Precision Threshold and Noise: An Alternative Framework of Sensitivity Measures

Darren Gray(✉)

Statistics Canada, Ottawa, Canada
darren.gray@canada.ca

Abstract. At many national statistical organizations, linear sensitivity measures such as the prior-posterior and dominance rules provide the basis for assessing statistical disclosure risk in tabular magnitude data. However, these measures are not always well-suited for issues present in survey data such as negative values, respondent waivers and sampling weights. In order to address this gap, this paper introduces the Precision Threshold and Noise framework, defining a new class of sensitivity measures. These measures expand upon existing theory by relaxing certain restrictions, providing a powerful, flexible and functional tool for national statistical organizations in the assessment of disclosure risk.

Keywords: Statistical disclosure control · Linear sensitivity rules · Prior-posterior rule · pq rule · PTN sensitivity · Precision threshold · Noise

1 Introduction

Most, if not all National Statistical Organizations (NSOs) are required by law to protect the confidentiality of respondents and ensure that the information they provide is protected against statistical disclosure. For tables of magnitude data totals, established sensitivity rules such as the prior-posterior and dominance rules (also referred to as the pq and nk rules) are frequently used to assess disclosure risk. The status of a cell (with respect to these rules) can be assessed using a linear sensitivity measure of the form

$$S = \sum_r \alpha_r x_r \tag{1}$$

for a non-negative non-ascending finite input variable x_r (usually respondent contributions) and non-ascending finite coefficients α_r (determined by the choice of sensitivity rule). The cell is considered sensitive (i.e., at risk of disclosure) if $S > 0$ and safe otherwise.[1]

[1] Many NSOs have developed software to assess disclosure risk in tabular data; for examples please see [3,8]. For a detailed description of the prior posterior and dominance rules, we refer the reader to [4]; Chap. 4 gives an in-depth description of the rules, with examples. The expression of these rules as linear measures is given in [1] and [7, Chap. 6].

J. Domingo-Ferrer and M. Pejić-Bach (Eds.): PSD 2016, LNCS 9867, pp. 15–27, 2016.
DOI: 10.1007/978-3-319-45381-1_2

While powerful in their own right, these rules (and in general any sensitivity measure of the form above) were never designed to assess disclosure risk in the context of common survey issues such as negative values, respondent waivers and sampling weights. As an alternative, we introduce the Precision Threshold and Noise (PTN) framework of sensitivity measures. These measures require three input variables per respondent, which we collectively refer to as PTN variables: Precision Threshold (PT), Noise (N) and Self-Noise (SN). These variables are constructed to reflect the magnitude of protection required, and ambiguity provided, by a respondent contribution. The use of three input variables, instead of the single input variable present in linear sensitivity measures, allows for increased flexibility when dealing with survey data.

Along with these variables, the PTN sensitivity measures require two integer parameters, $n_t \geq 1$ and $n_s \geq 0$, to account for the variety of disclosure attack scenarios (or intruder scenarios) against which an NSO may wish to defend; the resulting measure is denoted $S_{n_s}^{n_t}$. In Sect. 2 we introduce S_1^1, the single target, single attacker sensitivity measure, and give a detailed definition of the PTN variables. Section 3 provides a demonstration of S_1^1 calculations, and explores other possible applications. A more detailed explanation of the parameters n_t and n_s is given in Sect. 4, along with some results on $S_{n_s}^{n_t}$ for arbitrary n_t, n_s.

2 *PTN* Pair Sensitivity

Within the PTN framework, S_1^1 is used to assess the risk of disclosure in a single target, single attacker scenario, and is referred to as PTN pair sensitivity. For a cell with two or more respondents, we assume that potential attackers have some knowledge of the size of the other contributions, in the form of upper and lower bounds. The concern is that this knowledge, combined with the publication of the cell total (and potentially other information) by the NSO may allow the attacker to estimate another respondent's contribution to within an unacceptably precise degree.

In this respect, S_1^1 can be considered a generalization of the prior-posterior rule. The prior-posterior rule (henceforth referred to as the pq rule) assumes that both the amount of protection required by, and attacker's prior knowledge of, a respondent contribution are proportional to the value of that contribution. In the PTN framework, we remove this restriction; we also allow for the possibility that attackers may not know the exact value of their own contribution to a cell total.

2.1 The Single Target, Single Attacker Premise

Let $T = \sum_r x_r$ represent the sum of respondent contributions $\{x_r\}$. We formulate a disclosure attack scenario whereby respondent s (the "suspect" or "attacker"; we use the two terms interchangeably) acting alone attempts to estimate the contribution x_t of respondent t (the "target") via the publication of total T. The suspect can derive bounds on x_t depending on their knowledge of the remainder $\sum_{r \neq t} x_r$, which includes their own contribution. Let $LB_s(\sum_{r \neq t} x_r)$

and $UB_s(\sum_{r\neq t} x_r)$ denote lower and upper bounds on this sum from the point of view of respondent s; they can then derive the following bounds on the target contribution:

$$T - UB_s\left(\sum_{r\neq t} x_r\right) \leq x_t \leq T - LB_s\left(\sum_{r\neq t} x_r\right) \tag{2}$$

Precision threshold is defined by the assumption that contribution x_t must be protected to within an interval $\left[x_t - \underline{PT}(t), x_t + \overline{PT}(t)\right]$ for some lower precision threshold $\underline{PT}(t) \geq 0$ and upper precision threshold $\overline{PT}(t) \geq 0$. The attack scenario formulated above is considered successful if this interval is not fully contained within the bounds defined in (2), in which case we refer to the target-suspect pair (t, s) as sensitive. A cell is considered sensitive if it contains any sensitive pairs, and safe otherwise.

2.2 Assumption: Suspect-Independent, Additive Bounds

To determine cell status (sensitive or safe) using (2) one must in theory determine $LB_s(\sum_{r\neq t} x_r)$ and $UB_s(\sum_{r\neq t} x_r)$ for every possible respondent pair (t, s). The problem is simplified if we make two assumptions:

1. For every respondent, there exist suspect-independent bounds $LB(r)$ and $UB(r)$ such that $LB_s(x_r) = LB(x_r)$ and $UB_s(x_r) = UB(x_r)$ for $r \neq s$.
2. Upper and lower bounds are additive over respondent sets.

Using the first assumption, we define lower noise $\underline{N}(r) = x_r - LB(x_r)$ and upper noise $\overline{N}(r) = UB(x_r) - x_r$. Let $LB_r(x_r)$ and $UB_r(x_r)$ denote bounds on respondent r's contribution from their own point of view, and define lower and upper self-noise as $\underline{SN}(r) = x_r - LB_r(x_r)$ and $\overline{SN}(r) = UB_r(x_r) - x_r$ respectively.

In many cases, it is reasonable to assume that respondents know their own contribution to a cell total exactly, in which case $LB_r(x_r) = UB_r(x_r) = x_r$ and both self-noise variables are zero; in this case we say the respondent is **self-aware**. However, we also wish to allow for scenarios where this might not hold, e.g., when T represents a weighted total and respondent r does not know the sampling weight assigned to them.

The second assumption allows us to rewrite (2) in terms of the upper and lower PTN variables; an equivalent definition of pair and cell sensitivity is then given below.

Definition 1. *For target/suspect pair (t, s) we respectively define PTN upper and lower pair sensitivity as follows:*

$$\begin{aligned} \overline{S}(t,s) &= \overline{PT}(t) - \underline{SN}(s) - \sum_{r\neq s,t} \underline{N}(r) \\ \underline{S}(t,s) &= \underline{PT}(t) - \overline{SN}(s) - \sum_{r\neq s,t} \overline{N}(r) \end{aligned} \tag{3}$$

We say the pair (t, s) *is sensitive if either* $\overline{S}(t, s)$ *or* $\underline{S}(t, s)$ *is positive and safe otherwise. Upper and lower pair sensitivity for the cell is defined as the maximum sensitivity taken over all possible distinct pairs:*

$$\begin{aligned} \overline{S}_1^1 &= max\left\{\overline{S}(t, s) \mid t \neq s\right\} \\ \underline{S}_1^1 &= max\left\{\underline{S}(t, s) \mid t \neq s\right\} \end{aligned} \tag{4}$$

Similarly, a cell is sensitive if $\overline{S}_1^1 > 0$ *or* $\underline{S}_1^1 > 0$ *and safe otherwise.*

Readers familiar with linear sensitivity forms of the pq and $p\%$ rules (see Eqs. 3.8 and 3.4 of [1]) may notice the similarity of those measures with the expressions above. There are some important differences. First, those rules do not allow for the possibility of non-zero self-noise associated with the attacker. Second, they make use of the fact that a worst-case disclosure attack occurs when the second-largest contributor attempts to estimate the largest contribution. In the PTN framework, this is not necessarily true; we show how to determine the worst-case scenario in the next section.

2.3 Maximal Pairs

Both upper and lower pair sensitivity take the form

$$S(t, s) = PT(t) - SN(s) - \sum_{r \neq s,t} N(r), \tag{5}$$

which we refer to as the general form. The general form for cell sensitivity can be similarly written as $S_1^1 = max\left\{S(t, s) \mid t \neq s\right\}$. For simplicity we will use these general forms for most discussion, and all proofs; any results on the general form apply to both upper and lower sensitivity as well. When $\underline{PT}(r) = \overline{PT}(r)$, $\underline{N}(r) = \overline{N}(r)$ and $\underline{SN}(r) = \overline{SN}(r)$ for each respondent we say that sensitivity is **symmetrical**; in this case the general form above can be used to describe both upper and lower sensitivity measures.

We define pair (t, s) as **maximal** if $S_1^1 = S(t, s)$, i.e., if the pair maximizes sensitivity within a cell. There is a clear motivation for finding maximal pairs: if both the upper and lower maximal pairs are safe, then the cell is safe as well. If either of the two are sensitive, then the cell is also sensitive.

Clearly, one can find maximal pairs (they are not necessarily unique) by simply calculating pair sensitivity over every possible pair. For n respondents, this represents $n(n-1)$ calculations (one for each distinct pair). This is not necessary, as we demonstrate below. To begin, we define target function f_t and suspect function f_s on respondent set $\{r\}$ as follows:

- Target function $f_t(r) = PT(r) + N(r)$
- Suspect function $f_s(r) = N(r) - SN(r)$

Re-arranging (5) such that the sum does not depend on (t, s) and substituting f_t and f_s gives

$$S(t, s) = f_t(t) + f_s(s) - \sum_r N(r), \tag{6}$$

which we refer to as **maximal form**. It is then clear that pair (t, s) is maximal if and only if $f_t(t) + f_s(s) = max\{f_t(i) + f_s(j) \mid i \neq j\}$.

We can find maximal pairs by ordering the respondents with respect to f_t and f_s. Let $\tau = \tau_1, \tau_2, \ldots$ and $\sigma = \sigma_1, \sigma_2, \ldots$ be ordered respondent indexes such that f_t and f_s are non-ascending, i.e., $f_t(\tau_1) \geq f_t(\tau_2) \geq \cdots$ and $f_s(\sigma_1) \geq f_s(\sigma_2) \geq \cdots$. We refer to τ and σ as **target** and **suspect orderings** respectively, noting they are not necessarily unique.

Theorem 1. *If $\tau_1 \neq \sigma_1$ (i.e., they do not refer to the same respondent) then (τ_1, σ_1) is a maximal pair. Otherwise, at least one of (τ_1, σ_2) or (τ_2, σ_1) is maximal.*[2]

The important result of this theorem is that it limits the number of steps required to find a maximal pair. Once respondents τ_1, τ_2, σ_1 and σ_2 are identified (with possible overlap), the number of calculations to determine cell sensitivity is at most two, not $n(n-1)$. By comparison, the pq rule requires only one calculation (once the top two respondents have been identified); calculating PTN pair sensitivity is at most twice as computationally demanding.

2.4 Relationship to the *pq* and *p*% Rules

The pq rule (for non-negative contributions) can be summarized as follows: given parameters $0 < p < q \leq 1$, the value of each contribution must be protected to within $p{*}100\,\%$ from disclosure attacks by other respondents. All respondents are self-aware, and can estimate the value of other contributions to within $q * 100\,\%$.

This fits the definition of a single target, single attacker scenario. The pq rule can be naturally expressed within the PTN framework using a symmetrical S_1^1 measure, and setting $PT(r) = px_r$, $N(r) = qx_r$ and $SN(r) = 0$ for all respondents. To show S_1^1 produces the same result as the pq rule under these conditions, we present the following theorem:

Theorem 2. *Suppose all respondents are self-aware. If there exists a respondent ordering $\eta = \eta_1, \eta_2, \ldots$ such that both PT and N are non-ascending, then (η_1, η_2) is maximal.*

Assuming $\{r\}$ is an ordered index such that the contributions x_r are non-ascending and applying Theorem 2 to our PTN interpretation of the pq rule, we determine that $(1, 2)$ must be a maximal pair. Then

$$S_1^1 = px_1 - \sum_{r \geq 3} qx_r,$$

which is exactly the pq rule as presented in [1], multiplied by a factor of q. (This factor does not affect cell status.)

A common variation on the pq rule is the $p\%$ rule, which assumes the only prior knowledge available to attackers about other respondent contributions is

[2] All theorem proofs appear in the Appendix.

that they are non-negative. Mathematically, the $p\%$ rule is equivalent to the pq rule with $q = 1$. Within the PTN framework, the $p\%$ rule can be expressed as an upper pair sensitivity measure $\overline{S}_1^1$ with $\overline{PT}(r) = px_r$, $\underline{N}(r) = x_r$ and $\underline{SN}(r) = 0$.

3 Pair Sensitivity Application

Having defined PTN pair sensitivity, we now demonstrate its effectiveness in treating common survey data issues such as negative values, waivers, and weights. For a good overview of the topic we refer readers to [6]; Tambay and Fillion provide proposals for dealing with these issues within G-Confid, the cell suppression software developed and used by Statistics Canada. Solutions are also proposed in [4] in a section titled *Sensitivity rules for special cases*, pp. 148–152.

In general, these solutions suggest some manipulation of the pq and/or $p\%$ rule; this may include altering the input dataset, or altering the rule in some way to obtain the desired result. We will show that many of these solutions can be replicated simply be choosing appropriate PTN variables.

3.1 S_1^1 Demonstration: Distribution Counts

To begin, we present a unique scenario that highlights the versatility of the PTN framework. Suppose we are given the following set of revenue data: {5000, 1100, 750, 500, 300}. Applying the $p\%$ rule with $p = 0.1$ to this dataset would produce a negative sensitivity value; the cell total would be considered safe for release. Should this result still apply if the total revenue for the cell is accompanied by the distribution counts displayed in Table 1? Clearly not; Table 1 provides non-zero lower bounds for all but the smallest respondent, contradicting the $p\%$ rule assumption that attackers only know respondent contributions to be non-negative.

Table 1. Revenue distribution and total revenue

Revenue range	Number of enterprises
$[0, 500)$	1
$[500, 1000)$	2
$[1000, 5000)$	1
$[5000, 10000)$	1
Total revenue:	\$7,650

The PTN framework can be used to apply the spirit of the $p\%$ rule in this scenario. We begin with the unmodified $\overline{S}_1^1$ interpretation of the $p\%$ rule given at the end of Sect. 2.4. To reflect the additional information available to potential attackers (i.e., the non-zero lower bounds), we set $\underline{N}(r) = x_r - LB(x_r)$ for each respondent, where $LB(x_r)$ is the lower bound of the revenue range containing x_r.

As the intervals $[x_r, (1+p)x_r]$ are fully contained within each contribution's respective revenue range, we leave $\overline{PT}(r)$ unchanged.

To apply Theorem 1, we calculate f_t and f_s for each respondent and rank them according to these values (allowing ties). These calculations, along with each respondent's contribution and relevant PTN variables, are found in Table 2. Applying the theorem, we determine that respondent pair $(01, 05)$ must be a maximal, giving

$$\overline{S}_1^1 = \overline{S}_1^1(01, 05) = f_t(01) + f_s(05) - \sum_r \underline{N}(r) = 150$$

and indicating that the cell is sensitive.

Table 2. Calculation of S_1^1

Respondent index	Contribution	Upper PT	Lower N	f_t	f_s	f_t rank	f_s rank
01	5000	500	0	500	0	1	4
02	1100	110	100	210	100	4	3
03	750	75	250	325	250	3	2
04	500	50	0	50	0	5	4
05	350	30	350	380	300	2	1

In addition to illustrating the versatility of the PTN framework, this example also demonstrates how Theorem 1 can be applied to quickly and efficiently find maximal pairs.

3.2 Negative Data

While the PTN variables are non-negative by definition, no such restriction is placed on the actual contributions x_r, making PTN sensitivity measures suitable for dealing with negative data. With respect to the pq rule, a potential solution consists of applying a symmetrical S_1^1 rule with $PT(r) = p|x_r|$, $N(r) = q|x_r|$ and $SN(r) = 0$ for each respondent. This is appropriate if we assume that each contribution must be protected to within $p * 100\,\%$ of its magnitude, and that potential attackers know the value of each contribution to within $q * 100\,\%$. Theorem 2 once again applies, this time ordering the set of respondents in terms of non-ascending magnitudes $\{|x_r|\}$. Then cell sensitivity S_1^1 is equal to

$$p|x_1| - \sum_{r \geq 3} q|x_r|,$$

which is exactly the pq rule applied to the absolute values. This is identical to a result obtained by Daalmans and de Waal in [2], who also provide a generalization of the pq rule allowing for negative contributions.

The assumptions about PT and N above may not make sense in all contexts. Tambay and Fillion bring up this exact point ([6, Sect. 4.3]), stating that the use of absolute values "may be acceptable if one thinks of the absolute value for a respondent as indicative of the level of protection that it needs as well as of the level of protective noise that it can offer to others" but that this is not always the case: for example, "if the variable of interest is profits then the fact that a respondent with 6 millions in revenues has generated profits of only 32,000 makes the latter figure inadequate as an indicator of the amount of protection required or provided". In this instance, they discuss the use of a proxy variable that incorporates revenue and profit into the *pq* rule calculations; the same result can be achieved within the PTN framework by incorporating this information into the construction of PT and N.

3.3 Respondent Waivers

In [6], Tambay and Fillion define a waiver as "an agreement where the respondent (enterprise) gives consent to a statistical agency to release their individual information". With respect to sensitivity calculations, they suggest replacing x_r by zero if respondent r provides a waiver. This naturally implies that the contribution neither requires nor provides protection; within the PTN framework this is equivalent to setting all PTN variables to zero, which provides the same result.

This method implicitly treats x_r as public knowledge; if this is not true, the method ignores a source of noise and potentially overestimates sensitivity. With respect to the *pq* and $p\%$ rules, an alternative is obtained by altering the PTN variables described in Sect. 2.4 in the presence of waivers: for respondents who sign a waiver, we set precision threshold to zero, but leave noise unchanged. To determine cell sensitivity, we make use of the suspect and target orderings (σ and τ) introduced in Theorem 1. In this context σ_1 and σ_2 represent the two largest contributors. If σ_1 has not signed a waiver, then it is easy to show that $\tau_1 = \sigma_1$ and (τ_1, σ_2) is maximal. On the other hand, suppose $\tau_1 \neq \sigma_1$; in this case (τ_1, σ_1) is maximal. If τ_1 has signed a waiver, then $S(\tau_1, \sigma_1) \leq 0$ and the cell is safe. Conversely, if the cell is sensitive, then τ_1 must *not* have signed a waiver; in fact they must be the largest contributor not to have done so.

In other words, if the cell is sensitive, the maximal target-suspect pair consists of the largest contributor without a waiver (τ_1) and the largest remaining contributor (σ_1 or σ_2). With respect to the $p\%$ rule, this is identical to the treatment of waivers proposed on page 148 of [4].

The following result shows that we do not need to identify τ_1 to determine cell status; we need only identify the two largest contributors.

Theorem 3. *Suppose all respondents are self-aware and that $PT(r) \leq N(r)$ for all respondents. Choose ordering η such that N is non-ascending, i.e., $N(\eta_1) \geq N(\eta_2) \geq \ldots$. If the cell is sensitive, then one of (η_1, η_2) or (η_2, η_1) is maximal.*

If $\{x_r\}$ are indexed in non-ascending order, the theorem above shows that we only need to calculate $S(1,2)$ or $S(2,1)$ to determine whether or not a cell is sensitive, as all other target-suspect pairs are safe.

3.4 Sampling Weights

The treatment of sampling weights is, in the author's opinion, the most complex and interesting application of PTN sensitivity. As this paper is simply an introduction to PTN sensitivity, we explore a simple scenario: a PTN framework interpretation of the $p\%$ rule assuming all unweighted contributions are non-negative, and all weights are at least one. We also consider two possibilities: attackers know the weights exactly, or only know that they are greater than or equal to one.

Cell total T now consists of weighted contributions $x_r = w_r y_r$ for respondent weights w_r and unweighted contributions y_r. As $LB(y_r) = 0$ for all respondents (according to the $p\%$ rule assumptions), it is reasonable that $LB(w_r y_r)$ should be zero as well, even if w_r is known. This gives $\underline{N}(r) = w_r y_r$. Self-noise is a different matter: it would be equal to zero if the weights are known, but $(w_r - 1)y_r$ if respondents only know that the weights are greater or equal to one.

Choosing appropriate precision thresholds can be more difficult. We begin by assuming the unweighted values y_r must be protected to within $p * 100\,\%$. If respondent weights are known exactly, then we suggest setting $\overline{PT}(r) = p * w_r y_r$. Alternatively, if they are not known, $\overline{PT}(r) = p * y_r - (w_r - 1)y_r$ is not a bad choice; it accounts for the fact that the weighted portion of $w_r y_r$ provides some natural protection.

Both scenarios (weights known vs. unknown) can be shown to satisfy the conditions of Theorem 2. When weights are known, the resulting cell sensitivity S_1^1 is equivalent to the $p\%$ rule applied to x_r. When weights are unknown, S_1^1 is equivalent to the $p\%$ rule applied to y_r and reduced by $\sum_r (w_r - 1)x_r$. The latter coincides with a sensitivity measure proposed by O'Malley and Ernst in [5].

Tambay and Fillion point out in [6] that this measure can have a potentially undesirable outcome: cells with a single respondent are declared safe if the weight of the respondent is at least $1 + p$. They suggest that protection levels remain constant at $p * y_r$ for $w_r < 3$, and are set to zero otherwise (with a bridging function to avoid any discontinuity around $w_r = 3$). The elegance of PTN sensitivity is that such concerns can be easily addressed simply by altering the PTN variables.

4 Arbitrary n_t and n_s

We briefly discuss the more general form of PTN sensitivity, allowing for arbitrary $n_t \geq 1$ and $n_s \geq 0$. Let T be a set of n_t respondents, and let $PT(T) \geq 0$ indicate the amount of desired protection for the group's aggregate contribution $\sum_{t \in T} x_t$. Let S be a set of n_s respondents that does not intersect T, and let

$SN(S) \geq 0$ indicate the amount of self-noise associated with their combined contribution to the total.

Suppose group S (the "suspect" group) wishes to estimate the aggregate contribution of group T (the "target" group). Expanding on the assumptions of Sect. 2.2, we will assume that PT and SN are also suspect-independent and additive over respondent sets, i.e., there exist $PT(r)$ and $SN(r)$ for all respondents such that $PT(T) = \sum_{t \in T} PT(t)$ for all possible sets T and $SN(S) = \sum_{s \in S} SN(s)$ for all possible sets S. Then we define set pair sensitivity as follows:

$$S(T, S) = \sum_{t \in T} PT(t) - \sum_{s \in S} SN(s) - \sum_{r \notin T \cup S} N(r) \tag{7}$$

Suppose we wished to ensure that *every* possible aggregated total of n_t contributions was protected against *every* combination of n_s colluding respondents. (When $n_s = 0$, the targeted contributions are only protected against external attacks.) We accomplish this by defining $S^{n_t}_{n_s}$ as the maximum $S(T, S)$ taken over all non-intersecting sets T, S of size n_t and n_s respectively. We say the set pair (T, S) is maximal if $S^{n_t}_{n_s} = S(T, S)$.

With this definition we can interpret all linear sensitivity measures (satisfying some conditions on the coefficients α_r) within the PTN framework; we provide details in the appendix. In particular the nk rule as described in Eq. 3.6 of [1] can be represented by choosing parameters $n_t = n$, $n_s = 0$ and setting $\overline{PT}(r) = ((100 - k)/k)x_r$, $\underline{N}(r) = x_r$ and $\underline{SN}(r) = 0$ for non-negative contributions x_r.

We do not present a general algorithm for finding maximal set pairs with respect to $S^{n_t}_{n_s}$ in this paper. However, we do present an interesting result comparing cell sensitivity as we allow n_t and n_s to vary:

Theorem 4. *For a cell with at least $n_t + n_s + 1$ respondents, suppose the PTN variables are fixed and that $SN(r) \leq N(r)$ for all respondents. Then the following relationships hold:*

$$S^{n_t}_{n_s} \leq S^{n_t}_{n_s+1} \leq S^{n_t+1}_{n_s} \tag{8}$$

In particular, we note two corollaries: that $S^1_0 \leq S^1_1$ and $S^1_{n_s} \leq S^{n_t}_0$ whenever $n_s \leq n_t - 1$. This demonstrates often-cited properties of the pq and nk rules: protecting individual respondents from internal attackers protects them from external attackers as well, and if a group of n_t respondents is protected from an external attack, every individual respondent in that group is protected from attacks by $n_t - 1$ (or fewer) colluding respondents.

5 Conclusion

We hope to have convinced the reader that the PTN framework offers a versatile tool in the context of statistical disclosure control. In particular, it offers potential solutions in the treatment of common survey data issues, and as we showed in Sect. 3, many of the solutions currently proposed in the statistical disclosure community can be implemented within this framework via the construction of

appropriate PTN variables. As treatments rely solely on the choice of PTN variables, implementing and testing new methods is simplified, and accessible to users who may have little to no experience with linear sensitivity measures.

Acknowledgments. The author is very grateful to Peter Wright, Jean-Marc Fillion, Jean-Louis Tambay and Mark Stinner for their thoughtful feedback on this paper and the PTN framework in general. Additionally, the author thanks Peter Wright and Karla Fox for supporting the author's interest in this field of research.

Appendix

Proof of Theorem 1

Proof. We start with the first statement, assuming $\tau_1 \neq \sigma_1$. As $f_t(\tau_1) \geq f_t(t)$ for any t and $f_s(\sigma_1) \geq f_s(s)$ for any s, it should be clear from (6) that

$$S(\tau_1, \sigma_1) \geq S(t, s)$$

for any pair (t, s), proving the first part of the theorem.

For the second part, we begin with the condition that $\tau_1 = \sigma_1$. Now, suppose (τ_1, σ_2) is not maximal. Then there exists maximal (τ_i, σ_j) where $(i, j) \neq (1, 2)$ such that $f_t(\tau_i) + f_s(\sigma_j) > f_t(\tau_1) + f_s(\sigma_2)$. As $f_t(\tau_1) \geq f_t(\tau_i)$ by definition, it follows that $f_s(\sigma_j) > f_s(\sigma_2)$ and we can conclude that $j = 1$. Then $(\tau_i, \sigma_j) = (\tau_i, \sigma_1)$ for some $i \neq 1$. But we know that $f_t(\tau_2) \geq f_t(\tau_i)$ and so $S(\tau_2, \sigma_1) \geq S(\tau_i, \sigma_1)$ for any $i \neq 1$. This shows that if (τ_1, σ_2) is not maximal, (τ_2, σ_1) must be, completing the proof. □

Proof of Theorem 2

Proof. When all respondents are self-aware, $f_t = PT + N$ and $f_s = N$, and consequently any ordering that results in non-ascending PT, N also results in non-ascending f_t, f_s. Setting $\tau = \sigma = \eta$ and applying Theorem 1, we conclude that one of (η_1, η_2) or (η_2, η_1) is maximal. From (6) we can see that

$$S(\eta_1, \eta_2) - S(\eta_2, \eta_1) = PT(\eta_1) - PT(\eta_2) \geq 0$$

showing $S(\eta_1, \eta_2) \geq S(\eta_2, \eta_1)$ and (η_1, η_2) is maximal. □

Proof of Theorem 3

Proof. The proof is self-evident for cells with two or fewer respondents, so we will assume there are at least three. Applying Theorem 1 and noting $f_s = N$ we can conclude that there exists a maximal pair of the form (η_i, η_j) for $j \leq 2$. As this pair is maximal it can be used to calculated cell sensitivity:

$$S_1^1 = S(\eta_i, \eta_j) = PT(\eta_i) - \sum_{r \neq i,j} N(\eta_r)$$

As $j \leq 2$, if $i \geq 3$ then exactly one of $N(\eta_1)$ or $N(\eta_2)$ is included in the summation above. Both of these are $\geq N(\eta_i)$ by ordering η, which is $\geq PT(\eta_i)$ by assumption. This means $S_1^1 < 0$ and the cell is safe. Conversely, if the cell is sensitive, there must exist a maximal pair of the form (η_i, η_j) with both $i, j \leq 2$, completing the proof. □

Interpreting Arbitrary Linear Sensitivity Measures in $S_{n_s}^{n_t}$ Form

All linear sensitivity measures of the form $\sum_r \alpha_r x_r$ can be expressed in PTN form, provided they satisfy the following conditions:

- Finite number of non-negative coefficients
- All positive coefficients have the same value, say α_+
- All negative coefficients have the same value, say α_-.

Assuming these conditions are met, an equivalent PTN sensitivity measure can be defined as follows:

- Set n_t equal to the number of positive coefficients
- Set n_s equal to the number of coefficients equal to zero
- Set $PT(r) = \alpha_+ x_r$ for all r
- Set $N(r) = |\alpha_-| x_r$ and $SN(r) = 0$ for all r

We show that the resulting PTN cell sensitivity measure is equivalent to $\sum_r \alpha_r x_r$ by first writing (7) as follows:

$$S(T, S) = \sum_{t \in T} (PT(t) + N(t)) + \sum_{s \in S} (N(s) - SN(s)) - \sum_r N(r) \tag{9}$$

Substituting in the appropriate PTN values gives

$$S(T, S) = \sum_{t \in T} (\alpha_+ + |\alpha_-|) x_t + \sum_{s \in S} |\alpha_-| x_s - \sum_r |\alpha_-| x_r. \tag{10}$$

It is easy to see that T and S should be selected from the largest $n_t + n_s$ respondents to maximize $S(T, S)$. If they are already indexed in non-ascending order, then sensitivity is maximized when $T = \{1, \ldots, n_t\}$ and $S = \{n_t + 1, \ldots, n_t + n_s\}$. Then cell sensitivity is given by

$$S_{n_s}^{n_t} = \sum_{r=1}^{n_t} \alpha_+ x_r - \sum_{r > n_t + n_s} |\alpha_-| x_r \tag{11}$$

which is exactly $\sum_r \alpha_r x_r$.

Proof of Theorem 4

We begin with a simple lemma:

Lemma 1. *Let T and S be non-intersecting sets of respondents. Let k be a respondent in neither, and assume $SN(k) \leq N(k)$. Then*

$$S(T,S) \leq S(T,S \cup k) \leq S(T \cup k,S). \tag{12}$$

Proof. We write (7) in maximal form, substituting in the target and suspect functions:

$$S(T,S) = \sum_{t \in T} f_t(t) + \sum_{s \in S} f_s(s) - \sum_r N(r) \tag{13}$$

Then $S(T,S \cup k) - S(T,S) = f_s(k)$. As $SN(k) \leq N(k)$ by assumption (we expect this to be true anyway, as a respondent should never know less about their own contribution than the general public), $f_s \geq 0$ proves the first inequality. The second inequality holds because $f_t \geq f_s$ for all respondents, including k. □

With this lemma, the proof of Theorem 4 is almost trivial:

Proof. Let (T,S) be maximal with respect to $S^{n_t}_{n_s}$. We know there exists at least one respondent $k \notin T \cup S$, and by Lemma 1, $S(T,S) \leq S(T,S \cup k)$, proving that $S^{n_t}_{n_s} \leq S^{n_t}_{n_s+1}$.

For the second inequality, we note that any set pair that is maximal with respect to $S^{n_t}_{n_s+1}$ can be written in the form $(T,S \cup k)$ for some T of size n_t, S of size n_s and single respondent k. Once again applying Lemma 1 we see that $S(T,S \cup k) \leq S(T \cup k,S)$ and consequently $S^{n_t}_{n_s+1} \leq S^{n_t+1}_{n_s}$. □

References

1. Cox, L.H.: Disclosure risk for tabular economic data. In: Doyle, P., Lane, J., Theeuwes, J., Zayatz, L. (eds.) Confidentiality, Disclosure and Data Access, Chap. 8. North-Holland, Amsterdam (2001)
2. Daalmans, J., de Waal, T.: An improved formulation of the disclosure auditing problem for secondary cell suppression. Trans. Data Priv. **3**(3), 217–251 (2010)
3. Hundepool, A., van de Wetering, A., Ramaswamy, R., de Wolf, P., Giessing, S., Fischetti, M., Salazar-Gonzalez, J., Castro, J., Lowthian, P.: τ-argus users manual. Version 3.5. Essnet-project (2011)
4. Hundepool, A., Domingo-Ferrer, J., Franconi, L., Giessing, S., Nordholt, E.S., Spicer, K., De Wolf, P.P.: Statistical Disclosure Control. John Wiley & Sons, Hoboken (2012)
5. O'Malley, M., Ernst, L.: Practical considerations in applying the pq-rule for primary disclosure suppressions. http://www.bls.gov/osmr/abstract/st/st070080.htm
6. Tambay, J.L., Fillion, J.M.: Strategies for processing tabular data using the g-confid cell suppression software. In: Joint Statistical Meetings, Montréal, Canada, pp. 3–8 (2013)
7. Willenborg, L., De Waal, T.: Elements of Statistical Disclosure Control. Lecture Notes in Statistics, vol. 155. Springer, New York (2001)
8. Wright, P.: G-Confid: Turning the tables on disclosure risk. Joint UNECE/Eurostat work session on statistical data confidentiality. http://www.unece.org/stats/documents/2013.10.confidentiality.html

http://www.springer.com/978-3-319-45380-4

Privacy in Statistical Databases
UNESCO Chair in Data Privacy, International Conference, PSD 2016, Dubrovnik, Croatia, September 14–16, 2016, Proceedings
Domingo-Ferrer, J.; Pejic-Bach, M. (Eds.)

2016, X, 273 p. 45 illus., Softcover
ISBN: 978-3-319-45380-4