Privacy for Sale: Anarchy, Ban, or Market?

Piotr Dworczak
Northwestern and GRAPE

Jacopo Perego
Columbia

Laura Veldkamp
Columbia and NBER

August 15, 2026

Abstract

Third-party data sharing is a widespread practice whereby firms acquire consumer data as a by-product of ordinary transactions and then sell it to other market participants, potentially imposing privacy costs on consumers. We study this practice in a framework with a monopolistic seller, a privately informed consumer, and a third party who chooses how intensively to utilize the data acquired from the seller. Our analysis highlights the role of the *privacy* β, which captures whether consumers with a higher willingness to pay for the seller's product are more ($\beta > 0$) or less ($\beta \leq 0$) averse to greater data use. When $\beta > 0$, third-party data sharing is coarse: the seller collects and sells at most one data record, pooling consumer types. In this case, greater data use lowers consumer rents. Conversely, when $\beta \leq 0$, third-party data sharing is granular: the seller collects and sells multiple records that reveal consumer types. In this case, greater data use increases consumer rents. These forces govern the welfare comparison among three policy regimes: data anarchy, a data ban, and a data market.

JEL Classification Numbers: D82, L86, D83

Keywords: Data privacy, data markets, mechanism design, information design, digital regulation.

1 Introduction

Third-party data sharing is a widespread feature of the digital economy. Firms acquire consumer data as a by-product of ordinary transactions and sell the resulting records to third parties, who use them for advertising, pricing, credit, insurance, and other decisions. These markets developed amid limited transparency about what was collected, where records traveled, and how they were ultimately used (FTC, 2014). These records are valuable, but their sale may impose privacy costs on consumers. When there are no enforceable limits on this practice, a seller that cannot credibly promise to withhold a valuable record will monetize it after the transaction. This state of *data anarchy* raises obvious concerns about consumer welfare.

A natural policy response is to prohibit third-party data sales. If consumers dislike having their data sold and used, a ban on data sale appears to remove the harm directly. An economist's response would be that this argument holds product prices fixed. If data sales can generate revenue, a seller may reduce prices to induce more purchases and acquire more data. A *data ban* may therefore harm consumers by making the product more expensive.

An alternative regulation—informed by basic economic intuition—is to create a market for data by giving consumers control over whether records generated by their transactions may be sold. In practice, such control need not require consumers to trade their records directly. It can instead take the form of enforceable promises by firms about whether a transaction record will be transferred to third parties. These promises allow a seller to offer both a private product and a version that permits data sale, effectively letting consumers exchange permission to use their data for a product discount. If prices adjust appropriately, such a *data market* could in principle compensate consumers for their loss of privacy, without suppressing a valuable transaction.

Yet the usual economic intuition in favor of market-based solutions is insufficient here. Product sellers have market power, consumers possess private information, and downstream data use is chosen by parties whose incentives do not reflect the harm they impose on consumers. Enforceable privacy promises create an additional screening instrument that a monopolistic seller can use to extract surplus. Does a data market empower consumers through choice and compensation, or enable the seller to charge a premium for privacy? Should policymakers permit data anarchy, impose data bans, or introduce laws that create a data market? Answering these questions requires a model in which

prices, disclosure, and downstream data use are jointly determined.

In this paper, we make two contributions. First, we develop a framework for studying third-party data sharing that draws on mechanism and information design, and use it to derive predictions about market outcomes under alternative policy regimes. Second, we compare welfare across data anarchy, a data ban, and a data market. The framework reveals a genuine policy puzzle: absent further structure, each of the three regimes may be best for *all* consumers. At the same time, our analysis identifies structural parameters and sufficient-statistics tests that sharpen welfare comparisons and, under interpretable conditions, deliver unambiguous policy rankings. The data market emerges as a robust, though imperfect, compromise: it weakly Pareto dominates either data anarchy or a data ban, and is never itself Pareto dominated when the welfare of all market participants is considered.

To explain these findings, we first describe our model, which features three players interacting over two stages. In the first stage, a monopolistic *seller* chooses a mechanism to sell a single product to a *consumer* who has private information about her valuation. The mechanism maps the consumer's messages into trade outcomes. If trade occurs, the message sent by the consumer becomes a *data record*. In the second stage, the seller may sell this record to a *third party*, such as an advertiser, at a take-it-or-leave-it price. If (and only if) the third party acquires the data record, she gains access to the consumer, updates beliefs about the consumer's type based on the equilibrium meaning of the record, and chooses a costly action, *data utilization*, such as targeted advertising intensity. The third party's optimal data utilization increases in the consumer's type, and imposes a cost on the consumer, a *privacy loss*, that increases with the intensity of data use. Importantly, the seller designs both the product mechanism and the information contained in the resulting records, but cannot dictate how the third party uses a record once acquired.

The three policy regimes differ in the seller's ability to commit to the treatment of these records. Under data anarchy, the seller cannot credibly commit to withhold a record after a product purchase; consumers must either purchase the product and forfeit privacy or not purchase at all. Under a data ban, third-party data sales are prohibited, and consumers never suffer a privacy loss. Under a data market, promises to sell or withhold records are enforceable. The seller can therefore unbundle purchase of the product from permission to sell the data it generates.

Each policy regime directly affects the extent to which consumer data is utilized by the third party. As emphasized above, the welfare consequences are mediated by equilibrium

price adjustments in the product market: the key question is whether price changes compensate consumers for changes in privacy losses. Tracking these price responses directly is difficult because, depending on the regime and model parameters, the seller may offer one, several, or even a continuum of differently priced options. Moreover, an observed price change can conflate two distinct forces: compensation for a change in privacy loss and an adjustment that expands or contracts the set of consumers who purchase. Our mechanism- and information-design approach allows us to disentangle these two forces. First, holding fixed the set of consumers who purchase the product, we ask how equilibrium data utilization affects their utilities; we call this the *data-compensation margin*. Second, we ask how the seller adjusts the set of consumers who purchase; we call this the *participation margin*.

This separation of forces brings into focus a key structural parameter: the *privacy* β, which captures the relationship between a consumer's willingness to pay for the product and her marginal privacy loss from data use. A positive privacy β means that higher-valuation consumers are more sensitive to data utilization; a negative privacy β means that they are less sensitive. Both cases arise naturally. Data-enabled price discrimination may expose high-valuation consumers to greater losses, producing a positive privacy β. Conversely, if low product valuations are associated with low income, data use that exposes poorer consumers to worse credit or insurance terms can produce a negative privacy β.

The data-compensation margin immediately reveals the importance of the privacy β. Holding the set of purchasers fixed, a pointwise increase in data utilization weakly lowers every consumer's utility when $\beta \geq 0$ and weakly raises it when $\beta \leq 0$. To see why, suppose that data utilization rises, and the seller lowers the product price to keep the marginal purchaser indifferent. Every inframarginal purchaser benefits from that price reduction but suffers a privacy loss. When $\beta > 0$, higher types suffer more from privacy loss than the marginal type and are therefore under-compensated. When $\beta < 0$, higher types suffer less and are over-compensated (Proposition 2).[1] Thus, holding participation fixed, consumers benefit from higher data utilization if and only if the privacy β is

[1]Equivalently, data utilization compresses the consumer rent profile when $\beta > 0$ and steepens it when $\beta < 0$. This intuition is reminiscent of the classic bundling logic in multiproduct monopoly: the correlation of valuations across goods determines how much surplus the monopolist can extract. In the symmetric two-good case, perfect negative correlation can make the bundle value constant across consumers, allowing full surplus extraction, whereas perfect positive correlation doubles the rent. See Adams and Yellen (1976) and McAfee, McMillan and Whinston (1989).

negative. This insight gives a first rough intuition for policy: since the three regimes can be ordered by the data utilization they induce, their welfare ranking along the data-compensation margin *reverses* with the sign of the privacy β.

This simple comparative-static logic is complicated by the participation margin: policy changes also affect the set of consumers purchasing the product through the seller's choice of the profit-maximizing mechanism. To understand these equilibrium adjustments, we must first characterize the structure of the market under each policy regime. Here, too, the sign of the privacy β plays a central role. When $\beta > 0$, consumer incentives and third-party incentives conflict. The third party wants data utilization to increase with the consumer type inferred from the record, while higher-type consumers have a stronger incentive to avoid data use. These forces can be reconciled only through coarse disclosure: all consumers whose data are sold must be pooled into a single record and face the same level of data utilization. When $\beta < 0$, by contrast, consumer and third-party incentives are aligned. Disclosure can then be granular, and in the seller-optimal mechanism each sold record fully reveals the consumer's type (Proposition 1).

The sign of the privacy β therefore governs two different economic objects: whether price adjustments compensate consumers for data use and whether market incentives produce coarse or granular disclosure.

When the privacy β is positive (Section 4), the seller sells the data of every purchaser in a single pooled record under data anarchy, while no data are sold under a ban. Under a data market, the seller generally offers two versions of the product: a cheaper version that permits data sale and a more expensive private version. Because data use lowers consumer rents through the data-compensation margin in this case, less disclosure protects consumers. If β is sufficiently small, the participation margin does not overturn this force, and the data market weakly Pareto dominates data anarchy (Proposition 4). The comparison with a data ban is distributional because the two margins work in opposite directions. On the one hand, the ban protects consumers from data use through the data-compensation margin, with the benefits concentrated among high-valuation consumers. On the other hand, the data market makes it profitable to serve lower-valuation consumers, expanding access through the participation margin. As a result, there is a cutoff below which consumers prefer the data market and above which they prefer the ban.

Granular disclosure and personalized pricing in the negative β case (Section 5) may look like the most exploitative outcome. Yet this is precisely the case in which consumers benefit from higher data utilization through the data-compensation margin. Under data

anarchy, every purchaser's type is revealed to the third party. A data market may instead give privacy to lower purchasing types while selling the fully revealing records of higher types, who receive type-dependent data discounts. Because price adjustments more than compensate inframarginal consumers for data use, a ban eliminates a source of consumer rents, and is therefore weakly Pareto dominated by a data market (Proposition 6). The comparison between the data market and data anarchy is again distributional. By separating product access from data sale, the data market can expand participation among lower-valuation consumers, while data anarchy gives high-valuation consumers the largest data-based discounts. There is a cutoff below which consumers prefer the data market and above which they prefer data anarchy.

The opening policy intuitions are therefore incomplete. Data anarchy cannot be dismissed solely because it permits pervasive data sales: when $\beta < 0$, the resulting discounts can benefit consumers, especially high-valuation ones. Nor is a ban necessarily an economically naive response: when $\beta > 0$, price adjustments may fail to compensate the consumers who suffer the greatest privacy losses. Selecting between these extremes requires knowing how privacy sensitivity covaries with product valuations and how regulation changes participation.

A data market is an imperfect but robust compromise. Its flexibility always weakly benefits the seller, who can reproduce the outcomes available under a ban or data anarchy. For small positive β, the data market Pareto dominates anarchy; for negative β, it Pareto dominates the ban. Whether the data market also weakly Pareto dominates the remaining regime can be determined through sufficient-statistics tests that we develop. With positive β, it weakly Pareto dominates the ban if and only if the price of its private option is no higher than the product price under the ban; with negative β, it weakly Pareto dominates data anarchy if and only if its cheapest product version is no more expensive than the cheapest version under data anarchy. When these conditions fail and the regimes cannot be Pareto ranked, lower-valuation consumers are the consistent beneficiaries of the data market's ability to separate product access from data sale. If willingness to pay for the product increases with income, then data markets—perhaps surprisingly—tend to benefit the poor.

Related Literature. We contribute to the growing literature on the economics of data and privacy.[2] We study a model in which data are collected as a by-product of a

[2]Acquisti, Taylor and Wagman (2016), Bergemann and Bonatti (2019), Bergemann and Ottaviani (2021), Goldfarb and Que (2023), Goldfarb and Tucker (2024), and Baley and Veldkamp (2025) provide

product transaction and then sold to a third-party data user. This framework allows us to analyze how product sales and data sales jointly shape the optimal selling mechanism and how third-party data sharing regulation affects market outcomes.

Several papers study related questions. Rayo (2013) and Krähmer and Strausz (2023) study environments in the spirit of Mussa and Rosen (1978) in which product differentiation is distorted because buyers care about the beliefs about their types induced by their purchase decisions. In both papers, the informational externality is intrinsic. In Rayo (2013), consumers value social status and therefore want to be perceived as high types; in Krähmer and Strausz (2023), some privacy-sensitive consumers incur a disutility whenever information about their type is revealed. In both models, the seller does not directly monetize consumer information by selling it to a downstream party. In Doval and Skreta (2026), the externality is instrumental and monetized: the consumer's purchase decision is observed by a downstream seller, who can use the implied information for price discrimination. While the basic tri-party structures of their model and ours are similar, the research questions are different: Doval and Skreta (2026) ask how the downstream data leakage distorts product line differentiation upstream, with the undistorted benchmark being the rich product line as in Mussa and Rosen (1978). Our setup is the opposite, with a simple posted price mechanism being the baseline, and data sharing introducing a potential need for versioning the product; our key notion of the privacy β has no parallel in Doval and Skreta's analysis.[3]

Bisceglia, Bonatti and Scott Morton (2025) is closely related, but differs along three main dimensions. First, they focus on *first-party data use*: firms collect data through their interactions with consumers and either use those data directly or fully contract over their downstream use. We instead study third-party data sharing, where the firm that collects the data neither uses them directly nor can fully control how they are subsequently used. Second, this distinction in the economic environment leads to a substantially different modeling approach. Our framework derives the third-party value and the use of data from a Bayesian information-design problem. Third, our papers emphasize different economic effects. Ours emphasizes how the sign of the privacy β shapes disclosure, the mechanism, and policy welfare; theirs studies how data practices can disadvantage naive consumers.

introductions to this literature.

[3]More distantly related, Lu (2025) studies a signaling problem in which consumers bear the cost of signaling and care about market beliefs because those beliefs affect wages.

Ichihashi, Jeon and Kim (2025) study a monopolist ad-supported platform that screens consumers through menus that jointly determine content access, advertising exposure, and price. Their model is richer than ours on the product-design side, since the platform allocates heterogeneous items and heterogeneous ads and chooses how to match them across consumers. By contrast, it is more reduced-form on the data-privacy side and features no explicit aftermarket. These modeling differences answer a different set of questions. Ali, Lewis and Vasserman (2023) study voluntary disclosure of verifiable information by consumers to a seller who subsequently personalizes prices. Their setting is further from ours, but shares the feature that both price adjustments and the granularity of disclosure shape consumer surplus.

More broadly, our paper contributes to the literature on mechanism design with aftermarkets. Calzolari and Pavan (2006) and Dworczak (2020) study general mechanism-design problems with aftermarkets and provide sufficient conditions under which privacy is optimal in single-agent settings. We do not impose those conditions here. Dworczak (2020) and Dworczak (2021) highlight how the distinction between submodular and supermodular aftermarkets—which resembles our distinction between positive and negative privacy β cases—shapes the form of the optimal mechanism and the desirability of privacy. Relatedly, Rivera Mora (2025) shows that, in a broader class of after-games, related submodularity and supermodularity conditions determine whether information sharing is feasible.

2 Model

Evaluating third-party data policy requires an environment in which product prices, information disclosure, and downstream data utilization are determined jointly. The seller designs the product mechanism and decides whether a transaction record is sold. But the seller's lack of ability to contract on the third party's data use makes the third party a meaningfully different actor and distinguishes this model from a two-party product-design problem in which privacy is simply another contractible product attribute. We then introduce the three policy regimes, which differ only in the seller's ability to sell or commit to withholding data records. The next section characterizes the feasible outcomes and shows how they depend on the sign of the privacy β.

2.1 Main Ingredients

Consider a game with three players: a consumer, a monopolistic seller, and a third party. The seller and the consumer first trade a product. Conditional on trade, the seller may sell to the third party the consumer data generated by the initial transaction. If the data is sold, the third party takes an action that affects both her own payoff and the consumer's payoff.

More formally, the seller has a single indivisible good for sale. The consumer's willingness to pay for this good is $\theta \in \Theta \triangleq [0,1]$, which we refer to as her private *type*. We assume θ is drawn from a distribution F with a density f that is Lipschitz continuous and strictly positive on Θ. Additionally, let the virtual value $\theta - (1 - F(\theta))/f(\theta)$ be strictly increasing.

In the first period, the seller publicly announces an indirect mechanism (x, t) that maps a message $m \in M$ sent by the consumer into a (deterministic) allocation of the product $x(m) \in \{0,1\}$ and a monetary transfer $t(m) \in \mathbb{R}$ paid by the consumer to the seller. The mechanism includes a non-participation message $m = \varnothing$, with $x(\varnothing) = t(\varnothing) = 0$. If the consumer chooses a message m such that $x(m) = 0$, then no transaction occurs, the game ends, and all payoffs are normalized to zero. If instead $x(m) = 1$, trade takes place and the game proceeds to the second period. In this case, we refer to the message m as the consumer's *data record*.

In the second period, the seller may sell record m to the third party through a take-it-or-leave-it offer at price $p(m)$. The third party can verify the value of the record, but she obtains access to the consumer only upon buying it. If record m is not sold, the third party cannot interact with the consumer, and the game ends. If it is sold, instead, the third party gets access to the consumer, updates her beliefs about her type, and chooses an action $a \in \mathbb{R}$, interpreted as the intensity of data utilization.

Next, we specify the payoffs of the three players. Fixing $(\tilde{x}, \tilde{t}, \tilde{a}, \tilde{p}) \in \{0,1\} \times \mathbb{R}^3$, the payoff of a consumer of type θ is

$$u(\theta \mid \tilde{x}, \tilde{t}, \tilde{a}) = \tilde{x} \cdot \big(\theta - \tilde{a} \cdot \ell(\theta)\big) - \tilde{t}. \tag{1}$$

The first term inside the parentheses, θ, captures the consumer's willingness to pay for the seller's product. The second term captures the privacy loss from data utilization at intensity $\tilde{a}$, where $\ell(\theta)$ denotes the consumer's marginal *privacy loss*, which we will discuss in Section 2.2. Note that the consumer's type θ jointly determines both her valuation for the product and her sensitivity to data utilization.

The seller's payoff is $\tilde{t} + \tilde{p}$, namely, the sum of the monetary transfer collected from the consumer and the price paid by the third party for the data record.

Finally, the third party's payoff from choosing intensity $\tilde{a}$ when the consumer's type is θ is $v(\theta, \tilde{a}) = 2\left(\theta\tilde{a} - \frac{1}{2}\tilde{a}^2\right)$.[4] Note that, for any record m, the third party's optimal action is $a^*(m) = \mathbb{E}[\theta \mid m]$. Thus, records that induce a higher posterior mean about the consumer's type lead to more intensive data utilization. Moreover, the value of record m to the third party is $\mathbb{E}\left[v(\theta, a^*(m)) \mid m\right] = a^*(m)^2$, and hence is always weakly positive.

In the game described above, a strategy for the seller consists of a choice of a first-period (deterministic) mechanism $(x, t) : M \to \{0, 1\} \times \mathbb{R}$—with the mechanism being publicly announced and committed to at the beginning of the game—and a second-period data-pricing strategy $p : M \to \mathbb{R}$. A consumer's strategy is a deterministic mapping from θ to a message in that mechanism. A strategy for the third party is to accept or reject the offer to acquire the data record m at price p from the seller, and—conditional on the purchase—to choose the intensity a as a function of m.

Our goal is to characterize the seller-optimal pure-strategy Perfect Bayesian equilibria (henceforth, equilibria) of this game and their welfare properties.[5]

2.2 The Privacy β

We assume that the privacy loss function ℓ is nonnegative and affine in θ:

$$\ell(\theta) \triangleq \alpha + \beta\theta \geq 0. \tag{2}$$

Intuitively, we can think of the affine structure in (2) as resulting from a linear regression of privacy concerns on willingness to pay for the good, θ. We will correspondingly refer to the slope of the "regression" as the *privacy* β.[6]

[4] The factor 2 is a convenient normalization that keeps the third party's payoff bounded between 0 and 1.

[5] Our focus on pure strategies is mainly for tractability and realism. In particular, we rule out any explicit randomization by the seller (see Dworczak (2020) for a discussion of how randomization complicates the analysis of mechanisms involving information disclosure). Online Appendix C.1 discusses in greater detail the underlying extensive form of our game and provides a formal definition of the equilibrium.

[6] Indeed, under this linear specification, $\beta = \frac{\mathrm{Cov}(\ell(\theta),\theta)}{\mathrm{Var}(\theta)}$, which is the coefficient of an OLS regression $\ell_i = \alpha + \beta\theta_i + \epsilon_i$. The consumer knows θ, but does not have private information about ϵ_i relative to what the seller can predict.

We will see that the sign of the privacy β plays a central role in shaping both the optimal mechanism and its welfare consequences. A *positive* privacy β indicates that consumers with a higher willingness to pay for the product also suffer a higher marginal disutility from data use; that is, they are more concerned about their privacy. A *negative* privacy β, instead, captures cases in which consumers with a higher willingness to pay for the product suffer a *lower* marginal disutility from data use; that is, they are less concerned about their privacy.

Note that a data record plays a dual role in our framework: it gives the third party access to the consumer *and* provides information about the consumer's type. In practice, a data record is valuable precisely because it combines identifiers that allow the third party to recognize or track the consumer (e.g., email address or device ID) with information about the consumer's preferences (e.g., the fact that she bought a particular product). When we refer to "consumer privacy," we primarily mean the former role: a consumer's privacy is compromised when the seller makes it possible for third parties to access her. The privacy β determines how the payoff consequences of privacy loss vary with the consumer's type. This perspective differs from approaches such as Eilat, Eliaz and Mu (2021), where privacy loss is modeled directly as a cost of third-party learning about the agent's type; in particular, a data record that is completely uninformative about the consumer's type still induces a privacy loss in our framework, as long as it is sold to the third party.

We discuss two simple examples to illustrate the concept of privacy and the economic interpretation of the privacy β.

Example 1 (*Advertising versus price discrimination*) Suppose that the third party sells a product that is complementary to the one sold by the seller. Sale of the consumer's data record to the third party enables her to target the consumer with customized ads. The consumer's value for the complementary product is 1 with probability θ and 0 with probability $1-\theta$. Thus, θ captures both the consumer's valuation for the seller's product and the probability that she needs the complementary product.[7] We consider two cases.

First, suppose that the price of the complementary product is fixed at $r < 1$, and interpret data utilization a as advertising intensity. Advertising makes the consumer aware of the complementary product with probability γa, for some given $\gamma \in (0,1)$. Advertising also imposes a nuisance cost αa. Then, the consumer's marginal loss from

[7] A similar structure, in which θ is both the value for the upstream good and the *probability* of having a high value for the downstream good, has been used by Dworczak (2021) and Doval and Skreta (2026).

data use is $\ell(\theta) = \alpha - \gamma(1-r)\theta$, so this example corresponds to our model with a negative privacy β: higher-type consumers have a lower marginal disutility from data use because advertising is more likely to steer them toward a product they need.

More generally, a negative privacy β does not require data use to benefit higher types. It can also arise when lower product valuations are associated with lower income and data use exposes poorer consumers to worse credit or insurance terms, so that lower types bear greater losses.

Second, suppose the consumer can buy the complementary product from the third party at price a or, with probability $1-\gamma$, from an offline vendor at a default price normalized to 0. Thus, a is a personalized surcharge. It is costly for the third party due to reputational or regulatory concerns. In this case, the consumer's marginal disutility from data use is then $\ell(\theta) = \gamma\theta$. This example corresponds to our model with a positive privacy β: higher-type consumers have a higher willingness to pay for privacy because price discrimination by the third party is more likely to hurt them. △

Example 2 (*Inequality in marginal values for money*). Suppose that consumers have homogeneous preferences over the seller's product and privacy, with a constant rate of substitution γ. They differ in income (or wealth), and thus in their marginal value of a dollar λ (as in Weitzman (1977) or Dworczak ⓡ al. (2021)). The consumer's utility is then $\tilde{x} \cdot (1 - \gamma\tilde{a}) - \lambda\tilde{t}$. After defining $\theta \equiv 1/\lambda$, we obtain our model with $\alpha = 0$ and $\beta = \gamma$, and thus we are in the case of positive privacy β. This example makes the simple point that—in the presence of inequality in opportunity cost of money—wealthier consumers tend to have a higher willingness to pay for all goods, including privacy. △

2.3 Policy Regimes

We consider three variants of the framework just introduced. They differ in the seller's ability to commit, at the beginning of the game, to the second-period data-pricing strategy $p : M \to \mathbb{R}$, and hence in which data records are sold to the third party.

Data Anarchy (denoted A). Under data anarchy, the seller cannot commit to the data-pricing strategy p. Instead, the data price must be chosen in a sequentially rational way. Since, as discussed above, the third party assigns positive value to any data records, the seller finds it optimal to sell any record she collects in any equilibrium. In other words, she cannot credibly commit to keep the consumer's records private. We envision this regime as resulting from the fact that, in many jurisdictions, there are no legal mechanisms or

institutional infrastructures to enforce such commitments.

Data Ban (denoted B). Under a data ban, the seller is legally prohibited from selling consumer data to third parties. This regime is easily accommodated in our model by assuming that the seller is exogenously committed to setting a price for data so high that the third party never purchases.

Data Market (denoted M). In a data market, the seller can commit to the data-pricing strategy p. That is, she can sell the product to the consumer with a promise that her data will not be sold to the third party.

2.4 Discussion

Our model makes several simplifying assumptions.

The most substantive assumption is that the consumer type is one-dimensional: θ jointly determines the consumer's valuation for the seller's product, her attitude toward data utilization through $\ell(\theta)$, and her value to the third party through $v(\theta, a)$. This simplification avoids the well-known challenge of multi-dimensional private information, which would make our model intractable. One interpretation is that consumers do not know their true disutility from privacy loss and hence use an approximation that takes the form assumed in equation (2).

We also made several assumptions about the way in which the seller interacts with the third party, which we see as less essential. First, disclosure is all-or-nothing: given a record m, the seller can either disclose m in full or withhold it entirely. Our results obtain under a richer disclosure model that allows for partial disclosure, provided no randomization is involved (see Appendix C). Second, disclosure is verifiable: Appendix C explains how our results change if disclosure is unverifiable. Third, the seller has full bargaining power vis-à-vis the third party. This assumption is not qualitatively important: because the third party has no private information, assigning some bargaining power to her would only change the split of surplus between the two parties. Fourth, data utilization a is not contractible, i.e., the seller cannot control the way in which the third party utilizes the data. This assumption turns the second stage of our framework into an information-design problem (the seller controls the third party's action only via disclosure of information). The assumption is realistic in most applications, although contracts giving the third party purpose-limited access to the database are also prevalent.[8]

[8]Making a contractible would further strengthen the commitment power of the seller in the data

Finally, we make several functional-form assumptions to keep the analysis tractable, two of which are worth highlighting. First, we assume that ℓ is affine. All our results extend to monotone privacy loss functions (see Appendix B.1). Second, we assume that the third party's payoff is linear in θ and quadratic in a. Our results extend to a broader class of third-party decision problems in which the optimal action is increasing and concave in the posterior mean and the third party strictly values information (see Appendix B.2).

3 The Role of the Privacy β: Implementability and Welfare

The paper's policy comparisons rest on two claims: the privacy β determines what forms of data sharing are feasible and whether data use creates or dissipates consumer rents. This section establishes both claims before solving for equilibrium.

We first characterize implementable outcomes and show how data anarchy, a data ban, and a data market impose different restrictions on disclosure. We then show that when $\beta > 0$, consumer incentive compatibility and third-party obedience conflict, forcing data sharing to be coarse; when $\beta \leq 0$, they align, permitting granular disclosure. Finally, holding product allocation fixed, the envelope formula shows that greater data utilization lowers consumer rents when $\beta \geq 0$ and raises them when $\beta \leq 0$. Sections 4 and 5 combine these results with the seller's participation and disclosure choices to derive the policy rankings.

3.1 Implementable Outcomes

Let an *outcome* (or direct mechanism) be described by a triple (x, t, a) where (i) $x : \Theta \to \{0,1\}$ is a deterministic allocation rule for the seller's product, (ii) $t : \Theta \to \mathbb{R}$ is the monetary transfer from the consumer to the seller, and (iii) $a : \Theta \to [0,1]$ is a deterministic recommendation rule for the third party—which also can be interpreted as an information structure. Without loss of generality, we assume $a(\theta) = 0$ whenever $x(\theta) = 0$, since the third party cannot use the data if the product was not sold. To simplify exposition, we will not distinguish between outcomes that differ on a measure-

market regime, likely leading to similar results.

zero set of types.[9]

We start by introducing three basic constraints imposed by private-information and moral-hazard frictions in our framework. Fix an outcome (x, t, a). Let $u(\hat{\theta}, \theta) \triangleq u\big(\theta \mid x(\hat{\theta}), t(\hat{\theta}), a(\hat{\theta})\big)$ denote the utility of a type-θ consumer who reports $\hat{\theta}$, and let $U(\theta) \triangleq u(\theta, \theta)$ denote the consumer's utility under truthful reporting. We say that an outcome is *individually rational* (IR) if for all θ, $U(\theta) \geq 0$ (we also, without loss of generality, impose that $U(0) = 0$). The outcome is *incentive compatible* (IC) if for all θ and $\hat{\theta}$, $U(\theta) \geq u(\theta, \hat{\theta})$. Finally, the outcome is *obedient* (OB) if, conditional on trade taking place and recommendation $\tilde{a} > 0$, it is optimal for the third party to follow such recommendation, that is, $\tilde{a} \in \arg\max_{a' \in A} \mathbb{E}\big(v(\theta, a') \mid a(\theta) = \tilde{a},\ x(\theta) = 1\big)$. Finally, we say that an outcome is *implementable* in a policy regime if it can be induced as an equilibrium of the continuation game after the selection of the seller's mechanism.

These three constraints characterize the set of outcomes the seller can implement across the policy regimes that we consider in this paper.

Remark 1. *An outcome (x, t, a) is implementable in*

1. *the data-market regime iff it satisfies IR, IC, and OB;*
2. *the data-anarchy regime iff it satisfies IR, IC, OB, and, additionally, $a(\theta) > 0$ for all θ such that $x(\theta) = 1$;*
3. *the data-ban regime iff it satisfies IR, IC, OB, and, additionally, $a(\theta) = 0$ for all θ.*

The first bullet follows from a standard application of the revelation principle. In particular, a data-market outcome may feature a positive measure of consumers who trade but whose data is not sold, since the seller can credibly commit to withholding their records. In the data-anarchy regime, by contrast, the seller lacks this commitment power. Thus, conditional on trade, any on-path recommendation must be sequentially optimal for the third party, including the recommendation $a = 0$. Since the third party's best response is the posterior mean of θ, a zero action can be sustained on a positive-probability trading set only if the posterior mean is zero, which can occur only on a null set. Finally, under a data ban, the seller is exogenously prevented from selling data, so $a(\theta) = 0$.

[9] Correspondingly, statements "for all types θ" should be interpreted as "for almost all types θ."

3.2 How Privacy β Affects Third-Party Data Sharing

In this section, we study how effectively the seller can monetize the data records generated by first-period transactions. Can these records reveal detailed information about individual consumers, or must data sales remain coarse? The answers to these questions determine whether the data collected as a by-product of goods transactions can support a rich market for consumer information. Our next result shows that the answer depends critically on the sign of the privacy β.

Proposition 1. *Consider any policy regime in which data sales are permitted.*

- *If $\beta > 0$, at most one data record is sold in any implementable outcome; this record pools all types of consumers whose data is sold.*
- *If $\beta \leq 0$, there exists an implementable outcome in which a continuum of data records is sold and each data record identifies a unique consumer type.*

Proposition 1 establishes that the sign of the privacy β determines whether data sales must be coarse or can be granular. When $\beta > 0$, the seller creates a single data record that pools consumer types, the third party always takes the same action conditional on purchasing any data, and the consumer receives a constant compensation for data sharing. In contrast, when $\beta \leq 0$, consumer records can be monetized in a granular way: Each data record may correspond to a different agent type, the third party can take customized actions depending on which record it purchases, and consumers' compensation varies with the intensity of data use.

The intuition for Proposition 1 comes from the way the IC and OB constraints interact and how this interaction depends on the privacy β.

Suppose first that the privacy β is positive. Consumers whose data the third party values most—namely, those with high θ—are precisely those who have the highest willingness to pay to protect their privacy. To satisfy IC, the seller must ensure that data utilization $a(\theta)$ is *non-increasing* on the set of trading types (Lemma A1 in the Appendix). Otherwise, under any monetary compensation for data sharing, either some high type would mimic a lower type to reduce their privacy loss or some lower type would mimic a higher type to obtain a lower price.[10] Conversely, because the third party would like to

[10]Indeed, suppose that the seller attempts to implement a higher data utilization $a' > a$ for a higher type $\theta' > \theta$. The option a' must be cheaper than a by at least $\beta\theta'(a' - a)$ to deter θ' from mimicking θ. But since the privacy loss for θ is lower than this discount, θ would then want to mimic θ'.

choose higher data utilization for higher consumer types, the OB constraint requires data utilization $a(\theta)$ to be *non-decreasing* on the set of trading types. Since these two forces pull in exactly opposite directions, the seller is forced to make data utilization constant in θ on the set of trading types. As a result, at most one strictly positive level $\tilde{a} > 0$ can be recommended to the third party, and any sold data reveal only *coarse* information about the consumer's type.

By contrast, suppose that the privacy β is negative. In this case, high-θ consumers are less sensitive to data utilization. To satisfy IC, the seller must ensure that data utilization $a(\theta)$ is *non-decreasing* on the set of trading types (Lemma A1 in the Appendix). Moreover, any non-decreasing $a(\theta)$ can be made incentive-compatible by an appropriate price schedule. The OB constraint does not depend on β and thus still requires data utilization to be *non-decreasing*. Since IC and OB constraints are aligned, the seller can implement mechanisms in which data utilization strictly increases with the consumer's type. In particular, sold records can fully reveal the type of each trading consumer, and information disclosure can be granular.

In a data market regime without other frictions, the granularity of the product menu might reveal the sign of β. However, fairness concerns and the technical complexity of such a menu could make such inference unreliable.

3.3 How Privacy β Affects Consumer Welfare

In this section, we analyze the second important consequence of the sign of the privacy β—its impact on consumer welfare. Although all consumers dislike data utilization, outcomes with more data use are not necessarily worse for consumers in equilibrium. This is because data use will affect prices that consumers face in the mechanism. Our next result shows that by fixing the set of trading types we can obtain a clean comparative static.

Proposition 2. *Among implementable outcomes with the same allocation rule x, the sign of the privacy β determines how data utilization affects consumer welfare:*

- *If $\beta \geq 0$, a pointwise increase in data utilization* weakly decreases *consumer welfare $U(\theta)$ for every θ.*
- *If $\beta \leq 0$, a pointwise increase in data utilization* weakly increases *consumer welfare $U(\theta)$ for every θ.*

To understand this result, recall that every consumer directly dislikes data utilization: the term $a(\theta)\ell(\theta) \geq 0$ enters negatively in the consumer's payoff in Equation (1). But equilibrium transfers adjust when data utilization changes. The relevant question is therefore whether this transfer adjustment more than offsets, exactly offsets, or only partially offsets the resulting privacy loss. Proposition 2 shows that the answer again depends critically on the sign of the privacy β.

The following thought experiment (illustrated in Figure 1) clarifies the basic logic behind Proposition 2. Suppose that the seller initially sells the product to every type above θ_g but does not sell any of their data, so that $a = 0$. Each trading consumer must then pay the same price for the product—the value θ_g of the marginal type. Now consider what happens when the seller induces a uniform level of data utilization $\bar{a} > 0$ while keeping the allocation rule, i.e., the cutoff θ_g, fixed. At optimum, the seller keeps the marginal type θ_g indifferent between buying or not, so she must reduce the price by $\bar{a}\ell(\theta_g)$—the privacy loss borne by the marginal type. What happens to other consumer types? If $\beta = 0$, so that privacy preferences do not depend on the consumer type, the price reduction exactly offsets the privacy loss for all consumers—utility is the same as in the case of no data use. If $\beta > 0$, the privacy loss $\ell(\theta)$ is increasing in type: Higher-valuation consumers are more privacy-sensitive. Thus, the price reduction only partially offsets the privacy loss of every type strictly above θ_g. Thus, data utilization lowers consumer surplus. Conversely, if $\beta < 0$, then $\ell(\theta)$ is decreasing in type, so the price reduction more than compensates the inframarginal consumers for their privacy loss. Thus, consumer welfare rises relative to the case with no data utilization.

The intuition from this simple case generalizes: When $\beta > 0$, high-valuation consumers are more privacy-sensitive, so exposing them to more data utilization compresses the rent profile. In particular, the marginal value of being a higher type falls from 1 to $1 - \beta a(\theta) < 1$, so higher data utilization dissipates consumer information rents. By contrast, when $\beta < 0$, high-valuation consumers are less privacy-sensitive. Data utilization then steepens the rent profile, increasing the marginal value of being a higher type from 1 to $1 - \beta a(\theta) > 1$. In this case, greater data utilization benefits consumers.

Proposition 2 holds fixed the set of consumers who buy the product. In equilibrium, however, data utilization and product allocation are chosen jointly, so regulation also changes the seller's incentive to serve the marginal consumer. This interaction is especially important when $\beta > 0$: because disclosed types are pooled, serving an additional consumer changes the information in the common record and hence the data uti-

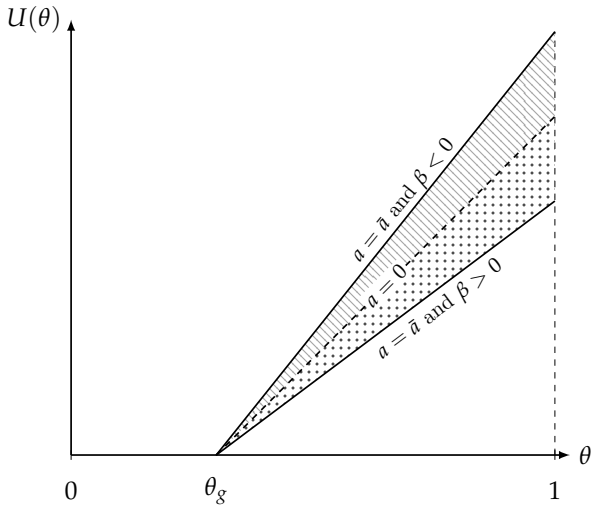

Figure 1: Consumer equilibrium utility $U(\theta)$ for a constant level of data utilization relative to no data utilization.

lization experienced by everyone in the pool, creating a disclosure externality familiar from information-design problems (see Galperti, Levkun and Perego (2024)). Some of our positive-β welfare results therefore require β to be sufficiently small, ensuring that the participation response does not overturn the rent effect identified above. Sections 4 and 5 analyze these equilibrium adjustments under each policy regime.

4 Positive Privacy β: Pooling and Privacy Premia

When $\beta > 0$, because disclosure must be coarse, the optimal mechanism can be summarized by two cutoffs: one determining who buys the product and another separating buyers whose data are sold from those who receive privacy. Characterizing these cutoffs and the resulting privacy premium allows us to study welfare when participation adjusts: regulation affects both the rents of existing buyers and the seller's incentive to serve the marginal consumer.

We find that less data utilization raises the rents of existing buyers, but data revenue may also make it profitable to serve lower types, which changes the composition and value of the pooled record. These forces explain why high-valuation consumers prefer a data ban, why the data market dominates anarchy when the data-compensation margin is not overturned, and how the observable price of two products with equal privacy can determine whether the Pareto-dominance case holds.

4.1 Equilibrium Outcomes

The first result characterizes the seller-optimal outcomes across the three policy regimes considered.

Proposition 3. *Suppose $\beta > 0$. Fix a policy regime $i \in \{A, B, M\}$. There exist cutoffs $\theta_g^i \le \theta_d^i$ such that a seller-optimal outcome (x^i, a^i, t^i) is as follows:*

- *All consumers whose type is above θ_g^i buy the product. That is, $x^i(\theta) = \mathbb{1}(\theta \ge \theta_g^i)$.*
- *Only consumers with types in $[\theta_g^i, \theta_d^i)$ have their data sold, and they are pooled into a single data record and thus receive the same treatment by the third party:*

$$a^i = \mathbb{E}(\theta \mid \theta_g^i \le \theta < \theta_d^i).$$

In particular, we have that:

- *Under data anarchy, $\theta_d^A = 1$, thus the data of* all *purchasing consumers is sold.*
- *Under data ban, $\theta_d^B = \theta_g^B$, thus the data of* no *purchasing consumers is sold.*
- *Under data market, $\theta_d^M \in [\theta_g^M, 1]$, thus the data of* some *purchasing consumers is sold.*

To begin, no data is sold in the data-ban regime, since third-party data sale is prohibited. The downstream data market therefore plays no role in the seller's problem, which reduces to the standard monopoly pricing problem of Myerson (1981). Hence, the seller posts a price equal to $\theta_g^{\mathrm{B}} = \theta_g^*$, where θ_g^* is the type at which the virtual product surplus, namely $\theta - (1 - F(\theta))/f(\theta)$, is zero.

In the data-anarchy regime, the optimal mechanism is also a posted price. In this case, however, the seller monetizes the data of every purchasing consumer. As anticipated by Proposition 1, data collection and monetization are coarse: only one record is sold to the third party, and it reveals only that the consumer bought the product, by pooling all types above θ_g^{A}.

Finally, in the data-market regime, the seller uses her commitment power to offer two versions of the product. One version comes with a privacy guarantee: the seller commits not to sell the data of consumers who purchase it. The other version comes with data monetization: the seller sells a pooled record revealing only that the consumer purchased the non-private version. Since $\beta > 0$, higher types are more privacy-sensitive, and therefore select the private version.

Price of privacy. In the data ban regime, privacy is imposed by regulation; in the data anarchy regime, the only way to ensure privacy is to not buy the product. The data market regime, however, potentially features a choice of the privacy treatment for the consumer, even conditional on buying the product.

Let $\bar{a} \triangleq \mathbb{E}(\theta \mid \theta_g^{\mathrm{M}} \leq \theta < \theta_d^{\mathrm{M}})$ be the equilibrium data utilization in the data market regime. Then, the consumer is charged

$$p_{\text{data}} = \theta_g^{\mathrm{M}} - \bar{a} \cdot \ell(\theta_g^{\mathrm{M}})$$

for the product if her data can then be used, but must pay

$$p_{\text{privacy}} = \theta_g^{\mathrm{M}} - \bar{a} \cdot \ell(\theta_g^{\mathrm{M}}) + \bar{a} \cdot \ell(\theta_d^{\mathrm{M}})$$

if she wants her data to be withheld. The price p_{data} is pinned down by the indifference of the marginal type θ_g^{M}. The difference in prices, $\bar{a} \cdot \ell(\theta_d^{\mathrm{M}})$, is increasing in θ_d^{M}—the highest type whose data is still disclosed. When all types between θ_g^{M} and θ_d^{M} are pooled, a higher θ_d^{M} raises data utilization conditional on sale of the data record. This makes it less attractive for data-sensitive high types, who choose the privacy-preserving option, to mimic lower types and allow their data to be sold. Hence, the seller can charge a higher price for the privacy-preserving option.

Consistent with Proposition 2, data sharing makes consumers worse off via the data-compensation margin. It is instructive to compare the privacy-preserving price p_{privacy} in the data market regime with the price the consumer pays for the product in the data ban regime. Holding fixed the participation margin, i.e., assuming that $\theta_g^{\mathrm{M}} = \theta_g^{\mathrm{B}}$, the consumer must pay $\bar{a}\beta(\theta_d^{\mathrm{M}} - \theta_g^{\mathrm{M}}) > 0$ more for the exact same outcome in the data market regime. We can interpret this quantity as a *privacy premium* paid by high types.

4.2 Welfare Effects and Policy Rankings

Having characterized the seller-optimal outcomes across the three policy regimes, we now discuss their welfare effects and policy rankings.

We begin with two observations. First, the seller is weakly better off under the data-market regime than under either data anarchy or a data ban. This is because data anarchy and data ban impose additional constraints relative to the data-market regime, limiting the seller's ability to extract surplus from the other players.[11] Second, the third

[11] Whether the seller prefers data anarchy to a data ban, or vice versa, depends on the primitives. Proposition A7 in the Appendix provides sufficient conditions for each case.

party is indifferent across regimes. Since she has no bargaining power and no private information, the seller extracts her entire surplus through the price charged for access to consumer data. Hence, the third party's payoff is always zero. In the remainder of this section, we focus on consumer welfare.

The following result summarizes our main findings.

Proposition 4. *Suppose $\beta > 0$.*

1. *If β is sufficiently small,*[12] *the data anarchy regime is Pareto dominated by the data market regime.*

2. *There exists a cutoff $\tilde{\theta}$ such that consumer types below $\tilde{\theta}$ weakly prefer the data market to the data ban, while consumer types above $\tilde{\theta}$ strictly prefer the data ban to the data market.*

Proposition 4 establishes that, when β is positive and sufficiently small, the data anarchy regime is weakly worse than the data market for *both* the seller and all consumers. The proof also shows that the data ban regime is weakly preferred to data anarchy by all consumers in this case. The fundamental reason is the one identified by Proposition 2: When $\beta > 0$, more data utilization hurts consumer welfare through the data-compensation margin.

As foreshadowed in Section 3.3, the assumption that β is small ensures that the participation margin (governed by the cutoff type θ_g from Proposition 3) goes in the same direction as the data-compensation margin. When data anarchy is replaced by data market, data utilization can only go down, which increases consumer welfare keeping θ_g constant (Proposition 2). But if the introduction of the data market increases θ_g, intermediate consumer types lose access to the product and are hence worse off. On one hand, the data market mitigates the negative disclosure externality present under data anarchy, where serving additional low types dilutes the pooled data record generated by all trading consumers—this pushes towards lower θ_g under the data market. On the other hand, the privacy premium charged to high types can be increased when the sold data record (pooling types between θ_g^{M} and θ_d^{M}) induces higher data utilization $\bar{a}$. When β is large relative to the value of data for the third party, the screening distortion becomes the dominant force, and the seller may raise θ_g^{M} to increase $\bar{a}$ and hence the privacy premium.

[12] It suffices that $\beta < 1/(1+L)$, where L is an upper bound on the absolute value of the slope of the inverse hazard rate of the distribution F.

The preceding discussion illustrates why giving more commitment power to the seller may hurt consumers. That said, we view the small-β case as natural: It says that willingness to pay for the product, while systematically related to marginal privacy losses, is not their primary determinant.

In the second part of Proposition 4, the trade-off between the data ban and the data market is driven by opposing data-compensation and participation margins. When $\beta > 0$, data utilization reduces consumer rents, so the data-compensation margin favors the data ban. This force is especially relevant for high types, who must pay a privacy premium in the data market to avoid disclosure. At the same time, the data market may benefit lower types through the participation margin: relative to the data ban, the seller values participation more on the margin because an additional sale generates data that can be sold to the third party. Thus, types above a cutoff $\tilde{\theta}$ prefer the data ban while types below it prefer the data market. The comparison of the data-compensation and participation margins gives a simple necessary and sufficient condition for when the data market weakly Pareto dominates the data ban.

Welfare Test 1. *Suppose $\beta > 0$. The data market weakly Pareto dominates the data ban if and only if the price $p_{privacy}$ in the data-market menu is no higher than the product price θ_g^B under the data ban.*

This is the first empirically implementable welfare test. The condition in Welfare Test 1 is useful because it is stated in terms of observable equilibrium objects. The relevant comparison is between the price of the private product version in the data-market menu and the product price under the data ban. Consumers who buy the private version under the data market receive the same downstream treatment as consumers under the data ban: their data are not sold. Hence, if the private version in the data-market menu is no more expensive than the data-ban product, every consumer who trades under the ban can obtain at least the same payoff under the data market, while consumers who do not trade can always opt out. Conversely, if the private version is more expensive than the data-ban product, sufficiently high types strictly prefer the data ban. Thus, the welfare comparison can be inferred from prices and privacy attributes of the offered products, without observing consumer types or the primitives of the model.

5 Negative Privacy β: Separation and Data Discounts

Section 3 showed that when $\beta < 0$, granular disclosure is feasible and greater data utilization raises consumer rents for a fixed set of buyers. We now apply these results to the seller's problem. The optimal mechanism can again be summarized by participation and disclosure cutoffs, but lower purchasing types may receive privacy while higher types sell their data in exchange for type-dependent discounts. The market therefore looks more like one in which consumers earn rents from selling their data. Allowing participation to adjust then reveals the central policy trade-off: a ban eliminates the consumer rents from data, while data anarchy preserves the data rents of high types but may make serving lower types unprofitable when their records must also be sold. A data market can expand product access, at the cost of reducing high types' data rents. We use this trade-off to compare the three regimes and derive an observable price test for the comparison between the data market and data anarchy.

5.1 Equilibrium Outcomes

We begin by characterizing the seller-optimal outcomes across the three policy regimes considered.

Proposition 5. *Suppose $\beta \leq 0$. Fix a policy regime $i \in \{A, B, M\}$. There exist cutoffs $\theta_g^i \leq \theta_d^i$ such that a seller-optimal outcome (x^i, a^i, t^i) is as follows:*

- *All consumers whose type is above θ_g^i buy the product. That is, $x^i(\theta) = \mathbb{1}(\theta \geq \theta_g^i)$.*
- *Only consumers with types in $[\theta_d^i, 1]$ have their data sold. Their types are fully revealed to the third party, whose action is therefore $a^i(\theta) = \theta$.*

  In particular, we have that:

  - *Under data anarchy, $\theta_d^A = \theta_g^A$, thus the data of* all *purchasing consumers is sold.*
  - *Under data ban, $\theta_d^B = 1$, thus the data of* no *purchasing consumers is sold.*
  - *Under data market, $\theta_d^M \in [\theta_g^M, 1]$, thus the data of* some *purchasing consumers is sold.*

The solution in the data ban case coincides with that discussed for the positive-β case. Since under a data ban third-party data sales are prohibited, there is no aftermarket and thus the sign of β does not affect the optimal outcome, which is a simple posted price.

In the data-anarchy regime, the seller cannot commit to withhold any data record generated by a purchasing consumer. Hence, the data of every consumer who buys the product are sold to the third party ($\theta_g^{\mathrm{A}} = \theta_d^{\mathrm{A}}$). Unlike in the positive-β case, however, data collection and monetization are granular rather than coarse (Proposition 1). The seller sells different records for different consumer types. Each record sold fully reveals the type of every purchasing consumer, and the third party's action is therefore $a^{\mathrm{A}}(\theta) = \theta$.

Finally, in the data-market regime, the seller uses her commitment power to offer a private option: types in $[\theta_g^{\mathrm{M}}, \theta_d^{\mathrm{M}})$ buy the product and experience no data utilization, i.e., $a(\theta) = 0$. Consumers with types in $[\theta_d^{\mathrm{M}}, 1]$, instead, buy the product and their data records are sold to the third party.

Proposition 5 differs from Proposition 3 in two qualitative ways. First, whenever data are monetized, monetization is granular rather than coarse, in line with Proposition 1. Second, the private option, when offered, caters to lower rather than higher types, within consumer types who buy the product. This is because, by definition when $\beta < 0$, it is the low types who have a higher willingness to pay for privacy.

Price of privacy. As in the positive β case, consumers receive privacy by design under a data ban and can only retain privacy by not buying the product in the data anarchy regime. In the data-market regime, because marginal types receive full privacy, the price for the privacy-preserving option is

$$p_{\text{privacy}} = \theta_g^{\mathrm{M}}.$$

There is no single compensation for data use, since every type whose data is sold experiences a different level of data utilization: For any $\theta \geq \theta_d^{\mathrm{M}}$, we have

$$p_{\text{data}}(\theta) = \theta_g^{\mathrm{M}} - \theta_d^{\mathrm{M}} \ell(\theta_d^{\mathrm{M}}) - \int_{\theta_d^{\mathrm{M}}}^{\theta} \ell(s) ds.$$

Thus, consumers whose data are sold, namely those with types $\theta \geq \theta_d^{\mathrm{M}}$, receive a discount relative to the private option. For the marginal type θ_d^{M}, the discount is exactly $\theta_d^{\mathrm{M}} \ell(\theta_d^{\mathrm{M}})$, which fully compensates this type for the marginal disutility she experiences from data utilization. For higher types, the discount is larger. Since $\beta < 0$, the marginal privacy loss is decreasing in type. As a result, this discount more than compensates types above θ_d^{M} for their privacy loss. Specifically, types whose data is utilized suffer a privacy loss of $\theta \ell(\theta)$ but the price they pay is lowered by

$$\theta_d^{\mathrm{M}} \ell(\theta_d^{\mathrm{M}}) + \int_{\theta_d^{\mathrm{M}}}^{\theta} \ell(s) ds \geq \theta \ell(\theta),$$

with the inequality being strict for all $\theta > \theta_d^{\mathrm{M}}$.

This calculation reinforces the intuition from Proposition 2 by showing that high types in the data market regime receive a *data premium*—a discount relative to the counterfactual treatment they would receive under a data ban. Unlike in the positive β case, consumers do not have to pay a privacy premium: Holding fixed the participation margin (i.e., assuming that $\theta_g^{\mathrm{M}} = \theta_g^{\mathrm{B}}$), consumers can secure the same privacy-preserving outcome at the same price in the data market and under a data ban. The fact that data utilization benefits consumers in the negative-β case underlies the flipped welfare results (relative to the positive-β case) to which we now turn.

5.2 Welfare Effects and Policy Rankings

The structure of the comparison parallels the one in Section 4.2. As before, the seller is weakly better off under the data-market regime, and the third party is always indifferent.[13] Hence, the relevant comparison is again driven by consumer welfare.

When $\beta < 0$, data utilization raises consumers' equilibrium rents. Nonetheless, the welfare effects of a policy can be complex for the same reasons discussed in Section 4.2. A policy introduces both a data-compensation margin—which data are being sold and how—and a participation margin—how this affects the sale of the product.

Proposition 6. *Suppose $\beta < 0$.*

1. *The data market weakly Pareto dominates the data ban.*
2. *There exists a cutoff $\tilde{\theta}$ such that consumer types below $\tilde{\theta}$ weakly prefer the data market, while consumer types above $\tilde{\theta}$ strictly prefer data anarchy.*

Proposition 6 shows that, when $\beta < 0$, one policy can be ruled out: a data ban. The reason is the mirror image of the positive-β case. Since data utilization increases consumers' equilibrium rents, prohibiting data sales removes a channel through which consumers can be compensated. A data market allows the seller to make extra profits through data sale, with the gains shared with consumers through discounts for data, as explained in the preceding subsection.

[13] Whether the seller prefers data anarchy to a data ban, or vice versa, depends on the primitives. Proposition A8 in the Appendix provides sufficient conditions for each case.

Unlike in the case of Proposition 4, no further assumption on β is needed: The participation margin can only reinforce the data-compensation margin. The structure of equilibrium revealed by Proposition 5 explains why. In the typical case when $\theta_g^{\mathrm{M}} < \theta_d^{\mathrm{M}}$, the marginal consumers are offered full privacy under the data market, so the incentives of the seller are identical on the margin across the two regimes—the prices are the same: $\theta_g^{\mathrm{M}} = \theta_g^{\mathrm{B}}$. In the boundary case $\theta_g^{\mathrm{M}} = \theta_d^{\mathrm{M}}$, in which the seller does not offer privacy in the mechanism under the data market, the marginal consumer offers valuable data and full disclosure rules out any data externalities. Thus, on the margin, the seller wants to sell the product to more types, and θ_g^{M} falls relative to θ_g^{B}—further benefiting consumers. Incidentally, in the boundary case, the data market and data anarchy regimes coincide, so the data ban is also Pareto dominated by data anarchy.

The comparison between the data market and data anarchy is more complicated because the data-compensation and participation margins go in opposite directions. Under data anarchy, consumers benefit from data utilization through data discounts, but privacy is not offered. The data market, by contrast, may introduce a private product version for intermediate types. This expands participation at the bottom of the market but reduces the data-utilization rents that high types obtain. As a result, types below some cutoff $\tilde{\theta}$ prefer the data market, while types above $\tilde{\theta}$ prefer data anarchy.

Because intermediate types benefit from the expanded access to the product under the data-market regime (whenever the two regimes differ), we can show that the data anarchy regime never Pareto dominates the data market:

Remark 2. *Suppose $\beta < 0$. The data anarchy regime never dominates the data market regime for consumers: The cutoff type $\tilde{\theta}$ can be taken to be strictly above θ_g^A.*

The opposite conclusion is not true: Even though high types receive rents from expanded data utilization under the data anarchy regime, they also receive rents from the expanded access to the product under the data market. As long as the latter effect dominates, all consumer types prefer the data market:

Welfare Test 2. *Suppose $\beta < 0$. The data market weakly Pareto dominates data anarchy ($\tilde{\theta} = 1$) if and only if the cheapest product version offered in the data-market menu is no more expensive than the cheapest product version offered under data anarchy.*

This is the second implementable welfare test. The condition in Remark 2 is stated in terms of observable equilibrium objects. When $\beta < 0$, the cheapest product version

is the one associated with the greatest data monetization and is selected by the highest types. Under both the data market and data anarchy, these consumers receive the same downstream treatment: their data are sold and their records fully reveal their types. Hence, if the cheapest data-market version is no more expensive than the cheapest data-anarchy version, even the consumers who benefit most from data monetization are weakly better off under the data market; conversely, if this condition fails, sufficiently high types strictly prefer data anarchy. Thus, the welfare comparison can be inferred from prices and privacy attributes of the offered products, without observing consumer types or the primitives of the model.

6 Conclusion

"For a struggling family, that could mean choosing between paying for privacy and paying for groceries" (Fulton, 2016). The spread of "consent-or-pay" models has made this concern increasingly concrete (European Data Protection Board, 2024). Our framework shows why the concern is important but incomplete. A privacy price may help the seller extract surplus, but data-funded discounts may also compensate consumers and expand access to the product. The privacy β determines how data use changes consumer rents and what forms of data sharing can be sustained.

Policy evaluation should therefore begin with observed prices and participation. Granular menus may initially appear troubling from a fairness perspective, but in our model they arise precisely when consumers benefit from data use through type-dependent discounts. More generally, the data market's flexibility can strengthen the seller's ability to screen consumers while also allowing transactions that would be excluded under a ban or data anarchy. Across both signs of the privacy β, lower-valuation consumers are the most consistent beneficiaries of this flexibility. If product valuations rise with income, then when private and non-private transactions are both available, it is poorest consumers who stand to gain the most. Ironically, it is the existence of a privacy premium that may make groceries affordable for the poor.

References

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

Adams, William James, and Janet L. Yellen. 1976. "Commodity Bundling and the Burden of Monopoly." *The Quarterly Journal of Economics*, 90(3): 475–498.

Ali, S Nageeb, Greg Lewis, and Shoshana Vasserman. 2023. "Voluntary Disclosure and Personalized Pricing." *The Review of Economic Studies*, 90(2): 538–571.

Baley, Isaac, and Laura L. Veldkamp. 2025. *The Data Economy: Tools and Applications.* Princeton, NJ:Princeton University Press.

Bergemann, Dirk, and Alessandro Bonatti. 2019. "Markets for Information: An Introduction." *Annual Review of Economics*, 11(1): 85–107.

Bergemann, Dirk, and Marco Ottaviani. 2021. "Information Markets and Nonmarkets." In *Handbook of Industrial Organization*. Vol. 4, Chapter 8, 593–672. Elsevier.

Bisceglia, Michele, Alessandro Bonatti, and Fiona Scott Morton. 2025. "Regulating Privacy Policies on Digital Platforms." Cowles Foundation for Research in Economics, Yale University 2474.

Calzolari, Giacomo, and Alessandro Pavan. 2006. "On the optimality of privacy in sequential contracting." *Journal of Economic Theory*, 130(1): 168–204.

Doval, Laura, and Vasiliki Skreta. 2026. "Purchase History and Product Personalization." *The RAND Journal of Economics*, 57(1): 78–102.

Dworczak, Piotr. 2020. "Mechanism Design With Aftermarkets: Cutoff Mechanisms." *Econometrica*, 88(6): 2629–2661.

Dworczak, Piotr. 2021. "Mechanism Design with Aftermarkets: Optimal Mechanisms under Binary Actions." Working paper, Northwestern University.

Dworczak, P. ⓡ S. D. Kominers ⓡ M. Akbarpour. 2021. "Redistribution through markets." *Econometrica*, 89(4): 1665–1698.

Eilat, Ran, Kfir Eliaz, and Xiaosheng Mu. 2021. "Bayesian Privacy." *Theoretical Economics*, 16(4): 1557–1603.

European Data Protection Board. 2024. "EDPB: 'Consent or Pay' Models Should Offer Real Choice."

FTC. 2014. "Data Brokers: A Call for Transparency and Accountability: A Report of the Federal Trade Commission." Federal Trade Commission.

Fulton, Sandra. 2016. "Pay-for-Privacy Schemes Put the Most Vulnerable People at Risk." *Free Press*.

Galperti, Simone, Aleksandr Levkun, and Jacopo Perego. 2024. "The Value of Data Records." *Review of Economic Studies*, 91(2): 1007–1038.

Goldfarb, Avi, and Catherine E. Tucker, ed. 2024. *The Economics of Privacy.* University of Chicago Press.

Goldfarb, Avi, and Verina F. Que. 2023. "The Economics of Digital Privacy." *Annual Review of Economics*, 15(1): 267–286.

Ichihashi, Shota, Doh-Shin Jeon, and Byung-Cheol Kim. 2025. "Mechanism Design for Ad-Supported Platforms." Working Paper.

Krähmer, Daniel, and Roland Strausz. 2023. "Optimal Nonlinear Pricing with Data-Sensitive Consumers." *American Economic Journal: Microeconomics*, 15(2): 80–108.

Lu, Zhuoran. 2025. "Selling Signals." *Journal of Economic Theory*, 224: 105966.

McAfee, R. Preston, John McMillan, and Michael D. Whinston. 1989. "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values." *The Quarterly Journal of Economics*, 104(2): 371–383.

Mussa, M., and S. Rosen. 1978. "Monopoly and Product Quality." *Journal of Economic Theory*, 18(2): 301–317.

Myerson, Roger B. 1981. "Optimal auction design." *Mathematics of operations research*, 6(1): 58–73.

Rayo, Luis. 2013. "Monopolistic Signal Provision." *The B.E. Journal of Theoretical Economics*, 13(1): 27–58.

Rivera Mora, Ernesto. 2025. "Neutral Mechanisms: On the Feasibility of Information Sharing." Job market paper / working paper.

Weitzman, Martin L. 1977. "Is the price system or rationing more effective in getting a commodity to those who need it most?" *Bell Journal of Economics*, 8(2): 517–524.

Appendix

A Proofs

A.1 Proofs for Section 3

Proof of Remark 1. Consider the data-market regime. If an outcome is implementable, consumer optimality implies IC, the opt-out message $\varnothing$ with $x(\varnothing) = t(\varnothing) = 0$ implies IR, and sequential optimality of the third party implies OB for every positive on-path recommendation. Conversely, let (x, t, a) satisfy IR, IC, and OB. Consider the direct mechanism with message space $\Theta \cup \{\varnothing\}$, where $\varnothing$ gives payoff zero. After report θ, the seller implements $(x(\theta), t(\theta))$ and commits to disclose the record iff $a(\theta) > 0$. If the record is disclosed, OB implies that the third party optimally chooses $a(\theta)$. IR and IC imply that truthful participation is optimal. Hence (x, t, a) is implementable.

Consider the data-anarchy regime. Necessity of IR, IC, and OB follows as above. It remains to characterize the additional restriction imposed by the absence of commitment. Let $Z \triangleq \{\theta \mid x(\theta) = 1,\ a(\theta) = 0\}$. Suppose by way of contradiction that $F(Z) > 0$. Then sequential rationality requires that the third party be willing to choose action zero after the records generated by types in Z. However, since third-party optimality requires that $a^* = \mathbb{E}(\theta|Z)$, this is possible only if $F(Z) = 0$, a contradiction. Conversely, suppose that (x, t, a) satisfies IR, IC, OB, and $F(Z) = 0$. Implement the outcome by the direct mechanism and let the seller sell every on-path trading record. Since $F(Z) = 0$, every on-path recommendation is positive, and OB guarantees third-party optimality. Hence the outcome is implementable.

Finally, consider the data-ban regime. Since data sales are prohibited, any implementable outcome must satisfy $a(\theta) = 0$, hence, $F\big(\{\theta \mid a(\theta) = 0\}\big) = 1$. Necessity of IR and IC follows as above, and OB is vacuously satisfied because no positive recommendation is made on path. Conversely, fix an outcome satisfying IR, IC, OB, and $a(\theta) = 0$ for all θ. The direct mechanism with message space $\Theta \cup \{\varnothing\}$ implements (x, t), with $\varnothing$ yielding payoff zero, and the ban ensures that no record is ever sold. IR and IC guarantee truthful participation. Since $a(\theta) = 0$ for all θ, OB is vacuous. Hence the outcome is implementable in the data-ban regime. □

Lemma A1. *An outcome (x, t, a) is IC, IR, and OB if and only if*

1. *There exists* $\theta_g \in [0,1]$ *such that* $x(\theta) = \mathbb{1}(\theta \geq \theta_g)$*, for all* θ*;*

2. $\beta a(\theta)$ *is non-increasing in* θ*, for all* $\theta \geq \theta_g$*;*

3. $\tilde{a} = \mathbb{E}\big(\theta \mid \tilde{a} = a(\theta)\big)$ *for all* $\tilde{a} > 0$ *in the image of* a*;*

4. *Transfers are given by*

$$t(\theta) = \begin{cases} 0 & \theta < \theta_g, \\ \theta - a(\theta)(\alpha + \beta\theta) - \int_{\theta_g}^{\theta} \big(1 - \beta a(s)\big) ds & \theta \geq \theta_g. \end{cases} \tag{A1}$$

Proof of Lemma A1. Fix an outcome (x,t,a). For each report θ', define $q(\theta') \triangleq x(\theta')\big(1 - \beta a(\theta')\big)$ and $b(\theta') \triangleq -\alpha x(\theta')a(\theta') - t(\theta')$. Then the payoff of type θ from reporting θ' can be written as $u(\theta', \theta) = \theta q(\theta') + b(\theta')$. Note that $q(\theta) \geq 0$ since $\beta < 1$ and $a \in [0,1]$.

Necessity. Suppose (x,t,a) is IC, IR, and OB. Fix $\theta > \theta'$. IC implies $U(\theta) = u(\theta,\theta) \geq u(\theta',\theta) = \theta q(\theta') + b(\theta')$ and $U(\theta') = u(\theta',\theta') \geq u(\theta,\theta') = \theta' q(\theta) + b(\theta)$. Adding these inequalities yields $(\theta - \theta')\big(q(\theta) - q(\theta')\big) \geq 0$. Since $\theta > \theta'$, we obtain $q(\theta) \geq q(\theta')$, so q is weakly increasing. Thus, the deterministic allocation rule x must be a cutoff rule: there exists $\theta_g \in [0,1]$ such that $x(\theta) = \mathbb{1}\{\theta \geq \theta_g\}$ (part 1). Moreover, on the trading region $[\theta_g, 1]$, we must have $q(\theta) = 1 - \beta a(\theta)$ nondecreasing (part 2).

Moreover, IC implies $U(\theta) = \max_{\hat{\theta} \in \Theta} \big(\theta q(\hat{\theta}) + b(\hat{\theta})\big)$. That is, $U(\theta)$ is the pointwise maximum of affine functions of θ. Thus, U is convex and, by the envelope theorem, $U'(\theta) =_{\text{a.e.}} q(\theta)$. Since $q \geq 0$, U is weakly increasing. Individual rationality requires $U(\theta) \geq 0$ for all θ. This pins down U up to the constant $U(0)$

$$U(\theta) = U(0) + \int_0^{\theta} q(s) ds.$$

Using that $U(\theta) \triangleq x(\theta)\big(\theta - a(\theta)\big(\alpha + \beta\theta\big)\big) - t(\theta)$ gives Equation (A1). From this, it is obvious that any revenue maximizing mechanism must satisfy $U(0) = 0$.

Finally, fix a recommendation $\tilde{a} > 0$ in the image of a. Since by assumption $a(\theta) = 0$ whenever $x(\theta) = 0$, observing $\tilde{a} > 0$ implies that trade has occurred. Given the definition of $v(\theta, a)$, the unique maximizer of $\max_{a' \in A} \mathbb{E}\big(v(\theta, a') \mid a(\theta) = \tilde{a},\ x(\theta) = 1\big)$ is $\tilde{a} = \mathbb{E}[\theta \mid a(\theta) = \tilde{a}]$. This proves necessity of condition 3.

Sufficiency. Conversely, suppose that conditions 1–4 hold.

Define $q(\theta) = x(\theta)\big(1 - \beta a(\theta)\big)$. Part 1 implies that $q(\theta) = 0$ for $\theta < \theta_g$, while condition 2 implies that $q(\theta) = 1 - \beta a(\theta)$ is weakly increasing on $[\theta_g, 1]$. Hence q is weakly

increasing on Θ. Let $U(\theta) \triangleq \int_0^\theta q(s)\,ds$ and $t(\theta) \triangleq x(\theta)\big(\theta - a(\theta)(\alpha+\beta\theta)\big) - U(\theta)$ (where we used $U(0)=0$). Then, U is convex and differentiable almost everywhere with $U'(\theta) =_{\text{a.e.}} q(\theta)$. Moreover, $U(\theta) \geq 0$ for all θ, and thus IR holds. Moreover, for any pair (θ,θ'),

$$\begin{aligned} u(\theta',\theta) &= x(\theta')\big(\theta - a(\theta')(\alpha+\beta\theta)\big) - t(\theta') \\ &= x(\theta')\big(\theta' - a(\theta')(\alpha+\beta\theta')\big) - t(\theta') + (\theta-\theta')x(\theta')\big(1-\beta a(\theta')\big) \\ &= U(\theta') + (\theta-\theta')q(\theta'), \end{aligned}$$

where we added and subtracted $x(\theta')(\theta' - a(\theta')(\alpha+\beta\theta'))$ in the second line. On the other hand, $U(\theta) - U(\theta') = \int_{\theta'}^{\theta} q(s)ds \geq (\theta-\theta')q(\theta')$, since $q(s) \geq q(\theta')$ on $[\theta',\theta]$. Thus, $U(\theta) \geq U(\theta') + (\theta-\theta')q(\theta') = u(\theta',\theta)$. The case with $\theta < \theta'$ follows by similar steps. We conclude that $U(\theta) \geq u(\theta,\theta')$ for all θ, hence the mechanism is IC.

Finally, condition 3 is exactly the third party's optimality condition derived above, so OB holds. Therefore the outcome is IC, IR, and OB. □

Proof of Proposition 1. Fix an implementable outcome in a regime in which data sales are permitted. By Remark 1, the outcome satisfies IC and OB.

Consider first $\beta > 0$. By Lemma A1, $a(\theta)$ is non-increasing on the set of trading types. Suppose, toward a contradiction, there are $\tilde{a}_H > \tilde{a}_L > 0$ such that the sets $a^{-1}(\tilde{a}_H)$ and $a^{-1}(\tilde{a}_L)$ have positive measure. Since a is non-increasing, every type in $a^{-1}(\tilde{a}_H)$ lies weakly below every type in $a^{-1}(\tilde{a}_L)$. Hence $\mathbb{E}(\theta \mid \theta \in a^{-1}(\tilde{a}_H)) \leq \mathbb{E}(\theta \mid \theta \in a^{-1}(\tilde{a}_L))$. By OB, however, $\tilde{a}_H = \mathbb{E}(\theta \mid \theta \in a^{-1}(\tilde{a}_H))$ and $\tilde{a}_L = \mathbb{E}(\theta \mid \theta \in a^{-1}(\tilde{a}_L))$, which contradicts $\tilde{a}_H > \tilde{a}_L$. Thus at most one positive data record can be sold.

Now suppose $\beta \leq 0$. Let $x(\theta) = 1$ and $a(\theta) = \theta$ for all θ. Then $\beta a(\theta) = \beta\theta$ is non-increasing, and OB holds because each positive record fully reveals the consumer's type. With transfers given by (A1), Lemma A1 implies that the outcome is IC, IR, and OB. Hence full disclosure is implementable in any regime in which data sales are permitted. □

Lemma A2. *Fix any IC, IR, and OB outcome (x,a,t). There exists $\theta_g \in [0,1]$ such that*

1. *The seller's expected profit is*

$$\int_{\theta_g}^{1} \Big(\underbrace{\phi_g(\theta)}_{\substack{\text{virtual} \\ \text{product surplus}}} + a(\theta) \underbrace{\phi_d(\theta)}_{\substack{\text{virtual} \\ \text{data surplus}}} \Big) f(\theta)\,d\theta.$$

where $\phi_g(\theta) \triangleq \theta - \frac{1-F(\theta)}{f(\theta)}$ *and* $\phi_d(\theta) \triangleq \theta - (\alpha + \beta\theta) + \beta\frac{1-F(\theta)}{f(\theta)} = \theta - \ell(\phi_g(\theta))$

2. *The third party's expected profit is* 0.

3. *Utility of consumer with type* θ *is*

$$U(\theta) = \begin{cases} 0 & \theta < \theta_g, \\ \underbrace{(\theta - \theta_g)}_{\substack{\text{surplus from} \\ \text{product}}} - \underbrace{\int_{\theta_g}^{\theta} \beta a(s)\, ds}_{\substack{\text{surplus from} \\ \text{data use}}}, & \theta \geq \theta_g. \end{cases}$$

Proof of Lemma A2. Fix any IC, IR, and OB outcome (x, a, t). Using Equation (A1) from Lemma A1, we obtain that there is $\theta_g \in [0,1]$ such that

$$\begin{aligned} \mathbb{E}(t(\theta)) &= \int_0^1 t(\theta) f(\theta)\, d\theta \\ &= \int_{\theta_g}^1 \Big(\theta - a(\theta)(\alpha + \beta\theta)\Big) f(\theta)\, d\theta - \int_{\theta_g}^1 \big(1 - \beta a(\theta)\big)\big(1 - F(\theta)\big)\, d\theta \\ &= \int_{\theta_g}^1 \Big(\phi_g(\theta) - a(\theta)\big(\alpha + \beta\phi_g(\theta)\big)\Big) f(\theta)\, d\theta. \end{aligned}$$

To obtain the second line, we used the fact that:

$$\int_{\theta_g}^1 \left(\int_{\theta_g}^{\theta} (1 - \beta a(s))\, ds\right) f(\theta)\, d\theta = \int_{\theta_g}^1 (1 - \beta a(s))\big(1 - F(s)\big)\, ds.$$

The expected price charged to the third party is $\mathbb{E}(v(\theta, a(\theta)))$. Thus, the seller's expected payoff is

$$\begin{aligned} \mathbb{E}(t) + \mathbb{E}(v) &= \int_{\theta_g}^1 \Big(\phi_g(\theta) - a(\theta)\big(\alpha + \beta\phi_g(\theta)\big) + 2\theta a(\theta) - a(\theta)^2\Big) f(\theta)\, d\theta \\ &= \int_{\theta_g}^1 \Big(\phi_g(\theta) + a(\theta)\phi_d(\theta) + a(\theta)\big(\theta - a(\theta)\big)\Big) f(\theta)\, d\theta. \end{aligned}$$

The result follows from $\int_{\theta_g}^1 a(\theta)\big(\theta - a(\theta)\big) f(\theta) d\theta = 0$. To see this, denote by $\tilde{\theta}$ the random variable of θ, which is distributed according to F. Then

$$\begin{aligned} \int_{\theta_g}^1 a(\theta)\big(\theta - a(\theta)\big) f(\theta)\, d\theta &= \big(1 - F(\theta_g)\big)\mathbb{E}\Big(a(\tilde{\theta})\big(\tilde{\theta} - a(\tilde{\theta})\big) \mid \tilde{\theta} \geq \theta_g\Big) \\ &= \big(1 - F(\theta_g)\big)\mathbb{E}\Big(a(\tilde{\theta})\Big(\mathbb{E}(\tilde{\theta} \mid a(\tilde{\theta}), \tilde{\theta} \geq \theta_g) - a(\tilde{\theta})\Big) \mid \tilde{\theta} \geq \theta_g\Big) \\ &= 0. \end{aligned}$$

The second equality uses the law of iterated expectations. The last equality holds because, if $a(\tilde{\theta}) = 0$, the entire expression is zero, while if $a(\tilde{\theta}) > 0$, OB implies $\mathbb{E}(\tilde{\theta} \mid a(\tilde{\theta}), \tilde{\theta} \geq \theta_g) = \mathbb{E}(\tilde{\theta} \mid a(\tilde{\theta})) = a(\tilde{\theta})$.

With regard to point 2, note that whenever a data record is sold, the seller makes a take-it-or-leave-it offer that extracts the entire value of the record from the third party. Hence, the third party's expected payoff is always zero.

With regard to point 3, by the proof of Lemma A1, there exists $\theta_g \in [0,1]$ such that

$$U(\theta) = \int_0^\theta x(s)\big(1 - \beta a(s)\big)\, ds = \begin{cases} 0, & \theta < \theta_g, \\ \displaystyle\int_{\theta_g}^\theta \big(1 - \beta a(s)\big)\, ds, & \theta \geq \theta_g. \end{cases}$$

The expression in the lemma then follows immediately. □

Proof of Proposition 2. Consider two implementable outcomes with the same allocation rule, and let a and $\hat{a}$ denote their data-utilization rules, with $\hat{a}(\theta) \geq a(\theta)$ for all θ. Since the allocation rule is the same, Lemma A2 applies with the same cutoff θ_g. For $\theta < \theta_g$, both utilities are equal to zero. For $\theta \geq \theta_g$, point 3 of Lemma A2 implies $\hat{U}(\theta) - U(\theta) = -\beta \int_{\theta_g}^\theta (\hat{a}(s) - a(s))\, ds$. The integral is nonnegative. Hence, if $\beta \geq 0$, $\hat{U}(\theta) \leq U(\theta)$ for every θ, while if $\beta \leq 0$, $\hat{U}(\theta) \geq U(\theta)$ for every θ. □

A.2 Proofs for Section 4

Proof of Proposition 3. Suppose $\beta > 0$ and fix a regime $i \in \{A, B, M\}$. By Remark 1, any outcome implementable in regime i must satisfy IR, IC, and OB, together with the additional regime-specific restriction as listed in the remark. By Lemma A1, any such outcome has a cutoff product allocation: there exists $\theta_g \in [0,1]$ such that $x(\theta) = \mathbb{1}\{\theta \geq \theta_g\}$. Lemma A1 also pins down the transfers through equation (A1). Thus, it remains only to characterize the data-utilization rule a.

Consider first the data ban. Remark 1 implies that $a(\theta) = 0$ for all θ. Hence the data-sale region is empty, which corresponds to $\theta_d^{\mathrm{B}} = \theta_g^{\mathrm{B}}$. The cutoff θ_g^{B} is chosen to maximize the seller's payoff in Lemma A2 subject to this restriction, which leads to choosing θ_g^{B} such that $\phi_g(\theta_g^{\mathrm{B}}) = 0$.

Now consider a regime in which data sales are permitted. Since $\beta > 0$, Proposition 1 implies that at most one positive data record can be sold. Let the corresponding recommendation be $\bar{a}$. By Lemma A1, a is non-increasing on the trading region $[\theta_g, 1]$.

Hence the set of trading types whose data are sold must be an initial interval of the trading region. Therefore, for some $\theta_d \in [\theta_g, 1]$, up to null sets,

$$a(\theta) = \bar{a}\,\mathbb{1}\{\theta \in [\theta_g, \theta_d)\}, \qquad \bar{a} = \mathbb{E}(\theta \mid \theta_g \le \theta < \theta_d)$$

whenever $\theta_d > \theta_g$, while $a(\theta) = 0$ when $\theta_d = \theta_g$. The equality defining $\bar{a}$ is exactly the obedience constraint.

The remaining regime-specific restrictions determine the admissible values of θ_d. Under data anarchy, Remark 1 rules out a positive-measure set of trading types with $a(\theta) = 0$. Since the trading region is $[\theta_g, 1]$, this implies $\theta_d^A = 1$. Under a data market, no further restriction is imposed, so θ_d^M can be any element of $[\theta_g^M, 1]$.

Finally, the cutoffs are chosen to maximize the seller's payoff from Lemma A2. For any feasible pair (θ_g, θ_d), this payoff is

$$\int_{\theta_g}^{1} \varphi_g(\theta) f(\theta) d\theta + \bar{a}(\theta_g, \theta_d) \int_{\theta_g}^{\theta_d} \varphi_d(\theta) f(\theta) d\theta,$$

with the convention that the second term is zero when $\theta_d = \theta_g$. In each regime, the feasible set for (θ_g, θ_d) is compact, and the objective is continuous under this convention. Hence a maximizer exists; denote it by (θ_g^i, θ_d^i). The corresponding allocation, data-utilization rule, and transfers have the form stated in the proposition. □

Lemma A3. *Suppose $0 < \beta \le \frac{1}{1+L}$. Let (θ_g^M, θ_d^M) be the thresholds in the data-market regime from Proposition 3. Then, $\theta_d^M \in \{\theta_g^M, 1\}$. Therefore, the data-market solution coincides with either the data-ban solution or the data-anarchy one.*

Proof of Lemma A3. By Proposition 3, any data-market outcome with $\beta > 0$ can be represented by two cutoffs $\theta_g \le \theta_d$. Types in $[\theta_g, \theta_d)$ are pooled into a single data record. If $\theta_d > \theta_g$, the induced data-utilization level is $\bar{a}(\theta_g, \theta_d) = \mathbb{E}[\theta \mid \theta_g \le \theta < \theta_d]$; if $\theta_d = \theta_g$, no data are sold.

By Lemma A2, for fixed θ_g, the only part of the seller's payoff that depends on θ_d is

$$D(\theta_d) \triangleq \bar{a}(\theta_g, \theta_d) \int_{\theta_g}^{\theta_d} \varphi_d(\theta) f(\theta) d\theta,$$

with the convention that $D(\theta_g) = 0$.

The small-β assumption implies that φ_d is nondecreasing. Indeed, writing $r(\theta) = (1 - F(\theta))/f(\theta)$, we have $\varphi_d(\theta) = \theta - (\alpha + \beta\theta) + \beta r(\theta)$, and therefore $\varphi_d'(\theta) \ge 1 - \beta(1 + L) \ge 0$ a.e.

We claim that, for every fixed θ_g, $D(\theta_d)$ is maximized at either $\theta_d = \theta_g$ or $\theta_d = 1$. To see this, take any $\theta_d \in (\theta_g, 1)$. If $\int_{\theta_g}^{\theta_d} \varphi_d(\theta) f(\theta) d\theta < 0$, then $D(\theta_d) < 0 = D(\theta_g)$, so θ_d is not optimal. If instead this integral is nonnegative, monotonicity of φ_d implies $\varphi_d(\theta_d) \geq 0$; otherwise φ_d would be negative throughout $[\theta_g, \theta_d]$. Hence extending the pool from θ_d to 1 weakly increases the integral. It also weakly increases $\bar{a}(\theta_g, \theta_d)$. Therefore $D(1) \geq D(\theta_d)$. Hence an optimal data-market outcome can be selected with $\theta_d^M \in \{\theta_g^M, 1\}$, as claimed. □

To state the next result, we first introduce the following terminology. Recall that θ_g^{B} is the unique type at which the virtual product surplus is zero, that is, $\phi_g(\theta_g^{\mathrm{B}}) = 0$. This is the marginal type who buys the product in the absence of an aftermarket. The ability to sell data will indirectly affect the optimal threshold θ_g^i. Ideally, the seller would like to sell the data records of all consumer types whose virtual data surplus ϕ_d is positive. We say that *data sales boost product sales* if $\phi_d(\theta_g^{\mathrm{B}}) \geq 0$: the seller wants to sell the data of the marginal consumer to whom she would sell the product absent an aftermarket. Notice that this condition is a restriction on primitives, not on endogenous variables. We say that *data sales depress product sales* if $\phi_d(\theta_g^{\mathrm{B}}) < 0$: the seller does not want to sell the data of the marginal consumer to whom she would sell the product absent an aftermarket. Intuitively, these distinctions determine whether allowing data sales changes the product-sales threshold θ_g in a way that benefits or harms consumers.

Proposition A7. *Suppose $\beta > 0$. The seller weakly prefers the data-market regime to other regimes. Moreover, if ϕ_d is increasing,*

1. *If data sales are never profitable, i.e., $\phi_d(1) \leq 0$, the optimal mechanism in the data-market regime coincides with that in the data ban regime, and $V^M = V^B \geq V^A$.*

2. *If data sale boosts product sales, i.e., $\phi_d(\theta_g^B) \geq 0$, then the optimal mechanism in the data market regime coincides with that in the data anarchy regime, and $V^M = V^A \geq V^B$.*

3. *In all other cases, namely, if data sales depress product sales, i.e., $\phi_d(\theta_g^B) < 0$, but data sales are sometimes profitable, i.e., $\phi_d(1) > 0$, then the optimal mechanism in the data market coincides with either that in the data ban regime or with that in the data anarchy regime, and $V^M = \max\{V^B, V^A\}$.*

Proof of Proposition A7. The seller's weak preference for the data-market regime follows directly from Remark 1. By Lemma A3, the data-market outcome coincides with

either the data-ban outcome or the data-anarchy outcome. Hence $V^{\mathrm{M}} = \max\{V^{\mathrm{B}}, V^{\mathrm{A}}\}$, and it remains only to compare V^{B} and V^{A}.

By Lemma A2, the seller's payoff under the data ban is $V^{\mathrm{B}} = \max_{\theta_g \in [0,1]} \int_{\theta_g}^{1} \phi_g(\theta) f(\theta) d\theta$, and the maximizer is θ_g^{B}, such that $\phi_g(\theta_g^{\mathrm{B}}) = 0$. Under data anarchy, Proposition 3 implies that the seller chooses a cutoff θ_g, sells the data of all purchasing consumers, and induces $\bar{a}(\theta_g) = \mathbb{E}[\theta \mid \theta \geq \theta_g]$. Therefore,

$$V^{\mathrm{A}} = \max_{\theta_g \in [0,1]} \int_{\theta_g}^{1} \left[\phi_g(\theta) + \bar{a}(\theta_g)\phi_d(\theta)\right] f(\theta) d\theta.$$

First, suppose that $\phi_d(1) \leq 0$. Since β is small, ϕ_d is nondecreasing. This implies $\phi_d(\theta) \leq 0$ for all θ. Hence, for every cutoff θ_g, the data-sale term in the data-anarchy payoff is weakly negative. Therefore the data-anarchy payoff at any cutoff is weakly below the data-ban payoff evaluated at the same cutoff, and so $V^{\mathrm{A}} \leq V^{\mathrm{B}}$. Since $V^{\mathrm{M}} = \max\{V^{\mathrm{B}}, V^{\mathrm{A}}\}$, we obtain $V^{\mathrm{M}} = V^{\mathrm{B}} \geq V^{\mathrm{A}}$.

Next, suppose that $\phi_d(\theta_g^{\mathrm{B}}) \geq 0$. Since ϕ_d is nondecreasing, $\phi_d(\theta) \geq 0$ for all $\theta \geq \theta_g^{\mathrm{B}}$. Evaluating the data-anarchy payoff at the data-ban cutoff θ_g^{B}, the data-sale term is weakly positive. Hence $V^{\mathrm{A}} \geq V^{\mathrm{B}}$. Therefore $V^{\mathrm{M}} = V^{\mathrm{A}} \geq V^{\mathrm{B}}$.

The other case is residual. □

Proof of Proposition 4. By assumption, $\beta < 1/(1+L)$, where L is an upper bound on the absolute value of the slope of the inverse hazard rate (Footnote 12). Letting $r(\theta) \equiv (1 - F(\theta))/f(\theta)$, we then have $\phi_d'(\theta) = 1 - \beta + \beta r'(\theta) \geq 1 - \beta(1+L) > 0$. That is, ϕ_d is nondecreasing. By Lemma A3, the data-market outcome coincides with either the data-ban outcome or the data-anarchy outcome. Since both outcomes are feasible in the data-market regime, the data-market outcome coincides with the one that gives the seller the larger payoff.

Let θ_g^{A} and θ_g^{B} be the product cutoffs under data anarchy and data ban. Under data ban, $U^{\mathrm{B}}(\theta) = (\theta - \theta_g^{\mathrm{B}})^+$. Under data anarchy, all purchasing consumers' records are pooled, so, letting $\bar{a}^{\mathrm{A}} \triangleq \bar{a}(\theta_g^{\mathrm{A}}, 1)$ and $\lambda^{\mathrm{A}} \triangleq 1 - \beta\bar{a}^{\mathrm{A}}$, Lemma A2 gives $U^{\mathrm{A}}(\theta) = \lambda^{\mathrm{A}}(\theta - \theta_g^{\mathrm{A}})^+$. Notice that $\lambda^{\mathrm{A}} \in (0,1)$.

We first compare the data market with data anarchy. If the data-market outcome coincides with the data-anarchy outcome, then $U^{\mathrm{M}} = U^{\mathrm{A}}$. Suppose instead that the data-market outcome coincides with the data-ban outcome. Then $V^{\mathrm{B}} \geq V^{\mathrm{A}}$. Write $\Phi_g(\tau) \triangleq \int_{\tau}^{1} \phi_g(\theta) f(\theta)\, d\theta$ and $\Phi_d(\tau, 1) \triangleq \int_{\tau}^{1} \phi_d(\theta) f(\theta)\, d\theta$. Since $V^{\mathrm{B}} = \Phi_g(\theta_g^{\mathrm{B}})$ and

$V^{\mathrm{A}} \geq \Phi_g(\theta_g^{\mathrm{B}}) + \bar{a}(\theta_g^{\mathrm{B}}, 1)\Phi_d(\theta_g^{\mathrm{B}}, 1)$, this implies

$$\Phi_g(\theta_g^{\mathrm{B}}) \geq \Phi_g(\theta_g^{\mathrm{B}}) + \bar{a}(\theta_g^{\mathrm{B}}, 1)\Phi_d(\theta_g^{\mathrm{B}}, 1),$$

and therefore $\Phi_d(\theta_g^{\mathrm{B}}, 1) \leq 0$. Since ϕ_d is nondecreasing, it follows that $\phi_d(\theta_g^{\mathrm{B}}) \leq 0$, and hence $\Phi_d(\tau, 1) \leq 0$ for every $\tau < \theta_g^{\mathrm{B}}$. Because Φ_g is maximized at θ_g^{B}, no maximizer of the data-anarchy problem can lie below θ_g^{B}. Thus $\theta_g^{\mathrm{A}} \geq \theta_g^{\mathrm{B}}$. Since $\lambda^{\mathrm{A}} < 1$, this implies $U^{\mathrm{B}}(\theta) \geq U^{\mathrm{A}}(\theta)$ for all θ. Hence $U^{\mathrm{M}}(\theta) \geq U^{\mathrm{A}}(\theta)$ for all θ, proving the first claim.

We now compare the data market with the data ban. If the data-market outcome coincides with the data-ban outcome, take $\tilde{\theta} = 1$. Suppose instead that it coincides with the data-anarchy outcome, so $U^{\mathrm{M}} = U^{\mathrm{A}}$. If $\theta_g^{\mathrm{B}} \leq \theta_g^{\mathrm{A}}$, then $U^{\mathrm{B}}(\theta) \geq U^{\mathrm{M}}(\theta)$ for all θ, with strict inequality for every $\theta > \theta_g^{\mathrm{B}}$; in this case take $\tilde{\theta} = \theta_g^{\mathrm{B}}$. If $\theta_g^{\mathrm{B}} > \theta_g^{\mathrm{A}}$, define

$$\tilde{\theta} \triangleq \min\left\{1, \theta_g^{\mathrm{A}} + \frac{\theta_g^{\mathrm{B}} - \theta_g^{\mathrm{A}}}{\beta \bar{a}^{\mathrm{A}}}\right\}.$$

For $\theta < \theta_g^{\mathrm{A}}$, both utilities are zero, and for $\theta \in (\theta_g^{\mathrm{A}}, \theta_g^{\mathrm{B}})$, $U^{\mathrm{M}}(\theta) > U^{\mathrm{B}}(\theta)$. For $\theta \geq \theta_g^{\mathrm{B}}$, the difference $U^{\mathrm{M}}(\theta) - U^{\mathrm{B}}(\theta) = \theta_g^{\mathrm{B}} - \lambda^{\mathrm{A}}\theta_g^{\mathrm{A}} - (1 - \lambda^{\mathrm{A}})\theta$ is strictly decreasing. Therefore $U^{\mathrm{M}}(\theta) \geq U^{\mathrm{B}}(\theta)$ if and only if $\theta \leq \tilde{\theta}$. Thus consumer types below $\tilde{\theta}$ weakly prefer the data market, while consumer types above $\tilde{\theta}$ strictly prefer the data ban. □

Proof of Welfare Test 1. Let $p^{\mathrm{B}} = \theta_g^{\mathrm{B}}$ be the data-ban price, and let $p^{\mathrm{M}}_{\mathrm{priv}}$ be the price of the private product version in the data-market menu. If the data-market outcome coincides with the data-ban outcome, then $p^{\mathrm{M}}_{\mathrm{priv}} = p^{\mathrm{B}}$ and the claim is immediate. Suppose instead that the data-market outcome coincides with the data-anarchy outcome. By the transfer formula in Proposition 3, the price of the private product version, possibly degenerate at $\theta = 1$, is $p^{\mathrm{M}}_{\mathrm{priv}} = \theta_g^{\mathrm{A}} + \beta \bar{a}^{\mathrm{A}}(1 - \theta_g^{\mathrm{A}})$.

From the proof of Proposition 4, the data market weakly Pareto dominates the data ban if and only if the highest type weakly prefers the data market to the data ban, that is, if and only if $(1 - \beta\bar{a}^{\mathrm{A}})(1 - \theta_g^{\mathrm{A}}) \geq 1 - \theta_g^{\mathrm{B}}$. Rearranging gives $\theta_g^{\mathrm{A}} + \beta\bar{a}^{\mathrm{A}}(1 - \theta_g^{\mathrm{A}}) \leq \theta_g^{\mathrm{B}}$, which is exactly $p^{\mathrm{M}}_{\mathrm{priv}} \leq p^{\mathrm{B}}$. □

A.3 Proofs for Section 5

Proof of Proposition 5. Suppose $\beta \leq 0$. By Remark 1, any implementable outcome must satisfy IR, IC, and OB, together with the regime-specific restrictions stated there. By Lemma A1, any such outcome has a cutoff product allocation: there exists $\theta_g \in [0, 1]$

such that $x(\theta) = 1(\theta \geq \theta_g)$. Lemma A1 also pins down transfers through equation (A1). It remains to characterize the data-utilization rule a.

Consider first the data ban. Remark 1 implies that $a(\theta) = 0$ for all θ. Hence the data-sale region is empty, which corresponds to $\theta_d^{\mathrm{B}} = 1$. The cutoff θ_g^{B} is chosen to maximize the seller's payoff in Lemma A2 subject to this restriction.

Now consider a regime in which data sales are permitted. For $\beta < 0$, Lemma A1 implies that a is nondecreasing on the trading region $[\theta_g, 1]$; for $\beta = 0$, incentive compatibility imposes no restriction on a. In either case, an optimal mechanism can be selected so that no positive data record pools a positive-measure set of types. To see this, suppose that a positive record pools an interval I and induces action $\bar{a} = \mathbb{E}(\theta \mid \theta \in I)$. Replacing this pooled record by full revelation on I, so that $\hat{a}(\theta) = \theta$ for $\theta \in I$, preserves OB. It also preserves IC: when $\beta < 0$, the replacement of a flat segment of a nondecreasing rule by the identity keeps the rule nondecreasing, and when $\beta = 0$, IC is unaffected by a. The change in the seller's payoff is

$$\int_I (\theta - \bar{a})\phi_d(\theta) f(\theta)\, d\theta = \int_I (\theta - \bar{a})\big(\phi_d(\theta) - \phi_d(\bar{a})\big) f(\theta)\, d\theta,$$

where the equality uses $\bar{a} = \mathbb{E}(\theta \mid \theta \in I)$. Since ϕ_d is nondecreasing in the negative-β case, this expression is weakly positive. Hence full revelation weakly improves the seller's payoff, and an optimal outcome can be chosen with no positive pooling. Thus, whenever a positive record is sold, it fully reveals the consumer's type and the third party's action is $a(\theta) = \theta$.

Under data anarchy, Remark 1 rules out a positive-measure set of trading types with $a(\theta) = 0$. Hence every purchasing consumer's record is sold, and the preceding paragraph implies $a^{\mathrm{A}}(\theta) = \theta$ for all $\theta \geq \theta_g^{\mathrm{A}}$. This corresponds to $\theta_d^{\mathrm{A}} = \theta_g^{\mathrm{A}}$. The cutoff θ_g^{A} is chosen to maximize

$$\int_{\theta_g}^1 \Big(\phi_g(\theta) + \theta\phi_d(\theta)\Big) f(\theta)\, d\theta.$$

Under the data-market regime, the seller can commit to withhold records. Since positive sold records fully reveal types, the seller chooses which trading types' records to sell. For $\beta < 0$, IC requires the resulting data-utilization rule to be nondecreasing; for $\beta = 0$, since ϕ_d is increasing, the set of types whose fully revealing records are sold can be selected as an upper interval. Thus there exists $\theta_d^{\mathrm{M}} \in [\theta_g^{\mathrm{M}}, 1]$ such that $a^{\mathrm{M}}(\theta) = 0$ for $\theta \in [\theta_g^{\mathrm{M}}, \theta_d^{\mathrm{M}})$ and $a^{\mathrm{M}}(\theta) = \theta$ for $\theta \geq \theta_d^{\mathrm{M}}$. For any feasible pair (θ_g, θ_d), the seller's payoff is

$$\int_{\theta_g}^1 \phi_g(\theta) f(\theta)\, d\theta + \int_{\theta_d}^1 \theta\phi_d(\theta) f(\theta)\, d\theta.$$

The feasible set is compact and the objective is continuous, so an optimal pair exists; denote it by $(\theta_g^{\mathrm{M}}, \theta_d^{\mathrm{M}})$.

The corresponding allocation, data-utilization rule, and transfers have the form stated in the proposition. □

Lemma A4. *Suppose $\beta < 0$. Recall that θ_g^* is defined by $\phi_g(\theta_g^*) = 0$, and let (θ_g^M, θ_d^M) denote an optimal pair of cutoffs in the data market regime, and let θ_g^A denote an optimal cutoff in the data anarchy regime. Then:*

- *If $\phi_d(\theta_g^*) < 0$, then $\theta_g^B = \theta_g^M < \theta_g^A \leq \theta_d^M$, with the last strict inequality if $\theta_d^M < 1$.*
- *If $\phi_d(\theta_g^*) \geq 0$, then $\theta_g^M = \theta_g^A = \theta_d^M \leq \theta_g^B$.*

Proof. Recall that $Q(\theta) \triangleq \phi_g(\theta) + \theta\phi_d(\theta)$. Since $\beta < 0$, ϕ_d is strictly increasing. By Proposition 5, the seller's problem in the data market reduces to choosing cutoffs (θ_g, θ_d) with $\theta_g \leq \theta_d$ so as to maximize

$$\Pi^{\mathrm{M}}(\theta_g, \theta_d) = \int_{\theta_g}^1 \phi_g(\theta) f(\theta)\, d\theta + \int_{\theta_d}^1 \theta\, \phi_d(\theta) f(\theta)\, d\theta.$$

Write $\Phi_g(\theta_g) \triangleq \int_{\theta_g}^1 \phi_g(\theta) f(\theta)\, d\theta$ and $\Phi_d(\theta_d) \triangleq \int_{\theta_d}^1 \theta\, \phi_d(\theta)\, f(\theta)\, d\theta$. Then $\Pi^{\mathrm{M}}(\theta_g, \theta_d) = \Phi_g(\theta_g) + \Phi_d(\theta_d)$. By regularity, Φ_g is uniquely maximized at θ_g^*. Note, $\Phi_d'(\theta_d) = -\theta_d\, \phi_d(\theta_d)\, f(\theta_d)$.

We now distinguish two cases.

Case 1: $\phi_d(\theta_g^*) < 0$. In this case, $\Phi_d'(\theta_g^*) = -\theta_g^* \phi_d(\theta_g^*) f(\theta_g^*) > 0$. Therefore no maximizer of Φ_d can lie weakly below θ_g^* (since ϕ_d is increasing). Let $\hat{\theta}_d \in \arg\max \Phi_d$. Then necessarily $\hat{\theta}_d > \theta_g^*$. Since Φ_g is uniquely maximized at θ_g^*, the feasible pair $(\theta_g, \theta_d) = (\theta_g^*, \hat{\theta}_d)$ simultaneously maximizes both terms of the separable objective, and therefore is an optimal pair in the data market regime. It follows that every optimal pair satisfies $\theta_g^{\mathrm{M}} = \theta_g^{\mathrm{B}} < \theta_d^{\mathrm{M}}$.

Now consider the data anarchy regime. By Proposition 5, the seller chooses $\tau \in [0, 1]$ to maximize $\Pi^{\mathrm{A}}(\tau) \triangleq \int_\tau^1 Q(\theta) f(\theta)\, d\theta$. For every $\theta \leq \theta_g^*$, we have $\phi_g(\theta) \leq 0$ and, since ϕ_d is increasing with $\phi_d(\theta_g^*) < 0$, also $\phi_d(\theta) \leq \phi_d(\theta_g^*) < 0$. Hence $Q(\theta) = \phi_g(\theta) + \theta\phi_d(\theta) < 0$ for all $\theta \leq \theta_g^*$. Therefore Π^{A} is strictly increasing on $[0, \theta_g^*]$, and by continuity also on $[0, \theta_g^* + \varepsilon]$ for some $\varepsilon > 0$. Thus, no maximizer of Π^{A} can lie weakly below $\theta_g^{\mathrm{B}} = \theta_g^{\mathrm{M}}$.

Assume now in addition that $\theta_d^{\mathrm{M}} < 1$. Since θ_d^{M} is then an interior maximizer of the differentiable function Φ_d, we must have $\Phi_d'(\theta_d^{\mathrm{M}}) = 0$, hence $\phi_d(\theta_d^{\mathrm{M}}) = 0$. Because

$\theta_d^{\mathrm{M}} > \theta_g^*$ and ϕ_g is strictly increasing with unique zero at θ_g^*, it follows that $Q(\theta_d^{\mathrm{M}}) = \phi_g(\theta_d^{\mathrm{M}}) + \theta_d^{\mathrm{M}}\phi_d(\theta_d^{\mathrm{M}}) = \phi_g(\theta_d^{\mathrm{M}}) > 0$. Moreover, for every $\theta > \theta_d^{\mathrm{M}}$, we have $\phi_g(\theta) > 0$ and $\phi_d(\theta) > 0$, hence $Q(\theta) > 0$. Therefore Π^{A} is strictly decreasing on $[\theta_d^{\mathrm{M}}, 1]$, so no maximizer of Π^{A} can lie weakly above θ_d^{M}. Thus $\theta_g^{\mathrm{A}} < \theta_d^{\mathrm{M}}$. Combining the two inequalities yields $\theta_g^{\mathrm{B}} = \theta_g^{\mathrm{M}} < \theta_g^{\mathrm{A}} < \theta_d^{\mathrm{M}}$.

Case 2: $\phi_d(\theta_g^*) \geq 0$. Since ϕ_d is strictly increasing, we have $\phi_d(\theta) \geq 0$ for all $\theta \geq \theta_g^*$, and therefore $\Phi_d'(\theta) = -\theta\phi_d(\theta)f(\theta) \leq 0$ for all $\theta \geq \theta_g^*$. Hence Φ_d is weakly decreasing on $[\theta_g^*, 1]$. Take any feasible pair (θ_g, θ_d). If $\theta_d \geq \theta_g^*$, then the best feasible choice of θ_g is $\theta_g = \theta_g^*$, since Φ_g is uniquely maximized at θ_g^*. But because Φ_d is weakly decreasing on $[\theta_g^*, 1]$, replacing θ_d by θ_g^* weakly raises profit. Hence no optimum with $\theta_d \geq \theta_g^*$ can satisfy $\theta_d > \theta_g^*$. If instead $\theta_d \leq \theta_g^*$, then feasibility implies $\theta_g \leq \theta_d \leq \theta_g^*$. Since $\phi_g < 0$ on $[0, \theta_g^*)$, we have $\Phi_g'(\theta) = -\phi_g(\theta)f(\theta) > 0$ on $[0, \theta_g^*)$, so Φ_g is strictly increasing on $[0, \theta_g^*]$. Therefore, holding θ_d fixed, the best feasible choice is $\theta_g = \theta_d$. Combining the two observations, every optimal pair must lie on the diagonal and satisfy $\theta_g^{\mathrm{M}} = \theta_d^{\mathrm{M}} \leq \theta_g^{\mathrm{B}}$.

Finally, on the diagonal $\theta_g = \theta_d = \tau$, the objective in the data market case becomes

$$\Phi_g(\tau) + \Phi_d(\tau) = \int_\tau^1 (\phi_g(\theta) + \theta\phi_d(\theta))f(\theta)\, d\theta = \int_\tau^1 Q(\theta)f(\theta)\, d\theta,$$

which is exactly the objective in the data anarchy case. Since the optimal pair in the data market regime lies on the diagonal, its cutoff θ_g^{M} is also optimal in the data anarchy regime. Hence, if the cutoff in the data anarchy regime is generically unique, then $\theta_g^{\mathrm{M}} = \theta_g^{\mathrm{A}} = \theta_d^{\mathrm{M}} \leq \theta_g^{\mathrm{B}}$. This completes the proof. □

Proposition A8. *The seller weakly prefers the data market regime to other regimes. Moreover,*

1. *If data sales are never profitable, i.e., $\phi_d(1) \leq 0$, the optimal mechanism in the data market regimes coincides with that in the data ban regime, and $V^M = V^B \geq V^A$.*

2. *If data sales depress product sales, i.e., $\phi_d(\theta_g^*) < 0$, but data sales are sometimes profitable, i.e., $\phi_d(1) > 0$, then the optimal mechanism in the data market contains an option with a privacy guarantee and $V^M > \max\{V^B, V^A\}$;*

3. *If data sale boosts product sales, i.e., $\phi_d(\theta_g^*) \geq 0$, then the optimal mechanism in the data market regime coincides with that in the data anarchy regime, and $V^M = V^A \geq V^B$.*

Proof of Proposition A8. Using the same notation as in the proof of Lemma A4, the seller's objective in regime M can be written as $\Pi^{\mathrm{M}}(\theta_g, \theta_d) = \Phi_g(\theta_g) + \Phi_d(\theta_d)$ such that $\theta_g \le \theta_d$. Moreover, $V^{\mathrm{B}} = \Phi_g(\theta_g^{\mathrm{B}})$ and $V^{\mathrm{A}} = \max_{\tau\in[0,1]}\{\Phi_g(\tau) + \Phi_d(\tau)\}$. Since the data ban and data anarchy mechanisms are feasible for the seller in the data market regime, $V^{\mathrm{M}} \ge \max\{V^{\mathrm{B}}, V^{\mathrm{A}}\}$.

If $\phi_d(\theta_g^*) \ge 0$, Lemma A4 implies that $\theta_g^{\mathrm{M}} = \theta_d^{\mathrm{M}} = \theta_g^{\mathrm{A}}$. Thus, optimal mechanisms in the data market and the data anarchy regimes coincide, and $V^{\mathrm{M}} = V^{\mathrm{A}} \ge V^{\mathrm{B}}$.

Next, suppose that $\phi_d(1) \le 0$. Since ϕ_d is strictly increasing, $\phi_d(\theta) \le 0$ for all $\theta \in [0,1]$. Therefore $\Phi_d'(\theta) = -\theta\,\phi_d(\theta) f(\theta) \ge 0$, so Φ_d is increasing and is maximized at $\theta = 1$, where $\Phi_d(1) = 0$. Hence in the data market regime, the optimal choice is $\theta_d^{\mathrm{M}} = 1$, and then optimality of θ_g^{B} for Φ_g gives $V^{\mathrm{M}} = \Phi_g(\theta_g^{\mathrm{B}}) = V^{\mathrm{B}}$. Also, for every $\tau \in [0,1]$, $\Phi_g(\tau) + \Phi_d(\tau) \le \Phi_g(\tau) + \Phi_d(1) = \Phi_g(\tau) \le \Phi_g(\theta_g^{\mathrm{B}}) = V^{\mathrm{B}}$, so $V^{\mathrm{A}} \le V^{\mathrm{B}}$. Therefore, optimal mechanisms in the data market and the data ban regimes coincide, and $V^{\mathrm{M}} = V^{\mathrm{B}} \ge V^{\mathrm{A}}$.

Finally, suppose that $\phi_d(\theta_g^*) < 0$ and $\phi_d(1) > 0$. Since ϕ_d is strictly increasing, there exists a unique $\hat{\theta} \in (\theta_g^*, 1)$ such that $\phi_d(\hat{\theta}) = 0$. Hence $\Phi_d'(\theta) = -\theta\,\phi_d(\theta) f(\theta)$ is positive on $[\theta_g^*, \hat{\theta})$ and negative on $(\hat{\theta}, 1]$, so $\hat{\theta}$ uniquely maximizes Φ_d on $[\theta_g^*, 1]$. Therefore, the optimal pair in regime M is $(\theta_g^{\mathrm{M}}, \theta_d^{\mathrm{M}}) = (\theta_g^*, \hat{\theta})$, and $V^{\mathrm{M}} = \Phi_g(\theta_g^*) + \Phi_d(\hat{\theta}) > \Phi_g(\theta_g^*) = V^{\mathrm{B}}$. Since $\hat{\theta} < 1$, Lemma A4 yields $\theta_g^* = \theta_g^{\mathrm{M}} < \theta_g^{\mathrm{A}} < \theta_d^{\mathrm{M}} = \hat{\theta}$. In particular, the optimal mechanism in the data market regime contains an option with a privacy guarantee. Because Φ_g is uniquely maximized at θ_g^*, we have $\Phi_g(\theta_g^{\mathrm{A}}) < \Phi_g(\theta_g^*)$, and because Φ_d is strictly increasing on $[\theta_g^*, \hat{\theta}]$, $\Phi_d(\theta_g^{\mathrm{A}}) < \Phi_d(\hat{\theta})$. Thus $V^{\mathrm{A}} = \Phi_g(\theta_g^{\mathrm{A}}) + \Phi_d(\theta_g^{\mathrm{A}}) < \Phi_g(\theta_g^*) + \Phi_d(\hat{\theta}) = V^{\mathrm{M}}$. Hence $V^{\mathrm{M}} > \max\{V^{\mathrm{B}}, V^{\mathrm{A}}\}$. □

Proof of Proposition 6. By Lemma A1, in any IC and IR mechanism the buyer's interim utility satisfies $U^i(\theta) = \int_0^\theta q^i(s)\,ds$, with $q^i(\theta) \triangleq x^i(\theta)\big(1 - \beta a^i(\theta)\big)$, for $i \in \{\mathrm{A}, \mathrm{B}, \mathrm{M}\}$. From Proposition 5, the optimal mechanisms imply $q^{\mathrm{B}}(\theta) = \mathbb{1}\{\theta \ge \theta_g^{\mathrm{B}}\}$, $q^{\mathrm{A}}(\theta) = \mathbb{1}\{\theta \ge \theta_g^{\mathrm{A}}\}(1 - \beta\theta)$, and $q^{\mathrm{M}}(\theta) = \mathbb{1}\{\theta_g^{\mathrm{M}} \le \theta\} - \beta\theta\mathbb{1}\{\theta \ge \theta_d^{\mathrm{M}}\}$.

For part 1, Lemma A4 implies that either $\theta_g^{\mathrm{M}} = \theta_g^{\mathrm{B}} < \theta_d^{\mathrm{M}}$ when $\phi_d(\theta_g^*) < 0$, or $\theta_g^{\mathrm{M}} = \theta_d^{\mathrm{M}} = \theta_g^{\mathrm{A}} \le \theta_g^{\mathrm{B}}$ when $\phi_d(\theta_g^*) \ge 0$. In either case, $q^{\mathrm{M}}(\theta) \ge q^{\mathrm{B}}(\theta)$ for every θ. Hence $U^{\mathrm{M}}(\theta) \ge U^{\mathrm{B}}(\theta)$ for all θ. Moreover, by Proposition A8, $V^{\mathrm{M}} \ge V^{\mathrm{B}}$.

For part 2, suppose first that $\phi_d(\theta_g^*) \ge 0$. By Lemma A4, $\theta_g^{\mathrm{M}} = \theta_g^{\mathrm{A}} = \theta_d^{\mathrm{M}}$. Hence $q^{\mathrm{M}}(\theta) = q^{\mathrm{A}}(\theta)$ for all θ, so $U^{\mathrm{M}}(\theta) = U^{\mathrm{A}}(\theta)$ for all $\theta \in [0,1]$. Moreover, Proposition A8 implies $V^{\mathrm{M}} = V^{\mathrm{A}}$. Thus, the two regimes lead to the same outcome for all players, and

the desired conclusion holds with $\tilde{\theta} = 1$.

Now, suppose that $\phi_d(\theta_g^*) < 0$. Then Lemma A4 gives $\theta_g^{\mathrm{M}} = \theta_g^{\mathrm{B}} < \theta_g^{\mathrm{A}} \le \theta_d^{\mathrm{M}}$. Define $D(\theta) \triangleq U^{\mathrm{M}}(\theta) - U^{\mathrm{A}}(\theta)$. Since $D'(\theta) = q^{\mathrm{M}}(\theta) - q^{\mathrm{A}}(\theta)$ almost everywhere, we obtain

$$q^{\mathrm{M}}(\theta) - q^{\mathrm{A}}(\theta) = \begin{cases} 0, & \theta < \theta_g^{\mathrm{M}}, \\ 1, & \theta_g^{\mathrm{M}} \le \theta < \theta_g^{\mathrm{A}}, \\ \beta\theta, & \theta_g^{\mathrm{A}} \le \theta < \theta_d^{\mathrm{M}}, \\ 0, & \theta \ge \theta_d^{\mathrm{M}}. \end{cases}$$

Integrating and using $D(0) = 0$, it follows that $D(\theta) = 0$ for $\theta < \theta_g^{\mathrm{M}}$, $D(\theta) = \theta - \theta_g^{\mathrm{M}} > 0$ for $\theta \in [\theta_g^{\mathrm{M}}, \theta_g^{\mathrm{A}})$,

$$D(\theta) = \theta_g^{\mathrm{A}} - \theta_g^{\mathrm{M}} + \frac{\beta}{2}\left(\theta^2 - (\theta_g^{\mathrm{A}})^2\right) \qquad \text{for } \theta \in [\theta_g^{\mathrm{A}}, \theta_d^{\mathrm{M}}),$$

and

$$D(\theta) = D(\theta_d^{\mathrm{M}}) = \theta_g^{\mathrm{A}} - \theta_g^{\mathrm{M}} + \frac{\beta}{2}\left((\theta_d^{\mathrm{M}})^2 - (\theta_g^{\mathrm{A}})^2\right) \qquad \text{for } \theta \in [\theta_d^{\mathrm{M}}, 1].$$

Therefore D is strictly decreasing on $(\theta_g^{\mathrm{A}}, \theta_d^{\mathrm{M}})$ and constant on $[\theta_d^{\mathrm{M}}, 1]$. If $D(\theta_d^{\mathrm{M}}) \ge 0$, then $D(\theta) \ge 0$ for all θ, and the desired conclusion holds with $\tilde{\theta} = 1$. If instead $D(\theta_d^{\mathrm{M}}) < 0$, then $D(\theta_g^{\mathrm{A}}) = \theta_g^{\mathrm{A}} - \theta_g^{\mathrm{M}} > 0$, while D is continuous and strictly decreasing on $[\theta_g^{\mathrm{A}}, \theta_d^{\mathrm{M}}]$. Hence there exists a unique $\tilde{\theta} \in (\theta_g^{\mathrm{A}}, \theta_d^{\mathrm{M}})$ such that $D(\tilde{\theta}) = 0$, and $D(\theta) \ge 0$ if and only if $\theta \le \tilde{\theta}$.

This completes the proof. □

Proof of Welfare Test 2. Let $p_{\min}^{\mathrm{M}}$ and $p_{\min}^{\mathrm{A}}$ denote the cheapest product-version prices in the data market and data anarchy regimes. Since the third party's payoff is zero in both regimes and Proposition A8 gives $V^{\mathrm{M}} \ge V^{\mathrm{A}}$, the data market weakly Pareto dominates data anarchy if and only if $U^{\mathrm{M}}(\theta) \ge U^{\mathrm{A}}(\theta)$ for all θ.

If $\phi_d(\theta_g^*) \ge 0$, Lemma A4 implies that the two regimes coincide, so $p_{\min}^{\mathrm{M}} = p_{\min}^{\mathrm{A}}$ and the claim is immediate. Suppose instead that $\phi_d(\theta_g^*) < 0$. By Proposition 5 and the transfer formula, the cheapest product version in both regimes is the version selected by the highest type. That type receives the product and has her data fully revealed in both regimes. Hence $U^{\mathrm{M}}(1) \ge U^{\mathrm{A}}(1)$ if and only if $p_{\min}^{\mathrm{M}} \le p_{\min}^{\mathrm{A}}$.

From the proof of Proposition 6, $D(\theta) \triangleq U^{\mathrm{M}}(\theta) - U^{\mathrm{A}}(\theta)$ is positive up to θ_g^{A}, strictly decreasing on $(\theta_g^{\mathrm{A}}, \theta_d^{\mathrm{M}})$, and constant on $[\theta_d^{\mathrm{M}}, 1]$. Therefore $D(\theta) \ge 0$ for every θ if and only if $D(1) \ge 0$. Combining this with the previous paragraph proves the claim. □

Online Appendix

B Extensions

B.1 Monotone Privacy Loss

This section illustrates how the main results in Sections 3–5 extend beyond the affine specification of Equation (2). Suppose that a consumer facing data utilization a incurs privacy loss $a\ell(\theta)$, where $\ell : [0,1] \to (0,\infty)$ is absolutely continuous and strictly monotone. We maintain the usual single-crossing condition that $\theta - a\ell(\theta)$ is nondecreasing in θ for every $a \in [0,1]$, equivalently $\ell'(\theta) \leq 1$ a.e.

For this extension, we need to assume that both the virtual product surplus and the virtual data surplus are nondecreasing. Let $r(\theta) \triangleq (1 - F(\theta))/f(\theta)$, $\phi_g(\theta) \triangleq \theta - r(\theta)$, and $\phi_d^\ell(\theta) \triangleq \theta - \ell(\theta) + r(\theta)\ell'(\theta)$.

Assumption 1. *For the results in this section, we assume that ϕ_g and ϕ_d^ℓ are nondecreasing.*

The following remark gives sufficient conditions on F and ℓ that guarantee this assumption is satisfied.

Remark A1. *Suppose that ϕ_g is nondecreasing, and let L be an upper bound on $|r'|$.*

- *Suppose ℓ is increasing. If ℓ is weakly convex and $\ell' \leq \frac{1}{1+L}$ a.e., then ϕ_d^ℓ is nondecreasing.*
- *Suppose ℓ is decreasing. If ℓ is weakly convex, then ϕ_d^ℓ is nondecreasing.*

The proof of these sufficient conditions is immediate and thus omitted.

The implementability results are unchanged. The regime characterization in Remark 1 does not use the affine form of ℓ. IC and IR again imply a cutoff allocation, and on the trading region IC requires $(a(\theta_H) - a(\theta_L))(\ell(\theta_H) - \ell(\theta_L)) \leq 0$ for all $\theta_L < \theta_H$. Hence a is nonincreasing when ℓ is increasing and nondecreasing when ℓ is decreasing. The envelope formula becomes $U(\theta) = \theta - \theta_g - \int_{\theta_g}^{\theta} a(s)\ell'(s)\,ds$ for $\theta \geq \theta_g$, with the corresponding transfer pinned down by $t(\theta) = \theta - a(\theta)\ell(\theta) - U(\theta)$. The seller's payoff representation is the one above: replace ϕ_d everywhere with ϕ_d^ℓ. The third party's payoff remains zero.

It follows that the Section 3 conclusions go through with the sign of β replaced by the monotonicity of ℓ. If ℓ is increasing, IC and OB are opposed, so at most one positive record can be sold. If ℓ is decreasing, IC and OB are aligned, and full disclosure is implementable. Moreover, holding the allocation fixed, increasing data utilization changes consumer utility by $-\int_{\theta_g}^{\theta} \Delta a(s)\ell'(s)\,ds$, so it weakly lowers consumer welfare when ℓ is increasing and weakly raises it when ℓ is decreasing.

Consider next the positive-privacy results in Section 4. When ℓ is increasing, the optimal-mechanism result in Proposition 3 survives: sold data are pooled into a single record, and the sold-data region is an initial interval $[\theta_g, \theta_d)$. The assumption that ϕ_d^ℓ is nondecreasing gives the same binary-optimality argument as before, so the data-market optimum coincides with either the data-ban outcome or the data-anarchy outcome. Proposition 4 then follows from the same comparison of the two cutoffs. The affine formulas are replaced by ℓ-differences; for example, the price of the private version in the data-market menu is $p^{\mathrm{M}}_{\mathrm{priv}} = \theta_g^{\mathrm{M}} + \bar{a}^{\mathrm{M}}(\ell(\theta_d^{\mathrm{M}}) - \ell(\theta_g^{\mathrm{M}}))$. Therefore Welfare Test 1 is unchanged in content: the data market weakly Pareto dominates the data ban if and only if this private-version price is no higher than the data-ban price.

Finally, consider the negative-privacy results in Section 5. When ℓ is decreasing, IC and OB are aligned. Since ϕ_d^ℓ is nondecreasing, replacing any positive pooled record by full revelation weakly raises the seller's payoff, exactly as in the affine proof. Hence Proposition 5 survives: sold records fully reveal type, and the data-sale region is an upper interval $[\theta_d, 1]$. The seller-profit comparisons in Proposition A8 go through after replacing ϕ_d with ϕ_d^ℓ. The welfare result also has the same structure. The data market weakly Pareto dominates the data ban, and the comparison between the data market and data anarchy has a single-crossing form. In the case in which data sales depress product sales, the affine condition for the data market to Pareto dominate anarchy becomes $\theta_g^{\mathrm{A}} - \theta_g^{\mathrm{M}} + \int_{\theta_g^{\mathrm{A}}}^{\theta_d^{\mathrm{M}}} s\ell'(s)\,ds \geq 0$. Welfare Test 2 is unchanged in content: the data market weakly Pareto dominates data anarchy if and only if the cheapest product version in the data-market menu is no more expensive than the cheapest product version under data anarchy.

B.2 More General Third-Party Payoffs

This section shows that the results extend beyond the quadratic third-party payoff used in the main text. Let $\mu \in \Delta([0,1])$ denote the third party's posterior belief about the consumer's type, let $\mathbb{E}_\mu[\theta]$ denote its posterior mean, let $a^\star(\mu)$ denote the third party's

optimal action, and let $V(\mu)$ denote the associated indirect value.

Assumption 2. *The third party's payoff is such that:*

- *for every posterior belief, the optimal action is unique and is a continuous, strictly increasing, and concave function* $h : [0,1] \to [0,1]$ *of the posterior mean, with* $h(0) = 0$*;*
- *the indirect value is a continuous and strictly convex function* $W : [0,1] \to \mathbb{R}_+$ *of the posterior mean, with* $W(0) = 0$ *and* $W(z) > 0$ *for every* $z > 0$*.*

Example. Assumption 2 is satisfied by the following class of primitives. Suppose the third party's payoff is $v_c(\theta, a) = \lambda\big(\theta a - c(a)\big)$, where $a \in [0,1]$, $\lambda > 0$, and c is three times continuously differentiable, $c(0) = c'(0) = 0$, $c'(1) \geq 1$, $c'' > 0$, and $c''' \geq 0$ on $[0,1]$. For a posterior with mean z, the third party solves $\max_a \lambda(za - c(a))$, so the unique optimal action is $h(z) = (c')^{-1}(z)$, which is well defined on $[0,1]$ because $c'(0) = 0$, $c'(1) \geq 1$, and $c'' > 0$. The indirect value is $W(z) = \lambda\big(zh(z) - c(h(z))\big)$. The envelope theorem gives $W'(z) = \lambda h(z)$, so W is strictly convex. Differentiating h gives

$$h'(z) = \frac{1}{c''(h(z))} > 0, \qquad h''(z) = -\frac{c'''(h(z))}{[c''(h(z))]^3} \leq 0,$$

so h is increasing and concave. Also, $W(0) = 0$ and $W(z) > 0$ for every $z > 0$, since a small positive action yields a strictly positive payoff whenever the posterior mean is strictly positive. The quadratic specification in the paper is the special case $\lambda = 2$ and $c(a) = a^2/2$, for which $h(z) = z$ and $W(z) = z^2$.

For later use, define $\rho(z) \triangleq W(z)/h(z)$ for $z > 0$, and extend ρ continuously to $z = 0$. Under Assumption 2, ρ is increasing. Indeed, strict convexity of W and $W(0) = 0$ imply that $W(z)/z$ is increasing, while concavity of h and $h(0) = 0$ imply that $h(z)/z$ is decreasing. Hence $\rho(z) = [W(z)/z][z/h(z)]$ is increasing on $(0,1]$. For the class of primitives above, writing $z = c'(a)$,

$$\frac{1}{\lambda}\rho(z) = c'(a) - \frac{c(a)}{a} = \frac{1}{a}\int_0^a sc''(s)\,ds,$$

which is increasing in a because c'' is nondecreasing. More explicitly, using $a = h(z)$,

$$\rho'(z) = \lambda\left[1 - \frac{zh(z) - c(h(z))}{h(z)^2 c''(h(z))}\right],$$

and

$$zh(z) - c(h(z)) = \int_0^{h(z)} sc''(s)\,ds \leq \frac{1}{2}h(z)^2 c''(h(z)).$$

Thus $\rho'(z) \geq \lambda/2$ for this class. In particular, under the paper's normalization $\lambda = 2$, we have $\rho'(z) \geq 1$, and in the quadratic case $\rho(z) = z$.

For later use, also write $k(\theta) \triangleq \alpha + \beta\phi_g(\theta)$ and $\psi(\theta) \triangleq \rho(\theta) - k(\theta)$. The function ψ is the generalized virtual data surplus. In the quadratic baseline, $\psi = \phi_d$.

Section 3. Under Assumption 2, the obedience constraint becomes

$$\tilde{a} = h(\mathbb{E}[\theta \mid a(\theta) = \tilde{a},\ x(\theta) = 1])$$

for every positive on-path recommendation $\tilde{a} > 0$. The characterization of implementable outcomes in Remark 1 is unchanged after replacing the baseline obedience condition with the one above. The consumer-side part of Lemma A1 is also unchanged. Thus IC and IR imply that there is a cutoff θ_g such that $x(\theta) = \mathbb{1}\{\theta \geq \theta_g\}$, that $\beta a(\theta)$ is non-increasing on the trading region, and that transfers are pinned down by Equation (A1).

Proposition 1 therefore extends without change in its qualitative content. If $\beta > 0$, IC requires $a(\theta)$ to be non-increasing over participating types. Since h is strictly increasing, obedience orders positive recommendations in the same direction as posterior means, so at most one positive data-utilization level can be sold. If $\beta \leq 0$, IC and OB are aligned, and full disclosure is feasible: the record generated by type θ induces posterior mean θ and action $a(\theta) = h(\theta)$.

Lemma A2 changes only in the seller's data-revenue term. Fix an IC, IR, and OB outcome (x, t, a). The expected transfer paid by consumers is still

$$\int_{\theta_g}^{1} \left[\phi_g(\theta) - a(\theta)\left(\alpha + \beta\phi_g(\theta)\right)\right] f(\theta)\, d\theta,$$

because this calculation uses only the consumer's envelope formula. Define

$$\Gamma(\tilde{a}) \triangleq \begin{cases} 0, & \tilde{a} = 0, \\ W(h^{-1}(\tilde{a})), & \tilde{a} > 0. \end{cases}$$

The seller's expected payoff is therefore

$$\Pi_S = \int_{\theta_g}^{1} \left[\phi_g(\theta) - a(\theta)\left(\alpha + \beta\phi_g(\theta)\right) + \Gamma(a(\theta))\right] f(\theta)\, d\theta.$$

The third party's expected profit remains zero because the seller extracts the full value of each sold record through a take-it-or-leave-it price. Finally, Proposition 2 is unchanged:

$$U(\theta) = \begin{cases} 0, & \theta < \theta_g, \\ \displaystyle\int_{\theta_g}^{\theta} \left(1 - \beta a(s)\right) ds, & \theta \geq \theta_g. \end{cases}$$

Section 4. When $\beta > 0$, Proposition 3 extends directly after replacing the pooled action by $\bar{a} = h\big(\mathbb{E}[\theta \mid \theta_g \le \theta < \theta_d]\big)$. Thus sold data are pooled into a single record on an initial interval $[\theta_g, \theta_d)$ of the trading region. Data anarchy imposes $\theta_d^{\mathrm{A}} = 1$, data ban imposes $\theta_d^{\mathrm{B}} = \theta_g^{\mathrm{B}}$, and the data market allows any $\theta_d^{\mathrm{M}} \in [\theta_g^{\mathrm{M}}, 1]$. For example, in the data-market regime,

$$t^{\mathrm{M}}(\theta) = \begin{cases} 0, & \theta < \theta_g^{\mathrm{M}}, \\ \theta_g^{\mathrm{M}} - \bar{a}^{\mathrm{M}} \ell(\theta_g^{\mathrm{M}}), & \theta_g^{\mathrm{M}} \le \theta < \theta_d^{\mathrm{M}}, \\ \theta_g^{\mathrm{M}} + \bar{a}^{\mathrm{M}} \big(\ell(\theta_d^{\mathrm{M}}) - \ell(\theta_g^{\mathrm{M}})\big), & \theta \ge \theta_d^{\mathrm{M}}. \end{cases}$$

Since $\ell(\theta) = \alpha + \beta\theta$, the price of the private version is $p_{\mathrm{priv}}^{\mathrm{M}} = \theta_g^{\mathrm{M}} + \beta \bar{a}^{\mathrm{M}} (\theta_d^{\mathrm{M}} - \theta_g^{\mathrm{M}})$.

For the welfare comparison with data anarchy, the small-β assumption in Footnote 12 is replaced by the following condition: ρ is absolutely continuous and

$$\beta < \frac{\inf_{z \in (0,1)} \rho'(z)}{1 + L}, \tag{A2}$$

where L denotes an upper bound on the absolute value of the slope of the inverse hazard rate $(1 - F)/f$. In the quadratic baseline, (A2) reduces to the assumption in Footnote 12, $\beta < 1/(1 + L)$.

The seller's data payoff from a pool $I = [\theta_g, \theta_d)$ is

$$\big(F(\theta_d) - F(\theta_g)\big) h(z_I) \left[\rho(z_I) - \mathbb{E}_I k(\theta)\right], \qquad z_I \triangleq \mathbb{E}_I \theta.$$

Hence the binary-optimality argument used for Proposition 4 requires the bracketed term to be increasing as the pool expands upward. Condition (A2) is sufficient. To see this, fix θ_g and write $m(d) = F(d) - F(\theta_g)$, $z(d) = \mathbb{E}[\theta \mid \theta_g \le \theta < d]$, and $\bar{k}(d) = \mathbb{E}[k(\theta) \mid \theta_g \le \theta < d]$. Then

$$\frac{d}{dd}\big(\rho(z(d)) - \bar{k}(d)\big) = \frac{f(d)}{m(d)} \left[\rho'(z(d))(d - z(d)) - (k(d) - \bar{k}(d))\right].$$

Since $k'(\theta) = \beta \phi_g'(\theta)$ and $\phi_g'(\theta) \le 1 + L$, we have $k(d) - \bar{k}(d) \le \beta(1 + L)(d - z(d))$. Condition (A2) therefore implies that this derivative is nonnegative. If the bracketed term is negative at some θ_d, selling no data is better; if it is nonnegative, expanding the pool to $\theta_d = 1$ weakly raises both the bracketed term and $m(d)h(z(d))$. Thus, for fixed θ_g, the seller optimally chooses either $\theta_d = \theta_g$ or $\theta_d = 1$. The seller-profit comparisons in Proposition A7 then follow from the same binary comparison between the data-ban and data-anarchy outcomes.

It follows that Proposition 4 and Welfare Test 1 extend under the same replacement of the pooled action, with (A2) replacing the small-β condition in Footnote 12 for the comparison with data anarchy. In particular, the data market weakly Pareto dominates the data ban if and only if $p^{\mathrm{M}}_{\mathrm{priv}} \leq p^{\mathrm{B}}$. If the data-market outcome has no positive-measure private version, this condition should be read as the corresponding high-type utility comparison, or equivalently as a privacy-adjusted price comparison.

Section 5. Now suppose $\beta < 0$. The additional content needed for Proposition 5 is that positive pooled records can be unpooled profitably. Consider a positive pooling interval I, let $z_I = \mathbb{E}_I\theta$, and compare the pooled record with full revelation on I. The seller's payoff gain from full revelation is

$$\Delta(I) = \mathbb{E}_I[W(\theta)] - W(z_I) - \mathbb{E}_I[(h(\theta) - h(z_I))k(\theta)]\,.$$

The first term is nonnegative by convexity of W. The second expectation is nonpositive. Indeed, h is increasing and concave, while k is decreasing because ϕ_g is increasing and $\beta < 0$. Moreover, k is nonnegative because $k(\theta) = \ell(\theta) - \beta(1 - F(\theta))/f(\theta)$ and $\beta < 0$; equivalently, $\phi_g(\theta) \leq \theta$ implies $k(\theta) = \alpha + \beta\phi_g(\theta) \geq \alpha + \beta\theta = \ell(\theta) \geq 0$. Formally,

$$\mathbb{E}_I[(h(\theta) - h(z_I))k(\theta)] = \mathrm{Cov}_I(h(\theta), k(\theta)) + \big(\mathbb{E}_I h(\theta) - h(z_I)\big)\mathbb{E}_I k(\theta) \leq 0,$$

where the covariance term is nonpositive by opposite monotonicity and the second term is nonpositive by concavity of h. Therefore $\Delta(I) \geq 0$, with strict inequality for every nondegenerate interval under Assumption 2. Thus an optimal outcome can be selected in which every positive record is fully revealing, and the third party takes action $a(\theta) = h(\theta)$ whenever data are sold.

In the data-market regime, the corresponding transfers are

$$t^{\mathrm{M}}(\theta) = \begin{cases} 0, & \theta < \theta^{\mathrm{M}}_g, \\ \theta^{\mathrm{M}}_g, & \theta^{\mathrm{M}}_g \leq \theta < \theta^{\mathrm{M}}_d, \\ \theta^{\mathrm{M}}_g - h(\theta)\ell(\theta) + \beta \displaystyle\int_{\theta^{\mathrm{M}}_d}^{\theta} h(s)\,ds, & \theta \geq \theta^{\mathrm{M}}_d. \end{cases}$$

Equivalently,

$$U^{\mathrm{M}}(\theta) = \begin{cases} 0, & \theta < \theta^{\mathrm{M}}_g, \\ \theta - \theta^{\mathrm{M}}_g, & \theta^{\mathrm{M}}_g \leq \theta < \theta^{\mathrm{M}}_d, \\ \theta - \theta^{\mathrm{M}}_g - \beta \displaystyle\int_{\theta^{\mathrm{M}}_d}^{\theta} h(s)\,ds, & \theta \geq \theta^{\mathrm{M}}_d. \end{cases}$$

Under full revelation, the seller's incremental data payoff from type θ is $W(\theta) - h(\theta)k(\theta) = h(\theta)\psi(\theta)$. Under Assumption 2, ψ is strictly increasing when $\beta < 0$, because ρ is increasing and k is decreasing. Hence the sold-data region is an upper interval. The data-market problem becomes

$$\max_{\theta_g \le \theta_d} \left\{ \int_{\theta_g}^{1} \phi_g(\theta) f(\theta)\, d\theta + \int_{\theta_d}^{1} h(\theta)\psi(\theta) f(\theta)\, d\theta \right\}.$$

Data anarchy imposes $\theta_d^{\mathrm{A}} = \theta_g^{\mathrm{A}}$, and data ban imposes $\theta_d^{\mathrm{B}} = 1$. This is the same cutoff problem as in Section 5 after replacing ϕ_d by ψ and $\theta\phi_d(\theta)$ by $h(\theta)\psi(\theta)$.

The cutoff rankings and seller-profit comparisons therefore extend verbatim with that replacement. In particular, letting θ_g^* solve $\phi_g(\theta_g^*) = 0$, if $\psi(1) \le 0$, then no data sale is profitable and the data-market outcome coincides with the data-ban outcome. If $\psi(\theta_g^*) \ge 0$, then the data-market and data-anarchy outcomes coincide: $\theta_g^{\mathrm{M}} = \theta_g^{\mathrm{A}} = \theta_d^{\mathrm{M}} \le \theta_g^{\mathrm{B}}$. Finally, if $\psi(\theta_g^*) < 0 < \psi(1)$, then the data-market outcome contains a private option and, generically, $\theta_g^{\mathrm{M}} = \theta_g^{\mathrm{B}} < \theta_g^{\mathrm{A}} < \theta_d^{\mathrm{M}}$, where θ_d^{M} is the unique zero of ψ.

Proposition 6 also extends. The envelope slopes are $q^{\mathrm{B}}(\theta) = \mathbb{1}\{\theta \ge \theta_g^{\mathrm{B}}\}$, $q^{\mathrm{A}}(\theta) = \mathbb{1}\{\theta \ge \theta_g^{\mathrm{A}}\}\big(1 - \beta h(\theta)\big)$, and $q^{\mathrm{M}}(\theta) = \mathbb{1}\{\theta \ge \theta_g^{\mathrm{M}}\} - \beta h(\theta)\mathbb{1}\{\theta \ge \theta_d^{\mathrm{M}}\}$. The same single-crossing argument gives that the data market weakly Pareto dominates the data ban, and that the comparison between the data market and data anarchy is cutoff-shaped. In the case $\psi(\theta_g^*) < 0$, the condition for the data market to weakly Pareto dominate data anarchy is

$$\theta_g^{\mathrm{A}} - \theta_g^{\mathrm{M}} + \beta \int_{\theta_g^{\mathrm{A}}}^{\theta_d^{\mathrm{M}}} h(s)\, ds \ge 0.$$

Finally, the observable condition in Welfare Test 2 extends with the same boundary qualification as in the alternative formulation. If $\theta_d^{\mathrm{M}} < 1$, the highest types receive the same downstream treatment under the data market and data anarchy: their records are sold, fully reveal their types, and induce action $h(\theta)$. Hence the data market weakly Pareto dominates data anarchy if and only if the cheapest fully revealing product version in the data-market menu is no more expensive than the cheapest fully revealing product version under data anarchy; equivalently, $t^{\mathrm{M}}(1) \le t^{\mathrm{A}}(1)$. If $\theta_d^{\mathrm{M}} = 1$, the data-market outcome has no positive-measure fully revealing version. In that boundary case, the raw price comparison must be replaced by the privacy-adjusted comparison

$$t^{\mathrm{M}}(1) + a^{\mathrm{M}}(1)\ell(1) \le t^{\mathrm{A}}(1) + h(1)\ell(1),$$

where $a^{\mathrm{M}}(1) = 0$ if the highest type buys a private version under the data market. When

$\theta_d^{\mathrm{M}} < 1$, $a^{\mathrm{M}}(1) = h(1)$, so the privacy terms cancel and the condition reduces to the raw price comparison.

C Other Materials

C.1 Equilibrium Definitions

In this appendix we spell out an extensive-form version of the games introduced in Section 2, define the equilibrium, and discuss some of the assumptions we have made.

C.2 Data anarchy regime

Given a consumer strategy $\sigma : \Theta \to M$, a realized record m identifies the subset of types $\Theta(m) \triangleq \sigma^{-1}(m) = \{\theta \in \Theta : \sigma(\theta) = m\}$. Selling record m means that the third party gets access to the consumer and verifiably learns that $\theta \in \Theta(m)$. Fix a first-period mechanism (x, t). After observing the mechanism, the consumer sends a message $m \in M$. The outcome $(m, x(m), t(m))$ is observed by the consumer and the seller but not by the third party. If $x(m) = 0$, the game ends and there is no data sale. If $x(m) = 1$, the seller observes the realized verifiable record m and chooses whether to sell it. Let $d : M \to \{0, 1\}$ denote the seller's continuation strategy, where $d(m) = 1$ means that record m is sold and $d(m) = 0$ means that it is withheld. By convention, $d(m) = 0$ whenever $x(m) = 0$. If $d(m) = 0$, no data are sold and the third party's action is exogenously set to 0. If $d(m) = 1$, the third party verifiably learns m and chooses an action in A to maximize her payoff. The induced equilibrium action is $a(m) \triangleq \mathbb{E}(\theta \mid \Theta(m))$.

Given a sold record m, the third party's willingness to pay is

$$W(m) \triangleq \max_{\hat{a} \in A} \mathbb{E}\left(v(\theta, \hat{a}) \mid \Theta(m)\right) = \frac{\delta}{2} a(m)^2.$$

Since the seller can make a take-it-or-leave-it offer, seller optimality implies that if record m is sold then the price is set equal to $W(m)$. Accordingly, we can omit the price as an equilibrium choice variable and write the seller's second-period continuation payoff directly as $W(m)$.

Definition 1. *Consider the data anarchy regime. Fix an indirect mechanism (x, t). A continuation perfect Bayesian equilibrium given (x, t) is a tuple (σ, d, a) consisting of a*

consumer strategy $\sigma : \Theta \to M$*, a seller continuation strategy* $d : M \to \{0,1\}$*, and a third party action* $a : M \to \alpha(A)$*, such that:*

1. *For every* θ*,*

$$\sigma(\theta) \in \arg\max_{m \in M} \{x(m)\left[\theta + d(m)(\beta\theta - \alpha)a(m)\right] - t(m)\}.$$

2. *For every* m *such that* $x(m) = 1$*,*

$$d(m) \in \arg\max_{\hat{d} \in \{0,1\}} \hat{d}\, W(m).$$

3. *For every* m *such that* $x(m)d(m) = 1$*,* $a(m) = \mathbb{E}(\theta \mid \sigma(\theta) = m)$*.*

The belief system is trivial given that data records are verifiable and thus it is omitted from the definition. Additionally, because the action after no sale is exogenously fixed at 0*, beliefs following no sale are irrelevant and omitted from the definition.*

A richer set of verifiable statements about the data record m**.** The same equilibrium outcomes would obtain in a richer verifiable-disclosure model in which, after observing m, the seller could send any verifiable statement $E \subseteq M$ such that $m \in E$. Such a statement would reveal only that the realized record belongs to E, and therefore would generate some continuation value $W(E)$. These continuation values can be ordered, and this ranking is independent of the true m: Thus, conditional on feasibility, the seller always prefers the statement that induces the highest value. For the same reason as in standard unraveling arguments, coarse verifiable statements cannot be sustained, except possibly for payoff-equivalent pooling across records that generate the same value. Thus this richer model delivers the same outcome as full disclosure of m.

C.2.1 Unverifiable Data Records

In the data anarchy regime, we assume that the seller can transfer to the third party verifiable data records. Verifiability means that if the seller has obtained a record m through her interaction with the consumer, she can disclose m or remain silent, but she cannot instead report a different record m'. In particular, she cannot fabricate data that were not actually generated in the upstream transaction. This assumption matters only in the data anarchy regime. Under a data ban, records are never traded and therefore

verifiability is irrelevant. Under a data market, instead, the seller commits ex ante to a data policy, so that the relevant disclosures are effectively verifiable by construction.

It is useful to briefly consider what happens if the verifiability assumption is relaxed. If communication between the seller and the third party is unverifiable, then after observing any consumer the seller has an incentive to send the message that induces the highest posterior mean. The reason is immediate: the third party's continuation payoff is increasing in the posterior mean, and so is the seller's payoff from downstream data monetization. Thus, regardless of the record actually generated in the interaction, the seller would always like to portray it as if it came from the highest pool of consumers. But then disclosure loses credibility altogether. In the continuation game between the seller and the third party, the only equilibrium is babbling, so that the seller cannot credibly transmit any informative record. As a result, in the data anarchy regime the optimal mechanism collapses to a posted price.

This observation sharpens the comparison across regimes. When $\beta > 0$, nothing changes relative to our baseline analysis: in the data anarchy regime the optimal mechanism is already a posted price, while the solutions under a data ban and under a data market are unchanged. The more interesting case is $\beta < 0$. There, under verifiable disclosure, the data anarchy may sustain a continuum of differentiated products through privacy-based versioning. By contrast, if records are unverifiable, this possibility disappears: the data anarchy regime also collapses to a posted price, while the allocations under a data ban and under a data market remain unchanged. Hence, privacy-based versioning can arise only when privacy concerns are decreasing and the seller has access to a data market with commitment.

The economic intuition is that commitment in the data market resolves two distinct credibility problems at once. First, as emphasized in the main text, it resolves the seller's lack of credibility vis-à-vis the consumer. Second, when records are otherwise unverifiable, it also resolves the seller's lack of credibility vis-à-vis the third party. For this reason, the introduction of a data market becomes even more consequential when raw data records are not independently verifiable. At the same time, focusing on the verifiable benchmark remains useful because it isolates the first friction without conflating it with the second. Still, the unverifiable case is informative: it shows that the paper's main qualitative message not only survives without verifiability, but in some respects becomes even stronger.

C.3 Equilibrium Definition under Commitment

Fix a first-period mechanism (x, t). We maintain the baseline assumption that realized records are verifiable. Thus, if record m is sold, the third party verifiably learns m and therefore chooses $a(m) \triangleq \mathbb{E}(\theta \mid \sigma(\theta) = m)$. As before, let $W(m) \triangleq \frac{\delta}{2} a(m)^2$ denote the third party's willingness to pay for record m.

The only difference relative to the data anarchy regime case is that the seller now chooses a deterministic disclosure rule $d : M \to \{0, 1\}$ ex ante, before the buyer sends a message.

Definition 2. *Consider the data market. Fix an indirect mechanism (x, t). A continuation perfect Bayesian equilibrium given (x, t) is a tuple (σ, d, a) consisting of a consumer strategy $\sigma : \Theta \to M$, a disclosure strategy $d : M \to \{0, 1\}$, and a third party action $a : M \to \alpha(A)$, such that:*

1. *For every θ,*

$$\sigma(\theta) \in \operatorname*{argmax}_{m \in M} \{x(m) [\theta + d(m)(\beta\theta - \alpha) a(m)] - t(m)\}.$$

2. *The seller's disclosure rule is ex ante optimal:*

$$d \in \operatorname*{argmax}_{\hat{d}: M \to \{0,1\}} \Pi(\hat{d}, \sigma_{\hat{d}}),$$

where $\sigma_{\hat{d}}$ is a buyer best response to $\hat{d}$, and

$$\Pi(\hat{d}, \sigma_{\hat{d}}) \triangleq \int_{\Theta} \left[t(\sigma_{\hat{d}}(\theta)) + x(\sigma_{\hat{d}}(\theta)) \hat{d}(\sigma_{\hat{d}}(\theta)) W(\sigma_{\hat{d}}(\theta)) \right] f(\theta) d\theta.$$

3. *For every m such that $x(m)d(m) = 1$, $a(m) = \mathbb{E}(\theta \mid \sigma(\theta) = m)$.*

Beliefs following no sale are irrelevant because the third party's action is exogenously fixed at 0.

Deterministic Disclosure Rules. Suppose that, instead of committing to a binary disclosure rule $d : M \to \{0, 1\}$, the seller could commit ex ante to a richer *deterministic* disclosure rule $\pi : M \to \text{A}$, where $0 \in \text{A}$ denotes no disclosure. Under such a rule, the third party observes only the realized signal $s \in \text{A}$ and therefore updates about the underlying record m from the induced partition of M. We claim that allowing this richer set of verifiable statements does not change the set of equilibrium outcomes. To see why,

fix a deterministic map $\pi : M \to \text{A}$. If two records m_1, m_2 are mapped to the same signal and both satisfy $x(m_1) = x(m_2) = 1$, then they induce the same aftermarket treatment. If $t(m_1) \neq t(m_2)$, the record with the higher transfer is weakly dominated for every type, so an IC mechanism cannot place positive probability on both records on path. Hence, whenever m_1 and m_2 are used on path and $\pi(m_1) = \pi(m_2)$, it must be that $t(m_1) = t(m_2)$. In that case one can quotient the mechanism by the partition induced by π, replacing m_1 and m_2 with a single coarser record without changing allocations, transfers, or the induced aftermarket action.

This deterministic-disclosure formulation implicitly maintains the assumption that records are verifiable, but enlarges the set of verifiable statements: after observing $m \in M$, the seller can verifiably assert membership in any event $E \subset M$ such that $m \in E$. The argument above shows that this added generality is without consequence for the seller-optimal equilibrium outcome.

Importantly, the seller-optimal equilibrium outcome would change if we considered stochastic disclosure rules, $\pi : M \to \Delta\text{A}$.