Bayesian Adaptive Choice Experiments

Marshall Drake
EEOC

Fernando Payró
UAB, BSE & CODE

Neil Thakral
Brown University

Linh T. Tô
Boston University

ESEM

Data collection in economics

Old-generation data collection: Often captures quantities now more readily available in administrative datasets

- E.g., income, employment, education, benefit take-up, etc.
- Other inputs to economic decisions are harder to measure
  - E.g., preferences, beliefs, expectations, risk attitudes, etc.

New-generation data collection can be more

- *customizable, controllable, interactive*
- promises to do much more than capturing coarse measures of these important economic inputs

Getting high-quality data on preferences

Reliably estimating valuations generally requires eliciting a series of responses to proposed choice sets (Discrete Choice Experiment):

- Direct questions about valuations? anchoring effects, highly uncertain, greater potential for hypothetical bias, ...

- Multiple price lists? order effects, reference point effects, range bias, complexity, difficulty accommodating multiple dimensions ...

- Eliciting multiple pairwise choices helps respondents think through tradeoffs more clearly (Feld, Nagy, and Osman, 2020)

Discrete choice experiments in economics

Broad range of applications in recent years:

- Labor economics: Preferences for job attributes (Eriksson & Kristensen 2014, Mas & Pallais 2017, Wiswall & Zafar 2018, Mas & Pallais 2019,Johnston 2021, Feld et al. 2022, Folke & Rickne 2022, Maestas et al. 2023), for discrimination (Kuhn & Osaki 2023)
- Other examples in many fields:
  - Public economics
  - Health economics
  - Development economics
  - Environmental economics
  - Urban, real estate, transportation
  - Measuring time and risk preferences in psychology and economics

Discrete choice experiments (DCE)

- DCEs often require asking each respondent multiple questions to narrow down their preferences. But trade-offs from having more questions...
  - (-) Less thoughtful responses (survey fatigue)
  - (-) Non-response bias (higher attrition rates)
  - (-) Higher cost of data collection (can be expensive)
- A shortcut researchers often have to make:
  - **Pool choice data** across respondents to estimate "average preferences" of an entire group
  - A less often discussed point: this shortcut can lead to biased estimates (will revisit)

Bayesian Adaptive Choice Experiments (BACE)

We propose the use of **dynamic adaptive choice experiments** to replace existing static choice experiments widely used to elicit preferences which **simultaneously**:

- Efficiently obtain **individual-level preference estimates** while accommodating flexible underlying utility functions
- **Overcome statistical biases when aggregating data** across respondents to obtain average group-level estimates

Adaptative Choice Experiments

Optimal dynamic experimental design to estimate parameters efficiently dates back to Wald (1950)

Designs are chosen often some information criterion:

- (Negative) mean squared error of the model's predictions (Fedorov 1972; Cohn et al. 1996; Schein 2005)
- Entropy of the responses (Bates et al. 1996)
- Mutual information between the response and the parameters (**this paper**, and also used in psychology, neuroscience)
- Introduce their own information criterion like EC^2 (Wang, Filiba, and Camerer 2010; Imai and Camerer 2017)
- Norm of the Hessian of the posterior at its mode (Toubia et al. 2013)

This Paper

New theoretical results

- Characterize the designs the procedure using Shannon mutual information criterion generates and document its advantages
- Establish convergence properties

Computational aspects of BACE

- Population Monte Carlo (PMC) methods for computing posteriors and mutual information (numerical integration)
- Bayesian Optimization for finding the most-informative design
- Simulations
- Package `github.com/tt-econ/bace` and Manual `tt-econ.github.io/BACE`

1. Motivation ✓
2. Related Literature
3. BACE Procedure
4. Theoretical Results

Related literature

Decision Theory

- C. P. Chambers, F. Echenique, and N. S. Lambert. Recovering preferences from finite data. Econometrica, 89(4):1633–1664, 2021

Optimal Experimental Design

- Dennis J Aigner. 1979. A brief introduction to the methodology of optimal experimental design
- Steffen Andersen, Glenn W Harrison, Morten Igel Lau, and E Elisabet Rutström. 2006. Elicitation using multiple price list formats
- O Cappé, A Guillin, J. M Marin, and C. P Robert. 2004. Population Monte Carlo. Journal of Computational and graphical statistics 13, 4 (2004), 907–929
- Daniel R Cavagnaro, Richard Gonzalez, Jay I Myung, and Mark A Pitt. 2013. Optimal decision stimuli for risky choice experiments: An adaptive approach. Management science 59, 2 (2013), 358–375

BACE general framework

- $\theta \in \Theta$: compact and convex set of preferences parameters
- $D \in \mathcal{D}$: set of designs/designs the experimenter can show the respondent
- $x \in X$: respondent's choices in compact and connected metric space
- $t \in \{1, \dots, T\}$: time period
- $(d^{1:t}, x^{1:t}) \equiv ((d_1, x_1), \dots, (d_t, x_t))$: set of past designs and answers

The Procedure

We propose the analyst choose menu choose design d that maximizes

$$u(d) = \int_\theta \int_x \log \left[\frac{\Pr(\theta|x,d)}{\Pr(\theta)} \right] \Pr(x|\theta,d) dx \Pr(\theta) d\theta$$

- Roughly, $\log \left[\frac{\Pr(\theta|x,d)}{\Pr(\theta)} \right]$ measures the difference between the posterior and the prior after observing (x, d)
- $\int_x \log \left[\frac{\Pr(\theta|x,d)}{\Pr(\theta)} \right] \Pr(x|\theta,d) dx$ is the expected distance if the agent has preference parameter θ
- Therefore, u captures the expected movement of the posterior

Dynamic Procedure

Prior
$\Pr(\theta \mid (x, d)^{(1:t)})$

$t + 1 \leftarrow t$

Posterior
$\Pr(\theta \mid (x, d)^{(1:t+1)})$

θ—parameter vector of interest
d_t—the question design presented at time t
$(x, d)^t$—answer input at time t

Optimal question design
$\{d_{t+1}\} = \operatorname{argmax}_{d_t} u(d_t)$

Observe answer choice
$(x, d)^{(t+1)}$

Example: Job amenity bundles

Two job scenarios differing by: earnings, work-from-home, control over schedule

- Jobs $j \in \{0, 1\}$ consist of earnings y_j and amenity $a_j \in \{0, 1\}$ and amenity $b_j \in \{0, 1\}$
- $u_j = \log(y_j) + \alpha a_j + \beta b_j + \gamma a_j b_j$
- Probability of making an error when facing $\{j, j'\}$ is a Gumbel distribution with scale parameter σ
  - Also consider the case in which the probability of choosing at random is fixed $p \in [0, 1]$

Estimation procedure:

- Posterior mean
- MLE

How does it look like

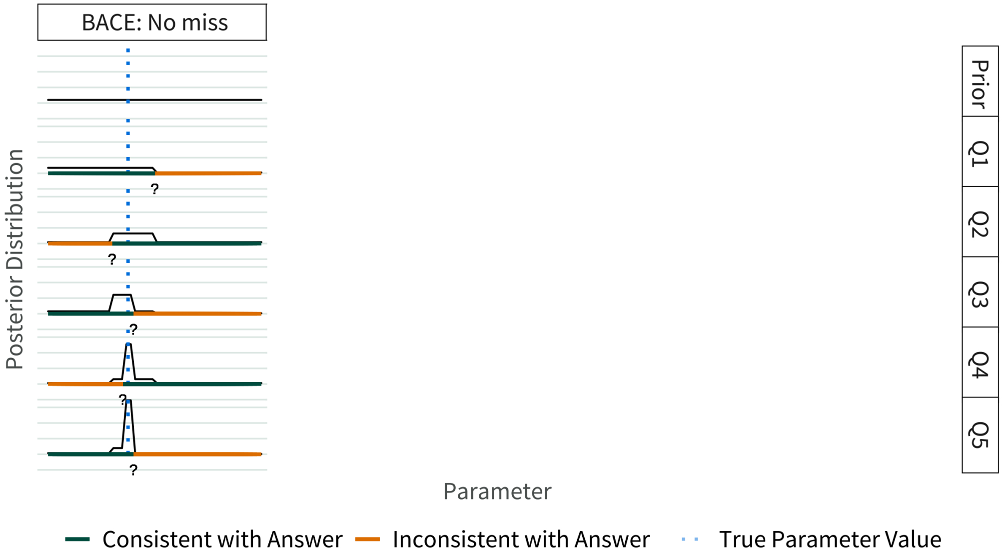

How does it look like

How does it look like

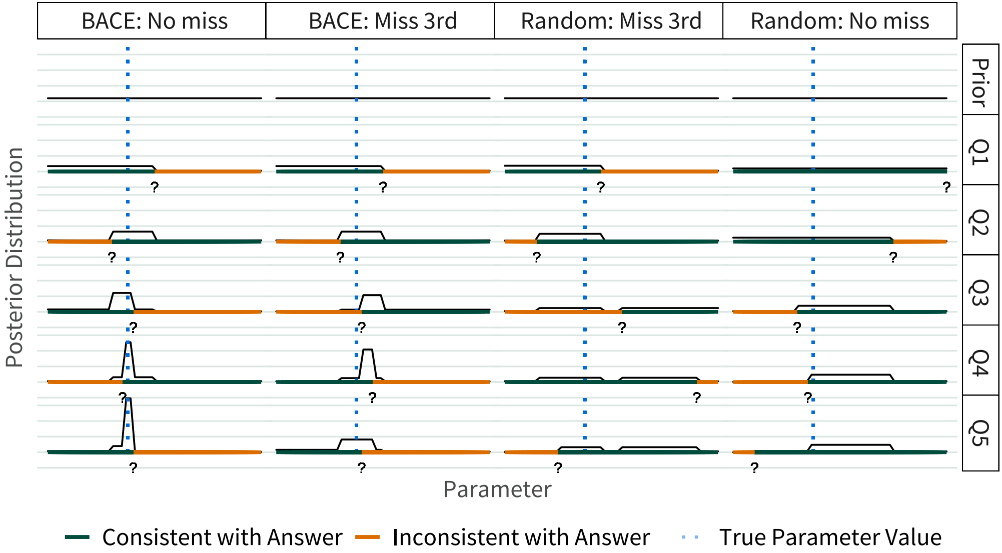

Simulations: Interaction term γ—true versus estimated

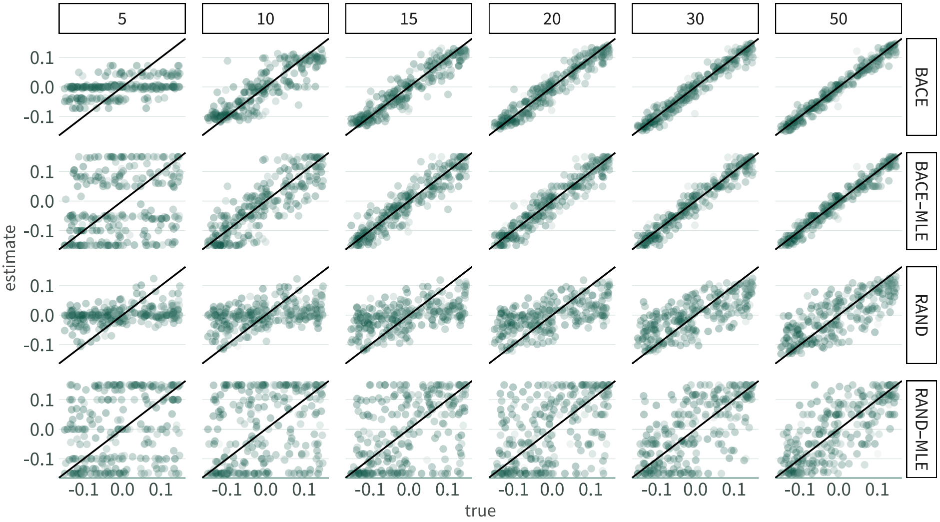

Simulations: Interaction term γ—average MSE

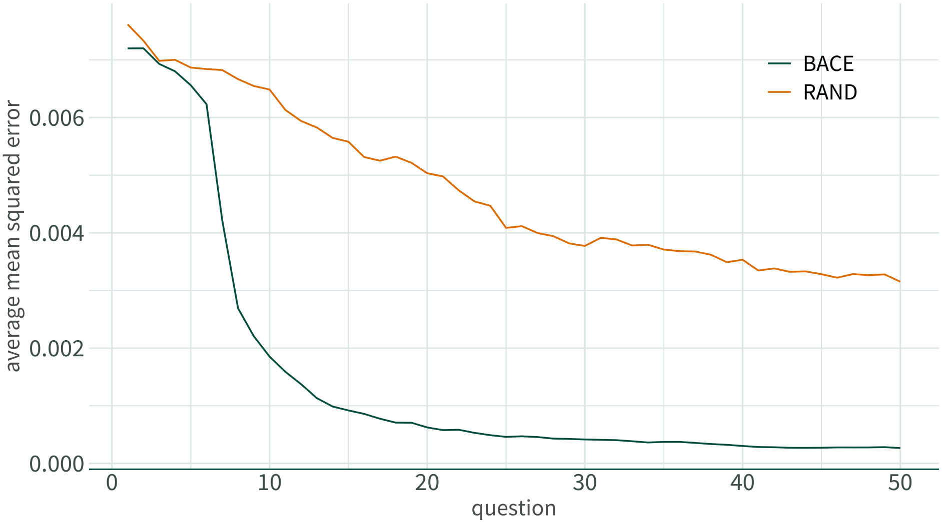

1. Motivation ✓
2. Related Literature ✓
3. BACE Procedure ✓
4. Theoretical Results

BACE formula—binary choice

The general BACE framework above accommodates any type of choice set: binary, multiple, or continuous

- For the theoretical results, we first focus on the binary choice case, and later extend the results
- We simplify the notations in this section:
  - Denote the current prior as $\Pi(\cdot) \equiv \Pr(\cdot|x^{(1:t)}, d^{(1:t)})$ with density f with full support
  - A design d is a set of two alternatives (a, b) (symmetric), $a, b \in A$
  - A binary answer x is $\langle a, b \rangle$ if alternative a is chosen over b and $\langle b, a \rangle$ otherwise
  - The likelihood function is then $\ell \langle a, b | \theta \rangle$
  - As a shorthand, denote the posterior when design (a, b) is presented and a is chosen over b as $f(\theta | \langle a, b \rangle)$

BACE formula—binary choice

We want to choose a design $d = (a, b)$ to maximize a utility function $u(a, b)$ defined by the mutual information. The formula is elegant but does not reveal much

$$u(a, b) = \int_{\Theta} \left(\ell(\langle a, b\rangle|\theta) \log \left(\frac{f(\theta|\langle a, b\rangle)}{f(\theta)} \right) + \ell(\langle b, a\rangle|\theta) \log \left(\frac{f(\theta|\langle b, a\rangle)}{f(\theta)} \right) \right) d\Pi$$

Our first result opens the black box:

Theorem

Suppose there exists (a, b) *such that* $\ell(\langle a, b\rangle|\theta) = 1$ *if and only if* $x \succ_\theta y$ *and* $\Pi(\theta|x \succ_\theta y) = \Pi(\theta|y \succ_\theta x) = \frac{1}{2}$. *Then,* (a, b) *maximizes* u

Non-Continuous Errors

Without any assumptions on ℓ, such menu may not exist... Moreover, ℓ may be such that the optimization problem is not well defined:

1. $\ell(\langle a, b\rangle|\theta) = 1$ whenever $x \succ_\theta y$
2. $\ell(\langle a, b\rangle|\theta) = \overline{q} > \frac{1}{2}$ whenever $x \succ_\theta y$,

Problem is well defined if the experimenter only considers finitely many alternatives:

Proposition

Let $\hat{X} \subset X$ be a finite set of alternatives. Assume $\ell(\theta, x, y) = q$ whenever $x \succ_\theta y$ for some $q \in (\frac{1}{2}, 1]$. Then (x, y) maximizes u if and only if

$$(x, y) \in \operatorname*{argmin}_{(x',y') \in \hat{X} \times \hat{X}} \left| \Pi(\theta | x' \succ_\theta y') - \frac{1}{2} \right| + \left| \Pi(\theta | y' \succ_\theta x') - \frac{1}{2} \right|$$

Non-Continuous Errors

Result implies that if there exists (a, b) such that $\Pi(\theta | a \succ_\theta b) = \frac{1}{2}$, then the menu $\{a, b\}$ is consistent with BACE. When does such pair **always** exist?

- Condition is technical but intuitive
- Amounts to stating that there exists a menu $\{x, y\}$ from which no matter what the subject chooses, it would only be consistent with "50%" of the preferences

Testing Assumption: let λ be the Lebesgue measure and assume that

1. For any $\alpha \in (0, \lambda(\Theta))$, there exists a menu $\{x, y\}$ such that $\lambda(\{\theta | x \succ_\theta y\}) = \alpha$
2. $\alpha_n \to \alpha$ implies $x_{\alpha_n} \to x_\alpha$ and $y_{\alpha_n} \to y_\alpha$

Consistency

We are interested in answering the question: If an experimenter uses BACE, will she learn the true preference? If so, how fast?

- Assume the subject does not make mistakes: $\ell(\theta, x, y) = 1$ whenever $x \succ_\theta y$
- Testing condition
- $\theta^* \in \Theta$ is the true preference

Characterization and consistency results for the case in which the subject is allowed to make mistakes will come later

Intuition

Our results imply that if an experimenter uses BACE, she will only observe sequences $(x_t, y_t)_{t=1}^{\infty}$ of revealed preferences such that

1. $x_T \succeq_{\theta^*} y_T$ for all T
2. $\Pi_T(\theta | x \succ_\theta y) = \frac{1}{2}$ for all T

where $\Pi_T = \Pi(\cdot | (x_t, y_t)_{t=1}^{T})$

- Define a sequence $(x_t, y_t)_{t=1}^{\infty}$ as BACE-compatible if it satisfies (1) and (2)
- Notice that at each T, observing (x_T, y_T) **rules out half of the mass** of Π_T
- Formally $\Pi_{T-1}(\text{supp}(\Pi_T)) = \frac{1}{2}$
- Implies Π_T degenerates as fast as $\frac{1}{2^T} \to 0$

Characterization

Proposition

If $(x_t, y_t)_{t=1}^{\infty}$ *is BACE-compatible, then*

1. $\Pi_T \to \delta_{\theta *}$ *weakly*
2. *for any* T*,*

$$\Pi(\mathsf{supp}(\Pi(\cdot|(x_t, y_t)_{t=1}^{T}))) = \frac{1}{2^T}.$$

Continuous Errors

What if the problem is well defined (ℓ is continuous) but the perfect menu does not exist?

- Turns out that even when it does not exist, the optimal menu is characterized by similar properties

Continuous Errors

What if the problem is well defined (ℓ is continuous) but the perfect menu does not exist?

- Turns out that even when it does not exist, the optimal menu is characterized by similar properties

Theorem

Assume ℓ is continuous. Then $d = (a, b)$ maximizes u if and only if

$$(a, b) \in \underset{(a', b') \in \mathcal{H}}{\operatorname{argmax}} \int_{\Theta} \Big[\ell(\langle a', b' \rangle | \theta) \log(\ell(\langle a', b' \rangle | \theta)) + \\ \ell(\langle b', a' \rangle | \theta) \log(\ell(\langle b', a' \rangle | \theta)) \Big] d\Pi,$$

where $\mathcal{H}$ represents the set of all pairs (a, b) satisfying

$$\int_{\Theta} \ell(\langle a, b \rangle | \theta) d\Pi = \frac{1}{2}.$$

BACE formula remarks: Half-space partitioning designs

$$\mathcal{H} = \{(a, b) | \int_\Theta \ell(\langle a, b\rangle | \theta) d\Pi = \frac{1}{2}\}$$

$\mathcal{H}$ are the designs that partition the space of preference in $\frac{1}{2}$ according to the current prior

- Intuitively, $\int_\Theta \ell(\langle a, b\rangle | \theta) d\Pi$ is the ex-ante probability of observing the respondent choose a over b
  - If $\int_\Theta \ell(\langle a, b\rangle | \theta) d\Pi > \frac{1}{2}$, then it is more likely to observe a being chosen
  - which means you will not learn much if you indeed observe a being chosen.

BACE formula remarks: Entropy

Objective function has a nice interpretation:

$$\int_{\Theta} \Big[\ell(\langle a', b' \rangle | \theta) \log(\ell(\langle a', b' \rangle | \theta)) + \ell(\langle b', a' \rangle | \theta) \log(\ell(\langle b', a' \rangle | \theta)) \Big] d\Pi$$

Recall that for a Bernoulli random variable with probability p, its entropy is given by

$$-p \log(p) - (1-p) \log(1-p)$$

Corollary

BACE generates the design that minimizes entropy over the half-space-partitioning designs. Thus:

1. *Designs that are dominated in terms of error probabilities are never chosen.*
2. *If mistakes are not preference dependent, then BACE generates the design from $\mathcal{H}$ that minimizes the error probability.*

Consistency

Assume a continuous ℓ. Does the procedure work?

- Same questions as before
- A more involved approach is needed to define convergence of beliefs
  1. Need to show that if the experimenter uses BACE, each θ induces a well defined probability measure P_θ over $(X \times X)^\infty$
  2. $\theta \sim \Pi$ and $(X_1, Y_1), (X_2, Y_2), \dots | \theta \sim P_\theta$ jointly define non-i.i.d. Bayesian inference problem
  3. The experimenter is *consistent* at $\theta \in \Theta$ if for every neighborhood U of θ, we have
  $$\Pi(U | (X_t, Y_t)_{t=1}^T) \to 1 \quad \text{almost surely under } P_\theta,$$
- Goal: Show that an experimenter who uses BACE is consistent

BACE Consistency

Theorem

There exists a set Θ_Π *with* $\Pi(\Theta_\Pi) = 1$ *such that the experimenter is consistent at all* $\theta \in \Theta_\Pi$.

Intuition:

- By Martingale Convergence Theorem, beliefs will converge
- Since BACE only generates half-space-partitioning menus, in the limit, the posterior distribution cannot assign positive weight to more than one parameter
- Finally, each θ is more likely to generate a sequence that is more consistent with $\succeq_\theta$ than with another $\succeq_{\theta'}$ so beliefs cannot be mistaken in the limit

Remarks

In the proof we only use the Half-Space-Partitioning property of BACE

- Thus result holds for **any** procedure that also generates Half-Space-Partitioning menus

- Unfortunately, no rate of convergence (in general)
    - If people make mistakes at a fixed probability $(1-q)$, then the rate is $2^t[(q,(1-q))]^{\frac{t}{2}} \to 0$

Concluding remarks

In the paper we also prove

- Achieves the highest rate of convergence when subjects do not make mistakes
- Works for uncertain **parametric** error
- There exists an incentive compatible payment procedure
- Works for probabilistic Data

Thank You!

Implementation details: Priors

```python
# Preference parameters (theta_params)
# Dictionary where each preference parameter has a prior distribution specified by a scipy.stats distribution
# All entries must have a .rvs() and .log_pdf() method
theta_params = dict(
    blue_ink=scipy.stats.uniform(loc=-2, scale=4),
    fountain_pen=scipy.stats.norm(loc=1, scale=1),
    p = scipy.stats.uniform(loc=0.75, scale=0.24)
)
```

Implementation details: Question space

```python
# Design parameters (design_params)
# Dictionary where each parameter specifies what designs can be chosen for a characteristic
# See https://github.com/ARM-software/mango#DomainSpace for details on specifying designs
design_params = dict(
    price_a = scipy.stats.uniform(0.5, 5),
    price_b = scipy.stats.uniform(0.5, 5),
    color_a = ['Black', 'Blue'],
    color_b = ['Black', 'Blue'],
    type_a = ['Fountain', 'Ballpoint'],
    type_b = ['Fountain', 'Ballpoint'],
)
```

Implementation details: Utility function

```python
# Specify likelihood function
# Returns Prob(answer | theta, design) for each answer in answers
def likelihood_pdf(answer, thetas,
                   # All keys in design_params here
                   price_a, price_b,
                   color_a, color_b,
                   type_a, type_b):

    eps = 1e-10

    u_a = -price_a + thetas['blue_ink'] * (color_a == "Blue") + thetas['fountain_pen'] * (type_a == 'Fountain')
    u_b = -price_b + thetas['blue_ink'] * (color_b == "Blue") + thetas['fountain_pen'] * (type_b == 'Fountain')
    utility_diff = u_b - u_a

    # Choose higher utility option with probability p. Choose randomly otherwise.
    u = utility_diff>0
    u[utility_diff==0] = 1/2
    likelihood = u * thetas['p'] + (1/2) * (1-thetas['p'])
```

Implementation details: Qualtrics 1

Implementation details: Qualtrics 2

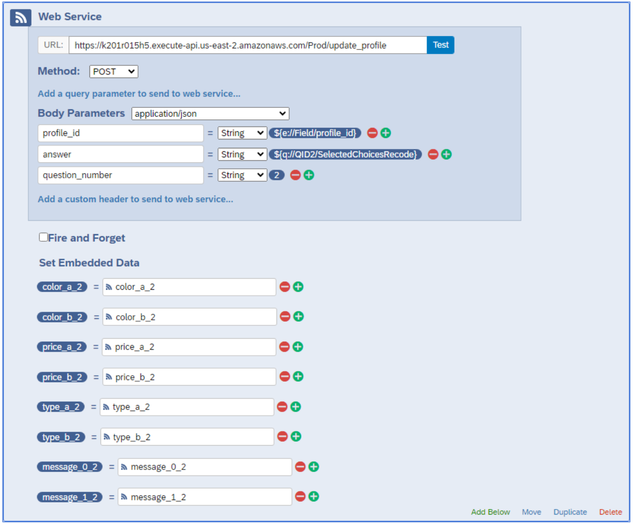

Implementation

Main challenges

- Computational burden of the optimal next-based scenario
- How to implement outside of the lab

Implementation

1. Pre-calculate the "optimal tree"
   - Tree size grows exponentially and include many unlikely paths
   - Fixed tree by setting and prior
2. Calculate in real time and adapt to any setting and prior ✓
   - Design and answer spaces can be discretized $\implies$ calculation of $I(d)$ boils down to having a reliable estimate for the posterior
   - Use Population Mote Carlo for "posterior" calculation
   - Employ Bayesian optimization to find the best next-scenario

Implementation

Population Mote Carlo (PMC)

1. Given (d^t, x^t), sample θ's according to the current prior
2. Estimate posterior distribution for each θ_i to get $P(\theta_i | (d^t, x^t))$
3. Select a proposal distribution q and sample $\theta_i' \sim q(\theta_i)$
4. Calculate weights $w_i = \frac{\ell(x|d^t, \theta_i')}{q_i(\theta_i'|\theta_i)}$ and normalize them to sum to 1
5. Resample θ_i' with replacement using the weights
6. Build a discrete measure $\sum_{i=1}^{N} w_i \delta_{\theta_i'}$

Discrete measure can be used to approximate $I(d_{t+1})$

Implementation

Bayesian Optimization: $\max_{d_{t+1}} I(d_{t+1})$

1. Evaluate I at N random points to get $d_1, ..., d_N$
2. Fit the observed data with a Gaussian Process
3. For each d we can calculate the expected improvement with respect to the best d_i in the sample
   - Has a closed-form solution
4. Choose the d with the highest expected improvement

Implementation

Backend Framework

- Create a profile in PostgreSQL database that will store relevant information for that user, such as the current estimate of the posterior distribution
- REST API using the FastAPI framework calculates the optimal next scenario and transmits it to Qualitrics (or Survey Monkey)
- User observes question and answers accordingly
- Question and answer are transmitted to the Server which updates the posterior using PMC
- Posterior, answer, and question are stored in PostgreSQL

Takes 1-2 seconds

Implementation

https://github.com/tt-econ/bace

Fig. 2. Implementation schematic

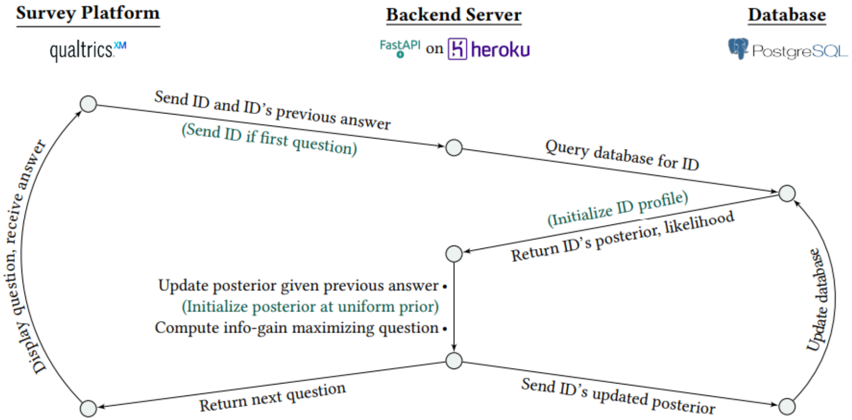

Note: A simplified schematic of the interactions between the survey platform and the backend computation and database servers.