Module 13: Final Exam Prep

DAV-6300-1: Experimental Optimization

David Sweet // 20241217

Overview

- A/B testing I: Measurement
- A/B testing II: Design, Measure, Analyze
- A/B testing III: Practical Concerns
- Multi-armed bandits I
- Multi-armed bandits II: Thompson sampling
- Contextual bandits
- Response surface methodology
- Gaussian process regression
- Bayesian optimization

Review: Business Metric (BM)

- Observe business metric, y_i
- Ex: incident of fraud, clicking “like”, sharing a video, clicking on an ad, skipping a song, swiping left or right, lingering on a photo, watching a video until the end
- Ex, daily: revenue, active users, shares traded, pnl

Examples?

Review: Law of Large Numbers

- N observations, y_i, the business metric
- w/mean $\mu = \frac{\Sigma_i^N y_i}{N}$
- As $N \rightarrow \infty$, $\mu \rightarrow E[y]$
- IOW: Our measurement (μ) estimates the "true" business metric

Review: Central Limit Theorem

- As $N \to \infty$, $\mu \sim \mathcal{N}(E[y], VAR[y]/N)$
- IOW: Measurement (μ) is normally distributed
  - …even if observations (y_i) are not
  - …when we have enough observations

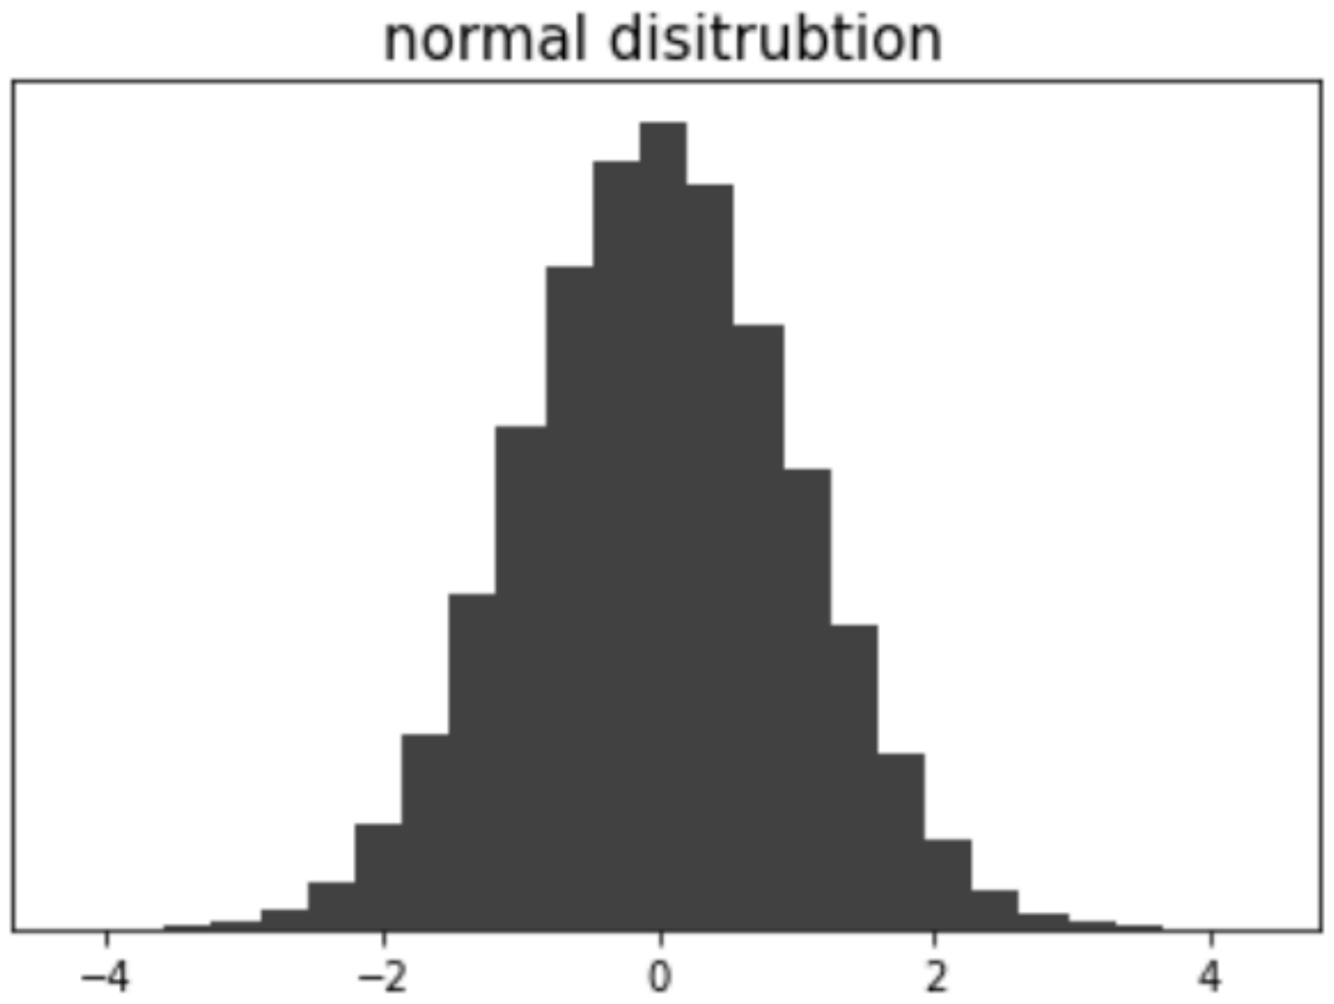

Review: Central Limit Theorem

- Observations, y_i, may have any distribution:

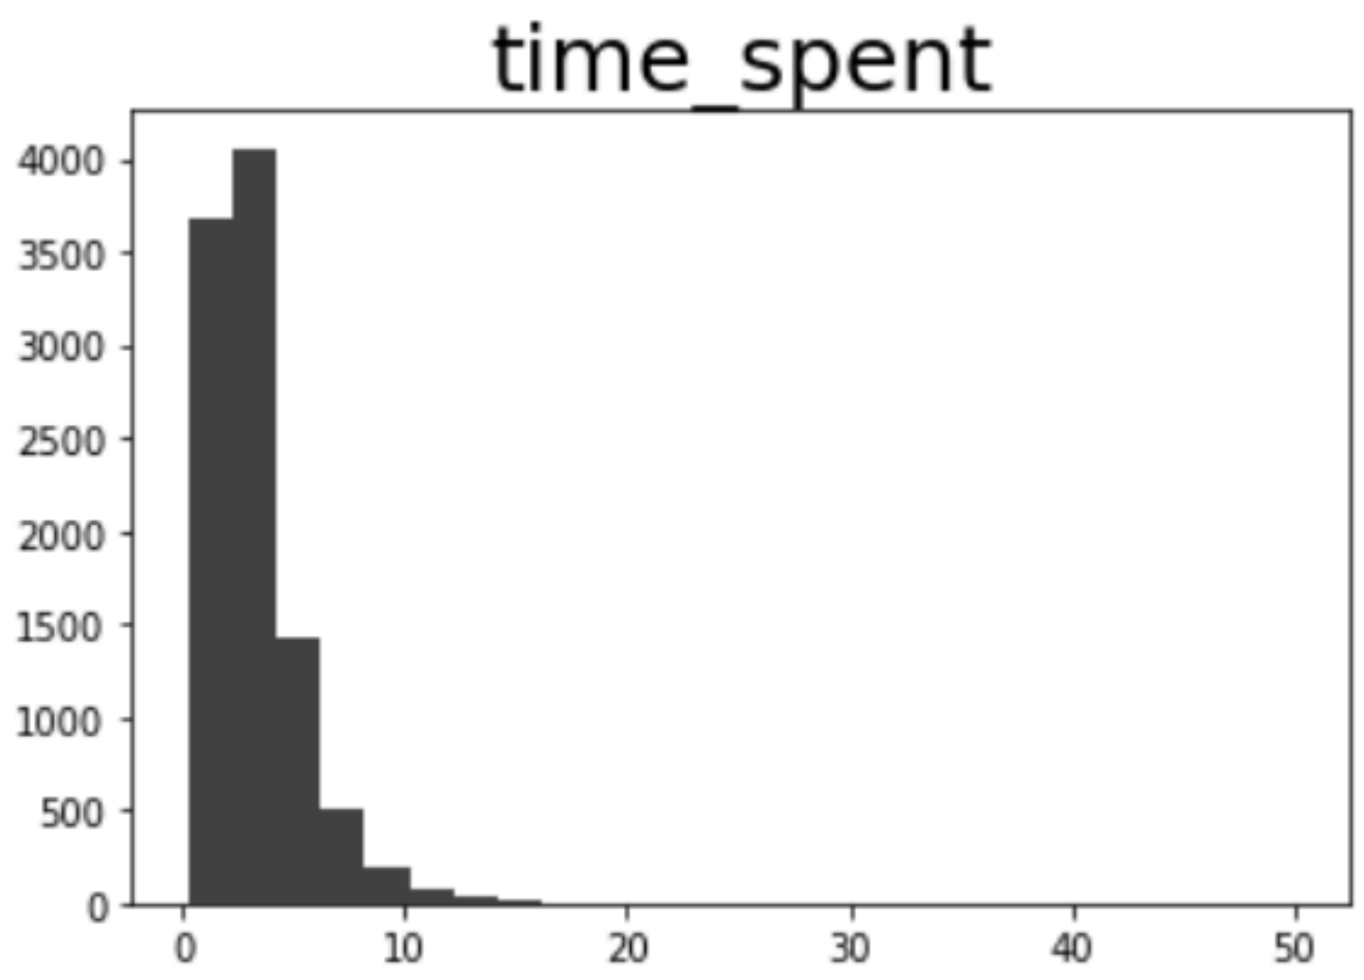

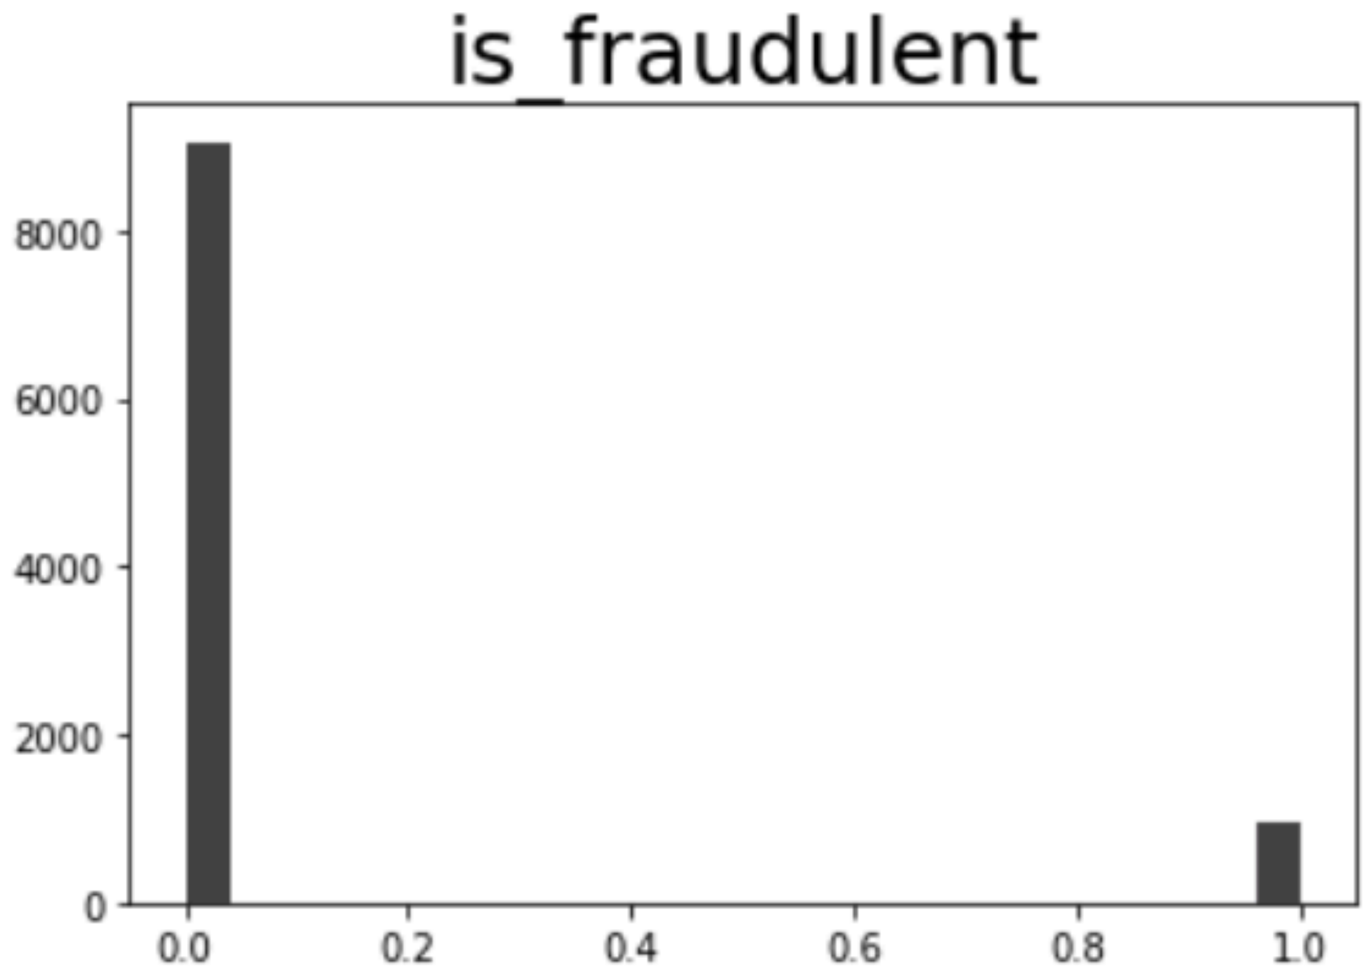

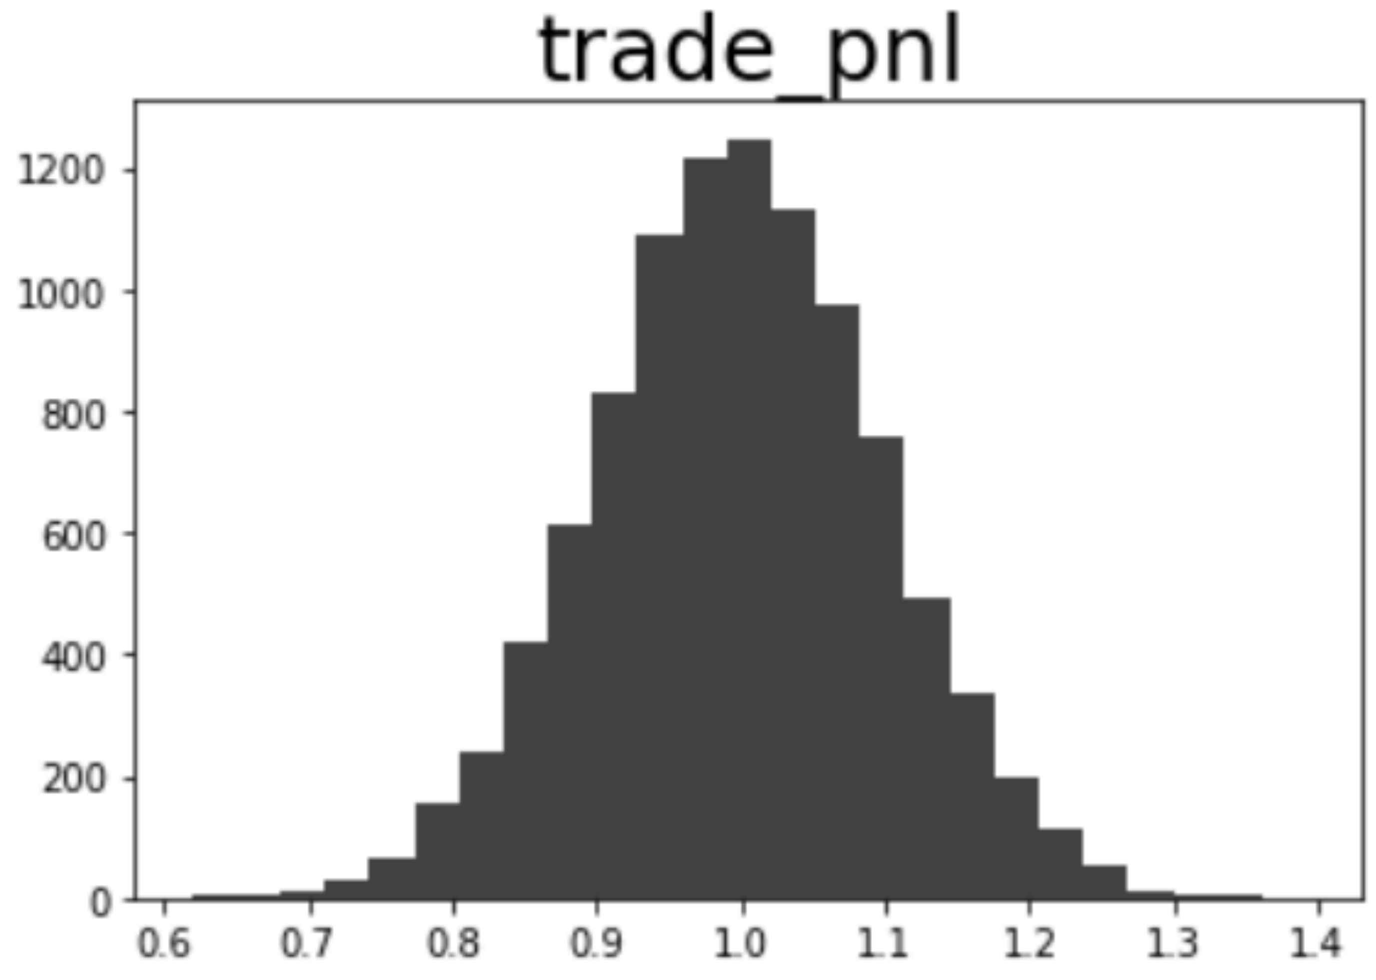

- still, $\mu \sim \mathcal{N}(E[y], VAR[y]/N)$ for large N

$\mu = y$ when $N = 1$

Review: Law of Large Numbers

- N observations, y_i, the business metric
- w/mean $\mu = \dfrac{\Sigma_i^N y_i}{N}$
- As $N \to \infty$, $\mu \to E[y]$
- IOW: Our measurement (μ) estimates the “true” business metric

Review: Central Limit Theorem

- As $N \to \infty$, $\mu \sim \mathcal{N}(E[y], VAR[y]/N)$
- IOW: Measurement (μ) is normally distributed
  - ...even if observations (y_i) are not
  - ...when we have enough observations
- σ = estimates $STD(y)$
- $se = \sigma/\sqrt{N}$ = estimates $STD(\mu)$

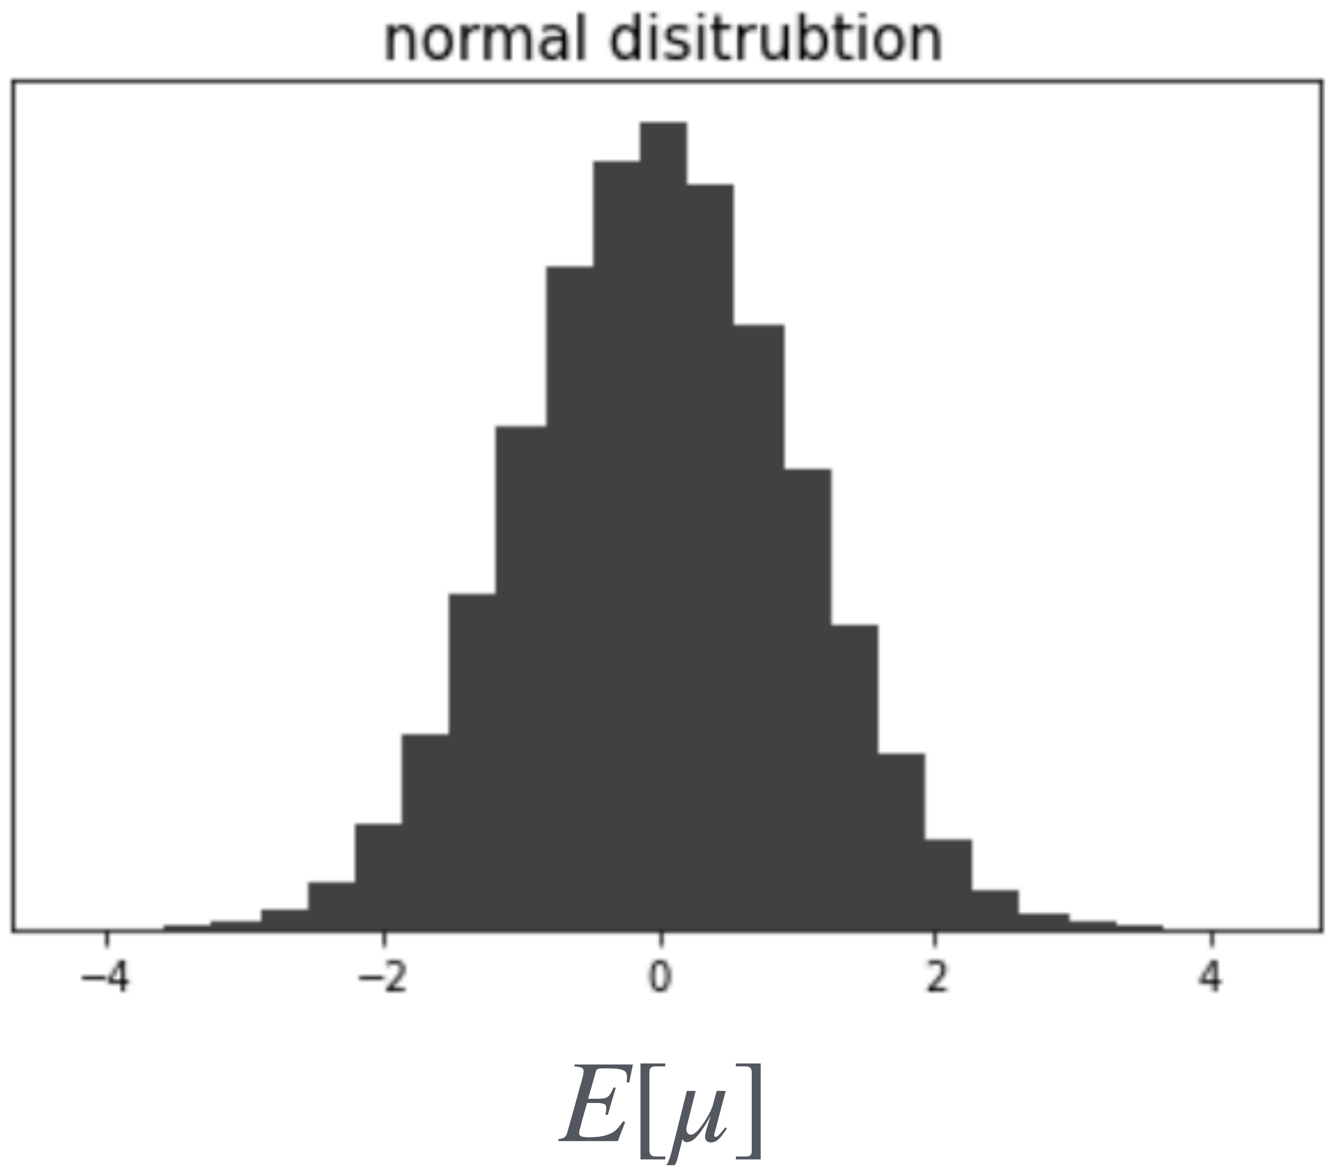

Review: A/B Test

- Goal: Accept or reject B
- Design: $N \geq (\frac{2.5\hat{\sigma}_\delta}{PS})^2$
- Measure: Replicate (reduce variance), Randomize (reduce bias)
- Analyze:

Criterion 1: $\delta > 1.6se \;\; (t > 1.6)$
Criterion 2: $\delta > PS$

Key Terms

- Optimism Bias
- Early Stopping
- Familywise error
- Bonferroni Correction

Review: LLN, CLT, A/B Testing

- As $N \to \infty$
  - LLN: $\mu \to E[y]$, estimate approaches “true” BM
  - CLT: $\mu \sim \mathcal{N}(E[y], VAR[y]/N)$, normal, narrows w/N
- **Design**: $N \geq \left(\frac{2.5\hat{\sigma}_\delta}{PS}\right)^2$
- **Measure**: Randomize, $\delta = \mu_B - \mu_A$, $se = \sigma_\delta/\sqrt{N}$
- **Analyze**: If $\delta > PS$ and $\frac{\delta}{se} \geq 1.64$, then accept B.

Review: False Positive Traps

- **Don’t stop early**, even if t-stat looks good
- **Beware multiple comparisons** in A/B/C/… tests
  - Use Bonferroni correction: p = 0.05 / (K-1)
  - Then accept if: $\mu > PS$ and $t = \dfrac{\delta}{se} \geq 1.64$

Review: Randomization

- A/B test: A=old ad, B=new ad
- Business metric is ad revenue/day
- A/B test design says N=10,000
- The A/B test has been running for three days, and you've collected 4,000 observations each of A and B so far. You calculate t from the 4,000 ind. meas:
  - $t = \frac{\mu}{se} = 8.3$ <== 8.3 is large. What does this tell you?

Review: Early Stopping

- $t = \dfrac{\mu}{se} = 8.3$ <== What does this tell you?
- Note: $se = \dfrac{\sigma_\delta}{\sqrt{4000}}$ > $se = \dfrac{\sigma_\delta}{\sqrt{10000}}$
- Therefore μ_B must be *much* larger than μ_A

Review: Early Stopping

- Stop now, capture extra revenue from B
  - I.e., reduce opportunity cost
- But, early stopping leads to false positives
- What could we do?

Key Terms

- Exploration
- Exploitation
- Arm
- Multi-armed bandit

Review: Randomization

- A/B test: 50/50 between A & B
  - N observations each
- Epsilon-greedy: 90% to best (so far) arm, ε=10% to other arms
  - Decay ε: $\varepsilon_n = \dfrac{kc(BM_0/PS)^2}{n}$ <== $\varepsilon_n \sim 1/n$
  - Stop when ε is small, $\varepsilon < \varepsilon_{\text{stop}}$

Key Terms

- Allocation
- Meta-parameters
- Thompson sampling
- Exploration vs. exploitation

Review: Thompson sampling

- Allocate observations to arms in proportion to the probability each arm is best
  - $p_{\text{arm}} \propto p_{\text{best}}$
- Stop when $\max\{p_{\text{best}}\} > 0.95$

Review: Predictor-in-controller

- Predictor: Estimates a target, ex., P{click} on an ad
- Controller: Uses predictions to make a decision / choose an action
  - Ex., "Of the 1000 ads available, show the one with the highest P{click}"

Review: Thompson sampling

- Allocate observations to arms in proportion to the probability each arm is best
  - $p_{\text{arm}} \propto p_{\text{best}}$
- Stop when $\max\{p_{\text{best}}\} > 0.95$

Review: A/B Test

- Goal: Accept or reject B
- Design: $$N \geq \left(\frac{2.5\hat{\sigma}_\delta}{PS}\right)^2$$
- Measure: Replicate (reduce variance), Randomize (reduce bias)
- Analyze:

Criterion 1: $\delta > 1.6se \ \ (t > 1.6)$

Criterion 2: $\delta > PS$

Review: Law of Large Numbers

- N observations, y_i, the business metric
- w/mean $\mu = \frac{\Sigma_i^N y_i}{N}$
- As $N \to \infty$, $\mu \to E[y]$
- IOW: Our measurement (μ) estimates the true, unobservable business metric

Review: A/B Test

- Goal: Accept or reject B
- Design: $$N \geq (\frac{2.5\hat{\sigma}_\delta}{PS})^2$$
- Measure: Replicate (reduce variance), Randomize (reduce bias)
- Analyze:

Criterion 1: $\delta > 1.6se \;\; (t > 1.6)$
Criterion 2: $\delta > PS$

Review: Thompson sampling

- Allocate observations to arms in proportion to the probability each arm is best
  - $p_{\text{arm}} \propto p_{\text{best}}$
- Stop when $\max\{p_{\text{best}}\} > 0.95$

Recap: Thompson Sampling

- Method:
  - **Draw once** from each arm's model of BM: $\{m_a \sim \mathcal{N}(\mu_a, se_a)\}$
  - Run arm $a* = \arg\max_a \{m_a\}$
  - Allocate $\propto p_a$
  - Stop when $p_{a*} > p_{stop}$

Review: Response Surface Methodology

- Surrogate: Model (regression)
  - Maps parameters, x, to measurements, y
- Analogy
  - $E[BM]$ is to observation y
  - as response function, $f(x)$, is to surrogate, $y(x)$

Key Terms

- Surrogate (again)
- Gaussian Process
- Gaussian Process Regression (GPR)
- Non-parametric
- Aleatoric (measurement) & epistemic (model) uncertainty

Review: Response Surface Methodology

- Surrogate: Model (regression)
  - Maps parameters, x, to measurements, y
- Analogy
  - $E[BM]$ is to observation y
  - as response function, $f(x)$, is to surrogate, $y(x)$

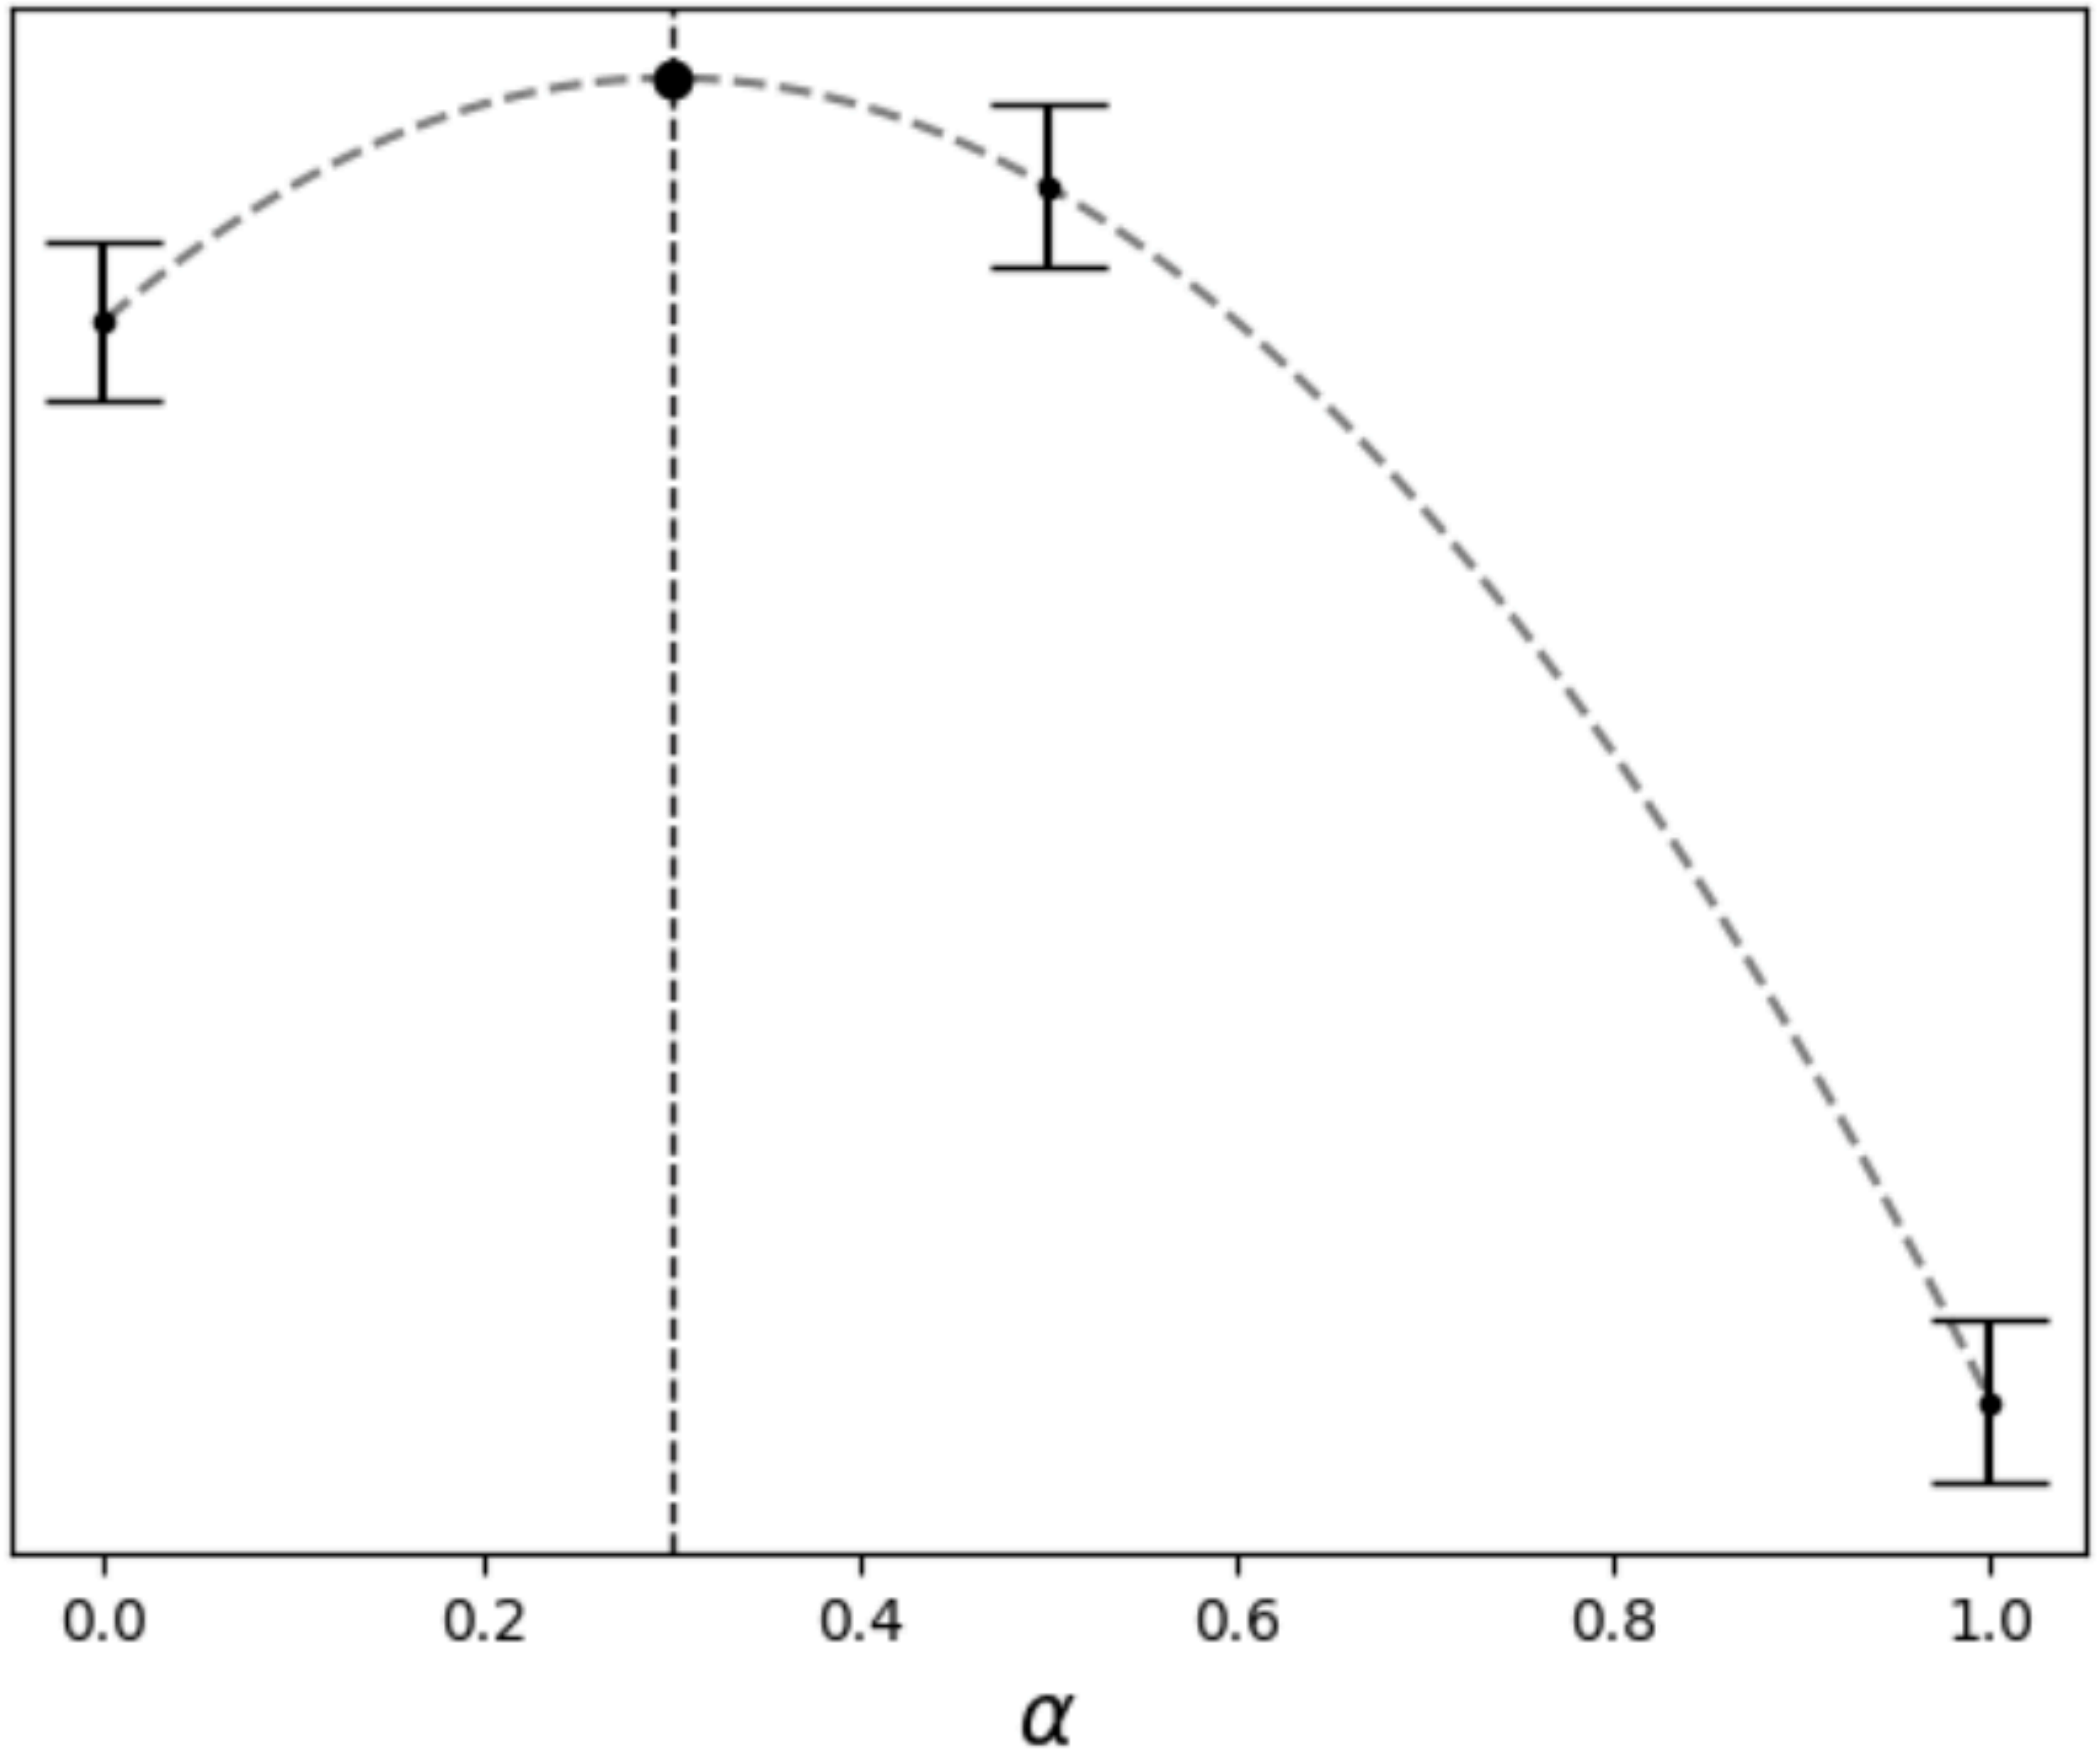

Review: Gaussian process regression

- $y(x) \sim \mathcal{N}(\mu(x), se^2(x))$
- Kernel function: $k(x, x') = e^{-d(x,x')^2/(2s^2)}$

$$\mu(x) = K_x^T(K_{xx} + se_0^2 I)^{-1}\mathbf{y}$$

$$se^2(x) = 1 - K_x^T(K_{xx} + se_0^2 I)^{-1}K_x$$

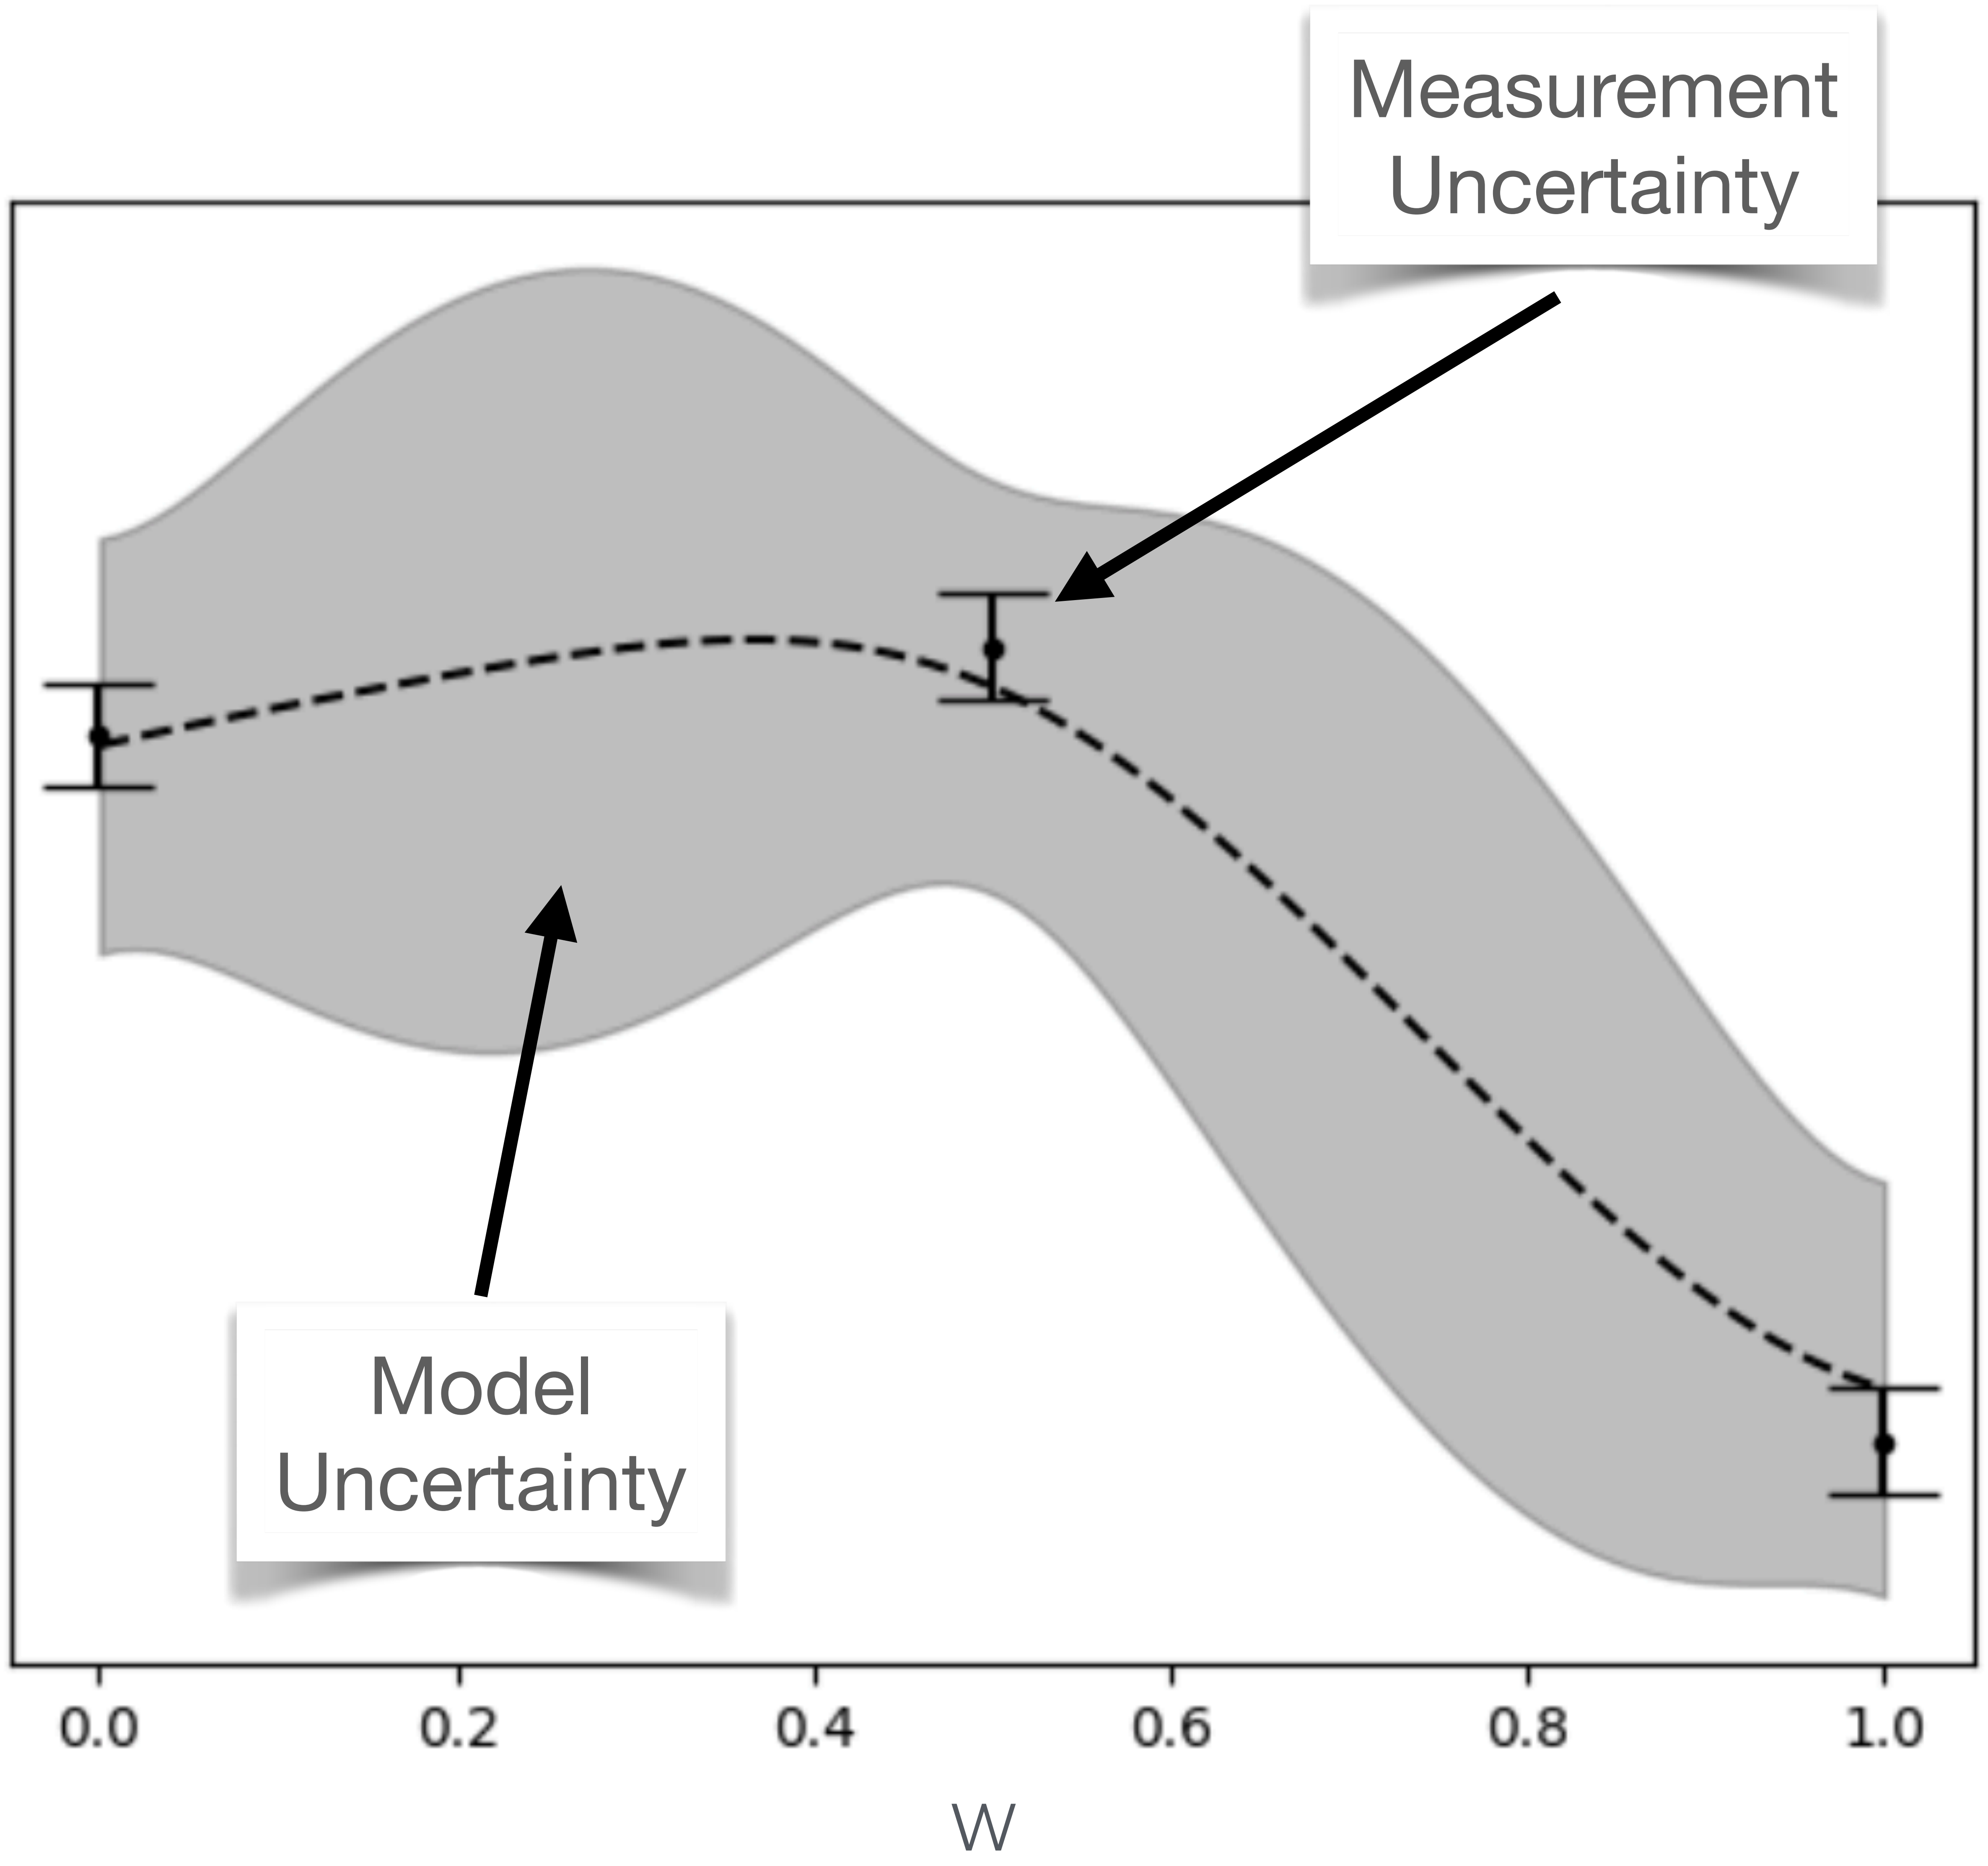

Review: Multi-armed bandits

- Model observations so far:
  - ε-greedy: μ_a
  - TS: $y_a \sim \mathcal{N}(\mu_a, se_a^2)$

$$\mu_a = \sum_i y_{a,i}/N$$
$$se_a = \left[\sqrt{\sum_i (y_{a,i} - \mu_a)^2/N}\right]/\sqrt{N}$$

- Select arm:
  - ε-greedy: 90% $\arg\max_a \mu_{a'}$ 10% random
  - TS:
    - Draw $m_a \sim \mathcal{N}(\mu_a, se_a^2)$
    - Arm $\arg\max_a m_a$

BO: Connections

1. A/B testing: Take a low-se, low-bias measurement

2. RSM: Build a surrogate

3. MAB: Balance exploration & exploitation in design

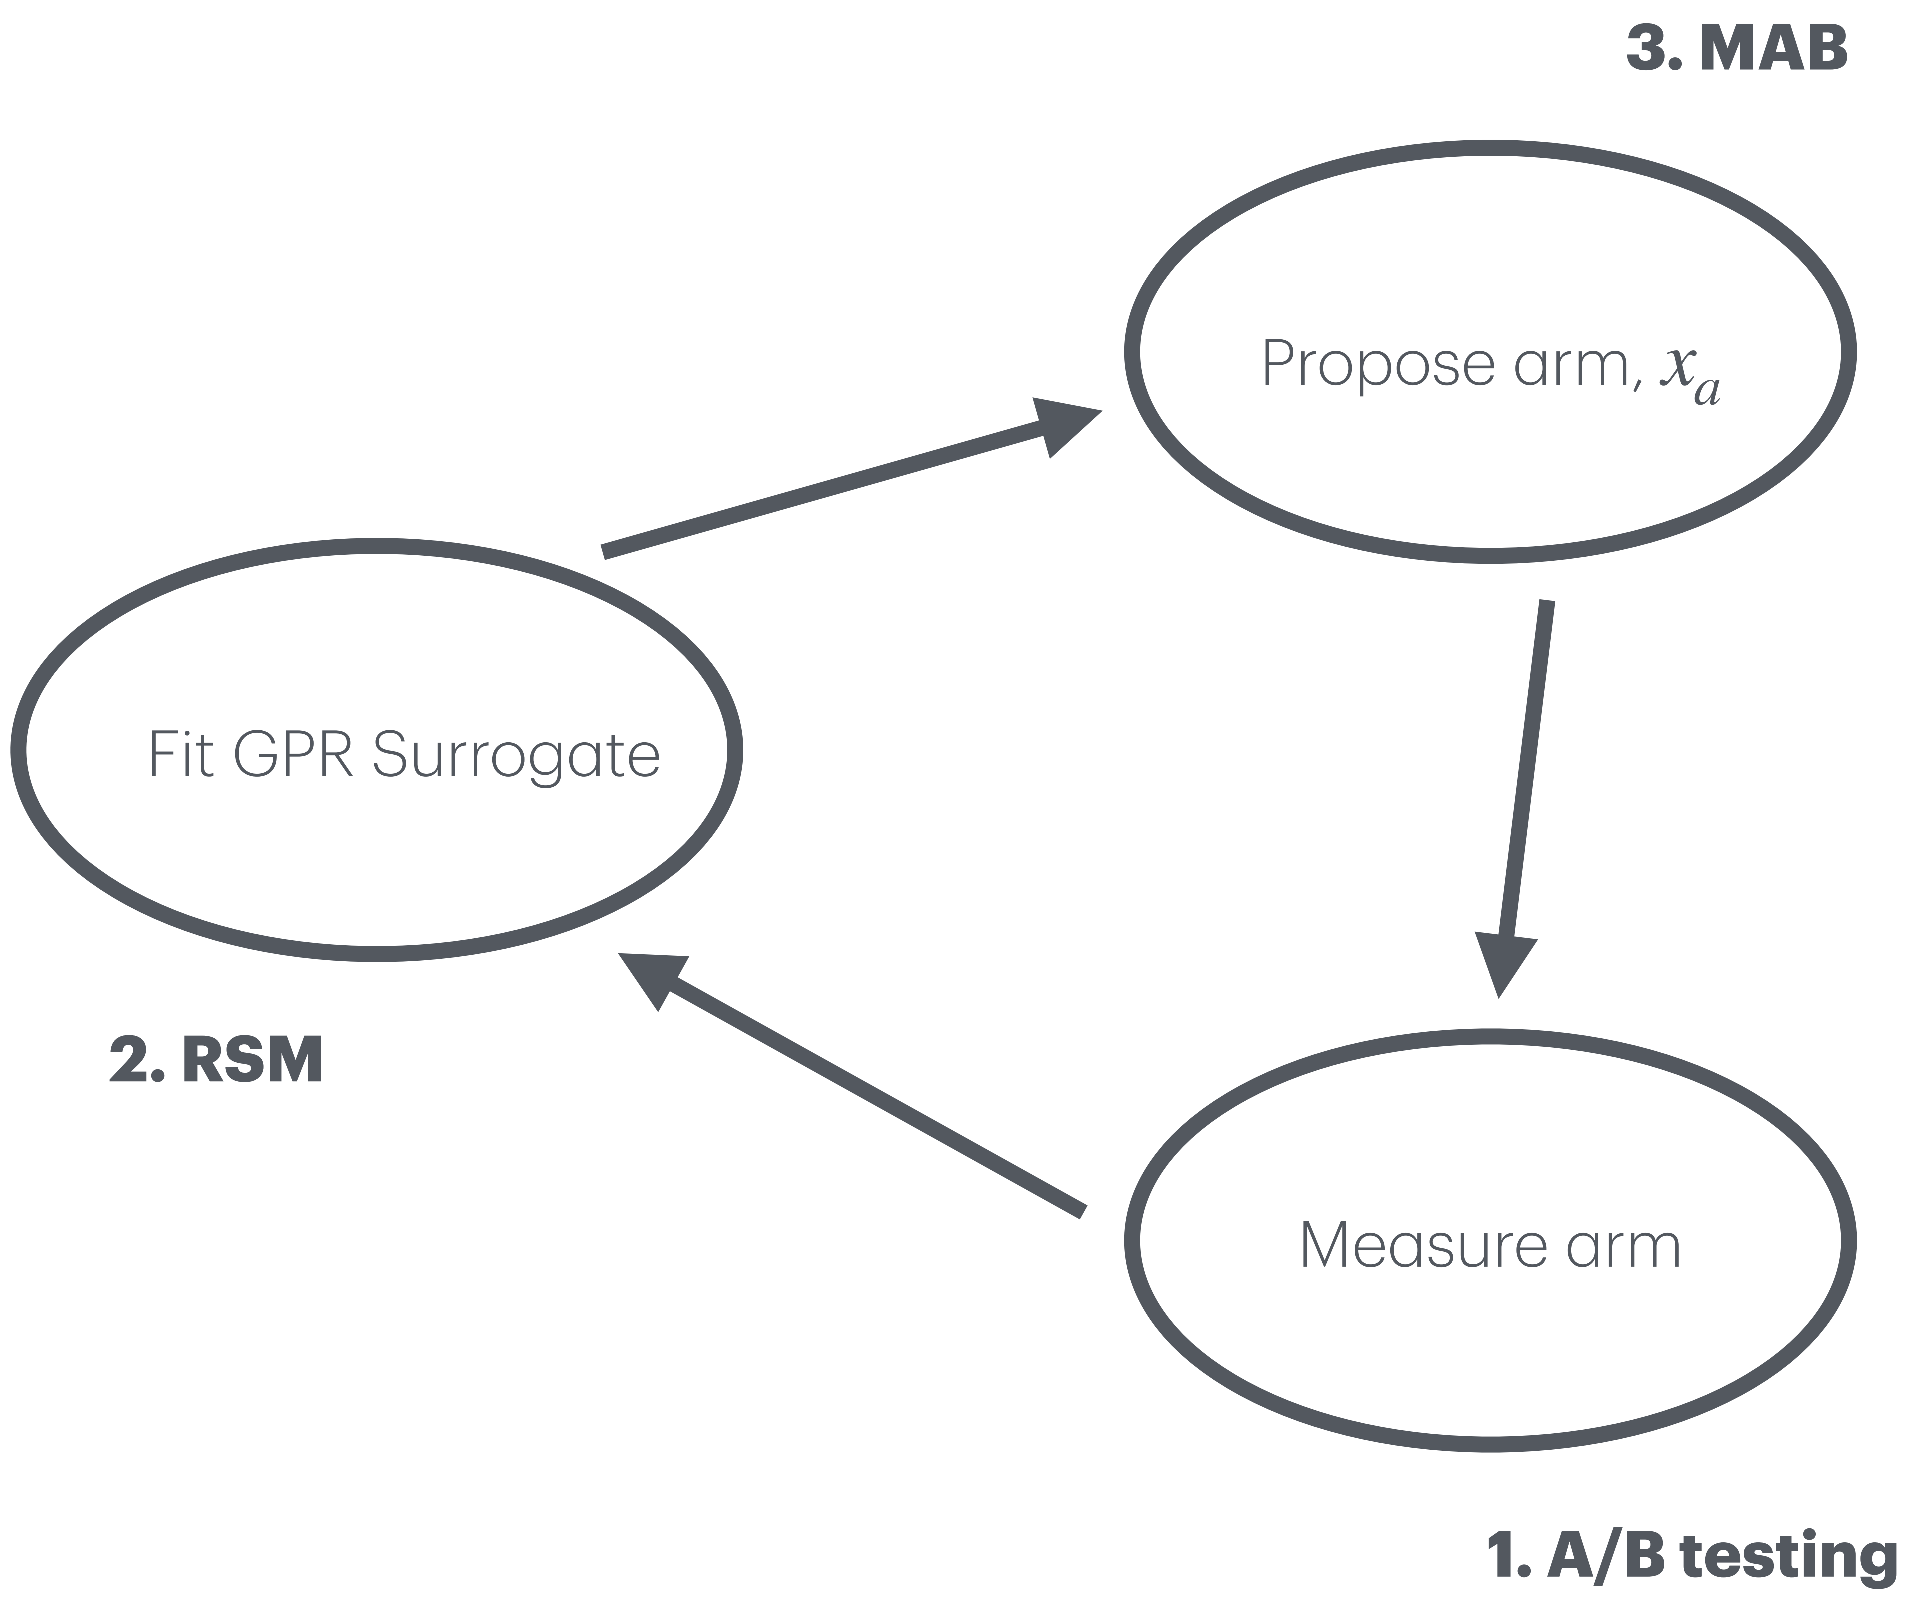

Good Luck!