

Anytime-valid inference in N-of-1 trials

Ivana Malenica

Harvard University, United States

IMALENICA@FAS.HARVARD.EDU

Yongyi Guo

University of Wisconsin-Madison, United States

GUO98@WISC.EDU

Kyra Gan

Cornell Tech, United States

KYRAGAN@CORNELL.EDU

Stefan Konigorski

Harvard University, United States; Hasso Plattner Institute, Germany

STEFAN.KONIGORSKI@HPI.DE

Abstract

App-based N-of-1 trials offer a scalable experimental design for assessing the effects of health interventions at an individual level. Their practical success depends on the strong motivation of participants, which, in turn, translates into high adherence and reduced loss to follow-up. One way to maintain participant engagement is by sharing their interim results. Continuously testing hypotheses during a trial, known as “peeking”, can also lead to shorter, lower-risk trials by detecting strong effects early. Nevertheless, traditionally, results are only presented upon the trial’s conclusion. In this work, we introduce a potential outcomes framework that permits interim peeking of the results and enables statistically valid inferences to be drawn at any point during N-of-1 trials. Our work builds on the growing literature on *valid confidence sequences*, which enables anytime-valid inference with uniform type-1 error guarantees over time. We propose several causal estimands for treatment effects applicable in an N-of-1 trial and demonstrate, through empirical evaluation, that the proposed approach results in valid confidence sequences over time. We anticipate that incorporating anytime-valid inference into clinical trials can significantly enhance trial participation and empower participants.

Keywords: N-of-1 trials, anytime-valid inference, design-based, confidence sequence, causal inference, personalized medicine

many clinical and biomedical applications. In particular, understanding which treatment and dosage are most effective for a particular patient lies at the heart of personalized medicine. To this aim, different methodologies have been proposed. One popular approach involves the collection and analysis of extensive population-level datasets with the goal of estimating effects at the individual level. With suitable covariates and a clear understanding of disease mechanisms, it becomes feasible to potentially acquire individual-level effects (Shalit et al., 2017; Bica et al., 2020; Smith et al., 2020; Verstraete et al., 2021; Diemert et al., 2021). Nevertheless, the practicality of this approach is often limited to specific applications. For instance, in cases like cancer, personalized signatures can be derived from genetic mutations. However, more often, causal effects can only be identified within specific subgroups (van Kruijsdijk et al., 2014; Zhang et al., 2017; van Amsterdam et al., 2022; Msaouel et al., 2022). As a second approach, stemming from recent biotechnological advancements, personalized treatments have been developed directly for selected rare target diseases in personalized drug development (Kim et al., 2019; Seydel, 2023). However, such approaches are still restricted to a selected class of drug targets and limited by resources and costs.

As a third approach, experimental studies can be designed to directly evaluate and compare the effectiveness of one or multiple treatments in a given person. These so-called N-of-1 trials have been established as the gold standard for inferring individual-level effects, and different guidelines have been proposed for their standardized application (Nikles and Mitchell, 2015; Vohra et al., 2015; Porcino

1. Introduction

The statistical inference of individual causal effects of health interventions holds great importance for

et al., 2020). More formally, N-of-1 trials are multicrossover randomized controlled trials (RCTs) in one person, hence one or more treatments are administered over time in a predefined, potentially randomized, sequence. Also, if the same N-of-1 trial is performed in multiple persons in a so-called series of N-of-1 trials, the trials can be jointly analyzed to yield efficient population-level treatment effect estimates (Zucker et al., 2010). In order to achieve sufficient statistical power for inference, N-of-1 trials require frequent longitudinal measurements which has hindered their application in practice. Only recently, digital tools have been developed for setting up and executing N-of-1 trials. This, in turn, enables scalability across various studies, patients, and providers (Taylor et al., 2018; Daskalova et al., 2020; Zenner et al., 2022; Konigorski et al., 2022). For example, the StudyU platform (<https://www.studyu.health>) contains several open studies that are currently running with patients and healthy participants in different countries, and are investigating different health interventions on outcomes such as chronic pain, antidepressant discontinuation symptoms, and sleep quality.

Digital app-based N-of-1 trials provide a straightforward and expandable means for evaluating patient outcomes, whether actively reported by patients or passively collected sensor data. However, their success still relies on the retention and high adherence of the patients to the trial. To keep the trial participants engaged, apart from developing user-friendly interfaces for the apps, it is also important to provide frequent feedback. A classical study design involves collecting all trial data, conducting analysis, and subsequently reporting the results back to the participants. However, if patients are required to provide daily outcomes in a trial that runs for weeks or even months, without any intermediate feedback, patients might lose interest. On the other hand, “peeking” at the intermediate results (performing a hypothesis test before the end of the trial) introduces statistical biases. Thus, a framework for valid statistical inference is required at all time points of an N-of-1 trial in order to enable intermediate analysis.

Contributions. In this work, we provide a statistical framework that allows anytime-valid inference in N-of-1 trials with intermediate analysis of the results. Our contributions are to provide a (i) potential outcomes framework for N-of-1 trials, (ii) formal definition of several causal estimands of interest in N-of-1 trials, and (iii) construction of confidence se-

quences that allow anytime-valid inference. Finally, we validate our approach empirically in simulation studies and compare it to existing state-of-the-art approaches.

1.1. Related Work

The field of N-of-1 trials has originated from the medical domain, and continues to be largely driven by its clinical applications. As such, the existing literature has almost exclusively focused on an applied presentation, and omitted a more formal statistical definition of the study setup and estimands of interest. There exist some recent contributions: Yang et al. (2021) give a formalization of different potential study designs and their consequences on design parameters. Daza (2018, 2019) and Daza and Schneider (2022) present a counterfactual framework for single case studies which include both observational and experimental N-of-1 settings. They define individual-level target estimands and estimators in their work, which can include time trends and carryover effects. As a difference to these previous approaches, we consider design-based estimands based on the immediate causal effects conditional on the accrued history, and focus on constructing anytime-valid confidence sequences for the proposed target parameters.

Other related work has been published in the traditional RCTs literature that is relevant to our aim of enabling anytime-valid inference. In RCTs, interim analyses are often performed to evaluate safety and side effects, as well as to assess intermediate treatment effects that might warrant early termination of the trial based on predefined criteria for treatment inferiority or superiority. Common approaches for such interim analyses include the O’Brien-Fleming (O’Brien and Fleming, 1979), Haybittle-Peto (Haybittle, 1971; Peto et al., 1976) and Pocock (Pocock, 1977) methods, which aim to control the overall type I error across all interim tests by adjusting the respective critical test statistic values. While the Pocock method chooses the same alpha level at all interim tests, the O’Brien-Fleming and Haybittle-Peto method require stronger evidence at earlier interim points. In the CONSORT extension with guidelines for harm-related stopping rules, the O’Brien-Fleming method is recommended (Ioannidis et al., 2004). Demets and Lan (1994) and others have generalized these ideas through an “alpha-spending function” which generates critical values such that the sum of probabilities of exceeding those values

across the interim tests equal the type I error rate alpha. We note that all these approaches have been designed for classical population-level RCTs, and have not been evaluated in N-of-1 trials. As only exception, [Selukar \(2021\)](#) considers a very specific situation, when a series of N-of-1 trials is sequentially monitored, interim analyses are few and pre-defined, and only summary statistics are available of each trial.

Third, relevant work has originated from causal inference literature on sequential designs for single time series ([Bojinov and Shephard, 2019](#); [Malenica et al., 2021](#); [Ham et al., 2022](#)). [Malenica et al. \(2021\)](#) define conditional estimand classes of interest and provide inference on them in an adaptive setting of a single time series. [Bojinov and Shephard \(2019\)](#) provide a potential outcome framework for single time series and define estimands of interest from the design-based perspective. Finally, we adapt and build on work by [Ham et al. \(2022\)](#), who provide design-based confidence sequences for very long time series in a setting where treatment is randomized at each time point. In this work we focus specifically on the setup of N-of-1 trials as crossover experiments in treatment blocks, with N-of-1 trial-specific estimands.

2. Statistical Formulation

2.1. Notation and Observed Data

We consider the trajectory of a single individual in an N-of-1 trial with K treatment periods (also denoted as “periods” or “blocks”). Suppose for each period $k \in [K] := \{1, \dots, K\}$, there are T_k time points. We write (k, t) to indicate the t -th time point of the treatment period k . We emphasize that, despite our focus on a single N-of-1 trial, our results generalize to a series of N-of-1 trials with multiple patients.

Let $O_{k,t} := (A_{k,t}, Y_{k,t}, W_{k,t})$ denote data for a single individual at time point (k, t) including the treatment, outcome, and covariates. Specifically, at each time point (k, t) , one assigns a binary treatment $A_{k,t} \in \mathcal{A} := \{0, 1\}$ to a patient, where $A_{k,t} = 1$ denotes the treatment and $A_{k,t} = 0$ the control (or alternative treatment). In N-of-1 trials, the same treatment is assigned throughout a period (e.g., for block k , control is given at all time points $(1, \dots, T_k)$). In an extreme case, one might randomize at each follow-up time (corresponding to the chronological observation), so $T_k = 1$ for all $k \in [K]$. Notice that the number of time points within a block, T_k , can vary by k , allowing for different lengths of each treatment pe-

riod. Once the treatment is assigned for block k , post-treatment outcome of interest $Y_{k,t} \in \mathcal{Y}$ and possibly a vector of other time-varying covariates $W_{k,t} \in \mathcal{W}$ are collected. Therefore, at each time point t , we assign treatment $A_{k,t}$ according to the period k , then collect $Y_{k,t}$, followed by $W_{k,t}$. See [Figure 1](#) for an illustration of the N-of-1 trial path.

Let $O_{k,1:t} = (O_{k,1}, \dots, O_{k,t})$ denote the data collected in period k up to time t . As each k contains T_k data points, we define $O_k := O_{k,1:T_k} = (O_{k,1}, \dots, O_{k,T_k})$ as the full data collected in the k -th treatment period. Similarly, we write $A_k = (A_{k,1}, \dots, A_{k,T_k})$ for the full and $A_{k,1:t} = (A_{k,1}, \dots, A_{k,t})$ for the cropped sequence of treatments in block k . To clarify, we illustrate the proposed notation with a simple example where $K = 2$ and $T_1 = T_2 = 2$. Without putting any assumptions on the design of treatment blocks (and therefore treatment assignment), all possible treatment sequences for a single individual are as follows: $\{(1, 1), (0, 0)\}$, $\{(0, 0), (1, 1)\}$, $\{(1, 1), (1, 1)\}$ and $\{(0, 0), (0, 0)\}$. Symbolically, we write the resulting treatment sequence as $\{(A_{1,1}, A_{1,2}), (A_{2,1}, A_{2,2})\} = \{A_{1,1:2}, A_{2,1:2}\}$.

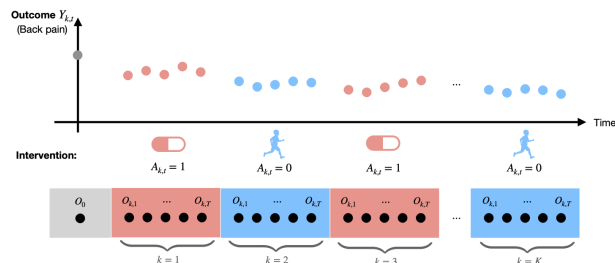


Figure 1: Illustration of a N-of-1 trial path for a single individual.

It follows that data collected for a whole block k is then $O_k = (A_k, Y_k, W_k)$, where $Y_k = (Y_{k,1}, \dots, Y_{k,T_k})$ and $W_k = (W_{k,1}, \dots, W_{k,T_k})$. When no time-varying covariates are gathered beyond the outcome of interest, we represent this as $W_k = \emptyset$. We emphasize that W_k (and the entire history, or some function of it, for that matter) can be used to inform treatment allocation in the subsequent block, $k + 1$, allowing for adaptive treatment assignment. Finally, let O_0 be the baseline covariates obtained before the start of the trial. Without loss of generality, we assume $T_k = T$ for all blocks $k \in [K]$ ($T > 1$), which aligns with the

design of a canonical N-of-1 trial. We, however, emphasize that all of our results generalize to the setting where the T_k values vary in each treatment period.

Below, we provide some essential definitions regarding time and block-specific variable collections used throughout the manuscript. For a single individual, let $\bar{O}_{k,t} = (O_0, O_{1,1}, \dots, O_{k,t})$ denote the full history up to time t of block k (including baseline covariates), while $\underline{O}_{k,t} = (O_{k,t}, \dots, O_{K,T})$ denotes all future variables from index (k, t) to (K, T) . Then, the full *trajectory* (or a *time-series*) is presented by $O^{KT} := \bar{O}_{K,T} = \{O_0, O_{1,1}, \dots, O_{K,T}\}$. Similarly, we define all past and future collections of A , Y , and W . For example, we write $\bar{A}_{k,t} = (A_{1,1}, \dots, A_{k,t})$ for the sequence of treatments until the t -th time point of period k . Lastly, let $H_{k,t}^A := \bar{O}_{k,t-1}$ be the full history of all variables until $A_{k,t}$. It then follows that $H_{k,t}^Y := (A_{k,t}, \bar{O}_{k,t-1})$ and $H_{k,t}^W := (A_{k,t}, Y_{k,t}, \bar{O}_{k,t-1})$ are the full variable histories until $Y_{k,t}$ and $W_{k,t}$.

2.2. Likelihood of the Trajectory

We let $O^{KT} \sim P_0$, where P_0 denotes the true probability distribution of O^{KT} . Throughout the remainder of the text, we use capital letters to indicate random variables, and lowercase for their realizations. We use the naught subscript to denote *true* probability distributions or components thereof. Let \mathcal{M} denote the *statistical model* for the probability distribution of the data, which is nonparametric, beyond possible knowledge of the treatment mechanism (i.e., known randomization probabilities). We note that the true probability distribution of the data is an element of \mathcal{M} , and denote P as any probability distribution such that $P \in \mathcal{M}$. Suppose that P_0 admits a density p_0 w.r.t. a dominating measure μ over \mathcal{O} which can be written as the product measure $\mu = \times_{k=1, t=1}^{k=K, t=T} (\mu_A \times \mu_Y \times \mu_W)$, with μ_A , μ_Y , and μ_W measures over \mathcal{A} , \mathcal{Y} , and \mathcal{W} . The likelihood of O^{KT} can be factorized according to the time-ordering as

$$p_0(O^{KT}) = p_{0, O_0}(o_0) \prod_{k=1}^K \prod_{t=1}^{T_k} p_{0, A}(a_{k,t} | h_{k,t}^A) \quad (1)$$

$$p_{0, Y}(y_{k,t} | h_{k,t}^Y) p_{0, W}(w_{k,t} | h_{k,t}^W),$$

where $a_{k,t} \mapsto p_{0, A}(a_{k,t} | h_{k,t}^A)$, $y_{k,t} \mapsto p_{0, Y}(y_{k,t} | h_{k,t}^Y)$, and $w_{k,t} \mapsto p_{0, W}(w_{k,t} | h_{k,t}^W)$ are conditional densities w.r.t. dominating measures μ_A , μ_Y , μ_W .

3. Causal Effects for N-of-1 Trials

The main design components of an N-of-1 trial include the (1) number of blocks K , (2) length of each block, T , and (3) choice of treatment allocation for treatment periods (pre-specified or randomized, type of randomization) (Yang et al., 2021). If the treatment sequence is pre-specified before a trial, then an individual is assigned a specific treatment sequence deterministically (e.g., control-treatment-control-treatment). This is useful when one is interested in the effect of a specific treatment sequence, or wishes to avoid randomly generating unwanted treatment sequences (e.g., giving the same treatment across consecutive periods). Alternatively, one might generate treatment sequences by randomizing treatment allocation across blocks. In this work, we focus specifically on randomized treatment sequences.

There are numerous randomization schemes for designing an N-of-1 trial (Yang et al., 2021). In this work, we rely on (1) *pairwise randomization*, where the order of two different treatments in a consecutive pair of treatment periods is randomized; (2) *restricted randomization*, where treatment is randomly assigned with the restriction that the number of treatment and control periods is approximately the same (but treatment probability is never zero) and (3) *unrestricted randomization*, where treatment is randomly assigned at each period. All of the listed schemes randomize on the block-level, and assign the same treatment at all time points within a block.

3.1. Time-Series Potential Outcomes

We define $\bar{a}_k = (a_1, \dots, a_k)$ as the *treatment path* until period k , where we remind that $a_k = a_{k,1:T}$. In a point treatment setting, the treatment path is of length 1, and each study participant has 2^1 potential outcomes. In an N-of-1 trial, however, we follow a single individual over time and administer K different treatments, each of length T (or more generally, T_k). For a binary treatment and unrestricted randomization, we then have 2^K different treatment paths that could have been observed. In Table 1, we include the total number of possible treatment sequences for each considered randomization scheme, at both odd and even number of treatment periods.

We define $Y_{k,t}(\bar{a}_{k,t})$ as the potential outcome at time point (k, t) , which may depend on the full history of assigned treatments up until (k, t) . Note that we don't make assumptions on carryover or other time-dependent effects. Conse-

Table 1: Number of possible treatment paths generated under different randomization schemes.

Randomization	Odd K	Even K
Pairwise	$2^{\frac{K+1}{2}}$	$2^{\frac{K}{2}}$
Restricted	$\sim 2^{\binom{K}{(K-1)/2}}$	$\sim 2^{\binom{K}{K/2}}$
Unrestricted	2^K	2^K

quently, we denote $\bar{Y}_{k,t}(\bar{a}_{k,t})$ as the collection of potential outcomes up until (k, t) : $\bar{Y}_{k,t}(\bar{a}_{k,t}) = (Y_{1,1}(\bar{a}_{1,1}), \dots, Y_{1,T}(\bar{a}_{1,T}), \dots, Y_{k,1}(\bar{a}_{k,1}), \dots, Y_{k,t}(\bar{a}_{k,t}))$. Further, let $Y_k(\bar{a}_k)$ denote a summary potential outcome of treatment period k . In particular, $Y_k(\bar{a}_k)$ depends on the potential outcomes for all time points t within block k , i.e., $Y_k(\bar{a}_k) := f(Y_{k,1}(\bar{a}_{k,1}), \dots, Y_{k,T}(\bar{a}_{k,T}))$, where f is any function that takes as input the potential outcomes in block k . Finally, let $\bar{Y}_k(\bar{a}_k) = (Y_1(\bar{a}_1), \dots, Y_k(\bar{a}_k))$ denote the counterfactual outcomes that would have been observed over time under treatment path $\bar{A}_k = \bar{a}_k$. When f is the average function across the respective block, we have that $\bar{Y}_k(\bar{a}_k) = (Y_1(\bar{a}_1), \dots, Y_k(\bar{a}_k)) = (1/T \sum_{t=1}^T Y_{1,t}(\bar{a}_{1,t}), \dots, 1/T \sum_{t=1}^T Y_{K,t}(\bar{a}_{K,t}))$.

Note that we assume that future potential outcomes do not cause past treatments (Granger, 1969). We formalize this in Assumption 1. Finally, we do not make assumptions on the dimension of any of the defined potential outcomes.

Assumption 1 (Granger Causality) Let $\bar{W}_T = \emptyset$. For each $k \in [K]$, we have that

$$\begin{aligned} P(A_{k,1} = a_{k,1} \mid \bar{A}_{k-1} = \bar{a}_{k-1}, \bar{Y}_K(\bar{a}_K)) \\ = P(A_{k,1} = a_{k,1} \mid \bar{A}_{k-1} = \bar{a}_{k-1}, \bar{Y}_{k-1}(\bar{a}_{k-1})). \end{aligned}$$

To illustrate the notation, consider a simple design with unrestricted randomization, $A \in \{0, 1\}$, $K = 2$ and $T = 2$. For the first period $k = 1$, there are $2^1 = 2$ potential outcomes: $Y_{1,2}(0, 0)$ and $Y_{1,2}(1, 1)$. For the second treatment period, there are $2^2 = 4$ potential outcomes at the end of the treatment block: $Y_{2,2}(0, 0, 0, 0)$, $Y_{2,2}(0, 0, 1, 1)$, $Y_{2,2}(1, 1, 0, 0)$, and $Y_{2,2}(1, 1, 1, 1)$, and the total number of potential outcomes at the end of the two treatment periods under unrestricted randomization is $2 + 4 = 6$ (more generally, $2(2^K - 1)$ for a trial consisting of K treatment periods). The total number of treatment paths, however, is $2^K = 2^2 = 4$, out of which we observe only

one. See Bojinov and Shephard (2019) for more examples. Let's consider the block-average function for f , so that each $Y_k(\bar{a}_k)$ is the average of all potential outcomes in block k . For the first period, we then have that $Y_1(\bar{a}_1) = 1/2 \sum_{t=1}^{T=2} Y_{1,t}(\bar{a}_{1,t})$, and $Y_1(\bar{a}_1)$ is an average of $Y_{1,1}(0)$ and $Y_{1,2}(0, 0)$, or $Y_{1,1}(1)$ and $Y_{1,2}(1, 1)$. Similarly, for the second treatment period $k = 2$, $Y_2(\bar{a}_2) = 1/2 \sum_{t=1}^{T=2} Y_{2,t}(\bar{a}_{2,t})$. Therefore, the potential outcomes that would have been observed over time are $\bar{Y}_2(\bar{a}_2) = (Y_1(\bar{a}_1), Y_2(\bar{a}_2))$.

3.2. Target Parameter

In an N-of-1 trial, treatments are assigned based on the current treatment period. In this section, we establish several causal estimands that might be of interest in an N-of-1 trial, using the potential outcomes as defined in Subsection 3.1. The first target parameter is the *immediate causal effect* (ICE) of treatment, as opposed to control, at period k . It is defined as the short-term (contemporaneous) effect of administering treatments during period k , assessed at time (k, T) , right after the last treatment in period k , conditional on the observed past. A formal definition is provided in Definition 1. Another parameter of interest is the time t -specific ICE, which represents the causal effect of assigning treatment from time $(k, 1)$ to (k, t) . This target parameter hints at the effect of administering treatment for t time points in a treatment period. We note that ICE is a special case of the time t -specific ICE where $t = T$.

Definition 1 (Immediate Causal Effect)

$$\psi_k(\bar{a}_{k-1}) = Y_k(\bar{a}_{k-1}, a_k = \mathbf{1}) - Y_k(\bar{a}_{k-1}, a_k = \mathbf{0})$$

for any $k \in [K]$. $\mathbf{1}$, $\mathbf{0}$ are vectors of dimension $T \times 1$.

Definition 2 (Time t -specific ICE)

$$\begin{aligned} \psi_{k,t}(\bar{a}_{k-1}) = Y_{k,t}(\bar{a}_{k-1}, a_{k,1:t} = \mathbf{1}_t) \\ - Y_{k,t}(\bar{a}_{k-1}, a_{k,1:t} = \mathbf{0}_t) \end{aligned}$$

for any $k \in [K]$, where $\mathbf{1}_t$ and $\mathbf{0}_t$ are vectors of dimension $t \times 1$, and $a_{k,1:t} = (a_{k,1}, \dots, a_{k,t})$.

Notice that the causal estimands in both Definition 1 and Definition 2 are functions of the entire treatment path. As such, they include the *carry-over effect* from the past treatment period assignment in addition to the effect of period k . We also emphasize that they are *data-adaptive* parameters — the estimand changes as a function of the observed

past and/or treatment path. Defining causal effects conditional on history might seem unusual from the classical causal inference perspective. However, such data-adaptive approach allows us to define causal effects (1) for long time-series, (2) with valid inference (as the central limit theorem still holds), and (3) without any additional assumptions on the time-series structure (Bojinov and Shephard, 2019). Data-adaptive target parameters in longitudinal settings have been previously described from both super population and design-based perspectives (Bojinov and Shephard, 2019; Malenica et al., 2021). Lastly, we define another target parameter of interest in N-of-1 trials in Definition 3, the Average Immediate Causal Effect (AICE) — the running average of treatment effects over blocks. A similar parameter can also be defined for the average over time t -specific ICE.

Definition 3 (AICE) For any $k \in [K]$,

$$\psi_{\text{AICE}}^k = \frac{1}{k} \sum_{j=1}^k \psi_j(\bar{a}_{j-1}).$$

3.3. Estimation

In this paper, we focus on the design-based approach to causal inference. Within the design-based paradigm, the full set of potential outcomes is fixed and always conditioned on; as such, the only source of randomness comes from the treatment assignment. Let \mathcal{F}_k denote the filtration which contains all observed data up to time (k, T) conditional on $\bar{Y}_k(\bar{a}_k)$ ($\bar{O}_k = \{(A_{j,t}, Y_{j,t}, W_{j,t})\}_{j=1, t=1}^{j=k, t=T}$), and all the potential outcomes ($\bar{Y}_K(\bar{a}_K)$). The filtration \mathcal{F}_k obeys the nesting property where $\mathcal{F}_k \subset \mathcal{F}_{k+1}$ for all k . At the beginning of each period k , we randomly assign treatment with probability $g(A_{k,1}) = P(A_{k,1} | \mathcal{F}_{k-1})$. Note that $g(A_{k,1}) = g(A_{k,2}) = \dots = g(A_{k,T})$, as the probability of treatment is the same for each time point t in block k . By Assumption 1, it follows that $g(A_{k,1}) = P(A_{k,1} | \mathcal{F}_{k-1}) = P(A_{k,1} | \bar{A}_{k-1} = \bar{a}_{k-1}, \bar{Y}_{k-1}(\bar{a}_{k-1}))$. We emphasize that if treatment is assigned independently of the past, then $g(A_{k,1}) = P(A_{k,1})$.

To estimate ICE and AICE, we focus on the time-series version of the Horvitz-Thomson and the stabilized IPTW (Hájek) estimator in this work (Horvitz and Thompson, 1952; Hájek, 1971; Robins et al., 2000; Hirano et al., 2003; Imbens and Rubin, 2015). To enable statistical inference, we assume there is a positive probability of treatment and control at every

period k . Formally stated in Assumption 2, the positivity assumption excludes the pre-determined treatment periods occasionally used in N-of-1 trials.

Assumption 2 (Positivity) For every $k \in [K]$,

$$0 < g(A_{k,1}) < 1.$$

Under Assumption 2, we can estimate ICE and AICE using the observed data. The IPTW estimator of $\psi_k(\bar{a}_{k-1})$, denoted as $\hat{\psi}_k$, is then defined as

$$\hat{\psi}_k := \frac{\mathbf{1}(A_{k,1} = 1)f(Y_{k,1:T})}{g(A_{k,1})} - \frac{\mathbf{1}(A_{k,1} = 0)f(Y_{k,1:T})}{1 - g(A_{k,1})}.$$

Recall f is any function that takes data of block k as input (in the estimator, observed data at time points $(k, 1 : T)$). The variance estimator is defined as

$$\hat{\sigma}_k^2 := \frac{\mathbf{1}(A_{k,1} = 1)f(Y_{k,1:T})^2}{g(A_{k,1})^2} + \frac{\mathbf{1}(A_{k,1} = 0)f(Y_{k,1:T})^2}{(1 - g(A_{k,1}))^2}.$$

Lemma 1 establishes that the proposed estimator is unbiased, and derives its variance over the randomization (proof in Appendix A).

Lemma 1 (Properties of the ICE Estimator)

Under Assumption 2, it follows that

$$\mathbb{E}(\hat{\psi}_k - \psi_k(\bar{a}_{k-1}) | \mathcal{F}_{k-1}) = 0$$

and

$$\text{Var}(\hat{\psi}_k - \psi_k(\bar{a}_{k-1}) | \mathcal{F}_{k-1}) \leq \mathbb{E}(\hat{\sigma}_k^2 | \mathcal{F}_{k-1}).$$

The running average immediate effect over treatment periods, i.e. the AICE, can be estimated by

$$\hat{\psi}_{\text{AICE}}^k = 1/k \sum_{j=1}^k \hat{\psi}_j. \quad (2)$$

The unbiasedness of $\hat{\psi}_{\text{AICE}}^k$ follows trivially from Lemma 1. To stabilize the variance of the IPTW, we also investigate the Hájek estimator of AICE, denoted as $\tilde{\psi}_{\text{AICE}}^k$ and presented in Equation (3). We allocate the study of the Hájek estimator of AICE to Appendix B.

$$\tilde{\psi}_{\text{AICE}}^k := \frac{\sum_{j=1}^k \mathbf{1}(A_{j,1} = 1)f(Y_{j,1:T})/g(A_{j,1})}{\sum_{j=1}^k \mathbf{1}(A_{j,1} = 1)/g(A_{j,1})} - \frac{\sum_{j=1}^k \mathbf{1}(A_{j,1} = 0)f(Y_{j,1:T})/(1 - g(A_{j,1}))}{\sum_{j=1}^k \mathbf{1}(A_{j,1} = 0)/(1 - g(A_{j,1}))}, \quad (3)$$

with the corresponding variance estimator:

$$\begin{aligned} \tilde{\sigma}_{\text{AICE}}^2 := & \frac{\sum_{j=1}^k \mathbb{1}(A_{j,1} = 1) f(Y_{j,1:T})^2 / g(A_{j,1})^2}{\sum_{j=1}^k \mathbb{1}(A_{j,1} = 1) / g(A_{j,1})^2} \\ & + \frac{\sum_{j=1}^k \mathbb{1}(A_{j,1} = 0) f(Y_{j,1:T})^2 / (1 - g(A_{j,1}))^2}{\sum_{j=1}^k \mathbb{1}(A_{j,1} = 0) / (1 - g(A_{j,1}))^2} \end{aligned} \quad (4)$$

4. Confidence Sequences

We now introduce design-based asymptotic confidence sequences for N-of-1 trials. First, we define a *confidence sequence* as a sequence of confidence intervals that are uniformly valid over time (also known as *anytime-valid*). We say $(I_k)_{k=1}^K$ is a valid confidence sequence with type-1 error α (or level $1 - \alpha$) for the target parameter $(\psi_{\text{AICE}}^k)_{k=1}^K$ if for any data-dependent stopping rule at $1 \leq \tau \leq K$,

$$P(\exists k \in \{1, \dots, \tau\} \text{ s.t. } \psi_{\text{AICE}}^k \notin I_k) \leq \alpha. \quad (5)$$

Under Equation (5), we can perform valid inference through each I_k . Furthermore, one can terminate a trial as soon as a statistically significant effect is detected with $\psi_{\text{AICE}}^k \notin I_k$, allowing for "peeking" during the trial duration.

Anytime-valid inference was first introduced by Wald (1945). Since then, significant advancements have been made in developing confidence sequences under minimal regularity conditions (Howard et al., 2021; Bibaut et al., 2023). Two of the key contributions include the idea of time-uniform analogues of asymptotic confidence intervals (known as *asymptotic confidence sequences*), and their extension to design-based framework for anytime-valid causal inference (Waudby-Smith et al., 2023; Ham et al., 2022). For clarity, we provide a semi-formal definition as Definition 4. Informally, asymptotic confidence sequences are valid confidence sequences as the number of time-points grows. While practically this might mean we don't have valid coverage at early times, this concern is alleviated by the (i) N-of-1 design, where each period is of length $T > 1$, and by the (ii) upper bound variance estimator introduced in Section 3.3.

Definition 4 (Asymptotic Confidence Sequence)

A sequence of intervals $(I_k)_{k=1}^K$ is a level $1 - \alpha$ asymptotic confidence sequence for the target parameter sequence $(\psi_{\text{AICE}}^k)_{k=1}^K$ if there exists a non-asymptotic confidence sequence $(I'_k)_{k=1}^K$ of level $1 - \alpha$ such that each interval I_k shares the center with I'_k ,

and that $\text{width}(I_k) / \text{width}(I'_k) \rightarrow 1$ a.s.. Moreover, we say $(I_k)_k$ has approximation rate R if $\text{width}(I_k) - \text{width}(I'_k) = O_{\text{a.s.}}(R)$.

We now formally introduce the asymptotically valid confidence sequences for the target parameter sequence of the running average immediate causal effect, $(\psi_{\text{AICE}}^k)_k$. We emphasize however, that all mentioned estimands in Section 3.2 are just special cases of the results derived in Theorem 1. In particular, Theorem 1 applies to a weighted mean version of AICE, which simplifies to the Immediate Causal Effect when the weight is zero for all treatment periods except the time t block. Before stating Theorem 1 (proof in Appendix C), we need two more assumptions. In Assumption 3, we assume there is an unknown, possibly extreme constant M which bounds the realized potential outcomes. As M can be arbitrarily large and realizations are bounded, we consider Assumption 3 a mild regularity condition. Assumption 4 concerns the behavior of the variance, and is satisfied as long as potential outcomes do not vanish over time.

Assumption 3 (Bounded Potential Outcomes)

There exists a constant $M \in \mathbb{R}$ such that for any $k \in [K]$, and any treatment path \bar{a}_k , $|Y_k(\bar{a}_k)| \leq M$.

Assumption 4 (Non-vanishing Variance)

Let $\tilde{S}_k := \sum_{j=1}^k \sigma_j^2$, where $\sigma_j^2 := \frac{Y_j(\bar{a}_{j-1}, \mathbf{1}_T)^2}{g(A_{j,1})} + \frac{Y_j(\bar{a}_{j-1}, \mathbf{0}_T)^2}{1 - g(A_{j,1})}$. Then, $\tilde{S}_k \rightarrow \infty$ as $k \rightarrow \infty$ a.s.

Theorem 1 Let $S_k = \sum_{j=1}^k \hat{\sigma}_j^2$. Under Assumptions 2, 3 and 4, and for any constant $\eta > 0$,

$$\frac{1}{k} \sum_{j=1}^k \hat{\psi}_j \pm \frac{1}{k} \sqrt{\frac{\eta^2 S_k + 1}{\eta^2} \log \left(\frac{\eta^2 S_k + 1}{\alpha^2} \right)}$$

forms a valid $(1 - \alpha)$ asymptotic confidence sequence for the target parameter sequence $(\psi_{\text{AICE}}^k)_k$, with approximation rate $o(\sqrt{\tilde{S}_k \log \tilde{S}_k / k})$.

5. Experiments

In the first experiment, we generate 1000 independent N-of-1 trials for an individual under the null hypothesis and illustrate the need for novel approaches to construct valid confidence sequences. In each independent trial, we set $K = 30$ blocks and $T = 10$ time points within each block. For simplicity, there is no treatment effect or covariates. We consider unrestricted randomization with 50% probability for each treatment.

First, we demonstrate the inflated type-1 error of a naive t-test (which constructs a confidence interval according to the two-sample t-test with level $\alpha = 0.05$ at each block K), as well as the O’Brien-Fleming approach (O’Brien and Fleming, 1979). In Figure 3, we observe that the naive approach quickly accumulates type I errors, reaching 0.6 only after 6 treatment blocks. A naive application of the O’Brien-Fleming approach yields better performance than the simple t-test, but still result in inflated type I errors that increase over time. In contrast, our proposed anytime-valid confidence interval achieves the target type 1 error level and does not yield any false positives.

The second experiment aims to visualize the performance of the proposed confidence sequences and demonstrate empirically that the sequences achieve both early stopping and time uniform coverage. For this purpose, we generate 1000 independent N-of-1 trials, each with a decreasing treatment effect size of $5 + 1/k$, no carry-over effects, $K = 100$ treatment blocks, and $T = 10$ data points within each block. We consider both the unrestricted and pairwise randomization schemes, where we construct confidence sequences of AICE using our proposed IPTW (Theorem 1 derived from Equation (2)) and stabilized IPTW estimators (Equation (3)). In Figure 2 and Figure 4 in Appendix D, we illustrate the confidence sequences constructed by the proposed algorithms in a single run.

In Table 2, we compute the average stopping time and time-uniform coverage proportion among the 1000 independent trials. Here the stopping time is set as the earliest block at which the confidence interval excludes 0, up to the 100th block. The time-uniform coverage refers to the proportion among the random experiments where *all* confidence intervals cover the true treatment effect. We observe that the proposed confidence sequences enjoy high coverage and their widths are decreasing over time as expected. The stabilized estimator shows a slightly smoother interval

sequence and slightly longer average stopping times, though overall both estimators show similar empirical behavior in our considered scenarios. These insights hold for both randomization schemes. In contrast, the O’Brien Fleming method has an average stopping time of 5.54 (2.93) and a uniform time coverage of 0. While the average stopping time of O’Brien Fleming is lower than for our proposed method, the reason for that is due to accumulated bias from testing many hypotheses (alas, better than the naive t-test, as seen in Figure 1). As expected, the coverage of O’Brien Fleming method is poor.




Figure 2: Anytime valid confidence intervals of AICE obtained by IPTW in a single run at $\alpha = 0.05$. The dashed line represents the zero (null) line. Top row: unrestricted randomization scheme. Bottom row: pairwise randomization scheme.

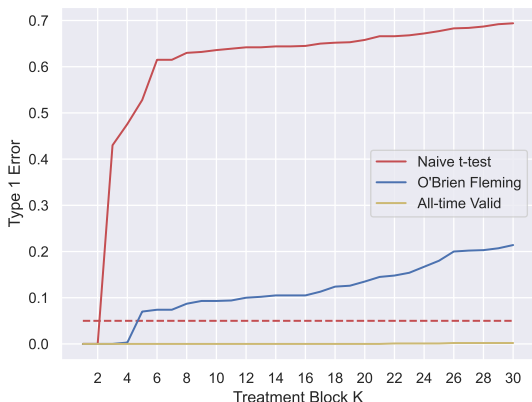


Figure 3: Empirical type I error of a naive t-test, the O’Brien-Fleming approach, and the anytime-valid approach, of 1000 N-of-1 trials in the unrestricted randomization setting. The dashed line represents $\alpha = 0.05$.

	Avg. Stop Time	Coverage
O’Brien Fleming	5.54 (2.93)	0
IPTW	32.50 (7.40)	0.959
S-IPTW	31.28 (7.41)	1.0
Pair IPTW	31.96 (2.98)	1.0
Pair S-IPTW	33.32 (3.96)	1.0

Table 2: Average stopping time and uniform time coverage proportion over 1000 independent trials for the proposed confidence sequences under the unrestricted and pairwise randomization setting. ‘S-IPTW’ denotes the stabilized IPTW. ‘Pair IPTW’ and ‘Pair S-IPTW’ represent the IPTW and stabilized IPTW, respectively, in the pairwise randomization setting.

6. Discussion

In this work, we provide a statistical framework that enables anytime-valid inference in N-of-1 trials for intermediate peeking at the results and analyzing them. We validate our approach in simulation studies and compare it to existing state-of-the-art approaches. The results indicate that recommended methods for

population-level RCTs provide invalid confidence sequences for N-of-1 trials. Our proposed approach, however, results in valid confidence sequences.

This contribution adds to the literature on interim analysis and anytime-valid inference, with a specific focus on N-of-1 trials. Our proposed estimands allow traditional N-of-1 trials to make use of our developed methodology, and allow participants of digital N-of-1 trials to start looking at the results while the trial is ongoing. There is a high need for a user-friendly solution that allows peeking, and we expect that this will enable the further widespread use of N-of-1 trials.

As some concrete examples, we consider two studies that have recently been conducted using the StudyU platform. In a first study run in a hospital setting in Brisbane, Australia, patients designed their personalized trial by selecting their non-pharmacological interventions of interest and tested their effect on reducing their chronic pain. The default study design was 2 weeks baseline, 6 weeks interventions, followed by 2 or more weeks of baseline assessments. There was a large wish of study participants to obtain intermediate feedback on their results and to stop the trial as early as possible. In a follow-up trial with the same population, we are planning to implement anytime-valid inference to allow this feature. As a second example, students in Germany recently ran a series of N-of-1 trials by similarly choosing their own intervention of interest. They evaluated the effect of non-pharmacological interventions on sleep quality and designed a study in an ABAB design with daily measurements over 1 month. Without separate sample size calculations for each personalized trial, the optimal study length was unclear and most trials ended up not showing an effect of the intervention, potentially because the studies were too short. Implementing anytime-valid inference might also be beneficial in this instance, allowing the study to continue to run.

Follow-up work can evaluate our developed methodology across a broader range of scenarios and apply it to a clinical N-of-1 trial. Further, it can be extended to population-level analyses of (adaptive) series of N-of-1 trials.

References

Aurelien Bibaut, Nathan Kallus, and Michael Lindon. Near-optimal non-parametric sequential tests and confidence sequences with possibly dependent observations, 2023. arXiv:2212.14411.

- Ioana Bica, Ahmed M. Alaa, Craig Lambert, and Mihaela van der Schaar. From real-world patient data to individualized treatment effects using machine learning: Current and future methods to address underlying challenges. *Clinical Pharmacology & Therapeutics*, 109(1):87–100, June 2020. doi: 10.1002/cpt.1907. URL <https://doi.org/10.1002/cpt.1907>.
- 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.
- Nediyana Daskalova, Jina Yoon, Yibing Wang, Cintia Araujo, Guillermo Beltran, Nicole Nugent, John McGeary, Joseph Jay Williams, and Jeff Huang. SleepBandits: Guided flexible self-experiments for sleep. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, April 2020. doi: 10.1145/3313831.3376584. URL <https://doi.org/10.1145/3313831.3376584>.
- Eric J. Daza. Causal analysis of self-tracked time series data using a counterfactual framework for n-of-1 trials. *Methods of Information in Medicine*, 57(S 01):e10–e21, 2018. URL <https://doi.org/10.3414/me16-02-0044>.
- Eric J. Daza and Logan Schneider. Model-twin randomization (motr): A monte carlo method for estimating the within-individual average treatment effect using wearable sensors, 2022. URL <https://arxiv.org/abs/2208.00739>.
- Eric Jay Daza. Person as population: A longitudinal view of single-subject causal inference for analyzing self-tracked health data, 2019. URL <https://arxiv.org/abs/1901.03423>.
- David L. Demets and K. K. Gordon Lan. Interim analysis: The alpha spending function approach. *Statistics in Medicine*, 13(13-14):1341–1352, July 1994. doi: 10.1002/sim.4780131308. URL <https://doi.org/10.1002/sim.4780131308>.
- Eustache Diemert, Artem Betlei, Christophe Renaudin, Massih-Reza Amini, Théophile Gregoir, and Thibaud Rahier. A large scale benchmark for individual treatment effect prediction and uplift modeling, 2021. URL <https://arxiv.org/abs/2111.10106>.
- Clive W.J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/1912791>.
- Dae Woong Ham, Iavor Bojinov, Michael Lindon, and Martin Tingley. Design-based confidence sequences for anytime-valid causal inference. *arXiv preprint arXiv:2210.08639*, 2022.
- J. L. Haybittle. Repeated assessment of results in clinical trials of cancer treatment. *The British Journal of Radiology*, 44(526):793–797, October 1971. doi: 10.1259/0007-1285-44-526-793. URL <https://doi.org/10.1259/0007-1285-44-526-793>.
- Keisuke Hirano, Guido W. Imbens, and Geert Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003. doi: <https://doi.org/10.1111/1468-0262.00442>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00442>.
- Daniel G. Horvitz and Donovan J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952. ISSN 01621459. URL <http://www.jstor.org/stable/2280784>.
- 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. doi: 10.1214/20-AOS1991. URL <https://doi.org/10.1214/20-AOS1991>.
- Jaroslav Hájek. Comment on “an essay on the logical foundations of survey sampling” by basu, d. *The foundations of survey sampling*, (236):236, 1971.
- Guido W. Imbens and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, 2015. doi: 10.1017/CBO9781139025751.
- John P.A. Ioannidis, Stephen J W Evans, Peter C Gøtzsche, Robert T O’Neill, Douglas G Altman, Kenneth Schulz, David Moher, and CONSORT Group. Better reporting of harms in randomized trials: An extension of the

- CONSORT statement. *Annals of Internal Medicine*, 141(10):781, November 2004. doi: 10.7326/0003-4819-141-10-200411160-00009. URL <https://doi.org/10.7326/0003-4819-141-10-200411160-00009>.
- Jinkuk Kim, Chunguang Hu, Christelle Moufawad El Achkar, Lauren E. Black, Julie Douville, Austin Larson, Mary K. Pendergast, Sara F. Goldkind, Eunjung A. Lee, Ashley Kuniholm, Aubrie Soucy, Jai Vaze, Nandkishore R. Belur, Kristina Fredriksen, Iva Stojkowska, Alla Tsytsykova, Myriam Armand, Renata L. DiDonato, Jaejoon Choi, Laura Cornelissen, Luis M. Pereira, Erika F. Augustine, Casie A. Genetti, Kira Dies, Brenda Barton, Lucinda Williams, Benjamin D. Goodlett, Bobbie L. Riley, Amy Pasternak, Emily R. Berry, Kelly A. Pflock, Stephen Chu, Chantal Reed, Kimberly Tyndall, Pankaj B. Agrawal, Alan H. Beggs, P. Ellen Grant, David K. Urion, Richard O. Snyder, Susan E. Waisbren, Annapurna Poduri, Peter J. Park, Al Patterson, Alessandra Biffi, Joseph R. Mazzulli, Olaf Bodamer, Charles B. Berde, and Timothy W. Yu. Patient-customized oligonucleotide therapy for a rare genetic disease. *New England Journal of Medicine*, 381(17):1644–1652, October 2019. doi: 10.1056/nejmoa1813279. URL <https://doi.org/10.1056/nejmoa1813279>.
- Stefan Konigorski, Sarah Wernicke, Tamara Slosarek, Alexander M Zenner, Nils Strelow, Darius F Ruether, Florian Henschel, Manisha Manaswini, Fabian Pottbäcker, Jonathan A Edelman, Babajide Owoyele, Matteo Danieletto, Eddy Golden, Micol Zweig, Girish N Nadkarni, and Erwin Böttinger. StudyU: A platform for designing and conducting innovative digital n-of-1 trials. *Journal of Medical Internet Research*, 24(7):e35884, July 2022. doi: 10.2196/35884. URL <https://doi.org/10.2196/35884>.
- Ivana Malenica, Aurelien F. Bibaut, and Mark J. van der Laan. Adaptive sequential design for a single time-series. 2021. arXiv:2102.00102.
- Pavlos Msaouel, Juhee Lee, Jose A. Karam, and Peter F. Thall. A causal framework for making individualized treatment decisions in oncology. *Cancers*, 14(16):3923, August 2022. doi: 10.3390/cancers14163923. URL <https://doi.org/10.3390/cancers14163923>.
- Jane Nikles and Geoffrey Mitchell, editors. *The Essential Guide to N-of-1 Trials in Health*. Springer Netherlands, 2015. doi: 10.1007/978-94-017-7200-6. URL <https://doi.org/10.1007/978-94-017-7200-6>.
- Peter C. O'Brien and Thomas R. Fleming. A multiple testing procedure for clinical trials. *Biometrics*, 35(3):549, September 1979. doi: 10.2307/2530245. URL <https://doi.org/10.2307/2530245>.
- R Peto, M C Pike, P Armitage, N E Breslow, D R Cox, S V Howard, N Mantel, K McPherson, J Peto, and P G Smith. Design and analysis of randomized clinical trials requiring prolonged observation of each patient. i. introduction and design. *British Journal of Cancer*, 34(6):585–612, December 1976. doi: 10.1038/bjc.1976.220. URL <https://doi.org/10.1038/bjc.1976.220>.
- S. J. Pocock. Group sequential methods in the design and analysis of clinical trials. *Biometrika*, 64(2):191–199, August 1977. doi: 10.1093/biomet/64.2.191. URL <https://doi.org/10.1093/biomet/64.2.191>.
- Antony J Porcino, Larissa Shamseer, An-Wen Chan, Richard L Kravitz, Aaron Orkin, Salima Punja, Philippe Ravaud, Christopher H Schmid, and Sunita Vohra. SPIRIT extension and elaboration for n-of-1 trials: SPENT 2019 checklist. *BMJ*, page m122, February 2020. doi: 10.1136/bmj.m122. URL <https://doi.org/10.1136/bmj.m122>.
- Herbert Robbins. Statistical methods related to the law of the iterated logarithm. *The Annals of Mathematical Statistics*, 41(5):1397–1409, 1970. ISSN 00034851. URL <http://www.jstor.org/stable/2239848>.
- James M. Robins, Miguel A. Hernán, and Babette Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, Sep 2000.
- Subodh R. Selukar. *Practical Considerations for Modern Clinical Trials: Three Projects in Clinical Trial Design, Conduct and Analysis*. PhD thesis, University of Washington, 2021. URL <https://digital.lib.washington.edu/researchworks/handle/1773/48208>.
- Caroline Seydel. Personalized medicine is having its day. *Nature Biotechnology*, 41(4):441–446, March 2023. doi:

- 10.1038/s41587-023-01724-9. URL <https://doi.org/10.1038/s41587-023-01724-9>.
- Uri Shalit, Fredrik D. Johansson, and David Sonntag. Estimating individual treatment effect: Generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, page 3076–3085. JMLR.org, 2017.
- Bevan I. Smith, Charles Chimedza, and Jacoba H. Bührmann. Global and individual treatment effects using machine learning methods. *International Journal of Artificial Intelligence in Education*, 30(3):431–458, July 2020. doi: 10.1007/s40593-020-00203-5. URL <https://doi.org/10.1007/s40593-020-00203-5>.
- Volker Strassen. Almost sure behavior of sums of independent random variables and martingales. 1967. URL <https://api.semanticscholar.org/CorpusID:117021204>.
- Sara Taylor, Akane Sano, Craig Ferguson, Akshay Mohan, and Rosalind Picard. QuantifyMe: An open-source automated single-case experimental design platform. *Sensors*, 18(4):1097, April 2018. doi: 10.3390/s18041097. URL <https://doi.org/10.3390/s18041097>.
- Wouter A. C. van Amsterdam, Joost. J. C. Verhoeff, Netanja I. Harlianto, Gijs A. Bartholomeus, Aahlad Manas Puli, Pim A. de Jong, Tim Leiner, Anne S. R. van Lindert, Marinus J. C. Eijkemans, and Rajesh Ranganath. Individual treatment effect estimation in the presence of unobserved confounding using proxies: a cohort study in stage III non-small cell lung cancer. *Scientific Reports*, 12(1), April 2022. doi: 10.1038/s41598-022-09775-9. URL <https://doi.org/10.1038/s41598-022-09775-9>.
- R.C.M. van Kruijsdijk, F.L.J. Visseren, A.C. Dingemans, H.J.M. Groen, J.G. Aerts, Y. van der Graaf, and E.F. Smit. Prediction of treatment effects for individual cancer patients using data from randomized clinical trials. *Annals of Oncology*, 25:iv423, September 2014. doi: 10.1093/annonc/mdu348.21. URL <https://doi.org/10.1093/annonc/mdu348.21>.
- Kenneth Verstraete, Nilakash Das, Iwein Gyselinck, Maarten De Vos, and Wim Janssens. Machine learning for estimating individual treatment effects in randomized controlled trials. *European Respiratory Journal*, 58(suppl 65), 2021. ISSN 0903-1936. doi: 10.1183/13993003.congress-2021.PA3450. URL https://erj.ersjournals.com/content/58/suppl_65/PA3450.
- Jean Ville. *Étude critique de la notion de collectif*. 1939. URL <http://eudml.org/doc/192893>.
- S. Vohra, L. Shamseer, M. Sampson, C. Bukutu, C. H. Schmid, R. Tate, J. Nikles, D. R. Zucker, R. Kravitz, G. Guyatt, D. G. Altman, and D. Moher and. CONSORT extension for reporting n-of-1 trials (CENT) 2015 statement. *BMJ*, 350(may14 17):h1738–h1738, May 2015. doi: 10.1136/bmj.h1738. URL <https://doi.org/10.1136/bmj.h1738>.
- A. Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16(2):117–186, 1945. ISSN 00034851. URL <http://www.jstor.org/stable/2235829>.
- Ian Waudby-Smith, David Arbour, Ritwik Sinha, Edward H. Kennedy, and Aaditya Ramdas. Time-uniform central limit theory and asymptotic confidence sequences, 2023. arXiv:2103.06476.
- Jiabei Yang, Jon A. Steingrímsson, and Christopher H. Schmid. Sample size calculations for n-of-1 trials. *arXiv preprint arXiv:2110.08970*, 2021.
- Alexander M. Zenner, Erwin Böttinger, and Stefan Konigorski. StudyMe: a new mobile app for user-centric n-of-1 trials. *Trials*, 23(1), December 2022. doi: 10.1186/s13063-022-06893-7. URL <https://doi.org/10.1186/s13063-022-06893-7>.
- Weijia Zhang, Thuc Duy Le, Lin Liu, Zhi-Hua Zhou, and Jiuyong Li. Mining heterogeneous causal effects for personalized cancer treatment. *Bioinformatics*, 33(15):2372–2378, March 2017. doi: 10.1093/bioinformatics/btx174. URL <https://doi.org/10.1093/bioinformatics/btx174>.
- Deborah R. Zucker, Robin Ruthazer, and Christopher H. Schmid. Individual (n-of-1) trials can be combined to give population comparative treatment effect estimates: methodologic considerations. *Journal of Clinical Epidemiology*, 63(12):1312–1323, December 2010. doi: 10.1016/j.jclinepi.2010.04.020. URL <https://doi.org/10.1016/j.jclinepi.2010.04.020>.

Appendix A. Proof of Lemma 1

The proof of Lemma 1 follows from classical results in design-based causal inference and results in Appendix A of [Bojinov and Shephard \(2019\)](#).

Proof [Lemma 1] Let $g^0(A_{k,1}) = 1 - g(A_{k,1})$, $\mathbb{1}_1 = \mathbb{1}(A_{k,1} = 1)$, and $\mathbb{1}_0 = \mathbb{1}(A_{k,1} = 0)$. Then,

$$\begin{aligned} & \mathbb{E}(\hat{\psi}_k \mid \mathcal{F}_{k-1}) \\ &= \mathbb{E}\left(\frac{\mathbb{1}_1 f(Y_{k,1:T})}{g(A_{k,1})} - \frac{\mathbb{1}_0 f(Y_{k,1:T})}{g^0(A_{k,1})} \mid \mathcal{F}_{k-1}\right) \\ &= \frac{g(A_{k,1})Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})}{g(A_{k,1})} \\ &\quad - \frac{g^0(A_{k,1})Y_k(\bar{a}_{k-1}, a_k = \mathbf{0})}{g^0(A_{k,1})} \\ &= Y_k(\bar{a}_{k-1}, a_k = \mathbf{1}) - Y_k(\bar{a}_{k-1}, a_k = \mathbf{0}) \\ &= \psi_k(\bar{a}_{k-1}). \end{aligned}$$

The second equality follows as $Y_k(\bar{a}_k) = f(Y_{k,1}(\bar{a}_{k,1}), \dots, Y_{k,T}(\bar{a}_{k,T}))$ and $Y_k(\bar{a}_{k-1}, a_k = \mathbf{1}) = f(Y_{k,1}(\bar{a}_{k-1,1}), \dots, Y_{k,T}(\bar{a}_{k-1,1}))$. Therefore we have that $\mathbb{E}(\hat{\psi}_k - \psi_k(\bar{a}_{k-1}) \mid \mathcal{F}_{k-1}) = 0$ and $\mathbb{E}(|\hat{\psi}_k - \psi_k(\bar{a}_{k-1})|) < \infty$. Note that the errors form a martingale difference sequence, and are uncorrelated through time. We now proceed to derive an upper bound on the variance of the proposed estimator. First, we derive the closed form expression for $\text{Var}(\hat{\psi}_k - \psi_k(\bar{a}_{k-1}) \mid \mathcal{F}_{k-1})$. Note that

$$\begin{aligned} & \text{Var}(\hat{\psi}_k - \psi_k(\bar{a}_{k-1}) \mid \mathcal{F}_{k-1}) \\ &= \mathbb{E}\left[\left(\frac{f(Y_{k,1:T})}{g(A_{k,1})}(\mathbb{1}_1 - g(A_{k,1}))\right.\right. \\ &\quad \left.\left. - \frac{f(Y_{k,1:T})}{g^0(A_{k,1})}(\mathbb{1}_0 - g^0(A_{k,1}))\right)^2 \mid \mathcal{F}_{k-1}\right] \\ &= \mathbb{E}\left[\underbrace{\left(\frac{f(Y_{k,1:T})}{g(A_{k,1})}(\mathbb{1}_1 - g(A_{k,1}))\right)^2}_{\text{Term 1}}\right. \\ &\quad \left. + \underbrace{\left(\frac{f(Y_{k,1:T})}{g^0(A_{k,1})}(\mathbb{1}_0 - g^0(A_{k,1}))\right)^2}_{\text{Term 2}}\right. \\ &\quad \left. - 2 \underbrace{\frac{f(Y_{k,1:T})}{g(A_{k,1})}(\mathbb{1}_1 - g(A_{k,1})) \frac{f(Y_{k,1:T})}{g^0(A_{k,1})}(\mathbb{1}_0 - g^0(A_{k,1}))}_{\text{Term 3}} \mid \mathcal{F}_{k-1}\right]. \end{aligned}$$

We look at each term separately. It follows that Term 1 equals

$$\begin{aligned} & \mathbb{E}\left[\left(\frac{f(Y_{k,1:T})}{g(A_{k,1})}(\mathbb{1}_1 - g(A_{k,1}))\right)^2 \mid \mathcal{F}_{k-1}\right] \\ &= \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})^2}{g(A_{k,1})^2} \mathbb{E}[(\mathbb{1}_1 - g(A_{k,1}))^2 \mid \mathcal{F}_{k-1}] \\ &= \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})^2}{g(A_{k,1})^2} (g(A_{k,1}) - g(A_{k,1})^2). \end{aligned}$$

Similarly, we have that Term 2 equals

$$\begin{aligned} & \mathbb{E}\left[\left(\frac{f(Y_{k,1:T})}{g^0(A_{k,1})}(\mathbb{1}_0 - g^0(A_{k,1}))\right)^2 \mid \mathcal{F}_{k-1}\right] \\ &= \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{0})^2}{g^0(A_{k,1})^2} (g^0(A_{k,1}) - g^0(A_{k,1})^2). \end{aligned}$$

Finally, we obtain the Term 3, which is as follows:

$$\begin{aligned} & \mathbb{E}\left[-2 \frac{f(Y_{k,1:T})}{g(A_{k,1})}(\mathbb{1}_1 - g(A_{k,1}))\right. \\ &\quad \left. \frac{f(Y_{k,1:T})}{g^0(A_{k,1})}(\mathbb{1}_0 - g^0(A_{k,1})) \mid \mathcal{F}_{k-1}\right] \\ &= -2 \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})}{g(A_{k,1})} \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{0})}{g^0(A_{k,1})} \\ &\quad \mathbb{E}[(\mathbb{1}_1 - g(A_{k,1}))(\mathbb{1}_0 - g^0(A_{k,1})) \mid \mathcal{F}_{k-1}] \\ &= 2Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})Y_k(\bar{a}_{k-1}, a_k = \mathbf{0}). \end{aligned}$$

Combining all terms and using the fact that $a^2 + b^2 \geq 2ab$ (where $a = Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})$ and $b = Y_k(\bar{a}_{k-1}, a_k = \mathbf{0})$), we obtain the following upper bound on the variance:

$$\begin{aligned} & \text{Var}(\hat{\psi}_k - \psi_k(\bar{a}_{k-1}) \mid \mathcal{F}_{k-1}) \\ &= \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})^2}{g(A_{k,1})^2} (g(A_{k,1}) - g(A_{k,1})^2) \\ &\quad + \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{0})^2}{g^0(A_{k,1})^2} (g^0(A_{k,1}) - g^0(A_{k,1})^2) \\ &\quad + 2Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})Y_k(\bar{a}_{k-1}, a_k = \mathbf{0}) \\ &= \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})^2}{g(A_{k,1})} + \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{0})^2}{g^0(A_{k,1})} \\ &\quad - (Y_k(\bar{a}_{k-1}, a_k = \mathbf{1}) - Y_k(\bar{a}_{k-1}, a_k = \mathbf{0}))^2 \\ &\leq \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})^2}{g(A_{k,1})} + \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{0})^2}{g^0(A_{k,1})}. \end{aligned}$$

Note that

$$\begin{aligned}
 & \mathbb{E}(\hat{\sigma}_k^2 | \mathcal{F}_{k-1}) \\
 &= \mathbb{E} \left(\frac{\mathbb{1}_1 f(Y_{k,1:T})^2}{g(A_{k,1})^2} + \frac{\mathbb{1}_0 f(Y_{k,1:T})^2}{g^0(A_{k,1})^2} \mid \mathcal{F}_{k-1} \right) \\
 &= \frac{g(A_{k,1}) Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})^2}{g(A_{k,1})^2} \\
 &\quad + \frac{g^0(A_{k,1}) Y_k(\bar{a}_{k-1}, a_k = \mathbf{0})^2}{g^0(A_{k,1})^2} \\
 &= \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{1})^2}{g(A_{k,1})} + \frac{Y_k(\bar{a}_{k-1}, a_k = \mathbf{0})^2}{g^0(A_{k,1})},
 \end{aligned}$$

and it follows that $\text{Var}(\hat{\psi}_k - \psi_k(\bar{a}_{k-1}) \mid \mathcal{F}_{k-1}) \leq \mathbb{E}(\hat{\sigma}_k^2 | \mathcal{F}_{k-1}) = \sigma_k^2$. \blacksquare

Appendix B. Hájek estimator of the AICE

In the following, we study statistical properties of the Hájek estimator of the running average immediate effect over treatment periods, i.e. the AICE.

Lemma 2 (Hájek estimator of the AICE)

Under Assumption 2, it follows that

$$\mathbb{E}(\tilde{\psi}_{\text{AICE}}^k - \psi_{\text{AICE}}^k | \mathcal{F}_{k-1}) = 0$$

and

$$\text{Var}(\tilde{\psi}_{\text{AICE}}^k - \psi_{\text{AICE}}^k | \mathcal{F}_{k-1}) \leq \mathbb{E}(\tilde{\sigma}_{\text{AICE}}^2 | \mathcal{F}_{k-1}).$$

Proof [Lemma 2] Let $g^0(A_{j,1}) = 1 - g(A_{j,1})$, $\mathbb{1}_1 = \mathbb{1}(A_{j,1} = 1)$, and $\mathbb{1}_0 = \mathbb{1}(A_{j,1} = 0)$. Then,

$$\begin{aligned}
 & \mathbb{E}(\tilde{\psi}_{\text{AICE}}^k \mid \mathcal{F}_{k-1}) \\
 &= \mathbb{E} \left(\frac{\sum_{j=1}^k \mathbb{1}_1 f(Y_{j,1:T}) / g(A_{j,1})}{\sum_{j=1}^k \mathbb{1}_1 / g(A_{j,1})} \right. \\
 &\quad \left. - \frac{\sum_{j=1}^k \mathbb{1}_0 f(Y_{j,1:T}) / g^0(A_{j,1})}{\sum_{j=1}^k \mathbb{1}_0 / g^0(A_{j,1})} \mid \mathcal{F}_{k-1} \right) \\
 &= \frac{\sum_{j=1}^k g(A_{j,1}) Y_j(\bar{a}_{j-1}, a_j = \mathbf{1}) / g(A_{j,1})}{\sum_{j=1}^k g(A_{j,1}) / g(A_{j,1})} \\
 &\quad - \frac{\sum_{j=1}^k g^0(A_{j,1}) Y_j(\bar{a}_{j-1}, a_j = \mathbf{0}) / g^0(A_{j,1})}{\sum_{j=1}^k g^0(A_{j,1}) / g^0(A_{j,1})} \\
 &= \frac{1}{k} \sum_{j=1}^k (Y_j(\bar{a}_{j-1}, a_j = \mathbf{1}) - Y_j(\bar{a}_{j-1}, a_j = \mathbf{0})) \\
 &= \psi_{\text{AICE}}^k.
 \end{aligned}$$

Obtaining the upper bound on the variance estimator tracks derivation already presented in the proof of Lemma 1, therefore we omit the details. It follows that

$$\begin{aligned}
 & \text{Var}(\tilde{\psi}_{\text{AICE}}^k - \psi_{\text{AICE}}^k | \mathcal{F}_{k-1}) \\
 &= \frac{\sum_{j=1}^k Y_j(\bar{a}_{j-1}, a_k = \mathbf{1})^2 / g(A_{j,1})}{\sum_{j=1}^k 1 / g(A_{j,1})} \\
 &\quad + \frac{\sum_{j=1}^k Y_j(\bar{a}_{j-1}, a_k = \mathbf{0})^2 / g^0(A_{j,1})}{\sum_{j=1}^k 1 / g^0(A_{j,1})} \\
 &\quad - \left(\sum_{j=1}^k Y_j(\bar{a}_{j-1}, a_k = \mathbf{1}) - \sum_{j=1}^k Y_j(\bar{a}_{j-1}, a_k = \mathbf{0}) \right)^2 \\
 &\leq \frac{\sum_{j=1}^k Y_j(\bar{a}_{j-1}, a_k = \mathbf{1})^2 / g(A_{j,1})}{\sum_{j=1}^k 1 / g(A_{j,1})} \\
 &\quad + \frac{\sum_{j=1}^k Y_j(\bar{a}_{j-1}, a_k = \mathbf{0})^2 / g^0(A_{j,1})}{\sum_{j=1}^k 1 / g^0(A_{j,1})}.
 \end{aligned}$$

Since we have that

$$\begin{aligned}
 & \mathbb{E}(\tilde{\sigma}_{\text{AICE}}^2 | \mathcal{F}_{k-1}) \\
 &= \mathbb{E} \left(\frac{\sum_{j=1}^k \mathbb{1}_1 f(Y_{j,1:T})^2 / g(A_{j,1})^2}{\sum_{j=1}^k \mathbb{1}_1 / g(A_{j,1})^2} \right. \\
 &\quad \left. - \frac{\sum_{j=1}^k \mathbb{1}_0 f(Y_{j,1:T})^2 / g^0(A_{j,1})^2}{\sum_{j=1}^k \mathbb{1}_0 / g^0(A_{j,1})^2} \mid \mathcal{F}_{k-1} \right) \\
 &= \frac{\sum_{j=1}^k Y_j(\bar{a}_{j-1}, a_k = \mathbf{1})^2 / g(A_{j,1})}{\sum_{j=1}^k 1 / g(A_{j,1})} \\
 &\quad + \frac{\sum_{j=1}^k Y_j(\bar{a}_{j-1}, a_k = \mathbf{0})^2 / g^0(A_{j,1})}{\sum_{j=1}^k 1 / g^0(A_{j,1})},
 \end{aligned}$$

it follows that $\text{Var}(\tilde{\psi}_{\text{AICE}}^k - \psi_{\text{AICE}}^k | \mathcal{F}_{k-1}) \leq \mathbb{E}(\tilde{\sigma}_{\text{AICE}}^2 | \mathcal{F}_{k-1})$. \blacksquare

Appendix C. Proof of Theorem 1

The proof of Theorem 1 can be obtained as in Ham et al. (2022) (Theorem 5.2) and Waudby-Smith et al. (2023) (Theorem 2.3), adapted to the N-of-1 setting and proposed estimators. First, we state the Ville's maximal inequality in Lemma 3, as it is used in the proof (Ville, 1939). We also use Theorem 4.4 of Strassen (1967), adapted to our setting (same adaptation as in Ham et al. (2022)).

Lemma 3 (Ville's Maximal Inequality) *Let M_t denote a non-negative martingale with respect to a filtration \mathcal{F} . Then,*

$$P(\exists t \in \mathbb{N}_0 : M_t \geq \frac{1}{\alpha}) \leq \alpha M_0.$$

For $M_0 = 1$, we obtain the uniform type-1 error guarantee.

Proof [Theorem 1] Let X_k denote an i.i.d standard normal random variable. As in Robbins (1970), we construct a non-negative martingale with initial value one for any $\lambda \in \mathbb{R}$ with respect to the filtration:

$$M_k(\lambda) := \exp\left(\sum_{j=1}^k (\lambda \sigma_j X_j - \frac{\lambda^2 \sigma_j^2}{2})\right),$$

where σ_j^2 is the true variance $\sigma_j^2 := \frac{Y_j(\bar{a}_{j-1}, \mathbf{1}_T)^2}{g(A_{j,1})} + \frac{Y_j(\bar{a}_{j-1}, \mathbf{0}_T)^2}{1-g(A_{j,1})}$. By Robbins (1970), for any probability distribution $F(\lambda)$ on \mathbb{R} , $\int_{\lambda} M_k(\lambda) dF(\lambda)$ is still a non-negative martingale with initial value one. For a normal mixing distribution with mean λ and variance η^2 , the resulting martingale M_k is written as

$$\frac{1}{\eta\sqrt{2\pi}} \int_{\lambda} \exp\left(\sum_{j=1}^k (\lambda \sigma_j X_j - \frac{\lambda^2 \sigma_j^2}{2})\right) \exp\left(\frac{-\lambda^2}{2\eta^2}\right) d\lambda.$$

Let $Z_k = \sum_{j=1}^k \sigma_j X_j$ and $\bar{\sigma}_k^2 = 1/k \sum_{j=1}^k \sigma_j^2$. Then,

$$\begin{aligned} M_k &= \frac{1}{\eta\sqrt{2\pi}} \int_{\lambda} \exp\left(\lambda Z_k - \frac{k\lambda^2 \bar{\sigma}_k^2}{2}\right) \exp\left(\frac{-\lambda^2}{2\eta^2}\right) d\lambda \\ &= \frac{1}{\eta\sqrt{2\pi}} \int_{\lambda} \exp\left(\frac{2\eta^2 \lambda Z_k - \lambda^2(1 + k\eta^2 \bar{\sigma}_k^2)}{2\eta^2}\right) d\lambda \\ &= \frac{1}{\eta\sqrt{2\pi}} \int_{\lambda} \exp\left(\frac{2\lambda b - \lambda^2 a}{2\eta^2}\right), \end{aligned}$$

where $a = k\eta^2 \bar{\sigma}_k^2 + 1$ and $b = \eta^2 Z_k$. We proceed to complete the square by adding and subtracting $(b/a)^2$:

$$\begin{aligned} \exp\left(\frac{2\lambda b - \lambda^2 a}{2\eta^2}\right) &= \exp\left(\frac{b}{\eta^2} \lambda - \frac{a}{2\eta^2} \lambda^2\right) \\ &= \exp\left(-\frac{a}{2\eta^2} \left(\lambda^2 - \frac{2b}{a} \lambda\right)\right) \\ &= \exp\left(-\frac{a}{2\eta^2} \left(\lambda - \frac{b}{a}\right)^2 + \left(\frac{b^2}{2a\eta^2}\right)\right). \end{aligned}$$

Substituting back into M_k , and replacing a and b for their values, we have that

$$\begin{aligned} M_k &= \frac{1}{\sqrt{2\pi\eta^2/a}} \int_{\lambda} \exp\left(-\frac{a}{2\eta^2} \left(\lambda - \frac{b}{a}\right)^2\right) d\lambda \\ &\quad \frac{1}{\sqrt{a}} \exp\left(\frac{b^2}{2a\eta^2}\right) \\ &= \frac{1}{\sqrt{a}} \exp\left(\frac{b^2}{2a\eta^2}\right) \\ &= \frac{1}{\sqrt{k\eta^2 \bar{\sigma}_k^2 + 1}} \exp\left(\frac{\eta^2 Z_k^2}{2(k\eta^2 \bar{\sigma}_k^2 + 1)}\right). \end{aligned}$$

By Robbins (1970), we know that M_k is a non-negative martingale. Applying Lemma 3, it follows that

$$P(\forall k \geq 1, M_k < \frac{1}{\alpha}) \geq 1 - \alpha.$$

By basic algebra manipulation, we can show that $\frac{1}{\sqrt{k\eta^2 \bar{\sigma}_k^2 + 1}} \exp\left(\frac{\eta^2 Z_k^2}{2(k\eta^2 \bar{\sigma}_k^2 + 1)}\right) < \frac{1}{\alpha}$ simplifies to

$$P\left(\forall k \geq 1, \left|\frac{1}{k} Z_k\right| < C_k\right) \geq 1 - \alpha, \quad (6)$$

where $C_k = \sqrt{\frac{2(k\eta^2 \bar{\sigma}_k^2 + 1)}{k^2 \eta^2} \log\left(\frac{\sqrt{k\eta^2 \bar{\sigma}_k^2 + 1}}{\alpha}\right)}$.

Let $u_k = \hat{\psi}_k - \psi_k(\bar{a}_{k-1})$. In Appendix A, we discuss that $\mathbb{E}(\hat{\psi}_k - \psi_k(\bar{a}_{k-1}) | \mathcal{F}_{k-1}) = 0$ and $\mathbb{E}(|\hat{\psi}_k - \psi_k(\bar{a}_{k-1})|) < \infty$, so $(u_k)_k$ is a martingale difference sequence with respect to \mathcal{F}_{k-1} (and thus, uncorrelated through time). We note that the same applies for martingale difference sequence formed by the Hájek estimator, as shown in Lemma 2. We proceed to utilize the strong approximation theorem from Strassen (1967) (Theorem 4.4), which applies to martingale difference sequences of the form $\mathbb{E}(X_n | \sigma(X_1, \dots, X_{n-1})) = 0$ (notation as in Strassen (1967)). In the Appendix D, Ham et al. (2022) edit Theorem 4.4 in order to accommodate martingale difference sequences we also have, of the form u_k . The edited formulation of Theorem 4.4 follows due to Assumption 3, and results in the following approximation:

$$\frac{1}{k} \sum_{j=1}^k u_j = \frac{1}{k} \sum_{j=1}^k \sigma_j X_j + o\left(\frac{\tilde{S}_k^{3/8} \log \tilde{S}_k}{k}\right) a.s.$$

where we remind that $\tilde{S}_k = \sum_{j=1}^k \sigma_j^2$. Substituting back into Equation (6), we have that

$$P(\forall k \geq 1, \left| \frac{1}{k} \sum_{j=1}^k u_j \right| < C_k + o\left(\frac{\tilde{S}_k^{3/8} \log \tilde{S}_k}{k}\right)) \geq 1 - \alpha. \quad (7)$$

Note that Equation (7) is the non-asymptotic confidence width, denoted as $(I'_k)_{k=1}^K$ in Definition 4. By Assumption 4, the asymptotic confidence width is

$$P\left(\forall k \geq 1, \left| \frac{1}{k} \sum_{j=1}^k u_j \right| < C_k\right) \geq 1 - \alpha. \quad (8)$$

Also by Assumption 4, the $(1 - \alpha)$ asymptotic confidence sequence for the target parameter sequence $(\psi_{\text{AICE}}^k)_k$ is then

$$\frac{1}{k} \sum_{j=1}^k \hat{\psi}_j \pm \sqrt{\frac{2(k\eta^2\tilde{\sigma}_k^2 + 1)}{k^2\eta^2} \log\left(\frac{\sqrt{k\eta^2\tilde{\sigma}_k^2 + 1}}{\alpha}\right)}. \quad (9)$$

Finally, note that the $(1 - \alpha)$ asymptotic confidence sequence in Equation (9) relies on the true variance. In step 3 of Appendix D in Ham et al. (2022) and Appendix A.2 of Waudby-Smith et al. (2023), they show that under further assumption of $\tilde{\sigma}_k^2 \xrightarrow{a.s.} \bar{\sigma}_k^2$, with $\bar{\sigma}_k^2 = 1/k \sum_{j=1}^k \hat{\sigma}_j^2$, we have that $C_k = \sqrt{\frac{2(k\eta^2\bar{\sigma}_k^2 + 1)}{k^2\eta^2} \log\left(\frac{\sqrt{k\eta^2\bar{\sigma}_k^2 + 1}}{\alpha}\right)}$. We include the argument for clarity in what follows. Let $\bar{\sigma}_k^2 - \tilde{\sigma}_k^2 = o(\bar{\sigma}_k^2)$. Then we have that

$$\begin{aligned} & \sqrt{\frac{2(k\eta^2\tilde{\sigma}_k^2 + 1)}{k^2\eta^2} \log\left(\frac{\sqrt{k\eta^2\tilde{\sigma}_k^2 + 1}}{\alpha}\right)} \\ &= \sqrt{\frac{2(k\eta^2(\tilde{\sigma}_k^2 + o(\tilde{\sigma}_k^2)) + 1)}{k^2\eta^2} \log\left(\frac{\sqrt{k\eta^2(\tilde{\sigma}_k^2 + o(\tilde{\sigma}_k^2)) + 1}}{\alpha}\right)} \\ &= \sqrt{\frac{k\eta^2(\tilde{\sigma}_k^2 + o(\tilde{\sigma}_k^2)) + 1}{k^2\eta^2} \log\left(\frac{k\eta^2(\tilde{\sigma}_k^2 + o(\tilde{\sigma}_k^2)) + 1}{\alpha^2}\right)} \\ &= \underbrace{\sqrt{\left(\frac{k\eta^2\tilde{\sigma}_k^2 + 1}{k^2\eta^2} + o(\tilde{\sigma}_k^2/k)\right)}}_{\text{Term 1}} \underbrace{\sqrt{\left(\frac{k\eta^2\tilde{\sigma}_k^2 + o(k\tilde{\sigma}_k^2) + 1}{\alpha^2}\right)}}_{\text{Term 2}}. \end{aligned}$$

We focus on Term 2 in the rest of the proof.

$$\begin{aligned} & \log\left(\frac{k\eta^2\tilde{\sigma}_k^2 + o(k\tilde{\sigma}_k^2) + 1}{\alpha^2}\right) \\ &= \log\left(\frac{k\eta^2\tilde{\sigma}_k^2 + 1}{\alpha^2}\right) + \log(1 + o(1)) \\ &= \log\left(\frac{k\eta^2\tilde{\sigma}_k^2 + 1}{\alpha^2}\right) + o(1). \end{aligned}$$

Therefore, valid $(1 - \alpha)$ asymptotic confidence sequence for $(\psi_{\text{AICE}}^k)_k$ is

$$\frac{1}{k} \sum_{j=1}^k \hat{\psi}_j \pm \frac{1}{k} \sqrt{\frac{\eta^2 \sum_{j=1}^k \hat{\sigma}_j^2 + 1}{\eta^2} \log\left(\frac{\eta^2 \sum_{j=1}^k \hat{\sigma}_j^2 + 1}{\alpha^2}\right)}.$$

■

Appendix D. Supplementary Figures




Figure 4: Anytime valid confidence intervals of AICE obtained by IPTW (left) and stabilized IPTW (right), respectively, in a single run at $\alpha = 0.05$. The dashed line represents the zero (null) line. Top row: unrestricted randomization scheme. Bottom row: pairwise randomization scheme.