Data-Driven and Game-Theoretic Approaches for Privacy

by

Chong Huang

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosphy

Approved August 2018 by the
Graduate Supervisory Committee:

Lalitha Sankar, Chair
Oliver Kosut
Angelia Nedich
Lei Ying

ARIZONA STATE UNIVERSITY

August 2018

ABSTRACT

In the past few decades, there has been a remarkable shift in the boundary between public and private information. The application of information technology and electronic communications allow service providers (businesses) to collect a large amount of data. However, this "data collection" process can put the privacy of users at risk and also lead to user reluctance in accepting services or sharing data. This dissertation first investigates privacy sensitive consumer-retailers/service providers interactions under different scenarios, and then focuses on a unified framework for various information-theoretic privacy and privacy mechanisms that can be learned directly from data.

Existing approaches such as differential privacy or information-theoretic privacy try to quantify privacy risk but do not capture the subjective experience and heterogeneous expression of privacy-sensitivity. The first part of this dissertation introduces models to study consumer-retailer interaction problems and to better understand how retailers/service providers can balance their revenue objectives while being sensitive to user privacy concerns. This dissertation considers the following three scenarios: (i) the consumer-retailer interaction via personalized advertisements; (ii) incentive mechanisms that electrical utility providers need to offer for privacy sensitive consumers with alternative energy sources; (iii) the market viability of offering privacy guaranteed free online services. We use game-theoretic models to capture the behaviors of both consumers and retailers, and provide insights for retailers to maximize their profits when interacting with privacy sensitive consumers.

Preserving the utility of published datasets while simultaneously providing provable privacy guarantees is a well-known challenge. In the second part, a novel context-aware privacy framework called generative adversarial privacy (GAP) is introduced. Inspired by recent advancements in generative adversarial networks, GAP allows the

data holder to learn the privatization mechanism directly from the data. Under GAP, finding the optimal privacy mechanism is formulated as a constrained minimax game between a privatizer and an adversary. For appropriately chosen adversarial loss functions, GAP provides privacy guarantees against strong information-theoretic adversaries. Both synthetic and real-world datasets are used to show that GAP can greatly reduce the adversary's capability of inferring private information at a small cost of distorting the data.

ACKNOWLEDGMENTS

Above all, I would like to express my sincere thanks and gratitude to my advisor Dr. Lalitha Sankar for her guidance, inspiration and support throughout the course of this research. She was always willing to discuss a problem and always encouraging and helpful in suggesting new directions.

I would like to acknowledge Dr. Anand Sarwate and Dr. Peter Kairouz for all the inspiring discussions and guidance. I am also grateful to my committee members Dr. Oliver Kosut, Dr. Angelia Nedich, and Dr. Lei Ying for their insightful and valuable suggestions and comments.

In addition, I would also like to express my gratitude to the National Science Foundation (NSF) and Power System Engineering Research Center (PSERC) for the financial supports in this research.

Finally, I must give special thanks to my family and friends for their encouragements and support.

TABLE OF CONTENTS

CHAPTER Page

CHAPTER Page

CHAPTER Page

LIST OF TABLES

LIST OF FIGURES

Figure Page

Figure Page

Figure Page

Chapter 1

INTRODUCTION

1.1 How to Incentivize and Interact with Privacy Sensitive Consumer?

Programs such as retailer "loyalty cards" allow companies to automatically track a customer's financial transactions, purchasing behavior, and preferences. They can then use this information to offer customized incentives, such as discounts on related goods. Consumers may benefit from retailer's knowledge by using more of these targeted discounts or coupons while shopping. However, in some cases the coupon offer implies that the retailer has learned something sensitive or private about the consumer. Nevertheless, consumers also want companies to be transparent about what information they collect and how it will be used. In some cases the coupon offer implies that the retailer has learned something sensitive or private about the consumer. Due to predictions from machine learning algorithms, retailers seem to know more about the consumer than they expect or feel they have disclosed. For example, a retailer could infer a consumer's pregnancy [1]. Such violations may make consumers skittish about purchasing from that retailer.

Consumers are more willing to share broad demographic data and information about their usage of media content. Those types of data are generally considered to be less personal and can be anonymous. However, data or information that implies their private interest is considered to be highly unwilling to share by consumers, such as web browsing history, information about their social lives and financial information. The increase of consumers' privacy concerns forces retailers to be more careful about designing incentive schemes for the consumers since perceived privacy violations may "creep out" consumers.

However, modeling the privacy-sensitivity of a consumer is not always straightforward: widely-studied models for quantifying privacy risk using differential privacy

or information theory do not capture the subjective experience and heterogeneous *expression* of consumer privacy. This section introduces a framework to model the consumer-retailer interaction problem and better understand how retailers can develop coupon-offering policies that balances their revenue objectives while being sensitive to consumer privacy concerns. The main challenge for the retailer is that the consumer's responses to coupons are not known *a priori*; Furthermore, consumers do not "add noise" to their purchasing behavior as a mechanism to stay private. Rather, the offer of a coupon may provoke a reaction from the consumer, ranging from "unaffected" to "ambiguous" or "partially concerned" to "creeped out." This reaction is mediated by the consumer's sensitivity level to privacy violations, and it is these levels that we seek to model via a Markov decision process. These privacy-sensitivity states of the consumers are often revealed to the retailer through their purchasing patterns. In the simplest case, they may accept or reject a targeted coupon.

1.1.1 Motivation

According to a report by The New York Times [2], an infuriated father went into a Target store in Minneapolis, demanding to talk to a manager. He claimed his daughter got coupons on baby product in the mail. "Shes still in high school, and you are sending her coupons for baby clothes and cribs? Are you trying to encourage her to get pregnant?", he said. The manager has no clue what the man was talking about. The coupon book was sent to the man's daughter's address and contained promotions on nursery furniture, maternity clothing and pictures of smiling infants. The manager apologized and then called a few days later to apologize again. To his surprise, the father was somewhat abashed over the phone. "I had a talk with my daughter," he said. "It turns out there has been some activities in my house I haven't been completely aware of. She's due in August. I owe you an apology." It turns out that Target has discovered the girl's pregnancy way before her father did.

The technology used behind the story is that Target creates consumer accounts that store a history of everything they have bought and any demographic information Target has collected from them or bought from other sources. Using that, by looking at historical purchasing data, analysts from Target identify about 25 products that, when analyzed together, allowed them to assign each shopper a "pregnancy prediction" score. Sending out coupons based on this score system does not violate the privacy law. However, this marketing strategy clearly makes consumers feel uncomfortable since they sense their private information have been leaked.

With the increasing privacy concerns from consumers, the retailers should be more careful when offering promotions or advertisements to consumers. Unwinding all of these complicated interactions between data mining and privacy loss is difficult, so we focus instead on the phenomenon of how retailers react to privacy-sensitive consumers in a simple two-party interaction.

1.1.2 Contributions

We propose a partially-observed Markov decision process (POMDP) model for this problem in which the consumer's state encodes their privacy sensitivity, and the retailer can offer different levels of privacy-violating coupons. The simplest instance of our model is one with two states for the consumer, denoted as "Normal" and "Alerted," and two types of coupons: untargeted *low privacy* (LP) or targeted *high privacy* (HP). At each time, the retailer may offer a coupon and the consumer transitions from one state to another according to a Markov chain that is independent of the offered coupon. The retailer suffers a cost that depends both on the type of coupon offered and the state of the consumer. The costs reflect the advantage of offering targeted HP coupons relative to untargeted LP ones while simultaneously capturing the risk of doing so when the consumer is already "Alerted".

Under the assumption that the retailer (via surveys or prior knowledge) knows the statistics of the consumer Markov process, i.e., the likelihoods of becoming "Alerted" and staying "Alerted", and a belief about the initial consumer state, we study the problem of determining the optimal coupon-offering policy that the retailer should adopt to minimize the long-term discounted costs of offering coupons. We extend the simple model above to multiple states and coupon-dependent transitions. We model the latter via two Markov processes for the consumer, one for each type (HP or LP) of coupon such that a persnickety consumer who is easily "Alerted" will be more likely to do so when offered an HP (relative to LP) coupon. Furthermore, for noisy costs, we propose a heuristic method to compute the decision policy. Moreover, if the initial belief state is unknown to the retailer, we use a Bayesian model to estimate the belief state. Our main results can be summarized as follows:

1. There exists an optimal, stationary, threshold-based policy for offering coupons such that a HP coupon is offered only if the belief of being in the "Alerted" state at each interaction time is below a certain threshold; this threshold is a function of all the model parameters. This structural result holds for multiple states and coupon-dependent transitions.

2. The threshold for offering a targeted HP coupon increases in the following cases:

   (a) once "Alerted", the consumer remains so for a while – the retailer is more willing to take risks since the the consumer takes a while to transition to "Normal";

   (b) the consumer is very unlikely to get "Alerted";

   (c) the cost of offering an untargeted LP coupon is high and close to the cost of offering a targeted HP coupon to an "Alerted" consumer; and

(d) when the retailer does not discount the future heavily (future rewards nearly as important as present), the retailer stands to benefit by offering HP coupons for a larger set of beliefs about the consumer's state. Conversely, when the retailer discounts the future heavily, it values the present rewards more than future rewards. Thus, the retailer tends to play conservatively so that it will not "creep out" the consumer in the present.

3. For the coupon-dependent Markov model for the consumer, the threshold is smaller than for the non-coupon dependent case which encapsulates the fact that highly sensitive consumers will force the retailers to behave more conservatively.

4. By adopting a heuristic threshold policy computed by the mean value of costs, the retailer can minimize the discounted cost effectively even if costs are noisy. Moreover, the Bayesian approach helps the retailer to estimate the consumer state when the initial belief state is unknown.

Our results use many fundamental tools and techniques from the theory of MDPs through appropriate and meaningful problem modeling. We briefly review the related literature in consumer privacy studies as well as MDPs.

1.1.3 Consumer Privacy Models

Several economic studies have examined consumer's attitudes towards privacy via surveys and data analysis including studies on the benefits and costs of using private data. Taylor [3] discovers that the market for consumer information provides companies with incentives to charge high experimental prices for studying consumer behavior and privacy aware customers strategically reduce their demand in order to protect their privacy. The authors of [4] explore how individual privacy will changes over time. Their results show that the amount of privacy for each individual will decline over time and it will be increasingly difficult to maintain privacy. Reference [5]

shows, somewhat surprisingly, that individuals are not strictly rational in privacy-related decision making. However, to date, no formal model has been proposed that captures consumer's privacy sensitivity.

In understanding the use of private information, several approaches have been taken to handle the tradeoff between the privacy of individuals (consumers) and the efficiency or utility of the data user (the retailer). The benefits and costs of using private data have been studied by Aquisti and Grossklags [6]. Their work suggests that solving the privacy problem means to find a balance between information sharing and information hiding that is in the interest of data subjects as well as of the society as a whole. Sankar et al. propose an information theoretic approach to capture privacy and utility tradeoff [7]. They use rate-distortion approach to develop a utility-privacy tradeoff region for *i.i.d.* data sources with known distribution in the database. Formal methods such as differential privacy [8, 9] are finding use in modeling the value of private data for market design [10] and for the problem of partitioning goods with private valuation function amongst the agents [11]. In these models the goal is to elicit private information from individuals. As more of the purchase transaction and data become electronic, consumers are becoming increasingly aware that their electronic purchases and other activities are being monitored. To the best of our knowledge, a formal model for consumer-retailer interactions and the related privacy issues has not been studied before; in particular, our work focuses on explicitly considering the consequence to the retailer of the consumers' awareness of privacy violations.

1.1.4 Markov Decision Processes

In this project we study the problem from the retailer's perspective: given a privacy-sensitive consumer who may become alerted to privacy violations, how should the retailer offer incentives (which we call *coupons*) to maximize its revenue? The problem has been modeled as a Markov decision process (MDP) in which the con-

sumer's state changes over time according to a Markov chain and the retailer can offer a coupon at each time step which involves either high or low privacy risk to the consumer. The retailer bears a cost that is dependent on the type of coupon offered and the state of the consumer. In this model, offering a high privacy risk coupon can provide information about the consumer state, but risks lower revenue if the consumer is alerted.

Markov decision processes (MDPs) are common discrete time mathematical models for decision making when observable outputs are partially depend on internal states and exterior inputs. It have been widely used for decades across many fields (see [12, 13]). We briefly list a few closely-related works. Lipsa and Martins [14] and Nayyar et al. [15] have used MDPs to model remote sensing and communication problems. In the Lipsa-Martins model, an observer has causal access to a first-order linear time-invariant system and must communicate with an estimator over a cost-constrained communication link. The goal is to minimize a joint cost given by the estimation error and the communication cost. Nayyar et al. study a similar model for a discrete state estimation model where the communication cost is dictated by an energy-harvesting process that constrains the sensor/observer. Both of these works study finite-horizon problems where the cost is a combination of two costs. The analogy to our problem is that offering a targeted coupon allows the retailer to estimate the state of the consumer. However, in our model the retailer goal is not to estimate the consumer state but to minimize cost. The model we use is most similar to Ross's model of product quality control with deterioration [16], which was more recently used by Laourine and Tong to study the Gilbert-Elliot channel in wireless communications [17], in which the channel has two states and the transmitter has two actions (transmit or not). They study an infinite-horizon problem with costs associated to different types of data transmission rates. They establish the stationary optimal pol-

icy in a model which has two states and three actions with a fixed transition matrix and perfect knowledge of belief states. We cannot apply their results directly due to our different cost structure, but use ideas from their proofs. Furthermore, we go beyond these works to study privacy-utility tradeoffs in consumer-retailer interactions with more than two states and action-dependent transition probabilities. We apply more general MDP analysis tools to address our formal behavioral model for privacy-sensitive consumers.

Classical target-search problems [18, 19] also use MDPs to develop optimal policies for tracking action. In reference [18], the authors study the problem of tracking a moving target. They formulate this problem as partially observable MDP (POMDP) by assuming the target is in one of many states (locations) and the movement of the target follows a Markov chain. The decision process is terminated when the target has been found. The action set is the possible state of the target, only one state can be searched at each time and there is a cost associated to searching each state. Also, the decision maker has an overlook probability which is the probability that it may not find the target even if the right state is searched. They prove that the optimal policy has threshold structure when the overlook probabilities and the search costs are all the same among different states. They also show that in the general overlook probability case with general cost, the optimality equation (value function) is satisfied by a piecewise linear function and the threshold property holds for a large proportion of the possible transition matrices, search costs and overlook probabilities. However, in our work, we look at a different problem in which we want to capture interaction between retailer and consumer with privacy concerns via offering coupons. Our major objective is to minimize cost of the retailer over an infinite horizon rather than track the state. Moreover, our cost not only depends on the state, but also on the action taken by the decision maker, i.e. we may have two costs in alerted state

depending on the action of the retailer. Similar to their work, we study the two state model and extend to multi-state problem. Moreover, we consider the case where the transition matrix is dependent on the action of the decision maker. Mansourifard et al. [19] consider a state tracking problem with Markov transition and develope bounds and a heuristic policy for state estimation. Our formulation is different: we consider cost minimization problems for retailers when consumers have privacy concerns and develop a closed form solution for the stationary optimal policy in different Markov transition models.

In the context of privacy, MDPs have been used by Venkitasubramaniam [20] to study privacy and utility trade-off in control systems with time-varying state by quantifying privacy via the information-theoretic equivocation function. However, in his paper, the state is really the state of a control system rather than the state of privacy sensitivity of a consumer. While this approach has some similarity to ours in terms of using MDP model, his method uses an average reward which is different from our problem because the cost of consumer privacy violation has a short-term effect. In our work we do not quantify *privacy loss* directly; instead we model *privacy-sensitivity states* and resulting user behavior via MDPs to determine interaction policies that can benefit both consumers and retailers. To the best of our knowledge, a formal model for consumer-retailer interactions and the related privacy issues has not been studied before; in particular, our work focuses on explicitly considering the consequence to the retailer of the consumers' awareness of privacy violations.

1.2 Incentive Mechanisms for Privacy-sensitive Electricity Consumers with Alternative Energy Sources

Alternative energy sources, especially rooftop photovoltaic (PV) systems, are getting more prevalent at the distribution level of the electric power grid. As a result, there is a need to monitor energy consumption patterns at the distribution level

for more efficient dispatch and stable/reliable system operations. However, such advanced metering infrastructure can create potential threats to consumer privacy since they have much higher sampling rate and data processing capability than traditional meters.

The ability to collect electricity consumption data from consumers benefits the electricity provider in many ways, including improving load forecasting and system dispatch efficiency. This can be achieved via the use of smart meters that provide fine grained energy usage information to the electricity provider. However, the collected information may be used by malicious users or third party data processing entities to analyze consumers' electricity consumption behaviors and make inferences about personal habits of consumers. Thus, privacy-sensitive consumers can use their alternative energy sources (e.g., battery and PV) to mask their consumption, or even refuse to use smart meters so that they can have some privacy. While alternative energy sources can provide some measure of privacy, it is not a reliable resource of energy due to uncertainties in weather conditions. Thus, consumers may also have to turn to the grid for energy when alternative energy sources cannot meet their demand.

From the electricity provider's perspective, the uncertainties in alternative energy sources may also cause provision issues. Thus, it is in the electricity provider's interest to incentivize consumers to consume a desired amount of energy to maintain power system's stability [21]. Also, it is in the consumers' interest to exploit these incentives while simultaneously ensuring a certain level of privacy. To this end, we monetize consumer privacy using a valuation function that captures the fact that privacy leakage of a consumer is directly a function of the power that it consumes from the grid. We address the problem of how electricity providers can incentivize privacy-sensitive consumers with access to alternative energy sources to consume a basic amount of electricity from the grid. In short, the goal of our price-based incen-

tive approach is to allow both parties, namely consumers and the electricity provider, to negotiate consumption and data sharing such that all parties can potentially profit from interactions.

1.2.1 Background and Related Work

The increasing number of smart meters deployed in business and residential buildings has raised concerns about privacy. The adversary can make inferences about consumers' energy consumption behavior via data collected from smart meters [22, 23, 24]. Multiple methods and metrics have been proposed to quantify and protect smart meter privacy including using battery to hide consumption [25, 26, 27, 28], distorting the metering data [29], using anonymization of smart meter [30] and reducing sampling rate of smart meters via contracts [31]. In [32], Denic *et.al.* show that privacy preserving algorithms which use battery to mask load behavior can affect consumers' demand for electricity from the grid and electricity prices. Thus, it is possible that privacy protection mechanisms can affect the reliable operation of the grid. Finally, from a demand response view point, using price as control signals have been studied extensively, e.g., [33, 34, 35] and references therein.

Our model considers consumers who can achieve privacy by masking their consumption profiles using alternative energy sources. In contrast to prior work, our pricing mechanism focuses on balancing needs of the electricity provider against consumer privacy. Furthermore, in place of an abstract privacy metric, we monetize privacy leakage via an arbitrary valuation function dependent on the amount of electricity the consumer consumes from the grid.

1.2.2 Contributions

The main contribution of this project is to propose a novel approach to study the trade-off between privacy and energy cost minimization for consumers via incentives offered by an electricity provider to consume power directly from the grid. To this

end, we formulate a multi-player non-cooperative game to model interactions between consumers and the electricity provider. In this game, the strategy of the electricity provider is the incentive price it offers to encourage consumers to compromise certain level of privacy by consuming a desired amount of power from the grid. On the other hand, the strategy of a consumer is to select the proportion of electricity it consumes from the grid to exploit its reward from the grid while ensuring that the corresponding valuation of privacy leakage is acceptable.

In this problem, one can consider a pure strategy where consumers and the electricity provider decide on a specific consumption and price value from a range of options. However, it is also possible that consumers may not make deterministic decisions but may choose out of various strategies with different likelihoods. As a result, the response of the electricity provider will also be a corresponding random choice. To make the analysis tractable, we look at discrete sets of options for both consumers and the electricity provider. In practice, since pricing is often a tiered model with discrete levels, this model also captures this practical setting. In this project, we study a more general problem by allowing for uncertainties in the behavior of consumers, and in response, the provider by considering mixed strategies. We focus on a two-player game with two levels of consumption strategies and a two-tiered pricing structure. For this model, we prove the existence and uniqueness of the non-degenerate mixed, i.e., non-pure strategy Nash equilibrium. The proposed incentive mechanism increases both the net profit of the electricity provider and the reward for the consumer for specific choices of reward and profit functions.

1.3 The Impact of Privacy on Free Online Service Markets

There has been a steady increase in online interactions between consumers and retailers, where the term retailer refers to entities who offer products for free (e.g., social media, search engines, free applications, to name a few). The advances in tech-

nology have enabled retailers (henceforth referred to as *service providers*) to collect, store, process, sell, and share customer-specific information for targeted advertising (ads) and tiered pricing tactics. In fact, many oft used online services are free and consumers implicitly accede to tracking for customized services. Targeted ads are a part of the emerging revenue/profit model for such service providers (SPs) offering free services. Consumers are delighted by free services until they begin encountering privacy violations on a daily/frequent basis. While such infractions taken individually could be ignored or discounted, the totality of data available about consumers with a variety of retailers and the resulting privacy consequences raise serious concerns [36].

Service providers are beginning to acknowledge that consumers are sensitive to privacy violations. For example, Google [37] and Apple [38] recently adopted differentially private mechanisms for collecting user data for statistical analyses. However, the details of these mechanisms are opaque and offer even less clarity on whether the consumer actually has a choice. In this context, it is worth understanding if privacy differentiated services can provide such choices for consumers. In a competitive marketplace, the aggregated weight of targeting may drive some consumers to seek more privacy-protective alternatives. The cost to the consumer of this action may be a lower quality of service (QoS) (e.g., poorer search engine capabilities). However, it could eventually lead to a more open model for consumer sharing of private information, i.e., one from implicit assent to informed consent [36, 39].

To understand the influence of consumers' heterogeneous privacy preference on SPs' behavior in a competitive market, we take a game-theoretic approach to model the interactions between SPs and consumers. In particular, we address the following questions:

- Can privacy-differentiated services lead to a sustainable marketplace? With consumers' heterogeneity in privacy preference, SPs can offer services with dif-

ferent QoS and privacy risks. This question deals with whether there exists a market equilibrium that sustains the competing SPs for various consumer/SP parameters.

- What are the equilibrium QoS-privacy risk strategies for the SPs? This question is related to calculating the equilibrium behavior of the SPs. For different market models, we examine the optimal strategies for SPs under competition.

- How do various consumer/SP parameters, such as consumers' privacy preference/valuation and SPs' profit/cost affect the equilibrium outcome? Given different model parameters, we investigate the effect of each one of these parameters on the equilibrium outcome of the competition between SPs.

We also generalize the model to multiple SPs (e.g., Google, DuckDuckGo, and Bing) and illustrate the instability of multi-competitor markets.

1.3.1 Related Work

Targeted advertising is a common method for service providers to exploit knowledge of consumers; this in turn can lead to privacy violations. Our work is informed by the literature on targeting strategies for retailers [40, 41, 42, 43, 44, 45, 46, 47, 48, 49], but rather than optimizing retailer strategies, we are interested in identifying how privacy differentiated services can address privacy concerns.

The problem of market segmentation is a classic and well-studied problem in microeconomics [50] with focus on how pricing and product differentiation can lead to a stable and competitive marketplace. However, the free online service market presents a new challenge wherein monetary quantification of both 'free' services and the data collected about consumers is not simple and straightforward. Equally challenging is the quantification of consumer privacy since it requires capturing the heterogeneous expressions of privacy sensitivity that can range from 'don't care' at one extreme

to 'hyper vigilant' at the other. However, some aspects of market models can be brought to bear to our problem; in particular, the oligopolistic market model with a small number of competitors, barriers to entry that are not as high as those for monopolies, and with differentiated products fits appropriately for the markets we are considering wherein two or (a few) more service providers offer products of the same type but differentiated by QoS and privacy risk.

A nuanced model that captures differentiation between two firms and consumer preferences is the Hotelling model [51]. It has been widely used for market analysis across many fields [52, 53]. This model captures differentiation between market players by mapping firms to positions on a unit length line such that the location is indicative of the firm's 'differentiation level', the total line length is reflective of the entire market, a consumer's privacy preference is a point on the line, and the optimal locations of the firm results from the simultaneous game between the players indicate the resulting segmentation. The model captures utility for consumer as both the advantage (price) from the firm as well as the 'transportation cost' from the consumer's location to that firm. Consumers choose the seller which gives them the highest utility (in terms of the price as well as the 'transportation' costs).

Privacy and market segmentation. An extensive body of literature on economic models for privacy was recently reviewed by Acquisti *et al.* [54]. These models illustrate the large semantic range covered by the word "privacy". Wang *et.al.* [9] study the value of privacy in a market that allows trading private data as commodity. By modeling the interaction between a single data collector and consumers as a game, they show that in a Nash equilibrium, the data collector offers a payment which equals to the monetary value of data privacy in the market for private data. Meanwhile, a consumer's best response is to report the data with the same value to the data collector. Different from their work, our paper focuses on the interac-

tions between competing free online SPs and consumers. Jentzsch *et.al.* [55] propose a model to study competitions between two service providers by taking consumer's privacy preference (binary choices: low privacy/high privacy) into account using a vertical Hotelling model. Thus, consumers select the service provider based on their privacy concerns and the amount of payment to the service provider. They provide analysis of equilibrium strategies for SPs. Lee *et al.* [56] study the influence of privacy protection on the segmentation of a duopoly. In their model, firms may offer standard and personalized products with personalized prices to three different types of privacy-sensitive consumers (the 'unconcerned' who always share information, 'pragmatic' who only share if a firm adopts privacy protection, and the 'fundamentalists' who never share). They show that a privacy-friendly firm can enlarge market share by attracting more 'pragmatists' to share personal information. From this expansion it can earn more profits rather than compete with its rival for the other consumers. In contrast to both above-mentioned models, our model differs in focusing on 'free' services, and thus, introduces new models for quantifying QoS- and privacy-based differentiators; furthermore, our model generalizes the discrete set of privacy sensitive consumers in [56] to a continuous set of privacy risks, thus allowing analysis over an entire range of privacy expression and a more nuanced view of how SPs should offer services to all types of consumers.

1.3.2 Contributions

We propose a novel model for the privacy differentiated market segmentation problem in which service providers offer free services differentiated by QoS and privacy risk. Our model captures a variety of free online services such as search engines, social networking sites, and software apps that are free, and therefore, use consumer data in a variety of ways for revenue generation. Each SP's gain from using consumer data is captured by a revenue function and its cost of doing so is captured by a

cost function. The goal of each SP is to choose a QoS and privacy risk tuple that maximizes its profit (difference of revenue and cost). We assume that consumers can map their heterogeneous privacy sensitivity to a quantitative scale. The SPs use this quantitative scale to differentiate themselves. Each consumer chooses the SP that maximizes a desired function of its privacy risk valuation and the QoS-privacy risk tuple offered by the SP.

Our model is built upon the classical 'spatial' Hotelling model [51] wherein the location is now proxy for privacy risk (that both SPs offer and consumers prefer). The QoS offered by the SP models the product price in the Hotelling model. In contrast to the classical Hotelling model in which there is a non-negative transportation cost irrespective of the locations of consumer and retailer, here consumers will always benefit from SPs that offer lower privacy risk than what they prefer. Thus, there is an asymmetry in the transportation cost. We model the interactions between SPs and consumers as a three-stage sequential game and compute the equilibrium QoS-privacy risk tuple as well as consumers' choices using backward induction. We use the equilibrium strategies to compute the resulting market share and profit for specific models of cost and revenue (to SPs), distribution of consumer heterogeneous privacy choices, as well as consumer privacy valuation.

We show that there does not exist any equilibrium in which both SPs offer the same privacy risk for the two-SP market with linear valuation function (cost, revenue, consumer utility). Furthermore, when the privacy preference of consumers follows a uniform distribution, we can obtain closed form solutions for the two-SP market with linear valuation functions. For this settings, our results highlight the following: (i) when consumers place a high value on privacy, it leads to a lower use of private data by SPs, i.e., their advertised privacy risk reduces; (ii) SPs offering high privacy risk services are sustainable only if they offer sufficiently high QoS; (iii) SPs that are ca-

pable of differentiating on services that do not directly use consumer data gain larger market share; and (iv) higher consumer privacy valuation "softens" the competition between SPs. We also study the case in which consumer's privacy preference follows a truncated Gaussian distribution. Since it is very difficult to obtain a closed form solution, we analyze the market numerically. Based on our numerical result, we observe similar behavior in the equilibrium strategies and market share compared to the uniform case. In extending the work to more than two SPs, we illustrate the instability of such markets and highlight the challenges of studying market segmentation for more than two participants (a problem acknowledged in economics [57]).

1.4 Generative Adversarial Privacy

The explosion of information collection across a variety of electronic platforms is enabling the use of *inferential machine learning* (ML) and artificial intelligence to guide consumers through a myriad of choices and decisions in their daily lives. In this era of artificial intelligence, data is quickly becoming the most valuable resource [58]. Indeed, large scale datasets provide tremendous *utility* in helping researchers design state-of-the-art machine learning algorithms that can learn from and make predictions on real-life data. Scholars and researchers are increasingly demanding access to larger datasets that allow them to learn more sophisticated models. Unfortunately, more often than not, in addition to containing *public* information that can be shared or published, large scale datasets also contain *private* information about participating individuals (see Figure 1.1). Thus, data collection and curation organizations are reluctant to release such datasets before carefully *sanitizing* them, especially in light of recent public policies on data sharing [59, 36].

To protect the privacy of individuals, datasets are typically anonymized before their release. This is done by stripping off personally identifiable information (e.g., first and last name, social security number, IDs, etc.) [60, 61, 62]. Anonymiza-

Original meter data X | Private data Y

	Meter data 10:00, 09/06/2010	…	Meter data 23:30, 05/06/2011	Income	Occupancy
Entry (row 1)	0.140	…	0.253	70,000	1
Entry (row 2)	0.108	…	0.371	60,000	3
Entry (row 3)	0.248	…	0.192	200,000	4
⋮	⋮	⋮	⋮	⋮	⋮
Entry (row n)	0.210	…	0.182	150,000	3

Database D

Perturbation

Perturbed meter data $\hat{X}$

Meter data 10:00, 09/06/2010	…	Meter data 23:30, 05/06/2011
0.231	…	0.302
0.158	…	0.350
0.226	…	0.176
⋮	⋮	⋮
0.179	…	0.202

Figure 1.1: An example privacy preserving mechanism for smart meter data

tion, however, does not provide immunity against correlation and linkage attacks [63, 64]. Indeed, several successful attempts to re-identify individuals from anonymized datasets have been reported in the past ten years. For instance, [63] is able to successfully de-anonymize watch histories in the Netflix Prize, a public recommender system competition. In a more recent attack, [65] showed that participants of an anonymized DNA study were identified by linking their DNA data with the publicly available Personal Genome Project dataset. Even more recently, [66] successfully designed re-identification attacks on anonymized fMRI imaging datasets. Other annoymization techniques, such as generalization [67, 68, 69] and suppression [70, 71, 72], also cannot prevent an adversary from performing the sensitive linkages or recover private information from published datasets [73].

Addressing the shortcomings of anonymization techniques requires data randomization. In recent years, two randomization-based approaches with provable *statistical privacy* guarantees have emerged: (a) context-free approaches that assume worst-case dataset statistics and adversaries; (b) context-aware approaches that explicitly model the dataset statistics and adversary's capabilities.

Context-free privacy. One of the most popular context-free notions of privacy is *differential privacy* (DP) [74, 75, 76]. DP, quantified by a leakage parameter ϵ *, re-

*Smaller $\epsilon \in [0, \infty)$ implies smaller leakage and stronger privacy guarantees.

stricts distinguishability between *any* two "neighboring" datasets from the published data. DP provides strong, context-free theoretical guarantees against worst-case adversaries. However, training machine learning models on randomized data with DP guarantees often leads to a significantly reduced utility and comes with a tremendous hit in sample complexity [77, 78, 79, 80, 81, 82] in the desired leakage regimes. For example, learning population level histograms under local DP suffers from a stupendous increase in sample complexity by a factor proportional to the size of the dictionary [83, 84, 81].

Context-aware privacy. Context-aware privacy notions have been so far studied by information theorists under the rubric of *information theoretic* (IT) privacy [85, 86, 87, 88, 89, 90, 91]. IT privacy has predominantly been quantified by mutual information (MI) which models how well an adversary, with access to the released data, can refine its belief about the private features of the data. Recently, Issa *et al.* [92] introduced *maximal leakage* (MaxL) to quantify leakage to a strong adversary capable of guessing any function of the dataset. They also showed that their adversarial model can be generalized to encompass local DP (wherein the mechanism ensures limited distinction for *any* pair of entries—a stronger DP notion without a neighborhood constraint [93, 94]) [95]. When one restricts the adversary to guessing specific private features (and not all functions of these features), the resulting adversary is a maximum *a posteriori* (MAP) adversary that has been studied by Asoodeh *et al.* in [96, 97].

Compared to context-free privacy notions, context-aware privacy notions achieve a better privacy-utility tradeoff by incorporating the statistics of the dataset and placing reasonable restrictions on the capabilities of the adversary. However, using information-theoretic quantities (such as MI) as privacy metrics requires learning the parameters of the privatization mechanism in a data-driven fashion that involves

minimizing an empirical information-theoretic loss function. This task is remarkably challenging in practice [98].

Generative adversarial privacy. Given the challenges of existing privacy approaches, we take a fundamentally new approach towards enabling private data publishing with guarantees on both privacy and utility. Instead of adopting worst-case, context-free notions of data privacy (such as differential privacy), we introduce a novel context-aware model of privacy that allows the designer to cleverly add noise where it matters. An inherent challenge in taking a context-aware privacy approach is that it requires having access to priors, such as joint distributions of public and private variables. Such information is hardly ever present in practice. To overcome this issue, we take a *data-driven approach* to context-aware privacy. We leverage recent advancements in generative adversarial networks (GANs) to introduce a unified framework for context-aware privacy called *generative adversarial privacy* (GAP). Under GAP, the parameters of a generative model, representing the privatization mechanism, are learned from the data itself.

1.4.1 Contributions

We investigate a setting where a data holder would like to publish a dataset $\mathcal{D}$ in a privacy preserving fashion. Each row in $\mathcal{D}$ contains both private variables (represented by Y) and public variables (represented by X). The goal of the data holder is to generate $\hat{X}$ in a way such that: (a) $\hat{X}$ is as good of a representation of X as possible, and (b) an adversary cannot use $\hat{X}$ to reliably infer Y. To this end, we present GAP, a unified framework for context-aware privacy that includes existing information-theoretic privacy notions. Our formulation is inspired by GANs [99, 100, 101] and error probability games [102, 103, 104, 105, 106]. It includes two learning blocks: a *privatizer*, whose task is to output a sanitized version of the public variables (subject to some distortion constraints); and an *adversary*, whose task is

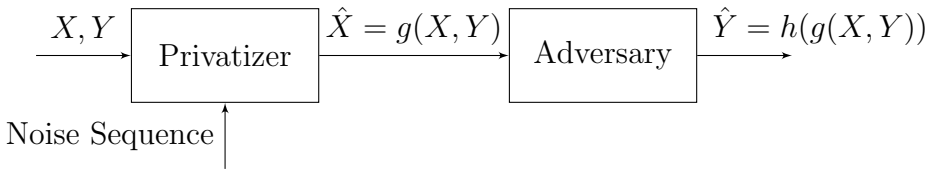

Figure 1.2: Generative adversarial privacy

to learn the private variables from the sanitized data. The privatizer and adversary achieve their goals by competing in a constrained minimax, zero-sum game. On the one hand, the privatizer (a conditional generative model) is designed to minimize the adversary's performance in inferring Y reliably. On the other hand, the adversary (a classifier) seeks to find the best inference strategy that maximizes its performance. This generative adversarial framework is represented in Figure 1.2.

We list our main contributions below.

1. We introduce GAP as a minimax game-theoretic formulation (see Figure 1.2) to design privacy mechanisms matched to an adversarial model.

2. We show that our framework captures a rich class of statistical adversaries. This allows us to compare data-driven approaches directly against strong inferential adversaries (e.g., a maximum *a posteriori* (MAP) probability maximizing adversary with access to dataset statistics).

3. We make precise connections between data-driven privacy methods and the minimax game-theoretic GAP formulation; this implies that when: (i) the neural networks used in the data-driven approach have sufficient capacity, (ii) the learning rate is sufficiently small, and (iii) the training data is sufficiently large, the learned privacy scheme converges to the game-theoretically optimal one.

4. To showcase the power of our data-driven framework, we investigate several simple, albeit canonical, datasets: binary data model in which X and Y are both

random variables, and Gaussian Mixture Model (GMM) where Y is binary and X is a conditionally multi-dimensional Gaussian vector. We derive and compare the performance of game-theoretically optimal privatization mechanisms with those that are directly learned in a data-driven fashion to show that the gap between theory and practice is negligible.

5. Finally, we demonstrate the performance of GAP on meaningful, widely used datasets. We first use the GENKI dataset [107], for which we identify the public (images of faces) and private (gender) features. Next we test our GAP framework on the MNIST dataset [108] for which we consider the images of hand-written digits and a binary variable which identifies whether there is a circular structure in the digit (e.g., digits $0, 6, 8, 9$ contain circular structure) as public and private features, respectively. Our results show that GAP can significantly reduce an adversary's capability of inferring private features with limited amount of distortion on the public features. Furthermore, we show that GAP allows data receivers to learn other non-private features from the privatized data.

1.4.2 Related Work

In practice, a context-free notion of privacy (such as DP) is desirable because it places no restrictions on the dataset statistics or adversary's strength. This explains why DP has been remarkably successful in the past ten years, and has been deployed in array of systems, including Google's Chrome browser [37] and Apple's iOS [109]. Nevertheless, because of its strong context-free nature, DP has suffered from a sequence of impossibility results. These results have made the deployment of DP with a reasonable leakage parameter practically impossible. Indeed, it was recently reported that Apple's DP implementation suffers from several limitations—most notable of which is Apple's use of unacceptably large leakage parameters [110].

Context-aware privacy notions can exploit the structure and statistics of the dataset to design mechanisms matched to both the data and adversarial models. In this context, information-theoretic metrics for privacy are naturally well suited. In fact, the adversarial model determines the appropriate information metric: an estimating adversary that minimizes mean square error is captured by χ^2-squared measures [111], a belief refining adversary is captured by MI [87], an adversary that can make a hard MAP decision for a specific set of private features is captured by the Arimoto MI of order ∞ [96, 97], and an adversary that can guess any function of the private features is captured by the maximal (over all distributions of the dataset for a fixed support) Sibson information of order ∞ [92, 95].

Information-theoretic metrics, and in particular MI privacy, allow the use of Fano's inequality and its variants [112] to bound the rate of learning the private variables for a variety of learning metrics, such as error probability and minimum mean-squared error (MMSE). Despite the strength of MI in providing statistical utility as well as capturing a fairly strong adversary that involves refining beliefs, in the absence of priors on the dataset, using MI as an empirical loss function leads to computationally intractable procedures when learning the optimal parameters of the privatization mechanism from data. Indeed, training algorithms with empirical information-theoretic loss functions is a challenging problem that has been explored in specific learning contexts, such as determining randomized encoders for the information bottleneck problem [98] and designing deep auto-encoders using a rate-distortion paradigm [113, 114]. Even in these specific contexts, variational approaches were taken to minimize/maximize a surrogate function instead of minimizing/maximizing an empirical mutual information loss function directly [115]. In an effort to bridge theory and practice, we present a general data-driven framework to design privacy mechanisms that can capture a range of information-theoretic privacy metrics as loss functions. We will show how

our framework leads to very practical (generative adversarial) data-driven formulations that match their corresponding theoretical formulations.

In the context of publishing datasets with privacy and utility guarantees, a number of similar approaches have been recently considered. We briefly review them and clarify how our work is different. In [116], the authors consider linear privatizer and adversary models by adding noise in directions that are orthogonal to the public features in the hope that the "spaces" of the public and private features are orthogonal (or nearly orthogonal). This allows the privatizer to achieve full privacy without sacrificing utility. However, this work is restrictive in the sense that it requires the public and private features to be nearly orthogonal. Furthermore, this work provides no rigorous quantification of privacy and only investigates a limited class of linear adversaries and privatizers.

DP-based obfuscators for data publishing have been considered in [117, 118]. The author in [117] considers a deterministic, compressive mapping of the input data with differentially private noise added either before or after the mapping. The mapping rule is determined by a data-driven methodology to design minimax filters that allow non-malicious entities to learn some public features from the filtered data, while preventing malicious entities from learning other private features. The approach in [118] relies on using deep auto-encoders to determine the relevant feature space to add differentially private noise to, eliminating the need to add noise to the original data. After noise adding, the original signal is reconstructed. These novel approaches leverage minimax filters and deep auto-encoders to incorporate a notion of context-aware privacy and achieve better privacy-utility tradeoffs while using DP to enforce privacy. However, DP will still incur an insurmountable utility cost since it assumes worst-case dataset statistics. Our approach captures a broader class of randomization-based mechanisms

via a generative model which allows the privatizer to tailor the noise to the statistics of the dataset.

Our work is also closely related to adversarial neural cryptography [119], learning censored representations [120], and privacy preserving image sharing [121], in which adversarial learning is used to learn how to protect communications by encryption or hide/remove sensitive information. Similar to these problems, our model includes a minimax formulation and uses adversarial neural networks to learn privatization schemes. However, in [120, 121], the authors use non-generative auto-encoders to remove sensitive information, which do not have an obvious generative interpretation. Instead, we use a GANs-like approach to learn privatization schemes that prevent an adversary from inferring the private data. Moreover, these papers consider a Lagrangian formulation for the utility-privacy tradeoff that the obfuscator computes. We go beyond these works by studying a game-theoretic setting with constrained optimization, which provides a specific privacy guarantee for a fixed distortion. We also compare the performance of the privatization schemes learned in an adversarial fashion with the game-theoretically optimal ones for some canonical data models.

We use conditional generative models to represent privatization schemes. Generative models have recently received a lot of attention in the machine learning community [122, 123, 100, 101, 99]. Ultimately, deep generative models hold the promise of discovering and efficiently internalizing the statistics of the target signal to be generated. State-of-the-art generative models are trained in an adversarial fashion [101, 99]: the generated signal is fed into a discriminator which attempts to distinguish whether the data is real (i.e., sampled from the true underlying distribution) or synthetic (i.e., generated from a low dimensional noise sequence). Training generative models in an adversarial fashion has proven to be successful in computer vision and enabled several exciting applications [124, 125, 126]. Analogous to how the generator

is trained in GANs, we train the privatizer in an adversarial fashion by making it compete with an attacker.

1.5 Outline of Dissertation

In the first part, this dissertation investigates the decision making problem in three scenarios: (i) designing incentive schemes for privacy sensitive users, in which a retailer seeks to offer incentives (coupons) to maximize its profit while minimally "creep out" consumers; (ii) game theoretic incentive schemes for encouraging privacy sensitive households to share energy consumption data with the grid; (iii) market segmentation for privacy differentiated free services. In the second part, this dissertation studies a unified framework for various information-theoretic privacy and privacy mechanisms that can be learned directly from data. A brief introduction with literature review is presented in Chapter 1. The rest of this dissertation is organized as follows.

In Chapter 2, we study how to design incentive schemes for consumers who are privacy sensitive. A detailed description of dynamic modeling of interactions between the retailer and privacy aware consumers is provided. A two-state, two-action POMDP model is used to capture the cost minimization problem of offering coupons to consumers with privacy concerns. Furthermore, extensions of this formulation including multi-level state and action dependent transition model are described. Models for studying coupon offering policies under noisy costs and unknown consumer state are also described in this chapter. The optimal coupon offering policy is derived to minimize the discounted cost associated to consumers' response to privacy sensitive coupon. Model analysis is then conducted and some interesting properties of the optimal coupon offering policy mode is identified. After that, extensions to multi-level consumer state and coupon dependent transition are studied. Based on the two-state, two-action model, a heuristic method is proposed to make decisions when the received

cost is noisy. Also, a Bayesian data analysis framework is proposed to estimate the consumer behavior when the initial state of the consumer is unknown to the retailer.

In Chapter 3, we propose a novel approach to study the tradeoff between privacy and energy cost minimization under the assumption that the utility company offers incentives to households to encourage data sharing through energy consumption. A non-cooperative game is formulated to model interactions between households and the utility company. Under certain assumptions on the utility functions and strategy sets, we prove that the mixed strategy Nash equilibrium exists and provide a closed form solution of the mixed strategy Nash equilibrium. For a specific choice of utility functions, we illustrate the influence of the proposed mechanism on the net profit, supply-demand imbalance of the electricity provider, and consumer benefits.

In Chapter 4, we seek to understand the effect of offering privacy- and QoS- differentiated online services on consumers with heterogeneous expressions of privacy sensitivity. We have quantified the influence of privacy differentiated services as the fraction of consumers that prefer each type of QoS and privacy risk tuple. We have presented an analysis built upon the classical Hotelling model to compute these fractions for both the two- and multi- SP problems. Similar to the classical segmentation models, our problem also involves parameters that capture cost, revenue, and consumer valuation functions. We study the market segmentation for relatively simple yet meaningful functions such as linear cost models and uniform (as well as truncated Gaussian) distribution. Furthermore, we extend our analysis to multi-SP case and discovered instability of market segmentation in a market for more than two SPs.

In Chapter 5, we formally present our GAP model. We also show how, as a special case, it can recover several information-theoretic notions of privacy. We then study several simple (but canonical) dataset models (e.g., binary data model and Gaussian mixture data model). In particular, we present theoretically optimal privacy

mechanisms, and demonstrate how privacy mechanisms can be learned from data using a generative adversarial network. For both models, we show that the privacy mechanisms learned from data match the theoretically optimal ones. Finally, we showcase the performance of GAP on the GENKI and MNIST dataset.

Chapter 2

HOW TO INCENTIVIZE AND INTERACT WITH PRIVACY SENSITIVE CONSUMER?

2.1 Problem Formulation for Consumer Retailer Interactions

We model interactions between a retailer and a consumer via a discrete-time system (Figure 2.1). At each time t, the consumer has a discrete-valued state and the retailer may offer one of two coupons: high privacy risk (HP) or low privacy risk (LP). We assume a sophisticated consumer who can distinguish whether a coupon is HP or LP and responds to the personalized coupon by imposing a cost on the retailer that depends on the coupon offered and its own state. For example, a consumer who is "alerted" (privacy-aware) may respond to an HP coupon by imposing a high cost to the retailer, such as reducing purchases at the retailer. The retailer's goal is to decide which type of coupon to offer at each time t to minimize its cost.

2.1.1 Consumer with Two States and Coupon Independent Transition.

Consumer Model

Modelling Assumption 1. ***(Consumer's state)*** *We model the consumer's response to coupons by assuming them to be in one of several states. Each state corresponds to a type of consumer behavior in terms of purchasing (privacy sensitivity).*

For this paper, we first focus on the two-state case; the consumer may be Normal or Alerted. Later we will extend this model to multiple consumer states, consumer with coupon dependent response, and unknown initial consumer state cases. The consumer state at time t is denoted by $G_t \in \{\mathsf{Normal}, \mathsf{Alerted}\}$. If a consumer is in Normal state, the consumer is less sensitive to coupons from the retailer in terms of privacy. However, in the Alerted state, the consumer is likely to be more sensitive to coupons offered by the retailer, since it is more cautious about revealing information

to the retailer. The evolution of the consumer state is modeled as a infinite-horizon discrete time Markov chain (Figure 2.1). The consumer starts out in a random initial state unknown to the retailer and the transition of the consumer state is independent of the action of the retailer. A *belief state* is a probability distribution over possible states in which the consumer could be. The belief of the consumer being in Alerted state at time t is denoted by p_t. We define $\lambda_{N,A} = Pr[G_t = \mathsf{Alerted} | G_{t-1} = \mathsf{Normal}]$ to be the transition probability from Normal state to Alerted state and $\lambda_{A,A} = Pr[G_t = \mathsf{Alerted} | G_{t-1} = \mathsf{Alerted}]$ to be the probability of staying in Alerted state when the previous state is also Alerted. The transition matrix $\mathbf{\Lambda}$ of the Markov chain can be written as

$$\mathbf{\Lambda} = \begin{pmatrix} 1 - \lambda_{N,A} & \lambda_{N,A} \\ 1 - \lambda_{A,A} & \lambda_{A,A} \end{pmatrix}. \tag{2.1}$$

We assume the transition probabilities are known to the retailer; this may come from statistical analysis such as a survey of consumer attitudes. The one step transition function, defined by

$$T(p_t) = (1 - p_t)\lambda_{N,A} + p_t \lambda_{A,A}, \tag{2.2}$$

which represents the belief that the consumer is in Alerted state at time $t + 1$ given p_t, the Alerted state belief at time t.

Modelling Assumption 2. ***(State transitions)*** *Consumers have an inertia in that they tend to stay in the same state. Moreover, once consumers feel their privacy is violated, it will take some time for them to come back to* Normal *state.*

The above assumption implies $\lambda_{A,A} \geq 1 - \lambda_{A,A}$, $1 - \lambda_{N,A} \geq \lambda_{N,A}$, and $\lambda_{N,A} \geq 1 - \lambda_{A,A}$. Thus, by combining the above three inequalities, we have $\lambda_{A,A} \geq \lambda_{N,A}$.

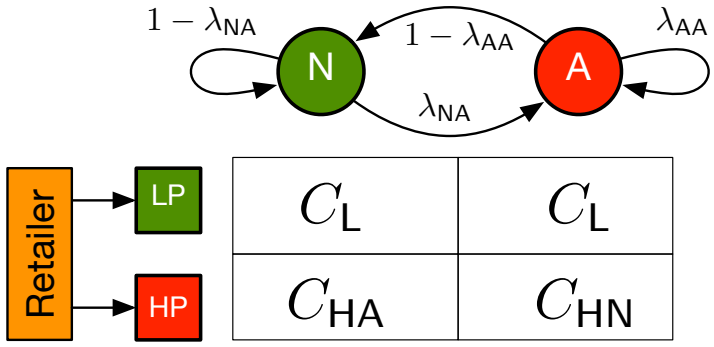

Figure 2.1: Costs to the retailer for offering LP/HP coupons in each privacy sensitive state of the consumer between which the state transitions under a Markov model

Retailer Model

At each time t, the retailer can take an *action* by offering a coupon to the consumer. We define the action at time t to be $u_t \in \{\mathsf{HP}, \mathsf{LP}\}$, where HP denotes offering a high privacy risk coupon (e.g. a targeted coupon) and LP denotes offering a low privacy risk coupon (e.g. a generic coupon). The retailer's utility is modeled by a *cost* (negative revenue) which depends on the consumer's state and the type of coupon being offered. If the retailer offers an LP coupon, it suffers a cost C_L independent of the consumer's state: offering LP coupons does not reveal anything about the state. However, if the retailer offers an HP coupon, then the cost is C_{HN} or C_{HA} depending on whether the consumer's state is Normal or Alerted. Offering an HP (high privacy risk, targeted) coupon to a Normal consumer should incur a low cost (high reward), but offering an HP coupon to an Alerted consumer should incur a high cost (low reward) since an Alerted consumer is privacy-sensitive. Thus, we assume $C_{HN} \leq C_L \leq C_{HA}$.

Under these conditions, the retailer's objective is to choose u_t at each time t to minimize the total cost incurred over the entire time horizon. The HP coupon reveals information about the state through the cost, but is risky if the consumer is alerted, creating a tension between cost minimization and acquiring state information.

Minimum Cost Function

We define $C(p_t, u_t)$ to be the expected cost acquired from an individual consumer at time t where p_t is the probability that the consumer is in $\mathsf{Alerted}$ state and u_t is the retailer's action:

$$C(p_t, u_t) = \begin{cases} C_L & \text{if } u_t = \mathsf{LP} \\ (1 - p_t)C_{HN} + p_t C_{HA} & \text{if } u_t = \mathsf{HP} \end{cases}. \tag{2.3}$$

Since the retailer knows the consumer state from the incurred cost only when an HP coupon is offered, the state of the consumer may not be directly observable to the retailer. Therefore, the problem is actually a Partially Observable Markov Decision Process (POMDP) [127].

We model the cost of violating a consumer's privacy as a short term effect. Thus, we adopt a discounted cost model with discount factor $\beta \in (0, 1)$. We define $\mathcal{P} = \{[0, 1]\}$ and $\mathcal{U} = \{\mathsf{LP}, \mathsf{HP}\}$ to be the belief space and the action space, respectively. At each time t, the retailer has to choose which action u_t to take in order to minimize the expected discounted cost over infinite horizon. A policy π for the retailer is a rule that selects a coupon to offer at each time, i.e. $\pi : \mathcal{P} \to \mathcal{U}$. Thus, given that the belief of the consumer being in $\mathsf{Alerted}$ state at time t is p_t and the policy is π, the infinite-horizon discounted cost starting from t is

$$V_\beta^{\pi,t}(p_t) = \mathbb{E}_\pi \left[\sum_{i=t}^{\infty} \beta^i C(p_i, u_i) | p_t \right], \tag{2.4}$$

where $\mathbb{E}_\pi$ indicates the expectation over the policy π. The objective of the retailer is equivalent to minimizing the discounted cost over all possible policies. Thus, we define the minimum cost function starting from time t over all policies to be

$$V_\beta^t(p_t) = \min_\pi V_\beta^{\pi,t}(p_t) \text{ for all } p_t \in [0, 1]. \tag{2.5}$$

We define $V^t_{\beta,u_t}(p_t)$ to be the infinite-horizon discounted cost starting from t with initial action u_t and p_{t+1} to be the belief of the consumer being in Alerted state at time $t+1$. The minimum cost function $V^t_\beta(p_t)$ satisfies the Bellman equation [127]:

$$V^t_\beta(p_t) = \min_{u_t \in \{\mathsf{HP},\mathsf{LP}\}} \{V^t_{\beta,u_t}(p_t)\}, \tag{2.6}$$

$$V^t_{\beta,u_t}(p_t) = \beta^t C(p_t, u_t) + V^{t+1}_\beta(p_{t+1}|p_t, u_t). \tag{2.7}$$

An optimal policy is *stationary* if it is a deterministic function of states, i.e., the optimal action at a particular state is the optimal action in this state at all times. In the context of our model, the optimal stationary policy is a deterministic and time invariant function mapping $\mathcal{P}$ into $\mathcal{U}$. Since the problem is an infinite-horizon, finite state and finite action POMDP with discounted cost, finding an optimal strategy to this problem is equivalent to solving an associated MDP problem in belief space [128], which is an infinite-horizon discounted MDP with finite action space and uncountably infinite state space. By Theorem 6.3 and its generalization in [129], there exists an optimal stationary policy π^* in the belief space such that starting from time t,

$$V^t_\beta(p_t) = V^{\pi^*,t}_\beta(p_t). \tag{2.8}$$

Thus, only the optimal stationary policy is considered because it is tractable and achieves the same minimum cost as any optimal non-stationary policy.

By (2.6) and (2.7), the minimum cost function evolves as follows. If an HP coupon is offered at time t, the retailer can perfectly infer the consumer state based on the incurred cost. Therefore,

$$V^t_{\beta,\mathsf{HP}}(p_t) = \beta^t C(p_t, \mathsf{HP}) + (1-p_t)V^{t+1}_\beta(\lambda_{N,A}) + p_t V^{t+1}_\beta(\lambda_{A,A}). \tag{2.9}$$

If an LP coupon is offered at time t, the retailer cannot infer the consumer state from the cost since both Normal and Alerted consumer impose the same cost C_L. Hence,

the discounted cost function can be written as

$$\begin{aligned} V^t_{\beta,\mathsf{LP}}(p_t) &= \beta^t C(p_t, \mathsf{LP}) + V^{t+1}_\beta(p_{t+1}) \\ &= \beta^t C_L + V^{t+1}_\beta(T(p_t)). \end{aligned} \tag{2.10}$$

Correspondingly, the minimum cost function is given by

$$V^t_\beta(p_t) = \min\{V^t_{\beta,\mathsf{LP}}(p_t), V^t_{\beta,\mathsf{HP}}(p_t)\}. \tag{2.11}$$

In the sequel, we also consider the following value functions in addition to those defined above. For notational clarity, we define them all here.

- $V^{t\sim k}_\beta(p)$: the minimum cost when the decision horizon starts from t and only spans k stages with initial belief p at time t.
- $V^{t\sim k}_{\beta,u_t}(p)$: the minimum cost when the decision horizon starts from t and only spans k stages with initial belief p and initial action u_t.
- $V_\beta(p)$: the minimum cost function starting from $t = 0$.

We now describe some simple extensions of this basic model.

2.1.2 Consumer with Multi-Level Alerted States

In this section, the case that the consumer has multiple $\mathsf{Alerted}$ states is studied. Without loss of generality, we define $G_t \in \{\mathsf{Normal}, \mathsf{Alerted}_1, \ldots \mathsf{Alerted}_K\}$ to be the consumer state at time t. If the consumer is in $\mathsf{Alerted}_k$ state, it is even more cautious about coupons than in $\mathsf{Alerted}_{k-1}$ state. Beliefs of the consumer being in Normal, $\mathsf{Alerted}_1, \ldots, \mathsf{Alerted}_K$ state at time t are defined by $\bar{\mathbf{p}}_t = (p_{N,t}, p_{A_1,t}, \ldots, p_{A_K,t})^T$. At each time t, the retailer can offer either an HP or an LP coupon. Costs of the retailer when an HP coupon is offered while the state of the consumer is Normal, $\mathsf{Alerted}_1, \ldots, \mathsf{Alerted}_K$ are defined by $\bar{\mathbf{C}} = (C_{HN}, C_{HA_1}, \ldots, C_{HA_K})^T$. If an LP coupon

is offered, no matter in which state, the retailer gets a cost of C_L. We assume that $C_{HA_K} \geq \cdots \geq C_{HA_1} \geq C_L \geq C_{HN}$. The minimum cost function evolves as follows:

$$V_\beta^t(\bar{\mathbf{p}}_t) = \min\{V_{\beta,\mathsf{LP}}^t(\bar{\mathbf{p}}_t), V_{\beta,\mathsf{HP}}^t(\bar{\mathbf{p}}_t)\}, \tag{2.12}$$

where $V_{\beta,\mathsf{LP}}^t(\bar{\mathbf{p}}_t) = \beta^t C_L + V_\beta^{t+1}(\bar{\mathbf{p}}_{t+1})$ and $V_{\beta,\mathsf{HP}}^t(\bar{\mathbf{p}}_t) = \beta^t \bar{\mathbf{p}}_t^T \bar{\mathbf{C}} + V_\beta^{t+1}(\bar{\mathbf{p}}_{t+1})$ represents the cost of offering an LP and an HP coupon, respectively. This model can be generalized to consumer with finitely many states.

2.1.3 Consumer with Coupon Dependent Transition

In the previous formulations, we assume that the consumer's state transition is independent of the retailer's action. A natural extension is the case where the action of the retailer can affect the dynamics of the consumer state evolution (Figure 2.2). Generally, a consumer's reactions to HP and LP coupons are different. For example, a consumer is likely to feel less comfortable when being offered a coupon on medication (HP) than food (LP). Thus, in Section 2.2.2, we assume that the Markov transition probabilities are dependent on the coupon offered with transition matrix given by $\mathbf{\Lambda}_{\mathsf{LP}}(\mathbf{\Lambda}_{\mathsf{HP}})$, where $\mathbf{\Lambda}_{\mathsf{LP}}$ and $\mathbf{\Lambda}_{\mathsf{HP}}$ are defined as:

$$\mathbf{\Lambda}_{\mathsf{LP}} = \begin{pmatrix} 1-\lambda_{N,A} & \lambda_{N,A} \\ 1-\lambda_{A,A} & \lambda_{A,A} \end{pmatrix}, \mathbf{\Lambda}_{\mathsf{HP}} = \begin{pmatrix} 1-\lambda'_{N,A} & \lambda'_{N,A} \\ 1-\lambda'_{A,A} & \lambda'_{A,A} \end{pmatrix}. \tag{2.13}$$

Thus, the minimum cost function is given by (2.11), where $V_{\beta,\mathsf{LP}}^t(p_t) = \beta^t C(p_t, \mathsf{LP}) + V_\beta^{t+1}(T(p_t))$ and $V_{\beta,\mathsf{HP}}^t(p_t) = \beta^t C(p_t, \mathsf{HP}) + (1-p_t)V_\beta^{t+1}(\lambda'_{N,A}) + p_t V_\beta^{t+1}(\lambda'_{A,A})$ denotes the cost function of using an LP coupon and an HP coupon, respectively. $T(p_t)$ is the one step transition given by $T(p_t) = \lambda_{N,A}(1-p_t) + \lambda_{A,A}p_t$.

2.1.4 Policies under Noisy Cost Feedback and Uncertain Initial Belief

Consider a setting in which the feedback regarding the cost may be noisy, e.g., the cost incurred by the consumer's response to the coupon is not deterministic. For each individual consumer, the state transition is independent of the action of the

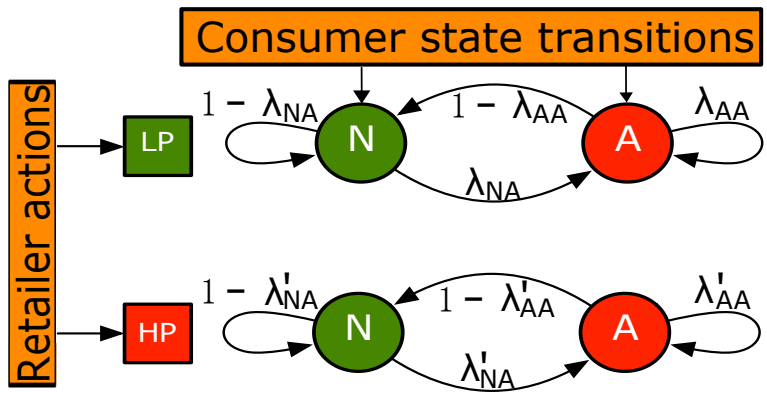

Figure 2.2: Coupon type (HP or LP) dependent Markov state transition model for the consumer.

retailer. For given state G_t and action u_t, define the distribution of observing a cost $C_t = c$ to be $f(c|G_t, u_t)$. In this case, the threshold policy computed using costs might not be optimal. Moreover, if the initial belief is unknown to the retailer, it has to estimate the consumer state before making decision. Thus, we propose some alternative approaches to decide which coupon to offer when those costs are random. A heuristic approach to deal with the randomized cost is to use the threshold τ computed by the mean value of costs. Furthermore, the estimation of consumer belief state p_t or the actual state G_t is updated by the maximum a posteriori rule [130]. After the estimation process, the retailer decides which coupon to offer based on the threshold policy given in Section 2.2.1.

2.1.5 Summary of Main Results

For the problems described in Subsection 2.1.1, 2.1.2, and 2.1.3, given all system parameters, we show the following:

- there exists an optimal stationary solution which has a single threshold property;
- the threshold only depends on the system parameters, i.e., transition probabilities and instantaneous cost associated with each type of coupon.

This means by adopting the optimal policy, the retailer will offer an HP coupon if p_t is less than some threshold and offer an LP if p_t is above the threshold.

For the model described in Subsection 2.1.3, we assume that cost feedbacks are noisy and consumer belief state is unknown to the retailer. For this model:

- we design a heuristic threshold policy when the received costs are noisy.
- a Bayesian estimation approach is proposed to estimate the actual state or the belief state of the consumer when the initial state is unknown to the retailer.

2.2 Optimal Policy for Retailers

For each consumer-retailer interaction model provided in section 3.1, we compute the optimal coupon offering policy for the retailer. We first consider the case in which the retailer knows the consumer statistics. Later, we study consumers with noisy cost feedback.

2.2.1 Optimal Policies with Known Consumer Statistics

In this subsection, we consider the basic formulation as well as the first three extensions. First, we assume that there are only one retailer and one consumer in the system and the state transition of the consumer is independent of the coupon offered. The evolution of the minimum cost function is given in (2.9), (2.10), and (2.11).

Properties of Minimum Cost Function

Lemma 1. *Notice that $V_\beta^{t\sim k}(p)$ is the minimum cost when the decision horizon starts from t and only spans k stages with initial belief p at time t, given a time invariant action set $u_i \in \mathcal{U} = \{\mathsf{LP}, \mathsf{HP}\}$, for any $i = 0, 1, \ldots$, $V_\beta^{t\sim k}(p) = \beta V_\beta^{t-1\sim k}(p)$.*

Proof. By (2.5) and $u_i \in \{\mathsf{LP}, \mathsf{HP}\}$ for any $i = 0, 1, \ldots$.

$$\begin{aligned} V_\beta^{t\sim k}(p) &= \min_\pi \mathbb{E}_\pi \left[\sum_{i=t}^{t+k-1} \beta^i C(p_i, u_i) | p_t = p \right] \\ &= \beta \min_\pi \mathbb{E}_\pi \left[\sum_{i=t-1}^{t+k-2} \beta^i C(p_i, u_i) | p_{t-1} = p \right] \\ &= \beta V_\beta^{t-1\sim k}(p). \end{aligned} \tag{2.14}$$

By using induction on t, we can easily prove $V_\beta^{t\sim k}(p) = \beta V_\beta^{t-1\sim k}(p) = \cdots = \beta^t V_\beta^{0\sim k}(p)$.

□

Lemma 2. *The minimum cost function $V_\beta^t(p)$ is a concave and non-decreasing function of p.*

Proof. We prove these properties by induction. Remember that $V_{\beta,u_t}^{t\sim k}(p)$ is the minimum cost when the decision horizon starts from t and only spans k stages with initial belief p and initial action u_t. For $k = 1$,

$$V_\beta^{t\sim k}(p) = \min\{C_L, (1-p)C_{HN} + pC_{HA}\}, \tag{2.15}$$

which is a concave function of p. For $k = n-1$, assume that $V_\beta^{t\sim k}(p)$ is a concave function. Then, for $k = n$, since $V_\beta^{t\sim n-1}(p)$ is concave and $V_{\beta,\mathsf{LP}}^{t\sim k}(p) = \beta^t C_L + V_\beta^{t+1\sim n-1}(T(p))$, by the definition of concavity and Lemma 1, we can conclude that $V_{\beta,\mathsf{LP}}^{t\sim k}(p)$ is concave. Furthermore, $V_{\beta,\mathsf{HP}}^{t\sim k}(p)$ is an affine function of p, so $V_\beta^{t\sim k}(p) = \min\{V_{\beta,\mathsf{LP}}^{t\sim k}(p), V_{\beta,\mathsf{HP}}^{t\sim k}(p)\}$ is a concave function of p. Taking $k \to \infty, V_\beta^{t\sim k}(p) \to V_\beta^t(p)$, which implies $V_\beta^t(p)$ is a concave function.

Next, we prove the non-decreasing property of the minimum cost function. For $k = 1$, as shown in equation (2.15), it is a non-decreasing function of p. Assume that $V_\beta^{t\sim k}(p)$ is a non-decreasing function for $k = n-1$. For $k = n$, Let $p_1 \geq p_2$,

$$V_{\beta,\mathsf{LP}}^{t\sim k}(p_1) - V_{\beta,\mathsf{LP}}^{t\sim k}(p_2) \tag{2.16}$$

$$
\begin{aligned}
&= \beta(V_\beta^{t\sim n-1}(T(p_1)) - V_\beta^{t\sim n-1}(T(p_2))) \\
&= \beta(V_\beta^{t\sim n-1}((\lambda_{A,A} - \lambda_{N,A})p_1 + \lambda_{N,A}) \\
&\qquad - V_\beta^{t\sim n-1}((\lambda_{A,A} - \lambda_{N,A})p_2 + \lambda_{N,A}))) \\
&\geq 0.
\end{aligned}
$$

By using the same technique, we can prove that given $p_2 - p_1 \leq 0, C_{HN} - C_{HA} \leq 0$ and $V_\beta^{t\sim k-1}(\lambda_{N,A}) - V_\beta^{t\sim k-1}(\lambda_{A,A}) \leq 0$,

$$
V_{\beta,\mathsf{HP}}^{t\sim k}(p_1) - V_{\beta,\mathsf{HP}}^{t\sim k}(p_2) \geq 0. \tag{2.17}
$$

Since $V_\beta^{t\sim k}(p_t) = \min\{V_{\beta,\mathsf{LP}}^{t\sim k}(p), V_{\beta,\mathsf{HP}}^{t\sim k}(p)\}$, it is the minimum of two non-decreasing functions. Therefore, $V_\beta^{t\sim k}(p)$ is non-decreasing. By taking $k \to \infty, V_\beta^{t\sim k}(p) \to V_\beta^t(p)$. Thus, $V_\beta^t(p)$ is a non-decreasing function. □

Lemma 3. *Let Φ_{HP} to be the set of values of p_t for which offering an* HP *coupon is the optimal action at time t. Then, Φ_{HP} is a convex set.*

Proof. Since $\Phi_{\mathsf{HP}} = \{p \in [0,1], V_\beta^t(p) = V_{\beta,\mathsf{HP}}^t(p)\}$, assume that $p_t = ap_{t,1} + (1-a)p_{t,2}$ in which $p_{t,1}, p_{t,2} \in \Phi_{\mathsf{HP}}$ and $a \in [0,1]$, $V_\beta^t(p_t)$ can be written as:

$$
\begin{aligned}
V_\beta^t(p_t) &= V_\beta^t(ap_{t,1} + (1-a)p_{t,2}) \qquad (2.18) \\
&\geq aV_\beta^t(p_{t,1}) + (1-a)V_\beta^t(p_{t,2}) \\
&= aV_{\beta,\mathsf{HP}}^t(p_{t,1}) + (1-a)V_{\beta,\mathsf{HP}}^t(p_{t,2}) \\
&= a[(1-p_{t,1})[\beta^t C_{HN} + \beta V_\beta^t(\lambda_{N,A})] + p_{t,1}[\beta^t C_{HA} + \beta V_\beta^t(\lambda_{A,A})]] \\
&\qquad + (1-a)[(1-p_{t,2})[\beta^t C_{HN} + \beta V_\beta^t(\lambda_{N,A})] + p_{t,2}[\beta^t C_{HA} + \beta V_\beta^t(\lambda_{A,A})]] \\
&= V_{\beta,\mathsf{HP}}^t(ap_{t,1} + (1-a)p_{t,2}).
\end{aligned}
$$

Thus, we have shown that:

$$
V_\beta^t(p_t) \geq V_{\beta,\mathsf{HP}}^t(ap_{t,1} + (1-a)p_{t,1}) = V_{\beta,\mathsf{HP}}^t(p_t). \tag{2.19}
$$

By the definition of $V_\beta^t(p_t)$ in (2.11), $V_\beta^t(p_t) \leq V_{\beta,\mathsf{HP}}^t(p_t)$. Therefore, $V_{\beta,\mathsf{HP}}^t(p_t) = V_\beta^t(p_t)$, which implies Φ_{HP} is convex. □

Optimal Stationary Policy Structure

Theorem 1. *There exists a threshold $\tau \in [0,1]$ such that*

$$\pi^*(p_t) = \begin{cases} \mathsf{LP} & \text{if } \tau \leq p_t \leq 1 \\ \mathsf{HP} & \text{if } 0 \leq p_t \leq \tau \end{cases}. \tag{2.20}$$

is optimal. More precisely, let $\delta \triangleq C_{HA} - C_{HN} + \beta(V_\beta(\lambda_{A,A}) - V_\beta(\lambda_{N,A}))$,

$$\tau = \begin{cases} \frac{C_L - (1-\beta)(C_{HN} + \beta V_\beta(\lambda_{N,A}))}{(1-\beta)\delta} & T(\tau) \geq \tau \\ \frac{C_L + \beta\lambda_{N,A}(C_{HA} + \beta V_\beta(\lambda_{A,A}))}{(1-(\lambda_{A,A}-\lambda_{N,A})\beta)\delta} - \frac{(1-\beta(1-\lambda_{N,A}))(C_{HN}+\beta V_\beta(\lambda_{N,A}))}{(1-(\lambda_{A,A}-\lambda_{N,A})\beta)\delta} & T(\tau) < \tau \end{cases}, \tag{2.21}$$

where for $\lambda_{N,A} \geq \tau$,

$$V_\beta(\lambda_{N,A}) = V_\beta(\lambda_{A,A}) = C_L/(1-\beta) \tag{2.22}$$

and for $\lambda_{N,A} < \tau$,

$$V_\beta(\lambda_{N,A}) = (1-\lambda_{N,A})[C_{HN} + V_\beta^1(\lambda_{N,A})] + \lambda_{N,A}[C_{HA} + V_\beta^1(\lambda_{A,A})], \tag{2.23}$$

$$V_\beta(\lambda_{A,A}) = \min_{n \geq 0}\{G(n)\}, \tag{2.24}$$

where

$$G(n) = \frac{C_L \frac{1-\beta^n}{1-\beta} + \beta^n[\bar{T}^n(\lambda_{A,A})(C_{HN} + C(\lambda_{N,A})) + T^n(\lambda_{A,A})C_{HA}]}{1 - \beta^{n+1}[\bar{T}^n(\lambda_{A,A})\frac{\lambda_{N,A}\beta}{1-(1-\lambda_{N,A})\beta} + T^n(\lambda_{A,A})]} \tag{2.25}$$

$$T^n(\lambda_{A,A}) = \frac{(\lambda_{A,A} - \lambda_{N,A})^{n+1}(1-\lambda_{A,A}) + \lambda_{N,A}}{1 - (\lambda_{A,A} - \lambda_{N,A})} \tag{2.26}$$

$$\bar{T}^n(\lambda_{A,A}) = 1 - T^n(\lambda_{A,A}) \tag{2.27}$$

$$C(\lambda_{N,A}) = \beta\frac{(1-\lambda_{N,A})C_{HN} + \lambda_{N,A}C_{HA}}{1-(1-\lambda_{N,A})\beta}. \tag{2.28}$$

The proof of Theorem 1 is provided in the Appendix A. An immediate consequence of this result is an upper bound on p_t for offering an HP coupon. We define κ to be the ratio between the gain from offering an HP coupon to a Normal consumer and the loss from offering an HP coupon to a consumer whom the retailer thinks is Normal but is actually Alerted. Thus,

$$\kappa = \frac{C_L - C_{HN}}{C_{HA} - C_{HN}}. \tag{2.29}$$

For fixed costs, the threshold can be bounded by the following two Corollaries.

Corollary 1. *If $p_t \leq \kappa$, then it is optimal for the retailer to offer an* HP *coupon.*

Corollary 2. *Fix coupon offering costs and $\lambda_{A,A}$, let $\lambda_1 = \frac{C_L - C_{HN}}{C_{HA} - C_{HN}}$ and λ_2 be the solution of $\frac{\lambda_2}{1-(\lambda_{A,A}-\lambda_2)} = \frac{\beta(C_L - C_{HA})\lambda_2 + C_L - C_{HN}}{(1-\beta)C_{HA} - C_{HN} + \beta C_L}$. When $\lambda_{N,A} \geq \lambda_2$, the threshold τ in the optimal stationary policy can be written as a closed form expression with respect to (w.r.t) $\lambda_{N,A}$: if $\lambda_{N,A} > \lambda_1$,*

$$\tau = \kappa; \tag{2.30}$$

if $\lambda_2 < \lambda_{N,A} < \lambda_1$,

$$\tau = \frac{\beta(C_L - C_{HA})\lambda_{N,A} + C_L - C_{HN}}{(1-\beta)C_{HA} - C_{HN} + \beta C_L}. \tag{2.31}$$

Moreover, if $\lambda_{N,A} < \lambda_2$, τ can be upperbounded by

$$\bar{\tau} = \frac{\lambda_2}{1 - (\lambda_{A,A} - \lambda_2)}. \tag{2.32}$$

A detailed proof of Corollary 1 and 2 are presented in the Appendix B and Appendix C, respectively.

To illustrate the performance of the proposed threshold policy, we compare the discounted cost resulted from the threshold policy with the greedy policy which minimizes the instantaneous cost at each decision epoch as well as a lazy policy which

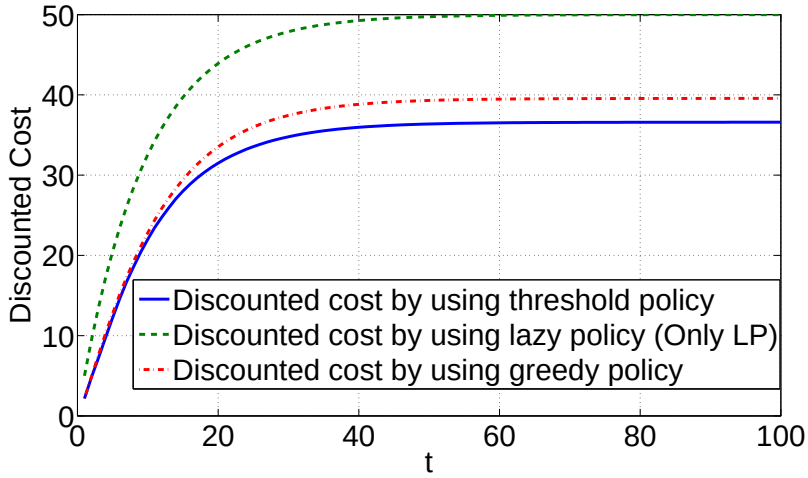

Figure 2.3: Discounted cost resulted by using different decision policies

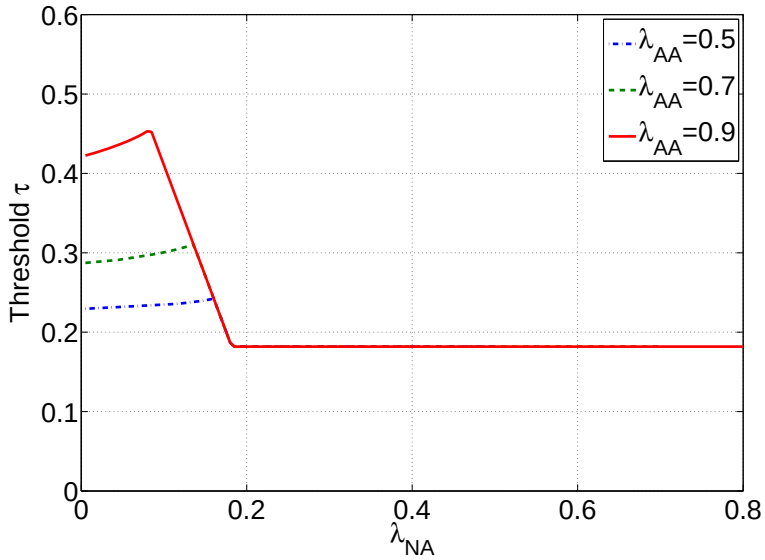

Figure 2.4: Threshold τ vs. $\lambda_{N,A}$ (Parameters: $\beta = 0.9, C_L = 3, C_{HN} = 1, C_{HA} = 12, \kappa = 0.18$).

a retailer only offers LP coupons. We plot the discounted cost averaged over 1000 independent MDPs *w.r.t.* time t for different decision policies in Figure 2.3. The illustration demonstrates that the proposed threshold policy performs better than the greedy policy and the lazy policy.

Figure 2.4 shows the optimal threshold policy *w.r.t* $\lambda_{N,A}$ for three fixed choices of $\lambda_{A,A}$. It can be seen that the threshold is increasing when $\lambda_{N,A}$ is small, this is

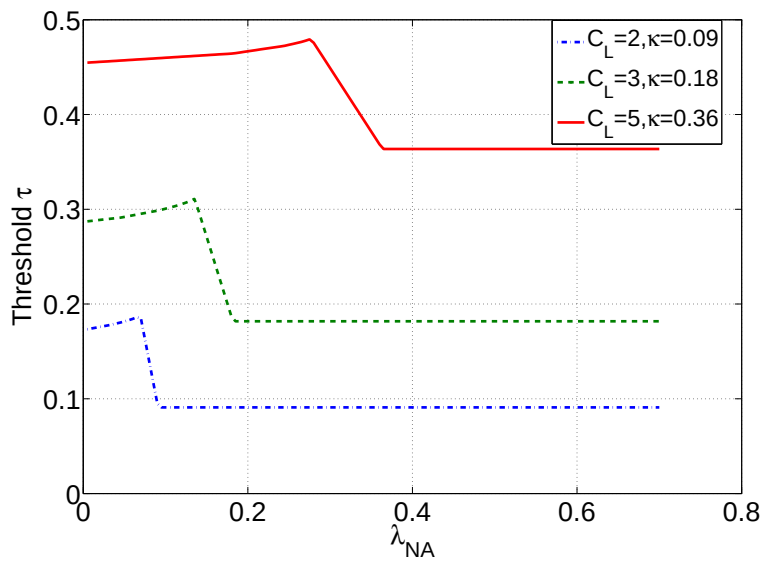

Figure 2.5: Threshold τ vs. $\lambda_{N,A}$ (Parameters: $\lambda_{A,A} = 0.7, \beta = 0.9, C_{HN} = 1, C_{HA} = 12$).

because for a small $\lambda_{N,A}$, the consumers is less likely to transition from Normal to Alerted. Therefore, the retailer tends to offer an HP coupon to the consumer. When $\lambda_{N,A}$ gets larger, the consumer is more likely to transition from Normal to Alerted. Thus, the retailer tends to play conservatively by decreasing the threshold for offering an LP coupon. When $\lambda_{N,A}$ is greater than κ, the retailer will just use κ to be the threshold for offering an HP coupon. One can also observe that with increasing $\lambda_{A,A}$, the threshold τ decreases. On the other hand, for fixed C_{HN} and C_{HA}, Figure 2.5 shows that the threshold τ increases as the cost of offering an LP coupon increases, making it more desirable to take a risk and offer an HP coupon.

The relationship between the discount factor β and the threshold τ as functions of transition probabilities is shown in Figure 2.6 and 2.7. It can be seen in Figure 2.6 that the threshold increases as β increases. This is because when β is small, the retailer values the present rewards more than future rewards. Therefore, the retailer tends to play conservatively so that it will not "creep out" the consumer in the present. Figure 2.7 shows that the threshold is high when $\lambda_{A,A}$ is large or $\lambda_{N,A}$ is small. A

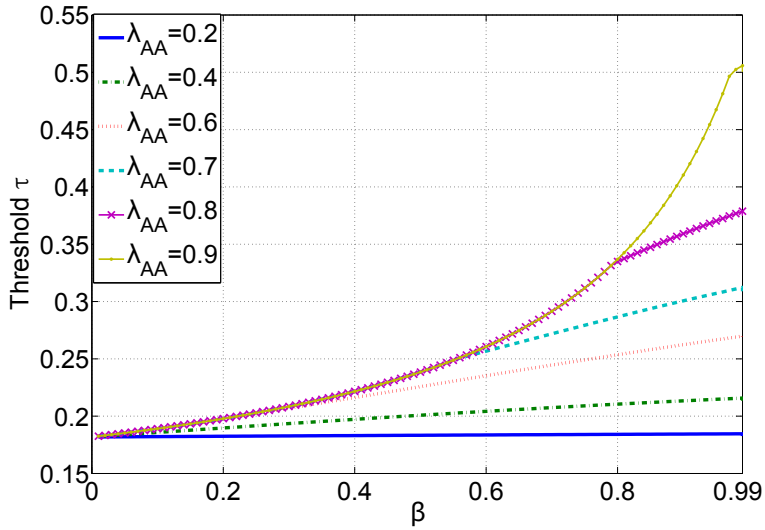

Figure 2.6: Threshold τ vs. β for different values of $\lambda_{A,A}$ (Parameters: $\lambda_{N,A} = 0.1, C_L = 3, C_{HN} = 1, C_{HA} = 12, \kappa = 0.18$).

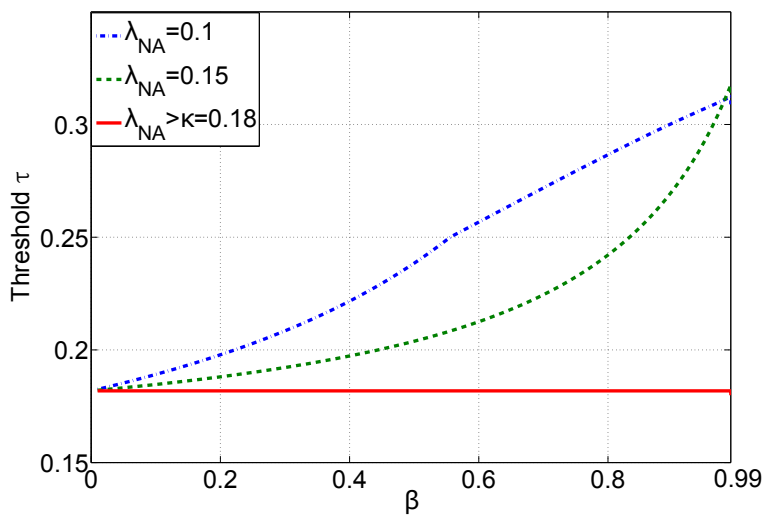

Figure 2.7: Threshold τ vs. β for different values of $\lambda_{N,A}$ (Parameters: $\lambda_{A,A} = 0.7, C_L = 3, C_{HN} = 1, C_{HA} = 12$).

high $\lambda_{A,A}$ value indicates that a consumer is more likely to remain in Alerted state. The retailer is willing to play aggressively since once the consumer is in alerted state, it can take a very long time to transition back to Normal state. A low $\lambda_{N,A}$ value implies that the consumer is not very privacy sensitive. Thus, the retailer tends to offer HP coupons to reduce cost. One can also observe in Figure 2.7 that the threshold

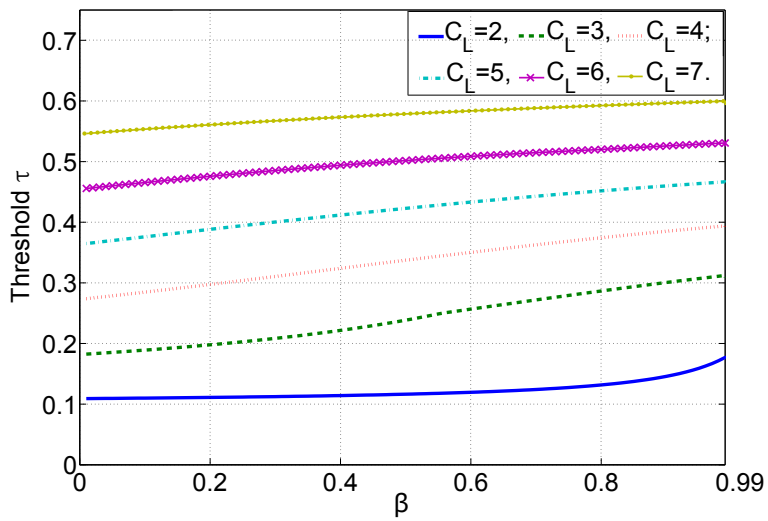

Figure 2.8: Threshold τ vs. β for different values of C_L (Parameters: $\lambda_{N,A} = 0.1, \lambda_{A,A} = 0.9, C_{HN} = 1, C_{HA} = 12$).

τ equals to κ after $\lambda_{N,A}$ exceeds the ratio κ. This is consistent with results shown in Figure 2.4 and 2.5.

The effect of an LP coupon cost on the threshold for different discount factors is plotted in Figure 2.8. It can be seen that a higher C_L will increase the threshold because the retailer is more likely to offer an HP coupon when the cost of offering an LP coupon is high.

Consumer with Multi-Level Alerted States

In this section, we study the case that the consumer has multiple Alerted states. Without loss of generality, we define the transition matrix to be

$$\mathbf{\Lambda} = \begin{pmatrix} \lambda_{N,N} & \lambda_{N,A_1} & \dots & \lambda_{N,A_K} \\ \lambda_{A_1,N} & \lambda_{A_1,A_1} & \dots & \lambda_{A_1,A_K} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{A_K,N} & \lambda_{A_K,A_1} & \dots & \lambda_{A_K,A_K} \end{pmatrix} \tag{2.33}$$

and $\bar{\mathbf{e}}_i$ to be the ith row of $\mathbf{\Lambda}$. The expected cost at time t, given belief $\bar{\mathbf{p}}_t$ and action u_t, has the following expression:

$$C(\bar{\mathbf{p}}_t, u_t) = \begin{cases} C_L & \text{if } u_t = \mathsf{LP} \\ \bar{\mathbf{p}}_t^T \bar{\mathbf{C}} & \text{if } u_t = \mathsf{HP} \end{cases}. \tag{2.34}$$

Assuming that the retailer has perfect information about the belief states, the cost function evolves as follows. By using an LP coupon at time t,

$$V_{\beta,\mathsf{LP}}^t(\bar{\mathbf{p}}_t) = \beta^t C_L + V_\beta^{t+1}(\bar{\mathbf{p}}_{t+1}) = \beta^t C_L + V_\beta^{t+1}(T(\bar{\mathbf{p}}_t)), \tag{2.35}$$

where $T(\bar{\mathbf{p}}_t) = \bar{\mathbf{p}}_t^T \mathbf{\Lambda}$ is the Markov transition operator generalizing (2.2). By using an HP coupon at time t,

$$V_{\beta,\mathsf{HP}}^t(\bar{\mathbf{p}}_t) = \beta^t \bar{\mathbf{p}}_t^T \bar{\mathbf{C}} + V_\beta^{t+1}(\bar{\mathbf{p}}_{t+1}) = \beta^t \bar{\mathbf{p}}_t^T \bar{\mathbf{C}} + \bar{\mathbf{p}}_t^T \begin{pmatrix} V_\beta^{t+1}(\bar{\mathbf{e}}_1) \\ V_\beta^{t+1}(\bar{\mathbf{e}}_2) \\ \vdots \\ V_\beta^{t+1}(\bar{\mathbf{e}}_{K+1}) \end{pmatrix}. \tag{2.36}$$

Therefore, by (2.11), we have $V_\beta^t(\bar{\mathbf{p}}_t) = \min\{V_{\beta,\mathsf{LP}}^t(\bar{\mathbf{p}}_t), V_{\beta,\mathsf{HP}}^t(\bar{\mathbf{p}}_t)\}$.

In this problem, since the instantaneous costs are nondecreasing with the state when the action is fixed and the evolution of belief state is the same for both LP and HP, the existence of an optimal stationary policy with threshold property is guaranteed by Proposition 2 in [131]. The optimal stationary policy for a three-state consumer model is illustrated in Figure 2.9. For fixed costs, the plot shows the partition of the belief space based on the optimal actions and reveals that offering an HP coupon is optimal when $p_{N,t}$, the belief of the consumer being in Normal state, is high.

2.2.2 Consumers with Coupon Dependent Transitions

Generally, consumers' reaction to HP and LP coupons are different. To be more specific, a consumer is likely to feel less comfortable when being offered a coupon

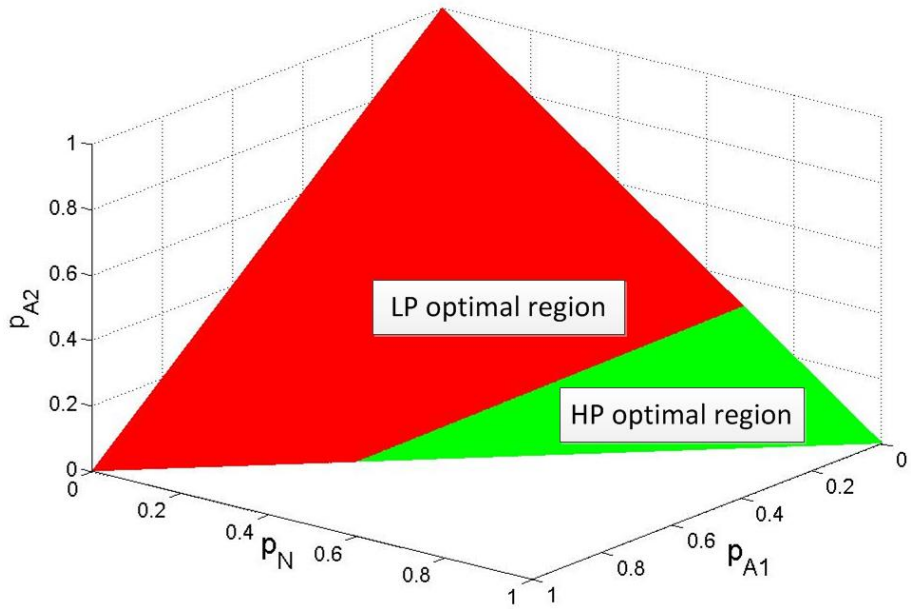

Figure 2.9: Example of the optimal policy region for three-state consumer. (Parameters: $\lambda_{\mathsf{N,N}} = 0.7, \lambda_{\mathsf{N},A1} = 0.2, \lambda_{\mathsf{N,A2}} = 0.1; \lambda_{\mathsf{A1,N}} = 0.2, \lambda_{\mathsf{A1,A1}} = 0.5, \lambda_{\mathsf{A1,A2}} = 0.3; \lambda_{\mathsf{A2,N}} = 0.1, \lambda_{\mathsf{A2,A1}} = 0.2, \lambda_{\mathsf{A2,A2}} = 0.7; \beta = 0.9, C_{\mathsf{L}} = 7, C_{\mathsf{HN}} = 1, C_{\mathsf{HA1}} = 10, C_{\mathsf{HA2}} = 20$).

on medication (HP) than food (LP). Thus, we assume that the Markov transition probabilities are dependent on the coupon offered. Let p_t denote the belief of a consumer being in the Alerted state at time t.

As shown in Figure 2.2, by offering an LP coupon, the state transition follows the Markov chain

$$\mathbf{\Lambda}_{\mathsf{LP}} = \begin{pmatrix} 1-\lambda_{N,A} & \lambda_{N,A} \\ 1-\lambda_{A,A} & \lambda_{A,A} \end{pmatrix}. \tag{2.37}$$

Otherwise, the state transition follows

$$\mathbf{\Lambda}_{\mathsf{HP}} = \begin{pmatrix} 1-\lambda'_{N,A} & \lambda'_{N,A} \\ 1-\lambda'_{A,A} & \lambda'_{A,A} \end{pmatrix}. \tag{2.38}$$

According to the model in Section 3.1, $\lambda_{A,A} > \lambda_{N,A}, \lambda'_{A,A} > \lambda'_{N,A}$. Moreover, we assume that offering an HP coupon will increase the probability of transition to or staying at Alerted state. Therefore, $\lambda'_{A,A} > \lambda_{A,A}$ and $\lambda'_{N,A} > \lambda_{N,A}$. The minimum

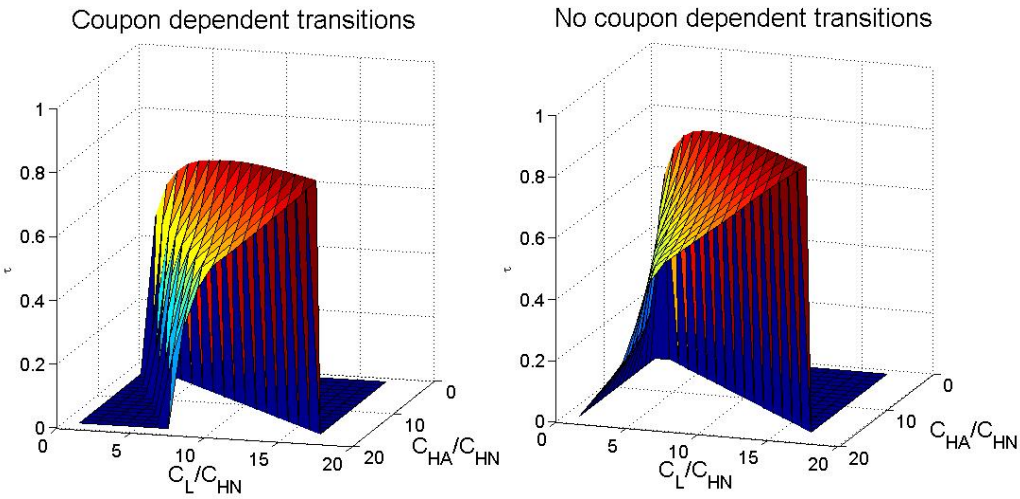

Figure 2.10: Optimal policy threshold for consumer with/without coupon dependent transition probabilities. (Parameters: $\lambda_{N,A} = 0.2, \lambda_{A,A} = 0.8, \lambda'_{N,A} = 0.5, \lambda'_{A,A} = 0.9, \beta = 0.9$).

cost function evolves as follows: for an HP coupon offered at time t, we have

$$V^t_{\beta,\mathsf{HP}}(p_t) = \beta^t C(p_t, \mathsf{HP}) + (1 - p_t)V^{t+1}_\beta(\lambda'_{N,A}) + p_t V^{t+1}_\beta(\lambda'_{A,A}).$$

Otherwise,

$$V^t_{\beta,\mathsf{LP}}(p_t) = \beta^t C_L + V^{t+1}_\beta(p_{t+1}) = \beta^t C_L + V^{t+1}_\beta(T(p_t)),$$

where $T(p_t) = \lambda_{N,A}(1 - p_t) + \lambda_{A,A}p_t$ is the one step transition defined in Section 3.1. In this case, the transition probability is just a deterministic function of the retailer action. Thus, finding an optimal strategy to this problem is equivalent to solving an associated MDP problem in belief space. Furthermore, Theorem 6.3 and its generalization in Ross (1992) still hold since the transition probability is a function of the action. Therefore, there exists an optimal stationary policy π^* in the belief space which minimizes the infinite horizon discounted cost.

Theorem 2. *Given action dependent transition matrices $\mathbf{\Lambda}_{\mathsf{LP}}$ and $\mathbf{\Lambda}_{\mathsf{HP}}$, the optimal stationary policy has threshold structure.*

The proof of Theorem 2 is provided in the Appendix D.

Figure 2.10 shows the effect of costs on the threshold τ. We observe that for a fixed C_L and C_{HA} pair, the threshold for LP coupons for consumers in this model is lower than our original model without coupon-dependent transition probabilities. The retailer can only offer an LP coupon with certain combination of costs; we call this the LP-only region. One can also see that the LP-only region for the coupon-independent transition case is smaller than that for the coupon-dependent transition case since for the latter, the likelihood of being in an $\mathsf{Alerted}$ state is higher for the same costs.

2.2.3 Policies under Noisy Cost Feedback and Uncertain Initial Belief

In the previous sections, if the retailer offers an HP coupon at time t, then it could learn the state of the consumer at time t based on whether the received cost was C_{HN} or C_{HA}. However, in reality, the cost observed by the retailer may not be deterministic. Thus, in this section, we study the case in which the received costs are modeled as a random variable. If the cost feedback is random, then the retailer may not be able to infer the consumer's state exactly. We describe policy heuristics for this setting that perform Bayesian estimation of the quantity p_t used in the threshold policy earlier. This approach is also useful when the initial value p_0 is not known to the retailer.

We model the noisy cost feedback by assuming the received cost C_t is random. The distribution of C_t is given by a conditional probability density $f(c|G_t, u_t)$ on a bounded subset of $\mathbb{R}$, where G_t is the state of the consumer and u_t is the action taken by the retailer at time t. To match the previous model, we further take $f(c|G_t = \mathsf{Alerted}, u_t = \mathsf{LP}) = f(c|G_t = \mathsf{Normal}, u_t = \mathsf{LP})$ to indicate that the received cost conveys no information about the state under an LP coupon. Let $f(c|u_t = \mathsf{LP}) = f(c|G_t = \mathsf{Alerted}, u_t = \mathsf{LP})$. For a given value $p_t = p$, define the likelihood of observing

a cost $C_t = c$ under the two coupons:

$$\ell(c|\mathsf{LP}, p) = f(c|\mathsf{Alerted}, \mathsf{LP}) \tag{2.39}$$

$$\ell(c|\mathsf{HP}, p) = f(c|\mathsf{Normal}, \mathsf{HP})(1-p) + f(c|\mathsf{Alerted}, \mathsf{HP})p \tag{2.40}$$

These likelihoods will be useful in defining the two estimators.

In both approaches in this section the retailer computes an estimate $\hat{p}_t$ of the probability p_t that $G_t = \mathsf{Alerted}$. It then uses (2.20) to decide which coupon to offer at time t by comparing $\hat{p}_t$ to a version of the threshold in (2.21). Define $\mathcal{C}_{\mathcal{L}}, \mathcal{C}_{\mathcal{HN}}$, and $\mathcal{C}_{\mathcal{HA}}$ to be the feasible cost sets $\{c : f(c|\mathsf{LP}) > 0\}$, $\{c : f(c|\mathsf{Alerted}, \mathsf{HP}) > 0\}$, and $\{c : f(c|\mathsf{Normal}, \mathsf{HP}) > 0\}$, respectively. Since τ involves costs C_L, C_{HN} and C_{HA}, there are several ways to compute an approximate threshold under the cost uncertainty.

Firstly, we can set C_L, C_{HN} and C_{HA} to be the expected costs:

$$C_L = \int_{\mathbb{R}} c f(c|\mathsf{LP}) dc \tag{2.41}$$

$$C_{HN} = \int_{\mathbb{R}} c f(c|\mathsf{Normal}, \mathsf{HP}) dc \tag{2.42}$$

$$C_{HA} = \int_{\mathbb{R}} c f(c|\mathsf{Alerted}, \mathsf{HP}) dc. \tag{2.43}$$

Plugging these into (2.21) gives the mean threshold τ_{avg}. Since τ is monotonically increasing in C_L and C_{HA} and monotonically decreasing in C_{HN}, we can compute an upper bound on τ by setting $C_L = \max\{c : c \in \mathcal{C}_{\mathcal{L}}\}$, $C_{HA} = \max\{c : c \in \mathcal{C}_{\mathcal{HA}}\}$, and $C_{HN} = \max\{c : c \in \mathcal{C}_{\mathcal{HN}}\}$. These values give the upper bound threshold τ_{max}. Similarly, by setting C_L and C_{HA} to the lower bounds on the support and C_{HN} to the upper bound, we obtain a lower bound threshold τ_{min}. Finally, we computed a robust version of threshold τ_{R} as $\tau_{\mathsf{R}} = \{\tau : \max_{C_L, C_{HN}, C_{HA}} \{\min_{\pi(p_t)} V_\beta^t(p_t)\}\}$, where $(C_L, C_{HN}, C_{HA}) \in \mathcal{C}_{\mathcal{L}} \times \mathcal{C}_{\mathcal{HN}} \times \mathcal{C}_{\mathcal{HA}}$. This threshold policy is the largest (cost case)

threshold over all possible combination of costs. Thus, it gives the $max - min$ value of the total discounted cost.

Estimation of the Consumer State

In the previous model, if $u_t = \mathsf{HP}$ the retailer could infer G_t based on C_t, so p_{t+1} is given by the state transitions of the Markov chain. With noisy costs this exact inference is no longer possible. A simple heuristic for the retailer is to try to infer G_t based on the random cost C_t, compute an estimate of p_t, and then use the previous strategy.

At time $t = 1$, given an initial p_0 we estimate $\hat{p}_1 = T(p_0)$. The retailer then applies the threshold policy (2.20) with input $\hat{p}_1$ to offer a coupon. For times $t = 2, 3, \ldots$ the retailer treats the estimate $\hat{p}_{t-1}$ as an estimate of the probability that $G_{t-1} = \mathsf{Alerted}$. If $u_{t-1} = \mathsf{LP}$, then the retailer sets $\hat{p}_t = T(\hat{p}_{t-1})$. If $u_{t-1} = \mathsf{HP}$ then the retailer uses a maximum a posteriori probability (MAP) detection rule to estimate the state G_{t-1} based on the received cost C_{t-1}. That is, it sets $\hat{G}_{t-1} = \mathsf{Normal}$ if

$$\frac{f(C_{t-1}|\mathsf{Normal}, \mathsf{HP})(1 - \hat{p}_{t-1})}{f(C_{t-1}|\mathsf{Alerted}, \mathsf{HP})\hat{p}_{t-1}} > 1 \tag{2.44}$$

and $\hat{G}_{t-1} = \mathsf{Alerted}$ otherwise, where C_{t-1} is the received cost at time $t - 1$. It then uses the following estimate p_t at time t:

$$\hat{p}_t = \begin{cases} \lambda_{N,A} & \text{if } \hat{G}_t = \mathsf{Normal} \\ \lambda_{A,A} & \text{if } \hat{G}_t = \mathsf{Alerted} \end{cases}. \tag{2.45}$$

Essentially, the retailer uses MAP estimation to infer G_{t-1} after receiving the cost C_{t-1} from the action $u_{t-1} = \mathsf{HP}$. If the densities $f(c|\mathsf{Normal}, \mathsf{HP})$ and $f(c|\mathsf{Alerted}, \mathsf{HP})$ have disjoint supports, then the inference of G_{t-1} is error free, so $\hat{G}_{t-1} = G_{t-1}$ and the estimate $\hat{p}_t$ is correct. Figure 2.11 shows the discounted cost as a function of

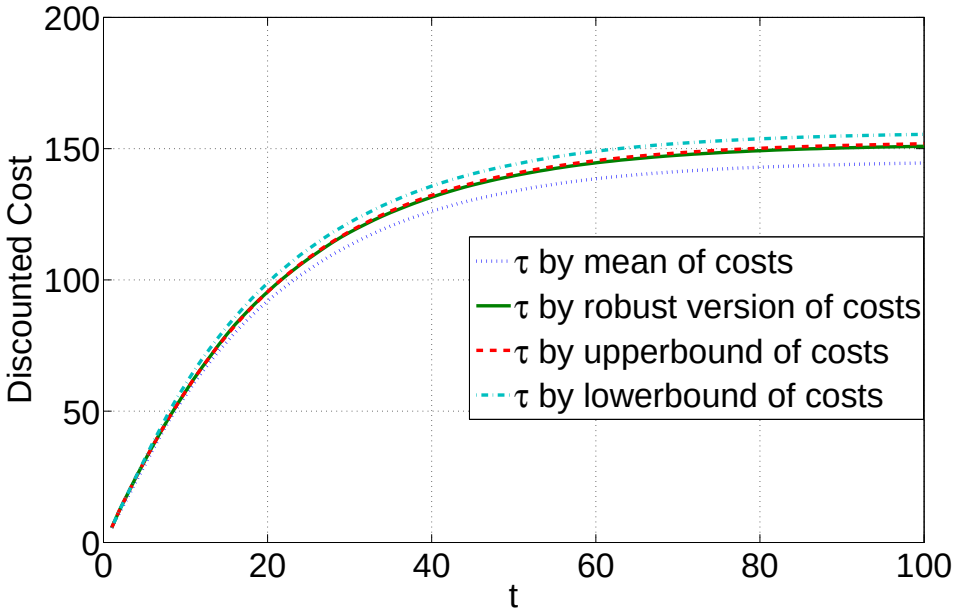

Figure 2.11: Temporal discounted costs for different heuristics on computing thresholds. (Parameters: $\lambda_{N,A} = 0.2$, $\lambda_{A,A} = 0.8$, $p_0 = 0.2$,$\beta = 0.95$, $f(c|\mathsf{LP}) = \mathsf{Unif}[6, 10]$, $f(c|\mathsf{Normal}, \mathsf{HP}) = \mathsf{Unif}[0.2, 5.8]$, and $f(c|\mathsf{Alerted}, \mathsf{HP}) = \mathsf{Unif}[12, 20]$). The discounted cost is averaged over 1000 independent runs.

time for some different variants of the threshold in (2.21). In this example the cost distributions are uniformly distributed in disjoint intervals. The plot shows that the mean threshold yields a total discounted cost that is slightly less than the upper and lower bound thresholds.

Bayesian Estimation of State Probabilities

In the previous approach, the retailer estimates the underlying state and then uses this to form an estimate of the probability p_t that $G_t = \mathsf{Alerted}$. A different approach is to form a Bayes estimate of p_t: the retailer computes a probability distribution on $[0, 1]$ representing its uncertainty about p_t. To choose an action u_t it can use a point estimate of p_t to plug into (2.20) with one of the thresholds described before.

In this formulation, the estimator of p_t is a probability distribution. Let $q_{t-1}(p)$ be the estimator of p_{t-1}. The retailer treats this as a prior distribution. Upon receiving

the cost C_{t-1} it computes a posterior estimate on p_{t-1} using Bayes rule. If $u_{t-1} = \mathsf{HP}$, it sets

$$q_{t-1}(p|C_{t-1}) = \frac{\ell(C_{t-1}|\mathsf{HP}, p)q_{t-1}(p)}{\int_0^1 \ell(C_{t-1}|\mathsf{HP}, p')q_{t-1}(p')dp'}. \tag{2.46}$$

If $u_{t-1} = \mathsf{LP}$ then from (2.39) we can see that $\ell(C_{t-1}|\mathsf{LP}, p)$ does not depend on p, so the posterior $q_{t-1}(p|C_{t-1}) = q_{t-1}(p)$ in this case. Given the posterior estimate $q_{t-1}(p|C_{t-1})$, the retailer then evolves the state distribution through the Markov chain governing the state to form the prior distribution $q_t(p)$ for estimating p_t at time t. That is, if P_{t-1} is a random variable with distribution $q_{t-1}(p|C_{t-1})$, then $q_t(p)$ is the distribution of $T(P_{t-1})$. Let $Q_{t-1}(p|C_{t-1}) = \int_0^p q_{t-1}(p'|C_{t-1})dp'$ be the cumulative distribution function of P_{t-1}. Then

$$\mathbb{P}\left(T(P_{t-1}) \leq p\right) = \mathbb{P}\left(P_{t-1} \leq \frac{p - \lambda_{N,A}}{\lambda_{A,A} - \lambda_{N,A}}\right) = Q_{t-1}\left(\frac{p - \lambda_{N,A}}{\lambda_{A,A} - \lambda_{N,A}}|C_{t-1}\right), \tag{2.47}$$

$$q_t(p) = \frac{1}{\lambda_{A,A} - \lambda_{N,A}} q_{t-1}\left(\frac{p - \lambda_{N,A}}{\lambda_{A,A} - \lambda_{N,A}}|C_{t-1}\right). \tag{2.48}$$

The retailer then uses $q_t(p)$ to form a point estimate $\hat{p}_t$ of p_t suitable for applying the threshold policy in (2.20) and (2.21). We consider two such point estimates which we call the mean and max estimators, respectively:

$$\hat{p}_{t,\mathsf{mean}} = \int_0^1 p q_t(p)dp, \tag{2.49}$$

$$\hat{p}_{t,\mathsf{MAP}} = argmax_{p\in[0,1]} q_t(p). \tag{2.50}$$

Figure 2.12 shows the discounted cost versus time for uniformly distributed costs with overlapping support. The decision is made by following the optimal stationary policy computed by the mean threshold in Section 2.2.3. We illustrate the result for four algorithms: the solid curve and the dash-dot curve are the MAP and mean strategies described above, respectively; the dashed curve is a policy in which costs

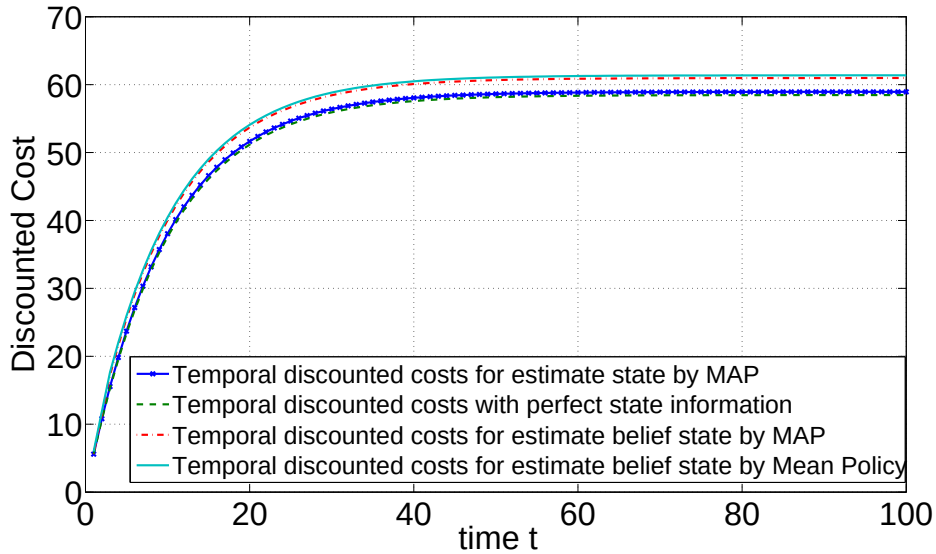

Figure 2.12: Temporal discounted costs for different estimation mechanisms. (Parameters: $\lambda_{N,A} = 0.2$,$\lambda_{A,A} = 0.8$, $p_0 = 0.2$, $\hat{p}_0 = 0.1$, $\beta = 0.9$, $f(c|\mathsf{LP}) = \mathsf{Unif}[3, 9]$, $f(c|\mathsf{Normal}, \mathsf{HP}) = \mathsf{Unif}[0.25, 7.75]$, $f(c|\mathsf{Alerted}, \mathsf{HP}) = \mathsf{Unif}[6, 18]$). The discounted cost is averaged over 1000 independent runs.

are random but the algorithm is given side information about G_t after choosing $u_t = \mathsf{HP}$ (perfect state information); finally, the curve with cross is the MAP estimate of actual state G_t described in Section 2.2.3. In this example, as one can expect, decision making with perfect state information has the minimum discounted cost. MAP estimation of G_t results in an 0.82% increase in total discounted cost compared to the case in which the retailer receives perfect information about consumer state. However, the MAP and mean policy to estimate belief state p_t only have 2.9% and 4.29% increase, respectively. Thus, the MAP for estimating belief performs slightly better than the Mean policy. Effectively, the lack of initial belief knowledge does not affect the discounted cost very much on average. This is because offering an HP coupon allows the retailer to learn the actual state from the cost feedback, thus, reset the belief state.

Chapter 3

INCENTIVE MECHANISMS FOR PRIVACY SENSITIVE ELECTRICITY CONSUMERS WITH ALTERNATIVE ENERGY SOURCES

3.1 System Model

Consider a distribution system with M privacy-sensitive consumers as shown in Figure 3.1. Each consumer has an installed smart meter, an alternative energy source (PV), and an energy storage device (battery). A consumer can determine how much electricity it needs to consume from the grid at any time intelligently using the battery and PV as alternative energy sources to simultaneously obtain certain level of privacy and reduce the cost of electricity.

3.1.1 Consumer Model

We consider a discrete-time model with a set of consumers $\mathcal{H} = \{1, 2, ..., M\}$ in one electricity provider's network. For consumer $i \in \mathcal{H}$, let $D_{i,t}$ be the inelastic net electricity demand from appliances that belong to consumer i at time t. With PV and battery installed, consumer i meets this demand using both alternative energy sources and the grid while maximizing use of its alternative energy sources. We denote

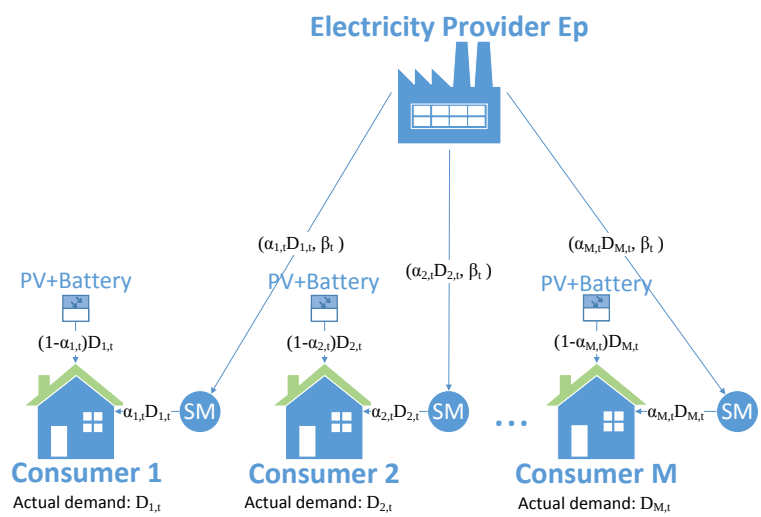

Figure 3.1: Consumer-electricity provider interaction diagram

$\alpha_{i,t} \in \mathcal{A}_{i,t} \triangleq \{\alpha_{i,t}^1, ..., \alpha_{i,t}^K\}$ to be the fraction of the net electricity demand $D_{i,t}$ that is consumed directly from the grid. We assume that without any incentive from the electricity provider, consumer i uses an intelligent privacy preserving algorithm to compute the amount of energy consumption from the grid to balance its net energy demands consistently with its privacy requirements. Such an algorithm will require consumer i to consume $\alpha_{i,t}^0 D_{i,t}$ to ensure maximal use of its alternative energy sources and consume from the grid only when needed. Specifically, consuming either more or less than this fraction from the grid can cause loss in privacy. However, this privacy preserving consumption may not meet the supply and demand requirements of the electricity provider. At the beginning of time t, consumer i reports its privacy preserving consumption $\alpha_{i,t}^0 D_{i,t}$ to the electricity provider. To ensure a desired total consumption X_t, after receiving all consumption information from consumers, the electricity provider decides $\beta_t \in \mathcal{B}_t \triangleq \{\beta_t^1, ..., \beta_t^J\}$ to be the incentive price for each KW of deviation from each consumer's privacy preserving consumption. Meanwhile, each consumer i adjusts its actual consumption to $\alpha_{i,t} D_{i,t}$.

When the electricity provider uses incentives to compensate consumers for changing electricity consumption behaviors, each privacy-sensitive consumer will match its privacy leakage to a monetized valuation of privacy loss. For fixed value of $D_{i,t}$, we assume that there exists a convex and bounded function $f(\alpha_{i,t} - \alpha_{i,t}^0, D_{i,t})$ with minimum value located at $\alpha_{i,t}^0$. This function maps consumer i's deviation from its privacy preserving consumption that it consumes directly from the grid to a monetized loss of its privacy leakage. This assumption is motivated by the observation that for consumers with access to alternative energy sources and a privacy preserving algorithm, $\alpha_{i,t}^0 D_{i,t}$ gives the minimum amount of privacy leakage and the valuation of privacy is increasingly higher as electricity consumption deviates increasingly from the optimal consumption profile computed by the privacy preserving algorithm. De-

fine $\delta_{i,t} \triangleq (\alpha_{i,t} - \alpha^0_{i,t})$ to be the change in fractional consumption when consumer i is willing to consume $\alpha_{i,t}D_{i,t}$ from the electricity provider. For consuming $\alpha_{i,t}D_{i,t}$ from the electricity provider, consumer i suffers a loss $f_{i,t}(\delta_{i,t}, D_{i,t})$ due to the privacy leakage resulting from deviating from $\alpha^0_{i,t}$. This compensation that consumer i receives for compromising its privacy is given by

$$U_{i,t}(\alpha_{i,t}, \beta_t, D_{i,t}) = \begin{cases} \beta_t(\alpha_{i,t} - \alpha^0_{i,t})D_{i,t} \text{ if } X_t \geq \sum\limits_{i \in \mathcal{H}} \alpha^0_{i,t}D_{i,t} \\ -\beta_t(\alpha_{i,t} - \alpha^0_{i,t})D_{i,t} \text{ otherwise} \end{cases}. \tag{3.1}$$

The incentive mechanism in (3.1) indicates that when the aggregated privacy preserving consumption over all consumers is lower (higher) than the amount required by the electricity provider, it is in the consumer's interest to exploit the incentives and increase (decrease) its consumption to meet the needs of the electricity provider. Our model also assumes that the electricity provider can penalize the consumer for not adhering to the consumption requirement of the electricity provider via negative incentives for each KW of deviation from $\alpha^0_{i,t}D_{i,t}$.

We denote $R_{i,t}(\alpha_{i,t}, \beta_t) = U_{i,t}(\alpha_{i,t}, \beta_t, D_{i,t}) - f_{i,t}(\alpha_{i,t} - \alpha^0_{i,t}, D_{i,t})$ to be the reward of consumer i for fixed demand $D_{i,t}$. Recall that for mixed strategies, each consumer and the electricity provider can choose from some convex combinations over their pure strategies. Let $\boldsymbol{p}_t = (\boldsymbol{p}_{1,t}, ..., \boldsymbol{p}_{M,t})$ be the vector of mixed strategies for all consumers. We also define $\boldsymbol{p}_{i,t} \in \boldsymbol{\mathcal{P}}_{i,t}$ to be the vector of probability distribution over pure strategies $\alpha_{i,t}$ and $\boldsymbol{\mathcal{P}}_t \triangleq \boldsymbol{\mathcal{P}}_{1,t} \times \cdots \boldsymbol{\mathcal{P}}_{M,t}$ to be the feasible set of $\boldsymbol{p}_t$. Furthermore, we denote $\boldsymbol{q}_t(\beta_t)$ to be the mixed strategy for the electricity provider whose feasible set of mixed strategies is defined to be $\boldsymbol{\mathcal{Q}}_t$.

For a given mixed strategy, the expected reward of each consumer i is given by

$$R^E_{i,t}(\boldsymbol{p}_{i,t}, \boldsymbol{q}_t) = \sum_{\alpha_{i,t} \in \mathcal{A}_{i,t}, \beta_t \in \mathcal{B}_t} (p_{i,t}(\alpha_{i,t})q_t(\beta_t))R_{i,t}(\alpha_{i,t}, \beta_t). \tag{3.2}$$

To maximize its expected reward at time t, consumer i solves the following optimization problem

$$\max_{\boldsymbol{p}_{i,t} \in \boldsymbol{\mathcal{P}}_{i,t}} \quad R^E_{i,t}(\boldsymbol{p}_{i,t}, \boldsymbol{q}_t). \tag{3.3}$$

This optimization captures both the economic benefit to the consumer and the monetized consequence of privacy leakage due to consuming energy from the grid.

3.1.2 Electricity Provider Model

Each smart meter installed by the electricity provider samples electricity consumption at a fixed rate and transmits the actual consumption data from each consumer ($\alpha_{i,t} D_{i,t}$) back to the electricity provider instantly. We assume that the electricity provider requires X_t amount of energy to be consumed by consumers so that it can ensure reliable power grid operations. If the consumption differs from X_t, it suffers a loss $L(X_t - \sum_{i \in \mathcal{H}} \alpha_{i,t} D_{i,t})$, which is assumed to be a continuous and convex function of $X_t - \sum_{i \in \mathcal{H}} \alpha_{i,t} D_{i,t}$.

We denote $\boldsymbol{D}_t = (D_{1,t}, D_{2,t}, ..., D_{M,t})$ and $\boldsymbol{\alpha}_t = (\alpha_{1,t}, \alpha_{2,t}, ..., \alpha_{M,t})$ to be the vector of net electricity demand and the fraction of net demand consumed directly from the grid, respectively. The total payment by the electricity provider to consumers when offering β_t at time t is given by

$$PMT_t(\beta_t, \boldsymbol{\alpha}_t \boldsymbol{D}_t^T) = \sum_{i \in \mathcal{H}} U_{i,t}(\alpha_{i,t}, \beta_t, D_{i,t}). \tag{3.4}$$

We define the profit of the electricity provider when offering incentive price β_t and supplying $\sum_{i \in \mathcal{H}} \alpha_{i,t} D_{i,t}$ amount of electricity to consumers as a bounded function $G(\beta_t, \boldsymbol{\alpha}_t \boldsymbol{D}_t^T)$. This profit function accounts for the cost of offering incentives as well as the gain from better load forecasting and efficient system operation. As a result, the net profit of the electricity provider after offering incentive β_t is given by:

$$V_t(\beta_t, \boldsymbol{\alpha}_t) = G(\beta_t, \boldsymbol{\alpha}_t \boldsymbol{D}_t^T) - L(X_t - \boldsymbol{\alpha}_t \boldsymbol{D}_t^T). \tag{3.5}$$

The expected net profit of the electricity provider is

$$V_t^E(\boldsymbol{p}_t, \boldsymbol{q}_t) = \sum_{\boldsymbol{\alpha}_t \in \mathcal{A}_t, \beta_t \in \mathcal{B}_t} (\prod_{i=1}^{M} p_{i,t}(\alpha_{i,t})) q_t(\beta_t) V_t(\beta_t, \boldsymbol{\alpha}_t). \tag{3.6}$$

The objective of the electricity provider is to choose q_t such that it maximizes its expected net profit from offering incentives. Thus, the electricity provider solves the following optimization problem

$$\max_{\boldsymbol{q}_t \in \boldsymbol{\mathcal{Q}}_t} \quad V_t^E(\boldsymbol{p}_t, \boldsymbol{q}_t). \tag{3.7}$$

3.2 Consumer-Electricity Provider Game

For a given set of electricity demand $\boldsymbol{D}_t$, the amount of electricity that each consumer consumes from the grid and the incentive price strongly impact both the profit for electricity provider and rewards for consumers. Moreover, the strategy of each consumer also affects other consumers' strategies indirectly by influencing the strategy of the electricity provider. To capture the trade-off between privacy costs and incentive costs, we use non-cooperative game theory to study the provider's incentive pricing policy and consumers' energy consumption fractions over time. The structure of the consumer-electricity provider game is given as follows:

- Set of players: $\{(\mathcal{H}, Ep)\}$ is the set of players in which consumers belong to set $\mathcal{H}$ and the electricity provider is denoted by Ep.

- Set of strategies: $\{(\mathcal{A}_t, \mathcal{B}_t)\}$ is the tuple of strategy sets for consumers and the electricity provider, where the strategy of consumer i (consumption fraction $\alpha_{i,t}$) lies in $\mathcal{A}_{i,t}$ and the strategy of the electricity provider Ep (incentive price β_t) belongs to $\mathcal{B}_t$.

- Payoff functions: $\{(\{R_{i,t}\}_{i \in \mathcal{H}}, V_t)\}$ is the tuple of payoff functions in which we denote $R_{i,t}$ to be the reward for consumer i and V_t to be the net profit for the electricity provider Ep.

3.2.1 Mixed Strategy Nash Equilibrium

The resulting strategic game can be written as $\{(\mathcal{H}, Ep), (\mathcal{A}_t, \mathcal{B}_t), (\{R_{i,t}\}_{i\in\mathcal{H}}, V_t)\}$. It has one well-studied solution called the mixed strategy Nash equilibrium. A mixed strategy Nash equilibrium is a probabilistic strategy tuple (i.e., set of probability distributions on pure strategies of each player) in which none of the players can be more profitable by unilaterally deviating to any pure strategy from this equilibrium strategy. It presents a stable outcome of interactions between consumers and the electricity provider. The mixed strategy Nash equilibrium is defined as follows:

Definition 1. *Consider the strategic game given by* $\{(\mathcal{H}, Ep), (\mathcal{A}_t, \mathcal{B}_t), (\{R_{i,t}\}_{i\in\mathcal{H}}, V_t)\}$, *a mixed strategy tuple* $(\{\boldsymbol{p}^*_{i,t}\}_{i\in\mathcal{H}}, \boldsymbol{q}^*_t)$ *is a mixed strategy Nash equilibrium if and only if* $R^E_{i,t}(\boldsymbol{p}^*_{i,t}, \boldsymbol{p}^*_{-\mathbf{i},t}, \boldsymbol{q}^*_t) \geq R^E_{i,t}(p_{i,t}(\alpha_{i,t}) = 1, \boldsymbol{p}^*_{-\mathbf{i},t}, \boldsymbol{q}^*_t)$, $\forall \alpha_{i,t} \in \mathcal{A}_{i,t}, i \in \mathcal{H}$ *and* $V^E_t(\boldsymbol{p}^*_t, \boldsymbol{q}^*_t) \geq V^E_t(\boldsymbol{p}^*_t, q_t(\beta_t) = 1)$, $\forall \beta_t \in \mathcal{B}_t$, *where the vector* $\boldsymbol{p}_{-\mathbf{i},t}$ *denotes the mixed strategies of all other consumers.*

Proposition 1. *There exists at least one mixed strategy Nash equilibrium for the above consumer-electricity provider game.*

The proof of Proposition 1 is provided in [132]. Generally, finding the Nash equilibrium is not easy. One method to do so is using best response [133]. The best response is a function which captures the behavior of each player by making other players' strategies fixed. By Definition 1, in a mixed strategy Nash equilibrium, each player plays the best response *w.r.t* other players' strategies.

3.3 A Two-Player Example

In this section, we study the case in which there is only one consumer and one electricity provider in the game. We denote $\alpha^0_t D_t$ to be the electricity consumption decided by the privacy preserving algorithm. The consumer can deviate either above or below $\alpha^0_t D_t$ and we consider two levels for each such deviation. Correspondingly,

we assume that the incentive price provided by the electricity provider has a two tier pricing structure. Thus, we define $\mathcal{A}_t \triangleq \{\alpha_t^0 - \delta_t^H, \alpha_t^0 - \delta_t^L, \alpha_t^0 + \delta_t^L, \alpha_t^0 + \delta_t^H\}$ and $\mathcal{B}_t = \{\beta_t^L, \beta_t^H\}$ to be the strategy set of the consumer and the electricity provider, respectively. Furthermore, we assume that $0 < \delta_t^L < \delta_t^H$, $0 < \beta_t^L < \beta_t^H$.

Definition 2. *A mixed strategy Nash equilibrium is nondegenerate if each player plays more than one of its pure strategies with non-zero probability.*

Theorem 3. *For the two-player consumer-electricity provider game defined by* $\{(\mathcal{H}, Ep), (\mathcal{A}_t, \mathcal{B}_t), (\{R_t\}, V_t)\}$*, there exists a unique nondegenerate mixed strategy Nash equilibrium if*

$$\begin{cases} \quad sign(G(\beta_t^H, (\alpha_t^0 + \delta_t^H)D_t) - G(\beta_t^L, (\alpha_t^0 + \delta_t^H)D_t)) \\ = sign(G(\beta_t^L, (\alpha_t^0 + \delta_t^L)D_t) - G(\beta_t^H, \alpha_t^0 + \delta_t^L D_t)); \\ \quad sign(G(\beta_t^H, (\alpha_t^0 - \delta_t^H)D_t) - G(\beta_t^L, (\alpha_t^0 - \delta_t^H)D_t)) \\ = sign(G(\beta_t^L, (\alpha_t^0 - \delta_t^L)D_t) - G(\beta_t^H, (\alpha_t^0 - \delta_t^L)D_t)) \end{cases} \tag{3.8}$$

and

$$\begin{cases} \beta_t^L(\delta_t^H - \delta_t^L) < \frac{f(\delta_t^H, D_t) - f(\delta_t^L, D_t)}{D_t} < \beta_t^H(\delta_t^H - \delta_t^L) \\ \beta_t^L(\delta_t^H - \delta_t^L) < \frac{f(-\delta_t^H, D_t) - f(-\delta_t^L, D_t)}{D_t} < \beta_t^H(\delta_t^H - \delta_t^L) \end{cases}, \tag{3.9}$$

where $sign(\cdot)$ *denotes the algebraic sign function.*

The intuition behind (3.8) is that for different deviation levels of the consumer (δ_t^H, δ_t^L), the difference in the profit for offering β_t^H and β_t^H should have different signs, otherwise, the electricity provider has an incentive to deviate to one of the pure strategies unilaterally. Similarly, the intuition behind (3.9) is that the difference in the monetized privacy leakage loss when playing δ_t^H and δ_t^L should be within the limits of the difference in reward a consumer can get from the electricity provider. Otherwise, the consumer has an incentive to deviate to one of the pure strategies unilaterally.

Proof. Due to the variation of power supply and demand, the power supply X_t can be either higher or lower than $\alpha_t^0 D_t$.

Case 1: If $X_t \geq \alpha_t^0 D_t$, the electricity supply is more than the amount that the consumer is willing to consume by using the privacy preserving algorithm. Thus, the electricity provider can use incentives to encourage the consumer to increase its consumption in order to balance power supply with demand. By (3.1) and (3.3), $\alpha_t^0 - \delta_t^H$ and $\alpha_t^0 - \delta_t^L$ are strictly dominated strategies since playing these strategies will result in a penalty from the electricity provider. Thus, we assume that the consumer plays $\alpha_t^0 + \delta_t^L$ and $\alpha_t^0 + \delta_t^H$ with probability q_t^* and $1 - q_t^*$ in the mixed strategy Nash equilibrium, respectively. If the electricity provider best-responds with a mixed strategy, it must be indifferent between playing β_t^L and β_t^H. Otherwise, the electricity provider has no reason to play a mixed strategy. Thus,

$$\begin{aligned} & q_t^*[G(\beta_t^L, (\alpha_t^0 + \delta_t^L)D_t) - L(X_t - (\alpha_t^0 + \delta_t^L)D_t)] \\ & \quad + (1 - q_t^*)[G(\beta_t^L, (\alpha_t^0 + \delta_t^H)D_t) - L(X_t - (\alpha_t^0 + \delta_t^H)D_t)] \\ & = q_t^*[G(\beta_t^H, (\alpha_t^0 + \delta_t^L)D_t) - L(X_t - (\alpha_t^0 + \delta_t^L)D_t)] \\ & \quad + (1 - q_t^*)[G(\beta_t^H, (\alpha_t^0 + \delta_t^H)D_t) - L(X_t - (\alpha_t^0 + \delta_t^H)D_t)]. \end{aligned} \tag{3.10}$$

In (3.11), we have the solution for (3.10).

$$q_t^* = \frac{G(\beta_t^H, (\alpha_t^0 + \delta_t^H)D_t) - G(\beta_t^L, (\alpha_t^0 + \delta_t^H)D_t)}{G(\beta_t^H, (\alpha_t^0 + \delta_t^H)D_t) + G(\beta_t^L, (\alpha_t^0 + \delta_t^L)D_t) - G(\beta_t^L, (\alpha_t^0 + \delta_t^H)D_t) - G(\beta_t^H, (\alpha_t^0 + \delta_t^L)D_t)} \tag{3.11}$$

Similarly, if the electricity provider plays β_t^L and β_t^H with probability p_t^* and $1 - p_t^*$ in the mixed strategy Nash equilibrium, the consumer's mixed strategy best response must be indifferent to $\alpha_t^0 + \delta_t^L$ and $\alpha_t^0 + \delta_t^H$. Therefore, we have

$$\begin{aligned} & p_t^*[\beta_t^L \delta_t^L D_t - f(\delta_t^L, D_t)] + (1 - p_t^*)[\beta_t^H \delta_t^L D_t - f(\delta_t^L, D_t)] \\ & = p_t^*[\beta_t^L \delta_t^H D_t - f(\delta_t^H, D_t)] + (1 - p_t^*)[\beta_t^H \delta_t^H D_t - f(\delta_t^H, D_t)], \end{aligned}$$

which implies

$$p_t^* = \frac{\beta_t^H(\delta_t^H - \delta_t^L)D_t - (f(\delta_t^H, D_t) - f(\delta_t^L, D_t))}{(\beta_t^H - \beta_t^L)(\delta_t^H - \delta_t^L)D_t}. \tag{3.12}$$

Case 2: If $X_t < \alpha_t^0 D_t$, the electricity provider can use incentives to decrease electricity consumption. Thus, $\{\alpha_t^0 + \delta_t^H, \alpha_t^0 + \delta_t^L\}$ are strictly dominated strategies since playing these strategies will result in a penalty from the electricity provider. We assume that the consumer plays $\alpha_t^0 - \delta_t^L$ and $\alpha_t^0 - \delta_t^H$ with probability q_t^* and $1 - q_t^*$ in the mixed strategy Nash equilibrium, respectively. Furthermore, we assume that the electricity provider plays β_t^L with probability p_t^*. By following the same argument for the case in which $X_t \geq \alpha_t^0 D_t$, in the mixed Nash equilibrium, we have q_t^* similar to (3.11) with δ_t^H, δ_t^L replaced by $-\delta_t^H, -\delta_t^L$ and

$$p_t^* = \frac{\beta_t^H(\delta_t^H - \delta_t^L)D_t - (f(-\delta_t^H, D_t) - f(-\delta_t^L, D_t))}{(\beta_t^H - \beta_t^L)(\delta_t^H - \delta_t^L)D_t}. \tag{3.13}$$

Thus, by (3.11), (3.12), and (3.13), we have Theorem 3. Moreover, since (3.11)-(3.13) are all computable for fixed pure strategy sets and reward/profit functions, there exists a unique nondegenerate mixed strategy Nash equilibrium if (3.8) and (3.9) hold. □

For given electricity supply-demand profile and reward-loss functions, Theorem 3 provides a relationship between the incentive price (β_t) and the fraction of total electricity demand from the grid (α_t), such that the players achieve a unique nondegenerate mixed strategy Nash equilibrium.

3.4 Illustration of Results

In this section, we illustrate our model and results. For simplicity, we consider the following net profit function:

$$V_t(\beta_t, \alpha_t) = g(\beta_t, \alpha_t D_t) - PMT_t(\beta_t, \alpha_t D_t) - r_L(X_t - \alpha_t D_t)^2, \tag{3.14}$$

where

$$g(\beta_t, \alpha_t D_t) = \begin{cases} 1/3|\alpha_t - \alpha_t^0|D_t & \text{if } \beta_t = \beta_t^H, \delta_t = \delta_t^H \\ 1/6|\alpha_t - \alpha_t^0|D_t & \text{if } \beta_t = \beta_t^L, \delta_t = \delta_t^H \\ 0.2|\alpha_t - \alpha_t^0|D_t & \text{if } \beta_t = \beta_t^L, \delta_t = \delta_t^L \\ 0.3|\alpha_t - \alpha_t^0|D_t & \text{if } \beta_t = \beta_t^H, \delta_t = \delta_t^L \end{cases} . \tag{3.15}$$

The first term of (3.14) denotes the benefit for offering incentives. The intuition behind this function is that offering β_t^H (β_t^L) indicates that the electricity provider desires the consumer to deviate a large (small) amount from α_t^0 and is willing to pay a high (low) incentive price. Thus, the gain for strategy δ_t^H (δ_t^L) should be larger than δ_t^L (δ_t^H). The second term is the incentive payment to the consumer and the third term is the loss function caused by supply-demand imbalance. We use r_L as the supply-demand imbalance loss factor assigned by the electricity provider. The privacy valuation function of the consumer is given by

$$f_t(\alpha_t - \alpha_t^0, D_t) = f_t(\delta_t, D_t) = r_P(\delta_t D_t)^2, \tag{3.16}$$

in which privacy loss is captured by how much does the strategy deviate from $\alpha_t^0 D_t$. The constant r_P indicates the rate of privacy leakage loss $w.r.t$ the deviation from the optimal privacy preserving consumption $\alpha_t^0 D_t$.

We assume the PV system can provide approximately 70% of total electricity consumption of a consumer. At time t, we choose the strategy of the consumer when there is no incentive offered by the electricity provider as $\alpha_t^0 = 0.3$. We choose the following deviation fractions and incentive prices: $\delta_t^L = 0.1, \delta_t^H = 0.3, \beta_t^L = 0.05\$/KW, \beta_t^H = 0.2\$/KW$.

In Figure 3.2, we plot the relationship between the mixed strategy Nash equilibrium (p_t^*, q_t^*) and the privacy leakage loss factor r_P. We assume $D_t = 10KW$ and $X_t = 4.2KW$. In the mixed strategy Nash equilibrium, the strategy of the consumer

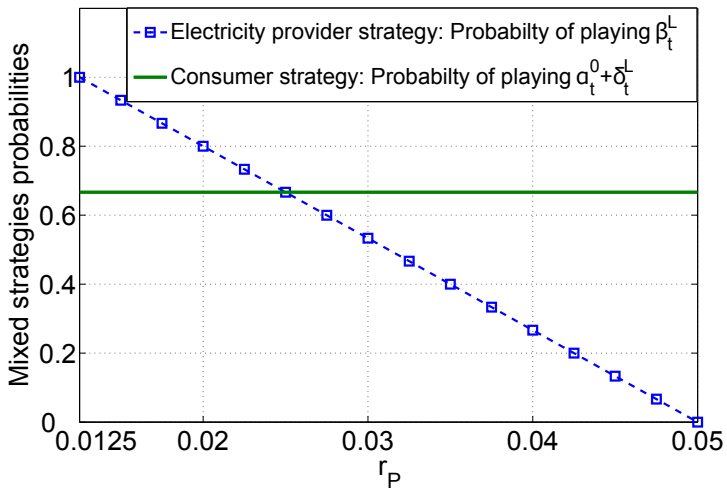

Figure 3.2: Mixed strategy Nash equilibrium vs. privacy leakage cost r_P.

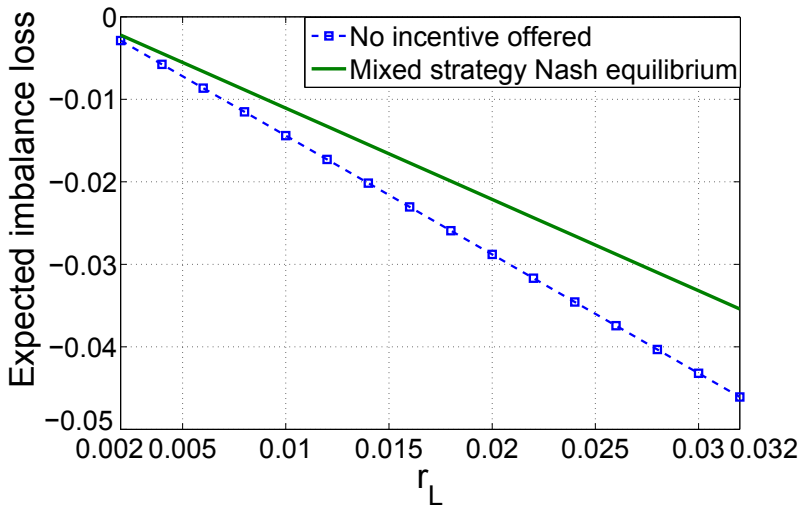

Figure 3.3: Supply-demand imbalance loss with/without incentives vs. imbalance loss factor r_L. ($r_P = 0.03\$/KW^2$)

does not change with r_P. However, for the electricity provider, the probability of playing β_t^L decreases with r_P. This is due to the fact that the electricity provider has to offer a high incentive price for consumers who have a high valuation of their privacy. Figure 3.3 illustrates the loss due to supply-demand imbalance with/without incentives for different imbalance loss factors (r_L). The figure shows that offering incentives can help reduce supply-demand imbalance loss.

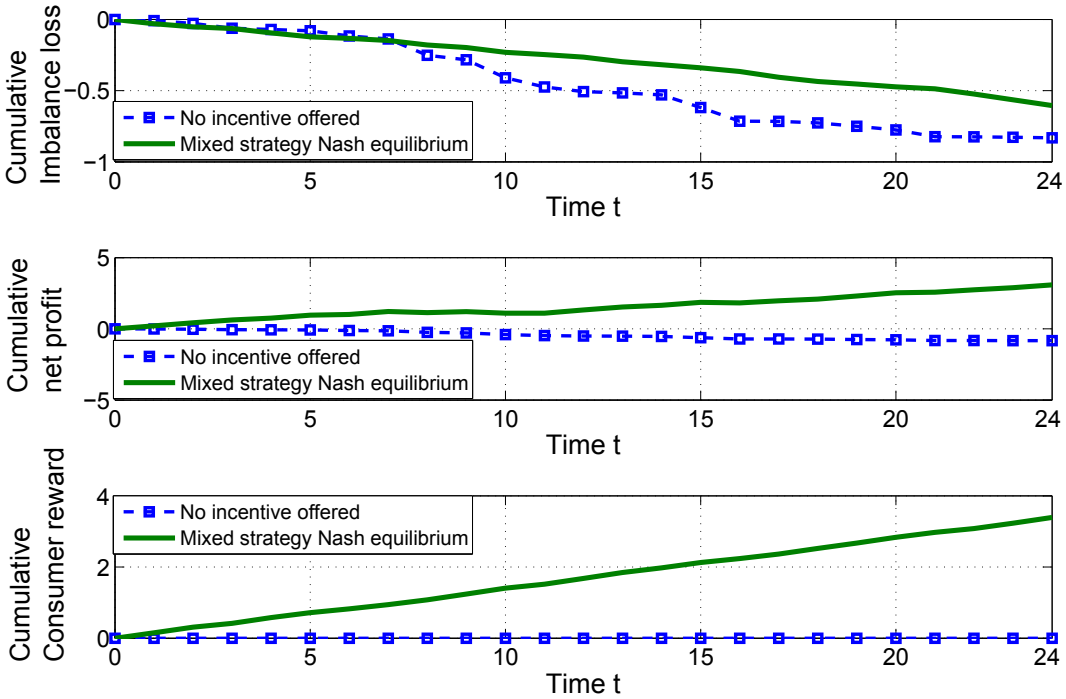

Figure 3.4: Cumulative imbalance loss, net profit of the electricity provider as well as reward of the consumer

To study the performance of the proposed incentive mechanism, we model the demand of the consumer (D_t) and the privacy preserving consumption fraction (α_t^0) as random variables which follow truncated standard normal distributions. We assume $D_t \in [9, 11]$ and $\alpha_t^0 \in [0, 0.6]$ with mean value equal to $10KW$ and 0.3. Furthermore, we assume the requirement of electricity provider $X_t \in [0, 6]$ also follows a truncated standard normal distribution with mean value $3KW$. We divide the time horizon into 24 hours and use the following parameters: $r_L = 0.015\$/KW^2, r_P = 0.03\$/KW^2$.

The cumulative net profit, loss of the electricity provider as well as the reward of the consumer at time t are defined to be the summation of net profit, loss of the electricity provider and reward of the consumer from time 1 to t, respectively. Figure 3.4 shows that the proposed mechanism can successfully incentivize data sharing from privacy-sensitive consumers to both increase net profit of the electricity provider and reduce loss incurred by supply-demand imbalance.

Chapter 4

THE IMPACT OF PRIVACY ON FREE ONLINE SERVICE MARKETS

4.1 Problem Model and Game Formulation

In this section, we introduce a game-theoretic model for two SPs that offer the same type of free online services (e.g., search engine, social network) and infinitely many consumers. Each SP offers the free services with a quantified privacy risk guarantee ε and quality of service (QoS) v. Just as Google at present advertises RAPPOR [37] with a certain level of differential privacy risk, in the future, it is possible that SPs will adopt one or more metrics to quantify their privacy risks. This paper makes such an assumption of privacy risk quantifiability. Furthermore, we assume SPs advertise their quantified privacy risk and QoS to consumers. Thus, both ε and v are observable to consumers. The observable privacy risk value could be the ε value in differential privacy adopted by Google RAPPOR and the QoS could be the accuracy of search results. An SP differentiates its service by a tuple (v, ε) that it advertises to all consumers. A consumer's preference of privacy differentiated service is modeled by a utility function which depends on its privacy risk valuation and the QoS-privacy risk tuple offered by the SP. In reality, it is natural to assume that consumers prefer high QoS and low privacy risk. Thus, in our model, a consumer will have a higher utility if he or she receives higher QoS or lower privacy risk. Finally, consumer privacy heterogeneity is modeled as a distribution.

4.1.1 Two-SP Market Model

SP Model

We consider two rational (i.e., profit maximization entities) SPs, denoted by SP_1 and SP_2. Both SPs provide the same kind of free service; but they differ in the QoS offered. Thus, SP_1 and SP_2 offer QoS v_1 and v_2, respectively, where in general

$v_1 \neq v_2$. Furthermore, SP_1 and SP_2 guarantee that the privacy risk for using their services is at most ε_1 and ε_2, respectively, where $\varepsilon_1, \varepsilon_2 \in [0, \bar{\varepsilon}]$. Without loss of generality, we assume $\varepsilon_2 \geq \varepsilon_1$. Under this assumption, SP_2 must offer a higher QoS ($v_2 \geq v_1$). Otherwise, its strategy will be dominated by its opponent since SP_1 will offer both higher QoS and lower privacy risk. For example, SP_1 and SP_2 could be Duckduckgo and Google, respectively, in the search engine market, with the QoS given by the accuracy of search results. On the other hand, the privacy risk can correspond to different guarantees they provide on consumer data use; e.g., whether they will use consumer data only for statistical purposes or target consumers with tailored ads. We model this privacy risk guarantee as a variable taking values over a continuous range. In practice, such guarantees may be coarse granular choices; for example, between completely opting out of the targeting or allowing data use only for statistical purposes or complete data use only by SP or all possible data usage and sale. We assume that the SPs generate revenue in two ways: (i) by exploiting the private data of consumers to offer *targeted* ads and other services to consumers; and (ii) by providing interested advertisers an online platform to reach consumers. This latter revenue is independent of private data and simply derived from the revenue capability of the platform.

Let $R_P(\varepsilon_i)$ denote the revenue of SP_i, $i \in \{1, 2\}$, resulting from using the private data of consumers and let $R_{\text{NP},i}$ denote the revenue generated without using consumers' private information (e.g., from interested advertisers). The total revenue, $R(\varepsilon_i)$, of SP_i from offering privacy guaranteed service is thus

$$R(\varepsilon_i) = R_P(\varepsilon_i) + R_{\text{NP},i}, i \in \{1, 2\}. \tag{4.1}$$

Notice that in reality, through spillovers and externalities associated with using consumers' private data, the revenue generating capabilities for firms can increase even

from sources that don't directly use consumer personal information. However, it is very hard to capture these externalities precisely since they are highly data and service model dependent. We start with a simple model in which we assume that SPs will not use consumers' private data for services that do not require private data. Our proposed model provides an intuition on the equilibrium strategies of SPs and market. Furthermore, it is useful to note that even this relatively simple revenue decoupled setting is highly parameterized. Our analysis allows us to understand the dependencies on the various parameters in the problem.

Offering free services to consumers often comes with a cost to the SPs, such as the cost of service, online platform creation, and continued operations. Furthermore, we note that free online services profit from using consumer data and therefore incur data processing related costs. Let $C(v_i; \varepsilon_i)$ denote the cost of offering free services with privacy risk level ε_i. We model $C(v_i; \varepsilon_i)$ as sum of two non-negative costs: (i) $C_{\text{QoS}}(v_i)$ of providing services with QoS v_i; and (ii) $C_P(\varepsilon_i)$ as the processing (data analytics) cost of exploiting private data to the privacy risk level of ε_i such that

$$C(v_i; \varepsilon_i) = C_{\text{QoS}}(v_i) + C_P(\varepsilon_i), i \in \{1, 2\}. \tag{4.2}$$

We assume $R_P(\varepsilon_i) - C_P(\varepsilon_i) > 0$. Otherwise, SP_i will not have incentives to exploit consumers' private information since the cost of processing consumers' private information is higher than the revenue it gains from using consumers' private information. Thus, via (4.1) and (4.2), our model captures the fact that the benefit of using private data by each SP involves both cost and revenue.

Consumer Model

We formulate both consumer utility and the resulting consumer-SP game based on the classical Hotelling model. The Hotelling model maps retailers to two locations

(x_1, x_2) on a $[0, 1]$ line such that the strategy of each retailer is to determine the best location-price tuple that maximizes its profit. The location (see Figure 4.1a) is a proxy for a specific product differentiator. A consumer with its own product differentiator preference (traditionally assumed to be uniformly distributed over $[0, 1]$) is mapped to a location $x \in [0, 1]$ on the line as shown in Figure 4.1a. Such a spatial model allows computing the market segment by identifying both the optimal locations of the retailers and an indifferent threshold between the two optimal retailer locations at which both retailers are equally desirable. For such a uniform consumer preference model, the segmentation for each retailer is simply its distance to the indifference point (see Figure 4.1a). Consumers choose the retailer with the least product price and "transportation cost" (modeled as a linear function of location) for a desired consumer valuation of the product. Note that transportation costs are metaphorical for any non-price-based differentiation of the two retailers.

For our problem, we obtain a Hotelling model by: (i) introducing a *normalized privacy risk* and mapping it to spatial location; and (ii) by viewing the QoS as the net valuation of service by the consumer. Note that since we study a free services market, we use QoS as a measure of consumer satisfaction. We note that in the classical Hotelling model, the consumer pays a non-negative transportation cost for any retailer whose location is different from its own. However, our problem departs from this model in that higher and lower privacy risks offered by SPs relative to a consumer preferred privacy risk choice are not viewed similarly.

We assume there exists infinitely many rational consumers that are interested in the services provided by the SPs. In keeping the standard game-theoretic definition, rational refers to consumers interested in maximizing some measure of utility via interactions with the SPs. We use a random variable $E \in [0, \bar{\varepsilon}]$ to denote the heterogeneous privacy preferences of consumers; such a model assumes that the privacy

preferences of consumers are independent and identically distributed, a reasonable assumption when the consumer set is very large. Let $E = \varepsilon$ denote the privacy risk preference of a consumer. If SP_i offers a privacy risk guarantee ε_i higher than ε, then using its service will result in a privacy cost to the consumer due to perceived privacy risk violation. On the other hand, the consumer gains from choosing an SP_i that offers an $\varepsilon_i < \varepsilon$ as a result of the extra privacy protection offered. Let $x = F_E(\varepsilon) \in [0, 1]$ be a differentiable cumulative distribution function of ε. Thus, $x = F_E(\varepsilon)$ can be considered as a normalized privacy risk tolerance (i.e., restricted to $[0, 1]$) which indicates the proportion of the consumers with a privacy risk tolerance of at most ε. Since the privacy risk ε can be over an arbitrary range $[0, \bar{\varepsilon}]$, the normalized spatial privacy risk is given by the cumulative distribution function (CDF) $F_E(\varepsilon)$. Thus, for a consumer whose normalized privacy risk tolerance is located at $x \in [0, 1]$, its actual privacy risk tolerance is $\varepsilon = F_E^{-1}(x)$. We can similarly map the privacy risks offered by the SPs to *normalized locations* $x_1 = F_E(\varepsilon_1)$ and $x_2 = F_E(\varepsilon_2)$ on the $[0, 1]$ line as shown in Figure 4.1b.

Analogous to the Hotelling model, we let $u_i(x)$ denote the utility (in units of QoS) from SP_i as perceived by a consumer with a normalized privacy preference (location) x. Our model for $u_i(x)$ contains two parts: (i) a positive QoS v_i offered by SP_i; and (ii) the gain or loss in the perceived QoS as a result of a mismatch between consumer privacy preference and SP_i's privacy risk offering. We introduce a gain factor t that allows mapping the privacy mismatch $t(x - x_i)\varepsilon_i$ to a QoS quantity. This mismatch utility indicates that when the SP offers a service with privacy risk lower than the consumer's tolerance, the consumer receives a positive utility due to extra privacy protection. However, if the service offered has a higher privacy risk than the consumer's tolerance, the consumer will receive negative utility for privacy risk violation. In other words, given the same level of QoS, the better the privacy

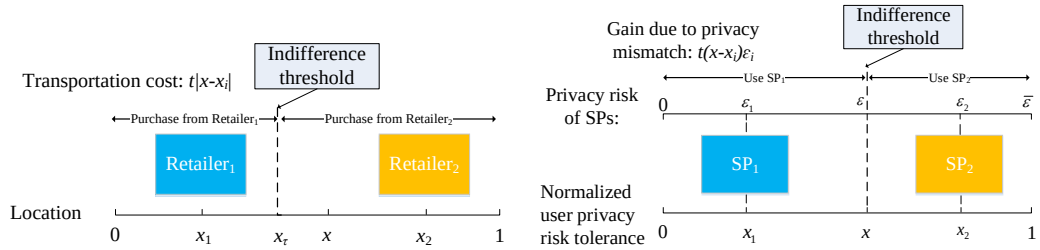

(a) Classical hotelling model (b) Our modified hotelling model

Figure 4.1: User choice model for using different SPs

risk guarantee an SP offers, the more the consumer prefers the SP. We now write the utility or profit function for consumers and SPs.

Consumer utility and SP profits

For the consumer located at x, the overall perceived utility for choosing services provided by SP_1 and SP_2 are

$$u_i(x) = v_i + t(x - x_i)\varepsilon_i, i \in \{1, 2\}. \tag{4.3}$$

For each $SP_i, i \in \{1, 2\}$, let $(v_{-i}, \varepsilon_{-i})$ be its competitor's strategy. For the revenue and cost models in (4.1) and (4.2), the profit of SP_i is simply the difference

$$\pi_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i}) = [R(\varepsilon_i) - C(v_i; \varepsilon_i)]n_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i}), \tag{4.4}$$

where $n_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i})$ denotes the fraction of consumers who choose SP_i.

Modelling Assumption 3. *We assume that the services provided by both SPs have non-negative QoS.*

Since consumers are rational, they expect to have positive utility through the interactions with the SPs. It is reasonable to assume that SPs have no incentive to offer services with a negative QoS. In other words, we assume $v_1 \geq 0$ and $v_2 \geq 0$.

Modelling Assumption 4. *We assume the model parameters are chosen such that they ensure the market is completely covered by SP_1 and SP_2.*

The above assumption implies that each consumer must choose one of the SPs. Such an assumption is implicitly built into the classical Hotelling model to ensure competition between SPs and our model continues to do so too. Later we provide a sufficient condition for sustaining the equilibrium market segmentation under these assumptions.

4.1.2 Two-SP Non-cooperative Game Formulation

We note that the SPs compete against each other through their distinct QoS and privacy risk offerings, which in turn affects consumer choices and helps determine the stable market segmentation. Thus, the interactions between SPs can be formulated as a non-cooperative game in which the strategy of each SP is a (QoS, privacy risk) tuple and that of the consumer is choosing an SP. Furthermore, we assume that the SPs are rational and have perfect and complete information. They play to maximize their own profits and know the exact profit function for any given strategy.

The Game: the interactions between retailers and consumers in the Hotelling model can be viewed as a sequential game [51]. For our model, such a sequential game involves three stages (Figure 4.2). In the first stage, the differentiator, i.e., the normalized privacy risk ε_i, is advertised by SP_i. This is followed by each SP determining its QoS for the advertised risk. Finally, the consumers choose the preferred SP based on the (v_i, ε_i) tuple that maximizes its utility. Our sequential game assumes that the selection of privacy risk happens before the selection of the QoS. This is due to the fact that SPs first advertise their privacy risks to differentiate themselves from their competitors (e.g., Google advertises RAPPOR while Duckduckgo advertises not using private data of the consumers) and then adjust their QoS strategies based on the advertised privacy risks and the privacy preferences of the consumers.

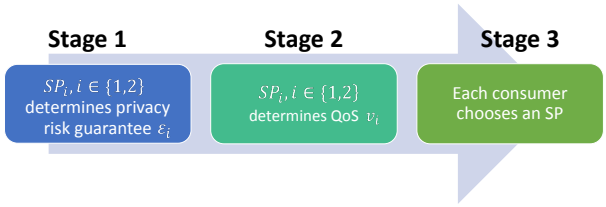

Figure 4.2: The three-stage sequential game

The game can be formally described as follows: (i) a set of players $\{1, 2, \mathcal{C}\}$, where 1 and 2 denote SP_1 and SP_2, respectively, and the set $\mathcal{C}$ contains infinitely many consumers; (ii) a collection of strategy tuples $(v_i, \varepsilon_i) \in \mathcal{V}_i \times \mathcal{E}_i$ for SP_i and a collection of binary choices (strategies) for the consumer $b \in \mathcal{B} = \{1, 2\}$; and (iii) a profit function π_i for each SP_i and a utility function u_i for each consumer for choosing SP_i.

4.2 The Subgame Perfect Nash Equilibrium for the Two-SP Game

In a sequential game, each stage is referred to as a subgame [133]. One often associates a strategy profile with a sequential game. A strategy profile is a vector whose i^{th} entry is the strategy for all players at the ith stage of the sequential game. A non-cooperative sequential game has one well-studied solution: the Subgame Perfect Nash Equilibrium (SPNE). A strategy profile is an SPNE if its entries are the Nash equilibria of the subgame resulting at each stage of the sequential game. The SPNE of a sequential game captures an equilibrium solution such that no player can make more profit by unilaterally deviating from this strategy in every subgame.

Since the above non-cooperative game is a game with finite number of stages and perfect information, it can be solved using backward induction. Backward induction is the process of reasoning backwards in time (or sequence), starting from the last stage of the sequential game, to determine a sequence of optimal strategies. It proceeds by first determining the optimal strategies in the last stage. Using this information, one can then decide the optimal strategies for the second-to-last stage of the game.

This process continues backwards until the optimal strategies for every stage has been determined. We apply backward induction to the three-stage game as follows.

Stage 3, Users' decisions: Each consumer located at $x \in [0, 1]$ can choose the services provided by either SPs based on its valuation function in (4.3). The resulting optimal strategy for the consumer is to choose the SP whose index is given by

$$arg \max_{i \in \{1,2\}} v_i + t(x - x_i)\varepsilon_i. \tag{4.5}$$

Since the consumer's utility is a linear function of the normalized privacy risk x and the market is completely covered by the SPs, there exists a threshold x_τ such that the consumer located at x_τ is indifferent to using services provided by SP_1 or SP_2. Thus, at the indifference threshold x_τ, we have

$$u_2(x_\tau) = u_1(x_\tau) \implies v_2 + t(x_\tau - x_2)\varepsilon_2 = v_1 + t(x_\tau - x_1)\varepsilon_1. \tag{4.6}$$

Simplifying further leads to the indifference threshold for choosing between the SPs

$$x_\tau = \frac{v_1 - v_2 + t(F_E(\varepsilon_2)\varepsilon_2 - F_E(\varepsilon_1)\varepsilon_1)}{t(\varepsilon_2 - \varepsilon_1)}, \tag{4.7}$$

where x_1 and x_2 have been replaced by their corresponding normalized privacy risk values. Thus, given the SPs' strategies $(v_i, \varepsilon_i), i \in \{1, 2\}$, the optimal strategy of a consumer located at x is to use the service of SP_1 if $x \leq x_\tau$ and SP_2 otherwise. If $v_1 = v_2$ and $\varepsilon_1 = \varepsilon_2$, consumers are indifferent between SP_1 and SP_2. In this case, we assume the consumers use the following tie-breaking rule:

Modelling Assumption 5. *If $v_1 = v_2$ and $\varepsilon_1 = \varepsilon_2$, consumers choose either SPs with probability $\frac{1}{2}$.*

Stage 2, SPs determine QoS: In the second stage, for a given privacy risk guarantee ε_i, SP_i chooses its QoS v_i to maximize its profit π_i. Since a consumer's normalized privacy risk tolerance denotes the fraction of the population whose privacy

risk tolerance is at most ε, x_τ determines the proportion of consumers who choose SP_1, i.e., n_1. As a result, the profit functions of SP_1 and SP_2 can be written as

$$\pi_1(v_1;\varepsilon_1;v_2;\varepsilon_2) = [R(\varepsilon_1) - C(v_1;\varepsilon_1)]x_\tau, \tag{4.8}$$

$$\pi_2(v_1;\varepsilon_1;v_2;\varepsilon_2) = [R(\varepsilon_2) - C(v_2;\varepsilon_2)](1 - x_\tau). \tag{4.9}$$

To find the SPNE in this stage, we use the best response method [134]. The best response is a function which captures the behavior of each player while fixing the strategies of the other players. For any $v_{-i} \in \mathcal{V}_{-i}$, we define $BR_i(v_{-i})$ as the best strategy of SP_i such that

$$BR_i(v_{-i}) = arg \max_{v_i} \pi_i(v_i;\varepsilon_i;v_{-i};\varepsilon_{-i}), i \in \{1,2\}. \tag{4.10}$$

In the SPNE, each player plays the best response $w.r.t$ other players' strategies. Thus, a Nash equilibrium in this stage is a profile $\boldsymbol{v}^* = (v_i^*, v_{-i}^*)$ for which

$$v_i^* \in BR_i(v_{-i}), \forall i \in \{1,2\}. \tag{4.11}$$

To find the Nash equilibria, we first calculate the best response function of each SP, then find a strategy profile $\boldsymbol{v}^*$ for which $v_i^* \in BR_i(v_{-i}), \forall i \in \{1,2\}$. For a given set of privacy risk guarantees $\{\varepsilon_1, \varepsilon_2\}$, the optimal QoS v_i^* of $SP_i, i \in \{1,2\}$ in the SPNE is then determined by the solution to the following set of simultaneous equations

$$v_1^* = arg \max_{v_1} \pi_1(v_1;\varepsilon_1;v_2;\varepsilon_2), \tag{4.12}$$

$$v_2^* = arg \max_{v_2} \pi_2(v_1;\varepsilon_1;v_2;\varepsilon_2). \tag{4.13}$$

Stage 1, SPs determine privacy risk guarantee: In the first stage, we compute equilibrium strategies ε_1 and ε_2 that the two SPs should advertise for optimal market share. Note that v_1^*, v_2^*, and x_τ have been computed in stages 2 and 3 for a fixed ε_1 and ε_2, and therefore, are functions of ε_1 and ε_2. The objective functions π_1 and π_2 are thus also functions of ε_1 and ε_2; this in turn implies they can be maximized to find the equilibrium strategy ε_1^* and ε_2^* using the best response method.

4.3 Two-SP Market with Linear Cost and Revenue Functions

Thus far, we have considered a general model for consumer privacy preference. To obtain better intuition and meaningful analytical solutions, we consider a linear cost and revenue model for each SP. The cost function of SP_i is modeled as

$$C(v_i; \varepsilon_i) = cv_i + c\lambda\varepsilon_i, i \in \{1, 2\}, \tag{4.14}$$

where c and λ are constant scale factors in units of cost/QoS and QoS/privacy risk, respectively. We model the revenue of each SP from offering a privacy guaranteed service by a linear function

$$R(\varepsilon_i) = r\varepsilon_i + p_i, i \in \{1, 2\}, \tag{4.15}$$

where r is the revenue per unit privacy risk for using consumers' private data. The parameters p_1 and p_2 model the fixed revenues of the SPs that are independent of consumers' private data.

Theorem 4. *There does not exist any SPNE in which both SPs offer the same privacy risk.*

Proof sketch: The detailed proof of Theorem 4 is in Appendix E; we briefly outline the proof. First, we assume there exists an SPNE where both SPs offer the same privacy risk $\tilde{\varepsilon}$. Then, using backward induction, we show that one of the SPs is better off by unilaterally deviating from the equilibrium strategy $\tilde{\varepsilon}$; implying that there does not exist an SPNE in which both SPs offer the same privacy risk.

Remark: Note that the result of Theorem 4 does not exhibit the minimal differentiation behavior (i.e., both firms place themselves close to each other) observed in [51]. This is due to the fact that higher and lower privacy risks offered by SPs relative to a consumer preferred privacy risk choice are not viewed similarly; that is, the symmetric transportation cost no longer holds in our model, and thus resulting an asymmetric gain due to privacy mismatch in (4.3).

4.3.1 Uniform Consumer Privacy Risk Tolerance

We assume consumers have uniformly distributed privacy risk tolerance between 0 and $\bar{\varepsilon}$. The resulting normalized privacy risk of each SP is given by

$$x_i = F_E(\varepsilon_i) = \frac{\varepsilon_i}{\bar{\varepsilon}}, i \in \{1, 2\}.$$

We define

$$\alpha = \frac{r}{c} - \lambda, \quad \tilde{C} = ct\bar{\varepsilon}. \tag{4.16}$$

Note that α is the ratio of net profit from using consumer data for a unit of privacy risk to the cost for providing a unit of QoS. Furthermore, $\tilde{C}$ is the cost of providing non-zero utility to the consumer with a maximal mismatch of privacy risk (relative to SP).

By using the backward induction method, the computed SPNE of the two-SP non-cooperative game is presented in the following theorem.

Theorem 5. *There exists an SPNE given by*

$$\varepsilon_2^* = \frac{12\bar{\varepsilon}c\alpha + 15ct\bar{\varepsilon} - 16(p_2 - p_1)}{24tc}, \tag{4.17}$$

$$v_2^* = \frac{(2\alpha + t)c\alpha 6\bar{\varepsilon} + (\alpha - t)9ct\bar{\varepsilon} + (t - 2\alpha)8p_2 + (\alpha + t)16p_1}{24ct}, \tag{4.18}$$

$$\varepsilon_1^* = \varepsilon_2^* - \frac{3\bar{\varepsilon}}{4}, \tag{4.19}$$

$$v_1^* = v_2^* - \frac{3\bar{\varepsilon}}{4}\alpha + \frac{p_2 - p_1}{3c}, \tag{4.20}$$

if the model parameters $\{c, r, \lambda, t, \bar{\varepsilon}, p_1, p_2\}$ *satisfy*

$$-1 \leq \frac{16(p_2 - p_1)}{9ct\bar{\varepsilon}} \leq 1, \tag{4.21}$$

$$\frac{4\alpha - 3t}{3} \leq \frac{16(p_2 - p_1)}{9c\bar{\varepsilon}} \leq \frac{4\alpha - t}{3}, \tag{4.22}$$

$$(12(r - c\lambda)\bar{\varepsilon})^2 - (15ct\bar{\varepsilon})^2 + 288ct\bar{\varepsilon}(p_2 + p_1) \geq [16(p_2 - p_1)]^2. \tag{4.23}$$

At this SPNE, the market indifference threshold which determines the market segmentation is given by

$$x_\tau^* = \frac{1}{2} - \frac{8(p_2 - p_1)}{9ct\bar{\varepsilon}} = \frac{1}{2} - \frac{8(p_2 - p_1)}{9\tilde{C}}, \tag{4.24}$$

and the total profits of both SPs are

$$\pi_1^* = \frac{4c}{27t\bar{\varepsilon}}(\frac{9t\bar{\varepsilon}}{8} - \frac{2(p_2 - p_1)}{c})^2 = \frac{1}{3}(\frac{3}{4}\sqrt{\tilde{C}} - \frac{4(p_2 - p_1)}{3\sqrt{\tilde{C}}})^2 \tag{4.25}$$

$$\pi_2^* = \frac{4c}{27t\bar{\varepsilon}}(\frac{9t\bar{\varepsilon}}{8} + \frac{2(p_2 - p_1)}{c})^2 = \frac{1}{3}(\frac{3}{4}\sqrt{\tilde{C}} + \frac{4(p_2 - p_1)}{3\sqrt{\tilde{C}}})^2, \tag{4.26}$$

where α and $\tilde{C}$ are defined in (4.16).

Proof sketch: The proof of Theorem 5 is provided in Appendix F. We briefly sketch the proof details here. Our approach involves using a three-stage backward induction to compute equilibrium strategies starting from the third stage; at each stage, the equilibrium strategies are computed using those computed from future stages. In the third stage, for a fixed pair of strategies of each SP, the consumer makes the choice, this in turn helps determining the indifference threshold x_τ. This x_τ is now used in the second stage to compute the equilibrium QoS (v_1^*, v_2^*) for a fixed set of risk $(\varepsilon_1, \varepsilon_2)$. Finally, the first stage involves computing the equilibrium privacy risk for these choice of v_1^*, v_2^* and x_τ by solving the corresponding best response functions, thereby obtaining the solutions in (4.17)–(4.20). The conditions in (4.21)–(4.23) result from requiring the equilibrium strategies as well as the equilibrium market segmentation to satisfy the following: (i) feasible threshold: $0 \leq x_\tau^* \leq 1$; (ii) feasible risk: $0 \leq \varepsilon_1^*, \varepsilon_2^* \leq \bar{\varepsilon}$; (iii) non-zero consumer utility: $v_1^* - tx_1^*\varepsilon_1^* \geq 0$ or $v_2^* - tx_2^*\varepsilon_2^* \geq 0$. Substituting equilibrium strategies (4.17)–(4.20) and (4.16) into (4.7) yields the market share of SP_1 in (4.24). The profits for both SPs (see (4.25) and (4.26)) result from using equilibrium strategies (4.17)–(4.20) to compute (4.8) and (4.9).

Remark: Note that the equilibrium solution in (4.17)–(4.20) is highly parametrized. For a given set of parameters that satisfy conditions in (4.21)–(4.23), the sequential game yields an SPNE. By (4.24), the SP with higher privacy-independent revenue owns a larger market share, leading to a higher total profit in the SPNE (see (4.25) and (4.26)). Note that p_1 and p_2 are the only differentiator of SPs in the set of model parameters. For a fixed $p_2 - p_1$, both π_1^* and π_2^* are decreasing functions of $\tilde{C}$ when $\tilde{C} \in [0, \frac{16|p_2-p_1|}{9}]$ and increasing afterwards. On the other hand, by (4.21), we have $\frac{16|p_2-p_1|}{9} \leq ct\bar{\varepsilon}$, which implies $\tilde{C} \geq \frac{16|p_2-p_1|}{9}$. Thus, both π_1^* and π_2^* are increasing functions of $\tilde{C}$ in the SPNE. In the following, based on Theorem 5, we highlight the effect of each one of these model parameters on the SPNE while keeping all other parameters fixed.

1. *Heterogeneity of consumer privacy preferences* ($\bar{\varepsilon}$): for the SPNE presented in Theorem 5, observe that v_i^* and $\varepsilon_i^*, i \in \{1, 2\}$ are linear functions of $\bar{\varepsilon}$. Furthermore, $\varepsilon_2^* = \varepsilon_1^* + \frac{3}{4}\bar{\varepsilon}$; this implies that at the SPNE, the SP that offers the higher privacy risk (i.e., ε_2^*) offers exactly $\frac{3}{4}\bar{\varepsilon}$ higher than that of its competitor. For all other parameters fixed, as $\bar{\varepsilon}$ increases, SP_2's privacy risk increases linearly. On the other hand, SP_1's privacy risk increases linearly with $\bar{\varepsilon}$ only if the model parameters are such that $4(r - c\lambda) > ct$; otherwise, it decreases linearly (see (F.15) in Appendix F). To further understand the dependency, we consider the following two cases:

   - If $p_2 - p_1 > 0$, as $\bar{\varepsilon}$ increases, SP_2 can increase its revenue by increasing its privacy risk offerings. As a result, more consumers who have low privacy risk preferences will choose SP_1. Therefore, the market share of SP_2 decreases. However, the profits of both SPs increases as $\bar{\varepsilon}$ increases. The intuition behind this is that when $\bar{\varepsilon}$ increases, both SPs can exploit con-

sumers' private information from a larger range of privacy risk preferences. As a result, their revenue from exploiting consumers' private information also increases at the SPNE, which in turn leads to an increase in both SPs' profits.

- If $p_2 - p_1 < 0$, as $\bar{\varepsilon}$ increases, SP_2 increases its advertised privacy risk to exploit more private information from consumer. Despite this, the market share of SP_2 increases with $\bar{\varepsilon}$. This is because as $\bar{\varepsilon}$ increases, SP_2 provides a higher utility than SP_1 to some consumers who prefer SP_1 before. Furthermore, each SP's profit increases as $\bar{\varepsilon}$ increases.

2. *Operation cost* (c): when c increases, by (4.24), the SP with lower privacy-independent revenue benefits since its market share increases. Observe from (4.17) and (4.18) that if $p_2 - p_1 > 0$, both SPs increase their privacy risk strategies in the SPNE as c increases. They do so because SPs can use consumers' private information to increase its privacy dependent revenue, thereby offsetting their cost. Otherwise, they decrease their privacy risks. As a result of these strategies, when c increases, both SPs' profits also increase.

3. *Privacy independent revenue* (p_1, p_2): as the difference in the privacy-independent revenues ($p_2 - p_1$) increases, both SPs offer lower privacy risks to attract consumers. From (4.24)–(4.26) and condition (4.21), we see that as $p_2 - p_1$ increases, the market share and profit of SP_1 decreases while SP_2's market share and profit increases. This is because a larger difference in the revenue independent of consumer's private data gives SP_2 more market power in the competition. As a result, SP_2's profit increases while SP_1's profit decreases.

4. *Consumer privacy valuation or skittishness* (t): when t increases, by (4.17)–(4.20), we have $\varepsilon_1^* = \frac{12\bar{\varepsilon}c\alpha - 16(p_2 - p_1)}{24tc} - \frac{1}{8}\bar{\varepsilon} \geq 0$. Therefore, both SPs decrease their

privacy risks as t increases. Furthermore, by (4.24), the SP with lower privacy-independent revenue benefits since its market share increases. For this linear model considered, as t increases, both SPs decrease their privacy risks. This results in a decrease in the cost and revenue of both SPs but cost supersedes revenue, thereby leading to a profit for both SPs. In other words, a higher privacy valuation from consumers "softens" the competition between SPs.

4.3.2 Truncated Gaussian Consumer Privacy Risk Tolerance

In this section, we model a consumer's privacy tolerance as a random variable E that follows a Gaussian distribution $\mathcal{N}(\frac{\bar{\varepsilon}}{2}, \sigma^2)$ with a mean of $\frac{\bar{\varepsilon}}{2}$ and a standard deviation of σ. Since $E \in [0, \bar{\varepsilon}]$, we restrict the Gaussian distribution to lie within the interval $[0, \bar{\varepsilon}]$. Thus, E follows a truncated Gaussian distribution with CDF

$$F_E(\varepsilon) = \begin{cases} \frac{\Phi(\frac{\varepsilon - \frac{\bar{\varepsilon}}{2}}{\sigma}) - \Phi(-\frac{\bar{\varepsilon}}{2\sigma})}{\Phi(\frac{\bar{\varepsilon}}{2\sigma}) - \Phi(-\frac{\bar{\varepsilon}}{2\sigma})} & \varepsilon \in [0, \bar{\varepsilon}] \\ 0 & \varepsilon \in [-\infty, 0] \\ 1 & \varepsilon \in [\bar{\varepsilon}, +\infty] \end{cases}, \tag{4.27}$$

where $\Phi(y)$ denotes the CDF of the standard Gaussian distribution.

In contrast to the uniform distribution case, the CDF in (4.27) is not amenable to a closed form solution. Thus, we characterize the equilibrium numerically. To find the SPNE, we first compute the SPNE QoS in the second stage as functions of privacy risk guarantees by solving (4.12). Then, we use an iterated best response method to find the optimal privacy risk guarantee of an SP by fixing its competitor's strategy in each iteration. When the process converges, we have found an SPNE in which no SP is better off by unilaterally deviating from the equilibrium.

4.3.3 Illustration of Results

In this section, we illustrate our model and results. First, we assume consumers have uniformly distributed privacy risk tolerance. We plot each SP's SPNE strategy,

market share, and total profit $w.r.t$ consumers' maximum privacy risk tolerance $\bar{\varepsilon}$ for different values of consumer privacy risk valuation t. Later, we study the model in which consumers' privacy risk tolerance follows a Gaussian distribution $\mathcal{N}(\frac{\bar{\varepsilon}}{2}, 1)$ truncated between 0 and $\bar{\varepsilon}$. The model parameters are given as follows:

Parameter	c	λ	r	p_1	p_2
Value	0.5	0.75	0.7	0.4	0.8

Table 4.1: Numerical example model parameters

Consumers with Uniformly Distributed Privacy Risk Tolerance

In this section, we vary $\bar{\varepsilon}$ from 3 to 5 to study properties of SPNE. Our choice of values in Table 4.1 is one set of parameters for which we can determine a meaningful range of t values. However, there exists many such combinations of parameter values. Note that by (4.21)–(4.23) in Theorem 5, t must belong to $[0.58, 0.85]$ when other parameters are given in Table 4.1 for sustaining the SPNE. In Figure 4.3, the equilibrium strategies of different SPs are plotted. As $\bar{\varepsilon}$ increases, both SPs increase their privacy risk offerings. Furthermore, it can be seen that as t, the valuation of privacy by consumer, decreases, each SP increases its privacy risk to generate more profit from using private data. Correspondingly, the SPs will have to provide higher QoS to attract consumers. On the other hand, if t increases, both SPs reduce their privacy risks to avoid violating consumers' privacy.

It is worth noting that for the special case of $t = 0.7$, we observe that SP_1 caters to smaller set of privacy sensitive consumers. The reason for this is as follows: indeed, one generally expects SP_1 to offer a larger privacy risk as $\bar{\varepsilon}$ increases. However, for a large enough privacy valuation (in this case $t = 0.7$), what we observe is that

since consumers highly value privacy, the cost of offering a high QoS proportionally increases for SP_1. The resulting profit is insufficient to justify the cost.

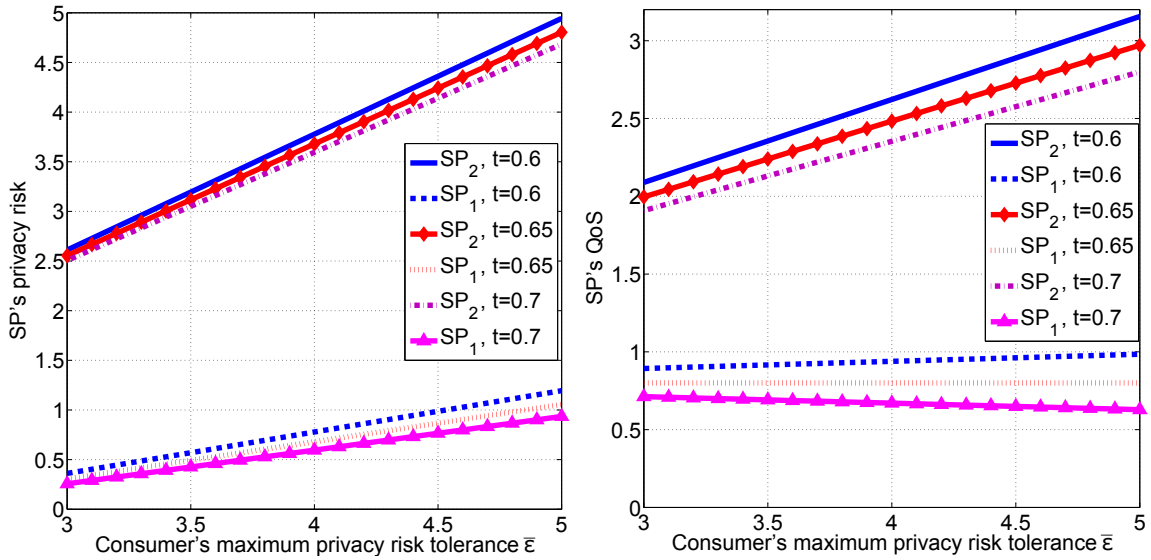

(a) Privacy risk (b) QoS

Figure 4.3: SPNE strategies of SPs under uniform consumer privacy risk

The market shares of different SPs in the SPNE are presented in Figure 4.4. We observe that the equilibrium market share of SP_2 decreases as t increases. The intuition behind this is that if t increases, the consumer's valuation of privacy mismatch also increases. Thus, it is more difficult for SP_2 to attract consumers with privacy tolerance lower than ε_2. As a result, its market share decreases. Notice that in Figure 4.4, as $\bar{\varepsilon}$ decreases, the equilibrium market share of SP_2 increases. This is because consumers experience a lower negative utility from the mismatch between their preferred and the offered privacy risk when the net range is smaller (recall that the utility from mismatch is given by $t(x - x_i)\varepsilon_i, \varepsilon_i \in [0, \bar{\varepsilon}]$). As a result, more consumers will choose the SP with a higher privacy risk to enjoy a higher QoS.

In Figure 4.5, we plot the total profit at the SPNE for each SP as a function of the maximum consumer privacy risk tolerance $\bar{\varepsilon}$ for different values of t. As shown in the

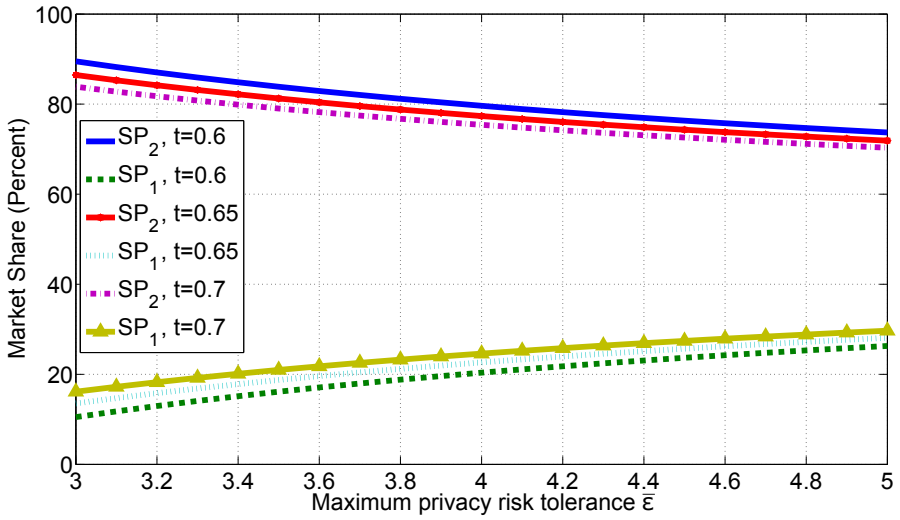

Figure 4.4: Market shares of SPs at SPNE under uniform consumer privacy risk

figure, the total profit of both SPs at SPNE increase as $\bar{\varepsilon}$ increases. This is due to the fact that a larger $\bar{\varepsilon}$ indicates a larger range of consumer preferences, and then, more possibilities for the SPs to exploit private information. Thus, both SPs can benefit from using private data of consumers that have a higher privacy risk tolerance. As t increases, both SPs decrease their privacy risks. As a result, the cost and revenue of both SPs decrease. However, in this case, cost supersedes revenue. Therefore, both SPs make more profit. In other words, a higher privacy valuation from consumers "softens" the competition.

Truncated Gaussian Consumer Privacy Risk Tolerance

We now consider the case in which consumers' privacy risk tolerance follows a truncated Gaussian distribution with a mean of $\frac{\bar{\varepsilon}}{2}$ and a standard deviation of 1. The equilibrium strategies of different SPs are shown in Figure 4.6. As with the uniform distribution scenario, here too we observe that the privacy risk and the QoS offered by each SP are linear functions of $\bar{\varepsilon}$. We also notice that in this SPNE, SP_2 will always provide service with maximum privacy risk (Figure 4.6a) for the set of parameters in

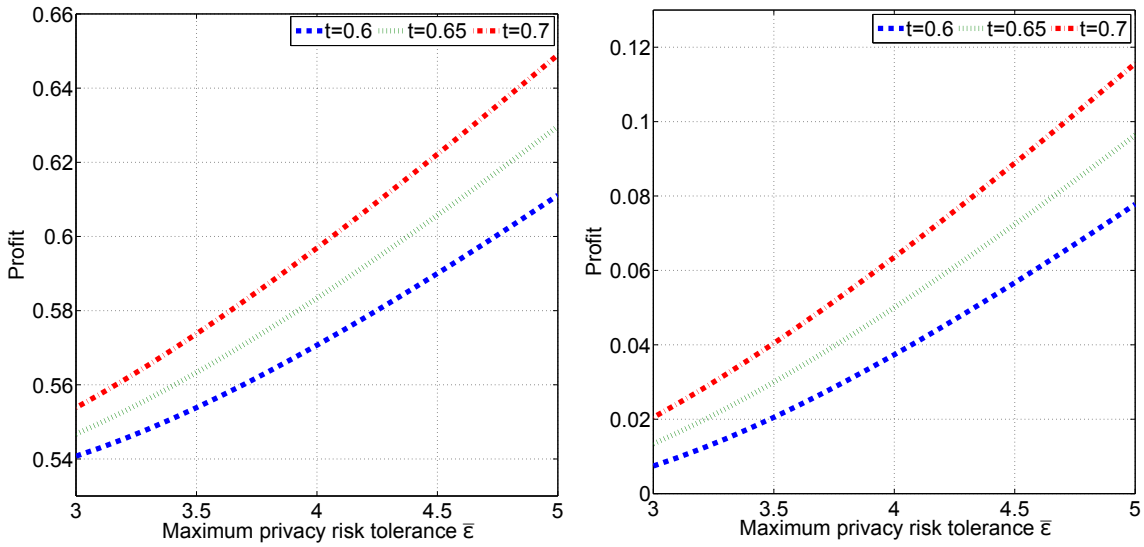

(a) SP_2 (b) SP_1

Figure 4.5: Profit of SPs at SPNE under uniform consumer privacy risk

Table 4.1. This is because in contrast to the uniform distribution, for the truncated Gaussian distribution, there are a relatively smaller number of consumers concentrated in the tail end of $[0, \bar{\varepsilon}]$. Thus, for SP_1 to make a profit, it has to offer a higher privacy risk so that it can capture a large number of consumers in the middle of the $[0, \bar{\varepsilon}]$ range. This in turn forces SP_2 to increase to its privacy risk to differentiate its QoS offering and thus have a higher profit. Since the privacy risk preference is bound by $[0, \bar{\varepsilon}]$, SP_2 can only offer the highest privacy risk in this example.

Figure 4.7 shows market shares of different SPs at SPNE vs. consumers' maximum privacy risk tolerance for different values of t under truncated Gaussian privacy tolerance distribution. As t decreases, the market share of SP_2 at SPNE increases, and vice versa. Also, when $\bar{\varepsilon}$ decreases, the equilibrium market share of SP_2 increases. Furthermore, it can be seen that for the same $\bar{\varepsilon}$, the market share of SP_2 (SP_1) is smaller (larger) when consumers' privacy tolerance follows the truncated Gaussian distribution compared to uniform distribution. Our numerical analysis shows that at SPNE, SP_2 is forced to provide service with maximum privacy risk. We argue that

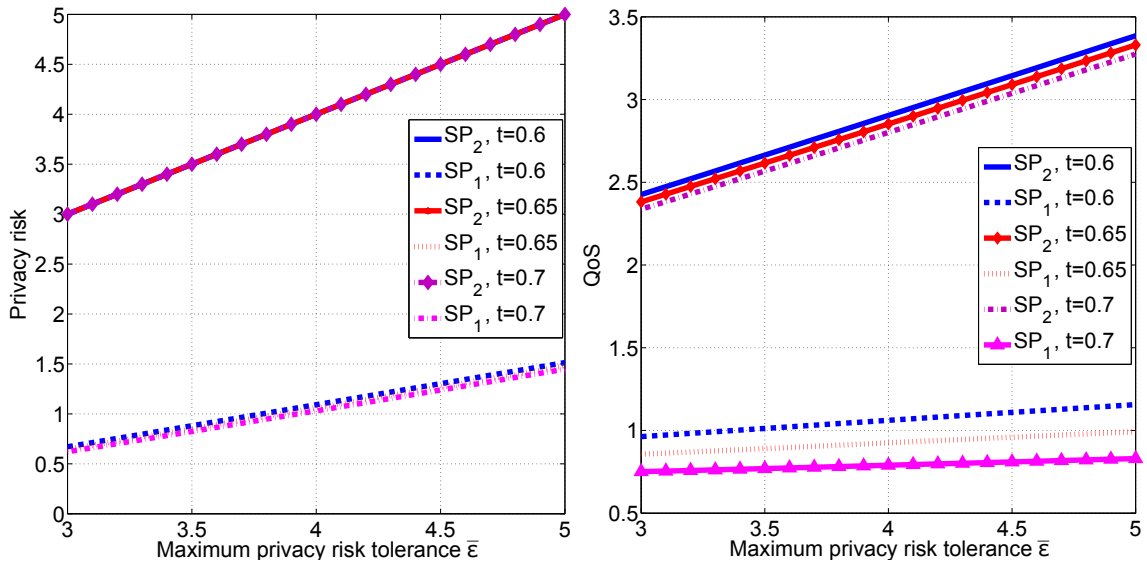

(a) Privacy risk (b) QoS

Figure 4.6: SPNE strategies of SPs under truncated Gaussian consumer privacy risk

this is due to the shape of the distribution that limits the number of consumers at the two extremes thus compelling the two SPs to compete for the large bulk of consumers distributed around $\bar{\varepsilon}/2$. Given the ability of SP_2 to make more profit on untargeted services relative to SP_1, the SPNE solution leads to SP_1 increasing its market share to be profitable and SP_2 achieving profitability with a smaller market share.

The relationship between total profit of different SPs at SPNE vs. consumers' maximum privacy risk tolerance for different values of t is shown in Figure 4.8. Similar to Figure 4.5, both SPs' total profit increase as $\bar{\varepsilon}$ increases. However, in contrast to Figure 4.5, as t decreases, the total profit of SP_2 increases. This is because SP_2 always offers $\bar{\varepsilon}$ in the SPNE. Notice that SP_2's equilibrium QoS is also a linear function of $\bar{\varepsilon}$ (see Figure 4.6b). On the other hand, SP_2's market share increases as t decreases (see Figure 4.7). By (4.4), (4.14), and (4.15); the total profit of SP_2 increases as t decreases.

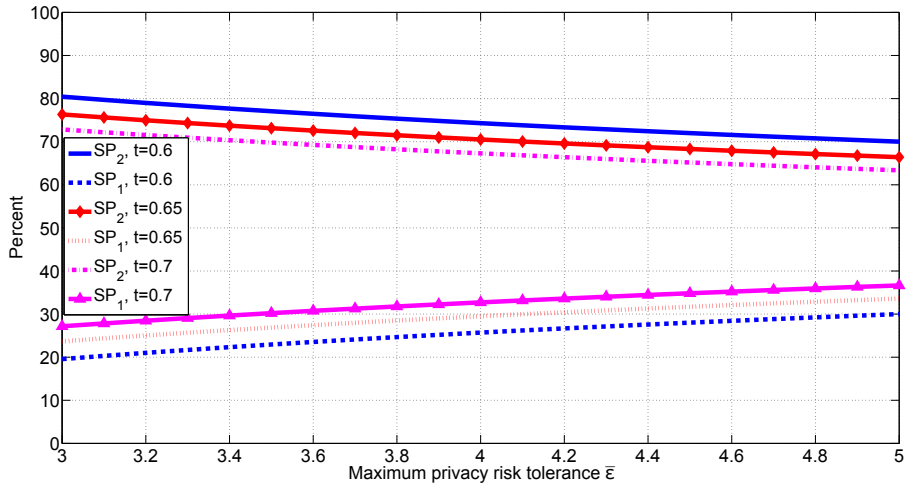

Figure 4.7: Market shares of SPs at SPNE under truncated Gaussian consumer privacy risk

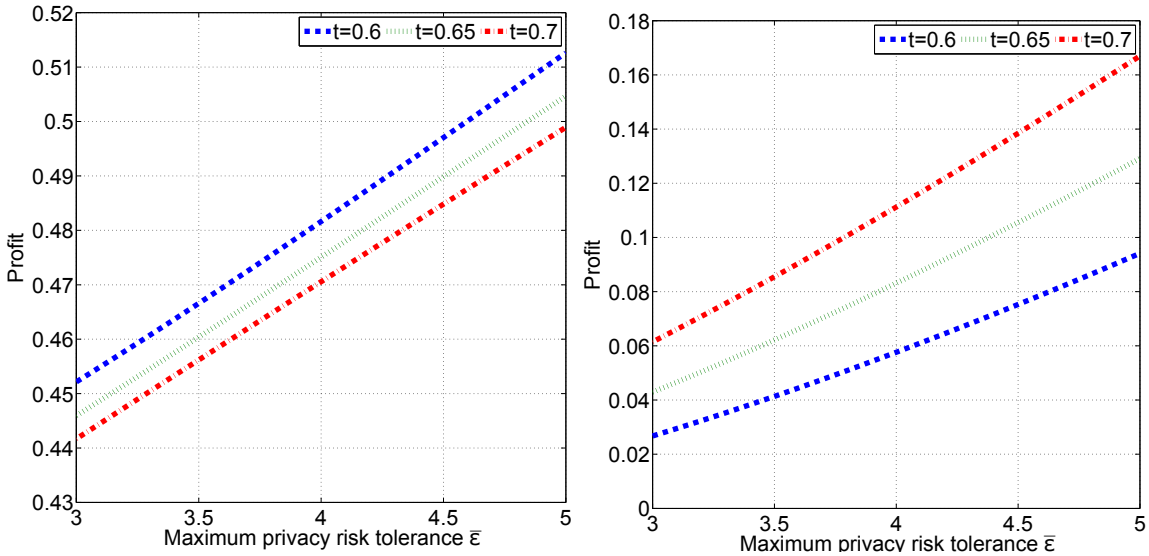

(a) SP_2 (b) SP_1

Figure 4.8: Profit of SPs at SPNE under truncated Gaussian consumer privacy risk

4.4 Market with Multiple Service Providers

In the previous sections, we studied the market with two SPs. In this section, we examine a generalized model with multiple SPs (Figure 4.9). We allow for a finitely arbitrary number of SPs, each of which offers the same type of free service but with different QoS and privacy risk guarantee to consumers. In particular, we assume there

are m SPs in the market. Our models for cost, revenue and utility for each SP as well as the consumers are the same as for the two-SP model described in Section 4.1.1. Furthermore, we assume a consumer's privacy risk tolerance is uniformly distributed between $[0, \varepsilon]$. Analogous to the two SP model, the interactions between the m SPs and consumers can also be viewed as a non-cooperative sequential game. The m-SP game proceeds in three stages.

In the first stage, each of the m SPs chooses its own privacy risk guarantee resulting in a vector $\boldsymbol{\varepsilon} = (\varepsilon_1, \varepsilon_2, ..., \varepsilon_m)$ (on the interval $[0, \bar{\varepsilon}]$). Without loss of generality, we assume $\varepsilon_1 \leq \varepsilon_2 \leq ... \leq \varepsilon_m$. At the second stage, given the privacy risk $\boldsymbol{\varepsilon}$ determined in the first stage, the SPs simultaneously determine their QoS values to obtain a vector $\boldsymbol{v} = \{v_1, v_2, ..., v_m\}$. At the last stage, each consumer chooses the SP that yields the maximal perceived utility for the consumer.

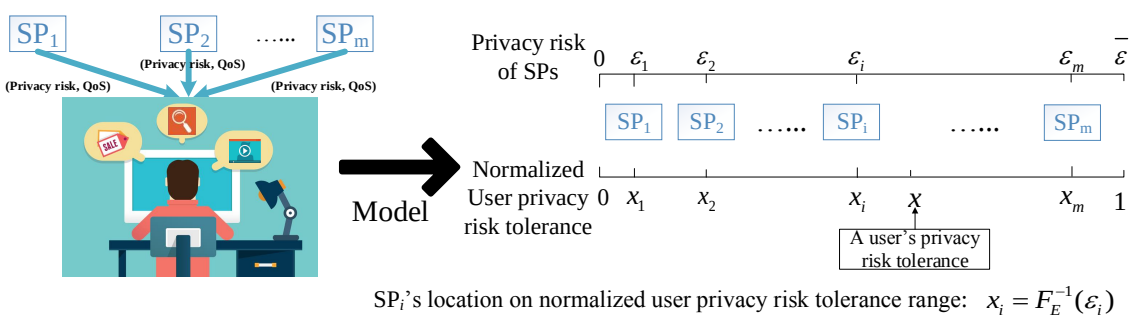

Figure 4.9: Market model for multiple SPs offering services with privacy guarantee

To find the SPNE, we apply backward induction to the three stage game described above as follows. In the last stage of the game, for fixed QoS and privacy risk guarantee strategies of the SPs, consumers' choices of SPs are determined by their privacy risk tolerances. In the two-SP case, the consumer located at x_τ divides the set of consumers into two convex subsets where the consumers in the left subset will choose SP_1 and vice versa. However, for the multiple SP case, the market share of SP_i ($i \in \{1, ..., m\}$) is not necessarily a convex set between the indifference threshold in which consumers are indifferent to choosing SP_{i-1} or SP_i and the threshold in

which consumers are indifferent to choosing SP_i or SP_{i+1}. This is due to the fact that in general the problem requires each SP_i to compete with all other SPs, even if their privacy risk offerings are very different (e.g., SPs with a large difference in locations in Figure 4.9). We note that this will not happen in the equilibrium since an SP with zero market share would be better off by either improving its QoS to attract some consumers or just exit the market. Therefore, in the equilibrium, each SP only competes directly with its two closest neighbors. For given QoS profile $\boldsymbol{v} = \{v_1, ..., v_m\}$ and privacy risk profile $\boldsymbol{\varepsilon} = (\varepsilon_1, ..., \varepsilon_m)$, the market segmentation is

$$
\begin{aligned}
n_1 &= \frac{v_1 - v_2 + t(F_E(\varepsilon_2)\varepsilon_2 - F_E(\varepsilon_1)\varepsilon_1)}{t(\varepsilon_2 - \varepsilon_1)}, \\
n_i &= \frac{v_i - v_{i+1} + t(F_E(\varepsilon_{i+1})\varepsilon_{i+1} - F_E(\varepsilon_i)\varepsilon_i)}{t(\varepsilon_{i+1} - \varepsilon_i)} - \frac{v_{i-1} - v_i + t(F_E(\varepsilon_i)\varepsilon_i - F_E(\varepsilon_{i-1})\varepsilon_{i-1})}{t(\varepsilon_i - \varepsilon_{i-1})}, \\
&\qquad i \in \{2, ..., m-1\}, \\
n_m &= 1 - \frac{v_{m-1} - v_m + t(F_E(\varepsilon_m)\varepsilon_N - F_E(\varepsilon_{m-1})\varepsilon_{m-1})}{t(\varepsilon_m - \varepsilon_{m-1})}.
\end{aligned}
$$

Furthermore, we define the objective function of SP_i to be

$$\pi_i(\boldsymbol{\varepsilon}; \boldsymbol{v}) = [R(\varepsilon_i) - C(v_i; \varepsilon_i)]n_i(\boldsymbol{\varepsilon}; \boldsymbol{v}), i \in \{1, ..., m\}.$$

For a given privacy risk guarantee profile $\boldsymbol{\varepsilon}$, the optimal QoS of SP_i ($i \in \{1, ..., m\}$) is determined by

$$arg \max_{v_i} \pi_i(\boldsymbol{\varepsilon}; \boldsymbol{v}), i \in \{1, ..., m\} \tag{4.28}$$

while fixing all other players' strategies.

We note that the cost function $C(v_i; \varepsilon_i)$ and the market segmentation computed in the first stage are both linear functions of v_i. Thus, for a fixed privacy risk guarantee profile $\boldsymbol{\varepsilon}$, the objective function of SP_i in this stage is a concave function $w.r.t$ its own strategy v_i. Furthermore, the feasible set of each SP's strategy is a convex set. Thus, the non-cooperative game among the SPs in this stage is a m-player concave game.

By [135], there exists a Nash equilibrium. We define $\delta_i \triangleq \frac{1}{t(\varepsilon_{i+1}-\varepsilon_i)}$, $y_1 \triangleq r(\varepsilon_1)+p_1-c\lambda\varepsilon_1-ctx_2\varepsilon_2+ctx_1\varepsilon_1$, $y_N \triangleq r(\varepsilon_N)+p_N-c\lambda\varepsilon_N-ct(1-x_N)\varepsilon_N+ct(1-x_{N-1})\varepsilon_{N-1}$ and $y_i \triangleq \frac{r(\varepsilon_i)+p_i-c\lambda\varepsilon_i+ctx_i\varepsilon_i-ctx_{i+1}\varepsilon_{i+1}}{t(\varepsilon_{i+1}-\varepsilon_i)} + \frac{r(\varepsilon_i)+p_i-c\lambda\varepsilon_i-ctx_{i-1}\varepsilon_{i-1}+ctx_i\varepsilon_i}{t(\varepsilon_i-\varepsilon_{i-1})} \quad \forall i \in \{2,...,m\}$. Applying the first order condition to SPs' profit functions (solving simultaneous linear equations obtained from $\frac{\partial \pi_i(\boldsymbol{\varepsilon};\boldsymbol{v})}{\partial v_i}=0, i \in \{1,2,...,m\}$) yields the equilibrium strategies

$$v_1^* = \frac{v_2}{2} + \frac{y_1}{2c}, \tag{4.29}$$

$$v_i^* = \frac{cv_{i+1}\delta_i + cv_{i-1}\delta_{i-1} + y_i}{2c[\delta_i + \delta_{i-1}]}, i \in \{2,..,m\}, \tag{4.30}$$

$$v_m^* = \frac{v_{m-1}}{2} + \frac{y_m}{2c}. \tag{4.31}$$

In the last stage, the SPs determine their privacy risk guarantees $\boldsymbol{\varepsilon}$ by considering equilibrium strategies in previous stages (n_i and $v_i^* \quad \forall i \in \{1,...,m\}$) as functions of $\boldsymbol{\varepsilon}$. Therefore, the optimal privacy risk strategy of SP_i is determine by

$$arg \max_{\varepsilon_i} \pi_i(\boldsymbol{\varepsilon};\boldsymbol{v}), i \in \{1,...,m\} \tag{4.32}$$

while fixing all other players' strategies. For reasons of intractability (solving highly parameterized high order polynomial equations), a full characterization of privacy risk equilibria could not be achieved. Thus, we characterize the SPNE numerically by using the iterated best response method. We consider a three-SP market and adopt the model parameters presented in Table 4.1. Furthermore, we assume $t = 0.7$ and $\bar{\varepsilon} = 5$. The initial privacy risk of SP_i is given by $\frac{i\bar{\varepsilon}}{i+1}$ for $i \in \{1,2,3\}$. Although there exists an SPNE in the second stage of the sequential game for fixed privacy guarantees, the existence of an equilibrium in the first stage can not be guaranteed.

The best response strategies of the SPs for different values of SP_2's privacy independent revenue are plotted in Figure 4.10. It can be seen that the two SPs with lower privacy risks proceed to jump over each other in each round of best response iteration, attempting to lower its privacy risk to attract more consumers from its competitor.

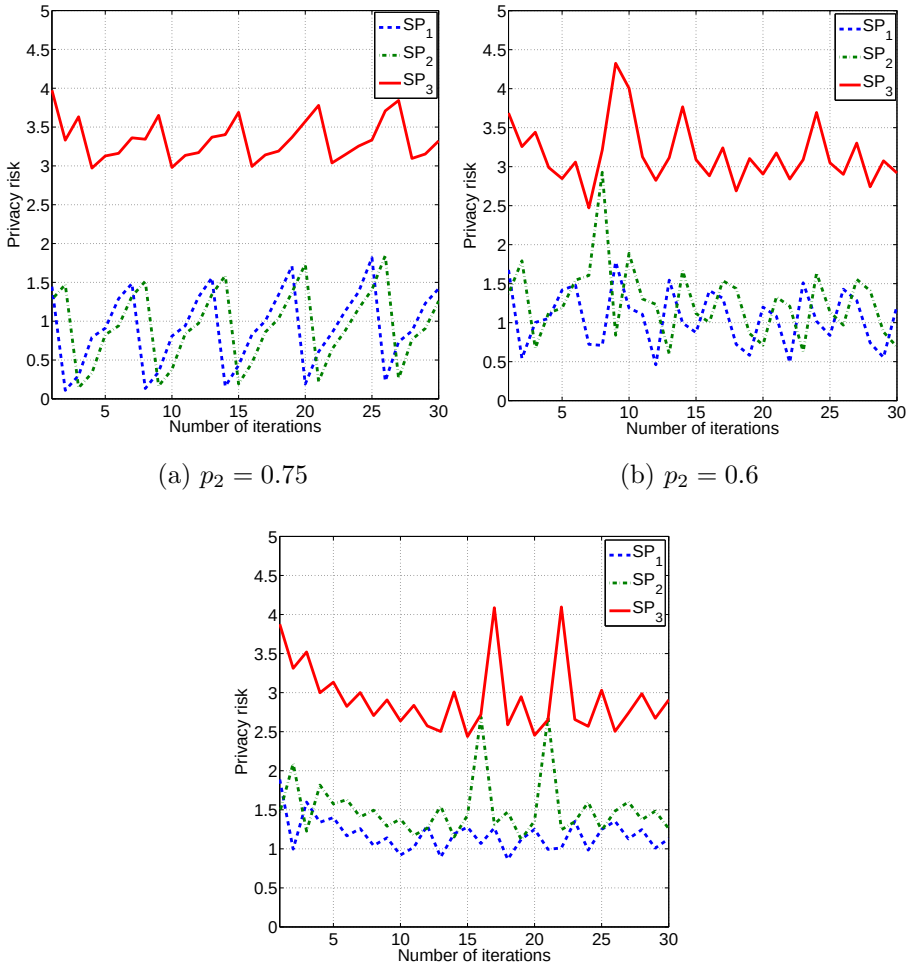

(a) $p_2 = 0.75$

(b) $p_2 = 0.6$

(c) $p_2 = 0.45$

Figure 4.10: Best response of each SP's privacy risk for different values of SP_2's revenue independent of using private data

The SP with the highest private data independent revenue adopts a high privacy risk strategy to focus on consumers with high privacy risk tolerance and exploiting their private data extensively. We observe that when p_2 is large, SP_3's privacy risk strategy is also higher on average. On the other hand, SP_1's best response strategy is lower. The intuition behind is that a larger p_2 allows SP_2 to set a higher privacy risk to make more profit from using consumer data. This forces SP_3 to increase its privacy risk to differentiate itself from SP_2. On the other hand, a higher privacy risk of SP_2 will encourage SP_1 to lower its privacy risk to attract more consumers.

Chapter 5

GENERATIVE ADVERSARIAL PRIVACY

5.1 Generative Adversarial Privacy Model

We consider a dataset $\mathcal{D}$ which contains both public and private variables for n individuals (see Figure 1.1). We represent the public variables by a random variable $X \in \mathcal{X}$, and the private variables (which are typically correlated with the public variables) by a random variable $Y \in \mathcal{Y}$. Each dataset entry contains a pair of public and private variables denoted by (X, Y). Instances of X and Y are denoted by x and y, respectively. We assume that each entry pair (X, Y) is distributed according to $P(X, Y)$, and is independent from other entry pairs in the dataset. Since the dataset entries are independent of each other, we restrict our attention to memoryless mechanisms: privacy mechanisms that are applied on each data entry separately. Formally, we define the privacy mechanism as a randomized mapping given by

$$g(X, Y) : \mathcal{X} \times \mathcal{Y} \to \mathcal{X}.$$

We consider two different types of privatization schemes: (a) private data dependent (PDD) schemes, and (b) private data independent (PDI) schemes. A privatization mechanism is PDD if its output is dependent on both Y and X. It is PDI if its output only depends on X. PDD mechanisms are naturally superior to PDI mechanisms. We show, in sections 5.2 and 5.3, that there is a sizeable gap in performance between these two approaches.

In our proposed GAP framework, the privatizer is pitted against an adversary. We model the interactions between the privatizer and the adversary as a non-cooperative game. For a fixed g, the goal of the adversary is to reliably infer Y from $g(X, Y)$ using a strategy h. For a fixed adversarial strategy h, the goal of the privatizer is to design g in a way that minimizes the adversary's capability of inferring the private

variable from the perturbed data. The optimal privacy mechanism is obtained as an equilibrium point at which both the privatizer and the adversary can not improve their strategies by unilaterally deviating from the equilibrium point.

5.1.1 Formulation

Given the output $\hat{X} = g(X, Y)$ of a privacy mechanism $g(X, Y)$, we define $\hat{Y} = h(g(X, Y))$ to be the adversary's inference of the private variable Y from $\hat{X}$. To quantify the effect of adversarial inference, for a given public-private pair (x, y), we model the loss of the adversary as

$$\ell(h(g(X = x, Y = y)), Y = y) : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}.$$

Therefore, the expected loss of the adversary *w.r.t.* X and Y is defined to be

$$L(h, g) \triangleq \mathbb{E}[\ell(h(g(X, Y)), Y)], \tag{5.1}$$

where the expectation is taken over $P(X, Y)$ and the randomness in g and h.

Intuitively, the privatizer would like to minimize the adversary's ability to learn Y reliably from the published data. This can be trivially done by releasing an $\hat{X}$ independent of X. However, such an approach provides no utility for data analysts who want to learn non-private variables from $\hat{X}$. To overcome this issue, we capture the loss incurred by privatizing the original data via a distortion function $d(\hat{x}, x) : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$, which measures how far the original data $X = x$ is from the privatized data $\hat{X} = \hat{x}$. Thus, the average distortion under $g(X, Y)$ is $\mathbb{E}[d(g(X, Y), X)]$, where the expectation is taken over $P(X, Y)$ and the randomness in g.

On the one hand, the data holder would like to find a privacy mechanism g that is both privacy preserving (in the sense that it is difficult for the adversary to learn Y from $\hat{X}$) and utility preserving (in the sense that it does not distort the original data too much). On the other hand, for a fixed choice of privacy mechanism g, the

adversary would like to find a (potentially randomized) function h that minimizes its expected loss, which is equivalent to maximizing the negative of the expected loss. To achieve these two opposing goals, we model the problem as a constrained minimax game between the privatizer and the adversary:

$$\begin{aligned} \min_{g(\cdot)} \max_{h(\cdot)} \quad & -L(h,g) \\ s.t. \quad & \mathbb{E}[d(g(X,Y),X)] \leq D, \end{aligned} \tag{5.2}$$

where the constant $D \geq 0$ determines the allowable distortion for the privatizer and the expectation is taken over $P(X,Y)$ and the randomness in g and h.

5.1.2 GAP under Various Loss Functions

The above formulation places no restrictions on the adversary. Indeed, different loss functions and decision rules lead to different adversarial models. In what follows, we will discuss a variety of loss functions under hard and soft decision rules, and show how our GAP framework can recover several popular information theoretic privacy notions.

Hard Decision Rules. When the adversary adopts a hard decision rule, $h(g(X,Y))$ is an estimate of Y. Under this setting, we can choose $\ell(h(g(X,Y)),Y)$ in a variety of ways. For instance, if Y is continuous, the adversary can attempt to minimize the difference between the estimated and true private variable values. This can be achieved by considering a squared loss function

$$\ell(h(g(X,Y)),Y) = (h(g(X,Y)) - Y)^2, \tag{5.3}$$

which is known as the ℓ_2 loss. In this case, one can verify that the adversary's optimal decision rule is $h^* = \mathbb{E}[Y|g(X,Y)]$, which is the conditional mean of Y given $g(X,Y)$. Furthermore, under the adversary's optimal decision rule, the minimax problem in

(5.2) simplifies to

$$\min_{g(\cdot)} -\text{mmse}(Y|g(X,Y)) = -\max_{g(\cdot)} \text{mmse}(Y|g(X,Y)),$$

subject to the distortion constraint. Here $\text{mmse}(Y|g(X,Y))$ is the resulting minimum mean square error (MMSE) under $h^* = \mathbb{E}[Y|g(X,Y)]$. Thus, under the ℓ_2 loss, GAP provides privacy guarantees against an MMSE adversary. On the other hand, when Y is discrete (e.g., age, gender, political affiliation, etc), the adversary can attempt to maximize its classification accuracy. This is achieved by considering a 0-1 loss function [136]

$$\ell(h(g(X,Y)),Y) = \begin{cases} 0 & \text{if } h(g(X,Y)) = Y \\ 1 & \text{otherwise} \end{cases}. \tag{5.4}$$

In this case, one can verify that the adversary's optimal decision rule is the *maximum a posteriori probability* (MAP) decision rule: $h^* = \text{argmax}_{y\in\mathcal{Y}} P(y|g(X,Y))$, with ties broken uniformly at random. Moreover, under the MAP decision rule, the minimax problem in (5.2) reduces to

$$\min_{g(\cdot)} -(1 - \max_{y\in\mathcal{Y}} P(y, g(X,Y))) = \min_{g(\cdot)} \max_{y\in\mathcal{Y}} P(y, g(X,Y)) - 1, \tag{5.5}$$

subject to the distortion constraint. Thus, under a 0-1 loss function, the GAP formulation provides privacy guarantees against a MAP adversary.

Soft Decision Rules. Instead of a *hard decision* rule, we can also consider a broader class of *soft decision* rules where $h(g(X,Y))$ is a distribution over $\mathcal{Y}$; i.e., $h(g(X,Y)) = P_h(y|g(X,Y))$ for $y \in \mathcal{Y}$. In this context, we can analyze the performance under a log-loss

$$\ell(h(g(X,Y)),y) = \log \frac{1}{P_h(y|g(X,Y))}. \tag{5.6}$$

In this case, the objective of the adversary simplifies to

$$\max_{h(\cdot)} -\mathbb{E}[\log \frac{1}{P_h(y|g(X,Y))}] = -H(Y|g(X,Y)),$$

and that the maximization is attained at $P_h^*(y|g(X,Y)) = P(y|g(X,Y))$. Therefore, the optimal adversarial decision rule is determined by the true conditional distribution $P(y|g(X,Y))$, which we assume is known to the data holder in the game-theoretic setting. Thus, under the log-loss function, the minimax optimization problem in (5.2) reduces to

$$\min_{g(\cdot)} -H(Y|g(X,Y)) = \min_{g(\cdot)} I(g(X,Y);Y) - H(Y),$$

subject to the distortion constraint. Thus, under the log-loss in (5.6), GAP is equivalent to using MI as the privacy metric [86].

The 0-1 loss captures a strong guessing adversary; in contrast, log-loss or information-loss models a belief refining adversary. Next, we consider a more general α-loss function [137] that allows continuous interpolation between these extremes via

$$\ell(h(g(X,Y)),y) = \frac{\alpha}{\alpha-1}\left(1 - P_h(y|g(X,Y))^{1-\frac{1}{\alpha}}\right), \tag{5.7}$$

for any $\alpha > 1$. As shown in [137], for very large α ($\alpha \to \infty$), this loss approaches that of the 0-1 (MAP) adversary. As α decreases, the convexity of the loss function encourages the estimator $\hat{Y}$ to be probabilistic, as it increasingly rewards correct inferences of lesser and lesser likely outcomes (in contrast to a hard decision rule by a MAP adversary of the most likely outcome) conditioned on the revealed data. As $\alpha \to 1$, (5.7) yields the logarithmic loss, and the optimal belief $P_{\hat{Y}}$ is simply the posterior belief. Denoting $H_\alpha^{\mathrm{a}}(Y|g(Y,X))$ as the Arimoto conditional entropy of order α, one can verify that [137]

$$\max_{h(\cdot)} -\mathbb{E}\left[\frac{\alpha}{\alpha-1}\left(1 - P_h(y|g(X,Y))^{1-\frac{1}{\alpha}}\right)\right] = -H_\alpha^{\mathrm{a}}(Y|g(X,Y)),$$

which is achieved by a 'α-tilted' conditional distribution [137]

$$P_h^*(y|g(X,Y)) = \frac{P(y|g(X,Y))^\alpha}{\sum_{y\in\mathcal{Y}} P(y|g(X,Y))^\alpha}.$$

Under this choice of a decision rule, the objective of the minimax optimization in (5.2) reduces to

$$\min_{g(\cdot)} -H_\alpha^{\mathrm{a}}(Y|g(X,Y)) = \min_{g(\cdot)} I_\alpha^{\mathrm{a}}(g(X,Y);Y) - H_\alpha(Y), \tag{5.8}$$

where I_α^{a} is the Arimoto mutual information and H_α is the Rényi entropy. Note that as $\alpha \to 1$, we recover the classical MI privacy setting and when $\alpha \to \infty$, we recover the 0-1 loss.

5.1.3 Data-driven GAP

So far, we have focused on a setting where the data holder has access to $P(X,Y)$. When $P(X,Y)$ is known, the data holder can simply solve the constrained minimax optimization problem in (5.2) (theoretical version of GAP) to obtain a privatization mechanism that would perform best against a chosen type of adversary. In the absence of $P(X,Y)$, we propose a data-driven version of GAP that allows the data holder to learn privatization mechanisms directly from a dataset of the form $\mathcal{D} = \{(x_{(i)}, y_{(i)})\}_{i=1}^n$. Under the data-driven version of GAP, we represent the privacy mechanism via a conditional generative model $g(X,Y;\Theta_p)$ parameterized by Θ_p. This generative model takes (X,Y) as inputs and outputs $\hat{X}$. In the training phase, the data holder learns the optimal parameters Θ_p by competing against a *computational adversary*: a classifier modeled by a neural network $h(g(X,Y;\Theta_p);\Theta_a)$ parameterized by Θ_a. After convergence, we evaluate the performance of the learned $g(X,Y;\Theta_p^*)$ by computing the maximal probability of inferring Y under the MAP adversary studied in the theoretical version of GAP.

We note that in theory, the functions h and g can (in general) be arbitrary; i.e., they can capture all possible learning algorithms. However, in practice, we need to

restrict them to a rich hypothesis class. Figure 5.1 shows an example of the GAP model in which the privatizer and adversary are modeled as multi-layer "randomized" neural networks. For a fixed h and g, we quantify the adversary's *empirical loss* using a continuous and differentiable function

$$L_{\text{EMP}}(\Theta_p, \Theta_a) = \frac{1}{n}\sum_{i=1}^{n} \ell(h(g(x_{(i)}, y_{(i)}; \Theta_p); \Theta_a), y_{(i)}), \tag{5.9}$$

where $(x_{(i)}, y_{(i)})$ is the ith row of $\mathcal{D}$ and $\ell(h(g(x_{(i)}, y_{(i)}; \Theta_p); \Theta_a), y_{(i)})$ is the adversary loss in the data-driven context. The optimal parameters for the privatizer and adversary are the solution to

$$\begin{aligned} \min_{\Theta_p} \max_{\Theta_a} \quad & -L_{\text{EMP}}(\Theta_p, \Theta_a) \\ s.t. \quad & \mathbb{E}_{\mathcal{D}}[d(g(X, Y; \Theta_p), X)] \leq D, \end{aligned} \tag{5.10}$$

where the expectation is taken over the dataset $\mathcal{D}$ and the randomness in g.

In keeping with the now common practice in machine learning, in the data-driven approach for GAP, one can use the empirical log-loss function [138, 139] given by (5.9) with

$$\begin{aligned} \ell(h(g(x_{(i)}, y_{(i)}; \Theta_p); \Theta_a), y_{(i)}) = & -y_{(i)} \log h(g(x_{(i)}, y_{(i)}; \Theta_p); \Theta_a) \\ & - (1 - y_{(i)}) \log(1 - h(g(x_{(i)}, y_{(i)}; \Theta_p); \Theta_a)), \end{aligned} \tag{5.11}$$

which leads to a minimum cross-entropy adversary. As a result, the empirical loss of the adversary is quantified by the cross-entropy

$$\begin{aligned} L_{\text{XE}}(\Theta_p, \Theta_a) = -\frac{1}{n}\sum_{i=1}^{n} [& y_{(i)} \log h(g(x_{(i)}, y_{(i)}; \Theta_p); \Theta_a) \\ & + (1 - y_{(i)}) \log(1 - h(g(x_{(i)}, y_{(i)}; \Theta_p); \Theta_a))]. \end{aligned} \tag{5.12}$$

An alternative loss that can be readily used in this setting is the α-loss introduced in Section 5.1.2. In the data-driven context, the α-loss can be written as

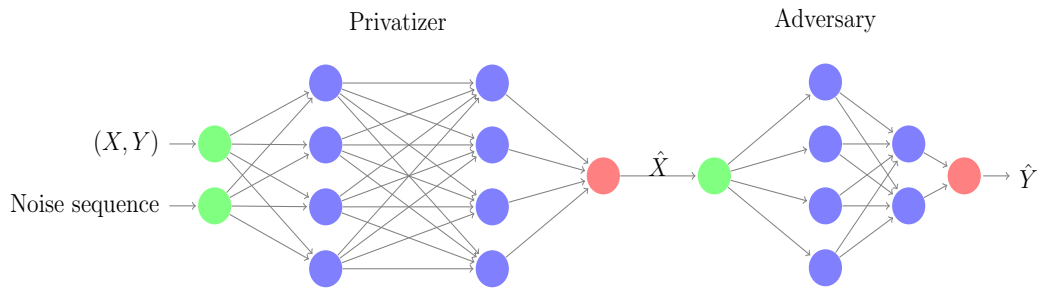

Figure 5.1: A multi-layer neural network model for the privatizer and adversary

$$\ell(h(g(x_{(i)}, y_{(i)}; \Theta_p); \Theta_a), y_{(i)}) = \frac{\alpha}{\alpha - 1} \Big(y_{(i)}(1 - h(g(x_{(i)}, y_{(i)}; \Theta_p); \Theta_a)^{1-\frac{1}{\alpha}}) \\ + (1 - y_{(i)})(1 - (1 - h(g(x_{(i)}, y_{(i)}; \Theta_p); \Theta_a))^{1-\frac{1}{\alpha}}) \Big), \quad (5.13)$$

for any constant $\alpha > 1$. As discussed in Section 5.1.2, the α-loss captures a variety of adversarial models and recovers both the log-loss (when $\alpha \to 1$) and 0-1 loss (when $\alpha \to \infty$). Futhermore, (5.13) suggests that α-leakage can be used as a surrogate (and smoother) loss function for the 0-1 loss (when α is relatively large).

The minimax optimization problem in (5.10) is a two-player non-cooperative game between the privatizer and the adversary. The strategies of the privatizer and adversary are given by Θ_p and Θ_a, respectively. Each player chooses the strategy that optimizes its objective function *w.r.t.* what its opponent does. In particular, the privatizer must expect that if it chooses Θ_p, the adversary will choose a Θ_a that minimizes the negative of its own loss function based on the choice of the privatizer. The optimal privacy mechanism is given by the equilibrium of the privatizer-adversary game.

In practice, we can learn the equilibrium of the game using an iterative algorithm presented in Algorithm 1. We first maximize the negative of the adversary's loss function in the inner loop to compute the parameters of h for a fixed g. Then, we minimize the privatizer's loss function, which is modeled as the negative of the adversary's loss function, to compute the parameters of g for a fixed h. To avoid

over-fitting and ensure convergence, we alternate between training the adversary for k epochs and training the privatizer for one epoch. This results in the adversary moving towards its optimal solution for small perturbations of the privatizer [101].

To incorporate the distortion constraint into the learning algorithm, we use the *penalty method* [140] and *augmented Lagrangian method* [141] to replace the constrained optimization problem by a series of unconstrained problems whose solutions asymptotically converge to the solution of the constrained problem. Under the penalty method, the unconstrained optimization problem is formed by adding a penalty to the objective function. The added penalty consists of a penalty parameter ρ_t multiplied by a measure of violation of the constraint. The measure of violation is non-zero when the constraint is violated and is zero if the constraint is not violated. Therefore, in Algorithm 1, the constrained optimization problem of the privatizer can be approximated by a series of unconstrained optimization problems with the loss function

$$\begin{aligned}\ell(\Theta_p^t, \Theta_a^{t+1}) = &- \frac{1}{M}\sum_{i=1}^{M} \ell(h(g(x_{(i)}, y_{(i)}; \Theta_p^t); \Theta_a^{t+1}), y_{(i)}) \\ &+ \rho_t \max\{0, \frac{1}{M}\sum_{i=1}^{M} d(g(x_{(i)}, y_{(i)}; \Theta_p^t), x_{(i)}) - D\},\end{aligned} \tag{5.14}$$

where ρ_t is a penalty coefficient which increases with the number of iterations t. For convex optimization problems, the solution to the series of unconstrained problems will eventually converge to the solution of the original constrained problem [140].

The augmented Lagrangian method is another approach to enforce equality constraints by penalizing the objective function whenever the constraints are not satisfied. Different from the penalty method, the augmented Lagrangian method combines the use of a Lagrange multiplier and a quadratic penalty term. Note that this method is designed for equality constraints. Therefore, we introduce a slack variable δ to convert the inequality distortion constraint into an equality constraint. Using the augmented Lagrangian method, the constrained optimization problem of the privatizer can be

Algorithm 1 Alternating minimax privacy preserving algorithm

Input: dataset $\mathcal{D}$, distortion parameter D, iteration number T

Output: Optimal privatizer parameter Θ_p^*

procedure ALERNATE MINIMAX($\mathcal{D}, D, T$)

Initialize Θ_p^1 and Θ_a^1

for $t = 1, ..., T$ **do**

Random minibatch of M datapoints $\{x_{(1)}, ..., x_{(M)}\}$ drawn from full dataset

Generate $\{\hat{x}_{(1)}, ..., \hat{x}_{(M)}\}$ via $\hat{x}_{(i)} = g(x_{(i)}, y_{(i)}; \Theta_p^t)$

Update the adversary parameter Θ_a^{t+1} by stochastic gradient ascend for k epochs

$$\Theta_a^{t+1} = \Theta_a^t + \alpha_t \nabla_{\Theta_a^t} \frac{1}{M} \sum_{i=1}^{M} -\ell(h(\hat{x}_{(i)}; \Theta_a^t), y_{(i)}), \quad \alpha_t > 0$$

Compute the descent direction $\nabla_{\Theta_p^t} l(\Theta_p^t, \Theta_a^{t+1})$, where

$$\ell(\Theta_p^t, \Theta_a^{t+1}) = -\frac{1}{M} \sum_{i=1}^{M} \ell(h(g(x_{(i)}, y_{(i)}; \Theta_p^t); \Theta_a^{t+1}), y_{(i)})$$

subject to $\frac{1}{M} \sum_{i=1}^{M} [d(g(x_{(i)}, y_{(i)}; \Theta_p^t), x_{(i)})] \leq D$

Perform line search along $\nabla_{\Theta_p^t} l(\Theta_p^t, \Theta_a^{t+1})$ and update

$$\Theta_p^{t+1} = \Theta_p^t - \alpha_t \nabla_{\Theta_p^t} \ell(\Theta_p^t, \Theta_a^{t+1})$$

Exit if solution converged

return Θ_p^{t+1}

replaced by a series of unconstrained problems with the loss function given by

$$\ell(\Theta_p^t, \Theta_a^{t+1}, \delta) = -\frac{1}{M}\sum_{i=1}^{M} \ell(h(g(x_{(i)}, y_{(i)}; \Theta_p^t); \Theta_a^{t+1}), y_{(i)}) \tag{5.15}$$
$$+ \frac{\rho_t}{2}(\frac{1}{M}\sum_{i=1}^{M} d(g(x_{(i)}, y_{(i)}; \Theta_p^t), x_{(i)}) + \delta - D)^2$$
$$- \lambda_t(\frac{1}{M}\sum_{i=1}^{M} d(g(x_{(i)}, y_{(i)}; \Theta_p^t), x_{(i)}) + \delta - D),$$

where ρ_t is a penalty coefficient which increases with the number of iterations t and λ_t is updated according to the rule $\lambda_{t+1} = \lambda_t - \rho_t(\frac{1}{M}\sum_{i=1}^{M} d(g(x_{(i)}, y_{(i)}; \Theta_p^t), x_{(i)}) + \delta - D)$. For convex optimization problems, the solution to the series of unconstrained problems formulated by the augmented Lagrangian method also converges to the solution of the original constrained problem [141].

5.1.4 Outline of Work

Our GAP framework is very general and can be used to capture many notions of privacy via various decision rules and loss funcitons. In the rest of this chapter, we investigate GAP under 0-1 loss for two simple yet canonical dataset models: (a) the binary data model (Section 5.2), and (b) the binary Gaussian mixture model (Section 5.3). Under the binary data model, both X and Y are binary. Under the binary Gaussian mixture model, Y is binary whereas X is conditionally Gaussian. We use these results to validate that the data-driven version of GAP can discover "theoretically optimal" privatization schemes.

In the data-driven approach of GAP, since $P(X, Y)$ is typically unknown in practice and our objective is to learn privatization schemes directly from data, we have to consider the empirical (data-driven) version of (5.5). Such an approach immediately hits a roadblock because taking derivatives of a 0-1 loss function *w.r.t.* the parameters of h and g is ill-defined. To circumvent this issue, similar to the common practice in the ML literature, we use the empirical log-loss (5.12) as the loss function for the

adversary. We derive game-theoretically optimal mechanisms for the 0-1 loss function, and use them as a benchmark against which we compare the performance of the data-driven GAP mechanisms. Finally, we demonstrate the performance of GAP on two meaningful, widely used dataset: GENKI and MNIST.

5.2 Binary Data Model

In this section, we study a setting where both the public and private variables are binary-valued random variables. Let $p_{i,j}$ denote the joint probability of $(X, Y) = (i, j)$, where $i, j \in \{0, 1\}$. To prevent an adversary from correctly inferring the private variable Y from the public variable X, the privatizer applies a randomized mechanism on X to generate the privatized data $\hat{X}$. Since both the original and privatized public variables are binary, the distortion between x and $\hat{x}$ can be quantified by the Hamming distortion; i.e. $d(\hat{x}, x) = 1$ if $\hat{x} \neq x$ and $d(\hat{x}, x) = 0$ if $\hat{x} = x$. Thus, the expected distortion is given by $\mathbb{E}[d(\hat{X}, X)] = P(\hat{X} \neq X)$.

5.2.1 Theoretical Approach for Binary Data Model

The adversary's objective is to correctly guess Y from $\hat{X}$. We consider a MAP adversary who has access to the joint distribution of (X, Y) and the privacy mechanism. The privatizer's goal is to privatize X in a way that minimizes the adversary's probability of correctly inferring Y from $\hat{X}$ subject to the distortion constraint. We first focus on private-data dependent (PDD) privacy mechanisms that depend on both Y and X. We later consider private-data independent (PDI) privacy mechanisms that only depend on X.

PDD Privacy Mechanism

Let $g(X, Y)$ denote a PDD mechanism. Since X, Y, and $\hat{X}$ are binary random variables, the mechanism $g(X, Y)$ can be represented by the conditional distribution $P(\hat{X}|X, Y)$ that maps the public and private variable pair (X, Y) to an output $\hat{X}$

given by

$$P(\hat{X}=0|X=0,Y=0)=s_{0,0},\quad P(\hat{X}=0|X=0,Y=1)=s_{0,1},$$

$$P(\hat{X}=1|X=1,Y=0)=s_{1,0},\quad P(\hat{X}=1|X=1,Y=1)=s_{1,1}.$$

Thus, the marginal distribution of $\hat{X}$ is given by

$$P(\hat{X}=0)=\sum_{X,Y}P(\hat{X}=0|X,Y)P(X,Y)=s_{0,0}p_{0,0}+s_{0,1}p_{0,1}+(1-s_{1,0})p_{1,0}+(1-s_{1,1})p_{1,1},$$

$$P(\hat{X}=1)=\sum_{X,Y}P(\hat{X}=1|X,Y)P(X,Y)=(1-s_{0,0})p_{0,0}+(1-s_{0,1})p_{0,1}+s_{1,0}p_{1,0}+s_{1,1}p_{1,1}.$$

If $\hat{X}=0$, the adversary's inference accuracy for guessing $\hat{Y}=1$ is

$$P(Y=1,\hat{X}=0)=\sum_{X}P(X,Y=1)P(\hat{X}=0|X,Y=1)=p_{1,1}(1-s_{1,1})+p_{0,1}s_{0,1}, \tag{5.16}$$

and the inference accuracy for guessing $\hat{Y}=0$ is

$$P(Y=0,\hat{X}=0)=\sum_{X}P(X,Y=0)P(\hat{X}=0|X,Y=0)=p_{1,0}(1-s_{1,0})+p_{0,0}s_{0,0}. \tag{5.17}$$

Let $\mathbf{s}=\{s_{0,0},s_{0,1},s_{1,0},s_{1,1}\}$. For $\hat{X}=0$, the MAP adversary's inference accuracy is

$$P_{\mathrm{d}}^{(\mathrm{B})}(\mathbf{s},\hat{X}=0)=\max\{P(Y=1,\hat{X}=0),P(Y=0,\hat{X}=0)\}. \tag{5.18}$$

Similarly, if $\hat{X}=1$, the MAP adversary's inference accuracy is given by

$$P_{\mathrm{d}}^{(\mathrm{B})}(\mathbf{s},\hat{X}=1)=\max\{P(Y=1,\hat{X}=1),P(Y=0,\hat{X}=1)\}, \tag{5.19}$$

where

$$P(Y=1,\hat{X}=1)=\sum_{X}P(X,Y=1)P(\hat{X}=1|X,Y=1)=p_{1,1}s_{1,1}+p_{0,1}(1-s_{0,1}),$$

$$P(Y=0,\hat{X}=1)=\sum_{X}P(X,Y=0)P(\hat{X}=1|X,Y=0)=p_{1,0}s_{1,0}+p_{0,0}(1-s_{0,0}).$$

As a result, for a fixed privacy mechanism $\mathbf{s}$, the MAP adversary's inference accuracy can be written as

$$P_{\mathrm{d}}^{(\mathrm{B})}=\max_{h(\cdot)}P(h(g(X,Y))=Y)=P_{\mathrm{d}}^{(\mathrm{B})}(\mathbf{s},\hat{X}=0)+P_{\mathrm{d}}^{(\mathrm{B})}(\mathbf{s},\hat{X}=1). \tag{5.20}$$

Thus, the optimal PDD privacy mechanism is determined by solving

$$\begin{aligned}\min_{\mathbf{s}}\quad & P_{\mathrm{d}}^{(\mathrm{B})}(\mathbf{s},\hat{X}=0)+P_{\mathrm{d}}^{(\mathrm{B})}(\mathbf{s},\hat{X}=1) \\ s.t.\quad & P(\hat{X}=0,X=1)+P(\hat{X}=1,X=0)\leq D \\ & \mathbf{s}\in[0,1]^4.\end{aligned} \tag{5.21}$$

Notice that the above constrained optimization problem is a four dimensional optimization problem parameterized by $\mathbf{p}=\{p_{0,0},p_{0,1},p_{1,0},p_{1,1}\}$ and D. Interestingly, we can formulate (5.21) as a linear program (LP) given by

$$\begin{aligned}\min_{s_{1,1},s_{0,1},s_{1,0},s_{0,0},t_0,t_1}\quad & t_0+t_1 \\ s.t.\quad & 0\leq s_{1,1},s_{0,1},s_{1,0},s_{0,0}\leq 1 \\ & p_{1,1}(1-s_{1,1})+p_{0,1}s_{0,1}\leq t_0 \\ & p_{1,0}(1-s_{1,0})+p_{0,0}s_{0,0}\leq t_0 \\ & p_{1,1}s_{1,1}+p_{0,1}(1-s_{0,1})\leq t_1 \\ & p_{1,0}s_{1,0}+p_{0,0}(1-s_{0,0})\leq t_1 \\ & p_{1,1}(1-s_{1,1})+p_{0,1}(1-s_{0,1})+p_{1,0}(1-s_{1,0})+p_{0,0}(1-s_{0,0})\leq D,\end{aligned} \tag{5.22}$$

where t_0 and t_1 are two slack variables representing the maxima in (5.18) and (5.19), respectively. The optimal mechanism can be obtained by numerically solving (5.22) using any off-the-shelf LP solver.

PDI Privacy Mechanism

In the previous section, we considered PDD privacy mechanisms. Although we were able to formulate the problem as a linear program with four variables, determining a closed form solution for such a highly parameterized problem is not analytically tractable. Thus, we now consider the simple (yet meaningful) class of PDI privacy mechanisms. Under PDI privacy mechanisms, the Markov chain $Y \to X \to \hat{X}$ holds. As a result, $P(Y, \hat{X} = \hat{x})$ can be written as

$$\begin{aligned} P(Y, \hat{X} = \hat{x}) &= \sum_X P(Y, \hat{X} = \hat{x}|X)P(X) = \sum_X P(Y|X)P(\hat{X} = \hat{x}|X)P(X) \\ &= \sum_X P(Y, X)P(\hat{X} = \hat{x}|X), \end{aligned} \tag{5.23}$$

where the second equality is due to the conditional independence property of the Markov chain $Y \to X \to \hat{X}$.

For the PDI mechanisms, the privacy mechanism $g(X, Y)$ can be represented by the conditional distribution $P(\hat{X}|X)$. To make the problem more tractable, we focus on a slightly simpler setting in which $Y = X \oplus N$, where $N \in \{0, 1\}$ is a random variable independent of X and follows a Bernoulli distribution with parameter q. In this setting, the joint distribution of (X, Y) can be computed as

$$P(X = 1, Y = 1) = P(Y = 1|X = 1)P(X = 1) = p(1 - q), \tag{5.24}$$

$$P(X = 0, Y = 1) = P(Y = 1|X = 0)P(X = 0) = (1 - p)q, \tag{5.25}$$

$$P(X = 1, Y = 0) = P(Y = 0|X = 1)P(X = 1) = pq, \tag{5.26}$$

$$P(X = 0, Y = 0) = P(Y = 0|X = 0)P(X = 0) = (1 - p)(1 - q). \tag{5.27}$$

Let $\mathbf{s} = \{s_0, s_1\}$ in which $s_0 = P(\hat{X} = 0|X = 0)$ and $s_1 = P(\hat{X} = 1|X = 1)$. The joint distribution of $(Y, \hat{X})$ is given by

$$P(Y = 1, \hat{X} = 0) = p(1 - q)(1 - s_1) + (1 - p)qs_0,$$

$$P(Y=0,\hat{X}=0)=pq(1-s_1)+(1-p)(1-q)s_0,$$

$$P(Y=1,\hat{X}=1)=p(1-q)s_1+(1-p)q(1-s_0),$$

$$P(Y=0,\hat{X}=1)=pqs_1+(1-p)(1-q)(1-s_0).$$

Using the above joint probabilities, for a fixed $\mathbf{s}$, we can write the MAP adversary's inference accuracy as

$$P_{\mathrm{d}}^{(\mathrm{B})}=\max_{h(\cdot)}P(h(g(X,Y))=Y)=\max\{P(Y=1,\hat{X}=0),P(Y=0,\hat{X}=0)\} \quad (5.28)$$
$$+\max\{P(Y=1,\hat{X}=1),P(Y=0,\hat{X}=1)\}.$$

Therefore, the optimal PDI privacy mechanism is given by the solution to

$$\begin{aligned} \min_{\mathbf{s}} \quad & P_{\mathrm{d}}^{(\mathrm{B})} \\ s.t. \quad & P(\hat{X}=0,X=1)+P(\hat{X}=1,X=0)\leq D \\ & \mathbf{s}\in[0,1]^2, \end{aligned} \quad (5.29)$$

where the distortion in (5.29) is given by $(1-s_0)(1-p)+(1-s_1)p$. By (5.28), $P_{\mathrm{d}}^{(\mathrm{B})}$ can be considered as a sum of two functions, where each function is a maximum of two linear functions. Thus, it is convex in s_0 and s_1 for different values of p,q and D.

Theorem 6. *For fixed p,q and D, there exists infinitely many PDI privacy mechanisms that achieve the optimal privacy-utility tradeoff. If $q=\frac{1}{2}$, any privacy mechanism that satisfies $\{s_0,s_1|ps_1+(1-p)s_0\geq 1-D, s_0,s_1\in[0,1]\}$ is optimal. If $q\neq\frac{1}{2}$, the optimal PDI privacy mechanism is given as follows:*

- *If $1-D>\max\{p,1-p\}$, the optimal privacy mechanism is given by $\{s_0,s_1|ps_1+(1-p)s_0=1-D, s_0,s_1\in[0,1]\}$. The adversary's accuracy of correctly guessing the private variable is*

$$\begin{cases} (1-2q)(1-D)+q & \text{if } q<\frac{1}{2} \\ (2q-1)(1-D)+1-q & \text{if } q>\frac{1}{2} \end{cases}. \quad (5.30)$$

- *Otherwise, the optimal privacy mechanism is given by* $\{s_0, s_1 | \max\{\min\{p, 1-p\}, 1-D\} \leq ps_1 + (1-p)s_0 \leq \max\{p, 1-p\}, s_0, s_1 \in [0,1]\}$ *and the adversary's accuracy of correctly guessing the private variable is*

$$\begin{cases} p(1-q) + (1-p)q & \textit{if } p \geq \frac{1}{2}, q < \frac{1}{2} \textit{ or } p \leq \frac{1}{2}, q > \frac{1}{2} \\ pq + (1-p)(1-q) & \textit{if } p \geq \frac{1}{2}, q > \frac{1}{2} \textit{ or } p \leq \frac{1}{2}, q < \frac{1}{2} \end{cases}. \tag{5.31}$$

Proof sketch: The proof of Theorem 6 is provided in Appendix G. We briefly sketch the proof details here. For the special case $q = \frac{1}{2}$, the solution is trivial since the private variable Y is independent of the public variable X. Thus, the optimal solution is given by any s_0, s_1 that satisfies the distortion constraint $\{s_0, s_1 | ps_1 + (1-p)s_0 \geq 1-D, s_0, s_1 \in [0,1]\}$. For $q \neq \frac{1}{2}$, we separate the optimization problem in (5.29) into four subproblems based on the decision of the adversary and compute the optimal privacy mechanism of the privatizer in each subproblem. Summarizing the optimal solutions to the subproblems for different values of p, q and D yields Theorem 6.

Remark: Note that if $1 - D > \max\{p, 1-p\}$, i.e., $D < \min\{p, 1-p\}$, the privacy guarantee achieved by the optimal PDI mechanism (the MAP adversary's accuracy of correctly guessing the private variable) decreases linearly with D. For $D \geq \min\{p, 1-p\}$, the optimal PDI mechanism achieves a constant privacy guarantee regardless of D. However, in this case, the privatizer can just use the optimal privacy mechanism with $D = \min\{p, 1-p\}$ to optimize privacy guarantee without further sacrificing utility.

5.2.2 Data-driven Approach for Binary Data Model

In practice, the joint distribution of (X, Y) is often unknown to the data holder. Instead, the data holder has access to a dataset $\mathcal{D}$, which is used to learn a good privatization mechanism in a generative adversarial fashion. In the training phase, the data holder learns the parameters of the conditional generative model (representing the privatization scheme) by competing against a computational adversary

represented by a neural network. The details of both neural networks are provided later in this section. When convergence is reached, we evaluate the performance of the learned privatization scheme by computing the accuracy of inferring Y under a strong MAP adversary that: (a) has access to the joint distribution of (X, Y), (b) has knowledge of the learned privacy mechanism, and (c) can compute the MAP rule. The MAP adversary corresponds to the 0-1 loss function that is effectively looking at the inference error rate of the adversary. Ultimately, the data holder's hope is to learn a privatization scheme that matches the one obtained under the game-theoretic framework, where both the adversary and privatizer are assumed to have access to $P(X, Y)$. To evaluate our data-driven approach, we compare the mechanisms learned in an adversarial fashion on $\mathcal{D}$ with the game-theoretically optimal ones.

Since the private variable Y is binary, we use the empirical log-loss function for the adversary (see (5.12)). For a fixed Θ_p, the adversary learns the optimal Θ_a^* by maximizing $-L_{\text{XE}}(h(g(X, Y; \Theta_p); \Theta_a), Y)$ given in (5.12). For a fixed Θ_a, the privatizer learns the optimal Θ_p^* by minimizing $-L_{\text{XE}}(h(g(X, Y; \Theta_p); \Theta_a), Y)$ subject to the distortion constraint (see (5.10)). Since both X and Y are binary variables, we can use the privatizer parameter Θ_p to represent the privacy mechanism $\mathbf{s}$ directly. For the adversary, we define $\Theta_a = (\Theta_{a,0}, \Theta_{a,1})$, where $\Theta_{a,0} = P(Y = 0|\hat{X} = 0)$ and $\Theta_{a,1} = P(Y = 1|\hat{X} = 1)$. Thus, given a privatized public variable input $g(x_{(i)}, y_{(i)}; \Theta_p) \in \{0, 1\}$, the output belief of the adversary guessing $y_{(i)} = 1$ can be written as $(1 - \Theta_{a,0})(1 - g(x_{(i)}, y_{(i)}; \Theta_p)) + \Theta_{a,1} g(x_{(i)}, y_{(i)}; \Theta_p)$.

For PDD privacy mechanisms, we have $\Theta_p = \mathbf{s} = \{s_{0,0}, s_{0,1}, s_{1,0}, s_{1,1}\}$. Given the fact that both $x_{(i)}$ and $y_{(i)}$ are binary, we use two simple neural networks to model the privatizer and the adversary. As shown in Figure 5.2, the privatizer is modeled as a two-layer neural network parameterized by $\mathbf{s}$, while the adversary is modeled as a two-layer neural network classifier. From the perspective of the privatizer, the belief

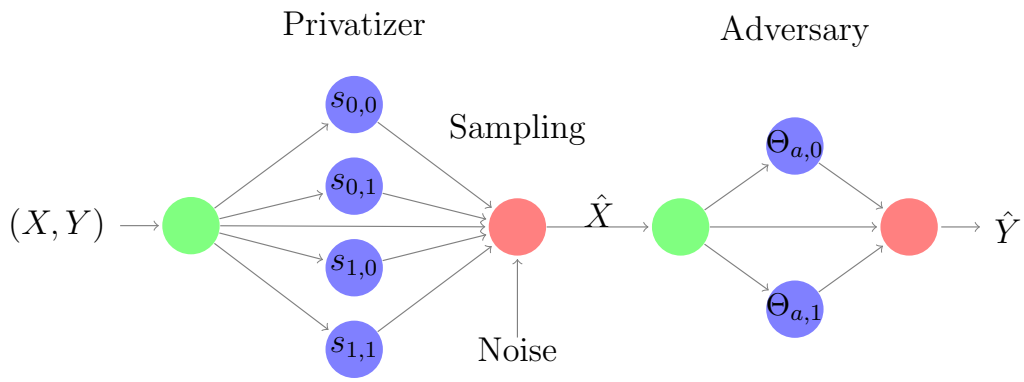

Figure 5.2: Neural network structure of the privatizer and adversary for binary data model

of an adversary guessing $y_{(i)} = 1$ conditioned on the input $(x_{(i)}, y_{(i)})$ is given by

$$h(g(x_{(i)}, y_{(i)}; \mathbf{s}); \Theta_a) = \Theta_{a,1} P(\hat{x}_{(i)} = 1) + (1 - \Theta_{a,0}) P(\hat{x}_{(i)} = 0), \tag{5.32}$$

where

$$\begin{aligned} P(\hat{x}_{(i)} = 1) =& x_{(i)} y_{(i)} s_{1,1} + (1 - x_{(i)}) y_{(i)} (1 - s_{0,1}) \\ &+ x_{(i)} (1 - y_{(i)}) s_{1,0} + (1 - x_{(i)})(1 - y_{(i)})(1 - s_{0,0}), \\ P(\hat{x}_{(i)} = 0) =& x_{(i)} y_{(i)} (1 - s_{1,1}) + (1 - x_{(i)}) y_{(i)} s_{0,1} \\ &+ x_{(i)} (1 - y_{(i)})(1 - s_{1,0}) + (1 - x_{(i)})(1 - y_{(i)}) s_{0,0}. \end{aligned}$$

Furthermore, the expected distortion is given by

$$\begin{aligned} \mathbb{E}_{\mathcal{D}}[d(g(X, Y; \mathbf{s}), X)] =& \frac{1}{n} \sum_{i=1}^{n} [x_{(i)} y_{(i)} (1 - s_{1,1}) + x_{(i)} (1 - y_{(i)})(1 - s_{1,0}) \\ &+ (1 - x_{(i)}) y_{(i)} (1 - s_{0,1}) + (1 - x_{(i)})(1 - y_{(i)})(1 - s_{0,0})]. \end{aligned} \tag{5.33}$$

Similar to the PDD case, we can also compute the belief of guessing $y_{(i)} = 1$ conditional on the input $(x_{(i)}, y_{(i)})$ for the PDI schemes. Observe that in the PDI case, $\Theta_p = \mathbf{s} = \{s_0, s_1\}$. Therefore, we have

$$h(g(x_{(i)}, y_{(i)}; \mathbf{s}); \Theta_a) = \Theta_{a,1} [x_{(i)} s_1 + (1 - x_{(i)})(1 - s_0)] \tag{5.34}$$

$$+ (1 - \Theta_{a,0})[(1 - x_{(i)})s_0 + x_{(i)}(1 - s_1)].$$

Under PDI schemes, the expected distortion is given by

$$\mathbb{E}_{\mathcal{D}}[d(g(X, Y; \mathbf{s}), X)] = \frac{1}{n} \sum_{i=1}^{n} [x_{(i)}(1 - s_1) + (1 - x_{(i)})(1 - s_0)]. \tag{5.35}$$

Thus, we can use Algorithm 1 proposed in Section 5.1.3 to learn the optimal PDD and PDI privacy mechanisms from the dataset.

5.2.3 Illustration of Results

We now evaluate our proposed GAP framework using synthetic datasets. We focus on the setting in which $Y = X \oplus N$, where $N \in \{0, 1\}$ is a random variable independent of X and follows a Bernoulli distribution with parameter q. We generate two synthetic datasets with (p, q) equal to $(0.75, 0.25)$ and $(0.5, 0.25)$, respectively. Each synthetic dataset used in this experiment contains $10,000$ training samples and $2,000$ test samples. We use Tensorflow [142] to train both the privatizer and the adversary using Adam optimizer with a learning rate of 0.01 and a minibatch size of 200. The distortion constraint is enforced by the penalty method provided in (5.14).

Figure 5.3a illustrates the performance of both optimal PDD and PDI privacy mechanisms against a strong theoretical MAP adversary when $(p, q) = (0.5, 0.25)$. It can be seen that the inference accuracy of the MAP adversary reduces as the distortion increases for both optimal PDD and PDI privacy mechanisms. As one would expect, the PDD privacy mechanism achieves a lower inference accuracy for the adversary, i.e., better privacy, than the PDI mechanism. Furthermore, when the distortion is higher than some threshold, the inference accuracy of the MAP adversary saturates regardless of the distortion. This is due to the fact that the correlation between the private variable and the privatized public variable cannot be further reduced once the distortion is larger than the saturation threshold. Therefore, increasing distortion will not further reduce the accuracy of the MAP adversary. We also observe that

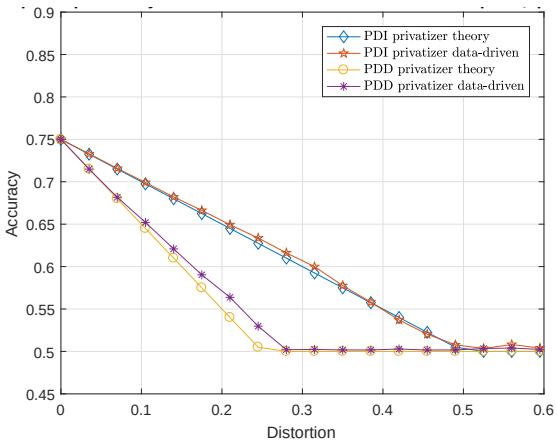

(a) Performance of privacy mechanisms against MAP adversary for $p = 0.5$

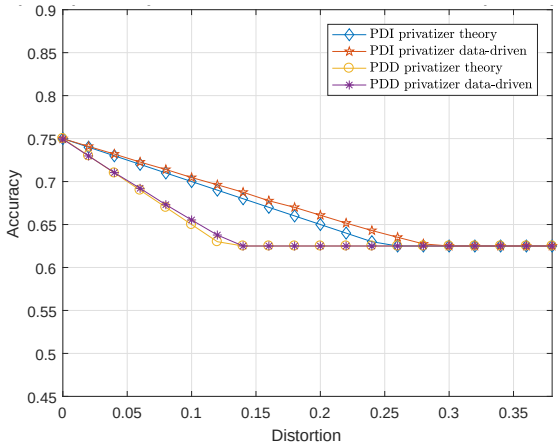

(b) Performance of privacy mechanisms against MAP adversary for $p = 0.75$

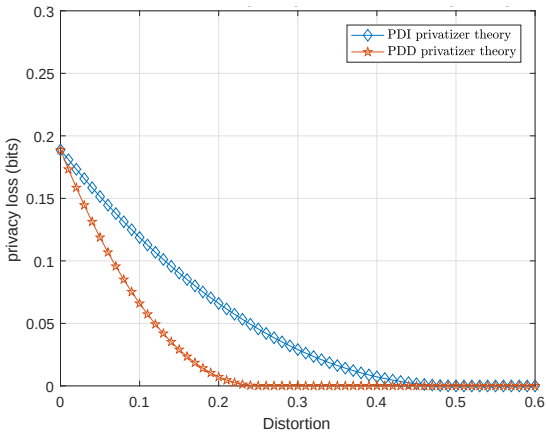

(c) Performance of privacy mechanisms under MI privacy metric for $p = 0.5$

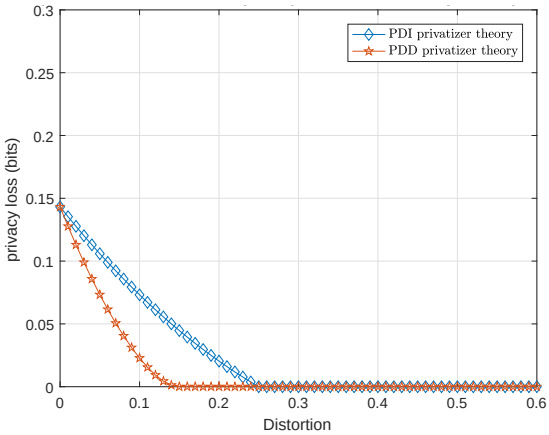

(d) Performance of privacy mechanisms under MI privacy metric for $p = 0.75$

Figure 5.3: Privacy-distortion tradeoff for binary data model

the privacy mechanism obtained via the data-driven approach performs very well when pitted against the MAP adversary (maximum accuracy difference around 3% compared to the theoretical approach). In other words, for the binary data model, the data-driven version of GAP can yield privacy mechanisms that perform as well as the mechanisms computed under the theoretical version of GAP, which assumes that the privatizer has access to the underlying distribution of the dataset.

Figure 5.3b shows the performance of both optimal PDD and PDI privacy mechanisms against the MAP adversary for $(p, q) = (0.75, 0.25)$. Similar to the equal prior case, we observe that both PDD and PDI privacy mechanisms reduce the accuracy of the MAP adversary as the distortion increases and saturate when the distortion goes above a certain threshold. It can be seen that the saturation thresholds for both PDD and PDI privacy mechanisms in Figure 5.3b are lower than the "equal prior" case plotted in Figure 5.3a. The reason is that when $(p, q) = (0.75, 0.25)$, the correlation between Y and X is weaker than the "equal prior" case. Therefore, it requires less distortion to achieve the same privacy. We also observe that the performance of the GAP mechanism obtained via the data-driven approach is comparable to the mechanism computed via the theoretical approach.

The performance of the GAP mechanism obtained using the log-loss function (i.e., MI privacy) is plotted in Figure 5.3c and 5.3d. Similar to the MAP adversary case, as the distortion increases, the mutual information between the private variable and the privatized public variable achieved by the optimal PDD and PDI mechanisms decreases as long as the distortion is below some threshold. When the distortion goes above the threshold, the optimal privacy mechanism is able to make the private variable and the privatized public variable independent regardless of the distortion. Furthermore, the values of the saturation thresholds are very close to what we observe in Figure 5.3a and 5.3b.

5.3 Binary Gaussian Mixture Model

Thus far, we have studied a simple binary dataset model. In many real datasets, the sample space of variables often takes more than just two possible values. It is well known that the Gaussian distribution is a flexible approximate for many distributions [143]. Therefore, in this section, we study a setting where $Y \in \{0, 1\}$ and X is a Gaussian random variable whose mean and variance are dependent on Y. Without

loss of generality, let $\mathbb{E}[X|Y=1] = -\mathbb{E}[X|Y=0] = \mu$ and $P(Y=1) = \tilde{p}$. Thus, $X|Y=0 \sim \mathcal{N}(-\mu, \Sigma_0)$ and $X|Y=1 \sim \mathcal{N}(\mu, \Sigma_1)$.

5.3.1 GAP for Single-dimensional Gaussian Mixture Model

In this section, we consider the setting where the public variable is a single-dimensional Gaussian random variable conditional on the private variable, i.e., $X|Y = 0 \sim \mathcal{N}(-\mu, \sigma_0)$ and $X|Y = 1 \sim \mathcal{N}(\mu, \sigma_1)$. Similar to the binary data model, we study two privatization schemes: (a) private-data independent (PDI) schemes (where $\hat{X} = g(X)$), and (b) private-data dependent (PDD) schemes (where $\hat{X} = g(X, Y)$). In order to have a tractable model for the privatizer, we assume $g(X, Y)$ is realized by adding an affine function of an independently generated random noise to the public variable X. The affine function enables controlling both the mean and variance of the privatized data. In particular, we consider $g(X, Y) = X + (1-Y)\beta_0 - Y\beta_1 + (1-Y)\gamma_0 N + Y\gamma_1 N$, in which N is a one dimensional random variable and $\beta_0, \beta_1, \gamma_0, \gamma_1$ are constant parameters. The goal of the privatizer is to sanitze the public data X subject to the distortion constraint $\mathbb{E}_{\hat{X},X}||\hat{X} - X||_2^2 \leq D$.

To make the problem more tractable, let us consider a slightly simpler setting in which $\sigma_0 = \sigma_1 = \sigma$. We will relax this assumption later when we take a data-driven approach. We further assume that N is a standard Gaussian random variable. One might, rightfully, question our choice of focusing on adding (potentially Y-dependent) Gaussian noise. Though other distributions can be considered, our approach is motivated by the following two reasons:

- (a) Even though it is known that adding Gaussian noise is not the worst case noise adding mechanism for non-Gaussian X [103], identifying the optimal noise distribution is mathematically intractable. Thus, for tractability and ease of analysis, we choose Gaussian noise.

- (b) Adding Gaussian noise to each data entry preserves the conditional Gaussianity of the released dataset.

In what follows, we will analyze a variety of PDI and PDD mechanisms.

PDI Gaussian Noise Adding Privacy Mechanism

We consider a PDI noise adding privatization scheme which adds an affine function of the standard Gaussian noise to the public variable. Since the privacy mechanism is PDI, we have $g(X,Y) = X + \beta + \gamma N$, where β and γ are constant parameters and $N \sim \mathcal{N}(0,1)$. Using the classical Gaussian hypothesis testing analysis [144], it is straightforward to verify that the optimal inference accuracy (i.e., probability of detection) of the MAP adversary is given by

$$P_{\mathrm{d}}^{(\mathrm{G})} = \tilde{p}Q\left(-\frac{\alpha}{2} + \frac{1}{\alpha}\ln\left(\frac{1-\tilde{p}}{\tilde{p}}\right)\right) + (1-\tilde{p})Q\left(-\frac{\alpha}{2} - \frac{1}{\alpha}\ln\left(\frac{1-\tilde{p}}{\tilde{p}}\right)\right), \tag{5.36}$$

where $\alpha = \frac{2\mu}{\sqrt{\gamma^2+\sigma^2}}$ and $Q(x) = \frac{1}{\sqrt{2\pi}}\int_x^\infty \exp(-\frac{u^2}{2})du$. Moreover, since $\mathbb{E}_{\hat{X},X}[d(\hat{X},X)] = \beta^2 + \gamma^2$, the distortion constraint is equivalent to $\beta^2 + \gamma^2 \leq D$.

Theorem 7. *For a PDI Gaussian noise adding privatization scheme given by $g(X,Y) = X + \beta + \gamma N$, with $\beta \in \mathbb{R}$ and $\gamma \geq 0$, the optimal parameters are given by*

$$\beta^* = 0, \gamma^* = \sqrt{D}. \tag{5.37}$$

Let $\alpha^* = \frac{2\mu}{\sqrt{D+\sigma^2}}$. For this optimal scheme, the accuracy of the MAP adversary is

$$P_d^{(G)*} = \tilde{p}Q\left(-\frac{\alpha^*}{2} + \frac{1}{\alpha^*}\ln\left(\frac{1-\tilde{p}}{\tilde{p}}\right)\right) + (1-\tilde{p})Q\left(-\frac{\alpha^*}{2} - \frac{1}{\alpha^*}\ln\left(\frac{1-\tilde{p}}{\tilde{p}}\right)\right). \tag{5.38}$$

The proof of Theorem 7 is provided in Appendix H. We observe that the PDI Gaussian noise adding privatization scheme which minimizes the inference accuracy of the MAP adversary with distortion upper-bounded by D is to add a zero-mean Gaussian noise with variance D.

PDD Gaussian Noise Adding Privacy Mechanism

For PDD privatization schemes, we first consider a simple case in which $\gamma_0 = \gamma_1 = 0$. Without loss of generality, we assume that both β_0 and β_1 are non-negative. The privatized data is given by $\hat{X} = X + (1-Y)\beta_0 - Y\beta_1$. This is a PDD mechanism since $\hat{X}$ depends on both X and Y. Intuitively, this mechanism privatizes the data by shifting the two Gaussian distributions (under $Y = 0$ and $Y = 1$) closer to each other. Under this mechanism, it is easy to show that the adversary's MAP probability of inferring the private variable Y from $\hat{X}$ is given by $P_{\mathrm{d}}^{(\mathrm{G})}$ in (5.36) with $\alpha = \frac{2\mu-(\beta_1+\beta_0)}{\sigma}$. Observe that since $d(\hat{X}, X) = ((1-Y)\beta_0 - Y\beta_1)^2$, we have $\mathbb{E}_{\hat{X},X}[d(\hat{X}, X)] = (1-\tilde{p})\beta_0^2 + \tilde{p}\beta_1^2$. Thus, the distortion constraint implies $(1-\tilde{p})\beta_0^2 + \tilde{p}\beta_1^2 \leq D$.

Theorem 8. *For a PDD privatization scheme given by $g(X,Y) = X+(1-Y)\beta_0-Y\beta_1$, $\beta_0, \beta_1 \geq 0$, the optimal parameters are given by*

$$\beta_0^* = \sqrt{\frac{\tilde{p}D}{1-\tilde{p}}}, \quad \beta_1^* = \sqrt{\frac{(1-\tilde{p})D}{\tilde{p}}}. \tag{5.39}$$

For this optimal PDD privatization scheme, the accuracy of the MAP adversary is given by (5.36) *with* $\alpha = \frac{2\mu-(\sqrt{\frac{(1-\tilde{p})D}{\tilde{p}}}+\sqrt{\frac{\tilde{p}D}{1-\tilde{p}}})}{\sigma}$.

The proof of Theorem 8 is provided in Appendix I. When $P(Y = 1) = P(Y = 0) = \frac{1}{2}$, we have $\beta_0 = \beta_1 = \sqrt{D}$, which implies that the optimal privacy mechanism for this particular case is to shift the two Gaussian distributions closer to each other equally by $\sqrt{D}$ regardless of the variance σ^2. When $P(Y = 1) = \tilde{p} > \frac{1}{2}$, the Gaussian distribution with a lower prior probability, in this case, $X|Y = 0$, gets shifted $\frac{\tilde{p}}{1-\tilde{p}}$ times more than $X|Y = 1$.

Next, we consider a slightly more complicated case in which $\gamma_0 = \gamma_1 = \gamma \geq 0$. Thus, the privacy mechanism is given by $g(X,Y) = X + (1-Y)\beta_0 - Y\beta_1 + \gamma N$, where $N \sim \mathcal{N}(0,1)$. Intuitively, this mechanism privatizes the data by shifting the

two Gaussian distributions (under $Y = 0$ and $Y = 1$) closer to each other and adding another Gaussian noise $N \in \mathcal{N}(0,1)$ scaled by a constant γ. In this case, the MAP probability of inferring the private variable Y from $\hat{X}$ is given by (5.36) with $\alpha = \frac{2\mu-(\beta_1+\beta_0)}{\sqrt{\gamma^2+\sigma^2}}$. Furthermore, the distortion constraint is equivalent to $(1-\tilde{p})\beta_0^2 + \tilde{p}\beta_1^2 + \gamma^2 \leq D$.

Theorem 9. *For a PDD privatization scheme $g(X,Y) = X + (1-Y)\beta_0 - Y\beta_1 + \gamma N$ with $\beta_0, \beta_1, \gamma \geq 0$, the optimal parameters $\beta_0^*, \beta_1^*, \gamma^*$ are given by the solution to*

$$\begin{aligned} \min_{\beta_0,\beta_1,\gamma} \quad & \frac{2\mu - \beta_0 - \beta_1}{\sqrt{\gamma^2+\sigma^2}} \\ s.t. \quad & (1-\tilde{p})\beta_0^2 + \tilde{p}\beta_1^2 + \gamma^2 \leq D \\ & \beta_0, \beta_1, \gamma \geq 0. \end{aligned} \tag{5.40}$$

Using this optimal scheme, the accuracy of the MAP adversary is given by (5.36) *with* $\alpha = \frac{2\mu-\beta_0^*-\beta_1^*}{\sqrt{(\gamma^*)^2+\sigma^2}}$.

Proof. Similar to the proofs of Theorem 7 and 8, we can compute the derivative of $P_{\mathrm{d}}^{(\mathrm{G})}$ *w.r.t.* α. It is easy to verify that $P_{\mathrm{d}}^{(\mathrm{G})}$ is monotonically increasing with α. Therefore, the optimal mechanism is given by the solution to (5.40). Substituting the optimal parameters into (5.36) yields the MAP probability of inferring the private variable Y from $\hat{X}$. □

Remark: Note that the objective function in (5.40) only depends on $\beta_0 + \beta_1$ and γ. We define $\beta = \beta_0 + \beta_1$. Thus, the above objective function can be written as

$$\min_{\beta,\gamma} \frac{2\mu - \beta}{\sqrt{\gamma^2+\sigma^2}}. \tag{5.41}$$

It is straightforward to verify that the determinant of the Hessian of (5.41) is always non-positive. Therefore, the above optimization problem is non-convex in β and γ.

Finally, we consider the PDD Gaussian noise adding privatization scheme given by $g(X,Y) = X + (1-Y)\beta_0 - Y\beta_1 + (1-Y)\gamma_0 N + Y\gamma_1 N$, where $N \sim \mathcal{N}(0,1)$. This PDD mechanism is the most general one in the Gaussian noise adding setting and includes the two previous mechanisms. The objective of the privatizer is to minimize the adversary's probability of correctly inferring Y from $g(X,Y)$ subject to the distortion constraint given by $\tilde{p}((\beta_1)^2 + (\gamma_1)^2) + (1-\tilde{p})((\beta_0)^2 + (\gamma_0)^2) \leq D$. As we have discussed in the remark after Theorem 9, the problem becomes non-convex even for the simpler case in which $\gamma_0 = \gamma_1 = \gamma$. In order to obtain the optimal parameters for this case, we first show that the optimal privacy mechanism lies on the boundary of the distortion constraint.

Proposition 1. *For the privacy mechanism given by* $g(X,Y) = X + (1-Y)\beta_0 - Y\beta_1 + (1-Y)\gamma_0 N + Y\gamma_1 N$*, the optimal parameters* $\beta_0^*, \beta_1^*, \gamma_0^*, \gamma_1^*$ *satisfy* $\tilde{p}((\beta_1^*)^2 + (\gamma_1^*)^2) + (1-\tilde{p})((\beta_0^*)^2 + (\gamma_0^*)^2) = D$.

Proof. We prove the above statement by contradiction. Assume that the optimal parameters satisfy $\tilde{p}((\beta_1^*)^2+(\gamma_1^*)^2)+(1-\tilde{p})((\beta_0^*)^2+(\gamma_0^*)^2) < D$. Let $\tilde{\beta}_1 = \beta_1^* + c$, where $c > 0$ is chosen so that $\tilde{p}((\tilde{\beta}_1)^2+(\gamma_1^*)^2)+(1-\tilde{p})((\beta_0^*)^2+(\gamma_0^*)^2) = D$. Since the inference accuracy is monotonically decreasing with β_1, the resultant inference accuracy can only be lower for replacing β_1^* with $\tilde{\beta}_1$. This contradicts with the assumption that $\tilde{p}((\beta_1^*)^2 + (\gamma_1^*)^2) + (1-\tilde{p})((\beta_0^*)^2 + (\gamma_0^*)^2) < D$. Using the same type of analysis, we can show that any parameter that deviates from $\tilde{p}((\beta_1^*)^2+(\gamma_1^*)^2)+(1-\tilde{p})((\beta_0^*)^2+(\gamma_0^*)^2) = D$ is suboptimal. $\square$

Let $e_0^2 = (\beta_0^*)^2 + (\gamma_0^*)^2$ and $e_1^2 = (\beta_1^*)^2 + (\gamma_1^*)^2$. Since the optimal parameters of the privatizer lie on the boundary of the distortion constraint, we have $\tilde{p}e_1^2 + (1-\tilde{p})e_0^2 = D$. This implies (e_0, e_1) lies on the boundary of an ellipse parametrized by $\tilde{p}$ and D. Thus, we have $e_1 = \sqrt{\frac{D}{\tilde{p}}}\frac{1-\epsilon^2}{1+\epsilon^2}$ and $e_0 = 2\sqrt{\frac{D}{1-\tilde{p}}}\frac{\epsilon}{1+\epsilon^2}$, where $\epsilon \in [0,1]$. Therefore, the optimal

parameters satisfy

$$(\beta_0^*)^2 + (\gamma_0^*)^2 = \left[2\sqrt{\frac{D}{1-\tilde{p}}}\frac{\epsilon}{1+\epsilon^2}\right]^2, \quad (\beta_1^*)^2 + (\gamma_1^*)^2 = \left[\sqrt{\frac{D}{\tilde{p}}}\frac{1-\epsilon^2}{1+\epsilon^2}\right]^2. \tag{5.42}$$

This implies $(\beta_i^*, \gamma_i^*), i \in \{0,1\}$ lie on the boundary of two circles parametrized by $D, \tilde{p}$ and ϵ. Thus, we can write $\beta_0^*, \beta_1^*, \gamma_0^*, \gamma_1^*$ as

$$\begin{aligned} \beta_0^* &= 2\sqrt{\frac{D}{1-\tilde{p}}}\frac{\epsilon}{1+\epsilon^2}\frac{1-w_0^2}{1+w_0^2}, \quad \beta_1^* = \sqrt{\frac{D}{\tilde{p}}}\frac{1-\epsilon^2}{1+\epsilon^2}\frac{1-w_1^2}{1+w_1^2}, \\ \gamma_0^* &= 4\sqrt{\frac{D}{1-\tilde{p}}}\frac{\epsilon}{1+\epsilon^2}\frac{w_0}{1+w_0^2}, \quad \gamma_1^* = 2\sqrt{\frac{D}{\tilde{p}}}\frac{1-\epsilon^2}{1+\epsilon^2}\frac{w_1}{1+w_1^2}, \end{aligned} \tag{5.43}$$

where $\epsilon, w_0, w_1 \in [0,1]$. The optimal parameters $\beta_0^*, \beta_1^*, \gamma_0^*, \gamma_1^*$ can be computed by a grid search in the cube parametrized by $\epsilon, w_0, w_1 \in [0,1]$ that minimizes the accuracy of the MAP adversary. In the following section, we will use this general PDD Gaussian noise adding privatization scheme in our data-driven simulations and compare the performance of the privacy mechanisms obtained by both theoretical and data-driven approaches.

Data-driven Approach

To illustrate our data-driven GAP approach, we assume the privatizer only has access to the dataset $\mathcal{D}$ but does not know the joint distribution of (X, Y). Finding the optimal privacy mechanism becomes a learning problem. In the training phase, we use the empirical log-loss function $L_{\text{XE}}(h(g(X, Y; \Theta_p); \Theta_a), Y)$ provided in (5.12) for the adversary. Thus, for a fixed privatizer parameter Θ_p, the adversary learns the optimal parameter Θ_a^* that maximizes $-L_{\text{XE}}(h(g(X, Y; \Theta_p); \Theta_a), Y)$. On the other hand, the optimal parameter for the privacy mechanism is obtained by solving (5.10). After convergence, we use the learned data-driven GAP mechanism to compute the accuracy of inferring the private variable under a strong MAP adversary. We evaluate

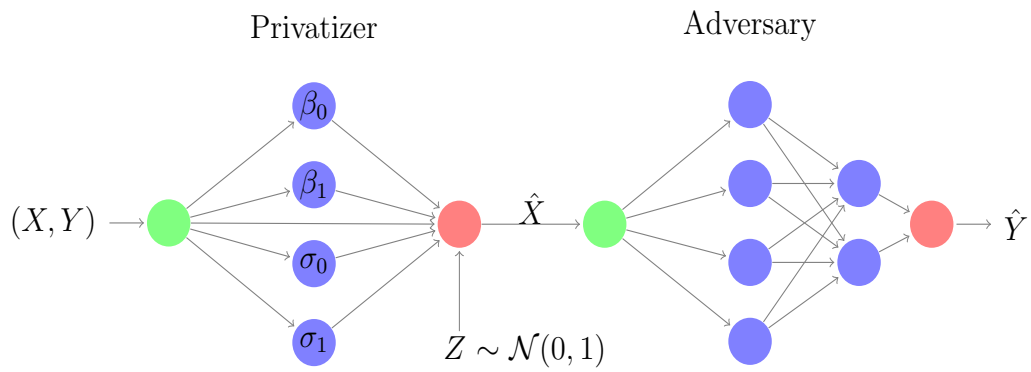

Figure 5.4: Neural network structure of GAP for single-dimensional Gaussian mixture data

our data-driven approach by comparing the mechanisms learned in an adversarial fashion on $\mathcal{D}$ with the game-theoretically optimal ones in which both the adversary and privatizer are assumed to have access to $P(X, Y)$.

We consider the PDD Gaussian noise adding privacy mechanism given by $g(X, Y) = X + (1 - Y)\beta_0 - Y\beta_1 + (1 - Y)\gamma_0 N + Y\gamma_1 N$. Similar to the binary setting, we use two neural networks to model the privatizer and the adversary. As shown in Figure 5.4, the privatizer is modeled by a two-layer neural network with parameters $\beta_0, \beta_1, \gamma_0, \gamma_1 \in \mathbb{R}$. The adversary, whose goal is to infer Y from privatized data $\hat{X}$, is modeled by a three-layer neural network classifier with leaky ReLU activations. The random noise is drawn from a standard Gaussian distribution $N \sim \mathcal{N}(0, 1)$.

In order to enforce the distortion constraint, we use the augmented Lagrangian method to penalize the learning objective when the constraint is not satisfied. In the binary Gaussian mixture model setting, the augmented Lagrangian method uses two parameters, namely λ_t and ρ_t to approximate the constrained optimization problem by a series of unconstrained problems. Intuitively, a large value of ρ_t enforces the distortion constraint to be binding, whereas λ_t is an estimate of the Lagrangian multiplier. To obtain the optimal solution of the constrained optimization problem, we solve a series of unconstrained problems given by (5.15).

Dataset	$P(Y=1)$	$X\|Y=0$	$X\|Y=1$
1	0.5	$\mathcal{N}(-3,1)$	$\mathcal{N}(3,1)$
2	0.5	$\mathcal{N}(-3,4)$	$\mathcal{N}(3,1)$
3	0.75	$\mathcal{N}(-3,1)$	$\mathcal{N}(3,1)$
4	0.75	$\mathcal{N}(-3,4)$	$\mathcal{N}(3,1)$

Table 5.1: Synthetic datasets for binary Gaussian mixture model

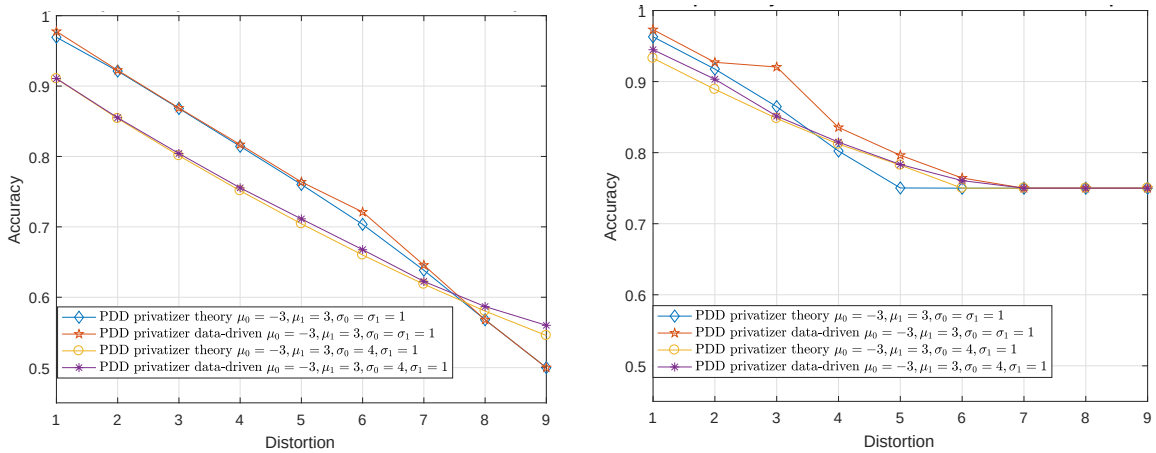

(a) $\tilde{p}=0.5$ (b) $\tilde{p}=0.75$

Figure 5.5: Performance of PDD mechanisms against MAP adversary*

Illustration of Results

We use synthetic datasets to evaluate our proposed GAP framework. We consider four synthetic datasets shown in Table 5.1. Each synthetic dataset used in this experiment contains $20,000$ training samples and $2,000$ test samples. We use Tensorflow to train both the privatizer and the adversary using Adam optimizer with a learning rate of 0.01 and a minibatch size of 200.

Figure 5.5a and 5.5b illustrate the performance of the optimal PDD Gaussian noise adding mechanisms against the strong theoretical MAP adversary when $P(Y=1)=0.5$ and $P(Y=1)=0.75$, respectively. It can be seen that the optimal mecha-

nisms obtained by both theoretical and data-driven approaches reduce the inference accuracy of the MAP adversary as the distortion increases. Similar to the binary data model, we observe that the accuracy of the adversary saturates when the distortion crosses some threshold. Moreover, it is worth pointing out that for the binary Gaussian mixture setting, we also observe that the privacy mechanism obtained through the data-driven approach performs very well when pitted against the MAP adversary (maximum accuracy difference around 6% compared with theoretical approach). In other words, for the binary Gaussian mixture model, the data-driven approach for GAP can generate privacy mechanisms that are comparable, in terms of performance, to the theoretical approach, which assumes the privatizer has access to the underlying distribution of the data.

5.3.2 GAP for Multi-dimensional Gaussian Mixture Models

In this section, we focus on a setting where $Y \in \{0, 1\}$ and X is an m-dimensional Gaussian mixture random vector whose mean is dependent on Y. Let $P(Y = 1) = \tilde{p}$, $X|Y = 0 \sim \mathcal{N}(-\mu, \Sigma)$, and $X|Y = 1 \sim \mathcal{N}(\mu, \Sigma)$, where $\mu = (\mu_1, ..., \mu_m)$. Without loss of generality, we assume that $X|Y = 0$ and $X|Y = 1$ have the same covariance Σ.

We consider a MAP adversary who has access to $P(X, Y)$ and the privacy mechanism. The privatizer's goal is to privatize X in a way that minimizes the adversary's probability of correctly inferring Y from $\hat{X}$. In order to have a tractable model for the privatizer, we mainly focus on linear (precisely affine) GAP mechanisms $\hat{X} = g(X) = X + Z + \beta$, where Z is an independently generated noise vector. This linear GAP mechanism enables controlling both the mean and covariance of the privatized data. To quantify utility of the privatized data, we use the ℓ_2 dis-

*This simulation is completed with Peter Kairouz and Xiao Chen from Stanford University

tance between X and $\hat{X}$ as a distortion measure to obtain a distortion constraint $\mathbb{E}_{X,\hat{X}}\|X-\hat{X}\|^2 \leq D$.

Theoretical Approach

Consider the setup where both the privatizer and the adversary have access to $P(X,Y)$. Further, let Z be a zero-mean multi-dimensional Gaussian random vector. Although other distributions can be considered, we choose additive Gaussian noise for tractability reasons.

Without loss of generality, we assume that $\beta = (\beta_1, ..., \beta_m)$ is a constant parameter vector and $Z \sim \mathcal{N}(0, \Sigma_p)$. Following similar analysis in [144], we can show that the adversary's probability of detection is given by

$$P_d^{(G)} = \tilde{p}Q\left(-\frac{\alpha}{2} + \frac{1}{\alpha}\ln\left(\frac{1-\tilde{p}}{\tilde{p}}\right)\right) + (1-\tilde{p})Q\left(-\frac{\alpha}{2} - \frac{1}{\alpha}\ln\left(\frac{1-\tilde{p}}{\tilde{p}}\right)\right), \tag{5.44}$$

where $\alpha = \sqrt{(2\mu)^T(\Sigma+\Sigma_p)^{-1}2\mu}$. Furthermore, since $\mathbb{E}_{X,\hat{X}}[d(\hat{X},X)] = \mathbb{E}_{X,\hat{X}}\|X-\hat{X}\|^2 = \mathbb{E}\|Z+\beta\|^2 = \|\beta\|^2 + tr(\Sigma_p)$, the distortion constraint implies that $\|\beta\|^2 + tr(\Sigma_p) \leq D$. To make the problem more tractable, we assume both X and Z are independent multi-dimensional Gaussian random vectors with diagonal covariance matrices. In this case, the optimal privacy mechanism is given by the solution of

$$\begin{aligned} \min_{\beta,\Sigma_p} \quad & (2\mu)^T(\Sigma+\Sigma_p)^{-1}2\mu \\ s.t. \quad & \|\beta\|^2 + tr(\Sigma_p) \leq D. \end{aligned} \tag{5.45}$$

Theorem 10. *Consider GAP mechanisms given by $g(X) = X + Z + \beta$, where X and Z are multi-dimensional Gaussian random vectors with diagonal covariance matrices Σ and Σ_p. Let $\{\sigma_1^2, ..., \sigma_m^2\}$ and $\{\sigma_{p_1}^2, ..., \sigma_{p_m}^2\}$ be the diagonal entries of Σ and Σ_p, respectively. The parameters of the minimax optimal privacy mechanism are*

$$\beta_i^* = 0, \quad {\sigma_{p_i}^*}^2 = \left(\frac{|\mu_i|}{\sqrt{\lambda_0^*}} - \sigma_i^2, 0\right)^+, \forall i = \{1, 2, ..., m\},$$

where λ_0^* *is chosen such that* $\sum_{i=1}^{m}\left(\frac{|\mu_i|}{\sqrt{\lambda_0^*}}-\sigma_i^2\right)^+ = D$. *For this optimal mechanism, the accuracy of the MAP adversary is given by* (5.36) *with* $\alpha = 2\sqrt{\sum_{i=1}^{m}\frac{\mu_i^2}{\sigma_i^2+\left(\frac{|\mu_i|}{\sqrt{\lambda_0^*}}-\sigma_i^2\right)^+}}$.

The proof of Theorem 10 is provided in Appendix J. We observe that the when σ_i^2 is greater than some threshold $\frac{|\mu_i|}{\sqrt{\lambda_0^*}}$, no noise is added to the data on this dimension due to the high variance. When σ_i^2 is smaller than $\frac{|\mu_i|}{\sqrt{\lambda_0^*}}$, the amount of noise added to this dimension is proportional to $|\mu_i|$; this is intuitive since a large $|\mu_i|$ indicates the two conditionally Gaussian distributions are further away on this dimension, and thus, distinguishable. Thus, more noise needs to be added in order to reduce the MAP adversary's inference accuracy.

Data-driven Approach

For the data-driven linear GAP mechanism, we assume the privatizer only has access to the dataset $\mathcal{D}$ with n data samples but not the actual distribution of (X, Y). Computing the optimal privacy mechanism becomes a learning problem. In the training phase, the data holder learns the parameter of the GAP mechanism by competing against a computational adversary modeled by a multi-layer neural network. When convergence is reached, we evaluate the performance of the learned mechanism by comparing with the one obtained from the game-theoretic approach. To quantify the performance of the learned GAP mechanism, we compute the accuracy of inferring Y under a strong MAP adversary that has access to both the joint distribution of (X, Y) and the privacy mechanism.

Since the private variable Y is binary, we measure the training loss of the adversary network by the empirical log-loss function (5.12) For a fixed privatizer parameter Θ_p, the adversary learns the optimal Θ_a^* by maximizing (5.12). For a fixed Θ_a, the

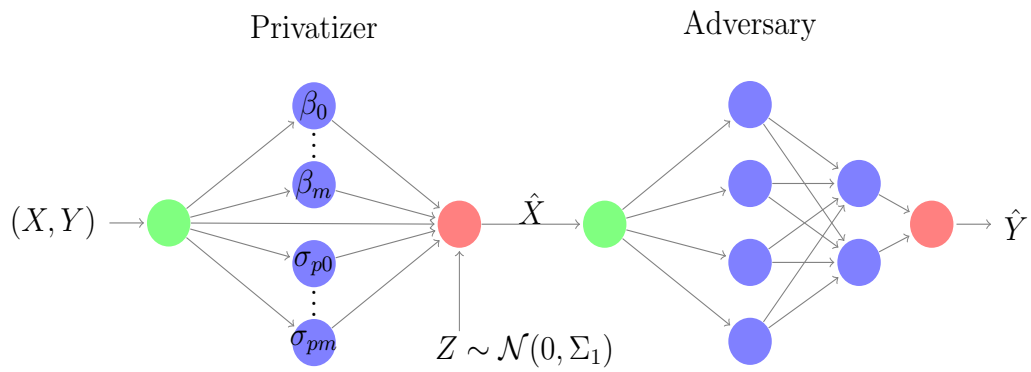

Figure 5.6: Neural network structure of GAP for multi-dimensional Gaussian mixture data

privatizer learns the optimal Θ_p^* by minimizing $-L_n(h(g(X;\Theta_p);\Theta_a), Y)$ subject to the distortion constraint $\mathbb{E}_{X,\hat{X}}\|X - \hat{X}\|^2 \leq D$.

As shown in Figure 5.6, the privatizer is modeled by a two-layer neural network with parameters $\Theta_p = \{\beta_0, ..., \beta_m, \sigma_{p0}, ..., \sigma_{pm}\}$, where β_j and σ_{pj} represent the mean and standard deviation for each dimension $j \in \{1, ..., m\}$, respectively. The random noise Z is drawn from a m-dimensional independent zero-mean standard Gaussian distribution with covariance Σ_1. Thus, we have $\hat{X}_j = X_j + \beta_j + \sigma_{pj} Z_j$. The adversary, whose goal is to infer Y from privatized data $\hat{X}$, is modeled by a three-layer neural network classifier with leaky ReLU activations.

As shown in Figure 5.6, the privatizer is modeled by a two-layer neural network with parameters $\Theta_p = \{\beta_0, ..., \beta_m, \sigma_{p0}, ..., \sigma_{pm}\}$. The adversary, whose goal is to infer Y from privatized data $\hat{X}$, is modeled by a three-layer neural network classifier with leaky ReLU activations. The random noise Z is drawn from a m-dimensional independent zero-mean standard Gaussian distribution with covariance Σ_1.

To incorporate the distortion constraint into the learning process, we add a penalty term to the objective of the privatizer. Thus, the training loss function of the priva-

tizer is given by

$$L(\Theta_p, \Theta_a) = L_n(\Theta_p, \Theta_a) + \rho_t \max\{0, \frac{1}{n}\sum_{i=1}^{n} d(g(x_{(i)}; \Theta_p), x_{(i)}) - D\}, \tag{5.46}$$

where ρ_t is a penalty coefficient which increases with the number of iterations t. The added penalty consists of a penalty parameter ρ_t multiplied by a measure of violation of the constraint. This measure of violation is non-zero when the constraint is violated. Otherwise, it is zero.

Illustration of Results

We use synthetic datasets to evaluate the performance of the learned GAP mechanisms. Each dataset contains $20,000$ training samples and $2,000$ test samples. Each data entry is sampled from an independent multi-dimensional Gaussian mixture model. We consider two categories of synthetic datasets with $P(Y = 1)$ equals to 0.75 and 0.5, respectively. Both the privatizer and the adversary in the GAP framework are trained on Tensorflow [142] using Adam optimizer with a learning rate of 0.005 and a minibatch size of $1,000$. The distortion constraint is enforced by the penalty method as detailed in (5.14).

Figure 5.7 illustrates the performance of the learned GAP mechanism against a strong theoretical MAP adversary for $\tilde{p} = 0.75$. It can be seen that the inference accuracy of the MAP adversary reduces as the distortion increases and asymptotically approaches (as expected) the prior on the private variable. This is because noise adding mechanisms cannot further reduce the accuracy of the MAP adversary than the prior on Y. We also observe that the privacy mechanism obtained via the data-driven approach performs very well when pitted against the MAP adversary (maximum accuracy difference around 0.3% compared to the theoretical approach).

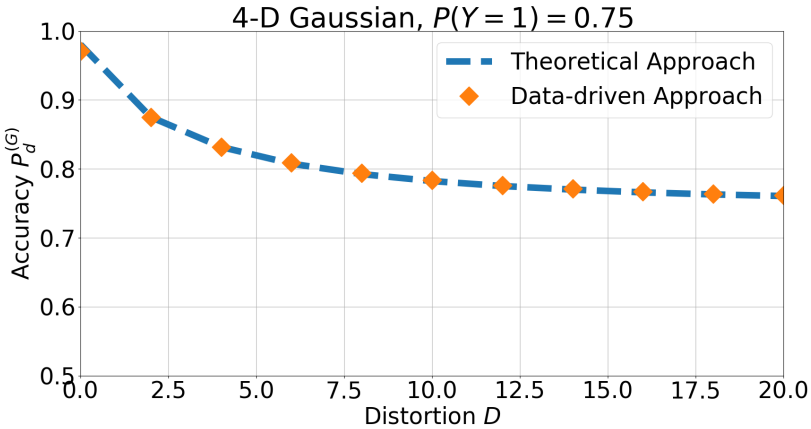

(a) 4-D Gaussian mixture

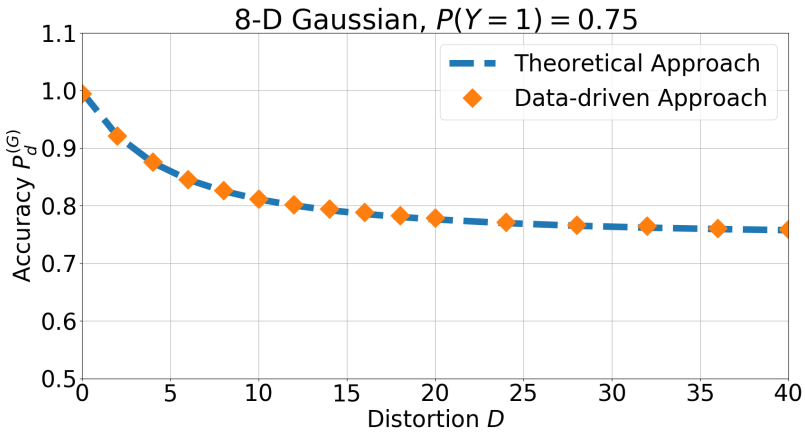

(b) 8-D Gaussian mixture

Figure 5.7: Performance of learned GAP mechanisms against MAP adversary ($\tilde{p} = 0.75$)

The performance of the learned GAP mechanism against a strong theoretical MAP adversary for $\tilde{p} = 0.5$ is illustrated in Figure 5.8. Similar to the case in which $\tilde{p} = 0.75$, we also observe that the privacy mechanism obtained via the data-driven approach performs very well when pitted against the MAP adversary (maximum accuracy difference around 0.8% compared to the theoretical approach). In other words, for the Gaussian mixture data model with binary private variable, the data-driven version of GAP can learn privacy mechanisms that perform as well as the

mechanisms computed under the theoretical version of GAP, which assumes that the privatizer has access to the underlying distribution of the dataset.

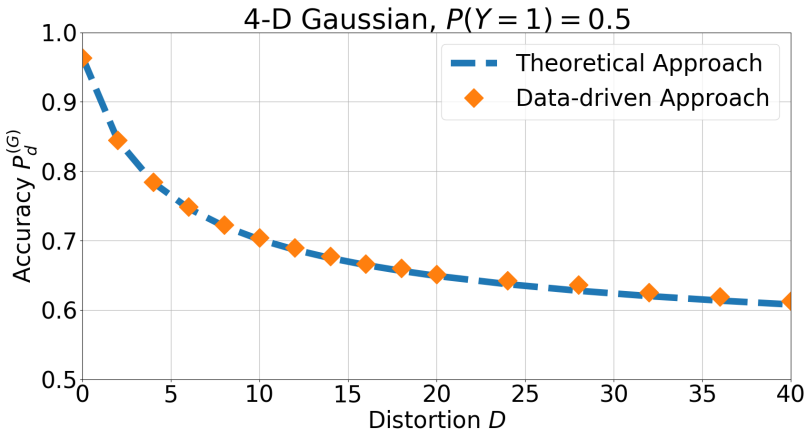

(a) 4-D Gaussian mixture

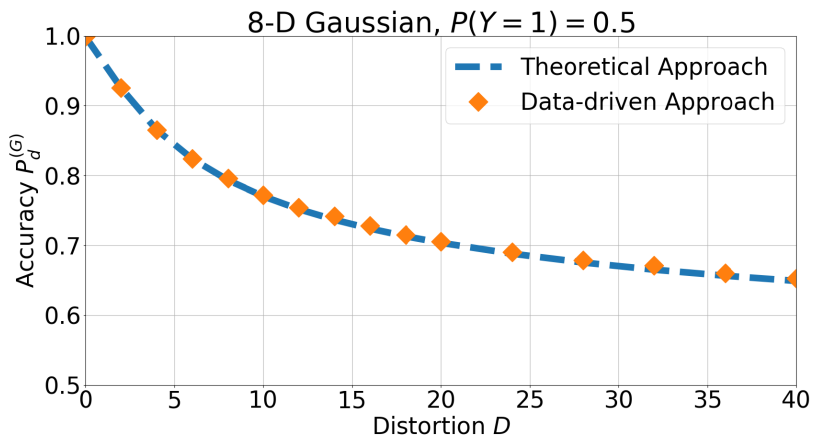

(b) 8-D Gaussian mixture

Figure 5.8: Performance of learned GAP mechanisms against MAP adversary ($\tilde{p} = 0.5$)

5.4 GAP for Real Datasets

We apply the proposed GAP framework to two different datasets to demonstrate its capabilities. First of all, we apply the data-driven GAP to the GENKI dataset [107] which contains $1,940$ greyscale face images. Then, we consider the MNIST dataset [108] which contains $70,000$ images of hand-written digits. We choose cross entropy to be the loss function for the adversary and use the penalty method introduced in

Section 5.1.3 to enforce the distortion constraint. The privatizer is trained and tested in an adversarial fashion using Tensorflow.

5.4.1 The GENKI Dataset

The GENKI dataset consists of $1,940$ face images with different facial expressions. Each data sample is a 16×16 greyscale image. We choose $N = 1,740$ training samples (50% male and 50% female). Among each gender group, we have 50% smile and 50% non-smile faces. The test dataset contains 200 samples (50% male and 50% female; 50% smile and 50% non-smile). We consider gender as private variable Y and the image pixels as public variable X. Our goal is to learn a GAP mechanism that restricts inferences on gender with limited distortion.

Privatizer Model

In this experiment, we consider two different privatizer architectures: the feedforward neural network privatizer (FNNP) and the transposed convolutional neural network privatizer (TCNNP). The FNNP architecture uses a multi-layer feedforward neural network to combine the low-dimensional random noise and the original image together (Figure 5.9). The TCNNP takes a low-dimensional random noise vector and generates a high-dimensional noise mask using multi-layer transposed convolutional neural networks. The noise mask is added to the original image to generate the privatized image (Figure 5.9).

The FNNP is modeled by a 4-layer feedforward neural network. We first reshape each image to a long vector (256×1), and then concatenate it with a 100×1 Gaussian random noise vector. Each entry in the noise vector is sampled independently from a standard Gaussian distribution. We feed the entire vector to a 4-layer fully connected (FC) neural network. Each layer has 256 neurons with leaky ReLU activation. Finally, we reshape the output of the last layer to a 16×16 greyscale image.

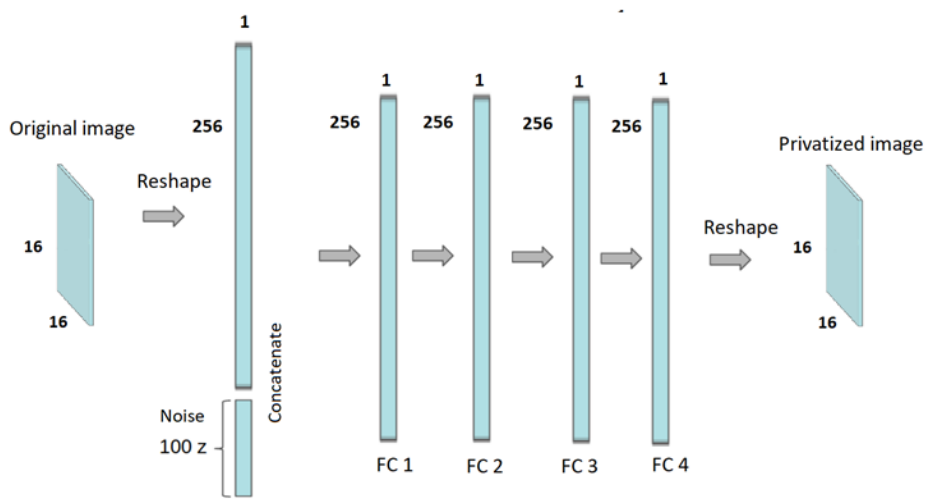

Figure 5.9: Feedforward neural network privatizer

To model the TCNNP, we first generate a 100×1 dimension standard Gaussian noise vector and use linear projection to map the noise vector to a $4 \times 4 \times 256$ feature tensor. The feature tensor is then fed to an initial transposed convolutional layer (DeCONV) with 128 filters (filter size 3×3, stride 2) and ReLU activation, followed by another transposed convolutional layer with 1 filter (filter size 3×3, stride 2) and tanh activation. We add batch normalization [145] to each hidden layer to prevent covariance shift and help gradients to flow. The output of the second transposed convolutional layer is added to the original image to generate the privatized data.

Adversary Model

In our data-driven GAP, we model the adversary using state-of-the-art convolutional neural networks (CNNs). This architecture outperforms most of other models for image classification [146, 147, 148, 149]. In this experiment, the adversary is model by a 7-layer CNN (Figure 5.11). The privatized images are fed to two convolutional layers (CONV) whose sizes are $3 \times 3 \times 32$ and $3 \times 3 \times 64$, respectively. Each convolutional layer is followed by batch normalization and ReLU activation. The output of each

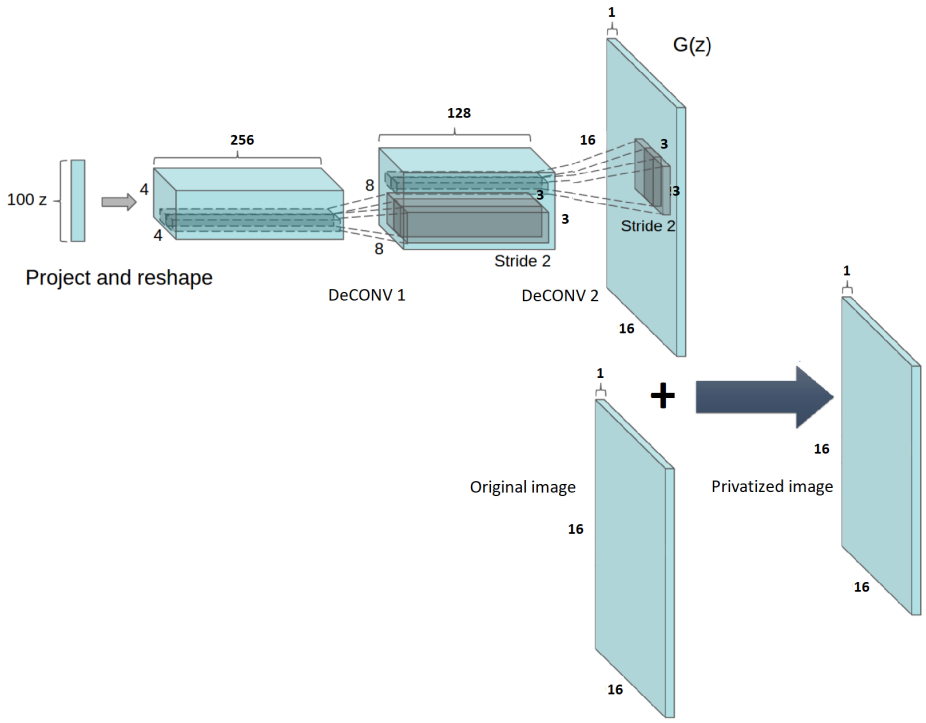

Figure 5.10: Transposed convolutional neural network privatizer

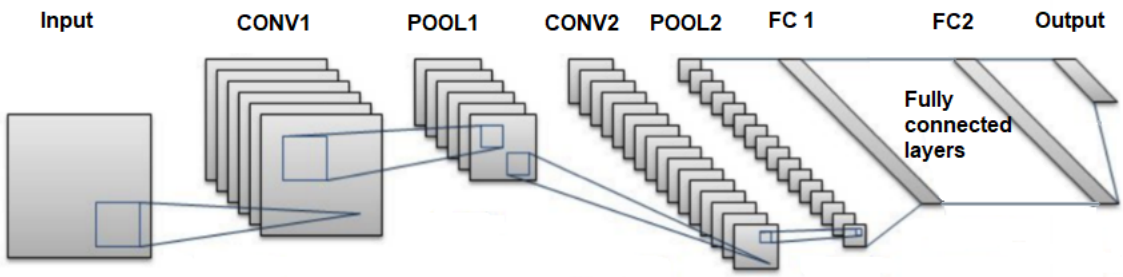

Figure 5.11: Convolutional neural network adversary

convolutional layer is then fed to a 2×2 maxpool layer (POOL) to generate features for classification. The second maxpool layer is followed by two fully-connected layers, which contain 1024 neurons with batch normalization and ReLU activation. Finally, the output of the fully-connected layers are mapped to the output layer, which contains two neurons capturing the belief of the subject being a male or a female.

Illustration of Results

Figure 5.12 illustrates the gender classification accuracy of the adversary for different values of distortion. It can be seen that the adversary's accuracy of classifying the private label (gender) decreases progressively as the distortion increases. Given the same distortion value, FNNP achieves better privacy compared to TCNNP: when the distortion is small (0.0039 per pixel), the adversary's classification accuracy is already reduced to 80% and 61% by using the TCNNP and the FNNP architecture, respectively. When we increase the distortion value to 0.0195, the classification accuracy further decreases to 60% and 50.5%, respectively. The intuition behind this is that the FNNP uses both the noise vector and the original image to generate the privatized image. However, the TCNNP generates the noise mask independent of the original image pixels and add the noise mask to the original image in the final step. To demonstrate the effectiveness of the learned GAP mechanisms, we plot the gender classification accuracy for the dataset privatized by the learned GAP mechanisms as well as adding independent uniform or Laplace noise. It can be seen that for the same distortion, the learned GAP mechanisms achieve much lower gender classification accuracy than using uniform or Laplace noise.

To study the influence of GAP on other non-private classification tasks, we train another CNN (see Figure 5.11) to perform facial expression classification on datasets privatized by different privacy mechanisms. Figure 5.13 illustrates the facial expression classification accuracy for different values of distortion. It can be seen that the accuracy of the expression classification decreases slowly as the distortion increases. However, even for a large distortion value (0.019 per pixel), the expression classification accuracy only decreases by 13% at most. Furthermore, we observe that given

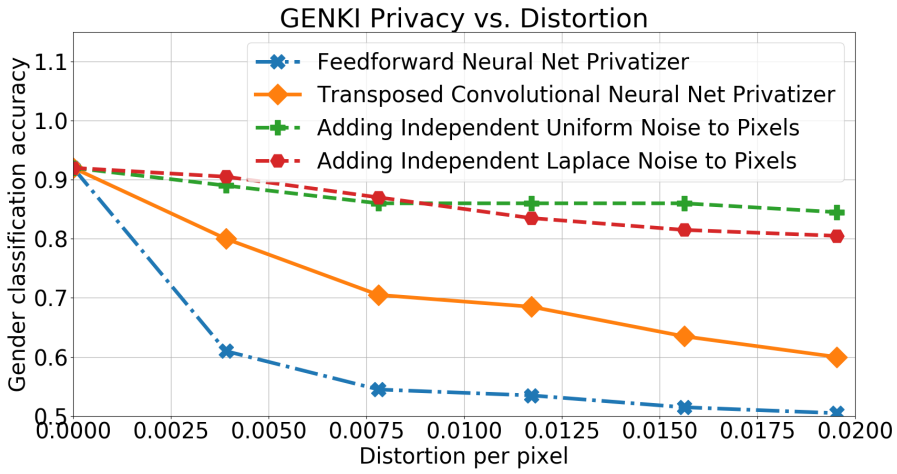

Figure 5.12: Gender classification accuracy for different distortion values

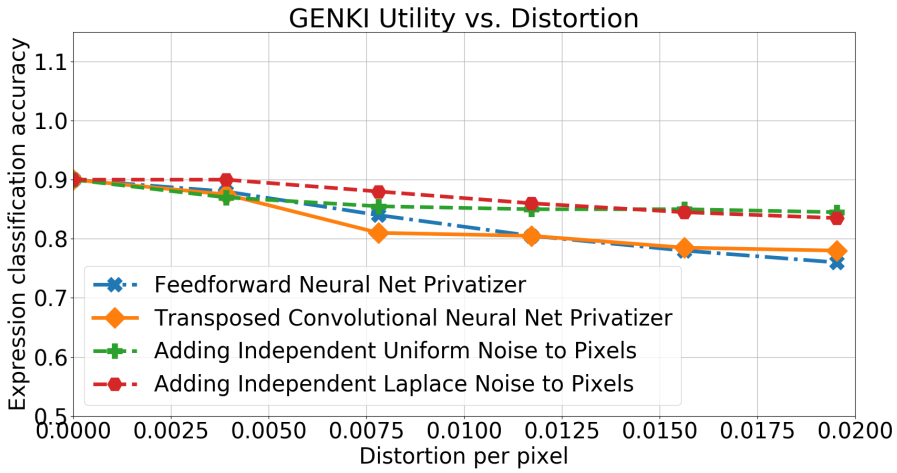

Figure 5.13: Facial expression classification accuracy for different distortion values

the same distortion value, the FNNP and TCNNP achieve similar facial expression classification accuracy.

In both Figure 5.12 and Figure 5.13, we observe that when the distortion value is small, adding Laplace noise yields higher accuracy than uniform noise in both gender and expression classification. However, when the distortion value becomes large, the uniform noise yields higher accuracy in both gender and facial expression classification. This is due to the fact that for the same distortion (variance of the noise), the Laplace distribution is very spiky when the distortion value is small. As

a result, the noise values added to pixels are concentrated around a small region centered at 0. However, when the distortion value is large, the Laplace noise becomes more spread out. As a result, larger noise values are more likely to be added to the pixels and thus help reduce the classification accuracy.

The privatized images using FNNP and TCNNP architectures under different distortion values are shown in Figure 5.14 and Figure 5.15. We observe that given the same distortion value, the adversary makes more mistakes when the data is privatized by the FNNP architecture. Furthermore, both privatizers change mostly eyes, nose, mouth, beard, and hair. We also observe that the outputs of the FNNP look more smooth.

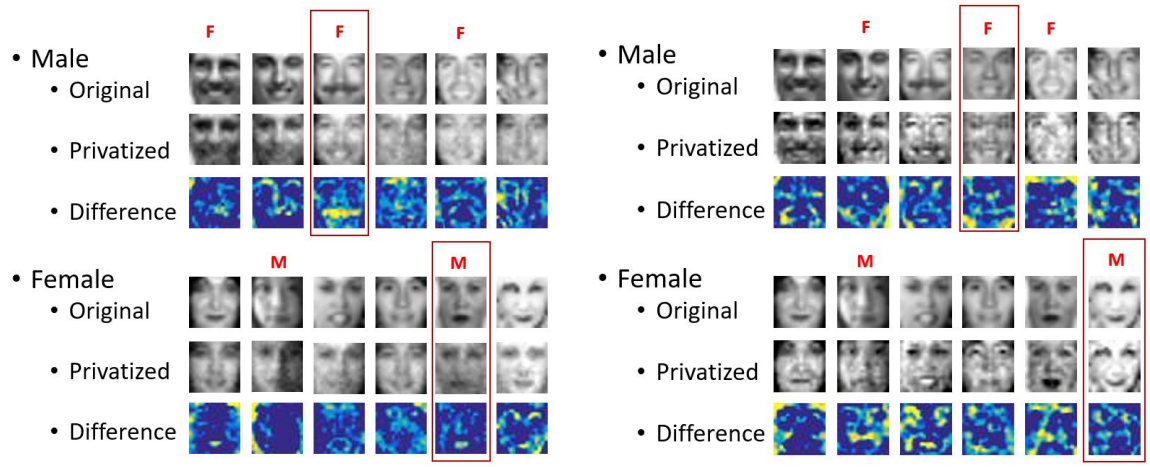

(a) Feedforward nerual network privatizer

(b) Transposed convolutional nerual network privatizer

Figure 5.14: Privatized images with 0.0117 per pixel distortion

5.4.2 The MNIST Dataset

The MNIST dataset consists of $70,000$ images of hand-written digits. Each data sample is a 28×28 greyscale image. The dataset is divided into $60,000$ training samples and $10,000$ test samples. We model the private feature as a binary variable

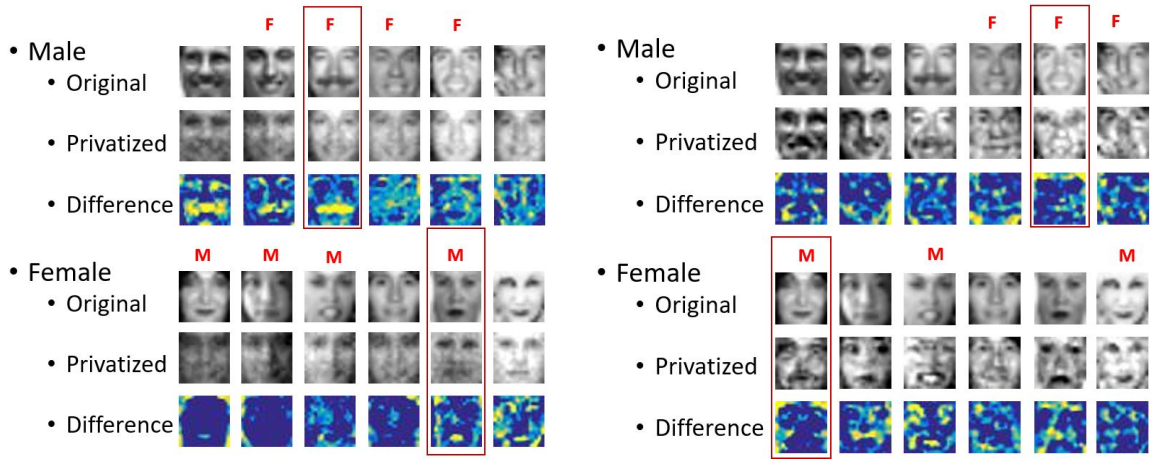

(a) Feedforward nerual network privatizer

(b) Transposed convolutional nerual network privatizer

Figure 5.15: Privatized images with 0.0195 per pixel distortion

Y which identifies whether there is a circular structure in the digits (e.g., $0, 6, 8, 9$ contain circular structure). The image pixels are considered as public variable X.

Privatizer and Adversary Models

Note that the images in MNIST have more pixels than GENKI (28×28 *vs.* 16×16). However, due to the nature of hand-written digits, we assume that the images are concentrated on a much lower dimension manifold [126]. In this experiment, we choose the CNN shown in Figure 5.11 as the adversary. For the privatizer, we first use a CNN to map each image to a low-dimensional feature vector. Then, we use a multi-layer feedforward neural network to combine the low-dimensional random noise with the feature vector. Finally, we generate high-dimensional images using a multi-layer transposed convolutional neural network.

Figure 5.16 illustrates the architecture of the privatizer. The original images are fed to two convolutional layers whose sizes are $3 \times 3 \times 32$ and $3 \times 3 \times 64$, respectively. The output of each convolutional layer is then fed to a 2×2 maxpool layer. The

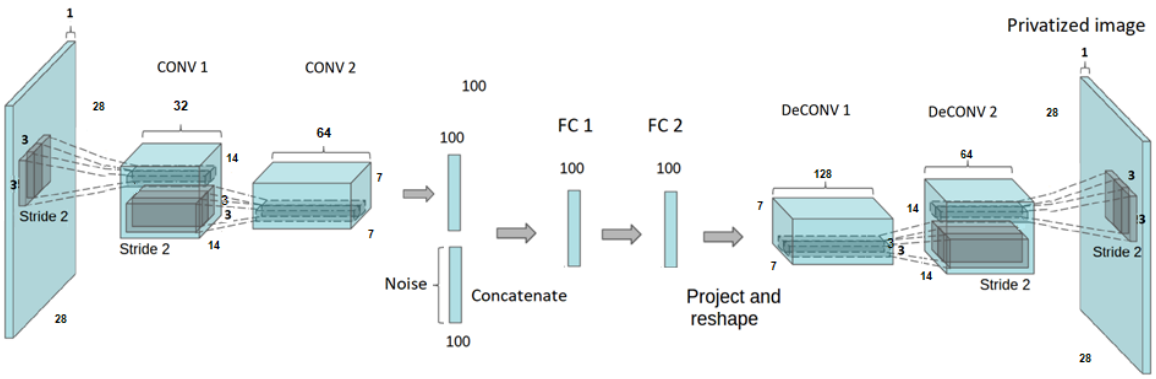

Figure 5.16: MNIST privatizer structure

second maxpool layer is followed by three fully-connected layers, which contain 1024, 256 and 100 neurons, respectively. We use the output of the third fully-connected layer as the low-dimensional feature vector for the image. Each convolutional and fully-connected layer uses leaky ReLU activation. After obtaining the feature vector, we concatenate it with a 100×1 independently generated Gaussian random noise and feed the entire vector to a 2-layer fully-connected neural network. Each layer has 100 hidden neurons. Then, we use linear projection to map the noised feature vector to a $7 \times 7 \times 128$ feature tensor. The feature tensor is then fed to an initial transposed convolutional layer with 64 filters (filter size 3×3, stride 2) and leaky ReLU activation, followed by another transposed convolutional layer with 1 filter (filter size 3×3, stride 2) and tanh activation. We add batch normalization to each hidden layer to prevent covariance shift and help gradients to flow. The output of the second transposed convolutional layer is the privatized data.

Illustration of Results

Figure 5.17 illustrates the private variable classification accuracy of the adversary for different values of distortion. It can be seen that for the dataset privatized by the GAP mechanism, the adversary's accuracy of classifying the private label, i.e., whether there is a circular structure in the digit, decreases progressively as the distortion

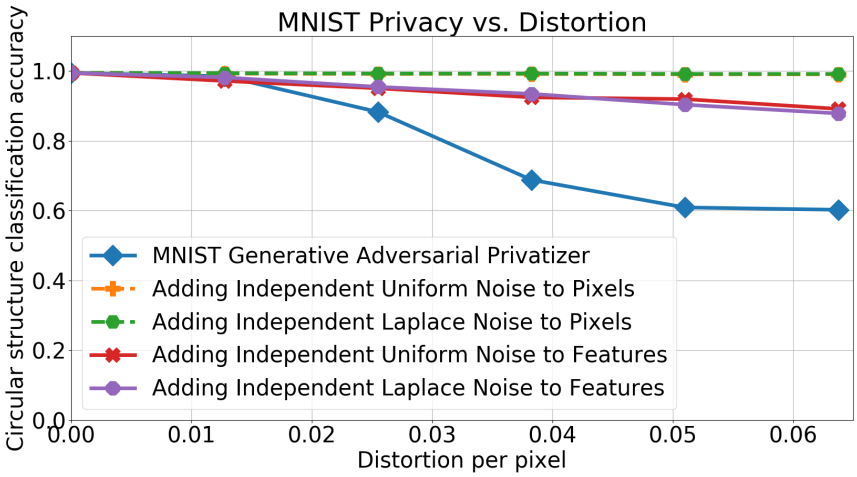

Figure 5.17: Circular structure classification accuracy for different distortion values

increases. We observe that the private variable classification drops to 60% even for a very small per pixel distortion (0.051). Furthermore, adding uniform or Laplace noise to either each pixel or the extracted feature vector does not prevent the adversary from learning the private feature effectively.

To study the influence of the learned GAP mechanism on non-private classification tasks, we train another CNN (see Figure 5.11) to classify the value of each digit using the privatized dataset. Figure 5.18 shows the digit value classification accuracy for different privacy mechanisms. We also observe that for a dataset privatized by the GAP mechanism, the digit value classification accuracy decreases as the distortion increases. Even if the private variable classification accuracy drops to 60% at 0.0663 per pixel distortion, the CNN trained on the privatized dataset can still achieve 58% digit value classification accuracy.

The privatized images under different distortion values are shown in Figure 5.19. We observe that the privatizer successfully extracts the features of each digit and adds noise selectively to reduce the inference capability of the adversary. We notice that if the digit contains a circular structure, the privatizer tends to break it. Otherwise, the privatizer tries to complete a circular structure.

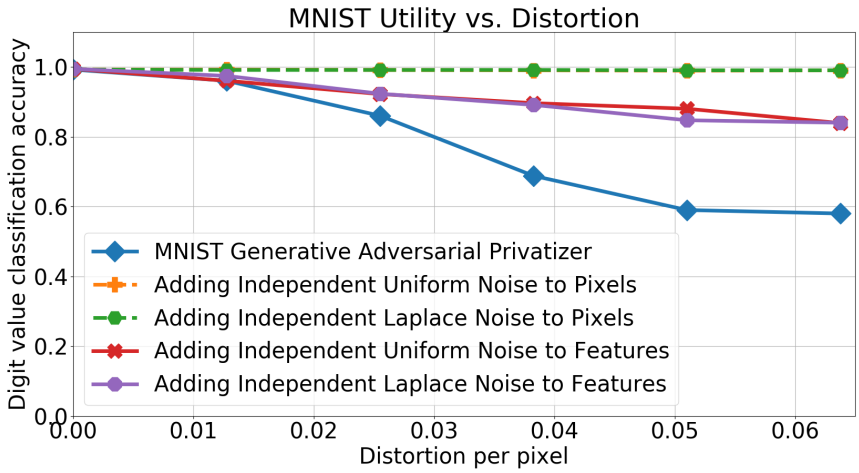

Figure 5.18: Digit value classification accuracy for different distortion values

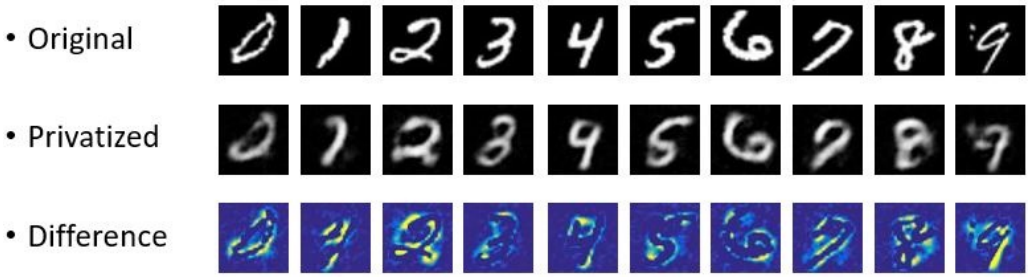

(a) Privatized images with 0.0255 per pixel distortion

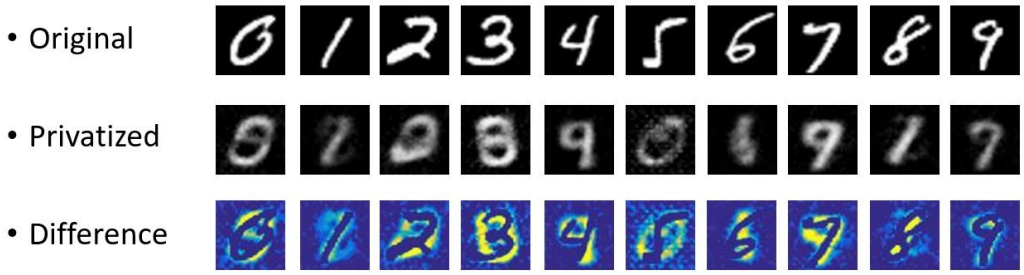

(b) Privatized images with 0.0510 per pixel distortion

Figure 5.19: MNIST privatized images for different distortion values

Chapter 6

CONCLUSIONS AND FUTURE WORK

This dissertation studies two fundamental problems: (i) decision making for interactions between retailers/service providers and privacy sensitive users; (ii) a unified, data-driven framework for various information-theoretic privacy. In the first problem, we have studied privacy-utility tradeoff in different scenarios. The most significant contribution is the idea that retailers/service providers can develop strategies to encourage users to interact with them (e.g., via using coupons distributed by retailers or using services provided by service providers) while taking their privacy sensitivities into account. We have also investigated the influence of privacy on free online service market. For the second problem, we have proposed a novel context-aware privacy framework called generative adversarial privacy (GAP). GAP captures a variety of information-theoretic privacy notions via a minimax game and allows the data holders to learn privacy mechanisms from data directly. We now provide some more specific comments on conclusions and future directions in each of the problems we have studied.

6.1 How to Incentivize and Interact with Privacy Sensitive Consumer?

We have proposed a POMDP model to capture the interactions between a retailer and a privacy-sensitive consumer in the context of personalized shopping. The retailer seeks to minimize the expected discounted cost of violating the consumer's privacy. We have shown that the optimal coupon-offering policy is a stationary policy that takes the form of an explicit threshold that depends on the model parameters. In summary, the retailer offers an HP coupon when the Normal to Alerted transition probability is low or the probability of staying in Alerted state is high. Furthermore, the threshold optimal policy also holds for consumers whose privacy sensitivity can be captured via multiple alerted states as well as for the case in which consumers

exhibit coupon-dependent transition. For the case in which the cost feedbacks from the consumer are noisy, we have introduced a heuristic method using the mean value of costs to compute the decision threshold. Furthermore, under noisy cost feedbacks scenario, we have introduced a Bayesian data analysis approach for decision making by estimating consumerbelief state when the initial belief state is unknown to the retailer.

6.2 Incentive Mechanisms for Privacy-Sensitive Electricity Consumers

We have introduced a novel approach to study the tradeoff between privacy and energy cost minimization for consumers under the assumption that the electricity provider offers incentives to consumers for encouraging them to compromise a certain level of privacy for stable and economic grid operation. A non-cooperative game-theoretic model has been developed to capture interactions between consumers and the electricity provider. With access to alternative energy sources, privacy-sensitive consumers can choose the fraction of electricity they consume from the grid to mask their consumption behavior. On the other hand, the strategy of the electricity provider is to use incentives to encourage consumers to consume a desired amount of electricity consistent with its supply. In particular, we have studied the mixed strategy Nash equilibrium. In the two-player scenario, we have proved the existence and uniqueness of the nondegenerate mixed strategy Nash equilibrium. For a specific choice of profit and valuation functions, our illustrations have shown that the proposed incentive mechanism both increases the net profit and reduces supply-demand imbalance loss of the electricity provider. Furthermore, consumers also benefit from this mechanism for electricity cost reduction.

6.3 Impact of Privacy on Free Online Service Markets

Our work seeks to understand the effect of offering privacy- and QoS- differentiated online services on consumers with heterogeneous privacy sensitivities. We have

quantified this effect as the fraction of consumers that prefer lower privacy risks with the accompanying lower QoS to the alternative of higher risks and higher QoS. We have presented an analysis built upon the classical Hotelling model to compute equilibrium QoS-privacy risk strategies and market segmentation for the two-SP problem. Analogous to the classical segmentation models, our problem also involves parameters that capture cost, revenue, and consumer valuation functions that are dependent and independent of privacy risks. While such a parametrized model can make the analysis challenging, our results for relatively simple yet meaningful functions such as linear cost models and uniform (as well as truncated Gaussian) distribution of consumer preferences suggest that SPs that have higher profits from untargeted services have an edge in the market. SPs competing on offering higher privacy risk services have to offer better QoS or use other means of increasing untargeted revenue to gain market share. Our work also shows the instability of such market with more than two SPs.

6.4 Generative Adversarial Privacy

We have presented a unified framework for context-aware privacy called generative adversarial privacy (GAP). GAP allows the data holder to learn the privatization mechanism directly from the dataset (to be published) without requiring access to the dataset statistics. Under GAP, finding the optimal privacy mechanism is formulated as a game between two players: a privatizer and an adversary. An iterative minimax algorithm is proposed to obtain the optimal mechanism under the GAP framework.

To evaluate the performance of the proposed GAP model, we first focus on two types of data models: (i) binary data model; and (ii) binary Gaussian mixture model. For both cases, the optimal GAP mechanisms are learned using an empirical log-loss function. For each type of dataset, both private-data dependent and private-data independent mechanisms are studied. These results are cross-validated against the privacy guarantees obtained by computing the game-theoretically optimal mechanism

under a strong MAP adversary. In the MAP adversary setting, we have shown that for the binary data model, the optimal GAP mechanism is obtained by solving a linear program. For the binary Gaussian mixture model, the optimal additive Gaussian noise privatization scheme is determined. Simulations with synthetic datasets for both types (i) and (ii) show that the privacy mechanisms learned via the GAP framework perform as well as the mechanisms obtained from theoretical computation. We have also validated the performance of GAP on real datasets such as GENKI and MNIST.

6.5 Future Work

Although methods for designing incentive schemes for privacy-sensitive users and context-aware privacy preserving mechanisms are studied in this report, the work accomplished so far has just provided a few possible solutions to addressing such interesting problems. To better understand the tradeoff between acquiring information and maximizing revenue and approaches for privacy preserving data sharing/publishing, more work needs to be done. We propose to pursuit the following directions for the problems considered in this dissertation.

For the retailer-consumer interaction problem, one straightforward extension of our work is to model uncertainties in the statistical model for the consumer transition probabilities. Further a field, one can also develop game-theoretic models to study the interaction between a retailer and strategic consumers and develop methods to test those models in practice.

For the electricity provider-consumer problem, one of the interesting directions is to develop dynamic game models to capture interactions between consumers and the electricity provider over a certain period of time. Another avenue is to use prospect theory to study subjective behavior of the electricity provider and consumers.

The market segmentation model assumes at least two or more SPs were able to overcome the barrier to entry and differentiate themselves. An immediate question

we will address going forward is whether such barriers to entry are in fact surmountable when competitors use privacy as a differentiator. Also, extending the model to capture externalities of using private data could lead to interesting insights into real-world market interactions. Another challenge to address is to develop models to capture privacy risks that are not directly observable to consumers. These analyses are crucial for developing better privacy policies to effectively enable safe and secure online commerce.

For the generative adversarial privacy problem, there are several fundamental questions that we seek to address. An immediate one is to develop techniques to rigorously benchmark data-driven results for large datasets against computable theoretical guarantees. The proposed data-driven version of GAP is a learning-based approach trained on finitely many training samples. Thus, generalization bounds for the data-driven GAP which provide guarantees on the performance of the learned mechanism on unseen test samples are needed. Furthermore, it will be interesting to investigate the influence of different privatizer and adversary structures on the performance of GAP. Finally, it will be also interesting to compare our approach to a context-free notion of privacy such as differential privacy.

REFERENCES

[1] K. Hill, "How Target figured out a teen girl was pregnant before her father did," *[online] Availabe at: http://www. forbes. com/sites/kashmirhill/2012/02/16/how-target-figured-outa-teen-girl-was-pregnant-before-her-father-did/(Accessed July 4th, 2012)*, 2012. 1.1

[2] C. Duhigg, "How companies learn your secrets," *The New York Times*, vol. 16, p. 2012, 2012. 1.1.1

[3] C. R. Taylor, "Consumer privacy and the market for customer information," *RAND Journal of Economics*, pp. 631–650, 2004. 1.1.3

[4] R. T. Rust, P. Kannan, and N. Peng, "The customer economics of internet privacy," *Journal of the Academy of Marketing Science*, vol. 30, no. 4, pp. 455–464, 2002. 1.1.3

[5] A. Acquisti and J. Grossklags, "Privacy and rationality in individual decision making," *IEEE Security and Privacy*, vol. 2, pp. 24–30, 2005. 1.1.3

[6] A. Acquisti, "The economics of personal data and the economics of privacy," *Background Paper for OECD Joint WPISP-WPIE Roundtable*, vol. 1, 2010. 1.1.3

[7] L. Sankar, S. Kar, R. Tandon, and H. V. Poor, "Competitive privacy in the smart grid: An information-theoretic approach," in *Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on*, pp. 220–225, IEEE, 2011. 1.1.3

[8] C. Dwork, "Differential privacy," in *Encyclopedia of Cryptography and Security*, pp. 338–340, Springer, 2011. 1.1.3

[9] W. Wang, L. Ying, and J. Zhang, "The value of privacy: Strategic data subjects, incentive mechanisms and fundamental limits," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 44, pp. 249–260, ACM, 2016. 1.1.3, 1.3.1

[10] A. Ghosh and A. Roth, "Selling privacy at auction," *Games and Economic Behavior*, 2013. 1.1.3

[11] J. Hsu, Z. Huang, A. Roth, T. Roughgarden, and Z. S. Wu, "Private matchings and allocations," *arXiv preprint arXiv:1311.2828*, 2013. 1.1.3

[12] E. A. Feinberg, A. Shwartz, and E. Altman, *Handbook of Markov decision processes: methods and applications*. Kluwer Academic Publishers Boston, MA, 2002. 1.1.4

[13] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*, vol. 414. John Wiley & Sons, 2009. 1.1.4

[14] G. M. Lipsa and N. C. Martins, "Remote state estimation with communication costs for first-order lti systems," *Automatic Control, IEEE Transactions on*, vol. 56, no. 9, pp. 2013–2025, 2011. 1.1.4

[15] A. Nayyar, T. Basar, D. Teneketzis, and V. V. Veeravalli, "Optimal strategies for communication and remote estimation with an energy harvesting sensor," *Automatic Control, IEEE Transactions on*, vol. 58, no. 9, pp. 2246–2260, 2013. 1.1.4

[16] S. M. Ross, "Quality control under markovian deterioration," *Management Science*, vol. 17, no. 9, pp. 587–596, 1971. 1.1.4

[17] A. Laourine and L. Tong, "Betting on gilbert-elliot channels," *Wireless Communications, IEEE Transactions on*, vol. 9, no. 2, pp. 723–733, 2010. 1.1.4

[18] I. MacPhee and B. Jordan, "Optimal search for a moving target," *Probability in the Engineering and Informational Sciences*, vol. 9, no. 02, pp. 159–182, 1995. 1.1.4

[19] P. Mansourifard and T. Javidi, "Tracking of real-valued continuous markovian random processes with asymmetric cost and observation." [online] Available at: `http://anrg.usc.edu/www/papers/Mansourifard_ACC2015.pdf/`, 2014. 1.1.4

[20] P. Venkitasubramaniam, "Privacy in stochastic control: A Markov decision process perspective.," in *Proc. Allerton Conf.*, pp. 381–388, 2013. 1.1.4

[21] J. D. Glover, M. Sarma, and T. Overbye, *Power System Analysis & Design, SI Version*. Cengage Learning, 2011. 1.2

[22] F. Sultanem, "Using appliance signatures for monitoring residential loads at meter panel level," *Power Delivery, IEEE Transactions on*, vol. 6, no. 4, pp. 1380–1385, 1991. 1.2.1

[23] G. W. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992. 1.2.1

[24] M. L. Marceau and R. Zmeureanu, "Nonintrusive load disaggregation computer program to estimate the energy consumption of major end uses in residential buildings," *Energy Conversion and Management*, vol. 41, no. 13, pp. 1389–1403, 2000. 1.2.1

[25] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. Lewis, R. Cepeda, *et al.*, "Privacy for smart meters: Towards undetectable appliance load signatures," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pp. 232–237, IEEE, 2010. 1.2.1

[26] D. Varodayan and A. Khisti, "Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, pp. 1932–1935, IEEE, 2011. 1.2.1

[27] L. Yang, X. Chen, J. Zhang, and H. V. Poor, "Optimal privacy-preserving energy management for smart meters," in *INFOCOM, 2014 Proceedings IEEE*, pp. 513–521, IEEE, 2014. 1.2.1

[28] C. Huang and L. Sankar, "Incentive mechanisms for privacy-sensitive electricity consumers with alternative energy sources," in *Information Science and Systems (CISS), 2016 Annual Conference on*, pp. 175–180, IEEE, 2016. 1.2.1

[29] L. Sankar, S. R. Rajagopalan, S. Mohajer, and H. V. Poor, "Smart meter privacy: A theoretical framework," *smart grid, IEEE transactions on*, vol. 4, no. 2, pp. 837–846, 2013. 1.2.1

[30] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pp. 238–243, IEEE, 2010. 1.2.1

[31] L. J. Ratliff, C. Barreto, R. Dong, H. Ohlsson, A. Cardenas, and S. S. Sastry, "Effects of risk on privacy contracts for demand-side management." arXiv preprint arXiv:1409.7926, 2014. 1.2.1

[32] S. Denic, G. Kalogridis, and Z. Fan, "Privacy vs pricing for smart grids," in *First IARIA International Conference on Smart Grids, Green Communications and IT Energyaware Technologies*, 2011. 1.2.1

[33] S. Borenstein, M. Jaske, and A. Rosenfeld, "Dynamic pricing, advanced metering, and demand response in electricity markets," 2002. 1.2.1

[34] A.-H. Mohsenian-Rad, V. W. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid," *Smart Grid, IEEE Transactions on*, vol. 1, no. 3, pp. 320–331, 2010. 1.2.1

[35] C. Huang and S. Sarkar, "Dynamic pricing for distributed generation in smart grid," in *Green Technologies Conference, 2013 IEEE*, pp. 422–429, IEEE, 2013. 1.2.1

[36] EUGDPR, "The EU general data protection regulation (GDPR)." http://www.eugdpr.org/, 2017. http://www.eugdpr.org/. 1.3, 1.4

[37] Úlfar Erlingsson, V. Pihur, and A. Korolova, "RAPPOR: Randomized aggregatable privacy-preserving ordinal response," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS '14)*, pp. 1054–1067, 2014. 1.3, 1.4.2, 4.1

[38] Apple Inc., "About privacy and location services in ios 8 and later." https://support.apple.com/en-is/HT203033, Sep 2016. 1.3

[39] C. Cadwalladr and E. Graham-Harrison, "Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach," 03 2018. 1.3

[40] G. Shaffer and Z. J. Zhang, "Competitive coupon targeting," *Marketing Science*, vol. 14, no. 4, pp. 395–416, 1995. 1.3.1

[41] Y. Chen and G. Iyer, "Research note consumer addressability and customized pricing," *Marketing Science*, vol. 21, no. 2, pp. 197–208, 2002. 1.3.1

[42] Z. Tang, Y. Hu, and M. D. Smith, "Gaining trust through online privacy protection: Self-regulation, mandatory standards, or caveat emptor," *Journal of Management Information Systems*, vol. 24, no. 4, pp. 153–173, 2008. 1.3.1

[43] J. Campbell, A. Goldfarb, and C. Tucker, "Privacy regulation and market structure," *Journal of Economics & Management Strategy*, vol. 24, no. 1, pp. 47–73, 2015. 1.3.1

[44] V. Conitzer, C. R. Taylor, and L. Wagman, "Hide and seek: Costly consumer privacy in a market with repeat purchases," *Marketing Science*, vol. 31, no. 2, pp. 277–292, 2012. 1.3.1

[45] Y. Chen, C. Narasimhan, and Z. J. Zhang, "Individual marketing with imperfect targetability," *Marketing Science*, vol. 20, no. 1, pp. 23–41, 2001. 1.3.1

[46] R. K. Chellappa and S. Shivendu, "Mechanism design for "free" but "no free disposal" services: The economics of personalization under privacy concerns," *Management Science*, vol. 56, no. 10, pp. 1766–1780, 2010. 1.3.1

[47] A. Datta, M. C. Tschantz, and A. Datta, "Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination," in *Proceedings on Privacy Enhancing Technologies*, vol. 2015, pp. 92–112, apr 2015. 1.3.1

[48] C. Huang, L. Sankar, and A. D. Sarwate, "Incentive schemes for privacy-sensitive consumers," in *International Conference on Decision and Game Theory for Security*, pp. 358–369, Springer, 2015. 1.3.1

[49] C. Huang, L. Sankar, and A. D. Sarwate, "Designing incentive schemes for privacy-sensitive users," *Journal of Privacy and Confidentiality*, vol. 7, no. 1(5), 2016. Available at https://journalprivacyconfidentiality.org/index.php/jpc/article/view/646. 1.3.1

[50] J. M. Perloff, *Microeconomics: theory and applications with calculus*. Pearson, 2016. 1.3.1

[51] H. Hotelling, "Stability in competition," in *The Collected Economics Articles of Harold Hotelling*, pp. 50–63, Springer, 1990. 1.3.1, 1.3.2, 4.1.2, 4.3

[52] Z. Yu, S. Li, and L. Tong, "Market dynamics and indirect network effects in electric vehicle diffusion," *Transportation Research Part D: Transport and Environment*, vol. 47, pp. 336–356, 2016. 1.3.1

[53] M. H. Lotfi, G. Kesidis, and S. Sarkar, "Network nonneutrality on the internet: Content provision under a subscription revenue model," *ACM SIGMETRICS Performance Evaluation Review*, vol. 42, no. 3, pp. 44–44, 2014. 1.3.1

[54] A. Acquisti, C. Taylor, and L. Wagman, "The economics of privacy," *Journal of Economic Literature*, vol. 54, pp. 442–492, jun 2016. 1.3.1

[55] N. Jentzsch, S. Preibusch, and A. Harasser, "Study on monetising privacy: An economic model for pricing personal information," *ENISA, Feb*, 2012. 1.3.1

[56] D.-J. Lee, J.-H. Ahn, and Y. Bang, "Managing consumer privacy concerns in personalization: a strategic analysis of privacy protection," *Management Information Systems Quarterly*, vol. 35, no. 2, pp. 423–444, 2011. 1.3.1

[57] S. Brenner, "Hotelling games with three, four, and more players," *Journal of Regional Science*, vol. 45, no. 4, pp. 851–864, 2005. 1.3.2

[58] The Economist, "The world's most valuable resource is no longer oil, but data," *The Economist*, 2017. 1.4

[59] National Science and Technology Council Networking and Information Technology Research and Development Program, "National privacy research strategy," tech. rep., Executive Office of the President of The United States, June 2016. 1.4

[60] P. Samarati and L. Sweeney, "Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and specialization," in *Technical Report SRI-CSL-98-04*, (SRI Intl.), 1998. 1.4

[61] L. Sweeney, "k-anonymity: A model for protecting privacy," *Intl. J. Uncertainty, Fuzziness, and Knowledge-based Systems*, vol. 10, no. 5, pp. 557–570, 2002. 1.4

[62] N. Li, T. Li, and S. Venkatasubramanian, "t-closeness: Privacy beyond k-anonymity and l-diversity," in *Data Engineering, 2007. ICDE 2007. IEEE 23rd International Conference on*, pp. 106–115, IEEE, 2007. 1.4

[63] A. Narayanan and V. Shmatikov, "Robust de-anonymization of large sparse datasets," in *Security and Privacy, 2008. SP 2008. IEEE Symposium on*, pp. 111–125, IEEE, 2008. 1.4

[64] A. Harmanci and M. Gerstein, "Quantification of private information leakage from phenotype-genotype data: linking attacks," *Nat Meth*, vol. 13, no. 3, pp. 251–256, 2016. Article. 1.4

[65] L. Sweeney, A. Abu, and J. Winn, "Identifying participants in the personal genome project by name (a re-identification experiment)," 2013. 1.4

[66] E. S. Finn, X. Shen, D. Scheinost, M. D. Rosenberg, J. Huang, M. M. Chun, X. Papademetris, and R. T. Constable, "Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity," *Nat Neurosci*, vol. 18, no. 11, pp. 1664–1671, 2015. Article. 1.4

[67] K. LeFevre, D. J. DeWitt, and R. Ramakrishnan, "Incognito: Efficient full-domain k-anonymity," in *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, pp. 49–60, ACM, 2005. 1.4

[68] R. J. Bayardo and R. Agrawal, "Data privacy through optimal k-anonymization," in *Data Engineering, 2005. ICDE 2005. Proceedings. 21st International Conference on*, pp. 217–228, IEEE, 2005. 1.4

[69] B. C. Fung, K. Wang, and S. Y. Philip, "Anonymizing classification data for privacy preservation," *IEEE transactions on knowledge and data engineering*, vol. 19, no. 5, 2007. 1.4

[70] V. S. Iyengar, "Transforming data to satisfy privacy constraints," in *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 279–288, ACM, 2002. 1.4

[71] P. Samarati, "Protecting respondents identities in microdata release," *IEEE transactions on Knowledge and Data Engineering*, vol. 13, no. 6, pp. 1010–1027, 2001. 1.4

[72] K. Wang, B. C. Fung, and S. Y. Philip, "Handicapping attacker's confidence: an alternative to k-anonymization," *Knowledge and Information Systems*, vol. 11, no. 3, pp. 345–368, 2007. 1.4

[73] B. Fung, K. Wang, R. Chen, and P. S. Yu, "Privacy-preserving data publishing: A survey of recent developments," *ACM Computing Surveys (CSUR)*, vol. 42, no. 4, p. 14, 2010. 1.4

[74] C. Dwork, "Differential privacy," in *Proc. 33rd Intl. Colloq. Automata, Lang., Prog.*, (Venice, Italy), July 2006. 1.4

[75] C. Dwork, "Differential privacy: A survey of results," in *Theory and Applications of Models of Computation: Lecture Notes in Computer Science*, New York:Springer, Apr. 2008. 1.4

[76] C. Dwork, A. Roth, *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014. 1.4

[77] S. E. Fienberg, A. Rinaldo, and X. Yang, *Differential Privacy and the Risk-Utility Tradeoff for Multi-dimensional Contingency Tables*, pp. 187–199. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010. 1.4

[78] Y. Wang, J. Lee, and D. Kifer, "Differentially private hypothesis testing, revisited," *arXiv preprint arXiv:1511.03376*, 2015. 1.4

[79] F. Yu, S. E. Fienberg, A. B. Slavković, and C. Uhler, "Scalable privacy-preserving data sharing methodology for genome-wide association studies," *Journal of biomedical informatics*, vol. 50, pp. 133–141, 2014. 1.4

[80] V. Karwa and A. Slavković, "Inference using noisy degrees: Differentially private β-model and synthetic graphs," *The Annals of Statistics*, vol. 44, no. 1, pp. 87–112, 2016. 1.4

[81] J. Duchi, M. Wainwright, and M. Jordan, "Minimax optimal procedures for locally private estimation," *arXiv preprint arXiv:1604.02390*, 2016. 1.4

[82] P. Kairouz, S. Oh, and P. Viswanath, "Extremal mechanisms for local differential privacy," *J. Mach. Learn. Res.*, vol. 17, pp. 492–542, Jan. 2016. 1.4

[83] M. Ye and A. Barg, "Optimal schemes for discrete distribution estimation under local differential privacy," in *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 759–763, June 2017. 1.4

[84] P. Kairouz, K. Bonawitz, and D. Ramage, "Discrete distribution estimation under local privacy," in *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, pp. 2436–2444, JMLR.org, 2016. 1.4

[85] D. Rebollo-Monedero, J. Forne, and J. Domingo-Ferrer, "From t-Closeness-Like Privacy to Postrandomization via Information Theory," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, pp. 1623–1636, Nov. 2010. 1.4

[86] F. P. Calmon and N. Fawaz, "Privacy against statistical inference," in *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*, pp. 1401–1408, 2012. 1.4, 5.1.2

[87] L. Sankar, S. R. Rajagopalan, and H. V. Poor, "Utility-privacy tradeoffs in databases: An information-theoretic approach," *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 6, pp. 838–852, 2013. 1.4, 1.4.2

[88] S. Salamatian, A. Zhang, F. P. Calmon, S. Bhamidipati, N. Fawaz, B. Kveton, P. Oliveira, and N. Taft, "Managing your private and public data: Bringing down inference attacks against your privacy," vol. 9, no. 7, pp. 1240–1255, 2015. 1.4

[89] Y. O. Basciftci, Y. Wang, and P. Ishwar, "On privacy-utility tradeoffs for constrained data release mechanisms," in *2016 Information Theory and Applications Workshop (ITA)*, pp. 1–6, Jan 2016. 1.4

[90] C. Huang, P. Kairouz, X. Chen, L. Sankar, and R. Rajagopal, "Context-aware generative adversarial privacy," *Entropy*, vol. 19, no. 12, p. 656, 2017. 1.4

[91] C. Huang, P. Kairouz, X. Chen, L. Sankar, and R. Rajagopal, "Generative adversarial privacy," *arXiv preprint arXiv:1807.05306*, 2018. 1.4

[92] I. Issa, S. Kamath, and A. B. Wagner, "An operational measure of information leakage," in *2016 Annual Conference on Information Science and Systems, CISS 2016, Princeton, NJ, USA, March 16-18, 2016*, pp. 234–239, 2016. 1.4, 1.4.2

[93] S. L. Warner, "Randomized response: A survey technique for eliminating evasive answer bias," *Journal of the American Statistical Association*, vol. 60, no. 309, pp. 63–69, 1965. 1.4

[94] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," in *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pp. 429–438, IEEE, 2013. 1.4

[95] I. Issa and A. B. Wagner, "Operational definitions for some common information leakage metrics," in *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 769–773, June 2017. 1.4, 1.4.2

[96] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Estimation efficiency under privacy constraints." arXiv:1707.02409, 2017. 1.4, 1.4.2

[97] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Privacy-aware guessing efficiency," in *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 754–758, June 2017. 1.4, 1.4.2

[98] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy, "Deep variational information bottleneck," *arXiv preprint arXiv:1612.00410*, 2016. 1.4, 1.4.2

[99] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014. 1.4.1, 1.4.2

[100] J. H. Schmidhuber, "Learning factorial codes by predictability minimization.," *Neural Computation*, 1992. 1.4.1, 1.4.2

[101] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, pp. 2672–2680, 2014. 1.4.1, 1.4.2, 5.1.3

[102] J. Morris, "On single-sample robust detection of known signals with additive unknown-mean amplitude-bounded random interference," *IEEE Transactions on Information Theory*, vol. 26, no. 2, pp. 199–209, 1980. 1.4.1

[103] S. Shamai and S. Verdu, "Worst-case power-constrained noise for binary-input channels," *IEEE Transactions on Information Theory*, vol. 38, no. 5, pp. 1494–1511, 1992. 1.4.1, 5.3.1

[104] J. Morris, "On single-sample robust detection of known signals with additive unknown-mean amplitude-bounded random interference–ii: The randomized decision rule solution (corresp.)," *IEEE Transactions on Information Theory*, vol. 27, no. 1, pp. 132–136, 1981. 1.4.1

[105] J. M. Morris and N. E. Dennis, "A random-threshold decision rule for known signals with additive amplitude-bounded nonstationary random interference," *IEEE Transactions on Communications*, vol. 38, no. 2, pp. 160–164, 1990. 1.4.1

[106] W. L. Root, "Communications through unspecified additive noise," *Information and Control*, vol. 4, no. 1, pp. 15–29, 1961. 1.4.1

[107] J. Whitehill and J. Movellan, "Discriminately decreasing discriminability with learned image filters," in *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pp. 2488–2495, IEEE, 2012. 5, 5.4

[108] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998. 5, 5.4

[109] WWDC 2016, "Engineering privacy for your user." `https://developer.apple.com/videos/play/wwdc2016/709/`, 2016. 1.4.2

[110] J. Tang, A. Korolova, X. Bai, X. Wang, and X. Wang, "Privacy loss in Apple's implementation of differential privacy on MacOS 10.12." arXiv:1709.02753, Sep. 2017. 1.4.2

[111] F. P. Calmon, M. Varia, M. Médard, M. M. Christiansen, K. R. Duffy, and S. Tessaro, "Bounds on inference," in *Proc. 51st Annual Allerton Conf. on Commun., Control, and Comput.*, pp. 567–574, IEEE, 2013. 1.4.2

[112] S. Verdú, "α-mutual information," in *2015 Information Theory and Applications Workshop (ITA)*, 2015. 1.4.2

[113] Y. Zhang, M. Ozay, Z. Sun, and T. Okatani, "Information potential autoencoders." arXiv:1706.04635, 2017. 1.4.2

[114] L. Theis, W. Shi, A. Cunningham, and F. Huszár, "Lossy image compression with compressive autoencoders," in *International Conference on Learning Representations*, 2017. 1.4.2

[115] M. Sugiyama and K. M. Borgwardt, "Measuring statistical dependence via the mutual information dimension," *dim*, vol. 10, p. 1, 2013. 1.4.2

[116] K. Xu, T. Cao, S. Shah, C. Maung, and H. Schweitzer, "Cleaning the null space: A privacy mechanism for predictors," in *Proc. AAAI Conference on Artificial Intelligence*, 2017. 1.4.2

[117] J. Hamm, "Minimax filter: Learning to preserve privacy from inference attacks," *arXiv preprint arXiv:1610.03577*, 2016. 1.4.2

[118] C. Liu, S. Chakraborty, and P. Mittal, "Deeprotect: Enabling inference-based access control on mobile sensing applications." arXiv:1702.06159, 2017. 1.4.2

[119] M. Abadi and D. G. Andersen, "Learning to protect communications with adversarial neural cryptography," *arXiv preprint arXiv:1610.06918*, 2016. 1.4.2

[120] H. Edwards and A. Storkey, "Censoring representations with an adversary," *arXiv preprint arXiv:1511.05897*, 2015. 1.4.2

[121] N. Raval, A. Machanavajjhala, and L. P. Cox, "Protecting visual secrets using adversarial nets," in *CVPR Workshop Proceedings*, 2017. 1.4.2

[122] P. Smolensky, "Information processing in dynamical systems: Foundations of harmony theory," tech. rep., Colorado University at Boulder Department of Computer Science, 1986. 1.4.2

[123] G. E. Hinton, "Deep belief networks," *Scholarpedia*, vol. 4, no. 5, p. 5947, 2009. 1.4.2

[124] I. Goodfellow, "Nips 2016 tutorial: Generative adversarial networks," *arXiv preprint arXiv:1701.00160*, 2016. 1.4.2

[125] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of gans for improved quality, stability, and variation," *arXiv preprint arXiv:1710.10196*, 2017. 1.4.2

[126] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein gan," *arXiv preprint arXiv:1701.07875*, 2017. 1.4.2, 5.4.2

[127] D. P. Bertsekas, *Dynamic programming and optimal control*, vol. 1, 2. Athena Scientific Belmont, MA, 1995. 2.1.1, 2.1.1

[128] B. Bonet and J. Pearl, "Qualitative mdps and pomdps: An order-of-magnitude approximation," in *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*, pp. 61–68, Morgan Kaufmann Publishers Inc., 2002. 2.1.1

[129] S. M. Ross, *Applied probability models with optimization applications*. Courier Dover Publications, 2013. 2.1.1

[130] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian data analysis*, vol. 2. Taylor & Francis, 2014. 2.1.4

[131] W. S. Lovejoy, "Some monotonicity results for partially observed markov decision processes," *Operations Research*, vol. 35, no. 5, pp. 736–743, 1987. 2.2.1

[132] J. Nash, "Non-cooperative games," *Annals of mathematics*, pp. 286–295, 1951. 3.2.1

[133] D. Fudenberg and J. Tirole, *Game theory*. MIT press, MA, 1991. 3.2.1, 4.2

[134] M. J. Osborne and A. Rubinstein, *A course in game theory*. MIT press, 1994. 4.2

[135] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games," *Econometrica: Journal of the Econometric Society*, pp. 520–534, 1965. 4.4, F

[136] T. Nguyen and S. Sanner, "Algorithms for direct 0–1 loss optimization in binary classification," in *International Conference on Machine Learning*, pp. 1085–1093, 2013. 5.1.2

[137] J. Liao, O. Kosut, L. Sankar, and F. du Pin Calmon, "A general framework for information leakage: Privacy utility trade-offs." Sept. 2017. 5.1.2, 5.1.2

[138] G. P. Zhang, "Neural networks for classification: a survey," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 30, no. 4, pp. 451–462, 2000. 5.1.3

[139] Y. Tang, "Deep learning using linear support vector machines," *arXiv preprint arXiv:1306.0239*, 2013. 5.1.3

[140] W. E. Lillo, M. H. Loh, S. Hui, and S. H. Zak, "On solving constrained optimization problems with neural networks: A penalty method approach," *IEEE Transactions on neural networks*, vol. 4, no. 6, pp. 931–940, 1993. 5.1.3, 5.1.3

[141] J. Eckstein and W. Yao, "Augmented lagrangian and alternating direction methods for convex optimization: A tutorial and some illustrative computational results," *RUTCOR Research Reports*, vol. 32, 2012. 5.1.3, 5.1.3

[142] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*, 2016. 5.2.3, 5.3.2

[143] E. W. Weisstein, "Normal distribution," 2002. 5.3

[144] R. G. Gallager, *Stochastic processes: theory for applications*. Cambridge University Press, 2013. 5.3.1, 5.3.2

[145] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167*, 2015. 5.4.1

[146] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, pp. 1097–1105, 2012. 5.4.1

[147] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014. 5.4.1

[148] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016. 5.4.1

[149] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, *et al.*, "Going deeper with convolutions," Cvpr, 2015. 5.4.1

APPENDIX A

PROOF OF THEOREM 1

Proof. Let p_F be the stationary distribution of the Markov transition. Then $p_F = \lambda_{A,A}p_F + (1-p_F)\lambda_{N,A}$, which implies $p_F = \frac{\lambda_{N,A}}{1-\lambda_{A,A}+\lambda_{N,A}}$. If $V^t_{\beta,\mathsf{LP}}(p_t) > V^t_{\beta,\mathsf{HP}}(p_t)$ Remember that the threshold is the solution to $V^t_{\beta,\mathsf{LP}}(p_t) = V^t_{\beta,\mathsf{HP}}(p_t)$. Let τ be the threshold value, we have:

$$\begin{aligned}&\beta^t C_L + V^{t+1}_\beta(T(\tau))\\&= (1-\tau)[\beta^t C_{HN} + V^{t+1}_\beta(\lambda_{N,A})] + \tau[\beta^t C_{HA} + V^{t+1}_\beta(\lambda_{A,A})].\end{aligned} \tag{A.1}$$

By the definition of $V^t_\beta(p_t)$, we know that $V^t_\beta(p_t) = \beta^t V_\beta(p_t)$. Thus $V^t_\beta(\lambda_{N,A}) = \beta^t V_\beta(\lambda_{N,A})$ and $V^t_\beta(\lambda_{A,A}) = \beta^t V_\beta(\lambda_{A,A})$.

If $T(\tau) \geq \tau$, which is equivalent to $p_F \geq \tau$, then $V^{t+1}_\beta(T(\tau)) = V^{t+1}_{\beta,\mathsf{LP}}(T(\tau))$. Therefore, $V^t_{\beta,\mathsf{LP}}(\tau) = \lim_{n\to\infty}\{\beta^t\frac{1-\beta^n}{1-\beta}C_L + \beta^n V^{t+1}_\beta(T^n(\tau))\}$ where $T^n(\tau) = T(T^{n-1}(\tau)) = p_F(1-(\lambda_{A,A}-\lambda_{N,A})^n) + (\lambda_{A,A}-\lambda_{N,A})^n\tau$. Taking $n\to\infty$, we have $V^t_{\beta,\mathsf{LP}}(\tau) = \beta^t\frac{C}{1-\beta}$. Substitute this into (A.1) yields:

$$\frac{C_L}{1-\beta} = (1-\tau)C_{HN} + \tau C_{HA} + \beta(\tau V_\beta(\lambda_{A,A}) + (1-\tau)V_\beta(\lambda_{N,A})). \tag{A.2}$$

By rearranging terms in the above expression, we have

$$\tau = \frac{\frac{C_L}{1-\beta} - C_{HN} - \beta V_\beta(\lambda_{N,A})}{(C_{HA}-C_{HN}) + \beta(V_\beta(\lambda_{A,A}) - V_\beta(\lambda_{N,A}))}. \tag{A.3}$$

If $p_F \leq \tau$, then $T(\tau) \leq \tau$. Therefore $V^{t+1}_\beta(T(\tau)) = V^{t+1}_{\beta,\mathsf{HP}}(T(\tau))$, which implies

$$V^t_{\beta,\mathsf{LP}}(\tau) = \beta^t C_L + V^{t+1}_\beta(T(\tau)) = \beta^t C_L + V^{t+1}_{\beta,\mathsf{HP}}(T(\tau)) = V^t_{\beta,\mathsf{HP}}(\tau). \tag{A.4}$$

In this case,

$$C_L + \beta V_{\beta,\mathsf{HP}}(T(\tau)) = V_{\beta,\mathsf{HP}}(\tau). \tag{A.5}$$

Substitute (2.1) and (2.9) into (A.5), we have

$$\begin{aligned}\tau = &\frac{C_L - (1-\beta(1-\lambda_{N,A}))(C_{HN} + \beta V_\beta(\lambda_{N,A}))}{(1-(\lambda_{A,A}-\lambda_{N,A})\beta)(C_{HA}-C_{HN}+\beta(V_\beta(\lambda_{A,A}) - V(\lambda_{N,A})))}\\&+\frac{\beta\lambda_{N,A}(C_{HA}+\beta V_\beta(\lambda_{A,A}))}{(1-(\lambda_{A,A}-\lambda_{N,A})\beta)(C_{HA}-C_{HN}+\beta(V_\beta(\lambda_{A,A}) - V(\lambda_{N,A})))}.\end{aligned} \tag{A.6}$$

Next, we present how to compute $V_\beta(\lambda_{N,A})$ and $V_\beta(\lambda_{A,A})$.

Case 1: If $\lambda_{N,A} \geq \tau$, then by Model Assumption 2, $\lambda_{A,A} \geq \lambda_{N,A} \geq \tau$ and $p_F \geq \lambda_{N,A} \geq \tau$. Thus, both $\lambda_{A,A}$ and $\lambda_{N,A}$ are in Φ_{LP}, therefore,

$$V_\beta(\lambda_{N,A}) = V_\beta(\lambda_{A,A}) = \frac{C_L}{1-\beta}. \tag{A.7}$$

Case 2: If $\lambda_{N,A} \leq \tau$, we have $V_\beta(\lambda_{N,A}) = V_{\beta,\mathsf{HP}}(\lambda_{N,A})$. Therefore,

$$V_\beta(\lambda_{N,A}) = (1-\lambda_{N,A})[C_{HN} + V_\beta^1(\lambda_{N,A})] + \lambda_{N,A}[C_{HA} + V_\beta^1(\lambda_{A,A})]. \tag{A.8}$$

$$V_\beta(\lambda_{A,A}) = \min_{u_t \in \{\mathsf{HP},\mathsf{LP}\}} V_{\beta,u_t}(\lambda_{A,A}) \tag{A.9}$$

$$= \min\{C_L + V_\beta^1(T(\lambda_{A,A})), V_{\mathsf{HP}}(\lambda_{A,A})\} \tag{A.10}$$

$$= \min\{C_L \frac{1-\beta^N}{1-\beta}, \min_{0 \leq n \leq N-1}\{C_L \frac{1-\beta^n}{1-\beta} + V_{\beta,\mathsf{HP}}^n(T^n(\lambda_{A,A}))\}\}. \tag{A.11}$$

Since $N \to \infty$ and $0 \leq \beta \leq 1$,

$$V_\beta(\lambda_{A,A}) = \min_{n>0}\{C_L \frac{1-\beta^n}{1-\beta} + \beta^n V_{\beta,\mathsf{HP}}(T^n(\lambda_{A,A}))\}. \tag{A.12}$$

we have:

$$V_\beta(\lambda_{A,A}) = \min_{n \geq 0}\{\frac{C_L \frac{1-\beta^n}{1-\beta} + \beta^n[\bar{T}^n(\lambda_{A,A})(C_{HN} + C(\lambda_{N,A})) + T^n(\lambda_{A,A})C_{HA}]}{1 - \beta^{n+1}[\bar{T}^n(\lambda_{A,A})\frac{\lambda_{N,A}\beta}{1-(1-\lambda_{N,A})\beta} + T^n(\lambda_{A,A})]}\}. \tag{A.13}$$

where

$$T^n(\lambda_{A,A}) = T(T^{n-1}(\lambda_{A,A})) = \frac{(\lambda_{A,A} - \lambda_{N,A})^{n+1}(1-\lambda_{A,A}) + \lambda_{N,A}}{1 - (\lambda_{A,A} - \lambda_{N,A})}, \tag{A.14}$$

$$\bar{T}^n(\lambda_{A,A}) = 1 - T^n(\lambda_{A,A}) \tag{A.15}$$

$$C(\lambda_{N,A}) = \beta \frac{(1-\lambda_{N,A})C_{HN} + \lambda_{N,A}C_{HA}}{1-(1-\lambda_{N,A})\beta}. \tag{A.16}$$

Next, we prove the uniqueness of τ. Noticing that $V^t_{\beta,\mathsf{LP}}(p)$ is a concave and non-decreasing function of p and $V^t_{\beta,\mathsf{HP}}(p)$ is an affine and non-decreasing function of p (see Lemma 2). Thus, both $V^t_{\beta,\mathsf{LP}}(p)$ and $V^t_{\beta,\mathsf{HP}}(p)$ are continuous functions (every concave/affine function is continuous). Furthermore, if $p = 0$, the optimal action will be offering HP since the retailer is sure that the state of consumer is Normal and $C_{HN} < C_L$. This implies

$$V^t_{\beta,\mathsf{LP}}(p = 0) > V^t_{\beta,\mathsf{HP}}(p = 0). \tag{A.17}$$

Likewise, when $p = 1$, the optimal action will be offering LP since the retailer is sure that the state of consumer is $\mathsf{Alerted}$ and $C_L < C_{HA}$. Thus, we have

$$V^t_{\beta,\mathsf{LP}}(p = 1) < V^t_{\beta,\mathsf{HP}}(p = 1). \tag{A.18}$$

Thus, no action is uniformly better than the other in this model. Therefore, by (A.17),(A.18) and continuity and concavity of $V^t_{\beta,\mathsf{LP}}(p)$ and $V^t_{\beta,\mathsf{HP}}(p)$, there is a unique solution to $V^t_{\beta,\mathsf{LP}}(p) = V^t_{\beta,\mathsf{HP}}(p)$ for $p \in [0, 1]$. □

APPENDIX B

PROOF OF COROLLARY 1

Proof. By setting $V_{\mathsf{LP}}(p_t) \leq V_{\mathsf{HP}}(p_t)$, we have

$$\beta^t C_L + \beta V_\beta^t(T(p_t)) \leq (1-p_t)[\beta^t C_{HN} + \beta V_\beta^t(\lambda_{N,A})] + p_t[\beta^t C_{HA} + \beta V_\beta^t(\lambda_{A,A})]. \quad \text{(B.1)}$$

By Lemma 2, $V_\beta^t(p_t)$ is a concave function. Thus,

$$\begin{aligned} V_\beta^t(T(p_t)) &= V_\beta^t(\lambda_{N,A}(1-p_t) + \lambda_{A,A}p_t) \\ &\geq (1-p_t)V_\beta^t(\lambda_{N,A}) + p_t V_\beta^t(\lambda_{A,A}). \end{aligned} \quad \text{(B.2)}$$

By substituting (B.2) into (B.1), we can simplify inequality (B.1) to $(1-p_t)C_{HN} + p_t C_{HA} \geq C_L$, which implies $p_t \geq \frac{C_L - C_{HN}}{C_{HA} - C_{HN}} = \kappa$ when $V_{\mathsf{LP}}^t(p_t) \leq V_{\mathsf{HP}}^t(p_t)$. Thus, $p_t < \kappa$ implies $V_{\mathsf{LP}}(p_t) > V_{\mathsf{HP}}(p_t)$. □

APPENDIX C

PROOF OF COROLLARY 2

Proof. Assume that $\lambda_{N,A} \geq \tau$, we have $\lambda_{A,A} > p_F = \frac{\lambda_{N,A}}{1-(\lambda_{A,A}-\lambda_{N,A})} > \lambda_{N,A} \geq \tau$. In this case, By (A.3) and (A.7), we have

$$\tau = \frac{C_L - C_{HN}}{C_{HA} - C_{HN}} = \kappa. \tag{C.1}$$

Thus, $\tau = \kappa$ if $\lambda_{N,A} > \kappa$. Assume that $\lambda_{N,A} < \tau$, then there are two cases for p_F:

Case 1: $p_F > \tau$, then $\lambda_{A,A} > p_F > \tau$, which implies

$$V_\beta(\lambda_{A,A}) = V_{\beta,\mathsf{LP}}(\lambda_{A,A}) = \frac{C_L}{1-\beta}. \tag{C.2}$$

By (A.3), (A.8), and (C.2), we have

$$\tau = \frac{\beta(C_L - C_{HA})\lambda_{N,A} + C_L - C_{HN}}{(1-\beta)C_{HA} - C_{HN} + \beta C_L}. \tag{C.3}$$

Therefore, $\tau = \frac{\beta(C_L-C_{HA})\lambda_{N,A}+C_L-C_{HN}}{(1-\beta)C_{HA}-C_{HN}+\beta C_L}$ if $p_F = \frac{\lambda_{N,A}}{1-(\lambda_{A,A}-\lambda_{N,A})} \geq \tau = \frac{\beta(C_L-C_{HA})\lambda_{N,A}+C_L-C_{HN}}{(1-\beta)C_{HA}-C_{HN}+\beta C_L}$ and $\lambda_{N,A} < \frac{\beta(C_L-C_{HA})\lambda_{N,A}+C_L-C_{HN}}{(1-\beta)C_{HA}-C_{HN}+\beta C_L}$.

Case 2: $p_F < \tau$, τ can be computed by (A.6), (A.8), and (A.13). Moreover, for fixed $\lambda_{A,A}$, (A.6) is a non-decreasing function w.r.t. $\lambda_{N,A}$. Thus, let $\tau^+ = \frac{\lambda_{N,A}}{1-(\lambda_{A,A}-\lambda_{N,A})} = \frac{\beta(C_L-C_{HA})\lambda_{N,A}+C_L-C_{HN}}{(1-\beta)C_{HA}-C_{HN}+\beta C_L}$, $\tau \leq \tau^+$ in Case 2. Therefore, τ^+ is an upperbound for the optimal action in Case 2.

Since (A.6) is non-decreasing, (C.3) is decreasing and intersects with (C.1) at $\lambda_{N,A} = \frac{C_L-C_{HN}}{C_{HA}-C_{HN}}$, we have proved Corollary 2. □

APPENDIX D

PROOF OF THEOREM 2

Proof. Assume that τ is the threshold of offering either HP or LP coupons, then we have $V^t_{\beta,\mathsf{LP}}(\tau) = V^t_{\beta,\mathsf{HP}}(\tau)$. Noticing that the state of the consumer is revealed to the retailer through cost when an HP coupon is offered, we have

$$\begin{aligned}
&V^t_{\beta,\mathsf{LP}}(\tau) - V^t_{\beta,\mathsf{HP}}(\tau) \\
&= \beta^t(C_L - (1-\tau)C_{HN} - \tau C_{HA}) + [V^{t+1}_\beta(T(\tau)) - (1-\tau)V^{t+1}_\beta(\lambda'_{N,A}) - \tau V^{t+1}_\beta(\lambda'_{A,A})] \\
&= 0.
\end{aligned} \tag{D.1}$$

The above equation is similar to (A.1) with $V^{t+1}_\beta(\lambda_{N,A})$ and $V^{t+1}_\beta(\lambda_{A,A})$ replaced by $V^{t+1}_\beta(\lambda'_{N,A})$ and $V^{t+1}_\beta(\lambda'_{A,A})$, respectively. Thus, Lemmas 1-3 still hold. Therefore, the proof follows the same argument for proving Theorem 1; we omit it for brevity. $\square$

APPENDIX E

PROOF OF THEOREM 4

Proof. We prove by contradiction. Suppose that both SPs offer the same privacy risk $\tilde{\varepsilon}$, we prove that there is no unilateral profitable deviation in the subgames using backward induction. Without loss of generality, we assume $p_1 \leq p_2$. We now prove when both SPs choose the same $\tilde{\varepsilon}$, one of the SPs will be better off by unilaterally deviate from offering $\tilde{\varepsilon}$. We start at the third stage wherein each consumer chooses the SP which maximizes its utility (4.5). Since $\varepsilon_1^* = \varepsilon_2^* = \tilde{\varepsilon}$, every consumer will choose the SP that offers the highest QoS. At the second stage, given the privacy risk strategy $\varepsilon_1^* = \varepsilon_2^* = \tilde{\varepsilon}$ and the equilibrium strategy in the third stage, each SP determines its QoS offering by solving (4.12). Finally, we show that SP_2 will be better off if it deviates from $\tilde{\varepsilon}$ unilaterally. By Assumption 5, each SP has equal share of the market if $v_1 = v_2$ and $\varepsilon_1 = \varepsilon_2$. The profit of SP_i can be written as

$$\begin{aligned}\pi_i &= [R(\varepsilon_i) - C(v_i;\varepsilon_i)]n_i(v_i;\varepsilon_i;v_{-i};\varepsilon_{-i}) \qquad \text{(E.1)}\\ &= \begin{cases} r\varepsilon_i + p_i - c(v_i + \lambda\varepsilon_i) & \text{if } v_i > v_{-i} \\ \frac{r\varepsilon_i + p_i - c(v_i+\lambda\varepsilon_i)}{2} & \text{if } v_i = v_{-i} \\ 0 & \text{if } v_i < v_{-i} \end{cases}.\end{aligned}$$

As argued in section 4.1.1, we assume that the net profit from using consumers' private data is non-negative $R_P(\varepsilon_i) - C_P(\varepsilon_i) > 0$. Thus, $(r - c\lambda)\varepsilon_i > 0 \ \forall \varepsilon_i \in [0, \bar{\varepsilon}]$, which indicates $r - c\lambda > 0$. Since every consumer will choose the SP that offers the highest QoS, each SP's best response strategy with respect to its competitor is to increase v_i until one of the SPs realizes it is not profitable to increase QoS anymore. Therefore, by (E.1), both SPs will increase v_i until $R(\varepsilon_i) - C(v_i, \varepsilon_i) = 0$ for one of the SPs. Since we assume $p_1 \leq p_2$, we prove the theorem for the following two cases:

Case 1: $p_1 = p_2 = p$, i.e., both SPs have the same privacy-independent revenue. In this case, given $\varepsilon_1^* = \varepsilon_2^* = \tilde{\varepsilon}$, each SP will increase its QoS to beat its competitor until $R(\varepsilon_i) - C(v_i, \varepsilon_i) = r\tilde{\varepsilon} + p - c(v_i + \lambda\tilde{\varepsilon}) = 0$. As a result, both SPs' equilibrium

strategies at this stage are given by

$$v_1^* = v_2^* = \frac{(r - c\lambda)\tilde{\varepsilon} + p}{c}. \tag{E.2}$$

At the first stage, the SPs determine their privacy risks based on the equilibrium strategies in the second and the third stages. Given the equilibrium strategies in the second stage (E.2), both SPs have zero profit. Since we assume $\varepsilon_1 \leq \varepsilon_2$, SP_1 can only reduce its privacy risk from $\tilde{\varepsilon}$ and SP_2 can only increase from it. We now prove that it is a non-profitable deviation for SP_1 to decrease its privacy risk to $\tilde{\varepsilon}_1$ unilaterally. Since SP_1's QoS strategy is given by $v_1^* = \frac{(r-c\lambda)\tilde{\varepsilon}+p}{c}$, its profit is given by

$$R(\tilde{\varepsilon}_1) - C(v_1^*, \tilde{\varepsilon}_1) = r\tilde{\varepsilon}_1 + p - c(v_1^* + \lambda\tilde{\varepsilon}_1) = (r - c\lambda)(\tilde{\varepsilon}_1 - \tilde{\varepsilon}) < 0.$$

Thus, SP_1 does not have incentives to deviate from playing $\tilde{\varepsilon}$ unilaterally. On the other hand, if SP_2 increases its privacy risk from $\tilde{\varepsilon}$ to $\tilde{\varepsilon}_2$ unilaterally, its profit is given by

$$R(\tilde{\varepsilon}_2) - C(v_2^*, \tilde{\varepsilon}_2) = r\tilde{\varepsilon}_2 + p - c(v_2^* + \lambda\tilde{\varepsilon}_2) = (r - c\lambda)(\tilde{\varepsilon}_2 - \tilde{\varepsilon}) > 0.$$

Therefore, SP_2 is better off by changing its privacy risk from $\tilde{\varepsilon}$ to $\tilde{\varepsilon}_2$ unilaterally. Thus, there is no SPNE such that both SPs offer the same privacy risk when $p_1 = p_2$.

Case 2: $p_1 < p_2$, i.e., SP_2 has a higher privacy-independent revenue than SP_1. In this case, since $p_1 < p_2$, both SPs will keep increasing its QoS until SP_1 has zero profit. Thus, by solving $R(\tilde{\varepsilon}) - C(v_1, \tilde{\varepsilon}) = 0$, SP_1 will play $v_1^* = \frac{(r-c\lambda)\tilde{\varepsilon}+p_1}{c}$ at the equilibrium. On the other hand, SP_2 will offer an QoS slightly higher than v_1^* and captures the entire market. At the first stage, given the equilibrium strategy of the second stage described above, both SPs choose their privacy risk offerings. We now prove that it is a non-profitable deviation for SP_1 to decrease its privacy risk to $\tilde{\varepsilon}_1$ unilaterally. Since SP_1 offers $v_1^* = \frac{(r-c\lambda)\tilde{\varepsilon}+p_1}{c}$ at the second stage, its profit is

$$R(\tilde{\varepsilon}_1) - C(v_1^*, \tilde{\varepsilon}_1) = r\tilde{\varepsilon}_1 + p_1 - c(v_1^* + \lambda\tilde{\varepsilon}_1)$$

$$
\begin{aligned}
&= r\tilde{\varepsilon}_1 + p_1 - c(\frac{(r - c\lambda)\tilde{\varepsilon} + p_1}{c} + \lambda\tilde{\varepsilon}_1) \\
&= (r - c\lambda)(\tilde{\varepsilon}_1 - \tilde{\varepsilon}) \\
&< 0.
\end{aligned}
$$

Thus, SP_1 does not have incentives to deviate from playing $\tilde{\varepsilon}$ unilaterally. On the other hand, if SP_2 increases its privacy risk from $\tilde{\varepsilon}$ to $\tilde{\varepsilon}_2$ unilaterally, its profit is given by

$$
\begin{aligned}
R(\tilde{\varepsilon}_2) - C(v_2^*, \tilde{\varepsilon}_2) &= r\tilde{\varepsilon}_2 + p_2 - c(v_2^* + \lambda\tilde{\varepsilon}_2) \\
&= r\tilde{\varepsilon}_2 + p_2 - c(\frac{(r - c\lambda)\tilde{\varepsilon} + p_1}{c} + \lambda\tilde{\varepsilon}_2) \\
&= (r - c\lambda)(\tilde{\varepsilon}_2 - \tilde{\varepsilon}) + p_2 - p_1 \\
&= (r - c\lambda)(\tilde{\varepsilon}_2 - \tilde{\varepsilon}) + R(\tilde{\varepsilon}) - C(v_2^*, \tilde{\varepsilon}) \\
&> R(\tilde{\varepsilon}) - C(v_2^*, \tilde{\varepsilon}).
\end{aligned}
$$

Thus, SP_2 has incentives to deviate from offering the same privacy risk. Therefore, playing $\varepsilon_1^* = \varepsilon_2^* = \tilde{\varepsilon}$ is not an SPNE when $p_1 < p_2$. $\square$

APPENDIX F

PROOF OF THEOREM 5

Proof. Starting form the last stage in which consumers choose different SPs, we use backward induction to find the SPNE of the sequential game. In the last stage, each consumer located at $x \in [0, 1]$ chooses an SP which maximize its utility function (4.3). By (4.7) and the assumption that consumers' privacy risk tolerances are uniformly distributed, the indifference threshold x_τ is given by

$$x_\tau = \frac{v_1 - v_2 + \frac{t(\varepsilon_2^2 - \varepsilon_1^2)}{\bar{\varepsilon}}}{t(\varepsilon_2 - \varepsilon_1)} = n_1(v_1; \varepsilon_1; v_2; \varepsilon_2). \tag{F.1}$$

At the second stage, the optimal strategy of each SP is determined by the solution of (4.12). For fixed privacy risk guarantees ε_2 and ε_1, the objective function of $SP_i, i \in \{1, 2\}$ in this stage, i.e. $\pi_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i})$, is a concave function with respect to its own strategy v_i. Furthermore, the feasible set of SP_i's strategy is a convex set ($v_i \in [0, +\infty]$). Thus, the non-cooperative subgame between SP_2 and SP_1 in this stage can be considered as a two-player concave game. By Theorem 1 and 2 in [135], we can establish

Lemma 4. *For fixed privacy risk strategies, there exists a unique Nash equilibrium in the game between SP_2 and SP_1 at the second stage.*

To compute the equilibrium strategy of the second stage, we first substitute (4.14), (4.15), and (F.1) into (4.9) and (4.8). Then, we apply the first order condition to SPs' profit functions and solve the simultaneous equations given by

$$\frac{\partial \pi_i(v_i; \varepsilon_i; v_{-i}; \varepsilon_{-i})}{\partial v_i} = 0 \quad \forall i \in \{1, 2\}. \tag{F.2}$$

Solving the above simultaneous equations yields

$$v_1 = \frac{r\varepsilon_1 + p_1}{2c} + \frac{v_2 - \lambda\varepsilon_1 - tx_2\varepsilon_2 + tx_1\varepsilon_1}{2}, \tag{F.3}$$

$$v_2 = \frac{r\varepsilon_2 + p_2}{2c} + \frac{v_1 - \lambda\varepsilon_2 - t(1 - x_2)\varepsilon_2 + t(1 - x_1)\varepsilon_1}{2}. \tag{F.4}$$

For given privacy guarantees ε_1, and ε_2, solving the simultaneous linear equations above by substituting (F.3) into (F.4) yields the equilibrium strategies

$$v_1^*(\varepsilon_2, \varepsilon_1) = \frac{2(r\varepsilon_1 + p_1) + r\varepsilon_2 + p_2}{3c} + \frac{t(1+x_1)\varepsilon_1 - \lambda(\varepsilon_2 + 2\varepsilon_1) - t(1+x_2)\varepsilon_2}{3}, \quad \text{(F.5)}$$

$$v_2^*(\varepsilon_2, \varepsilon_1) = \frac{2(r\varepsilon_2 + p_2) + r\varepsilon_1 + p_1}{3c} + \frac{t(2-x_1)\varepsilon_1 - \lambda(2\varepsilon_2 + \varepsilon_1) - t(2-x_2)\varepsilon_2}{3}. \quad \text{(F.6)}$$

At the first stage, the SPs determine their optimal privacy risk by considering the QoS of each SP and the market segmentation computed in previous stages as functions of privacy risks offered by the SPs. By substituting (F.6) and (F.5) into (4.9) and (4.8), the profit functions of the SPs can be written as

$$\pi_2 = \frac{c}{9t(\varepsilon_2 - \varepsilon_1)}[\frac{p_2 - p_1}{c} + (\frac{r}{c} - \lambda + t\frac{2\bar{\varepsilon} - \varepsilon_2 - \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1)]^2, \quad \text{(F.7)}$$

$$\pi_1 = \frac{c}{9t(\varepsilon_2 - \varepsilon_1)}[-\frac{p_2 - p_1}{c} + (-\frac{r}{c} + \lambda + t\frac{\bar{\varepsilon} + \varepsilon_2 + \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1)]^2. \quad \text{(F.8)}$$

Next, we apply the first order condition to SPs' profit functions to compute the equilibrium strategies. Taking the derivatives of π_2 and π_1 with respect to ε_2 and ε_1 and set both of their values to 0 yields

$$\frac{\partial \pi_2}{\partial \varepsilon_2} = \frac{c[(\frac{r}{c} - \lambda + t\frac{2\bar{\varepsilon} - \varepsilon_2 - \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) + \frac{p_2 - p_1}{c}][(\frac{r}{c} - \lambda + t\frac{2\bar{\varepsilon} - 3\varepsilon_2 + \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) - \frac{p_2 - p_1}{c}]}{9t(\varepsilon_2 - \varepsilon_1)^2} = 0, \quad \text{(F.9)}$$

$$\frac{\partial \pi_1}{\partial \varepsilon_1} = \frac{c[(-\frac{r}{c} + \lambda + t\frac{\bar{\varepsilon} + \varepsilon_2 + \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) - \frac{p_2 - p_1}{c}][(\frac{r}{c} - \lambda - t\frac{\bar{\varepsilon} - \varepsilon_2 + 3\varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) - \frac{p_2 - p_1}{c}]}{9t(\varepsilon_2 - \varepsilon_1)^2} = 0. \quad \text{(F.10)}$$

Solving the two simultaneous equations above yields

$$(\frac{r}{c} - \lambda + t\frac{2\bar{\varepsilon} - \varepsilon_2 - \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) + \frac{p_2 - p_1}{c} = 0 \quad \text{(F.11)}$$

or

$$(\frac{r}{c} - \lambda + t\frac{2\bar{\varepsilon} - 3\varepsilon_2 + \varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2 - \varepsilon_1) - \frac{p_2 - p_1}{c} = 0 \quad \text{(F.12)}$$

and

$$(-\frac{r}{c}+\lambda+t\frac{\bar{\varepsilon}+\varepsilon_2+\varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2-\varepsilon_1)-\frac{p_2-p_1}{c}=0 \tag{F.13}$$

or

$$(\frac{r}{c}-\lambda-t\frac{\bar{\varepsilon}-\varepsilon_2+3\varepsilon_1}{\bar{\varepsilon}})(\varepsilon_2-\varepsilon_1)-\frac{p_2-p_1}{c}=0. \tag{F.14}$$

We note that the strategies given by (F.11) and (F.13) result in 0 profits in (F.7) and (F.8). This indicates the privacy risk determined by (F.11) and (F.13) are strictly dominated by the strategies given by the solution of (F.12) and (F.14). Solving (F.12) and (F.14) yields the equilibrium privacy risk (4.17) and

$$\varepsilon_1^*=\frac{12\bar{\varepsilon}c\alpha-3ct\bar{\varepsilon}-16(p_2-p_1)}{24tc}. \tag{F.15}$$

By subtracting (F.15) from (4.17), we have (4.19). Substitute the solution of ε_2^* and ε_1^* to (F.6) and (F.5), we have (4.18) and

$$v_1^*=\frac{(2\alpha-t)c\alpha 6\bar{\varepsilon}+(\alpha-3t)3ct\bar{\varepsilon}+(t-\alpha)16p_2+(2\alpha+t)8p_1}{24ct}. \tag{F.16}$$

Subtracting (F.16) from (4.18) yields (4.20).

Next, we prove the sufficient condition for the existence of the above SPNE. First of all, the model parameters must sustain a competitive market environment. Thus, in the equilibrium, each SP must have non-zero market share. This indicates the parameters must satisfy $0\leq x_\tau^*=\frac{v_1^*-v_2^*+t(x_2^*\varepsilon_2^*-x_1^*\varepsilon_1^*)}{t(\varepsilon_2^*-\varepsilon_1^*)}\leq 1$. Substitute (4.17), (4.18), (4.19), and (4.20) into the above inequality, we have (4.21). Furthermore, in the SPNE, the QoS of each SP must be non-negative (QoS feasibility) and the privacy risk guarantees must be bounded between 0 and $\bar{\varepsilon}$ (privacy risk feasibility). By the model assumption in Section 4.1.1, we have $\varepsilon_1\leq\varepsilon_2$. Thus, we only requires $\varepsilon_2\leq\bar{\varepsilon}$ and $\varepsilon_1\geq 0$. Substitute (4.17) and (4.19) into the two inequalities above yields (4.22).

Let $x_i^* = \frac{\varepsilon_i^*}{\bar{\varepsilon}}, i \in \{A, B\}$ denote the normalized privacy risk of each SP in the SPNE. The equilibrium strategies must satisfy the complete market coverage condition given by $u_i(x) = v_i^* - t(x - x_i^*)\varepsilon_i^* \geq 0 \quad \forall x \in [0, 1]$ for at least one $i \in \{A, B\}$.

Substituting (4.22) into (4.20), we have $v_2^* - v_1^* = \frac{3\bar{\varepsilon}}{4}\alpha - \frac{p_2 - p_1}{3c} \geq \frac{3t\bar{\varepsilon}}{16} + \frac{2(p_2 - p_1)}{3c} > 0$, thus we only need $v_1 \geq 0$ for QoS feasibility. Furthermore, the Hotelling model feasibility condition implies $v_1^* - tx_1^*\varepsilon_1^* \geq v_2^* - tx_2^*\varepsilon_2^*$. Since $u_i(x)$ is an increasing function of x, complete market coverage condition can be simplified to $u_1(0) \geq 0$. As a result, the QoS feasibility condition and the complete market coverage condition can be simplified to $v_1^* - tx_1^*\varepsilon_1^* \geq 0$. Therefore, the sufficient condition for the existence of SPNE is given by:

1. $0 \leq \frac{v_1^* - v_2^* + t(x_2^*\varepsilon_2 - x_1^*\varepsilon_1)}{t(\varepsilon_2^* - \varepsilon_1^*)} \leq 1$,

2. $0 \leq \varepsilon_1^*, \varepsilon_2^* \leq \bar{\varepsilon}$,

3. $v_1^* - tx_1^*\varepsilon_1^* \geq 0$.

Solving the above three inequalities yield (4.21), (4.22), and (4.23). The equilibrium market share and profits of the SPs are obtained by substituting (4.17), (4.18), (4.19), and (4.20) into (4.7), (4.8), and (4.9). □

APPENDIX G

PROOF OF THEOREM 6

Proof. If $q = \frac{1}{2}$, X is independent of Y. The optimal solution is given by any (s_0, s_1) that satisfies the distortion constraint $(\{s_0, s_1 | ps_1 + (1-p)s_0 \geq 1 - D, s_0, s_1 \in [0,1]\})$ since X and Y are already independent. If $q \neq \frac{1}{2}$, since each maximum in (5.29) can only be one of the two values (i.e., the inference accuracy of guessing $\hat{Y} = 0$ or $\hat{Y} = 1$), the objective function of the privatizer is determined by the relationship between $P(Y = 1, \hat{X} = i)$ and $P(Y = 0, \hat{X} = i), i \in \{0, 1\}$. Therefore, the optimization problem in (5.29) can be decomposed into the following four subproblems:

Subproblem 1: $P(Y = 1, \hat{X} = 0) \geq P(Y = 0, \hat{X} = 0)$ and $P(Y = 1, \hat{X} = 1) \leq P(Y = 0, \hat{X} = 1)$, which implies $p(1-2q)(1-s_1) - (1-p)(1-2q)s_0 \geq 0$ and $(1-p)(1-2q)(1-s_0) - p(1-2q)s_1 \geq 0$. As a result, the objective of the privatizer is given by $P(Y = 1, \hat{X} = 0) + P(Y = 0, \hat{X} = 1)$. Thus, the optimization problem in (5.29) can be written as

$$\begin{aligned}
\min_{s_0, s_1} \quad & (2q-1)[ps_1 + (1-p)s_0] + 1 - q \\
s.t. \quad & 0 \leq s_0 \leq 1 \\
& 0 \leq s_1 \leq 1 \\
& p(1-2q)s_1 + (1-p)(1-2q)s_0 \leq p(1-2q) \\
& p(1-2q)s_1 + (1-p)(1-2q)s_0 \leq (1-p)(1-2q) \\
& -ps_1 - (1-p)s_0 \leq D - 1.
\end{aligned} \tag{G.1}$$

- If $1 - 2q > 0$, i.e., $q < \frac{1}{2}$, we have $ps_1 + (1-p)s_0 \leq p$ and $ps_1 + (1-p)s_0 \leq 1 - p$. The privatizer must maximize $ps_1 + (1-p)s_0$ to reduce the adversary's probability of correctly inferring the private variable. Thus, if $1 - D \leq \min\{p, 1-p\}$, the optimal value is given by $(2q-1)\min\{p, 1-p\} + 1 - q$; the corresponding optimal solution is given by $\{s_0, s_1 | ps_1 + (1-p)s_0 = \min\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$. Otherwise, the problem is infeasible.

- If $1-2q<0$, i.e., $q>\frac{1}{2}$, we have $ps_1+(1-p)s_0 \geq p$ and $ps_1+(1-p)s_0 \geq 1-p$. In this case, the privatizer has to minimize $ps_1+(1-p)s_0$. Thus, if $1-D \geq \max\{p, 1-p\}$, the optimal value is given by $(2q-1)(1-D)+1-q$; the corresponding optimal solution is $\{s_0, s_1 | ps_1+(1-p)s_0 = 1-D, 0 \leq s_0, s_1 \leq 1\}$. Otherwise, the optimal value is $(2q-1)\max\{p, 1-p\}+1-q$ and the corresponding optimal solution is given by $\{s_0, s_1 | ps_1+(1-p)s_0 = \max\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$.

Subproblem 2: $P(Y=1, \hat{X}=0) \leq P(Y=0, \hat{X}=0)$ and $P(Y=1, \hat{X}=1) \geq P(Y=0, \hat{X}=1)$, which implies $p(1-2q)(1-s_1)-(1-p)(1-2q)s_0 \leq 0$ and $(1-p)(1-2q)(1-s_0)-p(1-2q)s_1 \leq 0$. Thus, the objective of the privatizer is given by $P(Y=0, \hat{X}=0)+P(Y=1, \hat{X}=1)$. Therefore, the optimization problem in (5.29) can be written as

$$\begin{aligned}
\min_{s_0, s_1} \quad & (1-2q)[ps_1+(1-p)s_0]+q \\
s.t. \quad & 0 \leq s_0 \leq 1 \\
& 0 \leq s_1 \leq 1 \\
& -p(1-2q)s_1-(1-p)(1-2q)s_0 \leq -p(1-2q) \\
& -p(1-2q)s_1-(1-p)(1-2q)s_0 \leq -(1-p)(1-2q) \\
& -ps_1-(1-p)s_0 \leq D-1.
\end{aligned} \tag{G.2}$$

- If $1-2q>0$, i.e., $q<\frac{1}{2}$, we have $ps_1+(1-p)s_0 \geq p$ and $ps_1+(1-p)s_0 \geq 1-p$. The privatizer needs to minimize $ps_1+(1-p)s_0$ to reduce the adversary's probability of correctly inferring the private variable. Thus, if $1-D \geq \max\{p, 1-p\}$, the optimal value is given by $(1-2q)(1-D)+q$; the corresponding optimal solution is $\{s_0, s_1 | ps_1+(1-p)s_0 = 1-D, 0 \leq s_0, s_1 \leq 1\}$. Otherwise, the optimal value is $(1-2q)\max\{p, 1-p\}+q$ and the corresponding optimal solution is given by $\{s_0, s_1 | ps_1+(1-p)s_0 = \max\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$.

- If $1-2q<0$, i.e., $q>\frac{1}{2}$, we have $ps_1+(1-p)s_0 \leq p$ and $ps_1+(1-p)s_0 \leq 1-p$. In this case, the privatizer needs to maximize $ps_1+(1-p)s_0$. Thus, if $1-D \leq \min\{p, 1-p\}$, the optimal value is given by $(1-2q)\min\{p, 1-p\}+q$; the corresponding optimal solution is given by $\{s_0, s_1 | ps_1+(1-p)s_0 = \min\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$. Otherwise, the problem is infeasible.

Subproblem 3: $P(Y=1, \hat{X}=0) \geq P(Y=0, \hat{X}=0)$ and $P(Y=1, \hat{X}=1) \geq P(Y=0, \hat{X}=1)$, we have $p(1-2q)(1-s_1)-(1-p)(1-2q)s_0 \geq 0$ and $(1-p)(1-2q)(1-s_0)-p(1-2q)s_1 \leq 0$. Under this scenario, the objective function in (5.29) is given by $P(Y=1, \hat{X}=0)+P(Y=1, \hat{X}=1)$. Thus, the privatizer solves

$$\begin{aligned}
\min_{s_0,s_1} \quad & p(1-q)+(1-p)q \\
s.t. \quad & 0 \leq s_0 \leq 1 \\
& 0 \leq s_1 \leq 1 \\
& p(1-2q)s_1+(1-p)(1-2q)s_0 \leq p(1-2q) \\
& -p(1-2q)s_1-(1-p)(1-2q)s_0 \leq -(1-p)(1-2q) \\
& -ps_1-(1-p)s_0 \leq D-1.
\end{aligned} \tag{G.3}$$

- If $1-2q>0$, i.e., $q<\frac{1}{2}$, the problem becomes infeasible for $p<\frac{1}{2}$. For $p \geq \frac{1}{2}$, if $1-D>\max\{p, 1-p\}$, the problem is also infeasible; if $\min\{p, 1-p\} \leq 1-D \leq \max\{p, 1-p\}$, the optimal value is given by $p(1-q)+(1-p)q$ and the corresponding optimal solution is $\{s_0, s_1 | 1-D \leq ps_1+(1-p)s_0 \leq \max\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$; otherwise, the optimal value is $p(1-q)+(1-p)q$ and the corresponding optimal solution is given by $\{s_0, s_1 | \min\{p, 1-p\} \leq ps_1+(1-p)s_0 \leq \max\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$.

- If $1-2q<0$, i.e., $q>\frac{1}{2}$, the problem is infeasible for $p>\frac{1}{2}$. For $p \leq \frac{1}{2}$, if $1-D>\max\{p, 1-p\}$, the problem is also infeasible; if $\min\{p, 1-p\} \leq 1-D \leq \max\{p, 1-p\}$, the optimal value is given by $p(1-q)+(1-p)q$

and the corresponding optimal solution is $\{s_0, s_1 | 1 - D \leq ps_1 + (1-p)s_0 \leq \max\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$; otherwise, the optimal value is $p(1-q)+(1-p)q$ and the corresponding optimal solution is given by $\{s_0, s_1 | \min\{p, 1-p\} \leq ps_1 + (1-p)s_0 \leq \max\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$.

Subproblem 4: $P(Y = 1, \hat{X} = 0) \leq P(Y = 0, \hat{X} = 0)$ and $P(Y = 1, \hat{X} = 1) \leq P(Y = 0, \hat{X} = 1)$, which implies $p(1-2q)(1-s_1) - (1-p)(1-2q)s_0 \leq 0$ and $(1-p)(1-2q)(1-s_0) - p(1-2q)s_1 \geq 0$. Thus, the optimization problem in (5.29) is given by

$$\begin{aligned} \min_{s_0, s_1} \quad & pq + (1-p)(1-q) \\ s.t. \quad & 0 \leq s_0 \leq 1 \\ & 0 \leq s_1 \leq 1 \\ & -p(1-2q)s_1 - (1-p)(1-2q)s_0 \leq -p(1-2q) \\ & p(1-2q)s_1 + (1-p)(1-2q)s_0 \leq (1-p)(1-2q) \\ & -ps_1 - (1-p)s_0 \leq D - 1. \end{aligned} \tag{G.4}$$

- If $1 - 2q > 0$, i.e., $q < \frac{1}{2}$, the problem becomes infeasible for $p > \frac{1}{2}$. For $p \leq \frac{1}{2}$, if $1 - D > \max\{p, 1-p\}$, the problem is also infeasible; if $\min\{p, 1-p\} \leq 1 - D \leq \max\{p, 1-p\}$, the optimal value is given by $pq + (1-p)(1-q)$ and the corresponding optimal solution is $\{s_0, s_1 | 1 - D \leq ps_1 + (1-p)s_0 \leq \max\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$; otherwise, the optimal value is $pq+(1-p)(1-q)$ and the corresponding optimal solution is given by $\{s_0, s_1 | \min\{p, 1-p\} \leq ps_1 + (1-p)s_0 \leq \max\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$.

- If $1 - 2q < 0$, i.e., $q > \frac{1}{2}$, the problem becomes infeasible for $p < \frac{1}{2}$. For $p \geq \frac{1}{2}$, if $1 - D > \max\{p, 1-p\}$, the problem is also infeasible; if $\min\{p, 1-p\} \leq 1 - D \leq \max\{p, 1-p\}$, the optimal value is given by $pq + (1-p)(1-q)$ and the corresponding optimal solution is $\{s_0, s_1 | 1 - D \leq ps_1 + (1-p)s_0 \leq$

$\max\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$; otherwise, the optimal value is $pq + (1-p)(1-q)$ and the corresponding optimal solution is given by $\{s_0, s_1 | \min\{p, 1-p\} \leq ps_1 + (1-p)s_0 \leq \max\{p, 1-p\}, 0 \leq s_0, s_1 \leq 1\}$.

Summarizing the analysis above yields Theorem 6. $\square$

APPENDIX H

PROOF OF THEOREM 7

Proof. Let us consider $\hat{X} = X + \beta + \gamma N$, where $\beta \in \mathbb{R}$ and $\gamma \geq 0$. Given the MAP adversary's optimal inference accuracy in (5.36), the objective of the privatizer is to

$$\begin{aligned} \min_{\beta,\gamma} \quad & P_{\mathrm{d}}^{(\mathrm{G})} \\ s.t. \quad & \beta^2 + \gamma^2 \leq D \\ & \gamma \geq 0. \end{aligned} \tag{H.1}$$

Define $\frac{1-\tilde{p}}{\tilde{p}} = \eta$. The gradient of $P_{\mathrm{d}}^{(\mathrm{G})}$ *w.r.t.* α is given by

$$\begin{aligned} \frac{\partial P_{\mathrm{d}}^{(\mathrm{G})}}{\partial \alpha} =& \tilde{p}\left(-\frac{1}{\sqrt{2\pi}} e^{-\frac{\left(-\frac{\alpha}{2}+\frac{1}{\alpha}\ln\eta\right)^2}{2}}\right)\left(-\frac{1}{2} - \frac{1}{\alpha^2}\ln\eta\right) \qquad \text{(H.2)} \\ &+ (1-\tilde{p})\left(-\frac{1}{\sqrt{2\pi}} e^{-\frac{\left(-\frac{\alpha}{2}-\frac{1}{\alpha}\ln\eta\right)^2}{2}}\right)\left(-\frac{1}{2} + \frac{1}{\alpha^2}\ln\eta\right) \\ =& \frac{1}{2\sqrt{2\pi}}\left(\tilde{p} e^{-\frac{\left(-\frac{\alpha}{2}+\frac{1}{\alpha}\ln\eta\right)^2}{2}} + (1-\tilde{p}) e^{-\frac{\left(-\frac{\alpha}{2}-\frac{1}{\alpha}\ln\eta\right)^2}{2}}\right) \qquad \text{(H.3)} \\ &+ \frac{\ln\eta}{\alpha^2\sqrt{2\pi}}\left(\tilde{p} e^{-\frac{\left(-\frac{\alpha}{2}+\frac{1}{\alpha}\ln\eta\right)^2}{2}} - (1-\tilde{p}) e^{-\frac{\left(-\frac{\alpha}{2}-\frac{1}{\alpha}\ln\eta\right)^2}{2}}\right). \end{aligned}$$

Note that

$$\frac{\tilde{p} e^{-\frac{\left(-\frac{\alpha}{2}+\frac{1}{\alpha}\ln\eta\right)^2}{2}}}{(1-\tilde{p}) e^{-\frac{\left(-\frac{\alpha}{2}-\frac{1}{\alpha}\ln\eta\right)^2}{2}}} = \frac{\tilde{p}}{1-\tilde{p}} e^{\frac{\left(-\frac{\alpha}{2}-\frac{1}{\alpha}\ln\eta\right)^2 - \left(-\frac{\alpha}{2}+\frac{1}{\alpha}\ln\eta\right)^2}{2}} = \frac{\tilde{p}}{1-\tilde{p}} e^{\frac{2\ln\eta}{2}} = \frac{\tilde{p}}{1-\tilde{p}} e^{\ln\eta} = 1. \tag{H.4}$$

Therefore, the second term in (H.3) is 0. Furthermore, the first term in (H.3) is always positive. Thus, $P_{\mathrm{d}}^{(\mathrm{G})}$ is monotonically increasing in α. As a result, the optimization problem in (H.1) is equivalent to

$$\begin{aligned} \max_{\beta,\gamma} \quad & \sqrt{\gamma^2 + \sigma^2} \\ s.t. \quad & \beta^2 + \gamma^2 \leq D \\ & \gamma \geq 0. \end{aligned} \tag{H.5}$$

Therefore, the optimal solution is given by $\beta^* = 0$ and $\gamma^* = \sqrt{D}$. Substituting the optimal solution back into (5.36) yields the MAP probability of correctly inferring the private variable Y from $\hat{X}$. $\square$

APPENDIX I

PROOF OF THEOREM 8

Proof. Let us consider $\hat{X} = X + (1-Y)\beta_0 - Y\beta_1$, where β_0 and β_1 are both non-negative. Given the MAP adversary's optimal inference accuracy $P_{\mathrm{d}}^{(\mathrm{G})}$, the objective of the privatizer is to

$$\begin{aligned} \min_{\beta_0,\beta_1} \quad & P_{\mathrm{d}}^{(\mathrm{G})} \\ s.t. \quad & (1-\tilde{p})\beta_0^2 + \tilde{p}\beta_1^2 \leq D \\ & \beta_0, \beta_1 \geq 0. \end{aligned} \tag{I.1}$$

Recall that $P_{\mathrm{d}}^{(\mathrm{G})}$ is monotonically increasing in $\alpha = \frac{2\mu-(\beta_1+\beta_0)}{\sigma}$. As a result, the optimization problem in (I.1) is equivalent to

$$\begin{aligned} \max_{\beta_0,\beta_1} \quad & \beta_1 + \beta_0 \\ s.t. \quad & (1-\tilde{p})\beta_0^2 + \tilde{p}\beta_1^2 \leq D \\ & \beta_0, \beta_1 \geq 0. \end{aligned} \tag{I.2}$$

Note that the above optimization problem is convex. Therefore, using the KKT conditions, we obtain the optimal solution

$$\beta_0^* = \sqrt{\frac{\tilde{p}D}{1-\tilde{p}}}, \quad \beta_1^* = \sqrt{\frac{(1-\tilde{p})D}{\tilde{p}}}. \tag{I.3}$$

Substituting the above optimal solution into $P_{\mathrm{d}}^{(\mathrm{G})}$ yields the MAP probability of correctly inferring the private variable Y from $\hat{X}$. □

APPENDIX J

PROOF OF THEOREM 10

Proof. The objective function in (5.45) can be written as

$$2\begin{bmatrix}\mu_1 & \mu_2 & \dots & \mu_m\end{bmatrix}\begin{bmatrix}\frac{1}{\sigma_1^2+\sigma_{p_1}^2} & 0 & \dots & 0\\ 0 & \frac{1}{\sigma_2^2+\sigma_{p_2}^2} & \dots & 0\\ \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & \dots & \frac{1}{\sigma_m^2+\sigma_{p_m}^2}\end{bmatrix}2\begin{bmatrix}\mu_1\\ \mu_2\\ \vdots\\ \mu_m\end{bmatrix}=\sum_{i=1}^{m}\frac{4\mu_i^2}{\sigma_i^2+\sigma_{p_i}^2}.$$

Thus, the optimization problem in (5.45) is equivalent to

$$\begin{aligned}\min_{\beta,\sigma_{p_1}^2,\dots,\sigma_{p_m}^2} \quad & \sum_{i=1}^{m}\frac{\mu_i^2}{\sigma_i^2+\sigma_{p_i}^2} \\ s.t. \quad & \|\beta\|^2+tr(\Sigma_p)\le D \\ & \sigma_{p_i}^2\ge 0 \quad \forall i\in\{1,2,...m\}.\end{aligned} \tag{J.1}$$

Since a non-zero β does not affect the objective function but result in positive distortion, the optimal mechanism satisfies $\beta=(0,...,0)$. Furthermore, the Lagrangian of the above optimization problem is given by

$$L(\sigma_{p_1}^2,...,\sigma_{p_m}^2,\lambda)=\sum_{i=1}^{m}\frac{\mu_i^2}{\sigma_i^2+\sigma_{p_i}^2}+\lambda_0(\sum_{i=1}^{m}\sigma_{p_i}^2-D)-\sum_{i=1}^{m}\lambda_i\sigma_{p_i}^2, \tag{J.2}$$

where $\lambda=\{\lambda_0,...,\lambda_m\}$ denotes the Lagrangian multipliers associated with the constraints. Taking the derivatives of $L(\sigma_{p_1}^2,...,\sigma_{p_m}^2,\lambda)$ with respect to $\sigma_{p_i}^2,\forall i\in\{1,...,m\}$, we have

$$\frac{\partial L(\sigma_{p_1}^2,...,\sigma_{p_m}^2,\lambda)}{\partial\sigma_{p_i}^2}=-\frac{\mu_i^2}{(\sigma_i^2+\sigma_{p_i}^2)^2}+\lambda_0-\lambda_i. \tag{J.3}$$

Notice that the objective function in (5.45) is decreasing in $\sigma_{p_i}^2,\forall i\in\{1,...,m\}$. Thus, the optimal solution ${\sigma_{p_i}^*}^2$ satisfies $\sum_{i=1}^{m}{\sigma_{p_i}^*}^2=D$. By the KKT conditions, we have

$$\frac{\partial L(\sigma_{p_1}^2,...,\sigma_{p_m}^2,\lambda)}{\partial\sigma_{p_i}^2}\Big|_{\sigma_{p_i}^2={\sigma_{p_i}^*}^2,\lambda=\lambda^*}=-\frac{\mu_i^2}{(\sigma_i^2+{\sigma_{p_i}^*}^2)^2}+\lambda_0^*-\lambda_i^*=0. \tag{J.4}$$

Since $\lambda_i^*, i \in \{0, 1, ..., m\}$ is dual feasible, we have $\lambda_i^* \geq 0, i \in \{0, 1, ..., m\}$. Therefore

$$\lambda_0^* \geq \frac{\mu_i^2}{(\sigma_i^2 + {\sigma_{p_i}^*}^2)^2}.$$

If $\lambda_0^* > \frac{\mu_i^2}{\sigma_i^4}$, we have $\lambda_0^* > \frac{\mu_i^2}{(\sigma_i^2 + {\sigma_{p_i}^*}^2)^2}$. This implies $\lambda_i^* > 0$. Thus, by complementary slackness, ${\sigma_{p_i}^*}^2 = 0$. On the other hand, if $\lambda_0^* < \frac{\mu_i^2}{\sigma_i^4}$, we have ${\sigma_{p_i}^*}^2 > 0$. Furthermore, by the complementary slackness condition, $\lambda_i^* {\sigma_{p_i}^*}^2 = 0, \forall {\sigma_{p_i}^*}^2$. This implies $\lambda_i^* = 0, \forall {\sigma_{p_i}^*}^2 > 0$. As a result, for all ${\sigma_{p_i}^*}^2 > 0$, we have

$$\frac{|\mu_i|}{\sqrt{\lambda_0^*}} = \sigma_i^2 + {\sigma_{p_i}^*}^2. \tag{J.5}$$

Therefore, ${\sigma_{p_i}^*}^2 = \max\{\frac{|\mu_i|}{\sqrt{\lambda_0^*}} - \sigma_i^2, 0\} = \left(\frac{|\mu_i|}{\sqrt{\lambda_0^*}} - \sigma_i^2\right)^+$ with $\sum_{i=1}^{m} {\sigma_{p_i}^*}^2 = D$. Substitute this optimal solution into (5.36) with $\alpha = \sqrt{(2\mu)^T(\Sigma + \Sigma_p)^{-1} 2\mu}$, we obtain the accuracy of the MAP adversary. $\square$