

Identification Risks Evaluation of Partially Synthetic Data with the IdentificationRiskCalculation R Package

Ryan Hornby*, Jingchen Hu**

*Vassar College, Box 2785, 124 Raymond Ave, Poughkeepsie, NY 12604, United States.

**Vassar College, Box 27, 124 Raymond Ave, Poughkeepsie, NY 12604, United States.

E-mail: rhornby@vassar.edu, jihu@vassar.edu

Received 6 April 2021; received in revised form 12 April 2021; accepted 17 April 2021

Abstract. We extend a general approach to evaluating identification risk of synthesized variables in partially synthetic data. For multiple continuous synthesized variables, we introduce the use of a radius r in the construction of identification risk probability of each target record, and illustrate with working examples. We create the `IdentificationRiskCalculation` R package to aid researchers and data disseminators in performing these identification risks evaluation calculations. We demonstrate our methods through the R package with applications to a data sample from the Consumer Expenditure Surveys, and discuss the impacts on risk and data utility of 1) the choice of radius r , 2) the choice of synthesized variables, and 3) the choice of the number of synthetic datasets. We give recommendations for statistical agencies for synthesizing and evaluating identification risk of continuous variables.

Keywords. Identification disclosure risks, Privacy protection, Synthetic data, Utility-risk trade-off

1 Introduction

Data synthesis has been actively researched and practiced for microdata dissemination, where record-level data of individuals or business establishments, is collected and released to the public. Specifically, statistical agencies build statistical models on the original, confidential data, and generate records based on the model estimations (Rubin, 1993; Little, 1993; Raghunathan *et al.*, 2003; Reiter and Raghunathan, 2007). These generated data, called synthetic data, undergo evaluations of utility and disclosure risks, and is released to the public when the utility level and the risk level are satisfactory (Drechsler, 2011).

Many work has been done on various utility measures and evaluations, including Karr *et al.* (2006); Woo *et al.* (2009); Snoke *et al.* (2018), and there seems to be some consensus: statistical agencies should check global utility of synthetic data, where similarity between the original data distribution and the synthetic data distribution is evaluated (e.g. propensity scores metric (Woo *et al.*, 2009; Snoke *et al.*, 2018)); moreover, statistical agencies should

Variable	Type	Description
Age	Continuous	Age of reference person.
Urban	Categorical	Whether this CU located in an urban or rural area.
Tenure	Categorical	Housing tenure.
Educ	Categorical	Education level of reference person.
Expenditure	Continuous	Total expenditure last quarter.
Marital	Categorical	Marital status of reference person.
Income	Continuous	Total amount of family income before taxes in the last 12 months.

Table 1: Variables used from the CE data sample.

check analysis-specific utility of synthetic data, such as estimation of quantities of interest (e.g. mean, quantiles, and regression coefficients), and compare them between the original and the synthetic data through metric such as the interval overlap (Karr *et al.*, 2006).

By contrast, proposed approaches to identification disclosure risk evaluation in the literature are mostly specific to the applications, especially for continuous synthesized variables: Reiter (2005) uses matching-based approaches based on the most frequent and average value. Domingo-Ferrer *et al.* (2001) and Kim *et al.* (2015) calculate and report the percentages of records for which the correct link is the closest, second, or third closest match. Torra *et al.* (2006) proposes approaches of distance-based record linkage for risk evaluation of partially synthetic data. Wang and Reiter (2012) uses a Bayesian approach where they calculate the average Euclidean distance that the intruder’s prior would be away from the actual location data, leading to counting and reporting the number of other data points that are within a circle with the calculated average Euclidean distance. See Hu (2019) for an extensive review.

Reiter and Mitra (2009) proposed a general matching-based approach to evaluating the identification risks of partially synthetic data, mainly for categorical variables, as demonstrated in applications of Hu and Savitsky (2018); Drechsler and Hu (forthcoming). Hu *et al.* (forthcoming) proposed the use of radius r for identification risks evaluation of one continuous synthesized variable. In this work, we extend this radius-based approach to cases where multiple continuous variables are synthesized. For each synthesized continuous variable x_i of record i , we cast a range $R(x_i, r) = [x_i - r, x_i + r]$ around x_i with radius r , and declare a match if a record j ’s synthetic value x_j^* falls into this range, i.e. $\mathbb{I}(x_j^* \in R(x_i, r)) = 1$. Using intuitive working examples, we illustrate how the range works for cases with multiple continuous synthesized variables. Our approach of the radius parameter r shares some similarities with the swapping parameter p in data swapping proposed by Nin *et al.* (2008).

We create the `IdentificationRiskCalculation` R package to facilitate these identification risks evaluation calculations (Hornby and Hu, 2020). We demonstrate our proposed identification risk evaluation methods using our R package with applications to a data sample of the Consumer Expenditure Surveys (CE), published by the U.S. Bureau of Labor Statistics¹. Our data sample comes from the 3rd quartile in 2018. It includes sensitive continuous variables, such as income and expenditure, which will be synthesized and evaluated for identification disclosure risks; variable details are in Table 1.

In our applications, we discuss the effect of the choice of radius r on identification risk,

¹For information about the CE public-use microdata (PUMD), visit: <https://www.bls.gov/cex/pumd.htm>

the effect of variable choices on the utility-risk trade-off, and the effect of the number of synthetic datasets on the utility-risk trade-off. In all applications, we use CART models to perform data synthesis with the default settings in the `synthpop` package (Nowok *et al.*, 2016).

The remainder of the article is organized as follows. In Section 2 we describe our proposed identification risk evaluation methods of synthesized variables of different data types, with definition, illustrations, and discussions of implications. Section 3 presents information about the `IdentificationRiskCalculation` R package. Section 4 presents our applications to the CE data sample, with three subsections focusing the effects of radius r , the effects of synthesized variables, and the effects of the number of synthetic datasets. We end with a few concluding remarks in Section 5.

2 Identification Risk Evaluation of Synthesized Variables

The main evaluation approach to identification risk evaluation of partially synthetic data is proposed by Reiter and Mitra (2009), where three summaries of identification risk measures are commonly used: 1) the expected match risk, 2) the true match rate, and 3) the false match rate. Most applications of this approach to identification risk evaluation only consider categorical variables (Hu and Hoshino, 2018; Hu and Savitsky, 2018; Drechsler and Hu, forthcoming), with the exception of Hu *et al.* (forthcoming), who considers one continuous variable.

We follow the general approach in Reiter and Mitra (2009) and focus on the expected match risk as the identification risk. Specifically, we define the record-level identification risk, IR_i , as:

$$IR_i = \frac{T_i}{c_i}, \quad (1)$$

where c_i is the number of records with the highest match probability for the target record i , and $T_i = 1$ if the target is among the c_i number of matched records, and 0 otherwise. When $m > 1$ synthetic datasets are simulated, $IR_i^{(l)}$ should be calculated for each synthetic dataset $l = 1, \dots, m$ and the average $IR_i = 1/m IR_i^{(l)}$ can be reported.

The calculation of c_i and T_i depends on the variable type of synthesized variables, which we will introduce in detail in each subsection with an illustrative example. Overall, the expression of IR_i in Equation (1) captures the probability of record i being correctly identified, and this setup ensures that $IR_i \in [0, 1]$. We can also obtain a file-level summary of identification disclosure risk of n records as below, a summary we use in our applications in Section 4:

$$IR = \sum_{i=1}^n IR_i. \quad (2)$$

We note that for microdata undergone other Statistical Disclosure Control (SDC) procedures, such as masking, record-level identification risk IR_i defined in Equation (1) can also be used, as long as the data disseminators have an approach to determining records with the highest match probability for record i , from which c_i and T_i are calculated. Nevertheless, the general approach originally proposed by Reiter and Mitra (2009) relies on Bayesian methodology of setting up the risk estimation process that is specific to partially synthetic data.

2.1 Categorical variables

For categorical variables, to obtain the set of records with the highest match probability, we assume the intruder has access to a combination of known, unsynthesized variables. For simplicity, consider all known, unsynthesized variables as categorical, though our proposed method generalizes to continuous known, unsynthesized variables in a straightforward manner.

We assume the intruder knows the *true confidential* values of the synthesized categorical variables of the target record i —it is reasonable to assume an intruder trying to identify a record in the synthetic data with the knowledge of the true confidential values of the synthesized variables. Equipped with information of the known, unsynthesized variables *and* the true values of the synthesized categorical variables of record i , the intruder will then search for records in the synthetic data, who share the same available information with the target record i . Formally, we compute c_i by:

$$c_i = \sum_{j=1}^n K_i(j)S_i(j), \quad (3)$$

where $K_i(j)$ and $S_i(j)$ are binary indicators. $K_i(j) = 1$ if the known, unsynthesized variables of record j are the same as the target record i , and 0 otherwise. Similarly, $S_i(j) = 1$ if the synthesized categorical variables of record j are the same as the target record i , and 0 otherwise.

After finding c_i matched records, the intruder proceeds to figure out which of the matched records is the target record i , and we need to evaluate how likely a correct match happens. On the one hand, if synthesized categorical variables of target record i result in $S_i(i) = 1$ (i.e. $T_i = 1$), then the target record is among the matched records. In this case, the intruder will simply randomly guess which of the c_i records is i , therefore the identification risk probability is $IR_i = 1/c_i$. On the other hand, if synthesized categorical variables of target record i result in $S_i(i) = 0$ (i.e. $T_i = 0$), then the target record is not among the matched records. In this case, the intruder will have a 0 probability of finding the true identify i , therefore $IR_i = 0/c_i = 0$.


Given this basic setup, we now turn to our identification disclosure risk evaluation of synthesized continuous variables, where we modify the approach to adapt to their continuous nature. Specifically, instead of looking for an exact match as in the categorical case, we assume matching within a distance, defined through a radius r .

2.2 One continuous variable

We start with one synthesized continuous variable. As before, we assume the intruder knows the *true confidential* value of the synthesized variable, x_i of record i , when trying to find the identity of this record. We define a range $R(x_i, r) = [x_i - r, x_i + r]$ around x_i with a radius r . We then define c_i , the number of records with the highest match probability for target record i as:

$$c_i = \sum_{j=1}^n K_i(j)\mathbb{I}(x_j^* \in R(x_i, r)), \quad (4)$$

where $K_i(j)$ is the binary indicator of whether record j shares the same known variables as record i . x_j^* denotes the synthetic value of record j , and $\mathbb{I}(\cdot)$ is a binary indicator, where $\mathbb{I}(x_j^* \in R(x_i, r)) = 1$ if $x_j^* \in [x_i - r, x_i + r]$, and 0 otherwise.



As can be seen in our setup, we do not perform an exact match for continuous x_j^* as $x_j^* = x_i$. Rather, we perform matching based on a range of value of x_i , specifically, $[x_i - r, x_i + r]$. We recommend a percentage radius r to reflect the magnitude of x_i . For example, $r = 20\%$ results in a range of $[x_i - 0.2x_i, x_i + 0.2x_i] = [0.8x_i, 1.2x_i]$.

We now illustrate our approach with two examples of the CE data sample, following the legends above. We assume the continuous variable income is synthesized. Our matching for record i is based on the range $R(\text{Income}_i, r)$, which is indicated by the two vertical blue bars in the two cases in Figure 1 and Figure 2. As can be seen, for one continuous variable, the range results in an interval.




Figure 1: Case 1 - $IR_i = T_i/c_i = 1/5$.



Figure 2: Case 2 - $IR_i = T_i/c_i = 0/4$.

For case 1 in Figure 1, with 5 synthetic records inside the interval, $c_i = 5$. Moreover, the synthetic income of record i (the square) is in the interval, $T_i = 1$, resulting in $IR_i = T_i/c_i = 1/5$. For case 2 in Figure 2, with 4 synthetic records inside the interval, $c_i = 4$. However, in this case the synthetic income of record i (the square) is *not* in the interval, $T_i = 0$, resulting in $IR_i = T_i/c_i = 0/4 = 0$.

2.3 Two continuous variables

We now move to two synthesized continuous variables, and the *true confidential* values are x_i and y_i , respectively, for record i . We therefore have two ranges, $R(x_i, r) = [x_i - r, x_i + r]$ and $R(y_i, r) = [y_i - r, y_i + r]$. We then define c_i , the number of records with the highest match probability for target record i as:

$$c_i = \sum_{j=1}^n K_i(j) \mathbb{I}(x_j^* \in R(x_i, r) \cap y_j^* \in R(y_i, r)), \tag{5}$$

where $K_i(j)$ is the binary indicator of whether record j shares the same known variables as record i . x_j^* and y_j^* denote the synthetic values of record j of these two continuous variables. As before, $\mathbb{I}(\cdot)$ is a binary indicator, where $\mathbb{I}(x_j^* \in R(x_i, r) \cap y_j^* \in R(y_i, r)) = 1$ if $x_j^* \in [x_i - r, x_i + r]$ and $y_j^* \in [y_i - r, y_i + r]$, and 0 otherwise. That is, we declare a match only when the synthetic values of *both* variables of record j fall into their associated ranges. We recommend a percentage radius r as in Section 2.2, and one can use different percentage values for different continuous variables.

We now illustrate our approach with two examples of the CE data sample, following the same legends in Section 2.2. We assume income and expenditure are the two synthesized

continuous variables. Therefore, our matching for record i is based on the two ranges $R(\text{Income}_i, r)$ and $R(\text{Expenditure}_i, r)$, which is indicated by the blue rectangle in Figure 3 and Figure 4. As can be seen, for two continuous variables, the combination of two ranges forms a rectangle.

For case 3 in Figure 3, with 4 synthetic records inside the rectangle, $c_i = 4$. Moreover, the synthetic income *and* expenditure of record i are in their corresponding ranges, so that record i (the square) is *in* the rectangle, $T_i = 1$, resulting in $IR_i = T_i/c_i = 1/4$. For case 4 in Figure 4, with 3 synthetic records inside the rectangle $c_i = 3$. However, in this case the synthetic income is in its range but the synthetic expenditure is not, resulting record i (the square) being *outside* of the rectangle. Therefore, $T_i = 0$ resulting in $IR_i = T_i/c_i = 0/3 = 0$.




Figure 3: Case 3 - $IR_i = T_i/c_i = 1/4$.




Figure 4: Case 4 - $IR_i = T_i/c_i = 0/3$.




Figure 5: Case 5 - $IR_i = T_i/c_i = 1/3$.

Our proposed identification risk evaluation methods can be extended to more than two continuous variables in a straightforward manner. For example, when working with three continuous synthesized variables, the combination of three ranges forms a cuboid.

It is worth illustrating how our identification risk evaluation setup adapts to cases where additional categorical variables are synthesized. We introduce case 5 in Figure 5, which compared to case 3 in Figure 3, an additional categorical variable, for example tenure from the CE sample, is synthesized. As can be seen in the differences between case 3 and case 5, the effects of synthesizing an additional categorical variable are that fewer synthetic records remain matched overall, because we have one more variable to match with. The

effects might impact both T_i and c_i . In the illustrative example of case 5, we have $IR_i = T_i/c_i = 1/3$.

3 The `IdentificationRiskCalculation` R package

We have created an R package, `IdentificationRiskCalculation`, to facilitate the identification risk evaluation computation of our proposed methods (Hornby and Hu, 2020). In this section we describe how to use the function `IdentificationRisk()` and what outputs it produces. The package can be installed with the following R code:

```
library(devtools)
install_github("RyanHornby/IdentificationRiskCalculation")
```

The `IdentificationRisk()` function in this R package computes the identification risks for *all* records in the confidential dataset. Its outputs include matrices listing the values for c_i , T_i , and IR_i for each record (row) and each synthetic dataset (column). In addition, vectors of the true and false match rates are produced, if users are interested in reporting them (Reiter and Mitra, 2009).

Below is an example call of the `IdentificationRisk()` function:

```
IdentificationRisk(origdata,
                  syndata,
                  known,
                  syn,
                  r,
                  percentage = true,
                  euclideanDist = false)
```

The first required argument is `origdata` which is a dataframe of the confidential dataset. The next argument is `syndata` which is a list of m synthetic datasets. The next two arguments, `known` and `syn`, are both vectors. These vectors contain the names of the columns corresponding to the known and synthesized variables respectively. The final required argument is `r` which is radius to evaluate the continuous variables at. This radius can be either one value or a vector of length equal to the total number of continuous variables in the known and synthetic vectors.

There are a couple of optional arguments that work with this radius. First is `percentage`, which by default is set to `true` such that the radius value(s) are calculated based on the confidential values. Next is `euclideanDist`, which by default is set to `false` such that each radius is treated independently leading to a rectangle in the 2D case instead of an ellipse, if there are two synthesized continuous variables.

We now proceed to demonstrate the usage of this R package with CE applications.

4 Applications to the CE Data Sample

In this section, we present three applications with four scenarios to the CE data sample to demonstrate our identification risk evaluation methods, available in our R package. We use the `synthpop` R package to perform data synthesis (Nowok *et al.*, 2016). Specifically, we use CART synthesis models to synthesize several variables in the CE data sample, some

are continuous and some are categorical (Reiter, 2005). All three applications are partially synthetic data, where a subset of variables are synthesized (Little, 1993). This CE sample is available in our `IdentificationRiskCalculation` R package.

In all three applications, we assume the intruder knows the values of {Age, Urban, Marital} of each record in the CE data sample. These variables are not synthesized in any application. Note that Age is a continuous variable, and we use the percentage radius method in Section 2.2 for its matching. Given these assumptions, we start with the following R code:

```
CEdata <- IdentificationRiskCalculation::CEdata
CEdata$Urban <- as.factor(CEdata$Urban)
CEdata$Marital <- as.factor(CEdata$Marital)
CEdata$Tenure <- as.factor(CEdata$Tenure)

knownvars <- c("Age", "Urban", "Marital")
r_age <- 0.1
```

Each of the four scenarios is associated with a set of synthesized variables. Table 2 lists the detail of each scenario, including the nature (continuous vs categorical; “con” stands for continuous and “cat” stands for categorical) of each variable and the synthesis order (from left to right).

Scenario	Synthesized variables
(1)	Income (con)
(2)	Tenure (cat), Income (con)
(3)	Expenditure (con), Income (con)
(4)	Tenure (cat), Expenditure (con), Income (con)

Table 2: Detail of the four scenarios of synthesized variables.

We provide sample R code for synthesizing with the `synthpop` package and calculating the risks with our package for each of these scenarios using the following code, starting with scenario 1:

```
synvars1 <- c("Income")
syndata1 <- synthpop::syn(CEdata,
                          m = 20,
                          visit.sequence = synvars1)
r_income <- 0.1
riskList1 <- IdentificationRisk(origdata = CEdata,
                               syndata = syndata1$syn,
                               known = knownvars,
                               syn = synvars1,
                               r = c(r_age, r_income))
IR1 <- riskList1$exp.risk_vector
```

Next for scenario 2:

```
synvars2 <- c("Tenure", "Income")
syndata2 <- synthpop::syn(CEdata,
                          m = 20,
```



```

                                visit.sequence = synvars2)
r_income <- 0.1
riskList2 <- IdentificationRisk(origdata = CEdata,
                               syndata = syndata2$syn,
                               known = knownvars,
                               syn = synvars2,
                               r = c(r_age, r_income))
IR2 <- riskList2$exp.risk_vector

```

Now for scenario 3:

```

synvars3 <- c("Expenditure", "Income")
syndata3 <- synthpop::syn(CEdata,
                          m = 20,
                          visit.sequence = synvars3)
r_income <- 0.1
r_expenditure <- 0.1
riskList3 <- IdentificationRisk(origdata = CEdata,
                               syndata = syndata3$syn,
                               known = knownvars,
                               syn = synvars3,
                               r = c(r_age, r_expenditure,
                                     r_income))
IR3 <- riskList3$exp.risk_vector

```

Finally for scenario 4:

```

synvars4 <- c("Tenure", "Expenditure", "Income")
syndata4 <- synthpop::syn(CEdata,
                          m = 20,
                          visit.sequence = synvars4)
r_income <- 0.1
r_expenditure <- 0.1
riskList4 <- IdentificationRisk(origdata = CEdata,
                               syndata = syndata4$syn,
                               known = knownvars,
                               syn = synvars4,
                               r = c(r_age, r_expenditure,
                                     r_income))
IR4 <- riskList4$exp.risk_vector

```

Note that in all the sample R code in this section, we use a radius of 10%, however care should be taken in choosing a radius as we explore in Section 4.1. Also, to adequately explore our proposed identification risk evaluation methods and avoid the impact of high variability when simulating only $m = 1$ synthetic dataset, we simulate $m = 20$ synthetic datasets in all three applications except for the last one. In the third application in Section 4.3, we evaluate the choice of m , the number of synthetic datasets and its impact on identification disclosure risk and utility.

4.1 Effect of radius r on identification risk

Our identification risk evaluation approach for synthesized continuous variables in Section 2.2 and Section 2.3 relies on the choice of radius r . In our first application, we explore the choice of r and its impact on the identification risk results.

Figure 6 illustrates the effect of $r = \{1\%, 2.5\%, 5\%, 10\%, 20\%, 30\%\}$ on the identification risk. We consider all four scenarios of synthesized variables listed in Table 2. In each scenario, we evaluate the file-level identification risk in Equation (2) with a choice of r for each of the $m = 20$ synthetic datasets. Each boxplot in Figure 6 represents the distribution of the file-level identification risk across the $m = 20$ synthetic datasets under that scenario.




Figure 6: Effects of radius r on identification risk in four scenarios.

Figure 6 shows in each scenario, there exists one r value that maximizes the identification risk. For example, in scenario (1) where only income is synthesized, $r = 2.5\%$ is the radius value that expresses highest identification risk. After the peak as radius r increases, on the one hand, records with $T_i = 1$ will remain with $T_i = 1$ (we are using an wider interval with a larger r , and any value in a shorter interval will remain in a wider interval when both intervals have the same center). However, as r increases, c_i increases as well, since a wider interval will catch more synthetic records, making the record-level identification risk T_i/c_i to decrease. On the other hand, records with $T_i = 0$ might flip to $T_i = 1$ as the interval becomes wider, although the impact of increasing c_i remains. Similar discussion can be made about before the peak radius value.

Overall, in each scenario we can identify a radius r that maximizes the identification risk of the simulated synthetic datasets. For scenarios (1) and (2), the maximizing r is 2.5%, while for scenarios (3) and (4), the maximizing r is 10%. We recommend the practice of

finding the maximizing radius value r for every scenario considered in real data applications. Moreover, we recommend that results should be reported based on the maximum identification risk, as those are most conservative: if identification risk is acceptable at its maximum with a r value, it will be acceptable at any other r values with lower risks.

For illustration purpose, we use the same value of r for each continuous variable in scenarios where more than one continuous variables are synthesized (e.g. scenario (3) and (4)). In practice, statistical agencies can use different combinations of percentage radius values to explore the maximizing radius combinations, e.g. r_{income} for income and $r_{\text{expenditure}}$ for expenditure in scenarios (3) and (4).

4.2 Effects of synthesized variables on identification risk and utility

In addition to illustrating the maximizing radius r for each scenario, Figure 6 also shows that as more variables are synthesized, the identification risk of simulated synthetic datasets will decrease. Since different scenarios reach their maximum identification risks with different radius r values, we present Figure 7 where boxplots of identification risk distributions for each scenario at its maximizing radius r are plotted and compared.




Figure 7: Identification risk of four scenarios of synthesized variables.

Figure 7 shows that as more variables are synthesized, identification risk decreases, a result has been demonstrated in many applications, including Reiter and Mitra (2009); Drechsler and Hu (forthcoming). In fact, synthesizing more variables is a typical strategy when statistical agencies need a higher privacy protection.

What is interesting here is that between scenarios (2) and (3), both synthesize two variables, the effects of the type of variables to be synthesized on identification risk are quite different: synthesizing expenditure, a continuous variable, will provide a higher level of privacy protection compared to synthesizing tenure, a categorical variable. Note that we

use the maximizing radius r , suggesting that synthesizing an additional continuous variable indeed provides a higher level of privacy protection. Moreover, if we compare scenarios (3) and (4), given synthesized income and expenditure, to additionally synthesize tenure does not provide much further privacy protection. The biggest risk reduction comes with the choice of synthesizing a continuous variable.

It is well known that a higher privacy protection does not come for free. As risk decreases, the usefulness of the synthetic data gets compromised, a phenomenon known as the utility-risk trade-off of synthetic data (Duncan and Stokes, 2012; Drechsler and Hu, forthcoming; Hu *et al.*, forthcoming; Savitsky *et al.*, 2020). To illustrate in our applications, we consider the global utility measure of propensity scores, outlined in Woo *et al.* (2009) and Snoke *et al.* (2018):

$$U_p = \frac{1}{2n} \sum_{i=1}^{2n} \left(\hat{p}_i - \frac{1}{2} \right)^2, \quad (6)$$

where n is the number of records in the original dataset (and each generated synthetic dataset in our application), and \hat{p}_i is the propensity score estimated from a classification model for record i belonging to the synthetic dataset. Smaller and closer-to-0 values of U_p indicate high utility, since the estimated propensity score for a large number of combined records is around $1/2$, meaning that the chosen classification model cannot differentiate the original dataset and the synthetic dataset. In our evaluations, we use logistic regression models for classification.

Figure 8 shows the utility-risk trade-off of scenarios (1), (2), and (4), and Figure 9 shows that of scenarios (1), (3), and (4). In each figure, we plot the identification risk and propensity score of each of the $m = 20$ synthetic datasets generated under the three scenarios. Moreover, in each scenario, we add a rectangle showing the 1st quartile and the 3rd quartile on both dimensions: utility on the x-axis and identification risk on the y-axis. The length and the height of each rectangle show the variability of each dimension across $m = 20$ synthetic datasets.




Figure 8: Utility-risk trade-off in scenarios (1), (2), and (4).




Figure 9: Utility-risk trade-off in scenarios (1), (3), and (4).

In each figure, we confirm the utility-risk trade-off: when we synthesize more variables, for example from scenario (1) to scenario (2) in Figure 8, identification risk decreases and

the utility decreases (increased propensity scores indicate lower utility). Similar observations can be made for any pair of scenarios in either figure.

However, comparing scenario (2) and scenario (3) across Figure 8 and Figure 9, we observe a case for synthesizing an additional continuous variable, as in scenario (3), over an additional categorical variable, as in scenario (2). The choice of the two continuous variables being synthesized (income and expenditure) not only expresses lower identification risk, as we have seen in Figure 7, but also higher utility: the propensity score of scenario (3) are substantially lower than that of scenario (2), indicating higher utility. Moreover, the results of utility and risk express smaller variability across synthetic datasets in scenario (3), evident in the smaller size of the rectangle, on both dimensions. This suggests the robustness of our proposed radius-based matching strategy of identification risk evaluation for continuous variables.

4.3 Effects of number of synthetic datasets on identification risk and utility

The synthetic data literature has traditionally recommended the release of $m > 1$ synthetic datasets, mainly from the utility perspective. Releasing multiple synthetic datasets provides stabilized point estimates and standard error estimates of quantities of interest by accounting for the uncertainty in the synthesis process. To obtain valid inferences for quantities of interest, data users should follow appropriate combining rules (Raghunathan *et al.*, 2003; Drechsler, 2011). More recently, recommendations of releasing $m = 1$ synthetic dataset have been made, where synthetic datasets not simulated from posterior predictive distributions (Klein and Sinha, 2015; Raab *et al.*, 2016). New variance estimate formula are therefore proposed to obtain valid inferences. From the disclosure risk perspective, releasing fewer number of synthetic datasets has also been demonstrated in practice to have lower disclosure risks (Reiter and Mitra, 2009; Kinney *et al.*, 2011; Klein and Sinha, 2015).

To evaluate the effects of m on identification risk and utility of continuous synthesized variables, we focus on scenario (1) where income is synthesized, and experiment with $m = \{1, 10, 20\}$. For each value of m , we repeat the synthetic data generation process for 1000 times, and take the average of identification risks and the average of propensity score utilities across m synthetic datasets, and visualize them in Figure 10. We propose the use of what we call a “2D-boxplot” where for a given m , the sides of the rectangle represent the 1st quartile and the 3rd quartile across the 1000 simulations. Moreover, the lines sticking out from the rectangle end at the minimum and the maximum values across the 1000 simulations.

Our risk results resonate with the aforementioned works recommending $m = 1$ for higher privacy protection, as the average identification risk is the smallest for $m = 1$, and increases as m increases to $m = 10$ and $m = 20$. Figure 10 also shows that the median identification risk is the smallest for $m = 1$, and that of $m = 10$ is slightly smaller than that of $m = 20$.

However, what Figure 10 also reveals is that there is a lot more variability across the identification risk evaluations when $m = 1$, or more generally, the variability increases as m decreases. As evident in the size of the three rectangles in Figure 10, we could get an outcome of identification risk as small as 0.0076, and we could also get an outcome as large as 0.0126, while the minimum and the maximum values are 0.0087 and 0.0011 for $m = 10$ and 0.0090 and 0.0108 for $m = 20$. These suggest that while $m = 1$ produces the smallest average identification risk (i.e. highest privacy protection), it could severely overestimate and / or underestimate the identification risk outcome given its large spread of values.




Figure 10: 2D-boxplot of utility-risk trade-off in scenario (1) with different values of m .

The utility results present a similar picture: $m = 1$ produces the smallest average propensity score, indicating highest utility level, at the cost of a larger spread. In fact, if we only rely on the average identification risk and the average propensity score, $m = 1$ is the best choice with lowest risk and highest utility. However, our use of a 2D-boxplot highlights that using $m = 1$ could present an inaccurate picture of the utility-risk trade-off of generated synthetic datasets. We also experiment with the other three scenarios in Table 2 and reach the same conclusion.

5 Concluding Remarks

In this article, we extend a general approach to evaluating the identification risk of multiple continuous synthesized variables. We cast a range with a radius r around the true value of each synthesized continuous variable, and declare a match based on a record falling into that range. Our applications to the CE data sample recommend the practice of finding and reporting based on the maximizing radius r of identification risk.

We have created and made public the `IdentificationRiskCalculation` R package. Interested researchers and practitioners are encouraged to implement the general approach to identification risk evaluation for categorical variables, and try out and experiment with our proposed methods for multiple continuous synthesized variables.

We have seen that compared to synthesizing categorical variables, synthesizing continuous variables can be efficient at reducing identification risk with minimal impacts on the data utility. We also saw that generating $m > 1$ synthetic datasets achieves a more accurate picture of the utility-risk trade-off of generated synthetic datasets. We believe future research, such as using different synthesis models, different datasets, and different utility measures, are valuable.

Another type of disclosure is attribute disclosure, referring to the intruder accurately guesses the confidential values of the synthetic variables given the released synthetic data (Hu, 2019). For readers interested in Bayesian estimation of attribute disclosure risks in synthetic data generated using Bayesian synthesis models, we have a manuscript (Hornby and Hu, 2021b) providing computational details with an R package (Hornby and Hu, 2021a) for further reading.

References

- Domingo-Ferrer, J., Mateo-Sanz, J., and Torra, V. (2001). Comparing SDC methods for microdata on the basis of information loss and disclosure risk. *Pre-proceedings of ENKNTTS* 807–826.
- Drechsler, J. (2011). *Synthetic Datasets for Statistical Disclosure Control*. Springer: New York.
- Drechsler, J. and Hu, J. (forthcoming). Synthesizing geocodes to facilitate access to detailed geographical information in large scale administrative data. *Journal of Survey Statistics and Methodology*.
- Duncan, G. T. and Stokes, S. L. (2012). Disclosure risk vs. data utility: The R-U confidentiality map as applied to topcoding. *CHANCE* **17**, 16–20.
- Hornby, R. and Hu, J. (2020). *IdentificationRiskCalculation: Calculating the Identification Risk in Partially Synthetic Microdata*. R package, <https://github.com/RyanHornby/IdentificationRiskCalculation>.
- Hornby, R. and Hu, J. (2021a). *AttributeRiskCalculation: Calculating Attribute Disclosure Risks Synthetic Microdata*. R package, <https://github.com/RyanHornby/AttributeRiskCalculation>.
- Hornby, R. and Hu, J. (2021b). Bayesian estimation of attribute disclosure risks in synthetic data with the `AttributeRiskCalculation` R package arXiv:2103.09805.
- Hu, J. (2019). Bayesian estimation of attribute and identification disclosure risks in synthetic data. *Transactions on Data Privacy* **12**, 61–89.
- Hu, J. and Hoshino, N. (2018). The Quasi-Multinomial synthesizer for categorical data. In J. Domingo-Ferrer and F. Montes, eds., *Privacy in Statistical Databases*, vol. 11126 of *Lecture Notes in Computer Science*, 75–91. Springer.
- Hu, J. and Savitsky, T. D. (2018). Bayesian data synthesis and disclosure risk quantification: An application to the Consumer Expenditure Surveys. arxiv: 1809.10074.
- Hu, J., Savitsky, T. D., and Williams, M. R. (forthcoming). Risk-efficient Bayesian data synthesis for privacy protection. *Journal of Survey Statistics and Methodology*.
- Karr, A. F., Kohnen, C. N., Oganian, A., Reiter, J. P., and Sanil, A. P. (2006). A framework for evaluating the utility of data altered to protect confidentiality. *The American Statistician* **60**, 224–232.
- Kim, H. J., Karr, A. F., and Reiter, J. P. (2015). Statistical disclosure limitation in the presence of edit rules. *Journal of Official Statistics* **31**, 1, 121–138.

- Kinney, S. K., Reiter, J. P., Reznek, A. P., Miranda, J., Jarmin, R. S., and Abowd, J. M. (2011). Towards unrestricted public use business microdata: The synthetic Longitudinal Business Database. *International Statistical Review* **79**, 363–384.
- Klein, M. and Sinha, B. (2015). Likelihood based finite sample inference for singly imputed synthetic data under the multivariate normal and multiple linear regression models. *The Journal of Privacy and Confidentiality* **7**, 43–98.
- Little, R. J. A. (1993). Statistical analysis of masked data. *Journal of Official Statistics* **9**, 407–426.
- Nin, J., Herranz, J., and Torra, V. (2008). Rethinking rank swapping to decrease disclosure risk. *Data and Knowledge Engineering* **64**, 346–364.
- Nowok, N., Raab, G. M., and Dibben, C. (2016). synthpop: Bespoke creation of synthetic data in R. *Journal of Statistical Software* **74**.
- Raab, G. M., Nowok, B., and Dibben, C. (2016). Practical data synthesis for large samples. *The Journal of Privacy and Confidentiality* **7**, 67–97.
- Raghunathan, T. E., Reiter, J. P., and Rubin, D. B. (2003). Multiple imputation for statistical disclosure limitation. *Journal of Official Statistics* **19**, 1–16.
- Reiter, J. P. (2005). Using cart to generate partially synthetic public use microdata. *Journal of Official Statistics* **21**, 441–462.
- Reiter, J. P. and Mitra, R. (2009). Estimating risks of identification disclosure in partially synthetic data. *The Journal of Privacy and Confidentiality* **1**, 99–110.
- Reiter, J. P. and Raghunathan, T. E. (2007). The multiple adaptations of multiple imputation. *Journal of the American Statistical Association* **102**, 1462–1471.
- Rubin, D. B. (1993). Discussion statistical disclosure limitation. *Journal of Official Statistics* **9**, 461–468.
- Savitsky, T. D., Williams, M. R., and Hu, J. (2020). Bayesian pseudo posterior mechanism under differential privacy arXiv:1909.11796.
- Snoke, J., Raab, G. M., Nowok, B., Dibben, C., and Slavkovic, A. (2018). General and specific utility measures for synthetic data. *Journal of the Royal Statistical Society, Series A* **181**, 663–688.
- Torra, V., Abowd, J. M., and Domingo-Ferrer, J. (2006). Using Mahalanobis distance-based record linkage for disclosure risk assessment. In *Privacy in Statistical Databases*, vol. 4302 of *Lecture Notes in Computer Science*, 233–242. Springer.
- Wang, H. and Reiter, J. P. (2012). Multiple imputation for sharing precise geographies in public use data. *Annals of Applied Statistics* **6**, 229–252.
- Woo, M.-J., Reiter, J. P., Oganian, A., and Karr, A. F. (2009). Global measures of data utility for microdata masked for disclosure limitation. *Journal of Privacy and Confidentiality* **1**, 1.