

# Randomization Tests for Adaptively Collected Data

Yash Nair and Lucas Janson

Department of Statistics, Harvard University

## Abstract

Randomization tests (including permutation tests) are one of the most fundamental methods in statistics, enabling a range of inferential tasks such as testing for (conditional) independence of random variables, constructing confidence intervals in semiparametric location models, and constructing (by inverting a permutation test) model-free prediction intervals via conformal inference. Randomization tests are intuitive, easy to implement, and exactly valid for any sample size, but their use is generally confined to independent and/or exchangeable data. Yet in many applications including clinical trials, online education, online advertising, and protein design, data is routinely collected *adaptively*, meaning that the aspects of the data under the data collector’s control (e.g., treatment assignments) are assigned at each time step via a (possibly randomized) algorithm that depends on all the data observed so far; such assignment algorithms include (contextual) bandit and reinforcement learning algorithms as well as adaptive experimental designs. In this paper we present a general framework for randomization testing on adaptively collected data (despite its non-exchangeability), encompassing (and in some cases improving) the few existing results on randomization testing and conformal inference for adaptively collected data, as well as many other important settings. The key to our framework is the ability to compute likelihood-ratio-based weights involving known quantities based purely on the known adaptive assignment algorithm, as long as a certain proportionality condition is met. These weights can then be accounted for in our framework to conduct an exact randomization test, but in order for the test to be powerful, resamples need to be diverse yet have weights as close to equal as possible. Thus, we additionally present novel computationally tractable resampling algorithms for various popular adaptive assignment algorithms, data-generating environments, and types of inferential tasks. Finally, we demonstrate via a range of simulations our framework’s power (in the case of hypothesis testing) and narrow widths (in the case of confidence or prediction intervals produced by inverting randomization tests).

## 1 Introduction

### 1.1 Motivation

Randomization tests form an important methodological framework that enjoys a wide array of uses across statistics, ranging from testing equality of distributions to testing (conditional) independence between covariates and response in supervised learning settings. Beyond these classical uses of randomization tests, the framework also encompasses permutation testing and (inversely) conformal inference [Vovk et al., 2005]. We briefly note here that, throughout this paper, we use the phrase “randomization test” not to refer to a test applied on data obtained via the physical act of randomization taken by an experimenter, but rather the randomization used by the test itself to compute a p-value. Specifically, we view a randomization test as a randomized procedure, taking the observed data set  $D$  as input, that samples  $m$  copies  $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)}$  that are jointly exchangeable with  $D$  under the null, and then computes a (valid) p-value by taking a simple average of indicators comparing the resampled data to the observed data via a given test statistic. This interpretation of a randomization test has been referred to as a *quasi-randomization* test by Zhang and Zhao [2022].

While randomization tests can make very weak assumptions on the marginal distribution of observations, they generally make quite restrictive assumptions on the *joint* distribution of the data.

In particular, they usually require independent and/or exchangeable data [e.g., Candes et al., 2018, Edgington and Ongheena, 2007, Fisher et al., 1937, Lehmann et al., 2005, Pitman, 1937, Vovk et al., 2005], an assumption that is often violated in many real-world settings in which data are gathered in an adaptive fashion. One such example is drug discovery research, in which a scientist adaptively submits experiments serially, where the  $t^{\text{th}}$  experiment’s design is based upon the results of the first  $t-1$  [Popova et al., 2018]. This complex dependency will, in general, violate the exchangeability assumption when analyzing the data from all experiments at once.

Similar situations arise in experimental designs in the natural sciences more broadly in addition to mobile health, adaptive clinical trials, online education, and online advertising, to name a few.

Scenarios like those above naturally lead to a number of important inferential tasks that need to be performed on adaptively collected data, all of which we show can be performed via a randomization testing-based framework:

- Testing if two or more treatments induce the same distribution over outcomes
- Detecting non-stationarity in the data (i.e., testing whether the outcome depends not only on the treatment, but also on time)
- Using past data to predict the outcomes of future samples with quantifiable uncertainty
- Constructing confidence intervals for the difference in locations of outcome distributions corresponding to different treatments in semiparametric models

We now briefly describe various settings in which the above tasks and generalizations thereof are both difficult and important problems and motivate them with real-world examples.

**Problem Domain 1** (Comparing response distributions of different covariates). *An interesting and challenging problem in adaptive data collection settings is to test whether two or more different treatments amongst a total of  $K$  give rise to the same response distribution, i.e., for some pair  $i, j \in \{1, \dots, K\}$ , is  $(Y \mid X = i) \stackrel{d}{=} (Y \mid X = j)$ ? In the most extreme case, one might ask if there is even any dependency at all between  $X$  and  $Y$  (i.e., whether the treatment has any effect on the outcome). Beyond the mobile health examples delineated below in Problem Domain 2, for which it is important to test if two different treatments actually have the same effect (or whether treatments have any effect at all), such testing is also useful in recommender systems [Li et al., 2010, Schafer et al., 1999], to determine if different content displays actually produce different outcomes. Further complicating this setting is that different users may react differently (e.g., as measured by clickthrough rate) to different content. Thus, one might ask: Given the ‘context’ surrounding the user’s preferences, do they react the same way to different advertisements?*

**Problem Domain 2** (Testing non-stationarity). *An important problem that arises in domains in which actions are adaptively chosen and outcomes observed, is that of non-stationarity detection. That is, one may ask if the conditional distributions  $Y_t \mid X_t$  in data gathered via some adaptive decision-making procedure change with  $t$ . One such scenario in which non-stationarity testing is important is in mobile health applications. For example, the Just-in-Time Adaptive Intervention (JITAI) [Nahum-Shani et al., 2016] is a mobile health framework designed to provide appropriate support to patients with dynamic, time-varying states (e.g., mood, location, etc.). JITAIs have been applied to, for instance, suicide prevention [Coppersmith et al., 2021], smoking relapse prevention [Battalio et al., 2021], and addiction science [Carpenter et al., 2020]. In situations like these, it is vital to be able to know whether or not a patient’s response,  $Y_t$ , to the treatment,  $X_t$ , is varying over time,  $t$ , so as to facilitate administration of the most effective and appropriate care possible. Tests of non-stationarity provide a principled method of achieving this goal.*

**Problem Domain 3** (Predicting with high confidence). *In experimental design settings, providing prediction intervals for the results of future experiments can be instrumental in guiding researchers in decision-making. A motivating example comes from Alvarsson et al. [2021]. In their work, the authors introduce conformal inference [Burnaev and Vovk, 2014, Shafer and Vovk, 2008, Vovk, 2013, Vovk et al., 2005, 2009] to researchers in the field of drug discovery and go on to give an example of how the methodology—which also assumes i.i.d. or exchangeable data—can be applied to classify various types of ATP-Binding Cassette transporters*

at any user-specified uncertainty (i.e., significance) level. However, in light of the sequential and adaptive nature of many experimental designs in this field and in related fields like drug development [e.g., Godfrey et al., 2013, Mahajan and Gupta, 2010, Pallmann et al., 2018], there is need for provably valid test-time prediction intervals for these adaptively collected data.

**Problem Domain 4** (Relating parameters in semiparametric models). *A final difficult problem in scenarios like those of the previous Problem Domains is to estimate how the parameters of different response distributions belonging to a semiparametric model relate to one another. For example, suppose that  $Y \mid X = x \sim p_{\theta_x}$ , where  $\{p_{\theta} : \theta \in \mathcal{R}^d\}$  is a location family with parameter indexed by the treatment. How can we estimate  $\theta_x - \theta_{x'}$  for treatments  $x$  and  $x'$ ? This problem setting is ubiquitous, for example, in the multi-armed bandit literature, where it is common to assume that the reward distributions of all arms are distributed according to a location family like the Normal [Lai and Robbins, 1985]. In such settings, being able to estimate parametric relationships between different treatments in finite samples is useful for a variety of real-world problems. As an example, as described in Zhang et al. [2020], performing post hoc inference can be useful in many settings ranging from recommender systems [Mary et al., 2015] to anomaly detection [Ban and He, 2020] to finance and portfolio selection [Huo and Fu, 2017]. In such settings, after-the-fact inference may be helpful to researchers and practitioners who wish to understand the differences between different treatments and potentially utilize such differences to guide future decision-making.*

A number of tasks that we consider in this paper—and, indeed, some of those presented in the Problem Domains above—involve the construction of prediction or confidence intervals. However, as we show, constructing such intervals reduces to hypothesis testing by inverting the corresponding test. In particular, although not usually framed this way, conformal prediction is the inversion of a permutation test, a specific type of randomization test (see, e.g., Chernozhukov et al. [2018]). More generally, although it is unusual to invert a randomization test, we show that by applying certain transformations to the data, our randomization tests can be inverted to construct confidence intervals in semiparametric models (see, e.g., Rabideau and Wang [2021]). Section 3.3 delineates precisely how our tests can be inverted to produce conformal/confidence intervals. Due to this inverse relationship between prediction/estimation and testing, we generally (with the exception of Section 3.3) restrict our attention *solely* to hypothesis testing for clarity of presentation, and only state results in terms of the validity of such tests (and, in particular, the stochastic domination of the uniform distribution by their p-values).

Finally, all of the problem domains which motivate this work are centered around adaptively collected data. Indeed, as we show in this paper, we will be able to answer all of these inferential questions on adaptively collected data via a general framework for randomization testing *as long as* the adaptivity is known to the analyst. That is, we assume that the analyst has access to the (potentially non-deterministic) decision-making algorithm used to assign treatments in the data. We now introduce some notation to define this setup and formalize these notions.

## 1.2 Notation

The adaptive data collection settings we consider in this paper—of which popular settings like bandits, Markov Decision Processes (MDPs), reinforcement learning, and adaptive experimental designs are all special cases—comprise a data collection horizon of  $T$ , the number of timesteps of data collected by the practitioner (we will treat  $T$  as deterministic in this paper). At each timestep  $t$ , a context (sometimes called a *state*)  $C_t \in \mathcal{C}$  is observed, an action (also referred to as a *treatment*)  $X_t \in \mathcal{X}$  is selected, and a response (called an *outcome* or *reward* in some situations)  $Y_t \in \mathcal{Y}$  is subsequently observed.  $\mathcal{Z} := \mathcal{C} \times \mathcal{X} \times \mathcal{Y}$  is the (assumed time-invariant) sample space of the triple  $Z_t := (C_t, X_t, Y_t)$ —in this paper, we assume that  $\mathcal{Z}$  is discrete; see Remark 1 for an explanation of why as well as how to handle the case of continuous  $\mathcal{Z}$ . More formally, define the history at time  $t$ ,  $H_t$ , to be the sequence  $((C_1, X_1, Y_1), \dots, (C_t, X_t, Y_t))$ ;  $H_0$  is simply defined to be empty  $\emptyset$ . Then the action  $X_t$  is selected *conditionally* on  $H_{t-1}$  and  $C_t$ . These action-selection probabilities are encoded in the *known* decision-making (i.e., adaptive assignment) algorithm  $\mathcal{A}$ , where the probability of selecting action  $x \in \mathcal{X}$  at the  $t^{\text{th}}$  timestep given the  $t^{\text{th}}$  context and prior history of context-action-response triples up until time  $t$  is given by  $\mathbb{P}_{\mathcal{A}}(x|C_t, H_{t-1})$ . The joint density of the full

data sequence  $H_T = ((C_1, X_1, Y_1), \dots, (C_T, X_T, Y_T))$  is denoted by  $f$ . Unless otherwise noted, we always consider Markovian systems, and thus assume that:

$$Y_t \perp\!\!\!\perp H_{t-1} \mid (C_t, X_t).^1 \tag{1}$$

In this paper, we also sometimes consider domains without contexts, in which case  $C_t = \emptyset$  for all  $t \in [T]$ . For a given data sequence  $(C_1, X_1, Y_1), \dots, (C_T, X_T, Y_T)$ , set  $D := H_T$  to be their concatenation; take  $\mathcal{D} := \mathcal{Z}^T$  to be the sample space of the random variable  $D$ . Finally, as a note on general notation, we use  $[n]$  to refer to the set  $\{1, \dots, n\}$  and write  $[n : m]$  to denote  $\{n, \dots, m\}$  for integers  $n \leq m$ .

### 1.3 Contribution

In this paper, we provide exactly valid randomization-based inference for adaptively collected data. We make the following contributions, outlined below:

We derive a weighted Monte Carlo (MC) randomization testing framework. This test is able to resample data from essentially any arbitrary resampling distribution (as long as a certain proportionality hypothesis is satisfied) and, by weighting the samples according to certain likelihood ratios, produce a p-value. Along the way, we show that we are also able to utilize an unweighted Markov Chain Monte Carlo (MCMC) randomization procedure developed in Barber and Janson [2022+], Bates et al. [2021], Berrett et al. [2020], Besag and Clifford [1989], and applied in other settings, to perform inferential tasks on adaptively collected data. Both approaches offer exactly valid Type-I error control at essentially whatever computational cost the user desires (although, more computation generally results in higher power). Our novel weighted MC approach, however, is empirically no less powerful than the MCMC procedure and is, in fact, often more powerful in our simulations, especially when the number of resamples is small.

We use both the weighted MC and unweighted MCMC frameworks to test against a number of general null hypotheses on adaptively collected data:

1. Letting  $g$  be any function of  $X$ , we can test

$$Y_t \perp\!\!\!\perp X_t \mid (C_t, g(X_t)) \tag{2}$$

in particular settings by resampling copies of  $X_t$  for each  $t \in [T]$  from *any* distribution conditional on  $g(X_t)$  and weighting these resamples by certain likelihood-based ratios. In the special case in which  $g$  is a constant function,  $X_t$  can sometimes be sampled in such a way such that no weighting is required. In particular, the prior works of Bojinov and Shephard [2019], Ham and Qiu [2022], Pocock and Simon [1975], Simon [1979] are all able to handle this special setting by simply resampling the  $X_t$ 's by fixing the contexts and responses and re-running  $\mathcal{A}$ . In general, however, when  $g$  is non-constant, this unweighted procedure cannot be applied, yet our procedure always can be. One example of such a non-constant  $g$  might map treatments  $x$  in a multidimensional treatment space  $\mathcal{X}$  to their first dimension  $x_1$ . The hypothesis being tested in equation (2) would then be of conditional independence of the response with all other dimensions of the treatment given the context and first treatment dimension. Additionally, our framework applies in challenging settings like MDPs, whereas that of prior work does not. Beyond this added generality of our testing procedure, it offers various improvements (in terms of power and potential computational efficiency) over the existing work in certain settings.

2. We can test for non-stationarity in the conditional distributions  $Y_t \mid X_t, C_t$  in various complex adaptive data collection settings.
3. Under the semiparametric assumption that  $Y_t \mid X_t, C_t$  follows a location model with location parameter additive in  $X_t$ , our first test can be inverted to construct confidence regions for the location differences between actions as well. Similarly, our non-stationarity test can be inverted to construct conformal

---

<sup>1</sup>This assumption will only become relevant in Section 3. Additionally, while this Markovianity assumption is standard in the adaptive data collection literature [e.g., Bibaut et al., 2021, Hahn et al., 2011, Zhang et al., 2020], we discuss in Remark 4 how our procedure can be applied to test slightly different hypotheses when no Markovianity assumption is made.

prediction regions for  $Y_T$  under any adaptive assignment algorithm, considerably generalizing existing work which can only perform conformal inference under certain very specific adaptive assignment algorithms involving only a single round of adaptivity [Fannjiang et al., 2022, Tibshirani et al., 2020] or can only do so approximately [Barber et al., 2022, Chernozhukov et al., 2018, Gibbs and Candès, 2021].

Within our framework, it is more desirable (in terms of power) to have the weight of each resample as close as possible to that of the observed data; developing resampling procedures which achieve this for the various inferential tasks and environments discussed above is challenging. As such, we devise a number of novel and computationally tractable resampling procedures to be used for our tests.

Finally, to demonstrate the practicality of our randomization tests (which can be deployed at essentially any user-specified computational load) as well as the statistical efficiency of our resampling procedures, we perform a simulation study. By considering the aforementioned inferential tasks in a variety of data-generating environments, and by using data collected by a number of decision-making algorithms—including both deterministic and randomized ones—we demonstrate that our resampling procedures produce a test which is both statistically efficient (in terms of power and average confidence/prediction interval length) and computationally tractable, while, of course, also guaranteeing exact validity.

## 1.4 Related work

### 1.4.1 Randomization tests for adaptively collected data

As noted in Rosenberger et al. [2019], the works of Pocock and Simon [1975] and Simon [1979] consider randomization testing for adaptively collected data consisting of actions and outcomes in the Markovian setting of equation (1) under a certain restricted set of assignment algorithms. Extending their approach, Ham and Qiu [2022] are able to additionally handle context variables. Bojinov and Shephard [2019] also develop a randomization test that generalizes that of Pocock and Simon [1975], Simon [1979] primarily by relaxing the Markovianity assumption. The key to the approach put forth in all of the aforementioned papers is that they restrict the adaptivity of the data collection environment and consider null hypotheses such that simply fixing the contexts and responses and resampling the treatments via the known adaptive assignment algorithm produces exactly exchangeable copies of the data, thereby enabling them to conduct a standard *unweighted* randomization test. On the other hand, our main result, the *weighted* MC randomization test (Theorem 2.1) makes no such assumptions about the data-generating distribution nor the resampling distribution other than that they satisfy a certain proportionality hypothesis. Also, as we empirically demonstrate in Section 5, our approach can be powerful even on data generated by deterministic assignment algorithms in the Markovian setting, while that of Bojinov and Shephard [2019], Ham and Qiu [2022], Pocock and Simon [1975], Simon [1979] is provably powerless. For more details regarding the comparison of our work to prior work see Remarks 3 and 4.

Other existing works for performing inference in reinforcement learning and adaptive data collection settings are either asymptotic [e.g., Bugni et al., 2018, Deshpande et al., 2018, Hadad et al., 2021, Lai and Wei, 1982, Ma et al., 2015, Zhang et al., 2021, 2022, 2020] or, if not, are conservative in their finite-sample bounds [e.g., Howard et al., 2021, Kaufmann and Koolen, 2021]. As opposed to these works, all our hypothesis tests are able to maintain exact and non-conservative validity in finite samples.

### 1.4.2 MCMC, conditional and importance-weighted permutation/randomization tests

The connection between sampling exchangeable samples and MCMC was first developed by Besag and Clifford [1989] and has more recently been utilized by Barber and Janson [2022+], Bates et al. [2021], Berrett et al. [2020], for the purposes of efficiently running a conditional permutation test, generating knockoffs, and approximate co-sufficient sampling, respectively. The MCMC randomization test developed by Besag and Clifford [1989] which we outline in Section 2 takes the same approach: it generates exchangeable samples

using base draws from a Metropolis–Hastings Markov Chain. However, to the best of our knowledge, our work is the first that applies this MCMC approach to adaptively collected data.

Another related line of work is that of Harrison [2012] (which was brought to our attention following the preliminary online release of an earlier version of this paper). An importance special case of our main theorem (Theorem 2.1) discussed in Remark 9 regarding the validity of our weighted MC randomization test can be derived from their (conditional) importance-weighted randomization test [Harrison, 2012, Theorem 2]; this connection is made precise in Appendix A.1.

### 1.4.3 Conformal inference

Our weighted MC randomization test, when applied to non-stationarity testing and inverted, has close connections with conformal inference. The works of Tibshirani et al. [2020] and Fannjiang et al. [2022] are most related to our weighted MC randomization test when applied to non-stationarity testing and conformal inference. Tibshirani et al. [2020] first extend conformal inference from i.i.d. data to the setting of covariate shift, under the assumption that the likelihood ratio between test and train covariates is known by developing a more general *weighted* conformal prediction (WCP) procedure. Relatedly, Hu and Lei [2020] apply WCP to give an asymptotically powerful non-stationarity test in the same setting. The work of Fannjiang et al. [2022] extends Tibshirani et al. [2020] and shows how to handle dependence between train and test sets in the setting of feedback covariate shift. These prior methods handle only a single round of adaptivity (i.e., the first  $T - 1$  rounds of “train-time” data is fully i.i.d., but the “test-time” covariate distribution at timestep  $T$  may either be different [Hu and Lei, 2020, Tibshirani et al., 2020] or depend on the first  $T - 1$  datapoints via feedback covariate shift [Fannjiang et al., 2022]). In contrast, our test for non-stationarity (the inversion of which yields a conformal prediction interval) is able to handle *any number* of rounds of arbitrary analyst adaptivity. More similar to the permutation version of the weighted randomization test we propose in Section 2, Roach and Valdar [2018] propose a type of weighted Monte Carlo permutation test for conditionally i.i.d. resamples. Their test, however, does not assume knowledge of the resampling distribution and hence may be inexact in finite samples (though they show that it is asymptotically equal to a full weighted permutation test, similar to the one defined in Fannjiang et al. [2022]) and is also less general than our Theorem 2.1 in that it assumes conditionally i.i.d. resamples.

Nettasinghe et al. [2022] also develop an exact conformal inference procedure for non-exchangeable data, in the setting of HMMs, by using the notion of partial exchangeability [Diaconis and Freedman, 1980] and restricting to a subgroup of permutations which preserve certain transition structure. In the application of our (weighted) conformal prediction procedure to MDPs (described in Section 3.2), the same subgroup arises. Our method is, however, more general, as their work does not allow for adaptivity in treatment assignments and hence does not require the use of the likelihood ratio-based weights described in Section 1.3.

A final related line of work is in handling fully non-exchangeable data; that is, in contrast to Tibshirani et al. [2020] and Fannjiang et al. [2022], no exchangeability is assumed within any particular train set. In particular, Chernozhukov et al. [2018], Gibbs and Candès [2021], Xu and Xie [2021], Barber et al. [2022] all consider various versions of non-exchangeable data and extend conformal inference to these regimes. The methods presented in these works, however, are anti-conservative with bounds degrading as the degree of non-exchangeability grows. In contrast, the methods presented in this paper (using the adaptivity known to the analyst) provide exact validity in the adaptive (and hence potentially highly non-exchangeable) settings in which they apply.

## 1.5 Outline

In Section 2 we introduce the weighted MC randomization test. We also briefly describe how to perform an unweighted randomization test by using MCMC sampling. We then, in Section 3, show how our approach can be applied to solve a wide range of difficult inferential problems on adaptively collected data. Then in Section 4 we introduce a number of novel algorithms for generating resamples that can be used for a variety of inferential tasks and adaptive data collection environments. Finally, in Section 5, we perform a simulation

study of the methods introduced in Sections 2 and 3, using the resampling procedures from Section 4, that both empirically validates our methods by illustrating their validity and computational tractability, and demonstrates their power in a variety of challenging settings.

## 2 The weighted MC randomization test

In this section we introduce our main approach to randomization testing: a novel weighted MC randomization test. Additionally, at the end of this section, we briefly review an unweighted MCMC randomization testing framework introduced by Besag and Clifford [1989], but, to the best of our knowledge, never before applied to adaptively collected data (as we do in Section 3). The key to both approaches is the ability to define likelihood-ratio-based weights that satisfy a certain proportionality hypothesis, which, as we show in Section 3, is possible in a range of inferential tasks in various adaptive data collection settings. Empirically, our weighted MC approach never performs worse than the unweighted MCMC test, and indeed often dominates it, especially when the number of resamples is small.

Our weighted MC randomization test is a randomized procedure, taking the data set  $D$  as input, which samples  $m$  conditionally i.i.d. MC draws  $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)}$  given  $D$ . Define  $\mathfrak{D}$  to be the list  $(\tilde{D}^{(0)}, \tilde{D}^{(1)}, \dots, \tilde{D}^{(m)})$  where  $\tilde{D}^{(0)} := D$  and let  $q(\tilde{D}^{(i)}|D)$  denote the conditional probability of sampling  $\tilde{D}^{(i)}$  from this conditional distribution given that the dataset  $D$  was observed. Setting  $\hat{f}(\tilde{D}^{(i)}) := \prod_{t=1}^T \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)}|\tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)})$ , the procedure then calculates likelihood-ratio-based weights for each resample  $\tilde{D}^{(i)}$  in  $\mathfrak{D}$ :

$$w_{\mathfrak{D}}^q(\tilde{D}^{(i)}) = \frac{\hat{f}(\tilde{D}^{(i)}) \prod_{k \in [0:m] \setminus \{i\}} q(\tilde{D}^{(k)}|\tilde{D}^{(i)})}{\sum_{j=0}^m \hat{f}(\tilde{D}^{(j)}) \prod_{k \in [0:m] \setminus \{j\}} q(\tilde{D}^{(k)}|\tilde{D}^{(j)})}, \quad (3)$$

where  $\tilde{C}_t^{(i)}$ ,  $\tilde{X}_t^{(i)}$  and  $\tilde{H}_t^{(i)}$  are, respectively, the context at, action at, and history up until timestep  $t$  of the  $i^{\text{th}}$  resample  $\tilde{D}^{(i)}$ . Finally, the procedure outputs the p-value

$$p := \sum_{i=0}^m w_{\mathfrak{D}}^q(\tilde{D}^{(i)}) \mathbf{1}[S(\tilde{D}^{(i)}) \geq S(D)], \quad (4)$$

where  $S : \mathcal{D} \rightarrow \mathbb{R}$  is any test statistic. We summarize this procedure in the pseudocode of Algorithm 1

---

**Algorithm 1:** Weighted Monte Carlo randomization test

---

**Input:** Data sequence  $D$ , resampling distribution  $q$ , number of samples  $m$ , test statistic  $S$

- 1 Sample  $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)} \mid D \stackrel{\text{i.i.d.}}{\sim} q(\cdot|D)$
- 2  $\mathfrak{D} \leftarrow (\tilde{D}^{(0)}, \tilde{D}^{(1)}, \dots, \tilde{D}^{(m)})$  where  $\tilde{D}^{(0)} := D$
- 3 Compute  $w_{\mathfrak{D}}^q(\tilde{D}^{(i)})$  for each  $i \in [0 : m]$  via equation (3)

**Output:**  $p$ , calculated via equation (4)

---

The above weighting scheme may not always yield a valid p-value. We show, however, that as long as the hypothesis

$$\mathcal{H}_0^\infty : \hat{f}(\tilde{D}^{(i)}) = K f(\tilde{D}^{(i)}), \forall i \in [0 : m] \text{ for some constant } K \text{ not depending on } i, \quad (5)$$

holds, then  $p$  is a valid p-value. Importantly, note that the left hand side of the proportionality (5) is *always* computable (since it depends only on  $\mathcal{A}$ , which is known), whereas the right hand side is not in general since the density  $f$  is unknown. As we show in Section 3,  $\mathcal{H}_0^\infty$  holds for a number of null hypotheses in various adaptive data collection settings, thereby allowing for the valid use of the above to test against such nulls. We now present the main result of this paper, which is that, under  $\mathcal{H}_0^\infty$ ,  $p$  is a valid p-value.

**Theorem 2.1.** *The p-value defined in equation (4) stochastically dominates the uniform distribution under equation (5) for any resampling distribution  $q$  satisfying  $\mathcal{H}_0^\infty$ .*

While we relegate the proof of Theorem 2.1 to Appendix A, we note here that the proof is quite different than those of randomization tests in the standard unweighted case. For example, the proof of validity of the standard (unweighted) randomization test for exchangeable data and resamples—which is essentially an immediate consequence of the joint exchangeability of the original data and all resamples (see, e.g., Bates et al. [2021], Berrett et al. [2020], Besag and Clifford [1989], Candes et al. [2018], Edgington and Onghena [2007])—clearly will not apply in our case, since the data and resamples are not jointly exchangeable and the weighting scheme further makes it unclear how any such exchangeability could be used to prove validity of  $p$ . A second, more recent proof technique, given by Hemerik and Goeman [2018] in the context of the MC permutation test on exchangeable data, is to condition on the set  $\text{orb}(D)$  of all possible permutations of the observed data and use the fact that, conditionally on  $\text{orb}(D)$ , the original data is uniform over  $\text{orb}(D)$  and the set of resampled permutations is also uniform to show that the test is valid. In our setting, however, due to the non-exchangeability, this conditional distribution need not be uniform.

Rather, the proof of Theorem 2.1 is via the application of a simple yet new (to the best of our knowledge) application of Bayes’ theorem and proceeds by showing the conditional validity of  $p$  given the list  $\mathfrak{D}$ ; see Appendix A for details. Furthermore, as discussed in Remark 8, the assumption that the resamples  $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)}$  are drawn conditionally i.i.d. given  $D$  can be relaxed and a more general test (with generalized weights) can be used; see Appendix E.1 for details.

Finally, as discussed in Section 1.4.2, a special case of the above procedure discussed in Remark 9 can be derived from an importance-weighted MC randomization test presented in Harrison [2012] (although their proof technique is different and we were unable to extend it to prove our more general result). See Remark 9 and Appendix A.1 for details.

**Remark 1.** *The proof of Theorem 2.1 given in Appendix A applies only to discrete data distributions  $f$  and resampling distributions  $q$  (recall we assumed all data was discrete in Section 1.2). We focus only on such distributions as in any real-world application of our test, the computer on which our test is being run on has only finite precision, rendering both the original dataset as well as all resamples discrete. That is, any practitioner running our test on a computer is necessarily dealing with discrete data, not due to any restrictions of the test itself, but rather the finite precision limitation of the computer. This is not to say, however, that continuous distributions should never be considered in any situation. For example, in an asymptotic theory, rates of convergence for discrete distributions may degrade as the resolution of discretization becomes finer. In such a case, a continuous approximation of a computer-rendered discrete distribution may be considerably more useful than directly handling the discrete distribution itself. The key difference in our case, however, is that our theory is exact in finite samples for any discrete distribution. Hence, the guarantee of Theorem 2.1 does not worsen at finer discrete resolutions, but rather remains intact, thereby permitting our consideration of only discrete distributions.*

*We note, however, that our procedure should apply much more broadly to continuous distributions, but perhaps not completely generally. See [Huang and Janson, 2020, Figure 5] for one such example in which our theory does not hold because Bayes’ theorem is violated, by defining a dataset with  $T = 2$  as well as a resampling distribution with  $m = 1$  for which the marginal distribution of the resample  $\mathbb{P}(\tilde{D}^{(1)})$  is not equal to  $\tilde{q}(\tilde{D}^{(1)}|D)f(D)$  by taking  $X_1 \sim \text{Unif}(0, 1)$  and defining  $q$  to sample  $(\tilde{X}_1^{(1)}, \tilde{X}_1^{(2)}) \sim L_{X_1}$  where  $L_{X_1}$  is the segment of length 1 orthogonal to the  $x$ -axis, intersecting it at the point  $X_1$ , with an angle of  $(1 - X_1)^{10}\pi/2$  with the  $y$ -axis.*

**Remark 2** (Randomizing for exact Type-I error control). *Just as with a standard randomization test and, as a special case, conformal inference in which the procedure is called smoothing [e.g., Vovk et al., 2005], one can define a “lower”  $p$ -value*

$$p^- := \sum_{i=0}^m w_{\mathfrak{D}}^q(\tilde{D}^{(i)}) \mathbf{1}[S(\tilde{D}^{(i)}) > S(D)]$$

*and randomize to obtain a test for Theorem 2.1 which controls Type-I error at exactly the nominal level. Specifically, the test which fails to reject when  $p^- > \alpha$ , rejects with probability  $\frac{\alpha - p^-}{p - p^-}$  when  $p^- \leq \alpha < p$ , and otherwise always rejects, controls Type-I error at exactly the nominal level  $\alpha$ .*

While Theorem 2.1 gives one great freedom and generality in computing valid randomization  $p$ -values by

allowing for many choices of the conditional distribution  $q$ , two aspects in particular remain, at present, quite unclear: (a) in what situations and under what null hypotheses can we ensure that the proportionality hypothesis  $\mathcal{H}_0^\propto$  holds, as well as (b) *which* choices of  $q$  one should sample from to obtain a powerful test. We address the first question in the context of adaptively collected data in Section 3 and discuss the second in Sections 4 and 5.

**The unweighted MCMC randomization test** Above, we discussed a *weighted* MC randomization test which weights resamples based on certain likelihood ratios. Here, we briefly outline an *unweighted* MCMC randomization test developed by Besag and Clifford [1989]. While the test itself is not new (see, e.g., Barber and Janson [2022+], Bates et al. [2021], Berrett et al. [2019] for more recent utilizations and extensions of the test), our application of it to adaptively collected data, to the best of our knowledge, is. The randomized procedure for the unweighted MCMC randomization test repeatedly takes Metropolis–Hastings steps, starting at  $D$  by, at each step  $i$ , sampling a proposal  $\tilde{D}$  given  $\tilde{D}^{(i-1)}$  from  $q(\cdot|\tilde{D}^{(i-1)})$ , and additionally computing an acceptance ratio under the stationary distribution  $f$  to decide if the Metropolis–Hastings step accepts the proposal or remains at  $\tilde{D}^{(i-1)}$ . Just as with Algorithm 1, the acceptance ratio calculation actually only involves known quantities depending on  $\hat{f}$  and  $q$  and hence the validity of the MCMC test is again implied by the proportionality hypothesis  $\mathcal{H}_0^\propto$ . See Besag and Clifford [1989] for a delineation of the general test. We also note here that, in all of our simulations, our weighted MC test performs no worse than—and often dominates—the unweighted MCMC test. As such, exploring the utility of our weighted MC test in the settings of Barber and Janson [2022+], Bates et al. [2021], Berrett et al. [2020], all of which use the unweighted MCMC procedure, may be of interest for future work.

### 3 Randomization testing for adaptively collected data

In this section, we apply the randomization tests from Section 2 to solve a host of challenging inferential tasks in adaptive data collection settings. In Section 2, we saw that to perform a weighted MC randomization test we needed to be able to run Algorithm 1 which in turn required us to choose a  $q$  that satisfies the hypothesis  $\mathcal{H}_0^\propto$ . In this section, we show that, for a number of null hypotheses in various adaptive data collection settings,  $\mathcal{H}_0^\propto$  can be satisfied with many non-trivial choices of  $q$ , thus allowing for the application of the weighted MC randomization test. We note here that, since the focus of this section is ensuring that  $\mathcal{H}_0^\propto$  holds (under various null hypotheses and adaptive data collection settings), all results also immediately apply to the unweighted MCMC test.

We address two broad classes of tasks and describe each—as well as the adaptive data collection environments in which we consider them—in detail in the sections that follow:

1. Testing conditional independence between treatment and response conditional on context and some function of the treatment (Section 3.1). Applying certain transformations to the data and inverting such tests allows us to construct confidence intervals for response distribution parameters in certain semiparametric models (Section 3.3.1).
2. Non-stationarity testing (Section 3.2) and its application to conformal inference (Section 3.3.2)

In Sections 3.1 and 3.2 below, we discuss a number of adaptive data collection environment setups. Each of these setups assumes some combination/subset of equations (1) (Markovianity), (6), (7a), (7b):

$$Y_t \mid C_t, X_t \text{ does not depend on } t \qquad (Y\text{-Stationarity}) \qquad (6)$$

This first assumption asserts that the environment is *Y-stationary* so that the conditional response distribution does not vary with time; this distinction will be important in whether or not additional randomization can be employed for a more powerful conditional independence test.

$$C_t \perp\!\!\!\perp (X_1, \dots, X_{t-1}) \mid (C_1, Y_1, \dots, C_{t-1}, Y_{t-1}) \quad (\text{Weak non-reactivity}) \quad (7a)$$

$$C_t \mid H_{t-1} \text{ depends neither on } t \text{ nor } H_{t-1} \quad (C\text{-stationarity \& strong non-reactivity}) \quad (7b)$$

The assumptions of equations (7a) and (7b) are related and describe the environment through the behavior of contexts. In particular, the weak non-reactivity assumption (equation (7a)) says that the actions prior to time  $t$  do not affect  $C_t$ , conditionally on the previous contexts and responses; i.e., the future states of the environment are essentially (conditionally) *non-reactive* to prior actions. This weak non-reactivity assumption is also made in [Ham and Qiu, 2022, Assumption 1]. The  $C$ -stationarity & strong non-reactivity assumption (equation (7b)) is a stricter version of this, and stipulates that each context is generated by the environment independently from the past and also that the distribution from which it is generated does not change with  $t$ .

Before proceeding, we pause here to emphasize that all results that ensue in this section hold for *any* adaptive assignment algorithm  $\mathcal{A}$ , as long as it is known. We present all results by first describing the null hypothesis being tested, then explaining the various environments in which it can be tested, and finally delineating any constraints on the resampling distribution  $q$  in each of these environments so as to ensure that  $\mathcal{H}_0^\infty$  holds under the null.

### 3.1 Conditional independence testing

**Null hypothesis** Let  $g$  be any function of  $x \in \mathcal{X}$ , with unspecified codomain. We wish to perform a conditional independence test between treatment and response given both context and the  $g$ -evaluation of the treatment:

$$\mathcal{H}_0^{\perp\!\!\!\perp, g} : Y_t \perp\!\!\!\perp X_t \mid (C_t, g(X_t)), \forall t \in [T].$$

One example in which this hypothesis might be employed is in the problem of A/B testing in online advertising, where a company presents users with the same ad but with differing hues, saturations, and brightnesses (i.e., so the treatment space is 3-dimensional). One concrete hypothesis that may be tested in this situation is if saturation and brightness are enough to predict user response alone (i.e., is user response independent of hue given both saturation and brightness?). Setting  $g(X_t) = (X_{t, \text{saturation}}, X_{t, \text{brightness}})$  to map  $X_t$  to its saturation and brightness components,  $\mathcal{H}_0^{\perp\!\!\!\perp, g}$  is equivalent to this hypothesis. A second scenario is in testing if a particular subset of treatments induces the same response distribution. That is, letting  $\mathcal{X} = \{0, 1, 2\}$  be the treatment space, the null hypothesis that  $Y_t \perp\!\!\!\perp X_t \mid (C_t, X_t \in \{0, 1\})$  is equivalent to  $\mathcal{H}_0^{\perp\!\!\!\perp, g}$  with  $g(X_t) = \mathbb{I}(X_t = 2)$ . Finally, in the special case in which  $g$  is a constant function, the hypothesis being tested is that of simple conditional independence between treatment and response given context, and, for this particular choice of  $g$ , the unweighted test of prior work [Bojinov and Shephard, 2019, Ham and Qiu, 2022, Pocock and Simon, 1975, Simon, 1979] can be employed in Environments 1 and 3 (but not in Environments 2 or 4); as we describe below, however, our framework allows for a significant improvement in power over this prior work when employed in Environment 3.

**Environment 1** (Non-reactive environment). *The non-reactive environment is one in which the Markovianity (equation (1)) and weak non-reactivity (equation (7a)) assumptions hold. Examples of a non-reactive environment include (contextual) bandits—a setting that is common in the mobile health literature [e.g., Tewari and Murphy, 2017]—as well as many common adaptive experimental designs, such as adaptive crop yield experiments [e.g., Alesso et al., 2021].*

**Environment 2** (Markov Decision Process). *The next environment we consider is an MDP, which assumes only the Markovianity assumption (equation (1)) holds, and also stipulates that  $Y_t := C_{t+1}$ ; note that this implies that Assumption (7a) also holds. One example of MDP data arises in electronic health records [Li et al., 2022].*

Environments 3 and 4 below are  $Y$ -stationary counterparts of Environments 1 and 2, respectively.

**Environment 3** (Stationary non-reactive environment). *The most restricted of the environments we consider, the stationary non-reactive environment, assumes Markovianity (equation (1)),  $Y$ -stationarity (equation (6)), and  $C$ -stationarity & strong non-reactivity (equation (7b)). Examples of this environment arise in reinforcement learning as a stationary (contextual) bandit as well as in many common adaptive experimental designs—one example being in adaptive clinical trials [e.g., Giles et al., 2003].<sup>2</sup>*

**Environment 4** (Stationary MDP). *With the additional assumption of stationarity, the stationary MDP is a special type of MDP in which  $Y$ -stationarity (equation (6)) also holds, thereby assuming that the MDP’s transition distribution does not change across timesteps. An example of stationary MDP data is in patient admissions to the emergency room during surge demand [e.g., Lee and Lee, 2018].<sup>3</sup> A second example arises in robotics, where the robot’s dynamics and interactions with its environment can be modeled as a stationary MDP [e.g., Suárez-Hernández et al., 2019].*

**Constraints on  $q$**  The resampling procedure  $q$  may incorporate two sources of randomization: (a) randomizing the order of the data by some permutation in a certain environment-specific set  $\Pi$ , and (b) randomizing the actions conditional on their  $g$ -evaluations. More specifically,  $q$  must be such that the sequences of contexts, responses, and  $g$ -evaluations of the treatment in each resample are equal to some (random) permutation—albeit, the same one— $\pi \in \Pi$  of their respective sequences in the original data:

$$\left(\tilde{C}_t^{(i)}, g(\tilde{X}_t^{(i)}), \tilde{Y}_t^{(i)}\right) = \left(C_{\pi_i(t)}, g(X_{\pi_i(t)}), Y_{\pi_i(t)}\right), \quad \forall t \in [T], \text{ for some } \pi_i \in \Pi. \quad (8)$$

Note that equation (8) does not force  $\tilde{X}_t^{(i)} = X_{\pi_i(t)}$ , and thus allows the treatments to be further randomized over  $\mathcal{X}$ , so long as the restriction  $g(X_{\pi_i(t)}) = g(\tilde{X}_t^{(i)})$  is met. Finally, the restricted set  $\Pi$  of allowable permutations is environment-specific; we first focus on Environments 1–3, where the non-stationarity of Environments 1 and 2 prevent us from permuting the data at all, while the stationarity of Environment 3 allows for any permutation:

**Proposition 3.1.** *Let  $\Pi_i, i \in [3]$  denote the set of allowable permutations for Environments 1–3, above. Then if  $\Pi_1 = \Pi_2 = \{\text{id}\}$ , where  $\text{id}$  denotes the identity permutation, and  $\Pi_3 = \Pi_{[T]}$ , the set of all permutations on  $[T]$ , the proportionality hypothesis  $\mathcal{H}_\delta^\infty$  is satisfied under  $\mathcal{H}_0^{\perp, g}$ .*

*Proof.* In Environments 1 and 2, the restriction that  $\Pi_1 = \Pi_2 = \{\text{id}\}$  ensures that

$$\begin{aligned} \hat{f}(\tilde{D}^{(i)}) &= \prod_{t=1}^T \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \\ &= \frac{\prod_{t=1}^T \mathbb{P}_t(Y_t | C_t, g(X_t), \tilde{X}_t^{(i)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \mathbb{P}(C_t | C_1, Y_1, \dots, C_{t-1}, Y_{t-1})}{\prod_{t=1}^T \mathbb{P}_t(Y_t | C_t, g(X_t), X_t) \mathbb{P}(C_t | C_1, Y_1, \dots, C_{t-1}, Y_{t-1})} && \text{by } \mathcal{H}_0^{\perp, g} \text{ and that } g(\tilde{X}_t^{(i)}) = g(X_t) \\ &= \frac{\prod_{t=1}^T \mathbb{P}_t(\tilde{Y}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{X}_t^{(i)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \mathbb{P}(\tilde{C}_t^{(i)} | \tilde{C}_1^{(i)}, \tilde{Y}_1^{(i)}, \dots, \tilde{C}_{t-1}^{(i)}, \tilde{Y}_{t-1}^{(i)})}{\prod_{t=1}^T \mathbb{P}_t(Y_t | C_t, g(X_t), X_t) \mathbb{P}(C_t | C_1, Y_1, \dots, C_{t-1}, Y_{t-1})} && \text{as } \Pi_1 = \Pi_2 = \{\text{id}\} \text{ \& equation (8)} \\ &= \frac{1}{\prod_{t=1}^T \mathbb{P}_t(Y_t | C_t, g(X_t), X_t) \mathbb{P}(C_t | C_1, Y_1, \dots, C_{t-1}, Y_{t-1})} f(\tilde{D}^{(i)}), && \text{due to equation (7a)} \end{aligned}$$

where the subscript  $t$  on  $\mathbb{P}$  is used to emphasize that the conditional distribution  $Y_t | C_t, X_t$  may depend on  $t$ . Hence, the proportionality hypothesis is satisfied with  $K = \frac{1}{\prod_{t=1}^T \mathbb{P}_t(Y_t | C_t, g(X_t), X_t) \mathbb{P}(C_t | C_1, Y_1, \dots, C_{t-1}, Y_{t-1})}$ .

<sup>2</sup>To distinguish this setting from the adaptive crop yield experiment example, note that in an adaptive clinical trial, patients are selected i.i.d. from some population, thus obeying the  $C$ -stationarity & strong non-reactivity (equation (7b)). On the other hand, in an adaptive crop yield experiment, due to weather fluctuations (such as temperature and precipitation) over time, this assumption is violated. For the same reason, we expect adaptive clinical trial settings, but not adaptive crop yield experiments, to be  $Y$ -stationary.

<sup>3</sup>Again, this example differs from that of electronic health records in Environment 2 because here, we do not expect the transition dynamics to change over time, whereas in the previous example, administering certain medication to a patient may in fact change their internal state and thus alter the way in which they react to the same medication later on.

On the other hand, in Environment 3,  $\Pi_3 = \Pi_{[T]}$  and any permutation is allowed since

$$\begin{aligned}
\hat{f}(\tilde{D}^{(i)}) &= \prod_{t=1}^T \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \\
&= \frac{\prod_{t=1}^T \mathbb{P}(Y_{\pi_i(t)} | C_{\pi_i(t)}, X_{\pi_i(t)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \mathbb{P}(C_{\pi_i(t)})}{\prod_{t=1}^T \mathbb{P}(Y_t | C_t, X_t) \mathbb{P}(C_t)} && \text{due to equations (6) and (7b)} \\
&= \frac{\prod_{t=1}^T \mathbb{P}(Y_{\pi_i(t)} | C_{\pi_i(t)}, \tilde{X}_t^{(i)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \mathbb{P}(C_{\pi_i(t)})}{\prod_{t=1}^T \mathbb{P}(Y_t | C_t, X_t) \mathbb{P}(C_t)} && \text{by } \mathcal{H}_0^{\perp, g} \text{ and that } g(\tilde{X}_t^{(i)}) = g(X_{\pi_i(t)}) \\
&= \frac{\prod_{t=1}^T \mathbb{P}(\tilde{Y}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{X}_t^{(i)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \mathbb{P}(\tilde{C}_t^{(i)})}{\prod_{t=1}^T \mathbb{P}(Y_t | C_t, X_t) \mathbb{P}(C_t)} && \text{by equation (8)} \\
&= \frac{1}{\prod_{t=1}^T \mathbb{P}(Y_t | C_t, X_t) \mathbb{P}(C_t)} \cdot f(\tilde{D}^{(i)}), && \text{due to equation (7b)}
\end{aligned}$$

where we have dropped the subscript  $t$  on  $\mathbb{P}$  used above (and also do not need to subscript  $\mathbb{P}(C_{\pi_i(t)})$  due to the  $Y$ -stationarity assumption of equation (6) (resp. the  $C$ -stationarity assumption of equation (7b)). Hence, once again, the proportionality hypothesis holds, but now with  $K = \frac{1}{\prod_{t=1}^T \mathbb{P}(Y_t | C_t, X_t) \mathbb{P}(C_t)}$ .  $\square$

The precise definition of the allowable set of permutations  $\Pi_4$  (also introduced in Nettasinghe et al. [2022]) in Environment 4 is somewhat more complicated due to the serial dependence between consecutive timesteps (i.e.,  $C_{t+1}$  is generated conditionally on  $(C_t, X_t)$ ) that is not present in Environment 3. In particular, only permutations which preserve the property that the  $t^{\text{th}}$  response is equal to the  $(t+1)^{\text{th}}$  context are allowed. We formally state a Proposition here, but relegate its proof to Appendix B:

**Proposition 3.2.** *Setting*

$$\Pi_4 := \{\pi \in \Pi_{[T]} : Y_{\pi(t)} = C_{\pi(t+1)}, \forall t \in [T-1], \text{ and } C_{\pi(1)} = C_1\},$$

the proportionality hypothesis  $\mathcal{H}_0^\infty$  is satisfied under  $\mathcal{H}_0^{\perp, g}$  in Environment 4.

We briefly note here that unless the MDP's state space  $\mathcal{C}$  is relatively small, then the (data-dependent) permutation set  $\Pi_4$  will typically be quite small and virtually no randomization in permutations can be performed. There are, however, settings in which the state space  $\mathcal{C}$  is small, such as in the *restless multi-armed bandit* considered by Mate et al. [2020, 2021] in the context of patient adherence monitoring and well-being, in which our test can be used effectively.

**Remark 3** (Comparison with existing work). *In Environment 1 (and Environment 3, which is a special case), when  $g$  is a constant function, the unweighted randomization test of Bojinov and Shephard [2019], Ham and Qiu [2022], Pocock and Simon [1975], Simon [1979] applies. However, the testing procedure presented above has two advantages over those in prior work. First, by randomizing timesteps in addition to treatments, our procedure is more powerful than theirs in the stationary setting of Environment 3, especially for deterministic assignment algorithms for which their procedure is powerless; we empirically demonstrate this in Section 5.1.1. Second, their procedure requires that each resample rerun  $\mathcal{A}$  independently  $m$  times with the same sequence of  $C_t$  and  $Y_t$  as in  $D$ .<sup>4</sup> While this would seem to be the most natural and statistically powerful approach, it may not always be computationally feasible. In such a case, our method, by using a more easily computable resampling procedure, provides a computationally tractable workaround (albeit perhaps at the cost of degraded statistical efficiency per resample). As a concrete example, the adaptive assignment algorithm  $\mathcal{A}$  used to generate the original data could be based on Thompson sampling in a complex Bayesian model, thus rendering it too computationally burdensome to run for any but a very small number of MC samples. On the other hand, unnormalized densities—which are all that our procedure requires, due to the proportionality assumption  $\mathcal{H}_0^\infty$ —are generally easy to compute, thus rendering computation of the weights in equation (3) tractable under a less computationally intensive resampling procedure  $q$ .*

<sup>4</sup>Pocock and Simon [1975] and Simon [1979] only describe how to do so for certain adaptive assignment algorithms, and both, as well as Bojinov and Shephard [2019], only consider the non-contextual case.

**Remark 4** (Relaxing structural assumptions). *In Appendix C, we show the above testing procedure can be generalized to arbitrary adaptive data collection processes in which none of the assumptions of equations (6),(7a),(7b) nor the Markovianity assumption of equation (1) are made, and the adaptive data collection environment is assumed to be completely general. In such an environment, we show that one can, somewhat analogously, consider a sequence of functions  $g_1, \dots, g_T$  and test simultaneous (over  $t$ ) conditional independence between the  $t^{\text{th}}$  context (as well as the  $t^{\text{th}}$  response) and the prior sequence of actions given the prior sequences of contexts and responses as well as the sequence comprising the  $g_s$ -evaluation of the  $s^{\text{th}}$  action for  $s \in [t]$ . The special case of this hypothesis wherein all the  $g_t$  are constant allows for unweighted randomization testing in the completely general environment described in this Remark as shown by the prior work of Bojinov and Shephard [2019].*

## 3.2 Testing for non-stationarity

**Null hypothesis.** *For our non-stationarity test, the null hypothesis  $\mathcal{H}_0^S$  is that the response distribution is stationary (but unknown). That is, the conditional distribution of response given the context and treatment is the same across all timesteps:*

$$\mathcal{H}_0^S : Y_t \mid (C_t, X_t) \text{ does not depend on } t.$$

In this section, we consider two environments. The first is a special type of non-reactive environment (Environment 1):

**Environment 5** (*C-stationary strongly non-reactive environment*). *The C-stationary strongly non-reactive environment is an environment in which equation (7b) holds, in addition to the Markovianity assumption of equation (1). Once again, (contextual) bandits and various adaptive experimental designs are special cases. One concrete example is in adaptive experimental designs studying the effects of job search assistance programs on helping job seekers find work [Caria et al., 2020] over short periods of time since, in these brief time intervals, we expect the “context” surrounding each individual (e.g., background, credentials, etc.) to be roughly i.i.d. One additional example of a C-stationary strongly non-reactive environment for which our theory also holds is an episodic MDP, in which each episode is viewed as a single time step, the reward sequence a single response, and the action sequence a single (high-dimensional) action.<sup>5</sup>*

The second environment we consider is simply the MDP of Environment 2.

**Constraints on  $q$**  By a proof similar to that of Proposition 3.1 for Environment 3, it is straightforward to see that, as long as each draw from the resampling distribution  $q$  is (any) permutation of the original data  $D$ , the proportionality hypothesis  $\mathcal{H}_0^\infty$  holds under  $\mathcal{H}_0^S$  in Environment 5. That is, as long as we set  $\Pi = \Pi_{[T]}$  and do not allow for any other randomization, then  $\hat{f}(\tilde{D}^{(i)}) \propto f(\tilde{D}^{(i)})$ . Under the MDP setting of Environment 2, we may use the same randomization set of permutations  $\Pi_4$  stated in Proposition 3.2 and again, do not include any additional randomization. We summarize this below:

**Proposition 3.3.** *If  $\Pi = \{\pi \in \Pi_{[T]} : Y_{\pi(t)} = C_{\pi(t+1)}, \forall t \in [T-1], \text{ and } C_{\pi(1)} = C_1\}$  in Environment 2 and  $\Pi = \Pi_{[T]}$  in Environment 5, and the only randomization in  $q$ 's resampling consists solely of permutations drawn from  $\Pi$ , then the proportionality hypothesis  $\mathcal{H}_0^\infty$  is satisfied under  $\mathcal{H}_0^S$ .*

## 3.3 Inverting tests to construct confidence and prediction intervals

### 3.3.1 Confidence regions in semiparametric models

Recall that the hypothesis tests in Section 3.1 were all nonparametric and focused on testing (conditional) independence and distributional equality. In this section, we turn to the problem of exact parametric

<sup>5</sup>This episodic MDP setting also falls under the category of a (stationary) non-reactive environment, and hence our tests of conditional independence presented in Section 3.1 also apply.

inference in semiparametric models—focusing on semiparametric location models as a case study—and use a technique previously used in standard randomization testing to construct confidence intervals [e.g., Rabideau and Wang, 2021]. Before proceeding, we emphasize that parametric inference for adaptively collected data is a challenging problem and, as far as we are aware, all examples in the literature are either asymptotic [e.g., Deshpande et al., 2018, Hadad et al., 2021, Zhang et al., 2020], or conservative (see, e.g., Howard et al. [2021], Kaufmann and Koolen [2021]).

Now, consider the setting in which the response distribution is distributed according to a location family with locations determined by action. More precisely, letting  $h_0$  denote a (unknown) base density, we assume that at each time-step  $t \in [T]$ ,  $Y_t \mid (X_t = x) \sim h_0(y - \theta_x)$ , where  $x \mapsto \theta_x$  is some mapping of actions to location parameters. In such a setting it is quite natural to ask: how much better or worse is action  $x$  compared to  $x'$ , in terms of their location parameters? More formally, how can we test against the null hypothesis  $\mathcal{H}_0^{\text{Loc}, \delta, x, x'}$  that  $\theta_{x'} - \theta_x = \delta$  for actions  $x, x' \in \mathcal{X}$  and  $\delta \in \mathbb{R}$ ?

To test against  $\mathcal{H}_0^{\text{Loc}, \delta, x, x'}$ , we modify the dataset  $D$  and then perform a conditional independence test. Specifically, by modifying the dataset by replacing  $Y_t$  with  $Y_t + \delta \cdot \mathbf{1}[X_t = x]$ , we get that  $\mathcal{H}_0^{\text{Loc}, \delta, x, x'}$  implies that  $x$  and  $x'$  induce the same distribution over rewards in the modified data generating distribution. Hence

a test against  $\mathcal{H}^{\perp, g}$  with  $g(X_t) = \begin{cases} \{x, x'\} & \text{if } X_t \in \{x, x'\} \\ X_t & \text{otherwise} \end{cases}$  on this modified dataset serves as a test against

$\mathcal{H}_0^{\text{Loc}, \delta, x, x'}$ . These same ideas can also be applied to scale families; see Appendix D for details.

Most importantly, these tests can all be inverted to construct confidence regions for the parameter in question, namely  $\theta_{x'} - \theta_x$ . For example, consider constructing a confidence region for the difference in locations of treatments  $x$  and  $x'$  at nominal miscoverage rate  $\alpha$ . One can construct the acceptance region by simply running the test against  $\mathcal{H}_0^{\text{Loc}, \delta, x, x'}$  at level  $\alpha$  at each  $\delta$  in the parameter space  $\Delta$ : those  $\delta$  for which the test fails to reject make up the acceptance region. In cases where the parameter space  $\Delta$  is either continuous or too large to iterate over completely, approximate confidence regions can be constructed as follows: (a) grid the space into a small finite set of discrete points  $\Delta' \subseteq \Delta$ , (b) run the above procedure at each  $\delta \in \Delta'$  to obtain an accepted set of grid points  $A'$ , and (c) include all  $\delta \in \Delta$  within a certain user-specified distance of any of the accepted grid points of  $A'$  in the confidence region.

### 3.3.2 Prediction regions for $Y_T$

Similar to the previous section, the tests of Section 3.2 can be applied to construct prediction intervals for  $Y_T$  before it is observed. In particular, again, one can construct the acceptance region for  $Y_T$  by running the non-stationarity tests described in Section 3.2 on the almost-fully realized dataset  $((C_1, X_1, Y_1), \dots, (C_T, X_T, y))$ , with  $y$  ranging over  $\mathcal{Y}$ ; just as with confidence intervals, in the case of large  $\mathcal{Y}$ , the gridding/rounding procedure described at the end of the previous section can be used to construct approximate prediction intervals by simply selecting a small discrete set  $\mathcal{Y}' \subseteq \mathcal{Y}$ . We do note, however, that it is not uncommon for  $\mathcal{Y}$  to be small or even categorical in many (challenging) settings and thus, for such problems, we can grid  $\mathcal{Y}$  directly to obtain exactly valid prediction regions. We now make a few remarks about some broader connections of our procedure when used in this fashion to conformal inference, as well as challenges and speedups in constructing intervals.

**Remark 5** (Conditional conformal inference). *While conditional conformal inference is impossible in general, even in the case of i.i.d. data [Lei and Wasserman, 2014], Theorem 2.1 admits a conditional version which may be useful when the covariate space  $\mathcal{X}$  is finite and small. In particular, the validity of the test still holds when we replace  $f$  with the conditional density  $f(\cdot | X_T)$ , which conditions on the “test-time” covariate  $X_T = x_T$ ; thus, when inverting the hypothesis test and constructing a conformal prediction interval, one can guarantee valid coverage conditional on  $X_T$  (i.e.,  $\mathbb{P}(Y_T \in C_T^\alpha(X_T) | X_T) \geq 1 - \alpha$ , where  $C_m^\alpha(X_T)$  denotes conformal band at  $X_T$  with nominal miscoverage rate  $\alpha$ ). In practice, this is possible as long as each resample*

$\tilde{D}^{(i)}$  has  $\tilde{X}_T^{(i)} = X_T$  (in addition to  $\tilde{D}^{(i)}$  being an allowable permutation of  $D$ ), so that

$$\frac{f(\tilde{D}^{(i)}|X_T) \prod_{k \in [0:m] \setminus \{i\}} q(\tilde{D}^{(k)}|\tilde{D}^{(i)})}{\sum_{j=0}^m f(\tilde{D}^{(j)}|X_T) \prod_{k \in [0:m] \setminus \{i\}} q(\tilde{D}^{(k)}|\tilde{D}^{(j)})} = \frac{f(\tilde{D}^{(i)}) \prod_{k \in [0:m] \setminus \{i\}} q(\tilde{D}^{(k)}|\tilde{D}^{(i)})}{\sum_{j=0}^m f(\tilde{D}^{(j)}) \prod_{k \in [0:m] \setminus \{i\}} q(\tilde{D}^{(k)}|\tilde{D}^{(j)})}.$$

In turn,  $p$ 's validity once again ensues so long as the proportionality hypothesis  $\mathcal{H}_0^\infty$  holds.

**Remark 6** (Challenges with explicit construction of conformal bands). *Split conformal inference [Papadopoulos et al., 2007] is a variant of conformal inference which involves data splitting in order to construct a conformal prediction region which can be explicitly and efficiently computed. Unfortunately, such an explicit construction of a conformal band using data splitting does not carry over using our procedure to test non-stationarity, since the weights, through  $\hat{f}(\tilde{D}^{(i)})$ , may depend on the test-time response grid values  $y$ . Similarly, the discretization procedure of Chen et al. [2018] used to construct explicit bands in the i.i.d. setting by rounding the response to a small discrete set is not valid in our framework since, in general, probabilities under  $\mathbb{P}_A$  involving the rounded data are either unknown or not efficiently computable.*

**Remark 7** (Sharing samples). *To address the challenges of Remark 6, one can grid the space into a small finite set  $\mathcal{Y}'$  and run the test at each  $y \in \mathcal{Y}'$ , as discussed above in the case of continuous or large  $\mathcal{Y}$ . Naively, however, this involves drawing  $m$  resamples for each  $y \in \mathcal{Y}'$ . To reduce the total number of resamples needed, we can however share resamples between different values of  $y$ . That is, for  $y_1, y_2 \in \mathcal{Y}'$ , resamples drawn in association with  $y_1$  can be used to determine whether or not to accept  $y_2$  and vice versa, thereby more effectively using each resample drawn. See Appendix E.2 for details.*

## 4 Resampling procedures

In Section 3, we focused on an information-theoretic question: in what data regimes can we define a sampling procedure  $q$  for which our testing procedure is valid? Here, we turn to the second key question posed at the end of Section 2: which choices of  $q$  are *statistically* best? Thus, while the last section specified how to choose the support of  $q$  (i.e., which variables should be randomized over and which should be held fixed) depending on the null hypothesis being tested against and the adaptive data collection environment in which it is being tested, this section explores what *distributions* to choose over these supports. In the remainder of this section, we provide a partial answer to this challenging question by proposing a number of resampling procedures that are compatible with the constraints outlined in Section 3 and will be shown to be powerful and produce short confidence/prediction regions in Section 5.

Recall that we are focusing solely on conditionally i.i.d. resampling from procedures  $q$  and thus all resampling/proposal distributions considered in this section can be applied to both the weighted MC and unweighted MCMC tests (with  $q$  as the proposal distribution for the MCMC test). However, as mentioned in Section 2, our weighted MC procedure does not in general require conditionally i.i.d. resampling and hence the following remark describing how our test can be generalized in this case is in order:

**Remark 8** (Non-i.i.d. resampling). *When the resamples  $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)}$  are not generated in a conditionally i.i.d. manner, the test in Algorithm 1 can be slightly generalized. In particular, letting  $\Sigma$  be any subset of  $\Pi_{[0:m]}$ , the set of permutations on  $[0 : m]$ , and  $\tilde{q}((\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)})|D)$  denote the conditional probability of sampling  $(\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)})$  given that  $D$  was observed, the procedure can be generalized by redefining the weights to be*

$$w_{\Sigma}^{\tilde{q}}(\tilde{D}^{(i)}) = \frac{\hat{f}(\tilde{D}^{(i)}) \sum_{\pi \in \Sigma: \pi(0)=i} \tilde{q}((\tilde{D}^{(\pi(1))}, \dots, \tilde{D}^{(\pi(m))})|\tilde{D}^{(i)})}{\sum_{j=0}^m \hat{f}(\tilde{D}^{(j)}) \sum_{\pi' \in \Sigma: \pi'(0)=j} \tilde{q}((\tilde{D}^{(\pi'(1))}, \dots, \tilde{D}^{(\pi'(m))})|\tilde{D}^{(j)})}.$$

We relegate a description of this generalized procedure, as well as the proof of its validity, to Appendix E.1.

**Remark 9** (Computational issues with conditionally i.i.d. resampling). *We also note here that even with a conditionally i.i.d. resampling scheme  $q$ , the computation of the  $p$ -value  $p$  can sometimes take  $\Omega(m^2)$*

computations. This is because, even with a conditionally i.i.d. resampling procedure, all pairs of conditional densities  $q(\tilde{D}^{(i)}|\tilde{D}^{(j)})$  must be computed and incorporated into the  $p$ -value computation. On the other hand, if the conditional density draws samples  $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)}$  such that  $q(\cdot|\tilde{D}^{(i)})$  is the same for all  $i \in [0 : m]$ , then

$$\prod_{k \in [0:m] \setminus \{i\}} q(\tilde{D}^{(k)}|\tilde{D}^{(i)}) = \frac{\prod_{k=0}^m q(\tilde{D}^{(k)}|\tilde{D}^{(i)})}{q(\tilde{D}^{(i)}|\tilde{D}^{(i)})} \propto (q(\tilde{D}^{(i)}|\tilde{D}^{(i)}))^{-1} \quad (9)$$

and so the  $p$ -value can be computed much more quickly in only  $O(m)$  computations by replacing  $\prod_{k \in [0:m] \setminus \{i\}} q(\tilde{D}^{(k)}|\tilde{D}^{(i)})$  with  $(q(\tilde{D}^{(i)}|\tilde{D}^{(i)}))^{-1}$  in the weight equation (3); as a result, we focus on resampling procedures that have this property.

An equivalent way of characterizing this special case procedure of Algorithm 1 is that the conditional distribution  $q(\cdot|D)$  depends only on  $D$  through some statistic  $A(D)$  which is invariant under all resamples (i.e.,  $A(\tilde{D}^{(i)}) = A(D)$ , for all  $i \in [m]$ ) and so the resulting  $p$ -value from Algorithm 1 is equal, under  $\mathcal{H}_0^\infty$ , to

$$\sum_{i=0}^m \frac{f(\tilde{D}^{(i)}|A(D))/q(\tilde{D}^{(i)}|A(D))}{\sum_{j=0}^m f(\tilde{D}^{(j)}|A(D))/q(\tilde{D}^{(j)}|A(D))} \mathbf{1}[S(\tilde{D}^{(i)}) \geq S(D)],$$

which can be reduced, after conditioning on  $A(D)$ , to the importance-weighted  $p$ -value of [Harrison, 2012, Theorem 2]. For a more comprehensive explanation of the above connection as well as further details, see Appendix A.1.

The setup in this section is that we first consider resampling procedures used to test non-stationarity. We then go on to discuss resampling algorithms for testing the conditional independence null hypothesis  $\mathcal{H}_0^{\perp, g}$ , some of which use some of the non-stationarity testing resampling algorithms as subprocedures in their sampling process. All resampling procedures presented in this section are evaluated in our simulations in Section 5. Additionally, more detailed pseudocode outlines of the resampling procedures in this section can be found in Appendix F. Lastly, we make a brief note about the computation of probabilities under the resampling distribution  $q$ .

**Remark 10.** All resampling procedures we consider in this section—and in our simulations in Section 5—either sample uniformly or involve some sort of sequential sampling procedure. In either case, computation of conditional probabilities under  $q$  are tractable as they are constant in the former and, in the latter, can be calculated as the sample is generated by serially multiplying together the corresponding conditional probabilities of each sequential sample as it is generated. Of course, probabilities of the form  $q(D|\tilde{D}^{(i)})$  for  $i \in [m]$  still must be computed (as the original dataset  $D$  is never resampled from  $q$ ) from scratch, but this is simply done in the same manner as above: behaving as though  $D$  had indeed been sampled from  $q$  conditionally on  $\tilde{D}^{(i)}$  and sequentially multiplying conditional probabilities of each resampled timestep.

## 4.1 Non-stationarity testing in a $C$ -stationary strongly non-reactive environment

We first describe four types of resampling procedures that can be used for non-stationarity testing in a  $C$ -stationary strongly non-reactive environment (i.e., Environment 5). As discussed in Section 3.2, all such distributions must only randomize timesteps by permuting the data sequence. Apart from uniform permutations, the other three procedures discussed in this section randomly permute the data in a way intended to mimic  $\mathcal{A}$  while also ensuring diverse random samples. We thus call these three resampling procedures  $\text{imitation}_\pi$ ,  $\text{re-imitation}_\pi$ , and  $\text{cond-imitation}_\pi$ , the common word *imitation* referencing the mimicking of  $\mathcal{A}$  (the prefixes *re* and *cond* will be explained later on in this section). All three procedures randomly permute the data by serially sampling not-yet-sampled timesteps from the original data sequence.

**uniform $_\pi$  sampling** The  $\text{uniform}_\pi$  sampling procedure simply selects a permutation of the data uniformly at random. Although simple, intuitively it may result in a diverse set of resamples.

**imitation $_{\pi}$  sampling** This sampling procedure samples permutations by sequentially resampling timesteps from the original data, where the sampling distribution at time  $t$  acts as though the first  $t - 1$  already-resampled timesteps were drawn according to the true data generating-process and selects the  $t^{\text{th}}$  timestep proportionally to the probabilities dictated by  $\mathbb{P}_{\mathcal{A}}$ . In other words, at each timestep  $t$ , letting  $R$  denote the set of not-yet-sampled timesteps, the imitation $_{\pi}$  distribution draws a timestep  $t'$  from  $R$ , where the probability of drawing any given  $s \in R$  is proportional to  $\mathbb{P}_{\mathcal{A}}(X_s|C_s, \tilde{H}_{t-1})$ ; if  $\mathbb{P}_{\mathcal{A}}(X_s|C_s, \tilde{H}_{t-1}) = 0$  for all  $s \in R$ , the procedure is ended and an attempt at a new resample can be begun using the same process. Finally,  $(C_{t'}, X_{t'}, Y_{t'})$  is appended to  $\tilde{H}_{t-1}$  and  $t'$  is removed from  $R$ . Intuitively, the imitation $_{\pi}$  distribution behaves as  $\mathcal{A}$  would, feeding in the already-realized sequence of responses  $(Y_1, \dots, Y_T)$ , except that it may only sample amongst actions which correspond to not-yet-selected timesteps. See Algorithm 3 for pseudocode.

The re-imitation $_{\pi}$  and cond-imitation $_{\pi}$  distributions which we describe below are intended *only* for randomized decision-making algorithms  $\mathcal{A}$ . When applied to a deterministic algorithm, they are both the same as imitation $_{\pi}$ .

**re-imitation $_{\pi}$  sampling** This distribution is similar to the imitation $_{\pi}$  distribution, except that, to incorporate more diversity, it independently regenerates the exogenous randomness that the randomized decision-making algorithm  $\mathcal{A}$  makes as it mimics it to draw resamples; it thus *rerandomizes* the randomness of  $\mathcal{A}$ . More specifically, the re-imitation $_{\pi}$  distribution views  $\mathcal{A}$  as a sequence of decision rules  $\delta_t$  which take as input the tuple  $(C_t, H_{t-1}, U_1, \dots, U_t)$ , where  $U_t$  is the exogenous random variable generated by  $\mathcal{A}$  at time  $t$ , and output which action to take:

1. Sample a permutation of  $D$  by sequentially resampling timesteps from the data, where the sampling distribution at time  $t$  uses the  $t - 1$  already-resampled timesteps as well as the resampled exogenous randomness  $\tilde{U}_1, \dots, \tilde{U}_{t-1}$ , and then generates the random variable  $\tilde{U}_t$  from  $U_t$ 's distribution, but conditional on the remaining timesteps. Specifically,  $\tilde{U}_t$  is sampled from the conditional distribution of  $U_t$  given that

$$X_s = \delta_t(C_s, \tilde{H}_{t-1}, \tilde{U}_1, \dots, \tilde{U}_{t-1}, U_t) \text{ for at least one not-yet-selected timestep } s.$$

If, however, this conditional distribution is degenerate (i.e., because there does not exist  $\tilde{U}_t$  for which  $X_s = \delta_t(C_s, \tilde{H}_{t-1}, \tilde{U}_1, \dots, \tilde{U}_{t-1}, \tilde{U}_t)$  for any remaining timesteps  $s$ ), the process is terminated and sampling for a new resample is begun.

2. Select the  $t^{\text{th}}$  timestep uniformly over all those not-yet-sampled triples  $(C_s, X_s, Y_s)$  with

$$X_s = \delta_t(C_s, \tilde{H}_{t-1}, \tilde{U}_1, \dots, \tilde{U}_t). \tag{10}$$

The motivation behind the re-imitation $_{\pi}$  distribution, as opposed to imitation $_{\pi}$ , is that, by incorporating more of the randomness used in the decision-making algorithm one may be able to obtain more diverse samples while also better imitating the decision-making mechanisms of  $\mathcal{A}$ . See Algorithm 4 for pseudocode.

**cond-imitation $_{\pi}$  sampling** The cond-imitation $_{\pi}$  distribution is the same as re-imitation $_{\pi}$  except that instead of resampling the  $\tilde{U}_t$ 's, it *conditions* on them and thus uses the same sequence of exogenous randomness as re-imitation $_{\pi}$  does; this, of course, requires knowing the original sequence  $U_1, \dots, U_t$ . Intuitively, this conditioning that cond-imitation $_{\pi}$  sampling performs should bias the weights closer to  $(m + 1)^{-1}$ , resulting in more powerful resampling. See Algorithm 5 for pseudocode.

## 4.2 Non-stationarity testing in an MDP

We finally briefly discuss the resampling procedures for non-stationarity testing in an MDP (Environment 2). We consider essentially the same four types of resampling procedures used for non-stationarity testing in a  $C$ -stationary strongly non-reactive environment described in the last section. The only difference however, is that, as discussed in Section 3.2, only a subset of permutations are allowed in the MDP setting so as to ensure that the  $\tilde{Y}_t^{(i)} = \tilde{C}_{t+1}^{(i)}$  condition that holds for  $i = 0$  (i.e., in the observed data) also holds for each

resample  $\tilde{D}^{(i)}$ . Thus, while the four procedures which we consider here—also called  $\text{uniform}_\pi$ ,  $\text{imitation}_\pi$ ,  $\text{re-imitation}_\pi$ , and  $\text{cond-imitation}_\pi$ —sample timesteps serially without replacement according to  $\mathcal{A}$  as their analogs did in the previous section, they do so over only a restricted subset of not-already-sampled timesteps at each round.<sup>6</sup>

For this reason, the MDP data that we consider includes an additional action  $X_{T+1}$ ; hence the permutations which we consider can be viewed as permuting the  $T+1$  state-action pairs in this augmented dataset.<sup>7</sup> More specifically, suppose that at the end of round  $t-1$ , the state-action pair  $(C_{s'}, X_{s'})$  had just been selected and appended to  $\tilde{H}_{t-2}$ . Then, the set of allowable timesteps which can be sampled at round  $t$  are only those not-yet-sampled state-action pairs  $(C_s, X_s)$  for which the state-action-next state triple  $(C_{s'}, X_{s'}, C_s)$  is present in the original data sequence  $D$ . The precise way in which the timestep is sampled is then in exact accordance with the weighting, randomization, and conditioning that correspond to the  $\text{uniform}_\pi$ ,  $\text{imitation}_\pi$ ,  $\text{re-imitation}_\pi$ , and  $\text{cond-imitation}_\pi$  sampling procedures described in the previous section. See Algorithms 6, 7, 8, and 9 for pseudocode.

### 4.3 Conditional independence testing

We now describe four types of resampling procedures that can be used in a stationary non-reactive environment (Environment 3). Three of these resampling procedures randomize both timesteps and the action  $X_t$  conditional on  $g(X_t)$  allowed by the stationarity, as discussed in Section 3.1. The procedure which does not both randomize timesteps and actions simply randomizes the actions alone and is called  $\text{imitation}_X$ ; as such, it can also be applied to the non-reactive environment (Environment 1) as well as an MDP (Environment 2). Of the three procedures which randomize both timesteps and actions, two randomize the timesteps and actions in two separate stages and therefore use some of the resampling procedures discussed in the previous section (as well as one more) in the first stage; we call these resampling procedures  $\text{uniform}_\pi + \text{imitation}_X$  and  $\text{restricted-uniform}_\pi + \text{imitation}_X$  (the latter of these two uses a permutation scheme that involves  $g$  and was thus not discussed in the previous two sections). These resampling procedures can all be applied to a stationary MDP (Environment 4), by using the analogous MDP permutation distribution as described in Section 4.2. We note here that we do not consider the combinations  $\text{imitation}_\pi + \text{imitation}_X$ ,  $\text{re-imitation}_\pi + \text{imitation}_X$ , and  $\text{cond-imitation}_\pi + \text{imitation}_X$  because, while conditionally i.i.d., they violate the property discussed in Remark 9 that  $g(\cdot|\tilde{D}^{(i)})$  is the same for all  $i \in [0 : m]$ , and hence require  $\Omega(m^2)$  computations rendering them somewhat computationally burdensome. The fourth resampling scheme combines the two stages of permuting timesteps and randomizing  $X_t$  into a single procedure and is thus referred to as  $\text{combined}_{\pi, X}$ ; this procedure applies in the stationary non-reactive environment (Environment 3) and can be modified to work in a stationary MDP (Environment 4) by using the usual sequential permutation restriction described in Section 4.2. Finally, just as in the previous two sections, all of these distributions are based on the idea of trying to draw resamples in a way that mimics the behavior of  $\mathcal{A}$ .

**imitation<sub>X</sub> sampling** The  $\text{imitation}_X$  distribution, at each timestep  $t$ , conditions on the  $t-1$  already-resampled data points  $((C_1, \tilde{X}_1, Y_1), \dots, (C_{t-1}, \tilde{X}_{t-1}, Y_{t-1}))$  as well as  $C_t$  and, treating them as though they were true data points sampled by  $\mathcal{A}$ , samples  $\tilde{X}_t$  amongst all  $x \in \mathcal{X}$  with  $g(x) = g(X_t)$ , weighting proportionally to the action-selection probabilities of  $\mathbb{P}_\mathcal{A}$ —if all weights are 0, then the process is exited and an attempt at a new resample can be begun. The intuition behind this sampling procedure is that it attempts to mimic  $\mathcal{A}$  by essentially feeding in the already-realized context,  $g$ -evaluation, and response sequences to generate the sequence of actions (each sampled conditionally on the  $g$ -evaluation at the corresponding timestep), thereby mimicking the true data-generating distribution induced by  $\mathcal{A}$  (resulting in weights closer to  $(m+1)^{-1}$ ). We note here that the  $\text{imitation}_X$  resampling algorithm when applied in Environments 1 and 3 yields precisely the same (unweighted) test as the prior work of Bojinov and Shephard [2019], Ham and Qiu [2022], Pocock and Simon [1975], Simon [1979]. See Algorithm 11 for pseudocode.

<sup>6</sup>For  $\text{uniform}_\pi$  permutations, we sequentially sample indices uniformly at random from these restricted subsets.

<sup>7</sup>Note that, practically speaking, if the dataset  $D$  in consideration does *not* have this additional action  $X_{T+1}$ , then the analyst can add it with ease (because they know the adaptive assignment algorithm  $\mathcal{A}$ ) and can do so without affecting the null or alternative distribution from which the data was drawn (because the null/alternative distributions govern only the transition dynamics, and not action selection).

**uniform $_{\pi}$ +imitation $_X$  sampling** This resampling procedure proceeds in two stages: the first stage applies a uniform permutation using the uniform $_{\pi}$  sampling of Section 4.1 and the second randomizes the treatment conditional on its  $g$ -evaluation in the permuted data sequence using the imitation $_X$  resampling procedure. Similar to imitation $_X$ , we intuitively expect such a sampling procedure to have weights near  $(m+1)^{-1}$  while also incorporating significant diversity (resulting in more varied evaluated test statistics  $S(\tilde{D}^{(i)})$ ) due to the initial uniform permutation. See Algorithm 12 for pseudocode.

**restricted-uniform $_{\pi}$ +imitation $_X$  sampling** This procedure is identical to uniform $_{\pi}$ +imitation $_X$  sampling, except that it modifies the uniform sampling in step 1. In particular, for this resampling procedure, the randomly sampled permutation in the first stage is a *restricted uniform* permutation, which does *not* permute the sequence of  $g$ -evaluations. In other words, only permutations  $\pi$  for which  $g(X_t) = g(X_{\pi(t)})$  for all  $t \in [T]$  are allowed, and the sampling is uniform over this restricted set. The intuition behind this sampling scheme is similar to that of the uniform $_{\pi}$ +imitation $_X$  sampling procedure except that, by using restricted uniform permutations rather than fully uniform permutations, the sampled data appears more similar to the original data with the goal of making the weights closer to  $(m+1)^{-1}$ . See Algorithm 13 for pseudocode.

**combined $_{\pi,X}$  sampling** As opposed to the previous two sampling schemes which permute timesteps and randomize treatments conditional on their  $g$ -evaluations in two separate stages, this sampling approach *combines* both types of randomization into a single resampling stage. Specifically, at each timestep  $t$ , the combined $_{\pi,X}$  resampling procedure samples  $t'$  according to the imitation $_{\pi}$  distribution from Section 4.1 on  $g$ -evaluations over not-yet-selected timesteps, again, by conditioning on the already-resampled data, and then randomizes  $X_{t'}$  conditional on  $g(X_{t'})$ . In other words, at timestep  $t$ , the combined $_{\pi,X}$  resampling proceeds by:

1. Selecting a not-already-selected timestep  $t'$  conditionally on  $\tilde{H}_{t-1}$  via the imitation $_{\pi}$  distribution of Section 4.1 applied to  $g$ -evaluations of actions (instead of simply the actions themselves). If  $\mathbb{P}_{\mathcal{A}}(g(X_s)|\tilde{H}_{t-1}, C_s) = 0^8$  for all not-yet-selected timesteps  $s$ , then no such sample  $t'$  can be generated and so the sampling process is terminated and a new one may be begun.
2. Once the timestep  $(C_{t'}, X_{t'}, Y_{t'})$  is selected,  $\tilde{X}_{t'}$  is sampled via  $\mathbb{P}_{\mathcal{A}}(\cdot|\tilde{H}_{t-1}, C_{t'}, g(X_{t'}))$ .

Similar to restricted-uniform $_{\pi}$ +imitation $_X$  sampling, the intuition behind combined $_{\pi,X}$  sampling is that both randomization across timesteps and the treatments, conditional on their  $g$ -evaluations, are incorporated, except that here they are combined and both the timestep and first action component randomization mimic  $\mathcal{A}$ . See Algorithm 14 for pseudocode.

## 5 Empirical Results

In Section 4 we proposed a number of resampling procedures  $q$ . In this section, we implement these resampling procedures in a variety of adaptive data collection environments—which are all special cases of the more general environments discussed in Section 3 for which our theory holds—and both demonstrate their validity and evaluate their statistical efficiency.

In our experiments, we first study the performance of our tests as the horizon  $T$  increases. Thus, we plot both average Type-I error (to empirically validate the Type-I error control guarantee of Theorem 2.1) as well as power under an alternative distribution, for values of  $T$  ranging from 10 to 100. This power plot, however, does not paint the whole picture. In each simulation, a combination of two factors influences the power of our method: 1. the *intrinsic* difficulty of the environment and 2. the quality of our resampling algorithm. As an attempt to disentangle these two components, we additionally measure and plot the power of a standard randomization test on data collected via a baseline uniform i.i.d. adaptive assignment algorithm

<sup>8</sup>By  $\mathbb{P}_{\mathcal{A}}(g(X_s)|\tilde{H}_{t-1}, C_s)$  we mean the probability (induced by  $\mathbb{P}_{\mathcal{A}}$ ) that the  $g$ -evaluation at the  $t^{\text{th}}$  timestep is equal to  $g(X_s)$  given the history  $\tilde{H}_{t-1}$  and context  $C_s$ .

Contextual stationary non-reactive environment	$\epsilon$ -greedy, LinUCB
Contextless stationary non-reactive environment	$\epsilon$ -greedy, UCB
Contextless $C$ -stationary strongly non-reactive environment	$\epsilon$ -greedy, UCB
Contextual $C$ -stationary strongly non-reactive environment	$\epsilon$ -greedy, LinUCB
Markov decision process	$\epsilon$ -greedy $Q$ -learning, greedy $Q$ -learning

Table 1: Table of environments and adaptive assignment algorithms considered in each.

which selects actions uniformly at random from  $\mathcal{X}$  at each timestep  $t$  independently from  $H_{t-1}$ . This baseline measurement captures the intrinsic difficulty of the environment: a less powerful test on data gathered by this baseline i.i.d. treatment assignment indicates a more difficult environment. Now, while the power plots do illustrate the second factor regarding the quality of our resampling algorithms, this factor can be further decomposed into two more contributing components to paint a clearer picture: 1. the effective sample size of the resampling procedure (i.e., ideally having weights close to  $(m+1)^{-1}$ ) and 2. the diversity of the resamples (i.e., if all resamples are equal—and equal to the original data—then all weights will be  $(m+1)^{-1}$ , but our test will be powerless). We disentangle these two components—and thereby give a more complete explanation of the accompanying power plots—by measuring effective sample size,

$$m_{\text{eff}} := \frac{\left(\sum_{i=0}^m w_{\mathfrak{D}}^q(\tilde{D}^{(i)})\right)^2}{\sum_{i=0}^m w_{\mathfrak{D}}^q(\tilde{D}^{(i)})^2},$$

and plotting  $\frac{m_{\text{eff}}}{m}$ , the fractional effective sample size, under the alternative at the same increments as in the power plot. In Type-I error, power, and fractional effective sample size plots, we take  $m$ , the number of resamples drawn by our test, to be 100. We also present plots showing how the power of our procedure grows with  $m$  (taking values  $10^2$ ,  $10^3$ , and  $10^4$ ) for fixed  $T = 100$ .

We note here that all randomization tests performed in our simulations use smoothed  $p$ -values (see Remark 2) and thus provably control Type-I error (whose nominal rate is set at  $\alpha = 0.05$  in all our experiments) at *exactly* the nominal rate. In addition to testing, we also construct approximate confidence and conformal prediction intervals (at nominal miscoverage rate 0.05), using the procedures described in Section 3.3; as these inversion procedures involve non-smoothed  $p$ -values, we expect the coverage to be above 0.95. All results are averaged over 1000 independent trials and plotted with  $\pm 2$  standard error bars. Finally, we only show plots involving the weighted MC version of our test (and its inversion for interval construction) in this section, because for each resampling procedure we consider, our weighted MC test *is never outperformed by (in terms of power and length)*—and often in fact dominates—the unweighted MCMC test in all of our simulations. For analogous plots illustrating results for the unweighted MCMC test, see Appendix G.1.

**Computation** We also plot the computation times for each of our resampling algorithms in the various environments and adaptive data assignment algorithms considered here in Appendix G.2. In nearly all environments and assignment algorithms, we are able to run a powerful test on datasets of length  $T = 100$  within a matter of minutes (and often just seconds) in terms of serial computation time, and our test is of course embarrassingly parallelizable if desired. Generally, the computation times for each resampling procedure scale roughly linearly with  $T$ . All code for our simulations is publicly available at <https://github.com/Yashnair123/RTs-for-AdaptiveData>.

**Environments and adaptive assignment algorithms** We briefly describe the environments in and adaptive assignment algorithms for which we conduct our simulations. We consider a total of five different environments: two are examples of a stationary non-reactive environment (Environment 3)—one with contexts and one without contexts—two are instantiations of a  $C$ -stationary strongly non-reactive environment (Environment 5)—again, both contextual and contextless—and the last is an instance of an MDP (Environment 2); recall that Environments 3 and 5 are special cases of Environment 1. The adaptive assignment algorithms we consider in these environments are summarized in Table 1; note that, in each environment, we consider both a randomized and a deterministic adaptive assignment algorithm.

We now briefly describe each of these adaptive assignment algorithms.  $Q$ -learning [Watkins, 1989] maintains an estimate of the state-action value  $Q$  function and selects actions at each timestep based on the current

estimate; we consider one version in which this action selection is (deterministically) greedy and one in which it is  $\epsilon$ -greedy. The  $\epsilon$ -greedy algorithm in the contextless stationary non-reactive environment and contextless  $C$ -stationary strongly non-reactive environment both, at each timestep, determine the empirically best action and then select an action  $\epsilon$ -greedily. In the contextual stationary non-reactive environment, the  $\epsilon$ -greedy algorithm behaves similarly by maintaining a linear regressor  $L_x$  for each action  $x \in \mathcal{X}$ , and progressively updating  $L_x$  with the context-response (i.e., input-output) pair  $(C_t, Y_t)$  when  $x$  is selected at time  $t$  (upon seeing  $C_t$ ); the algorithm selects, at time  $t$  after seeing  $C_t$ , the action  $x$  with highest predicted response by the  $L_x$ 's. Finally, UCB [Auer et al., 2002] is a deterministic bandit algorithm which additively inflates each action's empirical value by a bound on its error with respect to the true value and selects the highest value action. The contextual analogue in which the response depends linearly on the context is LinUCB [Li et al., 2010].

We note here that essentially all of the adaptive assignment algorithms above are typically used in reinforcement learning as they all enjoy quite low regret. However, in comparison to much of the literature on asymptotic inference in reinforcement learning, which impose clipping constraints on these assignment algorithms that stipulate that action-selection probabilities cannot be too close to 0 or 1 [Deshpande et al., 2018, Hadad et al., 2021, Zhang et al., 2020], our procedure makes no such assumptions and even allows for deterministic adaptive assignment algorithms.

## 5.1 Conditional independence testing

In this section we apply the conditional independence testing framework and corresponding resampling algorithms from Section 4.3 to two environments: a stationary strongly non-reactive environment (Environment 3) with contexts and one without contexts. Our simulations in the former environment demonstrate the power gain our framework has over prior work [Bojinov and Shephard, 2019, Ham and Qiu, 2022, Pocock and Simon, 1975, Simon, 1979] in a stationary environment by incorporating the additional randomization over permutations described in Section 3.1. On the other hand, our simulations in the latter environment focus on the problem of testing if a subset of treatments induce the same response and thus demonstrates our test's power on an inferential test for which, to the best of our knowledge, no prior exact test exists.

### 5.1.1 Conditional independence testing with constant $g$

Our first series of simulations is in a contextual stationary strongly non-reactive environment, and involves testing against the null hypothesis  $\mathcal{H}^{\perp, g}$  with  $g(X_t) := \emptyset$  (i.e.,  $Y_t \perp\!\!\!\perp X_t \mid C_t$  for all  $t \in [T]$ ). While, as mentioned in Remark 3, this setting has been covered in prior work, our framework offers a more powerful test, as our simulation results illustrate. Additionally, we note that in this setting, the weighted MC and unweighted MCMC tests are identical. This is because all weights in the MC test are  $(m+1)^{-1}$  and all acceptance ratios in the MCMC test are 1 since our resampling procedures all involve the  $\text{imitation}_X$  distribution which, under a constant  $g$ , is the same as sampling, conditionally on the context and response sequences, from  $\mathcal{A}$ . For this reason, we do not plot the fractional effective sample sizes, as they too are all equal to 1.

The precise environment in which our simulations in this section are conducted has treatment space  $\mathcal{X} = \{0, 1\}$ . Letting  $I_k$  denote the  $k \times k$  identity matrix and  $\vec{1}_k$  the length- $k$  vector of all 1's, we consider null and alternative distributions involving 2-dimensional contexts  $C_t$  sampled i.i.d. from  $\mathcal{N}\left(\begin{pmatrix} 1 \\ -1 \end{pmatrix}, I_2\right)$  and whose conditional response distributions are, respectively, given by:

$$Y_t \mid (C_t, X_t) \sim \mathcal{N}(C_t^\top \vec{1}_2, 1) \text{ and } Y_t \mid (C_t, X_t) \sim \mathcal{N}(C_t^\top \vec{1}_2 + X_t, 1).$$

Finally, the test statistic used is simply the absolute value of the  $t$ -test statistic against  $\beta_2 = 0$  in a Normal linear model with design matrix comprising rows of the form  $(1, X_t, C_{t,1}, C_{t,2})$  and response vector  $(Y_t)_{t=1}^T$ .

Figure 1 demonstrates that the added diversity of incorporating uniform permutations—as opposed to sampling only from the  $\text{imitation}_X$  distribution—indeed results in an increase in power. Notably, when comparing

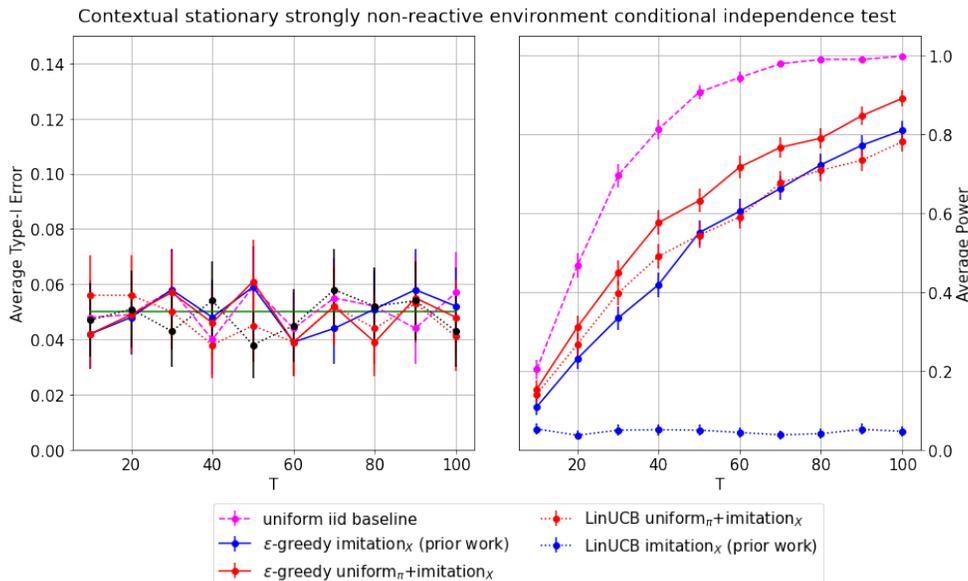


Figure 1: Type-I error (left) and power (right) of randomization tests at fixed  $m = 100$  and varying  $T$  in a contextual stationary strongly non-reactive environment on data gathered via  $\epsilon$ -greedy and LinUCB.

our resampling scheme,  $\text{uniform}_\pi + \text{imitation}_X$ , against the  $\text{imitation}_X$  resampling of prior work [Bojinov and Shephard, 2019, Ham and Qiu, 2022, Pocock and Simon, 1975, Simon, 1979] we observe that our approach is more powerful than theirs on data collected via an  $\epsilon$ -greedy treatment assignment. Perhaps even more strikingly, our approach when applied to LinUCB, a deterministic adaptive assignment algorithm, has high and increasing (with  $T$ ) power. This is in contrast to the usual  $\text{imitation}_X$  resampling of the prior work, which is powerless. We relegate the power curves for increasing  $m$  but finite  $T$  to Appendix G.3 as they are all approximately constant.

### 5.1.2 Conditional independence testing with non-constant $g$

We now discuss our simulations in the contextless stationary non-reactive setting of Environment 3, given as an example in the beginning of Section 3.1. In particular, we consider a contextless instantiation of a stationary non-reactive environment with action space  $\mathcal{X} = \{0, 1, 2\}$  and are testing against the null hypothesis  $\mathcal{H}_0^{\perp\perp, g}$  with  $g(X_t) := \mathbb{I}(X_t = 2)$  (i.e., that  $Y_t \perp\!\!\!\perp X_t \mid X_t \in \{0, 1\}$  for all  $t \in [T]$ ).

In this environment, we specify the null and alternative distributions as

$$Y_t \mid X_t \sim \begin{cases} \mathcal{N}(0, 1) & \text{if } X_t \in \{0, 1\} \\ \mathcal{N}(2, 1) & \text{if } X_t = 2 \end{cases} \quad \text{and} \quad Y_t \mid X_t \sim \begin{cases} \mathcal{N}(0, 1) & \text{if } X_t = 0 \\ \mathcal{N}(3, 1) & \text{if } X_t = 1 \\ \mathcal{N}(2, 1) & \text{if } X_t = 2 \end{cases},$$

respectively.

In all our simulations, the test statistic  $S(D)$  we use is, similar to as in Section 5.1.1, simply the absolute value of the  $t$ -test statistic for the test against  $\beta_2 = 0$  in a Normal linear model whose design matrix has  $(1, \mathbb{I}(X_t = 0), \mathbb{I}(X_t = 2))$  as its  $t^{\text{th}}$  row and  $(Y_t)_{t=1}^T$  as the response vector. Figure 2 summarizes the results in this domain.

First, the Type-I error rates in Figure 2 are controlled at exactly the nominal level, validating the theoretical guarantee of Theorem 2.1. With regards to power, we notice that restricted- $\text{uniform}_\pi + \text{imitation}_X$  sampling performs better on data gathered via  $\epsilon$ -greedy while combined- $\text{uniform}_\pi, X$  sampling performs better under UCB. The fractional effective sample size plots in the same figure offer an explanation for why this is the case: the fractional effective sample sizes under  $\epsilon$ -greedy are larger with restricted- $\text{uniform}_\pi + \text{imitation}_X$  sampling than

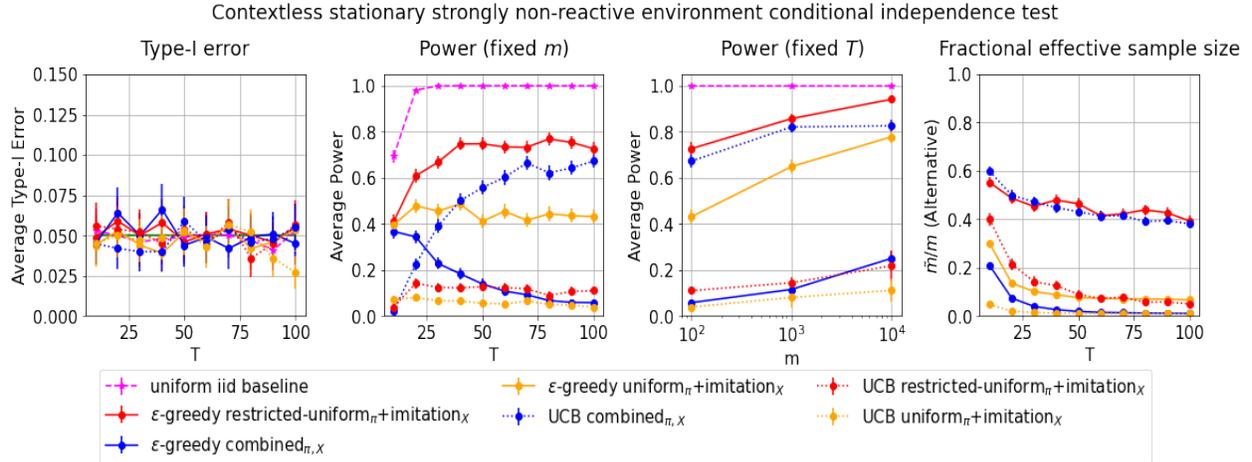


Figure 2: Type-I error rate (leftmost) and power (second from left) of the MC randomization test at fixed  $m = 100$  and varying  $T$  as well as power at fixed  $T = 100$  and varying  $m$  (third from left) and fractional effective sample size plots at fixed  $m = 100$  and varying  $T$  (rightmost) in a contextless stationary strongly non-reactive environment on data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

with combined $_{\pi,X}$  and, conversely, under UCB they are larger with combined $_{\pi,X}$  than restricted-uniform $_{\pi}$  + imitation $_X$ . We note here that this testing problem is much harder on data gathered via  $\epsilon$ -greedy and UCB than on the uniform i.i.d. baseline (thus explaining the power gap between the latter and all resampling procedures applied on data gathered by the former two, especially for large  $T$ ). This is because  $\epsilon$ -greedy and UCB are low-regret adaptive assignment algorithms and thus, under the alternative, will select action 1 very frequently, and all other actions much less frequently. Hence the problem of detecting if the response distributions induced by  $X_t = 0$  and  $X_t = 1$  are the same becomes much more challenging, as action 0 is sampled rarely under these two low-regret algorithms. On the other hand, it is sampled more frequently and at the same rate as action 1 under the uniform i.i.d. baseline. Despite this challenge, however, our method is still able to attain quite high power under both  $\epsilon$ -greedy and UCB adaptive assignment algorithms.

## 5.2 Non-stationarity testing

In this section, we empirically evaluate our test of non-stationarity on three different environments: the first two are examples of the  $C$ -stationary strongly non-reactive setting of Environment 5—one with contexts and one without—and the third is an MDP (Environment 2).

### 5.2.1 Testing non-stationarity in a $C$ -stationary strongly non-reactive environment

Our simulations in this section are performed on a contextless  $C$ -stationary strongly non-reactive environment with action space  $\mathcal{X} = \{-1, 1\}$  and Gaussian rewards: the reward distribution for the first  $T - 1$  steps is given by  $Y_t | X_t \sim \mathcal{N}(X_t, 1)$ . Under the null hypothesis  $\mathcal{H}_0^S$ , the reward distribution at the  $T^{\text{th}}$  timestep is unchanged and hence  $Y_T | X_T \sim \mathcal{N}(X_T, 1)$ . We analyze power under an alternative distribution that samples  $Y_T | X_T \sim \mathcal{N}(4X_T, 1)$ . Finally, as we are testing for non-stationarity, our test statistic  $S$  is simply a non-conformity score, and we choose it to be the absolute residual:  $S(D) = |Y_T - \hat{\mu}_D(X_T)|$ , where  $\hat{\mu}_D$  is the fitted ordinary least squares (OLS) model to  $D$  with an intercept term.

Figure 3 shows the results for our simulations in this section. Indeed the Type-I error plots once again demonstrate the validity our randomization tests. In terms of power, Figure 3 shows that cond-imitation $_{\pi}$  sampling performs best under an  $\epsilon$ -greedy treatment assignment and imitation $_{\pi}$  performs best under UCB and both attain power close to that of the uniform i.i.d. baseline for nearly all values of  $T$ . Figure 3, however, also shows that, with large enough  $m$ , re-imitation $_{\pi}$  sampling eventually performs better than cond-imitation $_{\pi}$  at  $T = 100$ . Combining this information with the fractional effective sample size plots of

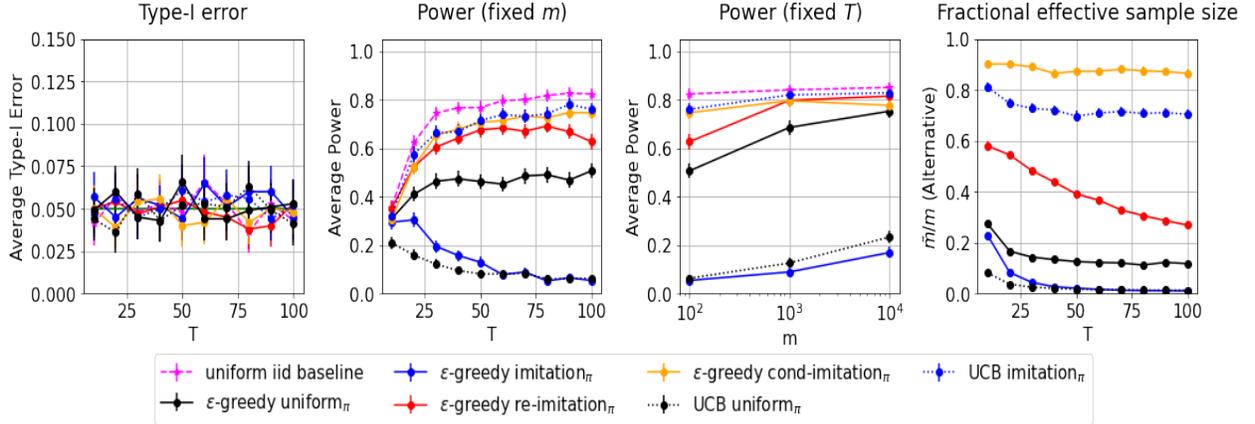
Contextless  $C$ -stationary strongly non-reactive environment, non-stationarity test

Figure 3: Type-I error rate (leftmost) and power (second from right) of the MC randomization test at fixed  $m = 100$  and varying  $T$  as well as power for fixed  $T = 100$  and varying  $m$  (third from right) and fractional effective sample size at fixed  $m = 100$  and varying  $T$  (rightmost) in a contextless  $C$ -stationary strongly non-reactive environment with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

Figure 3, we thus see that while  $\text{cond-imitation}_\pi$  has greater effective sample size than  $\text{re-imitation}_\pi$ , it has less diverse samples, leading to a more powerful procedure when  $m$  is small (since, in this regime, the diversity plays a greater role), but a (slightly) less powerful one when  $m$  is large.

## 5.2.2 Testing non-stationarity in a contextual $C$ -stationary strongly non-reactive environment

We now describe our simulations in a contextual  $C$ -stationary strongly non-reactive environment. The specific environment we consider in this section involves a action space  $\mathcal{X} = \{-1, 1\}$  and 100-dimensional contexts sampled i.i.d. from  $\mathcal{N}(\vec{1}_{100}, I_{100})$ . The conditional response distribution during the first  $T - 1$  timesteps is a sparse linear combination of the context vector and is given by

$$Y_t | (C_t, X_t) \sim \mathcal{N}(-5X_t + \sum_{j=1}^{10} C_{t,j}, 1).$$

The null conditional response distribution at time  $T$  is of course the same as the above whereas the alternative distribution we evaluate power under swaps the effects of the two treatments and is given by

$$Y_T | (C_T, X_T) \sim \mathcal{N}(5X_T + \sum_{j=1}^{10} C_{T,j}, 1).$$

We regularize the regressor  $L_x$  used in the  $\epsilon$ -greedy adaptive assignment by using Lasso regression with with penalty parameter 10 as opposed to OLS. Finally, the test statistic used is the same non-conformity score as in the previous section, except that we again use Lasso—this time with penalty parameter chosen through 5-fold cross-validation—instead of OLS, as  $\hat{\mu}_D$ .

Figure 4 shows the Type-I error, power, and effective sample size curves for our simulations in this environment. The Type-I error is again controlled at the nominal level. In terms of power, similar conclusions to those drawn in the contextless  $C$ -stationary strongly non-reactive environment of the previous section apply here, too. In particular, while  $\text{cond-imitation}_\pi$  sampling performs best under an  $\epsilon$ -greedy adaptive assignment at  $m = 100$  (and again exhibits, for large  $T$ , power quite close to—and, in fact greater than, for small  $T$ —that of the uniform i.i.d. baseline) it has essentially the same power as  $\text{re-imitation}_\pi$  at both  $m = 10^3$  and  $m = 10^4$ . Using the fractional effective sample size plots, we may therefore infer, once again, that while  $\text{cond-imitation}_\pi$  sampling has a higher effective sample size than  $\text{re-imitation}_\pi$  does under  $\epsilon$ -greedy, its samples exhibit less diversity than those of  $\text{re-imitation}_\pi$ . Under LinUCB, the contextual analog

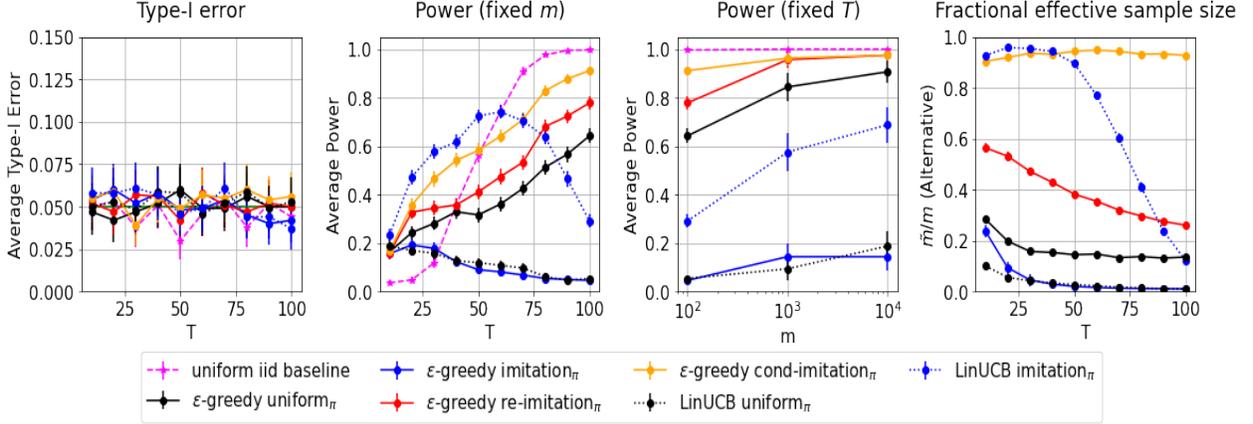
Contextual  $C$ -stationary strongly non-reactive environment, non-stationarity test

Figure 4: Type-I error rate (leftmost) and power (second from right) of the MC randomization test at fixed  $m = 100$  and varying  $T$  as well as power for fixed  $T = 100$  and varying  $m$  (third from right) and fractional effective sample size at fixed  $m = 100$  and varying  $T$  (rightmost) in a contextless  $C$ -stationary strongly non-reactive environment with data gathered via  $\epsilon$ -greedy, LinUCB, and the uniform i.i.d. baseline.

of the deterministic UCB assignment, however, we see that while imitation $_{\pi}$  has quite high power for small to moderate values of  $T$ , the power drops for larger  $T$ . The fractional effective sample size curve explains that this is caused by a corresponding drop in the effective sample size as  $T$  grows. We hypothesize that this is due to the extremely uneven action selection exhibited by LinUCB due to its very low regret (even as compared to  $\epsilon$ -greedy); we discuss why this low regret and the uneven action selection it causes—which we again emphasize is extreme in the case of LinUCB, even in comparison to  $\epsilon$ -greedy—renders the testing problem harder in the paragraph below. As such, we leave the problem of developing a powerful resampling scheme for deterministic algorithms like LinUCB in this type of contextual  $C$ -stationary strongly non-reactive environment to future work.

Lastly, we hypothesize that the shift in relative power of the uniform i.i.d. baseline (from the least powerful adaptive assignment-resampling algorithm combination for small  $T$  to the most powerful for large  $T$ ) is again an artifact of the low-regret properties of  $\epsilon$ -greedy and LinUCB. This low regret results in the outlier context-action-response triple sampled at the  $T^{\text{th}}$  timestep to often have action agreeing with the action taken during most of the first  $T - 1$  timesteps. Thus, for small  $T$ , these first  $T - 1$  context-treatment-response triples may outweigh the effect of the outlier at time  $T$  when training  $\hat{\mu}_D$  via Lasso with cross-validation, thereby resulting in better outlier detection than that under the uniform i.i.d. base, wherein only around half of the first  $T - 1$  timesteps (and thus very few, in total) will have action agreeing with that at the  $T^{\text{th}}$ . The effect however, is diminished with large  $T$ , most likely because in that regime even half of the data generated via  $Y_t$ 's true (null) conditional distribution, which agrees with the action taken at time  $T$ , is enough to offset the effect of the outlier in training  $\hat{\mu}_D$ . In particular, the remaining timesteps whose corresponding action differs from the one taken at time  $T$ —of which there are more under the uniform i.i.d. baseline than under  $\epsilon$ -greedy (and *many* more under the baseline than under LinUCB) for large  $T$ —may allow for a better fit  $\hat{\mu}_D$  through more effective learning of the dependence of  $Y_t$  on  $X_t$ . As noted above, this may also explain the difficulty in attaining high power for large  $T$  under LinUCB in this environment. We empirically validate this hypothesis in Figure 24 of Appendix G.3 by showing that the uniform i.i.d. baseline exhibits precisely this same relative performance, with respect to  $T$ , in comparison to a biased i.i.d. assignment algorithm that selects action  $-1$  with probability 0.9 and otherwise selects action 1.

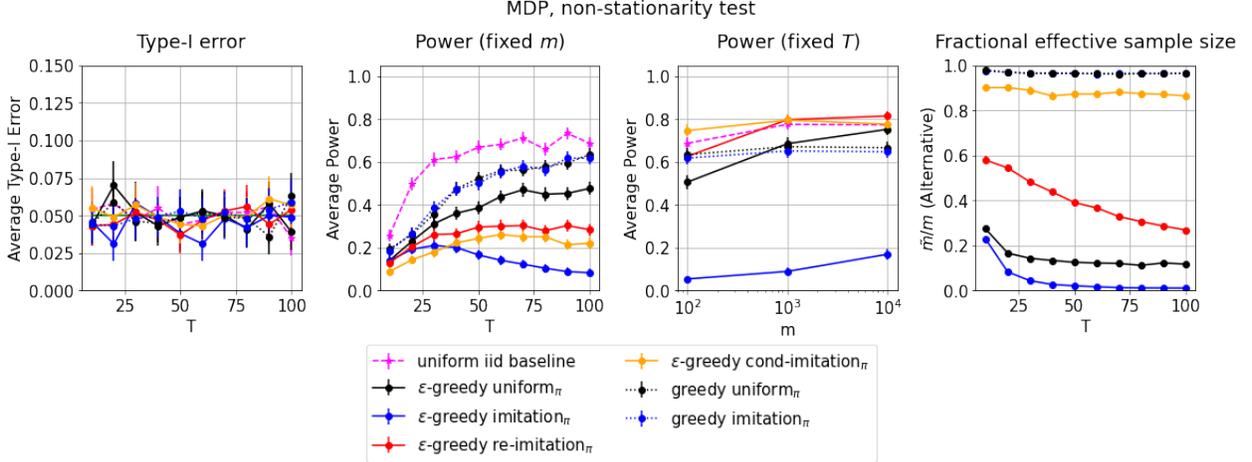


Figure 5: Type-I error rate (leftmost) and power (second from right) of the MC randomization test at fixed  $m = 100$  and varying  $T$  as well as power for fixed  $T = 100$  and varying  $m$  (third from right) and fractional effective sample size at fixed  $m = 100$  and varying  $T$  (rightmost) in a contextless  $C$ -stationary strongly non-reactive environment with data gathered via  $\epsilon$ -greedy and greedy  $Q$ -learning.

### 5.2.3 Testing non-stationarity in an MDP

Our final set of hypothesis testing simulations is in an MDP, where recall that  $Y_t = C_{t+1}$  and hence we drop the notation  $Y_t$  and refer to  $C_t$  as states instead of contexts.<sup>9</sup> The precise specifications of environment involve a state space of  $\mathcal{C} = \{0, 1, 2\}$ , action space  $\mathcal{X} = \{-1, 1\}$ , and transition kernel during the first  $T - 1$  transitions given by

$$C_{t+1} \mid (C_t, X_t) \sim \begin{cases} C_t + X_t \pmod{3} & \text{with probability } 0.95 \\ C_t - X_t \pmod{3} & \text{with probability } 0.05 \end{cases}.$$

Under the null hypothesis, the transition distribution at time  $T$  remains the same as above, but under the alternative, we swap the role of the two actions so that the alternative distribution is given by

$$C_{T+1} \mid (C_T, X_T) \sim \begin{cases} C_T - X_T \pmod{3} & \text{with probability } 0.95 \\ C_T + X_T \pmod{3} & \text{with probability } 0.05 \end{cases}.$$

Finally, the test statistic that we use in these simulations trains a decision tree classifier on the data  $D$  and outputs the negative log likelihood loss of the trained model on the triple  $(C_T, X_T, C_{T+1})$  at time  $T$ .

Figure 5 summarizes the results for our simulations in this setting. Once again, Type-I error is controlled at exactly the nominal level, as guaranteed by our theory. In a departure from the results of the previous two sections (most likely due to the sequential dependence in this environment not present in the last two as well as the modified sampling process that it requires), uniform $_{\pi}$  resampling has the greatest power for  $Q$ -learning with  $\epsilon$ -greedy action selection. Additionally, both uniform $_{\pi}$  and imitation $_{\pi}$  resampling perform similarly well for  $Q$ -learning with greedy action selection. Again, using both the fractional effective sample size plots and power plots with fixed  $T$  but varying  $m$ , we see that under  $\epsilon$ -greedy, re-imitation $_{\pi}$  resampling has lower effective sample size than uniform $_{\pi}$ , but higher diversity, leading it to perform worse for smaller  $m$  but better for large  $m$ . Finally, we note that, in comparison to the uniform i.i.d. baseline<sup>10</sup>, uniform $_{\pi}$

<sup>9</sup>As discussed in Section 4.2, the dataset we consider in these simulations is the usual dataset  $D$  along with one final action taken at time  $T + 1$ :  $((C_1, X_1, C_2), \dots, (C_T, X_T, C_{T+1}), X_{T+1})$  as this allows us to simply permute the state-action pairs  $(C_t, X_t)$ .

<sup>10</sup>In contrast to the last section, the uniform i.i.d. data used in this section is not actually gathered in the same environment as the other adaptive assignment algorithms, and in particular, the uniform i.i.d. data is *not* MDP data. Rather, the data comprises state-action-next state triples  $(C, X, C')$  sampled i.i.d. from a distribution in which both  $C$  and  $X$  are independent and uniform over  $\mathcal{C}$  and  $\mathcal{X}$  respectively, and  $C'$  is sampled from the transition distribution described above, conditional on  $(C, X)$ .

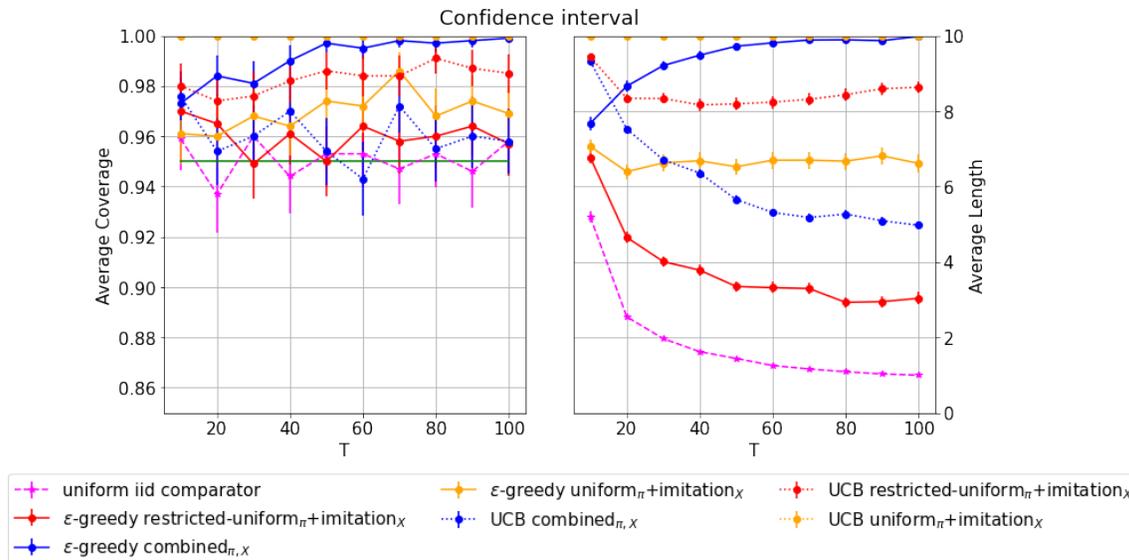


Figure 6: Coverage and average length of confidence intervals for location difference between  $Y_t | (X_t = 0)$ , and  $Y_t | (X_t = 1)$  (i.e.,  $b_0 = 4$ ) using the MC randomization test with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

and  $\text{imitation}_\pi$  are very competitive under the greedy  $Q$ -learning adaptive assignment, and  $\text{imitation}_\pi$  also exhibits quite high power under  $\epsilon$ -greedy.

### 5.3 Constructing confidence and prediction intervals

In this section we apply our framework to constructing confidence and prediction intervals using the inversion procedures described in Section 3.3. In particular, we plot both the coverage of the intervals as well as their average length, averaged over 1000 trials for varying  $T$  at a fixed number of MC samples  $m = 100$ . For construction of conformal intervals, we also apply the sample sharing described in Remark 7 (in addition to the standard gridding procedure described in Section 3.3) at both  $m = 10$  and  $m = 100$ . Finally, we note that some of the resampling procedures in certain environments and adaptive assignment algorithms considered in this section exhibit overcoverage, in contrast to the last section, in which all tests attained Type-I error control at exactly the nominal level. We remind the reader that this is due to the fact that the gridding procedure described in Section 3.3 is based on the (approximate) inversion of a conservative test, rather than that of the smoothed test described in Remark 2.

#### 5.3.1 Confidence intervals

We now discuss our simulations constructing confidence intervals using the gridding procedure described at the end of Section 3.3. We consider essentially the same environment as the contextless stationary strongly non-reactive environment discussed in Section 5.1.2, except that here we have

$$Y_t | X_t \sim \begin{cases} \mathcal{N}(0, 1) & \text{if } X_t = 0 \\ \mathcal{N}(b_0, 1) & \text{if } X_t = 1 \\ \mathcal{N}(2, 1) & \text{if } X_t = 2 \end{cases},$$

with  $b_0 = 4$ ; we construct a confidence interval for the location difference between  $Y_t | (X_t = 0)$ , and  $Y_t | (X_t = 1)$ , which simply corresponds to  $b_0$ , using a gridding set  $\mathcal{Y}' = [-1 : 9]$ .

Figure 6 summarizes the results for these simulations. We note there is no evidence in our simulations that the necessary (and in fact rather coarse) gridding results in any undercoverage. In terms of length, our results

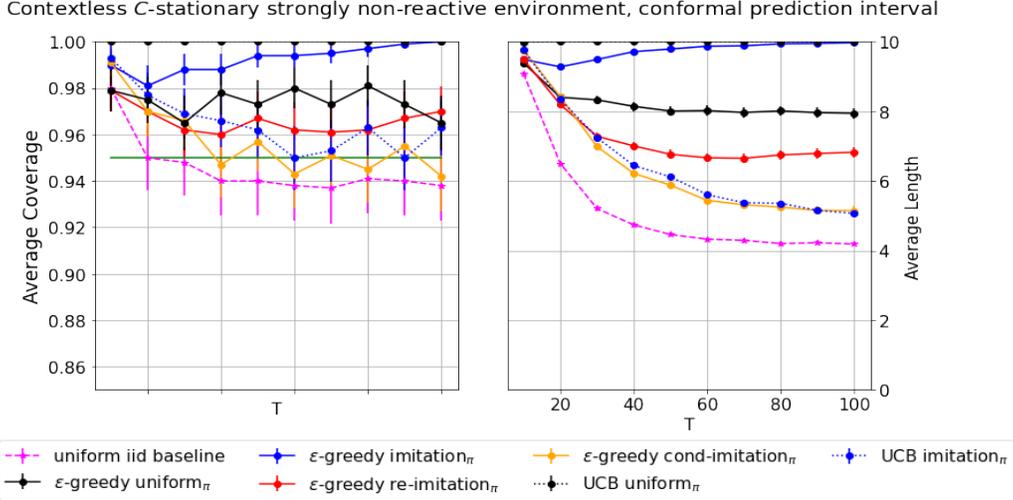


Figure 7: Coverage and average length of conformal prediction intervals for  $Y_T$  using the MC randomization test with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

align with the power results presented in Section 5.1.2. In particular, under an  $\epsilon$ -greedy adaptive assignment, restricted-uniform $_{\pi}$ +imitation $_X$  sampling produces the shortest average length intervals, whereas under UCB, combined $_{\pi,X}$  resampling does.

### 5.3.2 Conformal prediction intervals

Finally, we discuss our simulations constructing conformal prediction intervals. The environment considered in this section is the same as contextless  $C$ -stationary strongly non-reactive environment considered in Section 5.2.1 except that we do not have access to  $Y_T$ ; it is the quantity for which we construct the conformal prediction region.

Our simulations in this section use the gridding procedure described at the end of Section 3.3 and assess the benefit of the sharing of samples described in Remark 7 and Appendix E.2. Without sample sharing, we fix  $m = 100$  and set  $\mathcal{Y}' = [-5 : 5]$  and plot coverage and length at varying  $T$ . With sample sharing, we use the same gridding set  $\mathcal{Y}'$  and study both  $m = 10$  and  $m = 100$  number of MC samples<sup>11</sup> and varying  $T$ . As discussed in Appendix E.2 this sample sharing results in non-conditionally-i.i.d. draws because the resamples generated corresponding to  $y_1 \in \mathcal{Y}'$  may come from (and indeed do in our simulations) a different distribution than those sampled according to  $y_2 \in \mathcal{Y}'$  for  $y_1 \neq y_2$ . As such, following Remark 8 (and as discussed in further detail in Appendix E.1), we choose a subset of permutations  $\Sigma$  to be the set of  $m + 1$  permutations swapping 0 and  $i$  for each  $i \in [0 : m]$  and employ the test of Algorithm 1 with weights  $w_{\mathcal{D}}^{\tilde{q}, \Sigma}(\tilde{D}^{(i)})$ .

Figures 7 and 8 summarize our simulation results constructing approximate conformal prediction intervals. Figure 7 illustrates both coverage and length as  $T$  grows by simply using the standard gridding procedure described at the end of Section 3.3. With regards to average length, our results mirror those in Section 5.2.1, in which cond-imitation $_{\pi}$  resampling is best for an  $\epsilon$ -greedy adaptive assignment, and imitation $_{\pi}$  is best for UCB. In Figure 8, we observe that any undercoverage, attributable to the gridding of  $\mathcal{Y}$ , is relatively minor, and should be fixable by using a finer-resolution grid; the reason that the uniform i.i.d. baseline tends to exhibit the greatest undercoverage is most likely explained by that fact that we are not smoothing in this set of experiments, so that whereas all other procedures are conservative, the uniform i.i.d. baseline is not. Additionally, we see essentially the same ranking of resampling procedures, although re-imitation $_{\pi}$  does perform slightly better than cond-imitation $_{\pi}$  when 100 MC samples are used. Finally, the shortest average length curves in Figure 8 using  $m = 10$  are quite competitive in comparison to their counterparts in

<sup>11</sup>For each  $y \in \mathcal{Y}'$ ,  $m$  samples are generated, but all  $m|\mathcal{Y}'|$  samples are used to determine the membership of  $y \in \mathcal{Y}'$  in the prediction interval

Contextual  $\mathcal{C}$ -stationary strongly non-reactive environment, conformal prediction interval, shared samples

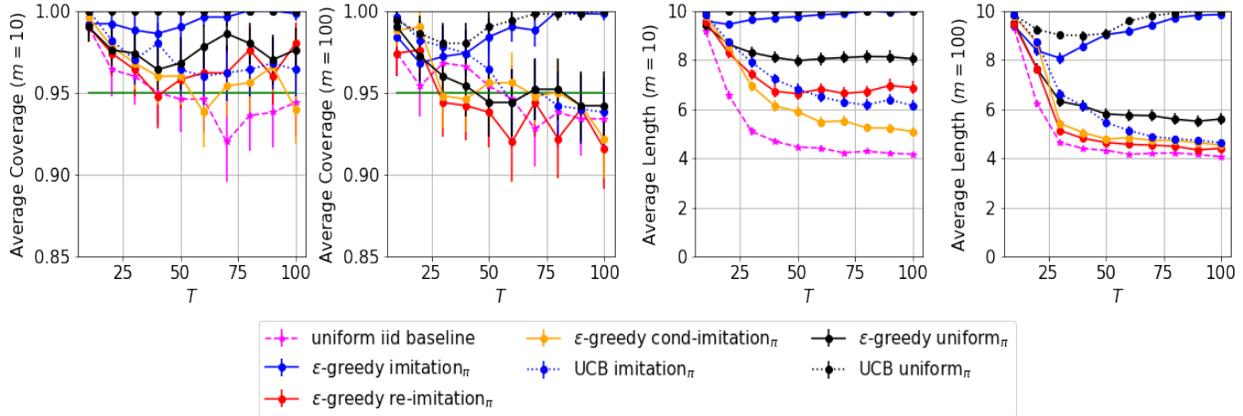


Figure 8: Coverage and average length of approximate conformal prediction intervals, with sample sharing, for  $Y_T$  using the MC randomization test with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

Figure 7. In particular,  $\text{imitation}_\pi$  under UCB produces only slightly wider intervals with  $m = 10$  and sample sharing than with  $m = 100$  and no sample sharing; for  $\epsilon$ -greedy,  $\text{cond-imitation}_\pi$  performs essentially the same in both settings. This demonstrates that, by utilizing sample sharing, we are able to obtain conformal prediction intervals of nearly the same length at a fraction of the computational cost (because in the left hand column of Figure 8 we generate a *total* of 110 samples, whereas in Figure 7, we generate 100 samples *for each*  $y \in \mathcal{Y}$ ). Indeed, as illustrated by Figures 21 and 22 of Appendix G, the computation time required by the procedure using  $m = 10$  and shared sampling is, for most resampling schemes and adaptive assignments, more than 5 times faster than its  $m = 100$  non-sample sharing counterpart.

## 6 Discussion

In this paper, we study the problem of performing various challenging inferential tasks on adaptively collected data. Through the development of a weighted MC randomization test (along with the novel application of the unweighted MCMC randomization test of Besag and Clifford [1989]), we show that such tasks can be performed with *exact* Type-I error control with a great degree of flexibility in the choice of resampling procedure as long as the proportionality hypothesis  $\mathcal{H}_0^\infty$  is satisfied.

The question, however, of how *best* to perform these tests still remains, and is an interesting one for future research. In particular, while we have discussed some preliminary empirical results demonstrating powerful resampling algorithms in a number of different challenging environments, there are a number of interesting and potentially fruitful directions forward. In particular, as mentioned in Section 5.2.2, one important direction for future work is to develop resampling procedures that are powerful for deterministic algorithms like LinUCB in a contextual  $\mathcal{C}$ -stationary strongly non-reactive setting. But more generally, it may be interesting to formally investigate any commonalities between powerful resampling algorithms in different environments in order to develop a more principled method for deciding which procedure to apply to a particular problem (i.e., combination of adaptive assignment algorithm, environment, and inferential task).

Along this same vein of how best to resample, it may be useful to also consider different resampling procedures which incorporate non-conditionally-i.i.d. sampling. Utilization of these non-conditionally-i.i.d. resampling schemes under the unweighted MCMC randomization testing framework has the potential to be especially advantageous as the full MCMC test will still require only  $O(m)$  number of computations. On the other hand, as discussed in Remark 8 and Appendix E.1, when applying such sampling schemes to the weighted MC test, we may generalize and choose a subset  $\Sigma$  of the full permutation set  $\Pi_{[0:m]}$  different than  $\Sigma_{\text{swap}}$ , and condition only on the data permutations induced by these sets rather than the full set of permutations.

As such, one final direction for future work that may also be of interest is to investigate how the power of

the weighted MC test changes under different choices of  $\Sigma$ . Specifically, while we expect the test to be more powerful for choices of  $\Sigma$  which are larger, we also expect such larger choices to require a greater amount of computation time. Hence, it may be worthwhile to study this tradeoff between the size of the sets  $\Sigma$  and computation time. Just as in the potential investigation of more powerful resampling procedures mentioned above, exploring if there are more well-structured connections between not-too-large choices of  $\Sigma$  that yield a powerful test and the specific problem at hand may allow for the development of a more coherent theory—or at least a more formal set of guidelines—for how best to select/develop a resampling algorithm for any given problem.

**Acknowledgements** YN and LJ were partially supported by NSF CBET-21120 and LJ was partially supported by DMS-2045981. The authors would also like to thank Subhabrata Sen, Joe Blitzstein, and Henry Bosch for valuable discussions.

## References

- Carlos Agustín Alesso, Pablo Ariel Cipriotti, Germán Alberto Bollero, and Nicolas Federico Martin. Design of on-farm precision experiments to estimate site-specific crop responses. *Agronomy Journal*, 113(2):1366–1380, 2021. doi: <https://doi.org/10.1002/agj2.20572>. URL <https://access.onlinelibrary.wiley.com/doi/abs/10.1002/agj2.20572>.
- Jonathan Alvarsson, Staffan Arvidsson McShane, Ulf Norinder, and Ola Spjuth. Predicting with confidence: Using conformal prediction in drug discovery. *Journal of Pharmaceutical Sciences*, 110(1):42–49, 2021. ISSN 0022-3549. doi: <https://doi.org/10.1016/j.xphs.2020.09.055>. URL <https://www.sciencedirect.com/science/article/pii/S002235492030589X>.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.
- Yikun Ban and Jingrui He. Generic outlier detection in multi-armed bandit. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 913–923, 2020.
- Rina Barber, Emmanuel Candes, Aaditya Ramdas, and Ryan Tibshirani. Conformal prediction beyond exchangeability. *arXiv preprint arXiv:2202.13415*, 2022.
- Rina Foygel Barber and Lucas Janson. Testing goodness-of-fit and conditional independence with approximate co-sufficient sampling. *Annals of Statistics*, 2022+. To Appear.
- Stephen Bates, Emmanuel Candès, Lucas Janson, and Wenshuo Wang. Metropolized knockoff sampling. *Journal of the American Statistical Association*, 116(535):1413–1427, 2021. doi: 10.1080/01621459.2020.1729163. URL <https://doi.org/10.1080/01621459.2020.1729163>.
- Samuel L. Battalio, David E. Conroy, Walter Dempsey, Peng Liao, Marianne Menictas, Susan Murphy, Inbal Nahum-Shani, Tianchen Qian, Santosh Kumar, and Bonnie Spring. Sense2stop: A micro-randomized trial using wearable sensors to optimize a just-in-time-adaptive stress management intervention for smoking relapse prevention. *Contemporary Clinical Trials*, 109:106534, 2021. ISSN 1551-7144. doi: <https://doi.org/10.1016/j.cct.2021.106534>. URL <https://www.sciencedirect.com/science/article/pii/S1551714421002706>.
- Thomas B. Berrett, Yi Wang, Rina Foygel Barber, and Richard J. Samworth. The conditional permutation test for independence while controlling for confounders. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2019.
- Thomas B. Berrett, Yi Wang, Rina Foygel Barber, and Richard J. Samworth. The conditional permutation test for independence while controlling for confounders. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(1):175–197, 2020. doi: <https://doi.org/10.1111/rssb.12340>. URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/rssb.12340>.

- Julian Besag and Peter Clifford. Generalized monte carlo significance tests. *Biometrika*, 76(4):633–642, 1989.
- Aurélien Bibaut, Nathan Kallus, Maria Dimakopoulou, Antoine Chambaz, and Mark van der Laan. Risk minimization from adaptively collected data: Guarantees for supervised and policy learning. *Advances in Neural Information Processing Systems*, 34:19261–19273, 2021.
- Iavor Bojinov and Neil Shephard. Time series experiments and causal estimands: exact randomization tests and trading. *Journal of the American Statistical Association*, 114(528):1665–1682, 2019.
- Federico A Bugni, Ivan A Canay, and Azeem M Shaikh. Inference under covariate-adaptive randomization. *Journal of the American Statistical Association*, 113(524):1784–1796, 2018.
- Evgeny Burnaev and Vladimir Vovk. Efficiency of conformalized ridge regression. In *Conference on Learning Theory*, pages 605–622. PMLR, 2014.
- Emmanuel Candes, Yingying Fan, Lucas Janson, and Jinchi Lv. Panning for gold: ‘model-x’ knockoffs for high dimensional controlled variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(3):551–577, 2018.
- Stefano Caria, Maximilian Kasy, Simon Quinn, Soha Shami, Alex Teytelboym, et al. An adaptive targeted field experiment: Job search assistance for refugees in jordan. 2020.
- Stephanie Carpenter, Marianne Menictas, Inbal Nahum-Shani, David Wetter, and Susan Murphy. Developments in mobile health just-in-time adaptive interventions for addiction science. *Current Addiction Reports*, 7, 09 2020. doi: 10.1007/s40429-020-00322-y.
- Wenyu Chen, Kelli-Jean Chun, and Rina Foygel Barber. Discretized conformal prediction for efficient distribution-free inference. *Stat*, 7(1):e173, 2018. doi: <https://doi.org/10.1002/sta4.173>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sta4.173>. e173 sta4.173.
- Victor Chernozhukov, Kaspar Wüthrich, and Zhu Yinchu. Exact and robust conformal inference methods for predictive machine learning with dependent data. In *Conference On Learning Theory*, pages 732–749. PMLR, 2018.
- Daniel D Coppersmith, Walter Dempsey, Evan Kleiman, Kate Bentley, Susan Murphy, and Matthew Nock. Just-in-time adaptive interventions for suicide prevention: Promise, challenges, and future directions, Jan 2021. URL <https://arxiv.org/abs/2012.03613>.
- Yash Deshpande, Lester Mackey, Vasilis Syrgkanis, and Matt Taddy. Accurate inference for adaptive linear models. In *International Conference on Machine Learning*, pages 1194–1203. PMLR, 2018.
- P. Diaconis and D. Freedman. De Finetti’s Theorem for Markov Chains. *The Annals of Probability*, 8(1): 115 – 130, 1980. doi: 10.1214/aop/1176994828. URL <https://doi.org/10.1214/aop/1176994828>.
- Eugene Edgington and Patrick Onghena. *Randomization tests*. Chapman and Hall/CRC, 2007.
- Clara Fannjiang, Stephen Bates, Anastasios N. Angelopoulos, Jennifer Listgarten, and Michael I. Jordan. Conformal prediction for the design problem. *CoRR*, abs/2202.03613, 2022. URL <https://arxiv.org/abs/2202.03613>.
- Ronald Aylmer Fisher et al. The design of experiments. *The design of experiments.*, (2nd Ed), 1937.
- Isaac Gibbs and Emmanuel Candès. Adaptive conformal inference under distribution shift. 2021. Accepted.
- Francis J Giles, Hagop M Kantarjian, Jorge E Cortes, Guillermo Garcia-Manero, Srdan Verstovsek, Stefan Faderl, Deborah A Thomas, Alessandra Ferrajoli, Susan O’Brien, Jay K Wathen, et al. Adaptive randomized study of idarubicin and cytarabine versus troxacitabine and cytarabine versus troxacitabine and idarubicin in untreated patients 50 years or older with adverse karyotype acute myeloid leukemia. *Journal of clinical oncology*, 21(9):1722–1727, 2003.

- Alexander G. Godfrey, Thierry Masquelin, and Horst Hemmerle. A remote-controlled adaptive medchem lab: an innovative approach to enable drug discovery in the 21st century. *Drug Discovery Today*, 18(17):795–802, 2013. ISSN 1359-6446. doi: <https://doi.org/10.1016/j.drudis.2013.03.001>. URL <https://www.sciencedirect.com/science/article/pii/S135964461300069X>.
- Vitor Hadad, David A Hirshberg, Ruohan Zhan, Stefan Wager, and Susan Athey. Confidence intervals for policy evaluation in adaptive experiments. *Proceedings of the National Academy of Sciences*, 118(15), 2021.
- Jinyong Hahn, Keisuke Hirano, and Dean Karlan. Adaptive experimental design using the propensity score. *Journal of Business & Economic Statistics*, 29(1):96–108, 2011.
- Dae Woong Ham and Jiase Qiu. Hypothesis testing in sequentially sampled data: Adaprt to maximize power beyond iid sampling, 2022. URL <https://arxiv.org/abs/2205.02430>.
- Matthew T. Harrison. Conservative hypothesis tests and confidence intervals using importance sampling. *Biometrika*, 99(1):57–69, 2012. ISSN 00063444, 14643510. URL <http://www.jstor.org/stable/41720672>.
- Jesse Hemerik and Jelle Goeman. Exact testing with random permutations. *TEST*, 27(4):811–825, 2018. doi: 10.1007/s11749-017-0571-1. URL <https://doi.org/10.1007/s11749-017-0571-1>.
- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021.
- Xiaoyu Hu and Jing Lei. A distribution-free test of covariate shift using conformal prediction. *arXiv preprint arXiv:2010.07147*, 2020.
- Dongming Huang and Lucas Janson. Relaxing the assumptions of knockoffs by conditioning. *Ann. Statist.*, 48(5):3021–3042, 10 2020. doi: 10.1214/19-AOS1920. URL <https://doi.org/10.1214/19-AOS1920>.
- Xiaoguang Huo and Feng Fu. Risk-aware multi-armed bandit problem with application to portfolio selection. *Royal Society open science*, 4(11):171377, 2017.
- Emilie Kaufmann and Wouter M Koolen. Mixture martingales revisited with applications to sequential tests and confidence intervals. *Journal of Machine Learning Research*, 22(246):1–44, 2021.
- T.L Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985. ISSN 0196-8858. doi: [https://doi.org/10.1016/0196-8858\(85\)90002-8](https://doi.org/10.1016/0196-8858(85)90002-8). URL <https://www.sciencedirect.com/science/article/pii/0196885885900028>.
- Tze Leung Lai and Ching Zong Wei. Least squares estimates in stochastic regression models with applications to identification and control of dynamic systems. *The Annals of Statistics*, 10(1):154–166, 1982.
- Hyun-Rok Lee and Taesik Lee. Markov decision process model for patient admission decision at an emergency department under a surge demand. *Flexible Services and Manufacturing Journal*, 30(1):98–122, 2018.
- Erich Leo Lehmann, Joseph P Romano, and George Casella. *Testing statistical hypotheses*, volume 3. Springer, 2005.
- Jing Lei and Larry Wasserman. Distribution-free prediction bands for non-parametric regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):71–96, 2014.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.
- Tianhao Li, Zhishun Wang, Wei Lu, Qian Zhang, and Dengfeng Li. Electronic health records based reinforcement learning for treatment optimizing. *Information Systems*, 104:101878, 2022.

- Wei Ma, Feifang Hu, and Lixin Zhang. Testing hypotheses of covariate-adaptive randomized clinical trials. *Journal of the American Statistical Association*, 110(510):669–680, 2015.
- Rajiv Mahajan and Kapil Gupta. Adaptive design clinical trials: Methodology, challenges and prospect. *Indian journal of pharmacology*, 42(4):201, 2010.
- J er mie Mary, Romaric Gaudel, and Philippe Preux. Bandits and recommender systems. In *International Workshop on Machine Learning, Optimization and Big Data*, pages 325–336. Springer, 2015.
- Aditya Mate, Jackson Killian, Haifeng Xu, Andrew Perrault, and Milind Tambe. Collapsing bandits and their application to public health intervention. *Advances in Neural Information Processing Systems*, 33: 15639–15650, 2020.
- Aditya Mate, Lovish Madaan, Aparna Taneja, Neha Madhiwalla, Shresth Verma, Gargi Singh, Aparna Hegde, Pradeep Varakantham, and Milind Tambe. Field study in deploying restless multi-armed bandits: Assisting non-profits in improving maternal and child health, 2021. URL <https://arxiv.org/abs/2109.08075>.
- Inbal Nahum-Shani, Shawna Smith, Bonnie Spring, Linda Collins, Katie Witkiewitz, Ambuj Tewari, and Susan Murphy. Just-in-time adaptive interventions (jitais) in mobile health: Key components and design principles for ongoing health behavior support. *Annals of Behavioral Medicine*, 52, 09 2016. doi: 10.1007/s12160-016-9830-8.
- Buddhika Nettasinghe, Samrat Chatterjee, Ramakrishna Tipireddy, and Mahantesh Halappanavar. Extending conformal prediction to hidden markov models with exact validity via de finetti’s theorem for markov chains, 2022. URL <https://arxiv.org/abs/2210.02271>.
- Philip Pallmann, Alun W. Bedding, Babak Choodari-Oskooei, Munyaradzi Dimairo, Laura Flight, Lisa V. Hampson, Jane Holmes, Adrian P. Mander, Lang’o Odondi, Matthew R. Sydes, Sof a S. Villar, James M. S. Wason, Christopher J. Weir, Graham M. Wheeler, Christina Yap, and Thomas Jaki. Adaptive designs in clinical trials: why use them, and how to run and report them. *BMC Medicine*, 16(1):29, 2018. doi: 10.1186/s12916-018-1017-7. URL <https://doi.org/10.1186/s12916-018-1017-7>.
- Harris Papadopoulos, Volodya Vovk, and Alex Gammerman. Conformal prediction with neural networks. In *19th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2007)*, volume 2, pages 388–395, 2007. doi: 10.1109/ICTAI.2007.47.
- E. J. G. Pitman. Significance tests which may be applied to samples from any populations. *Supplement to the Journal of the Royal Statistical Society*, 4(1):119–130, 1937. doi: <https://doi.org/10.2307/2984124>. URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.2307/2984124>.
- Stuart J Pocock and Richard Simon. Sequential treatment assignment with balancing for prognostic factors in the controlled clinical trial. *Biometrics*, pages 103–115, 1975.
- Mariya Popova, Olexandr Isayev, and Alexander Tropsha. Deep reinforcement learning for de novo drug design. *Science advances*, 4(7):eaap7885, 2018.
- Dustin J Rabideau and Rui Wang. Randomization-based confidence intervals for cluster randomized trials. *Biostatistics*, 22(4):913–927, 2021.
- Jeffrey Roach and William Valdar. Permutation tests of non-exchangeable null models. *arXiv preprint arXiv:1808.10483*, 2018.
- William F. Rosenberger, Diane Uschner, and Yanying Wang. Randomization: The forgotten component of the randomized clinical trial. *Statistics in Medicine*, 38(1):1–12, 2019. doi: <https://doi.org/10.1002/sim.7901>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.7901>.
- J Ben Schafer, Joseph Konstan, and John Riedl. Recommender systems in e-commerce. In *Proceedings of the 1st ACM conference on Electronic commerce*, pages 158–166, 1999.

- Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3), 2008.
- Richard Simon. Restricted randomization designs in clinical trials. *Biometrics*, pages 503–512, 1979.
- Alejandro Suárez-Hernández, Carme Torras, and Guillem Alenya. Practical resolution methods for mdps in robotics exemplified with disassembly planning. *IEEE Robotics and Automation Letters*, 4(3):2282–2288, 2019.
- Ambuj Tewari and Susan A Murphy. From ads to interventions: Contextual bandits in mobile health. In *Mobile Health*, pages 495–517. Springer, 2017.
- Ryan J. Tibshirani, Rina Foygel Barber, Emmanuel J. Candès, and Aaditya Ramdas. Conformal prediction under covariate shift, 2020.
- Vladimir Vovk. Transductive conformal predictors. In *IFIP International Conference on Artificial Intelligence Applications and Innovations*, pages 348–360. Springer, 2013.
- Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer-Verlag, Berlin, Heidelberg, 2005. ISBN 0387001522.
- Vladimir Vovk, Ilia Nouretdinov, and Alex Gammerman. On-line predictive linear regression. *The Annals of Statistics*, 37(3):1566 – 1590, 2009. doi: 10.1214/08-AOS622. URL <https://doi.org/10.1214/08-AOS622>.
- Christopher John Cornish Hellaby Watkins. Learning from delayed rewards. 1989.
- Chen Xu and Yao Xie. Conformal prediction interval for dynamic time-series. In *International Conference on Machine Learning*, pages 11559–11569. PMLR, 2021.
- Kelly Zhang, Lucas Janson, and Susan Murphy. Statistical inference with m-estimators on adaptively collected data. December 2021.
- Kelly Zhang, Lucas Janson, and Susan Murphy. Statistical inference after adaptive sampling in non-markovian environments. *arXiv preprint arXiv:2202.07098*, 2022.
- Kelly W. Zhang, Lucas Janson, and Susan Murphy. Inference for batched bandits. December 2020.
- Yao Zhang and Qingyuan Zhao. What is a randomization test?, 2022. URL <https://arxiv.org/abs/2203.10980>.

## A Proof of Theorem 2.1

We consider the case of discrete data  $D$  and discrete resamples  $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)}$  and apply Bayes’ theorem. Define  $\mathfrak{d} = (d_0, \dots, d_m)$  to be some list of data sets,  $\mathfrak{d}_{-0}$  to be the list  $(d_1, \dots, d_m)$ ,  $\mathfrak{D}_{-0}$  to denote the list  $(\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)})$ , and  $\text{orb}(\mathfrak{d}) := \{\mathfrak{d}^\pi : \pi \in \Pi_{[0:m]}\}$ , where  $\mathfrak{d}^\pi := (d_{\pi(0)}, \dots, d_{\pi(m)})$ ; we also use  $(\mathfrak{d}_{-0})^\pi$  to denote the permuted list  $(d_{\pi(1)}, \dots, d_{\pi(m)})$  for any  $\pi \in \Pi_{[0:m]}$ . Then, by Bayes’ theorem, we have that

$$\begin{aligned}
\mathbb{P}(D = d_i | \mathfrak{D} \in \text{orb}(\mathfrak{d})) &= \frac{\mathbb{P}(\mathfrak{D} \in \text{orb}(\mathfrak{d}) | D = d_i) f(d_i)}{\sum_{\text{distinct } d_j \in \mathfrak{d}} \mathbb{P}(\mathfrak{D} \in \text{orb}(\mathfrak{d}) | D = d_j) f(d_j)} \\
&= \frac{\mathbb{P}(\mathfrak{D}_{-0} \in \{(\mathfrak{d}_{-0})^\pi : \pi \in \Pi_{[0:m]} \text{ such that } d_{\pi(0)} = d_i\} | D = d_i) f(d_i)}{\sum_{\text{distinct } d_j \in \mathfrak{d}} \mathbb{P}(\mathfrak{D}_{-0} \in \{(\mathfrak{d}_{-0})^\pi : \pi \in \Pi_{[0:m]} \text{ such that } d_{\pi(0)} = d_j\} | D = d_j) f(d_j)} \\
&= \frac{|\{(\mathfrak{d}_{-0})^\pi : \pi \in \Pi_{[0:m]} \text{ such that } d_{\pi(0)} = d_i\}| \left( \prod_{k \in [0:m] \setminus \{i\}} q(d_k | d_i) \right) f(d_i)}{\sum_{\text{distinct } d_j \in \mathfrak{d}} |\{(\mathfrak{d}_{-0})^\pi : \pi \in \Pi_{[0:m]} \text{ such that } d_{\pi(0)} = d_j\}| \left( \prod_{k \in [0:m] \setminus \{j\}} q(d_k | d_j) \right) f(d_j)}.
\end{aligned}$$

where we say that  $d \in \mathfrak{d}$  if  $d$  is an element of the list  $\mathfrak{d}$  and where the last equality follows from the conditional i.i.d.-ness of the resamples given  $D$ .

Letting  $m_{d_i}(\mathfrak{d})$  denote the number of times  $d_i$  appears in the list  $\mathfrak{d}$  we have that

$$\begin{aligned} |\{(\mathfrak{d}_{-0})^\pi : \pi \in \Pi_{[0:m]} \text{ with } d_{\pi(0)} = d_i\}| &= \frac{m!}{(m_{d_i}(\mathfrak{d}) - 1)! \prod_{\substack{\text{distinct } d \in \mathfrak{d} \\ \text{not equal to } d_i}} m_d(\mathfrak{d})!} \\ &= m_{d_i}(\mathfrak{d}) \cdot \frac{m!}{\prod_{\text{distinct } d \in \mathfrak{d}} m_d(\mathfrak{d})!} \end{aligned}$$

and hence,

$$\begin{aligned} \mathbb{P}(D = d_i | \mathfrak{D} \in \text{orb}(\mathfrak{d})) &= \frac{m_{d_i}(\mathfrak{d}) f(d_i) \prod_{k \in [0:m] \setminus \{i\}} q(d_k | d_i)}{\sum_{j=0}^m f(d_j) \prod_{k \in [0:m] \setminus \{j\}} q(d_k | d_j)} \\ &= \frac{m_{d_i}(\mathfrak{d}) \hat{f}(d_i) \prod_{k \in [0:m] \setminus \{i\}} q(d_k | d_i)}{\sum_{j=0}^m \hat{f}(d_j) \prod_{k \in [0:m] \setminus \{j\}} q(d_k | d_j)} \\ &= m_{d_i}(\mathfrak{d}) w_{\mathfrak{d}}^q(d_i), \end{aligned}$$

where the second-to-last inequality is because of the proportionality hypothesis  $\mathcal{H}_0^\infty$ . Equivalently, we can think of  $D$ 's conditional distribution given  $\text{orb}(\mathfrak{D})$  as a draw from the  $m + 1$  elements of  $\mathfrak{D}$ , with weight on each element given by the  $w_{\mathfrak{D}}^q$  function applied to that element. Defining the set  $\mathfrak{S} := \{S(\tilde{D}) : \tilde{D} \in \mathfrak{D}\}^{12}$ ,  $S(D)$  can be viewed as a draw from the weighted distribution on  $\mathfrak{S}$  with the total weight of each element  $S(\tilde{D}^{(i)})$  equal to

$$\sum_{j:S(\tilde{D}^{(j)})=S(\tilde{D}^{(i)})} \frac{\hat{f}(\tilde{D}^{(j)}) \prod_{k \in [0:m] \setminus \{j\}} q(\tilde{D}^{(k)} | \tilde{D}^{(j)})}{\sum_{j'=0}^m \hat{f}(\tilde{D}^{(j')}) \prod_{k' \in [0:m] \setminus \{j'\}} q(\tilde{D}^{(k')} | \tilde{D}^{(j')})} = \sum_{j:S(\tilde{D}^{(j)})=S(\tilde{D}^{(i)})} w_{\mathfrak{D}}^q(\tilde{D}^{(j)}).$$

Finally, note that

$$\begin{aligned} \mathbb{P}(p \leq \alpha) &= \mathbb{P}\left(\sum_{i=0}^m w_{\mathfrak{D}}^q(\tilde{D}^{(i)}) \mathbf{1}[S(\tilde{D}^{(i)}) \geq S(D)] \leq \alpha\right) \\ &= \mathbb{P}\left(S(D) > \inf\{s \in \mathfrak{S} : \sum_{\tilde{D} \in \mathfrak{D}: S(\tilde{D}) \leq s} w_{\mathfrak{D}}^q(\tilde{D}) \geq 1 - \alpha\}\right). \end{aligned}$$

Then, we have that

$$\mathbb{P}\left(S(D) > \inf\{s \in \mathfrak{S} : \sum_{\tilde{D} \in \mathfrak{D}: S(\tilde{D}) \leq s} w_{\mathfrak{D}}^q(\tilde{D}) \geq 1 - \alpha\} \mid \mathfrak{S}\right) \leq \alpha, \quad (11)$$

since 1. the infimum in equation (11) is the  $(1 - \alpha)^{\text{th}}$  conditional quantile of the weighted distribution over  $\mathfrak{S}$  and 2. as noted above,  $S(D)$  is conditionally a draw from this weighted distribution. Inequality (11) then also holds unconditionally, as desired.

## A.1 Connections to Harrison [2012]

### A.1.1 Equivalence of the tests

Harrison [2012] consider a problem in which the analyst has obtained a sample  $D \sim P$  from some (null) target distribution  $P$  and draws samples  $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)}$  *i.i.d.*  $Q$  from some known proposal distribution  $Q$ ,

<sup>12</sup>Just as above, we say that  $\tilde{D} \in \mathfrak{D}$  if  $\tilde{D}$  is an element of the list  $\mathfrak{D}$ .

independently from  $D$ . Then, assuming that the importance weights  $v(d) := \frac{dP}{dQ}(d)$  are known, it is shown that the p-value

$$\tilde{p}_*(X, \mathfrak{D}) = \sum_{i=0}^m \frac{v(\tilde{D}^{(i)})}{\sum_{j=0}^m v(\tilde{D}^{(j)})} \mathbf{1}[S(\tilde{D}^{(i)}) \geq S(\tilde{D}^{(0)})],$$

where, as usual,  $\tilde{D}^{(0)} := D$  and  $S$  is any test statistic is valid [Harrison, 2012, Theorem 2]. It is further shown in their work that a conditional version of this test, which first conditions on a statistic  $A(D)$  of the data, replaces  $P$  with its corresponding conditional distribution given  $A(D)$ , and allows  $Q$  to be any distribution depending on  $A(D)$  (but independent from  $D$  given  $A(D)$ ), is still valid (in fact, conditionally valid given  $A(D)$ ).

This is precisely the form of the special case of our test presented in Remark 9 with  $A(D)$  being a statistic of the observed data for which  $A(\tilde{D}^{(i)}) = A(D)$  and  $q(\cdot|\tilde{D}^{(i)}, A(\tilde{D}^{(i)})) = q(\cdot|A(D))$ , for all  $i \in [m]$ . Conditional on such a choice of  $A(D)$ ,  $D$  follows the distribution  $P_{A(D)} := f(\cdot|A(D))$ , and  $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)}$  are i.i.d. draws from  $Q_{A(D)} := q(\cdot|A(D))$  which are (jointly) independent from  $D$ . Thus, as discussed in Remark 9, using the fact that the conditional probability  $q(\tilde{D}^{(i)}|\tilde{D}^{(i)})$  may now be instead written as  $q(\tilde{D}^{(i)}|A(D))$ , the weight  $w_{\mathfrak{D}}^q(\tilde{D}^{(i)})$  involved in our p-value calculation is equal (under  $\mathcal{H}_0^\infty$ ) to

$$\frac{f(\tilde{D}^{(i)})/q(\tilde{D}^{(i)}|A(D))}{\sum_{j=0}^m f(\tilde{D}^{(j)})/q(\tilde{D}^{(j)}|A(D))} = \frac{f(\tilde{D}^{(i)}|A(D))/q(\tilde{D}^{(i)}|A(D))}{\sum_{j=0}^m f(\tilde{D}^{(j)}|A(D))/q(\tilde{D}^{(j)}|A(D))} = \frac{v_{A(D)}(\tilde{D}^{(i)})}{\sum_{j=0}^m v_{A(D)}(\tilde{D}^{(j)})},$$

where  $v_{A(D)}(d)$  is the conditional importance weight  $\frac{dP_{A(D)}}{dQ_{A(D)}}(d)$ . Thus our p-value  $p$  is precisely equal to the p-value  $\tilde{p}_*(X, \mathfrak{D})$  resulting from Harrison [2012]’s conditional importance sampling procedure.

### A.1.2 Practical considerations

Using Harrison [2012]’s conditional importance sampling procedure, we may characterize the resampling (i.e., proposal) distributions  $q$  discussed in Section 4 by simply specifying a statistic  $A(D)$  that is invariant under all resamples and for which  $q(\cdot|\tilde{D}^{(i)}, A(\tilde{D}^{(i)}))$  depends only on  $A(\tilde{D}^{(i)}) = A(D)$ . In particular, for non-stationarity testing, each of  $\text{uniform}_\pi$ ,  $\text{imitation}_\pi$ ,  $\text{re-imitation}_\pi$ , and  $\text{cond-imitation}_\pi$  depend only on the observed  $D$  through the unordered tuples  $(C_i, X_i, Y_i)$  (which are the same as in the original data in all resamples), and so our method can be viewed as applying Harrison [2012]’s procedure with  $A(D)$  equal to the *multiset*  $\{(C_1, X_1, Y_1), \dots, (C_T, X_T, Y_T)\}$ . For conditional independence testing,  $\text{imitation}_X$  depends only on the  $g$ -evaluations of the actions (which too are unchanged during resampling) and so in this case  $A(D) = ((C_1, g(X_1), Y_1), \dots, (C_T, g(X_T), Y_T))$ , whereas  $\text{uniform}_\pi + \text{imitation}_X$ ,  $\text{restricted-uniform}_\pi + \text{imitation}_X$ , and  $\text{combined}_{\pi, X}$  are also further agnostic to the order of the data (where we replace the actions with their  $g$ -evaluations) and do not change the multiset of unordered datapoints, so that for these three resampling distributions we have that  $A(D) = \{(C_1, g(X_1), Y_1), \dots, (C_T, g(X_T), Y_T)\}$ , where we once again use the brackets  $\{\cdot\}$  to denote multiset.

### A.1.3 Comparison of proof techniques

Finally, we comment on differences in proof techniques between the special case of our Theorem 2.1 discussed in Remark 9 and Theorem 2 of Harrison [2012].

The main difference is that our proof is via the proof of Theorem 2.1, which is more general than Theorem 2 of Harrison [2012]. In particular, the key to our proof given above is to *condition* on the multiset  $\mathfrak{S} := \{S(\tilde{D}) : \tilde{D} \in \mathfrak{D}\}$  (or equivalently, on the unordered elements of  $\mathfrak{D}$  prior to their  $S$ -evaluations), and to show, via Bayes’ theorem that, conditional on  $\mathfrak{S}$ , the event that  $p \leq \alpha$  is equal to the event that  $S(D)$  lies above its  $(1 - \alpha)$  conditional quantile, which occurs with probability no greater than  $\alpha$ .

On the other hand, the statement of Theorem 2 given in Harrison [2012] is essentially weaker than that of Theorem 2.1 and so, rather than by conditioning, their proof is via a change of measure argument. This proof technique, however, allows their Theorem 2 to be slightly more general than our statement in Remark 9 in

that they need not restrict attention to discrete distributions to avoid measure-theoretic pathologies involving (regular) conditional probability distributions, but rather need only assume that  $P \ll Q$ .

Restricting to the case in which the proposal distribution  $Q$  is independent from  $P$  (from which the conditional importance sampling result that is equivalent to our Theorem 2.1 is a corollary), Harrison [2012] show that

$$\frac{dP_{\mathfrak{D}}}{dP_{\mathfrak{U}}}(\mathfrak{d}) = \frac{(m+1) \frac{dP}{dQ}(d_0)}{\sum_{j=0}^m \frac{dP}{dQ}(d_j)} \quad (12)$$

where  $\mathfrak{U}$  is a uniform permutation of  $\mathfrak{D}$ ,  $P_{\mathfrak{D}}$  and  $P_{\mathfrak{U}}$  denote the laws of  $\mathfrak{D}$  and  $\mathfrak{U}$ , respectively, and  $\mathfrak{d} := (d_0, \dots, d_m)$ . Using equation (12) as well as the fact that  $\mathfrak{U}$ 's law is invariant under permutation, probabilities of events under  $\mathfrak{D}$ 's distribution (and in particular, those involving the p-value  $\tilde{p}_*(X, \mathfrak{D})$ , which is a deterministic function of  $\mathfrak{D}$ ) may be transformed into events involving (permutations of)  $\mathfrak{U}$  which, combined with a deterministic inequality that may be viewed as an unconditional version of inequality (11) (the argument for which is the same as it is simply a statement about weighted quantiles) then shows that the p-value is valid. We refer the reader to Harrison [2012] for details.

## B Conditional independence testing in a stationary MDP

In this section, we give a proof of Proposition 3.2 regarding the restricted permutation set in Environment 4:

*Proof.* Let  $\pi_i$  be any permutation in

$$\Pi_4 := \{\pi \in \Pi_{[T]} : Y_{\pi(t)} = C_{\pi(t+1)}, \forall t \in [T-1], \text{ and } C_{\pi(1)} = C_1\}.$$

Then

$$\begin{aligned} \hat{f}(\tilde{D}^{(i)}) &= \prod_{t=1}^T \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \\ &= \frac{\prod_{t=1}^T \mathbb{P}(Y_{\pi_i(t)} | C_{\pi_i(t)}, X_{\pi_i(t)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)})}{\prod_{t=1}^T \mathbb{P}(Y_t | C_t, X_t)} && \text{because of equation (6)} \\ &= \frac{\mathbb{P}(C_{\pi_i(1)}) \prod_{t=1}^T \mathbb{P}(C_{\pi_i(t+1)} | C_{\pi_i(t)}, X_{\pi_i(t)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)})}{\mathbb{P}(C_1) \prod_{t=1}^T \mathbb{P}(C_{t+1} | C_t, X_t)} && \text{since } \pi_i \in \Pi_4 \\ &= \frac{\mathbb{P}(C_{\pi_i(1)}) \prod_{t=1}^T \mathbb{P}(C_{\pi_i(t+1)} | C_{\pi_i(t)}, \tilde{X}_t^{(i)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)})}{\mathbb{P}(C_1) \prod_{t=1}^T \mathbb{P}(C_{t+1} | C_t, X_t)} && \text{by } \mathcal{H}_0^{\perp, g} \text{ and that } g(\tilde{X}_t^{(i)}) = g(X_{\pi_i(t)}) \\ &= \frac{\mathbb{P}(\tilde{C}_1^{(i)}) \prod_{t=1}^T \mathbb{P}(\tilde{C}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{X}_t^{(i)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} | \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)})}{\mathbb{P}(C_1) \prod_{t=1}^T \mathbb{P}(C_{t+1} | C_t, X_t)} && \text{by equation (8)} \\ &= \frac{1}{\mathbb{P}(C_1) \prod_{t=1}^T \mathbb{P}(C_{t+1} | C_t, X_t)} \cdot f(\tilde{D}^{(i)}), && \text{due to equation (1)} \end{aligned}$$

and so the proportionality hypothesis is satisfied with  $K = \frac{1}{\mathbb{P}(C_1) \prod_{t=1}^T \mathbb{P}(C_{t+1} | C_t, X_t)}$ .  $\square$

## C Inference in a general adaptive data collection process

In this section, we relax all structural assumptions made in Section 3 and assume that the data are gathered according to any adaptive data collection process (i.e., we make no assumptions on the joint distribution of  $(C_t, X_t, Y_t)_{t=1}^T$  other than that the sequence  $(X_t)_{t=1}^T$  is non-anticipating with respect to the filtration  $\mathcal{F}_t := \sigma(\mathcal{H}_{t-1} \cup \{C_t\})$ ). The hypothesis we are interested in testing is related to that discussed in Section 3.1, except here it involves *sequences* of actions. In fact, we consider  $T$  functions  $g_1, \dots, g_T$ , where  $g_t$  takes as

input the sequence of first  $t$  actions  $X_{1:t} := (X_1, \dots, X_t)$ ; the null hypothesis we are interested in testing is that of simultaneous independence, conditional on these  $g$ -evaluations:

$$\mathcal{H}_0^{\perp, g_1, \dots, g_T} : [C_t \perp\!\!\!\perp X_{1:t-1} \mid (C_{1:t-1}, g_{t-1}(X_{1:t-1}), Y_{1:t-1})] \text{ and } [Y_t \perp\!\!\!\perp X_{1:t} \mid (C_{1:t}, g_t(X_{1:t}), Y_{1:t-1})], \forall t \in [T],$$

where we define

$$C_{1:0} = X_{1:0} = g_0(X_{1:0}) = Y_{1:0} = \emptyset.$$

Informally,  $\mathcal{H}_0^{\perp, g_1, \dots, g_T}$  states that the actions  $X_t$  only influence future contexts and responses via their values filtered through the functions  $g_t$ . As noted in Remark 4, the setting in which all  $g_t$  are constant has been covered in prior work by Bojinov and Shephard [2019]; however, to the best of our knowledge, the case of non-constant  $g_t$  is novel. We now describe and prove the validity of the testing procedure for the null hypothesis  $\mathcal{H}_0^{\perp, g_1, \dots, g_T}$ .

**Proposition C.1.** *Suppose that the resampling distribution  $q$  randomizes only the action sequence  $X_{1:T}$  conditional on  $(g_t(X_{1:t}))_{t=1}^T$ . That is, in each resample  $\tilde{D}^{(i)}$ , we have that*

$$\left( \tilde{C}_t^{(i)}, g_t(\tilde{X}_{1:t}^{(i)}), \tilde{Y}_t^{(i)} \right) = (C_t, g_t(X_{1:t}), Y_t), \quad \forall t \in [T]. \quad (13)$$

Then the null hypothesis  $\mathcal{H}_0^{\perp, g_1, \dots, g_T}$  implies the proportionality hypothesis  $\mathcal{H}_0^\times$ .

*Proof.* We have that

$$\begin{aligned} & \hat{f}(\tilde{D}^{(i)}) \\ &= \prod_{t=1}^T \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} \mid \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \\ &= \frac{\prod_{t=1}^T \mathbb{P}(\tilde{Y}_t^{(i)} \mid \tilde{C}_{1:t}^{(i)}, g_t(\tilde{X}_{1:t}^{(i)}), \tilde{Y}_{1:t-1}^{(i)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} \mid \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \mathbb{P}(\tilde{C}_t^{(i)} \mid \tilde{C}_{1:t-1}^{(i)}, g_{t-1}(\tilde{X}_{1:t-1}^{(i)}), \tilde{Y}_{1:t-1}^{(i)})}{\prod_{t=1}^T \mathbb{P}(Y_t \mid C_{1:t}, g_t(X_{1:t}), Y_{1:t-1}) \mathbb{P}(C_t \mid C_{1:t-1}, g_{t-1}(X_{1:t-1}), Y_{1:t-1})} && \text{equation (13)} \\ &= \frac{\prod_{t=1}^T \mathbb{P}(\tilde{Y}_t^{(i)} \mid \tilde{C}_{1:t}^{(i)}, g_t(\tilde{X}_{1:t}^{(i)}), \tilde{Y}_{1:t-1}^{(i)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} \mid \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \mathbb{P}(\tilde{C}_t^{(i)} \mid \tilde{C}_{1:t-1}^{(i)}, g_{t-1}(\tilde{X}_{1:t-1}^{(i)}), \tilde{Y}_{1:t-1}^{(i)}, \tilde{X}_{1:t-1}^{(i)})}{\prod_{t=1}^T \mathbb{P}(Y_t \mid C_{1:t}, g_t(X_{1:t}), Y_{1:t-1}) \mathbb{P}(C_t \mid C_{1:t-1}, g_{t-1}(X_{1:t-1}), Y_{1:t-1})} && \text{due to } \mathcal{H}_0^{\perp, g_1, \dots, g_T} \\ &= \frac{\prod_{t=1}^T \mathbb{P}(\tilde{Y}_t^{(i)} \mid \tilde{C}_{1:t}^{(i)}, \tilde{X}_{1:t}^{(i)}, \tilde{Y}_{1:t-1}^{(i)}) \mathbb{P}_{\mathcal{A}}(\tilde{X}_t^{(i)} \mid \tilde{C}_t^{(i)}, \tilde{H}_{t-1}^{(i)}) \mathbb{P}(\tilde{C}_t^{(i)} \mid \tilde{C}_{1:t-1}^{(i)}, \tilde{X}_{1:t-1}^{(i)}, Y_{1:t-1}^{(i)})}{\prod_{t=1}^T \mathbb{P}(Y_t \mid C_{1:t}, g_t(X_{1:t}), Y_{1:t-1}) \mathbb{P}(C_t \mid C_{1:t-1}, g_{t-1}(X_{1:t-1}), Y_{1:t-1})} \\ &= \frac{1}{\prod_{t=1}^T \mathbb{P}(Y_t \mid C_{1:t}, g_t(X_{1:t}), Y_{1:t-1}) \mathbb{P}(C_t \mid C_{1:t-1}, g_{t-1}(X_{1:t-1}), Y_{1:t-1})} \cdot f(\tilde{D}^{(i)}), \end{aligned}$$

which satisfies the proportionality hypothesis with  $K = \frac{1}{\prod_{t=1}^T \mathbb{P}(Y_t \mid C_{1:t}, g_t(X_{1:t}), Y_{1:t-1}) \mathbb{P}(C_t \mid C_{1:t-1}, g_{t-1}(X_{1:t-1}), Y_{1:t-1})}$ .  $\square$

One concrete setting in which the above randomization procedure could be used to test against the hypothesis  $\mathcal{H}_0^{\perp, g_1, \dots, g_T}$  is in a completely general mobile health setting in which no environmental assumptions are made and there are only two treatments: a control and an experimental treatment. A hypothesis which may be of interest to test is that, conditional on the prior sequence of contexts and responses, each of the next response and context is independent of the action sequence given the *number* of experimental treatments taken up until that time. In such a setting, one could use our framework under the above randomization scheme to test the null hypothesis  $\mathcal{H}_0^{\perp, g_1, \dots, g_T}$  where  $g_t(X_{1:t})$  is an indicator that the experimental treatment appears at least half of the time in the treatment sequence  $X_{1:t}$ ; i.e.,  $g_t(X_{1:t}) = \mathbf{1}[\sum_{s=1}^t X_s \geq t/2]$  where  $X = 1$  denotes the experimental treatment and  $X = 0$  denotes the control.

## D Inference in scale families

In this section, we describe how the test discussed in Section 3.3.1 can also be used for inference in semiparametric scale families.

Instead of assuming that each reward distribution is a member of the same location family, it may be more natural in some situations to assume a semiparametric scale family. As such, letting  $\theta_x$  now denote action  $x$ 's scale parameter (i.e., we assume that  $Y_t \mid X_t = x \sim h_0(y/\theta_x)$ <sup>13</sup>), we may instead test against the null hypothesis  $\mathcal{H}^{\text{Scale}, \delta, x, x'}$  that  $\frac{\theta_{x'}}{\theta_x} = \delta$ . To do so, we simply revise the reward-modification trick discussed in Section 3.3.1 for the location family to instead replace  $Y_t$  with  $Y_t \cdot \delta \cdot \mathbf{1}[X_t = x]$  and again perform the usual test for conditional independence from Section 3.1 using  $g(X_t) = \begin{cases} \{x, x'\} & \text{if } X_t \in \{x, x'\} \\ X_t & \text{otherwise} \end{cases}$ .

Similar to as discussed in Section 3.3.1, the above test can be inverted to construct (simultaneous) confidence intervals.

## E Non-conditionally-i.i.d. resampling and sharing samples

In this section, we discuss how to handle non-conditionally-i.i.d. resampling procedures and how this more general procedure can be used to share samples between different values  $y \in \mathcal{Y}$  in the construction of conformal prediction regions.

### E.1 Non-conditionally-i.i.d. resampling

In this section, we describe a more general version of the weighted MC randomization test of Algorithm 1 that can be used in the setting of non-conditionally-i.i.d. resampling, as discussed in Remark 8. In particular, the Remark states that, letting  $\tilde{q}$  denote the joint conditional distribution of the resamples  $(\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)})$ , one may redefine the weights to be

$$w_{\mathfrak{D}}^{\tilde{q}, \Sigma}(\tilde{D}^{(i)}) = \frac{\hat{f}(\tilde{D}^{(i)}) \sum_{\pi \in \Sigma: \pi(0)=i} \tilde{q}((\mathfrak{D}_{-0})^\pi \mid \tilde{D}^{(i)})}{\sum_{j \in [0:m]} \hat{f}(\tilde{D}^{(j)}) \sum_{\pi' \in \Sigma: \pi'(0)=j} \tilde{q}((\mathfrak{D}_{-0})^{\pi'} \mid \tilde{D}^{(j)})}, \quad (14)$$

and the p-value

$$p := \sum_{i=0}^m w_{\mathfrak{D}}^{\tilde{q}, \Sigma}(\tilde{D}^{(i)}) \mathbf{1}[S(\tilde{D}^{(i)}) \geq S(D)]$$

will be valid.

Recall that the key idea behind Algorithm 1—and the proof of Theorem 2.1—was to condition on the event  $(\tilde{D}^{(0)}, \dots, \tilde{D}^{(m)}) \in \{(d_{\pi(0)}, \dots, d_{\pi(m)}) : \pi \in \Pi_{[0:m]}\}$  for some list  $\mathfrak{d} = (d_0, \dots, d_m)$  and to apply Bayes' theorem as well as the proportionality hypothesis to derive the weight formula (equation (3)). The entire permutation set  $\Pi_{[0:m]}$  need not, however, be considered. Our generalization involves using any arbitrary subset of permutations  $\Sigma$ ; when  $\Sigma$  is small, our method is computationally tractable for non-conditionally-i.i.d. resampling schemes.

More specifically, define

$$\text{pseudo-orb}_{\Sigma}(\mathfrak{d}) := \{(d_{\pi(0)}, \dots, d_{\pi(m)}) : \pi \in \Sigma\}$$

to be the  $\Sigma$ -pseudo-orbit<sup>14</sup> of the list  $\mathfrak{d}$ . Then, by conditioning on  $\text{pseudo-orb}_{\Sigma}(\mathfrak{D})$ , we will be able to show that the weighting given by equation (14) will be valid. We first prove a result similar to Theorem 2.1 in the case of distinct data  $D$  and resamples  $\tilde{D}^{(1)}, \dots, \tilde{D}^{(m)}$  using exactly the same proof strategy:

<sup>13</sup>Recall that we only consider discrete random variables, per Remark 1, and hence no Jacobian is required.

<sup>14</sup>We use the prefix “pseudo” since  $\Sigma$  need not be a group.

**Lemma E.1.** Let  $\Sigma$  be any subset of the full set of permutations on  $[0 : m]$ ,  $\Pi_{[0:m]}$ . Define  $\Sigma_i$  to be  $\{\pi \in \Sigma : \pi(0) = i\}$  and assume that  $\tilde{q}$  is such that  $\tilde{D}^{(0)}, \dots, \tilde{D}^{(m)}$  are all distinct. Then, with weights

$$w_{\mathfrak{D},i}^{\tilde{q},\Sigma}(\tilde{D}^{(i)}) := \frac{\hat{f}(\tilde{D}^{(i)}) \sum_{\pi \in \Sigma_i} \tilde{q}((\mathfrak{D}_{-0})^\pi | \tilde{D}^{(i)})}{\sum_{j=0}^m \hat{f}(\tilde{D}^{(j)}) \sum_{\pi' \in \Sigma_j} \tilde{q}((\mathfrak{D}_{-0})^{\pi'} | \tilde{D}^{(j)})},$$

we have that

$$p := \sum_{i=0}^m w_{\mathfrak{D},i}^{\tilde{q},\Sigma}(\tilde{D}^{(i)}) \mathbf{1}[S(\tilde{D}^{(i)}) \geq S(D)]$$

stochastically dominates the uniform distribution under  $\mathcal{H}_0^\infty$ .

*Proof sketch.* We give a sketch of the proof as much of it is the same as that of Theorem 2.1. Using Bayes' theorem, we have that

$$\begin{aligned} & \mathbb{P}(D = d_i | (\tilde{D}^{(0)}, \dots, \tilde{D}^{(m)}) \in \text{pseudo-orb}_\Sigma(\mathfrak{D})) \\ &= \frac{\mathbb{P}((\tilde{D}^{(0)}, \dots, \tilde{D}^{(m)}) \in \text{pseudo-orb}_\Sigma(\mathfrak{D}) | D = d_i) f(d_i)}{\sum_{j=0}^m \mathbb{P}((\tilde{D}^{(0)}, \dots, \tilde{D}^{(m)}) \in \text{pseudo-orb}_\Sigma(\mathfrak{D}) | D = d_j) f(d_j)} \\ &= \frac{f(d_i) \sum_{\pi \in \Sigma_i} \mathbb{P}((\mathfrak{D}_{-0}) = (d_{\pi(1)}, \dots, d_{\pi(m)}) | D = d_i)}{\sum_{j=0}^m f(d_j) \sum_{\pi' \in \Sigma_j} \mathbb{P}((\mathfrak{D}_{-0}) = (d_{\pi'(1)}, \dots, d_{\pi'(m)}) | D = d_j)}. \end{aligned}$$

Thus, conditional on  $\text{pseudo-orb}_\Sigma(\mathfrak{D})$ ,  $S(D)$ 's distribution is indeed the weighted distribution on the multiset  $\{S(\tilde{D}) : \tilde{D} \in \mathfrak{D}\}$  with weight of the  $i^{\text{th}}$  element equal to  $w_{\mathfrak{D},i}^{\tilde{q},\Sigma}(\tilde{D}^{(i)})$  under  $\mathcal{H}_0^\infty$ . The rest of the proof, in this setting of distinct resamples, follows in precisely the same manner as that of Theorem 2.1.  $\square$

We now extend to the case of repeated samples via a coupling argument. Define  $\mathcal{D}' := \mathcal{D} \times [0 : m]$  so that, for each point  $d$  in  $\mathcal{D}$  it has  $m + 1$  "clones" in  $\mathcal{D}'$  given by  $(d, 0), \dots, (d, m)$ . We also define a density  $f'$  on  $\mathcal{D}'$  given by  $f'((d, i)) := (m + 1)^{-1} f(d)$  for all  $d \in \mathcal{D}, i \in [0 : m]$ . Lastly, we extend the test statistic  $S$  to a test statistic  $S'$  on  $\mathcal{D}'$  which also takes each clone to the value of its original:  $S'((d, i)) = S(d)$  for all  $d \in \mathcal{D}, i \in [0 : m]$ .

We will couple  $p$  to a random variable  $p'$  obtained by a procedure on the space  $\mathcal{D}'$  which draws distinct elements  $\tilde{D}'^{(i)}$ , to be defined below. We call the original process (that may draw repeated elements and acts only on  $\mathcal{D}$ )  $\mathcal{P}$  and the coupled process  $\mathcal{P}'$ . The observed dataset that the coupled process  $\mathcal{P}'$  will see and use is a cloned version  $(D, j)$  of the true observed data  $D$ , where  $j$  is sampled uniformly at random from the set  $[0 : m]$ ; importantly  $\mathcal{P}'$  treats this cloned dataset  $D' := (D, j)$  as though it were the observed dataset. Finally, we use  $\tilde{q}'$  to denote the conditional resampling distribution of  $\mathcal{P}'$ , induced by  $\tilde{q}$ . That is, for a list  $\mathfrak{D}'_{-0} := (\tilde{D}'^{(1)}, \dots, \tilde{D}'^{(m)})$  of resamples in  $\mathcal{D}'$ ,  $\tilde{q}'(\mathfrak{D}'_{-0} | D')$  denotes the joint conditional probability of  $\mathcal{P}'$  having obtained the list of resamples  $\mathfrak{D}'_{-0}$  given that it observed the dataset  $D'$ . Algorithm 2 describes how  $\mathcal{P}'$  is defined with respect to  $\mathcal{P}$ .

At a high level, the coupled process  $\mathcal{P}'$  draws distinct elements by sampling an element's clone number  $j$  uniformly at random from all remaining yet-unselected values; using these distinct resamples, it then calculates a p-value in precisely the same way as Algorithm 1 with weights  $w_{\mathfrak{D}',i}^{\tilde{q},\Sigma}(\tilde{D}'^{(i)})$ .

Since the draws of  $\mathcal{P}'$  are guaranteed to be distinct, Lemma E.1 will guarantee  $p'$ 's stochastic domination of the uniform distribution as long as  $f'$  is the true density of  $D'$ , the observed data for  $\mathcal{P}'$  (since  $\tilde{q}'$  is the conditional resampling distribution and because  $f'$  is trivially proportional to itself, thereby satisfying the proportionality hypothesis in this coupled model). For  $d' \in \mathcal{D}'$ , let  $C(d')$  denote the second component of  $d'$  (i.e., the *clone number* of  $d'$ ) and  $d$  be the first component (i.e., the uncloned version of  $d'$  in  $\mathcal{D}$ ); then we see that  $f'$  is  $D'$ 's true density:

$$\begin{aligned} \mathbb{P}(D' = d') &= \mathbb{P}(C(D') = C(d') | D = d) \mathbb{P}(D = d) \\ &= \mathbb{P}(C(D') = C(d')) \mathbb{P}(D = d) && \text{since the clone number } j \text{ is sampled independently of } D \\ &= (m + 1)^{-1} f(d) && \text{because } j \sim \text{Unif}([0 : m]) \text{ to clone } D \\ &= f'(d') \end{aligned}$$

---

**Algorithm 2:** Coupling of  $\mathcal{P}'$  to  $\mathcal{P}$ 


---

**Input:**  $\mathcal{P}'$ 's observed dataset  $D' := (D, j)$  where  $j \sim \text{Unif}([0 : m])$

- 1  $\tilde{D}'^{(0)} \leftarrow D'$
- 2  $\mathfrak{D}'^{(0)} \leftarrow (\tilde{D}'^{(0)})$
- 3 **for**  $i = 1, \dots, m$  **do**
- 4      $\tilde{D}^{(i)} \leftarrow i^{\text{th}}$  element of  $\mathfrak{D}$  in the process  $\mathcal{P}$
- 5      $j \leftarrow \text{Unif}(\{j' \in [0 : m] : (\tilde{D}^{(i)}, j') \notin \mathfrak{D}'^{(i-1)}\})$
- 6      $\tilde{D}'^{(i)} \leftarrow (\tilde{D}^{(i)}, j)$
- 7      $\mathfrak{D}'^{(i)} \leftarrow$  concatenation of  $\mathfrak{D}'^{(i-1)}$  with  $\tilde{D}'^{(i)}$
- 8  $\mathfrak{D}' \leftarrow \mathfrak{D}'^{(m)}$

**Output:**

$$p' := \frac{\sum_{i=0}^m \mathbf{1}[S'(\tilde{D}'^{(i)}) \geq S'(D')] f'(\tilde{D}'^{(i)}) \sum_{\pi \in \Sigma_i} \tilde{q}'((\mathfrak{D}'_{-0})^\pi | \tilde{D}'^{(i)})}{\sum_{j=0}^m f'(\tilde{D}'^{(j)}) \sum_{\pi' \in \Sigma_j} \tilde{q}'((\mathfrak{D}'_{-0})^{\pi'} | \tilde{D}'^{(j)})},$$


---

as desired.

We complete the coupling by showing that

$$\tilde{q}'((\mathfrak{d}'_{-0})^\pi | d'_i) = \kappa \tilde{q}((\mathfrak{d}_{-0})^\pi | d_i), \forall i \in [0 : m], \pi \in \Sigma_i \text{ for some } \kappa > 0 \text{ depending on neither } i \text{ nor } \pi, \quad (15)$$

where  $\mathfrak{d}$  is some list  $(d_0, \dots, d_m)$  of elements in  $\mathcal{D}$  and  $\mathfrak{d}' := (d'_0, \dots, d'_m)$  is a list of elements of  $\mathcal{D}'$  comprising distinct clones of the  $d_i$  (i.e., the first component of  $d'_i$  is equal to  $d_i$ , and each element of  $\mathfrak{d}'$  is distinct),  $(\mathfrak{d}'_{-0})^\pi := (d'_{\pi(1)}, \dots, d'_{\pi(m)})$ , and  $(\mathfrak{d}_{-0})^\pi := (d_{\pi(1)}, \dots, d_{\pi(m)})$ . To see why this is true, note that if  $\mathcal{P}'$  observes  $d'_i$  as the original dataset, then  $\mathcal{P}$  must have seen its uncloned version:  $d_i$ . Thus, applying the cloning number function  $C$  defined above to  $(\mathfrak{d}'_{-0})^\pi$  elementwise, we have that

$$\begin{aligned} \tilde{q}'((\mathfrak{d}'_{-0})^\pi | d'_i) &= \mathbb{P}(\mathfrak{D}'_{-0} = (\mathfrak{d}'_{-0})^\pi | D' = d'_i) \\ &= \mathbb{P}(C(\mathfrak{D}'_{-0}) = C((\mathfrak{d}'_{-0})^\pi), \mathfrak{D}_{-0} = (\mathfrak{d}_{-0})^\pi | C(D') = C(d'_i), D = d_i) \\ &= \mathbb{P}(C(\mathfrak{D}'_{-0}) = C((\mathfrak{d}'_{-0})^\pi) | \mathfrak{D}_{-0} = (\mathfrak{d}_{-0})^\pi, D' = d'_i) \mathbb{P}(\mathfrak{D}_{-0} = (\mathfrak{d}_{-0})^\pi | D = d_i) \quad \text{as } \mathfrak{D}_{-0} \perp\!\!\!\perp C(D') | D. \end{aligned}$$

To see that this is proportional to  $\tilde{q}((\mathfrak{d}_{-0})^\pi | d_i)$  (which is precisely the second term in the last line above), let  $m_d((\mathfrak{d}_{-0})^\pi)$  denote the number of elements in  $(\mathfrak{d}_{-0})^\pi$  with first component equal to  $d$ , and observe that we may write the first term in the last line above as

$$\begin{aligned} &\mathbb{P}(C(\mathfrak{D}'_{-0}) = C((\mathfrak{d}'_{-0})^\pi) | \mathfrak{D}_{-0} = (\mathfrak{d}_{-0})^\pi, D' = d'_i) \\ &= \left( \prod_{\substack{\text{distinct } d \in \mathfrak{d} \\ \text{not equal to } d_i}} \frac{1}{\prod_{k=1}^{m_d((\mathfrak{d}'_{-0})^\pi)} (m+1-k+1)} \right) \cdot \frac{1}{\prod_{k=1}^{m_{d_i}((\mathfrak{d}'_{-0})^\pi)} (m+1-k)} \\ &= \left( \prod_{\substack{\text{distinct } d \in \mathfrak{d} \\ \text{not equal to } d_i}} \frac{(m+1-m_d((\mathfrak{d}'_{-0})^\pi))!}{(m+1)!} \right) \cdot \frac{(m-m_{d_i}((\mathfrak{d}'_{-0})^\pi))!}{m!}. \end{aligned}$$

Now note that  $m_d((\mathfrak{d}'_{-0})^\pi)$  does not depend on the choice of  $\pi \in \Sigma_i$  and so we need only consider the permutation  $\pi_i^* \in \Sigma_i$  given by

$$\pi_i^* : k \mapsto \begin{cases} i & \text{if } k = 0 \\ k-1 & \text{if } 0 < k \leq i \\ k & \text{if } i < k \leq m \end{cases}$$

so that, for any  $\pi \in \Sigma_i$ ,  $m_d((\mathfrak{d}'_{-0})^\pi) = m_d((\mathfrak{d}'_{-0})^{\pi_i^*}) = m_d(\mathfrak{d}'_{-i})$ , where  $\mathfrak{d}'_{-i} := (d'_0, \dots, d'_{i-1}, d'_{i+1}, \dots, d'_m)$ . As such, combining this fact with what has been shown above, gives that

$$\mathbb{P}(C(\mathfrak{D}'_{-0}) = C((\mathfrak{d}'_{-0})^{\pi_i^*}) | \mathfrak{D}_{-0} = (\mathfrak{d}_{-0})^{\pi_i^*}, D' = d'_i) = \left( \prod_{\substack{\text{distinct } d \in \mathfrak{d} \\ \text{not equal to } d_i}} \frac{(m+1-m_d(\mathfrak{d}'_{-i}))!}{(m+1)!} \right) \cdot \frac{(m-m_{d_i}(\mathfrak{d}'_{-i}))!}{m!}.$$

Note that the right-hand side above depends on  $d_i$  only through its value, not its subscript, and hence the equation above takes the same value for any  $j \neq i$  for which  $d_j = d_i$ . In the case of distinct  $i, j$  with  $d_i \neq d_j$ , the above gives that

$$\begin{aligned}
& \mathbb{P}(C(\mathfrak{D}'_{-0}) = C((\mathfrak{d}'_{-0})^{\pi_i^*}) | \mathfrak{D}_{-0} = (\mathfrak{d}_{-0})^{\pi_i^*}, D' = d'_i) \\
&= \left( \prod_{\text{distinct } d \in \mathfrak{d} \text{ equal to neither } d_i \text{ nor } d_j} \frac{(m+1 - m_d(\mathfrak{d}'_{-i}))!}{(m+1)!} \right) \cdot \frac{(m+1 - m_{d_j}(\mathfrak{d}'_{-i}))!}{(m+1)!} \cdot \frac{(m - m_{d_i}(\mathfrak{d}'_{-i}))!}{m!} \\
&= \left( \prod_{\text{distinct } d \in \mathfrak{d} \text{ equal to neither } d_i \text{ nor } d_j} \frac{(m+1 - m_d(\mathfrak{d}'_{-j}))!}{(m+1)!} \right) \cdot \frac{(m+1 - m_{d_j}(\mathfrak{d}'_{-i}))!}{(m+1)!} \cdot \frac{(m - m_{d_i}(\mathfrak{d}'_{-i}))!}{m!} \\
&= \left( \prod_{\text{distinct } d \in \mathfrak{d} \text{ equal to neither } d_i \text{ nor } d_j} \frac{(m+1 - m_d(\mathfrak{d}'_{-j}))!}{(m+1)!} \right) \cdot \frac{(m - m_{d_j}(\mathfrak{d}'_{-j}))!}{(m+1)!} \cdot \frac{(m+1 - m_{d_i}(\mathfrak{d}'_{-j}))!}{m!} \\
&= \left( \prod_{\text{distinct } d \in \mathfrak{d} \text{ not equal to } d_j} \frac{(m+1 - m_d(\mathfrak{d}'_{-j}))!}{(m+1)!} \right) \cdot \frac{(m - m_{d_j}(\mathfrak{d}'_{-j}))!}{m!} \\
&= \mathbb{P}(C(\mathfrak{D}'_{-0}) = C(\mathfrak{d}'_{-j}) | \mathfrak{D}_{-0} = \mathfrak{d}_{-j}, D' = d'_j) \\
&= \mathbb{P}(C(\mathfrak{D}'_{-0}) = C((\mathfrak{d}'_{-0})^{\pi_j^*}) | \mathfrak{D}_{-0} = (\mathfrak{d}_{-0})^{\pi_j^*}, D' = d'_j),
\end{aligned}$$

as desired, where the third equality follows from the fact that  $m_{d_j}(\mathfrak{d}'_{-j}) + 1 = m_{d_j}(\mathfrak{d}'_{-i})$  for any pair  $i, j$  with  $d_i \neq d_j$ . By the same argument made for  $\pi_i^*$ , the last line of the above remains true if we replace  $\pi_j^*$  with any  $\pi \in \Sigma_j$  and thus, we have shown that not only does  $\mathbb{P}(C(\mathfrak{D}'_{-0}) = C((\mathfrak{d}'_{-0})^\pi) | \mathfrak{D}_{-0} = (\mathfrak{d}_{-0})^\pi, D' = d'_i)$  not depend on the choice of  $\pi \in \Sigma_i$ , for any given  $i$ , but it also does not depend on  $i \in [0 : m]$ , and hence we have the desired proportionality of equation (15).

Finally, since  $f'(\tilde{D}'^{(i)}) = (m+1)^{-1} f(\tilde{D}^{(i)})$  by definition, and  $\hat{f}(\tilde{D}^{(i)}) = K f(\tilde{D}^{(i)}) \forall i \in [0 : m]$  for some  $K$  not depending on  $i$  under the proportionality hypothesis, we have that  $\hat{f}(\tilde{D}^{(i)}) = K(m+1) f'(\tilde{D}'^{(i)})$ ,  $\forall i \in [0 : m]$  and hence the two are proportional. As such, we have that

$$\begin{aligned}
& \frac{\sum_{i=0}^m \mathbf{1}[S'(\tilde{D}'^{(i)}) \geq S'(D')] f'(\tilde{D}'^{(i)}) \sum_{\pi \in \Sigma_i} \tilde{q}'((\mathfrak{D}'_{-0})^\pi | \tilde{D}'^{(i)})}{\sum_{j=0}^m f'(\tilde{D}'^{(j)}) \sum_{\pi' \in \Sigma_j} \tilde{q}'((\mathfrak{D}'_{-0})^{\pi'} | \tilde{D}'^{(j)})} \\
&= \frac{\sum_{i=0}^m \mathbf{1}[S(\tilde{D}^{(i)}) \geq S(D)] \hat{f}(\tilde{D}^{(i)}) \sum_{\pi \in \Sigma_i} \tilde{q}((\mathfrak{D}_{-0})^\pi | \tilde{D}^{(i)})}{\sum_{j=0}^m \hat{f}(\tilde{D}^{(j)}) \sum_{\pi' \in \Sigma_j} \tilde{q}((\mathfrak{D}_{-0})^{\pi'} | \tilde{D}^{(j)})}
\end{aligned}$$

and so  $p' = p$ . This combined with the fact shown above that  $p'$  is a valid p-value implies the desired validity of  $p$  under  $\mathcal{H}_0^\infty$ .

## E.2 Sharing samples between $y \in \mathcal{Y}$

We now show how the above framework can be used to share samples between different grid values  $y \in \mathcal{Y}$  as discussed in Remark 7. We describe this sample sharing in the absence of the rounding procedure of Remark 7 which considers only a small grid  $\mathcal{Y}' \subseteq \mathcal{Y}$  (i.e., we will obtain samples for *each*  $y$  in the full support  $\mathcal{Y}$  and share these samples amongst all other elements of  $\mathcal{Y}$ ), but both of these methods can be combined to construct conformal prediction regions even more efficiently.

In more detail, recall that the naive interval construction described in Section 3.3 runs  $|\mathcal{Y}|$  *independent* non-stationarity tests, using our weighted MC randomization testing framework for conditionally i.i.d. re-samples using  $q$ , at each  $y \in \mathcal{Y}$ . To formalize this notion, set  $(D_{[1:T-1]}, C_T, X_T)$  to denote the entire dataset except the last response and let  $U_y$  denote the exogenous randomness used by the weighted MC randomization test when determining if  $y$  is in the acceptance region. Then we can define an acceptance

function  $\varphi(U_y, (D_{[1:T-1]}, C_T, X_T), y)$  which is 1 (i.e., accepts) if the test, using  $U_y$  as exogenous randomness, states that  $y$  is in the acceptance region upon seeing  $(D_{[1:T-1]}, C_T, X_T)$ , and is otherwise 0 (i.e., rejects). In the naive gridding procedure, the  $U_y$  are all jointly independent and are distributed identically to the exogeneous randomness  $U$  that is used in the usual non-stationarity test. As such, it is clear that  $\mathbb{E}[\varphi(U_{Y_T}, (D_{[1:T-1]}, C_T, X_T), Y_T)] \geq 1 - \alpha$  due to the validity of the non-stationarity test proved in Section 3.2 and so the corresponding prediction region attains coverage at least that of the nominal rate. The  $U_y$ , however, need not be independent. Instead, for example, if we have  $U_y = U$  for all  $y \in \mathcal{Y}$ , where, again,  $U$  is the independent exogenous randomness used by the non-stationarity test, then our prediction region would once again control miscoverage at the nominal rate. More generally, as long as each  $U_y$  (and  $U$ ) is generated independently from  $D$  and we have the following equality in distribution:

$$(U_{Y_T}, Y_T) \stackrel{d}{=} (U, Y_T),$$

the resulting prediction regions will be valid. The process of sharing samples between  $y \in \mathcal{Y}$  is based upon this idea.

While described as resampling datasets  $\tilde{D}^{(i)}$  conditional on  $D$  in Section 3.2, the non-stationarity tests we describe really sample permutations  $\pi_i$ , conditionally on  $D$ , and then set  $\tilde{D}^{(i)}$  equal to  $D^{\pi_i} := ((C_{\pi_i(1)}, X_{\pi_i(1)}, Y_{\pi_i(1)}), \dots, (C_{\pi_i(T)}, X_{\pi_i(T)}, Y_{\pi_i(T)}))$ ; we use  $q_\Pi$  to denote the corresponding (to  $q$ ) resampling distribution over permutations. In the naive interval construction procedure which uses independent exogenous randomness, let  $\mathfrak{D}_\Pi^y$  denote the (multi)set of permutations sampled for determining if  $y$  is in the acceptance region. We claim that all these samples can be shared so that each  $y \in \mathcal{Y}$  instead uses the entire set of shared permutations  $\mathfrak{D}_\Pi := \bigcup_{y \in \mathcal{Y}} \mathfrak{D}_\Pi^y$  as opposed to just  $\mathfrak{D}_\Pi^y$ :

**Proposition E.1.** *Let  $q$  be any conditionally i.i.d. resampling distribution and let  $q_\Pi$  denote the corresponding resampling distribution over permutations. Furthermore define, for each  $y \in \mathcal{Y}$ , a conditional distribution  $q'_{\Pi,y}$  over permutations which draws its samples conditionally independently from  $Y_T$  given  $(D_{[1:T-1]}, C_T, X_T)$  and is defined by*

$$q'_{\Pi,y}(\cdot | ((C_1, X_1, Y_1), \dots, (C_T, X_T, Y_T))) = q_\Pi(\cdot | ((C_1, X_1, Y_1), \dots, (C_T, X_T, y))).$$

To construct a prediction region, suppose that, for each  $y \in \mathcal{Y}$ , a sampled (multi)set  $\mathfrak{D}_\Pi^y = \{\pi_{1,y}, \dots, \pi_{m,y}\}$  of permutations is generated by sampling each permutation conditionally i.i.d. from  $q'_{\Pi,y}(\cdot | D)$ <sup>15</sup>, but the full (multi)set  $\mathfrak{D}_\Pi$  of  $m|\mathcal{Y}|$  permutations is used to determine  $y$ 's membership in the acceptance region. Specifically, defining  $D_y := ((C_1, X_1, Y_1), \dots, (C_T, X_T, y))$ , writing  $\mathcal{Y} = \{y_1, \dots, y_{|\mathcal{Y}|}\}$ , and defining  $\Lambda := \{(0, 0)\} \cup ([m] \times [|\mathcal{Y}|])$ , we construct an acceptance region  $\mathfrak{A}$  from  $\mathfrak{D}_\Pi$  via the rule:

$$y \in \mathfrak{A} \iff \sum_{(i,j) \in \Lambda} w_{\mathfrak{D}_\Pi, i, y_j}^{q,y} ((D_y)^{\pi_{i,y_j}}) \mathbf{1}[S((D_y)^{\pi_{i,y_j}}) \geq S(D_y)] > \alpha, \quad (16)$$

where

$$w_{\mathfrak{D}_\Pi, i, y_j}^{q,y} ((D_y)^{\pi_{i,y_j}}) := \frac{\hat{f}((D_y)^{\pi_{i,y_j}}) q'_{\Pi, y_j}(\pi_{i,y_j}^{-1} | (D_y)^{\pi_{i,y_j}}) \prod_{(\tilde{i}, \tilde{j}) \in \Lambda \setminus \{(i,j)\}} q'_{\Pi, y_{\tilde{j}}}(\pi_{\tilde{i}, y_{\tilde{j}}} \circ \pi_{i,y_j}^{-1} | (D_y)^{\pi_{i,y_j}})}{\sum_{(i', j') \in \Lambda} \hat{f}((D_y)^{\pi_{i', y_{j'}}}) q'_{\Pi, y_{j'}}(\pi_{i', y_{j'}}^{-1} | (D_y)^{\pi_{i', y_{j'}}}) \prod_{(\tilde{i}', \tilde{j}') \in \Lambda \setminus \{(i', j')\}} q'_{\Pi, y_{\tilde{j}'}}(\pi_{\tilde{i}', y_{\tilde{j}'}} \circ \pi_{i', y_{j'}}^{-1} | (D_y)^{\pi_{i', y_{j'}}})},$$

and we take  $\pi_{0, y_0}$  to simply be the identity permutation.

Then the resultant prediction region  $\mathfrak{A}$  controls miscoverage at the nominal rate under the proportionality hypothesis  $\mathcal{H}_0^\propto$ .

<sup>15</sup>Since  $q'_{\Pi,y}$  does not depend on  $Y_T$ , we can indeed sample from this distribution during prediction region construction wherein  $Y_T$  is unobserved

Before giving a proof we make a brief computational/procedural note that, similar to equation (9) of Remark 9, when  $q(\cdot|(D_y)^{\pi_{i,y_j}})$  does not depend on  $(i,j) \in \Lambda$  for any  $y \in \mathcal{Y}$ , we have that

$$\begin{aligned}
w_{\mathfrak{D}_{\Pi,i,y_j}}^{q,y}((D_y)^{\pi_{i,y_j}}) &= \frac{\hat{f}((D_y)^{\pi_{i,y_j}})q'_{\Pi,y_j}(\pi_{i,y_j}^{-1}|(D_y)^{\pi_{i,y_j}}) \prod_{(\tilde{i},\tilde{j}) \in \Lambda \setminus \{(i,j)\}} q'_{\Pi,y_{\tilde{j}}}(\pi_{\tilde{i},y_{\tilde{j}}} \circ \pi_{i,y_j}^{-1}|(D_y)^{\pi_{i,y_j}})}{\sum_{(i',j') \in \Lambda} \hat{f}((D_y)^{\pi_{i',y_{j'}}})q'_{\Pi,y_{j'}}(\pi_{i',y_{j'}}^{-1}|(D_y)^{\pi_{i',y_{j'}}}) \prod_{(\tilde{i}',\tilde{j}') \in \Lambda \setminus \{(i',j')\}} q'_{\Pi,y_{\tilde{j}'}}(\pi_{\tilde{i}',y_{\tilde{j}'}} \circ \pi_{i',y_{j'}}^{-1}|(D_y)^{\pi_{i',y_{j'}}})} \\
&= \frac{\hat{f}((D_y)^{\pi_{i,y_j}})q'_{\Pi,y_j}(\pi_{i,y_j}^{-1}|(D_y)^{\pi_{i,y_j}})/q'_{\Pi,y_j}(\pi_{i,y_j} \circ \pi_{i,y_j}^{-1}|(D_y)^{\pi_{i,y_j}})}{\sum_{(i',j') \in \Lambda} \hat{f}((D_y)^{\pi_{i',y_{j'}}})q'_{\Pi,y_{j'}}(\pi_{i',y_{j'}}^{-1}|(D_y)^{\pi_{i',y_{j'}}})q'_{\Pi,y_{j'}}(\pi_{i',y_{j'}} \circ \pi_{i',y_{j'}}^{-1}|(D_y)^{\pi_{i',y_{j'}}})} \\
&= \frac{\hat{f}((D_y)^{\pi_{i,y_j}})q'_{\Pi,y_j}(\pi_{i,y_j}^{-1}|(D_y)^{\pi_{i,y_j}})/q'_{\Pi,y_j}(\text{id}|(D_y)^{\pi_{i,y_j}})}{\sum_{(i',j') \in \Lambda} \hat{f}((D_y)^{\pi_{i',y_{j'}}})q'_{\Pi,y_{j'}}(\pi_{i',y_{j'}}^{-1}|(D_y)^{\pi_{i',y_{j'}}})q'_{\Pi,y_{j'}}(\text{id}|(D_y)^{\pi_{i',y_{j'}}})},
\end{aligned}$$

where  $\text{id}$  denotes the identity permutation. Hence, computation of  $w_{\mathfrak{D}_{\Pi,i,y_j}}^{q,y}((D_y)^{\pi_{i,y_j}})$  across all  $(i,j) \in \Lambda$  becomes tractable in only  $O(m|\mathcal{Y}|)$  operations.

*Proof.* Consider a non-stationarity test in which we draw  $m|\mathcal{Y}|$  samples and the resampling distribution  $\check{q}$  we consider ignores  $Y_T$ , and instead draws these  $m|\mathcal{Y}|$  samples in  $|\mathcal{Y}|$  groups, the  $j^{\text{th}}$  of which consists of  $m$  resamples  $D^{\pi_{1,y_j}}, \dots, D^{\pi_{m,y_j}}$  where the  $\pi_{i,y_j}$  are drawn conditionally i.i.d. from  $q'_{\Pi,y_j}(\cdot|D)$  for each  $j \in [|\mathcal{Y}|]$ , and where  $\mathcal{Y} = \{y_1, \dots, y_{|\mathcal{Y}|}\}$ . It is important to note that  $\check{q}$  is *not* a conditionally i.i.d. resampling scheme and thus, as discussed in Section 5.3.2, the workaround discussed in Section E.1 must be employed. Defining

$$\Sigma = \{(0 \ \ell) : \ell \in [0 : m|\mathcal{Y}]\} \text{ and letting } i \in [0 : m] \text{ index with any given group and } j \in \begin{cases} [|\mathcal{Y}|] & \text{if } i > 0 \\ \{0\} & \text{if } i = 0 \end{cases}$$

index over groups, notice that

$$w_{\mathfrak{D}_{\Pi,i,y_j}}^{q,Y_T}(D^{\pi_{i,y_j}}) = w_{\mathfrak{D}_{\Sigma,(i,j)}}^{\check{q},\Sigma}(\tilde{D}^{(i,j)}),$$

where  $\mathfrak{D}$  denotes the full set of  $m|\mathcal{Y}|$  resamples and we double index resamples as  $D^{(i,j)} = D^{\pi_{i,y_j}}$  for all  $(i,j)$  in  $\Lambda$ . This is because, letting  $q'_{y_j}$  denote the induced (by  $q'_{\Pi,y_j}$ ) distribution on each resample given by  $q'_{y_j}(D^\pi|D) = q'_{\Pi,y_j}(\pi|D)$  and looking at the rightmost terms in the numerator of  $w_{\mathfrak{D}_{\Pi,i,y_j}}^{q,y}((D_y)^{\pi_{i,y_j}})$ 's definition, we have that

$$\begin{aligned}
& q'_{\Pi,y_j}(\pi_{i,y_j}^{-1}|D^{\pi_{i,y_j}}) \prod_{(\tilde{i},\tilde{j}) \in \Lambda \setminus \{(i,j)\}} q'_{\Pi,y_{\tilde{j}}}(\pi_{\tilde{i},y_{\tilde{j}}} \circ \pi_{i,y_j}^{-1}|D^{\pi_{i,y_j}}) \\
&= q'_{y_j}(D|D^{\pi_{i,y_j}}) \prod_{(\tilde{i},\tilde{j}) \in \Lambda \setminus \{(i,j)\}} q'_{y_{\tilde{j}}}(D^{\pi_{\tilde{i},y_{\tilde{j}}}}|D^{\pi_{i,y_j}}) \\
&= q'_{y_j}(D|\tilde{D}^{(i,j)}) \prod_{(\tilde{i},\tilde{j}) \in \Lambda \setminus \{(i,j)\}} q'_{y_{\tilde{j}}}(\tilde{D}^{(\tilde{i},\tilde{j})}|\tilde{D}^{(i,j)}) \\
&= \sum_{\pi \in \Sigma_{(i,j)}} \check{q}(((\tilde{D}^{(0,0)}, \tilde{D}^{(1,1)}, \dots, \tilde{D}^{(m,1)}, \dots, \tilde{D}^{(1,|\mathcal{Y}|)}, \dots, \tilde{D}^{(m,|\mathcal{Y}|)})^\pi)_{-0}|\tilde{D}^{(i,j)}),
\end{aligned}$$

where  $\Sigma_{(i,j)}$  is the single element subset of  $\Sigma$  containing solely the permutation on  $[0 : m|\mathcal{Y}]$  that swaps 0 and  $m(j-1) + i$ . Thus, the numerators of  $w_{\mathfrak{D}_{\Pi,i,y_j}}^{q,Y_T}(D^{\pi_{i,y_j}})$  and  $w_{\mathfrak{D}_{\Sigma,(i,j)}}^{\check{q},\Sigma}(\tilde{D}^{(i,j)})$  are equal, and precisely the same argument shows that the denominators are also equal and thus  $w_{\mathfrak{D}_{\Pi,i,y_j}}^{q,Y_T}(D^{\pi_{i,y_j}}) = w_{\mathfrak{D}_{\Sigma,(i,j)}}^{\check{q},\Sigma}(\tilde{D}^{(i,j)})$ .

Hence, the p-value

$$\sum_{(i,j) \in \Lambda} w_{\mathfrak{D}_{\Pi,i,y_j}}^{q,y}(D^{\pi_{i,y_j}}) \mathbf{1}[S(D^{\pi_{i,y_j}}) \geq S(D)],$$

corresponding to line (16) is a valid p-value by the main result of Section E.1. Letting  $U$  denote the exogenous randomness used by this test, notice that the exogenous random variables  $U_y$  used by each  $y \in \mathcal{Y}$

in the prediction region construction of line (16) are all (deterministically) equal—since the only randomness in determining if  $y \in \mathfrak{A}$  is in the random sampling of permutations  $\mathfrak{D}_\Pi$ , which does not depend on the specific choice of  $y \in \mathcal{Y}$ —and marginally identically distributed to  $U$ . Since each of  $U$  and  $U_y$  is generated independently from  $Y_T$ , we have that  $(U_{Y_T}, Y_T) \stackrel{d}{=} (U, Y_T)$  and hence the corresponding prediction region  $\mathfrak{A}$  is valid, as desired.  $\square$

## F Pseudocode for resampling distributions

In this appendix section, we provide pseudocode for all resampling procedures described in Section 4. Throughout this section, we use the notation  $\text{Cat}(p)$ , for any  $p \in [0, 1]^d$  with  $\sum_{i=1}^d p_i = 1$ , to denote the categorical distribution on  $[d]$  with probability of sampling  $i$  equal to  $p_i$ .

### F.1 Non-stationarity testing in a $C$ -stationary strongly non-reactive environment

We begin with the resampling procedures for non-stationarity testing in a  $C$ -stationary strongly non-reactive environment (Environment 5).

---

#### Algorithm 3: imitation $_\pi$

---

**Input:** Data sequence  $D$

1 Set  $\tilde{D}$  to the empty list and set  $R \leftarrow [T]$

2 **for**  $t = 1, \dots, T$  **do**

3     Sample

$$i \sim \text{Cat} \left( \left( \frac{\mathbb{P}_{\mathcal{A}}(X_j | \tilde{D}, C_j) \mathbf{1}[j \in R]}{\sum_{j'=1}^T \mathbb{P}_{\mathcal{A}}(X_{j'} | \tilde{D}, C_{j'}) \mathbf{1}[j' \in R]} \right)_{j=1}^T \right),$$

if  $\mathbb{P}_{\mathcal{A}}(\cdot | \tilde{D})$ 's support intersects  $R$ ; otherwise, terminate the sampling procedure

4     Append  $Z_i$  to  $\tilde{D}$  and set  $R \leftarrow R \setminus \{i\}$

**Output:**  $\tilde{D}$

---



---

#### Algorithm 4: re-imitation $_\pi$

---

**Input:** Data sequence  $D$ ; probability distribution  $\mathbb{P}_{U_t}$  denoting the distribution of the  $t^{\text{th}}$  exogenous random variable  $U_t$  generated by  $\mathcal{A}$

1 Set  $\tilde{D}$  to the empty list and set  $R \leftarrow [T]$

2 **for**  $t = 1, \dots, T$  **do**

3     Sample

$$\tilde{U}_t \sim \mathbb{P}_{U_t}(\cdot | \exists s \in R : X_s = \delta_t(C_s, \tilde{H}_{t-1}, \tilde{U}_1, \dots, \tilde{U}_t))$$

if the conditioning event is non-empty; otherwise, terminate the sampling procedure

4     Sample  $i \sim \text{Unif}(\{s \in R : X_s = \delta_t(C_s, \tilde{H}_{t-1}, \tilde{U}_1, \dots, \tilde{U}_t)\})$

5     Append  $Z_i$  to  $\tilde{D}$  and set  $R \leftarrow R \setminus \{i\}$

**Output:**  $\tilde{D}$

---

---

**Algorithm 5:** cond-imitation $_{\pi}$ 

---

**Input:** Data sequence  $D$  as well as the exogenous randomness  $U_1, \dots, U_T$  used to generate it

- 1 Set  $\tilde{D}$  to the empty list and set  $R \leftarrow [T]$
- 2 **for**  $t = 1, \dots, T$  **do**
- 3     Sample  $i \sim \text{Unif}(\{s \in R : X_s = \delta_t(C_s, \tilde{H}_{t-1}, U_1, \dots, U_t)\})$  if the set is non-empty; otherwise, terminate the sampling procedure
- 4     Append  $Z_i$  to  $\tilde{D}$  and set  $R \leftarrow R \setminus \{i\}$

**Output:**  $\tilde{D}$

---

## F.2 Non-stationarity testing in an MDP

As discussed in Section 4.2, the datasets used in non-stationarity testing in an MDP are augmented with an additional action  $X_{T+1}$ , and thus we may view these datasets as a list of  $T + 1$  state action pairs:  $D = ((C_1, X_1), \dots, (C_{T+1}, X_{T+1}))$ . Under this framework, all resampling procedures in this section utilize the following function  $\phi$ , which takes the state-action pair  $(c, x)$  to the set of indices which follow it:

$$\phi(c, x) = \{t \in [2 : T + 1] : (C_{t-1}, X_{t-1}) = (c, x)\}.$$

Additionally, using this view of the dataset  $D$ , we use  $Z_t$  to denote the  $t^{\text{th}}$  state-action pair  $(C_t, X_t)$ ;  $\tilde{Z}_t$  denotes the  $t^{\text{th}}$  state-action pair of the resampled dataset  $\tilde{D}$ . Using the function  $\phi$ , we now present the corresponding uniform $_{\pi}$ , imitation $_{\pi}$ , re-imitation $_{\pi}$ , and cond-imitation $_{\pi}$  for an MDP under this setup of the dataset  $D$ .

---

**Algorithm 6:** uniform $_{\pi}$  in an MDP

---

**Input:** Data sequence  $D$

- 1 Set  $\tilde{D} \leftarrow ((C_1, X_1))$  and set  $R \leftarrow [1 : T + 1]$
- 2 **for**  $t = 2, \dots, T + 1$  **do**
- 3     Sample
$$i \sim \text{Unif}(R \cap \phi(\tilde{Z}_{t-1}))$$
- 4     if the set is non-empty; otherwise, terminate the sampling procedure
- 4     Append  $Z_i$  to  $\tilde{D}$  and set  $R \leftarrow R \setminus \{i\}$

**Output:**  $\tilde{D}$

---

---

**Algorithm 7:** imitation $_{\pi}$  in an MDP

---

**Input:** Data sequence  $D$

- 1 Set  $\tilde{D} \leftarrow ((C_1, X_1))$  and set  $R \leftarrow [1 : T + 1]$
- 2 **for**  $t = 2, \dots, T + 1$  **do**
- 3     Sample
$$i \sim \text{Cat}\left(\left(\frac{\mathbb{P}_{\mathcal{A}}(X_j | \tilde{D}, C_j) \mathbf{1}[j \in R \cap \phi(\tilde{Z}_{t-1})]}{\sum_{j'=2}^{T+1} \mathbb{P}_{\mathcal{A}}(X_{j'} | \tilde{D}, C_{j'}) \mathbf{1}[j' \in R \cap \phi(\tilde{Z}_{t-1})]}\right)_{j=2}^{T+1}\right),$$
- 4     if  $\exists j \in R \cap \phi(\tilde{Z}_{t-1})$  such that  $\mathbb{P}_{\mathcal{A}}(X_j | \tilde{D}, C_j) > 0$ ; otherwise, terminate the sampling procedure
- 4     Append  $Z_i$  to  $\tilde{D}$  and set  $R \leftarrow R \setminus \{i\}$

**Output:**  $\tilde{D}$

---

---

**Algorithm 8:** re-imitation $_{\pi}$  in an MDP

---

**Input:** Data sequence  $D$ ; probability distribution  $\mathbb{P}_{U_t}$  denoting the distribution of the  $t^{\text{th}}$  exogenous random variable  $U_t$  generated by  $\mathcal{A}$

- 1 Set  $\tilde{D} \leftarrow ((C_1, X_1))$  and set  $R \leftarrow [1 : T + 1]$
- 2 **for**  $t = 2, \dots, T + 1$  **do**
- 3     Sample
  - 4          $\tilde{U}_t \sim \mathbb{P}_{U_t}(\cdot | \exists s \in R \cap \phi(\tilde{Z}_{t-1}) : X_s = \delta_t(C_s, \tilde{H}_{t-1}, \tilde{U}_1, \dots, \tilde{U}_t))$
  - 5         if the conditioning event is non-empty; otherwise, terminate the sampling procedure
- 4     Sample  $i \sim \text{Unif}(\{s \in R \cap \phi(\tilde{Z}_{t-1}) : X_s = \delta_t(C_s, \tilde{H}_{t-1}, \tilde{U}_1, \dots, \tilde{U}_t)\})$
- 5     Append  $Z_i$  to  $\tilde{D}$  and set  $R \leftarrow R \setminus \{i\}$

**Output:**  $\tilde{D}$

---

---

**Algorithm 9:** cond-imitation $_{\pi}$  in an MDP

---

**Input:** Data sequence  $D$  as well as the exogenous randomness  $U_1, \dots, U_{T+1}$  used to generate it

- 1 Set  $\tilde{D} \leftarrow ((C_1, X_1))$  and set  $R \leftarrow [1 : T + 1]$
- 2 **for**  $t = 2, \dots, T + 1$  **do**
- 3     Sample  $i \sim \text{Unif}(\{s \in R \cap \phi(\tilde{Z}_{t-1}) : X_s = \delta_t(C_s, \tilde{H}_{t-1}, U_1, \dots, U_t)\})$  if the set is non-empty; otherwise, terminate the sampling procedure
- 4     Append  $Z_i$  to  $\tilde{D}$  and set  $R \leftarrow R \setminus \{i\}$

**Output:**  $\tilde{D}$

---

### F.3 Conditional independence testing

We now give pseudocode for our resampling procedures for conditional independence testing discussed in Section 4.3. We first show pseudocode for restricted-uniform $_{\pi}$  resampling, which, although not used on its own for conditional independence testing, makes up the first stage of the restricted-uniform $_{\pi}$ +imitation $_X$  resampling scheme. We then go on to present the imitation $_X$  resampling procedure, which is also a key ingredient that is used in the uniform $_{\pi}$ +imitation $_X$  and restricted-uniform $_{\pi}$ +imitation $_X$  resampling procedures, which are presented subsequently. Finally, we show pseudocode for combined $_{\pi, X}$  sampling, which combines the permutation and randomization of  $X_t$ 's into a single stage.

---

**Algorithm 10:** restricted-uniform $_{\pi}$ 

---

**Input:** Data sequence  $D$

- 1 Set  $\tilde{D}$  to the empty list
- 2 Set  $\Gamma \leftarrow \{\pi' \in \Pi_{[T]} : g(X_{\pi'(t)}) = g(X_t), \forall t \in [T]\}$
- 3 Sample  $\pi \sim \text{Unif}(\Gamma)$
- 4  $\tilde{D} \leftarrow (Z_{\pi(t)})_{t=1}^T$

**Output:**  $\tilde{D}$

---

## G Supplementary simulation results

### G.1 MCMC plots

In this section, we present the power, Type-I error, coverage, and length plots for the unweighted MCMC randomization test and its inversion, for the inferential tasks discussed in Section 5, in Figures 9–14.

---

**Algorithm 11:** imitation<sub>X</sub>

---

**Input:** Data sequence  $D$ 

- 1 Set  $\tilde{D}$  to the empty list
- 2 **for**  $t = 1, \dots, T$  **do**
- 3     Sample
  - if there exists  $x \in \mathcal{X}$  with  $\mathbb{P}_{\mathcal{A}}(x|\tilde{D}, C_t, g(X_t)) > 0$ ; otherwise, terminate the sampling procedure
- 4     Append  $\tilde{Z}_t := (C_t, \tilde{X}_t, Y_t)$  to  $\tilde{D}$

$$\tilde{X}_t \sim \mathbb{P}_{\mathcal{A}}(\cdot|\tilde{D}, C_t, g(X_t)),$$

**Output:**  $\tilde{D}$ 

---

---

**Algorithm 12:** uniform <sub>$\pi$</sub> +imitation<sub>X</sub>

---

**Input:** Data sequence  $D$ 

- 1 Sample  $D'$  according to the uniform <sub>$\pi$</sub>  distribution applied to  $D$
- 2 Sample  $\tilde{D}$  according to the imitation<sub>X</sub> distribution (Algorithm 11) applied to  $D'$

**Output:**  $\tilde{D}$ 

---

---

**Algorithm 13:** restricted-uniform <sub>$\pi$</sub> +imitation<sub>X</sub>

---

**Input:** Data sequence  $D$ 

- 1 Sample  $D'$  according to the restricted-uniform <sub>$\pi$</sub>  distribution (Algorithm 10) applied to  $D$
- 2 Sample  $\tilde{D}$  according to the imitation<sub>X</sub> distribution (Algorithm 11) applied to  $D'$

**Output:**  $\tilde{D}$ 

---

---

**Algorithm 14:** combined <sub>$\pi, X$</sub> 

---

**Input:** Data sequence  $D$ 

- 1 Set  $\tilde{D}$  to the empty list and set  $R \leftarrow [T]$
- 2 **for**  $t = 1, \dots, T$  **do**
- 3     Sample

$$i \sim \text{Cat} \left( \left( \frac{\sum_{x \in \mathcal{X}: g(x)=g(X_j)} \mathbb{P}_{\mathcal{A}}(x|\tilde{D}, C_j) \mathbf{1}[j \in R]}{\sum_{j'=1}^T \sum_{x' \in \mathcal{X}: g(x')=g(X_{j'})} \mathbb{P}_{\mathcal{A}}(x'|\tilde{D}, C_{j'}) \mathbf{1}[j' \in R]} \right)_{j=1}^T \right),$$

- if there exists  $j \in R$  such that  $\sum_{x \in \mathcal{X}: g(x)=g(X_j)} \mathbb{P}_{\mathcal{A}}(x|\tilde{D}, C_j) > 0$ ; otherwise, terminate the sampling procedure
- 4     Set  $R \leftarrow R \setminus \{i\}$
- 5     Sample
  - if there exists  $x \in \mathcal{X}$  with  $\mathbb{P}_{\mathcal{A}}(x|\tilde{D}, C_i, g(X_i)) > 0$ ; otherwise, terminate the sampling procedure
- 6     Append  $\tilde{Z}_i := (C_i, \tilde{X}_i, Y_i)$  to  $\tilde{D}$

$$\tilde{X}_i \sim \mathbb{P}_{\mathcal{A}}(\cdot|\tilde{D}, C_i, g(X_i)),$$

**Output:**  $\tilde{D}$ 

---

## G.2 Computation times

In this section, we plot the computation time curves (to compute  $p$ ) for all resampling algorithms, environments, adaptive assignment algorithms, and types of randomization test (i.e., weighted MC or unweighted MCMC) discussed in Section 5 in Figures 15–22.

Contextless stationary strongly non-reactive environment conditional independence test

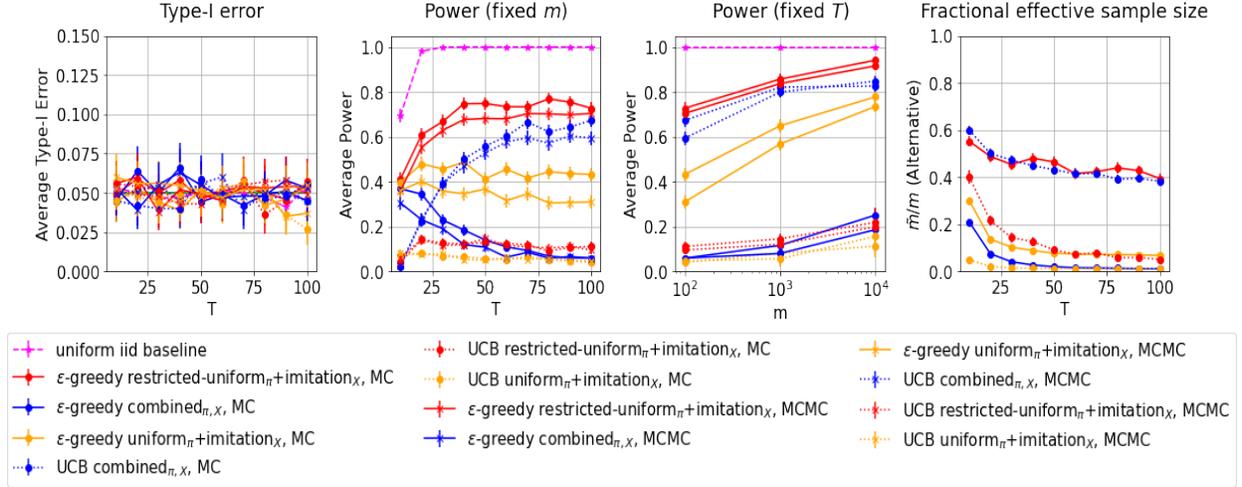


Figure 9: Type-I error rate (leftmost) and power (second from left) of both weighted MC and unweighted MCMC randomization tests at fixed  $m = 100$  and varying  $T$  as well as power at fixed  $T = 100$  and varying  $m$  (third from left) and fractional effective sample size plots at fixed  $m = 100$  and varying  $T$  (rightmost) in a contextless stationary strongly non-reactive environment on data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

Contextless C-stationary strongly non-reactive environment, non-stationarity test

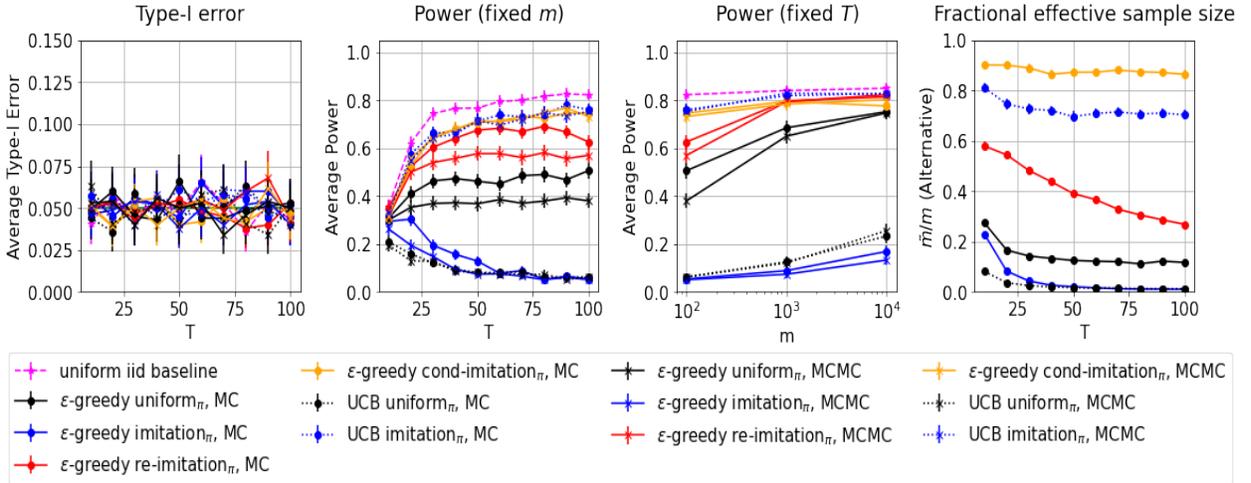


Figure 10: Type-I error rate (leftmost) and power (second from right) of both weighted MC and unweighted MCMC randomization tests at fixed  $m = 100$  and varying  $T$  as well as power for fixed  $T = 100$  and varying  $m$  (third from right) and fractional effective sample size at fixed  $m = 100$  and varying  $T$  (rightmost) in a contextless C-stationary strongly non-reactive environment with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

### G.3 Auxiliary simulation results

In this section we present all remaining auxiliary simulation results. Figure 23 displays the power plot at fixed  $T = 100$  and varying  $m$  for the conditional independence test in the contextual stationary strongly non-reactive environment on data gathered via  $\epsilon$ -greedy and LinUCB discussed in Section 5.1.1. On the other hand, Figure 24 illustrates that the phenomenon of shift in relative performance of the uniform i.i.d. baseline in comparison to both  $\epsilon$ -greedy and LinUCB, described in Section 5.2.2 also occurs when the baseline is compared to a biased i.i.d. adaptive assignment algorithm which selects actions at each timestep independently from  $2\text{Bern}(0.1) - 1$ .

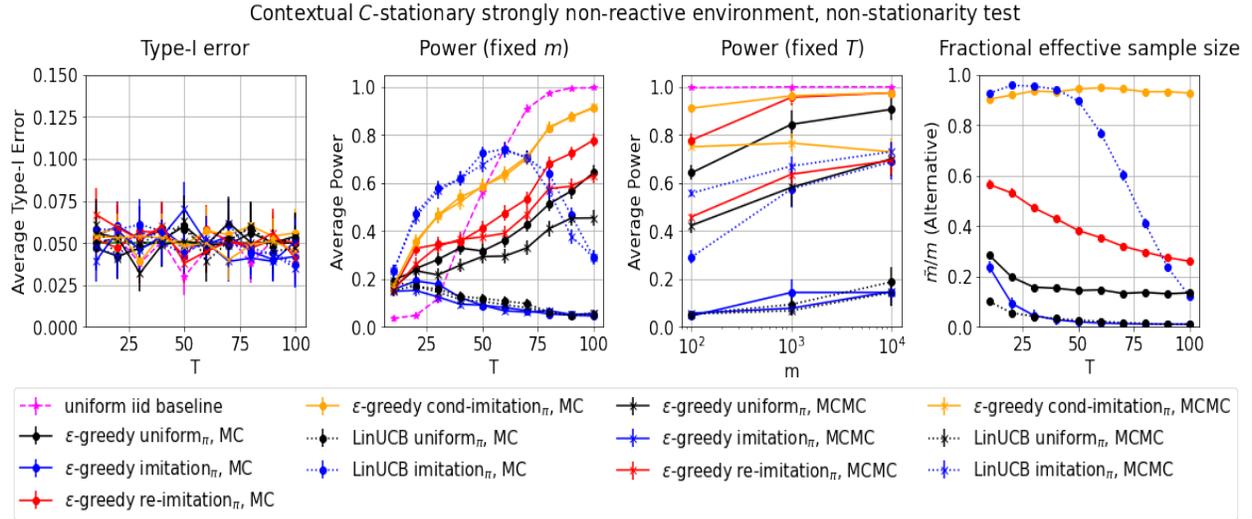


Figure 11: Type-I error rate (leftmost) and power (second from right) of both weighted MC and unweighted MCMC randomization tests at fixed  $m = 100$  and varying  $T$  as well as power for fixed  $T = 100$  and varying  $m$  (third from right) and fractional effective sample size at fixed  $m = 100$  and varying  $T$  (rightmost) in a contextless  $C$ -stationary strongly non-reactive environment with data gathered via  $\epsilon$ -greedy, LinUCB, and the uniform i.i.d. baseline.

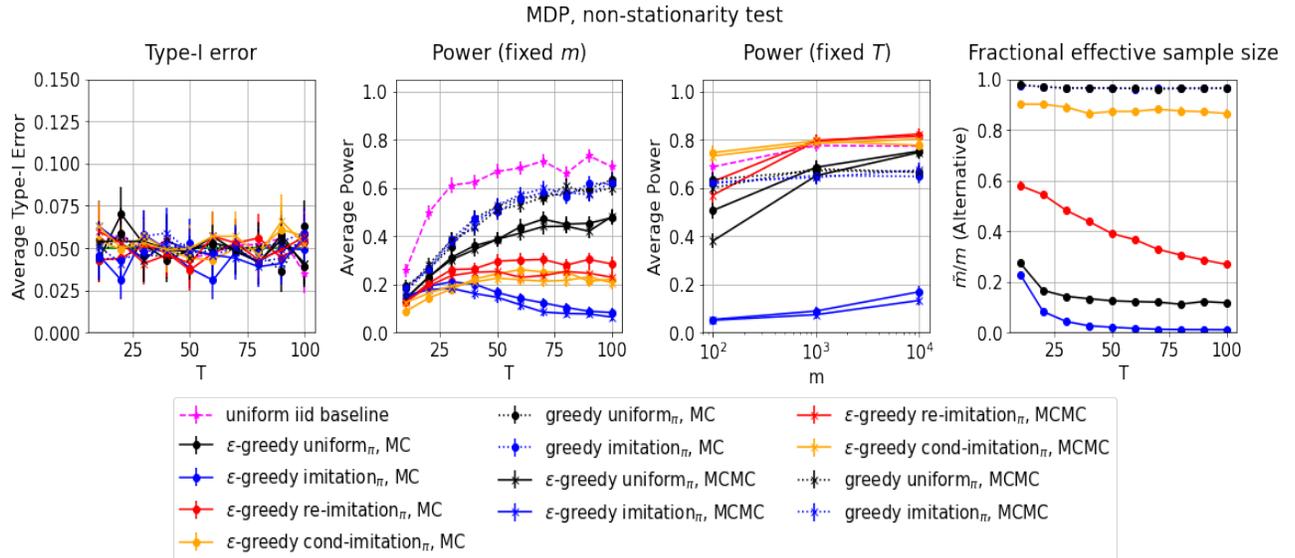


Figure 12: Type-I error rate (leftmost) and power (second from right) of both weighted MC and unweighted MCMC randomization tests at fixed  $m = 100$  and varying  $T$  as well as power for fixed  $T = 100$  and varying  $m$  (third from right) and fractional effective sample size at fixed  $m = 100$  and varying  $T$  (rightmost) in a contextless  $C$ -stationary strongly non-reactive environment with data gathered via  $\epsilon$ -greedy and greedy  $Q$ -learning.

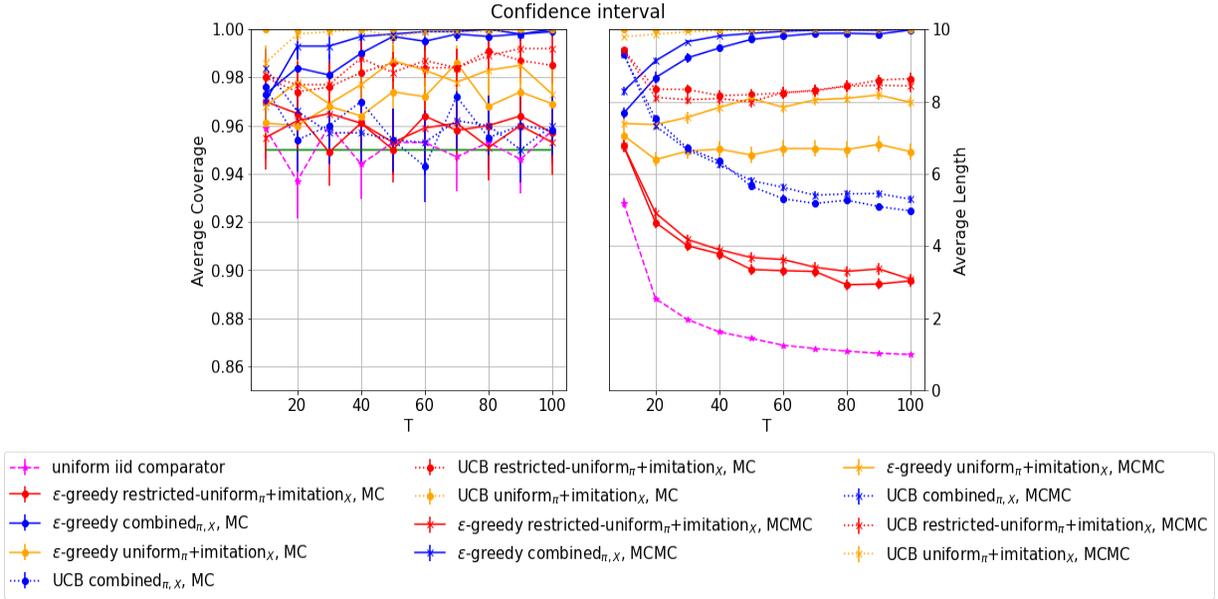


Figure 13: Coverage and average length of confidence intervals for  $b_0$  using both weighted MC and unweighted MCMC randomization tests with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

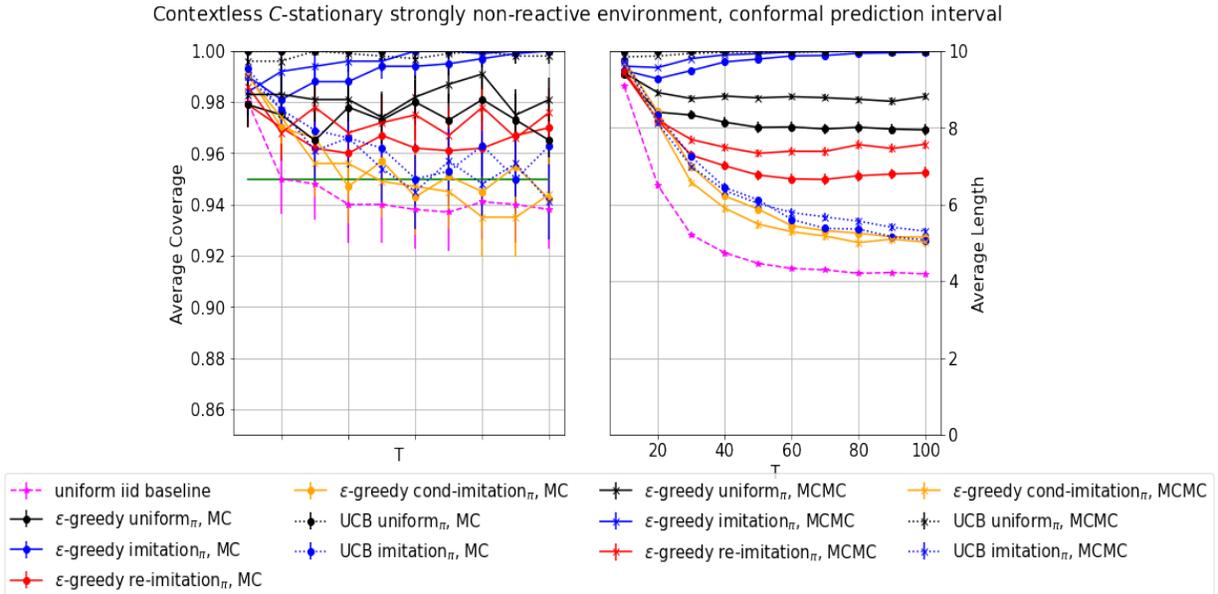


Figure 14: Coverage and average length of conformal prediction intervals for  $Y_T$  using both weighted MC and unweighted MCMC randomization tests with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

Contextual stationary strongly non-reactive environment conditional independence test, computation times

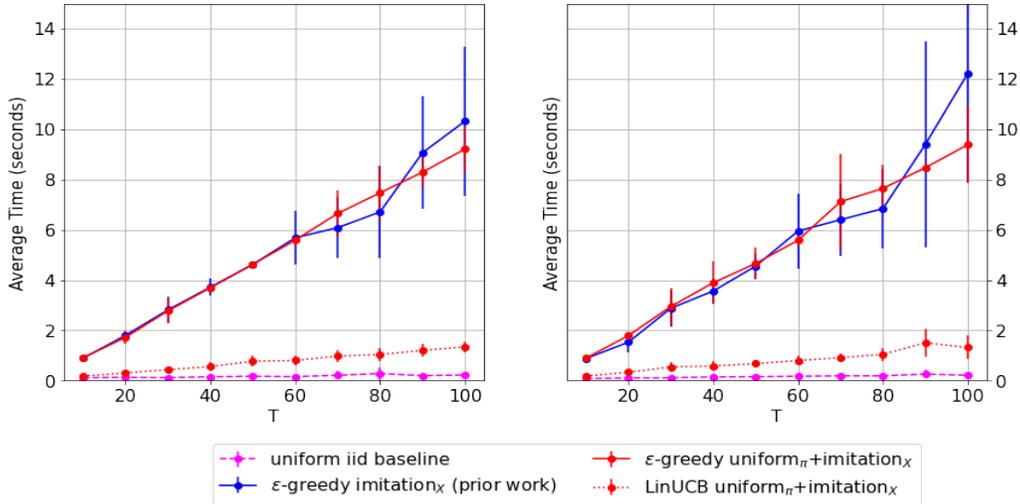


Figure 15: Computation times under the null (left) and alternative (right) distributions of randomization tests at fixed  $m = 100$  and varying  $T$  in a contextual stationary strongly non-reactive environment on data gathered via  $\epsilon$ -greedy and LinUCB. Note that the computation times for LinUCB imitation $_X$  are omitted as both Type-I error and power curves are simply generated i.i.d. Bern(0.05).

Contextless stationary strongly non-reactive environment conditional independence test, average computation times

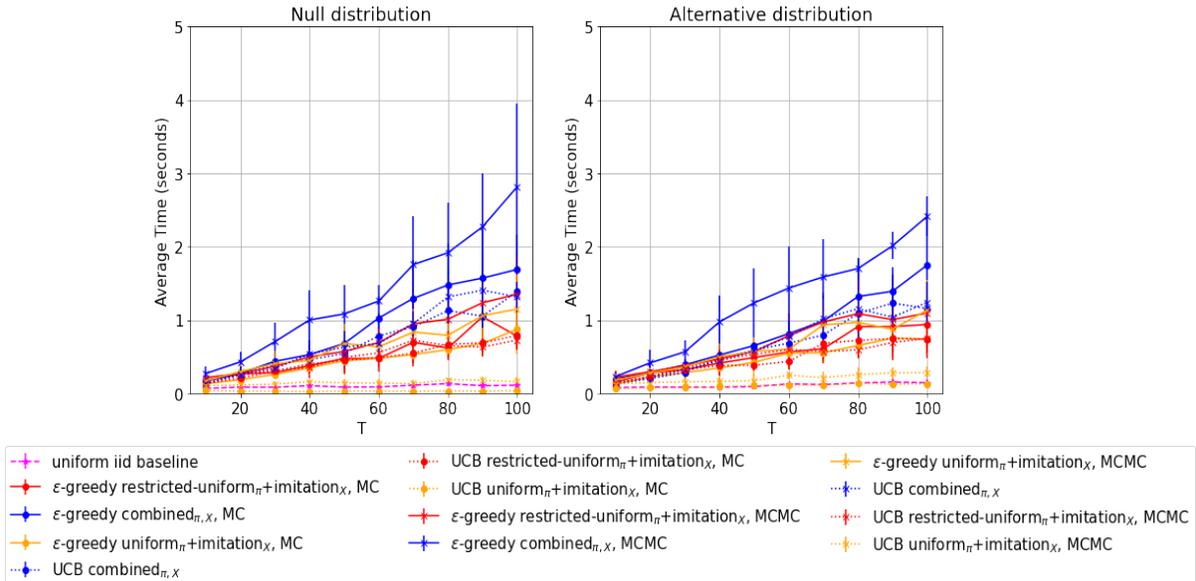


Figure 16: Computation times under the null (left) and alternative (right) distributions of both weighted MC and unweighted MCMC randomization tests at fixed  $m = 100$  and varying  $T$  in a contextless stationary strongly non-reactive environment on data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

Contextless  $C$ -stationary strongly non-reactive environment, non-stationarity test, average computation times

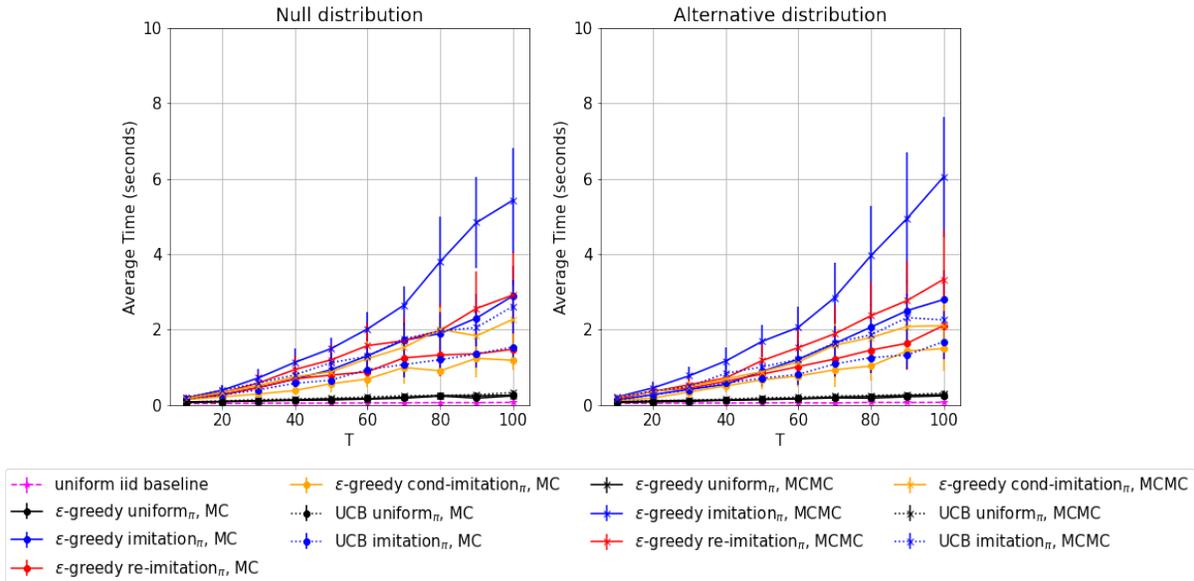


Figure 17: Computation times under the null (left) and alternative (right) distributions of both weighted MC and unweighted MCMC randomization tests at fixed  $m = 100$  and varying  $T$  in a contextless  $C$ -stationary strongly non-reactive environment with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

Contextual  $C$ -stationary strongly non-reactive environment, non-stationarity test, computation times

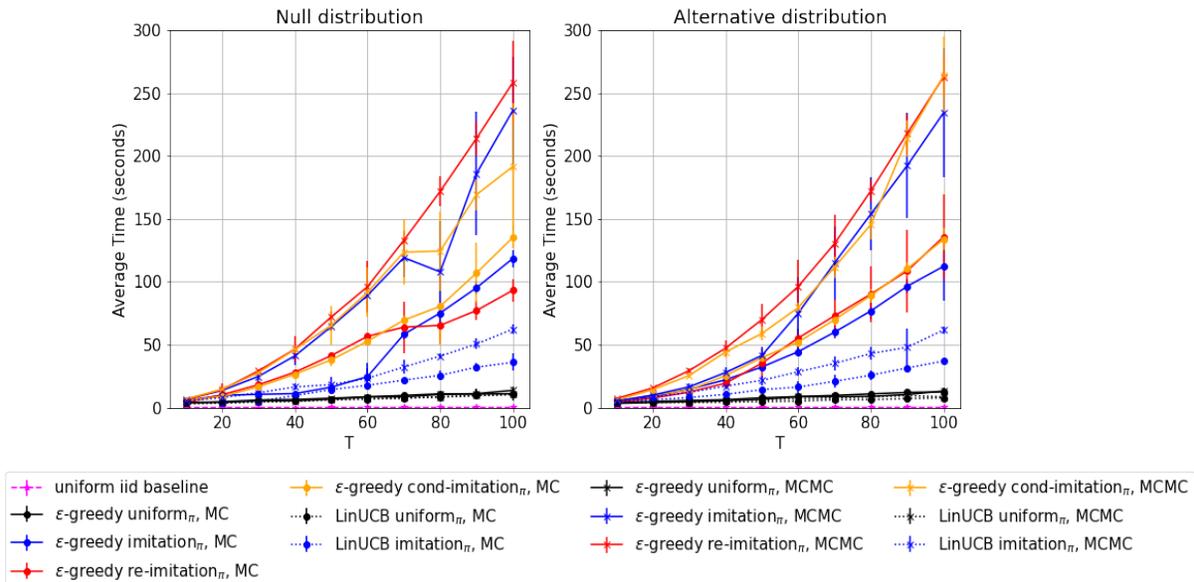


Figure 18: Computation times under the null (left) and alternative (right) distributions of both weighted MC and unweighted MCMC randomization tests at fixed  $m = 100$  and varying  $T$  in a contextless  $C$ -stationary strongly non-reactive environment with data gathered via  $\epsilon$ -greedy, LinUCB, and the uniform i.i.d. baseline.

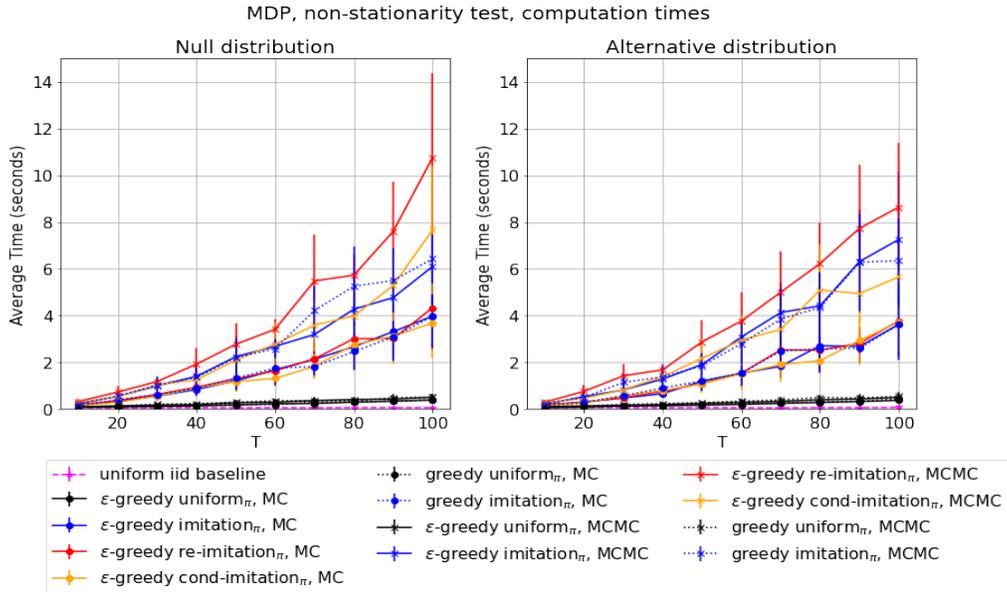


Figure 19: Computation times under the null (left) and alternative (right) distributions of both weighted MC and unweighted MCMC randomization tests at fixed  $m = 100$  and varying  $T$  in a contextless  $C$ -stationary strongly non-reactive environment with data gathered via  $\epsilon$ -greedy and greedy  $Q$ -learning.

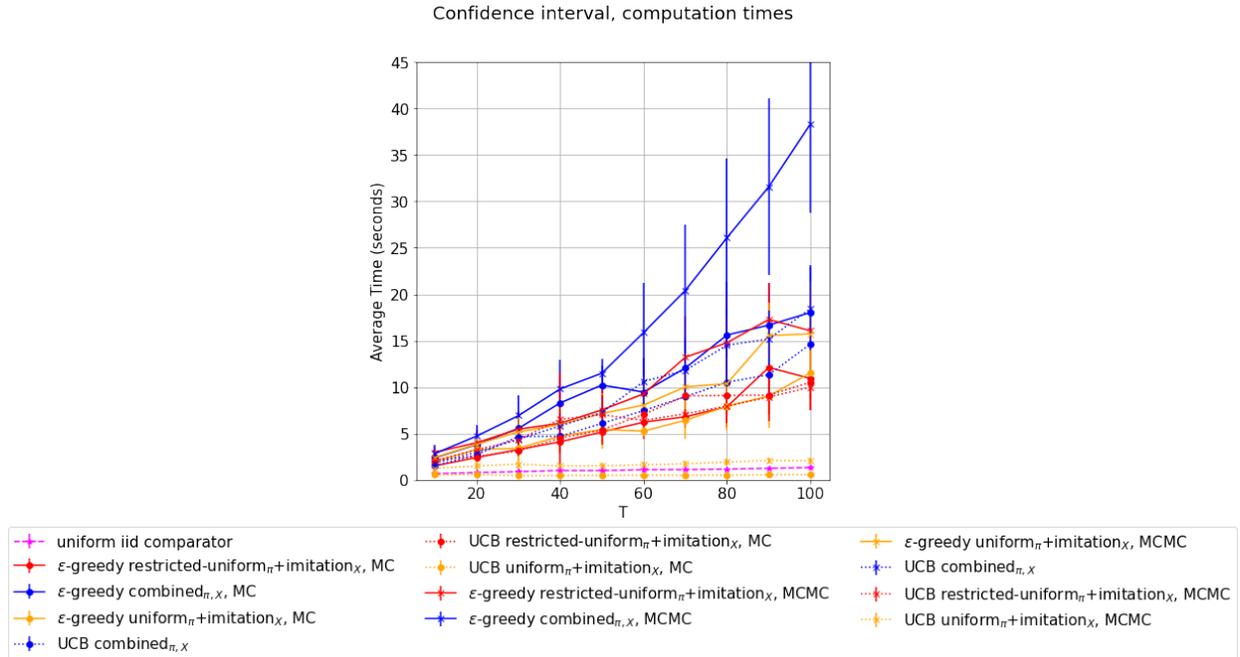


Figure 20: Computation time of construction of confidence interval for  $b_0$  using both weighted MC and unweighted MCMC randomization tests with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

Contextless  $C$ -stationary strongly non-reactive environment, conformal prediction interval, computation times

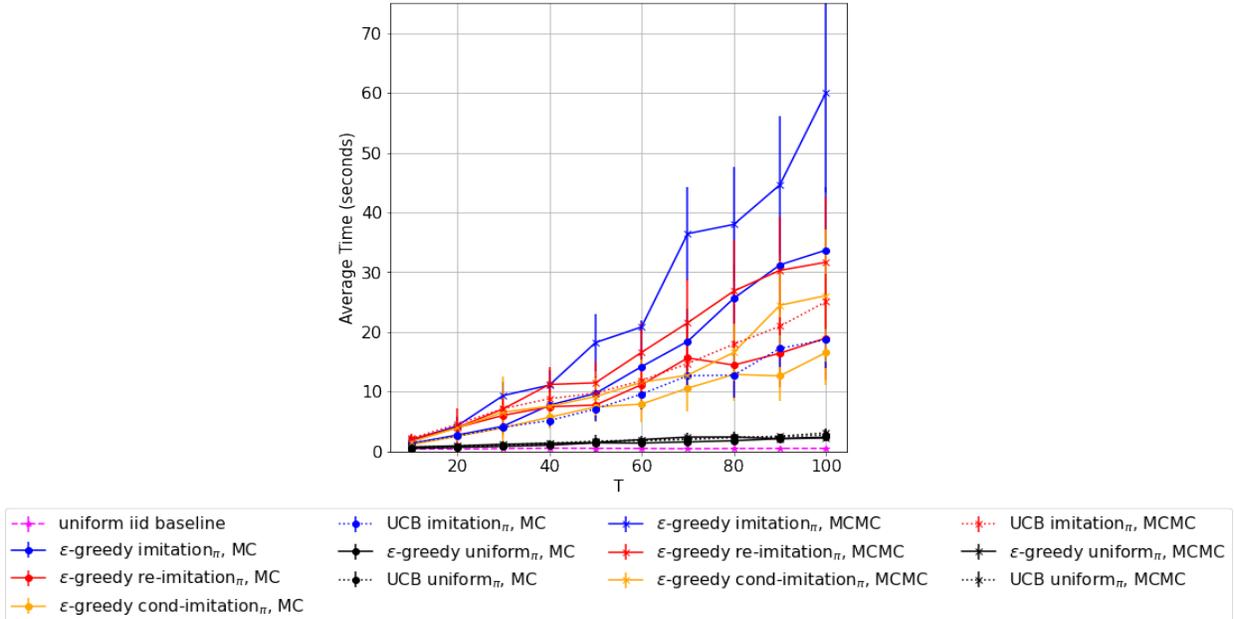


Figure 21: Computation time of construction of conformal prediction interval for  $Y_T$  using both weighted MC and unweighted MCMC randomization tests with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

Contextual  $C$ -stationary strongly non-reactive environment, conformal prediction interval, shared samples, computation times

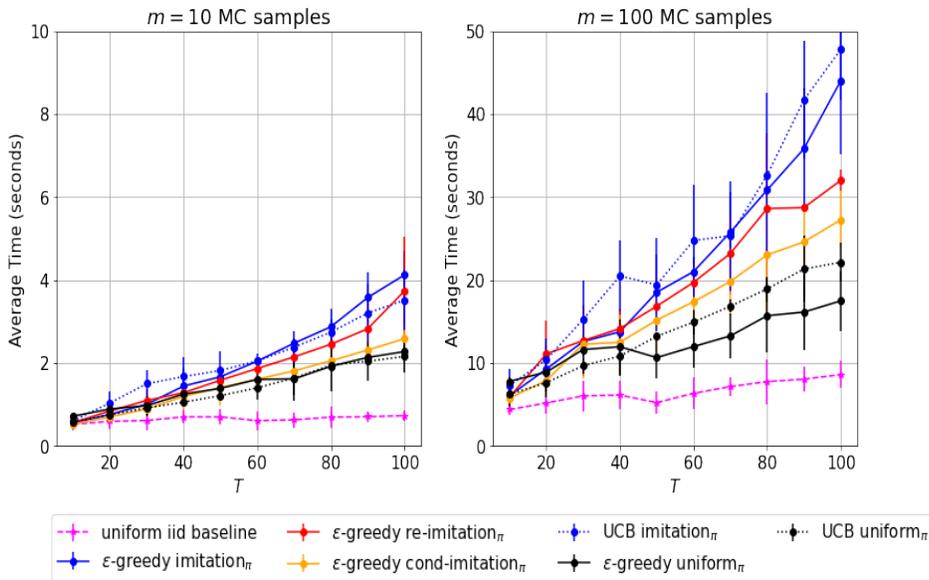


Figure 22: Computation time of construction of conformal prediction interval for  $Y_T$  using the MC randomization test with data gathered via  $\epsilon$ -greedy, UCB, and the uniform i.i.d. baseline.

Contextual stationary strongly non-reactive environment conditional independence test

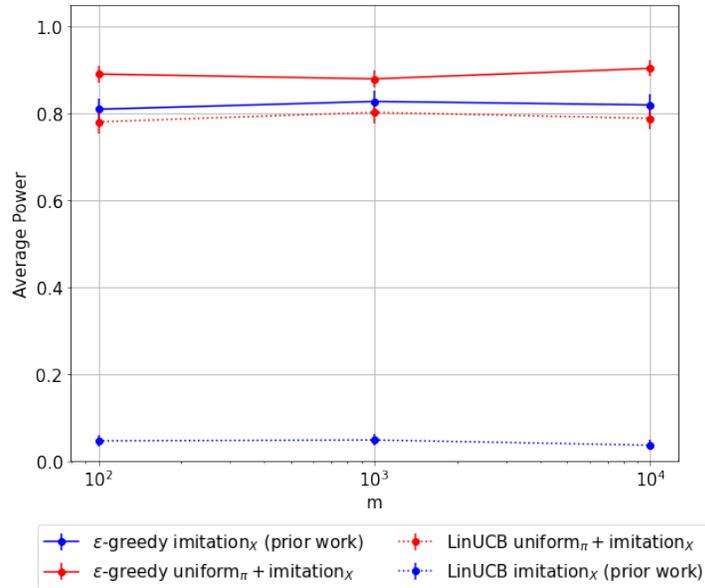


Figure 23: Power of randomization tests at fixed  $T = 100$  and varying  $m$  in a contextual stationary strongly non-reactive environment on data gathered via  $\epsilon$ -greedy and LinUCB.

Contextual C-stationary strongly non-reactive environment, non-stationarity test, comparison of iid assignments: power

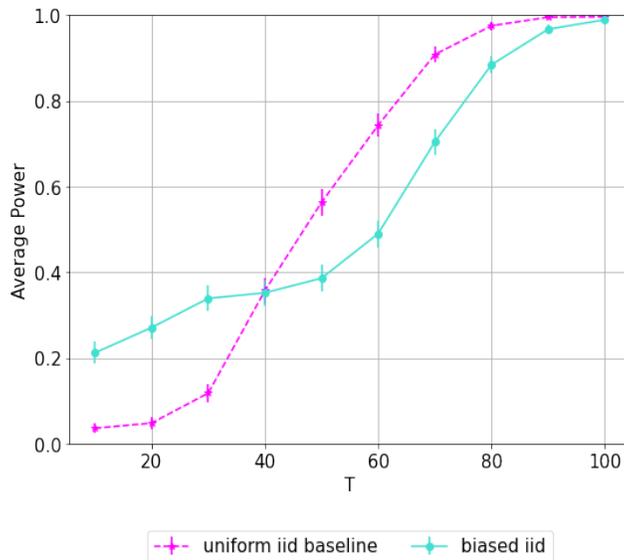


Figure 24: Power comparison of uniform i.i.d. baseline versus biased i.i.d. baseline that selects actions independently from  $2\text{Bern}(0.1) - 1$