The DARC Toolbox: automated, flexible, and efficient delayed and risky choice experiments using Bayesian adaptive design^{*}

Benjamin T. Vincent

*School of Social Sciences
University of Dundee, UK*

B.T.VINCENT@DUNDEE.AC.UK

Tom Rainforth

*Department of Statistics
University of Oxford, UK*

RAINFORTH@STATS.OX.AC.UK

Abstract

Delayed and risky choice (DARC) experiments are a cornerstone of research in psychology, behavioural economics and neuroeconomics. By collecting an agent's preferences between pairs of prospects we can characterise their preferences, investigate what affects them, and probe the underlying decision making mechanisms. We present a state-of-the-art approach and software toolbox allowing such DARC experiments to be run in a highly efficient way. Data collection is costly, so our toolbox automatically and adaptively generates pairs of prospects in real time to maximise the information gathered about the participant's behaviours. We demonstrate that this leads to improvements over alternative experimental paradigms. The key to realising our real time and automatic performance is a number of advances over current Bayesian adaptive design methodology. In particular, we derive an improved estimator for discrete output problems and design a novel algorithm for automating sequential adaptive design. We provide a number of pre-prepared DARC tools for researchers to use, but a key contribution is an adaptive experiment toolbox that can be extended to virtually any 2-alternative-choice task. In particular, to carry out custom adaptive experiments using our toolbox, the user need only encode their behavioural model and design space – both the subsequent inference and sequential design optimisation are automated for arbitrary models the user might write.

Keywords: Financial psychophysics, decision making, risky choice, discounting, magnitude effect, inter-temporal choice, time preference, two-alternative forced choice, Bayesian adaptive design, valuation, Monte Carlo methods, sequential design optimisation, maximum marginal likelihood estimation, automated inference

^{*}. Accompanying software is available to download from <https://github.com/drbenvincent/darc-experiments-matlab> or <https://github.com/drbenvincent/darc-experiments-python>

1. Introduction

One problem faced by researchers trying to understand human and animal decision making is to appropriately construct choice environments so as to learn an agent’s decision preferences, what influences these preferences, and what mechanisms underlie these. A bewildering array of experimental paradigms have been created by researchers in order to do this, but one central strategy has been to locate points in decision space where an agent is indifferent between two choices. For example, *if* we were to discount future rewards exponentially as they become more delayed into the future and we discover indifference between £100 now and £60 in 5 years, then we can infer an annual discount rate of approximately 10%.

In order to find these unknown indifference points, agents express preferences between sets of pairs of prospects, most commonly. They must evaluate each prospect which involves a *reward* (monetary or otherwise), a *delay* between choice and receipt of reward, a *risk* of never receiving the reward, and potentially some cognitive or physical *effort* required to obtain the reward. If our experimental procedure has gone well, then the pairs of prospects and the agent’s responses are sufficient for us to make inferences about the parameters of a model, or which of multiple models, best account for the data.

The experimenter’s task of designing choice environments is not trivial however. In the next section, we describe a spectrum of different approaches that researchers have taken, ranging from simple fixed questionnaires to adaptive approaches in which the participant undergoes a customised experience, akin to a choose your own adventure.

The tasks we set ourselves in this work is to a) improve upon the current state of the art Bayesian adaptive design methods in order to run near-optimal experiments, and b) to demonstrate this with a variety of delayed and risky choice (DARC) experiments. Being able to near-optimally select experimental stimuli in real time, on a trial-to-trial basis, has a number of significant advantages. We will maximise the rate at which we obtain information about a participant, allowing researchers to achieve a set level of measurement precision in fewer trials, or to achieve higher precision in the same number of trials, compared to existing methods. This is particularly appealing when experiments are costly in time or resources, such as when real monetary rewards are given, special participant populations are being tested, or when using fMRI or electrophysiological measurements.

1.1 Types of experimental procedures

The experimental tasks that have been applied to study issues in DARC are particularly broad. Many of these are complex and highly domain-specific, and so here we present a brief high-level review of various experimental methods that have been used to study inter-temporal choice tasks which allow us to learn how rewards are discounted as a function of delay.

Experimental methods to study discounting can be categorised in a number of ways, one of which is whether they are passive or adaptive. The simplest and most widely used passive approach to inter-temporal choice is the Kirby (2009) 27-item questionnaire. This is easy to administer, but the majority of the questions are far from indifference points for any given experimental participant and so are uninformative.

There are a number of adaptive approaches however, which seek greater efficiency by determining which discounting questions to pose to a participant, in real time while an experiment is in progress. These in turn can be sub-categorised into heuristic vs. model-based approaches.

The vast majority of adaptive methods use heuristic procedures. Simple rules are created to stepwise increase or decrease one of the decision parameters (reward, delay, risk, etc) in an attempt to find an indifference point. This approach is exactly the same as the commonly used staircase procedures in visual psychophysics (Kingdom and Prins, 2009). Many variations on this theme have been proposed and used in empirical studies, (e.g. Mazur, 1988; Richards, Mitchell, de Wit, and Seiden, 1997), and a particularly clear algorithm is provided by Frye, Galizio, Friedel, DeHart, and Odum (2016) (see Appendix E.1) who provide software for experimenters to use. The result of their procedure is a set of estimated indifference points at a number of delays specified by the experimenter. This would then be scored, to produce estimated discounting parameters, by fitting a parametric discount function (e.g. exponential, hyperbolic, etc) to the raw data or estimated indifference points, either by error minimisation, maximum likelihood (Wileyto, Audrain-McGovern, Epstein, and Lerman, 2004), or Bayesian methods (Vincent, 2016).

This kind of heuristic procedure has been pushed to the limits by (Koffarnus and Bickel, 2014, see Appendix E.2) who present a procedure with only 5 trials and can be conducted in less than 1 minute. In order to do this, the authors take a parametric approach. If we assume a priori that people discount hyperbolically, then it is easy to see that the discount rate can be fully determined by finding just a single indifference point. This approach still uses simple heuristic rules to determine which question (out of a set of 31) will be posed next. Even though the estimated discount rates are correlated with longer adjusting amount procedures ($r^2 = 0.45$, based on their Figure 1), the trade-off between speed of testing and accuracy of results becomes clear. For example, fewer trials will result in less reliable measures, and we are less able to detect any deviations from the assumption of hyperbolic discounting.

Model-based methods are another class of adaptive experimental approach and consist of either maximum likelihood or Bayesian methods. Maximum likelihood-based approaches construct a model of response probability to a set of questions for a given set of parameter values. However, these approaches have more often been used in the data analysis stage (e.g. Peters, Miedl, and Büchel, 2012) rather than the data collection stage. A Bayesian approach to data collection in discounting tasks was recently proposed by Pooseh, Bernhardt, Guevara, Huys, and Smolka (2017). The primary benefit of doing this is that it allows for prior knowledge of discounting parameters to be incorporated. However, while the method presented by Pooseh et al. (2017) would be expected to be reasonably efficient, it also includes a number of heuristic elements. For example, choice of questions to pose to participants is based on the intuition that asking questions at currently estimated indifference points will be the most informative. While this might seem reasonable, it is in fact suboptimal. For example, in the context of estimating a psychometric function in the field of visual psychophysics, Kontsevich and Tyler (1999) demonstrate that at certain points in the experiment it becomes optimal to place experimental design variables either side of the predicted indifference point in order to learn about parameters relating to response variability for example (see their Figure 2). This is particularly important in order to not get fooled when participant behaviours deviates from your model (termed model misspecification), and is discussed further in the paper and Appendix G.

The gold standard in terms of efficient experimentation comes from the field of Bayesian adaptive design (Cavagnaro, Aranovich, McClure, Pitt, and Myung, 2016; Chaloner and Verdinelli, 1995; Myung, Cavagnaro, and Pitt, 2013a; Sebastiani and Wynn, 2000a). Also known as sequential Bayesian experimental design, Bayesian adaptive design is another class of adaptive methods that will be described in detail in the next section.

1.2 Overview of our approach

Our goal is to collect data as efficiently as possible, so that we learn as much as we can about a participant’s behaviour with each question posed. The overall approach to achieve this is to use Bayesian adaptive design (see Figure 1) which has been applied to achieve similar efficiency aims in a wide variety of areas (Cavagnaro et al., 2016; Kontsevich and Tyler, 1999; Myung, Karabatsos, and Iverson, 2005; Prins, 2013; Watson, 1983). There are 4 main steps required to initially set up an adaptive experiment:

1.2.1 STEP 1: DEFINE A DESIGN SPACE

We as experimenters need to define the design space $\mathcal{D}$ (i.e. set of possible questions we could ask) appropriate for the chosen experiment. This design space can be thought of as a large table where each row d represents one possible question that could be posed to a participant. The number of columns represent the dimensionality of the experiment.

1.2.2 STEP 2: DEFINE A RESPONSE SPACE

In this paper we confine ourselves to experiments where participants make binary responses, that is they can respond with 1 of 2 possible responses on each trial, $y \in \{0, 1\}$. We note, however, that our Bayesian adaptive design approach (though not, at present, the provided toolbox) also applies equally well to cases where there are any (finite) number of possible responses.

1.2.3 STEP 3: DEFINE A GENERATIVE MODEL

We need a participant response model (i.e. a likelihood $p(y|d, \theta)$, see Figure 2) which relates latent parameters θ to observable responses y for a given design d . For discounting tasks, such a model could be based upon one of many possible discounting functions (see Doyle, 2013). It does not matter for our purposes whether this model is descriptive or explanatory, but it must provide us with the likelihood of a response given an experimental design and set of model parameters. Assuming the model is a reasonable way to capture participant’s behaviours, then we better characterise a participant’s behaviour by reducing our uncertainty over the parameters.

While we have focussed on DARC experiments with binary choices (i.e. 2 prospects), our approach is very general. It can be adapted to virtually any binary-choice task.

1.2.4 STEP 4: DEFINE A PRIOR OVER MODEL PARAMETERS, $p(\theta)$

We can make experiments even more efficient by providing any existing knowledge about participants in the form of priors over parameters. Very often we will have this kind of prior knowledge because of the particular participant population being tested. By doing so, we can avoid asking questions which result in responses that would not have surprised us. Prior knowledge can not only be used in the data collection stage, but also in the data analysis stage of DARC experiments (e.g. Nilsson, Rieskamp, and Wagenmakers, 2011; Vincent, 2016). It does however require experimenters to have given thought to not just what behaviour they may expect, but to produce a prior belief over parameters which results in this expected behaviour.

After these preparatory steps have been completed, we have 3 repeating steps that form the core of an adaptive experiment:

Figure 1: The steps used in Bayesian adaptive design. The initial setup steps for different models and experimental paradigms require some attention. We provide a number of pre-constructed tools in the toolbox, but the toolbox is modular and expandable for users to construct their own. Once set up, using these in an adaptive experiment is automatic. The toolbox provides parameter estimates as well as raw data to use in more advanced analysis methods.

1.2.5 STEP 5: QUESTION SELECTION (AKA DESIGN OPTIMISATION)

On any given trial, we must select which design to present to the participant out of the set of possible designs. Static, non-adaptive, procedures simply select the next question in the list, regardless of any previous responses, and so will waste experiment time by asking questions which provide us

Figure 2: The general form of the data generating process in 2-choice tasks. Our approach is very general and can be applied to any binary response situation where a data likelihood function $P(y|\theta, d)$ can be defined. Here Ψ represents a generic model which maps experimental design and model parameters to response probability. Continuous and discrete variables are represented as circles and squares, respectively. Single-bordered nodes represent stochastic variables, and double-bordered nodes represent deterministic functions of parent nodes.

with little information with which to update our beliefs. Adaptive methods however, can use the model and current posterior over parameters in order to maximise expected information gain about parameters, given that we can only guess how the participant will respond. The optimal design d^* is chosen on the basis of maximising the expectation of a utility function over all possible experimental outcomes. In our case, we seek to maximise expected information gain about model parameters (see Section 2.4). It can be shown that this procedure corresponds to the optimal decision under uncertainty (under the assumption that our model is correct) for choosing the next design. A key novel contribution of our work is in using this principled information-theoretic approach for the design optimisation and in developing an efficient, provably correct, and automated process for carrying it out.

1.2.6 STEP 6: COLLECT PARTICIPANT RESPONSE

Our task now is to take the optimal design, d^* , and present this to the participant in an appropriate manner, such that the participant can understand and choose between the prospects. In the software toolbox, this involves a simple function call – the input is a the chosen design d^* which is simply a row vector representing the stimulus presented to the participant, and the output is the response which could consist of the choice as well as other useful information such as reaction time. This function can be adapted by experimenters to suit their particular preferences for stimulus presentation, choice framing and response recording they require.

1.2.7 STEP 7: INFERENCE STEP TO UPDATE OUR BELIEFS.

Non-probabilistic approaches traditionally use heuristic scoring methods to arrive at a single estimated indifference point. However, here we conduct inference in order to update our beliefs about parameters given the new response data. This is conceptually rather easy as we can just apply Bayes' rule, but in practice this is not trivial. A key part of our approach is doing this automatically and efficiently for a wide selection of models in a manner that is able to scale to models with many parameters, unlike the commonly used grid approximation methods employed by other adaptive approaches (Kontsevich and Tyler, 1999; Prins, 2013; Treutwein and Strasburger, 1999). Because we evade the heuristic estimation of indifference points, we are not tied to design strategies which take 'slices' along delay, reward, or probability values in search of indifference points. Instead, our use of Bayesian inference allows us to use a statistically grounded decision-theoretic approach and also to consider a much larger design space to test a greater range of possible delays or reward values. Furthermore, once the experiment is complete, we are able to immediately return a full posterior over the participant's parameter values, conveying significantly more information than a simple point estimate.

Steps 5–7 are repeated until we meet our experiment termination criteria. A simple scheme could be when we reach a predefined number of trials, but more advanced schemes are also possible such as terminating when the entropy (i.e. uncertainty) in the participant's parameters falls below a desired threshold.

1.3 Our contributions

Many of the practical advances which will interest the experimenter, rely upon a number of important theoretical and computational advancements in Bayesian adaptive design which are essential in both fully automating the Bayesian adaptive design process and in ensuring that it can be carried out fast enough for the toolbox to be run in real-time. In the next section we present our novel and very general Bayesian adaptive design optimisation approach. It can be applied to almost all research contexts where participants make binary choices, but we go on to thoroughly describe how the Bayesian adaptive design approach can be applied to risky inter-temporal choice tasks. At the heart of this approach is a novel Monte Carlo based method for the inference and optimisation required for Bayesian adaptive design. To the best of our knowledge, this is the first fully automated approach to *sequential* Bayesian adaptive design for a general class of problems – the user need only specify their model and set of allowable designs, everything else is automated. Achieving this involves many core innovations as follows:

- We employ a general purpose inference scheme, meaning that we can carry out inference efficiently for a large range of possible models the user might specify.
- Previous common Bayesian adaptive design methods have typically used either inefficient grid search (Kontsevich and Tyler, 1999; Prins, 2013; Treutwein and Strasburger, 1999) or naïve nested Monte Carlo approaches (Myung et al., 2013a) to calculate the expected gain in Shannon information used as the target for the design optimisation (Chaloner and Verdinelli, 1995). We instead introduce a novel non-nested Monte Carlo estimation scheme with a substantially improved convergence rate (Rainforth, Cornish, Yang, Warrington, and Wood, 2017), leading to improvements of orders of magnitude in the efficiency of the estimation.

- By using a *particle* based inference scheme, namely population Monte Carlo (PMC) (Cappé, Guillin, Marin, and Robert, 2004), we are able to propagate our distribution of beliefs about model parameters using a set of samples known as particles, the number of which can be scaled according to available computer power. This propagation allows us to both transfer information from one inference to the next and modularise the problems of inference and optimisation. Not only does this dramatically improve the speed, accuracy and memory requirements of the toolbox compared with existing methods, making our software practical for real-time usage, it is also key to our ability to fully automate the sequential Bayesian adaptive design process.
- We use a novel optimisation scheme that only requires a set of samples representing the posterior and a set of valid designs as input. This means we are able to fully automate both the inference and optimisation processes for any user specified model by inputting the particles output from the former into our novel scheme for solving the latter. Furthermore, our design optimisation scheme builds on ideas from Amzal, Bois, Parent, and Robert (2006) to produce a novel and highly efficient scheme for optimisation of intractable expectations in its own right, adaptively allocating computational resources between the estimators for different designs.

These advances are described in detail in Section 2. However those primarily interested in the application to specific experimental methods can feel free to skip this technical section, progressing to Section 3 onwards. We describe how we apply Bayesian adaptive design in the context of DARC experiments, provide a detailed worked example and evaluation of adaptive temporal discounting experiments, and demonstrate the adaptability by applying our toolbox to other DARC experiments.

This work also makes a number of novel practical contributions, allowing experimenters to easily run efficient DARC experiments. More specifically:

- We are the first to provide a comprehensive set of efficient experimental tools to run DARC experiments, with easy extendability to many other 2-choice experiments.
- We make the software freely available (see end of document) for researchers to collect data using these methods in their own experiments. The code is written in a modular fashion to allow simple extensions such as using different time discounting or probability weighting functions, and to construct entirely new models for different experimental tasks.
- Through the example of temporal discounting, we demonstrate accurate recovery of known delay discounting parameters for a simulated participant with much higher precision (lower variance) than alternative fixed and adaptive methods.
- We demonstrate the generality and ease of extending the approach to other DARC experiments by describing a) temporal discounting the the magnitude effect, b) probability discounting (risky choice), and c) combined delayed and risky choices.

2. Our approach to finding the optimal design

From tuning microscopes to designing surveys on politics, all problems of experimental design can be reduced to the same mathematical abstraction: choosing a design that maximizes the expected

information gained from the experiment. Bayesian adaptive design (Chaloner and Verdinelli, 1995; Sebastiani and Wynn, 2000b), also known as Bayesian experimental design or Bayesian-optimal design, provides a powerful framework for this abstraction. In Bayesian adaptive design one constructs a likelihood model for the experiment outcome and then uses this, along with a prior on the parameters one wishes to learn about, to find the experimental setup that maximizes the expected information gathered by the experiment. If the model is correct, this forms a design strategy that is optimal from an information-theoretic viewpoint (Sebastiani and Wynn, 2000b). The general nature of its formulation means that Bayesian adaptive design can, at least in principle, be applied to almost any experimental design situation. It can be particularly useful in *sequential* design settings, where it provides a means of constructing adaptive strategies that use previous experience in order to make decisions which optimize the sequential learning process. In the context of DARC experiments, this means we can adapt the questions asked to each participant, resulting in questions tailored to match their behaviour and maximize the information gained.

In this section we will introduce a method for automating Bayesian adaptive design in arbitrary sequential design problems. Our approach only requires the user to specify their model and design space, with all subsequent calculations automated for any arbitrary model the user might write. Not only does this allow for rapid development of new models and experiments, it also means users require no expertise in Bayesian adaptive design to setup and run their own efficient adaptive experiments. The framework is also extremely general and can be applied to any problem with discrete responses.¹ As alluded to in Section 1.3, there are a number of technical innovations underlying our approach. Roughly speaking, these can be grouped into the development of a new estimator for the Bayesian adaptive design equations and a mixed inference-optimisation algorithm for using this estimator to carry out the sequential adaptive design. After providing a more technical introduction to Bayesian adaptive design, we will now explain these in turn.

2.1 Bayesian-optimal design

In order to carry out adaptive design, we need some metric for how good a design is. By optimizing this metric we can then find the best design. The key to Bayesian adaptive design is a particular metric that we now derive. For simplicity, we will for now only consider the problem of selecting a single design and ignore conditioning on previous design-response pairs.

Let the parameters we wish to learn about be denoted by $\theta \in \Theta$, for which we defined a prior distribution $p(\theta)$. Let the probability of achieving outcome (i.e. response) $y \in \mathcal{Y}$, given parameters θ and a design $d \in \mathcal{D}$, be defined by the likelihood model $p(y|\theta, d)$. Under our model, the outcome of the experiment given a chosen d is distributed according to

$$p(y|d) = \int_{\Theta} p(y, \theta|d) d\theta = \int_{\Theta} p(y|\theta, d) p(\theta) d\theta \quad (1)$$

where we have used the fact that $p(\theta) = p(\theta|d)$ because θ is independent of the design.

To derive our metric for how good a design is, we first define a utility function $U(y, d)$, representing the utility of choosing design d and getting response y . Typically our aim is to maximize the information about the parameters θ , which can be done by maximizing the gain in Shannon

1. Though our code implementation currently only supports binary responses, this is not inherent to the approach and we plan to extend the code at a later date.

information between the prior and the posterior

$$U(y, d) = \int_{\Theta} p(\theta|y, d) \log(p(\theta|y, d)) d\theta - \int_{\Theta} p(\theta) \log(p(\theta)) d\theta. \quad (2)$$

In other words, we want to minimize the entropy after observing y , relative to the entropy before. However, we are still uncertain about the outcome y . The Bayesian adaptive design metric is, therefore, defined as the expected utility, i.e. the expectation of $U(y, d)$ with respect to $p(y|d)$:

$$\bar{U}(d) = \int_{\mathcal{Y}} p(y|d) \left(\int_{\Theta} p(\theta|y, d) \log(p(\theta|y, d)) d\theta - \int_{\Theta} p(\theta) \log(p(\theta)) d\theta \right) dy. \quad (3)$$

Note that this corresponds to the mutual information between the parameters θ and the observations y . As $\bar{U}(d)$ corresponds to the Bayesian adaptive design metric, the Bayesian-optimal design is now given by

$$d^* = \operatorname{argmax}_{d \in \mathcal{D}} \bar{U}(d). \quad (4)$$

We can intuitively interpret d^* as being the design that most reduces the uncertainty in θ on average over possible experimental results. If our likelihood model is correct, i.e. if experimental outcomes are truly distributed according to $p(y|\theta, d)$ for a given θ and d , then it is easy to see from the above definition that d^* is the true optimal design, in terms of information gain, given our current information about the parameters $p(\theta)$. In practice, our likelihood model is an approximation of the real world. Nonetheless, Bayesian adaptive design remains a very powerful and statistically principled approach that is typically significantly superior to other, more heuristic-based, alternatives.

However, a major drawback to the Bayesian adaptive design approach is that it is typically difficult and computationally intensive to carry out. Not only does it represent an optimisation of an intractable expectation, the integrand is itself intractable because the $p(\theta|y, d)$ term cannot be calculated analytically. Consequently, (3) forms a so-called nested expectation which cannot be directly estimated using, for example, conventional Monte Carlo estimation. Instead one typically must rely on nested estimation schemes which are usually very computationally expensive (Rainforth, 2018; Rainforth et al., 2017).

Remarkably, however, our novel approach, allows the accurate estimation of d^* with a computational complexity similar to that of finding an indifference point. From a practical perspective, the time taken in estimating d^* is similar to that of the inference used to keep track of the distribution over the parameters through the experiment and can easily be run in real-time on a small laptop.

2.2 Estimating $\bar{U}(d)$

As we explained in the last section, estimating d^* is in general challenging because the posterior $p(\theta|y, d)$ is rarely known in closed form. One popular method within the visual psychophysics literature (Kontsevich and Tyler, 1999; Prins, 2013; Treutwein and Strasburger, 1999) is to use a grid-search approach to the required inference, which subsequently provides (biased) estimates for $p(\theta|y, d)$. However, this is highly unsatisfactory for a number of reasons. Firstly, it is well documented that grid-search is typically a substantially inferior inference method, compared with, for example, Monte Carlo approaches (Robert, 2004). Its many shortcomings include an exponential degradation of performance with dimension, failure to deal with unbounded variables and spending most of its computational budget evaluating points with very low posterior probability. Secondly, it

is a highly inflexible approach and requires careful, problem specific, human input – the grid must be chosen up front. Both these make it unsuited to our needs.

Alternatively, as shown by, for example, [Myung et al. \(2013a\)](#), one can instead use the following nested Monte Carlo estimator (see Appendix A for a derivation)

$$\bar{U}(d) \approx \hat{U}_{NMC}(d) := \frac{1}{N} \sum_{n=1}^N \left[\log(p(y_n | \theta_{n,0}, d)) - \log \left(\frac{1}{M} \sum_{m=1}^M p(y_n | \theta_{n,m}, d) \right) \right] \quad (5)$$

where each $\theta_{n,m} \sim p(\theta)$, $y_n \sim p(y | \theta = \theta_{n,0}, d)$, and N and M are a number of outer and inner samples respectively. The total cost of this estimator is $T = NM$ and the best achievable convergence rate in the mean squared error is order $1/T^{2/3}$ ([Fort, Gobet, and Moulines, 2017](#); [Hong and Juneja, 2009](#); [Rainforth et al., 2017](#)) – substantially slower than the order $1/T$ rate of conventional Monte Carlo.

To address the inefficiency of this nested estimator, we have developed a method of estimating the Bayesian adaptive design metric (3) that maintains the convergence rate of conventional Monte Carlo whenever y can only take on one of C possible values ($y_1, \dots, y_C$), providing a substantial improvement. We will omit the derivation of this estimator here, referring the reader to Appendix A, quoting only the final result as follows

$$\begin{aligned} \bar{U}(d) \approx \hat{U}(d) := & \frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C p(y_c | \theta_n, d) \log(p(y_c | \theta_n, d)) \\ & - \sum_{c=1}^C \left[\left(\frac{1}{N} \sum_{n=1}^N p(y_c | \theta_n, d) \right) \log \left(\frac{1}{N} \sum_{n=1}^N p(y_c | \theta_n, d) \right) \right] \end{aligned} \quad (6)$$

where each $\theta_n \sim p(\theta)$ and y_c are fixed constants. The critical difference to the nested estimator (5) is that C ($= 2$ for our DARC experiments) is a fixed constant rather than being a number of samples that needs to be increased to increase the accuracy of the estimator; we only here need N samples of θ rather than NM samples. Therefore, if our estimator uses N samples, then using (5) will require on the order of $N^{3/2}$ samples to achieve the same accuracy. Given N is typically quite large, ($\gtrsim 10^4$) this means (5) usually needs multiple orders of magnitude more samples to achieve the same accuracy as our estimator.

We finish the description for our estimator of $\bar{U}(d)$ by noting that it requires only the ability to sample θ_n and evaluate $p(y_c | \theta, d)$ for a given θ , d , and y_c . The latter of these is always trivial as it simply corresponds to the model we have defined. The former is also trivial in the absence of previous experiment-response pairs, but, as we will discuss next, requires more detailed consideration in sequential experiments.

2.3 Sequential experiments and parameter inference

We have thus far assumed that there is no previous data (i.e. design-response pairs), but in practise the experiment is sequential and previous data must be incorporated in choosing future designs. This previous data is incorporated into the model through conditioning θ . As such, at experiment iteration t , we replace $p(\theta)$ with the posterior $p(\theta | d_{1:t-1}, y_{1:t-1})$, where $d_{1:t-1}$ and $y_{1:t-1}$ are respectively the designs and responses at previous iterations. This posterior now represents our beliefs about the participants parameters given the information from previous questions and the ultimate aim of our experiment is to refine it. Making changes to our estimator to operate in this sequential setting is

Algorithm 1 Sequential Bayesian adaptive design

Inputs: Prior $p(\theta)$, likelihood $p(y|\theta, d)$, number of experiment iterations T , candidate designs $\mathcal{D}$
Outputs: Design-response pairs $\{d_t, y_t\}_{t=1:T}$, posterior samples $\theta_{1:M}$

```

1:  $\theta_m \sim p(\theta) \quad \forall m = 1, \dots, M$ 
2: for  $t = 1 : T$  do
3:    $d_t \leftarrow \text{DESIGNOPTIMISATION}(\theta_{1:M}, p(y|\theta, d), \mathcal{D})$  ▷ Find next step optimal design
4:    $y_t \leftarrow \text{RUNEXPERIMENT}(d_t)$  ▷ Get response from using design  $d_t$ 
5:    $\mathcal{D} \leftarrow \mathcal{D} \setminus d_t$  ▷ Eliminate new design from candidate set (optional)
6:    $\theta_{1:M} \leftarrow \text{INFERENCE}(p(\theta), p(y|\theta, d), \{d_i, y_i\}_{i=1:t}, \theta_{1:M})$  ▷ Update the posterior
7: end for
8: return  $\{d_t, y_t\}_{t=1:T}, \theta_{1:M}$ 

```

conceptually remarkably simple – we just sample θ from the posterior instead of $p(\theta)$. In other words, we sample $\theta_n \sim p(\theta|d_{1:t-1}, y_{1:t-1})$ and then use (6).

Unfortunately, generating these samples can be challenging as it requires *Bayesian inference*. Despite its clear shortfalls, this is perhaps why the grid-search estimation approach has previously been employed so readily in the literature as it, at first glance, seems to sidestep this issue. However, the key point here is that grid-search *is* an inference method and so it is not side-stepping the issue at all, but instead just carrying out this required inference very inefficiently. In fact, the grid-search approximation is even more inefficient than the naïve nested estimator (5).

Our approach instead exploits two key features of the estimator $\hat{U}(d)$ – that it only requires provision of samples $\hat{\theta}_n$ and that the same samples can be used to evaluate multiple designs. This means that we can modularize the problems of inference and of optimizing the Bayesian adaptive design metric to choose the best design. Namely, we can carry out inference to generate a “reservoir” of samples $\theta_{1:M}$ (for some large M) and then pass these on to an algorithm that uses them to find the optimal design, drawing samples from the reservoir as an when required, rather than generating directly from the posterior on-demand. A high level summary of our approach is given in Algorithm 1. As input we require the model (in the form of the prior and the likelihood), the set of candidate designs, and a number of experiment iterations T . For the first design, we just generate our reservoir simply by sampling from the prior (line 1). We then iterate between using our reservoir to choose the next design (line 3), evaluating that design (line 4), and updating the posterior to include this new design-response pair (line 6).

Given this framework, we see that there are two key components to the approach which we are yet to introduce – the design optimisation algorithm and the method for carrying out inference. While the latter is introduced in the next section, the former builds on ideas from probabilistic programming (Goodman, Mansinghka, Roy, Bonawitz, and Tenenbaum, 2008; Rainforth, 2017; Wood, Meent, and Mansinghka, 2014) in order to automate the inference problem for a wide array of problems. The core philosophy of probabilistic programming systems (PPSs) is to decouple model specification and inference, with the latter carried out by an inference engine capable of operating on arbitrary models. While our approach is certainly not a PPS in its own right, it builds on the same idea of automating inference for a wide range of models to remove this burden from the user. Specifically, we require the user to provide only the prior, likelihood, and data, with the subsequent inference fully automated (though there are parameters users can tweak if they desire

Algorithm 2 Design optimisation**Inputs:** Reservoir $\theta_{1:M}$, candidate designs $\mathcal{D} = d_{1:K}$, number rounds L and samples per round N **Outputs:** Chosen design d^*

```

1:  $\hat{U}_{1:K} \leftarrow 0$  ▷ Initialize estimates
2: for  $\ell = 1 : L$  do
3:    $n_{1:K} \leftarrow \text{ALLOCATESAMPLES}(N, \hat{U}_{1:K})$  ▷ Allocate a number of samples  $n_k$  to each design
4:   for  $k = 1 : K$  do
5:      $\hat{\theta}_{1:n_k} \sim \text{DISCRETE}(\theta_{1:M})$  ▷ Draw  $n_k$  samples from reservoir
6:      $\hat{U}_k \leftarrow \text{UPDATEESTIMATE}(\hat{U}_k, \hat{\theta}_{1:n_k}, d_k)$  ▷ Use new samples to refine estimate
7:   end for
8: end for
9:  $k^* \leftarrow \text{argmax}_{k \in \{1, \dots, K\}} \hat{U}_k$  ▷ Best design is one with highest  $\hat{U}_k$ 
10: return  $d_{k^*}$ 

```

such as the computational budget). Our inference scheme is based on *population Monte Carlo* (PMC) (Cappé et al., 2004), technical details for which are given in Appendix B.

There are a number of advantages to using PMC for the problem at hand. Firstly, the fact that it is based on a population of samples representing the target distribution allows for transferring of considerable information from the inference problem at one experiment iteration to the next, compared with say a conventional MCMC sampler. Secondly, PMC naturally outputs the required sample reservoir, corresponding the final population of samples. Finally, because it uses localized proposals, it is less prone to sample degeneracy than importance sampling based algorithms such as SMC (Doucet, Godsill, and Andrieu, 2000). In particular, such methods would scale exponentially badly in the dimensionality of θ , while PMC can be noticeably more robust if the proposal is chosen appropriately. Though θ is quite low dimensional in the models we have considered, using PMC means that, at least in principle, our toolbox will be able to scale gracefully to higher dimensions.

2.4 Optimisation

Having demonstrated how to estimate $\bar{U}(d)$, we will now look at how to optimise for d . In all but the most trivial scenarios, simply evaluating the utility for each candidate design using the full reservoir of samples and then choosing the best will not be practical due to its excessive computational cost, particularly given there will often be thousands of possible designs or even more. Instead we will need to be clever about how we allocate our finite computational resources to achieve the most reliable estimate for which design is optimal. To do this we introduce a novel approach that uses adaptive resource allocation between different design, sharing some parallels with sequential Monte Carlo search (Amzal et al., 2006), Thompson sampling (Thompson, 1933), adaptive Monte Carlo via bandit allocation (Neufeld, György, Schuurmans, and Szepesvári, 2014), and the Hyperband algorithm (Li, Jamieson, DeSalvo, Rostamizadeh, and Talwalkar, 2016). We emphasise though that our method is a new algorithm in its own right.

The underlying idea of our approach is that we are willing to accept higher variance in the Bayesian adaptive design metric estimates for lower utility designs as this error is unlikely to lead to an incorrect estimation of which design is optimal. Therefore, we desire to take more Monte Carlo samples, and thus achieve a more accurate estimate, for the designs we believe to have a higher chance of being the optimal design. An overview of our algorithm is given in Algorithm 2, a

more detailed version of which is given in Appendix C. We again highlight the importance of using a reservoir of samples to provide modularization between the inference and optimisation processes, and thereby the automation of our pipeline.

Rather than calculating a full estimate $\hat{U}_k$ for a design d_k in one go, we exploit the intermediate information from incomplete evaluations to reallocate computational resources. Given a total budget of NL samples, we carry out L sampling rounds, allocating N samples between the K designs at each round, assigning more samples to more promising designs. These new samples are used to refine the estimates $\hat{U}_k$, which in turn leads to a refinement of the sample allocation strategy at the next round. At the end of our L rounds, we will have used more samples, and thus achieved better estimates, for the better designs. As we are only concerned with which design is best, this in turn increases the chance we will choose the true optimal design. Further explanation and technical details are given in Appendix C.

2.5 Theoretical result

We finish the introduction of approach with the following theoretical result that confirms its statistical correctness, in the sense that the method is guaranteed to return the optimal design at each iteration in the limit of large computational resources.

Theorem 1. *Assume that the PMC proposal $q(\theta'|\theta)$ satisfies the convergence requirements laid out in Cappé et al. (2004) and that each $\bar{U}_t$ (defined as per (23)) is a measurable function. Given a finite number of candidate designs $d_{1:K}$, fixed parameters $\rho > 0$ and $N \in \mathbb{N}^+$, and a function $\tau : \mathbb{R} \rightarrow \mathbb{R}$ such that $\tau(L) \geq NL \forall L$, then using Algorithms 1 and 3 to choose design outcome pairs with $M = \tau(L)$ ensures that the chosen design at each iteration t satisfies*

$$\bar{U}_t(d_t) = \max_{d \in \mathcal{D}} \bar{U}_t(d) \quad (7)$$

almost surely (i.e. occurs with probability 1) in the limit $L \rightarrow \infty$.

Proof. See Appendix D. □

Remark 1. This result holds when both the number of particles in the design optimisation algorithm N and the number of iterations of the PMC inference are fixed finite values. The convergence is in L (and by proxy M) only.

3. Bayesian adaptive design applied to DARC experiments

The advances we have made in the previous section can be applied to virtually any 2-choice experiment that can be modelled. Given this is inherently very general we envisage that these methods (embodied in our software toolbox) can be used and adapted by researchers for a wide variety of experimental contexts. Obviously, it is important to demonstrate our approach and to compare it to others, and we do this by choosing the general domain of delayed and risky choice experiments. The DARC domain is still incredibly broad however (see Figure 3), so while it is impossible to embody all possible decision making experiments in one fell swoop, we target a broad class of models and a reasonably broad design space that will be easy for researchers to adapt to their particular needs.

Figure 3: A schematic of some of the core modelling themes (rows) in the utility-based decision making approach. There are a wide family of models under the utility-based approach, from more normative approaches (with linear probability weighting functions and exponential time discounting) to more descriptive approaches such as the Prospect theory family of models (which does not address delayed rewards), to the discounting approach (which typically assumes a linear utility function).

3.1 Utility-based models of DARC behaviour

Readers are referred to Table 1 for a summary of the notation we use for modelling behaviour in DARC experiments. First, we assume that participants compare two prospects $\mathcal{P}$ which consist of rewards R , to be received with a delay D , with probability P , i.e. $\mathcal{P} = (R, D, P)$. For the present work, we assume a simple reward/no-reward structure such that there is a $1 - P$ probability of obtaining zero reward. As mentioned, we also restrict ourselves to experiments where choices are made between 2 prospects, $\mathcal{P}^a$, and $\mathcal{P}^b$. The preferred prospect is that which has highest present subjective value $V(\mathcal{P})$, also termed utility, but we describe the stochastic nature of choices in more

Term	Explanation
$V(\mathcal{P})$	present subjective value of a prospect $\mathcal{P}$
$u(R)$	subjective value function of a reward R
$d(D)$	discounting of a delay D
$\pi(P)$	risk function for probability P
θ	parameters
$p(\theta)$	prior over parameters
$p(y d, \theta)$	likelihood of response y for a given design d and set of parameters θ
Ψ	a generic participant response model for binary choices
$L()$	a choice (or link) function
Φ	cumulative standard normal distribution
y	responses, $\in \{0, 1\}$
t	trial number

Table 1: Notation used for modelling behaviour in DARC experiments.

detail below. The present subjective value of a prospect, in this class of models, is determined by

$$V(\mathcal{P}) = u(R) \cdot \pi(P) \cdot d(D) \quad (8)$$

where $u(R)$ is a function which maps objective rewards to subjective value (i.e. utility), $\pi(P)$ is a probability weighting function (equally a probability discounting function), and $d(D)$ is a discounting factor accounting for delay of the prospect².

In this utility-based approach, there are of course many model variations (see Figure 3), each focussing on possible functions for utility $u(R)$, risk $\pi(P)$, and delay $d(D)$, each of which will have parameters which need to be inferred from behavioural data. The approach presented here will allow highly efficient experiments to be conducted to maximise information gained about these parameters, for a given model.

Determining the nature of the *value function* has been the focus of sustained theoretical and empirical work (e.g. [Kahneman and Tversky, 1979](#); [Stott, 2006](#)). The DARC toolbox is suited to continue this exploration, but in the present work we focus on the discounting approach which traditionally assumes a linear utility function, $u(R) = R$.

The *probability weighting function* is normally composed of some function $\pi(P)$ that over- or under-weights low and high probabilities in line with the central ideas of prospect theory ([Kahneman and Tversky, 1979](#); [Tversky and Kahneman, 1992](#)) and derivative models. However, the discounting literature typically refers to the probability weighting function as ‘probability discounting’, and considers hyperbolic discounting (or similar) of odds against receiving a prospect. This mismatch between approaches is only superficial; odds and probabilities can be trivially calculated from each other, and there are interesting connections and equivalences between probability weighting and probability discounting ([Prelec and Loewenstein, 1991](#); [Rachlin, Raineri, and Cross, 1991](#); [Takahashi, 2011](#); [Takemura and Murakami, 2016](#)).

2. Note that when dealing with the more general case of a full payoff distribution, then R , P , and D are now vectors where the sum of elements in P equals 1. In this case, Equation 8 becomes $V(\mathcal{P}) = \sum_i u(R[i]) \cdot \pi(P[i]) \cdot d(D[i])$. Our provided software toolbox easily extends to payoff distributions.

Term	Explanation
$\mathcal{D}$	Design space, a 2-dimensional matrix
d	A particular design, one row of $\mathcal{D}$
d^*	The optimal design given our model and beliefs
$\mathcal{P}$	A prospect, $\mathcal{P} = (R, D, P)$
a	Prospect a , also referred to as $\mathcal{P}^a$
b	Prospect b , also referred to as $\mathcal{P}^b$
R	reward magnitude
D	delay of reward
P	probability of obtaining reward

Table 2: Summary of the notation used for the design space

The *delay discounting function* $d(D)$ is typically not included in the Prospect Theory class of models, but allows us to predict the present subjective value of delayed rewards. Given that many real-world decisions do not result in instant reward delivery, understanding agents' temporal preferences is important. This has been done (under the utility-based approach) by investigating different forms of the discount function (Cavagnaro et al., 2016; Doyle, 2013; McKerchar, Green, Myerson, Pickford, Hill, and Stout, 2009; Peters et al., 2012).

In order to conduct design optimisation we need a likelihood model, which can be thought of as a putative model of participant responses $P(y|\theta, d)$ for a given design d and set of parameter values θ (see Figure 2). We define this as

$$P(y|\theta, d) \sim \text{Bernoulli}(\Psi(V(\mathcal{P}^a), V(\mathcal{P}^b), \theta)) \quad (9)$$

where the probability of preferring $\mathcal{P}^b$ (coded as $y = 1$), $\Psi(V(\mathcal{P}^a), V(\mathcal{P}^b), \theta)$, is defined by a psychometric function (see Vincent, 2016, for more details) for which we take as a default

$$\Psi(V(\mathcal{P}^a), V(\mathcal{P}^b), \theta) = L(V(\mathcal{P}^a), V(\mathcal{P}^b), \theta) \quad (10)$$

$$= \varepsilon + (1 - 2\varepsilon) \cdot \Phi\left(\frac{V(\mathcal{P}^b) - V(\mathcal{P}^a)}{\alpha}\right) \quad (11)$$

where Φ is the cumulative standard normal distribution. The comparison acuity α can be thought of as capturing noise in the participant's process of calculating present subjective value, and the lapse rate ε is the baseline rate of randomly choosing a response. These are useful parameters in terms of accounting for participant response errors (see Appendix F).

3.2 A design space for DARC experiments

In the previous section, we have already indirectly introduced our design space. It consists of pairs of prospects, each consisting of a reward, a probability of obtaining the reward, and a delay. In this section we describe issues relating to the design space in more detail so that readers can better understand the Bayesian adaptive design approach, understand what we have done, and be better able to modify our code or extend it to fit their own needs. Readers are referred to Table 2 for the notation used in relation to the design space.

The 27 questions from Kirby (2009) represent one particular design space that has seen heavy use in delay discounting studies. It has 27 fixed questions (each represented by a row in a table), with 4 question variables (sooner reward amount R^a , delay of sooner reward D^a , later reward amount R^b and delay of later reward D^b), with rewards considered to be certain. The questions asked using this model are pre-set, such that information gathered during the experiment is wastefully ignored.

Our design space $\mathcal{D}$ can also be represented by a table, with 6 question variables R^a , D^a , P^a , R^b , D^b and P^b . Again, any row in this table represents a particular design, $d = [\mathcal{P}^a, \mathcal{P}^b]$. However, unlike in Kirby (2009), $\mathcal{D}$ contains hundreds or thousands of possible questions, only a small proportion of which will be used in any particular experiment and which are adaptively chosen to be maximally informative for the particular participant.

The software provided is very flexible, so we can adjust the design space to the needs of an experiment (see Table 3). We can choose to explore either delayed choice³, risky choice⁴, or both. For example, we can apply this to inter-temporal choice tasks by setting $P^a = P^b = 1$ and specifying a range of allowable delays such as $D^a = 0$ for no front-end delays, and D^b spanning many values from minutes to hours, to days or years. Likewise, we can adapt this for inter-temporal choice with a magnitude effect by ensuring we offer a wide range of delayed reward values (e.g. $R^b = [10, 100, 1000]$). We can turn this from a time discounting experiment into a risky choice experiment simply by setting $D^a = D^b = 0$ such that both prospects are immediate. We can also optionally fix $\mathcal{P}^a$ to be certain and $\mathcal{P}^b$ as risky by setting $P^a = 1$ and $0 < P^b < 1$.

While the toolbox is modular, and highly automated, there are some complexities which arise from trying to model real behaviour with simple parametric models. Until we find the ‘true’ model describing how agents behave in experiments, we will have an element of model misspecification. We may also have some degree of redundancy in how different parameters can account for different aspects of the data. To deal with these inevitable issues that affect all cognitive modelling efforts, we sometimes need to partially restrict the candidate designs available to the design selector on each experimental trial to ensure a good spread in the questions over features that are of no consequence to the model and therefore do not effect the theoretical optimality of the questions. When needed, these are dynamically chosen for each experimental trial in isolation using a heuristic described in Appendix G.

4. Worked example: temporal discounting experiments

Our approach is exceedingly general, and could be applied to any two-choice decision making task that we can also provide a generative model for. In the previous section we begun to define a specific class of models and appropriate design space for our topic of interest, DARC experiments. In this section we provide a worked example for temporal discounting, before showing how this can easily be applied to other DARC experiments and models.

4.1 Building the adaptive experiment

In order to make the most efficient inferences about a participant’s discount rate, we assume a parametric form for the discount function, namely the widely used 1-parameter hyperbolic function

3. Delay discounting only is a 4-dimensional design space, $d = (R^a, D^a, 1, R^b, D^b, 1)$ where $D^a \geq 0$, $D^b > 0$, $D^b > D^a$, and $R^b > R^a$ tends to be imposed.

4. Probability discounting only is a 4-dimensional design space, $d = (R^a, 0, P^a, R^b, 0, P^b)$ where $R^b > R^a$, $0 < P^b < 1$ and typically $P^a = 1$.

Experiment type	$u(R)$	$d(D)$	$\pi(P)$	R^a	D^a	P^a	R^b	D^b	P^b
Inter-temporal choice	R	$\frac{1}{1+k \cdot D}$	P	*	*	1	*	*	1
Inter-temporal choice	R	$\frac{1}{1+(\exp(m \cdot \log(R)+c)) \cdot D}$	P	*	*	1	*	*	1
Risky choice	R	D	$\frac{1}{1+h \cdot \frac{1-P}{P}}$	*	0	*	*	0	*
Risky inter-temporal choice	R	$\frac{1}{1+k \cdot D}$	$\frac{1}{1+h \cdot \frac{1-P}{P}}$	*	*	*	*	*	*

Table 3: A summary of the models and design space restrictions for each of the experiment types highlighted in this paper. The modular nature of the approach should be clear, there are specific pieces of code corresponding to elements in the table to implement a specific experiment (row). Entries labelled (*) are free variables that can be either single values or vectors of experimenter-defined allowed values. Note that the actual parameters estimated are $\log(k)$ and h . We focus on the 1-parameter hyperbolic discount function, but our toolbox can implement alternatives (see the toolbox software for examples).

(Mazur, 1987, see Table 3). We stress that it is trivial to adapt our software implementation to use alternative discount functions. The parameters of this model are the log discount rate $\log(k)$, the error rate ϵ , and comparison acuity α .

For this model, we define specific restrictions on the design space, appropriate for estimating a participant’s discount rate. We use no front-end delay $D^a = 0$. We also do not want to estimate the magnitude effect so we restrict ourselves to a single reward magnitude $R^b = £100$. Rewards are certain $P^a = P^b = 1$. We define a set of possible delays for the later reward as $D^b = (1, 2, 3, 4, 5, 6, 7, 8, 9, 12 \text{ hours}; 1, 2, 3, 4, 5, 6, 7 \text{ days}; 2, 3, 4 \text{ weeks}; 3, 4, 5, 6, 8, 9 \text{ months}; 1, 2, 3, 4, 5, 6, 7, 8, 10, 15, 20, 25 \text{ years})$, where a month is defined as 30 days.

The prior over $\log(k)$ allows experimenters to utilise prior knowledge about their participants, and because of this there is no single correct prior that can be recommended. For example, if we are testing a population suffering from anorexia nervosa, we would want our prior over $\log(k)$ to be centred on -5.9 (based upon Decker, Figner, and Steinglass, 2015), but if we were testing those suffering from major depressive disorder, our prior should be centred on -3.2 (based upon Pulcu, Trotter, Thomas, McFarquhar, Juhasz, Sahakian, Deakin, Zahn, Anderson, and Elliott, 2013). Rather than being an inconvenience, we can utilise our domain knowledge in order to avoid wasting time and resources during an experiment, which is the entire aim of our adaptive approach. Our software makes it easy for experimenters to provide their own prior knowledge about discount rates of their testing population.

The mode of the prior over $\log(k)$ also needs to take into account the reward magnitude on offer. There is a body of evidence suggesting our discount rates are strongly affected by the reward magnitude (see for example Johnson and Bickel, 2002), in a manner that is well modelled by a power law (ie. a straight line in $\log(\text{reward}), \log(k)$ space, Vincent, 2016). So even if we do not want to characterise the magnitude effect, we should not ignore that it exists. We recommend researchers pay particular attention to this issue to avoid erroneous priors based on previous research conducted on the right population, but at a different reward magnitude.

Our recommended uninformative prior, for rewards in the order of £100, was based on a brief survey of the literature including [Green, Myerson, Lichtman, Rosen, and Fry \(1996\)](#), [Pulcu et al. \(2013\)](#), and [Decker et al. \(2015\)](#). Values of $\log(k)$ range from approximately -2.5 (bipolar disorder) to approximately -6 (anorexia nervosa). Therefore the prior used was centred on the middle of this range, so our recommended uninformative prior is $\log(k) \sim \text{Normal}(-4.25, 1.5)$.

4.2 Comparison to alternative fixed and adaptive methods

Figure 4 demonstrates 4 different experimental protocols (rows top to bottom; [Kirby \(2009\)](#), [Koffarnus and Bickel \(2014\)](#), [Frye et al. \(2016\)](#), and our approach) to measuring delay discounting of simulated participants. We simulated participants with high, medium and low temporal discounting (rows). Discount rates used correspond to those with major depressive disorder ($k = 0.04$, half life of 25 days, [Pulcu et al., 2013](#)), upper income older adults ($k = 0.01$, half life of 100 days, [Green et al., 1996](#)), and those with anorexia nervosa ($k = 0.0028$, half life of 357 days, [Decker et al., 2015](#)). These examples provide an overview of the different methods in terms of how their experimental designs (points) unfold. Intuitively, we can interpret designs which fall far from the indifference curve to be less informative than those which lie closer to the indifference curve. We also represent the uncertainty in discount rates by plotting a random subset of 100 draws from the posterior over $\log(k)$. High uncertainty is bad, and is shown as a broad range of credible discount functions.

4.3 Efficiency of our methods compared to others

Figure 5 shows the result of a large parameter estimation simulation. We simulated 500 participants whose log discount rates were drawn from our prior over $\log(k)$. Each of these simulated participants responded to the 3 competitor delay discounting experimental protocols and our own. The figure shows how our uncertainty (measured as posterior entropy over $\log(k)$) decreases over trials of each of the experiments.

Of these approaches, the widely used method of ([Kirby, 2009](#), Figure 5) was the worst in that the rate at which our uncertainty decreases is lowest, and the total amount of uncertainty we can remove is limited. The limitations should be clear when examining the design space, represented by data points in Figure 4, top row.

The approach of ([Koffarnus and Bickel, 2014](#), Figure 5) is promising. It has a relatively rapid decrease in uncertainty, indicating that each trial results in learning a lot about discount rates. The main limitation however is that the approach does not scale; the procedure is only defined for a maximum of 5 trials.

The approach of ([Frye et al., 2016](#), Figure 5) is also promising. It is more efficient than the approach of [Kirby \(2009\)](#), and can be scaled up to high numbers of trials by adding more delays and more trials per delay. The approach is not quite as efficient as [Koffarnus and Bickel \(2014\)](#) however, but can result in more precise discount rate estimates once more than 10 trials worth of data have been obtained. This approach is quite variable however, on average the approach decreases uncertainty at an acceptable level, but on occasion the approach can be as inefficient as the method of [Kirby \(2009\)](#).

Our approach however (Figure 5) is, on average, the most efficient. We see a very rapid decline in uncertainty, which beats the method of [Koffarnus and Bickel \(2014\)](#). While the method of [Frye et al. \(2016\)](#) can also go beyond 5 trials, our method can do this in a more efficient manner. For example, there are times when the method of [Frye et al. \(2016\)](#) would require double the amount

Figure 4: Illustration of the different experimental approaches (rows) for delay discounting, applied to simulated observers with high, medium, and low discount rates (columns). Solid black lines show point estimates (posterior means) of the hyperbolic discount function, along with a representation of our uncertainty (light grey lines are samples from the posterior over $\log(k)$). Circles represent designs posed to the simulated participant. Filled circles represent choosing the immediate reward, empty circles the later reward. Simulated participants had true parameters $\alpha = 2$ and $\varepsilon = 0.01$ were used, but ε was fixed, while inference was conducted over α . All x-axes are cropped to 0-365 days delay for ease of visual comparison.

of trials as our approach in order to reach the same level of precision in the estimated discount rate. To emphasise the benefits of our approach; the amount of information gained, about a simulated participant, using 27 trials of the Kirby (2009) can be achieved in approximately 3 of our adaptive trials and is comfortably exceeded with 4 trials.

Figure 5: Comparison of how our uncertainty about log delay discount rates is reduced as a function of trials. The plot is based upon 500 simulated observers, each of which have a log discount rate sampled from our prior over $\log(k)$. Each simulated participant was subjected to the 4 experimental methods under consideration. Parameters α and ε as in Figure 4. Solid lines represent the median posterior entropy over the simulated experiments, shaded regions represent 50% credible regions for visual clarity.

4.4 Parameter recovery simulations

Figure 6 shows how each of the experimental approaches are able to recover known time discount rates of simulated participants. The Kirby (2009) method provides reasonable estimates, but the 95% credible regions are quite broad given that 27 designs were used. However we can see extreme quantisation in the inferred values. Because of the design space used in this method, it is not possible to make very fine-grained inferences about a participant’s discount rate. Instead, it is more appropriate to think of putting them into one of a number of discounting categories.

The 5-trial protocol of Koffarnus and Bickel (2014) is interesting in that it overall provides good estimates, but there are some interesting failure cases. We believe these are caused by an implicit assumption that answers are accurate – when incorrect responses are given it is likely to lead to wildly inaccurate estimates. This will have implications for data collected using this paradigm. The same quantisation issues reoccur here, so fine-grained evaluation of discount rates are not possible.

Figure 6: Parameter recovery of true time discounting rates for the existing approaches (top row) and our approach (bottom row). Each point and error bar represents the mean and 95% credible intervals of the posterior over $\log(k)$ of one simulated discounting experiment. Parameters α and ε as in Figure 4.

The method of [Frye et al. \(2016\)](#) run with 20 trials is impressive for a heuristic approach. The estimates are both accurate, and precise. We should certainly favour this approach over that of [Kirby \(2009\)](#).

The parameter recovery performance for our method is, however, even better. After 20 trials we can again see that inferred discount rates are accurate and precise. We rely on Figure 5 however for a clearer comparison of the two approaches, which shows our method is clearly more efficient than the heuristic method of [Frye et al. \(2016\)](#).

Figure 7: The role of prior beliefs on the estimated $\log(k)$ using our method. In the low data regime, the prior has a strong influence upon our beliefs, but as more data is obtained, the influence of the prior diminishes. Parameters α and ε as in Figure 4.

4.5 Role of the prior

Figure 7 shows the influence of the prior on our posterior beliefs of $\log(k)$ for low to high amounts of trial data. We can see that in the low data regime the prior has the effect of drawing our estimates toward the prior mean. If a participant has an extremely high discount rate of $\log(k) = -1$, then after 2–4 trials worth of data, our estimate is attenuated toward the prior mean. However, as more trial data arrives, our posterior shifts and extremely high discount rates become more credible.

This influence of the prior on both our beliefs, and the subsequent influence on what question is deemed most informative is a desirable property. Ours is the only method where our prior beliefs that extreme discount rates are actually rare are taken into account. This is a desirable effect if our prior truly reflects our beliefs about how discount rates are distributed across our participant

population. However, if our participant population does deviate from our prior, then it does not take many trials for the data to assign appropriate credibility to the participant’s true discount rate.

5. Implementing other DARC experiments

So far we have presented advances to Bayesian adaptive design, and outlined a flexible approach of applying this to DARC experiments. We have thoroughly demonstrated the benefits of this approach in the context of delay discounting experiments, as compared to some existing approaches. However the real power of this toolbox is its generality – in this section we demonstrate how our approach can easily be applied to 3 additional experiment types. Namely, delayed (i.e. inter-temporal) choice with magnitude effect, risky choice (i.e. probability discounting), and combined delayed and risky choice experiments. Table 3 lists the 4 experimental approaches (models and design spaces) we focus upon in this paper which essentially plug-in to the core toolbox. At the end of this section we briefly describe our successful parameter recovery simulations to highlight that the toolbox does indeed function very well not only for temporal discounting, but for a range of experimental paradigms.

We implement these with the popular hyperbolic discount functions, but again stress that it is trivial to extend these to alternative discount functions. Indeed, we provide examples in our software to demonstrate how to, for example, extend to a 2-parameter hyperboloid model of risky choice.

5.1 Delayed choice, with the magnitude effect

The DARC Toolbox allows for extensions to the utility-based approach. We can, for example, extend Equation 8 to

$$V(\mathcal{P}) = u(R) \cdot \pi(P, R) \cdot d(D, R). \quad (12)$$

By making $d(\cdot)$ a function of both the delay and reward, we can capture the magnitude effect where present bias decreases for higher magnitude delayed rewards (Kirby and Maraković, 1996; Vincent, 2016). Doing the same for probability weighting, we could capture the peanuts effect where present bias increases for higher magnitude probabilistic rewards (Brown and Cross, 2000; Green, Myerson, and O’Donoghue, 1999).

The second experimental procedure presented addresses the question, “How does a participant’s discount rate vary with reward magnitude?” Again, we assume people discount according to the 1-parameter hyperbolic function (see Table 3), however, we define that the log of the discount rate varies linearly with the log of the reward magnitude (as in Vincent, 2016). The parameters used to estimate the magnitude effect are: m for the slope of how log discount rate changes with log reward magnitude, c for the intercept of this line, and again ϵ and α as before.

The form of the design space is similar, but testing values of R^b over a number of orders of magnitude clearly allows much better estimation of the slope of the magnitude effect. For the present work we defined, $R^b = (10, 100, 1000)$.

When estimating the magnitude effect, we set the priors over the slope and intercept as $m \sim \text{Normal}(-0.243, 5)$, and $c \sim \text{Normal}(-4.716, 100)$, respectively. The mean slope was chosen as -0.243 and mean intercept as -4.716 because it was the most probable value from a survey of previous studies presented in Johnson and Bickel (2002), analysed by Vincent (2016). The standard deviations were chosen to reflect our increased uncertainty in the intercept.

5.2 Risky choice (probability discounting)

Rachlin et al. (1991) showed that discounting of the probability of receiving a delayed reward (compared to a certain immediate reward) can also be described using the 1-parameter hyperbolic discount function. The process of estimating the probability discounting parameter h (see Table 3) is therefore identical to that of discount rates $\log(k)$. The descriptive adequacy of the hyperbolic discount functions for probabilistic outcomes (Blackburn and El-Deredy, 2013; Bruce, Bruce, Catley, Lynch, Goggin, Reed, Lim, Strober, Glusman, Ness, and Jarmolowicz, 2015; Estle, Green, Myerson, and Holt, 2007; Lawyer, Williams, Prihodova, Rollins, and Lester, 2010; Rasmussen, Lawyer, and Reilly, 2010) will allow researchers to more efficiently collect data which can contribute to more mechanism-focussed questions. For example, there is an ongoing debate about whether delay and probability discounting is driven by a single process, or two independent processes (e.g. Myerson, 2004), but in terms of our task here of data collection, we can remain impartial to this mechanistic debate.

The response model remains essentially the same as for time discounting, however the discount function is changed to represent that participants are now discounting odds against receiving the later reward $\frac{1-p^b}{p^b}$, as opposed to time delay D^b (see Table 3).

A slightly different approach was taken to arrive at a prior over probability discount rate h than the time discount rate k . We should note that when $h = 1$ a participant would be risk-neutral, using a maximum expected value strategy (Green et al., 1999). For example, they would be indifferent between £10 and 10% of £100. However, when $0 < h < 1$, a participant would be risk-seeking, preferring the latter option, and when $h > 1$ they are risk-averse preferring the former option. Here we will restrict $h > 0$. There are a number of studies showing people vary along this risk-aversion risk-seeking dimension. We propose using as prior where h is Gamma distributed with parameters mode = 1 and $\sigma^2 = 4$.

5.3 Delayed and risky choice

Yi, de la Piedad, and Bickel (2006) provided the first evidence that discounting of delayed *and* probabilistic rewards can be well described by a hyperbolic model of discounting according to a summative overall discount rate. However, Vanderveldt, Ariana, Green, Leonard, and Myerson, Joel (2015) provide evidence that the subjective value function is well described by a multiplicative interaction between probability and delay discount rates (as in Equation 8). They used a 2-parameter hyperboloid model, but in keeping with our attempt to retain simplicity we apply this multiplicative interaction to the 1-parameter hyperbolic function (see Table 3, bottom row).

In theory, we could freely vary all 6 aspects of the design space $(R^a, D^a, P^a, R^b, D^b, P^b)$. However, here we constrain prospect a to be immediate and certain, but the toolbox can vary the reward level. Prospect b is set to have a reward value of £100, but the toolbox can vary the delay and the probability.

The parameters, $\log(k)$ and h have the same meaning as the previously discussed time and probability discounting models, and so the same priors for each parameter were used.

5.4 Parameter recovery simulations

One of the strengths of our approaches is that we can apply it to multiple different delayed and risky choice experiments. While two of the alternative methods we have examined (Frye et al., 2016;

(Koffarnus and Bickel, 2014) could be adapted to work with risky choice (probability discounting), they are not general enough to be applied to all models. In particular, they would not work with models which have a discount *surface* that varies over 2 dimensions such as with combined delay and probability discounting, nor delay discounting where discount rate varies with reward magnitude. Therefore it is not possible to compare our method against these alternative methods. But the performance of our approach in general has already been established in the previous section with the basic time discounting experiment.

However, to confirm that the usefulness of the methods extend to these other experimental procedures (time discounting with magnitude effect, probability discounting, and combined probability and time discounting), we conducted a similar parameter recovery procedure on our remaining discounting models (see Figure 8). We again find that our approach is able to do well at recovering known discounting parameters of simulated participants. For each of the discounting experiment types, we examine a cross section of true parameter values that span a reasonable range for that model, and these are shown on the x-axes. The role of the prior with low trial numbers still occurs (as we showed for the time discounting model, Figure 7) but are not shown here; we present parameter estimates after a total of 30 simulated trials.

6. Discussion

6.1 Example use cases

To emphasise the generality of the toolbox we have presented, we will briefly discuss a few example use cases. These examples, and more practical implementation advice is provided along with the toolbox at <https://github.com/drbenvincent/darc-experiments-matlab> or <https://github.com/drbenvincent/darc-experiments-python>.

Thus far we have seen examples with positive rewards. It is trivial to run DARC experiments with losses – one need simply alter the design space, for example such that R^a and R^b consist of negative values. This has been done and one of us have been using the toolbox to collect data about behaviours in both the gain and loss domains.

The toolbox will provide optimal designs, d^* , and the user can use this to present pairs of prospects to participants using whatever presentation method they would like. By default, we assume experiments are being run with monetary rewards, and options are presented as text in clickable buttons. Various built-in options are provided (see Figure 9) including money (in various commodities) and non-monetary rewards, and provide either delay or date framing of delays (DeHart, 2015; Read, Frederick, Orsel, and Rahman, 2005). This flexibility has meant that the present methods have already been used for data collection for a range of different commodity types (Skrynka and Vincent, 2017). Researchers can of course update the existing code to present in whatever manner is appropriate for their experiment. For example, one could test the effects of verbal versus graphical portrayal of risk as in (Drichoutis and Lusk, 2017). Or if testing non-humans, the experimenter would need to write some interface code with their particular experimental apparatus.

It is also possible to simultaneously run multiple experiments. The software is built such that researchers could interleave trials from a delay discounting experiment and a risky choice experiment, or a discounting experiment with an exponential discounting function and a hyperbolic discounting, or any such combination. Examples of how to do this are shown along with the code.

It is also entirely possible to inject manually derived experimental designs into the sequence of automatically generated designs. This might be useful if experimenters wish to display catch trials,

Figure 8: Parameter recovery for remaining discounting tasks (rows). Many simulations (points and 95% credible intervals) of 30 trials were run to see how close the estimated parameters (y axes) match the true parameters of a simulated observer (x axes). Our method can correctly infer known parameter values of simulated observers for hyperbolic discounting with magnitude effect (top row; parameter sweep of $(m, c) = (-1, -0.5)$ to $(-0.5, -2.5)$), probability discounting (middle row) and combined probability and time discounting (bottom row; parameter sweep of $(\log(k), h) = (-8, 0.1)$ to $(-1, 10)$). Parameters α and ϵ as in Figure 4.

Figure 9: Examples of existing built-in ways in which the experimenter can modify framing of the optimal designs, d^* produced by the toolbox. Rewards can be either monetary (in various currencies) or non-monetary in nature. Delays can be framed as delays from now, or as specific dates in the future.

or present pairs of prospects that for a given decision difficulty or reward conflict, that is, distance from indifference, $V(\mathcal{P}^a) - V(\mathcal{P}^b)$. This could be useful in situations where researchers are trying to map behaviour (or EEG or fMRI measures) as a function of distance from indifference points. For example, [Eppinger, Heekeren, and Li \(2017\)](#) found differences in BOLD signals for low vs high decision difficulty. Similarly, [Lin, Saunders, Hutcherson, and Inzlicht \(2017\)](#) find pupil dilation and midfrontal theta signals (from EEG) systematically vary as a function of decision difficulty. Using our adaptive methods would allow very rapid real-time assessment of a participant’s present behaviour (as captured by the posterior over θ) and allow trials to be injected with systematic decision difficulty / reward conflict levels.

There are many more possible imaginable use cases. Our approach has been to make the toolbox as flexible as possible, and our documentation and code should allow for researchers to modify it for their own custom purposes.

6.2 Analysis of collected experimental data

After running an adaptive experiment, our toolbox provides raw behavioural data, point estimates (e.g. posterior median) and full posterior distributions over parameters. While it is entirely rea-

sonable to enter the point estimates of parameters into one’s larger experimental dataset, it may be advantageous to use more sophisticated methods.

The recommended workflow is to use the adaptive experimental methods presented in this paper in order to obtain raw behavioural data, and then to use hierarchical Bayesian analysis methods (e.g. [Ahn, Haines, and Zhang, 2017](#); [Nilsson et al., 2011](#); [Vincent, 2016](#)) with study-specific priors in order to draw research conclusions.

The next best approach would be to use non-hierarchical but still Bayesian analysis methods – which corresponds to taking the estimated posterior estimates from the toolbox (if the user has customised the priors to be study-specific). However, our toolbox uses 50,000 particles to represent the posterior, and if greater precision is needed, then MCMC-based methods in the packages mentioned above are suitable for generating hundreds of thousands of samples from the posterior.

In turn, the next best approach would be to use maximum likelihood methods ([Franck, Kofarnus, House, and Bickel, 2015](#); [Gilroy, Franck, and Hantula, 2017](#); [Myung et al., 2013a](#); [Young, 2018](#)). These approaches can work very well, and are convenient when researchers are not able or willing to formalise any prior beliefs, however sometimes the data is not sufficient (in the absence of priors) to uniquely identify the maximum likelihood parameters.

6.3 Cautions

While the adaptive methods we have outlined here represent a significant advance in running efficient experiments, there are a couple of caveats which have to be kept in mind.

We assumed each experimental trial to be an independent event. This is an often-made assumption in the modelling of human behaviour, but we should bear in mind that trial-to-trial and order effects are observed ([Robles and Vargas, 2007,0](#)).

In virtually all work of this type, the response model makes an assumption that the participant’s responses are essentially fixed, in that the parameters that describe their behaviour are unchanging (i.e. stationary). This assumption may not in fact be true: it could be, for example, that participants are not entirely clear about their internal preferences, and running through an experiment may lead to refinement of their discount rates.

If real or hypothetical monetary rewards are offered, then experimenters must take care with priors over parameters when using different currencies. Many of the studies which report discount rates, or log discount rates have a meaning in the currency in which those studies were undertaken, often in U.S. dollars. Therefore, experimenters should either attempt to work out an exchange rate (which could involve complications) or to simply provide less informative priors.

6.4 Conclusion

We have presented methodological advances to the Bayesian adaptive design approach, along with a software toolbox implementation. This is a powerful and extendable resource which can be applied to any 2-choice task, however here we focussed on evaluating our approach using a variety of DARC experiments. We provide adaptive experiments for delay discounting, delay discounting with magnitude effects, probability discounting (risky choice), and combined delay discounting and risky choice. Further, these procedures are computationally efficient enough to run in real time on a mid-range laptop. We found that our method is superior to a number of other procedures in terms of quickly reducing our uncertainty about participant’s behaviours (as captured by model parameters) in low numbers of experimental trials. Experimenters could then maintain accurate measures

even when running experiments with fewer trials, which is highly desirable when testing special populations or when using EEG, fMRI, or other imaging methods. Alternatively, even more precise measures of behaviour can be obtained in a given number of trials. Together with the hierarchical Bayesian data analysis methods (Nilsson et al., 2011; Vincent, 2016), we now have a state-of-the-art, fully Bayesian data collection and analysis pipeline to aid researchers reach robust and reliable research conclusions.

Acknowledgements

The authors thank Daniel Baker for comments on an early draft of this manuscript. Tom Rainforth's research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) ERC grant agreement no. 617071. However, the majority of this work was undertaken while he was in the Department of Engineering Science, University of Oxford, and was supported by a BP industrial grant. We also thank Aaron Johnson for the invitation to run a workshop at Concordia University on the methods presented in this paper, which gave rise to useful feedback for the paper and software.

Software

The adaptive experiments are freely available to download at <https://github.com/drbenvincent/darc-experiments-matlab> or <https://github.com/drbenvincent/darc-experiments-python>. Examples are provided online to demonstrate how to run experiments with a few simple commands. We are interested in hearing about user-experiences, feature requests, and bug reports so that we can improve the toolbox over time. The DARC Toolbox is written in a modular and expandable fashion to allow researchers to construct and contribute additional models.

Appendix A. Derivations for Bayesian adaptive design metric estimators

Our starting point to derive both the nested Monte Carlo (NMC) and our new non-nested estimator is the expected utility given in (3) of the main paper and repeated below for convenience

$$\bar{U}(d) = \int_{\mathcal{Y}} p(y|d) \left(\int_{\Theta} p(\theta|y, d) \log(p(\theta|y, d)) d\theta - \int_{\Theta} p(\theta) \log(p(\theta)) d\theta \right) dy. \quad (13)$$

We start by deriving the nested estimator, for which the first step, as shown by [Myung et al. \(2013a\)](#), is to use Bayes rule to rearrange (13) into a form that only contains known terms as follows (remembering that $p(\theta) = p(\theta|d)$)

$$\begin{aligned} \bar{U}(d) &= \int_{\mathcal{Y}} p(y|d) \left(\int_{\Theta} p(\theta|y, d) \log(p(\theta|y, d)) d\theta - \int_{\Theta} p(\theta) \log(p(\theta)) d\theta \right) dy \\ &= \int_{\mathcal{Y}} \int_{\Theta} p(y, \theta|d) \log \left(\frac{p(\theta|y, d)}{p(\theta)} \right) d\theta dy \\ &= \int_{\mathcal{Y}} \int_{\Theta} p(y, \theta|d) \log \left(\frac{p(y|\theta, d)}{p(y|d)} \right) d\theta dy \\ &= \int_{\mathcal{Y}} \int_{\Theta} p(y, \theta|d) \log(p(y|\theta, d)) d\theta dy - \int_{\mathcal{Y}} p(y|d) \log(p(y|d)) dy. \end{aligned} \quad (14)$$

Though still problematic, this reformulation offers a number of advantages. One of these is that the first of the two terms now takes the form of a standard inference problem and is amenable to, for example, Monte Carlo estimation as follows

$$\bar{V}_1(d) = \int_{\mathcal{Y}} \int_{\Theta} p(y, \theta|d) \log(p(y|\theta, d)) d\theta dy \approx \frac{1}{N} \sum_{n=1}^N \log(p(y_n|\theta_n, d)) \quad (15)$$

where $\theta_n \sim p(\theta)$ and $y_n \sim p(y|\theta = \theta_n, d)$. We can do this because, although the full integral is not analytic, the integrand in this term can now be evaluated exactly. As such, the integral of this term can be evaluated to a given accuracy orders of magnitude faster than the utility evaluation for grid-search approaches.

By contrast, the second term of (14) is not directly amenable to standard MC estimation as the marginal $p(y|d)$ represents an expectation and taking its logarithm represents a nonlinear functional mapping. However, we can still use a NMC estimator as follows

$$\bar{V}_2(d) = \int_{\mathcal{Y}} p(y|d) \log(p(y|d)) dy \approx \frac{1}{N} \sum_{n=1}^N \log \left(\frac{1}{M} \sum_{m=1}^M p(y_n|\theta_{n,m}, d) \right) \quad (16)$$

where $\theta_{n,m} \sim p(\theta)$ and $y_n \sim p(y|d)$. We can sample the latter by first sampling an otherwise unused $\theta_{n,0} \sim p(\theta)$ and then sampling a single corresponding $y_n \sim p(y|\theta_{n,0}, d)$.

Putting (15) and (16) together (and renaming θ_n from (15) as $\theta_{n,0}$ for notational consistency with (16)) we now have the following complete estimator, given in (5) in the main paper and implicitly used by [Myung, Cavagnaro, and Pitt \(2013b\)](#) amongst others,

$$\bar{U}(d) \approx \hat{U}_{NMC}(d) := \frac{1}{N} \sum_{n=1}^N \left[\log(p(y_n|\theta_{n,0}, d)) - \log \left(\frac{1}{M} \sum_{m=1}^M p(y_n|\theta_{n,m}, d) \right) \right] \quad (17)$$

where $\theta_{n,m} \sim p(\theta) \forall m \in 0 : M, n \in 1 : N$ and $y_n \sim p(y|\theta = \theta_{n,0}, d) \forall n \in 1 : N$.

As explained in the main paper, this NMC estimator is substantially less efficient than a standard Monte Carlo estimator (Rainforth et al., 2017) and so we seek a more efficient means of estimating $\bar{V}_2$. To do this, we note that if y can only take on one of C possible values ($y_1, \dots, y_C$), we can achieve significant improvements in the convergence rate by calculating the outer integral in $\bar{V}_2$ analytically through enumeration, instead of using Monte Carlo. Namely we have

$$\int_{\mathcal{Y}} p(y|d) \log(p(y|d)) dy = \sum_{c=1}^C p(y_c|d) \log(p(y_c|d)), \quad (18)$$

such that we no longer have to worry about an integral of non-analytic term at all. Instead, we need only estimate

$$\hat{P}_c = \frac{1}{N} \sum_{n=1}^N p(y_c|\theta_n, d) \approx p(y_c|d) \quad (19)$$

for $c = 1, \dots, C$ and substitute $p(y_c|d)$ for $\hat{P}_c$ in (18), noting that this represents a deterministic mapping of our estimates and that the y_c are a finite number of fixed constants rather than Monte Carlo samples.

We can also improve upon (15) by carrying out the integration over y analytically. This process is known as Rao-Blackwellisation (Casella and Robert, 1996) and is (in terms of variance of the estimate) never worse than vanilla Monte Carlo. Namely we have

$$\bar{V}_1(d) \approx \hat{V}_1(d) = \frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C p(y_c|\theta_n, d) \log(p(y_c|\theta_n, d)) \quad (20)$$

where $\theta_n \sim p(\theta)$ as before. Putting this together leads to our final estimator as per (6) in the main paper

$$\begin{aligned} \bar{U}(d) \approx \hat{U}(d) := & \frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C p(y_c|\theta_n, d) \log(p(y_c|\theta_n, d)) \\ & - \sum_{c=1}^C \left[\left(\frac{1}{N} \sum_{n=1}^N p(y_c|\theta_n, d) \right) \log \left(\frac{1}{N} \sum_{n=1}^N p(y_c|\theta_n, d) \right) \right] \end{aligned} \quad (21)$$

It is straightforward to show (see, for example, (Rainforth, 2017, Chapter 13)) that the mean squared error of $\hat{U}(d)$ to $\bar{U}(d)$ converges at a rate $O(1/N)$, i.e.

$$\mathbb{E} \left[(\hat{U}(d) - \bar{U}(d))^2 \right] \leq \frac{\sigma}{N} \quad (22)$$

as per standard Monte Carlo convergence, with σ being a constant equal to the standard deviation resulting from an the estimator using a single sample. Consequently, we have that as $N \rightarrow \infty$, the error in $\hat{U}(d)$ as an approximation of $\bar{U}(d)$ tends to zero and that this error decreases faster than that of $\hat{U}_{NMC}(d)$.

Appendix B. Sequential Experiments and Inference Procedure

From a mathematical perspective, we can generalize to the sequential design setting by incorporating data in the standard Bayesian fashion such that at experiment iteration t , we replace $p(\theta)$ with $p(\theta|d_{1:t-1}, y_{1:t-1})$, where $d_{1:t-1}$ and $y_{1:t-1}$ are respectively the designs and outcomes at previous iterations. The likelihood $p(y_t|\theta, d_t)$, on the other hand, is unchanged (presuming it is a parametric distribution) as, conditioned on θ and d , the current outcome is independent of the previous data. Putting this together, we get that the expected information gain criteria for the sequential case is

$$\begin{aligned} \bar{U}_t(d) = & \int_{\mathcal{Y}} \int_{\Theta} p(\theta|d_{1:t-1}, y_{1:t-1}) p(y_t|\theta, d_t) \log(p(y_t|\theta, d_t)) d\theta dy_t \\ & - \int_{\mathcal{Y}} p(y_t|y_{1:t-1}, d_{1:t}) \log(p(y_t|y_{1:t-1}, d_{1:t})) dy_t \end{aligned} \quad (23)$$

where

$$p(y_t|y_{1:t-1}, d_{1:t}) = \int_{\Theta} p(\theta|d_{1:t-1}, y_{1:t-1}) p(y_t|\theta, d_t) d\theta. \quad (24)$$

We can now see that these terms are the same as in the non-sequential case, except that expectations are taken with respect to $p(\theta|d_{1:t-1}, y_{1:t-1})$ rather than $p(\theta)$. Therefore, we can use the same Monte Carlo estimators (19), (20), and (6), simply changing how the θ_n are sampled.

Unlike $p(\theta)$, it is typically not possible to evaluate, or sample from, $p(\theta|d_{1:t-1}, y_{1:t-1})$ exactly – we need to perform Bayesian inference to estimate it. Noting the implicit assumption of the likelihood model that observations are independent of one other given the parameters we have

$$p(\theta|d_{1:t-1}, y_{1:t-1}) \propto p(\theta) \prod_{i=1}^{t-1} p(y_i|\theta, d_i) \quad (25)$$

which is a typical Bayesian inference problem. We now remember that other than the likelihood, which is unchanged, the estimation scheme introduced in Appendix A requires only samples of θ as input. Therefore we can use a Monte Carlo inference method to produce these required samples θ_n .

There are many suitable candidates for this inference such as Metropolis Hastings (Hastings, 1970), sequential Monte Carlo (SMC) (Doucet et al., 2000), and Hamiltonian Monte Carlo (Duane, Kennedy, Pendleton, and Roweth, 1987). The relative merit of these will depend on characteristics of the problem such as multi-modality, dimensionality of θ , and the total number of questions that will be asked. In this work we have decided to use population Monte Carlo (PMC) (Cappé et al., 2004). At a high level, the PMC algorithm targeting an unnormalized distribution $\pi(\theta)$ ($= p(\theta) \prod_{i=1}^{t-1} p(y_i|\theta, d_i)$ in our scenario) comprises of independently propagating a *population* of, say M , samples known as particles. At each iteration j each particle is propagated in a manner similar to Metropolis Hastings using a proposal $q(\theta^{j+1}|\theta^j)$.⁵ However, instead of applying an individual accept/reject step to each sample, they are each assigned an importance weight

$$w_m^{j+1} = \frac{\pi(\theta_m^{j+1})}{q(\theta_m^{j+1}|\theta_m^j)} \quad (26)$$

5. Though it is permitted to use a different form of proposal for each sample in the population, we do not do so in our approach and so omit this case for clarity.

where $m \in \{1, \dots, M\}$ is the sample index. The population is then *resampled* by sampling M new particles with replacement from the original population, with the probability of sampling each particle proportional to its weight w_m^{j+1} . This creates a new population of *unweighted* particles, where particles that had high weight before the resampling will typically be replicated multiple times, whilst low weight particles will typically have been removed. The rationale of this resampling is to reallocate computational resources to where they are needed – each of the duplicated particles will be propagated to different points at the next proposal step and thus form distinct samples. As the proposals are generally localized, replicating good particles will increase the number of samples in high probability regions, while eliminating those with negligible probability mass. Over time the particles will converge to the target distribution, such that they constitute samples from π . We refer the reader to [Cappé et al. \(2004\)](#) for further details.

In order to run this inference, six things need to be specified: the prior $p(\theta)$, the likelihood for a single response $p(y|\theta, d)$, the question response pairs $\{y_i, d_i\}_{i=1:t-1}$, the proposal $q(\theta^{j+1}|\theta^j)$, the population size M , and the number of inference iterations $j_{\max}$. The first two correspond to specifying the model, as detailed in the main paper, and are the main input provided by the user. The question response pairs are, of course, collected during the running of the experiment. For the proposal we use a non-isotropic Gaussian, with length scales set adaptively at run time using the previous samples of θ . Though we found this worked well in practice across the range of models we consider, it may be necessary to specify this in a more problem specific manner for more challenging models. The modularity of our approach means that this can easily be changed if required. The final two things to be set, M and $j_{\max}$, are parameters that dictate the accuracy of the inference and are set based on computational restrictions. We take as default $M = 5 \times 10^4$ and $j_{\max} = 5$ for which inference generally takes less than a second on a mid-range laptop. We note though, that they may need to be raised for problems where the inference is more challenging, or lowered in scenarios where time restrictions are even stricter.

Appendix C. Design optimisation details

Imagine that we are considering K designs with a total sample budget $\bar{N}K$ and decide to use the same number of samples, $\bar{N}$, to calculate an estimate $\hat{U}_k$ for the utility of each design d_k . After we have taken say n samples for each estimate, where $1 \leq n < \bar{N}$, then we will have a set of intermediate estimates $\hat{U}_{1:K}^n$ for each utility. Now remembering that our aim is to establish which design is best (i.e. which $\bar{U}_k$ is largest), we note that our intermediate estimates convey valuable information – they show that some designs are more likely to be optimal than others. Consequently, sticking to our original plan and using the same number of samples, $\bar{N} - n$, to complete our final estimates will be inefficient as some designs will be highly unlikely to be optimal and therefore it is pointless to spend a noticeable proportion of our computational budget to carefully refine their estimates. It is clearly more efficient, in terms of final optimisation problem, to take relatively more samples for the promising designs, so that we can better distinguish between their relative utilities, even if this results in very poor estimates for the low utility designs. This is the key idea of our approach - to adaptively change the number of samples used to estimate the utilities for different designs, based on the probability the design may end up being optimal given its intermediate estimates.

More specifically, given a total budget of $NL = \bar{N}K$ samples for all our estimates, then we will carry out L rounds of sampling (we take $L = 50$ as default), where at each round we adaptively allocate N samples (we take $N = 5 \times 10^4$ as default) between the different estimators in proportion

Algorithm 3 Design optimisation

Inputs: Reservoir $\theta_{1:M}$, likelihood $p(y|\theta, d)$, candidate designs $\mathcal{D} = d_{1:K}$, number of rounds L , number of samples per round N , annealing schedule $\gamma: \mathbb{Z}^+ \rightarrow \mathbb{R}$, minimum selection probability ratio ρ

Outputs: Chosen design d^*

```

1:  $\hat{U}_{1:K} \leftarrow 0$  ▷ Initialize estimates
2: for  $\ell = 1 : L$  do
3:    $\tilde{p}_{1:K} \leftarrow \hat{U}_{1:K}^{\gamma(\ell+1)}$  ▷ Unnormalized selection probability = annealed  $\hat{U}_k$ 
4:    $Z \leftarrow \sum_{k=1}^K \tilde{p}_k$  ▷ Normalization constant.
5:    $p_k \leftarrow \max\left(\frac{\rho}{N}, \frac{\tilde{p}_k}{Z}\right) \quad \forall k = 1, \dots, K$  ▷ Ensure probabilities are all at least minimum value
6:    $p_k \leftarrow \frac{p_k}{\sum_{k=1}^K p_k} \quad \forall k = 1, \dots, K$  ▷ Renormalize
7:    $n_{1:K} \sim \text{MULTINOMIAL}(N, p_{1:K})$  ▷ Sample number of samples to take for design
8:   for  $k = 1 : K$  do
9:      $\hat{\theta}_{1:n_k} \sim \text{DISCRETE}(\theta_{1:M})$  ▷ Draw  $n_k$  samples from reservoir
10:     $\hat{U}_k \leftarrow \text{UPDATEESTIMATE}(\hat{U}_k, \hat{\theta}_{1:n_k}, d_k)$  ▷ Use new samples to refine estimate
11:   end for
12: end for
13:  $k^* \leftarrow \text{argmax}_{k \in \{1, \dots, K\}} \hat{U}_k$  ▷ Best design is one with highest  $\hat{U}_k$ 
14: return  $d_{k^*}$ 

```

to how relatively promising their respective designs appear. To do this, we *sample a number of samples* to take for each design at each round in proportion to an annealed version of its utility estimate (line 7 of Algorithm 3). Namely, if $p_{k,\ell}$ is the probability each sample from our budget is assigned to design k at round ℓ and $\hat{U}_{k,\ell}$ is the corresponding running estimate for the utility of that design, then we have (line 3 of Algorithm 3)

$$p_{k,\ell} \propto (\hat{U}_{k,\ell})^{\gamma(\ell)} \quad (27)$$

where $\gamma: \mathbb{Z}^+ \rightarrow \mathbb{R}$ is an annealing function (see next paragraph). Thus we will, on average, add $\mathbb{E}[n_{k,\ell}] = p_{k,\ell}N$ samples of θ to the estimate for $\bar{U}_k$ at round ℓ (lines 9 and 10 of Algorithm 3). Specifically, the algorithm stores running estimates (not shown in Algorithm 3) for each $(\hat{V}_1)_k$, $(\hat{P}_c)_k$, and the number of estimates taken so far N_k . Using the likelihood $p(y|\theta, d)$ and the design d_k , each of these can then be updated with the n_k new samples from the reservoir $\{\hat{\theta}_j\}_{j=1:n_k}$ at each round. This corresponds to the UPDATE function in Algorithm 3.

The purpose of the annealing function is to control how aggressive our algorithm is in its allocation of computational resources to promising designs. At the early rounds, we are quite uncertain about the relative utilities and therefore wish to allocate samples fairly evenly. At the later rounds, we become more certain about the estimates so we become more aggressive about concentrating our efforts on promising designs. Therefore $\gamma(\ell)$ is set to increase with ℓ so that the algorithm becomes more aggressive as our estimates become better refined. By default, we take $\gamma(\ell)$ to vary linearly between $\gamma(1) = 0$ and $\gamma(L) = 50$.

We finish by noting a small, but important, subtlety of our method. As introduced so far, our algorithm is not guaranteed to converge. For example, it might be that $\hat{U}_{k^*} = 0$ after the first round of sampling where k^* indicates the true optimal design and $\bar{U}_{k^*} \neq 0$. Presuming that another design has a non zero estimate at this point, sampling naively from (27) would mean that no additional samples are ever added to $\hat{U}_{k^*}$, even if infinite rounds are undertaken, meaning that the true optimum will

be missed. To guard against this, we introduce a new parameter $0 < \rho \leq 1$ (by default we take $\rho = 0.01$). After calculating $p_{k,\ell}$ according to (27) and normalizing so that $\sum_{k=1}^K p_{k,\ell} = 1$, we update the probabilities by setting

$$p_{k,\ell} \leftarrow \max\left(\frac{\rho}{N}, p_{k,\ell}\right) \quad (28)$$

and finally renormalizing (lines 3 to 6 in Algorithm 3). This ensures that each $p_{k,\ell} > \frac{\rho}{2N}$ as the maximum possible value for $\sum_{k=1}^K p_{k,\ell}$ is $1 + \frac{N-1}{N}$ before renormalizing (which occurs when all but one $p_{k,\ell}$ equals 0 prior to applying (28)) and so the minimum possible value of $p_{k,\ell}$ after renormalizing is $\frac{\rho}{N(1+\frac{N-1}{N})} > \frac{\rho}{2N}$. Consequently,

$$\mathbb{E}\left[\sum_{\ell=1}^L n_{k,\ell}\right] > \frac{\rho L}{2} \quad (29)$$

and so for any finite value of ρ , each estimate will, in expectation, be assigned infinitely many samples as $L \rightarrow \infty$, regardless of the true utilities and the values of $\gamma(\ell)$.⁶ As we show in the next section, we can use this to prove the convergence of the algorithm, in the sense of guaranteeing the optimal design will be found given infinite computational budget.

Appendix D. Proof of Theorem 1

We start our proof by noting that by the definition of $\tau(\cdot)$ then $L \rightarrow \infty$ also predicates that $M \rightarrow \infty$, while $M \geq NL$ also ensures the reservoir is large enough to not need to reuse samples. By the convergence results shown for PMC in [Chopin et al. \(2004\)](#) this means that Monte Carlo estimates of measurable functions made with respect to our infinite reservoir will converge almost surely.

As the $n_{k,\ell}$ are not independent, we break each down into two terms $n_{k,\ell} = m_{k,\ell} + r_{k,\ell}$ where $m_{k,\ell} \sim \text{BINOMIAL}(N, \frac{\rho}{2N})$ and $r_{k,\ell} \sim \text{BINOMIAL}(N, p_{k,\ell} - \frac{\rho}{2N})$ (which induces the appropriate distribution on $n_{k,\ell}$). For any given k , all the $m_{k,\ell}$ are mutually independent and $P(m_{k,\ell} \geq 1) = 1 - (1 - \frac{\rho}{2N})^N \rightarrow 1 - \exp(-\rho/2)$ as $N \rightarrow \infty$. Therefore, for any possible N (including $N \rightarrow \infty$) we have $\sum_{\ell=1}^{\infty} P(m_{k,\ell} \geq 1) = \infty$ and so by the second Borel-Cantelli lemma, the event $m_{k,\ell} \geq 1$ occurs infinitely often with probability 1. Thus, as each $r_{k,\ell} \geq 0$, it must also be the case that $P(\lim_{L \rightarrow \infty} \sum_{\ell=1}^L n_{k,\ell} = \infty) = 1$.

Together these two results ensure almost sure convergence of each of the K estimates $\hat{U}_k$ (i.e. all the estimates become exact) at each iteration and so we must choose the optimal design (or one the equally best designs if there is a tie) almost surely in the limit $L \rightarrow \infty$.

Appendix E. Alternative delay discounting procedures

E.1 The method of [Frye et al. \(2016\)](#)

Algorithm 4 describes the approach taken by [Frye et al. \(2016\)](#) (their Supplementary File 2), focussing upon the selection of the experimental design (D^a, D^b, R^a, R^b) .

6. Note that one should, if desired, be able to instead use an upper confidence bounding strategy ([Auer, 2002](#)) to produce a zero-regret strategy whilst maintaining convergence

Algorithm 4 Method of [Frye et al. \(2016\)](#)**Inputs:** Vector of delays, delayed reward magnitude R^b , trials per delay T **Outputs:** Vector of responses for each trial

```

1:  $D^a \leftarrow 0$ 
2: for each  $d$  in delays do
3:    $D^b \leftarrow d$ 
4:   for  $t = 1 : T$  do
5:     if  $t = 1$  then
6:        $\kappa \leftarrow 0.25, \quad R^a \leftarrow 0.5R^b$ 
7:     else if responses[ $t - 1$ ] = sooner then
8:        $R^a \leftarrow R^a - R^b \kappa$  ▷ decrease sooner reward
9:     else
10:       $R^a \leftarrow R^a + R^b \kappa$  ▷ increase sooner reward
11:    end if
12:    responses[ $t$ ]  $\leftarrow$  ASKQUESTION( $D^a, D^b, R^a, R^b$ )
13:     $\kappa \leftarrow \kappa/2$ 
14:  end for
15:  indiff[ $t$ ]  $\leftarrow R^a$  ▷ Estimated indifference point
16: end for

```

E.2 The method of [Koffarnus and Bickel \(2014\)](#)

In the 5-trial approach of [Koffarnus and Bickel \(2014\)](#), a pre-defined series of 31 delays (D^b) are chosen (their Table 1), ranging from 1 hour to 25 years. All sooner rewards are immediate ($D^a = 0$), and the delayed reward is always double the immediate reward ($R^b = 2R^a$). This design space can be thought of as a series of 31 questions of different D^b values which takes a cross section at a discount fraction of 0.5. So by assuming a hyperbolic discount function, we can determine the half life (k^{-1}), which is the delay in which the present subjective value of a future reward is half of its objective value. Their heuristic algorithm is given in [Algorithm 5](#).

Appendix F. Nuisance parameters: error rate ε and comparison acuity α

The parameters α and ε can be described as nuisance parameters, in that they are an important part of the model, but they are not of key theoretical interest. In other words, these parameters can capture something useful about participant's response errors, but for the present purposes we are primarily interested in the discounting parameters; either $\log(k)$, (m, c) , or $\log(h)$, depending upon the model.

Our software allows for these parameters to be either set to fixed values (as is done for demonstrative purposes in the shown experiments $\alpha = 2$, and $\varepsilon = 0.01$) or to be included as parameters by introducing priors. In practise we recommend setting ε , as its accurate inference would require hundreds if not thousands of data-points⁷.

Placing a prior over α on the other hand can be eminently useful by reducing the model misspecification. This comes with a caveat though – if not fixed, α will be treated as a parameter for which

7. For example consider the case were we have one highly dubious result in an experiment of 30 questions. Any values of ε between say 0.02 to 0.1 remain very plausible in the absence of other information.

Algorithm 5 Method of [Koffarnus and Bickel \(2014\)](#)**Inputs:** Delayed reward value R^b **Outputs:** Vector of responses for each trial

```

1: delays  $\leftarrow$  As per Table 1 in Koffarnus and Bickel \(2014\)
2:  $D^a \leftarrow 0$ ,  $R^a \leftarrow R^b/2$ ,  $n \leftarrow 8$ 
3: for  $t = 1 : 5$  do
4:   if  $t = 1$  then ▷ Determine index  $i$  into delay list
5:      $i \leftarrow 16$  ▷ First delay corresponds to  $D^b$  of 3 weeks
6:   else
7:     if responses[ $t - 1$ ] = sooner then
8:        $i \leftarrow i - n$  ▷ decrease delay of later reward
9:     else
10:       $i \leftarrow i + n$  ▷ increase delay of later reward
11:    end if
12:     $n \leftarrow n/2$ 
13:  end if
14:   $D^b \leftarrow \text{delays}[i]$ 
15:  responses[ $t$ ]  $\leftarrow \text{ASKQUESTION}(D^a, D^b, R^a, R^b)$ 
16: end for

```

we wish to minimize uncertainty. The rationale for this is that treating it explicitly as a nuisance parameter would require the replacement of $p(y|\theta, d)$ with $\int p(y|\theta, d) d\alpha$. This would massively increase the computational cost of the estimating $\bar{U}$ as it would restrict one to only using a substantially less efficient NMC based estimator.⁸ Furthermore, we argue we should still look to minimize the entropy with respect to α even if we are not interested in the “true” value of α and in the absence of computational considerations. There are two core reasons for this: a) the optimal design equations are myopic ([González, Osborne, and Lawrence, 2016](#)) – they only consider one iteration in isolation – and b) they presume that the responses will be generated exactly by $p(y_t|y_{1:t-1}, d_{1:t})$, whereas in practice model-misspecification means this is not the case. Reducing uncertainty in α alleviates both these issues, reducing the level of model-misspecification and improving our estimate of the predictive marginal. Therefore although treating α as a nuisance variable when we expect to make no further questions might be preferable, information gathered about α will be helpful for future questions and so, even if it is a nuisance variable, it should be included in the entropy reduction to improve performance of the experiment as a whole.

Appendix G. Heuristics for coping with model misspecification and redundancy

The true underlying model for a participant is clearly far more complicated than can be modelled using a simple parametric model. Remaining tractable and ensuring interpretability of the parameters makes such simple models necessary, while practical limitations on the information that can be extracted from the experiment mean that more complicated models are not likely to gather significantly more information. Unfortunately, misspecification can be problematic when inferring parameters. For example, the presence of a very low likelihood observation in a misspecified model

8. Note that the resulting estimator is distinct to the $\hat{U}_{NMC}(d)$ defined in (17).

can result in a posterior that is far-fetched to a human observer, but still technically correct in the given model. This can be particularly damaging for adaptive design as the resulting severe overconfidence in current predictions – the model has no concept it might be wrong – can cause the process to get stuck repeatedly asking similar questions.

Our particular model also has a large degree of redundancy, as our discount factors depend only on the difference between the rewards and are thus effectively independent of the sum of the rewards $V(\mathcal{P}^a) + V(\mathcal{P}^b)$. This, along with the issue of model misspecification, means that it is necessary for us to sometimes incorporate heuristic components to the design optimisation to artificially force a good spread of questions, rather than just ensuring that all our questions are pertinent.

There are a number of subtleties to this heuristic and we will only provide the high level idea here, referring the reader to the `generate_designs.m` function in the code for full details. The desire of our heuristic is to ensure a good spread in values of $V(\mathcal{P}^a) + V(\mathcal{P}^b)$ as this value does not effect the outcome probability in our model. To do this, we will manipulate the set of candidate designs $\mathcal{D}$ at each step of the algorithm. Our heuristic works by first taking a point estimate for our parameters $\bar{\theta}$ (for which we use the mean) that can be used to cheaply estimate subjective values, which we then calculate for all candidate and previous questions. We further eliminate any points from both the candidate and previous questions where the response is effectively certain as these points will clearly not make for helpful questions. We then use a kernel density estimator on $V(\mathcal{P}^a) + V(\mathcal{P}^b)$ for the previous questions and use this as a measure of how often we have asked questions in that region of $V(\mathcal{P}^a) + V(\mathcal{P}^b)$ before. Minimizing this gives us a target value of $V(\mathcal{P}^a) + V(\mathcal{P}^b)$. To select the candidate designs, we then split the candidate designs into an number of evenly spaced bins (by default 100) of values for the probability of preferring $\mathcal{P}^b$, i.e. $\psi(V(\mathcal{P}^a), V(\mathcal{P}^b), \bar{\theta})$. Our set of candidate designs is then taken as the point in each bin which has $V(\mathcal{P}^a) + V(\mathcal{P}^b)$ closest to the target. There are also various fall-back strategies for when most bins are empty etc, for which we refer the reader to the code for full details.

References

- Woo-Young Ahn, Nathaniel Haines, and Lei Zhang. Revealing Neurocomputational Mechanisms of Reinforcement Learning and Decision-Making With the hBayesDM Package. *Computational Psychiatry*, 1:24–57, October 2017.
- Billy Amzal, Frédéric Y Bois, Eric Parent, and Christian P Robert. Bayesian-optimal design via interacting particle systems. *Journal of the American Statistical Association*, 101(474):773–785, 2006.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *JMLR*, 2002.
- Marianna Blackburn and Wael El-Deredy. The future is risky: Discounting of delayed and uncertain outcomes. *Behavioural Processes*, 94:9–18, March 2013.
- Jay Brown and David Cross. Discounting in judgments of delay and probability. *Journal of Behavioral Decision Making*, 13(2):145–159, 2000.
- Jared M Bruce, Amanda S Bruce, Delwyn Catley, Sharon Lynch, Kathleen Goggin, Derek Reed, Seung-Lark Lim, Lauren Strober, Morgan Glusman, Abigail R Ness, and David P Jarmolowicz. Being Kind to Your Future Self: Probability Discounting of Health Decision-Making. 50(2):297–309, December 2015.
- Olivier Cappé, Arnaud Guillin, Jean-Michel Marin, and Christian P Robert. Population Monte Carlo. *Journal of Computational and Graphical Statistics*, 13(4):907–929, 2004.
- George Casella and Christian P Robert. Rao-blackwellisation of sampling schemes. *Biometrika*, 83(1):81–94, 1996.
- Daniel R Cavagnaro, Gabriel J Aranovich, Samuel M McClure, Mark A Pitt, and Jay I Myung. On the functional form of temporal discounting: An optimized adaptive test. *Journal of Risk and Uncertainty*, 52(3):1–22, September 2016.
- Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 1995.
- Nicolas Chopin et al. Central limit theorem for sequential monte carlo methods and its application to bayesian inference. *The Annals of Statistics*, 32(6):2385–2411, 2004.
- Johannes Hugo Decker, Bernd Figner, and Joanna E Steinglass. On Weight and Waiting: Delay Discounting in Anorexia Nervosa Pretreatment and Posttreatment. *Biological Psychiatry*, pages 1–9, January 2015.
- William Brady DeHart. The effects of the framing of time on delay discounting. *Journal of the experimental analysis of behavior*, 103(1):10–21, January 2015.
- Arnaud Doucet, Simon Godsill, and Christophe Andrieu. On sequential monte carlo sampling methods for bayesian filtering. *Statistics and computing*, 10(3):197–208, 2000.
- J R Doyle. Survey of time preference, delay discounting models. *Judgment and Decision Making*, 8(2):116–135, March 2013.
- Andreas C Drichoutis and Jayson L Lusk. What can multiple price lists really tell us about risk preferences? *Journal of Risk and Uncertainty*, 53(2-3):1–18, March 2017.
- Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987.
- Ben Eppinger, Hauke R Heekeren, and Shu-Chen Li. Age Differences in the Neural Mechanisms of Intertemporal Choice Under Subjective Decision Conflict. *Cerebral Cortex*, pages 1–11,

September 2017.

- Sara J Estle, Leonard Green, Joel Myerson, and Daniel D Holt. Discounting of Monetary and Directly Consumable Rewards. 18(1):58–63, 2007.
- Gersende Fort, Emmanuel Gobet, and Eric Moulines. Mcmc design-based non-parametric regression for rare-event. application to nested risk computations. *Monte Carlo Methods Appl*, 2017.
- Christopher T Franck, Mikhail N Koffarnus, Leanna L House, and Warren K Bickel. Accurate characterization of delay discounting: a multiple model approach using approximate Bayesian model selection and a unified discounting measure. *Journal of the experimental analysis of behavior*, 103(1):218–233, January 2015.
- Charles C J Frye, Ann Galizio, Jonathan E Friedel, W Brady DeHart, and Amy L Odum. Measuring Delay Discounting in Humans Using an Adjusting Amount Task. *Journal of visualized experiments : JoVE*, (107), January 2016.
- Shawn P Gilroy, Christopher T Franck, and Donald A Hantula. The discounting model selector: Statistical software for delay discounting applications. *Journal of the experimental analysis of behavior*, 107(3):388–401, May 2017.
- Javier González, Michael Osborne, and Neil Lawrence. Glasses: Relieving the myopia of bayesian optimisation. In *Artificial Intelligence and Statistics*, pages 790–799, 2016.
- Noah D Goodman, Vikash K Mansinghka, Daniel Roy, Keith Bonawitz, and Joshua B Tenenbaum. Church: A language for generative models. In *In Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*. Citeseer, 2008.
- L Green, Joel Myerson, D Lichtman, S Rosen, and A Fry. Temporal discounting in choice between delayed rewards: the role of age and income. *Psychology and Aging*, 11(1):79–84, March 1996.
- Leonard Green, Joel Myerson, and Pawel Ostaszewski. Amount of reward has opposite effects on the discounting of delayed and probabilistic outcomes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(2):418–427, 1999.
- W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- L Jeff Hong and Sandeep Juneja. Estimating the mean of a non-linear function of conditional expectation. In *Winter Simulation Conference*, 2009.
- Matthe W Johnson and Warren K Bickel. Within-subject comparison of real and hypothetical money rewards in delay discounting. *Journal of the experimental analysis of behavior*, 77(2):129–146, March 2002.
- Daniel Kahneman and Amos Tversky. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2):263–292, March 1979.
- Frederick A A Kingdom and Nicolaas Prins. *Psychophysics. A Practical Introduction*. Academic Press, September 2009.
- K N Kirby and N N Maraković. Delay-discounting probabilistic rewards: Rates decrease as amounts increase. *Psychonomic Bulletin & Review*, 3(1):100–104, March 1996.
- Kris N Kirby. One-year temporal stability of delay-discount rates. *Psychonomic Bulletin & Review*, 16(3):457–462, June 2009.

- Mikhail N Koffarnus and Warren K Bickel. A 5-trial adjusting delay discounting task: accurate discount rates in less than one minute. *Experimental and Clinical Psychopharmacology*, 22(3):222–228, June 2014.
- L L Kontsevich and C W Tyler. Bayesian adaptive estimation of psychometric slope and threshold. *Vision research*, 39(16):2729–2737, August 1999.
- Steven R Lawyer, Sonja A Williams, Tereza Prihodova, Jason D Rollins, and Anita C Lester. Probability and delay discounting of hypothetical sexual outcomes. *Behavioural Processes*, 84(3):687–692, July 2010.
- Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *arXiv preprint arXiv:1603.06560*, 2016.
- Hause Lin, Blair Saunders, Cendri A Hutcherson, and Michael Inzlicht. Midfrontal theta and pupil dilation parametrically track subjective conflict (but also surprise) during intertemporal choice. *bioRxiv*, pages 1–56, August 2017.
- James E Mazur. An adjusting procedure for studying delayed reinforcement. In M L Commons, J A Nevin, and Howard Rachlin, editors, *Quantitative Analyses of Behavior*, pages 55–73. Erlbaum, Hillsdale, NJ, 1987.
- James E Mazur. Estimation of Indifference Points with an Adjusting-Delay Procedure. *Journal of the experimental analysis of behavior*, 49(1):37–47, January 1988.
- Todd L McKerchar, Leonard Green, Joel Myerson, T Stephen Pickford, Jade C Hill, and Steven C Stout. A comparison of four models of delay discounting in humans. *Behavioural Processes*, 81(2):256–259, June 2009.
- Joel Myerson. A discounting framework for choice with delayed and probabilistic rewards. *Psychological bulletin*, 130(5):769–792, September 2004.
- Jay I Myung, George Karabatsos, and Geoffrey J Iverson. A Bayesian approach to testing decision making axioms. *Journal of Mathematical Psychology*, 49(3):205–225, June 2005.
- Jay I Myung, Daniel R Cavagnaro, and Mark A Pitt. A tutorial on adaptive design optimization. *Journal of Mathematical Psychology*, 57(3-4):53–67, June 2013a.
- Jay I Myung, Daniel R Cavagnaro, and Mark A Pitt. A tutorial on adaptive design optimization. *Journal of mathematical psychology*, 57(3):53–67, 2013b.
- James Neufeld, András György, Dale Schuurmans, and Csaba Szepesvári. Adaptive Monte Carlo via bandit allocation. *arXiv preprint arXiv:1405.3318*, 2014.
- Håkan Nilsson, Jörg Rieskamp, and Eric-Jan Wagenmakers. Hierarchical Bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology*, 55(1):84–93, February 2011.
- Jan Peters, Stephan Franz Miedl, and Christian Büchel. Formal Comparison of Dual-Parameter Temporal Discounting Models in Controls and Pathological Gamblers. *PLoS ONE*, 7(11):e47225–12, November 2012.
- Shakoor Pooseh, Nadine Bernhardt, Alvaro Guevara, Quentin J M Huys, and Michael N Smolka. Value-based decision-making battery: A Bayesian adaptive approach to assess impulsive and risky behavior. *Behavior Research Methods*, 82(1):1–14, March 2017.
- Drazen Prelec and George Loewenstein. Decision Making Over Time and Under Uncertainty: A Common Approach. *Management Science*, 37(7):770–786, July 1991.

- N Prins. The psi-marginal adaptive method: How to give nuisance parameters the attention they deserve (no more, no less). *Journal of Vision*, 13(7):1–17, June 2013.
- E Pulcu, P D Trotter, E J Thomas, M McFarquhar, G Juhasz, B J Sahakian, J F W Deakin, R Zahn, I M Anderson, and R Elliott. Temporal discounting in major depressive disorder. *Psychological Medicine*, 44(09):1825–1834, November 2013.
- Howard Rachlin, Andres Raineri, and David Cross. Subjective probability and delay. *Journal of the experimental analysis of behavior*, 55(2):233–244, March 1991.
- Tom Rainforth. Automating inference, learning, and design using probabilistic programming. *Doctoral Dissertation*, 2017.
- Tom Rainforth. Nesting probabilistic programs. *arXiv preprint arXiv:1803.06328*, 2018.
- Tom Rainforth, Robert Cornish, Hongseok Yang, Andrew Warrington, and Frank Wood. On the opportunities and pitfalls of nesting Monte Carlo estimators. *arXiv preprint arXiv:1709.06181*, 2017.
- Erin B Rasmussen, Steven R Lawyer, and William Reilly. Percent body fat is related to delay and probability discounting for food in humans. *Behavioural Processes*, 83(1):23–30, January 2010.
- Daniel Read, Shane Frederick, Burcu Orsel, and Juwaria Rahman. Four Score and Seven Years from Now: The Date/Delay Effect in Temporal Discounting. *Management Science*, 51(9):1326–1335, September 2005.
- J B Richards, S H Mitchell, H de Wit, and L S Seiden. Determination of discount functions in rats with an adjusting-amount procedure. *Journal of the experimental analysis of behavior*, 67(3):353–366, May 1997.
- Christian P Robert. *Monte carlo methods*. Wiley Online Library, 2004.
- Elias Robles and Perla A Vargas. Functional parameters of delay discounting assessment tasks: Order of presentation. *Behavioural Processes*, 75(2):237–241, June 2007.
- Elias Robles and Perla A Vargas. Parameters of delay discounting assessment: Number of trials, effort, and sequential effects. *Behavioural Processes*, 78(2):285–290, June 2008.
- P Sebastiani and H P Wynn. Maximum entropy sampling and optimal Bayesian experimental design. 62(1):145–157, January 2000a.
- Paola Sebastiani and Henry P Wynn. Maximum entropy sampling and optimal bayesian experimental design. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(1):145–157, 2000b.
- Jordan Skrynka and Benjamin T Vincent. Subjective hunger, not blood glucose, influences domain general time preference. pages 1–12, July 2017.
- Henry P Stott. Cumulative prospect theory’s functional menagerie. *Journal of Risk and Uncertainty*, 32(2):101–130, March 2006.
- Taiki Takahashi. Psychophysics of the probability weighting function. *Physica A: Statistical Mechanics and its Applications*, 2011.
- Kazuhisa Takemura and Hajime Murakami. Probability Weighting Functions Derived from Hyperbolic Time Discounting: Psychophysical Models and Their Individual Level Testing. 7:778, 2016.

- William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- B Treutwein and H Strasburger. Fitting the psychometric function. *Perception & Psychophysics*, 61(1):87–106, January 1999.
- Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992.
- Vanderveldt, Ariana, Green, Leonard, and Myerson, Joel. Discounting of monetary rewards that are both delayed and probabilistic: delay and probability combine multiplicatively, not additively. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(1):148–162, January 2015.
- Benjamin T Vincent. Hierarchical Bayesian estimation and hypothesis testing for delay discounting tasks. *Behavior Research Methods*, 48(4):1608–1620, December 2016.
- A B Watson. QUEST: a Bayesian adaptive psychometric method. *Perception & psychophysics*, 33(2):113–120, February 1983.
- E Paul Wileyto, Janet Audrain-McGovern, Leonard H Epstein, and Caryn Lerman. Using logistic regression to estimate delay-discounting functions. *Behavior Research Methods, Instruments, & Computers*, 36(1):41–51, February 2004.
- Frank Wood, Jan Willem Meent, and Vikash Mansinghka. A new approach to probabilistic programming inference. In *Artificial Intelligence and Statistics*, pages 1024–1032, 2014.
- Richard Yi, Xochitl de la Piedad, and Warren K Bickel. The combined effects of delay and probability in discounting. *Behavioural Processes*, 73(2):149–155, September 2006.
- Michael E Young. Discounting: A practical guide to multilevel analysis of choice data. *Journal of the experimental analysis of behavior*, 9(1):476–20, February 2018.