ERASMUS UNIVERSITEIT ROTTERDAM

Name student: Duco Twan Mülder

Student ID number: 434081

Differential Privacy in the Context of Min-Max Optimization Problems

Bachelor Thesis

to achieve the university degree of

Bachelor of Science

Econometrics and Operations Research

Economics and Business Economics

Bachelor Squared

submitted to

ERASMUS UNIVERSITY ROTTERDAM
Erasmus School of Economics

Supervisor: prof. dr. S.I. Birbil

Second assessor: U.C. Cakmak

Rotterdam, July 7th 2019

The views stated in this thesis are those of the author and not necessarily those of Erasmus School of Economics or Erasmus University Rotterdam

Table of Contents

Differential Privacy in the Context of Min-Max Optimization Problems

Duco Twan Mülder

Erasmus University Rotterdam 434081dm@student.eur.nl

Abstract. People's profound concern for sharing information due to privacy loss, decreases the accuracy of statistics. Differential privacy provides a mathematical definition of privacy and potentially increases privacy levels while keeping sufficient accuracy. In this research, we implement differentially private mechanisms in the context of min-max functions. First, we evaluate the performance of Laplacian methods, the exponential method and a private version of the subgradient method, in optimizing a min-max function that generates a convex, piecewise affine optimization problem. We show that the private subgradient method attains best results under most circumstances. Second, we present an algorithm based on the Johnson-Lindenstrauss transformation, to find a min-max shortest path in a differentially private way. Our research shows that the algorithm maintains a relatively good level of optimality, particularly for large graphs. This paper additionally contains an extended related work section about the applications of differential privacy in economic theory and economic applications.

Keywords: differential privacy · min-max optimization · Laplace noise · exponential mechanism · private subgradient method · Johnson-Lindenstrauss Transform · min-max shortest path.

Introduction

With the increase in available data and advancements in technology, data analysis and its applications gain a more prominent role in society. In public transport, for example, the suspension system of a bus can be analyzed in order to estimate the number of passengers in a bus. With that information, one could optimize the bus driving schedule or compute the optimal route (Szauter, Istenes, & Rödönyi, 2016). At the same time, the bus company can acquire information that the passengers did not wish to share, which might include sensitive information such as the passenger's daily activities or even the body weight of the passengers. As data analysis techniques are advancing expeditiously, companies and governments are better able to obtain information on an individual level from aggregate statistics or large data sets. For that reason, society's and policy makers' interest towards privacy has recently grown tremendously.

People's main concern for sharing information is that by joining a dataset, personal information can *leak* and potentially have unfavourable consequences. Narayanan and Shmatikov (2006) found that even when a dataset is anonymized, analysts can observe the change in statistical parameters, and hence derive personal information of the joining participant. Welfare economists show that the loss of statistical accuracy due to too much privacy protection can lead to serious welfare losses to society (Abowd & Schmutte, 2015). For that reason, we investigate mechanisms that assure that joining a dataset cannot lead to exposure of personal information, while maximizing the accuracy of statistical methods.

Privacy as such can be guaranteed by mechanisms that satisfy *differential privacy*; a meaningful and mathematically rigorous definition of privacy useful for quantifying and bounding privacy loss (Dwork, 2011). Following the explanation of Dwork (2008), mechanisms satisfying differential privacy address the concerns that any participant might have about the leakage of her personal information. Thus, when a (potential) participant joins or leaves a dataset, no outputs would become significantly more or less likely, and thus consequences of outputs do not change. Mathematically:

Definition 1. *(Dwork, 2008) A randomized function $\mathcal{K}$ gives ϵ-differential privacy if for all data sets D_1 and D_2 differing on at most one element, and all $S \subseteq$ Range(K),*

$$Pr[\mathcal{K}(D_1) \in S] \leq exp(\epsilon) \cdot Pr[\mathcal{K}(D_2) \in S] \tag{1}$$

Differential privacy has yet shown to be a powerful tool to guarantee privacy and safety in various daily used applications. For example, since the introduction of iOS 10, Apple uses differential privacy in all Iphones to improve QuicktType, Emoji suggestions and Lookup Hints in Notes among others, without compromising individual privacy. [1]. Another example is Google's use of differential privacy in its traffic tracker in Google maps. [2]

Scientific interest in the field of differential privacy started in the beginning of this millennium. As a starting point, Dinur and Nissim (2003) concluded that all (private) information in a database can be revealed by the results of a relatively small number of random queries. This lead to the *fundamental law of information recovery*: a too accurate estimation of too many statistics is generally non-private. From that perspective, researchers have tried to develop methods where the added noise is minimal, while guaranteeing privacy.

In this research, we build on the research from Han, Topcu, and Pappas (2014), who used existing differentially private mechanisms to solve optimization problems with *min-max* objective functions. More specifically, Han et al. (2014) used differential privacy in the context of convex, piecewise affine objective functions. We find that a private version of the subgradient method is most suitable under most circumstances. Next to further evaluation of these methods, this research is extended by finding a differentially-private optimum to a min-max shortest path problem. Our algorithm performs well, particularly for larger shortest path problems.

In the remainder of this paper, we will first explain and further formalize the research problem. Then, in the related work section, we will introduce the implemented mechanisms as well as provide an in-depth economic analysis on how differential privacy can be applied in economic theory and applications. Next, we will describe the methodology applied to answer the research question. The paper continues with the results

[1] https://www.apple.com/newsroom/2016/06/apple-previews-ios-10-biggest-ios-release-ever.html
[2] https://europe.googleblog.com/2015/11/tackling-urban-mobility-with-technology.html

of this research. In the remainder of this paper, conclusions are drawn from the results and the research methodology is discussed. We will evaluate our results and provide suggestions for further research.

Problem Description

In this research, we investigate to what extent optimization problems with a min-max objective function can be solved in a differentially private fashion. The first part of this research is devoted to evaluating differentially private mechanisms in convex, piecewise affine optimization problems with a min-max objective function. The second part investigates to what extend we can apply a differential private mechanism to a min-max shortest path problem. In this section, we formalize the research problem in both contexts. The section ends by summarizing the research problem in a research question.

Min-max objectives: convex piecewise affine functions

In this research, we start by implementing differentially private mechanisms in the context of a convex, piecewise affine objective function. The use of studying this class of functions can be illustrated by the fact that many problems, such as finding the ℓ_∞ or ℓ_2 and also the standard resource allocation problem, can be rewritten as convex, piecewise linear functions. In particular, we try to find a differentially private solution to a minimization problem with a convex, piecewise affine objective function as such:

$$f(x) = \max_{i=1,2,\ldots,m} \{a_i^T x + b_i\} \tag{2}$$

for some constants $\{a_i \in \mathbb{R}^d, b_i \in \mathbb{R}\}_{i=1}^m$. Like the research from Han et al., 2014, we restrict our attention to the case where user information is in $\{b_i\}_{i=1}^m$, such that the database $D = \{b_i\}_{i=1}^m$. We assume both $\{a_i\}_{i=1}^m$ and any constraint set to be public information and fixed.

Like Han et al. (2014), we add constraints in the form of convex Polytope $\mathcal{P}$ creating an optimization problem as such:

$$\min_x f(x) \quad \text{s.t.} \quad x \in \mathcal{P} \tag{3}$$

With the definition of a dataset D as $\{b_i\}_{i=1}^m$ in mind, we can formalize the phrase *"differing on at most one element"* from definition 1. If this is the case, two data sets are considered *adjacent*. When one thinks of a database as a set of rows, databases D_1 and D_2 are adjacent if one is a proper subset of the other and the larger database contains just one additional row (Dwork, 2008). Formally,

Definition 2. *(Han et al., 2014) We define two databases $D = \{b_i\}_{i=1}^m$ and $D' = \{b_i'\}_{i=1}^m$ to be adjacent, denoted by $Adj(\cdot, \cdot)$ if they satisfy the following requirement:*

$$Adj(D, D') \quad \textit{if and only if} \quad \max_{i \in \{1,2,\ldots,m\}} \mid b_i - b_i' \mid \leq b_{max} \tag{4}$$

Finding a private min-max shortest path

After evaluation, we will use the differential private methods for min-max objective functions in the context of graph theory. Like datasets, (directed) graphs could well include sensitive information, such as financial transactions, romantic relations or communication patterns. Differential privacy has therefore been a key subject in graph analysis (Costea, Barbu, & Rughinis, 2013).

In this context, we will try to solve a version of the well-known shortest path problem in a differentially private version. Given a directed network $\mathcal{G} = (V, E, w)$, where V denotes the vertices, E denotes the edges and w the weights belonging to the edges, the classical shortest path problem aims to find the shortest path from source vertex k to all other vertices. In this research however, we construct a min-max version of the shortest path problem similar to the optimization problem as in equation 3.

Though the classical version of the shortest path problem assumes a fixed adjacency matrix, in many applications, the exact value of edge weights is not generally known. In for example evacuation or transport planning, there are multiple scenarios of edge weights w (Ruzika & Thiemann, 2012). Let i be a scenario in the set $1, ..., S$. Each scenario $i \in S$, has different edge weights $w_{e,i}$, hence a different adjacency matrix A_i. The min-max shortest path problem aims to find the shortest path having the worst case weights among all scenarios (Ruzika & Thiemann, 2012). Formally, the objective function becomes:

$$\min_{e \in E} \max_{i=1,\ldots,m} \sum_{e \in E} w_{e,i} \cdot x_e \tag{5}$$

where x_e equals 1 if edge e is chosen in the path and 0 otherwise.

For differential privacy in graphs, we adhere to standard edge differential privacy as formulated by Kasiviswanathan, Nissim, Raskhodnikova, and Smith (2013). First, we define adjacency within graphs; two graphs are edge neighbours if they differ in one edge. Formally:

Definition 3. *(Kasiviswanathan et al., 2013) Graphs $\mathcal{G} = V, E$ and $\mathcal{G}' = V', E'$, are edge-neighbors if $V = V'$ and $E' = E - \{e\}$ for some edge $e \in E'$.*

Thus, within the context of graphs, we define differential privacy by adjusting equation 1 in definition 4 to:

$$Pr[\mathcal{K}(\mathcal{G}_1) \in S] \leq exp(\epsilon) \cdot Pr[\mathcal{K}(\mathcal{G}_2) \in S] \tag{6}$$

Research question

Now that we have formalized the research problem, we can formulate a research question. This research aims to evaluate differentially private mechanisms in the context of min-max functions. First, we evaluate differentially private mechanisms in a min-max function that generates a convex, piecewise linear optimization problem. Second, we evaluate a differentially private mechanism in a min-max function in the context of graph theory. The research question therefore becomes:

To what extent are differentially private mechanisms able to accurately solve min-max optimization problems? That is, what is the optimal way of solving problems with objectives as formulated in equation 3 and equation 5, while keeping differential privacy as defined in definition 1 and 3?

Related work

In this section, we will introduce the differentially private mechanisms as implemented in this research, followed by reporting the current evaluation results. Then, we will introduce a differential private mechanism in graph theory and one of its applications in marketing. Next to the mathematical and algorithmic side of privacy in min-max optimization, economic researchers have tried to place this study in a more applied context; we will discuss privacy as an economic good, followed by an economic interpretation of the min-max function as specified in equations 3 and 5.

Standard mechanisms in differential privacy In this research, we will build on the theory and methods as introduced by Han et al. (2014). In this paper, a framework is demonstrated for three differentially private mechanisms. First, the *Laplace mechanism*, in which noise is added to either the problem data or the solution, is introduced. Second, they describe the exponential mechanism; it guarantees ϵ-differential privacy by randomly reporting the solution according to a probability density function. Finally, Han et al. (2014) implement an ϵ-differentially private version of the classical subgradient method. All their methods are evaluated on a convex, min-max piecewise affine optimization problem.

In their simulations, Han et al. (2014) try to obtain relations between the differentially private optima and two variables; the size of the constraint set c, and the number of affine functions m. They discover for all privacy preserving mechanisms, the expected optimal value grows (worsens) as c increases. Similarly, the optimal value grows when the number of affine function m increases. In these simulations, Han et al. (2014) obtain $\{a_i\}_{i=1}^{m}$ and $\{b_i\}_{i=1}^{m}$ from i.i.d. Gaussian distributions and take a d-dimensional hypercube centered at the origin as a constraint set $\mathcal{P}$

Differential privacy in graphs Within differentially private analysis of graphs, Blocki, Blum, Datta, and Sheffet (2012) have developed an algorithm that generates a differentially private version of a graph. In addition to standard definition of differential privacy as in definition 4, they however add a term δ that weakens the definition. They define $\epsilon - \delta$-differential privacy as such:

Definition 4. *(Blocki et al., 2012) A randomized function $\mathcal{K}$ gives $\epsilon - \delta$-differential privacy if for all data sets D_1 and D_2 differing on at most one element, and all $S \subseteq$ Range(K),*

$$Pr[\mathcal{K}(D_1) \in S] \leq exp(\epsilon) \cdot Pr[\mathcal{K}(D_2) \in S] + \delta \tag{7}$$

Setting δ close to zero approaches the regular definition as in 4. In their research, Blocki et al. (2012) prove that their algorithm outputs a $\epsilon - \delta$-differential private version of a graph. They show that this version obtains good results for $(S, \bar{S})$-cut queries.

Min-max shortest path in marketing The optimization problem as specified by equation 5, is particularly interesting to marketing, as the graph could represent a *social recommendation system.* A system as such describes has a node for each individual i in the system and an edge weight $w_{e_{i,j}}$ that indicates the similarity (or difference) between individuals i and j. Based on the level of similarity, the system can suggest items belonging to individual i to individual j. When the suggestion is done, the system can measure the true similarity by clicking behaviour. Guy (2015) concludes that systems as such increasingly play a role in (digital) marketing.

By minimizing the 'difference' between two individuals i and j in case the clicking behaviour shows maximum difference, we arrive back at the optimization problem 5. In the regular case an individual i has only shown interested in a small amount of items, and/or has a very limited amount of 'relatives' in the system, privacy becomes an issue (Machanavajjhala, Korolova, & Sarma, 2011). Thus, finding an ϵ differentially private version of this algorithm would give a new valuable tool to marketeers.

Privacy as an economic good Since the introduction of Akerlof's famous *Market for Lemons* (Akerlof, 1978), economists naturally consider information as an economic good, for which the law of marginal utility holds similarly to other products. Privacy, the ability to exclude your information from others, can therefore also be considered an economic good.

For example, Acquisti, Taylor, and Wagman (2016) show that the act of not sharing information could have economical value on both an individual level as well as on a level of welfare economics. More particular, Daughety and Reinganum (2010) use the example of an individual checking into a drug or alcohol rehab. When the individual is not able to keep his search for rehab private, the stigma associated with doing so could deter him from seeking treatment, generating costs to both individual and society. Another call for privacy protection could be the fact that privacy as an economic good, might not be allocated efficiently *due to* imperfect (asymmetric) information and bounded rationality; McDonald and Cranor (2010) show that a very limited share of internet users is aware of the extent to which their personal information is collected and identified online. On the other hand, Posner (1977) argues that protection of privacy actually *causes* information asymmetry, causing a market inefficiency. Then, using privacy mechanisms solely transfers the cost of less available information to other market participants.

Differential privacy opportunities in game theory The concept of differential privacy has the potential to change the views and conclusions in standard situations from Game Theory. For example, in the the standard prisoner's dilemma, we conclude that two rational individuals might not cooperate, while it would be in both of their best interest to do so. Corfman and Lehmann (1994) show the existence of this effect in the price setting of advertisement budgets among companies; all companies are better off by lower advertisement budgets, but all companies do have higher advertisement budgets than desired, fearing the competitor to raise budgets and overtake customers.

However, with differential privacy, a recommender mechanism as described by Kearns, Pai, Roth, and Ullman (2014) can be introduced. In a mechanism as such, a third party has the power to suggest strategies on the basis of voluntary participation. Kearns et al. (2014) show that for games with a large number of players, the recommender mechanism could change the equilibrium of the game such that all parties are better off. The logic behind their reasoning is that under differential privacy, players risk less utility loss due to information sharing. After all, under differential privacy, other players are not able to infer private information based on the answer to the query, in this case the recommender's suggestion, whereas they could in the normal situation. Thus, under the assumption that the recommender (or *curator* as often called in the field of differential privacy) is a trustworthy source, a better equilibrium is potentially attained.

Minmax optimization in Microeconomics The most natural interpretation of the minmax optimization problem as defined in both equation 2 and equation 5 is by using a loss function. In both optimization problems, we let $l_i(x)$ be the loss function in scenario i. For optimization problem 2, we have m scenarios, each with a loss function $l_i(x) = a_i^T x + b_i$. For the shortest path optimization problem 5, the loss function equals $l_i(x) = \sum_{e \in E} w_{e,i} \cdot x_e$. As regularly done in microeconomic problems, we assume agents to be *risk-averse*, causing them to consider the loss function in which the loss is maximal, in their planning. Then by minimizing the maximization, we obtain the optimization problems we investigate in this research.

The assumption that agents show a level of risk-aversion that they minimize a *Murphy's law* (worst-case) scenario might seem strong at first. However, in practice, research in several fields shows that agents often behave according to a *min-max-strategy* (or *minimax-strategy*) as such. For example, Sawyer and MacRae (1962) show that agents use a min-max-strategy in certain voting decisions. In sports, Walker and Wooders (2001) show that top tennis players on average choose to serve to the side they believe minimizes the potential loss the opponent could cause with the opponent's best return shot. Within game theory, a minimax-strategy could also include minimizing the maximum utility/profit of the other player, rather than just minimizing one's own maximum loss.

Methodology

In this section, we explain what mechanisms are used to compute differentially private solutions to the min-max problems specified in the problem description. Please see the appendix for a description of the Python code.

Differential privacy in convex piecewise affine functions

The research starts by implementing differentially private mechanisms to the min-max problem constructed by Han et al. (2014). Like in their research, the performance of the mechanisms will be evaluated for different values of the constraint size c and the number of affine functions m. Additionally, we will investigate its performance for different values of ϵ and different locations of the hypercube center. Finally, we will test whether their mechanisms will work when we add a quadratic factor to the objective function.

Laplace mechanism The first mechanism we will evaluate simply adds noise to a query q. In this mechanism, we first define the ℓ_2-sensitivity $\mathit{\Delta}_2$ of q as follows:

Definition 3 (Sensitivity Han et al., 2014)

$$\mathit{\Delta}_2 := \max_{D,D'} \|q(D) - q(D')\|_2 \tag{8}$$

with $\|\cdot\|$ denoting the ℓ_2-norm (equivalent to the Euclidean distance). Then, Han et al. (2014) show that for a given query q, it holds that the mechanism $\mathcal{K}(D) = q(D) + w$, where w is a random vector with a probability distribution proportional to $exp(-\epsilon\|w\|_2/\mathit{\Delta}_2)$, guarantees ϵ-differential privacy. In two different approaches, we add noise to either the private data b or to the problem solution x.

In case we add Laplacian noise to the private data b, privacy is guaranteed by the post-processing rule: once the problem is privatized, obtaining the optimal solution does not change the level of privacy (Han et al., 2014). As in this case, the query is b, the ℓ_2-sensitivity equals $\mathit{\Delta} := \max_{b,b'} \|b - b'\|_2 = \sqrt{m}b_{max}$. We can therefore draw the noise w_P from $exp(-\epsilon\|w\|_2/\sqrt{m}b_{max})$. Alternatively, when we add the Laplacian noise to the problem solution $M_s(D) = x_{opt}(D) + w_s$, the exact value of $\mathit{\Delta}$ is often unavailable. However when sticking to a compact constraint set $\mathcal{P}$, $diam(\mathcal{P})$, provides an upper bound for δ. In the case of the hypercube constraint set used by Han et al. (2014), we can simply use $diam(\mathcal{P}) = \sqrt{d}$ as an approximation of $\mathit{\Delta}$. Note that after adding the noise to the problem solution, we have to re-verify the feasibility of the solution with respect to the given constraint set.

To obtain w from the distribution proportional to $exp(-\epsilon\|w\|_2)$, we draw the magnitude $\bar{w}$ and direction $\hat{e}$ separately. Chaudhuri, Monteleoni, and Sarwate (2011) show that $\bar{w}$ follows a Gamma distribution $\mathit{\Gamma}(d, \lambda)$ and the distribution of $\hat{e}$ is isotropic. As the sum of d exponential(λ) random variables has a $\mathit{\Gamma}(d, \lambda)$ distribution, we draw d i.i.d. samples $w_1, w_2, ..., w_d$ from the exponential distribution $\lambda \cdot exp(-\lambda \cdot w_i)(w_i \geq 0)$ to obtain

$\bar{w} = \sum_{i=1}^{d} w_i$ (Han et al., 2014). We multiply $\bar{w}$ by a direction $\hat{e}$ that can be generated by drawing from the d-dimensional standard Gaussian distribution, followed by normalization.

Exponential mechanism Second, we will implement the *exponential mechanism*, in which ϵ-differential privacy is satisfied by randomly reporting q based on the following probability density function:

$$\frac{exp(\epsilon u(q,D)/2\Delta_u)}{\int_{q'\in Q} exp(\epsilon u(q',D)/2\Delta_u)dq'} \tag{9}$$

where $u(q,D)$ is a required scoring function and Δ_u equals a (global) sensitivity of scoring function u (McSherry & Talwar, 2007). As a natural choice, we propose to use the negative objective function $-f$ as a scoring function. By choosing this scoring function, we can use the lemma shown by Han et al. (2014) that when $u(x,D) = -f(x,D)$, the sensitivity of u for the adjacency relation as defined in definition 2 equals b_{max}. Under those restrictions, equation 9 simplifies to:

$$\frac{exp(-\epsilon f(\tilde{x}_{opt},D)/2b_{max})}{\int_{x\in\mathcal{P}} exp(-\epsilon f(\tilde{x}_{opt},D)/2b_{max}))dx} \tag{10}$$

which is also ϵ-differentially private (Han et al., 2014). This mechanism requires drawing samples from a distribution proportional to a non-negative function. In this research, we will use a Monte Carlo Markov Chain (MCMC) to draw samples by simulating a Markov chain with a stationary distribution equal to the target distribution.

Private subgradient method The last mechanism we wish to implement in this context is the ϵ-differentially private subgradient algorithm. Traditionally, a function g is a subgradient of f at x_0 if and only if for all x:

$$f(x) \geq f(x_0) + g^T(x - x_0) \tag{11}$$

(Han et al., 2014). Then, one can find a subgradient of a piecewise linear function f by finding $k \in \{1, 2, ..., m\}$ such that:

$$a_k^T x_0 + b_k = \max_{i=1,2,\ldots,m} \{a_i^T x_0 + b_i\} \tag{12}$$

To guarantee ϵ-differential privacy however, we apply the exponential mechanism as described above with a scoring function $u_{sub}(i; x_0, D) = a_i^T x_0 + b_i$. We then select the index i^* using the exponential mechanism using formula 10. We then use the computed ϵ-differentially private subgradient in a modified subgradient method that preserves ϵ/k differential privacy in each iteration. As this algorithm iterates k times, the final solution satisfies ϵ-differential privacy. Likewise Han et al. (2014), we fix the number of iterations at 100. A pseudocode of the algorithms can be found below:

In addition to the work of Han et al. (2014), we will optimize over the learning rates α. We will apply a *diminishing step size rule* as recommended by Boyd, Xiao, and Mutapcic (2003); a certain constant α will be

Algorithm 1 ϵ-differentially private subgradient (Han, Topcu, & Pappas, 2014)

```
1. Choose the scoring function $u : \{1, 2, ..., m\} \to \mathbb{R}$ as $u_{sub}(i; x_0, D) = a_i^T x_0 + b_i$.
2. Select the index $i*$ using the exponential mechanism:
       $Pr[i^* = i] \propto exp(-\epsilon u_{sub}(i; x_0)/2b_{max})$
3. Output $a_{i*}$ as the approximate subgradient at $x_0$
```

Algorithm 2 ϵ-differentially private subgradient method (Han, Topcu, & Pappas, 2014)

```
1. Choose the number of iterations $k$, learning rate base $\alpha$ and $x^{(1)} \in \mathcal{P}$
2. For $i = 1, 2, ..., k$, repeat:
   (a) Obtain an $(\epsilon/k)$-private subgradient $g^{(i)}$ using Algorithm 1
   (b) Update $x^{i+1} := x^{(i)} - \alpha_i^i g^{(i)}$
3. Output $x^{k+1}$ as the solution
```

raised to to the power of the current iteration i, so the step size becomes α^i. In this method, we determine the optimal value of the constant by looping over α.

Simulation details In evaluating the Laplace, exponential and subgradient mechanisms, we will adhere to the experiment set-up from Han et al. (2014). In their research, $\{a_i\}_{i=1}^m$ and $\{b_i\}_{i=1}^m$ are both generated from i.i.d. Gaussian distributions. The constraint set is a d-dimensional hypercube centered at the origin, with diameter $2\sqrt{d}c$. The level of privacy ϵ equals 0.1 and the expected objective value is approximated by the average of 1,000 runs.

The MCMC used for the exponential mechanism has a multivariate Gaussian proposal distribution. The covarince matrix of the Gaussian distribution equals $\Sigma = \eta c I_{d \times d}$ where $I_{d \times d}$ equals the identity matrix and $\eta = 0.1$. Note that the covariance matrix is propotional to the size of the constraint set c. The MCMC is considered to have converged after 5000 MCMC steps. Unless stated otherwise, we set $d = 4$, $c = 3$, $m = 50$, $\epsilon = 0.1$ and the location of the center of the hypercube constraint set at the origin.

Extended evaluation After implementing the methods proposed by Han et al. (2014) on the research problem described by 3, we will evaluate the methods effectiveness by changing hyperparameters and generalizing their research problem.

As a starting point, we will verify the results of Han et al. (2014) for the performance of the mechanisms with different values of the size of the constraint set c and the number of affine functions m. In this research, we will extend this evaluation by investigating the methods for different values of ϵ and for different locations of the hypercube center o. Research in investigating the effect of ϵ on the performance of the mechanisms is particularly interesting, since there seems to be a discrepancy between values used in academic literature versus values used in the industry. For example, Han et al. (2014) set $\epsilon = 0.1$, whereas Apple applies differential privacy with ϵ-values equal to 4 or even 8. [3]

[3] www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf

Additionally, we will evaluate the performance of the proposed methods under piece-wise quadratic, rather than piece-wise affine functions. In that case, the optimization problem still has a convex shape. Then, function $f(x)$ in optimization problem 3 changes to:

$$f(x) = \max_{i=1,2,\ldots,m} \{a_i^T x + q_i x^2 + b_i\} \tag{13}$$

Differential privacy in min-max shortest path problem

In order to find a differentially private solution to the min-max shortest path problem, we use the algorithm from Blocki et al. (2012), as specified in 3. The input of this algorithm requires an *Edge matrix* of a graph rather than the adjacency representation of a directed graph. Element $a_{i,j}$ of an adjacency matrix represents the weight between node i and j and is strictly positive in our implementation. An edge matrix E_G, size $(2\cdot\binom{n}{2}\times n)$, has rows belonging to every edge and columns to every vertex in a graph. The entries of E_G are:

$$(E_G)_{(\{u,v\},x)} = \{\sqrt{a_{u,v}} \text{ if } u \sim v \text{ and } x = u \text{ ; } -\sqrt{a_{u,v}} \text{ if } u \sim v \text{ and x=v ;0 else}\} \tag{14}$$

where $a_{u,v}$ represents element u, v from the adjacency matrix.

Once we have obtained the edge matrix E_G, we input it in algorithm 3 to obtain ϵ-differential privacy. Note that we set $\delta = 0.001$ such that our original definition of differential privacy is approached. We initialize the other parameters $\eta = 0.7$ and $\nu = 0.8$

Algorithm 3 Outputting a Graph under Differential Privacy (Blocki, Blum, Datta, & Sheffet, 2012)

```
1. Set $r = \frac{8ln(2/\nu)}{\eta^2}$ and $w = \frac{\sqrt{32rln(2/\delta)}}{\epsilon} ln(4r/\delta)$
2. For every $u \neq v$, set $w_{u,v} = \frac{w}{n} + (1 - \frac{w}{n}) \cdot w_{u,v}$
3. Pick a matrix M of size $r\times$ $2\cdot\binom{n}{2}$, whose entries are i.i.d. samples of $\mathcal{N}(0,1)$
4. Output $\tilde{L} = \frac{1}{r} E_G^T M^T M E_G$
```

Theorem 1. *Algorithm 3 outputs a perturbed Laplacian $\tilde{L_G}$ of the graph which is differentially private as defined in definition 4.*

Proof. Let G and G' be two graphs differing only on one edge according to definition 3. Let (a, b) be the edge in G' that is not present in G. Let $O = ME_G$. In algorithm 3, the output equals $\tilde{L} = \frac{1}{r}O^T O$. This matrix O is composed of r identically distributed rows, as algorithm 3 creates each row by obtaining a $2\cdot\binom{n}{2}$-dimensional vector Y with entries obtained from the standard normal distribution and then taking $Y^T E_G$. Each row maintains (ϵ_0, δ_0)-differential privacy with the right values of ϵ_0 and δ_0.

Claim. Let $\epsilon_0 = \frac{\epsilon}{\sqrt{4rln(2/\delta)}}$, $\delta_0 = \frac{\delta}{2r}$. Then,

$$: P[E_G^T Y(x)] \leq e^{\epsilon_0} P[E_{G'}^T Y(x)] \tag{15}$$

Let then $S = \{x : P[E_G^T Y(x)] \geq e^{-\epsilon_0} P[E_{G'}^T Y(x)]$. Then

$$P[S] \geq 1 - \delta_0 \tag{16}$$

This claim has been proven by Blocki et al. (2012) using the Singular value Decomposition theory of L_G and $L_{G'}$ and Weyl's inequality. We then apply the sequential composition theorem as formulated by Dwork, Roth, et al. (2014) to prove the theorem:

Theorem 2. *(Dwork, Roth, et al., 2014)) Suppose a mechanism M_1 preserves ϵ_1-differential privacy and another mechanism M_2 preserves ϵ_2-differential privacy. Then a new mechanism $M(D) := (M_1(D), M_2(D))$ preserves $(\epsilon_1 + \epsilon_2)$-differential privacy.*

When we apply this theorem for r i.i.d. samples, each preserving $\epsilon_0 - \delta_0$-differential privacy, to the above mentioned claim, the proof is completed. ■

Intuitively, the algorithm takes the parameters ϵ and δ, combined with some other parameters ν and η (not important in our analysis, but used in approximation queries by Blocki et al. (2012)). These parameters insert perturbation in the edge matrix in step 2 of the algorithm. Then, randomness is inserted by a matrix M. The output matrix, a Laplacian, finally preserves $\epsilon - \delta$-differential privacy and can easily be converted back to an adjacency matrix to compute the min-max shortest path.

A non-perturbed Laplacian L_G, which equals $E_G^T E_G$, is the $n \times n$ matrix whose diagonal entries are $(L_G)_{u,u} = \sum_{x \sim u} w_{x,u}$ and non-diagonal entries are $(L_G)_{u,v} = -w_{u,v}$. As we are interested in a version of the shortest path, possible edges from node u to the same node u are irrelevant. We therefore continue by taking the absolute value of all entries of the perturbed Laplacian $\tilde{L_G}$ to solve the optimization problem as specified in equation 5. For fast computation, we can divide each entry by a large positive constant $\mathcal{M}$. Changing the scale of the weights (intuitively, from miles to kilometers for example) does not change the solution of the shortest path, assuming all weights are non-negative.

Similar to the set-up of the previous simulation experiment, we will set the level of privacy ϵ at 0.1 and compute the expected costs of the min-max shortest path by averaging 1,000 runs. Each run, the simulation generates a adjacency matrix, where each entry equals the absolute value of a number generated from the standard Gaussian distribution, assuring non-negative entries. We will investigate the effect of changing the number of nodes n and the number of scenarios S.

Results

In this research, we fill first evaluate the methods proposed by Han et al. (2014). In addition to evaluating them solely over the values of the constraint set c and the number of affine functions m, we will also include an evaluation of the methods for different centers of the hypercube and different values of epsilon. Additionally, we will evaluate the proposed methods under piecewise quadratic functions, rather than linear functions.

Different values of constraint set *c*

Figure 1 plots the objective values of all proposed mechanism over the size of the constraint set c. All differentially private mechanisms increase with c, hence their performance decreases for a larger constraint set. This result shows that a larger c increases the amount of perturbation in the mechanisms. Though the pattern of these results are in line with the results from Han et al. (2014), our implementation outperforms them in the absolute sense.

The large increase in perturbation for the Laplacian methods (yellow and green lines) and the exponential mechanism (violet line) have different causes. The latter, performs worse for higher c, because the distribution from which the solution is drawn is less concentrated around the optimal solution as c grows. For the Laplacian mechanisms, like Han et al. (2014), we see a steep increase with c. Though Han et al. (2014) did not provide explanation for this growth, we have performed further inspection of the results in order to explain this phenomenon.

When inspecting the solutions of the Laplacian mechanisms, it stands out that in both cases, the proposed solutions exceed the boundaries of the hypercube constraint. For that reason, the Laplacian method described by Han et al. (2014) in many cases yields corner solutions. Due to the shape of a convex function, a corner solution of a problem with a larger set of constraints will automatically be larger than the one with a smaller set of constraints, which explains the steep increase of the Laplacian mechanisms over c.

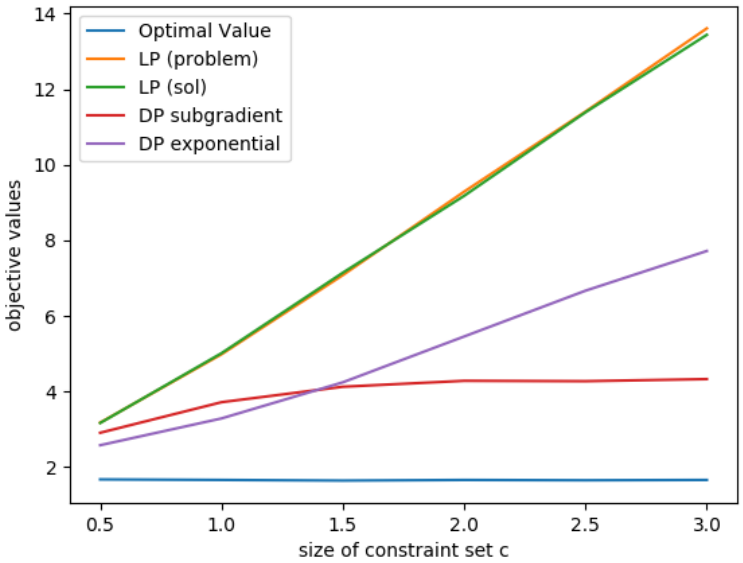

Fig. 1: Objective values plotted over the size of the constraint set c

Different number of affine functions m

Like Han et al. (2014), we also investigate the performance of the methods for different values of the number of affine functions m. The results of our implementation are plotted in figure 2. Again, the results of our implementation follows the same pattern as Han et al. (2014), but yields better results in the absolute sense.

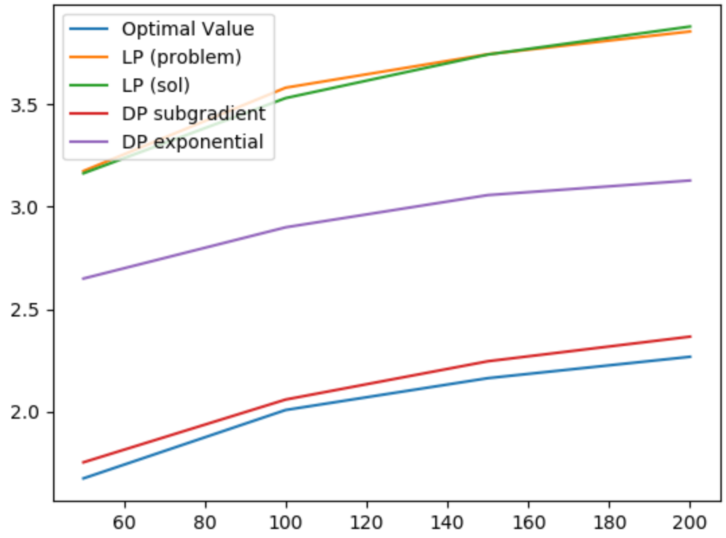

Fig. 2: Objective values plotted over the number of affine functions m

A first observation from figure 2, is that without any privacy mechanism, the optimal value increases over m (blue line). This observation is logical, as for each 'piece' in the function, different values for $a_i^T x + b_i$ are chosen by taking the maximum. As a_i and b_i are generated from a normal distribution, a higher number of generated values will increase the probability of a larger value for both a_i ad b_i, so the mean optimal values will increase over m. The increase logically slows down for larger m, as the probability that $a_i^T x + b_i$ is maximum decreases for larger values of m.

Similarly to Han et al. (2014), we observe that the gap between the differentially private mechanisms and the nonprivate optimal value is relatively constant. We can theoretically verify this suboptimality gap using bounds for the expected suboptimality. For the Laplacian mechanism on the solution, the expected suboptimality is bounded as:

$$E[f(M_s(D)) - f(x_{opt})] \leq \max_i \|a_i\|_2^{\frac{3}{2}} \Delta/\epsilon \tag{17}$$

Han et al., 2014 where Δ is the sensitivity as defined in definition 3, where the query q is x_{opt}. From 17, we see that a bound on the expected suboptimality only depends on m through the objective value, causing the proportional growth. Finally, we will report the solutions of the algorithm that generates a privately computes a min-max shortest path.

Different values of privacy level ϵ

In figure 3, we evaluate the privacy mechanisms for different values of the privacy level of ϵ. Note that the first evaluated point is 0.01 rather than 0. As should be expected, the amount of perturbation in the mechanisms generally decreases for larger ϵ. Note again from equation 4, that a smaller value of ϵ indicates a larger level of privacy.

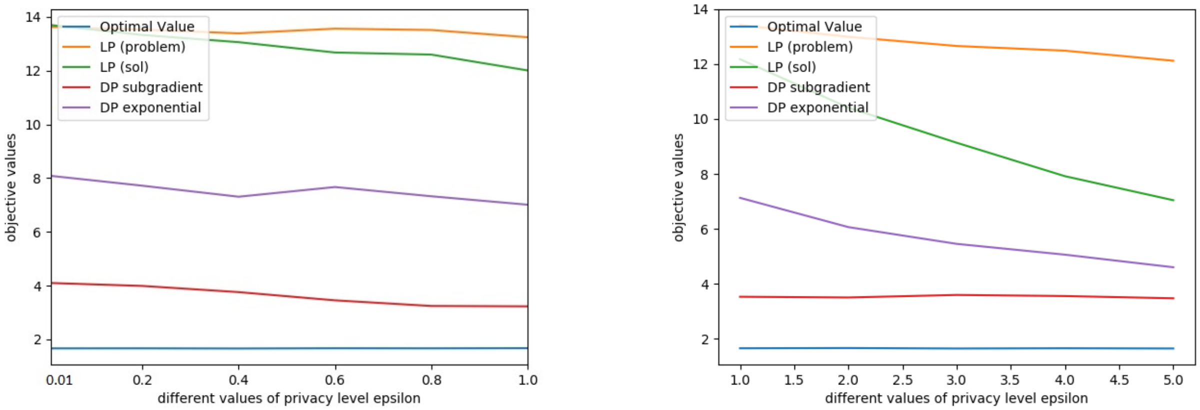

(a) $\epsilon \in [0.01, ..., 1]$

(b) $\epsilon \in [1, ..., 5]$

Fig. 3: Objective values plotted over the level of privacy ϵ

We observe that the private subgradient method is least sensitive for higher values of ϵ. Also note the drastic increase in accuracy for the Laplacian methods for values of $\epsilon \geq 1$.

Different locations of the hypercube constraint

To further generalize the performance evaluation of Han et al. (2014), we have investigated to what extent the differentially private mechanisms differ when we change the center of the hypercube constraint set. Figure 4 plots the (differentially private) objective values over the center of the constraint set. Note from the x-axis that figure 4b reflects the same information, but is a zoomed in version compared to figure 4a. Instead of $c = 3$, we have set $c = 0.5$ in order to increase the visibility of changes in optimal value caused by the change of the hypercube center. This is the case as the difference between objective values is smaller for $c = 0.5$ (please see figure 1 for this result).

We observe a parabola-shaped results, in which we find that the lowest value for the objective functions for a hypercube centered at zero. Due to the min-max setup of our optimization problem, the fact that the true optimum is lowest for a hypercube centered at zero is not surprising, as the maximum value over all a_i

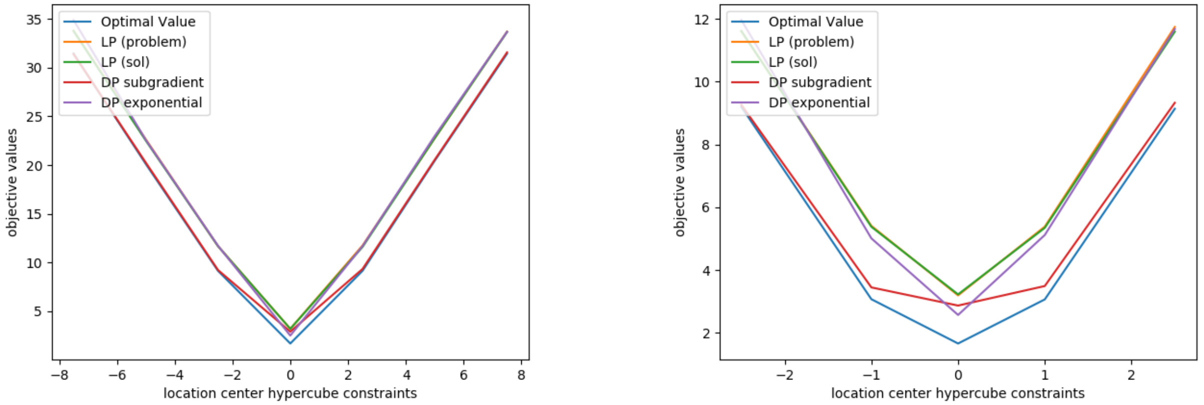

(a) location center $\in [-8, 8]$

(b) location center $\in [-2.5, 2.5]$

Fig. 4: Objective values plotted over different centers of the hypercube constraint set

is multiplied by a smaller value of x. More interesting in figure 4 is the fact that the perturbation by the differentially private mechanisms is relatively constant for different locations of the hypercube.

Including a quadratic term

Rather than the piecewise linear function as tested in Han et al. (2014), we will evaluate the mechanisms for a piecewise quadratic function as defined in equation 13. As the shape of this problem is still convex, the nonprivate problem compared to the tests without a quadratic term yield similar results. The results are plotted in figure 5.

For the same reason, we have set $c = 0.5$ in figure 5b. The parameters (other than c) in figure 5a are according to the simulation specification in the previous section.

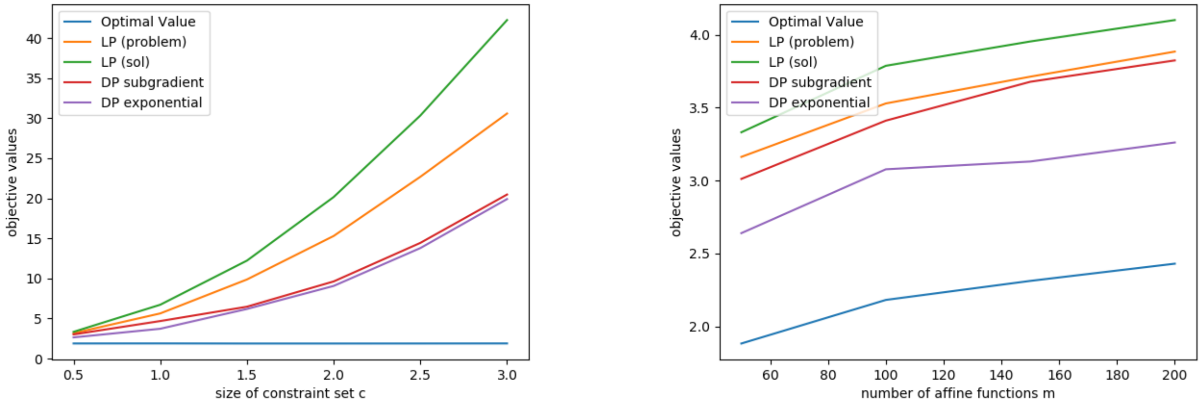

(a) Plotted over the size of the constraint set c

(b) Plotted over the number of affine functions m

Fig. 5: Objective values plotted with quadratic term included

For the private methods, we see that all methods are monotonically increasing for c. However, the shape as well as the size of the perturbation differs from the tests without the quadratic term for all mechanisms. Though both the exponential as well as the private subgradient method yield higher optimal values than before, the exponential mechanism performs relatively better than the private subgradient method, making it the best private mechanism for all values of c in this case.

Differentially private min-max shortest path

In this section, we will show results for implementing differential privacy in a min-max shortest path optimization problem. We will analyze the effect of the number of nodes n and the number of scenarios S. Figure 6 shows our findings. The blue line, 'optimal costs' indicates the costs of the *true* shortest path as computed by Dijkstra's algorithm (mean value over 1000 runs). The orange line, 'private costs', indicates the costs computed based on the original adjacency graph, while taking the route that the Dijkstra's algorithm suggests based on the differentially private version of the adjacency graph as outputted by algorithm 3.

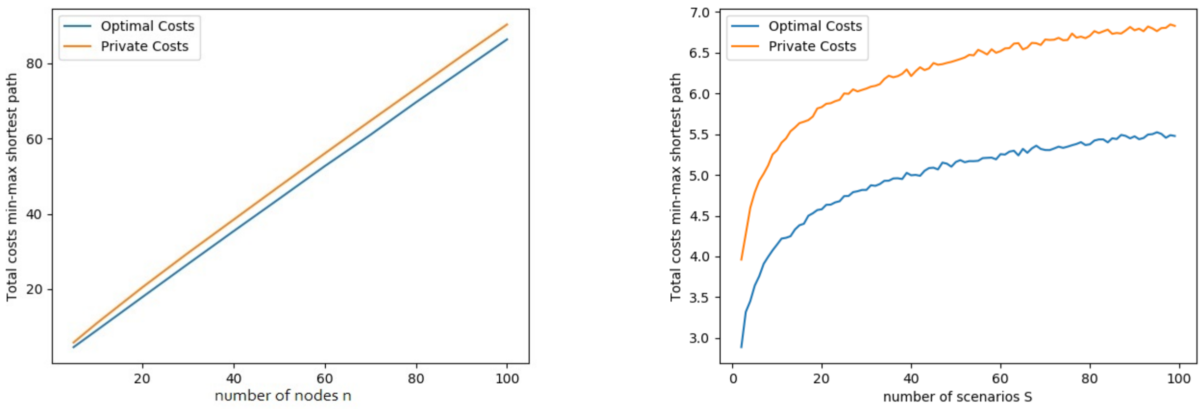

(a) Plotted over the number of nodes n

(b) Plotted over the number of scenarios S

Fig. 6: Costs of the shortest path plotted over the number of nodes n and the number of scenarios S

In figure 6a, we plot the values of the shortest paths over the number of nodes n, where we fix the number of scenarios $S = 20$. Naturally, the costs of optimal path monotonically increase for the number of nodes n. More interesting is the extra costs of taking route based on a differentially private version of the graph remains rather constant for the amount of nodes n.

The effect of increasing the number of scenarios S, while fixing the number of nodes at $n = 10$, is summarized in figure 6b. We see a steep increase for both the value of the true optimal costs as well as the value of the private version. This result for the nonprivate solution makes intuitive sense; as we take the 'worst case scenario' from randomly generated scenarios, the probability that scenario S is worse than all other scenarios $i = 1, .., S-1$ decreases for larger values of S. After all, we draw scenario S from the same distribution as the other $S-1$ scenarios, but the number of scenarios it is compared to, is larger for higher values of S, explaining the increase of costs to slow down. Similar to plotting over n, we again see that the extra costs due to privatization remains relatively constant for when increasing S.

Conclusion & Discussion

First of all, based on the results of this research, we hold a critical attitude towards the use of Laplacian privacy mechanisms in the context of convex, piecewise linear optimization problems. Due to the large noise that these mechanisms add to either the problem data or problem solution, the proposed solution most often exceeds the boundaries of the constraint set c. Because in that case, corner solutions are chosen and the proposed solution is discarded, a linear relation between the objective value and c is observed in figure 1. As this algorithm often simply sets x equal to a corner solution, the Laplacian mechanisms seem of little use in this context. It should also be noted that the research by Han et al. (2014) is limited by implementing existing mechanisms to an optimization problem with a very limited complexity of the constraint set, which they assume to be a d-dimension hypercube.

Next to that, the evaluation of Han et al. (2014) is limited in the sense that the implemented mechanisms are only evaluated over the variables c and m. For those variables, we generally observe the same pattern in results. First, we also see all mechanisms perform worse for higher levels of c. Second, we observe that the perturbation by the mechanisms is constant over number of affine functions m.

In their research, Han et al. (2014) do not refer to the level of dimension d that is used in their analysis. This is likely to be the cause of the slight differences with our results. Our implementation has improved the absolute values of the objectives for all mechanisms. Another slight difference is the result in figure 1 that shows that the exponential mechanism performs better than the differentially private subgradient method for $c \leq 1.0$, whereas Han et al. (2014) report overall supremacy of the latter. Further research could be performed to investigate why and how the dimension d influences the effect that c has on the overall objective.

Our analysis over the different values of the level of privacy ϵ is consistent with the expectation from theory; indeed we see that a higher level of privacy, hence a lower ϵ, causes more perturbation in all differentially private mechanisms. Quite noteworthy is the fact that the private subgradient method is least sensitive to larger values of ϵ, making it the most promising mechanism to guarantee higher levels of privacy. We see that for some large values of ϵ used in the industry, the Laplacian mechanisms drastically increase in performance power. This explains why the industry - most often using a version of the Laplace mechanism - opts for high values of ϵ. We conclude that privacy promises as such by the industry should be well reviewed by academic standards.

Further (economic) analysis could help finding a generic optimal value for ϵ in which we should balance the trade-off between privacy and information sharing. For example, Gupte and Sundararajan (2010) define a loss function depending on the non-private solution, the output of the private mechanism and the degree of privacy they call α, in which a higher value of α indicates a higher level of privacy. Taking $\alpha = \frac{1}{\epsilon}$ would implement differential privacy in their analysis. Then, by researching what loss function should be considered in a certain situation, one could work towards minimizing the loss function with respect to ϵ, creating an optimum value for ϵ.

The results for adding a quadratic term shows that in the new optimization problem, the effects of both c and m are generally similar to the situation without the quadratic term. However, it is interesting to see that for all values of c and m, now the exponential mechanism outperforms the private subgradient method. We have not been able to theoretically explain this observation. The result however could be a starting point for further research in differentially privacy for different objective functions.

Next to the extended evaluation of existing methods, this research has shown that differential privacy can also be implemented for min-max shortest path problems. The results of our implementation are promising, certainly for larger values of the number of nodes n and the number of scenarios S. As the costs of the optimal path increase over both n and S, the amount of perturbation as a percentage of the true solution decreases, making the mechanism more powerful for large values of n and S. However, the running time increases rapidly for larger values of n and S, because the running time of Dijkstra's algorithm is (bounded by) $O(\|n^2\|)$. Further research in differential privacy in (min-max) shortest path problems should therefore include methods to increase the efficiency of Dijkstra's algorithm by for example search trees, Fibonacci heaps and priority queues as presented by Nannicini, Delling, Liberti, and Schultes (2008) among others.

Our implementation of the min-max shortest path problem is limited, due to the relatively simple generation of the adjacency matrix. Rather than solely generating a random positive matrix, it is interesting to investigate how the differentially private mechanism behaves for graphs with a certain specification; one could for example pre-specify the in- and out-degree, or increase the variance of edge weights. The set-up in this research could launch research in this field.

The differential privacy as assured in the implementation of our mechanism is applied in various other (mathematical) fields. In particular, Abadi et al. (2016) apply a differentially private version of stochastic gradient descent in deep learning models. They achieve high training accuracy statistics and therefore demonstrate the high potential of differential privacy in this field.

A last interesting point of further research would be investigating the effect of private information that is not i.i.d. distributed. Though this research shows good simulation results for i.i.d. Gaussian distributions, Kifer and Machanavajjhala (2011) criticize this assumption. They admit that deleting a data point is indeed equivalent to hiding evidence of participation when data is believed to be independently generated. If this is not the case however, they show that differential privacy is not a suitable definition to use. As the independence assumption is a strong one in particular, investigation of privacy mechanisms for correlated data sets should be considered.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 acm sigsac conference on computer and communications security* (pp. 308–318). ACM.

Abowd, J. M. & Schmutte, I. M. (2015). Revisiting the economics of privacy: Population statistics and confidentiality protection as public goods.

Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, *54*(2), 442–92.

Akerlof, G. A. (1978). The market for "lemons": Quality uncertainty and the market mechanism. In *Uncertainty in economics* (pp. 235–251). Elsevier.

Blocki, J., Blum, A., Datta, A., & Sheffet, O. (2012). The johnson-lindenstrauss transform itself preserves differential privacy. *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*. doi:10.1109/focs.2012.67

Boyd, S., Xiao, L., & Mutapcic, A. (2003). Subgradient methods. *lecture notes of EE392o, Stanford University, Autumn Quarter*, *2004*, 2004–2005.

Chaudhuri, K., Monteleoni, C., & Sarwate, A. D. (2011). Differentially private empirical risk minimization. *Journal of Machine Learning Research*, *12*(Mar), 1069–1109.

Corfman, K. P. & Lehmann, D. R. (1994). The prisoner's dilemma and the role of information in setting advertising budgets. *Journal of Advertising*, *23*(2), 35–48.

Costea, S., Barbu, M., & Rughinis, R. (2013). Qualitative analysis of differential privacy applied over graph structures. In *2013 11th roedunet international conference* (pp. 1–4). IEEE.

Daughety, A. F. & Reinganum, J. F. (2010). Public goods, social pressure, and the choice between privacy and publicity. *American Economic Journal: Microeconomics*, *2*(2), 191–221.

Dinur, I. & Nissim, K. (2003). Revealing information while preserving privacy. *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems - PODS 03*. doi:10.1145/773153.773173

Dwork, C. (2008). Differential privacy: A survey of results. *Theory and Applications of Models of Computation*, 1–19. doi:10.1007/978-3-540-79228-4_1

Dwork, C. (2011). Differential privacy. *Encyclopedia of Cryptography and Security*, 338–340. doi:10.1007/978-1-4419-5906-5_752

Dwork, C., Roth, A. et al. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, *9*(3–4), 211–407.

Gupte, M. & Sundararajan, M. (2010). Universally optimal privacy mechanisms for minimax agents. In *Proceedings of the twenty-ninth acm sigmod-sigact-sigart symposium on principles of database systems* (pp. 135–146). ACM.

Guy, I. (2015). Social recommender systems. In *Recommender systems handbook* (pp. 511–543). Springer.

Han, S., Topcu, U., & Pappas, G. J. (2014). Differentially private convex optimization with piecewise affine objectives. *53rd IEEE Conference on Decision and Control*. doi:10.1109/cdc.2014.7039718

Kasiviswanathan, S. P., Nissim, K., Raskhodnikova, S., & Smith, A. (2013). Analyzing graphs with node differential privacy. In *Theory of cryptography conference* (pp. 457–476). Springer.

Kearns, M., Pai, M., Roth, A., & Ullman, J. (2014). Mechanism design in large games: Incentives and privacy. In *Proceedings of the 5th conference on innovations in theoretical computer science* (pp. 403–410). ACM.

Kifer, D. & Machanavajjhala, A. (2011). No free lunch in data privacy. In *Proceedings of the 2011 acm sigmod international conference on management of data* (pp. 193–204). ACM.

Machanavajjhala, A., Korolova, A., & Sarma, A. D. (2011). Personalized social recommendations: Accurate or private. *Proceedings of the VLDB Endowment*, *4*(7), 440–450.

McDonald, A. & Cranor, L. F. (2010). Beliefs and behaviors: Internet users' understanding of behavioral advertising. Tprc.

McSherry, F. & Talwar, K. (2007). Mechanism design via differential privacy. In *Focs* (Vol. 7, pp. 94–103).

Nannicini, G., Delling, D., Liberti, L., & Schultes, D. (2008). Bidirectional a search for time-dependent fast paths. In *International workshop on experimental and efficient algorithms* (pp. 334–346). Springer.

Narayanan, A. & Shmatikov, V. (2006). How to break anonymity of the netflix prize dataset. *arXiv preprint cs/0610105*.

Posner, R. A. (1977). The right of privacy. *Ga. L. Rev. 12*, 393.

Ruzika, S. & Thiemann, M. (2012). Min-max quickest path problems. *Networks*, *60*(4), 253–258.

Sawyer, J. & MacRae, D. (1962). Game theory and cumulative voting in illinois: 1902–1954. *American Political Science Review*, *56*(4), 936–946.

Szauter, F., Istenes, G., & Rödönyi, G. (2016). Spectral analysis of suspension system of a commercial city bus. In *2016 ieee 14th international symposium on intelligent systems and informatics (sisy)* (pp. 67–72). IEEE.

Walker, M. & Wooders, J. (2001). Minimax play at wimbledon. *American Economic Review*, *91*(5), 1521–1538.

Appendix

The programming code can be found in the attached Zip-file. The Zip-file contains five files. Please see a brief explanation about the files in table 1.

Table 1: Files in attached Zip-file and brief explanation

README.txt	Instructions about the code.
main.py	Main program containing initialization of variables, boolean specification of which experiments wil be performed and plotting of output graphs.
helper_functions.py	Helper functions that allow computing the nonprivate, Laplacian and private subgradient solution of the min-max optimization problem specified in equation 2.
exponential.py	Helper functions that perform the Monte Carlo Markov Chain and finally computes the solution of the exponential mechanism to the min-max optimization problem specified in equation 2.
graphs.py	Helper functions to all functions regarding the min-max shortest path optimization problem as specified in equation 5, containing creation of a random graph, making an Edge matrix from an adjacency matrix, solving for non-privacy and differential privacy and Dijkstra's algorithm.