DEPARTMENT OF STATISTICS
AHMADU BELLO UNIVERSITY, ZARIA
2nd International Conference on Statistical Sciences and Data Analytics (ICSSDA 2026) Proceedings Paper

Semiparametric Estimation of Mann-Whitney-Type Causal Effects Under Non-Normal Economic Outcomes: Evidence from Nigeria's Fuel Subsidy Reform

Safiya Sada Sani[1,2,*], Mohammed Usman[2], Ibrahim Adamu[2], Jibril Yahaya Kajuru[2], Ibrahim Abubakar Sadiq[2], Abdulhamid Ado Osi[3]

[1] *Department of Agronomy, Ahmadu Bello University (ABU), Zaria, Kaduna, Nigeria.*

[2] *Department of Statistics, Ahmadu Bello University (ABU), Zaria, Kaduna, Nigeria.*

[3] *Department of Statistics, Aliko Dangote University of Science &Tech. Wudil, Kano, Nigeria.*

[*] Correspondence: sssani.sabo0@gmail.com.

Abstract

Economic policy reforms often generate outcome distributions that violate the assumptions underlying conventional mean-based causal estimands, necessitating more robust semiparametric approaches. This article evaluates the welfare consequences of Nigeria's 2023 fuel subsidy removal using the semiparametric outcome-regression estimator for the Mann-Whitney-type causal effect developed by Sani *et al*. (2026). Household welfare is measured by per-capita consumption expenditure, observed in pre- and post-reform waves of the Nigeria General Household Survey (GHS). Reform exposure is proxied by household fuel dependence, with households reporting positive fuel expenditure classified as highly exposed and households reporting no fuel expenditure classified as less exposed. For comparison, regression-adjustment models are estimated for both the raw and log-transformed outcomes. To capture distributional effects under skewness and heterogeneity, the Mann-Whitney-type causal effect is estimated as the probability that a randomly selected highly exposed household has a higher welfare outcome than a comparable less exposed household after adjusting for observed covariates. Analyses are conducted on a 20% random subsample of the pooled GHS data to facilitate semiparametric estimation and bootstrap inference. The semiparametric analysis reveals substantial distributional differences between exposure groups, yielding an estimated Mann-Whitney-type causal effect of 0.687 (95% CI: 0.673 - 0.701), whereas mean-based conclusions vary depending on outcome scaling. Subgroup analyses indicate modest heterogeneity, with larger estimated effects among urban households and households located in southern regions. These findings demonstrate that semiparametric Mann-Whitney-type causal effects estimators provide a more informative framework for evaluating policy impacts in the presence of heterogeneous and non-normal outcomes than conventional mean-based methods.

Keywords: Fuel subsidy removal; Semiparametric causal inference; Mann-Whitney-type causal effect; Distributional treatment effects; Welfare analysis.

1. Introduction

Evaluating the welfare consequences of major economic reforms can be challenging when outcome distributions are highly skewed, heavy-tailed, or otherwise non-Gaussian. In such settings, conventional mean-based estimands may be sensitive to extreme observations and may fail to capture important distributional differences across affected populations. These limitations have motivated growing interest in rank-based causal estimands, such as the Mann-Whitney-type causal effect, which characterises treatment effects through pairwise outcome comparisons rather than moments of the outcome distribution.

Nigeria's 2023 fuel subsidy removal provides an important context for examining this issue. Recent research highlights that the removal of the fuel subsidy in Nigeria has broad economic and welfare implications. Akpan and Uford (2024), using an exploratory qualitative approach, show that subsidy removal leads to higher fuel pump prices, which then raise the prices of goods and services across the economy because fuel is a key input in production and distribution. Their findings point to rising household living costs driven by economy-wide price transmission. Complementing this perspective, Olatunbosun *et al*. (2025) argue through discourse analysis that although subsidy removal may strengthen fiscal sustainability and improve market efficiency, it can also heighten inflation, erode purchasing power, deepen poverty, and generate welfare losses, especially for vulnerable households and informal-sector workers. They emphasise that the effects are uneven across socioeconomic groups, underscoring the need to assess not only average outcomes but also the distributional consequences of subsidy reforms.

The Mann-Whitney framework introduced by Mann and Whitney (1947) has been extended to increasingly complex settings. Hart (2001) emphasised that it captures differences in entire distributions rather than merely differences in central tendency. Bilker and Wang (1996) extended the approach to semiparametric settings involving truncated data, demonstrating improved performance relative to existing nonparametric alternatives. Later, Zhang *et al*. (2019) developed outcome-regression, inverse-probability-weighted, and doubly robust estimators for Mann-Whitney-type causal effects in observational studies. Building on this work, Sani *et al*. (2026) proposed a semiparametric cumulative probability model (CPM) estimator that combines covariate adjustment with the desirable robustness properties of rank-based inference.

This study applies the CPM estimator of Sani *et al*. (2026) to evaluate the welfare consequences of Nigeria's fuel subsidy removal using data from the Nigeria General Household Survey. Household welfare is measured by per-capita consumption expenditure, and fuel dependence is used to classify households according to their level of exposure to the reform. The analysis compares conventional regression-adjustment estimates with the Mann-Whitney-type causal effect to assess whether conclusions differ between mean-based and distributional perspectives.

The study contributes to the literature by applying a semiparametric rank-based causal estimator to a major economic reform and demonstrating its usefulness for evaluating policy impacts under highly non-normal welfare outcomes.

2. Materials and Methods

2.1 Sample Size Determination

The empirical analysis used pooled observations from Wave 4 (2018–2019) and Wave 5 (2023–2024) of the Nigeria General Household Survey (GHS). The predetermined nationally representative sample size was relied upon. To reduce the computational burden associated with semiparametric estimation and bootstrap inference, a 20% simple random sample of the pooled dataset was selected using a fixed random seed, resulting in a final analytical sample of 1,924 households.

2.2. Study Design

This study used a cross-sectional design to evaluate the welfare effects of Nigeria's fuel subsidy removal using a causal inference framework based on the potential outcomes approach. Household per-capita consumption expenditure was used as the outcome variable and serves as a proxy for household welfare.

2.2.1 Causal Framework

Let Y denote household per-capita consumption expenditure, defined as

$$Y = \frac{Total\ Household\ Consumption\ Expenditure}{Household\ Size} \tag{1}$$

where total consumption expenditure is the sum of food, non-food, energy, and fuel expenditures. Higher values indicate greater household consumption expenditure per person and therefore higher welfare costs. Consumption expenditure is widely used in welfare analysis because it reflects a household's ability to maintain living standards in the face of economic shocks. For robustness, both the raw measure and its logarithmic transformation $\log(Y+1)$ were analyzed.

Let $A \in \{0,1\}$ indicate household exposure to the fuel subsidy removal. This exposure was proxied by household fuel dependence, measured as the share of total consumption expenditure devoted to fuel,

$$A = \frac{Fuel\ Expenditure}{Total\ Expenditure}. \tag{2}$$

Households having a non-zero value on the fuel dependence scale were classified as highly exposed (A=1), reflecting direct vulnerability to fuel-price increases. Households with zero fuel dependence value, on the other hand, were classified as less exposed (A=0). This binary classification captures heterogeneity in household dependence on fuel and provides a transparent measure of reform exposure. Let $T \in \{0,1\}$ indicate the pre-reform (2018–2019) and post-reform (2023–2024) survey periods, respectively.

Finally, let $\boldsymbol{X}$ denote the covariate vector of observed household characteristics. For this study, these are household size, sex of household head (male or female), age of household head, sector of residence (urban or rural), and geographic zone of residence (North-Central, North-East, North-West, South-East, South-South and South-West).

The causal analysis seeks to compare welfare outcomes between households with different levels of exposure to the fuel subsidy removal while accounting for differences in observed covariates. Identification of causal effect relies on the assumptions of conditional unconfoundedness $\{Y(1), Y(0)\} \perp \mathbf{X},$ and overlap, $0 < P(A = 1 \mid \mathbf{X}) < 1.$ Under these assumptions, causal effect can be identified from observed data.

2.3. Method of Data Collection

This study is a secondary analysis of two waves of the General Health Survey (GHS) conducted by the National Bureau of Statistics (NBS) in collaboration with the World Bank. The GHS employed a stratified multistage sampling and interviewer-administered questionnaires to collect nationally representative household data, which were downloaded from the official microdata portal for analysis.

2.4. Method of Data Analysis

Data were coded, cleaned, and derived variables were obtained. All analyses

$$Y = \frac{Total\, household\, Consumption\, Expenditure}{Household\, Size}, \tag{3}$$

were conducted using RStudio (version 2026.05.0+218, "Golden Wattle" Release, RStudio Team, 2026) on a Windows 11 Pro platform (Microsoft Corporation, 2021). The statistical techniques employed in this study were as follows:

2.4.1. Descriptive Statistics and Baseline Analysis

Frequencies and percentages were used to summarise categorical variables, while means and standard deviations were used for continuous variables. Baseline differences between the two exposure groups were assessed using Chi-squared tests for categorical variables and the Wilcoxon rank-sum test for continuous variables. All statistical tests were two-tailed, and a p-value of less than 0.05 was considered statistically significant.

2.4.2 Regression-Adjustment

As a benchmark, a conventional regression-adjustment approach was employed to estimate the relationship between exposure to the fuel subsidy removal and household consumption expenditure using the covariates chosen for the study. The regression model is specified as:

$$Y_i = \beta_0 + \beta_1 A_i + \beta_2 T_i + \beta_3 (A_i \times T_i) + \mathbf{X}_i^T \gamma + \varepsilon_i \qquad {}_{i=1,\ldots,n} \tag{4}$$

where Y_i is household per-capita consumption expenditure, A_i represents exposure status, T_i denotes the survey period, $\boldsymbol{X}_i$ is a vector of control variables (the covariates), and ε_i is a stochastic error term.

The interaction term ($A_i \times T_i$) captures whether welfare outcomes evolved differently over time for highly exposed households relative to less exposed households, after adjustment for observed covariates. The model yields an estimate of the average treatment effect on welfare costs and provides a useful comparison with the Mann–Whitney-type estimator. The same model was fit for the log-transformed outcome.

2.4.2. Mann-Whitney-Type Causal Effect

The primary estimand is the Mann–Whitney-type causal effect,

$$\theta = P\{Y(1) > Y(0)\} = E\left[h(Y(1), Y(0)\right], \tag{5}$$

a contrast between the two potential outcomes. This parameter can be viewed as a causal analogue of the probability underlying the classical Mann–Whitney framework (Mann & Whitney, 1947), extended to potential outcomes and confounding adjustment. It has a natural probabilistic interpretation. It represents the probability that a randomly selected highly exposed household has a higher welfare outcome than a randomly selected less exposed household with comparable covariates. Here, $h(\cdot)$ is represented by an indicator function defined as

$$h(y_1, y_0) = I(y_1 > y_0) + \frac{1}{2} I(y_1 = y_0), \qquad I(\cdot) = \{^{1,\ \ \textit{if condition holds}}_{0,\ \ \textit{otherwise}} \tag{6}$$

For continuous outcomes, the second part of the RHS of the indicator function vanishes. A value of $\theta = 0.5$ indicates no distributional difference between exposure groups. Values greater than 0.5 indicate that welfare outcomes tend to be higher among highly exposed households, whereas values less than 0.5 indicate that welfare outcomes tend to be higher among less exposed households.

2.4.2.1 Semiparametric Estimation

Following Sani et al. (2026), a semiparametric model-based estimator of θ uses the conditional outcome distribution $F_{Y|A,T,X}(y \mid A, T, \mathbf{X})$ modelled using a cumulative probability model (CPM) defined as

$$G\left[P(Y \le y \mid A, T, \mathbf{X}\right] = \boldsymbol{\alpha}_i - \beta^T \mathbf{X} - \tau_1 A - \tau_2 T - \tau_3 (A \times T) \qquad {}_{i=1,\ldots,n} \tag{7}$$

where $G[\cdot]$ is a link function and $\boldsymbol{\alpha}_i$ a semiparametrically estimated intercept function. The fitted CPM is then used to estimate marginal potential outcome distributions $\widehat{F}_{\mathrm{a}}(y), a = (0,1)$. These are used to evaluate the Mann–Whitney-type causal effect:

$$\hat{\theta} = \sum_{i=1}^{n} \sum_{j=1}^{n} h(y_1, y_0) \times \left[\hat{F}_1(y_i) - \hat{F}_1(y_{i-1})\right] \times \left[\hat{F}_0(y_i) - \hat{F}_0(y_{j-1})\right] \tag{8}$$

$\widehat{F}_a(y_i) - \widehat{F}_a(y_{i-1})$ are finite difference approximations to the marginal densities, and i and j denote independent observations randomly drawn from the treated and control groups, respectively. Due to its rank-based nature, it is invariant to monotonic transformations of the outcome variable.

2.4.3. Subgroup Analyses

To investigate treatment-effect heterogeneity, subgroup-specific Mann–Whitney-type causal effects were estimated from the fitted CPM by restricting predictions to the observed covariate distribution within each subgroup, and computing standardised marginal potential-outcome distributions under high and low exposure. Subgroup analyses were conducted for:

i. Urban and rural households;

ii. Male- and female-headed households; and

iii. Northern and Southern regions of Nigeria.

These analyses provide insight into whether the welfare effects of the fuel subsidy removal vary across different population groups.

2.4.3. Statistical Uncertainty Assessment via Nonparametric Bootstrap

Statistical uncertainty was quantified using nonparametric bootstrap procedures, from which standard errors and confidence intervals were constructed.

3. Results and Discussions

Table 1 presents summary statistics and tests of differences between high- and low-exposure households in terms of household welfare expenditure and the covariates. Note that p-values are from Wilcoxon rank-sum tests (continuous variables) and Pearson's chi-squared tests (categorical variables). Figure 1 displays the distribution of household welfare expenditure across exposure groups. Figure 2 presents the distribution of the log-transformed welfare outcome.

Table 1. Household Characteristics by Fuel Exposure Status

Characteristic	**Overall** N = 1,924[1]	**Low fuel exposure** N = 1,352[1]	**High fuel exposure** N = 572[1]	**p-value**[2]
Per Capita Expenditure (₦)	6,477 (10,264)	4,998 (7,813)	9,971 (13,891)	<0.001
Household size	5.9 (3.6)	5.8 (3.6)	6.2 (3.4)	0.029
Age of household head	51 (15)	52 (16)	50 (14)	0.031
Sex of household head				<0.001
Male	1,495 (78%)	996 (74%)	499 (87%)	
Female	429 (22%)	356 (26%)	73 (13%)	
Sector				0.007
Urban	628 (33%)	416 (31%)	212 (37%)	
Rural	1,296 (67%)	936 (69%)	360 (63%)	
Zone				<0.001
North-Central	336 (17%)	240 (18%)	96 (17%)	
North-East	325 (17%)	267 (20%)	58 (10%)	
North-West	320 (17%)	237 (18%)	83 (15%)	
South-East	321 (17%)	207 (15%)	114 (20%)	
South-South	315 (16%)	180 (13%)	135 (24%)	
South-West	307 (16%)	221 (16%)	86 (15%)	

[1]Mean (SD); n (%); [2]Wilcoxon rank sum test; Pearson's Chi-squared test

Highly exposed households exhibited substantially higher average per-capita consumption expenditure (9,971; SD = 13,891) than low-exposure households (4,998; SD = 7,813), with the difference being statistically significant ($p < 0.001$). Significant differences were also observed across household size, age of household head, sex of household head, sector of residence, and geographic zone, indicating substantial heterogeneity between exposure groups.

Figure 1. Histogram of household per-capita consumption expenditure by exposure group.

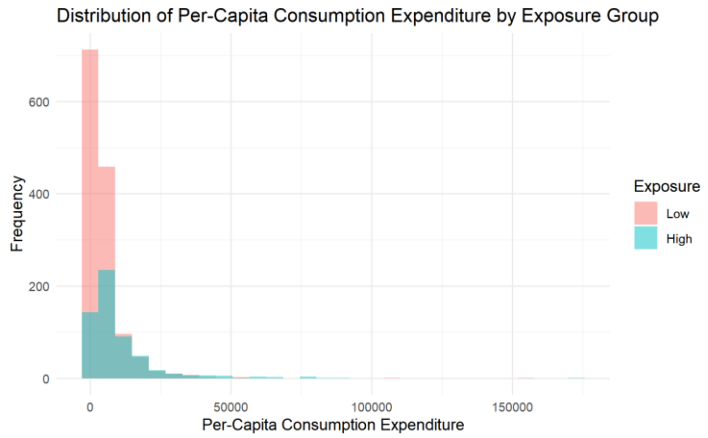

The distributions in both low and high exposure groups are characterised by strong positive skewness, several extreme observations, and deviations from normality. Highly exposed households generally occupy higher positions in the expenditure distribution, although considerable overlap exists between the groups. These distributional features suggest that mean-based measures alone may not fully characterise welfare differences associated with fuel subsidy removal.

Figure 2. Histogram of log-transformed household per-capita consumption expenditure.

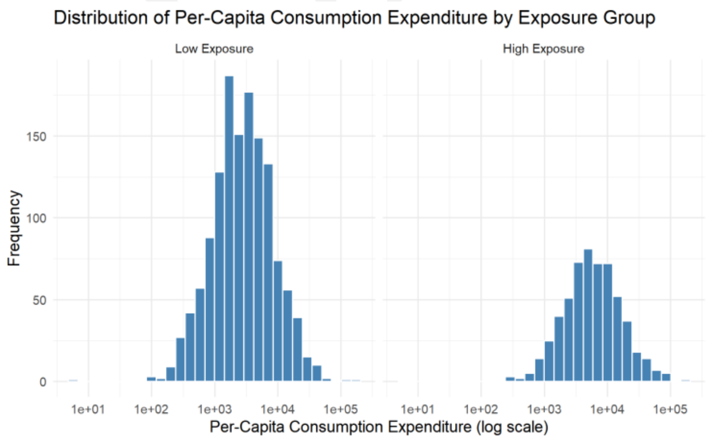

Relative to the raw distribution shown in Figure 1, the logarithmic transformation substantially reduces the influence of extreme observations and produces a more symmetric distribution. Although some departures from normality remain, the transformation attenuates the strong right-skewness observed in the original data and provides a useful benchmark for assessing the sensitivity of regression-based results to outcome scaling. The impact of this transformation becomes evident in the regression-adjustment results presented in Section 3.2, where inferences regarding the exposure-by-period interaction differ between the raw and log-transformed specifications, illustrating the sensitivity of mean-based methods to outcome scaling.

3.2 Mean-Based versus Distributional Effects

Table 2 presents the estimates from both the regression adjustment and cumulative probability models.

Regression-adjustment models indicated higher average expenditure among highly exposed households following the subsidy removal. However, conclusions regarding the exposure -by-period interaction differed between the raw log-transformed outcomes, indicating sensitivity to mean-based inference, outcome scaling and extreme observations. The CPM produced similar conclusions regarding the main effects of exposure and post-reform period but indicated a small and statistically insignificant exposure-by-reform interaction.

The estimated Mann-Whitney-type causal effect was $\hat{\theta} = 0.687$ (95% CI: 0.673-0.701). Because $\theta = 0.5$ indicates no distributional difference, this estimate provides strong evidence of substantial distributional separation between highly exposed and less exposed households. Specifically, the estimate implies a 68.7% probability that a randomly selected highly exposed household has a higher per-capita consumption expenditure than a randomly selected less exposed household with comparable characteristics.

Table 2. Comparison of Mean-Based and Distributional Estimates

Measure	Estimate	95% CI	p-value
Regression Adjustment (Raw Expenditure)			
Exposure x Post-Reform Interaction	4,158.11	-	<0.001
Post-Reform Average Treatment Effect (₦)	6,668.28	4,136.94 - 9,368.40	<0.001
Regression Adjustment (Log Expenditure)			
Exposure x Post-Reform Interaction	-0.037	-	0.648
Post-Reform Average Treatment Effect (log points)	0.677	0.541 - 0.817	<0.001
Cumulative Probability Model (CPM)			
Exposure x Post-Reform Interaction	-0.036	-	0.720
Mann-Whitney-Type Causal Effect (θ)	0.687	0.673 - 0.701	<0.001

Note: The regression-adjustment models estimate mean-based effects on household per-capita expenditure using the raw and untransformed outcomes. CPM estimates distributional effects without requiring a parametric specification of the outcome distribution, and θ is interpreted as the probability that a randomly selected exposed household has a higher welfare outcome than a comparable less exposed household. Main effects of exposure and the post-reform period (not indicated in the table) were found to be highly significant in all three models.

Taken together, the findings highlight important differences between mean-based and distributional approaches. While conclusions regarding average effects depend on outcome transformation, the Mann-Whitney-type causal effect provides stable evidence of substantial distributional differences between exposure groups. These results illustrate the value of rank-based semiparametric estimands for evaluating policy interventions when outcomes are highly skewed and heterogeneous.

3.3 Subgroup Heterogeneity

Table 3 reports the subgroup-specific estimates of the Mann-Whitney-type causal effect and corresponding 95% confidence intervals.

Table 3. Subgroup Estimates of the Mann-Whitney-Type Causal Effect

Subgroup	$\hat{\theta}$	95% CI
Urban sector	0.699	0.649, 0.749
Rural sector	0.688	0.627, 0.750
Male-headed household	0.675	0.624, 0.726
Female-headed household	0.693	0.642, 0.743
North	0.684	0.607, 0.761
South	0.695	0.647, 0.744

Positive effects of fuel subsidy removal were observed across all subgroups considered. The estimated effect was slightly larger among urban households $(\hat{\theta} = 0.699)$than rural households $(\hat{\theta} = 0.688)$, indicating a stronger degree of distributional separation in urban areas. Similarly, modest differences were observed between male- and female-headed households, suggesting that household structure may influence vulnerability to fuel-price shocks.

Geographic variation was also evident, with somewhat larger effects observed in southern regions than in northern regions. Although the magnitude of these differences was moderate, the results suggest that the welfare consequences of subsidy removal were not uniformly distributed across the country.

Overall, the subgroup analyses reinforce the central finding that the welfare effects of fuel subsidy removal are heterogeneous and that distributional measures provide valuable information beyond average treatment effects. Differences in fuel dependence, market integration, and local economic conditions may shape household responses to subsidy reforms.

4. Conclusion

This study applied the semiparametric cumulative probability model estimator of Sani et al. (2026) to evaluate the welfare consequences of Nigeria's 2023 fuel subsidy removal. Using household per-capita consumption expenditure from the Nigeria General Household Survey, the analysis compared a conventional regression-adjustment benchmark with the Mann-Whitney-type causal effect. It demonstrates that semiparametric rank-based estimators provide robust insights into Nigeria’s fuel subsidy removal, highlighting distributional welfare differences that mean-based methods obscure.

The results showed that conclusions based on mean outcomes were sensitive to outcome scaling and transformation. In contrast, the semiparametric analysis consistently indicated substantial distributional differences between high- and low-exposure households, yielding an estimated Mann-Whitney-type causal effect of 0.687 (95% CI: 0.673–0.701). Subgroup analyses revealed modest heterogeneity across demographic and geographic groups, with somewhat larger effects observed among urban households and households located in southern regions.

Methodologically, the study demonstrates the practical value of semiparametric rank-based causal estimands for analysing highly skewed and heterogeneous welfare outcomes. By focusing on distributional dominance rather than average effects, the Mann-Whitney-type causal effect provides a more reliable framework and an interpretable and robust measure for evaluating policy impacts under non-normal, skewed and heterogeneous outcome distributions.

These findings contribute to the growing literature on semiparametric causal inference and suggest that distributional methods can play an important role in evaluating major economic reforms. Future research may extend this framework to alternative welfare indicators, richer exposure definitions, and other large-scale policy interventions in developing economies.

References

Akpan, S. J., & Uford, I. C. (2024). Effect of fuel subsidy removal on purchase of consumer goods in Nigeria: A literature review. *International Journal of Economics, Business and Social Science Research*, 2(4), 30–38.

Bilker, W. B., & Wang, M. C. (1996). A semiparametric extension of the Mann-Whitney test for randomly truncated data. *Biometrics*, 10-20.

Hart, A. (2001). Mann–Whitney test is not just a test of medians: Differences in spread can be important. *BMJ, 323*(7309), 391–393. https://doi.org/10.1136/bmj.323.7309.391

Mann, H. B., & Whitney, D. R. (1947). On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, 18(1), 50–60.

Odusanya, O. S. E., Ferguson, O. and Chuke, P. I. (2025). Implications of fuel subsidy removal on the Nigerian economy. *International Journal of Innovative Social Sciences & Humanities Research* 13(2):200-206.

Sani, S. S., Blette, B. S., Li, C., Yahaya, A., Dikko, H. G., Usman, A., ... & Shepherd, B. E. (2026). Semiparametric outcome regression-based estimator of Mann–Whitney-type causal effect. *BMC Medical Research Methodology*.

Zhang, Z., Liu, C., Ma, S., & Zhang, M. (2019). Estimating Mann–Whitney-type causal effects for right-censored survival outcomes. *Journal of Causal Inference*, *7*(1), 20180010.