Where's Waldo? Bigger Data, Less Privacy Risk

August 30, 2022

Abstract

To protect individuals' privacy, the General Data Protection Regulation requires firms, researchers, and policy makers to minimize data collection. We propose a framework that combines generative adversarial networks (GANs) to sample data from any marketing data set's complex joint distribution and differential privacy to quantify customers' privacy risk. We apply our framework to two privacy-sensitive marketing applications and consistently find that maximizing data collection enables a reduction in customers' privacy risk while maintaining the ability to derive meaningful insights (a "Where's Waldo" effect). This "Where's Waldo" effect has a reciprocal shape; customers' privacy risk is inversely proportional to an analyst's sample size. Subsequently, we simulate a likelihood-based privacy attack to benchmark our framework against existing methods and find that it outperforms in terms of privacy protection. Collectively, our findings can help firms control and communicate the level of customer privacy risk they allow with stakeholders, reduce customers' privacy risk by maximizing data collection, and share data even under strict privacy regulations. In doing so, this study helps researchers accelerate scientific cooperation and progress both within and outside the field of marketing.

Keywords:
privacy, differential privacy, generative adversarial networks.

INTRODUCTION

Every analysis that a firms perform to derive meaningful insights (i.e., utility) introduces privacy risk for its customers (Dinur and Nissim 2003). Consumers are aware of this privacy risk and perceive privacy concerns (Acquisti, Brandimarte, and Loewenstein 2015; Acquisti, Taylor, and Wagman 2016). These privacy concerns may lead to less willingness to disclose information, less effective personalization, and a decrease in firm and industry performance, as well as an increase in regulatory oversight (Goldfarb and Tucker 2011; Martin, Borah, and Palmatier 2017). To address these privacy concerns, the European Commission (2012) has introduced the General Data Protection Regulation (GDPR), which outlines two data protection directives: data minimization and anonymization.

Existing data protection methods (e.g., removing privacy-sensitive variables, swapping, aggregation) that are used to comply with these directives remain privacy-sensitive. An exciting development is Anand and Lee's (2022) novel deep learning approach that samples "synthetic" data that are only close in distribution to real data but random on an individual level. Unfortunately, even this approach has "negative foreseen impacts on data protection" (European Data Protection Supervisor 2022). For example, Chen et al. (2019) show how intensive care unit patients can still be re-identified from synthetic data. This privacy risk is a direct consequence of Dinur and Nissim (2003)'s Fundamental Law of Information Recovery (FLIR), which states that each time an analyst performs an analysis, they inevitably introduce privacy risk for the individuals included in the data set, and with enough analyses, every individual can be completely reconstructed (e.g., Garfinkel, Abowd, and Martindale 2018). The FLIR teaches us that there exists an inherent trade-off between privacy and utility; the best we can do is to control the privacy risk an individual might incur when being subjected to an analysis. Dwork and Roth (2014) introduce a mathematical definition of privacy that controls this inherent privacy risk: differential privacy.

In this paper, we combine generative adversarial networks (GANs) with the current gold

standard of privacy protection: differential privacy (Dwork and Roth 2014; Gupta, Moutafis, and Schneider 2022). We mathematically limit customers' privacy risk and navigate the privacy-utility trade-off in two increasingly complex marketing applications. In contrast to GDPR's data minimization directive, we show that firms in fact can reduce their customers' privacy risk by maximizing data collection. Our paper makes three contributions to the marketing literature that addresses privacy:

- **Conceptually**, we contribute in two ways: (1) currently, GDPR requires data minimization to protect individuals' privacy. In contrast, we find that firms need to increase a data protection method's underlying sample size (i.e., data maximization). We call this a "Where's Waldo" effect; customers can reduce their privacy risk by hiding in a large crowd. The "Where's Waldo" effect is of a reciprocal shape; customers' privacy risk is inversely proportional to an analyst's sample size. (2) We stress with the FLIR how each time an analyst performs an analysis, they inevitably introduce privacy risk for the customers included. Existing data protection methods (e.g., Danaher and Smith 2011; Schneider et al. 2018; Anand and Lee 2022) fail to limit this privacy risk and are vulnerable to complete reconstruction of a firm's customers. Although these methods do introduce noise to the data, they do not mathematically limit a customer's privacy risk. As a result, firms that use existing methods remain vulnerable to privacy scandals as its customers' privacy risk might grow uncontrollably over time. Our framework proposes a way to navigate the privacy-utility trade-off; bridging the gap between privacy officers and analysts. Ultimately, increasing the sample size aligns the goals of these two parties.

- **Methodologically**, our contribution is twofold: (1) we build on the work of Anand and Lee (2022) by combining GANs with differential privacy. Our framework enables analysts to control, quantify, and interpret the privacy risk that customers pay (or the level of privacy protection) as a consequence of marketing analytics. The majority of the marketing literature on privacy addresses perceptions of privacy that are difficult

to quantify (e.g., Beke et al. 2022; Goldfarb and Tucker 2011; Gupta, Moutafis, and Schneider 2022; Martin, Borah, and Palmatier 2017). Using our framework, firms can transparently communicate the level of consumers' privacy risk that they allow with policy makers (e.g., European Commission 2021). Alternatively, policy makers could consider mandating a universal level of privacy risk. On the academic side, researchers can use our framework to publish data together with a (*Journal of Marketing Research*) publication to allow reproducibility of a study.[1] (2) We develop a (likelihood-based) privacy attack that allows analysts to obtain an empirical estimate of differential privacy's privacy risk for any data protection method. We benchmark our framework against existing methods and find that it outperforms in terms of privacy protection.

- **Empirically**, we provide extensive evidence for the "Where's Waldo" effect in two marketing applications. We navigate the privacy-utility trade-off in a churn setting, where typically privacy-sensitive explanatory variables are used to explain churn, and a pharmaceutical marketing application, where we show that policy makers can benefit from sharing patient data to monitor public policy's effectiveness. In the latter case, the panel nature of the data introduces additional complexity in sampling data of high utility. We estimate a variety of typical marketing models and calculate the deviation between the protected and real parameter estimates to measure utility.

The remainder of this paper is organized as follows. In the following section, we discuss related work on the development of data protection methods in marketing. We formally introduce GANs, differential privacy and give theoretical arguments for the "Where's Waldo" effect in the second section. Next, we apply GANs in the context of two marketing applications and then discuss the privacy-utility trade-off. We conclude with a discussion of limitations and directions for future research.

[1]The data and code are prepared in a Github repository and will be released upon publication for replication.

EXISTING DATA PROTECTION METHODS

To protect consumers' privacy, GDPR and Wedel and Kannan (2016) outline two key principles relating to the processing of personal data: data minimization and anonymization (European Commission 2012, Article 5). Data minimization pertains to limiting the collection and disposal of data (e.g., Holtrop et al. 2017; Zhou, Lu, and Ding 2020), which may impede the goal of academics to derive generalizable results. Data anonymization can be accomplished with non-model-based and model-based approaches (Wieringa et al. 2021). Non-model-based approaches include removing personal identifiable information, recoding, swapping, randomizing data, or k-anonymization (Grewal, Gupta, and Hamilton 2021). The FLIR demonstrates that non-model-based methods either remain privacy-sensitive or are too extreme to derive any utility from the data (Dinur and Nissim 2003). To conserve space, we do not review non-model based approaches and instead refer the reader to Wieringa et al.'s (2021) comprehensive overview. To facilitate a comparison of our approach with existing data protection methods, we present a summary of model-based data anonymization (see Table 1 for an overview of the relevant marketing literature).

Table 1: An overview of the model-based data protection methods from the marketing literature.

Study	Method	Multivariate	Assumption(s) Free	Privacy Definition	Privacy Risk Control
Danaher and Smith (2011)	MCMC	✓	✗	✗	✗
Schneider et al. (2017)	MCMC	✗	✗	κ-differential privacy	✗
Schneider et al. (2018)	MCMC	✗	✗	Bayesian parameter shrinkage	✗
Anand and Lee (2022)	GANs	✓	✗	Maximum loss of privacy (MLP)	✗
This study	GANs	✓	✓	ε-differential privacy	✓

Note. MCMC, Markov Chain Monte Carlo.

Marketing literature

An early attempt to model-based data generation is Danaher and Smith's (2011) introduction of copula models to the marketing literature. Copula models allow us to learn complex joint distributions, especially to account for the complex mixture of continuous

and discrete variables that characterize marketing data. More formally, if we have m random variables $(X_1, \ldots, X_m)$ with distribution functions $(p_1(X_1), \ldots, p_m(X_m))$, we want to learn the joint distribution $p(X_1, \ldots, X_m)$. Sklar (1959) proposes that there always exists a copula function C that captures the dependence between the random variables, that is, $p\left(X_1 = x_1, \ldots, X_m = x_m\right) = C\left(p_1\left(x_1\right), \ldots, p_m\left(x_m\right)\right)$ for all $x_i \in X_i$ (for $i = 1, \ldots, m$). Importantly, the functional form of this copula function is independent from the distributional family of the marginals' distribution; it only dictates the dependence between the marginals. Danaher and Smith's (2011) introduction of copula models represents a seminal development because previous methods could not accurately account for such combinations of complex distributional families. The authors rely on Bayesian estimation with Markov chain Monte Carlo (MCMC) sampling to sample from a posterior distribution (see Table 1). A downside of this Bayesian procedure is that it requires us to make assumptions (priors) and introduces difficulties with respect to its convergence for complex dgps.

Schneider et al. (2017) lay the foundation of model-based privacy research in marketing by explicitly studying the privacy-utility trade-off in a specific case: segmentation. The authors aim to segment customer data with theoretical privacy guarantees. They draw samples from a Dirichlet-multinomial model with a privacy parameter κ (kappa) to generate protected data for each segment. A high κ results in a small level of privacy protection but high utility, and a low κ results in a high level of privacy protection but low utility. The study is highly innovative in that it is inspired by differential privacy, but it does not exactly satisfy differential privacy (see "Privacy Definition" in Table 1). Schneider et al. (2017) provide protection for a very specific marketing case and must rely on MCMC (see Table 1); in addition, during identification of the segment of customers, a low parameter κ substantially reduces the ability to derive meaningful insights.

Schneider et al. (2018) extend this earlier work by developing a Bayesian shrinkage model with a privacy protection parameter κ. They argue that this model provides privacy protection by shrinking (with parameter κ) the variance of the parameters' posterior distribution.

For a range of values of κ, the authors obtain a posterior distribution of parameter estimates using MCMC sampling and multiply these estimates with the original data to obtain a posterior distribution of store sales. Consequently, the posterior distribution of store sales is protected through shrinkage of the parameter estimates and can be shared. The authors note that the study's main limitation is that only a small number of variables can be included with their approach. Furthermore, the methodology only allows for sampling a single protected variable. To sample this single protected variable, the methodology requires careful specification of a functional form. For example, in the specific case of store sales, Schneider et al. (2018) assume that the dgp of store sales follows a SCAN*PRO specification (Leeflang et al. 2015); however, if the functional form is inaccurate, sampling from the posterior distribution of parameter estimates might lead to inaccurate data. Furthermore, the parameter estimates that result in the posterior distribution are not robust to violations of econometric assumptions.

Anand and Lee (2022) propose a novel deep learning approach to privacy protection. They use GANs to learn a marketing data set's dgp and solve marketing problems. This allows them to sample customers that are random on an individual level but equal in distribution to the real data. Our paper differentiates itself in the following ways: (1) the authors assume that because a GAN's generator does not have direct access to the real data, it protects privacy. However, there is nothing that restricts the generator from overfitting the real data; risking the re-identification of real individuals (e.g., Chen et al. 2019). To measure this re-identification risk (2), they operationalize privacy with the re-identification probability of a customer in the original data set. To estimate this probability, the authors make a strong assumption that re-identification risk (or privacy protection) is a (linear) function of the variables that are in the data set (see "Assumption(s) Free" in Table 1). In doing so, the authors implicitly ignore the FLIR's implications, privacy risk that might come from external data sets or variables (e.g., Narayanan and Shmatikov 2008), increasingly powerful privacy attacks (e.g., Carlini et al. 2021), and events in the future that potentially increase

this re-identification probability (see "Privacy Risk Control" in Table 1). For example, Chen et al. (2019) consider Anand and Lee's 2022 scenario where only a generator is shared and find that individuals can still be re-identified, especially, with increasingly large synthetic samples from the generator. Ultimately, Chen et al. (2019) finds that differential privacy successfully limits this privacy risk and protects privacy.

More importantly, (3) if firms would rely on their method to share data, a firm's customers' privacy risk is unbounded. For example, if the generator is shared, there is no control over the number of synthetic samples. Therefore, there is nothing that restricts the privacy risk to go to infinity over time as more synthetic samples are released. This is just one example of a privacy attack which shows how firms remain vulnerable to privacy scandals (e.g., Facebook and Cambridge Analytica or Netflix's prize). Clearly, GANs without differential privacy have "negative foreseen impacts on data protection" and "remain vulnerable to privacy attacks" as the European Data Protection Supervisor (2022) remarks.

In contrast, our framework mathematically proves a limit on the amount of customer's privacy risk and addresses all of these issues mentioned above (also see the "Differential privacy" section). In addition, our framework allows entities to control, interpret and transparently communicate the customers' privacy risk. Importantly, we provide guidelines to further reduce customers' privacy risk, such as data maximization (see "Privacy Risk Control" in Table 1). Also, we provide a privacy attack that allows us to estimate differential's privacy risk in practice. Finally, in terms of utility, Anand and Lee (2022) do not modify a GAN to learn customer dynamics, which we address in our (panel) pharmaceutical marketing application.

Computer science literature

The computer science literature largely focuses on developing algorithms that satisfy differential privacy. Abadi et al. (2016) develop differentially private stochastic gradient descent (DPSGD) to limit individuals' privacy risk in training neural networks. DPSGD

samples for each training iteration a mini-batch from the training set, clips the gradients of each individual in the l_2-norm to control the individuals' contribution to the gradient, and subsequently carefully adds Gaussian noise (also see Algorithm 1 in Abadi et al. 2016). Compared with existing approaches, they allow us to train machine learning models (e.g., neural networks) and reduce the utility costs of increasing privacy protection. In other words, they obtain a tighter estimate of the privacy risk for the same number of training iterations (the privacy risk ε over training iterations grows less than linear). Another contribution is that we can set, a priori to training, a desired level of privacy risk ε_0.

Xie et al. (2018) develop GANs that satisfy differential privacy. They do not use DPSGD but do use some of the underlying ideas of DPSGD (e.g., gradient clipping). The authors are interested in the effect of varying levels of differential privacy ($\varepsilon = \infty, 29, 14, 9.6$) on utility. In two experiments, they measure utility by (1) the visual appearance of handwritten digits, (2) the GAN's training progress, and the predictive accuracy of (3) handwritten digits and (4) a disease. Importantly, the main contribution of Xie et al. (2018) is to prove that their GAN can satisfy differential privacy.

Papernot et al. (2018) introduce Private Aggregation of Teacher Ensembles (PATE) in an attempt to improve both privacy protection and utility. The idea is to first partition the data set one wishes to protect and train a classification model on each partition. Papernot et al. (2018) call these classification models "teachers". To satisfy differential privacy, the authors carefully introduce calibrated noise into the teachers' outputs. Subsequently, PATE requires a public data set (that is assumed to be private because it is already public) and uses the teachers to label the public data. Because the labels satisfy differential privacy, any model that is trained on the differentially private labels and public data also satisfies differential privacy (Papernot et al. 2018). The main contribution is that by changing the location where we introduce noise (teachers' predictions instead of the gradients in Abadi et al. 2016), we can both improve privacy protection and utility. Yoon, Jordon, and van der Schaar (2019) translate this PATE framework to GANs. However, to use this framework

analysts must rely on the availability of public data and assume that this does not increase the privacy risk.

In this study, we use DPSGD to satisfy differential privacy. This prevents analysts from relying on publicly available data, which can be extremely difficult or even impossible to obtain in marketing applications (e.g., physicians' prescriptions to patients) or because of privacy regulations. Our study differs from the aforementioned studies in that we investigate the ability to derive meaningful insights with typical marketing models and the managerial implications of differential privacy. For example, we find that data maximization implies less privacy risk (or stronger privacy protection). In addition, we simulate a privacy attack in the context of a marketing application to empirically verify differential privacy's promises.

METHODOLOGY

Dinur and Nissim (2003)'s FLIR shows how an adversary can completely reconstruct a database by simply requesting a number of random statistics (e.g., Garfinkel, Abowd, and Martindale 2018). Importantly, this remains possible even if some noise is added to the calculated statistics. This noise can be introduced by existing (non-)model-based data protection methods such as swapping, k-anonymity, Bayesian shrinkage, or GANs without differential privacy. Although these methods do introduce noise into the data, it is impossible to control this noise in such a way that we can formally prove privacy protection and thus impossible to quantify the privacy risk (European Data Protection Supervisor 2022; Lin, Sekar, and Fanti 2021). Dwork et al. (2006) develop differential privacy, which provides a way of adding noise in a controlled way such that we can prove privacy protection.

Differential privacy

Differential privacy considers two hypothetical scenarios: one in which a person is included in a data set versus one in which only this person is excluded from the data set. Informally, differential privacy introduces a distribution of noise (or uncertainty) around

the possible outcomes (from both scenarios) of any analysis. Noise is carefully added in such a way that the distance between the distributions of the outcomes of the two scenarios can be controlled. For a smaller distance, all the possible outcomes that might arise are almost equal, and thus, it becomes increasingly difficult to infer from the outcomes whether an individual's data were used in the analysis. Importantly, this remains equally difficult regardless of how many statistics one might compute (i.e., it controls the privacy risk from FLIR), what kind of privacy attacks are developed (e.g., Carlini et al. 2018; Fredrikson, Jha, and Ristenpart 2015; Liu, Tan, and Garg 2020) or what external information might become available (e.g., McSherry and Mironov 2009; Narayanan and Shmatikov 2008). Formally, differential privacy is defined as follows:

Definition 3.1. (Differential Privacy) Any randomized algorithm $\mathcal{A} : \mathcal{D} \rightarrow \mathcal{O}$ with a space of all possible data sets $\mathcal{D}$ for any subset O of outputs of $\mathcal{A}$ is differentially private if for any two adjacent data sets D and D' (D, $D' \in \mathcal{D}$), the following holds:

$$\mathbb{P}[\mathcal{A}(D) \in O] \leq e^{\varepsilon}\mathbb{P}[\mathcal{A}(D') \in O] + \delta, \tag{1}$$

where $\mathbb{P}[\mathcal{A}(D) \in O]$ is the probability that the output of algorithm $\mathcal{A}$ using data set D is in O using data set D and D' is an adjacent data set that differs at most for one individual (see Dwork and Roth (2014) for a textbook reference).

The most important parameter in Equation 1 is ε, which can be interpreted as the privacy risk. Differential privacy promises that an individual's privacy risk can increase *at most* with a factor e^{ε}. A small ε implies a small difference between the two outcome distributions (from both scenarios). Hence, the privacy risk for an individual is small, but the utility for an analyst may be low due to the addition of noise that is required to satisfy a small ε. A large ε implies a large difference between both scenarios and thus a relatively high data utility at the cost of a large privacy risk. Therefore, we can view ε as a quantification of the privacy-utility trade-off (Wieringa et al. 2021). Importantly, differential privacy does not

imply a complete elimination of the privacy risk of an individual, but rather a bound on the privacy risk that results from the FLIR (Dinur and Nissim 2003). In empirical applications, a complete elimination of privacy risk ($\varepsilon = 0$) is considered useless, because the data utility is completely destroyed by the addition of noise to satisfy such a level of differential privacy. The original definition was introduced without δ. The introduction of δ relaxes ε-differential privacy.[2]

GANs

A GAN can be explained by the idea of a game between two players. For illustrative purposes, we explain the general principle of a GAN using a typical application: image generation. The game is characterized by a competition between the generator G who tries to create images from random noise, and the discriminator D, who acts as a detective to classify the images as either false or real (see Figure 1).

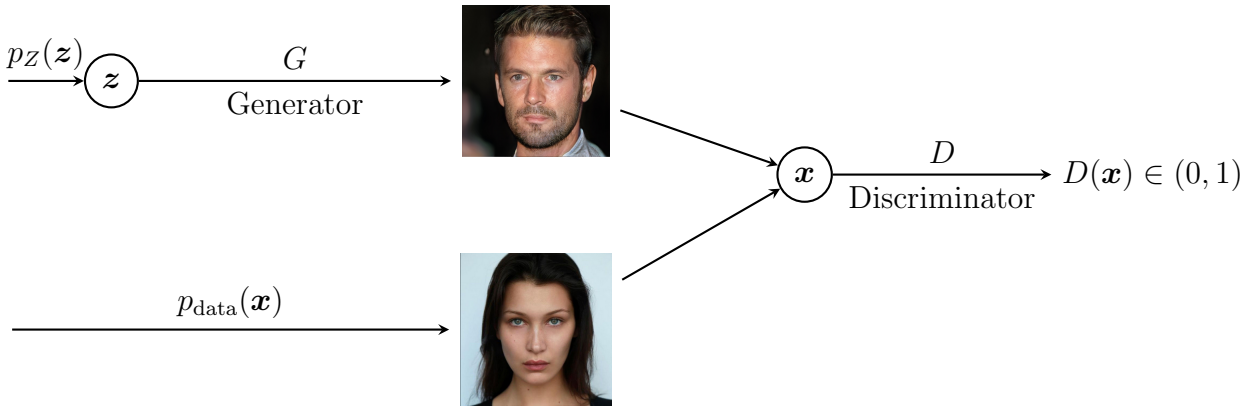

Figure 1: A visualization of the two adversarial players that define a GAN. The generator receives feedback from the discriminator on how to transform its weights to improve the image quality.

The goal of the game for the generator is to create images of sufficient quality that fool the discriminator into thinking they are real. The generator learns to transform random

[2]To illustrate why δ is introduced in Equation 1, consider a scenario in which an extremely rare outcome occurs, which implies a small value for $\mathbb{P}[\mathcal{A}(D') \in \mathcal{O}]$. In such a scenario, we would require a very large ε to satisfy the definition. δ allows us to protect privacy with a lower ε in such rare occasions. In line with the literature, we treat δ as a constant and set it to the inverse of the number of observations n in our sample (Abadi et al. 2016).

noise into high-quality images by maximizing the discriminator's probability of making classification errors. Hence, the generator's ability to generate high-quality images affects the discriminator's ability to classify the images as either false or real. Similarly, the discriminator's classification ability affects the generator's ability to generate high-quality images. Thus, this competition drives both players to improve their respective performances. When the performance of the two players converges as a result of the competition between them, the generator's images are indistinguishable from the real images.

Formal objective

Formally, both players are functions represented by neural networks. Generally speaking, a neural network $N_{\boldsymbol{\theta}}$ admits n input units and maps this into k output units, and the network is parameterized by its weights $\boldsymbol{\theta}$. We thus consider a function $N_{\boldsymbol{\theta}} : \mathbb{R}^n \to \mathbb{R}^k$. Let Z be an i.i.d. standard normal random variable with the distribution $p_Z(\boldsymbol{z})$. The generator is defined as a neural network G that maps noise samples $\boldsymbol{z}$ from the distribution $p_Z(\boldsymbol{z})$ into samples $G(\boldsymbol{z})$ aimed to fool the discriminator. We can think of G as a random variable based on a function $G : \mathbb{R}^n \to \mathbb{R}^k$, where n is the dimensionality of the noise sample and k is the dimensionality that the discriminator requires. In empirical applications, n is at least as large as k, and if we expect that the dgp is very complex, we may make n substantially larger than k. The distribution of G is defined as p_G.

Let p_{data} denote the distribution of the dgp, real distribution or ground truth. In practice, we do not have direct access to p_{data}, but only to a number of samples from p_{data}. The discriminator is defined as a neural network D that takes samples of equal size from p_{data} labeled with 1 and p_G labeled with 0 from the generator as inputs. Similarly, we can think of the discriminator as a random variable D that is a function $D : \mathbb{R}^k \to (0, 1)$, where k is the dimensionality of the data set. The discriminator predicts the probability of a sample to be from p_{data} or from p_G. The distribution of D is defined as p_D. These functions compete

in a zero-sum minimax game with the following objective (Goodfellow et al. 2014):

$$\min_G \max_D V(G, D), \tag{2}$$

where the value function is defined as

$$V(G, D) := \mathbb{E}_{\boldsymbol{x} \sim p_{\text{data}}(\boldsymbol{x})}[\log(D(\boldsymbol{x}))] + \mathbb{E}_{\boldsymbol{z} \sim p_Z(\boldsymbol{z})}[\log(1 - D(G(\boldsymbol{z}))]. \tag{3}$$

In the first part of Equation 3, $\mathbb{E}_{\boldsymbol{x} \sim p_{\text{data}}(\boldsymbol{x})}$ refers to the expected samples $\boldsymbol{x}$ from p_{data}. The term $\log(D(\boldsymbol{x}))$ refers to the discriminator's log-probability estimates that samples are from p_{data}. In the second part of Equation 3, $\mathbb{E}_{\boldsymbol{z} \sim p_Z(\boldsymbol{z})}$ refers to the expected samples $\boldsymbol{z}$ from a multivariate standard normal, $G(\boldsymbol{z})$ refers to a mapping from noise samples $\boldsymbol{z}$ to generator's samples, and $D(G(\boldsymbol{z}))$ refers to the probabilities that the generator's samples $G(\boldsymbol{z})$ are from p_{data}.

Graphically, we can interpret the value function $V(G, D)$ as a three-dimensional space with a loss surface that depends on the generator's weights $\boldsymbol{\theta}^{(G)}$ and the discriminator's weights $\boldsymbol{\theta}^{(D)}$. On this surface, we train G and D in an iterative manner with gradient descent-based optimization methods to arrive at a critical point.

To obtain a loss function for the discriminator, we can observe that D's objective is to maximize the log-likelihood of classifying a sample from p_{data} as true. Simultaneously, its objective is to maximize the log-likelihood of *not* classifying a sample from p_G as true. To learn the discriminator's distribution p_D, D maximizes the following loss function:

$$J^{(D)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)}) = \mathbb{E}_{\boldsymbol{x} \sim p_{\text{data}}}[\log(D(\boldsymbol{x}))] + \mathbb{E}_{\boldsymbol{z} \sim p_Z}[\log(1 - D(G(\boldsymbol{z}))]. \tag{4}$$

We let D (i.e., the gradients of Equation 4) satisfy differential privacy with DPSGD (Abadi et al. 2016). By the post-processing property of differential privacy (Dwork and Roth (2014), Proposition 2.1), we only have to let the discriminator satisfy a level of differential

privacy (also see Yoon, Jordon, and van der Schaar 2019). It follows that the discriminator's predictions are differentially private, which implies that under post-processing the GAN's outcomes are equally differentially private (i.e., differentially private data).

The second term from Equation 3 expresses G's objective to minimize the probability that the generator's samples evaluated by D are in fact from the generator. In other words, the objective of G is to minimize the maximum attainable of D. To learn the generator's distribution p_G, Goodfellow (2016) reformulates the objective of G from Equation 3 to minimize the following loss function:

$$J^{(G)}(\boldsymbol{\theta}^{(D)}, \boldsymbol{\theta}^{(G)}) = -\mathbb{E}_{\boldsymbol{z}\sim p_Z}[\log(D(G(\boldsymbol{z})))]. \tag{5}$$

In the limit, the minimax game results in p_G converging in probability to p_{data}. As a result, the generator maps random noise into high-quality samples that closely resemble samples from the dgp.

The "Where's Waldo?" effect

Abadi et al. (2016) prove that DPSGD provides $O(q\varepsilon_0\sqrt{T})$-differential privacy, where $q = 1/n$ is a customer's sampling probability from the training data set of size n, ε_0 is an analyst's desired level of privacy risk a priori to training an algorithm, and T denotes the number of training iterations.[3] Ultimately, $\frac{1}{n}\varepsilon_0\sqrt{T}$ is the function that governs the observed privacy risk (ε) after training an algorithm (i.e., $\varepsilon = \frac{1}{n}\varepsilon_0\sqrt{T}$). This implies that if the number of individuals in the data set (n) grows, the sampling probability ($1/n$) decreases and, as a result, the privacy risk (ε) decreases, ceteris paribus. In other words, to provide stronger privacy protection a manager needs to maximize data collection and increase the sample size to train our framework.

To isolate the effect of the sample size on privacy risk, we need to assume values for an

[3] We ignore δ for simplicity. This does not affect the managerial implications.

analyst's training iterations T and the desired level of privacy risk ε_0. We assume that an analyst needs 100,000 training iterations (T = 100,000 following Anand and Lee (2022)) to obtain reasonable utility and assume the following range of possible values an analyst may choose a priori to training: $\varepsilon_0 = (.01, 50)$. We use these values to obtain the privacy risk and average the privacy risk per samples size:

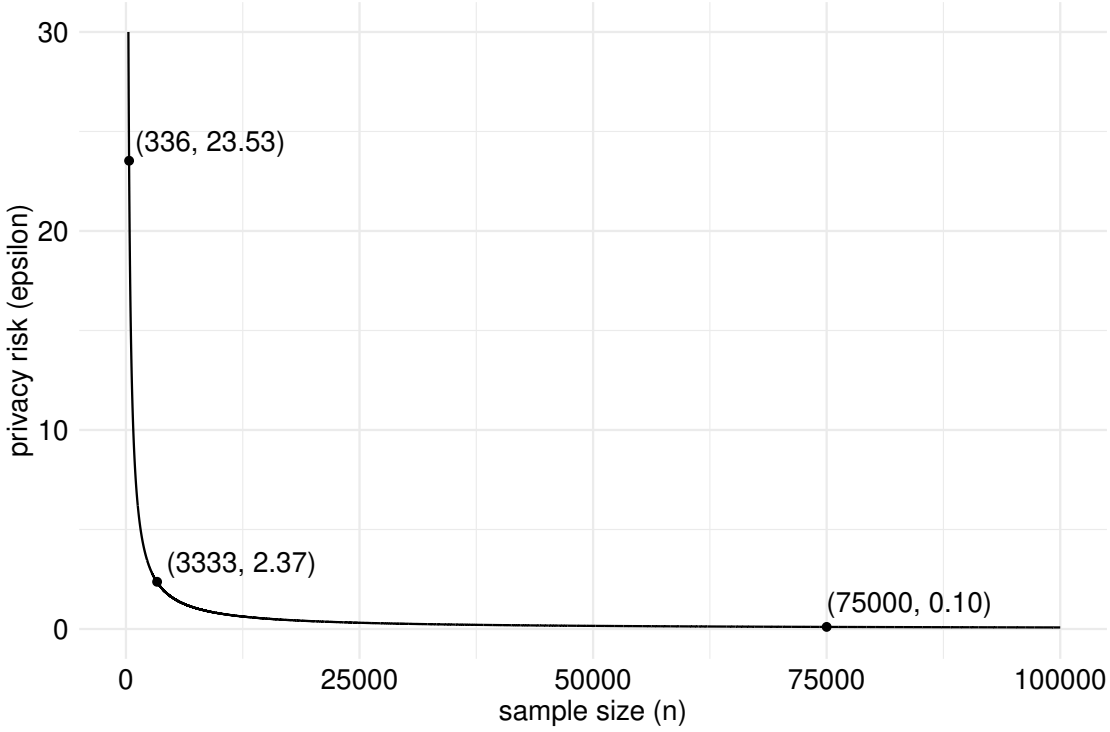

Figure 2: The relationship between privacy risk and sample size.

Visually, the "Where's Waldo" effect is reciprocal shaped; the customers' privacy risk is inversely proportional to the sample size (see Figure 2). We highlight three scenarios to illustrate the "Where's Waldo" effect. Imagine that an analyst has a total sample size of 100,000 customers, but the analyst starts its analysis with 336 customers. Consequently, the privacy risk of customers that are included increases with a factor of $e^{23.53} = 16,555,761,416$. If the analyst decides to increase the sample size to 3,333 customers, the customers' privacy risk decreases to a $e^{2.37} = 10$-factor increase. Finally, for a sample size of 75,000 customers, the privacy risk decreases substantially to a $e^{.1} = 1.1$-factor increase. We present empirical evidence for the "Where's Waldo" effect next.

MARKETING APPLICATIONS

We apply GANs with differential privacy to two marketing applications in which data protection is important. In each application, we limit the sample size to study its effect on the individuals' privacy risk. For our first application, we study the drivers of churn. We use a cross-sectional churn data set of 3,333 customers with 17 variables (see Table 2). The data set describes the arguably privacy-sensitive calling behavior of customers at a telecommunication service provider, such as minutes called during the day or night.

Variable	Min.	1st Q.	Mean	3rd Q.	Max.
Tenure (in months)	1	74	101	127	243
International plan (no/yes)	0	0	.09	0	1
Voicemail plan (no/yes)	0	0	.28	1	1
Day					
- Minutes	0	143.7	179.8	216.4	350.8
- Number of calls	0	87	100.4	114	165
- Charge (in dollars)	0	24.4	30.6	36.8	59.6
Evening					
- Minutes	0	166.6	201	235.3	363.7
- Number of calls	0	87	100.1	114	170
- Charge (in dollars)	0	14.2	17.1	20	30.9
Night					
- Minutes	23.2	167	200.9	235.3	395
- Number of calls	33	87	100.1	113	175
- Charge (in dollars)	1	7.52	9.0	10.6	17.8
International					
- Minutes	0	8.5	10.2	12.1	20
- Number of calls	0	3	4.5	6	20
- Charge (in dollars)	0	2.3	2.8	3.3	5.4
Customer service calls	0	1	1.6	2	9
Churn (no/yes)	0	0	.15	0	1

Table 2: Descriptive statistics of the churn data set.

Our second application is a pharmaceutical marketing application in which we study the effect of detailing on physician's drug prescription behavior. Recent disclosure laws aimed

to increase the transparency of detailing efforts make data sharing of particular importance (Guo, Sriram, and Manchanda 2020). We use a balanced prescription panel data set that consists of 46,593 prescriptions, written by 336 physicians over a period of 52 weeks (see Table 3). The panel nature of the data set poses a challenge to learn the relationships that characterize such data sets. We modify our framework to specifically learn these relationships and successfully sample differentially private panel data.

Variable	Min.	1st Q.	Mean	3rd Q.	Max.
Physician No.	4	213.2	482.4	800.5	999
Week	1	13.8	26.5	39.3	52
Physician's gender (male = 1)	0	1	.77	1	1
Practice size	1	4	5.37	7	8
Prescriptions (in units)	0	1	2.65	4	33
Patient's age (in years)	40.64	60.08	62.73	65.92	82
Detailing (in units)	0	0	.97	0	36

Table 3: Descriptive statistics of the pharmaceutical data set.

Application 1: Sharing cross-sectional churn data

For many firms, customer retention is a top priority (Ascarza 2018; Lemmens and Gupta 2020). Consequently, the marketing literature has given considerable attention to developing methodologies to explain and predict churn. To do so, managers and researchers are required to use privacy-sensitive data. Holtrop et al. (2017), for example, find that customer age, whether a customer has children, education level, income, social class, and prosperity level are predictors of churn. When dealing with such privacy-sensitive data, firms' data protection (or the lack thereof) can lead to privacy concerns and have a negative effect on firm performance and the industry as a whole (Martin, Borah, and Palmatier 2017).

To analyse the effect of privacy risk (ε) on utility, we train a deep convolutional GAN (DCGAN) with DPSGD to satisfy differential privacy (Radford, Metz, and Chintala 2015).[4]

[4]Compared with GANs, DCGANs use convolutional layers instead of fully connected layers (see Radford, Metz, and Chintala 2015). Before we apply differential privacy, we compare alternative architectures of GANs to determine which type of GAN generates the highest quality of samples. These results are available upon request.

We vary the privacy risk for the following values $\varepsilon = \{.01, .05, .1, .5, 1, 3, 5, 7, 13, \infty\}$. Differential privacy allows us to argue about the privacy risk of each individual customer as follows: for the lowest level of privacy risk $\varepsilon = .01$, any initial knowledge an adversary might have of a customer can increase *at most* or *in a worst-case scenario* with a multiplicative factor of $e^{.01} \approx 1.01$, or 1%. For a modest privacy risk $\varepsilon = .1$, a customer's privacy risk can increase at most with a factor of $e^{.1} \approx 1.10$ or 10%. For increasingly larger values of ε, the privacy risk increases exponentially. To illustrate, when ε is set to 13, an individual's privacy risk increases with a factor of $442,413$.

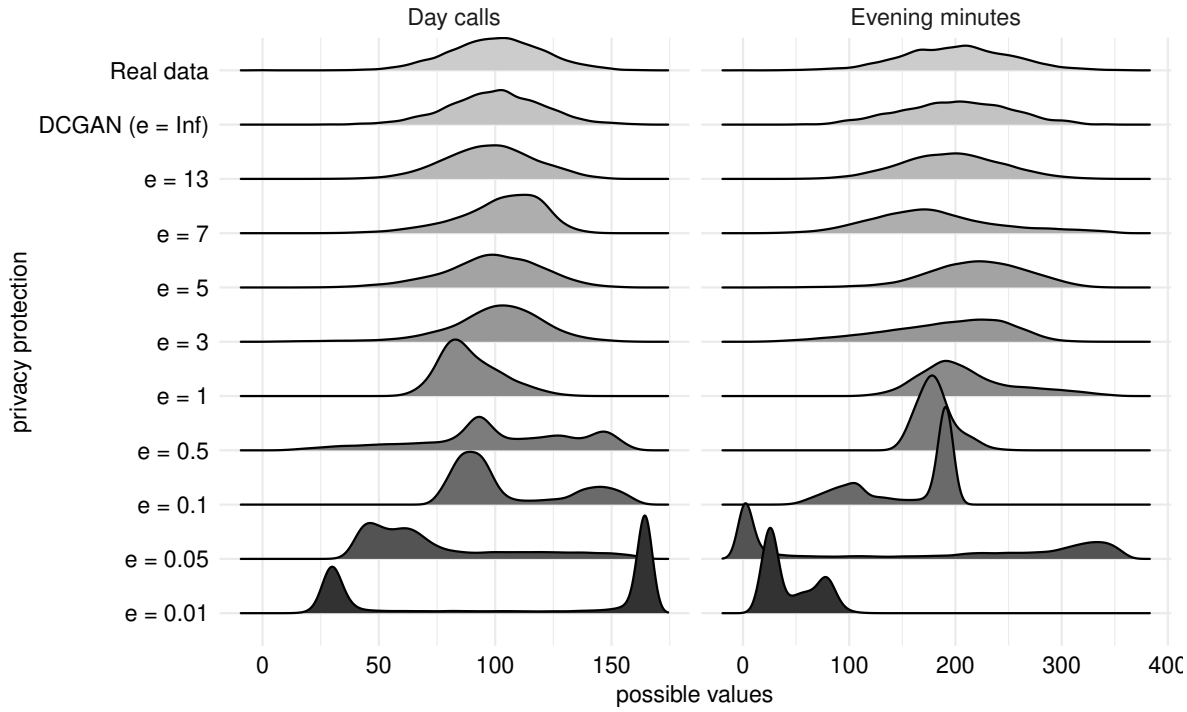

Figure 3: The marginal distributions for *Day calls* and *Evening minutes* with a varying levels of privacy protection (e = ε).

In Figure 3, we visualize the marginal distribution of the variables *Day calls* and *Evening minutes* with varying levels of privacy risk. Visually, the DCGAN's data without differential privacy (which equals $\varepsilon \approx \infty$) learns the real marginal distributions with considerable accuracy. Consequently, we keep the DCGAN's architecture and only vary the amount of privacy risk (ε). Interestingly, lower privacy risk (ε) and thus stronger protection implies

increasingly higher probability on values that in reality do not occur. Somewhat impressively, the differentially private probability densities start to resemble the real probability density around $\varepsilon = 1$. This implies that we can accurately learn, visualize, and share the probability densities at the cost of a factor of 2.72 increase in privacy risk. An alternative interpretation is that: by publicly releasing the probability density that satisfies a level of $\varepsilon = 1$, we increase any initial suspicion that an individual was included in the analysis with a factor of 2.72.

What are the drivers of churn behavior among customers? To prevent customers from churning, it is important to understand the reasons why they churn. To gain insight into these drivers, we use both data sets to estimate a logistic regression. Specifically, we specify the utility y_i^* for a customer to be a function of the following:

$$y_i^* = \mathbf{x}_i'\boldsymbol{\beta} + \nu_i := \beta_0 + \beta_1 \text{AccountLength}_i + \beta_2 \text{IntlPlan}_i + \beta_3 \text{VMailPlan}_i + \beta_4 \text{DayMins}_i + \beta_5 \text{DayCalls}_i + \beta_6 \text{EveMins}_i + \beta_7 \text{EveCalls}_i + \beta_8 \text{NightMins}_i + \beta_9 \text{NightCalls}_i + \beta_{10} \text{IntlMins}_i + \beta_{11} \text{IntlCalls}_i + \beta_{12} \text{CustServCalls}_i + \nu_i, \quad (6)$$

where β_0 denotes the intercept; AccountLength_i describes customer i's tenure; IntlPlan_i is a dummy indicating whether a customer i has an international plan; VMailPlan_i denotes whether a customer i has a voicemail plan; CustServCalls_i denotes customer i's number of calls to customer service; the other variables describe calling behavior during the Day, Eve, Night and International; and ν_i is an error term following an extreme value distribution.

We estimate Equation 6 using differentially private data and the real data. To measure utility, we assume that the entity, say a firm, shares a data set with another party. The objective of the receiving party is to derive real insights from the data that are shared. We define the loss of utility of the data in terms of the mean absolute percentage deviation

(MAPD) between protected and real parameter estimates:

$$\text{MAPD} = \frac{1}{J}\sum_{j=1}^{J}\left|\frac{\hat{\beta}_j - \tilde{\beta}_j}{\hat{\beta}_j}\right|,$$

where J is the total number of parameters, $\hat{\beta}_j$ is an estimated unprotected or real parameter, and $\tilde{\beta}_j$ is an estimated protected parameter. We assume that the betas from the estimations on the real data $\hat{\beta}_j$ have zero deviation from the true population parameters β_j. This allows us to visualize the trade-off between privacy risk (ε) and loss of utility (MAPD) in Figure 4.

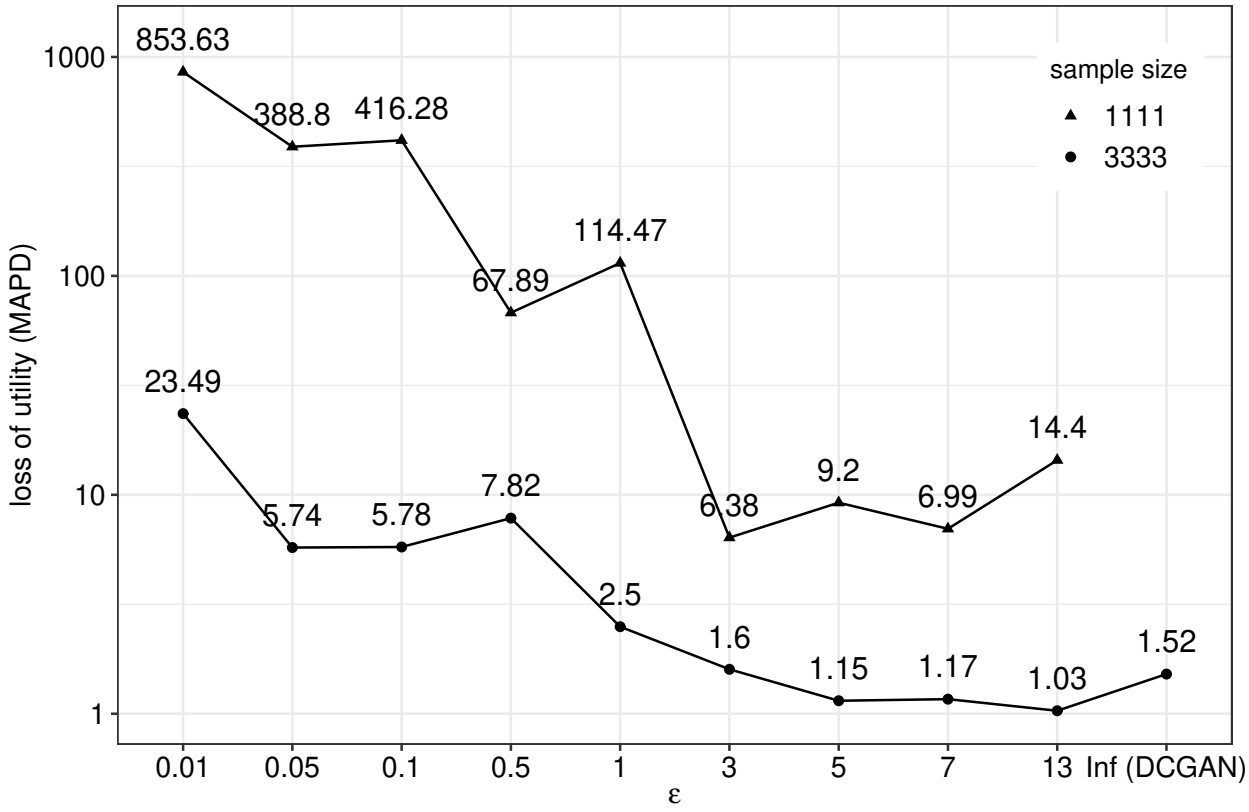

Figure 4: The trade-off between loss of utility (MAPD on a logarithmic scale) and privacy risk (ε) for a scenario in which a manager uses logistic regression to explain churn.

To capture the effect of sample size on privacy risk, we first use every observation that is available in the data set (n = 3,333). From Figure 4, we observe that for relatively strong privacy protection $\varepsilon \leq 1$, the ability to derive meaningful insights is limited. For $\varepsilon \geq 3$, the MAPD is almost equal to the situation in which we share data from the DCGAN without

differential privacy ($\varepsilon = \infty$). Consequently, a manager can share or derive meaningful insights (or the differentially private parameter estimates in Figure 5) at the cost of an increase of a factor of $e^3 \approx 20$ in customer's privacy risk. In other words, by being included in the churn analysis, a customer's privacy risk increases with a factor of 20 versus the scenario in which the customer would have been excluded. Second, when we decrease the sample size to 1,111 observations, a manager has to pay a higher price in terms of utility to maintain the same level of privacy risk. For the same level of privacy risk ($\varepsilon = 3$), the manager's ability to derive insights decreases substantially (MAPD increases from 1.6 to 6.38). Alternatively, a manager can decide to increase the customers' privacy risk to obtain the same level of utility.

Henceforth, we only consider the scenario in which a manager has access to every observation in the churn data set (n = 3,333). In Figure 5, we observe the parameter estimates that increase customers' privacy risk with a factor of $e^3 \approx 20$. The parameter estimates are almost all in the same direction as the real parameter estimates, which would lead to similar marketing insights. For example, a manager would find from all data sets that having an international plan or making more customer service calls increases a customer's probability of churning.

In conclusion, a manager has to pay an increase of $e^3 = 20$ in privacy risk to derive meaningful insights. To provide a benchmark, Abadi et al. (2016) derive meaningful utility at the cost of a privacy risk increase of $e^8 \approx 2{,}980$. Importantly, we provide empirical evidence for our "Where's Waldo" effect; to reduce the price in utility for stronger privacy protection (or lower privacy risk), we encourage analysts to maximize data collection.

Application 2: Sharing sensitive panel drug prescription data

In 2020, pharmaceutical firms invested approximately $2 billion to promote physicians' prescription behavior, also referred to as detailing (Guo, Sriram, and Manchanda 2020).[5]

[5]See https://openpaymentsdata.cms.gov/summary.

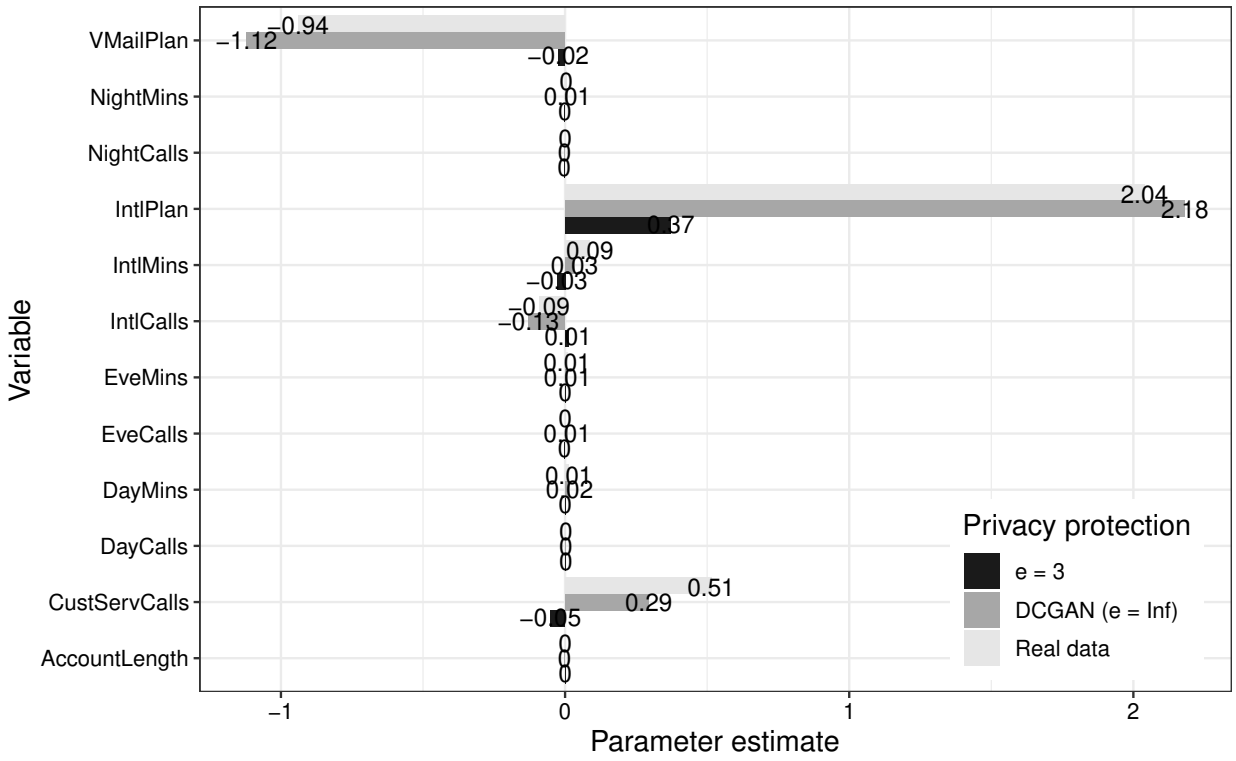

Figure 5: Parameter estimates of logistic regression with differential privacy ($\varepsilon = 3$), DCGAN without differential privacy and the real data.

Such detailing efforts have proved to be the most effective marketing instrument in the pharmaceutical sector, but they are also subject to public debate (Kremer et al. 2008). Public policy makers are concerned about the increasing use of detailing efforts and have enacted disclosure laws to provide full transparency of such efforts (Guo, Sriram, and Manchanda 2020). This might require sharing sensitive patient data to demonstrate the legitimacy of certain prescriptions to patients. As a result, data sharing of such privacy-sensitive patient records is of importance to policy makers, physicians, and pharmaceutical firms to derive empirical generalizations on the effect of detailing on prescription behavior or the effect of disclosure laws and future regulation.

DCGAN. To sample and share a panel data set, we need to modify our framework to learn the heterogeneity in behavior among physicians and the potential relationships over time.[6] These aspects make the GAN's architecture considerably more complex. To alleviate

[6]To satisfy differential privacy, we have experimented with a large variety of architectures that are available

these difficulties, we treat our panel data set such that a single physician with 7 variables over 52 weeks is a single observation. In other words, a single observation or physician is a matrix $\boldsymbol{X}_i \in \mathbb{R}^{t \times k}$ where i indicates a physician, t indicates the 52 weeks, and k denotes the 7 variables present in the data set. In total, we have access to a sample of 336 physicians. This implies that we satisfy differential privacy for the physician's entire prescription behavior (Dwork, Kohli, and Mulligan 2019). Ultimately, we use the DCGAN's generator (see Web Appendix 7) to sample 336 matrices and obtain differentially private data.

We find that due to the relatively small sample size ($n = 336$ versus $n = 3{,}333$), we are required to introduce more noise to obtain the same level of privacy protection (Abadi et al. 2016). Intuitively, differential privacy promises to reduce every single individual's contribution to the GAN's outcomes. Counterintuitively, a smaller sample size forces us to introduce more noise to hide this contribution. To understand this phenomena in greater detail, consider the situation in which we would have access to a population sample of physicians and obtain a significant effect of detailing on prescriptions (outcome). Now if we exclude a single physician from the population, this effect (outcome) would only change with a very small amount and would require little noise to hide. In our case, we only have 336 physicians, and thus, exclusion of one physician might lead to a relatively large change in the effect (outcome). Consequently, a smaller sample size requires significantly more noise to provide stronger levels of differential privacy. We decide to increase the privacy risk (ε), such that our main goal becomes to investigating which level of privacy protection provides the ability to derive utility.

We operationalize data utility in this case as the ability to learn the relationships between the variables for each physician over time. The kernel matrices from the convolutional layers convolve, in both the time and variable dimensions, over a batch of physician matrices to learn the relationships over time and heterogeneity in behavior among physicians (see Web Appendix 7). For early layers in the generator, we define a larger kernel matrix to learn

for GANs. Unfortunately, all of the alternatives introduce too many parameters, which increases the gradient's dimensionality and leads to excessive noise injection (Abadi et al. 2016). This hinders the ability to derive utility.

dependencies over a longer time period. For increasingly deep layers in the network, we shrink the kernel matrix size to learn dependencies over a shorter time period.

What is the effect of detailing on physicians' prescription behavior? In light of disclosure laws, policy makers might be interested in whether detailing has a statistically significant effect on the number of physicians' sensitive prescriptions. Therefore, we investigate whether we can derive a similar detailing effect from the differentially private panel data. Let Rx_{it} be the prescriptions of a physician i within a week t. We specify the probability of observing a number of Rx_{it} prescriptions as

$$\mathbb{P}\left(Rx_{it}\right) = \frac{\exp(-\tilde{\lambda}_{it})\tilde{\lambda}_{it}^{Rx_{it}}}{Rx_{it}!}, \tag{7}$$

where

$$\tilde{\lambda}_{it} = \lambda_{it}c_i = \exp(\beta_0 + \beta_1 \text{Detailing}_{it} + \delta \boldsymbol{X}_{it}); \tag{8}$$

Detailing_{it} contains the detailing efforts such as actual calls or visits to physicians; $\boldsymbol{X}_{it}$ is a matrix that contains physician-specific control variables such as gender, a patient's age and practice size; and c_i denotes the physician specific unobserved heterogeneity. We estimate a pooled Poisson specification, random- and fixed-effects Poisson specifications, and a pooled negative binomial specification.[7] In Figure 6, we only visualize the parameter estimates from the pooled Poisson specification because our results are robust to alternative specifications.

Now consider the situation in which we want to share data with policy makers who are interested in the effect of detailing on prescriptions. First, from the real data, we observe that for a one-unit increase in detailing, the expected prescriptions increase by a factor of 1.03. Second, if we use a DCGAN ($\varepsilon = \infty$), a policy maker would derive similar insights: a one-unit increase in detailing leads to a 1.02 factor increase in prescriptions. If we introduce

[7] The pooled estimator treats the panel as a cross-section and assumes the unobserved heterogeneity c_i to be 1 for all i. The random-effects Poisson assumes that c_i is Gamma distributed. The fixed-effects estimator is estimated with a sufficient statistic for c_i (Hausman, Hall, and Griliches 1984).

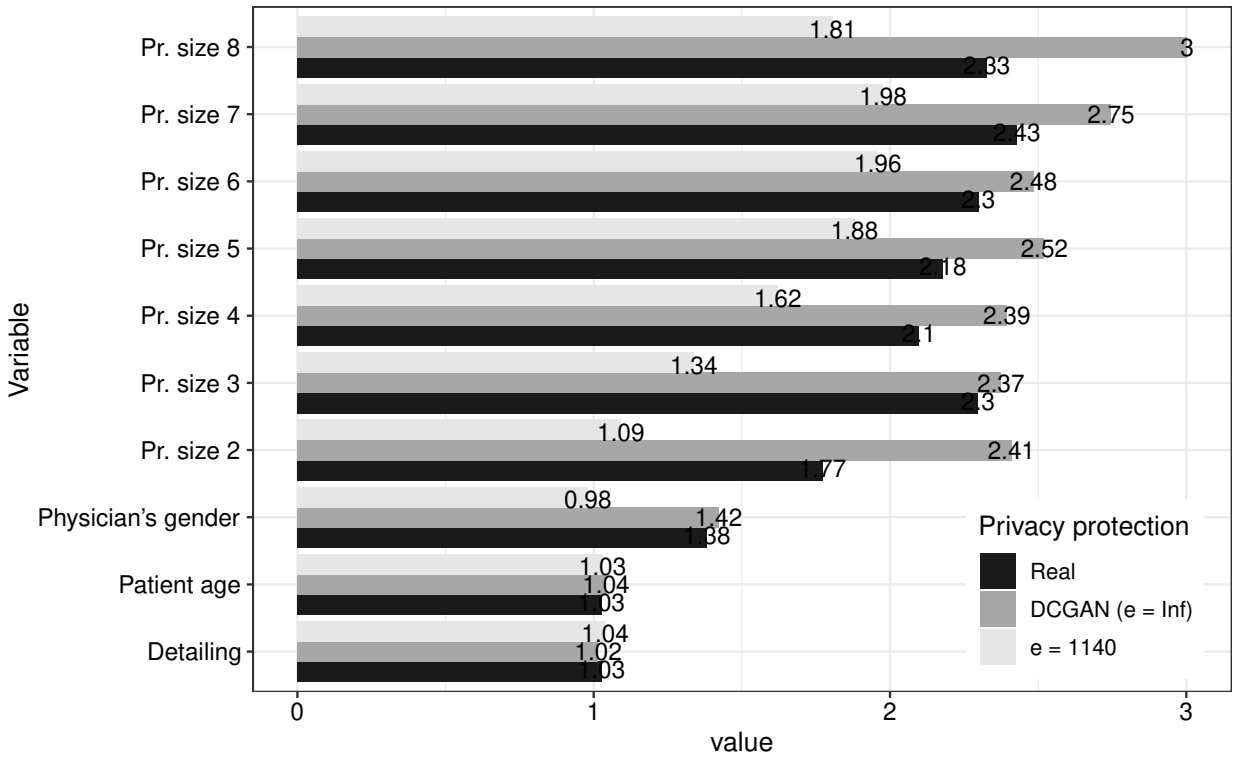

Figure 6: Parameter estimates of the pooled Poisson specification with differential privacy ($\varepsilon = 1{,}140$), $\varepsilon = \infty$ and the real data.

differential privacy into the DCGAN, detailing's observed effect becomes stronger. At the cost of an increase in privacy risk of $\varepsilon = 1{,}140$, we observe a factor increase in expected prescriptions of 1.04.

To visualize the privacy-utility trade-off, we operationalize loss of utility as the average MAPD over all estimators. First, we consider the situation in which an analyst has access to every physician's prescription behavior ($n = 336$). When sharing data with infinite privacy risk, we obtain MAPDs of .31, .32, .37, and .32 for the pooled, negative binomial, random-effects, and fixed-effects Poisson, respectively. Hence, we observe an average MAPD of .33. When we limit the privacy risk ($\varepsilon = 1{,}140$), we obtain an average MAPD of .57 (with MAPDs of .43, .40, .69, and .77 for the pooled, negative binomial, random effects, and fixed-effects Poisson, respectively). Next, we limit the sample size to 168 physicians and once more study the privacy-utility trade-off. In Figure 7, we observe that the price of privacy protection in utility consistently increases. For example for the same level of privacy risk ($\varepsilon = 1{,}140$), the

ability to derive meaningful insights is reduced substantially (average MAPD increases to 1.13). Alternatively, one is forced to increase the privacy risk to improve the utility.

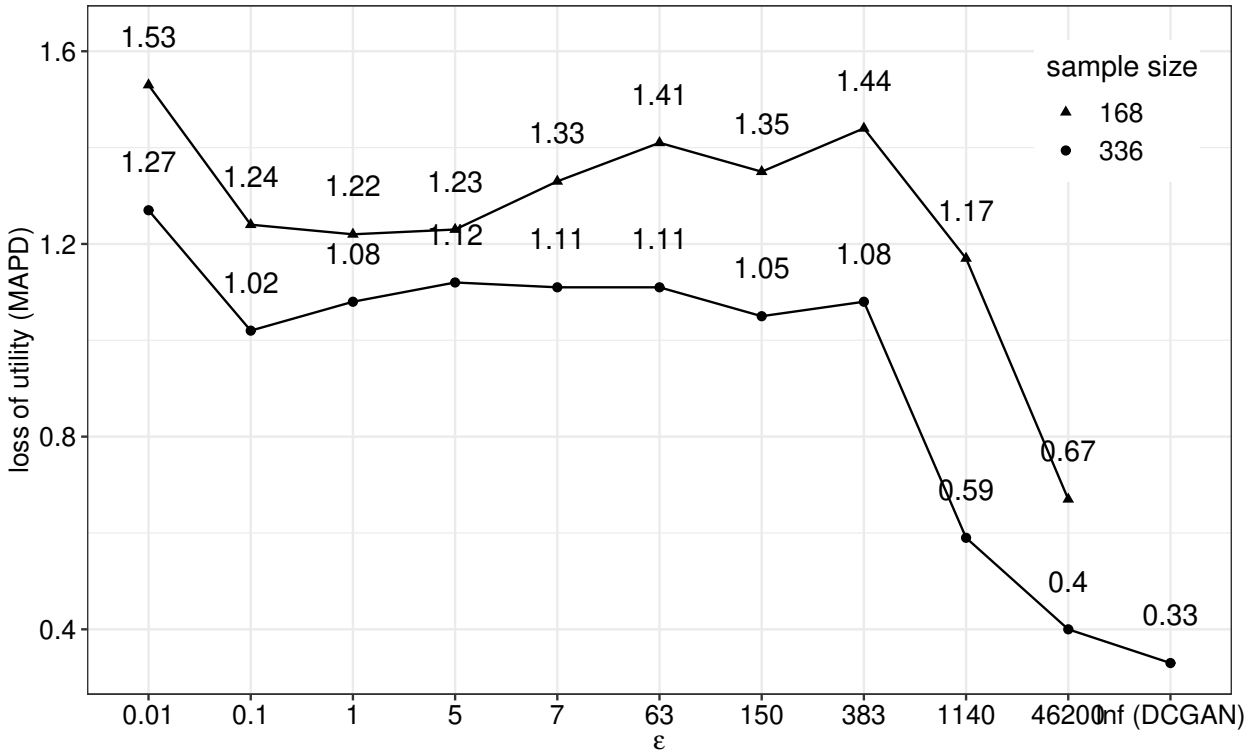

Figure 7: The trade-off between loss of utility (MAPD) and privacy risk (ε) when using a variety of estimations to explain physicians' prescriptions. To obtain a single MAPD, we average over the pooled Poisson, negative binomial, random-effects, and fixed-effects estimation MAPDs.

In conclusion, we show that the insight of a statistically positive effect of detailing on prescriptions is robust to assumptions we make during model specification. Although a physician's privacy risk increases substantially (ε = 1,140), a policy maker who does not have access to the real data can derive similar insights while limiting the privacy risk.

Does detailing have a persistent effect on a physician's prescriptions? To provide empirical evidence for the long-term effects of marketing actions, extant marketing literature has focused on whether marketing actions have a persistent or only temporary effect (Dekimpe and Hanssens 1995; Nijs et al. 2001). Often, such studies use unit-root tests to investigate whether time series are evolutionary or stationary followed by VAR models to derive

an impulse-response function to interpret a possible persistent long-term effect. Similarly, we are interested in whether the differentially private time series can provide meaningful marketing insights with respect to the long-term effect of detailing. We first aggregate the real and differentially private panel data sets on a weekly level by summing the number of prescriptions and detailing efforts in Figure 8. Visually, the DCGAN with $\varepsilon = \infty$ seems to learn the dynamics in the data well, whereas the introduction of differential privacy leads to excessive noise.

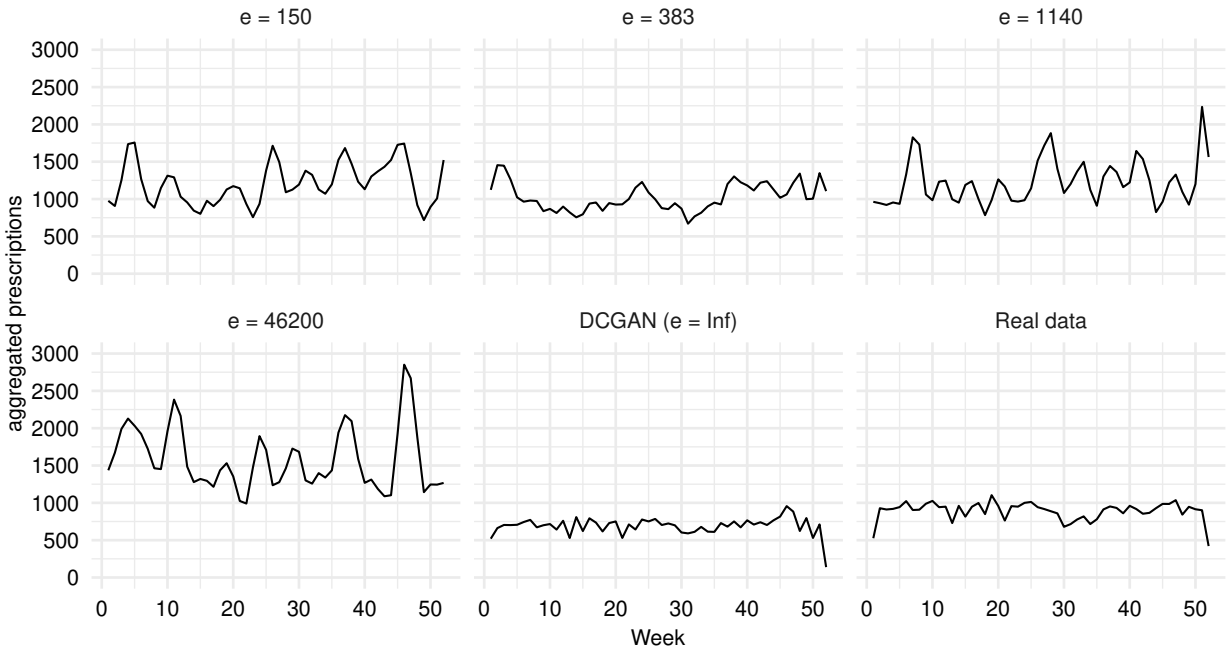

Figure 8: Time series of prescriptions for physicians aggregated on a weekly level (e = ε).

To study the possible long-term effect of detailing on prescription behavior of physicians, we specify a VAR model and derive an impulse-response function. For the real data, the first step is to take the first difference of prescriptions, which leads to stationary prescriptions over time. Subsequently, we specify a VAR(2) where we set the number of lags based on the Bayesian information criterion (BIC) from the real data:

$$\mathbf{y}_t = \boldsymbol{\beta}_0 + \boldsymbol{\beta}_1 \mathbf{y}_{t-1} + \boldsymbol{\beta}_2 \mathbf{y}_{t-2} + \mathbf{u}_t, \tag{9}$$

where $\mathbf{y}_t$, $\mathbf{y}_{t-1}$, and $\mathbf{y}_{t-2}$ are 2×1 vectors containing the first differences of prescriptions and detailing efforts at time t, $t-1$, and $t-2$; $\boldsymbol{\beta}_0$ is a 2×1 vector denoting the two constants; $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ are 2×2 matrices measuring the effect of prescriptions and detailing at $t-1$ and $t-2$ on prescriptions and detailing at time t; and $\mathbf{u}_t$ is a 2×1 disturbance term vector.

We simulate the impulse-response function over time in Figure 9 (Dekimpe and Hanssens 1995). Overall, the real impulse-response function shows that detailing has no lasting effect on physicians' prescriptions. Clearly, a privacy protection of $\varepsilon \leq 46{,}200$ introduces too much noise to resemble the real impulse-response function. A manager is forced to pay ∞ privacy risk to study the effect of detailing on prescriptions over time.

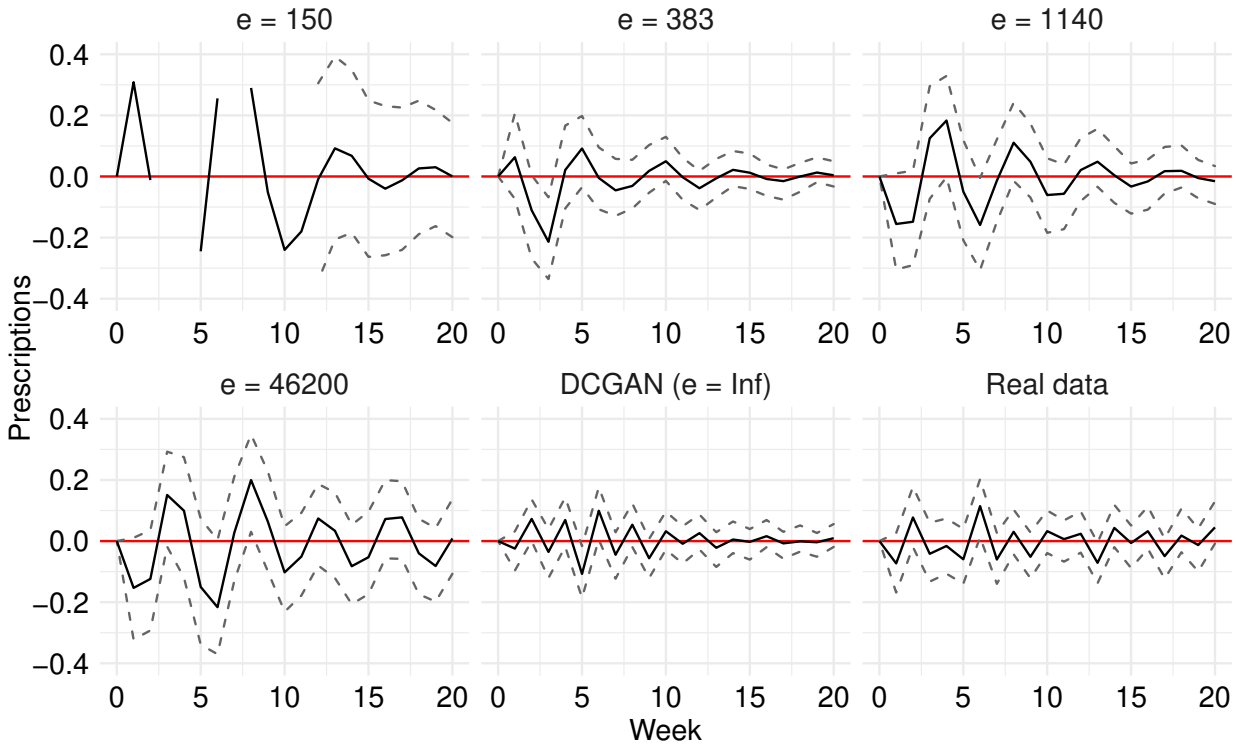

Figure 9: Impulse-response function of detailing on the number of prescriptions based on a VAR(2) model. We obtain confidence intervals by applying bootstrapping. We scale the y-axis according to the real data (e = ε).

Next, consider the situation in which a pharmaceutical firm shares data with a policy maker to study the effect of detailing on prescriptions over time. In terms of utility, a policy maker would at first glance derive that detailing does have a positive effect on prescriptions

of physicians (see Figure 6). When we study the effect of detailing over time, we observe that detailing has a insignificant effect on prescriptions (Figure 9). In terms of privacy, we consistently show that a larger sample size increases the ability to derive meaningful insights. When we compare the two applications, we also observe that the privacy risk increases due to the small sample size (ε = 1,140 vs. $\varepsilon = 3$). In addition, we find that for an increasing granularity in insights (e.g., pooled Poisson vs. impulse-response function), we observe that individuals have to pay a higher privacy risk price (ε = 1,140 vs. $\varepsilon = \infty$).

THE PRIVACY-UTILITY TRADE-OFF IN MARKETING

In theory, differential privacy provides a upper bound for the privacy risk an individual can incur. For example, in our churn application we restrict the customers' privacy risk to not further increase than a factor 20 ($\varepsilon = 3$). However, in a practical marketing context an adversary may never have the capabilities to fully exploit the privacy risk that is admissible. For example, DPSGD from Abadi et al. (2016) assume that an adversary has access to a data protection method's weights during training, which might never occur in marketing practice.

In this section, we return to the churn application and consider a privacy attack to investigate whether differential privacy's upper bound overstates the privacy risk possible in marketing practice. Initial evidence suggests that a large gap may exist between the theoretical bound and its empirical estimate when the capabilities of an adversary are reduced by its context (Nasr et al. 2021). For example, for a theoretical upper bound of $\varepsilon = 4$, Nasr et al. (2021) obtain an empirical estimate of $\hat{\varepsilon} = .31$. In our privacy attack, we compare our proposed method (GANs with differential privacy) with four alternative data protection methods: 5%, 20%, and 50% swapping; a Gaussian copula from Danaher and Smith (2011); GANs without differential privacy from Anand and Lee (2022); and, finally, the real data.[8]

[8] We do not compare our method with Bayesian shrinkage from Schneider et al. (2018) because their method only generates a single variable instead of a multivariate data set.

Benchmarks

Swapping. For each variable that is present in the data set, we randomly select 5%, 20%, and 50% of the observations and shuffle them randomly.

Gaussian copula. On the entire data set, we estimate a Gaussian copula model from Danaher and Smith (2011). Once we have fitted the Gaussian copula, we can sample data using MCMC sampling from the posterior distribution.

GANs without differential privacy. As a final benchmark, we use the DCGANs without differential privacy (which implies unbounded privacy risk) from Anand and Lee (2022). The European Commission and, specifically, European Data Protection Supervisor (2022) note that synthetic data is vulnerable to privacy attacks. To offer a fair benchmark, the GANs' architectures *with* differential privacy are an exact copy of the architectures *without* differential privacy. We only vary the discriminators' optimizer where we apply DPSGD to satisfy differential privacy (Abadi et al. 2016).

Likelihood-based privacy attack

To obtain an estimate of the empirical privacy risk, we employ a likelihood-based membership inference attack in case of the churn application that we studied previously in this paper. The objective of our attack is to infer whether an individual was included in a model's training set (also see Carlini et al. 2021; Chen et al. 2019; Jayaraman et al. 2020; Nasr et al. 2021; Yeom, Fredrikson, and Jha 2017). The membership attack set up is as follows:

1. A first marketer samples a training data set D from the dgp $\mathbb{D}$ and trains a data protection method $f(\hat{\theta})$ (e.g., a GAN).

2. A second marketer samples another data set $\tilde{D}$ from $\mathbb{D}$ and trains the same data protection method. This likely leads to different weights in the model: $f(\tilde{\theta})$.

3. An adversary requests protected samples from both $f(\hat{\theta})$ (i.e., training sample) and $f(\tilde{\theta})$ (i.e., adversary sample) and estimates the probability density functions of both

samples (i.e., p_{training} and $p_{\text{adversary}}$).

4. The adversary receives a data point x_i that is a sample from the training set D or from the dgp $\mathbb{D}$ (i.e., the observations in D and $\tilde{D}$ are distinct).

5. The adversary scores these data points on both p_{training} and $p_{\text{adversary}}$. If $p_{\text{training}}(x_i) > p_{\text{adversary}}(x_i)$, then the adversary decides that x_i is from the training set D.

By following this procedure, we can estimate the empirical epsilon as follows (Carlini et al. 2021):

$$\hat{\varepsilon} = \max\left(\log\frac{1-FPR}{FNR}, \log\frac{1-FNR}{FPR}\right) \tag{10}$$

where FPR measures the false positive rate (or probability of membership given non-membership of the training data) and FNR denotes the false negative rate (or probability of membership given non-membership of the training data). We simulate the procedure 100 times and present the estimates of the average privacy risk in Figure 10.[9] We measure loss of utility with the MAPDs from the churn application described in the "Marketing Applications" section.

From Figure 10, we conclude that the DCGANs with differential privacy (e.g., $\varepsilon = 1$ or 5) outperform existing data protection methods, including GANs without differential privacy. For the real data, we find an empirical epsilon of ∞. For 5%, 20%, and 50% swapping, we find empirical epsilons of 1.28, .51, and .13, respectively. This translates to a 3.59-, 1.67- and 1.14-factor increase in customers' privacy risk. For the Gaussian copula and GAN without differential privacy, we observe an empirical epsilon of .12 and .06 (or 1.13- and 1.06-factor increases in customers' privacy risk), respectively. In line with Nasr et al. (2021), we find a large gap between the theoretical upper bound of privacy risk (ε) and empirical privacy risk $\hat{\varepsilon}$. For example, when we set the theoretical upper bound to 5, we find an empirical epsilon of .02. Instead of a 14,741% increase in customers' privacy risk, a privacy attacker

[9] We simulate the GANs and Gaussian copula models 20 times.

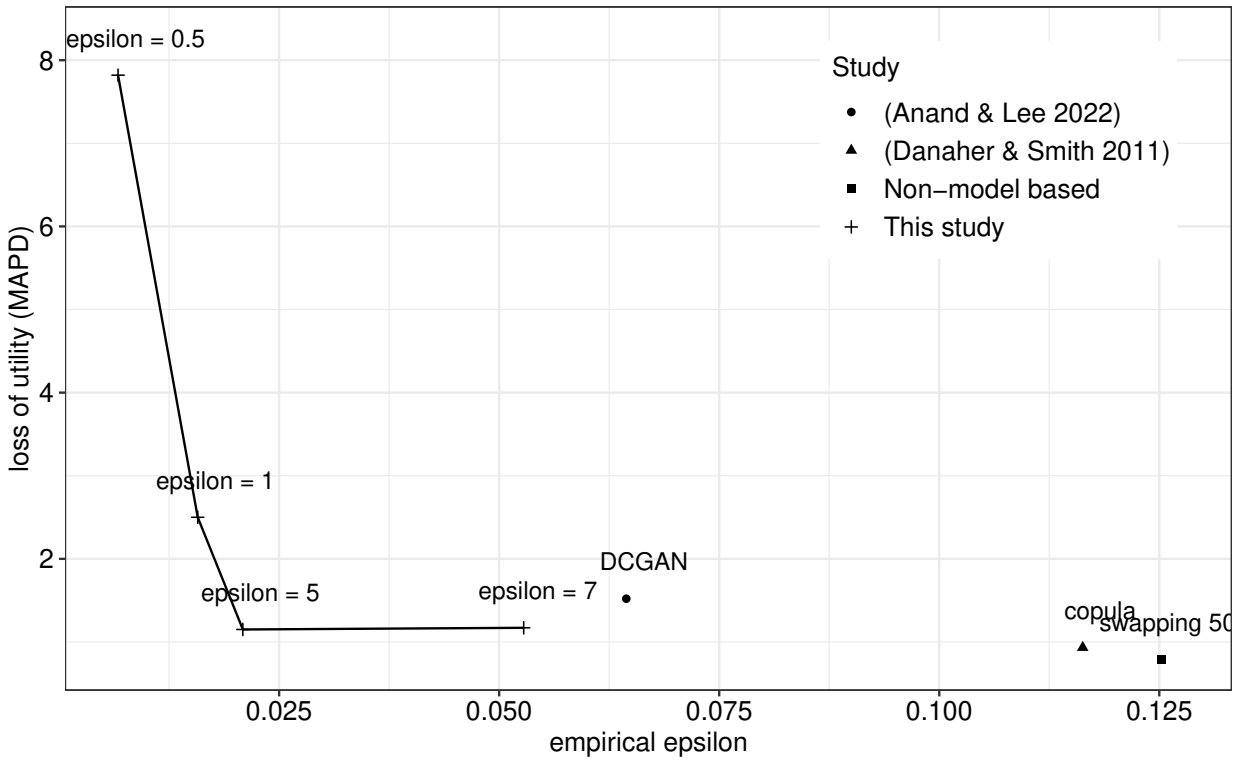

Figure 10: The trade-off between privacy (empirical epsilon) and loss of utility (MAPD) from the churn application. We limit the x-axis at .125.

in a marketing context is only capable of increasing the privacy risk 2%. For additional (technical) reasons why such a gap may exist, we refer the reader to Nasr et al. (2021).

DISCUSSION

The majority of the marketing literature focuses on measuring and improving customers' privacy perceptions that are difficult to quantify (e.g., Beke et al. 2022; Goldfarb and Tucker 2011; Martin, Borah, and Palmatier 2017). This paper proposes a framework that combines differential privacy and GANs to learn any marketing data set's complex joint distribution and protect consumers' privacy risk in a mathematically rigorous way. A key distinguishing feature from earlier work is that our method allows us to quantify and interpret the privacy risk that comes with deriving marketing insights (i.e., utility). We apply our method to two privacy-sensitive marketing applications and visualize the trade-off a manager, researcher,

or policy maker must make between privacy risk and utility. We find that they have to pay a high price in terms of privacy risk to derive meaningful insights. For example, if a customer is included in a churn analysis its privacy risk increases by a factor of 20. In a pharmaceutical marketing case, we find that for insights at an increasingly granular level, we must pay an even higher privacy risk price. To reduce this privacy risk, GDPR requires data minimization. In contrast to current beliefs, we find that the way to further reduce customers' privacy risk is to maximize data collection. Specifically, a larger sample size to train our framework reduces consumers' privacy risk. Intuitively, one can think of this as a "Where's Waldo" effect; customers can improve their privacy by hiding in a large crowd.

Although our findings imply a substantial privacy risk cost in exchange for utility, we end our paper with a comforting message. Due to the restrictions imposed by a marketing context, a privacy attacker does not have the capabilities to fully exploit the privacy risk that is admissible by differential privacy. We develop a likelihood-based membership inference attack which allows analysts to estimate differential privacy's privacy risk. We observe a large gap between the theoretical privacy risk (from differential privacy) and our customers' privacy risk estimates. In line with existing literature (e.g., Nasr et al. 2021), we find that customers' privacy risk only increases with a factor of 1.02 instead of a theoretical 148-factor increase in privacy risk. In terms of existing data protection methods that do not satisfy differential privacy, we do not find a high empirical privacy risk, but we stress that this risk is unbounded and thus may grow to infinity over time.

Our study has two major managerial implications. First, our findings suggest that firms need to be extremely careful when relying on existing data protection methods that do not satisfy differential privacy. When using such methods, firms run the risk of their customers being re-identified, and the potential privacy risk that customers run is theoretically unbounded (Dinur and Nissim 2003). Existing methods allow for the possibility that customers' privacy risk may grow to be extremely large. Therefore, firms that do not rely on differential privacy to limit their customers' privacy risk remain vulnerable to privacy

scandals (e.g., Netflix, Facebook and Cambridge Analytica). Our findings emphasize the importance for firms to adopt differential privacy and transparently communicate, control, and protect their customers' privacy (Martin, Borah, and Palmatier 2017). Second, we provide a framework to navigate the trade-off between privacy risk and utility. We demonstrate that to simultaneously improve customers' privacy protection and the ability to derive insights, firms should maximize data collection. This "Where's Waldo" effect is strongest for a small sample size and becomes weaker as the sample size becomes increasingly large.

Our findings also have implications for public policy makers. First, our study corroborates the European Data Protection Supervisor (2022)'s idea that existing data protection methods remain vulnerable to privacy attacks and that, currently, differential privacy is the only method that provides theoretical privacy protection guarantees. Second, our findings indicate a need to reflect on (or further specify) GDPR's directive to minimize data collection. Currently, GDPR considers large-scale data collection as a threat to one's privacy and that no data collection at all implies perfect privacy protection. In contrast, our findings show that customers' privacy protection is actually strengthened by data collection maximization. This new perspective on data collection could have positive consequences for not only privacy protection but also, among others, economic growth, and societal and scientific progress (European Commission 2021).

Finally, an implication for the scientific community is that researchers who use our framework can publish differentially private data to increase their studies' replicability. Alternatively, it promotes data sharing among researchers, which potentially improves scientific progress both within and outside the field of marketing.

Limitations and future research

One limitation of our study is that there a literature stream in computer science that mathematically derives ways to reduce the cost of privacy risk (from differential privacy) in exchange for utility (e.g., Abadi et al. 2016; Ghazi et al. 2021; Papernot et al. 2018). In this

study, we focus on putting the currently state-of-the-art DPSGD algorithm to marketing practice. However, as the computer science literature develops ways to satisfy stronger levels of differential privacy without hurting utility, we expect our privacy-utility trade-offs to shift closer to the origin. Simultaneously, the results from our privacy attack questions whether the marketing literature needs further improvements, as we already observe a very low privacy risk in (marketing) practice. We urge future research to further examine the empirical privacy risk by varying the assumptions on what an attacker may have access to (e.g., the generator's weights or auxiliary data).

Another limitation is that we do not apply GANs to unstructured data. Recently, the marketing literature has started to uncover the value of unstructured data to inform marketing actions. Given that GANs are able to approximate any dgp, they should also be able to sample unstructured data such as high-resolution images or text. Given the vast literature on image generation, we do not pursue such an application (Denton et al. 2015; Goodfellow et al. 2014; Karras et al. 2017; Salimans et al. 2016). Another fruitful area of research would be to sample artificial text using GANs. Text data are often represented by discrete sequences, which introduce difficulties with respect to the convergence of GANs (see Goodfellow 2016). To circumvent such difficulties, the AI literature (e.g., Kusner and Hernández-Lobato 2016) has proposed numerous solutions. We suggest the application of GANs with differential privacy to preserve individuals' privacy in text data as an area for future research.

REFERENCES

Abadi, Martin, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang (2016), "Deep Learning with Differential Privacy," *in Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security - CCS'16* http://dx.doi.org/10.1145/2976749.2978318.

Acquisti, Alessandro, Laura Brandimarte, and George Loewenstein (2015), "Privacy and human behavior in the age of information," *Science*, 347 (6221), 509–514 https://science.sciencemag.org/content/347/6221/509.

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman (2016), "The economics of privacy," *Journal of Economic Literature*, 54 (2), 442–492 https://ideas.repec.org/a/aea/jeclit/v54y2016i2p442-92.html.

Anand, Piyush and Clarence Lee (2022), "Using Deep Learning to Overcome Privacy and Scalability Issues in Customer Data Transfer," *Marketing Science*, 0 (0), null https://doi.org/10.1287/mksc.2022.1365.

Ascarza, Eva (2018), "Retention Futility: Targeting High-Risk Customers Might Be Ineffective," *Journal of Marketing Research*, 55 (1), 80–98 https://doi.org/10.1509/jmr.16.0163.

Beke, Frank T., Felix Eggers, Peter C. Verhoef, and Jaap E. Wieringa (2022), "Consumers' privacy calculus: The PRICAL index development and validation," *International Journal of Research in Marketing*, 39 (1), 20–41 Funding Information: The authors acknowledge the Customer Insight Center from the University of Groningen for funding. They are grateful for receiving valuable comments from participants of EMAC in Oslo, the Marketing Science Conference in Shanghai, seminar participants at IESEG Business School Lille, EDHEC Business School Lille, and University of Stavanger. Funding Information: The authors acknowledge the Customer Insight Center from the University of Groningen for funding. They are grateful for receiving valuable comments from participants of EMAC in Oslo, the Marketing Science Conference in Shanghai, seminar participants at IESEG Business School Lille, EDHEC Business School Lille, and University of Stavanger. Publisher Copyright: © 2021 Elsevier B.V.

Carlini, Nicholas, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramèr (2021), "Membership Inference Attacks From First Principles," *CoRR*, abs/2112.03570 https://arxiv.org/abs/2112.03570.

Carlini, Nicholas, Chang Liu, Jernej Kos, Úlfar Erlingsson, and Dawn Song (2018), "The Secret Sharer: Measuring Unintended Neural Network Memorization & Extracting Secrets," *CoRR*, abs/1802.08232 http://arxiv.org/abs/1802.08232.

Chen, Dingfan, Ning Yu, Yang Zhang, and Mario Fritz (2019), "GAN-Leaks: A Taxonomy of Membership Inference Attacks against GANs," *CoRR*, abs/1909.03935 http://arxiv.org/abs/1909.03935.

Danaher, Peter J. and Michael S. Smith (2011), "Modeling Multivariate Distributions Using Copulas: Applications in Marketing," *Marketing Science*, 30 (1), 4–21 https://doi.org/10.1287/mksc.1090.0491.

Dekimpe, Marnik G. and Dominique M. Hanssens (1995), "Empirical Generalizations about Market Evolution and Stationarity," *Marketing Science*, 14 (3), G109–G121 http://www.jstor.org/stable/184153.

Denton, Emily L., Soumith Chintala, Arthur Szlam, and Robert Fergus (2015), "Deep Generative Image Models using a Laplacian Pyramid of Adversarial Networks," *CoRR* http://arxiv.org/abs/1506.05751.

Dinur, Irit and Kobbi Nissim (2003), "Revealing Information While Preserving Privacy," *in Proceedings of the Twenty-Second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, page 202–210 https://doi.org/10.1145/773153.773173.

Dwork, Cynthia, Nitin Kohli, and Deirdre Mulligan (2019), "Differential Privacy in Practice: Expose your Epsilons!," *Journal of Privacy and Confidentiality*, 9 (2) https://journalprivacyconfidentiality.org/index.php/jpc/article/view/689.

Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith (2006), "Calibrating Noise to Sensitivity in Private Data Analysis," *in Proceedings of the Third Conference on Theory of Cryptography*, page 265–284 https://doi.org/10.1007/11681878_14.

Dwork, Cynthia and Aaron Roth (2014), "The Algorithmic Foundations of Differential Privacy," *Foundations and Trends in Theoretical Computer Science*, 9 (3–4), 211–407 http://dx.doi.org/10.1561/0400000042.

European Commission (2012), *Regulation of the European parliament and of the council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)* European Commission, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52012PC0011.

European Commission (2021), "European data strategy," *European data strategy: Making the EU a role model for a society empowered by data.* https://ec.europa.eu/info/strategy/priorities-2019-2024/europe-fit-digital-age/european-data-strategy_en.

European Data Protection Supervisor (2022), "Synthetic Data," *European Data Protection Supervisor* https://edps.europa.eu/press-publications/publications/techsonar/synthetic-data_en.

Fredrikson, Matt, Somesh Jha, and Thomas Ristenpart (2015), "Model Inversion Attacks That Exploit Confidence Information and Basic Countermeasures," *in in Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, page 1322–1333 https://doi.org/10.1145/2810103.2813677.

Garfinkel, Simson, John M. Abowd, and Christian Martindale (2018), "Understanding Database Reconstruction Attacks on Public Data: These Attacks on Statistical Databases Are No Longer a Theoretical Danger.," *Queue*, 16 (5), 28–53 https://doi.org/10.1145/3291276.3295691.

Ghazi, Badih, Noah Golowich, Ravi Kumar, Pasin Manurangsi, and Chiyuan Zhang (2021), "On Deep Learning with Label Differential Privacy," *CoRR*, abs/2102.06062 https://arxiv.org/abs/2102.06062.

Goldfarb, Avi and Catherine E. Tucker (2011), "Privacy Regulation and Online Advertising," *Management Science*, 57 (1), 57–71 https://doi.org/10.1287/mnsc.1100.1246.

Goodfellow, Ian "NIPS 2016 Tutorial: Generative Adversarial Networks," (2016) https://arxiv.org/abs/1701.00160.

Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio "Generative Adversarial Nets," Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, "Advances in Neural Information Processing Systems 27," pages 2672–2680, Curran Associates, Inc. (2014) http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf.

Grewal, Rajdeep, Sachin Gupta, and Rebecca Hamilton (2021), "Marketing Insights from Multimedia Data: Text, Image, Audio, and Video," *Journal of Marketing Research*, 58 (6), 1025–1033 https://doi.org/10.1177/00222437211054601.

Guo, Tong, Srinivasaraghavan Sriram, and Puneet Manchanda (2020), ""Let the Sunshine In": The Impact of Industry Payment Disclosure on Physician Prescription Behavior," *Marketing Science*, 39 (3), 516–539 https://doi.org/10.1287/mksc.2019.1181.

Gupta, Sachin, Panos Moutafis, and Matthew Schneider (2022), "The marketer at the privacy table," *American Marketing Association* https://www.ama.org/2022/03/17/the-marketer-at-the-privacy-table/.

Hausman, Jerry, Bronwyn H. Hall, and Zvi Griliches (1984), "Econometric Models for Count Data with an Application to the Patents-R&D Relationship," *Econometrica*, 52 (4), 909–938 http://www.jstor.org/stable/1911191.

Holtrop, Niels, Jaap E. Wieringa, Maarten J. Gijsenberg, and Peter C. Verhoef (2017), "No future without the past? Predicting churn in the face of customer privacy," *International Journal of Research in Marketing*, 34 (1), 154 – 172 http://www.sciencedirect.com/science/article/pii/S0167811616300805.

Jayaraman, Bargav, Lingxiao Wang, David Evans, and Quanquan Gu (2020), "Revisiting Membership Inference Under Realistic Assumptions," *CoRR*, abs/2005.10881 https://arxiv.org/abs/2005.10881.

Karras, Tero, Timo Aila, Samuli Laine, and Jaakko Lehtinen "Progressive Growing of GANs for Improved Quality, Stability, and Variation," (2017) https://arxiv.org/abs/1710.10196.

Kremer, Sara T.M., Tammo H.A. Bijmolt, Peter S.H. Leeflang, and Jaap E. Wieringa (2008), "Generalizations on the effectiveness of pharmaceutical promotional expenditures," *International Journal of Research in Marketing*, 25 (4), 234–246 https://www.sciencedirect.com/science/article/pii/S0167811608000566.

Kusner, Matt J. and José Miguel Hernández-Lobato "GANs for Sequences of Discrete Elements with the Gumbel-softmax Distribution," (2016) https://arxiv.org/abs/1611.04051.

Leeflang, Peter, Tammo Bijmolt, Koen Pauwels, and Jaap Wieringa (2015), *Modeling Markets: Analyzing Marketing Phenomena and Improving Marketing Decision Making* International Series in Quantitative Marketing, Berlin: Springer.

Lemmens, Aurélie and Sunil Gupta (2020), "Managing Churn to Maximize Profits," *Marketing Science*, 39 (5), 956–973 https://doi.org/10.1287/mksc.2020.1229.

Lin, Zinan, Vyas Sekar, and Giulia Fanti (2021), "On the Privacy Properties of GAN-generated Samples," *in Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, 130, 1522–1530 https://inProceedings.mlr.press/v130/lin21b.html.

Liu, Kang, Benjamin Tan, and Siddharth Garg (2020), "Subverting Privacy-Preserving GANs: Hiding Secrets in Sanitized Images," *CoRR*, abs/2009.09283 https://arxiv.org/abs/2009.09283.

Martin, Kelly D., Abhishek Borah, and Robert W. Palmatier (2017), "Data Privacy: Effects on Customer and Firm Performance," *Journal of Marketing*, 81 (1), 36–58 https://doi.org/10.1509/jm.15.0497.

McSherry, Frank and Ilya Mironov (2009), "Differentially Private Recommender Systems: Building Privacy into the Netflix Prize Contenders," *in Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 627–636 https://doi.org/10.1145/1557019.1557090.

Narayanan, Arvind and Vitaly Shmatikov (2008), "Robust De-Anonymization of Large Sparse Datasets," *in in Proceedings of the 2008 IEEE Symposium on Security and Privacy*, page 111–125 https://doi.org/10.1109/SP.2008.33.

Nasr, Milad, Shuang Song, Abhradeep Thakurta, Nicolas Papernot, and Nicholas Carlini (2021), "Adversary Instantiation: Lower Bounds for Differentially Private Machine Learning," *CoRR*, abs/2101.04535 https://arxiv.org/abs/2101.04535.

Nijs, Vincent R., Marnik G. Dekimpe, Jan-Benedict E.M. Steenkamp, and Dominique M. Hanssens (2001), "The Category-Demand Effects of Price Promotions," *Marketing Science*, 20 (1), 1–22 https://doi.org/10.1287/mksc.20.1.1.10197.

Papernot, Nicolas, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Úlfar Erlingsson "Scalable Private Learning with PATE," (2018) https://arxiv.org/abs/1802.08908.

Radford, Alec, Luke Metz, and Soumith Chintala "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks," (2015) https://arxiv.org/abs/1511.06434.

Salimans, Tim, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen "Improved Techniques for Training GANs," (2016) https://arxiv.org/abs/1606.03498.

Schneider, Matthew J., Sharan Jagpal, Sachin Gupta, Shaobo Li, and Yan Yu (2017), "Protecting customer privacy when marketing with second-party data," *International Journal of Research in Marketing*, 34 (3), 593–603 https://ideas.repec.org/a/eee/ijrema/v34y2017i3p593-603.html.

Schneider, Matthew J., Sharan Jagpal, Sachin Gupta, Shaobo Li, and Yan Yu (2018), "A Flexible Method for Protecting Marketing Data: An Application to Point-of-Sale Data," *Marketing Science*, 37 (1), 153–171 https://doi.org/10.1287/mksc.2017.1064.

Sklar, M (1959), "Fonctions de repartition an dimensions et leurs marges," *Publications de l'Institut de statistique de l'Université de Paris*, 8, 229–231.

Wedel, Michel and P.K. Kannan (2016), "Marketing Analytics for Data-Rich Environments," *Journal of Marketing*, 80 (6), 97–121 https://doi.org/10.1509/jm.15.0413.

Wieringa, Jaap, P.K. Kannan, Xiao Ma, Thomas Reutterer, Hans Risselada, and Bernd Skiera (2021), "Data analytics in a privacy-concerned world," *Journal of Business Research* http://www.sciencedirect.com/science/article/pii/S0148296319303078.

Xie, Liyang, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou (2018), "Differentially Private Generative Adversarial Network," *CoRR*, abs/1802.06739 http://arxiv.org/abs/1802.06739.

Yeom, Samuel, Matt Fredrikson, and Somesh Jha (2017), "The Unintended Consequences of Overfitting: Training Data Inference Attacks," *CoRR*, abs/1709.01604 http://arxiv.org/abs/1709.01604.

Yoon, Jinsung, James Jordon, and Mihaela van der Schaar (2019), "PATE-GAN: Generating Synthetic Data with Differential Privacy Guarantees," *International Conference on Learning Representations* https://openreview.net/forum?id=S1zk9iRqF7.

Zhou, Yinghui, Shasha Lu, and Min Ding (2020), "Contour-as-Face Framework: A Method to Preserve Privacy and Perception," *Journal of Marketing Research*, 57 (4), 617–639 https://doi.org/10.1177/0022243720920256.

APPENDIX: THE GAN'S GENERATOR AND CONVOLUTION OPERATION.

The generator's architecture is as follows:

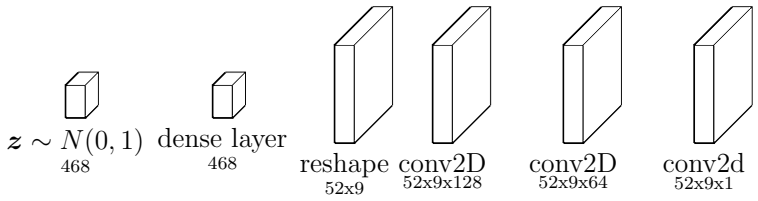

Figure 11: Topology of the GAN's generator. Sizes at the layers refer to output shape of the layer. For the convolutional layers (conv2D), the last dimension refers to the number of kernel matrices (see Radford, Metz, and Chintala 2015). Each kernel matrix learns its own weights.

To learn the dynamics in a physician's matrix $\boldsymbol{X}_i$, we use convolutional layers with a kernel $\boldsymbol{K}$ that work as follows:

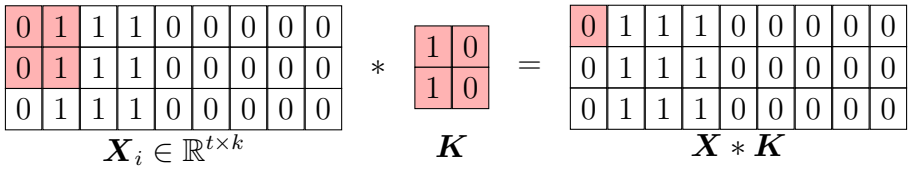

Figure 12: Example of a 2D convolution operation. To keep the same size of the result $\boldsymbol{X} * \boldsymbol{K}$, we can use zero padding, which adds zeros to the vector $\boldsymbol{X}$ to ensure that the outcome remains of the same dimensionality.