Algorithms and Forward Externalities

Andrew Kosenko*

May 2024

Abstract

We study a model of insurance provision by a monopolist with incomplete information à la Stiglitz (1977), augmenting it with additional time periods, and assuming that in addition to the different (unobserved by the monopolist) types, consumers have a persistent, observed by the monopolist, characteristic that is related to their type; the characteristic may be used by the monopolist to change subsequent contracts. Information provided by earlier consumers is used to extract the consumer surplus from later consumers - a *forward externality*; the possibility of a pooling component in equilibrium modulates this effect. We also survey the burgeoning recent literature, and discuss the importance of anonymizing consumer data (by merging data records, or by sustaining pooling equilibria) for consumer privacy and welfare.

Keywords: pooling, pooling externalities, learning externalities, forward externalities, monopoly, screening, data, price discrimination.

JEL Codes: D42, D81, D82, D83, D86.

1 Introduction

The proliferation of human-computer interactions and the consequent creation of vast amounts of data is a pervasive feature of today's economic environments. The most minute of the decisions we make - which cup of coffee to order, which website to visit, which item to search for, or which route to take - are recorded, stored, and sold for further use. Individual decisions no longer have only immediate consequences that flow from that particular decision, and are limited in their impact. Our choices are recorded, stored, and used by other agents in interaction with us, creating a novel situation - much of our choice behavior is now set against a backdrop of our previous behavior. This (more or less passive) creation of enormous amounts of data, and the value of *not* creating that data - that is, privacy - has led to a myriad new issues.

*Assistant Professor of Economics, School of Management, Marist College. andrewkosenko.com, `kosenko.andrew@gmail.com`.
I am deeply humbled to contribute to this volume. First and foremost, I would like to thank Joseph Stiglitz - his work and thought has had a profound impact on the way I think about economics, and the way I think about the world, and I could not be more grateful for the chance to contribute my work to this volume. I would also like to thank Domenico Delli Gatti, Eleni Papageorgiou, and Martín Guzman for their organization efforts, Nate Neligh for helpful discussions, and an anonymous referee for very helpful suggestions. Devon De Sanctis provided excellent research assistance. Any remaining errors are my own.

We use the Stiglitz (1977) model of monopoly price discrimination with incomplete information to study data externalities that are created when individuals participate in transactions that record and store observable characteristics about them, across time. Individuals who, through their purchases, create data that is linked to their preferences, exert novel externalities on individuals who make choices at a later time, because the menus faces by these individuals will be modified, using the information gathered earlier, and thus, will face lower consumer surplus; these *forward externalities* arise endogenously. As the literature on the topics of data pricing, privacy, and externalities has grown quickly, with several major conclusions, the second half of this paper (section 3) is a review of this recent body of work.

For an example of such a data externality, consider the problem facing a ride-sharing app. The price for a ride is set by balancing supply of drivers with demand for rides, locations, and perhaps, demand estimation motives. Suppose that the algorithm receives a request for a ride to a prominent hospital from one (say, the *first*) rider at an unusual time (for example, late at night), and that the posted price is accepted. The individual controlling the algorithm may reasonably conclude that riders who request such a ride (at a specific time and location) may have a high willingness to pay; as such, the price may be increased, perhaps significantly so. The next time another, *second*, individual requests a ride to the same location at the same time, the quoted price will reflect the value gleaned from a transaction with the first individual. Thus, individuals who interact with the algorithm at one time exert an externality on individuals who interact with the same algorithm at a later time, *even if those later individuals have never interacted with that algorithm before*. This is a novel informational externality; because of the way it operates - using information provided by earlier consumers to extract consumer surplus from later consumers - we use the term "forward externality."

Financial innovation offers a more sophisticated setting: consider the example of fintech firms (such as Klarna, Affirm, or Afterpay) providing instant credit, or "buy now, pay later" (BNPL) services. These firms offer instant credit (often in the form of additional options at the checkout stage), offering both interest-based and interest-free loans to customers. These firms operate by interacting with the consumer at the checkout stage, obtain consumer consent to verify the consumer's credit history (and perhaps additional elements, such as checking and savings account information), and, contingent on approval, extend loans to consumers.[1] Because BNPL firms are, in effect, insurers, this situation serves as an example of the model of insurance with observable characteristics discussed in section 2.

To study the fintech setting, we add several features to the Stiglitz (1977) model: first, there are two periods, with transactions taking place in both; the exposition, following the original, is mostly graphical.

[1] In essence, these fintech firms "buy" the product from the seller at the time of purchase, and recoup the cost when the consumer repays the loan. The lender takes on all of the default risk, with a number of consequences (one of which is that to be able to make an interest-free loan, the price of the good or service is higher, because of the presence of the online lender).

We study repetitions of this setting in section 2.4. Second, consumers may be "naive" or "sophisticated." The naifs purchase contracts without internalizing the fact that their purchases may be matched with their observable characteristics to identify their type. The sophisticates do internalize this fact. Third, the insurer observes "data" - a signal that is strongly correlated with the type of the consumer before offering contracts. With such data externalities, the asymmetry in information disappears over time, thus enabling essentially full surplus extraction and highly personalized pricing.

Both the share of the economy, and the amount of time and money individuals spend shopping online and interacting with algorithms of various forms (including social media) is rising. In 2022 for example, Americans spent over a record $1 trillon online (Nasdaq, 2023) with the share of retail done online hovering between 15% and 20% (International Monetary Fund, 2022). All of this commerce generates data. In 2014, Axciom, one of the largest brokers at the time, had data on 700 million active consumers worldwide, with thousands of data points per person (Federal Trade Commission, 2014). This number has surely grown in the intervening decade. In 2022, the Federal Trade Commission alleged in a lawsuit that Kochava, another data broker, was offering for sale data on locations of electronic devices, with "125 million monthly active users" (Federal Trade Commission (2022)). Such data often includes not just prices and quantities, but additional metrics (time of transaction, characteristics of the device used, etc.: the "metadata"); browsing history, and cross-device tracking that matches patterns of activity on different sites or devices to the same user provides yet another layer of data on users. The fact that this data can be combined, sold, and used to identify individual consumer patterns (and often in fact, even individual consumers) is at the heart of the problem. The presumption of relative anonymity of markets for ordinary services and goods has been broken. While efforts to control, or limit, the use of cookie data, and thus restore some degree of anonymity, are underway, this process has not yet produced a widely accepted alternative to using cookies for interest-based advertising (IBA).

To be sure, this issue is not entirely new.[2] What is new, however, is the scale of data creation, its personalization, and its granularity. What we suggest here is that with data externalities, this kind of discrimination may be both more pervasive, and more effective, operating on both the extensive and the intensive margins.

Among the insights that have emerged from this literature (and, to preview the results, that will appear in the present model) is the *value of pooling*. In typical models of information economics pooling occurs when more than one type chooses the same contract or action (thus, "pooling" with other types, and making it impossible to distinguish between the types based purely on their actions).

Importantly, in recent theoretical work (surveyed below, as well as the model discussed here), and in

[2] *Inter alia*, Stiglitz, 1975b ; Stiglitz, 1975c, and Leitzinger and Stiglitz, 1984 have raised this issue in several contexts. Greenwald and Stiglitz (1986) address the issue of pecuniary externalities with imperfect and asymmetric information. Phelps, 1972, Arrow, 1973, Rothschild and Stiglitz, 1982, Stiglitz 1973, 1974b also discuss these issues in the context of discrimination.

empirical settings (*inter alia*, Google's "Topics" approach) the ability to pool consumers into groups has emerged as a key tool. Critically, it's the (signals, information, records, or messages, depending on the setting) of the *low* type that are important to have on hand, to be able to pool the high(er) types with, thus "hiding" the higher types. One interpretation of this is that pooling is a way of restoring anonymity and wresting back control from the sellers, advertisers, and data intermediaries.[3]

Pooling of this kind has emerged in the industry as well. Google has experimented with "Federated Learning of Cohorts," (FLoC) a cookie-free method of IBA advertising (although FLoC was sunset in 2022 after concerns that this may still reveal consumer identities). As of December 2023, Google is working on "Topics", a similar initiative. Instead of cookies, "Topics" classifies users according to human-curated topics of interest; examples of such "Topics" include "/Autos & Vehicles/Vehicle Parts & Accessories/High Performance & Aftermarket Auto Parts" and "/Travel & Transportation/Hotels & Accommodations/Vacation Rentals & Short-Term Stays." The topics of interest to a consumer are disclosed to advertisers, without disclosing more granular information (such as browsing history, which could be used to identify the user).

In addition to the theoretical model discussed in section 2, we also survey (in section 3) the recent literature on these issues, from which a number of general insights has emerged. One is that the potential gains for consumers from targeting are essentially nil - whatever surplus increase from personalization obtains, is typically appropriated by the seller side. Second is that markets for data may fail (nor should we be surprised that it be so - data are rife with externalities, and the rivalry and excludability properties vary widely). Third is that decreasing the amount of data that is created and shared with the firms (whether sellers or data brokers) is generally welfare increasing; the mechanisms for doing so (opt-in policies, merging of data, and outright bans on data use) are varied. "Compensation for data" policies may be insufficient, because the amount of compensation for any consumer may not reflect the externality that consumer's data has on future consumers. Finally, pooling (whether by consumers - with different types choosing the same consumption bundle, by data intermediaries - by merging and anonymizing data, or by firms - by providing pooling contracts in equilibrium) has emerged as a critical feature. Pooling consumers together, merging data records, anonymizing web histories have a salutary effect on consumer welfare. While privacy is intrinsically valuable, it will generally increase consumer surplus; pooling (or otherwise anonymizing) data records addresses both motives.

[3]The General Data Protection Regulation (GDPR) "right to be forgotten" provision is one way of maintaining anonymity online; the examples above suggest that it is not the only way.

2 Forward Externalities: a Stiglitz (1977) Approach

2.1 Model

We use the Stiglitz (1977) model of a monopolist designing insurance contracts for sale to consumers of different risk types to illuminate this setting. The main insight that emerges is that if there is a sufficient amount of data, algorithms may allow the seller to go from standard price discrimination to first degree price discrimination using sequential screening.

Following the original work, notation, and interpretation, a consumer has a payoff-relevant type that is private information to her, and persistent across the two periods. The possible types are p_H and p_L, each reflecting a probability of accident or default, with $p_H > p_L$. The proportion of high-risk types is $\eta \in (0, 1)$. Individuals have wealth W_0; W_{NA} denotes wealth in the "no accident" state, while W_A denotes wealth if an accident does occur. The only difference between the types of consumers is the probability of accident; all consumers start with a level of wealth W_0, and, in case of an accident, suffer "damages" d. A monopolist insurer offers contracts to both types, characterized by a price α paid to the insurer in the no accident case, and a (net) premium β that the individual receives from the insurer in case of an accident. All individuals have a strictly concave utility function U over wealth. The monopoly equilibrium concept (or just "equilibrium') is a set of contracts such that no other set can generate higher profits, consumers maximize their utility by choice of contracts, and each consumer obtains at least the level of utility they would have if they purchased no insurance.

We suppose that interaction (contracts offers and purchases) is sequential, at one of two time periods ($t = 1, 2$); there is no discounting. This repetition of interactions with consumers, as well as the repetition of interaction with different consumers having obtained some information, is at the heart of the problem we study here.

Suppose that the fraction of high-risk individuals is low enough so that it it optimal to serve the low-risk types.[4] Denote the value functions by $V_H(\alpha, \beta) = U(W_0 - d + \beta)p_H + U(W_0 - \alpha)(1 - p_H)$, and $V_L(\alpha, \beta) = U(W_0 - d + \beta)p_L + U(W_0 - \alpha)(1 - p_L)$; in the (rare) cases where the times of the contract will be necessary to emphasize, they will be denoted by superscripts.

Correlated with the type is an observable, payoff-irrelevant, characteristic $\omega \in \{\omega', \omega''\}$ (for example, the location, the dates of travel, the time of purchase request). While the monopolist cannot see the type, it can see the observable characteristic, but does not know the mapping between the observable characteristic and the type. The observable characteristics are generated according to one of two distributions described by the following assumption:

[4] Without this assumption it becomes easier to support price discrimination.

Assumption 1. *Either* $\mathbb{P}(\omega'|p = p_H) = 1$, $\mathbb{P}(\omega''|p = p_H) = 0$, *or* $\mathbb{P}(\omega'|p = p_L) = 1$, $\mathbb{P}(\omega''|p = p_H) = 0$.

The monopolist doesn't know (and assigns some probability, a hyperparameter) which of the two characteristic generating distributions is the true one. Assumption 1 is questionable (although in light of the discussion on the amounts of kinds of data in the introduction, where we noted that the data are enough to identify not just the "type" of the individual, but the very identity of the individual, it may not be as extreme as it seems), but may be relaxed.

The timing is as follows:

1. Nature chooses the type of the consumer
2. Seller observes first time period ($t = 1$) characteristic, ω_1
3. Seller offers $t = 1$ contract(s)
4. Consumer chooses which contract, if any, to purchase
5. Seller observes the second time period characteristic ω_2
6. Seller offers $t = 2$ contracts
7. Consumer chooses which contract, if any, to purchase
8. Uncertainty about default is resolved, and payoffs are realized

The monopolist will also perform inference about the characteristic-generating process component, although the inference will be of a very simple kind. Once the characteristic is observed, the monopolist will offer contracts, and observe which individuals with a particular characteristic choose which contracts. If the contracts are separating, the monopolist can immediately infer the true type-characteristic map. If the contract is pooling, the monopolist will act à la Stiglitz (1977) in the second period.

It is straightforward to show that if the two types buy different contracts, the high type will be fully insured (Property 1 in Stiglitz (1977)), and that the low-risk type's utility level will be the same as if they had not purchased insurance (Property 2 in Stiglitz (1977)).

In the second period, one of two situations occurs. Either there is incomplete information (in which case the monopolist offers the separating contracts from Stiglitz (1977)), or there is complete information (in which case the monopolist offers each type a contract that fully extracts the surplus). It is easy to show that the latter case generates more profits for the monopolist. Figure 1 illustrates the full surplus-extracting contracts while figure 2 illustrated an example of the Stiglitz (1977) separating contracts.

There are several important "frictions" at play in this model: asymmetric information and market power (as in Stiglitz (1977)), as well as incomplete information (the inference made by the monopolist), and repeated interaction. The result we show stems from the conjunction of these four features. For instance, without repeated interactions (but retaining the other three frictions), because the monopolist does not know the characteristic-generating process, and thus, ω_1 becomes, in effect, an irrelevant signal, we are back in the Stiglitz (1977) world. Without market power (with two or more sellers) a new possibility arises - that the informational advantage the monopolist obtains in our setting will be competed away. While a full analysis of the repeated competitive setting is beyond the scope of this paper (which focuses on the monopoly case), the work of Ichihashi (2021a), Jin and Vasserman (2021), and Ali, Lewis, and Vasserman (2022), discussed more fully below, provides some caution to this intuition - all three papers show that *broadly speaking*, without additional assumptions, consumers may be worse off with data revelation (in our notation, the ω) than without it; competition alone may not help.

Shutting down the incomplete information (the type-characteristic map) component, and assuming that a monopolist interacts with consumers of unknown types over two (or more periods) leads to one of two situations. If the monopolist is able to identify individual consumers in the second period (i.e. the monopolist is able to record not just the contract purchased, but the very identity of the consumer in the first period), we would obtain the same result; this is, perhaps, not as enlightening, since if a monopolist is able to record the identity of a consumer, the full-surplus-extraction outcome is not terribly surprising. If, on the other hand, the monopolist is unable to record the identity of the consumer, we are in a world of simple repetition of the Stiglitz (1977) contracts, without any links across time periods. And finally, shutting down the asymmetric information friction (and supposing that neither the consumer nor the monopolist know p) leads to an odd situation: neither side has an informational advantage, and in addition, there is uncertainty about the type-characteristic map. It is hard to think of an application of such a model, and we do not develop it here. To sum up,[5] all of these frictions are necessary for the result developed here.

2.2 Naive consumers

Suppose first, following the literature on bounded rationality in industrial organization and behavioral economics,[6] that the consumer is *naive* - they make *statically* optimal choices, without internalizing the effect they have on their future options (menus of contracts, in this case). Naifs purchase the most attractive contract to them at every point in time, and do not internalize the externality they are exerting on their

[5] This discussion was suggested by a very helpful anonymous referee.

[6] The distinction between "naive" and "sophisticated" consumers of this kind - the former do not internalize the fact that their preferences may be time-inconsistent, while the latter do - first appeared in Strotz (1957). O'Donoghue and Rabin (1999, 2001) discuss these ideas in the context of behavioral economics, while Spigler (2011) discusses applications in industrial organization.

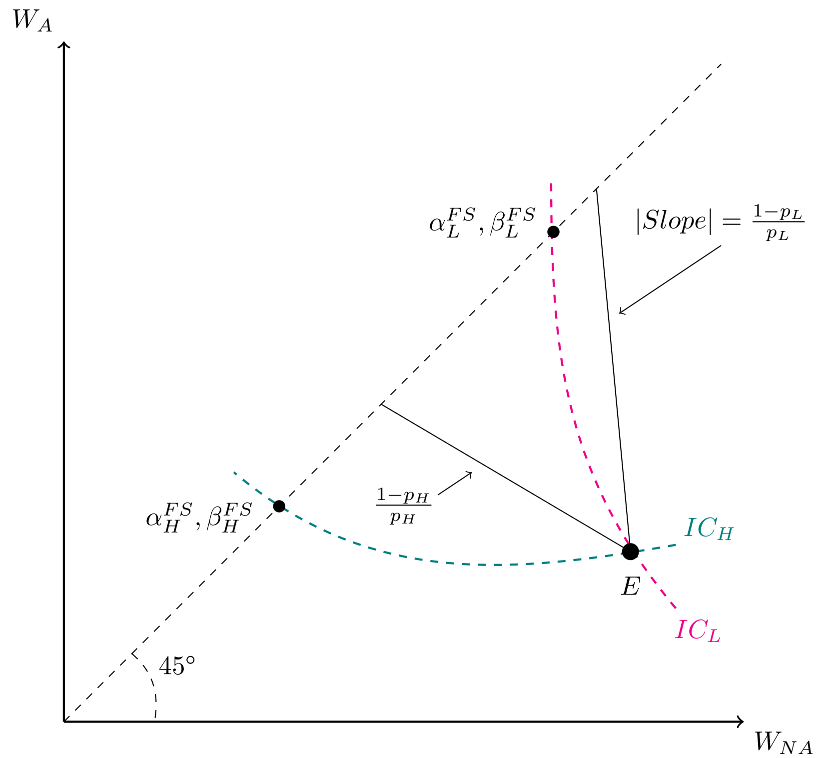

Figure 1: Full surplus extraction contracts

future selves.

In the first period then, the monopolist, knowing this, has to balance two considerations - profit-maximizing contract design, and the information value of the contracts. However, because the statically optimal contracts are separating to begin with, the monopolist can offer the same contracts, identify the types of consumers, and extract full surplus in the second period. Thus, naiveté on the part of consumers enables first-degree price discrimination.

2.3 Sophisticated consumers

Suppose now, that consumers are *sophisticated* - they are aware of the fact that the choices they make in the first period will affect the menu of contracts they will be offered in the second period.

In the second period the situation remains the same - the monopolist will offer either the Stiglitz (1977) contracts, or the full surplus extraction, personalized, first-degree price discriminatory contracts, depending on whether there is incomplete information, or complete information, respectively. Observe that in either case, the low-risk type obtains no surplus - either it is fully extracted (and they are fully insured), or they are partially insured, at terms which make her indifferent to not purchasing any insurance.

In the first period the situation is more complicated. Suppose that the Stiglitz (1977) contracts are also

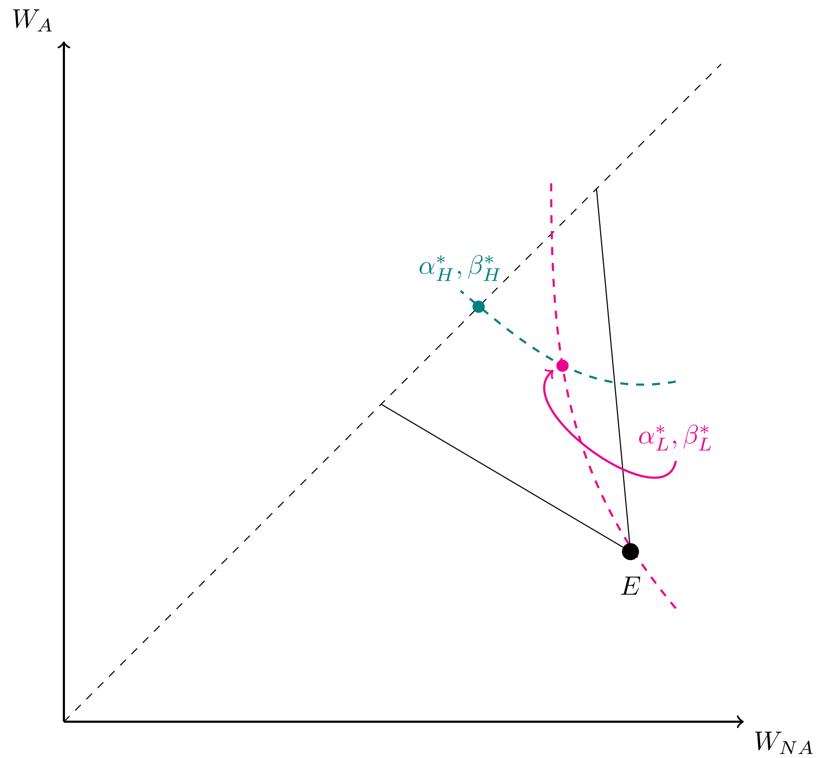

Figure 2: Stiglitz (1977) separating contracts

offered in the first period. If, as we have supposed, there are enough low-risk individuals to provide insurance for both types in the second period, the high-risk type earns some surplus in the second period. Thus, if the value of the surplus in the first period is high enough for the high type, they may prefer to switch to the low-risk type's contract in the first period (thus pooling in the first period, obtaining lower yet still positive surplus than from (α_H^*, β_H^*), and preserving the situation of incomplete information), and obtaining the surplus (α_H^*, β_H^*) today, getting identified as the high type, and obtaining zero surplus tomorrow. Figure 3 illustrates this situation. The incentive compatibility constraint ensuring separation for the high-type, provided that the Stiglitz (1977) contracts are available, is

$$V_H(\tilde{\alpha}_H^1, \tilde{\beta}_H^1) + V(\alpha_H^{FS}, \beta_H^{FS}) \geq V_H(\alpha_L^*, \beta_L^*) + V_H(\alpha_H^*, \beta_H^*) \tag{1}$$

The following proposition shows that a pooling component may be sustained as part of an equilibrium of the two-period game.

Proposition 1. *If consumers are sophisticated, there may be pooling in the first period.*

Proof. In Stiglitz (1977) the pooling is broken in the case where the pooling contract is not fully insuring, by finding a separating deviation that increases profits, and leaves both types indifferent between it and

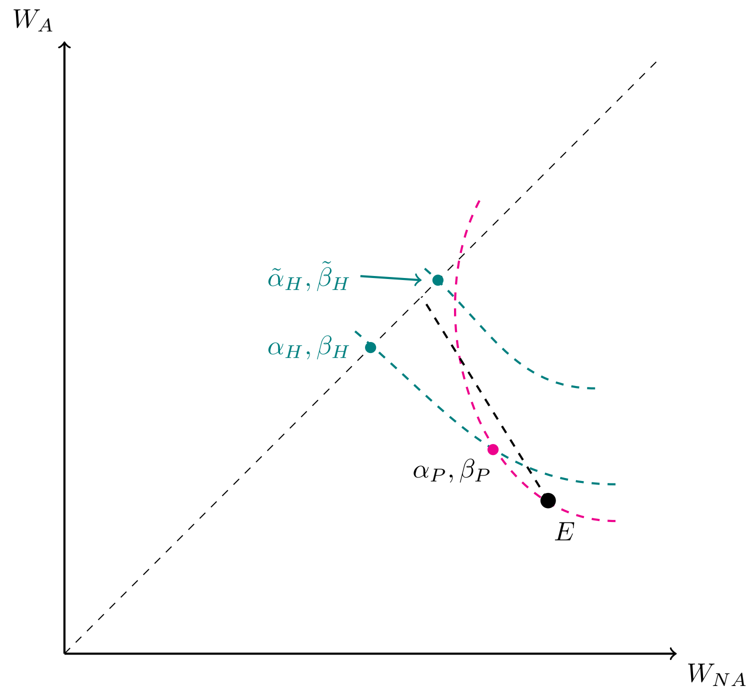

Figure 3: A pooling outcome may no longer be broken in the two-period model. In the Stiglitz (1977) model, a putative pooling equilibrium contract (α_P, β_P) is broken by offering (α_H, β_H) for the high type, and letting $(\alpha_P, \beta_P) = (\alpha_L, \beta_L)$. With two period, however, a sophisticated high-risk type would realize that separating in this period implies a loss in utility in the next period, and is no longer indifferent between (α_H, β_H) and (α_P, β_P), preferring to pool at (α_P, β_P). To generate separation, the contract for the high type has to give him a higher level of utility. However, if $(\tilde{\alpha}_H, \tilde{\beta}_H)$, with the property that $V(\tilde{\alpha}_H, \tilde{\beta}_H) + V(\tilde{\alpha}_H^{FS}, \tilde{\beta}_H^{FS}) = V(\tilde{\alpha}_P, \tilde{\beta}_P) + V(\tilde{\alpha}_H^*, \tilde{\beta}_H^*)$ is required, profits on $(\tilde{\alpha}_H, \tilde{\beta}_H)$ may be negative, and overall profits are ambiguous. A separating contract no longer breaks the pooling contract, and pooling may be sustained.

the pooling contract. In our setting, however, such a deviation is insufficient: it exists, but because of the informational effect of separation, the sophisticated high type would prefer to pool (and not separate), because while in the first period the separating contract leaves that type indifferent, in the second period the consequences of the separating contract are worse.

Denote by (α_P^1, β_P^1) the candidate pooling contract. If there exists $(\tilde{\alpha}_H^1, \tilde{\beta}_H^1)$ and $(\tilde{\alpha}_L^1, \tilde{\beta}_L^1)$ such that

$$V_H(\tilde{\alpha}_H^1, \tilde{\beta}_H^1) \geq V_H(\tilde{\alpha}_P^1, \tilde{\beta}_P^1) + V_H(\alpha_H^*, \beta_H^*) \tag{2}$$

$$V_L(\tilde{\alpha}_L^1, \tilde{\beta}_L^1) \geq V_L(\tilde{\alpha}_P^1, \tilde{\beta}_P^1) \tag{3}$$

$$
\begin{aligned}
&\eta\left[-p_H^2(\tilde{\beta}_H^1+\beta_H^{FS})+p_H(1-p_H)(\alpha_H^{FS}-\tilde{\beta}_H^1)+p_H(1-p_H)(\tilde{\alpha}_H^1-\beta^{FS})+(1-p_H)^2(\tilde{\alpha}_H^1+\alpha_H^{FS})\right]+ \\
&+(1-\eta)\left[-p_L^2(\tilde{\beta}_L^1+\beta_L^{FS})+p_L(1-p_L)(\alpha_L^{FS}-\tilde{\beta}_L^1)+p_L(1-p_L)(\tilde{\alpha}_L^1-\beta_L^{FS})+(1-p_L)^2(\tilde{\alpha}_L^1+_L^{FS})\right]\geq \\
&\geq\eta\left[-p_H^2(\beta_P^1+\beta_H^*)+p_H(1-p_H)(\alpha_H^*-\beta_P^1)+p_H(1-p_H)(\alpha_P^1-\beta_H^*)+(1-p_H)^2(\alpha_P^1+\alpha_H^*)\right]+ \\
&+(1-\eta)\left[-p_L^2(\beta_P^1+\beta_L^*)+p_L(1-p_L)(\alpha_L^*-\beta_P^1)+p_L(1-p_L)(\alpha_P^1-\beta_L^*)+(1-p_L)^2(\alpha_P^1+\alpha_L^*)\right]
\end{aligned}
\tag{4}
$$

i.e. constructing a separating contract following the observations made in Property 3 of Stiglitz (1977), such that the profits on the separating deviation plus the full surplus extraction are higher than the profits on the pooling contract plus the profits on the separating contract in the second period, then there is a separating contract in the first period. Eq. 4 express the need for lower profits today for the firm, compensated for by higher, fully extracting, profits tomorrow, to be greater than pooling profits today plus the incomplete information separating profits tomorrow, broken down by type. In essence, the monopolist is forced to offer rent (a better contract) to the sophisticated high-risk types today, to compensate them for the losses from surplus extraction tomorrow.

Eqs. 2, 3, and 4, if they are satisfied, generate a profitable deviation that constructs a separating set of contracts along the same lines as Stiglitz (1977): the low-risk type obtains a contract along her indifference curve through $(\tilde{\alpha}_P^1, \tilde{\beta}_P^1)$, to the north-west of $(\tilde{\alpha}_P^1, \tilde{\beta}_P^1)$. The high-risk type obtains a contract that is on a higher indifference curve than that through $(\tilde{\alpha}_P^1, \tilde{\beta}_P^1)$, reflecting the informational rent of this type. The monopolist is offering the high-risk types a better deal (thus its profits on this type fall), while offering the low-risk types a different deal that they like just as much as the putative pooling contract. Because the indifference curves equal the odds of the risk type at the 45° line and are flatter than that to the south-east of the 45° line, the profits on the low-risk types increase. Depending on parameters, the profit level may be greater than at the best pooling contract (determined below), and a separating equilibrium may be sustained. □

Thus, the monopolist will either offer a pooling contract at $t = 1$ (with a small positive, or zero, level of profits), followed by a separating contract at $t = 2$ (with positive profits), or a separating contract at $t = 1$ (with zero or negative profits), and a fully surplus-extracting contract at $t = 2$, with higher profits. Figure 4 illustrates the case with a zero-profit pooling contract at $t = 1$.

The conditions for separation in proposition 1 are more difficult to satisfy than the analogous conditions in Stiglitz (1977); one needs to raise the high type's utility (lowering $\tilde{\alpha}_H^1$, while making sure the monopolist's profits (which are increasing in $\tilde{\alpha}_H^1$) do not fall. If such an $(\tilde{\alpha}_H^1, \tilde{\beta}_H^1)$ does not exist, there may be pooling (on

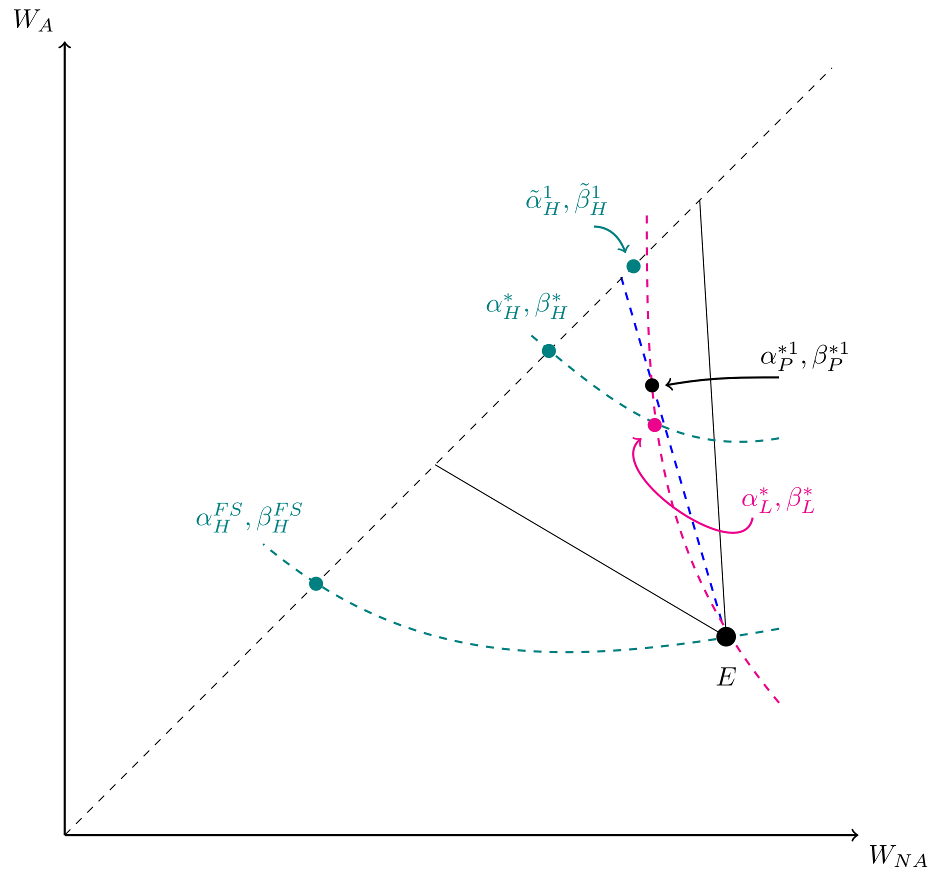

Figure 4: Two-period case with sophisticated consumers. (α^*_H, β^*_H) and (α^*_L, β^*_L) are (an example of) the Stiglitz (1977) separating contracts. With two periods, these contracts no longer separate. (α^*_L, β^*_L) and $(\tilde{\alpha}^1_H, \tilde{\beta}^1_H)$ *may* be equilibrium separating contracts. If profits from $\alpha^*_L, \beta^*_L)$ and $(\tilde{\alpha}^1_H, \tilde{\beta}^1_H)$ are lower than profits from a putative separating contract, there is a pooling equilibrium with $(\alpha^{*1}_P, \beta^{*1}_P)$; this would be the case if, as drawn, most individuals are low-risk.

a less-than-full insurance contract) in the first period. In this case, the high-risk type is also better off than in the static model, but this time the effect is purely mechanical, due to the fact that there are two periods, rather than one.

To find the optimal less-than-full insurance pooling contract (note that if the participation constraint for the low type is satisfied, it is also satisfied for the high type) we maximize

$$\max_{\tilde{\alpha}^1_P, \tilde{\beta}^1_P} (1-\overline{p})\tilde{\alpha}^1_P - \overline{p}\tilde{\beta}^1_P \tag{5}$$

$$\text{s.t. } V_L(\tilde{\alpha}^1_P, \tilde{\beta}^1_P) \geq V_L(0,0) \tag{6}$$

which are the same as the corresponding ones in Stiglitz (1977). The first-order condition for the problem is:

$$U'(W_0 - d + \beta)\frac{\frac{p_L}{1-p_L}}{\frac{\overline{p}}{1-\overline{p}}} = U'(W_0 - \alpha) \quad (7)$$

Thus, if a pooling contract is optimal (i.e., it generates higher profits for the monopolist), that contract is the $(\alpha_P^{*1}, \beta_P^{*1})$ that solves this problem.

2.4 Repeated Interactions: Forward Externalities and Surplus Extraction

Suppose now that either the situation in part 2.2 (with naive consumers) or the situation in part 2.3 (with sophisticated consumers) is repeated. There is a long-lived seller that lives for a (finite or infinite) number of periods T, who collects and can use the information about the mapping between types and observable characteristics - that is, data gleaned from transactions. Each "generation" of consumers lives for two periods; "new" consumers (i.e., those who have never interacted with this particular seller before), enter at $t = 3, 5, ...$, etc. The situation is identical in all other respects (i.e. in the risk and preferences parameters, in the timing of contract offers, and in the response), except for the fact that the monopolist now has additional information about the type-characteristic map.

If consumers are naive, as we saw in section 2.2, there is full surplus extraction in the second period. If the monopolist is able to use the information from the first two periods in subsequent periods, they will; thus, *all* future naive consumers will face first-degree price discrimination, as a result of the information gathered in the first period.

With sophisticates, the situation is somewhat more complex. Certainly, if there was separation in the first period with these consumers, there was full surplus extraction in the second period, and armed with this information, the monopolist will continue to extract surplus. Note that in this case, even a "new" sophisticated consumer - one who has never interacted with this seller before - will face a personalized price, unlike a sophisticated consumer at time $t = 1$. Thus, the earlier sophisticates exert a forward externality on later ones. Furthermore, while being a sophisticated type increased the payoff of the high type of the consumer, while leaving the low type as well off, thus "sophistication benefit" completely fades with repeated interactions: even the most sophisticated of consumers are offered personalized prices.

Finally, consider the situation with sophisticated consumers where there was pooling in the first period (and separation in the second). Once the types reveal themselves by choosing separating contracts in the second period, the monopolist observes the link between choices and observable characteristics. *It is at this point that the addition of observable characteristics plays a role.* Because the map between types and characteristics is persistent and deterministic, from time period three onward, the seller no longer has to offer the pooling contracts (even if one was optimal in the first period).[7]

Thus, even sophisticated consumers who would have obtained a pooling contract, had they interacted with the seller at time $t = 1$ no longer do so at $t = 3$. The monopolist will offer fully surplus extracting contracts, basing the offer on the characteristic it observed before choosing contracts. The outcome facing the consumer from time $t = 3$ onward is the same as the outcome that faced the naive consumer at time $t = 2$ onward. The fact that observable characteristics of one agent are informative about the type of another agent has enabled full surplus extraction.

3 Review of the Literature

The recent literature on the issue of data, privacy, and data-related price discrimination has developed quickly; we survey this literature here, focusing on more recent work. Bergemann and Bonatti (2019), who survey the literature on markets for data, and Acquisti, Taylor, and Wagman (2016), who survey the literature on the economics of privacy, provide excellent recent surveys, and we make no attempt to subsume their work, reviewing only the relevant work since then.

Virtually all of the recent work (Courty Hand Li (2000), Acquisti and Varian (2005), Conitzer, Taylor and Wagman (2012), Belleflamme and Vergote (2016), Choi, Jeon, and Kim (2019), Bergemann, Brooks, and Morris (2015), Ichihashi (2019), Bergemann, Bonatti, and Gan (2022), Galperti, Levkun, and Perego (2022), Yang (2022)) focuses on the case of a monopolist seller; there may or may not also be a data intermediary. Acemoglu et al. (2022) study a monopolist data intermediary (without explicitly including a seller). By contrast, Chen, Choe, and Matsushima (2020) study competition among firms who can provide individual price offers to consumers and consumers who may manipulate this fact, Anderson, Baik, and Larson (2022) study Bertrand competition, and Bergemann and Bonatti (2015) study a competitive setting with a single data intermediary. Ali, Lewis, and Vasserman (2022) study both the monopoly and competitive benchmarks, while Jann and Schottmüller (2019) study a game-theoretic model (without players readily analogous to firms and consumers). Ichihashi (2021) studies competition between data intermediaries. Acemoglu et al. (2022), Bergemann, Bonatti, and Gan (2022) Galperti, Levkun, and Perego (2022), Ichihashi (2021a), and Yang (2022) explicitly model prices for data.

[7] The assumption of a deterministic mapping between types θ and observables ω may be relaxed. In this case, the monopolist will not offer the sully surplus extracting contracts in subsequent periods. However - as the monopolist's beliefs about the map between the characteristic and the type are converging, (one of) the optimal contracts will converge to the fully surplus-extracting contract for that type. Thus, if it is the probability of the low type that is high and increasing, the optimal separating contract will move along the indifference curve of the low type through E up and to the left, toward the point where the indifference curve crosses the 45°line; because the slope of the indifference curve is equal to $\frac{1-p_L}{p_l}$ there, and is lower than that at less than full insurance, profits on the low type are increasing. Of course, the intersection of these two lines is the contract that would be offered to the low type with complete information.

Bergemann, Brooks, and Morris (2015) foreshadow some of the literature to follow (although they do not focus on data or externalities); they study the problem of a monopolist armed with additional information about which "segment" of the market a consumer belongs to; the monopolist then engages in third-degree price discrimination. They prove an anything-goes result for the monopoly setting - virtually any reasonable split of the surplus from trade (with a monopolist, so that the producer surplus cannot be less than the monopoly outcome) is possible.

Most similar to the present discussion[8] is the work of Choi, Jeon, and Kim (2019) and Conitzer, Taylor, and Wagman (2012). Motivated by discussions of privacy Choi, Jeon, and Kim (2019) consider a monopolist seller who collects data and sets prices for content that it sells to consumers. They compare the social planner's solution to this setting with the monopoly outcome, also finding that if informational externalities are large enough, the monopolist will collect too much data relative to the socially optimal benchmark (and serve too many consumers).

Relative to Choi, Jeon, and Kim (2019), who *assume* that data presents externalities (their "nuisance cost" function, capturing the disutility of data loss in case of a breach, or the intrinsic preference for privacy), the result here obtains endogenously. Echoing the findings of Choi, Jeon, and Kim (2019), the present analysis also shows that it is not enough for consumers to be "sophisticated," or to be provided with additional information to help guide their choices: the externalities persist even in these cases. Thus, simply "opting out" would not sufficiently mitigate the negative data externalities; an outright ban (akin to the ban on the use of genetic information, or strong protections against the use of data by third parties that are included in the current General Data Protection Regulation of the European Union) may be a better alternative.

Conitzer, Taylor, and Wagman (2012) study a monopolist selling to individuals with different privately known valuations; their focus is on whether anonymizing past purchase history (possibly at a cost) has an effect on outcomes. A pooling equilibrium exists in their model as well, and they find that privacy can improve welfare. Belleflamme and Vergote (2016) also study a setting where individuals can hide their data from the monopolist; they find that although individually, some consumers are better off without data tracking, the total consumer surplus may fall, because the monopolist raises its prices for those who do hide their data (and perhaps, excludes some individuals from purchasing). Thus, their work is suggesting that bans on tracking technology may be counterproductive (although participating in such a ban is not individually rational for all consumers).

Hidir and Vellodi (2021) study a monopolist interacting with consumers who can endogenously generate segmentation, finding that the consumer-optimal segmentation is the "least informative segmentation

[8] Although most, if not all, of the papers discussed here posit uncertainty about consumer valuations, as opposed to the probability of an accident, with the corresponding benefit payment.

that guarantees trade." They also illustrate how a cheap talk environment can implement this outcome. Thus, provided that trade takes place, the outcome that is best for consumers is the outcome that provides the seller with the least amount of information, and this outcome is achieved by consumers pooling together. Again, pooling arises as a pro-consumer force.

Yang (2022) reports a finding that is complementary to the results discussed above - in a (mechanism design) setting with a monopolist and consumers with differing valuations the *data intermediary-optimal* outcome results in (essentially) full surplus extraction - consumers pay their values. Thus, if consumer-optimal outcomes provide as little information as possible (conditional on trade), data intermediary-optimal outcomes are the "flip side of the coin" - they provide enough information for full surplus extraction by the seller. Ichihashi (2020) also shows that in a setting where a firm can recommend products to consumers, consumer and producer surplus crowd each other out - the consumer is better off whenever the monopolist is worse off.

Ali, Lewis, and Vasserman (2022) study whether with consumer control of data, which they model as the ability to disclose verifiably something about their true type, can improve welfare. They study both the monopoly and the competitive market settings with heterogeneity in the buyers' valuations. In the monopolist case, consumers never benefit from disclosure; with more flexibility in disclosure (for instance, if lower value types can credibly distinguish themselves from higher value types), disclosure may result in welfare gains. The sufficient condition for existence of such a disclosure system is relatively demanding - a joint condition on the set of messages that the types can and cannot send, sa well as the distribution of types, requiring that some higher value types cannot send messages associated with low value types, and that upon observing the message associated with higher types, the seller does not raise the price for those types relative to the optimal uniform price. In the competitive setting, Ali, Lewis, and Vasserman (2022) construct equilibria that are beneficial for consumers, and, under relatively strong assumptions on the distribution of consumer types, all types benefit in these equilibria, relative to uniform pricing. The construction and the welfare results rely on the ability of consumers to disclose different information to different firms (such selective disclosure will drive competition between firms); many current tracking systems do not allow for such disclosure. Indeed, the presence of data intermediaries may nullify this result; as the authors also point out, so would sharing data between firms.

Galperti, Levkun, and Perego (2022) study an information design model with a monopolist seller, heterogeneous buyers, and observable characteristics that are revealed to a data intermediary, characterizing the value of "data records" - collections of observable characteristics - for the data intermediary. They, too, identify a novel externality of data: for the platform, the value of "lower" valuation records is valuable (even if it generates zero transaction value), because the presence of these records allows the intermediary to "hide" the higher value records, and thus provide value to the seller, without fully revealing the valuations. In Appendix F of their work, Galperti, Levkun, and Perego (2022) present an example that is similar the setting discussed here; in their work, in a static example, if the data intermediary acts in the interests of the seller and maximizes the seller's profits, they obtain the same conclusion - such correlation strictly increases the profits of the seller, at the cost of lowering the surplus of the consumer. Galperti, Levkun, and Perego (2022) also argue in favor of compensation for data for the individuals who generated it.

Also working within an information design framework, Bonatti and Cisternas (2023) show that better information on the part of the firms harms naive consumers, but may help sophisticates (the intuition is very similar to the result discussed here). Their results imply that not only should information that firms have on consumers be made available to the latter, but that to the extent that "sophistication" may be promoted, it should be.

Ichihashi (2021b) studies competition between data intermediaries, and provides and important caveat to the discussion around data - even with such competition, data intermediaries may sustain monopoly outcomes in the market for data, while firms that sell services to the consumer use the data to extract surplus from their transactions. Ichihashi (2021a) studies purchases of data from consumers who have no ex-ante additional information but can commit (in the same sense) to release future information; the firm "learns" some information, and consumers obtain utility from their information sales decision, as well as a function of the firm's information. The welfare implications of this are ambiguous, turning on whether data are complements or substitutes.

Anderson, Baik and Larson (2022) support banning IBA targeting (under some conditions), and show that opt-in (rather than opt-out) policies are welfare improving. Belleflamme and Vergoute (2016) and Chen, Choe, and Matsushima (2020) reach an opposite conclusion from Anderson, Baik, and Larson (2022). In the latter paper more consumer control is generally better for consumer surplus, in the former two it is generally worse. Jones and Tonetti (2020) also argue for consumer control, in a general equilibrium model. Jann and Shottmüller (2019) argue strongly in favor of privacy in a game-theoretic model, showing that privacy dominates a lack of privacy in the Pareto sense (under mild conditions).

On the empirical side, Aridor, Che, and Salz (2023), and Jin and Vasserman (2021) investigate the impact of data on welfare. Aridor, Che, and Salz (2023) look at the data provided by an intermediary who contracts with online travel agencies and meta-search engines. They find that consumer who opt out exert an externality on those who opt in; their interpretation of the effect of this is consistent with the work discussed here - while opt-in consumers *may* be better off (due to the improvement in services they receive), if data is used for surplus extraction, they will be worse off. Jin and Vasserman (2021) use data that has both screening, moral hazard, and competitive elements - a program where consumers may opt in to be tracked by their auto insurer, in exchange for possible discounts. Opt-in rates are low, and their structural estimate is that the average consumer incurs a *loss* of \$93 from being monitored. (Interestingly, safer drivers not only self-select into monitoring, but become significantly safer still while enrolled.)

4 Concluding Remarks

There is an argument in law and economics that privacy can be welfare decreasing (see, *nter alia*, Posner (1981) and Stigler (1980)) since more information results in better outcomes. The discussion in the economics literature since then, and more specifically, in the literature on the economics of information, surveyed most recently in Stiglitz and Kosenko (2024a, 2024b), has clarified that this argument is incorrect.

Contemporary proponents of data creation argue that better data allows providers to personalize their services, thus increasing the value consumers derive from their services. In this view, data facilitates high-value matches (say, in a two-sided market). While it may be true that better data may enable the firms to find good matches for their consumers, the recent literature, as well as the present analysis, have shown that the same force allows the seller to appropriate the (increased) surplus from the transaction. Sellers do find the "right" buyer - but they also exploit their information, leaving the buyer little, if any, of the (increased) value of the match.

It is indeed, rare, that a multitude of papers, each approaching the topic from a somewhat different angle, with different assumptions and setting, reach, on the whole, quite similar conclusions. In much the same way that the analysis of markets with imperfect and asymmetric information has upended economists' understanding of market economies, leading to an appreciation that markets on their own are virtually never efficient, the analysis of markets for data, the analysis of privacy, and the resulting price discrimination has established that markets for data may fail drastically. Privacy is not only intrinsically valuable, but increases consumer surplus; policies that provide consumer control ("opt-in/opt-out" choices), or compensate consumers for data, are beneficial but may not restore efficiency; the gains from personalization for consumers are more than offset by the losses due to surplus extraction. Finally, merging or otherwise anonymizing data, and sustaining pooling equilibria, have emerged as a critical tools on the arsenal of consumers and regulators.

References

[1] 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. DOI: 10.1257/mic.20200200

[2] Acquisti, Alessandro, Curtis Taylor, and Liad Wagman. (2016). The Economics of Privacy. *Journal of Economic Literature*, 54 (2): 442-92. DOI: 10.1257/jel.54.2.442

[3] Acquisti, Alessandro and Hal R. Varian, (2005). Conditioning Prices on Purchase History. *Marketing Science*, 24(3):367-381. https://doi.org/10.1287/mksc.1040.0103

[4] Ali, S. Nageeb, Greg Lewis, Shoshana Vasserman. (2023). Voluntary Disclosure and Personalized Pricing, *The Review of Economic Studies*, 90(2), Pages 538-571, https://doi.org/10.1093/restud/rdac033

[5] Simon Anderson, Alicia Baik, Nathan Larson, Price Discrimination in the Information Age: Prices, Poaching, and Privacy with Personalized Targeted Discounts, *The Review of Economic Studies*, 90(5), Pages 2085-2115, https://doi.org/10.1093/restud/rdac073

[6] Aridor, G., Che, Y.-K. and Salz, T. (2023), The effect of privacy regulation on the data industry: empirical evidence from GDPR. *The RAND Journal of Economics*, 54: 695-730. https://doi.org/10.1111/1756-2171.12455

[7] Belleflamme, Paul, and Wouter Vergote (2016). Monopoly price discrimination and privacy: The hidden cost of hiding. *Economics Letters*, 142, Pages 141-144. https://doi.org/10.1016/j.econlet.2016.10.027.

[8] Bergemann, Dirk, and Alessandro Bonatti. (2015). "Selling Cookies." *American Economic Journal: Microeconomics*, 7 (3): 259-94. DOI: 10.1257/mic.20140155

[9] Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin. (2018). The Design and Price of Information. *American Economic Review*, 108 (1): 1-48. DOI: 10.1257/aer.20161079

[10] Bergemann, D., Bonatti, A. and Gan, T. (2022), The economics of social data. *The RAND Journal of Economics*, 53: 263-296. https://doi.org/10.1111/1756-2171.12407

[11] Bonatti, Alessandro and Gonzalo Cisternas. (2020). Consumer Scores and Price Discrimination, *The Review of Economic Studies*, 87(2), Pages 750-791, https://doi.org/10.1093/restud/rdz046

[12] Zhijun Chen, Chongwoo Choe, Noriaki Matsushima (2020) Competitive Personalized Pricing. *Management Science* 66(9):4003-4023. https://doi.org/10.1287/mnsc.2019.3392

[13] Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim. (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics*, 173. https://doi.org/10.1016/j.jpubeco.2019.02.001

[14] Conitzer, Vincent, Curtis R. Taylor, Liad Wagman. (2012). Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases. *Marketing Science* 31(2):277-292. https://doi.org/10.1287/mksc.1110.0691

[15] Consumer Financial Protection Bureau (2022). "Buy Now, Pay Later: Market trends and consumer impacts." https://files.consumerfinance.gov/f/documents/cfpb_buy-now-pay-later-market-trends-consumer-impacts_report_2022-09.pdf

[16] Courty, P., & Li, H. (2000). Sequential Screening. *The Review of Economic Studies*, 67(4), 697-717. http://www.jstor.org/stable/2695944

[17] Elliott, Matthew and Galeotti, Andrea and Koh, Andrew and Li, Wenhao, Market Segmentation Through Information. http://dx.doi.org/10.2139/ssrn.3432315

[18] Fainmesser, Itay P., Andrea Galeotti. (2016). Pricing Network Effects, *The Review of Economic Studies*, 83(1), Pages 165-198, https://doi.org/10.1093/restud/rdv032

[19] Fainmesser, Itay P., and Andrea Galeotti. 2020. "Pricing Network Effects: Competition." *American Economic Journal: Microeconomics*, 12 (3): 1-32. DOI: 10.1257/mic.20170226

[20] Fainmesser, Itay Perah, Galeotti, Andrea, and Momot, Ruslan, Digital Privacy (September 25, 2019). HEC Paris Research Paper No. MOSI-2019-1351. http://dx.doi.org/10.2139/ssrn.3459274

[21] Farboodi, Maryam and Veldkamp, Laura (2022). Data and Markets. MIT Sloan Research Paper No. 6887-22. http://dx.doi.org/10.2139/ssrn.4284192

[22] Federal Trade Commission. (2022). "Case No. 2:22-cv-377 Complaint for Permanent Injunction and Other Relief". https://www.ftc.gov/system/files/ftc_gov/pdf/1.%20Complaint.pdf

[23] Galperti, Simone, Aleksandr Levkun, Jacopo Perego. (2023). The Value of Data Records, *The Review of Economic Studies*. https://doi.org/10.1093/restud/rdad044

[24] Google, Inc. (2023). Topics API overview. https://developers.google.com/privacy-sandbox/relevance/topics

[25] Sinem Hidir, Nikhil Vellodi. (2021). Privacy, Personalization, and Price Discrimination, *Journal of the European Economic Association*, 19(2), Pages 1342â1363, https://doi.org/10.1093/jeea/jvaa027

[26] Ichihashi, Shota. 2020. "Online Privacy and Information Disclosure by Consumers." *American Economic Review*, 110 (2): 569-95. DOI: 10.1257/aer.20181052

[27] Ichihashi, Shota. (2021a). The Economics of Data Externalities. *Journal of Economic Theory* 196C. https://doi.org/10.1016/j.jet.2021.105316.

[28] Ichihashi, S. (2021b), Competing data intermediaries. *The RAND Journal of Economics*, 52: 515-537. https://doi.org/10.1111/1756-2171.12382

[29] International Monetary Fund. (2022). "Pandemicâs E-commerce Surge Proves Less Persistent, More Varied" https://www.imf.org/en/Blogs/Articles/2022/03/17/pandemics-e-commerce-surge-proves-less-persistent-more-varied

[30] Jann, Ole, and Christoph Schottmüller. (2020). An Informational Theory of Privacy, *The Economic Journal*, 130(526), Pages 93â124, https://doi.org/10.1093/ej/uez045

[31] Jin, Yizhou and Shoshana Vasserman. (2021). Buying Data from Consumers The Impact of Monitoring Programs in U.S. Auto Insurance. Working paper.

[32] Jones, Charles I., and Christopher Tonetti. 2020. "Nonrivalry and the Economics of Data." *American Economic Review*, 110 (9): 2819-58. DOI: 10.1257/aer.20191330

[33] Leitzinger, J. & J. E. Stiglitz. (1984). Information Externalities in Oil and Gas Leasing. *Contemporary Economic Policy Issues*, 5, pp. 44-57.

[34] Miklós-Thal, Janine, Avi Goldfarb Avery M. Haviv Catherine Tucker. (2023). Digital Hermits. NBER Working Paper series 30920. http://www.nber.org/papers/w30920

[35] https://www.nasdaq.com/articles/e-commerce-spending-tops-$1-trillion-in-2022

[36] O'Donoghue, T., and Rabin, M. (1999). Doing It Now or Later. The American Economic Review, 89(1), 103-124. http://www.jstor.org/stable/116981

[37] O'Donoghue, Ted and Matthew Rabin. (2001). Choice and Procrastination, *The Quarterly Journal of Economic*, 116(1), Pages 121-160, https://doi.org/10.1162/003355301556365

[38] Rothschild, M., and Stiglitz, J. (1976). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. The Quarterly Journal of Economics, 90(4), 629-649. https://doi.org/10.2307/1885326

[39] Office of the Comptroller of the Currency (2023a). "Consumer Use of Buy Now, Pay Later: Insights from the CFPB Making Ends Meet Survey". https://files.consumerfinance.gov/f/documents/cfpb_consumer-use-of-buy-now-pay-later_2023-03.pdf

[40] Office of the Comptroller of the Currency (2023b). "Retail Lending: Risk Management of 'Buy Now, Pay Later' Lending". https://www.occ.gov/news-issuances/bulletins/2023/bulletin-2023-37.html

[41] Posner,R.A (1981). The economics of privacy. *American Economic Review*, 71(2), pp.405â9.

[42] Spigler, Ran. (2011). *Bounded Rationality and Industrial Organization*. Oxford University Press.

[43] Stigler, G. J. (1980). An Introduction to Privacy in Economics and Politics. *The Journal of Legal Studies*, 9(4), 623-644. http://www.jstor.org/stable/724174

[44] Stiglitz, J. E. (1974a). Theories of Discrimination and Economic Policy. *Patterns of Racial Discrimination*, G. von Furstenberg, et al. (eds.), D.C. Heath and Company (Lexington Books), pp. 5-26

[45] Stiglitz, J. E. (1974b). Alternative Theories of Wage Determination and Unemployment in L.D.C.'s: the Labor Turnover Model. *Quarterly Journal of Economics*, 88(2), pp. 194-227.

[46] Stiglitz, J. E. (1977). Monopoly, Non-Linear Pricing and Imperfect Information: The Insurance Market. *The Review of Economic Studies*, 44(3), 407-430. https://doi.org/10.2307/2296899

[47] Stiglitz, Joseph E. (2001). Biographical. https://www.nobelprize.org/prizes/economic-sciences/2001/stiglitz/biographical/

[48] Stiglitz, Joseph E. and Andrew Kosenko (2024a). Robust Theory and Fragile Practice: Endogenous Information in a World of Disinformation Part 1: Indirect Communication (with Joseph E. Stiglitz). *Elgar Companion to Information Economics*, chapter 2. Eds. Daphne R. Raban and Julia Wlodarczuk. Edward Elgar Publishing.

[49] Stiglitz, Joseph E. and Andrew Kosenko (2024b). Robust Theory and Fragile Practice: Endogenous Information in a World of Disinformation Part 2: Direct Communication (with Joseph E. Stiglitz). *Elgar Companion to Information Economics*, chapter 3. Eds. Daphne R. Raban and Julia Wlodarczuk. Edward Elgar Publishing.

[50] Strotz, R. H. (1957). The Empirical Implications of a Utility Tree. *Econometrica*, 25(2), 269-280. https://doi.org/10.2307/1910254

[51] Yang, Kai Hao. (2022). Selling Consumer Data for Profit: Optimal Market-Segmentation Design and Its Consequences. *American Economic Review*, 112 (4): 1364-93. DOI: 10.1257/aer.20210616