

PROFITING FROM CORRELATIONS: ADJUSTED ESTIMATORS FOR CATEGORICAL DATA

TOBIAS NIEBUHR AND MATHIAS TRABS

ABSTRACT. To take sample biases and skewness in the observations into account, practitioners frequently weight their observations according to some marginal distribution. The present paper demonstrates that such weighting can indeed improve the estimation. Studying contingency tables, estimators for marginal distributions are proposed under the assumption that another marginal is known. It is shown that the weighted estimators have a strictly smaller asymptotic variance whenever the two marginals are correlated. The finite sample performance is illustrated in a simulation study. As an application to traffic accident data the method allows for correcting a well-known bias in the observed injury severity distribution.

1. INTRODUCTION

Categorical data analysis is fundamental for many research fields and applications. When applying strategies to real-world data, it seems to be the rule rather than the exception that datasets are lacking to some degree. This is usually due to under-reportings of subgroups of the population, to non-representative study participants, or miscoded observations (see e.g. Lu et al. (2000), Hazell and Shakir (2006), Yamayoto et al. (2008) or Wang et al. (2015), and the applications therein). All these scenarios have in common that they result in a somehow skewed observation distribution. A major task in applied studies is to address the data's skewness at first before proceeding with any other investigation.

The most common way to correct for such a skewness is the application of weighting factors. However, a weighted or cloned sample does not give any additional information on the underlying random mechanisms, a fact that is often ignored by practitioners. Statistical methods have to incorporate the weighting step instead of treating the data clones as new and independent observations. This cloning of dependencies makes the use of weighted samples rather challenging.

In the univariate case, it can be easily seen that weighting or cloning of data can only increase the (asymptotic) variance of distribution estimators compared to an unweighted approach. In higher dimensions the situation turns out to be different. We will consider a two-dimensional discrete distribution from which one marginal is assumed to be known or can be estimated with higher accuracy. Such additional information in one category is reasonable in many applications, for instance, in our leading example from traffic accident research, where some properties of the underlying population are detailed reported in large databases and statistical surveys, while data on other categories are collected only in specific (and smaller) studies. The intriguing question is whether the standard unweighted empirical probability estimators can be improved using an additional amount of information.

We will settle this question to the positive. The constructed estimators are asymptotically unbiased and adapt to the known marginal distribution as desired in practice. They also satisfy a central limit theorem with a substantially smaller asymptotic covariance matrix (in the sense of positive semi-definite matrices) than the classical estimators. The information theoretic gain is linked to the dependence between the two marginals.

Date: November 5, 2018.

Key words and phrases. contingency table; categorical data; empirical frequencies; traffic accident data.

Our weighted estimation approach is related to the classical problem of ranked contingency tables where a two-way table should be adjusted to the two fixed marginals. Following Deming and Stephan (1940), there is a series of papers on this topic, for instance, Ireland and Kullback (1968) and Little and Wu (1991). However, we have a different target since we aim for the estimation of an unknown marginal instead of the estimation of the contingency table itself.

In addition to our asymptotic analysis, a simulation study reveals a good finite sample performance of the proposed method. Already for small sample sizes considerable improvements for dependent marginals are reported. As a little price to pay, we observe a small finite sample bias.

Our motivating application comes from traffic accident research. The official national statistics, as, e.g., the National Highway Traffic Safety Administration for the United States or the Statistisches Bundesamt for Germany, provide highly accurate information about the accident situation. However, the administrations usually provide only very low dimensional data. To evaluate new driver assistance systems, the automobile producers require further information on more than one accident characteristic. Therefore, they retrieve accident studies of more depth, e.g. the National Automotive Sampling System or the German In-Depth Accident Study (GIDAS). Based on GIDAS data, we estimate the (discretized) distribution of the speed reduction due to collision where the estimators are adjusted for the injury severity distribution from the national statistic. As a result our method allows for correcting a well-known bias in the observed injury severity distribution.

The paper is organized as follows: In Section 2 we briefly discuss the one-dimensional case while the two-dimensional analysis is presented in Section 3. Section 4 contains the simulation study. The real data example is investigated in Section 5. Section 6 concludes the paper. All technical proofs are deferred to Section 7.

2. ONE-DIMENSIONAL DISTRIBUTIONS

Before we come to our main results in the next section, we recall some basic facts for the simple univariate case. We observe an independent and identically distributed (i.i.d.) sample $X_1, \dots, X_n \in \{1, \dots, I\}$ stemming from a discrete probability distribution $\mathbb{P} = (p_1, \dots, p_I)$ with $\sum_{i=1}^I p_i = 1$ and

$$P(X_t = i) = p_i \quad \text{for all } t = 1, \dots, n \text{ and } i = 1, \dots, I.$$

Here and throughout the $I \in \mathbb{N}$ possible values, or categories, are labeled by $1, \dots, I$ without loss of generality. The aim is to estimate the probability distribution \mathbb{P} . The usual estimates are computed by the observed relative frequencies

$$\hat{p}_i := \frac{1}{n} \sum_{t=1}^n \mathbb{1}_{\{X_t=i\}} \quad \text{for } i = 1, \dots, I.$$

These estimators satisfy, as $n \rightarrow \infty$, the well-known central limit theorem

$$(2.1) \quad \sqrt{n}((\hat{p}_1, \dots, \hat{p}_I)^\top - (p_1, \dots, p_I)^\top) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \Sigma),$$

where $\xrightarrow{\mathcal{D}}$ denotes weak convergence and where the asymptotic covariance matrix $\Sigma = (\Sigma_{rs})_{r,s=1,\dots,I}$ is given by

$$(2.2) \quad \Sigma_{rs} = \begin{cases} -p_r p_s & , r \neq s \\ p_r(1 - p_r) & , r = s \end{cases}.$$

The estimators \hat{p}_i , $i = 1, \dots, I$, can be interpreted as uniformly weighted (or unweighted) means of the observations. Changing the weights from $1/n$ to some general weights $w_t \geq 0$, $t = 1, \dots, n$, with $\sum_{t=1}^n w_t = 1$, yields the more general estimators

$$(2.3) \quad \tilde{p}_i := \sum_{t=1}^n w_t \mathbb{1}_{\{X_t=i\}}, \quad i = 1, \dots, I,$$

for which we observe a slightly different limiting result whose proof is straightforward:

Lemma 2.1. *Let $X_1, \dots, X_n \in \{1, \dots, I\}$ be an i.i.d. sample stemming from a discrete probability distribution $\mathbb{P} = (p_1, \dots, p_I)$, $I \in \mathbb{N}$. Then the estimators $(\tilde{p}_1, \dots, \tilde{p}_I)^\top$ from (2.3) satisfy, as $n \rightarrow \infty$,*

$$(2.4) \quad \frac{1}{\sqrt{\sum_{t=1}^n w_t^2}} ((\tilde{p}_1, \dots, \tilde{p}_I)^\top - (p_1, \dots, p_I)^\top) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \Sigma),$$

where Σ is given by (2.2).

By the Cauchy-Schwartz inequality we have $1 = \sum_{i=1}^n w_i \leq \sqrt{\sum_{i=1}^n w_i^2} \sqrt{n}$ and thus $\frac{1}{\sqrt{\sum_{i=1}^n w_i^2}} \leq \frac{1}{\sqrt{n}}$. Equality holds true if and only if the weights are chosen as $w_t = 1/n$ for all $t = 1, \dots, n$. Consequently, in a univariate i.i.d. framework non-uniformly weighting of observations always leads to an increased variance and slower convergence rates resulting in wider confidence interval approximations for finite sample sizes. We should emphasize at this point that this property might change when leaving the i.i.d. framework and dependencies between the observations are allowed.

Remark 2.2. *In some situations appliers prefer to use data clones instead of (relatively) weighted observations. This may be due to easier implementations when working with databases since it is relatively simple to copy cases. Lemma 2.1 applies immediately to cloned data which can be seen as follows: Let $Z_t \in \mathbb{N}$ for all $t = 1, \dots, n$, be a deterministic sequence with $\sum_{t=1}^n Z_t = N$. Any observation X_t may be cloned Z_t times such that the sample size artificially increases. The new (cloned) sample then reads Y_1, \dots, Y_N , where $Y_{1+\sum_{k=1}^{t-1} Z_k} = \dots = Y_{\sum_{k=1}^t Z_k} = X_t$ for all $t = 1, \dots, n$. The corresponding estimator for p_i is defined by*

$$\check{p}_i = \frac{1}{N} \sum_{r=1}^N \mathbb{1}_{\{Y_r=i\}} = \frac{1}{N} \sum_{t=1}^n Z_t \mathbb{1}_{\{X_t=i\}}.$$

We recover the estimators from (2.3) with weights $w_t := \frac{Z_t}{N}$. Since the minimum variance is obtained for non-cloned observations, i.e. $Z_t = 1$ for all $t = 1, \dots, n$, we conclude that cloning increases the variability in a one-dimensional model.

3. CONTINGENCY TABLES

The findings of the previous section directly transfer to the multivariate case. However, we will now investigate whether the situation changes if we have some extra information. Considering a two-dimensional contingency table, we study the estimation of one marginal distribution, say the first, under the assumption that the other marginal distribution, the second, is fully known.

Let $(X_1, Y_1), \dots, (X_n, Y_n)$ be a two-dimensional sample taking values in $\{1, \dots, I\} \times \{1, \dots, J\}$ for some $I, J \in \mathbb{N}$ and being distributed according to the law $\mathbb{P} = (p_{ij})_{i=1, \dots, I; j=1, \dots, J}$. Having I, J fixed, we consider a multinomial sampling scheme. The first and second marginals, or equivalently row and column marginal, are denoted by (p_{1, \dots, p_I}) and $(p_{\cdot 1}, \dots, p_{\cdot J})$, respectively. Without loss of generality, we can assume $p_{\cdot j} > 0$ for all $j = 1, \dots, J$. This setting also includes higher dimensional distributions since we can represent all coordinates with known (joint) distribution as some multi-dimensional random variable $(Y_r)_{r=1, \dots, R}$, $R \in \mathbb{N}$, in a finite state space while the remaining coordinates with unknown marginal distribution can be written as some $(X_s)_{s=1, \dots, S}$, $S \in \mathbb{N}$.

The commonly used estimators for the two-dimensional probabilities p_{ij} are given by

$$\hat{p}_{ij} = \frac{1}{n} \sum_{t=1}^n \mathbb{1}_{\{X_t=i, Y_t=j\}} \quad \text{for all } i = 1, \dots, I, j = 1, \dots, J.$$

The resulting estimator for the first marginal distribution is then defined via

$$\hat{p}_{i\cdot} := \frac{1}{n} \sum_{t=1}^n \mathbb{1}_{\{X_t=i\}} = \sum_{j=1}^J \hat{p}_{ij}, \quad i = 1, \dots, I.$$

In view of (2.1), these estimators satisfy, as $n \rightarrow \infty$,

$$(3.1) \quad \sqrt{n} \left((\hat{p}_{1\cdot}, \dots, \hat{p}_{I\cdot})^\top - (p_{1\cdot}, \dots, p_{I\cdot})^\top \right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \Sigma)$$

where $\Sigma = (\Sigma_{rs})_{r,s=1,\dots,I}$ is given by

$$(3.2) \quad \Sigma_{rs} = \begin{cases} p_{r\cdot}(1 - p_{r\cdot}) & , r = s \\ -p_{r\cdot}p_{s\cdot} & , r \neq s \end{cases}.$$

The estimators $(\hat{p}_{1\cdot}, \dots, \hat{p}_{I\cdot})$ ignore any specific information on the second components of the sample. In other words, the estimators treat the two-dimensional sample as a one-dimensional sample. In order to incorporate the additional information to our estimates, we weight the estimators \hat{p}_{ij} such that the column-wise marginal distribution of the estimators coincide with the true and known marginal distribution $(p_{\cdot 1}, \dots, p_{\cdot J})$. More precisely, we introduce the weighted estimators as

$$(3.3) \quad \tilde{p}_{ij} := \frac{p_{\cdot j}}{\hat{p}_{\cdot j}} \hat{p}_{ij} \quad \text{for any } i = 1, \dots, I; j = 1, \dots, J.$$

Owing to the assumption $p_{\cdot j} > 0$, $j = 1, \dots, J$, we have $\hat{p}_{\cdot j} > 0$ with probability one for a sufficiently large sample size. Hence, due to the above definition we indeed obtain $\tilde{p}_{\cdot j} := \sum_{i=1}^I \tilde{p}_{ij} = p_{\cdot j}$ for all $j = 1, \dots, J$.

Following the discussion in Remark 2.2, the weighted estimators \tilde{p}_{ij} can be implemented via careful cloning. For instance, if every entry in the database with $Y_t = j$ is copied $np_{\cdot j}$ times, then the relative frequencies in the modified dataset coincide with \tilde{p}_{ij} .

Remark 3.1. *The assumption that $\mathbb{P}^Y = (p_{\cdot 1}, \dots, p_{\cdot J})$ is known can be considerably relaxed. The following results can be extended to the case where the second marginal distribution can be estimated with higher accuracy such that the additional estimation error is negligible, for instance, if we have another i.i.d. sample $Y'_1, \dots, Y'_m \sim \mathbb{P}^Y$ at hand where the sample size $m \in \mathbb{N}$ satisfies $\frac{n}{m} \rightarrow 0$.*

While the adjusted estimators are fairly simple to implement, they are not anymore linear in the data owing to the data dependent weights. The non-linearity causes a small bias which is negligible compared to the parametric rate $n^{-1/2}$.

Lemma 3.2. *Let $(X_1, Y_1), \dots, (X_n, Y_n) \in \{1, \dots, I\} \times \{1, \dots, J\}$ be an i.i.d. sample of a two-dimensional distribution $\mathbb{P} = (p_{ij})_{i=1,\dots,I; j=1,\dots,J}$, with some $I, J \in \mathbb{N}$. Then the estimators \tilde{p}_{ij} from (3.3) satisfy $|\mathbb{E}[\tilde{p}_{ij}] - p_{ij}| \leq 2/n$ for all $1 \leq i \leq I, 1 \leq j \leq J$.*

Using the weighted estimators \tilde{p}_{ij} , the modified contingency table reads

$$\begin{array}{ccc|c} \tilde{p}_{11} & \cdots & \tilde{p}_{1J} & \tilde{p}_{1\cdot} \\ \vdots & & \vdots & \vdots \\ \tilde{p}_{I1} & \cdots & \tilde{p}_{IJ} & \tilde{p}_{I\cdot} \\ \hline p_{\cdot 1} & & p_{\cdot J} & 1 \end{array}$$

where the estimates of interest are given by

$$\tilde{p}_{i\cdot} := \sum_{j=1}^J \tilde{p}_{ij} \quad \text{for all } j = 1, \dots, J.$$

It is worth to mention that the adjusted estimators \tilde{p}_{ij} do not change the cross-product ratios

$$\frac{p_{ij}p_{rs}}{p_{rj}p_{is}}, \quad i, r = 1, \dots, I; j, s = 1, \dots, J, \quad i \neq r, j \neq s,$$

by construction. These ratios describe the degree of association in a contingency table. If, for instance, all cross-product ratios in a contingency table are equal to 1, the table yields independence of the marginals, while values of zero or infinity describe dependence. Since the weighted estimators maintain the two-dimensional inner dependence structure for raw data, cross-product ratio-based tests, e.g. for independence (Fisher (1962), Goodman (1964)), can be applied to the modified contingency table without any restrictions and changes in interpretation.

Remark 3.3. *The choice of the weights is in line with a single iteration of the so-called iterative proportional fitting (IPF) procedure by Deming and Stephan (1940). By taking marginal distributions of lower dimension, the IPF algorithm computes a higher dimensional distribution fulfilling the given marginal restrictions (cf. Stephan (1942)). The algorithm works conditional on a specific initialization table. This situation can be interpreted as a two-dimensional sample from which both marginals are known. The procedure aims at adapting the sample characteristics to the given marginals. In the present setup, the weighted sample can be regarded as the result of a single IPF iteration such that the following considerations also hold for any iteration step of the IPF algorithm. For further details on the IPF we refer to Fienberg (1970) and Rüschemdorf (1995).*

Intuitively, the inclusion of further information should at least not worsen the estimators' features. To rigorously evaluate whether there is an improvement from the original estimators $(\hat{p}_{i\cdot})_{i=1,\dots,I}$ to the modified estimators $(\tilde{p}_{i\cdot})_{i=1,\dots,I}$, we will compare their limiting variances.

Theorem 3.4. *Let $(X_1, Y_1), \dots, (X_n, Y_n) \in \{1, \dots, I\} \times \{1, \dots, J\}$ be an i.i.d. sample of a two-dimensional distribution $\mathbb{P} = (p_{ij})_{i=1,\dots,I; j=1,\dots,J}$, with some $I, J \in \mathbb{N}$. Then the estimators $\tilde{p}_{i\cdot} = \sum_{j=1}^J \tilde{p}_{ij}$, $i = 1, \dots, I$, with \tilde{p}_{ij} from (3.3) satisfy, as $n \rightarrow \infty$, the limiting result*

$$\sqrt{n}((\tilde{p}_{1\cdot}, \dots, \tilde{p}_{I\cdot})^\top - (p_{1\cdot}, \dots, p_{I\cdot})^\top) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \Gamma),$$

where $\Gamma = (\Gamma_{k,l})_{k,l=1,\dots,I}$ is given by

$$(3.4) \quad \Gamma_{k,l} = p_{k\cdot} \mathbf{1}_{k=l} - \sum_{j=1}^J \frac{p_{kj}p_{lj}}{p_{\cdot j}}.$$

The quite simple and explicit structure of the asymptotic variance Γ is charming. It can be rigorously understood in terms of the conditional (co-)variance. The latter is defined by

$$\text{Cov}(V, W|Z) := \mathbb{E}[(V - \mathbb{E}[V|Z])(W - \mathbb{E}[W|Z])|Z]$$

for arbitrary random variables V, W and Z on the same probability space. In the following, the notation $\text{Var}(V|Z) := \text{Cov}(V, V|Z)$ also is used. As verified by the following lemma, the asymptotic covariance matrix Γ of the weighted estimators for the first marginal is given by the covariance of the unweighted estimators conditional on the second component. Hence, we recover in the asymptotic scale that the second margin is known.

Lemma 3.5. *Grant the assumption of Theorem 3.4. We then have for any $k, l \in \{1, \dots, I\}$ and $\Gamma_{k,l}$ from (3.4) that*

$$\frac{1}{n}\Gamma_{k,l} = \mathbb{E}[\text{Cov}(\hat{p}_{k\cdot}, \hat{p}_{l\cdot} | Y_1, \dots, Y_n)].$$

This lemma implies together with the law of total variance that for any $i = 1, \dots, I$

$$(3.5) \quad \frac{\Sigma_{i,i}}{n} = \text{Var}(\hat{p}_{i\cdot}) = \mathbb{E}[\text{Var}(\hat{p}_{i\cdot} | Y_1, \dots, Y_n)] + \text{Var}(\mathbb{E}[\hat{p}_{i\cdot} | Y_1, \dots, Y_n]) \geq \frac{\Gamma_{i,i}}{n}.$$

Therefore, the asymptotic variance of the adjusted estimators \tilde{p}_i is indeed less or equal to the asymptotic variance of the classical estimators \hat{p}_i . The (asymptotic) information theoretic gain is given by $\text{Var}(\mathbb{E}[\hat{p}_i | Y_1, \dots, Y_n])$. This fact can be generalized to the joint limit law of $(\tilde{p}_{1\cdot}, \dots, \tilde{p}_{I\cdot})$.

Corollary 3.6. *Grant the assumptions of Theorem 3.4. Then the asymptotic covariance matrices Σ and Γ given in (3.2) and (3.4), respectively, satisfy $\Sigma \geq \Gamma$ in the sense of positive semi-definite matrices: For any vector $c \in \mathbb{R}^I$, we have*

$$c^\top (\Sigma - \Gamma)c = \text{Var} \left(\mathbb{E} \left[\sum_{i=1}^I c_i \mathbb{1}_{\{X_1=i\}} \middle| Y_1 \right] \right) \geq 0.$$

In particular, $\Sigma = \Gamma$ holds if and only if (X_1, \dots, X_n) and (Y_1, \dots, Y_n) are independent.

Corollary 3.6 shows that the modification of the estimators by the additional information on the known marginal asymptotically improves the estimation of the target marginal. The degree of improvement depends on the degree of association of the two marginals. To be more precise, the following bound for the overall relative variance reduction can be easily deduced from (3.5):

$$\sum_{i=1}^I \frac{\Sigma_{ii} - \Gamma_{ii}}{\Sigma_{ii}} \geq \sum_{i=1}^I \frac{\Sigma_{ii} - \Gamma_{ii}}{p_{i\cdot}} = \sum_{i=1}^I \sum_{j=1}^J \frac{(p_{ij} - p_{i\cdot} p_{\cdot j})^2}{p_{i\cdot} p_{\cdot j}}.$$

On the right-hand side we recover a natural measure for dependence being the population counterpart of the well known χ^2 -test statistic.

4. NUMERICAL RESULTS

Theorem 3.4 and Corollary 3.6 state that using the modified estimators is asymptotically advantageous over using estimators which ignore the additional information. We have seen that the degree of improvement is related to the degree of association. Focusing first on this effect in a finite sample situation, we investigate the gain of the modification by simulations in 2x2 contingency tables with fixed marginals by varying the degree of association in terms of the cross-product ratio $cpr = (p_{11}p_{22})/(p_{12}p_{21})$ and several sample sizes n .

We consider three combinations of marginals, namely

- (I) $(p_{1\cdot}, p_{2\cdot}) = (0.5, 0.5)$ and $(p_{\cdot 1}, p_{\cdot 2}) = (0.5, 0.5)$,
- (II) $(p_{1\cdot}, p_{2\cdot}) = (0.9, 0.1)$ and $(p_{\cdot 1}, p_{\cdot 2}) = (0.7, 0.3)$,
- (III) $(p_{1\cdot}, p_{2\cdot}) = (0.2, 0.8)$ and $(p_{\cdot 1}, p_{\cdot 2}) = (0.7, 0.3)$.

Let us focus on the estimation of $p_{1\cdot}$. To compare the performance of the estimators $\hat{p}_{1\cdot}$ and $\tilde{p}_{1\cdot}$, we simulated for any sample size $n \in \{20, 50, 200, 500\}$ and any given cross-product ratio 100,000 samples from which the estimators were computed and noted. Since the theory has shown that the variance of $\hat{p}_{1\cdot}$ should be larger than the variance of $\tilde{p}_{1\cdot}$, we denote the average relative proportion of the variance which is removed by the modified estimator $\tilde{p}_{1\cdot}$. The simulation results are given in Figures 1-3. Positive values in the figures indicate whenever the modification of the estimator is -on average- advantageous while negative values indicate that the modification was misleading. Recall that $\log(cpr) = 0$ corresponds to independence of the marginals, while the degree of association increases with the distance from the origin in the horizontal axis.

The simulation results show that for already small and moderate sample sizes the modified estimators lead to performance improvements. It seems that an exception is given in case of independent marginals. In these cases there is a very slight disadvantage for the modified estimator. This effect vanishes when the sample size increases as expected from the theory. Even more important the modified estimators lead to substantial improvements when there is some degree of dependence.

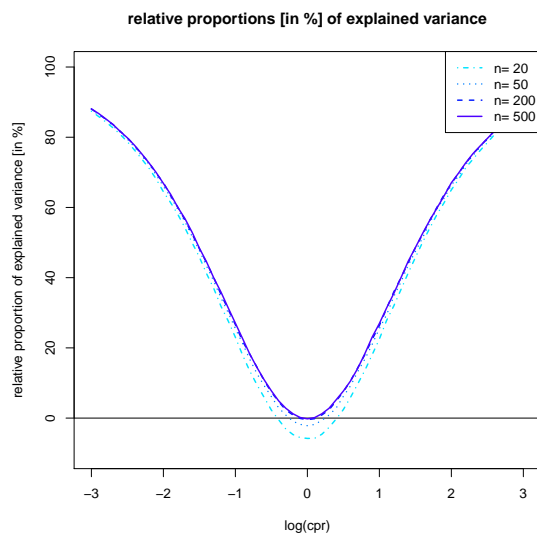


FIGURE 1. Relative proportions (in %) of explained variances conditional on the sample size n and the cross-product ratio cpr for the fixed marginals $(p_{1\cdot}, p_{2\cdot}) = (0.5, 0.5)$ and $(p_{\cdot 1}, p_{\cdot 2}) = (0.5, 0.5)$.

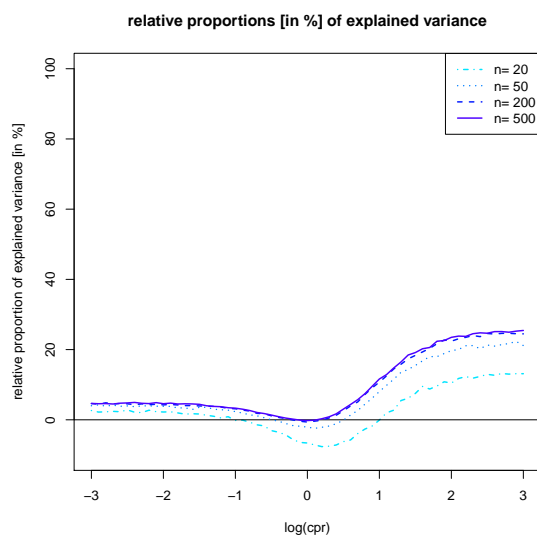


FIGURE 2. Relative proportions (in %) of explained variances conditional on the sample size n and the cross-product ratio cpr for the fixed marginals $(p_{1\cdot}, p_{2\cdot}) = (0.9, 0.1)$ and $(p_{\cdot 1}, p_{\cdot 2}) = (0.7, 0.3)$.

We will now consider a second simulation setup in order to investigate the finite sample performance for large contingency tables. To this end, we generate two independent binomial $Bin(J, p)$ -random variables (Y, Z) with success probability $p = 0.5$. Setting $X = Y + Z \sim Bin(2J, p)$, we obtain a two dimensional random vector $(X, Y) \in \{0, \dots, 2J\} \times \{0, \dots, J\}$ with fixed correlation $Cor(X, Y) = 1/\sqrt{2}$ for any dimension J . Based on n independent copies of (X, Y) , we use the classical as well as the weighted estimators for the marginal distribution $(p_{0\cdot}, \dots, p_{2J\cdot})$ of X . We

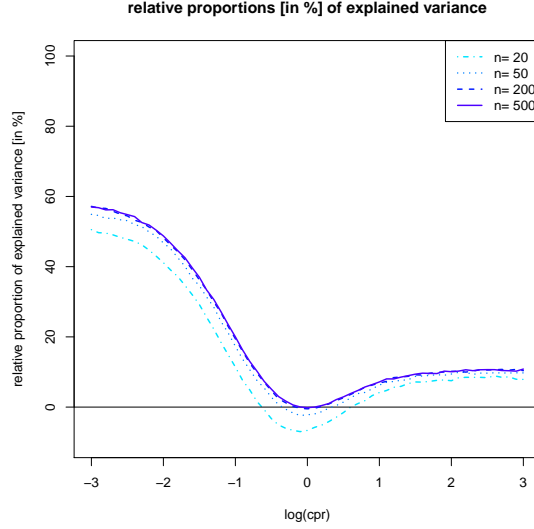


FIGURE 3. Relative proportions (in %) of explained variances conditional on the sample size n and the cross-product ratio cpr for the fixed marginals $(p_{1\cdot}, p_{2\cdot}) = (0.2, 0.8)$ and $(p_{\cdot 1}, p_{\cdot 2}) = (0.7, 0.3)$.

I	Bias(\hat{p})	Var(\hat{p})	MSE(\hat{p})	Bias(\tilde{p})	Var(\tilde{p})	MSE(\tilde{p})
5	$0.731 \cdot 10^{-4}$	$2.459 \cdot 10^{-4}$	$2.459 \cdot 10^{-4}$	$0.691 \cdot 10^{-4}$	$2.035 \cdot 10^{-4}$	$2.035 \cdot 10^{-4}$
11	$1.008 \cdot 10^{-4}$	$3.188 \cdot 10^{-4}$	$3.188 \cdot 10^{-4}$	$0.984 \cdot 10^{-4}$	$2.895 \cdot 10^{-4}$	$2.895 \cdot 10^{-4}$
21	$0.879 \cdot 10^{-4}$	$3.651 \cdot 10^{-4}$	$3.652 \cdot 10^{-4}$	$2.257 \cdot 10^{-4}$	$3.449 \cdot 10^{-4}$	$3.449 \cdot 10^{-4}$
41	$0.815 \cdot 10^{-4}$	$3.991 \cdot 10^{-4}$	$3.991 \cdot 10^{-4}$	$2.934 \cdot 10^{-4}$	$3.869 \cdot 10^{-4}$	$3.869 \cdot 10^{-4}$
61	$1.205 \cdot 10^{-4}$	$4.184 \cdot 10^{-4}$	$4.118 \cdot 10^{-4}$	$3.460 \cdot 10^{-4}$	$4.104 \cdot 10^{-4}$	$4.105 \cdot 10^{-4}$
101	$1.258 \cdot 10^{-4}$	$4.354 \cdot 10^{-4}$	$4.354 \cdot 10^{-4}$	$4.348 \cdot 10^{-4}$	$4.322 \cdot 10^{-4}$	$4.323 \cdot 10^{-4}$

TABLE 1. Monte Carlo approximation of bias $\text{Bias}(\tilde{p}) := \|(\mathbb{E}[\tilde{p}_{i\cdot}] - p_{i\cdot})_i\|_{\ell^2}$, the variance $\text{Var}(\tilde{p}) := \sum_{i=1}^I \text{Var}(\tilde{p}_{i\cdot})$ and the mean squared error $\text{MSE}(\tilde{p}) := \mathbb{E}[\|\tilde{p}_{i\cdot} - p_{i\cdot}\|_{\ell^2}^2]$ of the estimators $(\hat{p}_{i\cdot})_{i=0,\dots,I-1}$ and $(\tilde{p}_{i\cdot})_{i=0,\dots,I-1}$, respectively.

thus have $I = 2J + 1$ parameters. Using 100,000 Monte Carlo iterations, we approximate the ℓ^2 -norm of bias, the overall variance and mean squared error of the estimated vectors $(\hat{p}_{i\cdot})_{i=0,\dots,I-1}$ and $(\tilde{p}_{i\cdot})_{i=0,\dots,I-1}$.

The simulation results for $n = 1000$ and $J \in \{2, 5, 10, 20, 30, 50\}$ are reported in Table 1, confirming our theory. We observe a slightly larger bias of the modified estimators for moderate and large dimensions due to the small non-asymptotic bias of $\tilde{p}_{i\cdot}$ for any $i = 1, \dots, I$. Nevertheless, the squared bias is only of the order of 10^{-4} of the variance. The contribution of this bias to the MSE is thus negligible. Consequently, the weighted estimators are in advantage over the classical approach in all considered scenarios. The relative improvement $(\text{MSE}(\hat{p}) - \text{MSE}(\tilde{p})) / \text{MSE}(\hat{p})$ ranges from 17.2% for dimension 5 to 0.7% for dimension 101.

5. A CASE STUDY FROM ACCIDENT RESEARCH

The assessment of the effectiveness of advanced driver assistance systems plays a crucial role in traffic accident research. For reliable analysis detailed information on the pre-crash phase of an accident has to be known to predict possible benefits of future driver assistant systems. The

type of injury	slightly injured	severely injured	fatally injured
number of cases	106,181 [89.6%]	11,898 [10.0%]	423 [0.4%]

TABLE 2. Injury severity distribution in passenger car to passenger car collisions with two accident participants for Germany in 2014 (from Deutsches Bundesamt (2016), p. 100).

	slightly injured	severely injured	fatally injured	total	adjusted	relative difference
$0 \leq \Delta v \leq 10$	346	24	2	372 [11.4%]	12.0%	+5.06%
$11 \leq \Delta v \leq 20$	935	118	1	1054 [32.4%]	33.6%	+3.67%
$21 \leq \Delta v \leq 30$	739	192	4	935 [28.7%]	28.9%	+0.43%
$31 \leq \Delta v \leq 40$	335	154	7	496 [15.2%]	14.7%	-3.49%
$41 \leq \Delta v \leq 50$	124	92	7	223 [6.9%]	6.3%	-7.63%
$51 \leq \Delta v \leq 60$	41	50	6	97 [3.0%]	2.6%	-12.36%
$61 \leq \Delta v \leq 70$	16	25	6	47 [1.4%]	1.2%	-15.75%
$71 \leq \Delta v$	2	21	7	30 [0.9%]	0.7%	-25.89%
total	2,538 [78.0%]	676 [20.8%]	40 [1.2%]	3,254 [100%]		

TABLE 3. Two-dimensional data pattern (kind of injury and Δv) from GIDAS in passenger car to passenger car collisions with two accident participants together with marginal estimates for both approaches and their relative difference.

German In-Depth Accident Study (GIDAS) data contains hundreds of categories which carefully have to be reported for every single accident. Accidents are reported to GIDAS teams by police, if and only if injured participants are to be expected. This consequently leads to a substantial bias in the injury severity (e.g. Otte et al. (2003), Pfeiffer and Schmidt (2006)). Collision speed or the speed reduction due to the collision, say Δv , is found to be a major correlate to injury severity in traffic accidents.

In this case study we estimate the (discretized) distribution of Δv in passenger car to passenger car collisions with two accident participants. The corresponding GIDAS subsample contains 8,753 cases. To address the skewness in the injury severity of the GIDAS sample, we queried the German National Statistic of the year 2014 (Statistisches Bundesamt (2016)) from which we found the true injury severity distribution for Germany for these accidents (cf. Table 2).

Table 2 can be interpreted as a known marginal distribution for the GIDAS sample. Hence, having the two-dimensional distribution (about the injury severity and Δv) from GIDAS at hand, see Table 3, we can apply the estimation strategy as developed in Section 3. The estimates for the distribution of Δv for the purely GIDAS-based approach, i.e. $\hat{p}_{i\cdot}$, are contained in the column *total* of Table 3. The National Statistic-aided weighted estimators $\tilde{p}_{i\cdot}$ are given in the column *adjusted* besides its relative difference to the purely GIDAS-based estimates.

From Table 3 we see that the adapted estimates -compared to the ordinary estimates- are increased for lower collision speeds and decreased for higher collision speeds. On a relative scale, the estimate for the highest Δv interval is reduced to around three-quarters of the original estimate. This is not too surprising. It was already stated that GIDAS is biased towards the more severely injured traffic participants, cf. Otte et al. (2003), or compare Table 2 to the marginal distribution in Table 3. Hence, it is to be expected that accidents with higher collision speeds -which are clearly

associated with more severe injuries– will be down-weighted. Our proposed method gives precise weights to adjust the more severely injured cases leading to the final outcome in Table 3.

6. CONCLUSIONS

The paper investigated how weighting affects the estimation of a discrete probability distribution. While for one-dimensional data relative weighting will always increase estimation variances, it has been shown that additional information on a marginal distribution in a contingency table allows for estimation improvements of further marginals if there is some degree of association between the two categories. The gain in terms of the asymptotic (co)variance increases with the degree of dependence. The weighting causes a small bias, which is however negligible compared to the improved variance. For independent marginals the weighted estimators have the same asymptotic behavior as their classical unweighted counterparts. The simulations indicate a clear gain when the marginals are substantially associated. Therefore, from theory and from the simulations perspective we suggest to use the adjusted estimators in applications whenever it cannot be assumed that the marginals are independent.

7. PROOFS

Proof of Lemma 3.2: We decompose

$$\frac{1}{\hat{p}_{\cdot j}} - \frac{1}{p_{\cdot j}} = \frac{p_{\cdot j} - \hat{p}_{\cdot j}}{p_{\cdot j}\hat{p}_{\cdot j}} = \frac{p_{\cdot j} - \hat{p}_{\cdot j}}{p_{\cdot j}^2} + \frac{(p_{\cdot j} - \hat{p}_{\cdot j})^2}{p_{\cdot j}^2\hat{p}_{\cdot j}}.$$

Therefore,

$$\begin{aligned} \mathbb{E}[\tilde{p}_{ij}] - p_{ij} &= p_{\cdot j}\mathbb{E}\left[\left(\frac{1}{\hat{p}_{\cdot j}} - \frac{1}{p_{\cdot j}}\right)\hat{p}_{ij}\right] \\ &= \frac{1}{p_{\cdot j}}\mathbb{E}[(p_{\cdot j} - \hat{p}_{\cdot j})\hat{p}_{ij}] + \mathbb{E}\left[\frac{(p_{\cdot j} - \hat{p}_{\cdot j})^2}{p_{\cdot j}} \frac{\hat{p}_{ij}}{\hat{p}_{\cdot j}}\right] =: T_1 + T_2. \end{aligned}$$

Since $\hat{p}_{\cdot j}$ is unbiased, we obtain for the first term

$$\begin{aligned} T_1 &= \frac{1}{p_{\cdot j}}\mathbb{E}[(p_{\cdot j} - \hat{p}_{\cdot j})(\hat{p}_{ij} - p_{ij})] \\ &\leq \frac{1}{p_{\cdot j}} \text{Var}(\hat{p}_{\cdot j})^{1/2} \text{Var}(\hat{p}_{ij})^{1/2} \leq \frac{1}{n} \end{aligned}$$

using the Cauchy-Schwarz inequality, $\text{Var}(\hat{p}_{\cdot j}) = \frac{1}{n}p_{\cdot j}(1 - p_{\cdot j})$ (and analogously for \hat{p}_{ij}) and $p_{ij} \leq p_{\cdot j}$. For the second term the property $\hat{p}_{ij} \leq \hat{p}_{\cdot j}$ yields

$$T_2 \leq \frac{1}{p_{\cdot j}} \text{Var}(\hat{p}_{\cdot j}) \leq \frac{1}{n}.$$

□

Proof of Theorem 3.4: We first note that the matrix $(\tilde{p}_{ij})_{i=1,\dots,I,j=1,\dots,J}$ can be understood as a function of all \hat{p}_{ij} . Since a limiting result for the joint distribution of all \hat{p}_{ij} is given in (3.1), the asymptotic features of the joint distribution of the weighted estimators \tilde{p}_{ij} can be determined using the delta method.

We define the matrices

$$Z = (\tilde{p}_{ij})_{i=1,\dots,I;j=1,\dots,J} \quad \text{and} \quad W = (\hat{p}_{ij})_{i=1,\dots,I;j=1,\dots,J}$$

and calculate

$$\begin{aligned} Z &= W \cdot \text{diag}\left(\frac{p_{\cdot 1}}{\hat{p}_{\cdot 1}}, \dots, \frac{p_{\cdot J}}{\hat{p}_{\cdot J}}\right) \\ &= W \cdot \text{diag}(\hat{p}_{\cdot 1}, \dots, \hat{p}_{\cdot J})^{-1} \cdot \text{diag}(p_{\cdot 1}, \dots, p_{\cdot J}) \end{aligned}$$

$$= W \cdot \text{diag}(\mathbb{1}_I^\top \cdot W)^{-1} \cdot \text{diag}(p_{\cdot 1}, \dots, p_{\cdot J}),$$

where $\mathbb{1}_I \in \mathbb{R}^I$ represents a vector of I repetitions of the value 1. Now interpret

$$\begin{pmatrix} \tilde{p}_{1\cdot} \\ \vdots \\ \tilde{p}_{I\cdot} \end{pmatrix} = Z \mathbb{1}_J = W \cdot \text{diag}(\mathbb{1}_I^\top \cdot W)^{-1} \cdot \begin{pmatrix} p_{\cdot 1} \\ \vdots \\ p_{\cdot J} \end{pmatrix}$$

as a function in W or equivalently $\text{vec}(W) = \text{vec}((\hat{p}_{ij})_{i=1, \dots, I; j=1, \dots, J})$. Hence, $\tilde{P} := (\tilde{p}_{i\cdot})_{i=1, \dots, I}$ can be defined by $\tilde{P} : \mathbb{R}^{IJ} \rightarrow \mathbb{R}^I$, $\text{vec}(W) \mapsto (\tilde{p}_{i\cdot})_{i=1, \dots, I}$. The entries of the associated Jacobi matrix $J_{\tilde{P}} \in \mathbb{R}^{I \times (IJ)}$ are computed in the sequel.

To this end, we denote by \mathbb{I}_{ij} the matrix of zeros with a single 1 at the i -th row and j -th column. We set $u := u_{ij} := i + (j - 1)I$ for any $i = 1, \dots, I$ and $j = 1, \dots, J$ and calculate the u -th column of $J_{\tilde{P}}$ at the point $\text{vec}(W)$ (for brevity we write $\frac{\partial}{\partial W_u}$ instead of $\frac{\partial}{\partial \text{vec}(W)_u}$):

$$\begin{aligned} & \frac{\partial}{\partial W_u} (\tilde{p}_{1\cdot}, \dots, \tilde{p}_{I\cdot})^\top \\ &= \left(\frac{\partial}{\partial W_u} W \right) \cdot \text{diag}(\mathbb{1}_I^\top \cdot W)^{-1} \cdot (p_{\cdot 1}, \dots, p_{\cdot J})^\top + W \cdot \frac{\partial}{\partial W_u} \text{diag}(\mathbb{1}_I^\top \cdot W)^{-1} \cdot (p_{\cdot 1}, \dots, p_{\cdot J})^\top \\ &= \mathbb{I}_{ij} \cdot \text{diag}(\mathbb{1}_I^\top \cdot W)^{-1} \cdot (p_{\cdot 1}, \dots, p_{\cdot J})^\top \\ &\quad + W \cdot \left(\frac{\partial}{\partial W_{i+(j-1)I}} \text{diag} \left(\frac{1}{(\mathbb{1}_I^\top \cdot W)_1}, \dots, \frac{1}{(\mathbb{1}_I^\top \cdot W)_J} \right) \right) \cdot (p_{\cdot 1}, \dots, p_{\cdot J})^\top \\ &= \mathbb{I}_{ij} \cdot \text{diag}(\mathbb{1}_I^\top \cdot W)^{-1} \cdot (p_{\cdot 1}, \dots, p_{\cdot J})^\top - W \cdot \text{diag} \left(\left(\frac{\mathbb{1}_{r=j}}{(\mathbb{1}_I^\top \cdot W)_r^2} \right)_{r=1, \dots, J} \right) \cdot (p_{\cdot 1}, \dots, p_{\cdot J})^\top \\ &= \mathbb{I}_{ij} \cdot \text{diag}(\mathbb{1}_I^\top \cdot W)^{-1} \cdot (p_{\cdot 1}, \dots, p_{\cdot J})^\top - W \cdot \left(\frac{p_{\cdot 1} \mathbb{1}_{j=1}}{(\mathbb{1}_I^\top \cdot W)_1^2}, \dots, \frac{p_{\cdot J} \mathbb{1}_{j=J}}{(\mathbb{1}_I^\top \cdot W)_J^2} \right)^\top \\ &= \left(\frac{p_{\cdot j} \mathbb{1}_{r=i}}{(\mathbb{1}_I^\top \cdot W)_j} \right)_{r=1, \dots, I} - \frac{p_{\cdot j}}{(\mathbb{1}_I^\top \cdot W)_j^2} (W_{rj})_{r=1, \dots, I} \\ &= \frac{p_{\cdot j}}{\hat{p}_{\cdot j}} (\mathbb{1}_{i=r})_{r=1, \dots, I} - \frac{p_{\cdot j}}{\hat{p}_{\cdot j}^2} (W_{rj})_{r=1, \dots, I}. \end{aligned}$$

For the delta method to apply, we do not need $J_{\tilde{P}}$ at $\text{vec}(W)$, but at its non-estimated target value, i.e. we require $J_{\tilde{P}}$ at $\text{vec}(P) := \text{vec}((p_{ij})_{i=1, \dots, I; j=1, \dots, J})$. By the above computation, the m -th row and $(i + (j - 1)I)$ -th column value is

$$J_{\tilde{P}, m, i+(j-1)I}(\text{vec}(P)) = -\frac{p_{m,j}}{p_{\cdot j}} + \mathbb{1}_{i=m}.$$

Hence, having all derivatives at hand, the limiting covariance matrix reads

$$\Gamma = J_{\tilde{P}} \Sigma J_{\tilde{P}}^\top$$

which we can be calculated component-wise to

$$\begin{aligned} \Gamma_{m,n} &= \sum_{r=1}^I \sum_{s=1}^J \sum_{t=1}^I \sum_{v=1}^J J_{\tilde{P}, m, r+(s-1)I} \Sigma_{r+(s-1)I, t+(v-1)I} J_{\tilde{P}, n, t+(v-1)I} \\ &= \sum_{r=1}^I \sum_{s=1}^J \sum_{t=1}^I \sum_{v=1}^J \left(-\frac{p_{m,s}}{p_{\cdot s}} + \mathbb{1}_{r=m} \right) (-p_{r,s} p_{t,v} + p_{r,s} \mathbb{1}_{r+(s-1)I=t+(v-1)I}) \left(-\frac{p_{n,v}}{p_{\cdot v}} + \mathbb{1}_{t=n} \right) \\ &= -\sum_{r=1}^I \sum_{s=1}^J \sum_{t=1}^I \sum_{v=1}^J p_{r,s} p_{t,v} \left(-\frac{p_{m,s}}{p_{\cdot s}} + \mathbb{1}_{r=m} \right) \left(-\frac{p_{n,v}}{p_{\cdot v}} + \mathbb{1}_{t=n} \right) \end{aligned}$$

$$\begin{aligned}
& + \sum_{r=1}^I \sum_{s=1}^J p_{r,s} \left(-\frac{p_{m,s}}{p_{\cdot s}} + \mathbb{1}_{r=m} \right) \left(-\frac{p_{n,s}}{p_{\cdot s}} + \mathbb{1}_{r=n} \right) \\
= & - \left(\sum_{r=1}^I \sum_{s=1}^J p_{r,s} \left(-\frac{p_{m,s}}{p_{\cdot s}} + \mathbb{1}_{r=m} \right) \right) \left(\sum_{t=1}^I \sum_{v=1}^J p_{t,v} \left(-\frac{p_{n,v}}{p_{\cdot v}} + \mathbb{1}_{t=n} \right) \right) \\
& + \sum_{r=1}^I \sum_{s=1}^J \frac{p_{rs} p_{ms} p_{ns}}{p_{\cdot s}^2} - 2 \sum_{s=1}^J \frac{p_{ms} p_{ns}}{p_{\cdot s}} + \sum_{s=1}^J p_{ns} \mathbb{1}_{m=n} \\
= & - \left(-\sum_{s=1}^J p_{ms} + p_{m\cdot} \right) \left(-\sum_{r=1}^J p_{nr} + p_{n\cdot} \right) \\
& + \sum_{r=1}^I \sum_{s=1}^J \frac{p_{rs} p_{ms} p_{ns}}{p_{\cdot s}^2} - 2 \sum_{s=1}^J \frac{p_{ms} p_{ns}}{p_{\cdot s}} + \sum_{s=1}^J p_{ns} \mathbb{1}_{m=n} \\
= & \sum_{s=1}^J \frac{p_{ms} p_{ns}}{p_{\cdot s}} - 2 \sum_{s=1}^J \frac{p_{ms} p_{ns}}{p_{\cdot s}} + p_{n\cdot} \mathbb{1}_{m=n} \\
= & p_{n\cdot} \mathbb{1}_{m=n} - \sum_{s=1}^J \frac{p_{ms} p_{ns}}{p_{\cdot s}}.
\end{aligned}$$

□

Proof of Lemma 3.5. Since $(X_t, Y_t)_{t=1, \dots, n}$ is an i.i.d. sample, it suffices to consider $n = 1$. We have for any $k, l = 1, \dots, I$

$$\begin{aligned}
\text{Cov}(\hat{p}_{k\cdot}, \hat{p}_{l\cdot} | Y_1) &= \mathbb{E}[\hat{p}_{k\cdot} \hat{p}_{l\cdot} | Y_1] - \mathbb{E}[\hat{p}_{k\cdot} | Y_1] \mathbb{E}[\hat{p}_{l\cdot} | Y_1] \\
&= \sum_{s=1}^J \sum_{r=1}^J \left(\mathbb{E}[\hat{p}_{ks} \hat{p}_{lr} | Y_1] - \mathbb{E}[\hat{p}_{ks} | Y_1] \mathbb{E}[\hat{p}_{lr} | Y_1] \right).
\end{aligned}$$

Owing to $\hat{p}_{ks} = \mathbb{1}_{\{X_1=k\}} \mathbb{1}_{\{Y_1=s\}}$, we obtain

$$\begin{aligned}
\text{Cov}(\hat{p}_{k\cdot}, \hat{p}_{l\cdot} | Y_1) &= \sum_{s=1}^J \mathbb{1}_{\{Y_1=s\}} \left(\mathbb{E}[\mathbb{1}_{\{X_1=k\}} \mathbb{1}_{\{X_1=l\}} | Y_1] - \mathbb{E}[\mathbb{1}_{\{X_1=k\}} | Y_1] \mathbb{E}[\mathbb{1}_{\{X_1=l\}} | Y_1] \right) \\
&= \sum_{s=1}^J \mathbb{1}_{\{Y_1=s\}} \left(\frac{p_{k,s}}{p_{\cdot s}} \mathbb{1}_{k=l} - \frac{p_{k,s} p_{l,s}}{p_{\cdot s}^2} \right).
\end{aligned}$$

Therefore,

$$\mathbb{E}[\text{Cov}(\hat{p}_{k\cdot}, \hat{p}_{l\cdot} | Y_1)] = p_{k\cdot} \mathbb{1}_{k=l} - \sum_{s=1}^J \frac{p_{k,s} p_{l,s}}{p_{\cdot s}}.$$

□

Proof of Corollary 3.6: In view of (3.2) and (3.4) we have

$$\Sigma - \Gamma = - \left(p_{m\cdot} p_{n\cdot} - \sum_{s=1}^J \frac{p_{ms} p_{ns}}{p_{\cdot s}} \right)_{m,n=1, \dots, I}.$$

We thus obtain for any $c \in \mathbb{R}^I$ that

$$\begin{aligned}
c^\top (\Sigma - \Gamma) c &= \sum_{i=1}^I \sum_{j=1}^I c_i c_j \left(\sum_{s=1}^J \frac{p_{is} p_{js}}{p_{\cdot s}} - p_{i\cdot} p_{j\cdot} \right) \\
&= \sum_{s=1}^J \frac{1}{p_{\cdot s}} \left(\sum_{i=1}^I c_i p_{is} \right)^2 - \left(\sum_{i=1}^I c_i p_{i\cdot} \right)^2 = \sum_{s=1}^J p_{\cdot s} \left(\sum_{i=1}^I c_i \frac{p_{is}}{p_{\cdot s}} \right)^2 - \left(\sum_{i=1}^I c_i p_{i\cdot} \right)^2.
\end{aligned}$$

In terms of the function $f_c: \{1, \dots, I\} \rightarrow \{c_1, \dots, c_I\}, i \mapsto c_i$ the last line can be written as

$$\begin{aligned}
 c^\top (\Sigma - \Gamma) c &= \mathbb{E}[\mathbb{E}[f_c(X_1)|Y_1]^2] - \mathbb{E}[f_c(X_1)]^2 \\
 &= \mathbb{E}[\mathbb{E}[f_c(X_1)|Y_1]^2] - \mathbb{E}[\mathbb{E}[f_c(X_1)|Y_1]]^2 \\
 (7.1) \qquad &= \text{Var}(\mathbb{E}[f_c(X_1)|Y_1]) \\
 &\geq 0,
 \end{aligned}$$

where equality holds if and only if the conditional expectation $\mathbb{E}[f_c(X_1)|Y_1]$ is constant for any function such function f_c .

It remains to verify the equivalence of $\Sigma = \Gamma$ and the independence of X_1 and Y_1 . Let X_1 and Y_1 be independent. Then $\mathbb{E}[\mathbb{E}[f(X_1)|Y_1]] = \mathbb{E}[f(X_1)]$ almost surely for any finite function $f: \{1, \dots, I\} \rightarrow \mathbb{R}$. Therefore, $c^\top (\Sigma - \Gamma) c = 0$ for any $c \in \mathbb{R}^I$, due to (7.1), and thus $\Sigma - \Gamma = 0$. For the other direction suppose (7.1) is zero for any $c \in \mathbb{R}^I$ and let $i \in \{1, \dots, I\}$ and $j \in \{1, \dots, J\}$ be arbitrary. Choosing $c \in \mathbb{R}^I$ as the vector with 1 in the i th row and all other entries zero, we obtain $\mathbb{E}[\mathbb{1}_{\{X_1=i\}}|Y_1] = P(X_1 = i)$ almost surely. We conclude

$$P(X_1 = i, Y_1 = j) = \mathbb{E}[\mathbb{1}_{\{Y_1=j\}} \mathbb{E}[\mathbb{1}_{\{X_1=i\}}|Y_1]] = P(X_1 = i)P(Y_1 = j)$$

for all $i = 1, \dots, I$ and $j = 1, \dots, J$, i.e., X_1 and Y_1 are independent. \square

ACKNOWLEDGEMENTS

We are indebted to Dr. Mirko Junge and the Volkswagen AG for providing the traffic accident data and for fruitful discussions on the case study. The authors also thank two anonymous referees for helpful comments.

REFERENCES

- Deming, W.E. and Stephan, F.F. (1940). On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *The Annals of Mathematical Statistics*, Vol. 11, No. 4, pp. 427-444.
- Fienberg, S.E. (1970): *An iterative procedure for estimation in contingency tables*. The Annals of Mathematical Statistics, Vol. 41, No. 3, pp. 907-917.
- Fisher, R.A. (1962): *Confidence limits for a cross-product ratio*. Australian Journal of Statistics, Vol. 4, p. 41.
- Goodman, L.A. (1964): *Simultaneous confidence limits for cross-product ratios in contingency tables*. Journal of the Royal Statistical Society, Series B, Vol. 26, No. 1, pp. 86-102.
- Hazell, L. and Shakir, S.A.W. (2006). Under-Reporting of Adverse Drug Reactions. *Drug Safety*, Vol. 29, No. 5, pp. 285-396.
- Lu, T.-H., Lee, M.-C. and Chou, M.-C. (2000): Accuracy of cause-of-death coding in Taiwan: types of miscoding and effects on mortality statistics *International Journal of Epidemiology*, Vol. 29, No. 2, pp. 336-343.
- Otte, D., Krettek, C., Brunner, H. and Zwipp, H. (2003): Scientific approach and methodology of a new in-depth-investigation study in Germany so called GIDAS. *Proceedings of the 18th International Technical Conference on the Enhanced Safety of Vehicles*, paper no. 0204, Nagoya (Japan).
- Pfeiffer, M. and Schmidt, J. (2006): Statistical and methodological foundations of the GIDAS accident survey system. *Proceedings of the 2nd International Conference on Expert Symposium on Accident Research*, Hanover, Germany, pp. 8187.
- Rüschendorf, L. (1995): *Convergence of the iterative proportional fitting procedure*. The Annals of Mathematical Statistics, Vol. 23, No. 4, pp. 1160-1174.
- Statistisches Bundesamt (2016): *Verkehr: Verkehrsunfälle 2014*. Statistisches Bundesamt Wiesbaden, Fachserie 8, Reihe 7 (in German), 352 pages.
- Stephan, F.F. (1942): *An iterative method of adjusting sample frequency tables when expected marginal totals are known*. The Annals of Mathematical Statistics, Vol. 13, No. 2, pp. 166-178.
- Wang, W., Rothschild, D., Goel, S. and Gelman, A. (2015): Forecasting elections with non-representative polls. *International Journal of Forecasting*, Vol. 31, No. 3, pp. 980-991.
- Yamamoto, T., Hashiji, J. and Shankar, V.N. (2008): Underreporting in traffic accident data, bias in parameters and the structure of injury severity models. *Accident Analysis and Prevention*, Vol. 40, No. 4, pp. 1320-1329.
- Ireland, C. T. and Kullback, S. (1968): Contingency tables with given marginals. *Biometrika*, Vol. 55, pp. 179-188.

Little, R. J., and Wu, M. M. (1991): Models for contingency tables with known margins when target and sampled populations differ. *Journal of the American Statistical Association*, Vol. 86, No. 413, pp. 87-95.

UNIVERSITÄT HAMBURG, FACHBEREICH MATHEMATIK, BUNDESSTRASSE 55, D-20146 HAMBURG, GERMANY. EMAIL: TOBIAS.NIEBUHR@UNI-HAMBURG.DE; MATHIAS.TRABS@UNI-HAMBURG.DE