

October 2026

Data Sharing, Privacy, and Innovation in Healthcare

Chiara Canta, Anna D'Annunzio and Yassine Lefouili

Data Sharing, Privacy, and Innovation in Healthcare*

Chiara Canta[†]

Anna D’Annunzio[‡]

Yassine Lefouili[§]

October 2026

Abstract

We study the impact of privacy regulation on data-driven innovation and welfare in the healthcare sector. A platform providing medical services (e.g., a hospital) collects patient data and may sell non-rival access to downstream innovators. Without data externalities, consent can induce inefficiently low data sharing because the platform cannot identify privacy-concerned patients and captures only part of the social value of innovation. Consequently, requiring consent can reduce welfare for intermediate privacy costs. When one patient’s disclosure reveals information about others, consent decisions become strategic complements. These externalities reduce the compensation required for disclosure and can instead produce excessive sharing, potentially justifying additional restrictions on compensation or data sharing. If the platform controls innovator access, consent can induce excessive restrictions without data externalities, whereas sufficiently strong externalities can lead the platform to admit too many innovators.

Keywords: Data sharing; privacy; consent policy; innovation; healthcare.

JEL Codes: D83, L15, L51.

*We thank Bertrand Chopard, Sébastien Mitraillé, and audiences at the Digital Health Workshop (Toulouse, 2024) and Paris Descartes seminar (2025). Funding from the Agence Nationale de la Recherche (ANR) under the TEPREME project (ANR-21-CE26-0014-01) and under grant ANR-17-EURE-0010 (Investissements d’Avenir program) is gratefully acknowledged.

[†]TBS Business School, Toulouse, France.

[‡]University of Rome - Tor Vergata. Other Affiliations: CSEF and Toulouse School of Economics.

[§]Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France.

1 Introduction

The healthcare sector is undergoing a deep transformation, driven by the proliferation of health platforms and the growth of health data collection. There is a wide range of data sources, ranging from traditional ones, such as hospitals and testing laboratories, to new products, such as health-related apps and wearable devices. Health data is collected by many different firms, either as part of their primary activity or as a byproduct of their core activities.

Health data is a critical input for R&D activities undertaken by pharmaceutical companies and medical-device manufacturers. This has led to the emergence of data-sharing partnerships between holders of medical data and healthcare innovators. For instance, several major U.S. hospital networks have established such partnerships. A well-known case is the Mayo Clinic, a large and prestigious institution treating 1.5 million patients a year. In 2020, the Mayo Clinic had 16 data-sharing partnerships (Ross, 2020). In this case, patients gave consent at the moment of treatment and were not financially compensated.¹ The Mayo Clinic went on to create a data platform allowing innovators to analyze secure medical data from multiple hospitals and health systems.²

However, besides the opportunities for innovation presented by the availability of health data, there are significant concerns regarding the protection of sensitive information. In Europe, the General Data Protection Regulation (GDPR) aims to address these concerns. The GDPR imposes stringent requirements on the handling of sensitive data, like health data. In particular, explicit consent is typically required for sharing health data for commercial purposes. Moreover, the European Union has recently created the European Health Data Space (EHDS), which is meant to establish a framework for the use and exchange of electronic health data (EU, 2025).³ The EHDS requires health data holders to make data available and enables health data users to

¹As Ross (2020) points out, “[...] compensating patients for their data raises a potential flashpoint with academic medical centers. Mayo executives said it could slow innovation and undermine development of new treatments and digital services — a need whose urgency has been reinforced by the Covid-19 pandemic.”

²See <https://www.mayoclinicplatform.org/>. Another similar platform is Epic Cosmos (see <https://cosmos.epic.com/>).

³The EHDS regulation will apply in stages starting 2027.

access it in secure processing environments, based on permits issued by Health Data Access Bodies (HDABs). For pharmaceutical and medtech companies, this represents a significant opportunity to legally access and utilize health data from EU citizens in a transparent and regulated manner. These opportunities are paired with strict compliance requirements, including duties to respect patient opt-out preferences, protect intellectual property, and ensure data security.

While privacy regulation is intended to protect consumer privacy, it may pose challenges for R&D activities in sectors that require access to data. A well-known example of the tension between privacy regulation and innovation is the DeepMind case. The Royal Free London NHS Foundation Trust shared data on 1.6 million patients so that the company could build an app for acute kidney injury detection. The UK Information Commissioner’s Office later found that the Trust had shared sensitive data without adequate transparency and without giving patients a proper choice.

In this paper, we study the interplay between data-driven innovation in the health-care sector and privacy regulation. We consider a platform that provides a medical service to patients and can collect their personal data in a setting where privacy costs are heterogeneous across patients. The platform can share this data with innovators in exchange for a fee. We also assume that patient data constitutes an input for welfare-enhancing innovation and is non-rival, meaning that it can be shared with multiple innovators without reducing its value to each of them.

We assume that a fraction of patients are privacy-concerned, while the others are privacy-unconcerned. We consider two regulatory regimes, one in which consent is required and one in which it is not. Under the consent regime, data can be collected and shared with third parties only with patients’ agreement. In this case, privacy-concerned patients agree to data sharing only if they receive financial compensation that is at least equal to their privacy cost. Under the no-consent regime, all data can be collected and shared by the platform irrespective of whether patients have agreed to this. We characterize the equilibrium outcomes under these two scenarios and investigate their welfare implications.

We first consider a baseline model with no data externalities. The desirability of the consent regime relative to the no-consent regime is driven by the comparison of

the choice made by the profit-maximizing platform and the corresponding welfare-maximizing choice of the regulator. There are cases where the platform’s strategy is perfectly aligned with a social planner’s choice. This holds if privacy costs are sufficiently low relative to the innovation value of data. In such a scenario, it is socially optimal for all patient data to be shared with innovators and, under the consent regime, the platform is willing to offer sufficient compensation to induce privacy-concerned patients to share their data. Hence, requiring that the platform secures its patients’ consent before sharing their data is inconsequential from a social welfare perspective; the only effect of such a policy is to generate a monetary transfer from the platform to the patients.

However, for intermediate values of privacy costs, the consent regime is dominated by the no-consent regime from a social perspective. The reason is that the former results in less data sharing than the socially desirable level because the platform lacks incentives to compensate patients for their privacy costs. This is due to two factors. First, the platform is not able to discriminate between privacy-concerned and privacy-unconcerned patients, which implies that the total monetary transfers made to patients to induce the privacy-concerned ones to share their data are larger than the total privacy costs that they incur. Second, the platform is unable to capture the whole social value of innovations.

Finally, we find that the consent regime dominates the no-consent regime when privacy costs are sufficiently high relative to the innovation value of the data. In this case, the platform does not find it optimal to compensate privacy-concerned patients for the use of their data, and the consent requirement serves to mitigate excessive data sharing that would otherwise arise under the no-consent regime.

We then examine the case with data externalities, i.e., where privacy-concerned patients incur a privacy cost also when other patients agree to data sharing. For instance, confidential data shared by some patients may be used to infer confidential information about other patients using exclusively their disclosed (non-confidential) information. An obvious example is genetic data, which can also reveal sensitive information about the biological relatives of patients who share their data. More generally, third parties could use shared data to infer health information about patients who never consented

to data collection, potentially leading to discrimination or privacy violations.

Because data externalities affect patients even if they do not share data, their willingness to accept data sharing increases and consent decisions become strategic complements. In this context, we show that, because of the lack of coordination among patients, under the consent regime the platform may obtain data sharing for a transfer that is lower than the social cost of sharing. If the effects of this coordination problem are relatively small, then the results are similar to those obtained without data externalities. However, when data externalities are large and the number of privacy-concerned patients is high, the coordination problem becomes more severe, and may lead to oversharing even under the consent regime. In this case, imposing consent may not be sufficient to reduce data sharing to its optimal level.

We also consider stricter regulation aimed at correcting for excessive sharing. We find that, if oversharing occurs under the consent regime, forbidding transfers from the platform to patients (i.e., imposing a consent-with-no-payment regime) may be optimal. A full prohibition of data sharing may also be a solution in extreme cases.

We finally extend the model to allow the platform to control the number of innovators accessing patient data. Under the no-consent regime, the platform admits all potential innovators. Under the consent regime, however, it may restrict access to reduce the compensation required to induce all patients to share their data. Consent therefore introduces an additional trade-off between lower privacy costs and lower innovation, even when all patients agree to share their data. Without data externalities, the platform may admit an inefficiently low number of innovators. Data externalities can lead to the opposite distortion: when they are sufficiently strong, the platform may admit too many innovators, so that excessive innovator access becomes an additional source of inefficiency under the consent regime.

The rest of the paper is organized as follows. Section 2 reviews the relevant literature. Section 3 presents the baseline model without data externalities and characterizes the socially optimal privacy regime. Section 4 extends the analysis to a setting with data externalities. Section 5 considers an extension in which the platform controls the number of innovators accessing patient data. Section 6 concludes.

2 Related literature

There is a vast literature on privacy (see, for instance, Acquisti et al., 2016 for a review of the literature). Our paper is related to the recent strand of that literature that studies the impact of privacy regulation on investment in innovation and quality.

Lefouili et al. (2024) analyze the impact of a privacy regulation taking the form of a cap on information disclosure on quality-enhancing innovation. They show that, if the share of privacy-concerned consumers is sufficiently small, privacy regulation has a negative effect on innovation and may harm users. However, if the share of privacy-concerned consumers is sufficiently large, privacy regulation may have a positive effect on innovation. In this case, there is no trade-off between privacy and innovation and users always benefit from privacy regulation. There are two main differences between this paper and ours. First, the former considers innovation that is not enabled by data, but instead aims to increase demand for the service, thereby increasing the amount of data collected. In contrast, we consider a model where data is an input for innovation. Second, the privacy regulations studied in the two papers are different.

Conti and Reverberi (2021) investigate how a privacy regulation that imposes an opt-in policy affects a firm's investment in quality in a model featuring price discrimination. The main trade-off for users is the following: while sharing information leads to a better consumption experience, it also allows the firm to offer consumers personalized prices based on their willingness to pay. The authors find that quality may be higher under an opt-out regime than under an opt-in regime if there are strong complementarities between information and quality. A major way in which this paper differs from ours is that, as Lefouili et al. (2024), the authors consider investment in quality-enhancing innovation that does not require data as an input.

There are few papers analyzing the relationship between privacy and investment in healthcare. Conti et al. (2024) consider the interplay between pharmaceutical innovation and privacy regulation in the context of precision medicine. They show that there is underinvestment in targeted treatments and that privacy protection measures reducing the perceived risk of sharing personal data can lead to more innovation and higher social welfare. Conti et al. (2026) generalize the above model along many dimensions.

In particular, they analyze how a public agency that manages access to health data can address market failures. A key difference between these papers and ours is that they consider a (monopolistic) pharmaceutical firm that directly collects data from users and makes investment decisions, while we consider a two-sided platform that collects patient data and chooses to share it with innovators, in a context where (some) patients suffer negative network externalities from innovators because of their taste for privacy.

The impact of privacy regulation on innovation has also been studied empirically. Goldfarb and Tucker (2012) analyze several empirical studies in the healthcare and the online advertising sectors, finding that privacy regulations may raise the costs and/or lower the benefits associated with data-driven innovation, hence weakening firms' investment incentives. More specifically, Miller and Tucker (2009) find that the adoption of Electronic Medical Records (EMRs) by health providers is affected by privacy legislation. The latter reduces the adoption of EMRs, and consequently the positive network effects deriving from exchanging information with other health providers through EMRs. Miller and Tucker (2012) find that increased information transparency reduced adoption of EMRs in states with more malpractice litigation. Miller and Tucker (2018) study how different genetic privacy laws may affect data disclosure. Exploiting variation in U.S. state laws over time, they find that giving individuals control over the redisclosure of their genetic information encourages genetic testing, unlike notification of potential privacy risks.

Finally, this paper is related to the literature on data externalities. Choi et al. (2019) highlight that when a firm learns about non-participants from participants' data, individuals ignore the privacy costs imposed on others, leading to excessive data collection. Acemoglu et al. (2022) show that similar externalities depress the price of data in competitive data markets and generate too much data sharing relative to the social optimum. Ichihashi (2021) examines how data externalities affect surplus division between a firm and consumers, as well as the firm's incentives to acquire information. He shows that, when the firm can choose what information to purchase from consumers, it acquires weakly more information than what is socially efficient. Bergemann et al. (2022) emphasize the social nature of data: each user's data contain information about "similar" users, creating economies of scale in data acquisition for

intermediaries. They show that a data broker can partially internalize externalities through aggregation and anonymization and, under certain conditions, implement outcomes closer to the social optimum. Overall, this literature shows that laissez-faire data markets usually fail to internalize informational spillovers and provides a foundation for regulatory tools such as consent rules, anonymization requirements, and mediated data sharing.

3 Baseline model

We consider a platform operating as a monopolist, connecting two groups of agents: patients and innovators. The platform provides a medical service to patients and can collect their personal data generated as a byproduct of the consumption of the medical service. We assume that some patients are privacy-concerned, and that data sharing can occur under two different regimes depending on whether patient consent is required or not. Once collected by the platform, data can be shared with innovators, who use it as an input in the innovation process to develop new drugs, medical devices, diagnostic tools, and treatments. These innovations have a social value that can be partially captured by innovators. In this section we consider a baseline setup with no data externalities. In Section 4 we consider a scenario with data externalities, where data shared by patients allow innovators to infer information about other patients who did not share their data.

3.1 Setup

Platform. The platform provides a medical service to patients at an exogenous price p for a unit cost c . The platform can collect and share data with innovators under a consent or a no-consent regime. If patient consent is required, the platform offers a transfer $t \geq 0$ to patients who agree to their data being collected and shared with innovators. If patient consent is not required, the platform can collect and share patients' data without making any transfer to them, i.e., $t = 0$. In both cases, collected data is shared with the innovators joining the platform in exchange for a fixed (membership) fee f .

Patients. There is a unit mass of patients. We assume that all patients have the same valuation v for the medical service offered by the platform. We further suppose that v is high enough to ensure that all patients join the platform irrespective of the privacy regime. This assumption is characteristic of the healthcare context, where health treatments may have a strong impact on patients' well-being, and can even be life-saving.

If their data is shared, patients may suffer a privacy cost. We assume that there is a share $\lambda \in (0, 1)$ of patients who are privacy-concerned and incur a privacy cost $D(n_I)$ when their data is shared with n_I innovators.⁴ We suppose that this cost is increasing in the number of innovators (i.e., $D' > 0$). The remaining share $1 - \lambda$ of patients are privacy-unconcerned, meaning that they incur no privacy cost when their data is shared with third parties.

The utility of privacy-unconcerned patients is $v - p + t$ because they always accept that their data is shared. The utility of privacy-concerned patients depends on whether their data is shared: it is equal to $v - p$ if data is not shared, and to $v - p + t - D(n_I)$ if it is.

Innovators. There are $n_I > 0$ symmetric innovators. They value patients' personal data because the latter is an input for innovation. Specifically, the probability that a given innovator succeeds is denoted by $q(n_S) \in [0, 1]$ where n_S is the number (mass) of patients whose data is shared with innovators. We suppose that $q(0) = 0$, $q'(n_S) > 0$ and $q''(n_S) \geq 0$.⁵ We further assume that innovators' probabilities of success are independent, and that each successful innovation generates the same additive social value V (regardless of the number of successful innovators). This assumption allows us to abstract from externalities among innovators (arising from competition or complementarities), and focus instead on the externalities between innovators and patients generated by data sharing. The expected welfare generated when n_I innovators obtain access to the data of n_S patients is $n_I q(n_S) V$. Moreover, assuming that each innovator

⁴In the baseline setup, there is no privacy cost for patients whose data is not shared. This assumption will be relaxed in Section 4.

⁵Weak convexity includes the linear case in which additional patient data have a constant marginal effect on the probability of innovation. It also allows for increasing marginal effects, which can arise from complementarities between patient data.

captures a profit $0 < R < V$ from its innovation, the expected profit of an innovator is $q(n_S)R - f$.⁶

Regulator. The regulator chooses the privacy policy to maximize social welfare, that is, the sum of the social surplus generated by the medical treatment and the social surplus generated by innovation net of privacy costs. The privacy policy is such that the regulator chooses whether patients' consent is required for their data to be collected by the platform and shared with innovators. These privacy regimes are denoted by superscripts C (Consent) and NC (No Consent), respectively.

Timing. The timing of the game is as follows:

1. The regulator sets the privacy regime, which can be either the consent regime or the no-consent regime.
2. The platform sets a payment t for patients who agree to their data being shared and sets a fixed fee f for innovators who wish to access patient data.
3. Innovators decide whether to pay the fixed fee in exchange for data.
4. Under the consent regime, each patient decides whether to consent to the sharing of their data; under the no-consent regime, the platform can collect and share data from all patients.

3.2 First-best

Before solving the game, we characterize the first-best, obtained by maximizing social welfare assuming that the regulator directly controls each patient's sharing decision.

Note that for a given number n_S of patients sharing their data, it is socially optimal to prioritize the sharing of data by privacy-unconcerned patients. Specifically, if $n_S \leq 1 - \lambda$, all sharing patients are privacy-unconcerned. If $n_S > 1 - \lambda$, then all privacy-unconcerned patients share their data, and the remaining sharing patients are privacy-concerned. The reason is that sharing the data of a privacy-unconcerned patient entails

⁶We assume that there are no other sources of patient data.

a lower social privacy cost than sharing the data of a privacy-concerned patient, while generating the same social innovation benefit. Taking this into account, social welfare can be written as follows:

$$W(n_I, n_S) = \begin{cases} v - c + n_I q(n_S) V & \text{if } n_S \in [0, 1 - \lambda], \\ v - c + n_I q(n_S) V - (n_S - 1 + \lambda) D(n_I) & \text{if } n_S \in (1 - \lambda, 1]. \end{cases}$$

When no privacy-concerned patients share their data (i.e., $n_S \in [0, 1 - \lambda]$), social welfare is the sum of the benefit of the treatment for patients net of its cost ($v - c$), and the expected social value of innovation when the data of n_S patients is shared with the n_I innovators ($n_I q(n_S) V$). Note that social welfare is increasing in n_S over $[0, 1 - \lambda]$: in the absence of data externalities, sharing the data of additional privacy-unconcerned patients increases the probability of innovation at no social cost. Therefore, a first-best outcome necessarily involves sharing the data of all privacy-unconcerned patients, implying that $n_S \geq 1 - \lambda$.

When at least some privacy-concerned patients share their data (i.e., $n_S \in (1 - \lambda, 1]$), social welfare also includes the privacy costs $(n_S - 1 + \lambda) D(n_I)$ that they incur. Since the innovation probability $q(\cdot)$ is weakly convex, social welfare is also weakly convex over $(1 - \lambda, 1]$. This implies that the first-best outcome is achieved by only privacy-unconcerned patients sharing their data (i.e., $n_S = 1 - \lambda$) or all patients sharing their data (i.e., $n_S = 1$).⁷ In the former case, social welfare is given by

$$W_{\text{partial}} = v - c + n_I q(1 - \lambda) V,$$

while in the latter case it is given by

$$W_{\text{all}} = v - c + n_I q(1) V - \lambda D(n_I).$$

⁷If these two allocations generate the same welfare, both are socially optimal. Also, if $q(n_S)$ is linear on $[1 - \lambda, 1]$, then social welfare is linear in n_S on that interval, with slope $n_I q' V - D(n_I)$. If $q' = \alpha$ and $n_I \alpha V = D(n_I)$, then any n_S in the interval is also optimal. We ignore this case, as one of the two corner allocations is still optimal.

It is (strictly) socially optimal that all patients share their data if and only if

$$W_{all} > W_{partial} \iff \frac{\lambda D(n_I)}{n_I} < V\Delta q, \quad (1)$$

where $\Delta q := q(1) - q(1 - \lambda)$. The condition above means that the per innovator privacy cost of sharing the data of privacy-concerned patients (i.e., $\lambda D(n_I)/n_I$) is smaller than the increase in the expected value of an innovation generated by making those data available to innovators (i.e., $V\Delta q$). If the reverse strict inequality holds, the first best is such that only privacy-unconcerned patients share their data.

Condition (1) can be rewritten as

$$n_I \Delta q > A_0 := \lambda \frac{D(n_I)}{V}. \quad (2)$$

The first-best can be achieved if the regulator knows the privacy cost of each patient, and can control the number and type of patients sharing data. We assume throughout the rest of the section that the regulator knows only the distribution of privacy costs but can impose restrictions on the sharing of health data. This assumption is consistent with the recent establishment of the European Health Data Space (EHDS) by the European Union (EU, 2025). The EHDS is meant to establish a European framework for the use and exchange of electronic health data, both for primary (e.g., access for healthcare delivery) and secondary use (e.g., research and innovation). The EHDS enables the regulator to control access to health data for research purposes, even if the regulator may not have precise information on the privacy preferences of patients.

3.3 Equilibrium analysis

In the baseline model, we assume that the platform can only set a fixed fee f for innovators, but cannot prevent innovators from joining the platform as long as they pay the fee (in other words, the platform cannot discriminate across innovators). Then, the platform can either set the fixed fee such that all innovators join the platform, or set it high enough such that no innovators join it. In this latter case, the platform does not make any profit by selling data, and its profit only derives from providing the medical service and is equal to $p - c$. In Section 5, we consider an extension in which

the platform can restrict the number of innovators.

3.3.1 No Consent

Under the no-consent regime, the platform can collect data from patients and share it with innovators without seeking patients' consent. Hence, it sets $t = 0$ and collects the data of all patients ($n_S = 1$). Moreover, because innovators are symmetric, the optimal fixed fee f is the maximum fee they accept, which extracts all their surplus from trade, i.e., $f^{NC} = q(1)R$. Hence, the platform's profit is

$$\Pi^{NC}(n_I) = n_I q(1)R + p - c. \quad (3)$$

Social welfare under the no-consent regime is given by

$$W^{NC}(n_I) = v - c - \lambda D(n_I) + n_I q(1)V.$$

This corresponds to the welfare under full sharing. In this case, the social value of innovation is the highest because $n_S = 1$ and $q'(\cdot) > 0$. However, because data on privacy-concerned patients are shared, the latter incur privacy costs.

3.3.2 Consent

Now, consider the case where, at stage 4, patients must give their consent for their data to be shared with innovators. All privacy-unconcerned patients agree to share their data irrespective of the financial compensation offered by the platform. Privacy-concerned patients agree to share their data if and only if the utility they get from sharing ($v - p + t - D(n_I)$) is greater than or equal to the utility from not sharing ($v - p$). Hence, they share data if and only if the platform compensates them for their privacy cost, i.e., $t \geq D(n_I)$.

Therefore, at stage 2, the platform sets a transfer to patients that is either (i) $t = 0$, which induces only privacy-unconcerned patients to agree to data sharing (i.e., $n_S = (1 - \lambda)$), or (ii) $t = D(n_I)$, which induces all patients to agree to data sharing (i.e., $n_S = 1$). As in the previous scenario, the platform sets a fee $f = q(n_S)R$ extracting all surplus from innovators.

Therefore, the profit of the platform is

$$\Pi^C(n_I) = \begin{cases} n_I q(1 - \lambda)R + p - c, & \text{if } t = 0, \\ n_I q(1)R - D(n_I) + p - c, & \text{if } t = D(n_I). \end{cases}$$

Hence, setting $t = 0$ and sharing data only from privacy-unconcerned patients is (strictly) optimal from the platform's perspective if and only if

$$n_I q(1 - \lambda)R + p - c > n_I q(1)R - D(n_I) + p - c \iff n_I \Delta q < B_0 := \frac{D(n_I)}{R}. \quad (4)$$

The platform chooses not to induce data sharing from privacy-concerned patients when the impact of their data on the total private value of successful innovations $n_I \Delta q R$ (which it captures) is lower than the total payment it must make to all patients to access that data (that is, $D(n_I)$).⁸ Note that the total payment exceeds the aggregate privacy costs, as the platform is unable to distinguish between privacy-concerned and unconcerned patients and, therefore, cannot discriminate between them.

Given the result above, the equilibrium profit of the platform under the consent regime is

$$\Pi^C(n_I) = \begin{cases} n_I q(1 - \lambda)R + p - c & \text{if } n_I \Delta q < \frac{D(n_I)}{R} \\ n_I q(1)R - D(n_I) + p - c & \text{if } n_I \Delta q \geq \frac{D(n_I)}{R}. \end{cases} \quad (5)$$

Furthermore, social welfare under the consent regime is

$$W^C(n_I) = \begin{cases} n_I q(1 - \lambda)V + v - c & \text{if } n_I \Delta q < \frac{D(n_I)}{R} \\ n_I q(1)V - \lambda D(n_I) + v - c & \text{if } n_I \Delta q \geq \frac{D(n_I)}{R}. \end{cases} \quad (6)$$

3.3.3 Regulator's choice

We now solve the first stage of the game to determine the choice of the privacy regime made by the regulator. The regulator chooses between the two regimes by comparing welfare under the consent and no-consent regimes. For the platform, profits are always greater under the no-consent regime because the platform can share data with inno-

⁸We assume that whenever the platform is indifferent between inducing full and partial data sharing, it chooses full sharing.

vators without making any payment to patients. However, from a welfare perspective, the comparison of the two regimes is not clear-cut.

Consider first the case where, under the consent regime, the platform sets $t = 0$ and only privacy-unconcerned patients agree to data sharing (i.e., $n_S = (1 - \lambda)$), which occurs if $n_I \Delta q < D(n_I)/R = B_0$. Comparing welfare levels in this region yields

$$W^C(n_I) - W^{NC}(n_I) = -n_I V \Delta q + \lambda D(n_I),$$

implying that welfare is (weakly) higher under the consent regime if $n_I \Delta q \leq \lambda D(n_I)/V = A_0$. Conversely, the no-consent regime dominates if

$$n_I \Delta q \in (A_0, B_0).$$

Now, consider the case where under the consent regime, the platform sets $t = D(n_I)$ and all patients share data (i.e., $n_S = 1$), which occurs if $n_I \Delta q \geq B_0$. In this case, we find

$$W^C(n_I) - W^{NC}(n_I) = 0,$$

implying that the two regimes are equivalent from a social welfare perspective.

We summarize these findings in the following:

Proposition 1. *From a social perspective:*

- i) the consent regime (weakly) dominates the no-consent regime if $n_I \Delta q \leq A_0$;*
- ii) the no-consent regime dominates the consent regime if $n_I \Delta q \in (A_0, B_0)$;*
- iii) the consent and no-consent regimes are equivalent if $n_I \Delta q \geq B_0$.*

To explain the above result, it is useful to compare social welfare under each regime with the first-best. This comparison, summarized in Figure 1, highlights potential divergences between the trade-off faced by the platform under each regime and the corresponding social trade-off, depending on both the level of privacy costs and the value of data for innovation.




Figure 1: First-best scenario and welfare comparisons of the consent and no-consent regimes.

More precisely, if privacy costs are low relative to the impact of data on innovation (i.e., $n_I \Delta q \geq A_0$), it is socially optimal that all patients share data. If $n_I \Delta q \geq B_0$, at equilibrium the platform induces full data sharing under both the consent and no-consent regimes, implying that welfare is the same (and maximized) under the two scenarios. The reason is that, under the consent regime, the platform is willing to offer sufficient compensation for privacy-concerned patients to agree to share their data. Note, however, that while imposing a consent regime has no impact on social welfare in this scenario, it does have distributional consequences, as it shifts surplus from the platform to patients.

For intermediate levels of privacy costs (i.e., $A_0 < n_I \Delta q < B_0$), the level of data sharing in equilibrium under the consent regime is too low: while it is socially optimal for all patients to share their data (because the increase in the expected social value of innovation due to full data sharing is higher than the aggregate privacy costs), under the consent regime the platform is not willing to offer a payment that convinces privacy-concerned patients to share their data. This divergence between private and social incentives reflects two frictions. First, the platform cannot discriminate between privacy-concerned and privacy-unconcerned patients, so the compensation required to induce full data sharing exceeds the associated privacy costs. Second, the platform does not internalize the full social value of innovation, as it only captures part of it. In this scenario, social welfare is lower under the consent regime than under the no-consent one, where instead all patient data is collected and shared with innovators.

By contrast, if privacy costs are high enough relative to the value of data for innovation (i.e., $n_I \Delta q < A_0$), it is not socially optimal for privacy-concerned patients to share data, because the privacy costs they incur are not compensated from a social point of

view by the expected social value of additional innovation due to additional data from privacy-concerned patients. Welfare is then higher under the consent regime.

4 Data externalities

We now consider a richer model where the choice of a patient to agree to data sharing can generate privacy externalities. This may occur when data shared by some patients allow the innovators to infer information about patients who did not share their data (see, e.g., Choi et al., 2019 and Acemoglu et al., 2022). DNA is an obvious example, since it inherently reveals information about relatives. More broadly, when data is used for machine learning, innovators can exploit patterns from similar patients to infer the health status of individuals who never shared their own data.

4.1 Setup

With data externalities, privacy-concerned patients incur a privacy cost even if they do not agree to share their data, as long as other patients agree to data sharing. Conversely, privacy-unconcerned patients do not incur any cost, whether their data is shared or not. Moreover, we assume that privacy-concerned patients sharing their data do not incur an additional cost related to externalities.

Formally, the privacy cost incurred by privacy-concerned patients when they do not agree to data sharing is denoted as $\tilde{D}(n_I, n_S)$, which we assume to be increasing in the number of innovators n_I and the number of patients sharing their data n_S , with $\tilde{D}(n_I, 0) = 0$. The assumption that $\tilde{D}(n_I, n_S)$ increases with n_S captures the fact that non-sharing patients are negatively affected by a wider sharing of other patients' data that can be used to infer information about them. We further assume that the privacy cost due to data externalities is lower than the privacy cost incurred when the patient agrees to having their data shared, that is, $\tilde{D}(n_I, n_S) < D(n_I)$ for any $n_I > 0$ and $n_S \in [0, 1]$.

Thus, the utility of privacy-unconcerned patients is again $v - p + t$, since they always agree to data sharing and get the transfer t from the platform. The utility of privacy-concerned patients is equal to $v - p - \tilde{D}(n_I, n_S)$ if their own data is not shared,

and to $v - p + t - D(n_I)$ if it is.

We keep assuming that the probability of innovation depends only on the mass n_S of patients who directly share their data and is therefore given by $q(n_S)$. This allows us to isolate the effect of informational inference on patients' privacy costs. Allowing inferred information on non-sharing patients to contribute to innovation would not substantially alter our main results. For example, the probability of innovation could be written as $q_\theta(n_S)$, where $\theta \in [0, 1]$ measures the informational value of inferences about non-sharing patients, and $q_\theta(n_S)$ is continuous and (weakly) increasing in θ for any n_S . Defining $\Delta q_\theta := q(1) - q_\theta(1 - \lambda)$, the analysis would still have the same structure after replacing q with q_θ and Δq with Δq_θ .⁹ In particular, Proposition 2 would continue to hold with $n_I \Delta q_\theta$ replacing $n_I \Delta q$.¹⁰

4.2 First-best

We first analyze the first-best scenario, where the regulator knows patients' privacy costs and chooses the level of data sharing that maximizes social welfare.

Recall that for a given number n_S of patients sharing their data, it is socially optimal to prioritize the sharing of data by privacy-unconcerned patients. Taking this into account, social welfare can be written as follows:

$$W(n_I, n_S) = \begin{cases} v - c + n_I q(n_S) V - \lambda \tilde{D}(n_I, n_S), & \text{if } n_S \in [0, 1 - \lambda] \\ v - c + n_I q(n_S) V - (n_S - 1 + \lambda) D(n_I) - (1 - n_S) \tilde{D}(n_I, n_S) & \text{if } n_S \in (1 - \lambda, 1]. \end{cases}$$

For the sake of simplicity and comparability with the baseline model, we assume that $W(n_I, n_S)$ is weakly convex in n_S over each of the two intervals defined above ($[0, 1 - \lambda]$ and $(1 - \lambda, 1]$).¹¹ This ensures that it cannot be strictly socially optimal

⁹This requires preserving the weak convexity property of the social welfare function.

¹⁰When $\theta = 0$ this specification coincides with our baseline innovation technology. By continuity, whenever the relevant inequalities are strict at $\theta = 0$, the qualitative comparison with the baseline model continues to hold for sufficiently small values of θ . For larger values, however, the comparison may change because inference increases the information available under partial sharing but not under full sharing, thereby reducing the incremental innovation benefit of moving from partial to full sharing.

¹¹For example, it is straightforward to show that this holds if $\tilde{D}(n_I, n_S)$ is weakly concave in n_S (because $q(n_S)$ is weakly convex in n_S).

for only a strict subset of the privacy-unconcerned patients to share their data ($n_S \in (0, 1 - \lambda)$) or only a strict subset of the privacy-concerned patients to share their data ($n_S \in (1 - \lambda, 1)$). In particular, conditional on some data sharing being socially optimal, social welfare is maximized at either $n_S = 1 - \lambda$ or $n_S = 1$ (as in the baseline model).

If the data of all patients is shared (i.e., $n_S = 1$), social welfare is given by

$$W_{all} = v - c + n_I q(1)V - \lambda D(n_I).$$

If only privacy-unconcerned patients have their data shared (i.e., $n_S = 1 - \lambda$), social welfare is given by

$$W_{partial} = v - c + n_I q(1 - \lambda)V - \lambda \tilde{D}(n_I, 1 - \lambda).$$

Unlike in the baseline model, the presence of negative data externalities can make no sharing (strictly) optimal. This occurs if and only if

$$v - c > \max\{W_{all}, W_{partial}\}.$$

This is the case if and only if the privacy costs of data sharing (including those arising from externalities) are sufficiently high relative to the expected social value of the innovation it enables, that is, if and only if

$$\max\{n_I q(1 - \lambda)V - \lambda \tilde{D}(n_I, 1 - \lambda), n_I q(1)V - \lambda D(n_I)\} < 0. \quad (7)$$

If the above condition is not satisfied, then some sharing (full or partial) is socially optimal. More specifically, full sharing is socially preferred to partial sharing (i.e., $W_{all} > W_{partial}$) if and only if

$$n_I \Delta q > A := \frac{\lambda \left(D(n_I) - \tilde{D}(n_I, 1 - \lambda) \right)}{V}. \quad (8)$$

This holds when the incremental privacy costs from sharing for privacy-concerned patients are low relative to the beneficial social impact of such data sharing stemming from the resulting increase in the probability of innovation. The incremental privacy

costs are low when the direct privacy cost incurred by a privacy-concerned patient from sharing their data is close in magnitude to the data externality they receive when privacy-unconcerned patients share their data (which may occur when innovators possess large datasets and strong inference capabilities), or when the share of privacy-concerned patients λ is low enough. Otherwise, the first-best is such that only privacy-unconcerned patients share their data.

The threshold A increases with λ for two reasons. First, as in the baseline model, this occurs because, as λ increases, the share of patients who incur an additional cost when moving from partial to full sharing increases. The second effect is novel: if λ is high, the data externality under partial sharing is small (as the share of privacy-unconcerned patients is low).

Furthermore, it is easy to show that $A < A_0$, where A_0 is defined in condition (2). This implies that the region in which full data sharing socially dominates partial data sharing (i.e., restricted to privacy-unconcerned patients) expands when data externalities are present. The reason is that, in this case, privacy-concerned patients incur a cost even if they do not disclose their data.

4.3 Equilibrium analysis

We now derive the equilibrium of the game in the presence of data externalities.

4.3.1 No Consent

Under the no-consent regime, the platform collects data from all patients (i.e., $n_S = 1$) without any compensation. Privacy-concerned patients incur a privacy cost $D(n_I)$. The profit of the platform in this scenario is $n_I q(1)R + p - c$, and welfare is $W^{NC} = v - c + n_I q(1)V - \lambda D(n_I)$. The results are the same as in the baseline model.

4.3.2 Consent

Consider now the case where patients have to give their consent for the platform to share their data. We maintain in this section the assumption that patients who are indifferent between sharing and not sharing their data agree to sharing in the consent regime. This implies that privacy-unconcerned patients always agree to data sharing.

By contrast, privacy-concerned patients compare the utility of agreeing to data sharing, $v - p + t - D(n_I)$, with the utility of refusing, $v - p - \tilde{D}(n_I, n_S)$. Let $x \in [0, \lambda]$ denote the mass of privacy-concerned patients who agree to data sharing. Since all privacy-unconcerned patients agree to share their data, the total mass of patients whose data is shared is $n_S = 1 - \lambda + x$. Hence, the gain from consenting for a privacy-concerned patient is

$$G(x, t) := t - D(n_I) + \tilde{D}(n_I, 1 - \lambda + x).$$

A privacy-concerned patient therefore consents if and only if $G(x, t) \geq 0$. Since $\tilde{D}(n_I, n_S)$ is increasing in n_S , the incentive to consent is increasing in the mass of other patients who consent: patients' consent decisions are strategic complements. This arises because data externalities affect patients' outside option. As more patients share their data, a patient who refuses to share suffers a larger informational privacy externality, whereas their privacy cost from sharing remains equal to $D(n_I)$. Refusal therefore becomes less attractive, making consent more attractive. This strategic complementarity, which is absent from the baseline model, can generate multiple self-fulfilling consent outcomes.

We now characterize the resulting equilibria. Consider the two candidate symmetric outcomes.¹² If no privacy-concerned patient agrees to data sharing, then $x = 0$ and

$$G(0, t) = t - D(n_I) + \tilde{D}(n_I, 1 - \lambda) = t - \bar{t},$$

where $\bar{t} := D(n_I) - \tilde{D}(n_I, 1 - \lambda)$. Thus, an outcome in which no privacy-concerned patient agrees is an equilibrium if $t < \bar{t}$. If instead all privacy-concerned patients agree, then $x = \lambda$ and

$$G(\lambda, t) = t - D(n_I) + \tilde{D}(n_I, 1) = t - \underline{t},$$

where $\underline{t} := D(n_I) - \tilde{D}(n_I, 1)$. Thus, an outcome in which all privacy-concerned patients agree to share their data is an equilibrium if $t \geq \underline{t}$. Notice that $\underline{t} \leq \bar{t}$, because

¹²The assumption that an indifferent patient consents rules out asymmetric equilibria in which a strict subset of privacy-concerned patients consent while the others do not. Such equilibria would require privacy-concerned patients to be indifferent between consenting and not consenting, with some of them nevertheless refusing to share their data.

$\tilde{D}(n_I, 1) \geq \tilde{D}(n_I, 1 - \lambda)$, with the inequality being strict whenever $\lambda > 0$.

We can therefore distinguish three cases. First, if $t \in [0, \underline{t})$, then $G(x, t) < 0$ for all $x \in [0, \lambda]$. Consequently, no privacy-concerned patient agrees to data sharing, and this is the unique equilibrium.

Second, if $t \in [\underline{t}, \bar{t})$, there are two equilibria: one in which none of the privacy-concerned patients consent and one in which they all do. Indeed, if a privacy-concerned patient expects no other privacy-concerned patient to agree to data sharing, refusing is optimal because $G(0, t) < 0$, and these expectations are self-fulfilling. If a privacy-concerned patient expects all privacy-concerned patients to consent, agreeing is optimal because $G(\lambda, t) \geq 0$, and these expectations are also self-fulfilling.¹³

Third, if $t \geq \bar{t}$, then $G(x, t) \geq 0$ for every $x \in [0, \lambda]$. Hence, all privacy-concerned patients consent, and this is the unique equilibrium.

For $t \in [\underline{t}, \bar{t})$, there are two equilibria: one in which all privacy-concerned patients consent and one in which none of them does. We assume that the former is selected to capture patients' inability to coordinate on their collectively preferred outcome in the presence of data externalities. Indeed, privacy-concerned patients collectively prefer that none of them consents because $t < \bar{t}$. Selecting the equilibrium where no privacy-concerned patients consent would rule out the coordination-failure mechanism that we seek to analyze.¹⁴

This coordination failure is consistent with consent being requested privately and individually. Patients generally cannot observe or condition their decisions on the choices of others and lack a mechanism for committing collectively to refusal. By contrast, the platform controls the consent procedure and can make broad participation appear likely.

¹³As previously mentioned, without the assumption that an indifferent patient consents, partial-sharing equilibria could also arise. An interior equilibrium would be characterized by a mass $x^* \in (0, \lambda)$ satisfying $t = D(n_I) - \tilde{D}(n_I, 1 - \lambda + x^*)$.

¹⁴If we select this equilibrium, the equilibrium analysis would have the same structure as in the baseline model without data externalities. The key difference would be that, to induce full sharing, the platform must offer a transfer $\bar{t} = D(n_I) - \tilde{D}(n_I, 1 - \lambda)$.

Therefore, the platform's profit is given by

$$\Pi^C(n_I, t) = \begin{cases} n_I q(1 - \lambda)R + p - c - (1 - \lambda)t, & \text{if } t \in [0, \underline{t}), \\ n_I q(1)R + p - c - t, & \text{if } t \in [\underline{t}, +\infty). \end{cases}$$

Given the shape of the platform's profit function, it is maximized at either $t = 0$ or $t = \underline{t}$. Therefore, as in our baseline model, in equilibrium either all patients agree to data sharing, or only the privacy-unconcerned agree to it.

The platform chooses $t = \underline{t}$, thereby inducing all patients to agree to their data being shared if and only if

$$n_I q(1 - \lambda)R + p - c \leq n_I q(1)R - D(n_I) + \tilde{D}(n_I, 1) + p - c,$$

which is equivalent to

$$n_I \Delta q \geq B := \frac{D(n_I) - \tilde{D}(n_I, 1)}{R}.$$

Full data sharing is more likely to occur under the consent regime when the incremental privacy cost of agreeing to data sharing for privacy-concerned patients (i.e., $D(n_I) - \tilde{D}(n_I, 1)$) is low. This incremental cost is small if the data externality $\tilde{D}(n_I, 1)$ suffered by a non-sharing privacy-concerned patient when all other patients share their data is large.

Straightforward comparisons with condition (4) yield $B < B_0$, implying that sharing is more likely to occur in equilibrium when there are externalities than in the benchmark model. This holds because the compensation needed to induce patients to share data is lower in the presence of (negative) data externalities.

Thus, social welfare under the consent regime is

$$W^C(n_I) = \begin{cases} v - c + n_I q(1 - \lambda)V - \lambda \tilde{D}(n_I, (1 - \lambda)), & \text{if } n_I \Delta q < \frac{D(n_I) - \tilde{D}(n_I, 1)}{R}, \\ v - c + n_I q(1)V - \lambda D(n_I), & \text{if } n_I \Delta q \geq \frac{D(n_I) - \tilde{D}(n_I, 1)}{R}. \end{cases}$$

4.3.3 Regulator's choice

We can now establish a welfare comparison between the consent and no-consent regimes to determine the regime chosen by the regulator. When $n_I \Delta q \geq B$, both regimes lead

to full sharing and they are equivalent from a welfare perspective. In this case, the regulator is indifferent between the two regimes.

When $n_I \Delta q < B$, under the no-consent regime the data of all patients is shared, while under the consent regime only the data of the $(1 - \lambda)$ privacy-unconcerned patients is shared. Full sharing socially dominates partial sharing if and only if $n_I \Delta q \geq A$, with A defined in condition (8).

Hence, when $n_I \Delta q < B$, we should consider two cases, depending on the comparison between thresholds A and B . In the baseline scenario, we found unambiguously that $A_0 < B_0$. By contrast, the presence of externalities makes it possible for B to be smaller than A . In this case, full sharing may arise in equilibrium under the consent regime even though partial sharing is socially optimal. Note that such oversharing never occurs in our baseline model with no data externalities.

In the case where $A < B$, the comparison is similar to that in Section 3. Welfare is higher under the consent regime than under the no-consent regime when $n_I \Delta q < A < B$, because the platform then chooses not to pay patients so that only the $(1 - \lambda)$ privacy-unconcerned patients share data, which is socially optimal. If $A \leq n_I \Delta q < B$, under the consent regime the platform again only gets data from $(1 - \lambda)$ patients, while full sharing is socially optimal. In this case, the absence of a consent requirement enhances welfare. Finally, if $A < B \leq n_I \Delta q$, the platform induces all patients to share, which is socially optimal. Since the no-consent regime leads to the same outcome, both regimes are equivalent from a welfare perspective.

As mentioned above, the case where $A > B$ is new. In this scenario, the comparison of the consent and no-consent regimes yields only two possible outcomes. If $n_I \Delta q \geq B$, both regimes lead to full sharing and they are equivalent from a welfare perspective. However, if $n_I \Delta q < B$, the consent regime dominates the no-consent regime.

We now provide conditions under which the threshold A is below or above the threshold B :

Lemma 1. *If $\tilde{D}(n_I, 1)/D(n_I) < (V - R)/V$, then $A < B$ for all λ . If $\tilde{D}(n_I, 1)/D(n_I) \geq (V - R)/V$, there exists a threshold $\tilde{\lambda} \in (0, 1]$ such that $A < B \iff \lambda < \tilde{\lambda}$.*

Proof. After rearrangements, we can show that $A < B$ if and only if

$$\frac{\lambda \left(D(n_I) - \tilde{D}(n_I, 1 - \lambda) \right)}{D(n_I) - \tilde{D}(n_I, 1)} < \frac{V}{R}. \quad (9)$$

The left-hand side of this inequality is increasing in λ and equal to 0 if $\lambda = 0$. Therefore, the inequality is satisfied for all λ if it is satisfied for $\lambda = 1$. This is the case if and only if

$$\frac{D(n_I)}{D(n_I) - \tilde{D}(n_I, 1)} < \frac{V}{R} \iff \frac{\tilde{D}(n_I, 1)}{D(n_I)} < \frac{V - R}{V}.$$

Conversely, if $\tilde{D}(n_I, 1)/D(n_I) \geq 1 - R/V$, there exists a threshold $\tilde{\lambda} \in (0, 1]$ such that condition (9) is satisfied for $\lambda < \tilde{\lambda}$. $\square$

Lemma 1 states that the scenario where $A < B$ occurs when the magnitude of privacy cost from data externalities over direct privacy costs ($\tilde{D}(n_I, 1)/D(n_I)$) is lower compared to the incremental value of innovation over the total social value ($(V - R)/V$), or this relative measure is higher but the fraction of privacy-concerned patients λ is small. However, if they are both high, then the scenario where $A > B$ occurs.

We can now compare the first-best scenario with the equilibrium one. First, if no data sharing at all is socially optimal (i.e., if condition (7) holds), then there is always oversharing in equilibrium. The reason is that the platform always collects and shares data of privacy-unconcerned patients. If some data sharing is socially optimal (i.e., if condition (7) does not hold), the analysis of the case where $A < B$ is qualitatively equivalent to the benchmark case with no data externalities. Note that data externalities make full data sharing more appealing to the platform because they reduce the outside option of privacy-concerned patients, thereby lowering the compensation the platform must pay them to share their data. They also make full data sharing relatively more appealing from a social perspective, because non-sharing patients also incur a cost.

However, as noted above, data externalities open up the possibility that the equilibrium involves full data sharing under the consent regime, while partial data sharing is the first-best outcome, which requires that $A > B$. Such oversharing is due to the fact that the platform now needs to compensate patients only $D(n_I) - \tilde{D}(n_I, 1)$, and

not the incremental privacy cost per privacy-concerned patient, $D(n_I) - \tilde{D}(n_I, 1 - \lambda)$. In words, a single patient who does not agree to data sharing while all others agree still suffers from an externality, which is relatively high when $A > B$ (see Lemma 1). As a consequence, individual patients are willing to share their data for an amount lower than their direct privacy costs under the consent regime. The social cost of sharing is higher than the compensation required to make each individual patient share their data, because the social planner considers a situation where *all* the privacy-concerned patients simultaneously refuse to share. The discrepancy between the “private” and the “social” data externality increases with the number of privacy-concerned patients. When the latter is large enough, there is oversharing even under the consent regime.




Figure 2: Welfare comparisons of the consent and no-consent regimes. Case $A < B$




Figure 3: Welfare comparisons of the consent and no-consent regimes. Case $A > B$

Our results are summarized in the following proposition, which is visually represented in Figures 2 and 3.

Proposition 2. *In the presence of data externalities, from a social perspective:*

- i) the consent regime dominates the no-consent regime if $n_I \Delta q < \min\{A, B\}$. If*

condition (7) holds, consent leads to suboptimal oversharing. Otherwise, it decentralizes the first-best;

- ii) the no-consent regime dominates the consent regime if $A < B$ and $A \leq n_I \Delta q < B$. If condition (7) holds, the no-consent regime leads to suboptimal oversharing. Otherwise, it decentralizes the first-best;
- iii) the consent and no-consent regimes are equivalent if $n_I \Delta q \geq \max\{A, B\}$. They lead to oversharing if condition (7) holds, and they are first-best optimal otherwise;
- iv) the consent and no-consent regimes are equivalent and lead to suboptimal oversharing when $A > B$ and $B \leq n_I \Delta q < A$.

To sum up, when condition (7) is satisfied, from a welfare perspective neither the consent regime nor the no-consent regime decentralizes the first-best allocation, which entails no data sharing.

The no-consent regime leads to the optimal outcome when (7) is violated and $n_I \Delta q \geq A$, because in this case it is optimal that all data is shared. If $n_I \Delta q < A$, it is optimal that only privacy-unconcerned patients have their data shared, implying that the no-consent regime does not lead to the optimal outcome.

4.3.4 More restrictive regulations

Our results suggest that, when there are many privacy-concerned patients and data externalities are strong enough compared to the relative social value of innovation, standard consent policies may not be sufficient to correct for oversharing. This calls for stricter regulations limiting data sharing.

Consent with no payment. The possibility of oversharing in the scenario with data externalities that occurs when $A > B$ may justify a regulation that discourages data sharing. The regulator could require the platform to obtain patients' consent while prohibiting financial compensation for those who agree to data sharing. This could be viewed as a special (extreme) case of a broader regulation that imposes a cap on the compensation the platform can offer to patients for their consent.¹⁵

¹⁵In our model, any binding cap on the compensation offered by the platform to patients who agree

Under a regulation that mandates consent but forbids any payment to patients who agree to data sharing, privacy-unconcerned patients will share, while privacy-concerned patients will not. Focusing on the scenario $A > B$, it follows from the analysis above that this regulation—referred to as *consent with no payment*—is equivalent from a welfare perspective to the consent-with-payment regime if $n_I \Delta q < B$, dominates both the consent-with-payment regime and the no-consent regime if $B \leq n_I \Delta q < A$, and is dominated if $n_I \Delta q > A$. Intuitively, the consent-with-no-payment regime is a strictly optimal policy if and only if both the consent-with-payment regime and the no-consent regime lead to oversharing, which occurs for intermediate values of $n_I \Delta q$.

Prohibition of data sharing. The consent-with-no-payment regime can increase welfare relative to both the consent-with-payment and no-consent regimes by preventing privacy-concerned patients from agreeing to data sharing for intermediate benefits from innovation. However, when the expected benefits from overall innovation are small (in absolute terms) relative to the privacy costs associated with own-data sharing or data externalities, it may be socially optimal to prohibit data sharing altogether for both privacy-concerned and privacy-unconcerned patients.

Specifically, a full prohibition of data sharing achieves the first-best outcome and thus (weakly) outperforms all other regimes if condition (7) holds, that is, if data sharing leads to privacy costs that are high relative to the social value generated by all resulting innovation. Similar prohibitions exist in other contexts. For instance, the EU Digital Services Act prohibits the use of sensitive personal data (such as health data) for online behavioral advertising, irrespective of whether the user has consented or received compensation.

to share their data has the same effect as forbidding payment entirely (i.e., setting the cap to zero), because the model assumes only two types of patients.

5 Extension: The platform controls the number of innovators

5.1 No data externalities

In the baseline model, the number of innovators is fixed and the platform cannot prevent an innovator from accessing patient data as long as the innovator pays the membership fee. We now return to the setting without data externalities described in Section 3 and allow the platform to restrict the number of innovators receiving access. Let $\bar{n}_I$ denote the total number of potential innovators and let $n_I \in [0, \bar{n}_I]$ denote the number admitted by the platform. We assume that the platform can commit to this choice.¹⁶

Under the no-consent regime, the platform does not compensate patients, and its profit is

$$\Pi^{NC}(n_I) = n_I Rq(1) + p - c,$$

which is increasing in n_I . It therefore admits all potential innovators:

$$n_I^{NC} = \bar{n}_I.$$

Under the consent regime, for any given n_I , the platform chooses between two pricing strategies, leading to either full sharing (F) or partial sharing (P). It can set $t = 0$, in which case only privacy-unconcerned patients share their data, yielding profit

$$\Pi_P(n_I) = n_I Rq(1 - \lambda) + p - c.$$

The profit associated with partial sharing is increasing in n_I and is therefore maximized at $\bar{n}_I$. Alternatively, it can set $t = D(n_I)$, in which case all patients share their data, yielding profit

$$\Pi_F(n_I) = n_I Rq(1) - D(n_I) + p - c.$$

¹⁶For tractability, we allow n_I to be continuous.

In the case with full sharing an interior solution may occur. Let

$$n_F \in \operatorname{argmax}_{n_I \in [0, \bar{n}_I]} \{n_I Rq(1) - D(n_I)\}.$$

We assume that $D'(0) < Rq(1)$, which ensures that $n_F > 0$.

The platform's choice under consent can then be characterized by comparing partial sharing with all innovators and full sharing with a potentially restricted number n_F of innovators.

Proposition 3. *Suppose that the platform controls the number of innovators accessing patient data. Under the no-consent regime, it admits all innovators, so that $n_I^{NC} = \bar{n}_I$. Under the consent regime,*

i) the platform sets $t = D(n_F)$, admits n_F innovators, and induces all patients to share their data if

$$n_F Rq(1) - D(n_F) > \bar{n}_I Rq(1 - \lambda); \quad (10)$$

ii) if the reverse strict inequality holds, the platform sets $t = 0$, admits all $\bar{n}_I$ innovators, and obtains data only from privacy-unconcerned patients.

Proposition 3 implies that the shape of the privacy-cost function plays a key role in the characterization of the platform's optimal strategy. If D is linear or concave, $n_I Rq(1) - D(n_I)$ is linear or convex, so its maximum is attained at an endpoint. Choosing $n_I = 0$ yields no data-related profit and is strictly dominated by partial sharing with $\bar{n}_I$ innovators. Hence, the platform never restricts access when D is linear or concave. By contrast, if D is strictly convex, the full-sharing profit is strictly concave and may have an interior maximizer. In particular, if $Rq(1) < D'(\bar{n}_I)$ then $n_F \in (0, \bar{n}_I)$ is uniquely characterized by

$$D'(n_F) = Rq(1). \quad (11)$$

The platform restricts access if, in addition, condition (10) holds. Thus, access restriction requires privacy costs to increase sufficiently rapidly with the number of innovators and the full-sharing strategy to be more profitable than partial sharing with all innovators.

We now compare welfare across the two privacy regimes. Under no consent, all patients share their data and all innovators are admitted, so welfare is

$$W^{NC}(\bar{n}_I) = v - c + \bar{n}_I V q(1) - \lambda D(\bar{n}_I).$$

If the platform chooses partial sharing under consent, Proposition 3 implies that it admits all innovators. In this case, welfare is

$$W^C(\bar{n}_I) = v - c + \bar{n}_I V q(1 - \lambda).$$

Hence, the welfare difference can be rewritten as

$$W^C - W^{NC} = \lambda D(\bar{n}_I) - \bar{n}_I V \Delta q.$$

Consequently, the welfare comparison is the same as in the baseline model, and consent is socially preferable if and only if

$$\frac{D(\bar{n}_I)}{\bar{n}_I} > \frac{V \Delta q}{\lambda}.$$

If instead the platform chooses full sharing with n_F innovators under consent, welfare is

$$W^C(n_F) = v - c + n_F V q(1) - \lambda D(n_F).$$

If $n_F = \bar{n}_I$, the two regimes generate the same allocation and the same welfare. If $n_F < \bar{n}_I$, their welfare difference is

$$W^C(n_F) - W^{NC}(\bar{n}_I) = \lambda [D(\bar{n}_I) - D(n_F)] - (\bar{n}_I - n_F) V q(1).$$

It follows that consent is socially preferable if and only if

$$\lambda(D(\bar{n}_I) - D(n_F)) > V q(1)(\bar{n}_I - n_F). \quad (12)$$

The left-hand side is the privacy cost saved due to the excluded innovators, whereas the right-hand side is the social innovation value lost due to the excluded innovators.

Interestingly, when n_F is interior, condition (11) implies that the marginal social

welfare from admitting an additional innovator while maintaining full sharing is

$$Vq(1) - \lambda D'(n_F) = q(1)(V - \lambda R) > 0.$$

Conditional on full sharing, the platform therefore restricts innovator access more than a social planner would choose if $W^C(n_F)$ is strictly concave, which is the case if D is strictly convex. Nevertheless, consent may still generate higher welfare than no consent because the latter always admits all innovators and may generate excessive aggregate privacy costs, as characterized by condition (12).

To summarize, allowing the platform to restrict innovator access introduces an additional channel through which privacy regulation affects welfare. When consent induces partial sharing, the platform admits all innovators and the welfare comparison is unchanged relative to the baseline model. When consent induces full sharing, the platform may restrict the number of innovators to reduce the compensation required to induce patients to share their data. This generates a new trade-off between lower privacy costs and lower innovation.

5.2 Data externalities

We now extend the analysis to the case with data externalities. We maintain the equilibrium-selection rule introduced in Section 4. For a given number of innovators n_I , the lowest transfer that induces all privacy-concerned patients to consent is

$$\underline{t}(n_I) = D(n_I) - \tilde{D}(n_I, 1).$$

It is useful to define the platform's net compensation cost as

$$C(n_I) := D(n_I) - \tilde{D}(n_I, 1).$$

If the platform sets $t = 0$, only privacy-unconcerned patients share their data and its profit is

$$\Pi_P(n_I) = n_I R q(1 - \lambda) + p - c.$$

As before, this profit is increasing in n_I and is maximized at $n_I = \bar{n}_I$. If the platform

instead induces full sharing, its profit is

$$\Pi_F^E(n_I) = n_I Rq(1) - C(n_I) + p - c.$$

The platform chooses the number of innovators accessing the data maximizing the above equation. Let

$$n_F^E \in \operatorname{argmax}_{n_I \in [0, \bar{n}_I]} \{n_I Rq(1) - C(n_I)\}.$$

We assume that $C'(0) < Rq(1)$, which ensures that $n_F^E > 0$.

The platform induces full sharing and admits n_F^E innovators if

$$n_F^E Rq(1) - C(n_F^E) > \bar{n}_I Rq(1 - \lambda).$$

If the reverse strict inequality holds, it sets $t = 0$, admits all $\bar{n}_I$ innovators, and obtains data only from privacy-unconcerned patients. If the two expressions are equal, both strategies maximize profit. Thus, the characterization in Proposition 3 continues to apply after replacing the direct compensation cost $D(n_I)$ with the net compensation cost $C(n_I)$.

Data externalities make full sharing more attractive to the platform because they reduce the compensation required from $D(n_I)$ to $C(n_I)$. They may also affect the number of innovators admitted under full sharing. In particular, if C is strictly convex and the optimal number of innovators is interior, it satisfies

$$C'(n_F^E) = D'(n_F^E) - \frac{\partial \tilde{D}(n_F^E, 1)}{\partial n_I} = Rq(1). \quad (13)$$

By reducing the marginal compensation cost, the data externality gives the platform stronger incentives to grant access to additional innovators.¹⁷

We now turn to the welfare comparison. If consent produces partial sharing with all innovators, then

$$W^C - W^{NC} = \lambda \left[D(\bar{n}_I) - \tilde{D}(\bar{n}_I, 1 - \lambda) \right] - \bar{n}_I V \Delta q.$$

¹⁷Recall that the marginal externality $\partial \tilde{D}(n_I, 1)/\partial n_I$ is positive.

If consent instead produces full sharing with $n_F^E < \bar{n}_I$, then

$$W^C - W^{NC} = \lambda [D(\bar{n}_I) - D(n_F^E)] - (\bar{n}_I - n_F^E)Vq(1).$$

Hence, conditional on full sharing, consent is socially preferable if the privacy costs saved by restricting innovator access exceed the associated loss in innovation.

Data externalities also alter the comparison between the platform's choice and the socially optimal number of innovators. At an interior n_F^E , equation (13) implies that the marginal social welfare from admitting another innovator while maintaining full sharing is

$$Vq(1) - \lambda D'(n_F^E) = q(1)(V - \lambda R) - \lambda \frac{\partial \tilde{D}(n_F^E, 1)}{\partial n_I}.$$

Unlike the case without data externalities, this expression need not be positive. Conditional on full sharing, when the marginal data externality is relatively small, the platform restricts access more than a social planner would prefer if the welfare function is strictly concave.¹⁸ When it is sufficiently large, the platform may instead admit too many innovators. Data externalities can therefore generate excessive access by innovators in addition to excessive data sharing.

6 Conclusion

This paper sheds light on the desirability of a consent-based privacy regulation in the context of data-driven innovation in the healthcare sector. Our analysis examines the trade-off faced by a profit-maximizing two-sided platform that collects patient data as a byproduct of providing medical services and can share it with innovators. We compare this trade-off with the corresponding social trade-off.

In our baseline model, we assume away data externalities. In this setting, the platform's incentives are aligned with those of a social planner if privacy costs are sufficiently low relative to the value of data for innovation. In such a scenario, the consent and no-consent regimes are equivalent from a social welfare perspective. However,

¹⁸A sufficient condition for this is that D is strictly convex.

when privacy costs increase, the two regimes are no longer equivalent, as the platform’s incentives diverge from those of a social planner. More specifically, for intermediate levels of privacy costs, the no-consent regime dominates the consent regime, whereas the reverse holds when privacy costs are sufficiently high.

When data externalities that harm patients even when they do not share their own data are present, additional effects come into play. Since patients cannot coordinate their decisions to accept or reject data sharing, a single patient who refuses to share their data can expect a high data externality. This reduces the compensation required to induce privacy-concerned patients to agree to data sharing under the consent regime. In this case, oversharing may occur even under the consent regime, which may call for more restrictive regulations, such as requiring the platform to obtain patients’ consent while prohibiting financial compensation for those who agree to share their data, or even a full prohibition of data sharing.

Then, we allow the platform to control access by innovators. In this case, it admits all potential innovators under no consent, but may restrict access under consent. Without data externalities, an interior platform choice involves fewer innovators than a social planner would prefer, conditional on full sharing. Data externalities reduce the compensation required to induce sharing and may therefore encourage the platform to admit more innovators. If the marginal externality is sufficiently strong, this may result in excessive access by innovators.

Our analysis could be further extended in various directions. First, as already mentioned, our framework abstracts from same-side network externalities on the innovators’ side. A natural extension would be to incorporate the latter by allowing for competition or complementarities among innovators. Second, we assume that the platform cannot discriminate across patients, and must offer the same payment in exchange for consent. Extending the analysis to allow for more complex contracts including differentiated payments for different levels of data sharing is left for further investigation. Finally, we assume that the platform has a monopoly on the data-sharing side. Although the market is characterized by large platforms, such as the Mayo Clinic platform in the U.S., allowing for competition among platforms for innovators would be an interesting avenue for future research.

References

- Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asu Ozdaglar (2022) “Too Much Data: Prices and Inefficiencies in Data Markets,” *American Economic Journal: Microeconomics*, 14 (4), 218–56, 10.1257/mic.20200200.
- Acquisti, Alessandro, Curtis Taylor, and Liad Wagman (2016) “The Economics of Privacy,” *Journal of Economic Literature*, 54 (2), 442–92, 10.1257/jel.54.2.442.
- Bergemann, Dirk, Alessandro Bonatti, and Tan Gan (2022) “The economics of social data,” *The RAND Journal of Economics*, 53 (2), 263–296.
- Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim (2019) “Privacy and personal data collection with information externalities,” *Journal of Public Economics*, 173, 113–124, <https://doi.org/10.1016/j.jpubeco.2019.02.001>.
- Conti, Chiara, Anna D’Annunzio, and Pierfrancesco Reverberi (2026) “Access to personal data for healthcare innovation,” Mimeo.
- Conti, Chiara, Anna D’Annunzio, and Pierfrancesco Reverberi (2024) “Data-Driven Health Innovation and Privacy Regulation,” *The B.E. Journal of Economic Analysis & Policy*, 24 (1), 329–338, doi:10.1515/bejeap-2023-0020.
- Conti, Chiara and Pierfrancesco Reverberi (2021) “Price discrimination and product quality under opt-in privacy regulation,” *Information Economics and Policy*, 55, 100912, <https://doi.org/10.1016/j.infoecopol.2020.100912>.
- EU (2025) “Regulation (EU) 2025/327 of the European Parliament and of the Council on the European Health Data Space,” Available at https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L_202500327.
- Goldfarb, Avi and Catherine Tucker (2012) “Privacy and Innovation,” *Innovation Policy and the Economy*, 12 (1), 65–90, 10.1086/663156.
- Ichihashi, Shota (2021) “The economics of data externalities,” *Journal of Economic Theory*, 196, 105316.

- Lefouili, Yassine, Leonardo Madio, and Ying Lei Toh (2024) “Privacy Regulation and Quality-Enhancing Innovation,” *The Journal of Industrial Economics*, 72 (2), 662–684, <https://doi.org/10.1111/joie.12374>.
- Miller, Amalia R. and Catherine Tucker (2009) “Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records,” *Management Science*, 55 (7), 1077–1093, 10.1287/mnsc.1090.1014.
- (2018) “Privacy Protection, Personalized Medicine, and Genetic Testing,” *Management Science*, 64 (10), 4648–4668, 10.1287/mnsc.2017.2858.
- Miller, Amalia R. and Catherine E. Tucker (2012) “Electronic Discovery and the Adoption of Information Technology,” *The Journal of Law, Economics, and Organization*, 30 (2), 217–243, 10.1093/jleo/ews038.
- Ross, Casey (2020) “At Mayo Clinic, sharing patient data with companies fuels AI innovation — and concerns about consent,” Available at <https://www.statnews.com/2020/06/03/mayo-clinic-patient-data-fuels-artificial-intelligence-consent-concerns/>.