Cost over Content: Information Choice in Trade

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

August 2025

Abstract

A privately-informed buyer takes an action that impacts the seller's information. The available actions differ in both content (what information is revealed and what remains hidden) and costs. For a large class of trading environments where buyers can choose from an arbitrary set of signal-generating processes and the costs of these processes can depend on the buyer's private information we establish a "cost-over-content" theorem: buyers will only choose least expensive processes. We explore implications for data trade, voluntary disclosure, the scope for signaling, the privacy puzzle, and the importance of default information for regulation.

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^{TM} 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, Nemanja Antic, Itay Ashlagi, Jean-Michel Benkert, Laura Doval, Ernst Fehr, Amanda Friedenberg, Tan Gan, Bob Gibbons, Marina Halac, Jason Hartline, Emir Kamenica, Navin Kartik, Shengwu Li, Elliot Lipnowski, Meg Meyer, Benny Moldovanu, Stephen Morris, Wojciech Olszewski, Marco Ottaviani, Alessandro Pavan, Agathe Pernoud, Ferdinand Pieroth, Andrea Prat, Armin Schmutzler, Vasiliki Skreta, Joel Sobel, Peter Norman Sørensen, Jakub Steiner, Bruno Strulovici, Balazs Szentes, Steve Tadelis, Curtis Taylor, Cedric Wasser, 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 Northwestern Communication, Persuasion, and Disclosure Conference, and FORC 2025. The project received funding from the European Research Council under EU Horizon 2020 grant agreement No 866376.

1 Introduction

Exogenous private information may profoundly impact trade: information affects the distribution of surplus among trading parties, can cause market failure, or enhance trade surplus by limiting monopoly power.[1] At the same time, a privately-informed party can often impact what her trading partner knows. Consumers may choose between different privacy and data-sharing options when shopping online on a platform.[2] Suppliers may choose between different warranty lengths to convince the purchaser of the quality of their product. Employees seeking promotion may signal their privately-known talent by getting certifications; they can then also decide whether to voluntarly disclose their exam scores.

What information would a privately informed party endogenously choose for their counterpart and how much would she be willing to pay to impact what the other party knows? Would she be willing to protect her privacy at a cost, engage in costly signaling or pay to voluntarily disclose information? Understanding such choices is fundamental for analyzing trade outcomes. They are also key for evaluating policies aimed at privacy, disclosure, or consumer protection. Their importance is further highlighted by well-known market failures in the presence of fixed information structures and the resulting selection problems.[3]

To illustrate the forces at play let us consider an example of a privacy choice. A buyer privately knows her value for an object sold by a seller; the two parties have common knowledge of the finite-support distribution from which the buyer's value is drawn as well as the seller's value of the object. The buyer can choose between two privacy options: (i) allow the seller to observe her online browsing history or (ii) incur a small cost and ensure that the seller has no access to her browsing history. The seller sees which option the buyer picks, and if the buyer allows her online history to be observed, the seller also updates their belief about the buyer's preferences based on this history. Finally, the seller makes a take-it-or-leave-it offer to the buyer.

In this example, the buyer never pays to hide her online history. Indeed, there

[1] For canonical developments of these ideas see, e.g., Akerlof (1970), Coase (1972), and Myerson and Satterthwaite (1983).

[2] These choices are key for the many business models that are based on access to consumer data. For example, the effectiveness of consumer advertising greatly depends on access to consumer data, Deisenroth et al. (2024) and on privacy laws regulating such access, Goldfarb and Tucker (2011).

[3] 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.

is a perfect Bayesian equilibrium in which the buyer chooses not to pay.[4] The optimality of this choice is ensured by the seller's belief that a buyer who hides her browsing history has high value; hence the price set by the seller is high and in effect the buyer would realize no gain from trade following the deviation to costly privacy protection, making this deviation unprofitable. Crucially, in *no* weak perfect Bayesian equilibrium the buyer ever chooses to pay to hide her online history. To see it, suppose by way of contradiction that there is an equilibrium in which buyers with some values and some online histories choose to protect their privacy with positive probability. Although the seller may or may not be able to perfectly infer the buyer's value from her realized online history, in equilibrium the seller knows a tight lower bound on the values of buyers choosing to hide their histories. In effect, the seller sets the price to be weakly higher than this lower bound. In turn, any buyer type with minimum value has a profitable deviation to allow her online history to be observed. This saves the protection costs without loosing any gains from trade. This contradiction shows that there is no equilibrium where the buyer would pay to hide her browsing history.

One might conjecture that choosing the cheapest option and not paying to affect the seller's information is specific to this example. Instead, we show that this is a universal feature of *any* weak perfect Bayesian equilibrium in a large class of static and dynamic monopoly problems.

In our setup, a price-setting seller sells an object to a buyer (or a continuum thereof) over one or more periods. The object has value to the buyer and the seller. We impose no restrictions on the link between the traders' values nor any restrictions on the initial distribution of private information between them, except that the buyer *knows* her own value (an assumption that we partially relax later). Values may be private, trade may be subject to adverse or advantageous selection, as in insurance problems with privately known risk types, 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.[5]

The buyer has some degree of control over the seller's information: she can choose from an arbitrary finite set of dynamic signal-generating processes or information structures, which we call experiments. The buyer's choice is observed by the

[4] In the paper we focus on weak perfect Bayesian equilibria, a choice which makes our cost-over-content insight stronger. At the same time, this and all other example equilibria we construct also satisfy the more stringent requirement of sequential equilibria. Cf. Appendix A for discussion.

[5] See, e.g., Einav and Finkelstein (2021) and Horton, Johari, and Kircher (2024).

seller and each experiment determines the information that the seller and the buyer receive over time. The information they respectively receive may be the same, correlated, or conditionally independent. 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 regarding 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. This cost may be the buyer's private information, and different buyer types may differ in the incremental cost of choosing more expensive experiments, as in Spence (1973). For instance, the buyer 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.

In particular, the environments we study allow for many modes of information provision or communication. In the above example, the buyer decides whether to protect their privacy as in, e.g., Acquisti, Taylor, and Wagman (2016). Costly signaling of Spence (1973) (and cheap talk of Crawford and Sobel 1982) can be seen as a special case of the above example in which the browsing history is uninformative, and all information is carried via equilibrium choices of the buyer. Beyond the example, the buyer may choose to disclose evidence, as in the classic analysis of voluntary disclosure of Grossman and Hart (1980), Milgrom (1981), or Verrecchia (1983). The buyer may also be able to obfuscate, destroy, or hide evidence, as in, e.g., Dye (1985). In general, we allow the buyer (sender) to *impact* (i) what the seller (receiver) learns about the buyer's value, (ii) what the seller learns about his own value, and (iii) what the buyer learns about the seller's value. We allow the learning to be private or public and we allow the information transmitted to impact what the parties learn about each other's learning as well as about each others' arbitrary higher-order beliefs.

In this general setting—with dynamic pricing, flexible observation structures, and minimal assumptions on information and cost structures, integrating disclosure and signaling—we determine what experiments are chosen.[6] We show that in all equilibria, the buyer opts for a cheapest experiment at each point in time, irrespective of its

[6]Our findings are novel already in simple cases of our model such as static privacy choice, voluntary disclosure, or signaling. In particular, the present paper subsumes our analysis of privacy, Madarasz and Pycia (2020).

content and that of the more costly alternatives. Furthermore, for each sequence of cheapest experiments, there is an equilibrium in which these experiments are chosen on path.[7] 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, there is *no* equilibrium where the sender chooses a non-cheapest experiment. Given any arbitrary set of dynamic information processes, informed buyers select among these on the basis of costs and not on content. The underlying logic resembles that in Akerlof (1972) and Spence (1973): in making their decision the buyer takes into account the (equilibrium) inferences the seller will make. This logic is illustrated in our discussion of the privacy example above.[8]

A corollary of our results is that easy-to-change *information defaults* may have powerful economic consequences, which underscores defaults' relevance for regulation. For the many economic settings in which there is some initial or default distribution of information that a buyer can change incurring some, potentially arbitrarily small, cost, our paper provides general sufficient conditions that ensure that such a change will never occur in equilibrium irrespective of what the default is. Instead, it is this default which pins down the efficiency of and the distribution of benefits from trade. In effect, small changes in the relative costs of different information structures may have large impacts on trade outcomes. We arrive at these conclusions without imposing any equilibrium refinement or bounded-rationality assumptions like the presence of some form of choice overload or directly assumed 'default bias.' Relatedly, in a wide range of contexts, our logic provides an equilibrium rationale for the exogenous information assumption commonly imposed in economic models.

In Section 4 and Appendix A, we discuss the key assumptions that drive the cost-over-content result: observability of the buyer's choice, the presence of a default experiment or a class of experiments that are cheapest for all buyer types, seller's bargaining power, and buyer's knowing their own value.[9] We provide examples

[7] In particular, when experiments do not differ in their costs, then *any* any of them might be selected in an equilibrium, and, with more than one experiment available to the buyer, there are multiple equilibria. E.g., if there is an information structure under which trade is efficient, selecting this information structure, and thus efficient trade, is an equilibrium. Similarly, 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. For expositional reasons, we present the result in three theorems of increasing generality: Theorem 1 covers the static case, Theorem 2 the dynamic case, while Theorem 3 adds history dependence.

[8] We illustrate this logic further with two examples in Section 2. The proof in the general case is fairly involved and we present it in Appendix B.

[9] A fifth assumption built into our model is that the seller cannot condition their price offer on

showing why the result breaks if we relax any of these assumptions completely and we also show that the main thrust of the cost-over-content result obtains under substantially weakened assumptions. This analysis thus provides guidance on when non-cheapest experiment choices may be observed and impact trade in equilibrium and when they do not.

Propositions 1 and 2 show that the result only relies on minimal observability of the buyer's information choice; for example, it suffices that the seller observes whether or not a cheapest experiment is chosen by the buyer without the receiver needing to observe which non-cheapest experiment is chosen. Proposition 3 gives an analogue of the cost-over-content result with no default assumptions. Proposition 5 shows that in a natural subclass of problems, bargaining power can be allocated randomly and switch from the seller to the buyer.[10] Our logic thus extends to cases where the equilibrium rent of each buyer type is always strictly positive and changes stochastically over time. Theorem 3 shows that we can allow the buyer to learn about their own value provided the signals from which the buyer learns are public and the learning preserves the value-ranking of different types of the buyer.

In Section 5, we turn to applications. While the above results make no direct predictions about the content of the information the buyer ends up choosing, they allow us to draw content inferences in the many contexts where there is a natural link between costs and content. In Section 5.1, we show that the signaling theory of Spence (1973) hinges on his assumptions on bargaining power. If Spencian job applicant faced a monopolistic employer, the applicant would choose the least expensive level of education.

In Section 5.2, we propose an explanation of the observed lack of voluntary disclosure in settings in which the presence of hard information is commonly known. Voluntary disclosure is predicted by the unravelling logic of Grossman and Hart (1980) and Milgrom (1981), but, as we discuss in Conclusion, it is rarely observed in markets. While our logic could also be conceptualized in terms of unravelling—higher-value buyers do not pay to impact information because lower-value buyers do not—the two logics are distinct. What drives unravelling in Grossman, Hart, and

buyer's experiment choice, cf. Appendix A.

[10]See our earlier drafts for analysis of general price setting mechanisms, such as, e.g., Double Auctions. In a symmetric environment in which both the buyer and the seller control information and the price is set by a general mechanisms, we have shown that the properties of the cheapest experiments determine market efficiency, effectively endogenizing traders' information in Myerson and Satthertwaite (1983) analysis of bilateral trade.

Milgrom and the subsequent literature is the content of hard information disclosed; indeed, in this literature content drives unravelling even when cost differences are present, as in Verecchia (1983). In contrast, in our logic, cost is the dominant force. The reason is that we focus on environments in which the receiver (seller) has price-setting power, while the voluntary-disclosure literature allowed the sender (buyer) to set the price.

In Section 5.3, we provide an explanation why consumers may support regulation that protects their privacy even though they will never pay to protect it themselves. Our result shows that this so-called *privacy paradox* is an equilibrium phenomenon driven by market power.

As a final application, we study markets, in which the link between cost and content emerges endogenously. We consider a provider of a platform for online shopping who chooses from a set of technologically feasible experiments—e.g., privacy or tracking options—and decides which options to offer to the buyer and at what price or subsidy. The provider 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 the information flows (or lack thereof), we leverage the logic of our cost-over-content theorem to show that the platform charges the seller and neither charges nor compensates the buyer for choosing an information structure. In effect, the platform maximizes the value of the information passed to the seller, as measured in the seller's direct profit from interacting with the buyer (Proposition 5). As a corollary, even though the platform has bargaining power over the seller and buyers have full property rights over their data, buyers may be strictly better off when *data trade* is banned. The ban benefits the buyers in absence of interdependencies between their and seller's values, but, as we also show, in lemon markets à la Akerlof (1972), data trade might benefit the buyers.

In Section 6 we show that our analysis fully extends to dynamic trade environments (Theorem 2).[11] We further show that the analysis extends even if (i) we

[11] For expositional reasons we first present the cost-over-content force in an environment in which the buyer makes only one choice among experiments and the seller makes only one price offer, and then show that these restriction can be fully relaxed.

partially relax the assumption that the buyer knows her own value, and (ii) we allow history-dependent dynamics where the buyer's choice of experiments and the realizations of public signals not only affect the traders' information, as in our baseline setup, but can also affect:

- future sets of experiments and their costs,
- traders' values, and
- the buyer's learning about her own value.

traders' values, and In the presence of these history dependencies, the cost comparison of experiments at any history needs to take into account the experiment's impact on future experiments, both their availability and costs. Theorem 3 provides sufficient conditions under which the buyer choose dynamically cheapest experiments in all equilibria.

This final theorem hinges on the *public observability* of those of buyer's choices and resulting signal realizations that affect future experiments, values, and buyer's learning. If these choices or signal realizations were private, then the buyer may well choose non-cheapest experiments. Indeed a special case of our history-dependent setup is where the buyer's choice generates no information and only affects her own value; many instances of this environment were extensively studied in the context of the hold-up problem. In particular, Gul (2001) showed that relaxing full observability of the buyer's investment choice or the returns thereof while letting the discount factor between periods approximate 1 eliminates hold up and leads to efficient investment choices. This feature of hold up is in stark contrast to our Theorems 1 and 2 in which the buyer's choices impact only the traders' information, and not their values, and in which the above-discussed minimal observability is sufficient for our cost-over-content result to hold and for the buyer to always choose a cheapest sequence of experiments.

1.1 Related Literature

Our paper is related 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 5, our analysis relies on different 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 agents who signal receive payoffs that are commensurate with the receiver's value.[12]

Our results are also related to the literature on voluntary disclosure which considers particular evidence structures that the privately informed trader can disclose to her counterpart. A central insight of this literature is the full disclosure benchmark: in the class of environments they study, Grossman and Hart (1980), Grossman (1981) and Milgrom (1981) show that there is a unique perfect equilibrium, in which all types choose full disclosure. The main departure from the full-disclosure benchmark was established by Dye (1985) for the case when the receiver does not know whether the sender is able to provide evidence. We contribute to this literature by showing that, in our model, full disclosure does not obtain if it is not the cheapest. In Section 5, we describe how the evidence structures of voluntary disclosure are special cases of our setup, and we explain the stark contrast between our results and those of the earlier literature.[13]

Third, we contribute to the literature on privacy.[14] 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 the voluntary disclosure literature discussed above. 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 5.

Our paper is also related 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

[12]Cf. Kreps and Sobel (1994), Riley (2001), and Starkov (2023) for surveys of signaling models all of which share this feature.

[13]In particular, note the contrast between our results and Verrechia (1983), who shows that when disclosure is costly, then low types do not disclose, but high types do, and that the equilibrium converges to full disclosure as this cost goes to zero.

[14]Privacy was the main focus of an earlier draft of the present paper, Madarasz and Pycia (2020).

also hold in the presence of constraints on information transmission. Our approach is similar to 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. In this aspect, our work is closer to Koessler and Skreta (2023) who consider *interim* information design in a context where all information designs are costless and describe environments where ex ante optimal mechanisms are also interim optimal.

We contribute to the study of market failures 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 à la Akerlof (1970).[15] 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. Our Theorem 3 is related to the literature on the interaction of private information and the holdup problem, e.g. Gibbons (1992), and Gul (2001). In Section 6, we discuss differences between information and investment choice.

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 Examples and Base Model

Before describing our base model, we consider two examples.

[15] In earlier drafts we also analyzed endogenous information in Myerson and Sattherwaite (1983). For the impact these theorems had on economic policy, see Milgrom (2004), Tirole (2012), and Loertscher, Marx, and Wilkening (2015).

Example 1. Privacy Protection under Private Values

Suppose that a buyer's value for an object is drawn from a uniform distribution over $[0, 1]$ and its realization is the buyer's private information. The seller has a commonly known value of 0. The timing of their interaction is as follows:

1. The buyer chooses between two actions: *doing nothing* and *protecting privacy.*

2. Having observed signal realization $s = v$, if the buyer chose to do nothing, or $s = \emptyset$, if the buyer protected their privacy, the seller chooses price $p \in \mathbb{R}$ making a take-it-or-leave offer to the buyer.

3. The buyer accepts or rejects the offer. The seller's payoff is $p\chi_{\text{offer accepted}}$ where $\chi_{\text{condition}}$ is the indicator function that takes value 1 if the condition is satisfied and 0 otherwise. The buyer's payoff is $(v - p)\,\chi_{\text{offer accepted}} - c\chi_{\text{privacy protected}}$, where $c \in (0, \frac{1}{8})$ is the cost of privacy protection.

In this example, there is an equilibrium in which all buyer types do nothing. Having learned the buyer's value, the seller then sets $p = v$, and the buyer accepts.[16] Crucially, this equilibrium is essentially unique: in no weak perfect Bayesian equilibrum does the buyer ever protect her privacy. Indeed, suppose by way of contradiction that some buyer type v protects their privacy with positive probability, and let $\underline{v}$ be the infimum of types that do so. In equilibrium the seller understands this and, seeing signal realization $\emptyset$ the seller would set the price p that is weakly higher than $\underline{v}$.[17] Buyer types in $[\underline{v}, \underline{v} + c)$, would then all strictly preferred the initial choice of doing nothing instead of protecting their privacy, which contradicts that $\underline{v}$ is the infimum of types that protect privacy, and hence showing that privacy protection cannot be sustained in equilibrium.

This example gives a simple equilibrium explanation of the well-documented privacy paradox, which we discuss in more detail in Sections 5.3 and 7. Ex ante, not knowing her value, the buyer would prefer all buyer types to choose privacy protection.[18] However, the equilibrium forces make privacy protection impossible.

[16]The equilibrium is supported by the seller's belief that ascribes the deviation to protecting privacy to buyer with value $v = 1$. We focus on weak perfect Bayesian equilibria as this makes our main results stronger, but a sequential equilibrium also exists here and in our other examples.

[17]Indeed, the seller can strictly improve upon any price $p' < \underline{v}$ by setting the price $\frac{p'+\underline{v}}{2} > p'$ because all buyer types choosing protection still buy at this higher price.

[18]Indeed, the seller would then keep their prior belief about buyer's value and optimally set the price $p = \frac{1}{2}$, resulting in the buyer's expected utility being $\frac{1}{8} - c > 0$.

Example 2. Evidence and Interdependent Values.

Suppose now that the buyer's valuation v is drawn from a uniform distribution on $[3,5]$ and again let its realization be the buyer's private information. Let the seller's cost be $2v-4$ so there is adverse selection à la the lemons problem of Akerlof (1970). The timing is as in the previous example, expect the buyer can now choose between three signals (evidence reports): $s^{\text{no}}, s^{\text{partial}}, s^{\text{full}} : [3,5] \to \{\emptyset, L, H\} \cup [3,5]$. Let us fix a parameter $\hat{v} \in (3,4)$ and let resulting realization of these signals observed by the seller before they make the price offer be as follows

- $s^{\text{no}}(v) = \emptyset$ for all $v \in [3,5]$;
- $s^{\text{partial}}(v) = L$ if $v \leq \hat{v}$ and $s^{\text{partial}}(v) = H$ otherwise;
- $s^{\text{full}}(v) = v$ for all $v \in [3,5]$.

The buyer incurs a cost $c(s)$ that depends on the signal s chosen. To simplify discussion, suppose there is a unique cheapest signal.

Analogously to the previous example, in any weak perfect Bayesian equilibrium the buyer chooses the cheapest signal and there is an equilibrium in which the cheapest signal is chosen. We thus have three possibilities:

- If no disclosure s^{no} is cheapest, then the buyer does not disclose any information and there is no trade in equilibrium.
- If partial disclosure s^{partial} is cheapest, the buyer chooses it in equilibrium. Having learned that $v \leq \hat{v}$, the seller then sets the price $p = 3$; the buyer accepts this price and all buyer types $v \in (3, \hat{v})$ receive strictly positive surplus. Having learned that $v > \hat{v}$, the seller sets a high price that is rejected by the buyer, and there is no trade.
- If full disclosure s^{full} is cheapest, then buyer chooses it, the seller sets the price $p = v$, trade is ex post efficient, and each buyer type receives zero surplus.

In effect, the trade outcome is determined by the signal costs. Arbitrarily small changes of these costs can have big impact on traders' payoffs and they can shift the market between efficiency and complete breakdown.

2.1 Model

A buyer (sender) and a seller (receiver) can trade an asset. Each trader's valuation for the asset is determined by the state of nature $\omega \in \Omega$ distributed according to the commonly known probability measure $\mathbb{P}$ which is measurable with respect to a σ-algebra on measurable subsets of Ω. For ease of exposition, below we describe a static trade environment in which the buyer knows their value at the outset. We present the general setup with dynamic games in Section 5.

Initial Information. We allow for an arbitrary initial distribution of information. In particular, we impose no restriction on what the parties may or may not know, nor on higher-order uncertainty. Private information may be correlated or independent. Each trader's information at the beginning of the game 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.

Values. The buyer's value is denoted by $b(\omega) \in \mathbb{R}$, the seller's value by $v(\omega) \in \mathbb{R}$. Both traders receive the value of the object at the time of their transaction. We thus interpret the seller's value as arising from avoiding the cost of servicing the buyer. We assume that the buyer's value $b(\omega)$ is determined by buyer's type $\theta_b \in \Theta_b$, which allows us to write the buyer's value as $b(\theta_b)$; we also assume that b is continuous as a function of θ_b. We impose *no* other assumptions on $b(\omega)$ and $v(\omega)$ or the relationship between them.

Values in our setup may then 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.[19] They may also have arbitrary non-monotonic relationships. For instance, different seller types may have different costs of servicing different buyer's types allowing for match-specific trade surplus.

Information Choice. The buyer chooses from a non-empty finite set S of arbitrary information-generating processes or Blackwell experiments (experiments henceforth). Each experiment s generates three different signals $s = (s^v, s^b, s^p)$. Each signal s^o (where $o \in \{v, b, p\}$) is a mapping $s^o : \Omega \to \Delta(Z)$ from the states of nature

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

Ω to the set of distributions over feasible signal realizations Z, where Z is a compact set. The realization of s^v is observed privately by the seller, the realization of s^b is observed privately by the buyer, and the realization of s^p is observed publicly by both traders.[20] One natural case of this environment is when the parties only learn from the public signal s^p, e.g., because the private signals s^v, s^b always take a constant value.

Special cases of this information environment include Spence (1973) signaling, Grossman and Hart (1980), Milgrom (1981), and Dye (1985) evidence structures, and Crawford and Sobel (1982) cheap talk (see Section 5 and Appendix B for details). The three signals can be arbitrarily correlated. In particular, we allow both public and private learning as well as *second-* and *higher-order* uncertainty that the players can have about each other's beliefs (e.g., the seller can privately learn about the buyer's cost, or the buyer's belief about her valuation, etc). A special case is common or public learning where the buyer and seller update their beliefs only on the basis of the public signal s^p, e.g., because all private signals always take a constant value, say $(s^v, s^b) = (z_0^s, z_0^b)$ for some fixed $(z_0^s, z_0^b) \in Z \times Z$. This is the case in many sender-receiver setups. Another special case is where only the seller learns and does so privately and so $(s^b, s^p) = (z_0^b, z_0^p)$ for some fixed $(z_0^b, z_0^p) \in Z \times Z$.

Experiment Costs. The key aspect of our setup is that choosing different experiments may have different costs for the buyer. We denote by $c(s, \omega) \in \mathbb{R}$ the bounded cost that the buyer incurs when choosing experiment s. We assume that $c(s, \omega)$ is determined by θ_b, which allows us to denote it by $c(s, \theta_b)$. We assume that Θ_b is endowed with a topology and that the mapping $\Theta_b \ni \theta_b \to c(s, \omega)$ is continuous.

Our other assumption on the cost function is that there exists a subset of experiments $\underline{S} \subseteq S$ such that, for each type θ_b, the cost $c(s^*, \theta_b) = \min_{s \in S} c(s, \theta_b)$ if and only if $s^* \in \underline{S}$. We call these experiments the *default experiments*. This assumption means the class of lowest-costs experiments is common across buyer types. It is always satisfied when the costs are independent of ω. 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

[20] We could equivalently represent each signal s^o as a mapping from Ω to a realization space Z. Allowing for lotteries on Z whose resolutions are not determined by the state of nature just simplifies the presentation in the applications. To economize on notation we assume that all signal realizations come from set Z, we could equivalently distinguish between three realization sets Z^v, Z^b and Z^p.

(e.g., in privacy considerations), there is indeed an initial distribution of information the players are endowed with, and choosing some alternative information structure carries some potentially type-dependent cost. The assumption is also satisfied in settings such as Spence (1973) signaling model, in which acquiring no education is the lowest-cost signaling action for all productivity types. We impose no other ordinal or cardinal restriction on the cost function; e.g., allow both cardinal cost differences and the ordinal cost-based ranking of experiments to be type dependent.

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.

Timing. In the one-period base setup, the buyer first chooses the experiment $s \in S$. The seller then observes the buyer's experiment choice and the resulting signal realizations (s^s, s^p), and the seller makes a price offer p. Finally, the buyer observes the signal realizations (s^b, s^p) and either accepts or rejects the offer. (We introduce the dynamic model in Section 6).

Payoffs and equilibrium. When an offer is accepted, the buyer's payoff is $b(\omega) - p - c(s, \omega)$ and the seller's payoff is $p - v(\omega)$. In case there is no trade, payoffs from trade are normalized to zero, but the buyer still incurs the costs of the chosen experiment $-c(s, \omega)$.

We employ the standard notion of weak perfect Bayesian equilibrium (PBE or equilibrium henceforth) as our solution concept.[21] We choose this weak equilibrium notion because it strengthens our main message: that only cheapest experiments are chosen. As this choice weakens our existence results, in Appendix A, we show under mild additional assumptions that the existence of equilibria obtains for stronger equilibrium concepts, including the sequential equilibrium when the seller is uninformed.

[21] The players update their beliefs via Bayes rule whenever possible, and the buyer does not update their belief over their own value based on seller's actions because the seller has no information that is relevant for buyer's value updating. See, e.g., Fudenberg and Tirole 1991.

3 Information Choice: Cost over Content

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

Theorem 1 (Static Case)

Existence: *For any experiment $s \in \underline{S}$, there exists a weak PBE in which all buyer types choose s.*

Cost over Content*: In no weak PBE does the buyer ever choose any non-cheapest experiment, $s \notin \underline{S}$.*

The main message of this theorem is that, given any cost structure, and any initial information, in *no* equilibrium does the buyer ever choose a non-cheapest experiment. The buyer (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 that is $\underline{S}$ is a singleton, then there is an equilibrium in which all buyer types choose this cheapest experiment and in all equilibria all buyer types do so. On the other hand, if there are no cost differences between the experiments, that is, $\underline{S} = S$ then pooling on *any* given sequence of experiments can be supported in equilibrium.[22] Buyer types may also separate and choose different experiments, but they will only choose cheapest ones.

The proof of this theorem and all our other results is in Appendix B. The logic behind Theorem 1 resembles that in Example 1. To prove the first part of the theorem we fix a cheapest experiment $s \in \underline{S}$ and construct a pooling equilibrium in which all buyer types choose s on path. The optimality of this choice is ensured as off path the seller believes that he faces the highest value buyer and sets the price accordingly. To prove the second part of the theorem we show that, as in Example 1, following each experiment choice there is always a buyer type whose expected equilibrium surplus from trading is not much higher than trading surplus this type would achieve following a deviation to a cheapest experiments. In consequence, there cannot be a perfect Bayesian equilibrium in which some types chose a non-cheapest experiment, as (irrespective of the content of the experiments) some of these types would benefit from a deviation to a cheapest 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

[22] For additional comments on this case, see Appendix 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. In particular, 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).

The theorem also tells us 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.

Finally, small costs associated with changing this structure can have significant welfare implications. Example 2 illustrates this issue. In this example, small cost differences are sufficient to move the market from a complete trade failure to fully efficient trade.[23]

We next discuss several applications of the cost-over-content force established in the above theorem (Section 5), and then discuss the key assumptions and limitations of this force (Section 6).

[23]Theorem 1 implies that the correspondence from the profiles of experiment costs to experiments that can be chosen in equilibrium is upper hemicontinous but not lower hemicontinuous, when the topology on cost profiles is Euclidean and the topology on S is discrete. The upper hemicontinuity obtains because any experiment s that is an equilibrium choice along a convergent sequence of cost profiles is among the lowest-cost experiments along this sequence, and hence in the limit of this sequence, and hence it is an equilibrium choice at the limit cost profile. The lower hemicontinuity fails because, when two experiments both have the lowest cost then for each of them there is an equilibrium in which it is chosen, then a tiny perturbation of costs is sufficient for the now more expensive of the two experiment to stop being an equilibrium choice.

4 Limits of Cost over Content

The cost-over-content force determines equilibrium outcomes under the set of assumptions we imposed in Section 2. If any of our main assumptions is fully relaxed, our main finding (Theorem 1) would not hold. At the same time, Theorem 1 remains valid when these assumptions are partially relaxed. In particular, we show that cost over content only relies on a minimal observability of whether a default action was chosen; indeed, a weaker versions of the result remain true with no observability assumptions. Cost over content extends to any random allocation of bargaining power when seller's value is commonly known. Cost over content also obtains when the buyer can publicly learn about her value provided the value updates do not reverse expected value comparisons. For ease of exposition, we maintain the focus on static, one-shot, interactions. In Section 7, we also show that these results extend to general dynamic settings.

4.1 Minimal Observability of Experiment Choice

We assumed that the seller observes the experiment the buyer chose. This is a common assumption in information economics, and, as we show below, we can substantially relax it. Theorem 1, however, does require some minimal observability.

To see why some observability is needed, consider a variant of Example 1 in which the buyer with a privately known value v_b drawn uniformly from $[0, 1]$, chooses between two experiments: (i) a no protection default at cost 0, (ii) and full protection at cost $c \in (0, \frac{1}{2})$. Suppose that the choice is covert, so that the seller does not directly observe which of the two experiments was chosen. When the buyer chooses no protection then, as in Example 1, the seller observes v_b.If the buyer chooses full protection, the seller observes a realization drawn at random from the uniform distribution on $[0, 1]$. There then does not exist a perfect Bayesian equilibrium in which all buyer types choose the cheapest experiment (no protection). Indeed, in such an equilibrium, the seller would set the price equal to the realization of the signal and the buyers with values $v_b > \sqrt{2c}$ would have a profitable deviation to the more costly full protection. Note that the cost over content force would still prevail, just as in Example 1, if the experiment choice was covert but the support of signal realizations generated by the full protection experiment was disjoint with the support of realizations generated by no protection.[24]

[24]For $c = \frac{1}{4}$, there is then an equilibrium in which buyers with values $v_b < \frac{1}{2}$ choose no protection,

How much observability is needed for the cost-over-content insight of Theorem 1? The following proposition shows that it suffices that the receiver observes whether or not the buyer chose a cheapest experiment.[25] Which of the non-cheapest experiments was chosen (if any) can be fully *unobservable*.

Proposition 1 *As long as the seller observes whether or not the buyer chooses an element of $\underline{S}$, there is no perfect Bayesian equilibrium in which the buyer chooses some $s \notin \underline{S}$ on the equilibrium path.*

We can further allow for a flexible observation structure regarding the buyer's choice of experiment. Fix an arbitrary partition on S and, for any $s \in S$ denote by $O(s)$ the element of the partition to which s belongs. Suppose that, given any choice of experiment by the buyer, $s \in S$, the seller observes that the buyer has chosen an experiment from $O(s)$, but not which experiment within $O(s)$.

Proposition 2 *Suppose that for some experiment $\hat{s} \in S$ we have $O(\hat{s}) \cap \underline{S} = \emptyset$. Then there is no perfect Bayesian equilibrium in which the buyer chooses $\hat{s}$ 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.

4.2 Defaults

Our second key assumption is that there is a class of cheapest experiments — defaults — which is common across buyer types. Theorem 1 hinges also on this assumption. In the absence of such a default class, pooling on a cheapest experiment loses meaning (because for different types different experiments may be cheapest), and furthermore, some buyer types might in equilibrium choose experiments that are not cheapest for them.[26]

buyers with values $v_b \geq \frac{1}{2}$ choose full protection, and the seller sets the price $p = \frac{1}{2}$ when seeing signal realization in $[0, \frac{1}{4}) \cup [\frac{1}{2}, 1]$ and the price equal to the signal realization when the realized signal is in $[\frac{1}{4}, \frac{1}{2})$. Similar equilibria exists for any $c \in (0, \frac{1}{2})$. For more involved examples in which, in a context of the Akerlof's lemon model, the sender chooses an experiment covertly without being subject to our cost-over-content force, see Pavan and Tirole (2025).

[25] 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.

[26] As an illustration consider an analogue of Example 1 modifying it only so that: full-privacy-protection is cheaper than no-protection for types $b \in [0, \frac{3}{4}]$ while costing $\varepsilon \in (0, \frac{1}{4})$ more than

In absence of defaults, we can still consider experiments that are uniformly more costly then other experiments. Let us fully relax the assumption that the class of cheapest experiments is common across types; our logic delivers then the following implication as to which experiments are chosen:

Proposition 3 *Suppose that* $s', s'' \in S$, *and* $E[c(s', \omega)|\theta_b] < E[c(s'', \omega)|\theta_b]$ *for all* $\theta_b \in \Theta_b$. *Then,* s'' *is never chosen on path of a PBE.*

In absence of any common default, it is 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.

4.3 Bargaining power

Our third key assumption is that the seller sets the price, as in the classic monopoly problem. If the bargaining power were in the hands of the buyer, some buyer types could have wanted to choose a more expensive experiment. This is familiar from the costly signaling model of Spence (1973) and the classic disclosure models, where it is the sender who sets the price, as we discuss in Section 5.1.[27]

The cost-over-content force, however, does extend to general distributions of bargaining power in environments where there is no learning about the seller's value. If now the buyer has the bargaining power and makes a price offer that the receiver can accept or reject, then obviously the buyer would choose one of the cheapest experiments. More generally, we may then allow the bargaining power to be assigned *randomly*.[28] We present our result on the random allocation of bargaining power (Proposition 5) in the dynamic model of Section 6.1.

4.4 Buyer's Learning

The fourth key assumption we are making is that the buyer chooses the experiment already knowing her own benefit from trade. The crux of this assumption is that the

no-protection for types $b \in (\frac{3}{4}, 1]$. Then there is a perfect Bayesian equilibrium, in which protection is chosen by all buyer types, including those for whom it is not the cheapest. In this equilibrium the seller sets the price $p = \frac{1}{2}$ after privacy is protected, while he sets the price $p = v$ otherwise.

[27] In Spence, the worker engaged in costly signaling has bargaining power because she faces competitive employers; cf. footnote 31.

[28] For general trading protocols such as, e.g., Double Auctions, our earlier drafts additionally show that when information structure is endogenous in bilateral trade à la Myerson and Satthertwaite (1983), then only the cheapest information structures determines whether efficient trade is possible.

buyer does not update her estimate of her own valuation. This assumption again can be relaxed partially, but not fully.[29] As we will show in Theorem 3 as long as the buyer learns about her valuation publicly and such learning satisfies a monotonicity condition, our cost-over-content result remains intact. Otherwise it may well fail.

5 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 questions and we discuss related empirical evidence in the Conclusion.

5.1 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,

[29] To see why we cannot fully relax this assumption, consider the following example with non-monotonic public learning. The buyer privately knows her type θ_b, drawn from a uniform distribution on $[2, 4]$. At the beginning of the game, the buyer chooses between two experiments, one costs 0 and generates a public signal realization that always takes value 0, the other experiment costs $c \in (0, \frac{1}{2})$ and generates public signal s whose realization is distributed uniformly at random on $\{-1, 1\}$. The buyer's value conditional on θ_b and the realization z of the signal s is $v_b(\theta_b, z) = \theta_b + 2(\theta_b - 3)z$. In particular, the buyer's type θ_b is equal to the expected valuation of the buyer, $E_s[v_b(\theta_b, z)] = \theta_b$, but the buyer can learn about her valuation from the more costly experiment. After seeing the realization of the signal, the seller makes a price offer p; seeing the signal realization of z and the price p, and the buyer either accepts or rejects the offer. There is then a perfect Bayesian equilibrium in which all types pay c for the more expensive experiment. In this equilibrium, upon seeing signal realization $z = -1$, the seller infers the value distribution is uniform on $[2, 4]$ and sets price 2; upon seeing signal realization $z = 1$, the seller infers the value distribution is uniform on $[0, 6]$ and sets price 3. Given these prices and the seller's belief that assigns to any deviator to the cheaper experiment type $\theta_b = 4$, all buyer types indeed prefer the more expensive experiment over the cheaper experiment. Private learning by the buyer may also overturn the cost-over-content insight as can be inferred from Roesler and Szentes (2017): in their model the buyer, with no private information about their value, may strictly benefit from costless learning.

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.[30]

In this Spencian case of our model 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; note that in Spence (1973), the bargaining power resides with the sender because she faces a competitive market of receivers.[31] 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 has bargaining power	cost over content	cost over content
Sender has bargaining power	cost over content	Spence (1973)

Table 1: Costly Signalling

Theorem 1 establishes the first row and Proposition 5 below 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

[30] 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 $c(s, \omega) - c(s', \omega) > c(s, \omega') - c(s', \omega')$.

[31] In Spence's analysis, the worker engaged in costly signaling faces competitive employers who offer wages that give the worker her entire contribution to firm's surplus. If the worker made wage offers, she would have asked for the same wage as offered by the competitive market.

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).

5.2 Verifiable Disclosure

The disclosure setting is another special case of our model in which an experiment represents arbitrary soft or hard evidence, i.e., the realization of s may be correlated with the payoff state. In particular, our concept of experiment contains the canonical settings of verifiable disclosure, including both the evidence structures introduced by Grossman and Hart (1980), Grossman (1981), and Milgrom (1981) (GHM below), in which the receiver knows the kind of hard evidence available to the sender, and the evidence structures introduced by Dye (1985), in which the availability of evidence is private information of the sender.

Disclosure with evidence. GHM 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.

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 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.

Theorem 1 contrasts with the seminal results of 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. The substantive difference between GHM and our Theorem 1 and Proposition 5 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 (given the strategy of the highest type) fully disclose her value, etc. This drives their classic unraveling logic.[32] 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 in our setup is cost not content.

Disclosure with uncertain evidence. 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. The uncertainty allows for partial disclosure in equilibrium. The seminal model of partial disclosure, due to Dye (1985), hinges on the uncertainty of the receiver as to whether the sender has evidence.[33] In Dye's

[32]When the sender has bargaining power, the same logic drives disclosure in the presence of costs, as shown by Verrechia (1983).

[33]Content drives 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.

model partial disclosure is also driven by *content*—bad news is suppressed, good news is disclosed—and we can also map Dye's evidence structure into our model.

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).

5.3 Privacy, Data Trade, and Platform Design

We now consider applications of our result to the problem of consumer data. In the application of our model to the analysis of privacy, the experiments the buyer chooses among may correspond to 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.[34]

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 3.

5.3.1 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,

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

Catalini, and Tucker (2017), Johnson, Shriver, and Du (2020). Firms appear to use online consumer 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.[35]

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.[36] 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.

[35] 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."

[36] 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.

5.3.2 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.[37] The uninformed platform provider has access to a finite set S of technologically feasible experiments (privacy designs). Each experiment $s \in S$ is a mapping from the set of states of nature Ω to some space of signal realizations Z observed by the platform. We interpret each s as a privacy design. For simplicity, all experiments are equally costly for the provider to supply, and we normalize this cost to zero. As in the base model, the buyer has type θ_b and value $b(\theta_b)$ from the object while the seller has value $v(\omega)$. The parties share a common prior. The timing is as follows.

First, the platform provider selects some subset $S_P \subseteq S$, assigns to each $s \in S_P$ some fee or subsidy $c(s) \in \mathbb{R}$, and makes a take-it-or-leave-it offer to the seller. If accepted, the offer determines payment the seller makes to the platform and the information the seller has when setting the price in the online channel (see below).

Second, the seller can accept or reject the platform's offer:

- If the seller accepts the offer, the buyer and the seller have access to both online and offline sale channels. The buyer then chooses whether to shop online, or interact with the seller offline.

- If the seller rejects the offer, the platform obtains a payoff of 0, while the seller and the buyer interact offline.

If the seller accepted the platform's offer and the buyer chose to shop online, then the buyer chooses an experiment $s \in S_P$; having observed the information determined by the seller's contract with the platform, the buyer's experiment choice, and the signal realization, the seller then makes a price offer, which the buyer can accept or reject.

The offline interaction of the buyer and seller when the seller rejected the platform's offer is the same as when the seller accepted the offer and buyer chose to shop offline. In both cases, the seller receives no further information about the buyer and makes a take-it-or-leave-it offer that the buyer can accept or reject.

[37] 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.

The resulting provider's payoff is the sum of the payments the provider received from the seller and the buyer (with negative payments allowed). The seller's payoff is the profit from the sale to the buyer minus the payment to the provider (if any); where the profit from the sale is $v(\omega)$ plus the price paid by the buyer if the purchase took place and 0 otherwise. The payoff of the buyer who makes an online purchase is $b(\omega)$ minus the payments to the platform provider and the seller. If the buyer makes an offline purchase, her payoff is $b(\omega)$ minus the payment to the seller. If she chooses online option s and then rejects the seller's offer, her payment is $-c(s)$, and if she goes offline and does not buy, her utility is normalized to zero.

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.

Proposition 4 *In any PBE,*

1. *The provider chooses $S_P \subseteq S$ that maximizes the seller's expected profit when selling to the buyer.*

2. *All privacy options $s \in S_P$ chosen with positive probability by the buyer are free, $c(s) = 0$.*

In effect, 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 profits. 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 S). 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 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 2 from Section 2 and S consisting of the experiments from the first two bullet points of this example.

6 General Dynamics

In this section we consider dynamic extensions of our model. For ease of exposition, first, we extend to the baseline setup to simple dynamic trade. Next, we introduce general history-dependent dynamics.

6.1 Simple Dynamics

Values and initial information are just as in the baseline setup from Section 2, while trade and experiment choice can now occur over time.

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). The timing within each period is analogous to that in the base model.

In the beginning of each period $t \geq 1$, the buyer first chooses a dynamic signal process which generates signal realizations in period $t' \geq t$. We continue to 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 experiment of period t, generates three signals $s_{t,t'} = (s^v_{t,t'}, s^b_{t,t'}, s^o_{t,t'})$, two private ones (one observed by the buyer, one by the seller) and a public (observed by both) in each subsequent periods $t' \geq t$. Analogously to the base model, each $s^0_{t,t'}$ (where $o \in \{v, b, p\}$) is a mapping $s^o : \Omega \to \Delta(Z_{t'})$ from the states of nature Ω to the set of distributions over feasible signal realizations in period t', $Z_{t'}$, where $Z_{t'}$ is a compact set.

Following the buyer's choice from S_t in period t, the seller observes the buyer's current experiment choice and the realizations of relevant signals. He then makes a price offer p_t which the buyer, after observing the realizations of relevant signals, 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. 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 and adopt an even significantly weaker assumption later in this section (Theorem 3).

Payoffs and equilibrium. Each player's payoff is the discounted sum of payoffs from each period of their interaction; the discount factor is $\delta \in [0, 1)$.[38] When an offer is accepted in period t, the buyer's period-t payoff is $b(\omega) - p_t - c(s_t, \omega)$, and the seller's period-t payoff is $p_t - v(\omega)$. In case there is no trade in period t, the buyer period-t payoff is $-c(s_t, \omega)$, while the seller's period-t payoff is 0.

In this dynamic environment, the weak perfect Bayesian equilibria (PBE) still satisfy the cost-over-content property:

Theorem 2 (Simple Dynamics)

Existence: *For any sequence s_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.*

To illustrate the above consider the canonical problem of dynamic trade based on the model of Gul, Sonneschein and Wilson (1982), the domain of the Coase conjecture.

Example: 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

[38] In all theorems we may allow possibly different discount factors $\delta_b, \delta_s \in [0, 1)$ and when the horizon T is finite, we can allow $\delta_b, \delta_s \in [0, 1]$.

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.

By using Theorem 2 in the proof of the data-trade Proposition 4, we can show that the analogue of this proposition holds true if, following the buyer's choice of experiments, the seller and the buyer interact over any number of periods T, including $T = \infty$. This dynamic analogue of Proposition 4 implies an economic force which is in contrast to a key aspect of the classic Coase conjecture, the value of being privately informed at the time of trade. Recall the data trade model of Section 5 and suppose that there is no data trade (e.g. S contains only one uninformative experiment). Comparing a buyer who enters object trade being privately informed about her value 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 who enters data trade being privately informed about her value (interim data trade), to one who enters data trade but becomes informed right after data trade, that is, when object trade begins (ex ante data trade). Proposition 4 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.[39]

Limits and Bargaining. The assumptions we discussed in Section 4 that are necessary for Theorem 1 to hold, a fortiori are also necessary for Theorem 2. The exact analogues of Propositions 1, 2, & 3 continue to hold in the dynamic environment.

We now describe the random allocation of bargaining power result mentioned in Section 4. Consider a setting where there is no private information about the seller's value, an often-imposed assumption. We may then allow the bargaining power to

[39]This 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).

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 5 *[Bargaining Power Switch] Suppose that the bargaining power switches from the seller to the buyer at a time determined by $(\hat{t}, \hat{m})$ and that there is no learning about $v(\omega)$.*

- *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 2, 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.

6.2 History-dependent Dynamics.

In the simple dynamic model above, the traders' values were stationary and the information choice problems were separable across time. We now relax these assumptions. The traders' values, 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 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, then the seller sees some signal realizations and makes his offer, then the buyer sees some signal realizations. By h_{t-1} 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 $t-1$ and (ii) the realized history up to the point of the buyer's choice of experiment in the beginning of period t. 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, the game only continues to period t if the buyer rejects the seller's offer in period $t-1$ (or if $t=1$, in which case only period t is relevant), and the information the buyer has at both these histories is the same, as is the information the seller has. 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.[40] We assume that $b(\omega, h_t^P)$ is measurable with respect to (θ_b, h_t^P), thus we can denote it by $b(\theta_b, h_t^P)$.

- 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.

[40] Extending the notation we denote by $b(\omega, \emptyset)$ the value the buyer has before the start of the game, where discounting implies that the buyer is indifferent between, say, $b(\omega, \emptyset) = \delta$ before the start of the game and $b(\omega, h_1^P) = 1$ at the end of the first period. Note also 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.

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:

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

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 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:

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

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,

$$C_0\left(\theta_b; \emptyset\right) \equiv \delta \min_{s_1 \in S_1(\emptyset)} C_1\left(s_1, \theta_b; \emptyset\right).$$

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 environments of Theorems 1 and 2 satisfy these assumptions. 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.[41]

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, one experiment $s_t \in S_t(h_{t-1}^P)$ 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.[42]

[41] 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.

[42] 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.

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 $b\left(\theta_b, h_t^P\right) = b_1\left(\theta_b, \emptyset\right) + b_2(h_t^P)$ 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 Theorems 1 and 2 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 $\cup_{h_{t-1}^P \in \mathcal{H}^P} S_t(h_{t-1}^P)$ such that, for any $h_{t-1}^P \in \mathcal{H}^P$ we have $\sigma(h_{t-1}^P) \in S_t(h_{t-1}^P)$.[43] The public experiment schedule σ induces an experiment choice $\sigma(h_{t-1}^P)$ 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 $\sigma(h_{t-1}^P)$ is dynamically cheapest.[44]

Theorem 3 (History Dependence)

Existence*: For any experiment schedule σ consisting of dynamically cheapest experiments, there exists a PBE such that all buyer types pool on $\sigma(h_{t-1}^P)$ 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 3 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 3 establishes that the buyer always

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

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

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 3 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: Observability of Choice. Theorems 1 and 2 remained valid under minimal assumptions on the observability of the buyer's actions (see Proposition 1). Similarly, the validity (and proof) of Theorem 3 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 $S_t(h_{t-1}^P)$, 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.[45]

Theorem 3 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 3 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.

[45] 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.

Remark: Costly Information versus Costly Investment. For Theorems 1 and 2 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, $\delta \to 1$, her investment choice becomes fully efficient.[46]

Unobservable return from investment. Even when the buyer's investment choice is itself observable, the lack of public observability of the return from her 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 shows, when $T = \infty$ and $\delta \to 1$, 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.[47]

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)

[46] Relatedly, Gibbons (1992) points out that PBEs in the hold-up problem with unobservable investment might require mixed investment strategies. As usual, $\delta \to 1$ can be interpreted as the shortening of time that elapses between two offers.

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

does not matter for our Theorems 1 and 2. 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.

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.[48]

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 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

[48]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.

data that can be collected, and the extent to which they can be exclusively purchased by firms, may be more effective.[49] 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.

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.

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

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.

Fudenberg, D. and J. Tirole (1991). *Game Theory.* MIT Press.

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.

Starkov, E. (2023). Only time will tell: Credible dynamic signaling. Journal of Mathematical Economics, 109, 102894.

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.

Appendix A. Corollaries, Extensions, and Limits

A.1. Constant Costs

A key aspect of our setup is that different information structures *may* be differentially costly. In particular, our setup includes as special case the cheap talk environment of Crawford and Sobel (1982), in which each experiment (or message) $s \in S$ is in itself uninformative. Here, all experiments have the same costs, i.e., $c(s,\omega) = c(s',\omega)$ for *all* (s,s') and ω.

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.

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.[50]

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.[51]

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

Appendix A.2. 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

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

[51] 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.

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 3.

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 3 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 3, 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.

A.3. Two-sided Information Choice

In the analysis above only the buyer chose an experiment. However, our proofs remain valid (with no substantive adjustments needed) if the seller can also choose an experiment simultaneously, or after the buyer chooses one. Such an experiment choice can allow the seller to direct 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.

All our results remain valid under general two-sided information choice. See our earlier draft, Madarasz and Pycia (2025) for more details.

A.4. Further Comments on Limits of Cost over Content

In Section 4, we discussed fourth assumptions that drive our cost-over-content insight. A fifth assumption built into our model is that the seller cannot commit to condition the price offer on the experiment choice of the buyer. Our model hence generalizes signaling and disclosure (as discussed in Section 5) but it does not allow the seller screen the buyers.[52]

To see why result would fail if the seller was able to commit to a price conditional on the experiment, let us consider a variant of Example 2. The seller moves first and publicly commits to an offer that conditions the price on the experiment choice and signal realization. After observing the seller's commitment, the buyer chooses between full disclosure experiment at cost $c \in (0,1)$ and no disclosure at cost 0; the partial disclosure is too costly to be chosen. Given the price scheme to which the seller is committed, the choice of experiment and its realization determines the price at which the buyer can buy the object. The buyer then decides whether to accept or reject the purchase at this price.

In this example, there is a perfect Bayesian equilibrium in which the seller commits to the price $v_B - c$ if the buyer chooses full disclosure and their value v_B is revealed to be in $[3, 4-c]$; the seller also commits to a high, trade-precluding, price if the buyer does not choose full disclosure or if the buyer's value is disclosed to be above $4-c$. Given this price scheme, the buyer with value in $[3, 4-c]$ chooses the

[52]The screening offers that condition on the experiment choice break the cost-over-content force. However, this force still determines equilibria in environments in which seller sells goods of varied quality and screens the various types of the buyer through a menu of price-quality pairs as in Mussa and Rosen (1978), provided the offered menu is not conditioned on the buyer's choice of experiment nor its realization.

full disclosure and accepts the resulting seller's price offer.

Note that the key assumption behind the cost-over-content force is the lack of conditioning on the buyer's choice, not the timing of offers. As Section 6 shows the cost-over-content insight is valid in dynamic environments in which the seller makes sequential price offers over time both before and after the buyer's experiment choices.

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

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

Appendix B. Proofs

Proof of Theorems 1, 2, 3

We focus on Theorem 3 as Theorems 1 and 2 are effectively its special cases. 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.[53]

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, let's denote it by $\theta_{max} \in \Theta_b$, and have the seller attribute all 0-probability deviations to this type.[54] 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 $\theta_{\max}$

[53] 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.

[54] 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.

and the seller's value.[55] 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.[56] 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 $w(\theta_b, h_{t'-1}) \geq w(h_{t'-1}) > 0$ 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 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 $\underline{B}_t(h^P_t) - p_t(h_t)$ 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 $\underline{B}_{t''}(h^P_{t''}) - p_{t''}(h_{t''}) > \frac{1+\delta}{2}R$. All on-path types θ_b would

[55] 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.

[56] 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.

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 $C_{t''}(\theta_b; h^P_{t''})$ 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 $\delta^{t'''-t''}(\underline{B}_{t'''}(h^P_{t'''}) - p_{t'''}(h_{t'''})) \leq \delta^{t'''-t''}R \leq \delta R$ (if buying at time t'''). By definition of $C_{t''}(\theta_b; h^P_{t''})$ and the law of iterated expectations, the expected dynamic cost of future experiments is equal to $C_{t''+1}(\theta_b; h^P_{t''})$. 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 $p_t(h_t) \geq \underline{B}_t(h^P_t)$ 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 $-C_{t'}(s'_{t'}, \theta_b; h^P_{t'-1})$ plus the type's continuation equilibrium expectation of $B_t\left(\theta_b, h^P_t\right) - \underline{B}_t(h^P_t)$ 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 $\underline{B}_{t'-1}(h^P_{t'-1})$ 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 $\underline{B}_{t'-1}(h^P_{t'-1})$ at $h_{t'-1}$ are within ϵ of $\underline{B}_t(h^P_t)$ at any continuation history $h_t \supseteq h_{t'-1}$ for any $t = t', ..., T$. 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: Adding Learning Opportunities in Theorems 1, 2, and 3. With straightforward notational adjustments, the above proof establishes the validity of analogues of Theorems 1, 2, and 3 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.

Proof of Propositions 1 and 2

The proof follows the same steps as the proof of the second parts of Theorems 1, 2, and 3.

Proof of Proposition 3

The proof of this proposition follows the same steps as the proof of the second part of Theorems 1, 2, and 3. 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''.

Proof of Proposition 4

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 Proposition 5

We derive Proposition 5 from Theorem 2.[57] 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 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 5.

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

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

if such maximum exists and is positive; else, set $t^* = \infty$.[58] 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

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

[58] 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^*.

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 $\delta^{t^*-\hat{t}}(b(\theta_b) - v) - \sum_{t=\hat{t},\ldots,t^*} \delta^{t-\hat{t}} \min_{s\in S_t} c(s_t, \theta_b)$. 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 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'.[59]

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.[60] 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

[59] 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.

[60] 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.

"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, 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.[61]
  - 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.

[61] 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.