Information Choice: Cost over Content

Kristof Madarasz (LSE) and Marek Pycia (U Zurich)[*]

Abstract

When supplying information, agents choose between options that differ both in their contents and in their costs. We establish a "cost-over-content" theorem for a large class of dynamic trading environments where buyers choose from arbitrary sets of processes (experiments) that reveal information to the seller. When all experiments are equally costly, choosing any given experiment is a perfect equilibrium. However, when experiments differ in costs, there is a unique equilibrium: all buyers choose the cheapest experiment, regardless of the information it provides. We explore implications for market performance, privacy, data sale, and defaults in market regulation.

Keywords: Information Design, Signalling, Dynamic Pricing, Platforms, Privacy Paradox, Defaults.

[*]June 2023. First presentation 2016. Previously titled "Towards a Resolution of the Privacy Paradox," online on bepress™ (Digital Commons) archive, https://works.bepress.com/kristof_madarasz/47/, since 2020, also CEPR Discussion Paper 16873. For their comments, we would like to thank Dilip Abreu, Jean-Michel Benkert, Ernst Fehr, Elisabetta Iossa, Stephen Morris, Michaela Pagel, Heather Sarsons, Armin Schmutzler, Balázs Szentes, Rakesh Vohra, Cédric Wasser, and the audiences at LSE Job Talk Lunch Seminar, UCLA, U Zurich, and SSES. We drafted this paper while we were both visiting the William S. Dietrich II Economic Theory Center at Princeton University, and we would like to thank the Center for its hospitality.

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program grant agreement No 866376.

1 Introduction

A central insight of microeconomics is that private information can have a profound impact on the efficiency of trade and on the allocation of the surplus between buyers and sellers. For instance, when deciding what prices to charge, both in the context of static or dynamic price discrimination, profit-maximizing sellers face a trade-off between the efficient allocation of their products and the provision of information rents to their consumers. This trade-off, and thus the volume of trade and consumer welfare, critically depends on just how much private information buyers have. Similarly, the owner of an asset may secure a sale if she can credibly signal that her asset is of sufficiently high quality. At the same time, trade may well break down in the absence of such disclosure even if it is common knowledge that there are benefits from trade.

A privately informed party can, however, often affect what her trading partner learns about her information both before and during trade. For example, buyers can decide how much information to disclose about themselves by selecting how much of their online search is traceable by the seller. Similarly, as in pure signalling, a worker can exert costly but economically wasteful effort to credibly communicate her ability, and also her outside option, to an employer. It is then natural to ask how shall a privately informed party choose the information she supplies to her trading partner and what will be the result of such endogenous determination of the distribution of information on the terms of trade.

Choosing the information others receive is also the question at the center of the rapidly growing and influential literature on information design. Information design focuses on a sender's optimal choice of information within a given game. The sender (e.g., the buyer) chooses an information structure or experiment to determine what the receiver (e.g., the seller) knows before they take an action, cf. Kamenica and Gentzkow (2012) and Bergemann and Morris (2019). This approach has been applied to a variety of problems including a class of sender-receiver games where it is often referred to as Bayesian persuasion. The focus of information design, however, is on a sender who chooses an experiment before possessing any private information about the payoff state.[1] While this assumption holds true in many contexts, in other contexts, a sender frequently possesses private information prior to deciding on the information to supply.

Our paper revisits information design in the classic monopoly problem where a seller faces a buyer at one or more time periods. The buyer has private information about her valuation for the seller's product and the seller makes price offers. Before starting their interaction, the buyer can choose from an arbitrary finite set of dynamic information structures or experiments. Each experiment determines the information that the seller receives over time. We also permit experiments to deliver news, either privately or publicly, regarding the timing of information arrival itself. We further allow different experiments to have different costs. For instance, the buyer may need to pay a fee to prevent an app from tracking her, thereby preserving her privacy. Alternatively, disclosing some information correlated with her preferences may be costly for her. Despite the complexity of this setting with dynamic pricing, flexible observation structures, and arbitrary information designs, where a given information design may help some (or all) consumers, while another could hurt them, we present two robust insights on the endogenous supply of information.

First, we show that when experiments do not differ in their costs, then anything goes. The buyer choosing *any* one of the experiments irrespective of her private information is a perfect Bayesian equilibrium. Specifically, the buyer's ex ante choice of experiment, as examined in classic information design, remains an equilibrium information structure even if she makes the choice at the interim stage, that is, when she knows the payoff state.

Second, we show, however, that when experiments do differ in their costs, there is instead a unique perfect Bayesian equilibrium. The buyer opts for the least expensive experiment, regardless of its information content. This finding implies that all buyer types select the least ex-

[1] This is referred as the commitment assumption in that the sender picks a state-contingent signal distribution before learning the payoff state, even though the actual message sent is fully conditioned on the payoff state.

pensive experiment, regardless of whether it fully reveals the buyer's preferences, provides information about some types but not others, or maintains perfect privacy. Given an arbitrary set of potential information designs, informed buyers select experiments purely on the basis of their costs and not on their informational content.

Third, we show that these two results remain true in a variety of extensions of our base setup: two-sided payoff uncertainty, random allocation of the bargaining power, dynamic choice of information design, general observation structures, and the presence of common values.

For instance, given common values, the information supplied by the buyer directly impacts the seller's preferences over trade, e.g., the seller's cost of serving the buyer might directly depend on the buyer's type. With common values there may be no trade in the absence of information revelation due to classic selection problems (as, e.g. in Akerlof 1970), while there can be substantial trade if more information was offered, for example if the buyer chose some form of noisy or partial information disclosure. However, the buyer still invariably opts for the least expensive experiment in equilibrium.

Similarly, when the sender makes information design choices over time, we show that she always chooses the cheapest design in any given period where such a choice may arise. Consequently, small inter-temporal redistributions of information costs may result in dramatic shifts in the supply of information and the efficiency of trade. In sum, we establish an informational irrelevance result for a large class of trading environments. When prices respond flexibly to information, senders will always choose the lowest cost option irrespective of its content.

Although our main result says nothing about the kind of information design the sender will end up choosing, in many contexts there is a natural link between the information content of an experiment and its cost. For instance, considering the literature on voluntary disclosure (e.g., Grossman 1981, Milgrom 1981, Verrecchia 1983), our findings imply that if partial or full disclosure incurs costs while non-disclosure is free, senders will opt not to disclose in the unique equilibrium. Conversely, when there is a cost associated with privacy protection, and the default is full transparency, the equilibrium preferences fully reverse — even though the sender is fully aware of the potential large difference in her equilibrium payoffs.

In Section 4, we then consider an application of the above by embedding it into a simple market setting where the buyer's privacy options, including cost differences, emerge endogenously. In this context we first point out that even when the buyer has the right to choose what information to share with the seller, unless full privacy protection is the directly cheapest option, the buyer will never choose that.

We then endogenize the supply and the pricing of experiments. We consider privacy platforms (experiments) offered by a platform provider for online shopping. The platform provider can choose from an arbitrary set of technologically feasible privacy platforms. She decides which platforms to offer and at what price each. The platform provider also contracts with the seller and agrees on some profit-sharing, that is, on a contract where in exchange for the information provided to the seller, the seller's transfer to the platform provider is increasing in the seller's realized profit. Finally, buyers have the choice to select a platform, to shop offline directly from the seller, or to exit the market. In this context we show that in equilibrium the platform provider offers platforms that maximize the value of the information passed to the seller, as measured in the seller's profit, and offers these platforms for free.

Our paper then finds that in a basic market setting buyers lack the proper equilibrium incentives to protect their privacy. At the same time, parties that benefit from acquiring information are willing to incur substantial costs to do so, provided competition does not fully restrict their ability to extract surplus from consumers. This sharp asymmetry has various consequences. It implies that the opportunity to engage in data trade can often cause buyers to be worse off compared to where such data trade is prohibited. This observation may then shed light on how the widespread ability to track consumers can have very significant impact on the distribution of gains from trade even when consumers seemingly have access to cheap methods of protecting such informational rents. Furthermore, policies that directly regulate information gathering, or information sale may be considered to safeguard consumer welfare. In the presence of transaction costs or regulated trade, our result also points to the power of defaults in information choice and their impact both on the efficiency of trade and the distribution of the benefits arising from it. We discuss theoretical extensions and such policy considerations in the Conclusion.

1.1 Related Literature

Our paper relates to various strands in the literature.

First, as mentioned, our paper links to the expanding literature on information design, e.g., Bergemann and Morris (2019), which considers one party's (the information designer's) observable choice about the information structure in the game. An instance of this is the 'Bayesian persuasion' problem where a sender picks the experiment for a receiver who then takes an action affecting the sender's payoff. In this context, when the full space of experiments is available, Kamenica and Gentzkow (2012) canonized the solution method of concavification and Gentzkow and Kamencia (2014) considered costly information design, e.g., with the cost of each experiment linked to its respective entropy reduction, and showed that key insights of this literature continue to hold.

Bergemann, Brooks, and Morris (2015) applied the information-design approach to the classic static monopoly problem. They study the welfare implications of the information the seller obtains about the buyer's private valuation. Their results imply that if the buyer picked what the seller learned about the buyer preferences, before the buyer herself knew anything about her own preferences, then the maximal payoff the buyer can achieve corresponds to the difference between the total surplus from trade minus the seller's uniform monopoly profit (the one achievable by the seller posting a single price knowing only the prior). Instead, if the seller picked what he learned about the buyer's preferences, naturally the seller would want to learn everything to extract the full surplus from trade. An important implication of their results is that ex ante the buyer is willing to pay a substantial amount to control the seller's information. In contrast, our results imply that interim the buyer's willingness to pay for informational control drops to null.

In this strand, we also relate to Roesler and Szentes (2017) who also study the classic static monopoly problem with private values. They determine the optimal information structure that an uninformed buyer would ex ante choose to learn about her *own* valuation. The seller observes this choice of information structure, but not the signal generated from this observation structure which become the buyer's private information. Our analysis includes the static monopoly problem but we differ in that we focus on the situation in which the buyer is already informed and chooses the seller's information.

Second, we relate to the literature on voluntary disclosure of verifiable information. This literature considers particular evidence structures that the privately informed party can disclose to an uninformed party, e.g., Grossman and Hart (1980). A key insight of this literature is a full disclosure result. Specifically, Grossman (1981) and Milgrom (1981) show that in the environments they study there is a unique perfect equilibrium, in which all types choose full disclosure. Verrechia (1983) shows that when disclosure is costly, the equilibrium is given by a cutoff where low types do not disclose while high types do, and as the cost of disclosure decreases, the set of types that choose disclosure increases and the equilibrium converges to full disclosure as this cost goes to zero. Given such a fixed cost of disclosure, Jovanovich (1982) shows that there may be too much disclosure from a social efficiency perspective.[2] In Section 3.5 we describe how our setup effectively captures the evidence structure of voluntary disclosure, and provide a detailed comparison of our results and the logic of these papers.

Third, our approach is related to the literature on signaling. Our results highlight the importance of the common assumption built into signaling models: that by signalling, agents receive a payoff change com-

[2]For a review see, e.g., Milgrom (2008). For more recent work in this tradition cf. e.g. Ali, Lewis, and Vasserman (2023) who consider costless verifiable disclosure of a privately informed buyer to a monopolist seller in a static setting using the evidence structure of Grossman (1981); they show that the buyer may choose and benefit from partial disclosure of her information. See also Strausz (2017), who establishes the revelation principle for mechanism design with costless verifiable evidence.

mensurate with their types (cf. Kreps and Sobel (1994) for a list of possible assumptions, all of which share this feature).[3] For instance, in Spence (1973), the signaling agents' payoffs correspond to their types because the market to whom they signal is competitive. In contrast, we focus on a buyer who faces a monopolistic seller and who might obtain no gain in the rent she receives when conveying her type.

Fourth, we contribute to the literature on privacy. Acquisti, Taylor, and Wagman (2016) provide an example where the seller can elicit full disclosure by a buyer for a small discount linking the argument to Grossman and Milgrom as discussed above.[4] Pram (2021) considers an adverse selection problem where both the seller and the buyer's payoff depends directly on the buyer's private information and shows that both parties may benefit from disclosure. Calzolari and Pavan (2006) and Board and Liu (2018) study how firms may disclose consumer information to each other and track it dynamically in the context of price discrimination. A growing number of papers explore the ways in which consumers may benefit from their private information being transmitted to improve the match between sellers and buyers, e.g., Kim and Kircher (2015), Mirkin and Pycia (2015), de Corniere and de Nijs (2016), Hidir and Vellodi (2021). We discuss other work on privacy in Section 4.

We also relate to the broader economic literature on the role of defaults in determining economic outcomes. In this literature, defaults matter because of consumer inertia and procrastination as, e.g., in Madrian and Shea (2001). Defaults might also anchor strategic reasoning (cf., e.g., Crawford and Iriberri 2007). In contrast, in our environment the presence of a default fully determines the equilibrium choices without any behavioral considerations.

[3] Cf. also the survey by Riley (2001) and the recent general signaling analysis of Starkov (2023), which also implicitly incorporate commensurate-payoff assumptions.

[4] In our earlier draft, Madarasz and Pycia (2020), now subsumed by the current paper, we relatedly showed that whenever there is a default privacy choice that is cheapest for all types, then this default choice is made in every equilibrium.

2 Setup

Consider the classic dynamic monopoly problem. The seller (receiver) owns an object whose value to the seller we now normalize to be 0. The buyer (sender) privately knows her value for this object which is denoted by $\theta \in [\underline{\theta}, \overline{\theta}] = \Theta \subset \mathbb{R}^+$ and is drawn from a commonly known cdf $F(\theta)$.

In each period $t \geq 1$ of the game, the seller makes a price offer. The buyer can accept or reject this offer. If she accepts some offer p_t in period t, her payoff in period t is $\theta - p_t$ while the seller's payoff in this period is p_t. Accepting an offer ends the game. If she rejects an offer, the game continues or ends without trade in which case payoffs are null. The players have von Neumann-Morgenstern preferences over these payoffs. Players discount their future payoffs at their own, potentially different, fixed discount rates, each of which is an element of $(0, 1)$. The horizon of the game T can be finite ($t \in \{1, ..., T\}$) or infinite ($t \in \{1, ...\}$).

We study the information design choice of the privately-informed party in this classic problem. Specifically, at period $t = 0$, the buyer chooses from a finite set of dynamic experiments $S = \{s_1, ...s_N\}$. Each experiment,

$$s_j : \Theta \to \Delta(Z_1 \times Z_2 \times ....),$$

is a Borel measurable function of the payoff state and, for each t, Z_t denotes the set of possible signal realizations in period t. An experiment may immediately reveal the payoff state, may reveal it gradually, may convey only partial or partitional information, or none at all. Signal realization can also be correlated over time. We also allow for these realizations to be either public or observed privately by the seller. Finally, we employ the standard notion of Perfect Bayesian Nash equilibrium (PBE or equilibrium henceforth) as our solution concept.

Remark 1. We can expand the set of experiments and allow for more complex information processes. Formally, let Ω be the space of the states of Nature with a common prior. Each state resolves all uncertainty about the environment; in particular, θ is also determined by the state of nature (and we maintain the standard assumption that the dependence is Borel). Let $\mathbb{P}$ be a probability measure on Ω and $F(\theta)$ be the resulting cdf on buyer's values. Each experiment,

$$s_j : \Omega \to \Delta(Z_1 \times Z_2 \times ....)$$

is a Borel function of the state of nature $\omega \in \Omega$ which assigns to each state a probability distribution over a set of, possibly correlated, signal realizations over time. This more general formulation also allows for auto-correlation between signals and for the players to receive information about the arrival of payoff information per se, etc. All our results will continue to hold under this general formulation.

3 Information Choice

We turn to the characterization of our results. We start with the more restrictive case where each available experiment has the same direct cost or benefit, then consider the case where experiments have heterogeneous direct costs or benefits. Finally, we consider a number of extensions of our setup that demonstrate the robustness and the force of our results.

3.1 Costless Experiments

First, suppose that there are *no cost* differences associated with choosing any element of S. For example, experiments may all be free, or if the choice of an experiment per se carries some cost $c \in \mathbb{R}$, or direct utility benefit $b \in \mathbb{R}$, incurred by the sender, then we assume that this direct cost or benefit is the same across *all* experiments in S. Our first result claims that under these conditions the buyer pooling on any given experiment is a PBE.

Proposition 1 *For any given $s^* \in S$, the buyer always choosing s^* is a Perfect Bayesian Equilibrium.*

Our first result claims that when all experiments are equally costly, then *anything goes* in the sense that the buyer choosing any experiment s is an equilibrium. Suppose that one experiment reveals the buyer's type perfectly and immediately while the other one only reveals some partitional information, given any partitioning of the buyer's types. Then the choice of either is an equilibrium.[5]

To emphasize an implication of this result, consider the classic information design setting where the buyer picks an experiment *ex ante*, the seller's information, before privately learning her preferences, the so-called commitment assumption. The above implies that even if the commitment assumption is relaxed and the buyer picks an experiment at the *interim* stage, as long as these experiments are equally costly, the choice of each experiment can be supported in equilibrium. In other words, whatever choice the buyer would make from S ex ante, this choice remains an equilibrium even if the buyer has to choose ex interim, i.e., once she already knows her preferences. In turn, the buyer who controls the seller's information, can achieve the same payoff as in the classic information design problem under the ex ante choice assumption.[6]

Corollary 1 *The buyer can realize the same ex ante expected payoff as the one she could realize if she could commit to a choice from S ex ante.*

Suppose one of the experiments $s_p \in S$ is a no-revelation experiment: the distribution of signal realizations is independent of the payoff state θ. If the buyer chooses this experiment, the seller cannot infer anything about the payoff state from observing the signals per se. The above result then implies that full privacy protection is always an equilibrium, that is, there is an equilibrium where all buyer types choose s_p. If s_p is available, the buyer can ensure at least the consumer surplus that she would in the classic monopoly setting with exogenous private information. She can do potentially much better depending on the elements of the set S, as in Bergemann, Brooks, and Morris (2015).

[5] The very special case of the present model in which all experiments are free and none of them generates any information corresponds to cheap talk. In the cheap talk context, Proposition 1 simplifies to asserting that any message can constitute the support of an uninformative bubbling equilibrium, cf. Crawford and Sobel (1982).

[6] In the case the buyer commits to a choice ex ante without having any private information, and $T = 1$, this is analyzed by Bergemann, Brooks and Morris (2015).

3.2 Costly Experiments

We now turn to the case where different experiments in S have different costs. Let $c(s) \in \mathbb{R}$ be the cost of *choosing* experiment s incurred by the sender at $t = 0$. We do not require this cost to be positive. A negative cost can correspond to some direct benefit of choosing a given experiment (e.g., the entertainment or convenience value of a search platform on which the buyer purchases from the seller). The cost of an experiment may also be related to its informativeness or entropy reduction.

Suppose that there exists a unique cheapest experiment. In various settings, we can also interpret this as the default choice. A motivation for this terminology is that defaults are typically less costly to choose than other options, but we do not require this choice to be free or of no direct utility benefit. Below we call s with the lowest $c(s)$ as the cheapest experiment.

Theorem 1 ***[Private Values]*** *In equilibrium the buyer always chooses the cheapest experiment from S irrespective of its information content.*

In contrast with the previous observation, the above main result shows that when different experiments are differentially costly, then there is a unique equilibrium. Crucially, the information content of the chosen experiment and the content of the unchosen experiments do not matter. This informational irrelevance result means that when the sender makes the information design choice at the interim stage, then the informational characteristics of these designs are irrelevant. Instead, all that matters is the cost associated with each design.[7]

To provide intuition, note first that a pooling equilibrium on the lowest cost experiment is a perfect equilibrium just as before. The proof is then based on the fact that, since the seller's prices respond rationally to information in each period, the equilibrium surplus from trading obtained by the extremal buyer type choosing any given experiment is

[7]For example, in the context of costly persuasion of Gentzkow and Kamenica (2014), when the cost of an experiment is proportional to its entropy reduction, then the unique equilibrium in our setting is full privacy protection.

weakly bounded from above by the surplus from trading this type obtains by deviating to any other experiment. In turn, a separating equilibrium cannot exist. If it did, then from the set of types choosing the more costly experiment with positive probability, there would always be some type for whom the difference between the trading surplus attained when sticking to the more costly experiment, versus the trading surplus attained when deviating to a cheaper experiment, must be smaller than the difference between the direct costs of these two experiments. In other words, there would always be types who could save more by deviating to a cheaper experiment than how much they could possibly lose in terms of the surplus they obtained from trade. This logic on the impossibility of separation also implies that pooling on an experiment with non-lowest cost can not be a perfect equilibrium.

3.3 Discussion of Assumptions

Our Proposition 1 and Theorem 1 remain valid in many extensions of our model. We discuss possible extensions in a series of remarks.

Remark 2. Type-dependent Costs. Above, the cost (direct benefit) of an experiment was independent of the buyer's type. The above argument, however, does not depend on this. The cost of each experiment can be type dependent $c(s, \theta) \in \mathbb{R}$. The only assumption which we need on $c(s, \theta)$ is that there is a commonly cheapest one for all types, that is, there exist s^* such that $c(s^*, \theta) < c(s', \theta)$ for any θ and $s' \in S$. This way, our setup can encompass aspects of costly signalling where sending a certain message is more costly to some types than to others (even if the experiment itself reveals no direct information, as in some classic examples). Such costly signalling fits into our setting as long as there is a common lowest cost (default) choice.[8]

Remark 3. Above, there was also a unique least costly experiment in S. In case there are multiple least costly ones, the buyer pooling on any one of them is an equilibrium. Furthermore, just as above, the sender

[8] For a further discussion of this point see the earlier version of our paper, Madarasz and Pycia (2020).

choosing any more costly experiment with positive probability is not an equilibrium.

Remark 4: Observation Structure. We also assumed that signal realisations were either public or the receiver's private information. This ruled out more general forms of *second*- and *higher-order* uncertainty that the players can have about each other's beliefs or updates, e.g., on the timing of the arrival of information under the more general experiments of the form described in Remark 1 of Section 2. Our results are, however, robust to considering general observational structures involving complex uncertainties that the players' can have about each other's beliefs.

Specifically, suppose that each experiment s_j generates two signal realizations: x_t^s and x_t^b, in each period t, one for the seller and one for the buyer respectively. When $x_t^s = x_t^b$, the signal realizations in period t are public. When x_t^s and x_t^b are independent, signals are purely private. Given the more general definition of an experiment discussed in Remark 1, we can allow for $x_t^s \neq x_t^b$ and these two to be imperfectly correlated as determined by the, by the players unobserved, state of Nature. We can then allow much more general observation structures including, e.g., the buyer to learn privately over time about the dynamic unfolding of the seller's information (including about what he learns about such learning).

Remark 5: Sequential Information Choice. In our setup the buyer made a design choice only at $t = 0$ about the seller's (dynamic) information. In effect, she commits to an experiment. This assumption can also be relaxed. Suppose that at the beginning of each round t, the buyer can make a choice of experiment $s_{j,t}$ from some set of experiments S_t. Each experiment $s_{j,t} \in S_t$ determines a type-dependent distribution of signals

$$s_{j,t} : \Theta \to \Delta(Z_t^t \times Z_{t+1}^t \times ....),$$

where Z_k^t denotes the set of possible signal realizations of time-t experiments in period $k \geq t$ and the mapping is measurable. Each experiment

$s_{j,t} \in S_t$ is associated with some flow cost $c(s_{j,t}) \in \mathbb{R}$ incurred at time t. To ensure the optimization problems are bounded, we assume that there is a cap $\overline{c}$ such that $|c(s_{j,t})| \leq \overline{c}$ for each experiment $s_{j,t}$ in each period t, but we again impose no sign restrictions on these values. Finally, to simplify the statement, we also assume that the set of period-t experiments and their costs do not depend on prior choices and signal realizations and that in each period there is a unique cheapest experiment s_t. As the sender makes information choices sequentially, we allow the choice in period t to be conditioned on the entire realized history of the game till the beginning of period t.

In this sequential choice environment, our result remains true and the sender still chooses the least costly experiment in each period.

Proposition 2 *[Sequential Experiment Choice] In any equilibrium, in each period t, the buyer always chooses the cheapest experiment $s_{j,t} \in S_t$.*

The proof of this result, which is described in the Appendix, hinges on the dynamic consistency of intertemporal choice. This is ensured in our setting given exponential discounting. Since the cost of an experiment is allowed to be negative, e.g., the entertainment value of a platform, the setting also allows for the buyer to want to postpone trade simply to collect such positive values from engagement. Nevertheless such a motive for equilibrium delay does not change our insight that equilibrium information choice is independent of the information revelation such choice generates. Note also that the above implies that small inter-temporal redistribution of costs across experiments can have major impact on the outcome of trade.

Remark 6: Bargaining Power. We followed the classic assumption that the seller (receiver) has the bargaining power. This assumption can also be relaxed. Our conclusion directly extends to the case where it is the buyer (sender) instead of the seller who has the bargaining power. If the buyer has the bargaining power, the conclusion holds almost immediately. Since then, her surplus from trading is independent of what experiment she chooses, she can force the seller to sell at his reservation value.

In fact our result extends to the case where up until some round $\hat{t} \geq 1$ it is the seller who has the bargaining power, but following some (possibly stochastic) $\hat{t}$, the bargaining power switches, and it is the buyer who can make offers thereon. Under such general stochastic bargaining procedures, all positive types of the buyer obtain a strictly positive rent in equilibrium. Nevertheless, the buyer always chooses a cheapest experiment irrespective of her type or the experiment's content. We summarize the above in the following proposition whose proof is again in the Appendix.

Proposition 3 *[Bargaining Power Switch] Let $\hat{t}$ be a random variable with support contained in $\{$ 1,2,3,...$\}$ according to a commonly known distribution and suppose that the bargaining power switches from the seller to the buyer at period $\hat{t}$. Then, (i) for any cheapest experiment s there is an equilibrium in which all types of the buyer choose s and (ii) in any equilibrium, all types of the buyer choose a cheapest experiment.*

In the proposition above, we allow different experiments to have the same costs, thus extending both Proposition 1 and Theorem 1. We could relax the assumption that the bargaining power switches at time $\hat{t}$ with support contained in $\{$ 1,2,3,...$\}$ by allowing the bargaining power to always stay with the seller with positive probability. Similarly to our sequential choice proposition, these results also hinge on the dynamic consistency of intertemporal choice.

Remark 7: Two-sided Payoff Uncertainty. Finally, we also assumed that the seller's valuation was common knowledge. Suppose instead that the seller privately knows her valuation of the object, or the cost of producing it, e.g., as in the classic setup of Myerson and Satterthwaite (1983).[9] Such uncertainty about the seller's valuation faced by the buyer does not change our result. Since the seller will never offer the object for less than how much he values it, it is still true that a buyer type being an extremal type amongst those that choose a particular experiment will not receive a greater equilibrium surplus from trade

[9]Unlike Myerson and Satterthwaite (1983), we do not need to assume that valuations are independently distributed.

than if she deviated to a cheaper experiment and the proof of Theorem 1 applies.

3.4 Common Values

In the base model, we considered the classic private value environment where the buyer's private information did not affect the seller's valuation of the object per se. As we now show, our insights extend to general common value environments. We maintain the same assumptions on the timing and costs of experiments as in the base model (SSecton 3.2), but we relax the assumptions on the traders' valuations. We now assume that the buyer's (or sender's) valuation is given by some non-negative and bounded $b(\theta) \in \mathbb{R}^+$ and the seller's (or receiver's) valuation is given by some $v(\theta) \in \mathbb{R}$, which might represent, e.g., the cost of producing, servicing, or delivering the good.

These assumptions are satisfied in many economic environments. The valuations and delivery costs might be increasing in some common factor so that serving a buyer who values the object more may be more costly, as in the case of insurance for example. Valuations may also reflect a common quality or earning potential of the object that both parties value, as in the case of asset trade. The valuations might also be negatively correlated or have no monotonic relationship at all. The parties might initially have common knowledge that trade is mutually beneficial or they might not. The setting thus nests many environments with classic adverse selection.

Theorem 2 ***[Common Values]*** *In equilibrium the buyer always chooses the cheapest experiment from S irrespective of its information content.*

In the presence of common values, the value of trade for the seller at a given price often directly depends on the buyer's type; the seller faces uncertainty about his payoff even conditional on the buyer agreeing on trade. In turn, depending on what he knows about the buyer, the seller may prefer not to trade at some (or any) positive price. Nevertheless, the logic of Theorem 1 continues to apply and the buyer, in equilibrium, never wants to choose any information structure, other than the cheapest one.

As an example, consider a selection problem where $v(\theta)$ is the cost of serving a given type, and this cost is increasing in θ. Suppose that in the absence of disclosure, there is a break-down of trade. Trade could be possible, however, with a finer partitioning. Our result implies that when supplying information occurs endogenously and such a finer partitioning of the payoff state is not the cheapest option, trade still does not occur in equilibrium.

Example 3 ***[Efficiency or Market Failure.]*** *Suppose that $b(\theta) = \theta$ is distributed on $[0, 1]$ and $v(\theta)$ is non-negative and sufficiently steeply increasing in θ. Let there be one period in which information can be disclosed and in which the seller makes a price offer, $T = 1$. If non-disclosure is the cheapest experiment, there is no trade in equilibrium. Any price offer $p \leq 1$ will attract the types the seller would not like to trade with. If full disclosure is the cheapest experiment, trade is efficient. There are many intermediate cases. For instance, suppose that the cheapest experiment discloses whether or not $\theta \leq a$, for some $a < 1$. If $b(a) > v(a)$, then adverse selection can be alleviated and there could be trade between the seller and buyer types below a.*

This example demonstrates that a minor change in the cost of supplying information can substantially change the overall efficiency of trade. The loss of trade surplus can dwarf compared to the cost difference between the relevant experiments, or the transaction cost associated with supplying information.

Finally, note that Theorem 2 and our comments on efficiency remain valid if we relax many of our baseline assumptions just as before. Specifically, Remarks 2-5 continue to hold as stated. An analogue of Remark 6 also holds again and the results do not hinge on the seller having the full bargaining power, e.g. if state θ is revealed by or at the time of the switch of bargaining power.

One further implication of Theorem 2 is that the buyer is unable to signal her type by choosing an experiment. For example, even if each experiment was uninformative per se, it delivered only noise, one might suggest that the fact that different experiments have different costs, allows the buyer to engage in meaningful costly signalling in a common value setting. Instead, our result implies the non-existence of separating or partially separating equilibria.

3.5 Verifiable Disclosure

In our setup, we allow for arbitrary experiments. The results can then be applied to the classic problem of verifiable disclosure. Simultaneously, our findings are in contrast with some seminal insights of this literature, such as those presented by Grossman and Hart (1980), Grossman (1981), or Milgrom (1981). To establish a connection between their setup and ours, we first discuss the difference in evidence structure between our setting and theirs. We then discuss the difference in the payoff structures.

In terms of the evidence structure, in our setting, the sender can choose any experiment from S irrespective of her type. In contrast, in Grossman and Milgrom, the set of messages from which the sender makes a choice, corresponds to the set of all subsets of Θ. At the same time, a given message $M \subset \Theta$ can only be chosen by types θ that are in this set M, i.e., $\theta \in M$. Types outside of this set M *can not* choose message M. This assumption is often underpinned by the premise that verification is cost-free; hence if a type tried to falsify information, such a deception would be immediately detected and fail verification, e.g., Milgrom and Roberts (1986).

However, we can directly map this evidence structure into our setup in terms of the induced posteriors of the receiver. Consider a discrete set Θ and message $M \subseteq \Theta$. Suppose that any type can choose an experiment s_M defined as follows:

- if $\theta \in M$, this experiment immediately reveals that $\theta \in M$,
- if $\theta \notin M$, this experiment immediately reveals that the sender's type resides in the complement of M.[10]

[10]Signal realizations in Z can be represented as M and $\Theta \setminus M$.

Intuitively, types can verify that they are in M, and any lie is detected immediately. Experiment s_M, as defined above, induces a receiver posterior identical to that obtained by selecting message M with the no-lying restriction in place a la Grossman and Milgrom. At the same time, this experiment satisfies all aspects of our setup. We can analogously define experiment s'_M for any $M' \subseteq \Theta$, and let S be the collection of all such experiments. Consequently, given the above mapping, we can describe their evidence structure as a special case of ours.[11] In turn, all our results apply.

Allowing each type to pick any experiment along the lines of the above mapping has some implications. For example, in the environment of Ali, Lewis, and Vasserman (2023), who consider the classic static monopoly problem with disclosure a la Grossman, if one allowed each type of the buyer to pick any experiment in this fashion, then, by Proposition 1, not only full disclosure and partitional information revelation, but pooling on *any* available experiment would become a perfect equilibrium. Furthermore, if the set of experiments the buyer can choose from also contained the experiment which generates the buyer-optimal ex ante information structure, as identified by Bergemann, Brooks, and Morris (2015), then there would be an ex-post Pareto optimal perfect equilibrium, in which all buyer types choose this experiment. At the same time, if there are differences in costs, then our Theorem 1 applies.

There is a substantive difference in terms of the strategic assumptions between our setup and that considered by Grossman, Hart, and Milgrom.[12] As noted, in their setting, when it comes to choosing an experiment, there is always a sender type who has a strictly dominant

[11] Alternatively, we can also approximate this classic evidence structure with a set of experiments where the realization of each experiment generates a full support posterior over Θ. Specifically, consider an experiment s_M which leads to signal realization M with probability $1-\epsilon$ and to some other signal realization M^c with probability ϵ for all types in M. For all types outside of M, the possible signal realizations are the same, but with the reverse assignment of the probabilities. In terms of the induced posteriors of the receiver, the difference between this experiment and the message M a la Grossman (1981) can be arbitrarily close, e.g., when measured in terms of the Kullback-Leibler divergence, vanishing as ϵ vanishes.

[12] Their analyses were extended to costly disclosure by Jovanovic (1982) and Verrechia (1983); see our discussion of Related Literature above.

strategy and this strictly dominant strategy is to choose full disclosure. Any type of the sender who is maximal in the support of the receiver's belief always achieves a strictly greater payoff, irrespective of the choices of the other types, if she fully discloses herself. This drives their classic unravelling logic, including in the case of costly disclosure. Instead, in our setting no type, maximal or not, needs to have a strictly dominant choice from the set of experiments. No choice of experiment guarantees a strictly higher payoff than another irrespective of how other types choose. In turn, there is no unravelling in our setting and the strategic forces at play are clearly distinct.

In our context of trade, where prices emerge endogenously in response to the distribution of information, we can further classify situations along two dimensions. First, depending on whether or not the receiver's payoff depends *directly* on the sender's private information, as in the case of private versus common values. Second, whether it is the sender or the receiver who has the power to make price offers. Four possible scenarios arise in this classification, and our logic applies to three of these, while the logic of Grossman, Hart, and Milgrom's applies to the remaining one (under additional assumptions they impose on payoffs). The table below summarizes these observations.

	R doesn't directly care about θ	R does directly care about θ
R sets the price	current paper	current paper
S sets the price	current paper	G-H-M

In the boxes marked above, our results imply that the sender's choice of the receiver's information depends purely on the cost of supplying information and not on the content of the information supplied. This is true both when the sender's payoff directly depends on her type and also when it does not.

4 Privacy Paradox and Platform Design

Although our irrelevance result says nothing about the kind of information design the sender will end up choosing, in many contexts there is a natural link between the information content of an experiment and its cost. Below, we first consider the link between our results and an empirical phenomenon termed the privacy paradox, then embedd our result in a simple market setting where cost differences emerge endogenously.

4.1 Privacy Paradox

The idea of a privacy paradox refers to the observation that, despite expressing concern about the loss of their privacy, consumers seem to take minimal action to protect it, such as concealing their behavioral patterns, paying extra for apps or products that do not track or verify them, or request access to various aspects of their data. Empirical research has documented such a general discrepancy between people's stated preferences versus their revealed preferences regarding privacy, e.g., Berendt, Gunther, and Spiekermann (2005), Athey, Catalini, and Tucker (2017). For example, Johnson, Shriver, and Du (2020) describe that "though consumers express strong privacy concerns in surveys, we find that only 0.23 percent of American ad impressions arise from users who opted out of online behavioral advertising." As Barth and de Jong (2017) summarize in their survey of the literature, "while many users show theoretical interest in their privacy and maintain a positive attitude towards privacy-protection behavior, this rarely translates into actual protective behavior."[13]

Firms appear to use online consumer data for price and search discrimination, e.g., Milkians et al. (2012, 2013), with the value of consumer advertising greatly depending on consumers' privacy choices, e.g., Goldfarb and Tucker (2012). People also appear to be generally concerned

[13] In various settings consumers tend to share private information in exchange for small retail value and personalized services, e.g., Berseford et al. (2012), while being very concerned about their privacy. For example, Rose (2005) finds that although most survey respondents reported that they were concerned about their privacy, only 47 percent of them expressed a willingness to pay any amount to ensure the privacy of their information.

about firms' attempts to collect, store, and interpret information about them. For example, Turow et al. (2009) report, based on a representative survey, that "contrary to what many marketers claim, most adult Americans (66%) do not want marketers to tailor advertisements to their interests."[14] They also report that 63% believe advertisers should be required by law to immediately delete information about their internet activity.

To explain this discrepancy, various authors have argued that behavioral factors, such as a taste for immediate gratification, framing, or the miscalibration of probabilities, may well be at play, e.g., Acquisti et al (2013).[15]

As Athey, Catalini, and Tucker (2017) point out that this discrepancy could be consistent with either of two scenarios. It might reflect that people's stated preferences echo their normative preferences, or it could indicate a situation where, because expressing a preference for privacy is essentially costless in surveys, consumers are eager to state such a preference. However, when faced with even minimal costs, this interest in privacy quickly evaporates. Whether one or the other is the case matters for policy not least because privacy protection has been shown to impose significant losses, e.g., Kim and Wagman (2015).

While behavioral factors are likely important, and survey responses might not be reliable reflections of actual preferences, our result has a robust implication. Whenever there is a uniformly cheaper option than full privacy protection, then, in the scenarios considered in Section 3, buyers in equilibrium shall never opt for privacy protection. This is true despite the fact that buyers truly believe that they stand to lose a great deal from the loss of their privacy and would, therefore, have made a different choice ex ante—that is, prior to possessing private information.

Recall our notation of s_p from Section 3. In the proposition below, we

[14]Moreover, when Americans are informed of three common ways that marketers gather data about people in order to tailor ads, even higher percentages - between 73% and 86% - say they would not want such advertising.

[15]Another psychological force is the illusion of transparency whereby people believe that their private information may leak to others irrespective of the actions they take, e.g., Madarasz (2021).

allow for experiment costs to depend on the buyer's type, as in Remark 2, and otherwise maintain the assumptions of our common value (and hence also private value) model. We also allow there to be no uniformly cheapest experiment, that is, no $s \in S$ that is cheapest for all θ's.

Proposition 4 *Suppose that $s_p \in S$ and there exists $s' \in S$ such that $c(s', \theta) < c(s_p, \theta)$ for all $\theta \in \Theta$. Full privacy protection is never chosen in equilibrium.*

This proposition extends our cost-over-content results to the environment in which privacy protection is costly but no assumptions are otherwise imposed on the costs: the costs can depend on types and the cheapest default option does not need to exist. Any given experiment, not just s_p, that is uniformly more costly than some other experiment in S, can never be chosen in equilibrium.

To illustrate the above, suppose that at the beginning of each period t the buyer's valuation may privately leak to the seller with probability $\alpha_t(s)$ that depends on the buyer's investment (s) in protecting her privacy. For example, the seller (or the operator of an online platform on behalf of the seller) may be able to figure out the buyer's valuations for the seller's product from observing the buyer's activity online. The buyer can pick an investment level from a finite set S which contains the null investment option, s_0 at cost 0, while all other levels of investments carry strictly positive costs. The leakage probability $\alpha_t(s)$ is decreasing in some ordering of the investment s. In this environment, the above proposition implies that the buyer never invests any amount in decreasing the leakage probability. For a more detailed discussion, we refer the reader to Madarasz and Pycia (2020).

4.2 Platform Design

So far we kept the emergence of the buyer's privacy options to be an exogenous aspect of the environment. We now consider the case where these options and their costs to consumers arise endogenously. In addition to the buyer and the seller, we introduce a third agent: a platform provider.[16] The uninformed platform provider has access to a finite set P of technologically feasible experiments (platforms). For simplicity, all experiments in this set are equally costly for the provider to supply, and we normalize this to zero, but this assumption can be fully relaxed. The parties share a common prior.

The platform provider selects some subset $S \subseteq P$ and assigns to each $s \in S$ some fee $c(s) \in \mathbb{R}$ that the buyer incurs when selecting s. In other words, the provider selects a menu $\{s, c(s)\}_{s \in S \subseteq P}$. We interpret each s as a privacy platform. The signal generating process s determines what the seller may learn about the buyer. In this application, this can be a function both of the buyer's behavior on the platform and a contract between the platform provider and the buyer specifying what type of information the provider might pass on to the seller, e.g., selective tracking. We interpret the fee $c(s)$ associated with a given privacy platform s, as the price the buyer needs to pay (or the subsidy she receives) when participating on platform s. Note that we impose no sign restriction on $c(s)$ which can be positive, negative, or null.

The provider's payoff equals to the sum of the fee $c(s)$ paid by the buyer picking experiment s from S, and the payment that the seller will make to the provider. Although the details of the bargaining between the provider and the seller do not matter, to make the model fully specified, we assume that the platform provider makes a take-it-or-leave-it offer to the seller (committing to a menu $\{s, c(s)\}_{s \in S \subseteq P}$). If the offer is accepted, the game continues and the seller gets access to the information that the buyer's choice of experiment contractually allows. If it is rejected, the game ends.

Finally, the buyer makes her choice from the menu $\{s, c(s)\}_{s \in S \subseteq P}$. We also allow the buyer to choose to shop 'offline' directly from the profit-maximizing seller, i.e., not use any of the privacy platforms and fully protect her privacy for free. If the buyer chooses the offline option, and never buys, her utility is normalized to zero. The reminder of the

[16] We study a platform provider who extracts information from the buyer to provide it to the seller. For a complementary analysis of a platform provider who provides buyers with information, see, e.g., Terstiege and Wasser (2021).

environment is identical to our base model with common values, which includes the base private value model as a special case.

In this environment the selection of experiments and their costs is endogenous and determined by the platform provider in equilibrium. What experiments (platforms) will the provider offer? It turns out that we can make a tight prediction. To simplify its formulation, we say that a platform is *relevant* in an equilibrium if it is chosen by the buyer with positive probability. Note that the presence of privacy platforms that are not relevant has no impact on equilibrium outcomes.

Proposition 5 *In equilibrium all relevant platforms are provided for free, $c(s) = 0$. The set of platforms provided maximizes the seller's revenue among all feasible sets of platforms, and with probability one the buyer participates in one of the relevant platforms provided (thus never shops offline).*

5 Conclusion

This paper considers the information choice of a privately informed sender for a receiver in the context of classic trade. Our approach aligns with the information design literature, allowing the sender to choose from a large set of arbitrary experiments. We follow the literature on disclosure and signalling in that the sender is privately informed when picking the receiver's information. At this intersection our paper presents a general informational irrelevance result whereby the sender's choice of the receiver's information is determined purely by the cost of supplying this information and not by its content.

Theoretically, there are a number of directions for extending our results and establishing its limits. One may explore how more complex bargaining protocols impact our results beyond the discussion we offered. One can also extend our setup to consider cases where the receiver faces uncertainty about whether or not the sender is informed, as in, e.g., Dye (1985), Acharya, DeMarzo, and Kremer (2011), Ben-Porath, Dekel, and Lipman (2018).

Another extension of our results could consider settings where the buyer's choice of experiment impacts what both the seller and buyer learn about the payoff state. It would then be interesting to extend our result to the presence of buyer learning in the context of trade as in Roesler and Szentes (2017), Ravid, Roesler, and Szentes (2022), or Thereze (2022). Specifically, one may consider cases between the ex ante perspective, as adopted in information design, and the interim case, as adopted in this paper; the sender does have private information when choosing the experiment, but this experiment may also reveal information about the payoff state to the sender as well.

As an important application, we consider market settings in which buyers' privacy options arise endogenously. Here, our analysis highlights the potential significance of defaults in privacy regulation. Our results imply that allowing firms and consumers to trade data may broadly put consumers at a disadvantage because the sellers can acquire private data cheaply.[17] Instead, as we show, the direct regulation of the data collected may be a more effective way of safeguarding information that consumers would prefer to keep private. This might help inform such regulations as GDPR in the EU or the similar regulatory attempts in the US. More generally, in the presence of non-trivial transaction costs or regulation of data trade, our result points to the power of defaults. For example, changing the default from opt-in to opt-out can have significant impact both on the allocation of surplus and the efficiency of trade even if the default is non-biding in that switching to a different privacy option by design is relatively easy. [18]

Our results are also in contrast with earlier accounts of behavior-based price discrimination, and other accounts based on the Coasian dynamics, e.g., Hart and Tirole (1988), Laffont and Tirole (1988), Taylor (2004) or Acquisti and Varian (2005). Under classic price discrimination, given Coasian informational dynamics, the monopolist loses and consumers gain when the seller tracks prior purchasing decisions. Along

[17] Acquisti, Taylor, and Wagman (2016) recognized that the possibility of such cheap data acquisition in an example and our analysis implies that the problem is ubiquitous in trade.

[18] Note also that the logic we identify for why defaults matter is different from the procrastination or endowment effect they are usually associated with, e.g., Madrian and Shea (2001).

these lines, Taylor (2004) finds that, in the presence of tracking technologies that allow sellers to infer consumers' preferences and sell such information to others who can then engage in price discrimination, the usefulness of privacy regulatory protection depends on consumers' level of sophistication. In particular, regulation is not necessary if consumers are aware of how a company may use and sell their data and buyers can adapt their purchasing decisions accordingly, because it is in a company's best interest to protect customers' data even if there is no specific regulation that forces it to do so. Our results instead imply that when consumers can express their privacy preferences, the supply of information is itself endogenous, firms will greatly benefit from tracking and even if privacy protection remains a cheap option for consumers, e.g., ensured by a regulator, this may also come at a significant reduction in consumer surplus.

6 Appendix

Proof of Proposition 1. The following is a perfect Bayesian equilibrium. All buyer types chooses s^*. Following any deviation to some $s' \neq s^*$, the seller attributes this deviation to $\overline{\theta}$; note that this deviation has probability 0 and hence one is free to assign the seller's equilibrium beliefs. At each period t, the seller charges the price that is optimal given the seller's posterior belief at t. In particular, following any choice of $s' \neq s^*$, the seller charges $p_t = \overline{\theta}$ at any t in the continuation game. ■

Remark 8: Equilibrium Beliefs in the Proof of Proposition 1. While the above simple construction relies on extreme beliefs of the seller following a deviation by the buyer, our insights, including the proposition just proven, does not hinge on these extreme beliefs. To see this, suppose that the set of types is discrete.[19] For any s^* there is then a perfect equilibrium in which all buyer types select experiment s^*. Indeed, the construction of this equilibrium is as above except that

[19]The same argument applies to general Θ provided for any possible experiment s, at any time t, there is a well defined maximum type that has positive probability given the seller's equilibrium belief and the outcome of the experiment till time t.

seeing a deviation to experiment s', the seller believes, at time 0, that the deviating type is type θ_0 which is the highest type that has positive probability on the basis of outcomes of experiment s' till time 0. At each subsequent time t, the seller keeps his belief $\theta_t = \theta_{t-1}$ except if θ_{t-1} is no longer consistent with the outcomes of the experiment till time t; in the latter case the seller updates his belief and believes that the deviating type was type θ_t which is the highest type that has positive probability on the basis of outcomes of experiment s' till time t. At each time, the seller charges price p_t equal to the value of type θ_t. Note that no buyer type θ can gain strictly by deviating to any experiment s' because after such deviation the deviating type θ is in the set of types that are consistent with the outcomes of the experiment s' till any time t.

Proof of Theorem 1. First, note that all types pooling on the lowest cost experiment s is on path of an equilibrium. This on path behavior is supported in equilibrium when, following the 0 probability event that the buyer chose some experiment s' that costs $w > 0$ more than the cheapest experiment, at each time t the seller's belief puts probability 1 on a buyer's type θ_t that is higher than the supremum of types consistent with the outcome of s' till t minus $\frac{w}{2}$.

To ensure that the seller updates his beliefs via Bayes rule whenever possible, we assume that $\theta_{t+1} = \theta_t$, as long as θ_t remains consistent with the outcome of the experiment s' till time $t+1$. Given such beliefs, the seller's best response is to charge, at any time t, the price $p_t = \theta_t$.[20] Because $\theta \leq \theta_t + \frac{w}{2}$, the gain from the deviation is bounded from above by $\frac{w}{2}$ and is strictly lower than the additional cost w of choosing experiment s' rather than the cheapest experiment. Hence, the deviation is not profitable.

It remains to show that buyers with all types choose the least expensive experiment on the path of every equilibrium. By way of contradiction, suppose that there is an equilibrium in which some buyer types with positive probability choose some experiment s' with non-minimal

[20] In the second step of the proof below, we derive a more general bound on prices. This bound implies that $p_t \geq \theta_t$ and seller's impatience allows us to assume that indeed $p_t = \theta_t$ is a best response.

cost. As above, let w be the cost difference between s' and the least expensive experiment s. Let $\underline{\theta}_{s'}$ be the infimum of buyer types that choose s' with positive probability. In equilibrium, the seller then knows that $\underline{\theta}_{s'}$ is a lower bound on all buyer types the seller might observe choosing s'. Notice that this lower bound remains a lower bound on types in any time period irrespective of the outcomes of the experiment.

The price p_t charged by the seller in any period t is then bounded from below by $\underline{\theta}_{s'}$. Indeed, suppose not. Define $\underline{p}$ to be the infimum of p_t following s' over all signal realizations and subsequent time periods. Then $\underline{\theta}_{s'} - \underline{p} > 0$; let denote this difference by ε and let δ be the buyer's discount rate. Note that the non-negativity of prices implies that $\underline{\theta}_{s'} > 0$. By definition, there is some period t following experiment s' and some history of signals such that $p_t < \underline{p} + \frac{1-\delta}{3}\varepsilon$. All on-path types θ of the buyer buy at this history at p_t as it gives utility of at least $\theta - p_t \geq \theta - \underline{p} - \frac{1-\delta}{3}\varepsilon > 0$ while the utility from postponed purchase is at most $\delta(\theta - \underline{p})$, and hence the potential utility gain from waiting is bounded from above by:

$$(\delta - 1)(\theta - \underline{p}) + \frac{1-\delta}{3}\varepsilon \leq (\delta - 1)(\varepsilon) + \frac{1-\delta}{3}\varepsilon = \frac{2}{3}(\delta - 1)\varepsilon,$$

which is negative. Furthermore, the seller then has a profitable deviation at the history considered. If the seller raises the price at this history by $\frac{1-\delta}{3}\varepsilon$, then, in the conjectured PBE, the buyer would still buy immediately as an analogue of the above calculation shows that the potential utility gain from waiting is bounded from above by $-\frac{1-\delta}{3}\varepsilon$, which is again negative. This contradiction shows that $p_t \geq \underline{\theta}_{s'}$ at every history following the choice of experiment s'.

Type θ's surplus from buying the good following the choice of s' is hence bounded above by $\theta - \underline{\theta}_{s'}$ and hence types close to $\underline{\theta}_{s'}$ would strictly benefit from a deviation to the cheapest experiment s; a contradiction showing that in all PBEs all types choose the cheapest experiment s. ■

Proof of Proposition 2. The proof follows similar steps to the proof of Theorem 1. In the proof below, we denote by c_t the cost the cheapest experiment that can be chosen at time t. Let δ again denote the buyer's discount factor.

There is an equilibrium in which on path all types pool on the lowest cost experiment s_t at each period t. This on path behavior is supported by the belief that seeing the 0 probability deviation at some time t to some experiment s'_t that costs $w > 0$ more than s_t, the seller attributes this deviation, at each $t' \geq t$, to a type with value above the supremum of types consistent with the outcome of the experiments till t' minus $\frac{w}{2}$ (this type is not necessarily the supremum of types). Let us denote this type by θ'_t, and assume that $\theta_{t''+1} = \theta_{t''}$, as long as $\theta_{t''}$ remains consistent with the outcome of the experiment till time $t''+1$. Given such beliefs, the seller's best response is to charge at any time t' at least the price $p_{t'} = \max\{0, \theta'_t + \sum_{t''=t'+1,t'+2,\ldots} \delta^{t''-t'} c_{t''}\}$.[21] Thus, the gain from the deviation is bounded from above by $\frac{w}{2}$ and is strictly lower than the increase in the cost of the experiment w. Hence, the deviation is not profitable.

It remains to show that all buyer types choose the least expensive experiment in each t on the path of every equilibrium. By way of contradiction, suppose that there is an equilibrium in which some buyer types with positive probability choose some experiment with non-minimal cost. Let t' be the earliest time at which this happens, and let $s'_{t'}$ be a non-minimal cost experiment chosen with positive probability at that point. Similarly to the above, let w again be the difference between the cost of $s'_{t'}$ and $c_{t'}$. Let $\underline{\theta}_{s'_{t'}}$ denote the infimum of buyer types that choose $s'_{t'}$ at time t' with positive probability. In equilibrium, the seller then knows that $\underline{\theta}_{s'_{t'}}$ is a lower bound on all buyer types the seller might observe choosing $s'_{t'}$. Notice that this lower bound remains a lower bound on the types the seller is facing in any subsequent time period, irrespective of the dynamic outcomes of the experiment and that of future choices.

The price p_t charged by the seller in any period t is then bounded from below by $\underline{\theta}_{s'} + \sum_{t'''=t+1,t+2,\ldots} \delta^{t'''-t} c_{t'''}$. Indeed, suppose this bound fails at some period t''. Define $\underline{p}$ to be the infimum of $p_t - \sum_{t'''=t+1,t+2,\ldots} \delta^{t'''-t} c_{t'''}$

[21] A more generally applicable bound is established in the second step of the proof. In the case of a finite horizon game, the present bound follows immediately via backward induction showing that this price sequence is the unique (given the beliefs described above) best response of the seller.

where the infimum is taken over all periods $t \geq t''$ and all signal realizations in the continuation equilibrium starting at the history of realizations at which the above bound fails. The failure of the bound implies that $\underline{\theta}_{s'_{t'}} - \underline{p} > 0$; let us denote this positive difference by ε.

By definition of $\underline{p}$, there is some period $t''' \geq t''$ following $s'_{t'}$ and some history of signals, such that $p_{t'''} - (\sum_{t=t'''+1,t'''+2,\ldots} \delta^{t-t'''} c_t) < \underline{p} + \frac{1-\delta}{3}\varepsilon$. All on-path types θ of the buyer would buy at this history at $p_{t'''}$ as it would give utility (evaluated from the perspective of time t''') of at least $\theta - p_{t'''} > \theta - (\sum_{t=t'''+1,t'''+2,\ldots} \delta^{t-t'''} c_t) - \underline{p} - \frac{1-\delta}{3}\varepsilon$. Because any $p_{t''''}$ is bounded from below by $\underline{p} + \sum_{t=t''''+1,t''''+2,\ldots} \delta^{t-t''''} c_t$, the utility from postponing purchase and buying at some time $t'''' > t'''$ is at most $\delta^{t''''-t'''}(\theta - (\underline{p} + \sum_{t=t''''+1,t''''+2,\ldots} \delta^{t-t''''} c_t)) - \sum_{t=t'''+1,\ldots,t''''} \delta^{t-t'''} c_t$. Note that by taking $t'''' = \infty$, we obtain the case of not buying at all. Subtracting these bounds (and taking into account that $t'''' - t''' \geq 1$ and $\theta - \underline{p} \geq \varepsilon > 0$), we infer that the potential utility gain from waiting is bounded from above by:

$$(\delta^{t''''-t'''} - 1)(\theta - \underline{p}) + \frac{1-\delta}{3}\varepsilon \leq (\delta - 1)\varepsilon + \frac{1-\delta}{3}\varepsilon = \frac{2}{3}(\delta - 1)\varepsilon,$$

which is negative. Furthermore, the seller then has a profitable deviation at the time t''' history considered: if the seller raises the price at this history by $\frac{1-\delta}{3}\varepsilon$, then in the conjectured PBE the buyer would still buy immediately as an analogue of the above calculation shows that the potential utility gain from waiting is bounded from above by $\frac{1}{3}(\delta - 1)\varepsilon$ and is negative. This contradiction shows that $p_t \geq \underline{\theta}_{s'_{t'}} + \sum_{\tau=t+1,t+2,\ldots} \delta^{\tau-t} c_\tau$ at every history following the choice of experiment $s'_{t'}$. Type θ's surplus from buying the good following the choice of $s'_{t'}$ is hence bounded above by $\theta - \underline{\theta}_{s'_{t'}}$ and types close to $\underline{\theta}_{s'_{t'}}$ would strictly benefit from a deviation. This contradiction shows that in all PBEs, all types choose the cheapest experiments s_t at each time period t. ■

Proof of Proposition 3 . The choice of a cheapest experiment on path is still supportable in equilibrium and with the same seller's beliefs following a deviation as described above.

To show, by way of contradiction, that in all equilibria a cheapest experiment is chosen, suppose that there is an equilibrium in which a non-cheapest experiment s' is chosen by some types with strictly positive probability. Let $\underline{\theta}_{s'}$ be the infimum of types picking the experiment s' and let $w > 0$ be the cost difference between s' and the cheapest experiment.

At any time t at which the seller has the bargaining power, let $r_t(\theta) = \mathbb{E}_{\hat{t}>t}\delta^{\hat{t}-t}\theta$ be the expected rent of type θ (from the perspective of time t) from simply waiting till the possibly random time $\hat{t} > t$ at which the bargaining power switches. Note that following the bargaining power switch, standard arguments imply that the buyer will be able to extract all the rent from the seller. Let $q_t(\theta) = \theta - r_t(\theta)$ be the price that makes the buyer of type θ indifferent between buying at this price $q_t(\theta)$ and achieving surplus $r_t(\theta)$ by simply waiting to extract the rent when the bargaining power switches. In particular, the buyer never accepts any price greater than $q_t(\theta)$. Note that both $r_t(\theta)$ and $q_t(\theta)$ are increasing in θ.

Given a history till time period t, at which it is the seller who makes a price offer p_t, let $u_t(\theta)$ denote the utility (from the perspective of period t) of buyer type $\theta \geq \underline{\theta}_{s'}$ from rejecting the price offer and playing the equilibrium strategy (and hence optimally) from the next period on. We claim: (a) that $\underline{\theta}_{s'}$ is indifferent between buying at $q_t(\underline{\theta}_{s'})$ and rejecting this price, in other words, that $u_t(\underline{\theta}_{s'}) = r_t(\underline{\theta}_{s'}))$, (b) that

$$u_t(\theta) \leq u_t(\underline{\theta}_{s'}) + \theta - \underline{\theta}_{s'},$$

and (c) that any price p_t charged by the seller in any period $t < \hat{t}$ following the choice of experiment s' is bounded from below by $q_t(\underline{\theta}_{s'})$.

We first prove this claim by assuming that the support of $\hat{t}$ is bounded. We then show how to relax this assumption. Given the boundedness of the support of $\hat{t}$, its lowest upper bound, which we denote by $\hat{T}$, is a finite integer. We run backward induction starting at $\hat{T}-1$. As at $\hat{T}$ the bargaining power is with the buyer, the buyer is thus able to extract all rents at $\hat{T}$, and hence (a) and (b) from above hold. This implies that if the seller has the bargaining power at $t = \hat{T} - 1$, then the seller asks for at least $q_t(\underline{\theta}_{s'})$ because at any strictly lower price, all buyer types would buy for sure; thus (c) holds as well.

For the inductive step, suppose that the seller makes price offer at t^* and that (a), (b), and (c) hold true in periods $t = t^* + 1, ..., \hat{T} - 1$. The law of iterated expectations implies that the buyer of type $\underline{\theta}_{s'}$ (who might or might not be on the equilibrium path at this moment) would be indifferent between buying at price $q_{t^*}(\underline{\theta}_{s'})$ and rejecting (and playing optimally thereafter) because (a) and (b) holds in the next period, t^*+1, if the seller is still making offers then, and also if there is bargaining power switch in $t^* + 1$. Hence (a) is true also at time t^*. In particular,

$$u_{t^*}(\underline{\theta}_{s'}) = \delta\mathbb{P}(\hat{t} = t^* + 1)\underline{\theta}_{s'} + \delta\mathbb{P}(\hat{t} > t^* + 1)u_{t^*+1}(\underline{\theta}_{s'}) \tag{1}$$

is the utility of this buyer type under both alternatives. Further, buyer of type $\theta > \underline{\theta}_{s'}$ who rejects the price offer at t^*, has utility of at most:

$$\delta\mathbb{P}(\hat{t} = t^* + 1)\theta + \delta\mathbb{P}(\hat{t} > t^* + 1)\max(\theta - q_{t^*+1}(\underline{\theta}_{s'}), u_{t^*+1}(\theta))).$$

By the part (a) of the inductive assumption, $\theta - q_{t^*+1}(\underline{\theta}_{s'}) = u_{t^*+1}(\underline{\theta}_{s'}) + \theta - \underline{\theta}_{s'}$. Hence, the upper bound on the utility we just derived is (by 1) weakly lower than:

$$\delta\mathbb{P}(\hat{t} = t^*+1)\theta + \delta\mathbb{P}(\hat{t} > t^*+1)(u_{t^*+1}(\underline{\theta}_{s'}) + \theta - \underline{\theta}_{s'}) = u_{t^*}(\underline{\theta}_{s'}) + \delta(\theta - \underline{\theta}_{s'}).$$

This proves that (b) also holds true in period t^*.

By (a), which we already proved for time t^*, buying at price $q_{t^*}(\underline{\theta}_{s'}))$ gives type θ utility $u_{t^*}(\underline{\theta}_{s'}) + \theta - \underline{\theta}_{s'}$. Hence, this type θ weakly prefers to buy at time t^* at price $q_{t^*}(\underline{\theta}_{s'}))$ over rejecting this price. As a consequence, all types $\theta \geq \underline{\theta}_{s'}$ strictly prefer to buy at any price strictly lower than $q_{t^*}(\underline{\theta}_{s'})$, and hence, in any equilibrium, the seller's price is at least $q_{t^*}(\underline{\theta}_{s'})$ also in period t^*. This proves (c).

The prices the seller offers are bounded below by $q_t(\underline{\theta}_{s'})$ and the utility from waiting for type θ of the buyer, $u_t(\theta)$, is bounded from above by $u_t(\underline{\theta}_{s'}) + \theta - \underline{\theta}_{s'}$ also when the support of $\hat{t}$ is unbounded. The argument extends the inductive logic of the bounded-support case. The boundedness of the type space Θ implies that there is some $K > 0$ such that $u_t(\underline{\theta}_{s'})$ is bounded from above by $r_t(\underline{\theta}_{s'})+K$ and for any $\theta \geq \underline{\theta}_{s'}$, the utility $u_t(\theta)$ from waiting at any period $t = \hat{T}-1$ is bounded from above by $u_t(\underline{\theta}_{s'}) + \theta - \underline{\theta}_{s'} + K$. As above, we can then inductively infer that at any earlier time t these utilities are bounded from above as follows:

$$u_t(\underline{\theta}_{s'}) \leq r_t(\underline{\theta}_{s'}) + \delta^{\hat{T}-1-t}K,$$

and

$$u_t(\theta) \leq u_t(\underline{\theta}_{s'}) + \theta - \underline{\theta}_{s'} + \delta^{\hat{T}-1-t}K.$$

Hence, the price charged by the seller is bounded from below by $p_t \geq q_t(\underline{\theta}_{s'}) - \delta^{\hat{T}-1-t}K$. Having proven these inductive inequalities, we can take $\hat{T}$ to be arbitrary large, while keeping K constant, and infer that the all parts (a), (b), and (c) of the claim hold true.

The bounds on prices and the utility from waiting we established imply that any type θ that picks s' in the postulated equilibrium gets at most $\theta - \underline{\theta}_{s'}$ more than this type would obtain by picking s' and waiting till bargaining power shift at $\hat{t}$. The waiting strategy is available following any experiment choice s, and, thus, any type $\theta < \underline{\theta}_{s'} + w$ strictly prefers to deviate to a cheapest experiment s and then waiting till $\hat{t}$. This contradiction concludes the proof. ■

Proof of Theorem 2. The argument is analogous to the proof of Theorem 1 with two adjustments needed. First, $b(\theta)$ rather than θ is now the buyer's value; the price offered to the buyer is hence bounded from below by $\inf b(\theta)$ with inf taken over θ in the support of the seller's equilibrium belief. Second, we must allow for the possibility that the seller chooses not to sell at some period, e.g., by setting a prohibitively high price. This possibility does not break our argument because conditional on such prohibitively high price the buyer still prefers to choose the cheapest experiment. ■

Proof or Proposition 4. The proof follows the same steps as the proof of Theorem 1. Specifically, consider the infimum of types who choose s_p. This type has a strict incentive to deviate to s' given the argument developed in the proof of Theorem 1. Hence, s_p can never be chosen in equilibrium. ■

Proof of Proposition 5. The continuation game played by the buyer and the seller after the transfer from the seller to the provider is agreed on and the provider chooses the set of platforms S and their fees satisfies the assumptions of Theorem 2. We thus know from Theorem 2 that the buyer will choose the offline option if the lowest platform fee is strictly above 0; such high fees hence cannot be on equilibrium path as the platform provider would have profitable deviation in which the lowest fee is negative (subsidy). We also know from Theorem 2 that the buyer will choose the cheapest platform whenever the lowest platform fee is strictly below 0. Hence no fee strictly below 0 can be on equilibrium path as the provider would benefit by raising the fee while still keeping it strictly below 0. Theorem 2 also implies that if the lowest fee was 0 and the buyer would choose the offline option with positive probability then the provider would have a profitable deviation in which they slightly lower the fee below 0 .

In effect, we can conclude that in any equilibrium the lowest fee is exactly 0 and the buyer chooses a cheapest platform offered by the provider. The monotonicity of provider's profit in the expected seller's profit then implies that in any equilibrium the provider offers a set of platforms maximizing the seller's expected profits and one of these platforms is chosen by the buyer.

■

References

[1] Akerlof, G. (1970). The Market for Lemons: Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics*, 84(3): 488-500.

[2] Acharya, V., P. DeMarzo, and I. Kremer. (2011). Endogenous Information Flows and the Clustering of Announcements. *American Economic Review*, 101(7): 2955-79.

[3] Acquisti, A., L. John, and G. Loewenstein. (2013). What is Privacy Worth. *Journal of Legal Studies*, 42(2): 249-274.

[4] Acquisti, A. and H. Varian (2005). Conditioning Prices on Purchase History. *Marketing Science*, 24(3), 367-381.

[5] Acquisti, A., C. Taylor and L. Wagman. (2016). The Economics of Privacy. *Journal of Economics Literature*. 54(2): 442-92.

[6] Ali, N., G. Lewis, and S. Vasserman. (2023). Voluntary Disclosure and Personalized Pricing. *Review of Economic Studies*, 90(2): 538-571.

[7] Athey, S., C. Catalini, and C. Tucker. (2017). The Digital Privacy Pardox: Small Money, Small Cost, Small Talk. NBER WP. 23488.

[8] Barth, S. and M. de Jong. (2017). The privacy paradox – Investigating discrepancies between expressed privacy concerns and actual online behavior – A systematic literature review. *Telematics and Informatics*, 1038-1058.

[9] Ben-Porath, E., E. Dekel, and B. Lipman. (2018). Disclosure and choice. *Review of Economic Studies*, 85: 1471-1501.

[10] Berendt, B., O. Gunther, and S. Spiekermann. (2005). Privacy in e-commerce: stated preferences vs. actual behavior. Communications of the ACM, 48(4): 101-106.

[11] Bergemann, D., B. Brooks, and S. Morris. (2015). The Limits of Price Discrimination. *American Economic Review*, 105(3): 921-957.

[12] Bergemann, D., and S. Morris. (2019). Information Design: A Unified Perspective. *Journal of Economic Literature*, 57(1): 44–95.

[13] Board, S., and J. Lu (2018). Competitive Information Disclosure in Search Markets. *Journal of Political Economy*, 126(5): 1965-2010.

[14] Calzolari, G., and A. Pavan. (2006). On the optimality of privacy in sequential contracting. *Journal of Economic Theory*, 130(1): 168-204.

[15] de Corniere, A., and R. de Nijs. (2016). Online advertising and privacy. *Rand Journal of Economics*, 47(1): 48-72.

[16] Crawford V. P. and N. Iriberri. (2007). Level-K Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions? *Econometrica* 75(6), 1721-1770.

[17] Crawford, V., and J. Sobel. (1982). Strategic Information Trans-

mission. *Econometrica*, 50(6): 1431-1451.

[18] Dye, R. (1985). Disclosure of nonproprietary information. *Journal of Accounting Research*, 23: 123-145.

[19] Gentzkow, M. and E. Kamenica. (2017). Costly Persuasion. *American Economic Review*, 104(5): 457-462.

[20] Goldfarb, A. and C. Tucker. (2011). Privacy Regulation and Online Advertising. *Management Science*, 57(1): 57-71.

[21] Grossman, S. and O. Hart. (1980). Disclosure Laws and Takeover Bids. *Journal of Finance*, 35(2): 323-334.

[22] Grossman, S. (1981). The Informational Role of Warranties and Private Disclosure about Product Quality. *Journal of Law and Economics*, 24: 461-483.

[23] Hart, O. and J. Tirole. (1988). Contract Renegotiation and Coasian Dynamics. *Review of Economic Studies*, 55(4): 509-540.

[24] Hidir, S. and N. Vellodi. (2021). Privacy, Personalization, and Price Discrimination. *Journal of the European Economic Association*, 19(2): 1342-1363.

[25] Johnson, G., S. Shriver, and S. Du. (2020). Consumer Privacy Choice in Online Advertising: Who Opts Out and at What Cost to Industry? *Marketing Science*, 39(1): 33-51.

[26] Jovanovic, B. (1982). Truthful Disclosure of Information. *Bell Journal of Economics*, 13 (1): 36-44.

[27] Kamenica, E. and M. Gentzkow. (2011). Bayesian Persuasion. *American Economic Review*, 101(6): 2590-2615.

[28] Kamenica, E. (2019). Bayesian Persuasion and Information Design. *Annual Review of Economics*, 11: 249-272.

[29] Kim, J-H. and L. Wagman. (2015). Screening incentives and privacy protection in financial markets: a theoretical and empirical analysis, *RAND Journal of Economics*, 46(1): 1-22.

[30] Kim, K. and P. Kircher. (2015). Efficient Competition through Cheap Talk: The Case of Competing Auctions. *Econometrica*, 83(5): 1849-1875.

[31] Kreps, D., and J. Sobel. (1994). Signaling. Chapter 25 in *Handbook of Game Theory*, Volume 2, Edited by R.J. Aumann and S. Hart.

Elsevier Science B.V.

[32] Laffont, J.-J. and J. Tirole. (1988). The Dynamics of Incentive Contracts. *Econometrica*, 56(5): 1153-1175.

[33] Madarasz, K. and M. Pycia. (2020). Towards a Resolution of the Privacy Paradox. *https://works.bepress.com/kristof_madarasz/47/.*

[34] Madarasz, K. (2021). Bargaining under the Illusion of Transparency. *American Economic Review*, 111(11): 3500-3539.

[35] Madrian, B. and D. Shea. (2001). The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *Quarterly Journal of Economics*, 116(4): 1149-1187.

[36] Milgrom, P. (1981). Good News and Bad News: Representation Theorems and Applications. *Bell Journal of Economics*, 12: 380-391.

[37] Milgrom, P. and J. Roberts. (1986). Relying on the information of interested parties. *RAND Journal of Economics*, 17(1): 18-32.

[38] Milgrom, P. (2008). What the Seller Won't Tell You: Persuasion and Disclosure in Markets. *Journal of Economic Perspectives*, 22(2): 115-131.

[39] Mirkin, K. and M. Pycia. (2015). Directed Search and the Futility of Cheap Talk. UCLA Working Paper.

[40] Mikians, J., L. Gyarmati, V. Erramilli, and N. Laoutaris. (2012). Detecting Price and Search Discrimination on the Internet. *Proceedings of the 11th ACM Workshop on Hot Topics in Network*, 79-84. New York: Association for Computing Machinery.

[41] Mikians, J., L. Gyarmati, V. Erramilli, and N. Laoutaris. (2013). Crowd-Assisted Search for Price Discrimination in e-Commerce: First Results. *Proceedings of the 9th ACM Conference on Emerging Networking Experiments and Technologies*, 1-6. New York: Association for Computing Machinery.

[42] Myerson, R. and M. Satterthwaite. (1983). Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory*, 29(2): 265-281.

[43] Pram, K. (2021). Disclosure, welfare and adverse selection. *Journal of Economic Theory*, 197(C).

[44] Ravid, D., A.-K. Roesler, and B., Szentes. (2022). Learning before

trading: On the inefficiency of ignoring free information. *Journal of Political Economy*, 130(2): 346-387.

[45] Riley, J. (2001). Silver Signals: Twenty-Five Years of Screening and Signaling. *Journal of Economic Literature*, 39: 432-478.

[46] Roesler, A.-K. and B. Szentes. (2017). Buyer-Optimal Learning and Monopoly Pricing. *American Economic Review*, 107(7): 2072-80.

[47] Rose, E. (2005). Data Users versus Data Subjects: Are Consumers Willing to Pay for Property Rights to Personal Information? *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*. Washington, D.C.: IEEE Computer Society Press.

[48] Spence, M. (1973). Job Market Signaling. *Quarterly Journal of Economics*, 87(3), 355–79.

[49] Starkov, E. (2023). Only Time Will Tell: Credible Dynamic Signaling. Working Paper.

[50] Strausz, R. (2017). Mechanism Design with Partially Verifiable Information. Working Paper.

[51] Taylor, C. (2004). "Consumer Privacy and the Market for Customer Information." *Rand Journal of Economics*, 35(4) 631-650.

[52] Thereze, J. (2022). Screening Costly Information. mimeo, Princeton.

[53] Terstiege S. and C. Wasser. (2021). Product Information Provided by Sellers and Platforms. Working Paper.

[54] Turow, J., J. King, C. J. Hoofnatle, A. Bleakley, and M. Henessey. (2009). Americans Reject Tailored Advertising and Three Activities That Enable It. Departmental Papers (ASC), 137.

[55] Verrecchia, R. (1983). Discretionary disclosure. *Journal of Accounting and Economics*, (5): 179-194.