Sociological Methodology
2015, Vol. 45(1) 320–356

DOI: 10.1177/0081175015570097
http://sm.sagepub.com

7

SHRINKAGE ESTIMATION OF LOG-ODDS RATIOS FOR COMPARING MOBILITY TABLES

Xiang Zhou*

Abstract

Statistical analysis of mobility tables has long played a pivotal role in comparative stratification research. This article proposes a shrinkage estimator of the log-odds ratio for comparing mobility tables. Building on an empirical Bayes framework, the shrinkage estimator improves estimation efficiency by "borrowing strength" across multiple tables while placing no restrictions on the pattern of association within tables. Numerical simulation shows that the shrinkage estimator outperforms the usual maximum likelihood estimator (MLE) in both the total squared error and the correlation with the true values. Moreover, the benefits of the shrinkage estimator relative to the MLE depend on both the variation in the true log-odds ratio and the variation in sample size among mobility regimes. To illustrate the effects of shrinkage, the author contrasts the shrinkage estimates with the usual estimates for the mobility data assembled by Hazelrigg and Garnier for 16 countries in the 1960s and 1970s. For mobility tables with more than two categories, the shrinkage estimates of log-odds ratios can also be used to calculate summary measures of association that are based on aggregations of log-odds ratios. Specifically, the author constructs an adjusted estimator of the Altham index and, with a set of calibrated simulations, demonstrates its usefulness in

*University of Michigan, Ann Arbor, USA

Corresponding Author:
Xiang Zhou, University of Michigan, Institute for Social Research, 426 Thompson Street, Ann Arbor, MI 48104, USA
Email: xiangzh@umich.edu

enhancing both the precision of individual estimates and the accuracy of cross-table comparisons. Finally, using two real data sets, the author shows that in gauging the overall degree of social fluidity, the adjusted estimator of the Altham index agrees more closely with results from the Unidiff model than does the direct estimator of the Altham index.

Keywords

shrinkage estimator, log-odds ratio, mobility tables, empirical Bayes methods, log-linear models

1. INTRODUCTION

Comparative mobility analysis has long been at the core of social stratification research. To investigate how patterns of intergenerational mobility differ across countries or vary over time, stratification researchers typically compare a collection of mobility tables that cross-classify fathers and sons by their occupations or classes. To draw such comparisons, researchers until the 1970s had relied on simple calculations of inflow and outflow rates (e.g., Lipset and Zetterberg 1956; Miller 1960) or the construct of "mobility ratios" (e.g., Glass and Berent 1954; Rogoff 1953), both of which turned out to be inadequate to separate changes in relative mobility (also known as exchange mobility, circulation mobility, or social fluidity) from changes in marginal distributions (i.e., structural mobility).[1] Beginning in the late 1960s, thanks to the pioneering work of Leo Goodman (1968, 1969), it has been recognized that all associations in an $I \times J$ contingency table can be captured by a sufficient set of $(I-1)(J-1)$ odds ratios.[2] This fundamental discovery paved the way for the subsequent development of log-linear and log-multiplicative models (e.g., Duncan 1979; Goodman 1979; Hauser 1980), in which the natural logarithms of odds ratios are expressed as regression coefficients or their linear combinations.

Given the centrality of odds ratios in depicting the structure of row-column association, a natural approach to comparing mobility tables, as suggested by Goodman (1969), is to directly compare their corresponding (log) odds ratios in search of similarities and differences. Although mobility studies in sociology have been dominated by log-linear modeling since the 1970s, this older model-free approach has its own appeal because it allows a panoramic view of the association between origin and destination without invoking parametric assumptions (see Hout and

Guest [2013] for an illustration). Meanwhile, using log-odds ratios as building blocks, Altham (1970) proposed a number of aggregate measures of association for comparing contingency tables. One of these measures (see section 3) has been recently used to examine long-term trends in occupational mobility in Great Britain and the United States (Ferrie 2005; Long and Ferrie 2007, 2013).

Unlike log-linear modeling, the model-free approach to comparing mobility tables imposes no parametric constraints on the pattern of association between origin and destination. Instead, it requires that every log-odds ratio be estimated separately from data. Estimation of single log-odds ratios, however, can be highly imprecise in practice. Indeed, the usual maximum likelihood estimator (MLE) of the log-odds ratio (i.e., $\log\frac{n_{11}n_{22}}{n_{12}n_{21}}$) will be accompanied by a large standard error unless all of the associated cells contain many cases,[3] a condition that often fails for real mobility tables. As a result, direct comparisons in sample log-odds ratios across tables are prone to conflate true variations in relative mobility with sampling fluctuations. On one hand, if relative mobility is constant and trendless in all complex societies, as implied by the hypothesis of constant social fluidity (CSF; Erikson and Goldthorpe 1992; Featherman, Jones, and Hauser 1975; Grusky and Hauser 1984), the observed differences will stem entirely from sampling and measurement errors. On the other hand, if social fluidity does differ across countries and change over time, sampling variability may also contaminate empirical comparisons between mobility regimes. In particular, when the mobility tables under investigation vary greatly in sample size, the relatively sparse tables are more likely to be estimated at the extremes of the mobility spectrum because they are subject to larger sampling errors. Because sample size is presumably unrelated to the true amount of social fluidity, this statistical artifact may distort the rank order of mobility regimes in the size of origin-destination association. Such a distortion can be substantively significant unless sampling errors are negligible relative to systematic variations among mobility regimes. The latter condition, unfortunately, seldom holds in comparative mobility research.

In log-linear modeling, estimation uncertainty is partly alleviated through parametric assumptions. For example, the CSF model assumes no cross-table variation in all log-odds ratios, and the Unidiff model (Erikson and Goldthorpe 1992; Xie 1992) stipulates that the relative magnitudes of different log-odds ratios are uniform in all tables. These assumptions, however, may accord poorly with real data. In this article, I propose a shrinkage method for estimating log-odds ratios that attempts to enhance estimation efficiency without explicitly constraining the patterns of row-column association. Building on an empirical Bayes model (Efron and Morris 1973; Fay and Herriot 1979), the shrinkage estimator “borrows strength” across multiple tables while placing no restrictions on the structure of association within tables. As I will show by simulation, the shrinkage method leads to lower total squared errors than does the usual MLE of the log-odds ratio. More important, when tables vary greatly in sample size—a situation that we often encounter in comparative mobility analysis—the shrinkage estimates exhibit markedly higher correlations with the true log-odds ratios than do the usual estimates. Therefore, the shrinkage method can enhance the accuracy of cross-table comparisons in the degree of row-column association. Moreover, the shrinkage estimates of log-odds ratios can be used to calculate summary measures of association that are based on aggregations of log-odds ratios. To illustrate this point, I construct an adjusted estimator of the Altham index (Altham 1970; Altham and Ferrie 2007), and, with a set of calibrated simulations, demonstrate its usefulness in enhancing both the precision of individual estimates and the accuracy of cross-table comparisons. Finally, using two sets of real mobility tables, I show that in gauging the overall degree of social fluidity, the adjusted estimates of the Altham index agree more closely with results from the Unidiff model than do direct estimates of the Altham index.

2. SHRINKAGE ESTIMATION OF LOG-ODDS RATIOS

2.1. *Usual Estimator of the Log-odds Ratio*

Let us consider K 2×2 contingency tables, which, say, cross-classify fathers and sons according to nonmanual and manual classes in K countries. Denoting by n_{ijk} the cell frequency pertaining to the ith row and the jth column in country k, the observed log-odds ratios for these tables can be expressed as

$$Y_k = \log \frac{n_{11k} n_{22k}}{n_{12k} n_{21k}}, k = 1, 2, \cdots K. \tag{1}$$

Assuming a multinomial sampling distribution for each country, these sample log-odds ratios are also the maximum likelihood estimates of population log-odds ratios.[4] They are therefore asymptotically normal, that is,

$$\sqrt{n_{++k}}(Y_k - \theta_k) \xrightarrow{d} N(0, V_k),$$

where n_{++k} and θ_k represent the sample size and the population log-odds ratio for country k. Using the delta method, it is not hard to show that the asymptotic variance of Y_k is

$$\sigma_k^2 = \frac{V_k}{n_{++k}} = \frac{1}{n_{++k}\pi_{11k}} + \frac{1}{n_{++k}\pi_{12k}} + \frac{1}{n_{++k}\pi_{21k}} + \frac{1}{n_{++k}\pi_{22k}},$$

where the π_{ijk}'s denote the unknown cell probabilities (Agresti 2002:75–76).

Substituting the observed proportions for the π_{ijk}'s, we obtain a sample estimate of σ_k^2:

$$\widehat{\sigma_k^2} = \frac{1}{n_{11k}} + \frac{1}{n_{12k}} + \frac{1}{n_{21k}} + \frac{1}{n_{22k}}. \tag{2}$$

Because there is a finite, however small, probability that any of the four cells are zero, the observed log-odds ratio (equation 1) may equal ∞ or $-\infty$. In such cases, a common practice is to add one half to all of the four cell frequencies, yielding a modified estimator (Agresti 2002:71):

$$\widetilde{Y_k} = \log\frac{(n_{11k} + 0.5)(n_{22k} + 0.5)}{(n_{12k} + 0.5)(n_{21k} + 0.5)}.$$

Haldane (1956) showed that this modification reduces the sampling bias from the order of $O(n^{-1})$ to the order of $O(n^{-2})$. Moreover, Gart and Zweifel (1967) noted that the corresponding variance estimator

$$\widetilde{\sigma_k^2} = \frac{1}{n_{11k} + 0.5} + \frac{1}{n_{12k} + 0.5} + \frac{1}{n_{21k} + 0.5} + \frac{1}{n_{22k} + 0.5}$$

is an unbiased estimator of $\mathrm{Var}(\widetilde{Y_k})$ except for terms of $O(n^{-3})$. I therefore adopt these adjustments in the case of zero cells throughout the rest of the article.[5]

Since the observed log-odds ratio (equation 1) coincides with the MLE, it is consistent and asymptotically efficient. Nonetheless, the asymptotic variance estimator (equation 2) indicates that the MLE can be highly imprecise in small samples: Unless all of the four cells contain many cases, the standard error will be very large. As a result, if we directly compare the observed log-odds ratios from different tables, those from relatively sparse tables will be more likely to be ranked at the extremes. This is undesirable because sample size is presumably unrelated to the true degree of association. The shrinkage approach I present below aims to improve both the precision of estimates from sparse tables and the accuracy of ranking among different mobility regimes.

2.2. *Empirical Bayes Shrinkage*

To explicate the shrinkage approach, let us first accept the normal approximations of the observed log-odds ratios, that is,

$$Y_k|\theta_k \overset{indep}{\sim} N(\theta_k, \widehat{\sigma_k^2}). \tag{3}$$

Now consider a Bayes model in which the population log-odds ratios themselves follow a normal prior

$$\theta_k \overset{i.i.d.}{\sim} N(\mu, \tau^2), \tag{4}$$

where μ and τ^2 are hyperparameters representing the prior mean and the prior variance of the unknown θ_k's. It is easy to show that the posterior distribution of θ_k is also normal, and the Bayes estimator, that is, the posterior mean, can be written as

$$E(\theta_k|Y_k) = \mu + (1 - \frac{\widehat{\sigma_k^2}}{\tau^2 + \widehat{\sigma_k^2}})(Y_k - \mu). \tag{5}$$

Estimating the hyperparameters μ and τ^2 directly from the data, say, through maximizing the marginal likelihood, leads to an empirical Bayes estimator (Efron and Morris 1973, 1975)

$$\hat{\theta}_k^{EB} = \hat{\mu} + (1 - \frac{\widehat{\sigma_k^2}}{\widehat{\tau^2} + \widehat{\sigma_k^2}})(Y_k - \hat{\mu}). \tag{6}$$

In the statistics literature, $\hat{\theta}_k^{EB}$ has been described as a shrinkage estimator because it "shrinks" the observed outcome Y_k toward the estimated prior mean $\hat{\mu}$ with a shrinkage factor of $\frac{\widehat{\sigma_k^2}}{\widehat{\tau^2}+\widehat{\sigma_k^2}}$. The shrinkage factor, clearly, depends on the precision of the observation Y_k: the larger is the sampling variance $\widehat{\sigma_k^2}$, the stronger is the degree of shrinkage. Indeed, the empirical Bayes estimator can be expressed as a precision-weighted average between Y_k and $\hat{\mu}$ (Raudenbush and Bryk 1985, 2002):

$$\hat{\theta}_k^{EB} = \frac{1/\widehat{\tau^2}}{1/\widehat{\tau^2} + 1/\widehat{\sigma_k^2}}\hat{\mu} + \frac{1/\widehat{\sigma_k^2}}{1/\widehat{\tau^2} + 1/\widehat{\sigma_k^2}}Y_k,$$

where the weight accorded to Y_k is proportional to its sampling precision $1/\widehat{\sigma_k^2}$ and the weight accorded to $\hat{\mu}$ is proportional to $1/\widehat{\tau^2}$, a measure of the concentration of the unknown θ_k's around the prior mean μ.

Because the shrinkage factor in the posterior mean (equation 5) is a convex function of the prior variance τ^2, a substitution of a nearly unbiased estimate $\widehat{\tau^2}$ for τ^2 would produce an upward bias for the shrinkage factor $\frac{\widehat{\sigma_k^2}}{\widehat{\tau^2}+\widehat{\sigma_k^2}}$ (by Jensen's inequality). To alleviate this problem, Morris (1983) suggested that the estimator (equation 6) be replaced by

$$\hat{\theta}_k^{EB} = \hat{\mu} + [1 - \frac{(K-3)\widehat{\sigma_k^2}}{(K-1)\left(\widehat{\tau^2} + \widehat{\sigma_k^2}\right)}](Y_k - \hat{\mu}), \tag{7}$$

where the multiplying constant $\frac{K-3}{K-1}$ is used to offset the bias of $\frac{\widehat{\sigma_k^2}}{\widehat{\tau^2}+\widehat{\sigma_k^2}}$ as an estimate of the shrinkage factor $\frac{\sigma_k^2}{\tau^2+\sigma_k^2}$.

The empirical Bayes framework sketched above was initially proposed by Efron and Morris (1973, 1975) to interpret the James-Stein rule for estimating multivariate normal means. Indeed, Stein (1956) and James and Stein (1961) discovered that for simultaneous estimation of unrelated normal means, the usual MLE (i.e., Y_k's) can be inadmissible and dominated by a shrinkage estimator similar in form to the empirical

Bayes estimator (equation 7). On the other hand, the empirical Bayes method closely parallels the notion of best linear unbiased prediction (BLUP) in random-effects models (Robinson 1991). Specifically, when both the prior variance τ^2 and the sampling variances σ_k^2 are known, it can be shown that the following statistic minimizes the mean squared error between θ_k and any unbiased estimator of θ_k that is linear in the Y_k's (Harville 1976):

$$\hat{\theta}_k^{BLUP} = \hat{\mu} + (1 - \frac{\sigma_k^2}{\tau^2 + \sigma_k^2})(Y_k - \hat{\mu}). \tag{8}$$

Here $\hat{\mu} = \sum_{k=1}^{k} w_k Y_k / \sum_{k=1}^{k} w_k$ is the minimum variance unbiased estimator (MVUE) of μ, where $w_k = 1/(\tau^2 + \sigma_k^2)$. Replacing the variance components τ^2 and σ_k^2 with their estimates would yield the *empirical best linear unbiased predictor* (EBLUP) of θ_k, which coincides with the empirical Bayes estimator (equation 7), except for the lack of the multiplying constant $\frac{K-3}{K-1}$.

While the theoretical work by James and Stein (1961) demonstrates the advantage of shrinkage in a fixed-effects world, the concepts of BLUP and EBLUP justify the empirical Bayes estimator through a random-effects formulation. From either perspective, the key idea is to reduce the influence of sampling variability by "borrowing strength" from other observations (as reflected in $\hat{\mu}$). Because the shrinkage factor roughly equals the ratio of the sampling variance σ_k^2 to the overall variance of Y_k (i.e., $\tau^2 + \sigma_k^2$), the shrinkage rule may be interpreted as "purging" sampling errors from the estimation of true parameters. This procedure can be highly effective when sampling uncertainty is substantial relative to the true variation among the parameters of interest. As illustrated by Efron and Morris (1975), given data from the first 45 at-bats of 18 Major League Baseball players in the 1970 season, the shrinkage approach performs much better than the MLE in predicting their future batting averages. More recently, Savitz and Raudenbush (2009) showed that similar types of shrinkage estimators can improve the precision and predictive validity of ecometric measures in neighborhood studies. Considering that observed log-odds ratios frequently suffer from large sampling errors, we expect that the shrinkage approach can significantly enhance the estimation precision of log-odds ratios by pooling data from multiple mobility tables.

Meanwhile, we notice from equation (7) that the degree of shrinkage is higher for observations with larger sampling variances. This relationship is intuitive because the need for "borrowing strength" should be stronger for relatively imprecise estimates. Differences in the degree of shrinkage, moreover, can alter the rank order of the estimates; that is, the shrinkage estimates may rank the population log-odds ratios differently from the observed log-odds ratios. Efron and Morris (1975) noted that the empirical Bayes method typically outperforms MLE in ordering the true values. Therefore, besides improving the estimation precision of individual log-odds ratios, the shrinkage approach can also enhance the accuracy of cross-table comparisons.

2.3. *Estimation, Inference, and Implementation*

To empirically estimate μ and τ^2, a natural idea is to derive their MLE on the basis of the joint marginal distribution

$$Y_k \overset{indep}{\sim} N(\mu, \tau^2 + \widehat{\sigma_k^2}).$$

Unfortunately, the likelihood equation in this case defies an analytical solution. I now describe an alternative approach proposed by Carter and Rolph (1974), one that is closely related to the procedures used in Fay and Herriot (1979), Morris (1983), and Sidik and Jonkman (2005). As mentioned above, when τ^2 is known, the MVUE of μ is given by the weighted average of the Y_k's

$$\hat{\mu}(\tau^2) = \frac{\sum_{k=1}^{K} w_k(\tau^2) Y_k}{\sum_{k=1}^{K} w_k(\tau^2)},$$

where the weights are

$$w_k(\tau^2) = \frac{1}{\tau^2 + \widehat{\sigma_k^2}}.$$

Here $w_k(\tau^2)$ and $\hat{\mu}(\tau^2)$ highlight their dependence on τ^2. Meanwhile, we observe that the weighted sum of squared deviations of the Y_k's follows a chi-square distribution with $K - 1$ degrees of freedom, that is,

$$\sum_{k=1}^{K} w_k(\tau^2)\left(Y_k - \hat{\mu}(\tau^2)\right)^2 \sim \chi^2_{K-1}.$$

Thus we have

$$\mathrm{E}\left[\sum_{k=1}^{K} w_k(\tau^2)(Y_k - \hat{\mu}(\tau^2))^2\right] = K - 1.$$

Carter and Rolph (1974) suggested that τ^2 be estimated as the unique positive solution that satisfies

$$\sum_{k=1}^{K} w_k\left(\widehat{\tau^2}\right)\left(Y_k - \hat{\mu}\left(\widehat{\tau^2}\right)\right)^2 = K - 1.$$

In the case in which no positive solution exists, $\widehat{\tau^2}$ is set to be zero. To solve the above equation, a simple Newton-Raphson procedure was described by Fay and Herriot (1979:276), which typically converges in fewer than ten iterations. With the converged value of $\widehat{\tau^2}$, the prior mean μ is estimated accordingly as $\hat{\mu}\left(\widehat{\tau^2}\right)$. By plugging $\hat{\mu}$ and $\widehat{\tau^2}$ into equation (7), we obtain the empirical Bayes estimates of the unknown θ_k's.

To fully assess the uncertainty of the empirical Bayes estimator (equation 7), we must take into account the estimation of μ, τ^2, and σ_k^2's. To avoid analytical challenges, I now consider a naive estimator of the standard error of $\hat{\theta}_k^{EB}$ that treats the variance estimates $\widehat{\tau^2}$ and $\widehat{\sigma_k^2}$'s as the true underlying parameters. Denoting by B_k the shrinkage factor $\frac{(K-3)\widehat{\sigma_k^2}}{(K-1)\left(\widehat{\tau^2}+\widehat{\sigma_k^2}\right)}$ in equation (7), the mean squared error between $\hat{\theta}_k^{EB}$ and θ_k can be written as

$$\begin{aligned}
&E\left(\hat{\theta}_k^{EB} - \theta_k\right)^2 = E[(1 - B_k)Y_k + B_k\hat{\mu} - \theta_k]^2 \\
&= E[(1 - B_k)(Y_k - \theta_k) + B_k(\hat{\mu} - \theta_k)]^2 \\
&= (1 - B_k)\widehat{\sigma_k^2} + 2(1 - B_k)B_k\left(\frac{w_k}{\sum w_k}\right)\widehat{\sigma_k^2} + B_k^2(\widehat{\tau^2} - \frac{2w_k\widehat{\tau^2}}{\sum w_k} + \frac{1}{\sum w_k}).
\end{aligned}$$

Therefore, by taking the square root of the right-hand side, we obtain an estimator of the standard error of $\hat{\theta}_k^{EB}$. Alternatively, we can fit random-effects models using standard software for meta-analysis (such as the metafor package in R; see Viechtbauer 2010) and extract estimates of BLUPs and their standard errors, which should be very close to the empirical Bayes estimates.

The standard error derived above tends to underestimate the uncertainty of $\hat{\theta}_k^{EB}$'s because it ignores the estimation of τ^2 and σ_k^2's. A fully

Bayesian approach, as noted by Raudenbush and Bryk (2002), will take account of the estimation uncertainty of μ, τ^2, and θ_k's simultaneously. To build a full Bayes model, we may supply the hyperparameters μ and τ^2 with noninformative priors (e.g., by setting a normal prior with a variance of 10^6 for μ and a uniform prior from 0 to 10^4 for τ^2). Such a model can be easily implemented using standard Markov chain Monte Carlo software such as BUGS. In section 2.5, I illustrate both the empirical Bayes and the full Bayes methods using a set of 16 mobility tables.

2.4. *Usual Estimator versus Shrinkage Estimator in Simulation*

We now turn our attention back to the setting of K 2×2 mobility tables, each representing a country. As noted earlier, the shrinkage factor is decided by the sampling variance of the observed log-odds ratio relative to the true variation in log-odds ratio among the K countries. The influence of shrinkage, therefore, should be stronger when *the true variation in mobility* is relatively small compared with sampling errors. On the other hand, because sampling variance typically differs from country to country, the shrinkage estimates may exhibit a different rank order from that of the usual estimates. Clearly, the extent of this discrepancy should depend on *the extent of variation in sample size* among these countries. In this subsection, I use numerical simulation to examine how potential advantages of the shrinkage approach vary along these two dimensions. I compare the performance between the usual estimator (equation 1) and the shrinkage estimator (equation 7) in two aspects: (1) total squared error and (2) correlation with the true log-odds ratios.

Let us consider 100 2×2 mobility tables depicting, say, intergenerational mobility between white-collar and blue-collar occupations in 100 countries.[6] Following the convention in mobility table analysis, I represent father's occupation in rows and son's occupation in columns. In this simulation, I assume that these countries are at the same stage of industrial development such that 40% of the sample is from white-collar origin in all of the 100 mobility tables. In other words, the row marginal distribution is fixed to be (.4, .6). Despite the homogeneous origin distribution, I allow these countries to vary in the extent of relative mobility as measured by the log-odds ratio. In particular, I create three scenarios in which the true variation in log-odds ratio among these countries is small, medium, and large. Suppose that a son's occupation given a father's occupation follows a binomial distribution, and use $p^k_{1|1}$ and $p^k_{1|2}$

to denote the probabilities of working in a white-collar occupation respectively for a person from white-collar origin and for a person from blue-collar origin in country k. I assume that $p^k_{1|1}$ and $p^k_{1|2}$ are independently and uniformly distributed around .7 and .3, respectively, which means that the probability of being immobile (i.e., staying in the main diagonal of the table) is about .7 for both white-collar and blue-collar occupations. I then construct the three scenarios by letting the range of the two uniform distributions be .08, .16, and .24.[7] In other words, $p^k_{1|1}$ and $p^k_{1|2}$ are independently drawn from the following two distributions:

$$p^k_{1|1} \overset{i.i.d.}{\sim} \text{Uniform}(.7 - .04^*\alpha, .7 + .04^*\alpha), k = 1, 2, \cdots 100, \tag{9}$$

$$p^k_{1|2} \overset{i.i.d.}{\sim} \text{Uniform}(.3 - .04^*\alpha, .3 + .04^*\alpha), k = 1, 2, \cdots 100, \tag{10}$$

where the parameter α, which may take 1, 2, and 3, is used to generate settings in which the true variation in log-odds ratio is small, medium, and large.

The three scenarios above differ in the true variation of log-odds ratio and thus in the estimate of τ^2 in equation (7), which will affect the shrinkage factor uniformly for all countries. As mentioned earlier, the contrasts between the shrinkage estimator and the usual estimator may also depend on the amount of variation in sample size among the mobility tables, which shapes the variation among the $\widehat{\sigma^2_k}$'s. Therefore, I also compare the performance between the two estimators as variation in sample size changes from very small to very large. Specifically, I assume that the sample size follows a log-uniform distribution as below:

$$\log n_{++k} \overset{i.i.d.}{\sim} \mathit{Uniform}\left(\log 800^* 2^{\beta}, \log 1250^* 2^{\beta}\right), k = 1, 2, \cdots 100, \tag{11}$$

where n_{++k} denotes the sample size for country k. I vary the parameter β from 0 to 4 with a step size of 1, thereby generating five scenarios with a gradual change in the variation of sample size while fixing the medium sample size among these countries to be about 1,000. For example, sample size will range between 800 and 1,250 when β takes 0 but range between 50 and 20,000 when β takes 4.

In this simulation, I exhaust all possible combinations of α and β, resulting in $3 \times 5 = 15$ scenarios. For each of these scenarios, I generated the 100 mobility tables in the following steps:

1. For each table k, generate the sample size using $n_{++k} = \lfloor \exp(M) \rceil$, where M is a random draw from the uniform distribution shown in equation 11 and $\lfloor \exp(M) \rceil$ means taking the integer closest to exp(M).
2. Calculate the row marginals (n_{1+k}, n_{2+k}) by assigning 40% of the sample size n_{++k} to the first category (i.e., white collar).
3. Generate the transition probabilities $p^k_{1|1}$ and $p^k_{1|2}$ using the uniform distributions shown in equations 9 and 10.
4. Create the mobility table $(n_{11k}, n_{12k}, n_{21k}, n_{22k})$ using binomial draws for each row, that is, binomial $(n_{1+k}, p^k_{1|1})$ for the first row and binomial $(n_{2+k}, p^k_{1|2})$ for the second row.

Given the simulated tables, I applied both the usual estimator (equation 1) and the empirical Bayes estimator (equation 7) to estimate the log-odds ratios. I then evaluated the performance of the two estimators using two criteria: (1) total squared error, i.e., $\sum_{k=1}^{100} \left(\hat{\theta}_k - \theta_k\right)^2$, and (2) Pearson's correlation coefficient (among the 100 countries), that is, Cor($\hat{\theta}_k, \theta_k$). To smooth random fluctuations, I averaged these two measures over 500 iterations of the above procedures (data generation, estimation, and evaluation) for each of the 15 scenarios.

Figure 1 presents the results, with panel A for total squared errors and panel B for the correlation coefficients. In both panels, I represent the usual estimator in squares and the shrinkage estimator in triangles. The three scenarios in which the true variation in log-odds ratio is small, medium, and large are represented respectively by solid, dashed, and dotted lines. First, we observe that in virtually all of the 15 scenarios, the shrinkage estimator exhibits lower total squared errors and higher correlations with the true values than does the usual estimator. This is consistent with theoretical results on joint estimation of normal means as discussed by Efron and Morris (1973, 1975). Second, as shown by both panels, the benefits of the shrinkage estimator are greater when the true variation in log-odds ratio is smaller. This relationship is intuitive because the shrinkage approach is essentially pooling information across cases, which should be more effective when these cases are more similar to each other. We also note that for both estimators, the correlation with the true values increases as the true variation in log-odds ratio increases. This is because when the true differences are larger, they are less likely to be confounded by sampling fluctuations and thus more likely to be detected from the data. Finally, reading along the x-axis, we find that the advantage of the shrinkage estimator becomes more pronounced as

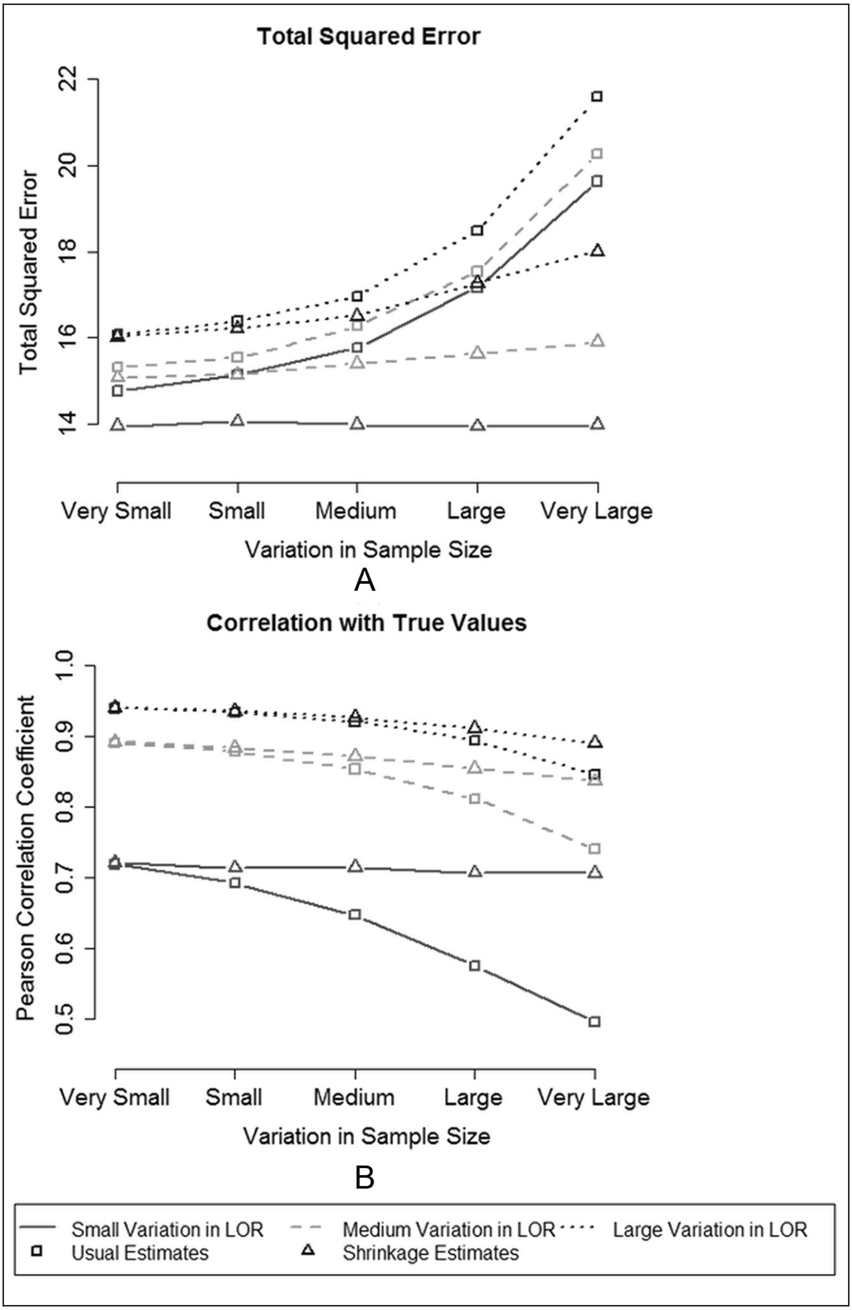

Figure 1. Usual estimator versus empirical Bayes estimator of the log-odds ratio (LOR) in total squared error (A) and Pearson's correlation with the true values (B) under different scenarios.

the variation in sample size increases. In fact, both estimators perform worse when there is greater variation in sample size. However, the shrinkage estimator is far more robust than the usual estimator in this aspect. For instance, in the case in which the true variation in log-odds ratio is small (solid lines), the correlation between the usual estimates and the true values declines from above .7 to below .5 as the variation in sample size changes from very small to very large, whereas the correlation between the shrinkage estimates and the true values stays roughly unchanged (about .71) regardless of the variation in sample size.

To sum up, this simulation study suggests that the shrinkage estimator almost always outperforms the usual estimator in joint estimation of multiple log-odds ratios, either in terms of total squared error or in terms of the correlation with the true values. Moreover, the advantage of the shrinkage estimator is more pronounced when there is less variation in the true log-odds ratio or more variation in sample size. In particular, the higher correlations with the true values exhibited by the shrinkage estimator reveal its great potential for enhancing the accuracy of cross-table comparisons.

2.5. *Shrinkage at Work: An Example*

I now apply the shrinkage method to the mobility data assembled by Hazelrigg and Garnier (1976), which provide 3×3 classifications of son's occupation by father's occupation for 16 countries in the 1960s and 1970s (henceforth referred to as HG-16). The data are displayed in Table 1. In each of the 16 tables, occupation is categorized as white collar, blue collar, or farm. Let us consider two sets of log-odds ratios that are of particular substantive interest: (1) the log-odds ratio pertaining to the 2×2 subtable of white-collar and blue-collar workers and (2) the log-odds ratio pertaining to the 2×2 subtable of blue-collar workers and farmers. We may perceive these two log-odds ratios as measuring the strengths of class boundaries between white collar and blue collar and between blue collar and farm. For each measure, I contrast the observed log-odds ratios with both the empirical Bayes estimates and the full Bayes estimates. To generate the full Bayes estimates, I ran five independent Markov chains, each containing 4,000 iterations, and retained the last 2,000 vectors from each run. The point estimates and the standard errors of the log-odds ratios were estimated respectively as the posterior means and the posterior standard deviations.

Table 1. Mobility Tables for 16 Countries, Father's Occupation by Son's Occupation

Australia			Belgium			France			Hungary		
292	170	29	497	100	12	2,085	1,047	74	479	190	14
290	608	37	300	434	7	936	2,367	57	1,029	2,615	347
81	171	175	102	101	129	592	1,255	1,587	516	3,110	3,751
Italy			Japan			Philippines			Spain		
233	75	10	465	122	21	239	110	76	7,622	2,124	379
104	291	23	159	258	20	91	292	111	3,495	9,072	597
71	212	320	285	307	333	317	527	3,098	4,597	8,173	14,833
United States			West Germany			West Malaysia			Yugoslavia		
1,650	641	34	3,634	850	270	406	235	144	61	24	7
1,618	2,692	70	1,021	1,694	306	176	369	183	37	92	13
694	1,648	644	1,068	1,310	1,927	315	578	2,311	77	148	223
Denmark			Finland			Norway			Sweden		
79	34	2	39	29	2	90	29	5	89	30	0
55	119	8	24	115	10	72	89	11	81	142	3
25	48	84	40	66	79	41	47	47	27	48	29

Note: The row and column categories are "white collar," "blue collar," and "farm" for all countries in the table.
Source: Grusky and Hauser (1983:56); see also Raftery (1995:115).

The results are shown in Table 2. On one hand, we observe that for countries with very large sample sizes, such as Spain, United States, and West Germany, both the point estimates and the standard errors are largely the same across different methods. Because within-sample precision is sufficiently high for these countries, the shrinkage factors assigned to the observed log-odds ratios are almost zero. The shrinkage estimates, therefore, closely resemble the MLE in both location and precision. On the other hand, for relatively sparse tables, such as Finland, Norway, and Sweden, both the point estimates and the standard errors are markedly changed under the shrinkage methods. However, the empirical Bayes approach and the full Bayes approach yield essentially identical point estimates, although the latter gives slightly larger standard errors as it incorporates the uncertainty of the prior variance τ^2.

Table 2. Point Estimates and Estimated Standard Errors for Two Sets of Log-odds Ratios in HG-16 under Different Estimation Methods

	LOR between White Collar and Blue Collar			LOR between Blue Collar and Farm		
	Observed	Empirical Bayes	Full Bayes	Observed	Empirical Bayes	Full Bayes
Australia	1.28	1.35	1.35	2.82	2.82	2.82
	(.12)	(.11)	(.12)	(.20)	(.19)	(.19)
Belgium	1.97	1.93	1.94	4.37	3.95	3.98
	(.13)	(.12)	(.13)	(.40)	(.34)	(.37)
France	1.62	1.62	1.62	3.96	3.91	3.90
	(.05)	(.05)	(.05)	(.14)	(.14)	(.14)
Hungary	1.86	1.85	1.85	2.21	2.21	2.21
	(.09)	(.09)	(.09)	(.06)	(.06)	(.06)
Italy	2.16	2.05	2.05	2.95	2.93	2.95
	(.18)	(.15)	(.16)	(.23)	(.22)	(.22)
Japan	1.82	1.81	1.81	2.64	2.66	2.66
	(.14)	(.13)	(.13)	(.25)	(.23)	(.23)
Philippines	1.94	1.89	1.90	2.74	2.74	2.74
	(.17)	(.14)	(.15)	(.12)	(.12)	(.12)
Spain	2.23	2.23	2.23	3.32	3.31	3.31
	(.03)	(.03)	(.03)	(.04)	(.04)	(.04)
United States	1.45	1.47	1.46	2.71	2.71	2.71
	(.06)	(.06)	(.06)	(.13)	(.13)	(.13)
West Germany	1.96	1.95	1.95	2.10	2.11	2.11
	(.05)	(.05)	(.05)	(.07)	(.07)	(.07)
West Malaysia	1.29	1.36	1.35	2.09	2.10	2.10
	(.12)	(.11)	(.12)	(.10)	(.10)	(.10)
Yugoslavia	1.84	1.80	1.80	2.37	2.45	2.43
	(.31)	(.21)	(.23)	(.31)	(.28)	(.30)
Denmark	1.61	1.67	1.67	3.26	3.14	3.13
	(.26)	(.19)	(.21)	(.41)	(.34)	(.37)
Finland	1.86	1.80	1.81	2.62	2.67	2.67
	(.33)	(.21)	(.23)	(.37)	(.32)	(.31)
Norway	1.34	1.52	1.52	2.09	2.27	2.26
	(.27)	(.19)	(.22)	(.38)	(.33)	(.33)
Sweden	1.65	1.69	1.69	3.35	3.11	3.09
	(.25)	(.19)	(.19)	(.63)	(.45)	(.49)

Note: Numbers in parentheses are estimated standard errors. HG-16 = the 16 3 × 3 mobility tables assembled by Hazelrigg and Garnier (1976); LOR = log-odds ratio.

Overall, shrinkage estimates based on either approach are more precise than the usual estimates.

To demonstrate the effects of shrinkage, I visualize the contrasts between the observed log-odds ratios and the empirical Bayes estimates in Figure 2, in which 9 of the 16 countries are marked for illustration: Belgium, France, Hungary, Italy, Spain, United States, West Malaysia, Norway, and Sweden. Panel A shows the log-odds ratio between white collar and blue collar. First, we find that most of the cross-country differences are consistent between the two sets of estimates: for example, according to either estimator, Spain and West Malaysia are respectively the least mobile (i.e., with the highest log-odds ratio) and the most mobile (i.e., with the lowest log-odds ratio) among the nine countries. However, because the observed log-odds ratios differ in sampling precision, the shrinkage estimator implies a slightly different rank order among these countries. In particular, Norway is more mobile than the United States according to the usual estimator (i.e., the observed odds ratio) but less mobile than the United States according to the shrinkage estimator. In other words, the empirical Bayes model suggests that the higher mobility of Norway exhibited by the raw data is due simply to its larger sampling variance, not because the barrier between white-collar and blue-collar jobs is more permeable in Norway than in the United States.[8]

Panel B demonstrates the effects of shrinkage for the log-odds ratio between blue collar and farm. Overall, these estimates are much higher than the estimates in panel A, indicating that the barrier between these two classes is much harder to cross than the barrier between white-collar and blue-collar jobs. Similar to panel A, the rankings among the nine countries are not much altered under the shrinkage approach, except that Norway is again “shrunk toward the mean.” We also find that the influence of shrinkage is the most pronounced for Belgium, which is markedly less mobile than France according to the observed log-odds ratio but closely resembles France in their shrinkage estimates. This is clearly related to the sparse cell of (blue collar, farm) in the Belgian table (see again Table 1).

3. ADJUSTED ESTIMATION OF THE ALTHAM INDEX

For mobility tables with more than two categories, we can use the shrinkage estimator (equation 7) to calculate summary measures of association that are based on aggregations of log-odds ratios. In this section, I construct an adjusted estimator of the Altham index, an aggregate measure of association that has been recently used for studying intergenerational occupational mobility (Ferrie 2005; Long and Ferrie 2007,

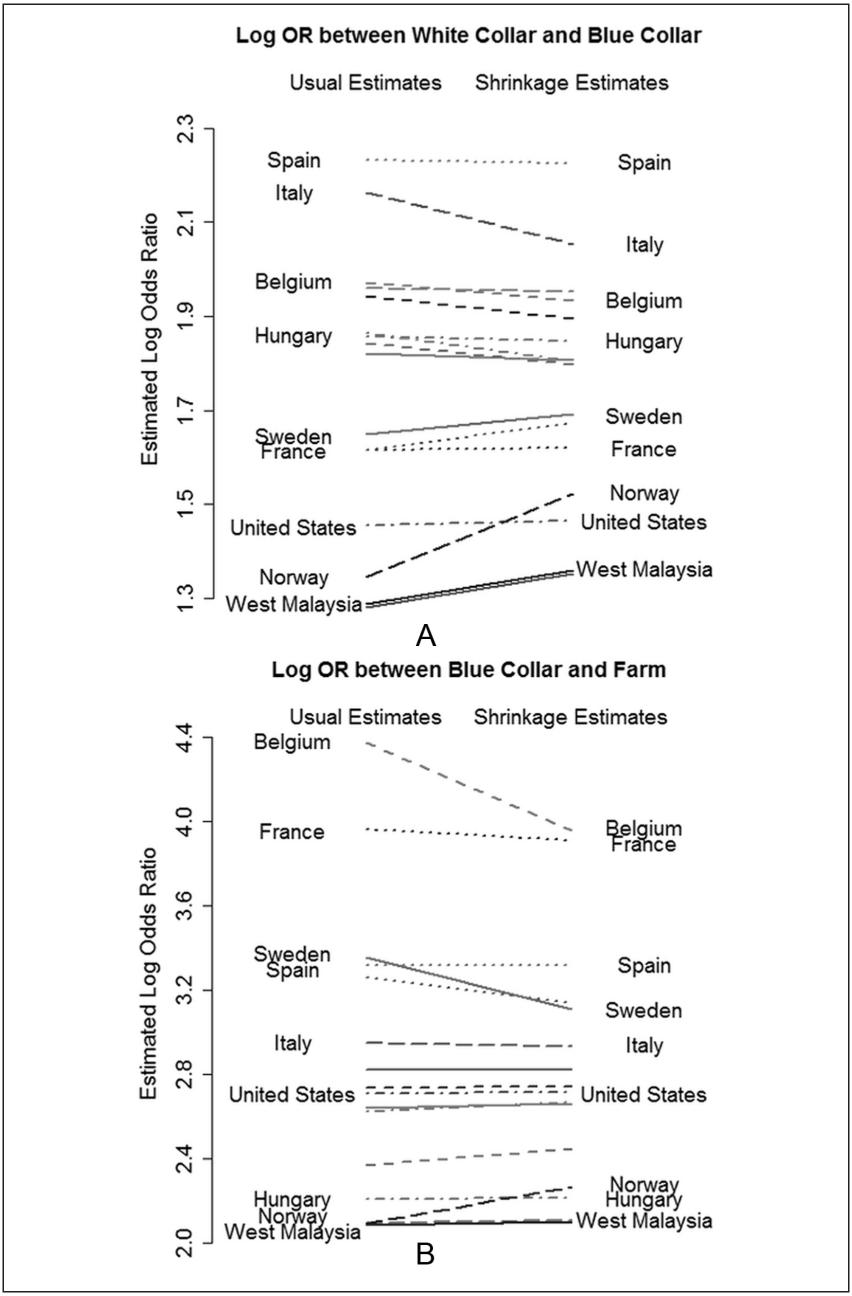

Figure 2. Usual estimates and shrinkage estimates for two sets of log-odds ratios in HG-16. HG-16 = the 16 3 × 3 mobility tables assembled by Hazelrigg and Garnier (1976); Log OR = log-odds ratio.

2013). Results from a set of calibrated simulations suggest that using shrinkage estimates of log-odds ratios can substantially improve the estimation precision of the Altham index.

3.1. *An Adjusted Estimator of the Altham Index*

To assess the total amount of association embodied in a two-way contingency table, Altham (1970) proposed a number of measures that are based on aggregations of log-odds ratios. One such measure is identical to the Euclidean distance between the full sets of log-odds ratios in two $I \times J$ tables P and Q, that is,

$$d(P,Q) = \left[\sum_{i=1}^{I}\sum_{j=1}^{J}\sum_{l=1}^{I}\sum_{m=1}^{J}\left|\log\frac{p_{ij}p_{lm}}{p_{im}p_{lj}} - \log\frac{q_{ij}q_{lm}}{q_{im}q_{lj}}\right|^{2}\right]^{1/2},$$

where p_{ij} and q_{ij} denote the probabilities associated with the cell (i,j) in table P and table Q. Although the metric $d(P,Q)$ gauges the distance between the row-column associations in tables P and Q, it does not tell us in which table the rows and the columns are more closely associated. To answer this question, we can compare $d(P,J)$ with $d(Q,J)$, where J denotes a contingency table in which the rows and columns are completely independent. Because all of the log-odds ratios are zero in an independent table, we have

$$d(P,J) = \left[\sum_{i=1}^{I}\sum_{j=1}^{J}\sum_{l=1}^{I}\sum_{m=1}^{J}\left|\log\frac{p_{ij}p_{lm}}{p_{im}p_{lj}}\right|^{2}\right]^{1/2}. \tag{12}$$

We can see that $d(P,J)$ is the square root of the sum of all squared log-odds ratios in table P. A larger value of $d(P,J)$ indicates a stronger association between the rows and columns. Hence, when P is a mobility table, a larger $d(P,J)$ corresponds to a more rigid class regime. Although this approach to comparing mobility tables is lesser known than log-linear models in comparative stratification research, it has been recently used by economic historians to study long-term trends in occupational mobility in Great Britain and the United States (Ferrie 2005; Long and Ferrie 2007, 2013). From here on, I use "the Altham index" to mean $d(P,J)$ for a contingency table P.

Now suppose we have a set of $I \times I$ mobility tables $M_1, M_2, \cdots M_k$ for K countries. For each country k, we can directly calculate the Altham index by substituting the observed log-odds ratios:

$$\hat{d}^{\text{Direct}}(M_k, J) = \left[\sum_{i=1}^{I} \sum_{j=1}^{J} \sum_{l=1}^{I} \sum_{m=1}^{J} \left| \log \frac{n_{ijk} n_{lmk}}{n_{imk} n_{ljk}} \right|^2 \right]^{1/2}, k = 1, 2, \cdots K, \tag{13}$$

where n_{ijk} denotes the observed frequency associated with the cell (i,j) in table k.[9] On the other hand, we can use the shrinkage estimator of the log-odds ratio for each row-column combination (i, j, l, m), yielding an adjusted estimator of the Altham index:

$$\hat{d}^{\text{Adjusted}}(M_k, J) = \left[\sum_{i=1}^{I} \sum_{j=1}^{J} \sum_{l=1}^{I} \sum_{m=1}^{J} \left| \hat{\theta}^{EB}_{(i,j,l,m),k} \right|^2 \right]^{1/2}, k = 1, 2, \cdots K, \tag{14}$$

where $\hat{\theta}^{EB}_{(i,j,l,m),k}$ denotes the shrinkage estimator (equation 7) of the log-odds ratio $\log \frac{p_{ij} p_{lm}}{p_{im} p_{lj}}$ in table k. Because the Altham index is not a linear function of the log-odds ratios, the adjusted estimator (equation 14) cannot be expressed as a weighted average between the direct estimator (equation 13) and a common mean as in equation (7). However, as we will see, the key effect of this adjustment is also "pulling" the direct estimates toward the middle, the extent of which depends on sample sizes of the corresponding tables.

3.2. *Direct Estimator versus Adjusted Estimator in Simulation*

Below, I use numerical simulation to evaluate the performance of the direct estimator (equation 13) and the adjusted estimator (equation 14) for the Altham index. As in the case of the log-odds ratio, I compare them in two aspects: (1) total squared error and (2) correlation with the true values. To mimic mobility regimes in the real world, I use HG-16 to motivate my simulation setup. First, I fitted the 16 3×3 mobility tables using four log-linear (or log-multiplicative) models: (1) quasi-perfect mobility, (2) uniform inheritance, (3) perfect blue-collar mobility, and (4) the Unidiff model with full row-column interaction. These models were proposed by Grusky and Hauser (1984) (a, b, c) and Xie

(1992) (d) to compare mobility regimes of the 16 countries.[10] I then treated the estimated parameters as the true parameters, yielding four data-generating models, that is, four simulation setups. For each of the four setups, I generated 1,000 independent samples of the 16 tables and, for each sample, obtained the direct and the adjusted estimates of the Altham index. With the "true" Altham indices readily available from the model parameters, I evaluated the two estimators using three criteria: (1) total squared error, (2) Pearson's correlation with the true values, and (3) Spearman's rank correlation with the true values. To smooth random fluctuations, each of the three measures was averaged over the 1,000 samples, thus producing the total mean squared error (total MSE) and the average correlation coefficients. The results are summarized in Table 3.

We first observe in this table that the adjusted estimator leads to a substantial reduction in total MSE in all of the four scenarios. For example, when data are generated from the Unidiff model, total MSE for the adjusted estimator is only about half of that for the direct estimator (38.8 / 77.0 = 50.4%). Moreover, the adjusted estimates compete well with the direct estimates in correlating with the true Altham indices. Specifically, the adjusted estimator (on average) brings an increase in Pearson's correlation in all of the four scenarios and an increase in Spearman's rank correlation in two of the four scenarios. Therefore, the shrinkage-based method for calculating the Altham index not only yields more precise individual estimates but may also enhance the accuracy of cross-table comparisons in the overall degree of association.

3.3. *An Illustration Using HG-16*

I now apply both estimators of the Altham index to the real data in HG-16. The results are shown in Figure 3A, in which the same nine countries as in section 2.5 are highlighted for illustration. Clearly, with the shrinkage estimates of log-odds ratios, the Altham index tends to be shrunk toward the middle, yet the degree of shrinkage varies considerably from country to country. For example, the adjusted estimate is very similar to the direct estimate for France, but the estimate for Sweden is heavily altered by the adjustment. According to the direct estimates, Sweden ranks as the least mobile (i.e., with the highest Altham index) among the 16 countries; but by the adjusted estimates, Sweden stands

Table 3. Direct Estimator versus Adjusted Estimator of the Altham Index in Simulation

Data-generating Model	Estimator	Total MSE	Average Correlation with $d(M_k, J)$ Pearson	Spearman's Rank
Quasi-perfect mobility	$\hat{d}^{\text{Direct}}(M_k, J)$	91.9	.916	.894
	$\hat{d}^{\text{Adjusted}}(M_k, J)$	73.5	.919	.886
Uniform inheritance	$\hat{d}^{\text{Direct}}(M_k, J)$	39.6	.904	.886
	$\hat{d}^{\text{Adjusted}}(M_k, J)$	22.3	.940	.918
Perfect blue-collar mobility	$\hat{d}^{\text{Direct}}(M_k, J)$	107.5	.894	.885
	$\hat{d}^{\text{Adjusted}}(M_k, J)$	74.0	.904	.873
Unidiff (full interaction)	$\hat{d}^{\text{Direct}}(M_k, J)$	77.0	.867	.855
	$\hat{d}^{\text{Adjusted}}(M_k, J)$	38.8	.906	.882

Note: MSE = mean squared error.

right in the middle, more mobile than Hungary, France, Belgium, Italy, and Spain. Such a sharp contrast suggests that the high (direct) estimate of the Altham index for Sweden is primarily a result of large sampling errors for some of the log-odds ratios in the Swedish data. As was shown in Table 1, the cell (white collar, farm) of the Swedish table contains no observation, which may have led to an incredibly high estimate of the Altham index.

We can also evaluate the Altham index for a subset of the mobility table. Figure 3B presents the results for the same set of tables with the farm-farm cells excluded. The uniqueness of the farm-farm cell has been emphasized by Xie and Killewald (2013), who argued that the extremely persistent degree of self-recruitment from farming among farmers (regardless of historical contexts) challenges the utility of odds ratio–based measures for comparing mobility regimes with very different levels of industrialization. Hence, calculating the Altham index without the farm-farm cell serves as a sensitivity check on the results in Figure 3A. Two findings emerge from this analysis. First, compared with panel A, we find that the exclusion of the farm-farm cell leads to significant changes in the positions of these countries along the mobility spectrum. For instance, when the full tables are analyzed, France and Hungary are fairly close to each other, both ranking among the least mobile regimes; when the farm-farm cells are excluded, France appears

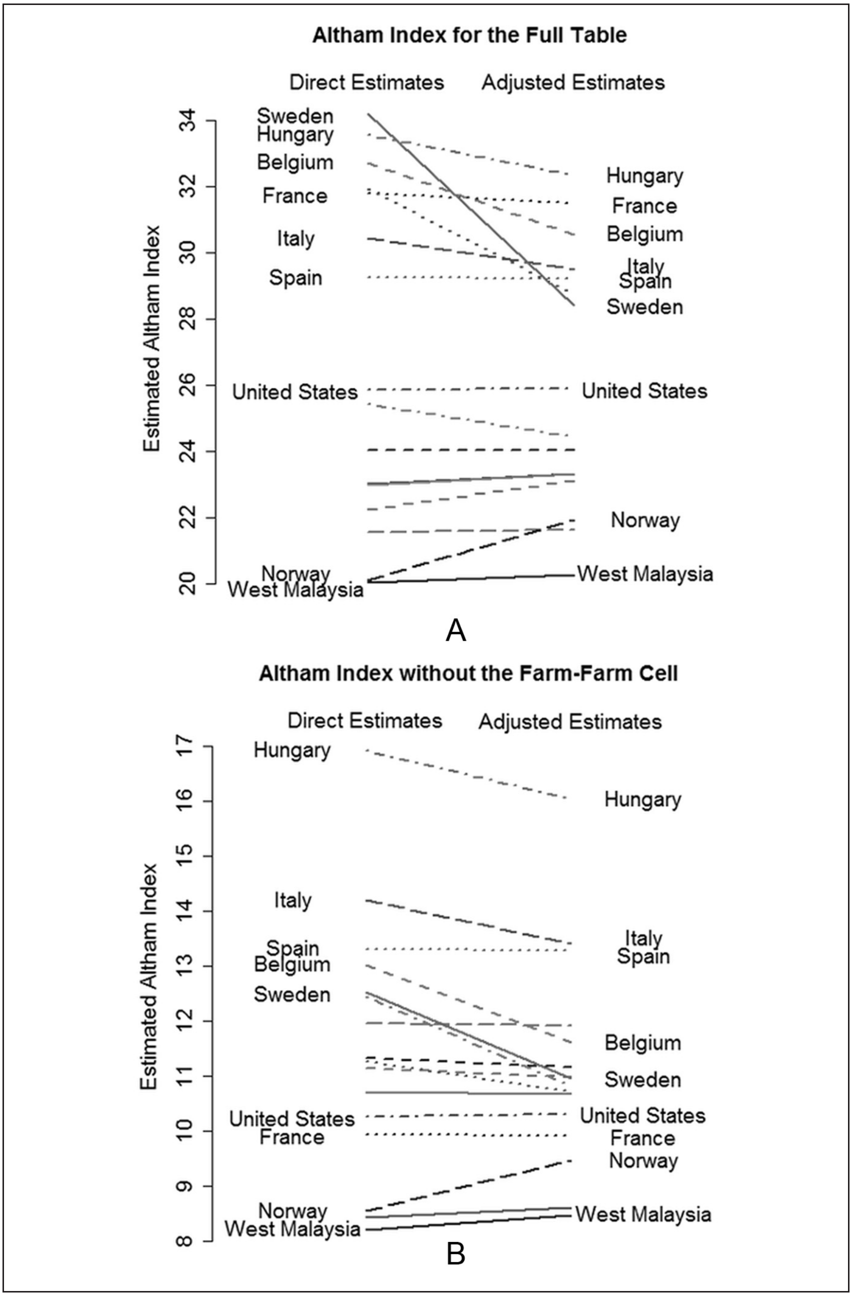

Figure 3. Direct estimates and adjusted estimates of the Altham index for HG-16 (A) and HG-16 without farm-farm cells (B). HG-16 = the 16 3 × 3 mobility tables assembled by Hazelrigg and Garnier (1976).

to be one of the most mobile countries, whereas Hungary stands out as the single most immobile regime, with an Altham index far higher than those of the others. Second, although the adjusted estimates have the same rank order as the direct estimates for the nine countries marked here, they differ substantially in relative positions. For example, according to the direct estimates (without the farm-farm cell), Norway and Sweden are far apart, one very close to West Malaysia and the other only slightly lower than Spain; however, with the shrinkage-based adjustment, these two Nordic countries are much more similar, with their Altham indices closer to France and the United States than to West Malaysia and Spain.

4. SHRINKING TOWARD CONVERGENCE: COMPARING THE ALTHAM INDEX WITH THE UNIDIFF MODEL

Although the Altham index provides a simple summary measure of the row-column association for a mobility table, log-linear modeling has been far more popular among sociological studies on intergenerational class mobility, in part because of its flexibility for accommodating fine-grained theoretical hypotheses (e.g., Erikson and Goldthorpe 1987; Hout 1984, 1988; Yamaguchi 1987). Among a plethora of log-linear and log-multiplicative models that have been proposed for studying mobility tables, the Unidiff model (also known as the log-multiplicative layer effect model) is particularly recognized for its ability to provide a single parameter that captures cross-table differences in social fluidity (Erikson and Goldthorpe 1992; Xie 1992). Hence, the Altham index and the Unidiff model constitute two different approaches to making overall comparisons between mobility tables. In this section, I first establish a theoretical equivalence between these two approaches in the ideal case in which the Unidiff model is the true data-generating model. Then, using two real data sets, I show that the adjusted estimates of the Altham index agree more closely with the layer effects estimated under the Unidiff model than do direct estimates of the Altham index.

4.1. *The Unidiff Model, the Layer Effect, and the Altham Index*

As in section 3.1, let us consider a set of $I \times I$ mobility tables $M_1, M_2, \ldots M_k$ for K countries. In a log-linear analysis, these tables are typically treated as a three-way table with I rows, I columns, and K layers. Denoting by F_{ijk} the expected frequency in the ith row, the jth column, and the kth layer (i.e., the kth country), the saturated log-linear model can be written as

$$\log F_{ijk} = \mu + \mu_i^R + \mu_j^C + \mu_k^L + \mu_{ik}^{RL} + \mu_{jk}^{CL} + \mu_{ij}^{RC} + \mu_{ijk}^{RCL}.$$

In this equation, the first six terms are used to saturate the marginal distributions of both the origin and the destination in each country, while the last two terms, μ_{ij}^{RC} and μ_{ijk}^{RCL}, capture variations in the origin-destination association across countries. However, because the saturated model exhausts all degrees of freedom, it would severely overfit the data. In practice, the researcher often wants to specify μ_{ij}^{RC} and μ_{ijk}^{RCL} in a parsimonious fashion. The Unidiff model, in particular, assumes that these countries share a common pattern of association between origin and destination while allowing the strength of association to vary across countries. As a result, the model can be written as

$$\log F_{ijk} = \mu + \mu_i^R + \mu_j^C + \mu_k^L + \mu_{ik}^{RL} + \mu_{jk}^{CL} + \psi_{ij}\phi_k. \tag{15}$$

Here, the parameter ψ_{ij} characterizes the common pattern of association, and the parameter ϕ_k, which is called the "layer effect," identifies the relative position of country k along the mobility spectrum.

According to equation (15), the expected log-odds ratio associated with the row-column combination (i, j, l, m) in table k can be calculated as

$$\theta_{(i,j,l,m),k} = \log F_{ijk} - \log F_{imk} - \log F_{ljk} + \log F_{lmk} = \theta^*_{i,j,l,m}\phi_k, \tag{16}$$

where $\theta^*_{i,j,l,m} = \psi_{ij} - \psi_{im} - \psi_{lj} + \psi_{lm}$. Therefore, under the Unidiff model, any log-odds ratio in a given table is the product of a common log-odds ratio $\theta^*_{i,j,l,m}$ and the layer effect ϕ_k. Clearly, a greater value of ϕ_k implies a lower degree of social fluidity. Substituting the above expression into equation (12), the Altham index becomes

$$d(M_k, J) = \left[\sum\nolimits_{i,j,l,m} \left|\theta_{(i,j,l,m),k}\right|^2\right]^{1/2} = \left[\sum\nolimits_{i,j,l,m} \left|\theta^*_{i,j,l,m}\right|^2\right]^{1/2} \phi_k. \tag{17}$$

Because the term $\left[\sum\nolimits_{i,j,l,m} \left|\theta^*_{i,j,l,m}\right|^2\right]^{1/2}$ does not depend on k, the Altham index $d(M_k, J)$ is directly proportional to the layer effect ϕ_k. In other words, these two measures of association are equivalent under the Unidiff model.

Real mobility data, however, may fail to support the assumptions of the Unidiff model. For example, according to the likelihood ratio test, the Unidiff model fits poorly for HG-16 (Xie 1992:390). In such cases, we may conclude that different mobility regimes exhibit different patterns of relative mobility, and proceed to develop more flexible models, such as the regression-type models proposed by Goodman and Hout (1998), to capture the nuances of cross-table differences. Nonetheless, tempted by such questions as "Overall, is country A more mobile than country B?" the researcher may still be interested in reducing subtle, multidimensional differences to simple, one-dimensional contrasts. In this regard, the Unidiff model and the Altham index constitute two reasonable yet distinct approaches. A natural question, then, is whether these two approaches would yield concordant results. Because the layer effect and the Altham index are equivalent when the Unidiff model is true, we would expect that they produce more similar results when data are more congruent with the Unidiff model. On the other hand, given the advantages of the adjusted estimator over the direct estimator for the Altham index, it is reasonable to conjecture that the adjusted estimator agrees more closely than the direct estimator with results from the Unidiff model. Below, I use two sets of real mobility tables to test these two hypotheses.

4.2. *Shrinking toward Convergence: Evidence from Two Data Sets*

I apply both estimators of the Altham index, along with the Unidiff model, to two data sets: (1) HG-16 (i.e., the 16 3×3 mobility tables assembled by Hazelrigg and Garnier [1976]) and (2) a collection of 149 6×6 mobility tables from 35 countries assembled by Ganzeboom, Luijkx, and Treiman (1989), henceforth GLT-149. Whereas occupation in HG-16 is crudely classified as white collar, blue collar, and farm, GLT-149 adopts the six-category version of the EGP class scheme (Erikson, Goldthorpe, and Portocarero 1979): the service class (I + II), routine nonmanual workers (III), petty bourgeoisie (IVa + IVb), farmers and agricultural laborers (IVc + VIIb), skilled manual workers (V + VI), and unskilled manual workers (VIIa).

Table 4. Correlations of Direct and Adjusted Estimates of the Altham Index with $\hat{\phi}_k^{\text{Unidiff}}$

Data Set	Estimator	Pearson's Correlation	Spearman's Rank Correlation
HG-16	$\hat{d}^{\text{Direct}}(M_k, J)$	.858	.832
	$\hat{d}^{\text{Adjusted}}(M_k, J)$	.852	.876
HG-15 (without Hungary)	$\hat{d}^{\text{Direct}}(M_k, J)$	.917	.846
	$\hat{d}^{\text{Adjusted}}(M_k, J)$	.939	.893
GLT-149	$\hat{d}^{\text{Direct}}(M_k, J)$	.817	.839
	$\hat{d}^{\text{Adjusted}}(M_k, J)$	.803	.899

To assess the extent to which different estimators of the Altham index accord with the layer effects estimated under the Unidiff model, I use Spearman's rank correlation as well as the Pearson correlation. Previous researchers analyzing HG-16 have pointed out that Hungary significantly deviates from the other 15 countries in patterns of interclass mobility (Grusky and Hauser 1984; Xie 1992). For this reason, I analyzed both the full set of HG-16 and the 15 tables excluding the Hungarian case (henceforth referred to as HG-15). The results are shown in Table 4. We can see that for all three data sets, the fitted layer effects $\hat{\phi}_k^{\text{Unidiff}}$ tend to correlate more strongly with the adjusted estimates of the Altham index than with the direct estimates, especially by Spearman's rank correlation. For example, the rank correlation for GLT-149 is .839 between $\hat{d}^{\text{Direct}}(M_k, J)$ and $\hat{\phi}_k^{\text{Unidiff}}$ but .899 between $\hat{d}^{\text{Adjusted}}(M_k, J)$ and $\hat{\phi}_k^{\text{Unidiff}}$.

On the other hand, we notice that when Hungary is excluded from HG-16, both estimates of the Altham index become more aligned with the fitted layer effects. The Pearson correlation, for example, increases from .858 to .917 between $\hat{d}^{\text{Direct}}(M_k, J)$ and $\hat{\phi}_k^{\text{Unidiff}}$ and from .852 to .939 between $\hat{d}^{\text{Adjusted}}(M_k, J)$ and $\hat{\phi}_k^{\text{Unidiff}}$. These results accord well with our first hypothesis: because Hungary contributes the lion's share to the model deviance (i.e., G^2), its exclusion considerably improves the fit between the data and the Unidiff model, thereby producing greater consistency between model-free (i.e., the Altham index) and model-based (i.e., the Unidiff model) inferences. To explore this relationship further, I examine how the above correlations change as the most poorly fitted cases are progressively excluded from the data sets. Specifically, for

HG-16, I performed a stepwise elimination of Hungary, France, West Germany, the United States, and Spain—in order of decreasing G^2 under the Unidiff model—and recalculated the correlations for each subset of the 16 tables. For GLT-149, the same procedures were followed except that five tables, rather than one table, were removed at a time and the correlation coefficients were recalculated until 40 tables were deleted.

Figure 4 shows the results, with panel A for HG-16 and panel B for GLT-149. In both panels, I represent Pearson's correlation in solid lines and Spearman's rank correlation in dashed lines. Meanwhile, squares and triangles denote direct and adjusted estimates of the Altham index, respectively. From the four contrasts between squares and triangles, we notice that the adjusted estimates of the Altham index almost always correlate more strongly with the fitted layer effects than do the direct estimates. On the other hand, reading along the *x*-axis, we find that the correlation coefficients generally increase as the most poorly fitted cases are excluded from the data sets. The upward drift, however, is more noticeable for the adjusted estimator than for the direct estimator. As a result, the gap between $\hat{d}^{\text{Direct}}(M_k, J)$ and $\hat{d}^{\text{Adjusted}}(M_k, J)$ in their correlations with $\hat{\phi}_k^{\text{Unidiff}}$ grows larger as data align more closely with the Unidiff model. For example, when the full set of GLT-149 is analyzed, the Pearson correlation between $\hat{d}^{\text{Adjusted}}(M_k, J)$ and $\hat{\phi}_k^{\text{Unidiff}}$ is .803, slightly lower than that between $\hat{d}^{\text{Direct}}(M_k, J)$ and $\hat{\phi}_k^{\text{Unidiff}}$ (.817; see again Table 4); but when the 40 tables with the largest deviances are excluded, the adjusted estimates of Altham indices correlate much more strongly with the fitted layer effects than do the direct estimates.

In short, these results suggest that in assessing the overall degree of social fluidity, the adjusted estimator of the Altham index accords more closely with the Unidiff model than does the direct estimator. Moreover, the contrast becomes more pronounced when data are more congruent with the Unidiff model. How do we understand these findings? First, we note that the adjusted estimator of the Altham index differs from the direct estimator only in its reliance on shrinkage estimates of the log-odds ratios. As mentioned earlier, the underlying principle of the shrinkage method is to borrow information from other cases, particularly through an empirical Bayes model with a normal prior. The adjusted estimator of the Altham index, therefore, may be considered as a semi-parametric method because it uses a normal Bayes model to smooth data across multiple tables but imposes no parametric constraints on the pattern of association within tables. In contrast, the direct estimator of the

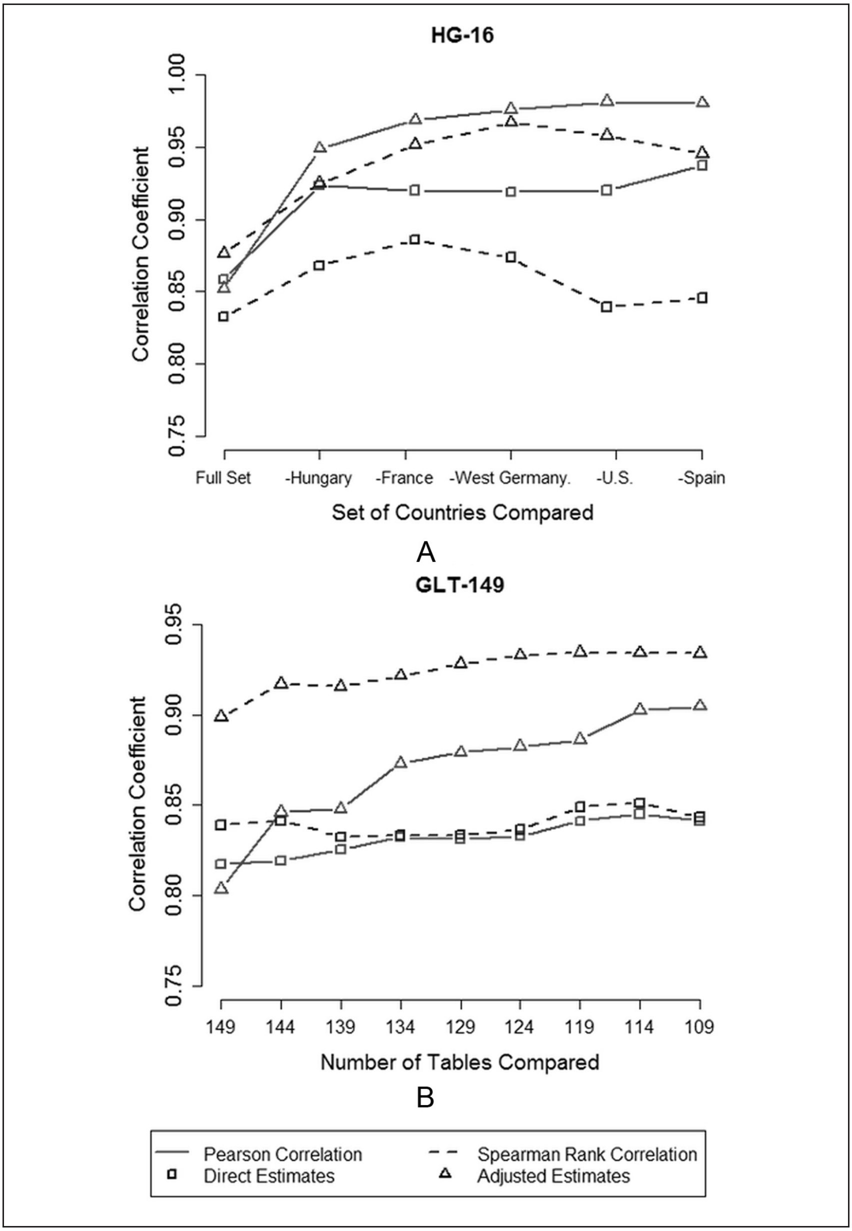

Figure 4. Direct estimates versus adjusted estimates of the Altham index in their correlations with $\hat{\phi}_k^{\text{Unidiff}}$ for varying subsets of HG-16 and GLT-149.
HG-16 = the 16 3 × 3 mobility tables assembled by Hazelrigg and Garnier (1976); GLT-149 = a collection of 149 6 × 6 mobility tables from 35 countries assembled by Ganzeboom, Luijkx, and Treiman (1989).

Altham index is fully nonparametric, involving no data smoothing either across or within tables. On the other hand, the Unidiff model stipulates that all log-odds ratios are determined as a product of a common pattern of association and table-specific effects. This multiplicative specification requires the Unidiff model to pool data both across tables (for estimating ψ_{ij}) and across cells within tables (for estimating ϕ_k). Hence, in the way that data are pooled to draw inferences, the adjusted estimator of the Altham index stands closer than the direct estimator to the Unidiff model, which probably explains why the shrinkage approach boosts convergence between a descriptive index and a parametric model in gauging social fluidity.

5. SUMMARY AND DISCUSSION

Building on an empirical Bayes framework, I have proposed a shrinkage estimator of the log-odds ratio for comparing mobility tables. This estimator enhances estimation precision by borrowing information across multiple tables while placing no restrictions on the pattern of association within tables. This approach stands in stark contrast to the usual MLE of the log-odds ratio, which involves no data pooling either across or within tables. Numerical simulation suggests that the shrinkage estimator outperforms the usual MLE in both the total squared error and the correlation with the true values. Moreover, the benefits of the shrinkage method are greater when there is less variation among the true log-odds ratios or more variation in sampling precision.

Furthermore, the shrinkage estimator of the log-odds ratio can be used to calculate the Altham index, an aggregate measure of association that has been recently adopted in comparative mobility research. Results from a set of calibrated simulations suggest that the adjusted estimator can substantially improve estimation precision while maintaining high correlations with the true values. Finally, using two real data sets, we find that the adjusted estimator of the Altham index accords more closely with the Unidiff model than does the direct estimator of the Altham index. This finding, as I have discussed, stems from the fact that both the Unidiff model and the shrinkage approach enforce information sharing across tables, albeit via apparently different mechanisms.

The shrinkage estimator (equation 7) derives from a Bayes model in which a common prior, that is, equation (4), is assumed for all cases. This assumption can easily be relaxed to incorporate our prior knowledge about the similarities and differences between mobility regimes. In particular, we can extend the prior distribution (equation 4) to

$$\theta_k \overset{indep}{\sim} N(\alpha + \beta^T X_k, \tau^2),$$

where X_k denotes a group of exogenous variables posited to affect the true log-odds ratio. The empirical Bayes estimator (equation 7) then becomes

$$\hat{\theta}_k^{EB} = \hat{\alpha} + \hat{\beta}^T X_k + [1 - \frac{(K-R-3)\widehat{\sigma_k^2}}{(K-R-1)\left(\widehat{\tau^2} + \widehat{\sigma_k^2}\right)}](Y_k - \hat{\alpha} - \hat{\beta}^T X_k),$$

where $\hat{\alpha}$ and $\hat{\beta}$ denote estimates of α and β, and R represents the dimension of X_k.[11] In this formulation, the usual estimate Y_k is shrunk not toward a common mean but toward the conditional mean $\hat{\alpha} + \hat{\beta}^T X_k$. For example, if we assume that economic development promotes social mobility, as the "thesis of industrialism" suggests (Treiman 1970), X_k could be a measure of the level of industrialization in country k. In this case, the shrinkage estimator borrows information not uniformly from all countries but mainly from countries at similar levels of industrialization. Note that if the number of tables K far exceeds the number of predictors R, the adjustment factor $\frac{K-R-3}{K-R-1}$ will be close to one and the empirical Bayes estimates can be approximated by EBLUPs from mixed-effects meta-analysis of log-odds ratios (see Viechtbauer [2010] for a guide to implementation).

For evaluating the overall degree of social fluidity, the Unidiff model and the Altham index constitute two valid yet distinctive approaches. The Unidiff model stipulates that all log-odds ratios are determined multiplicatively by a common pattern of association and layer-specific effects. This is a flexible but nontrivial assumption. Not only does it require that different log-odds ratios within a table are of the same relative magnitudes in all mobility regimes, but it also means that the rank order among mobility regimes does not depend on which log-odds ratio is being examined. For example, a Unidiff model for HG-16 would imply that the two sets of log-odds ratios in Figure 2 exhibit the same relative positions in the two panels, which is obviously at odds with the data. The Unidiff model, therefore, may incur a model specification bias if the true mobility regimes being compared do not comport with the "common-pattern" assumption. In contrast, the Altham index is fully nonparametric, thus being exempt from any type of model specification bias. For the same reason, however, direct calculation of the Altham index is susceptible to large sampling errors, especially for sparse tables. The shrinkage approach presented in this paper—which exploits a parametric Bayes model to "borrow strength" across tables but remains model-free within tables—serves as an eclectic formula for comparing mobility regimes, striking a balance between sampling variance and model specification bias. Clearly, this approach is applicable not only to comparative mobility analysis but to any area of research that calls for comparisons of multiple two-way contingency tables.

Acknowledgments

I thank Harry Ganzeboom for making the GLT-149 data available. I am grateful to Yu Xie, Hongwei Xu, and two anonymous reviewers for their helpful comments. The ideas expressed herein are those of the author.

Notes

1. The inadequacy of mobility ratio as a measure of association has been discussed by Blau and Duncan (1967:93–97), Tyree (1973), Hauser (1980:426–30), and Hout (1983:17–18).
2. Typically, the same occupational classification is used for origin and destination, that is, fathers and sons. In this case, $I = J$.
3. This observation derives from the fact that an estimated variance of the sample log-odds ratio can be expressed as $1/n_{11} + 1/n_{12} + 1/n_{21} + 1/n_{22}$(Agresti 2002:71). See also section 2.
4. The same conclusion holds when the sampling distribution is Poisson or product multinomial; see Powers and Xie (2008:79–80).
5. Clogg and Eliason (1987) noted that the practice of adding constants to all cells tends to shrink the data toward equiprobability. As we will see, this problem will be less relevant for the shrinkage estimator because the modified sample estimate Y_k is unlikely to receive much weight when there are zero cells.
6. For convenience, the agricultural sector is omitted in this simulation study.
7. The parameters for the row marginal distribution, the average transition probabilities, and the ranges of the transition probabilities are all chosen on the basis of the empirical mobility tables for 16 countries collected by Hazelrigg and Garnier (1976).
8. If we calculate the z score for the difference in observed log-odds ratio between Norway and the United States, we will find that it is not statistically significant.
9. As before, when any of the four cells are zero, one half is added to all of the four cells before calculation.

10. Models 1, 2, and 3 correspond respectively to models A2, A3, and A4 in Grusky and Hauser (1984:389); model 4 corresponds to model FI_x in Xie (1992:390).
11. The adjustment factor changes from $\frac{K-3}{K-1}$ to $\frac{K-R-3}{K-R-1}$ because R additional degrees of freedom are used to estimate the hyperparameters; see Morris (1983) for a more technical discussion.

References

Agresti, Alan. 2002. *Categorical Data Analysis*. Hoboken, NJ: John Wiley.

Altham, Patricia M. E.1970. "The Measurement of Association of Rows and Columns for an r×s Contingency Table."*Journal of the Royal Statistical Society, Series B* 32(1):63–73.

Altham, Patricia M. E., and Joseph P. Ferrie. 2007. "Comparing Contingency Tables Tools for Analyzing Data from Two Groups Cross-classified by Two Characteristics." *Historical Methods* 40(1):3–16.

Blau, Peter M., and Otis Dudley Duncan. 1967. *The American Occupational Structure*. New York: John Wiley.

Carter, Grace M., and John E. Rolph. 1974. "Empirical Bayes Methods Applied to Estimating Fire Alarm Probabilities." *Journal of the American Statistical Association* 69(348):880–85.

Clogg, Clifford C., and Scott R. Eliason. 1987. "Some Common Problems in Loglinear Analysis." *Sociological Methods and Research* 16(1):8–44.

Duncan, Otis Dudley. 1979. "How Destination Depends on Origin in the Occupational Mobility Table." *American Journal of Sociology* 84:793–803.

Efron, Bradley, and Carl Morris. 1973. "Stein's Estimation Rule and Its Competitors—An Empirical Bayes Approach." *Journal of the American Statistical Association* 68(341):117–30.

Efron, Bradley, and Carl Morris. 1975. "Data Analysis Using Stein's Estimator and Its Generalizations." *Journal of the American Statistical Association* 70(350):311–19.

Erikson, Robert, John H. Goldthorpe, and Lucienne Portocarero. 1979. "Intergenerational Class Mobility in Three Western European Societies: England, France and Sweden." *British journal of Sociology* 30(4): 415 -41.

Erikson, Robert, and John H. Goldthorpe. 1987. "Commonality and Variation in Social Fluidity in Industrial Nations. Part I: A Model for Evaluating the 'FJH Hypothesis.'" *European Sociological Review* 3(1):54–77.

Erikson, Robert, and John H. Goldthorpe. 1992. *The Constant Flux: A Study of Class Mobility in Industrial Societies*. New York: Oxford University Press.

Fay, Robert E., III, and Roger A. Herriot. 1979. "Estimates of Income for Small Places: An Application of James-Stein Procedures to Census Data." *Journal of the American Statistical Association* 74(366):269–77.

Featherman, David L., F. Lancaster Jones, and Robert M. Hauser. 1975. "Assumptions of Social Mobility Research in the US: The Case of Occupational Status." *Social Science Research* 4(4):329–60.

Ferrie, Joseph P. 2005. "History Lessons: The End of American Exceptionalism? Mobility in the United States since 1850." *Journal of Economic Perspectives* 19(3):199–215.

Ganzeboom, Harry B. G., Ruud Luijkx, and Donald J. Treiman. 1989. "Intergenerational Class Mobility in Comparative Perspective." *Research in Social Stratification and Mobility* 8:3–84.

Gart, John J., and James R. Zweifel. 1967. "On the Bias of Various Estimators of the Logit and Its Variance with Application to Quantal Bioassay." *Biometrika* 54(1–2): 181–87.

Glass, David Victor, and Jerzy Berent. 1954. *Social Mobility in Britain*. London: Routledge Kegan Paul.

Goodman, Leo A. 1968. "The Analysis of Cross-classified Data: Independence, Quasi-independence, and Interactions in Contingency Tables with or without Missing Entries." *Journal of the American Statistical Association* 63(324):1091–131.

Goodman, Leo A. 1969. "How to Ransack Social Mobility Tables and Other Kinds of Cross-classification Tables." *American Journal of Sociology* 75(1):1–40.

Goodman, Leo A. 1979. "Simple Models for the Analysis of Association in Cross-classifications Having Ordered Categories." *Journal of the American Statistical Association* 74(367):537–52.

Goodman, Leo A., and Michael Hout. 1998. "Statistical Methods and Graphical Displays for Analyzing How the Association between Two Qualitative Variables Differs among Countries, among Groups, or over Time: A Modified Regression-type Approach." Pp. 175–230 in *Sociological Methodology*, Vol. 28, edited by Adrian E. Raftery. Boston, MA: Blackwell.

Grusky, David B., and Robert M. Hauser. 1983. "Comparative Social Mobility Revisited: Models of Convergence and Divergence in Sixteen Countries." CDE Working Paper 83-6. Madison: Center for Demography and Ecology, University of Wisconsin.

Grusky, David B., and Robert M. Hauser. 1984. "Comparative Social Mobility Revisited: Models of Convergence and Divergence in 16 Countries." *American Sociological Review* 49(1):19–38.

Haldane, J.B.S. 1956. "The Estimation and Significance of the Logarithm of a Ratio of Frequencies." *Annals of Human Genetics* 20(4):309–11.

Harville, David. 1976. "Extension of the Gauss-Markov Theorem to Include the Estimation of Random Effects." *Annals of Statistics* 4(2):384–95.

Hauser, Robert M. 1980. "Some Exploratory Methods for Modeling Mobility Tables and Other Cross-classified Data." Pp. 413–58 in *Sociological Methodology*, Vol. 11, edited by Karl F. Schuessler. San Francisco, CA: Jossey-Bass.

Hazelrigg, Lawrence E., and Maurice A. Garnier. 1976. "Occupational Mobility in Industrial Societies: A Comparative Analysis of Differential Access to Occupational Ranks in Seventeen Countries." *American Sociological Review* 41:498–511.

Hout, Michael. 1983. *Mobility Tables*. Thousand Oaks, CA: Sage.

Hout, Michael. 1984. "Status, Autonomy, and Training in Occupational Mobility." *American Journal of Sociology* 89(6):1379–409.

Hout, Michael. 1988. "More Universalism, Less Structural Mobility: The American Occupational Structure in the 1980s." *American Journal of Sociology* 93(6): 1358–400.

Hout, Michael, and Avery M. Guest. 2013. "Intergenerational Occupational Mobility in Great Britain and the United States Since 1850: Comment." *American Economic Review* 103(5):2021–40.

James, William, and Charles Stein. 1961. "Estimation with Quadratic Loss." Pp. 361–79 in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1. Berkeley, CA: University of California Press.
Lipset, Seymour Martin, and Hans L. Zetterberg. 1956. *A Theory of Social Mobility*. New York: Columbia University Press.
Long, Jason, and Joseph Ferrie. 2007. "The Path to Convergence: Intergenerational Occupational Mobility in Britain and the U.S. in Three Eras." *Economic Journal* 117(519):C61–C71.
Long, Jason, and Joseph Ferrie. 2013. "Intergenerational Occupational Mobility in Great Britain and the United States Since 1850." *American Economic Review* 103(5): 2041–49.
Miller, Seymour Michael. 1960. "Comparative Social Mobility." *Current Sociology* 9(1):1–61.
Morris, Carl N. 1983. "Parametric Empirical Bayes Inference: Theory and Applications." *Journal of the American Statistical Association* 78(381):47–55.
Powers, Daniel A., and Yu Xie. 2008. *Statistical Methods for Categorical Data Analysis*. Bingley, UK: Emerald Group.
Raftery, Adrian E. 1995. "Bayesian Model Selection in Social Research." Pp. 111–64 in *Sociological Methodology*, Vol. 25, edited by Peter V. Marsden. Cambridge, MA: Blackwell.
Raudenbush, Stephen W., and Anthony S. Bryk. 1985. "Empirical Bayes Meta-analysis." *Journal of Educational and Behavioral Statistics* 10(2):75–98.
Raudenbush, Stephen W., and Anthony S. Bryk. 2002. *Hierarchical Linear Models: Applications and Data Analysis Methods*. 2nd ed.Thousand Oaks, CA: Sage.
Robinson, George K. 1991. "That BLUP Is a Good Thing: The Estimation of Random Effects." *Statistical Science* 6(1):15–32.
Rogoff, Natalie. 1953. *Recent Trends in Occupational Mobility*. New York: Arno.
Savitz, Natalya Verbitsky, and Stephen W. Raudenbush. 2009. "Exploiting Spatial Dependence to Improve Measurement of Neighborhood Social Processes." Pp. 151–83 in *Sociological Methodology*, Vol. 39, edited by Yu Xie. Boston: Wiley-Blackwell.
Sidik, Kurex, and Jeffrey N. Jonkman. 2005. "Simple Heterogeneity Variance Estimation for Meta-analysis." *Journal of the Royal Statistical Society, Series C (Applied Statistics)* 54(2):367–84.
Stein, Charles. 1956. "Inadmissibility of the Usual Estimator for the Mean of a Multivariate Normal Distribution." Pp. 197–206 in *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1. Berkeley, CA: University of California Press.
Treiman, Donald J. 1970. "Industrialization and Social Stratification." *Sociological Inquiry* 40(2):207–34.
Tyree, Andrea. 1973. "Mobility Ratios and Association in Mobility Tables." *Population Studies* 27(3):577–88.
Viechtbauer, Wolfgang. 2010. "Conducting Meta-analyses in R with the Metaphor Package." *Journal of Statistical Software* 36(3):1–48.
Xie, Yu. 1992. "The Log-multiplicative Layer Effect Model for Comparing Mobility Tables." *American Sociological Review* 57(3):380–95.

Xie, Yu,, and Alexandra Killewald. 2013. "Intergenerational Occupational Mobility in Britain and the U.S. Since 1850: Comment." *American Economic Review* 103(5): 2003–20.

Yamaguchi, Kazuo. 1987. "Models for Comparing Mobility Tables: Toward Parsimony and Substance." *American Sociological Review* 52(4):482–94.

Author Biography

Xiang Zhou is a doctoral candidate in sociology and statistics at the University of Michigan. His research interests include social inequality in China and quantitative methods. His recent publications appear in *Social Forces, Sociological Methodology, Sociological Methods and Research*, and the *Proceedings of the National Academy of Sciences*.