Cost over Content: Information Choice in Trade

Kristóf Madarász (LSE) and Marek Pycia (U Zurich)*

January 2025

Abstract

How much would buyers pay to have some control over what a seller knows about them? When deciding what information to provide to her counterpart, a privately-informed trader chooses between options that may differ both in their contents and in their costs. For a large class of static and dynamic trading environments where buyers choose from arbitrary sets of signal processes that reveal or obfuscate information to a seller, we establish a "cost-over-content" theorem. In equilibrium, buyers only choose cheapest processes, regardless of the information content they provide. Pooling on any cheapest process is an equilibrium. Our paper uncovers a general source of market failure linked to the direct cost of information choice with consequences for the role of information defaults. We explore applications to bargaining, signaling, disclosure, consumer privacy, and data trade.

Keywords: Communication Costs, Default-based Regulation, Data Trade, Dynamic Pricing, Privacy Paradox, Signaling, Disclosure, Efficiency.

*First presentation 2016. First posted draft 2019. Previously titled "Towards a Resolution of the Privacy Paradox," on bepress™ archive; "Information Choice: Cost over Content" CEPR DP18252; and "Information Choice in Trade: Cost over Content" SSRN 5078271. For their comments, we would like to thank Dilip Abreu, Jean-Michel Benkert, Laura Doval, Ernst Fehr, Amanda Friedenberg, Tan Gan, Bob Gibbons, Marina Halac, Jason Hartline, Shengwu Li, Meg Meyer, Benny Moldovanu, Stephen Morris, Pietro Ortoleva, Marco Ottaviani, Alessandro Pavan, Ferdinand Pieroth, Andrea Prat, Vasiliki Skreta, Peter Norman Sørensen, Balazs Szentes, Steve Tadelis, Curtis Taylor, Leeat Yariv, and audiences at LSE, UCLA, Zurich, Maastricht, CEU, Alpine Game Theory, UCL-LSE Theory Fest, Padua, Bonn Winter Theory Conference,CEPR Workshop of Incentives and Contracts, St. Gallen Workshop on Dynamic Contracts and Learning, ESSET Gerzensee, Swiss Theory, Conference on Institution and Mechanism Design in Honor of Vincent Crawford. The project received funding from the European Research Council under EU Horizon 2020 grant agreement No 866376.

1 Introduction

Private information has a profound impact on trade. When the distribution of private information is fixed, trade outcomes often critically depend on it, as established, e.g., by Akerlof (1970), Coase (1972), Myerson and Satterthwaite (1983). However, opportunities for a privately informed party to impact what her trading partner learns, either before or during trade, are ubiquitous. A consumer may take costly actions to protect her privacy when shopping online, e.g., pay for apps that track her less. A supplier may choose between different warranty lengths or assemble different pieces of evidence to convey the quality of its technology to a business interested in purchasing it. A homeowner may volunteer evidence to an insurance company. An employee seeking promotion may decide to take a course whose difficulty depends on her privately known talent, and then also decide whether to disclose her grade.

Such choices are often costly, but may benefit the sender by facilitating trade or improving its terms in her favor. How much would then a consumer pay to implement her preferred level of privacy? What warranty length would a supplier choose and how does she trade off the information communicated and the cost of providing this information? When does a homeowner want to disclose costly evidence voluntarily and how does it depend on what she thinks the insurance company already knows? How much effort would an employee put into revealing her grade?

The choice of information, such as communication or information provision, can take many forms. Spence (1973) introduced costly signaling in trade and showed that if different signals are *differently costly* to send, and such cost differences depend on the sender's private information, then a trader may credibly communicate her private information even in the absence of evidence. Crawford and Sobel (1982) showed that costless communication can credibly convey information if the preferences of the sender and the receiver are sufficiently congruent. A trader may also choose to disclose evidence, as in the classic work of Grossman (1981), Milgrom (1981), or Verrecchia (1983), who showed how voluntary disclosure may eliminate a trader's private information and restore efficient trade in the context of interdependent values. Alternatively, a trader may be able to obfuscate, destroy, or hide evidence, as in, e.g., Dye (1985), or protect their privacy as in, e.g., Acquisti, Taylor, and Wagman (2016).

Understanding such choices is potentially key for predicting and analyzing trade outcomes. This is especially so in light of well-known market failures in the presence of fixed information structures, e.g., Akerlof (1970).[1] How do these choices then affect the tendency of decentralized trade to achieve efficient outcomes?

Motivated by these questions, we revisit the general monopoly problem in which a seller makes price offers to a buyer, or a continuum thereof, at one or more time periods. We impose no restrictions on the link between the traders' valuations nor any restrictions on the initial distribution of private information between them, except that the buyer knows her own preferences. Trade can be subject to adverse or advantageous selection, as in insurance problems with privately known risk types, e.g., Einav and Finkelstein (2021), or there may be some match-specificity where different seller types may have different rankings over buyer types in terms of the cost of serving them, e.g., Horton, Johari, and Kircher (2024).

At any time period, the buyer can choose from an arbitrary finite set of dynamic signal-generating processes or information structures, which, for lack of a better term, we call experiments (also known as Blackwell experiments or tests). The buyer's choice is observed by the seller and each experiment determines the information that the seller and the buyer respectively receive over time. An experiment may *per se* be uninformative, as in classic signaling, or reveal arbitrary and potentially stochastic evidence. We permit experiments to deliver news regarding what information each player has and on the timing of information arrival itself, and to do so publicly or privately.

Different experiments may have different costs. The cost of an experiment can be related to its informativeness or to some external aspect. Buyer types may differ in the incremental cost of choosing more expensive experiments, as in Spence (1973). Alternatively, the buyer (or sender) may need to pay more to disclose or assemble more complex pieces of evidence. Such costs and cost differences between experiments may be the buyer's private information and may be arbitrarily correlated with her benefit from trade. Despite the generality of this setting, with dynamic pricing,

[1]For other analyses of endogenous information in Akerlof's lemon model, see Tirole (2012) on the impact of toxic asset repurchase programs, and Pavan and Tirole (2023) on equilibrium incentives to learn one's own value.

flexible observation structures, and arbitrary information structures, we present a number of robust insights that are also novel in simple cases of our model such as static privacy choice, voluntary disclosure, or signaling.[2]

First, when experiments do not differ in their costs, then *any* content goes and there is multiplicity of equilibria. The buyer pooling on any given sequence of experiments is a weak perfect Bayesian equilibrium. For example, the buyer's ex ante optimal choice of information design remains an equilibrium information structure even if she makes the choice at the interim stage when she already knows the payoff state. Similarly, if there is an information structure under which trade is efficient, selecting this information structure, and thus efficient trade, is always an equilibrium.

Second, however, when experiments differ in their costs, we show that in all weak perfect Bayesian equilibria, the buyer opts for a least expensive information structure at each point in time, irrespective of its content and that of more costly alternatives. Each buyer type selects some least expensive sequence of experiments regardless of whether this obfuscates or reveals information relative to other options. In other words, only "cheapest talk" matters. Given any arbitrary set of dynamic information processes, informed buyers select among these purely on the basis of costs and not on content; it is information costs and not content which selects among trade outcomes.

Third, we show that these insights remain true in a variety of extensions of our base setup: minimal assumptions on the observability of the buyer's information choice and more general allocations of the bargaining (price-setting) power. They also remain true in general history-dependent environments where the buyer's choice can also publicly impact the set of information choices she faces over time, the costs of different experiments, and the traders' valuations. In all these settings, our results imply that trade outcomes depend purely on the set of cheapest experiments. In turn, when it is costly for the buyer to change some initial (default) distribution of information, she will never do so in equilibrium.

The logic we present implies that *information defaults* may have key economic consequences and underscore their relevance for regulation. We arrive at this conclusion without any equilibrium refinement or the presence of some form of choice overload or directly assumed 'default bias.' In Section 4, we then apply our general

[2]In particular, the present paper subsumes our analysis of privacy, Madarasz and Pycia (2020).

insight to help explain why consumers may choose little protection of their private information even when protecting it would improve their welfare and show that this so-called *privacy paradox* is an equilibrium phenomenon driven by market power. We also contribute to the explanation of the observed lack of voluntary disclosure in settings with the clear presence of hard information. Methodologically, in a wide range of contexts, our logic provides an equilibrium rationale for the exogenous information distribution assumption commonly imposed in economic models.

Of course, classic results, e.g., Spence (1973), imply that there are many important trade environments where the opportunity for a privately informed trader to engage in costly information choice clearly matters. We show, however, that in a large class of environments this is not the case. Our results instead point to a novel source of market failure where this source is the direct cost of information provision or information manipulation. Small cost differences between the options the buyer can choose from can have large direct implications on decentralized trade outcomes. Since such outcomes may well affect market performance, for example under adverse selection, such cost differences may then be a direct source of market failure and impact market performance.

Our above results make no direct predictions about the content of the information the buyer ends up choosing. Still, they allow us to draw content inferences in the many contexts where there is a natural link between costs and content. In Section 4.3, we then also consider an application of our result to a simple market setting where this link emerges endogenously. We consider experiments offered by a platform provider for online shopping. The provider chooses from a set of technologically feasible experiments, e.g., privacy (tracking) options, and decides which options to offer to the buyer and at what price or subsidy each. The provider then makes a take-it-or-leave-it offer to the seller where, in exchange for a fixed fee from the seller, the provider transmits the information that the choice of the buyer will contractually permit. The buyer then has the choice to select any privacy option offered on the platform or to shop offline directly from the seller, in which case she freely preserves her existing privacy.

While the platform can charge both the buyer and the seller for its offering of information structures, following the logic of our cost-over-content theorem, we show that the platform neither charges nor compensates the buyer for choosing an information structure. At the same time, the platform charges the seller and maximizes the value of the information passed to the seller, as measured in the seller's direct profit from interacting with the buyer. A corollary of this result is that even though the platform has considerable bargaining power over the seller and buyers have full property rights over their data, in many cases buyers may be strictly better off when such *data trade* is banned, that is, they can only shop offline.

Finally, in Section 6 we apply our cost-over-content logic to the classic problem of Myerson and Satterthwaite (1983) considering trade with private values. We allow each trader to first choose any information structure. They then participate in any *fixed* direct bargaining protocol that is both budget balanced and ex post individually rational. Under mild assumptions, we show that efficient trade is possible if and only if efficient trade is possible under the cheapest (default) information structure. No matter what the bargaining protocol is or what soft or hard evidence the traders may be able to provide, our insight that the efficiency of trade hinges only on the cheapest information structure continues to hold. We discuss relations to the empirical literature and default-based regulation in the Conclusion.

1.1 Related Literature

Our paper relates to various strands in the literature. The literature on costly signaling emphasizes the direct role of type-dependent cost differences for information provision; see, e.g., Spence (1973). As we discuss in more detail in Section 4, we differ from this literature regarding the assumptions on the allocation of the bargaining power and/or the initial distribution of private information. Here, our results highlight the importance of the common assumption built into signaling models: that by signaling, agents receive a payoff change commensurate with the receiver's valuation.[3]

We also relate to the literature on voluntary disclosure which considers particular evidence structures that the privately informed trader can disclose to her counterpart; see, e.g., Grossman and Hart (1980), Dye (1985). A key insight of this literature is the full disclosure benchmark. In the class of environments they study, Grossman

[3]Cf. Kreps and Sobel (1994) and Riley (2001) for surveys of signaling models all of which share this feature.

(1981) and Milgrom (1981) show that there is a unique perfect equilibrium, in which all types choose full disclosure, see also, e.g., Seidmann and Winter (1997). Verrechia (1983) shows that when disclosure is costly, then low types do not disclose and that the equilibrium converges to full disclosure as this cost goes to zero. In Sections 2 and 4 we describe how our setup captures the evidence structure of voluntary disclosure (including Dye's model), and provide a detailed comparison to these results and their logic.

Third, we contribute to the literature on privacy. Acquisti, Taylor, and Wagman (2016) provide an example where a seller with a known cost can elicit full disclosure by a buyer for a small discount linking the argument to Grossman and Milgrom as discussed above.[4] Calzolari and Pavan (2006a, 2006b) 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 number of papers explore the ways in which consumers may benefit from their or the seller's private information being transmitted to improve match quality; see, e.g., Hidir and Vellodi (2021). We discuss other work on privacy in Section 4.

Our paper also links to the expanding literature on information design, e.g., Kamenica and Gentzkow (2011), which considers one party's observable choice of the information structure in the game. Gentzkow and Kamenica (2014) consider costly information design, e.g., with the cost of each experiment linked to its respective entropy reduction, and show that key insights of the problem they study continue to hold when allowing for such costs. Doval and Skreta (2024) show that these insights also hold in the presence of constraints on information transmission. Our approach aligns with the information design literature in that we allow the sender to choose from a large set of arbitrary experiments. We, however, allow arbitrary constraints on this set and, more importantly, we are closer to the literature on disclosure and signaling in that the sender is privately informed when picking the receiver's information, and there is a (privately known) cost to such information choice. We also refer to Koessler and Skreta (2023) who consider *interim* information design in a con-

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

text where all information designs are costless and describe environments where ex ante optimal mechanisms are also interim optimal. We also link to the literature on the interaction of private information and the holdup problem, e.g. Gibbons (1992), and Gul (2001). In Section 5, we discuss the key difference between information and investment choice.

We contribute to the study of market failures, both in environments with interdependent values (as in Akerlof 1970) and independent values (Myerson and Satthertwite 1983) by showing that the efficiency of trade hinges on the properties of the cheapest information structures. In particular, we show that the endogeneity of information structures does not resolve classical inefficiency insights.[5] We also show that the price dynamics in sales of durable goods (e.g., Coase 1972, Gul, Sonnenschein, and Wilson 1986) is determined by the cheapest information structures.

Finally, our results also relate to the broader economic literature on the role of easy-to-change defaults or *choice architectures* in determining economic outcomes. In this literature, such as in Madrian and Shea (2001) or Thaler and Sunstein (2008), defaults matter because of consumer inertia, choice overload, procrastination, or a so-called 'default bias', etc. Defaults might also anchor strategic reasoning (e.g., Crawford and Iriberri 2007. In contrast, in our environment, such defaults may fully determine equilibrium choices without any behavioral considerations.

2 Setup

A buyer, or a continuum thereof, and a seller can trade an asset. Each trader's valuation for the asset is determined by the state of nature $\omega \in \Omega$, where there is σ-algebra on measurable subsets of Ω and an associated and commonly known probability measure $\mathbb{P}$. In describing our setup below, we balance generality and tractability and thus introduce some extensions only later, in Sections 3 and 5.

Valuations. The buyer's valuation is denoted by $b(\omega) \in \mathbb{R}$, the seller's valuation by $v(\omega) \in \mathbb{R}$. Both traders receive the value of the object at the time of their transaction.

[5]For the impact these theorems had on economic policy, see Milgrom (2004), Tirole (2012), and Loertscher, Marx, and Wilkening (2015).

We thus interpret the seller's value as arising from avoiding the cost of servicing the buyer.[6] We impose *no* restrictions on the relationship between $b(\omega)$ and $v(\omega)$. Values may be private or interdependent, such, as in insurance where the seller's cost of servicing a buyer may depend on the buyer's privately known risk type. Values may be positively or negatively related, thus trade can be subject to either adverse or advantageous selection.[7] Values may also have no monotonic relationship at all. For instance, different seller types may have different costs of servicing different buyer's types allowing for match-specific trade surplus.

Initial Information. We allow for an arbitrary initial distribution of information and assume only that the buyer knows her own valuation. We impose *no* other restriction on what the parties may or may not know, nor on higher-order uncertainty. Private information may be correlated or independent.

Formally, at $t = 0$, each trader's information is captured by his or her type. The buyer's type is $\theta_b \in \Theta_b$, the seller's is $\theta_v \in \Theta_v$, where Θ_b and Θ_v are arbitrary partitions of Ω. The sets of types Θ_b and Θ_v are endowed with topologies and are compact in their respective topologies. The buyer knows her own valuation that is $b(\omega)$ is measurable with respect to Θ_b.[8] The seller may or may not know his own valuation or that of the buyer, may or may not know whether the buyer knows his valuation, etc. Similarly, the buyer may or may not know the seller's valuation, or what the seller knows about her valuation, etc.

Timing and experiments. Traders interact over periods $t \in \mathcal{T} = \{1, 2, 3, ..., T\}$ for some $T \geq 1$ or $t \in \mathcal{T} = \{1, 2, 3, ...\}$ (infinite T). In the beginning of each period $t \geq 1$, the buyer first chooses a dynamic signal process

$$s_t : \Omega \to \Delta(\times_{t' \in \mathcal{T}, t' \geq t} Z_{t',t}),$$

[6] With straightforward adjustments, we can alternatively treat the seller's (and buyer's) value as a discounted sum of flow payoffs from holding the object.

[7] As, e.g., in Einav and Finkelstein (2011).

[8] We relax this assumption in Section 5.

where $Z_{t',t}$ is the support of signal realizations of period t experiments in period $t' \geq t$.[9] We refer to these signal processes as *experiments* and denote the set of experiments that are feasible at time t as S_t. Each S_t is a non-empty finite set. In case S_t is a singleton set, the buyer has no real experiment choice in round t. Each signal realization is either public or observed privately by the seller. Each experiment can reveal information, to either player, not only about payoffs, but about the other player's information, about the timing of the arrival of payoff relevant information, the meaning of past or future experiments, etc. Section 2.1 provides examples of experiments.

Following the buyer's choice from S_t in period t, the seller observes the buyer's experiment choice and all the signal realizations generated by the chosen experiments until (and including) period t. He then makes a price offer p_t which the buyer can accept or reject. If she accepts it, the game terminates. If she rejects it, a new period starts iff $t < T$; a rejection at T also terminates the game. We denote by $\mathcal{H}$ the set of all histories in the game.

Experiment Costs. The key aspect of our setup is that different experiments may have different costs. Let $c(s_t, \omega) \in \mathbb{R}$ be the bounded cost of choosing experiment s_t in period t incurred by the buyer (sender) in period t. We assume that $c(s_t, \omega)$ is determined by $\theta_b \in \Theta_b$ and that the mapping $\Theta_b \ni \theta_b \to c(s_t, \omega)$ is continuous. In particular, the buyer knows the cost of each experiment at each time, allowing us to simplify the notation and denote the experiment costs by $c(s_t, \theta_b)$. Costs may be the buyer's private information.

We assume that for each t there is a subset of experiments $\underline{S}_t \subseteq S_t$ that have some state-independent lowest cost $\underline{c}_t \in \mathbb{R}$ among all experiments in S_t. In other words, the class of lowest-costs experiments is common across buyer types, and the lowest cost associated with this class is also common across buyer types. We impose no other ordinal or cardinal restriction on the cost function. The above assumption is only for expositional simplicity and we will adopt a significantly weaker assumption in Section 5 (Theorem 2).

[9] We could equivalently represent information structures as functions from states of nature to a realization space without needing to invoke lotteries.

Our assumption allows both cardinal cost differences to be type dependent and the ordinal cost-based ranking of experiments to be type dependent as well. It is automatically satisfied when the costs are independent of ω. Similarly, it is also always satisfied when changing the initial distribution of information has some positive, possibly type-dependent, cost, but not changing it is free. In many settings of economic interest, there is indeed an initial or default distribution of information the players are endowed with, and choosing some alternative information structure carries some potentially type-dependent cost.

The cost of an experiment may be related to its informativeness, such as entropy reduction, or to some external factor, such as the state-dependent effort needed to complete a course or certify a piece of evidence. The cost does not need to be positive. A negative cost may 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 can purchase online from the seller.

Payoffs and equilibrium. The players discount payoffs given possibly different interest rates $r_b, r_s > 0$ where $\Delta > 0$ is the length of the period. When an offer is accepted, at some period t^*, the buyer's payoff is $e^{-r_b\Delta(t^*-1)}(b(\omega)-p)-\sum_{t=1}^{t^*} e^{-r_b\Delta(t-1)}c(s_t,\omega)$ and the seller's payoff is $e^{-r_v\Delta(t^*-1)}(p-v(\omega))$. In case there is no trade, payoffs from trade are normalized to zero, but the buyer still incurs the costs of the chosen experiments. We employ the standard notion of weak perfect Bayesian equilibrium (PBE or equilibrium henceforth) as our solution concept. Our characterization of equilibria thus immediately applies to all more stringent equilibrium concepts. In Appendix B, we show that, under mild additional assumptions, not only weak perfect Bayesian equilibria but also more stringent equilibria exist in the environments we study.

2.1 Examples of Experiments

To illustrate the above, we briefly describe some classic and simple instances of information choice that are special cases of our setup. In these examples, we assume that $T = 1$ and that the buyer (sender) knows the state ω. A large literature focuses on signaling, where each experiment s in itself is uninformative.

Example 1: Costly Signaling. In the canonical analysis of Spence (1973), the worker (buyer) chooses the level of education (s) that has no impact on productivity and leaves $b(\omega)$ and $v(\omega)$ unaffected. Education is differentially costly for workers of different ability (ω): more schooling is always more expensive than less, but it is comparatively less expensive for a higher-ability worker, while no schooling is costless regardless of ability. In this environment, positive education (s) is a costly experiment that is, per se, uninformative: it provides no hard information, its realizations are independent of the ability ω.[10]

Example 2: Cheap Talk. In the cheap talk environment of Crawford and Sobel (1982), each experiment (or message) $s \in S$ is again in itself uninformative. Here, all experiments have the same costs, i.e., $c(s, \omega) = c(s', \omega)$ for *all* (s, s') and ω.

In our setup, an experiment can also represent arbitrary soft or hard evidence, i.e., the realization of s may be correlated with the payoff state. Indeed, canonical settings of verifiable disclosure, including the classic evidence structures introduced by Grossman and Hart (1980), Grossman (1981), and Milgrom (1981) (GHM below) can be mapped into our setup. G-H-M assume that each message that the sender can send corresponds to a subset of Θ_b and that a given message $M \subset \Theta_b$ can only be chosen by types within this set M, i.e., $\theta_b \in M$. Types outside of M can not choose this message ('no lying'). While in our setting, the sender (buyer) can choose any experiment from S, regardless of her type, we can map their messages into our experiments as follows.

Example 3: Disclosure with evidence. Let $Z = \{true, false\}$ and, for any given $M \subseteq \Omega$, let s_M be an experiment that reveals whether ω is in M or in its complement $\Omega - M$. As long as there is a separate experiment s_M for each $M \subseteq \Omega$, the sender can disclose any superset of the true state.

Experiment s_M sent by a type in M induces the receiver posterior identical to

[10] In this example, the cost $c(s, \omega)$ associated with a given experiment s may vary with the state ω, and different experiments $s \neq s'$ are differentially costly, $c(s, \omega) \neq c(s', \omega)$. The ordinal cost-based ranking of experiments is state independent and, given the order of experiments and an order on Ω, the cost function $c(s, \omega)$ satisfies the strict Spence-Mirrlees condition: if $s > s'$ and $\omega > \omega'$ then $c(s, \omega) - c(s', \omega) > c(s, \omega') - c(s', \omega')$.

that obtained by that type selecting message M in GHM. Similarly, experiment s_M sent by a type not in M induces the receiver posterior identical to that obtained by that type selecting message $\Theta_b - M$ in GHM. In both their and our models, the types can provide hard evidence that they are in M, and any lie is immediately detected.

In the classic disclosure setting, the receiver knows the kind of hard evidence available to the sender. A large literature in economics, finance, and accounting considered a setting where the receiver is uncertain as to whether or not the sender has evidence, see, e.g., Dye (1985). We can also map this structure into our setup.

Example 4: Disclosure with uncertain evidence. Suppose that $\Omega = [0,1] \times \{0,1\}^2$ and there are three experiments $\{s_0, s, s'\}$. The payoff state is $\hat{\omega} \in [0,1]$, and there is a two-dimensional auxiliary state $(\omega_1, \omega_2) \in \{0,1\}^2$. The first component of the auxiliary state, ω_1, determines whether the sender has evidence about $\hat{\omega}$. The second component, ω_2, affects what is revealed by experiments s and s'. If $\omega_2 = 1$, experiment s reveals the evidence, if and only if, the sender has evidence, while s' never reveals evidence. If $\omega_2 = 0$, then it is the exact reverse. Experiment s_0 never reveals any evidence. The state is the sender's private information. In turn, when the sender's choice of s or s' leads to no revelation, the receiver remains uncertain as to whether there is no evidence or the sender decided to hide it, as in Dye (1985).

Another special case is consumer privacy in the context of price discrimination and personalized pricing.[11]

Example 5: Privacy Choice. The buyer chooses between different privacy policies ($s \in S$) on a search engine. Each policy s specifies what aspects of the buyer's online behavior can be recorded and revealed to the seller (cookies). A buyer may choose to be fully tracked, selectively tracked, or tracked in a noisy fashion whereby her data are grouped with that of others.[12] Her choice may affect who she is grouped with even if her anonymity is always preserved. As trade is not zero-sum, sharing data can, in principle, both increase and decrease consumer surplus.

[11] The privacy application was at the core of early drafts of our paper, e.g., Madarasz and Pycia (2020), now subsumed by the present draft.

[12] As in differential privacy, e.g., Dwork and Roth (2014).

3 Information Choice: Cost over Content

Our central result on the buyer's endogenous information choice is as follows:

Theorem 1

Existence: *For any sequence $\{s_t\}_t$ such that $s_t \in \underline{S}_t$ for each t, there exists a PBE where all buyer types choose this sequence on the equilibrium path.*

Cost over Content*: In no PBE does the buyer ever choose any non-cheapest experiment, $s_t \notin \underline{S}_t$, in any t, on the equilibrium path.*

The main message of the theorem is that, given any cost structure, and any initial information, in *no* equilibrium does the buyer ever choose any non-cheapest experiment. The sender only engages in *cheapest talk* and the informational content of the experiments does not affect this conclusion. In particular, if there is a unique cheapest experiment in each period, that is, $\underline{S}_t$ is a singleton in each t, then there is an equilibrium in which all buyer types choose this cheapest experiment in each period and in all equilibria all buyer types do so. On the other hand, if there are no cost differences between the experiments, that is, $\underline{S}_t = S_t$ for all t, then pooling on *any* given sequence of experiments can be supported in equilibrium. Finally, buyer types may also separate and choose different experiments, but they will only choose cheapest ones.

Proving the first part of the theorem requires constructing a pooling equilibrium on a cost-minimizing sequence of experiments. The argument is straightforward, as off path, in a weak PBE the seller might believe that he faces the highest value buyer.[13] Proving the second part of the theorem takes more work. The central claim in the proof is that, because the seller's prices respond rationally to information in each period, following each experiment choice there is always a buyer type whose expected continuation equilibrium surplus is low relative to the minimum continuation surplus this type would achieve following a deviation to any other experiment. In consequence,

[13] In Appendix B, we show that, under mild assumptions on costs, such equilibria exist even if we restrict off-path beliefs to satisfy the sequentiality criteria or to be consistent with the information revealed by experiments.

if some types chose a non-cheapest experiment, then, irrespective of the content of this experiment, some of these types would benefit from a deviation to a cheaper experiment.

Within our setup, the theorem provides a foundation for the common modeling assumption that the information structure is exogenous in contexts where there is a default (unique cheapest) experiment and the buyer, at some potentially arbitrarily small cost, can choose another experiment, thus changing the information received by the seller. Our theorem implies that the presence of this choice does *not* change the outcome of trade. As long as the class of cheapest experiments is held constant, the set of non-cheapest experiments available to the sender does not matter for the possible equilibrium trade outcomes.

The theorem also provides a prediction on information choice. It describes what aspects of the problem need to be observed by an econometrician to be able to predict outcomes. The outcome depends only on the identity of the cheapest information structures, despite the great potential complexity of the environment. One may have conjectured that endogenous information choice would greatly complicate the econometric task in our setup by needing to know all available information structures and their costs. Instead, Theorem 1 implies that all that matters for prediction here is the knowledge of the set of cheapest information structures.

3.1 Observability of Information Choice

Above, we assumed that the seller always observed the buyer's choice from S_t. While the assumption that an agent observing a signal realization knows what experiment generated it is common in information economics, it is admittedly a strong assumption. It is also an assumption that we can substantially relax. For the cost-over-content insight of Theorem 1, it suffices that the receiver observes whether or not the buyer chose a cheapest experiment.[14] Which of the non-cheapest experiments was chosen (if any) can be fully *unobservable*.

[14] Relatedly, the existence of the equilibrium pooling on specific cheapest experiments requires the seller to observe only whether the buyer chose this specific cheapest experiment.

Proposition 1 *As long as the seller observes whether or not the buyer chooses an element of $\underline{S}_t$ in any given t, there is no PBE where the buyer chooses any $s_t \notin \underline{S}_t$ in any t on the equilibrium path.*

We can further allow for a flexible observation structure regarding the buyer's choice of experiment. Suppose that for each S_t there is a partition on S_t given by some $O_t : S_t \to 2^{S_t}$. For any choice of experiment by the buyer, $s_t \in S_t$, the seller observes that the buyer has chosen an experiment from $O_t(s_t)$, but not which experiment within $O_t(s_t)$.

Proposition 2 *Suppose that for some period t and experiment $\hat{s}_t \in S_t$ we have $O_t(\hat{s}_t) \cap \underline{S}_t = \emptyset$. Then there is no PBE in which the buyer chooses $\hat{s}_t$ in period t on the equilibrium path.*

Generalizing Proposition 1, this result shows that the buyer never chooses an experiment that belongs to a partition cell that does not contain a cheapest experiment.

3.2 Discussion

The key aspects of our setup are that (i) the buyer (in her capacity as sender) chooses interim, and (ii) that different information structures *may* be differentially costly. An implication of Theorem 1 is that in the absence of cost differences between different experiments, the ex ante and interim problems are closely linked. It is the presence of such cost differences that marks a sharp difference between the ex ante and interim problems. For simplicity, below, we focus on the one-period case ($T = 1$) of the base setup of Section 2.

Ex ante versus Interim Choice of Information. Suppose that experiment costs are homogeneous, that is, $c(s, \omega)$ is independent of $s \in S$. The cost can still depend on ω, but conditional on ω, this cost is the same across all experiments. The next corollary states that here any equilibrium outcome that the buyer can achieve by committing to a given experiment $s \in S$ ex ante can also be achieved as an equilibrium outcome when choosing from S interim. Below, we then compare the game where the buyer chooses from S ex ante, before learning her type θ_b (that is before having any private information vis-a-vis the seller), with the game where she chooses from S interim, that is after learning her type θ_b, as in our setup. In every other way, the two games are equivalent and we maintain the assumption that the buyer always knows her valuation before deciding to trade and that this is common knowledge.[15]

Corollary 1 *Suppose that $c(s, \omega)$ is independent of s. For any PBE of the ex ante game where the buyer chooses a given experiment $s \in S$ ex ante, there exists a PBE of the interim game where the buyer chooses the same experiment $s \in S$ interim. Furthermore, the ex ante expected payoff of each player is the same in the ex ante game and in the interim game.*

This corollary complements the inscrutability principle of Myerson (1983) and its general development by Koessler and Skreta (2023); the latter paper studies the relation between ex ante and interim implementation in general games without communication constraints we allow. The equivalence of Corollary 1 critically fails when we allow for more general cost structures, e.g., it fails in classic signaling.[16]

Corollary 1 tells us that interim information choice can support ex ante information design in the absence of cost differences. The static ex ante problem was studied, e.g., by Bergemann, Brooks, and Morris (2015) under the assumption that all experiments are free, $c(s, \omega) = 0$ for $\forall(s, \omega)$ and $v(\omega) = v$. They identify the information structure s_b^* that maximizes the buyer's ex ante expected surplus from trade. If the buyer made an observable choice ex ante, she would then choose this experiment. Corollary 1 implies that if $s_b^* \in S$, then all buyer types pooling on s_b^* interim is also a PBE.

In the presence of costs differences, the interim and ex ante problems become rather different. If the buyer chose *ex ante*, she would have a considerable willingness to pay for s_b^*. Furthermore, she would then be subject to a smooth trade-off between the expected consumer surplus associated with a given experiment's content and its cost. Instead when she chooses *interim*, this trade-off is mute. Theorem 1 implies that

[15]This corollary follows because our proof of the existence part of Theorem 1 only relies on assumptions that hold true in the corollary.

[16]In absence of communication costs, Mirkin and Pycia (2015) show that only pooling equilibria arise when sellers send cheap-talk messages to buyers to attract them to visit. When the sellers cannot freely set prices, Kim and Kircher (2015) construct separating equilibria.

her willingness to pay extra for any experiment is null. Our next example illustrates this further.

Example 6: Adverse Selection, Akerlof (1970). Let $\Omega = [1, 2]$; the buyer privately knows ω while the seller's prior is uniform. The buyer's value is $b(\omega) = \omega$ and the seller's cost is, say, $v(\omega) = \frac{2}{3}\omega^2$.

- If the buyer cannot disclose any hard information, there is no trade in equilibrium.
- If the buyer can freely disclose whether or not $\omega \leq \frac{3}{2}$, there is an equilibrium where she discloses this information. With disclosure, the seller sells only to buyer types $\omega \leq \frac{3}{2}$; these buyer types can all receive a positive surplus.
- However, if the above disclosure carries some arbitrarily small cost to the buyer, while no disclosure is costless, Theorem 1 implies that no buyer type will ever disclose. Hence, there is no trade in equilibrium.
- If full disclosure is the cheapest experiment, trade will always be ex post efficient with each buyer type receiving zero surplus even if any type is able to hide any information at an arbitrarily small cost.

Theorem 1 implies that the costly choice of endogenous information by the privately informed buyer will not alter the market failure present under exogenous information. Market outcomes depend here on the characteristics of the default information structure as if such a structure was the only choice (was exogenously imposed). Small costs associated with changing this structure can then have very significant welfare implications.

More formally, note that our result implies a form of upper hemicontinuity of the equilibrium correspondence of the information choice (adopting the discrete topology on S) in the experiment costs. An experiment s with the lowest cost is always an equilibrium choice and remains an equilibrium choice as the cost difference between experiments converges to zero. At the same time, this correspondence fails to be lower-hemicontinuous. While pooling on any experiment choice is an equilibrium when experiments are equally costly, any tiny cost difference ensures that the cheapest experiment is selected in any equilibrium.

We conclude this discussion by considering a classic infinite-horizon problem.

Example 7: Coase Conjecture. Let $T = \infty$. The seller's cost is normalized to zero, $v(\omega) = 0$, the buyer's valuation is $b(\omega) = \omega$, and the value of $\omega \in [1, 2]$ is the private information of the buyer. The celebrated Coase conjecture implies that as offers become frequent, $\Delta \to 0$, the seller sells immediately at a price which converges to the lowest possible buyer value, e.g., Gul, Sonnenschein and Wilson (1986). Instead, if the buyer's valuation leaks (or is just perceived to leak) to the seller, the result may well be the opposite with the seller appropriating the full surplus from trade, (e.g., Madarasz, 2021). Suppose now that the buyer has a single non-trivial information choice and this occurs at $t = 1$. The default is that the buyer's valuation leaks to the seller over time, but the buyer can pay to decrease or eliminate the frequency of such leakage. Proposition 1 implies that even if the *investment level itself is not observable* and the seller only observes whether the buyer invested or not, the buyer never invests any amount to limit such leakage in any equilibrium. We return to the discussion of the Coase conjecture in the context of data trade in Section 4.3.2.

3.3 Extensions

Above we have made a number of assumptions about the class of trading environments we consider. Our Theorem 1 remains valid after we relax many of these assumptions. In the following remarks, we consider separate extensions of our base setup.

Remark 1: Information Structure. We assumed that signal realizations from the chosen experiments were public or the seller's private information. This ruled out some forms of *second-* and *higher-order* uncertainty that the players can have about each other's belief updates, e.g., on the timing of the arrival of information. It also ruled out the buyer privately learning about the seller's cost, or his belief about her valuation, etc. Our results are, however, robust to considering general information structures involving complex uncertainties that the players can have about each other's beliefs. Specifically, suppose that each experiment s_t generates three signal realizations: $z^s_{t,t'}$, $z^b_{t,t'}$, and $z^p_{t,t'}$ in each period $t' \geq t$, one observed privately by the seller, one observed privately by the buyer, and one observed publicly by both, respectively. These realizations can be arbitrarily correlated. The proof of Theorem 1 implies that our result continues to hold under such generalization as well. We will return to this extension in Section 5.

Remark 2: Two-sided Information Choice. Above, only the buyer chose an experiment. However, our proofs remain valid if the seller can also choose an experiment at each t, before, simultaneously, or after the buyer chooses one. Such an experiment choice can allow the seller to driect what the parties learn about the state of the world. We can allow such an experiment choice to be costly for the seller just as it is for the buyer. Theorem 1 remains valid under general two-sided information choice. We return to two-sided information choice in Section 6.

Remark 3: Bargaining Power. We adapted the standard modeling assumption that the seller (receiver) has the bargaining power as in the classic monopoly problem. This assumption can also be relaxed. Consider now a setting where there is no private information about the seller's value, another usual assumption in studies of monopoly pricing. We may then allow the bargaining power to *randomly* switch from the seller to the buyer. This entails the case where it is the buyer rather than the seller who makes *all* offers and thus has the full bargaining power.

To formally model when the buyer gains the power to make price offers, consider a pair of random variables $(\hat{t}, \hat{m})$ distributed on $\mathcal{T} \times \{0, 1\}$. The realization of this pair of random variables determines (i) the period $\hat{t} \in \mathcal{T}$ in which the bargaining power shifts from the seller to the buyer, and (ii) when, *within this period*, the buyer learns that she can make the price offers. If $\hat{m} = 0$, then the buyer learns about the power shift before choosing the period $\hat{t}$ experiment; if $\hat{m} = 1$, she learns about the shift after choosing the period $\hat{t}$ experiment, but before the party with bargaining power makes the price offer.

We assume that the distribution of $(\hat{t}, \hat{m})$ is commonly known, its realization is commonly observed, and this realization is independent of the players' choices in the game, that is, it is an exogenous stochastic change in the bargaining power.

Proposition 3 *[Bargaining Power Switch] Suppose that the bargaining power switches from the seller to the buyer at a time determined by* $(\hat{t}, \hat{m})$*.*

- *For any sequence* $\{s_t\}_t$*, such that* $s_t \in \underline{S_t}$*, there exists a PBE such that all buyer types choose this sequence on path.*
- *In no PBE does the buyer ever choose any* $s_t \notin \underline{S}_t$ *in any* t *on path.*

Extending Theorem 1, this proposition underscores that the force we identify allows every buyer type to receive a strictly positive and possibly time-varying rent in equilibrium. Our cost-over-content insight does not hinge on equilibrium rents being stationary or necessarily zero for some buyer type.

Remark 4: No Default. Finally, our baseline setup assumes that there is a class of cheapest experiments in each period which is common across buyer types. This assumption plays an important role in Theorem 1. However, our logic has a clear implication for environments where this assumption is violated. While in the absence of such a default class, pooling on a cheapest experiment loses meaning (because for different types different experiments may be cheapest), one can still consider experiments that are uniformly more costly then other experiments.

Fully relaxing the assumption that the class of cheapest experiments is common across types, we obtain:

Proposition 4 *Suppose that* $t \in \mathcal{T}$*,* $s'_t, s''_t \in S_t$*, and* $E[c(s'_t, \omega)|\theta_b] < E[c(s''_t, \omega)|\theta_b]$ *for all* $\theta_b \in \Theta_b$*. Then,* s''_t *is never chosen on path of a PBE.*

In this setting, it is hence still true that if one experiment is uniformly more costly than some other experiment, the former is never chosen in equilibrium irrespective of its content or that of the alternatives.

4 Applications

Trade with endogenous information choice plays central role in many questions of economics, finance, and accounting. Given the scope for inefficiencies in trade with private information (selection problems), there are also broad discussions about regulatory interventions such as mandated disclosure, compulsory insurance, or privacy policies. The impact of such regulations rests on the equilibrium forces shaping endogenous information choice, such as costly signaling, voluntary disclosure, or privacy. We now detail applications of our results to these forces. For simplicity of exposition, we focus on static trade ($T = 1$). We discuss empirical evidence in the Conclusion.

4.1 Signaling

Return to Example 1 where each experiment is itself *uninformative*, but the cost of any given experiment $c(s, \omega)$ may well depend on the sender's private information. For such signaling to reveal information, a separating equilibrium in information choice need to arise. Our result implies that as long as there is a unique common cheapest message, e.g., no schooling as in Spence (1973), or no warranty, then *no* information will be conveyed via equilibrium signaling. This holds without imposing any further restrictions on off-equilibrium path beliefs, on the dynamic nature of signaling, or on how signaling costs depend on the payoff relevant state.

There is no contradiction between our results and those of classic signaling. To compare them, we can classify situations along two dimensions. First, distinguish between scenarios based on whether there is private information about the receiver's valuation $v(\omega)$. Second, distinguish between scenarios depending on whether it is the sender or the receiver who has the power to set the price. Four possible scenarios arise in this classification, and our logic applies to three of them, see Table 1.

	no private info about $v(\omega)$	private info about $v(\omega)$
Receiver sets the price	cost over content	cost over content
Sender sets the price	cost over content	Spence (1973)

Table 1: Costly Signalling

Theorem 1 establishes the first row and Proposition 3 completes the second row of the above table. The setup of Spence (1973), and a literature that followed it, e.g., Kreps and Sobel (1998), falls in the last box where signaling can lead to information transmission under further restrictions on the environment (such as conditions on costs $c(s, \omega)$, on the payoffs and the fact that there is no private information about $b(\omega)$, i.e., the sender's value from trade is commonly known).

4.2 Verifiable Disclosure

Theorem 1 also contrasts with the seminal results of voluntary disclosure, Grossman and Hart (1980), Grossman (1981), and Milgrom (1981) who predict full disclosure or, when disclosure is costly, as in Verrechia (1983), disclosure above a certain threshold which threshold converges to full disclosure as this cost approaches zero. In our setting, the continuity with respect to this cost fails, and disclosure only becomes an equilibrium if it is entirely free.

Like in the case of signaling, there is no contradiction between our Theorem 1 and the prior literature: the stark difference in outcomes is driven by differences in the environments studied. Although the above disclosure literature allows for different evidence structures than we do, this difference is not substantive, as we already discussed in Examples 3 and 4. The substantive difference between GHM and our Theorem 1 and Proposition 3 lies in the allocation of private information and / or bargaining power, analogously to the case of signaling, as reflected in Table 2.

	no private info about $v(\omega)$	private info about $v(\omega)$
Receiver sets the price	cost over content	cost over content
Sender sets the price	cost over content	GHM

Table 2: Evidence Disclosure

Let us comment further on strategic differences between our environment and that of GHM. In their setting, there is always a sender type, the highest type, and an information *content*, full disclosure, that this type *strictly prefers* to the content of all other experiments, irrespective of the choices of other sender types. Fixing this content choice of the highest type, the next highest type of the sender also strictly prefers messages that fully disclose her value, etc. This drives their classic unraveling logic (including in the presence of costly disclosure).[17] In contrast, in our environment, the sender types with lowest value are *indifferent among all experiment contents* in equilibrium. Furthermore, the higher-value sender types might want to pool with the lower-value sender types, rather than separate from some other types, and their choices are not necessarily driven by the content of the experiments chosen by other types. In case all experiments have the same costs, any given experiment content can be chosen in equilibrium in our setting while only full disclosure in theirs. What drives choice (and selection) in our setup is cost not content.

4.3 Data Trade: Privacy and Platform Design

We now consider applications of our result to the problem of consumer data. We first discuss consumer's incentives to take costly actions to protect their data, then endogenize the buyer's information choice set, S and $c(s,\omega)$, by embedding our result into a simple market setting of data trade. We maintain the general assumptions of our base model of Section 2.

4.3.1 Privacy Paradox

The idea of a privacy paradox refers to the observation that, despite expressing significant concerns about the loss of their privacy, consumers appear not to take even minimally costly actions to protect their privacy and in many contexts, such as selecting browser specifications, overwhelmingly stick the default privacy settings. Empirical research has documented this general discrepancy between people's stated versus revealed preferences regarding privacy, e.g., Barth and de Jong (2017), Athey, Catalini, and Tucker (2017), Johnson, Shriver, and Du (2020). Firms appear to use online con-

[17] In Dye (1985) where disclosure is partial, disclosure is also driven by *content*, the favorability of the news: bad news is suppressed, good news is disclosed. Content may drive disclosure also in settings with costly uncertain evidence, e.g., DeMarzo, Kremer, and Skrzypacz (2019). For partial disclosure see also, e.g., Ben Porath, Dekel, and Lipman (2018) who show that disclosure via costless stochastic experiments is inefficient, Ali, Lewis and Vasserman (2023) who show that limited costless disclosure may improve upon mandated non disclosure, or Strausz (2017), Ben Porath, Dekel, and Lipman (2019) for mechanism design with evidence.

sumer data for price and search discrimination, e.g., Mikians et al. (2013), and the effectiveness of consumer advertising greatly depends on privacy laws, Goldfarb and Tucker (2011), or access to consumer data, Deisenroth et al. (2024). People appear to be generally concerned about the attempts of firms to collect, store and interpret information about them, but do not choose to protect their privacy even when they can do so at little cost.[18]

To explain this discrepancy, various authors have argued that behavioral factors, such as consumer inertia, preference for immediate gratification, choice overload, or miscalibration of probabilities, may well be at play; e.g., Acquisti, John, and Loewenstein (2013), Madarasz (2021), Fletcher et al. (2023). Others argue that data externalities may contribute to a lower willingness to pay for privacy, e.g., Acemoglu et al. (2022). Athey, Catalini, and Tucker (2017) point out that this discrepancy might reflect that people's stated preferences may not be reliable measures of their actual preferences. Bird and Neeman (2023) argue that the lack of privacy might actually be beneficial for consumers who express concern about it. Although all of these factors are important, our result has a robust implication in addressing this apparent paradox.

Our Theorem 1 shows that the reported lack of privacy protection is an equilibrium phenomenon even if consumers care about protecting their data.[19] In particular, if preserving privacy is not directly the cheapest option, consumers will never choose it in equilibrium. By Proposition 1, this is true even if the level of privacy protection is unobservable, and the seller only observes whether or not the buyer incurred some additional cost to engage in privacy protection. Similarly, by Proposition 4, it is also true even if there is no uniformly cheapest experiment, but there is an experiment that is always cheaper than privacy protection.

[18] 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 "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."

[19] As discussed before, a previous equilibrium explanation was provided by Acquisti, Taylor, and Wagman (2016) in the context of an example in which the seller elicits full disclosure by offering the buyer a discount for disclosing her value.

4.3.2 Data Trade and Platform Design

So far we kept the buyer's options for information choice as *exogenous* aspects of the environment. We now endogenize them by introducing a third party, a platform provider with bargaining power.[20] The uninformed platform provider has access to a finite set P of technologically feasible experiments (privacy designs). For simplicity, all experiments are equally costly for the provider to supply, and we normalize this cost to zero. The parties share a common prior. The timing is as follows.

First, 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. We interpret each s as a privacy design. The signal-generating process s determines what the seller can learn about the buyer. This can be interpreted as a function both of the buyer's behavior on the platform and a contract between the provider and the buyer specifying what information the provider may pass on to the seller, e.g., selective tracking. We impose *no* sign restriction on $c(s)$, which can then be a fee paid by the buyer to the platform or a subsidy paid by the platform to the buyer when the buyer chooses s.

Second, having committed to a menu $\{s, c(s)\}_{s\in S}$, the provider makes a take-it-or-leave-it price offer to the seller. If the seller accepts the offer, the game continues, and the seller gets access to the information that the buyer's choice will contractually allow. If the seller rejects the offer, the platform obtains a payoff of 0, while the seller and the buyer interact offline (see below).

Third, if the seller accepts the offer, the buyer can choose an item from $\{s, c(s)\}_{s\in S}$ or can choose to shop 'offline' directly from the seller, i.e., not use any of the privacy designs offered by the platform. In case the buyer chooses to shop offline, the seller only observes that the buyer chose to shop offline but no signal realizations. We normalize the buyer's direct cost of offline shopping to $c_{\text{off}} = 0$. In either case, the seller subsequently makes a price offer which the buyer can accept or reject.

If the buyer buys online, her payoff is $b(\omega)$ minus the payment to the platform provider (which can be negative) and the seller. If the buyer chooses the offline option,

[20] Note that if the seller controlled data trade, or had the bargaining power over the provider, the seller could always structure the costs and contents of the buyer's options, to his advantage. He could perfectly and freely 'screen' the buyer. Thus, the interesting case, the one where our cost-over-content force is at play, is where it is the provider who structures this trade and has bargaining power over the seller.

her payoff is $b(\omega)$ minus the payment to the seller, or if she never buys, her utility is normalized to zero. The seller's payoff is $v(\omega)$ plus the price paid by the buyer minus his payment to the provider. The provider's payoff equals the fee $c(s)$ paid by the buyer picking experiment s from $\{s, c(s)\}_{s \in S}$ plus the payment by the seller.

What experiments will the provider offer and for what fee or subsidy each? Despite the generality of the setup, we can make a tight prediction. To simplify its formulation, we say that an experiment is *relevant* in equilibrium if the buyer chooses it with a positive ex ante probability. The presence of privacy options that are not relevant has no impact on the equilibrium outcomes.

Proposition 5 *In any PBE,*

1. *The provider chooses $S \subseteq P$ that maximizes the seller's expected revenue in the continuation game.*

2. *The provider offers all relevant privacy options for free, $c(s) = 0$ for all $s \in S$.*[21]

3. *The buyer chooses one of the privacy options offered on the platform and never shops offline.*

The platform never charges nor subsidizes the buyer for her choice of information, instead the platform charges the seller and provides him the information which maximizes the seller's revenue. As a consequence, the seller's profits from selling to the buyer weakly increase as more experiments become feasible for the platform.

The proposition implies that *allowing for data trade* may lower the buyer's surplus in some environments, while it may increase the surplus in others. To see how data trade may lower buyer surplus, suppose values are private and compare the scenario where the buyer can only shop offline to the above scenario where data trade via a platform is also possible. In the former scenario, the buyer's ex ante expected surplus can be significant. In the latter scenario, which still allows the buyer to shop offline and preserve her privacy, her consumer surplus may instead be minimal (depending on P). Thus, the option to engage in data trade can hurt the buyer despite her having full property rights over her data and the platform having the bargaining power over

[21] More generally, in the absence of normalizing c_{off}, we obtain that $c(s) = c_{\text{off}}$ for all $s \in S$.

the seller. To see how data trade may increase the buyer's surplus, consider the buyer's and seller's values from the adverse-selection Example 6 from Section 3 and P consisting of the experiments from the first two bullet points of this example.

Proposition 5—which holds true if, following the buyer's choice of experiments, the seller and the buyer interact for any number of periods T, including $T = \infty$— also implies an economic force which is in contrast with a key aspect of the Coase conjecture, the value of being privately informed at the time of trade. Suppose that there is no data trade (e.g. P contains only one uninformative experiment). Comparing a buyer who enters object trade being privately informed about her valuation to one without such private information. The Coase conjecture implies that the former can achieve a much higher surplus. All else equal, possessing private information here is a *blessing* for the buyer. Suppose now that data trade precedes object trade. Compare a buyer entering data trade being privately informed about her valuation (interim data trade), to one entering data trade before obtaining private information and becomes informed right after data trade, that is, when object trade begins (ex ante data trade). Proposition 5 implies that it is now the latter who may achieve a much higher surplus. Possessing private information at the time of data trade becomes a *curse* for the buyer.[22]

5 History-Dependence

In the base model, the traders' valuations were stationary and the information choice problems were separable across time. We now relax these assumptions. We do so in the context of our base model of Section 2 as extended by Remark 1. Here, the buyer's experiment choice is still observed by the seller, but each experiment can generate, in each subsequent period, a private signal observed only by the buyer, a private signal observed only by the seller, and a public signal observed by both.

The traders' valuations, the costs of experiments, and the information choice sets are now allowed to depend on the history of the buyer's experiment choices and the

[22]Our analysis of data trade also contrasts with accounts of behavior-based price discrimination based on Coasian dynamics, e.g., Hart and Tirole (1988). Under classic price discrimination, given Coasian informational dynamics, the monopolist loses and consumers gain when the seller tracks prior purchasing decisions, e.g., Taylor (2004), Fudenberg and Villas Boas (2006).

public signal realizations from the chosen experiments. In any history h, we call the prior choices of experiments and the realizations of public signals from experiments along h the **public history of experiments** and denote it by h^P. We denote the set of all these public histories of experiments by $\mathcal{H}^P$. Different histories h and h' in $\mathcal{H}$ may give rise to the same public history of experiments $h^P = (h')^P$ in $\mathcal{H}^P$.

Recall the within-period timing in our setup. In each period t, first the buyer chooses an experiment, this is followed by the signal realizations of the chosen experiments, followed by the seller's offer. By h_t we denote *both* (i) the realized history up to the point of the buyer's response to the seller's offer at the end of period t and (ii) the realized history up to the point of the buyer's choice of experiment in the beginning of period $t+1$. This slight abuse of notation simplifies our terminology without any loss in precision because it is the buyer who moves at both of these histories, and the game only continues to period $t+1$ if the buyer rejects the seller's offer in period t. The realized public history of experiments h_t^P is then the same along both of these histories. We allow for the following history dependencies:

- The set of experiments the buyer is choosing from in period t is determined by the public history of experiments h_{t-1}^P. We denote by $S_t(h_{t-1}^P)$ the finite set of experiments the buyer chooses from in period t following h_{t-1}^P.

- Experiment costs are determined by ω and the public history of experiments. We denote by $c\left(s_t, \omega; h_{t-1}^P\right)$ the cost of choosing $s_t \in S_t(h_{t-1}^P)$ in the beginning of period t following h_{t-1}^P. We assume that, for any experiment s_t, its cost $c\left(s_t, \omega; h_{t-1}^P\right)$ is measurable with respect to (θ_b, h_{t-1}^P). The cost of experiment s_t for all ω consistent with (θ_b, h_{t-1}^P) is hence the same and we can denote it by $c\left(s_t, \theta_b; h_{t-1}^P\right)$.

- The buyer's value $b(\omega, h_t^P)$, at the time she accepts or rejects the seller's offer at the end of period t, is also determined by ω and h_t^P.[23] Analogously to the base model, we assume that $b(\omega, h_t^P)$ is measurable with respect to (θ_b, h_t^P), thus we can denote it by $b(\theta_b, h_t^P)$.

[23]Note that h_t^P does not contain the seller's past offers, hence the buyer's value is not affected by seller's price demands.

- The seller's value $v(\omega, h_t^P)$, at the time the seller makes a price offer in period t, is also determined by ω and h_t^P.

A consequence of history dependence is that choices which minimize the expected dynamic cost of experiments do not necessarily minimize the period-by-period costs. Instead, the minimum expected dynamic cost of choosing $s_t \in S_t(h_{t-1}^P)$ at history h_{t-1} is defined by:

$$\mathbb{E}_{(\theta_b, h_{t-1})} \inf_{s_{t+1} \in S_{t+1}\left(h_t^P\right), s_{t+2} \in S_{t+2}\left(h_{t+1}^P\right), \dots} \left[c\left(s_t, \theta_b; h_{t-1}^P\right) + e^{-r_b \Delta} c\left(s_{t+1}, \theta_b; h_t^P\right) + \dots\right]$$

where the expectation is taken over all $h_{t+\ell}$ (with $l \geq 0$) histories that are consistent with the previous experiment choices and *along which the buyer always rejects the seller's offers*. In particular, the minimum expected dynamic cost is calculated under the assumption that the game continues till the final period T. This is effectively the value of the type-dependent stochastic outside option that the buyer can unilaterally achieve by never buying.

Although this minimum expected dynamic cost could in principle depend on s_t, θ_b and h_{t-1}, we restrict the dependence on history h_{t-1} to its public history of experiments component h_{t-1}^P, that is, we impose:

Assumption 0. The minimum expected dynamic cost of choosing $s_t \in S_t(h_{t-1}^P)$ at history h_{t-1} is always determined by θ_b and h_{t-1}^P for any t.

Given Assumption 0, we can denote the minimum expected dynamic cost of choosing s_t, following history h_{t-1}, by $C_t\left(s_t, \theta_b; h_{t-1}^P\right)$. We call an experiment choice $s_t \in S_t(h_{t-1}^P)$ **dynamically cheapest** if it minimizes $C_t\left(s_t, \theta_b; h_{t-1}^P\right)$ among all experiments in $S_t(h_{t-1}^P)$. Being finite, each $S_t(h_{t-1}^P)$ contains at least one dynamically cheapest experiment.

The **minimum expected dynamic cost of experiments** at the end of period t, at history h_t^P, is then given by:

$$C_t\left(\theta_b; h_t^P\right) \equiv e^{-r_b \Delta} \min_{s_{t+1} \in S_{t+1}(h_t^P)} C_{t+1}\left(s_{t+1}, \theta_b; h_t^P\right),$$

since the minimum dynamic cost presupposes that the buyer always rejects seller's price offers. We extend this last definition to the start of the game,

$$C_0\left(\theta_b;\emptyset\right)\equiv e^{-r_b\Delta}\min_{s_1\in S_1(\emptyset)}C_1\left(s_1,\theta_b;\emptyset\right).$$

We adopt the following assumptions for any h_t^P and any $\theta_b \in \Theta_b$:

Assumption 1. $\arg\min_{s_t\in S_t\left(h_{t-1}^P\right)} C_t\left(s_t,\theta_b;h_{t-1}^P\right)$ does not depend on θ_b.

Assumption 2. The mappings $\Theta_b \ni \theta_b \to C_t\left(s,\theta_b;h_{t-1}^P\right)$ and $\Theta_b \ni \theta_b \to b\left(\theta_b,h_t^P\right)$ are continuous in θ_b.

Assumption 3. The difference $b\left(\theta_b,h_t^P\right)-C_t\left(\theta_b;h_t^P\right)$ (the value net of avoided future costs) increases in $b\left(\theta_b,\emptyset\right)-C_1\left(\theta_b;\emptyset\right)$.

The environment of Theorem 1 satisfies these assumptions. Indeed, even in the absence of history-dependence, the above assumptions are more permissive than those we imposed in Section 2, e.g., the cost of cheapest experiments can depend on the buyer's privately known type. Other special cases of interest include the following.

Dynamic Commitments. Suppose that only choice sets $S_t(h_{t-1}^P)$ are history dependent, costs and valuations are stationary. This entails environments where the buyer can make dynamic commitments over information choices. For example, she may have a choice in S_1 that restricts her future information choice sets (by ensuring that they are singletons). The ability to commit may be critical here because the buyer might want to avoid suboptimal information choices in the future. Our setting now also captures such commitment options as well.[24]

Buyer Learning about her Value. Another special case is when only the buyer's value $b(\theta_b,h_t^P)$ depends on the realization of the public signals. For instance,

[24] More flexible commitments can be modeled through investments in future experiment costs. Suppose that the cost function $c(s_t;\theta_b,h_{t-1}^P)$ is history dependent, but valuations and the choice sets are not. The buyer's choice from S_t is then an investment which decreases or increases the costs of different experiments. The buyer might want to decrease the cost of a future experiment if she subsequently chooses it or if the cost change decreases the equilibrium prices. The buyer might want to increase the cost of future experiments as a way to ensure she will not choose them. Recall also our introductory example of an employee taking a course; when such course impacts the employee's outside options, this impact is reflected in our model as the change in the costs of future experiments.

one experiment $s_t \in S_t(h^P_{t-1})$ might imply that the value equals 1 at all continuation histories, while another experiment s'_t might, with probability $\frac{1}{2}$, generate public signals following which the value is 2 and, with probability $\frac{1}{2}$, public signals following which the value is 0. The buyer's choice from S_t can then be interpreted as a choice what to learn about her value.[25]

Investment in Value. Another special case is when only the players' valuations are history dependent and depend only on the buyer's *public* experiment choices. The buyer's choice from S_t can be understood as a public investment into her or the seller's valuation of the object. The impact of such investment must satisfy our monotonicity and continuity assumptions, but can otherwise be type-dependent. For example, the buyer's valuation at any history h_t may be given by $b\left(\theta_b, h^P_t\right) = b_1\left(\theta_b, \emptyset\right) + b_2(h^P_t)$ where b_1 and b_2 are arbitrary functions and the former is continuous in the buyer's type. An instance of this is where the buyer's education decision is both a signal and is also productive in that it changes both players' values.

The cost-over-content insight of Theorem 1 remains true in the above general class of environments. To formulate the analogue of its first part, we define a **public experiment schedule** to be any mapping σ from the set of public experiment histories $\mathcal{H}^P$ to the set of experiments $\cup_{h^P_{t-1}\in\mathcal{H}^P} S_t(h^P_{t-1})$ such that, for any $h^P_{t-1} \in \mathcal{H}^P$ we have $\sigma(h^P_{t-1}) \in S_t(h^P_{t-1})$.[26] The public experiment schedule σ induces an experiment choice $\sigma(h^P_{t-1})$ at any history h_{t-1} at which the buyer chooses an experiment. We say that a public experiment schedule σ consists of dynamically cheapest experiments if for each h_{t-1} at which the buyer chooses an experiment, the experiment $\sigma(h^P_{t-1})$ is dynamically cheapest.[27]

Theorem 2 *[History Dependence]*

[25]In our setup, while the buyer's valuation may always be privately known by the buyer, such learning *per se* is always *public*. Ravid, Roesler, and Szentes (2022) instead study an uninformed buyer's choice as to what to *privately* learn about her own value assuming the cost of this learning increases in the informativeness of the chosen experiment. They show that as such costs uniformly vanish, the players' payoffs converge to the lowest possible payoffs under costless experiments.

[26]Note that the time subscript t in S_t is determined by the history h_{t-1}.

[27]By our assumptions the same experiments are dynamically cheapest for all buyer types.

Existence: *For any experiment schedule σ consisting of dynamically cheapest experiments, there exists a PBE such that all buyer types pool on $\sigma(h^P_{t-1})$ at any on-path history h_{t-1} at which the buyer chooses an experiment.*

Cost over content: *In no PBE does the buyer ever choose a non-dynamically-cheapest experiment on the equilibrium path.*

The first part of the above result establishes existence for any cheapest public experiment schedule. The second and main part of Theorem 2 imposes *no* restriction on the buyer's strategy; it allows for all PBEs including mixed-strategy ones and ones in which the players condition their choices on *both* the public and private information they receive. This second part of Theorem 2 establishes that the buyer always chooses a dynamically cheapest experiment: rather than choosing the experiment that is cheapest in a given period per se, the buyer takes into account the impact on future expected experiment costs.

Theorem 2 implies that our cost-over-content insight remains valid also in environments where the buyer's choice publicly impacts her future choice sets, the traders' valuations, and the costs of experiments as well. For example, consider a context where the buyer's experiment choice privately impacts the seller's information and publicly her future experiment choice sets. If the unique dynamically cheapest experiment fully restricts her future choices, she will choose this experiment, effectively committing to a given information choice over time. At the same time, if such a commitment is dynamically costlier than full flexibility, she will never restrict her future choice sets.

Remark 6: Observability of Choice. Theorem 1 remained valid under minimal assumptions on the observability of the buyer's actions (Proposition 1). Similarly, the validity (and proof) of Theorem 2 only requires analogous minimal observability: it is sufficient (i) that the seller observes whether or not the buyer has chosen a dynamically cheapest experiment from each $S_t(h^P_{t-1})$, and (ii) that both the buyer and the seller observe those experiment choices and signal realizations that change future values, costs, or choice sets. As long as a chosen experiment does not impact these objects, for Theorem 2 to hold, the seller does not need to observe the identity of the chosen experiment only whether it was dynamically cheapest or not.[28]

Theorem 2 allows for the joint presence of information and investment choice. Minimal observability is sufficient for the former, but *not* for the latter. In fact, a special case of the setup is where the buyer's choice is purely an observable investment choice with observable returns in her valuation $b(h_t^P, \omega)$. This is the domain of the classic hold-up problem, e.g., Williamson (1979), Grossman and Hart (1986), and Gul (2001). Here, the classic insight is that the buyer under-invests into her valuation relative to the socially efficient level. Our Theorem 2 implies general conditions under which the buyer would not undertake any costly *public* investment into her valuation as long as this was (dynamically) costly. In contrast to pure information choice, however, here full observability is key.

Remark 7: Costly Information versus Costly Investment. For Theorem 1 to hold, the seller need *not observe* which non-cheapest experiment was chosen by the buyer *nor* any signal the buyer receives (for example, about the seller's value or his belief of her value, etc.) as a result of her choice. In contrast, the public observability of both investment choice and the return from such investment are central to the hold-up logic.

Unobservable investment. Gul (2001) shows that the lack of observability of the buyer's investment choice by the seller often fully overturns the hold-up problem. He considers the buyer's investment choice into her valuation that is *not* observed by the seller. The time horizon is $T = \infty$ and the parties have private values. He shows not only that the buyer undertakes costly investment to increase her valuation, but that as bargaining becomes smooth, $\Delta \to 0$, her investment choice becomes fully efficient.[29]

Unobservable return from investment. Even when the buyer's investment choice is itself observable, the lack of public observability of the return from her

[28]An analogue of Proposition 2 also remains true. With all observability assumptions relaxed, it is still true that the buyer never makes information structure choices that the seller can identify as not being cheapest; that is, in the formalism of Proposition 2, the buyer never chooses experiments from any observability partition cell that does not contain a cheapest experiment.

[29]Relatedly, Gibbons (1992) points out that PBEs in the hold-up problem with unobservable investment might require mixed investment strategies.

investment may well overturn or soften the hold-up problem. The following examples illustrate this. Suppose the buyer has a binary investment choice observable by the seller, but let the return from this investment be observed by the buyer only. The buyer initially makes a single investment choice; $S_1 = \{\text{invest, don't invest}\}$ with $c(\text{don't invest}) = 0$ and $c(\text{invest}) \in \mathbb{R}$. 'No investment' leaves the buyer's value at b_0, and 'Investment' stochastically increases it to some $\hat{b}_1$ distributed over the interval $[d, e]$. While the seller observes the investment decision, the realized return $\hat{b}_1$ is observed only by the buyer after her choice of investment. The buyer will then often choose 'Investment' in equilibrium. This is already true when $T = 1$ and also when $T > 1$. Indeed, as Gul (2001) also argues, when $T = \infty$ and $\Delta \to 0$, the buyer's observable investment choice again often becomes socially efficient. In these examples initially there is only a single buyer type. If instead there are multiple initial buyer types with different unobservable returns from costly investment, some may invest while others may not.[30]

In contrast, when the buyer's choice is purely about information, we showed that she never chooses a non-cheapest experiment under minimal observability. Whether the seller observes what non-cheapest experiment was chosen (or its realizations) does not matter for our Theorem 1. As Proposition 2 highlights, buyer types never choose non-cheapest information in our setup because in equilibrium they never want to separate themselves from other types by choosing more costly experiments. This "no separation" logic has no direct counterpart in hold up.

6 Efficient Bargaining under Endogenous Information

In the analysis above, we considered information choice in a general class of environments where price formation was the result of an offer made by one of the trading parties. Such price formation is the norm in many settings, but one may ask what happens when considering other forms of price formation such as, e.g., double auctions. In this Section we consider a broad class of bargaining protocols and show that our insights on efficiency being driven by informational defaults extends. Even though under some bargaining protocols a privately-informed party might not always

[30]For unobservable investment outcomes, see also Hermalin and Katz (2009) and Halac (2015).

choose the cheapest experiment, we show that nevertheless the efficiency of trade still hinges only on the properties of the cheapest experiments.

We revisit the seminal analysis of Myerson and Satthertwaite (1983), but allow for endogenous information choice. Like them, we focus on traders with private values which are drawn independently. We further assume that the traders' values are drawn from a common convex support which we normalize to be $[0, 1]$. For ease of exposition, we assume that traders chose experiments only once and then participate in some fixed trading mechanism. The timing is as follows.

First, the traders simultaneously select an experiment each, $s_b \in S_b$ and $s_v \in S_v$ respectively. Choosing each experiment is associated with some cost $c(s_b, \omega)$ for the buyer when she chooses s_b and $c(s_v, \omega)$ for the seller when he chooses s_v. Experimental costs are measurable and continuous in the respective player's type (as assumed before for the buyer). For simplicity, we assume that there is a unique cheapest experiment for each trader. Having chosen the experiments, the parties observe each other choices and the results of all chosen experiments.[31]

Second, both parties report their values to a direct mechanism $\varphi : [0, 1] \times [0, 1] \to \Delta([0, 1] \times \mathbb{R})$ that maps the pair of reports to a distribution over outcomes; each outcome $(\pi, p) \in [0, 1] \times \mathbb{R}$ states the probability π that the buyer receives the good and the transfer p from the buyer to the seller.[32] Analogously to the standard revelation principle argument, our focus on direct mechanisms is without loss of generality. We further restrict attention to mechanisms that assign the object to the trader who reported a higher value; in case of a tie, we assign the object to the buyer. We assume that the mechanism is ex post individually rational with respect to *reported values*, that is: the buyer only pays a positive amount when the buyer receives the good and the price the buyer then pays is weakly below the reported buyer's value; similarly, the seller never pays to the buyer and when the seller gives the good to the buyer the price the seller receives is weakly above the seller's reported value. We use the term *bargaining protocol* to refer to any φ satisfying the above properties. To illustrate, note that the classic double auction of Chatterjee and Samuelson (1983)

[31] It is sufficient that each party observes whether the other party chose the cheapest experiment (default) and observes the outcome of the experiment chosen by the other party.

[32] What the buyer pays is equal to what the seller receives, and hence φ is budget balanced.

satisfies these properties, and hence it is a bargaining protocol.[33]

We say that trade is efficient when the buyer receives the good iff $b(\omega) \geq v(\omega)$. The posterior full-support assumption below is satisfied, e.g., when no-disclosure is the cheapest experiment for at least one of the traders.

Theorem 3 ***[Efficient Bargaining]*** *Suppose that at least one of the trader's cheapest experiment has full-support posteriors. If a PBE implements efficient trade, then traders choose the cheapest experiments. Furthermore, there is always a PBE where players choose the cheapest experiments.*

The endogenous information choice can hence only lead to efficient trade if the traders choose the cheapest (default) information structure. In other words, the efficiency of trade is determined purely by what the cheapest information structure is. In our setting each trader type can choose an information structure which is different from the cheapest one and each information structure (including the cheapest ones) may contain hard evidence. Even if the default information structure does not allow for efficient trade, there may well be information structures which, if exogenously imposed, allow for efficient trade. One can then imagine that the parties may coordinate on such an information structure as long as it is not too costly to do so. Theorem 3 shows, however, that unless such an information structure is directly the cheapest one, the traders will never choose it. In turn, the overall efficiency loss here can be much greater than the (arbitrarily small) direct cost of moving from an inefficient information structure to one under which there is efficient trade.

Our proof hinges on the revenue equivalence theorem of Myerson (1981). A rough intuition is as follows. Suppose that it is the seller whose cheapest experiment leaves the buyer with full-support posteriors. In turn, unless some seller types choose a more expensive experiment, the revenue equivalence allows us to show that efficiency is only possible if the seller captures the entire surplus from trade. Our proof uses this property to show that if the equilibrium implements efficient trade then both parties must pool on their respective cheapest experiments as separation and efficiency

[33] The double auction assigns the object to the trader with higher reported value and if the object changes hands, the buyer pays a weighted average of the reported values. In the special case where the weight on the seller's reported value is null, the auction is equivalent to the buyer making a take-it-or-leave-it offer. If the weight on the buyer's reported value is null, the auction is equivalent to the seller making a take-it-or-leave-it offer.

become incompatible when information is costly. We can complement the above theorem and further show that under endogenous information choice efficient trade under any bargaining protocol is possible if and only if it is possible under all bargaining protocols.

Proposition 6 *Suppose that the cheapest experiment of at least one of the traders has full-support posteriors. There exists an equilibrium which implements efficient trade under any given bargaining protocol φ if and only if such an equilibrium exists under all bargaining protocols.*

When traders pick information endogenously, the efficiency of trade does not depend on the bargaining protocol per se. In turn, efficient bargaining under a given bargaining protocol is possible if and only if it is possible under the protocol where the seller sets the price. The proof uses the previously demonstrated fact that efficient trade depends only on the properties of the cheapest experiments. Given this fact we then show that if trade can be efficient under a given bargaining protocol, then the information structure supporting efficient trade under this bargaining protocol supports efficiency in all bargaining protocols.

7 Conclusion

Our paper considered the information choice problem of a privately informed buyer (sender) who influences the information of the seller (receiver) and herself before and/or during their trade. We establish that while trade outcomes critically depend on the content of information each party has, or learns during trade, the sender's choice of the receiver's information is determined purely by the cost of supplying this information. Our logic points to such costs being a direct source of market failure and also to the potential role of easy-to-change *informational defaults* in consumer protection, the regulation of selection markets, or the efficiency of bilateral trade.

The analysis suggests broad comparative static predictions. For example, the positive impact of education on earnings has long been associated with two distinct forces: human capital accumulation and equilibrium signaling. At the same time, the evidence on signaling is mixed and, e.g., Clark and Martorell (2014) find that when controlling for observation structures, obtaining a high school diploma is often *not* associated with any significant signaling motive. Our paper provides a robust rationale for why such signaling may play a comparatively smaller role in some settings. The more wage setting power an employer has, the lesser the equilibrium incentives for costly signaling via education may be. Indeed, if the labor market for minorities or women has fewer opportunities giving an employer greater power to set wages, signaling shall be a smaller force for them, a prediction which is consistent with the findings of Tyler, Murnane, and Willet (2000). Similarly, while the supply of different warranty lengths has been associated with costly signaling about product quality, Spence (1977), in our setup, such information provision is not an equilibrium and warranty length per se shall not serve as a means of information provision. Empirical findings are consistent with this prediction. For instance, in the market for computer servers where receivers are institutions with considerable bargaining power, Chu and Chintagunta (2011) find that quality and warranty length are uncorrelated. Instead the main driver of the supply of warranties appears to be risk-sharing not signaling.[34]

The analysis may help shed light on a 'no-disclosure puzzle.' Dranove and Jin (2010) describe the widespread lack of voluntary disclosure "even in markets with credible, low-cost mechanisms to disclose." Our result points to a robust equilibrium rationale for a lack of voluntary disclosure even when it is clear that senders have evidence, such as in the case of corporate records, mortality data, or the strategic grade non-disclosure of MBA students. If no disclosure is the directly cheapest or default option, then our results predict that senders may adapt this default. This prediction may help explain observations that information regulations, such as mandatory hygiene certifications, have significant impact on market outcomes; see, e.g., Jin and Leslie (2003). Our results predict that minimal disclosure requirements, such as those employed in asset or healthcare markets, will be sticky in that irrespective of the sender's private information and the opportunities to engage in more disclosure, a sender may follow this minimal requirement whatever its content may be.

Our results contribute to the regulatory discussions on trade via platforms, such as

[34]Such lack of signaling could also be due to consumers already being informed, however, consistent with our prediction, the authors find that new firms offer the same warranty duration initially as more established firms and do not change this once they become more established.

the EU's Digital Market Act or the Ending Platform Monopolies Act of 2021 in the US, e.g., Fletcher et al. (2023). While sellers, advertisers, and platforms derive significant benefits from the use of consumer data, consumers exhibit a very low willingness to pay for active privacy protection and overwhelmingly stick to default privacy settings, despite being concerned about their loss of privacy; e.g., Acquisti, Brandimarte, and Loewenstein (2015), Athey, Catalini, and Tucker (2017), CMA (2022). While it is common to explain this by consumer inertia or 'choice overload,' e.g., Fletcher et al. (2023), our cost-over-content theorem identifies this as a robust aspect of fully responsive equilibrium behavior. Our cost-over-content logic might then help inform privacy regulations, such as EU's GDPR. It suggests that active regulation of the content which is cheapest (easiest) to choose may be more consequential than simplifying choice or giving consumers more control over their data.

In particular, our logic implies that simply ensuring that buyers have full property rights over their data and that platforms and sellers are separate entities with the former having substantial bargaining power over the latter may not be effective tools of consumer protection in the context of data trade; direct regulation of the kind of data that can be collected, and the extent to which they can be exclusively purchased by firms, may be more effective.[35] Future research can explore the interaction of this force with competition between platforms and / or sellers, and in other contexts, such as information choice by bidders in auctions. It can also explore the extent to which our results would continue to hold when relaxing well-calibrated equilibrium beliefs allowing for miscalibrated expectations about the beliefs of others, or the link between others' beliefs and their actions.

References

Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2022). Too Much Data: Prices and Inefficiencies in Data Markets. *American Economic Journal: Microeconomics*, 14 (4): 218–56.

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

[35] For a more extensive discussion, see Madarasz and Pycia (2020).

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

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

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

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

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*, 34(7): 1038-1058.

Ben-Porath, E., E. Dekel, and B. Lipman (2018). Disclosure and Choice. *Review of Economic Studies*, 85(3): 1471-1501.

Ben-Porath, E., E. Dekel, and B. Lipman (2019). Mechanisms With Evidence: Commitment and Robustness. *Econometrica*, 87: 529-566

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

Bird, D. and Z. Neeman (2023). The Effect of Privacy on Market Structure and Prices. *The Journal of Law, Economics, and Organization*, ewad031.

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

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

Calzolari, G. and A. Pavan (2006b). Monopoly with resale. *The RAND Journal of Economics*, 37: 362-375.

Chatterjee, K. and W. Samuelson (1983). Bargaining under Incomplete Information. *Operations Research*, 31(5): 835–851.

Chu, J. and P. Chintagunta (2011). An Empirical Test of Warranty Theories in the U.S. Computer Server and Automobile Markets. *Journal of Marketing*, 75(2): 75-92.

Clark, D. and P. Martorell (2014). The Signaling Value of a High School Diploma. *Journal of Political Economy*, 122(2): 282–318.

Coase, R. (1972). Durability and Monopoly. *Journal of Law and Economics*, 15(1): 143–149.

Competition and Markets Authority, UK. (2022). Mobile ecosystems Market Study: Final report. https://www.gov.uk/cma-cases/mobile-ecosystems-market-study

Crawford, V. and J. Sobel (1982). Strategic Information Transmission. *Econometrica*, 50(6): 1431-1451.

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

Deisenroth, D., U. Manjeer, Z. Sohail, S. Tadelis, and N. Wernerfelt (2024). Digital Advertising and Market Structure: Implications for Privacy Regulation. NBER WP 32726.

DeMarzo, P., I. Kremer, and A. Skrzypacz (2019). Test Design and Minimum Standards. *American Economic Review*, 109 (6): 2173–2207.

Doval, L. and V. Skreta (2024). Constrained information design. *Mathematics of Operations Research* 49(1): 78-106.

Dranove, D. and G. Jin (2010). Quality Disclosure and Certification: Theory and Practice. *Journal of Economic Literature*, 48 (4): 935-63.

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

Dwork, C. and A. Roth (2014). The Algorithmic Foundations of Differential Privacy. Foundations and Trendsin Theoretical Computer Science, 9(3–4): 211–407.

Einav, L. and A. Finkelstein (2011). Selection in Insurance Markets: Theory and Evidence in Pictures. *Journal of Economic Perspectives*, 25 (1): 115–138.

Fletcher, A., G. Crawford, J. Crémer, D. Dinielli, P. Heidhues, M. Luca, T. Salz, M. Schnitzer, F. Scott Morton, K. Seim, and M. Sinkinson (2023). Consumer protection for online markets and large digital platforms. *Yale Journal on Regulation,*40:875-914.

Fudenberg, D. and M. Villas-Boas (2005). Behavior-Based Price Discrimination and

Customer Recognition. *Handbook on Economics and Information Systems.* Elsevier.

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

Gibbons, R. (1992). *Game Theory for Applied Economists.* PUP Princeton.

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

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

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

Grossman, S. and O. Hart (1986). The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of Political Economy*, 94: 691–719.

Gul, F., H. Sonnenschein, and R. Wilson (1986)."Foundations of Dynamic Monopoly and the Coase Conjecture."*Journal of Economic Theory.* 39(1): 155 - 90.

Gul, F. (2001). Unobservable Investment and the Hold-Up Problem. *Econometrica*, 69(2): 343–376.

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

Halac, M. (2015). Investing in a relationship. *RAND Journal of Economics*, 46(1): 165–185.

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

Hermalin, B. and M. Katz (2009). Information and the Hold-up Problem.*RAND Journal of Economics*, 40(3): 405-423.

Horton, J., R. Johari, and P. Kircher (2024). Sorting Through Cheap Talk: Theory and Evidence from a Labor Market. LIDAM DP 2024/13.

Jin, G. and P. Leslie (2003). The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards. *Quarterly Journal of Economics*, 118(2): 409–451.

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.

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

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

Koessler, F. and V. Skreta (2023). Informed Information Design. *Journal of Political Economy*, 131(11): 3186–3232.

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.

Kreps, D. and R. Wilson (1982). Sequential Equilibria. *Econometrica*, 50(4): 863–894.

Loertscher, S., L. Marx, and T. Wilkening (2015). A long way coming: Designing centralized markets with privately informed buyers and sellers. *Journal of Economic Literature* 53(4): 857-897.

Madarasz, K. and M. Pycia (2020). Towards a Resolution of the Privacy Paradox. Working Paper. *Bepress Archive*, https://works.bepress.com/kristof_madarasz/47/.

Madarasz, K. and M. Pycia (2023). Information Choice: Cost over Content. *CEPR Discussion Paper DP18252*, https://cepr.org/publications/dp18252/.

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

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.

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

Milgrom, P. (2004). *Putting Auction Theory to Work*. CUP, Cambridge.

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

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.

Myerson, R. (1983). Mechanism Design by an Informed Principal. *Econometrica*, 51(6): 1767–97.

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

Pavan, A. and J. Tirole (2024). Knowing your Lemon before you Dump It. WP.

Ravid, D., A.-K. Roesler, and B. Szentes (2022). Learning before trading: On the inefficiency of ignoring free information. *Journal of Political Economy*, 130: 346-387.

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

Seidmann, D. and E. Winter (1997). Strategic Information Transmission with Verifiable Messages. *Econometrica*, 65(1): 163–169.

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

Spence, M. (1977). Consumer Misperceptions, Product Failure and Producer Liability. *Review of Economic Studies*, 44(3): 561–572.

Strausz, R. (2017). Mechanism Design with Partially Verifiable Information. WP.

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

Thaler, R. and C. Sunstein (2008). Nudge. Yale University Press.

Tirole, J. (2012). Overcoming Adverse Selection: How Public Intervention Can Restore Market Functioning. *American Economic Review* 102 (1): 29–59.

Tyler, J., Murnane, R., and J. Willett (2000). Estimating the Labor Market Signaling Value of the GED. *Quarterly Journal of Economics*, 115(2): 431-468.

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

Online Appendix for: "Cost over Content: Information Choice in Trade"

Kristóf Madarász (LSE) and Marek Pycia (U Zurich)

Appendix A: Proofs

In the proofs below, we simplify notation and denote the buyer's discounting by δ, that is $\delta \equiv e^{-r_b \Delta}$.

Proof of Theorems 1 and 2

We focus on Theorem 2 as Theorem 1 is its special case. First, we fix a dynamically cheapest schedule of experiments σ and show that there is an equilibrium in which all buyer types pool on $\sigma(h_{t-1}^P)$ when choosing an experiment at any history h_{t-1}. Buying the object at period-t history h_t not only gives the buyer its value but also 'saves' the buyer the cost of future experiment choices. Accordingly, we denote by $B_t\left(\theta_b; h_t^P\right) = b\left(\theta_b; h_t^P\right) - C_t\left(\theta_b; h_t^P\right)$ the "full value" of the purchase at h_t. Note that this value is the same at the moment the seller makes an offer in period t and at the moment the buyer accepts or rejects this offer in period t.[36]

On path of the equilibrium we are constructing all buyer types pool on $\sigma(h_{t-1}^P)$ for each h_t. We can also assume that all buyer types pool on dynamically cheapest experiments off path. This buyer behavior is supported by the seller's belief that seeing at some history h_t, some 0-probability deviation to a more expensive experiment the seller attributes this deviation to a buyer type whose dynamic full value is equal to the supremum over the set of possible full values of types at the history where the seller sets the price. Note that our compactness and continuity assumptions imply that such a type exists at any history and our monotonicity assumption implies that the set of such types is the same at all histories. We can thus pick one of such types,

[36] An analogue of Theorem 2 remains true in the variation of the model we study in which the costs of experiments at subsequent times $t+1, t+2, ...$ are incurred whether or not the buyer buys the good at time t. The full value of the purchase is then simply $B_t\left(\theta_b; h_t^P\right)$ and, with this adjustment, the same argument delivers the result.

let's denote it by $\theta_{max} \in \Theta_b$, and have the seller attribute all 0-probability deviations to this type.[37] Given such beliefs, the seller's best response is to set, at any time $\tau \geq t$, the price that is weakly higher than both the full value of buyer type $\theta_{\max}$ and the seller's value.[38] Thus, the buyer does not gain from any deviation.

It remains to show that all buyer types choose the dynamically cheapest experiment at each history h_{t-1} on path of every equilibrium.[39] By way of contradiction, suppose that there is an equilibrium in which some buyer types, with positive probability, choose some experiment that is not dynamically cheapest. Let history $h_{t'-1}$, at the beginning of period t', be the earliest period at which this happens.

By $w(\theta_b, h_{t'-1})$, we denote the strictly positive lower bound on the additional dynamic cost type θ_b incurs by choosing a non-cheapest experiment at $h_{t'-1}$ rather than a cheapest one; that is $w(\theta_b, h_{t'-1})$ is set to equal:

$$\min_{s'_t \notin argmin_{s_{t'} \in S_{t'}(h^P_{t'-1})} C_{t'}\left(s_{t'}, \theta_b; h^P_{t'-1}\right)} C_{t'}\left(s'_{t'}, \theta_b; h^P_{t'-1}\right) - \min_{s_{t'} \in S_{t'}(h^P_{t'-1})} C_{t'}\left(s_{t'}, \theta_b; h^P_{t'-1}\right).$$

Recall that Θ_b is compact, the dynamic cost of each experiment is measurable and continuous in $\theta_b \in \Theta_b$, there is a finite number of experiments at history $h_{t'-1}$, and, by assumption, the same experiments are dynamically cheapest for all $\theta_b \in \Theta_b$. Hence, there is a strictly positive lower bound $w(h_{t'-1}) > 0$ such that $w\left(\theta_b, h_{t'-1}\right) \geq w(h_{t'-1}) > 0$ for all $\theta_b \in \Theta_b$.

Let $\underline{B}_t(h^P_t)$ denote the infimum of dynamic full values of buyer types (at the time the seller makes the offer at time $t \geq t'$) that, with positive probability, chose a non-dynamically cheapest experiment at time t'. The price $p_t(h_t)$ charged by the seller at some history h_t in any period $t \geq t'$ following the seller observing a non-cheapest experiment choice at $h_{t'}$ is then bounded from below by $\underline{B}_t(h^P_t)$ where h^P_t is the

[37]In Appendix B, we discuss how this step of the construction can be modified to ensure that the weak PBE we are now constructing satisfies sequential rationality and other refinements.

[38]The value for the seller lies in avoiding the costs of providing the product to the buyer. Note that we allow the seller's value to evolve over time and, in particular, the seller might set a price so high so as to avoid trade at time τ if the seller expects their value to go down in the future. 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.

[39]The same argument also establishes that there is no PBE in which, following a prior deviation by either of the players, the buyer ever chooses a not dynamically cheapest experiment.

public history of experiments at h_t. Indeed, by way of contradiction suppose that this bound fails at time t' or at a later time. Let R be the supremum of trading rents $\underline{B}_t(h_t^P) - p_t(h_t)$ taken over all periods $t \geq t'$ and all on-path continuation histories of $h_{t'}$; this supremum exists because the values and costs are bounded. Furthermore, the failure of the above bound implies that $R > 0$.

By definition of R, there is some time $t'' \geq t'$ and history $h_{t''}$ in the continuation game following $h_{t'}$, such that $\underline{B}_{t''}(h_{t''}^P) - p_{t''}(h_{t''}) > \frac{1+\delta}{2}R$. All on-path types θ_b would buy at this history at price $p_{t''}(h_{t''})$. Indeed, buying at this price would give the buyer the utility (evaluated from the perspective of time t'') of at least $C_{t''}(\theta_b; h_{t''}^P)$ plus the trading rent of at least $\frac{1+\delta}{2}R$. The expected utility from postponing the purchase and either never buying or buying at some continuation history $h_{t'''} \supsetneq h_{t''}$ at some time $t''' > t''$ also consists of the expected dynamic costs of future experiments and the trading rent of 0 (if never buying) or at most $\delta^{t'''-t''}(\underline{B}_{t'''}(h_{t'''}^P) - p_{t'''}(h_{t'''})) \leq \delta^{t'''-t''}R \leq \delta R$ (if buying at time t'''). By definition of $C_{t''}(\theta_b; h_{t''}^P)$ and the law of iterated expectations, the expected dynamic cost of future experiments is equal to $C_{t''+1}(\theta_b; h_{t''}^P)$. Thus, in the conjectured equilibrium, at history $h_{t''}$ the immediate purchase at price $p_{t''}(h_{t''})$ leads to utility at least $\frac{1-\delta}{2}R$ higher than any other course of actions by the buyer. Hence, in the equilibrium, all buyer types on path at $h_{t''}$ would buy at any price strictly lower than $p_{t''}(h_{t''}) + \frac{1-\delta}{2}R$. The seller then has a profitable deviation of slightly increasing the price at history $h_{t''}$. As such profitable deviations are not possible in a PBE, this is a contradiction that shows that $p_t(h_t) \geq \underline{B}_t(h_t^P)$ at every history following the choice of non-cheapest experiment at time t'.

In the equilibrium, the expected continuation utility of type θ_b from choosing some non-cheapest experiment $s'_{t'}$ at history $h_{t'-1}$ is hence bounded from above by $-C_{t'}(s'_{t'}, \theta_b; h_{t'-1}^P)$ plus the type's continuation equilibrium expectation of $B_t\left(\theta_b, h_t^P\right) - \underline{B}_t(h_t^P)$ where $h_t \supset h_{t'-1}$ is the (possibly stochastic) history at which, in the continuation equilibrium, this type buys the good. Then, types choosing $s'_{t'}$ with full values at history $h_{t'-1}$ sufficiently close to $\underline{B}_{t'-1}(h_{t'-1}^P)$ would strictly benefit from a deviation to a dynamically cheapest experiment. Indeed, take any $\epsilon \in (0, w(h_{t'-1}))$. If the game length T is finite, then our monotonicity and continuity assumptions ensure that types with full values sufficiently close to $\underline{B}_{t'-1}(h_{t'-1}^P)$ at $h_{t'-1}$ are within ϵ of $\underline{B}_t(h_t^P)$ at any continuation history $h_t \supseteq h_{t'-1}$ for any $t = t', ..., T$. If T is infinite, then take

$\hat{T}$ large enough so that $\delta^{\hat{T}-t'}$ times the supremum of possible differences of full values is smaller than ϵ. Then, our monotonicity and continuity assumptions ensure that for types with full values sufficiently close to $\underline{B}_{t'-1}(h^P_{t'-1})$ at history $h_{t'-1}$ are within ϵ of $\underline{B}_t(h^P_t)$ at any continuation history $h_t \supseteq h_{t'-1}$ for any $t = t', ..., \hat{T}$ (hence trading at any time till $\hat{T}$ gives them surplus smaller than ϵ), and the surplus from trading after $\hat{T}$ is also smaller than ϵ. In consequence, for these types, the expected continuation utility from choosing some non-cheapest $s'_{t'}$ at $h_{t'-1}$ is bounded from above by $-C_{t'}(s'_{t'}, \theta_b; h^P_{t'-1})+\epsilon$, which is bounded from above by $-C_{t'}(s_{t'}, \theta_b; h^P_{t'-1})-w(h_{t'-1})+\epsilon$, where $s_{t'} \in S_{t'}(h^P_{t'-1})$ is dynamically cheapest. By choosing $s_{t'}$ and never buying, these buyer types would obtain $-C_{t'}(s_{t'}, \theta_b; h^P_{t'-1})$, which would be a profitable deviation. This contradiction shows that in all equilibria, all types choose a dynamically cheapest experiment at each time period t.

Remark 8: Adding Learning Opportunities in Theorem 2. With straightforward notational adjustments, the above proof establishes the validity of an analogue of Theorem 2 for an environment in which each experiments generates signals (including public signals) at two moments within each period t, just before the buyer's experiment choice, and just before the seller's pricing decision. As in Theorem 2, the public history of experiments then consists of publicly observable experiment choices and all public signal realizations.

Proof of Propositions 1 and 2

The proof follows the same steps as the proof of the second parts of Theorems 1 and 2.

Proof of Proposition 3

We derive Proposition 3 from Theorem 2.[40] Recall that in this proposition the values are constant throughout the game and the costs of experiments do not depend on the history of actions and experiment realizations, but bargaining power might switch

[40]See Madarasz and Pycia (2023), subsumed by the present draft, for a longer direct proof of Proposition 3.

from the seller to the buyer at some random time $\hat{t}$; that time also does not depend on the traders' actions. We denote by Γ the game of Proposition 3.

The argument relies on two transformations of Γ. First, observe that the payoffs in the continuation game after the trading parties learned that the bargaining power switched at period $\hat{t}$ are uniquely determined because the buyer not only has the bargaining power (and knows her own value $b(\theta_b)$) but she also knows the seller's value v. Let $t^* \in \mathcal{T}, t^* \geq \hat{t}$ be a period that maximizes

$$\delta^{t^*-\hat{t}}(b(\theta_b) - v) - \sum_{t=\hat{t},\dots,t^*} \delta^{t-\hat{t}} \min_{s_t \in S_t} c(s_t, \theta_b)$$

if such maximum exists and is positive; else, set $t^* = \infty$.[41] If $t^* = \infty$, then in the essentially unique continuation equilibrium the buyer makes price requests that are never acceptable to the seller and the object is never sold; the continuation payoff of the seller is then 0 and the continuation payoff of the buyer (including the cost of experiment in period $\hat{t}$) is $-\sum_{t\in\mathcal{T},t\geq\hat{t}} \delta^{t-\hat{t}} \min_{s\in S_t} c(s_t, \theta_b)$. If $t^* < \infty$, then there also is an essentially unique continuation equilibrium in which before period t^* the buyer makes unacceptable price offers; in period t^* the buyer proposes to the seller the price exactly equal to the seller's value v and the seller always accepts; and in any subsequent period the buyer makes price offer weakly above v; the continuation payoff of the seller is then again 0 and the continuation payoff of the buyer (from the cost of experiment in period $\hat{t}$ onwards) equals $\delta^{t^*-\hat{t}}(b(\theta_b)-v)-\sum_{t=\hat{t},\dots,t^*} \delta^{t-\hat{t}} \min_{s\in S_t} c(s_t, \theta_b)$. The information generated by experiments has hence no impact on the payoffs and for this reason if, post-switch, the buyer has any choice of experiment, the buyer would choose the least expensive one. In particular, the claim of Proposition 3 holds true for any experiment choices made by the buyer after learning that the bargaining power is hers from the current period onwards.

To analyze pre-switch experiment choices, let us denote by Γ' the game that is

[41] The maximum always exists if T is finite. As the costs are bounded, the maximum also exists whenever $b(\theta_b) > v$; as the costs are nonnegative, this inequality also ensures that $t^* = \hat{t}$. In general, there might be multiple maximands t^*, and our argument does not depend on which one we choose. Also if the maximum exists and is equal to $-\sum_{t\in\mathcal{T},t\geq\hat{t}} \delta^{t-\hat{t}} \min_{s_t\in S_t} c(s_t, \theta_b)$ and then it does not matter whether we set t^* to be finite or $t^* = \infty$. In particular, the payoffs of the buyer and seller and the game transformations we discuss below are the same irrespective of the choice of t^*.

identical to Γ except that (i) it ends at the moment the buyer learns about the bargaining power switch in Γ (of course, both games might also end before the bargaining power switches) and (ii) if Γ' ends because of the bargaining power switch in Γ, then the traders receive in Γ' the above-derived payoffs. The analysis so far shows that the strategic choices faced by both traders at all pre-switch information sets of Γ are equivalent to choices at the corresponding information sets of Γ'.[42]

The key step of the proof is to construct an auxiliary game Γ'' that satisfies the assumptions of Theorem 2 and at pre-switch periods is strategically equivalent to Γ and Γ'. In particular, in Γ'' the bargaining power always stays with the seller. We construct this game so that it runs over $T+1$ periods whenever T is finite and over an infinite sequence of periods otherwise.[43] The construction starts with Γ' and modifies it as follows:

- We keep each pre-switch experiment as in Γ' except that each signal realization in Γ'' is replaced with a pair of messages: the first message in the pair is the signal realization in Γ' (observed by the same parties that observed it in Γ') and the second message is publicly observable and takes one of the following three values $\{\text{early switch}, \text{late switch}, \text{no switch}\}$; we call the first message the "original signal" and the second message the "switch flag". The marginal probability of each profile of original signal realizations is the same as in Γ and, conditional on the profile of realizations of the original signal: the switch flag takes value "no switch" with the probability that there was no power switch in the entire relevant period in game Γ; the switch flag takes value "early switch" with the probability that the bargaining power switched in Γ in the relevant period just before the buyer's experiment choice; the switch flag takes value "late switch" with the probability that the bargaining power switched in Γ in the relevant period just before the seller's pricing decision.

- Following the realization of either the early or the late switch flag in period-t,

[42] By strategic choices being equivalent we mean that there is an isomorphism—in which both traders make the same choices—between PBEs in Γ' and PBEs in Γ when restricted to choices before the bargaining power switch.

[43] For finite T, we make game Γ'' one period longer than T so as to ensure that the buyer has at least one more experiment choice following the bargaining power switch in period T.

game Γ'' continues with the seller having the bargaining power but with values and experiment sets and costs set as follows:

- Seller's value post switch is set to be higher than the highest price the buyer is willing to accept; note that this is possible given our assumptions that the values are bounded.
- Buyer's value post switch is set to 0.[44]
- Post switch the buyer is choosing from singleton sets of experiments whose expected dynamic costs at the period the switch is realized are equal to $\max\{0, b(\omega) - v\}$. The informational content of these experiments might be arbitrary.

The game Γ'' satisfies the assumptions of Theorem 2, and hence this theorem implies that in Γ'' there exist equilibria in which the buyer always chooses a cheapest experiment in all periods and that in all equilibria the buyer only chooses cheapest experiments on path. In particular, this implies that conditional on the switch flag taking the value "early switch" in some period t, the buyer's expected payoff in Γ'' is the same as in Γ'. Thus, pre-switch Γ'' is strategically equivalent to Γ' and hence to Γ. In effect, we conclude that in Γ there exist equilibria in which the buyer always chooses a cheapest experiment in all periods and that in all equilibria the buyer only chooses cheapest experiments on path.

Proof of Proposition 4

The proof of this proposition follows the same steps as the proof of the second part of Theorems 1 and 2. The role of the cost difference between the cheapest and more expensive experiments is now played by the cost difference between experiments s' and s''.

[44]The buyer's value can be set arbitrarily as long as it is sufficiently low so as that there is never any surplus from trade.

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 1. We thus know from Theorem 1 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 1 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 1 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.

Proof of Theorem 3

By symmetry, we can assume that the posteriors of the seller's cheapest experiment have full support. Note that below the trade surplus of a trader refers to the trader's payoff at the trade stage; it does not take into account the experiment cost. Theorem 3 then follows from the following:

Claim 1 *In any equilibrium that implements efficient trade the following obtain: (i) All types of the seller choose the cheapest experiment, (ii) All types of the seller receive full trade surplus, and (iii) All types of the buyer choose the cheapest buyer experiment.*

Proof of Claim 1. As in the proof of Theorem 1, compactness and continuity imply that we can select $w > 0$ to be a lower bound on the difference between the cost of the cheapest experiment $\underline{s}_v$ and any other experiment $s_v \neq \underline{s}_v$ for all seller types. The ex post individual rationality of φ implies that the trade surplus of each seller type is bounded from above by $1 - v$, where v is the seller's value. Hence, in any equilibrium, the seller types with values strictly higher than $1 - \Delta$ pick the cheapest experiment. Let $\bar{v} \leq 1 - \Delta$ be the supremum of values of types picking a more expensive experiment; if there are no such types let $\bar{v} = 0$. Hence all types in $(\bar{v}, 1]$ choose the cheapest experiment. By assumption that the posteriors of this cheapest experiment have full support, at the bargaining protocol stage, all seller types in $(\bar{v}, 1]$ are still possible for any realization of the signal $\underline{s}_v$.

We now show that the types in $(\bar{v}, 1]$ are given full trade surplus. We define an auxiliary direct one-agent mechanism M^S in which only the seller participates and reports their value; M^S generates the outcome that the equilibrium of the original game would have generated given the seller's reported value and the buyer's true value. By construction, this auxiliary direct mechanism is incentive-compatible. It also selects efficient object allocations because the original equilibrium implements efficient trade. In particular, the auxiliary mechanism M^S maps sellers' values to allocations monotonically.

Consider now another one-agent direct mechanism M^* that maps seller's value to efficient object allocation and sets the payment to the seller so that the seller receives the full surplus from efficient trade given buyer's true value. When efficiency allows the allocation of the object to either the seller or the buyer, we choose the same allocation that would obtain under M^S. This mechanism is incentive compatible for the sellers and it has the same object allocation rule as M^S. Because the allocations are monotonic in values, the celebrated payoff equivalence result of Myerson (1981) implies that the mapping from the seller's values in $(\bar{v}, 1]$ to expected payments in the one-agent mechanisms M^S differ by a constant from the analogous mapping in the one-agent mechanism M^*.[45]

This constant is 0. In the auxiliary mechanism M^S the highest type of the seller has trade surplus of 0 because the original bargaining protocol is ex post individual

[45] The type interval restriction is needed as Myerson's result relies on the support of values being an interval. We know that all types with values in $(\bar{v}, 1]$ are in the support at the trade stage by our assumption that the seller's cheapest experiment has full support posteriors.

rationality for the buyer. In the auxiliary mechanism, M^*, the highest type of the seller also has a trade surplus of 0. Hence, for types with values in $(\bar{v}, 1]$, the expected payments are the same between mechanisms M^S and M^*, and the object allocations in M^S and M^* are also the same. Hence, the utilities each seller's type receive are the same. This implies that in M^S, and hence in the original equilibrium, the seller with values in $(\bar{v}, 1]$ fully extracts the rents of efficient trade, just as they do in M^*.

This further implies that $\bar{v} = 0$. Indeed, if not then let θ_v be a seller type with value strictly above $\bar{v} - \varepsilon$. This type's trade surplus is bounded from above by ε plus the full trade surplus of a type with value $\bar{v}$. For $\varepsilon < \frac{\Delta}{2}$ this type would strictly benefit from deviating to the cheapest experiment and then reporting a type in $(\bar{v}, \bar{v} + \varepsilon)$, as this deviation saves Δ in experiment's cost while reducing the trade surplus by less than $2\varepsilon < \Delta$. By the same argument, any type with value 0 would select the cheapest experiment. Thus, all seller types choose the cheapest experiment (proving (i)).

We can further conclude that all seller types obtain the entire surplus from trade (proving (ii)). As no buyer type can then obtain any surplus from trade, all types of the buyer choose the cheapest experiment (proving (iii)), and concluding the proof of Claim 1, and hence also the proof of Theorem 3.

Proof of Proposition 6

We prove Proposition 6 by establishing the following stronger, though more technical, lemma.

Lemma 1 *Suppose that at least one of the trader's cheapest experiment has full-support posteriors. There exists an equilibrium which implements efficient trade if, and only if, the other trader's cheapest experiment has almost full disclosure, in the following sense, if*

$$\underline{s_b}(\omega) = \underline{s_b}(\omega') \implies b(\omega) = b(\omega') \text{ or } b(\omega) = \inf b \text{ or } b(\omega') = \inf b.$$

We prove this lemma in two steps. By symmetry, we can assume that seller's cheapest experiment has full-support posteriors.

Step 1. We show that if the cheapest experiment of the buyer leads to almost full disclosure, then there is an equilibrium that implements efficient trade.

Indeed, fix a bargaining protocol φ. As above, let $\underline{s}_b$ be the cheapest experiment of the buyer; which now by assumption leads to almost full disclosure. The following is then a perfect Bayesian equilibrium. All buyer types choose $\underline{s}_b$ and all seller types choose the cheapest experiment $\underline{s}_v$, which by assumption has full support posteriors. On path, both traders update via Bayes rule. Off path, we endow the buyer with an arbitrary posterior belief and we endow the seller with the posterior belief that the buyer has the highest value 1.[46] To the protocol φ the buyer reports their value truthfully; the seller reports the maximum of the seller's own value and of the supremum of the value the buyer might have given the seller's posterior belief after observing the buyer's experiment choice and its realization. On path the buyer chooses the almost full disclosure experiment, and hence at most two buyer values are possible following any realization.

This profile of beliefs and strategies is an equilibrium. The buyer's reporting strategy is a best response because it leads to the trade payoff of 0 and given that the seller's reported value is weakly above the buyer's true value (both on and off path of the PBE) and given the ex post individuality rationality of φ, no higher buyer's payoff is possible. The seller's reporting strategy is a best response off-path (straightforwardly), as well as on-path. To see the latter, notice that the almost full disclosure property of the buyer's experiment implies that one of the following two cases is obtained: (a) The seller's on-path report equals the buyer's true value, hence the buyer's report. Because φ assigns the good to the trader with higher reported value, and assigns it to the buyer in case of a tie, given the buyer's truthful reporting, the seller's report leads to trade and, by ex post individual rationality, the price equals the seller's report. The ex post individually rational for the buyer also implies that no seller's report could have given the seller a higher payoff. (b) The buyer's value

[46] As before, we can input beliefs that are consistent with the realization of the experiment. As the cheapest experiments are unique, following the other trader's deviation at the information choice stage, we can adapt the PBE construction so that the seller believes that the buyer has value sufficiently near the supremum of values consistent with the realization of the off-path s_b'' and the buyer believes that the seller has value sufficiently near the infimum of values consistent with the realization of off-path s_s''.

and report are equal to the seller's lowest possible value (0); in this case the seller's strategy leads to 0 payoff from trade and no higher payoff is possible given the ex post individual rationality for the buyer.

Given these reporting strategies and posterior beliefs, the experimental choices are also optimal. Were the buyer to deviate, the seller would report a value that is weakly above the buyer's value and the ex post individual rationality of the bargaining protocol would imply that the buyer would have no trade gain while the experiment the buyer deviated to has higher cost; thus no deviation at the experiment stage is profitable for the buyer. No deviation is profitable for the seller because on path the seller receives the entire trade surplus and pays the lowest possible experimental cost.

Step 2. We conclude the proof of the lemma by showing that if there is an equilibrium which implements efficient trade, then the unique cheapest buyer experiment must have almost full disclosure.

By way of contradiction, suppose that following the unique cheapest buyer's experiment, at the trade stage, two or more non-zero values of the buyer are still possible. Then there is a positive mass of types with values weakly above some b_h, and a positive mass of types with values weakly below some b_ℓ, where $b_h > b_\ell > 0$ and b_h and b_ℓ are in the support of the seller's posterior of the possible valuations of the buyer.

Recall that we showed in the proof of Theorem 3 that, in any equilibrium implementing efficient trade, the buyer does indeed choose the cheapest experiment and the seller captures all trade surplus. Thus, buyer types with values weakly above b_h get no trade surplus from reporting the truth and would strictly prefer to report b_ℓ as any such deviation gets strictly positive expected surplus by virtue of efficiency of the conjectured equilibrium. Indeed, the buyer's posterior belief about the seller's type assigns strictly positive probability to the seller having value strictly lower than b_ℓ and hence positive probability to trade after having reported b_ℓ; given the ex post individual rationality of φ, any such trade is profitable for the buyer's types with value above b_h.

Thus, in an efficient equilibrium, all types of buyers with values weakly above b_h are reporting values strictly lower than their true values. Ex post individual rationality of φ, hence implies that a buyer with value $b \geq b_h$ is not buying from seller with value lower but close to b_h; hence the equilibrium does not implement efficient trade.

This contradiction concludes the proof of the above Lemma and hence the proof of Proposition 6.

Appendix B: Equilibrium

In the main analysis we focus on weak perfect Bayesian equilibria. Thus our characterization results—most notably that cheapest experiments are chosen in all weak perfect Bayesian equilibria—apply to all stronger equilbrium concepts. We also proved the existence of weak perfect Bayesian equilibria, and we now show that existence extends to more demanding equilibrium concepts in environments in which the lowest cost experiment is unique at every history. To focus attention, we consider the existence parts of Theorems 1 and 2.

First note that we can require that, at each period τ, the off-path beliefs in the equilibrium we constructed in the proof of Theorems 1 and 2 are consistent with the outcome of the experiments till period τ. Indeed, when the lowest cost experiment is unique then the compactness of Θ_b, the continuity of each experiment's cost in $\theta_b \in \Theta_b$, and the finiteness of the set of experiments at each history h_t, imply that there is some $w(h_t^P) > 0$ such that for all buyer types the difference between the cost of any more expensive experiment and the cheapest experiment is at least $w(h_t^P)$. We can then endow the seller with the belief that the buyer is of type $\theta_{b,\tau}$ whose full value $B_t(\theta_{b,\tau}; h_\tau^P)$ is below the supremum of full values of types consistent with the outcome of the experiments observed by the seller till $h_{\tau+1}$ and above this supremum minus $\frac{w}{2}$; to make the beliefs consistent with the Bayes rule, we then assume that $\theta_{b,\tau+1} = \theta_{b,\tau}$ as long as $\theta_{b,\tau}$ remains consistent with the outcome of experiments observed by the seller till time $\tau + 2$. As in the proof of Theorems 1 and 2, the seller's best response is to set at any time-τ history h_τ the price that is weakly higher than both $B_\tau(\theta_\tau; h_\tau^P)$ and the seller's value; 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.

Second, the above equilibrium construction satisfies Kreps and Wilson's (1982) sequential-equilibrium refinement whenever this refinement is well-defined. As the buyer has only a finite number of actions at each history at which the buyer moves, the sequential-equilibrium consistency of beliefs is well-defined for the seller and the above equilibrium construction satisfies it. Indeed, in checking sequentiality it is enough to consider trembles in which buyer type $\theta_{b,\tau}$ selects the more expensive experiment with small probability that however is much larger than the probability of trembles by other buyer types. The sequentiality of seller's beliefs implies that, if the buyer knows the state of nature, and hence only the seller updates their beliefs during the game, then there exist sequential equilibria in the game we study (and our results are valid for them). If the buyer does not know the state of nature, then the sequentiality of buyer's beliefs is not well defined in our model in which the seller has a continuum of moves in each period.