

PLACEBO INFERENCE ON TREATMENT EFFECTS WHEN THE NUMBER OF CLUSTERS IS SMALL

ANDREAS HAGEMANN

ABSTRACT. I introduce a general, Fisher-style randomization testing framework to conduct nearly exact inference about the lack of effect of a binary treatment in the presence of very few, large clusters when the treatment effect is identified across clusters. The proposed randomization test formalizes and extends the intuitive notion of generating null distributions by assigning placebo treatments to untreated clusters. I show that under simple and easily verifiable conditions, the placebo test leads to asymptotically valid inference in a very large class of empirically relevant models. Examples discussed explicitly are (i) least squares regression with cluster-level treatment, (ii) difference-in-differences estimation, and (iii) binary choice models with cluster-level treatment. A simulation study and an empirical example are provided. The proposed inference procedure is easy to implement and performs well with as few as three treated and three untreated clusters.

JEL classification: C01, C21, C23

Keywords: cluster-robust inference, randomization, permutation

1. INTRODUCTION

It is standard practice in economics to conduct inference that is robust to within-cluster dependence. Units in the same cluster have to be expected to influence one another or are influenced by the same technical, political, or environmental shocks. Several analytical and bootstrap procedures are available to adjust inference for the presence of data clusters. However, as Bertrand, Duflo, and Mullainathan (2004) and others point out—typically in the context of inference on the treatment effect in a difference-in-differences model—the majority of these procedures perform poorly in situations where the number of clusters is small. Such situations are common in empirical practice. They arise, for example, in the analysis of policy reforms, where entire states are treated with the passage of a new law, or in a development context, where the introduction of a new technology affects entire villages. In this paper, I introduce a testing framework based on Fisher (1935) randomization that

Date: March 7, 2018. (First version online: November 14, 2017.) I would like to thank the co-editor, two referees, Federico Bugni, Matias Cattaneo, Kevin Lang, Sarah Miller, Azeem Shaikh, Elie Tamer, Matthew Webb, and several seminar audiences for useful comments and discussions. All errors are my own.

allows for nearly exact inference about the lack of effect of a treatment in the presence of a small number of large clusters. The number of clusters can either grow slightly with the sample size or remain fixed depending on the strength of other assumptions. The framework applies to situations where a binary treatment occurs in some but not all clusters and the treatment effect of interest is identified by between-cluster comparisons.

In a randomized trial, the average effect of a treatment is estimated by comparing the means of treatment and control groups. Computing this comparison of means for all possible ways in which individuals could have been assigned to the two groups generates “placebo” estimates. If treatment has no effect, the placebo estimates have the same distribution as the estimated treatment effect. A Fisher (1935) randomization or placebo test takes these observations as the null distribution to test the “sharp” hypothesis that there is no effect because the difference of treatment and control potential outcomes is zero for each individual. Such tests can be made exact under conventional assumptions. They are particularly attractive when only a small number of observations are available because the set of placebo estimates that have to be computed grows quickly with the sample size. More recently, placebo-type Monte Carlo experiments have been used in empirical economics as informal robustness exercises. I formalize and extend the notion of a placebo test to the cluster case by developing statistics that measure the size of a treatment effect of interest but are amenable to a placebo-like reassignment mechanism. Under simple and easily verifiable conditions, this placebo test leads to asymptotically valid, cluster-robust inference about conventional (non-sharp) null hypotheses in a very large class of empirically relevant models. The proofs rely in part on results of Neuhaus (1993) and Janssen (1997, 2005), who show that consistent permutation tests of certain hypotheses are possible even in situations where the joint distribution of the data is not invariant to permutations under the null hypothesis.

This paper complements recent work by Canay, Romano, and Shaikh (2014) and Ibragimov and Müller (2010, 2016). Canay et al. compute statistics for each cluster separately and obtain null distributions by permuting the signs of these statistics under an approximate symmetry assumption. The downside to their approach is that the parameter of interest has to be identified *within* each cluster. Hence, clusters have to be paired in an ad-hoc manner for difference-in-differences estimation, which reduces the (already small) number of clusters available for inference by half. Ibragimov and Müller (2010, 2016) develop a method that applies to inference about parameters that are identified either within or across clusters; it involves comparing a summary of statistics from each cluster and invoking a small sample result for the t distribution. The method of obtaining statistics used in Ibragimov and Müller (2016) is similar (but not identical) to the one in the present paper. However, as

Canay et al. point out, Ibragimov and Müller’s method tends to be overly conservative and can suffer from low power. As I show in my Monte Carlo study, the placebo test developed here tends to be less conservative and tends to have higher power than both the Canay et al. and Ibragimov and Müller tests when there are as few as three treated and three untreated clusters. In cases with six treated and six untreated clusters, all three methods and cluster-robust versions of the wild bootstrap (see, e.g. Cameron, Gelbach, and Miller, 2008; Webb, 2014) give similar results. The placebo test is therefore especially appropriate for situations with very small numbers of clusters.

Other methods of cluster-robust inference are surveyed in Cameron and Miller (2015) and MacKinnon and Webb (2016). They perform well for moderate numbers of clusters and are typically concerned with adjusting standard t and F tests in the linear regression model; Donald and Lang (2007) derive corrections to standard errors and degrees of freedom under random-effects assumptions. Hansen (2007) and Bester, Conley, and Hansen (2011) use standard cluster-robust covariance matrix estimators in a framework with a small number of clusters but adjust critical values using techniques originally developed by Kiefer and Vogelsang (2002, 2005) for time series. Imbens and Kolesár (2016) derive degrees-of-freedom and standard error corrections following the approach of Bell and McCaffrey (2002). Carter, Schnepel, and Steigerwald (2013) develop measures that can be used to determine degrees-of-freedom corrections. Early papers that recognize the necessity of corrections include Kloek (1981) and Moulton (1990). The present paper differs fundamentally from all of these approaches because it applies to a variety of models other than the linear regression model and derives critical values from the data that automatically account for within-cluster dependence instead of correcting the degrees of freedom of standard critical values.

The paper is organized as follows: Section 2 constructs a general class of test statistics and shows how it can be used to conduct placebo inference. Section 3 establishes the asymptotic validity of the placebo test under explicit regularity conditions. Section 4 verifies these conditions for several empirically relevant situations. Section 5 illustrates the finite sample behavior of the placebo test relative to other methods of inference in simulations, in data from the Current Population Survey, and in data from a large-scale experiment by Dal Bó and Fréchette (2011) on infinitely repeated games. Section 6 concludes. The appendix contains auxiliary results and proofs.

I use the following notation: $1\{\cdot\}$ is the indicator function, $|\cdot|$ is Euclidean norm, and $a \lesssim b$ means that a is bounded by an absolute constant times b . For a matrix A , $|A|$ denotes matrix Euclidean norm $\sqrt{\text{trace}(A'A)}$. The L_p -norm is $\|X\|_p = (\mathbb{E}|X|^p)^{1/p}$. Unless otherwise noted, limits are as $n \rightarrow \infty$. Convergence in distribution is denoted by \rightsquigarrow .

2. PLACEBO INFERENCE WHEN THE NUMBER OF CLUSTERS IS SMALL

This section introduces a general framework for Fisher-style placebo inference with a very small number of clusters. Applications of the framework to least squares regression with cluster-level treatment (Examples 2.1 and 2.3 below) and difference-in-differences estimation (Examples 2.2 and 2.4) are provided.

Consider a situation where data from several large clusters (e.g., counties, regions, schools, firms, or stretches of time) are available. Observations are possibly dependent within clusters but are independent across clusters. Some of the clusters received treatment but others did not. The quantity of interest is a treatment effect or an object related to a treatment effect that can be represented by a scalar parameter β . Because the entire cluster received treatment, this parameter is only determined up to a location shift θ_0 within a treated cluster and only the left-hand side of

$$\theta_1 = \theta_0 + \beta$$

can be identified from such a cluster. If the clusters have similar characteristics, θ_0 can be identified from an untreated cluster. Comparing the two clusters identifies β . As the following two examples illustrate, this situation is prevalent in modern empirical work.

Example 2.1 (Regression with cluster-level treatment). Consider a linear regression model

$$Y_{i,k} = \theta_0 + \beta D_k + \eta'_k X_{i,k} + U_{i,k}, \quad 1 \leq i \leq m_{n,k}. \quad (2.1)$$

Here i indexes individuals within a cluster so that cluster k has $m_{n,k}$ individuals. The goal is to conduct inference about the coefficient β on the treatment dummy D_k indicating whether cluster k received treatment or not. In addition, the regression includes covariates $X_{i,k}$ that vary within each cluster and have coefficients η_k that may vary across clusters such that $E(U_{i,k} | D_k, X_{i,k}) = 0$. Because D_k is either 1 or 0, the data identify $\theta_1 = \theta_0 + \beta$ within a treated cluster and θ_0 within an untreated cluster. \square

Example 2.2 (Difference in differences). Adjust the preceding example to the fixed-effects panel model

$$Y_{t,k} = \theta_0 I_t + \beta I_t D_k + \eta'_k X_{t,k} + \zeta_k + U_{t,k}. \quad (2.2)$$

Now k indexes individuals, t indexes time, the dummy $I_t = 1\{t > t_0\}$ indicates periods after an intervention at known time t_0 , and the dummy D_k indicates whether an intervention occurred. The $m_{n,k}$ observations of individual k now form the k -th cluster and the ζ_k are cluster fixed effects. Provided the components of $X_{t,k}$ vary before or after t_0 , the data again identify $\theta_1 = \theta_0 + \beta$ in a treated cluster and θ_0 in an untreated cluster. \square

The goal of this paper is to develop a simple permutation test of the null hypothesis

$$H_0: \beta = 0$$

(or, equivalently, $H_0: \theta_1 = \theta_0$) that bases its decision on a null distribution obtained by assigning “placebo treatments” to all possible combinations of clusters and leads to asymptotically valid inference when only very few clusters are available. For the asymptotic theory, the number of treated clusters $q_{1,n}$ and the number of untreated clusters $q_{0,n}$ grow with the total sample size $n = \sum_{k=1}^{q_n} m_{n,k}$ (where $q_n = q_{1,n} + q_{0,n}$) but in some circumstances can be fixed if other conditions are strengthened. If q_n grows with n , it does not have to do so at a specific rate. The number of clusters can therefore be very small relative to n . The parameter β does not need to be interpretable by itself and may possibly only determine whether the actual treatment effect of interest is zero or not. For example, suppose (2.1) is the latent model in a binary choice framework with symmetric link function F and $\eta_k \equiv \eta$; then $F(\theta_0 + \beta + \eta'x) - F(\theta_0 + \eta'x)$ for some x is typically the relevant treatment effect but $H_0: \beta = 0$ is still the appropriate hypothesis to test.

The permutation test developed in this paper is based on the idea that each cluster k can provide an estimate $\hat{\theta}_{n,k}$ of either θ_1 or θ_0 depending on whether k received treatment or not. Here, $\hat{\theta}_{n,k}$ can be any estimate that uses data only from cluster k and satisfies a mild regularity condition (Assumption 3.1 in the next section). The condition essentially states that either $\sqrt{n/q_n}(\hat{\theta}_{n,k} - \theta_1)$ or $\sqrt{n/q_n}(\hat{\theta}_{n,k} - \theta_0)$ can be approximated by a well-behaved random variable with positive variance. These variances need not be identical across clusters, so the $\hat{\theta}_{n,k}$ do not have to be standardized. Implicit in the assumption of a positive variance is also the requirement that the individual cluster sizes $m_{n,k}$ and the average cluster size n/q_n grow at a similar rate (in the sense that $m_{n,k}q_n/n$ converges to a positive constant) to ensure that each cluster provides similarly good estimates of either θ_1 or θ_0 . Convergence rates $\nu_n \rightarrow \infty$ other than $\sqrt{n/q_n}$ can also be accommodated as long as $\nu_n(\hat{\theta}_{n,k} - \theta_1)$ or $\nu_n(\hat{\theta}_{n,k} - \theta_0)$ have the required properties. Nonstandard estimators such as the smoothed maximum score estimator of Horowitz (1992), for which ν_n could be as slow as $(n/q_n)^{2/5}$, are therefore also included in the analysis.

I now give two examples of appropriate cluster-level statistics $\hat{\theta}_{n,k}$. In addition to the examples provided here, the methods discussed in this paper apply to estimates arising from, e.g., quantile regression, censored regression, and binary choice models such as probit. Section 4 contains further details and examples.

Example 2.3 (Regression with cluster-level treatment, continued). Suppose cluster k received treatment and cluster l did not. If each cluster is viewed as a separate regression, (2.1) can be written as

$$Y_{i,k} = \theta_1 + \eta'_k X_{i,k} + U_{i,k} \quad \text{and} \quad Y_{i,l} = \theta_0 + \eta'_l X_{i,l} + U_{i,l}.$$

Denote the least squares estimates of the constants θ_1 and θ_0 in these regressions by $\hat{\theta}_{n,k}$ and $\hat{\theta}_{n,l}$. If $\beta = 0$, then $\theta_1 = \theta_0$ and $\hat{\theta}_{n,k} - \hat{\theta}_{n,l} = (\hat{\theta}_{n,k} - \theta_0) - (\hat{\theta}_{n,l} - \theta_0) \approx 0$. Hence, even after rescaling by the convergence rates of the two least squares estimates, the difference $\sqrt{m_{n,k}}(\hat{\theta}_{n,k} - \theta_0) - \sqrt{m_{n,l}}(\hat{\theta}_{n,l} - \theta_0)$ will be bounded in probability, whereas if $\theta_1 > \theta_0$ the display diverges to positive infinity as $n \rightarrow \infty$. Under standard assumptions, $\sqrt{m_{n,k}}(\hat{\theta}_{n,k} - \theta_0)$ and $\sqrt{m_{n,l}}(\hat{\theta}_{n,l} - \theta_0)$ have asymptotic linear representations with positive variance. The same must be true for $\sqrt{n/q_n}(\hat{\theta}_{n,k} - \theta_0)$ and $\sqrt{n/q_n}(\hat{\theta}_{n,l} - \theta_0)$ if $m_{n,k}q_n/n$ converges to a positive constant for every $1 \leq k \leq q_n$. \square

Example 2.4 (Difference in differences, continued). Suppose k was treated but l was not. View each cluster as a separate regression and rewrite (2.2) as

$$Y_{t,k} = \theta_1 I_t + \eta'_k X_{t,k} + \zeta_k + U_{t,k} \quad \text{and} \quad Y_{t,l} = \theta_0 I_t + \eta'_l X_{t,l} + \zeta_l + U_{t,l}.$$

The least squares estimates $\hat{\theta}_{n,k}$ and $\hat{\theta}_{n,l}$ of the slope parameters θ_1 and θ_0 are again suitable cluster-level estimates. \square

Order the data such that indices $1 \leq k \leq q_{1,n}$ correspond to treated clusters and indices $q_{1,n} + 1 \leq k \leq q_n$ correspond to untreated clusters. Define the comparison-of-means function

$$(x_1, \dots, x_{q_n}) \mapsto \bar{T}_n(x_1, \dots, x_{q_n}) = \frac{1}{q_{1,n}} \sum_{k=1}^{q_{1,n}} x_k - \frac{1}{q_{0,n}} \sum_{k=1}^{q_{0,n}} x_{q_{1,n}+k} = \frac{1}{q_{1,n}q_{0,n}} \sum_{k=1}^{q_{1,n}} \sum_{l=1+q_{1,n}}^{q_n} (x_k - x_l).$$

A comparison of means of the cluster estimates $\hat{\theta}_n = (\hat{\theta}_{n,1}, \dots, \hat{\theta}_{n,q_n})$,

$$T_n := \bar{T}_n(\hat{\theta}_n) = \frac{1}{q_{1,n}q_{0,n}} \sum_{k=1}^{q_{1,n}} \sum_{l=1+q_{1,n}}^{q_n} (\hat{\theta}_{n,k} - \hat{\theta}_{n,l}), \quad (2.3)$$

summarizes all possible pairwise comparisons of treated and untreated clusters. The goal is now to construct a consistent permutation test based on T_n . This is possible because the scaled differences of cluster-level estimates $\nu_n(\hat{\theta}_{n,k} - \hat{\theta}_{n,l})$ remain stochastically bounded as n grows under the null hypothesis $\beta = 0$ and diverge to positive or negative infinity depending on whether $\beta > 0$ or $\beta < 0$. The results in Section 3 show that the summary statistic T_n inherits these features.

The idea underlying the permutation test based on (2.3) is that if $H_0: \theta_1 = \theta_0$ is true, then the behavior of the estimates of θ_1 and θ_0 from any pair of clusters should be approximately the same—regardless of whether any of these clusters actually received treatment or not. Hence, *all* differences $\hat{\theta}_{n,k} - \hat{\theta}_{n,l}$, $k \neq l \in \{1, 2, \dots, q_n\}$, should behave similarly under the null. The statistic in (2.3) could therefore be recomputed for every possible permutation of the indices $1, \dots, q_n$ and the distribution of these permuted statistics could be interpreted as an estimate of the null distribution of T_n . This is equivalent to assigning a placebo treatment indicator to every possible subset of $q_{1,n}$ clusters of the q_n available clusters and recording the corresponding placebo realization of T_n . The original statistic T_n could then be compared to the quantiles of the distribution of these placebo statistics. As I show in the next section, this intuition is correct in a very large class of models of practical interest if $q_{1,n}/q_n \rightarrow 1/2$ (see Corollary 3.4) or in situations where the clusters are very similar (Theorem 3.6). In general, however, the placebo realizations of T_n have to be adjusted with a correction factor based on the two-sample variance function

$$(x_1, \dots, x_{q_n}) \mapsto \hat{S}_n^2(x_1, \dots, x_{q_n}) = \frac{1}{q_{1,n}(q_{1,n} - 1)} \sum_{k=1}^{q_{1,n}} \left(x_k - \frac{1}{q_{1,n}} \sum_{l=1}^{q_{1,n}} x_l \right)^2 + \frac{1}{q_{0,n}(q_{0,n} - 1)} \sum_{k=1+q_{1,n}}^{q_n} \left(x_k - \frac{1}{q_{0,n}} \sum_{l=1+q_{1,n}}^{q_n} x_l \right)^2$$

to yield a consistent test (Theorem 3.3). The reason for this adjustment is that a naïve permutation test would require the joint distribution of $\hat{\theta}_n = (\hat{\theta}_{n,1}, \dots, \hat{\theta}_{n,q_n})$ to be—at least asymptotically—invariant to permutation, which is not necessarily satisfied here.

To define the placebo statistics, let Π^* be the set of permutations of $(1, 2, \dots, q_n)$. Each permutation $\pi \in \Pi^*$ depends on n whenever q_n does, but this is suppressed in the notation to prevent clutter. View the cluster-level estimates as maps $k \mapsto \hat{\theta}_{n,k}$ and, for each π , take $\pi(k)$ to be the k -th coordinate of π so that $\hat{\theta}_{n,\pi(k)}$ is computed with data from cluster $\pi(k)$. Denote the permuted vector of cluster-level estimates by

$$\pi \mapsto \pi \hat{\theta}_n = (\hat{\theta}_{n,\pi(1)}, \dots, \hat{\theta}_{n,\pi(q_n)}), \quad \pi \in \Pi^*.$$

This gives rise to the family of (adjusted) placebo statistics

$$T_n(\pi \hat{\theta}_n) = \bar{T}_n(\pi \hat{\theta}_n) \frac{\hat{S}_n(\hat{\theta}_n)}{\hat{S}_n(\pi \hat{\theta}_n)} \tag{2.4}$$

indexed by $\pi \in \Pi^*$. Also note that $T_n(\hat{\theta}_n) = \bar{T}_n(\hat{\theta}_n) = T_n$. I will occasionally use $T_n(\hat{\theta}_n)$ and $\bar{T}_n(\hat{\theta}_n)$ to emphasize the dependence of T_n on $\hat{\theta}_n$.

By construction, the ordering of the numbers $\pi(1), \dots, \pi(q_{1,n})$ and $\pi(q_{1,n} + 1), \dots, \pi(q_n)$ does not change the value of $T_n(\pi\hat{\theta}_n)$. Hence, it suffices to compute the $\binom{q_n}{q_{1,n}}$ placebo statistics indexed by the set $\Pi \subset \Pi^*$ for which the combination $\{\pi(1), \dots, \pi(q_{1,n})\}$ (and therefore the combination $\{\pi(q_{1,n} + 1), \dots, \pi(q_n)\}$) is unique. One way of representing this set is

$$\Pi = \left\{ \pi \in \Pi^* : \pi(1) < \dots < \pi(q_{1,n}) \text{ and } \pi(q_{1,n} + 1) < \dots < \pi(q_n) \right\}.$$

Denote by $\lceil a \rceil$ the smallest integer larger than a and let $|A|$ denote cardinality of a set A . The $1 - \alpha$ quantile

$$c_{n,\alpha} = c_{n,\alpha}(\hat{\theta}_n)$$

of $T_n(\pi\hat{\theta}_n)$ as π varies over Π is then the $\lceil |\Pi|(1 - \alpha) \rceil$ -th largest element of $\{T_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$.

The following procedure provides a generic test for placebo inference with cluster-level statistics. Throughout this paper I consider one-sided tests against the alternative $\beta > 0$ but all results below remain valid for the alternatives $\beta < 0$ and $\beta \neq 0$. See the Remarks immediately below for the necessary modifications.

Algorithm 2.5 (Placebo test). (i) Compute T_n as in (2.3).

(ii) For each permutation $\pi \in \Pi$, compute $T_n(\pi\hat{\theta}_n)$ as in (2.4).

(iii) Reject the null hypothesis $\beta = 0$ in favor of the alternative $\beta > 0$ if T_n exceeds $c_{n,\alpha}$, the $1 - \alpha$ quantile of $\{T_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$.

Remarks. (i) The adjustment factor in (2.4) converges to one if $q_{1,n}/q_n \rightarrow 1/2$. For the placebo test it is therefore sufficient to compare T_n to $\bar{c}_{n,\alpha} = \bar{c}_{n,\alpha}(\hat{\theta}_n)$, the $1 - \alpha$ quantile of $\{\bar{T}_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$, if $q_{1,n} = q_{0,n}$. The remaining parts of this Remark also apply to this version of the placebo test with $\bar{c}_{n,\alpha}$ in place of $c_{n,\alpha}$.

(ii) The test decision can also be made with the p -value

$$p_n(\hat{\theta}_n) = \inf\{p \in (0, 1) : T_n(\hat{\theta}_n) > c_{n,p}(\hat{\theta}_n)\} = \frac{1}{|\Pi|} \sum_{\pi \in \Pi} 1\{T_n(\pi\hat{\theta}_n) \geq T_n(\hat{\theta}_n)\} \quad (2.5)$$

because the statement $T_n > c_{n,\alpha}$ is equivalent to $p_n(\hat{\theta}_n) \leq \alpha$. See Appendix B for proofs of this assertion and the second equality in the preceding display.

(iii) For a one-sided test against $\beta < 0$, reject if $T_n(-\hat{\theta}_n) > c_{n,\alpha}(-\hat{\theta}_n)$ or, equivalently, if $T_n(\hat{\theta}_n)$ is smaller than the $\lceil |\Pi|\alpha \rceil$ -th largest element of $\{T_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$, where $\lceil a \rceil$ is the largest integer smaller than a . For a two-sided test against $\beta \neq 0$, reject if $T_n(\hat{\theta}_n) > c_{n,\alpha/2}(\hat{\theta}_n)$ or $T_n(-\hat{\theta}_n) > c_{n,\alpha/2}(-\hat{\theta}_n)$. A p -value for a two-sided test can be defined as $2 \min\{p_n(\hat{\theta}_n), p_n(-\hat{\theta}_n)\}$.

(iv) From a theoretical standpoint, there is no difference in tests based on Π or Π^* because they lead to identical test decisions at any sample size. However, in practice the permutation statistics should always be computed from Π in order to reduce the number of computations to $\binom{q_n}{q_{1,n}}$ statistics. For example, if there are 6 treated and 6 untreated clusters, then computing all permutations from Π^* requires evaluation of $12! \approx 497$ million statistics whereas Π requires only $\binom{12}{6} = 924$. If Π is too large to make computations feasible, then Π can be replaced with a set Π_M consisting of M draws from the uniform distribution on Π . It is easy to see that the resulting p -value $p_{n,M}(\hat{\theta}_n)$ approximates $p_n(\hat{\theta}_n)$ with arbitrary precision for M large enough in the sense that $p_{n,M}(\hat{\theta}_n) \xrightarrow{P} p_n(\hat{\theta}_n)$ as $M \rightarrow \infty$.

(v) If desired, the test decision in Algorithm 2.5 can be based on a randomized test instead of the nonrandomized test $t \mapsto 1\{t > c_{n,\alpha}\}$. The randomized test can be constructed with the help of a test function $t \mapsto 1\{\varphi_{n,\alpha}(t) \geq U\}$ defined by an independent variable U with a uniform distribution on $[0, 1]$ and the nonrandomized test function

$$\varphi_{n,\alpha}(t) = \begin{cases} 1 & t > c_{n,\alpha} \\ \delta_n & t = c_{n,\alpha} \\ 0 & t < c_{n,\alpha}, \end{cases} \quad \text{where } \delta_n = \frac{|\Pi|\alpha - |\{\pi \in \Pi : T_n(\pi\hat{\theta}_n) > c_{n,\alpha}\}|}{|\{\pi \in \Pi : T_n(\pi\hat{\theta}_n) = c_{n,\alpha}\}|}. \quad (2.6)$$

The randomized test has the advantage that if the distribution of the $\hat{\theta}_{n,1}, \dots, \hat{\theta}_{n,q_n}$ were invariant to permutation under the null hypothesis, then this test would be exact because $P(\varphi_{n,\alpha}(T_n) \geq U) = \alpha$ by a standard argument due to Hoeffding (1952). However, such a test is of little use in practice because rejecting the null if $\varphi_{n,\alpha}(T_n) \geq U$ bases the test decision on a *single* draw from the uniform distribution. Two researchers with identical data sets could therefore arrive at opposite test decisions simply because of two different draws of U or could draw until a desired conclusion was reached.

(vi) It should be noted that, under extreme circumstances, the power of a test that rejects if T_n exceeds a critical value obtained from a permutation distribution can be zero. This is because $T_n \in \{T_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$ and T_n cannot be larger than $c_{n,\alpha}$ whenever $[|\Pi|(1 - \alpha)] = |\Pi|$ or, equivalently, $|\Pi| < \alpha^{-1}$. (This does not change if Π^* is used instead of Π .) Hence, if $\alpha = .05$, one needs 3 treated and 3 untreated clusters to ensure that $|\Pi| = \binom{6}{3} = 20 \geq \alpha^{-1}$ in order for Algorithm 2.5 to have nonzero power. In contrast, the test of Canay et al. (2014) has $2^{\min\{q_1, q_0\}}$ possible permutations and therefore their nonrandomized test needs at least 5 treated and 5 untreated clusters in order to have nonzero power with a significance level of .05. This problem persists in randomized versions of these

tests because the test function $\varphi_{n,\alpha}$ in the preceding display is then bounded above by δ_n , which in turn is strictly smaller than 1 if $|\Pi| < \alpha^{-1}$. \square

The next two sections prove the consistency of Algorithm 2.5 under explicit regularity conditions and explain how these conditions can be verified. Section 5 illustrates the finite sample behavior of the placebo test relative to other methods of inference in three Monte Carlo experiments and an empirical application.

3. ASSUMPTIONS AND GENERIC ASYMPTOTICS

This section introduces a simple high-level condition (Assumption 3.1) that I use in Theorem 3.3 to prove the validity of the placebo test. I then discuss a version of the placebo test that—under stronger assumptions—can be consistent with a fixed number of clusters.

The standard proof strategy for establishing weak convergence of an estimator is to decompose it into a well-behaved leading term (to which distributional limit theory is applied) and a small remainder. I assume that such a decomposition is available for each of the cluster-level statistics $\hat{\theta}_{n,k}$. I discuss this condition in detail immediately below. The next section shows how it can be verified for several standard estimators.

Assumption 3.1. *There are constants θ_1, θ_0 and a sequence $\nu_n \rightarrow \infty$ such that*

$$\nu_n(\hat{\theta}_{n,k} - \theta_{1\{k \leq q_{1,n}\}}) = Z_{n,k} + R_{n,k}, \quad 1 \leq k \leq q_n, \quad (3.1)$$

where (i) $\min\{q_{1,n}, q_{0,n}\} \rightarrow \infty$ and $q_{1,n}/q_n \rightarrow \lambda \in (0, 1)$, (ii) $E Z_{n,k} = 0$ for each k , (iii) $E|Z_{n,k}|^p$ is uniformly bounded in n and k for some $p > 2$, (iv) $Z_{n,1}, \dots, Z_{n,q_n}$ are independent, (v) $\sum_{k=1}^{q_n} \text{Var} Z_{n,k}/q_n$ is bounded away from zero, and (vi)

$$\sum_{k=1}^{q_n} |R_{n,k}| = o_P(\sqrt{q_n}). \quad (3.2)$$

The leading term $Z_{n,k}$ in Assumption 3.1 is typically a byproduct of studying the large sample distribution of the estimator on which the $\hat{\theta}_{n,k}$ are based. Conditions (ii)-(v) impose a simple set of restrictions on this term. In particular, if $\hat{\theta}_{n,k}$ is an estimator with an asymptotic linear representation (such as any of the estimators mentioned in the previous section), then $Z_{n,k}$ plays the role of that representation and $E|Z_{n,k}|^p$ in (iii) can be simply bounded further by inequalities for dependent data; see, among many others, Wu (2005) and El Machkouri, Volný, and Wu (2013) for easy-to-use results that apply to very general dependence structures. If the cluster structure is generated by time or spatial dependence, an informative bound on $E|Z_{n,k}|^p$ will be related to the summability of the auto-covariance

function of the asymptotic linear representation. The condition on $E|Z_{n,k}|^p$ then guarantees that the data within all clusters exhibit short-range dependence, even though the precise form of the dependence does not need to be known to the researcher. This condition can be expected to fail, e.g., if individual time series in a panel (where cluster k contains the time series of the k -th individual) contain unit roots. Condition (iv) imposes that data across clusters are independent. Condition (v) ensures that the leading terms in (3.1) generally do not have a degenerate distribution. It is *not* required that the $Z_{n,k}$ converge in distribution, have the same variance, or that estimates of their variances are available.

Conditions (i) and (vi) are the only restrictions on the total number of clusters q_n and the remainder terms $R_{n,k}$. Together, they impose that the absolute sum of the remainder terms $R_{n,k}$ is small as q_n grows. Neither a specific rate of divergence of q_n nor closed-form expressions for $R_{n,k}$ are needed. However, as it stands, condition (vi) requires knowledge about the behavior of the collection of remainders $R_{n,1}, \dots, R_{n,q_n}$ as q_n grows, whereas in applications it is typically only known that $R_{n,k} \xrightarrow{P} 0$ for each fixed k . Fortunately, as the following result shows, convergence at each k is already enough to guarantee the existence of sequences $q_{1,n}$ and $q_{0,n}$ that satisfy conditions (i) and (vi). Inspection of its proof reveals that such sequences necessarily grow slowly with the total sample size. This can be gleaned from the fact that if q_n were fixed, convergence of the remainder terms to zero would imply (3.2) by the continuous mapping theorem. Hence, condition (vi) is particularly appropriate in the present situation where the number of clusters is small relative to the sample size.

Proposition 3.2. *Suppose that $R_{n,k} \xrightarrow{P} 0$ for each k . Then there are sequences $q_{1,n} \rightarrow \infty$ and $q_{0,n} \rightarrow \infty$ such that $q_{1,n}/(q_{1,n} + q_{0,n}) \rightarrow \lambda \in (0, 1)$ as $n \rightarrow \infty$ and (3.2) holds.*

Remark. The model at hand determines how large q_n can be relative to other quantities. For example, I show in the next section that (3.2) holds in a linear model with very general weak dependence assumptions under the mild condition $q_n/\inf_k m_{n,k}^2 \rightarrow 0$. (Recall that $m_{n,k}$ is the size of the k -th cluster.) \square

The following result establishes the consistency of the placebo test introduced in Algorithm 2.5 under Assumption 3.1. The theorem is stated as a one-sided test to the right but remains valid with obvious modifications as a one-sided test to the left or as a two-sided test.

Theorem 3.3. *Suppose Assumption 3.1 is satisfied. For all $\alpha \in (0, 1)$,*

- (i) *if $\beta = 0$, then $P(T_n > c_{n,\alpha}) \rightarrow \alpha$ and*
- (ii) *if $\beta > 0$, then $P(T_n > c_{n,\alpha}) \rightarrow 1$.*

Remark. The event $\{T_n > c_{n,\alpha}\}$ is invariant to multiplication of both T_n and the placebo statistics $\{T_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$ by the same positive scalar. The proof of Theorem 3.3 exploits this invariance to apply limit theory to properly scaled versions of both T_n and $T_n(\pi\hat{\theta}_n)$. The main issue is the dependence structure induced by the permutation procedure. The proof of Theorem 3.3(i) also appeals to a remarkable result due to Janssen (2005), who shows that consistent permutation tests of certain hypotheses are possible even in situations where the joint distribution of the data is not invariant to permutations under the null hypothesis. Janssen's results in turn rely on an insight of Neuhaus (1993). \square

If the number of treated and untreated clusters is approximately the same in the sense that $q_{1,n}/q_n \rightarrow 1/2$, then $\hat{S}_n(\hat{\theta}_n)/\hat{S}_n(\pi\hat{\theta}_n) \xrightarrow{P} 1$. In that case it suffices to compare T_n to $\bar{c}_{n,\alpha}$, the $1 - \alpha$ quantile of $\{\bar{T}_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$. (See the discussion above equation (2.3) for definitions.) A similar observation was made by Romano (1989) in a related context. This theoretical result is also confirmed by the Monte Carlo simulations in Section 5, which suggest that estimating the adjustment factor can be safely avoided if $q_{1,n} = q_{0,n}$.

Corollary 3.4. *In the situation of Theorem 3.3 with $q_{1,n}/q_n \rightarrow 1/2$,*

- (i) *if $\beta = 0$, then $P(T_n > \bar{c}_{n,\alpha}) \rightarrow \alpha$ and*
- (ii) *if $\beta > 0$, then $P(T_n > \bar{c}_{n,\alpha}) \rightarrow 1$.*

I now briefly discuss the situation where the number of treated clusters $q_{1,n} = q_1$ and untreated clusters $q_{0,n} = q_0$ is fixed. As the following results shows, the result in Theorem 3.3 under the alternative remains unchanged as long as the significance level is not too small. (See the Remarks below Algorithm 2.5 for a discussion why the power is zero if $\alpha < |\Pi|^{-1}$.)

Corollary 3.5. *Suppose there are constants θ_1 and θ_0 such that $\nu_n(\hat{\theta}_{n,k} - \theta_{1\{k \leq q_1\}}) = Z_{n,k} + R_{n,k}$, $1 \leq k \leq q$, where (i) $E Z_{n,k} = 0$ for each k , (ii) $E|Z_{n,k}|^p$ is uniformly bounded in n and k for some $p > 2$, (iii) $Z_{n,1}, \dots, Z_{n,q}$ are independent, (iv) $\text{Var } Z_{n,k}$ is bounded away from zero, $1 \leq k \leq q$, and (vi) $\sum_{k=1}^q |R_{n,k}| \xrightarrow{P} 0$. If $\beta > 0$, then $P(T_n > c_{n,\alpha}) \rightarrow 1\{\alpha \geq |\Pi|^{-1}\}$.*

However, with fixed $q_{1,n} = q_1$ and $q_{0,n} = q_0$, Assumption 3.1 does not seem to put enough conditions on the data to analyze the placebo test under the null hypothesis. To construct proper conditions, assume instead the joint convergence

$$\nu_n(\hat{\theta}_{n,1} - \theta_1, \dots, \hat{\theta}_{n,q_1} - \theta_1, \hat{\theta}_{n,q_1+1} - \theta_0, \dots, \hat{\theta}_{n,q} - \theta_0) \rightsquigarrow (Z_1, \dots, Z_q), \quad (3.3)$$

where the Z_1, \dots, Z_q are independent. In the canonical situation where $Z = (Z_1, \dots, Z_q)$ is mean-zero independent normal with standard deviations $\sigma = (\sigma_1, \dots, \sigma_q)$, extensive

numerical computations suggest that $\sup_{\sigma \geq 0} P_{\sigma}(T_n(Z) > c_{n,\alpha}(Z))$ slightly exceeds α . For example, for $q_1 = q_0 = 5$, I was unable to produce a rejection frequency above .077 for a test with $\alpha = .05$. This worst-case rejection frequency was achieved by combinations of near-infinite and zero variances. Less extreme values of σ produced tests that were, for all practical purposes, correct. Heavily imbalanced q_1 and q_0 lead to worse performance but, as can be expected from Theorem 3.3, the worst-case rejection frequencies were lower with larger q_1 and q_0 . For $q_1 = q_0 = 7$, they did not exceed .065 for tests with a .05 nominal level.

If the clusters are comparable enough such that the (Z_1, \dots, Z_q) in (3.3) are iid with a smooth but otherwise arbitrary distribution, then validity of the placebo test for fixed q can be restored in the sense that the placebo test is asymptotically conservative. Under the alternative, the same conditions are enough to guarantee that the test has power. The results here are given for the unadjusted critical values $\bar{c}_{n,\alpha}$. This is because the underlying asymptotics no longer rely on a normal approximation, which makes variances insufficient to describe the behavior of $(\hat{\theta}_{n,1}, \dots, \hat{\theta}_{n,q})$. Standardization is therefore no longer useful.

Theorem 3.6. *Suppose that (3.3) holds and the Z_k are iid with continuous distribution. For all $\alpha \in (0, 1)$,*

- (i) *if $\beta = 0$, then $\limsup_{n \rightarrow \infty} P(T_n > \bar{c}_{n,\alpha}) \leq \alpha$ and*
- (ii) *if $\beta > 0$, then $P(T_n > \bar{c}_{n,\alpha}) \rightarrow 1\{\alpha \geq |\Pi|^{-1}\}$.*

Remark. (i) The proof of part (i) of the theorem relies heavily on results of Canay et al. (2014), who deal with parameters that are identified within clusters. Their results focus on symmetric data. To the best of my knowledge, the results in part (ii) of the theorem and Corollary 3.5 are new.

(ii) See, e.g., Lemma 1 of Bester et al. (2011) for a result that is appropriate for establishing the distributional convergence (3.3). Results of this type usually require additional assumptions on how the space evolves as the number of data points increases and on the location of the data on that space. As the examples in the next section show, this difficulty can be avoided when working with Assumption 3.1.

(iii) If $q_1 = q_0$, then Corollary 3.4 and Theorem 3.6 together imply that the placebo test is “doubly robust” in the sense that it is consistent if the conditions of either Corollary 3.4 or Theorem 3.6 are met.

(iv) If α is a multiple of $1/q!$, the test $T_n(Z) > \bar{c}_{n,\alpha}(Z)$ with Z as in Theorem 3.6 is uniformly most powerful similar against $Z_1, \dots, Z_{q_1} \sim N(\mu_1, \sigma^2)$ and $Z_{q_1+1}, \dots, Z_q \sim N(\mu_0, \sigma^2)$ with $\mu_1 > \mu_0$, and σ^2 unknown. This “optimality” property follows from Lehmann and Stein (1949). If α is not a multiple of $1/q!$, then randomization can restore similarity.

The tests of Ibragimov and Müller (2010) and Canay et al. (2014) also have optimality properties in a one-sample setting for similarly restrictive classes of alternatives. \square

4. EXAMPLES

This section shows how the central conditions of Assumption 3.1 can be established in a given application. Examples treated explicitly are least squares regression (Example 4.2) and binary choice (Example 4.3) with cluster-level treatment, although the techniques described in this section apply more generally.

For simplicity, I represent each cluster k by a stationary random field

$$\xi_{s,k} = \xi_k(\varepsilon_{s-h,k} : h \in \mathbb{Z}^d), \quad s \in \mathbb{Z}^d,$$

where fields $(\xi_{s,k})_{s \in \mathbb{Z}^d}$ and $(\xi_{s,l})_{s \in \mathbb{Z}^d}$ are independent whenever $k \neq l$. Here, $(\varepsilon_{s,k})_{s \in \mathbb{Z}^d}$ is a field of iid copies of a random variable ε and ξ_k is a measurable, possibly unknown function with values in \mathbb{R}^{d_k} that transforms the input $(\varepsilon_{s,k})_{s \in \mathbb{Z}^d}$ into the output $\xi_{s,k}$. A very large class of commonly used time series ($d = 1$) and spatial models are of this form; see Shao and Wu (2007) and El Machkouri et al. (2013). Many other dependence structures are possible (see Dedecker et al., 2007, and the references therein), but this type of random field leads to particularly simple conditions.

For the class of fields given in the preceding display, the dependence within each cluster can be measured by comparing $\xi_{s,k}$ to a slightly perturbed version of itself. Let $\varepsilon_{0,k}^*$ be an iid copy of ε and let $\varepsilon_{s,k}^* = \varepsilon_{s,k}$ for $s \neq 0$, so that the only difference between $\xi_{s,k}$ and its coupled version $\xi_{s,k}^* = \xi_k(\varepsilon_{s-h,k}^* : h \in \mathbb{Z}^d)$ is the input at coordinate 0. El Machkouri et al. (2013) call a random field p -stable if $\sum_{s \in \mathbb{Z}^d} \|\xi_{s,k} - \xi_{s,k}^*\|_p < \infty$. Stability implies summability of the autocovariance function of $(\xi_{s,k})_{s \in \mathbb{Z}^d}$ whenever $p \geq 2$ and is therefore a “short-range” dependence condition. If, e.g., the process is linear such that $\xi_{s,k} = \sum_{h \in \mathbb{Z}^d} \alpha_h \varepsilon_{s-h,k}$ and $E|\varepsilon|^p < \infty$, then $\|\xi_{s,k} - \xi_{s,k}^*\|_p = |\alpha_s| \|\varepsilon_{0,k} - \varepsilon_{0,k}^*\|_p$ and p -stability is equivalent to $\sum_{h \in \mathbb{Z}^d} |\alpha_h| < \infty$.

The following result is an immediate consequence of the moment bound developed by El Machkouri et al. (2013). It relates moments of weighted sums of vectors at arbitrary coordinates in the field to the stability property. It is useful in the standard case where the leading term $Z_{n,k}$ in Assumption 3.1 is linear. Recall that $\xi_{s,k} \in \mathbb{R}^{d_k}$.

Proposition 4.1 (El Machkouri et al., 2013). *Let S be a finite subset of \mathbb{Z}^d and let $(a_s)_{s \in S}$ be real numbers. If $E|\xi_{0,k}|^p < \infty$ for some $p \geq 2$, then*

$$\left\| \sum_{s \in S} a_s (\xi_{s,k} - E\xi_{s,k}) \right\|_p \leq \sqrt[2p]{d_k} \left(2p \sum_{s \in S} a_s^2 \right)^{1/2} \sum_{s \in \mathbb{Z}^d} \|\xi_{s,k} - \xi_{s,k}^*\|_p.$$

In the following, if $\xi_{s,k}$ is p -stable and has a subvector $V_{s,k}$, then I denote the corresponding subvector of the perturbed version $\xi_{s,k}^*$ by $V_{s,k}^*$. Clearly, if $\xi_{s,k}$ is p -stable, then $V_{s,k}$ must be p -stable as well. I repeatedly use this result and the proposition above in the next two examples.

Example 4.2 (Regression with cluster-level treatment, continued). Write the regression model (2.1) from Examples 2.1 and 2.3 as

$$Y_{i,k} = \theta_{1\{k \leq q_{1,n}\}} + \eta'_k X_{i,k} + U_{i,k}.$$

Let $\tilde{X}_{i,k} = (1, X'_{i,k})'$. As before, view each cluster as a separate regression and take $\hat{\theta}_{n,k}$ to be the first entry of the least squares estimate

$$\left(\sum_{i=1}^{m_{n,k}} \tilde{X}_{i,k} \tilde{X}'_{i,k} \right)^{-1} \sum_{i=1}^{m_{n,k}} \tilde{X}_{i,k} Y_{i,k}$$

of $(\theta_{1\{k \leq q_{1,n}\}}, \eta'_k)'$. If $\Sigma_{n,k} = m_{n,k}^{-1} \sum_{i=1}^{m_{n,k}} \tilde{X}_{i,k} \tilde{X}'_{i,k}$ converges in probability to a positive definite limit Σ_k , then $\sqrt{n/q_n}(\hat{\theta}_{n,k} - \theta_{1\{k \leq q_{1,n}\}})$ can be represented as

$$e'_1 \Sigma_k^{-1} \left(\frac{m_{n,k} q_n}{n} \right)^{1/2} m_{n,k}^{-1/2} \sum_{i=1}^{m_{n,k}} \tilde{X}_{i,k} U_{i,k} + e'_1 (\Sigma_{n,k}^{-1} - \Sigma_k^{-1}) \left(\frac{m_{n,k} q_n}{n} \right)^{1/2} m_{n,k}^{-1/2} \sum_{i=1}^{m_{n,k}} \tilde{X}_{i,k} U_{i,k}$$

where e_1 is a conformable vector with first element equal to 1 and 0 otherwise. Denote the first term by $Z_{n,k}$ and the second term by $R_{n,k}$. This is the decomposition required for Assumption 3.1. Suppose $m_{n,k} q_n / n$ converges to a nonzero constant for every k and $\sup_{n,k} m_{n,k} q_n / n < \infty$ so that all clusters provide similarly good estimates; this holds, for example, if the cluster sizes are constant multiples of one another.

I now discuss conditions (i)-(vi) of Assumption 3.1 in the present context. Condition (ii) is satisfied because U_k is assumed to have conditional mean zero in Example 2.1. For (iii), suppose that for each k , the $(X'_{i,k}, U_{i,k})'$ come from a $2p$ -stable random field in the sense that each i is associated with a possibly unknown coordinate s in the field. Then $\tilde{X}_{i,k} U_{i,k}$ is also part of a stationary random field and an application of the Cauchy-Schwarz inequality after adding and subtracting yields

$$\|\tilde{X}_{i,k} U_{i,k} - \tilde{X}_{i,k}^* U_{i,k}^*\|_p \leq \|U_{i,k}^*\|_{2p} \|\tilde{X}_{i,k} - \tilde{X}_{i,k}^*\|_{2p} + \|\tilde{X}_{1,k}\|_{2p} \|U_{i,k} - U_{i,k}^*\|_{2p},$$

where $\|U_{i,k}^*\|_{2p} = \|U_{i,k}\|_{2p} = \|U_{1,k}\|_{2p}$ because $U_{i,k}^*$ and $U_{i,k}$ are identically distributed and $U_{i,k}$ is stationary in i . Conclude that if $(X'_{i,k}, U_{i,k})'$ is $2p$ -stable, then $\tilde{X}_{i,k} U_{i,k}$ is p -stable.

Proposition 4.1 now implies that there is a constant $C_{p,k}$ depending on only p and k with

$$\|Z_{n,k}\|_p \lesssim |e'_1 \Sigma_k^{-1}| \left\| m_{n,k}^{-1/2} \sum_{i=1}^{m_{n,k}} \tilde{X}_{i,k} U_{i,k} \right\|_p \leq C_{p,k}. \quad (4.1)$$

If the entire set of clusters consists of $2p$ -stable random fields and the eigenvalues of Σ_k are uniformly bounded away from zero, then $\sup_k C_{p,k} < \infty$ and condition (iii) is satisfied. Condition (iv) follows because the random fields are independent and (v) has to be assumed.

For (vi), let $\Sigma_k = \mathbb{E} \tilde{X}_{1,k} \tilde{X}'_{1,k}$ and apply first the Chebyshev inequality and then Proposition 4.1 to bound $P(|\Sigma_{n,k} - \Sigma_k| > \epsilon)$ by ϵ^{-2} times

$$\|\Sigma_{n,k} - \Sigma_k\|_2^2 = \left\| m_{n,k}^{-1} \sum_{i=1}^{m_{n,k}} (\tilde{X}_{i,k} - \mathbb{E} \tilde{X}_{i,k}) \right\|_4^4 \lesssim m_{n,k}^{-2}. \quad (4.2)$$

To transform this into a statement about inverses, write $\Sigma_{n,k}^{-1} - \Sigma_k^{-1} = \Sigma_{n,k}^{-1}(\Sigma_k - \Sigma_{n,k})\Sigma_k^{-1}$ to obtain the bound $|\Sigma_{n,k}^{-1} - \Sigma_k^{-1}| \leq |\Sigma_{n,k}^{-1}| |\Sigma_k^{-1}| |\Sigma_{n,k} - \Sigma_k|$. Because the eigenvalues of Σ_k^{-1} are uniformly bounded away from zero, there is a constant M such that $\sup_k |\Sigma_k^{-1}| \leq M$. Define the events $S_{n,k} = \{|\Sigma_{n,k}^{-1}| \leq M\}$ and $S_n = \cap_k S_{n,k}$. Apply the Markov inequality, then the Hölder inequality with exponents $p/(p-1)$ and p , and finally (4.1) to see that $P(\sum_{k=1}^{q_n} |R_{n,k}| > \epsilon \sqrt{q_n}, S_n)$ is bounded above by a constant multiple of

$$q_n^{-1/2} \sum_{k=1}^{q_n} \mathbb{E} |\Sigma_{n,k} - \Sigma_k| \left| m_{n,k}^{-1/2} \sum_{i=1}^{m_{n,k}} \tilde{X}_{i,k} U_{i,k} \right| \lesssim q_n^{-1/2} \sum_{k=1}^{q_n} \|\Sigma_{n,k} - \Sigma_k\|_{p/(p-1)}.$$

Because $\|\Sigma_{n,k} - \Sigma_k\|_{p/(p-1)} \leq \|\Sigma_{n,k} - \Sigma_k\|_2 \lesssim m_{n,k}^{-1/2}$, the right-hand side of the display is $O(\sqrt{q_n}/\inf_k m_{n,k})$. Now consider $P(S_n^c) \leq \sum_{k=1}^{q_n} P(|\Sigma_{n,k}^{-1}| > M)$. Denote by $\lambda_{\min}(A)$ and $\lambda_{\max}(B)$ the smallest and largest eigenvalues of symmetric matrices A and B . Let $r = \sup_k \sqrt{\text{rank } \Sigma_k}$. Then $P(|\Sigma_{n,k}^{-1}| > M) \leq P(\lambda_{\min}(\Sigma_{n,k}) < r/M)$ because the matrix Euclidean norm is also the 2-Schatten norm. By the Weyl inequality $|\lambda_{\min}(A) - \lambda_{\min}(B)| \leq \lambda_{\max}^{1/2}((A-B)'(A-B)) \leq |A-B|$, this is bounded by

$$P(\lambda_{\min}(\Sigma_k) < r/M + |\Sigma_{n,k} - \Sigma_k|) \leq 1\{\lambda_{\min}(\Sigma_k) < r/M + \delta\} + P(|\Sigma_{n,k} - \Sigma_k| > \delta)$$

for $\delta > 0$. Choose M large and δ small enough so that $\inf_k \lambda_{\min}(\Sigma_k) \geq r/M + \delta$. Now use the Chebyshev inequality and (4.2) to conclude

$$P(S_n^c) \leq \sum_{k=1}^{q_n} P(|\Sigma_{n,k} - \Sigma_k| > \delta) \lesssim \frac{q_n}{\inf_k m_{n,k}^2}.$$

It follows that $\sum_{k=1}^{q_n} |R_{n,k}| = o_P(\sqrt{q_n})$ as long as $\sqrt{q_n}/\inf_k m_{n,k} \rightarrow 0$, which allows for a wide range of sequences $q_n \rightarrow \infty$. \square

Example 4.3 (Binary choice with cluster-level treatment). Suppose the model in the preceding example is the latent model in a binary choice framework where $U_{i,k}$ is independent of treatment assignment and $X_{i,k}$, $U_{i,k}$ has a known, symmetric distribution function F , and $\eta_k \equiv \eta_0$. The treatment effect of interest is $F(\theta_1 + \eta'_0 x) - F(\theta_0 + \eta'_0 x)$ for some x so that $H_0: \theta_1 = \theta_0$ is the appropriate hypothesis to test. Assume that, for each k , the $(Y_{i,k}, X'_{i,k})'$ come from a stationary random field. Only $1\{Y_{i,k} > 0\}$, $X_{i,k}$, and treatment assignment are observed. Let $\psi_{\theta,\eta}(y, x) = (1, x')'(1\{y > 0\} - F(\theta + \eta'x))$ and suppose the moment condition $E\psi_{\theta_1\{k \leq q_{1,n}\}, \eta_0}(Y_{1,k}, X_{1,k}) = 0$ holds for every k . The corresponding Z -estimates $(\hat{\theta}_{n,k}, \hat{\eta}'_{n,k})'$ for the k -th cluster are zeros of

$$(\theta, \eta')' \mapsto \Psi_{n,k}(\theta, \eta) = m_{n,k}^{-1} \sum_{i=1}^{m_{n,k}} \psi_{\theta,\eta}(Y_{i,k}, X_{i,k}).$$

Denote the derivative of $\Psi_{n,k}$ with respect to (θ, η') by $\dot{\Psi}_{n,k}$ and let $\dot{\Psi}_k = E\dot{\Psi}_{n,k}$. If $\dot{\Psi}_k(\theta_{1\{k \leq q_{1,n}\}}, \eta_0)$ is non-singular, then standard arguments for Z -estimators (see, e.g., van der Vaart, 1998) suggest that one can take $Z_{n,k}$ as the first term on the right in

$$\sqrt{n/q_n}(\hat{\theta}_{n,k} - \theta_{1\{k \leq q_{1,n}\}}) = e'_1 \dot{\Psi}_k(\theta_{1\{k \leq q_{1,n}\}}, \eta_0)^{-1} \sqrt{n/q_n} \Psi_{n,k}(\theta_{1\{k \leq q_{1,n}\}}, \eta_0) + R_{n,k},$$

where $R_{n,k}$ is defined as the difference between the left-hand side and $Z_{n,k}$.

Let $\tilde{X}_{i,k} = (1, X'_{i,k})'$ and assume that F is Lipschitz. Using arguments as in the preceding example, apply the Cauchy-Schwarz inequality after repeated adding and subtracting to bound $\|\psi_{\theta_j,\eta}(Y_{i,k}, X_{i,k}) - \psi_{\theta_j,\eta}(Y_{i,k}^*, X_{i,k}^*)\|_p$ by a constant multiple of

$$\|\tilde{X}_{i,k} - \tilde{X}_{i,k}^*\|_p + \|\tilde{X}_{i,k} - \tilde{X}_{i,k}^*\|_{2p} + \|1\{Y_{i,k} \leq 0\} - 1\{Y_{i,k}^* \leq 0\}\|_{2p}.$$

Because $|1\{a < 0\} - 1\{b < 0\}| \leq 1\{|a| < |a - b|\}$ for $a, b \in \mathbb{R}$, the last term in the display is at most $\|1\{|Y_{i,k}| \leq |Y_{i,k} - Y_{i,k}^*|\}\|_{2p} \leq \|1\{|Y_{i,k}| \leq \epsilon\}\|_{2p} + \|1\{|Y_{i,k} - Y_{i,k}^*| > \epsilon\}\|_{2p}$ for every $\epsilon > 0$. Use Lipschitz continuity of F and the Chebyshev inequality to bound this by a constant multiple of $\epsilon^{1/(2p)} + \epsilon^{-1} \|Y_{i,k} - Y_{i,k}^*\|_{2p}$. Choose $\epsilon = \|Y_{i,k} - Y_{i,k}^*\|_{2p}^{2p/(2p+1)}$. Then the preceding display does not exceed a constant multiple of

$$\|\tilde{X}_{i,k} - \tilde{X}_{i,k}^*\|_p + \|\tilde{X}_{i,k} - \tilde{X}_{i,k}^*\|_{2p} + \|Y_{i,k} - Y_{i,k}^*\|_{2p}^{1/(2p+1)}.$$

Hence, $\psi_{\theta_j,\eta}(Y_{i,k}, X_{i,k})$ is p -stable as long as $X_{i,k}$ is $2p$ -stable and $Y_{i,k}$ satisfies the $2p$ -stability condition with $\|\cdot\|_{2p}$ strengthened to $\|\cdot\|_{2p}^{1/(2p+1)}$. (The last condition can be weakened by imposing a stronger stability condition on $X_{i,k}$.) Conclude from Proposition 4.1 that Assumption 3.1(iii) is satisfied. Condition (vi) can be established with similar arguments and Proposition 3.2. The remaining conditions follow as in Example 4.2. \square

5. NUMERICAL RESULTS

This section presents several Monte Carlo experiments to illustrate the small-sample properties of the placebo test in comparison to other methods of inference. I discuss linear regression with cluster-level treatment (Example 5.1), probit regression with cluster-level treatment (Example 5.2), and difference-in-differences estimation using placebo interventions in a sample from the Current Population Survey (Example 5.3). Finally, I use the placebo test to analyze data from an experiment on infinitely repeated games (Example 5.4).

Example 5.1 (Regression with cluster-level treatment, continued). The regression of interest is a simplified version of (2.1) with $\eta_k \equiv \eta$,

$$Y_{i,k} = \theta_0 + \beta D_k + \eta' X_{i,k} + U_{i,k}. \quad (5.1)$$

For the experiment, the errors $U_{i,k}$ have a “circular” dependence structure defined by

$$U_{i,k} = \frac{1}{h+1} \sum_{j=i}^{i+h} \tilde{U}_{c(j),k}, \quad 1 \leq i \leq m_{n,k} \text{ and } c(j) = j \pmod{m_{n,k}}, \quad (5.2)$$

where $h \in \{0, 1, \dots, m_{n,k} - 1\}$ and, for each k and i , $\tilde{U}_{i,k}$ is independently drawn from $N(0, 1)$ if $k \leq q_{1,n}$ and $N(0, 2)$ if $k > q_{1,n}$. This ensures that the $U_{i,k}$ are non-identically distributed and independent across clusters but identically h -dependent within clusters. Each element of the 5-vector of covariates $X_{i,k}$ is also h -dependent and is generated as in the preceding display but with independent draws from $N(0, 1)$ if $k \leq q_{1,n}$ and $\chi_2^2 - 2$ if $k > q_{1,n}$. The default degree of dependence is $h = 10$ and the cluster sizes $m_{n,k}$ are uniformly distributed on $\{15, \dots, 25\}$. For all simulations below I set $\theta_0 = 0$ and $\eta = (1, 1, 1, 1, 1)'$.

In this experiment, I test the hypothesis $H_0: \beta = 0$ against the one-sided alternative $H_1: \beta > 0$ for different values of the treatment effect β , the number of treated and untreated clusters $q_{1,n}$ and $q_{0,n}$, and the degree of dependence h . I consider (i) placebo inference as in Algorithm 2.5, (ii) two-sample inference with the Ibragimov and Müller (2016) t test, (iii) randomized and non-randomized inference with the permutation t test of Canay et al. (2014), (iv) cluster-robust wild bootstrap inference as described in Cameron et al. (2008) and Webb (2014), and (v) inference based on the cluster-robust regression t statistic of Bester et al. (2011). Several other methods for cluster-robust inference exist; see, e.g., MacKinnon and Webb (2017) for one such approach and an overview of the recent literature. They are omitted because these methods tend to be designed for very specific situations.

I now give brief descriptions of methods (i)-(v). For (i), I use the unadjusted critical values $\bar{c}_{n,\alpha}$ defined below Algorithm 2.5 whenever $q_{1,n} = q_{0,n}$ and use Algorithm 2.5 otherwise; the

results with adjusted critical values at $q_{1,n} = q_{0,n}$ were very similar and are therefore not reported.

For (ii), the Ibragimov-Müller test compares the absolute value of

$$\frac{\bar{T}_n(\hat{\theta}_n)}{\hat{S}_n(\hat{\theta}_n)}$$

to the $1 - \alpha$ quantile of a t distribution with $\min\{q_{1,n}, q_{0,n}\} - 1$ degrees of freedom. They show that this is a valid two-sided test for a range of fixed $q_{1,n}$ and $q_{0,n}$ if the $\sqrt{n}(\hat{\theta}_{n,k} - \theta_0)$ are asymptotically independent and normal, where the range is determined by α . Simulations suggest that the statistic in the preceding display can be used for a one-sided test, which is what I do here for ease of comparison.

For (iii), the Canay-Romano-Shaikh test only applies if $q_{1,n} = q_{0,n}$ and requires the researcher to match a treated cluster to a control cluster to identify β . Because both the group of treated clusters and the group of control clusters are relatively homogenous, I match treated and control clusters at random and estimate $\hat{\beta}_{n,k}$ by least squares in (2.1) for each of the $q_{1,n}$ matched pairs. Let $\bar{\beta}_n = q_{1,n}^{-1} \sum_{k=1}^{q_{1,n}} \hat{\beta}_{n,k}$. The Canay-Romano-Shaikh test compares

$$\frac{\bar{\beta}_n}{\sqrt{\sum_{k=1}^{q_{1,n}} (\hat{\beta}_{n,k} - \bar{\beta}_n)^2}}$$

to the $1 - \alpha$ quantile of the permutation distribution obtained by recomputing this statistic for all possible sign changes of the elements of $(\hat{\beta}_{n,1}, \dots, \hat{\beta}_{n,q_{1,n}})$. They show that their test is valid for fixed $q_{1,n} = q_{0,n}$ if the asymptotic distribution of $\sqrt{n}(\hat{\beta}_{n,1}, \dots, \hat{\beta}_{n,q_{1,n}})$ is asymptotically symmetric and some technical regularity conditions are satisfied. As mentioned in the Remarks below Algorithm 2.5, this test can be randomized. It should also be noted that the experimental design likely overstates the performance of the Canay-Romano-Shaikh test because there are $q_{1,n}!$ ways of matching pairs and hence there are equally many ways of computing the test statistic in the preceding display. In practice one may need to use a multiple testing adjustment to account for this indeterminacy, which can lead to significantly lower rejection rates.

For (iv), I estimate β via least squares in the pooled sample and standardize this estimate with the usual cluster-robust covariance matrix with degrees-of-freedom adjustment $(n - 1)q_n / ((n - d)(q_n - 1))$, where d is the number of controls in the pooled regression. I then compare this estimate to the bootstrap distribution of the t statistic obtained from 199 repetitions of the cluster-robust version of the wild bootstrap using the Webb (2014) 6-point

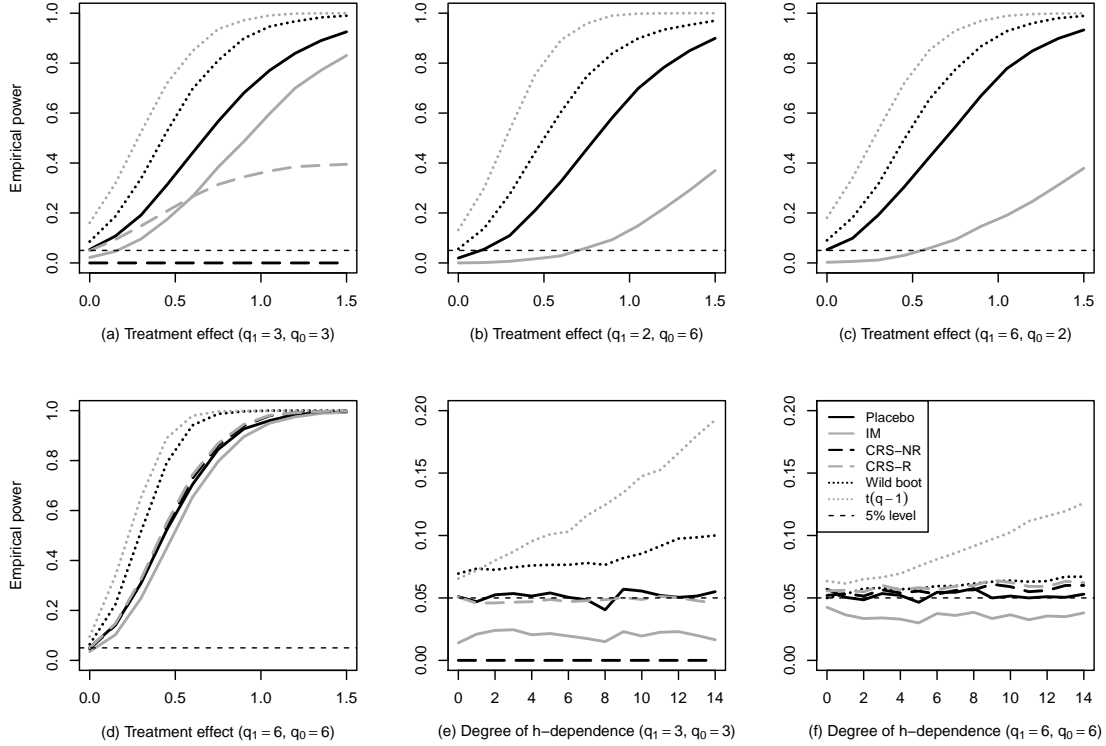


FIGURE 1. Empirical rejection frequencies of the hypothesis $H_0: \beta = 0$ in Example 5.1 using the placebo test (solid black lines), Ibragimov-Müller test (solid grey), Canay-Romano-Shaikh non-randomized (long-dashed black), and randomized permutation test (long-dashed grey), wild bootstrap (dotted black), and $t(q_n - 1)$ critical values (dotted grey) at the 5% level (short-dashed) as a function of the size of the treatment effect for (a) $q_{1,n} = q_{0,n} = 3$, (b) $q_{1,n} = 2, q_{0,n} = 6$, (c) $q_{1,n} = 6, q_{0,n} = 2$, (d) $q_{1,n} = q_{0,n} = 6$, and as a function of the degree of dependence h under H_0 for (e) $q_{1,n} = q_{0,n} = 3$ and (f) $q_{1,n} = q_{0,n} = 6$. The Canay-Romano-Shaikh test does not apply to (b) and (c).

distribution and with the null hypothesis imposed on the data. This procedure is outlined in detail in Cameron et al. (2008).

For (v), I compare the statistic from (iv) to the $1 - \alpha$ quantile of t distribution with $q_n - 1$ degrees of freedom. Bester et al. (2011) show that this test is valid for certain ranges of fixed q_n , where the range depends on α , if the distribution of the covariates is very similar across clusters and other technical regularity conditions are satisfied.

Panels (a)-(d) of Figure 1 show the empirical rejection frequencies of the hypothesis $H_0: \beta = 0$ for methods (i)-(v) at the 5% level (short-dashed line) as a function of $\beta \in \{0, .15, .3, \dots, 1.5\}$ for (a) $q_{1,n} = q_{0,n} = 3$, (b) $q_{1,n} = 2, q_{0,n} = 6$, (c) $q_{1,n} = 6, q_{0,n} = 2$, and (d) $q_{1,n} = q_{0,n} = 6$. Panels (e)-(f) plot the empirical rejection frequencies of the correct

hypothesis $H_0: \beta = 0$ for the same methods as a function of $h \in \{0, 1, \dots, 14\}$ for (e) $q_{1,n} = q_{0,n} = 3$ and (f) $q_{1,n} = q_{0,n} = 6$. Each horizontal coordinate was computed from 2,000 simulations and all five methods were faced with the same data.

As can be seen in panel (a), the placebo test (solid black line) performed well with three treated and three untreated clusters. It rejected in 5.35% of all cases under the null and its power increased quickly as a function of β . In comparison, the Ibragimov-Müller test (solid grey) was quite conservative (2.2%) and had lower power than the placebo test. As pointed out in the Remarks below Algorithm 2.5, the Canay-Romano-Shaikh non-randomized (long-dashed black) test has power equal to zero for all values of β at the 5% level if less than 5 treated clusters are available. The empirically less relevant randomized version of the permutation test (long-dashed grey) was nearly exact at $\beta = 0$ but had lower power than the placebo test. The wild bootstrap rejected a true null hypotheses in 8.55% of all cases. The t test over-rejected more severely with a rate of 16.05%. Because of this size distortion, their rejection frequencies for nonzero values of β should not be interpreted as estimates of their power. In (b), the placebo test behaved more conservatively (1.65%) when $q_{1,n} = 2$ and $q_{0,n} = 6$ but far outperformed the Ibragimov-Müller test for all values of β . The wild bootstrap clearly dominated all other methods of inference in this instance. However, as can be seen in (c), this appears to be an artifact of the experimental design. At $q_{1,n} = 6$ and $q_{0,n} = 2$, the wild bootstrap now over-rejected substantially with an empirical rejection frequency of 9.05% under the null. The power of the Ibragimov-Müller test increased slightly, whereas the placebo test was essentially at the nominal level (5.30%) and had high power. The t test was again less reliable than other tests both in (b) and (c). Once I increased the number of clusters to $q_{1,n} = q_{0,n} = 6$ in (d), the size distortion in the wild bootstrap disappeared and it provided a test with high power. The t test rejected in 9.50% of all cases. The placebo, Canay-Romano-Shaikh, and Ibragimov-Müller tests gave similar results, with the Ibragimov-Müller rejecting the fewest hypotheses. The consistent over-rejection of the wild bootstrap in this example can also be seen in (e), where the rejection frequency of the wild bootstrap increased with the degree of dependence. The over-rejection of the t test became much more severe with h . In contrast, the placebo test and the randomized Canay-Romano-Shaikh test were essentially at the nominal level for all values of h . The Ibragimov-Müller test was very conservative in all cases but this improved when I increased the number of clusters to $q_{1,n} = q_{0,n} = 6$ in (f). Here, the t test over-rejected but was less affected by the degree of dependence. All other methods were nearly exact.

I also experimented with the distribution of the errors and covariates (not shown). As can be expected from the Monte Carlo simulations of MacKinnon and Webb (2016, 2017),

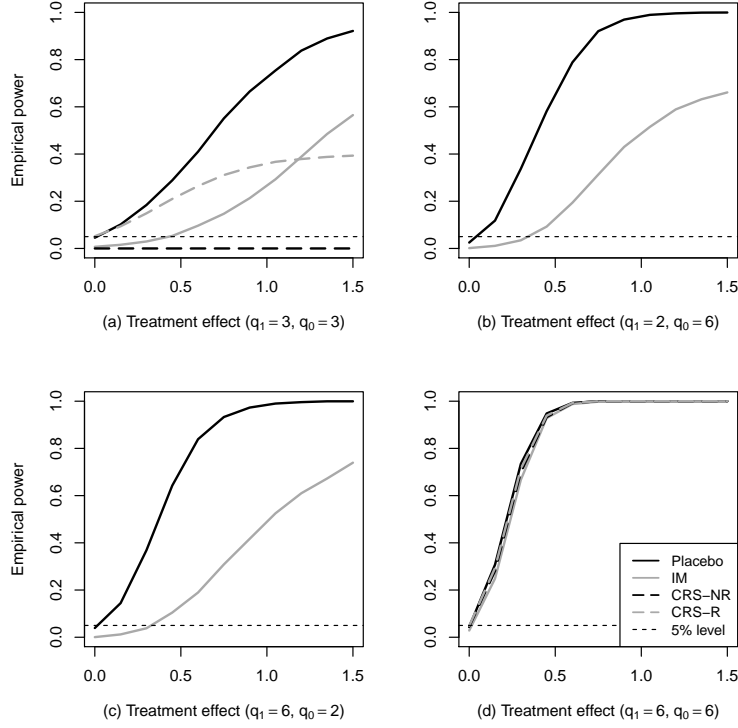


FIGURE 2. Empirical rejection frequencies of the hypothesis $H_0: \beta = 0$ in Example 5.2 with labels as in Figure 1 as a function of the size of the treatment effect for (a) $q_{1,n} = q_{0,n} = 3$, (b) $q_{1,n} = 2, q_{0,n} = 6$, (c) $q_{1,n} = 6, q_{0,n} = 2$, and (d) $q_{1,n} = q_{0,n} = 6$.

the performance of the wild bootstrap improved considerably if the clusters were more homogenous. Similarly, comparing the cluster-robust least squares t statistic to the $t(q_n - 1)$ distribution worked much better once the homogeneity conditions of Bester et al. (2011) were satisfied and the number of clusters was larger. The other methods were less affected by changes in the distribution of the errors and covariates. \square

Example 5.2 (Binary choice with cluster-level treatment, continued). This example uses (5.1) as the latent model in a probit regression. The data are as in Example 5.1 except that the $\tilde{U}_{i,k}$ variables used to generate the $U_{i,k}$ in (5.2) are now standard normal for all k . In addition, the cluster sizes are uniformly distributed on $\{350, \dots, 500\}$ to facilitate computation. Figure 2 repeats the experiments in Figure 1(a)-(d) in the probit case for $\beta \in \{0, .05, .1, \dots, 1.5\}$ with the placebo test, the Ibragimov-Müller test, and the Canay-Romano-Shaikh test. The wild bootstrap and the Bester-Conley-Hansen t test are primarily designed for linear models and are therefore not included here. The placebo test is




FIGURE 3. Empirical rejection frequencies of the hypothesis $H_0: \beta = 0$ in Example 5.3 with labels as in Figure 1 as a function of the size of the treatment effect for (a) $q_{1,n} = q_{0,n} = 3$ and (b) $q_{1,n} = q_{0,n} = 6$, and (c) as a function of the number of clusters under H_0 .

again seen to provide a test near nominal level at (a) $q_{1,n} = q_{0,n} = 3$, (b) $q_{1,n} = 2, q_{0,n} = 6$, and (c) $q_{1,n} = 6, q_{0,n} = 2$, whereas the Ibragimov-Müller test behaved considerably more conservatively. At $q_{1,n} = q_{0,n} = 6$, all methods controlled size and had high power. \square

Example 5.3 (Difference in differences, continued). For this experiment, I follow Bertrand et al. (2004) and use data from the Merged Outgoing Rotation Group of the Current Population Survey (CPS) of women in their fourth interview month. They consider women between the ages 25 and 50 with strictly positive earnings. I extract data on weekly earnings, age, education, and state of residence. The outcome of interest is log weekly earnings in a difference-in-differences regression as in (2.2). The controls in this regression are a quadratic polynomial in age, dummies for education (less than high school, high school, some college, and college or more), and state and year fixed effects. I focus on the years 1989-1991. This sample contains 86,006 observations with an average of approximately 573 women per state-year. For the experiment, I sampled q_n states at random from the 50 states and each state over time is viewed as a single cluster. I then assigned treatment in 1990 to half of these clusters even though no treatment was received in the actual data at this time.

Panels (a)-(b) of Figure 3 show the empirical rejection frequencies of the hypothesis $H_0: \beta = 0$ for the methods of inference used in Example 5.1 at the 5% level (short-dashed line) as a function of $\beta \in \{0, .015, .03, \dots, 0.15\}$ for (a) $q_{1,n} = q_{0,n} = 3$ and (b) $q_{1,n} = q_{0,n} = 6$. I generated the data for these two panels by adding β to the outcomes in treated states for years 1990 and 1991. Panel (c) plots the empirical rejection frequencies of the correct hypothesis $H_0: \beta = 0$ for the same methods as a function of $q_n \in \{6, 8, \dots, 16\}$ with $q_{1,n} = q_{0,n}$. Each horizontal coordinate was computed from 2,000 simulations and all five

methods were faced with the same data. I matched the clusters for the Canay-Romano-Shaikh test by size. As can be seen in panel (a), the CPS data was considerably less challenging than the simulated data from Example 2.1. The placebo test, the wild bootstrap, and the randomized Canay-Romano-Shaikh test provided an essentially exact test, although the placebo test had significantly higher power for most values of β . The Ibragimov-Müller test was conservative but surpassed the wild bootstrap and the Canay-Romano-Shaikh test in terms of power as β increased. The t test again over-rejected. All methods improved when I increased the number of clusters to 12 in (b). Notable here is the Bester-Conley-Hansen t test, which marginally over-rejected but provided a test with very high power. As panel (c) shows, all methods behaved quite similarly as the number of clusters increased further.

I also experimented with different years in this data set (not shown), but the results were similar. Furthermore, deviating from more standard tools such as the wild bootstrap and the Bester-Conley-Hansen t test did not seem necessary in the CPS data if a moderate and relatively balanced number of treatment and control clusters was available. \square

Example 5.4 (Infinitely repeated games; Dal Bó and Fréchette, 2011). In this empirical application, I follow Ibragimov and Müller (2016) and reanalyze a lab experiment of Dal Bó and Fréchette (2011), who measure cooperation between players in an infinitely repeated prisoner’s dilemma with a continuation probability of either $\delta = .5$ or $\delta = .75$. Cooperation resulted in a payoff of $R = 32, 40, \text{ or } 48$. Each (δ, R) combination was played in a separate experiment and three sessions per (δ, R) combination were held. A total of 37,042 games were completed with an average of approximately 2,058 games per session. I view each session as a cluster.

Cooperation can be supported as a subgame-perfect equilibrium action for all (δ, R) combinations other than $(.5, 32)$. (See Dal Bó and Fréchette’s discussion for a more precise statement.) These cooperative equilibria are Pareto efficient. The rate of cooperation $\theta(\delta, R)$ should therefore increase if either δ increases to δ' or R increases to R' . It is then natural to test hypotheses of the form $H_0: \beta = 0$ against $H_1: \beta > 0$ for $\beta = \theta(\delta', R) - \theta(\delta, R)$ or $\beta = \theta(\delta, R') - \theta(\delta, R)$, where each increase of R by eight for fixed δ or each increase in δ from .5 to .75 for fixed R is viewed as a treatment. This leads to seven separate null hypotheses. For each hypothesis, three treated sessions and three control sessions are available, which makes the data amenable to the placebo test.

The two panels of Table 1 reproduce the estimates of $\theta(\delta, R)$ from Dal Bó and Fréchette (2011, Table 3, bottom-right panel) but differ in the way inference was performed. For panel (a), I follow Dal Bó and Fréchette use probit regressions with a dummy for one of the two

TABLE 1. Empirical frequency of cooperation in Dal Bó and Fréchet (2011, Table 3, bottom-right panel) with tests of equal vs. more (“<” in rows, “^” in columns) cooperation between experimental setups using (a) cluster-robust standard errors and (b) Algorithm 2.5.

$\delta \backslash R$	(a) Dal Bó and Fréchet (2011)			(b) Placebo test		
	32	40	48	32	40	48
.50	.0982 <** ^**	.1798 <*** ^***	.3529 <*** ^***	.0982 <* ^*	.1798 <** ^**	.3529 <*** ^**
.75	.2025 <***	.5871 <*	.7642 <*	.2025 <***	.5871 <	.7642 <

Note: * indicates significance at 10%, ** at 5%, *** at 1%

(δ, R) combinations under consideration and compute cluster-robust standard errors. Dal Bó and Fréchet report results from two-sided tests. I compute one-sided tests, which makes two of their results more significant (as indicated by the number of asterisks). In either case, the finding is that higher continuation probabilities or payoffs lead to significantly higher rates of cooperation at conventional levels of significance. For panel (b), I apply the placebo test after computing probit regressions on a constant within each of the six sessions under consideration. As can be seen, in each case the difference in cooperation rates loses one level of significance but Dal Bó and Fréchet’s results remain largely intact even when a method designed for a small number of large clusters such as the placebo test is used. This also confirms Ibragimov and Müller’s (2016) finding that most of Dal Bó and Fréchet’s results continue to hold—albeit with larger p values—if their method is applied. \square

In summary, the placebo test performs well even in fairly extreme (but empirically relevant) situations where the number of clusters is very small, the within-cluster correlation is high, and the clusters are very heterogeneous. These findings hold both for simulated and real data sets in a variety of estimation problems.

6. CONCLUSION

I introduce a general, Fisher-style randomization testing framework to conduct nearly exact inference about the lack of effect of a binary treatment in the presence of very few, large clusters when the treatment effect is identified across clusters. The proposed randomization test formalizes and extends the intuitive notion of generating null distributions by assigning placebo treatments to untreated clusters. I show that under simple and easily verifiable conditions, the placebo test leads to asymptotically valid inference in a very large class of empirically relevant models. A simulation study and an empirical example are provided. The proposed inference procedure is easy to implement and performs well with as few as three treated and three untreated clusters.

APPENDIX A. TECHNICAL RESULTS

This section states several auxiliary results that are used in the proofs in the next section. I start with an “in probability” version of a well-known result for deterministic sequences.

Lemma A.1. *For each $\varepsilon > 0$, let $X_{n,\varepsilon}$ be random variable with $X_{n,\varepsilon} \xrightarrow{P} 0$. Then there is a deterministic sequence $\varepsilon_n \rightarrow 0$ such that $X_{n,\varepsilon_n} \xrightarrow{P} 0$.*

To formalize the concept of a permutation distribution, let π be a random draw from a uniform distribution on Π^* such that π is independent of the cluster-level statistics $\hat{\theta}_n = (\hat{\theta}_{n,1}, \dots, \hat{\theta}_{n,q_n})$. Denote probability computed with respect to the randomness in π by P_π . An object like $P_\pi(T_n(\pi\hat{\theta}_n) \leq t)$ is then the same as $P(T_n(\pi\hat{\theta}_n) \leq t \mid \hat{\theta}_n)$. The following result establishes that, after proper rescaling, the permutation distribution uniformly approximates the null distribution of T_n in probability, conditional on the statistics $\hat{\theta}_n$. This is the standard measure of consistency for resampling distributions; see van der Vaart (1998, p. 329). Let

$$(t_1, \dots, t_{q_n}) \mapsto S_n^2(t_1, \dots, t_{q_n}) = \frac{q_{1,n}q_{0,n}}{q_n} \hat{S}_n^2(t_1, \dots, t_{q_n}).$$

Lemma A.2. *If Assumption 3.1 holds with $\theta_1 = \theta_0$, then*

$$\sup_{t \in \mathbb{R}} \left| P \left(\sqrt{\frac{q_{1,n}q_{0,n}}{q_n}} \frac{T_n(\hat{\theta}_n)}{S_n(\hat{\theta}_n)} \leq t \right) - P_\pi \left(\sqrt{\frac{q_{1,n}q_{0,n}}{q_n}} \frac{T_n(\pi\hat{\theta}_n)}{S_n(\hat{\theta}_n)} \leq t \right) \right| \xrightarrow{P} 0.$$

Let g_n be nonrandom, measurable functions. A statement of the form “ $g_n(Y_n, X_n) \rightsquigarrow Z$ in probability, conditional on X_n ” means that the conditional distribution function of $g_n(Y_n, X_n)$ given X_n converges in probability (pointwise at continuity points) to the distribution function of Z . I frequently use the following results in conjunction with the conditional Slutsky lemma (Xiong and Li, 2008, Theorem 3.2) and with independent permutations $Y_n = \pi$ so that $P_\pi(g_n(\pi, X_n) \leq z) = P(g_n(\pi, X_n) \leq z \mid X_n)$. These results are known in the literature but included here for easy reference.

Lemma A.3. *Let g_n be nonrandom, measurable functions. The following are equivalent:*

- (i) $g_n(Y_n, X_n) \rightsquigarrow 0$ in probability, conditional on X_n ,
- (ii) $P(|g_n(Y_n, X_n)| > \varepsilon \mid X_n) \xrightarrow{P} 0$ for all $\varepsilon > 0$, and
- (iii) $g_n(Y_n, X_n) \xrightarrow{P} 0$.

The following result collects several useful consequences of Assumption 3.1 that will be used throughout the proofs in the next section. Let 1_{q_n} be a q_n -vector of ones and

$$\sigma_n^2 = \frac{q_{0,n}}{q_n q_{1,n}} \sum_{k=1}^{q_{1,n}} \mathbb{E} Z_{n,k}^2 + \frac{q_{1,n}}{q_n q_{0,n}} \sum_{k=1+q_{1,n}}^{q_n} \mathbb{E} Z_{n,k}.$$

Lemma A.4. *Suppose Assumption 3.1 holds with $\theta_1 = \theta_0$. Then*

- (i) $S_n^2(\nu_n(\hat{\theta}_n - \theta_0 1_{q_n})) - \sigma_n^2 \xrightarrow{P} 0$ and $S_n^2(\nu_n(\pi \hat{\theta}_n - \theta_0 1_{q_n})) - S_n^2(\pi Z_n) \xrightarrow{P} 0$,
- (ii) *there is a $\delta > 0$ such that $P(S_n(\pi Z_n) > \delta) \rightarrow 1$, and*
- (iii) $S_n^2(\nu_n(\pi \hat{\theta}_n - \theta_0 1_{q_n})) / S_n^2(\pi Z_n) \xrightarrow{P} 1$.

The following two theorems are consequences of deep results on arrays with exchangeable weights due to Janssen (2005). Bars denote averages over q_n in the next two statements.

Theorem A.5 (Janssen 2005, Theorem 4.1). *For arbitrary random arrays $X_{n,1}, \dots, X_{n,q_n}$ with $\min\{q_{1,n}, q_{0,n}\} \rightarrow \infty$ such that $0 < \liminf q_{1,n}/q_{0,n} \leq \limsup q_{1,n}/q_{0,n} < \infty$ and*

$$\frac{\max_{k \leq q_n} (X_{n,k} - \bar{X}_n)^2}{\sum_{k=1}^{q_n} (X_{n,k} - \bar{X}_n)^2} \xrightarrow{P} 0$$

we have, conditional on $X_{n,1}, \dots, X_{n,q_n}$,

$$\sqrt{\frac{q_{1,n} q_{0,n}}{q_n}} \frac{q_{1,n}^{-1} \sum_{k=1}^{q_{1,n}} X_{n,\pi(k)} - q_{0,n}^{-1} \sum_{k=1+q_{1,n}}^{q_n} X_{n,\pi(k)}}{S_n(X_{n,\pi(1)}, \dots, X_{n,\pi(q_n)})} \rightsquigarrow N(0, 1) \quad \text{in } P\text{-probability.}$$

Theorem A.6. *Consider an exchangeable triangular array of real-valued weights $W_{n,1}, \dots, W_{n,q_n}$ such that (i) $\max_{k \leq q_n} |W_{n,k} - \bar{W}_n| \xrightarrow{P} 0$, (ii) $\sum_{k=1}^{q_n} (W_{n,k} - \bar{W}_n)^2 \xrightarrow{P} 1$, (iii) $\sqrt{q_n}(W_{n,1} - \bar{W}_n) \rightsquigarrow W$ with $\mathbb{E}W = 0$ and $\text{Var}W = 1$. Also consider an arbitrary triangular array of real-valued random variables $X_{n,1}, \dots, X_{n,q_n}$ that is independent of the weights and satisfies (iv) $\max_{k \leq q_n} |X_{n,k}| \xrightarrow{P} 0$, (v) $\sum_{k=1}^{q_n} (X_{n,k} - \bar{X}_n)^2 \xrightarrow{P} \sigma^2$, and (vi) $\bar{X}_n \xrightarrow{P} 0$. Then $\sqrt{q_n} \sum_{k=1}^{q_n} W_{n,k} (X_{n,k} - \bar{X}_n)$ converges in distribution to $N(0, \sigma^2)$ in probability, conditional on $X_{n,1}, \dots, X_{n,q_n}$.*

Finally, the central limit theorems above do not require the $X_{n,k}$ to be independent. If they are, then the conditions of the theorem can be simplified using a result of A. Y. Khinchin. I include a proof for completeness.

Lemma A.7 (Khinchin). *A row-wise independent array $\{X_{n,k} : k \leq q_n, n \in \mathbb{N}\}$ satisfies $\max_{k \leq q_n} |X_{n,k}| \xrightarrow{P} 0$ if and only if $\sum_{k=1}^{q_n} P(|X_{n,k}| > \varepsilon) \rightarrow 0$ for every $\varepsilon > 0$.*

APPENDIX B. PROOFS

Proof of (2.5) and the equivalence of $T_n > c_{n,\alpha}$ and $p_n(\hat{\theta}_n) \leq \alpha$. To see the second equality in (2.5), let k_n be the index on the first order statistic of $\{T_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$ strictly smaller than T_n . If no such order statistic exists, then $T_n = \min\{T_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$ and the infimum in (2.5) must be 1. Suppose therefore that k_n exists and let $M = |\Pi|$. Then $p_n = p_n(\hat{\theta}_n)$ is the smallest p such that $(k_n - 1)/M < 1 - p \leq k_n/M$ (see, e.g., van der Vaart, 1998, p. 305) or, equivalently, $p_n = 1 - k_n/M$. Because $k_n = \sum_{\pi \in \Pi} 1\{T_n(\pi\hat{\theta}_n) < T_n\}$, the equality follows.

To see the equivalence of $T_n > c_{n,\alpha}$ and $p_n \leq \alpha$, suppose first that $T_n > c_{n,\alpha}$. Then $\alpha \in \{p : T_n > c_{n,p}\}$ and therefore $p_n = \inf\{p : T_n > c_{n,p}\} \leq \alpha$. Conversely, suppose $p_n \leq \alpha$. Note that $p \mapsto c_{n,p}$ is a right-continuous, decreasing step-function that moves through the order statistics of $\{T_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$ in reverse order as p increases. Hence, if $p_n \in [1/M, 1)$, there is a p such that $c_{n,p}$ is the first order statistic strictly smaller than T_n . The smallest such p is p_n and therefore $T_n > c_{n,p_n} \geq c_{n,\alpha}$. The extreme right-hand side of (2.5) contradicts $p_n < 1/M$ because $T_n \in \{T_n(\pi\hat{\theta}_n) : \pi \in \Pi\}$. Finally, $p_n = 1$ and $p_n \leq \alpha$ contradicts $\alpha \in (0, 1)$. \square

Proof of Lemma A.1. For each $m \in \mathbb{N}$, there is an n_m such that $P(|X_{n,1/m}| > 1/m) < 1/m$ for all $n > n_m$. Without loss of generality, take $n_1 < n_2 < \dots$. For each n , define implicitly m_n as the m that satisfies $n_m \leq n < n_{m+1}$ and let $m_n = 1$ for $n < n_1$. Then $\varepsilon_n = 1/m_n \rightarrow 0$ and, for any given $\delta > 0$, we eventually have $P(|X_{n,\varepsilon_n}| > \delta) \leq P(|X_{n,\varepsilon_n}| > \varepsilon_n) < \varepsilon_n$. \square

Proof of Proposition 3.2. Recall that $\lceil a \rceil$ is the smallest integer greater than a . Let $q_1(\varepsilon) = \lceil 1/\varepsilon \rceil$ and $q_0(\varepsilon) = \lceil q_1(\varepsilon)(1 - \lambda)/\lambda \rceil$. By the continuous mapping theorem,

$$X_{n,\varepsilon} := \frac{\sum_{k=1}^{q_1(\varepsilon)} |R_{n,k}| + \sum_{k=1+q_1(\varepsilon)}^{q_1(\varepsilon)+q_0(\varepsilon)} |R_{n,k}|}{\sqrt{q_1(\varepsilon) + q_0(\varepsilon)}} \xrightarrow{P} 0$$

for every fixed $\varepsilon > 0$. By Lemma A.1, there is a sequence $\varepsilon_n \rightarrow 0$ such that $X_{n,\varepsilon_n} \xrightarrow{P} 0$. Let $q_{1,n} = q_1(\varepsilon_n)$ and $q_{0,n} = q_0(\varepsilon_n)$. Note that $\min\{q_{1,n}, q_{0,n}\} \rightarrow \infty$ and $q_{1,n}/(q_{1,n} + q_{0,n}) \rightarrow \lambda$, as desired. \square

Proof of Lemma A.3. (ii) \Rightarrow (iii): Apply the Lebesgue dominated convergence theorem. (ii) \Leftarrow (iii): The Markov inequality gives

$$P(P(|g_n(Y_n, X_n)| > \varepsilon \mid X_n) > \delta) \leq \delta^{-1} P(|g_n(Y_n, X_n)| > \varepsilon) \rightarrow 0.$$

(i) \Rightarrow (ii): $P(|g_n(Y_n, X_n)| > \varepsilon \mid X_n)$ is at most

$$1 - P(g_n(Y_n, X_n) \leq \varepsilon \mid X_n) + P(g_n(Y_n, X_n) \leq -\varepsilon \mid X_n).$$

The definition of convergence in distribution in probability implies that the right-hand side converges to zero for every $\varepsilon > 0$. (i) \Leftarrow (iii): For every $z > 0$, $P(g_n(X_n, Y_n) \leq z \mid X_n) \xrightarrow{P} 1$ and $P(g_n(X_n, Y_n) \leq -z \mid X_n) \xrightarrow{P} 0$. Since z is arbitrary, $z = 0$ is a discontinuity point. The result follows. \square

Proof of Lemma A.4. Let $\bar{Z}_{1,n} = q_{1,n}^{-1} \sum_{k=1}^{q_{1,n}} Z_{n,k}$, $\bar{Z}_{0,n} = q_{0,n}^{-1} \sum_{k=1+q_{1,n}}^{q_n} Z_{n,k}$, $\bar{Z}_{1,n}(\pi) = q_{1,n}^{-1} \sum_{k=1}^{q_{1,n}} Z_{n,\pi(k)}$, and $\bar{Z}_{0,n}(\pi) = q_{0,n}^{-1} \sum_{k=1+q_{1,n}}^{q_n} Z_{n,\pi(k)}$. For (i), I start by approximating $S_n^2(\nu_n(\hat{\theta}_n - \theta_0 1_{q_n})) = \nu_n^2 S_n^2(\hat{\theta}_n)$. The absolute value of $\nu_n^2 \sum_{k=1}^{q_{1,n}} (\hat{\theta}_{n,k} - \bar{\theta}_{1,n})^2 - \sum_{k=1}^{q_{1,n}} (Z_{n,k} - \bar{Z}_{1,n})^2$ is bounded above by a constant multiple of

$$\left(\sum_{k=1}^{q_n} Z_{n,k}^2 \right)^{1/2} \left(\sum_{k=1}^{q_n} R_{n,k}^2 \right)^{1/2} + \sum_{k=1}^{q_n} R_{n,k}^2.$$

The same bound with $\sum_{k=1}^{q_{1,n}} (Z_{n,k} - \bar{Z}_{1,n})^2$ replaced by $\sum_{k=1+q_{1,n}}^{q_n} (Z_{n,k} - \bar{Z}_{0,n})^2$ applies to $\nu_n^2 \sum_{k=1+q_{1,n}}^{q_n} (\hat{\theta}_{n,k} - \bar{\theta}_{0,n})^2$; notice also that this bound is invariant to permutation. Use $\sum_{k=1}^{q_n} R_{n,k}^2 / q_n \leq (\sum_{k=1}^{q_n} |R_{n,k}|)^2 / q_n \xrightarrow{P} 0$ to conclude that $S_n^2(\nu_n(\hat{\theta}_n - \theta_0 1_{q_n})) - S_n^2(Z_n) \xrightarrow{P} 0$ as long as $q_n^{-1} \sum_{k=1}^{q_n} Z_{n,k}^2$ is bounded in probability. To this end, apply a weak law of large numbers for arrays (Durrett, 2010, Theorem 2.2.6, p. 52) to see that $q_n^{-1} \sum_{k=1}^{q_n} (Z_{n,k}^2 - \mathbb{E}Z_{n,k}^2) \xrightarrow{P} 0$. The moment conditions are easily satisfied in view of $\sum_{k=1}^{q_n} \mathbb{E}|Z_{n,k}|^p = O(q_n)$ and

$$\left| \frac{1}{q_n} \sum_{k=1}^{q_n} \mathbb{E}(Z_{n,k}^2 - \mathbb{E}Z_{n,k}^2) \mathbb{1}\{Z_{n,k}^2 - \mathbb{E}Z_{n,k}^2 > q_n\} \right| \leq \frac{2^p}{q_n^{p/2}} \sum_{k=1}^{q_n} \mathbb{E}|Z_{n,k}|^p = O(q_n^{1-(p/2)}). \quad (\text{B.1})$$

Chebyshev's inequality yields $\bar{Z}_{1,n} = q_{1,n}^{-1} \sum_{k=1}^{q_{1,n}} Z_{n,k} \xrightarrow{P} 0$ and $\bar{Z}_{0,n} = q_{0,n}^{-1} \sum_{k=1+q_{1,n}}^{q_n} Z_{n,k} \xrightarrow{P} 0$. To show the second part of (i), note that the bound for $\nu_n^2 \sum_{k=1}^{q_{1,n}} (\hat{\theta}_{n,k} - \bar{\theta}_{1,n})^2$ also applies to $\nu_n^2 \sum_{k=1}^{q_{1,n}} (\hat{\theta}_{n,\pi(k)} - \bar{\theta}_{1,n}(\pi))^2$ and $\nu_n^2 \sum_{k=1+q_{1,n}}^{q_n} (\hat{\theta}_{n,\pi(k)} - \bar{\theta}_{0,n}(\pi))^2$. That bound converges to zero unconditionally in probability.

For (ii), we have $\mathbb{E}(\bar{Z}_{1,n}(\pi) \mid \pi) = 0$ and, by the law of total variance, $\text{Var}(\bar{Z}_{1,n}(\pi))$ equals

$$q_{1,n}^{-2} \mathbb{E} \sum_{k=1}^{q_{1,n}} \text{Var}(Z_{n,\pi(k)} \mid \pi) \leq q_{1,n}^{-2} \mathbb{E} \sum_{k=1}^{q_n} \text{Var}(Z_{n,\pi(k)} \mid \pi) = q_{1,n}^{-2} \mathbb{E} \text{Var} \left(\sum_{k=1}^{q_n} Z_{n,\pi(k)} \mid \pi \right).$$

Because the sum on the far right is invariant to permutation, conclude that $\text{Var}(\bar{Z}_{1,n}(\pi)) \leq q_{1,n}^{-2} \sum_{k=1}^q \text{Var} Z_{n,k} \rightarrow 0$. It follows that $\bar{Z}_{1,n}(\pi) \xrightarrow{P} 0$ and, by a similar argument, $\bar{Z}_{0,n}(\pi) \xrightarrow{P} 0$ unconditionally. Use these two results, the law of large numbers for $q_n^{-1} \sum_{k=1}^{q_n} Z_{n,k}^2$ from above, and the fact that $S_n^2(\pi Z_n)$ is at least as large as

$$\min \left\{ \frac{q_{1,n}}{q_n(q_{0,n} - 1)}, \frac{q_{0,n}}{q_n(q_{1,n} - 1)} \right\} \left(\sum_{k=1}^{q_n} Z_{n,k}^2 - q_{1,n} \bar{Z}_{1,n}(\pi)^2 - q_{0,n} \bar{Z}_{0,n}(\pi)^2 \right)$$

to see that the probability that this expression differs by more than $\varepsilon/2$ from $\min\{(1 - \lambda)/\lambda, \lambda/(1 - \lambda)\} \sum_{k=1}^{q_n} \text{Var} Z_{n,k}/q_n$ converges to zero unconditionally for every $\varepsilon > 0$. By assumption, the latter expression is eventually larger than a small enough ε . Conclude that $P(S_n^2(\pi Z_n) > \varepsilon/2)$ approaches one unconditionally.

Arguing along the same lines, we have $S_n^2(\pi Z_n) \leq \max\{q_{0,n}/(q_{1,n} - 1), q_{1,n}/(q_{0,n} - 1)\} \sum_{k=1}^{q_n} Z_{n,k}^2/q_n$ and therefore $S_n^2(\pi Z_n)$ is unconditionally bounded in probability. Conclude that $1 - S_n^2(\nu_n(\hat{\theta}_n - \theta_0 1_{q_n}))/S_n^2(\pi Z_n)$ converges to zero conditionally in probability. \square

Proof of Lemma A.7. Apply the inequality $P(\max_{k \leq q_n} X_{n,k}^2 > \varepsilon) \leq \sum_{k=1}^{q_n} P(X_{n,k}^2 > \varepsilon)$, then independence, and finally $\log x \leq x - 1$ in

$$1 - \sum_{k=1}^{q_n} P(X_{n,k}^2 > \varepsilon) \leq P\left(\max_{k \leq q_n} X_{n,k}^2 \leq \varepsilon\right) = \prod_{k=1}^{q_n} (1 - P(X_{n,k}^2 > \varepsilon)) \leq e^{-\sum_{k=1}^{q_n} P(X_{n,k}^2 > \varepsilon)} \leq 1$$

to obtain the desired result. \square

Proof of Lemma A.2. By Assumption 3.1, $\sigma_n^2 = O(1)$ and $\sigma_n^2 \geq \min\{q_{0,n}/q_{1,n}, q_{1,n}/q_{0,n}\} \times \sum_{k=1}^{q_n} \text{Var} Z_{n,k}/q_n > 0$, which implies $1/\sigma_n = O(1)$. Let $w_{n,k} = \sqrt{q_{1,n}q_{0,n}}(q_{1,n}^{-1}1\{k \leq q_{1,n}\} - q_{0,n}^{-1}1\{k > q_{1,n}\})$. Use the fact that the limit superior of $\sup_{k \in \mathbb{N}} |w_{n,k}|$ is finite to deduce that

$$\left| \frac{1}{\sqrt{q_n}} \sum_{k=1}^{q_n} w_{n,k} R_{n,k} \right| \leq \frac{\sup_{k \in \mathbb{N}} |w_{n,k}|}{\sqrt{q_n}} \sum_{k=1}^{q_n} |R_{n,k}| \quad (\text{B.2})$$

converges to zero in probability and therefore

$$\sqrt{\frac{q_{1,n}q_{0,n}}{q_n}} \frac{T_n(\nu_n(\pi \hat{\theta}_n - \theta_0 1_{q_n}))}{\sigma_n} = \frac{1}{\sqrt{q_n}} \sum_{k=1}^{q_n} \frac{w_{n,k} \nu_n(\hat{\theta}_{n,k} - \theta_0)}{\sigma_n} = \frac{1}{\sqrt{q_n}} \sum_{k=1}^{q_n} \frac{w_{n,k} Z_{n,k}}{\sigma_n} + o_P(1).$$

Assumption 3.1(iii) ensures that the first term on the right of satisfies the Lyapunov condition. The Lindeberg-Feller central limit theorem and the Slutsky lemma then yield weak convergence of the display to $N(0, 1)$. By Lemma A.4,

$$\sqrt{\frac{q_{1,n}q_{0,n}}{q_n}} \left(\frac{T_n(\nu_n(\pi \hat{\theta}_n - \theta_0 1_{q_n}))}{S_n(\nu_n(\pi \hat{\theta}_n - \theta_0 1_{q_n}))} - \frac{T_n(\nu_n(\pi \hat{\theta}_n - \theta_0 1_{q_n}))}{\sigma_n} \right) \xrightarrow{P} 0.$$

I now turn to the permutation statistics. Define

$$\tilde{T}_n(\pi \hat{\theta}_n) = \frac{T_n(\pi \hat{\theta}_n)}{S_n(\hat{\theta}_n)} \frac{S_n(\pi \hat{\theta}_n)}{S_n(\pi Z_n)} = \sum_{k=1}^{q_n} \frac{w_{n,k}(\hat{\theta}_{n,\pi(k)} - \theta_0)}{S_n(\pi Z_n)}.$$

By Lemma A.4, we can work on the set $\{S_n(\pi Z_n) > \delta\}$. Just like the original statistic, $\sqrt{q_{1,n}q_{0,n}/q_n}\nu_n\tilde{T}_n(\pi\hat{\theta}_n)$ can be decomposed into

$$\frac{1}{\sqrt{q_n}} \sum_{k=1}^{q_n} \frac{w_{n,k}\nu_n(\hat{\theta}_{n,\pi(k)} - \theta_0)}{S_n(\pi Z_n)} = \frac{1}{\sqrt{q_n}} \sum_{k=1}^{q_n} \frac{w_{n,k}(Z_{n,\pi(k)} + R_{n,\pi(k)})}{S_n(\pi Z_n)}.$$

The absolute value of $\sum_{k=1}^{q_n} w_{n,k}R_{n,\pi(k)}/\sqrt{q_n}$ is bounded above the right-hand side of (B.2), which is invariant to permutation. Conclude that $\sqrt{q_{1,n}q_{0,n}/q_n}\nu_n\tilde{T}_n(\pi\hat{\theta}_n)$ is within $o_P(1)$ of

$$\frac{1}{\sqrt{q_n}} \sum_{k=1}^{q_n} \frac{w_{n,k}Z_{n,\pi(k)}}{S_n(\pi Z_n)}.$$

The conditional asymptotic distributions of the display and $\sqrt{q_{1,n}q_{0,n}/q_n}\tilde{T}_n(\pi\hat{\theta}_n)$ therefore coincide by the conditional Slutsky lemma. However, because permutation renders the $Z_{n,\pi(k)}$ dependent, the Lindeberg-Feller central limit theorem is no longer appropriate. I therefore apply multiplier central limit theory from Theorem A.5. In view of (B.1), it is enough to show $\bar{Z}_n \xrightarrow{P} 0$ and $\max_{k \leq q_n} Z_{n,k}^2/q_n \xrightarrow{P} 0$. The first condition is an immediate consequence of $\text{Var } \bar{Z}_n = O(q_n^{-1})$. By Lemma A.7, the second condition holds because $\sum_{k=1}^{q_n} P(Z_{n,k}^2 > q_n \varepsilon) \leq \sum_{k=1}^{q_n} E|Z_{n,k}|^p / (q_n \varepsilon)^{p/2} \rightarrow 0$. It follows from Lemma A.3 and the conditional Slutsky lemma that $\sqrt{q_{1,n}q_{0,n}/q_n}\nu_n\tilde{T}_n(\pi\hat{\theta}_n) \rightsquigarrow Z \sim N(0, 1)$ in probability. Apply Lemma A.4 and the conditional Slutsky lemma again to see that

$$\left(1 - \frac{S_n(\pi Z_n)}{S_n(\nu_n(\pi\hat{\theta}_n - \theta_0 \mathbf{1}_{q_n}))}\right) \sqrt{\frac{q_{1,n}q_{0,n}}{q_n}} \nu_n \tilde{T}_n(\pi\hat{\theta}_n) \rightsquigarrow 0 \times Z = 0$$

in probability conditionally and therefore $\sqrt{q_{1,n}q_{0,n}/q_n}\nu_n T_n(\pi\hat{\theta}_n) \rightsquigarrow N(0, 1)$ in probability conditionally.

Use Pólya's theorem and the triangle inequality to see that the desired conclusion follows if $\sup_{t \in \mathbb{R}} |P_\pi(\sqrt{q_{1,n}q_{0,n}/q_n}\nu_n T_n(\pi\hat{\theta}_n) \leq t) - \Phi(t)| \xrightarrow{P} 0$, where Φ is the standard normal distribution function. But this is already a consequence of the pointwise conditional convergence by a well known result (see, e.g, van der Vaart, 1998, Problem 23.1, p. 339); it can shown by constructing diagonal subsequences that ensure almost sure convergence uniformly on a countable dense subset of the real line. \square

Proof of Theorem 3.3. Under the null, $\sqrt{q_{1,n}q_{0,n}/q_n}c_{n,\alpha}/S_n(\hat{\theta}_n)$ converges in probability to $\Phi^{-1}(1 - \alpha)$ by the properties of quantile functions (van der Vaart, 1998, Lemma 21.2, p. 305) and Lemma A.2. Then

$$\sqrt{\frac{q_{1,n}q_{0,n}}{q_n}} \frac{T_n - c_{n,\alpha}}{S_n(\hat{\theta}_n)} \rightsquigarrow N(0, 1) - \Phi^{-1}(1 - \alpha)$$

by the Slutsky lemma and therefore $P(T_n > c_{n,\alpha}) \rightarrow \alpha$.

Consider the alternative $\beta > 0$. Suppose $\min\{q_{1,n}, q_{0,n}\} \rightarrow \infty$. Let $a_{n,k} = q_{1,n}^{-1}1\{k \leq q_{1,n}\} - q_{0,n}^{-1}1\{k > q_{1,n}\}$ and write $T_n = \beta + \sum_{k=1}^{q_n} a_{n,k}(Z_{n,k} + R_{n,k}) = \beta + o_P(1)$ by arguments as in the proof of Lemma A.4. Turning to the permutation statistics, notice that $q_{1,n}^{-1} \sum_{k=1}^{q_{1,n}} \hat{\theta}_{n,\pi(k)} - \theta_0 - \beta q_{1,n}/q_n = \beta q_{1,n}^{-1} \sum_{k=1}^{q_{1,n}} (1\{\pi(k) \leq q_{1,n}\} - q_{1,n}/q_n) + o_P(1)$ using arguments from before. The variance of the first term is bounded above by β^2 times

$$\frac{1}{q_{1,n}} + \left(\frac{q_{1,n}}{q_n} \frac{q_{1,n} - 1}{q_n - 1} - \left(\frac{q_{1,n}}{q_n} \right)^2 \right) \frac{q_{1,n} - 1}{q_{1,n}} \rightarrow 0.$$

The same argument applies to $q_{0,n}^{-1} \sum_{k=1+q_{1,n}}^{q_n} \hat{\theta}_{n,\pi(k)}$. Conclude that $S_n^2(\pi\hat{\theta}_n)$ is at least as large as

$$\min \left\{ \frac{q_{1,n}}{q_n(q_{0,n} - 1)}, \frac{q_{0,n}}{q_n(q_{1,n} - 1)} \right\} \left(\sum_{k=1}^{q_n} \left(\hat{\theta}_{n,k} - \theta_0 - \beta \frac{q_{1,n}}{q_n} \right)^2 - o_P(q_{1,n}) - o_P(q_{0,n}) \right).$$

Calculations as in the proof of Lemma A.4 show that this is within $o_P(1)$ of

$$\min \left\{ \frac{q_{1,n}}{q_{0,n} - 1}, \frac{q_{0,n}}{q_{1,n} - 1} \right\} \left(\frac{1}{q_n} \sum_{k=1}^{q_n} \mathbb{E} Z_{n,k}^2 + \beta^2 \frac{q_{1,n} q_{0,n}^2 + q_{1,n}^2 q_{0,n}}{q_n^3} \right),$$

which is bounded away from zero for n large enough. The right-hand side in

$$\frac{S_n^2(\hat{\theta}_n)}{S_n^2(\pi\hat{\theta}_n)} \leq \frac{\max\{q_{1,n}/(q_{0,n} - 1), q_{0,n}/(q_{1,n} - 1)\}}{\min\{q_{1,n}/(q_{0,n} - 1), q_{0,n}/(q_{1,n} - 1)\}} \frac{q_n^{-1} \sum_{k=1}^{q_n} (\hat{\theta}_{n,k} - \theta_0 - \beta q_{1,n}/q_n)^2}{q_n^{-1} \sum_{k=1}^{q_n} (\hat{\theta}_{n,k} - \theta_0 - \beta q_{1,n}/q_n)^2 + o_P(1)}$$

therefore converges in probability unconditionally either to $((1 - \lambda)/\lambda)^2$ or its reciprocal depending on whether $\lambda \leq .5$ or not. Conclude that

$$T_n(\pi\hat{\theta}_n) = \left(\frac{1}{q_{1,n}} \sum_{k=1}^{q_{1,n}} \hat{\theta}_{n,\pi(k)} - \frac{1}{q_{0,n}} \sum_{k=1+q_{1,n}}^{q_n} \hat{\theta}_{n,\pi(k)} \right) \frac{S_n(\hat{\theta}_n)}{S_n(\pi\hat{\theta}_n)} \xrightarrow{P} 0 \quad (\text{B.3})$$

unconditionally. It follows that for $0 < \varepsilon < \beta$ we have that the first term on the right of

$$P(T_n \leq c_{n,\alpha}) \leq P(T_n \leq \beta - \varepsilon) + P(c_{n,\alpha} > \beta - \varepsilon)$$

converges to zero. By the properties of quantile functions (van der Vaart, 1998, Lemma 21.1(i), p. 304), the second term on the right is equal to $P(P_\pi(T_n(\pi\hat{\theta}_n)) > \beta - \varepsilon) > \alpha$. Because this converges to 0 by (B.3) and Lemma A.3, conclude that $P(T_n > c_{n,\alpha}) \rightarrow 1$.

Now consider the alternative $\beta > 0$ with q_1 and q_0 fixed. Let $\pi_0 = (1, 2, \dots, q)$. We have $\hat{\theta}_{n,k} \xrightarrow{P} \theta_0 + \beta 1\{k \leq q_1\}$ and therefore $T_n \xrightarrow{P} \beta$. For $0 < \varepsilon < \beta$, use this result in $P(T_n \leq c_{n,\alpha}) \leq P(T_n \leq \beta - \varepsilon) + P(c_{n,\alpha} > \beta - \varepsilon)$ to see that the first term on the right

converges to zero. By the properties of quantile functions, the second term can be written as

$$P\left(\sum_{\pi \in \Pi} 1 \left\{ \frac{\hat{S}_n(\hat{\theta}_n)}{\hat{S}_n(\pi \hat{\theta}_n)} \sum_{k=1}^q a_{\pi(k)} \hat{\theta}_{n,k} > \beta - \varepsilon \right\} > \alpha | \Pi \right).$$

For each π , $\sum_{k=1}^q a_{\pi^{-1}(k)} \hat{\theta}_{n,k}$ is within $o_P(1)$ of $\beta \sum_{k=1}^{q_1} a_{\pi^{-1}(k)}$. Note that $\sum_{k=1}^{q_1} a_k = 1$ and $\beta \sum_{k=1}^{q_1} a_{\pi^{-1}(k)} < \beta$ for every $\pi \neq \pi_0$ because $\sum_{k=1}^{q_1} a_{\pi^{-1}(k)}$ can at most be $q_1^{-1}(q_1 - 1) - q_0^{-1}$. Let $p_1(\pi) = q_1^{-1} \sum_{k=1}^{q_1} 1\{\pi^{-1}(k) \leq q_1\}$ and $p_0(\pi) = q_0^{-1} \sum_{k=1+q_1}^q 1\{\pi^{-1}(k) \leq q_1\}$. Then

$$\hat{S}_n^2(\pi \hat{\theta}_n) \xrightarrow{P} \frac{p_1(\pi)(1 - p_1(\pi))}{q_1 - 1} + \frac{p_0(\pi)(1 - p_0(\pi))}{q_0 - 1}$$

and, in particular, $\hat{S}_n(\hat{\theta}_n) \xrightarrow{P} 0$. For $\pi \neq \pi_0$, the right-hand side of the display is strictly positive. Hence, $P(\sum_{k=1}^q a_{\pi^{-1}(k)} \hat{\theta}_{n,k} \hat{S}_n(\hat{\theta}_n) / \hat{S}_n(\pi \hat{\theta}_n) > \beta - \varepsilon)$ converges to $1\{\pi = \pi_0\}$ for every $\pi \neq \pi_0$ and the union bound gives $P(\cup_{\pi \neq \pi_0} \{\sum_{k=1}^q a_{\pi^{-1}(k)} \hat{\theta}_{n,k} \hat{S}_n(\hat{\theta}_n) / \hat{S}_n(\pi \hat{\theta}_n) > \beta - \varepsilon\}) \rightarrow 0$. Conclude that

$$P\left(\sum_{\pi \in \Pi} 1 \left\{ \frac{\hat{S}_n(\hat{\theta}_n)}{\hat{S}_n(\pi \hat{\theta}_n)} \sum_{k=1}^q a_{\pi^{-1}(k)} \hat{\theta}_{n,k} > \beta - \varepsilon \right\} = 1 \right) \rightarrow 1$$

and therefore $P(c_{n,\alpha} > \beta - \varepsilon) \rightarrow 1\{1 > \alpha |\Pi|\}$, which proves the result for $|\Pi|^{-1} \leq \alpha$. If $|\Pi|^{-1} > \alpha$, then $1\{T_n > c_{n,\alpha}\} \equiv 0$ and therefore $P(T_n > c_{n,\alpha}) \rightarrow 0$ trivially. The result follows. \square

Proof of Corollary 3.4. This can be shown following along the same lines as the proofs of Lemma A.2 and Theorem 3.3 with two modifications: both statistics can now be standardized by σ_n and the appeal to Theorem A.5 has to be replaced by an appeal to Theorem A.6. Given the results in the proofs of Lemma A.2 and Theorem 3.3, it suffices to verify conditions (i)-(iii) of Theorem A.6 under the null. Recall $w_{n,k} = \sqrt{q_{1,n} q_{0,n}} (q_{1,n}^{-1} 1\{k \leq q_{1,n}\} - q_{0,n}^{-1} 1\{k > q_{1,n}\})$ and let $W_{n,k} = w_{n,\pi(k)} / \sqrt{q_n}$. Then $\bar{W}_n = 0$ and (i) follows immediately. Condition (ii) is satisfied because $W_{n,k}^2 = 1\{\pi(k) \leq q_{1,n}\} q_{0,n} / (q_{1,n} q_n) + 1\{\pi(k) > q_{1,n}\} q_{1,n} / (q_{0,n} q_n)$ and therefore the $W_{n,k}^2$ sum to one. Because $w_{n,\pi(1)}$ converges weakly to a Rademacher variable, (iii) holds and the desired result follows. The result under the alternative is identical to the proof of the second part of Theorem 3.6 below. \square

Proof of Theorem 3.6. Let $Z = (Z_1, \dots, Z_q)$ and note that the elements of Π are now nonrandom. For $\pi \in \Pi$ and $x = (x_1, \dots, x_q) \in \mathbb{R}^q$, consider the map $x \mapsto T(x) = q_1^{-1} \sum_{k=1}^{q_1} x_k - q_0^{-1} \sum_{k=1+q_1}^q x_k$. Because $T(\pi\theta) = T(\pi\theta - \theta_0 1_q)$ and the test decision is invariant to multiplication by a positive scalar, I can work with $\nu_n(\hat{\theta}_n - \theta_0 1_q)$ in place of $\hat{\theta}_n$.

Suppose $\beta = 0$. Compared to the randomized test with a test function as in (2.6) (replace $c_{n,\alpha}$ by $\bar{c}_{n,\alpha}$), we have $1\{T_n > \bar{c}_{n,\alpha}\} \leq 1\{\varphi_{n,\alpha}(T_n) \geq U\}$ and therefore $P(T_n > \bar{c}_{n,\alpha}) \leq P(\varphi_{n,\alpha}(T_n) \geq U)$. The idea is now to apply Theorem 3.1 of Canay et al. (2014) to $P(\varphi_{n,\alpha}(T_n) \geq U)$. Their Assumptions 3.1(i) and (ii) hold by assumption. Their Assumption 3.1(iii) holds if for any two $\pi, \pi' \in \Pi$ with $\pi \neq \pi'$ we either have $T(\pi x) = T(\pi' x)$ for all x or $P(T(\pi Z) = T(\pi' Z)) = 0$. Arguing similarly to the proof of Canay et al. (2014, Lemma S.5.1), if that were not true, there would be $\pi, \pi' \in \Pi$ with $\pi \neq \pi'$ such that $T(\pi x) \neq T(\pi' x)$ for some t and $P(T(\pi Z) = T(\pi' Z)) > 0$. For this choice of $\pi \neq \pi'$, let $a_k = q_1^{-1}1\{k \leq q_1\} - q_0^{-1}1\{k > q_1\}$ and $b_k = a_{\pi^{-1}(k)} - a_{\pi'^{-1}(k)}$. Because $T(\pi Z) = T(\pi' Z)$ if and only if $\sum_{k=1}^q b_k Z_k = 0$, conclude that $P(T(\pi Z) = T(\pi' Z)) > 0$ is equivalent to a discontinuity of the distribution of $\sum_{k=1}^q b_k Z_k$ at zero,

$$P\left(\sum_{k=1}^q b_k Z_k \leq 0\right) > P\left(\sum_{k=1}^q b_k Z_k < 0\right).$$

A sum of independent variables is continuously distributed if at least one of the summands has that property. Because $\pi \neq \pi'$, we have $b_k \neq 0$ for at least one k . But every Z_k is continuously distributed, which contradicts the preceding display. Conclude from continuity of $x \mapsto T(\pi x)$ for every given $\pi \in \Pi$ that

$$\limsup_{n \rightarrow \infty} P(T_n > \bar{c}_{n,\alpha}) \leq \lim_{n \rightarrow \infty} P(\varphi_{n,\alpha}(T_n) \geq U) = \alpha$$

by Theorem 3.1 of Canay et al. (2014).

Now consider the alternative $\beta > 0$. As before, we have $\hat{\theta}_{n,k} \xrightarrow{P} \theta_0 + \beta 1\{k \leq q_1\}$ and therefore $T_n \xrightarrow{P} \beta$. For $0 < \varepsilon < \beta$, use this result in $P(T_n \leq \bar{c}_{n,\alpha}) \leq P(T_n \leq \beta - \varepsilon) + P(\bar{c}_{n,\alpha} > \beta - \varepsilon)$ to see that the first term on the right converges to zero. By the properties of quantile functions, the second term can be written as

$$P\left(\sum_{\pi \in \Pi} 1\left\{\sum_{k=1}^q a_{\pi^{-1}(k)} \hat{\theta}_{n,k} > \beta - \varepsilon\right\} > \alpha |\Pi|\right).$$

For each π , $\sum_{k=1}^q a_{\pi^{-1}(k)} \hat{\theta}_{n,k}$ is within $o_P(1)$ of $\beta \sum_{k=1}^{q_1} a_{\pi^{-1}(k)}$. Note that $\sum_{k=1}^{q_1} a_k = 1$ and, in particular, $\beta \sum_{k=1}^{q_1} a_{\pi^{-1}(k)} < \beta$ for every $\pi \neq \pi_0$ because $\sum_{k=1}^{q_1} a_{\pi^{-1}(k)}$ can at most be $q_1^{-1}(q_1 - 1) - q_0^{-1}$. Hence, $P(1\{\sum_{k=1}^q a_{\pi^{-1}(k)} \hat{\theta}_{n,k} > \beta - \varepsilon\} = 1)$ converges to $1\{\pi = \pi_0\}$ for every $\pi \neq \pi_0$, possibly after decreasing ε , and the union bound gives $P(\cup_{\pi \neq \pi_0} \{\sum_{k=1}^q a_{\pi^{-1}(k)} \hat{\theta}_{n,k} > \beta - \varepsilon\}) \rightarrow 0$. Conclude that

$$P\left(\sum_{\pi \in \Pi} 1\left\{\sum_{k=1}^q a_{\pi^{-1}(k)} \hat{\theta}_{n,k} > \beta - \varepsilon\right\} = 1\right) \rightarrow 1$$

and therefore $P(\bar{c}_{n,\alpha} > \beta - \varepsilon) \rightarrow 1\{1 > \alpha|\Pi|\}$, which proves the result for $|\Pi|^{-1} \leq \alpha$. If $|\Pi|^{-1} > \alpha$, then $1\{T_n > \bar{c}_{n,\alpha}\} \equiv 0$ and therefore $P(T_n > \bar{c}_{n,\alpha}) \rightarrow 0$ trivially. The result follows. \square

REFERENCES

- Bell, R. M. and D. F. McCaffrey (2002). Bias reduction in standard errors for linear regression with multi-stage samples. *Survey Methodology* 28, 169–181.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* 119, 249–275.
- Bester, C. A., T. G. Conley, and C. B. Hansen (2011). Inference with dependent data using cluster covariance estimators. *Journal of Econometrics* 165, 137–151.
- Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics* 90, 414–427.
- Cameron, A. C. and D. L. Miller (2015). A practitioner’s guide to cluster robust inference. *Journal of Human Resources* 50, 317–372.
- Canay, I., J. P. Romano, and A. M. Shaikh (2014). Randomization tests under an approximate symmetry assumption. Technical Report No. 2014-13, Stanford University.
- Carter, A. V., K. T. Schnepel, and D. G. Steigerwald (2013). Asymptotic behavior of a t test robust to cluster heterogeneity. University of California, Santa Barbara working paper.
- Dal Bó, P. and G. R. Fréchette (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review* 101, 411–429.
- Dedecker, J., P. Doukhan, G. Lang, J. R. León, S. Louhichi, and C. Prieur (2007). *Weak Dependence: With Examples and Applications*. Springer.
- Donald, S. G. and K. Lang (2007). Inference with difference-in-differences and other panel data. *Review of Economics and Statistics* 89, 221–233.
- Durrett, R. (2010). *Probability: Theory and Examples*. Cambridge University Press.
- El Machkouri, M., D. Volný, and W. B. Wu (2013). A central limit theorem for stationary random fields. *Stochastic Processes and their Applications* 123, 1–14.
- Fisher, R. A. (1935). “The coefficient of racial likeness” and the future of craniometry. *Journal of the Royal Anthropological Institute of Great Britain and Ireland* 66, 57–63.
- Hansen, C. B. (2007). Asymptotic properties of a robust variance matrix estimator for panel data when t is large. *Journal of Econometrics* 141, 597–620.
- Hoeffding, W. (1952). The large-sample power of tests based on permutations of observations. *Annals of Mathematical Statistics* 23, 169–192.

- Horowitz, J. L. (1992). A smoothed maximum score estimator for the binary response model. *Econometrica* 60, 505–531.
- Ibragimov, R. and U. Müller (2010). t -statistic based correlation and heterogeneity robust inference. *Journal of Business & Economic Statistics* 28, 453–468.
- Ibragimov, R. and U. Müller (2016). Inference with few heterogenous clusters. *Review of Economics and Statistics* 98, 83–06.
- Imbens, G. W. and M. Kolesár (2016). Robust standard errors in small samples: Some practical advice. *Review of Economics and Statistics*, forthcoming.
- Janssen, A. (1997). Studentized permutation tests for non-i.i.d. hypotheses and the generalized Behrens-Fisher problem. *Statistics & Probability Letters* 36, 9–21.
- Janssen, A. (2005). Resampling Student’s t -type statistics. *Annals of the Institute of Statistical Mathematics* 57, 507–529.
- Kiefer, N. M. and T. J. Vogelsang (2002). Heteroskedasticity autocorrelation robust testing using bandwidth equal to sample size. *Econometric Theory* 18, 1350–1366.
- Kiefer, N. M. and T. J. Vogelsang (2005). A new asymptotic theory for heteroskedasticity-autocorrelation robust tests. *Econometric Theory* 21, 1130–1164.
- Kloek, T. (1981). OLS estimation in a model where a microvariable is explained by aggregates and contemporaneous disturbances are equicorrelated. *Econometrica* 49, 205–207.
- Lehmann, E. L. and C. M. Stein (1949). On the theory of some non-parametric hypotheses. *Annals of Mathematical Statistics* 20, 28–45.
- MacKinnon, J. G. and M. D. Webb (2016). Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics* 32, 233–254.
- MacKinnon, J. G. and M. D. Webb (2017). The subcluster wild bootstrap for few (treated) clusters. Queen’s Economics Department Working Paper No. 1364.
- Moulton, B. R. (1990). An illustration of a pitfall in estimating the effects of aggregate variables on micro units. *Review of Economics & Statistics* 72, 334–338.
- Neuhaus, G. (1993). Conditional rank test for the two-sample problem under random censorship. *Annals of Statistics* 21, 1760–1779.
- Romano, J. P. (1989). On the behavior of randomization tests without a group invariance assumption. Technical Report No. 334, Department of Statistics, Stanford University.
- Shao, X. and W. B. Wu (2007). Asymptotic spectral theory for nonlinear time series. *Annals of Statistics* 35, 1773–1801.
- van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press.
- Webb, M. D. (2014). Reworking wild bootstrap based inference for clustered errors. Queen’s Economics Department Working Paper No. 1315.

- Wu, W. B. (2005). Nonlinear system theory: Another look at dependence. *Proceedings of the National Academy of Sciences* 102, 14150–14154.
- Xiong, S. and G. Li (2008). Some results on the convergence of conditional distributions. *Statistics and Probability Letters* 78, 3249–3253.

DEPARTMENT OF ECONOMICS, UNIVERSITY OF MICHIGAN, 611 TAPPAN AVE, ANN ARBOR, MI 48109,
USA. TEL.: +1 (734) 764-2355. FAX: +1 (734) 764-2769
E-mail address: `hagem@umich.edu`
URL: `umich.edu/~hagem`