

// NO.25-070 | 12/2025

DISCUSSION PAPER

// JAN KEMPER AND DAVUD ROSTAM-AFSCHAR

Inference for Batched Adaptive Experiments

Inference for Batched Adaptive Experiments[☆]

Jan Kemper^a, Davud Rostam-Afschar^{b,*}

^aUniversity of Mannheim, ZEW

^bUniversity of Mannheim, IZA, GLO, NeSt

Abstract

The advantages of adaptive experiments have led to their rapid adoption in economics, other fields, as well as among practitioners. However, adaptive experiments pose challenges for causal inference. This note suggests a BOLS (batched ordinary least squares) test statistic for inference of treatment effects in adaptive experiments. The statistic provides a precision-equalizing aggregation of per-period treatment-control differences under heteroskedasticity. The combined test statistic is a normalized average of heteroskedastic per-period z -statistics and can be used to construct asymptotically valid confidence intervals. We provide simulation results comparing rejection rates in the typical case with few treatment periods and few (or many) observations per batch.

Keywords: Adaptive experiments, Heteroskedasticity, Causal inference, Randomized controlled trial

JEL: C12, C13, C9, D83

1. Introduction

Adaptive experiments have become increasingly common because they allow for *earning while learning*. Such designs have been applied, for example, by [Kasy and Sautmann \(2021\)](#), [Caria et al. \(2023\)](#), [Offer-Westort et al. \(2021\)](#), [Avivi et al. \(2021\)](#), [Tabord-Meehan \(2023\)](#), [Hoffmann et al. \(2023\)](#), [Gaul et al. \(2025\)](#). They combine exploration and exploitation by updating treatment probabilities based on accumulated evidence. However, the dependence of assignment on past outcomes breaks the usual assumptions of random sampling and independent treatment assignment, complicating statistical inference. This is particularly problematic when there is no clear difference between outcomes under different treatments. For example, usual confidence intervals and bootstrap methods may overreject null hypotheses. [Hadad et al. \(2021\)](#) use large number-of-trials asymptotics to construct generally valid

[☆]We thank Leif Döring, Taro Fujita, Michael Knaus, Kei Hirano, Jack Porter, David Preinerstorfer, Christoph Rothe, Kelly W. Zhang for insightful comments and discussions, as well as participants at various seminars and conferences for helpful comments. Davud Rostam-Afschar thanks the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for financial support through CRC *TRR 266 Accounting for Transparency* (Project-ID 403041268). We declare that we have no interests, financial or otherwise, that relate to the research described in this paper.

*Corresponding author

Email addresses: jan.kemper@zew.de (Jan Kemper), rostam-afschar@uni-mannheim.de (Davud Rostam-Afschar)

confidence intervals. [Zhang et al. \(2020\)](#) note that typically the number of trials is limited but treatment assignment is adapted after each batch of observations arrives. For this important case, they derive valid frequentist inference procedures using large batch size asymptotics under homoskedasticity. This note extends their argument to the more general and empirically relevant case of heteroskedastic outcomes, deriving the corresponding BOLS (batched OLS) test statistic and explores its asymptotic distribution. The heteroskedastic case is relevant because researchers usually design experiments such that not only the outcome means but also their variances differ by treatment arm. Often the outcome is binary (success/failure), which results in heteroskedasticity by construction. The results are in line with [Hirano and Porter \(2025\)](#) who show that any limit distribution generated by joint choices of adaptive assignment rules and statistics can be represented within a unified Gaussian limit experiment.

2. Treatment Effects in Adaptive Experiments

Let periods be indexed by $t = 1, \dots, T$. In period t , there are $N_{1,t}$ treated and $N_{0,t}$ control units, with $n_t = N_{1,t} + N_{0,t}$ and treatment share $\pi_t = N_{1,t}/n_t$, so $\pi_t(1 - \pi_t)$ is a measure of balance. Let the per-period difference in sample means be

$$\hat{\Delta}_t = \bar{Y}_{1,t} - \bar{Y}_{0,t}. \quad (1)$$

For each treatment arm $a \in \{0, 1\}$, outcomes of individuals $i = 1, 2, \dots, N_{a,t}$ satisfy $Y_{i,t}(a) = \mu_{a,t} + \varepsilon_{i,t}(a)$. Within period t , the treated and control sample means are independent with possibly different variances. Thus, $\varepsilon_{i,t}(1)$ is independent of $\varepsilon_{j,t}(0)$ for all i, j with $\varepsilon_{i,t}(a) \stackrel{i.i.d.}{\sim} (0, \sigma_{a,t}^2)$ and

$$\text{Var}(\bar{Y}_{1,t}) = \frac{\sigma_{1,t}^2}{N_{1,t}}, \quad \text{Var}(\bar{Y}_{0,t}) = \frac{\sigma_{0,t}^2}{N_{0,t}}.$$

Hence, the variance of the period difference is

$$\text{Var}(\hat{\Delta}_t) = v_t \equiv \frac{\sigma_{1,t}^2}{N_{1,t}} + \frac{\sigma_{0,t}^2}{N_{0,t}} = \frac{1}{n_t} \left(\frac{\sigma_{1,t}^2}{\pi_t} + \frac{\sigma_{0,t}^2}{1 - \pi_t} \right). \quad (2)$$

3. Inference

3.1. Scaling weights

In adaptive experiments, the selection probability π_t is random because it depends on the realized history. Thus, the variance of the OLS estimator across periods depends on the selection probability which may result in asymptotic non-normality. Intuitively, if outcomes under two treatments are hard to distinguish, either treatment might get assigned more observations in repeated samples, and consequently the selection probability does not concentrate. [Zhang et al. \(2020\)](#) show that the selection probability is fixed, when conditioning on the history up to a given batch, and that the batchwise OLS, scaled by the selection probability, is asymptotically normal. We construct an estimator across periods scaled by the inverse standard error, such that each period's standardized mean difference has the same influence. Let

$$w_t = \frac{1}{\sqrt{v_t}}, \quad S = \sum_{s=1}^T w_s.$$

Define the weighted average effect estimate

$$\hat{\Delta} = \sum_{t=1}^T \frac{w_t}{S} \hat{\Delta}_t. \quad (3)$$

Remark 1. (i) When assignment probabilities and batch sizes are fixed and variances are time-invariant, all periods are weighted equally in the combined statistic with $1/T$.

3.2. Variance of the weighted estimator

By construction, $w_t^2 v_t = 1$, hence the conditional variance given the realized weights is

$$\text{Var}(\hat{\Delta}) = \sum_{t=1}^T \left(\frac{w_t}{S} \right)^2 v_t = \frac{1}{S^2} \sum_{t=1}^T w_t^2 v_t = \frac{T}{S^2}. \quad (4)$$

Therefore,

$$\text{SE}(\hat{\Delta}) = \sqrt{\text{Var}(\hat{\Delta})} = \frac{\sqrt{T}}{S}. \quad (5)$$

3.3. Heteroskedastic Z-statistic

Theorem 1 (Asymptotic Normality of the Heteroskedastic BOLS Statistic). *For testing $H_0 : \Delta = c$, define the period z-scores and the combined statistic*

$$\begin{aligned} z_{t,\text{het}} &= \frac{\hat{\Delta}_t - c}{\sqrt{v_t}}, & Z_{\text{het}} &= \frac{\hat{\Delta} - c}{\text{SE}(\hat{\Delta})} = \frac{1}{\sqrt{T}} \sum_{t=1}^T z_{t,\text{het}}. \\ Z_{\text{het}} &\xrightarrow{d} N(0, 1). \end{aligned} \quad (6)$$

Asymptotic normality for fixed T and large batch sizes n_t follows [Zhang et al. \(2020\)](#), cf. Theorem 3), large- T asymptotics are covered by [Hadad et al. \(2021\)](#), cf. Theorem 4). [Hirano and Porter \(2025\)](#) provide a unifying Gaussian limit representation. The asymptotic distribution can be used to approximate their finite-sample distribution when constructing confidence intervals.

3.4. Feasible implementation

In practice, the arm- and period-specific variances are unknown. Let $\hat{\sigma}_{a,t}^2$ be consistent estimators for $a \in \{0, 1\}$ and the feasible test statistic \hat{Z}_{het} .

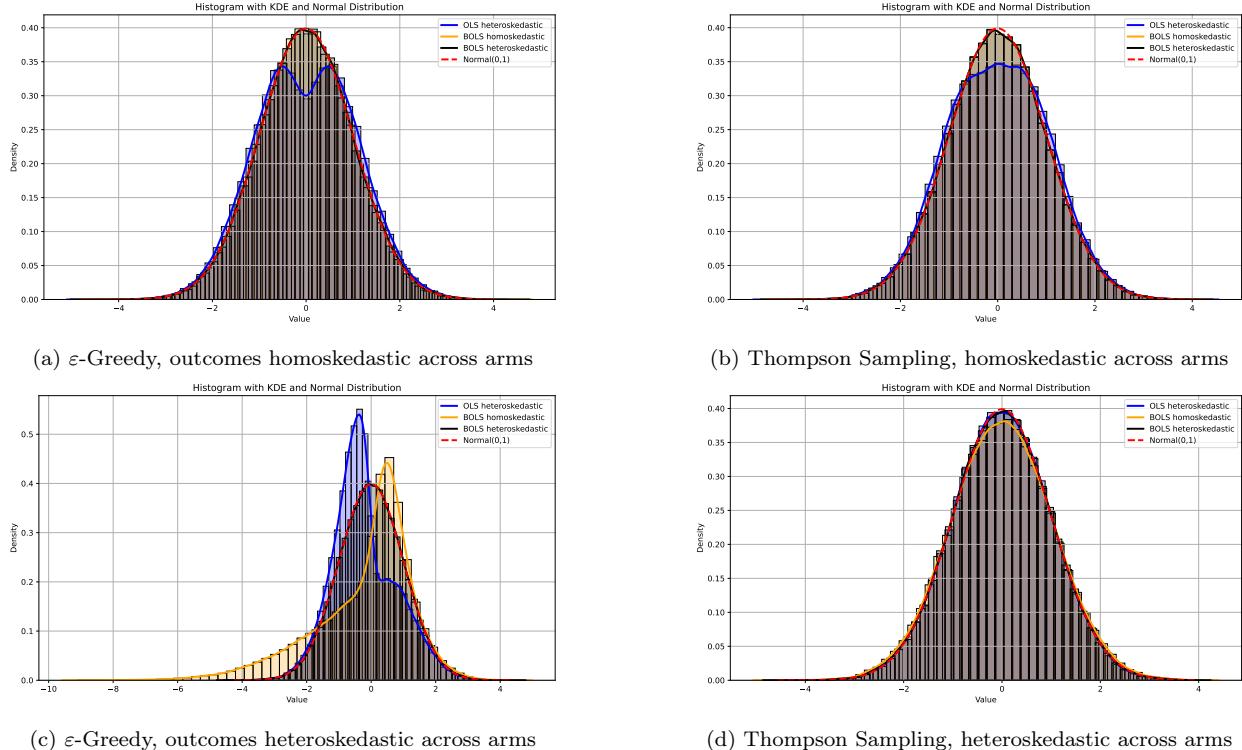
Corollary 1 (Asymptotically valid confidence interval). *Under the consistent variance estimation with feasible weighted estimator $\hat{\Delta}_F$, the two-sided $(1 - \alpha)$ confidence interval is*

$$\text{CI}_{1-\alpha}(\Delta) = \hat{\Delta}_F \pm \hat{Z}_{\text{het},1-\alpha/2} \text{SE}(\hat{\Delta}_F). \quad (7)$$

Remark 2 (Variance estimation). (i) For small samples, one may prefer finite-sample-adjusted within-period variance estimators (e.g., HC2/HC3). (ii) Under stationarity, arm specific variances pooled across batches $\hat{\sigma}_{a,t}^2 \xrightarrow{p} \sigma_a^2$ may be preferable (cf. corollary 4, [Zhang et al., 2020](#)). (iii) Under adaptivity, homoskedasticity ($\sigma_{1,t}^2 = \sigma_{0,t}^2 = \sigma^2$), and time-invariant batch size, the weights reduces to $\frac{w_t}{S} = \sqrt{N_{1,t}N_{0,t}} / \sum_{s=1}^T \sqrt{N_{1,s}N_{0,s}}$, so batches with more balanced sizes have larger weight. See Table A.1. (iv) it is straightforward to extend this to k treatment arms and contextual settings.

4. Monte Carlo Simulations

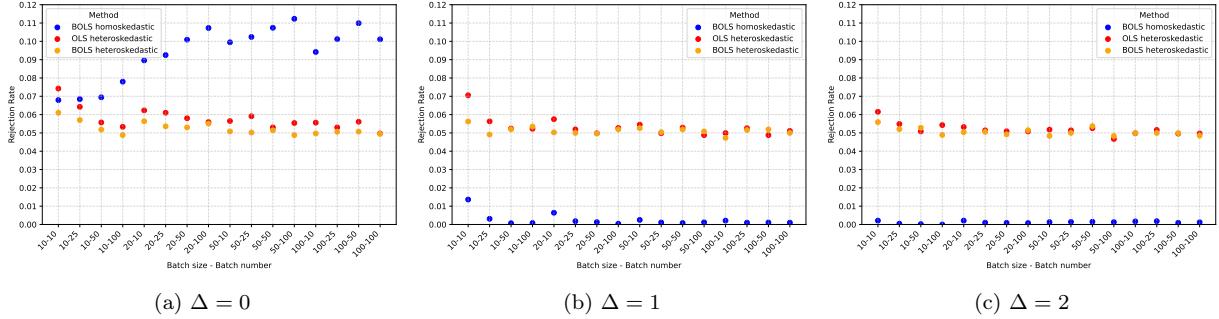
We compare three test statistics: the heteroskedasticity-robust OLS statistic, the BOLS statistic derived under homoskedasticity, and our heteroskedasticity-robust BOLS statistic. We report *rejection rates*, i.e., the proportion of simulated samples in which the null hypothesis is rejected. Under $H_0 : \Delta = 0$, this quantity measures the empirical size of the test (nominal 5%). Data are generated using two common adaptive sampling algorithms, ε -Greedy and Bernoulli Thompson Sampling in two-arm settings.



Notes: The figure shows the results of Monte Carlo simulations with 100,000 repetitions. Panels a) and c) show the distribution of the heteroskedasticity-robust OLS test statistic, the homoskedastic BOLS test statistic, and the heteroskedasticity-robust BOLS test statistic for data generated from an ε -Greedy experiment. The batch size is 500, the number of batches 25, the experiment consists of two arms, each with an expected value of 1. In panel a) the standard deviation is 1 for arm one and 1 for arm two (homoskedasticity). In panel c) the standard deviation for arm one is 4 and for arm two is 1, everything else remains equal. The red dotted line indicates the density of the standard normal distribution. Panel b) shows data generated from a Bernoulli Thompson algorithm. The batch size is 500, the number of batches 25, the experiment consists of two arms with an expected value of 0.5 and 0.5 (homoskedasticity). In panel d) the experiment consists of two arms with an expected value of 0.7 and 0.4 (heteroskedasticity).

Figure 1: Simulation I: Non-normality of OLS and homoskedastic BOLS

Simulation I: Non-normality of OLS and homoskedastic BOLS. Figure 1 shows the empirical distributions of the test statistics with 25 batches each with a size of 500 observations. For ε -Greedy (left panels), both arms have Gaussian outcomes (think of log incomes) with mean 1, with variances $(1^2, 1^2)$ in panel a) and $(4^2, 1^2)$ in panel c). The exploration rate is $\varepsilon = 0.2$. Each design is repeated 100,000 times. Consistent with [Zhang et al. \(2020\)](#), panel a) shows that OLS under zero-margin is non-normally distributed and homoskedastic BOLS recovers



Notes: This figure shows Monte Carlo simulation results for the ε -Greedy algorithm. Combinations of batch size and number of batches increase along the horizontal axis. Rejection rates are denoted on the vertical axis. The parameter Δ indicates the difference between the true expected values of both arms. Each circle shows the average rejection rate for the given test statistic which is indicated by color. For each batch size / number of batches combinations 10,000 repetitions were executed. In panel (a) the expected value μ_1 for arm 1 is 1 and μ_2 for arm 2 is also 1. The standard deviation in all panels for arm 1 is $\sigma_1 = 2$ and for arm 2 is $\sigma_2 = 1$. In panel (b) the expected value is $\mu_1 = 2$ and $\mu_2 = 1$. For panel (c) it is $\mu_1 = 3$ and $\mu_2 = 1$. In the batch size 10-10 scenario, on average, no test statistic could be calculated for one percent of the draws because one of the arms was never played.

Figure 2: ε -Greedy

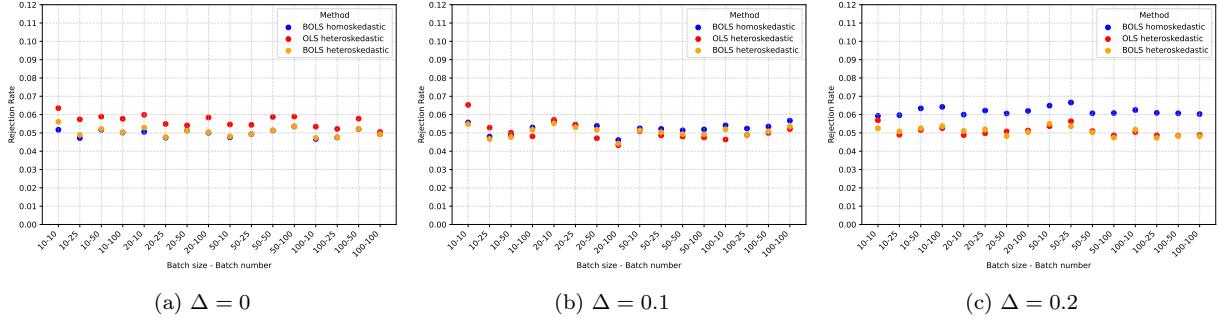
normality. But in the same setting under heteroskedasticity (panel c), both OLS and the homoskedastic BOLS statistic deviate markedly from normality in the zero-margin case. The homoskedastic BOLS statistic severely overrejects (17% instead of 5%). OLS yields approximately correct rejection rates but exhibits non-normal behavior. In contrast, our heteroskedasticity-robust BOLS statistic closely matches the standard normal distribution and delivers correct 5% rejection rates in both cases.

For Thompson Sampling (right panels), we set Bernoulli success probabilities to $(p_1, p_2) = (0.5, 0.5)$ in panel b) and to $(p_1, p_2) = (0.7, 0.4)$ in panel d). Panel b) confirms that OLS is non-normally distributed and the homoskedastic BOLS test statistics fixes the problem. But under mechanical heteroskedasticity in the non-zero margin case (panel d) the homoskedastic BOLS statistic overrejects ($\approx 6\%$) because it ignores heteroskedasticity. As the success probabilities differ substantially, both OLS and our heteroskedasticity-robust BOLS statistic closely match the standard normal distribution and deliver correct 5% rejection rates.

Simulation II: Rejection Rates at Small and Large Margins in Small and Large Samples. To study behavior in smaller samples, we vary the number of batches (10–100) and batch sizes (10–100). Each configuration is repeated 10,000 times. Figures 2 and 3 report rejection rates of a 5% significance level test of $H_0 : \Delta = 0$.

For ε -Greedy (Figure 2), the homoskedastic BOLS statistic is unreliable in all settings: it overrejects sharply at $\Delta = 0$ and underrejects for positive margins. The heteroskedasticity-robust OLS statistic performs reasonably well and improves as the margin grows. Across all designs, our heteroskedasticity-robust BOLS statistic maintains rejection rates close to 5%.

For Thompson Sampling (Figure 3), the zero-margin case exhibits no heteroskedasticity, so the heteroskedastic and the homoskedastic BOLS statistic perform well with rejection rates near 5%. The heteroskedasticity-robust OLS overrejects somewhat. When margins increase, inducing heteroskedasticity, the homoskedastic BOLS statistic begins to overreject, while both robust statistics remain close to nominal size.



Notes: This figure shows Monte Carlo simulation results for the Bernoulli Thompson Sampling algorithm. Combinations of batch size and number of batches increase along the horizontal axis. Rejection rates are denoted on the vertical axis. The parameter Δ indicates the difference between the true expected values of both arms. Each circle shows the average rejection rate for the given test statistic which is indicated by color. For each batch size / number of batches combinations 10,000 repetitions were executed. In panel (a) the success probability for arm 1 is $p_1 = 0.5$ and arm 2 is also $p_2 = 0.5$. In panel (b) it is $p_1 = 0.6$ and $p_2 = 0.5$. For panel (c) it is $p_1 = 0.7$ and $p_2 = 0.5$.

Figure 3: Bernoulli Thompson Sampling

5. Conclusion

Adaptive experiments have made selection-weighted inference increasingly relevant in sequential settings where treatment assignment depends on past outcomes. The BOLS selection-weighted statistic provides a simple, asymptotically valid procedure for inference under heteroskedasticity, extending previous results derived for the homoskedastic case.

References

AVIVI, H., P. KLINE, E. ROSE, AND C. WALTERS (2021): “Adaptive Correspondence Experiments,” *AEA Papers and Proceedings*, 111, 43–48.

CARIA, A. S. ET AL. (2023): “An Adaptive Targeted Field Experiment: Job Search Assistance for Refugees in Jordan,” *Journal of the European Economic Association*, 22(2), 781–836.

GAUL, J. J., F. KEUSCH, D. ROSTAM-AFSCHAR, AND T. SIMON (2025): “Invitation Messages for Business Surveys: A Multi-Armed Bandit Experiment,” *Survey Research Methods*, forthcoming.

HADAD, V., D. A. HIRSHBERG, R. ZHAN, S. WAGER, AND S. ATHEY (2021): “Confidence Intervals for Policy Evaluation in Adaptive Experiments,” *Proceedings of the National Academy of Sciences*, 118(15), e2014602118.

HIRANO, K., AND J. R. PORTER (2025): “Asymptotic Representations for Sequential Decisions, Adaptive Experiments, and Batched Bandits,” arXiv 2302.03117, econ.EM.

HOFFMANN, M., B. PICARD, C. MARIE, AND G. BIED (2023): “An Adaptive Experiment to Boost Online Skill Signaling and Visibility,” Draft, Accessed: 2024-04-10.

KASY, M., AND A. SAUTMANN (2021): “Adaptive Treatment Assignment in Experiments for Policy Choice,” *Econometrica*, 89(1), 113–132.

OFFER-WESTORT, M., A. COPPOCK, AND D. P. GREEN (2021): “Adaptive Experimental Design: Prospects and Applications in Political Science,” *American Journal of Political Science*, 65(4), 826–844.

TABORD-MEEHAN, M. (2023): “Stratification Trees for Adaptive Randomisation in Randomised Controlled Trials,” *The Review of Economic Studies*, 90(5), 2646–2673.

ZHANG, K., L. JANSON, AND S. MURPHY (2020): “Inference for Batched Bandits,” *Advances in Neural Information Processing Systems*, 33, 9818–9829.

Appendix A. Weight structure under homoskedasticity and heteroskedasticity

Table A.1: Weight structure under homoskedasticity and heteroskedasticity

	Non-adaptive	Adaptive
<i>Homoskedastic case: $\sigma_{1,t}^2 = \sigma_{0,t}^2 = \sigma^2$</i>		
Treatment share π_t	$\pi_t = \pi$ (fixed)	π_t depends on history
Variance v_t	$v_t = \frac{\sigma^2}{n_t \pi(1 - \pi)}$	$v_t = \frac{\sigma^2}{n_t \pi_t(1 - \pi_t)}$
Weight $\frac{w_t}{S}$	$\frac{\sqrt{n_t}}{\sum_{s=1}^T \sqrt{n_s}}$	$\frac{\sqrt{n_t \pi_t(1 - \pi_t)}}{\sum_{s=1}^T \sqrt{n_s \pi_s(1 - \pi_s)}}$
<i>Heteroskedastic case: $\sigma_{1,t}^2 \neq \sigma_{0,t}^2$</i>		
Treatment share π_t	$\pi_t = \pi$ (fixed)	π_t depends on history
Variance v_t	$v_t = \frac{1}{n_t} \left(\frac{\sigma_{1,t}^2}{\pi} + \frac{\sigma_{0,t}^2}{1 - \pi} \right)$	$v_t = \frac{1}{n_t} \left(\frac{\sigma_{1,t}^2}{\pi_t} + \frac{\sigma_{0,t}^2}{1 - \pi_t} \right)$
Weight $\frac{w_t}{S}$	$\frac{\sqrt{n_t} \left(\frac{\sigma_{1,t}^2}{\pi} + \frac{\sigma_{0,t}^2}{1 - \pi} \right)^{-1/2}}{\sum_{s=1}^T \sqrt{n_s} \left(\frac{\sigma_{1,s}^2}{\pi} + \frac{\sigma_{0,s}^2}{1 - \pi} \right)^{-1/2}}$	$\frac{\sqrt{n_t} \left(\frac{\sigma_{1,t}^2}{\pi_t} + \frac{\sigma_{0,t}^2}{1 - \pi_t} \right)^{-1/2}}{\sum_{s=1}^T \sqrt{n_s} \left(\frac{\sigma_{1,s}^2}{\pi_s} + \frac{\sigma_{0,s}^2}{1 - \pi_s} \right)^{-1/2}}$

Notes: $n_t = N_{1,t} + N_{0,t}$ is total batch size, and $\pi_t = N_{1,t}/n_t$ is the treatment share. Under homoskedasticity, weights depend on both sample size and balance $\pi_t(1 - \pi_t)$. Under heteroskedasticity, weights additionally adjust for treatment-specific outcome variances $\sigma_{a,t}^2$.



Download ZEW Discussion Papers:

<https://www.zew.de/en/publications/zew-discussion-papers>

or see:

<https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html>

<https://ideas.repec.org/s/zbw/zewdip.html>



IMPRINT

**ZEW – Leibniz-Zentrum für Europäische
Wirtschaftsforschung GmbH Mannheim**

ZEW – Leibniz Centre for European
Economic Research

L 7,1 · 68161 Mannheim · Germany

Phone +49 621 1235-01

info@zew.de · zew.de

Discussion Papers are intended to make results of ZEW
research promptly available to other economists in order
to encourage discussion and suggestions for revisions.
The authors are solely responsible for the contents which
do not necessarily represent the opinion of the ZEW.