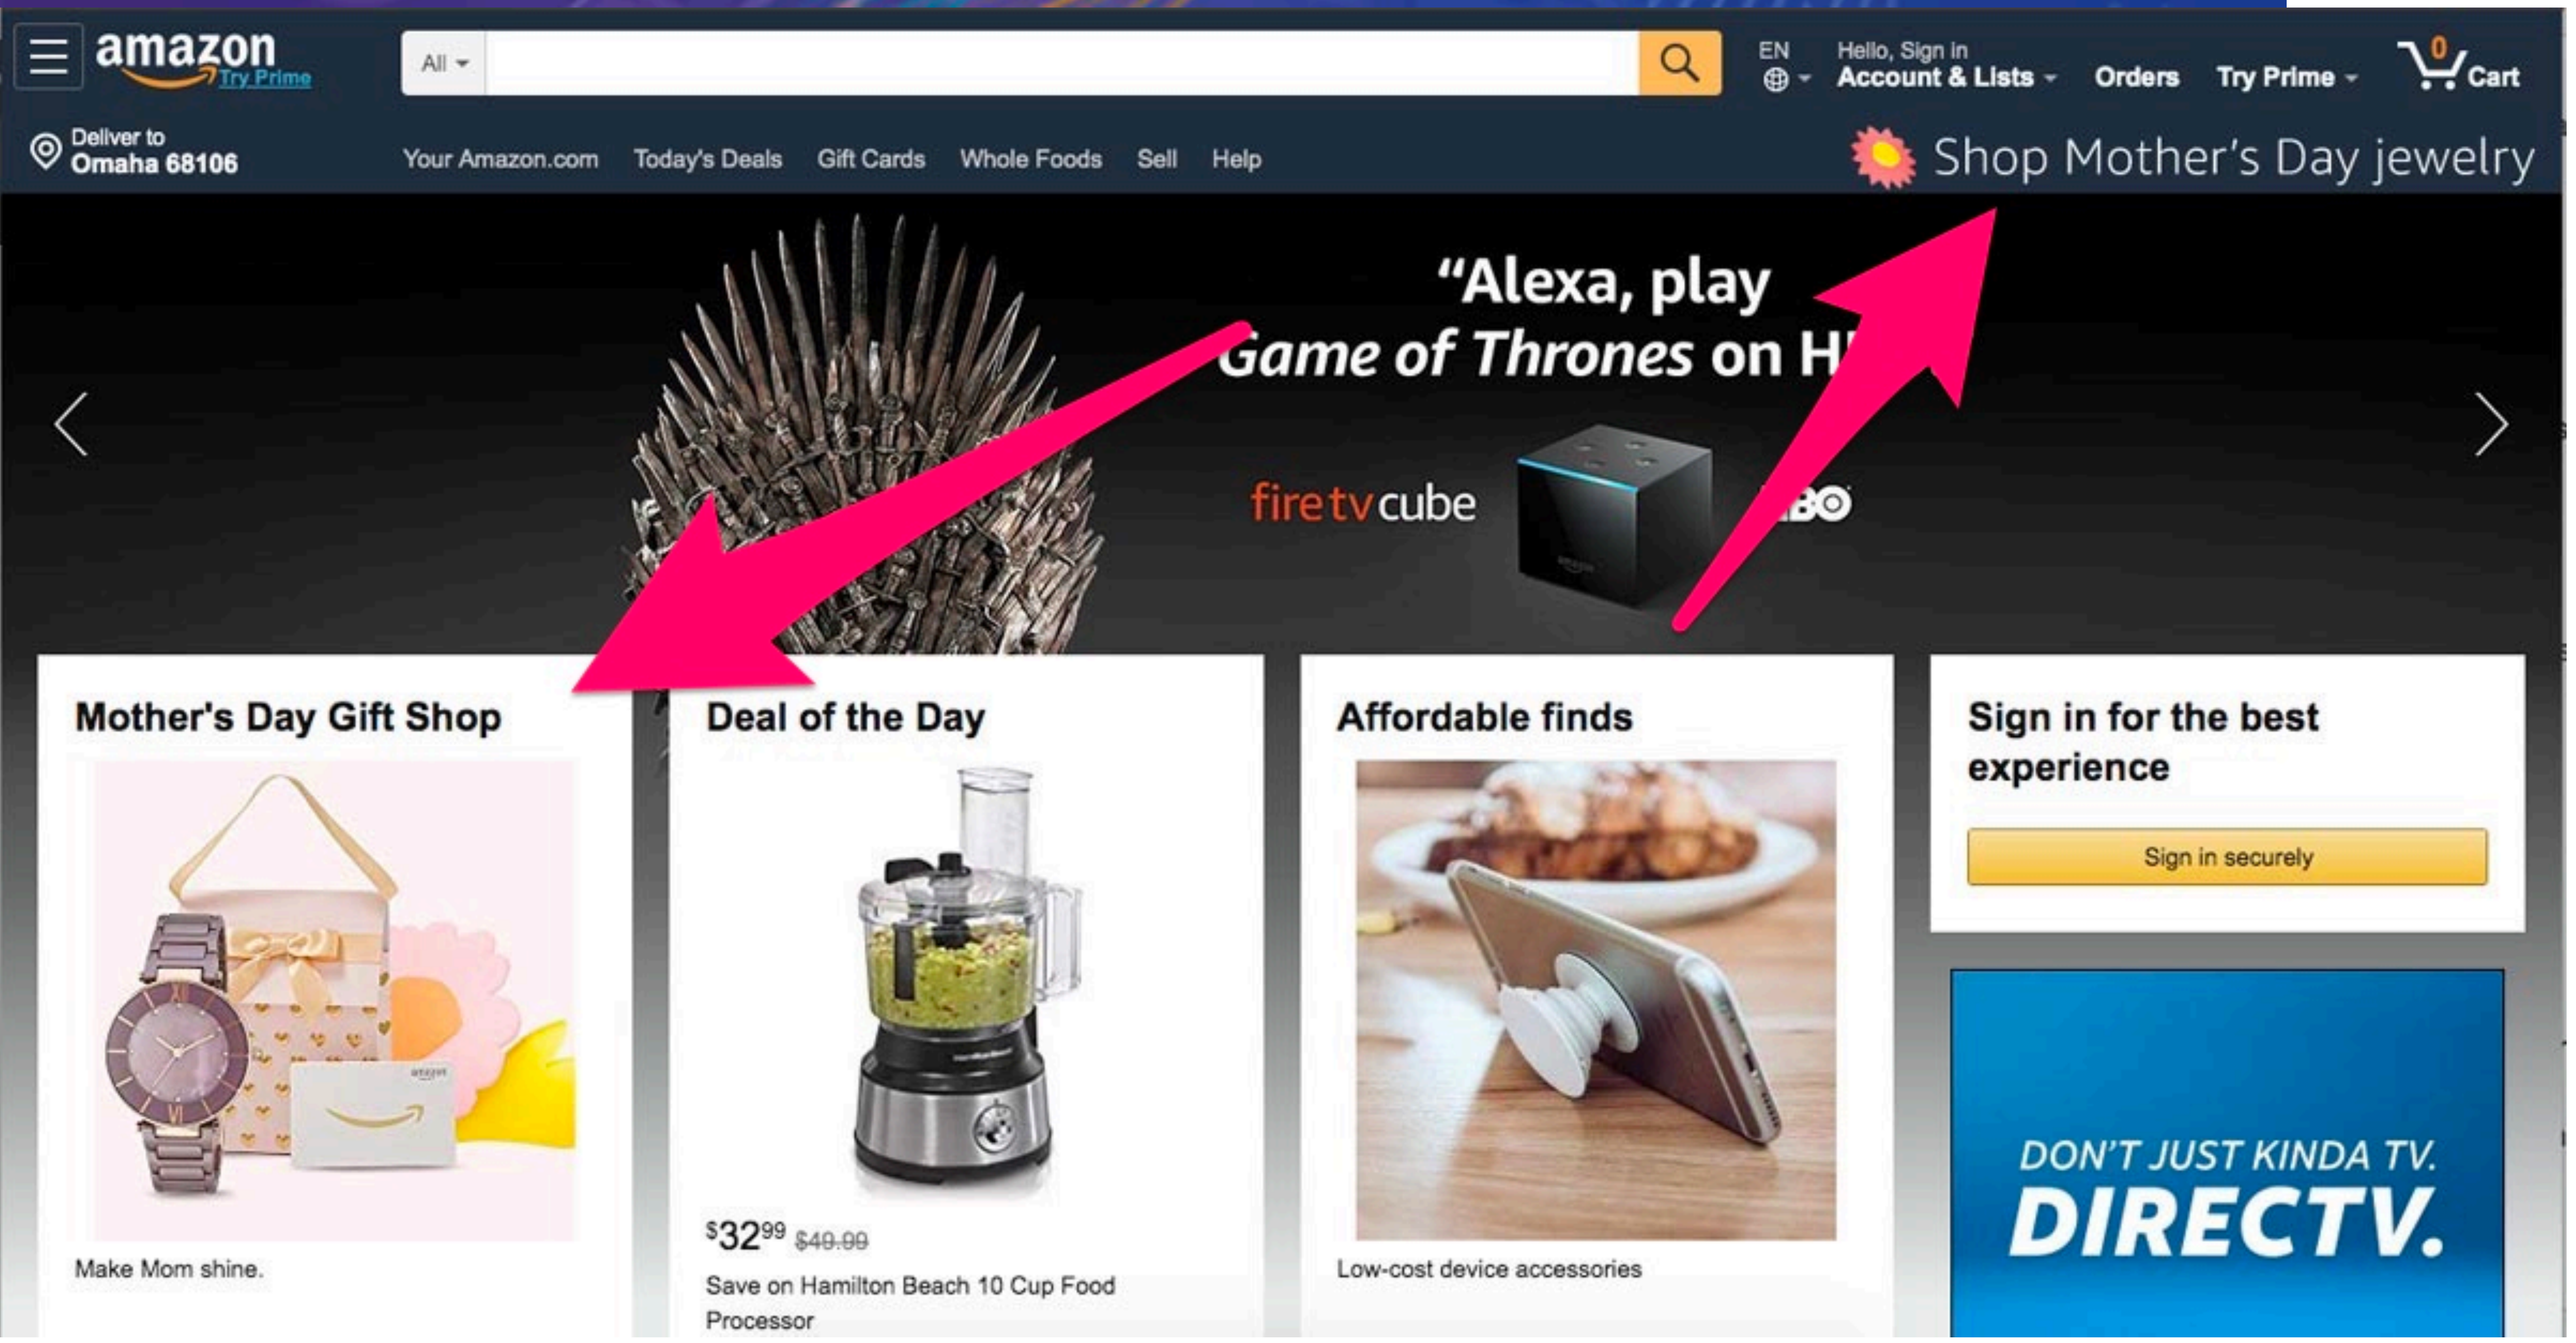

amazon
Try Prime
All
EN
Hello, Sign in
Account & Lists
Orders
Try Prime
0
Cart
Deliver to
Omaha 68106
Your Amazon.com
Today's Deals
Gift Cards
Whole Foods
Sell
Help
Shop Mother's Day jewelry
"Alexa, play
Game of Thrones on H
fire tv cube
Mother's Day Gift Shop
Make Mom shine.
Deal of the Day
$32.99 $49.99
Save on Hamilton Beach 10 Cup Food Processor
Affordable finds
Low-cost device accessories
Sign in for the best experience
Sign in securely
DON'T JUST KINDA TV.
DIRECTV.

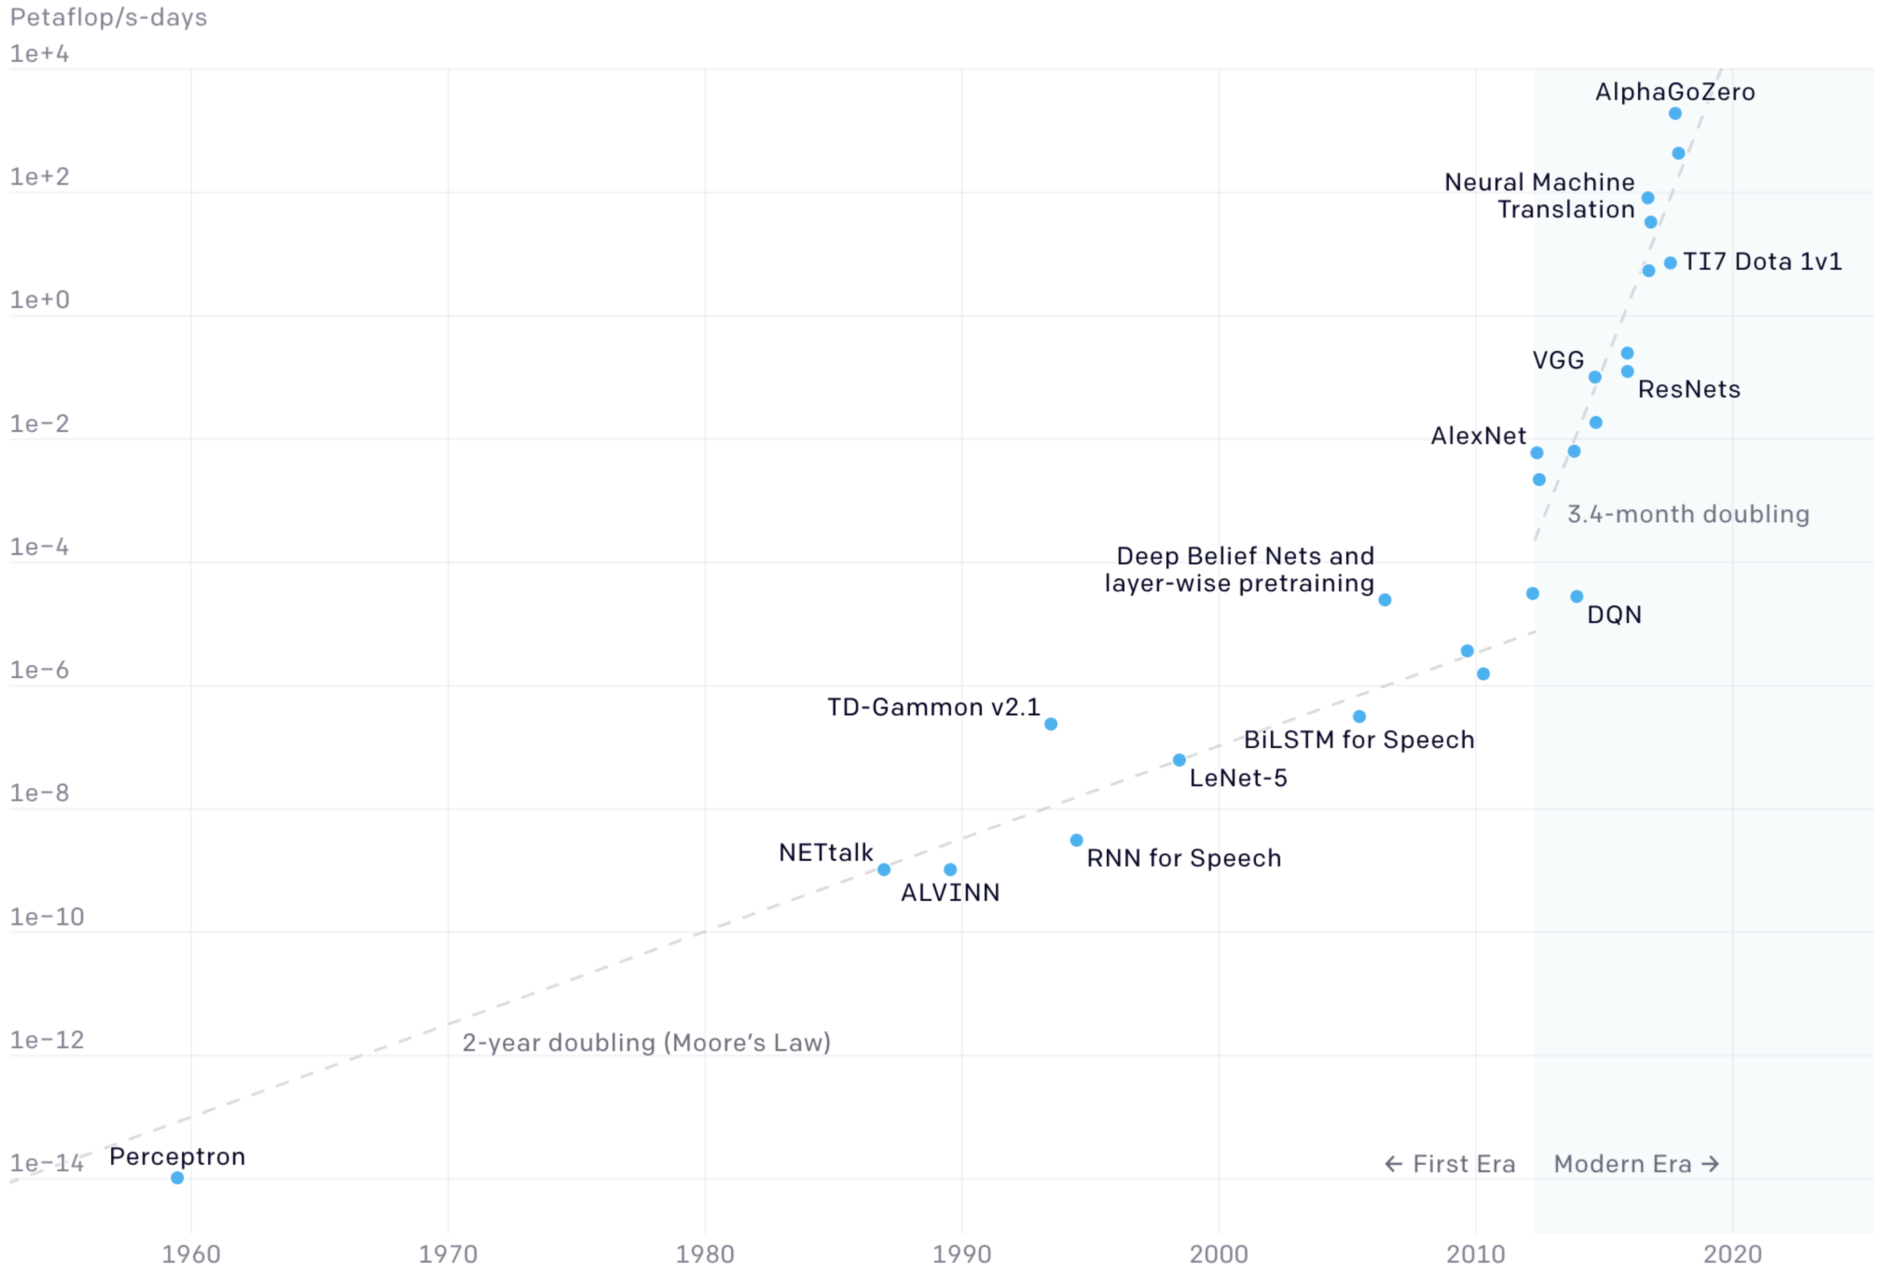

Petaflop/s-days
1e+4
1e+2
1e+0
1e-2
1e-4
1e-6
1e-8
1e-10
1e-12
1e-14
AlphaGoZero
Neural Machine Translation
TI7 Dota 1v1
VGG
ResNets
AlexNet
3.4-month doubling
Deep Belief Nets and layer-wise pretraining
DQN
TD-Gammon v2.1
BiLSTM for Speech
LeNet-5
NETtalk
ALVINN
RNN for Speech
2-year doubling (Moore's Law)
Perceptron
← First Era
Modern Era →
1960
1970
1980
1990
2000
2010
2020

Hyper-parameter optimizers

model_selection.GridSearchCV(estimator, ...)	Exhaustive search over specified parameter values for an estimator.
model_selection.HalvingGridSearchCV(...[, ...])	Search over specified parameter values with successive halving.
model_selection.ParameterGrid(param_grid)	Grid of parameters with a discrete number of values for each.
model_selection.ParameterSampler(...[, ...])	Generator on parameters sampled from given distributions.
model_selection.RandomizedSearchCV(...[, ...])	Randomized search on hyper parameters.
model_selection.HalvingRandomSearchCV(...[, ...])	Randomized search on hyper parameters.

Bayesian Optimization for a Better Dessert

Greg Kochanski, Daniel Golovin, John Karro, Benjamin Solnik, Subhodeep Moitra, and D. Sculley
{gpk, dgg, karro, bsolnik, smoitra, dsculley}@google.com; Google Brain Team

Abstract

We present a case study on applying Bayesian Optimization to a complex real-world system; our challenge was to optimize chocolate chip cookies. The process was a mixed-initiative system where both human chefs, human raters, and a machine optimizer participated in 144 experiments. This process resulted in highly rated cookies that deviated from expectations in some surprising ways – much less sugar in California, and cayenne in Pittsburgh. Our experience highlights the importance of incorporating domain expertise and the value of transfer learning approaches.

THE ACTIVE LEARNING LOOP

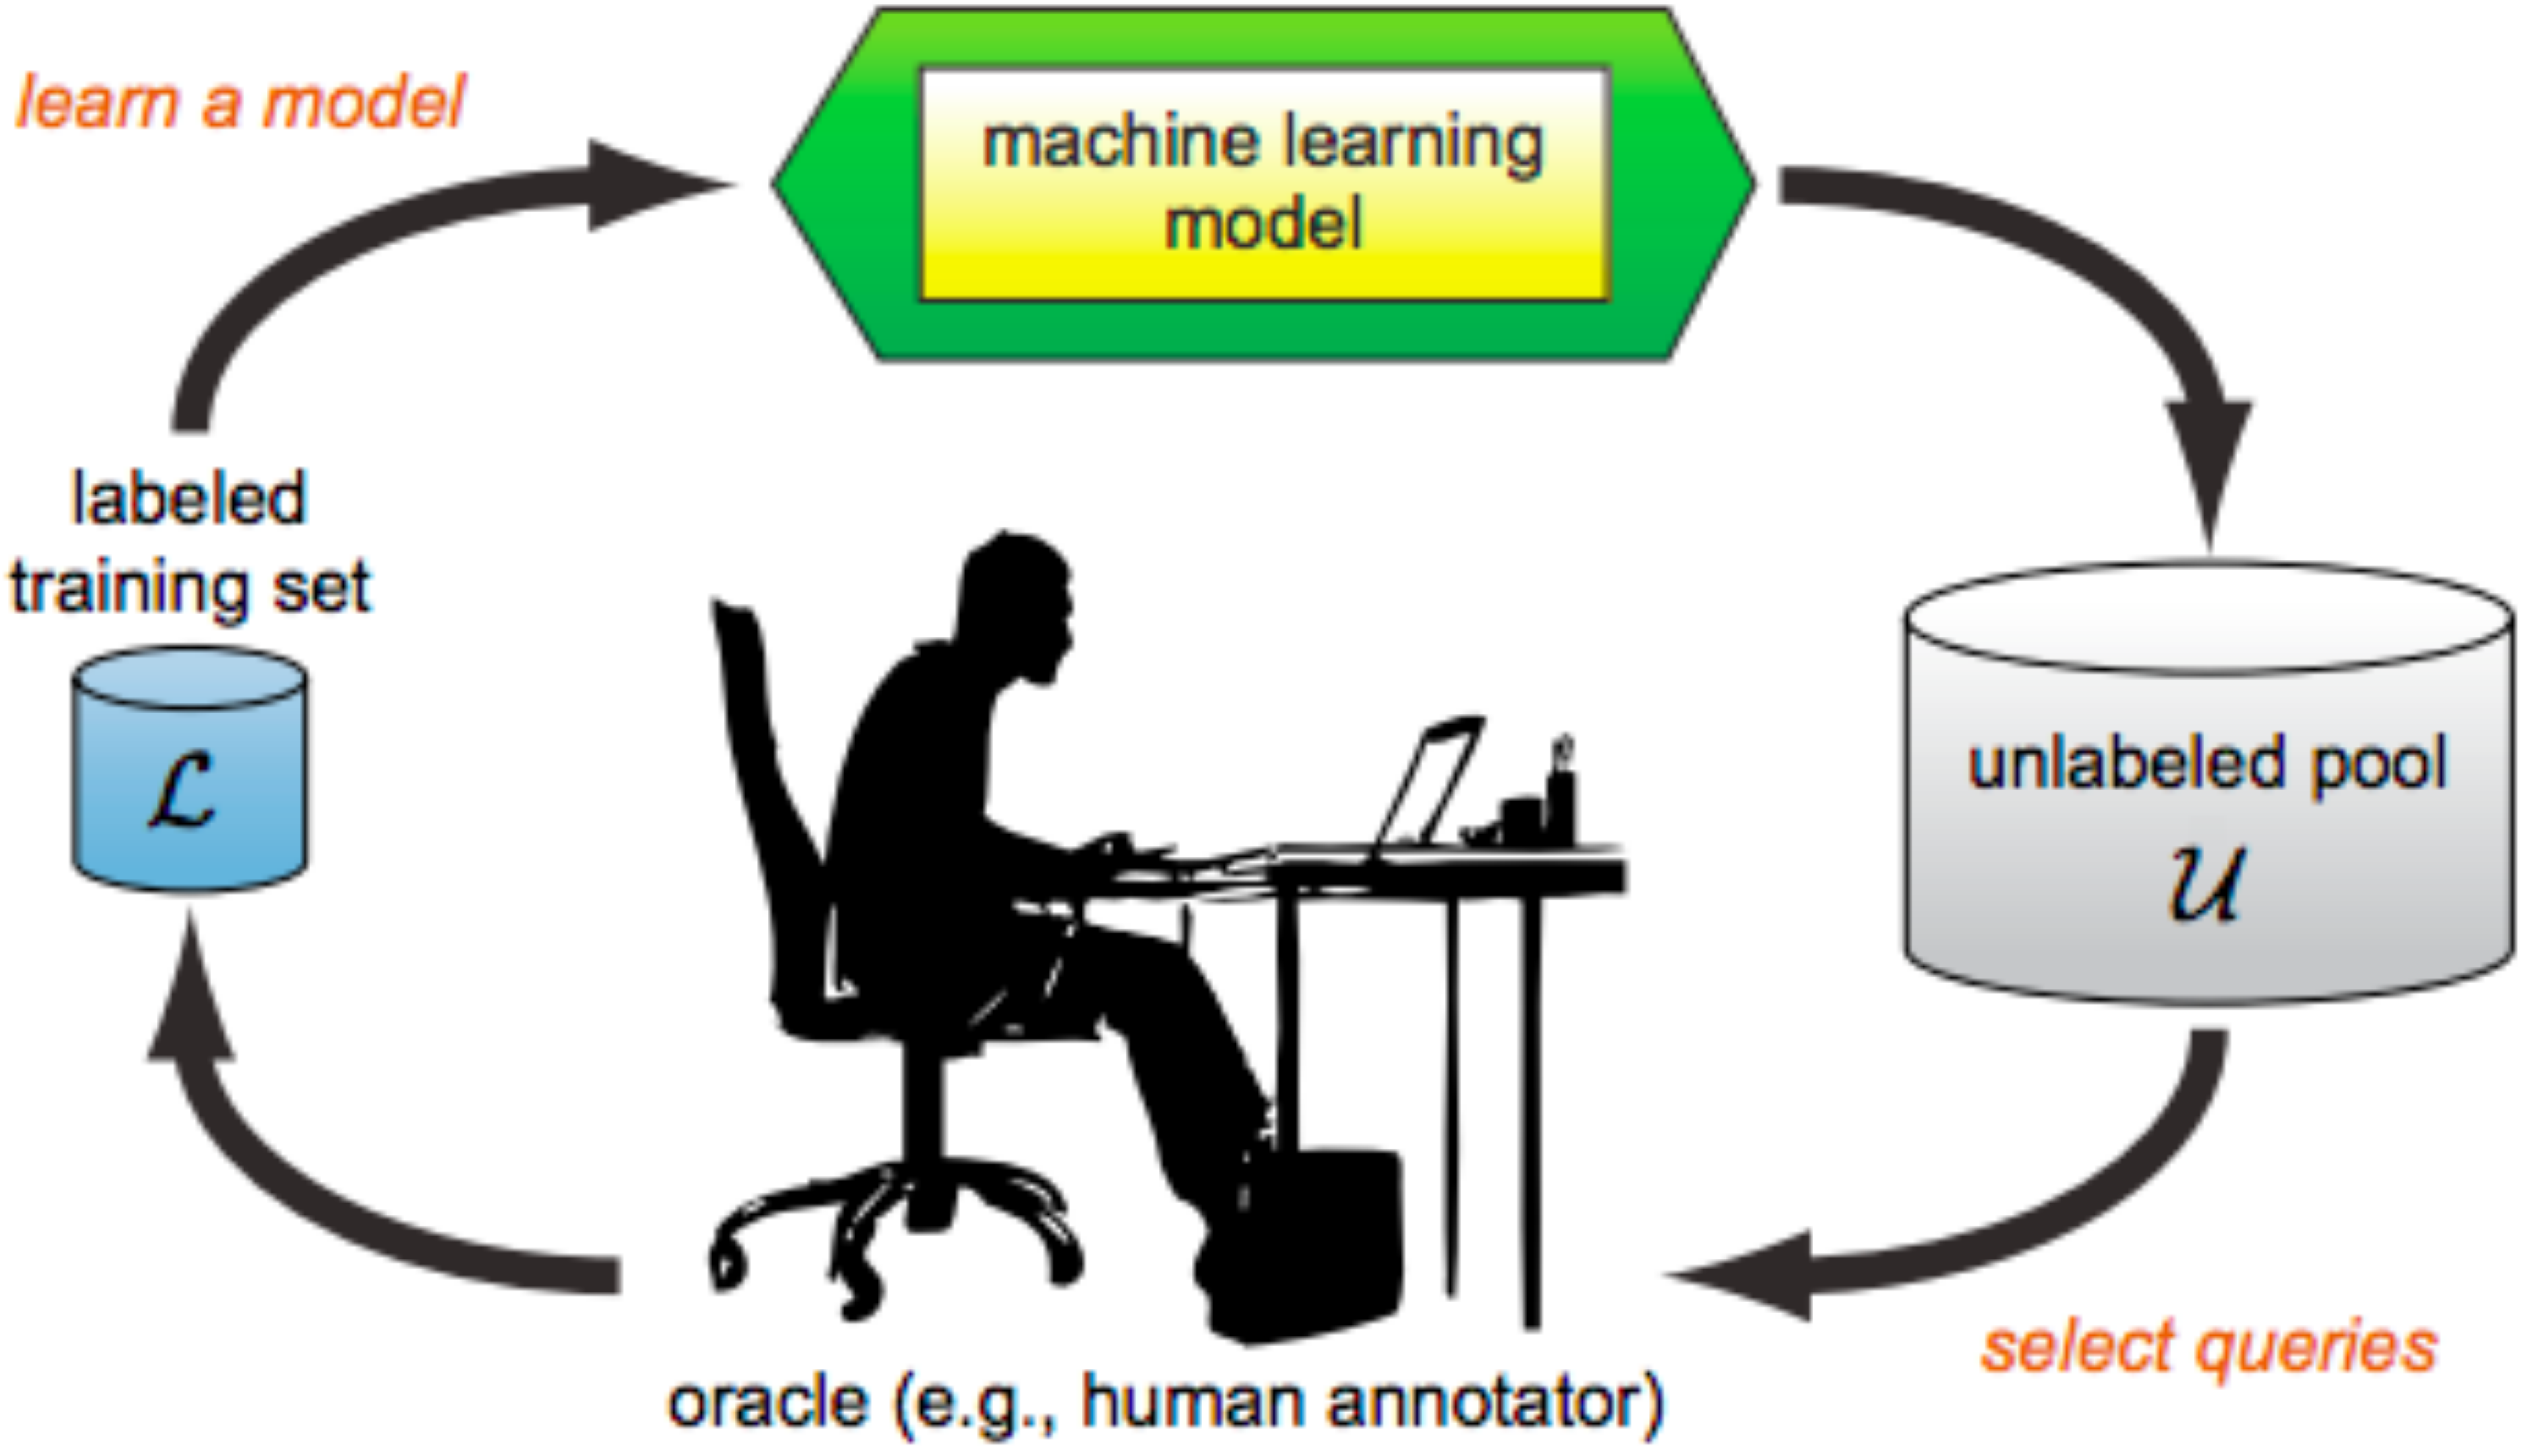

TWENTY QUESTIONS FOR OPTIMIZATION

TO BRING OR NOT TO BRING AN UMBRELLA

Given: chance of rain (outcome y) is equal to p

Question: whether or not to bring an umbrella to school (action a)

TO BRING OR NOT TO BRING AN UMBRELLA

Given: chance of rain (outcome y) is equal to p

Question: whether or not to bring an umbrella to school (action a)

	rain	no rain
umbrella	2	1
no umbrella	10	0

FANTASIZING ABOUT OUTCOMES

$\boxed{a_i}$

FANTASIZING ABOUT OUTCOMES

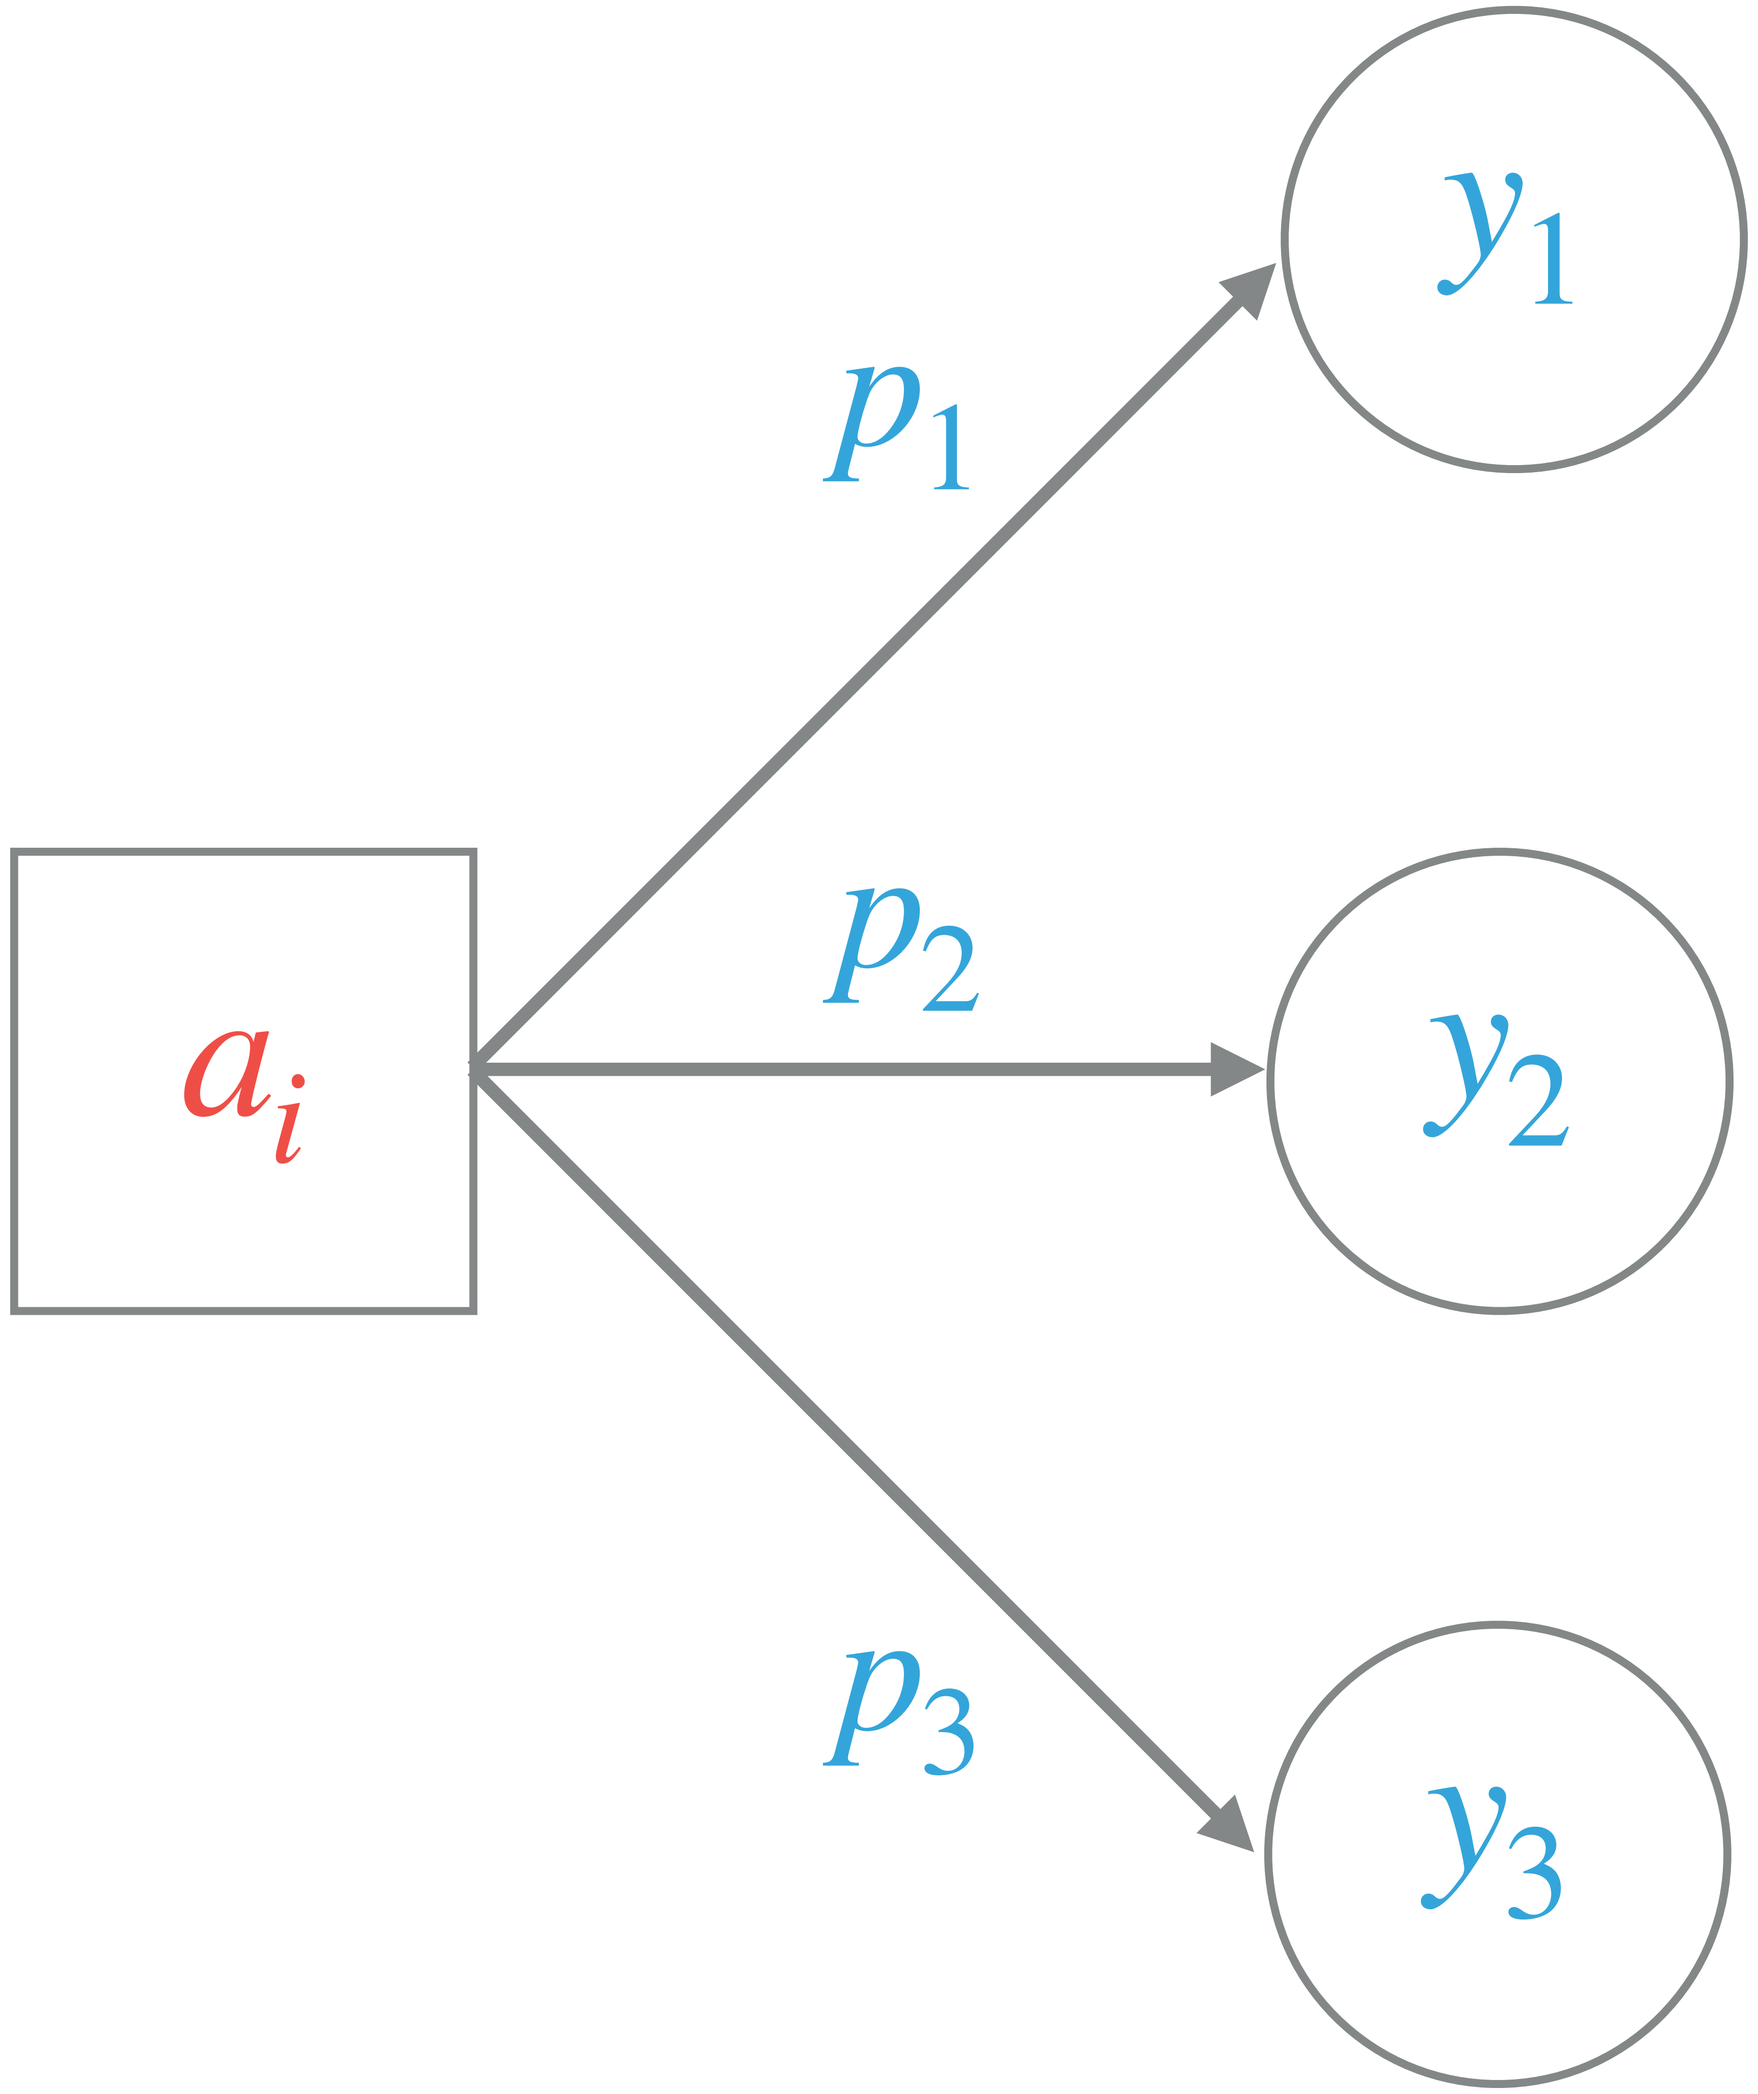

FANTASIZING ABOUT OUTCOMES

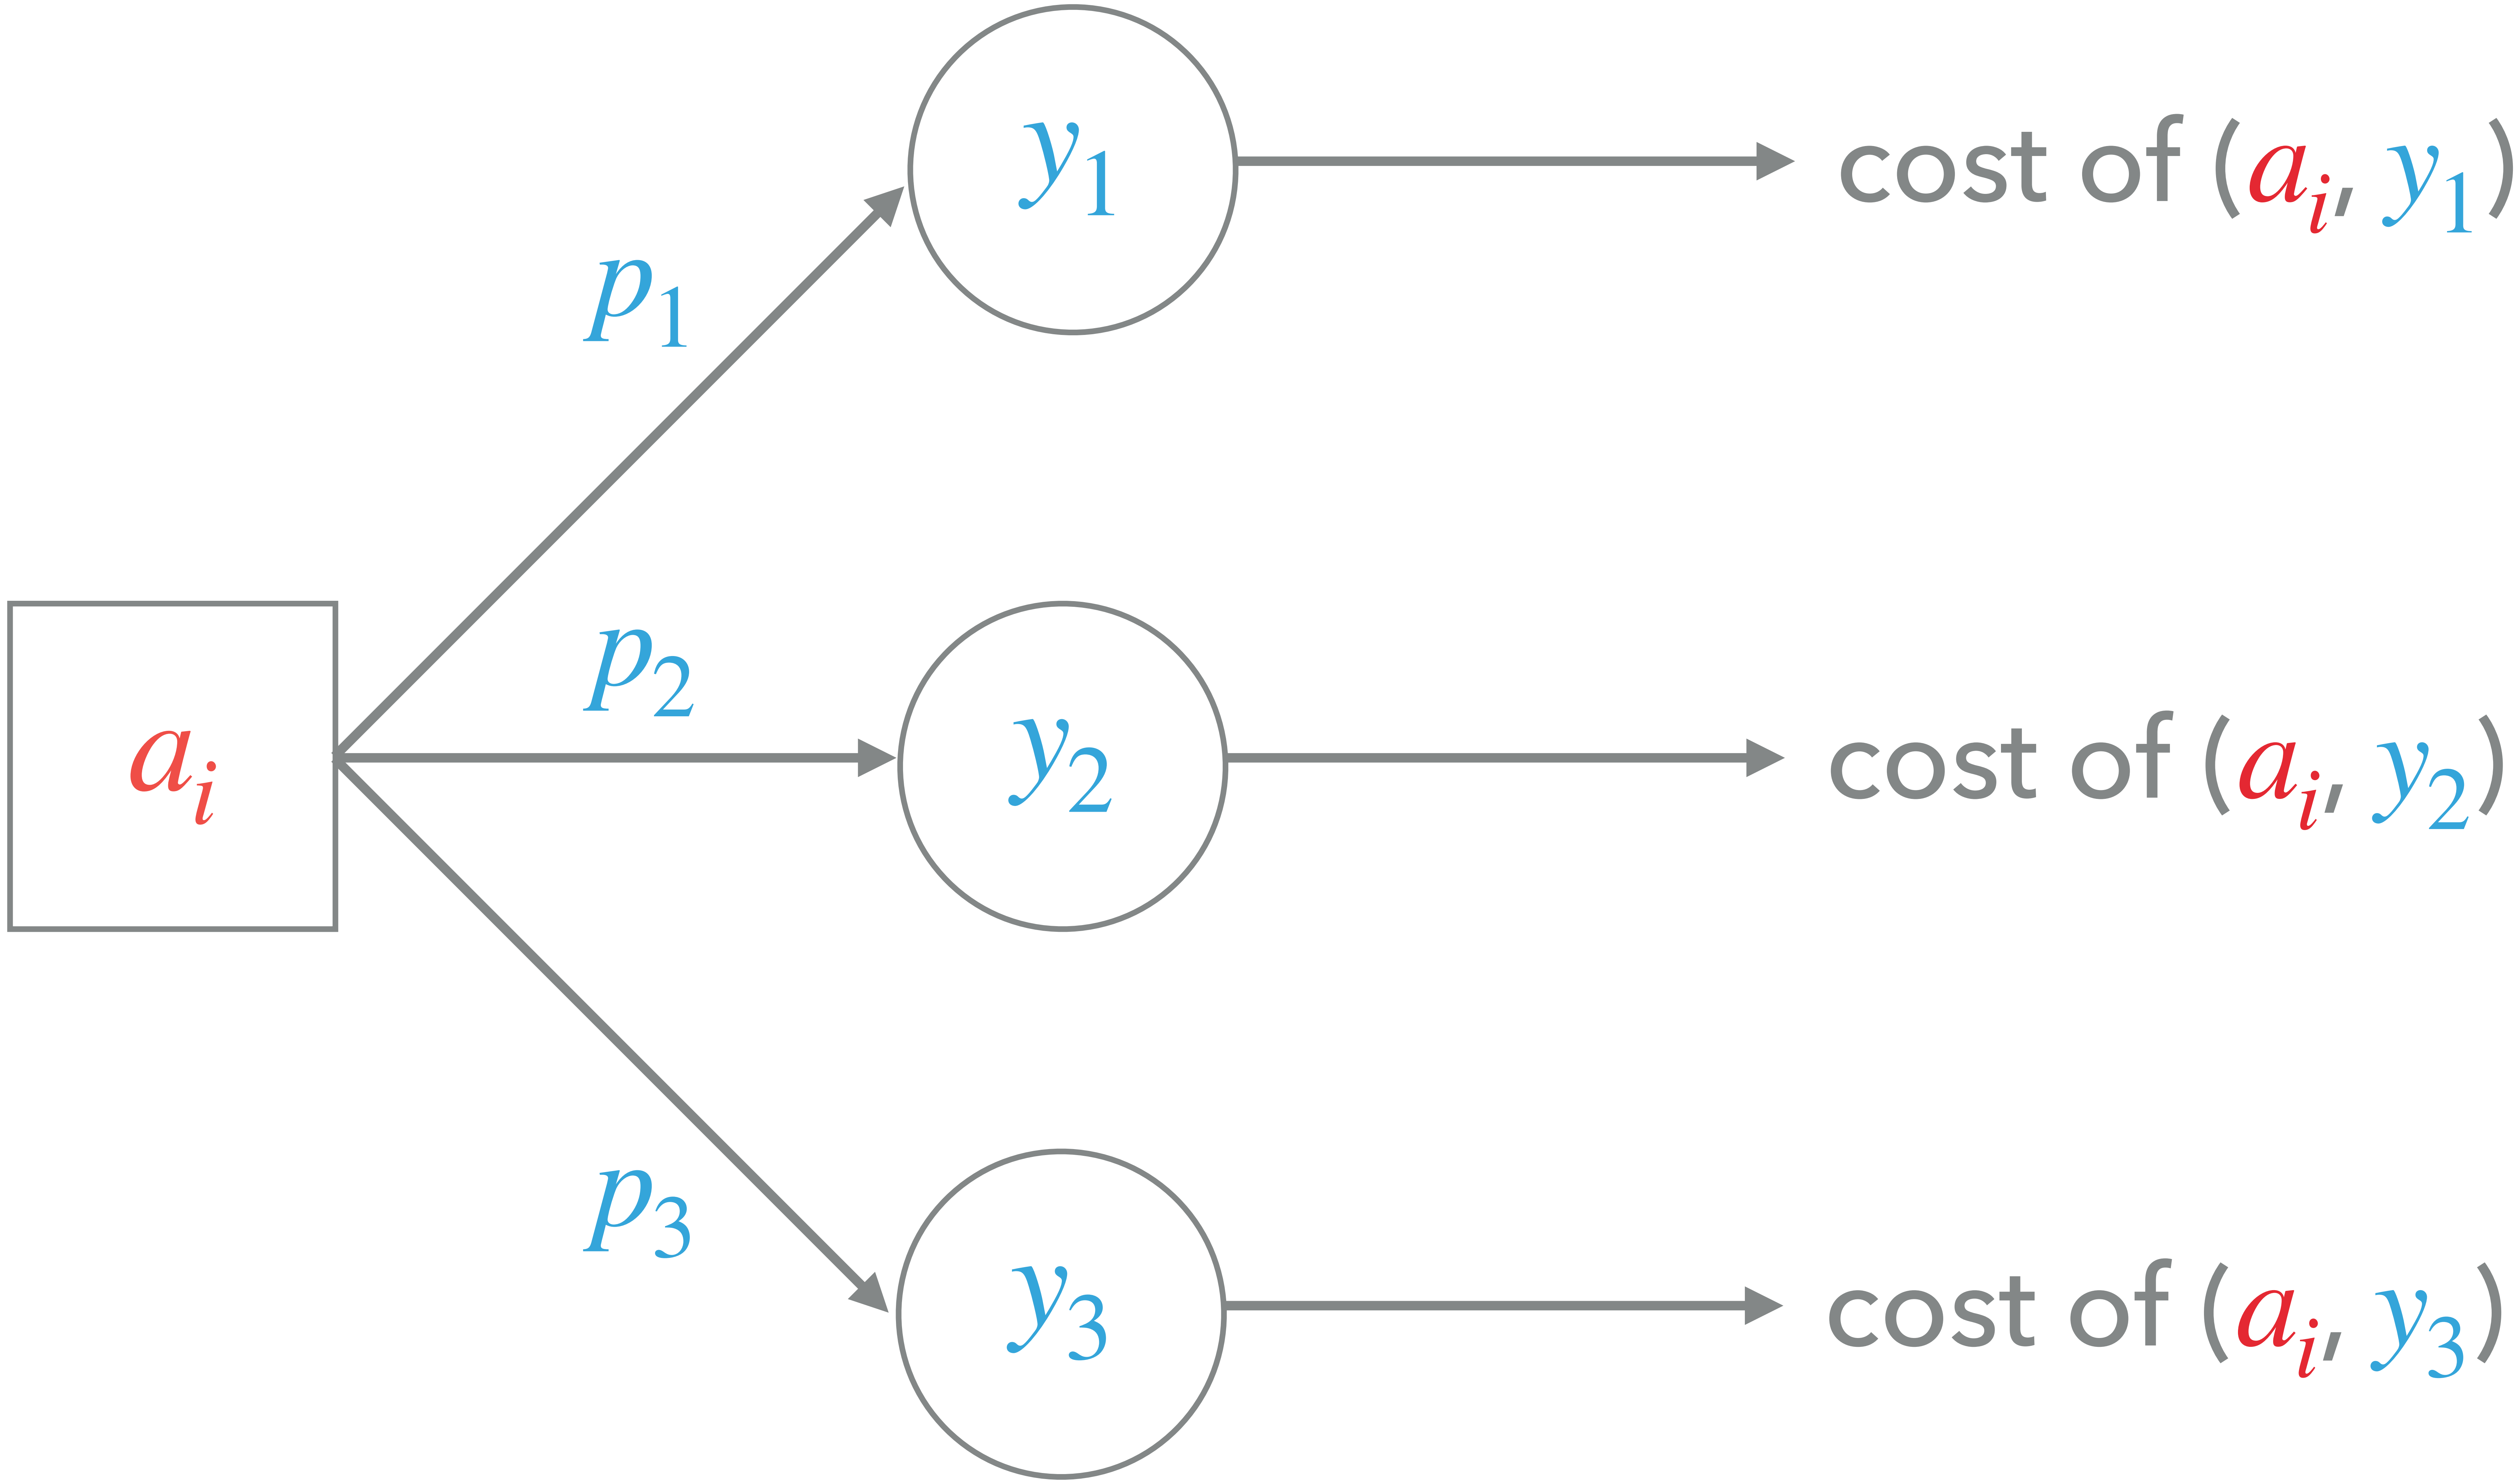

FANTASIZING ABOUT OUTCOMES

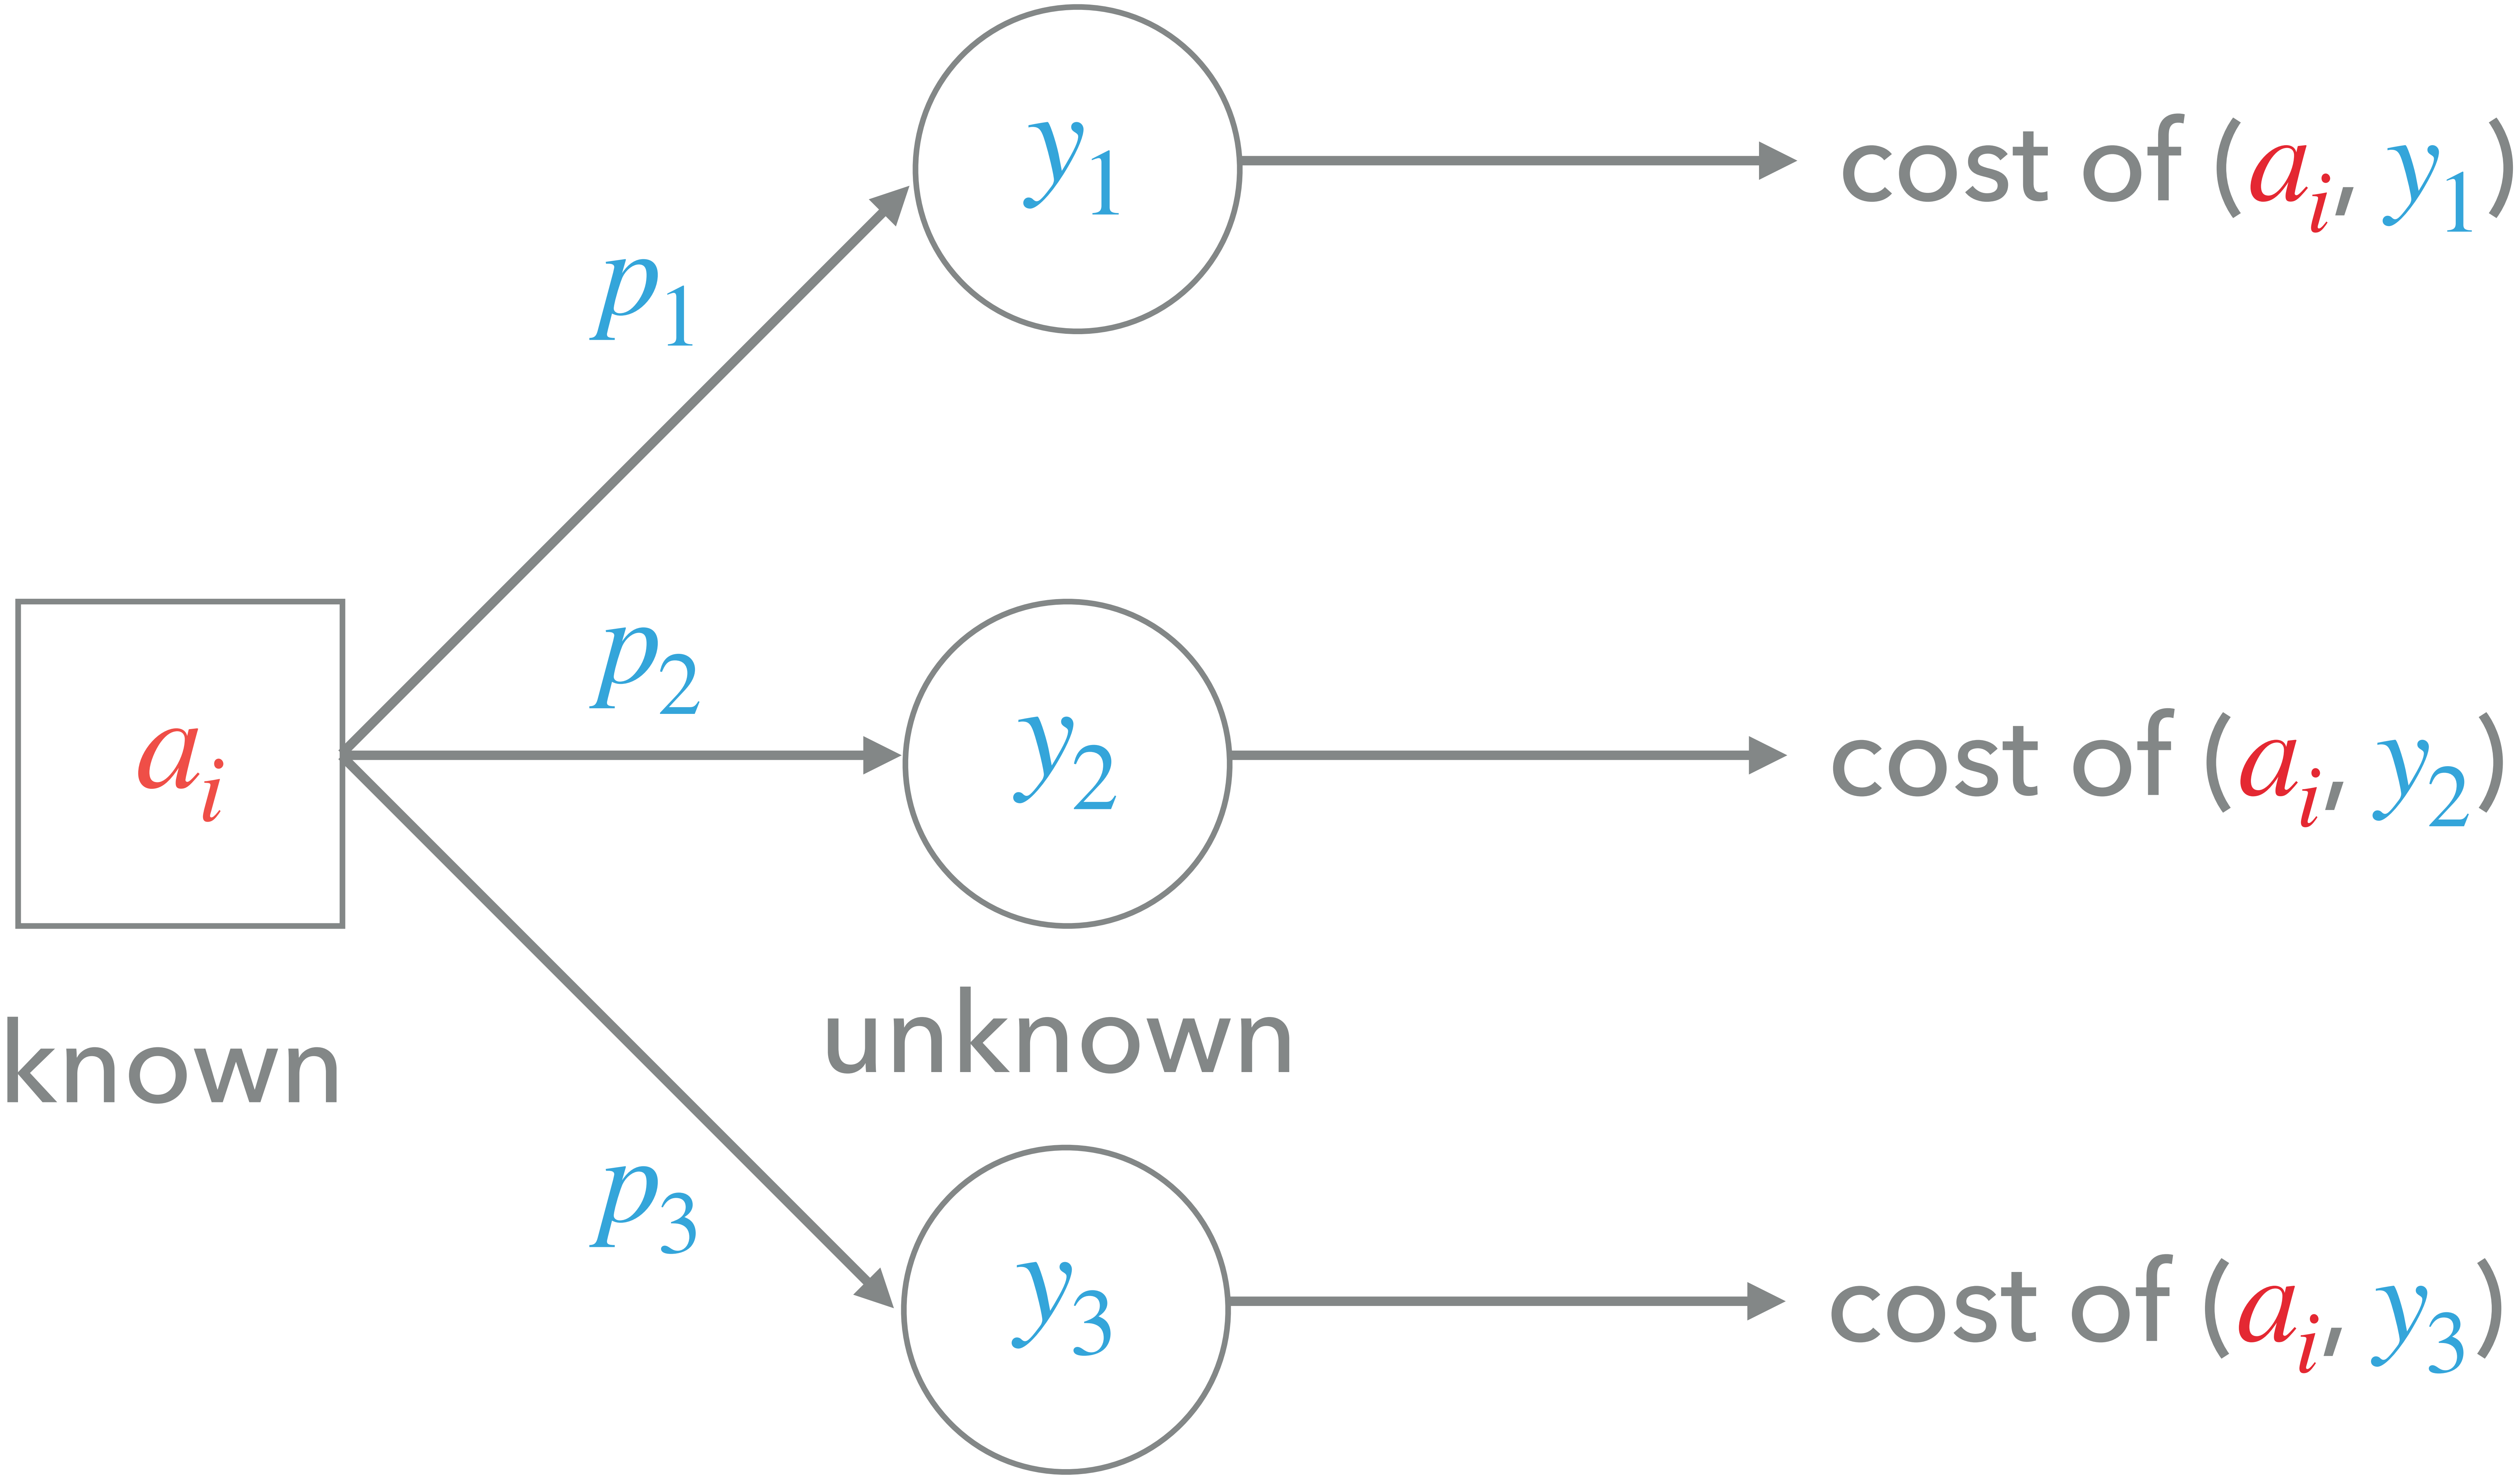

FANTASIZING ABOUT OUTCOMES

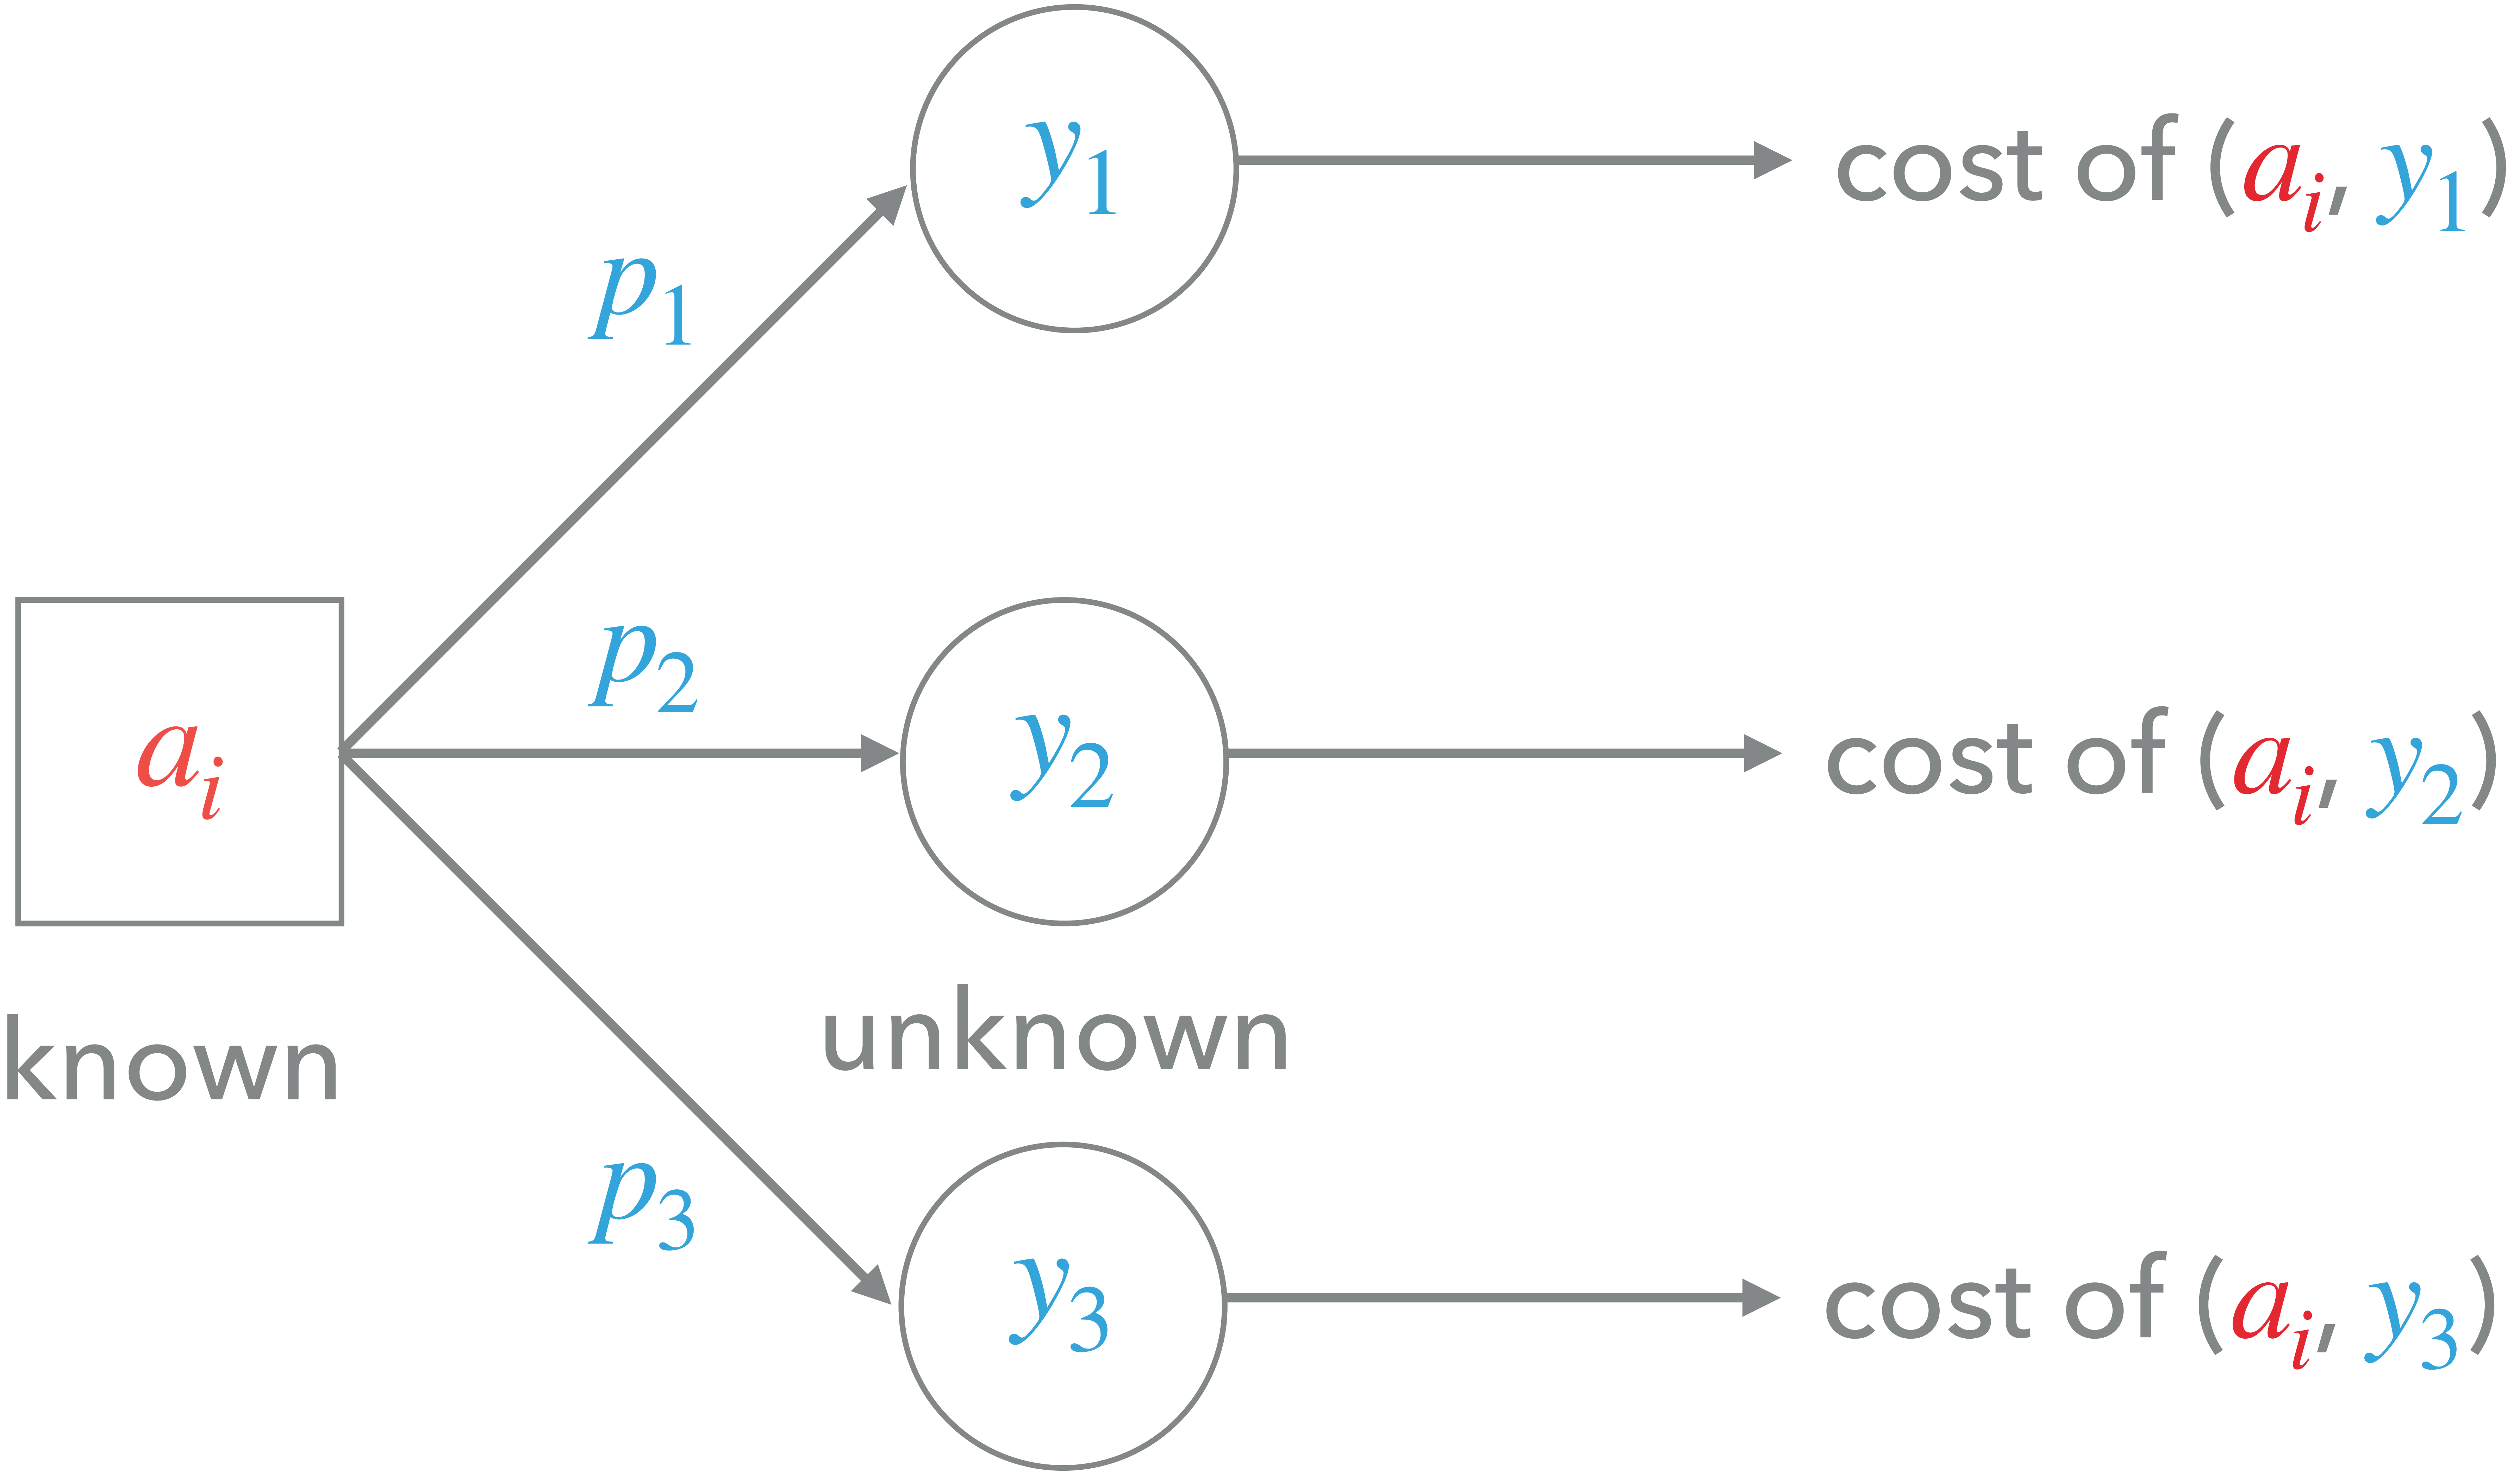

avg. cost of a_i

$\mathbb{E}_{y_j}\left[\text{cost of } (a_i, y_j)\right]$

–> pick the lowest-cost action

FANTASIZING ABOUT THE RAIN

	rain	no rain
umbrella	2	1
no umbrella	10	0

umbrella

FANTASIZING ABOUT THE RAIN

	rain	no rain
umbrella	2	1
no umbrella	10	0

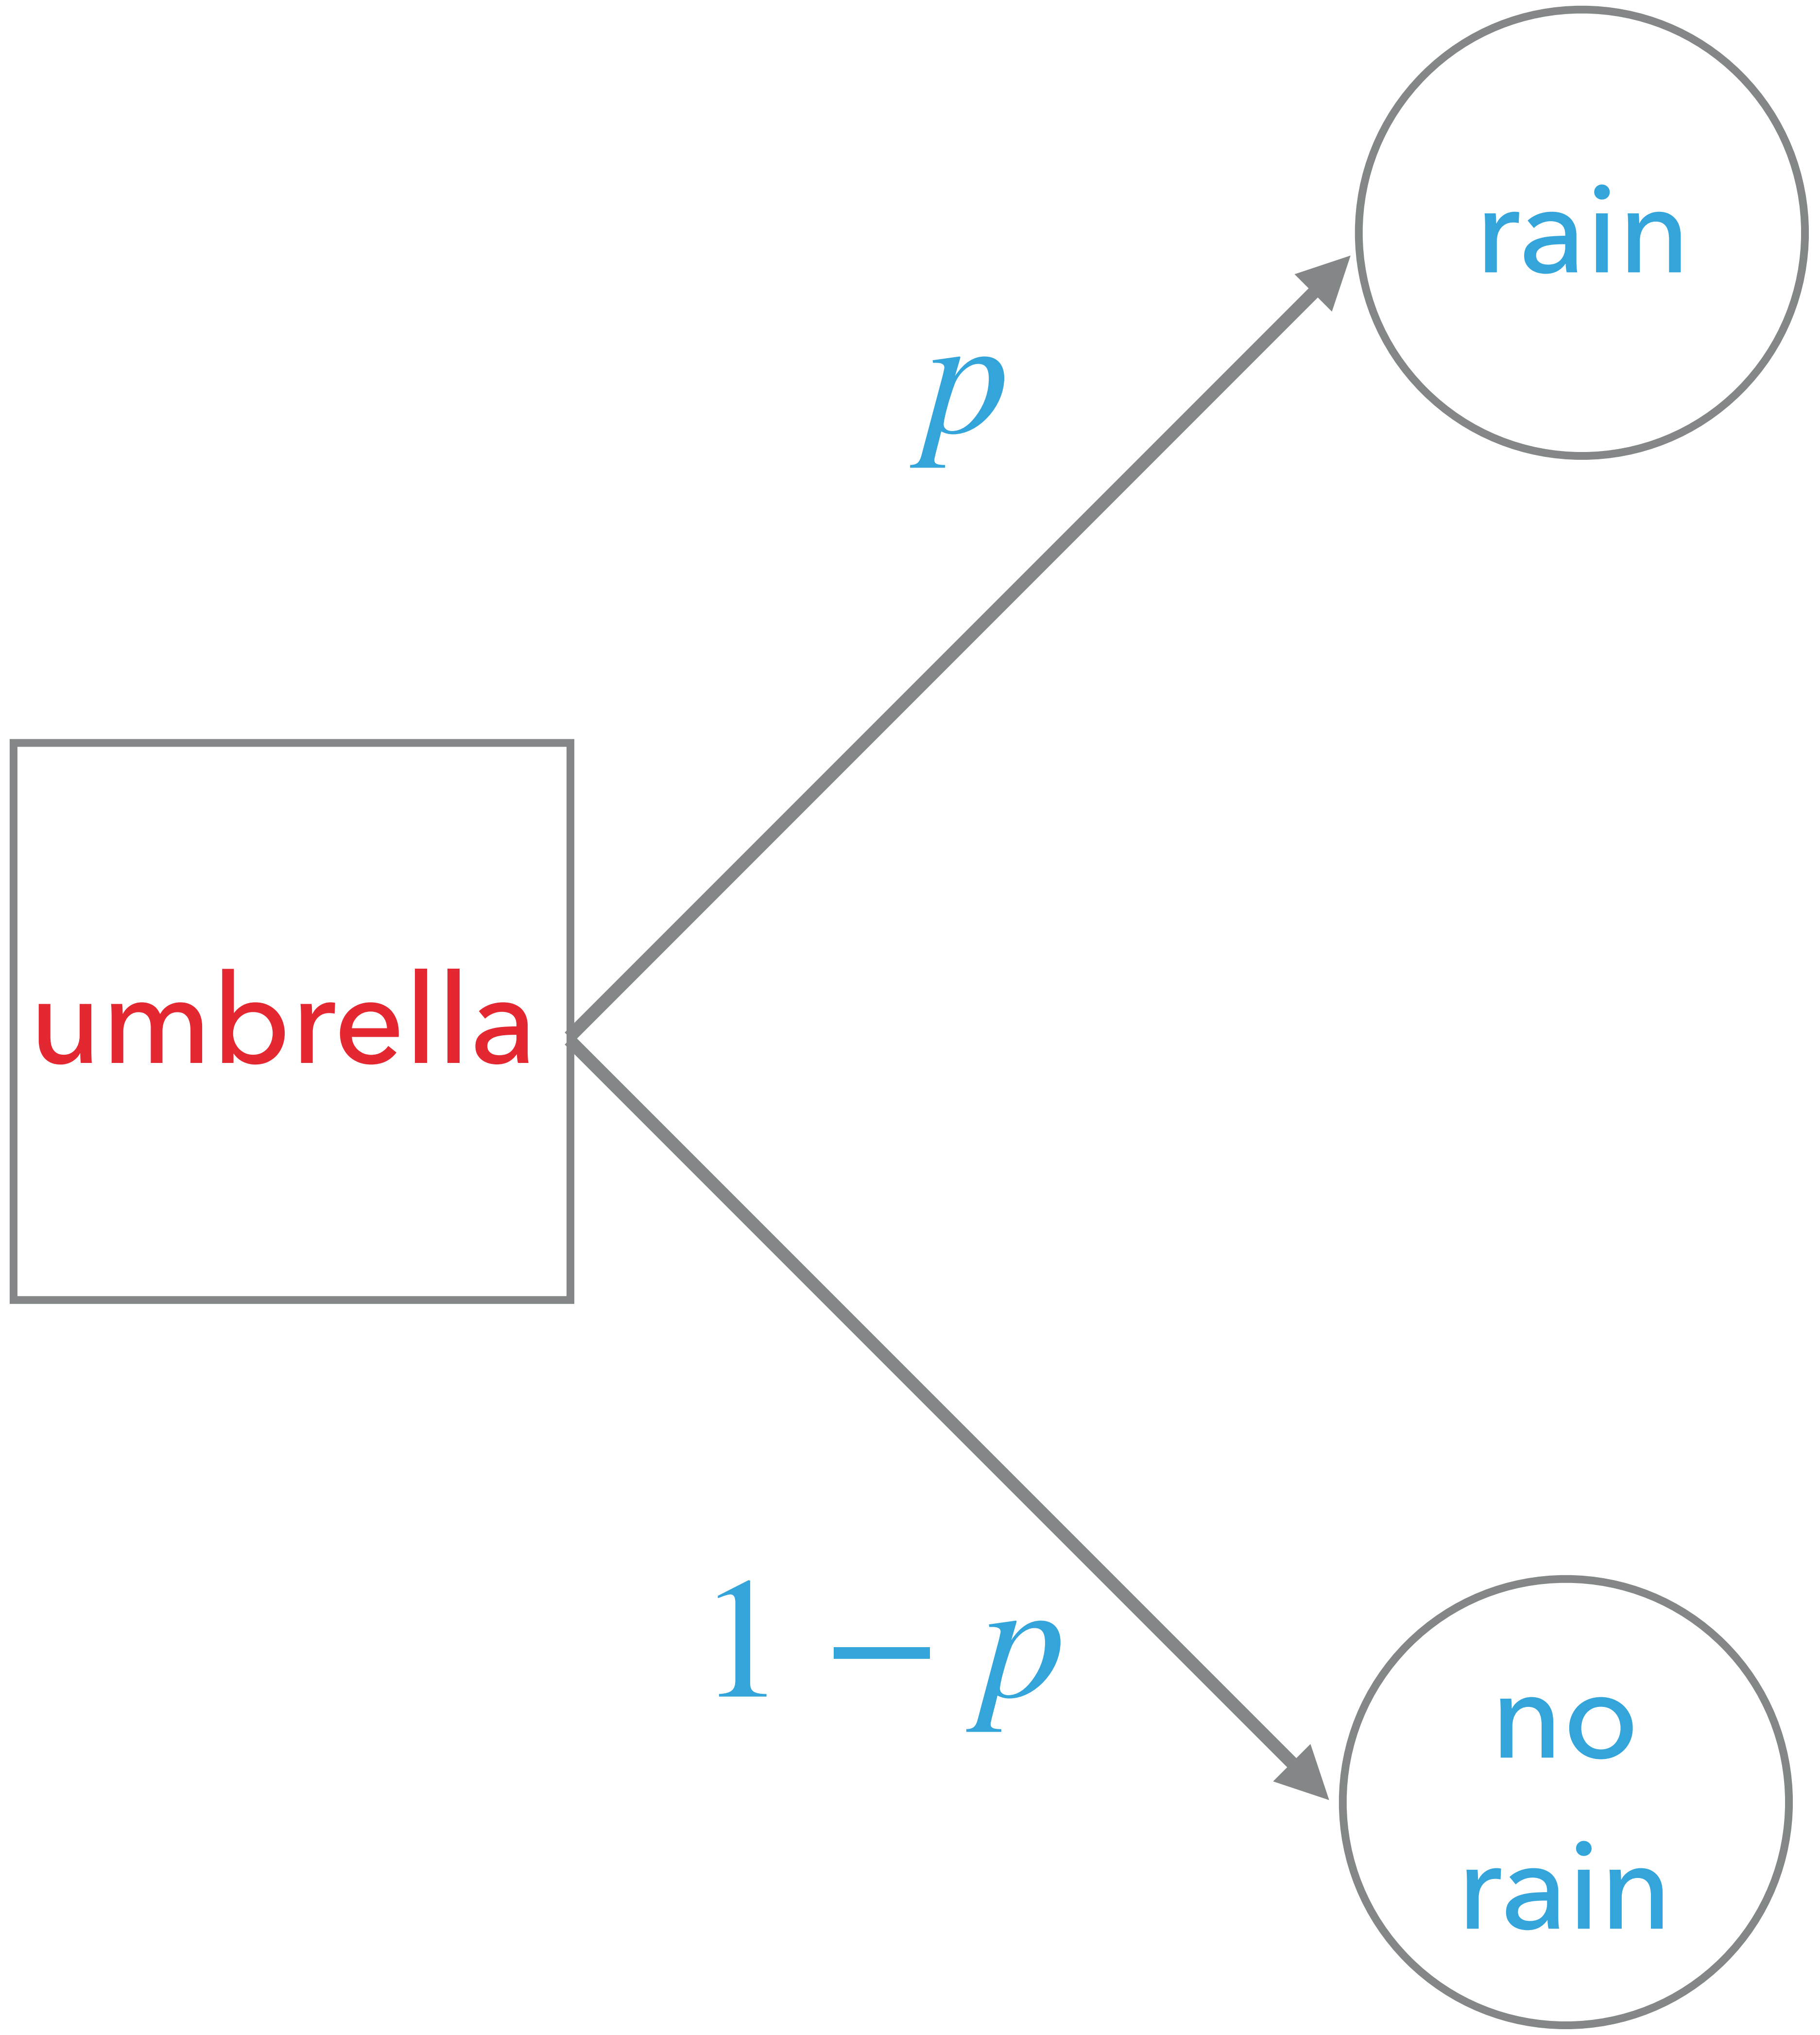

FANTASIZING ABOUT THE RAIN

	rain	no rain
umbrella	2	1
no umbrella	10	0

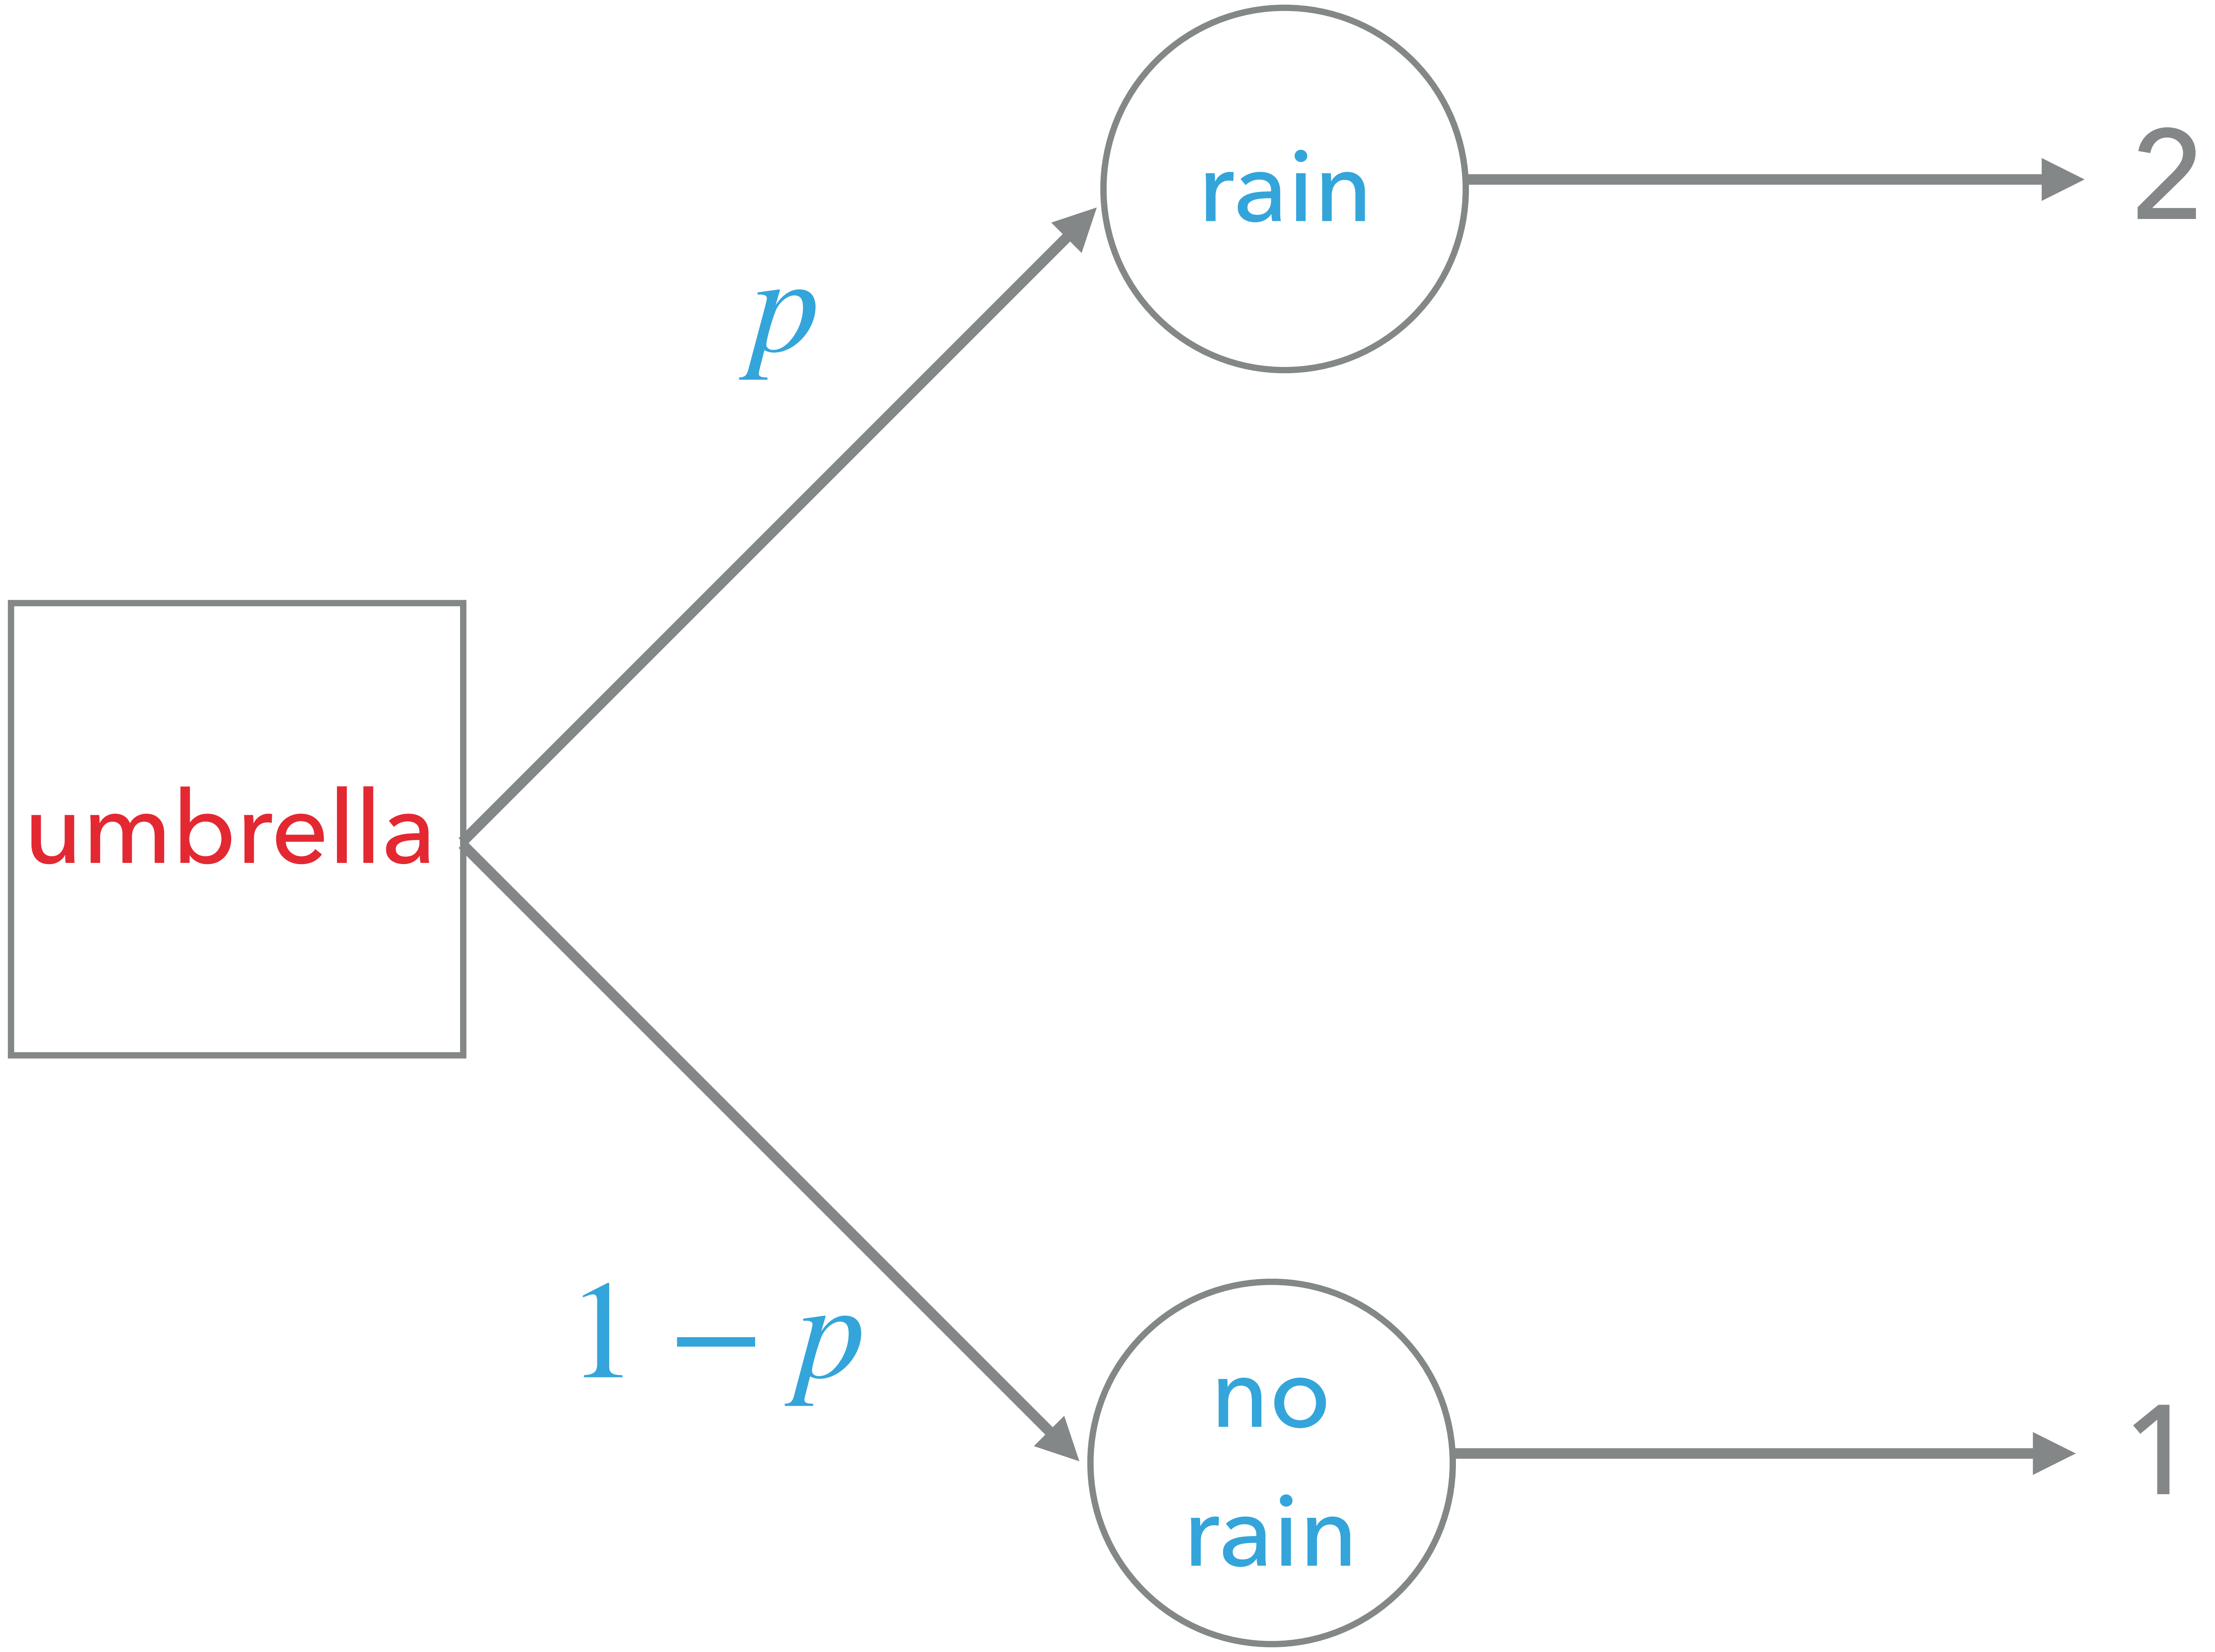

FANTASIZING ABOUT THE RAIN

	rain	no rain
umbrella	2	1
no umbrella	10	0

no umbrella

FANTASIZING ABOUT THE RAIN

	rain	no rain
umbrella	2	1
no umbrella	10	0

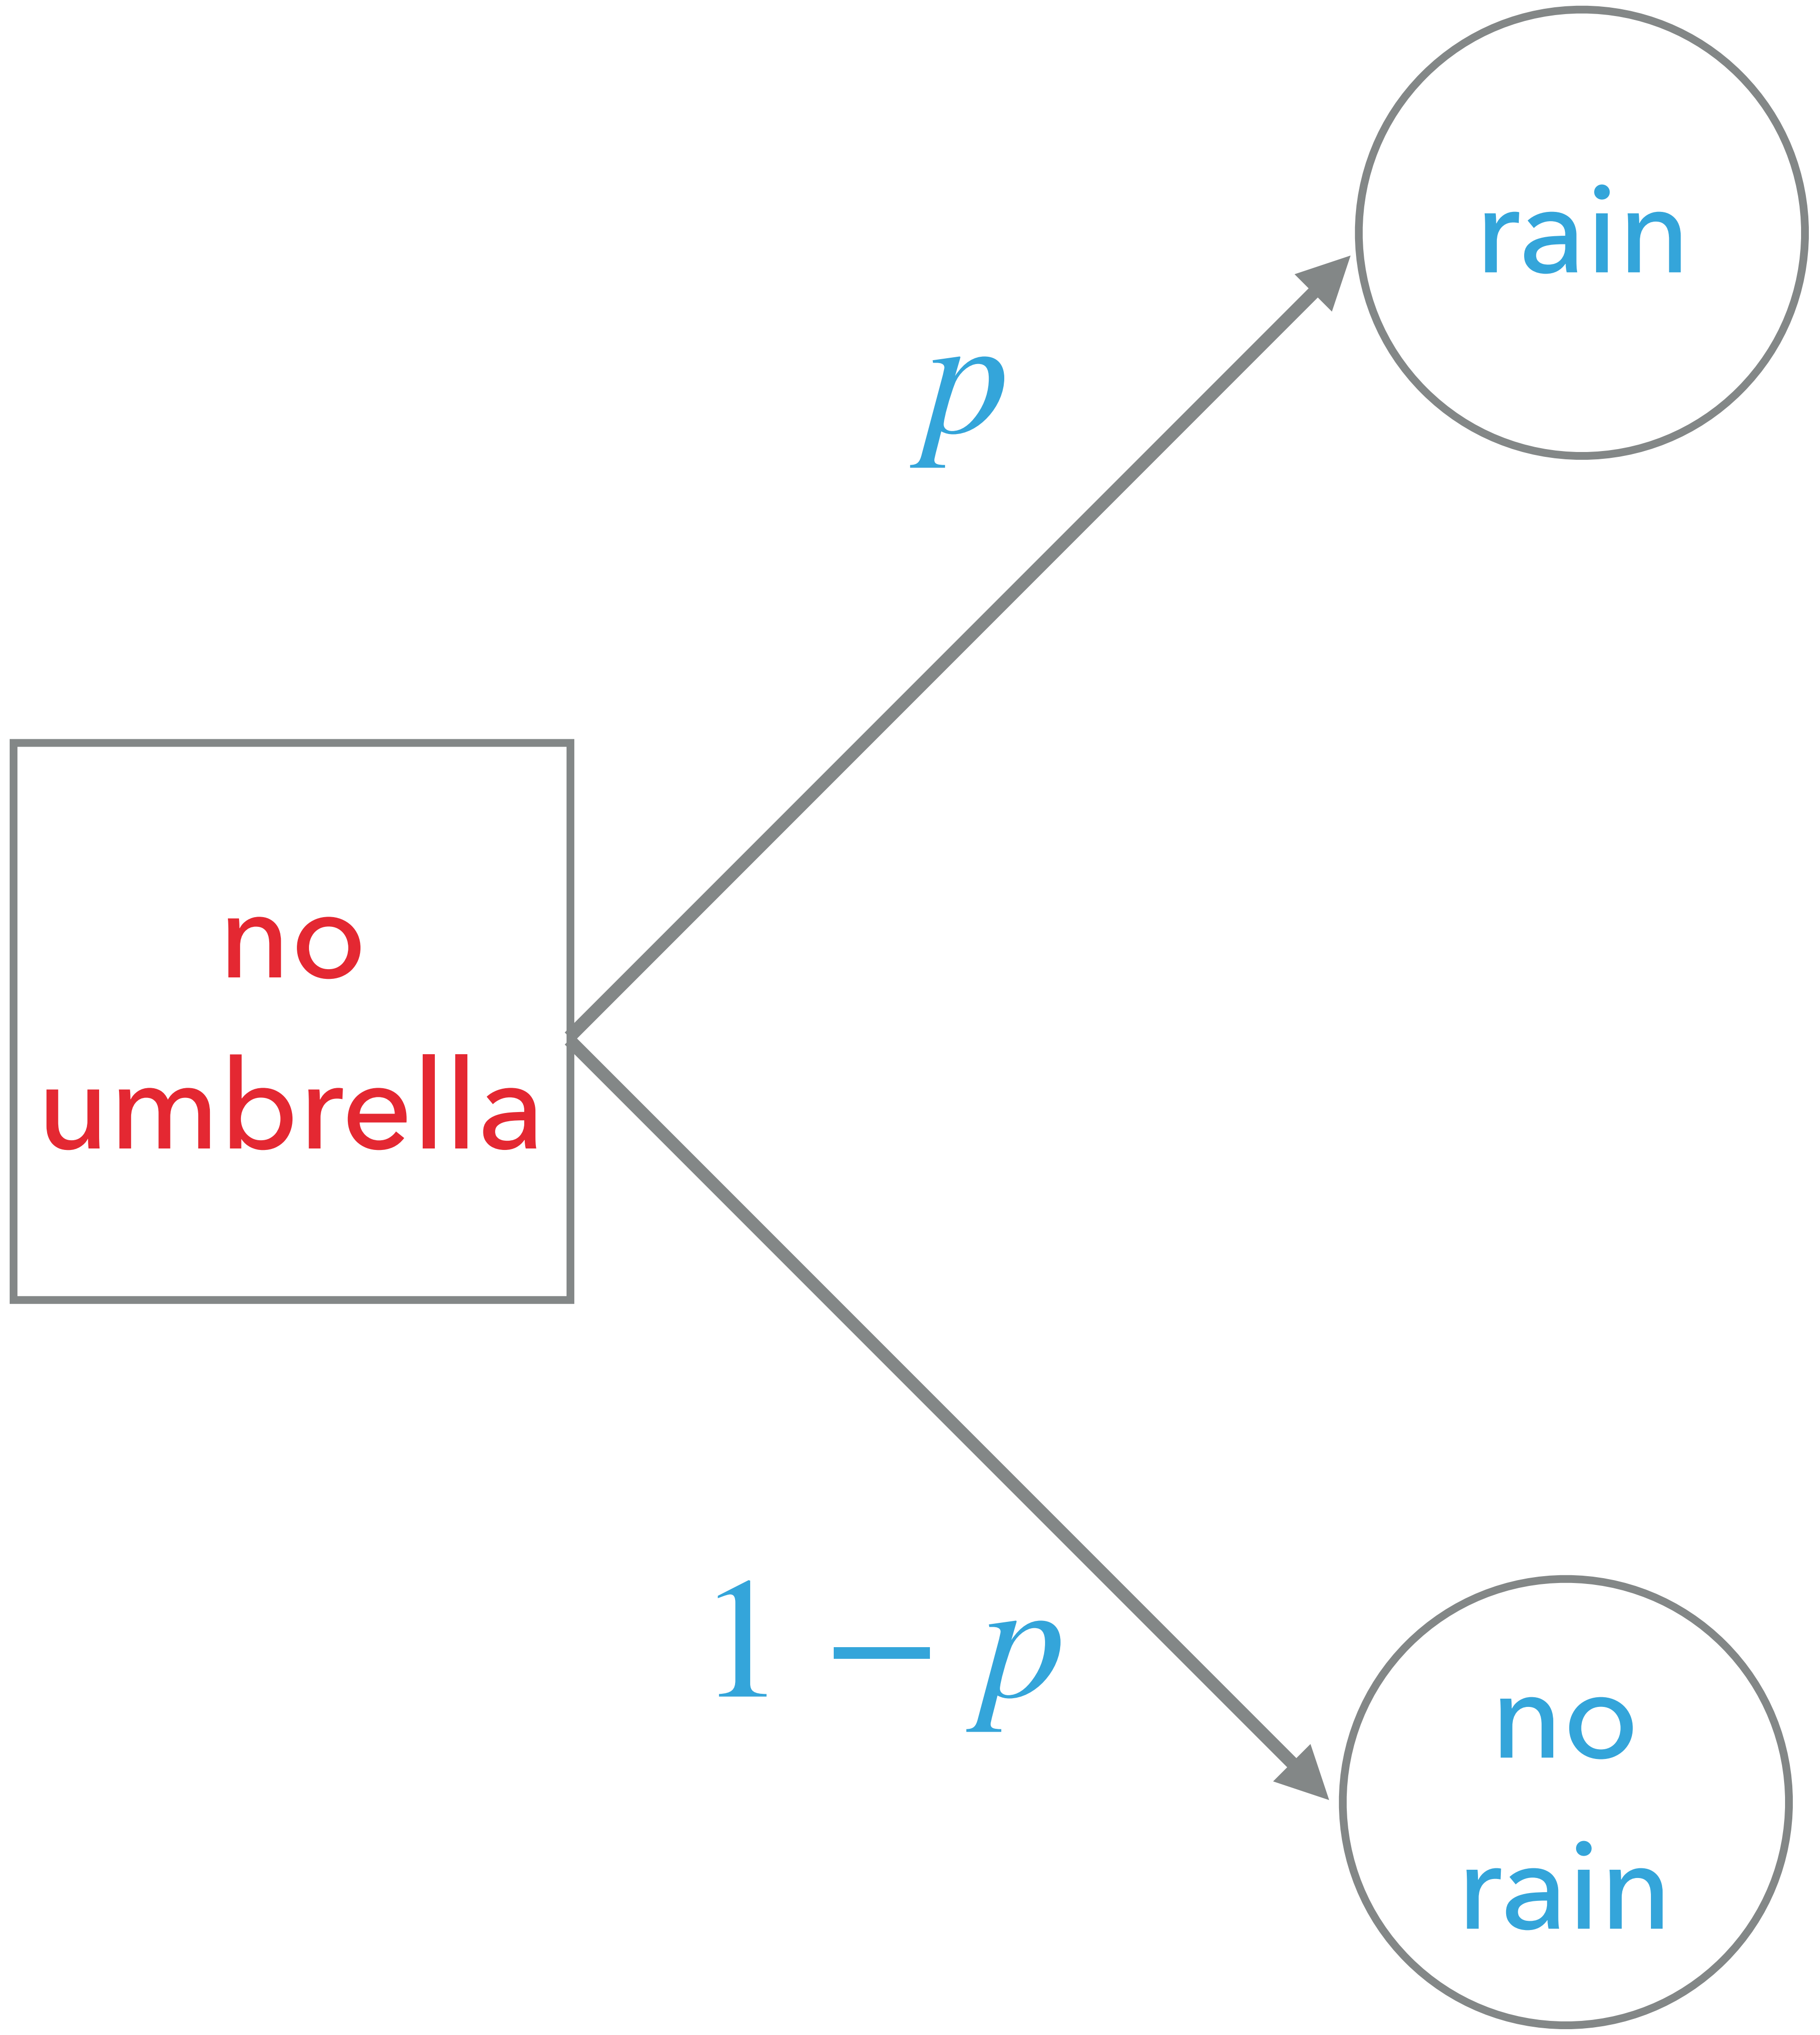

FANTASIZING ABOUT THE RAIN

	rain	no rain
umbrella	2	1
no umbrella	10	0

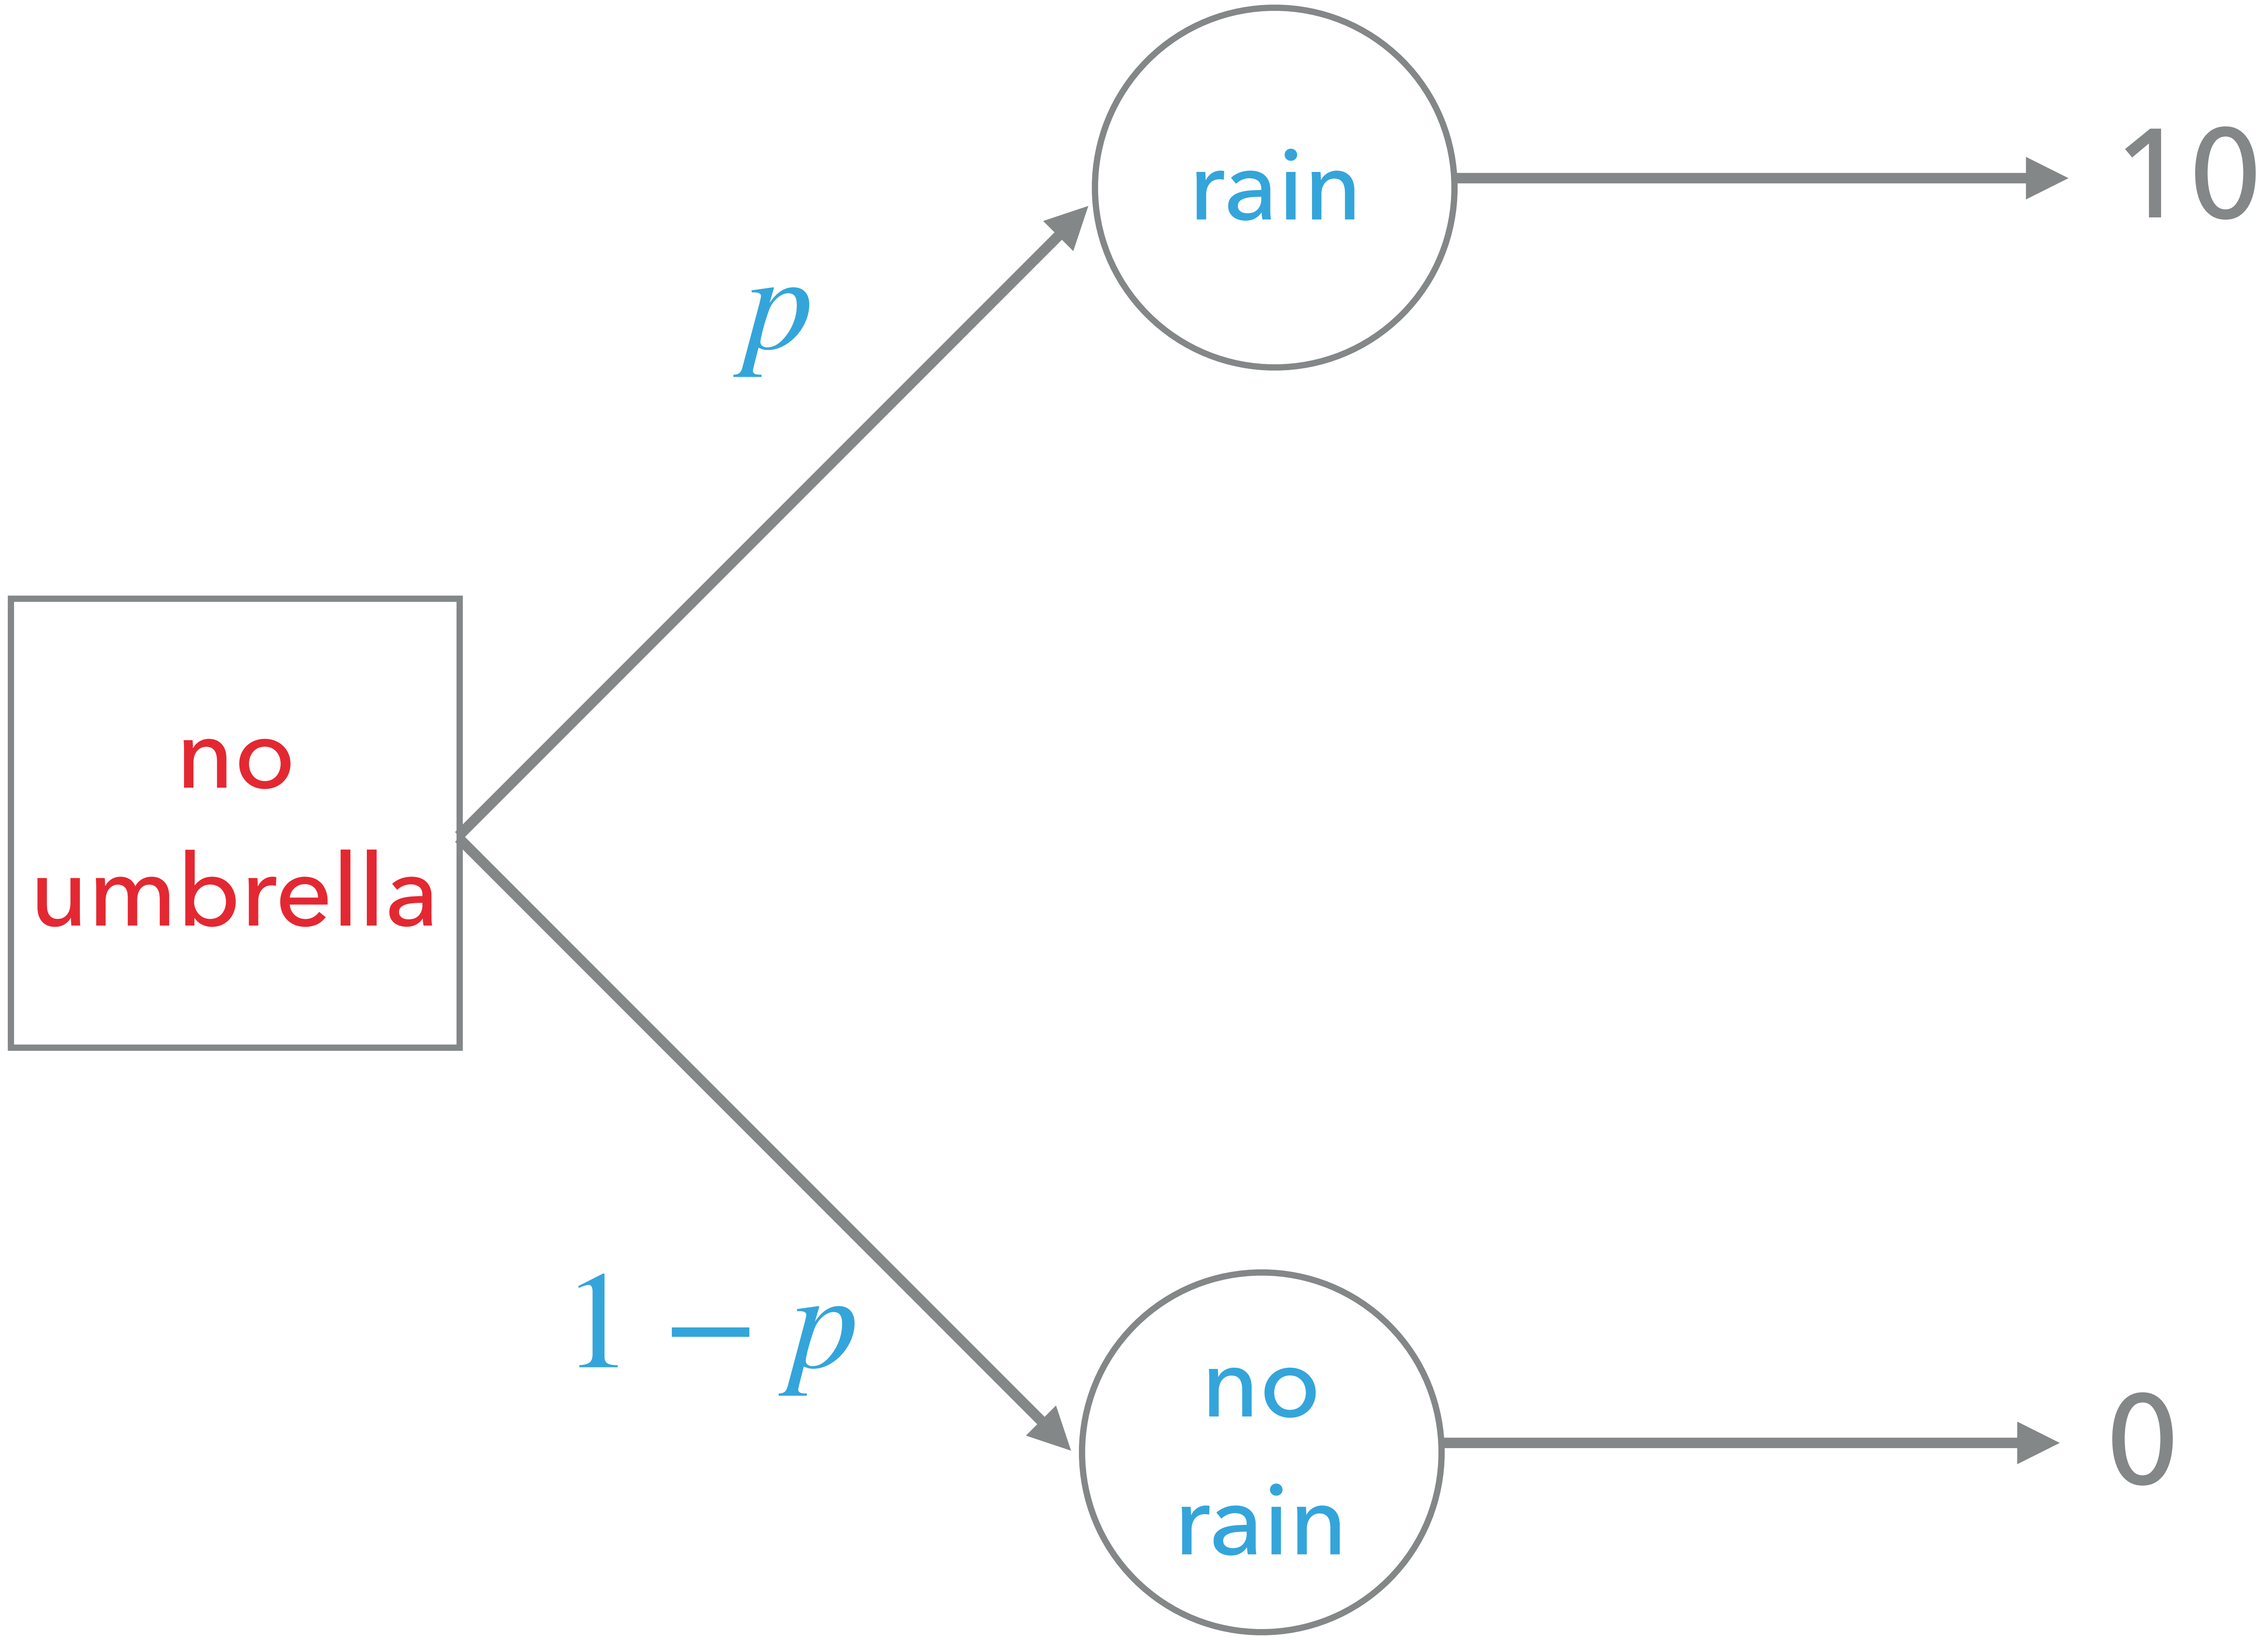

DECISION-MAKING IS CONTEXT-SPECIFIC

DECISION-MAKING IS CONTEXT-SPECIFIC

(Exp.) Costs: $p(2) + (1 - p)(1)$ vs.
$p(10)$

DECISION-MAKING IS CONTEXT-SPECIFIC

(Exp.) Costs: $p(2) + (1 - p)(1)$ vs. $p(10)$

- If $p = 0.9$, you **should** bring an umbrella

DECISION-MAKING IS CONTEXT-SPECIFIC

(Exp.) Costs: $p(2) + (1 - p)(1)$ vs. $p(10)$

- If $p = 0.9$, you **should** bring an umbrella
- If $p = 0.1$, you **should not** bring an umbrella

DECISION-MAKING IS CONTEXT-SPECIFIC

(Exp.) Costs: $p(2) + (1 - p)(1)$ vs. $p(10)$

- If $p = 0.9$, you **should** bring an umbrella
- If $p = 0.1$, you **should not** bring an umbrella

	rain	no rain
umbrella	2	1
no umbrella	1000	0

DECISION-MAKING IS CONTEXT-SPECIFIC

(Exp.) Costs: $p(2) + (1 - p)(1)$ vs. $p(10)$

- If $p = 0.9$, you **should** bring an umbrella
- If $p = 0.1$, you **should not** bring an umbrella

	rain	no rain
umbrella	1	1
no umbrella	0	0

COMPONENTS TO PROBABILISTIC DECISION-MAKING

Probabilistic predictive model

Decision-making policy

COMPONENTS TO PROBABILISTIC DECISION-MAKING

Probabilistic predictive model

Decision-making policy

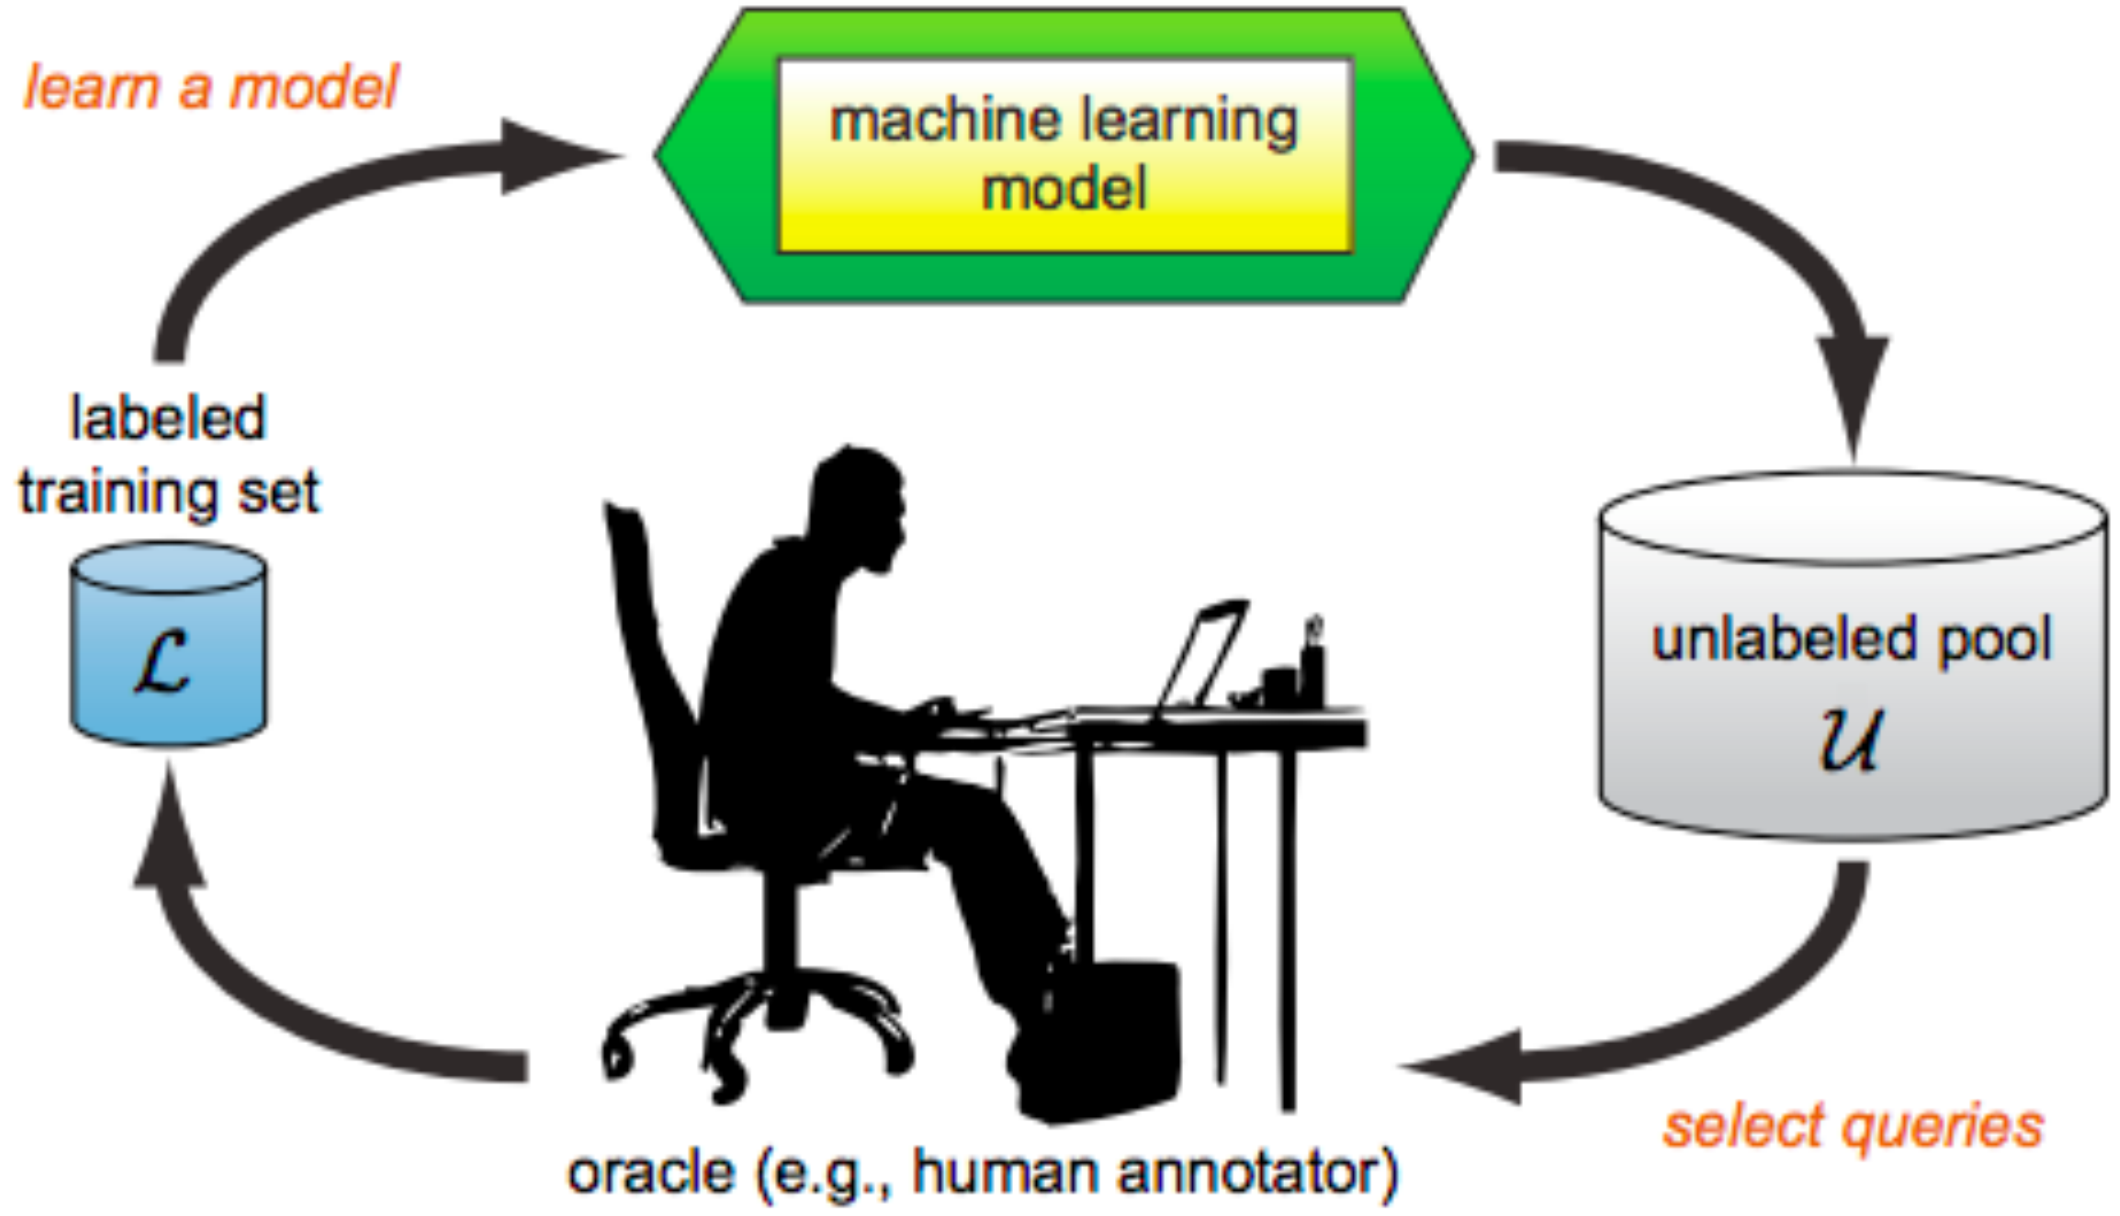

COMPONENTS TO PROBABILISTIC DECISION-MAKING

Probabilistic predictive model

Decision-making policy

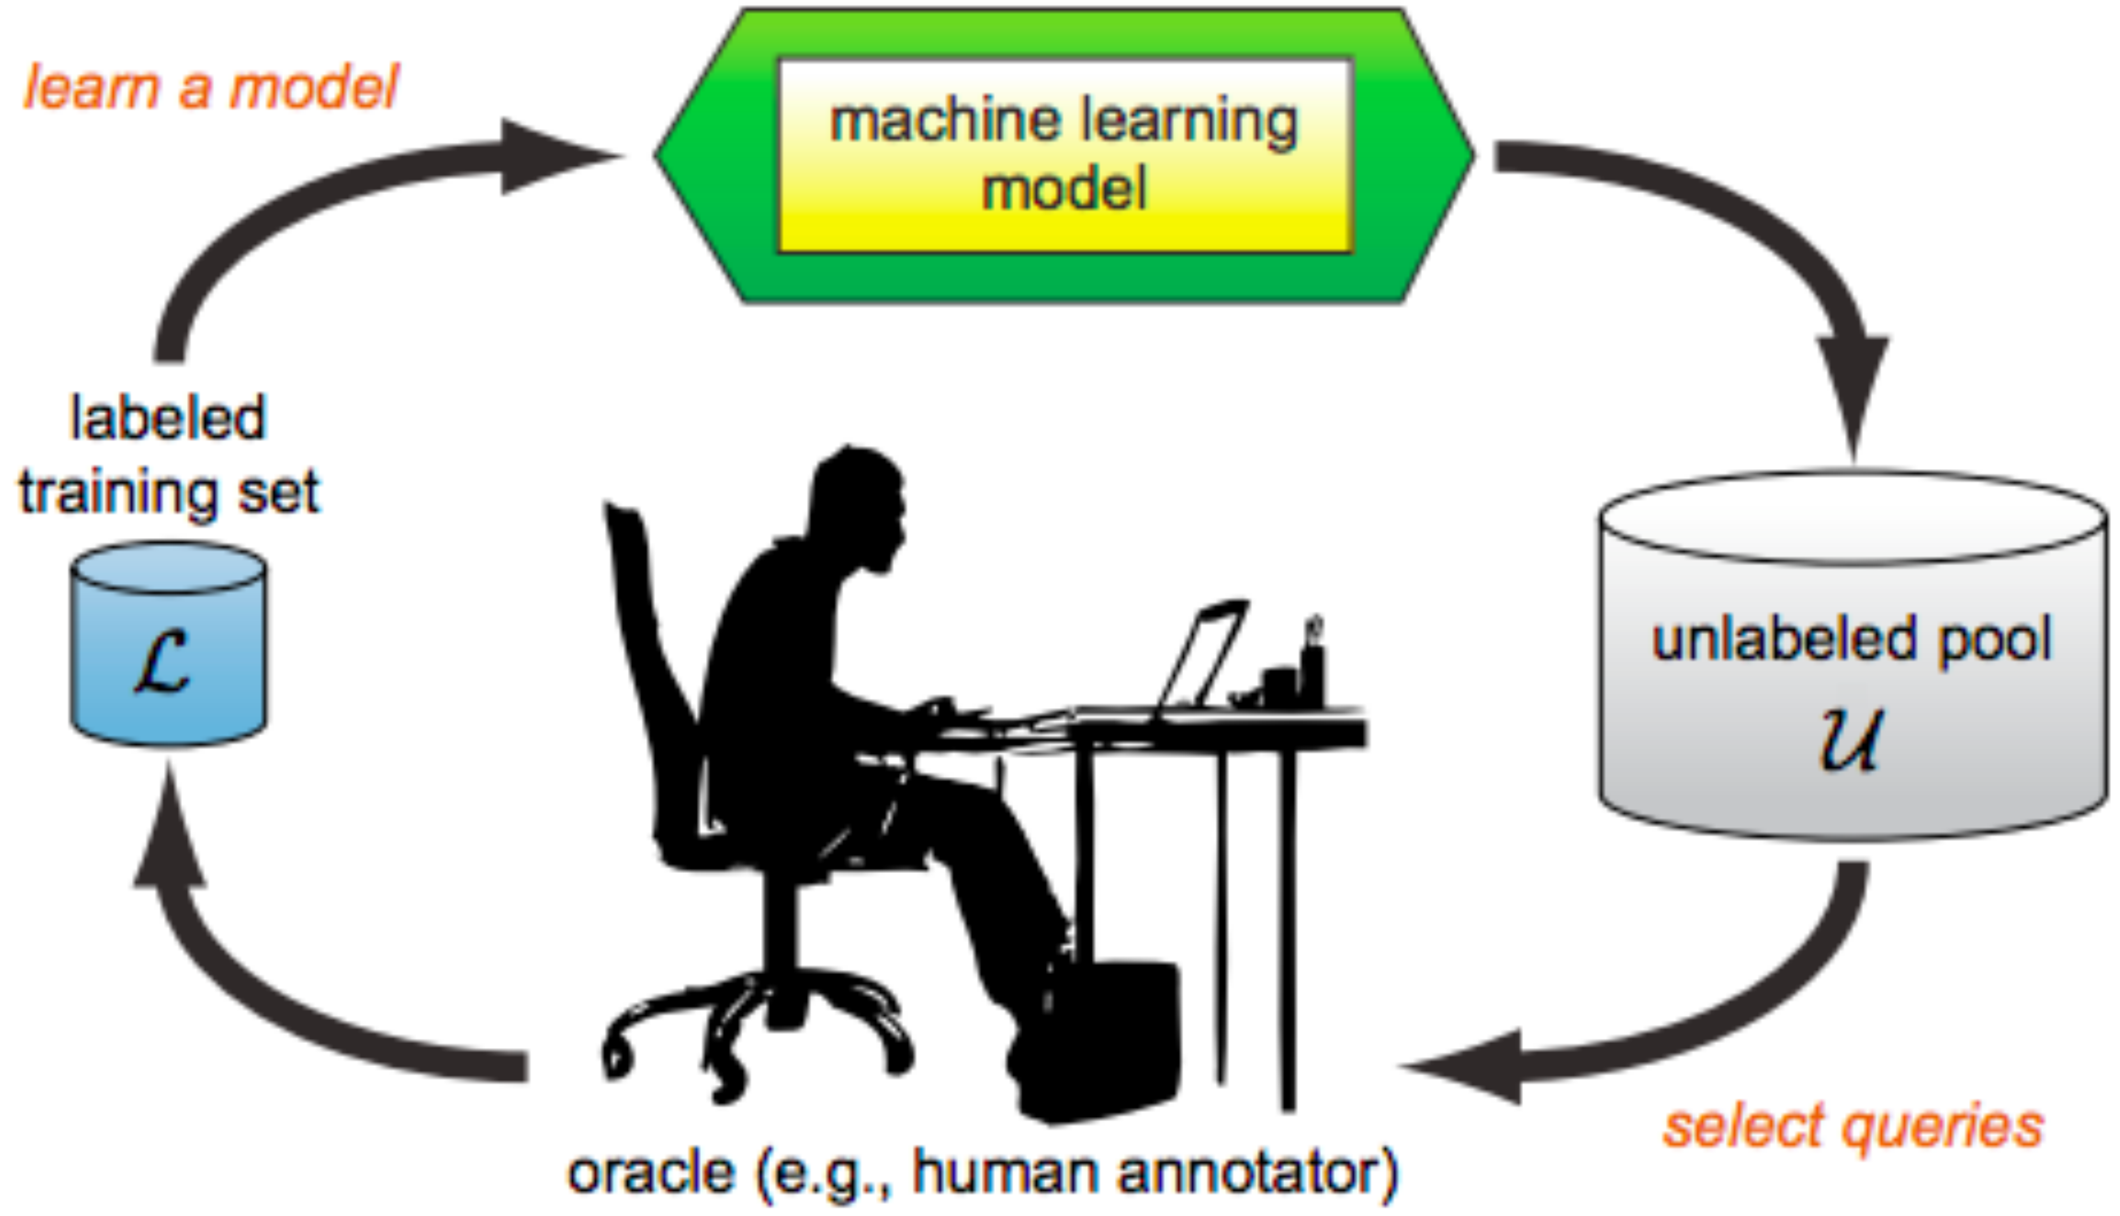

COMPONENTS TO PROBABILISTIC DECISION-MAKING

Probabilistic predictive model

Decision-making policy

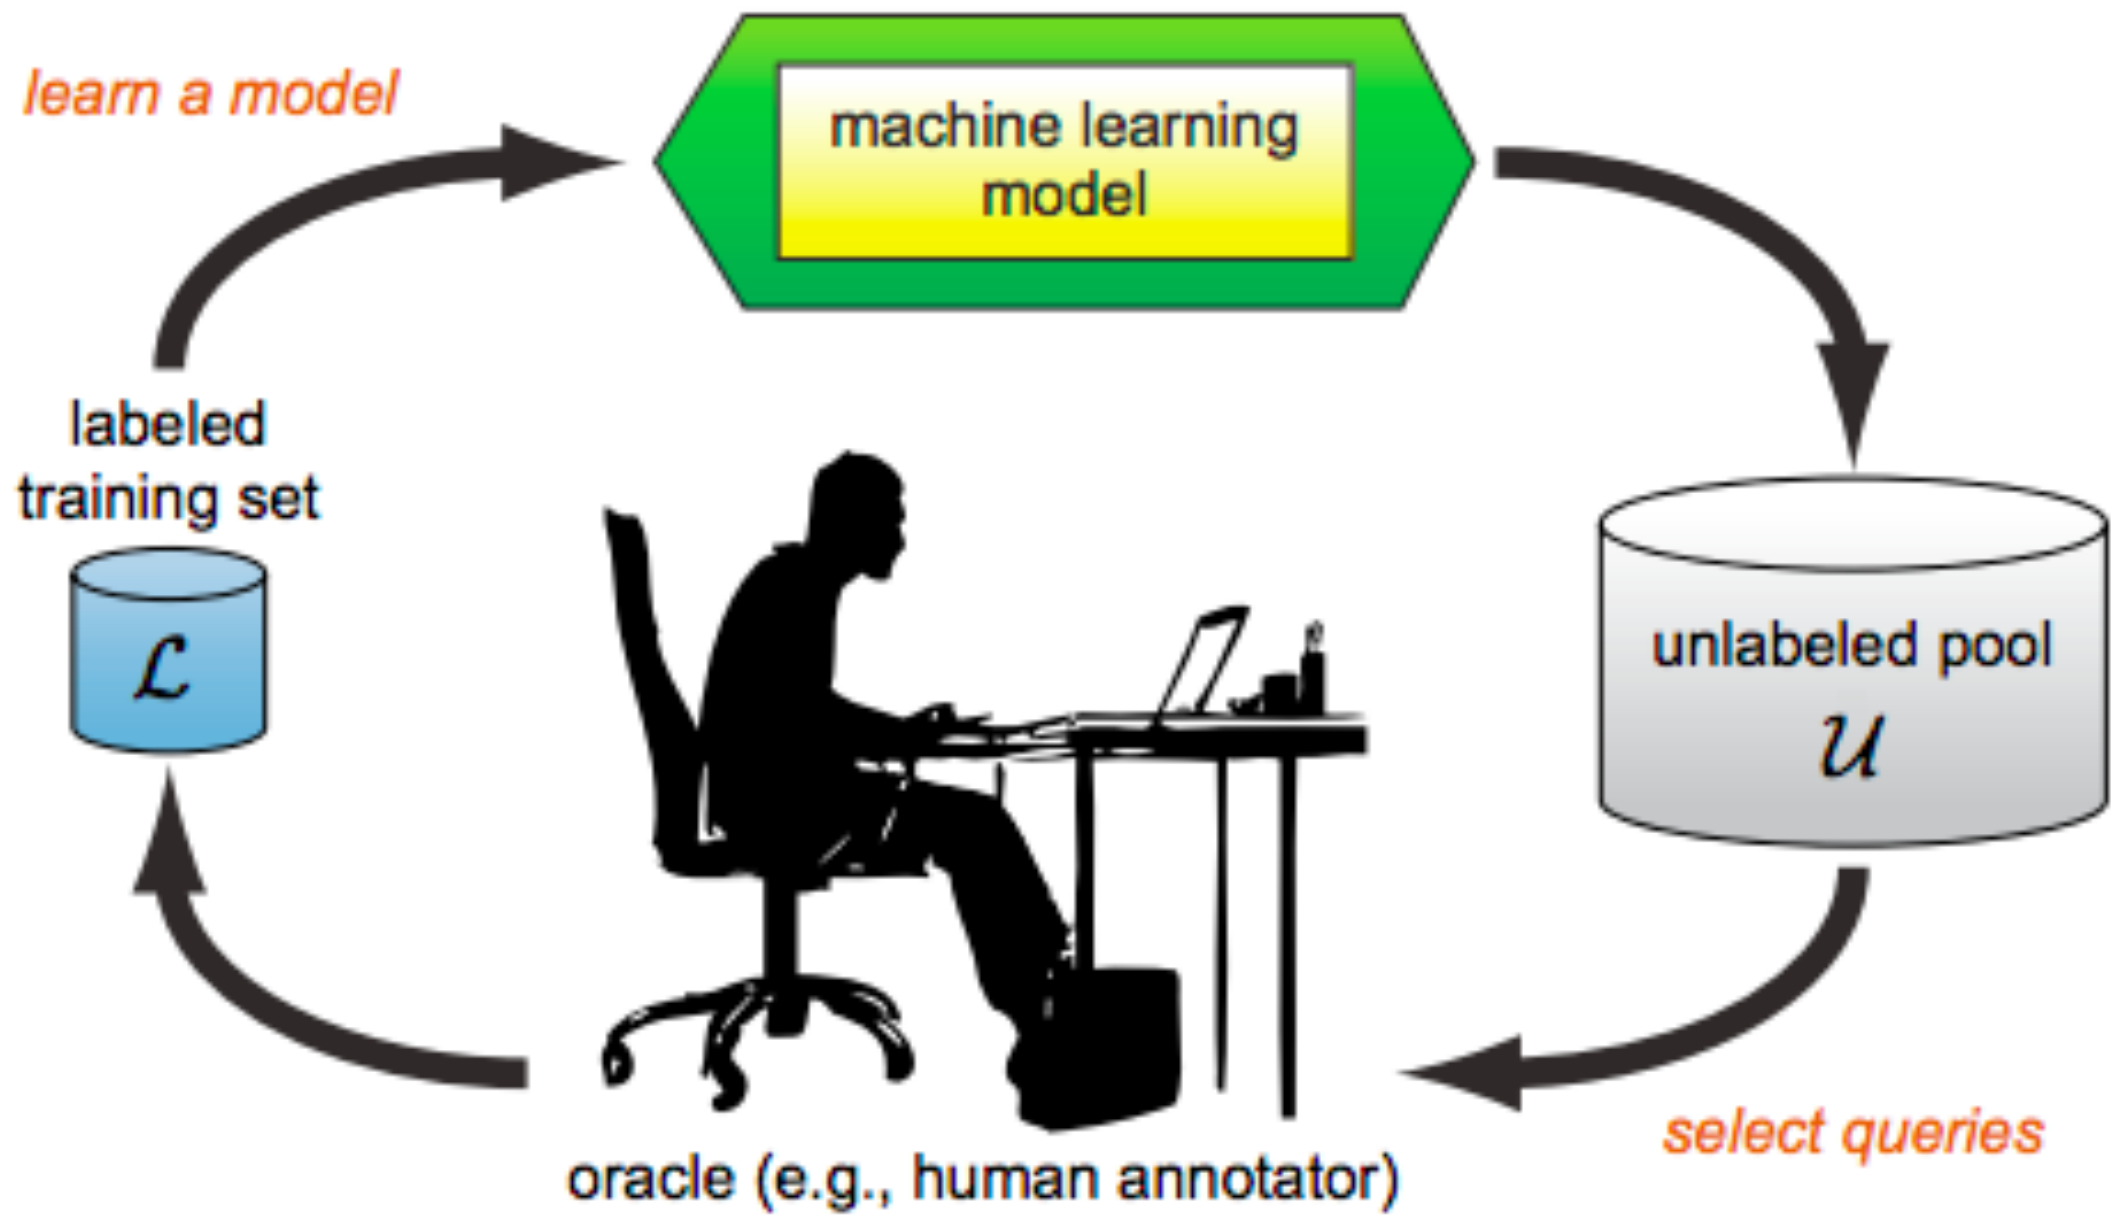

DECISION-MAKING IN BAYESIAN OPTIMIZATION

DECISION-MAKING IN BAYESIAN OPTIMIZATION

- Infinite number of **actions**

DECISION-MAKING IN BAYESIAN OPTIMIZATION

- Infinite number of **actions**
  - Which point x to query

DECISION-MAKING IN BAYESIAN OPTIMIZATION

- Infinite number of **actions**
  - Which point x to query
  - Search space is the domain D

DECISION-MAKING IN BAYESIAN OPTIMIZATION

- Infinite number of **actions**
  - Which point x to query
  - Search space is the domain D
- Infinite number of **possible outcomes**

DECISION-MAKING IN BAYESIAN OPTIMIZATION

- Infinite number of **actions**
  - Which point x to query
  - Search space is the domain D
- Infinite number of **possible outcomes**
  - The label y of the query x

DECISION-MAKING IN BAYESIAN OPTIMIZATION

- Infinite number of **actions**
  - Which point x to query
  - Search space is the domain D
- Infinite number of **possible outcomes**
  - The label y of the query x
  - Follows a normal distribution

DECISION-MAKING IN BAYESIAN OPTIMIZATION

- Infinite number of **actions**
  - Which point x to query
  - Search space is the domain D
- Infinite number of **possible outcomes**
  - The label y of the query x
  - Follows a normal distribution

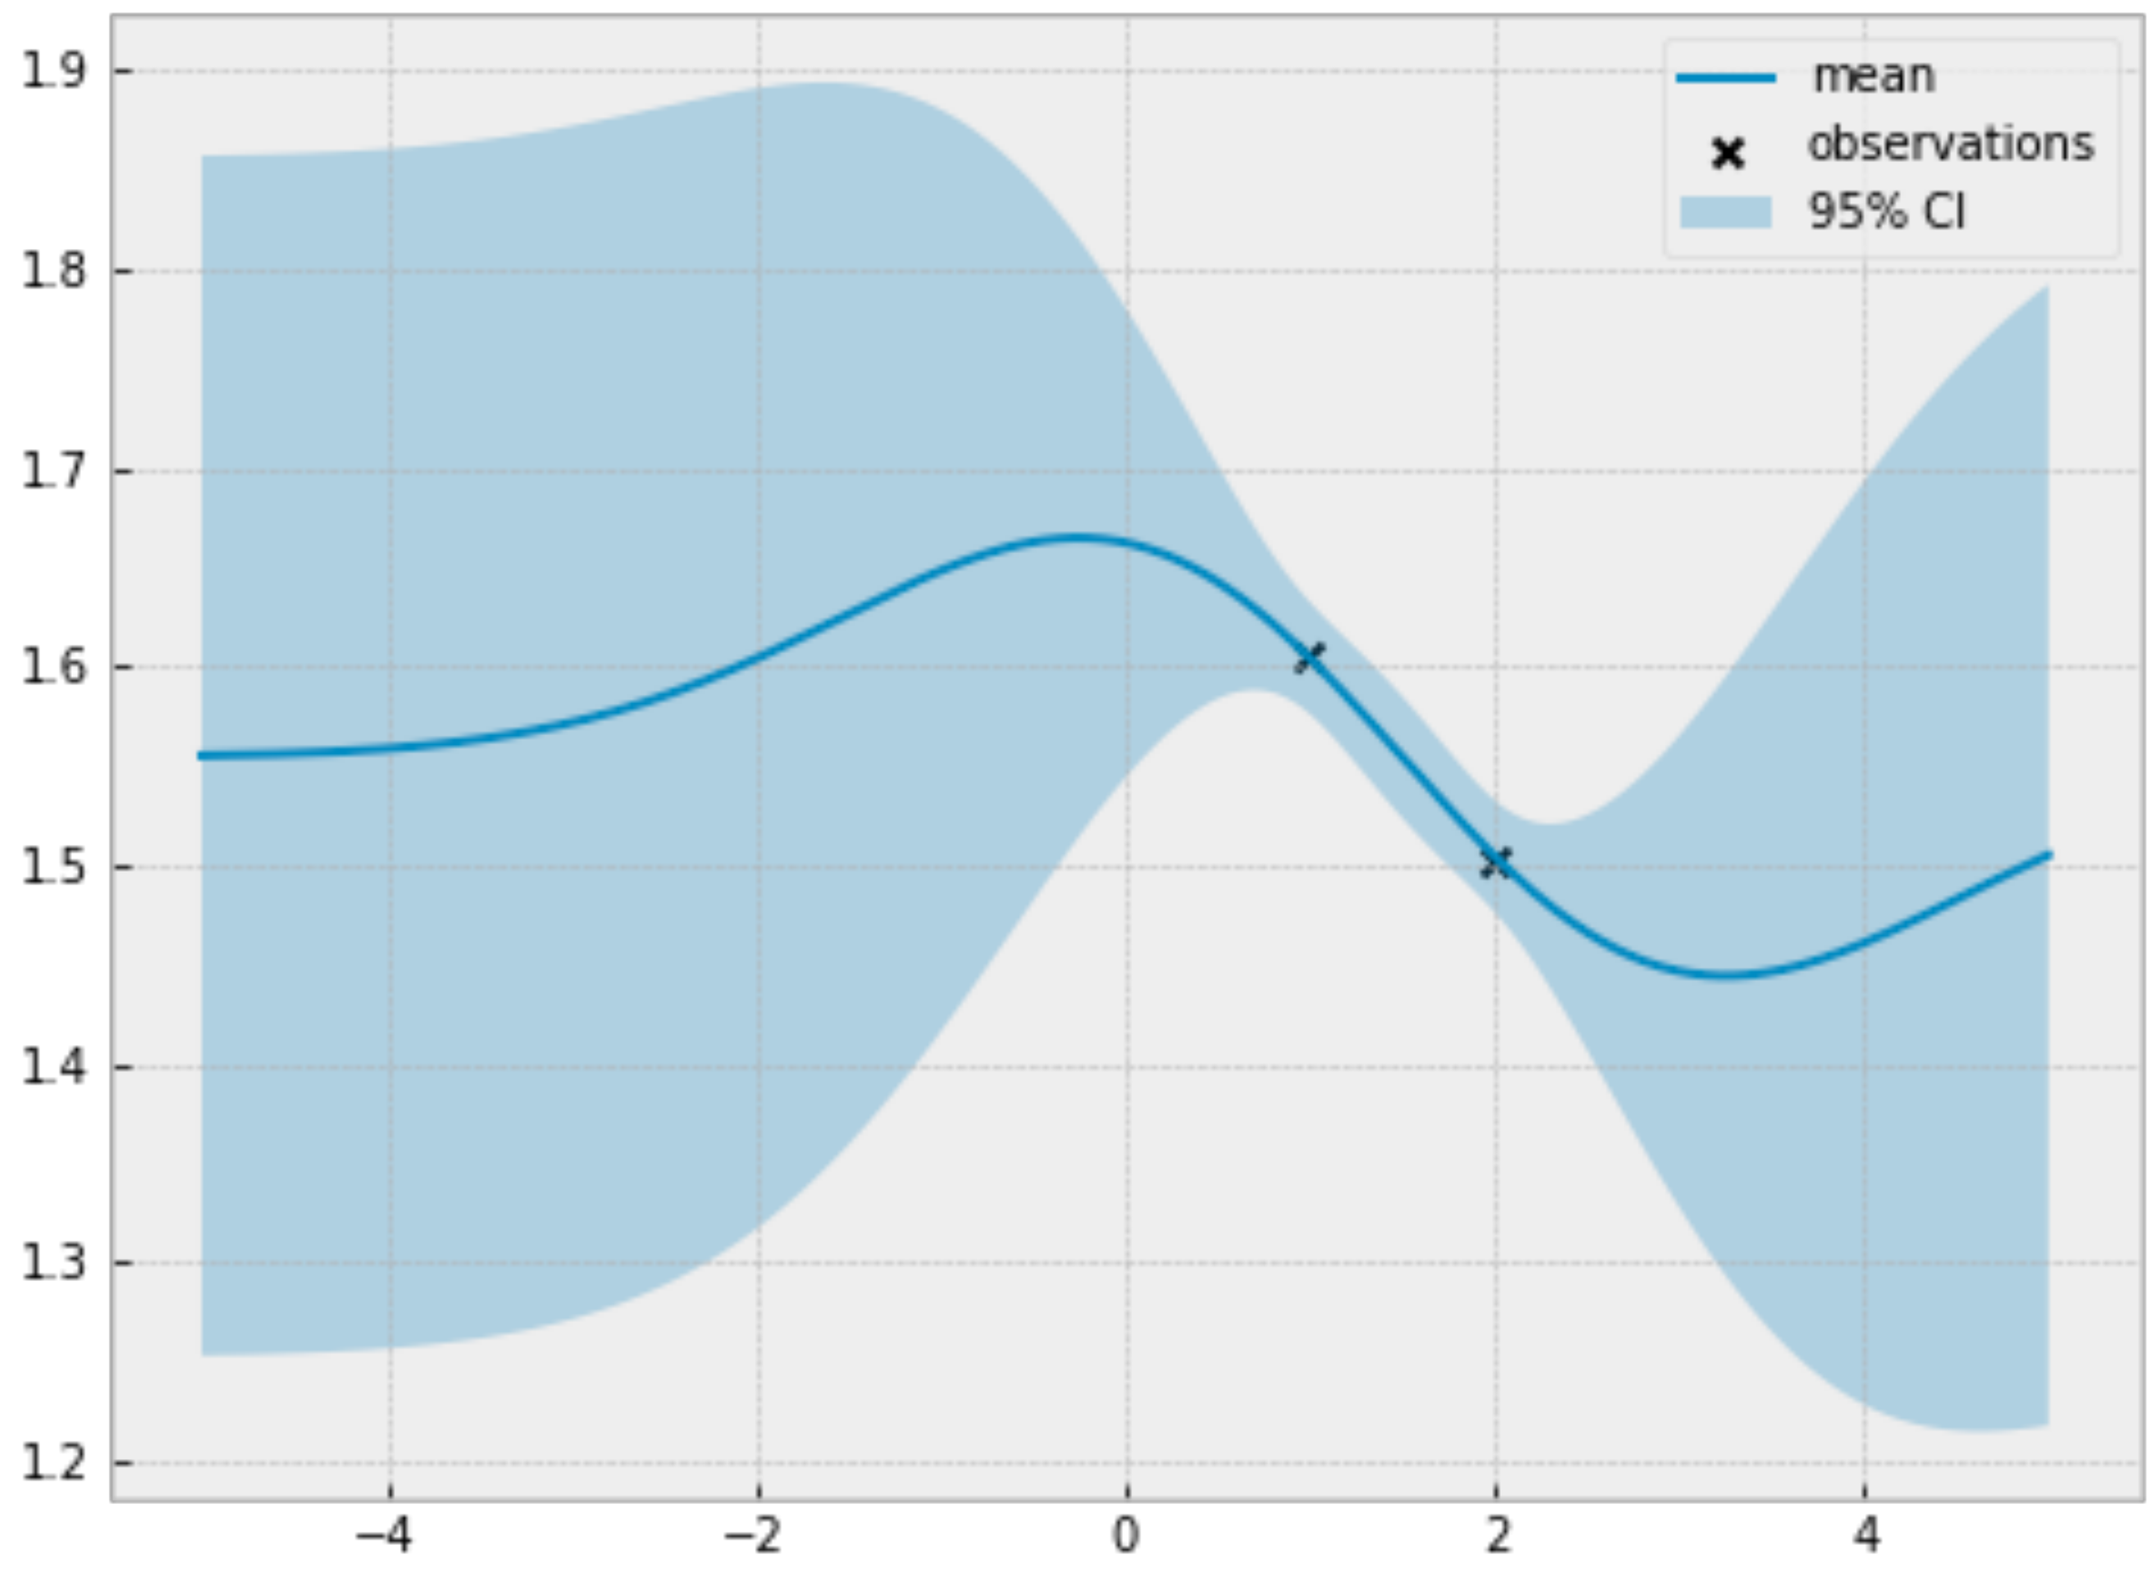

DEFINING UTILITY IN BAYESIAN OPTIMIZATION

DEFINING UTILITY IN BAYESIAN OPTIMIZATION

Goal: finding the function maximizer $x^* = \arg\max_D f(x)$

DEFINING UTILITY IN BAYESIAN OPTIMIZATION

Goal: finding the function maximizer $x^* = \arg\max_D f(x)$

Care about: uncovering a large $f(x)$ value

DEFINING UTILITY IN BAYESIAN OPTIMIZATION

Goal: finding the function maximizer $x^* = \arg\max_D f(x)$

Care about: uncovering a large $f(x)$ value

Concrete utility: improving from the best point seen so far (incumbent)

WHETHER TO IMPROVE FROM THE INCUMBENT

Utility: 1 if improve from the incumbent, 0 otherwise

WHETHER TO IMPROVE FROM THE INCUMBENT

Utility: 1 if improve from the incumbent, 0 otherwise

	…	-100	-99.9999	…	1.6054	1.6055	…
0	…	0	0	…	0	1	…
0.0001	…	0	0	…	0	1	…
0.0002	…	0	0	…	0	1	…
…	…	…	…	…	…	…	…

FANTASIZING ABOUT IMPROVEMENT

$$\boxed{x}$$

FANTASIZING ABOUT IMPROVEMENT

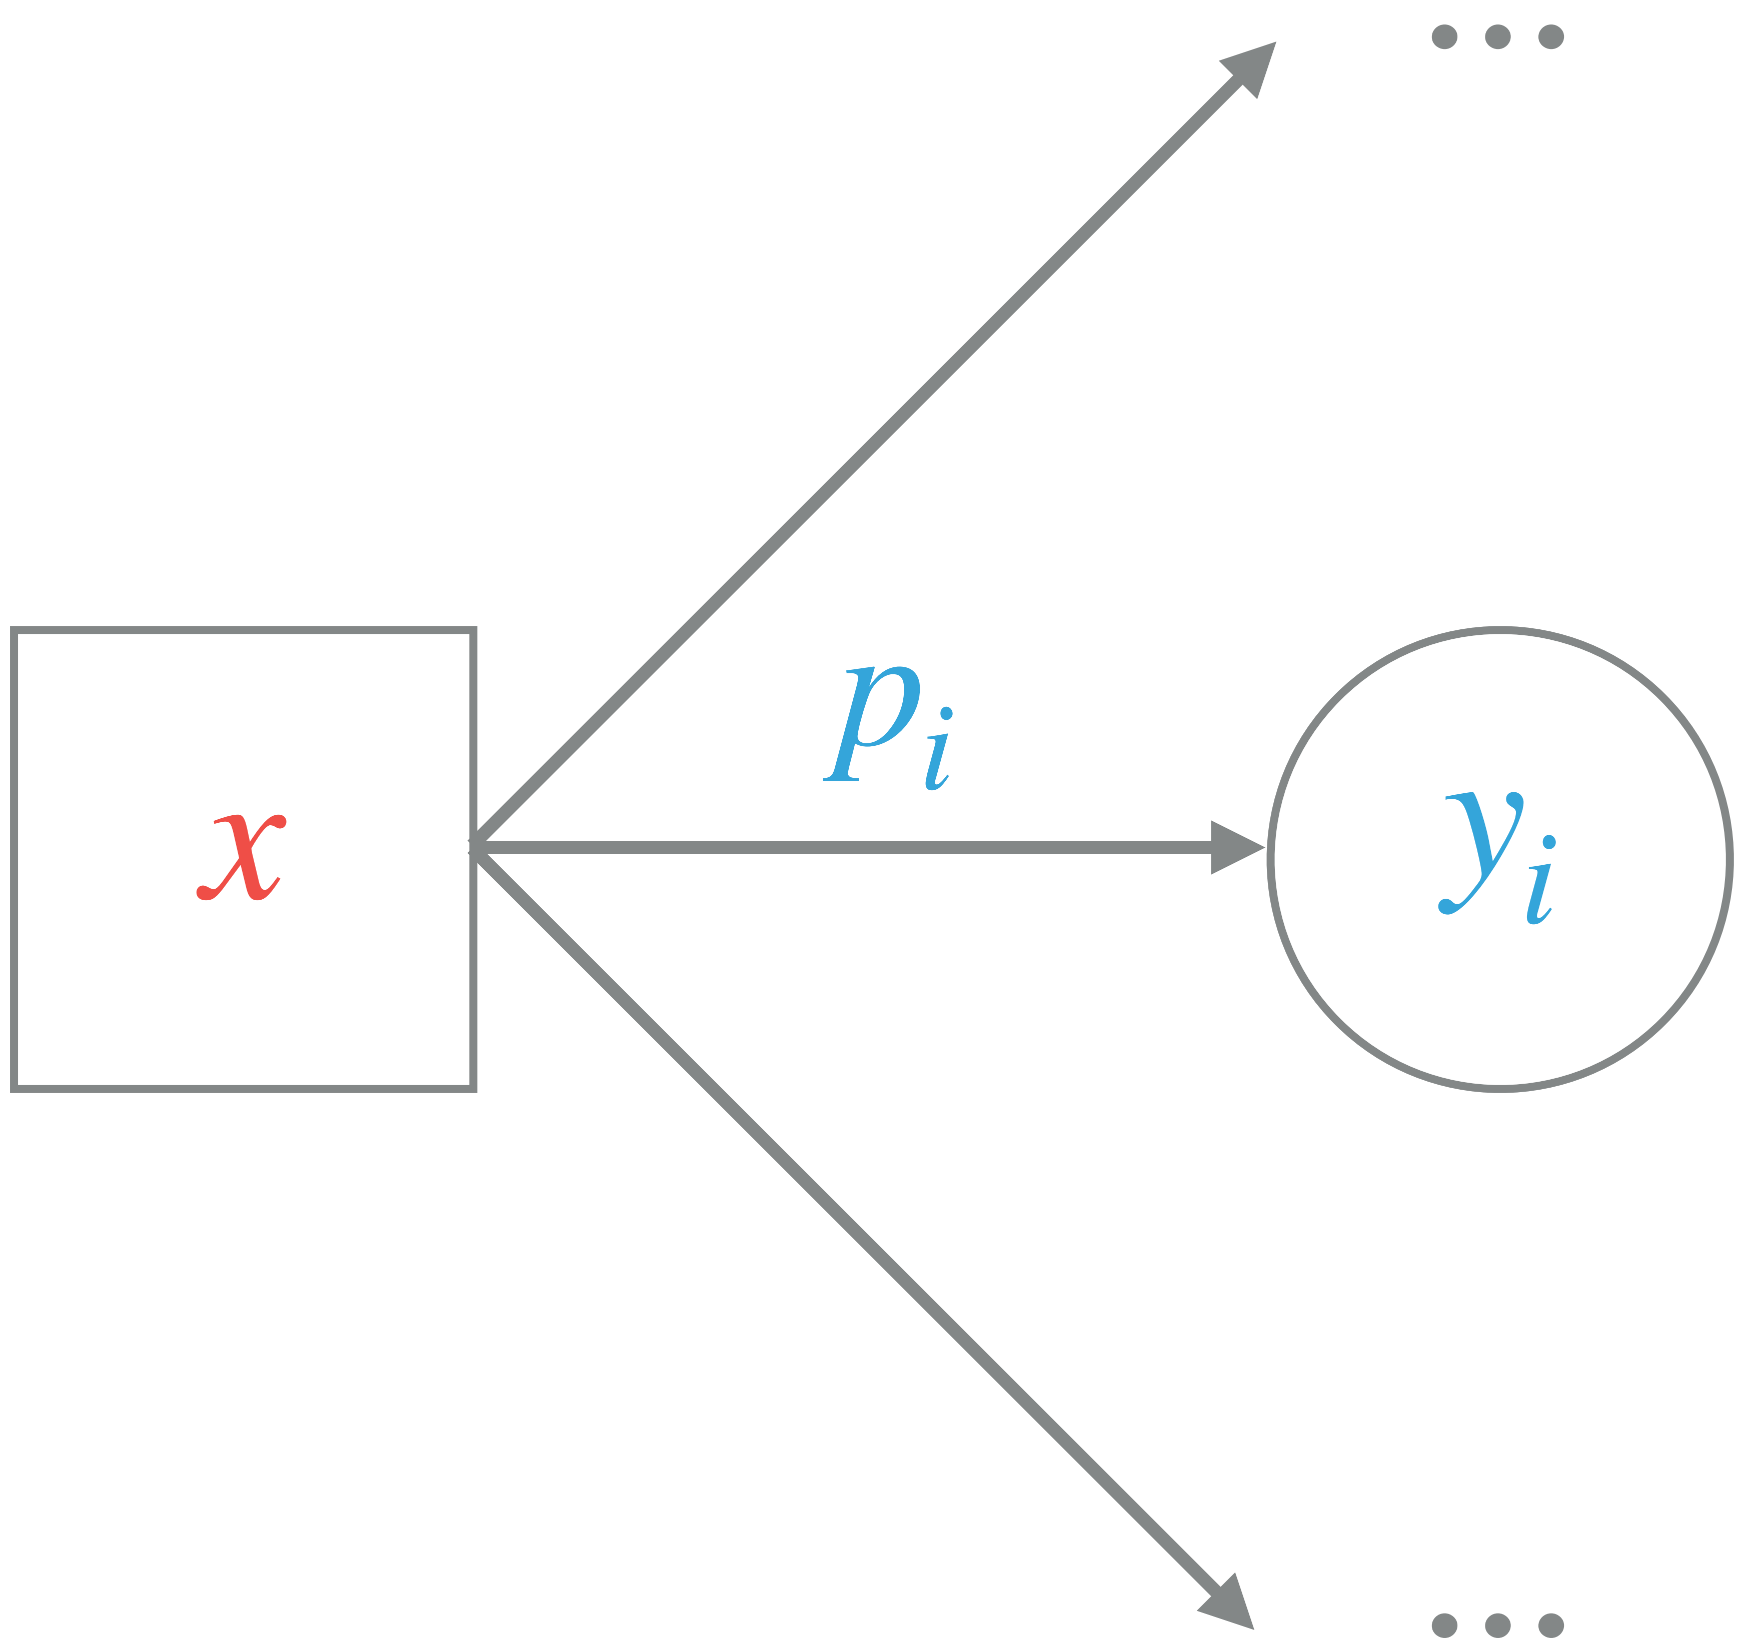

FANTASIZING ABOUT IMPROVEMENT

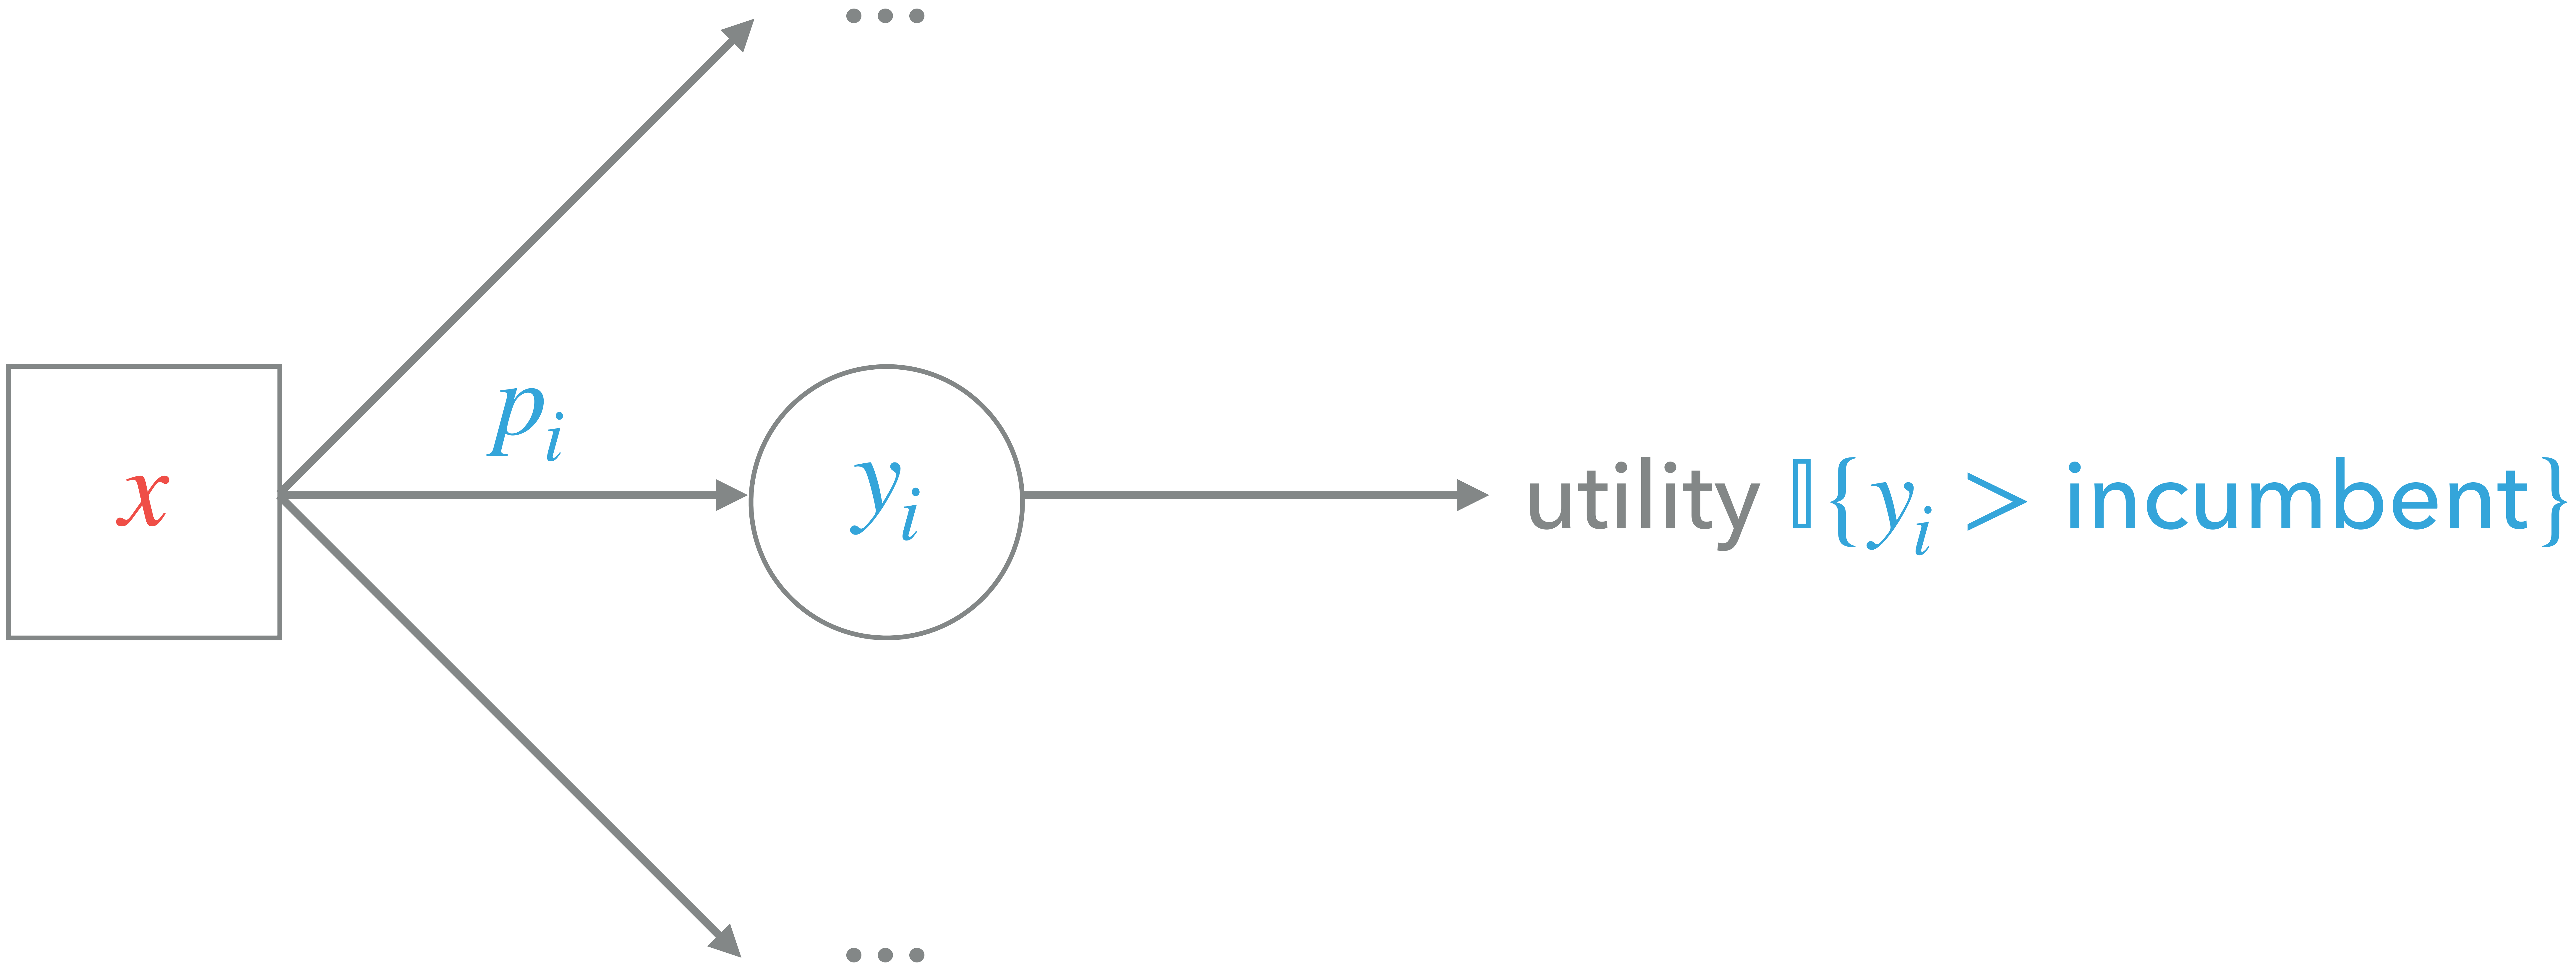

FANTASIZING ABOUT IMPROVEMENT

avg. utility of x

$$\mathbb{E}_{y_j}\left[\mathbb{1}\{y_i > \text{incumbent}\}\right]$$

$$= \Pr\left(y > \text{incumbent}\right)$$

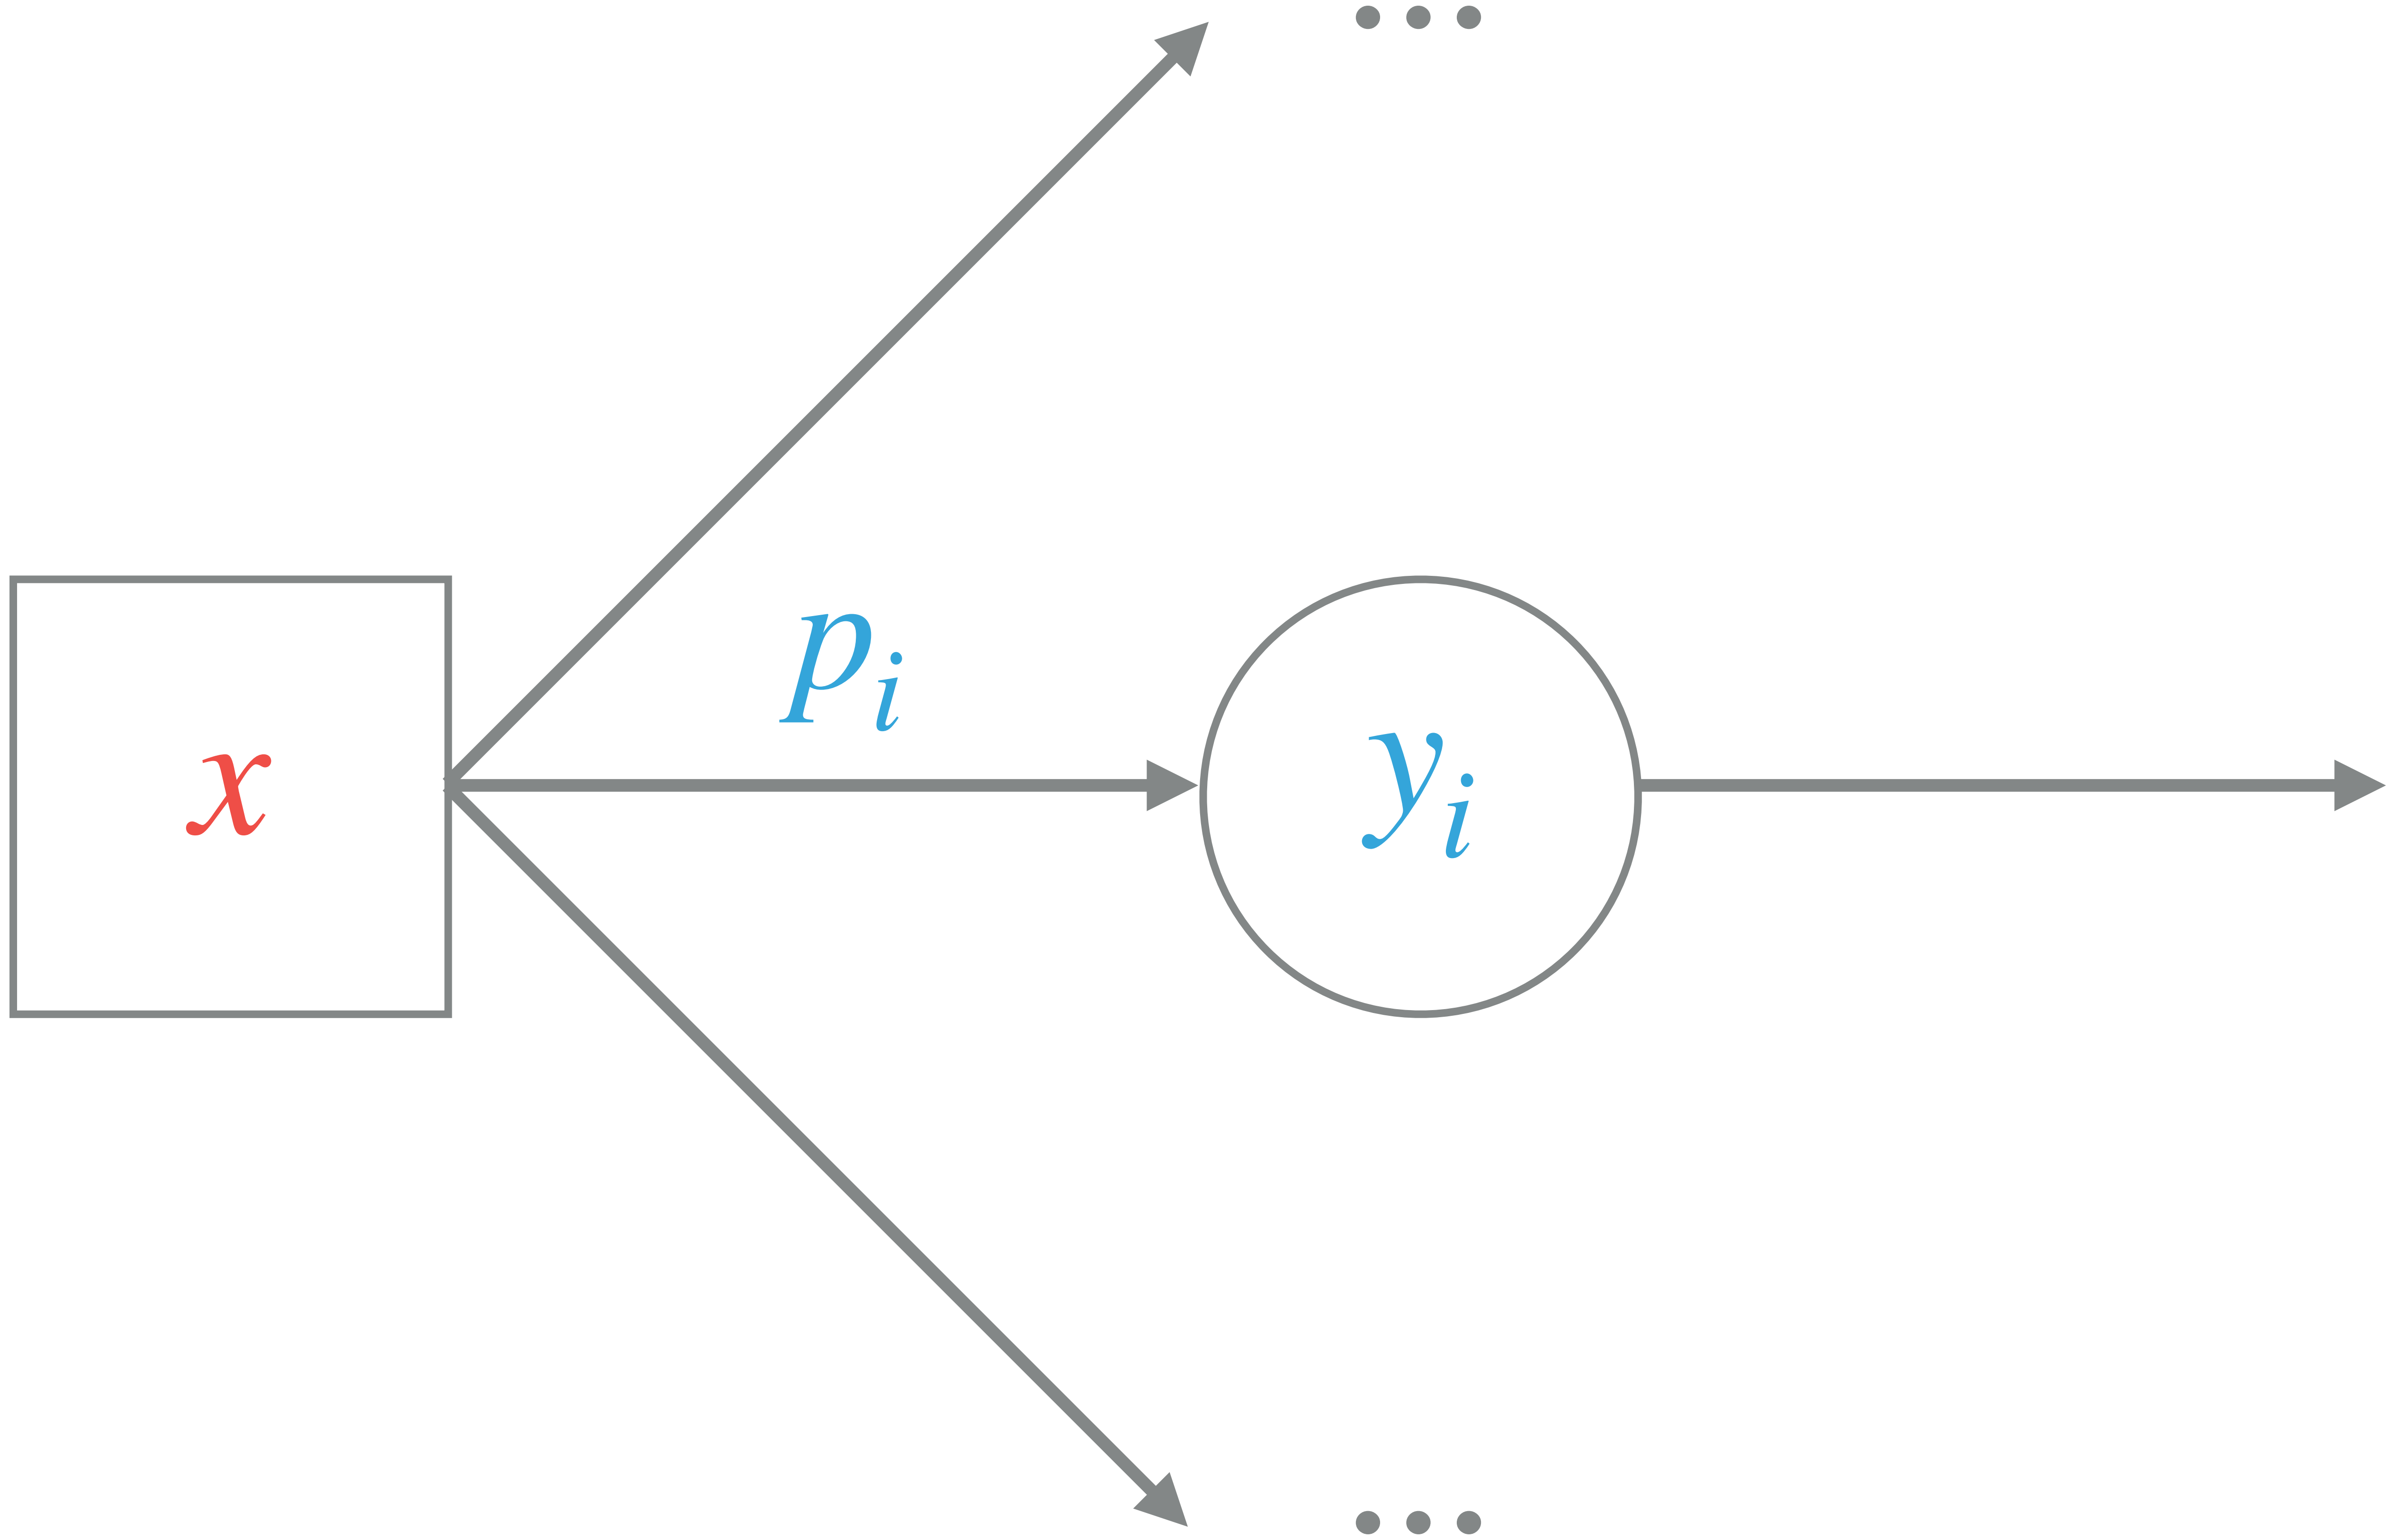

utility $\mathbb{1}\{y_i > \text{incumbent}\}$

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

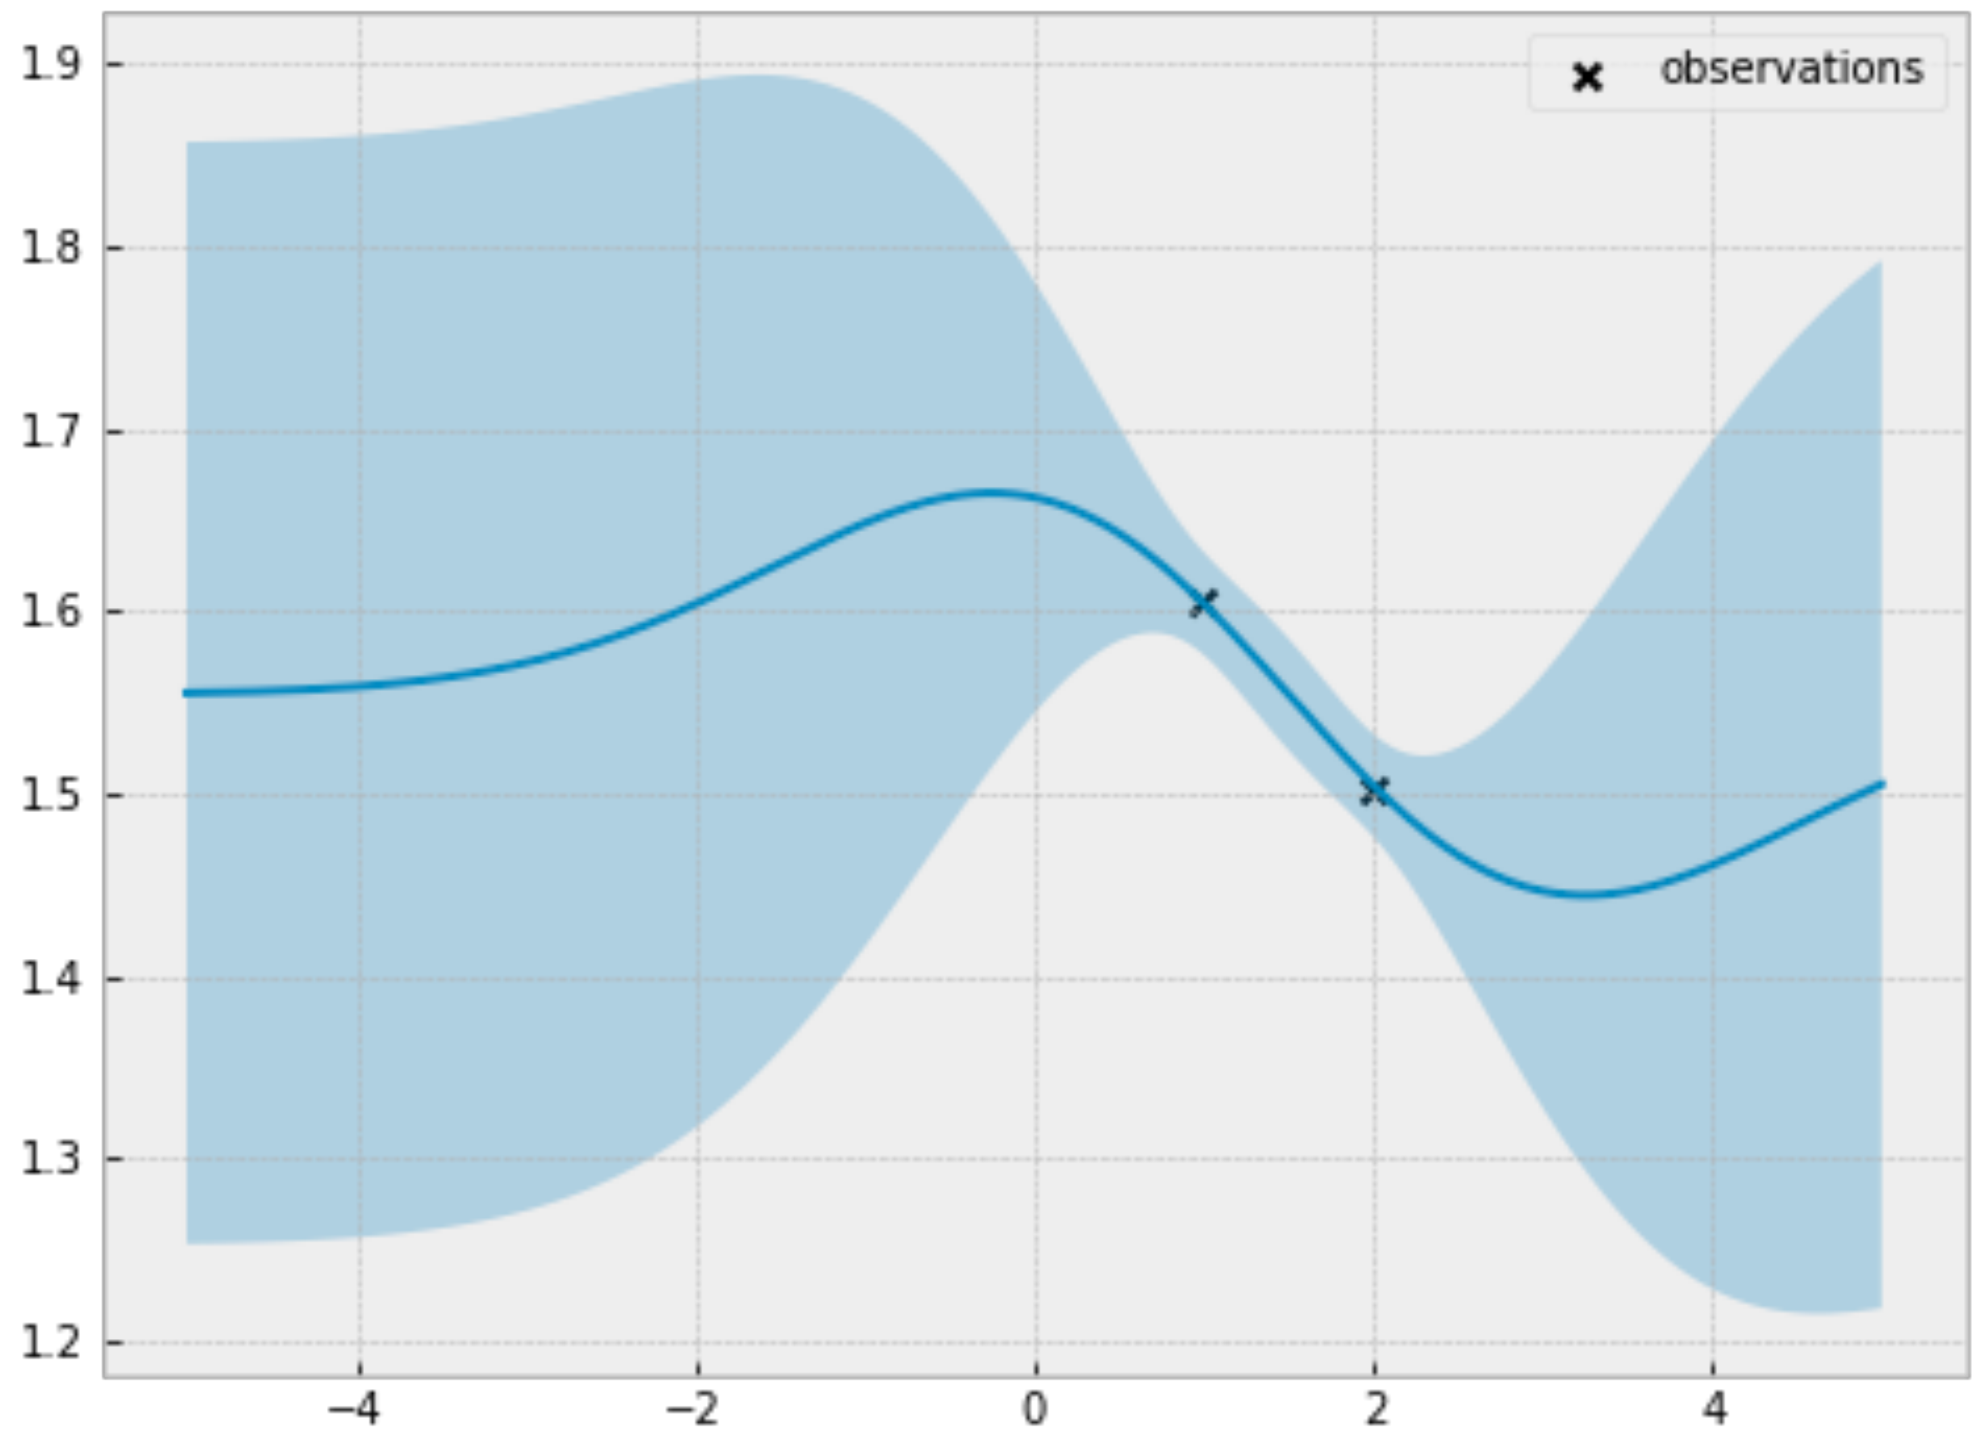

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

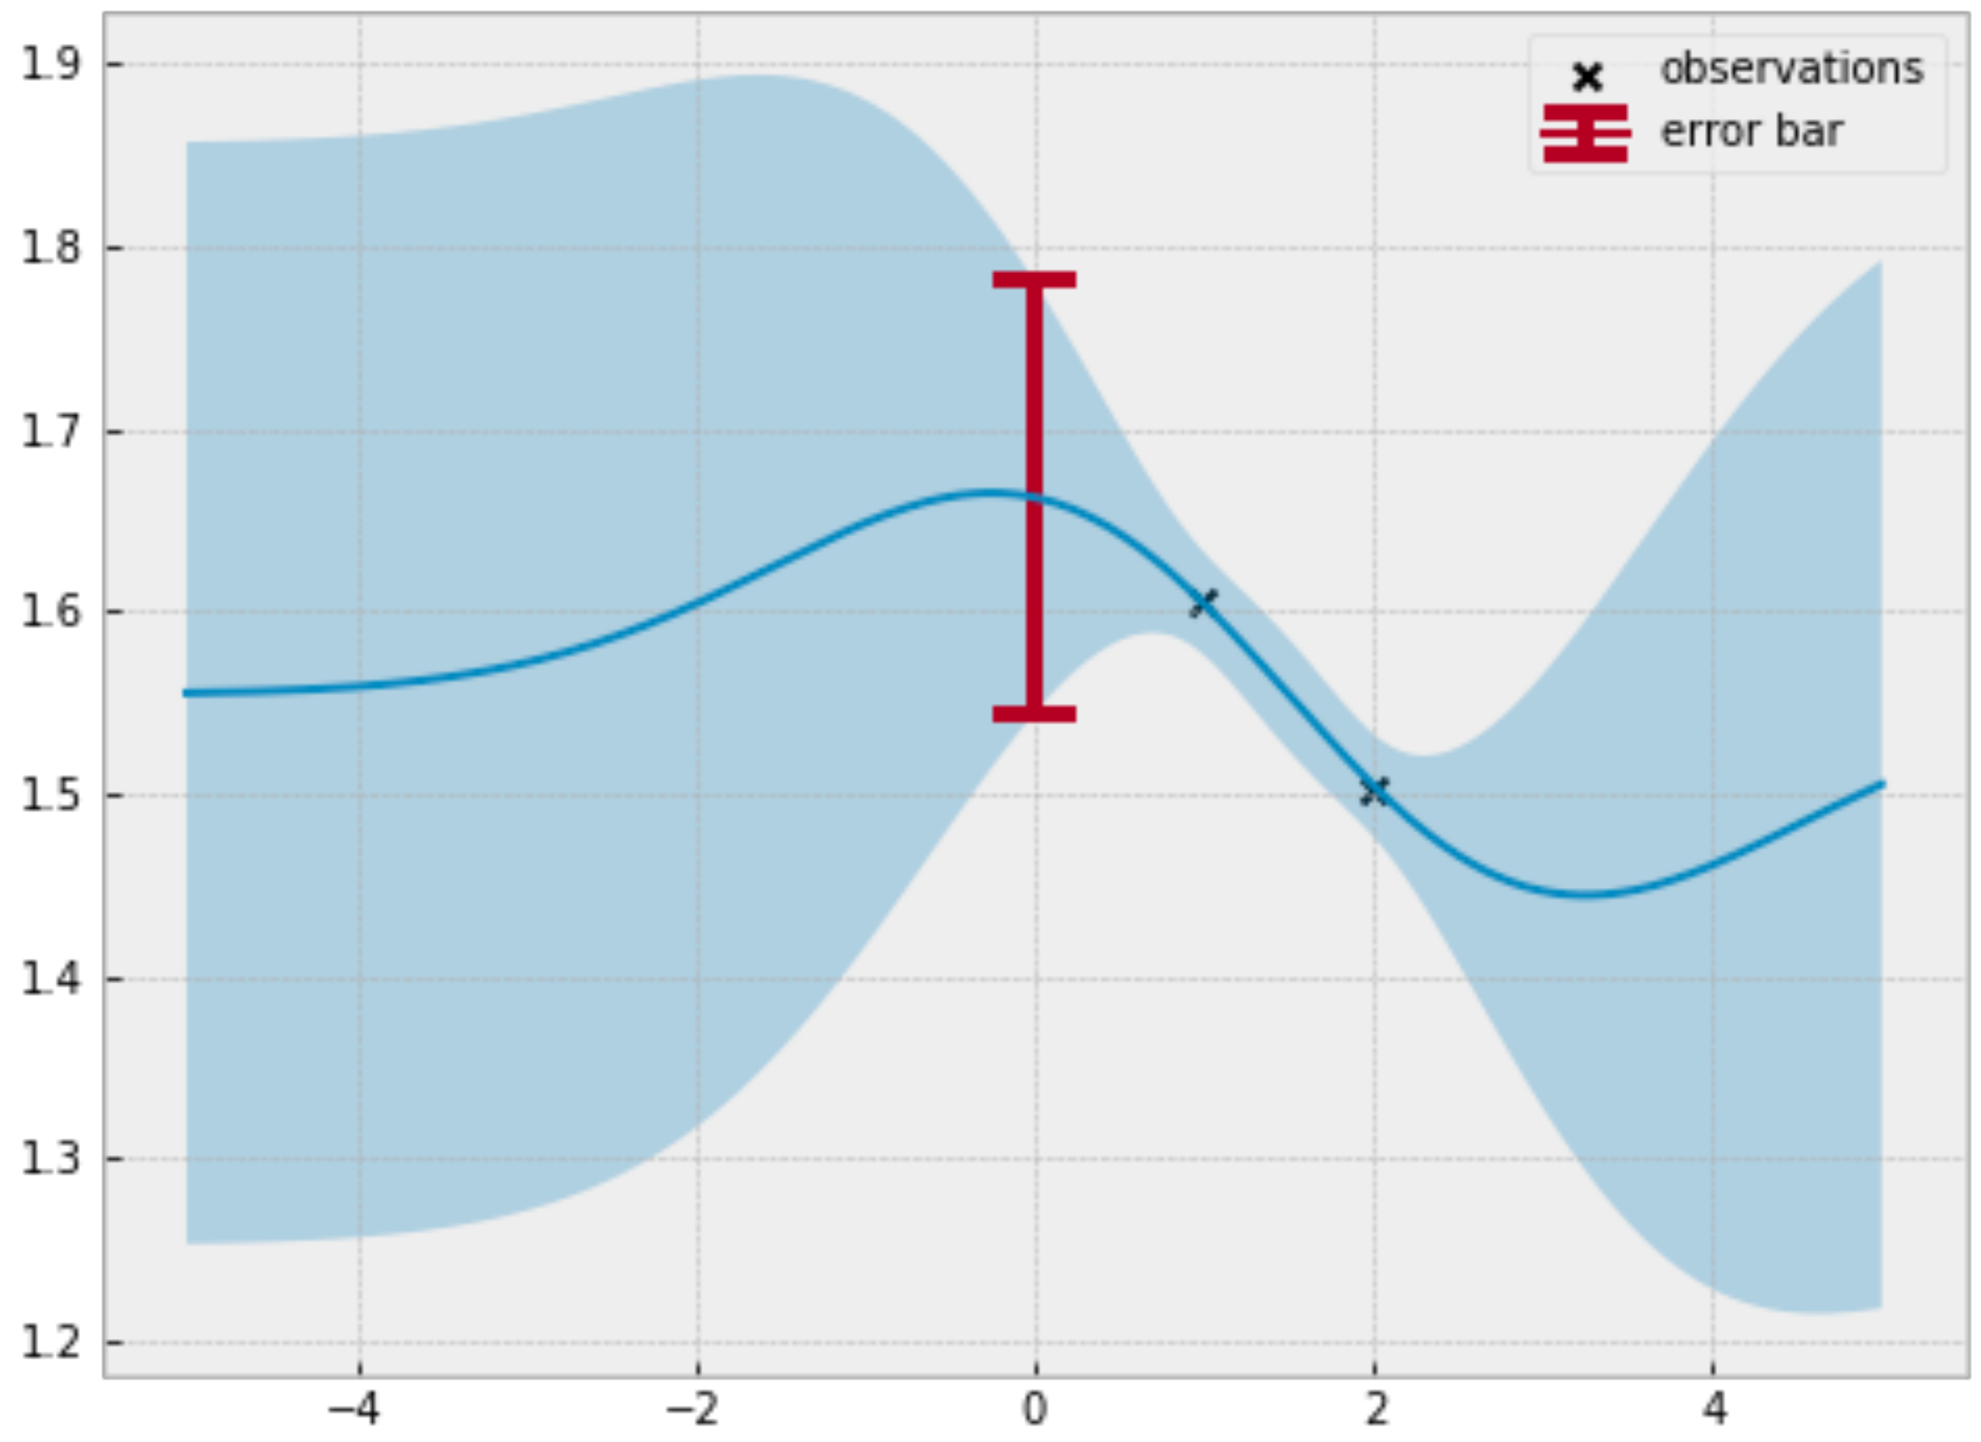

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

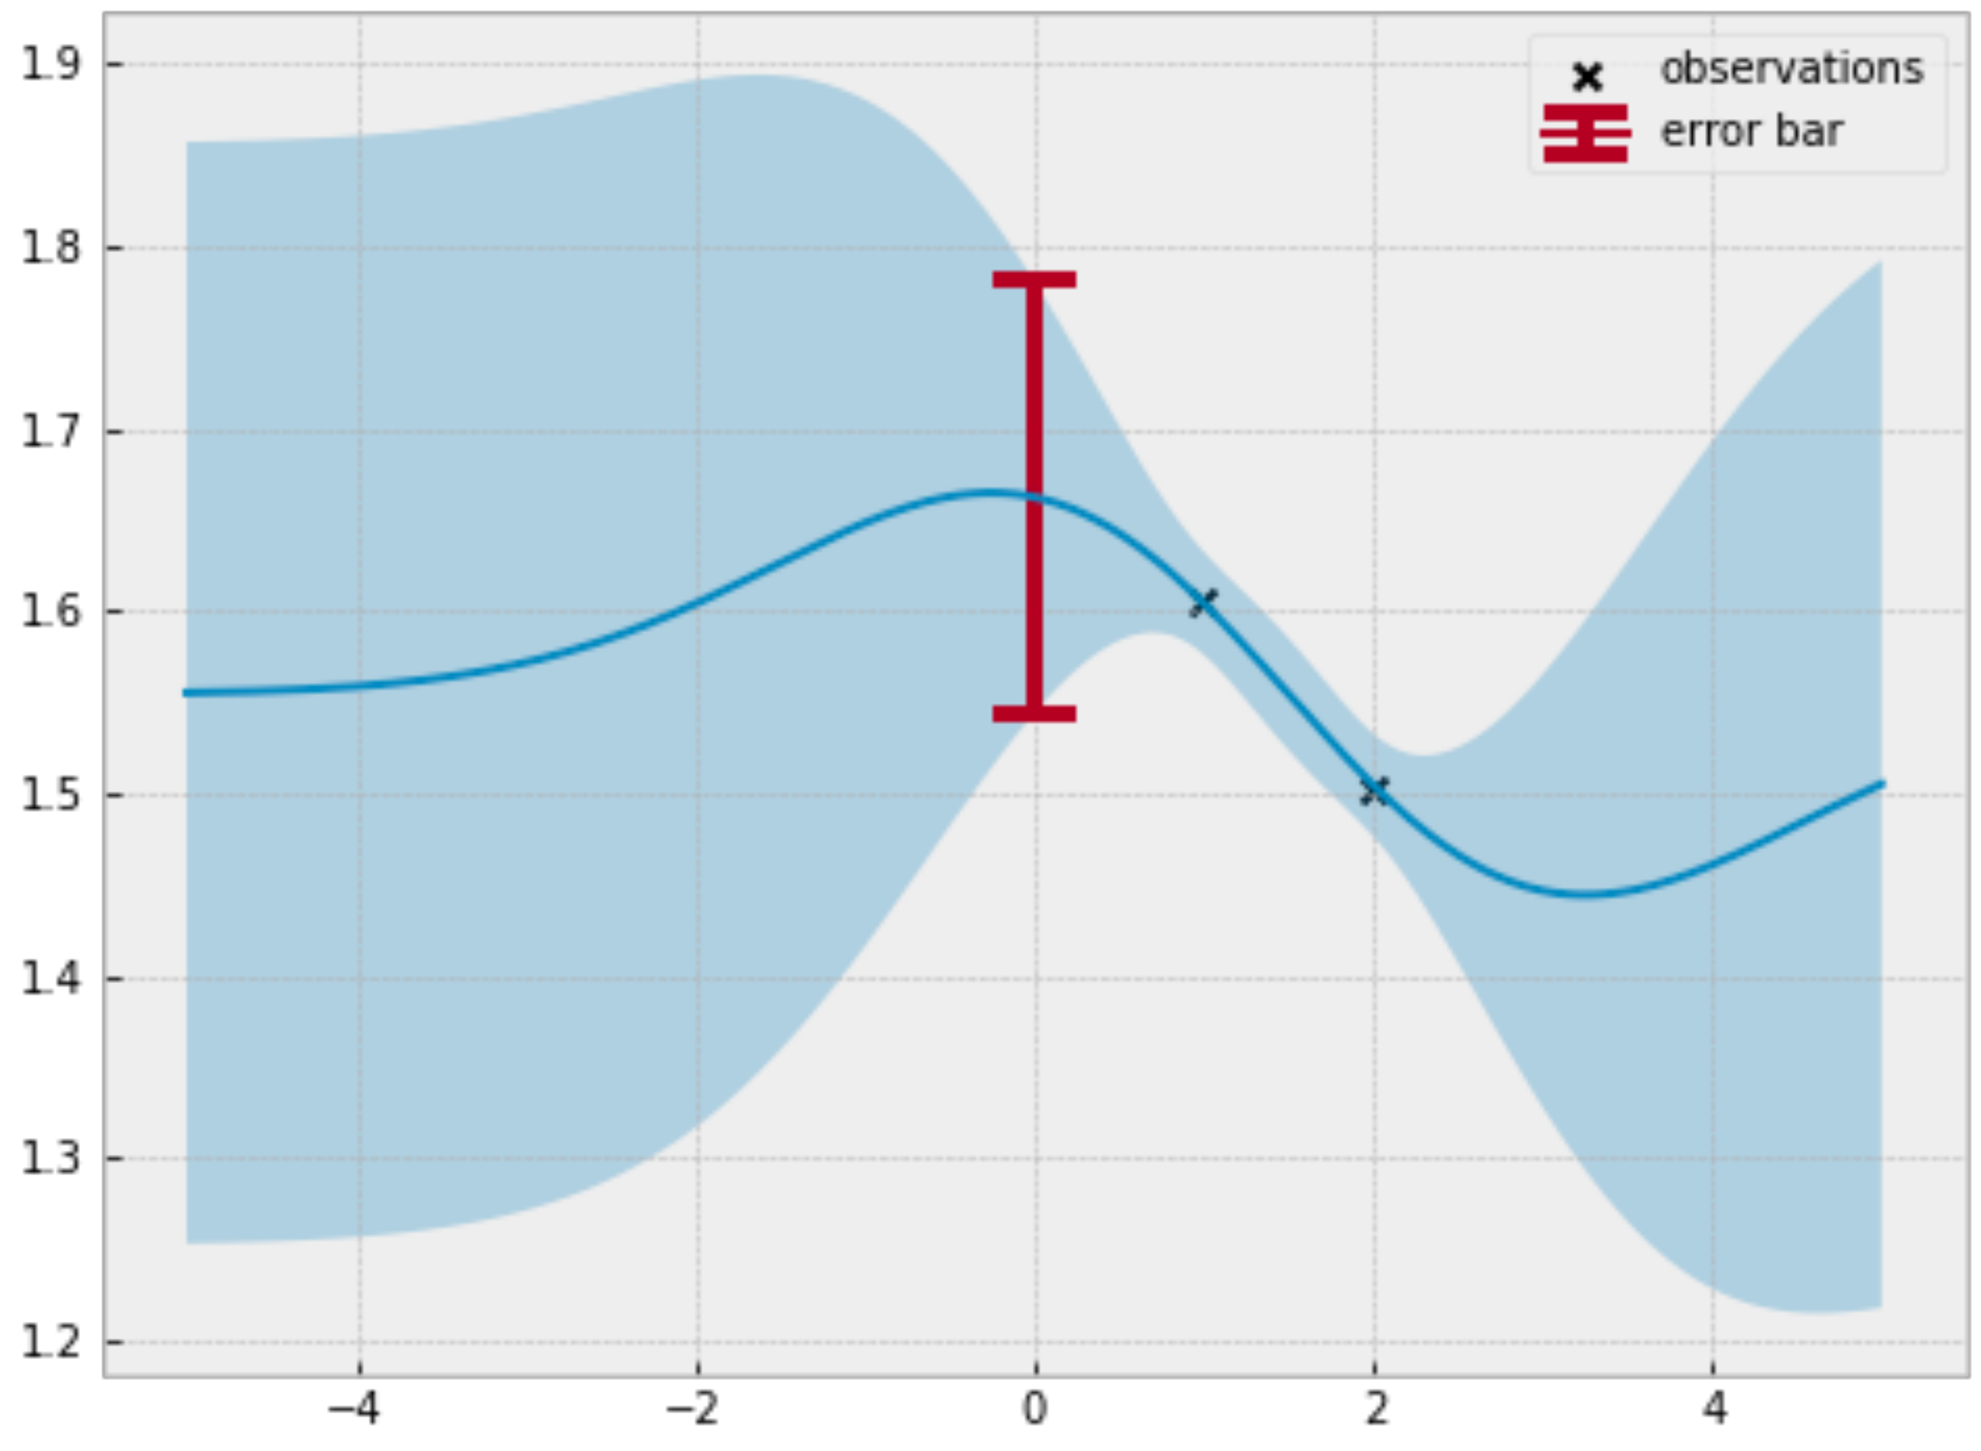

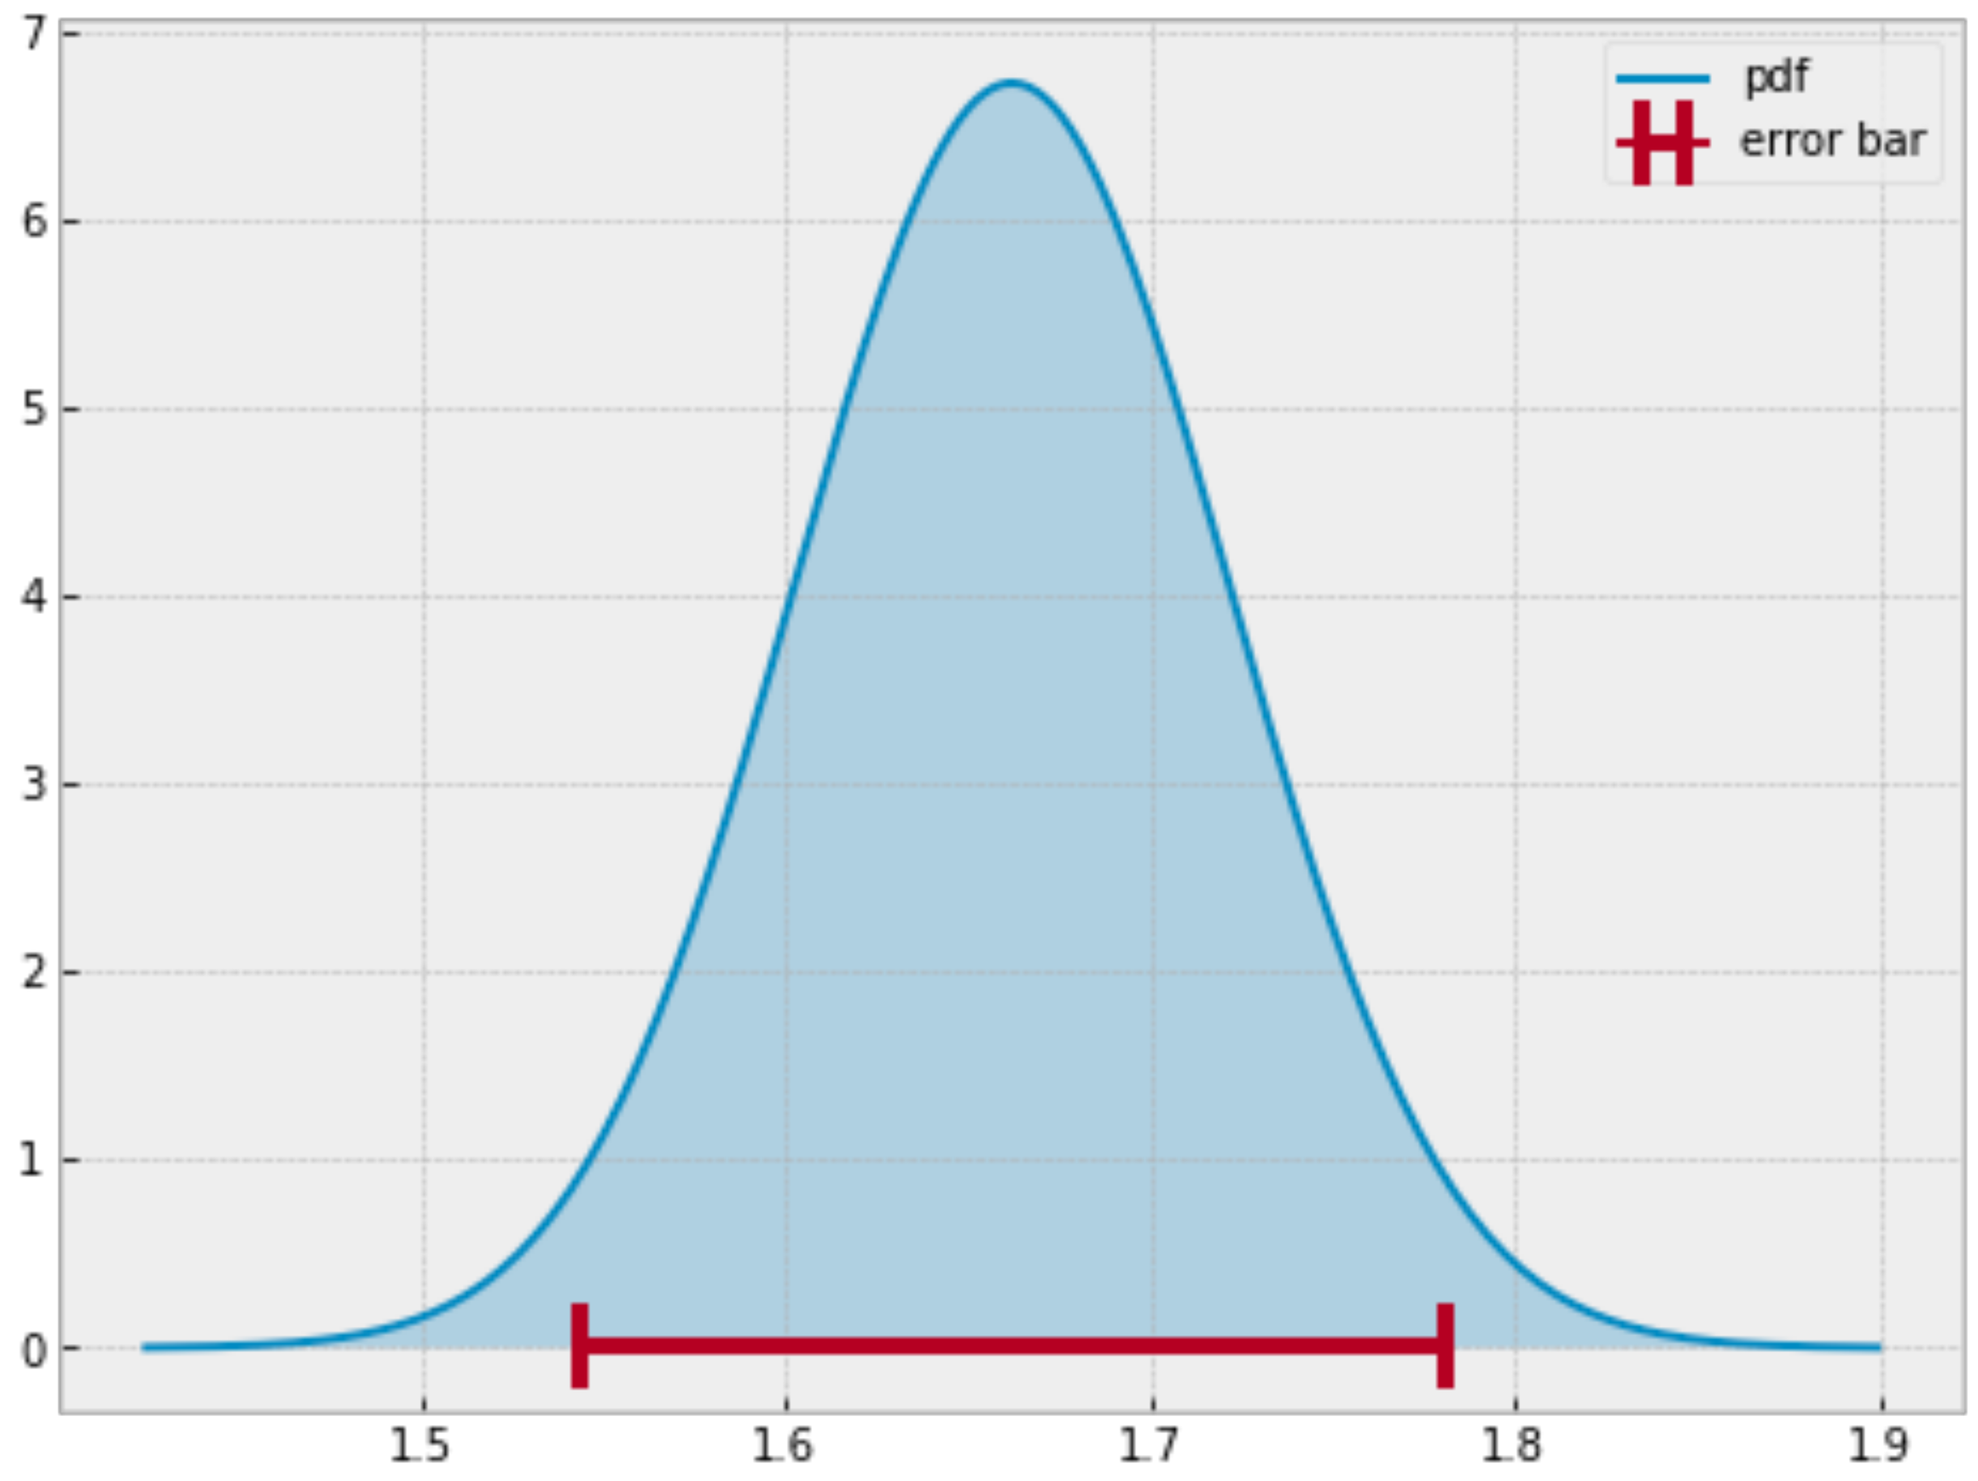

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

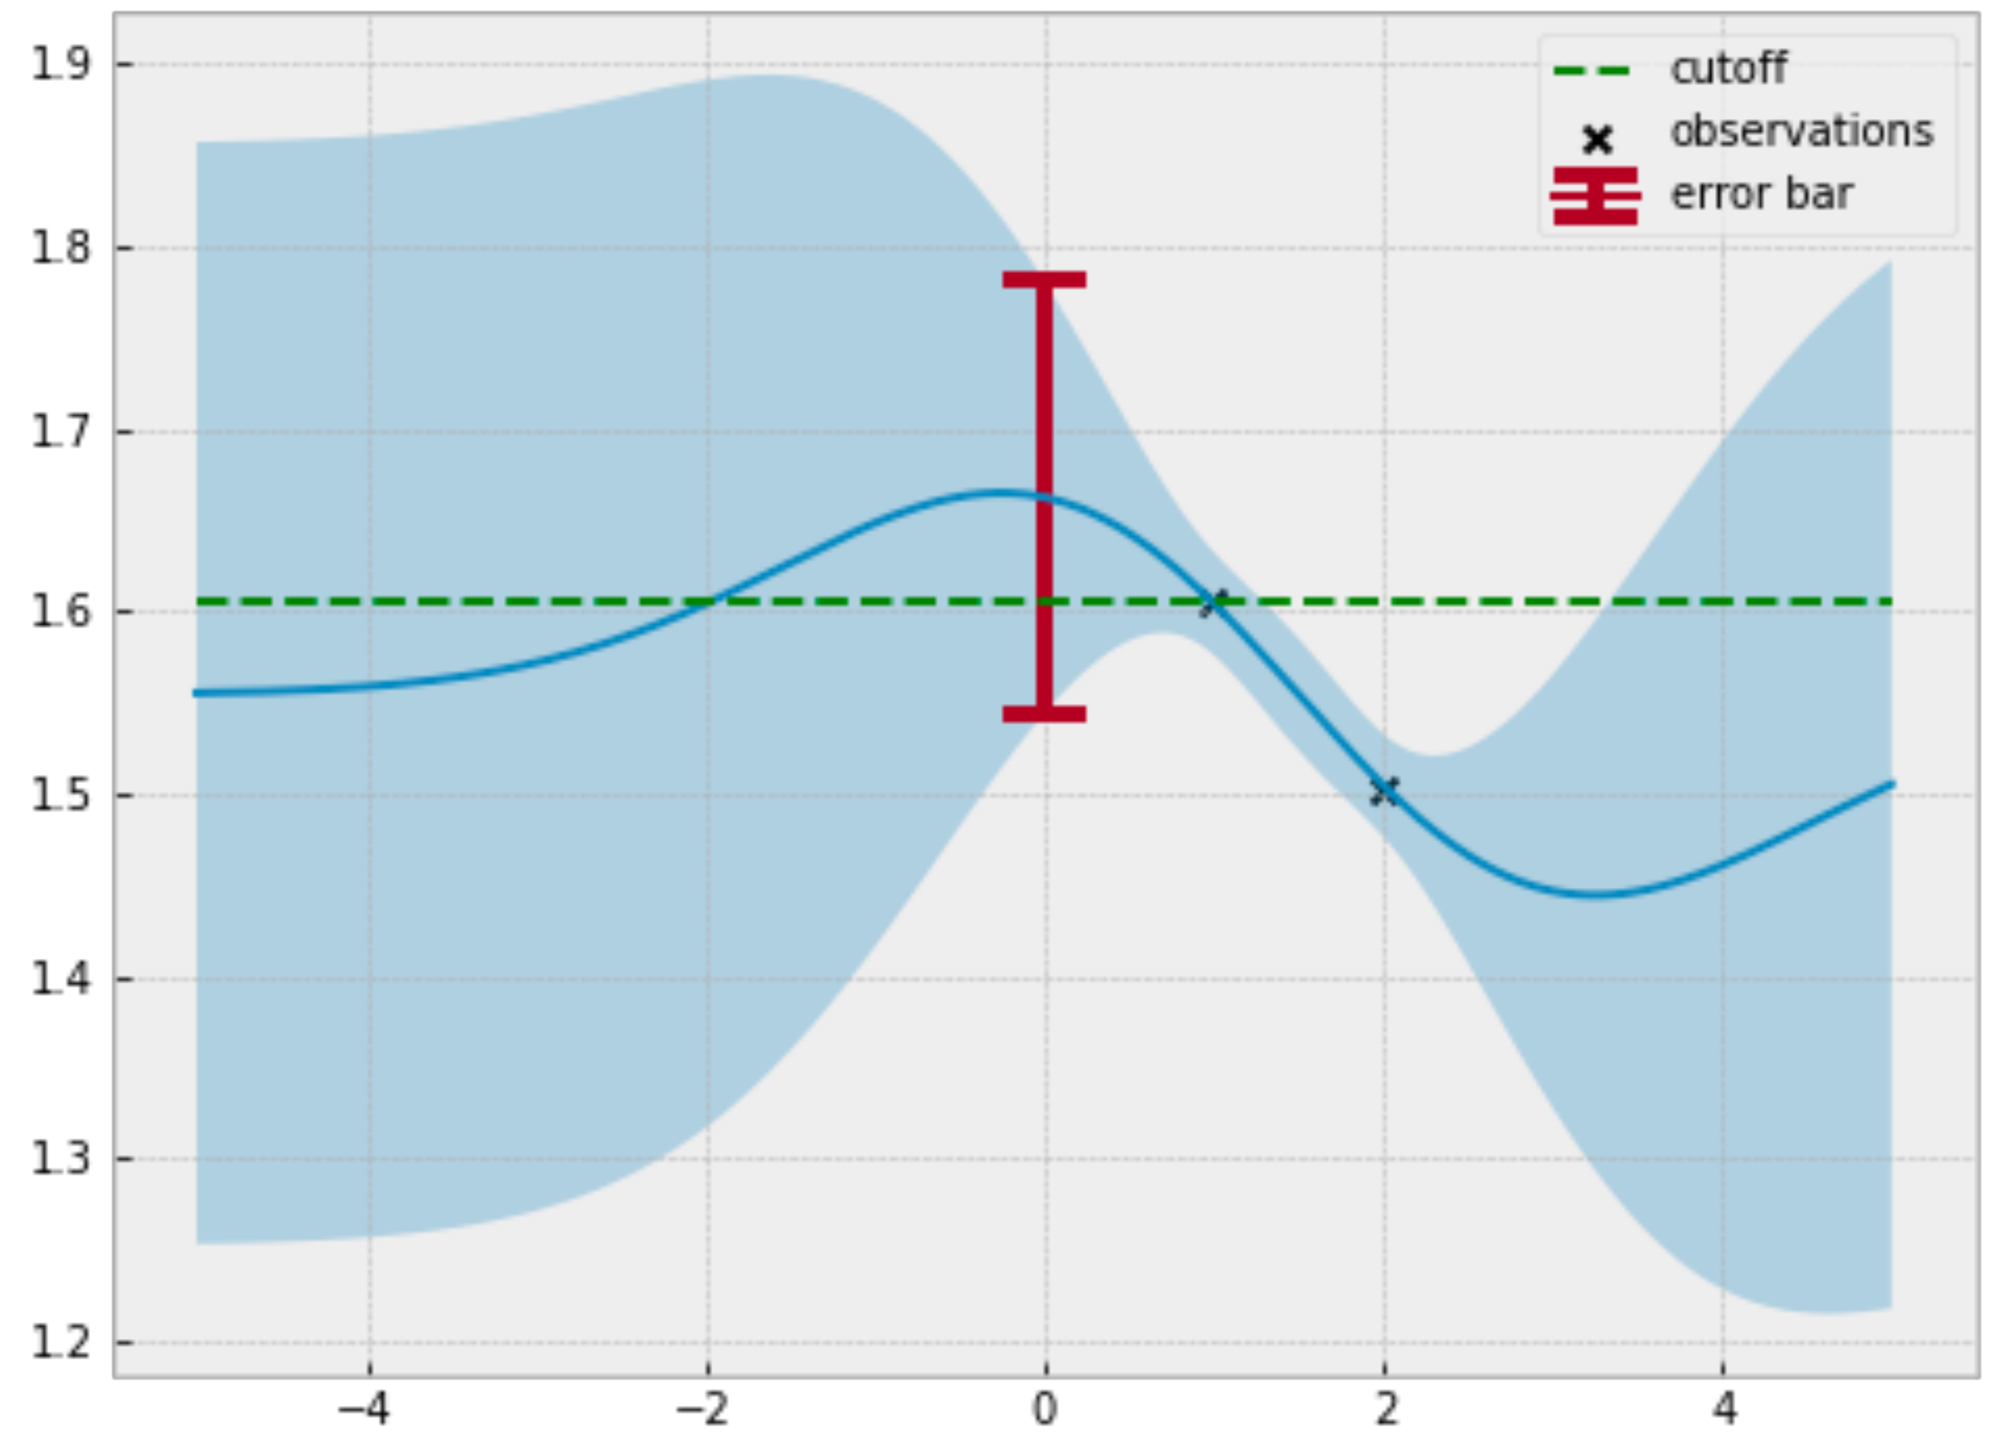

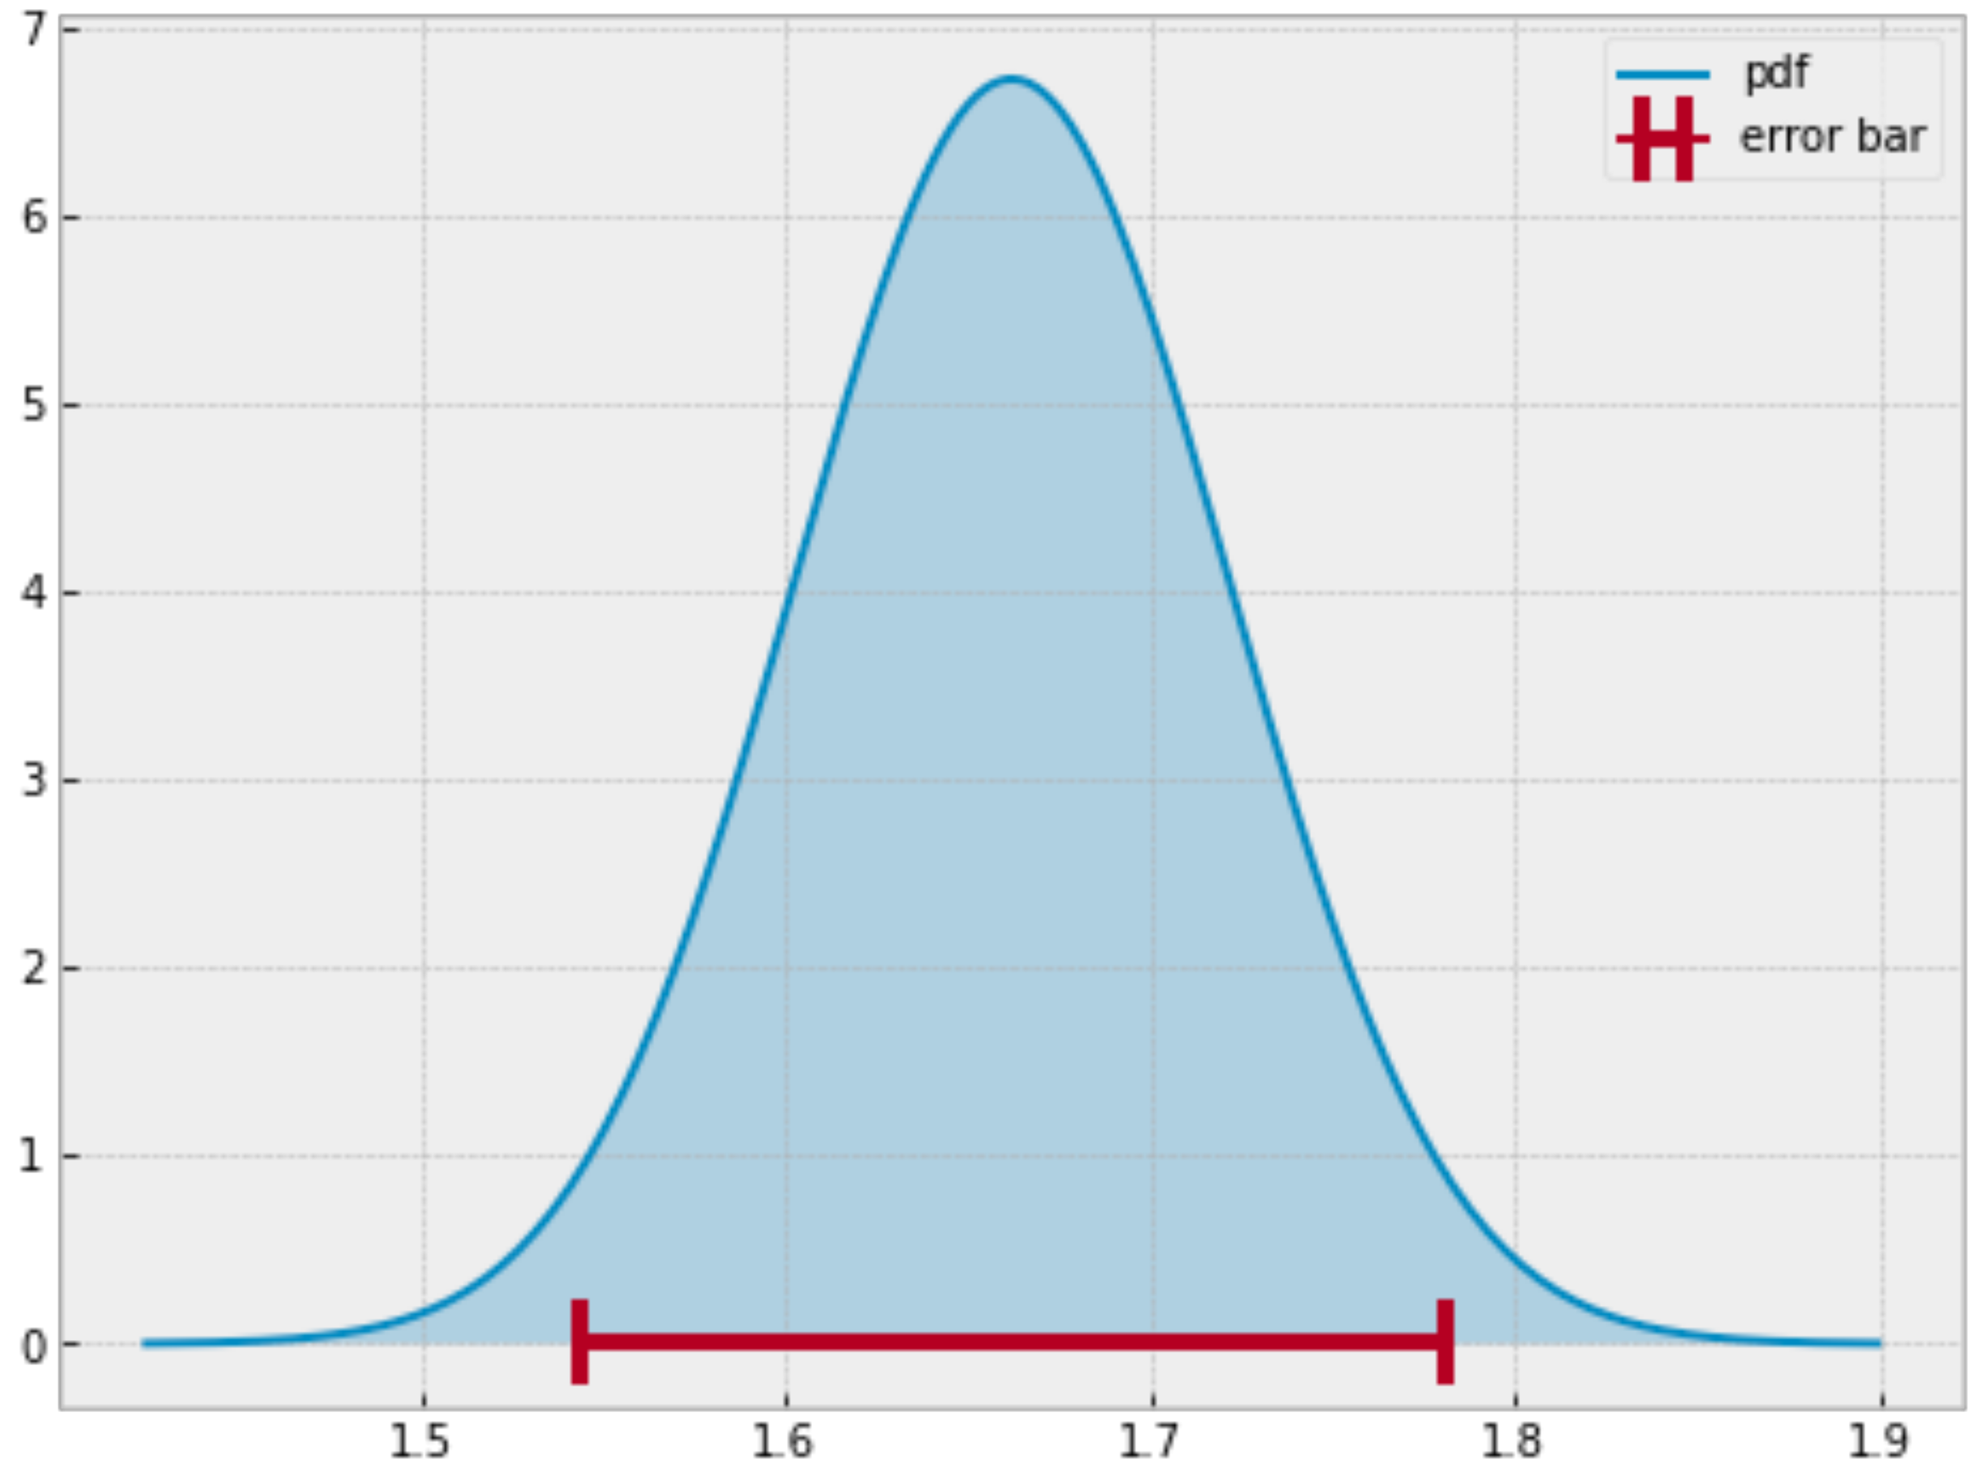

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

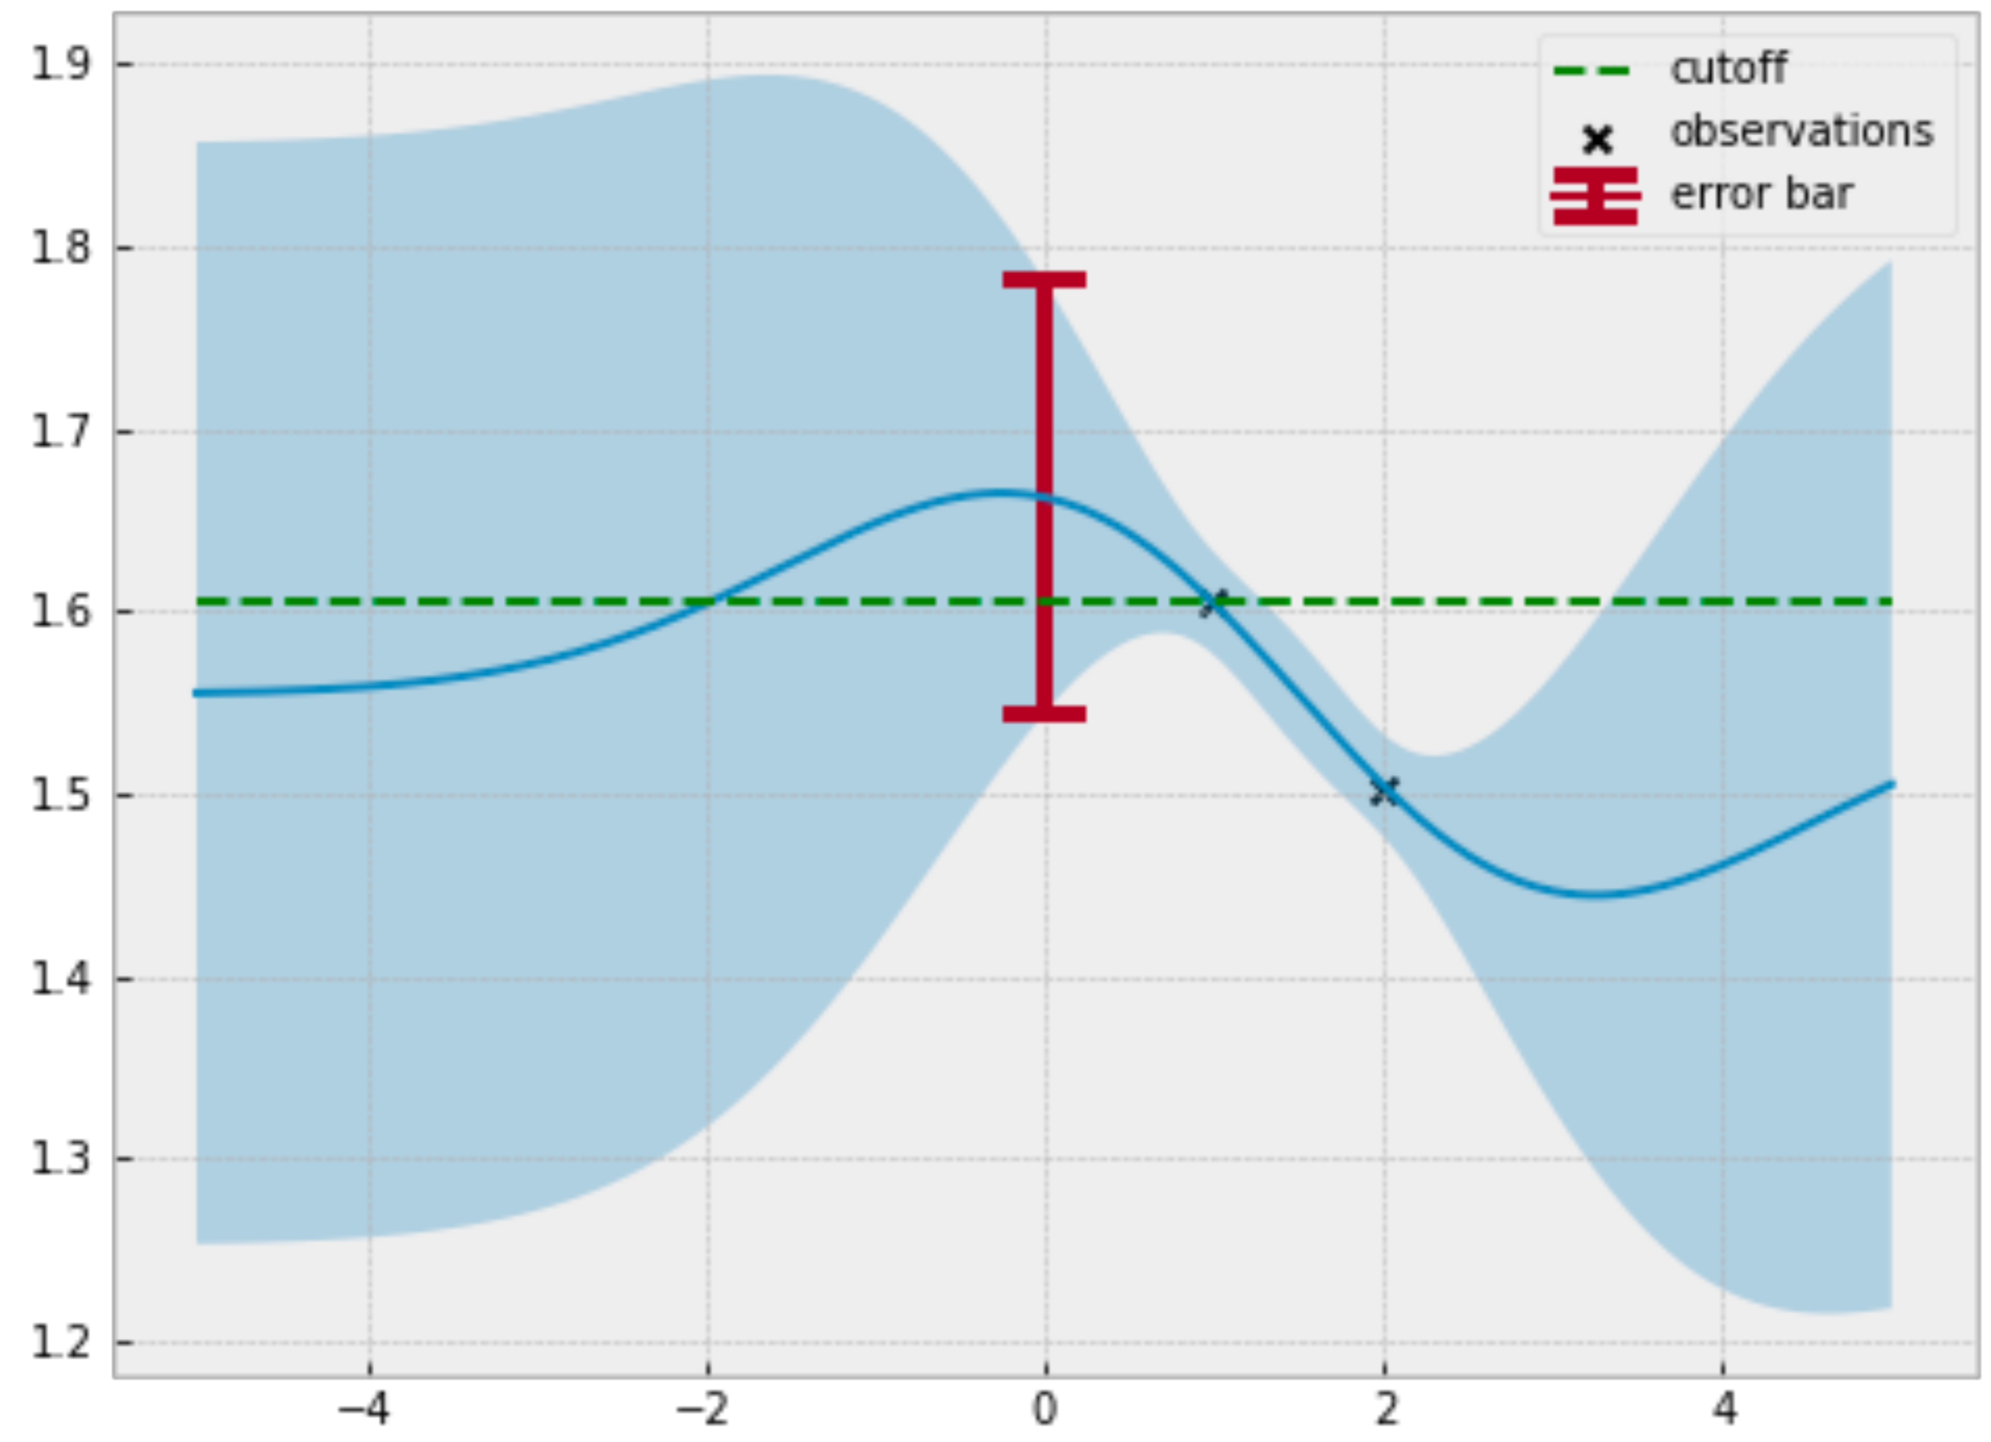

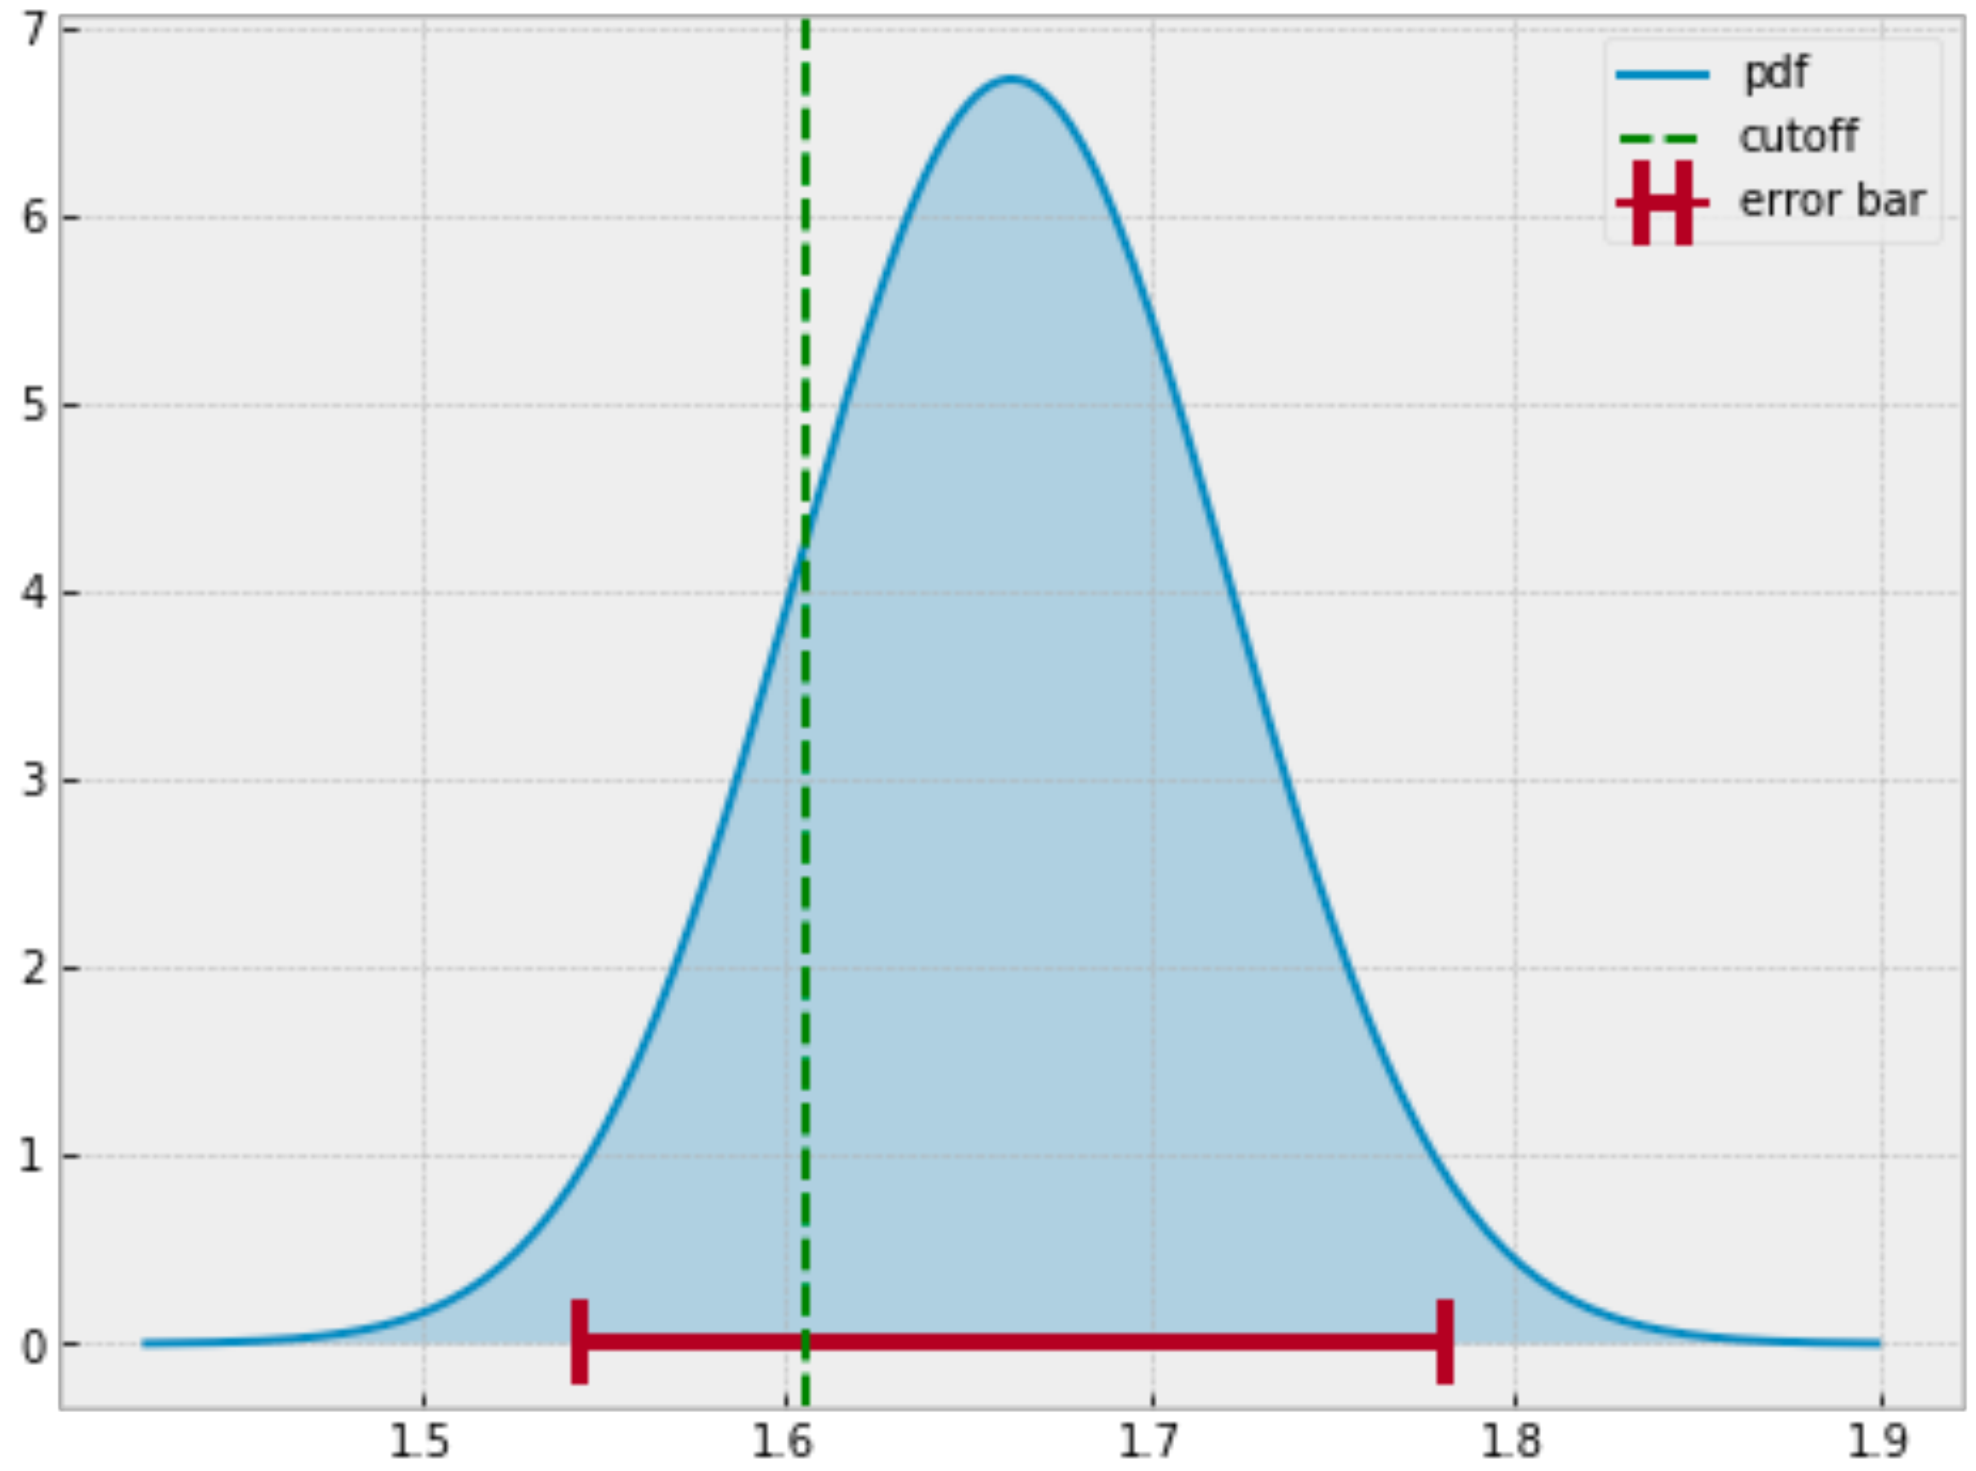

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

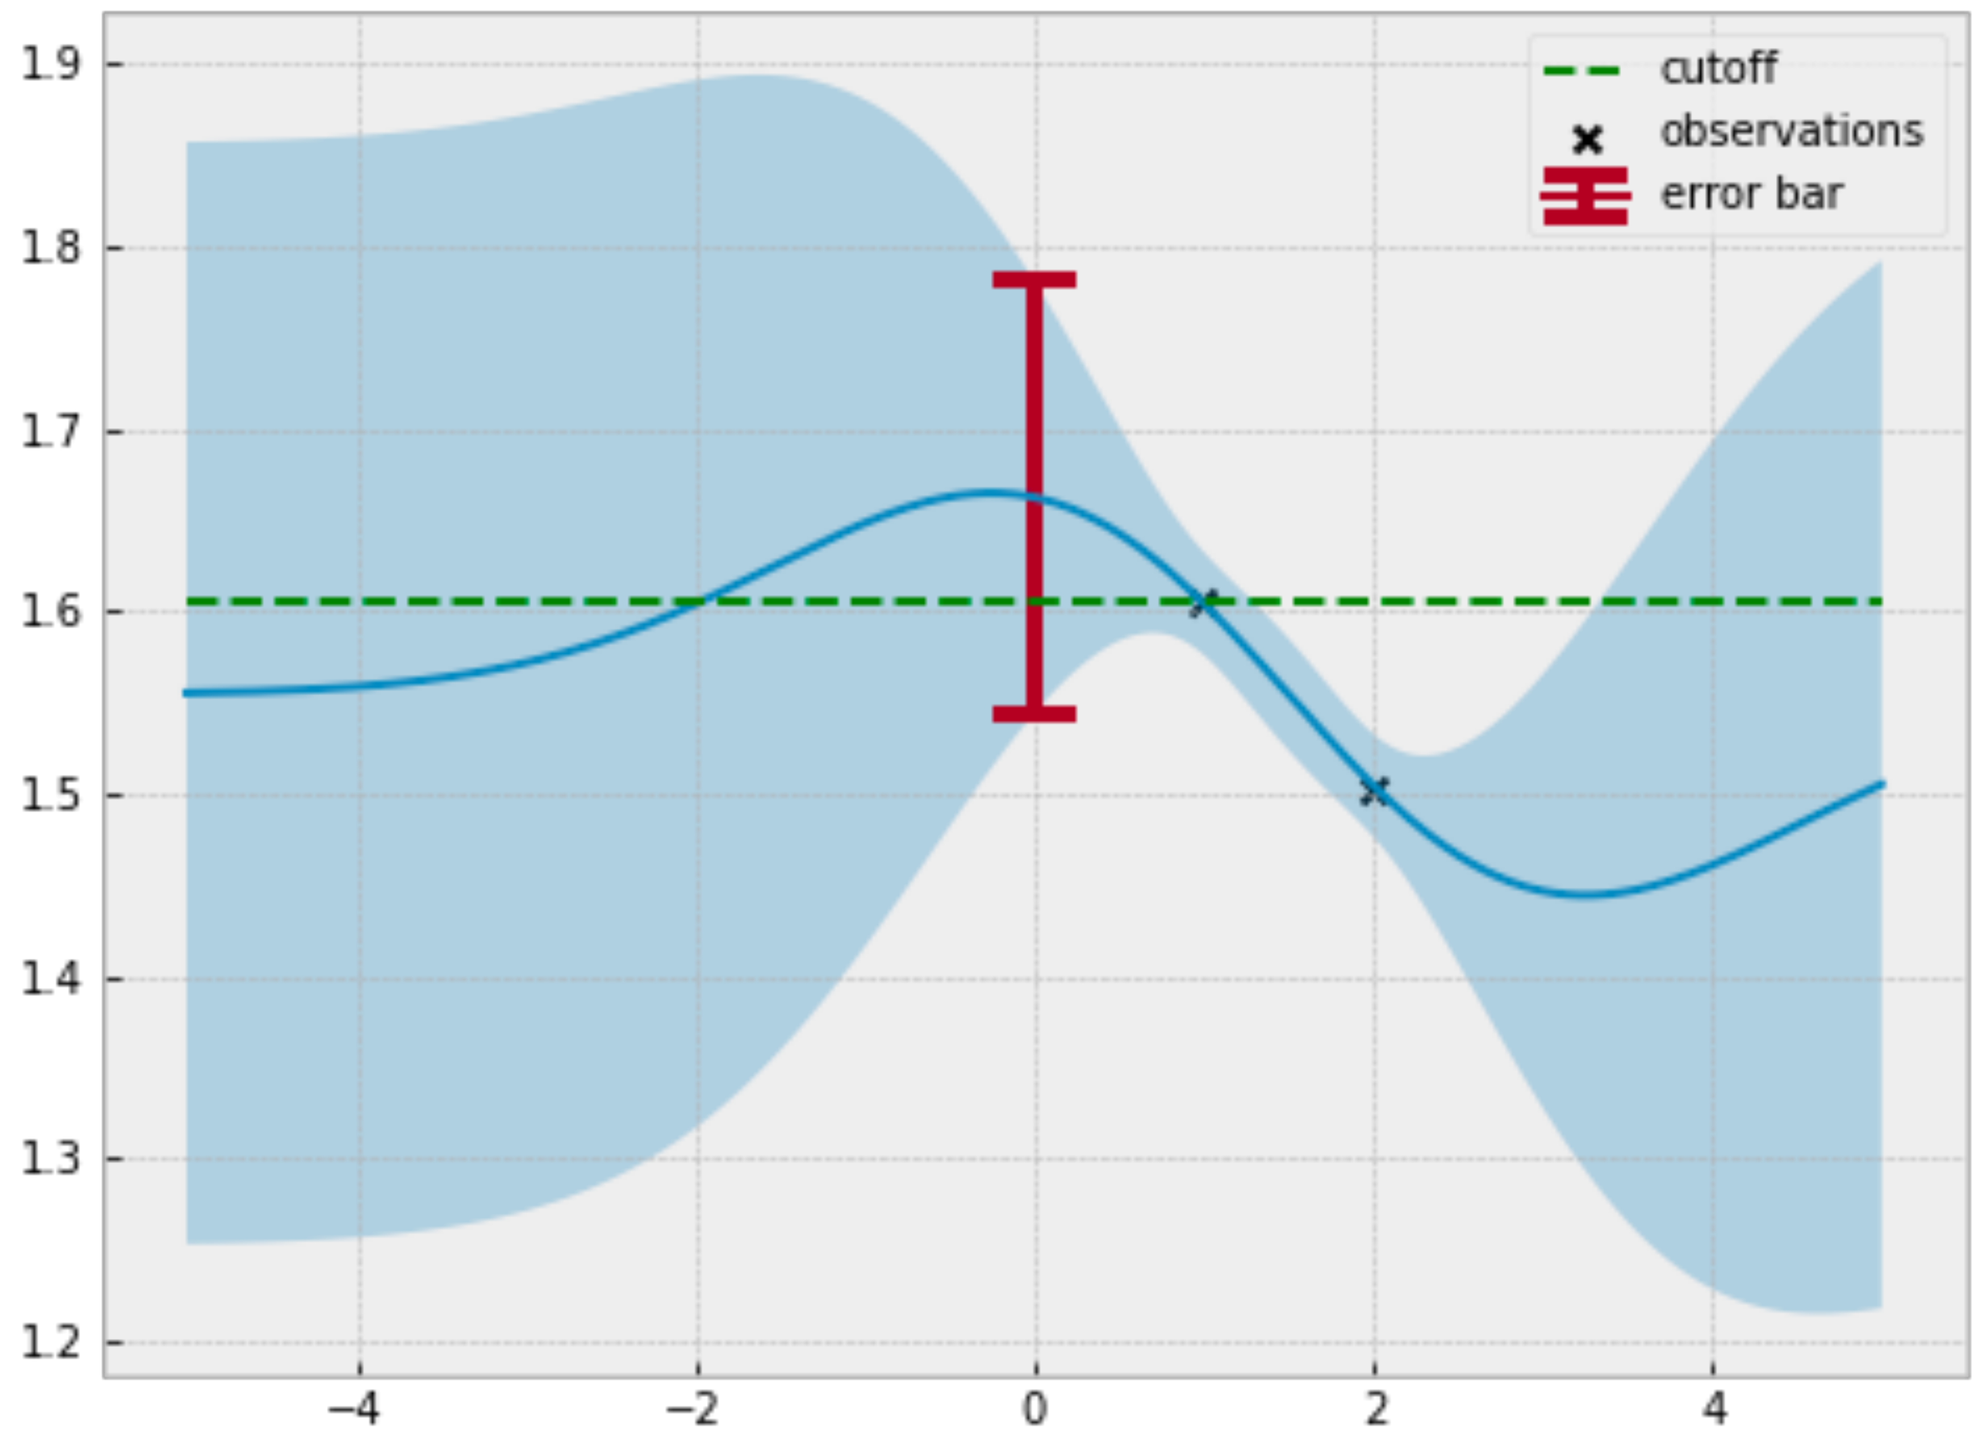

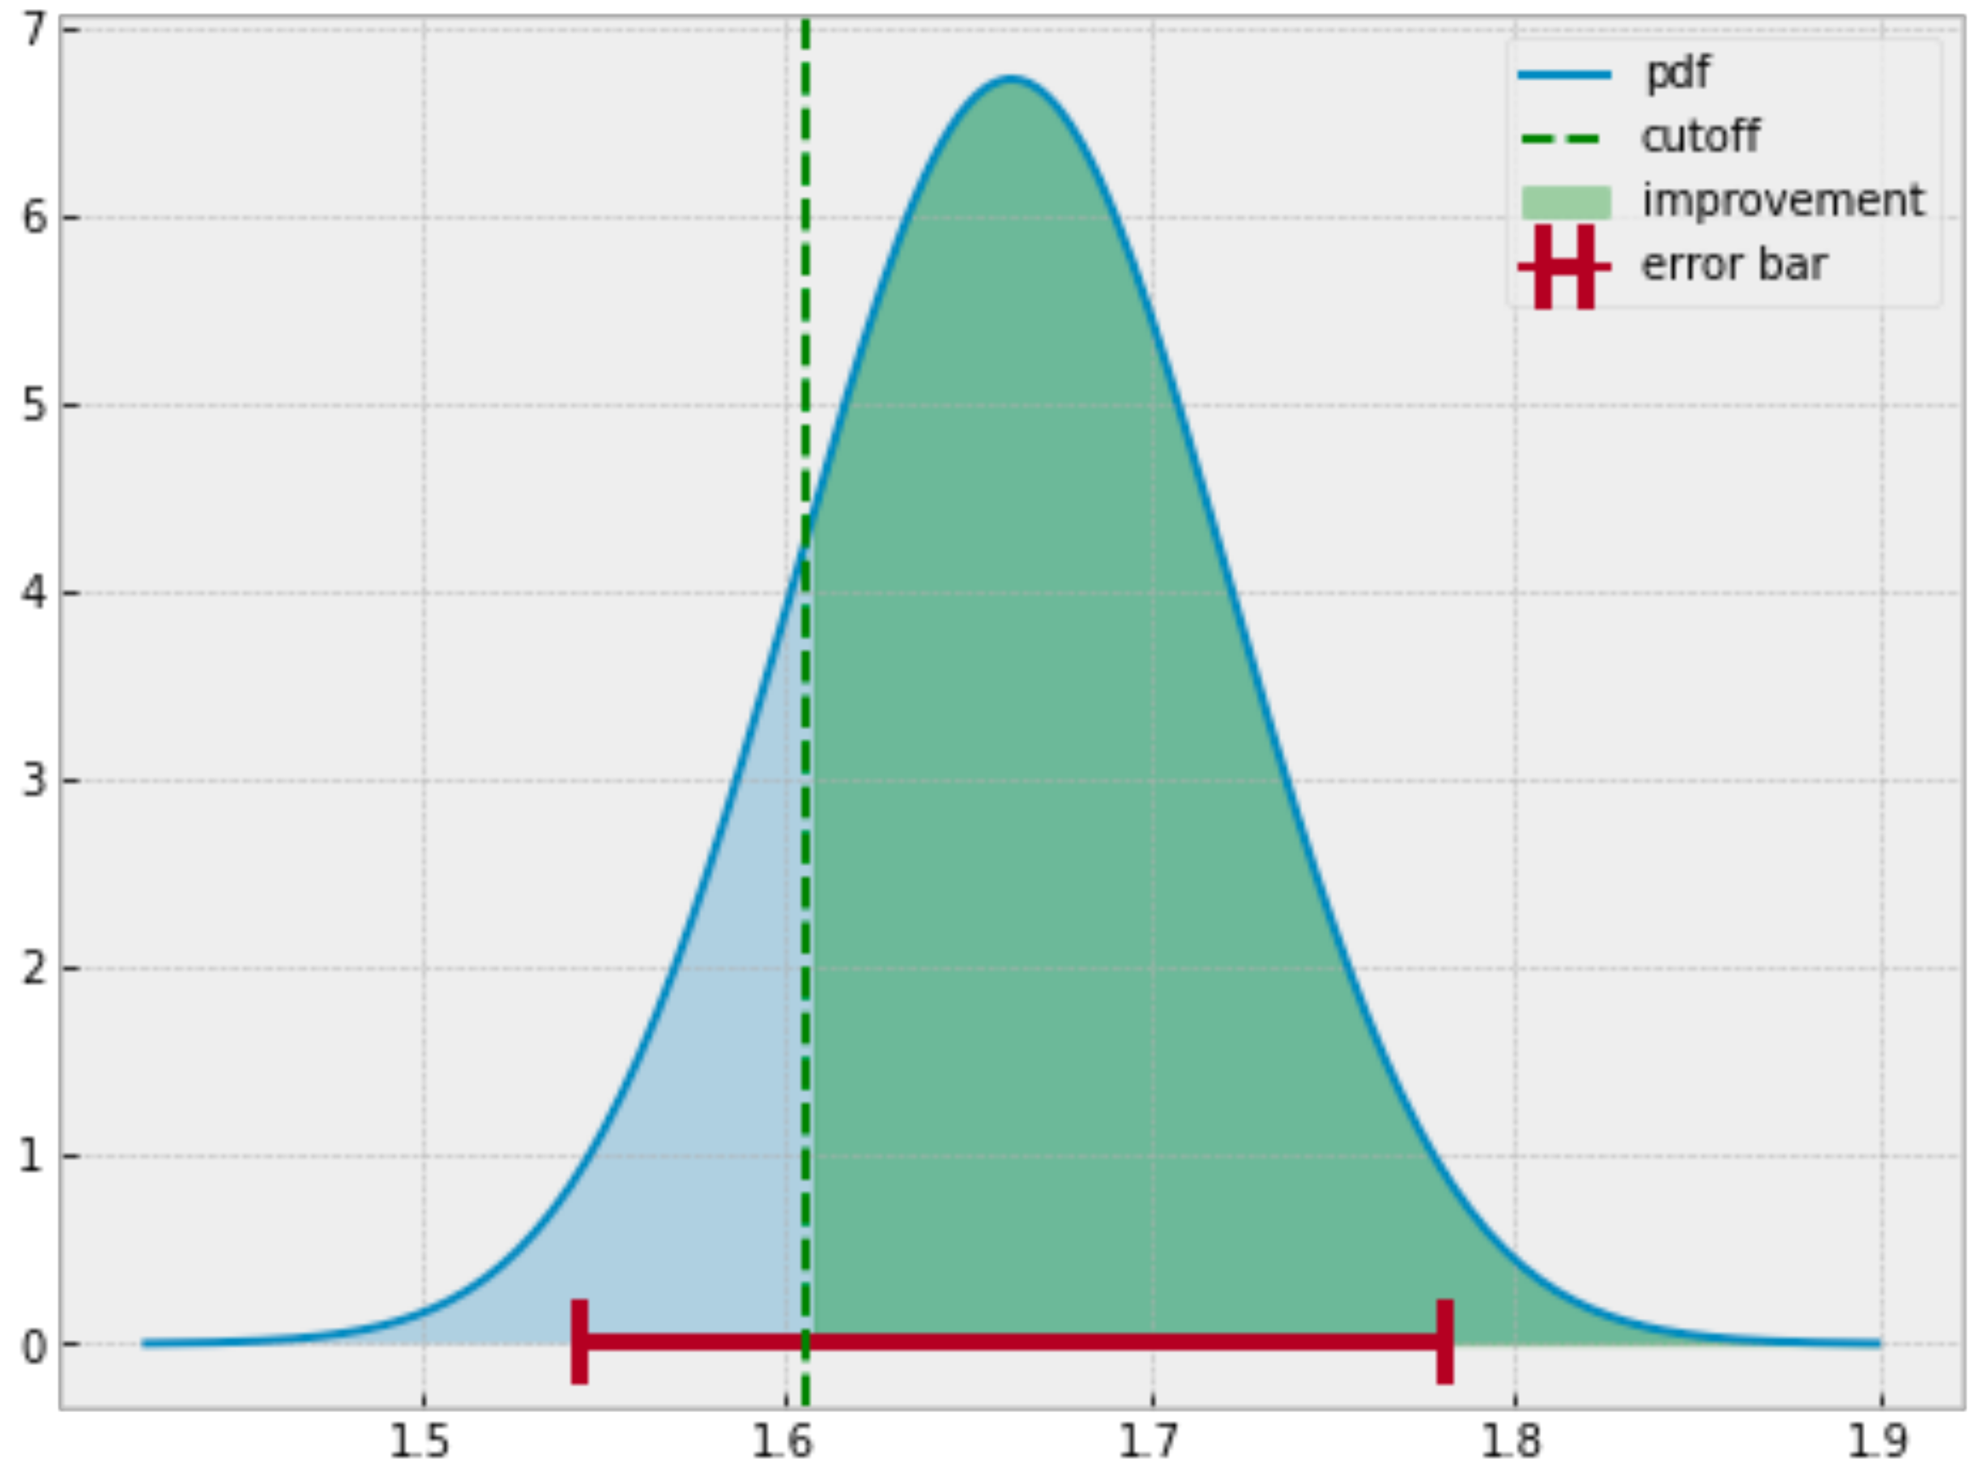

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

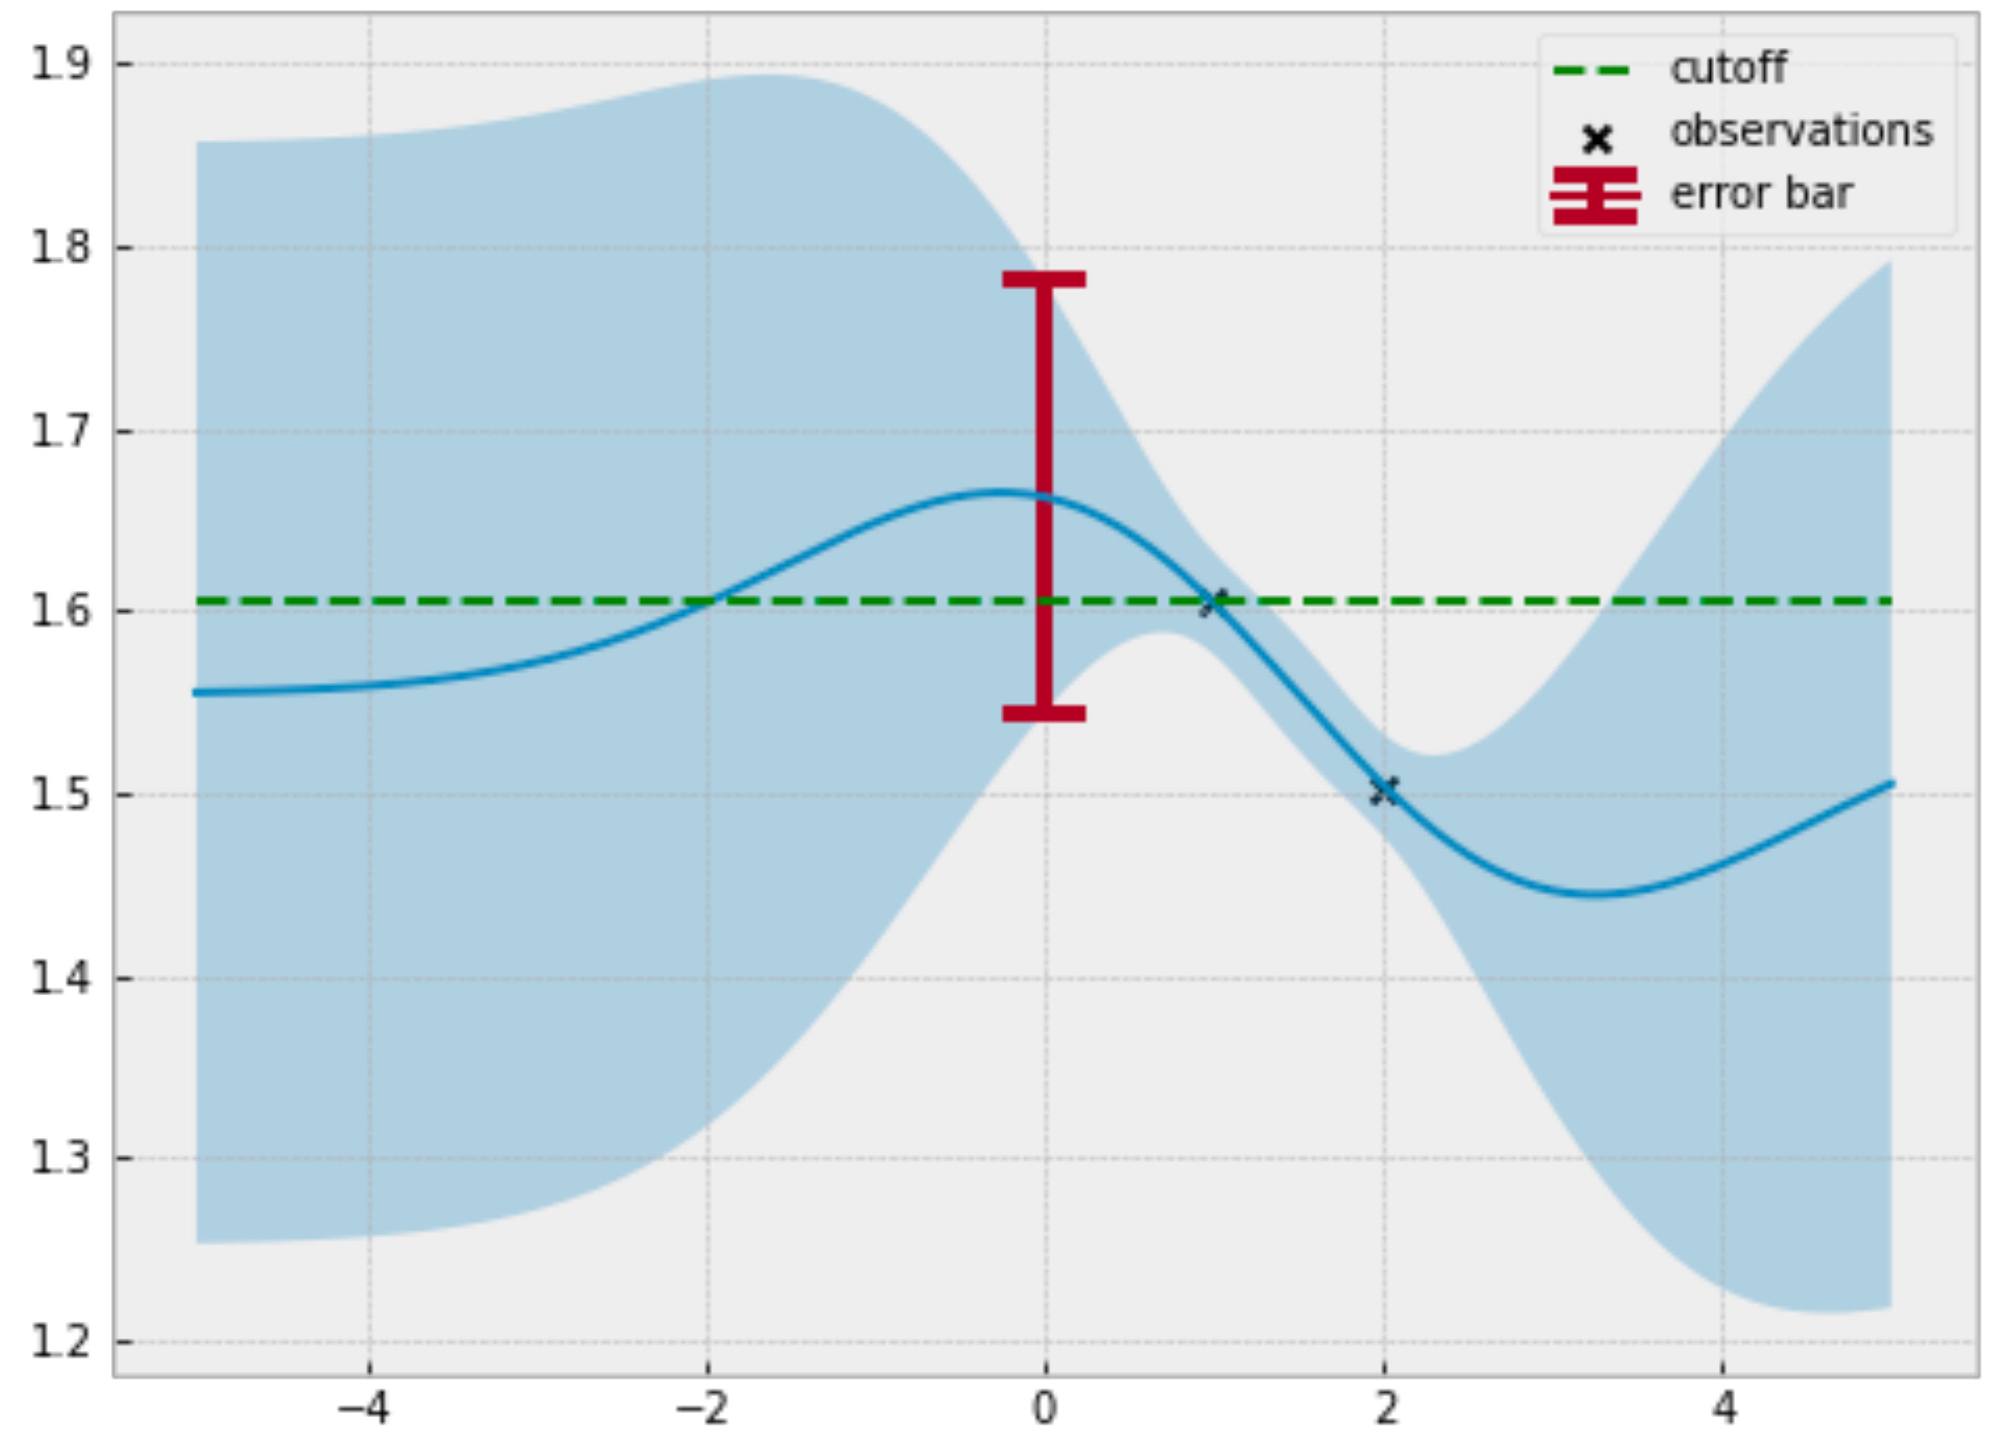

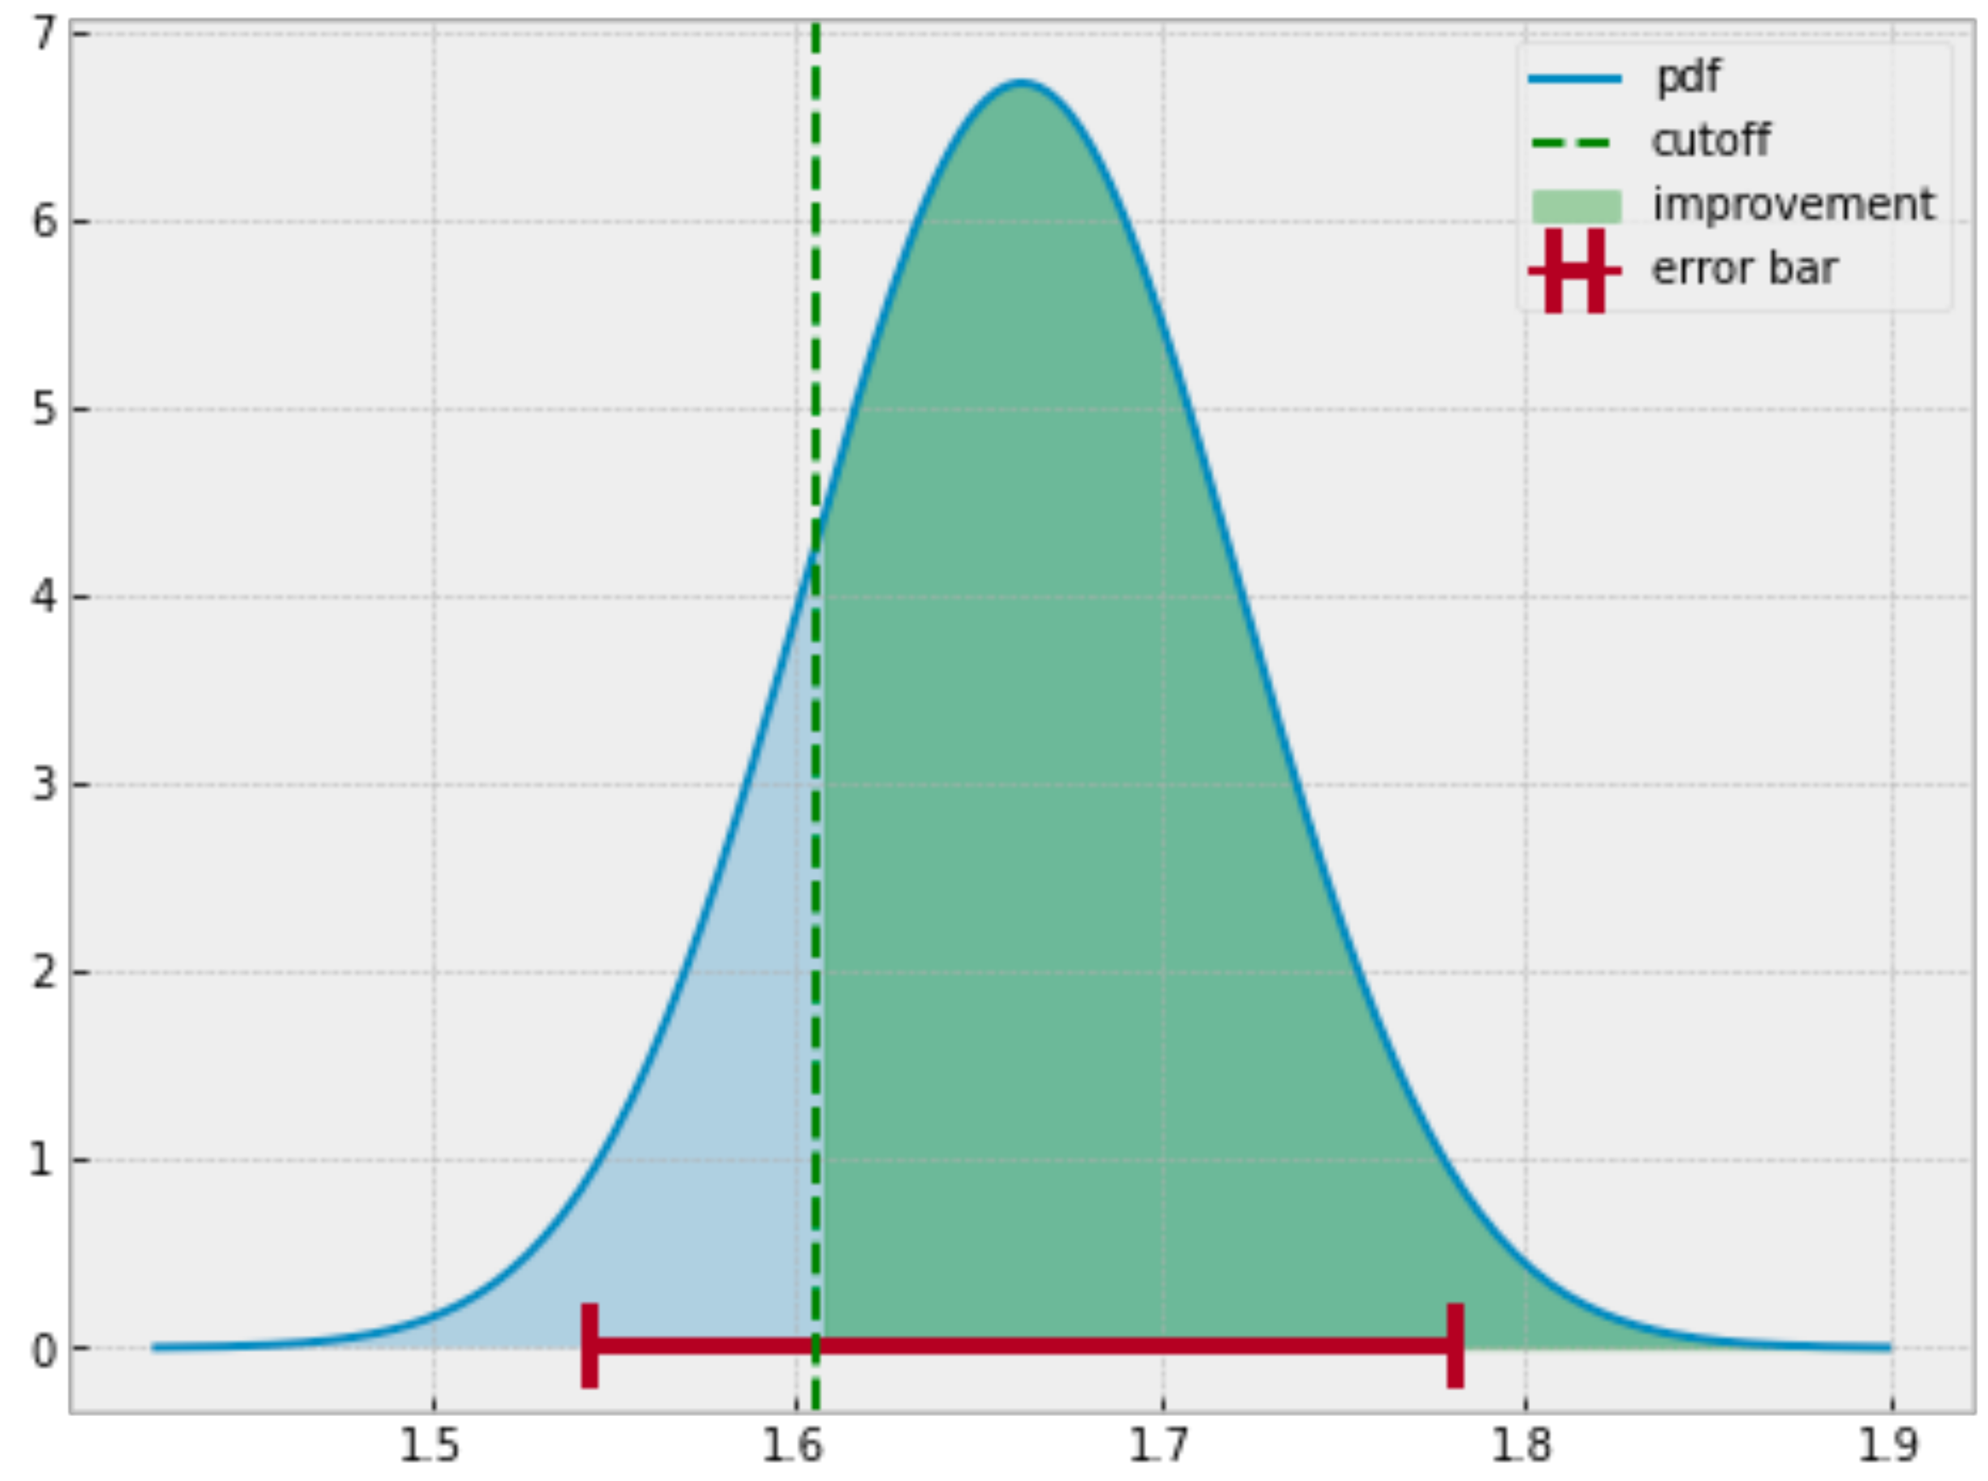

$$\Pr\left(y > \text{incumbent}\right) = \Phi\left(\frac{\mu - \text{incumbent}}{\sigma}\right)$$

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

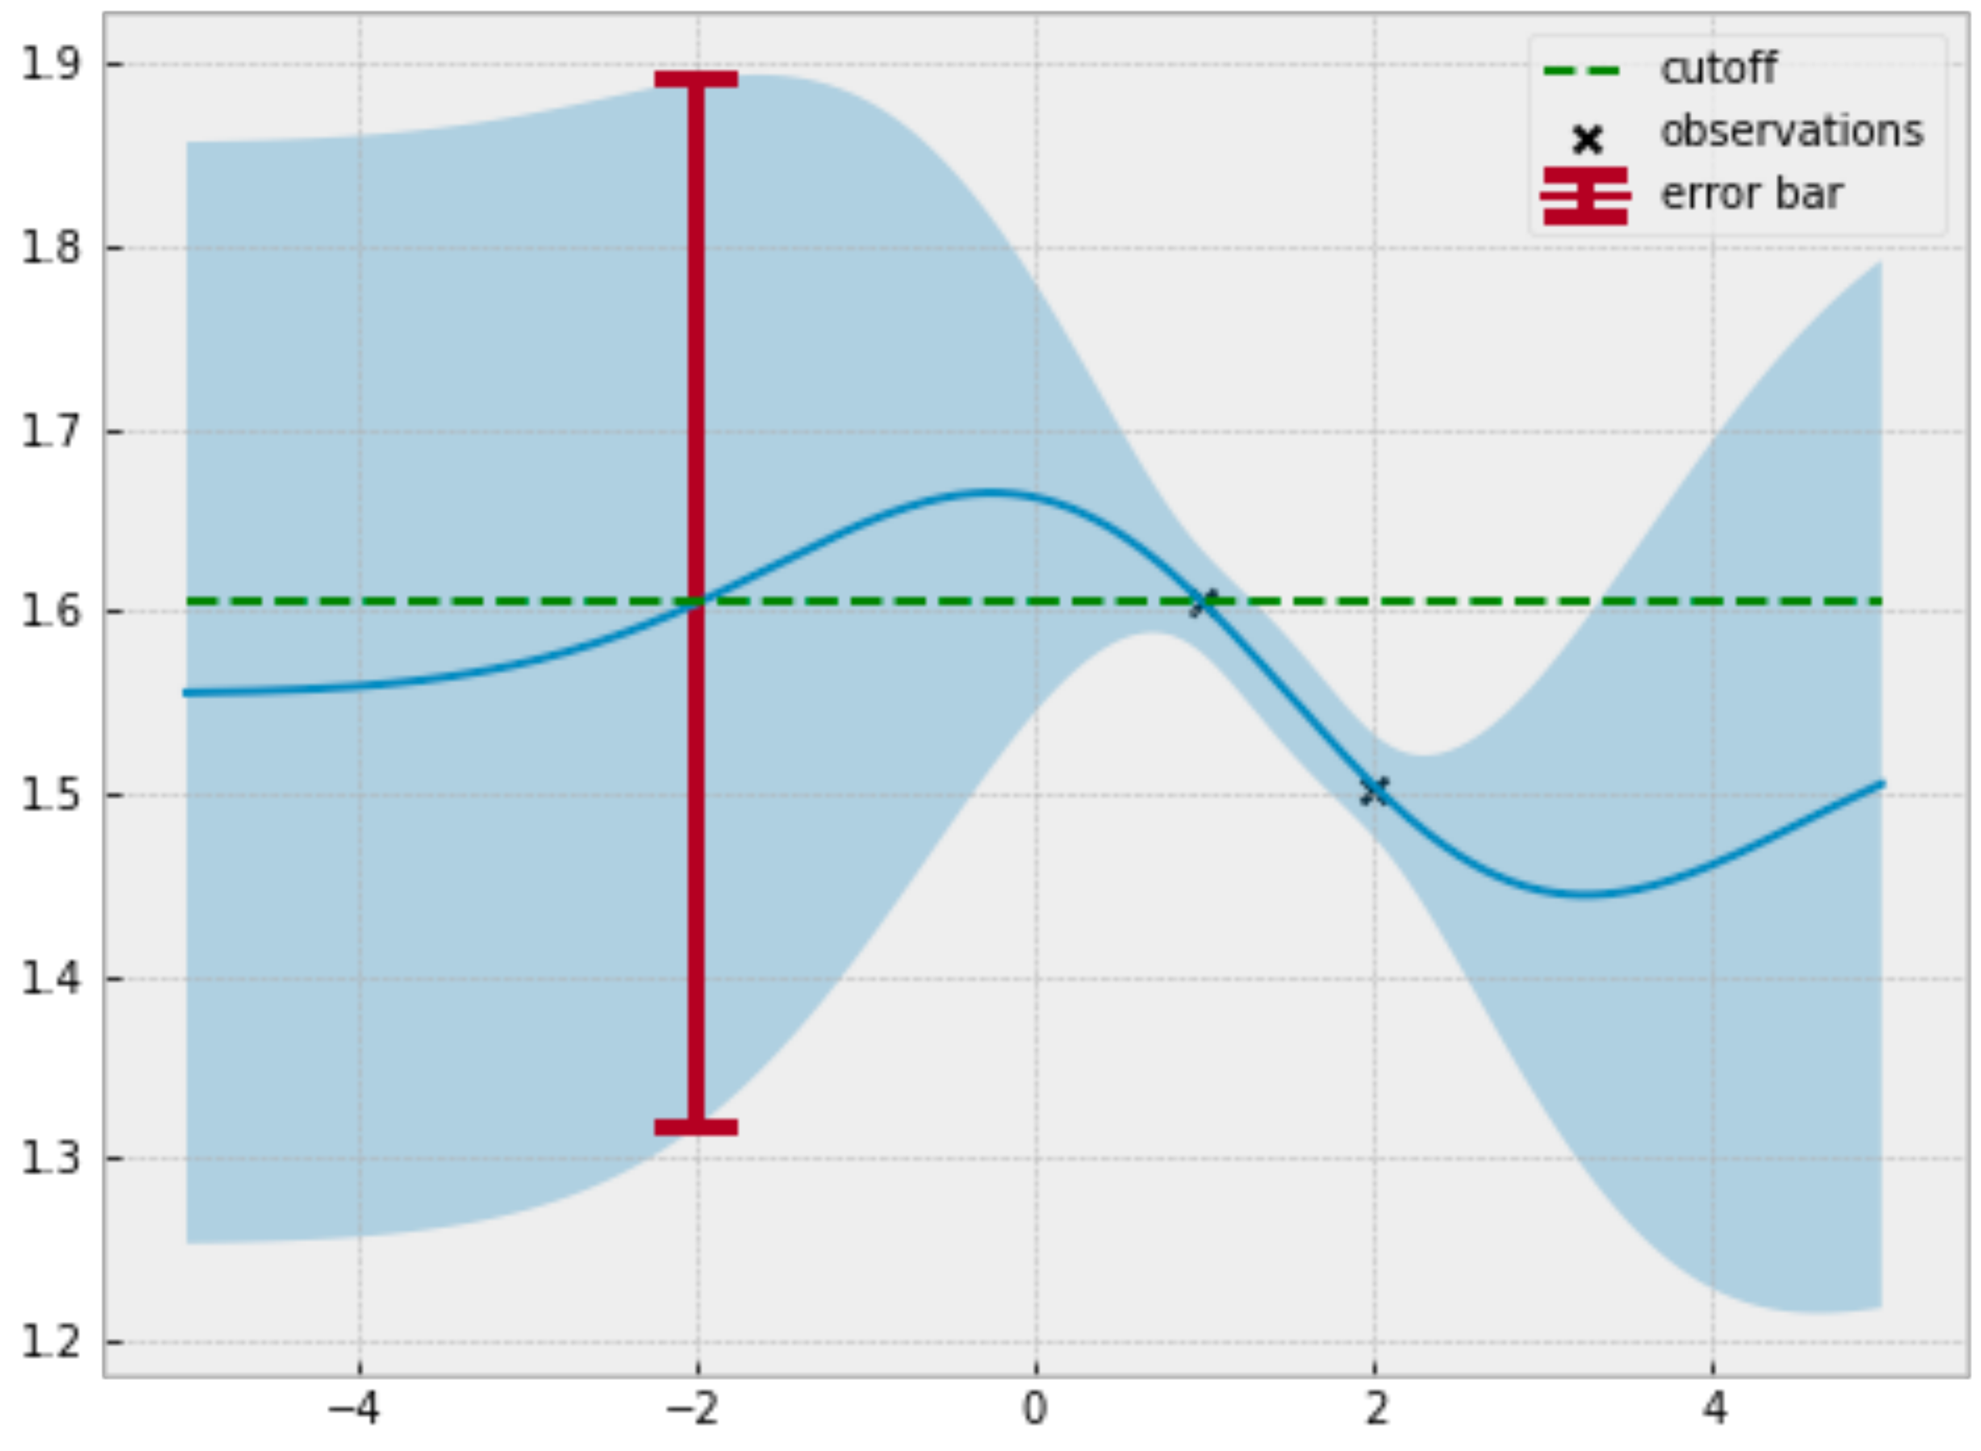

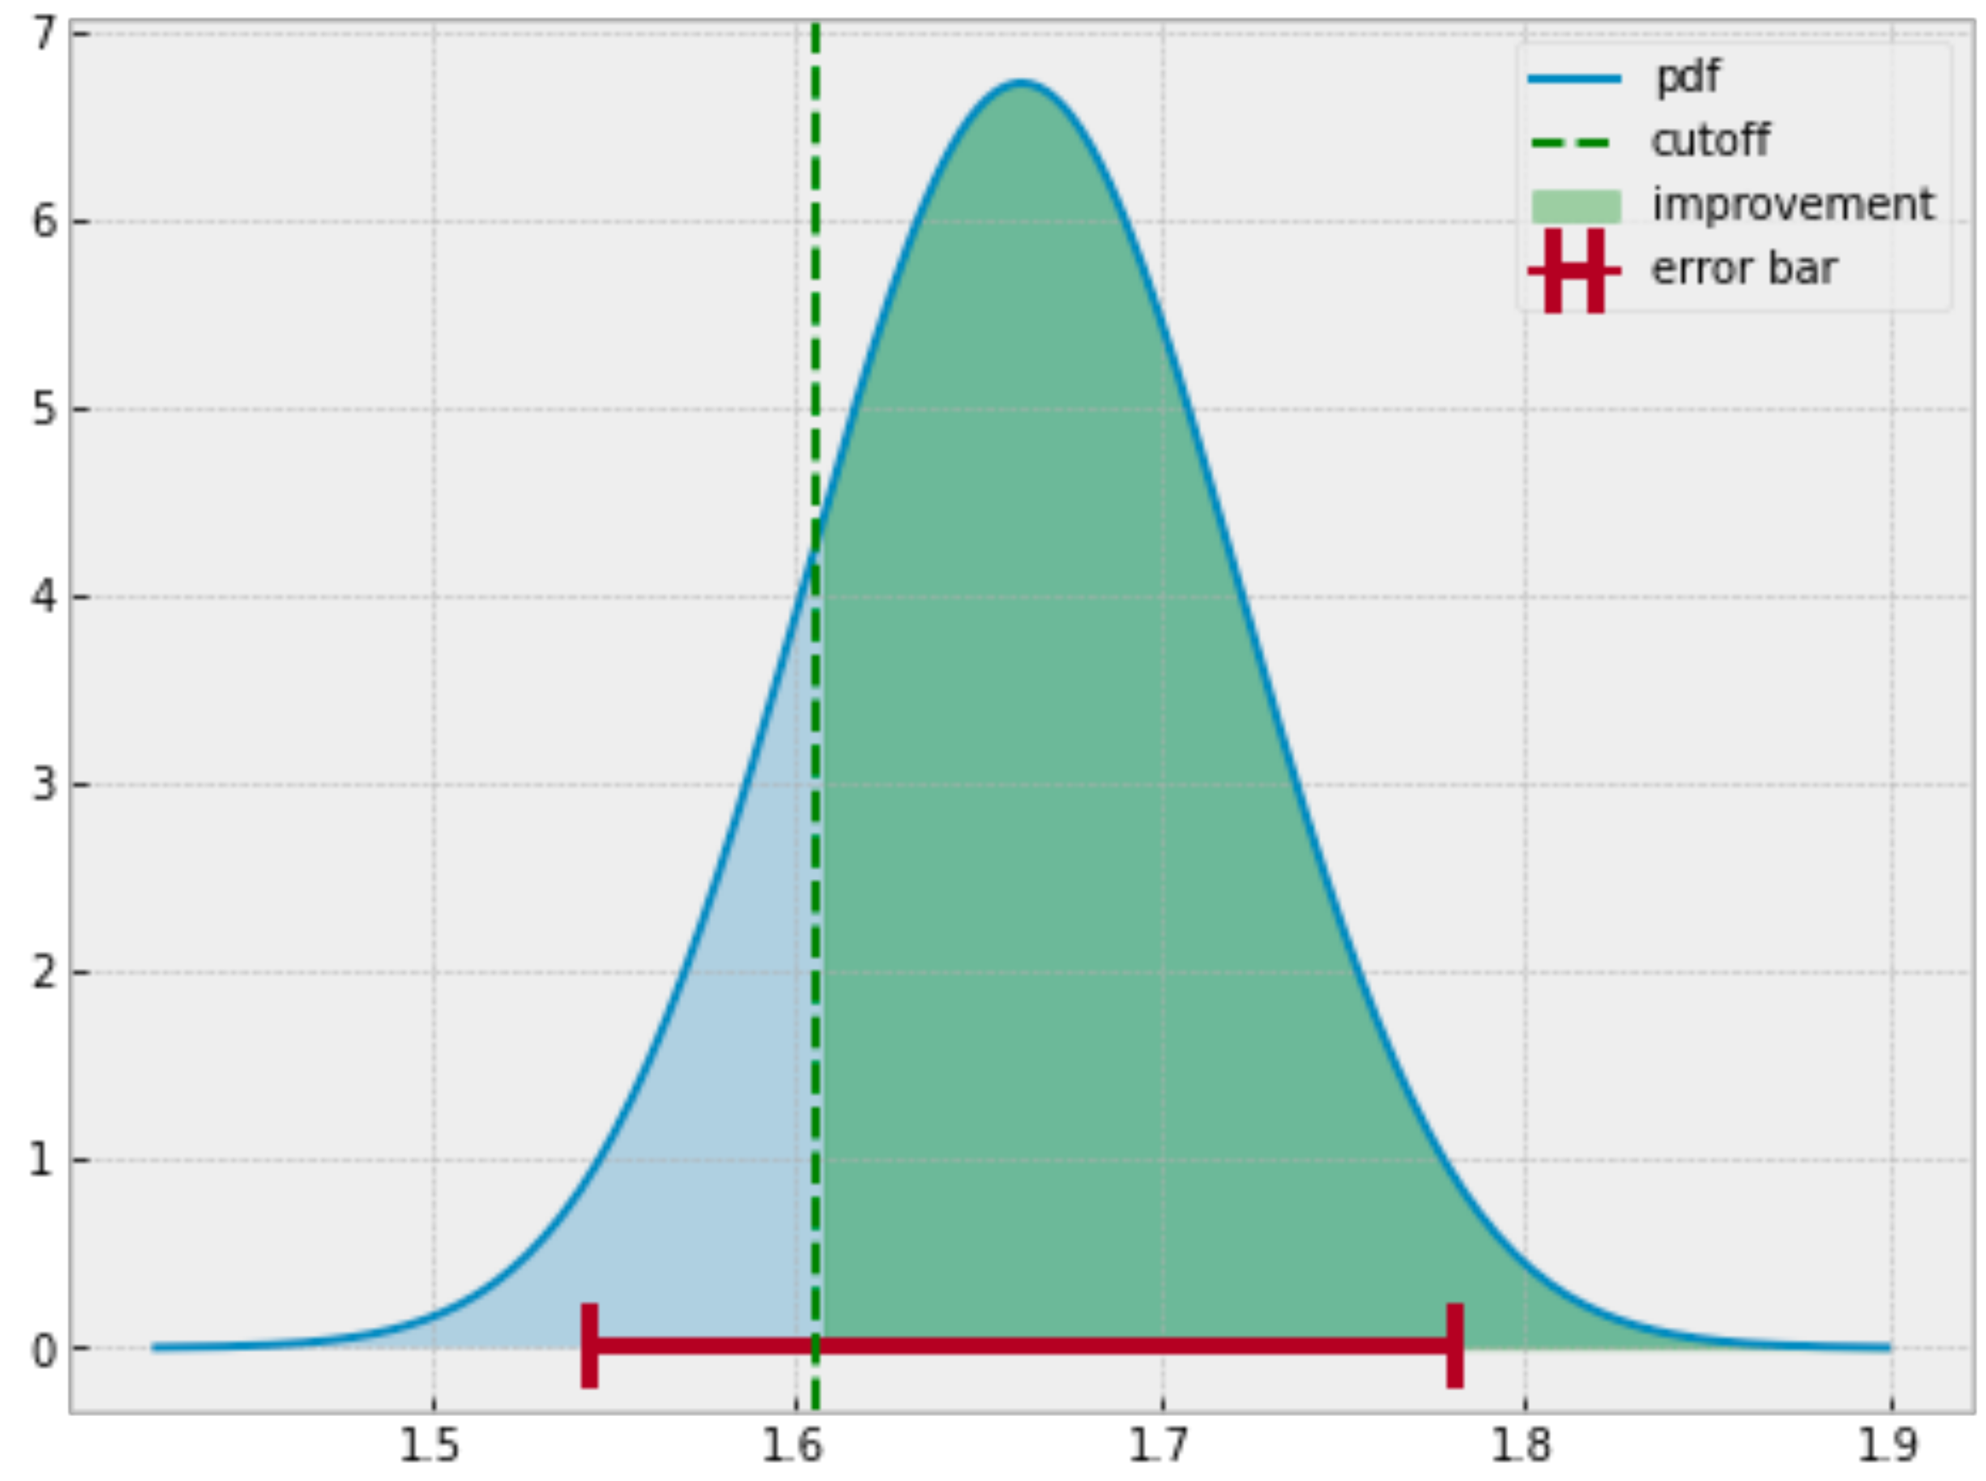

$$\Pr\left(y > \text{incumbent}\right) = \Phi\left(\frac{\mu - \text{incumbent}}{\sigma}\right)$$

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

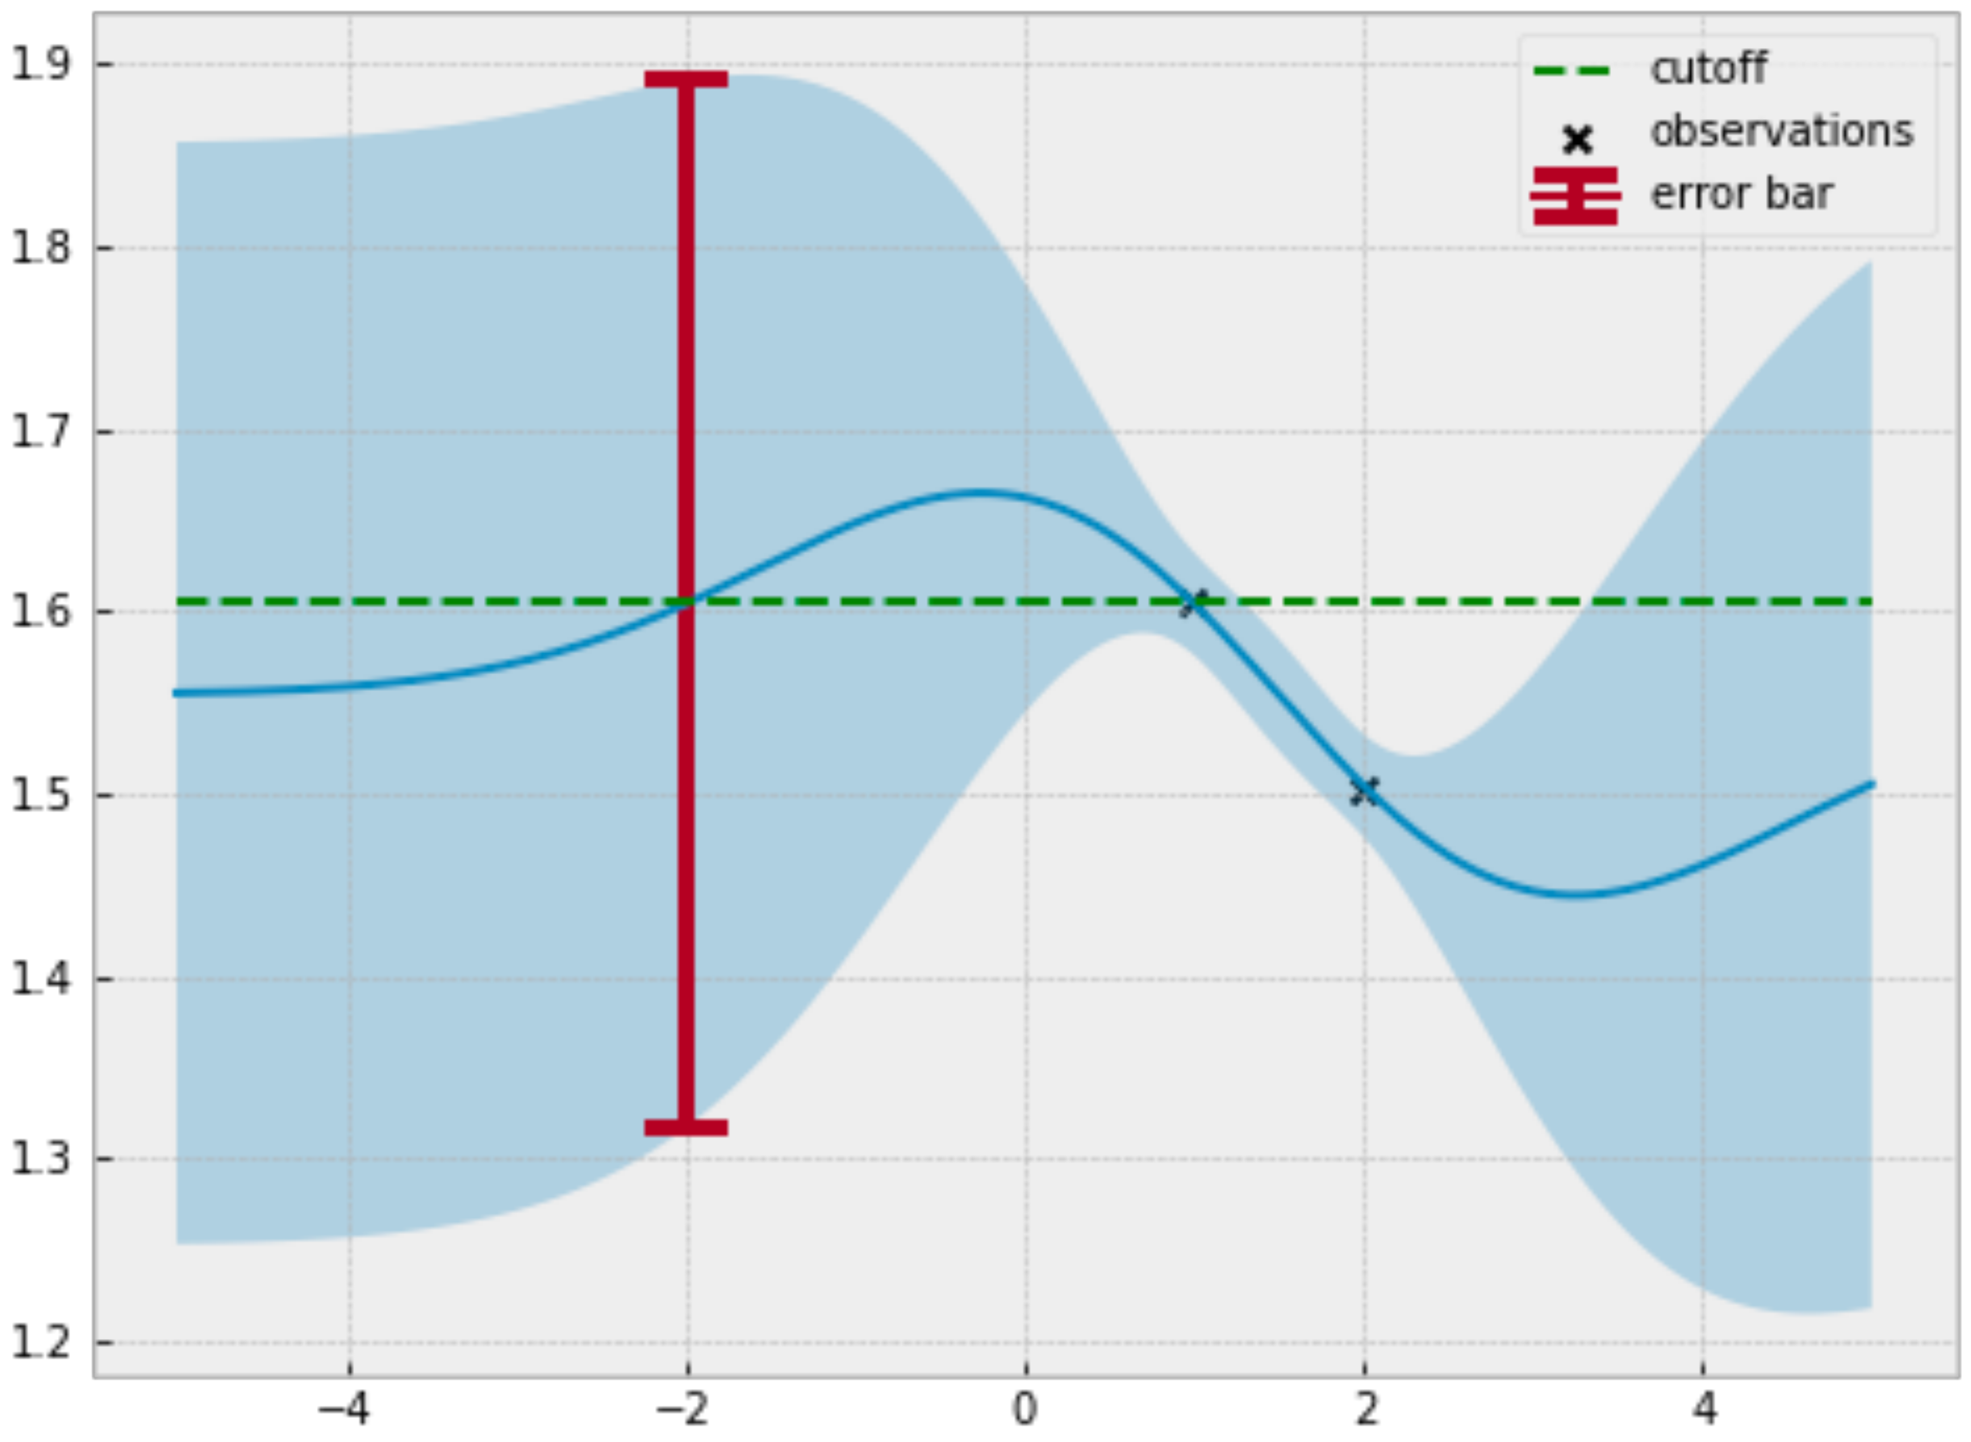

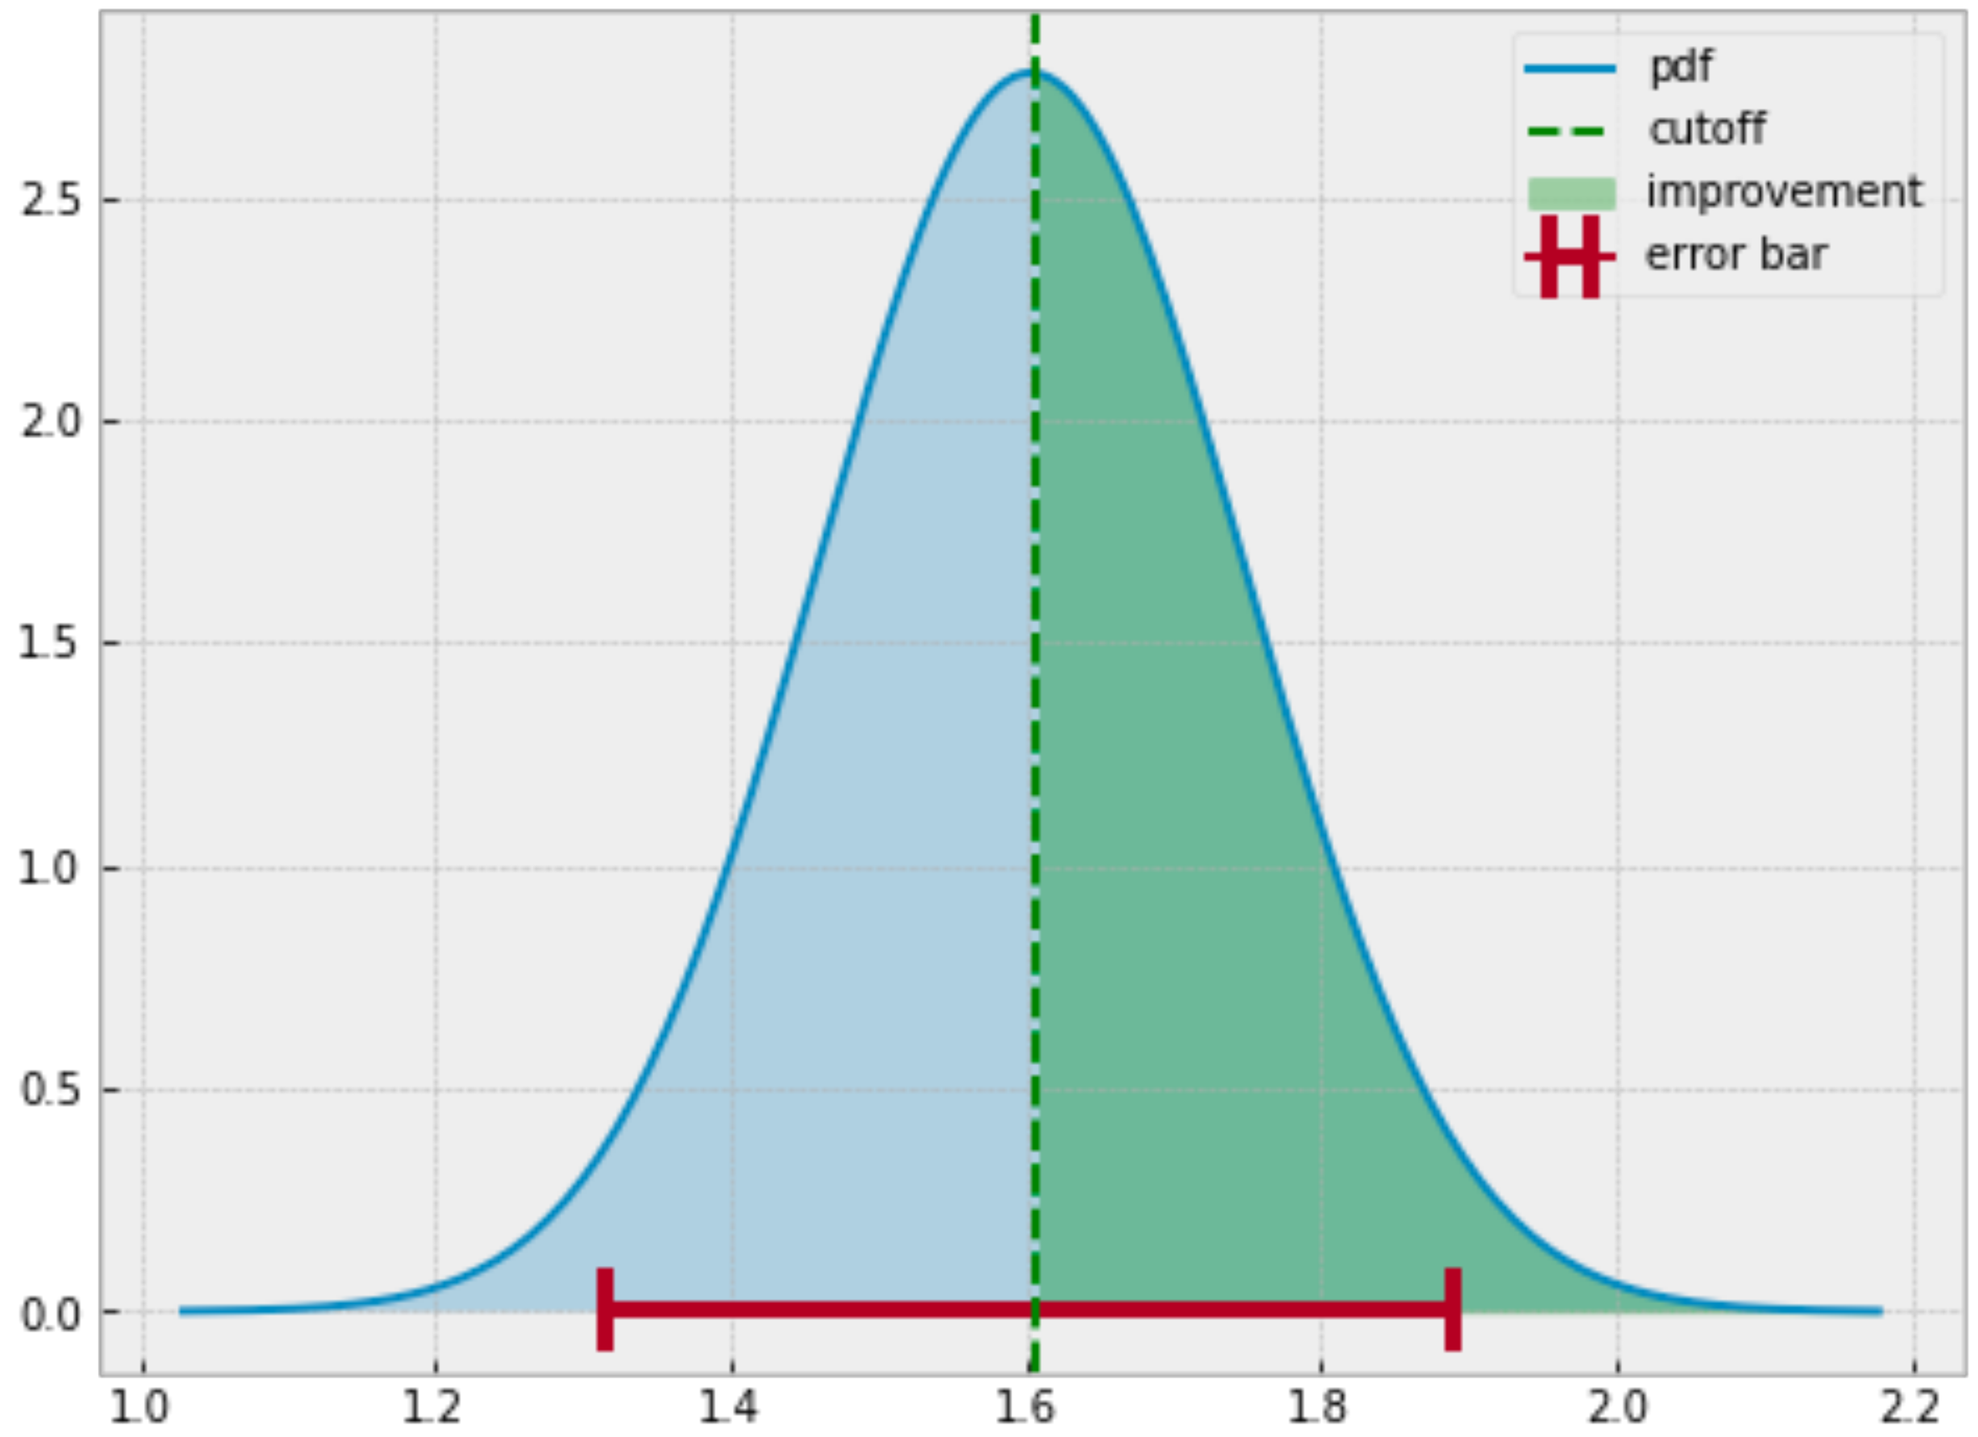

$$\Pr\left(y > \text{incumbent}\right) = \Phi\left(\frac{\mu - \text{incumbent}}{\sigma}\right)$$

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

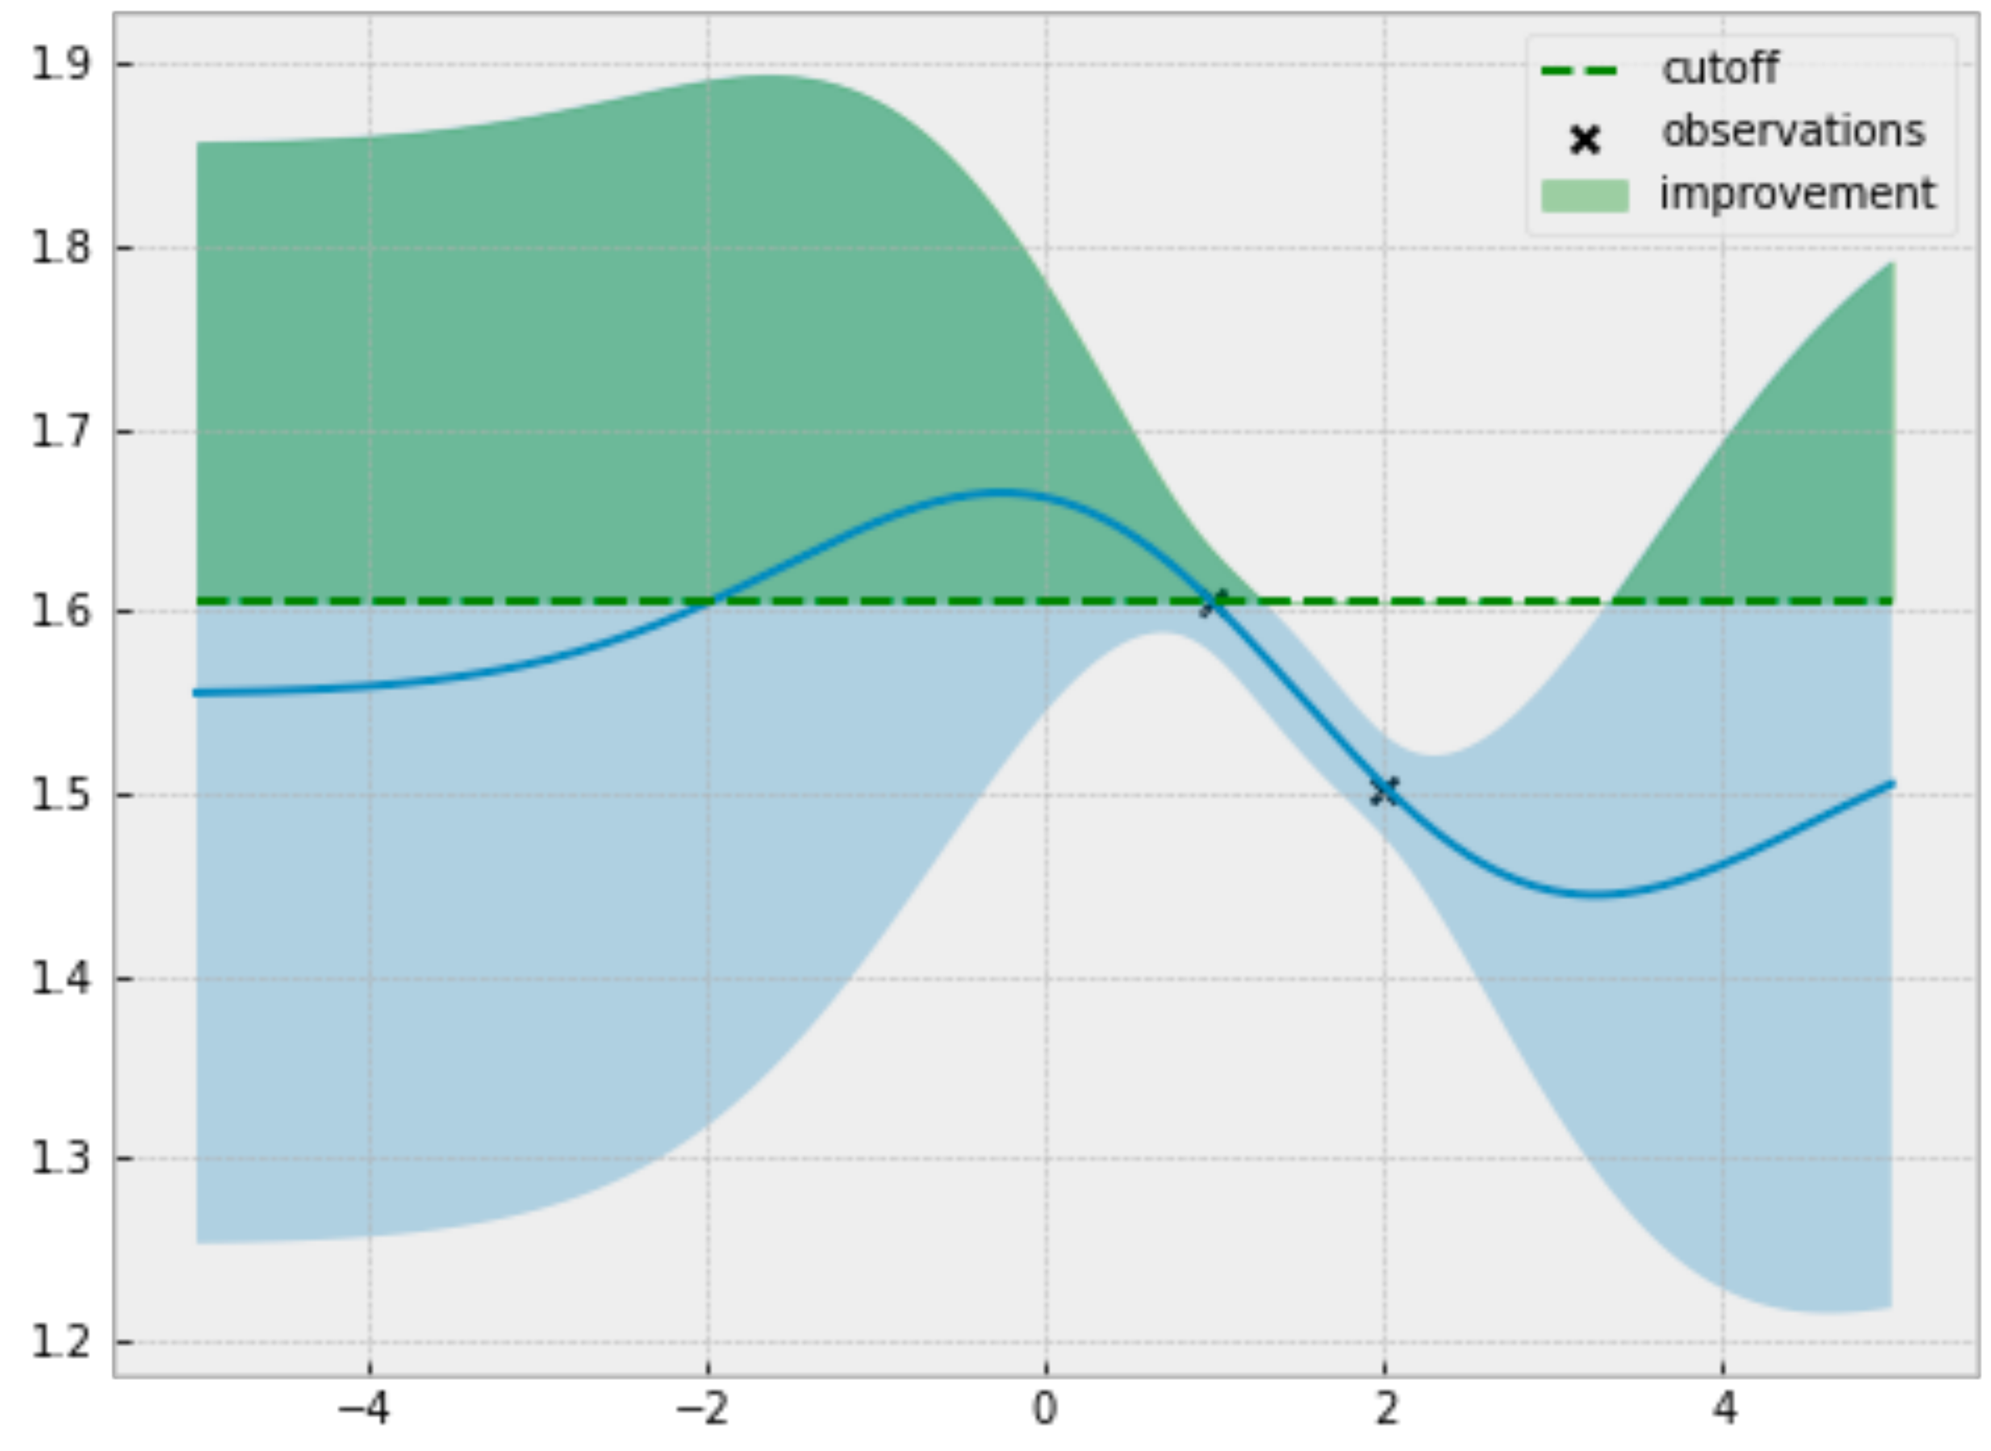

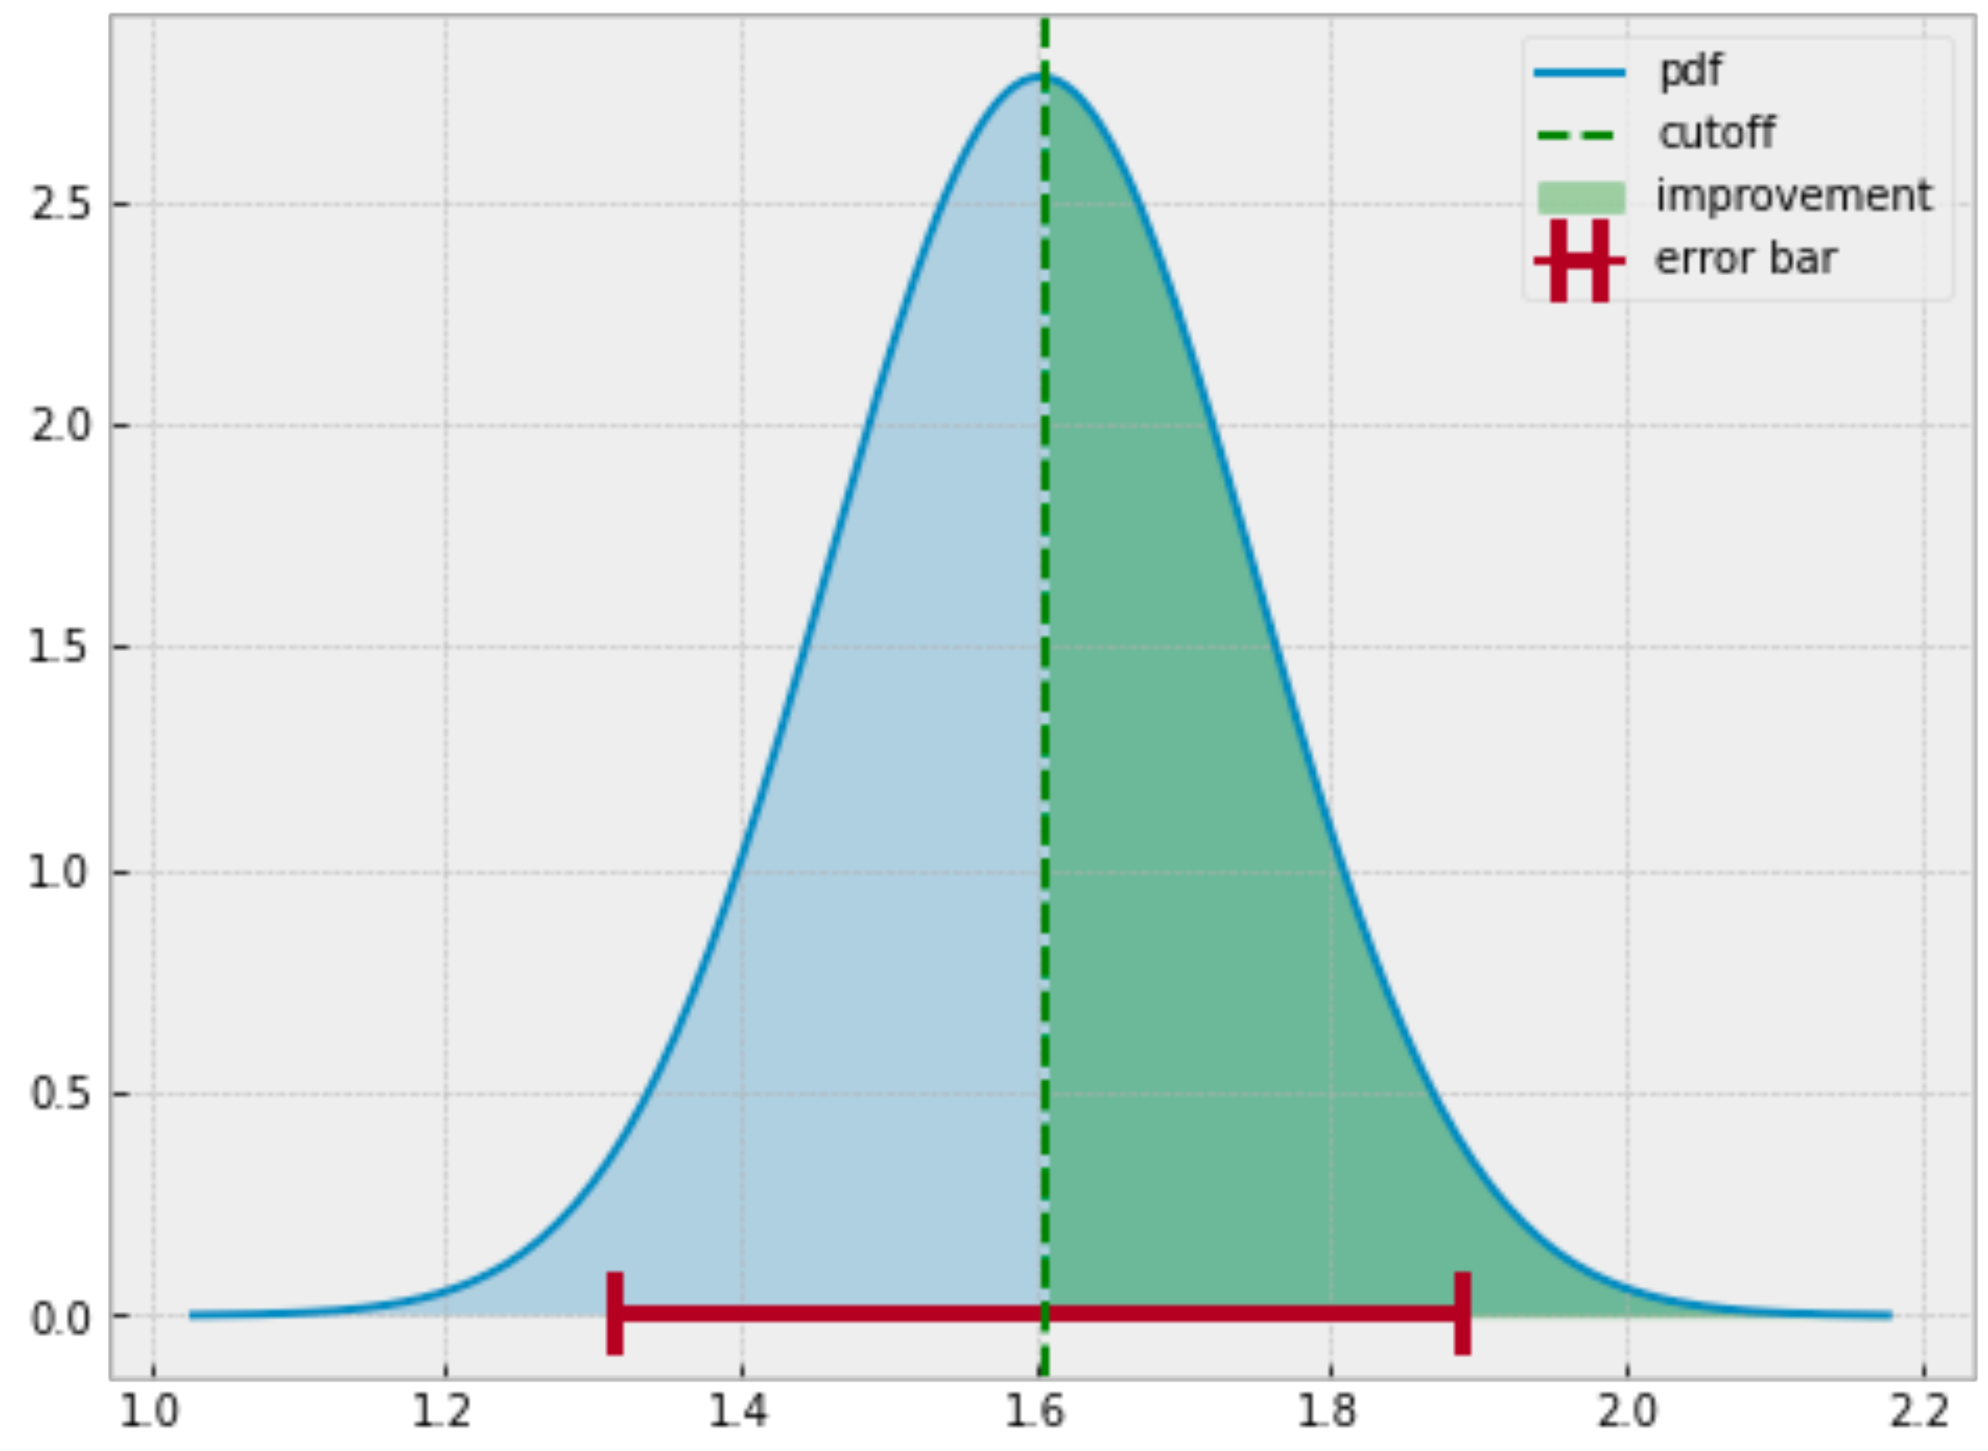

$$\Pr\left(y > \text{incumbent}\right) = \Phi\left(\frac{\mu - \text{incumbent}}{\sigma}\right)$$

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

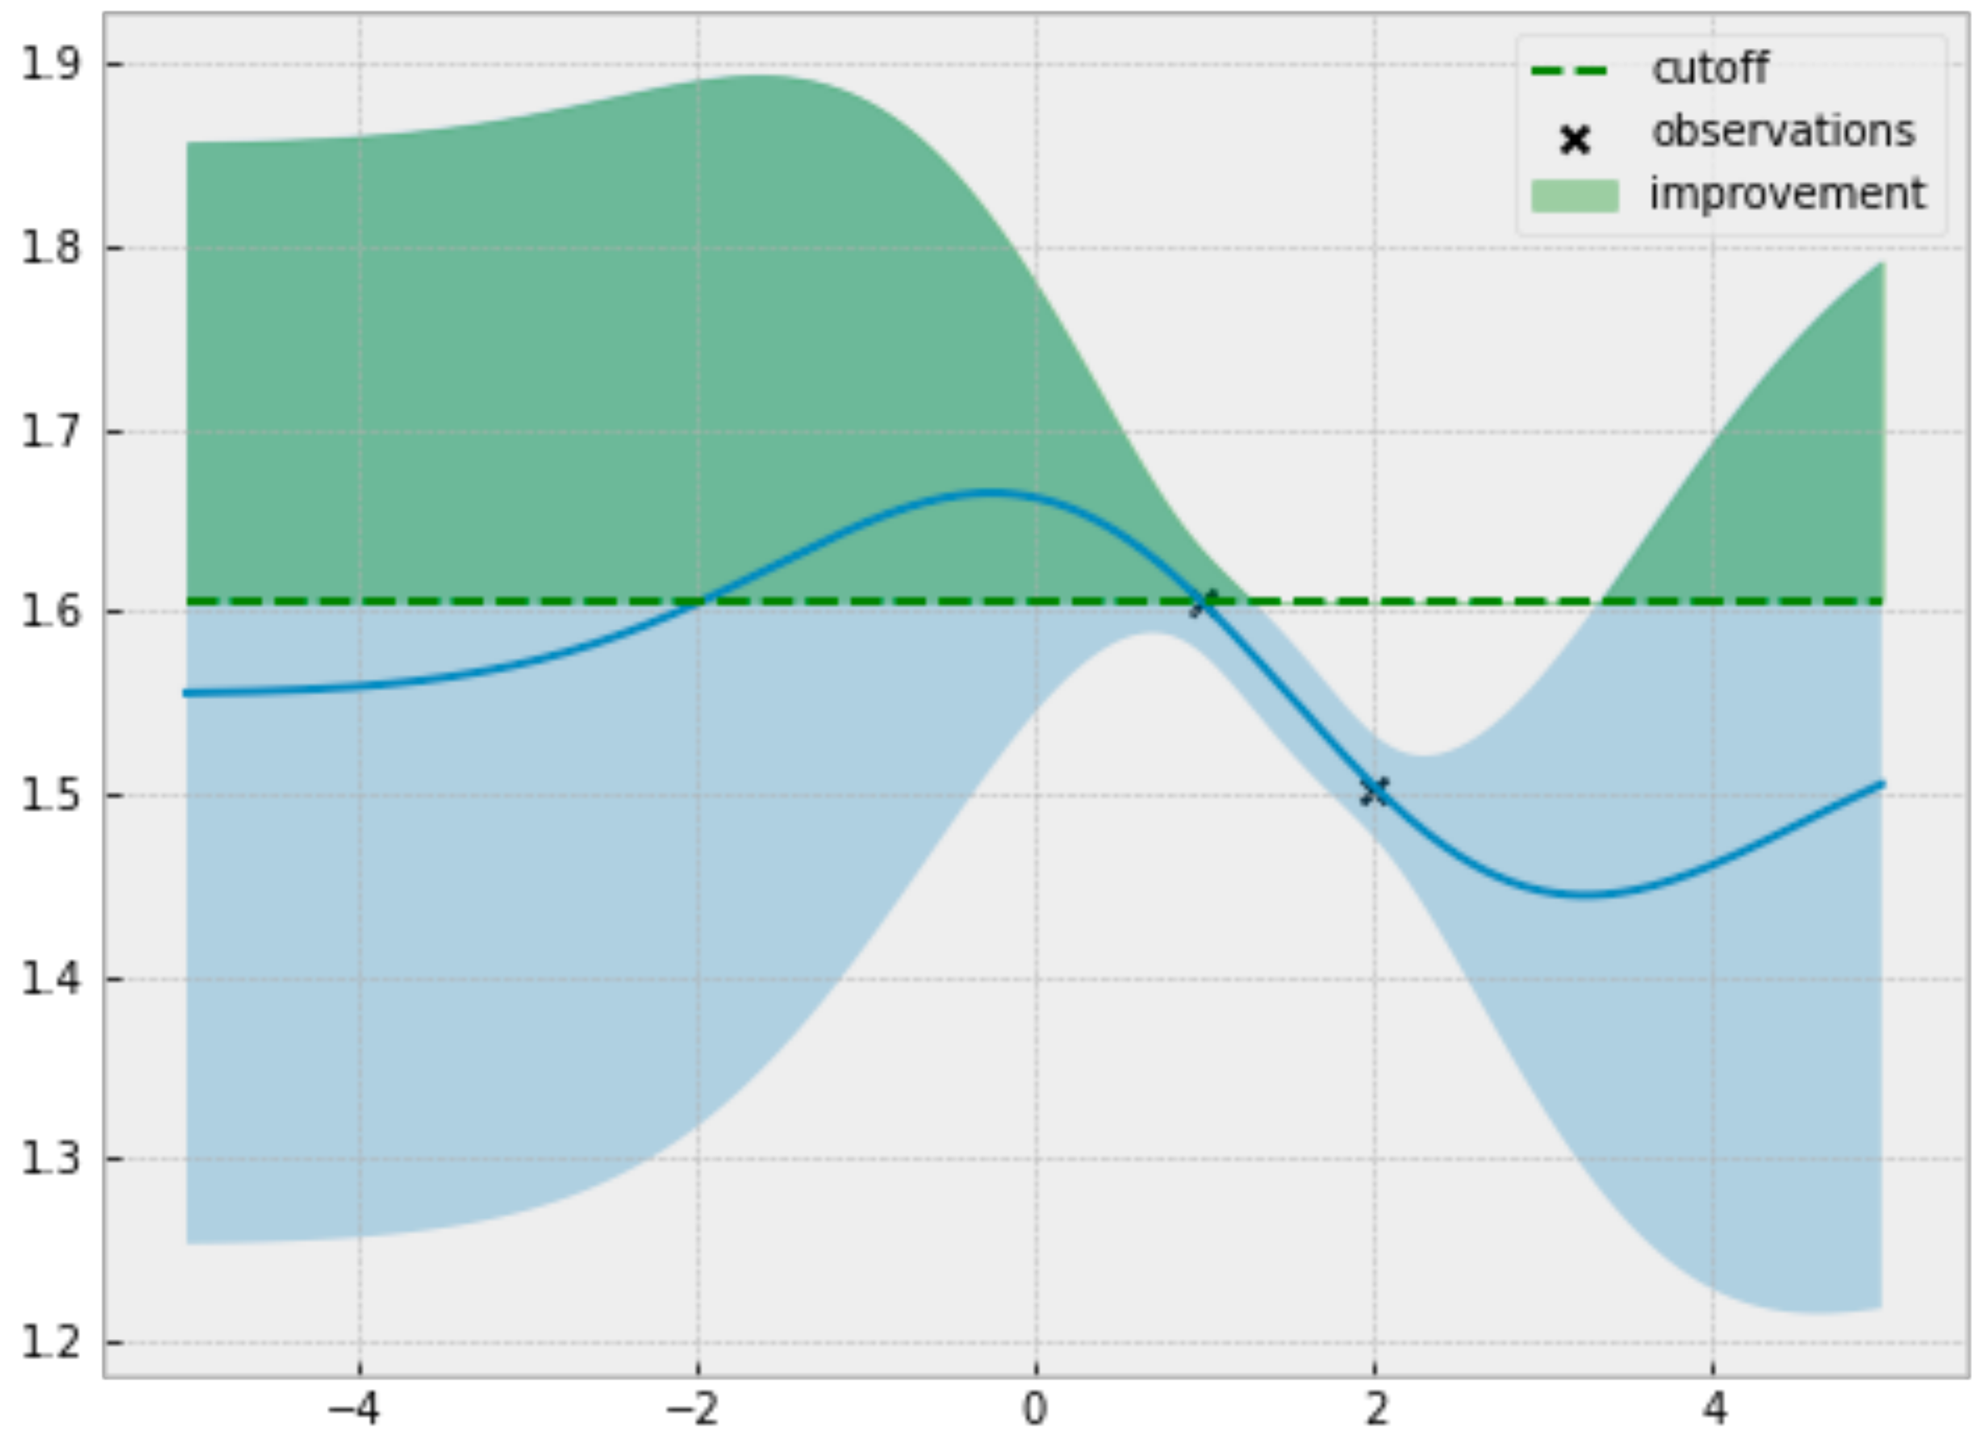

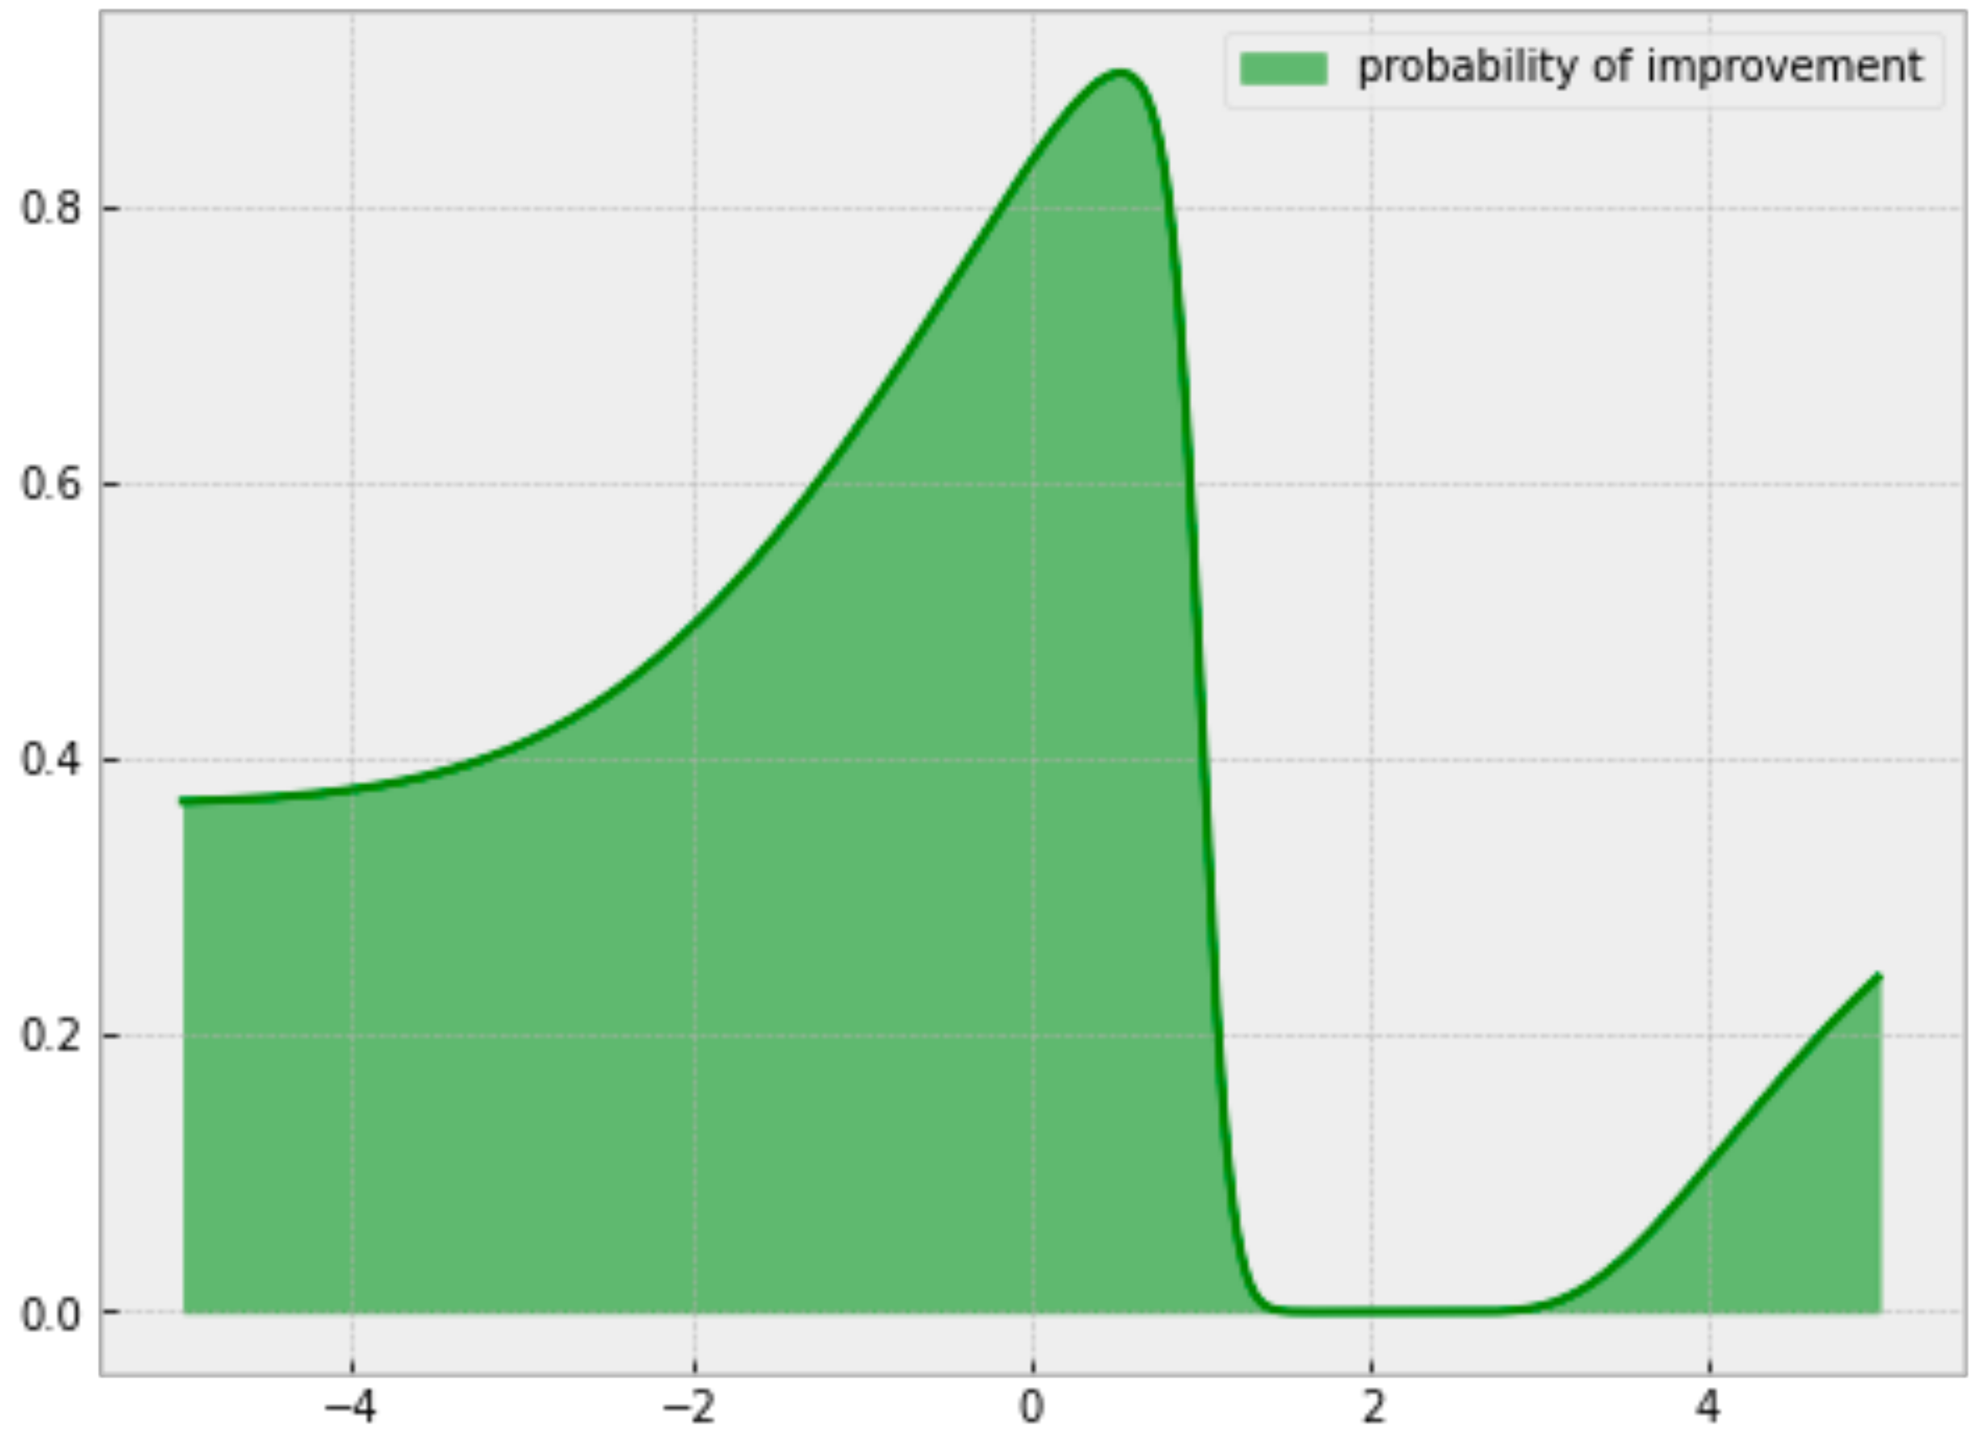

$$\Pr\left(y > \text{incumbent}\right) = \Phi\left(\frac{\mu - \text{incumbent}}{\sigma}\right)$$

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

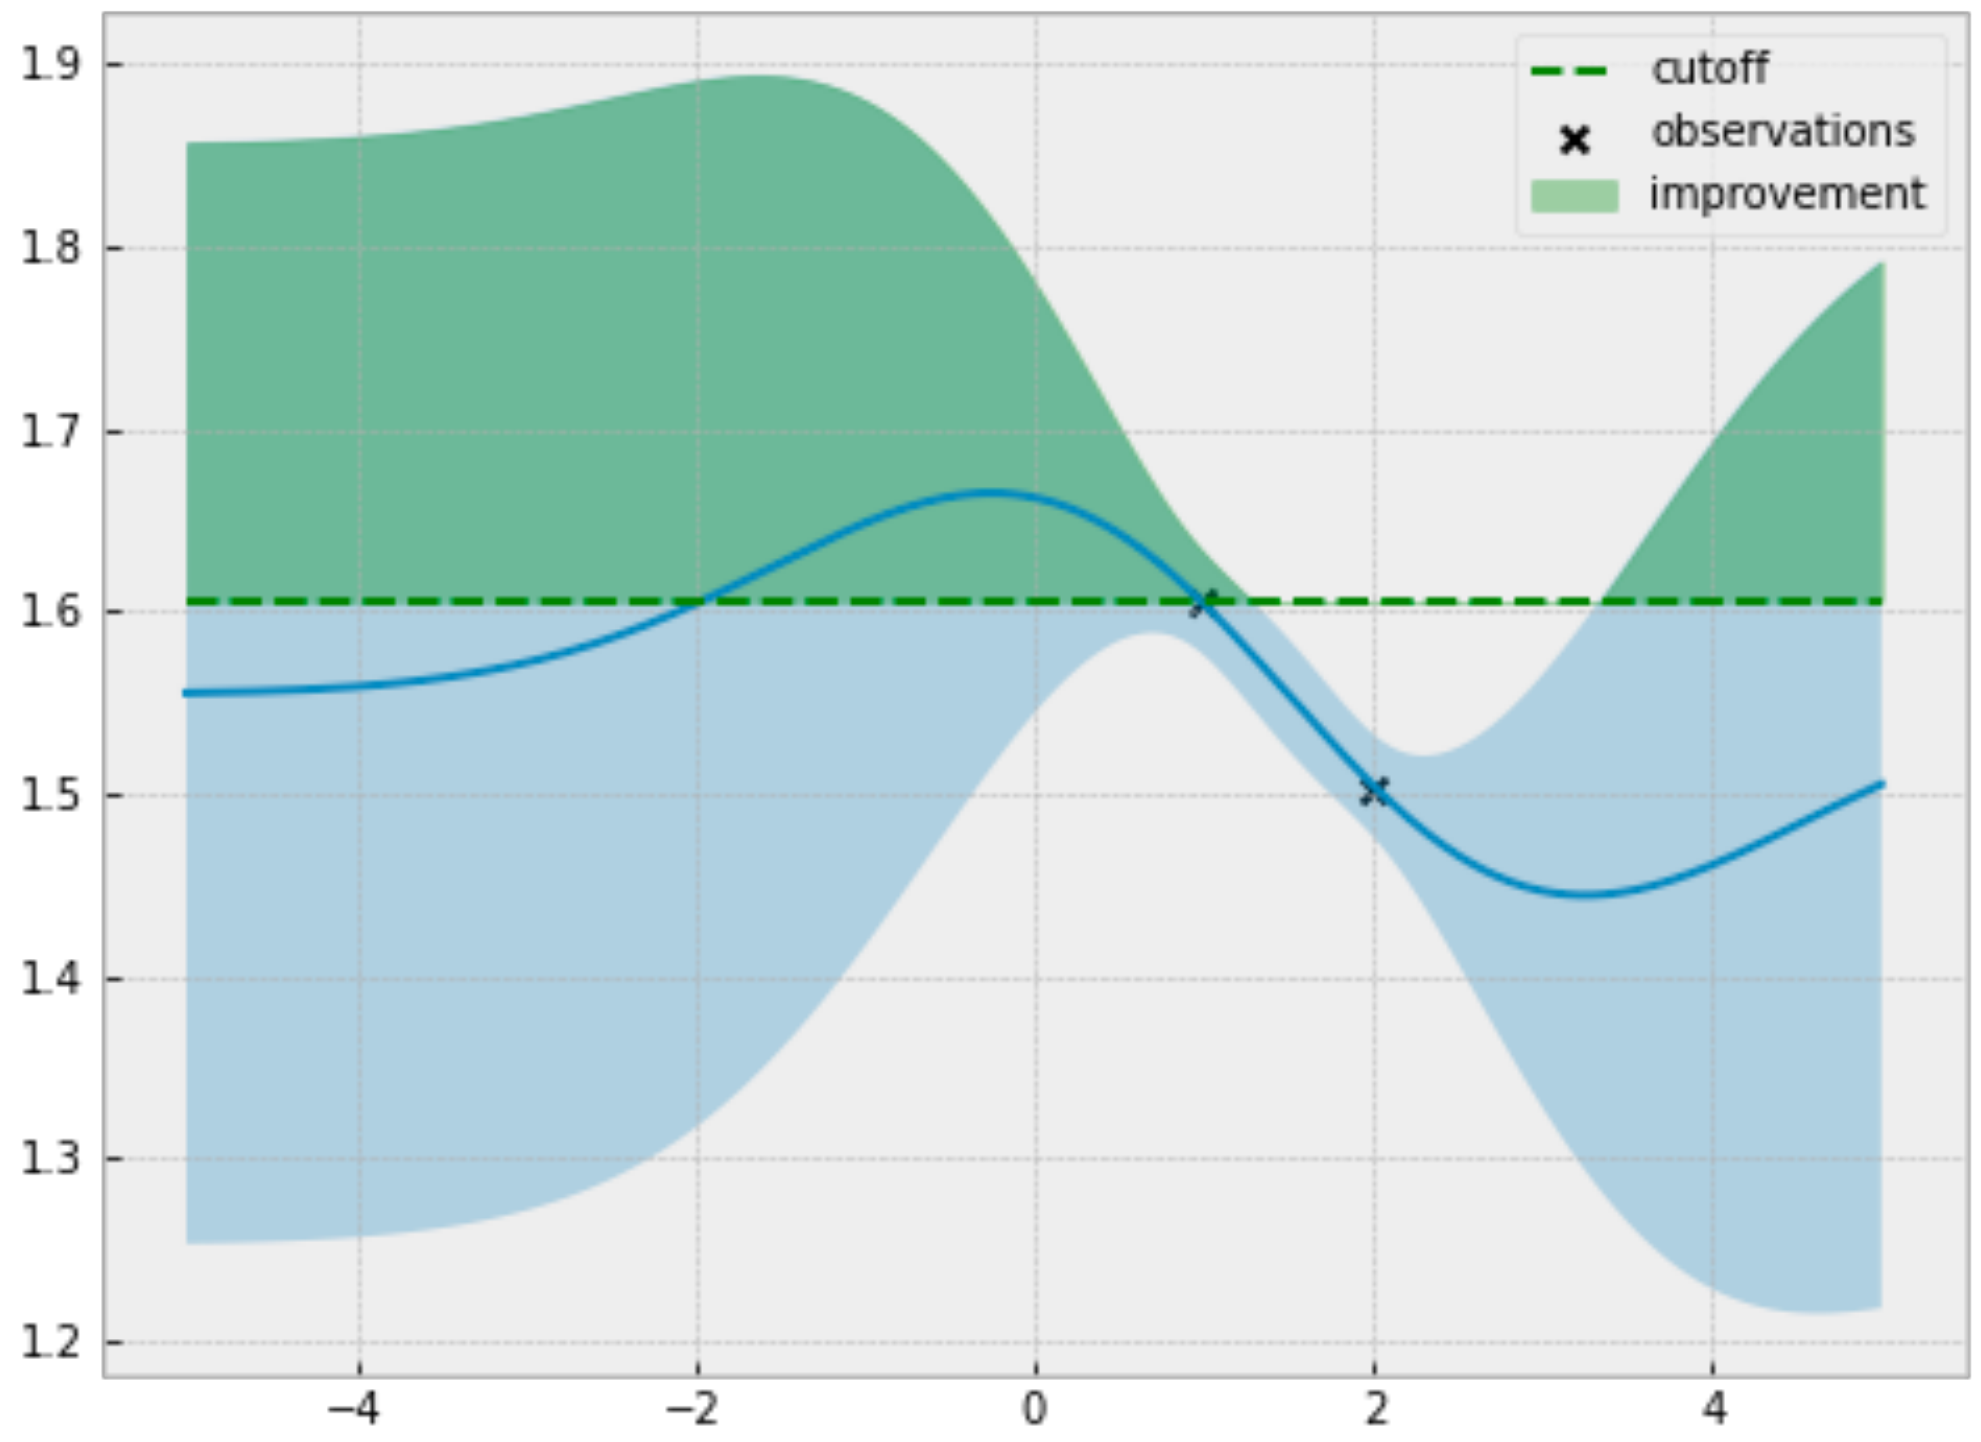

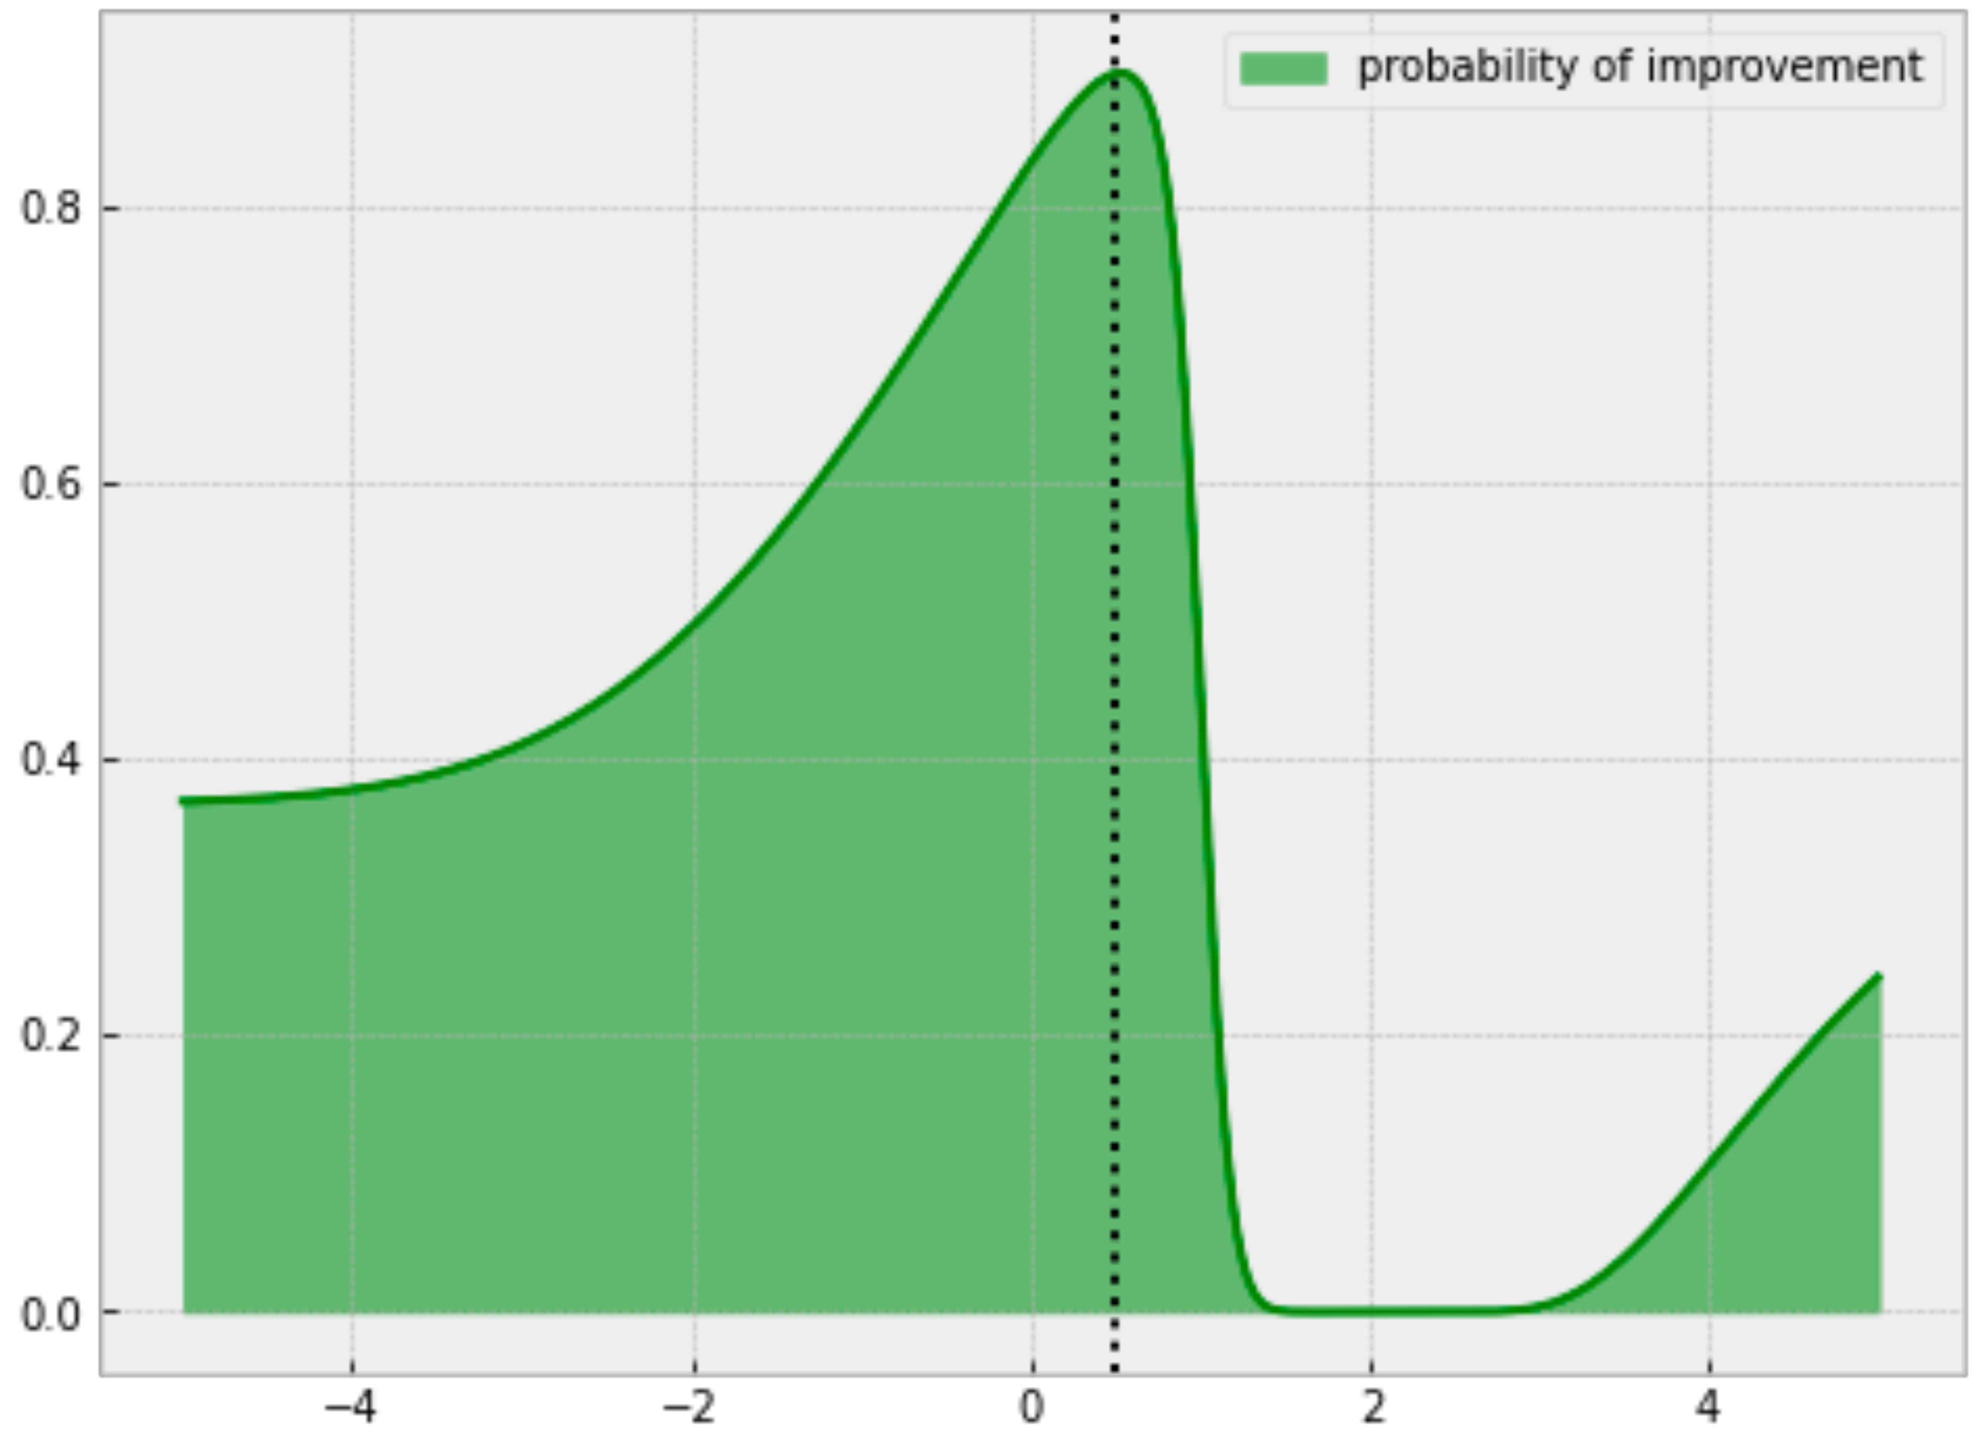

$$\Pr\left(y > \text{incumbent}\right) = \Phi\left(\frac{\mu - \text{incumbent}}{\sigma}\right)$$

THE ACTIVE LEARNING LOOP WITH PROBABILITY OF IMPROVEMENT

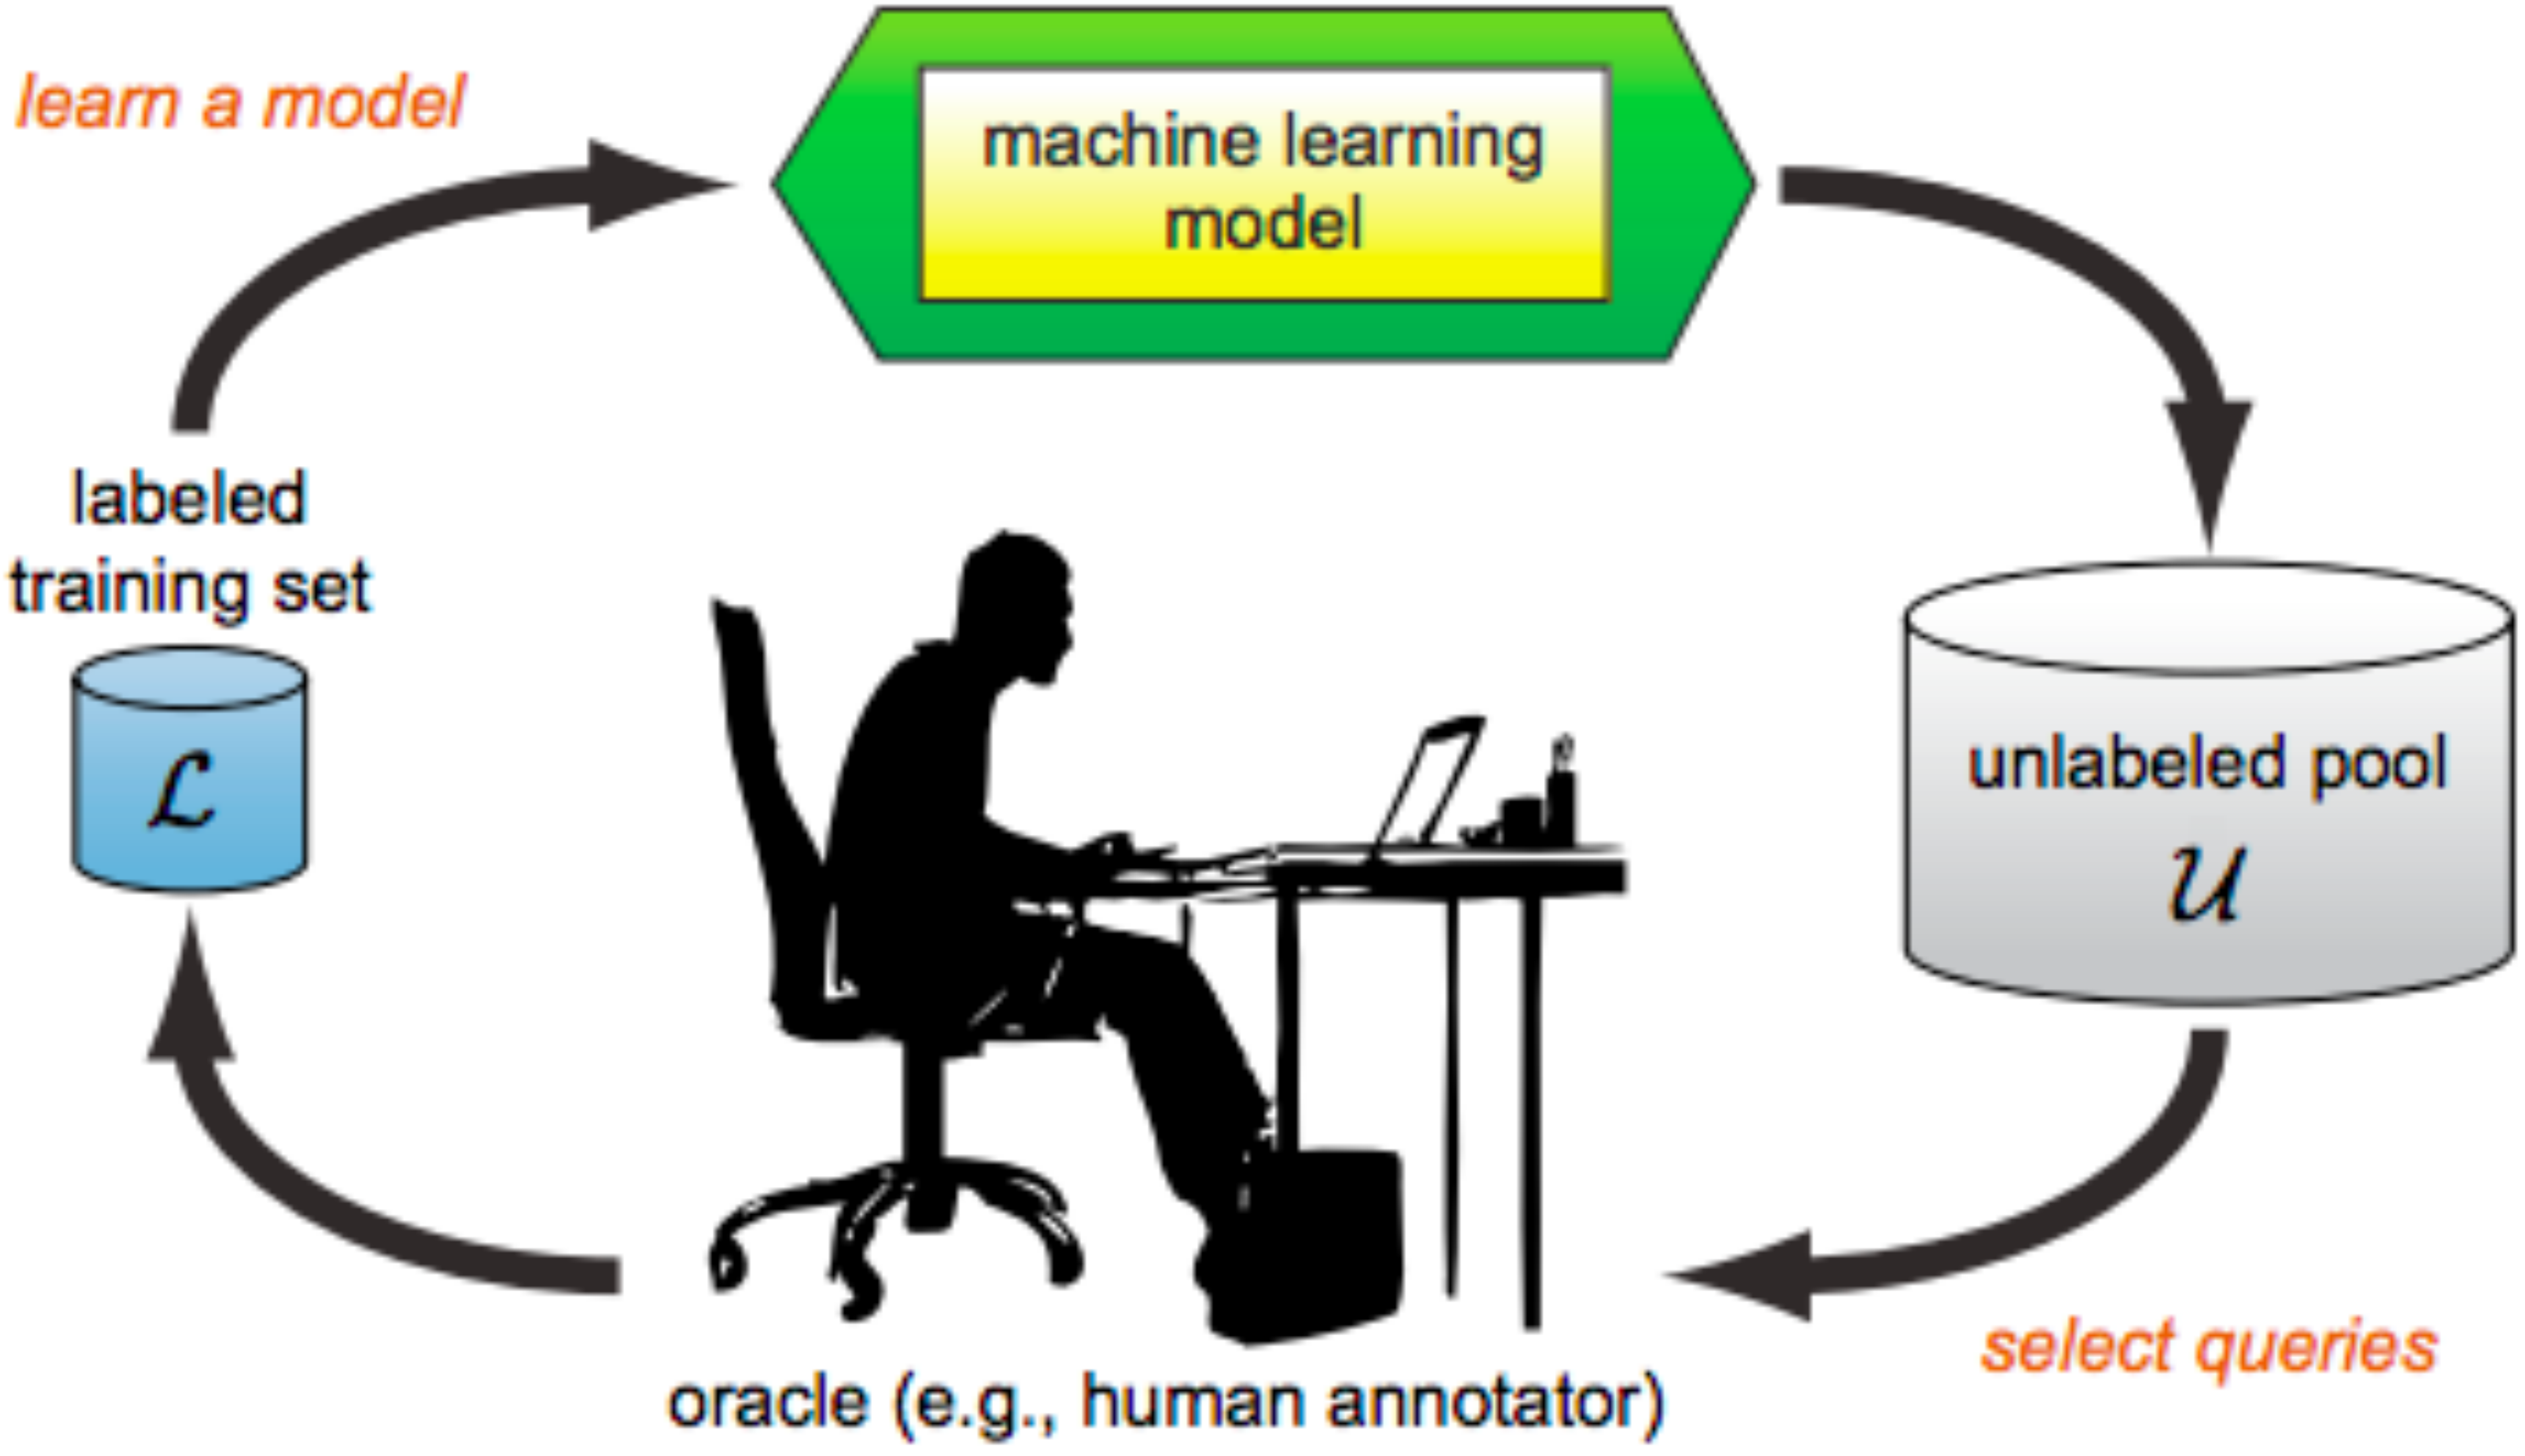

THE ACTIVE LEARNING LOOP WITH PROBABILITY OF IMPROVEMENT

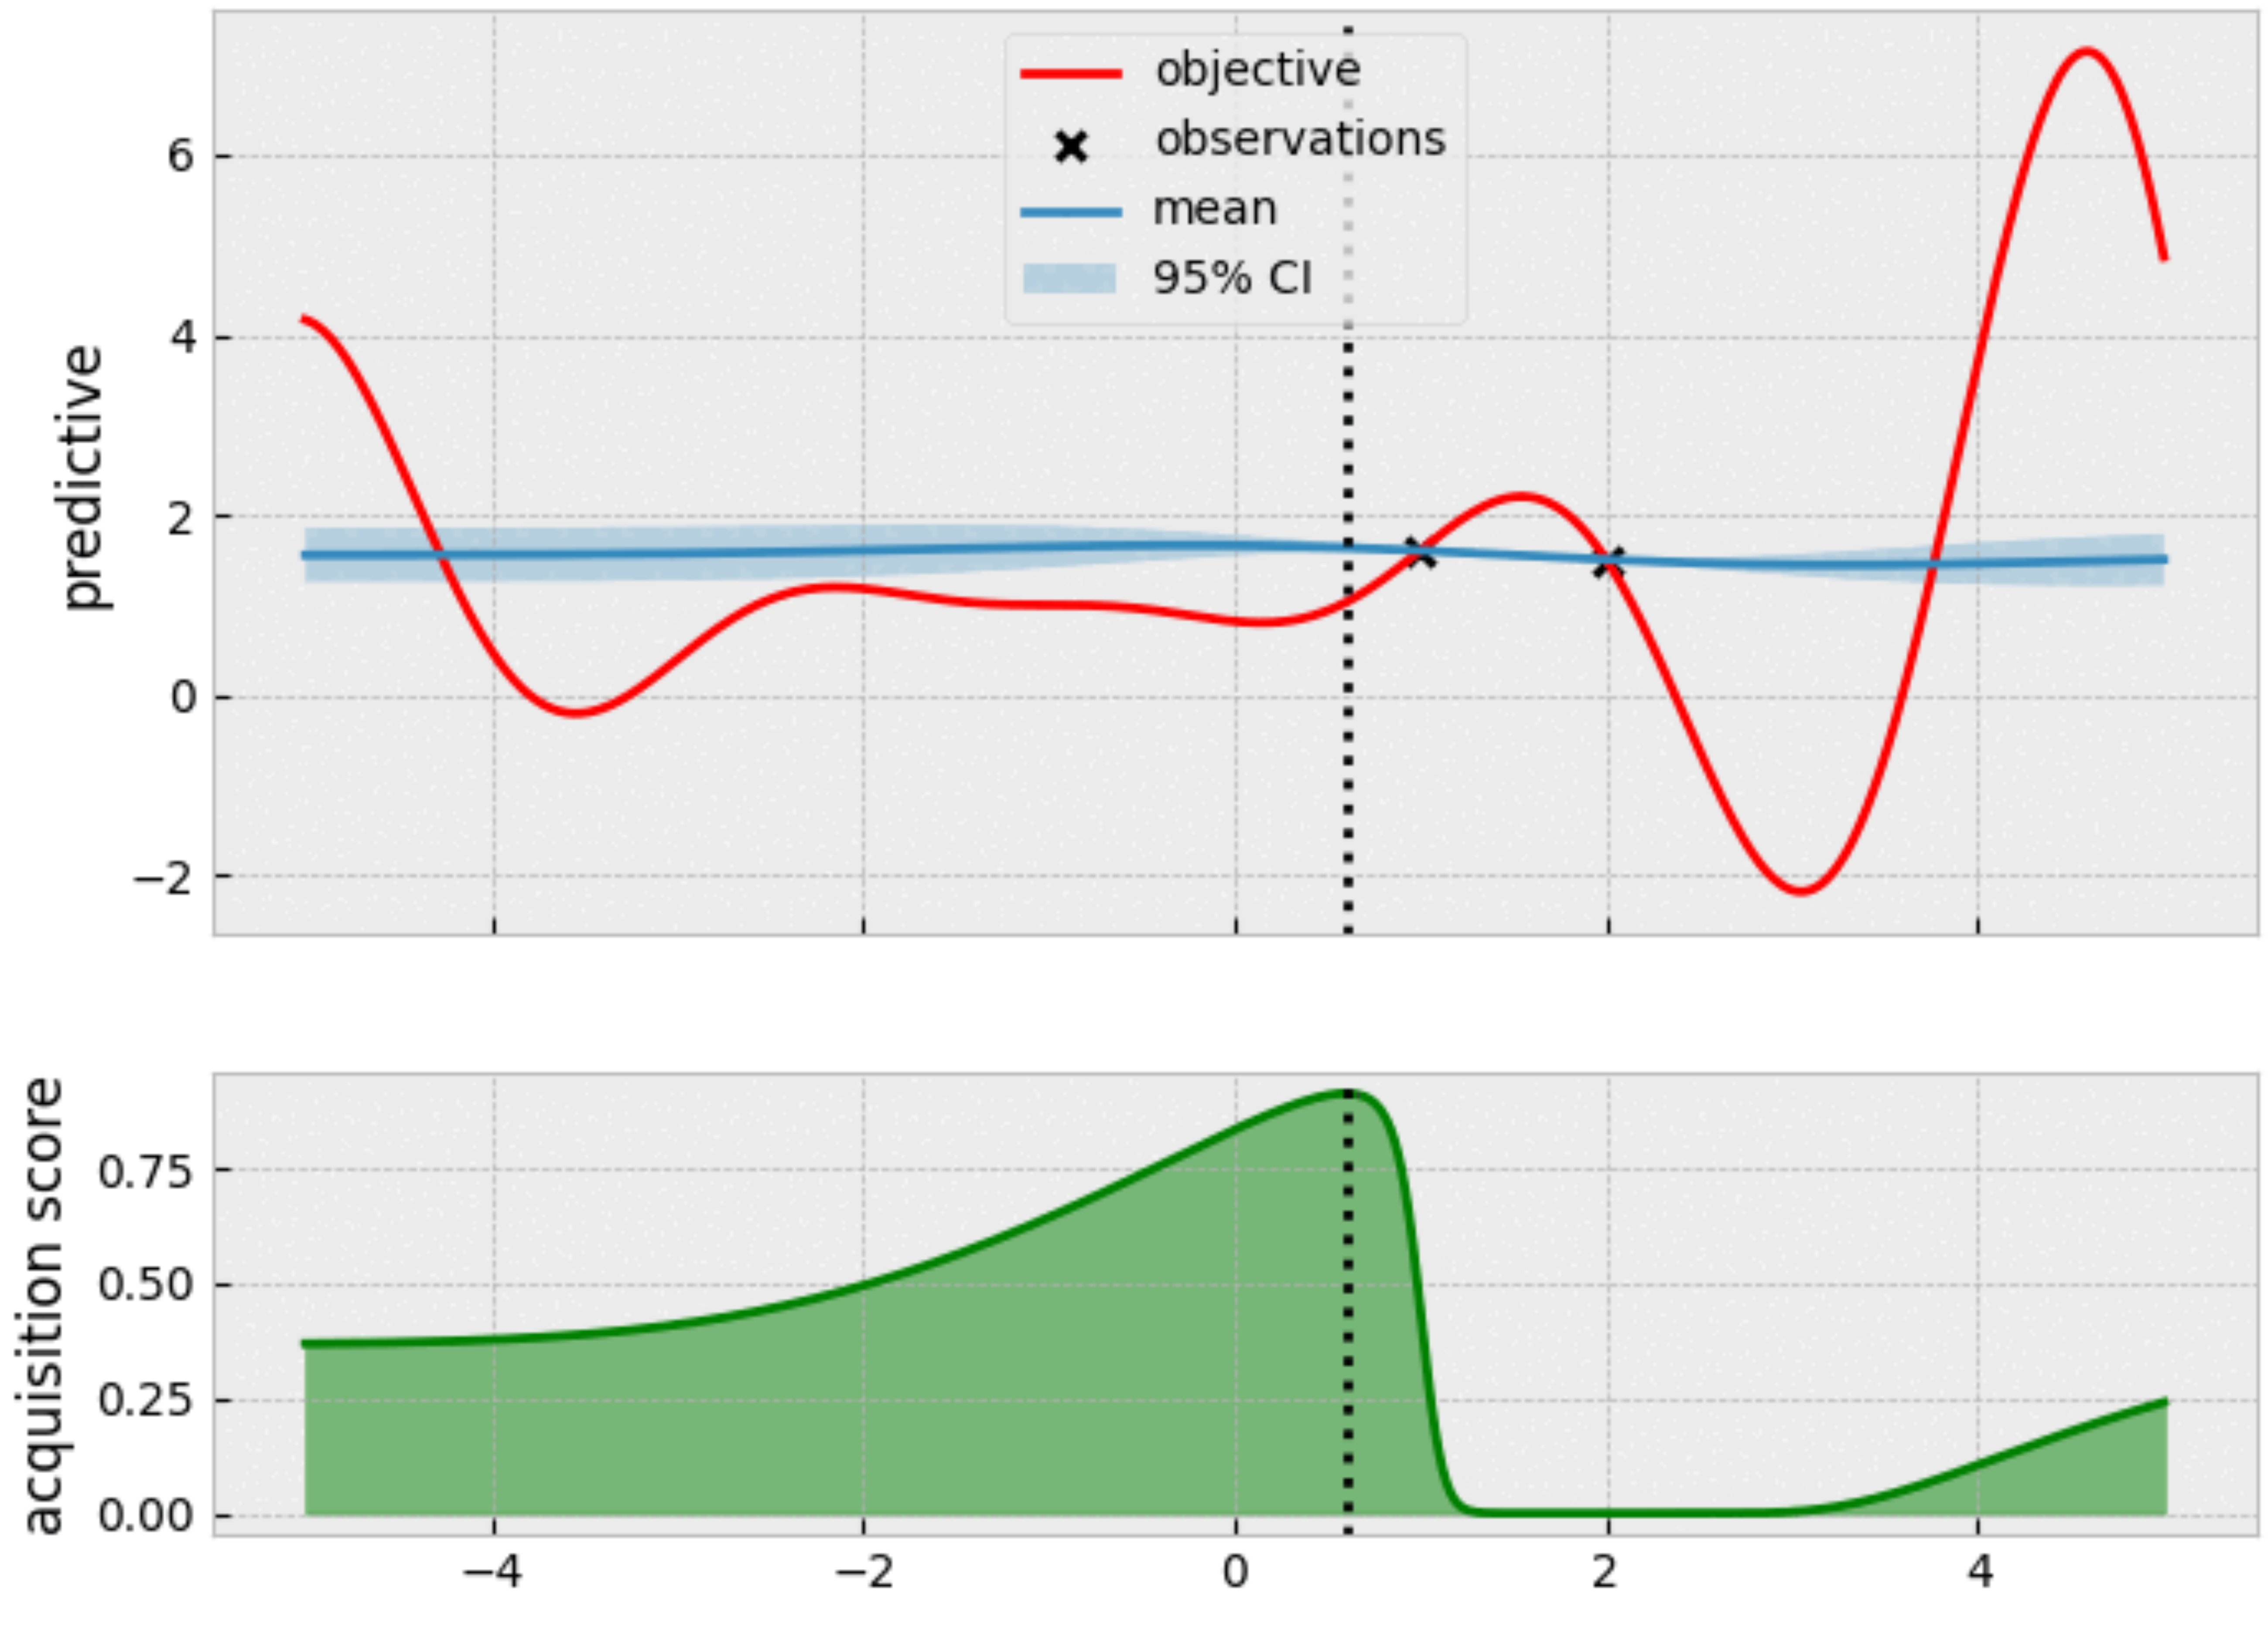

ENCOURAGING EXPLORATION

ENCOURAGING EXPLORATION

Simple improvement encourages exploitation.

ENCOURAGING EXPLORATION

Simple improvement encourages exploitation.

Solution:

ENCOURAGING EXPLORATION

Simple improvement encourages exploitation.

Solution:

1. Set a **stricter definition** of "improvement".

ENCOURAGING EXPLORATION

Simple improvement encourages exploitation.

Solution:

1. Set a **stricter definition** of "improvement".

…	1.61	1.62	…	1.61 + ε	…
…	0	0	…	1	…
…	0	0	…	1	…
…	0	0	…	1	…
…	…	…	…	…	…

ENCOURAGING EXPLORATION

Simple improvement encourages exploitation.

Solution:

1. Set a **stricter definition** of “improvement”.

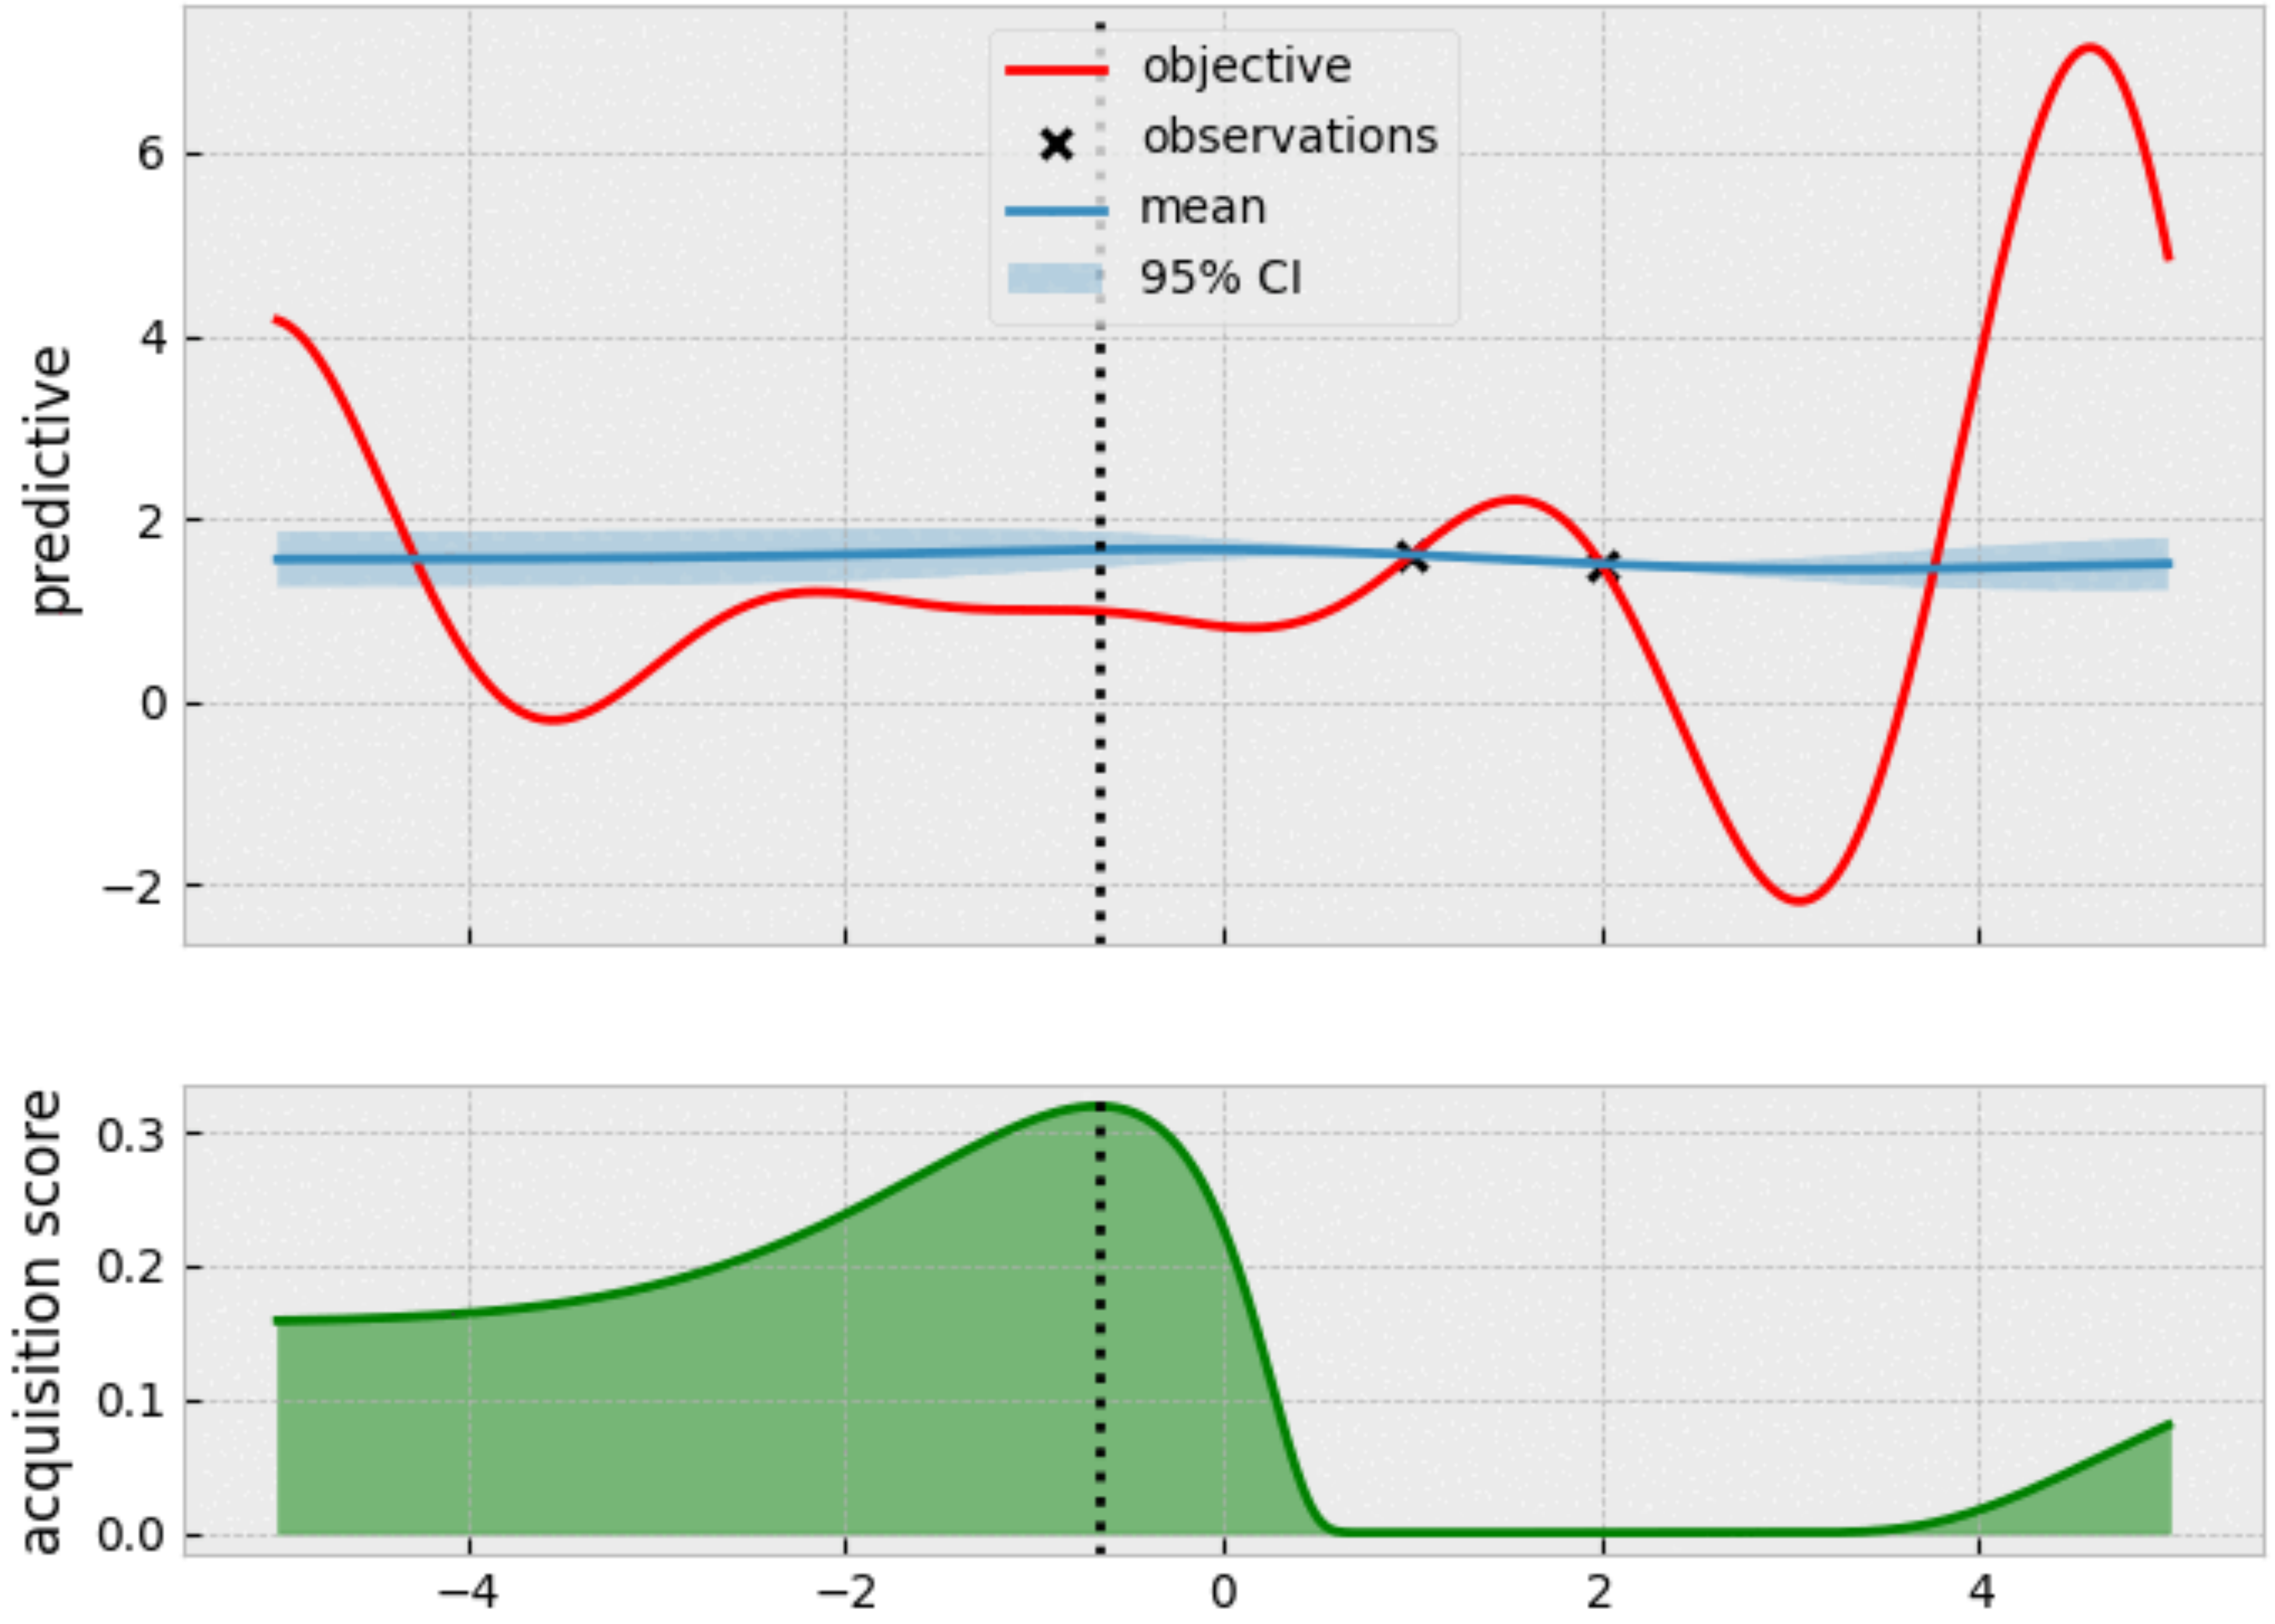

ENCOURAGING EXPLORATION

Simple improvement encourages exploitation.

Solution:

1. Set a **stricter definition** of "improvement".

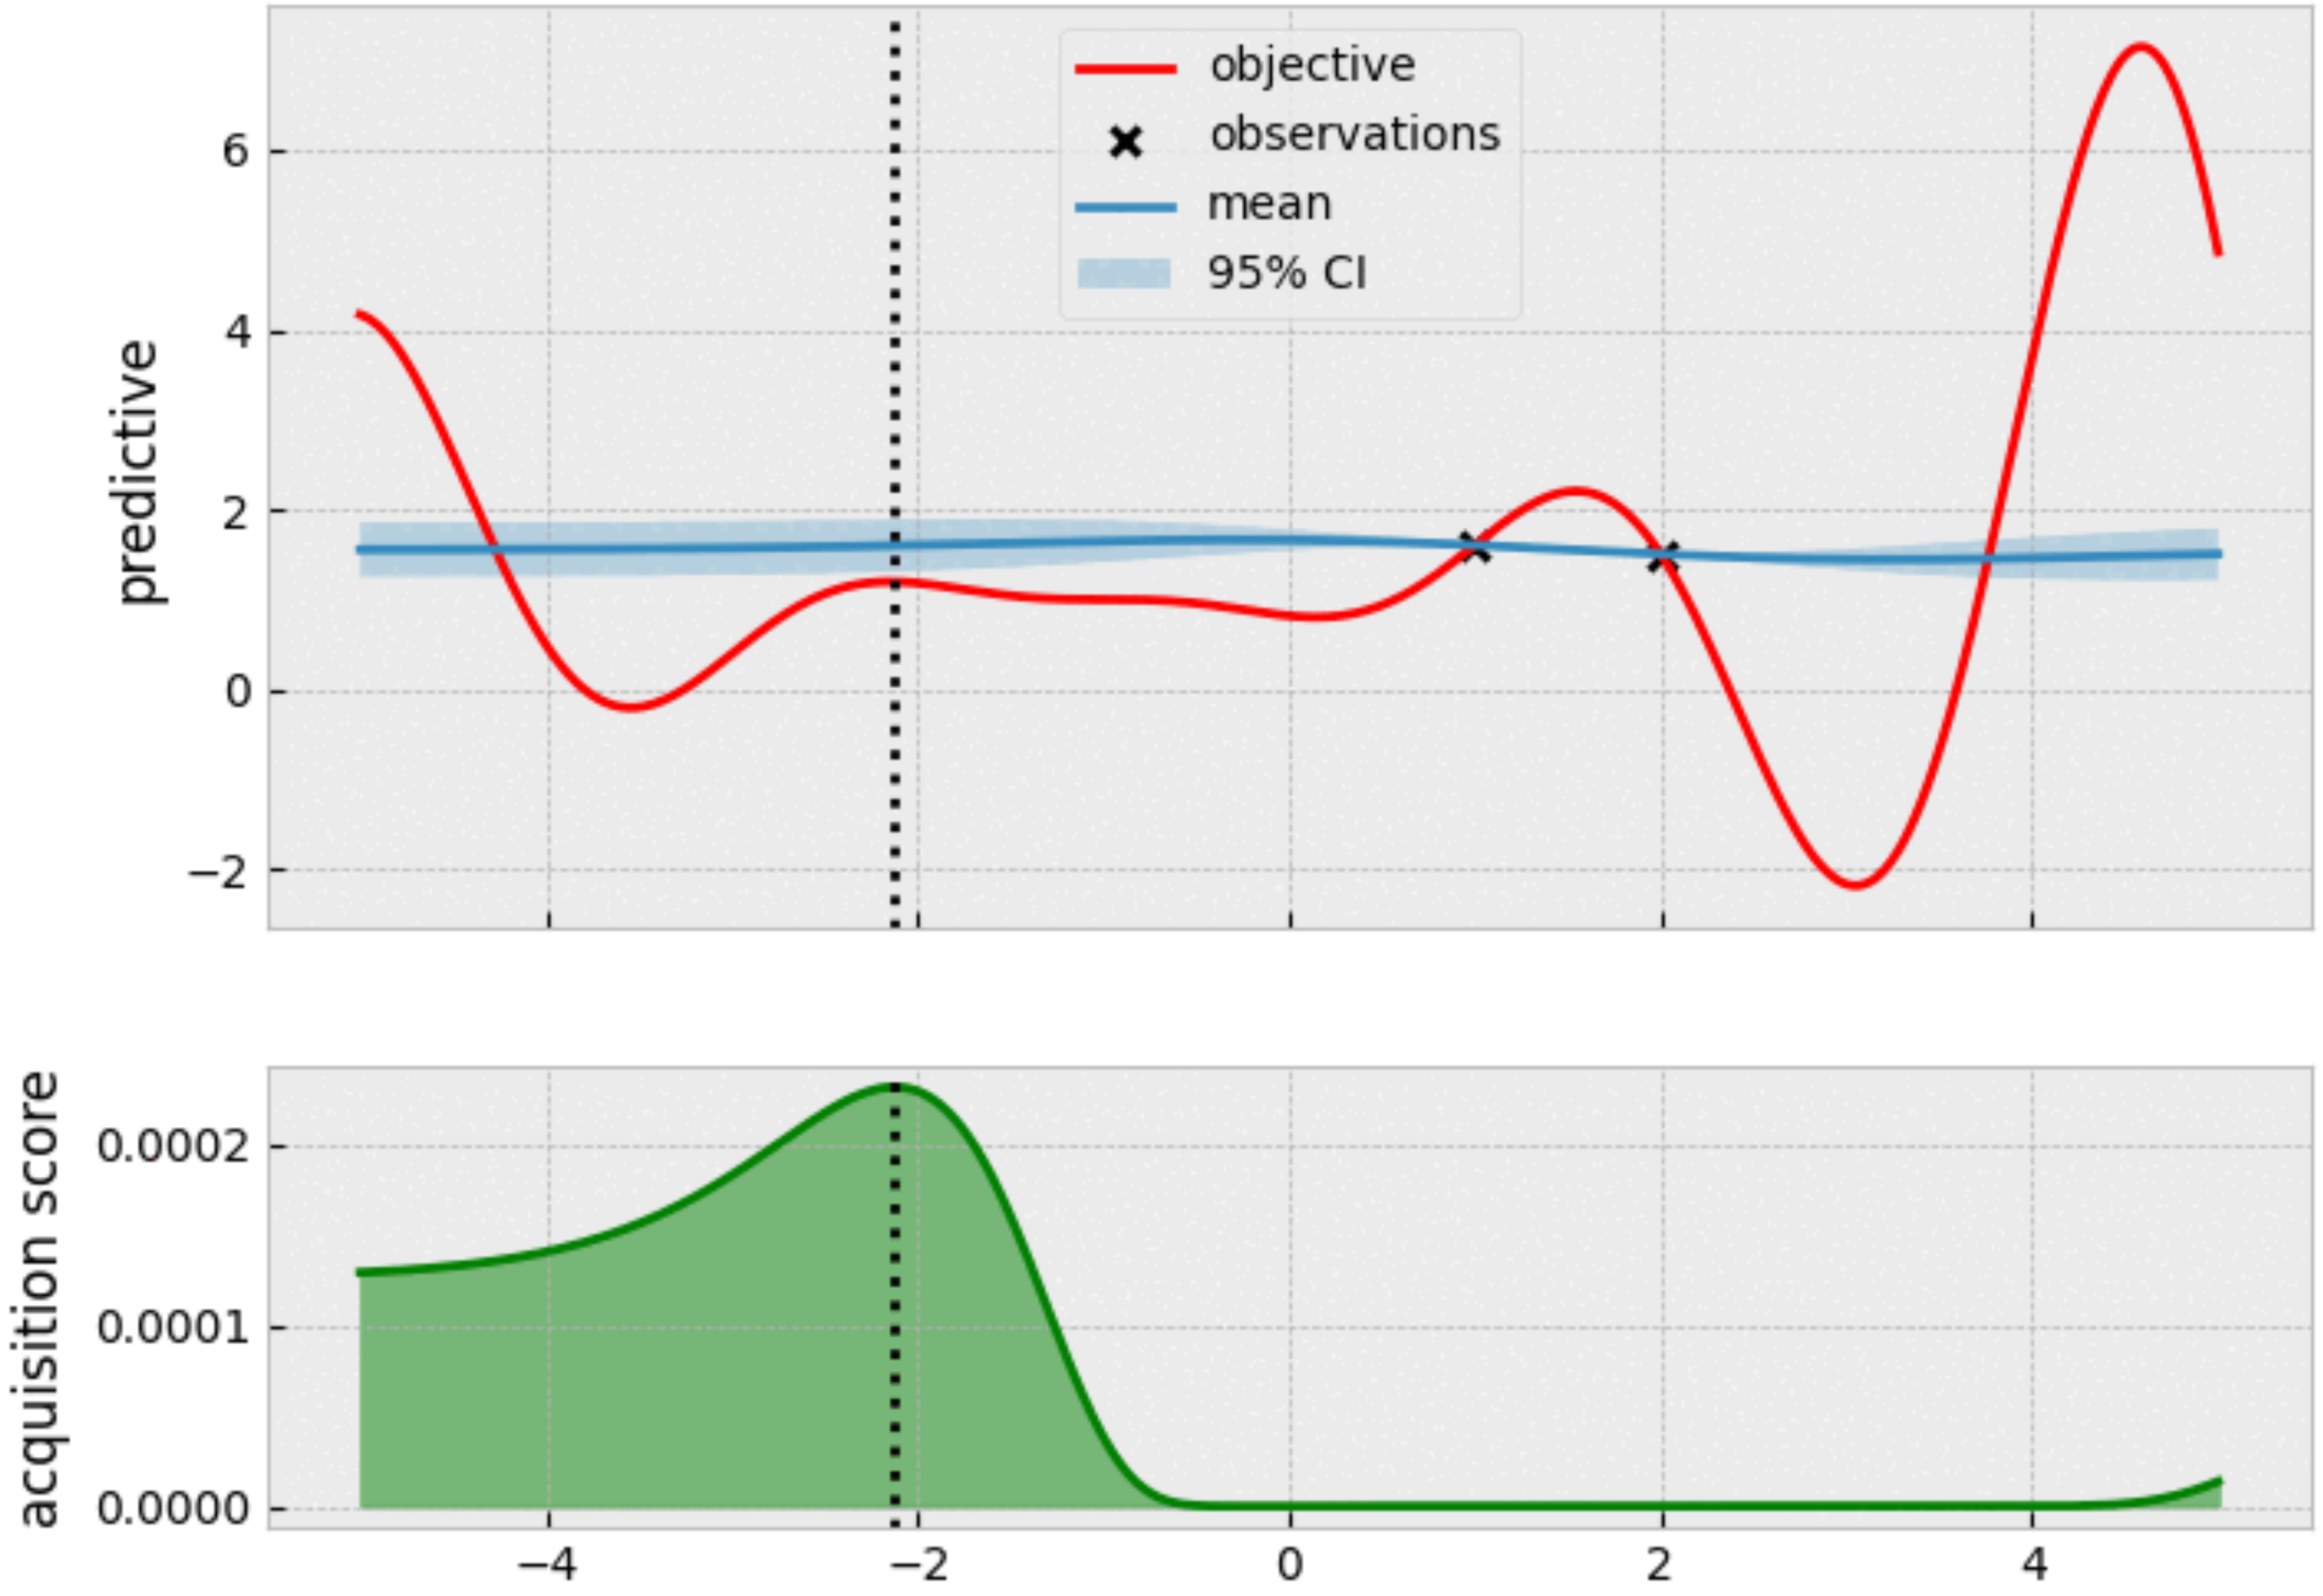

ENCOURAGING EXPLORATION

Simple improvement encourages exploitation.

Solution:

1. Set a **stricter definition** of "improvement".
2. **Redefine** utility

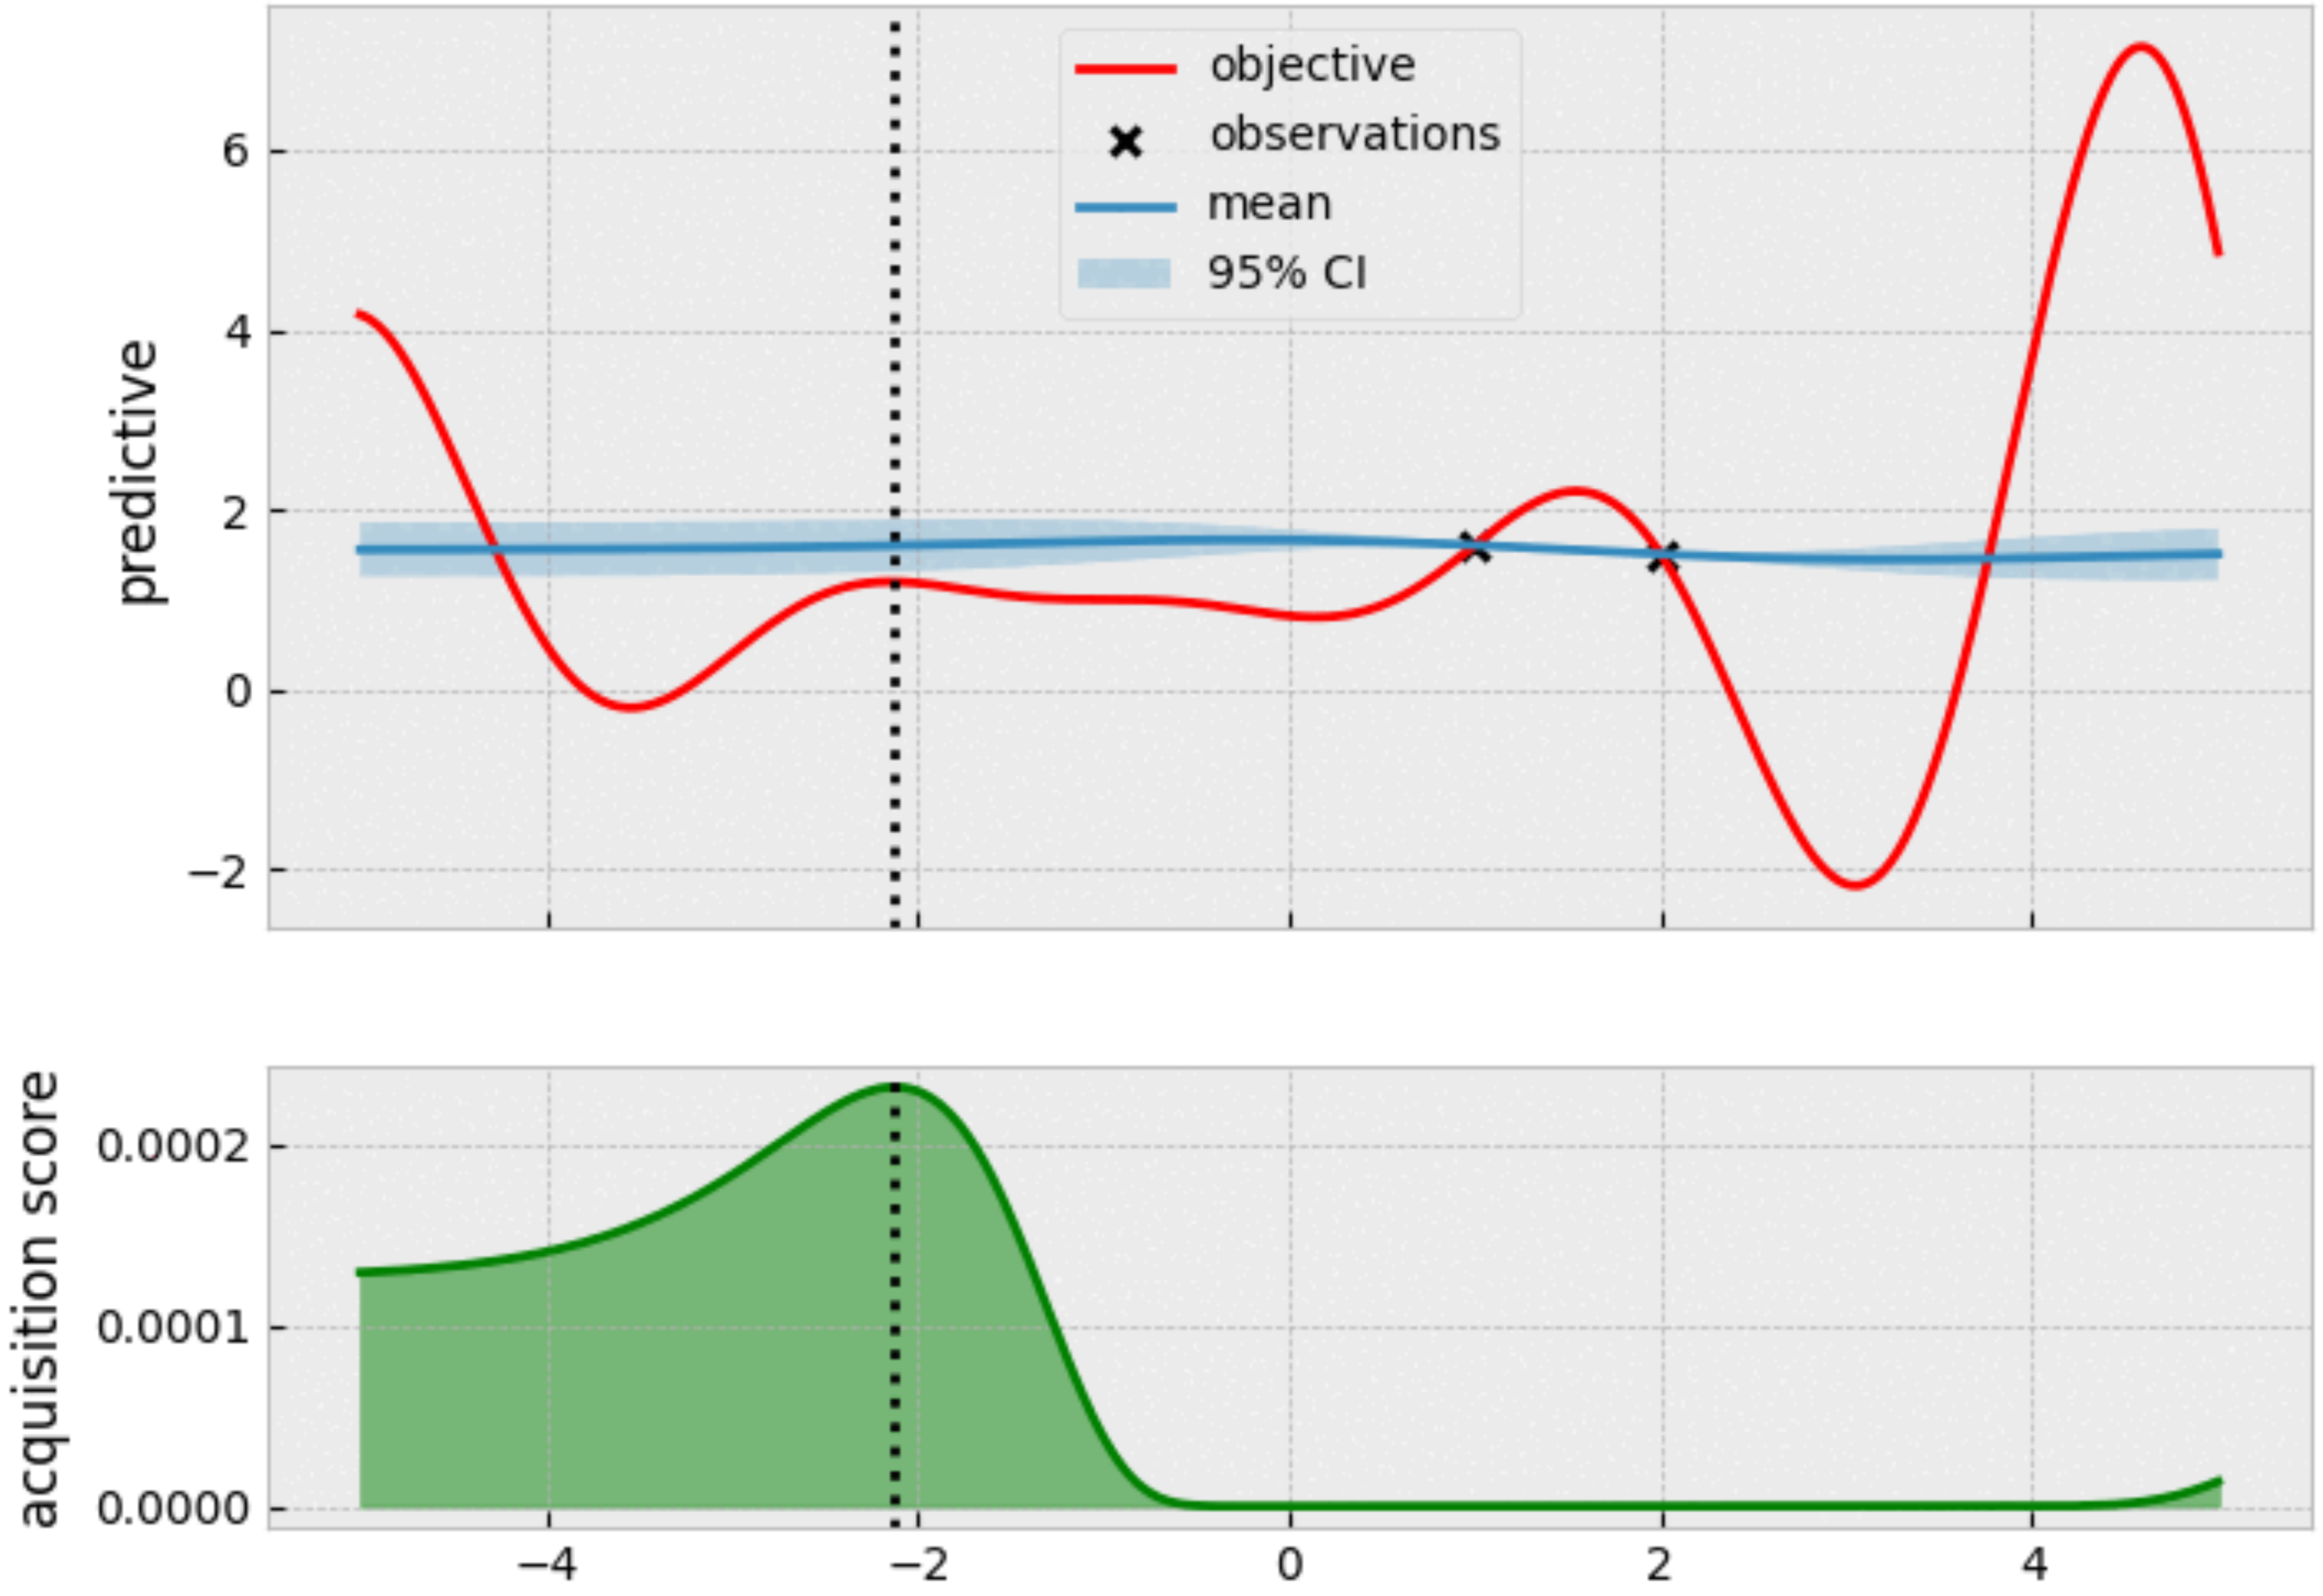

HOW MUCH TO IMPROVE FROM THE INCUMBENT

Utility: how much the incumbent improves

HOW MUCH TO IMPROVE FROM THE INCUMBENT

Utility: how much the incumbent improves

	…	1.6054	1.6055	1.6056	…	2	…
0	…	0	0.0001	0.0002	…	0.3946	…
0.0001	…	0	0.0001	0.0002	…	0.3946	…
0.0002	…	0	0.0001	0.0002	…	0.3946	…
…	…	…	…	…	…	…	…

FANTASIZING ABOUT IMPROVEMENT

$\boxed{x}$

FANTASIZING ABOUT IMPROVEMENT

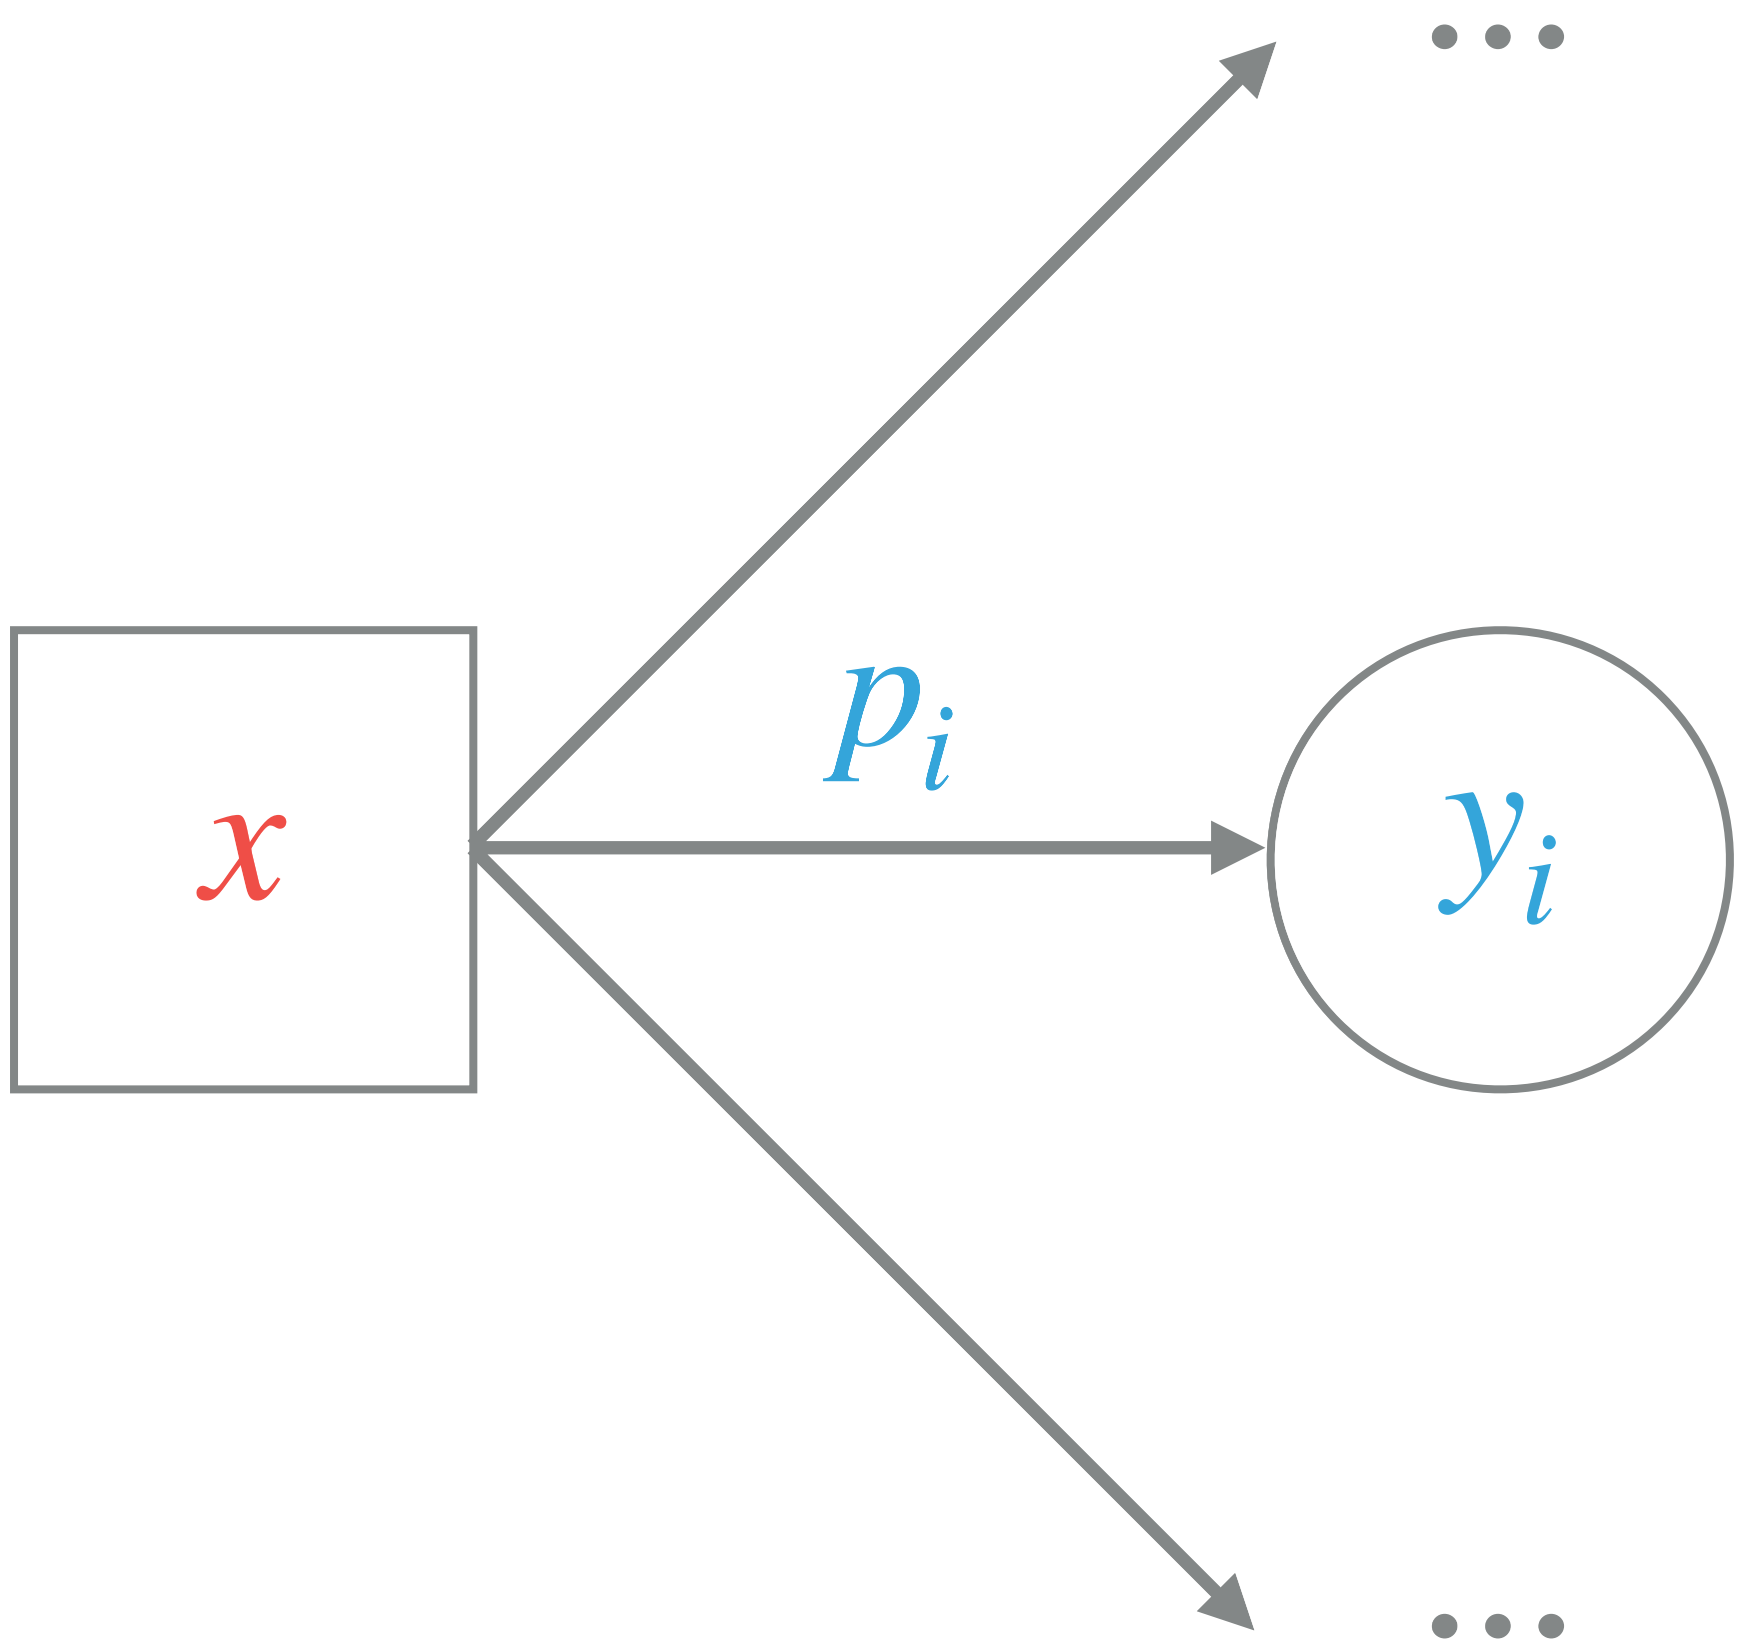

FANTASIZING ABOUT IMPROVEMENT

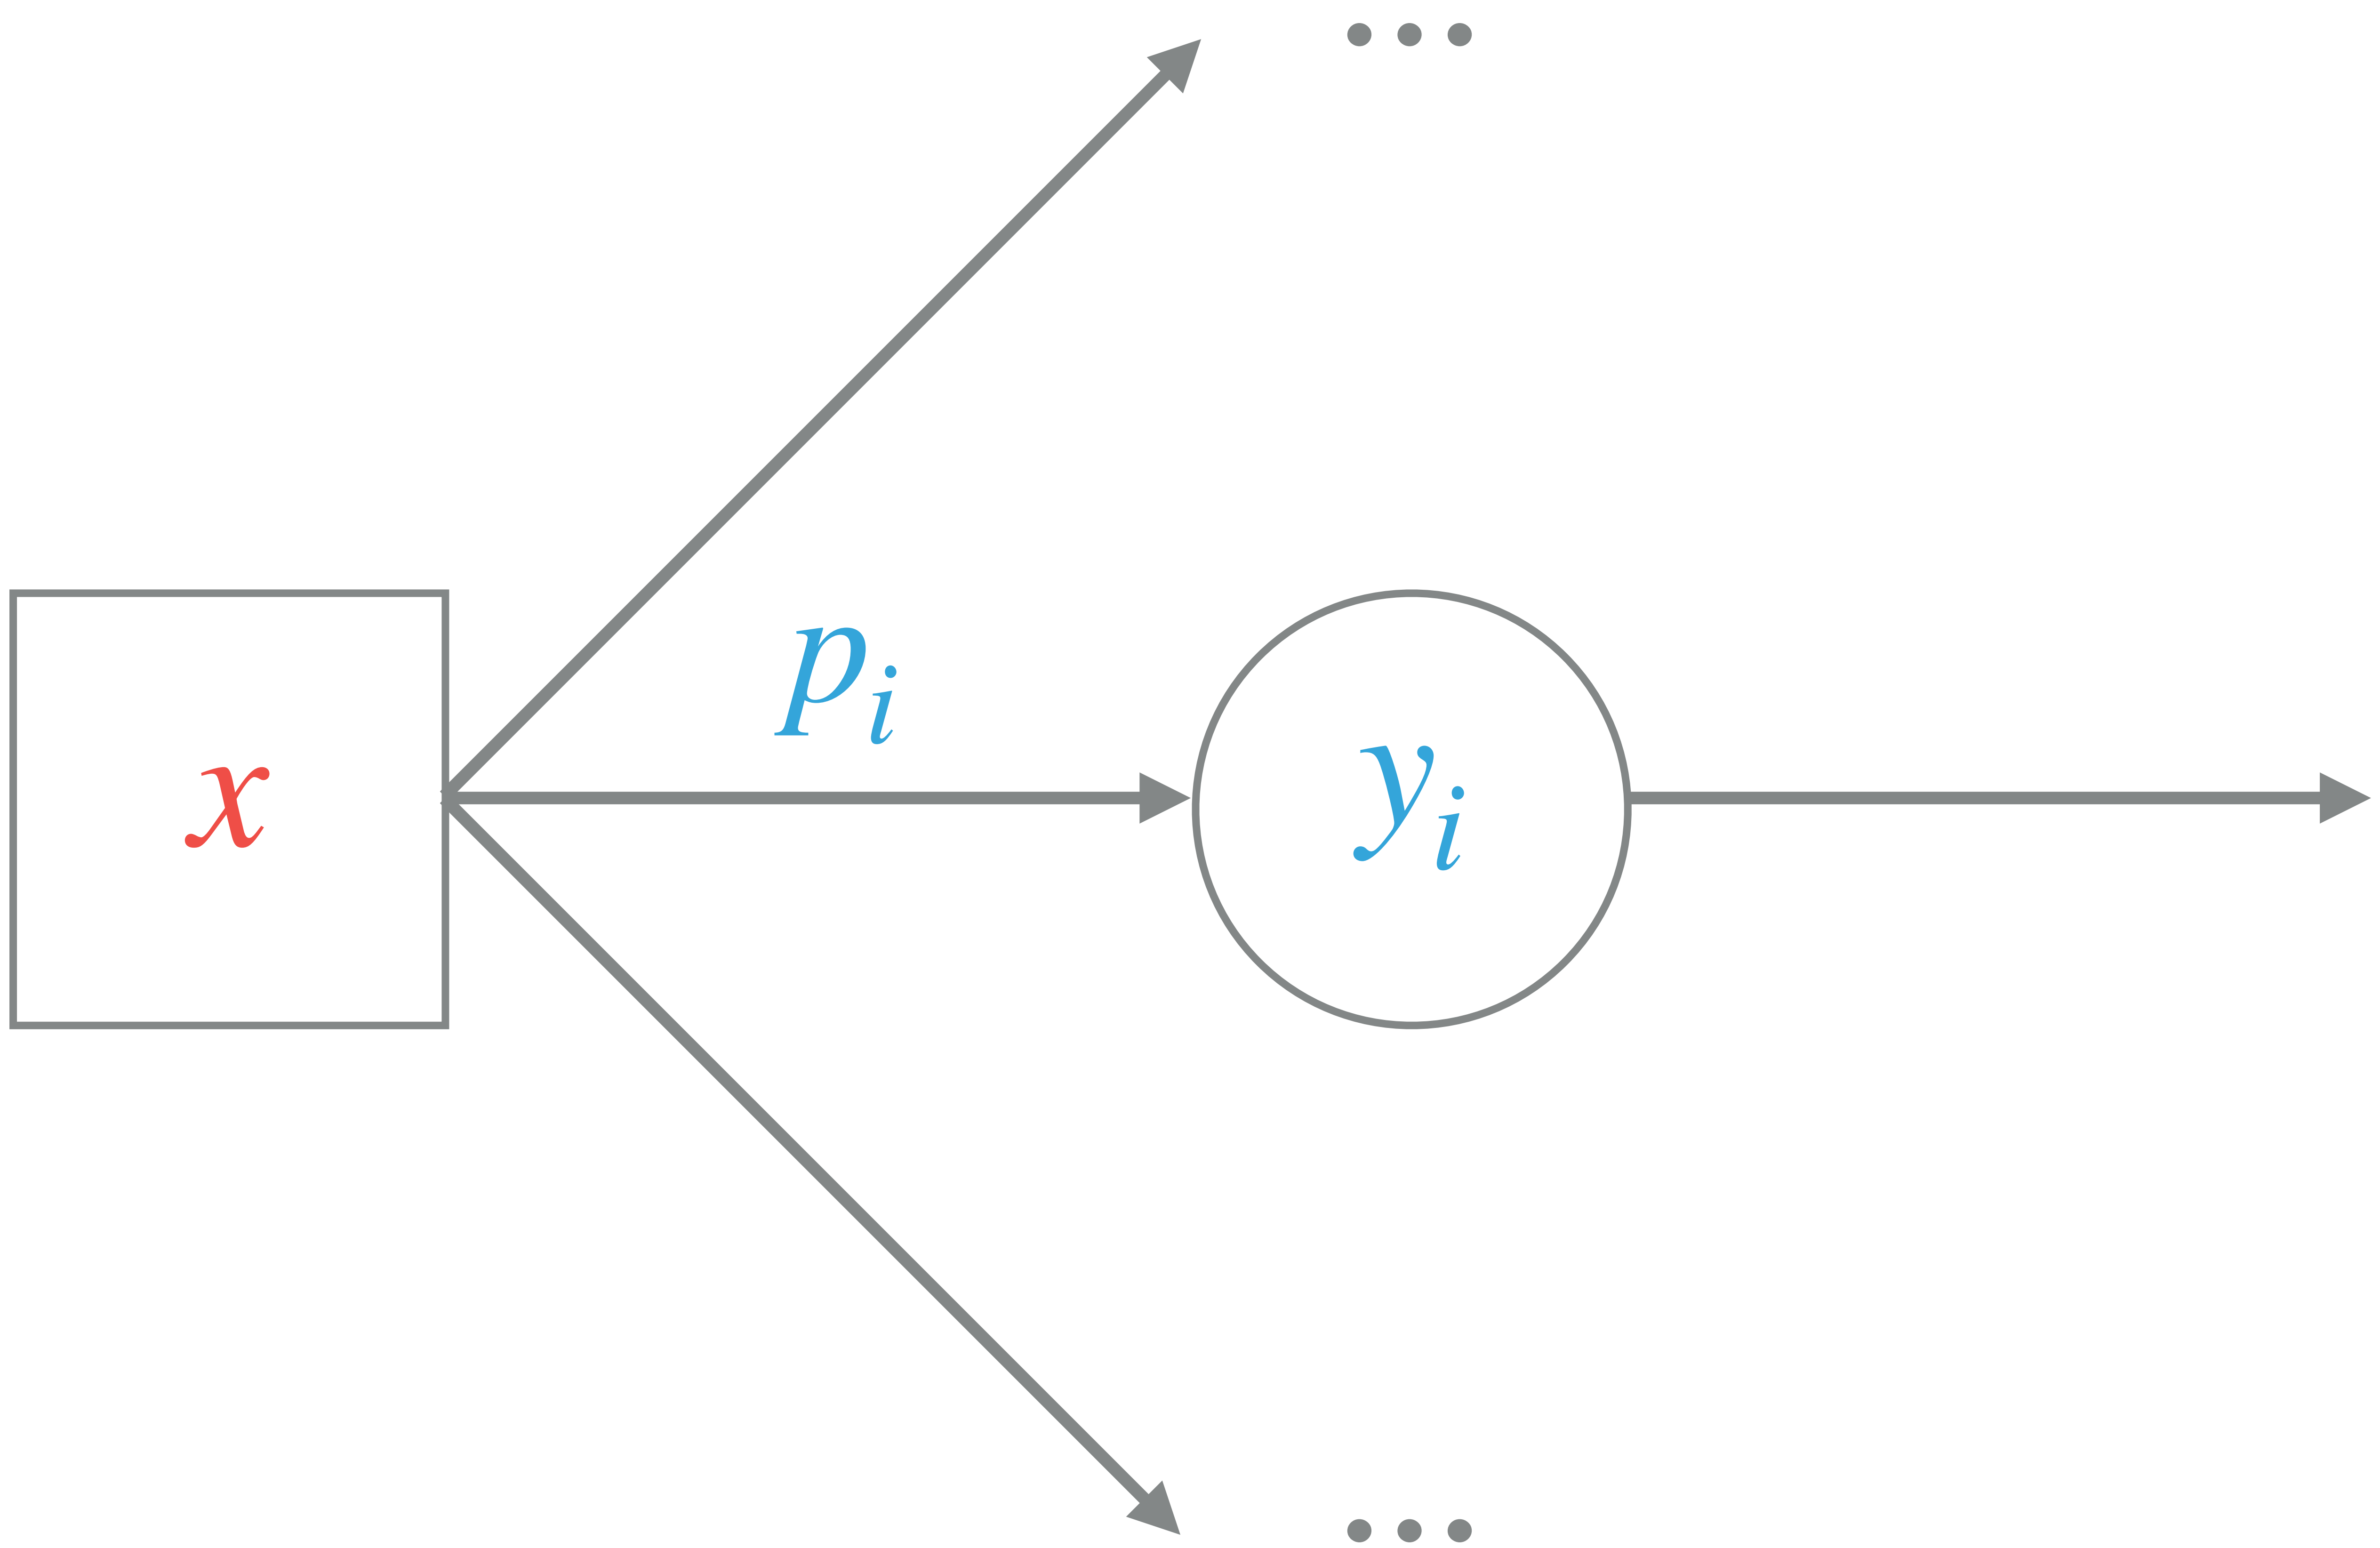

utility $\max\left(0,\ y_i - \text{incumbent}\right)$

FANTASIZING ABOUT IMPROVEMENT

avg. utility of x

$$\mathbb{E}_{y_j}\left[\max\left(0,\ y_i - \text{incumbent}\right)\right]$$

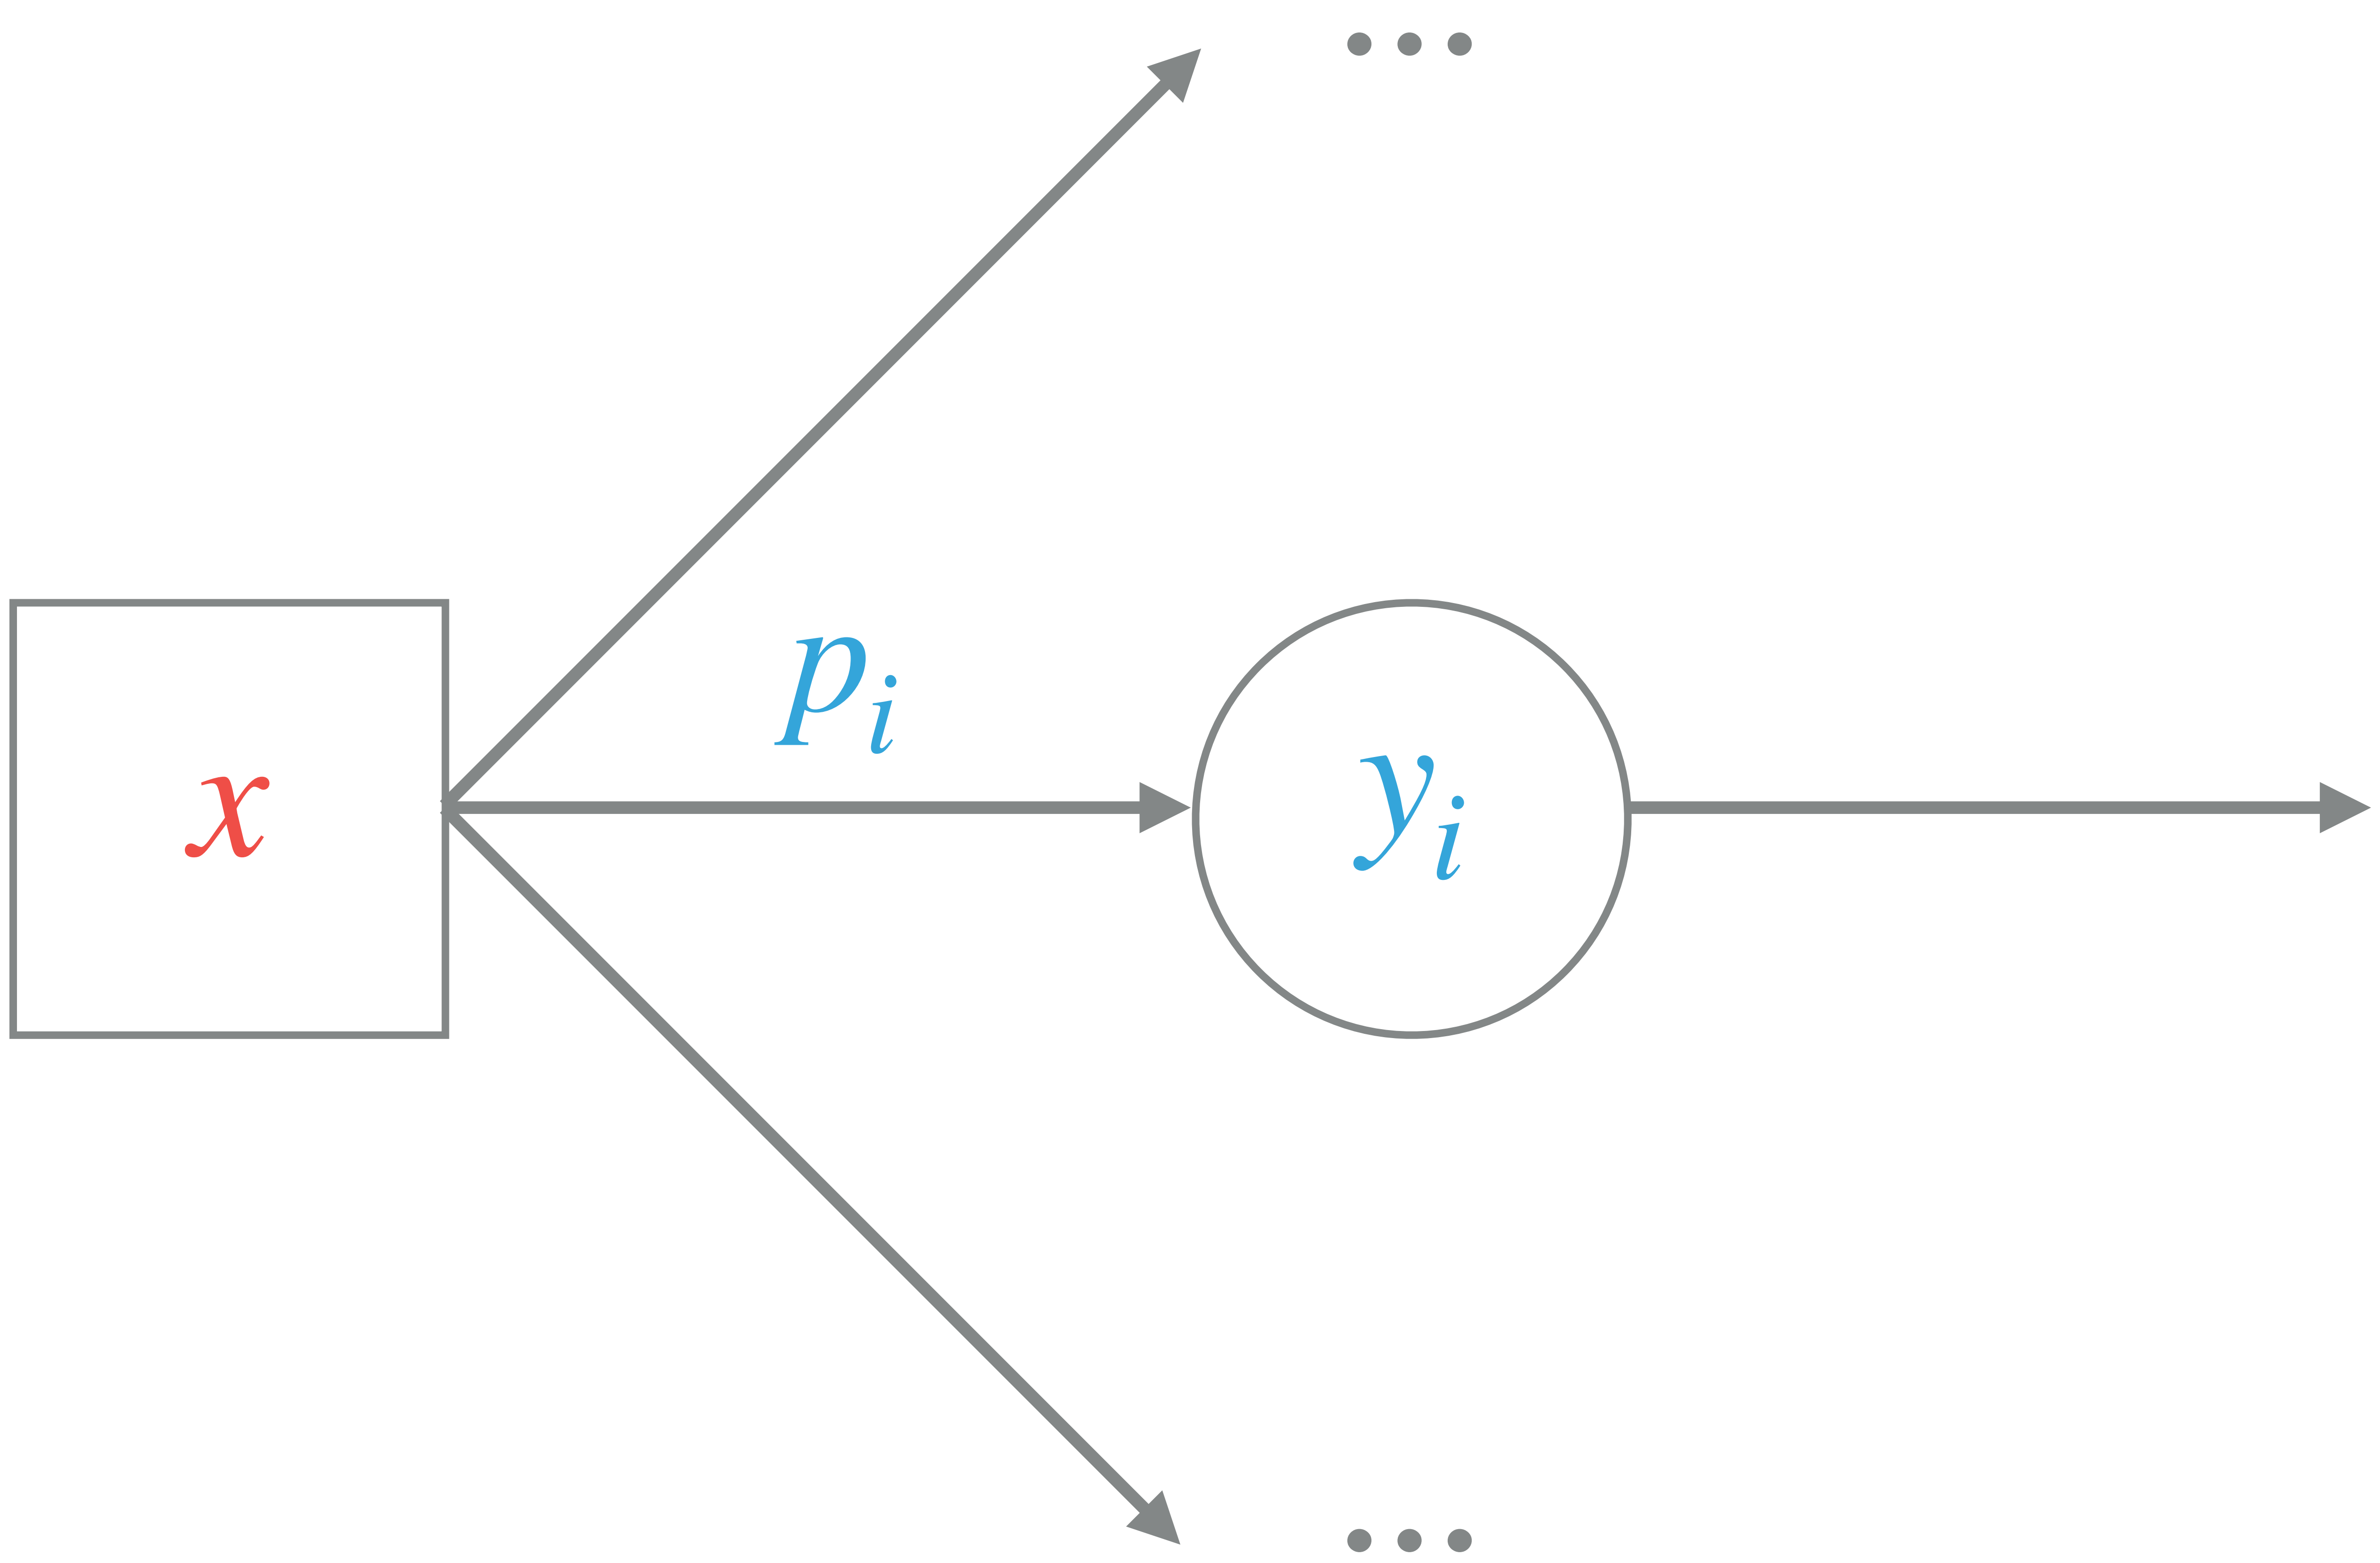

utility $\max\left(0,\ y_i - \text{incumbent}\right)$

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

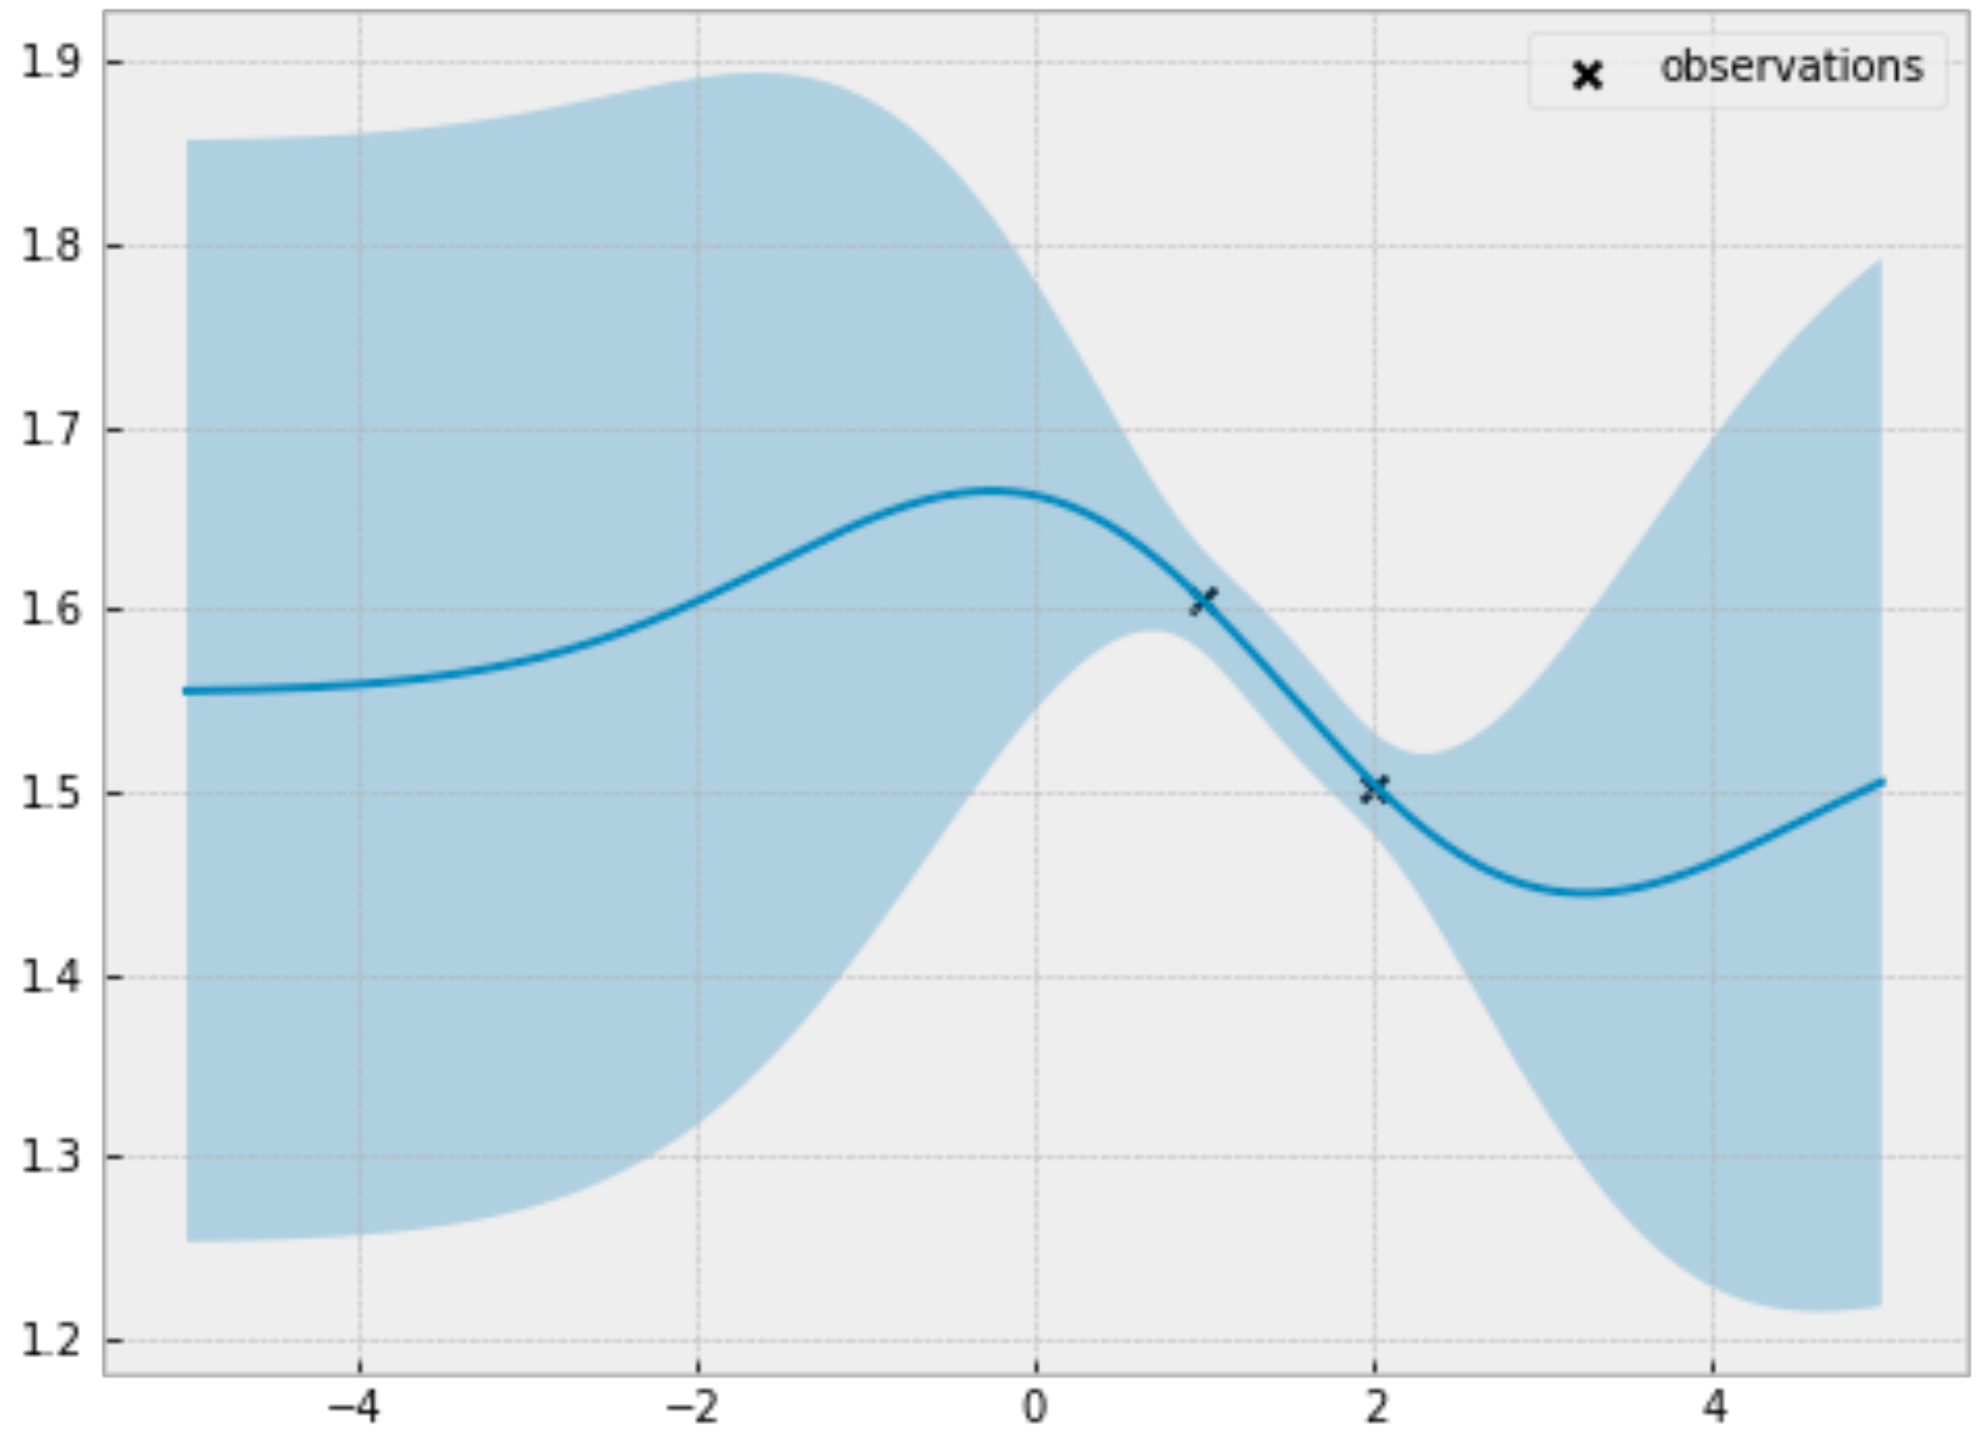

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

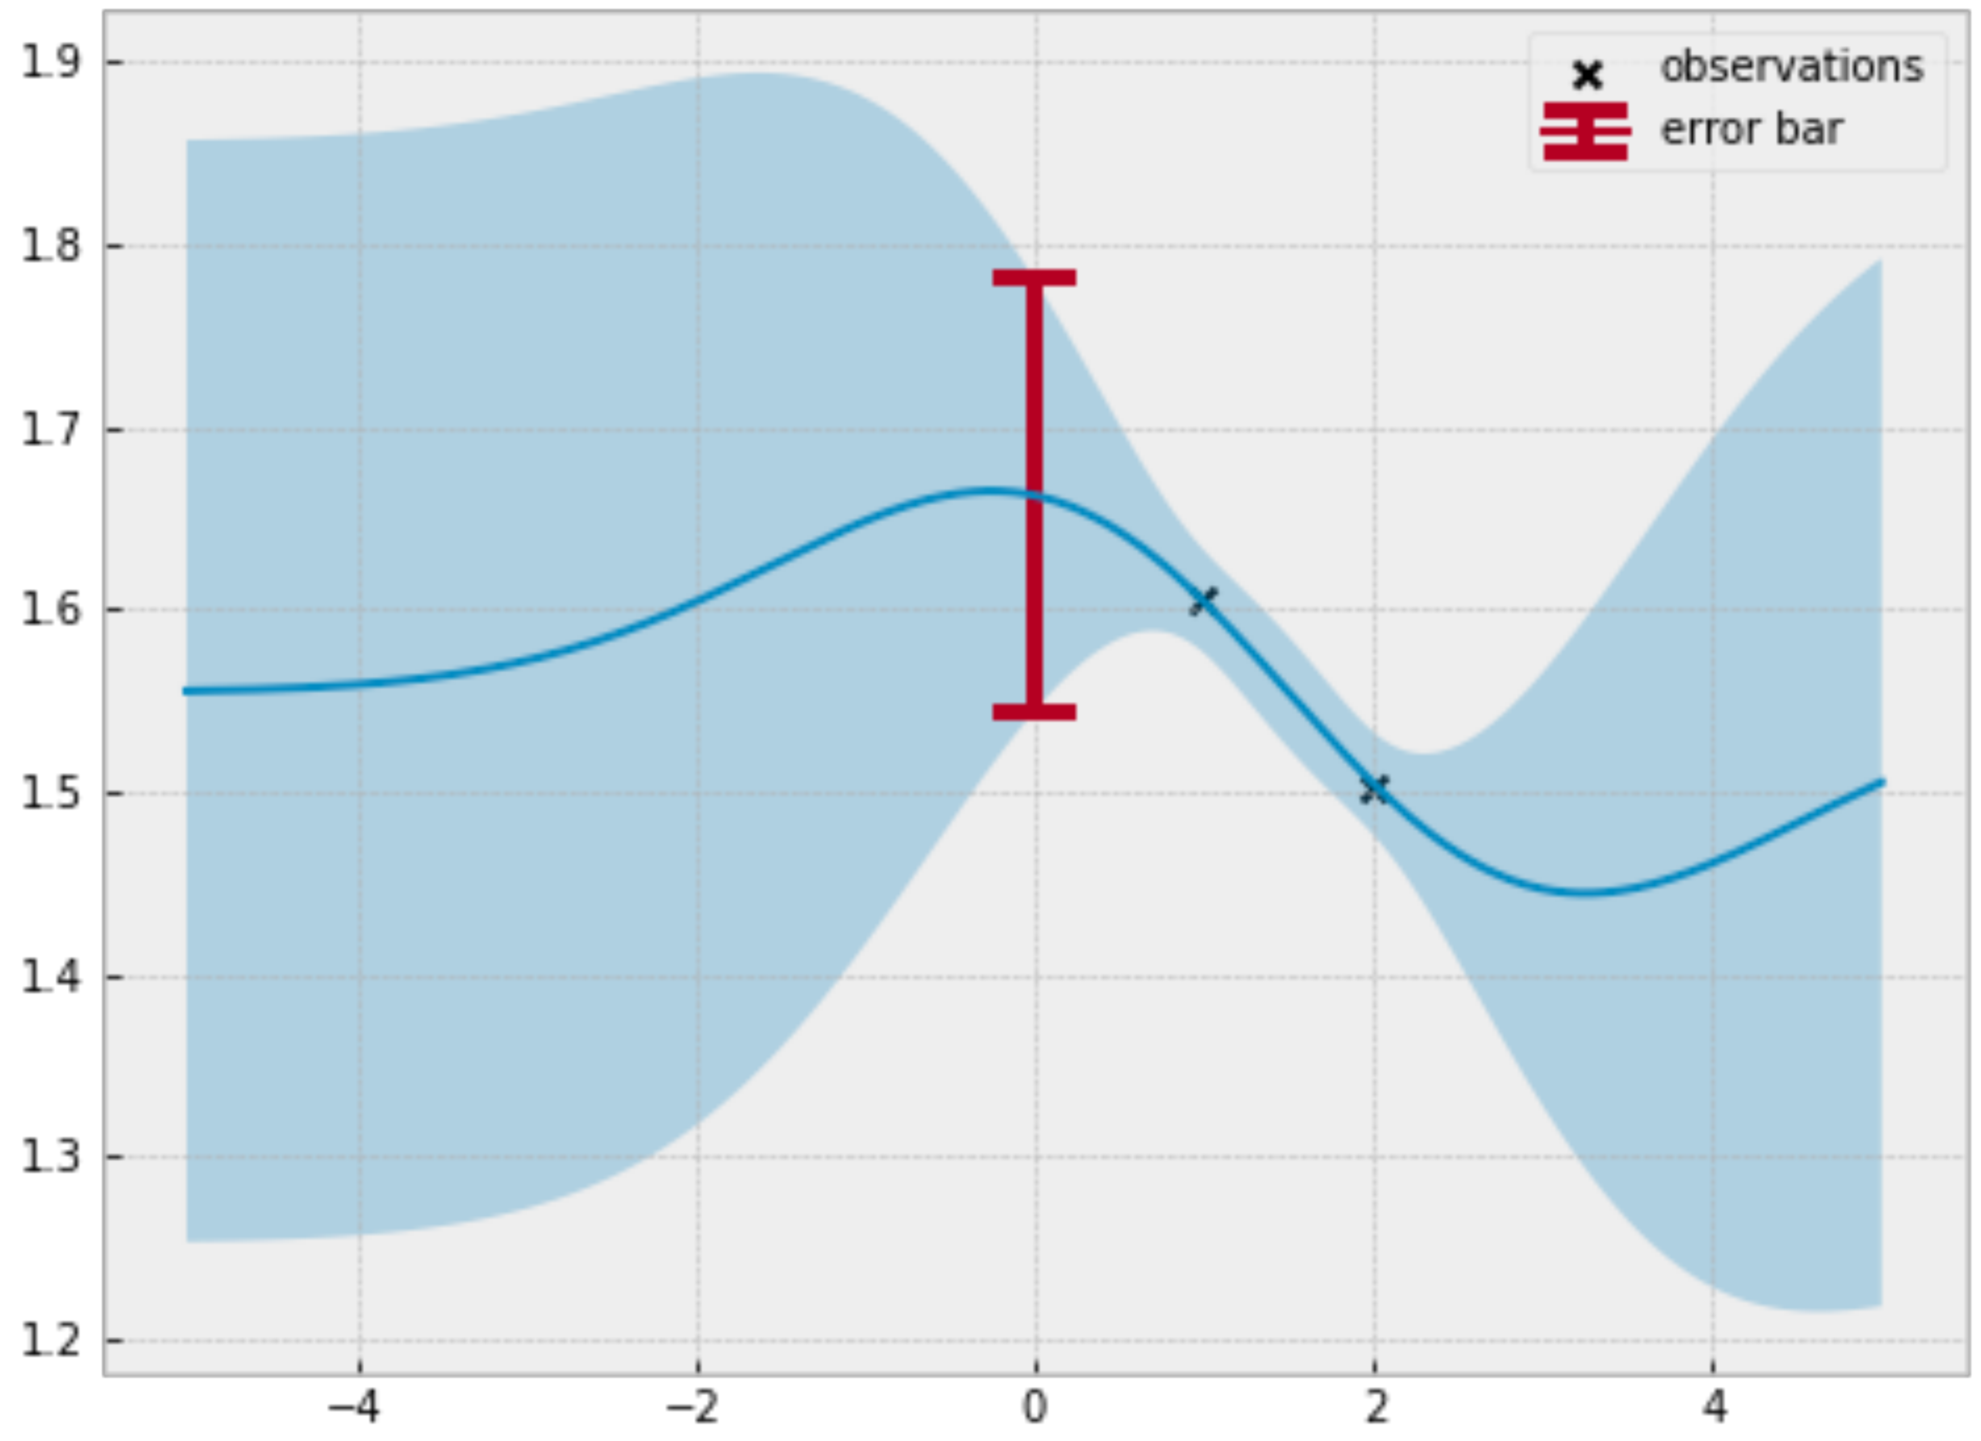

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

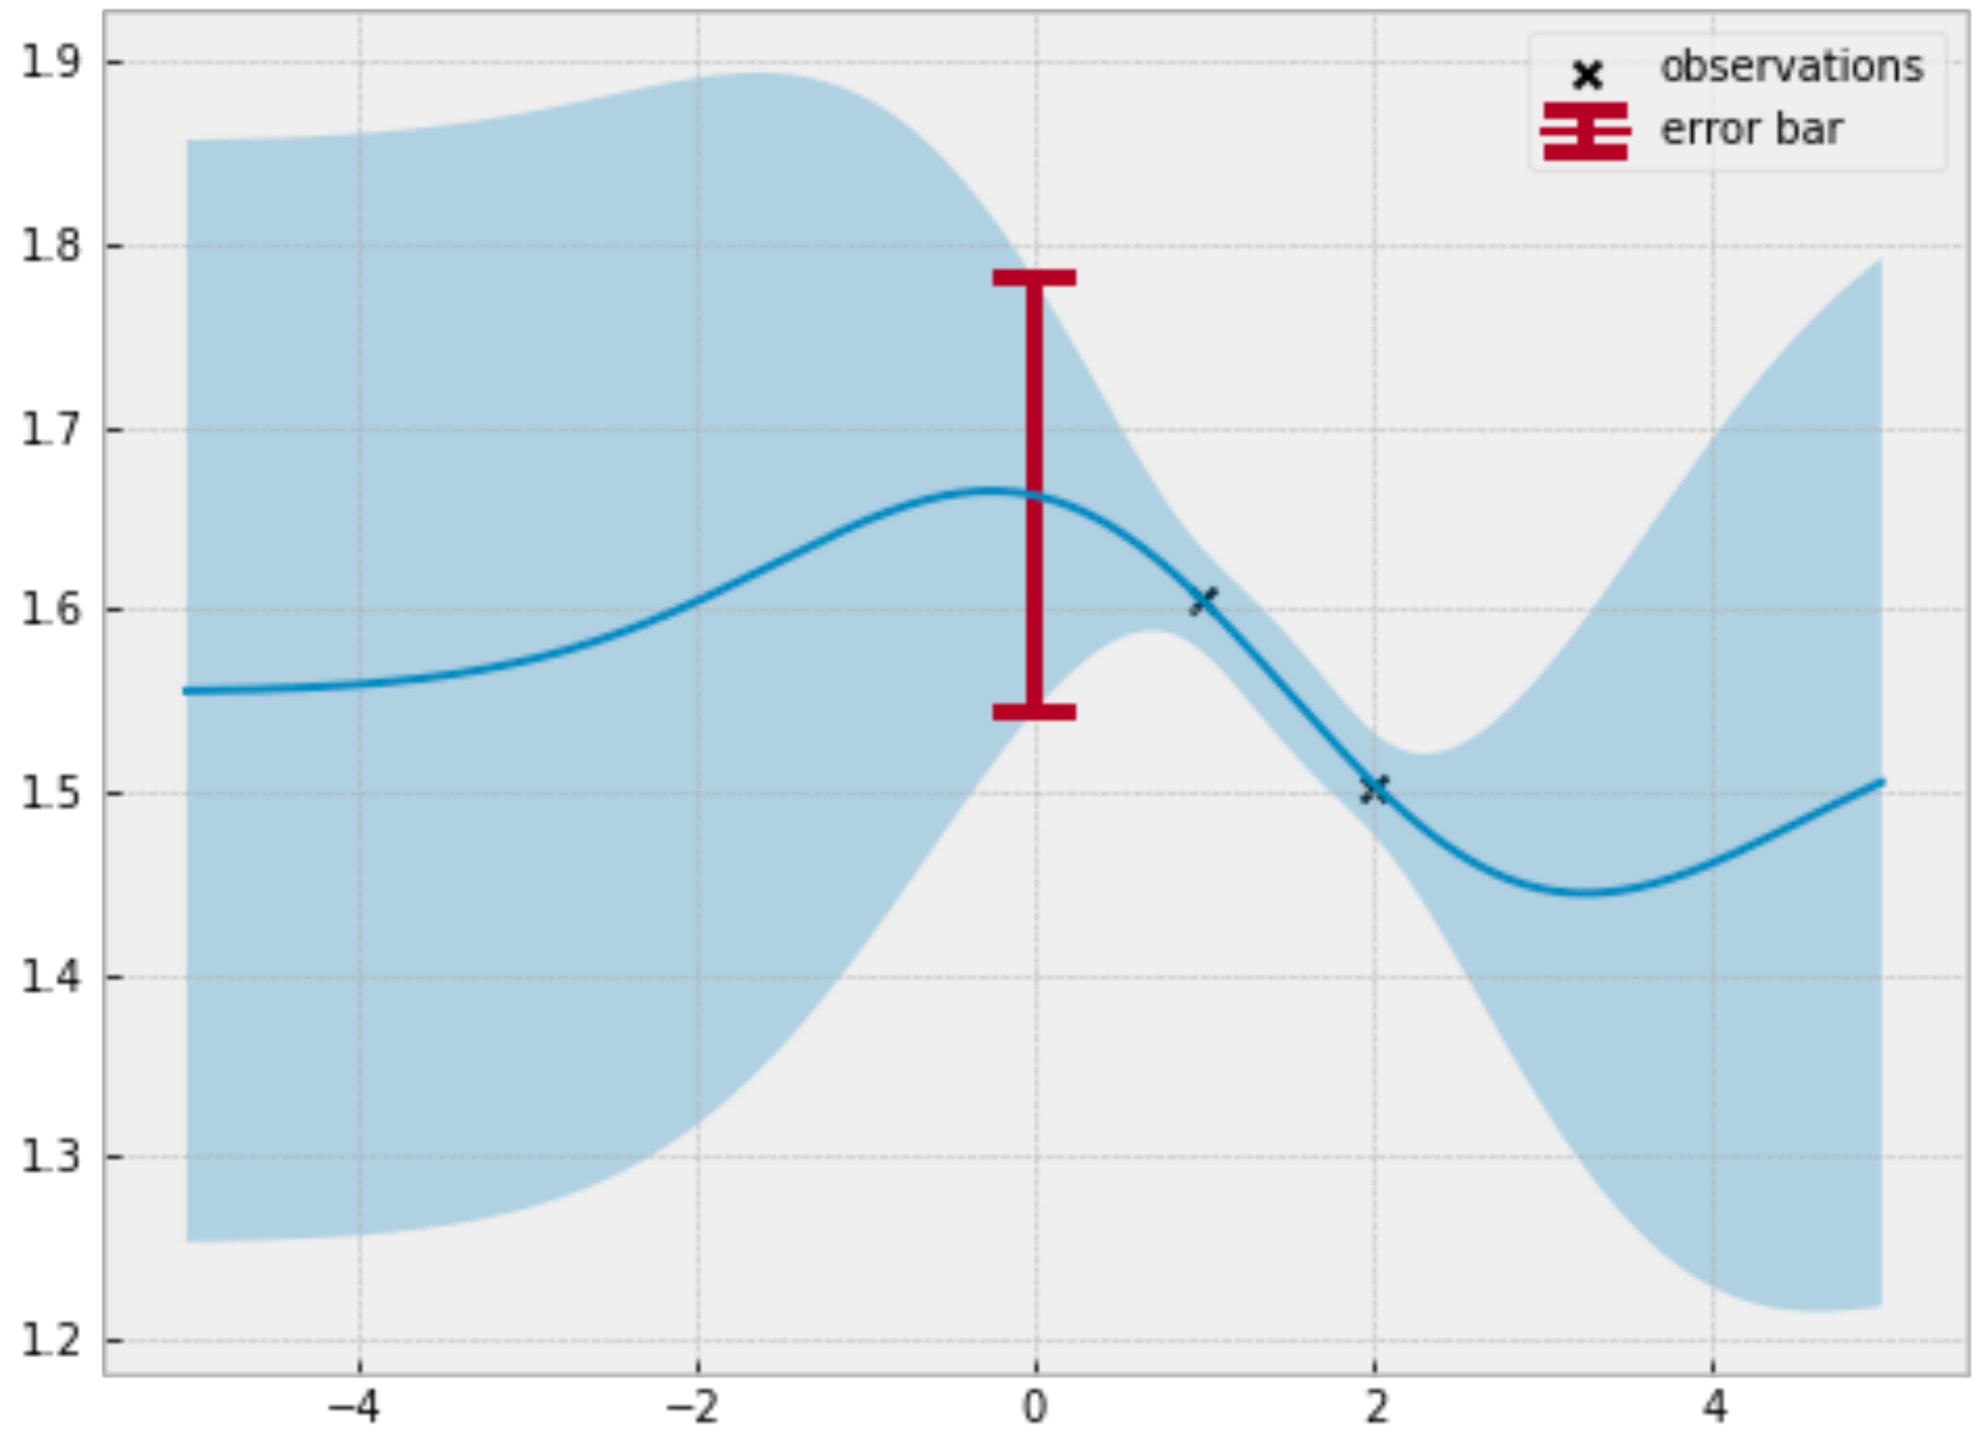

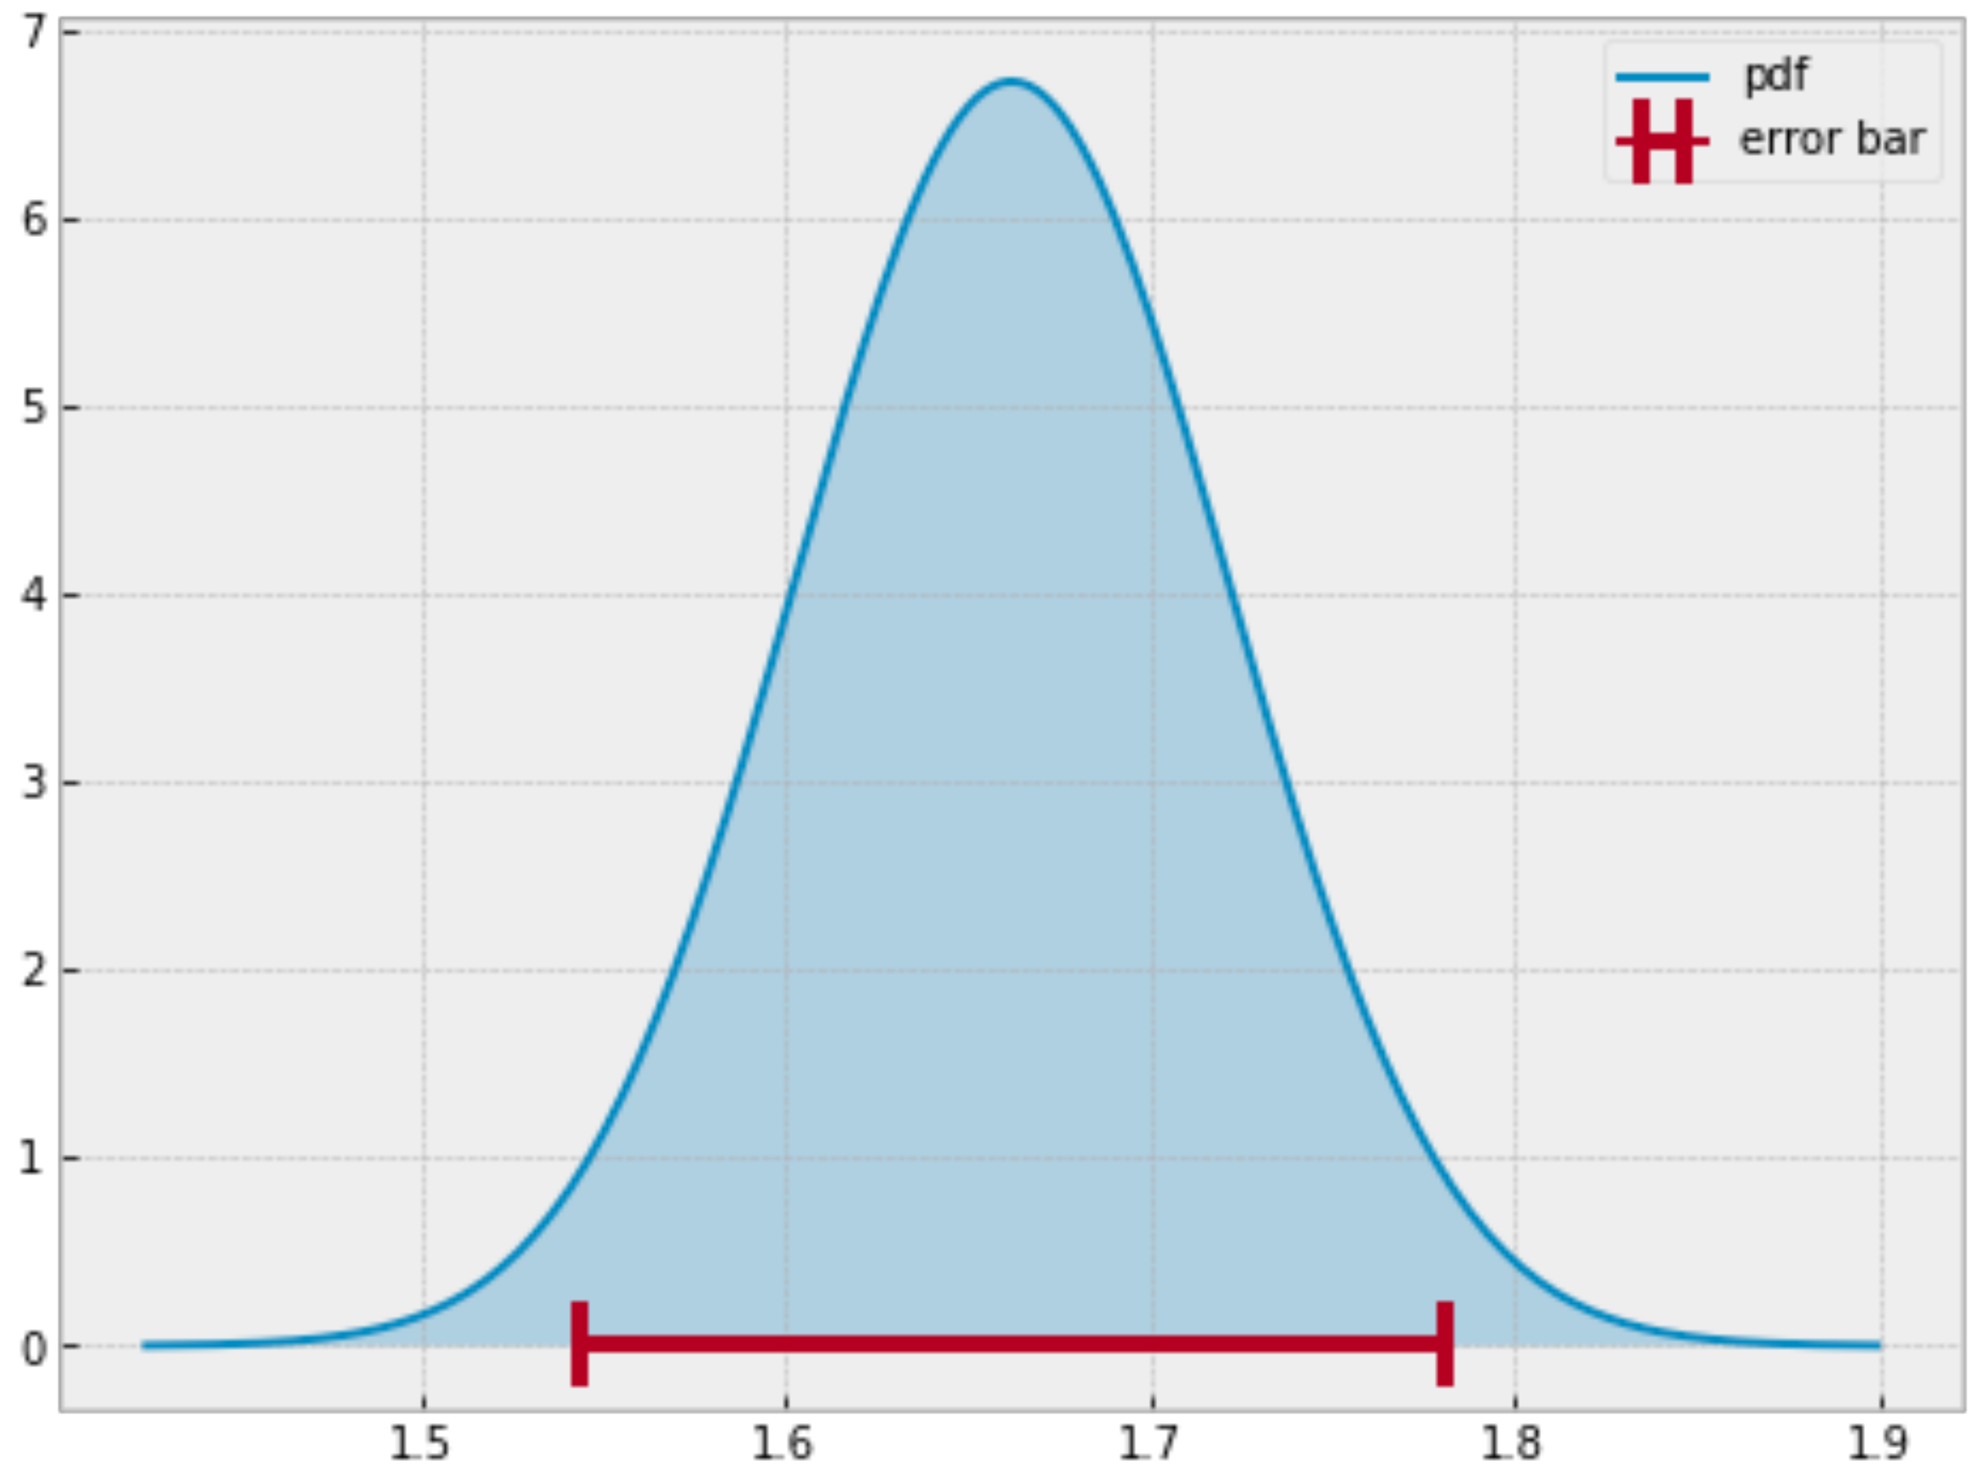

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

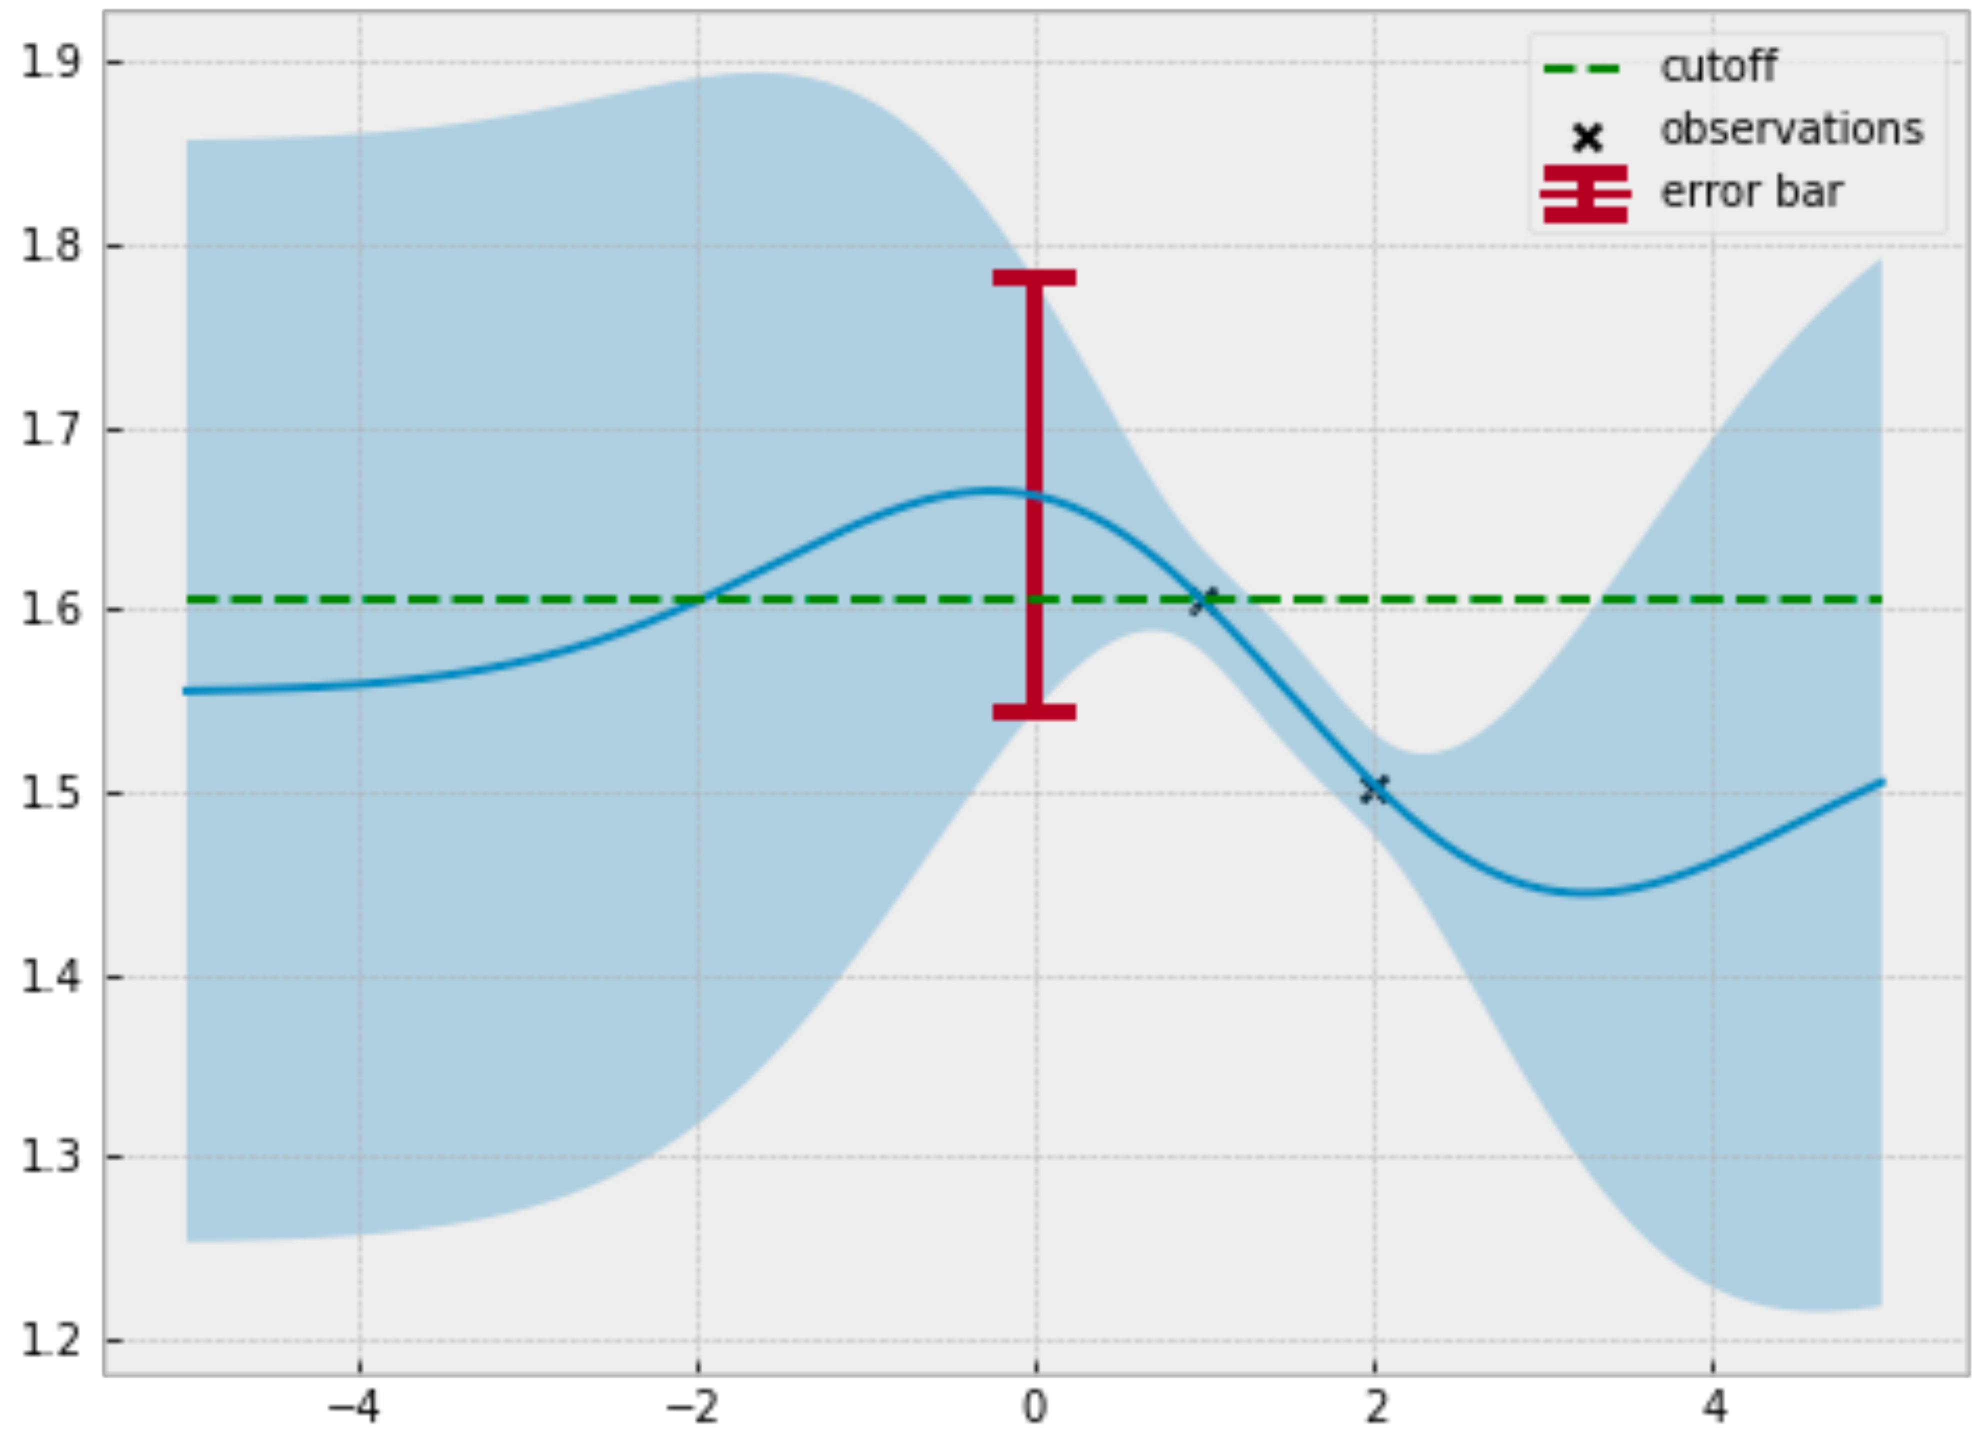

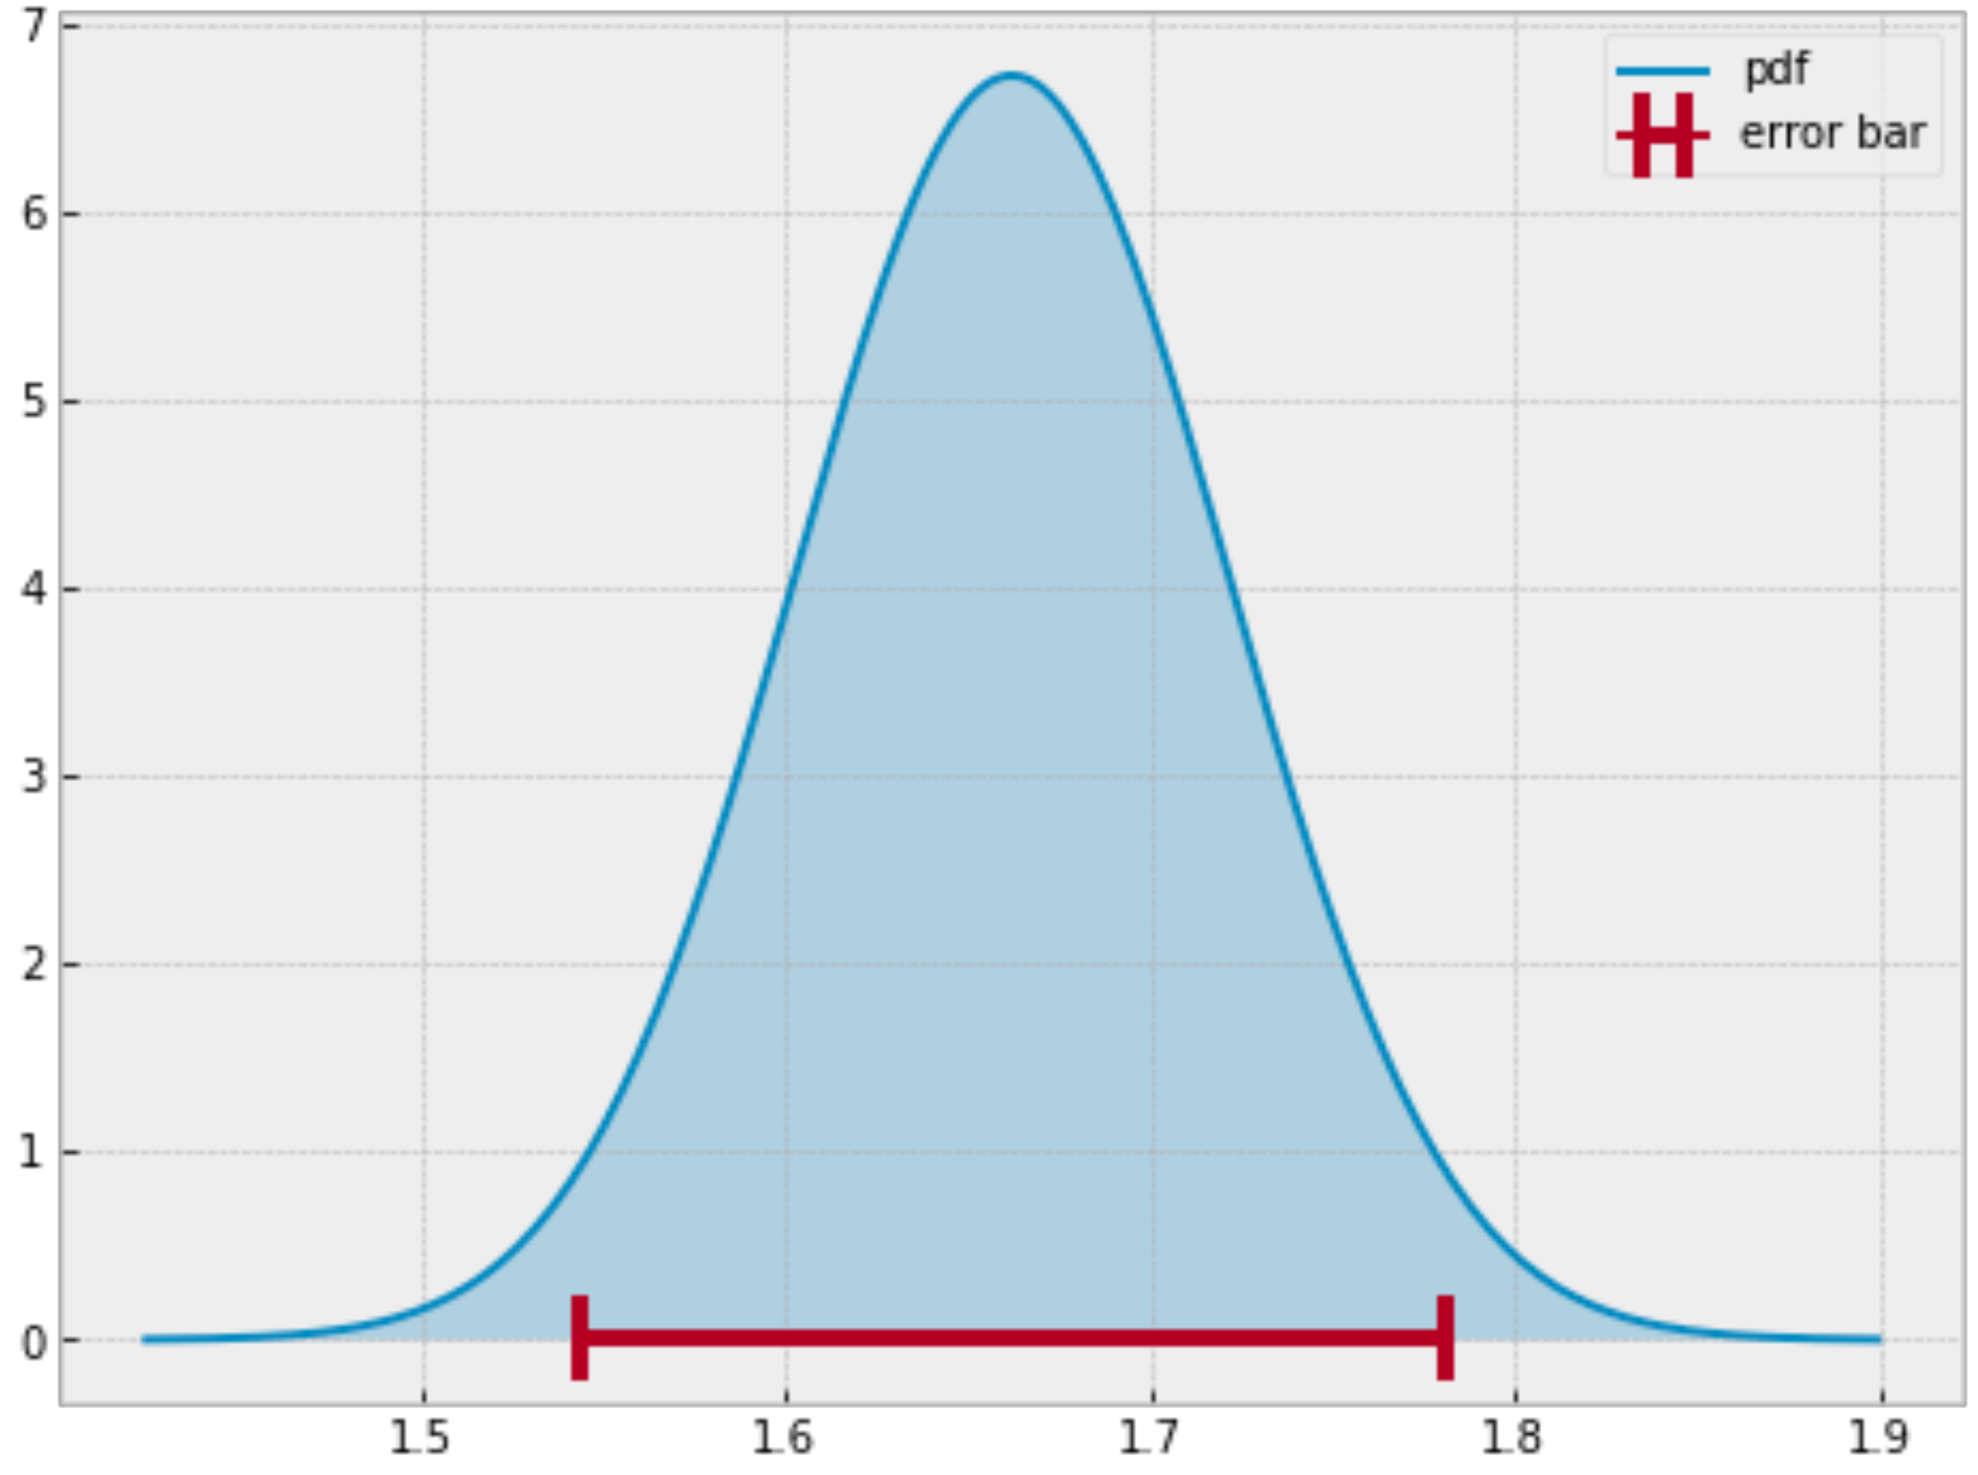

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

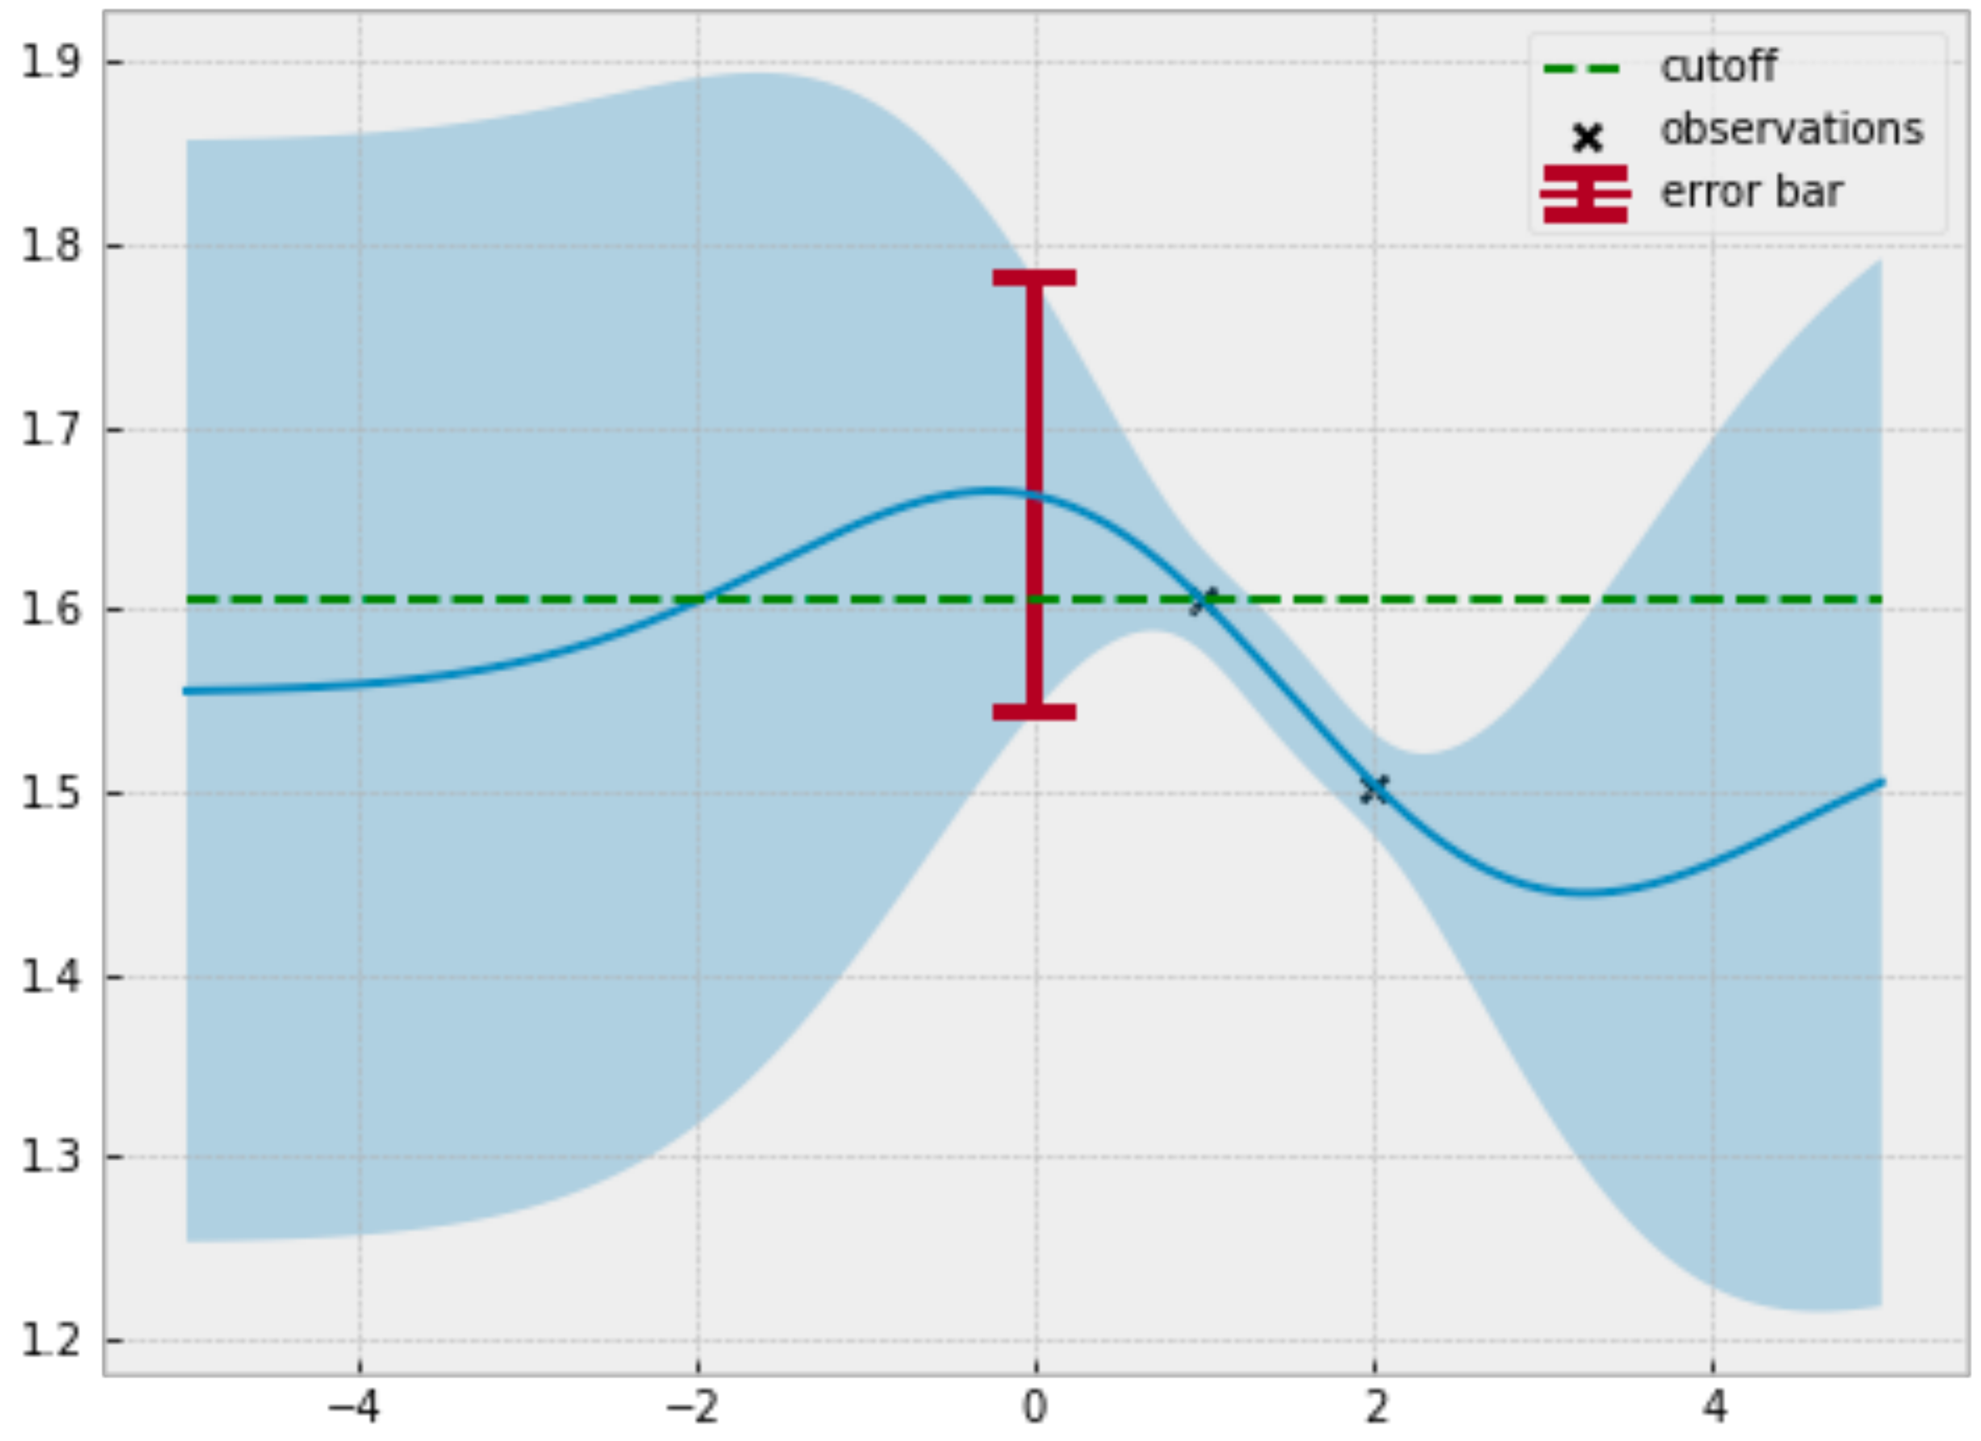

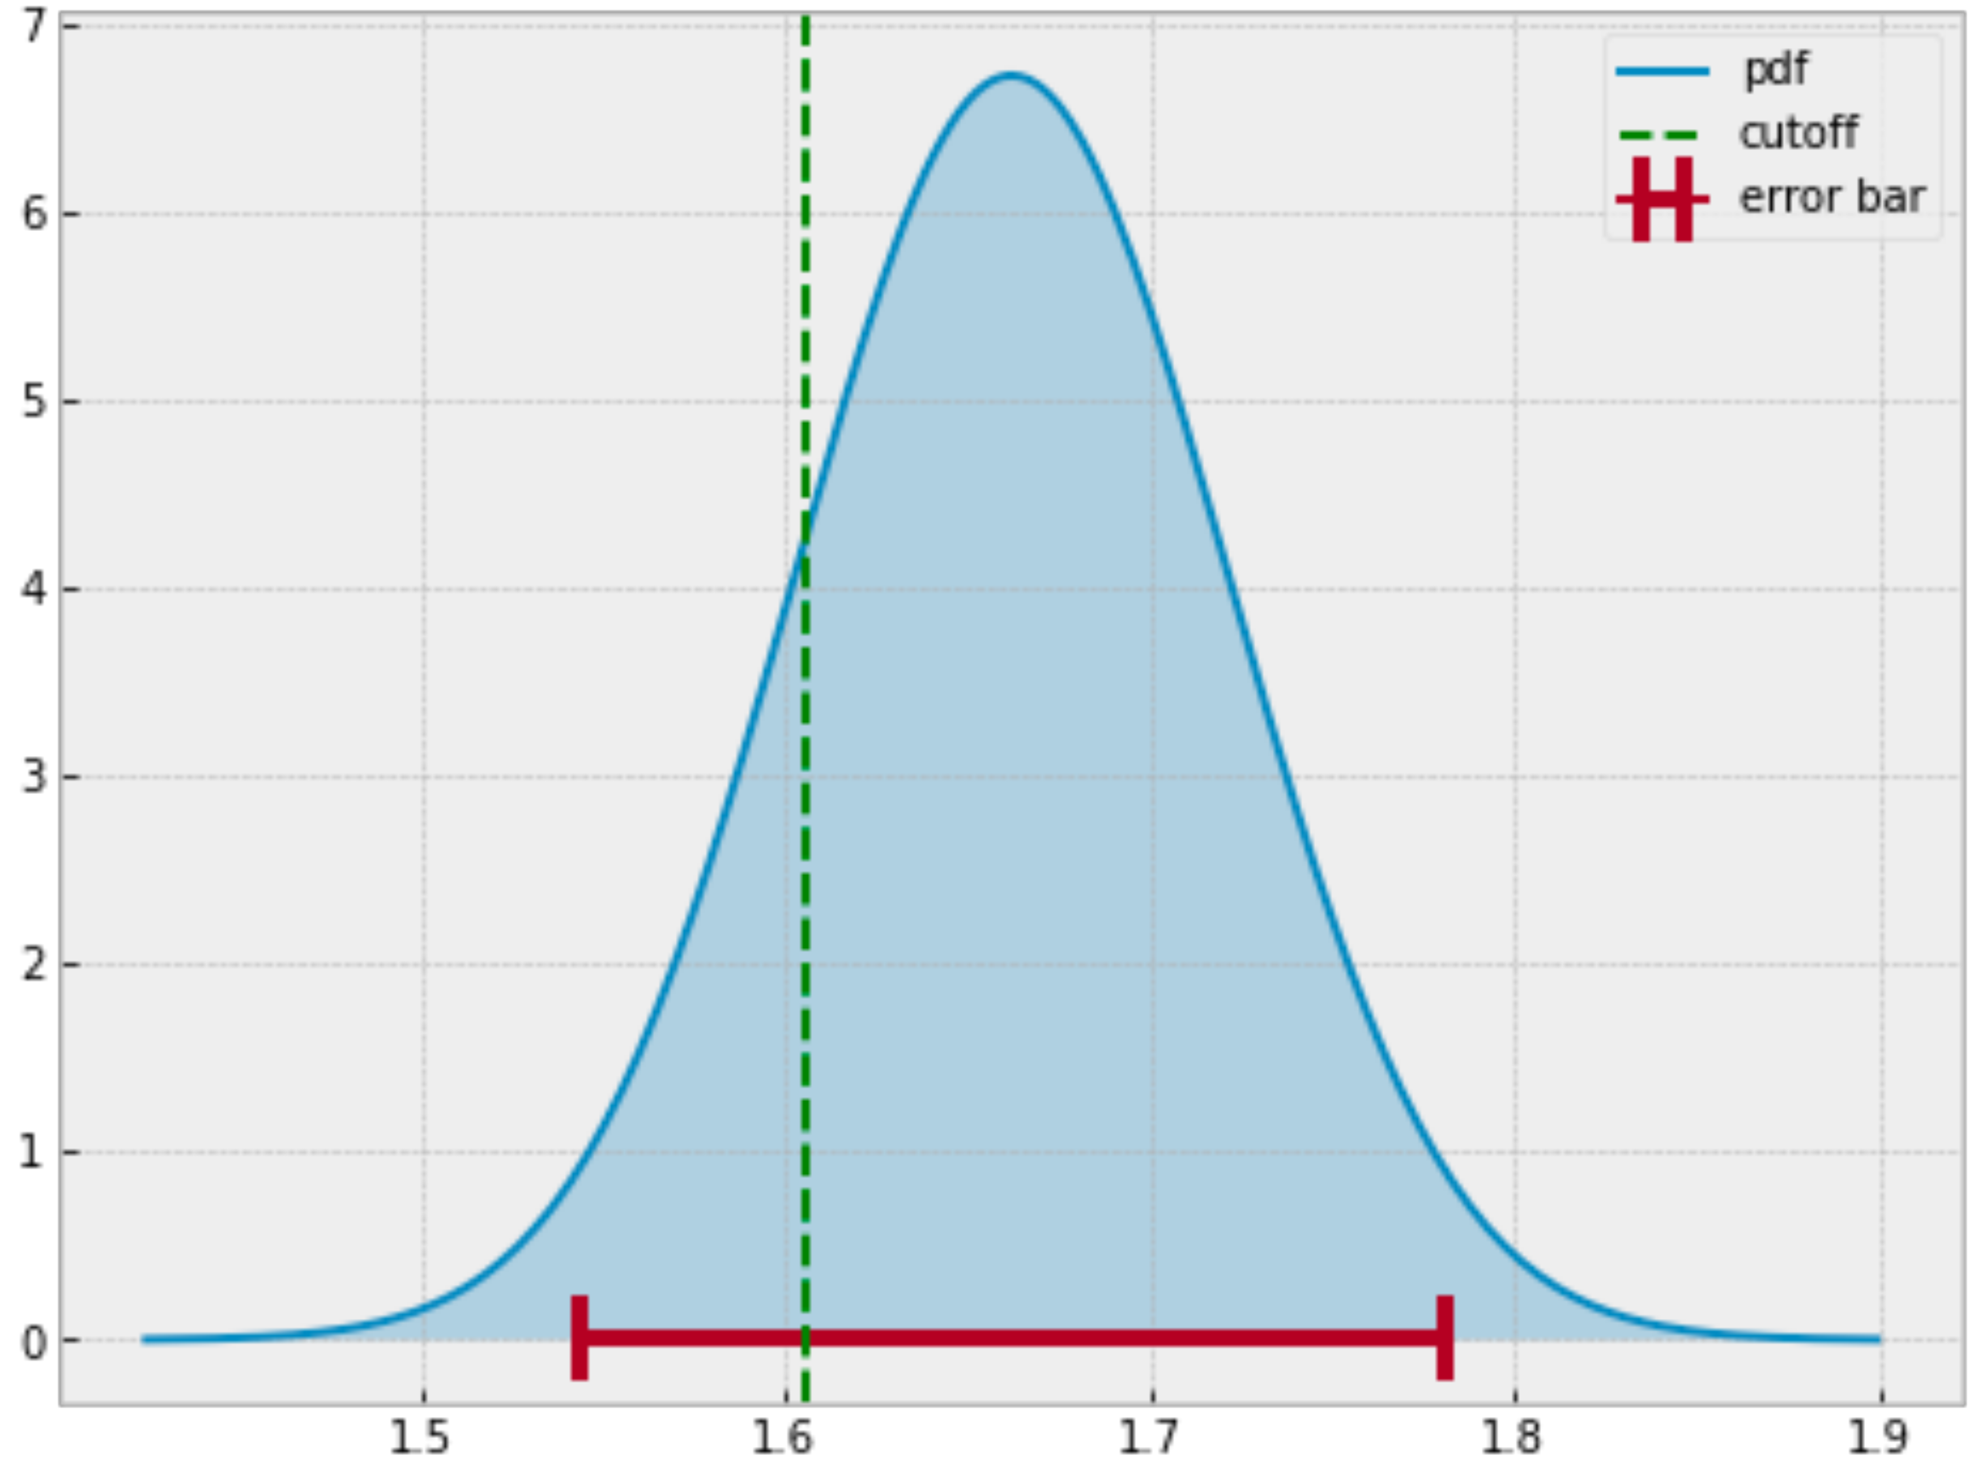

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

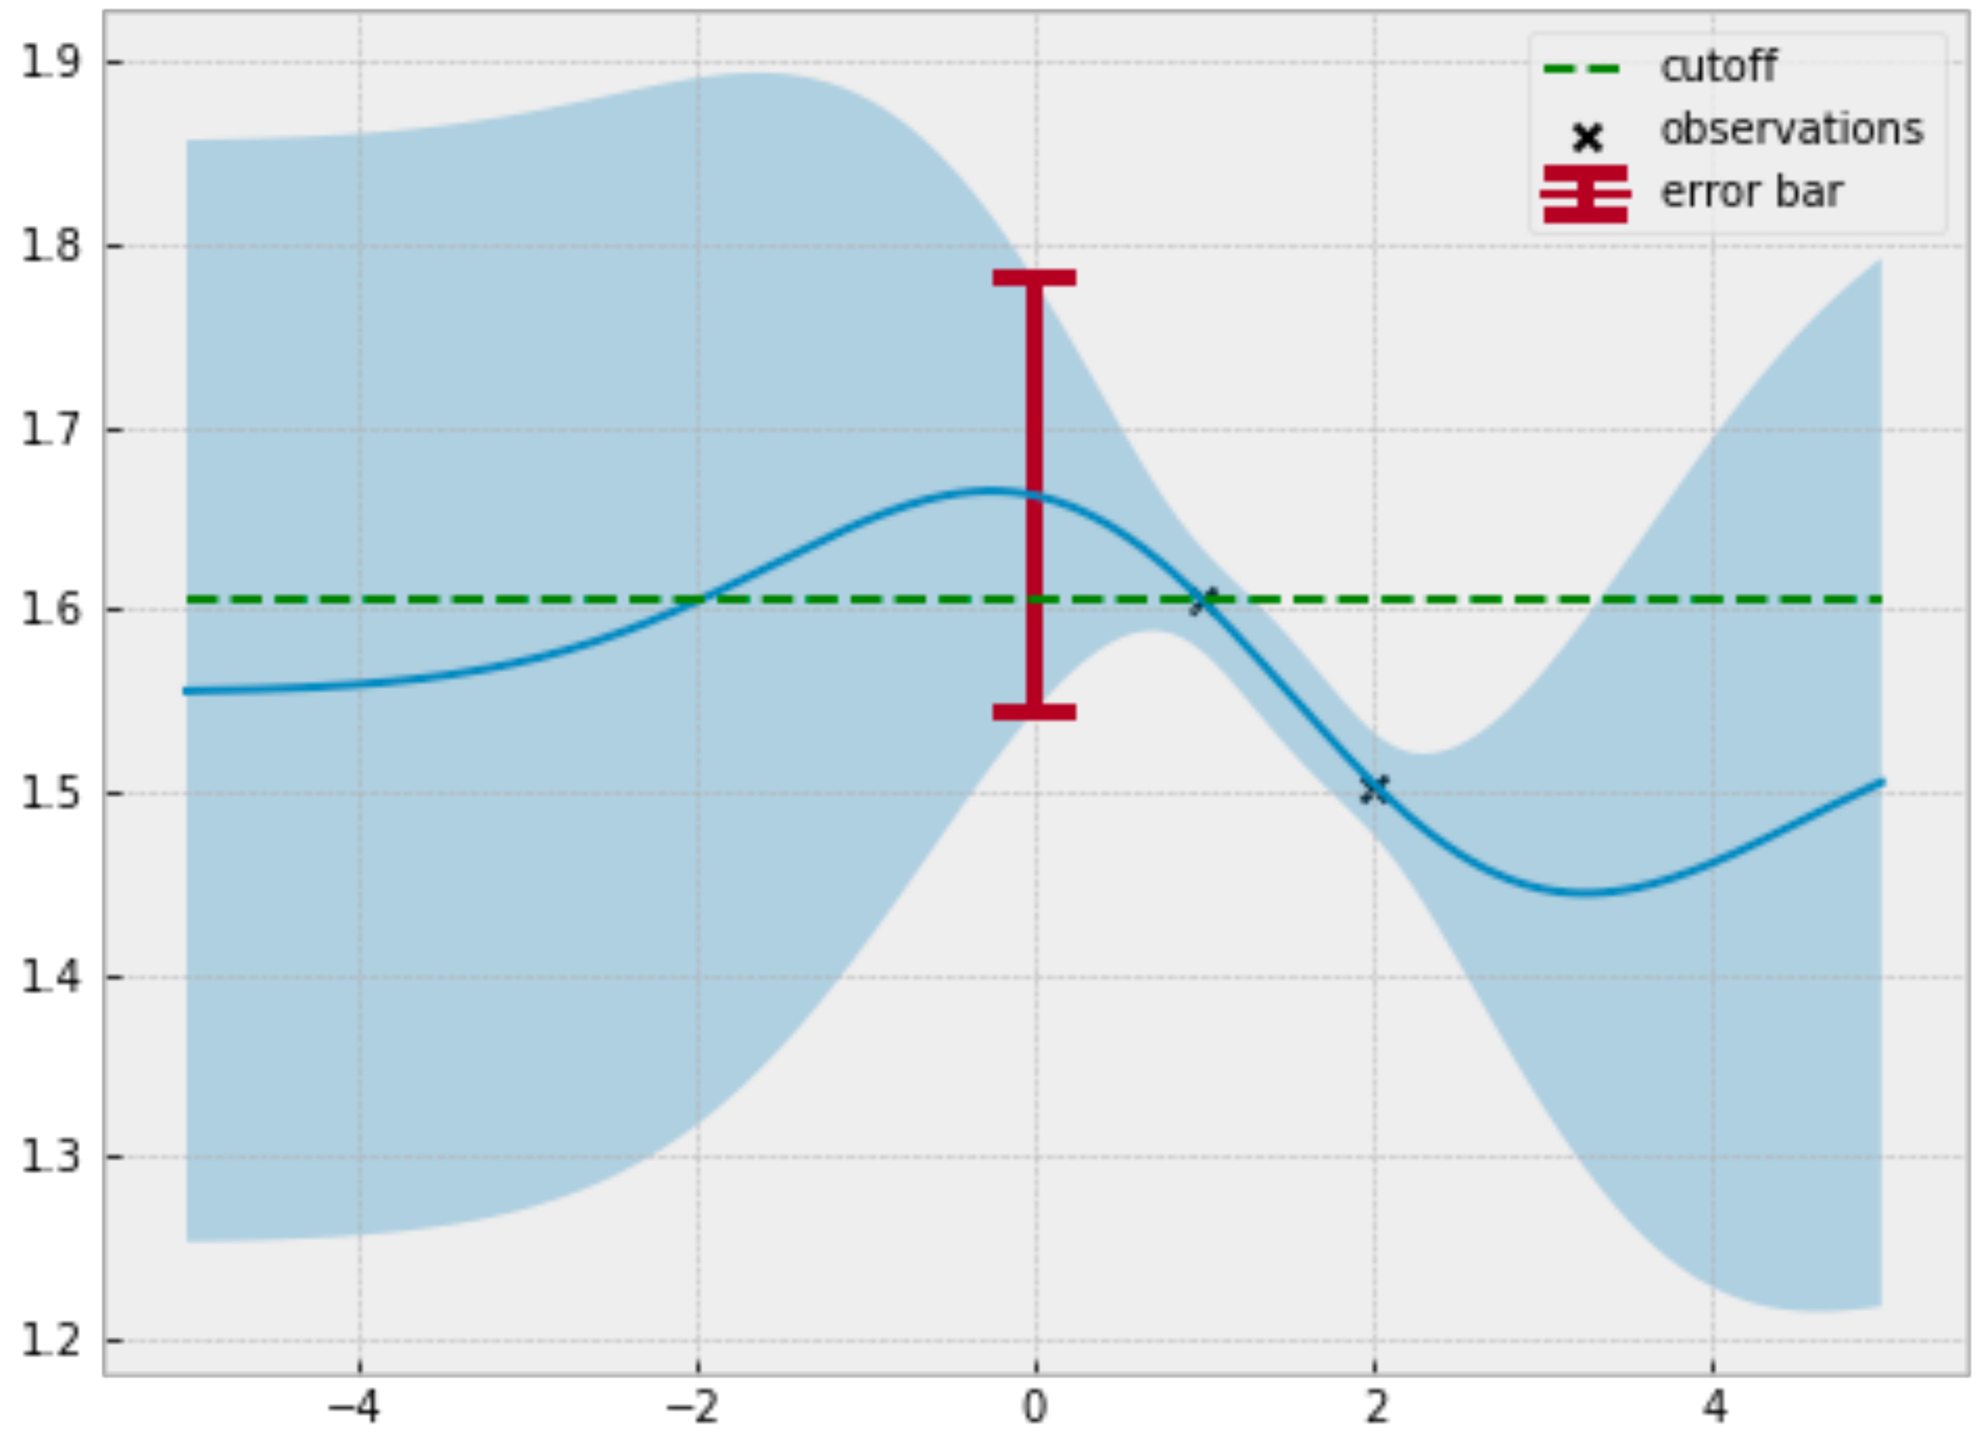

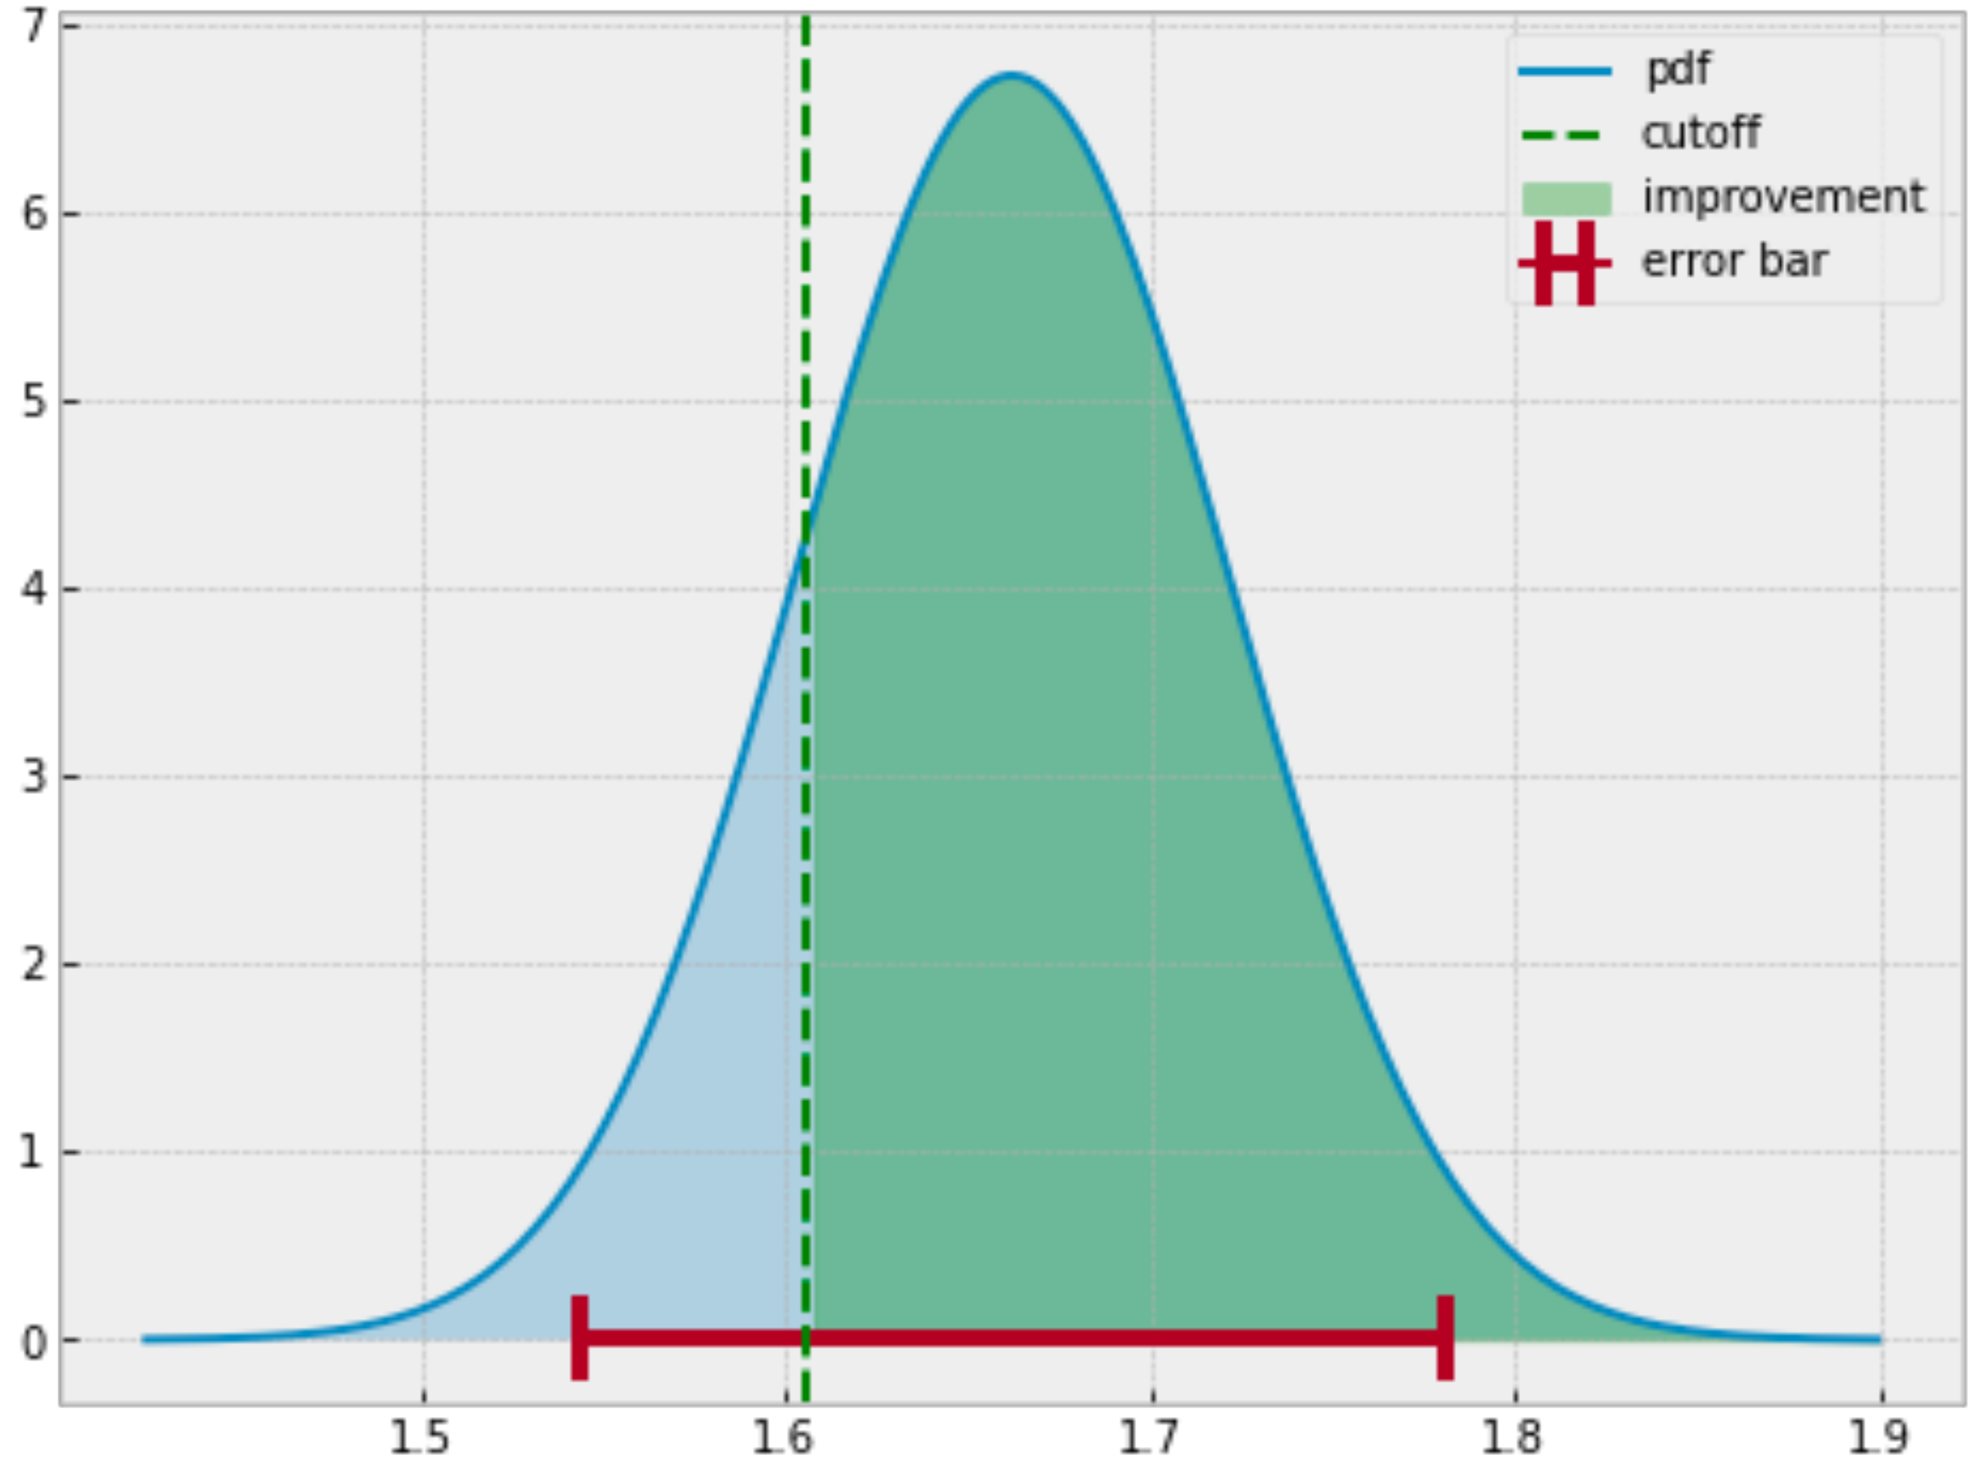

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

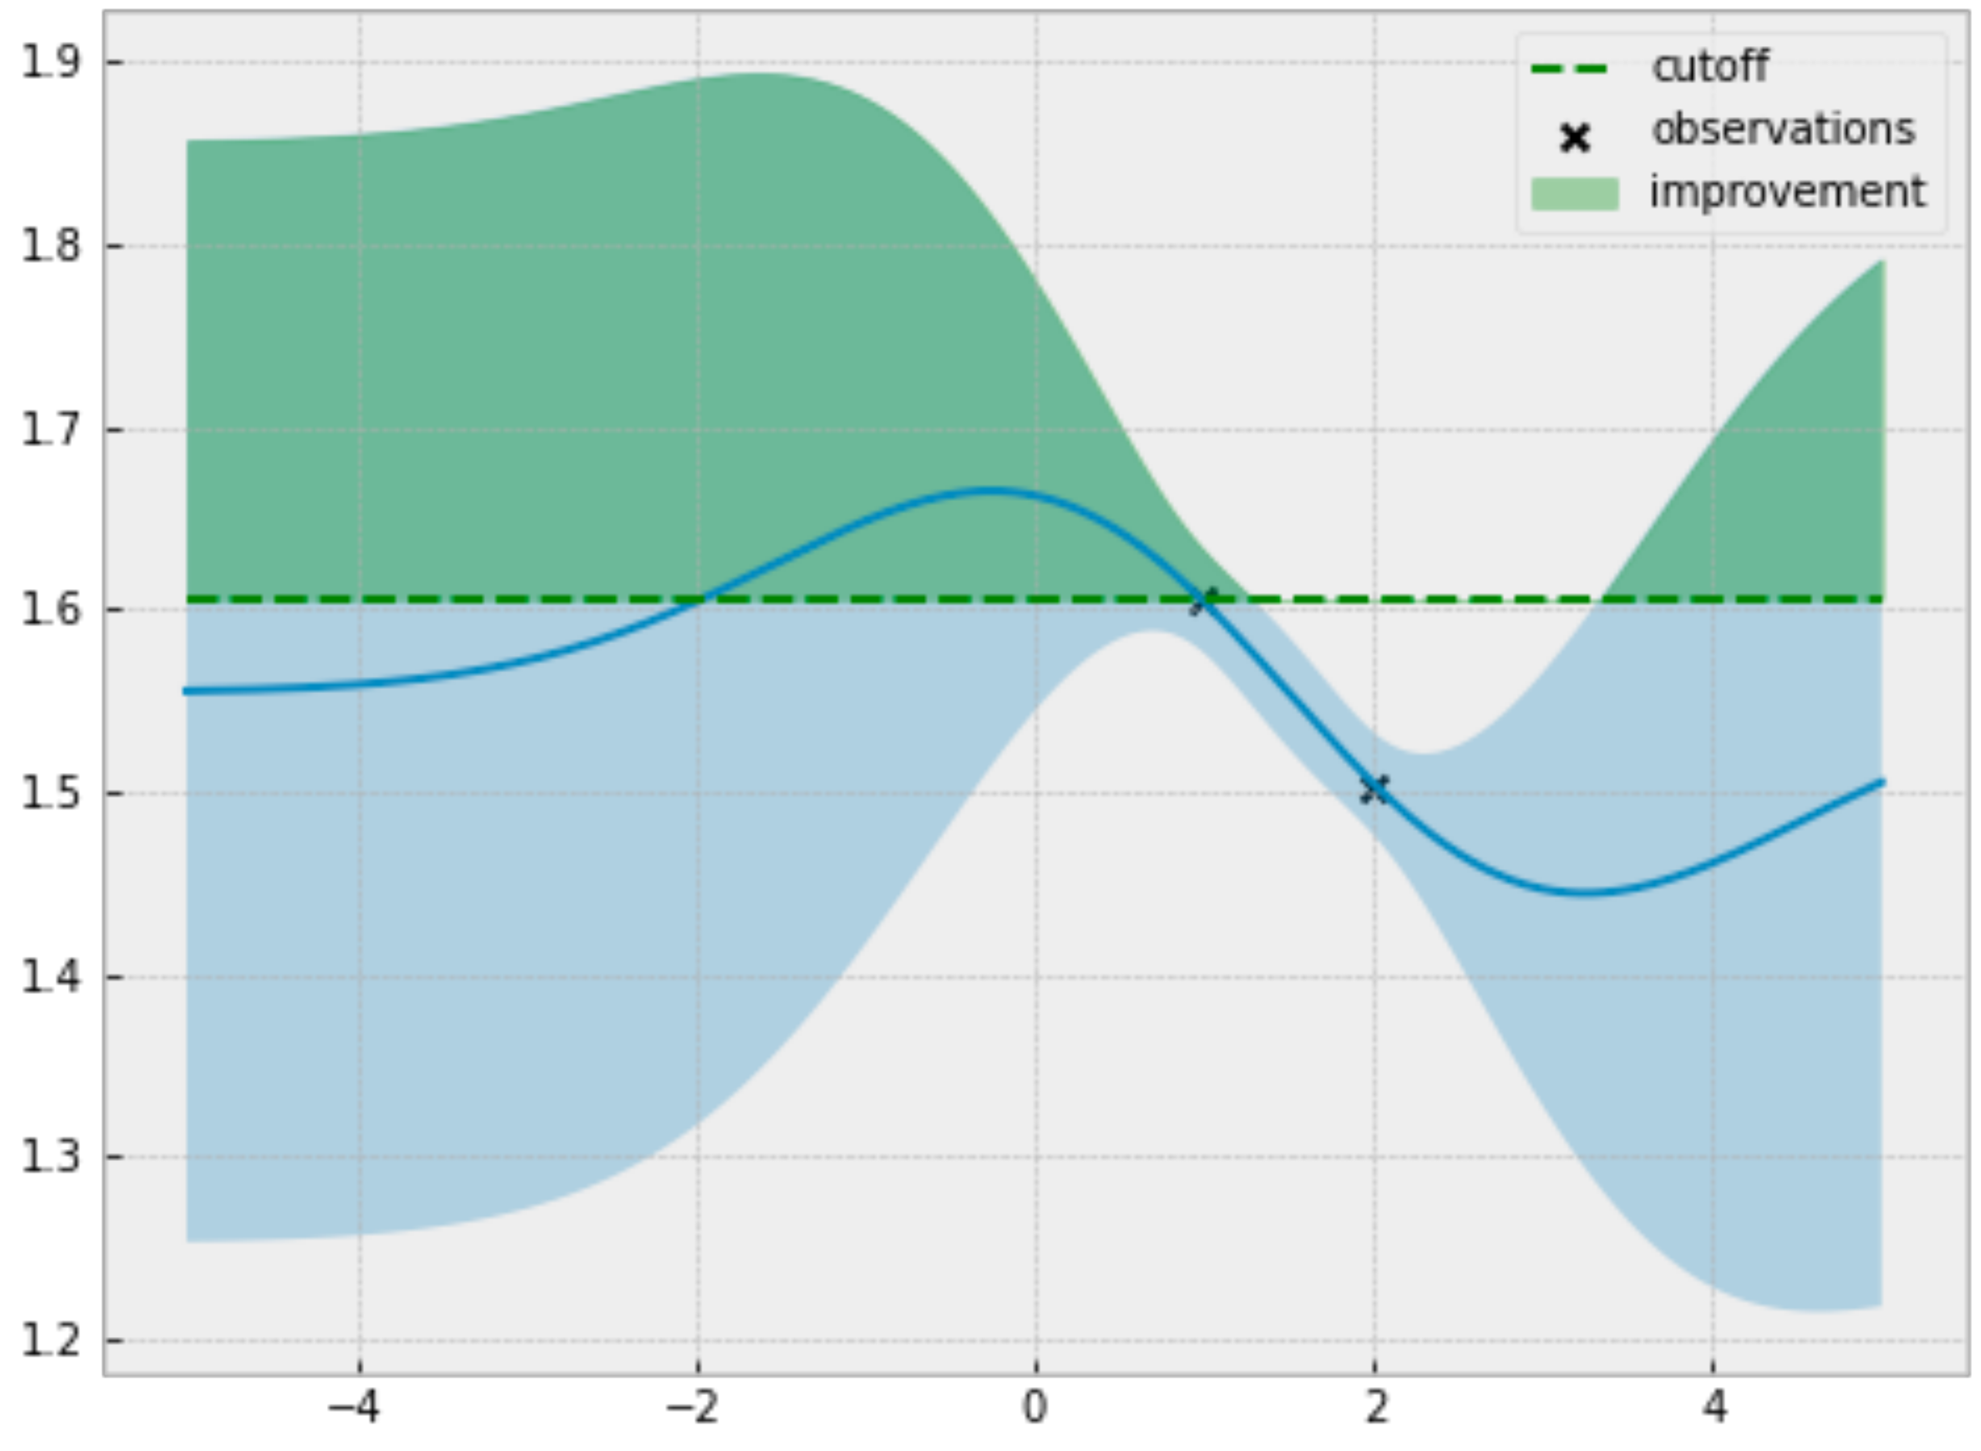

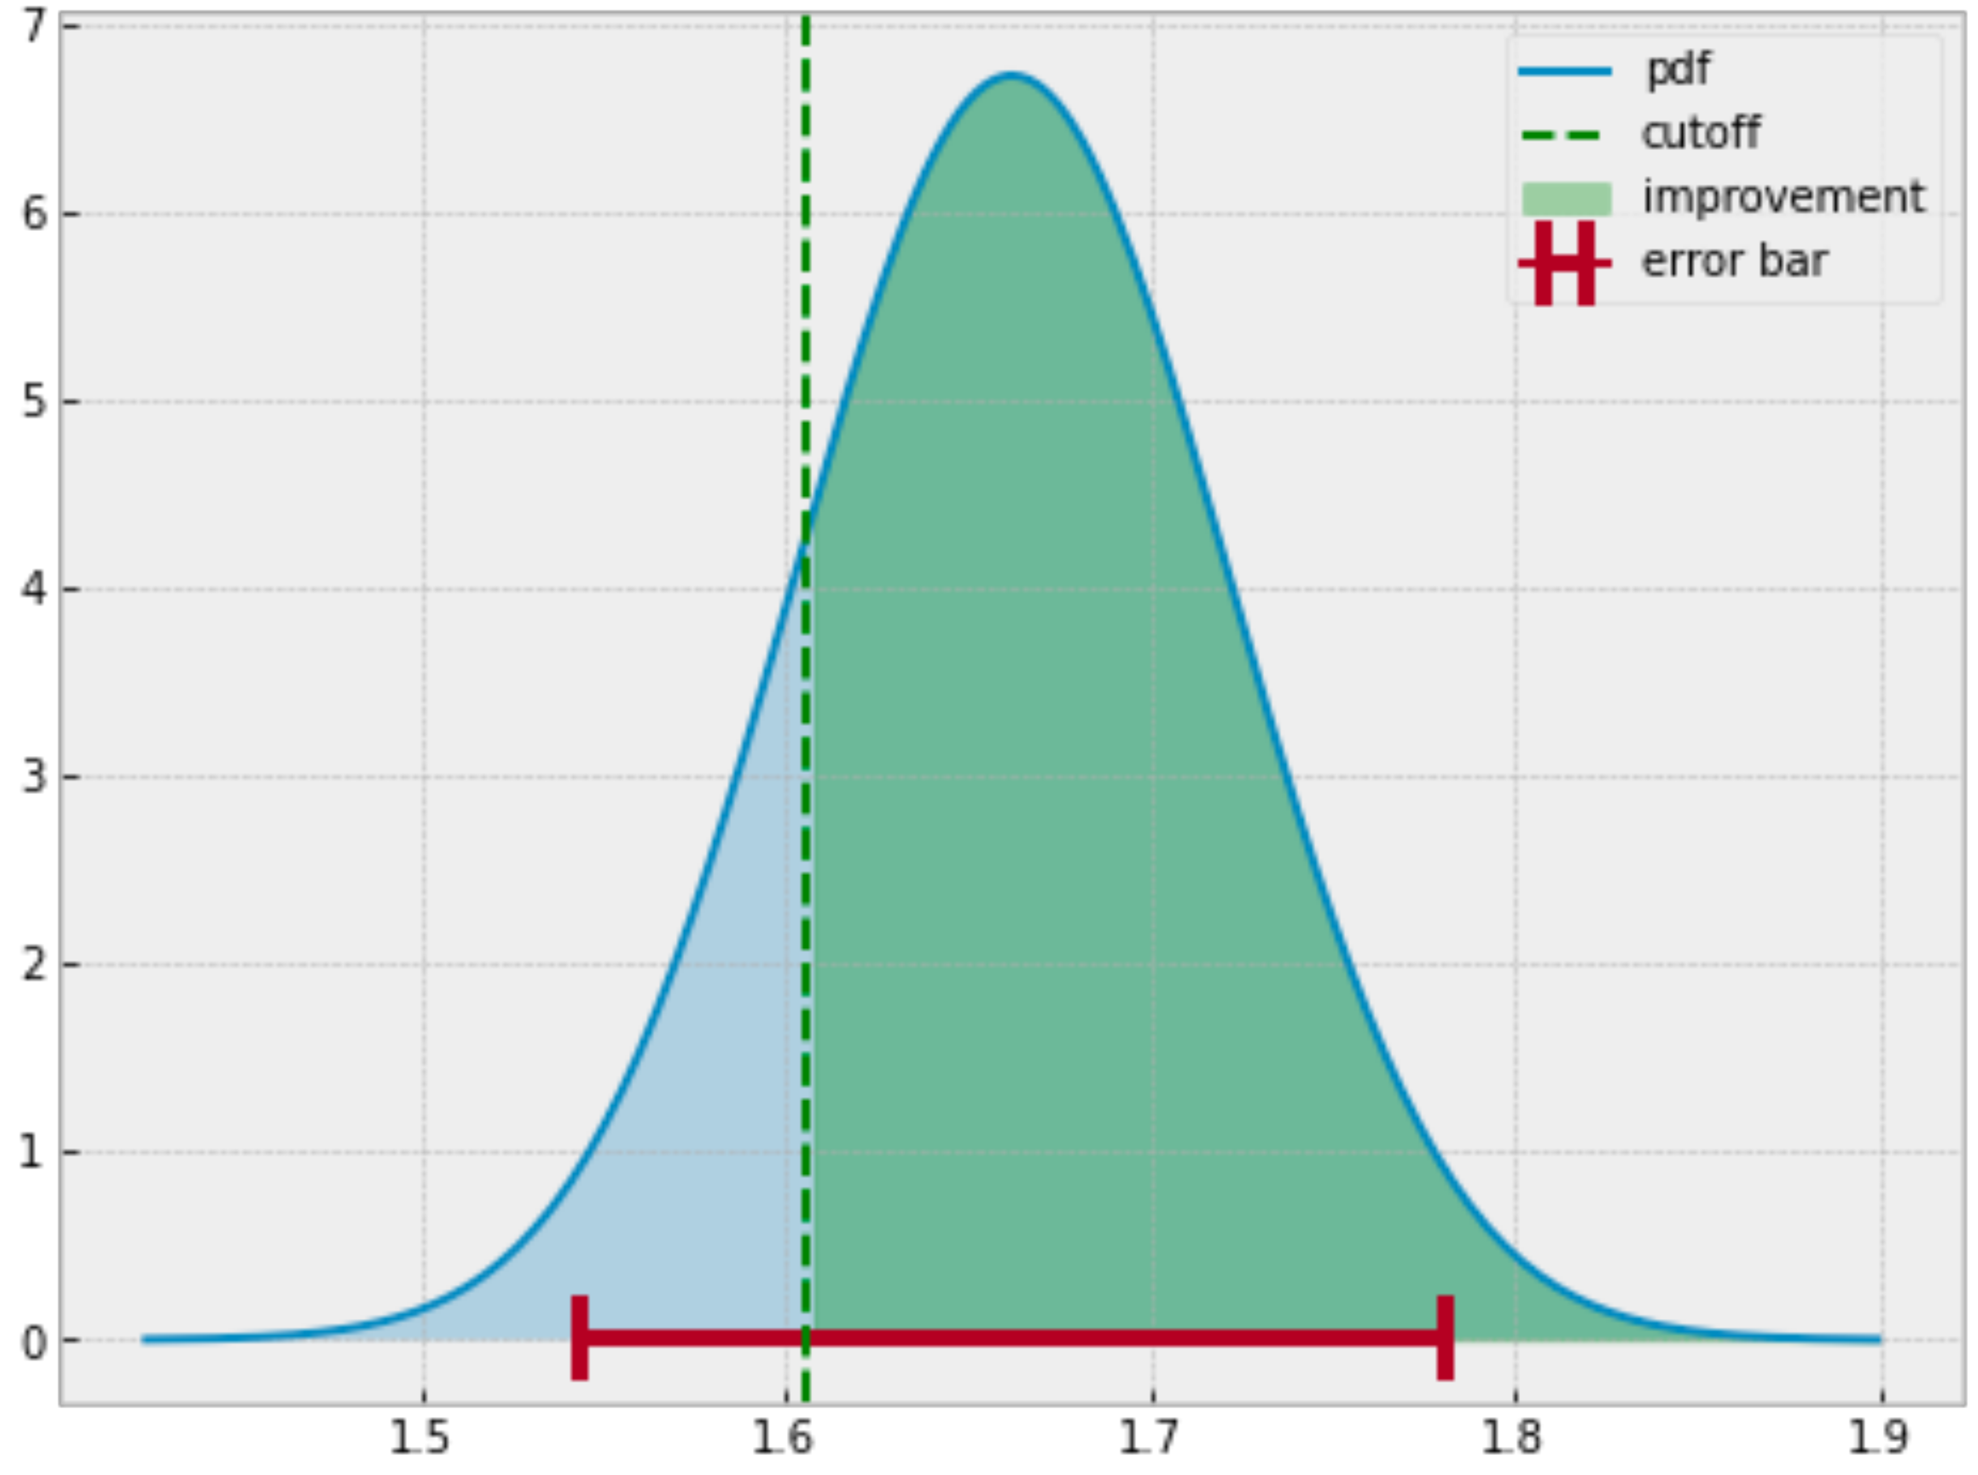

GAUSSIANITY IS AMENABLE TO IMPROVEMENT-RELATED CALCULATIONS

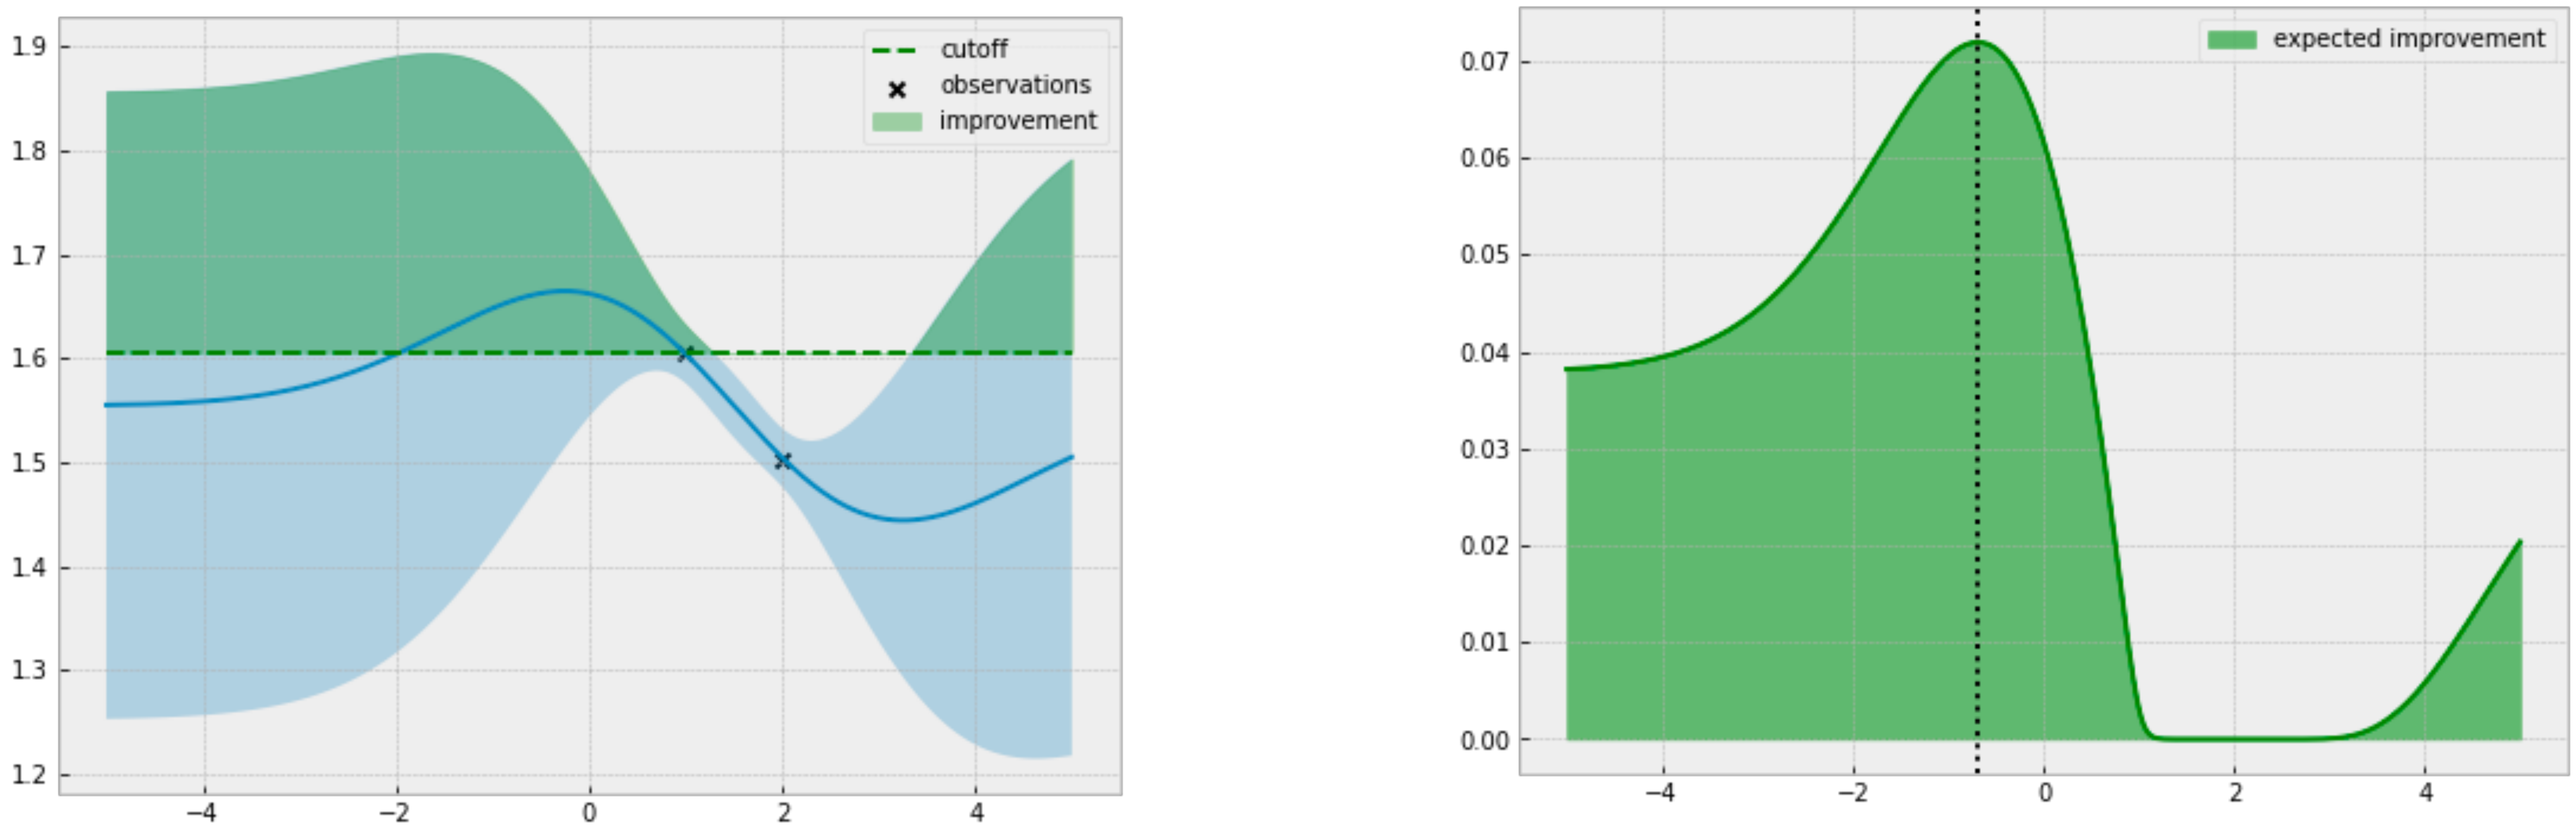

THE ACTIVE LEARNING LOOP WITH EXPECTED IMPROVEMENT

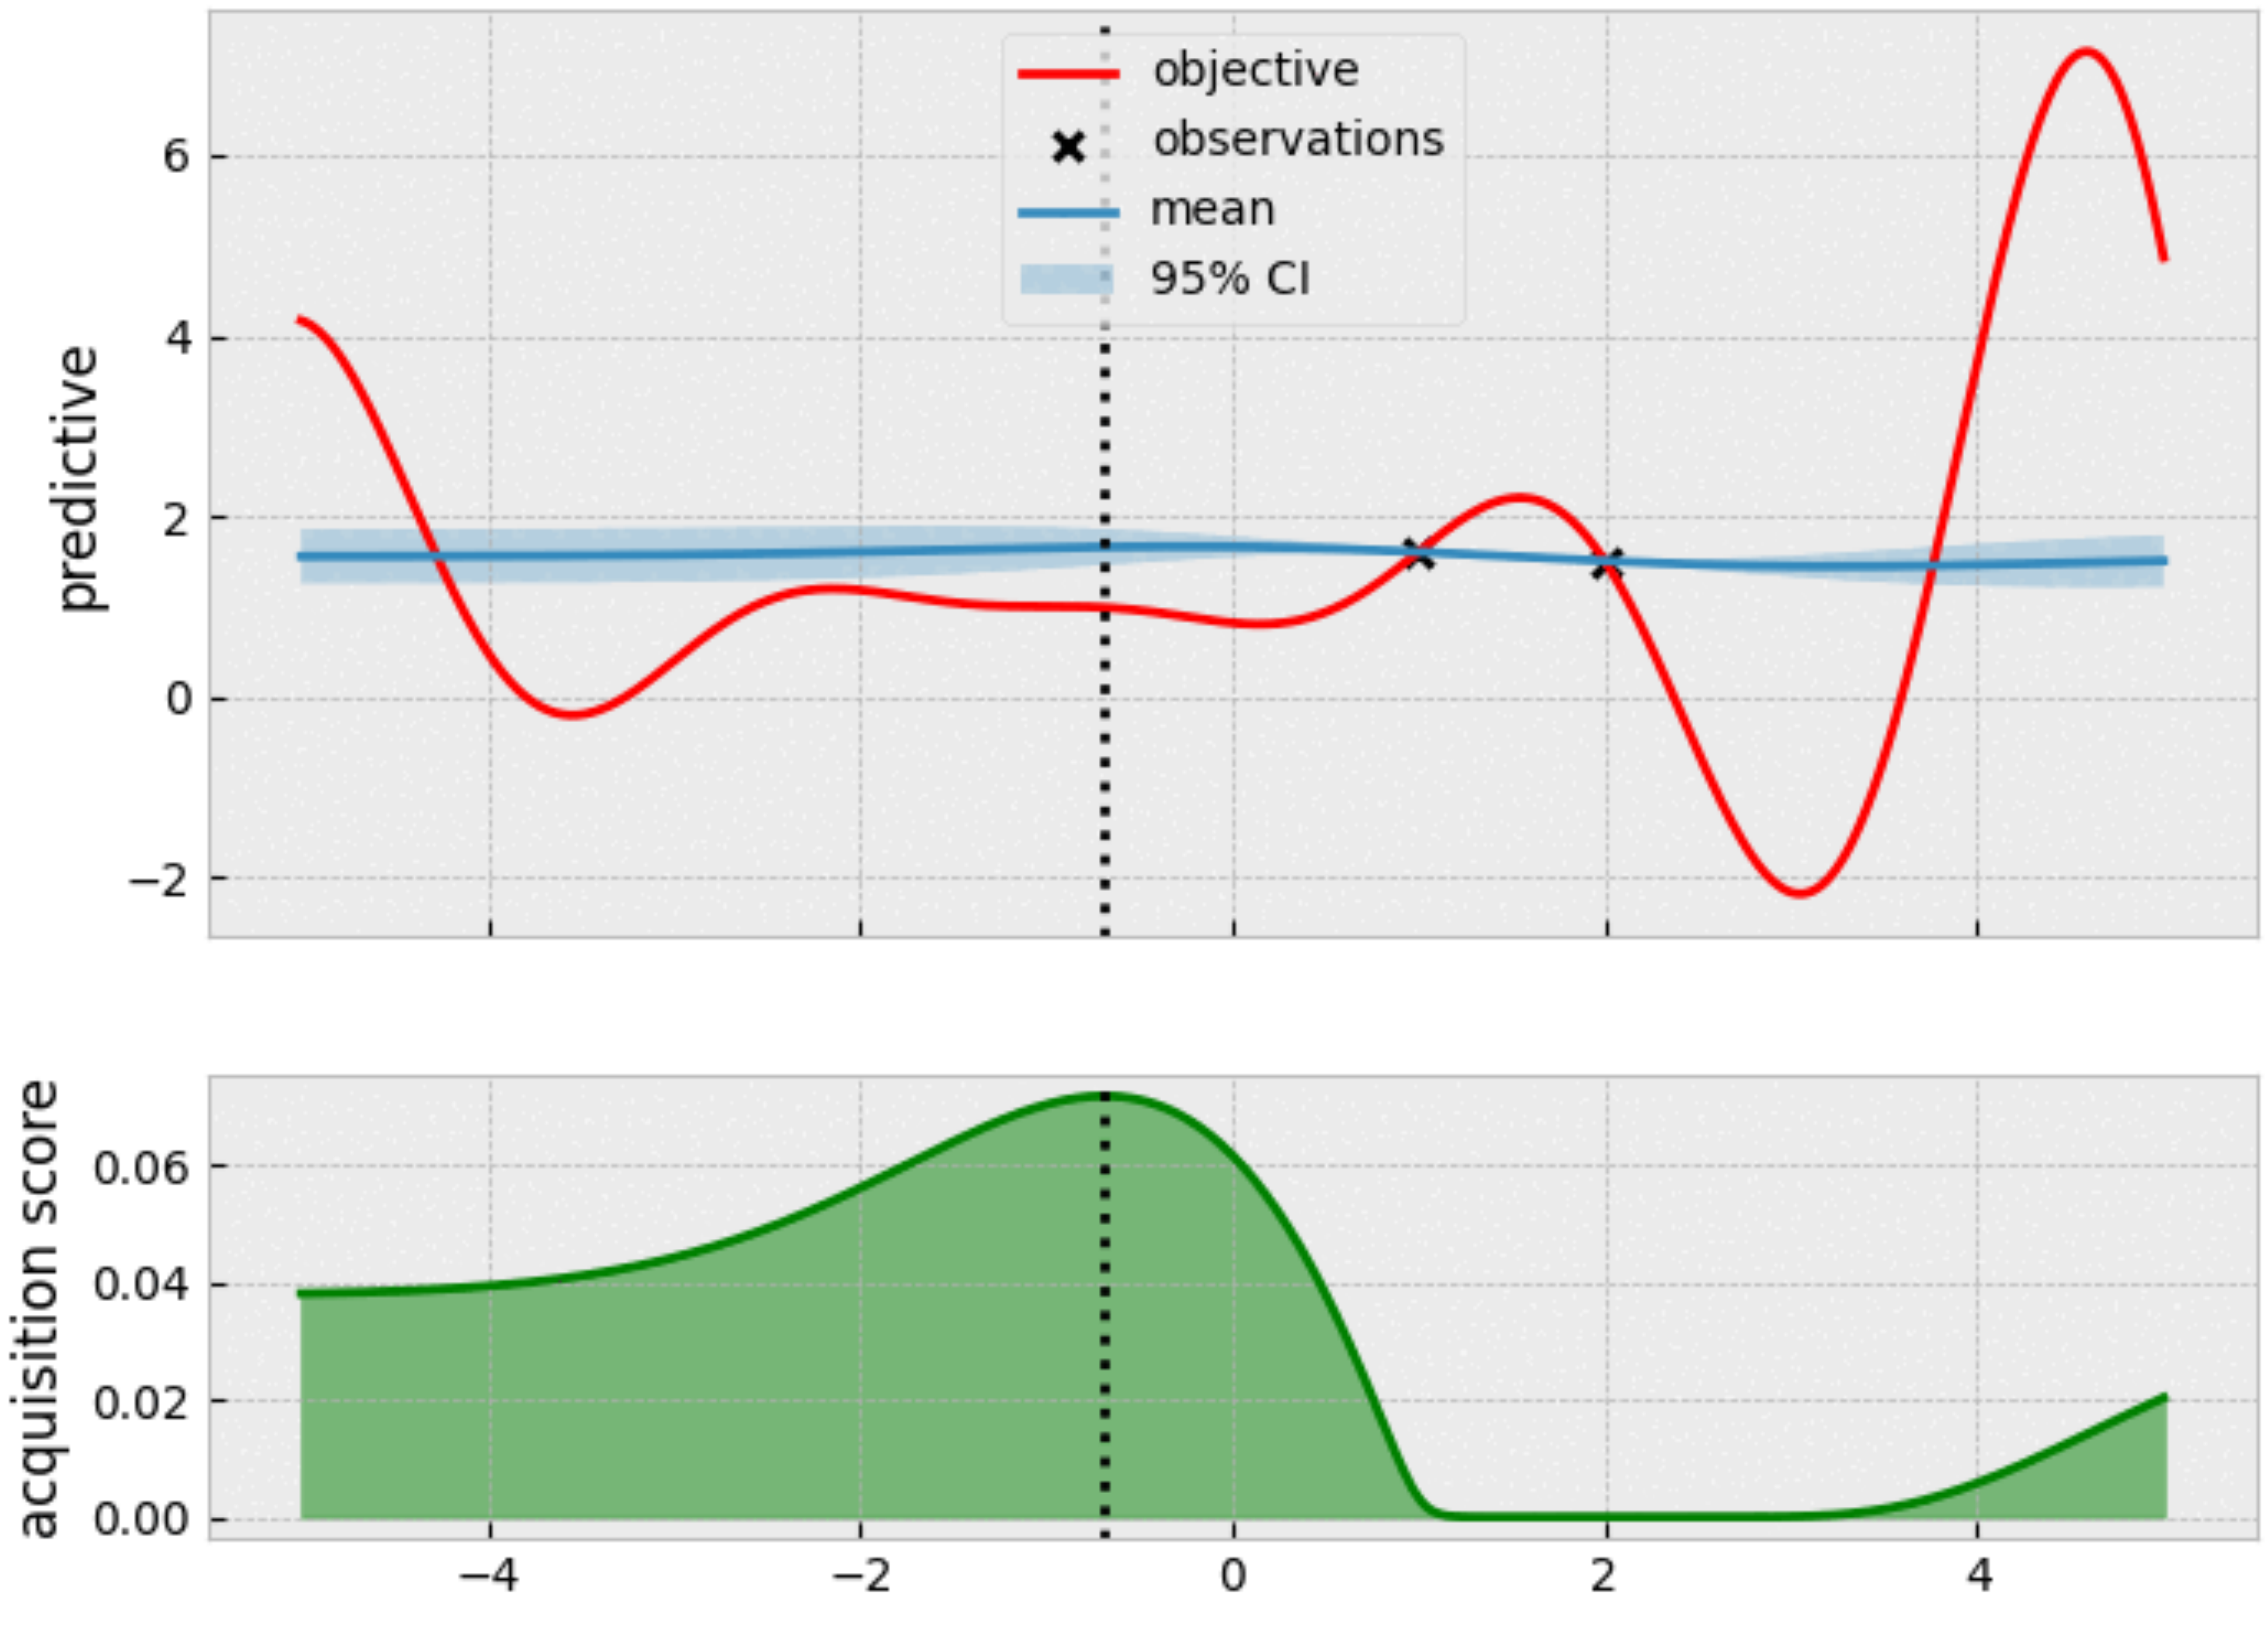

THE ACTIVE LEARNING LOOP WITH EXPECTED IMPROVEMENT

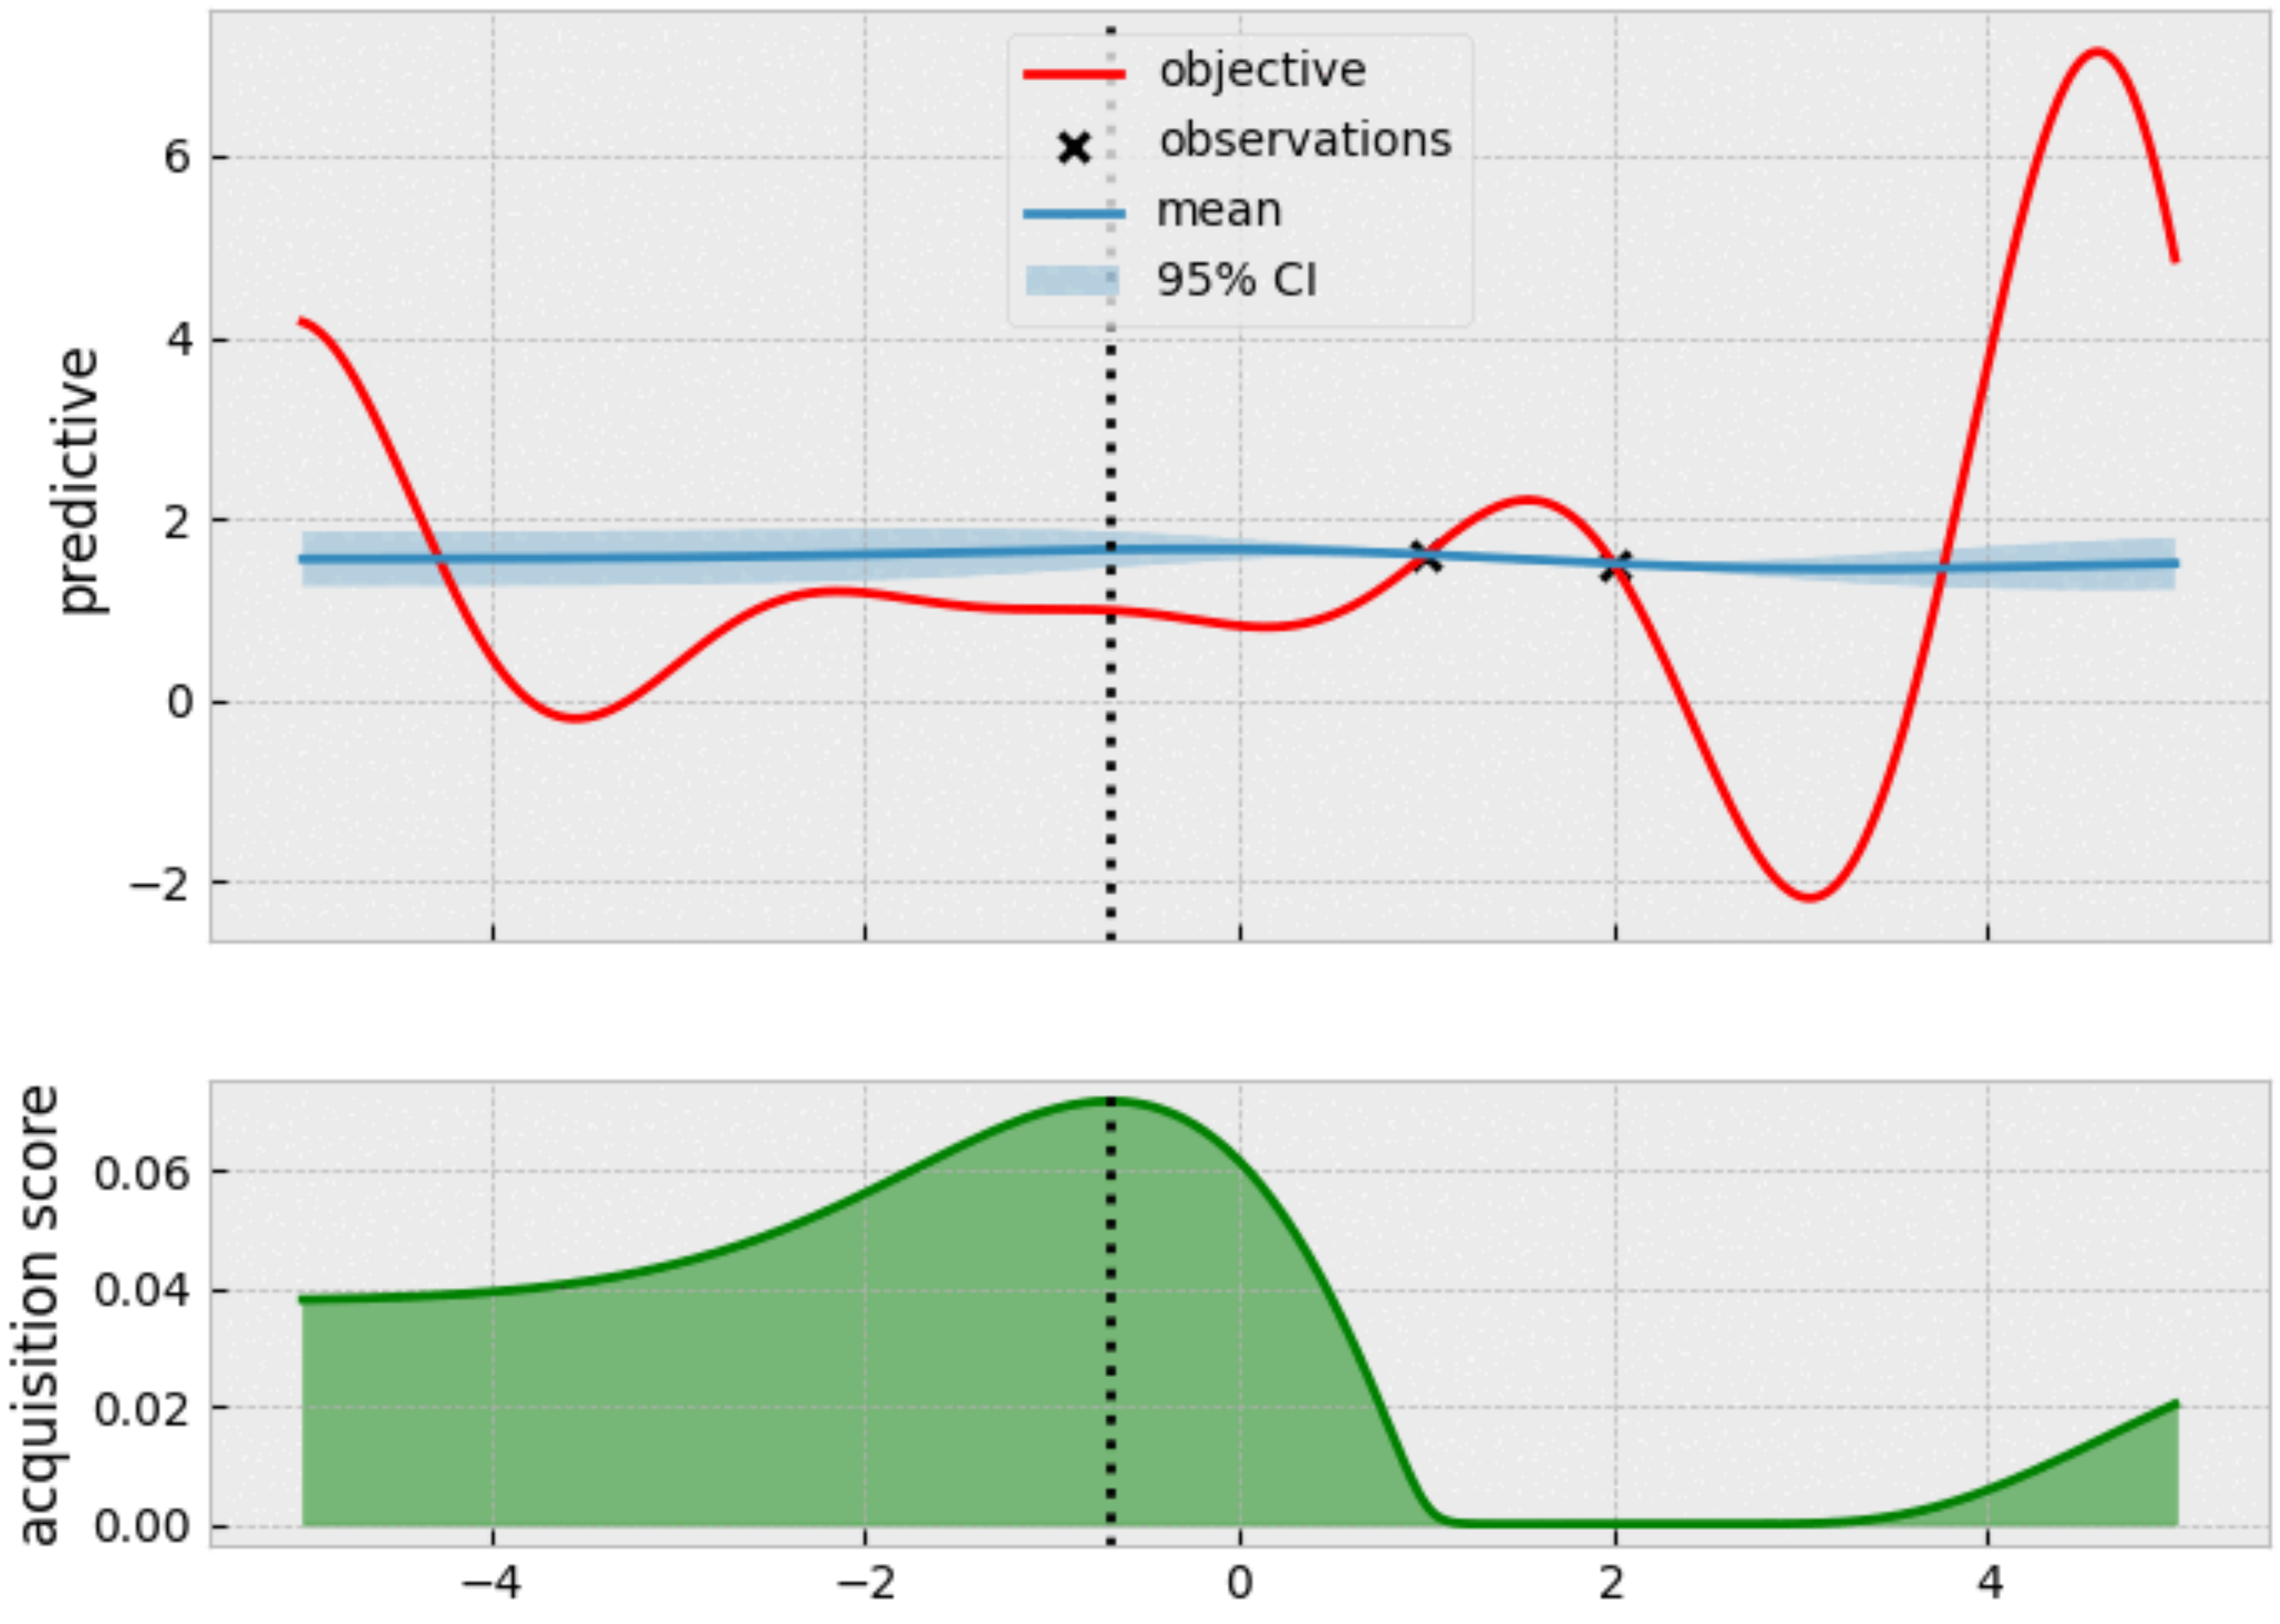

OTHER BAYESOPT POLICIES

OTHER BAYESOPT POLICIES

- Upper confidence bound uses the

  upper credible intervals

OTHER BAYESOPT POLICIES

- Upper confidence bound uses the upper credible intervals

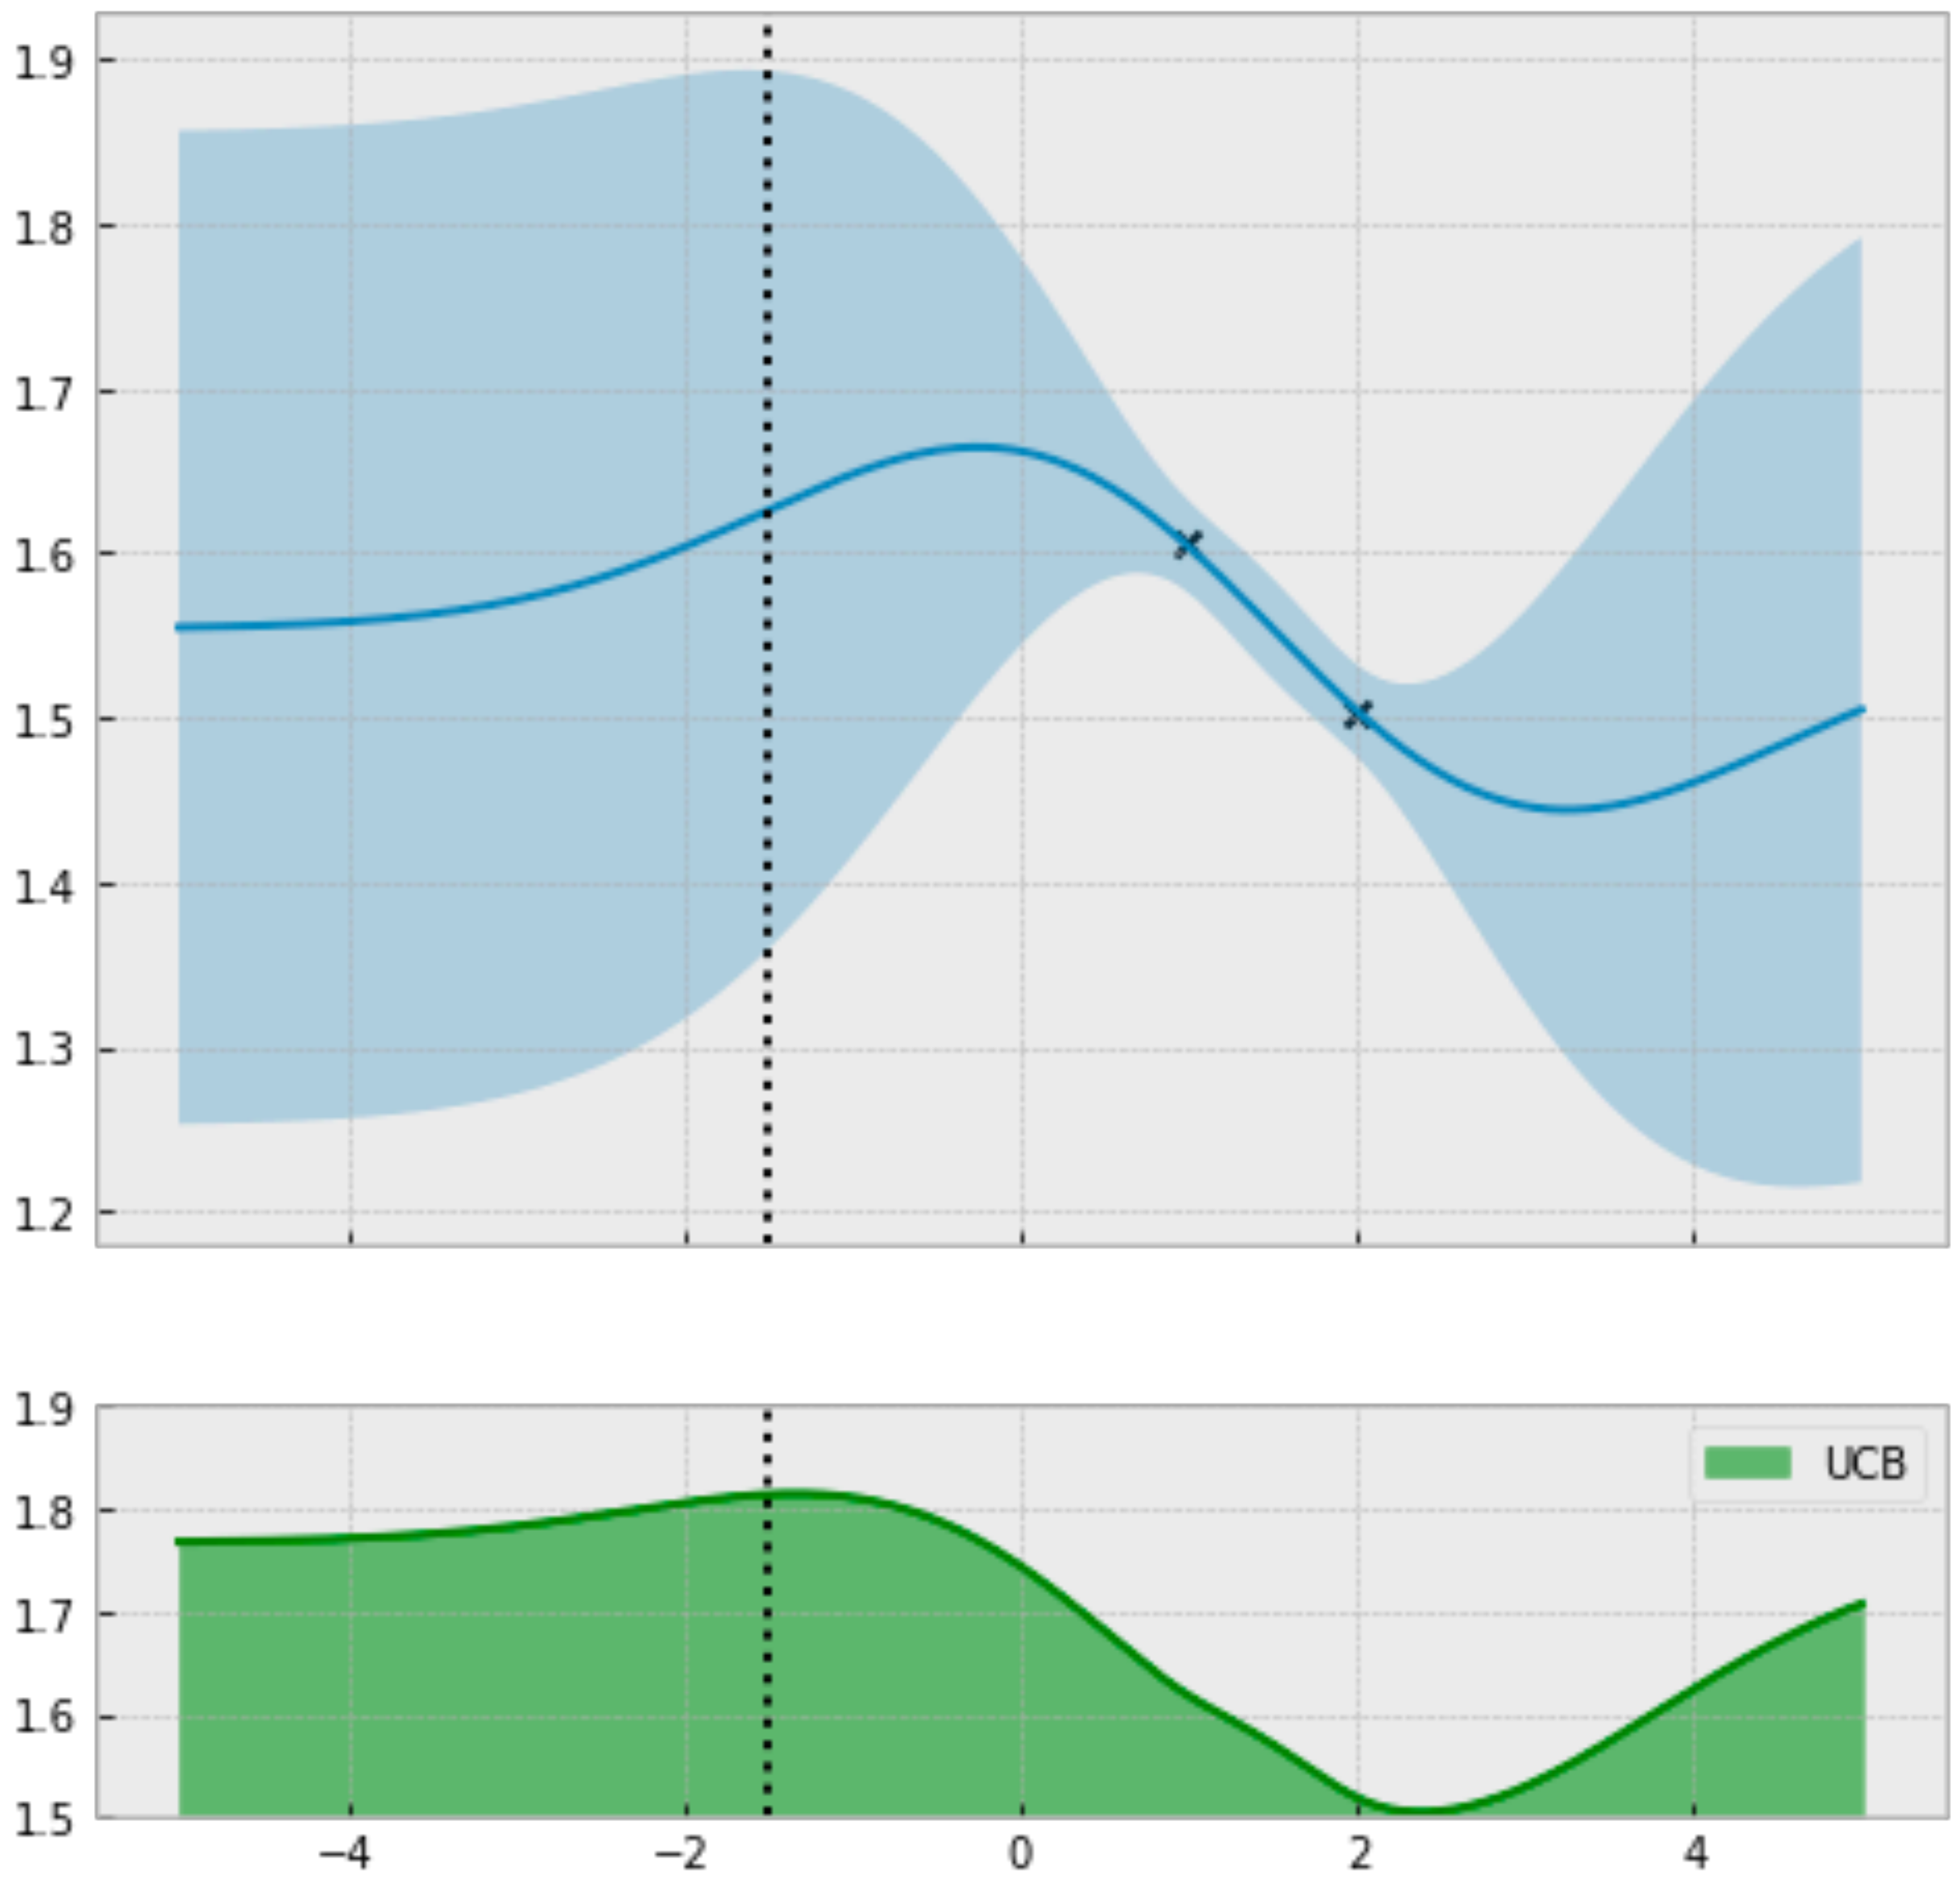

OTHER BAYESOPT POLICIES

- Upper confidence bound uses the upper credible intervals
- Thompson sampling maximizes a sample from the GP

OTHER BAYESOPT POLICIES

- Upper confidence bound uses the upper credible intervals
- Thompson sampling maximizes a sample from the GP

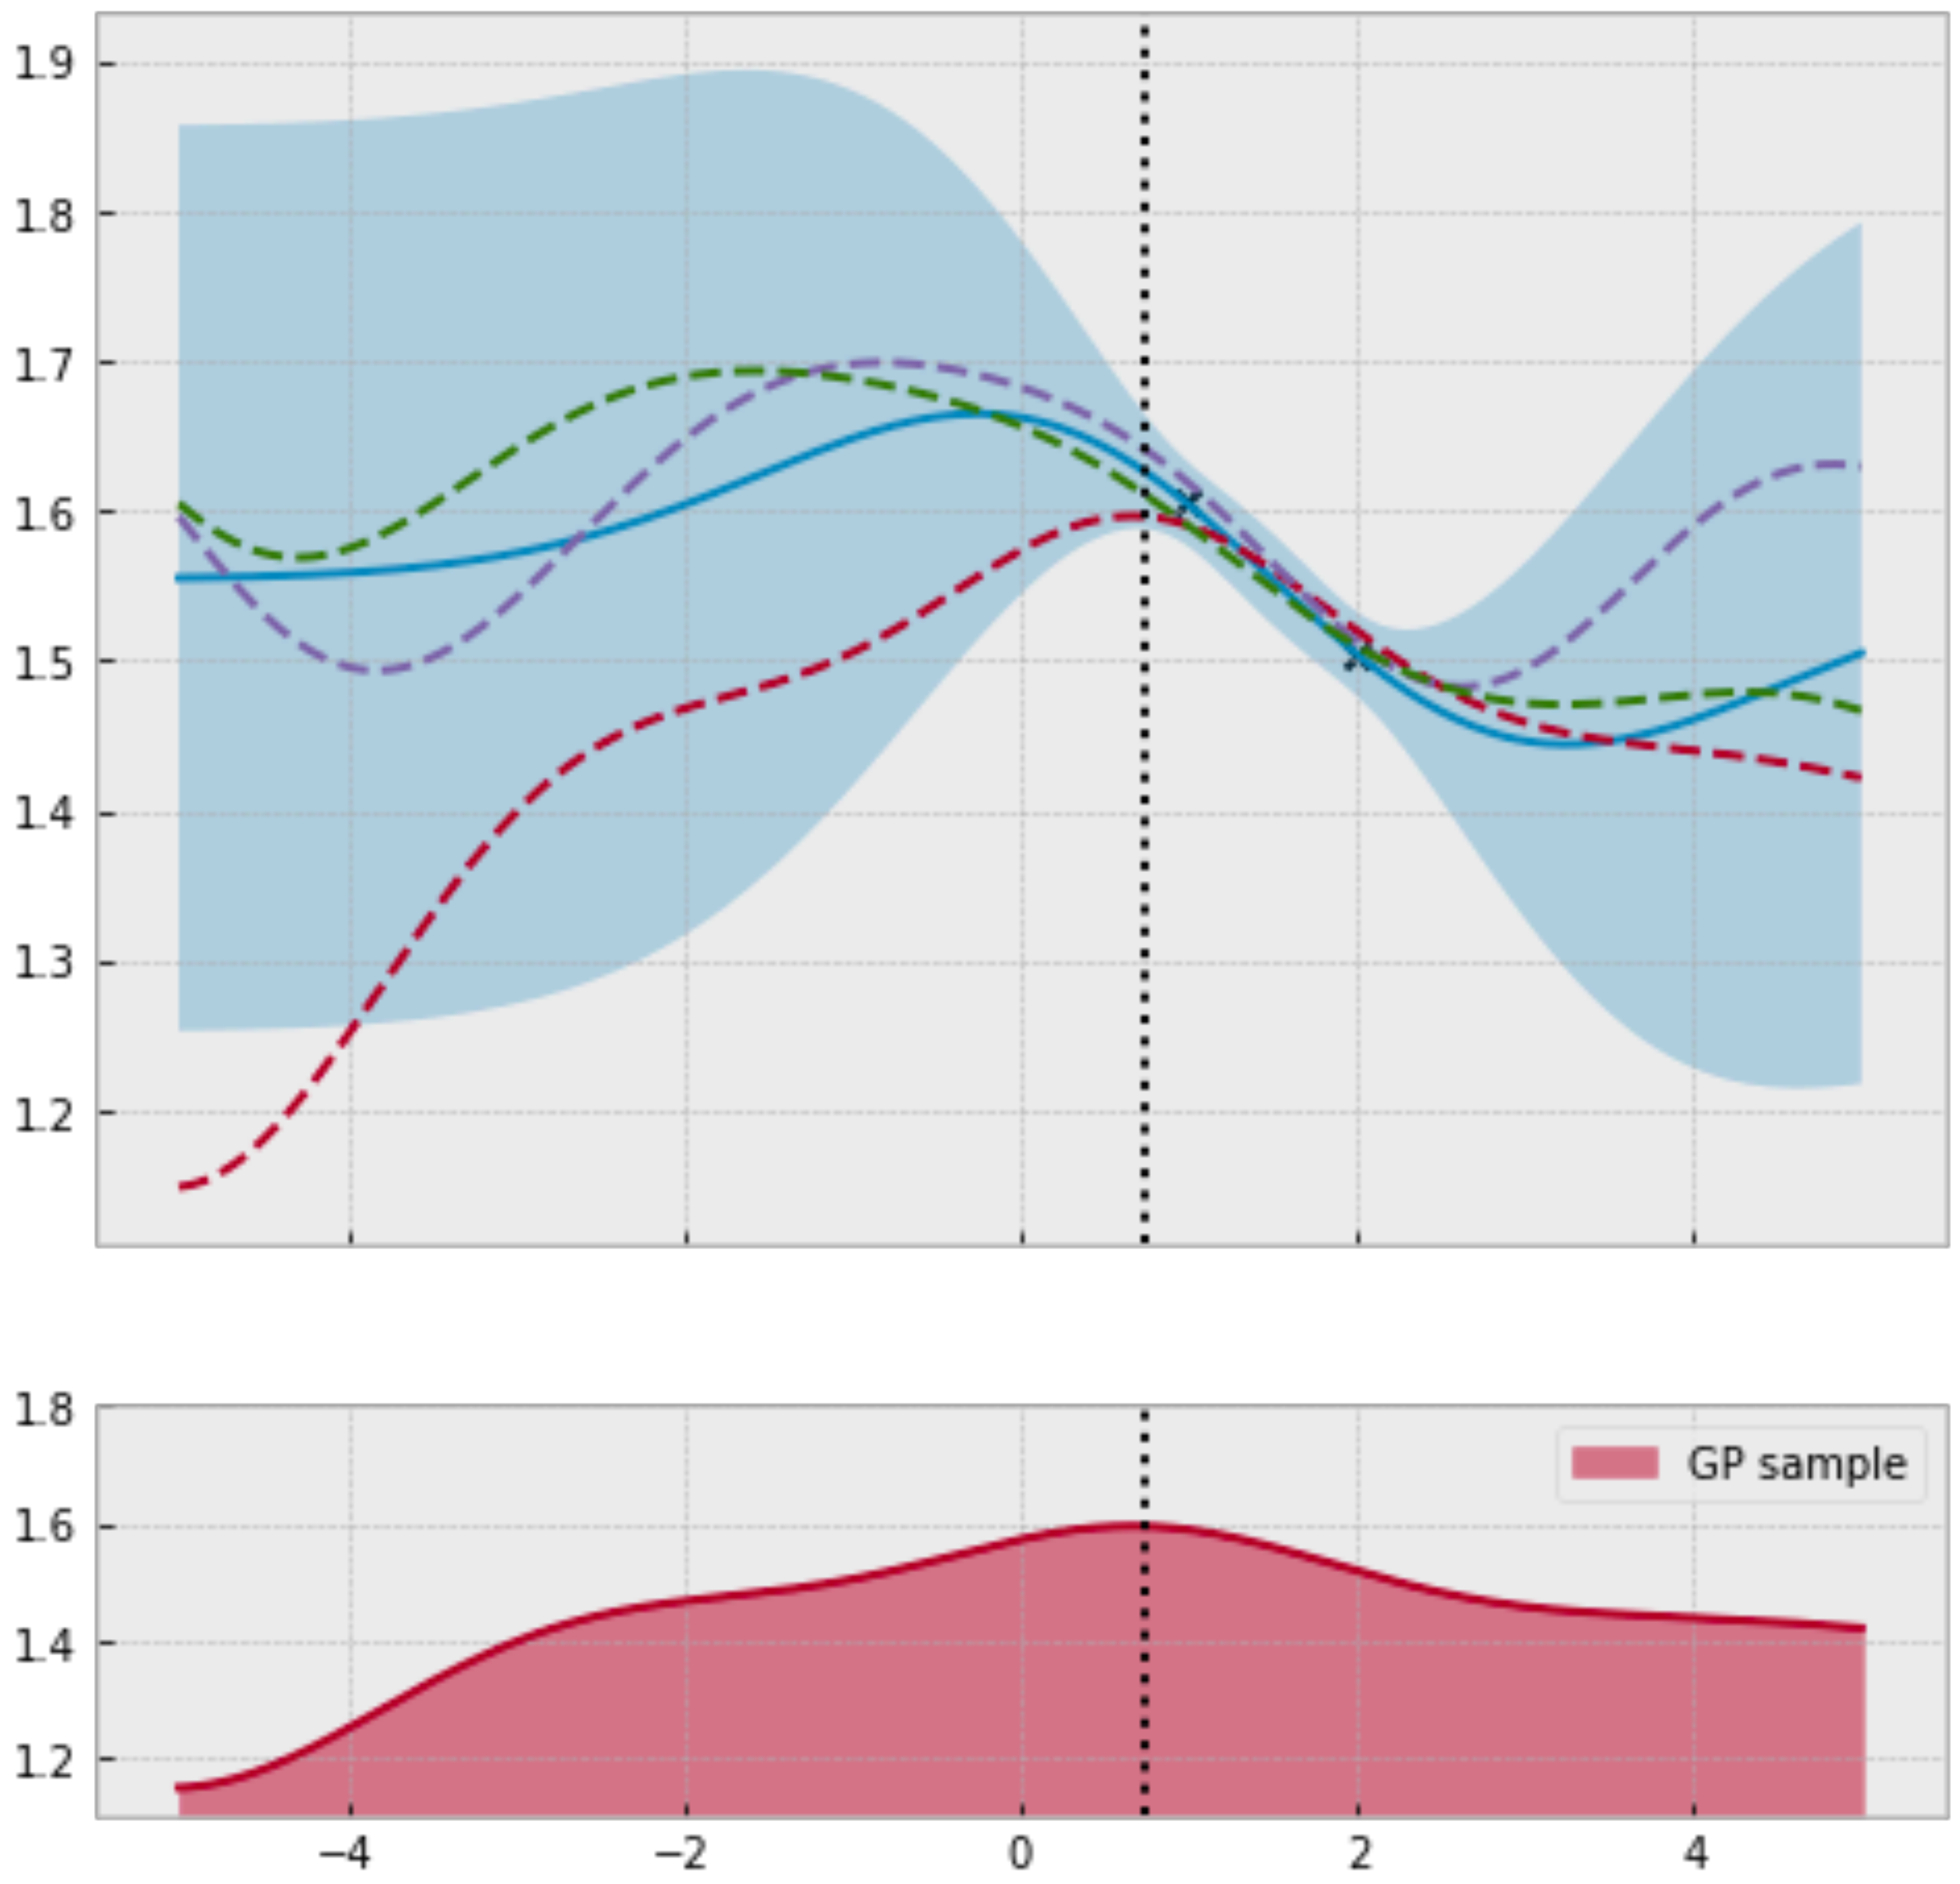

OTHER BAYESOPT POLICIES

- Upper confidence bound uses the upper credible intervals
- Thompson sampling maximizes a sample from the GP

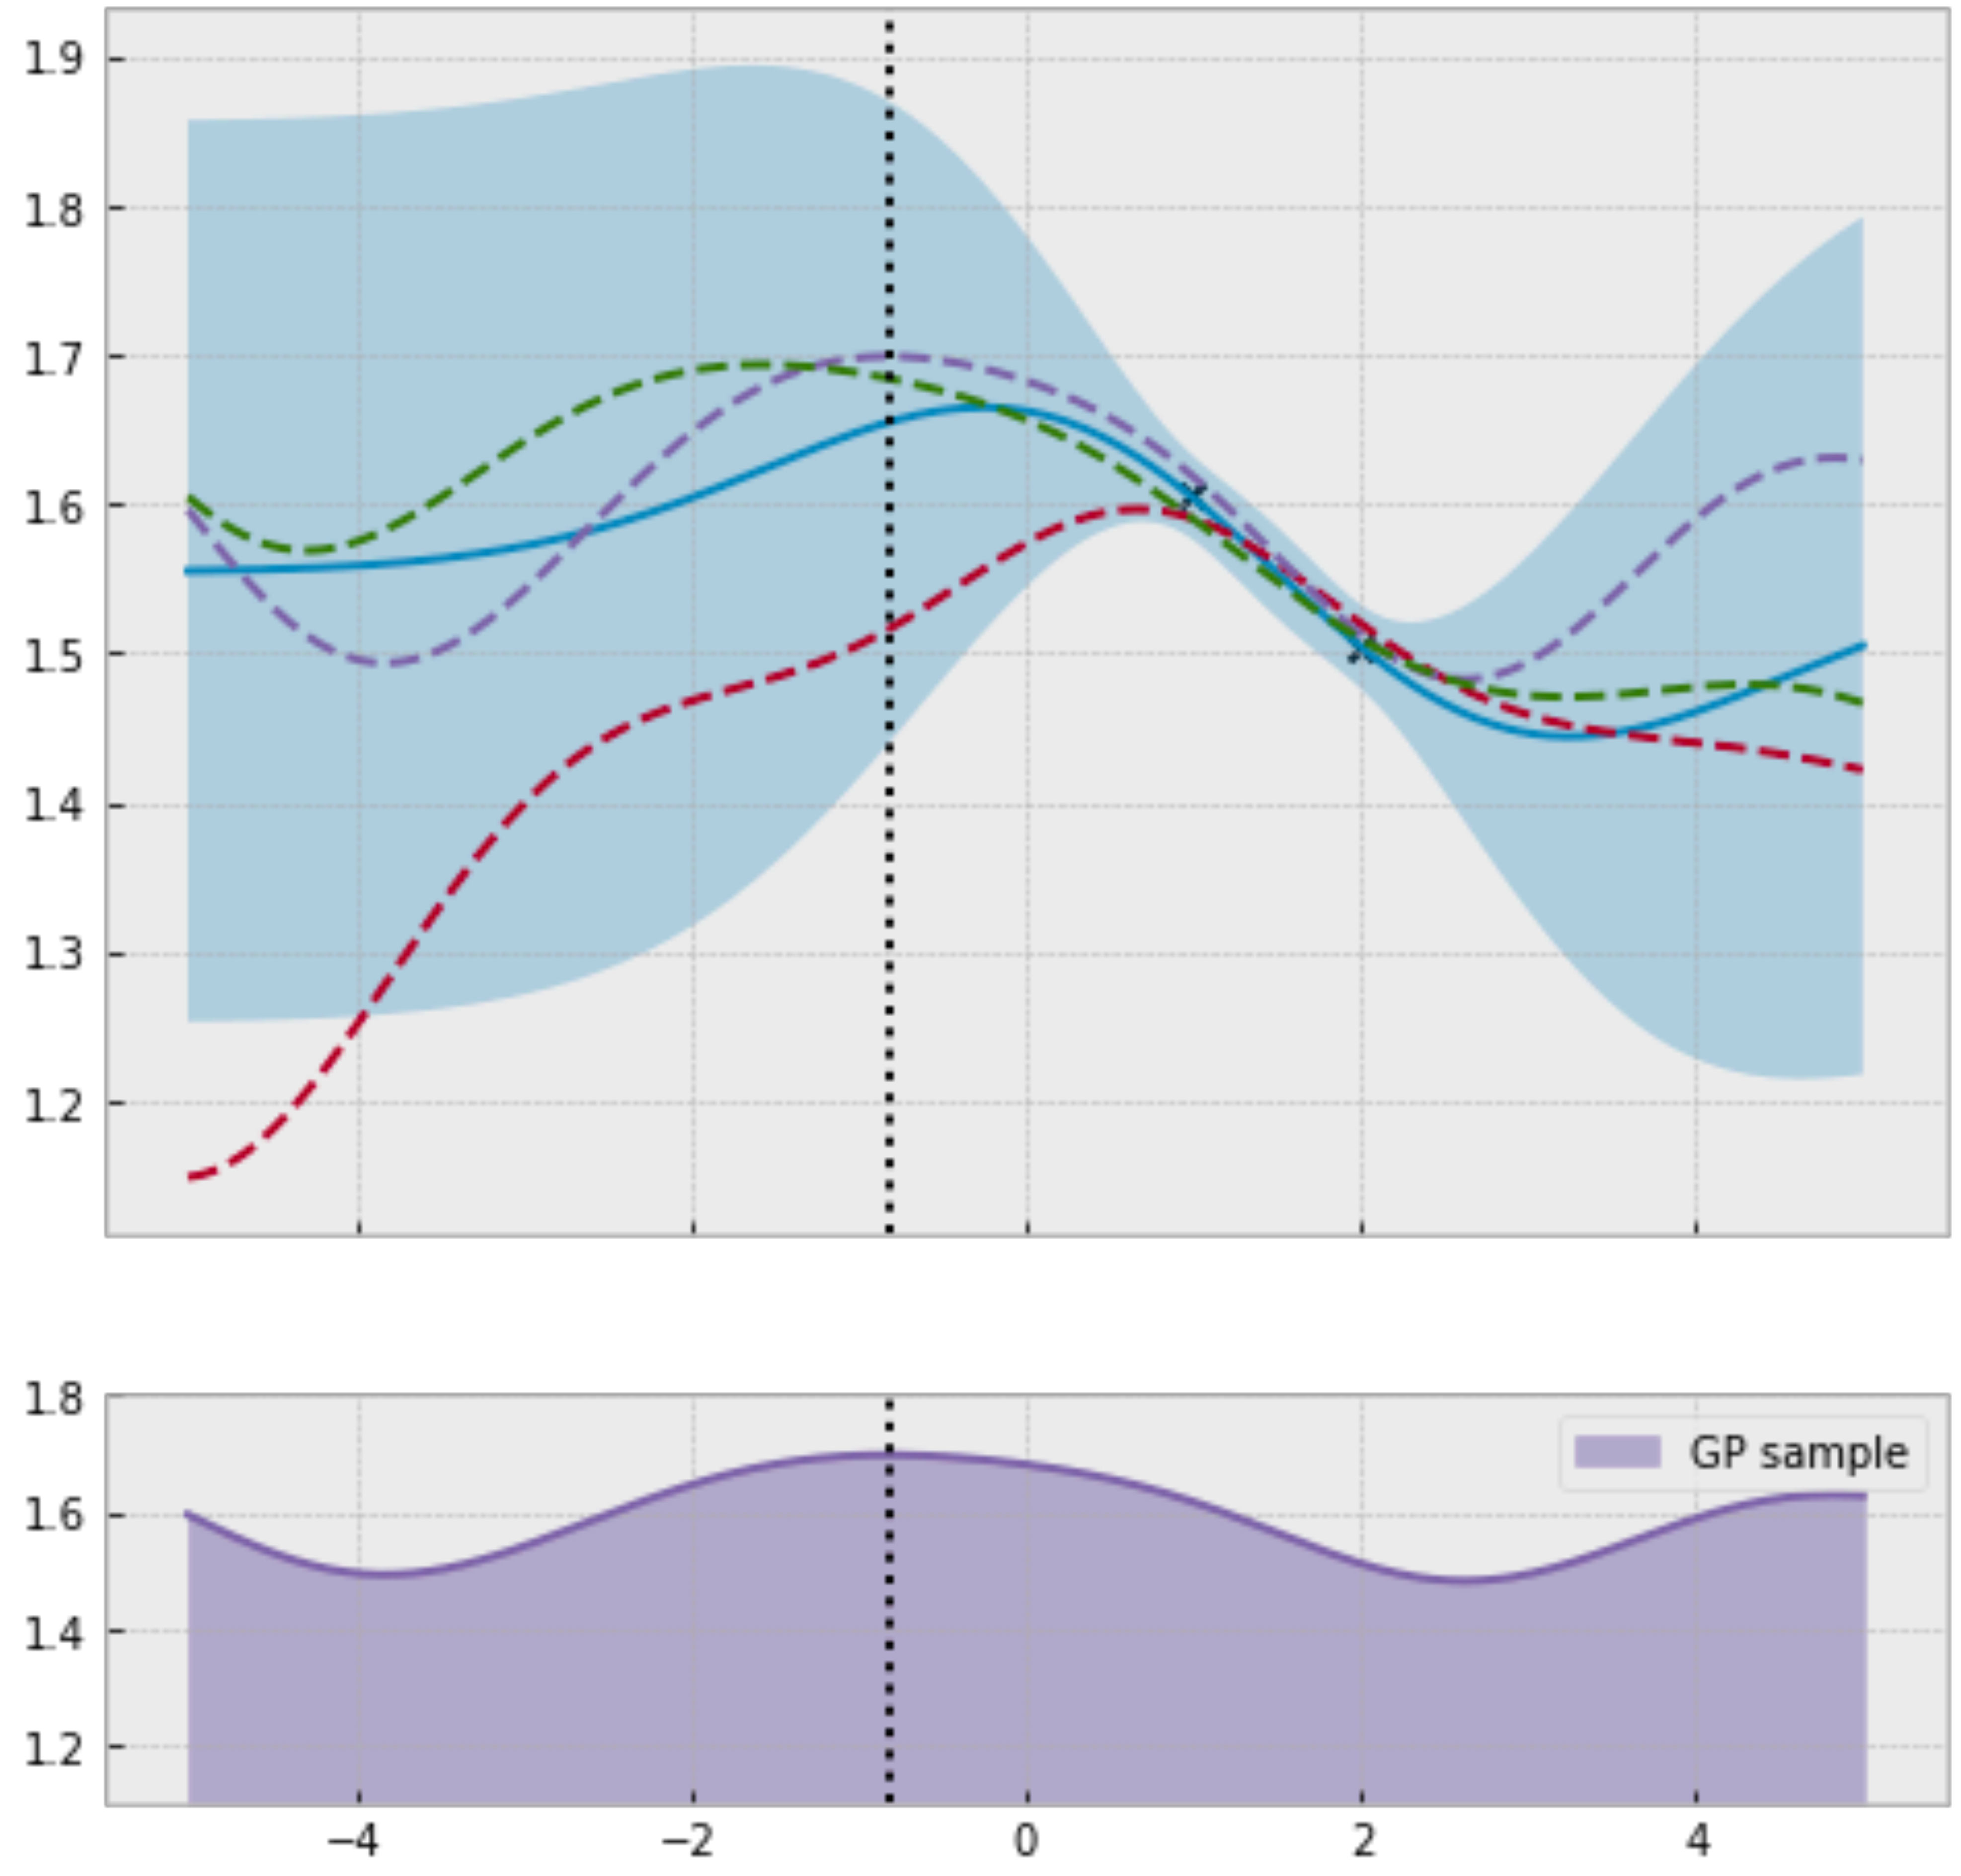

OTHER BAYESOPT POLICIES

- Upper confidence bound uses the upper credible intervals
- Thompson sampling maximizes a sample from the GP

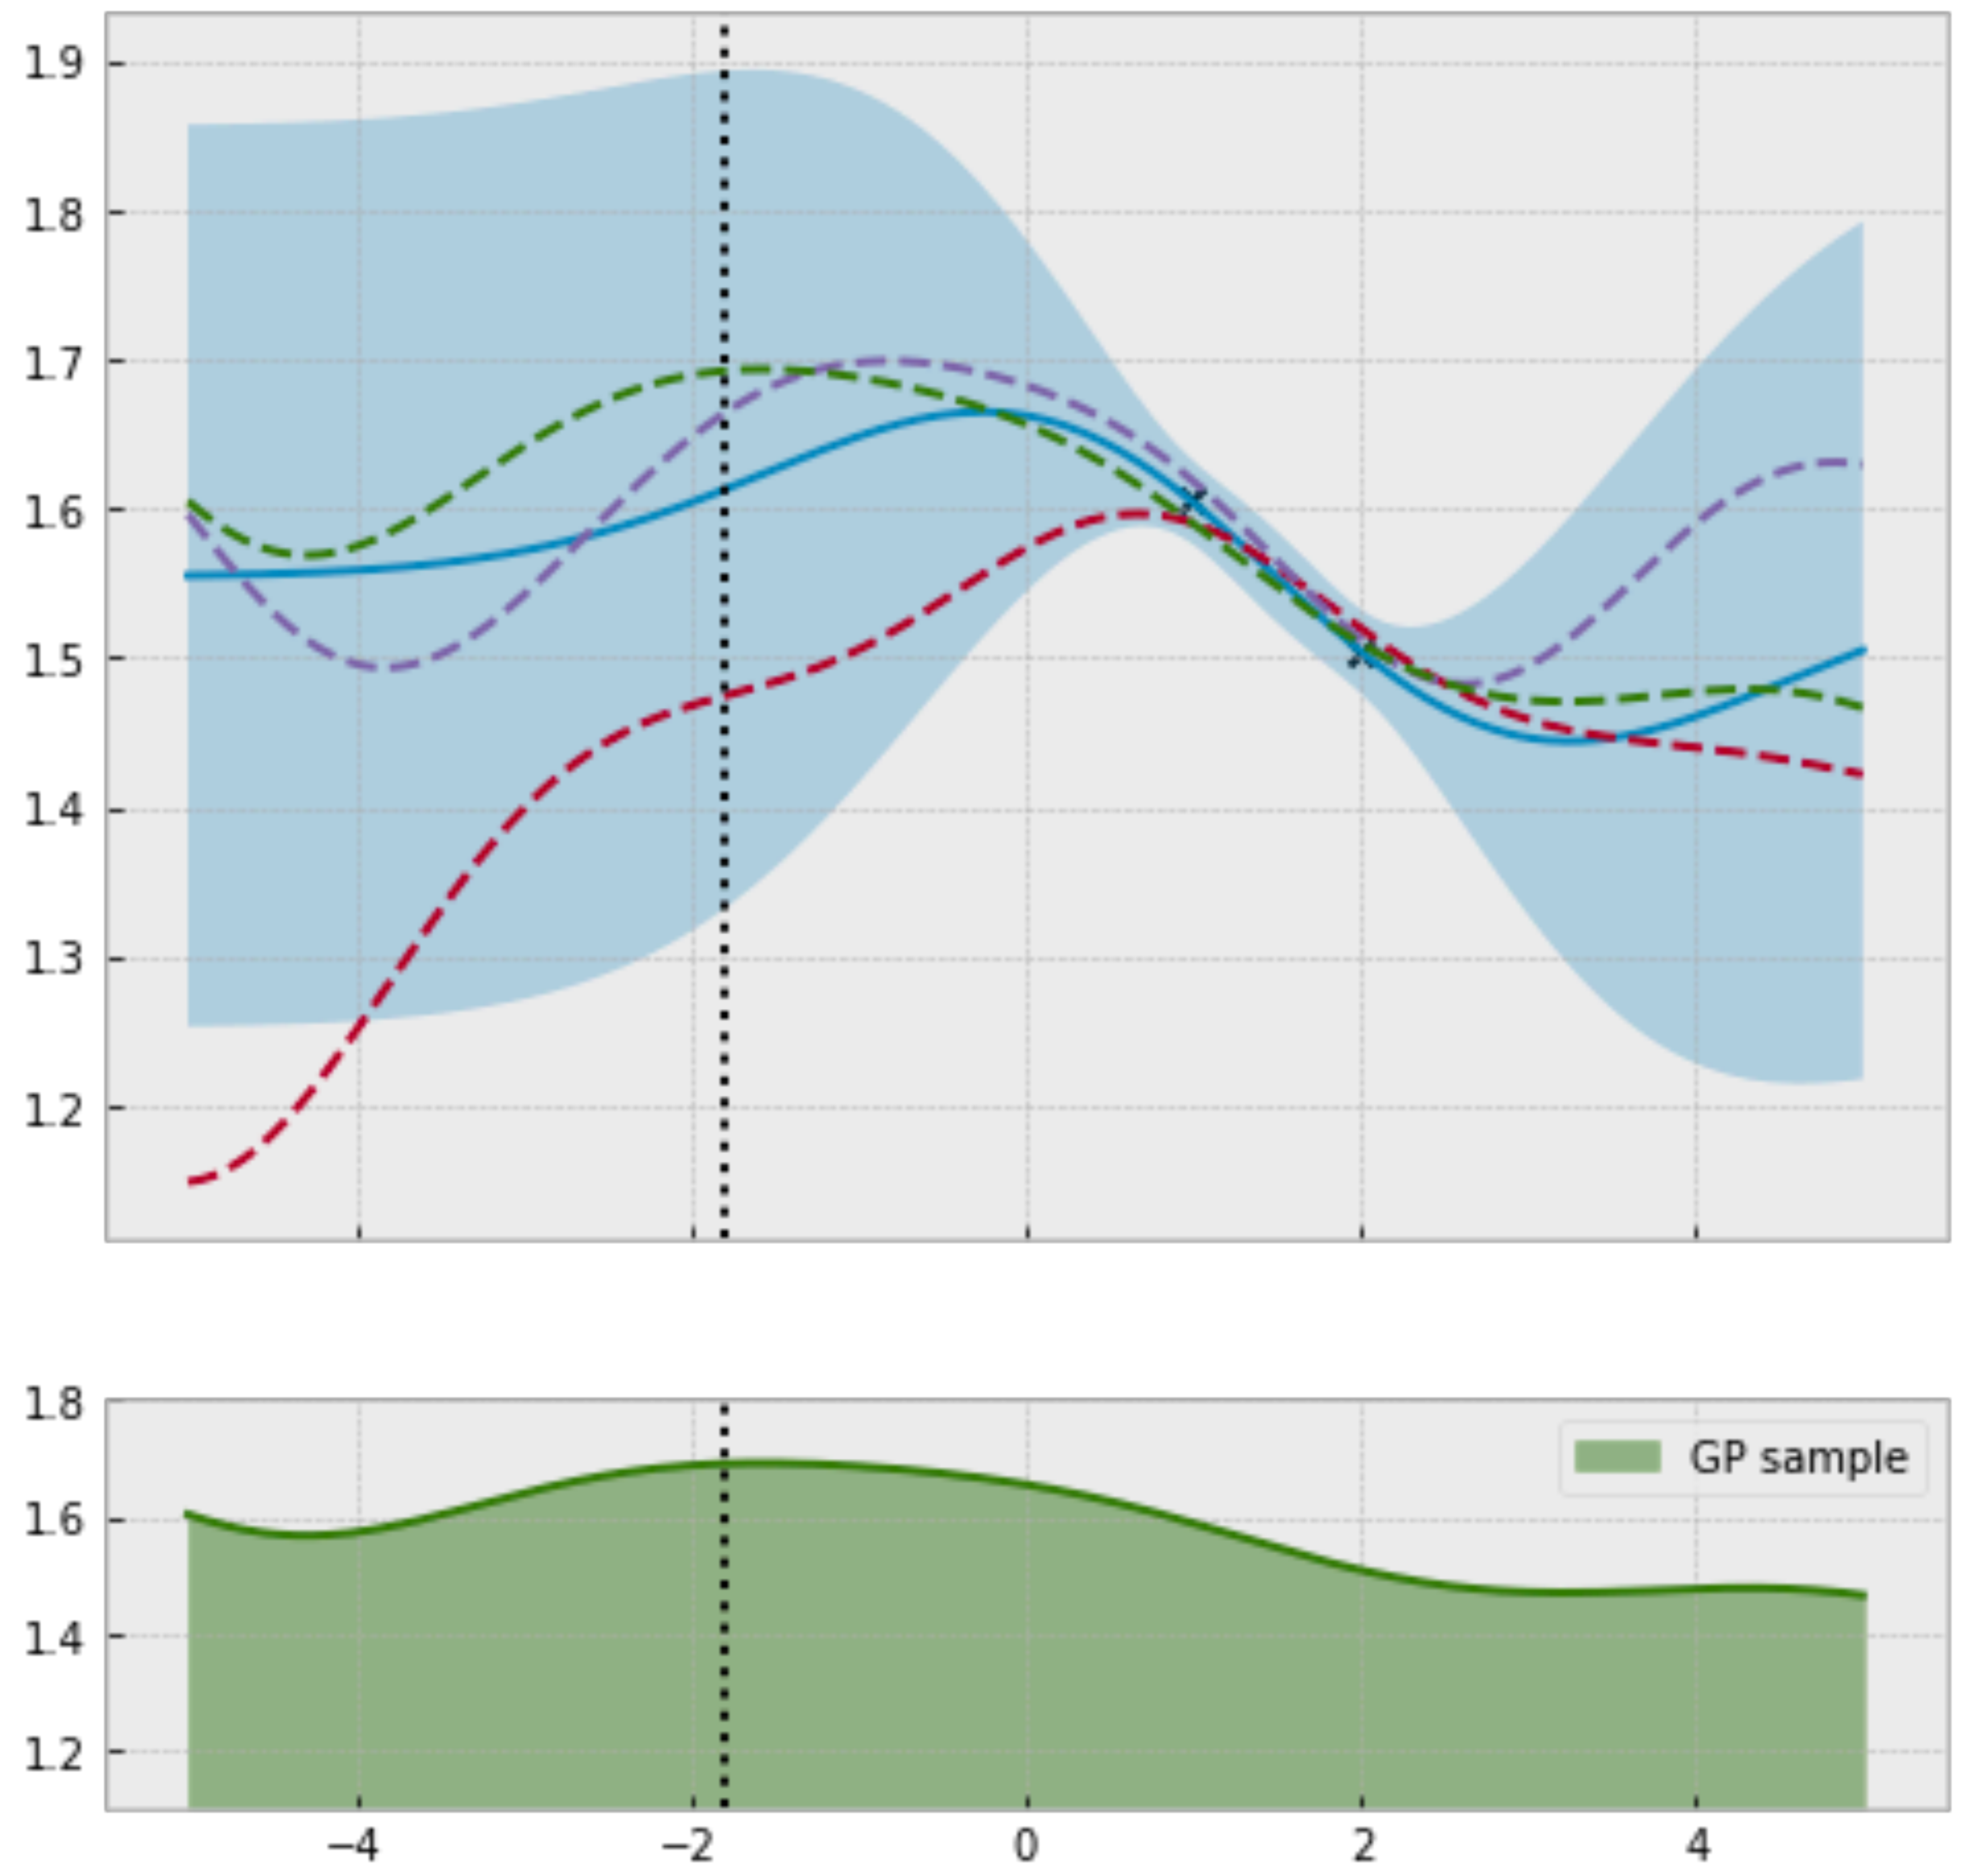

OTHER BAYESOPT POLICIES

- Upper confidence bound uses the upper credible intervals
- Thompson sampling maximizes a sample from the GP
- Entropy search maximizes information gain

OTHER BAYESOPT POLICIES

- Upper confidence bound uses the upper credible intervals
- Thompson sampling maximizes a sample from the GP
- Entropy search maximizes information gain

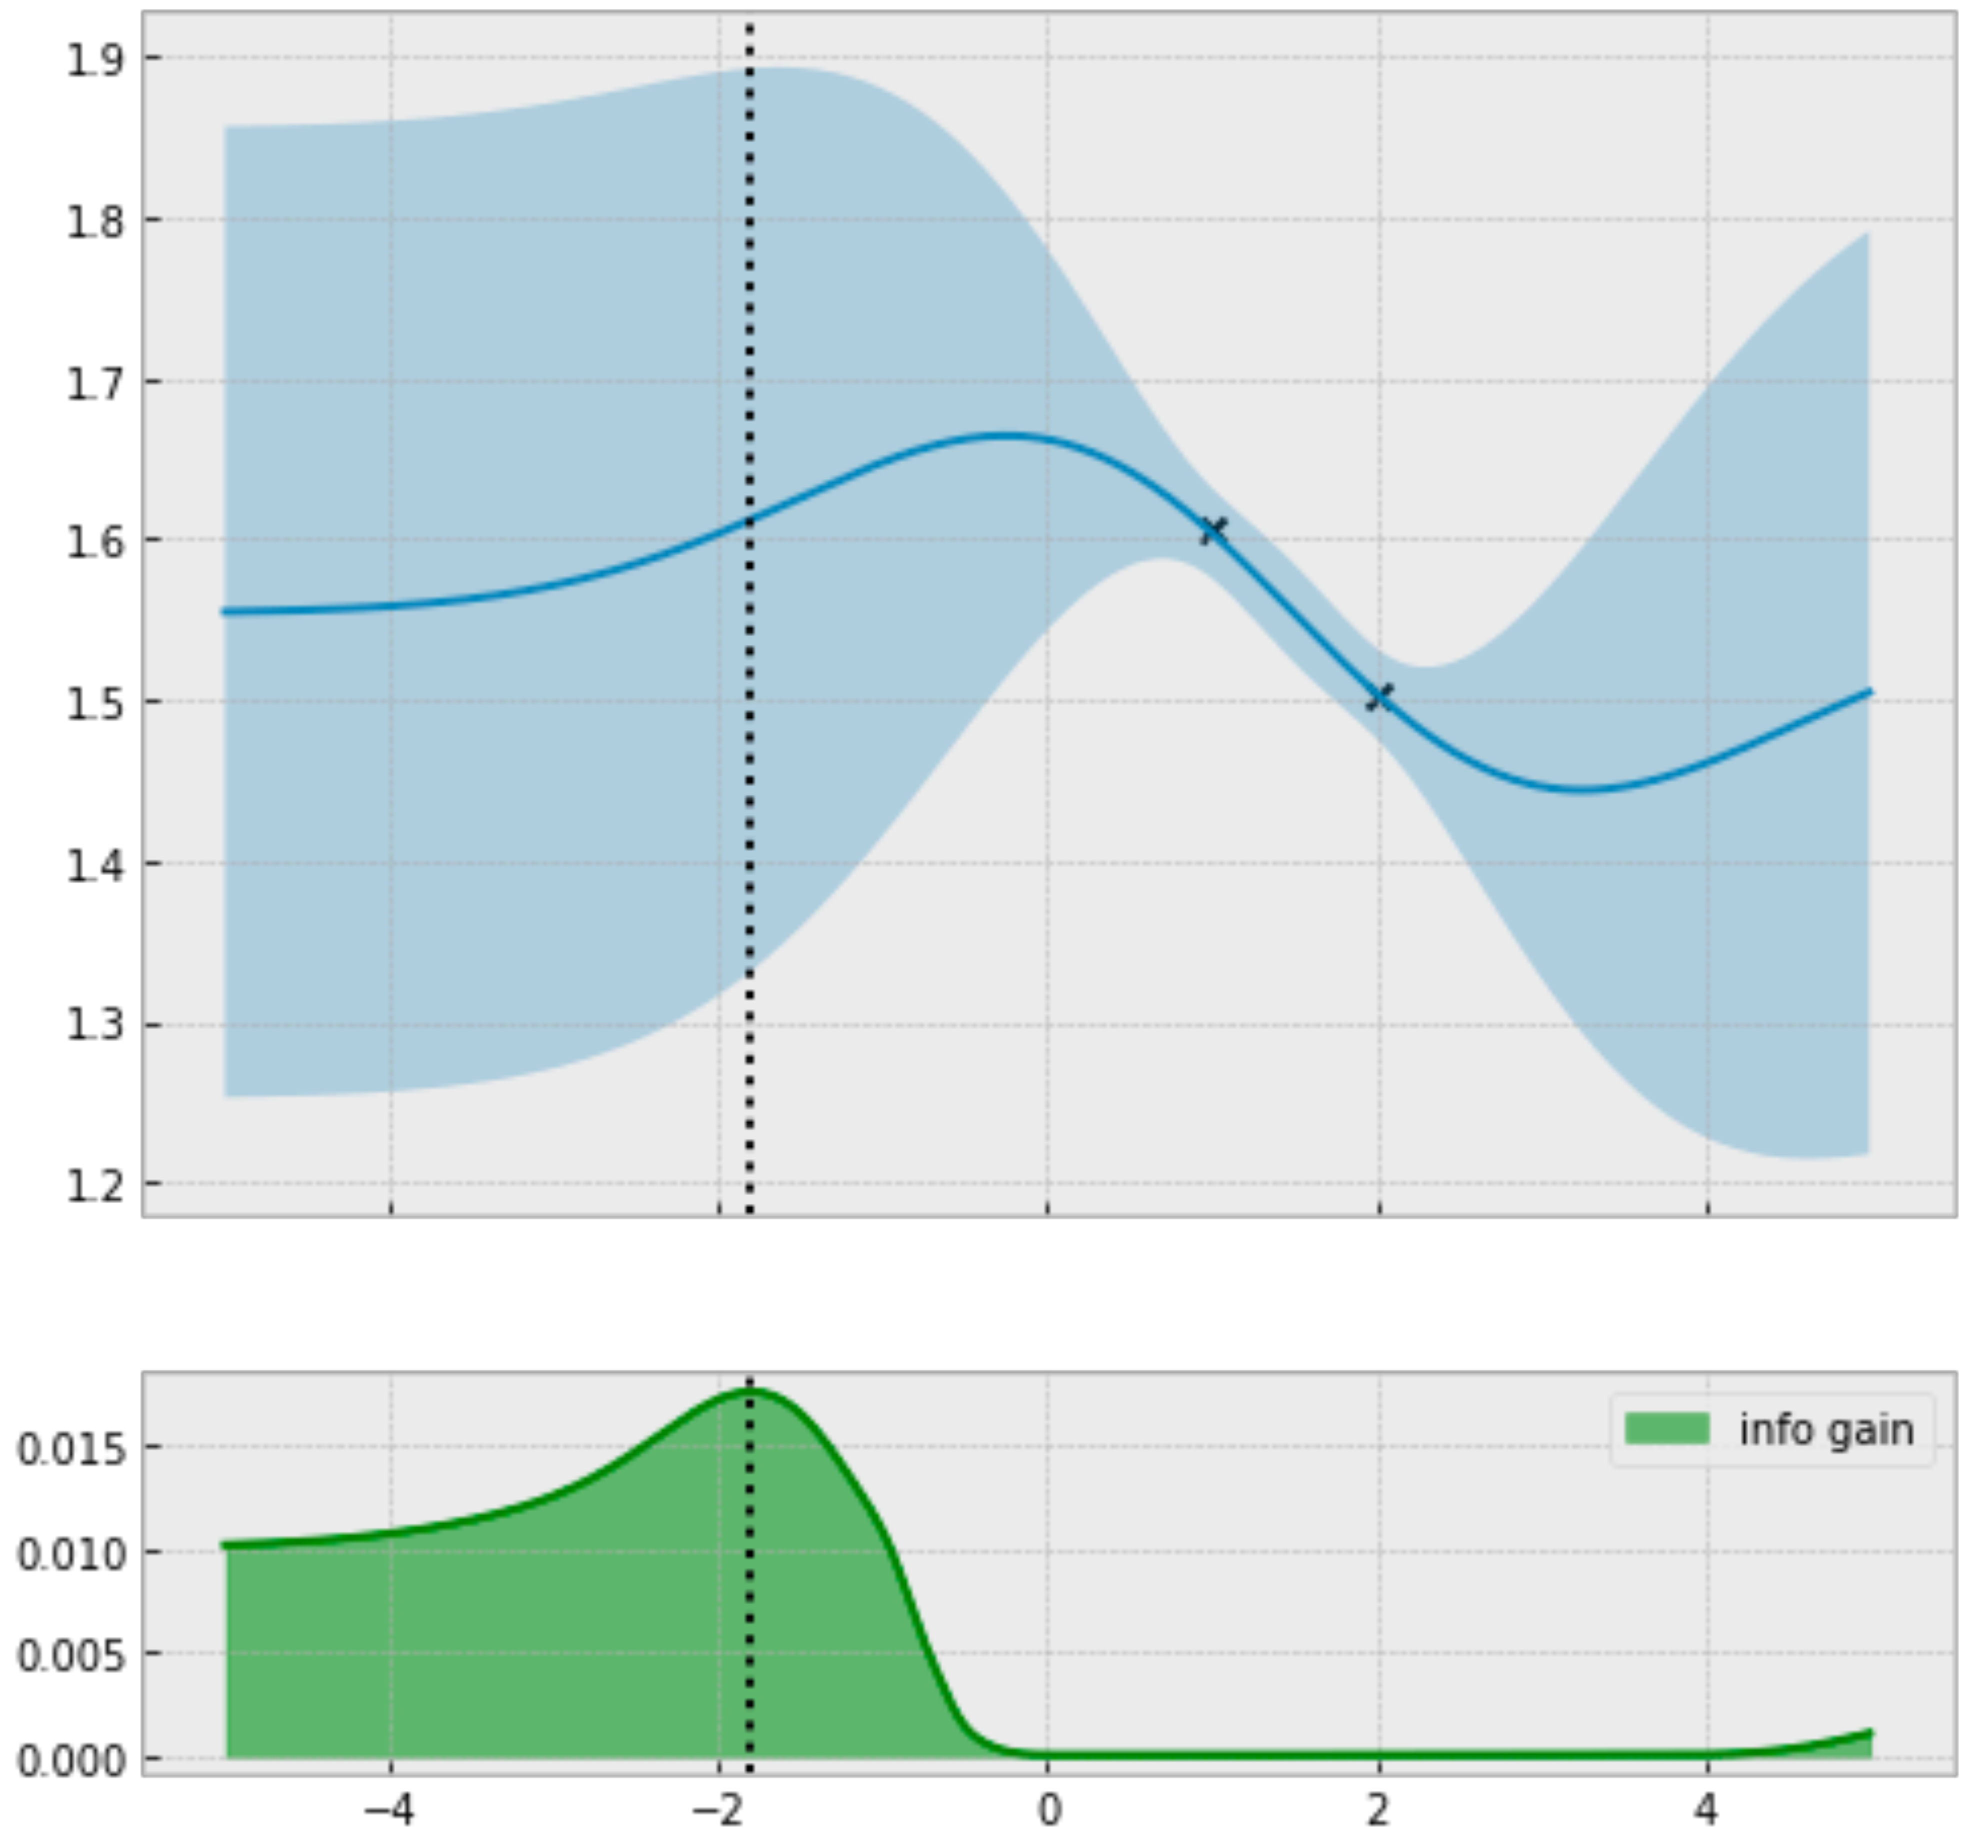

NONMYOPIA IN BAYESIAN OPTIMIZATION

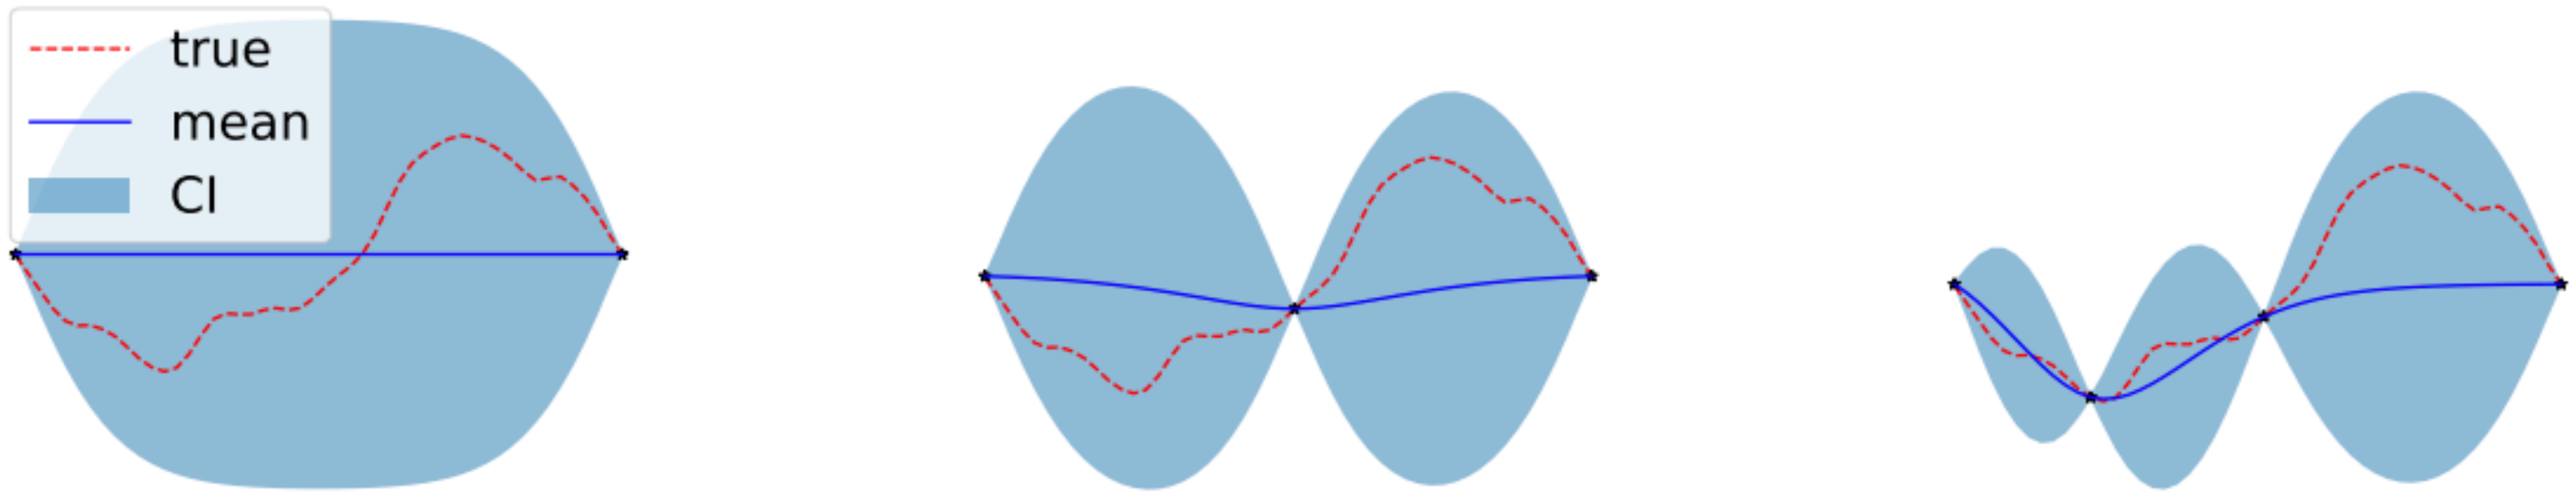

(a) initial state (b) EI iteration 1 (c) EI iteration 2

BINOCULARS for Efficient, Nonmyopic Sequential Experimental Design, Jiang et al., ICML 2020

NONMYOPIA IN BAYESIAN OPTIMIZATION

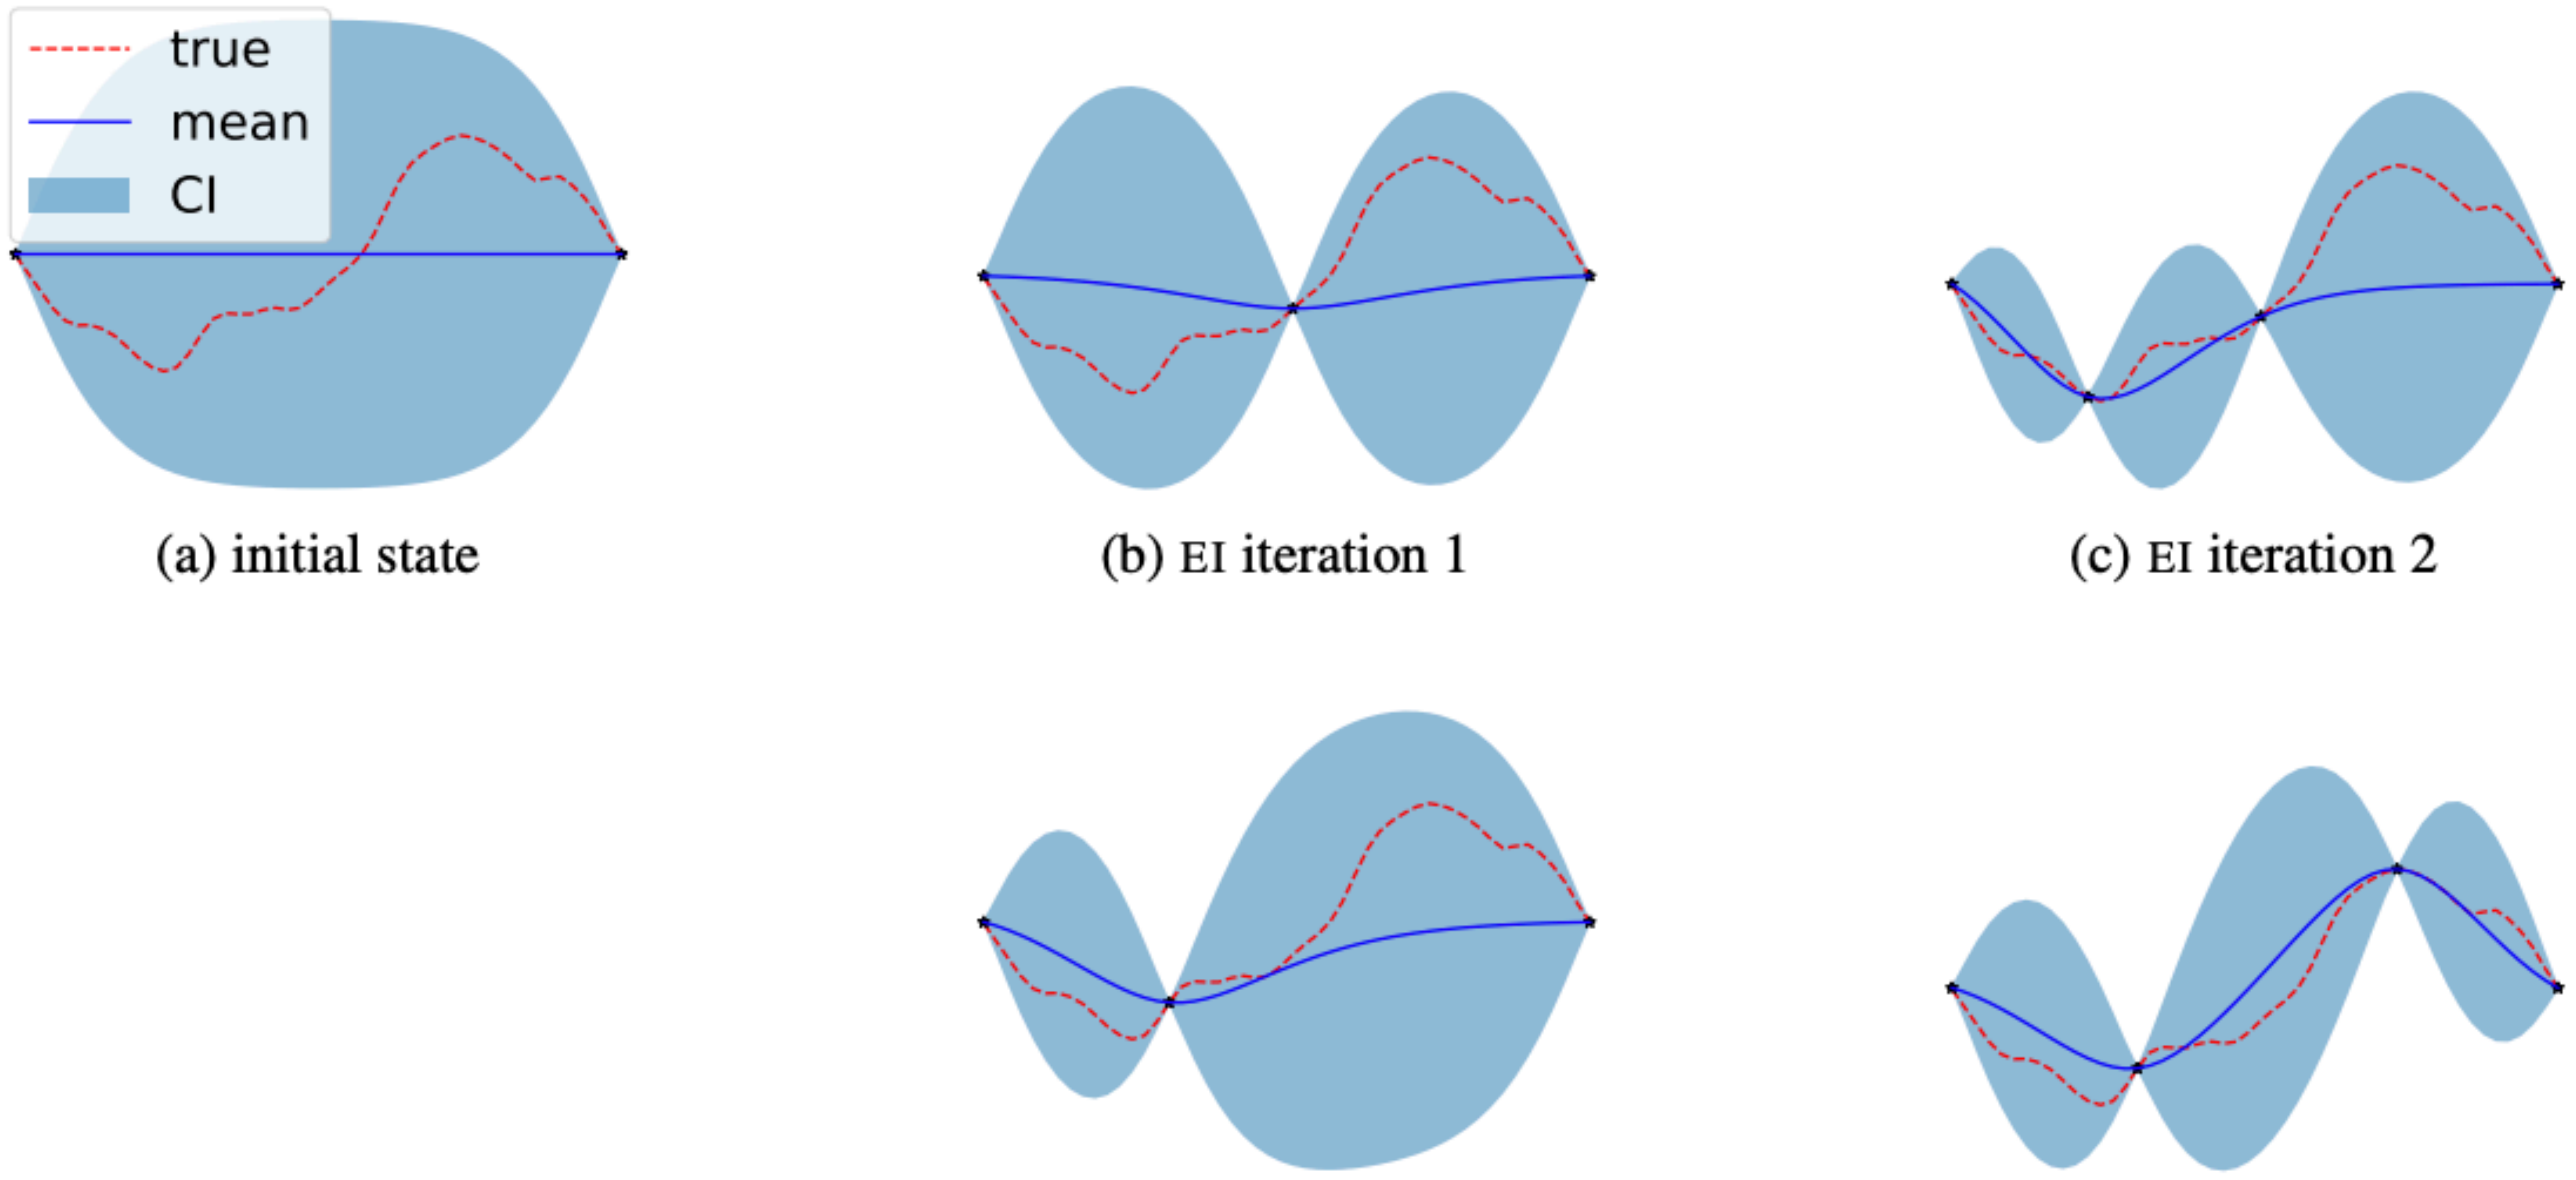

(a) initial state

(b) EI iteration 1

(c) EI iteration 2

(e) 2-EI iteration 1

(f) 2-EI iteration 2

BINOCULARS for Efficient, Nonmyopic Sequential Experimental Design, Jiang et al., ICML 2020