



Maximum likelihood estimation for ordinal randomized response data via the EM algorithm: identifiability and efficiency

A. F. Akintola^{a,*}, O. M. Olayiwola^b, A. A. Akintunde^b, I. A. Osinuga^c

^aDepartment of Mathematics and Statistics, Redeemer's University, Ede, Osun State, Nigeria

^bDepartment of Statistics, College of Physical Sciences, Federal University of Agriculture, Abeokuta, Ogun State, Nigeria

^cDepartment of Mathematics, College of Physical Sciences, Federal University of Agriculture, Abeokuta, Ogun State, Nigeria

Abstract

Surveys on sensitive matters, including security behaviour, substance use, and financial misconduct, are prone to insincere answers when respondents fear disapproval or consequences. Randomized response techniques address this by letting respondents answer through a privacy-preserving device, so no individual reply is ever revealed. The approach is well developed for dichotomous and quantitative questions but has seldom been adapted to ordinal Likert-type items, among the most common measurement tools in behavioural research. This article provides that extension. We propose a scrambling mechanism for four-category ordinal responses driven by a single six-sided die and show that the design is identifiable: the population category probabilities are uniquely determined by the population distribution of scrambled responses, and this can be checked before any data are collected. Because the true category is latent, we derive an expectation-maximization (EM) algorithm with closed-form updates that converges to the unique maximiser of the observed-data likelihood, and we give the asymptotic theory in the three free coordinates the simplex actually has. The maximum-likelihood and method-of-moments estimators are asymptotically equivalent, but they differ appreciably in finite samples. Across five distributions and sample sizes up to 1,000, the maximum-likelihood estimator reduces mean squared error by up to 27% at $n = 200$ for polarized distributions and by construction avoids the boundary violations that force the moment estimator to be projected back onto the simplex in over half of small polarized samples. The result is a practical and reliable tool for measuring sensitive ordinal outcomes without compromising respondent trust.

DOI: [10.46481/asr.2026.5.3.602](https://doi.org/10.46481/asr.2026.5.3.602)

Keywords: Randomized response, Ordinal data, Identifiability, EM algorithm, Maximum likelihood

Article History:

Received: 09 July 2026

Received in revised form: 09 September 2026

Accepted for publication: 17 September 2026

Available online: 02 October 2026

© 2026 The Author(s). Published by the Nigerian Society of Physical Sciences under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

1. Introduction

Attitudinal survey items on sensitive topics, such as academic dishonesty, drug use, tax evasion, and cybersecurity misconduct, are systematically under-reported when asked directly [1]. The randomized response technique (RRT) introduced by Warner [2] protects respondent privacy by placing a known random device between the respondent and the researcher, enabling population-level inference even when no individual reply can be linked to the truth. The body of theory that has grown around the idea is surveyed by

*Corresponding Author Tel. No.: +234-802-896-1466.

Email address: gboyegaa@run.edu.ng (A. F. Akintola)

Chaudhuri and Mukerjee [3], and a meta-analysis of thirty-five years of validation studies by Lensvelt-Mulders *et al.* [4] finds that randomized response designs recover higher prevalence estimates for sensitive behaviours than direct questioning does, which is the practical case for using them.

Over six decades, the methodology has matured for binary items [5, 6] and for continuous-quantity items [7, 8]. For ordinal Likert-scale items, by far the most common attitudinal format in applied survey research, no identifiable maximum-likelihood framework has been proposed, to the best of our knowledge, after searching the randomized response literature on ordinal, polychotomous, and misclassification likelihood models. This paper provides one and characterizes its behaviour relative to the standard method-of-moments alternative.

Two well-developed techniques fail to extend to the ordinal case. The first, which we shall term the unrelated-question approach, originates with Greenberg *et al.* [9] and is developed by Moors [10]. It requires the auxiliary distribution to be either *a priori* known or estimated from another sample, which adds additional free parameters when applied to ordinal settings, causing the model to become non-identifiable once more than two scale categories are present. A second branch of the ordinal scrambling literature proposes scrambling mechanisms but does not establish that the original category distribution can be recovered from them, nor does it prove identifiability.

This omission is common in early ordinal RRT proposals and leaves the consistency of any estimator built on these mechanisms open. A non-identifiable ordinal model produces estimating equations with multiple solutions, and no consistent estimator can exist, regardless of how large the sample is. Beyond identifiability, no published work has solved the likelihood inference problem in this setting: the observed-data likelihood is multinomial, but with a probability vector that is a linear transformation of the unknown parameter, and the resulting score equations have no closed-form solution.

This paper closes both gaps in a single setting using a standard six-sided die as the scrambling device. The proposed mechanism induces a 4×4 column-stochastic matrix P whose entries depend only on the known die probabilities and an explicit boundary rule, with no auxiliary unknowns. We prove identifiability by computing $\det(P) = -1/27$, which guarantees that the population category distribution π is uniquely determined by the population scrambled distribution λ . Identifiability is a population-model property: it establishes that the mapping $\pi \mapsto P\pi$ is injective, not that π can be recovered exactly from a finite sample.

Finite-sample estimates remain variable and, for the method-of-moments estimator, may fall outside the probability simplex. The identifiability condition can be verified before the survey is fielded. The integer entries of P^{-1} yield a method-of-moments estimator that is computable by hand. To overcome the latent-data structure, we treat the unobserved true category as missing and apply the expectation-maximization (EM) algorithm of Dempster *et al.* [11]; the E-step and M-step both have closed-form updates, and the M-step satisfies the simplex constraint by construction at every iteration. The bounded monotone EM sequence converges to the unique maximiser of the observed-data likelihood by Wu's [12] convergence theory and the strict concavity that follows from full column rank of P .

Asymptotic efficiency is then characterized through the Fisher information matrix, and the comparison with method-of-moments is made both analytically (via Cramér–Rao) and numerically (via Monte Carlo simulation across five distributions and five sample sizes). The four contributions of the paper can be summarized briefly. We provide an identifiable ordinal randomized response model that uses a single standard randomization device and whose identifiability can be checked before any data are collected. We derive the EM algorithm for the resulting scrambled multinomial likelihood, with closed-form updates and an interior starting value that keeps the iterates away from the simplex boundary.

We also give the asymptotic theory in the three free coordinates the simplex actually has, including the constrained Fisher information and delta-method covariance for ordinal contrasts, and we show that the maximum-likelihood and method-of-moments estimators are asymptotically equivalent for interior population distributions. Finally, we quantify, through a structured Monte Carlo study, the finite-sample difference that remains: mean-squared-error reductions of up to 27% over the projected method-of-moments at $n = 200$ for polarized distributions, a smaller but persistent advantage at $n = 1,000$, and EM convergence in 99.2–100.0% of replicates.

The rest of the paper is organized in the conventional way. Section 2 sets up the model, proves identifiability, and derives the EM algorithm and asymptotic theory. Section 3 presents the simulation study. Section 4 addresses what it shows, including an additional finite-sample feature of Fisher information confidence intervals that is not captured by standard asymptotic theory. Section 5 is the conclusion.

2. Materials and methods

2.1. Model setup

Let $K = 4$ represent the number of ordered response categories indexed by $\{1, 2, 3, 4\}$, representing a four-point Likert scale (Strongly Disagree, Disagree, Agree, Strongly Agree). For a total population of N respondents, each respondent i has a true latent preference $T_i \in \{1, 2, 3, 4\}$ for the sensitive attitude item. The parameter of interest is defined in Eq. (1) as

$$\pi = (\pi_1, \pi_2, \pi_3, \pi_4)^T \in \Delta^3, \quad \pi_k = \Pr(T_i = k), \quad \sum_{k=1}^4 \pi_k = 1. \quad (1)$$




Figure 1: One-step scrambling mechanism for ordinal randomized response. Every respondent rolls a six-sided die in private, which, along with the actual Likert category, determines the outcome through the known scrambling matrix P , $\det(P) = -1/27$.

The probability simplex is defined in Eq. (2) as

$$\Delta^3 = \left\{ \boldsymbol{\pi} \in \mathbb{R}^4 : \pi_k \geq 0, \sum_{k=1}^4 \pi_k = 1 \right\}. \quad (2)$$

2.2. Scrambling mechanism

For each sampled respondent, the protocol requires a private roll of a fair six-sided die before reporting, and the researcher does not observe the die outcome. The outcome D_i dictates the reported category $R_i \in \{1, 2, 3, 4\}$ according to Eq. (3):

$$R_i = \begin{cases} T_i - 1, & D_i \in \{1, 2\} \text{ and } T_i > 1 \quad (\text{shift down}), \\ T_i, & D_i \in \{3, 4\} \quad (\text{truthful}), \\ T_i + 1, & D_i \in \{5, 6\} \text{ and } T_i < 4 \quad (\text{shift up}). \end{cases} \quad (3)$$

Boundary (edge-protection) cases complete the mechanism: when $T_i = 1$ and $D_i \in \{1, 2\}$, a downward shift would yield $R_i = 0$, which lies outside the scale; edge protection applies, and $R_i = 1$ is reported. Symmetrically, when $T_i = 4$ and $D_i \in \{5, 6\}$, an upward shift would yield $R_i = 5$; edge protection applies, and $R_i = 4$ is reported. With these cases, the mechanism is fully specified, and the three events occur with equal probability as given in Eq. (4):

$$p_- = p_0 = p_+ = \frac{1}{3}. \quad (4)$$

Figure 1 summarizes the mechanism: the respondent's latent category and private die outcome jointly determine the reported category through the scrambling matrix P , whose construction is set out in Section 2.3.

2.3. The scrambling matrix and observed-data likelihood

Define the scrambling matrix P as the 4×4 matrix whose (j, k) entry is $P_{jk} = \Pr(R_i = j \mid T_i = k)$, the conditional probability that a respondent whose true category is k reports category j . Applying the mechanism in Eq. (3) gives

$$P = \frac{1}{3} \begin{bmatrix} 2 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 2 \end{bmatrix}. \quad (5)$$

The matrix in Eq. (5) is tridiagonal (one-step mechanism) and symmetric. Every column sums to one, proving that P is column-stochastic. Using the law of total probability, the marginal distribution of the observed scrambled response is given in Eq. (6):

$$\lambda_j = \sum_{k=1}^4 P_{jk} \pi_k, \quad \text{equivalently} \quad \lambda = P\pi. \quad (6)$$

For n respondents, observed counts $\mathbf{m}$ follow a multinomial distribution with parameters n and $\lambda = P\pi$, giving the observed-data log-likelihood in Eq. (7):

$$\ell(\pi | \mathbf{m}) = \text{const} + \sum_{j=1}^4 m_j \log(P\pi)_j. \quad (7)$$

2.4. Identifiability

Theorem 2.1 (Identifiability). The ordinal scrambling model is identified: P is invertible with $\det(P) = -1/27 \neq 0$.

Proof. Let $P = (1/3)A$, where A is a matrix with integer entries. Cofactor expansion along row 1 yields $\det(A) = -3$, hence

$$\det(P) = \frac{-3}{81} = -\frac{1}{27} \neq 0. \quad (8)$$

The nonzero determinant in Eq. (8) establishes invertibility. $\square$

Remark 1 (Implication). Since P can be calculated only by knowing the die probabilities, identifiability can be checked in theory before data are collected. This differs from unrelated-question designs.

The inverse scrambling matrix. Computing the inverse gives Eq. (9):

$$P^{-1} = \begin{bmatrix} 1 & 1 & -2 & 1 \\ 1 & -2 & 4 & -2 \\ -2 & 4 & -2 & 1 \\ 1 & -2 & 1 & 1 \end{bmatrix}. \quad (9)$$

Corollary 2.1 (Integer inverse). The inverse P^{-1} has integer entries, meaning the method-of-moments estimator is an integer linear combination of observed proportions, computable by hand.

2.5. Method-of-moments estimation

The method-of-moments (MoM) estimator is given in Eq. (10):

$$\widehat{\pi}^{\text{MoM}} = P^{-1} \widehat{\lambda}. \quad (10)$$

Asymptotic normality is stated in Eq. (11) and follows from the multivariate central limit theorem:

$$\sqrt{n}(\widehat{\pi}^{\text{MoM}} - \pi) \xrightarrow{d} N_4(\mathbf{0}, \Sigma^{\text{MoM}}), \quad \Sigma^{\text{MoM}} = P^{-1} \Sigma_{\lambda} P^{-T}, \quad (11)$$

where $\Sigma_{\lambda} = \text{diag}(\lambda) - \lambda\lambda^T$ is the multinomial covariance. The presence of negative entries in P^{-1} means that $\widehat{\pi}^{\text{MoM}}$ can produce components outside $[0, 1]$ in finite samples, with simplex projection frequency reaching 56.22% at $n = 200$ for polarized distributions in our simulation. The maximum-likelihood estimator (MLE) avoids this limitation by construction.

2.6. Maximum likelihood estimation via the EM algorithm

2.6.1. Complete-data formulation

The score equations in Eq. (12)

$$\frac{\partial \ell}{\partial \pi_k} = \sum_{j=1}^4 \frac{m_j P_{jk}}{\lambda_j}, \quad (12)$$

form a nonlinear system with no explicit solution. The EM algorithm augments the observed data with the unobserved true-category assignments. Define complete-data cell counts and column marginals (latent true-category counts). The complete-data log-likelihood is given in Eq. (13) as the standard multinomial:

$$\ell^c(\pi | \mathbf{n}) = \text{const} + \sum_{k=1}^4 n_k \log \pi_k. \quad (13)$$

2.6.2. The EM algorithm

E-step. By Bayes' theorem, the posterior probability that a respondent reporting category j has true category k is given in Eq. (14):

$$w_{jk}^{(t)} = \frac{P_{jk}\pi_k^{(t)}}{\lambda_j^{(t)}}. \quad (14)$$

Summing over observed categories gives the E-step update in Eq. (15):

$$n_k^{(t)} = \sum_{j=1}^4 \frac{m_j P_{jk} \pi_k^{(t)}}{\lambda_j^{(t)}}. \quad (15)$$

M-step. The M-step maximizes the Q -function subject to $\sum_k \pi_k = 1$ and yields Eq. (16):

$$\pi_k^{(t+1)} = \frac{n_k^{(t)}}{n}. \quad (16)$$

Theorem 2 assumes an interior starting value, and a projected method-of-moments estimate need not supply one. If it lands on a face of the simplex, some category carries zero mass, the M-step carries that zero forward at every iteration, and the algorithm can never recover the category. We therefore start from the raw moment estimate, projected onto the simplex when it falls outside, and then move it into the interior by flooring each component at a small positive value and renormalizing as in Eq. (17):

$$\pi_k^{(0)} = c \max\{\tilde{\pi}_k, \varepsilon\}, \quad \varepsilon = 10^{-3}. \quad (17)$$

where $\tilde{\pi}$ denotes the projected moment estimate and c is the normalizing constant that restores the sum to one. This guarantees $\pi^{(0)} \in \text{int}(\Delta^3)$ whatever the moment estimate returns, which is what Theorem 2 requires.

The floor is small relative to the category probabilities considered here, and because the log-likelihood is strictly concave on the interior, the limit does not depend on where in the interior the iteration begins.

2.6.3. Convergence

Theorem 2.2 (Convergence of the EM algorithm). Let $\pi^{(0)} \in \text{int}(\Delta^3)$. The sequence $\{\pi^{(t)}\}$ satisfies (i) $\ell(\pi^{(t+1)} | \mathbf{m}) \geq \ell(\pi^{(t)} | \mathbf{m})$ for all $t \geq 0$; and (ii) $\pi^{(t)} \rightarrow \hat{\pi}^{\text{MLE}}$ as $t \rightarrow \infty$, where $\hat{\pi}^{\text{MLE}}$ is the unique maximizer of $\ell(\pi | \mathbf{m})$ over Δ^3 .

Proof. Throughout, we assume the regularity condition that every observed count is strictly positive, $m_j > 0$ for all j . Since $\lambda = P\pi$ has strictly positive entries for any interior π , this holds with probability tending to one, and it supplies curvature in the argument below; without it, the log-likelihood can be flat in some directions, and the conclusion may fail.

Part (i) is the standard EM monotonicity property [11]: each iteration increases the observed-data log-likelihood or leaves it unchanged at a stationary point. Part (ii) requires two stages, since monotone convergence of the likelihood values does not by itself establish convergence of the parameter sequence. First, $\ell(\cdot | \mathbf{m})$ is continuous on the closed simplex Δ^3 , which is compact, and the iterates lie in it; by Wu [12], every limit point of $\{\pi^{(t)}\}$ is then a stationary point of $\ell(\cdot | \mathbf{m})$.

Second, that stationary point is unique. Because P has full column rank (Theorem 1), the map $\pi \mapsto P\pi$ is injective, and with every $m_j > 0$, the function $\pi \mapsto \sum_j m_j \log(P\pi)_j$ is a strictly concave function composed with an injective linear map, hence strictly concave on $\text{int}(\Delta^3)$. A strictly concave function on a convex set admits at most one maximizer, so the set of stationary points is a singleton, and the sequence converges to the unique global maximizer of $\ell(\cdot | \mathbf{m})$ over Δ^3 . $\square$

2.6.4. Asymptotic properties and efficiency

Because π is confined to the simplex, it carries only three free parameters, so the limit distribution must be stated in three coordinates. Let J be defined in Eq. (18):

$$J = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ -1 & -1 & -1 \end{bmatrix}. \quad (18)$$

The 4×3 contrast matrix obtained by writing $\pi_4 = 1 - \pi_1 - \pi_2 - \pi_3$, and let $\theta = (\pi_1, \pi_2, \pi_3)^T$ denote the free parameter. The ambient information matrix $I(\pi)$ has rank 3, with null space spanned by the all-ones vector, while the constrained information $I_c(\theta)$ is full rank whenever every λ_j is positive. The MLE attains the constrained Cramér–Rao bound.

Theorem 2.3 (Asymptotic normality and efficiency). As $n \rightarrow \infty$,

$$\sqrt{n}(\hat{\theta}^{\text{MLE}} - \theta) \xrightarrow{d} N_3(\mathbf{0}, I_c(\theta)^{-1}), \quad (19)$$

where the ambient and constrained information matrices are, respectively,

$$I(\pi) = P^T \text{diag}(\lambda)^{-1} P, \quad I_c(\theta) = J^T I(\pi) J. \quad (20)$$

Eq. (19) gives the asymptotic limit, while Eq. (20) defines the ambient and constrained information matrices. Comparing with the method-of-moments estimator in these same coordinates gives

$$\Sigma_c^{\text{MoM}} = I_c(\boldsymbol{\theta})^{-1} \quad \text{for every interior } \boldsymbol{\pi}. \quad (21)$$

Equation (21) records an equivalence rather than a strict ordering. Because P is invertible, the raw moment estimator $P^{-1}(\mathbf{m}/n)$ already maximizes the multinomial likelihood whenever it falls inside the simplex, and for any interior $\boldsymbol{\pi}$, the probability of that event tends to one. The two estimators therefore agree to first order, and no strict asymptotic efficiency claim is available for this design.

The advantage documented in Section 3 is a finite-sample one: the raw moment estimator leaves the simplex in more than half of replicates at $n = 200$ under a polarized distribution and must be projected back, whereas the M-step of the EM algorithm is simplex-valued by construction. The simulation finds an aggregate relative efficiency of 1.376 for the polarized distribution at $n = 200$, with a residual advantage of 1.076 at $n = 1,000$.

2.7. Inference for ordinal contrasts

A linear contrast $\gamma = \mathbf{c}^T \boldsymbol{\pi}$ captures the meaningful summaries listed in Eq. (22):

$$\begin{aligned} \gamma_1 &= \pi_3 + \pi_4 && \text{(agreement),} \\ \gamma_2 &= \pi_1 + \pi_4 && \text{(polarisation),} \\ \gamma_3 &= (\pi_3 + \pi_4) - (\pi_1 + \pi_2) && \text{(net agreement),} \\ \gamma_4 &= \sum_{k=1}^4 k\pi_k && \text{(weighted mean).} \end{aligned} \quad (22)$$

Remark 2 (Contrast variances). In finite samples, the estimated variance of a linear contrast under the projected method-of-moments estimator is generally at least as large as under the MLE, for the same reason given in Section 2.6.4: projection introduces boundary bias that the EM estimator avoids by construction. We do not claim a strict asymptotic ordering, since the two estimators are asymptotically equivalent for interior $\boldsymbol{\pi}$.

2.8. Simulation design

We conduct a structured Monte Carlo simulation with $B = 5,000$ replications. Five distributions are examined:

$$\boldsymbol{\pi}^{(1)} = (0.25, 0.25, 0.25, 0.25)^T \quad \text{(uniform),} \quad (23)$$

$$\boldsymbol{\pi}^{(2)} = (0.40, 0.30, 0.20, 0.10)^T \quad \text{(disagreement dominant),} \quad (24)$$

$$\boldsymbol{\pi}^{(3)} = (0.10, 0.20, 0.30, 0.40)^T \quad \text{(agreement dominant),} \quad (25)$$

$$\boldsymbol{\pi}^{(4)} = (0.40, 0.10, 0.10, 0.40)^T \quad \text{(polarized),} \quad (26)$$

$$\boldsymbol{\pi}^{(5)} = (0.10, 0.40, 0.40, 0.10)^T \quad \text{(moderate).} \quad (27)$$

The five simulation settings are specified in Eq. (23), Eq. (24), Eq. (25), Eq. (26), and Eq. (27). across sample sizes $n \in \{200, 400, 600, 800, 1,000\}$. The EM algorithm is initialized at the interior mixture described in Section 2.6.2 and iterated to a tolerance of $\varepsilon = 10^{-6}$. The Monte Carlo standard error for coverage at the 95% level is given in Eq. (28):

$$\sqrt{\frac{0.95 \times 0.05}{5,000}} \approx 0.003. \quad (28)$$

Simplex projection of the raw method-of-moments estimate uses the Euclidean nearest-point projection onto Δ^3 , computed by the sorting algorithm of Wang and Carreira-Perpiñán [13]. This is distinct from the Kullback–Leibler projection, which would coincide with the constrained MLE and would make the comparison circular. What follows, therefore, compares two different estimators, the Euclidean-projected method-of-moments estimator and the EM-based constrained MLE, not two algorithms for one objective.

As shown in Figure 2, five ordinal probability distributions were used in the Monte Carlo simulation, each defined on a four-point Likert scale ($k = 4$).




Figure 2: Five ordinal probability distributions used in the Monte Carlo simulation; each is defined on a four-point likert scale ($k = 4$).

3. Results

3.1. Primary estimation results

Table 1 reports the relative efficiency $RE = MSE(MoM)/MSE(MLE)$, averaged across the four category components, along with the 95% confidence-interval (CI) coverage of both estimators across all sample sizes. Coverage is reported componentwise and then averaged: for each category k , we record the proportion of replicates whose nominal 95% interval for π_k contains the true value, and Table 1 reports the mean of those four proportions.

Method-of-moments intervals use the delta method with covariance $P^{-1}\Sigma_k P^{-T}$; MLE intervals use the marginal standard errors obtained from the Moore–Penrose pseudoinverse of the Fisher information at the estimate. Because the simplex constraint makes that matrix rank 3, the pseudoinverse is the natural generalized inverse to use; the asymptotic theory of Section 2.6.4 is stated in the equivalent three-coordinate form.

RE values greater than 1 indicate an MLE advantage. Method-of-moments coverage uses the delta-method interval; MLE coverage uses the marginal interval obtained from the Moore–Penrose pseudoinverse of the Fisher information evaluated at the estimate, which is the quantity the simulation code computes. The Monte Carlo standard error for coverage at the 95% level is approximately 0.003.

Figure 3 plots mean relative efficiency against sample size for each of the five distributions. The polarized case yields the largest values throughout, reaching 1.376 at $n = 200$ and still 1.076 at $n = 1,000$; the uniform case sits closest to unity and is effectively indistinguishable from the projected moment estimator by $n = 1,000$.

Figure 4 shows 95% interval coverage for the two estimators. Method-of-moments intervals in panel (a) hold close to the nominal level from $n = 200$ onwards. MLE intervals in panel (b) over-cover for the polarized and moderate distributions when n is small, tightening towards nominal as the sample grows.

Figure 5 plots mean squared error for the MLE against the projected moment estimator on logarithmic axes, one panel per distribution. The MLE curve lies below its comparator at every sample size, with the widest separation in the polarized panel; panel (f) contrasts the relative efficiency at $n = 200$ with that at $n = 1,000$.

Every cell in the design improves the MLE over the simplex-projected MoM, with the improvement depending heavily on the shape of the true distribution. The polarized case is the most extreme; at $n = 200$, the aggregate RE is 1.376, with component-level relative efficiencies of 1.550 and 1.479 for the two boundary categories, corresponding to mean-squared-error reductions of 35.5% and 32.4%, compared with relative efficiencies of 1.333 and 1.292 for the interior categories. The gap becomes smaller with increasing n , but does not disappear.

By $n = 1,000$, the polarized RE is 1.076, but the other four distributions are much nearer unity, with RE between 1.000 and 1.021. Even for the uniform distribution, the gap is small for small n (RE = 1.074 at $n = 200$, dropping to 1.000 at $n = 1,000$), as expected by the theory: the efficiency gap is smallest when the scrambled distribution $\lambda = P\pi$ is approximately uniform, because in that regime the Fisher information matrix is well-conditioned and finite-sample boundary effects are minimal. Values of component-level RE in every cell are reported in supplementary file S3.

MLE coverage is conservative for small n , most extreme for the polarized distribution at $n = 200$, where the empirical 95% interval contains the truth 99.4% of the time rather than 95%. The effect originates in the variance estimate rather than in the point estimator. We report it here as an empirical finding; establishing its precise analytic source would require a separate treatment of likelihood inference near the simplex boundary, which is nonregular, and we do not attempt that here.

This effect vanishes as n increases; at $n = 1,000$, coverage in every cell is roughly nominal. None of this affects the conclusions about MSE, and for interval estimation at small n , the logical remedy is a parametric bootstrap of the EM estimator.

3.2. EM algorithm convergence

EM converged to the tolerance $\varepsilon = 10^{-6}$ within the maximum 1,000 iterations in 99.2%–100.0% of the replicates for every cell in the design. The few replicates that did not meet the convergence criterion reached the iteration cap; no divergence was observed.

Table 1: Mean relative efficiency (RE) and 95% CI coverage (%) by true distribution and sample size ($B = 5,000$ replications).

Distribution	n	Mean RE	MoM cov. (%)	MLE cov. (%)
$\pi^{(1)}$ Uniform	200	1.074	96.6	98.1
	400	1.021	94.8	95.6
	600	1.006	94.7	95.2
	800	1.002	95.1	95.5
	1,000	1.000	94.9	95.2
$\pi^{(2)}$ Disagreement dominant	200	1.205	97.7	98.9
	400	1.100	96.7	98.0
	600	1.046	96.6	97.4
	800	1.028	95.3	96.0
	1,000	1.015	94.9	95.5
$\pi^{(3)}$ Agreement dominant	200	1.215	97.7	99.0
	400	1.094	96.8	98.2
	600	1.047	96.6	97.5
	800	1.033	94.9	95.6
	1,000	1.017	95.0	95.6
$\pi^{(4)}$ Polarized	200	1.376	97.5	99.4
	400	1.239	97.0	99.1
	600	1.148	97.2	98.8
	800	1.109	97.1	98.6
	1,000	1.076	96.7	98.2
$\pi^{(5)}$ Moderate	200	1.260	96.1	97.9
	400	1.135	96.1	97.4
	600	1.062	96.4	97.1
	800	1.035	95.1	95.7
	1,000	1.021	94.8	95.1

As shown in Table 2, the mean iteration counts vary from 1.9 (uniform, $n = 1,000$) to 197.5 (polarized, $n = 200$); the polarized distribution always takes the longest, as expected, when probability mass is concentrated in boundary categories.

In that setting, the smallest eigenvalue of the Fisher information matrix shrinks and the local rate of EM convergence slows. Each iteration is $O(K^2)$, where $K = 4$, so the wall-clock time per replicate is far below 1 second, even for the slowest cell.

3.3. Simplex projection

Simplex projection is the routine in which a raw MoM estimate that falls outside the probability simplex is replaced by its closest simplex point. Projection rates depend strongly on both the distribution and the sample size. For the uniform distribution, the rate falls quickly with n : from 16.72% at $n = 200$ to 0.16% at $n = 1,000$. The polarized case is much harder. Projection is needed in 56.22% of replicates at $n = 200$, and the rate is still 17.04% at $n = 1,000$.

The projected MoM is well-defined and would be reported by any practitioner. It is also the comparator in our RE values throughout, but the underlying instability of the raw estimator is a real finite-sample feature that the MLE avoids by construction, since the M-step update is automatically simplex-valued.

EM convergence uses the criterion $\|\pi^{(t+1)} - \pi^{(t)}\|_1 < 10^{-6}$. MoM projection rate is the proportion of replicates in which the raw MoM estimator falls outside the probability simplex. Projection rates for the polarized distribution remain operationally significant




Figure 3: Mean relative efficiency by distribution and sample size.




Figure 4: 95% confidence interval coverage.

across the simulated sample sizes.

Figure 6 summarizes the operational diagnostics.

3.4. Contrast estimators

Coverage of the weighted-mean contrast γ_4 under the method-of-moments estimator is acceptable from $n = 200$ onwards, consistent with Remark 2. The MLE contrast intervals are again conservative at small n , over-covering the nominal 95% level




Figure 5: Mean squared error of MLE (solid, circles) versus MoM (dashed, squares).




Figure 6: Operational diagnostics. (a) MoM projection rate: the polarized distribution remains projection-prone (17.04%) even at $n = 1,000$. (b) Mean EM iterations decline rapidly with n ; the polarized case is slowest, averaging approximately 200 iterations at $n = 200$ and 86 at $n = 1,000$.

by three to four percentage points at $n = 200$ and converging to the nominal level as the sample grows. The pattern matches that of Section 3.1, and as there, we report it as an empirical finding rather than attributing it to a particular feature of the variance estimate.

For practitioners who need tight nominal coverage on this contrast at small n , a parametric bootstrap is the cleaner choice; for point-estimate accuracy, the MLE remains preferable.

Table 2: EM algorithm convergence and MoM projection statistics ($B = 5,000$ replications).

Distribution	n	EM conv.	Mean iter.	Max iter.	MoM proj.
$\pi^{(1)}$ Uniform	200	99.9%	69.2	1,000	16.72%
	400	100.0%	18.2	1,000	3.84%
	600	100.0%	7.8	897	1.20%
	800	100.0%	2.7	883	0.32%
	1,000	100.0%	1.9	886	0.16%
$\pi^{(2)}$ Disagreement dominant	200	99.7%	125.9	1,000	30.44%
	400	99.9%	60.7	1,000	13.64%
	600	100.0%	39.7	916	6.84%
	800	99.6%	22.3	1,000	4.06%
	1,000	99.8%	12.4	1,000	2.08%
$\pi^{(3)}$ Agreement dominant	200	99.8%	121.9	1,000	30.72%
	400	99.9%	59.3	1,000	12.96%
	600	100.0%	37.0	1,000	6.42%
	800	99.7%	21.8	1,000	4.08%
	1,000	99.8%	12.9	1,000	2.16%
$\pi^{(4)}$ Polarized	200	99.5%	197.5	1,000	56.22%
	400	99.9%	149.0	1,000	38.58%
	600	100.0%	127.6	883	28.88%
	800	100.0%	103.1	891	23.26%
	1,000	100.0%	85.8	887	17.04%
$\pi^{(5)}$ Moderate	200	99.2%	148.1	1,000	32.02%
	400	99.8%	66.3	1,000	14.82%
	600	100.0%	48.9	1,000	6.62%
	800	99.3%	24.3	1,000	4.10%
	1,000	99.6%	15.9	1,000	2.56%

3.5. Summary

Two things stand out from the simulation. The MLE is more efficient than the simplex-projected method-of-moments estimator in every cell of the design, the largest gain being a relative efficiency of 1.376 for the polarized distribution at $n = 200$ and a residual advantage of 1.076 at $n = 1,000$. The EM algorithm reached the convergence tolerance in 99.2%–100.0% of replicates within the 1,000-iteration cap.

Two finite-sample features go beyond the asymptotic theory and matter in practice. The confidence interval built from the constrained information is conservative at small n , over-covering the nominal 95% level, with coverage approaching nominal as n grows; we report this empirically and do not attribute it to a specific analytic cause. The raw method-of-moments estimator also falls outside the simplex in more than 17% of replicates at $n = 1,000$ for polarized distributions, so projection is unavoidable in practice for that case.

4. Discussion

Three findings from this work warrant individual discussion.

4.1. The polarized distribution benefits most from the MLE

The raw MoM estimator falls outside the probability simplex in 56% of replicates at $n = 200$ and 17% at $n = 1,000$ for the polarized distribution, and there is no sample size in our experiment at which the rate becomes negligible. The operational implication is concrete. An analyst who chooses MoM is committing in advance to applying simplex projection to a sizeable fraction of their replicates, and even after the projection, the resulting estimator is still less efficient than the MLE by 7–27% in MSE across $n \in [200, 1,000]$.

The MLE avoids the projection step entirely, since the M-step update is automatically simplex-valued; the advantage is finite-sample rather than asymptotic. The two estimators agree in the limit, but it is precisely in the sample sizes surveys actually use that it appears, which makes the MLE a practical default for polarized attitudinal items.

4.2. Conservative MLE coverage at small n

The MLE's 95% confidence interval is systematically conservative at small n , reaching 99.4% coverage for the polarized distribution at $n = 200$. This is a property of the variance estimate that produces the interval rather than of the point estimator, whose mean-squared-error behavior is unaffected. The ambient 4×4 information matrix is rank 3 because the probabilities sum to one, so the marginal standard errors are extracted using its Moore–Penrose pseudoinverse; this is equivalent to working in the three free coordinates θ of Section 2.6.4.

We record the over-coverage as an empirical finding and do not attribute it to any particular feature of that generalized inverse. The over-coverage diminishes as n grows and is essentially nominal in every cell at $n = 1,000$. Establishing its precise analytic source would require a treatment of likelihood inference near the simplex boundary, which is nonregular, and we do not attempt that here.

None of it affects the mean-squared-error comparisons in Table 1 or Figure 5. For practitioners needing tight nominal coverage at small n , a parametric bootstrap of the EM estimator is the appropriate remedy, since it respects the boundary geometry directly. A careful comparison of bootstrap variants is left for future work.

4.3. Pre-survey verifiability

Identifiability through $\det(P)$ is a more useful property than it might first appear. Because P is fully determined by the known die probabilities and the boundary rule, the analyst can verify identifiability by computing $\det(P)$ before any data are collected. Unrelated-question designs offer no comparable check, since identifiability there depends on auxiliary distributions that must be estimated. Recent ordinal-RRT work by Yunusa *et al.* [14] and Rueda *et al.* [15] reinforces the broader point that explicit identifiability checks should be standard in polychotomous extensions of the original Warner framework.

4.4. Limitations and extensions

The framework is developed for $k = 4$ ordered categories, the most common Likert format. Extending it to general k requires deriving the $k \times k$ scrambling matrix induced by the same one-step die mechanism and checking that its determinant is non-zero, and this must be done case by case: identifiability is not automatic. The tridiagonal pattern is singular for some category counts, $k = 3$ and $k = 6$ under a uniform three-way split being the simplest examples, so the extension is not a formality. A general characterization of the values of k for which this mechanism is identifiable would be a worthwhile companion result, and we leave it open.

A second limitation is that we assume full respondent compliance with the die. In practice, some respondents will not roll it, will misread the instruction card, or will override the device and answer as they please; quantifying the resulting bias and developing diagnostics for non-compliance are important extensions, and we intend to take them up separately. A third limitation is that the present design scrambles a single ordinal item.

Related scrambled-response models handle other structures. Yu *et al.* [16] and Tang *et al.* [17] develop non-randomized and two-sample designs for sensitive questions, and Azeem *et al.* [18] treat mean estimation under scrambled quantitative responses. Combining the identifiability argument given here with those designs is a natural direction. Finally, the present paper is methodological. An application of the framework to a real sensitive attitudinal item would test the practical considerations that simulation alone cannot probe.

5. Conclusion

This paper supplies what the randomized response literature has lacked for ordinal scrambled-response data: an identifiable model, an explicit EM algorithm, and an asymptotic theory stated in the right number of coordinates. Three points are worth carrying away.

The first is that identifiability here can be settled before a single questionnaire is distributed. Because the scrambling matrix is fixed by the die probabilities and the boundary rule, its determinant can be computed in advance, and a design that fails the check can be discarded before any respondent is inconvenienced. Unrelated-question designs offer no equivalent guarantee, since their identifiability rests on auxiliary distributions that must themselves be estimated.

The second concerns what the maximum-likelihood estimator actually buys. Asymptotically, it buys nothing: because the scrambling matrix is invertible, the moment estimator already maximizes the likelihood whenever it lands inside the simplex, and for an interior population distribution, it does so with probability tending to one. The gain is a finite-sample one, and it is largest exactly where survey practice is most likely to encounter it.

For polarized attitudes, the raw moment estimator leaves the simplex in more than half of the samples at $n = 200$ and in roughly one sample in six even at $n = 1,000$. Projection is then unavoidable, and the projected estimator carries a boundary bias that the EM estimator never incurs, since its M-step is simplex-valued by construction. For ordinal items expected to polarize, this makes the likelihood estimator a sensible default rather than a refinement.

The third is a caution about interval estimates. Intervals constructed from the constrained information are conservative at small samples, over-covering the nominal level by three to four percentage points, and this should be stated plainly when such intervals are reported. The effect disappears as the sample grows; where tight nominal coverage matters at small n , a parametric bootstrap of the EM estimator is the appropriate alternative. The R code, simulation output, per-component efficiency results, and S4 figure code are provided in supplementary materials S1–S4 so that every result reported here can be reproduced directly.

Data availability

This is a methodological study. No human participants were recruited, and no empirical survey data were collected or analyzed. Every numerical result reported in this article is produced by a Monte Carlo simulation. The complete simulation code (S1), the cell-level simulation output (S2), the per-component relative-efficiency results (S3), and the figure-generation script (S4) are provided as supplementary material with this article. S1 is seeded with `set.seed(20251029)` and runs deterministically within a given R installation. Small numerical differences may arise across R versions and computing platforms; they do not affect any of the comparisons or conclusions reported in the article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this manuscript.

Funding

The authors received no specific funding from any public, commercial, or not-for-profit organisation for this research.

Acknowledgment

We thank the Department of Mathematics and Statistics, Redeemer's University, Ede, Nigeria, for institutional support during the preparation of this manuscript. We are grateful to colleagues whose informal feedback at departmental seminars helped sharpen the framing of the simulation study.

References

- [1] R. Tourangeau & T. Yan, "Sensitive questions in surveys", *Psychological bulletin* **133** (2007) 859. <https://doi.org/10.1037/0033-2909.133.5.859>.
- [2] S. L. Warner, "Randomized response: a survey technique for eliminating evasive answer bias", *Journal of the American statistical association* **60** (1965) 63. <https://doi.org/10.1080/01621459.1965.10480775>.
- [3] A. Chaudhuri & R. Mukerjee, *Randomized response: theory and techniques*, Marcel Dekker, New York, USA, 1988. <https://archive.org/details/randomizedrespon0000chau>.
- [4] G. J. L. M. Lensvelt-Mulders, J. J. Hox, P. G. M. van der Heijden & C. J. M. Maas, "Meta-analysis of randomized response research: thirty-five years of validation", *Sociological methods & research* **33** (2005) 319. <https://doi.org/10.1177/0049124104268664>.
- [5] R. F. Boruch, "Assuring confidentiality of responses in social research: a note on strategies", *The American sociologist* **6** (1971) 308. <https://www.jstor.org/stable/27701807>.
- [6] J. A. Fox & P. E. Tracy, *Randomized response: a method for sensitive surveys*, SAGE Publications, Beverly Hills, USA, 1986.
- [7] B. H. Eichhorn & L. S. Hayre, "Scrambled randomized response methods for obtaining sensitive quantitative data", *Journal of statistical planning and inference* **7** (1983) 307. [https://doi.org/10.1016/0378-3758\(83\)90002-2](https://doi.org/10.1016/0378-3758(83)90002-2).
- [8] Diana, G. & Perri, P. F., "New scrambled response models for estimating the mean of a sensitive quantitative character", *Journal of applied statistics*, **37** (2010) 11, 1875-1890. <https://doi.org/10.1080/02664760903186031>.
- [9] B. G. Greenberg, A. L. A. Abul-El, W. R. Simmons & D. G. Horvitz, "The unrelated question randomized response model: theoretical framework", *Journal of the American statistical association* **64** (1969) 520. <https://doi.org/10.1080/01621459.1969.10500991>.
- [10] Moors, J. J. A. , "Optimization of the unrelated question randomized response model", *Journal of the American statistical association* **66** (1971) 629. <https://doi.org/10.1080/01621459.1971.10482320>.
- [11] A. P. Dempster, N. M. Laird & D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm", *Journal of the Royal statistical society series B* **39** (1977) 1. <https://doi.org/10.1111/j.2517-6161.1977.tb01600.x>.

- [12] C. F. J. Wu, "On the convergence properties of the EM algorithm", *The Annals of statistics* **11** (1983) 95. <https://doi.org/10.1214/aos/1176346060>.
- [13] W. Wang & M. Á. Carreira-Perpiñán, "Projection onto the probability simplex: an efficient algorithm with a simple proof, and an application", *arXiv* (2013). <https://arxiv.org/abs/1309.1541>.
- [14] M. A. Yunusa, A. Audu, U. Usman, K. O. Aremu & M. Aphane, "Calibrated-two optional randomized response techniques (C-TORRT) for the estimation of quantitative sensitive variable information", *PLOS ONE* **21** (2026) e0339271. <https://doi.org/10.1371/journal.pone.0339271>.
- [15] M. Rueda, B. Cobo & P. F. Perri, "Randomized response estimation in multiple frame surveys", *International journal of computer mathematics* **97** (2020) 189. <https://doi.org/10.1080/00207160.2018.1476856>.
- [16] J. W. Yu, G. L. Tian & M. L. Tang, "Two new models for survey sampling with sensitive characteristic: design and analysis", *Metrika* **67** (2008) 251. <https://doi.org/10.1007/s00184-007-0131-x>.
- [17] M. L. Tang, Q. Wu, G. L. Tian & J. H. Guo, "Two-sample non randomized response techniques for sensitive questions", *Communications in statistics, theory and methods* **43** (2014) 408. <https://doi.org/10.1080/03610926.2012.657323>.
- [18] Azeem, M., Asadullah, Ijaz, M., Hussain, S., Salahuddin, N., & Salam, A., "A novel randomized scrambling technique for mean estimation of a finite population", *Heliyon* **10** (2024) 11. <https://doi.org/10.1016/j.heliyon.2024.e31690>