The Bandit Problem: Are You Winning Enough?

A/B Testing: Problem?

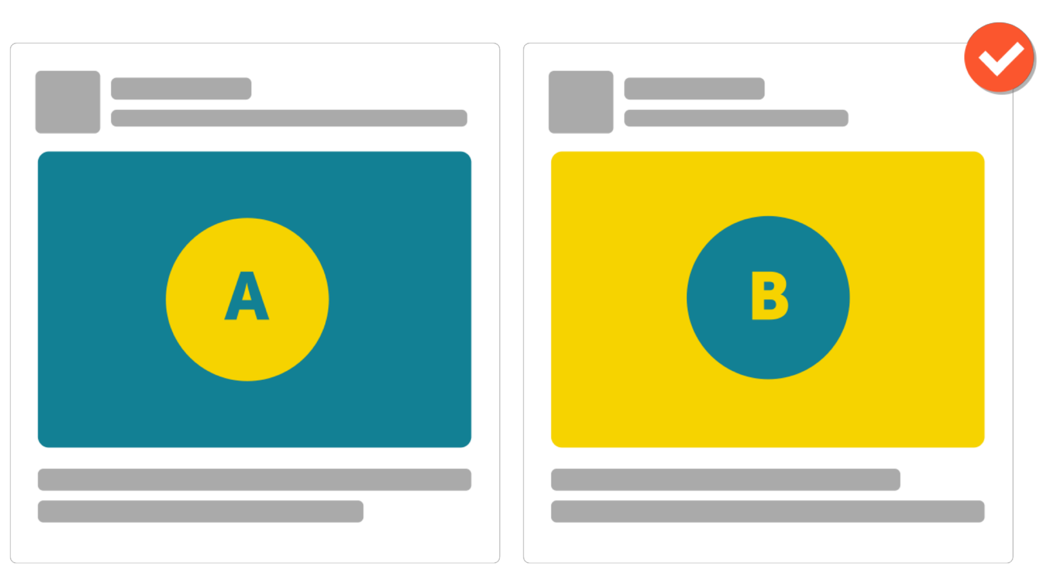

- ❑ Decide the level of statistical significance
- ❑ Send visitors to all the landing pages simultaneously, randomly
- ❑ All the pages get equal share of visitors
- ❑ Run the experiment for a long enough time
- ❑ Don't interrupt it; don't change anything while it runs
- ❑ Wait patiently, where it's advisable not to look at the intermediate results
- ❑ Throws away money on ineffective pages throughout the experiment

Proposed Solution: Multi-armed Bandits

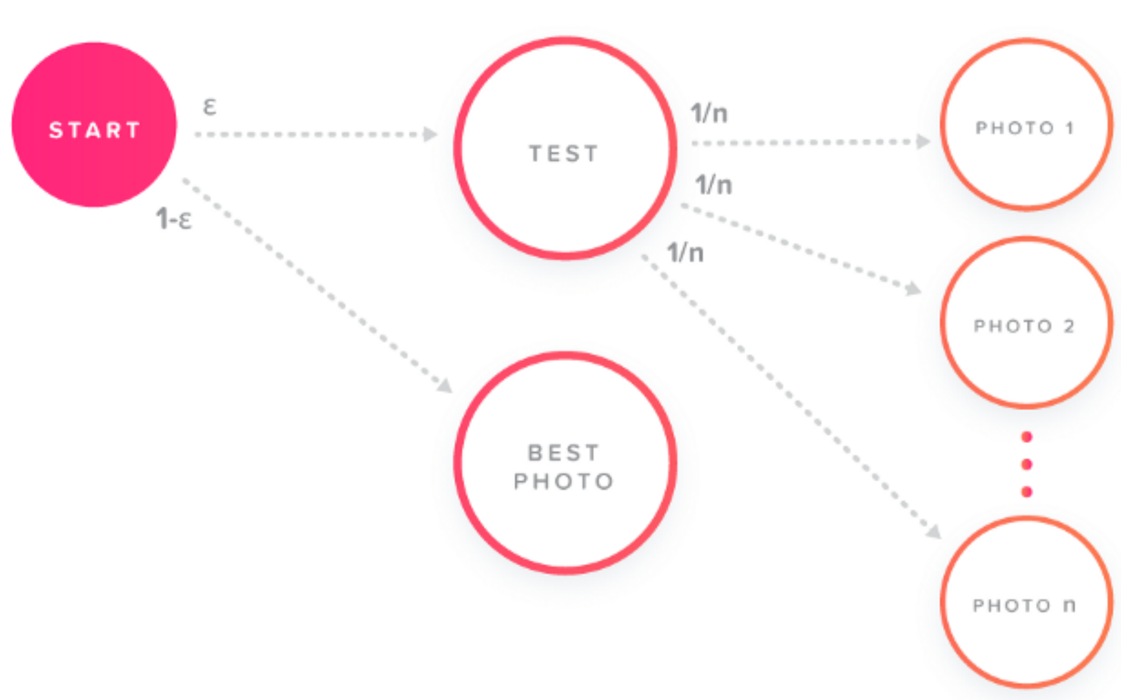

- ❑ Let's imagine a bandit machine where humans could choose to pull either the left or the right arm
- ❑ Each given a random payoff with the distribution of payoffs for each arm unknown
- ❑ Will you keep pulling the left arm and ignore the right or would you pull the right arm a few more times hoping that it might get better?
- ❑ Balance between exploration and exploitation
- ❑ Multi-armed bandit algorithms are more flexible and generally more efficient than A/B testing

Research & Simulations

- ❑ Multiple experiments conducted and documented to estimate performance, cumulative rewards, and probability of selecting the best arm, under various circumstances
- ❑ Monte-Carlo simulations performed to estimate the behavior in hypothetical worlds

- **Epsilon-Greedy, Annealing (naïve)**
- **Softmax, Annealing (structured exploration)**
- **Upper Confidence Bound (explicitly curious)**
- **Thompson Sampling (probability matching)**
- **Linear-UCB (contextual)**

Rewards Maximized!

- ❑ For instance, if we have two arms (variants) with 20% and 10% reward (conversion) probabilities, and 10,000 trials
- ❑ Then A/B testing would generate a reward for 1,500 trials
- ❑ While multi-armed bandits would generate a reward for ~ 1900 trials with epsilon-greedy strategy (ε = 0.2)
- ❑ Where the maximum achievable reward would be 2,000, if only we knew the best variant before the experiment

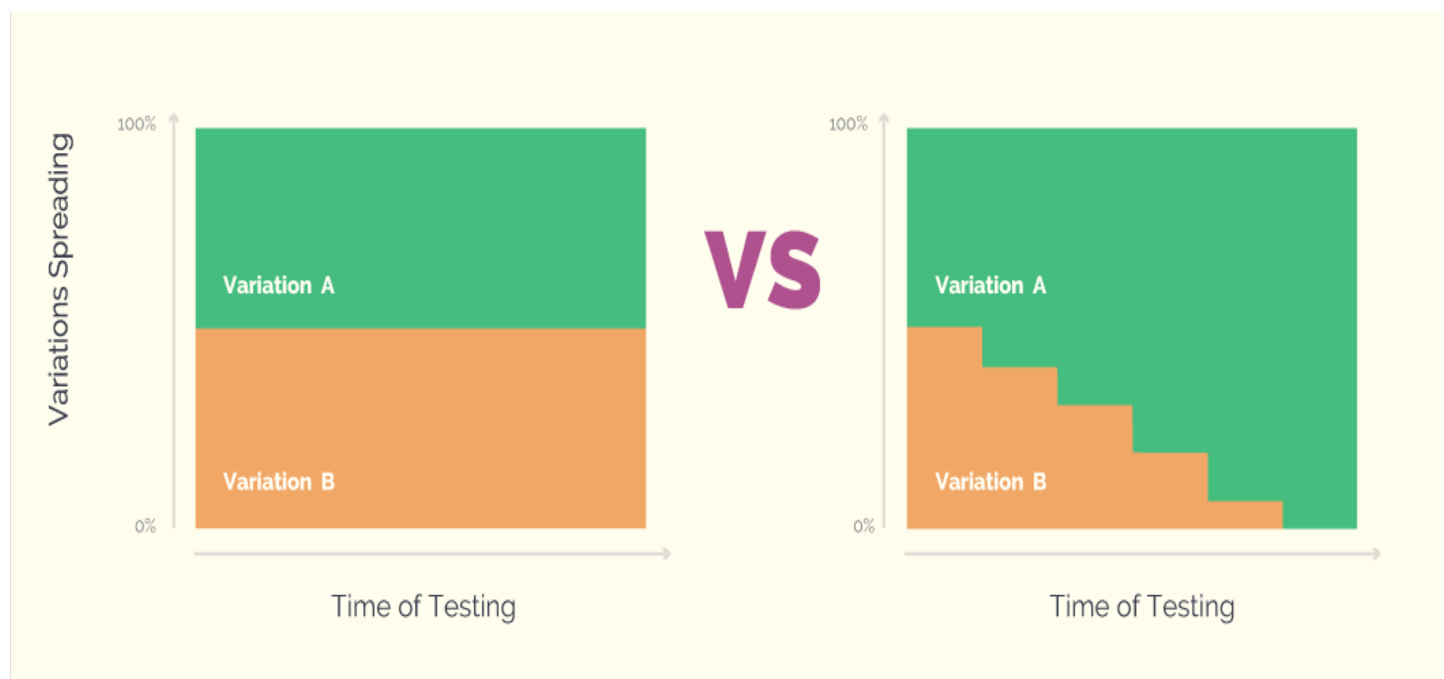

smartsheet

Harkar Talwar
Manisha Vyas
Prateek Tripathi
Sayali Chakradeo

Information School
UNIVERSITY of WASHINGTON