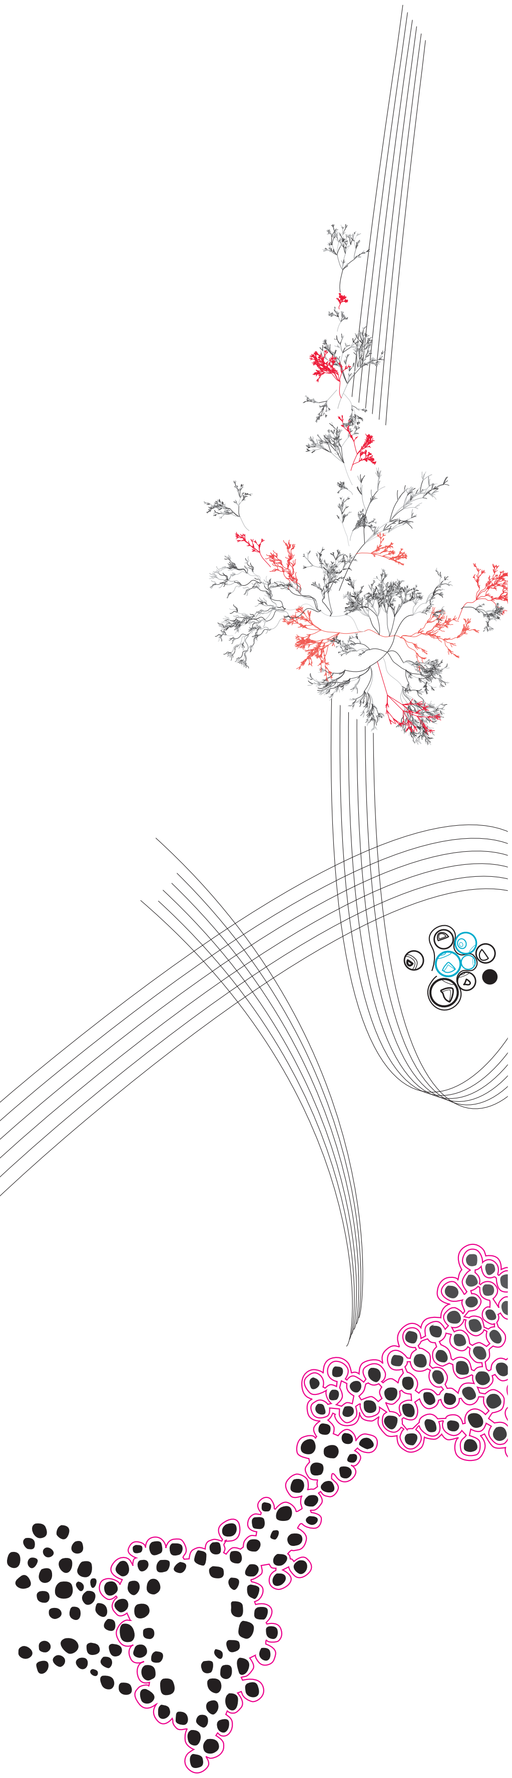

BSc Thesis Applied Mathematics

Pufferfish privacy when publishing on thematic maps

Martijn ter Steege

Supervisor: dr.ir. J. Goseling

January, 2021

Department of Applied Mathematics
Faculty of Electrical Engineering,
Mathematics and Computer Science

UNIVERSITY OF TWENTE.

Preface

This paper was written to fulfill the graduation requirements of the Bachelor Applied Mathematics at the University of Twente. The main research was performed from November 9, 2020 up to January 22, 2021.

I would like to thank J. Goseling for his guidance and his critical feedback. I do not believe that I would have enjoyed doing this research, if you were not by my side and motivating me to strive further. I think I would have created half this paper without your guidance, so I'm really grateful for your time and patience.

Next to that, I would like to thank my parents and grandpa for their support and ideal work atmosphere. Last but not least, I would like to thank my friends for their endless support.

Martijn,
22 January 2021

Contents

Pufferfish privacy when publishing on thematic maps

J.M. ter Steege*

January, 2021

Abstract

Statistical disclosure control is essential to ensure the privacy of individuals in a data set. Statistics of a data set can be visualized on thematic maps by colouring its geographic location. We will apply Pufferfish privacy to protect thematic grid maps by adding appropriate noise. We come up with absolute and relative error protection methods by applying the least necessary noise according to the Laplace distribution. This way, the highest utility remains, while privacy is guaranteed. After observing some unrealistic results, we have also developed a bounded privacy mechanism. Numerical experiments show how the mechanisms act on different input parameters.

Keywords: Pufferfish privacy, statistical disclosure control, differential privacy, thematic maps, grid maps, absolute error protection, relative error protection, bounded mechanism, Laplace distribution

1 Introduction

Statistical disclosure control is a critical technique in publishing data. It is of great value to publish useful data, while it is crucial that no individual information can be withdrawn from published data. Statistical disclosure control makes sure that sensitive data are adjusted accordingly in order to protect personal information, while maintaining the statistical utility. A logical step would be to add noise to or to remove outliers from raw data, but this may result in non-representative outcomes, which is undesirable. Therefore, a precise protection method is essential in order to optimize statistical publications.

Publication of data can be done in various ways. This research focuses on publishing on thematic maps. This can be done by dividing the map in grids and colouring the grids according to its statistical value. Figure 1 shows such a grid map of households in Enschede, used by Statistics Netherlands (Dutch: Centraal Bureau voor de Statistiek, CBS). Some of the cells in this figure are not coloured, due to its sensitivity: it is better to publish nothing than vulnerable information. Therefore, the utility of this grid is not optimal and thus can be improved. The goal of this research is to look into a different safe solution for publishing thematic grid maps.

1.1 Related work

In principle, the publication of a grid map is very similar to publishing a table [5]: all contributions in each grid cell get combined into one table value. After a table of u rows and v columns is filled, its values will get shown as coloured cells, which will produce a grid map of u cell rows and v cell columns. This means that tabular security is necessary for the sake of protecting the grid. In this section, we will explain some of the previous work and explain the possible improvements of preceding results.

*Email: j.m.tersteege@student.utwente.nl

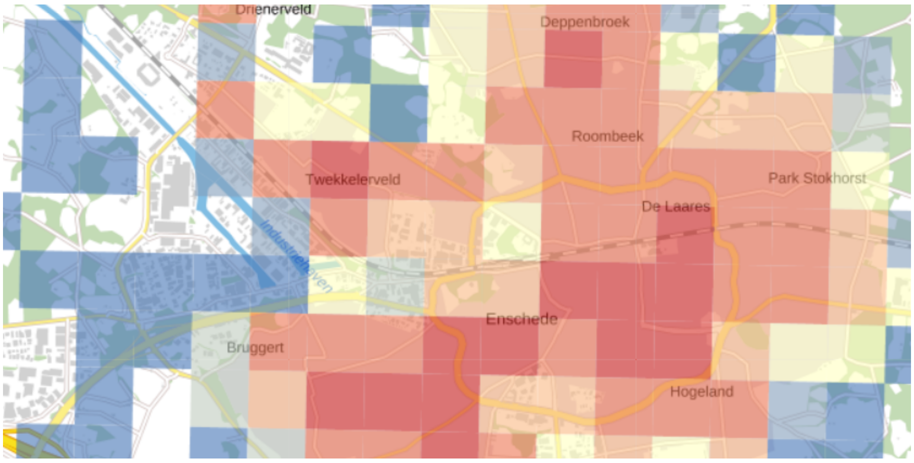

Figure 1: Number of private households in Enschede, averaged over grid cells of 500 meter square (`https://www.cbsinuwbuurt.nl`)

Detailed research has been done into tabular statistical disclosure control [4]. Several sensitivity rules were discovered and we will explain the most commonly used: the minimum frequency rule checks if each table cell has at least n contributors. It makes sense that a value that only has a few contributors is sensitive, so this rule eliminates that case. The (n, k) dominance rule checks whether the n highest contributions do not exceed $k\%$ of the total cell value. This one is comparable with the frequency rule, but now also checks if there are enough contributors that jointly contributes at least $k\%$. These rules are often used in combination with other sensitivity rules.

More in-depth attention should be paid to the $p\%$-rule. This rule indicates to what extent an attacker can estimate another contribution in a table cell. The attacker can get this estimation by subtracting his contribution from the total cell value, and use that result as approximation. The $p\%$-rule states that a cell may not be published if this approximation is within $p\%$ of another contribution. The attacker would get a too close estimation of another contributor, if this rule is not satisfied. Note that a $p\%$-attacker can only disclose information from the largest contributor in the cell, and that the attacker has to contribute to the cell in order to get a close estimation.

Lastly, the (p, q)-rule is an extension on the $p\%$ rule. In this case, the (p, q)-attacker has information about the lower bounds of each contributions, with relative error of at most $q\%$. The one exception is the contribution that the attacker wants to disclose. The attacker can get an approximation by subtracting his contribution and all lower bounds of the other contributions from the total cell value. If this approximation is within $p\%$ of the attacked contribution, then the (p, q)-rule is violated and the table value is not safe to publish.

The exact details of the $p\%$ and (p, q)-rule can be found in [5], but it should become clear that the conditions when these rules can be applied are limited: if an attacker does not contribute to a cell, attacking in that cell can become hard, if not impossible. Besides, only the largest contributor in each cell can be attacked. Lastly, a lower bound of the other contributions in a cell is necessary in the second sensitivity rule. Not all of these conditions are always applicable, so looking into other sensitivity rules will be really valuable.

Once sensitive cells are observed, it is essential to protect the table cells accordingly. Here are the most commonly used tabular protection methods: table restructuring makes sure that different table cells are combined in order to protect its cells. The downside of this method is that the total amount of cells decrease and less detailed information is given. Cell suppression is applied in Figure 1 and makes sure that cells that are unsafe are not published. This is even less desirable than table restructuring. The last noteworthy method is additive rounding: this method makes sure that each original value gets rounded to the closed value of a limited set of values. This way, it is only possible to know in which interval the original value was. Although these protection methods work well, we will look into other protection methods that do not use these, to hopefully increase the statistical utility.

In more recent research, investigation has been initiated into statistical disclosure control when publishing on thematic maps. The main goal of one of the previous research articles was to look into 'continuous visualisation on a geographic map, based on measurements that were taken at finitely many points' [5]. In this research, random noise from the normal distribution was added in order to protect sensitive data, according to the (p, q)-rule. While the research goal of continuous visualisation is an extension compared to publishing grid maps, it will be interesting to look into other sensitivity rules.

Another paper introduced the concept of differential privacy [3]. The goal of differential privacy is to make sure that each individual that contributes to the cell has approximately the same privacy that would follow as if that data point was excluded. That means that the published statistic should not excessively rely on any of the data contributions. The main result from this paper was to calculate the sensitivity of a data set, and to apply noise according to its sensitivity. This proceeds into noise with high variance for data with a few contributions and negligible noise for relatively safe information sources.

The Pufferfish framework can be seen as a generalisation of differential privacy [6]. It uses conditional probabilities to make it almost impossible to distinguish whether a statement about an individual is true or false, when considering published data. Research into the Pufferfish framework has barely taken place and thus Pufferfish framework applications are lacking.

1.2 Goal and Outline

As a consequence of the limitations in the currently-used sensitivity rules and insufficient research into implementing the Pufferfish framework to statistical databases, it is of great value to take a close look into this. Therefore, the goal of this research is to create protected thematic grid maps on general finite measurements, where the protection is based on the Pufferfish framework. The grid should not disclose detailed information about single contributions, since these measurements can be seen as sensitive information. It will be interesting to compare numerical results with other privacy techniques, in order to broaden the toolbox of statistical disclosure control.

First, in Section 2, we will introduce the Pufferfish framework and the mathematical model. In Section 3 we use theory to create mechanisms that generate protected grid maps. In Section 4, we will apply boundaries to previous mechanisms, which will result in a new bounded mechanism. Then, Section 5 shows the numerical results of the privacy mechanisms. We will make some final remarks in Section 6.

2 Model & preliminaries

First of all, the Pufferfish Framework is introduced. This framework will be applied as privacy definition. The privacy mechanism coupled to this framework will be introduced afterwards.

2.1 Privacy definition: the Pufferfish Framework

The Pufferfish Framework [6] relies on three parameters: a set of potential secrets $\mathbb{S}$, a set of discriminative pairs $\mathbb{S}_{pairs} \subseteq \mathbb{S} \times \mathbb{S}$ and a collection of data evolution scenarios $\mathbb{D}$. The potential secrets are in essence the set that a data curator tries to hide. A possible attacker should not be able to get information $s_i \in \mathbb{S}$, when evaluating the published data. An example of this could be 'individual h is in the data set'.

The discriminative pairs are the set that makes sure the potential secrets stay hidden. If $(s_i, s_j) \in \mathbb{S}_{pairs}$, then an attacker should not be able to distinguish whether statement s_i or s_j is true of the actual data set. Since s_i and s_j should be indistinguishable, it follows that s_i and s_j are mutually exclusive: it cannot happen that both s_i and s_j are true, since then it is impossible to distinguish them. However, s_i and s_j do not have to be exhaustive, since it might as well be that both s_i and s_j are not true. An example of a discriminative pair is ('individual h is not in the data set', 'individual h is in the data set').

The data evolution scenarios is the set of probability distributions that a possible attacker can apply to disclose the secrets $\mathbb{S}$. The data curator wants to make sure that an attacker is not able to retrieve the secrets, i.e. the attacker should not be able to distinguish between statements s_i and s_j ($(s_i, s_j) \in \mathbb{S}_{pairs}$) when using distribution $\theta \in \mathbb{D}$. This intuition is where the Pufferfish privacy is based on:

Definition 2.1 (ϵ-Pufferfish privacy)**.** Given a set of potential secrets $\mathbb{S}$, a set of discriminative pairs $\mathbb{S}_{pairs}$, a set of data evaluation scenarios $\mathbb{D}$, and a privacy parameter $\epsilon > 0$, a privacy mechanism $\mathbb{M}$ satisfies ϵ-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$ if

i for all possible outputs $\omega \in \text{range}(\mathbb{M})$,

ii for all pairs $(s_i, s_j) \in \mathbb{S}_{pairs}$ of potential secrets,

iii for all distributions $\theta \in \mathbb{D}$ for which $\mathrm{P}(s_i \mid \theta) \neq 0$ and $\mathrm{P}(s_j \mid \theta) \neq 0$, the following holds:

$$\begin{aligned} &P(\mathbb{M}(\mathbb{T}) = \omega \mid s_i, \theta) \leq e^{\epsilon} P(\mathbb{M}(\mathbb{T}) = \omega \mid s_j, \theta), \\ &P(\mathbb{M}(\mathbb{T}) = \omega \mid s_j, \theta) \leq e^{\epsilon} P(\mathbb{M}(\mathbb{T}) = \omega \mid s_i, \theta). \end{aligned} \tag{1}$$

In words, this means that the probability distribution of the output ω should be relatively stable, regardless of what potential secret is assumed to be true by an attacker. In this way, it should be nearly impossible for attackers to receive any information about any individual, which is the desired situation.

Another way of understanding Equations 1 is by looking at its odds ratio [6]. With some basic conditional probability theory, Equations 1 can be combined and rewritten as:

$$e^{-\epsilon} \leq \frac{P(s_i \mid \mathbb{M}(\mathbb{T}) = \omega, \theta)}{P(s_j \mid \mathbb{M}(\mathbb{T}) = \omega, \theta)} \bigg/ \frac{P(s_i \mid \theta)}{P(s_j \mid \theta)} \leq e^{\epsilon} \tag{2}$$

Close attention should be paid to Equation 2. It should become clear that strong privacy implies that the boundaries in Equation 2 are close to one. This means that the numerator and denominator should be relatively equal. Firstly, the fraction $\alpha = \frac{P(s_i|\theta)}{P(s_j|\theta)}$ is just the initial ratio that an attacker thinks s_i is α times more likely to be true than s_j. On the other hand, the fraction $\beta = \frac{P(s_i|\mathbb{M}(\mathbb{T})=\omega,\theta)}{P(s_j|\mathbb{M}(\mathbb{T})=\omega,\theta)}$ is the ratio that an attacker thinks s_i is β times more likely to be true than s_j, after observing published data ω. For Equation 2 to apply, both α and β should be relatively equal, thus publishing ω should provide barely any additional information for possible attackers. This way, no new information can be disclosed after observing protected output ω.

In the Pufferfish framework, θ can be seen as all information that an attacker will apply, in order to disclose personal data: think about prior knowledge and estimations of probability distributions. With all of this combined, we will be able to produce Pufferfish private mechanisms, but we will first describe the mathematical model.

2.2 Mathematical model

Throughout next sections, a lot of common notation will be used. In this section, some general symbols are introduced, while the Pufferfish mechanisms are introduced in the next section.

Let $\mathbb{T} = \{t_1, t_2, \ldots, t_n\} \in \mathbb{R}$ be the measurements of each individual in the determined data set. These measurements can vary from household income to gas consumption, and from inhabitants to age. It should not matter what topic is desired to be published, the corresponding output should be protected correspondingly. These measurements should be coupled to their geographic location.

The aimed result should consist of a grid $\mathbb{G}$. Let $\mathbb{G}$ consist of cells $c_{u,v}$, where u and v determine the geographic location of the cell. As mentioned, publishing a grid map is almost similar to publishing a table, so initially $\mathbb{G}$ can be seen as a table. Each cell should get a requested value, so let $g_{u,v}$ represent this value in each cell. These requested values can vary from the total sum to the average, depending on the publicist's preference, but our aim will be to safely publish the average contribution per cell. This general statistic will be a combination of each individual contribution t_i in that cell.

A data curator, an official authority that issues statistics, determines the privacy definition. This privacy definition controls the amplitude of the additional noise, depending on the value of protection. The privacy mechanism $\mathbb{M}$ applies appropriate noise in order to satisfy the privacy definition. The obtained output of this mechanism should be protected and is denoted as $\mathbb{M}(\mathbb{T}) \equiv \omega$.

3 Theoretical analysis: Mechanism $\mathbb{M}$

In this section, we will come up with two privacy mechanisms that both satisfy Pufferfish privacy. Its algorithms to produce thematic grid maps will be applied later.

3.1 Privacy mechanism with absolute error

In order to get a privacy mechanism of protecting grid values, a few assumptions are necessary. The first assumption is that the individual contributions $t_i \in c_{u,v}$ are non-negative and continuous. Another assumption is that a possible attacker does know whether individual i contributes to $g_{u,v}$, and is interested in its contribution. Since the contributions are continuous, the probability that an individual i has exactly value t_i is zero, so we have to work with probabilities that t_i is in an interval.

With this, we can already define potential secrets: Let $\sigma_{i,[x-c,x+c)}$ be the statement 'The contribution of individual i is in the range $[x-c, x+c)$', i.e. $t_i \in [x-c, x+c)$. This statement is somewhat broad and should be chosen carefully, since it uses an absolute interval. Imagine publishing grades of students, then an interval of one to ten covers all students, but publishing income of individuals per year is not sufficient on that same interval. We will first apply an absolute interval $[x-c, x+c)$ and then improve that to a relative interval.

As initial potential secret that we try to hide, we get:

$$\mathbb{S} = \left\{\sigma_{i,[x-c,x+c)} : i \in c_{u,v}, x > 0\right\} \tag{3}$$

As discriminative pairs, we will use neighbouring intervals, since these are mutually exclusive, and attackers should have a hard time to differentiate whether someones contribution is in one or the other neighbouring intervals.

Thus:

$$\mathbb{S}_{pairs} = \left\{\sigma_{i,[x-c,x+c)}, \sigma_{i,[x+c,x+3c)} : i \in c_{u,v}, x > 0\right\} \tag{4}$$

Then for the evolution scenarios $\mathbb{D}$, we'll choose the set of all probability distributions that distributes the measurements independently.

This results in:

$$\mathbb{D} = \begin{cases} \theta \equiv [f_1, \ldots, f_n], \\ P(\mathbb{T} \mid \theta) = \prod_{t_i \in \mathbb{T}} f_i(t_i), \end{cases} \tag{5}$$

With this choice of potential secrets $\mathbb{S}$, discriminative pairs $\mathbb{S}_{pairs}$ and data evolution scenarios $\mathbb{D}$, we let $|c_{u,v}| \equiv n$ denote the amount of contributors in cell $c_{u,v}$. Theorem 3.1 follows:

Theorem 3.1 (Average of contributions with absolute error)**.** *With* $\mathbb{S}$ *and* $\mathbb{S}_{pairs}$ *defined in Equations* (3) *and* (4)*, let* $\mathbb{D}$ *be the set of probability distributions specified in Equation* (5)*. The mechanism* $\mathbb{M}_{abs}$ *which returns* $X+\frac{\sum_{i\in c_{u,v}} t_i}{n}$*, where* X *has Laplace density* $\frac{n\epsilon}{8c}e^{-n\epsilon|x|/4c}$*, satisfies* ϵ*-Pufferfish*$(\mathbb{S},\mathbb{S}_{pairs},\mathbb{D})$*.*

Proof. We use t_i as individual contribution of individual i in cell $c_{u,v}$. Furthermore we use $f_i(t_i \in A)$ to represent the probability that t_i is in the set A and $f_i(t_i \mid A)$ the conditional probability density of t_i given $t_i \in A$. Consider individual i, real value x, and $\theta = [f_1, \dots, f_n]$ such that $f_1(t_1 \in [x-c, x+c)) \neq 0$ and $f_1(t_1 \in [x+c, x+3c)) \neq 0$. Then:

$$P(\mathbb{M}_{\text{abs}}(\mathbb{T}) = \omega \mid t_1 \in [x-c, x+c), \theta)$$

$$= \int \cdots \int \left(\begin{array}{c} \frac{n\epsilon}{8c}\exp\left(-\frac{n\epsilon}{4c}\left|\omega - \sum_{i=1}^{n}\frac{t_i}{n}\right|\right) \\ \times f_1(t_1 \mid t_1 \in [x-c, x+c)) f_2(t_2) \dots f_n(t_n) \end{array} \right) \mathrm{d}t_1\ \mathrm{d}t_2 \dots \mathrm{d}t_n$$

$$\geq \int \cdots \int \left(\begin{array}{c} \frac{n\epsilon}{8c}\exp\left(-\frac{n\epsilon}{4c}\left|\omega - \frac{x}{n} - \sum_{i=2}^{n}\frac{t_i}{n}\right| - \frac{n\epsilon}{4c}\left|\frac{x-t_1}{n}\right|\right) \\ \times f_1(t_1 \mid t_1 \in [x-c, x+c)) f_2(t_2) \dots f_n(t_n) \end{array} \right) \mathrm{d}t_1\ \mathrm{d}t_2 \dots \mathrm{d}t_n$$

(because: -$\left|A - \frac{t_1}{n}\right|$ = -$\left|A - \frac{x}{n} + \frac{x}{n} - \frac{t_1}{n}\right| \geq$ -$\left|A - \frac{x}{n}\right|$ - $\left|\frac{x}{n} - \frac{t_1}{n}\right|$ = -$\left|A - \frac{x}{n}\right|$ - $\left|\frac{x-t_1}{n}\right|$)

$$\geq \int \cdots \int \left(\begin{array}{c} \frac{n\epsilon}{8c}\exp\left(-\frac{n\epsilon}{4c}\left|\omega - \frac{x}{n} - \sum_{i=2}^{n}\frac{t_i}{n}\right| - \frac{n\epsilon}{4c}\frac{c}{n}\right) \\ \times f_1(t_1 \mid t_1 \in [x-c, x+c)) f_2(t_2) \dots f_n(t_n) \end{array} \right) \mathrm{d}t_1\ \mathrm{d}t_2 \dots \mathrm{d}t_n$$

(because: $t_1 \in [x-c, x+c)$, so $|x - t_1| \leq c$, thus -$\left|\frac{x-t_1}{n}\right| \geq -\frac{c}{n}$)

$$= \int \cdots \int \left(\begin{array}{c} \frac{n\epsilon}{8c}\exp\left(-\frac{n\epsilon}{4c}\left|\omega - \frac{x}{n} - \sum_{i=2}^{n}\frac{t_i}{n}\right| - \frac{\epsilon}{4}\right) \\ \times f_2(t_2) \dots f_n(t_n) \end{array} \right) \mathrm{d}t_2 \dots \mathrm{d}t_n$$

(we cancel out c and n and the integrand no longer depends on t_1.
It follows that we can cancel out $\int_{-\infty}^{\infty} f_1(t_1 \mid t_1 \in [x-c, x+c)) = 1$)

$$= \frac{n\epsilon}{8c}e^{-\epsilon/4} \int \cdots \int \left(\begin{array}{c} \frac{n\epsilon}{8c}\exp\left(-\frac{n\epsilon}{4c}\left|\omega - \frac{x}{n} - \sum_{i=2}^{n}\frac{t_i}{n}\right|\right) \\ \times f_2(t_2) \dots f_n(t_n) \end{array} \right) \mathrm{d}t_2 \dots \mathrm{d}t_n \tag{6}$$

Meanwhile,

$$P(\mathbb{M}_{\text{abs}}(\mathbb{T}) = \omega \mid t_1 \in [x+c, x+3c), \theta)$$

$$= \int \cdots \int \left(\begin{array}{r} \dfrac{n\epsilon}{8c} \exp\left(-\dfrac{n\epsilon}{4c} \left| \omega - \displaystyle\sum_{i=1}^{n} \dfrac{t_i}{n} \right| \right) \\ \times f_1(t_1 \mid t_1 \in [x+c, x+3c)) f_2(t_2) \ldots f_n(t_n) \end{array} \right) \mathrm{d}t_1 \; \mathrm{d}t_2 \ldots \mathrm{d}t_n$$

$$\leq \int \cdots \int \left(\begin{array}{c} \dfrac{n\epsilon}{8c} \exp\left(-\dfrac{n\epsilon}{4c} \left| \omega - \dfrac{x}{n} - \displaystyle\sum_{i=2}^{n} \dfrac{t_i}{n} \right| + \dfrac{n\epsilon}{4c} \left| \dfrac{x - t_1}{n} \right| \right) \\ \times f_1(t_1 \mid t_1 \in [x+c, x+3c)) f_2(t_2) \ldots f_n(t_n) \end{array} \right) \mathrm{d}t_1 \; \mathrm{d}t_2 \ldots \mathrm{d}t_n$$

(because: $\left|A - \frac{t_1}{n}\right| = \left|A - \frac{x}{n} + \frac{x}{n} - \frac{t_1}{n}\right| \geq \left|A - \frac{x}{n}\right| - \left|\frac{x}{n} - \frac{t_1}{n}\right|$
thus -$\left|A - \frac{t_1}{n}\right| \leq -\left|A - \frac{x}{n}\right| + \left|\frac{x - t_1}{n}\right|$)

$$\leq \int \cdots \int \left(\begin{array}{r} \dfrac{n\epsilon}{8c} \exp\left(-\dfrac{n\epsilon}{4c} \left| \omega - \dfrac{x}{n} - \displaystyle\sum_{i=2}^{n} \dfrac{t_i}{n} \right| + \dfrac{n\epsilon}{4c} \dfrac{3c}{n} \right) \\ \times f_1(t_1 \mid t_1 \in [x+c, x+3c)) f_2(t_2) \ldots f_n(t_n) \end{array} \right) \mathrm{d}t_1 \; \mathrm{d}t_2 \ldots \mathrm{d}t_n$$

(because: $t_1 \in [x+c, x+3c)$, so $|x - t_1| \leq 3c$, thus $\left|\frac{x - t_1}{n}\right| \leq \frac{3c}{n}$)

$$= \int \cdots \int \left(\begin{array}{r} \dfrac{n\epsilon}{8c} \exp\left(-\dfrac{n\epsilon}{4c} \left| \omega - \dfrac{x}{n} - \displaystyle\sum_{i=2}^{n} \dfrac{t_i}{n} \right| + \dfrac{3\epsilon}{4} \right) \\ \times f_2(t_2) \ldots f_n(t_n) \end{array} \right) \mathrm{d}t_2 \ldots \mathrm{d}t_n$$

(we cancel out c and n and the integrand no longer depends on t_1.
It follows that we can cancel out $\int_{-\infty}^{\infty} f_1(t_1 \mid t_1 \in [x - c, x + c)) = 1$)

$$= \frac{n\epsilon}{8c} e^{3\epsilon/4} \int \cdots \int \left(\begin{array}{r} \dfrac{n\epsilon}{8c} \exp\left(-\dfrac{n\epsilon}{4c} \left| \omega - \dfrac{x}{n} - \displaystyle\sum_{i=2}^{n} \dfrac{t_i}{n} \right| \right) \\ \times f_2(t_2) \ldots f_n(t_n) \end{array} \right) \mathrm{d}t_2 \ldots \mathrm{d}t_n \tag{7}$$

Comparing Equations (6) and (7), we see that the only difference between them is the constant multiplier outside the integral. Thus dividing, we get

$$\frac{P(\mathbb{M}_{\text{abs}}(\mathbb{T}) = \omega \mid t_1 \in [x - c, x + c), \theta)}{P(\mathbb{M}_{\text{abs}}(\mathbb{T}) = \omega \mid t_1 \in [x + c, x + 3c), \theta)} \geq \frac{\frac{n\epsilon}{8c} e^{\frac{-\epsilon}{4}}}{\frac{n\epsilon}{8c} e^{\frac{3\epsilon}{4}}} = e^{-\epsilon}$$

What follows is

$$P(\mathbb{M}_{\text{abs}}(\mathbb{T}) = \omega \mid t_1 \in [x+c, x+3c), \theta) \leq e^{\epsilon} P(\mathbb{M}_{\text{abs}}(\mathbb{T}) = \omega \mid t_1 \in [x-c, x+c), \theta) \tag{8}$$

A similar calculation results in

$$P(\mathbb{M}_{\text{abs}}(\mathbb{T}) = \omega \mid t_1 \in [x-c, x+c), \theta) \leq e^{\epsilon} P(\mathbb{M}_{\text{abs}}(\mathbb{T}) = \omega \mid t_1 \in [x+c, x+3c), \theta) \tag{9}$$

We can repeat this calculation for any individual contribution t_i, for other choices of $f_1, f_2, \ldots f_n$ and for other y, and so $\mathbb{M}_{\text{abs}}$ satisfies ϵ-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$. □

Note that $\mathbb{M}(\mathbb{T})$ is continuous, so its probability is just the height of its density function. Another observation is that the density function of the noise in Theorem 3.1 depends on ϵ, n and c. The variance of this Laplace distribution X is $\text{var}(X) = \frac{32c^2}{\epsilon^2 n^2}$. A relatively small privacy parameter ϵ should result into undistinguishability, so strong protection is necessary in order to control disclosure. For a relatively small amount of contributions n, it follows that this data is really sensitive. That also means that strong protection is essential. A small check by comparing small and large values of ϵ and n shows that the small values yields higher variance. This means that the noise for small values of ϵ and n will likely be more spread out, which should result in more total noise. The addition of more noise results in more randomness and hence stronger privacy.

For the interval parameter c, it is what harder to see. With some rational thinking, it makes sense that small values of c need strong protection, because the lower the c, the more accurate an attacker can discover personal data. However, the intervals in Equations (8) and (9) are the probability conditions, so an attacker is most likely to see similar outputs ω for the two small neighbouring intervals. Conversely, if the two neighbouring intervals are larger, than the probability that the outputs ω differ is higher, because more information is contained in a larger interval. This difference in outputs can disclose personal information and therefore, a large interval parameter c needs stronger protection. Therefore, large values of c result in higher variance compared to a small value of c.

Protecting data with absolute error yields additive noise. Every cell in the grid map has its own parameters, so any cell gets additive noise dependent on its sensitivity. This way, every cell gets minimal, but appropriate noise, so the grid's utility is as high as possible, while privacy is ensured. Algorithm 1 shows the pseudocode how each cell gets protected in order to satisfy Pufferfish privacy. Visual results of this algorithm will be shown in Section 5.1.

Algorithm 1: Absolute error Pufferfish protection

```
Require: $\epsilon$ and $c$
Require: $\mathbb{T} = \{t_1, t_2 \dots t_n\}$: the real data
Require: $r$: cell size
for all cells $c_{u,v}$ do
    n[u,v] ← count of $t_i$ in $c_{u,v}$ ;
    mean[u,v] ← mean of $t_i$ in $c_{u,v}$ ;
    noise[u,v] ← sample of distribution $\frac{n\epsilon}{8c} e^{-n\epsilon|x|/4c}$ ;
end
return mean + noise
```

3.2 Privacy mechanism with relative error

The result of Theorem 3.1 shows that the average of a data set is safe to publish, when adding appropriate noise to protect neighbouring absolute intervals. However, it would be more convenient to have neighbouring relative intervals, since that is applicable for any data set, while absolute intervals might have less utility for certain data.

Therefore, we will use the relative interval $[ky, y/k)$ with $k \in (0, 1)$. For example, investigating an interval with a relative error of 10% gives $k = 0.9$. With this interval, the potential secret and the discriminative pairs adjust according to this interval as well.

Adjusting the Pufferfish mechanism to relative error yields:

$$\mathbb{S} = \left\{ \sigma_{i,[ky,y/k)} : i \in c_{u,v}, y > 0 \right\} \tag{10}$$

and

$$\mathbb{S}_{pairs} = \left\{ \sigma_{i,[ky,y/k)}, \sigma_{i,[y/k,y/3k)} : i \in c_{u,v}, y > 0 \right\} \tag{11}$$

If $t_i \in [ky, y/k)$, then $\log(t_i) \in [\log(y) + \log(k), \log(y) - \log(k))$, so protecting y with a relative error is the same as protecting $\log(y)$ with an absolute error. With this intuition and the result of Theorem 3.1, Corollary 3.2 follows.

Corollary 3.2 (Average of contributions with relative error)**.** *With* $\mathbb{S}$ *and* $\mathbb{S}_{pairs}$ *defined in Equations* (10) *and* (11)*, let* $\mathbb{D}$ *be the set of probability distributions specified in Equation* (5)*. The mechanism which returns* $Y + \log(\frac{\sum_{i \in c_{u,v}} t_i}{n})$ *where* Y *has Laplace density* $\frac{-n\epsilon}{8\log(k)} e^{n\epsilon|y|/4\log(k)}$*, satisfies* ϵ*-Pufferfish*$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$*.*

One of the main benefits of Pufferfish privacy is that its definition is robust to post-processing [2]. That means that a Pufferfish private mechanism can be post-processed, and then it still satisfies Pufferfish privacy. Therefore, applying the exponential function of the mechanism of Corollary 3.2 also satisfies ϵ-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$. This results in the following:

$$\mathbb{M}_{\text{rel}} \equiv e^{Y + \log(\frac{\sum_{i \in c_{u,v}} t_i}{n})} = e^Y e^{\log(\frac{\sum_{i \in c_{u,v}} t_i}{n})} = e^Y \frac{\sum_{i \in c_{u,v}} t_i}{n} \tag{12}$$

Thus the mechanism $\mathbb{M}_{\text{rel}}$ that applies multiplicative noise to the average of a data set satisfies ϵ-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$ with relative error. Each cell will again get its own appropriate noise to optimize utility. The pseudocode of this mechanism is applied in Algorithm 2. Visual results of this algorithm will be shown in Section 5.2.

Algorithm 2: Relative error Pufferfish protection

```
Require: $\epsilon$ and $k$
Require: $\mathbb{T} = \{t_1, t_2 \dots t_n\}$: the real data
Require: $r$: cell size
for all cells $c_{u,v}$ do
    n[u,v] ← count of $t_i$ in $c_{u,v}$ ;
    mean[u,v] ← mean of $t_i$ in $c_{u,v}$ ;
    noise[u,v] ← sample of distribution $\frac{-n\epsilon}{8\log(k)} e^{n\epsilon|y|/4\log(k)}$ ;
end
return $e^{\text{noise}} \times$ mean
```

4 Theoretical analysis: Mechanism $\mathbb{B}$

As will become clear in Section 5, we will notice two remarkable effects in the outputs ω of mechanism $\mathbb{M}$: firstly, in the protection method with absolute error, some cells will have negative output, while it is assumed that all contributions are non-negative. Of course, this is unrealistic, so this should be solved. Then in the relative error case, some noise multipliers are significantly larger than others. This causes the grid to only show a few large measurements, while the protected information of the other cells are barely visible: the large measurements stretch out the legend bar which ensures that the majority of the results are in a smaller colour range, thus reducing the distinction between these results.

In order to resolve these peculiarities, we will look into the possibility of bounding the output results. That would mean that all protected negative values would be set to zero, and that an upper bound will make sure all high values get reduced.

4.1 More Pufferfish privacy definitions

The Pufferfish privacy definition in Definition 2.1 relied on probability distribution. This definition can be changed to its cumulative form, which makes it easier to work with when introducing boundaries. Therefore, the new privacy definition that will be applied is as follows:

Definition 4.1 (Cumulative ϵ-Pufferfish privacy)**.** Given a set of potential secrets $\mathbb{S}$, a set of discriminative pairs $\mathbb{S}_{pairs}$, a set of data evaluation scenarios $\mathbb{D}$, and a privacy parameter $\epsilon > 0$, a privacy mechanism $\mathbb{M}$ satisfies ϵ-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$ if

i for all possible outputs $\omega \in \text{range}(\mathbb{M})$,

ii for all pairs $(s_i, s_j) \in \mathbb{S}_{pairs}$ of potential secrets,

iii for all distributions $\theta \in \mathbb{D}$ for which $\text{P}(s_i \mid \theta) \neq 0$ and $\text{P}(s_j \mid \theta) \neq 0$, the following holds:

$$\begin{aligned} P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_i, \theta) &\leq e^{\epsilon} P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_j, \theta), \\ P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_j, \theta) &\leq e^{\epsilon} P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_i, \theta). \end{aligned} \tag{13}$$

A small adaptation of the aforementioned privacy definition is (ϵ, δ)-Pufferfish privacy [2]. This definition uses two privacy parameters and has a bit more freedom than the case with only one parameter. This results in less necessary noise which causes more utility, but also a weaker privacy protection:

Definition 4.2 (Cumulative (ϵ, δ)-Pufferfish privacy)**.** Given a set of potential secrets $\mathbb{S}$, a set of discriminative pairs $\mathbb{S}_{pairs}$, a set of data evaluation scenarios $\mathbb{D}$, and privacy parameters $\epsilon > 0$ and $\delta \in (0, 1)$, a privacy mechanism $\mathbb{M}$ satisfies (ϵ, δ)-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$ if

i for all possible outputs $\omega \in \text{range}(\mathbb{M})$,

ii for all pairs $(s_i, s_j) \in \mathbb{S}_{pairs}$ of potential secrets,

iii for all distributions $\theta \in \mathbb{D}$ for which $\text{P}(s_i \mid \theta) \neq 0$ and $\text{P}(s_j \mid \theta) \neq 0$, the following holds:

$$\begin{aligned} P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_i, \theta) &\leq e^{\epsilon} P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_j, \theta) + \delta, \\ P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_j, \theta) &\leq e^{\epsilon} P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_i, \theta) + \delta. \end{aligned} \tag{14}$$

4.2 Bounding the output ω

Take the protection mechanisms $\mathbb{M}$ from Section 3. A rough visualisation of its cumulative distribution of output ω will be as follows:

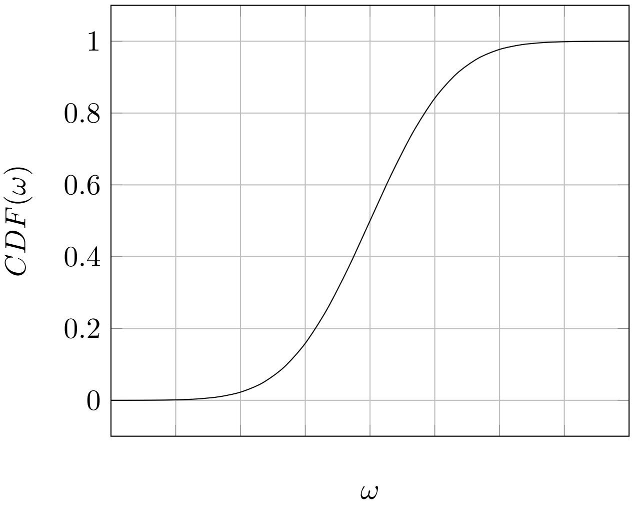

FIGURE 2: Cumulative distribution of mechanism $\mathbb{M}$

Now assume that output values ω cannot lie outside lower and upper boundaries L and U. Applying these boundaries results in the new mechanism $\mathbb{B}$ of Equation (15). The rough visualisation of its cumulative distribution is in Figure 3.

$$P(\mathbb{B}_{\text{abs/rel}}(\mathbb{T}) \leq \omega) = \begin{cases} 0 & \omega \leq L \\ P(\mathbb{M}_{\text{abs/rel}}(\mathbb{T}) \leq \omega) & L < \omega < U \\ 1 & \omega \geq U \end{cases} \tag{15}$$

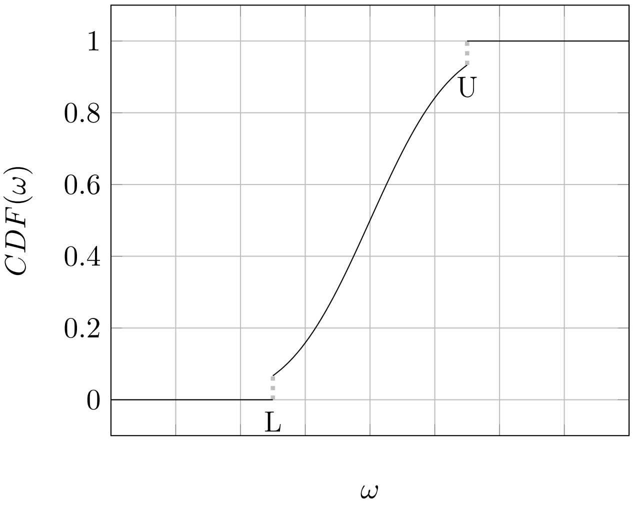

FIGURE 3: Cumulative distribution of mechanism $\mathbb{B}$

The cumulative distribution function of mechanism $\mathbb{B}$ is discontinuous and has jumps at $\omega = L$ and $\omega = U$. These jumps are visible in Figure 3 and we will define its lengths in Equations (16).

$$\begin{aligned} \Delta L &= P(\mathbb{M}(\mathbb{T}) \leq L) \\ \Delta U &= P(\mathbb{M}(\mathbb{T}) \geq U) \end{aligned} \tag{16}$$

Theorem 3.1 states that mechanism $\mathbb{M}$ satisfies Definition 2.1. This is equal to satisfying Definition 4.1, because in Theorem 3.1, the cumulative distribution functions of Definition 4.1 can be applied instead of the probability density functions. Mechanism $\mathbb{B}$ is a combination of this Pufferfish private mechanism and two boundaries. Therefore, the following theorem follows:

Theorem 4.3 (Bounded mechanism)**.** *Given a set of potential secrets $\mathbb{S}$, a set of discriminative pairs $\mathbb{S}_{pairs}$ and a set of data evaluation scenarios $\mathbb{D}$, let $\mathbb{M}$ be an ϵ-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$ mechanism satisfying Equations* (13) *for $\epsilon > 0$ and let δ be defined as* $\max\{\Delta U, e^{\epsilon}\Delta L\}$*. The mechanisms $\mathbb{B}_{abs}$ and $\mathbb{B}_{rel}$ that are defined in Equation* (15) *satisfy (ϵ, δ)-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$.*

Proof. Since the cumulative distribution function of mechanism $\mathbb{B}$ is discontinuous, we have to prove Pufferfish privacy for each continuous interval of ω, and then combine all cases into a general solution for all cases.

Case 1: $\boldsymbol{\omega \in (-\infty, L]}$
If $\omega \in (-\infty, L]$, then:

$$0 = P(\mathbb{B}(\mathbb{T}) \leq \omega) \leq P(\mathbb{M}(\mathbb{T}) \leq \omega)$$

and

$$P(\mathbb{M}(\mathbb{T}) \leq \omega) \leq P(\mathbb{B}(\mathbb{T}) \leq \omega) + \Delta L = \Delta L$$

Combining these using ϵ-Pufferfish privacy of $\mathbb{M}$ results in:

$$\begin{aligned} &P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_i, \theta) \leq P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_i, \theta) \\ &\leq e^{\epsilon} P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_j, \theta) \leq e^{\epsilon} P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_j, \theta) + e^{\epsilon}\Delta L \end{aligned}$$

Therefore:

$$P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_i, \theta) \leq e^{\epsilon} P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_j, \theta) + e^{\epsilon}\Delta L \tag{17}$$

A similar calculation results in

$$P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_j, \theta) \leq e^{\epsilon} P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_i, \theta) + e^{\epsilon}\Delta L \tag{18}$$

Case 2: $\boldsymbol{\omega \in (L, U)}$
If $\omega \in (L, U)$, then:

$$P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_i, \theta) = P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_i, \theta)$$

and

$$P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_j, \theta) = P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_j, \theta)$$

Therefore, using ϵ-Pufferfish privacy of $\mathbb{M}$ results in:

$$P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_i, \theta) \leq e^{\epsilon} P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_j, \theta) \tag{19}$$

and

$$P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_j, \theta) \leq e^{\epsilon} P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_i, \theta) \tag{20}$$

Case 3: $\boldsymbol{\omega \in [U, \infty)}$
If $\omega \in [U, \infty)$, then:

$$P(\mathbb{M}(\mathbb{T}) \leq \omega) \leq P(\mathbb{B}(\mathbb{T}) \leq \omega) = 1$$

and

$$1 - \Delta U = P(\mathbb{B}(\mathbb{T}) \leq \omega) - \Delta U \leq P(\mathbb{M}(\mathbb{T}) \leq \omega)$$

Combining these using ϵ-Pufferfish privacy of $\mathbb{M}$ results in:

$$P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_i, \theta) - \Delta U \leq P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_i, \theta)$$
$$\leq e^{\epsilon} P(\mathbb{M}(\mathbb{T}) \leq \omega \mid s_j, \theta) \leq e^{\epsilon} P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_j, \theta)$$

Therefore:

$$P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_i, \theta) \leq e^{\epsilon} P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_j, \theta) + \Delta U \tag{21}$$

A similar calculation results in

$$P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_j, \theta) \leq e^{\epsilon} P(\mathbb{B}(\mathbb{T}) \leq \omega \mid s_i, \theta) + \Delta U \tag{22}$$

Combining the solutions of all intervals of ω results in a general solution for all possible ω. This is possible for $\delta = \max\{\Delta U, e^{\epsilon}\Delta L\}$, so $\mathbb{B}$ satisfies (ϵ, δ)-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$. □

Using the result of Theorem 4.3, a lower and upper bound can be applied in order to output bounded results, while maintaining Pufferfish privacy. The downside of this mechanism is that its privacy definition is weaker compared to the original mechanism $\mathbb{M}$. Furthermore, the privacy parameter δ depends on the output ω. This causes a privacy leakage, which unfortunately can be exploited by a potential attacker.

4.3 Absolute error protection: worst case analysis

A solution to the privacy leakage from previous section is to determine an upper bound on δ that will definitely satisfy. This upper bound δ_{max} can be achieved by a worst case analysis, where the probability that the outcome of the original mechanism is outside the determined boundaries is the highest.

In Section 3.1, we made the assumption that all contributions t_i are non-negative. For now we will set the lower boundary at some $L_t \leq t_i$ for all i. Additionally, we will also assume that the contributions are upper bounded by some $U_t \geq t_i$ for all i, i.e. all contributions $t_i \in [L_t, U_t]$. Later in this section, it will become clear why these boundaries are chosen. Let the boundaries of mechanism $\mathbb{B}_{\text{abs}}$ be defined as $L \equiv L_t - \gamma < L_t$ and $U \equiv U_t + \gamma > U_t$ for $\gamma > 0$, as can be seen in Figure 4. Theorem 4.4 follows with this boundary parameter γ.

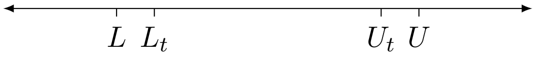

FIGURE 4: Visualisation of the boundaries

Theorem 4.4 ($\delta_{\max}$ with absolute error protection)**.** *With $\mathbb{S}$ and $\mathbb{S}_{pairs}$ defined in Equations* (3) *and* (4)*, let $\mathbb{D}$ be the set of probability distributions specified in Equation* (5)*, let $\mathbb{M}_{abs}$ be an absolute error ϵ-Pufferfish*$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$ *mechanism satisfying Equations* (13) *for $\epsilon > 0$ and let mechanism $\mathbb{B}_{abs}$ be defined in Equation* (15)*, which satisfies* (ϵ, δ)*-Pufferfish*$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$ *for $\delta = \max\{\Delta U, e^{\epsilon}\Delta L\}$. Then $\delta_{max} \equiv e^{\epsilon}P(X \leq -\gamma)$ is an upper bound of δ.*

Proof. We will first look into the worst case scenario for the lower bound of mechanism $\mathbb{B}_{\text{abs}}$, because this will yield the largest value of ΔL. This is the case when all individual contributions are equal to L_t, i.e. $t_i = L_t \ \forall i$. It immediately follows that the largest value of ΔU is achieved when $t_i = U_t \ \forall i$. We will apply this to mechanism $\mathbb{M}_{\text{abs}}(\mathbb{T}) = \frac{1}{n}\sum_{i=1}^{n} t_i + X$ where X has Laplace density $\frac{n\epsilon}{8c}e^{-n\epsilon|x|/4c}$. In the worst case scenario, the following happens with ΔL and ΔU from Equations (16):

$$\Delta L = P(\mathbb{M}_{\text{abs}}(\mathbb{T}) \leq L) = P(\frac{1}{n}\sum_{i=1}^{n} t_i + \leq L) \leq P(L_t + X \leq L) = P(X \leq L - L_t) = P(X \leq -\gamma)$$

and

$$\Delta U = P(\mathbb{M}_{\text{abs}}(\mathbb{T}) \geq U) = P(\frac{1}{n}\sum_{i=1}^{n} t_i + X \geq U) \leq P(U_t + X \geq U) = P(X \geq U - U_t) = P(X \geq \gamma)$$

Note that X is just the Laplace density of $\mathbb{M}_{\text{abs}}$, which in our case is symmetric around 0. That means that $\Delta L \leq P(X \leq -\gamma) = P(X \geq \gamma) \geq \Delta U$. Since $\delta = \max\{e^{\epsilon}\Delta L, \Delta U\}$ and $e^{\epsilon} > 1$ for $\epsilon > 0$, it follows that the upper bound on $e^{\epsilon}\Delta L$ will satisfy as an upper bound on δ. Therefore, $\delta_{\max} = e^{\epsilon}P(X \leq -\gamma)$ can be seen as an upper bound on δ, since this is the upper bound in the worst case scenario for both $e^{\epsilon}\Delta L$ and ΔU. □

Note that each cell in the grid has its own $\delta_{\max}$, since the amount of contributors n is not equal in each cell. Therefore, we will define $\Delta \equiv \max_{c_{u,v}} \delta_{\max}$ as the grid map upper bound of δ.

We will apply Theorem 4.4 to several grid maps that result from Algorithm 1. We will use $\gamma = \min_i t_i$, which results in $L = 0$ and $U = \max_i t_i + \min_i t_i$. Visualisation of the new mechanism $\mathbb{B}_{\text{abs}}$ can be seen in Section 5.1.

4.4 Relative error protection: worst case analysis

Looking at Theorem 4.4, a similar result will follow for the relative error protection. The same assumptions can be assumed with the exception of the defined L and U. Let the boundaries of mechanism $\mathbb{B}_{\text{rel}}$ be defined as $L \equiv \frac{L_t}{\lambda} < L_t$ and $U \equiv \lambda U_t > U_t$ for $\lambda > 1$, as still can be seen in Figure 4. The following corollary follows:

Corollary 4.5 ($\delta_{\max}$ with relative error protection)**.** With $\mathbb{S}$ and $\mathbb{S}_{pairs}$ defined in Equations (10) and (11), let $\mathbb{D}$ be the set of probability distributions specified in Equation (5), let $\mathbb{M}_{\text{rel}}$ be a relative error ϵ-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$ mechanism satisfying Equations (13) and let mechanism $\mathbb{B}_{\text{rel}}$ be defined in Equation (15), which satisfies (ϵ, δ)-Pufferfish$(\mathbb{S}, \mathbb{S}_{pairs}, \mathbb{D})$ for $\delta = \max\{\Delta U, e^{\epsilon}\Delta L\}$. Then $\delta_{\max} \equiv e^{\epsilon}P(Y \leq \log(\frac{1}{\lambda}))$ is an upper bound of δ.

The proof of Corollary 4.5 is exactly the same as the proof of Theorem 4.4, with applying mechanism $\mathbb{M}_{\text{rel}}$ instead of mechanism $\mathbb{M}_{\text{abs}}$. We will apply Corollary 4.5 to several grid maps that result from Algorithm 2. We will use $\lambda = 1.25$, which results in $L = \min_i t_i/1.25$ and $U = 1.25 \max_i t_i$. Visualisation of the new mechanism $\mathbb{B}_{\text{rel}}$ can be seen in Section 5.2.

5 Numerical results

In this section, we will simulate several grid maps that would follow from applying the protection methods from Sections 3 and 4. These privacy mechanisms satisfy Pufferfish privacy, so these results will be shown with different parameters. We will also compare the unprotected with the protected mappings. This comparison will be done using open data from the R package ‘sdcSpatial’ [1]. Figure 5 shows each contributor’s geographic location with its contribution, that will be plotted as a thematic grid map.

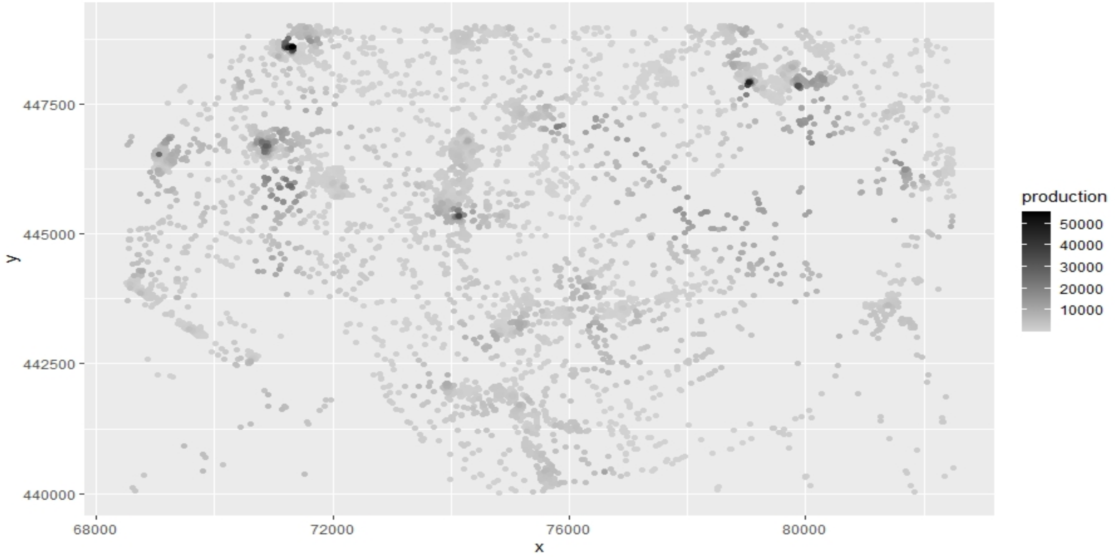

FIGURE 5: Used locations and measurement values

5.1 Absolute error protection

In Figures 6 and 7, we see several grid maps with cell length 500. Comparing Figure 6 and 7, the main result is that a larger interval parameter c, also results in higher variance noise. This result is also clearly visible with different cell sizes when comparing Figure 8 and 9 and comparing Figure 10 and 11. This observation was expected as mentioned in Section 3.1.

Comparing Figures 6, 8, 10 and comparing Figures 7, 9, 11, we see that larger cells, while keeping the interval parameter the same, result in less additional noise in general. This result was also expected, since larger cells result in more contributors per cell, so less protection is necessary.

Note that the bounded protection method $\mathbb{B}_{\text{abs}}$ is a useful addition in Figures 6, 7 and 9. In the other figures, the unbounded protection method $\mathbb{M}_{\text{abs}}$ barely results in negative values, so the bounded protection method is less important to apply in these cases: there is hardly any difference between the grid maps of $\mathbb{M}_{\text{abs}}$ and $\mathbb{B}_{\text{abs}}$.

A worrying trend in the bounded protection method is the Δ. In Figures 6, 7, 8 and 9 the value of Δ is quite high, which barely provides privacy according to Equations (14). The upper bound on δ has this value, because some cells still have zero contributions. This problem is resolved when applying larger cell sizes, which is visible in Figures 10 and 11.

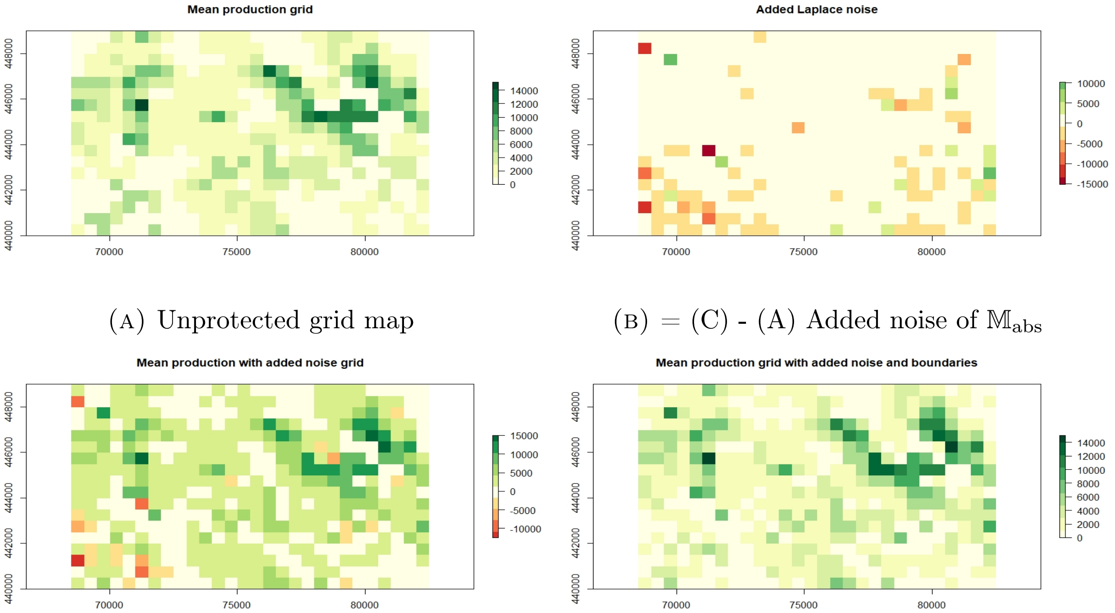

(A) Unprotected grid map

(B) = (C) - (A) Added noise of $\mathbb{M}_{\text{abs}}$

(C) Protected grid map of mechanism $\mathbb{M}_{\text{abs}}$

(D) Protected grid map of mechanism $\mathbb{B}_{\text{abs}}$

Figure 6: Different grid maps with variables r = 500, $\epsilon = 0.1$, c = 50, $\gamma = 59.46$ and $\Delta = 0.536$, getting protected to absolute error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{n\epsilon}{8c}e^{-n\epsilon|x|/4c}$

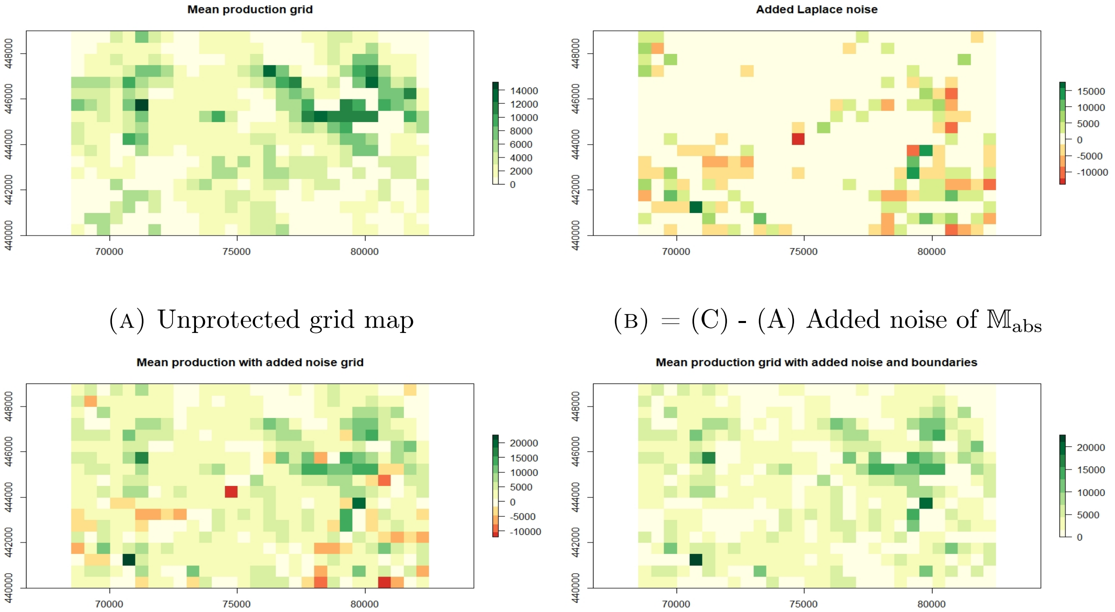

(A) Unprotected grid map

(B) = (C) - (A) Added noise of $\mathbb{M}_{\text{abs}}$

(C) Protected grid map of mechanism $\mathbb{M}_{\text{abs}}$

(D) Protected grid map of mechanism $\mathbb{B}_{\text{abs}}$

Figure 7: Different grid maps with variables r = 500, $\epsilon = 0.1$, c = 100, $\gamma = 59.46$ and $\Delta = 0.544$, getting protected to absolute error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{n\epsilon}{8c}e^{-n\epsilon|x|/4c}$

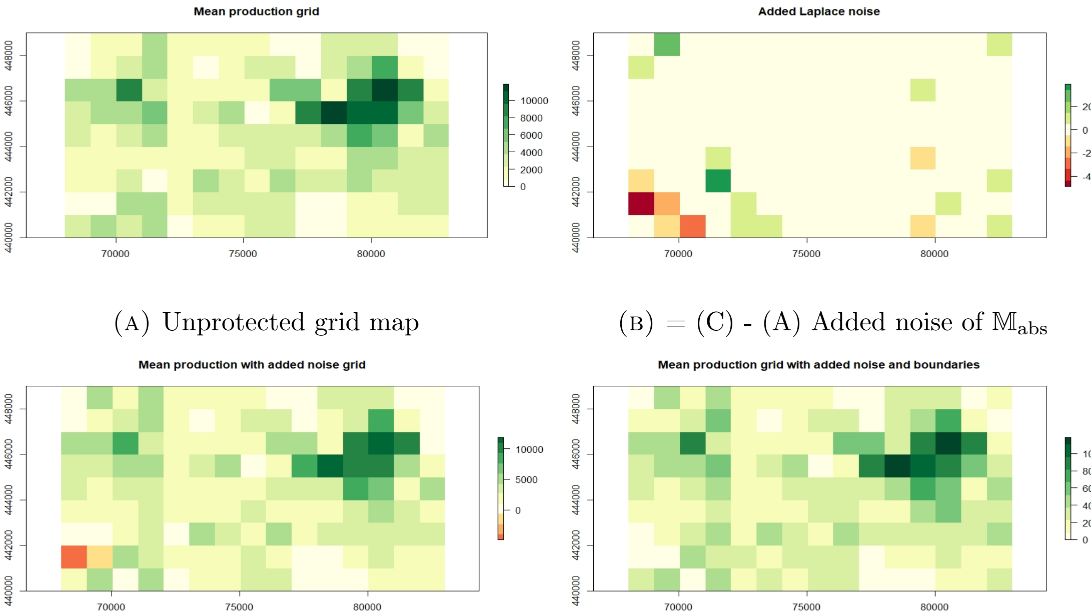

(A) Unprotected grid map

(B) = (C) - (A) Added noise of $\mathbb{M}_{\text{abs}}$

(C) Protected grid map of mechanism $\mathbb{M}_{\text{abs}}$

(D) Protected grid map of mechanism $\mathbb{B}_{\text{abs}}$

FIGURE 8: Different grid maps with variables r $= 1000$, $\epsilon = 0.1$, c $= 50$, $\gamma = 59.46$ and $\Delta = 0.536$, getting protected to absolute error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{n\epsilon}{8c}e^{-n\epsilon|x|/4c}$

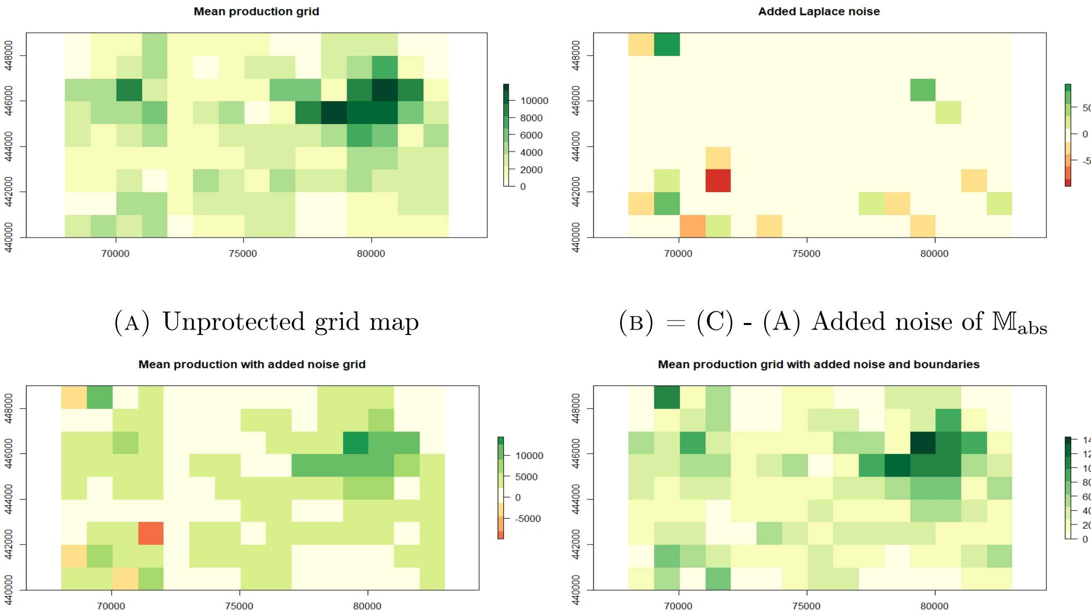

(A) Unprotected grid map

(B) = (C) - (A) Added noise of $\mathbb{M}_{\text{abs}}$

(C) Protected grid map of mechanism $\mathbb{M}_{\text{abs}}$

(D) Protected grid map of mechanism $\mathbb{B}_{\text{abs}}$

FIGURE 9: Different grid maps with variables r $= 1000$, $\epsilon = 0.1$, c $= 100$, $\gamma = 59.46$ and $\Delta = 0.544$, getting protected to absolute error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{n\epsilon}{8c}e^{-n\epsilon|x|/4c}$

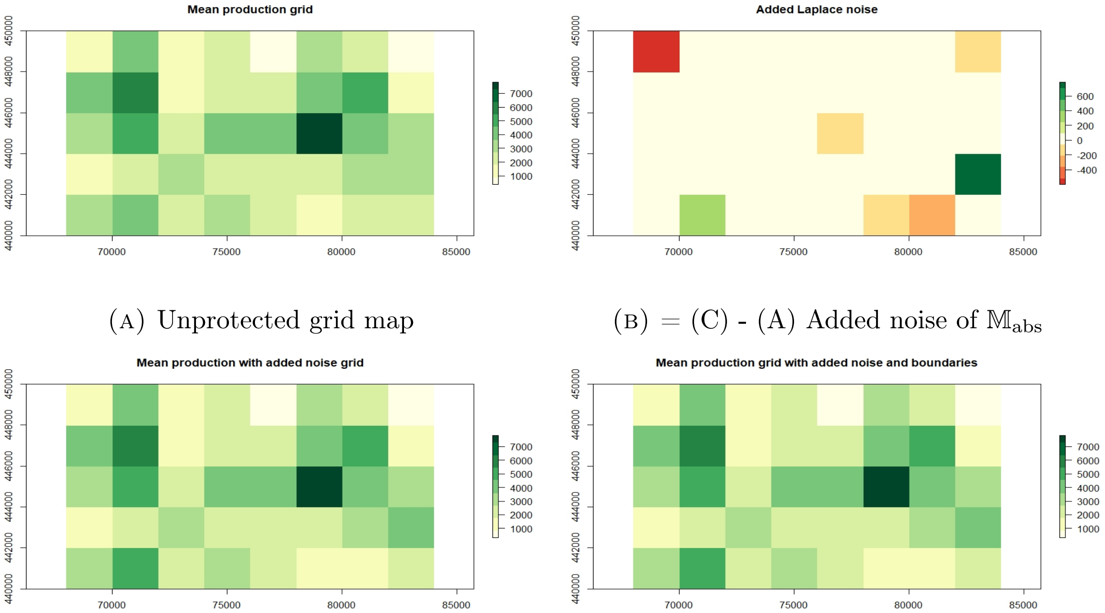
Mean production grid

Added Laplace noise

Mean production with added noise grid

Mean production grid with added noise and boundaries

(A) Unprotected grid map

(B) = (C) - (A) Added noise of $\mathbb{M}_{\text{abs}}$

(C) Protected grid map of mechanism $\mathbb{M}_{\text{abs}}$

(D) Protected grid map of mechanism $\mathbb{B}_{\text{abs}}$

Figure 10: Different grid maps with variables r $= 2000$, $\epsilon = 0.1$, c $= 50$, $\gamma = 59.46$ and $\Delta = 0.261$, getting protected to absolute error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{n\epsilon}{8c}e^{-n\epsilon|x|/4c}$

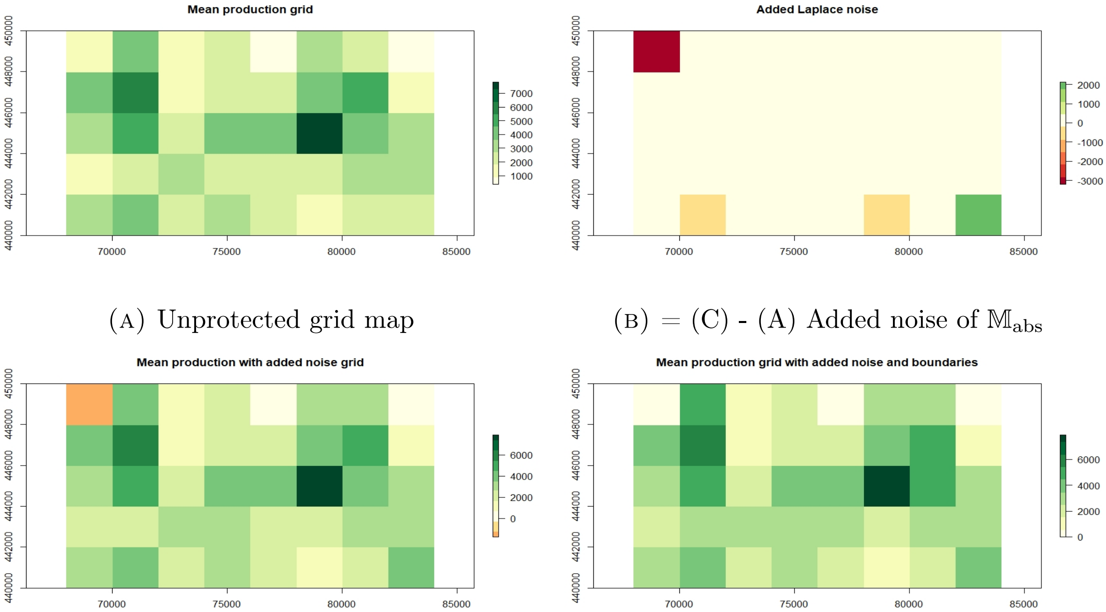
Mean production grid

Added Laplace noise

Mean production with added noise grid

Mean production grid with added noise and boundaries

(A) Unprotected grid map

(B) = (C) - (A) Added noise of $\mathbb{M}_{\text{abs}}$

(C) Protected grid map of mechanism $\mathbb{M}_{\text{abs}}$

(D) Protected grid map of mechanism $\mathbb{B}_{\text{abs}}$

Figure 11: Different grid maps with variables r $= 2000$, $\epsilon = 0.1$, c $= 100$, $\gamma = 59.46$ and $\Delta = 0.379$, getting protected to absolute error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{n\epsilon}{8c}e^{-n\epsilon|x|/4c}$

5.2 Relative error protection

In Figures 12 and 14, we can see that the relative error protection has a high chance to 'explode'. The result of this 'explosion' is clearly visible in Figures 12b and 14b, where some cells have a much larger noise multiplier than the rest. In the case of the relative error of 5%, this causes the protected grid map to be useless to publish. With a relative error of 1%, the noise multipliers are way lower, which results in the more useful protected grid maps of Figures 13 and 15.

For larger cells, the aforementioned problem also has a smaller impact, which also results in well-protected and useful grid maps, visible in Figures 16 and 17.

Note that the bounded protection method $\mathbb{B}_{\text{rel}}$ of Figures 12 and 14 has changed in comparison with its unbounded protection method $\mathbb{M}_{\text{rel}}$, but the utility of Figures 12d and 14d remains lower than preferred. In the other figures, the unbounded protection method $\mathbb{M}_{\text{abs}}$ does not result in extremely high values, so the bounded protection method $\mathbb{B}_{\text{abs}}$ is not necessary to apply in these cases.

In Figures 16 and 17, it follows that $\Delta = 0$, since the grid maps of mechanisms $\mathbb{M}_{\text{rel}}$ and $\mathbb{B}_{\text{rel}}$ are the same. This is better than the absolute protection case, because the same happens in Figure 10, but there it follows that $\Delta = 0.261$. This difference can be explained by the different mechanisms and the different methods of calculating $\delta_{\max}$, according to Theorem 4.4 and Corrollary 4.5.

A peculiarity occurs in Figures 13 and 15, where the unbounded and bounded grid maps are the same, but it follows that $\Delta \neq 0$. It appears that some cells in Figures 13c and 15c are zero, while the lower bound is higher. Therefore, these cells get valued as the lower bound and thus $\Delta > 0$.

Comparing the absolute noise protection with the relative noise protection, the main difference is that the absolute protection method is still significant with small cell sizes, especially after applying the boundaries. Both protection methods are useful with large cell sizes, and the bounded mechanism is barely necessary in that case. When the bounded mechanism is necessary, it appears that the upper bound on δ is quite large. This freedom in the privacy definition results in less protected contributors, and thus these grid maps are still quite sensitive to disclose individual information.

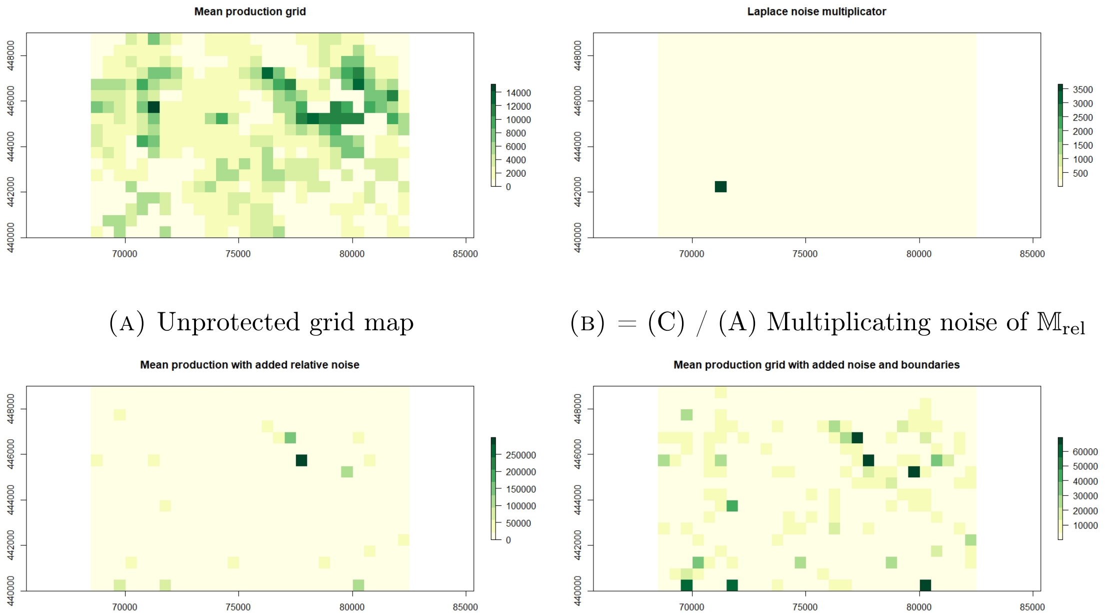

(A) Unprotected grid map

(B) = (C) / (A) Multiplicating noise of $\mathbb{M}_{\text{rel}}$

(C) Protected grid map of mechanism $\mathbb{M}_{\text{rel}}$

(D) Protected grid map of mechanism $\mathbb{B}_{\text{rel}}$

FIGURE 12: Different grid maps with variables r $= 500$, $\epsilon = 0.1$, k $= 95$, $\lambda = 1.25$ and $\Delta = 0.496$, getting protected to relative error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{-n\epsilon}{8\log k}e^{-n\epsilon|x|/4\log k}$

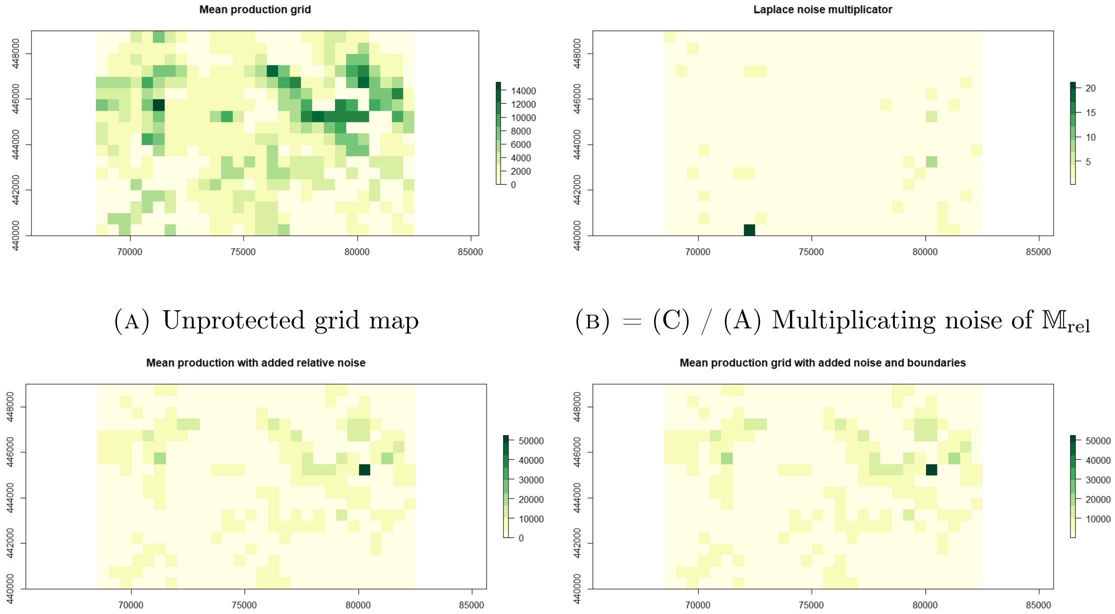

(A) Unprotected grid map

(B) = (C) / (A) Multiplicating noise of $\mathbb{M}_{\text{rel}}$

(C) Protected grid map of mechanism $\mathbb{M}_{\text{rel}}$

(D) Protected grid map of mechanism $\mathbb{B}_{\text{rel}}$

FIGURE 13: Different grid maps with variables r $= 500$, $\epsilon = 0.1$, k $= 99$, $\lambda = 1.25$ and $\Delta = 0.317$, getting protected to relative error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{-n\epsilon}{8\log k}e^{-n\epsilon|x|/4\log k}$

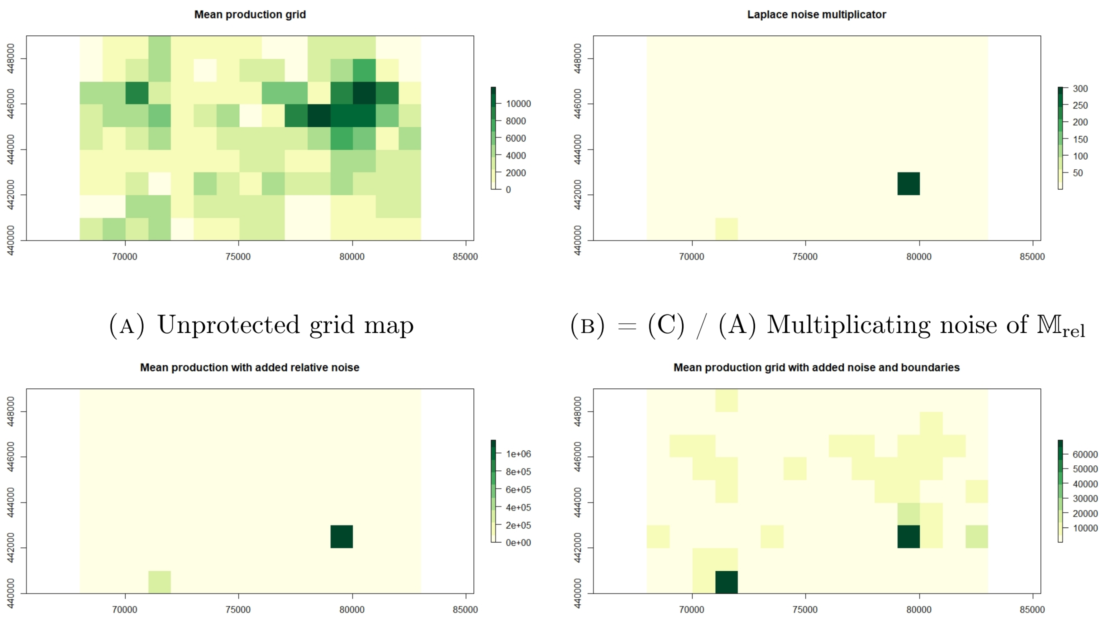

(a) Unprotected grid map

(b) = (C) / (A) Multiplicating noise of $\mathbb{M}_{\text{rel}}$

(c) Protected grid map of mechanism $\mathbb{M}_{\text{rel}}$

(d) Protected grid map of mechanism $\mathbb{B}_{\text{rel}}$

Figure 14: Different grid maps with variables r = 1000, $\epsilon = 0.1$, k = 95, $\lambda = 1.25$ and $\Delta = 0.496$, getting protected to relative error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{-n\epsilon}{8\log k} e^{-n\epsilon|x|/4\log k}$

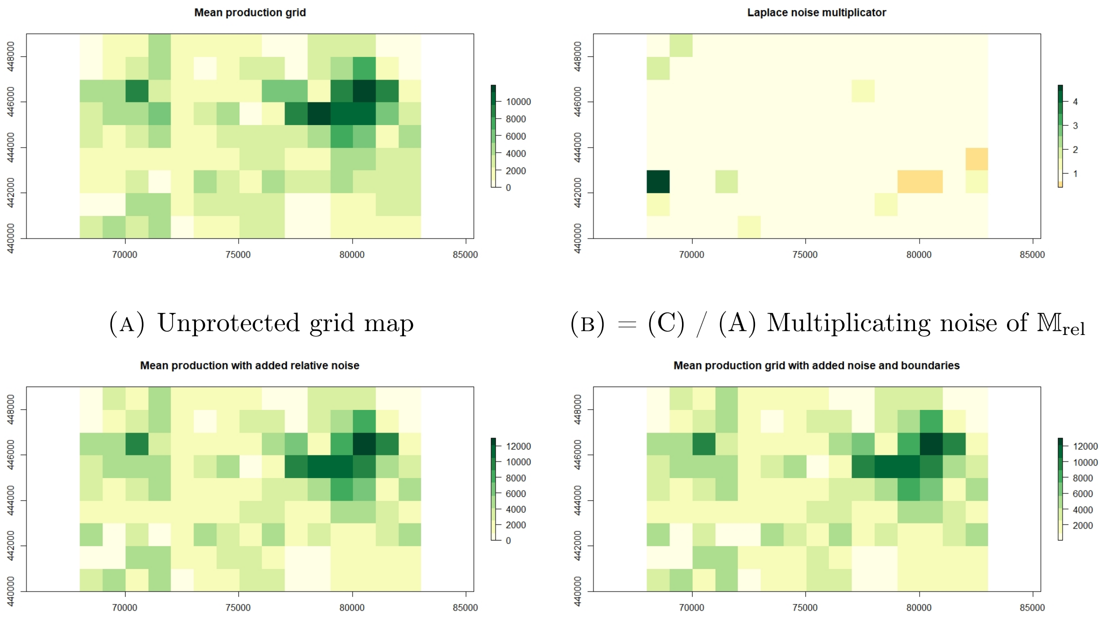

(a) Unprotected grid map

(b) = (C) / (A) Multiplicating noise of $\mathbb{M}_{\text{rel}}$

(c) Protected grid map of mechanism $\mathbb{M}_{\text{rel}}$

(d) Protected grid map of mechanism $\mathbb{B}_{\text{rel}}$

Figure 15: Different grid maps with variables r = 1000, $\epsilon = 0.1$, k = 99, $\lambda = 1.25$ and $\Delta = 0.317$, getting protected to relative error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{-n\epsilon}{8\log k} e^{-n\epsilon|x|/4\log k}$

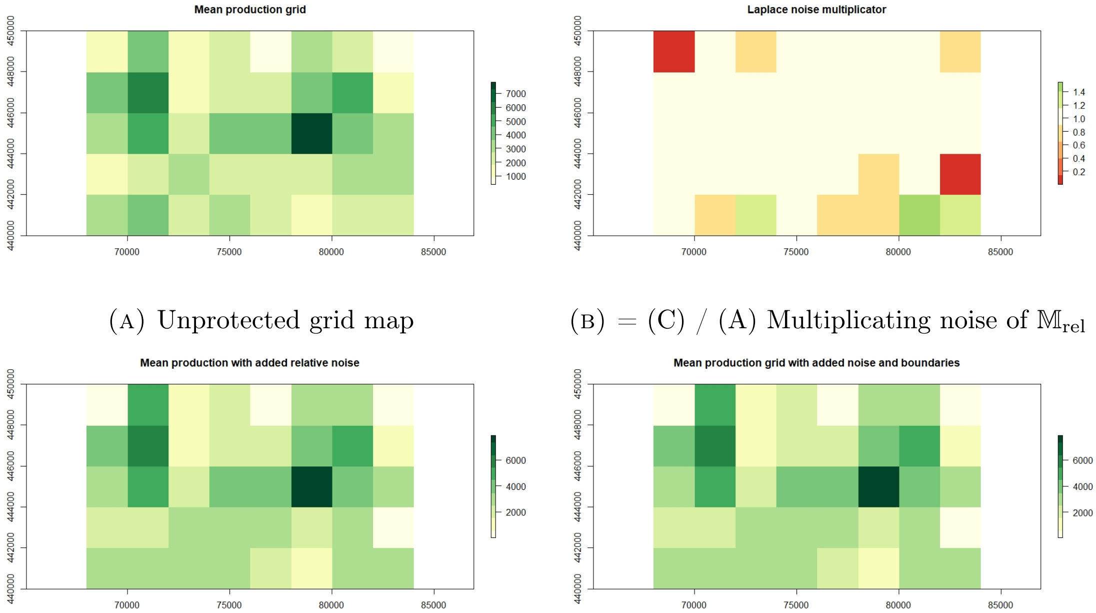

(A) Unprotected grid map

(B) = (C) / (A) Multiplicating noise of $\mathbb{M}_{\text{rel}}$

(C) Protected grid map of mechanism $\mathbb{M}_{\text{rel}}$

(D) Protected grid map of mechanism $\mathbb{B}_{\text{rel}}$

Figure 16: Different grid maps with variables r $= 2000$, $\epsilon = 0.1$, k $= 90$, $\lambda = 1.25$ and $\Delta = 0$, getting protected to relative error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{-n\epsilon}{8\text{log}k}e^{-n\epsilon|x|/4\text{log}k}$

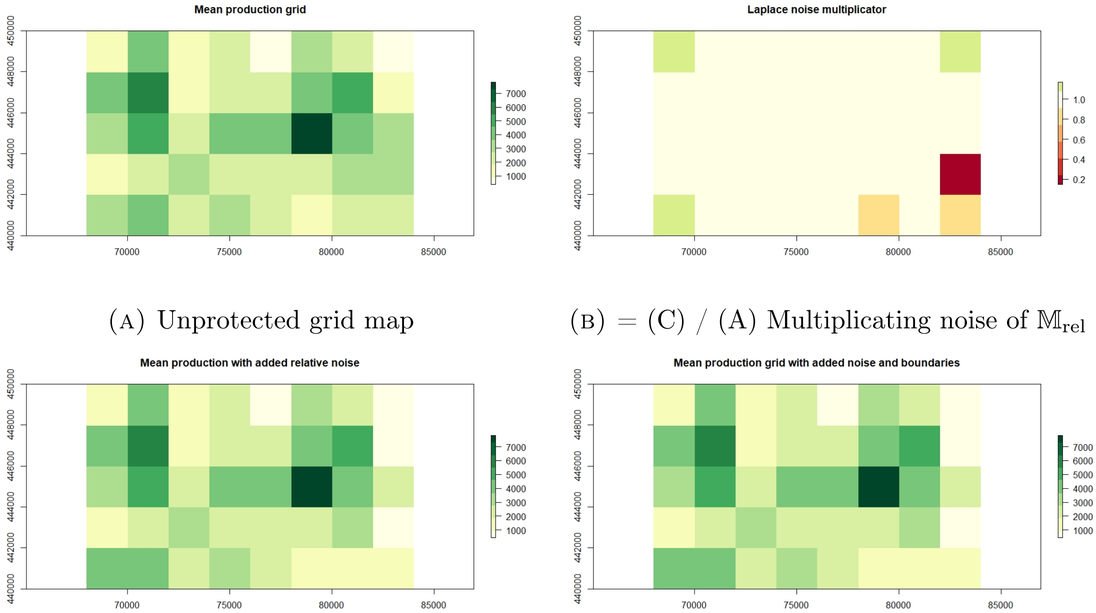
Mean production grid
Laplace noise multiplicator
Mean production with added relative noise
Mean production grid with added noise and boundaries
440000
442000
444000
446000
448000
450000
70000
75000
80000
85000
7000
6000
5000
4000
3000
2000
1000
1.0
0.8
0.6
0.4
0.2

(A) Unprotected grid map

(B) = (C) / (A) Multiplicating noise of $\mathbb{M}_{\text{rel}}$

(C) Protected grid map of mechanism $\mathbb{M}_{\text{rel}}$

(D) Protected grid map of mechanism $\mathbb{B}_{\text{rel}}$

Figure 17: Different grid maps with variables r $= 2000$, $\epsilon = 0.1$, k $= 95$, $\lambda = 1.25$ and $\Delta = 0$, getting protected to relative error, according to the Pufferfish privacy definition. The noise images satisfy density function $\frac{-n\epsilon}{8\text{log}k}e^{-n\epsilon|x|/4\text{log}k}$

6 Discussion & Conclusion

In this paper we looked into the application of the Pufferfish framework when publishing thematic grid maps. Using Laplace distributions, we found the minimal, but appropriate noise that is necessary to protect data. In this way, the utility of the data will remain as high as possible, while all contributing individuals are protected accordingly.

The first protection method made use of additive noise. With this noise, it would be impossible for an attacker to disclose any additional information. The output created with this method was protected well enough, that none of the published statistics overly depends on one of the contributions. This was done using conditional probabilities, where only the condition changed what statement about the data is assumed to be true. An attacker cannot gain additional information by observing the published data.

The main disadvantage of this method is that we have to choose the interval sizes carefully. This is dependent on what kind of data we want to publish, while it would be more ideal to have a general protection method. Hereby, the idea of protection with relative error arose. We noticed that relative protection can be rewritten as an additive case with logarithmic scale, so with a bit of rewriting, we found a method that secures with relative error. This method also satisfied Pufferfish privacy due its robustness to post-processing.

For both methods, we were able to create thematic grid maps. We compared different input parameters with each other, and we noticed that the protection methods worked as expected. For several combinations of input parameters, we saw that its protected grid had high utility, while being safe to publish.

Both absolute and relative protection has its setbacks. Neither negative, nor extremely positive values were expected, which caused some grid maps to be less realistic or useless. We created another privacy mechanism by bounding the results in order to tackle this problem. This new mechanism satisfied a weaker Pufferfish privacy definition. This bounded mechanism increased utility in the absolute error protection method by quite a bit. In the relative case, utility also got higher, but there is quite some room for improvement in this case.

Further research can be done into applying Pufferfish privacy with different data sets. Due to time constraints, we were only looking at one data set, while possibly other observations can be found by using other measurements.

Next to that, more investigation can be done into optimizing the utility of the bounded privacy mechanisms. Combining Pufferfish privacy with some sort of minimum frequency rule from Section 1.1 can improve privacy guarantees by quite a bit. Ideally, the relative error protection should be applicable to thematic maps with small cell sizes, so looking into this can be really valuable.

Finally, combining Pufferfish privacy with the tabular protection methods from Section 1.1 could be an improvement. This way, more global grid privacy can be guaranteed, while its utility reduction can be minimal. The current numerical results made use of additive rounding, but its additional protection was not taken into account. Therefore, further research can be done to investigate in this privacy gain.

References

[1] E. De Wolf P.-P. Han S. De Jonge. *sdcSpatial: Statistical Disclosure Control for Spatial Data.* URL: `https://CRAN.R-project.org/package=sdcSpatial`. (accessed: 17-11-2020).

[2] Cynthia Dwork, Aaron Roth, et al. "The algorithmic foundations of differential privacy." In: *Foundations and Trends in Theoretical Computer Science* 9.3-4 (2014), pp. 211–407. DOI: `http://doi.org/10.1561/0400000042`.

[3] Cynthia Dwork et al. "Calibrating noise to sensitivity in private data analysis". In: *Theory of cryptography conference*. Springer. 2006, pp. 265–284. DOI: `https://doi.org/10.1007/11681878_14`.

[4] A. Hundepool and P.-P. De Wolf. *Method Series - Statistical Disclosure Control. Statistics Netherlands.* 2012. URL: `www.cbs.nl/en-gb/our-services/methods/statistical-methods/output/output/statistical-disclosure-control`.

[5] D. A. Hut. "Statistical disclosure control when publishing on thematic maps." In: (2020). DOI: `http://dx.doi.org/10.1007/978-3-030-57521-2_14`.

[6] Daniel Kifer and Ashwin Machanavajjhala. "Pufferfish: A framework for mathematical privacy definitions." In: *ACM Transactions on Database Systems* 39.1 (2014), pp. 1–36. DOI: `http://dx.doi.org/10.1145/2514689`.