Bayesian Adaptive Choice Experiments

Marshall Drake — EEOC
Fernando Payró — UAB
Neil Thakral — Brown
Linh T. Tô — BU

February 2024*

Abstract

We propose the use of a dynamic choice experiment method, which we call Bayesian Adaptive Choice Experiment (BACE), to elicit preferences efficiently. BACE generates an adaptive sequence of menus from which subjects will make choices. Each menu is optimally chosen, according to the mutual information criterion, using the information provided by the subjects' previous choices. We provide sufficient conditions under which BACE achieves convergence and show that its convergence rate significantly improves upon existing discrete choice methods with randomly generated menus. We show that it achieves the highest possible rate of convergence whenever preferences are deterministic. Beyond efficiency gains, BACE addresses a bias in estimating population-level average preference parameters stemming from using combined data across individuals when individuals differ in their tendency to be inconsistent in their choices. Given that BACE requires the calculation of a Bayesian posterior as well as the solution to a non-trivial optimization problem, several computational challenges arise. We address such challenges by using Bayesian Monte Carlo techniques and provide a package for researchers to employ. The separation between a front-end survey interface and a back-end computational server allows the BACE package to be portable for research designs in a wide range of settings.

*Drake: Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215 (email: mhdrake@bu.edu). Payró: Departament d'Economia i d'Història Econòmica, Universitat Autònoma de Barcelona, Edifici B, Campus UAB 08193 (email: fpayrochew@gmail.com). Thakral: Department of Economics, Brown University, Box B, Providence, RI 02912 (email: neil_thakral@brown.edu). Tô: Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215 (email: linhto@bu.edu).

1 Introduction

Surveys can be a powerful tool to learn about important economic measures not present in other types of datasets. Old-generation surveys capture many of the quantities that are now more readily available in administrative datasets, such as measures of wages or public benefit take-up. But a new generation of surveys—which use customizable, controllable, and interactive methodologies—promises to do much more in helping us measure the important inputs for understanding a broad range of economic phenomena.

Perhaps the most central example is individual preferences, a fundamental ingredient for microeconomic analysis. While *revealed preference* has been a important tool for economists, this approach has proven to be limited in some important economic environments. Well-acknowledged limitations include the strong modeling assumptions and data availability required to infer individual preferences, the inability to learn about non-use values, and the unobserved factors, market imperfections, and behavioral biases that can present serious confounds when inferring preferences using observational data. This has led to a proliferation of research using *stated preference* approaches such as Discrete Choice Experiments (DCEs) to estimate individual preferences in a broad range of applications. These include studies in labor economics (Mas and Pallais, 2017), public economics (Neustadt and Zweifel, 2011), health economics (Ryan, Gerard and Amaya-Amaya, 2007), environmental economics (Carson and Czajkowski, 2014), development economics (Jeuland et al., 2009), agricultural economics (Schulz, Breustedt and Latacz-Lohmann, 2014), urban economics (Bullock, Scott and Gkartzios, 2011), education (Czajkowski et al., 2020), psychology (Ida and Goto, 2009), criminology (Picasso and Cohen, 2019), real estate (Glumac and Wissink, 2018), transportation (Bliemer and Rose, 2011), and marketing (Green, Krieger and Wind, 2001).

In economics, DCEs have become increasingly popular, with one of the more common uses in recent years being the measurement of workers' preferences for job attributes (Eriksson and Kristensen, 2014; Mas and Pallais, 2017; Wiswall and Zafar, 2018; Maestas et al., 2018; Mas and Pallais, 2019; Gelblum, 2020; Feld, Nagy and Osman, 2020), which is an application we revisit in this paper. Discrete choice experiments are also widely used in settings with incentivized choices. The proliferation of lab experiments also makes use of dynamic features of surveys to design incentivized experiments to measure important preference parameters such as risk and time preference or loss aversion (Andersen et al., 2006).[1] Beyond the academic literature, these methods feature prominently in policy analysis, regulation, and

[1]Multiple price lists are a form of discrete choice experiments. Experimentalists also use sequential binary lottery choices to estimate preference parameters.

litigation (Carson, 2012).

In this project we propose the use of dynamic choice experiments, a survey method which we call the Bayesian Adaptive Choice Experiment (BACE) framework, to replace existing static choice experiments widely used to elicit preferences. BACE makes it possible to efficiently obtain individual-level preference estimates while accommodating flexible underlying utility functions and a broad range of designs of the choice experiments. In addition, the method significantly improves efficiency and overcomes biases in approaches that rely on aggregating data across individuals to estimate an average population-level estimate.

A standard Discrete Choice Experiment (DCE) asks respondents to choose among a set of alternatives that vary along multiple dimensions. Without a dynamic framework, these alternatives consist of pre-generated, often randomized bundles of characteristics. The resulting lack of statistical power to infer preferences at the individual level typically necessitates a focus on estimating average preferences when using static approaches. This has several notable shortcomings: (1) it requires making assumptions about the preference distribution in the population as well as implicit assumptions about homogeneity in respondents' inconsistency in making choices; and (2) it can lead to biased estimates of average preferences, related to the mean-variance confound in estimating limited dependent variable models using maximum likelihood or minimum chi-square estimators.

BACE provides an efficient dynamic elicitation procedure for conducting choice experiments that overcomes these problems. It does so by generating an efficient sequence of choice scenarios based on a prior that gets updated with previous answers to obtain individual-level Bayesian posterior estimates. At each stage of experimentation, the next scenario to be presented is the one that will yield the greatest information gain about the parameter values, which can include a choice consistency parameter. The procedure thus allows for an efficient elicitation of preferences for each individual, taking into account heterogeneity in choice inconsistency.

The increasing use of hypothetical choice experiments in economics and related fields has helped provide evidence and support for the reliability of the method. Existing research shows that estimates from choice experiments are often in reasonable ranges and with expected signs (Mas and Pallais, 2017; Maestas et al., 2018); consistent across different subject pools (Mas and Pallais, 2017); consistent with subsequent choices (Wiswall and Zafar, 2018; Aucejo, French and Zafar, 2021); and superior to estimates from other types of survey questions such as open-ended questions and multiple price list, which tend to be noisy and inconsistent with basic economic theory (Feld, Nagy and Osman, 2020). In fact, when comparing four elicitation methods (discrete choice experiment, open-end questions, pay card / multiple price list, and double bounded dichotomous choice questions), Feld, Nagy and Osman (2020) find that only with the DCE is there no valuation that is inconsistent with economic theory. BACE then provides a timely and important improvement for a reliable method that has been proving its usefulness and can be applied broadly by many researchers.

While adaptive designs for choice experiments have been proposed in previous research, the biggest barrier to implementation outside of university computer laboratories using student subjects has been computational costs.

The idea of optimal experimental design to estimate parameters efficiently dates back to Peirce (1967), and Wald (1950) describes the idea of dynamic designs in statistics. While the concept is widespread in many fields in the physical and biological science, it is not often discussed and rarely implemented in economics (see Aigner 1979; Moffatt 2007; Chapman et al. 2018 for further discussion). The most common application in economics and psychology so far has been the elicitation of time and risk preferences (Cavagnaro et al., 2013; Toubia et al., 2013; Cavagnaro et al., 2016; Chapman et al., 2018; Imai and Camerer, 2018), though the implementations there are largely limited to small-scale within-the-lab versions or as coarse pre-computed approximations to the Bayesian-optimal dynamic elicitation.[2]

We contribute to this literature in three ways.

First, we formalize the method under a decision-theoretic framework. We show formal conditions under which preferences can be identified in finite choice data, along the lines of Chambers, Echenique and Lambert (2021). Their paper focuses on a setting with identically and independently distributed data in a frequentist framework. Our problem adds an additional layer of complexity when the menu choice is dynamically generated in a Bayesian framework. We discuss the convergence rate of our method both theoretically and in simulations.

Second, our implementation makes a step forward in allowing such procedures to be more widely adopted to study a wide range of other possible social science applications. We provide an implementation of BACE that is portable, scalable, and computationally feasible and provide a more detailed description of the properties of BACE. The code allows a survey platform (e.g., Qualtrics) to interact with on-the-cloud backend servers that can do large-scale computation of the next-best scenario simultaneously across survey subjects in real time. The computation uses Bayesian Monte Carlo methods to ensure that computational speed is practical when implementing live surveys for thousands of respondents simultaneously. This resulting package allows for other researchers to easily apply the method to different settings to address a wide range of questions.

[2]In the context of a correspondence study, Avivi et al. (2021) consider the efficiency gain from dynamically adapting the profiles of fictitious applicants sent to employers. A related but distinct literature explores adaptive procedures with a different objective, namely, to maximize the gain from experimental treatments while measuring treatment effects; recent papers include Caria et al. (2020) and Kasy and Sautmann (2021).

Third, we document a systematic bias that arises when attempting to estimate average preference parameters by pooling data across individuals. Pooling is common in practice when using less efficient elicitation approaches such as static designs. The bias result highlights the importance of using methods that provide precise individual-level estimates.

2 The BACE procedure

This section describes the main components that make up a Bayesian Adaptive Choice Experiment. We begin by introducing notation:

- $\theta \in \mathbb{R}^k$: Vector of preference parameters that the researcher is interested in estimating.
- D: Set of designs that can be shown to respondents. Each element $d \in D$ is a vector that represents a discrete choice question that can be shown to respondents.
- X: Discrete set of possible answers that can be observed. For example, an individual may choose between two options in a binary DCE. In that case, $X = [0, 1]$ or $X = [OptionA, OptionB]$.
- $t \in 1, \dots, T$: t represents the time period or the number of questions an individual has answered.
- The sets of past design scenarios and answers up to time t are given by $d^{(1:t)} \equiv \{d^1, \dots, d^t\}$ and $x^{(1:t)} \equiv \{x^1, \dots, x^t\}$, respectively.
- $\ell(x \mid \theta, d)$: The likelihood of observing the answer x to scenario d if an agent has preference parameters θ.

A standard discrete choice experiment involves asking subjects to choose between a series of hypothetical scenarios. Each scenario presents two (or more) options that vary along a set of characteristics. Subjects evaluate tradeoffs between the options and select their preferred option. The researcher's ultimate goal is to estimate a vector of preference parameters, θ, that captures how an individual values different characteristics of a good based on how subjects answer a series of scenarios. In practice, these scenarios are typically randomly generated across a range of values.[3]

The adaptive Bayesian approach to the standard discrete choice experiment generates an efficient sequence of hypothetical scenarios in real-time based on a prior that gets updated

[3] In many cases, scenarios that involve ex-ante pre-determined dominated choices are eliminated to increase statistical power, such as in Wiswall and Zafar (2018); Maestas et al. (2018).

with past answers. At each stage of experimentation, the next scenario to be presented is the one that will generate the highest amount of information gained about the parameter values. In a simple search model when subjects evaluate the value of one parameter against a benchmark one at a time, the most efficient approach is a binary search. The adaptive Bayesian approach has the same principle as a binary search, but it allows for complex search problems with multiple dimensions and allows for choices that are made inconsistently due to inattentiveness.

The researcher's overarching goal is to accurately estimate individual i's preferences, θ_i. To do so, the researcher estimates $\Pr\left(\theta_i \,\middle|\, x_i^{(1:T)}, d_i^{(1:T)}\right)$ for each individual, which represents the posterior belief over θ_i conditional on individual i's answers to T questions and the researcher's prior. In practice, the number of questions that each individual sees, T, is limited by the subject's attention and the researcher's budget. Standard processes may select a series of scenarios randomly or use a pre-determined sequence of questions that covers the parameter range of interest, e.g. standard multiple price lists. However, these methods are inefficient because they fail to incorporate information from previous questions, produce incorrect estimates when individuals make mistakes, and can only distinguish between a set of respondent types that are given by the chosen set of designs.[4] Intuitively, adaptive experimental approaches are able to estimate θ more accurately, and with fewer questions, than previous methods because future scenarios are chosen using information from how an individual made choices in the past.

Formally, the BACE procedure is as follows. At each period $t \in \{1, 2, \ldots, T\}$, a choice scenario d^t is presented, and the respondent's answer x^t is recorded. At time t, data consists of the set of questions that have been shown to respondents ($d^{(1:t)} \equiv \{d^1, \ldots, d^t\}$) and the observed answers ($x^{(1:t)} \equiv \{x^1, \ldots, x^t\}$) to those questions. This information acts as data entering the beginning of period $t+1$. Let θ be the parameter vector that the researcher wants to estimate. At the beginning of period $t+1$, our prior for θ is denoted as $\Pr\left(\theta \,\middle|\, x^{(1:t)}, d^{(1:t)}\right)$, which is calculated using Bayes' rule.

The problem at time $t+1$ is to find the optimal d^*_{t+1} among all possible scenarios d_{t+1}s. Intuitively, given all the possible answers that one could observe, which design d^*_t reveals the most information about the parameters we are interested in estimating, θ? The criterion chosen is based on information theory; we maximize the mutual information between the parameter random value Θ and the outcome random value X^{t+1} (the potential answer at $t+1$) conditional on the scenario d_{t+1} (Shannon, 1948). The interpretation is that we select

[4] For example, a multiple price list with N options can separate individuals into $N+1$ types. A binary decision tree with N questions separates individuals into 2^N types, and estimates will be inaccurate if individuals make mistakes, particularly during the early stages of an experiment.

the scenario that yields the largest information gain about the parameters of interest after observing a new answer.[5] Denote the mutual information as $U(d_{t+1}) := I(\Theta; X^{t+1}|d_{t+1})$:[6]

$$U(d_{t+1}) = \int_\theta \int_{x^{t+1}} \left[\log \frac{\Pr\big(\theta \,\big|\, x^{(1:t+1)}, d_{t+1}\big)}{\Pr\big(\theta \,\big|\, x^{(1:t)}\big)} \right] \Pr(x^{t+1} \,|\, \theta, d_{t+1}) \Pr\big(\theta \,\big|\, x^{(1:t)}\big) dx^{t+1} d\theta \quad (1)$$

In the language of the Bayesian experimental design literature (Chaloner and Verdinelli, 1995), this equation is the utility function of the researcher, and the objective is to find $d^*_{t+1} = \arg\max_{d_{t+1}} U(d_{t+1})$.

Note that $\Pr(x^{t+1} \,|\, \theta, d^{t+1})$ is the likelihood of the answer x^{t+1} given the presented scenario d^{t+1}, at parameter value θ, which can be computed from the utility function. In the case of testing across utility models, the formula above also needs to summarize over all the candidate models. The posterior $\Pr\big(\theta \,\big|\, x^{(1:t+1)}, d^{t+1}\big)$ can be computed given the likelihood function and the researcher's prior using Bayes' rule. Based on the realized answer, the posterior is then updated:

$$\Pr\big(\theta \,\big|\, x^{(1:t+1)}, d_{t+1}\big) = \frac{\Pr(x^{t+1} \,|\, \theta, d_{t+1}) \Pr\big(\theta \,\big|\, x^{(1:t)}\big)}{\int_{\theta'} \Pr(x^{t+1} \,|\, \theta', d_{t+1}) \Pr\big(\theta' \,\big|\, x^{(1:t)}\big) d\theta'} \quad (2)$$

Figure 1 shows the schematic illustration of the steps involved. The procedure starts with a prior distribution $p(\theta)$ that is defined by the researcher. This prior can be based on estimates from the literature, results from a pilot experiment, or the researcher's prior knowledge.[7] The posterior at time t constitutes the prior at time $t + 1$. The optimal design is chosen according to this prior and is presented to the respondent. The researcher observes the individual's answer and calculates a posterior distribution using the new data. If $t < T$, then the process is repeated using the posterior distribution as the new prior.

[5] Paninski (2005) shows that if the prior is absolutely continuous and has bounded density, the mutual information criterion can choose designs that lead to consistent and efficient parameter estimates. Alternative criteria include maximizing the inverse of the asymptotic covariance matrix of the maximum likelihood estimate as in Toubia et al. (2013) or the Equivalence Class Edge Cutting information criterion as in Imai and Camerer (2018), among others; see Ryan et al. (2016) for a review.

[6] The mutual information is the same as the Kullback-Leibler divergence between the joint distribution and the product of the marginal distributions of Θ and $X^{t+1}|d_{t+1}$.

[7] In practice, both the questions that are selected and the posterior estimates depend on the prior that the researcher selects, which is typical in Bayesian statistical methods. However, the influence of the prior distribution on posterior estimates will decrease as more data is accumulated. Nevertheless, the method can only produce parameter estimates that fall within the support of the prior distribution by design, and researchers should take care that the prior they specify covers the range of possible values they expect in the population.

3 Decision Theoretic Formulation of BACE

3.1 Primitives

BACE features an experimenter and a subject. The subject has a preference over a set of alternatives X, which is assumed to be a compact and connected metric space. The experimenter would like to learn the subject's preferences by observing her choose from menus.

By *preference*, we mean a binary relation $\succeq$ over X. Formally, $\succeq$ is a subset of $X \times X$. We write $x \succeq y$ whenever $(x, y) \in \succeq$. Given a preference $\succeq$, denote its asymmetric part by $\succ$ (i.e., $x \succ y$ if $(x, y) \in \succeq$ and $(y, x) \notin \succeq$) and its symmetric part by $\sim$ (i.e., $x \sim y$ if $(x, y) \in \succeq$ and $(y, x) \in \succeq$). We say x is strictly preferred (resp. indifferent) to y if $x \succ y$ (resp. $x \sim y$). Let $\mathcal{X}$ denote the set of all preferences over X. We endow $\mathcal{X}$ with the Hausdorff hemimetric.

The experimenter believes the subject's preferences conform to a parametric model. She considers a set of parameters $\Theta \subseteq \mathbb{R}^n$ and a function $\phi\colon \Theta \to \mathcal{X}' \subset \mathcal{X}$. We write $\succeq_\theta$ to denote $\phi(\theta)$ and refer to it as the preference of a subject conforming to θ. The experimenter's uncertainty over the subjects' preference is captured by a prior Π over the probability space $(\Theta, \mathcal{B}(\Theta))$, where $\mathcal{B}(\Theta)$ is the Borel σ-algebra.

In this framework, BACE would generate an adaptive sequence of menus from which the subject would make a choice. We restrict BACE to generate menus that contain no more than two alternatives but discuss generalizations of our results in Section 5.

When facing a menu, the subject's behavior is guided by his preferences. However, we allow the subject to make a mistake and choose an alternative that is not preferred with some probability. This is captured by the function $\ell\colon \Theta \times X \times X \to [0, 1]$. Assume that for all $x, y \in X$, the $\ell(\cdot, x, y)$ is measurable on Θ.

When making a choice from a menu $\{x, y\}$, a subject conforming to θ selects x over y with probability $\ell(\theta, x, y)$, and y over x with complementary probability. We assume $\ell(\theta, x, y) > \frac{1}{2}$ when $x \succ_\theta y$, and $\ell(\theta, x, y) = \frac{1}{2}$ when $x \sim_\theta y$.

Given the primitives—the set of alternatives X, the set of parameters Θ, the preference function ϕ, the prior Π, and the function ℓ—we refer to the tuple $(X, \Theta, \phi, \ell, \Pi)$ as a *parametric model*. We make the following assumptions on the model:

- Θ is a compact and convex set.

- $\phi\colon \Theta \to \mathcal{X}' \subset \mathcal{X}$ is one-to-one and onto, and $\mathcal{X}'$ is a set of transitive preferences such that $x \succeq y$ implies there exists $x', y' \in X$ that are arbitrarily close to x and y, respectively, such that $x' \succ y'$. Preferences that satisfy this property are called strict preferences.

- Π has full support and admits a density f.

3.2 BACE Procedure for Stochastic Choice

In order to apply BACE, the experimenter needs to be able to calculate a posterior after observing the subject make a choice. By Bayes' theorem, if the subject chooses x from $\{x, y\}$, resulting in the sample (x, y), then the posterior density is given by

$$f(\theta|(x,y)) = \frac{\ell(\theta,x,y)f(\theta)}{\int_\Theta \ell(\theta,x,y)f(\theta)d\theta},$$

whenever $\int_\Theta \ell(\theta,x,y)f(\theta)d\theta \neq 0$.

Define the function $u\colon X \times X \to \mathbb{R}$ by

$$u(x,y) = \int_\Theta \Big[\ell(\theta,x,y)\log(\frac{f(\theta|(x,y))}{f(\theta)}) + \ell(\theta,y,x)\log(\frac{f(\theta|(y,x))}{f(\theta)})\Big]f(\theta)d\theta.$$

Notice that $u(x,y) = u(y,x)$. Therefore, if (x,y) maximizes u, so does (y,x). An experimenter using BACE would choose the menu $\{x,y\}$ such that (x,y) maximizes u. Although elegant, the BACE formula does not shed any light on which menus are generated by the procedure. Our next result fully characterizes the properties of the menus BACE generates under the appropriate continuity assumption on ℓ.

Theorem 1. *Suppose ℓ is continuous in the product metric and $\ell(\theta,x,y) \in (0,1)$ for all $x, y \in X$ and $\theta \in \Theta$. Then the following are equivalent:*

$$(x,y) \in \underset{(x',y')\in X\times X}{\arg\max}\; u(x',y') \tag{3}$$

$$(x,y) \in \underset{(x',y')\in \mathcal{H}}{\arg\max} \int_\Theta \Big[\ell(\theta,x',y')\log(\ell(\theta,x',y')) + \ell(\theta,y',x')\log(\ell(\theta,y',x'))\Big]f(\theta)d\theta, \tag{4}$$

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

$$\int_\Theta \ell(\theta,x,y)f(\theta)d\theta = \frac{1}{2}.$$

We say $\{x,y\}$ is a *half-space-partitioning menu* if $(x,y) \in \mathcal{H}$ and $\Pi(\theta|x \succ_\theta y) > 0$. The restriction to half-space-partitioning menus establishes that BACE generates a menu $\{x,y\}$ in which the ex-ante probability of observing the subject choose x over y is equal to $\frac{1}{2}$. This alone does not characterize which of such menus will be generated. As we demonstrate next, BACE generates the menu in which the expected choice results in minimal entropy. Because entropy is a measure of stochasticity, we can infer that BACE generates the menu with the lowest expected stochasticity among the half-space-partitioning menus.

Observe that the tuple $(\theta, \{x, y\}, \ell)$ induces a Bernoulli distribution: Letting $\omega_1 = (x, y)$ and $\omega_2 = (y, x)$, we can define $p_\theta(\omega_1) = \ell(\theta, x, y)$ and $p_\theta(\omega_2) = \ell(\theta, y, x)$. Then p_θ is a Bernoulli probability measure. The entropy for Bernoulli measures is

$$H(p_\theta) = -p_\theta \log(p_\theta) - (1 - p_\theta) \log(1 - p_\theta).$$

Therefore, the objective function in Equation (4) can be equivalently written as

$$\int_\Theta -H(p_\theta) f(\theta) d\theta,$$

which implies that BACE generates the menu that minimizes the expected entropy among the half-space-partitioning menus.

Any procedure that minimizes entropy will never generate a menu that is dominated in terms of error probabilities. Intuitively, if there are two menus $\{x, y\}$ and $\{z, w\}$, and every possible preference θ is more likely to make a mistake when facing $\{z, w\}$ than when facing $\{x, y\}$, then the experimenter should not offer $\{z, w\}$. This observation implies that if the subject's probability of making a mistake is independent of their preference, then BACE will generate the menu from which they are less likely to make a mistake among the half-space-partitioning menus. The following corollary documents that BACE indeed has this property.

Corollary 1. *The following statements are true:*

1. *Assume $(x, y), (w, z) \in \mathcal{H}$ are such that $x \succ_\theta y \implies w \succ_\theta z$ and $\ell(\theta, x, y) > \ell(\theta, w, z)$. Then, $(w, z) \notin \underset{(x', y') \in X \times X}{\arg\max}\ u(x', y')$.*
2. *If $\ell(\theta, x, y) = q(x, y)$ whenever $x \succeq_\theta y$, then*

$$(x, y) \in \underset{(x', y') \in X \times X}{\arg\max}\ u(x', y') \iff (x, y) \in \underset{(x', y') \in \mathcal{H}}{\arg\max}\ q(x', y').$$

Theorem 1 relies on the fact that ℓ is continuous; however, this is not always an appropriate assumption. For example, if subjects do not make mistakes ($\ell(\theta, x, y) = 1 \iff x \succ_\theta y$) or the probability of making a mistake is constant ($\ell(\theta, x, y) = \overline{q} \iff x \succ_\theta y$), then ℓ will not be continuous. Without continuity of ℓ, we cannot ensure the existence of a maximizer for the BACE objective function and thus cannot provide a corresponding characterization.

Nevertheless, we can provide a characterization for the case in which the experimenter only considers a finite set of alternatives for the experiment.

Proposition 1. *Let $\hat{X} \subset X$ be a finite set of alternatives such that for all $x, y \in \hat{X}$, there is no open $C \subset \Theta$ for which $x \sim_\theta y$ for all $\theta \in C$ whenever $x \neq y$. If $\ell(\theta, x, y) = q \iff x \succ_\theta y$ for some $q \in (\frac{1}{2}, 1]$, then*

$$(x, y) \in \underset{(x',y') \in \hat{X} \times \hat{X}}{\arg\max}\ u(x', y') \iff (x, y) \in \underset{(x',y') \in \hat{X} \times \hat{X}}{\arg\min} |\Pi(\theta | x' \succ_\theta y') - \frac{1}{2}|.$$

Proposition 1 not only characterizes BACE when the experimenter can only consider finitely many alternatives but also provides a formal relationship between BACE and the binary search literature. Specifically, it illustrates that when the probability of mistakes is constant, BACE operates as a binary search algorithm: it first partitions the space into two sets and then tests them against each other. It also suggests a simplified way to implement BACE under specific assumptions.

Under a relatively mild additional assumption, we can restore the conclusion of Theorem 1. To give intuition, suppose the subject does not make any mistakes, and the experimenter's goal is to rule out "50%" of all the parameters Θ. The assumption 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 parameters. Formally, 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 $y_{\alpha_n} \to y_\alpha$.

We refer to this assumption as our *testing assumption*.

Proposition 2. *Assume $\ell(\theta, x, y) = q \iff x \succ_\theta y$ for some $q \in (\frac{1}{2}, 1]$. Then, under the testing assumption,*

$$(x, y) \in \underset{(x',y') \in X \times X}{\arg\max}\ u(x', y') \iff \Pi(\theta | x \succ_\theta y) = \Pi(\theta | y \succ_\theta x) = \frac{1}{2}.$$

4 Convergence Results

In this section, we show that an experimenter who uses BACE will eventually learn the true preference of the subject. We separate the results into two cases: one when the subject does not make mistakes and another when the subject does. We begin by introducing the necessary notation and concepts needed to formalize what we mean by "learning the true preference" in the case that the experimenter uses randomization to generate the menu to offer the subject.

4.1 Definition of "learning the true preference"

Chambers, Echenique and Lambert (2021) model the random menu-generating framework using a full-support probability measure ν over $(X, \mathcal{B}(X))$, where $\mathcal{B}(X)$ is the Borel σ-algebra on X, which they also use to denote the product measure on $(X \times X, \mathcal{B}(X) \times \mathcal{B}(X))$. Loosely speaking, $\nu(x, y) = \nu(x)\nu(y)$ represents the probability that the experimenter offers a menu consisting of x and y. They define

$$P_{\theta,\nu}(B) = \int_B \ell(\theta, x, y) d\nu(x, y) \quad \text{for } B \in \mathcal{B}(X) \times \mathcal{B}(X)$$

and interpret $P_\theta((x, y))$ as the probability that the experimenter observes the subject choosing x over y when she uses ν to generate the menus. Equivalently, the experimenter faces a random variable $(X, Y) \sim P_{\theta,\nu}$, and thus using ν infinitely many times to generate menus results in a sequence of independently and identically distributed random variables $(X_1, Y_1), (X_2, Y_2), (X_3, Y_3), \ldots \sim P_{\theta,\nu}$.

Within this framework, as Chambers, Echenique and Lambert (2021) show, a frequentist experimenter can estimate θ. However, given the novelty of the preference identification problem, there are no existing results for its Bayesian counterpart.

We take a distinct approach by considering a Bayesian experimenter, characterized by the prior Π. The experimenter's beliefs are captured by a distribution μ on $((X_1, Y_1), (X_2, Y_2), \ldots,)$, defined by letting

$$\begin{aligned} &\sim \Pi, \text{ and} \\ &(X_1, Y_1), (X_2, Y_2), (X_3, Y_3), \ldots | \theta \overset{\text{i.i.d.}}{\sim} P_{\theta,\nu}. \end{aligned}$$

Intuitively, the experimenter knows the sequence of random variables is conditionally i.i.d. according to some distribution $P_{\theta,\nu}$. However, she does not know which $P_{\theta,\nu}$ is the true distribution. This uncertainty is captured by Π.

Within this setting, we formalize the concept of "learning the truth" by adopting the notion of Bayesian consistency. The experimenter is *consistent* at $\theta \in \Theta$ if for every neighborhood U of θ, we have

$$\Pi(U | (X_t, Y_t)_{t=1}^T) \to 1 \quad \text{almost surely under } P_{\theta,\nu},$$

where $\Pi(U|(X_t, Y_t)_{t=1}^T)$ is calculated using Bayes' rule. For the case in which the subject does not make mistakes, leading to deterministic data, we need a definition of Bayesian consistency that does not rely on probability. We say that a sequence $(x_t, y_t)_{t=1}^\infty$ is *identified* at θ if for every neighborhood U of θ^*, the following convergence holds:

$$\Pi(U|(x_t, y_t)_{t=1}^T) \to 1.$$

These notions allow us to provide a preference identification result for the Bayesian experimenter in the i.i.d. setting under the same assumption as in Chambers, Echenique and Lambert (2021):

Theorem 2. *Suppose the measure ν on $X \times X$ satisfies $\nu(\{\{x, y\}|x \sim_\theta y\}) = 0$ for all $\theta \in \Theta$. Then there exists $\Theta_\Pi \subseteq \Theta$ with $\Pi(\Theta_\Pi) = 1$ such that the experimenter is consistent at all $\theta \in \Theta_\Pi$.*

4.2 Deterministic Choice

Assume throughout this section that the subject does not make any mistakes—$\ell(\theta, x, y) = 1$ whenever $x \succ_\theta y$—and that our testing condition holds. Proposition 2 implies that BACE will only generate sequences that satisfy

$$\Pi(\{\theta|x_T \succ_\theta y_T\}|(x_t, y_t)_{t=1}^T) = \Pi(\{\theta|y_T \succ_\theta x_T\}|(x_t, y_t)_{t=1}^T) = \frac{1}{2}. \tag{5}$$

Therefore, we say that a sequence $(x_t, y_t)_{t=1}^\infty$ is *BACE-compatible* if it satisfies Equation (5).

Because preferences are deterministic, each choice from a BACE-generated menu systematically eliminates half of the remaining possible preferences based on the current belief. This process results in the prior probability of the support of the posterior being exactly one-half. Therefore, the experimenter's posterior narrows down to the true parameter as fast as $\frac{1}{2^T} \to 0$. The following proposition formalizes this idea.

Proposition 3. *If $(x_t, y_t)_{t=1}^\infty$ is BACE-compatible, given a subject conforming to θ, then $(x_t, y_t)_{t=1}^\infty$ is identified at θ. Moreover, for any T,*

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

To understand how much faster this rate is compared to that from random menu generation, consider the case of a uniform prior Π over $\Theta = [0, 1]$, with $x \succeq_\theta y \implies x \succeq_{\theta'} y$ for all $\theta' \geq \theta$. This assumes the experimenter is dealing with one-dimensional, ordered preferences. For example, X could represent the set of all lotteries over some monetary interval, with $\theta \in \Theta$ indexing utility functions that exhibit constant relative risk aversion:

$$u(x) = \begin{cases} \frac{x^{1-\theta}}{1-\theta} & \text{if } \theta \neq 1 \\ \log(x) & \text{if } \theta = 1. \end{cases}$$

Further assuming the true preference is $\theta^* = 1$, we can calculate experimenter's posterior distribution from generating menus uniformly at random. Specifically, assume the experimenter uses the following menu generating procedure: $\{p, \delta_x\}$ where $x \sim \mathcal{U}(\text{CE}(0), \text{CE}(1))$ and $\text{CE}(\theta)$ is the certainty equivalent of lottery p for preferences θ. Then the posterior would be $[u, 1]$, where $u \sim \mathcal{U}(0, 1)$. Moreover, if the experimenter uses the uniform random menu generator T times, then the posterior becomes uniformly distributed over $[\overline{u}, 1]$, where

$$\overline{u} = \max\{u_1, \dots, u_T\}, \quad u_i \overset{\text{i.i.d.}}{\sim} \mathcal{U}(0, 1).$$

Given that $\mathbb{E}(\overline{u}) = \frac{T}{T+1}$, the expected size of the support reduces to $\frac{1}{T+1}$. This demonstrates a linear convergence for the experimenter's posterior, in contrast to the exponential convergence achievable through BACE.[8]

The above example illustrates a substantial improvement from using BACE relative to random menu generation. However, it leaves open the question of whether BACE achieves the best convergence rate possible.

We proceed to establish that no other menu-generating procedure can lead to faster convergence. Given a full support prior Π_0, define $\mathcal{S}_T$ as the set of all sequences $(x_t, y_t)_{t=1}^T$ such that each sequence conforms to some preference θ. Let $\mathcal{P}_T(\Pi)$ denote the set of all posteriors an experimenter might have after making the subject choose from $t \leq T$ menus. Formally,

$$\mathcal{S}_T = \{(x_t, y_t)_{t=1}^T | x_t \succeq_\theta y_t \text{ for all } t = 1, \dots, T, \theta \in \Theta\}$$
$$\mathcal{P}_T(\Pi_0) = \bigcup_{t \leq T} \{\Pi(\cdot | (x_{t'}, y_{t'})_{t'=1}^t) | (x_{t'}, y_{t'})_{t'=1}^t \in \mathcal{S}_t\}.$$

A menu-generating procedure maps beliefs into lotteries over binary menus using the current information, captured by the posterior, to determine the next menu, possibly at random. More specifically, define $\mathcal{M}_2$ as the set of all binary menus and $\Delta^*(\mathcal{M}_2)$ as the set of finite-

[8] One might worry that $\theta^* = 1$ represents the worst-case scenario for the uniform random menu generator, casting doubt on the rate's applicability to other values of θ^*. However, this concern is unwarranted; the same rate indeed applies across all $\theta^* \in [0, 1]$. The argument for this is less straightforward and thus omitted here for brevity.

support probability measures over these menus. A menu-generating procedure is a function $F\colon \mathcal{P}_T(\Pi) \to \Delta^*(\mathcal{M}_2)$, with the output denoted by F_Π. Let $\mathcal{F}$ represent the set of all such procedures.

As highlighted in previous sections, prevailing procedures generate menus at random, without accounting for subjects' choices. Regardless of the experimenter's beliefs, such procedures would produce the same distribution over menus:

$$F_\Pi = F_{\Pi'} \text{ for all } \Pi, \Pi' \in \mathcal{P}_T(\Pi_0).$$

In contrast, BACE only generates half-space-partitioning menus that directly reflect the experimenter's current belief. As the half-space-partitioning menu need not be unique, BACE defines a family of menu-generating procedures. Formally, let $F^{\mathcal{H}}$ be the set of procedures such that

$$\{x, y\} \in \text{supp}(F_\Pi) \implies \Pi(\theta | x \succ_\theta y) = \Pi(\theta | y \succ_\theta x) = \frac{1}{2} \text{ for all } \Pi \in \mathcal{P}_T(\Pi_0).$$

A tuple (θ, F, Π_0) induces a probability distribution $P_\theta(\cdot, F, \Pi_0)$ over $X \times X$:

$$P_\theta((x, y), F, \Pi_0) = F_\Pi(\{x, y\})\ell(\theta, x, y).$$

By definition, $P_\theta((x, y), F, \Pi_0)$ represents the probability of the experimenter observing the subject choose x over y, given that menus are generated via F under the prior Π_0. Therefore, $P_\theta(\cdot, F, \Pi_0)$ represents the distribution of possible data an experimenter with prior Π_0 may collect after applying the menu-generating procedure F a single time.

Consider a scenario in which the experimenter applies F once and observes (x_1, y_1). If the subject's true preference is θ, then applying F again would result in the distribution $P_\theta(\cdot, F, \Pi_0(\cdot|(x_1, y_1)))$. Notice that both $P_\theta(\cdot, F, \Pi)$ and $P_\theta(\cdot, F, \Pi_0(\cdot|(X_1, Y_1)))$, where $(X_1, Y_1) \sim P_\theta(\cdot, F, \Pi_0)$, jointly define a distribution over $(X \times X)^2$. This probability measure represents the distribution over possible choices an experimenter with prior Π_0 may observe after applying F twice. Iterating the reasoning above to apply F successively T times, we let $P_\theta^T(\cdot, F, \Pi)$ denote the induced probability measure over $(X \times X)^T$.

Given the context provided above, we can now define "fastest rate possible". Our definition is relevant to the current setting, characterized by deterministic choices. An observation (x, y) only affects the experimenter's posterior by assigning probability zero to the set of parameters θ for which $y \succ_\theta x$. A menu-generating procedure F can thus be deemed "better" than G if F rules out a greater mass of preferences compared to G. However, the expected mass reduction may depend on the true preference. A menu-generating procedure F may rule out a larger mass compared to G if θ^* is the true preference, but it may rule out a smaller mass if θ^{**} is the true preference. Since the experimenter does not know the true θ, we posit that F performs better than G if F rules out a higher expected mass of the prior in the worst-case scenario.

Formally, for a given prior Π_0, we say that F achieves the *fastest rate possible* after T choices if

$$\inf_{\theta\in\Theta} \mathbb{E}_{P_\theta^T(\cdot,F,\Pi_0)}[\Pi(\Theta\setminus \text{supp}(\Pi_T))] = \sup_{G\in\mathcal{F}} \inf_{\theta\in\Theta} \mathbb{E}_{P_\theta^T(\cdot,G,\Pi_0)}[\Pi(\Theta\setminus \text{supp}(\Pi_T))]$$

where $\Pi_T = \Pi_0(\cdot|(x_t, y_t)_{t=1}^T)$, and the expectation is taken with respect to the probability measure induced by the preference and the procedure.

Proposition 4. *For any full support prior Π_0, any procedure $F \in F^{\mathcal{H}}$ achieves the fastest rate possible.*

To understand why there cannot exist a procedure that outperforms BACE, consider deterministic procedures that are implemented only once. Suppose that, for a given parameter value θ, a procedure generates a menu $\{x, y\}$ that eliminates more than half of the possible preferences based on the prior. This implies that either $\Pi(\{\theta|x \succeq_\theta y\}) > \frac{1}{2}$ or $\Pi(\{\theta|y \succeq_\theta x\}) > \frac{1}{2}$. In both cases, there is a possible choice that leads the experimenter to rule out less than half of the preference space. Thus, in the worst-case scenario, the evaluation of such a procedure is determined by the choice that does not rule out more than half of the preference space.

4.3 Stochastic Choice

Suppose now that ℓ is continuous in the product metric, with $\ell(\theta, x, y) \in (0, 1)$ for all $x, y \in X$ and $\theta \in \Theta$. In the previous section, we presented convergence results that rely on the primary implication of BACE for deterministic preferences, which is that it generates half-space-partitioning menus. In this section, we extend the results to stochastic choice settings. As before, the results apply not only to BACE but to any procedure that generates half-space-partitioning menus.

Let $\Delta(\Theta)$ denote the set of all full support probabilities on $(\Theta, \mathcal{B}(\Theta))$. Define a *half-space-generating procedure* as a correspondence $\Gamma\colon \Delta(\Theta) \to X \times X$ such that $|\Gamma(\Pi)| = 2$, and $(x, y) \in \Gamma(\Pi)$ implies

1. $\int_\Theta \ell(\theta, x, y)d\Pi = \frac{1}{2}$
2. $(y, x) \in \Gamma(\Pi)$

3. $\Pi(\{\theta \in \Theta : x \succ_\theta y\}) > 0$.

A procedure Γ maps beliefs $\Delta(\Theta)$ into sets of potential observations (i.e., $\{(x, y), (y, x)\}$) that have the primary characteristic of BACE: they partition the belief space into two halves. Given such a procedure, we can construct the set of all possible sequences $(x_t, y_t)_{t=1}^\infty$, denoted Ω, that the experimenter could potentially observe. Moreover, each θ induces a probability measure P_θ over Ω (see Appendix D for details regarding this construction). As a result, the beliefs of an experimenter who plans to employ an *augmented* parametric model $(\Theta, \phi, \ell, \Gamma, \Pi)$ are defined by

$$\mu(E) = \int_\Theta P_\theta(E|_\Theta)d\Pi, \quad E \in \mathcal{S} \times \mathcal{B},$$

where $\Omega \subset (X \times X)^\infty$, $\mathcal{S}$ is a σ-algebra on Ω, $E|_\Theta = \{(x_t, y_t)_{t=1}^\infty : ((x_t, y_t)_{t=1}^\infty, \theta) \in E\}$. For any finite sequence $(x_t, y_t)_{t=1}^T$ with μ-positive probability, two key conditions hold:

1. $P_\theta((x_t, y_t)_{t=1}^T) = \prod_{t=1}^T \ell(\theta, x_t, y_t)$
2. For any $t \leq T$,

$$\int_\Theta \ell(\theta, x_t, y_t)d\Pi(\cdot|(x_{t'}, y_{t'})_{t'=1}^{t-1}) = \frac{1}{2}$$
$$\int_{B_t} \ell(\theta, x_t, y_t)d\Pi(\cdot|(x_{t'}, y_{t'})_{t'=1}^{t-1}) > \frac{1}{2}$$

   for some $B_1, \dots, B_t \in \mathcal{B}$.

Furthermore, for any sequence of data $(x_t, y_t)_{t=1}^T$, the marginal of $\mu(\cdot|(x_t, y_t)_{t=1}^T)$ over Θ is equivalent to the Bayesian posterior of the experimenter $\Pi(\cdot|(x_t, y_t)_{t=1}^T)$.

Theorem 3. *There exists a set Θ_Π with $\Pi(\Theta_\Pi) = 1$ such that the experimenter is consistent at all $\theta \in \Theta_\Pi$.*

The proof strategy of Theorem 3 relies on the Martingale Convergence Theorem. That result guarantees that beliefs will converge to a random variable that may exhibit "path dependency". It suffices to show that this variable is degenerate in the true parameter. To prove this, we leverage the property that BACE only generates half-space-partitioning menus. Roughly speaking, this property ensures that, in the limit, the posterior distribution cannot assign positive weight to more than one parameter. Finally, the fact that the beliefs are well-specified guarantees that the experimenter never rules out the true preference, and thus her beliefs concentrate around the true parameter.

While Theorem 3 guarantees consistency outside of a probability zero event, it remains silent on the rate of convergence. Unfortunately, providing a rate under the above generality remains elusive because of the known difficulties in identifying the rate of convergence for non-i.i.d. Bayesian inference problems. Very few results are known (see Ghosal and Van Der Vaart, 2007).

We can establish a convergence rate under additional assumptions on Θ and ℓ. First, we need Θ to be "ordered" in the sense that if a subject conforming to θ would choose x over y, then so should every θ' "greater than or equal to" θ. An example of a context where such an assumption holds is that of risk preferences. Second, we need the error to be independent of preferences and objects of choice. This implies that every θ is equally likely to make a mistake, regardless of the menu. Under these assumptions, we can derive a bound for the rate at which the experimenter's posterior concentrates on the true parameter.

Theorem 4. *Assume*

- *For all* $\theta \in \Theta$ *and* $x, y \in X$, $\ell(\theta, x, y) = \overline{q} > \frac{1}{2}$ *whenever* $x \succ_\theta y$.
- *There exists a weak order* $\geq$ *over* Θ *such that* $\theta' > \theta$ *implies* $x \succ_\theta y \implies x \succ_{\theta'} y$.

Then there exists a set Θ_Π *with* $\Pi(\Theta_\Pi) = 1$ *such that for any* $\theta \in \Theta_\Pi$ *and open* C *with* $\theta \in C$,

$$\Pi_T(C^c) \leq [2\overline{q}(1 - \overline{q})]^T$$

almost surely under P_θ *for large enough* T.

5 Discussion

Above we restricted the set of preferences that can be part of a parametric model. We assumed Θ is compact and convex and each θ maps into a unique locally strict preference. All of these assumptions are with loss of generality. Specifically, without convexity of Θ and local strictness, we cannot guarantee that our characterization results hold. Moreover, without the compactness of Θ, we cannot guarantee our consistency results hold.

Compactness and convexity of Θ are quite common in empirical settings. Specifically, compactness cannot be avoided as computers can only deal with finite sets. Moreover, priors over Θ are usually assumed to admit a density which requires convexity of Θ. The locally strict preference assumption does merit some discussion. As noted by Chambers, Echenique and Lambert (2021), it generalizes the notion of local non-satiation. Its key implication is that it guarantees that if choice is deterministic, then there exists a countable set of binary menus from which observing the subject make choices would provide enough information to identify the preference. This is a necessary condition for our consistency results to hold. Without it, there is no hope a Bayesian experimenter will learn the true parameter.

We also assumed the experimenter could only offer binary menus. Allowing the experimenter to offer more can help her in learning the preference faster. Suppose the subject's choice is deterministic and the experimenter is allowed to offer menus of three alternatives. Then, the BACE optimization problem would be:

$$\max_{(x,y,z)\in X\times X\times X} \Pi(\{\theta|x \succ_\theta y,z\})\log(\frac{1}{\Pi(\theta|x\succ_\theta y,z)}) + \Pi(\{\theta|y \succ_\theta x,z\})\log(\frac{1}{\Pi(\{\theta|y\succ_\theta x,z\})})$$
$$+\Pi(\{\theta|z \succ_\theta x,y\})\log(\frac{1}{\Pi(\{\theta|z\succ_\theta x,y\})})$$

Proposition 5. *Suppose there exists $x, y, z \in X$ such that*

$$\Pi(\theta|x\succ_\theta y,z) = \Pi(\theta|y\succ_\theta x,z) = \Pi(\theta|z\succ_\theta x,y) = \frac{1}{3}.$$

Then (x, y, z) solves the BACE optimization problem.

The proof of Proposition 5 is analogous to the proof of Proposition 2 and therefore omitted. The above proposition is straightforwardly generalized to menus with finitely many alternatives and has an important implication: Whenever n-space-partitioning menus exist, the experimenter can achieve a rate of $\frac{1}{n^T}$. The existence of such menus is not as rare as one would think. For instance, in the risk aversion setting one can always find n lotteries $p_1, ..., p_n$ such that they are ordered according to second-order stochastic dominance and

$$\Pi(\theta|p_1\succ_\theta p_2,...,p_n) = ... = \Pi(\theta|p_n\succ_\theta p_1,...,p_{n-1}) = \frac{1}{n}.$$

Finally, we conclude this section by discussing how BACE would behave under a different type of data. Experiments such as **?**Wiswall and Zafar (2018) present subjects with binary menus $\{x, y\}$ and ask them with what probability $p_{x,y}$ they would choose x over y. Allowing for such data means that a subject now is identified as a pair (θ, ℓ) where $\ell(\theta, x, y)$ is interpreted as the probability of choosing x the subject would report if she were to be offered a menu $\{x, y\}$.

Let $\mathcal{L} = \{\ell : \Theta\times X\times X \to [0,1] | \ell(\theta,x,y) > \frac{1}{2} \iff x\succ_\theta y$ and $\ell(\theta,x,y) = \frac{1}{2} \iff x\sim_\theta y\}$. The parameter space is now $\Theta\times\mathcal{L}$.

Specifying a prior over $\Theta\times\mathcal{L}$ is no easy task and involves several technicalities that do not provide further insight. To avoid this we make four assumptions:

1. There exists a compact and convex set of parameters $\Psi \subset \mathbb{R}^m$ and a one-to-one continuous function $\gamma : \Psi \to \mathcal{L}' \subset \mathcal{L}$. We write ℓ_ψ to denote $\gamma(\psi)$.

2. The agents prior Π is over $\Theta\times\Psi$ admits a product density $f = f_\Theta \cdot f_\Psi$.

3. For each $\theta, \theta' \in \Theta$ and $x, y \in X$ such that $x \succ_\theta y$ and $x \succ_{\theta'} y$,

$$\int_\Psi 1\{\psi \in \Psi | \ell_\psi(\theta, x, y) = p\} f_\Psi(\psi) d\psi = \int_\Psi 1\{\psi \in \Psi | \ell_\psi(\theta', x, y) = p\} f_\Psi(\psi) d\psi.$$

   for all $p \in (\frac{1}{2}, 1]$.

4. For each $\theta, \theta' \in \Theta$ and $x, y \in X$ such that $x \succ_\theta y$ and $y \succ_{\theta'} x$,

$$\int_0^1 \int_\Psi 1\{\ell_\psi(\theta, x, y) = p\} f_\Psi(\psi) d\psi dp = \int_0^1 \int_\Psi 1\{\ell_\psi(\theta', y, x) = p\} f_\Psi(\psi) d\psi dp$$

The first assumption effectively implies that the experimenter can express her uncertainty over how the subject makes mistakes over a parameter space. The second rules out the possibility that the experimenter has a theory that relates the agent's preference over alternatives to how she commits mistakes. Indeed, it implies that her uncertainty about mistakes and preferences are "independent". The third and fourth are meant to capture the richness of $\mathcal{L}$. The third assumption states that, in terms of the prior, the "size" of the set of error functions for two different preferences is the same that agree on a single menu, is the same. The fourth posits the same conclusion but for two preference that disagree on the menu.

Notice that for a given menu $\{x, y\}$, the agent can report any number between $[0, 1]$. Thus, if the subject reports $p_{x,y}$ when facing menu $\{x, y\}$ the experimenters posterior density is given by

$$f(\theta, \psi | p_{x,y}) = \frac{1\{\ell_\psi(\theta, x, y) = p_{x,y}\} f_\Psi(\psi) f_\Theta(\theta)}{\int_\Theta \int_\Psi 1\{\ell_\psi(\theta, x, y) = p_{x,y}\} f_\Psi(\psi) f_\Theta(\theta) d\psi d\theta}.$$

Then, the BACE objective function can be written as

$$u(x, y) = \int_\Theta \int_\Psi [\int_0^1 1\{\ell_\psi(\theta, x, y) = p\} \log(\frac{1}{\int_\Theta \int_\Psi 1\{\ell_\psi(\theta, x, y) = p\} f_\Psi(\psi) f_\Theta(\theta) d\psi d\theta}) dp] f_\Theta(\theta) f_\Psi(\psi) d\psi \theta$$

Theorem 5. *Assume there exists (x, y) such that*

$$\int_\Theta 1\{\theta | x \succ_\theta y\} f_\Theta(\theta) d\theta = \int_\Theta 1\{\theta | y \succ_\theta x\} f_\Theta(\theta) d\theta = \frac{1}{2}.$$

Then, $(x, y) \in \arg\max_{x,y \in X} u(x, y)$.

This theorem implies that the convergence rate of an experimenter that allows subjects to report probability distributions over alternatives will be the same as one that only observes deterministic choices when subjects do not make mistakes. Intuitively, for each $p_{x,y} > \frac{1}{2}$ there is a "large" number of ℓ's for which $\ell(\cdot, x, y) = p_{x,y}$. Hence, this extra information cannot be used to rule out any θ's that could be ruled out if the experimenter were to only learn x is strictly better than y. Therefore, applying BACE in this setting yields converge to the true parameter as fast as $\frac{1}{2^T} \to 0$.

The previous result relies heavily on assumptions 3 and 4 which reflect an experimenter that considers all possible ℓ's. If instead, we consider an experimenter who restricts her attention to a subset $\mathcal{L}' \subset \mathcal{L}$, then she could indeed get a faster convergence rate. For instance, suppose she only considers $\mathcal{L}' = \{\ell\}$ where ℓ is such that there exists a menu $\{x, y\}$ in which any two different preferences θ, θ' will report different probabilities: $\ell(\theta, x, y) \neq \ell(\theta', x, y)$. Then, from a single report one can perfectly identify the preference, and the convergence rate would be infinity.

6 Other Procedures

In recent years, some papers in economics have adopted adaptive experimental methods that employ concepts different from mutual information to estimate time and risk preferences. Specifically, Imai and Camerer (2018) employs the *Equivalence Class Edge Cutting* (EC^2) and Toubia et al. (2013) a method based on maximizing the Fisher Matrix. Neither of these papers provides a characterization of how the methods work. Indeed, to the best of our knowledge, Theorem 1 is the first characterizing theorem for adaptive experimental methods. Therefore, a theoretical comparison between these methods and BACE is not readily available. Ideally, if characterization of these methods existed we would use them to run a "horse" race between them. Since providing such results is outside of the scope of this paper, we only check if they are half-space-partitioning methods. Details for the claims in this section can be found in Appendix H.

We begin by studying the EC^2 method. Adapting the Δ_{EC^2} formula from Imai and Camerer (2018) to our setting yields

$$\begin{aligned}\Delta_{EC^2}(x,y) = &\left(\int_\Theta \ell(\theta,x,y)f(\theta)d\theta\right)\left(\int_\Theta \left[\ell(\theta,x,y)f(\theta|(x,y))\right]^2 d\theta\right)\\ &+\left(\int_\Theta \ell(\theta,y,x)f(\theta)d\theta\right)\left(\int_\Theta \left[\ell(\theta,y,x)f(\theta|(y,x))\right]^2 d\theta\right)\end{aligned}$$

As with BACE, inspection of the objective function does reveal it is a half-space-partitioning method. However, as opposed to BACE, the EC^2 method turns out to not be one in general. To illustrate, assume the subject does not make mistakes and consider the case in which $\Theta = [0,1]$ and $\Pi \sim U[0,1]$. Assume further that Θ is ordered in the sense that $x \succ_\theta y \implies x \succ_{\theta'} y$ for all $\theta' \geq \theta$. It turns out that under these assumptions, **any** pair (x,y) such that $x \succ_\theta y$ for some $\theta \in (0,1)$ maximizes Δ_{EC^2}. This means that the rate of convergence entirely depends on the tie-breaking rule the experimenter adopts. Moreover, it shows it is not a half-space-partitiong method.

A similar observation can be made for the Fisher matrix method. Adapting the objective function of Toubia et al. (2013) to our setting yields

$$F(x,y) = \left(\int_\Theta \ell(\theta,x,y) f(\theta) d\theta \right) \|H_{(f(\cdot|(x,y))}(\hat{\theta})\| + \left(\int_\Theta \ell(\theta,y,x) f(\theta) d\theta \right) \|H_{(f(\cdot|(y,x))}(\hat{\theta})\|$$

where H_f denotes the Hessian of f, $\hat{\theta}$ is the mode of f, and $\|\cdot\|$ is a norm (absolute value of the determinant).

Roughly, their method suggests to generate the menu that maximizes the expected value of the norm of the Hessian evaluated at the mode. Toubia et al. (2013) justify this by noting that it has been shown that under general conditions, the asymptotic covariance matrix of the Maximum Likelihood Estimator (MLE) is equal to the inverse of the Hessian of the log-likelihood function evaluated at the MLE. Therefore, reducing the asymptotic covariance matrix of the MLE is achieved by maximizing some norm of the Hessian of the likelihood function. Toubia et al. (2013) "apply" this insight to the Bayesian setting and use the Hessian of the posterior distribution at its mode instead of the MLE. They justfy this by noting that mode of the posterior distribution becomes the standard MLE in the case of a uniform prior or if data goes to infinity.

It is easy to see that maximizing F need not yield a half-space-partitioning menu. Consider the same example as with the EC^2 criteria. Under uniform prior and no mistakes, the Hessian of the posterior is zero. Hence, $F(x,y)$ is constant. This observation generalizes to the case in which mistakes are twice continuously differentiable. The reason is that under uniform prior, every parameter is the mode.

7 Incentive Compatibility

Overall we have shown that BACE works and has desirable properties. However, we have implicitly assumed that the subject always behaves truthfully in the sense that she is not strategic. In this section, we discuss to what degree this assumption is without loss.

Consider a subject who understands the inner workings of BACE. This effectively means that she knows all the sequences $(x_t, y_t)_{t=1}^T$ that are BACE-compatible in the sense that

$$(x_t, y_t) \in \arg\min_{(x_t,y_t)\in\mathcal{H}_t} \int_\Theta (\ell(\theta, x_t, y_t)\log(\ell(\theta, x_t, y_t)) + \ell(\theta, y_t, x_t)\log(\ell(\theta, y_t, x_t)))\, d\Pi_t$$
$$\mathcal{H}_t = \{(x, y)| \int_\Theta \ell(\theta, x, y) d\Pi_t = \frac{1}{2}\}$$

for all t.

If the experimenter were to tell the subject that she would get the last choice (x_T) as a reward for participating in the experiment, then the subject would have an incentive to lie. To illustrate, consider our risk aversion example without mistakes: $\Theta = [0, 1]$, $\theta^* = 1$, $\Pi \sim \mathcal{U}[0, 1]$, and

$$u(x) = \begin{cases} \frac{x^{1-\theta}}{1-\theta} & \text{if } \theta \neq 1 \\ \log(x) & \text{if } \theta = 1 \end{cases}.$$

Assume the experimenter uses BACE to choose the certainty equivalent of a lottery p from $(\text{CE}(1), \text{CE}(0))$. Then, if the subject chooses $T - 1$ times the lottery, the last menu she would be offered would be

$$\{p, \text{CE}(\frac{1}{2^T})\}.$$

However, if instead she answered truthfully the last menu she would be offered would be

$$\{p, \text{CE}(1 - \frac{1}{2^T})\}.$$

Because $\text{CE}(\cdot)$ is decreasing, the subject would strictly prefer $\text{CE}(\frac{1}{2^T})$ over $\text{CE}(1 - \frac{1}{2^T})$ and therefore has an incentive to lie.

Although one may view this example as a criticism of BACE, it relies entirely on the fact that the experimenter is offering the last choice as the reward for participation. We suggest the experimenter employs a random reward system based on BACE. To describe the system, we need a few preliminaries.

Given a BACE-compatible sample $(x_t, y_t)_{t=1}^T$, let Π_T be the corresponding posterior. Define $\hat{\theta}$ as the expectation of Π_T restricted to the set of θ's such that

$$\int_\Theta \ell(\theta, x_T, y_T) d\Pi_T = \frac{1}{2}.$$

Notice that if Θ is one-dimensional, then the experimenter does not need to take an expectation.

Moreover, if $\Pi(\cdot|(x_t, y_t)_{t=1}^T) \to \delta_{\theta^*}$ weakly, then $\hat{\theta} \to \theta^*$.

The reward system we propose works as follows: The experimenter tells the subject that after answering T questions, she will estimate her preferences using $\hat{\theta}$ and then generate a menu randomly and offer the reward that $\succeq_{\hat{\theta}_T}$ would choose from the menu. Moreover, if the experimenter wants to offer T rewards, then she can generate randomly T menus and then offer the choices that $\succeq_{\hat{\theta}_T}$ would make.

To show why our method is incentive-compatible consider a probability measure ν on $X \times X$. Then,

$$\rho_\nu(\succ_{\hat{\theta}} \setminus \succ_{\theta^*}) = \int_{X \times X} 1\{(x, y) \in \succ_{\hat{\theta}_T} \setminus \succ_{\theta^*}\} d\nu(x, y)$$

can be interpreted as the probability that a subject will get an alternative she would have not chosen from a random menu. Because the experimenter does not know θ^*, then she needs to choose ν in such a way that

$$\rho_\nu(\succ_{\theta'} \setminus \succ_\theta) = \int_{X \times X} 1\{(x, y) \in \succ_{\theta'} \setminus \succ_\theta\} d\nu(x, y)$$

is increasing in $d(\theta, \theta')$ for all $\theta, \theta' \in \Theta$. If she can do so, then the subject has the incentive to report accurately her preferences as it will make $\hat{\theta}$ the closest to θ^*.

How to construct ν depends on the specific setting. For instance, in our risk aversion example it is trivial. Indeed, one can just use a uniform random menu generator: $\{p, \delta_x\}$ where $x \sim \mathcal{U}(\text{CE}(0), \text{CE}(1))$. To see why it would work, fix $\theta, \theta', \theta'' \in (0, 1)$ such that $|\theta - \theta'| < |\theta - \theta''|$. Assume WLOG that $\theta' > \theta$. Let $\nu(x, y) = \frac{1}{2}(p, \delta_x) + \frac{1}{2}(\delta_x, p)$ where $x \sim \mathcal{U}(\text{CE}(0), \text{CE}(1))$. There are two cases: (1) $\theta'' > \theta'$ and $\theta'' < \theta$. In case 1,

$$P_\nu(\succ_{\theta'} \setminus \succ_\theta) = \theta' - \theta < \theta'' - \theta = P_\nu(\succ_{\theta''} \setminus \succ_\theta).$$

In case 2,

$$P_\nu(\succ_{\theta''} \setminus \succ_\theta) = \theta - \theta'' > \theta' - \theta = P_\nu(\succ_{\theta'} \setminus \succ_\theta).$$

More generally the following strategy can be applied: Fix an increasing sequence $(\epsilon_i)_{i=1}^n$ such that $\Theta \subset B_{\epsilon_n}(\theta)$ for all $\theta \in \Theta$. Given $\hat{\theta}$, find a sequence $(x_i, y_i)_{i=1}^n$ such that $\theta \in B_{\epsilon_i}(\hat{\theta})$ then $(x_i, y_i) \in \succeq_\theta$ and $(x_i, y_i) \notin \succ_{\theta'}$ for all $\theta' \in \Theta \backslash B_{\epsilon_i}(\hat{\theta})$. Take ν such that it samples uniformly from $(x_i, y_i)_{i=1}^n$. This procedure would then imply $P_\nu(\succ_{\theta'} \setminus \succ_\theta)$ is increasing in $d(\hat{\theta}, \theta)$. Hence, it would incentivize the agent to report truthfully as it would minimize $d(\hat{\theta}, \theta^*)$.

8 Extension to Uncertain Error

So far we have assumed the experimenter knows exactly how the subject makes mistakes given her preferences and is only uncertain about the subeject's preference. In this section, we discuss to what extend our results generalize to the case in which she is also uncertain about the error distribution.

We consider a setting similar to the one with probabilistic data. Specifically, we assume:

1. There exists a compact and convex set of parameters $\Psi \subset \mathbb{R}^m$ and a one-to-one continuous function $\gamma : \Psi \to \mathcal{L}' \subset \mathcal{L}$. We write ℓ_ψ to denote $\gamma(\psi)$.

2. The agents prior Π is over $\Theta \times \Psi$ admits a product density $f = f_\Theta \cdot f_\Psi$.

As in the probabilistic data setting, assumption one is made for tractability. Assumption two is made to rule out the possibility that the experimenter has a theory that relates the preferences to the errors.

Under these assumptions, the experimenter's posterior marginal density over preferences is given by

$$f_\Theta(\theta|(x,y)) = \frac{\int_\Psi \ell_\psi(\theta,x,y) f_\Psi(\psi) d\psi f_\Theta(\theta)}{\int_\Theta \int_\Psi \ell_\psi(\theta,x,y) f_\Psi(\psi) f_\Theta(\theta) d\psi d\theta}$$

whereas the marginal over errors is

$$f_\Psi(\psi|(x,y)) = \frac{\int_\Theta \ell_\psi(\theta,x,y) f_\Theta(\theta) d\theta f_\Psi(\psi)}{\int_\Psi \int_\Theta \ell_\psi(\theta,x,y) f_\Theta(\theta) d\theta f_\Psi(\psi) d\psi}$$

Therefore, a half-space partitioning-menu after (x,y) would satisfy

$$\begin{aligned} \frac{1}{2} &= \int_\Theta \int_\Psi \ell_\psi(\theta,x,y) f_\Psi(\psi|(x,y)) f_\Theta(\theta|(x,y)) d\psi d\theta \\ &= \int_\Theta \ell_{\overline{\psi}|(x,y)}(\theta,x,y) f_\Theta(\theta|(x,y)) d\theta \end{aligned}$$

where

$$\ell_{\overline{\psi}|(x,y)}(\theta,x,y) = \int_\Psi \ell_\psi(\theta,x,y) f_\Psi(\psi|(x,y)) d\psi.$$

This means that an experimenter that uses such a procedure will still be generating menus in which "half" of the preferences prefer one alternative to the other, the only difference is in the ℓ function. Indeed, with uncertain errors, it is as if the ℓ function were path-dependent.

However, as in many applied settings, if ℓ_ψ is assumed to not depend on (x, y), then such path-dependence partially disappears. Indeed, the fact that half-space-partitioning menus are being employed would not affect the experimenter's posterior over errors. Hence, as long as the error parameter is identified and ℓ is continuous in the product metric, we can obtain as a Corollary to Theorem 3 that the experimenter will learn the true preference.

9 BACE Implementation

There are a series of computational challenges that come with implementing adaptive online survey experiments at scale. It is compuationally intensive to solve the complex optimization problem that determines the optimal next-best scenario. This problem is compounded when the adaptive procedure is conducted flexibly for a wide range of subjects beyond the lab who may have poor Internet speeds or devices with limited computational power. We overcome these challenges and implement the method with thousands of survey respondents.

The computational difficulties can be seen from the formula in Equation (1). We are faced with a multi-dimensional integration problem that does not have a straightforward analytical solution. The general framework allows for complete flexibility over the space of scenarios, answers, and preference parameters, but the main challenge remains computational. In practice, the design and answer spaces can be discretized; for example, DCEs typically employ two answer choices given the simplicity of that framework for respondents, but our implementation of BACE can accommodate any number of discrete answer choices. The number of parameters that you want to estimate and the size of the design space then determine the complexity of the numerical integration problem.

When determining how to design software to deliver BACE, two main decisions must be made. We must select an optimization method for selecting designs and a method for calculating the posterior distribution over preference parameters. With four or fewer preference parameters, which can accommodate many standard applications, a grid-based optimization approach can work well. However, the size of the grid (and resulting computational burden) increases exponentially with the number of design paramers. We employ Bayesian Optimization, a userful tool for optimizing expensive objective functions, to select optimal designs.

Oftentimes, when calculating posterior estimates of an individual's preferences, the posterior is estimated using a predefined grid that covers the support of the prior. However, the required size of this grid grows exponentially as you estimate more preference parameters; moreover, resulting parameter estimates must be a convex combination of points on the initial grid, and the researcher cannot reliably estimate parameters outside of the initial grid.

Monte Carlo methods, therefore, can offer better solutions for parameter spaces with higher dimensions and can produce estimates outside of the initial grid (Press et al., 2007). In our current implementation, we use Population Monte Carlo (PMC), an adaptive importance sampling algorithm, to estimate the posterior distribution (Cappé et al., 2004); PMC combines ideas from importance sampling and sequential Monte Carlo methods to estimate a stationary target distribution (Smith, 2013).

Computing the next-best scenario and displaying it to the respondent quickly is particularly important within the context of online survey experiments, where researchers have limited control over participants' behavior and attention. Moreover, attrition caused by slow loading times presents a major concern. There are two approaches that make it possible to serve the next-best scenario to respondents within a reasonable time frame. On one hand, the researcher can precompute a decision tree that maps out the full range of questions that can be shown to a participant. This process requires computational resources upfront and communication between the survey interface and the look-up tree. Once the tree is computed, questions can quickly be delivered to respondents based on their previous answers. Practically, precomputing the decision tree is feasible for small T as the tree will have 2^T branches. However, even small changes to the decision environment would require the researcher to recompute the entire tree, which makes this method less flexible to changes. An alternative method computes the next-best scenario in real-time, which is feasible as long as such computation can be carried out within seconds. This method is more flexible as it can quickly accommodate changes to the design or parameter spaces and can handle experiments with many questions. Our current implementation uses the latter method but can be easily modified to accommodate the former.

An important innovation our package provides is the ability to perform all computations using backend servers and databases with simultaneous communication between the survey platform and the servers. This allows any user with reliable access to the Internet, regardless of the computing power of their device, to participate in one of our experiments. Furthermore, we use modern global optimization techniques to facilitate computation speeds. We develop code that allows a survey platform (e.g., Qualtrics) to interact with cloud-based backend servers that can compute the next-best scenario simultaneously for multiple subjects; Figure 2 provides a schematic summarizing the interactions between the survey platform, backend server, and database in our implementation. This enables a fully dynamic elicitation that can be administered outside of university laboratory settings. Our method can also be adapted to work with alternative platforms; for example, surveys could be delivered using text messages, which would remove the requirement that users be connected to the Internet. This would expand the reach of our framework to run experiments in other important contexts, such as developing countries or remote regions.

9.1 Design Optimization

The primary difficulty in performing adaptive experimentation in real-time is that computing the optimal design is computationally intensive. In each round, we must calculate the design that maximizes the mutual information: $d^*_{t+1} = \arg\max U(d_{t+1})$. Optimizing this function requires estimating a multi-dimensional integral that can be non-convex and rarely has an analytical solution. Typically, researchers select the optimal question by performing a brute force estimation over a grid of possible designs. However, a brute force search is computationally slow and inefficient, particularly when the researcher is considering a large design space that covers multiple parameters. Other approaches assume a functional form over the distribution of θ, which leads to an analytical solution (Paninski, 2005) or precompute the entire decision tree, which delivers designs quickly but requires computing 2^T branches and is infeasible for a large number of questions(Chapman et al., 2018). Vincent and Rainforth (2017) use an algorithm that allocates computational resources towards designs that are expected to perform better; the algorithm resembles a Thompson sampling or multi-armed bandit optimization.

Our approach to optimizing design selection is to use Bayesian Optimization—a sample efficient, sequential approach for optimizing expensive objective functions. Bayesian Optimization is a useful tool for optimizing black box functions, and it is frequently used for selecting hyperparameters in machine learning models or designing expensive experiments. It scales better than grid-based search, with standard Bayesian Optimization algorithms performing well for up to 15-20 parameters. Moreover, Bayesian Optimization provides the flexible optimization approach required to handle an objective function that depends on likelihood functions that we allow to take on any functional form. This feature is particularly important when designing a package that can be used flexibly by researchers asking a variety of questions. We provide an overview of the intuition behind Bayesian Optimization here. [9]

Bayesian Optimization works by constructing and optimizing a surrogate model of the true objective function using Gaussian process regression.[10] To start, a Gaussian process prior is assumed over the objective function. Initial data is constructed by evaluating the objective function at a small number of designs. A Gaussian process is fit to this data, and a well-behaved acquisition function is optimized to select the next design at which to evaluate

[9]See Shahriari et al. (2016) for a recent review of common applications of Bayesian Optimization and the theory, history, and intuition underlying this method. Frazier (2018) provides a helpful tutorial for understanding how Bayesian Optimization works.

[10]In our context, the objective function corresponds to the mutual information function defined in Equation (1).

the objective function. The objective function is evaluated at this design, and a new model is fit to the updated data. This process is repeated for a number of iterations, and the design that maximizes the objective function is returned.

Our current software implementation uses the Python package Mango to perform Bayesian Optimization within our application (Sandha et al., 2020, 2021). The user simply needs to specify values for each design parameter; which can be continuous distributions, a range of integers, or a discrete set of categories. Bayesian Optimization presents a flexible design optimization strategy that can handle the variety questions researchers hope to answer using BACE.

9.2 Computing the Posterior Distribution

This section describes our method for performing Bayesian Inference to estimate the posterior distribution based on an individual's history of responses. The researcher starts by defining a prior over the parameters of interest, $p(\theta)$. This prior can be informed by previous results from the literature, pilot experiments, or it can cover a generous range of possible estimates.[11] In round t, a question is selected and shown to the respondent. Based on the observed answer to d_t, we calculate the posterior estimate using Bayes' rule, which is given by Equation (2). At each stage of the experiment, we are not interested in a point estimate of the posterior; instead, we want to carry forward a population of samples that represents the full posterior distribution.

Typical adaptive experimental approaches use a grid search method or make functional form assumptions over the distribution of θ. However, grid search methods scale poorly with the number of parameters to be estimated and can produce biased estimates. In particular, standard grid search methods can only assign positive weight to parameter values that are a convex combination of points in the initial grid and suffer from particle degeneracy. To overcome some of the issues with grid search, we use a method called population Monte Carlo (PMC) to estimate the posterior distribution based on each question and answer. PMC is a sequential Monte Carlo method that uses adaptive importance sampling (AIS) to estimate a target distribution (Cappé et al., 2004).[12] Importantly, PMC is unbiased at each iteration and does not require the burn-in period or stopping rules typical for many Markov Chain Monte Carlo methods (Robert and Casella, 2004).

Let N be the number of particles sampled from the posterior distribution. We start by

[11] Importantly, the support of the prior distribution determines the range of possible posterior estimates, so researchers should ensure that the prior they specify covers the range of interest.

[12] Vincent and Rainforth (2017) develop an adaptive experimental framework for estimating time and risk preferences and also use PMC to estimate the posterior distribution. For a recent review of PMC and other AIS techniques, see Bugallo et al. (2017).

sampling $\theta = \{\theta_1, \ldots, \theta_N\}$ where each θ_i is drawn from the prior distribution $p(\cdot)$. A design is shown to the respondent, and the preferred option is chosen and recorded. Thus, $(d^{1:t}, x^{1:t})$ make up our data after period t. Our goal is to produce a population of particles that is sampled from the posterior distribution $\Pr\big(\theta \,\big|\, d^{(1:t)}, x^{(1:t)}\big)$. Since we cannot sample directly from this posterior distribution, we sample from a set of proposal functions, $q_i(\cdot)$ and use importance sampling techniques to correct for the fact that we are sampling from an alternate distribution.

The basic algorithm for PMC, which converges at a rate of $O(N^{-1/2})$, is described below.[13] The goal is to form an estimate of the posterior beliefs about an individual's preference parameters given their previous answers and a prior distribution over preferences: $\Pr(\theta \mid d^t, x^t)$. At time t, we have data $x^{(1:t)}$ and $d^{(1:t)}$, which constitute the observed responses for an individual. To estimate the posterior distribution, begin by sampling N initial points from the prior distribution over θ. For each point $\theta_{i,1}$, sample from a multivariate normal distribution centered at $\theta_{i,1}$ to get $\theta^{'}_{i,1}$. Importance weights are calculated as $w_{i,j} = \frac{l(x^{(1:t)} | d^{(1:t)}, \theta^{'}_{i,j}) \times p(\theta^{'}_{i,j})}{q(\theta^{'}_{i,j} | \theta_{i,j})}$. Weights are normalized to sum to one, and points $\theta^{'}_{i,j}$ are resampled with replacement using the normalized weights $w_{i,j}$. This process is repeated for J iterations.

While PMC produces an unbiased estimate of the posterior distribution at each iteration, the posterior estimates will be more precise as J and N increase. After each question and answer, the procedure produces a sample that is as if it was sampled from the posterior distribution $\Pr(\theta \mid d^{1:t}, x^{1:t})$. The proposal function governs the importance sampling process; we use a multivariate normal distribution centered at each existing point in the distribution. Importance sampling techniques require that the proposal distribution has fatter tails than the posterior distribution of interest.

In our application, PMC offers advantages over grid search, traditional Monte Carlo, and Markov Chain Monte Carlo (MCMC) methods. Under grid search, an initial grid is chosen, and the posterior is calculated by updating the weights on each point in the grid using Bayes' rule. This method offers no exploration; only points that are found in the initial grid are propagated forward. Thus, if the prior distribution is poorly specified and the individual's preferences are extreme, then only a limited number of grid points are used to calculate the resulting posterior.

In contrast, PMC uses importance sampling to explore points around the initial grid, which is particularly important if the prior distribution is poorly specified.[14] Under PMC, the

[13] This description captures the basic PMC algorithm from Cappé et al. (2004). For a comprehensive discussion of the key properties of PMC and recent adaptations of PMC algorithms, see Bugallo et al. (2017).

[14] When designing a software framework that is user-friendly across a variety of settings, exploration is particularly important as it makes the calculation of the posterior more robust to poorly specified priors. Nevertheless, researchers should be careful when defining the prior distribution to ensure that the support of

population of points will adapt to be located around regions of the posterior with high posterior likelihoods. By resampling with replacement, PMC also helps avoid particle degeneracy, which occurs when only a handful of points in the grid have non-negligible importance weights. Many MCMC methods also require a specified stopping rule or a burn-in period. In contrast, PMC is unbiased at each iteration and is, thus, preferred for our application, where we have a small number of effective iterations and are limited by computation time.

9.3 Monte Carlo Estimation of the Mutual Information

In this section, we describe how we estimate the mutual information using nested Monte Carlo. Recall the objective function in our Bayesian Optimization algorithm is given by Equation (1).

We want to estimate the mutual information for the next question:

$$U(d_{t+1}) = \int_\theta \int_{x^{t+1}} \{\log \frac{\Pr\big(\theta; x^{t+1} \,\big|\, x^{(1:t)}, d_{t+1}\big)}{\Pr\big(\theta \,\big|\, x^{(1:t)}, d_{t+1}\big) \Pr\big(x^{1:(t+1)} \,\big|\, d_{t+1}\big)}\} \Pr\big(\theta; x^{t+1} \,\big|\, x^{(1:t)}, d_{t+1}\big) dx^{t+1} d\theta \tag{6}$$

This can be rearranged into the form:

$$\begin{aligned} U(d_{t+1}) = &\int_{x^{t+1}} \int_\theta \Pr(x^{t+1} \,|\, \theta; d_{t+1}) \Pr\big(\theta \,\big|\, x^{(1:t)}\big) \log(\Pr(x^{t+1} \,|\, \theta; d_{t+1})) dx^{t+1} d\theta \\ &- \int_{x^{t+1}} \Pr(x^{t+1} \,|\, d_{t+1}) \log(\Pr(x^{t+1} \,|\, d_{t+1})) dx^{t+1} \end{aligned} \tag{7}$$

We estimate the integral above using nested Monte Carlo estimations. Let $\theta_i = \theta_1, \ldots, \theta_N$ be sampled from the posterior distribution $\Pr\big(\theta \,\big|\, x^{(1:t)}, d^{(1:t)}\big)$ after t questions. We use θ_i to estimate the integral above as follows:[15]

$$\begin{aligned} U(d_{t+1}) \approx &\frac{1}{N} \sum_{x_j} \sum_{\theta_i} \Pr\big(x_j \,\big|\, \theta_i, d_{t+1}\big) \log\big(\Pr\big(x_j \,\big|\, \theta_i, d_{t+1}\big)\big) \\ &- \sum_{x_j} \big(\frac{1}{N} \sum_{\theta_i} \Pr\big(x_j \,\big|\, \theta_i, d_{t+1}\big)\big) \log\left(\frac{1}{N} \sum_{\theta_i} \Pr\big(x_j \,\big|\, \theta_i, d_{t+1}\big)\right) \end{aligned} \tag{8}$$

where $N = |\theta|$ and $\Pr\big(x_j \,\big|\, \theta_i, d_{t+1}\big)$ is the likelihood of observing answer x_j to question d_{t+1}

the prior covers the range of the parameters of interest.

[15] Since we consider a discrete set of answer choices, we leverage the fact that probabilities sum to one to facilitate computation. If we have J discrete answer choices, then we need only compute the likelihood function for the first $J-1$ options. Thus, $\Pr\big(x_j \,\big|\, \theta_i, d\big) = 1 - \sum_{x_{k \neq j}} \Pr(x_k \,|\, \theta_i, d)$, and a similar result follows for the term inside of the log.

given the estimate of the individual's preferences at time t. For a discrete number of answer choices, the convergence rate of this estimator is $O(N^{-1})$, the standard error rate for Monte Carlo methods (Rainforth et al., 2018). Our package currently hands any number of discrete answer choices; however, we note that this can be extended to accommodate a continuous answer space.

9.4 Backend Framework

Ultimately, we want to create an experimental framework that can flexibly accommodate a large number of users in real-time. To make performance independent of the user's device, we use a backend server to handle all major computation. The current version of our software package creates a web application that is hosted remotely by Amazon Web Services (AWS).[16] A Python-based Flask application is hosted on AWS Lambda; a backend Amazon DynamoDB database stores profile-specific information that includes respondents' survey responses and treatment characteristics. The application can be set up to communicate with any standard survey framework (e.g. Qualtrics or Survey Monkey), deliver questions directly to respondents, or be adapted to work with any device that has an Internet or cellular connection.[17] By moving computation off of the user's device, the time that users face between questions depends minimally on the quality of the device they are using to take the survey, which may be correlated with income or other socio-demographic characteristics. Moreover, cloud services can be easily scaled, and the researcher can increase server-side resources to account for high numbers of users at launch.

The package architecture and process for selecting new questions is described in Figure 2. When a user starts the survey, a profile is created in the database that will store relevant information for that user. On the server, we compute the optimal design based on the history of designs that the user has seen. We store this in the database and transmit the selected design to the survey platform. The respondent views the question and chooses their preferred option. The respondent's ID and answer are sent back to the server application, which updates the posterior estimate using PMC. Based on the new posterior estimate, the next optimal design is chosen and sent to the respondent. This process is repeated for the duration of the experiment. At the end of the experiment, we have data that constitutes the full design history $d^{(1:T)}$, answer history $x^{(1:T)}$, and an estimate of the posterior distribution $\Pr\big(\theta \,\big|\, x^{(1:T)}, d^{(1:T)}\big)$. This distribution can be used to directly produce posterior estimates when the survey is complete. More generally, however, estimates can be recomputed after

[16]While we host the application using AWS, fhe framework can be adapted to alternative cloud providers.
[17]The framework can also be used to run experiments directly on a computer in a laboratory setting as well.

the experiment using alternative techniques, such as maximum likelihood estimation, on the data generated during the experiment.

10 Simulation validation

Our current simulations highlight two notable econometric advantages of the BACE method to elicit preferences. First, BACE allows for a much higher precision of the parameter estimates with fewer scenarios presented to each subject. Second, the standard approach of estimating the average preference parameter by pooling together scenarios and answers across subjects may result in biased estimates, which can be avoided when taking the average of the individual-level preferences from a BACE procedure.

We compare across two methods to generate the sequence of scenarios: the adaptive Bayesian approach (BACE), and a procedure where sequences are randomly generated (RAND). We also interact each method to generate the scenarios with one of two evaluation approaches: using Bayesian updating, and using a maximum likelihood estimator (MLE). Given the answer choices and the set of scenarios, one can Bayesian update the parameter distribution using Equation (2), or one can pool together the data and find the estimate that maximizes the likelihood of the data obtained based on Equation (9).

10.1 Efficiency

We start with the case when the two job scenarios only differ by one amenity. We first specify the utility function that determines choice. In each hypothetical scenario, two jobs $j \in \{0, 1\}$ are presented that the subject can choose from. The jobs consist of earnings y_j, and whether amenity a_j is at the base value ($a_j = 0$) or the alternative value ($a_j = 1$). Utility from job j is $u_j = \log(y_j) + \beta_i a_{ij}$. Willingness to pay for the alternative over the base value of amenity a_i as a fraction of y_0 can be easily derived to be $\exp(\beta_i) - 1$.

Without choice inconsistency, the individual always chooses the bundle with the higher utility. Since inconsistency may arise in practice, we consider two cases for modeling choices. In case 1, the probability of making a "mistake" is higher when the two bundles are closer in total utility; in this case, "mistake" is represented by an error term in the utility function that has a Gumbel distribution with scale parameter β (lower β represents higher consistency). In case 2, with a fixed probability p, the individual chooses randomly instead of choosing the higher-utility bundle (lower p represents higher consistency). In this case, subjects choose the job with the higher utility, but with $p \in [0, 1]$ chance of being inattentive, which we define as the probability of picking a choice at random instead.[18] The probability of choosing $x = j$ is then

$$\Pr(x = j \,|\, \theta \equiv \{\beta_i\}, d \equiv \{j, 1-j\}) = (1-p)\mathbb{1}_{\left\{\log(y_j/y_{1-j}) + \left(\beta_i(a_{ij} - a_{i(1-j)})\right) > 0\right\}} + p/2 \quad (9)$$

The formulas above can be easily extended to incorporate multiple amenities or interaction terms between amenity values.

Figure 3 shows that regardless of the evaluation method, the adaptive Bayesian approach yields more information about the utility function parameter faster. When the likelihood of being inattentive is below 50 percent, the procedure can lead to relatively precise estimates with 5 to 10 questions. Figure 4 shows that the correlation between the true parameters and the estimated ones approaches 1 much faster with the adaptive method. Between Bayesian updating and MLE, Bayesian updating has the edge, and it performs significantly faster.

In the next set of simulations, we consider a binary choice experiment under two different utility models, each with four parameters. The two utility models consist of the same preference parameters but differ based on how choice inconsistency operates. In both cases, the preference parameters are of the form $u_j = \log(y_j) + \beta_1 a_{1j} + \beta_2 a_{2j} + \gamma a_{1j} a_{2j}$ where j is the index of a chosen bundle of attributes described by $\left(y_j, a_{1j}, a_{2j}\right)$, $y \in \mathbb{R}$, $a \in \{0, 1\}$. Without choice inconsistency, the individual always chooses the bundle with the higher utility. Since inconsistency may arise in practice, we consider two cases for modeling choices as before.

In the two-amenity case, the story is much similar qualitatively, with the adaptive Bayesian approach gaining ground even more relative to the random approach. Figures 5 to 8 show that the amenity value coefficients continue to perform well with 10 questions or fewer. Figure 5 also shows that the Bayesian estimation procedure gives more precise estimates compared to using a standard maximum likelihood estimator, in which a numerical optimization procedure is required to compute the estimates. We also note that BACE yields a substantial improvement in mean squared error for estimating the interaction term when more scenarios are asked when there is a higher degree of choice consistency, but this is not the case for RAND (see Figures 6 and 8). The interaction coefficient also shows significant gains from using BACE relative to RAND. Figures 9 and 10 show that one would need about 15 to 20 questions to achieve high precision with the adaptive Bayesian approach, and would need at least 3 times more questions with the random approach for the same level of precision. According to the simulations, presenting 10 scenarios using BACE gives more precise estimates than presenting 50 or more scenarios in a randomly generated sequence.

[18] This formation mirrors Mas and Pallais (2017) who model an across-subject inattentive rate in a similar way.

Given that one can trade off between the interaction term and either of the two amenity coefficients, it is clear that the interaction term is harder to identify. The simulation results also show that Bayesian updating performs more consistently than MLE.

The online Appendix further shows that the method is robust to different ways to model inattentiveness. For example, even if the individual makes more mistakes when utility difference is smaller, for example, we are still able to recover the utility parameters well.

10.2 Bias

In Figure 11, we evaluate the importance of obtaining individual-level preferences even when the object of interest is the average preference in the population. We find that the standard approach of pooling together responses across individuals to estimate the average preference results in possible bias, even with large samples when the sample mean follows a normal distribution. This is true regardless of whether one uses scenarios and answers from BACE or RAND.

This is reminiscent of the mean-variance confound described in Yatchew and Griliches (1985). Intuitively, individuals may have heterogeneous tendencies to make inconsistent choices, and it is difficult to account for individual heterogeneity in choice inconsistency when combining all individual data in the estimation with a combined error term, leading to bias.

To elaborate, consider the simple case when each individual n make choices based on the following data generating process. The latent variable for choice i is $u_i = \alpha_n w_i + \beta_n z_i + \epsilon_i$, with w_i and z_i randomly drawn and ϵ_i being independent and identically distributed according to a logistic distribution. The outcome variable is $y_i = \mathbb{1}_{\{u_i>0\}}$. Because of the normalization involved (either of the variance of the error term, or one of the coefficients), we are interested in β_n/α_n for each individual n. An alternative normalization is $u_i = w_i + \frac{\beta_n}{\alpha_n} z_i + \frac{1}{\alpha_n}\epsilon_i$. Now u_i is measured in units of w_i and $\frac{1}{\alpha_n}$ is the scale of the error term (inconsistency in choices).

Assume that $i = 1, \ldots, I$ data points are collected for each individual n. Running a logit regression of y_i on w_i and z_i for each individual n should result in consistent estimates of α_n and β_n. However, when I is small, we sometimes pool together all data points across individuals and estimate y_i on w_i and z_i to obtain estimates α and β. When α_n and β_n vary across individuals: α and β do not recover the average of α_ns and average of β_ns, nor can we recover the average β_n/α_n. In some cases, the pooled estimate may even be outside the range of individual-level parameters. Of course, if α_n is the same across individuals, then we do recover the average of the β_ns

References

Aigner, Dennis J. 1979. "A brief introduction to the methodology of optimal experimental design." *Journal of Econometrics*, 11(1): 7–26. 3

Andersen, Steffen, Glenn W Harrison, Morten Igel Lau, and E Elisabet Rutström. 2006. "Elicitation using multiple price list formats." *Experimental Economics*, 9(4): 383–405. 1

Aucejo, Esteban M, Jacob F French, and Basit Zafar. 2021. "Estimating Students' Valuation for College Experiences." National Bureau of Economic Research. 2

Avivi, Hadar, Patrick Kline, Evan Rose, and Christopher Walters. 2021. "Adaptive Correspondence Experiments." Vol. 111, 43–48. 3

Bliemer, Michiel CJ, and John M Rose. 2011. "Experimental design influences on stated choice outputs: an empirical study in air travel choice." *Transportation Research Part A: Policy and Practice*, 45(1): 63–79. 1

Bugallo, Monica F., Victor Elvira, Luca Martino, David Luengo, Joaquin Miguez, and Petar M. Djuric. 2017. "Adaptive Importance Sampling: The past, the present, and the future." *IEEE Signal Processing Magazine*, 34(4): 60–79. 28, 29

Bullock, Craig H, Mark Scott, and Menelaos Gkartzios. 2011. "Rural residential preferences for house design and location: insights from a discrete choice experiment applied to Ireland." *Journal of environmental Planning and Management*, 54(5): 685–706. 1

Cappé, O, A Guillin, J. M Marin, and C. P Robert. 2004. "Population Monte Carlo." *Journal of computational and graphical statistics*, 13(4): 907–929. 26, 28, 29

Caria, Stefano, Maximilian Kasy, Simon Quinn, Soha Shami, Alex Teytelboym, et al. 2020. "An adaptive targeted field experiment: Job search assistance for refugees in jordan." 3

Carson, Richard T. 2012. "Contingent valuation: A practical alternative when prices aren't available." *Journal of economic perspectives*, 26(4): 27–42. 2

Carson, Richard T, and Mikołaj Czajkowski. 2014. "The discrete choice experiment approach to environmental contingent valuation." In *Handbook of choice modelling*. Edward Elgar Publishing. 1

Cavagnaro, Daniel R, Gabriel J Aranovich, Samuel M McClure, Mark A Pitt, and Jay I Myung. 2016. "On the functional form of temporal discounting: An optimized adaptive test." *Journal of Risk and Uncertainty*, 52(3): 233–254. 3

Cavagnaro, Daniel R, Richard Gonzalez, Jay I Myung, and Mark A Pitt. 2013. "Optimal decision stimuli for risky choice experiments: An adaptive approach." *Management science*, 59(2): 358–375. 3

Chaloner, Kathryn, and Isabella Verdinelli. 1995. “Bayesian experimental design: A review.” *Statistical Science*, 273–304. 6

Chambers, Christopher P, Federico Echenique, and Nicolas S Lambert. 2021. “Recovering preferences from finite data.” *Econometrica*, 89(4): 1633–1664. 3, 11, 12, 17

Chapman, Jonathan, Erik Snowberg, Stephanie Wang, and Colin Camerer. 2018. “Loss attitudes in the US population: Evidence from dynamically optimized sequential experimentation (DOSE).” National Bureau of Economic Research. 3, 27

Czajkowski, Mikołaj, Tomasz Gajderowicz, Marek Giergiczny, Gabriela Grotkowska, and Urszula Sztandar-Sztanderska. 2020. “Choosing the future: economic preferences for higher education using discrete choice experiment method.” *Research in Higher Education*, 61(4): 510–539. 1

Eriksson, Tor, and Nicolai Kristensen. 2014. “Wages or fringes? Some evidence on trade-offs and sorting.” *Journal of Labor Economics*, 32(4): 899–928. 1

Feld, Brian, AbdelRahman Nagy, and Adam Osman. 2020. “What Do Jobseekers Want? Comparing Methods to Estimate Reservation Wages and the Value of Job Attributes.” 1, 2

Frazier, Peter I. 2018. “A tutorial on Bayesian optimization.” *arXiv preprint arXiv:1807.02811*. 27

Gelblum, Madeleine. 2020. “Preferences for job tasks and gender gaps in the labor market.” *Job Market Paper*. 1

Ghosal, Subhashis, and Aad Van Der Vaart. 2007. “Convergence rates of posterior distributions for noniid observations.” *The Annals of Statistics*, 35(1): 192–223. 17

Glumac, Brano, and Thomas P Wissink. 2018. “Homebuyers’ preferences concerning installed photovoltaic systems: A discrete choice experiment.” *Journal of European Real Estate Research*. 1

Green, Paul E, Abba M Krieger, and Yoram Wind. 2001. “Thirty years of conjoint analysis: Reflections and prospects.” *Interfaces*, 31(3_supplement): S56–S73. 1

Ida, Takanori, and Rei Goto. 2009. “Interdependency among addictive behaviours and time/risk preferences: Discrete choice model analysis of smoking, drinking, and gambling.” *Journal of Economic Psychology*, 30(4): 608–621. 1

Imai, Taisuke, and Colin F Camerer. 2018. “Estimating time preferences from budget set choices using optimal adaptive design.” *Unpublished manuscript*. 3, 6, 20

Jeuland, Marc, Marcelino Lucas, John Clemens, and Dale Whittington. 2009. “A cost–benefit analysis of cholera vaccination programs in Beira, Mozambique.” *The World Bank Economic Review*, 23(2): 235–267. 1

Kasy, Maximilian, and Anja Sautmann. 2021. "Adaptive treatment assignment in experiments for policy choice." *Econometrica*, 89(1): 113–132. 3

Maestas, Nicole, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger. 2018. "The value of working conditions in the united states and implications for the structure of wages." National Bureau of Economic Research. 1, 2, 4

Mas, Alexandre, and Amanda Pallais. 2017. "Valuing alternative work arrangements." *American Economic Review*, 107(12): 3722–59. 1, 2, 33

Mas, Alexandre, and Amanda Pallais. 2019. "Labor supply and the value of non-work time: Experimental estimates from the field." *American Economic Review: Insights*, 1(1): 111–26. 1

Moffatt, Peter G. 2007. "Optimal experimental design in models of decision and choice." *Measurement in Economics*, 357. 3

Neustadt, Ilja, and Peter Zweifel. 2011. "Income redistribution: how to divide the pie?" 1

Paninski, Liam. 2005. "Asymptotic theory of information-theoretic experimental design." *Neural Computation*, 17(7): 1480–1507. 6, 27

Peirce, Charles S. 1967. "Note on the Theory of the Economy of Research." *Operations Research*, 15(4): 643–648. 3

Picasso, Emilio, and Mark A Cohen. 2019. "Valuing the public's demand for crime prevention programs: a discrete choice experiment." *Journal of experimental criminology*, 15(4): 529–550. 1

Press, William H, Saul A Teukolsky, William T Vetterling, and Brian P Flannery. 2007. *Numerical Recipes 3rd Edition: The Art of Scientific Computing.* Cambridge University Press. 26

Rainforth, Tom, Robert Cornish, Hongseok Yang, Andrew Warrington, and Frank Wood. 2018. "On Nesting Monte Carlo Estimators." *International Conference on Machine Learning (ICML).* 31

Robert, Christian P., and George Casella. 2004. *Monte Carlo statistical methods. Springer texts in statistics.* Second ed., New York:Springer. 28

Ryan, Elizabeth G, Christopher C Drovandi, James M McGree, and Anthony N Pettitt. 2016. "A review of modern computational algorithms for Bayesian optimal design." *International Statistical Review*, 84(1): 128–154. 6

Ryan, Mandy, Karen Gerard, and Mabel Amaya-Amaya. 2007. *Using discrete choice experiments to value health and health care.* Vol. 11, Springer Science & Business Media. 1

Sandha, Sandeep Singh, Mohit Aggarwal, Igor Fedorov, and Mani Srivastava. 2020. "Mango: A Python Library for Parallel Hyperparameter Tuning." 3987–3991, IEEE. 28

Sandha, Sandeep Singh, Mohit Aggarwal, Swapnil Sayan Saha, and Mani Srivastava. 2021. "Enabling Hyperparameter Tuning of Machine Learning Classifiers in Production." IEEE. 28

Schulz, Norbert, Gunnar Breustedt, and Uwe Latacz-Lohmann. 2014. "Assessing farmers' willingness to accept "greening": Insights from a discrete choice experiment in Germany." *Journal of Agricultural Economics*, 65(1): 26–48. 1

Shahriari, Bobak, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando de Freitas. 2016. "Taking the Human Out of the Loop: A Review of Bayesian Optimization." *Proceedings of the IEEE*, 104(1): 148–175. 27

Shannon, Claude Elwood. 1948. "A mathematical theory of communication." *The Bell system technical journal*, 27(3): 379–423. 5

Smith, Adrian. 2013. *Sequential Monte Carlo methods in practice.* Springer Science & Business Media. 26

Toubia, Olivier, Eric Johnson, Theodoros Evgeniou, and Philippe Delquié. 2013. "Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters." *Management Science*, 59(3): 613–640. 3, 6, 20, 21

Vincent, Benjamin T, and Tom Rainforth. 2017. "The DARC Toolbox: automated, flexible, and efficient delayed and risky choice experiments using Bayesian adaptive design." 27, 28

Wald, Abraham. 1950. "Statistical decision functions." 3

Wiswall, Matthew, and Basit Zafar. 2018. "Preference for the workplace, investment in human capital, and gender." *The Quarterly Journal of Economics*, 133(1): 457–507. 1, 2, 4, 18

Yatchew, Adonis, and Zvi Griliches. 1985. "Specification error in probit models." *The Review of Economics and Statistics*, 134–139. 34

Figure 1: BACE selection procedure of discrete choice scenarios

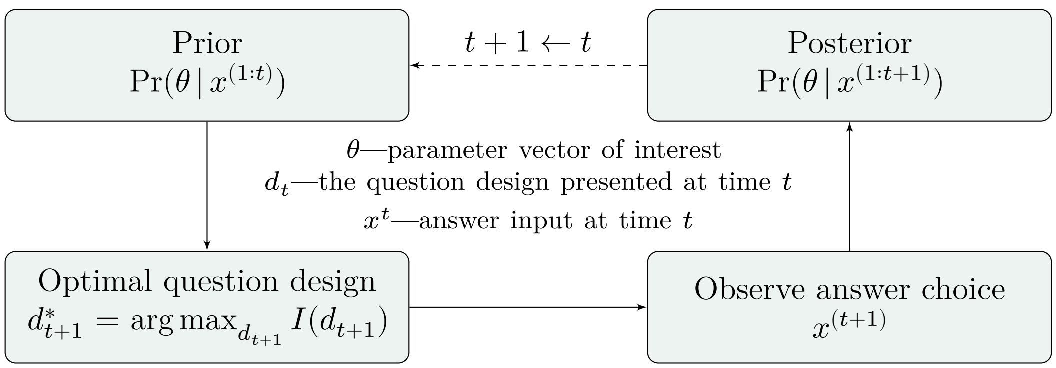

Note: See Section 2 for the notation. The initial prior for the parameter vector θ, initialized at $t = 0$, can be a uniform prior over the parameter space or chosen according to a pilot experiment. The Bayesian adaptive procedure chooses the question design that maximizes the *mutual information* between the parameter random value Θ and the outcome random value $X^{(t+1)}$, i.e., $I(d_{t+1}) := I(\Theta; X^{t+1}|d_{t+1})$, and then updates with the respondent's answer choice to the chosen question. The new data are used to update the posterior of θ using Bayes' rule, $\Pr\big(\theta \,\big|\, x^{(1:t+1)}, d_{t+1}\big) = \frac{\Pr(x^{t+1} \mid \theta, d_{t+1}) \Pr(\theta \mid x^{(1:t)})}{\int_{\theta'} \Pr(x^{t+1} \mid \theta', d_{t+1}) \Pr(\theta' \mid x^{(1:t)}) d\theta'}$, and the updated posterior is used as the prior in the next round.

Figure 2: Implementation schematic

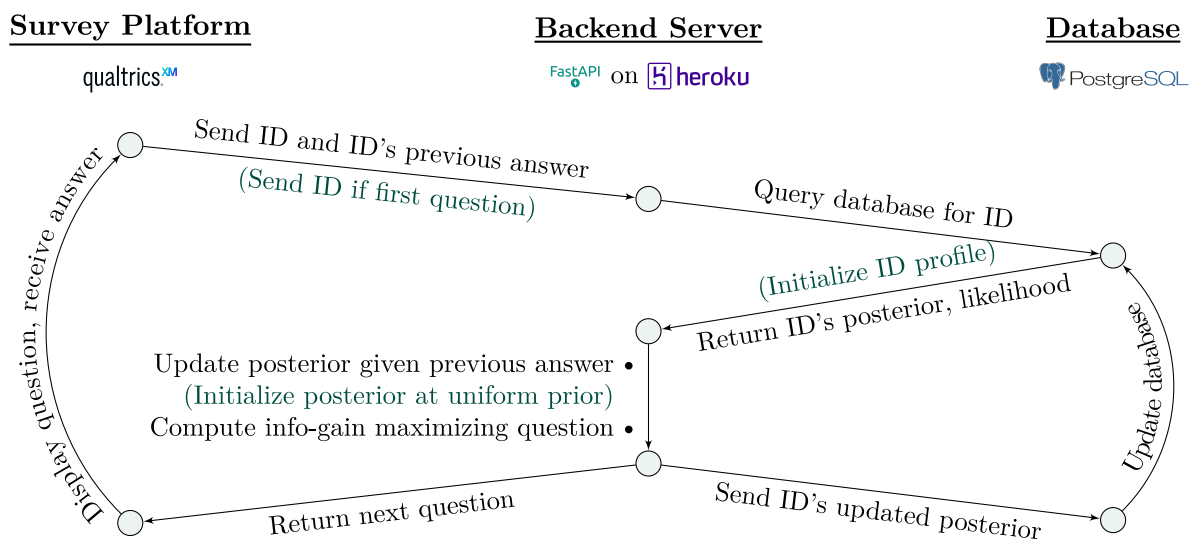

Note: A simplified schematic of the interactions between the survey platform and the backend computation and database servers.

Figure 3: Recovering amenity value coefficient when one amenity is presented

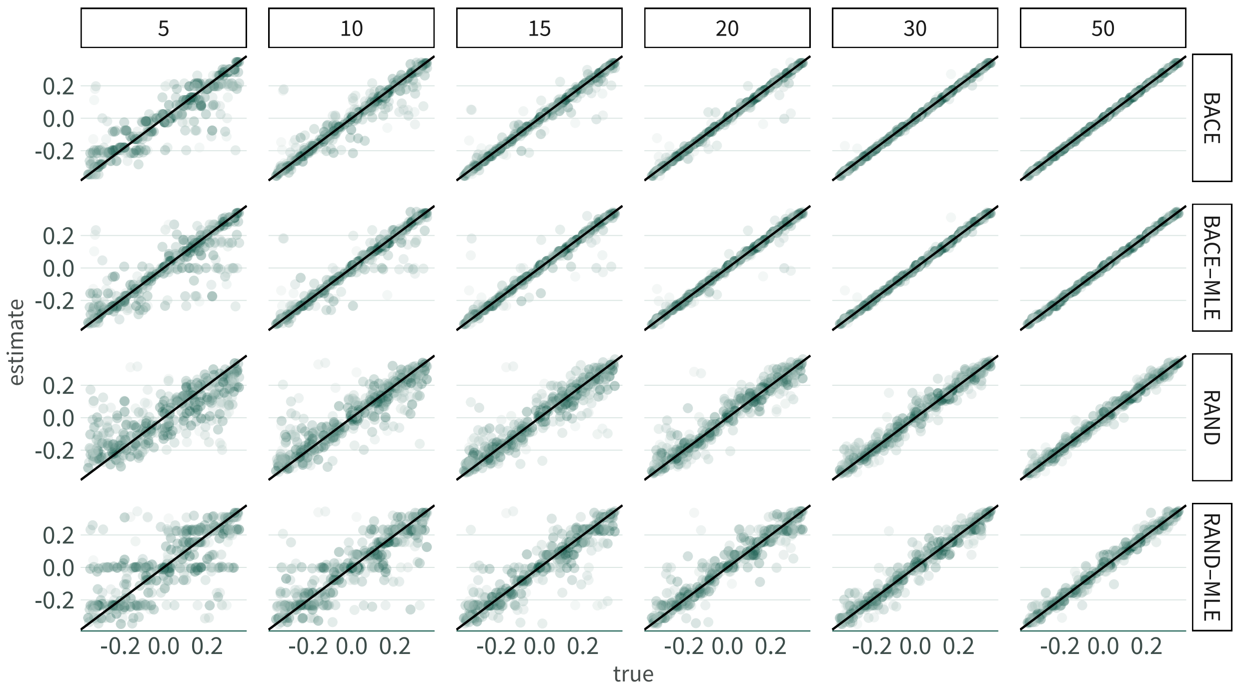

Note: The figure depicts the simulation results in the case when a binary choice is presented between two jobs which differ along two dimensions: earnings and the presence or absence of one amenity.
The y-axis is the true parameter value on the amenity coefficient.
The x-axis is the estimated parameter value on the amenity coefficient using four methods: when the sequence of questions are generated by the Bayesian Adaptive Choice Experiment (BACE) or randomly (RAND), and when the coefficients are recovered using the Bayesian approach or using a maximum likelihood estimator (MLE).
The lighter color corresponds to higher value of true p, the parameter that corresponds to the probability of choosing randomly rather than choosing the choice with higher utility.
Rows are the four methods (see above).
Columns are the estimates after the number of questions asked.

Figure 4: Average mean squared error between true and estimated amenity value coefficients when one amenity is presented

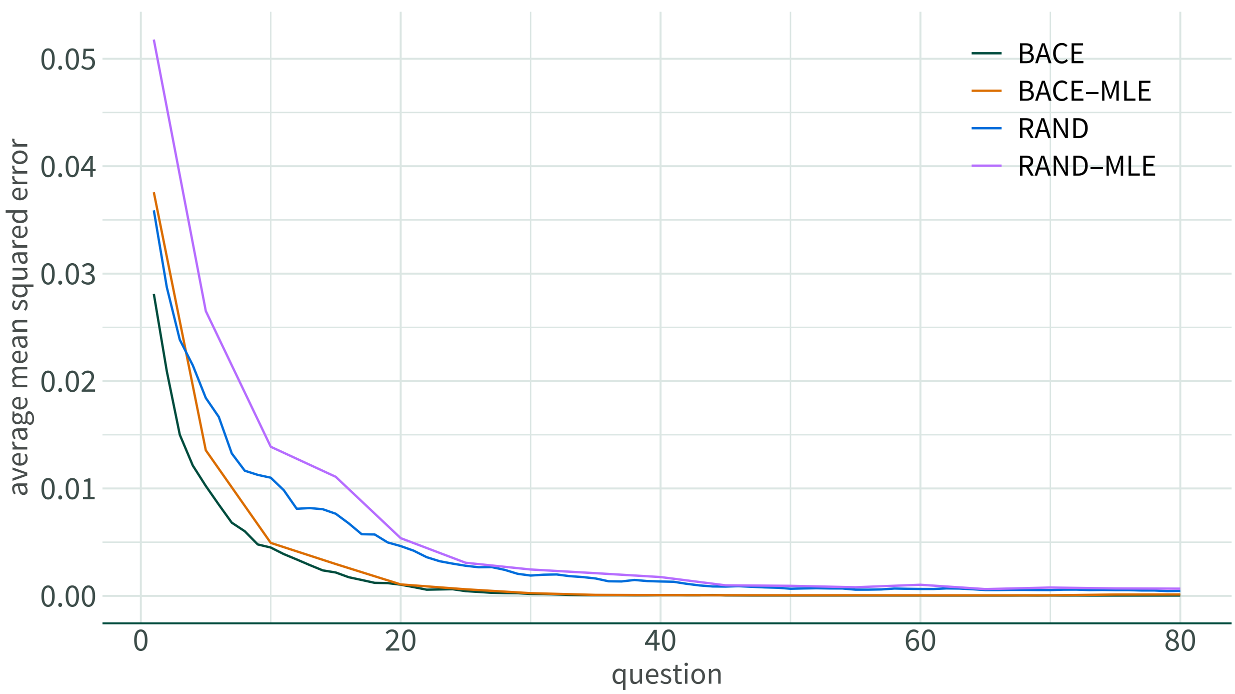

Note: The figure depicts the simulation results in the case when a binary choice is presented between two jobs which differ along two dimensions: earnings and the presence or absence of one amenity.
The x-axis is the number of questions used to obtain amenity coefficient estimates.
The y-axis is the average mean squared error between the estimates and the true values from the simulations.
The colors map to four methods: when the sequence of questions are generated by the Bayesian Adaptive Choice Experiment (BACE) or randomly (RAND), and when the coefficients are recovered using the Bayesian approach or using a maximum likelihood estimator (MLE).

Figure 5: Recovering first amenity value coefficient when two amenities are presented

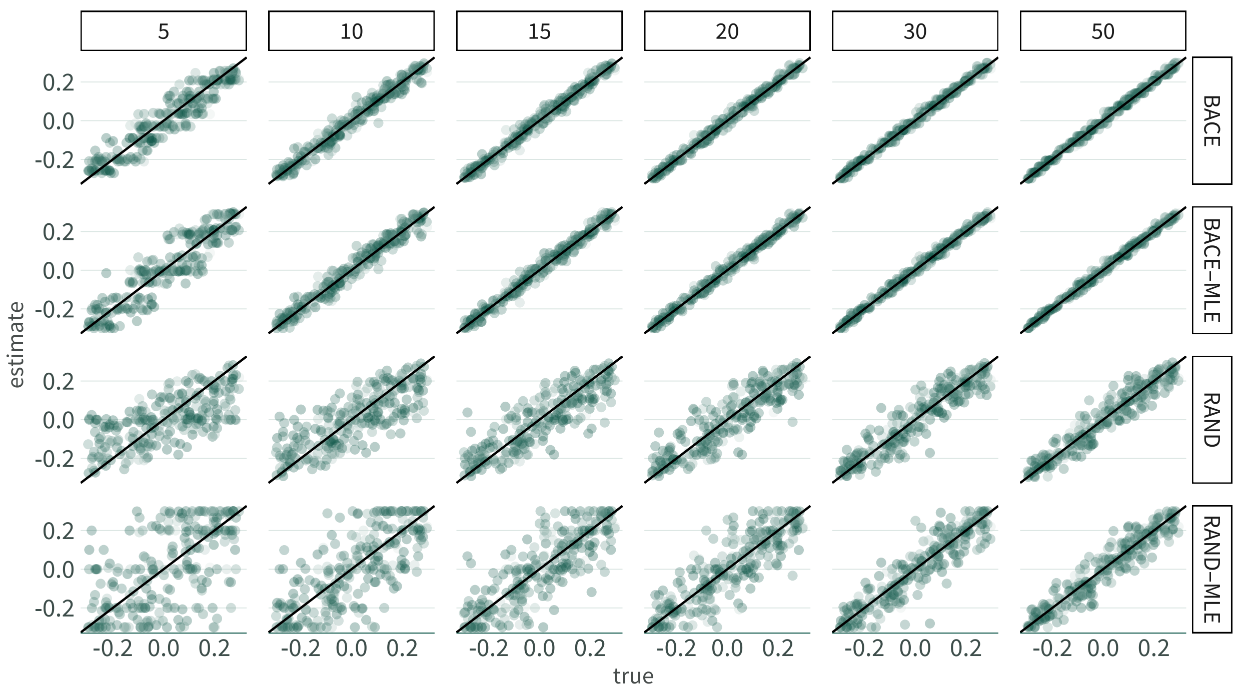

Note: The figure depicts the simulation results in the case when a binary choice is presented between two jobs which differ along three dimensions: earnings and the presence or absence of each of two amenities.
The y-axis is the true parameter value on the first amenity coefficient.
The x-axis is the estimated parameter value on the first amenity coefficient using four methods: when the sequence of questions are generated by the Bayesian Adaptive Choice Experiment (BACE) or randomly (RAND), and when the coefficients are recovered using the Bayesian approach or using a maximum likelihood estimator (MLE).
The lighter color corresponds to higher value of true p, the parameter that corresponds to the probability of choosing randomly rather than choosing the choice with higher utility.
Rows are the four methods (see above).
Columns are the estimates after the number of questions asked.

Figure 6: Average mean squared error between true and estimated first amenity value coefficients when two amenities are presented

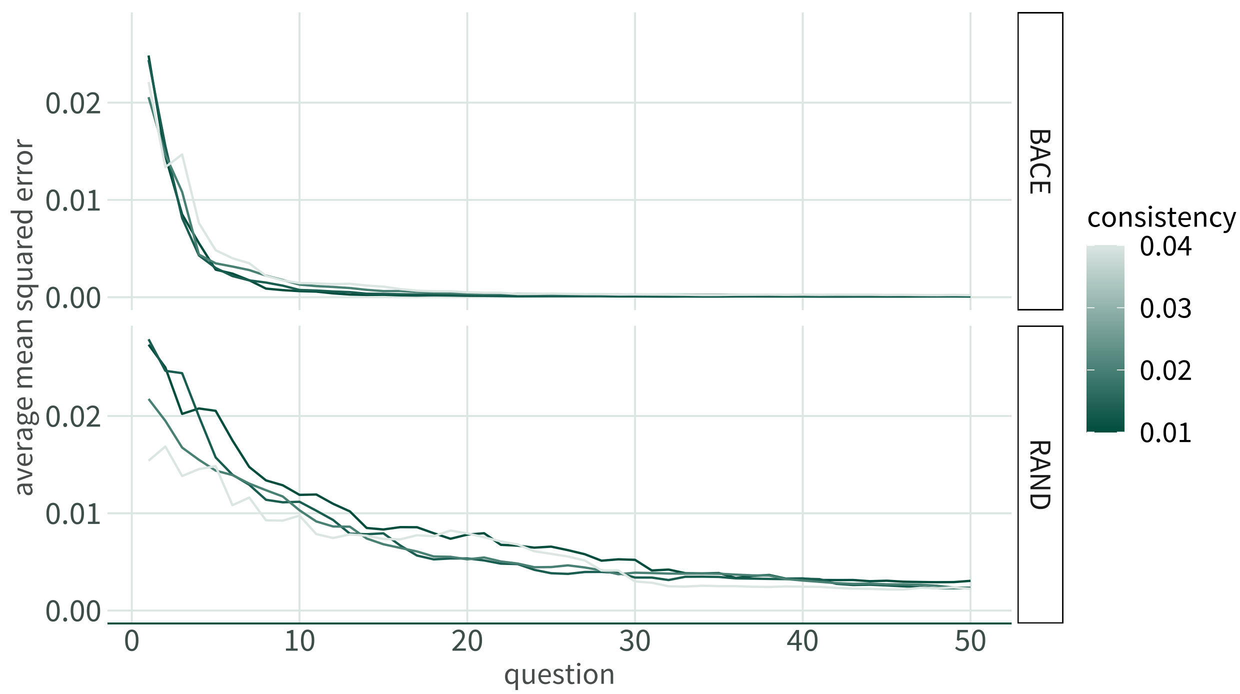

Note: The figure depicts the simulation results in the case when a binary choice is presented between two jobs which differ along three dimensions: earnings and the presence or absence of each of two amenities. The x-axis is the number of questions used to obtain amenity coefficient estimates.
The y-axis is the average mean squared error between the estimates and the true values from the simulations. The panels map to two methods: when the sequence of questions are generated by the Bayesian Adaptive Choice Experiment (BACE) or randomly (RAND). The colors correspond to different values of the choice inconsistency parameter.

Figure 7: Recovering second amenity value coefficient when two amenities are presented

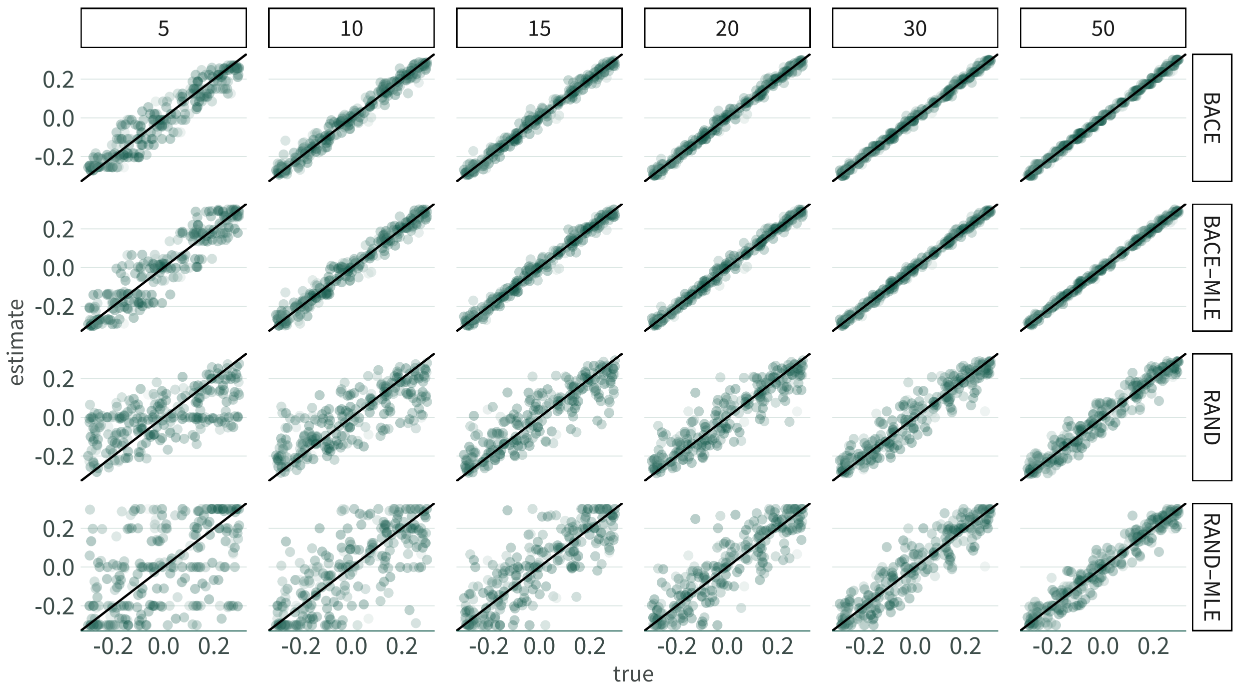

Note: The figure depicts the simulation results in the case when a binary choice is presented between two jobs which differ along three dimensions: earnings and the presence or absence of each of two amenities.
The y-axis is the true parameter value on the second amenity coefficient.
The x-axis is the estimated parameter value on the second amenity coefficient using four methods: when the sequence of questions are generated by the Bayesian Adaptive Choice Experiment (BACE) or randomly (RAND), and when the coefficients are recovered using the Bayesian approach or using a maximum likelihood estimator (MLE).
The lighter color corresponds to higher value of true p, the parameter that corresponds to the probability of choosing randomly rather than choosing the choice with higher utility.
Rows are the four methods (see above).
Columns are the estimates after the number of questions asked.

Figure 8: Average mean squared error between true and estimated second amenity value coefficients when two amenities are presented

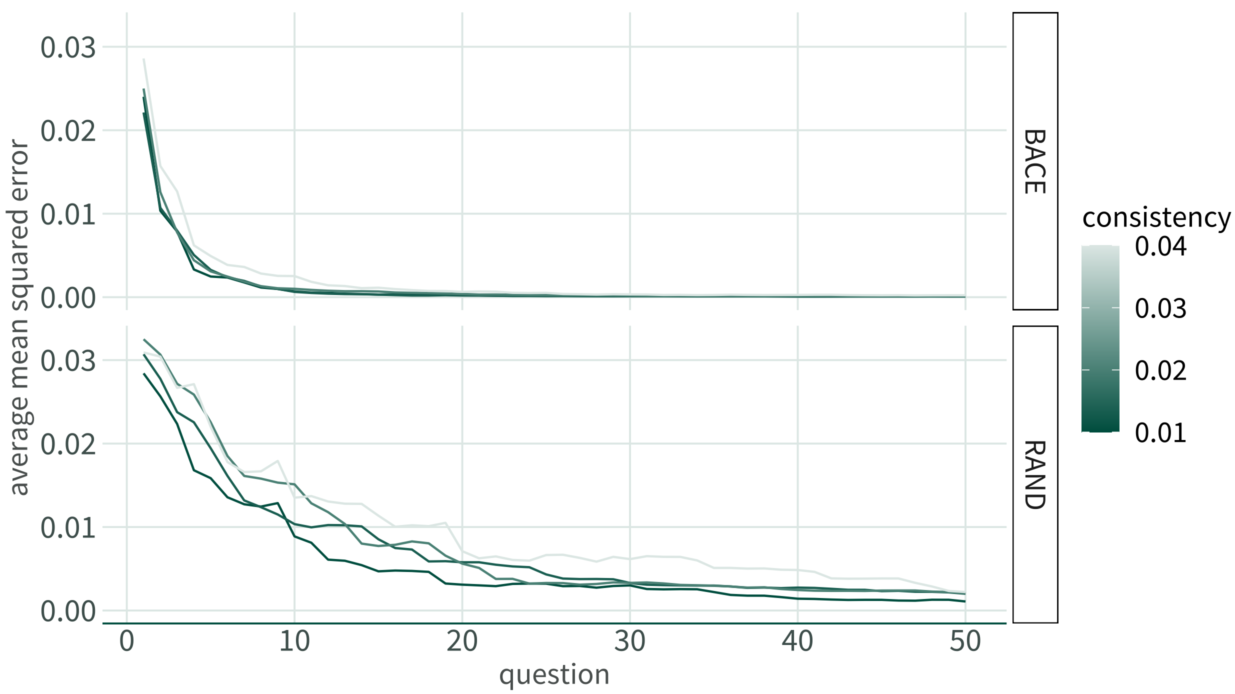

Note: The figure depicts the simulation results in the case when a binary choice is presented between two jobs which differ along three dimensions: earnings and the presence or absence of each of two amenities. The x-axis is the number of questions used to obtain amenity coefficient estimates.
The y-axis is the average mean squared error between the estimates and the true values from the simulations. The panels map to two methods: when the sequence of questions are generated by the Bayesian Adaptive Choice Experiment (BACE) or randomly (RAND). The colors correspond to different values of the choice inconsistency parameter.

Figure 9: Recovering the interaction coefficient when two amenities are presented

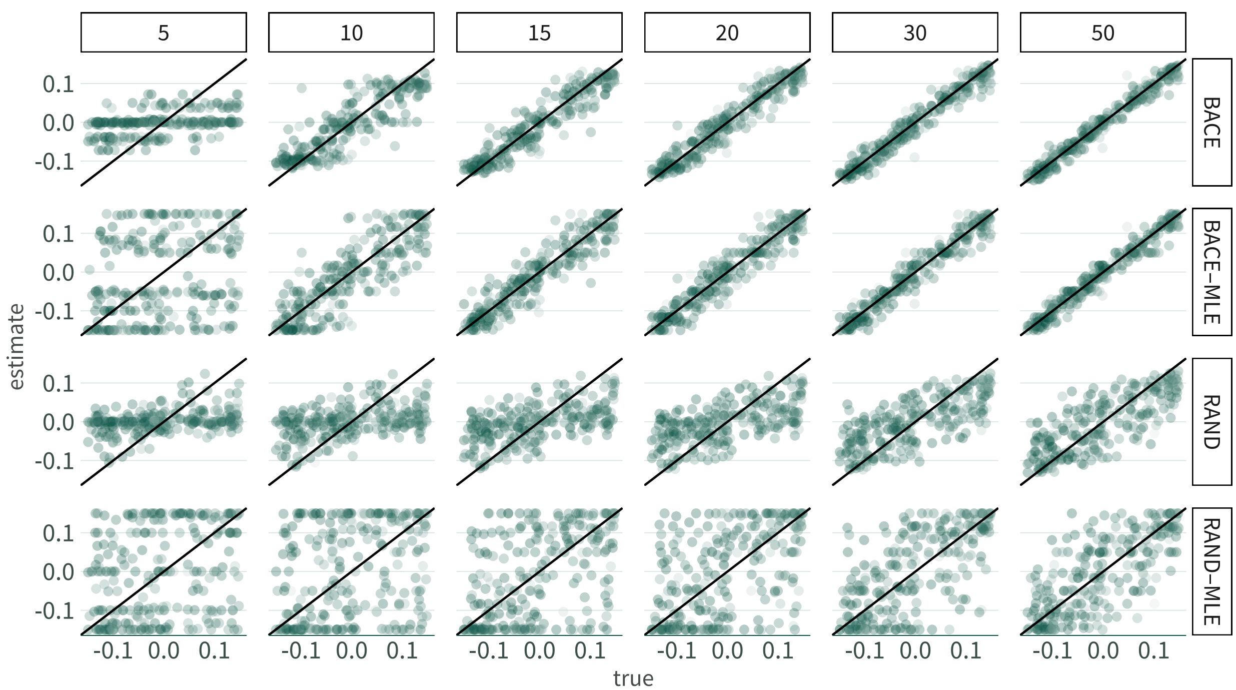

Note: The figure depicts the simulation results in the case when a binary choice is presented between two jobs which differ along three dimensions: earnings and the presence or absence of each of two amenities.
The y-axis is the true parameter value on the interaction term between the two amenities.
The x-axis is the estimated interaction coefficient using four methods: when the sequence of questions are generated by the Bayesian Adaptive Choice Experiment (BACE) or randomly (RAND), and when the coefficients are recovered using the Bayesian approach or using a maximum likelihood estimator (MLE).
The lighter color corresponds to higher value of true p, the parameter that corresponds to the probability of choosing randomly rather than choosing the choice with higher utility.
Rows are the four methods (see above).
Columns are the estimates after the number of questions asked.

Figure 10: Average mean squared error between true and estimated interaction coefficient when two amenities are presented

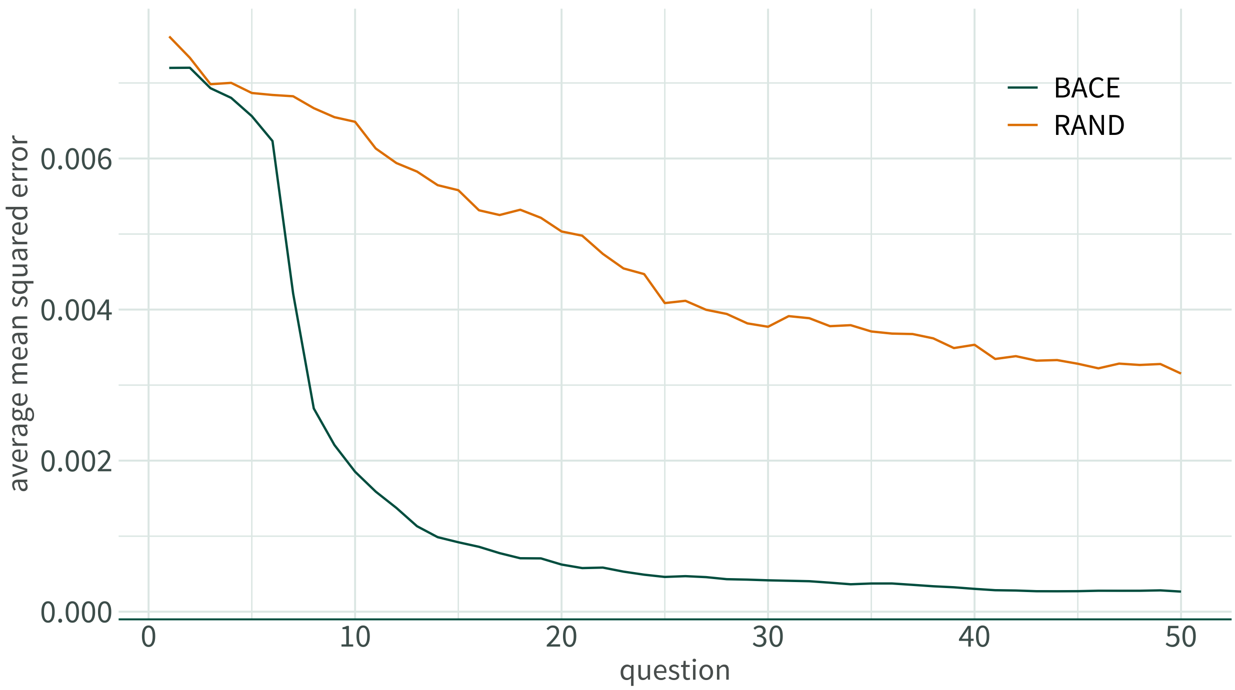

Note: The figure depicts the simulation results in the case when a binary choice is presented between two jobs which differ along three dimensions: earnings and the presence or absence of each of two amenities.
The x-axis is the number of questions used to obtain amenity coefficient estimates.
The y-axis is the average mean squared error between the estimates and the true values from the simulations.
The colors map to two methods: when the sequence of questions are generated by the Bayesian Adaptive Choice Experiment (BACE) or randomly (RAND).

Figure 11: Estimating population mean of γ from individual estimates vs. by pooling all answers

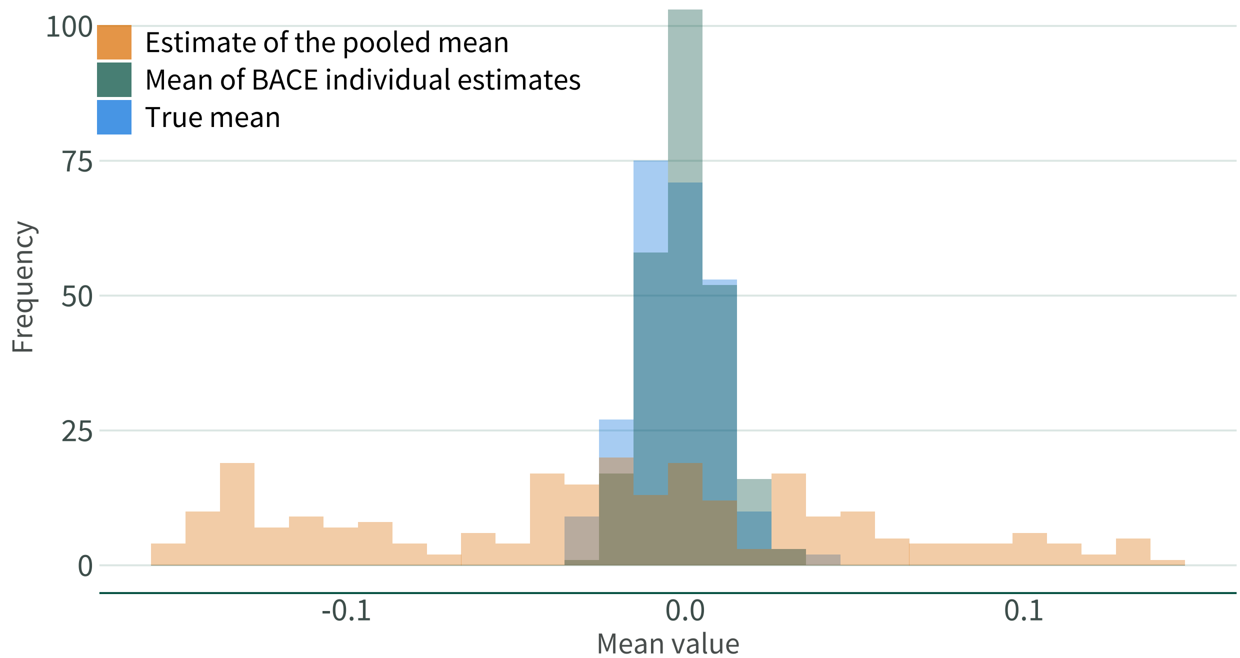

(a) Case 1

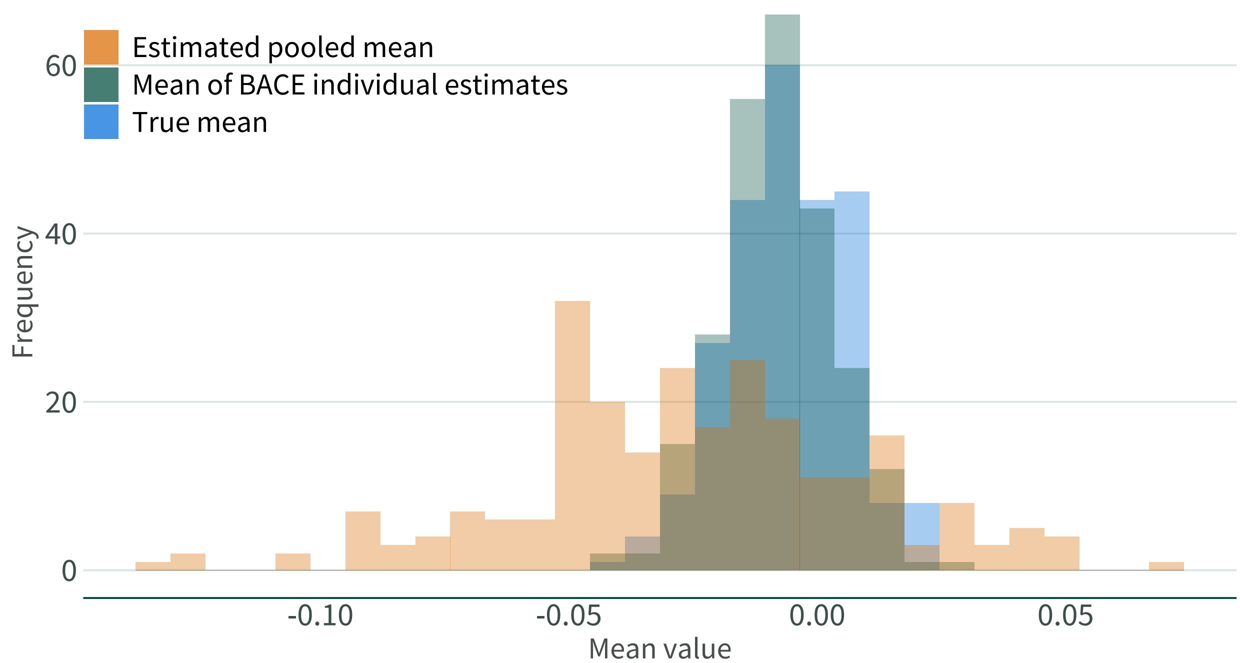

(b) Case 2

Note: Simulation result comparisons for estimating the mean γ in the generated data (blue) by averaging the BACE individual estimates (green) or by estimating the pooled data across all individual answers (orange). The number of scenarios per individual is held fixed at 20.

Figure 12: Mean parameter comparison to the standard pooling approach

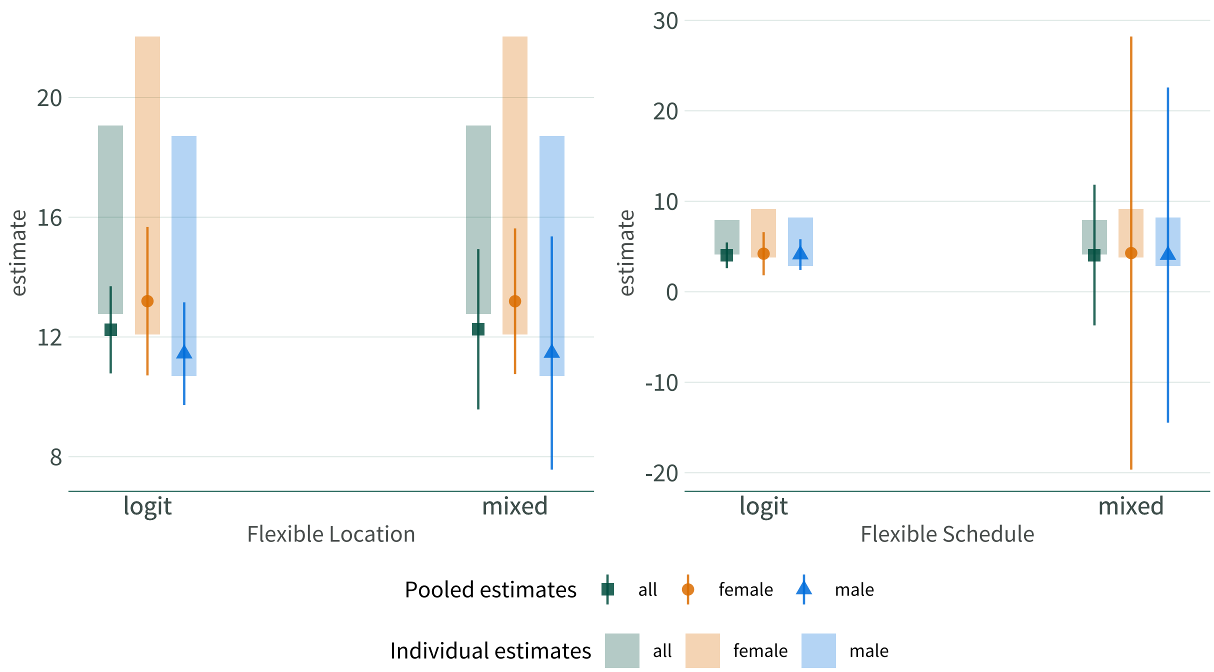

Note: This figure compares the mean of the individual-level WTP estimates to estimating the mean WTP when pooling together responses across all individuals. Individual estimates and bootstrapped confidence intervals are in the shaded regions. Pooled estimates are represented as points with error bars as confidence intervals.