

Randomization Inference When N Equals One

Tengyuan Liang¹ and Benjamin Recht²

¹University of Chicago

²University of California, Berkeley

October 25, 2023

Abstract

N-of-1 experiments, where a unit serves as its own control and treatment in different time windows, have been used in certain medical contexts for decades. However, due to effects that accumulate over long time windows and interventions that have complex evolution, a lack of robust inference tools has limited the widespread applicability of such N-of-1 designs. This work combines techniques from experiment design in causal inference and system identification from control theory to provide such an inference framework. We derive a model of the dynamic interference effect that arises in linear time-invariant dynamical systems. We show that a family of causal estimands analogous to those studied in potential outcomes are estimable via a standard estimator derived from the method of moments. We derive formulae for higher moments of this estimator and describe conditions under which N-of-1 designs may provide faster ways to estimate the effects of interventions in dynamical systems. We also provide conditions under which our estimator is asymptotically normal and derive valid confidence intervals for this setting.

Keywords— causal inference, system identification, potential outcomes, interference, time series.

1 Introduction

Randomized experiments are arguably the most important contribution of twentieth-century statistics, providing a program to establish the value of interventions to populations of individuals. By randomly assigning individuals to treatment and control, statistical analysis can simultaneously determine the magnitude of intervention effects and eliminate confounding explanations. However, a significant drawback of randomized experiments is that they yield conclusions about populations rather than individuals. When treatments present heterogeneous effects, a randomized experiment cannot inform individuals whether a treatment can work for them.

Is it possible to design experiments for individuals? *N-of-1 trials* attempt to solve this problem by having individuals apply randomized treatments in different time windows. The individual becomes the treatment and the control. N-of-1 trials are within-patient randomized controlled trials, where the treatment is randomized at every time period, with the subject potentially crossing over between treatment and control at each phase.

In medicine, the formalized N-of-1 trial is nearly as old as the randomized clinical trial itself. In 1950, only a few years after the famous randomized trials of Streptomycin, Quin and colleagues performed randomized within subject trials of an arthritis reduction agent where patients received a random treatment at multiple doctor visits [19]. D. D. Reid explained Latin Square designs for N-of-1 trials and explored these for testing

the effectiveness of laxatives [21] and arthritis ointments [20]. Snell and Armitage tested heroin as a cough suppressant with N-of-1 designs [26]. Bradford Hill even devotes a lengthy part of the Clinical Trials chapter (pages 260-263) in his *Principles of Medical Statistics* to the design of N-of-1 trials [10].

Since the 1950s, though never a dominant experimental design, N-of-1 trials have continued to find application. Louis et al. reported 28 self-controlled studies in the New England Journal of Medicine in 1978 and 1979 [13]. Guyatt et al. argued for using N-of-1 trials as part of general clinical practice with patients, helping individuals find the best treatment for them [8]. N-of-1 trials have been recently used successfully to test the side effects of statin treatments [30].

In the online experimentation community, N-of-1 experiments are often called switchback or interleaving designs. Netflix and LinkedIn use N-of-1 designs to test the time-varying effects of their recommendation and matching algorithms. For example, [3] reports interleaving online experiments where a client is randomly exposed to different app experiences for 30-minute intervals. N-of-1 designs have the potential to be tailored to personalized experiments to improve the experience of individual customers. Moreover, N-of-1 designs enable data scientists to mitigate interference effects from multiple AB tests running concurrently [27].

If the effect is immediate and transient, standard tools from randomized experiments can be used to estimate its magnitude and significance in N-of-1 trials. However, when the effect of an intervention persists over time, it becomes challenging to draw proper inferences. Though recent estimators have attempted otherwise, most N-of-1 designs try to reduce the problem to a randomized experiment where each time period is independent. In such a case, the treatment must take effect and cease rapidly. These designs also must assume an individual’s baseline is roughly constant. For these reasons, N-of-1 designs have primarily focused on chronic conditions treatable by acute medications.

Can we do better? In this work, we aim to combine techniques from experiment design in causal inference with system identification from control theory to provide a general inference framework for N-of-1 experiments. In what follows, we describe the foundational pillars of our framework and a view of how they might be combined for inferring effects in N-of-1 experiments. We then present a formal interference model for temporal effects drawn from the literature on dynamical systems. We define a family of estimands of causal effects of these models and then present a reasonable, unbiased estimator for this family. We demonstrate that this estimator is asymptotically normal and provide a plug-in estimator for its variance. Together, these enable the construction of valid confidence intervals. Along the way, we discuss the inferential challenges of this framework and why our approach, though more general than previous methods, remains limited.

1.1 Foundational Related Work

Our work draws and expands upon the literature on causal inference and system identification. In this section, we elaborate on the background and how we will attempt to reconcile these two different lines of research.

Causal inference The potential outcomes framework [15, 22] offers a systematic and rigorous approach to studying causal relationships in experimental and observational studies. In the simplest setting, a *unit* i has two potential outcomes under an intervention (also called a treatment). The observed outcome of unit i is denoted by y_i , and we distinguish the outcome under no treatment as $y_i(0)$ and under treatment as $y_i(1)$. If we let x_i denote the binary indicator equal to 1 when treatment is applied and 0 when treatment is not applied, then we have the expression

$$y_i := y_i(1) \cdot x_i + y_i(0) \cdot (1 - x_i)$$

This formula holds regardless of whether y_i are real-valued or binary outcomes.

Though this formula looks like it is tautological, it makes some hidden assumptions. In particular, this model assumes that the treatment applied to unit i only impacts unit i . This is called the *Stable Unit Treatment Values Assumption* (SUTVA). This assumption does not hold for within-unit trials. If a treatment

is applied in succession to a unit, care must be taken to ensure that the previous treatment does not influence the outcomes of future treatments.

When the x_i can be intentionally assigned, random assignment enables estimation of the average outcome under treatment and control. That is, if we are interested in estimating the quantity

$$\tau := \frac{1}{n} \sum_{i=1}^n y_i(1) - y_i(0),$$

we can do an experiment where x_i is assigned at random and then estimate τ from the observed outcomes.

Assume, for simplicity, that the x_i are assigned as independent Bernoulli coin flips. Denote the centered treatment variable $z_i = 2x_i - 1 \in \{\pm 1\}$. Then the *Horvitz-Thompson estimator* $\widehat{\tau}_{\text{HT}}$

$$\widehat{\tau}_{\text{HT}} := \frac{2}{n} \sum_{i=1}^n z_i \cdot y_i$$

is an unbiased estimator of the average treatment effect. When SUTVA applies, inference about the significance of causal effects is also elegant. $\widehat{\tau}_{\text{HT}}$ is asymptotically normal, and confidence intervals can be constructed using upper bounds on the variance of $\widehat{\tau}_{\text{HT}}$ (see, for example, [11]).

Beyond SUTVA For causal questions about time series data, SUTVA rarely applies [25, 7]. Suppose now that each unit is indexed by time. Suppose we have a sequence of interventions at each time x_1, x_2, \dots, x_T . Even if these interventions are binary-valued, the outcome at time t , y_t could potentially depend on all of the outcomes that have occurred before time t . Namely, y_t is potentially a function of x_1, \dots, x_t . We could represent y_t then as a sum over all possible interactions. That is,

$$y_t = \sum_{S: S \subseteq \{0, 1, \dots, t\}} \alpha_{t,S} \cdot \prod_{i \in S} z_i, \quad (1.1)$$

where again, $z_i = 2x_i - 1$. The standard Potential Outcomes framework is the special case when the coefficients $\alpha_{t,S} = 0$ for all $S \neq \emptyset$ or $\{t\}$. This basis expansion is fully general as it is the Fourier expansion of $y_t(x_0, x_1, \dots, x_t)$ [16]. But the expansion has a number of terms that grows exponentially with t , preventing any meaningful statistical analysis.

Since a full representation is intractable, more restricted interactions have been studied. The Granger-Sims causality framework [7, 25, 1] tests correlations in x 's that can explain y . For example, in its simplest form, Granger Sims studies how well linear expansions of the form

$$y_t = \alpha_{\emptyset} + \alpha_0 \cdot x_t + \alpha_1 \cdot x_{t-1} + \dots + \alpha_t \cdot x_0 + \text{error}_t$$

predict outcomes. Granger causality assumes *time-invariance* to provide practically useful answers to understand when time series can predict each other.

In this paper, we study this time-invariant model, under the assumption that the x_i are chosen inter-ventionally as part of an experiment.

System identification In control theory, modeling the input-output behavior of a dynamical system is a first step in designing a feedback policy. Such dynamical systems are assumed to obey dynamical equations of the form

$$s_{t+1} = f_t(s_t, x_t, \varepsilon_t), \quad y_t = h_t(s_t, x_t, \varepsilon_t).$$

x_t in this context is a controllable input to the dynamical system. y_t is the measured output and ε_t is an exogenous noise process. s_t is the *state* of the dynamical system. The state is the variable that makes future outputs conditionally independent of past outputs.

The goal in system identification is to find appropriate estimates of the function f_t and h_t so that inputs x_t can be planned to steer y_t to desired values.

Just as we have described above, the identification problem is only tractable if the class of models for f_t and h_t is simple enough. Several reduced model classes are useful for representing dynamical systems. One of the most fundamental models assumes that the functions f_t and h_t are linear (but might vary over time). In such a case, we can write the input-output map as

$$y_t = \sum_{k \leq t} G_{tk} x_k + e_t$$

where G_{tk} are scalars and e_t are linear functions of the ε_k variables and s_1 .

Under a further restriction, we can assume the system is linear and *time-invariant* so that $f_t = f$ and $h_t = h$ for all t . In this case, y_t can be written as a linear convolution of x_k and a sequence g_k :

$$y_t = \sum_{k \leq t} g_{t-k} x_k + e_t.$$

This model is the same as the 1-step Granger Causality model described in the previous section.

Identifying linear time-invariant systems has a rich literature (see the textbook by Ljung [12]). Importantly, the choice of x_k matters as this input design must make all g_k of interest identifiable. The most popular choice in theory and practice chooses x_k to be an i.i.d. sequence of zero mean random variables. In practice, it suffices that these variables take on two values, such as 1 and -1 [14, 29, 17]. When restricted to such signed inputs, the linear system identification problem looks like a Potential Outcomes estimation problem under a linear interference model.

Though there is extensive literature on system identification, only recently have researchers determined statistical error rates for the estimation of the sequence g_k from a single input sequence x_k . There are numerous reasons for this. First, since system identification of engineered systems is typically done in a controlled setting, sample complexity is not typically a constraint on engineers. Second, the statistical methods for system identification lean on random matrix results that have only been recently derived.

The analysis of Oymak and Ozay [18] shows that the parameters g_k can be estimated via least squares when x_t and e_t are zero-mean Gaussian random variables and $s_1 = 0$. They further assume that g_k obeys a decay condition that $|g_k| \leq C a^k$ for a constant C and $a < 1$. They show that the rate of error in the estimate \hat{g}_k satisfies

$$\sum_{k=1}^{\infty} |g_k - \hat{g}_k|^2 = O\left(\frac{1}{(1-a)^2 T}\right)$$

where T is the number of observed samples.

Later work by Simchowicz et al. [24] removed the dependence on a in exchange for another property of g called *phase rank*, but some restriction on the form of g is always required to estimate g . Similarly, work by Bakshi et al. [2] provides tighter bounds on this estimation error, but assumes the error signal is random and the input signal is zero mean. Moreover, the current system identification literature does not provide provable means of data-driven uncertainty quantification, a crucial component of causal inference.

A synthesis of causal inference and system identification The work in this paper builds on these three pillars to synthesize an inference framework for N-of-1 experiments. First, as we develop in the next section, we focus on the linear time-invariant outcome model studied in Granger-Sims causality, and in linear system identification

$$y_t = \sum_{s=0}^t x_s g_{t-s} + e_t. \tag{1.2}$$

This model generalizes the Neyman-Rubin framework to allow for a particular functional form of linear interference that is inspired by linear dynamical systems. In the next section, we develop a generalization of the Horvitz-Thompson estimator for linear estimands of the form

$$\tau(q) := \langle g, q \rangle.$$

In the case that q is the first standard basis vector, the generalized estimator will precisely be the Horvitz-Thompson estimator.

Unlike in the Granger-Sims and linear system identification frameworks, we do not assume the error signal e_t is random in (1.2). We will focus on estimating various linear functionals of the coefficients $g \in \mathbb{R}^T$. Focusing on linear estimands rather than the coefficients themselves should simplify the problem. We note, however, that results for this simplification cannot be derived from the prior art.

Additionally, in this work, we are interested in estimating g_k when e_t is a non-random signal. Simchowitz et al. [24] were the first to study estimation of g when e_t was non-random, and estimation under non-random conditions has been used in the study of online control [9]. However, none of these works provided specific forms of the error, nor did they quantify the variance sufficiently accurately for the construction of asymptotically-sharp confidence intervals.

However, the explicit form of the error is not provided. There are several practically relevant situations where such non-random noise arises. For instance, suppose we restrict our attention to input sequences that are constrained to be nonnegative. This departs from the zero-mean inputs in Granger-Sims and system identification. Non-zero-mean inputs can be modeled as zero-mean inputs plus a new, non-random error term.

Unlike system identification, our work is interested in estimating confidence intervals on linear functionals of g_k from finite experiments. None of the prior work has investigated such uncertainty quantification. In particular, we derive novel asymptotic normality guarantees for our estimator, which is closely related to the least-squares estimator of Oymak and Ozay [18] and the method-of-moments estimator of Bakshi et al. [2]. The proof of asymptotic normality is more delicate than what is typical of Horvitz-Thompson estimators. In particular, the asymptotic normality requires an intricate calculation balancing of eighth moments of the intervention variables.

Finally, we close this discussion with a discussion of recent complementary work in the context of online experiments. Bojinov et al. explored experiments on time series using the potential outcome framework for switchback experiments [4]. There, the authors consider a general time-varying outcome model with only short-term dependence on the experiment variable. Their experimental estimand focuses on average short-term effects, whereas our estimand considers any linear functional of the long-term response. The salient distinction from the present work is the focus on general, short-term effects versus linear, long-term effects. In this way, our papers complement each other and point to the potential of a general framework that can effectively handle general, long-term interference effects in N-of-1 trials.

2 Convolutions

The outcomes in linear time-invariant models can be represented as a convolution between an intervention sequence and a parameter vector known as an *impulse response*. Before expanding on our statistical theory, we first review the basic notation and theory of convolution. This notation will simplify many of the formulas, and eliminate the need to track subscripts.

Let u_t and v_t be two signals with indices in the nonnegative integers. The convolution of u_t and v_t is the signal

$$(u * v)_t = \sum_{s=0}^t u_s v_{t-s}$$

Note how this expression allows us to compactly rewrite the models of Granger-Sims causality and LTI systems.

Convolution is commutative ($u * v = v * u$) and linear. We can write the operator that takes v to $u * v$ in matrix form

$$u * v = \mathcal{T}_u v$$

where \mathcal{T}_u is the Toeplitz matrix

$$\mathcal{T}_u = \begin{bmatrix} u_0 & 0 & 0 & 0 & \dots & 0 \\ u_1 & u_0 & 0 & 0 & \dots & 0 \\ u_2 & u_1 & u_0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ u_{T-1} & u_{T-2} & u_{T-3} & \dots & \dots & u_0 \end{bmatrix}.$$

Note that we must have $\mathcal{T}_u v = \mathcal{T}_v u$.

The adjoint operator of \mathcal{T} maps matrices to vectors with components

$$\mathcal{T}^*(M)_s = \sum_t M_{t,t-s}.$$

Moreover, we have the composition $\mathcal{T}^*(\mathcal{T}_u)_s = (T - s + 1)u_s$.

In our theoretical analysis, we use *circular extensions* of signals that result in tractable closed-form expressions for many of the moments in our experiment designs. For two signals u and v of length T , the *circular convolution* is defined as

$$(u \otimes v)_t = \sum_{s=0}^{T-1} u_s v_{t-s \pmod{T}},$$

where the circular extension notation \pmod{T} means that for $s \in [0, t]$, $t - s \pmod{T} = t - s$ and for $s \in (t, T - 1]$, $t - s \pmod{T} = T + t - s$. Note here that the summation extends over all possible indices s , not just those that have values less than or equal to t . In matrix form, circular convolution corresponds to multiplication by a circulant matrix

$$u \otimes v = \mathcal{C}_u v$$

where \mathcal{C}_u is the matrix

$$\mathcal{C}_u = \begin{bmatrix} u_0 & u_{T-1} & u_{T-2} & u_{T-3} & \dots & u_1 \\ u_1 & u_0 & u_{T-1} & u_{T-2} & \dots & u_2 \\ u_2 & u_1 & u_0 & u_{T-1} & \dots & u_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ u_{T-1} & u_{T-2} & u_{T-3} & \dots & \dots & u_0 \end{bmatrix}.$$

For a variety of reasons, computation with circular convolutions is more elegant than with linear convolutions. For example, the Discrete Fourier Transform of a circular convolution is the product of the Discrete Fourier Transforms of the individual signals. Such properties will be used in depth in the theoretical analysis.

For the ease of writing subscripts in equations, we introduce the circular extension notation $v_t^\circ = v_{t \pmod{T}}$ for possible negative integer t . With this notation,

$$(u \otimes v)_t = \sum_{s=0}^{T-1} u_s v_{t-s}^\circ. \tag{2.1}$$

3 Estimands and Estimators of Time-invariant Treatment Effects

Suppose we have a scalar outcome sequence y_t that depends on a scalar treatment x_t via a linear dynamical system

$$y_t = \sum_{s=0}^t g_s x_{t-s} + e_t = (g * x)_t + e_t. \tag{3.1}$$

Here g is the *impulse response* of the linear system. For simplicity, we assume that the x_t are binary-valued, but we expect extensions to real-valued inputs to be relatively straightforward.

e_t is an exogenous signal that is not affected by the treatment e_t . But we are interested in analyzing e_t that are far from random. As a motivating example, suppose two treatments have treatment effects g_0 and g_1 . We are interested in which leads to overall better outcomes. In this case, $g = g_1 - g_0$ and $e = g_0$. These

signals are tightly correlated. Hence, we model e as an *adversarial error* oblivious to the randomization x . The error could be non-i.i.d., non-stationary, and even systematic, as long as it is oblivious to input x .

Which outcome properties might we be interested in? We could study the *average treatment effect*, which compares outcomes when x_t is equal to either all ones or all zeros:

$$\tau_{\text{ATE}} = \frac{1}{T} \sum_{t=1}^T y_t(1) - y_t(0)$$

Using (3.1), we can write τ_{ATE} in terms of the impulse response

$$\tau_{\text{ATE}} = \sum_{t=1}^T \frac{(T-t+1)}{T} g_t.$$

Hence, for full generality, we will study linear estimands of the form $\tau(q) = \langle D_T q, g \rangle$ where q is a vector and D_T is the diagonal matrix with t -th entry $\frac{T-t+1}{T}$. With this definition, $\tau_{\text{ATE}} = \tau(q)$ when $q = \mathbf{1}$.

For these designs, different estimands can highlight varied intervention effects. In a standard N-of-1 trial, the estimand would compare the effects of treatments at the end of each dosing period. That is, if a particular treatment is given for a week at a time, the estimand will measure outcomes at the end of each week. Standard N-of-1 designs use long periods to ensure that the interference between the effects of the different treatments have washed out before an outcome is measured. If a treatment has spillover, then measurements early in a treatment period will be influenced by the prior period's treatment.

But what if we are interested in the average treatment effect at the resolution of a day? We can likely measure outcomes each day, but can we measure the average differences? To get a feel for how these two estimands can be different, consider the following two treatments with two periods:

$$y_1^A = x_1, \quad y_2^A = x_1 + x_2,$$

and

$$y_1^B = 0, \quad y_2^B = 2x_1.$$

Both of these treatments correspond to linear time-invariant systems with $g^A = (1, 1)$ and $g^B = (0, 2)$.

For both treatments, the outcome if the treatment is taken for two time periods is 2. However, the average outcome for Treatment A is 1.5 and the average outcome for Treatment B is 1. If, for example, these treatments are pain medications, then Treatment A is preferable to Treatment B as the immediate relief is valuable. We thus aim to allow for a variety of linear estimands to capture the effect most aligned with preferred outcomes.

3.1 Estimation from Random Design

Consider the randomized experiment where we assign the input at time t to be zero or one with equal probability, independently at each $t \in [T]$. How can we estimate the treatment effect $\tau(q)$ from the observed outcomes of an experiment?

We now fix the notation for the remainder of the document to distinguish between random and nonrandom signals. We will use boldface font for random vectors to emphasize the stochastic nature of experimentation. In what follows, we use the normal font, say $x \in \mathbb{R}^T$ and $H \in \mathbb{R}^{T \times T}$ to denote deterministic vectors and matrices, and boldface say $\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{W}$ to denote vectors that are random due to experimental design. The lone exception is that we use $\mathbf{1}$ to denote the all ones vector and avoid confusion with the scalar 1.

With our fixed notation, the linear convolution model is written as

$$\mathbf{y} = \mathbf{x} * g + e \in \mathbb{R}^T, \tag{3.2}$$

and $g, \mathbf{x}, \mathbf{x} * g$ and e are all of length T with $\mathbf{x}_t \in \{0, 1\}, t \in [T]$.

A natural generalization of the Horvitz-Thompson estimator is the method-of-moments estimator discussed above.

$$\hat{\tau}(q) = \frac{1}{T} \langle (2\mathbf{x} - \mathbf{1}) * q, 2\mathbf{y} \rangle. \quad (3.3)$$

This choice is natural in the following sense. Suppose we restrict our attention to unbiased linear estimators $\langle w, \mathbf{y} \rangle$. Then unbiasedness is the linear constraint

$$\mathbb{E}_{\mathbf{x}}[\langle w, \mathbf{y} \rangle] = \langle D_T q, g \rangle.$$

Setting w as the Moore-Penrose pseudoinverse of this system of equations yields the estimator (3.3).

As we will do several times in what follows, we note that $\hat{\tau}(q)$ is a quadratic function of T independent Rademacher random variables. Throughout we will need to compute moments of such Rademacher chaos, and we provide explicit formulas for their first, second, and fourth moments in Appendix A.1.

Indeed, $\mathbf{z} = 2\mathbf{x} - \mathbf{1}$ is a vector of independent Rademacher random variables and we have that

$$\hat{\tau}(q) = \frac{1}{T} \langle \mathbf{z} * q, 2(\mathbf{x} * g + e) \rangle = \frac{1}{T} \langle \mathcal{T}_q \mathbf{z}, 2\mathcal{T}_g \mathbf{x} + 2e \rangle = \frac{1}{T} \langle \mathbf{z}, \mathcal{T}_q^\top \mathcal{T}_g \mathbf{z} \rangle + \frac{1}{T} \langle \mathbf{z}, \mathcal{T}_q^\top \mathcal{T}_g \mathbf{1} + 2\mathcal{T}_q^\top e \rangle.$$

From this calculation we can verify that $\hat{\tau}(q)$ is indeed unbiased as

$$\mathbb{E}[\hat{\tau}(q)] = \frac{1}{T} \text{tr}(\mathcal{T}_q^\top \mathcal{T}_g) = \sum_{t=1}^T \frac{T-t+1}{T} q_t g_t = \langle D_T q, g \rangle.$$

Moreover, the variance of $\hat{\tau}_q$ is given by

$$T^2 \text{var}(\hat{\tau}(q)) = \|\mathcal{T}_q^\top \mathcal{T}_g\|_F^2 + \text{tr}(\mathcal{T}_q^\top \mathcal{T}_g \mathcal{T}_q^\top \mathcal{T}_g) - 2 \sum_{t=1}^T \langle \mathcal{T}_q u_t, \mathcal{T}_g u_t \rangle^2 + \|\mathcal{T}_q^\top (g * \mathbf{1} + 2e)\|^2. \quad (3.4)$$

Here u_t denotes the t -th standard basis vector. Though this variance formula can be computed from the problem data, it is not wieldy or illuminating. In the later sections, we will compute an approximate variance formula that provides more insights into the dependence of the variance on g , q , and e .

Example utility of N-of-1 Design The simple variance calculations thus far have already shed some light on the utility of rapid switching in N-of-1 designs. If one wants to use a standard Horvitz-Thompson style treatment effect when implementing an N-of-1 design, the period at which the treatment is switched must be carefully chosen. Suppose a patient is trying to choose between two drugs, Drug A and Drug B. In a standard N-of-1 design, the drugs would be given over windows so that their effects wash out over time. That is, the half-life of the drug would determine the switching time.

In our N-of-1 design, the drugs can be interleaved more rapidly. The benefit of such an approach is that if there are decaying effects, the patient would experience the average response of the two treatments. But this averaging could come at the cost of having to run a longer experiment. Here, we show a toy example illustrating how the expected time required to distinguish two treatments is approximately the same.

Suppose that Drug A has an impulse response

$$g_t^{(1)} = A g_t^{(0)}$$

and Drug B has an impulse response

$$g_t^{(2)} = B g_t^{(0)}$$

Suppose A and B are positive scalars and, w.l.o.g., $A \geq B$.

The associated estimand for this treatment effect is

$$\tau_1 = g_0^{(1)} - g_0^{(2)} = (A - B) g_0^{(0)}$$

This corresponds to $\tau(q)$ with $q_1 = 1$ and all other entries equalling 0.

Consider the signal-to-noise ratio, equal to

$$SNR = \frac{\tau(q)}{\sqrt{\text{var}(\widehat{\tau}(q))}}.$$

Suppose we observe an observation length T . Using the variance formula (3.4), we have

$$\text{var}(\widehat{\tau}(q)) = \frac{(A-B)^2}{T} v(g^{(0)}) + \frac{(A+B)^2}{T} w(g^{(0)})$$

where

$$v(g^{(0)}) = \sum_{t=2}^T \frac{T-t+1}{T} g_t^{(0)2},$$

$$w(g^{(0)}) = \frac{1}{T} \|g^{(0)} * 1\|_2^2.$$

Note that both of these quantities are bounded for all t by $\|g^{(0)}\|_2^2$. A crude bound for the variance is thus

$$\text{var}(\widehat{\tau}(q)) \leq \frac{(A+B)^2}{T} \left\{ \left(\frac{A-B}{A+B} \right)^2 v(g^{(0)}) + w(g^{(0)}) \right\} \leq \frac{(A+B)^2}{T} \left(2\|g^{(0)}\|^2 - g_0^{(0)2} \right). \quad (3.5)$$

Thus, even with rapid interleaving of treatments, the SNR tends to infinity at a rate of \sqrt{T} .

We can say a bit more than this. If we do a standard experiment design where we wait ℓ time steps for the transient dynamics to end, measure a new outcome, and potentially switch the treatment, then the signal-to-noise ratio with random treatment assignments is

$$\frac{A-B}{A+B} \sqrt{T} \ell^{-1/2}.$$

We can compare the signal-to-noise ratios of the rapid design to the conservative design for particular models of $g^{(0)}$. As an illustrative example, suppose that $g_t^{(0)} = \beta^t$. That is, it's a simple exponential decay. Suppose t indexes days and $\beta = 0.5$. This means that the treatment has a half-life of one day. For an N-of-1 experiment that requires the effect to clear, this means you should wait 4-5 days between switching treatments. Let us choose 4 for convenience, setting $\ell^{-1/2}$ to be 0.5. Hence, the signal-to-noise ratio of the Horvitz-Thompson estimator is

$$0.5 \frac{A-B}{A+B} \sqrt{T},$$

where T is the total number of days the experiment is run.

On the other hand, using the convolutional estimator, the signal-to-noise ratio would be at least

$$0.77 \frac{A-B}{A+B} \sqrt{T}.$$

This bound is estimated using (3.5). This lower bound on the SNR is 1.5 times larger than that required for a standard design of the same duration. This shows that rapid interleaving of treatments can accelerate the time required to assess the difference between treatments with decaying effects.

3.2 Plug-in Estimators for the Variance

To construct confidence intervals using a normal approximation, we need to establish data-driven estimators for the variance of the estimator. Our approach is to provide direct, coarse estimates of the parameter vector g and the error signal e , and then to plug these values into the variance formula (3.4).

In this section, we assume q is only nonzero in the first K coordinates. As is evident from the variance formula (3.4), the variance of the estimator is potentially small only when the function q is only nonzero in the first few components. We will describe precisely how small K needs to be in more detail in Section 4.

For any integer k , let u_k denote the k th standard basis vector. Then $\widehat{\tau}(u_k)$ is an unbiased estimator of g_k . Consider the following truncated estimator of g :

$$\widehat{g}_{<K} = (\widehat{\tau}(u_0), \widehat{\tau}(u_1), \dots, \widehat{\tau}(u_{K-1}), 0, \dots, 0).$$

This estimator is unbiased in the first K components but approximates the remaining components as equal to zero. We can use this estimate of g to estimate the error signal e as well

$$\widehat{e}_t := \mathbf{y}_t - (\mathbf{x} * \widehat{g}_{<K})_t,$$

With these two estimates, we can estimate the variance. We simply plug these estimates into the variance formula (3.4):

$$\widehat{\mathcal{V}} := \|\mathcal{T}_q^\top \mathcal{T}_{\widehat{g}_{<K}}\|_F^2 + \text{tr}(\mathcal{T}_q^\top \mathcal{T}_{\widehat{g}_{<2K}} \mathcal{T}_q^\top \mathcal{T}_{\widehat{g}_{<2K}}) - 2 \sum_{t=1}^T \langle \mathcal{T}_q u_t, \mathcal{T}_{\widehat{g}_{<K}} u_t \rangle^2 + \|\mathcal{T}_q^\top (\widehat{g}_{<K} * \mathbf{1} + 2\widehat{e})\|^2. \quad (3.6)$$

Note that in the second term, we use a $2K$ -length estimate of g , but all other terms only use K terms. In the subsequent theory section, we provide evidence for this estimator being asymptotically normal, and hence yielding valid normal confidence intervals.

4 The Circular Convolution Model

Consider a circular convolution model for the impulse response with horizon size $T \in \mathbb{N}$,

$$\mathbf{y} = \mathbf{x} \otimes g + e \in \mathbb{R}^T. \quad (4.1)$$

How does this differ from linear convolution? As we described in Section 2, circular convolution models g and e as periodic signals, and such models do not describe the sorts of N-of-1 experiments typically considered in the medical context. However, on a long time horizon, when g decays quickly, circular convolutions and linear convolutions produce similar outputs. As we will see, shortly, circular convolutions will provide insights into the asymptotic behavior of linear convolution experiment designs, provided that the observation time is long enough.

Moreover, circular convolutions are interesting in their own right for modeling periodic phenomena. For instance, periodic behavior due to seasonality and other cyclic trends in financial times series (see, for example, Box [5] or Chapter 2.8 of Tsay[28]). Hence, we include full details for experiment design in such periodic systems.

In this section, we first develop the circular convolution model, providing the analogous estimand and estimator to the linear case. We then show that for large t , the estimators for the linear and circular models differ by a vanishing amount. We then establish asymptotic normality for the circular model, which in turn verifies the asymptotic normality of the linear design as well.

4.1 Estimand and Estimator

We are concerned with a family of estimands indexed by $q \in \mathbb{R}^T$

$$\tau(q) := \langle q, g \rangle. \quad (4.2)$$

Special cases include the *cumulative lag- K effects*, $K \in [1, T] \cap \mathbb{Z}$, where vector $q = \mathbf{1}_{<K} := \sum_{k=0}^{K-1} u_k$ is plugged in,

$$\tau_K := \tau(\mathbf{1}_{<K}) = \tau\left(\sum_{k=0}^{K-1} u_k\right) = \sum_{k=0}^{K-1} g_k. \quad (4.3)$$

Again, u_k is the k -th standard basis vector.

Accordingly, consider estimators of the following form, indexed by $q \in \mathbb{R}^T$ (and its circular extension q)

$$\widehat{\tau}(q) = \frac{1}{T} \langle (2\mathbf{x} - \mathbf{1}) \otimes q, 2\mathbf{y} \rangle. \quad (4.4)$$

As a special case, the estimator for the lag- K effects τ_K is

$$\widehat{\tau}_K := \widehat{\tau}(\mathbf{1}_{<K}) = \frac{1}{T} \langle (2\mathbf{x} - \mathbf{1}) \otimes \mathbf{1}_{<K}, 2\mathbf{y} \rangle. \quad (4.5)$$

For the linear convolution model, the estimator takes the same form as the (4.4), only changing the circular convolution to linear convolution,

$$\widehat{\tau}^{\text{lin}}(q) = \frac{1}{T} \langle (2\mathbf{x} - \mathbf{1}) * q, 2\mathbf{y} \rangle. \quad (4.6)$$

In this case, the estimand turns out to be slightly different than that in (4.2),

$$\tau^{\text{lin}}(q) := \sum_{t \in [T]} \frac{T-t}{T} q_t g_t. \quad (4.7)$$

However, one can already see that if q_t is 0 for $t \geq K$, then the linear and circular estimands approach each other as T tends to infinity.

Let us now consider the same simple design as before, when $\mathbf{x} \in \mathbb{R}^T$ is randomly treated at each time step. That is, $\mathbf{x}_t, 0 \leq t < T$ is i.i.d. Bernoulli. As before, we denote the centered randomization vector $\mathbf{z} := 2\mathbf{x} - \mathbf{1} \in \mathbb{R}^T$. The estimator $\widehat{\tau}(q)$ is a quadratic polynomial of the \mathbf{z} :

$$\begin{aligned} \widehat{\tau}(q) &= \frac{1}{T} \langle (2\mathbf{x} - \mathbf{1}) \otimes q, (2\mathbf{x} - \mathbf{1}) \otimes g \rangle + \frac{1}{T} \langle (2\mathbf{x} - \mathbf{1}) \otimes q, \mathbf{1} \otimes g + 2e \rangle, \\ &= \frac{1}{T} \langle \mathbf{z} \otimes q, \mathbf{z} \otimes g \rangle + \frac{1}{T} \langle \mathbf{z} \otimes q, h \rangle, \end{aligned}$$

where

$$h := \mathbf{1} \otimes g + 2e \in \mathbb{R}^T.$$

Observe that the difference between the estimator and the estimand, noted as \mathbf{W}_T , takes the form

$$\mathbf{W}_T := \widehat{\tau}(q) - \tau(q) = \frac{1}{T} \sum_{i \neq j \in [T]} \mathbf{z}_i \mathbf{z}_j H_{ij} + \frac{1}{T} \sum_{i \in [T]} \mathbf{z}_i L_i, \quad (4.8)$$

where

$$H_{ij} := \sum_{t \in [T]} q_{t-i}^\circ g_{t-j}^\circ, \quad L_i := \sum_{t \in [T]} q_{t-i}^\circ h_t. \quad (4.9)$$

Here the $^\circ$ notation denotes the circular extension defined in (2.1). This calculation reveals that the estimator is unbiased. The notation also allows us to connect back to the linear model.

4.2 Circular vs. Linear Convolutions

This section shows that circular and linear convolutions are intimately connected, at least for the type of impulse response we consider. Consider an impulse response signal of length $g \in \mathbb{R}^T$. As developed in Section 3, the estimator for the linear convolution model was given by

$$\mathbf{W}_T^{\mathcal{L}} := \widehat{\tau}_q - \tau_q = \frac{1}{T} \sum_{i \neq j \in [T]} \mathbf{z}_i \mathbf{z}_j H_{ij}^{\mathcal{L}} + \frac{1}{T} \sum_{i \in [T]} \mathbf{z}_i L_i^{\mathcal{L}},$$

where

$$\begin{aligned} H^{\mathcal{L}} &:= \mathcal{T}_q^\top \mathcal{T}_g, \\ L^{\mathcal{L}} &:= \mathcal{T}_q^\top h^{\mathcal{L}}, \text{ where } h^{\mathcal{L}} := \mathbf{1} * g + 2e. \end{aligned}$$

We can already see that the linear term $L^{\mathcal{L}}$ closely imitates the circular convolution model and exhibits asymptotic normality under mild conditions.

We now show that the matrix $H^{\mathcal{L}}$ and H are entrywise close.

Lemma 1. Recall H_{ij} 's defined in (4.9), and assume that q, g are supported only on the first K -entries with each entry bounded. Then

$$\sup_{z \in \{-1, 1\}^T} \frac{1}{T} |z^\top (H - H^\mathcal{L}) z| = O\left(\frac{K^3}{T}\right).$$

Proof. Note that we can write the entries of both matrices in a parallel manner

$$\begin{aligned} H_{ij} &= \langle \mathcal{C}_q u_i, \mathcal{C}_g u_j \rangle \\ H_{ij}^\mathcal{L} &= \langle \mathcal{T}_q u_i, \mathcal{T}_g u_j \rangle. \end{aligned}$$

These are dot products between columns of circulant and Toeplitz matrices, respectively. Let u be a vector supported only on the first K entries. Then $\mathcal{C}_u u_i = \mathcal{T}_u u_i$ whenever $i \leq T - K$. Hence, $H_{ij} = H_{ij}^\mathcal{L}$ unless either i or j is greater than $T - K$. We also have that both H_{ij} and $H_{ij}^\mathcal{L}$ equal 0 if $|i - j \pmod{T}| > K$ due to the fact that q, g are supported only on the first K -entries. By the same reasoning, $\max_{i,j} |H_{ij}|$ and $\max_{i,j} |H_{ij}^\mathcal{L}|$ are both $O(K)$. Hence $H_{ij} - H_{ij}^\mathcal{L}$ is only nonzero in at most $O(K^2)$ entries, and in these entries, the magnitude of the difference is at most $O(K)$. This proves the lemma. \square

In what follows, we shall establish $\frac{1}{\sqrt{T}}$ asymptotic normality for the circular convolution model where the higher order moment calculations and formulae are considerably simpler. Coupled with the above approximation result, asymptotic normality holds for the linear model.

4.3 Variance of the Circular Convolution Model

In the circular convolution model, the variance of $\hat{\tau}(q)$ yields a more interpretable expression than the linear model. Since the two estimators are close to each other for large T , we can draw further insights about how the variance depends on g and q .

Proposition 1 (Second Moment).

$$T \cdot \mathbb{E}_{\mathbf{x}} \left[\left(\hat{\tau}(q) - \tau(q) \right)^2 \right] = \mathcal{V}_Q + \mathcal{V}_L,$$

where

$$\mathcal{V}_Q = \text{var} \left(\frac{1}{\sqrt{T}} \sum_{i \neq j \in [T]} \mathbf{z}_i \mathbf{z}_j H_{ij} \right) = \|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2 \quad (4.10)$$

$$\mathcal{V}_L = \text{var} \left(\frac{1}{\sqrt{T}} \sum_{i \in [T]} \mathbf{z}_i L_i \right) = \frac{1}{T} \|(\mathbf{1} \otimes g + 2e) \otimes q\|_2^2. \quad (4.11)$$

The variance of this estimator thus is governed by the norm of the convolution of g with q and on the norm of the exogenous disturbance e . In order for this variance to be small, it is necessary that g and q have norms that grow sublinearly with T . In particular, this suggests that q should have sublinear support size. The estimable treatment effects must only concern a small number of components of the impulse response g . Under these assumptions, we can further show that the estimator is asymptotically normal. Asymptotic normality will serve as a building block for inference of the causal parameters.

4.4 Asymptotic Normality

We establish asymptotic normality individually for the linear and quadratic terms of (4.8). To establish the asymptotic normality for the quadratic term, we need one technical assumption on the decay of signal g and the vector q ; it is crucial to note that asymptotic normality does not hold for arbitrary g and q . When the signal $g \otimes q$ is approximately supported on the top $O(K)$ elements, the assumption holds for a wide range of $K = K(T)$.

Assumption 1. Let $q = \mathbf{1}_{<K}$, and assume there exist $K = K(T)$ such that

$$\|(|g| \otimes |q|) \otimes (|g| \otimes |q|)\|_2^2 \leq C \cdot T^{1-\epsilon} \cdot (\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2)^2$$

with some universal constants $C > 1$, $\epsilon > 0$.

This assumption can be satisfied by various means. One that most closely connects with assumptions in the dynamical systems literature is to assume g is asymptotically stable. For the purpose of this work, this just means that g_t can be upper bounded by an exponentially decaying function. Indeed, suppose $|g_t| \lesssim \rho^t$ with some positive $\rho \in (0, 1)$ and $\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2 = \Omega(1)$ (non-trivial limiting variance), then the above assumption holds with a wide range of $K = K(T) = \text{Polylog}(T)$. This is true because $\|(|g| \otimes |q|) \otimes (|g| \otimes |q|)\|_2^2 \lesssim \text{Poly}(K) = \text{Polylog}(T) \lesssim T^{1-\epsilon}$. The constant ρ is related to the mixing time of the dynamical system when driven by white noise [23].

With Assumption 1 established, we can now proceed to verify asymptotic normality. Denote

$$\mathbf{H}_T := \frac{1}{T} \sum_{i \neq j \in [T]} \mathbf{z}_i \mathbf{z}_j H_{ij}, \mathbf{L}_T := \frac{1}{T} \sum_{i \in [T]} \mathbf{z}_i L_i,$$

where H_{ij} and L_i are defined in (4.9).

Theorem 1 (Quadratic Term). Let $q = \mathbf{1}_{<K}$. Under Assumption 1,

$$\frac{\sqrt{T} \cdot \mathbf{H}_T}{\sqrt{\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2}} \Rightarrow \mathcal{N}(0, 1), \quad \text{as } T \rightarrow \infty.$$

In order to prove this theorem, we need to record two lemmas. These are both proven in the appendix. The first lemma controls the magnitude of terms in the second moment of \mathbf{H} . This is needed in multiple places in order to verify asymptotic normality.

Lemma 2. With H_{ij} defined as in (4.9),

$$\max_i \sum_{j:j \neq i} H_{ij}^2 + H_{ij} H_{ji} \leq \|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2.$$

The second lemma provides the critical moment condition bounding the ratio of the fourth moment and square of the second moment of \mathbf{H} .

Lemma 3 (Fourth Moment Estimate). Then

$$\left| \frac{\mathbb{E}[\mathbf{H}_T^4]}{(\mathbb{E}[\mathbf{H}_T^2])^2} - 3 \right| \leq \frac{4}{T} + \frac{16}{T} \frac{\|(|g| \otimes |q|) \otimes (|g| \otimes |q|)\|_2^2}{(\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2)^2}.$$

Together with Assumption 1, these two lemmas immediately yield a proof of Theorem 1 via the argument of DeJong [6].

Proof of Theorem 1. We first verify the assumption (a) in Theorem 2.1 of [6]. By Lemma 2,

$$\frac{\max_i \sum_{j:j \neq i} H_{ij}^2 + H_{ij} H_{ji}}{\sum_{i \neq j} H_{ij}^2 + H_{ij} H_{ji}} \leq \frac{1}{T},$$

which goes to zero as $T \rightarrow \infty$.

By Lemma 3

$$\left| \frac{\mathbb{E}[\mathbf{H}_T^4]}{(\mathbb{E}[\mathbf{H}_T^2])^2} - 3 \right| \leq \frac{4}{T} + \frac{16}{T} \frac{\|(|g| \otimes |q|) \otimes (|g| \otimes |q|)\|_2^2}{(\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2)^2}.$$

By Assumption 1, we have

$$\frac{\|(|g| \otimes |q|) \otimes (|g| \otimes |q|)\|_2^2}{(\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2)^2} \leq C \cdot T^{1-\epsilon}$$

and thus

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}[\mathbf{H}_T^4]}{(\mathbb{E}[\mathbf{H}_T^2])^2} = 3.$$

Now, we have verified the assumption (b) in Theorem 2.1 of [6]. Therefore, asymptotic normality is a direct consequence of Theorem 2.1 in [6]. \square

Asymptotic normality of the linear term follows from the standard Lindeberg-Feller Theorem.

Theorem 2 (Linear Term). *Let $q = \mathbf{1}_{<K}$. Assume that $g, e \in \mathbb{R}^T$ satisfy $|\langle \mathbf{1}, g \rangle| > 2\|e\|_\infty$. Then*

$$\frac{\sqrt{T} \cdot \mathbf{L}_T}{\sqrt{\frac{1}{T} \|(\mathbf{1} \otimes g + 2e) \otimes q\|_2^2}} \Rightarrow \mathcal{N}(0, 1), \quad \text{as } T \rightarrow \infty.$$

Proof of Theorem 2. We can write

$$\frac{\sqrt{T} \cdot \mathbf{L}_T}{\sqrt{\frac{1}{T} \|(\mathbf{1} \otimes g + 2e) \otimes q\|_2^2}} = \sum_{i \in [T]} \mathbf{z}_i \frac{\frac{1}{\sqrt{T}} L_i}{\sigma_T},$$

where

$$\begin{aligned} \sigma_T^2 &:= \frac{1}{T} \sum_{i \in [T]} L_i^2 = \frac{1}{T} \|(\mathbf{1} \otimes g + 2e) \otimes q\|_2^2 \\ &= \frac{1}{T} \|K \langle \mathbf{1}, g \rangle \mathbf{1} + 2e \otimes q\|_2^2 \end{aligned}$$

We know

$$\begin{aligned} \max_i |L_i| &\leq \|K \langle \mathbf{1}, g \rangle \mathbf{1} + 2e \otimes q\|_\infty \leq K |\langle \mathbf{1}, g \rangle| + 2\|e \otimes q\|_\infty \\ \sigma_T^2 &= K^2 \langle \mathbf{1}, g \rangle^2 + 4K \langle \mathbf{1}, g \rangle \frac{1}{T} \langle \mathbf{1}, e \otimes q \rangle + \frac{4}{T} \langle e \otimes q, e \otimes q \rangle \end{aligned}$$

Under the assumption, there exists a small constant ϵ such that (1) $K |\langle \mathbf{1}, g \rangle| > \frac{1}{(1-\epsilon)^2} 2\|e \otimes q\|_\infty$ and (2) $K^2 \langle \mathbf{1}, g \rangle^2 > \frac{1}{(1-\epsilon)^2} \frac{4}{T} \langle e \otimes q, e \otimes q \rangle$, and thus

$$\max_i |L_i| \leq (1 + (1-\epsilon)^2) K |\langle \mathbf{1}, g \rangle|.$$

Now, using the inequality

$$(1-\epsilon) K^2 \langle \mathbf{1}, g \rangle^2 + \frac{4}{1-\epsilon} \frac{1}{T} \langle e \otimes q, e \otimes q \rangle + 4K \langle \mathbf{1}, g \rangle \frac{1}{T} \langle \mathbf{1}, e \otimes q \rangle = \frac{1}{T} \|\sqrt{1-\epsilon} K \langle \mathbf{1}, g \rangle \mathbf{1} + 2 \frac{1}{\sqrt{1-\epsilon}} e \otimes q\|^2 \geq 0,$$

we can lower bound the variance

$$\begin{aligned} \sigma_T^2 &\geq K^2 \langle \mathbf{1}, g \rangle^2 + \frac{4}{T} \langle e \otimes q, e \otimes q \rangle - (1-\epsilon) K^2 \langle \mathbf{1}, g \rangle^2 - \frac{4}{1-\epsilon} \frac{1}{T} \langle e \otimes q, e \otimes q \rangle \\ &\geq \epsilon K^2 \langle \mathbf{1}, g \rangle^2 - \frac{\epsilon}{1-\epsilon} \frac{4}{T} \langle e \otimes q, e \otimes q \rangle \geq \epsilon^2 K^2 \langle \mathbf{1}, g \rangle^2. \end{aligned}$$

We hence have

$$\max_i \frac{\frac{1}{T} |L_i|^2}{\sigma_T^2} \leq \frac{1}{T} \frac{(1 + (1-\epsilon)^2)^2 K^2 \langle \mathbf{1}, g \rangle^2}{\epsilon^2 K^2 \langle \mathbf{1}, g \rangle^2} \rightarrow 0,$$

which implies for any fixed δ

$$\limsup_{T \rightarrow \infty} \frac{1}{\sigma_T^2} \sum_i \frac{1}{T} |L_i|^2 \mathbb{P}(|\mathbf{z}_i| \frac{1}{\sqrt{T}} L_i > \delta \sigma_T) = 0.$$

The last equation verifies the Lindeberg-Feller condition; thus, asymptotic normality holds. \square

5 Inference: Confidence Intervals

Equipped with the asymptotic normality, building confidence intervals is reduced to estimating the variance. In this section, we will extend the plug-in variance estimator of Section 3.2 to the circular case and show that it leads to valid confidence intervals. The argument for the linear convolution model is exactly the same, albeit with less compact formulas in the calculations.

We will use non-asymptotic concentration inequalities to derive that both terms can be estimated consistently with high probability. With the asymptotic normality established in the previous section we thus will derive data-driven confidence intervals on $\widehat{\tau}(q)$ when $q = \mathbf{1}_{<K}$. We will provide estimates $\widehat{\mathcal{V}}_Q$ and $\widehat{\mathcal{V}}_L$ for the quantities \mathcal{V}_Q and \mathcal{V}_L in the variance formula of Proposition 1. Then provided $K = O(\log(T))$,

$$\widehat{\tau}(q) \pm \mathcal{Z}_{1-\alpha/2} \sqrt{\frac{\widehat{\mathcal{V}}_Q + \widehat{\mathcal{V}}_L}{T}} .$$

will be valid confidence intervals with coverage $1 - \alpha$.

To start with, we state a uniform consistency result.

Theorem 3 (Uniform Consistency). *Assume $|e_t| \leq M$, $\forall t \geq 0$ and $\|g\|_1 \leq C$ for absolute constants $C, M > 0$. Let K be an integer that $K \ll T$, and recall $u_k, k \in [K]$ defined in (4.3). Then, with probability at least $1 - 4T^{-2}$*

$$\begin{aligned} \sup_{k < K} |\widehat{\tau}(u_k) - g_k| &\lesssim \frac{\log(T)}{\sqrt{T}} , \\ |\widehat{\tau}(\mathbf{1}_{<K}) - \langle \mathbf{1}_{<K}, g \rangle| &\lesssim \frac{K \log(T)}{\sqrt{T}} . \end{aligned}$$

Assume in addition that $|g_t| \lesssim c^t$ with some positive constant $c < 1$. Then with $K = K(T) = \Theta(\log(T))$,

$$|\widehat{\tau}(\mathbf{1}_{<K}) - \langle g, \mathbf{1} \rangle| \xrightarrow{a.s.} 0, \quad \text{as } T \rightarrow \infty .$$

This Theorem can be proven by invoking the Hanson-Wright and Bernstein inequalities. Before proving the theorem, we record a useful inequality that explicitly captures the error rate of our estimator.

Lemma 4. *With probability at least $1 - 4 \exp(-\gamma)$, for any fixed q ,*

$$\begin{aligned} |\widehat{\tau}(q) - \langle g, q \rangle| &\lesssim \sqrt{\frac{\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle}{T}} (\sqrt{\gamma} \vee \gamma) \\ &+ \sqrt{\frac{\|(\mathbf{1} \otimes g + 2e) \otimes q\|_2^2}{T^2}} \sqrt{\gamma} + \frac{\|q\|_1 \cdot \|\mathbf{1} \otimes g + 2e\|_\infty}{T} \gamma . \end{aligned}$$

Proof. Using (4.8), we have

$$|\widehat{\tau}(q) - \langle g, q \rangle| \leq \frac{1}{T} \left| \sum_{i \neq j \in [T]} \mathbf{z}_i \mathbf{z}_j H_{ij} \right| + \frac{1}{T} \left| \sum_{i \in [T]} \mathbf{z}_i L_i \right| .$$

We can bound each term separately. By the Hanson-Wright inequality, we have with probability at least $1 - 2 \exp(-\gamma)$

$$\left| \sum_{i \neq j} H_{ij} \mathbf{z}_i \mathbf{z}_j - \text{Tr}(H) \right| \lesssim \sqrt{\|H\|_{\text{F}}^2 \gamma} + \|H\|_{\text{op}} \gamma ,$$

As derived in the proof of Proposition 1,

$$\left\| \frac{1}{2} (H + H^\top) \right\|_{\text{F}}^2 = \frac{T}{2} (\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle g, q \rangle^2) .$$

and we always also have $\|H + H^\top\|_{\text{op}} \leq \|H + H^\top\|_{\text{F}}$.

For the second term, with probability at least $1 - 2\exp(-\gamma)$, Bernstein's inequality asserts

$$|\sum_i L_i \mathbf{z}_i| \lesssim \sqrt{\|L\|_2^2 \gamma} + \|L\|_\infty \gamma.$$

In this case, we again know from Proposition 1 that $\|L\|_2^2 = \|(\mathbf{1} \otimes g + 2e) \otimes q\|_2^2$ and again bound $\|L\|_\infty \leq \|q\|_1 \cdot \|\mathbf{1} \otimes g + 2e\|_\infty$, completing the proof. \square

Using this lemma, Theorem 3 is immediate

Proof of Theorem 3. For $q = u_k$, one can immediately verify that

$$\begin{aligned} \|g \otimes q\|_2^2 &= \sum_i (g_{i+k})^2 = \|g\|_2^2, \\ \langle g \otimes g, q \otimes q \rangle &\leq \sum_i g_i g_{2k-i} \leq \left(\sum_i g_i^2\right)^{1/2} \left(\sum_i (g_{2k-i})^2\right)^{1/2} = \|g\|_2^2. \end{aligned}$$

In addition,

$$\frac{\|(\mathbf{1} \otimes g + 2e) \otimes q\|_2^2}{T} = \frac{\|\mathbf{1} \otimes g + 2e\|_2^2}{T} \leq \|\mathbf{1} \otimes g + 2e\|_\infty^2.$$

Lemma 4 reads:

$$|\widehat{\tau}(u_k) - g_k| \lesssim \|g\|_2 \frac{\sqrt{\gamma} \vee \gamma}{\sqrt{T}} + \|\mathbf{1} \otimes g + 2e\|_\infty \left(\sqrt{\frac{\gamma}{T}} \vee \frac{\gamma}{T}\right),$$

with probability at least $1 - 4\exp(-\gamma)$. We complete the proof by applying union bound and plug in $\gamma = \log(T^2 K)$. To obtain the almost sure statement $\widehat{\tau}(\mathbf{1}_{<K}) - \langle g, \mathbf{1} \rangle \xrightarrow{\text{a.s.}} 0$, we take $K = \Theta(\log(T))$ and apply the Borel–Cantelli lemma since $\sum_T T^{-2} < \infty$. \square

Plug-in estimate of variance Let $q = \mathbf{1}_{<K}$, now we are ready to state the data-driven estimators for the terms \mathcal{V}_Q and \mathcal{V}_L in the case of convolution. The estimators are essentially the same as what was presented for linear convolution. Denote

$$\widehat{g}_{<K} = \overbrace{(\widehat{\tau}(u_0), \dots, \widehat{\tau}(u_{K-1}))}^{\text{first } K \text{ entries}}, 0, \dots, 0$$

and define the plug-in estimate of \mathcal{V}_Q

$$\widehat{\mathcal{V}}_Q = \|\widehat{g}_{<K} \otimes q\|_2^2 + \langle \widehat{g}_{<2K} \otimes \widehat{g}_{<2K}, q \otimes q \rangle - 2(\widehat{\tau}(\mathbf{1}_{<K}))^2 \quad (5.1)$$

For the term \mathcal{V}_L , we are going to design a plug-in estimate. We introduce the estimator for the error vector

$$\widehat{e}_t := \mathbf{y}_t - (\mathbf{x} \otimes \widehat{g}_{<K})_t,$$

and define

$$\widehat{\mathcal{V}}_L = \frac{1}{T} \|K \widehat{\tau}(\mathbf{1}_{<K}) \cdot \mathbf{1} + 2 \cdot \widehat{e} \otimes q\|_2^2. \quad (5.2)$$

Conceptually, the plug-in estimate of $\mathcal{V}_Q, \mathcal{V}_L$ is to replace $g \leftarrow \widehat{g}_{<K}$, and to replace $e \leftarrow \widehat{e} := \mathbf{y} - \mathbf{x} \otimes \widehat{g}_{<K}$.

Proposition 2 (Validity of the Plug-in Estimate of Variance). *Let $q = \mathbf{1}_{<K}$. Consider the same assumptions as in Theorem 3. Then with $K = K(T) = \Theta(\log(T))$, as $T \rightarrow \infty$*

$$\begin{aligned}\widehat{\mathcal{V}}_Q &\xrightarrow{a.s.} \mathcal{V}_Q, \\ \widehat{\mathcal{V}}_L &\xrightarrow{a.s.} \mathcal{V}_L.\end{aligned}$$

With the asymptotic normality established in the previous section and the above proposition asserting the validity of the plug-in estimate of variance, we can define the data-driven confidence interval with coverage $1 - \alpha$: choose $K = K(T) = \Theta(\log(T))$, let $q = \mathbf{1}_{<K}$, and

$$\widehat{\tau}(q) \pm C_\alpha \sqrt{\frac{\widehat{\mathcal{V}}_Q + \widehat{\mathcal{V}}_L}{T}}.$$

Proof. To show the consistency of $\widehat{\mathcal{V}}_Q \rightarrow \mathcal{V}_Q$, we will show consistency for each term in (5.1).

First we will show $\|\widehat{g}_{<2K} \otimes q\|_2^2 \xrightarrow{a.s.} \|g \otimes q\|_2^2$. For $t \geq 2K$

$$|(g \otimes q)_t| \leq \sum_{i \in [K]} |g_{t-i}| \leq K c^{t-K}$$

and thus $\|g \otimes q\|_2^2 - \sum_{t \in [2K]} (g \otimes q)_t^2 \leq K^2 c^{2K} = o_T(1)$ with $K = \Theta(\log(T))$. Note that for $t < 2K$, $(g \otimes q)_t = \sum_{i \in [K]} g_{t-i}^\circ$ and recall the uniform consistency for $\sup_{k < 2K} |\widehat{\tau}(u_k) - g_k| \lesssim \frac{\log(T)}{\sqrt{T}}$ (Theorem 3), we have with probability at least $1 - 4T^{-2}$

$$|(g \otimes q)_t - (\widehat{g}_{<K} \otimes q)_t| \lesssim \begin{cases} K \frac{\log(T)}{\sqrt{T}} + c^{T-1-K} = o_T(1) & \text{for } t \in [0, K) \\ K \frac{\log(T)}{\sqrt{T}} = o_T(1) & \text{for } t \in [K, 2K) \end{cases}$$

and thus, we have

$$\|\widehat{g}_{<K} \otimes q\|_2^2 - \sum_{t \in [2K]} (g \otimes q)_t^2 = o_T(1).$$

To sum up, the above derivation implies $\|\widehat{g}_{<K} \otimes q\|_2^2 \xrightarrow{a.s.} \|g \otimes q\|_2^2$.

Now we study the term $\langle \widehat{g}_{<2K} \otimes \widehat{g}_{<2K}, q \otimes q \rangle \xrightarrow{a.s.} \langle g \otimes g, q \otimes q \rangle$. First note that

$$\langle g \otimes g, q \otimes q \rangle = \sum_{t \in [2K]} (g \otimes g)_t (q \otimes q)_t,$$

and that $q \otimes q$ is supported on the first $2K$ -entries. For $t < 2K$, $(g \otimes g)_t - \sum_{i \in [2K]} g_i g_{t-i}^\circ \lesssim c^{2K}$, and by the uniform consistency in Theorem 3 and triangle inequality, we have

$$\begin{aligned} |(g \otimes g)_t - (\widehat{g}_{<2K} \otimes \widehat{g}_{<2K})_t| &\leq \left| \sum_{i \in [2K]} g_{t-i}^\circ g_i - \sum_{i \in [2K]} \widehat{g}_{t-i}^\circ \widehat{g}_i \right| + c^{2K} \\ &\leq \sum_{i \in [2K]} g_{t-i}^\circ \cdot \max_{i \in [2K]} |\widehat{g}_i - g_i| + \sum_{i \in [2K]} |\widehat{g}_i| \cdot \max_{i \in [2K]} |\widehat{g}_{t-i}^\circ - g_{t-i}^\circ| + c^{2K} \lesssim \text{Poly}(K) \frac{\log(T)}{\sqrt{T}} + c^{2K}. \end{aligned}$$

Therefore, $\langle \widehat{g}_{<2K} \otimes \widehat{g}_{<2K}, q \otimes q \rangle \xrightarrow{a.s.} \langle g \otimes g, q \otimes q \rangle$.

Lastly, Theorem 3 showed that

$$\langle \widehat{g}_K, q \rangle - \langle g, q \rangle \xrightarrow{a.s.} 0.$$

Put three terms together, we have shown $\widehat{\mathcal{V}}_Q \xrightarrow{a.s.} \mathcal{V}_Q$.

Now we move on to study the term \mathcal{V}_L .

$$\mathcal{V}_L = \frac{1}{T} \|(\langle \mathbf{1}, g \rangle \mathbf{1} + 2e) \otimes q\|_2^2 = \frac{1}{T} \sum_{t \in [T]} \left(K \langle \mathbf{1}, g \rangle + 2 \sum_{i \in [K]} e_{t-i}^\circ \right)^2$$

We introduce the estimator for the error vector

$$\widehat{e}_t := \mathbf{y}_t - (\mathbf{x} \otimes \widehat{g}_{<K})_t = e_t + \underbrace{(\mathbf{x} \otimes g_{<K})_t - (\mathbf{x} \otimes \widehat{g}_{<K})_t}_{(i)} + \underbrace{(\mathbf{x} \otimes g_{\geq K})_t}_{(ii)}.$$

To bound the term (i), we note that by Theorem 3

$$|(\mathbf{x} \otimes g_{<K})_t - (\mathbf{x} \otimes \widehat{g}_{<K})_t| = \left| \sum_{i \in [K]} \mathbf{x}_{t-i}^\circ (g_i - \widehat{g}_i) \right| \leq K \sup_{i \in [K]} |g_i - \widehat{g}_i| \lesssim \frac{K \log(T)}{\sqrt{T}}.$$

To bound the term (ii), we use the fact

$$|(\mathbf{x} \otimes g_{\geq K})_t| \leq \|g_{\geq K}\|_1 \|\mathbf{x}\|_\infty \lesssim c^K.$$

Thus for all $t \in [T]$,

$$\begin{aligned} \sup_{t \in [T]} |\widehat{e}_t - e_t| &\lesssim \frac{K \log(T)}{\sqrt{T}} + c^K, \\ \sup_{t \in [T]} \left| \sum_{i \in [K]} \widehat{e}_{t-i}^\circ - \sum_{i \in [K]} e_{t-i}^\circ \right| &\lesssim \frac{K^2 \log(T)}{\sqrt{T}} + Kc^K = o_T(1). \end{aligned}$$

So far, we have proved

$$\sup_{t \in [T]} \left| (K \widehat{\tau}(\mathbf{1}_{<K}) + 2(\widehat{e} \otimes q)_t) - (K \langle \mathbf{1}, g \rangle + 2(e \otimes q)_t) \right| = o_T(1)$$

and thus, as a direct consequence $\widehat{\mathcal{V}}_L \xrightarrow{a.s.} \mathcal{V}_L$. □

6 Empirics

In this section, we demonstrate a simple empirical example illustrating two theoretical results established, the asymptotic normality in Section 4.4, and the asymptotic equivalence between circular and linear convolution in Section 4.2. Besides validating the theory, the empirical example also demonstrates how asymptotic characterizations become accurate as T grows.

We fix an impulse response $g_t = 1.00 * 0.65^t - 1.60 * 0.50^t + 0.75 * 0.48^t$, $t \in [T]$ and set $K = 25$ in constructing the estimator $\widehat{\tau}_K$ defined in (4.5). For this impulse response, it is clear that the signals are negligible after $K = 25$. We consider two data-generating processes (DGP), the circular convolution model in (4.1) (see Fig. 1), and the linear convolution model in (3.2) (see Fig. 2), both with $e = 0$. For each DGP, we sample the impulse/treatment vector $\mathbf{x} \in \{0, 1\}^T$ with i.i.d. Bernoulli with $T = 200, 1000, 5000$ and then observe the response $\mathbf{y} \in \mathbb{R}^T$, to illustrate the asymptotic behavior. For each DGP and T , we run 20000 Monte Carlo simulations to show the histogram of $\widehat{\tau}_K$.

A few observations follow from the simulations in Fig. 1 and 2. First, the linear convolution and circular convolution are asymptotically equivalent, as shown in Section 4.2. Here for each T , the distributions of $\widehat{\tau}(\mathbf{1}_{<K})$ and $\widehat{\tau}^{\text{lin}}(\mathbf{1}_{<K})$ are indiscernible, especially for large T . Second, the variance formulae in circular convolution, as shown in Proposition 1, correctly provide the 95% coverage even in the linear convolution model, seen in Fig. 2. Even for $T = 200$, Fig. 2a demonstrate a 95.49% coverage, implying that the variance formula we derive under circular convolution remains an accurate approximation for the linear convolution model. Third, for each Figure, as T increases, the distributions approach an asymptotic normal distribution, validating the asymptotic normality results established in Section 4.4.




Figure 1: Circular convolution model $\mathbf{y} = \mathbf{x} \otimes g$: we plot the histogram of $\hat{\tau}(\mathbf{1}_{<K})$ defined in (4.5) based on 20000 Monte Carlo simulations, the red solid vertical line corresponds to $\tau(\mathbf{1}_{<K})$ defined in (4.3), the red dashed lines are two standard deviations from the expectation using the variance formula in Proposition 1.




Figure 2: Linear convolution model $\mathbf{y} = \mathbf{x} * g$: we plot the histogram of $\hat{\tau}^{\text{lin}}(\mathbf{1}_{<K})$ defined in (4.6) based on 20000 Monte Carlo simulations, the red solid vertical line corresponds to $\tau^{\text{lin}}(\mathbf{1}_{<K})$ defined in (4.7), the red dashed lines are two standard deviations from the expectation using the variance formula in Proposition 1.

7 Discussion and Future Work

Formalizing a statistical framework for N-of-1 designs using dynamical systems provides a convenient modeling test-bed for within-subject analyses. Our initial results here suggest that this project is feasible, and that it is relatively simple to write software to compute treatment effects under relatively simple dynamic interference patterns. However, it is mildly discouraging that even the most simple linear time-invariant model is already quite complicated. We hope that our initial work here encourages investigation into simpler designs with simpler analyses. Or, perhaps, theoretical arguments will be made as to why such simplicity is impossible.

As this paper is just a first step towards such a theory of N-of-1 designs, there are plenty of open questions that must be answered to solidify the statistical foundations. First, in this work, we only consider the simplest randomization design, toggling treatment randomly at each time step. An analysis of balanced designs would likely immediately reduce variance in simple settings. Moreover, more clever experiment designs could potentially mitigate the long-term dependencies inherent to the statistical model. Such designs could potentially reveal more information about the system from shorter observation times.

Another avenue for future work would be determining the optimal timing of interventions. Our statistical models are continuous processes sampled at a particular cadence of uniform discrete intervals. What is the best choice of this sampling time in order to most rapidly infer treatment effects? For example, in the case of medications, there will be time associated with the onset of the effect of treatment and time associated with the decay of the effect. Sampling too early or late will understate the impact of the treatment. In order to maximize the signal in an experiment, how can we determine the optimal intervention time? Can novel randomization schemes or adaptive designs help identify the optimal timing?

Finally, we hope that future work will investigate extensions beyond the linear time-invariant model studied in this work. We expect it would require little effort to extend the time-invariant treatment effects of this paper to study those where the response varies at each time step

$$y_t = \sum_{s=0}^{t-1} g_s(t-s)x_s.$$

This would make our framework a full generalization of SUTVA potential outcomes. That is, with this particular time-variant model, the SUTVA model is recovered when $g_k(t) = 0$ for all $t > 0$. The main stumbling block is how to generalize the estimate of the variance of the treatment effect for this model. Extensions to nonlinear models would also add to the generality of this framework. It could be useful to appeal to the control literature to understand how to model nonlinear dose responses and saturation effects between doses. Future work will benefit from such further synergies between causal inference and mathematical control.

Acknowledgements

The authors would like to thank Peng Ding, Avi Feller, Max Simchowitz, Panos Toulis, and Ruey Tsay for helpful suggestions and pointers to related work. This work was supported in part by ONR Award N00014-20-1-2497, NSF CCF CIF Award 2326498, and NSF Career Award (DMS-2042473). Additionally, TL acknowledges the generous support of the William Ladany Faculty Fellowship from the University of Chicago Booth School of Business and BR the generous support of the Miller Institute for Basic Research in Science.

References

- [1] Joshua D Angrist and Guido M Kuersteiner. Causal effects of monetary shocks: Semiparametric conditional independence tests with a multinomial propensity score. *Review of Economics and Statistics*, 93(3):725–747, 2011.

- [2] Ainesh Bakshi, Allen Liu, Ankur Moitra, and Morris Yau. A new approach to learning linear dynamical systems. *arXiv preprint arXiv:2301.09519*, 2023.
- [3] Iavor Bojinov, Guillaume Saint-Jacques, and Martin Tingley. Avoid the pitfalls of a/b testing make sure your experiments recognize customers’ varying needs. *Harvard Business Review*, 98(2):48–53, 2020.
- [4] Iavor Bojinov, David Simchi-Levi, and Jinglong Zhao. Design and Analysis of Switchback Experiments, April 2022.
- [5] George EP Box, Steven C Hillmer, and George C Tiao. Analysis and modeling of seasonal time series. In *Seasonal analysis of economic time series*, pages 309–344. NBER, 1978.
- [6] Peter de Jong. A central limit theorem for generalized quadratic forms. *Probability Theory and Related Fields*, 75(2):261–277, June 1987. ISSN 1432-2064. doi: 10.1007/BF00354037.
- [7] Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438, 1969.
- [8] Gordon Guyatt, David Sackett, D. Wayne Taylor, John Ghong, Robin Roberts, and Stewart Pugsley. Determining Optimal Therapy — Randomized Trials in Individual Patients. *New England Journal of Medicine*, 314(14):889–892, April 1986. ISSN 0028-4793, 1533-4406. doi: 10.1056/NEJM198604033141406. URL <http://www.nejm.org/doi/abs/10.1056/NEJM198604033141406>.
- [9] Elad Hazan, Sham Kakade, and Karan Singh. The nonstochastic control problem. In *Proceedings of the 31st International Conference on Algorithmic Learning Theory*, pages 408–421, 2020.
- [10] A. Bradford Hill. *Principles of Medical Statistics*. The Lancet Limited, London, 7th edition, 1961.
- [11] Xinran Li and Peng Ding. Exact confidence intervals for the average causal effect on a binary outcome. *Statistics in Medicine*, 35(6):957–960, 2016.
- [12] Lennart Ljung. *System Identification. Theory for the user*. Prentice Hall, Upper Saddle River, NJ, 2nd edition, 1998.
- [13] Thomas A. Louis, Philip W. Lavori, John C. Bailar, and Marcia Polansky. Crossover and Self-Controlled Designs in Clinical Research. *New England Journal of Medicine*, 310(1):24–31, January 1984. ISSN 0028-4793, 1533-4406. doi: 10.1056/NEJM198401053100106. URL <http://www.nejm.org/doi/abs/10.1056/NEJM198401053100106>.
- [14] Iven Mareels. Sufficiency of excitation. *Systems & control letters*, 5(3):159–163, 1984.
- [15] Jerzy Neyman. On the application of probability theory to agricultural experiments. essay on principles. *Ann. Agricultural Sciences*, pages 1–51, 1923.
- [16] Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- [17] Peter Van Overschee and Bart De Moor. N4sid: Subspace algorithms for the identification of combined deterministic– stochastic systems. *Automatica*, 30:75–93, 1994.
- [18] Samet Oymak and Necmiye Ozay. Non-asymptotic identification of lti systems from a single trajectory. In *2019 American control conference (ACC)*, pages 5655–5661. IEEE, 2019.
- [19] C. E. Quin, R. M. Mason, and J. Knowelden. Clinical Assessment of Rapidly Acting Agents in Rheumatoid Arthritis. *BMJ*, 2(4683):810–813, October 1950. ISSN 0959-8138, 1468-5833. doi: 10.1136/bmj.2.4683.810. URL <https://www.bmj.com/lookup/doi/10.1136/bmj.2.4683.810>.
- [20] D. D. Reid. The design of clinical experiments. *The Lancet*, 264(6852):1293–1296, 1954.
- [21] D. D. Reid. The patient as his own control. In *Controlled Clinical Trials*, London, 1960. The council for International Organizations of Medical Sciences, Blackwell Scientific Publications.

- [22] Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- [23] Max Simchowitz, Horia Mania, Stephen Tu, Michael I Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In *Conference On Learning Theory*, pages 439–473. PMLR, 2018.
- [24] Max Simchowitz, Ross Boczar, and Benjamin Recht. Learning linear dynamical systems with semi-parametric least squares. In *Conference on Learning Theory (COLT)*, 2019.
- [25] Christopher A Sims. Money, income, and causality. *The American economic review*, 62(4):540–552, 1972.
- [26] E.S. Snell and P. Armitage. CLINICAL COMPARISON OF DIAMORPHINE AND PHOLCODINE AS COUGH SUPPRESSANTS. *The Lancet*, 269(6974):860–862, April 1957. ISSN 01406736. doi: 10.1016/S0140-6736(57)91392-2. URL <https://linkinghub.elsevier.com/retrieve/pii/S0140673657913922>.
- [27] Stefan H Thomke. *Experimentation works: The surprising power of business experiments*. Harvard Business Press, 2020.
- [28] Ruey S Tsay. *Analysis of financial time series*. John wiley & sons, 2005.
- [29] Michel Verhaegen and Patrick Dewilde. Subspace model identification. *International Journal of Control*, 56(5): 1187–1210, 1992.
- [30] Frances A. Wood, James P. Howard, Judith A. Finegold, Alexandra N. Nowbar, David M. Thompson, Ahran D. Arnold, Christopher A. Rajkumar, Susan Connolly, Jaimini Cegla, Chris Stride, Peter Sever, Christine Norton, Simon A.M. Thom, Matthew J. Shun-Shin, and Darrel P. Francis. N-of-1 Trial of a Statin, Placebo, or No Treatment to Assess Side Effects. *New England Journal of Medicine*, 383(22):2182–2184, November 2020. ISSN 0028-4793. doi: 10.1056/NEJMc2031173. URL <https://doi.org/10.1056/NEJMc2031173>. Publisher: Massachusetts Medical Society _eprint: <https://doi.org/10.1056/NEJMc2031173>.

A Technical Proofs

A.1 Moments of Rademacher Chaos

Proposition 3. *Let W be a $d \times d$ symmetric matrix and let \mathbf{z} be a vector of d independent Rademacher random variables. Consider the random variable $\mathbf{G} := \sum_{i < j} \mathbf{z}_i \mathbf{z}_j W_{ij}$. Then*

$$\mathbb{E}[\mathbf{G}^2] = \sum_{i < j} W_{ij}^2$$

and

$$\mathbb{E}[\mathbf{G}^4] = 3(\mathbb{E}[\mathbf{G}^2])^2 - \sum_{i \neq j} W_{ij}^4 + 3 \sum_{i \neq j \neq k \neq l} W_{ij} W_{jk} W_{kl} W_{li}.$$

Proof of Proposition 3. We can define a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ on $|\mathcal{V}| = T$ indices, and the edge weight $W_e := W_{ij}$, $e \in \mathcal{E}$ denotes an undirected weight between nodes $i \leftrightarrow j$. Then

$$\mathbf{G} := \sum_{e \in \mathcal{E}} \mathbf{z}_e W_e$$

where for an edge $e : i \leftrightarrow j$, $\mathbf{z}_e := \mathbf{z}_i \mathbf{z}_j$. Note that $\mathbb{E}[\mathbf{z}_e \mathbf{z}_{e'}] = 0$ if $e \neq e'$. With this edge index notation, the second moment can be written as

$$\mathbb{E}[\mathbf{G}^2] = \sum_{e \in \mathcal{E}} W_e^2 + \sum_{e_1 \neq e_2 \in \mathcal{E}} \mathbb{E}[z_{e_1} z_{e_2}] W_{e_1} W_{e_2} = \sum_{e \in \mathcal{E}} W_e^2,$$

The edge indices prove convenient in evaluating the fourth moment. There are only three cases where $\mathbb{E}[\mathbf{z}_{e_1}\mathbf{z}_{e_2}\mathbf{z}_{e_3}\mathbf{z}_{e_4}]$ can be nonzero. First, if all of these edges are equal. Second, if there are two pairs of equal edges. Third, if there are four distinct edges that form a cycle.

Thus we have

$$\mathbb{E}[\mathbf{G}^4] = \sum_{e \in \mathcal{E}} W_e^4 + \frac{\binom{4}{2}\binom{2}{2}}{2!} \sum_{e_1 \neq e_2 \in \mathcal{E}} W_{e_1}^2 W_{e_2}^2 + \sum_{e_1 \neq e_2 \neq e_3 \neq e_4 \in \mathcal{E}: \text{ cycle}} W_{e_1} W_{e_2} W_{e_3} W_{e_4}.$$

The sum of the first two terms equals $\mathbb{E}[\mathbf{G}^4] = 3(\mathbb{E}[\mathbf{G}^2])^2 - \sum_{i \neq j} W_{ij}^4$. The last term is equal to

$$3 \sum_{i \neq j \neq k \neq l} W_{ij} W_{jk} W_{kl} W_{li}$$

To derive this, fix four distinct points i, j, k, l (say $i = 1, j = 2, k = 3, l = 4$). Then there will be $\frac{4!}{4 \times 2} = 3$ distinguishable edge arrangements ($1 \leftrightarrow 2 \leftrightarrow 3 \leftrightarrow 4 \leftrightarrow 1$, $1 \leftrightarrow 2 \leftrightarrow 4 \leftrightarrow 3 \leftrightarrow 1$, $1 \leftrightarrow 4 \leftrightarrow 2 \leftrightarrow 3 \leftrightarrow 1$) that form a cycle. In total, there will be $3 \times 4! = 72$ terms that involve node indices 1, 2, 3, 4, which can be grouped into

$$\sum_{\substack{e_1 \neq e_2 \neq e_3 \neq e_4 \in \mathcal{E}: \\ \text{cycle with indices } \{1,2,3,4\}}} W_{e_1} W_{e_2} W_{e_3} W_{e_4} = 24W_{12}W_{23}W_{34}W_{41} + 24W_{12}W_{24}W_{43}W_{31} + 24W_{14}W_{42}W_{23}W_{31}.$$

On the other hand, in the summation over indices, there will be $4! = 24$ terms with the index set $\{1, 2, 3, 4\}$:

$$\sum_{\substack{i \neq j \neq k \neq l: \\ \{i,j,k,l\} = \{1,2,3,4\}}} W_{ij} W_{jk} W_{kl} W_{li} = 8W_{12}W_{23}W_{34}W_{41} + 8W_{12}W_{24}W_{43}W_{31} + 8W_{14}W_{42}W_{23}W_{31}.$$

This verifies the desired moment formula. \square

The following Lemma specializes these moment calculations to a particular form of Rademacher chaos that appears several times in this work.

Lemma 5. *Let M be a matrix and \mathbf{z} be a vector of independent Rademacher random variables. Then*

$$\mathbb{E}[\langle \mathbf{z}, M\mathbf{z} \rangle] = \text{tr}(M)$$

and

$$\mathbb{E}[(\langle \mathbf{z}, M\mathbf{z} \rangle - \text{tr}(M))^2] = \|M\|_F^2 + \text{tr}(M^2) - 2 \sum_{i=0}^T M_{ii}^2$$

of Lemma 5. The first moment is an elementary calculation. For the second moment, we can write

$$\langle \mathbf{z}, M\mathbf{z} \rangle - \text{tr}(M) = \sum_{i < j} z_i z_j (M_{ij} + M_{ji}).$$

Apply Proposition 3 to the matrix $W = M + M^\top$ to complete the proof. \square

A.2 Properties of Circular Convolution

We collect a key fact about circular convolution that will be repeatedly used.

Proposition 4 (Circular Symmetry). *For any two sequence $a, b \in \mathbb{R}^T$, then for any $t, s \in [T]$*

$$\sum_{i \in [T]} a_{t-i}^\circ b_{s-i}^\circ = \sum_{i \in [T]} a_{i-s}^\circ b_{i-t}^\circ$$

Proof of Proposition 4. Define the correspondence $\sigma : i \mapsto t + s - i \pmod{T}$, note this is a bijective map from $\sigma : [T] \rightarrow [T]$.

$$\sum_{i \in [T]} a_{t-i}^\circ b_{s-i}^\circ = \sum_{i \in [T]} a_{t-\sigma(i)}^\circ b_{s-\sigma(i)}^\circ = \sum_{i \in [T]} a_{i-s}^\circ b_{i-t}^\circ.$$

\square

A.3 Proof of Proposition 1

Observe that

$$\mathbb{E}_{\mathbf{x}} \left[\left(\widehat{\Delta}(q) - \Delta(q) \right)^2 \right] = \mathbb{E}[\mathbf{W}_T^2] = \frac{1}{T^2} \sum_{i,j} \mathbb{E}[\mathbf{w}_i \mathbf{w}_j].$$

Note for $i \neq j$ (w.l.o.g. say $i > j$), due to the martingale difference property

$$\mathbb{E}[\mathbf{w}_i \mathbf{w}_j] = \mathbb{E}[\mathbf{w}_j \mathbb{E}[\mathbf{w}_i | \mathcal{F}_j]] = 0.$$

Therefore we need to calculate

$$\begin{aligned} \mathbb{E}[\mathbf{W}_T^2] &= \frac{1}{T^2} \sum_i \mathbb{E}[\mathbf{w}_i^2] = \frac{1}{T^2} \sum_i \sum_{j:j < i} (H_{ij} + H_{ji})^2 + L_i^2 \\ &= \frac{1}{T^2} \sum_{i \neq j} (H_{ij}^2 + H_{ij}H_{ji}) + \frac{1}{T^2} \sum_i L_i^2. \end{aligned}$$

The remaining argument is algebraic manipulation using properties of circular convolution. We will show

$$\begin{aligned} \sum_{i \neq j} H_{ij}^2 + H_{ij}H_{ji} &= T \cdot \left(\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2 \right), \\ \sum_i L_i^2 &= \|(\mathbf{1} \otimes g + 2e) \otimes q\|_2^2. \end{aligned}$$

First, for H we have

$$\begin{aligned} \sum_{i,j} H_{ij}^2 &= \sum_{i,j} \sum_{t \in [T]} q_{t-i}^\circ g_{t-j}^\circ \sum_{s \in [T]} q_{s-i}^\circ g_{s-j}^\circ \\ &= \sum_{t,s} \sum_i q_{t-i}^\circ q_{s-i}^\circ \left(\sum_j g_{t-j}^\circ g_{s-j}^\circ \right) \\ &= \sum_{t,s} \sum_i q_{t-i}^\circ q_{s-i}^\circ \left(\sum_j g_{j-s}^\circ g_{j-t}^\circ \right) \quad \text{by Prop. 4} \\ &= \sum_{i,j} \left(\sum_t g_{j-t}^\circ q_{t-i}^\circ \right) \left(\sum_s g_{j-s}^\circ q_{s-i}^\circ \right) \\ &= \sum_{i,j} (g \otimes q)_{j-i}^2 \pmod{T} \\ &= T \cdot \|g \otimes q\|_2^2. \end{aligned}$$

Similarly, we have

$$\begin{aligned} \sum_{i,j} H_{ij}H_{ji} &= \sum_{i,j} \sum_{t \in [T]} q_{t-i}^\circ g_{t-j}^\circ \sum_{s \in [T]} q_{s-j}^\circ g_{s-i}^\circ \\ &= \sum_{t,s} \sum_i q_{t-i}^\circ g_{s-i}^\circ \left(\sum_j g_{t-j}^\circ q_{s-j}^\circ \right) \\ &= \sum_{t,s} \sum_i q_{t-i}^\circ g_{s-i}^\circ \left(\sum_j g_{j-s}^\circ q_{j-t}^\circ \right) \quad \text{by Prop. 4} \\ &= \sum_{i,j} \left(\sum_t q_{j-t}^\circ q_{t-i}^\circ \right) \left(\sum_s g_{j-s}^\circ g_{s-i}^\circ \right) \\ &= \sum_{i,j} (g \otimes g)_{j-i} \pmod{T} (q \otimes q)_{j-i} \pmod{T} \\ &= T \cdot \langle g \otimes g, q \otimes q \rangle. \end{aligned}$$

And thus

$$\sum_{i \neq j} H_{ij}^2 + H_{ij}H_{ji} = T \cdot \left(\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2 \right).$$

For L , let $h = \mathbf{1} \otimes g + 2e$. We then have

$$\begin{aligned} \sum_i L_i^2 &= \sum_i \sum_{t,s} q_{t-i}^\circ h_t q_{s-i}^\circ h_s = \sum_{t,s} h_t h_s \sum_i q_{t-i}^\circ q_{s-i}^\circ \\ &= \sum_{t,s} h_t h_s \sum_i q_{i-s}^\circ q_{i-t}^\circ = \sum_i (h \otimes q)_i^2 = \|h \otimes q\|_2^2. \end{aligned}$$

A.4 Proof of Lemma 2

By the circular symmetry property in Proposition 4, we have

$$\begin{aligned} \max_{j:j \neq i} \omega_{ij}^2 &\leq \sum_j \omega_{ij}^2 - \omega_{ii}^2 = 2 \sum_{t,s} q_{t-i}^\circ q_{s-i}^\circ \sum_j g_{t-j}^\circ g_{s-j}^\circ + 2 \sum_{t,s} q_{t-i}^\circ g_{s-i}^\circ \sum_j g_{t-j}^\circ q_{s-j}^\circ - 4\langle g, q \rangle^2 \\ &= 2 \sum_{t,s} q_{t-i}^\circ q_{s-i}^\circ \sum_j g_{j-s}^\circ g_{j-t}^\circ + 2 \sum_{t,s} q_{t-i}^\circ g_{s-i}^\circ \sum_j g_{j-s}^\circ q_{j-t}^\circ - 4\langle g, q \rangle^2 \\ &= 2 \sum_j (g \otimes q)_{j-i}^2 \pmod{T} + 2 \sum_j (g \otimes g)_{j-i} \pmod{T} (q \otimes q)_{j-i} \pmod{T} - 4\langle g, q \rangle^2 \\ &= 2\|g \otimes q\|_2^2 + 2\langle g \otimes g, q \otimes q \rangle - 4\langle g, q \rangle^2. \end{aligned}$$

Note that the above holds for any i , we finish the proof.

A.5 Proof of Lemma 3

Define $\omega_{ij} := H_{ij} + H_{ji}$. we have

$$\omega_{ij} = \sum_{t \in [T]} q_{t-i}^\circ g_{t-j}^\circ + q_{t-j}^\circ g_{t-i}^\circ = \omega_{ji}.$$

Let $\mathbf{G} = \sum_{i < j} \mathbf{z}_i \mathbf{z}_j \omega_{ij} = T \cdot \mathbf{H}_T$. Invoking Proposition 3 yields the following equality

$$\frac{\mathbb{E}[\mathbf{H}_T^4]}{(\mathbb{E}[\mathbf{H}_T^2])^2} = 3 - \frac{\sum_{i \neq j} \omega_{ij}^4}{T^4 (\mathbb{E}[\mathbf{H}_T^2])^2} + \frac{\sum_{i \neq j \neq k \neq l} \omega_{ij} \omega_{jk} \omega_{kl} \omega_{li}}{T^4 (\mathbb{E}[\mathbf{H}_T^2])^2}. \quad (\text{A.1})$$

We already have by the argument in Proposition 1 that

$$\mathbb{E}[\mathbf{H}_T^2] = \frac{1}{T} \cdot (\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2).$$

Thus, to complete the theorem we need to bound the quantites

$$\sum_{i \neq j} \omega_{ij}^4 \quad \text{and} \quad \sum_{i \neq j \neq k \neq l} \omega_{ij} \omega_{jk} \omega_{kl} \omega_{li}. \quad (\text{A.2})$$

For the first term, we will show

$$\sum_{i \neq j} \omega_{ij}^4 \leq 4T \cdot (\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2)^2, \quad (\text{A.3})$$

To see this, we first establish

$$\sum_{i \neq j} \omega_{ij}^2 = 2 \sum_{i \neq j} H_{ij}^2 + H_{ij} H_{ji} = 2T \cdot (\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2).$$

By Lemma 2, we further have

$$\max_{i \neq j} \omega_{ij}^2 \leq 2\|g \otimes q\|_2^2 + 2\langle g \otimes g, q \otimes q \rangle - 4\langle q, g \rangle^2.$$

Put these two bounds together, we find

$$\sum_{i \neq j} \omega_{ij}^4 \leq \max_{i \neq j} \omega_{ij}^2 \cdot \sum_{i \neq j} \omega_{ij}^2 \leq 4T (\|g \otimes q\|_2^2 + \langle g \otimes g, q \otimes q \rangle - 2\langle q, g \rangle^2)^2,$$

which verifies the inequality (A.3).

Next, we aim to show

$$\left| \sum_{i \neq j \neq k \neq l} \omega_{ij} \omega_{jk} \omega_{kl} \omega_{li} \right| \leq 16T \cdot \|(|g| \otimes |q|) \otimes (|g| \otimes |q|)\|_2^2. \quad (\text{A.4})$$

To proceed, we first note

$$\begin{aligned} \left| \sum_{i \neq j \neq k \neq l} \omega_{ij} \omega_{jk} \omega_{kl} \omega_{li} \right| &\leq \sum_{i,j,k,l} |\omega_{ij} \omega_{jk} \omega_{kl} \omega_{li}| \\ &\leq \sum_{i,k} \left(\sum_j |\omega_{ij} \omega_{jk}| \right) \left(\sum_l |\omega_{kl} \omega_{li}| \right) \end{aligned} \quad (\text{A.5})$$

Now, the above Equation (A.5) can be break down into a sum of $2^4 = 16$ terms

$$\begin{aligned} \sum_j |\omega_{ij} \omega_{jk}| &\leq \sum_j \left(\sum_t |q_{t-i}^\circ g_{t-j}^\circ| + |q_{t-j}^\circ g_{t-i}^\circ| \right) \left(\sum_s |q_{s-j}^\circ g_{s-k}^\circ| + |q_{s-k}^\circ g_{s-j}^\circ| \right) \\ &= \sum_{t,s} \sum_j |q_{t-i}^\circ g_{t-j}^\circ| |q_{s-j}^\circ g_{s-k}^\circ| + \sum_{t,s} \sum_j |q_{t-j}^\circ g_{t-i}^\circ| |q_{s-k}^\circ g_{s-j}^\circ| \end{aligned} \quad (\text{A.6})$$

$$+ \sum_{t,s} \sum_j |q_{t-i}^\circ g_{t-j}^\circ| |q_{s-k}^\circ g_{s-j}^\circ| + \sum_{t,s} \sum_j |q_{t-j}^\circ g_{t-i}^\circ| |q_{s-j}^\circ g_{s-k}^\circ| \quad (\text{A.7})$$

Each of the above terms is handled similarly, and we will only showcase one. Denote $|g|$ to be the vector taking entrywise absolute value to the vector g , we know

$$\begin{aligned} \sum_{t,s} \sum_j |q_{t-i}^\circ g_{t-j}^\circ| |q_{s-j}^\circ g_{s-k}^\circ| &= \sum_{t,s} \sum_j |q_{t-i}^\circ g_{j-s}^\circ| |q_{j-t}^\circ g_{s-k}^\circ| \\ &= \sum_j (|q| \otimes |q|)_{j-i}^\circ (|g| \otimes |g|)_{j-k}^\circ \end{aligned}$$

where the first step is by the circular symmetry property in Proposition 4. Using the same idea on all 4-terms in (A.6)-(A.7), we have shown

$$\begin{aligned} \sum_j |\omega_{ij} \omega_{jk}| &= \sum_j (|q| \otimes |q|)_{j-i}^\circ (|g| \otimes |g|)_{j-k}^\circ + \sum_j (|q| \otimes |q|)_{j-k}^\circ (|g| \otimes |g|)_{j-i}^\circ \\ &\quad + \sum_j (|g| \otimes |g|)_{j-i}^\circ (|q| \otimes |q|)_{j-k}^\circ + \sum_j (|g| \otimes |g|)_{j-k}^\circ (|q| \otimes |q|)_{j-i}^\circ. \end{aligned}$$

Now one of the 16-terms in (A.5) takes the form

$$\begin{aligned} &\sum_{i,k} \sum_j (|g| \otimes |q|)_{j-i}^\circ (|g| \otimes |q|)_{j-k}^\circ \sum_l (|g| \otimes |q|)_{l-i}^\circ (|g| \otimes |q|)_{l-k}^\circ \\ &= \sum_{i,k} \sum_j (|g| \otimes |q|)_{k-j}^\circ (|g| \otimes |q|)_{i-j}^\circ \sum_l (|g| \otimes |q|)_{l-i}^\circ (|g| \otimes |q|)_{l-k}^\circ \\ &= \sum_{j,l} \sum_i (|g| \otimes |q|)_{i-j}^\circ (|g| \otimes |q|)_{l-i}^\circ \sum_k (|q| \otimes |q|)_{k-j}^\circ (|q| \otimes |q|)_{l-k}^\circ \\ &= \sum_{j,l} ((|g| \otimes |q|) \otimes (|g| \otimes |q|))_{l-j}^\circ ((|g| \otimes |q|) \otimes (|g| \otimes |q|))_{l-k}^\circ \\ &= T \langle (|g| \otimes |q|) \otimes (|g| \otimes |q|), (|g| \otimes |q|) \otimes (|g| \otimes |q|) \rangle \end{aligned}$$

where the first step uses the circular symmetry property in Proposition 4 again. Note that each of the 16 terms will look similar to the above, which looks like an 8-th moment with a balanced number of g and q 's (4 each).

To bound each of the 16 terms by the representative term, we need to recall Parseval's identity and convolution theorem: for a vector $g \in \mathbb{R}^T$, denote $\widehat{g} \in \mathbb{C}^T$ to denote its discrete-time Fourier transform (DFDT), then

$$\begin{aligned}\langle g, q \rangle &= \frac{1}{T} \sum_{k \in [T]} \overline{\widehat{g}_k} \widehat{q}_k . \\ \overline{g \otimes q} &= \widehat{g}_k \cdot \widehat{q}_k .\end{aligned}$$

Here, for a complex value $z \in \mathbb{C}$, we denote \bar{z} to be its complex conjugate and $\|z\|$ to denote its modulus. Take one of the 16 terms, say

$$\begin{aligned}&\langle (|g| \otimes |g|) \otimes (|g| \otimes |g|), (|q| \otimes |q|) \otimes (|q| \otimes |q|) \rangle \\ &= \frac{1}{T} \sum_{k \in [T]} \overline{\widehat{|g|}_k}^4 \widehat{|q|}_k^4 \\ &\leq \frac{1}{T} \sum_{k \in [T]} \|\widehat{|g|}_k\|^4 \|\widehat{|q|}_k\|^4 = \|(|g| \otimes |q|) \otimes (|g| \otimes |q|)\|_2^2 .\end{aligned}$$

Therefore, we have verified (A.4).

Using the inequalities (A.3) and (A.4) to upper bound (A.1) thus completes the proof.