The Foerder Institute for Economic Research
at Tel Aviv University

מכון למחקר כלכלי על שם
ד"ר ישעיהו פורדר
על יד אוניברסיטת תל אביב

The Eitan Berglas School of Economics

בית-הספר לכלכלה ע"ש איתן ברגלס

עמותה רשומה

What Should a Firm Know?
Protecting Consumers' Privacy Rents

Daniel Bird and Zvika Neeman

Working Paper No. 5-2020

The Foerder Institute for Economic Research
and
The Sackler Institute of Economic Studies

מכון פורדר, אוניברסיטת תל-אביב, תל אביב יפו 6997801 | טל': 03-6409255 | פקס: 03-6405815
Tel Aviv University, Tel Aviv Yafo 6997801 Israel | Tel. 972-3-6409255 | Fax 972-3-6405815 | email: foerder@tauex.tau.ac.il

What Should a Firm Know?
Protecting Consumers' Privacy Rents

Daniel Bird and Zvika Neeman*

Abstract

A monopolistic firm observes a signal about the state of the world and then makes a take-it-or-leave-it offer to an uninformed consumer who has recourse to some outside option. We provide a geometric characterization of the firm's information structure that maximizes the consumer's surplus: the optimal regime partitions the space of payoff states into polyhedral cones with disjoint interiors. We interpret our results in terms of the maximization of the consumer's "privacy rent." We illustrate and motivate our approach through the example of the regulation of the privacy of medical information in monopolistic health insurance markets.

JEL CODES: D42, D82, D83, L51.

KEYWORDS: Information Design, Privacy Design, Privacy Rents, Persuasion, Genetic Non-Discrimination.

*Eitan Berglas School of Economics, Tel Aviv University (e-mails: `dbird@tauex.tau.ac.il`, `zvika@tauex.tau.ac.il`).

We are very grateful to two anonymous referees and the Editor for constructive comments and suggestions. We have also benefited from comments by Yair Antler, Eddie Dekel, Dan Zeltzer, seminar participants at the University of Essex, Technion and Tel Aviv University, as well as conference participants at Israeli IO Day 2019 and the Econometric Society World Congress 2020. Bird acknowledges financial support from the Pinhas Sapir Center for Development and the Forder Institute.

1 Introduction

We consider a model in which a monopolistic firm observes a signal about the state of the world and then makes a take-it-or-leave-it offer to an uninformed consumer who has recourse to some outside option. We focus our attention on situations in which the state of the world reflects the consumer's personal characteristics. In such situations, we ask how it is possible to design the firm's information structure so as to maximize the consumer's surplus. In other words, we solve for the policy that maximizes the consumer's information rent, or "privacy rent."

Our study is motivated by the following example. Recent technological developments in personalized medicine allow health providers to tailor medical treatment to each consumer's personal genetic profile. Genetic bio-markers indicate susceptibility to disease, and so genetic testing can be used to predict the efficacy of various medical treatments. This type of information, although personal, is nevertheless proprietary to the health provider whose research and expertise generated it (in some cases, it is also protected by patent). No one else is able to perform the required diagnostic tests or interpret their results. In particular, the consumer does not know or would not even be able to understand this information if it were presented to him without proper explanation. This raises the concern that once the health provider learns a consumer's genetic profile, it may either bias the treatment selection it offers the consumer to reduce its cost at the expense of the consumer's welfare or deny coverage to the consumer altogether.[1]

The medical provider may be viewed as a firm that observes a signal about the state of the world, which is given by the consumer's genetic makeup, and offers the consumer a choice among several insurance or treatment options. We are interested in the question of how a designer should restrict medical providers' use of information to allow consumers to reap the full benefits of personalized medicine. Put more generally, we address the problem of how to balance firms' use of information so that they have enough information to facilitate mutually beneficial transactions, but are prevented from extracting too much surplus through exploitation of their informational advantage.

More concretely, suppose that a monopolistic insurance firm sells a standard medical insurance policy that covers treatment for some medical condition. Moreover, suppose that the firm is able to learn the state of the world (i.e., the consumer's genetic profile) through genetic testing. Suppose that there are three equally likely states of the world: the consumer is *immune* to the relevant medical condition, the consumer is *responsive* to treatment for this condition, or the consumer is *non-responsive* to treatment.[2]

[1]For example, insurance firms can, and do, deny life insurance to carriers of cancer genes such as BRCA. The Genetic Literacy Project website (https://geneticliteracyproject.org) lists many such examples.

[2]Genetic tests are typically divided into two classes: predictive and prognostic. Predictive tests pre-

We describe the consumer's payoffs and the firm's profit from the insurance policy in the table below. The payoffs to both the consumer and the firm from not offering the policy are normalized to zero.

State	Consumer's payoff	Firm's profit
immune	-1	1
responsive	2	-1
non-responsive	-2	-1

An immune consumer has no need for medical insurance, and so his payoff from the medical insurance policy is -1, a consumer who is responsive to the treatment that is provided through the policy has a payoff of 2, and a non-responsive consumer has a payoff of -2 because on top of the insurance premium, he also bears the costs associated with receiving treatment. The firm that offers this policy obtains a profit of 1 from an immune consumer (because such a consumer has no need of costly treatment), a profit of -1 from a responsive consumer (because such a consumer requires costly treatment), and a profit of -1 from a non-responsive consumer (because such a consumer requires the same costly treatment as a responsive one).

If the firm has no information about the consumer's genetic profile, possibly because genetic testing is prohibited by a genetic nondiscrimination law, then the firm will not offer the consumer any coverage because the expected profit from the standard insurance policy is negative. Moreover, even if the standard insurance policy were offered, then the consumer would reject it because it would provide him with a negative expected payoff.[3]

Suppose instead that genetic testing is unregulated and the firm knows the consumer's genetic profile. If the firm offers insurance only if the consumer is immune, then the consumer will deduce from being offered insurance that he is immune and will reject the firm's offer. The strategy that maximizes the firm's profit subject to the constraint that the consumer is willing to accept the firm's offer is to offer the insurance policy to an immune consumer and to a responsive consumer with probability one-half. Such a strategy generates an expected payoff of 0 for a consumer who is offered the insurance policy, and

dict susceptibility to disease (e.g., the BRCA gene is a predictive bio-marker for breast cancer, and polymorphisms in the gene encoding the enzyme MTHFR are predictive bio-markers for acute lymphocytic leukemia). Prognostic tests predict the efficacy of a treatment if such is needed (e.g., the TPMT genetic test screens for a gene variant of thiopurine methyltransferase to assess the effectiveness of the thiopurine class of drugs that are prescribed to patients suffering from leukemia and autoimmune disorders, and the Prometheus IBD Serology 7 test kit identifies a subset of inflammatory bowel disease patients who will benefit from budesonide).

[3]The fact that for *some* price both the firm and the consumer prefer no coverage to coverage implies that at least one of them would prefer no coverage at *any* price.

an ex-ante expected profit of $\frac{1}{6}$ for the firm. Hence, with no regulatory restrictions the consumer will not benefit from the possibility of genetic testing.

Suppose, however, that genetic privacy regulation implies that the firm can only learn if the consumer is non-responsive or not. In this case, the firm will not offer the policy if the consumer is non-responsive, but it will be willing to offer it if the consumer is not non-responsive. Such a strategy generates an ex-ante expected profit of 0 for the firm, and an ex-ante expected payoff of $\frac{1}{3}$ for the consumer.

This example demonstrates that judicious regulation of what firms may know (or how they can use their knowledge) can increase consumer surplus. We have illustrated our argument through the example of genetic or medical testing. But, the same argument applies to any "credence good" that is sold in a monopolistic market. Examples include medical treatment, repair services, other types of personal services, and expert advice.[4]

The popularity and success of products and services that require consumers to voluntarily share some of their personal information with firms, such as iPhone and Netflix, indicates that the voluntary sharing of private information can obviously benefit consumers in much broader settings than the one described above. Moreover, it is also widely appreciated that the information that is shared involuntarily through consumers' credit scores facilitates the smooth operation of credit markets. However, it is also at least as obvious that voluntary or involuntary sharing of some types of personal information may harm consumers, because it may permit price and other forms of discrimination, not to mention the scandalous manipulation of information such as that recently allegedly performed by Cambridge Analytica. This basic trade-off is illustrated by Varian's (2009) well-known example: a consumer who likes apples would like the apple seller to know whether he prefers Jonathan to Macintosh apples, but not his willingness to pay for his preferred type of apple.[5]

Lawmakers and regulators have appreciated this insight at least since the enacting of the Fair Credit Reporting Act of 1970. The act facilitates the sharing of individuals' recent financial history, but clearly states that certain financial events in a person's past cannot impact credit scores. The recent proliferation of data-gathering by firms has heightened the importance of such regulation and triggered a wave of new laws intended to protect the privacy of consumers.[6] Two laws that are especially relevant to this paper are

[4]A credence good is characterized by the fact that consumers can observe the utility they derive from the good ex post, but cannot tell if the type or quality of the good they have received is the ex-ante needed one; however, an expert seller is able to identify the type or quality that fits the consumer's needs by performing a diagnostic test (Darby and Karni 1973). For a survey of the economics of credence goods, see Dulleck and Kerschbamer (2006).

[5]For a recent survey on the economics of privacy, see Acquisti, Taylor and Wagman (2016).

[6]Examples include the European Union's General Data Protection Law of 2016 and India's proposed

the Genetic Information Nondiscrimination Act (GINA) of 2008 and the Preserving Employee Wellness Programs Act of 2017. GINA protects Americans from discrimination based on their genetic information in both health insurance and employment. In particular, GINA prohibits health insurers from discrimination based on the genetic information of enrollees. That is, health insurers may not use genetic information to make eligibility, coverage, underwriting, or premium-setting decisions. Moreover, health insurers may not request or require individuals (or their family members) to undergo genetic testing or to provide genetic information. The analysis in this paper (as in the example presented above) suggests that such a sweeping nondiscrimination law may be too blunt to be effective, and describes how a more nuanced approach may serve to promote social welfare. This concern is reflected by the Preserving Employee Wellness Programs Act of 2017 that limits the scope of GINA and allows employers to collect genetic data of employees and their families (and impose sanctions for noncompliance) for use in employee wellness programs.

We describe a model that generalizes the example presented above. In particular, we analyze how to regulate the interaction between a consumer that does not possess any relevant private information and a firm that requires sufficient information to facilitate welfare-enhancing trade. We describe the problem as one of how to structure the information of a monopolistic firm in a way that maximizes the consumer's expected payoff, subject to incentive compatibility constraints for both the firm and the consumer. This formulation is equivalent to a formulation in which the firm knows the state of the world and what is regulated is the firm's ability to condition its offer to the consumer on various characteristics of the consumer. In this sense, we describe a model of the optimal regulation of privacy.

The case in which the firm can make only two offers to the consumer (offer/not offer a standard contract) admits a particularly elegant formulation and solution. In this case, the payoff space for the problem can be reduced to a two-dimensional Euclidean space: payoffs and profits from the default offer are normalized to zero, and payoffs and profits from the other offer are measured relative to the origin. We show that optimal privacy requires that this two-dimensional space be divided into two half-spaces by a (straight) downward-sloping line that passes through the origin. Each half-space corresponds to an offer. The firm is only allowed to learn to which half-space the state of the world belongs, and the half-spaces are designed so that, in each half-space, the firm is induced to make the offer that is associated with this half-space, and this offer is in turn accepted by the

Personal Data Protection Bill of 2019, which aim to regulate privacy and data protection, or the United States' suggested Deceptive Experiences To Online Users Reduction Act, which aims to restrict large Internet platforms from manipulating consumers.

consumer.

If there are three or more offers that the firm can make, then the analysis is more involved. First, the incentive compatibility constraints for the firm become more complex. Intuitively, with more possible offers, the set of deviations is much richer: the firm may deviate after observing several rather than just one signal in order to make the deviation more palatable for the consumer. Second, the payoff space must be divided into more than two regions.

Nevertheless, we show that the aforementioned result generalizes in the following sense: notice that an alternative (but mathematically equivalent) description of the optimal solution for the case with two offers is that the payoff space is partitioned into polyhedral cones.[7] We show that this result continues to hold also in the case where the firm can make more than two offers. That is, optimal privacy is attained by partitioning the payoff space into polyhedral cones with disjoint interiors that are each associated with a specific offer.

This result provides a guideline for the evaluation of privacy policies by emphasizing the more relevant dimensions of the state and policy space. In particular, it suggests a measure of "closeness of states" that has the property that if the consumer receives a certain offer in some state, then he should, generally, also receive the same offer in all nearby states.[8] It is worth emphasizing that using this criterion implies that the optimal privacy regime may induce ex-post Pareto-dominated offers (see Example 3 below).

Related Literature

We present our model as a problem of optimal privacy design. However, in essence, we analyze a two-tier information design problem: a designer constructs the firm's information structure, and the firm, in turn, chooses the consumer's information structure. Accordingly, our paper is related to the literature on information design (see Bergemann and Morris 2019, for a recent survey) and, in particular, information design in monopoly-pricing problems (Bergemann, Brooks and Morris 2015; Roesler and Szentes 2017; Ichihashi 2020; Hidir and Vellodi 2021; Haghpanah and Siegel 2020).

The key difference between the aforementioned papers and ours is that in these papers the players do not transmit any information to one another. Roesler and Szentes (2017) excepted, the aforementioned papers assume that the buyer knows his type before deciding whether to purchase a good, and focus on designing the monopolist's information

[7]Each payoff vector is associated with the ray that emanates from the origin and passes through it. A polyhedral cone is a convex hull that contains a collection of such rays.

[8]We define this measure formally in Section 4.3.

about the buyer's type.[9] By contrast, Roesler and Szentes (2017) assume that the firm has no information about the buyer's type, and focus on the design of the buyer's information structure about his own valuation. A second distinction between the aforementioned papers and ours is that we allow for general preferences, while in those papers the firm's payoff is given by the price paid, and the buyer's utility is given by the difference between his valuation and the price. We discuss the above papers in more detail when we explain the relationship between our model and models of market segmentation in Section 6.3.

Another related paper is Garcia and Tsur (2021) who study information design in competitive insurance markets. They show that in order to reduce the welfare loss from adverse selection, a regulator should create risk pools that exhibit negative assortative matching. That is, in contrast to our results, the optimal solution implies that all individuals are covered by insurance (no exclusion), and high-cost individuals are necessarily bunched together with low-cost individuals.

In our model the firm's choice of an offer conveys information to the consumer about the state of nature. Hence, our paper is also related to the literature on Bayesian persuasion; see Kamenica and Gentzkow (2011) and the recent survey by Kamenica (2019). However, because the consumer in our model may only choose between what he was offered by the firm and the default, rather than choose any action from some predetermined set, our model is not a "pure persuasion" model. Within this literature, especially close to our work are those papers that study the impact of exogenous restrictions on the sender's ability to persuade, such as in the case with multiple receivers (e.g., Alonso and Camara 2016 and Galperti and Perego 2019), multiple senders (e.g., Gentzkow and Kamenica 2017 and Li and Norman 2021), or a privately informed receiver (e.g., Kolotilin et al. 2017 and Matyskova 2018). See Doval and Skreta (2021) for a general treatment of constrained information design problems.

The paper that is perhaps most closely related to ours is Ichihashi (2019) who analyzes the impact of coarsening the sender's information about the state of nature. Our work differs from his in three key aspects. First, whereas Ichihashi considers a model with only two possible actions, we allow for any finite number of actions or offers. As will become clear below, this difference has significant implications on the type of incentive compatibility constraints that the information designer must respect. Second, Ichihashi characterizes the set of payoffs that can be attained under some information structure, but does not consider how to attain each such payoff/profit pair. By contrast, our main result describes the geometry of the firm's information structure that maximizes the consumer's

[9] In related work, Ali, Lewis and Vasserman (2020) and Pram (2020) show that in order to improve his terms of trade, an informed buyer may wish to voluntarily renounce his privacy by disclosing hard information about himself.

expected payoff. Thus, our analysis provides an insight into the regulation of the firm's information that cannot be obtained from Ichihashi's work. And third, as mentioned above, Ichihashi considers persuasion. That is, upon receiving the sender's message, the receiver may choose either one of the two actions, whereas in our model, a receiver may choose the non-default action only if this action was offered by the firm.[10]

The rest of the paper proceeds as follows. Section 2 describes the model. The case with two offers is analyzed in Section 3, and the general case in Section 4. In Section 5 we discuss some economic implications of the model: in Section 5.1 we consider the specific case of the regulation of a monopolistic provider of medical insurance, and in Section 5.2 we consider the implications of our results on the effectiveness of regulation, under the special assumption that the firm's preferences are state independent. In Section 6 we consider the case of more general designer objectives, the case of pure persuasion, and explain how our basic model can be applied to the study of market segmentation and to the design of legislative procedures. Section 7 concludes. All proofs are relegated to the appendix.

2 Model

We consider the problem of a designer who attempts to protect an uninformed consumer from being manipulated by an interested monopolistic firm. Formally, the firm makes a take-it-or-leave-it offer to the consumer from a finite set $A = \{a_0, \ldots, a_N\}$, where a_0 is the default offer.[11] That is, a_0 denotes the outcome that prevails if the consumer rejects the firm's offer. The offers in the set A may be interpreted as different possible binding contracts(or terms of trade) between the consumer and the firm that have been approved by the designer. Note that two different offers in the set A may differ only in the specification of the payment made by the consumer to the firm.

The consumer's and firm's payoff from each offer depends on the state of the world $\theta \in \Theta = \{\theta_1, \ldots, \theta_M\}$ and is given by $u(a, \theta)$ and $v(a, \theta)$, respectively. The prior distribution of the state is denoted by π. We denote the vector $(\theta_1, \ldots, \theta_M)$ by $\overline{\theta}$ and the vector

[10]Our work also bears a formal resemblance to Wei (2020). In his model, a principal (the designer in our model) chooses an information upper bound for a rationally inattentive agent (the firm in our model). In a binary-action binary-state model, Wei shows that optimal disclosure involves information distortion, but to a lesser extent than in the case without learning costs. Forges and Renault (2021) also study a similar problem to ours, but in a cheap-talk setting. Their main result is that an informative equilibrium does not generally exist.

[11]Since the consumer is uninformed, offering the consumer a menu of alternatives from which to choose confers no advantage to the firm.

$(\pi(\theta_1), \ldots, \pi(\theta_M))$ by $\pi(\overline{\theta})$. The vectors $(u(a_i, \theta_1), \ldots, u(a_i, \theta_M))$ and $(v(a_i, \theta_1), \ldots, v(a_i, \theta_M))$ are denoted by $u(a_i, \overline{\theta})$ and $v(a_i, \overline{\theta})$, respectively. For simplicity, we normalize the payoffs from the default offer to zero in every state of the world for both the consumer and the firm (i.e., $u(a_0, \overline{\theta}) = v(a_0, \overline{\theta}) = 0$). This normalization implies that, in each state, a player's payoff is measured relative to his payoff from the default offer.

Before making its offer to the consumer, the firm observes a signal $s \in S$, upon which it can condition its offer. The firm's offer strategy $q(a|s)$ specifies the probability that the firm offers a after observing signal s. We depart from the literature by assuming that the information structure that gives rise to the signal that is observed by the firm is chosen strategically by the designer in order to maximize the consumer's expected payoff. We assume that the designer selects a finite set of signals S and a set of conditional distributions $\{p(s|\theta)\}_{\theta \in \Theta}$ that describe the probability that the firm observes signal $s \in S$ in state θ. We refer to $\langle p(s|\theta), S \rangle$ as a privacy regime.

In general information design problems, it is often not exactly clear how a given information structure can actually be implemented in practice. In the context of genetic testing information is obtained through diagnostic tests that the firm performs on the consumer. It follows that information design can only be implemented by restricting the set of tests the firm is allowed to perform. This also implies that the designer cannot simply inform (or require the firm to inform) the consumer about his type, because doing so is tantamount to requiring the firm to provide a free set of costly diagnostic tests to any consumer.[12]

Moreover, under this interpretation, the firm's offer strategy can be viewed as a contingent contract that specifies what treatments and services the consumer will be entitled to based on the outcome of the diagnostic tests the firm will perform. This perspective justifies both the firm's ability to commit to a recommendation strategy and the assumption that the consumer knows the firm's actual strategy.

The timing of the game is as follows. First, the designer selects the privacy regime. Then, before observing the realized signal, the firm selects its offer strategy $q(a|s)$. Finally, the state of the world and the signal are realized, the firm makes an offer to the consumer, and the consumer decides whether or not to accept the firm's offer. Note that even though the firm *can* (and does) commit to the strategy it employs to influence the consumer under the privacy regime imposed by the designer, it *cannot* commit to the strategy it would employ under other possible privacy regimes. That is, the firm does not have commitment power with respect to the designer and so it chooses an optimal

[12] It is rarely the case that one gene is a valid genetic bio-marker for more than one condition. This implies that any genetic information that the consumer already possesses about himself is unlikely to be relevant in his interaction with a firm that offers state-of-the-art personalized medicine.

strategy $q(\cdot|\cdot)$ for each privacy regime chosen by the designer. We use a perfect Bayesian equilibrium as our solution concept, and so the designer (correctly) believes that the firm will respond optimally to any privacy regime that the designer might choose. Importantly, the fact that the firm can commit implies that when the firm considers possible deviations from its offer strategy, it does so with the understanding that the consumer will be aware of the deviation, and that he will interpret the firm's offers correctly, given the prior and the firm's strategy.

Finally, while we focus on this single-consumer interpretation of the model, there is a well-known alternative interpretation according to which the monopolistic firm faces a continuum of individuals whose types are distributed according to the prior distribution π. The analysis is unchanged, and all of our results can be translated to this alternative interpretation. When we apply our results to markets for medical insurance in Section 5.1 below, we switch to this alternative interpretation.

3 The Case with Two Offers

We start by analyzing the simplest (nontrivial) version of our model in which the firm can make only two offers, the default offer a_0 and another offer a_1. It is instructive to present the consumer's payoff and the firm's profit as in Figure 1 below, which provides a geometric representation of the "payoff space" of the designer's problem. Each point in Figure 1 corresponds to a state of the world. A point's X-coordinate depicts $v(a_1, \theta)$, which, due to the normalization of payoffs, is also the difference between the firm's profits from offer a_1 and default offer a_0, and a point's Y-coordinate depicts $u(a_1, \theta)$, which is the difference between the consumer's payoffs from offer a_1 and default offer a_0. Note that Figure 1 contains no information about the probabilities of the different states.

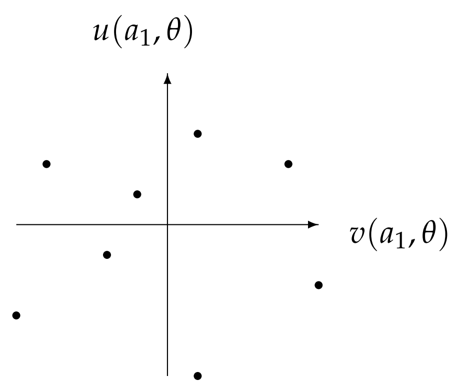

Figure 1: Geometric representation of payoff space

When there are only two offers, no loss of generality is implied by assuming that the designer selects an information structure with at most two signals: one that induces the firm to offer a_1 and one that induces the default offer[13] a_0. Thus, the designer's signals may be interpreted as instructions to the firm: "offer a_1" and "offer a_0," respectively. We refer to the states in which the designer instructs the firm to offer a_1 as the "acceptance region" and to the states in which the designer instructs the firm to offer the default offer a_0 as the "rejection region." Notice that the designer may send both signals in some states, and hence a state may belong to both regions.

Observe that if a state belongs to the first (top right) quadrant in Figure 1, then both the consumer and the firm prefer that this state belong to the acceptance region; if a state belongs to the third (bottom left) quadrant, then both the consumer and the firm prefer that this state belong to the rejection region; and if a state belongs to the second (top left) or fourth (bottom right) quadrant, then the consumer and firm have different preferences.

If it were possible, the designer (and consumer) would have liked the acceptance region to coincide with the area that lies above the X-axis in Figure 1. However, in this case, following the designer's instruction to offer a_1 is not necessarily a best response for the firm. For example, if states in the second quadrant are very likely, then under this "consumer-preferred partition" of the payoff space, a firm that is instructed by the designer to offer a_1 strictly prefers to offer the consumer a_0 instead.[14]

In general, when assigning states between the two regions, the designer must ensure that the firm is willing to forward the designer's suggestion to the consumer. That is, the firm must find it optimal to commit to offer a_1 if and only if it receives the signal that indicates that the state belongs to the acceptance region.

In this simple version of the model with two offers, this implies that the designer's choice of the optimal privacy regime must satisfy two incentive compatibility constraints:[15]

1. In the acceptance region the firm must weakly prefer a_1 to a_0.

2. In the rejection region the firm must either weakly prefer a_0 to a_1 or, if it prefers a_1 to a_0, then it must be the case that the consumer is unwilling to accept a_1 conditional on any convex combination of the consumer's induced beliefs in the acceptance and rejection regions.

The first point follows immediately from the fact that the firm cannot be forced to offer the consumer a_1. The second point is more subtle. If the firm prefers a_0 to a_1 in the rejection region, than obeying the designer's instruction to offer a_0 in the rejection

[13] We establish this "revelation principle" formally in Theorem 2 below.

[14] We slightly abuse terminology and refer to "partitions" of the payoff space, even though in some states more than one offer may be made.

[15] In Section 4.1 we formally derive the firm's incentive compatibility constraints for the general model.

region is obviously incentive compatible for the firm. But incentive compatibility is also possible if the firm prefers a_1 to a_0 in the rejection region. For this to be the case, it must be that any deviation to an offer strategy $q(a_1|rejection) = \epsilon$, $q(a_0|rejection) = 1 - \epsilon$, and $q(a_1|acceptance) = 1$, for some $\epsilon > 0$, induces the consumer to reject the firm's offer of a_1; that is, the new strategy violates the consumer's incentive compatibility constraint. In other words, the firm is willing to offer a_0 with certainty in the rejection region only if it is its preferred action there, or if it cannot induce the consumer to accept a_1 with a small probability in this region (note that the consumer will be aware that the firm is not following the designer's instructions, and so will understand that a_1 is offered both in the acceptance region and, with a small probability, also in the rejection region).

If a privacy regime is not incentive compatible for the firm, then the designer may "sweeten" it for the firm in one of two ways:

1. The designer can move states in the second quadrant, which give a positive payoff to the consumer but a negative payoff to the firm, from the acceptance to the rejection region.

2. The designer can move states in the fourth quadrant, which give a positive payoff to the firm but a negative payoff to the consumer, from the rejection to the acceptance region.

Note that both changes strengthen the firm's incentives to follow *both* of the instructions it can receive from the designer. For each state θ in the second and fourth quadrants, define the "slope of state θ" by the ratio

$$\rho(\theta) = \left| \frac{u(a_1, \theta)}{v(a_1, \theta)} \right|.$$

That is, $\rho(\theta)$ measures the ratio of the consumer's gain (loss) to the firm's loss (gain) from offer a_1 relative to the default offer.

To construct the optimal solution, the designer begins with the consumer-preferred partition and orders the states in both the second and fourth quadrants by their slopes $\rho(\theta)$. The designer then moves states with a small slope from one region to the other (from the acceptance to the rejection region in the second quadrant and from the rejection to the acceptance region in the fourth quadrant), until the firm is willing to make the suggested offer in both regions. Intuitively, states with a small slope $\rho(\theta)$ are more efficient sweeteners than states with a large slope because they generate a relatively large increase in the payoff to the firm at the expense of a relatively small loss in the payoff to the consumer. For example, assigning a state with a small slope in the second quadrant to the acceptance region provides the consumer with a small benefit relative to assigning this state to the rejection region. However, such an assignment induces a high relative cost to the firm.

We depict this idea graphically in Figure 2 below. Start with an acceptance region that coincides with the area above the X-axis. To make the problem nontrivial, suppose that the firm prefers a_0 to a_1 in both regions. Then begin progressively tilting the boundary of the acceptance region through the origin. Doing so moves the "average" state (depicted by hollow red circles) toward the south-east in the acceptance region, and toward the north-west in the rejection region. Because the average state in the acceptance region is initially positioned above the boundary line between the two regions, and because tilting moves this average state continuously, eventually the average state in the acceptance region will reach the boundary of the first quadrant. At this point, the firm is indifferent between the two offers in the acceptance region; i.e., the firm is willing to follow the designer's instruction in the acceptance region as well as in the rejection region.

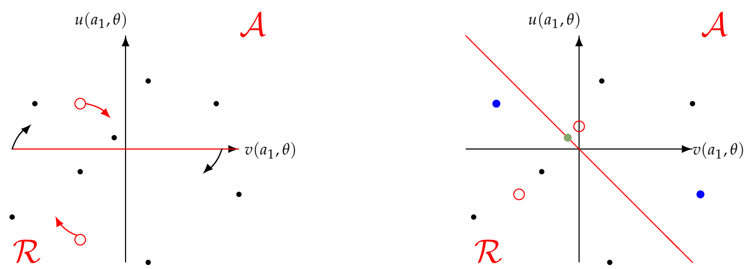

Figure 2: The optimal partition

In Figure 2, the states marked in blue are the states with the smallest slopes. These blue states are moved between the regions in order to sweeten the acceptance region for the firm. The green state is the state in which the designer sends both instructions with a positive probability.

We say that a privacy regime is a "half-plane regime" if it is characterized by a line with a negative slope that passes through the origin (as in Figure 2) such that states that lie above this line belong to the acceptance region and states that lie below this line belong to the rejection region. Generically, at most one state lies on the line, and this state may be split between the acceptance and rejection regions.

Notably, the tilting argument presented above was predicated on the assumption that both regions are nonempty under the optimal privacy regime, or, equivalently, that both offers are made with a positive probability. The argument would not work without this assumption. However, for some parameter values, this need not be the case. For example, if there are two equally likely states represented by $(-2, 1), (1, -2)$, then offer a_1 cannot be made and accepted in equilibrium. In this case, the optimal privacy regime induces only the default offer. The following sufficient condition ensures that both regions are nonempty and that both offers are made with a positive probability under the optimal privacy regime.

Lemma 1. *Suppose that the firm can make two offers and the first and third quadrants of the payoff space are nonempty. Then, there exists an optimal privacy regime in which both offers are made with a positive probability.*

Under the conditions of Lemma 1 the optimal privacy regime induces both offers and so the tilting argument presented above implies the following result.

Theorem 1. *Suppose that the firm can make only two offers. If the first and third quadrants of the payoff space are nonempty, then every optimal privacy regime is a half-plane regime.*

If no state is split between the acceptance and rejection regions in the optimal privacy regime, then the optimal privacy regime (defined in terms of which states are assigned to which regions) can be represented by many lines that pass through the origin of the payoff space. However, the tilting argument implies that the optimal privacy regime is (generically) unique in terms of which states are assigned to which regions.[16]

4 The General Case

We now return to the general case. We begin in Section 4.1 by deriving a "revelation principle" that characterizes the set of incentive-compatible "direct privacy regimes" from which the designer may choose. That is, we look at the privacy regimes in which the signal that is observed by the firm may in fact be interpreted as an instruction given by the designer to the firm regarding which offer the firm should make to the consumer. In Section 4.2 we present a few examples that illustrate the geometry of incentive-compatible regimes. Finally, in Section 4.3 we provide a geometric characterization of the optimal privacy regime.

4.1 A Revelation Principle

A "direct privacy regime" for the designer is given by a privacy regime $\langle p(s|\theta), S\rangle$, a firm's strategy $q(\cdot|\cdot)$, and a consumer's strategy, such that:

1. $S \subseteq A$.

[16] Uniqueness can only be established generically because if two states lie on the same line that passes through the origin, then only the weighted sum of the probabilities with which they are assigned to each region can be determined.

2. The firm follows the designer's instructions: $q(a|s = a) = 1$ for every "instruction" $a \in S$.

3. The consumer accepts the firm's offer.

A direct privacy regime is incentive compatible if the strategies described above are optimal for the firm and the consumer, given the designer's choice.

It is convenient to describe the incentive compatibility constraints through the restrictions they impose on the beliefs that a direct privacy regime induces for the firm and the consumer. Suppose that the designer uses the direct privacy regime $\langle p(s|\theta), S\rangle$. Denote the probability that the firm receives signal $a_i \in S$ by $p_{a_i} = \sum_{\theta \in \Theta} \pi(\theta)p(s = a_i|\theta)$, and the firm's conditional beliefs (over Θ) when that signal is received by μ_{a_i} . The beliefs μ_{a_i} are derived via Bayesian updating and are given by

$$\mu_{a_i}(\theta) = \frac{\pi(\theta)p(s = a_i|\theta)}{\sum_{\theta' \in \Theta} \pi(\theta')p(s = a_i|\theta')}.$$

Note that in a direct privacy regime, if the firm follows the designer's instructions, then the consumer's beliefs upon being offered a_i by the firm are also given by μ_{a_i}. Moreover, our assumption that the firm commits to its strategy implies that the firm's strategy is a general rule of behavior that is observable by the consumer. Therefore, were the firm to decide not to follow the instructions it receives, the consumer would be aware of that. Thus, for example, if the firm were to decide to offer a_j (with some probability) when it was instructed to offer a_i, then the consumer would realize that a_j is being offered in the event that it should be offered and also (with some probability) in the event that the firm was instructed by the designer to offer a_i. Hence, were the firm to deviate in such a way, the consumer's beliefs upon receiving offer a_i would be a convex combination of μ_{a_i} and μ_{a_j}.

In a direct privacy regime in which the firm obeys the designer's instructions, the consumer will accept the firm's offer if his payoff from doing so is no less than his payoff from the default offer. That is, the consumer will accept offer a_i if and only if

$$\mu_{a_i} \cdot u(a_i, \overline{\theta}) \geq 0. \tag{1}$$

Denote the set of beliefs under which the consumer accepts offer a_i by

$$\Xi_{a_i} = \{\mu \in \Pi : \mu \cdot u(a_i, \overline{\theta}) \geq 0\},$$

where Π is the space of probability distributions over Θ.

The consumer's incentive compatibility constraint (1) can therefore be written as

$$\mu_{a_i} \in \Xi_{a_i} \quad \forall a_i \in S. \tag{2}$$

The simplest incentive compatibility constraint for the firm captures the idea that when instructed to offer the consumer $a_i \in S$, there is no other offer that the firm (strictly) prefers to a_i that it can induce the consumer to accept instead. That is, if the firm strictly prefers offer a_j to offer a_i under its belief conditional on being instructed to make offer a_i, then it must be the case that the consumer will reject offer a_j under any non-degenerate convex combination of μ_{a_i} and μ_{a_j}. Observe that this requires, in turn, that (i) under the belief μ_{a_i}, the consumer reject a_j (otherwise, the consumer would accept the firm's offer of a_j instead of a_i), and (ii) under the belief μ_{a_j}, the consumer be indifferent between accepting and rejecting a_j (if the consumer's preference were strict, then the consumer would also accept a_j if it were also offered instead of a_i with a small probability).

To formalize this idea denote by $D(a_i)$ the set of offers that the firm (strictly) prefers to a_i under belief μ_{a_i}:

$$D(a_i) \equiv \{a_j \in A : \mu_{a_i} \cdot v(a_j, \overline{\theta}) > \mu_{a_i} \cdot v(a_i, \overline{\theta})\}.$$

To rule out the deviation described above, the following condition must hold:

$$\text{For every } a_i \in S, \text{ if } a_j \in D(a_i) \text{ then } \mu_{a_i} \notin \Xi_{a_j} \text{ and } \mu_{a_j} \in \partial\Xi_{a_j}, \tag{3}$$

where ∂X denotes the boundary of the set X. Note that the set $D(a_i)$ may include offers that are not suggested by the designer. For such an offer a_j, the condition $\mu_{a_j} \in \partial\Xi_{a_j}$ in (3) is vacuously satisfied, and so only the condition that $\mu_{a_i} \notin \Xi_{a_j}$ imposes a restriction.

The other type of incentive compatibility constraint for the firm deals with more complex deviations. Suppose that the firm would like to offer a_j after being instructed to offer a_i but knows that were it to change its strategy in that manner the consumer would reject a_j. One possible way in which the firm could still be able to get away with offering a_j instead of a_i in this case is if it "sweetens" offer a_j for the consumer as follows. The firm could offer a_j not only after being instructed to offer a_i, where the consumer would reject a_j under beliefs μ_{a_i}, but also after being instructed to offer a_k, where the consumer would accept a_j under beliefs μ_{a_k}.

Formally, for a given direct privacy regime, denote by $T(a_j)$ the set of signals that induce beliefs under which the consumer strictly prefers a_j to a_0:

$$T(a_j) \equiv \{a_k \in S : \mu_{a_k} \in int(\Xi_{a_j})\},$$

where $int(X)$ denotes the interior of a set X. In other words, $T(a_j)$ is the set of signals that the firm can use as a sweetener to induce the consumer to accept the firm's offer of a_j.

In order to incentivize the consumer to accept an offer of a_j not only when the designer instructs the firm to offer a_j, but also when, with some probability, the designer instructs the firm to offer a_i and a_k, the ratio between $p_{a_i}q(a_j|s = a_i)$ and $p_{a_k}q(a_j|s = a_k)$ must be sufficiently small.

Formally, for any such $a_j \in D(a_i)$ and $a_k \in T(a_j)$ define $\lambda(i, j, k)$ as the solution of the equation:

$$\lambda p_{a_i}\mu_{a_i} \cdot u(a_j, \overline{\theta}) + (1 - \lambda)p_{a_k}\mu_{a_k} \cdot u(a_j, \overline{\theta}) = 0.$$

That is, $\frac{\lambda(i;j,k)}{1-\lambda(i;j,k)}$ is the maximal ratio between $q(a_j|s = a_i)$ and $q(a_j|s = a_k)$ that induces the consumer to accept offer a_j after such a modification of the firm's strategy. Incentive compatibility for the firm requires that, in addition to constraint (3), there be no three offers $a_i \in S$, $a_j \in D(a_i)$, and $a_k \in T(a_j)$, such that the firm prefers to offer a_j after receiving the signals $s = a_i$ and $s = a_j$ (in the appropriate ratio) instead of following the designer's instructions, i.e.,

$$\begin{aligned}\lambda(i, j, k)p_{a_i}\mu_{a_i} \cdot v(a_i, \overline{\theta}) &+ (1 - \lambda(i, j, k))p_{a_k}\mu_{a_k} \cdot v(a_k, \overline{\theta}) \\ &\geq \Big(\lambda(i; j, k)p_{a_i}\mu_{a_i} + (1 - \lambda(i; j, k))p_{a_k}\mu_{a_k}\Big) \cdot v(a_j, \overline{\theta}) \\ &\forall a_i \in S, a_j \in D(a_i), a_k \in T(a_j). \end{aligned} \tag{4}$$

Solving for $\lambda(i; j, k)$ and simplifying shows that this constraint is equivalent to

$$\begin{aligned}\mu_{a_k} \cdot (v(a_k, \overline{\theta}) - v(a_j, \overline{\theta})) &\geq \frac{\mu_{a_k} \cdot u(a_j, \overline{\theta})}{-\mu_{a_i} \cdot u(a_j, \overline{\theta})}\mu_{a_i} \cdot (v(a_j, \overline{\theta}) - v(a_i, \overline{\theta})) \\ &\forall a_i \in S, a_j \in D(a_i), a_k \in T(a_j). \end{aligned} \tag{5}$$

This version of the constraint has the following intuitive interpretation as a cost-benefit analysis. The left-hand side represents the firm's cost from using signal a_k as a sweetener to induce the consumer to accept offer a_j. The right-hand side represents the firm's benefit from inducing the consumer to accept a_j rather than a_i under beliefs μ_{a_i}, scaled by a term that measures the effectiveness of signal a_k as a sweetener for offer a_j relative to the consumer's reluctance to accept a_j under beliefs μ_{a_i}.

The firm's problem of which take-it-or-leave-it offer to make to the consumer after each signal it observes is a linear optimization problem in the probabilities $q(\cdot, \cdot)$. Since linear problems are convex optimization problems, it follows that if the firm does not have a small improving deviation over a given offer strategy, then it does not have any improving deviation. This implies that we can derive a "revelation principle" that not only establishes the standard claim that restricting attention to direct incentive-compatible mechanisms involves no loss of generality, but also identifies a simple structure for the firm's incentive compatibility constraints. In particular, the two types of incentive compatibility constraints described above are necessary and sufficient for a direct mechanism to be incentive compatible for the firm.

Theorem 2. *If a state-contingent distribution of offers can be induced by some privacy regime to which both the firm and the consumer best respond, then the same distribution of state-contingent offers can also be induced by an incentive-compatible direct privacy regime. Furthermore, a direct privacy regime is incentive compatible if and only if it satisfies constraints* (2), (3), *and* (5).

Theorem 2 states that no loss of generality is entailed by restricting attention to direct privacy regimes that satisfy the consumer's incentive compatibility constraint (2) and the firm's incentive compatibility constraints (3) and (5). Hence, for the rest of the paper we restrict attention to such privacy regimes.

4.2 Examples of Incentive-Compatible Direct Privacy Regimes

In this section, we present two examples that illustrate the firm's incentive compatibility constraints in the belief simplex. In both examples we assume for simplicity that $v(a_i, \theta) = i$ for all θ. That is, the firm prefers higher-indexed offers in all states of the world.

Example 1: Two States, Two Offers

Suppose that $\Theta = \{\theta_1, \theta_2\}$ and $A = \{a_0, a_1\}$. If the consumer accepts a_1 under the prior belief ($\pi \in \Xi_{a_1}$), then if the firm always offers a_1 regardless of the designer's instruction, this offer will be accepted by the consumer. Hence, in this case, there is a single incentive-compatible direct privacy regime in which the designer instructs the firm to offer a_1 in both states of the world.

If the consumer rejects a_1 under the prior belief ($\pi \notin \Xi_{a_1}$), then in a direct privacy regime in which a_1 is offered with a positive probability, it must be the case that μ_{a_1} lies on the boundary of Ξ_{a_1}. This is because if μ_{a_1} belongs to the interior of Ξ_{a_1}, then the firm has the following profitable deviation: follow the designer's instructions, except, when instructed to offer a_0, offer a_1 with a small probability. The fact that μ_{a_1} belongs to the interior of Ξ_{a_1} implies that the consumer will accept the offer a_1 when it is suggested to him because if the probability is sufficiently small, then the consumer will still prefer a_1 to a_0 in the event that a_1 is offered to him.

Figure 3 depicts incentive-compatible and non-incentive-compatible direct privacy regimes (the figure displays the one-dimensional simplex representation of the state space).

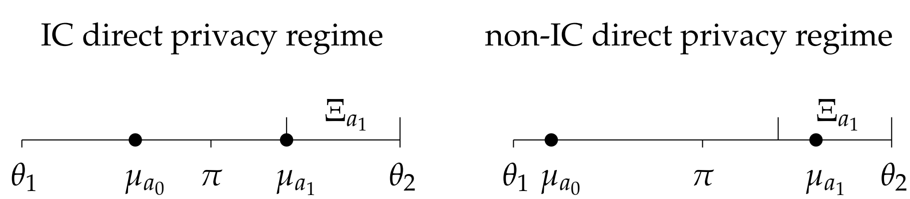

Figure 3: Incentive compatibility with two states and two offers

$\square$

Example 2: Three States, Three Offers

Suppose that $\Theta = \{\theta_1, \theta_2, \theta_3\}$ and $A = \{a_0, a_1, a_2\}$. We present two examples in which the designer instructs the firm to make all three offers with a positive probability.

First, we consider the case where the sets Ξ_{a_1} and Ξ_{a_2} are disjoint and the prior $\pi \notin \Xi_{a_1} \cup \Xi_{a_2}$. That is, under the prior beliefs, the consumer prefers a_0 to both a_1 and a_2. In this case, the firm's incentive compatibility constraint (3) has two implications: (i) the consumer must be indifferent between accepting and rejecting each of the offers a_1 and a_2 when they are made, and (ii) the consumer must reject both a_1 and a_2 under belief μ_{a_0}. Since the sets Ξ_{a_1} and Ξ_{a_2} are disjoint, in an incentive-compatible direct privacy regime signal a_1 cannot be used as a sweetener for offer a_2 (and vice versa), and so the firm has no profitable "complex deviations." Figure 4 depicts incentive-compatible and non-incentive-compatible direct privacy regimes in this case.

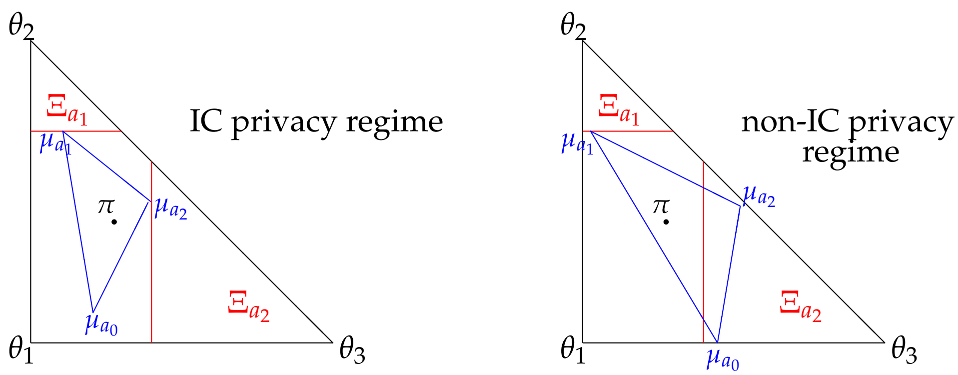

Figure 4: Incentive compatibility with three states and three offers

Note that the right-hand panel of Figure 4 contains two violations of the firm's simple incentive compatibility constraint. First, the fact that μ_{a_2} lies in the interior of Ξ_{a_2} implies that the firm can offer a_2 with a small probability also when it is instructed to offer a_1, and this offer would be accepted by the consumer. Second, the fact that μ_{a_0} belongs to Ξ_{a_2} implies that the firm can offer a_2 also when it is instructed to offer a_0, and, again, this offer would be accepted by the consumer.

Next, consider the case where $\Xi_{a_2} \subset \Xi_{a_1}$ and assume that the firm is instructed to make all offers with equal probabilities. To satisfy the firm's simple incentive compatibility constraint (3), the consumer's beliefs upon receiving offer a_i, $i \neq 0$, must be on the boundary of Ξ_{a_i}. In this example, we take these restrictions for granted, and focus on whether it is profitable for the firm to offer a_1 upon being instructed to offer both a_0 and a_2. That is, can the firm benefit from using signal a_2 as a sweetener to incentivize the consumer to accept offer a_1 under beliefs μ_{a_0}.

Suppose that the firm considers whether to offer the consumer a_1 with a small probability instead of offering a_0 and a_2 as instructed. Consider such a deviation where $\frac{q(a_1,s=a_0)}{q(a_1,s=a_2)} = \frac{\lambda}{1-\lambda}$. Constraint (4) requires that if such a deviation is incentive compatible for the consumer, then it is not profitable for the firm:

$$\lambda p_{a_0}\mu_{a_0}v(a_0,\overline{\theta}) + (1-\lambda)\mu_{a_2}p_{a_2}v(a_2,\overline{\theta}) \geq (\lambda p_{a_0}\mu_{a_0} + (1-\lambda)p_{a_2}\mu_{a_2})v(a_1,\overline{\theta}).$$

Given our assumptions on v and p the previous inequality can be simplified to

$$(1-\lambda)2 \geq (\lambda + (1-\lambda)) \text{ or } \lambda \leq \frac{1}{2}.$$

That is, any complex deviation that would be acceptable to the consumer must be such that the firm offers a_1 when instructed to offer a_0 with a weakly lower probability than the firm offers a_1 when instructed to offer a_2, which implies that the deviation is not beneficial for the firm. Geometrically, this implies that mixing μ_{a_2} and μ_{a_0} with equal weights must lead to a belief that is not in the interior of Ξ_{a_1}. The two examples in Figure 5 below depict two different utility functions for the consumer. In the example depicted in the left panel of Figure 5, a $\frac{1}{2}-\frac{1}{2}$ mix of μ_{a_2} and μ_{a_0} lies outside of the set Ξ_{a_1} and so the privacy regime described in the example satisfies incentive compatibility. In the example depicted in the right panel the same mixture lies in the interior of Ξ_{a_1} and so the privacy regime is not incentive compatible.

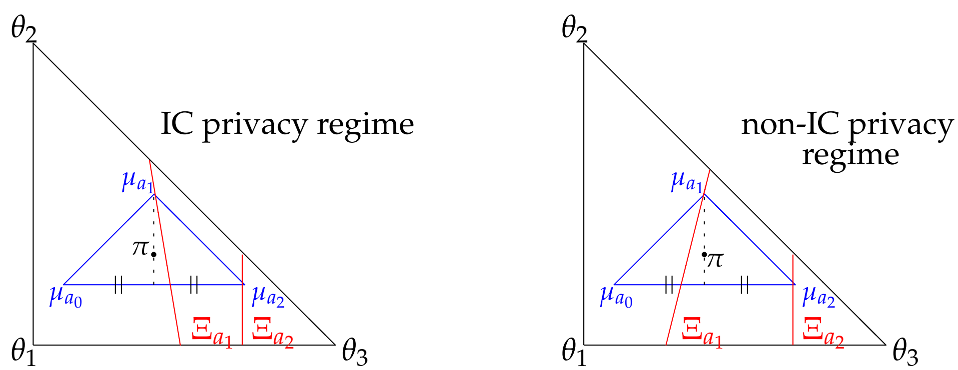

Figure 5: "Complex" incentive compatibility with three states and three offers

□

4.3 A Geometric Characterization of the Optimal Privacy Regime

Denote the vector of the consumer's and the firm's payoffs in state θ by

$$W(\theta) \equiv (u(a_1,\theta),\ldots,u(a_N,\theta),v(a_1,\theta),\ldots,v(a_N,\theta)).$$

We refer to the $2N$-dimensional Euclidean space that contains the set of vectors $\{W(\theta)\}_{\theta\in\Theta}$ as the payoff space.

Theorem 1 provides a geometric characterization of the optimal partition of the payoff space into acceptance and rejection regions for the case with two offers ($N = 1$). It shows that the two regions are given by two contiguous half-spaces that are separated by a hyperplane that passes through the origin. Clearly, with more than two possible offers it is impossible to separate the acceptance regions associated with different offers by a single hyperplane.

To generalize Theorem 1, we introduce the following definition:

Definition. *A set $C \in \mathbb{R}^n$ is a polyhedral cone if there exists a matrix $A \in \mathbb{R}^{m\times n}$ such that $C = \{x \in \mathbb{R}^n : A \cdot x \geq 0\}$.*

Note that a polyhedral cone is a convex set and a cone.

Theorem 1 implies that when there are two offers ($N = 1$) the acceptance and rejection regions are both given by polyhedral cones with disjoint interiors. Theorem 3 below shows that this property generalizes to any number of offers.

Theorem 3. *Suppose that the offers in the set $A^* \subseteq A$ are all made with a positive probability in an optimal privacy regime. Then, there exists a set of polyhedral cones $\{\mathcal{A}_{a_i}\}_{a_i\in A^*}$ with pairwise disjoint interiors that partition the payoff space such that $p(a_i, \theta) > 0$ only if $W(\theta) \in \mathcal{A}_{a_i}$.*

To understand the intuition behind Theorem 3 consider the problem faced by the designer. Ideally, the designer would like to assign each state to the acceptance region of the offer that provides the consumer with the highest payoff in that state. However, were the designer to try and do so the firm could use its information (i.e., the signals it receives) to induce the consumer to accept different offers in some states. Hence, in addition to considering the consumer's direct benefit from adding state θ to the acceptance region of a_i, the designer must also consider how sending signal $s = a_i$ in state θ impacts the firm's ability to use signal $s = a_i$ to induce the consumer to accept offers other than a_i. That is, for each signal $s = a_i$ and offer $a_j \neq a_i$ there is a shadow cost that measures the impact of changing the conditional belief after signal $s = a_i$ on the firm's incentive to follow the designer's instructions. In fact, the designer faces two additional shadow costs that are associated with the consumer's incentive compatibility constraint and the feasibility constraint $\Sigma_{s\in S}\, p(s, \theta) = 1$, which affect him in a similar way.

Now, consider the designer's problem in deciding whether to assign state θ to the acceptance region of a_i or the acceptance region of a_j. The value (incorporating the shadow costs) the designer obtains from assigning state θ to the acceptance region of any offer is linear in $W(\theta)$. Moreover, the shadow cost associated with the feasibility constraint is an additive term that does not depend on $W(\theta)$. Hence, there exists a hyperplane that passes through the origin that separates the payoff vectors that the designer would rather assign to the acceptance region of a_i from the payoff vectors that the designer would rather assign to the acceptance region of[17] a_j. Similarly, for every $a_k \neq a_i$ there exists a hyperplane that separates the payoff vectors in which a_i is better than a_k from the payoff vectors in which the opposite holds. It follows that the designer instructs the firm to offer a_i in state θ only if $W(\theta)$ lies above N different hyperplanes that pass through the origin. That is, the states in which a_i is offered are contained in a polyhedral cone.

Theorem 3 provides a guideline for evaluating privacy policies through the emphasis it places on the dimensions of the state space that should be considered. In particular, it suggests a measure of proximity of states that has the property that if two states are "close" to one another according to this measure, then they should also generally be assigned to the acceptance region of the same offer.

Definition. *For any two payoff vectors $W(\theta), W(\theta') \in \mathbb{R}^{2N}$, the cosine difference between $W(\theta)$ and $W(\theta')$ is given by*

$$\cos(W(\theta), W(\theta')) = \frac{W(\theta) \cdot W(\theta')}{||W(\theta)||\,||W(\theta')||},$$

where $||\cdot||$ denotes the l_2 norm.

Theorem 3 implies that if, under the optimal privacy regime, state θ belongs to the interior of the acceptance region of offer a_i, then the same offer will also be made in any state θ' that is sufficiently close to it (i.e., for which $\cos(W(\theta), W(\theta'))$ is sufficiently small).

Two states that have a small cosine difference may nevertheless be assigned to different offers under the optimal privacy regime if they lie on different sides of the boundary between their respective polyhedral cones. However, it is still possible to use the measure of cosine difference to pin down the allocation of states to offers. This can be done by introducing the "central ray of a polyhedral cone," that is, the ray that is equally distant from each of the hyperplanes that define the polyhedral cone. If we denote the central ray of polyhedral cone $\mathcal{A}_{a_i}$ by $C(\mathcal{A}_{a_i})$, then under an optimal privacy regime, state θ is assigned to $\mathcal{A}_{a_i}$ only if $W(\theta)$ is (weakly) closer to $C(\mathcal{A}_{a_i})$ than to any other central ray $C(\mathcal{A}_{a_j})$.

If we consider the alternative interpretation of the model in which the firm faces a continuum of individuals (and each state represents a different individual-type), then Theorem 3 suggests that an optimal privacy regime might generate discrimination between similar individual-types. Consider two individual-types θ and θ' such that (i) the

[17] This hyperplane passes through the origin because for the zero payoff vector the designer is indifferent between the two offers.

individuals' ordinal ranking of possible offers is identical for both individual-types, and (ii) the firm's ordinal ranking of possible offers is identical conditional on both individual-types. To an outside observer such individual-types may appear to be identical and, therefore, such an observer would expect these two individual-types to receive the same offer. However, Theorem 3 implies that these two individual-types may receive different offers. Thus, an optimal privacy regime might generate a perception of unjust discrimination to an outside observer.

Pareto-Efficiency of the Optimal Privacy Regime

It is well known that the provision of incentives may generally conflict with ex-post Pareto efficiency. Theorem 3 implies that the designer assigns states to offers based on the states' relative slope (as measured by the cosine difference), and not according to the vector of payoffs associated with each state. This suggests that the optimal design of privacy is also likely to violate Pareto efficiency.

The next example demonstrates that under the optimal privacy regime the firm may be instructed to make a strictly Pareto-dominated offer in some states.

Example 3: The optimal privacy regime induces ex-post Pareto-dominated offers Suppose that there are two equally likely states $\{\theta_1, \theta_2\}$ and three possible offers $\{a_0, a_1, a_2\}$. The firm strictly prefers a higher- to a lower-indexed offer in both states. The consumer's payoffs are

$$u(a_1, \overline{\theta}) = \begin{pmatrix} 1 \\ 1 \end{pmatrix}; \qquad u(a_2, \overline{\theta}) = \begin{pmatrix} -4 \\ 2 \end{pmatrix}.$$

In state θ_2, offer a_2 is the unique Pareto-efficient offer. Thus, ex-post Pareto efficiency requires that offer a_2 be made in state θ_2. However, we show that under the optimal privacy regime the firm is instructed to offer a_1 in both states of the world.

To see this, note first that the consumer's expected payoff from a_2, if it is offered, is zero under any incentive-compatible privacy regime. Clearly, the consumer's expected payoff from a_2 cannot be negative. And, if it is (strictly) positive, the consumer will still accept the firm's offer of a_2 if the firm also offers a_2 with a small probability when it is instructed to offer a_1. Because the firm prefers a_2 to a_1 in both states, such a deviation is profitable for the firm, which violates the firm's incentive compatibility constraint.

By construction, the consumer's payoffs from offers a_1 and a_0 are 1 and 0, respectively, regardless of the privacy regime. Therefore, the consumer's expected payoff under any incentive-compatible privacy regime is increasing in the probability that a_1 is made. A firm that is not given any information about the state will offer the consumer a_1 because it prefers it to a_0, and it realizes that the consumer will reject a_2 if the firm offers it with some (equal) probability in both states. It therefore follows that the direct privacy regime in which the designer instructs the firm to offer a_1 in both states is the optimal incentive-compatible privacy regime. $\square$

The existence of at least three possible offers is necessary for the optimal privacy regime to violate ex-post Pareto efficiency. If the designer modifies an incentive-compatible privacy regime by moving a state in which offer a_i is Pareto-dominant into a_i's acceptance region, then the consumer will *strictly* prefer to accept offer a_i when it is made. Hence, after such a change to the privacy regime, the firm can deviate and make offer a_i, with some probability, after it receives an instruction to make another offer a_j. If $N \geq 2$ and the consumer is already receiving a strictly positive payoff from offer a_j, then the attempt to make the privacy regime ex-post Pareto efficient may result in a reduction of the consumer's payoff, due to the implied change in the firm's behavior. This is exactly what happens in Example 3.

Example 3, continued
Observe that if the designer modifies the optimal but Pareto-inefficient privacy regime derived above by instructing the firm to make offer a_2 in state θ_2 with probability $p > 0$ (and offer a_1 otherwise), then the firm will respond by making offer a_2 not only when it is instructed to do so, but also with probability $\frac{2p}{3+p}$ when it is instructed to offer a_1. The consumer's gain from receiving offer a_2 (rather than a_1) in state θ_2 is $p + (1-p)\frac{2p}{3+p}$, while his loss from receiving a_2 (rather than a_1) in state θ_1 is $\frac{2p}{3+p} \cdot (1-(-4))$. It is straightforward to verify that the gain is smaller than the loss for any $p > 0$. $\square$

By contrast, when there is a single non-default offer, moving a state in which offer a_i is Pareto-dominated from the acceptance region of that offer to the acceptance region of the other offer cannot result in a reduction of the consumer's expected payoff because the firm will disobey the designer's instructions. To see this, distinguish between the following two cases. If under the original privacy regime the consumer is indifferent between accepting and rejecting offer a_1, then his expected utility under the original privacy regime is zero, which is the lowest possible payoff that the consumer can get in any incentive-compatible privacy regime. In this case, no modification of the privacy regime can possibly harm the consumer, which implies that any modification necessarily makes the consumer weakly better off. If, on the other hand, the consumer strictly prefers to accept offer a_1 when he receives it, then the firm must (weakly) prefer to make offer a_i when it is instructed to do so. This, in turn, implies that moving a state in which offer a_i is Pareto-dominated from the acceptance region of that offer to the acceptance region of the other offer slackens the incentive compatibility constraints of both the firm and the consumer. It therefore follows that if $N = 1$, then there exists an optimal privacy regime that is ex-post Pareto efficient.

5 Economic Implications

5.1 Application: Health Insurance Markets

Consider the following stylized example of an insurance market where a monopolistic insurance firm offers a single insurance contract. As in Handel, Hendel and Whinston (2015), each individual is indexed by his type, θ, which represents the individual's level of health. We assume that healthier types are more profitable to insure, and so we let θ represent the insurance firm's profit from insuring a type-θ individual. For simplicity, we consider the case where the distribution of types is continuous on some interval[18] $[\underline{\theta}, \overline{\theta}]$. We denote the utility of an individual with health type θ from buying insurance by $U(\theta)$, and assume that the function $U(\cdot)$ is differentiable and decreasing.

To make the problem interesting, we focus on the case where it is profitable for the firm to insure some individuals, but the firm's profit from insuring the entire population is negative. Formally, this implies that $\underline{\theta} < 0 < \overline{\theta}$ and $\mathbb{E}[\theta] < 0$. We also assume that $U(0) > 0$, which, together with the assumption that $U(\cdot)$ is decreasing, implies that for every individual-type in the population, *either* the firm or the individual benefits from insurance.

The designer's problem is to determine which types should receive insurance and which should not. A simple solution to this problem is to use an interval allocation rule in which individuals' types that are extremely costly to insure are excluded from the market and individuals' types that are not so costly to insure (or are profitable to insure) receive insurance. The natural candidate for the critical type that separates the two groups is the type $\hat{\theta}$ for which $\mathbb{E}[\theta | \theta \geq \hat{\theta}] = 0$. That is, the critical type $\hat{\theta}$ is such that the firm is just willing to insure all individuals who are at least as healthy as $\hat{\theta}$. We show that Theorem 1 can be used to determine if this simple allocation rule is in fact optimal.

Example 4: A Uniform Quadratic Example

Suppose that θ is distributed uniformly on the interval $[-2, 1]$. Note that for this distribution the firm is unwilling to insure the entire population, and that the critical type is given by $\hat{\theta} = -1$. Consider the two utility functions $U(\theta) = \frac{1}{4}\theta^2 - 3\theta + 1$ and $\tilde{U}(\theta) = 3\theta^2 - 6\theta + \frac{1}{2}$ that are depicted (in blue) in Figure 6 below. Figure 6 also depicts the

[18] In Section 2 we assumed that the type space is finite, but the results described in this section are simpler to describe with a continuum of types.

(dashed) lines that connect the origin with the points $(\hat{\theta}, U(\hat{\theta}))$ and $(\hat{\theta}, \tilde{U}(\hat{\theta}))$, respectively. Note that the two utility functions are drawn to different scales.

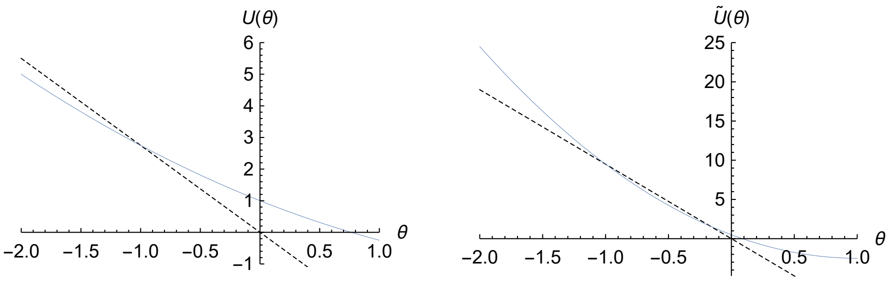

Figure 6: The payoff space in the health market

Note that for both utility functions the first quadrant of the payoff space is nonempty. Thus, the first part of the proof of Lemma 1 implies that some individual-types will receive insurance under the optimal privacy regime. On the other hand, the assumption that $\mathbb{E}(\theta) < 0$ implies that some individual-types are not insured under the optimal privacy regime. Therefore, Theorem 1 implies that for both of these utility functions, the optimal partition of the payoff space is given by a downward-sloping line that passes through the origin, such that all individual-types whose utility lies above this line receive insurance, and all individual-types whose utility lies below this line do not.

The simple interval allocation rule is optimal if and only if the set of types whose utility lies above the dashed line connecting the points $(\hat{\theta}, U(\hat{\theta}))$ and $(\hat{\theta}, \tilde{U}(\hat{\theta}))$ to the origin coincides with the set[19] $\{\theta : \theta \geq \hat{\theta}\}$. On the left-hand side of Figure 6 this holds true, but on the right-hand side it does not. Thus, the simple interval allocation rule is optimal for U but not for $\tilde{U}$. Moreover, in the latter case, the set of types that should optimally receive insurance are not connected.[20] □

Note that on the right-hand side of Figure 6 the slope $\frac{\tilde{U}(\underline{\theta})}{\underline{\theta}}$ is greater in absolute value than the slope $\frac{\tilde{U}(\hat{\theta})}{\hat{\theta}}$, while on the left-hand side the slope $\frac{U(\underline{\theta})}{\underline{\theta}}$ is smaller in absolute value than the slope of $\frac{U(\hat{\theta})}{\hat{\theta}}$. Thus, Example 4 suggests that whether the set of individual-types that receive insurance is an interval may be related to the relation between these two slopes. The next proposition establishes this formally.

[19]Since $U(\cdot)$ is convex and $U(0) > 0$, it follows that, for this allocation, individual-types that are offered insurance will purchase it.

[20]It can be shown that the set of types that receive insurance under the optimal privacy regime for $\tilde{U}$ is approximately given by $[-2, -1.735] \cup [-0.096, 1]$.

Proposition 4. *Suppose that θ is distributed continuously on $[\underline{\theta}, \overline{\theta}]$, where $\underline{\theta} < 0 < \overline{\theta}$ and $\mathbb{E}[\theta] < 0$. Suppose also that $U(\cdot)$ is differentiable, decreasing, convex, and is such that $U(0) > 0$. The optimal privacy regime is such that:*

1. *There exists $\epsilon > 0$ such that all individuals with types $\theta \geq -\epsilon$ receive insurance.*

2. *The set of individual-types that receive insurance is an interval if and only if*

$$\left| \frac{U\left(\hat{\theta}\right)}{\hat{\theta}} \right| \geq \left| \frac{U(\underline{\theta})}{\underline{\theta}} \right|.$$

The first part of Proposition 4 follows from the assumptions that $U(\theta)$ is convex, decreasing, and that $U(0) > 0$. These assumptions jointly imply that it is socially valuable to force individuals' types that are profitable to insure to buy insurance, even if they personally would prefer not to do so. Moreover, under these assumptions, the designer optimally uses these individuals' types to subsidize the slightly unhealthy types (i.e., individuals' types for which θ is negative but close to zero) for whom the cost of providing insurance is low and the value of receiving insurance is moderate.

The second part of the proposition addresses the question of whether the designer should also subsidize the extremely unhealthy individual-types (i.e., those types that lie at the lower end of the support) or the mildly unhealthy individual-types. The condition in this part of Proposition 4 is related to the curvature of $U(\cdot)$ on $[\underline{\theta}, \hat{\theta}]$, as captured by the ratio $\frac{U(\underline{\theta})}{U(\hat{\theta})}$: because $U(\cdot)$ is decreasing, this condition is violated only if that curvature is large. If $U(\cdot)$ is highly convex, then the designer's gain from insuring extremely unhealthy individual-types is very large relative to his gain from insuring mildly unhealthy individual-types. Thus, even though the insurance firm requires a greater subsidy to insure the former types than to insure the latter types, it is optimal for the designer to use the healthy individual-types to subsidize the extremely unhealthy individual-types. Hence, in this case, the designer will optimally select to provide insurance to an unconnected set of individual-types. On the other hand, if the curvature of $U(\cdot)$ is mild, the additional benefit of insuring the extremely unhealthy individual-types (rather than the mildly unhealthy ones) does not justify the larger subsidy that is required to incentivize the firm to do so. Hence, in this case, the designer will optimally select an interval allocation rule.

Notice that even in the stylized setting considered in this section (with a single insurance plan), the optimal allocation of individual-types to insurance plans does not always respect the natural ordering of types. In a more general setting, where the insurer offers multiple insurance plans, the payoff vector that represents each individual's type is a multidimensional vector that is determined by the cost and benefit of offering each insurance plan to that type. Even if the payoff function associated with each insurance plan is well behaved (e.g., monotone, convex, etc.), there is no reason to expect that the cosine difference between each type and the central rays of the polyhedral cones that represent the various insurance plans will be well behaved as well. Hence, the optimal allocation of individual-types to insurance plans need not generally partition the natural ordering of types into convex subsets.

5.2 The Effectiveness of Regulation with State-Independent Profit

Regulation in our model is generally effective. The expected payoff of the consumer under the optimal privacy regime is generally larger than under the two extreme benchmarks of a fully informed unregulated firm and an uninformed firm. However, when the firm's preferences are independent of the state of the world – a special case that has received much attention in the Bayesian persuasion literature – the ability of the designer to help the consumer is limited. In this section we analyze this special case. With no additional loss of generality, suppose that $v(a_j,\cdot) > v(a_i,\cdot)$ if[21] $j > i$.

Example 5: Benefit of Regulation with State-Independent Firm's Profits
Suppose that there are three equally likely states $\{\theta_1, \theta_2, \theta_3\}$ and that the consumer's payoffs are

$$u(a_1, \overline{\theta}) = \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}; \qquad u(a_2, \overline{\theta}) = \begin{pmatrix} -1 \\ 0 \\ 0 \end{pmatrix}.$$

If the firm knows the state and is unregulated, then it will offer the consumer a_2 in states θ_2 and θ_3, and a_1 in state θ_1. This provides the consumer with an expected payoff of $\frac{1}{3}$. An uninformed firm would offer a_1, which would also give the consumer an expected payoff of $\frac{1}{3}$. Under the optimal privacy regime the firm's induced information partition is given by $\{\theta_1, \theta_2\}, \{\theta_3\}$. In this case, the firm would offer a_1 in states $\{\theta_1, \theta_2\}$, and a_2 in state θ_3. This offer strategy gives the consumer an expected payoff of $\frac{2}{3}$. Notice that the consumer derives benefit only from offer a_1. □

Example 5 shows that although regulation benefits the consumer, this benefit is obtained in a very particular way. Namely, the consumer obtains a payoff that is higher than the default payoff only when the lowest-indexed offer (above the default) is made. The next proposition shows that this observation holds generally.

Proposition 5. *Suppose that the firm prefers higher-indexed to lower-indexed offers in all states of the world. Under any incentive-compatible direct privacy regime, the consumer is indifferent*

[21] Because the firm can impose the default offer on the consumer, no loss of generality is implied by assuming that the default offer is the worse one for the firm.

between the default offer and any offer $a_i \in A$ that is not the lowest-indexed offer that is made under that regime.

The limitation of the designer's ability to assist the consumer stems from the special form of the firm's incentive compatibility constraints when firm's preferences are state-independent. In particular, condition (3) requires that if the firm is instructed to offer both a_i and a_j, where $j > i$, then it must be the case that the consumer is indifferent between accepting and rejecting offer a_j.

Proposition 5 implies that when the firm's preferences are independent of the state of the world, the designer must select exactly one offer from which the consumer will benefit and the states in which that offer is made, and then allocate the remaining states to higher-indexed offers in a way that satisfies the incentive compatibility constraints. This insight not only simplifies the computation of optimal privacy regimes, but also justifies thinking about the designer's problem in the following way. In order to generate surplus for the consumer after offer a_i, the designer is willing to treat all signals $s = a_j$ for $j > i$ as "informational loss leaders." That is, the designer forgoes the opportunity to generate surplus for the consumer after signals $s = a_j$ for every $j > i$ in order to generate surplus for the consumer after signal $s = a_i$.

If the firm can make only one non-default offer, then Proposition 5 becomes starker and implies that regulation is futile. If the consumer is willing to accept offer a_1 under the prior, then the firm will offer a_1 in all states regardless of the choice of information structure. If offer a_1 is not made with probability one, then the firm's incentive compatibility constraint implies that the consumer must be indifferent between accepting and rejecting offer a_1. Thus, in such cases the consumer's expected utility is exactly the same as his utility from taking his outside option. Hence, regulation does not help the consumer.

Corollary 1. *Suppose that $N = 1$. If the firm prefers a_1 to a_0 in every state of the world, then all incentive-compatible privacy regimes provide the consumer with the same expected payoff.*

6 Extensions

6.1 General Designer's Preferences

Our baseline model focuses on the case in which the objective of the designer is to maximize the consumer's ex-ante utility. However, our main result, namely, the geometric characterization of the optimal privacy regime (Theorem 3), remains valid if the designer's objective is to maximize any weighted sum of the consumer's and firm's ex-ante payoffs. Hence, our insights into the design of optimal privacy regimes are applicable also in wider settings.

To see this, note that the incentive compatibility constraints of both the consumer and the firm (equations (2), (3), and (5)) are independent of the designer's preferences. Hence, the feasible set of direct privacy regimes induced by these constraints is also independent of the designer's objective function. Moreover, the transformations of these constraints used to derive the Lagrangian L in the proof of Theorem 3 are also independent of the designer's preferences.

The only place where the designer's preferences do play a role is in the first term of the Lagrangian L. Under the baseline model, the designer's direct value in the Lagrangian is

$$\sum_{a_i \in A^*} \sum_{\theta \in \Theta} \pi(\theta) p(a_i, \theta) u(a_i, \theta),$$

which implies that the derivative of the direct value with regard to $p(a_i, \theta)$ is a linear function of $W(\theta)$, namely, $\pi(\theta) u(a_i, \theta)$. If the designer's objective is to maximize the weighted sum of the consumer's and firm's payoffs

$$\sum_{a_i \in A^*} \sum_{\theta \in \Theta} \pi(\theta) p(a_i, \theta) \left(\beta u(a_i, \theta) + (1 - \beta) v(a_i, \theta) \right),$$

for some $\beta \in [0, 1]$, then the derivative of the direct value with regard to $p(a_i, \theta)$ still remains a linear function of $W(\theta)$. Specifically, it is equal to $\pi(\theta) \left(\beta u(a_i, \theta) + (1 - \beta) v(a_i, \theta) \right)$. Since Theorem 3 is a consequence of the linearity of the derivatives of the Lagrangian with respect to the probabilities $p(a_i, \theta)$, it will hold as long as the designer's preferences are a linear combination of $u(\cdot, \cdot)$ and $v(\cdot, \cdot)$.

6.2 Persuasion

We assumed that the firm makes a take-it-or-leave-it offer to the consumer. This implies that the firm can prevent the consumer from enjoying any non-null offer. However, in some settings, the firm may be unable to prevent the consumer from choosing any offer in set A, and must resort to influencing the consumer's choice via (Bayesian) persuasion. Our baseline model and its "persuasion variant" in which the consumer can choose any offer in set A following the firm's offer differ in the incentive compatibility constraints of both the consumer and the firm. In particular, when the consumer can select any offer in set A, the incentive compatibility constraint (2) is replaced by a constraint that stipulates that a_i is a best response to beliefs μ_{a_i}. That is, rather than requiring that the consumer's beliefs be above a certain hyperplane, the consumer's incentive compatibility constraints now require that the consumer's beliefs be in a polygon.[22] This more complicated struc-

[22] This is because whatever is chosen by the consumer has to be better than all the other options in set A, and comparison with every such option implies that the chosen option lies above some hyperplane.

ture of the consumer's incentive compatibility constraints can significantly complicate the firm's incentive compatibility constraints.

The designer's problem in the persuasion variant of our model with two possible offers ($N = 1$) is analyzed in Ichihashi (2019). Ichihashi characterizes the optimal privacy regime (the designer's strategy, in his terminology) implicitly as a function of the solution of an auxiliary persuasion game, and derives conditions under which the designer can help the consumer. In the case of $N = 1$, it is straightforward to adapt our techniques to the persuasion setting and augment Ichihashi's results by providing a direct geometric characterization of the optimal privacy regime. Namely, it is possible to show that Theorem 1 generalizes to this setting.

Even though the geometric structure of the optimal privacy regime is the same in the baseline model and the persuasion variant, the optimal privacy regimes in these two settings are generally different. Moreover, the consumer (weakly) prefers the optimal privacy regime in the persuasion variant to the optimal privacy regime in the baseline model.

Proposition 6. *Suppose that $N = 1$. The optimal privacy regime in the persuasion variant of the model is either a half-plane regime or it induces only one offer. Moreover, the consumer's expected payoff in the persuasion variant is weakly larger than his payoff in the baseline model.*

More generally, when $N > 1$, allowing the consumer to choose any offer in the set A has a big effect on the structure of the incentive compatibility constraints. In particular, the incentive compatibility constraints will depend not only on the angle between two beliefs, but also on the consumer's preferences along the line that connects these two beliefs. Moreover, the fact that the consumer's incentive compatibility constraints are given by polygons rather than hyperplanes implies that beliefs may lie on a kink of the set of incentive-compatible beliefs. Nevertheless, because our result relies on the linearity of the objective function and of the incentive compatibility constraints in the induced probabilities around the optimal privacy regime, which is also a feature of the persuasion variant of the model, we conjecture that our main results continue to hold in this case as well, even though a more complicated proof would be needed to establish these results.

6.3 Other Applications

6.3.1 Market Segmentation

Our framework can also be applied to the study of market segmentation. Consider a monopolistic firm that has the ability to engage in third-degree price discrimination. That is, the monopolist may be able to charge different prices from different sets of consumers.

Clearly, the ability of the monopolist to segment the market depends on the information it holds regarding consumers' characteristics.

Bergemann, Brooks and Morris (2015) identify the set of firm and consumer surplus pairs that result from all the possible segmentations of the market (under quasilinear preferences for the buyer and seller) when the firm sells a single good. Their results imply that any inefficient market can be segmented in a way that is Pareto improving for both the consumers and the firm.

Ichihashi (2020) and Hidir and Vellodi (2021) conduct welfare analysis when the monopolist is a multi-product firm that offers a single product to each consumer. Ichihashi (2020) considers a model where, *before* learning their types, consumers can decide whether or not to disclose information about their types to the firm, which then segments consumers into different markets based on the information it receives. He shows that it is profitable for the firm to commit not to use this information in order to price discriminate, and that such commitment reduces consumer welfare. By contrast, Hidir and Vellodi (2021) assume that *after* consumers learn their types, they can voluntarily disclose information to the firm. They show that the consumer-optimal segmentation is the least informative segmentation under which trade occurs with probability one.

Haghpanah and Siegel (2020) consider a more general model in which a multi-product monopolistic firm is also able to offer possibly different quantity/price menus of products and product bundles in different market segments. That is, in Haghpanah and Siegel's model, the firm is able to also engage in second-degree, on top of third-degree, price discrimination. The main result of Haghpanah and Siegel is that any inefficient non-segmented market can be generically segmented in a way that is Pareto improving for both the consumers and the firm. As noted by Haghpanah and Siegel, the main difficulty in applying Bayesian persuasion methods to the problem of market segmentation is that doing so requires that the firm's payoff be specified for any induced posterior belief or possible market segmentation; however, no characterization of the optimal menu for a multi-product monopolistic firm exists.

Our model can be recast to fit the scenario of market segmentation. In particular, our results can be used to derive the consumer optimal segmentation in markets where the firm's payoff from a given offer depends on the consumer's type. Such dependence arises naturally in insurance markets and markets for medical (or other types of) services.

Comparing our results to those of Bergemann, Brooks and Morris (2015) and Hidir and Vellodi (2021) highlights the effect of the fundamental difference between our two respective settings: in their settings, consumers know their types and the consumer-optimal market segmentation scheme is ex-post Pareto efficient. By contrast, in our setting consumers do not know their type and the optimal privacy regime may induce ex-post

Pareto-inefficient outcomes.[23]

To better understand the connection between the consumer's knowledge of his own type and the efficiency properties of optimal segmentation, consider the consumer's response to receiving a certain offer. In Bergemann, Brooks and Morris (2015) and Hidir and Vellodi (2021) consumers know their type, and so the firm's offer has no effect on consumers' inferences. It follows that the firm optimizes separately within each segment, and so the designer need not consider possible spillovers between different segments when he designs the optimal segmentation scheme. These papers show that a sophisticated form of market segmentation enables consumers to reap all the benefits from trade, and so the consumer-optimal segmentation is Pareto efficient.

By contrast, in our model, after receiving an offer the consumer makes an inference about his own type based on the firm's offer and general strategy. Thus, were the firm to make an identical offer to consumers in two different market segments, this would affect the consumers' inferences about their types in both of these segments. It follows that the firm may gain from merging two distinct market segments, and so the designer must consider the potential spillover between the different segments. Consequently, consumer-optimal segmentation may necessitate inefficient outcomes in some segments in order to limit the firm's ability to profit from merging segments.

6.3.2 Optimal Design of Legislative Procedures

Another application of our model is the design of efficient procedures within legislative bodies. The interaction between the consumer and the firm in our model is similar to the interaction between a chamber of a legislative body and a committee that proposes amendments under a restrictive procedure. The committee acquires information about the topic at hand, drafts a proposed amendment, and then the entire chamber decides whether to accept or reject the committee's proposal.[24] The committee may generally have different preferences than the chamber, and so it may attempt to exploit its informational advantage to its benefit. Our results show that regulation of the committee's

[23] Ichihashi (2020) considers a more complex setting where consumers learn their type between the time at which they disclose information and the time at which they decide whether to purchase the good. Similarly to us, he finds that consumer-optimal market segmentation is inefficient; however, his finding is due to a different mechanism. In his setting, the seller first commits to the prices at which he will sell each good, and then, after receiving information about the consumer's type, offers to sell the good that maximizes his expected revenue. Ichihashi shows that, in response, the consumer will disclose what is the good to which he assigns the maximal value, which, in turn, leads the seller to commit to a high price level that may prevent trade.

[24] See Gilligan and Krehbiel (1987) for a description of parliaments where such restrictive procedures are used, and an analysis of equilibrium outcomes under such procedures.

ability to acquire information will generally be beneficial. Moreover, our results provide a guideline for the geometric structure of the restrictions that should be considered.

7 Concluding Remarks

We study the optimal design of privacy in a setting where an informed firm makes a take-it-or-leave-it offer to an uninformed consumer. In particular, we are interested in settings where there is a real trade-off between providing the firm with enough information to facilitate welfare-enhancing trade and preventing it from extracting the surplus that is generated (e.g., medical insurance markets and markets for credence goods). Our main result provides a guideline for the design of optimal privacy regimes in such settings by providing a simple geometric condition that determines whether or not the firm should be allowed to distinguish between two consumer types.

More generally, our results also highlight the trade-offs that underlie optimal privacy laws and the potential benefit thereof. First, we show that when consumers do not know their type, then the optimal privacy regime may induce ex-post Pareto-inefficient outcomes in order to restrict the firm's ability to influence a consumer's inference about his own type. As explained in Section 6.3.1 above, this effect has not been addressed in previous work. Moreover, this suggests that the design of optimal privacy laws in markets where consumers do know their type is an inherently different problem than the design of privacy laws in markets where consumers do not know their type (and it is impossible to provide them with this information).

Second, Proposition 5 suggests that, when consumers do not know their types, if the firm's gain from trade does not depend on the consumer's type, then the design of privacy laws has a limited impact on consumer surplus. This suggests that in contexts in which the firm's payoff depends on the price of the transactions, but, conditional on the transaction occurring, is independent of the consumer's type (e.g., retail markets), privacy regulation is of limited value. By contrast, in contexts in which the firm's payoff does depend on the consumer's type (e.g., health care markets), privacy regulation can significantly affect consumer surplus. This comparison hints that privacy laws designed to regulate medical information (e.g., GINA) are likely to have a stronger impact on consumer welfare than laws designed to regulate retail markets (which is a primary goal of the General Data Protection Regulation).

Appendix: Proofs

Proof of Lemma 1. We establish this lemma in two steps. First, we show that if the third quadrant of the payoff space is nonempty, then a_0 is offered (with some probability) in any optimal privacy regime. Second, we show that if the first quadrant of the payoff space is nonempty, then there exists an optimal privacy regime in which a_1 is offered.

Suppose by way of contradiction that $\tilde{\theta}$ belongs to the third quadrant and a_1 is offered with probability one under an optimal privacy regime. Observe that moving $\tilde{\theta}$ to the rejection region increases the consumer's expected utility. Moreover, this modification relaxes the consumer's incentive compatibility constraint after the consumer receives offer a_1. Now, consider the impact of this change on the firm's incentive compatibility constraints. Since $\tilde{\theta}$ belongs to the third quadrant, in the modified rejection region the firm strictly prefers a_0 to a_1. By assumption, under the prior the firm prefers a_1 to a_0 (were this not the case, offering a_1 with probability one would not be an optimal strategy for the firm), and so conditional on $\theta \in \Theta \backslash \{\tilde{\theta}\}$ the firm also prefers a_1 to a_0. Hence, moving $\tilde{\theta}$ to the rejection region increases the payoffs of both players and is incentive compatible for both players.

Next, assume that there exists a state $\tilde{\theta}$ in the first quadrant and that the designer selects $S = \{s_0, s_1\}$, where s_1 is realized if and only if the state is $\tilde{\theta}$. After s_1 the consumer strictly prefers a_1 to a_0. Hence, offering a_i after s_i is incentive compatible for the consumer. Moreover, since after s_1 the firm also strictly prefers a_1 to a_0, the strategy of offering a_0 with probability one is suboptimal for the firm under this privacy regime. Because the firm faces a linear optimization problem, it has a best response to this privacy regime. Furthermore, under this best response it must make both offers. Since the consumer's expected utility against this strategy is at least as much as the utility from of receiving the default offer with certainty, it follows that having the firm make the default offer is not the unique optimal privacy regime. □

Proof of Theorem 1. By Lemma 1, under any optimal privacy regime the set of offers that are made is $A^* = \{a_0, a_1\}$. The more general Theorem 3 shows that in an optimal privacy regime the acceptance and rejection regions are separated by a line that passes through the origin. The discussion preceding the statement of Theorem 1 implies that this line is downward sloping. □

Proof of Theorem 2. First, we show that any outcome obtained under an arbitrary privacy regime can be replicated by a direct privacy regime.

Consider an arbitrary privacy regime $\langle \hat{S}, \hat{p}(\cdot|\cdot)\rangle$ and an offer strategy $\hat{q}(\cdot|\cdot)$ that is incentive compatible for the consumer and an optimal choice for the firm. Define a direct privacy regime that generates the same joint distribution of actions and states as follows:

$$\tilde{S} = \{a_i : \hat{q}(a_i|s) > 0 \text{ for some } s \in \hat{S}\}$$
$$\tilde{p}(a|\theta) = \sum_{s\in\hat{S}} \hat{q}(a|s)\hat{p}(s|\theta).$$

Note that the consumer's interim belief conditional on any offer made by the firm is the same under both regimes. Hence, this direct privacy regime is incentive compatible for the consumer. Moreover, selecting $q(a_i, s = a_i) = 1$ for all $a_i \in \tilde{S}$ is an optimal choice for the firm. To see this, assume by way of contradiction that there exists an offer strategy $\tilde{q}(\cdot|\cdot)$ that the firm strictly prefers to following the designer's instructions. This implies that under the privacy regime $\langle \hat{S}, \hat{p}(\cdot|\cdot)\rangle$, the firm strictly prefers the offer strategy

$$q(\cdot|s) = \tilde{q}(\cdot|a)\hat{q}(a|s) \quad \forall s \in \hat{S}$$

to $\hat{q}(\cdot|\cdot)$, a contradiction.

Next, we show that if a direct privacy regime satisfies the incentive compatibility constraints (2), (3), and (5), then it is incentive compatible. The fact that such a regime is incentive compatible for the consumer is immediate. Thus, assume to the contrary that there exists a direct privacy regime $\langle \tilde{S}, \tilde{p}(\cdot|\cdot)\rangle$ that satisfies the above constraints and for which the offer strategy $q(a_i|s = a_i) = 1$ for every $s \in \tilde{S}$ is not an optimal strategy for the firm. As the firm faces a linear optimization problem with at least one feasible solution, namely, making the default offer with probability one, it has an optimal strategy. Denote one such optimal strategy by $\tilde{q}(\cdot, \cdot)$.

Since this offer strategy is better for the firm than the strategy of following the designer's suggestions, there exist $a_j \in A$ and $a_i \in \tilde{S}$ for which both $\tilde{q}(a_j|s = a_i) > 0$ and $a_j \in D(a_i)$. Fix such an offer a_j, and let $B_D \equiv \{a_i \in \tilde{S}\backslash\{a_j\}|\tilde{q}(a_j|s = a_i) > 0$ and $a_j \in D(a_i)\}$. Since the privacy regime satisfies the incentive compatibility constraint (3), it must be the case that $\mu_{a_j} \in \partial\Xi_{a_j}$ and $\mu_{a_i} \notin \Xi_{a_j}$ for every $a_i \in B_D$. Since the offer strategy $\tilde{q}(\cdot, \cdot)$ is incentive compatible for the consumer, it follows that the set $B_T \equiv \{s \in T(a_i)|\tilde{q}(a_i|s) > 0\}$ is nonempty.

Select an arbitrary $a_i \in B_D$ and $a_k \in B_T$, and let $\epsilon_i, \epsilon_k > 0$ be such that

$$\epsilon_i p_{a_i}\mu_{a_i} \cdot u(a_j, \overline{\theta}) + \epsilon_k \mu_{a_k} \cdot u(a_j, \overline{\theta}) = 0.$$

Since the privacy regime satisfies incentive compatibility constraint (5), it follows that

$$\epsilon_i p_{a_i}\mu_{a_i} \cdot v(a_i, \overline{\theta}) + \epsilon_k \mu_{a_k} \cdot v(a_k, \overline{\theta}) \geq (\epsilon_i p_{a_i}\mu_{a_i} + \epsilon_k p_{a_k}\mu_{a_k}) \cdot v(a_j, \overline{\theta}).$$

If the inequality is strict, then it is both incentive compatible for the consumer and strictly profitable for the firm to increase $q(a_k|s = a_k)$ and $q(a_i|s = a_i)$ by ϵ_k and ϵ_i, respectively, and decrease $q(a_j|s = a_k)$ and $q(a_j|s = a_i)$ by ϵ_k and ϵ_i, respectively. Otherwise, since either $\epsilon_i = \tilde{q}(a_j|s = a_i)$ or $\epsilon_k = \tilde{q}(a_j|s = a_k)$ satisfies the above requirements, the firm's strategy provides it with the same profit as one in which there is one less signal after which the firm does not follow the designer's suggestion. Since $\tilde{S}$ is finite, an iterative application of the above argument implies that either $\tilde{q}(\cdot,\cdot)$ is suboptimal for the firm or that it is equivalent to following the designer's instructions, a contradiction. □

Proof of Theorem 3. Assume that there exists an optimal direct privacy regime that induces the offers $A^* \subset A$ via beliefs $\mu^*_{a_i}$ that are induced with probabilities $p^*_{a_i}$. Denote the offers that the consumer is indifferent to taking in this solution by

$$\overline{A} \equiv \{a_i \in A^* : \mu_{a_i} \in \partial\Xi_{a_i}\}.$$

Then define the constraint

$$\mu_{a_i} \in \partial\Xi_{a_i} \quad \forall a_i \in \overline{A}. \tag{6}$$

Since $\langle A^*, \{p^*_{a_i}, \mu^*_{a_i}\}\rangle$ is an optimal direct privacy regime, it follows that the solution to the designer's problem, subject to the consumer's incentive compatibility constraint (2) and the firm's incentive compatibility constraints (3) and (5), is also a local maximum to the designer's problem under constraints (6) and (5). Similarly, let $\overline{B} \subset A^* \times A^* \times A^*$ denote the set of triplets for which (5) is binding under $\{p^*_{a_i}, \mu^*_{a_i}\}$, i.e.,

$$\overline{B} \equiv \{(a_i, a_j, a_k) \in A^* \times A^* \times A^* : \mu_{a_k} \cdot (\overline{v}(a_k) - \overline{v}(a_i)) = \frac{\mu_{a_k} \cdot u(a_i, \overline{\theta})}{\mu_{a_j} \cdot u(a_i, \overline{\theta})} \mu_{a_j} \cdot (v(a_j, \overline{\theta}) - v(a_i, \overline{\theta}))\}. \tag{7}$$

Since $\langle A^*, \{p^*_{a_i}, \mu^*_{a_i}\}\rangle$ is an optimal direct privacy regime, it follows that any local maximum to the designer's problem with constraints (6) and (5) is also a local maximum under constraints (6) and (7). Because (6) are (7) equalities (rather than inequalities), the local maxima of this problem are given by the solution to the following Lagrangian:

$$\begin{aligned} L_1 = &\sum_{a_i \in A^*} p_{a_i}\Big(\mu_{a_i} \cdot u(a_i, \overline{\theta})\Big) + \sum_{i \in \overline{A}} \gamma_i \Big(\mu_{a_i} \cdot u(a_i, \overline{\theta})\Big) + \\ &\sum_{(a_i, a_j, a_k) \in \overline{B}} \gamma_{i,j,k} \Big(\mu_{a_k} \cdot (v(a_k, \overline{\theta}) - v(a_i, \overline{\theta}))(\mu_{a_j} \cdot u(a_i, \overline{\theta})) - \mu_{a_k} \cdot u(a_i, \overline{\theta}) \mu_{a_j} \cdot v(a_j, \overline{\theta}) - v(a_i, \overline{\theta}))\Big) \\ &+ \overline{\gamma} \cdot \left(\sum_{a_i \in A^*} p_{a_i} \mu_{a_i} - \pi(\overline{\theta})\right). \end{aligned}$$

Recall that $\{p^*_{a_i}, \mu^*_{a_i}\}$ is generated by some optimal direct privacy regime, and so we can rewrite the Lagrangian in terms of the signals $p(a_i|\theta_l), S$, where $S = A^*$ and the firm uses the offer function given by $R(s,s) = 1$; that is, the firm makes the offer that corresponds to the signal that it receives. Doing so gives the following Lagrangian:

$$
\begin{aligned}
L = & \sum_{a_i \in A^*} \sum_{\theta \in \Theta} \pi(\theta) p(a_i,\theta) u(a_i,\theta) \\
& + \sum_{a_i \in \overline{A}} \lambda_i \Big\{ \sum_{\theta} \pi(\theta) p(a_i,\theta) u(a_i,\theta) \Big\} \\
& + \sum_{(a_i,a_j,a_k) \in \overline{B}} \lambda_{i,j,k} \Bigg\{ \left(\sum_{\theta \in \Theta} \pi(\theta) p(a_k,\theta)(v(a_k,\theta) - v(a_i,\theta)) \right) \left(\sum_{\theta} \pi(\theta) p(a_j,\theta) u(a_i,\theta) \right) \\
& - \left(\sum_{\theta \in \Theta} \pi(\theta) p(a_k,\theta) u(a_i,\theta) \right) \left(\sum_{\theta \in \Theta} \pi(\theta) p(a_j,\theta)(v(a_j,\theta) - v(a_i,\theta) \right) \Bigg\} \\
& + \sum_{\theta \in \Theta} \lambda_\theta \left\{ \sum_{a_i \in A^*} p(a_i,\theta) - 1 \right\}.
\end{aligned}
$$

The derivative of L with respect to $p(a_l,\theta)$ is

$$
\begin{aligned}
\frac{\partial L}{\partial p(a_l,\theta)} = & \pi(\theta) u(a_l,\theta) + \lambda_i \pi(\theta) u(a_l,\theta) + \lambda_\theta \\
& + \sum_{a_i,a_k : (a_i,a_l,a_k) \in \bar{B}} \lambda_{i,l,k} \Bigg\{ \left(\sum_{\theta' \in \Theta} \pi(\theta') p(a_k,\theta')(v(a_k,\theta') - v(a_i,\theta')) \right) (\pi(\theta) u(a_i,\theta)) \\
& - \left(\sum_{\theta' \in \Theta} \pi(\theta') p(a_k,\theta') u(a_i,\theta') \right) (\pi(\theta)(v(a_l,\theta) - v(a_i,\theta)) \Bigg\} \\
& + \sum_{a_i,a_j : (a_i,a_j,a_l) \in \bar{B}} \lambda_{i,j,l} \Bigg\{ \pi(\theta) \left((a_l,\theta) - v(a_i,\theta)\right) \left(\sum_{\theta' \in \Theta} \pi(\theta') p(a_j,\theta') u(a_i,\theta') \right) \\
& - \pi(\theta) u(a_i,\theta) \left(\sum_{\theta' \in \Theta} \pi(\theta') p(a_j,\theta')(v(a_j,\theta') - v(a_i,\theta') \right) \Bigg\}.
\end{aligned}
$$

Denote by λ^* the value of the Lagrange multipliers in the solution. Evaluating this derivative for the optimal privacy regime gives

$$
\begin{aligned}
& u(a_l,\theta) + \quad \lambda_i^* \quad u(a_l,\theta) + \lambda_\theta^* \\
& + \sum_{a_i,a_k:(a_i,a_l,a_k)\in\bar{B}} \lambda_{i,l,k}^* \Big\{ p_{a_k}^* \mu_{a_k}^* \cdot (v(a_k,\overline{\theta}) - v(a_i,\overline{\theta})) u(a_i,\theta) \\
& \qquad - \left(p_{a_k}^* \mu_{a_k}^* \cdot u(a_i,\overline{\theta}) \right) ((v(a_l,\theta) - v(a_i,\theta)) \Big\} \\
& + \sum_{a_i,a_j:(a_i,a_j,a_l)\in\bar{B}} \lambda_{i,j,l}^* \Big\{ ((v(a_l,\theta) - v(a_i,\theta))) \left(p_{a_j}^* \mu_{a_j}^* \cdot u(a_i,\overline{\theta}) \right) \\
& \qquad - (u(a_i,\theta)) \left(p_{a_j}^* \mu_{a_j}^* \cdot (v(a_j,\overline{\theta}) - v(a_i,\overline{\theta}) \right) \Big\}.
\end{aligned}
$$

Observe that all the dot products do not depend on $W(\theta)$, and so this derivative can be expressed as $z_l \cdot W(\theta) + \lambda_\theta^*$ for some $z_{a_l} \in \mathbb{R}^{2N}$. Observe that in the solution to the designer's problem, it holds that $p(a_l,\theta) > 0$ only if $\frac{\partial L}{\partial p(a_l,\theta)} \geq 0$. Hence, if there exists a θ for which offers a_l and a_k are made with a (strictly) positive probability, then $\frac{\partial L}{\partial p(a_l,\theta)} = \frac{\partial L}{\partial p(a_k,\theta)} = 0$, and so $z_{a_l} \cdot W(\theta) = z_{a_k} \cdot W(\theta) = -\lambda_\theta^*$. Note that the set $\{W \in \mathbb{R}^{2N} : z_{a_l} \cdot W = z_{a_k} \cdot W\}$ is given by a hyperplane $Z_{a_l a_k}$ that contains the origin.

Moreover, for every state θ such that $W(\theta)$ is above $Z_{a_l a_k}$, we have that $\frac{\partial L}{\partial p(a_l,\theta)} > \frac{\partial L}{\partial p(a_k,\theta)}$. Therefore, it follows that if θ is below $Z_{a_l a_k}$ (for some $a_k \in A^*$), then $p(a_l,\theta) = 0$. Hence, the set of states in which offer a_l may be made is contained in the set $\mathcal{A}_l = \cup_{a_k \neq a_l} Y^+_{a_l a_k}$, where $Y^+_{a_l a_k}$ is the half-space above the hyperplane $Z_{a_l a_k}$. Moreover, $p(a_l,\theta) = 1$ if $W(\theta) \in int(\mathcal{A}_{a_l})$. Thus, the set of states in which offer a_l is made is contained in a polyhedral cone. Note that the union $\cup_{a_i:a_i\in A^*} \mathcal{A}_{a_i}$ must cover the entire payoff space. If there exists a payoff vector $W(\theta)$ that is not included in any $\mathcal{A}_{a_i}$, then an optimal offer does not exist for that payoff vector under the utility function given by the Lagrangian L. Since the set of offers is finite, this leads to a contradiction. □

***Proof of Proposition 4**.* The assumptions that $U(\cdot)$ is continuous, monotone, and that $U(0) > 0$ jointly imply that there are individual-types in the first quadrant of the payoff space and no individual-types in the third quadrant of the payoff space. This, in turn, implies that some individual-types receive insurance under an optimal privacy regime (see proof of Lemma 1). The assumption that $\mathbb{E}[\theta] < 0$ implies that under an incentive-compatible direct privacy regime some individual-types do not receive insurance. Hence, Theorem 1 implies that there exists $\alpha^* > 0$ such that (i) all individual-types in the first quadrant of the payoff space receive insurance, (ii) individual-types in the second quadrant of the payoff space receive insurance if and only if $\alpha(\theta) \equiv \left|\frac{U(\theta)}{\theta}\right| \geq \alpha^*$, and (iii) individual-types in the fourth quadrant of the payoff space receive insurance if and only if $\alpha(\theta) \leq \alpha^*$.

We start by showing that every individual-type for which $\theta \geq 0$ receives insurance under the optimal privacy regime. If the fourth quadrant of the payoff space is empty this is immediate. Otherwise, the assumptions that $U(\cdot)$ is convex and that $U(0) > 0$ jointly imply that the line with slope $U'(0)$ that passes through the origin lies strictly below $U(\cdot)$. Thus, the downward-sloping line with the smallest slope (in absolute value) that passes through the origin and is a tangent to $U(\cdot)$ must be such that the tangency point occurs at $\tilde{\theta} > 0$. If $\alpha^* \leq \alpha(\tilde{\theta})$, then all individuals with $\theta < 0$ will receive insurance and the firm will obtain a negative profit from selling insurance. Hence, in an incentive-compatible privacy regime, $\alpha^* > \alpha(\tilde{\theta})$. That is, individuals of type $\theta \geq 0$ receive insurance under the optimal privacy regime.

Furthermore, if only individual-types for which $\theta \geq 0$ receive insurance, then the firm's profit is strictly positive and its incentive compatibility constraints are slack. Thus, in optimum, some individual-types for which $\theta < 0$ will receive insurance. Because $\lim_{\theta \to 0} \alpha(\theta) = \infty$ there exists $\epsilon > 0$ such that, in optimum, individuals of types $[-\epsilon, 0]$ will receive insurance.

To establish the second part of the proposition, observe that $\alpha(\theta)$ is U-shaped on $[\underline{\theta}, 0]$. To see this, note that $\frac{\partial \alpha(\theta)}{\partial \theta} = \frac{U(\theta) - \theta U'(\theta)}{\theta^2}$ and so θ is an extremum of $\alpha(\theta)$ only if $\theta U'(\theta) - U(\theta) = 0$. That is, the straight line with slope $U'(\theta)$ that passes through the origin is a tangent of $U(\cdot)$ at θ. As $U(\theta)$ is convex there is at most one such extreme point, and, moreover, such a point is a minimum of $\alpha(\cdot)$.

First, consider the case where $\alpha(\underline{\theta}) \leq \alpha(\hat{\theta})$. The U-shape of $\alpha(\cdot)$ on $[\underline{\theta}, 0]$ implies that $\alpha(\theta) \geq \alpha(\hat{\theta})$ if and only if $\theta \geq \hat{\theta}$. Thus, if the designer selects $\alpha^* = \alpha(\hat{\theta})$ the individuals who receive insurance are exactly those for whom $\theta \geq \hat{\theta}$. Note that for this choice the firm's expected profit is zero. Moreover, for a lower α the firm will earn a negative profit, while for a higher α the consumer's welfare will be less than for α^*. Hence, $\alpha^* = \alpha(\hat{\theta})$ is the solution to the designer's problem.

Second, consider the case where $\alpha(\underline{\theta}) > \alpha(\hat{\theta})$. The U-shape and continuity of $\alpha(\cdot)$ imply that there exists an $\epsilon > 0$ such that $\alpha(\underline{\theta}) > \alpha(\theta)$ for all $\theta \in (\underline{\theta}, \hat{\theta} + \epsilon)$. It follows that if α^* is such that type $\underline{\theta}$ does not receive insurance, then neither do any individuals for whom $\theta \leq \hat{\theta} + \epsilon$. This implies that the firm is making a strictly positive profit from selling insurance. Hence, the designer can increase α^* and provide insurance to more types that benefit from it. Therefore, an individual of type $\underline{\theta}$ must receive insurance under the solution to the designer problem. Because an individual of type $\overline{\theta}$ also receives insurance, but the firm does not provide insurance to all individuals, the set of individuals that receive insurance is not an interval.

□

Proof of Proposition 5. The proof of the proposition follows immediately from the assumption that the firm prefers higher-indexed offers to lower-indexed ones and the firm's incentive compatibility constraint (3). To see this, note that under the assumed preferences, $D(a_i) = \{a_j : j > i\}$. Consequently, condition (3) requires that for every $a_i, a_j \in S$ such that $j > i$ it holds that $\mu_{a_j} \in \partial \Xi_{a_j}$. □

Proof of Proposition 6. When $N = 1$ the consumer's incentive compatibility constraint in the persuasion problem states that $\mu_{a_1} \in \Xi_1$ and $\mu_{a_0} \notin int(\Xi_1)$. Moreover, for this case there are only simple deviations for the firm, and, more importantly, the firm's incentive compatibility constraints for simple deviations remain unchanged. To see this, note that in this case if $a_1(a_0)$ is not a best response to belief μ, then $a_0(a_1)$ is the unique best response to that belief. Thus, to prevent the firm from making offer $a_1(a_0)$ when it should make offer $a_0(a_1)$, it must be the case that $a_0(a_1)$ is the unique best response to belief $\mu_{a_0}(\mu_{a_1})$ and the consumer is indifferent between the two offers at belief $\mu_{a_1}(\mu_{a_0})$. This is exactly what is stated in incentive compatibility constraint (3). As explained in Section 6.1, the additional (linear) component of the consumer's incentive compatibility constraint does not impact the derivation of Theorem 3, and so our geometric characterization remains valid for the persuasion problem with two offers.

Note that if under the solution to the baseline problem (i.e., when the firm makes a take-it-or-leave-it offer) only one offer is made, then the second part of the proposition is trivially satisfied. Hence, we assume that under the optimal privacy regime both offers are made. To establish the second part of the proposition, we distinguish between the following two cases in the solution of the baseline problem: (i) the consumer weakly prefers a_0 to a_1 in $\mathcal{A}_{a_0}$, and (ii) the consumer strictly prefers a_1 to a_0 in $\mathcal{A}_{a_0}$.

Case (i)*:* Note that for any given privacy regime chosen by the designer, the firm must respect more stringent consumer incentive compatibility constraints in the persuasion problem than in the baseline problem. Moreover, in this case, the solution of the baseline problem satisfies the consumer's incentive compatibility constraints of the persuasion problem. Hence, the solution to the original problem is incentive compatible in the persuasion problem. It follows that in the persuasion problem the designer can ensure that the consumer receives the expected utility that he receives under the optimal privacy regime in the baseline problem.

Case (ii)*:* Since under an incentive-compatible privacy regime, the consumer must weakly prefer a_1 to a_0 in $\mathcal{A}_{a_1}$, in this case, the consumer prefers offer a_1 to a_0 under the prior. Furthermore, as the optimal privacy regime is incentive compatible for the firm, it must strictly prefer a_0 to a_1 in $\mathcal{A}_{a_0}$. Assume that the designer were to provide the firm with complete information about the state of the world. If the firm induces offer a_0 in the persuasion problem with a positive probability, it must be the case that the consumer weakly prefers a_0 to a_1 conditional on receiving that recommendation. Moreover, similar reasoning to that used to describe the intuition for Theorem 1 implies that the firm's strategy can be described as a half-plane regime. In this case, the half-plane regime that is the solution to the baseline problem does not satisfy the consumer's incentive compatibility constraints in the persuasion problem. Hence, to restore incentive compatibility, the firm must choose a half-plane regime with a milder slope. This, in turn, increases the consumer's expected utility. Clearly, the designer can only further increase the consumer's expected utility by optimizing over privacy regimes. □

References

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman. 2016. "The Economics of Privacy." *Journal of Economic Literature*, 54(2): 442–492.

Ali, Nageeb, Greg Lewis, and Susana Vasserman. 2020. "Voluntary Disclosure and Personalized Pricing." *Working Paper*.

Alonso, Ricardo, and Odilon Camara. 2016. "Persuading Voters." *American Economic Review*, 106(11): 3590–3605.

Bergemann, Dirk, and Stephen Morris. 2019. "Information Design: A Unified Perspective." *Journal of Economic Literature*, 57(1): 44–95.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2015. "The Limits of Price Discrimination." *American Economic Review*, 105(3): 921–957.

Darby, Michael R., and Edi Karni. 1973. "Free Competition and the Optimal Amount of Fraud." *Journal of Law and Economics*, 16(1): 67–88.

Doval, Laura, and Vasiliki Skreta. 2021. "Constrained Information Design." *Working Paper*.

Dulleck, Uwe, and Rudolf Kerschbamer. 2006. "On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods." *Journal of Economic Literature*, 44(1): 5–42.

Forges, Françoise, and Jérôme Renault. 2021. "Strategic Information Transmission with Sender's Approval." *International Journal of Game Theory*, 1–28.

Galperti, Simone, and Jacopo Perego. 2019. "Belief Meddling in Social Networks: An Information-Design Approach." *Working Paper*.

Garcia, Daniel, and Matan Tsur. 2021. "Information Design in Competitive Insurance Markets." *Journal of Economic Theory*, 191: 105160.

Gentzkow, Matthew, and Emir Kamenica. 2017. "Competition in Persuasion." *Review of Economic Studies*, 84(1): 300–322.

Gilligan, Thomas W., and Keith Krehbiel. 1987. "Collective Decisionmaking and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures." *Journal of Law, Economics, & Organization*, 3(2): 287–335.

Haghpanah, Nima, and Ron Siegel. 2020. "Pareto Improving Segmentation of Multiproduct Markets." *Working Paper*.

Handel, Ben, Igal Hendel, and Michael D. Whinston. 2015. "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk." *Econometrica*, 83(4): 1261–1313.

Hidir, Sinem, and Nikhil Vellodi. 2021. "Privacy, Personalization, and Price Discrimination." *Journal of the European Economic Association*, 19(2): 1342–1363.

Ichihashi, Shota. 2019. "Limiting Sender's Information in Bayesian Persuasion." *Games and Economic Behavior*, 117: 276–288.

Ichihashi, Shota. 2020. "Online Privacy and Information Disclosure by Consumers." *American Economic Review*, 110(2): 569–595.

Kamenica, Emir. 2019. "Bayesian Persuasion and Information Design." *Annual Review of Economics*, 11(1): 249–272.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review*, 101(6): 2590–2615.

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. 2017. "Persuasion of a Privately Informed Receiver." *Econometrica*, 85(6): 1949–1964.

Li, Fei, and Peter Norman. 2021. "Sequential Persuasion." *Theoretical Economics*, 16: 639–675.

Matyskova, Ludmila. 2018. "Bayesian Persuasion with Costly Information Acquisition." *Working Paper*.

Pram, Kym. 2020. "Disclosure, Welfare and Adverse Selection." *Working Paper*.

Roesler, Anne-Katrin, and Balázs Szentes. 2017. "Buyer-Optimal Learning and Monopoly Pricing." *American Economic Review*, 107(7): 2072–2080.

Varian, Hal R. 2009. "Economic Aspects of Personal Privacy." In *Internet Policy and Economics: Challenges and Perspectives*. , ed. William H. Lehr and Lorenzo Maria Pupillo, 101–109. Boston, MA: Springer.

Wei, Dong. 2020. "Persuasion under Costly Learning." *Journal of Mathematical Economics*, 102451.