Booking.com

Subpopulation-aware experimentation platform

Christina Katsimerou
Principal Machine Learning Scientist

CWI, 25/05/2022

A

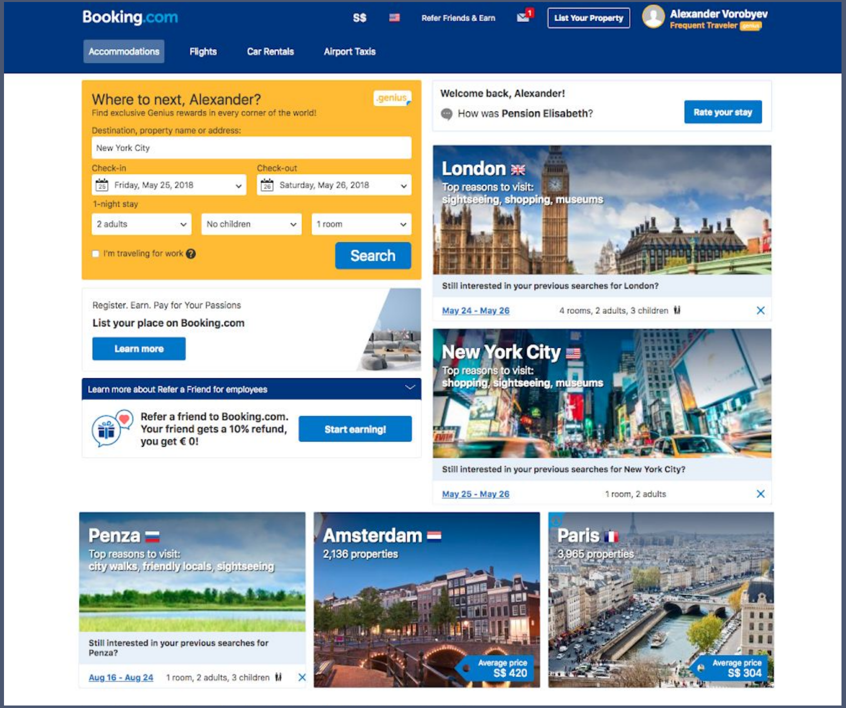
Booking.com
S$
Refer Friends & Earn
List Your Property
Alexander Vorobyev
Frequent Traveler
Accommodations
Flights
Car Rentals
Airport Taxis
Where to next, Alexander?
Find exclusive Genius rewards in every corner of the world!
Destination, property name or address:
New York City
Check-in
Friday, May 25, 2018
Check-out
Saturday, May 26, 2018
1-night stay
2 adults
No children
1 room
I'm traveling for work
Search
Welcome back, Alexander!
How was Pension Elisabeth?
Rate your stay
London
Top reasons to visit:
sightseeing, shopping, museums
Still interested in your previous searches for London?
May 24 - May 26
4 rooms, 2 adults, 3 children
Register. Earn. Pay for Your Passions
List your place on Booking.com
Learn more
New York City
Top reasons to visit:
shopping, sightseeing, museums
Learn more about Refer a Friend for employees
Refer a friend to Booking.com.
Your friend gets a 10% refund,
you get € 0!
Start earning!
Still interested in your previous searches for New York City?
May 25 - May 26
1 room, 2 adults
Penza
Top reasons to visit:
city walks, friendly locals, sightseeing
Still interested in your previous searches for Penza?
Aug 16 - Aug 24
1 room, 2 adults, 3 children
Amsterdam
2,136 properties
Average price
S$ 420
Paris
3,965 properties
Average price
S$ 304

B

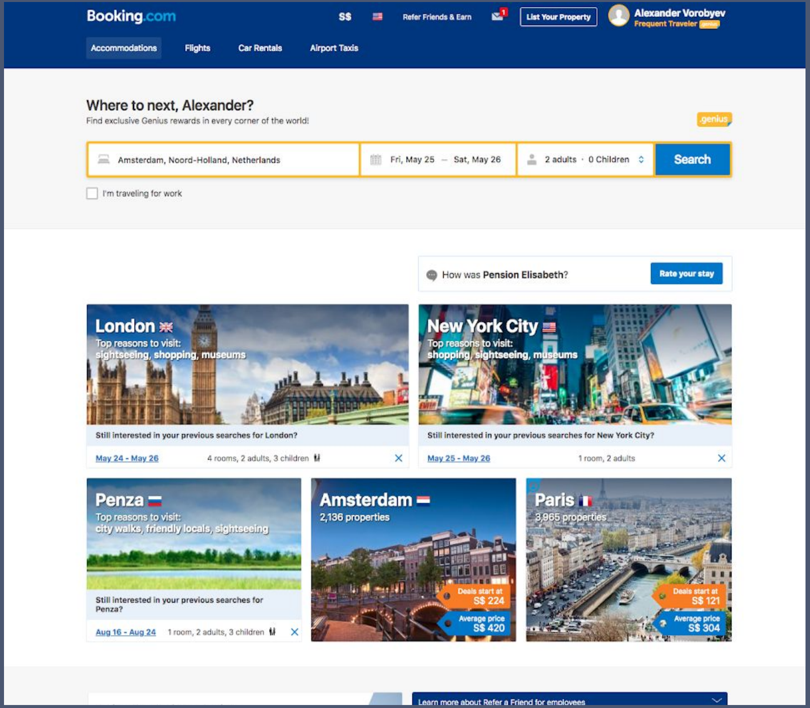
Booking.com
S$
Refer Friends & Earn
List Your Property
Alexander Vorobyev
Frequent Traveler
Accommodations
Flights
Car Rentals
Airport Taxis
Where to next, Alexander?
Find exclusive Genius rewards in every corner of the world!
genius
Amsterdam, Noord-Holland, Netherlands
Fri, May 25 — Sat, May 26
2 adults · 0 Children
Search
I'm traveling for work
How was Pension Elisabeth?
Rate your stay
London
Top reasons to visit:
sightseeing, shopping, museums
Still interested in your previous searches for London?
May 24 - May 26
4 rooms, 2 adults, 3 children
New York City
Top reasons to visit:
shopping, sightseeing, museums
Still interested in your previous searches for New York City?
May 25 - May 26
1 room, 2 adults
Penza
Top reasons to visit:
city walks, friendly locals, sightseeing
Still interested in your previous searches for Penza?
Aug 16 - Aug 24
1 room, 2 adults, 3 children
Amsterdam
2,136 properties
Deals start at
S$ 224
Average price
S$ 420
Paris
3,965 properties
Deals start at
S$ 121
Average price
S$ 304
Learn more about Refer a Friend for employees

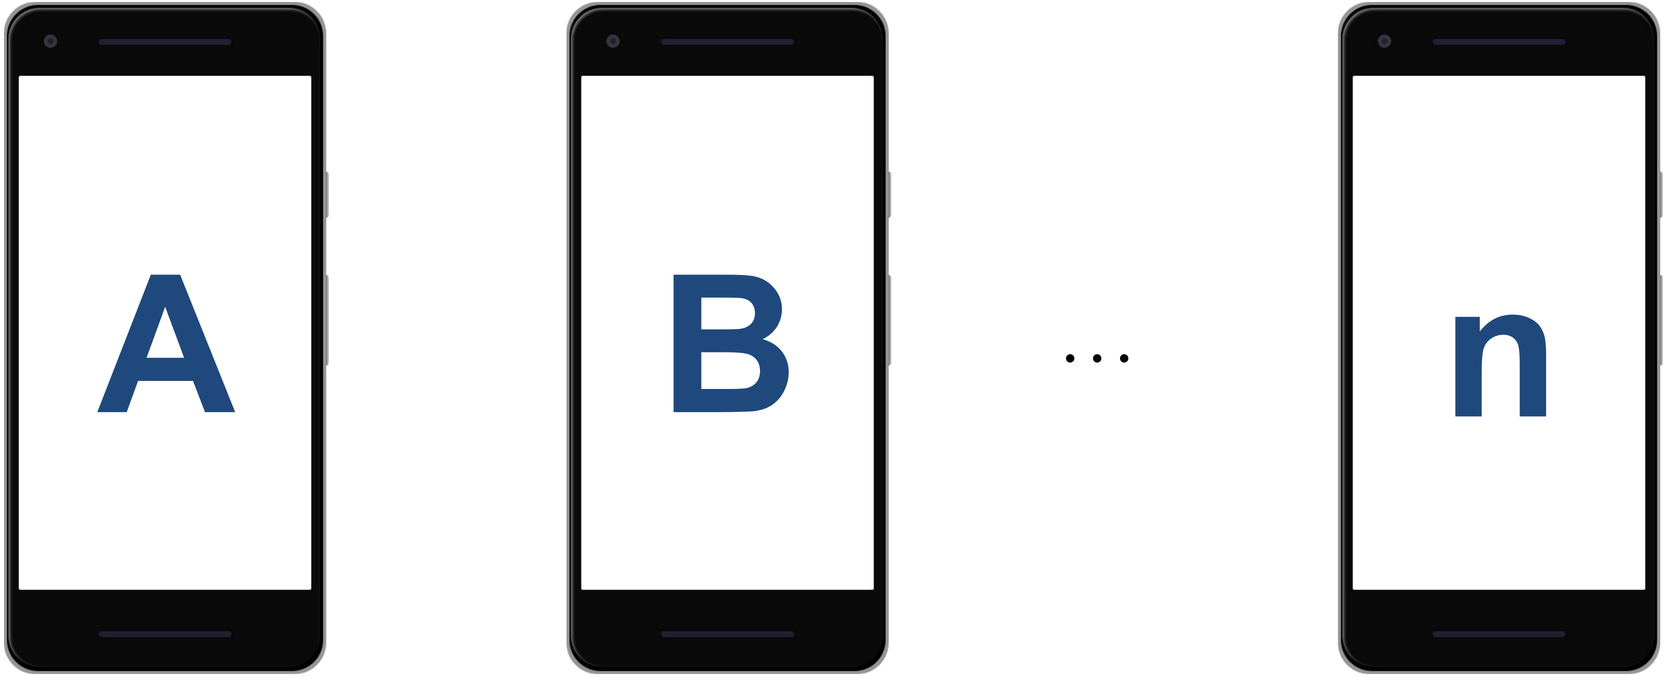

A/B/n testing

- Decision making
- Learning & Inspiration

Currently

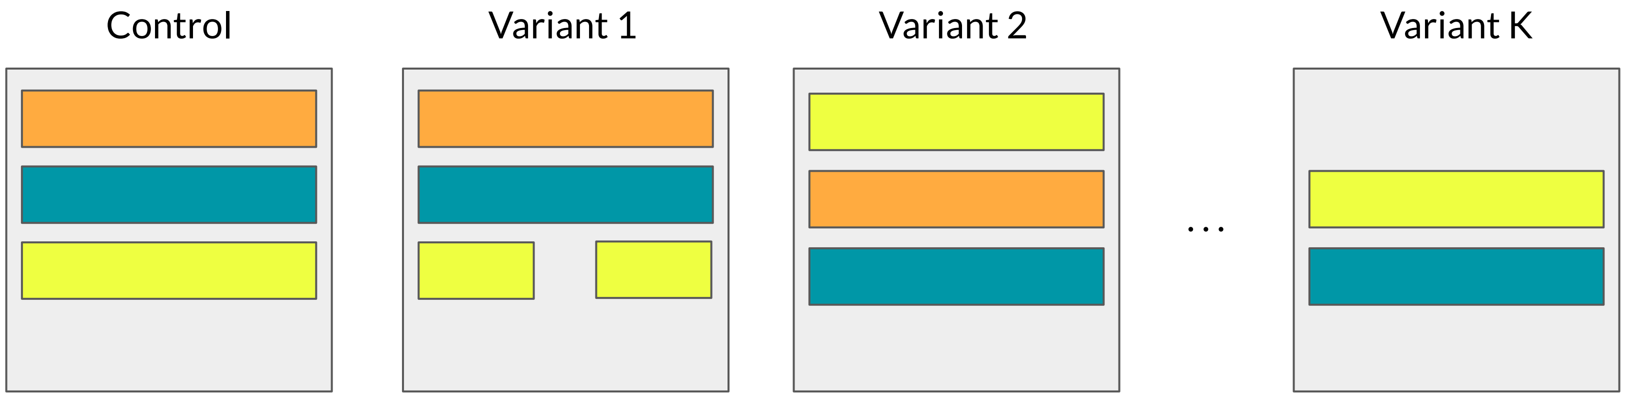

Currently

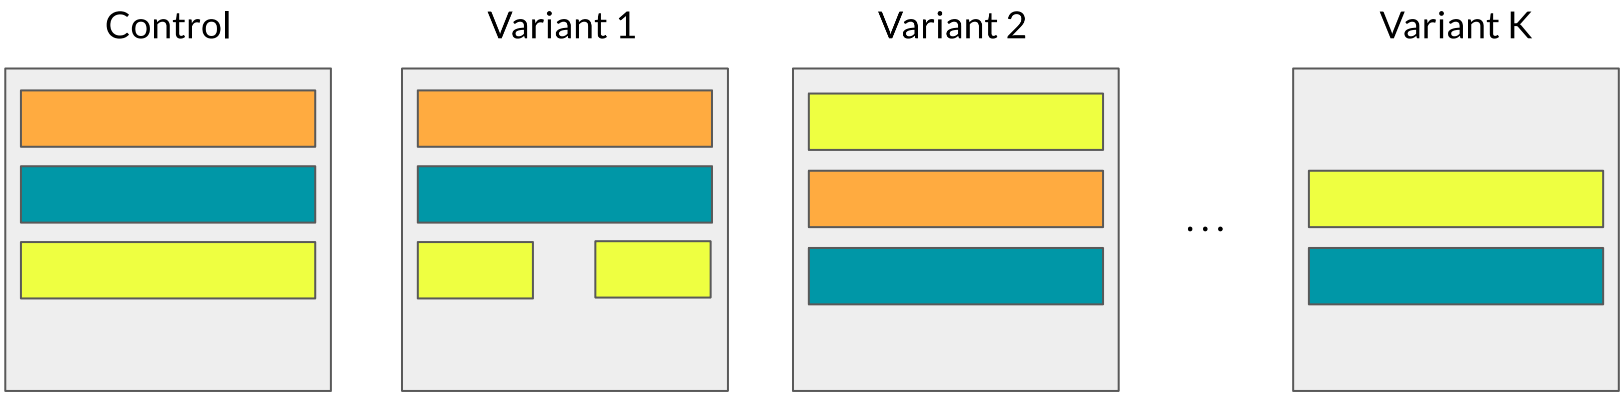

$$H_{0,k} : \mu_0 \leq \mu_k, k \in \{1, ..., K\}$$

Currently

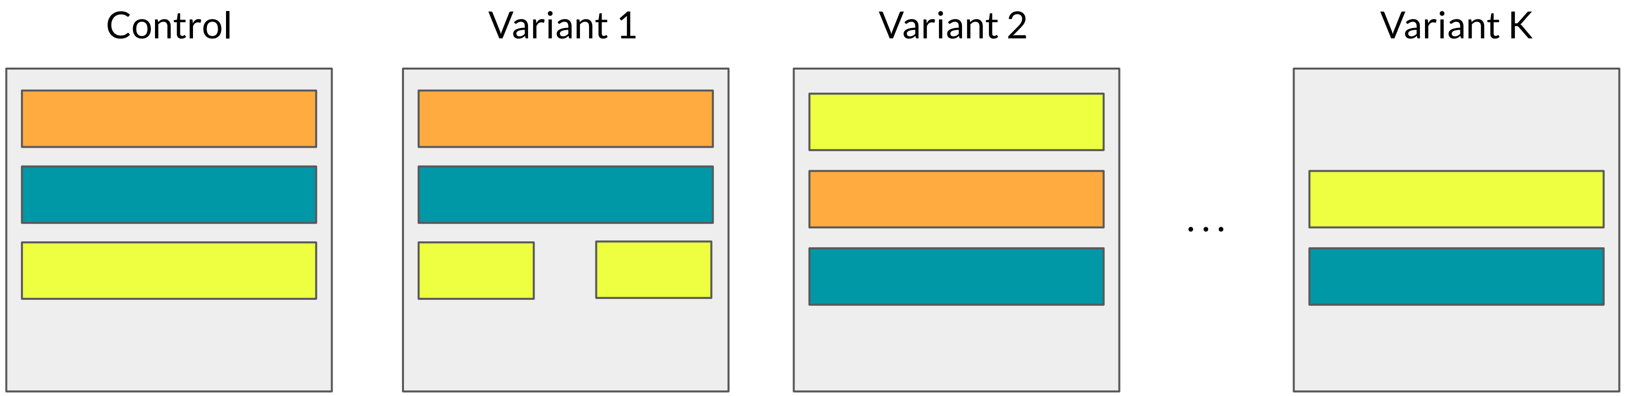

$$H_{0,k} : \mu_0 \leq \mu_k, k \in \{1, ..., K\}$$

For a fixed risk δ:

- Determine minimum desired detectable effect
- Compute sample size N

Currently

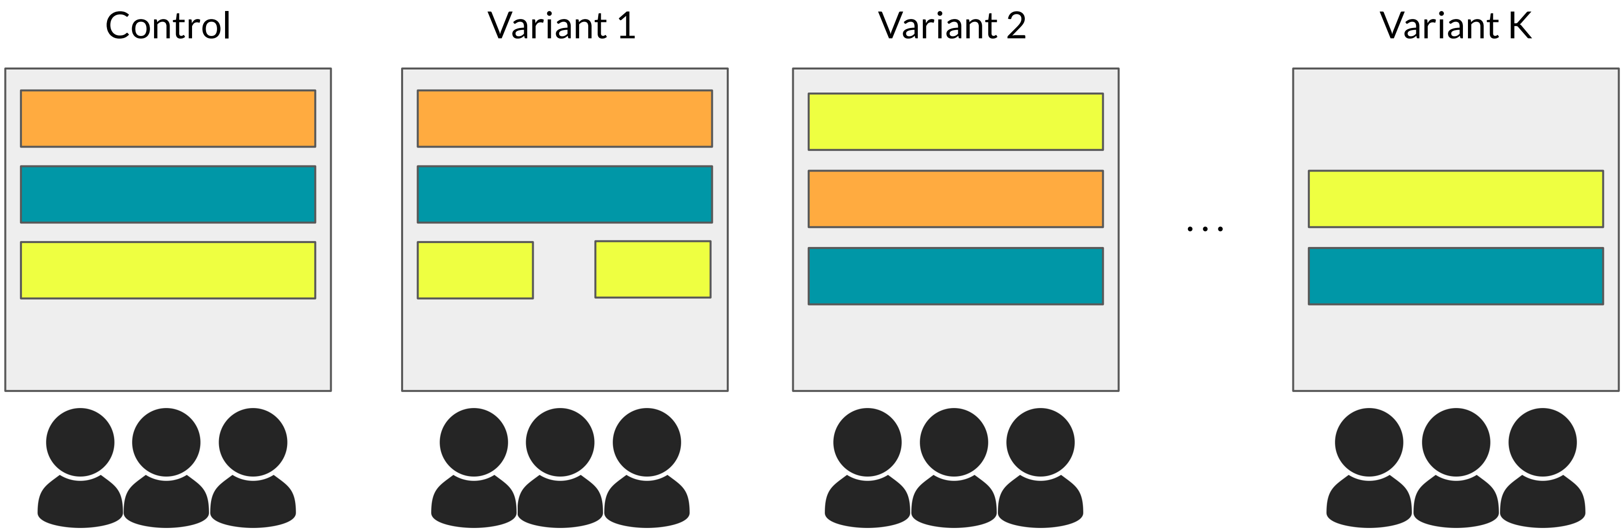

$$H_{0,k} : \mu_0 \leq \mu_k, k \in \{1, ..., K\}$$

For a fixed risk δ:

- Determine minimum desired detectable effect
- Compute sample size N
- Distribute traffic uniformly among variants

Currently

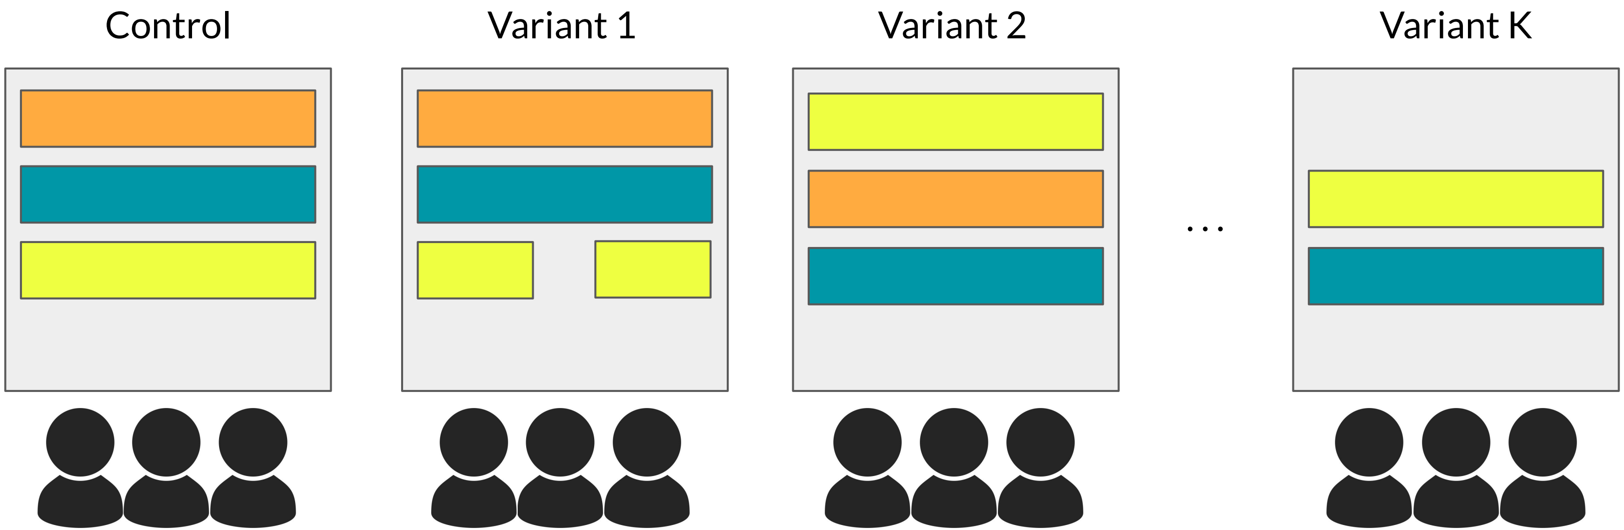

$$H_{0,k} : \mu_0 \leq \mu_k, k \in \{1, ..., K\}$$

For a fixed risk δ:

- Determine minimum desired detectable effect
- Compute sample size N
- Distribute traffic uniformly among variants
- At sample size N compute p-values p_k per hypothesis

Currently

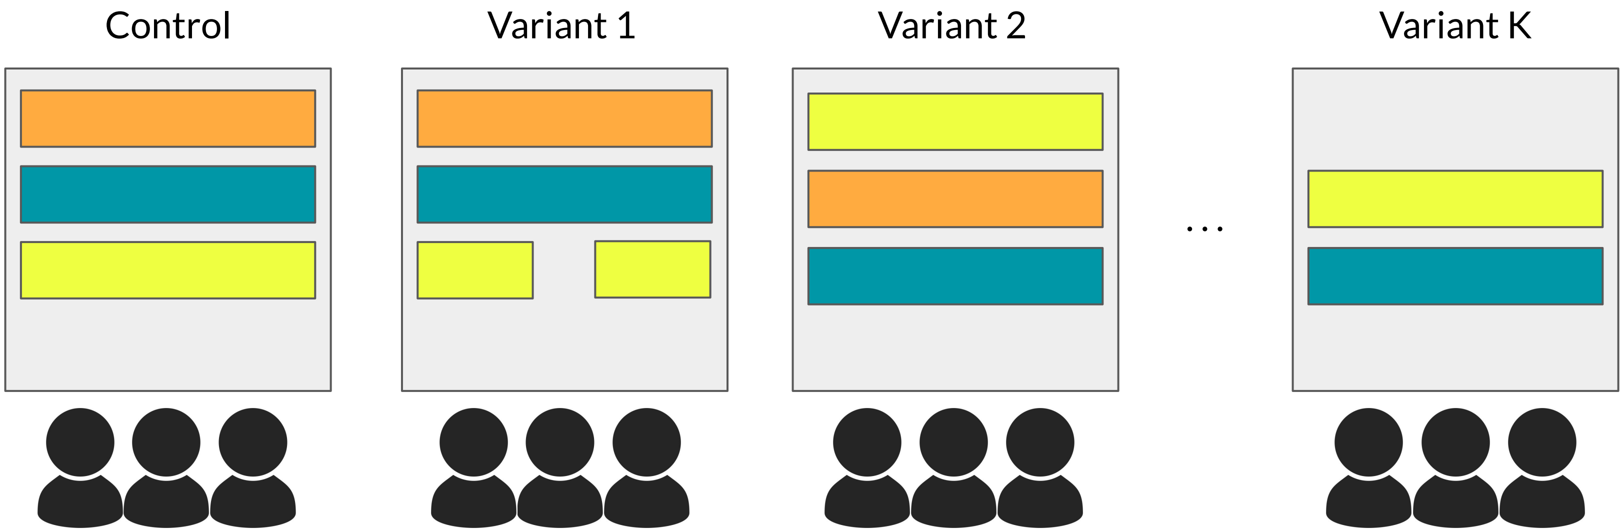

$$H_{0,k} : \mu_0 \leq \mu_k, k \in \{1, ..., K\}$$

For a fixed risk δ:

- Determine minimum desired detectable effect
- Compute sample size N
- Distribute traffic uniformly among variants
- At sample size N compute p-values p_k per hypothesis
- Reject $H_{0,k}$ if $p_k \leq \frac{\delta}{K}$

Problem #1: Uniform allocation

Control | Variant 1 | Variant 2 | … | Variant K

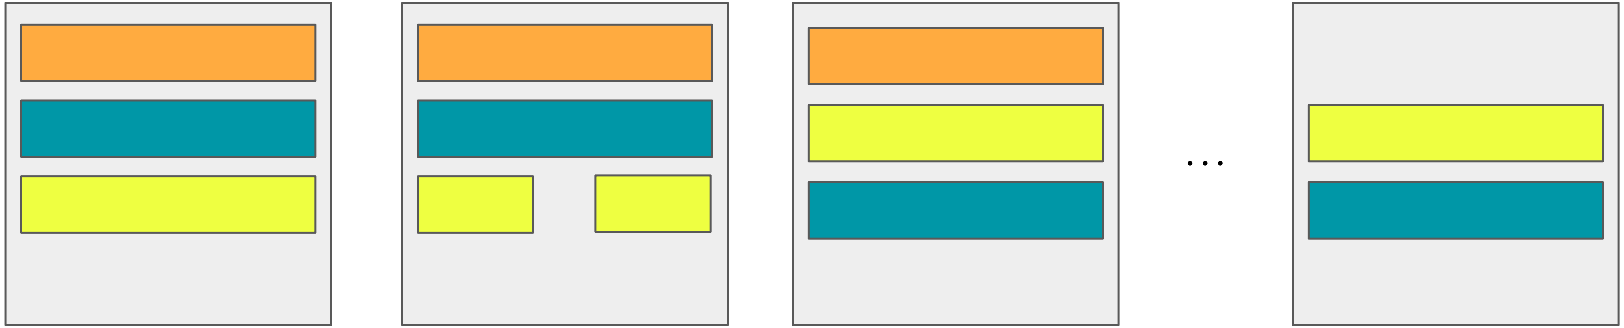

$\mu_0 = 0.1$ $\mu_1 = 0.1001$ $\mu_2 = 0.11$ $\mu_K = 0.08$

Problem #1: Uniform allocation

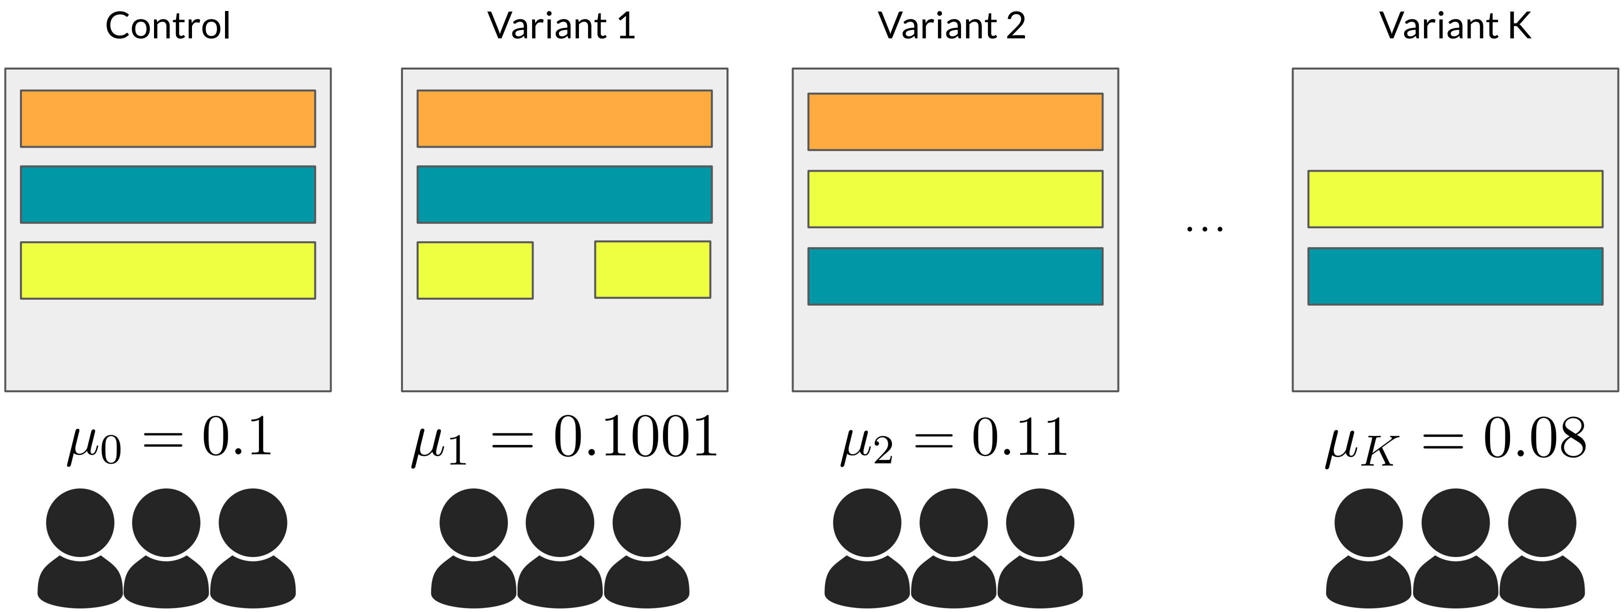

Problem #1: Uniform allocation

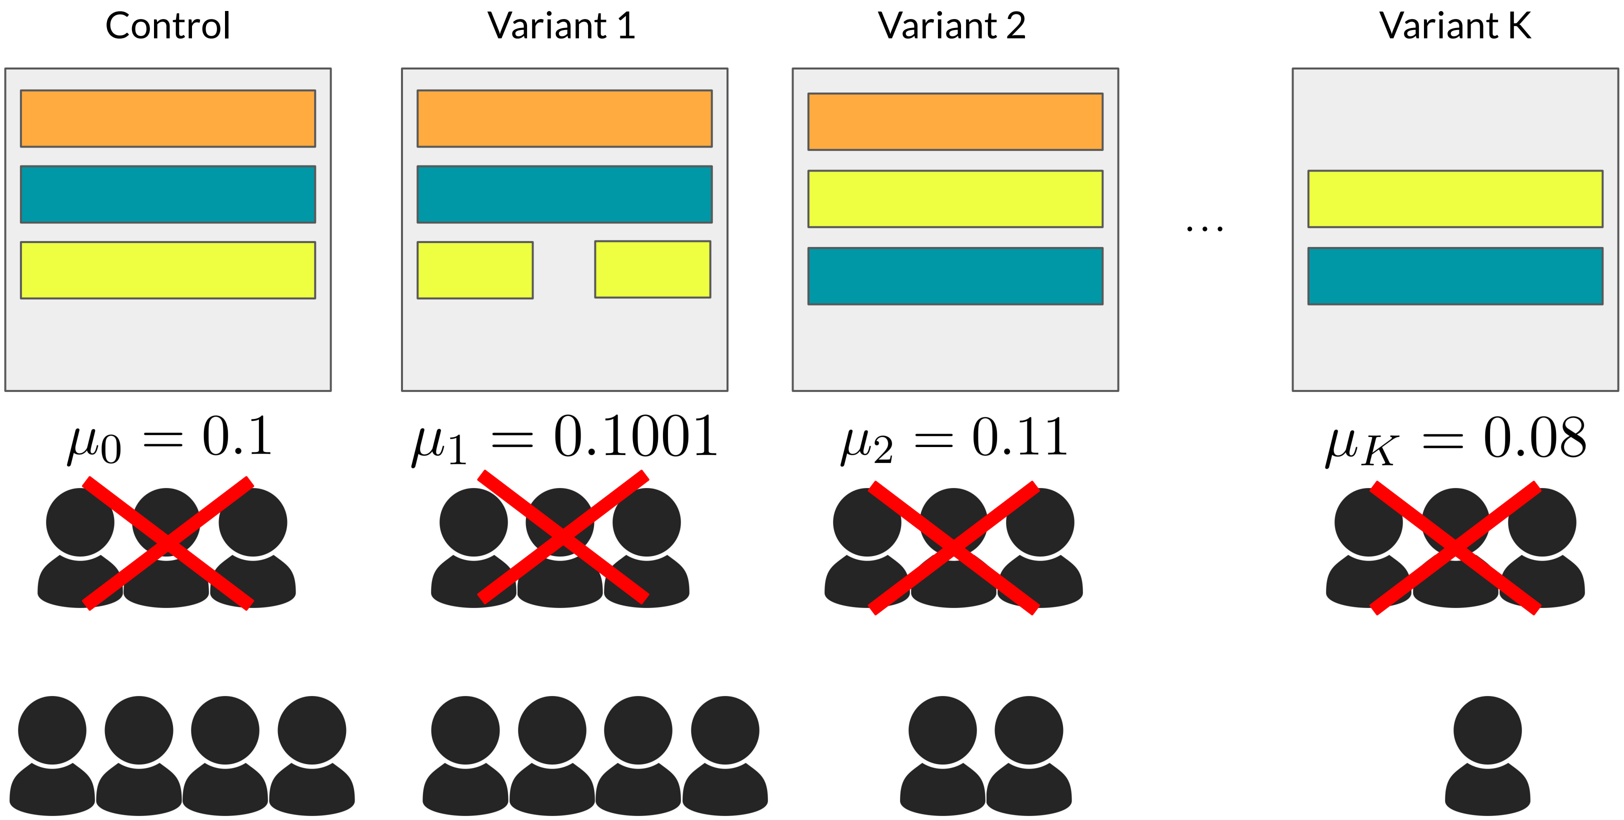

Problem #2: No monitoring

Experimenters want to **monitor continuously** their experiments to:

➔ Detect effects as quickly as possible
➔ Give up soon when there is no effect

Problem #2: No monitoring

Experimenters want to monitor continuously their experiments to:

➔ Detect effects as quickly as possible

➔ Give up soon when there is no effect

The platform does not allow **data dependent stopping times**

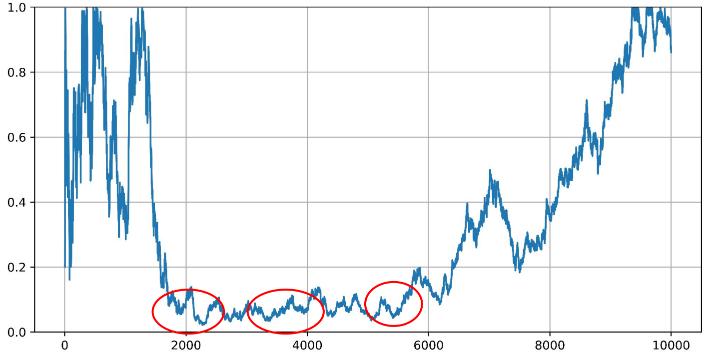

$p(t)$ under $H_0 : \mu_1 = \mu_2$

Problem #2: No monitoring

Experimenters want to monitor continuously their experiments to:

- ➔ Detect effects as quickly as possible
- ➔ Give up soon when there is no effect

The platform does not allow **data dependent stopping times**

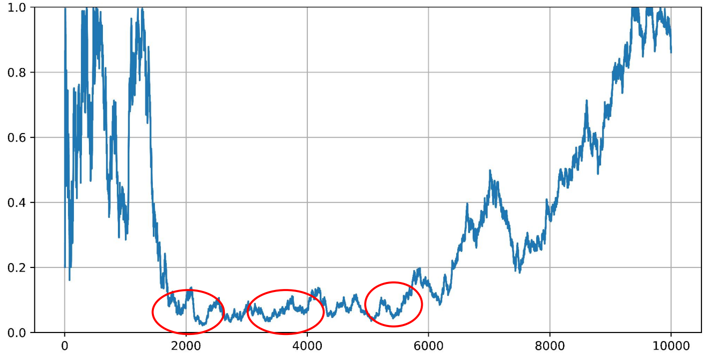

$p(t)$ under $H_0 : \mu_1 = \mu_2$

Monitoring without properly correcting for it inflates the false positive rate.

New formulation

We aim to identify the subset of **all variants that are better than the control,** while:

- optimizing adaptively the stopping time of the A/B/n experiment
- being ***δ*-PAC:** the probability of stopping and returning a wrong answer must be **≤ *δ***
- optimizing adaptively the allocation of variants to users

Seasonality

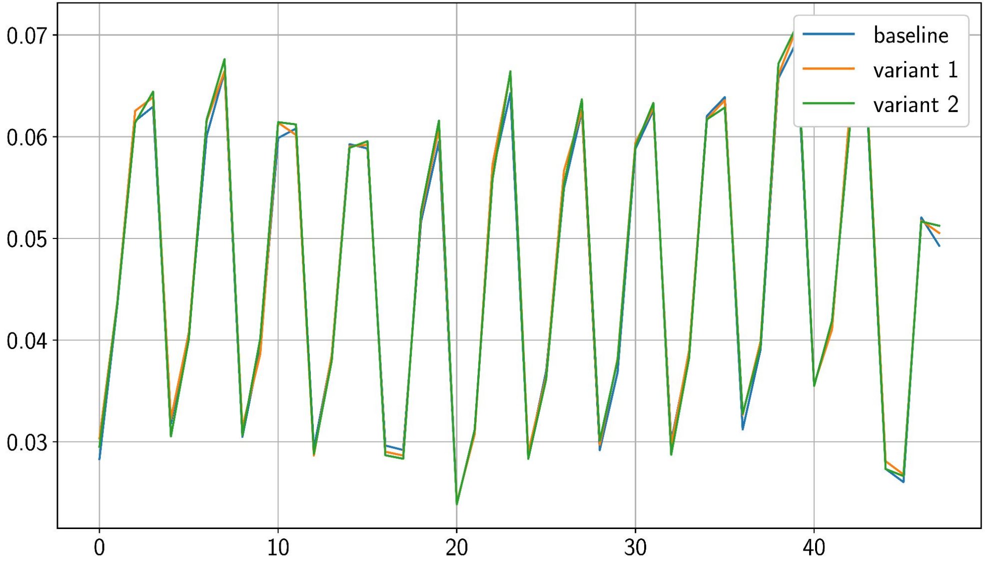

Click-through rate per 6 hours for 12 days

Seasonality

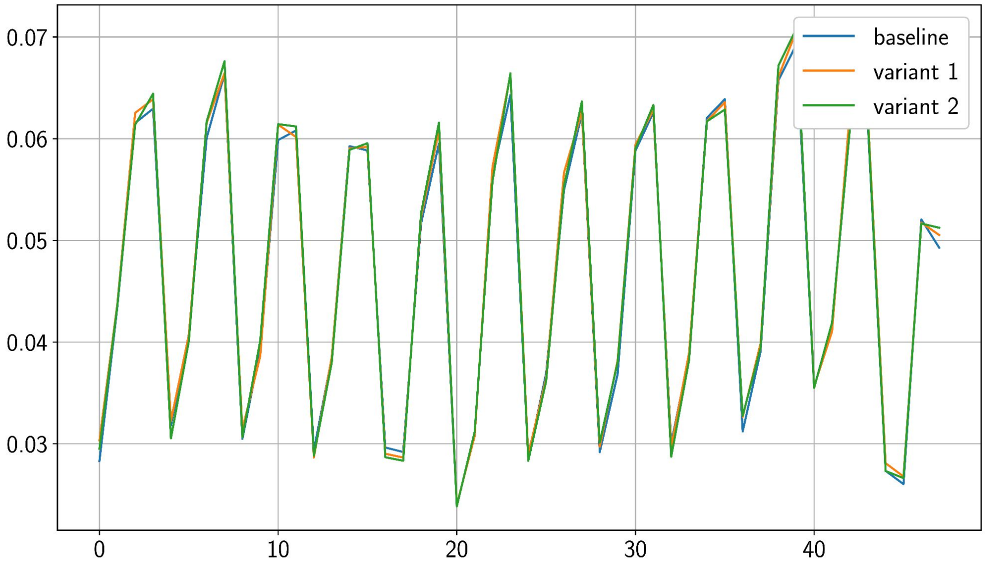

Click-through rate per 6 hours for 12 days

Seasons → subpopulations

With ***J*** subpopulations, each variant is a mixture of ***J*** distributions.

The value of the variant a is:

$$\mu_a = \sum_{i=1}^{J} \alpha_i \mu_{a,i}$$

Interaction with subpopulations

- ❑ Pick variant
- ❑ Pick subpopulation

(a) Active

- ❑ See subpopulation
- ❑ Pick variant

(b) Proportional

- ❑ Pick variant
- ❑ Don't see subpopulation

(c) Oblivious

Interaction with subpopulations

- ❏ Pick variant
- ❏ Pick subpopulation

(a) Active

- ❏ See subpopulation
- ❏ Pick variant

(b) Proportional

- ❏ Pick variant
- ❏ Don’t see subpopulation

(c) Oblivious

← more powerful ... less powerful

Model

A bandit with $K+1$ arms and J subpopulations is

- a $(K+1) \times J$ matrix $\boldsymbol{\mu}$ of reward distributions
- distribution $\boldsymbol{\alpha}$ of J subpopulations

Model

A bandit with $K+1$ arms and J subpopulations is

- a $(K+1) \times J$ matrix $\boldsymbol{\mu}$ of reward distributions
- distribution $\boldsymbol{\alpha}$ of J subpopulations

The correct answer for $(\boldsymbol{\mu}, \boldsymbol{\alpha})$ is

$$\mathcal{S}(\boldsymbol{\mu}) = \left\{ k \in \{1, \dots, K\} \text{ s.t } \sum_{j=1}^{J} \alpha_j \mu_{k,j} > \sum_{j=1}^{J} \alpha_j \mu_{0,j} \right\} .$$

Complexity

For any sampling strategy, the expected number of rounds $\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]$ is bounded from below.

$\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]$:

- depends on $\boldsymbol{\mu}$
- gets minimized when we sample with the optimal sampling proportions $w^\star(\boldsymbol{\mu})$, which depend on the mode of interaction with the subpopulations

Example

subpopulations

$$\boldsymbol{\mu} = \begin{bmatrix} 0.15 & 0.05 \\ 0.16 & 0.06 \\ 0.17 & 0.07 \end{bmatrix}$$

arms

$$\boldsymbol{\alpha} = [0.5, 0.5]$$

Example

subpopulations

$$\boldsymbol{\mu} = \begin{bmatrix} 0.15 & 0.05 \\ 0.16 & 0.06 \\ 0.17 & 0.07 \end{bmatrix}$$

arms

$$\boldsymbol{\alpha} = [0.5, 0.5]$$

	i=1	i=2
0	0.288	0.182
1	0.283	0.178
2	0.042	0.027
	0.613	**0.387**

$$\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]_{\text{active}} \geq 22370$$

Example

subpopulations

$$\boldsymbol{\mu} = \begin{bmatrix} 0.15 & 0.05 \\ 0.16 & 0.06 \\ 0.17 & 0.07 \end{bmatrix} \text{arms} \qquad \boldsymbol{\alpha} = [0.5, 0.5]$$

	i=1	i=2
0	0.288	0.182
1	0.283	0.178
2	0.042	0.027
	0.613	**0.387**

$$\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]_{\text{active}} \geq 22370$$

	i=1	i=2
0	0.236	0.231
1	0.230	0.234
2	0.034	0.035
	0.5	**0.5**

$$\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]_{\text{proportional}} \geq 23516$$

Example

subpopulations

$$\boldsymbol{\mu} = \begin{bmatrix} 0.15 & 0.05 \\ 0.16 & 0.06 \\ 0.17 & 0.07 \end{bmatrix} \quad \boldsymbol{\alpha} = [0.5, 0.5]$$

$$\boldsymbol{\mu} = \begin{bmatrix} 0.10 \\ 0.11 \\ 0.12 \end{bmatrix}$$

	i=1	i=2
0	0.288	0.182
1	0.283	0.178
2	0.042	0.027
	0.613	**0.387**

$$\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]_{\text{active}} \geq 22370$$

	i=1	i=2
0	0.236	0.231
1	0.230	0.234
2	0.034	0.035
	0.5	**0.5**

$$\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]_{\text{proportional}} \geq 23516$$

0	0.47
1	0.46
2	0.07
	1

$$\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]_{\text{oblivious}} \geq 24158$$

Track & Stop Algorithm

For $t \geq 1$:

- Sampling rule: given the current estimates $\hat{\boldsymbol{\mu}}(t)$

  ❶ estimate the target weights $\mathbf{w}_t(\hat{\boldsymbol{\mu}})$

  ❷ pick arm s.t. $\frac{N_{a,t}}{t} \to \mathbf{w}_t(\hat{\boldsymbol{\mu}})$

Track & Stop Algorithm

For $t \geq 1$:

- Sampling rule: given the current estimates $\hat{\boldsymbol{\mu}}(t)$

  1. estimate the target weights $\mathbf{w}_t(\hat{\boldsymbol{\mu}})$
  2. pick arm s.t. $\frac{N_{a,t}}{t} \to \mathbf{w}_t(\hat{\boldsymbol{\mu}})$

- Recommendation:
  $\mathcal{S}(\hat{\boldsymbol{\mu}}_t) = \{a \in \{1, \ldots, K\} : \hat{\mu}_a(t) > \hat{\mu}_0(t)\}$ with risk assessment $\hat{\delta}_t = \min\{\delta \in (0,1) | \Lambda(t) \geq \beta(t, \delta)\}$

A/B/n Booking.com

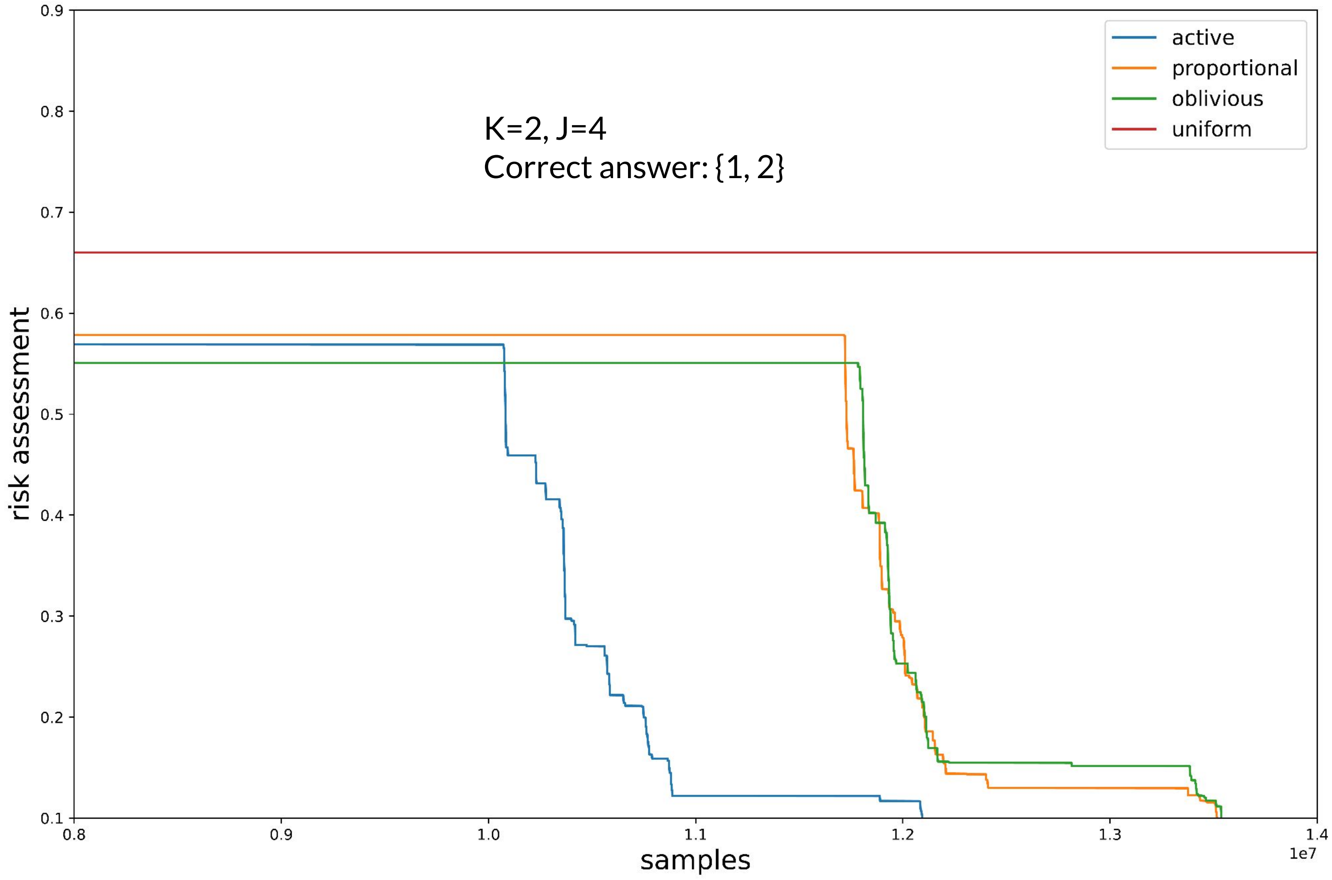

Conclusion

The sub-population aware experiment platform:

- Is more efficient allowing adaptive sampling and stopping in the presence of seasonality
- Experimenters can monitor continuously without errors

Conclusion

The sub-population aware experiment platform:

- Is more efficient allowing adaptive sampling and stopping in the presence of seasonality
- Experimenters can monitor continuously without errors

We have already run successfully A/B/n experiments, which we validated on the non adaptive platform

Conclusion

The sub-population aware experiment platform:

- Is more efficient allowing adaptive sampling and stopping in the presence of seasonality
- Experimenters can monitor continuously without errors

We have already run successfully A/B/n experiments which we validated on the non adaptive platform

Still work to be done on:

- Delayed rewards
- epsilon-better
- …

Thank you!

christina.katsimerou@booking.com

paper

booking.ai

https://jobs.booking.com/careers

References

- Garivier, A. and E. Kaufmann (2016). “Optimal best arm identification with fixed confidence”. In: Conference on Learning Theory. PMLR, pp. 998–1027.
- Russac, Y., C. Katsimerou, D. Bohle, O. Cappé, A. Garivier, and W. M. Koolen (Dec. 2021). “A/B/n Testing with Control in the Presence of Subpopulations”. In: Advances in Neural Information Processing Systems (NeurIPS) 34.

Appendix

Complexity

Theorem 1

For any strategy, the expected number of rounds for the ABC-S problem satisfies

$$\liminf_{\delta \to 0} \frac{\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]}{\ln(1/\delta)} \geq T^\star(\boldsymbol{\mu}) \tag{1}$$

where

$$T^\star(\boldsymbol{\mu})^{-1} = \sup_{\mathbf{w} \in \mathcal{C}} \min_{b \neq 0} \inf_{\lambda \in \mathcal{L} : \lambda_0 < \lambda_b} \sum_{a \in \{0,b\}} \sum_{i=1}^{J} w_{a,i} d(\mu_{a,i}, \lambda_{a,i}) \tag{2}$$

- A min-max optimization problem
- Optimal weights $\mathbf{w}^* =$ target relative frequencies of draws

Track & Stop Algorithm

For $t \geq 1$:

- Sampling rule: given the current estimates
  1. estimate the target weights $\mathbf{w}_t(\hat{\boldsymbol{\mu}})$
  2. pick arm s.t. $\frac{N_{a,t}}{t} \to \mathbf{w}_t(\hat{\boldsymbol{\mu}})$
- Recommendation:
  1. Stop at $\tau_\delta = \inf\{t \in \mathbb{N} | \text{test statistic} \geq \text{threshold}\}$
  2. Recommend $\mathcal{S}(\hat{\boldsymbol{\mu}}_t) = \{a \in \{1, \dots, K\} : \hat{\mu}_a(\tau_\delta) > \hat{\mu}_0(\tau_\delta)\}$

GLRT

Generalised Likelihood Ratio Statistic for exponential family bandit models:

$$\Lambda(t) = \min_{b \neq 0} \inf_{\lambda \in \mathcal{L}: \lambda_0 = \lambda_b} \sum_{a \in \{0,b\}} \sum_{i=1}^{J} N_{a,i}(t) d(\hat{\mu}_{a,i}(t), \lambda_{a,i})$$

Example

$$\boldsymbol{\mu} = \begin{bmatrix} 0.15 & 0.05 \\ 0.16 & 0.06 \end{bmatrix}, \boldsymbol{\alpha} = \begin{bmatrix} 0.5 & 0.5 \end{bmatrix}$$

T_{active}=16013

	i=1	i=2
0	0.309	0.194
1	0.305	0.192
	0.614	0.386

$T_{proportional}$=16840

	i=1	i=2
0	0.252	0.246
1	0.248	0.254
	0.5	0.5

$T_{agnostic}$=16841

	i=1	i=2
0	0.251	0.251
1	0.249	0.249
	0.5	0.5