The effect of combining categories on Bennett, Alpert and Goldstein's S

Matthijs J. Warrens *

Department of Methodology and Statistics, Tilburg University, P.O. Box 90153, 5000 LE, Tilburg, The Netherlands

ARTICLE INFO

ABSTRACT

Cohen's kappa is the most widely used descriptive measure of interrater agreement on a nominal scale. A measure that has repeatedly been proposed in the literature as an alternative to Cohen's kappa is Bennett, Alpert and Goldstein's S. The latter measure is equivalent to Janson and Vegelius' C and Brennan and Prediger's kappa$_n$. An agreement table can be collapsed into a table of smaller size by partitioning categories into subsets. The paper presents several results on how the overall S-value is related to the S-values of the collapsed tables.

It is shown that, if the categories are partitioned into subsets of the same size and if we consider all collapsed tables of this partition type, then the overall S-value is equivalent to the average S-value of the collapsed tables. This result illustrates that there are types of partitioning the categories that, on average, do not result in loss of information in terms of the S-value. In addition, it is proved that for all other partition types, the overall S-value is strictly smaller than the average S-value of the collapsed tables. A consequence is that there is always at least one way to combine categories such that the S-value increases. The S-value increases if we combine categories on which there exists considerable disagreement.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

In various fields of science, including behavioral sciences, the biomedical field and engineering sciences, it is frequently required that a group of objects (for example individuals) is rated by two or more experts on a set of mutually exclusive (nominal) categories. Examples are the psychiatric

* Tel.: +31 13 4663270; fax: +31 13 4663002.
 E-mail address: m.j.warrens@uvt.nl.

diagnosis of patients [13,48] or the classification of production faults [10,11]. Because there is often no golden standard, the reproducibility of the ratings is taken as an indicator of the quality of the category definitions and the raters' ability to apply them. Therefore, researchers typically require that the classification task is performed by at least two raters. A standard tool for measuring interrater agreement is Cohen's [8] kappa, denoted by κ [18,28,7,48,3,19,15,35,38,45]. The value of Cohen's κ is 1 when perfect agreement between the two raters occurs, 0 when agreement is equal to that expected under independence, and negative when agreement is less than expected by chance.

Although Cohen's κ has been used in thousands of applications [15], various authors have identified difficulties with kappa's interpretation [7,32,23,31,12,22,11,10]. The base rates are the marginal totals of the agreement table and reflect how often the categories were used by the raters. Because Cohen's κ is a function of the marginal totals, the measure is known to be marginal dependent or prevalence dependent [31,34,37]. A paradox associated with Cohen's κ is that, for a fixed value of the proportion of observed agreement between the raters, agreement tables with heterogeneous marginal totals (base rates) have higher κ-values than tables with homogeneous marginal totals. Hence, raters that use similar base rates are penalized compared to raters with different base rates [41]. Uebersax [32], and De Mast and Van Wieringen [11] provide several arguments as to why κ-values from samples with different base rates are not comparable. Ultimately, the problem is that many agreement tables with different marginal totals may produce the same κ-value [32, p. 144].

An agreement measure that has been proposed by various authors as an alternative to Cohen's κ is Bennett's et al. [5] S [33,25,42]. The measure is equivalent to the measure C in [17, p. 260] and the coefficient RE proposed in [16], and is denoted by κ_n in [7]. Furthermore, in the case of two categories, S is equivalent to measures discussed in, among others, Holley and Guilford [14], Maxwell [24], and Krippendorff [20]. A multi-rater version of S was proposed by Randolph [27] and Von Eye and Mun [36] (see also [10,43]). The measure S is a linear transformation of the proportion of observed agreement. Since S is not a function of the marginal totals (base rates), it does not suffer from the paradoxes associated with Cohen's κ [41]. Hsu and Field [15] argue that S has been proposed as an alternative to Cohen's κ to overcome kappa's marginal dependency. The measure is called the free-marginal kappa in [27].

In the literature, κ and S are usually presented as sample statistics, that is, as functions of the data. Landis and Koch [21], Kraemer [18], Bloch and Kraemer [6], De Mast and Van Wieringen [11] and De Mast [10] provided accounts of κ and/or S grounded in statistical modeling, making explicit the underlying premises and assumptions. De Mast [10] argues that the paradoxical behavior of Cohen's κ is explained from the fact that it is a measure of predictive association, rather than a pure measure of agreement (see also [6]). Furthermore, under the model discussed in [10,11], statistic S is the only measure of agreement that can be given some justification.

The number of categories used in various classification schemes varies from the minimum number of two to five in many practical applications. It is sometimes desirable to combine some of the categories and then calculate the κ-value or S-value of the collapsed agreement table [4]. It would be interesting to know how combining categories effects the values of κ and S. For example, when two categories are easily confused [28], the category definitions may overlap to some degree. In this case, one could ask if it is always possible to increase or decrease the κ-value or S-value by combining two categories. For Cohen's κ, the answer to this question is affirmative [44]. We will show in this paper that for the measure S there is always at least one way to combine categories such that the S-value increases. Furthermore, we will show that the S-value increases if we combine categories on which the two raters disagree considerably between themselves.

As a second example, if the scale of interest consists of too many categories, some of the categories can be combined to reduce the number of categories for the raters. The most extreme case would be to dichotomize the categorical variable. In this case, one could ask which type of partitioning of the categories results in the smallest loss or distortion of the information in terms of the S-value. We will show in this paper that if the categories are partitioned into subsets of the same size, the overall S-value is equal to the average S-value of all possible collapsed tables. In other words, it turns out that we can specify partition types that, on average, do not result in loss of information.

The paper is organized as follows. In the next section, we introduce notation and define Cohen's κ and Bennett et al. S. The main results of the paper are presented in Section 3. Numerical illustrations

of the main results are presented in Section 4. In Section 5, we present a necessary and sufficient condition for S to increase if categories are merged. Section 6 contains a discussion.

2. Cohen's kappa and Bennett et al. *S*

In this section, we introduce Cohen's κ and the measure S proposed in Bennett et al. [5]. We will study κ and S here as sample statistics, basically for notational convenience, since the results presented in Section 3 are of an algebraic nature. Alternatively, one could formulate population parameters that κ and S intend to estimate. Population parameters for κ and S can be found in, for example, DeMast [10], and DeMast and VanWieringen [11].

Suppose that two raters each independently assign the same $z \in \mathbb{N}_{\geq 1}$ objects (individuals, subjects) to the same set of $n \in \mathbb{N}_{\geq 2}$ unordered categories that are defined in advance. Let $\mathbf{F} = \{f_{ij}\}$ be a square contingency table where f_{ij} indicates the number of objects placed in category i by the first rater and in category j by the second rater $(i, j \in \{1, 2, \ldots, n\})$. We assume that the categories of the observers are in the same order, so that the diagonal elements f_{ii} of the agreement table $\mathbf{F}$ reflect the number of objects put in the same categories by both raters. For notational convenience, let $\mathbf{P}$ be the corresponding table of proportions with relative frequencies $p_{ij} = f_{ij}/z$. Row and column totals

$$p_i = \sum_{j=1}^{n} p_{ij} \quad \text{and} \quad q_i = \sum_{j=1}^{n} p_{ji}$$

are the marginal totals of $\mathbf{P}$. The marginal totals p_i and q_i are also called the base rates and they reflect how often the categories were used by raters 1 and 2, respectively.

A straightforward measure of agreement between the raters is the observed proportion of agreement

$$P = \sum_{i=1}^{n} p_{ii}.$$

There is some consensus in the literature [48,15,10,39,40,42] that the measure P should be corrected for agreement due to chance. A measure that incorporates chance-expected agreement is Cohen's κ [8]. The measure is defined as

$$\kappa = \frac{P - E}{1 - E}$$

where

$$E = \sum_{i=1}^{n} p_i q_i$$

is the proportion of agreement expected by chance alone [7,48,15]. As an example, we consider Table 8.17 in Agresti [2, p. 272] which reports the ratings by two neurologists of 149 subjects for categories (1) Certain multiple sclerosis, (2) Probable multiple sclerosis, (3) Possible multiple sclerosis, and (4) Doubtful, unlikely, or definitely not multiple sclerosis. The data were originally reported in [21]. Table 1 contains the corresponding relative frequencies of this 4×4 table. We have

$$P = 0.255 + 0.074 + 0.034 + 0.067 = 0.430$$
$$E = (0.295)(0.564) + (0.315)(0.248) + (0.235)(0.074) + (0.154)(0.114) = 0.280$$

and

$$\kappa = \frac{0.430 - 0.280}{1 - 0.280} = 0.208.$$

Bennett et al. [5] proposed the measure defined by

$$S = \frac{P - \frac{1}{n}}{1 - \frac{1}{n}} = \frac{nP - 1}{n - 1}$$

Table 1

Table of relative frequencies corresponding to Table 8.17 in Agresti [2, p. 272].

Neurologist 1	Neurologist 2				Totals
	1	2	3	4	
1	0.255	0.034	0.000	0.007	0.295
2	0.221	0.074	0.020	0.000	0.315
3	0.067	0.094	0.034	0.040	0.235
4	0.020	0.047	0.020	0.067	0.154
Totals	0.564	0.248	0.074	0.114	1.00

Table 2

Hypothetical agreement data for two raters.

Rater 1	Rater 2			Totals
	1	2	3	
1	0.10	0.00	0.05	0.15
2	0.10	0.25	0.05	0.40
3	0.05	0.05	0.35	0.45
Totals	0.25	0.30	0.45	1.00

where n is the number of categories. According to Hsu and Field [15], S is a popular alternative for Cohen's κ. Measure S is not a function of the marginal totals p_i and q_i and therefore not marginally dependent. Moreover, S is a linear transformation of the proportion of observed agreement P. The value of S is 1 when perfect agreement between the two raters occurs, and $-1/(n-1)$ when $P = 0$. For $n = 2$ categories the minimum value of S is -1. As the number of categories increases the minimum value goes to

$$\lim_{n \to \infty} \frac{-1}{n-1} = 0.$$

For the data in Table 1, we have $n = 4$ and $S = 0.239$.

Brennan and Prediger [7] argued that Cohen's κ on the one hand, and Bennett et al. S on the other hand, are appropriate in different contexts. Brennan and Prediger [7] make a distinction between studies where the marginal totals are fixed a priori or free to vary. Marginals are said to be "fixed" whenever the marginal totals (base rates) of the categories are known to the rater before classifying the objects. Brennan and Prediger [7] find Cohen's κ appropriate in reliability studies when marginal totals are fixed. When either or both of the marginal totals are free to vary, [7] proposed that κ is replaced by S.

The measure S has been criticized by various authors. Scott [29], and Zwick [48] noted that the value of S can be artificially increased by increasing the number of unused categories. In particular, since S is a decreasing function of $1/n$ (Lemma 1 in [42]), adding categories to which no objects are assigned by either rater, decreases $1/n$ and increases S. De Mast [10, p. 151] argues that this criticism seems misguided. Some other limitations of S are discussed in [15]. For example, since S does not depend on the marginal totals, its value can be large when the raters use very different base rates.

Warrens [39,42,43] showed how the values of Cohen's κ and S are related. An agreement table is called weakly marginal symmetric if the permutation that orders the marginal totals from lowest to highest is the same for the p_i and the q_i. Table 2 is an example of a weakly marginal symmetric table. Table 1 is not weakly marginal symmetric. The value of S is higher than that of Cohen's κ, that is, $S \geq \kappa$, if an agreement table is weakly marginal symmetric [42]. For the data in Table 2, we have $S = 0.55 > 0.531 = \kappa$. We also have $S \geq \kappa$ if both raters assign a certain minimum proportion of the objects to a specified category [43]. Both conditions are commonly observed in practice [42,43]. The two conditions are equivalent in the case of $n = 2$ categories [39,43]. The measure S has thus the tendency to produce higher values than Cohen's κ.

3. Main results

In this section, we present the main results. We first introduce some additional terminology and notation.

Let the $n \geq 3$ nominal categories of the agreement table be the elements of the set $C = \{c_1, c_2, \ldots, c_n\}$. A partition of C is a set of nonempty subsets of C such that every element in C is in exactly one of these subsets. Since the categories of an agreement table are nominal and the order in which the categories of a table are listed is irrelevant, combining categories of a $n \times n$ table is identical to partitioning C into $m \in \{2, 3, \ldots, n\}$ subsets. The $n \times n$ agreement table can be collapsed into a $m \times m$ table by combining categories that are in the same subset of a given partition. For $n = 4$, examples of partitions of $C = \{c_1, c_2, c_3, c_4\}$ are $\{\{c_1, c_2\}, \{c_3, c_4\}\}$ and $\{\{c_1, c_4\}, \{c_2\}, \{c_3\}\}$. The corresponding agreement tables have, respectively, sizes 2×2 and 3×3.

Let a_1 denote the number of subsets of size 1, a_2 the number of subsets of size 2, and so on, and a_n the subsets of size n. We have the identities

$$n = \sum_{i=1}^{n} i a_i = a_1 + 2a_2 + \cdots + na_n \tag{1a}$$

$$m = \sum_{i=1}^{n} a_i = a_1 + a_2 + \cdots + a_n. \tag{1b}$$

We are interested in all partitions of a certain type, that is, all partitions that there are for fixed values of $a_1, a_2, \ldots, a_n$. The type of a partition can be denoted by the $(n-1)$-tuple $(a_1, a_2, \ldots, a_{n-1})$. Note that by defining the type of a partition by a $(n-1)$-tuple instead of a n-tuple, we avoid the trivial partition with element $a_n = 1$ that combines all elements of C into a single subset. Three types of partitions for $n = 4$ categories are discussed in Section 4, namely, $(1, 0, 1)$, $(0, 2, 0)$ and $(2, 1, 0)$. The number of set partitions of C of the type $(a_1, a_2, \ldots, a_{n-1})$, that is, the number of set partitions with a_1 subsets of size 1, a_2 subsets of size 2, and so on, is given by [1, p. 823]

$$d(a_1, a_2, \ldots, a_{n-1}) = \frac{n!}{(1!)^{a_1}(a_1!)(2!)^{a_2}(a_2!)\cdots((n-1)!)^{a_{n-1}}(a_{n-1}!)}$$

$$= \frac{n!}{\prod_{i=1}^{n-1}(i!)^{a_i}\prod_{i=1}^{n-1}(a_i!)}.$$

Thus, the number of different $m \times m$ tables given a partition type $(a_1, a_2, \ldots, a_{n-1})$ of C is $d(a_1, a_2, \ldots, a_{n-1})$. For example, for $n = 4$ and partition types $(1, 0, 1)$, $(0, 2, 0)$ and $(2, 1, 0)$ we have, respectively, $d(1, 0, 1) = 4$, $d(0, 2, 0) = 3$ and $d(2, 1, 0) = 6$ (see Section 4).

If we consider all partitions ($m \times m$ tables) of a particular partition type, we may be interested in how often two categories occur together in the same subset. This number will be denoted by e. Theorem 1 gives the formula of e.

Theorem 1. *Consider all d partitions of the type $(a_1, a_2, \ldots, a_{n-1})$. The number of times two categories are in the same subset e is given by*

$$e = d\sum_{i=1}^{n-1}\frac{i(i-1)a_i}{n(n-1)}. \tag{2}$$

Proof. We first determine the number of set partitions that two categories are together in a subset of given size $i \geq 2$. This number will be denoted by e_i. It then follows that $e = \sum_{i=2}^{n-1} e_i$.

Consider an arbitrary subset of size $i \geq 2$ of the partition type $(a_1, a_2, \ldots, a_{n-1})$. For the $n - i$ categories that are not in this subset, we consider the number of partitions of the type $(a_1, \ldots, a_i - 1, \ldots, a_{n-1})$. This number is given by

$$d \cdot i! \, a_i \cdot \frac{(n-i)!}{n!}.$$

Given two fixed categories in the subset of size i, the number of ways to choose the remaining $i - 2$ categories in the subset from the remaining $n - 2$ categories is $\binom{n-2}{i-2}$. The number of partitions in which two categories occur together in a subset of size i is then given by

$$e_i = \frac{d \cdot i! a_i \cdot (n-i)!}{n!} \cdot \frac{(n-2)!}{(n-i)!(i-2)!} = d \cdot \frac{i(i-1)a_i}{n(n-1)}.$$

Using $e_1 = 0$, we obtain (2) by summing over all $i \in \{1, 2, \ldots, n-1\}$. □

The quantities d and e were also used in the proof of Theorem 1 in Warrens [46]. In that paper, no explicit formula for the number e was presented.

We are now ready to present the main result of the paper. Given a partition type $(a_1, a_2, \ldots, a_{n-1})$ of the n categories, there are d distinct $m \times m$ tables corresponding to the d partitions of this type. Theorem 2 shows that the overall S-value never exceeds the average S-value of the $m \times m$ tables.

Theorem 2. *Consider an agreement table with $n \geq 3$ categories and consider all d partitions of the type $(a_1, a_2, \ldots, a_{n-1})$. Let S denote the overall S-value of the $n \times n$ table, let P_ℓ for $\ell \in \{1, 2, \ldots, d\}$ denote the proportions of observed agreement of the $m \times m$ tables corresponding to the d partitions, and let*

$$S_\ell = \frac{P_\ell - \frac{1}{m}}{1 - \frac{1}{m}}$$

for $\ell \in \{1, 2, \ldots, d\}$. If $P < 1$, then

$$\frac{1}{d} \sum_{\ell=1}^{d} S_\ell \geq S. \tag{3}$$

Proof. We first determine the average of the P_ℓ. The proportion of observed agreement P_ℓ of a $m \times m$ table is equal to P, the proportion of observed agreement of the $n \times n$ table, plus a sum of the disagreements between the categories that are combined. If we consider all d partitions and the P_ℓ of the corresponding collapsed $m \times m$ tables, a pair of categories is combined a total of e times (Theorem 1). Hence, the average of the P_ℓ

$$\frac{1}{d} \sum_{\ell=1}^{d} P_\ell = \frac{1}{d} \left[dP + e \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} (p_{ij} + p_{ji}) \right]. \tag{4}$$

Since

$$\sum_{i=1}^{n-1} \sum_{j=i+1}^{n} (p_{ij} + p_{ji}) = \sum_{i=1}^{n} \sum_{j=1}^{n} p_{ij} - \sum_{i=1}^{n} p_{ii} = 1 - P,$$

(4) is equal to

$$\frac{1}{d} \sum_{\ell=1}^{d} P_\ell = P + \frac{e(1-P)}{d}. \tag{5}$$

Using (2) in (5), we obtain

$$\frac{1}{d} \sum_{\ell=1}^{d} P_\ell = P + \frac{1-P}{n-1} \cdot \frac{1}{n} \sum_{i=1}^{n-1} i(i-1)a_i = P + \frac{1-P}{n-1} \cdot w, \tag{6}$$

where

$$w = w(a_1, a_2, \ldots, a_{n-1}) = \frac{1}{n} \sum_{i=1}^{n-1} i(i-1)a_i.$$

Using (6), the average of the S_ℓ is then given by

$$\frac{1}{d}\sum_{\ell=1}^{d}S_\ell = \frac{m\left(\frac{1}{d}\sum_{\ell=1}^{d}P_\ell\right)-1}{m-1} = \frac{mw(1-P)+(n-1)(mP-1)}{(n-1)(m-1)}.$$

Hence, we have (3) if and only if

$$\frac{mw(1-P)+(n-1)(mP-1)}{m-1} \geq nP-1. \tag{7}$$

Using algebra it can be shown that for $1-P > 0$, inequality (7) is equivalent to the inequality $w \geq (n/m)-1$. The latter inequality can be written as

$$\frac{1}{n}\sum_{i=1}^{n-1}i^2 a_i - \frac{1}{n}\sum_{i=1}^{n-1}i a_i \geq \frac{n}{m}-1. \tag{8}$$

Using (1) and the fact that $a_n = 0$, inequality (8) is equal to

$$\left(\sum_{i=1}^{n-1}i^2 a_i\right)\left(\sum_{i=1}^{n-1}a_i\right) \geq \left(\sum_{i=1}^{n-1}i a_i\right)^2. \tag{9}$$

Finally, inequality (9) follows from using $u_i = i\sqrt{a_i}$ and $v_i = \sqrt{a_i}$ in the Cauchy–Schwarz inequality [1, p. 11]

$$\left(\sum_{i=1}^{n-1}u_i^2\right)\left(\sum_{i=1}^{n-1}v_i^2\right) \geq \left(\sum_{i=1}^{n-1}u_i v_i\right)^2. \tag{10}$$

This completes the proof. □

Theorem 2 shows that the overall S-value is equal to or smaller than the average S-value of the collapsed tables. Next, Theorem 3 specifies when the overall S-value is equivalent to the average S-value.

Theorem 3. *Consider the situation in Theorem 2. (3) is an equality if and only if $a_k = m$ for some $k \in \{1, 2, \ldots, n-1\}$ and $a_i = 0$ for $i \neq k$.*

Proof. The Cauchy–Schwarz inequality (10) becomes an equality if and only if one of the $(n-1)$-tuples $(u_1, u_2, \ldots, u_{n-1})$ and $(v_1, v_2, \ldots, v_{n-1})$ is a scalar multiple of the other [1, p. 11]. Hence, (3) becomes an equality if and only if the $(n-1)$-tuple $(\sqrt{a_1}, 2\sqrt{a_2}, \ldots, (n-1)\sqrt{a_{n-1}})$ is a multiple of $(\sqrt{a_1}, \sqrt{a_2}, \ldots, \sqrt{a_{n-1}})$. This is the case if and only if $a_k = m$ for some $k \in \{1, 2, \ldots, n-1\}$ and $a_i = 0$ for $i \neq k$. □

4. Numerical illustrations

In this section, we present numerical illustrations of Theorems 2 and 3 from Section 3. We will use the 4×4 agreement table presented in Table 1 as an example. Apart from keeping the agreement table intact or combining all categories into a single category, there are three non-trivial ways of collapsing a 4×4 table. We may combine all categories except one into a single category [13]; (for example $\{\{1, 2, 3\}, \{4\}\}$), combine 2 categories into one new category and the 2 other categories into a second new category (for example $\{\{1, 2\}, \{3, 4\}\}$), or combine 2 categories into a new category while leaving the others intact (for example $\{\{1, 2\}, \{3\}, \{4\}\}$). For each collapsed table, we may calculate the corresponding κ-value and S-value. In this section, it is shown how the κ-values and S-values of the collapsed tables are related to the overall κ-value and the overall S-value of the original 4×4 table.

An $n \times n$ agreement table can be collapsed into n distinct 2×2 tables by combining all categories except one into a single "all others" category [13]. Given a 4×4 table like the one in Table 1, there are four ways to combine all categories except one into a single category. Let $\kappa \{2, 3, 4\}$ and $S \{2, 3, 4\}$ denote the κ-value and S-value of the 2×2 table that is obtained by combining categories 2, 3 and 4. For the data in Table 1, we have

$$\kappa_1 = \kappa \{2, 3, 4\} = .337, \qquad S_1 = S \{2, 3, 4\} = .302,$$
$$\kappa_2 = \kappa \{1, 3, 4\} = -.022, \qquad S_2 = S \{1, 3, 4\} = .168,$$
$$\kappa_3 = \kappa \{1, 2, 4\} = .118, \qquad S_3 = S \{1, 2, 4\} = .517,$$
$$\kappa_4 = \kappa \{1, 2, 3\} = .424, \qquad S_4 = S \{1, 2, 3\} = .732.$$

The κ-value of the collapsed 2×2 table is an indicator of the degree of agreement for the individual category. It is called the category reliability in Fleiss [13]. Let w_1, w_2, w_3 and w_4 be the denominators of $\kappa_1, \kappa_2, \kappa_3$ and κ_4, respectively. Using the weights w_1, w_2, w_3 and w_4, we have

$$\frac{\sum_{i=1}^{4} w_i \kappa_i}{\sum_{i=1}^{4} w_i} = \frac{(0.526)(0.337) - (0.407)(0.022) + (0.274)(0.118) + (0.233)(0.424)}{0.526 + 0.407 + 0.274 + 0.233}$$

$$= 0.208 = \kappa.$$

Thus, the overall κ-value of the 4×4 table is equivalent to a weighted average of the category reliabilities, where the weights are the denominators of the 2×2 kappas [18,35]. Warrens [46] showed that given any partition type of the categories, the overall κ-value of the original table is a weighted average of the κ-values of the collapsed tables corresponding to all partitions of that type. Furthermore, we have

$$\frac{1}{4} \sum_{i=1}^{4} S_i = \frac{0.302 + 0.168 + 0.517 + 0.732}{4} = 0.430 > 0.239 = S.$$

Thus, if the categories are not partitioned into subsets of the same size, the overall S-value is strictly smaller than the average S-value of the collapsed tables.

A second possibility is that we combine only two categories into a single category while leaving the other two categories intact. By merging just two categories, the 4×4 table collapses into a 3×3 table. This can be done in six different ways. Let $S \{1, 2\}$ denote the S-value of the 3×3 table that is obtained by combining categories 1 and 2. For the data in Table 1, we have

$$S_5 = S \{1, 2\} = 0.527, \qquad S_8 = S \{2, 3\} = 0.315,$$
$$S_6 = S \{1, 3\} = 0.245, \qquad S_9 = S \{2, 4\} = 0.215,$$
$$S_7 = S \{1, 4\} = 0.185, \qquad S_{10} = S \{3, 4\} = 0.235$$

and

$$\frac{1}{6} \sum_{i=5}^{10} S_i = 0.287 > 0.239 = S.$$

Again this shows that if the categories are not partitioned into subsets of the same size, the overall S-value is strictly smaller than the average S-value of the collapsed tables.

Finally, since 2 is a divisor of 4, we may combine two categories into one new category and combine the other two categories into a second new category. This can be done in three different ways. Let $S \{1, 2\} \{3, 4\}$ denote the S-value of the 2×2 table that is obtained by combining categories 1 and 2, and 3 and 4. For the data in Table 1, we have

$$S_{11} = S \{1, 2\} \{3, 4\} = 0.490,$$
$$S_{12} = S \{1, 3\} \{2, 4\} = 0.087,$$
$$S_{13} = S \{1, 4\} \{2, 3\} = 0.141.$$

Furthermore, we have

$$\frac{1}{3} \sum_{i=11}^{13} S_i = 0.239 = S.$$

Thus, if the categories are partitioned into subsets of the same size, the overall S-value is equivalent to the average S-value of the collapsed tables. Furthermore, if the categories are partitioned into subsets of the same size, there is, on average, no loss of information in terms of the S-value. However, the above example also shows that the loss or gain of information may be substantial for a single partition (compare 0.087 and 0.490 to 0.239).

5. More on combining categories

Theorem 2 in Section 3 shows that the overall S-value is equal to or smaller than the average S-value of all collapsed tables corresponding to a certain partition type. Since an average of a set of elements is bounded by the maximum and minimum value of the elements, a direct consequence of Theorem 2 is that there is always at least one way to combine categories such that S increases. It would be interesting to know what type of categories, for example extreme or rarely used rating categories, should be combined if we want to increase S. In this section, it is shown that the S-value increases if we combine categories on which the two raters disagree considerably between themselves.

Since P is the proportion of observed agreement, we may interpret $1 - P$ as the proportion of observed disagreement. The quantity $1 - P$ is obtained by summing all off-diagonal elements of agreement table $\mathbf{P}$. Furthermore, because $\mathbf{P}$ has $n(n - 1)$ off-diagonal elements, the quantity $(1 - P)/n(n - 1)$ on the right-hand side of inequality (12) below is the average of the off-diagonal elements. In the following, we will call $(1 - P)/n(n - 1)$ the average disagreement.

The following result presents a necessary and sufficient condition for S to increase when $u \in \{1, 2, \ldots, n - 1\}$ categories are combined. The quantity U in (11) is the sum of all pairwise disagreements between the u categories.

Theorem 4. *Consider an agreement table with $n \geq 3$ categories and let S denote the corresponding S-value. Let S^* denote the S-value corresponding to the agreement table we obtain by combining $u \in \{1, 2, \ldots, n - 1\}$ categories, denoted by $i_1, i_2, \ldots, i_u$. Also define*

$$U = \sum_{j=1}^{u-1} \sum_{k=j+1}^{u} (p_{i_j i_k} + p_{i_k i_j}). \tag{11}$$

We have $S^* > S$ if and only if

$$\frac{n - u + 1}{n(u - 1)} \cdot U > \frac{1 - P}{n(n - 1)}. \tag{12}$$

Proof. If we combine u of the n categories, the proportion of observed agreement is increased by U. Furthermore, the collapsed table has $n - u + 1$ categories. We therefore have $S^* > S$ if and only if

$$\frac{P + U - \frac{1}{n-u+1}}{1 - \frac{1}{n-u+1}} > \frac{P - \frac{1}{n}}{1 - \frac{1}{n}}$$

$$\Updownarrow$$

$$\frac{(n - u + 1)(P + U) - 1}{n - u} > \frac{nP - 1}{n - 1}$$

$$\Updownarrow$$

$$(n - 1)(n - u + 1)(P + U) - (n - 1) > (nP - 1)(n - u)$$

$$\Updownarrow$$

$$(n - 1)(n - u + 1)U > (u - 1)(1 - P),$$

which is equivalent to inequality (12). $\square$

Inequality (12) can be seen as a Schouten-type inequality for the statistic S. This type of inequality was first studied by Schouten [28] for Cohen's κ. See also [47].

Although the quantity $(1 - P)/n(n - 1)$ on the right-hand side of inequality (12) is the average disagreement, the left-hand side of the inequality is more difficult to interpret. Inequality (13) below is the special case of inequality (12) for two categories.

Corollary. *Consider the situation in Theorem 4. If we combine two categories i and j we have $S^* > S$ if and only if*

$$\frac{n - 1}{n}(p_{ij} + p_{ji}) > \frac{1 - P}{n(n - 1)}. \tag{13}$$

The above corollary shows that the S-value increases if and only if $(n - 1)/n$ times the sum of the disagreements $p_{ij} + p_{ji}$ of categories i and j exceeds the average disagreement. The S-value thus increases if we combine categories on which there exists considerable disagreement. Let us illustrate the corollary using the 4×4 agreement table presented in Table 1. For Table 1, the average disagreement is

$$\frac{1 - P}{n(n - 1)} = \frac{1 - 0.430}{12} = 0.0475.$$

Furthermore, we have $n = 4$ and $n/(n - 1) = 1.333$. The corollary shows that the S-value increases if for two categories i and j the quantity $p_{ij} + p_{ji}$ exceeds the critical value $(1.333)(0.0475) = 0.0633$. Since $p_{21} = 0.221 > 0.0633$, it immediately follows that the value of S increases if we combine categories 1 and 2. Indeed, we have $S = 0.239$ and $S\{1, 2\} = 0.527$ (see Section 4). Furthermore, because $p_{31} = 0.067 > 0.0633$ and $p_{32} = 0.094 > 0.0633$ it also follows that S increases if we combine categories 1 and 3, and 2 and 3. Moreover, since $p_{12} + p_{21} > p_{23} + p_{32}$ the increase in S is more substantial if we combine categories 1 and 2. Finally, since

$$p_{14} + p_{41} = 0.007 + 0.020 = 0.027 < 0.0633$$
$$p_{24} + p_{42} = 0.000 + 0.047 = 0.047 < 0.0633$$
$$p_{34} + p_{43} = 0.040 + 0.020 = 0.060 < 0.0633,$$

the S-value decreases if we combine categories 1 and 4, 2 and 4, and 3 and 4.

6. Discussion

The most widely used descriptive measure for summarizing the cross-classification of two nominal variables with identical categories, for example, an agreement table, is the kappa statistic proposed by Cohen [8]. A measure that has been proposed by various authors as a 'better' alternative to Cohen's κ is [5] S, called $kappa_n$ in [7]. The measure S has been independently proposed by Janson and Vegelius [17] and Janes [16]. In this paper, we studied the effect of collapsing categories on the value of Bennett et al. S.

An agreement table with $n \geq 3$ categories can be collapsed into an agreement table of smaller size by combining some of the categories. Since the categories of an agreement table are nominal and the order in which the categories of a table are listed is irrelevant, combining categories of an agreement table is identical to partitioning the categories into subsets. So for each partition of the n categories into $m \in \{2, 3, \ldots, n - 1\}$ subsets there is a corresponding $m \times m$ table. In this paper, we proved that given a partition type of the categories the overall S-value of the $n \times n$ table never exceeds the average S-value of the $m \times m$ tables corresponding to all partitions of that type (Theorem 2). The overall S-value is equivalent to the average S-value if the n categories are partitioned into m subsets of equal size. In all other cases, the overall S-value is strictly smaller (Theorem 3). Numerical illustrations of Theorems 2 and 3 were presented in Section 4. In Section 5, we presented a necessary and sufficient condition for S to increase if categories are merged (Theorem 4).

The results have several practical and theoretical implications. If certain categories are easily confused or if the scale of interest consists of too many categories, it is desirable to merge some of the categories. How this affects the κ-value and S-value of the collapsed table usually depends on the data. It turns out that there exist partition types of the categories that on average do not result in loss of information as reflected in the overall S-value. If the categories are partitioned into subsets of the same size, the overall S-value is equal to the average S-value of all possible collapsed tables. However, the numerical illustrations in Section 4 do show that the loss of information can be substantial for an individual partition.

It also turns out that there is always at least one way to combine categories such that S increases. This follows from Theorem 2 and the fact that an average of a set of elements is bounded by the maximum and minimum value of the elements. The same property was proved for Cohen's κ in [44]. In Section 5, it was shown that the S-value increases if we combine categories on which there exists considerable disagreement. The S-value will increase if we combine categories on which the raters disagree the most. This is to be expected, since the statistic S is a linear transformation of the proportion of observed agreement. Thus, for increasing the S-value it is not important whether categories are rarely used or extreme rating categories, or what type of partition is used to partition the categories. We gain information in terms of the S-value if we combine categories on which the raters disagree (considerably).

With Cohen's κ there is also at least one way to combine categories such that κ decreases. This is however not the case for the measure S. As a counterexample, consider the hypothetical data for three categories in Table 2. For the data in Table 2, we have $P = 0.10 + 0.25 + 0.35 = 0.70$ and $S = 0.55$. We may combine two of the three categories into a new category and this can be done in three different ways. For all three collapsed 2×2 tables, we have $P = 0.80$ and $S = 0.60$, which illustrates that the S-value can sometimes only increase if we merge categories. Finally, since S is a decreasing function of $1/n$ and $1/n$ a decreasing function of the number of categories n, one might expect that the measure S tends to produce higher values for large tables than for smaller ones. However, Theorem 2 shows that the S-value tends to increase if we combine categories.

The results presented here show that Bennett et al. S and Cohen's κ differ in one more characteristic. A general problem with agreement measures like κ and S is that often only the extreme values (maximum and zero values) have a clear interpretation [10]. The κ-value of an agreement table with $n \geq 3$ categories can also be interpreted as a weighted average. Warrens [46] showed that given any partition type of the categories, the overall κ-value of the original table is a weighted average of the κ-values of the collapsed tables corresponding to all partitions of that type. The weights are the denominators of the kappas of the subtables. A consequence is that Cohen's κ can be interpreted as a weighted average of the κ-values of the agreement tables corresponding to all non-trivial partitions of the categories. Theorems 2 and 3 show that if the number of categories n is a composite number the S-value may be interpreted as an average. However, if n is a prime number, the overall S-value is strictly smaller than the average S-value of all collapsed tables corresponding to a particular partition type. Raju [26] derived a similar interpretation in terms of the number of items in subtests for Cronbach's [9] coefficient alpha [30]. Finally, one may be interested if there exist weights such that, similar to Cohen's κ, the overall S-value is a weighted average of the S-values of the collapsed tables. It turns out that such weights do not always exist. As a counterexample, we consider the data in Table 2. Since $S = 0.60$ for all three collapsed 2×2 tables, a weighted average

$$\frac{\sum_{i=1}^{3} w_i S_i}{\sum_{i=1}^{3} w_i} = \frac{(w_1 + w_2 + w_3)0.60}{w_1 + w_2 + w_3} = 0.60$$

where $w_i \in \mathbb{R}_{>0}$ for $i \in \{1, 2, 3\}$, always exceeds the overall S-value ($S = 0.55$).

Acknowledgments

References

[1] M. Abramowitz, I.A. Stegun, Handbook of Mathematical Functions with Formulas, Graphs and Mathematical Tables, Dover Publications, New York, 1965.
[2] A. Agresti, An Introduction to Categorical Data Analysis, John Wiley & Sons, Hoboken, New Jersey, 2007.
[3] M. Banerjee, M. Capozzoli, L. McSweeney, D. Sinha, Beyond kappa: a review of interrater agreement measures, The Canadian Journal of Statistics 27 (1999) 3–23.
[4] E. Bartfay, A. Donner, The effect of collapsing multinomial data when assessing agreement, International Journal of Epidemiology 29 (2000) 1070–1075.
[5] E.M. Bennett, R. Alpert, A.C. Goldstein, Communications through limited response questioning, Public Opinion Quarterly 18 (1954) 303–308.
[6] D.A. Bloch, H.C. Kraemer, 2×2 Kappa coefficients: measures of agreement or association, Biometrics 45 (1989) 269–287.
[7] R.L. Brennan, D.J. Prediger, Coefficient kappa: some uses, misuses, and alternatives, Educational and Psychological Measurement 41 (1981) 687–699.
[8] J. Cohen, A coefficient of agreement for nominal scales, Educational and Psychological Measurement 20 (1960) 213–220.
[9] L.J. Cronbach, Coefficient alpha and the internal structure of tests, Psychometrika 16 (1951) 297–334.
[10] J. De Mast, Agreement and kappa-type indices, The American Statistician 61 (2007) 148–153.
[11] J. De Mast, W.N. Van Wieringen, Measurement system analysis for categorical measurements: agreement and kappa-type indices, Journal of Quality Technology 39 (2007) 191–202.
[12] A.R. Feinstein, D.V. Cicchetti, High agreement but low kappa: I. The problems of two paradoxes, Journal of Clinical Epidemiology 43 (1990) 543–549.
[13] J.L. Fleiss, Statistical Methods for Rates and Proportions, Wiley, New York, 1981.
[14] J.W. Holley, J.P. Guilford, A note on the *G* index of agreement, Educational and Psychological Measurement 24 (1964) 749–753.
[15] L.M. Hsu, R. Field, Interrater agreement measures: comments on *kappa$_n$*, Cohen's kappa, Scott's π and Aickin's α, Understanding Statistics 2 (2003) 205–219.
[16] C.L. Janes, An extension of the random error coefficient of agreement to $N \times N$ tables, British Journal of Psychiatry 134 (1979) 617–619.
[17] S. Janson, J. Vegelius, On generalizations of the *G* index and the Phi coefficient to nominal scales, Multivariate Behavioral Research 14 (1979) 255–269.
[18] H.C. Kraemer, Ramifications of a population model for κ as a coefficient of reliability, Psychometrika 44 (1979) 461–472.
[19] H.C. Kraemer, V.S. Periyakoil, A. Noda, Kappa coefficients in medical research, Statistics in Medicine 21 (2002) 2109–2129.
[20] K. Krippendorff, Association, agreement, and equity, Quality and Quantity 21 (1987) 109–123.
[21] J.R. Landis, G.G. Koch, The measurement of observer agreement for categorical data, Biometrics 33 (1977) 159–174.
[22] C.A. Lantz, E. Nebenzahl, Behavior and interpretation of the κ statistic: resolution of the paradoxes, Journal of Clinical Epidemiology 49 (1996) 431–434.
[23] M. Maclure, W.C. Willett, Misinterpretation and misuse of the kappa statistic, Journal of Epidemiology 126 (1987) 161–169.
[24] A.E. Maxwell, Coefficients between observers and their interpretation, British Journal of Psychiatry 116 (1977) 651–655.
[25] G.J. Meyer, Assessing reliability. Critical corrections for a critical examination of the rorschach comprehensive system, Psychological Assesment 9 (1997) 480–489.
[26] N.S. Raju, A generalization of coefficient alpha, Psychometrika 42 (1977) 549–565.
[27] J.J. Randolph, 2005. Free-marginal multirater kappa multirater κ_{free}: an alternative to Fleiss' fixed-marginal multirater kappa, Paper presented at the Joensuu Learning and Instruction Symposium, Joensuu, Finland.
[28] H.J.A. Schouten, Nominal scale agreement among observers, Psychometrika 51 (1986) 453–466.
[29] W.A. Scott, Reliability of content analysis: the case of nominal scale coding, Public Opinion Quarterly 19 (1955) 321–325.
[30] K. Sijtsma, On the use, the misuse, and the very limited usefulness of Cronbach's alpha, Psychometrika 74 (2009) 107–120.
[31] W.D. Thompson, S.D. Walter, A reappraisal of the kappa coefficient, Journal of Clinical Epidemiology 41 (1988) 949–958.
[32] J.S. Uebersax, Diversity of decision-making models and the measurement of interrater agreement, Psychological Bulletin 101 (1987) 140–146.
[33] U.N. Umesh, R.A. Peterson, M.T. Sauber, Interjudge agreement and the maximum value of kappa, Educational and Psychological Measurement 49 (1989) 835–850.
[34] W. Vach, The dependence of Cohen's kappa on the prevalence does not matter, Journal of Clinical Epidemiology 58 (2005) 655–661.
[35] S. Vanbelle, A. Albert, Agreement between two independent groups of raters, Psychometrika 74 (2009) 477–491.
[36] A. Von Eye, E.Y. Mun, Analyzing Rater Agreement, Manifest Variable Methods, Lawrence Erlbaum Associates, 2006.
[37] A. Von Eye, M. Von Eye, On the marginal dependency of Cohen's κ, European Psychologist 13 (2008) 305–315.
[38] M.J. Warrens, On the equivalence of Cohen's kappa and the Hubert–Arabie adjusted Rand index, Journal of Classification 25 (2008) 177–183.
[39] M.J. Warrens, On similarity coefficients for 2×2 tables and correction for chance, Psychometrika 73 (2008) 487–502.
[40] M.J. Warrens, On association coefficients for 2×2 tables and properties that do not depend on the marginal distributions, Psychometrika 73 (2008) 777–789.
[41] M.J. Warrens, A formal proof of a paradox associated with Cohen's kappa, Journal of Classification 27 (2010) 322–332.
[42] M.J. Warrens, Inequalities between kappa and kappa-like statistics for $k \times k$ tables, Psychometrika 75 (2010) 176–185.
[43] M.J. Warrens, Inequalities between multi-rater kappas, Advances in Data Analysis and Classification 4 (2010) 271–286.
[44] M.J. Warrens, Cohen's kappa can always be increased and decreased by combining categories, Statistical Methodology 7 (2010) 673–677.
[45] M.J. Warrens, Weighted kappa is higher than Cohen's kappa for tridiagonal agreement tables, Statistical Methodology 8 (2011) 268–272.
[46] M.J. Warrens, Cohen's kappa is a weighted average, Statistical Methodology 8 (2011) 473–484.
[47] M.J. Warrens, A family of multi-rater kappas that can always be increased and decreased by combining categories, Statistical Methodology (2011), in press (doi:10.1016/j.stamet.2011.08.008).
[48] R. Zwick, Another look at interrater agreement, Psychological Bulletin 103 (1988) 374–378.