



Bilateral trade with strategic gradual learning [☆]

Kfir Eliaz ^{a,b,*}, Alexander Frug ^c

^a Tel-Aviv University, Israel

^b Aarhus University, Denmark

^c Universitat Pompeu Fabra and Barcelona Graduate School of Economics, Spain



ARTICLE INFO

Article history:

Received 4 April 2017

Available online 16 November 2017

JEL classification:

D82

D83

Keywords:

Strategic scheduling

Gradual learning

Endogenous asymmetric information

Gains from trade

ABSTRACT

We propose a model of bilateral trade in which private information about the quality of an asset can be acquired only gradually over time. An asset is characterized by a vector of binary i.i.d. attributes, and its worth to a player is equal to a weighted sum of the attributes' values (where weights differ across players). The seller is initially uninformed about the attributes' values, and each period he decides whether to make a price offer or to costlessly inspect an attribute's value. The buyer does not know the attributes' values, but he may or may not observe which inspections were performed (we consider both cases). We study the seller's strategic scheduling of inspections and its effect on the realized gains from trade in equilibrium. We identify the necessary and sufficient conditions under which the players can realize some gains from trade, and all gains from trade.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

Gathering information, identifying the relevant parts and internalizing their implication is a process that takes time. Even when information is readily accessible, decision-makers typically allocate their attention to a variety of tasks, and hence cannot devote their full resources to information acquisition. Consequently, information gathering and processing is often gradual. For example, the owner of an asset may only gradually learn all the various attributes of his asset; it takes time and experience for a worker to realize his ability in the various dimensions of his job; and a forecaster or analyst only gradually learns and processes the relevant characteristics of a state of nature.

The gradual nature of information acquisition is a form of friction that introduces new strategic considerations to economic environments where the participating agents may be asymmetrically informed. In particular, time, even if not costly per-se, may signal how much – and possibly what type of – information has been acquired and processed. Consequently, the *amount* of information that is learned (the degree of asymmetric information between agents), and the *scheduling* of information acquisition (which piece of information is learned each period) are determined *strategically*. Our goal is to understand how these strategic considerations affect trade.

We take a small step towards achieving this goal by focusing on bilateral trade, the basic building block of any trading environment. To isolate the pure effect of strategic gradual learning, we analyze the case in which only *one party* can acquire private information and information acquisition has *no direct cost*. We assume that all the bargaining power is in the hands

[☆] We thank Andrea Prat and Asher Wolinsky for valuable feedback. Eliaz acknowledges the financial support from ISF grant 374/16. Frug acknowledges the financial support from the Spanish Ministry of Economy and Competitiveness (projects ECO2014-56154-P and SEV-2011-0075).

* Corresponding author.

E-mail addresses: kfire@post.tau.ac.il (K. Eliaz), alexander.frug@upf.edu (A. Frug).

of the party that can acquire information, who makes a take-it-or-leave-it offer to the uninformed party. The side that can access information may be either the seller or the buyer – the exact identity does not affect our analysis. In some contexts it may be more natural for the buyer to be acquiring information.¹ In other contexts, only the party with the initial property rights over the asset can learn more about that asset. To fix ideas, we assume that it is the *seller* who acquires information.

The literature on bilateral trade with private information either begins its analysis when parties have *already acquired* their private information, or the parties decide in advance how much information to acquire and this information is acquired *instantaneously*.² We begin our analysis at a stage when both parties are symmetrically uninformed and then explicitly model the *gradual* nature in which private information is strategically acquired. Consequently, both the *extent* of asymmetry in information, and the *nature* of the asymmetry, are determined *endogenously*. How does this affect the parties' ability to realize potential gain from trade? How does it affect the frequency and timing of trade?

To address these questions we propose a simple model of bilateral trade over a multi-attribute asset. Each attribute is a binary random variable (an attribute is either available or not) and the attributes are i.i.d. The value of the asset to the buyer and seller is a weighted sum of attributes, where the weights of the buyer may differ from those of the seller. The different attributes may be interpreted as representing different potential benefits that can be reaped from the asset, and the two parties have different opinions about their contribution to the overall value of the asset. Both the buyer and seller are initially uninformed about the values of the attributes, but the seller can learn the realizations of the attributes gradually over time. In each period, the seller decides whether to make a take-it-or-leave-it offer to the buyer or to inspect another attribute – in which case he chooses an attribute and observes its realization at no cost. We first assume that the buyer observes neither the attributes that were inspected nor their realizations. The solution concept we employ is a refinement of perfect Bayesian equilibrium that imposes some “natural” restrictions on out-of-equilibrium beliefs.

Our model is meant to be a stylized representation of trading situations with the following key features. First, the buyer and seller disagree on the relative contribution of different attributes of the asset to its overall value. It is therefore unknown ex-ante which party can benefit most from the asset but it depends on the realized profile of attributes (in contrast to a standard lemons problem). Second, the time the asset has been with the seller signals to the buyer the extent to which the seller has already explored potential applications of the asset and learned whether they are worthwhile. If the seller is trying to sell the asset after holding on to it for a relatively long period, it is more likely that he has learned a lot about the asset and its potential.

To see what type of situations have the above features, think of an entrepreneur who develops a new technology or product, which may be used domestically or abroad. The entrepreneur can investigate the domestic demand for his invention, as well as the demand in various foreign countries. Such an investigation would take time and most likely would not yield results on all countries at the same time. Rather, the entrepreneur would only gradually learn about the potential demand in different countries. In addition, the logistics, regulations, costs and cultural differences involved in conducting business in different countries may lead the entrepreneur to value the demand of one country over another. However, a foreign entity interested in buying the entrepreneur's invention is likely to put different values than the entrepreneur on the potential demand in different countries. This may be due to greater expertise of the potential buyer in these markets, or because the potential buyer faces less regulations, or because his corporate culture is closer to the corporate culture in these countries. If the potential buyer were to be approached by the entrepreneur with a sales offer, it would try to infer from the timing of the offer, how much the entrepreneur has learned about the global potential of its invention. The longer the entrepreneur held on to its invention, the more likely it is that he researched the potential demand for it.

This paper studies how the scheduling of information acquisition by a seller affects the extent to which gains from trade can be realized. We first show that the equilibrium sequencing of inspections depends *only* on the seller's weights: attributes are inspected in *decreasing* order of importance to him. Intuitively, the more important an attribute is, the more informative it is to learn its value. This allows us to establish necessary and sufficient conditions for the existence of ex-post efficient equilibria where *all* gains from trade are realized. Roughly speaking, these conditions state that there is some integer $k \leq n$ such that for each of the k most valuable attributes to the seller, the seller's value is sufficiently above that of the buyer's. Interestingly, when an ex-post efficient equilibrium exists, ex-post efficiency is the unique equilibrium outcome. When it does not exist, then some gains from trade are realized in equilibrium only if the total weight that the seller gives to his k' least important attributes is lower than that of the buyer's for some k' . Conversely, if this condition is met, then some gains from trade are realized in any equilibrium.

When the buyer cannot observe which inspections the seller uses the timing of sale offers to signal this information. Even though in equilibrium the buyer infers what attributes the seller inspects, the fact that this information is not directly observed induces the seller to consider only his own valuation when deciding which attributes to inspect. We also study the case where the buyer observes which inspections are performed by the seller but does not observe the outcome of these inspections. In this case, the buyer can condition his willingness to pay on the actual inspections, and this can serve to incentivize the seller to take the buyer's weights into account and not only his. Hence, in equilibrium the seller strategically schedules his inspections by considering the difference between his and the buyer's valuation of each attribute. This allows

¹ For example, start-ups that come up with technological innovations may lack information on the potential applications and marketability of that innovation. This information, however, may be more easily accessible to an experienced firm that is interesting in buying the start-up.

² This literature analyzes the implication of *costly* information acquisition. See e.g., [Persico \(2000\)](#) and [Bergemann and Välimäki \(2002\)](#). Our analysis differs from these studies in that information is *costless* to acquire and information is gathered gradually over time.

us to show that it is always the case that in equilibrium *some* gains from trade are realized. *All* gains from trade are realized if and only if the seller's value is sufficiently above the buyer's for each attribute that is more valuable to the seller (specifically, the seller's excess valuation must exceed the aggregate excess valuation of the buyer over all attributes that the buyer values more). Observable inspection therefore enlarge the set of parameters for which ex-post efficiency can be attained. The rough intuition for this is that the seller schedules his inspections in a way that mitigates the adverse selection problem: When the buyer is offered to trade, he infers that this offer signals relatively more bad news to the seller than to the buyer.

The majority of the literature on bilateral trade considers buyers, who when offered a price, ask themselves “why did the seller make this particular price offer?”. Our contribution is to add a time dimension, which leads the buyer to also ask, “why is the buyer making me an offer *now*?”. If the buyer takes into account how long a good has been with the seller before it was put on sale, then the timing of a price offer may convey information on the amount and type of private information that the seller has. Hence, the seller can try and affect the information that the buyer gleans from the timing of offers by strategically choosing *what* to learn about his good and *when* to learn it. This can help mitigate the adverse selection problem that is inherent in trade under asymmetric information.

The paper is organized as follows. The next section discusses the related literature. Section 3 presents the basic model when inspections are unobservable, and Section 4 presents the analysis of this model. Section 5 studies the case of observable inspections. Concluding remarks are presented in Section 6. All proofs are in the text.

2. Related literature

The idea of strategic scheduling of learning was studied in Frug (2016) in the context of cheap talk. In that paper a sender, who is initially uninformed about the state of nature, has to decide each period which potential state to inspect and what cheap talk message to send. The author shows that full communication can be attained in equilibrium if the sender inspects states in opposite direction of his bias (e.g., if the sender's bias is towards high states, he inspects states from high to low).

The decision problem of which attributes to inspect when goods have several attributes was also analyzed in Klabjan et al. (2014). They consider the following single decision-maker problem. There is an object with several attributes, where each attribute's value is drawn from some distribution, and the decision-maker's payoff is a weighted average of the attributes' values. The decision-maker has to choose between the object and some fixed outside option. The decision-maker is initially uninformed about the attributes' values, but he can inspect each attribute at some cost. The authors consider the case in which the decision-maker can commit to inspect some set of attributes, and the no-commitment case. They show that when attributes have the same costs and are ranked by second-order stochastic dominance, no uninspected attribute will be dominated by an inspected attribute. The no-commitment case is analyzed for two special cases: binary distributions and equal costs and two independent and symmetric distributions with reservation utility equal to the expected value of the object. In the former case, the optimal order of inspection is decreasing in order of dominance and the latter case is in decreasing order of indices that depend on the distribution and the cost.

The major difference between the above paper and our own is that we consider a *two-player game*, where the seller's inspection decisions affect the buyer's beliefs, and hence the equilibrium price. This is precisely our innovation, to introduce gradual attainment of information into a strategic setting, where the sequencing of learning has strategic implications. In particular, a key feature of our model (which is absent in Klabjan et al., 2014) is the potential presence of adverse selection which the seller tries to influence through his inspection strategy. Thus, the efficient composition of the inspection set will be determined so as to mitigate the adverse selection problem, and will depend on both the seller's and buyer's valuations. While there are no direct costs of inspection in our model there are indirect costs through the effect of inspection on the continuation payoff, which is determined endogenously in equilibrium.

Our model has the features that a buyer tries to infer the seller's private information from the timing of a price offer, and the seller attempts to manipulate these inferences by deciding what to learn and when to make an offer. Similar features appear in Taylor (1999) and Kaya and Kim (2015), who study a lemons market where a seller sequentially meets buyers, each of whom observes some noisy signal about the quality of the good. The authors are interested in the question, how are the buyers' beliefs affected by the length of time in which a good was on the market but not sold? Taylor (1999) studies a two-period model of two-sided asymmetric information and compares the trading outcomes under different assumptions regarding the observability of past events. In general, buyers become more pessimistic about the quality of the asset over time. In Kaya and Kim (2015), however, the buyers may become more pessimistic or more optimistic over time, depending on the prior beliefs. On the one hand, a longer duration may signal poor quality since no one was willing to buy the good. On the other hand, a longer duration may actually signal high quality since the seller was unwilling to compromise on the price.³

In contrast to these papers, the seller in our model strategically decides when to learn (if at all) each piece of information. Since he meets only one buyer, the buyer's beliefs about the quality of the asset are based only on the information that the seller may have gathered during that period, and which persuaded him to put the asset for sale. Similar to

³ Both Kaya and Kim (2015) and our own paper deal with dynamic adverse selection. By now there is a huge literature that studies different aspects of this topic, many of which are not directly related to our model. We therefore discuss only those papers that are closest to our own.

Kaya and Kim (2015) a longer duration can in principle mean both positive and negative news. On the one hand, if the seller waited a long time, then he has accumulated more private information – including information which is more valuable to the buyer. If this information persuaded him to sell, then it should mean bad news for the buyer. On the other hand, if the seller did not sell at an earlier period it means that the information he gathered early on was probably good. In equilibrium the price asked at a late period must be lower than the price asked at an earlier period since otherwise, the seller would wait with the offer. This means that in equilibria where trade can occur in multiple periods, the buyers' beliefs become more pessimistic as time goes by (since equilibrium prices are less than or equal to the buyer's willingness-to-pay). Furthermore, when the buyer observes the seller's inspections (the case analyzed in Section 5), his beliefs about the quality of the asset is affected by both the price offer and by the actual learning schedule.

Our paper touches on the question of whether less private information can allow parties to realize more gains from trade. In the context of a lemons market Levin (2001) showed that more asymmetric information can either increase or decrease the amount of trade. We do not model a lemons market since the efficient allocation of the good is state-dependent. Nevertheless, as in Levin (2001) more asymmetry in information does not necessarily hinder trade. In particular, while the buyer and seller are symmetrically uninformed, all gains from trade can be realized when the seller gathers private information. In our model, in contrast to Levin (2001), both the nature and the extent of informational asymmetry develops gradually and endogenously.

A central theme of our work is that it is important not only what the seller *learns* but also when he learns each piece of information. A related theme, studied in Guttman et al. (2014), is that it is important not only what an agent *discloses* but also when he discloses it. They study a two period model where each period an agent may observe up to two signals of the quality of his asset and has to decide whether to reveal any of the signals he has observed so far and not disclosed yet. His decision affects each period's market price, which equals the expected value of the asset, conditional on the history. The authors show that the price at the end of the second period given disclosure of one signal is higher if the signal is disclosed later in the game. As their paper and ours analyze completely different frameworks, the only commonality is that both papers highlight how an informed agent can affect the inferences of an uninformed agent by strategically choosing the timing of an action that conveys information (in our case the action is a price offer, while in their case it is disclosure of hard information).

The idea that it takes time to gather private information about the value of a transaction was previously explored in the context of incomplete contracts in Bolton and Faure-Grimaud (2010). In that paper each side to a contract can privately learn about the value of the contract by repeatedly drawing a signal that may either fully reveal the information or reveal nothing. Each draw is costly in that it delays agreement and future payoffs are discounted. In contrast, the seller in our framework faces the dilemma of *what* to learn each period, taking into account that although learning is costless, the learning *path* has a strategic effect of influencing the buyer's beliefs.

We conclude by mentioning several recent papers that study endogenous information acquisition in environments that are substantially different from ours. Dang (2008) considers ultimatum bargaining in which there are always gains from trade, and where each side before taking his action (seller proposing, buyer responding) can incur a cost and become fully informed. Branco et al. (2012) analyze an optimal stopping problem in which a consumer decides how long to learn about a good and whether to buy it.⁴ Hwang (2016) studies a model in which a single seller faces infinitely many potential buyers who arrive according to a Poisson process. The seller is initially uninformed and learns exogenously from fully informative signals that arrive at random times, while buyers infer information from calendar time. Finally, Choi (2016) considers a dynamic durable good market where sellers learn from the trade decision of the previous sellers and from informative signals that arrive at random times.

3. A model

A given asset is characterized by n attributes, X_1, \dots, X_n . An attribute X_i is a binary random variable that gets the value $x_i = 1$ (the asset *has* the attribute X_i) with probability θ and value $x_i = 0$ (the asset does *not* have the attribute X_i) with probability $1 - \theta$ (a capital letter denotes a random variable while a small case letter denotes a realized value). Let $\Omega = \{0, 1\}^n$ be the set of all possible profiles of realizations of the n attributes. The random variables X_1, \dots, X_n are *i.i.d.*, and a profile of realizations $x = (x_1, \dots, x_n)$ is referred to as a state of nature. We denote by x_{-k} the profile of realizations of all attributes other than X_k , and by $\mathbf{1} = (1, 1, \dots, 1)$ and $\mathbf{0} = (0, 0, \dots, 0)$ the highest and lowest states, respectively.

There is a seller and a buyer whose valuations for the asset are given by the weighted sums $V_S(x) = \sum_{i=1}^n \alpha_i x_i$ and $V_B(x) = \sum_{i=1}^n \beta_i x_i$, respectively, where $\alpha_i, \beta_i > 0$ for each i . The key property assumed here is that the marginal contribution of every attribute X_i is *independent* of the realizations of X_{-i} . Section 6 discusses the implication of relaxing our assumptions that the X_i 's are distributed identically and that the marginal utility from X_i is independent of X_{-i} .

We denote the players' preferences by $\alpha \equiv (\alpha_1, \dots, \alpha_n)$ and $\beta \equiv (\beta_1, \dots, \beta_n)$ and normalize $\sum_{i=1}^n \alpha_i = 1$. Hence, there are ex-ante strict gains from trade if and only if $\sum_{i=1}^n \beta_i > 1$. We assume that there exist a pair of states, x and x' such that $V_S(x) > V_B(x)$ and $V_B(x') > V_S(x')$. Thus, the ex-post efficient allocation of the asset is state-dependent.⁵ We further

⁴ While the authors motivate their model by discussing multi-attribute goods, they assume these attributes are infinitely small (and there are infinitely many of them) and reduce the problem into one where the consumer's expected utility follows a Brownian motion.

⁵ In contrast to a pure lemons problem where the buyer values the asset more than the seller in every state of nature.

assume that the players have clear rankings over the importance of the different attributes, and $\alpha_1 > \dots > \alpha_n$. We will therefore refer to the i -th most valuable attribute, from the seller's perspective, as either attribute i or X_i .

The buyer does not observe the realized values of the attributes before purchasing the asset. The seller is initially uninformed about these values but can learn them over time in the following manner. At the end of each period $t \leq n$ the seller decides whether to make a take-it-or-leave-it price offer to the buyer, or to inspect one of the attributes that have not been previously inspected. Thus, if each period the seller inspects an attribute, then at the end of period t he would make a decision knowing the realizations of t attributes (at the end of period 0 the seller makes a decision without knowing any realization). If a price offer is made, then the game ends with trade if the buyer accepts the offer, and it ends with no trade if the buyer rejects the offer. If the seller inspects an attribute, then he observes the realized value of that attribute. The buyer observes neither the attributes that were inspected nor their realizations. If no price offer is made at the end of n periods, then the game ends with no trade.

There is no cost to inspecting attributes, hence the players' payoffs depend only on the realized state of nature and the terms of trade (if any). If the game ends with trade, then the seller receives the price, while the buyer receives his value of the asset minus the price paid. If the game ends with no trade, then the seller receives his value of the asset while the buyer receives a payoff of zero. Both players do not discount future payoffs. The distribution over Ω , the preferences α and β , and the seller's learning technology are commonly known.

A strategy for the buyer specifies for each pair consisting of a time period $t \leq n$ and a price p_t , a decision of whether to accept or reject the offer. In order for the seller to have a best response to the buyer's strategy we restrict attention to buyer's strategies with the property that for every period there is a maximal price that the buyer is willing to accept. We identify any buyer's strategy with the vector of maximal prices accepted by the buyer, $\bar{p} = (\bar{p}_0, \bar{p}_1, \dots, \bar{p}_n)$.

A nonterminal (seller's private) history h_t is given by a sequence of attributes inspected before period t , and the values of these attributes $(j_1, \dots, j_{t-1}; x_{j_1}, \dots, x_{j_{t-1}})$. Let H be the set of all nonterminal histories. A strategy for the seller is a pair of functions $(o(h), f(h))$, where $o(h)$ is an *inspection function* that assigns to every history $h \in H$ an attribute that was not inspected in h and $f(h)$ is a *trading function* that assigns to every history h either a price offer, which is a non-negative number, or no price offer, which is denoted \emptyset . Note that inspection and trading functions are defined *independently* of each other.

A history $h = (j_1, \dots, j_i; x_{j_1}, \dots, x_{j_i})$ is *consistent* with o if $o(j_1, \dots, j_t; x_{j_1}, \dots, x_{j_t}) = j_{t+1}$ for all $t < i$. Note that given any inspection function \hat{o} , and an index $t \in \{0, 1, \dots, n\}$, every sequence of t realizations $y \in \{0, 1\}^t$, uniquely determines an order of inspection \hat{o}_y until time t , and thus, determines a unique history $h = (\hat{o}_y; y)$ that is consistent with \hat{o} . We call the restriction of the inspection function o to histories that are consistent with o , an *inspection plan*.

Fix an inspection function o and a strategy for the buyer \bar{p} . Let $f(o, \bar{p})$ be an *optimal trading function* given o and \bar{p} .

Definition 1. An inspection function o is said to be weakly dominated if there exists an inspection function o' such that for every strategy of the buyer \bar{p} , the seller's expected payoff from $(o', f(o', \bar{p}))$ is at least as high as the expected payoff from $(o, f(o, \bar{p}))$, and for at least one buyer's strategy it is strictly higher.

The solution concept we employ is pure-strategy perfect Bayesian equilibrium (PBE) in which the seller uses an undominated inspection function, and at every history the buyer believes that seller is using an undominated inspection function. Henceforth, whenever we write "equilibrium" we mean this refinement of pure-strategy PBE.

We say that a given equilibrium is an *equilibrium with trade* if trade occurs in at least one state $x \neq \mathbf{0}$ in that equilibrium. An equilibrium is *efficient* if it maximizes the ex-ante expected surplus among all equilibria. An equilibrium is *ex-post efficient* if at every state of nature the trade-outcome of the equilibrium allocates the asset to the player with the highest valuation.

Discussion. The focus of this paper is on the *strategic* effect of gradual information acquisition. Hence, to isolate this effect we assume that inspection is costless. Adding costs of information would only confound the strategic motivation of acquiring information with its direct effect on payoffs.

We employ a novel equilibrium refinement that focuses only on the inspection component of the seller's strategy. A natural question that arises is why do we not use familiar refinements such as dominant strategies or undominated strategies? The answer is rooted in the observation that a seller's strategy is composed of two independent components, the inspection and trade functions, where the optimality of the latter crucially depends on the buyer's strategy. In particular, there is no dominant strategy because for any given seller's strategy, there is another strategy with higher price offers, that performs better against a buyer's strategy that accepts higher price offers.

In the next section we show that there exists a unique undominated inspection function. Thus, even though the optimal trade function depends on the price offers the buyer is willing to accept each period, we can identify the unique inspection plan that allows the seller to get the highest payoff against *any* buyer's strategy, by choosing an optimal trade function given the buyer's strategy. An alternative approach might have been to show that every undominated strategy for the seller consists of an inspection function with particular properties. However, this would not work since every inspection function is consistent with an undominated strategy, where, for example, after any history the seller asks for a price of one.

Our assumption that at most one price offer is made keeps the analysis tractable by reducing the possible histories on which a player can make his action contingent. Without this assumption, a player's action could depend on previous price

offers, and in constructing equilibrium strategies one would need to specify what each player would do after every possible history of rejected price offers. This complicates the analysis without adding substance.

Our decision to focus on PBE rather than on sequential equilibrium simplifies the construction of equilibria since there are many off equilibrium histories (which include both deviations from the equilibrium order of inspection and from the histories where price offers are supposed to be made) for which one has to specify a perturbation in order to satisfy the consistency requirement. More importantly, the notion of sequential equilibrium does not impose more discipline on out-of-equilibrium beliefs relative to the structural restrictions that we impose on beliefs.

4. Analysis

This section explores the implications of strategic gradual learning by the seller. We begin with an illustrative example that is still rich enough to present the main strategic aspects and results developed below.

4.1. An illustrative example

Suppose there are only two attributes and that $\beta_1 + \beta_2 = 1$. Using the notation defined above, X_1 is the most important attribute to the seller (thus, $\alpha_1 > \frac{1}{2}$).

A benchmark. Consider first the benchmark case in which the buyer does not observe how long the asset has been with the seller. In this case, the sequence of inspections that lead to a price offer cannot influence the buyer's willingness to pay. Since inspections are costless, it is always optimal for the seller to inspect all attributes before choosing whether to offer the asset for sale.⁶

The set of equilibrium outcomes (price and final allocation of the asset) in our benchmark is exactly the same as in the case where it is commonly known that the seller is fully informed: Equilibria in which the asset is not traded at any state other than $\mathbf{0}$ (no-trade equilibria) always exist.⁷ Moreover, if $\alpha_1 < \beta_1$, it is straightforward to show that only the no-trade equilibria exist. On the other hand, if $\alpha_1 > \beta_1$, gains from trade can be realized in an equilibrium if and only if $\theta\beta_2 \geq \alpha_2$. Thus, if θ is sufficiently small (a high likelihood that the asset is "a lemon"), gains from trade cannot be realized in equilibrium.

Our model. Suppose the buyer *does* observe how long the asset was with the seller but does not observe which attributes are inspected. This generates an opportunity for early sale offers that reveal to the buyer that only one attribute was inspected. Hence, in this case the sequence of inspections can influence the terms of trade. Clearly, any equilibrium where trade takes place only at the end of the last period is equivalent to some equilibrium of the fully informed seller benchmark. The more interesting equilibria are those with trade at the end of the first period when only a single attribute is inspected.

Which attribute to inspect? Note that if trade occurs in the first period, the realized value of the inspected attribute must be zero.⁸ Denoting by X_i the attribute inspected in period 1 and letting p be the equilibrium price, we can write the seller's expected payoff as,

$$\theta(\alpha_i + \theta\alpha_j) + (1 - \theta)p.$$

The key observation is that $\alpha_1 > \alpha_2$ implies that the seller can benefit from changing the order of inspections if and only if $i = 2$. Thus, in any equilibrium with trade at period 1, X_1 (the most important attribute to the seller) is inspected first.

When are gains from trade realized? If $\alpha_1 > \beta_1$ then in contrast to the benchmark, an ex-post efficient equilibrium exists for all $\theta \in (0, 1)$. To see this, note that the following is an equilibrium: At $t = 1$, the seller inspects X_1 and makes a price offer of $p = \theta\beta_2$ if and only if $x_1 = 0$. At $t = 2$, X_2 is inspected and no price offer is made. The buyer accepts any price offer (weakly) below $\theta\beta_2$ at the end of period 1, and rejects any price offer at period 2. By making an early price offer, the seller reveals that he is not fully informed about the qualities of the asset. The buyer, who recognizes that inspecting X_1 is the *unique* optimum for the seller, is willing to accept any price offer not exceeding $\theta\beta_2$ – his expected valuation conditional on $x_1 = 0$.

If $\alpha_1 < \beta_1$, then gains from trade cannot be realized in equilibrium (as in the benchmark). The seller's expected valuation given $x_1 = 0$ exceeds the buyer's willingness to pay, and hence, there cannot be an equilibrium with trade at the end of period 1. There cannot be an equilibrium with trade in the second period since at that point the seller is fully informed and we have already argued that when $\alpha_1 < \beta_1$ there is no trade in equilibrium. What precludes equilibria with trade when $\alpha_1 < \beta_1$ is the fact that the seller begins his inspections with X_1 . If, instead, he inspected X_2 , and the buyer knew that, all gains from trade could be realized. This illustrates that, perhaps not surprisingly, observability of inspections (but not

⁶ Inspecting all attributes, however, need not be the unique optimum for the seller. Suppose, for example, there exists an equilibrium in which the seller inspects both attributes, and sells the asset if and only if $x \in \{(0, 1), (0, 0)\}$. In this case, inspecting X_2 does not provide any instrumental information. In general, offering the asset for sale before all attributes are inspected is (also) optimal, if and only if the price is greater than the maximal element in the support of possible seller's valuations, given his information.

⁷ Such equilibria are supported by the buyer's most pessimistic on- and off-equilibrium beliefs where he assigns probability 1 to the state $(0, 0)$ whenever a price offer is made.

⁸ If the seller trades at some price when the value is one, he will also want to trade at that price when the value is zero. Therefore, the buyer would never pay a price above θ . However, the seller's valuation conditional on high realization for the inspected attribute exceeds θ .

the outcomes) facilitates trade. For the two-attribute case, observability of inspections guarantees ex-post efficiency as the unique equilibrium outcome for every α, β and θ . This is not true in general. We return to this case in Section 5.

4.2. The seller's equilibrium strategy

In this subsection we identify key properties of the seller's equilibrium strategy and explore their implications. We begin by characterizing the unique undominated inspection function.

Let o^α be an inspection function that assigns each non-terminal history h the attribute with the highest α_i among those attributes that were not inspected in h , i.e.,

$$o^\alpha(j_1, \dots, j_i; x_{j_1}, \dots, x_{j_i}) = \min\{k : k \in \{1, \dots, n\} \setminus \{j_1, \dots, j_i\}\}.$$

Thus, o^α induces an inspection plan according to which, at period t the seller inspects X_t , independently of the realizations along the inspection process.

Proposition 1. o^α is the unique undominated inspection function.

To get some intuition for this result note that there is a sense in which inspecting attributes in a decreasing order of their weights in the seller's payoff is akin to inspecting attributes in decreasing order of "informativeness". After the seller inspects an attribute in period t , his posterior valuation of the asset either increases or decreases, but the expected valuation is the same regardless of which attribute is inspected. Thus, the higher the weight of an attribute, the larger the spread between the posterior valuations. Hence, the unique undominated inspection function chooses at each point in time the attribute with the highest spread. This inspection plan may be interpreted as one in which at every point in time the seller has the best available information about the asset's value.

The proof of Proposition 1 proceeds in two steps. In the first step we show that for any strategy of the buyer, there is a best response consisting of a learning path that learns attributes in decreasing order of alphas, coupled with a corresponding optimal trading strategy. We show this by induction on the number of attributes that are left to inspect. In the inductive step, the seller may start with an attribute that does not have the highest alpha among the remaining attributes, but all remaining attributes will be learned in a decreasing order of alpha (this follows from the induction assumption). We then show that the seller can do better by inspecting the most important attribute first and behaving in a way that keeps the same probability of trade as in his original plan.

The second step requires us to show that there is at least one strategy for the buyer, against which the decreasing-alpha learning path is the unique best response. We show this by considering the shortest history where the decreasing-alpha path differs from other paths. Since there may be several such histories, we select the one where the seller's valuation of the good (given his information) is the lowest. This allows us to construct a strategy for the buyer for which the seller will not sell the good under the alternative learning path but will sell the good, and get a higher payoff, under the decreasing-alpha path.

Proof. The proof proceeds in two steps.

Step 1. For any strategy of the buyer \bar{p} , the seller's strategy $(o^\alpha, f(o^\alpha, \bar{p}))$ is a best response.

Let h^k be a history in which there are k attributes left to be inspected. Let $G(h^k)$ denote the continuation game that follows h^k . Given an inspection function o and a buyer's strategy \bar{p} , we denote by $o|_{h^k}$ and $(\bar{p}_{n-k}, \dots, \bar{p}_n)$ the inspection function and buyer's strategy restricted to $G(h^k)$.⁹ We denote by $f(o^\alpha|_{h^k}, (\bar{p}_{n-k}, \dots, \bar{p}_n))$ an optimal trading function in $G(h^k)$, given $(o^\alpha|_{h^k}, (\bar{p}_{n-k}, \dots, \bar{p}_n))$. We claim that for any non-terminal history h^k and for any buyer's strategy $(\bar{p}_{n-k}, \dots, \bar{p}_n)$, the seller's strategy $(o^\alpha|_{h^k}, f(o^\alpha|_{h^k}, (\bar{p}_{n-k}, \dots, \bar{p}_n)))$ is a best response in $G(h^k)$. This will prove step 1 since $(o^\alpha|_{h^n}, f(o^\alpha|_{h^n}, (\bar{p}_0, \bar{p}_1, \dots, \bar{p}_n))) = (o^\alpha, f(o^\alpha, \bar{p}))$.

The proof proceeds by induction on k , the number of remaining attributes. If only one attribute is left, the claim is trivially true. Assume that the claim is proven for some k . Consider some history h^{k+1} where there are $k+1$ attributes left to be inspected. Denote these attributes by x^1, \dots, x^{k+1} and assume that they are ordered according to the seller's order of importance (i.e., in a decreasing order of α_i). Let $(\bar{p}_{n-(k+1)}, \dots, \bar{p}_n)$ be some sequence of maximal prices that the buyer is willing to accept from now onwards. By the induction hypothesis, after h^{k+1} the seller has a best response to $(\bar{p}_{n-(k+1)}, \dots, \bar{p}_n)$ with the property that when there are only k remaining attributes, these are inspected in decreasing order of α_i independently of the realizations of the inspected attributes. Denote this best response by q_1 and let x^j be the first inspected attribute in q_1 . If $j = 1$, then the proof is complete. Assume that $j > 1$. Then q_1 induces the following order of inspection:

$$x^j, x^1, x^2, \dots, x^{j-1}, x^{j+1}, \dots, x^k, x^{k+1}. \quad (1)$$

⁹ Note that \bar{p}_{n-k} is the maximal price the buyer is willing to pay if the seller makes a sales offer at the beginning of $G(h^k)$, prior to any inspection.

We now show that there is a best response to $(\bar{p}_{n-(k+1)}, \dots, \bar{p}_n)$, which induces an order of inspection that swaps the first two attributes in the inspection order of q_1 :

$$x^1, x^j, x^2, \dots, x^{j-1}, x^{j+1}, \dots, x^k, x^{k+1}. \tag{2}$$

This will conclude the proof because after x^1 is inspected, only k attributes are left and we can apply the induction hypothesis again.

Denote the trading assignment under q_1 by $g : S \rightarrow \{hold, sell\}$, where S is the set of all sequences of $\{0, 1\}$ with length of at most $k + 1$. If $g(\{0\}) = g(\{1\})$, that is, the (immediate) trading decision after the inspection of x^j is independent of the value of x^j , the seller can sell exactly the same profiles of qualities under both inspection orders, at every period.

Assume that $g(\{0\}) = sell$ and $g(\{1\}) = hold$. Consider the seller's strategy under which he inspects the attributes according to the order (2), and his trade decisions are given by g . Denote this strategy by q_2 . The probability of trade at each period is the same under q_1 and q_2 . Thus, to compare between q_1 and q_2 we need to compare the (expected) value of the good, if it is not traded. Let $Hold(q)$ denote the set of states of nature in which the good is not traded under the strategy q . Note that we can obtain $Hold(q_2)$ from $Hold(q_1)$ by replacing every $x \in Hold(q_1)$ such that $x^j = 1$ and $x^1 = 0$ with $y \notin Hold(q_1)$ such that $y^j = 0$, $y^1 = 1$, and $y^i = x^i$ for all $i \notin \{1, j\}$. Since $\alpha^1 > \alpha^j$, the seller values y more than x . Therefore, $\mathbb{E}[V_S(x)|x \in Hold(q_1)] \leq \mathbb{E}[V_S(x)|x \in Hold(q_2)]$. This completes the proof of Step 1. ■

Step 2. Let $o \neq o^\alpha$. Then, there exist \bar{p} and h such that the seller's expected payoff after h , given o^α is strictly higher than his payoff given o .

Let h be a shortest history where $o^\alpha(h) \neq o(h)$ with the property that for each same-length history h' such that $o^\alpha(h') \neq o(h')$ the seller's expected valuation of the good after h is weakly lower than it is after h' . Let $|h|$ denote the length of h , and let $j = o(h)$ and $k = o^\alpha(h)$. Note that $k < j$. Let $V_S(h, x^k = 0)$ and $V_S(h, x^j = 0)$ be the seller's expected valuations of the good after h , given that $x^k = 0$ or $x^j = 0$, respectively. Note that $V_S(h, x^k = 0) < V_S(h, x^j = 0)$ because $\alpha_k > \alpha_j$. Let \bar{p} , be the buyer's strategy where $\bar{p}_{|h|+1} \in (V_S(h, x^k = 0), V_S(h, x^j = 0))$ and $\bar{p}_t = 0$ for all $t \neq |h| + 1$. Note that the seller's expected payoff after h under o^α is strictly higher than his payoff under o as under o the seller would never sell the good while under o^α he sells with positive probability (if after h the seller learns that $x_k = 0$), for a price that is above his expected valuation of the good. This completes the proof of Step 2. ■

Taken together, Steps 1 and 2 imply the statement of the Proposition. □

Proposition 1 implies that the seller's inspection function satisfies that after any history, and regardless of the realizations along that history, the attribute with the highest α_i among the uninspected attributes is inspected. It follows that along the equilibrium path the attributes are inspected in a decreasing order of α_i . Focusing on equilibria where the seller applies this unique undominated inspection function allows us to identify the following useful property of the seller's trade function.

Definition 2. A strategy for the seller is said to be **closed from below** if for every j and for every x_{-j} , if the seller makes a price offer at t when the state is $(x_{-j}, x_j = 1)$, then he makes a price offer at some $t' \leq t$ when the state is $(x_{-j}, x_j = 0)$.

This property is the dynamic counterpart of adverse-selection type of considerations on the part of the seller in an environment where his information evolves such that new pieces of information arrive *independently* of the information acquired in the past. In the present model, where learning is endogenous, this corresponds to cases where the inspection plan can be represented by one sequence of inspections that does not depend on past realizations (for example the inspection plan induced by o^α). To illustrate this definition consider the following simple example.

Example 1. Suppose there are only three attributes. Let (o^α, f) be a strategy where f is given in the table below.

h	f
$(1; 0)$	$\theta\beta_2 + \theta\beta_3$
$(1, 2; 1, 0)$	$\beta_1 + \theta\beta_3$
other	\emptyset

The inspection function o^α induces the order of inspection X_1, X_2, X_3 , regardless of the realizations observed. To see why this strategy is closed from below notice first that for all states (x_1, x_2, x_3) such that $x_1 = 0$, a price offer is made at the same period ($t = 1$). Now, consider the state $(x_1 = 1, x_2 = 0, x_3 = 1)$ in which there is a price offer at $t = 2$. Note that in the state $(x_1 = 1, x_2 = 0, x_3 = 0)$, the seller also offers the good for sale at $t = 2$, and if the state is $(x_1 = 0, x_2 = 0, x_3 = 1)$, the good has been offered for sale already at $t = 1$. ■

We now present two observations that will prove useful in the analysis below.

Observation 1. W.l.o.g., the seller's strategy in equilibrium is closed from below.

This result is driven by the fact that under the unique undominated inspection function σ^α , the set of inspected attributes in all periods does not depend on the observed realizations of the attributes. To see this, assume there was an equilibrium in which there are t, j and x_{-j} such that in state $(x_{-j}, x_j = 1)$ the asset is traded in period t , and in state $(x_{-j}, x_j = 0)$ the asset is not traded at any $\tau \leq t$. The set of inspected attributes at the end of period t is exactly $\{X_1, \dots, X_t\}$, and our assumptions imply that $j \leq t$ (i.e., the realized value of X_j is known by time t). Notice that the seller's expected valuation in state $(x_{-j}, x_j = 1)$, conditional on his information in period t , is strictly higher than it is in state $(x_{-j}, x_j = 0)$. Thus, conditional on the information at the end of period t in state $(x_{-j}, x_j = 0)$, the seller can strictly benefit from deviating and offering the asset for sale.

Observation 2. In any equilibrium with trade, prices are equal to at most $\theta \sum_{i=1}^n \beta_i$.

Fix an equilibrium and consider the earliest price offer (on the equilibrium path). The buyer takes into account that if the seller approached him after learning a given sequence of realizations, he would also approach him after observing a sequence in which any subset of ones in the original sequence is replaced with zeros. This brings the earliest on-equilibrium price (weakly) below the buyer's expected valuation ex ante. Finally, notice that on-equilibrium prices cannot increase over time (otherwise, it will be optimal for the seller to wait and sell at a higher price, contradicting that the "earliest" price offer is an on-equilibrium event). Thus, all equilibrium prices do not exceed the buyer's ex-ante valuation.

Notice that in order to bound the earliest on-equilibrium price, we employed a version of the standard "static" adverse selection argument for each attribute. This argument is valid in our dynamic environment precisely because, in any equilibrium, the set of inspected attributes *only depends on time* – this is a key property of the unique undominated inspection function.

4.3. Realizing gains from trade in equilibrium

Now that we understand how the seller schedules his inspections in equilibrium, we turn to identify conditions under which the buyer and seller realize gains from trade in equilibrium. We first ask whether there are equilibria in which the buyer and seller realize *all* the gains from trade, and if so, under what conditions.

We begin by introducing some useful notations. Let $I^S = \{i : \alpha_i > \beta_i\}$ denote the set of attributes that the seller values strictly more than the buyer. If the seller inspects the attributes according to σ^α , his information after t periods is represented by the vector $(x_1, \dots, x_t, X_{t+1}, \dots, X_n)$. The interpretation is that the seller knows the realizations of $\{X_1, \dots, X_t\}$ (hence the lower case letters), but does not know the realizations of $\{X_{t+1}, \dots, X_n\}$. We will sometimes summarize the seller's information at the end of time t by the t -prefix of the state of nature, $x(t) = (x_1, \dots, x_t)$. To save on notation, we suppress the conditional expectation symbol when there is no risk of confusion (for example, we write $V_S(x(t))$ instead of $\mathbb{E}[V_S(x)|x(t)]$). Let ε_i denote the realization of x such that $x_i = 1$ and $x_k = 0$ for all $k \neq i$.

Our first main result identifies two necessary and sufficient conditions for ex-post efficient equilibria. The first condition says that the attributes that are *relatively* more important to the seller than to the buyer are also the most important attributes to the seller in *absolute* terms in our environment. In other words, the weight that a seller puts on any attribute, which is more valuable to him than to the buyer, is higher than the weight he puts on any attribute, which is more valuable to the buyer. The second condition says that the gains from trade in a state where the asset has only attributes that are more important to the buyer are lower than the gains from trade in any state where the asset has only attributes that are more important to the seller. This implies that there is an ex-post efficient equilibrium only if there are no gains from trade ex-ante. Note that these conditions involve only conditions on α and β and are independent of θ .

To prove that the above conditions are necessary we do the following. We assume there is an ex-post efficient equilibrium and consider a realization in which only a single attribute, one that the seller values more than the buyer, has a value of one. Efficiency implies that in this state the seller keeps the good. Clearly, if some of the attributes that had zero valuation are replaced with ones, the seller's incentive to sell only decreases. However, if the necessary condition fails, we can find such attributes that when they have a realized value of one the good is worth more to the buyer. To prove that the conditions are sufficient we construct an equilibrium in which the seller inspects all attributes that are more important to him and sells if and only if they all have zero valuation.

Proposition 2. *There exists an ex-post efficient equilibrium if and only if there exists k such that $I^S = \{1, 2, \dots, k\}$ and $\alpha_j - \beta_j \geq \sum_{i>k} (\beta_i - \alpha_i)$ for all $j \leq k$.*

Proof. (Necessity). Assume by contradiction that there are indices $i < j$ such that $\alpha_i \leq \beta_i$ and $\alpha_j > \beta_j$. By ex-post efficiency, in state ε_j the good has to be allocated to the seller. Since $j > i$, for every $t \geq i$, $V_S(\varepsilon_j(t)) < V_S(\sum_{i \notin I^S} \varepsilon_i(t))$. Therefore, if, in an equilibrium, the seller chooses to keep the good at state ε_j , he would not sell the good in state $(\sum_{i \notin I^S} \varepsilon_i)$ in that equilibrium, in contradiction to ex-post efficiency. Thus, a necessary condition for an ex-post efficient equilibrium is that $I^S = \{1, 2, \dots, k\}$ for some k . Assume now that $\alpha_j - \beta_j < \sum_{i \notin I^S} (\beta_i - \alpha_i)$ for some $j \in I^S$. In this case, $V_B(\varepsilon_j + \sum_{i \notin I^S} \varepsilon_i) > V_S(\varepsilon_j + \sum_{i \notin I^S} \varepsilon_i)$, and so, by ex-post efficiency, the good has to be allocated to the buyer. On the other hand, ex-post efficiency implies that in state ε_j , the good has to stay with the seller. Since $V_S((\varepsilon_j + \sum_{i \notin I^S} \varepsilon_i)(t)) \geq V_S(\varepsilon_j(t))$ for all t ,

and the inequality is strict whenever $(\varepsilon_j + \sum_{i \neq j} \varepsilon_i)(t) \neq \varepsilon_j(t)$, these two trading decisions are not consistent with an equilibrium. Necessity follows.

(Sufficiency). We begin by specifying the buyer's strategy and beliefs. The buyer believes that the seller follows the unique undominated inspection function, and in case of a price offer at time t , he believes that $x_j = 0$ for all $j \leq t$. Hence, the buyer accepts any price offer up to $\theta \sum_{i>t} \beta_i$.

The seller's strategy is as follows. He inspects the attributes according to o^α . After histories consistent with o^α , the seller makes a price offer only at $t = k$ if and only if $x_j = 0$ for all $j \leq k$. Conditional on making a price offer at $t = k$, he asks for $\theta \sum_{i>k} \beta_i$. The remaining off-equilibrium components of the seller's strategy can be completed by backwards induction.

Given the buyer's beliefs, at each t he accepts price offers up to his expected value of the good. Hence, his strategy is sequentially rational.

Consider a seller who learns that $x_j = 0$ for all $j \leq k$ and sells at the end of that period. The seller's expected payoff is the same as if he would inspect all the attributes and sell at a price equal to the buyer's value of the asset – assuming the buyer knows the realized values of all the attributes. By our assumption on α and β , the seller would extract the maximal surplus, conditional on $x_j = 0$ for all $j \leq k$, since the value of the buyer is (weakly) higher than the seller's for every realization of (X_{k+1}, \dots, X_n) . Suppose instead that, he does not sell given $x_j = 0$ for all $j \leq k$. Then, in every realization of the attributes, either the seller sells at some $t > k$ for a price that is weakly lower than the buyer's expected value at time t (because the buyer believes that all realizations up to and including t are zero), or the seller keeps the good and receives a value, which is no greater than the buyer's value (since $\alpha_j \leq \beta_j$ for $j > k$). It follows that the seller cannot gain by not selling at $t = k$ when he learns that $x_j = 0$ for all $j \leq k$ in the first k periods.

We now show that if the seller follows o^α , he does not want to deviate and offer the good for sale after observing a prefix of zeros of length $l < k$. After such history, he can offer the good for sale and get $\theta \sum_{i>l} \beta_i$. Alternatively, he can wait another period, inspect X_{l+1} and sell the good at the end of next period (or never) if and only if $x_{l+1} = 0$. The seller's payoff from this behavior, given his information is

$$\theta(\alpha_{l+1} + \theta \sum_{i>l+1} \alpha_i) + (1 - \theta)\theta \sum_{i>l+1} \beta_i > \theta(\beta_{l+1} + \theta \sum_{i>l+1} \beta_i) + (1 - \theta)\theta \sum_{i>l+1} \beta_i = \theta \sum_{i>l} \beta_i$$

where the inequality follows from our assumption that $l < k$, which implies that $\alpha_l > \beta_l$ and $\sum_{i>l+1} \alpha_i > \sum_{i>l+1} \beta_i$.

To conclude, we show that if $x_j = 1$ for some $j \leq k$, the seller does not want to deviate and make a price offer at any $t \geq j$. The seller's expected utility from keeping the good, given the information at time t is at least

$$\alpha_j + \theta \sum_{i>t} \alpha_i > \theta(\alpha_j + \sum_{i>t} \alpha_i) > \theta(\beta_j + \sum_{i>t} \beta_i) > \theta \sum_{i>t} \beta_i$$

where the first inequality follows from $\theta < 1$, the second inequality follows from our assumption that for all $i < k$, $\alpha_i > \beta_i$, and $\alpha_j + \sum_{i>k} \alpha_i > \beta_j + \sum_{i>k} \beta_i$. The expression in RHS is the price the seller can get at period t if he offers the good for sale. Sufficiency follows. \square

Interestingly, when the conditions of Proposition 2 are met, ex-post efficiency is not only consistent with equilibrium but it is also the unique equilibrium outcome. To see this, notice that the seller can guarantee himself the payoff he obtained in the above ex-post efficient equilibrium by deviating to the strategy he plays in that equilibrium. Recall that under our equilibrium refinement, the buyer's most pessimistic beliefs given a (on or off-equilibrium) price offer at the end of period t are that $x_i = 0$ for all $i \leq t$. Thus at $t = k$, the buyer will accept any price offer up to $\theta \sum_{i>k} \beta_i$. Since in the equilibrium presented in the proof of Proposition 2 the whole expected surplus goes to the seller, any profile of strategies that generates a lower total surplus in expectation cannot be consistent with an equilibrium.

Another notable implication of Proposition 2 is that when ex-post efficiency is attainable in equilibrium, there is a unique period in which trade can take place. In other words, whenever an equilibrium involves multiple periods with trade on the equilibrium path, the trade-outcome of the equilibrium is not ex-post efficient. The timing of trade, however, is endogenous and it is determined so that the seller has just enough time to inspect all of the attributes he values more than the buyer, and only those attributes.

Since the result given in Proposition 2 holds for any value of θ , this stands in sharp contrast to a benchmark where the seller is fully informed. It is easy to see that if θ is sufficiently low, no gains from trade can be realized in the latter case.

While ex-post efficiency can be achieved in equilibrium, it is feasible only under somewhat narrow circumstances. A natural question that arises is, under what conditions does there exist an equilibrium with trade? The answer to this question turns out to depend on whether there is a "tail" in the seller's ranking of attributes, which is more valuable to the buyer. Fix some $j < n$ and consider the $n - j$ least important attributes to the seller (the " $n - j$ tail"). Define $R_j \equiv \sum_{i>j} (\beta_i - \alpha_i)$ as the buyer's excess valuation of $n - j$ least important attributes to the seller.

Proposition 3. *There exists an equilibrium with trade only if $R_j > 0$ for some j .*

Proof. Let e be an equilibrium with trade and assume by contradiction that for all $j \in \{0, \dots, n - 1\}$, $R_j \leq 0$. Let S the set of profiles that are sold in e and denote by $\eta_i = \frac{1}{|S|} \cdot |\{x \in S : x_i = 1\}|$ the proportion of elements $x \in S$ with $x_i = 1$. We now

show that if $i < j$, then $\eta_i \leq \eta_j$. Assume by contradiction that there exist $i < j$, such that $\eta_i > \eta_j$. Then, there is $x_{-i,-j}$ such that (1) $(x_{-i,-j}, x_i = 1, x_j = 0) \in S$ and (2) $(x_{-i,-j}, x_i = 0, x_j = 1) \notin S$. Since the seller's strategy is closed from below, from (1) we know that $(x_{-i,-j}, x_i = 0, x_j = 0) \in S$. Moreover, since $(x_{-i,-j}, x_i = 0, x_j = 1) \notin S$, it must be the case that the profile $(x_{-i,-j}, x_i = 0, x_j = 0)$ is offered for sale at the end of period $t \geq j$ (as the realization of X_j is known to the seller when $(x_{-i,-j}, x_i = 0, x_j = 0)$ is offered for sale). Applying closure from below again, we conclude that $(x_{-i,-j}, x_i = 1, x_j = 0)$ is offered for sale at $t \geq j$. Since $\alpha_i > \alpha_j$, trading decisions (1) and (2) cannot be optimal at the same time.

Therefore,

$$\begin{aligned} \mathbb{E}[V_S(x)|x \in S] - \mathbb{E}[V_B(x)|x \in S] &= \sum_{i=1}^n (\alpha_i - \beta_i) \eta_i \\ &= (\alpha_n - \beta_n) \eta_n + \sum_{i=1}^{n-1} (\alpha_i - \beta_i) \eta_i \\ &\geq \sum_{i=n-1}^n (\alpha_i - \beta_i) \eta_{n-1} + \sum_{i=1}^{n-1} (\alpha_i - \beta_i) \eta_i \\ &\vdots \\ &\geq \sum_{i=1}^n (\alpha_i - \beta_i) \eta_1 \\ &\geq 0 \end{aligned}$$

where the inequalities follow from our assumption that $R_j \leq 0$ for every $j \leq n - 1$ and from the fact that $\eta_1 \leq \dots \leq \eta_n$. It follows that $\mathbb{E}[V_S(x)|x \in S] \geq \mathbb{E}[V_B(x)|x \in S]$ implying that the trade under e does not realize any gains from trade, a contradiction. \square

The next proposition shows that if the necessary condition for an equilibrium with trade holds, then in any equilibrium the players realize some gains from trade. This follows from observing that when the “ $n - j$ tail” is positive, there cannot be an equilibrium with no trade because the seller can deviate and profitably sell the good to the buyer.

Proposition 4. *If $R_j > 0$ for some j , then any equilibrium involves trade.*

Proof. Suppose $R_j > 0$ for some j . Assume, by contradiction, that there exists an equilibrium in which no gains from trade are realized. Let j^* be the smallest index such that $R_{j^*} > 0$. Suppose the seller inspects his j^* most important attributes in the first j^* periods and if $x_i = 0$ for all $i \leq j^*$, he offers the asset for sale for a price $\theta \sum_{i>j^*} \beta_i$ (otherwise, he keeps the asset). By our refinement, the buyer would accept a price offer of $\theta \sum_{i>j^*} \beta_i$ at the end of period j^* . By the definition of j^* , the seller would make a positive profit from this sale. Since in the original equilibrium no gains from trade are realized, we have found a profitable deviation, a contradiction. \square

The following example illustrates an equilibrium with trade. It has the feature that trade may occur in multiple periods.

Example 2. Let $n = 4$, $\theta = 0.8$ and

i	1	2	3	4
α_i	0.4	0.3	0.2	0.1
β_i	0.1	0.2	0.3	0.4

Note that $\sum_{i=1}^4 \beta_i = 1$ so that there are no gains from trade ex-ante. The sufficient condition for (non-trivial) trade is satisfied because $R_1 = 0.3$. There exists an equilibrium with the following properties. The seller inspects the attributes in a decreasing order of α_i . If $x_1 = 0$, then the asset is traded in period 1 for a price of 0.72. If $x_1 = 1$ and $x_2 = 0$, then the asset is traded in period 2 for a price of 0.66.

Notice the strategic use of time in this equilibrium: The same trade-outcome can be achieved by inspecting the first two attributes during the first two periods, and offering the asset for sale only in the second period, if at least one of the inspected attributes equals zero. However, this one-sale strategy is inconsistent with equilibrium. If such an equilibrium existed, the good would be traded for $p_2 = \mathbb{E}[V_B(x)|x_1 = 0 \text{ or } x_2 = 0] \approx 0.69$. After observing $x_1 = 0$ in period 1, if the seller plays according to the suggested equilibrium, his payoff is p_2 (he will sell for sure at $t = 2$). However, by making an off-equilibrium price offer at the end of period 1, the seller can guarantee himself a payoff of $\mathbb{E}[V_B(x)|x_1 = 0] = 0.72$ (which corresponds to the highest price offer the buyer would accept under most pessimistic beliefs, consistent with our equilibrium refinement). \blacksquare

After identifying the conditions for an ex-post efficient equilibrium and the conditions under which some gains from trade can be realized in equilibrium, our third objective is to understand the structure of the most efficient equilibria for cases in which ex-post efficiency is not attainable. At present, we are unable to provide a full characterization of such equilibria for an arbitrary number of attributes and all α , β , and θ . As illustrated in [Example 2](#), an equilibrium can have complicated structures even for few attributes. In general, efficient equilibria heavily depend on the parameters of the environment and in particular, they can exhibit many periods with trade or strategic withholding of a (tradable) good.

However, we are able to characterize the efficient equilibrium of our model for the case where the difference between the fully informed seller benchmark and our gradual learning environment is probably the most prominent.

We say that an equilibrium exhibits “severe adverse selection” if the seller never sells the good if at least one of the inspected attributes has the high realization. That is, a sales offer in equilibrium means that the seller has the worst signal about the asset’s quality. Notice that if the seller is fully informed, then there cannot be an equilibrium with severe adverse selection and non-trivial trade. Hence, this type of equilibria is unique to an environment in which the seller can be partially informed. Do equilibria with severe adverse selection exist in our model? If so, are there environments in which these are the only equilibria?

Proposition 5. Let $m = \arg \max_i R_i$ and assume that $R_m > 0$.

(1) If $R_i \leq 0$ for all $i \neq m$, then for any $\theta \in (0, 1)$ there exists an equilibrium with severe adverse selection in which trade can occur only in period m .

(2) There exists $\hat{\theta} > 0$ such that for all $\theta < \hat{\theta}$ there exist only equilibria with adverse selection, and in these equilibria trade can occur only in period m .

Proof. Proof of (1). By assumption, $R_i \leq 0$ for all $i \neq m$. Consider the following profile of strategies and beliefs. The seller inspects the attributes according to o^α and makes a price offer of $\theta \sum_{i>m} \beta_i$ at time m if and only if $x_i = 0$ for all $i \leq m$. No price offer is made at any $t \neq m$. At every period t the buyer accepts any price below or equal to $\theta \sum_{i>t} \beta_i$. This is rationalized by the buyer’s most pessimistic beliefs.

We now show that the suggested profile of strategies and beliefs constitute an equilibrium. The buyer’s beliefs are consistent with our refinement and his actions are optimal given those beliefs. To verify whether the seller has a profitable deviation, suppose he sells the asset at some period $t \neq m$. By assumption, $R_t \leq 0$. It follows that

$$V_S(x(t)) \geq V_S(\mathbf{0}(t)) = \theta \sum_{i>t} \alpha_i \geq \theta \sum_{i>t} \beta_i,$$

where the RHS is the highest price the seller can get by offering the good for sale at period t , implying that the seller is weakly better off not selling the asset at t .

Assume that the seller’s information in period m is given by $(0, \dots, 0, X_{m+1}, \dots, X_n)$. The seller then prefers to sell the good and get $\theta \sum_{i>m} \beta_i$ rather than keep the good and get $\theta \sum_{i>m} \alpha_i$. To see that no other seller type at time m wants to sell note that it suffices to show that a seller of type $\tilde{x}(m) = (0, \dots, 0, 1, X_{m+1}, \dots, X_n)$ prefers to keep the good. This follows from

$$V_S(\tilde{x}(m)) > V_S(\mathbf{0}(m - 1)) = \theta \sum_{i>m-1} \alpha_i \geq \theta \sum_{i>m-1} \beta_i \geq \theta \sum_{i>m} \beta_i.$$

Proof of (2). We first find a threshold $\theta^* > 0$ such that for all $\theta < \theta^*$ any equilibrium satisfies that a price offer is made at time t if and only if $t \leq m$ and $x_i = 0$ for all $i \leq m$. By **Observation 2**, if $\theta < \theta^* = \alpha_n / \sum_{i=1}^n \beta_i$ then, in equilibrium, the seller is strictly better off keeping the asset whenever one or more attributes have a realization of one. Therefore, it must be a severe adverse-selection equilibrium.

From the definition of m it follows that if $\theta < \theta^*$, then there cannot be an equilibrium in which the asset is sold in period $t > m$. To see why, suppose there is an equilibrium with trade in period $t > m$. Consider the seller in period m after observing that $x_i = 0$ for all $i \leq m$. If he offers the asset for sale now, the buyer would accept a price of $\theta \sum_{i>m} \beta_i$, while keeping the good and following the equilibrium strategy yields an expected payoff of at most $\max_t [\theta \sum_{i>m} \alpha_i + (1 - \theta)^{t-m} \theta R_t]$. Since $R_m > R_t$, the seller strictly prefers to deviate and sell at time m for any $\theta > 0$.

We now show that for sufficiently low values of θ there does not exist an equilibrium in which the asset is sold at $t < m$. Since $R_m > R_t$ for all $t < m$ there exists $\theta^{**} \in (0, \theta^*)$ such that, for any $\theta < \theta^{**}$, and for all $t < m$, $(1 - \theta)^{m-t} R_m > R_t$. Assume $\theta < \theta^{**}$. Then, in any equilibrium with trade before period m , trade occurs after a unique history in which $x_i = 0$ for all $i \leq t$ for some $t < m$. Suppose the seller deviates from the proposed equilibrium by keeping the asset in period t , continuing to inspect attributes in decreasing order of α_i and offering the asset for sale in period m if and only if $x_i = 0$ for all $i \leq m$. Our equilibrium refinement implies that the buyer would accept an off-equilibrium price offer of $\theta \sum_{i>m} \beta_i$ in period m . Since $\theta < \theta^{**}$ the seller would strictly gain from this deviation, a contradiction.

Consider some $\theta < \theta^{**}$. From the previous paragraphs it follows that if there exists an equilibrium, then it must have the property that trade occurs after a unique history in which $x_i = 0$ for all $i \leq m$, and the price paid is $\theta \sum_{i>m} \beta_i$. To verify that such an equilibrium exists note that the pair of assessments from the proof of (1) constitutes an equilibrium. □

To illustrate part (1) of **Proposition 5** let $\alpha = (0.5, 0.3, 0.2)$ and $\beta = (0.3, 0.6, 0.1)$. Then $R_1 = 0.2$ while $R_i \leq 0$ for all $i \neq 1$. Hence, for any $\theta \in (0, 1)$ there is an equilibrium with the following properties: trade occurs only in period 1, it occurs if and only if $x_1 = 0$, and the price is $(0.7)\theta$. Note that if $\beta = (0.6, 0.3, 0.1)$, then by **Proposition 3**, there is no trade in equilibrium.

To illustrate part (2) consider **Example 2**. There is an equilibrium in which trade can occur either in period 1 if $x_1 = 0$, or in period 2 if $x_1 = 1$ and $x_2 = 0$. Indeed, condition (1) of **Proposition 5** is violated: $R_i > 0$ for all $i > 1$. However, if $\theta < 0.1$,

then all the equilibria will have the following properties: trade occurs only in period 2, it occurs if and only if $x_1 = x_2 = 0$, and the sale price is $(0.7)\theta$.

5. Observable inspections

Up until now we assumed that the inspection process is unobservable. This meant that the seller can only use time to signal the amount and type of information acquired before a price offer. It turns out that the unique undominated inspection function completely ignores the buyer’s valuations and inspects attributes only according to the seller’s order of importance.

There are many cases, however, where the process of information acquisition is observable. For instance, recall our entrepreneur example from the Introduction. Suppose that in order to estimate the demand in a given country, the firm has to send specialists to that country to conduct surveys. Even though the potential buyer will not observe the estimates collected by the firm, it may nevertheless be able to observe which markets were inspected.

In this section we explore the effects of strategic gradual learning when the buyer observes which attributes are inspected, but not their realizations. Intuitively, when the inspection is observable, the seller has an additional, more direct instrument to signal what information is available to him. The following example illustrates the seller’s strategic gradual learning when it is observable.

Example 3. Let $n = 4$, $\theta = 0.9$ and

i	1	2	3	4
α_i	0.50	0.25	0.13	0.12
β_i	0.71	0.14	0.08	0.07
$\alpha_i - \beta_i$	-0.21	0.11	0.05	0.05

Assume first that the buyer does not observe the seller’s inspections. Then by Proposition 3, there is no trade in equilibrium.

Assume next that the buyer does observe the seller’s inspections. The buyer’s beliefs about the inspected attributes are therefore pinned down, and hence, our equilibrium refinement has no bite. It can then be shown that there exists an equilibrium with the following properties. The seller begins by inspecting X_4 , his *least important* attribute. If $x_4 = 0$, he offers the asset for sale at a price equal to $\mathbb{E}[V_B(x)|x_4 = 0] = \theta(\beta_1 + \beta_2 + \beta_3) = 0.837$. If $x_4 = 1$, he continues by inspecting X_3 . If $x_3 = 0$, he offers the asset for sale at a price equal to $\mathbb{E}[V_B(x)|x_4 = 1, x_3 = 0] = \beta_4 + \theta(\beta_1 + \beta_2) = 0.835$. If $x_4 = x_3 = 1$, the seller then inspects X_2 and offers the asset for sale at a price $\mathbb{E}[V_B(x)|x_4 = x_3 = 1, x_2 = 0] = \beta_4 + \beta_3 + \theta\beta_1 = 0.789$ if $x_2 = 0$. If $x_4 = x_3 = x_2 = 1$ the seller keeps the asset. ■

Example 3 illustrates that if the buyer observes the seller’s inspections, then instead of inspecting his most important attributes, the seller strategically schedules his inspections to persuade the buyer that sale offers convey more bad news for the seller than to the buyer. This requires the seller to take both the buyer’s and his own valuations into account when scheduling his inspections. We are interested in understanding the implications of observing which inspections are carried out (but not their results) on the players’ ability to realize gains from trade.

Recall that when inspections are unobservable, ex-post efficient equilibrium exists only if the seller’s most important attributes are the ones that he values more than the buyer, and the difference between the players’ valuations of each such attribute is sufficiently large. The next result establishes that when inspections are observable, ex-post efficiency can be achieved in equilibrium for a larger set of parameters.

The main distinction from the unobservable case is that the order of inspection does not have to be decreasing in α_i . This affects both the proof of necessity and sufficiency. In the necessity part, it is no longer implied that the set of attributes that are more important for the seller are also those that the seller values the most. In the sufficiency part we now need to show that in the equilibrium candidate, the seller does not have an incentive to change the order of inspection.

Proposition 6. *There exists an ex-post efficient equilibrium if and only if $\alpha_j - \beta_j \geq \sum_{i \notin I^S} (\beta_i - \alpha_i)$ for all $j \in I^S$.*

Proof. (Necessity). Suppose there exists an ex-post efficient equilibrium. Assume, by contradiction, that there is a $j \in I^S$ such that $\alpha_j - \beta_j < \sum_{i \notin I^S} (\beta_i - \alpha_i)$. Ex-post efficiency implies that in state ε_j the good is allocated to the seller. Note that for any history of inspections h , the seller’s expected valuation in state ε_j is weakly lower than his expected valuation in state $(\varepsilon_j + \sum_{i \notin I^S} \varepsilon_i)$. Moreover, if given h , the seller can distinguish between ε_j and $(\varepsilon_j + \sum_{i \notin I^S} \varepsilon_i)$, this inequality is strict. This means that regardless of the seller’s inspection order, if he finds it optimal to keep the good at state ε_j , then he also prefers to keep the good at state $\varepsilon_j + \sum_{i \notin I^S} \varepsilon_i$. However, our negation assumption implies that $V_S(\varepsilon_j + \sum_{i \notin I^S} \varepsilon_i) < V_B(\varepsilon_j + \sum_{i \notin I^S} \varepsilon_i)$ in contradiction to the premise that an ex-post efficient equilibrium exists.

(Sufficiency). We construct an ex-post efficient equilibrium. We begin with the buyer’s strategy and beliefs. Whenever a price offer is made, the buyer believes that all inspected attributes had a realized value of zero. Thus, whenever a price offer is made, the buyer is willing to accept any price offer up to $\theta \sum_{i \notin INS} \beta_i$, where INS denotes the set of inspected attributes.

The seller’s plan of behavior is to inspect the attributes in a decreasing order of $(\alpha_i - \beta_i)$, make a price offer at the end of time $t = |I^S|$ if and only if all of the attributes inspected equal zero, and not to offer the good for sale at any $t \neq |I^S|$. To

obtain a full description of the seller’s strategy, the plan of behavior should be completed with a description of the seller’s actions after off-equilibrium histories. This can be done by backwards induction, given the buyer’s strategy.

Given the seller’s strategy, the buyer’s strategy is sequentially rational. We now show that the seller does not have a profitable deviation given the buyer’s strategy.

First, we show that after observing a realization of zeros for all $x_j, j \in I^S$, the seller prefers to sell the good immediately (for a price equal to the buyer’s expected valuation). By doing so, the seller obtains a payoff equal to what he would get if he would inspect all the attributes and sell at a price equal to the buyer’s value of the asset – assuming the buyer knows the realized values of all the attributes. By our assumptions on α and β , the seller would extract the maximal surplus, conditional on $x_j = 0$ for all $j \in I^S$, since the value of the buyer is (weakly) higher than the seller’s for every realization of attributes not in I^S . Suppose instead that the seller does not sell given $x_j = 0$ for all $j \in I^S$. Then in every realization of the attributes, he either sells at some $t > |I^S|$ for a price that is weakly lower than the buyer’s expected value at time t (because the buyer believes that all realizations up to and including t are zero), or the seller keeps the good and receives a value, which is no greater than the buyer’s value (since $\alpha_j \leq \beta_j$ for $j \notin I^S$). It follows that the seller cannot gain by not selling at $t = |I^S|$ when he learns that $x_j = 0$ for all $j \in I^S$.

We next show that the seller has no incentive to sell before inspecting all the attributes in I^S . Assume by contradiction that after inspecting only a subset of I^S the seller prefers to deviate and offer the good for sale. If such a deviation is profitable, then it is necessarily profitable when all the inspected attributes in I^S have a realized value of zero. Let t be the period in which the deviation occurs and denote by INS_t the set of inspected attributes by period t . The seller’s payoff from keeping the good at the end of period t is at least as high as his payoff from the following plan: Inspect the attribute $X_j, j = \arg \max_{i \notin INS_t} (\alpha_i - \beta_i)$, offer the good for sale at $t + 1$ if $x_j = 0$, and keep the good if $x_j = 1$. The seller’s expected payoff from this plan, given his information at the end of time t , is

$$\begin{aligned} & (1 - \theta) \left(\theta \sum_{i \notin INS_t} \beta_i - \theta \beta_j \right) + \theta (\alpha_j + \theta \sum_{i \notin INS_t} \alpha_i - \theta \alpha_j) > \\ & > (1 - \theta) \left(\theta \sum_{i \notin INS_t} \beta_i - \theta \beta_j \right) + \theta (\beta_j + \theta \sum_{i \notin INS_t} \beta_i - \theta \beta_j) = \\ & = \theta \sum_{i \notin INS_t} \beta_i, \end{aligned}$$

where the inequality follows from our assumption that $t < |I^S|$, which, in turn, implies that $\alpha_j > \beta_j$, and $\alpha_j - \beta_j \geq \sum_{i \notin I^S} (\beta_i - \alpha_i)$. The expression on the RHS corresponds to the seller’s payoff from selling the good at the end of period t , and the inequality shows that he prefers *not* to sell at the end of that period.

To conclude, we note that the seller never wants to sell the good after observing $x_j = 1$ for some $j \in I^S$. To see this, let j be the first attribute such that $x_j = 1$, and consider the seller’s payoff from deviating to a plan according to which he continues inspection (according to some inspection order) and offers the good for sale after some history of realizations. Denote by INS the set of attributes inspected when the seller makes a price offer. The seller’s payoff from selling the good after this history is $\theta \sum_{i \notin INS} \beta_i$. However, his payoff from keeping the good after this history is at least $\alpha_j + \theta \sum_{i \notin INS} \alpha_i$, which, by our assumptions on α and β , is greater than $\sum_{i \notin INS} \beta_i$. Sufficiency follows. \square

From the proof of Proposition 6 it follows that in any ex-post efficient equilibrium the seller sells the good if and only if every attribute that he values more than the buyer has zero value and the sale price is the same. The order of inspection can vary (in the proof we specify a particular order of inspections but there could be others). Thus, there is a *unique period* in which trade can occur in an ex-post efficient equilibrium.

When ex-post efficiency is unattainable in equilibrium, can parties realize some of the gains from trade?

Proposition 7. *Positive gains from trade are realized in every equilibrium.*

Proof. Assume, by contradiction, that there exists an equilibrium in which no gains from trade are realized. By assumption, $I^S \neq \{1, 2, \dots, n\}$. Consider the following deviation by the seller. He inspects all the attributes in I^S and offers the good for sale for the price of $\theta \sum_{j \notin I^S} \beta_j$ if and only if $\theta \sum_{j \notin I^S} \beta_j > \mathbb{E}[V_S(x)|h]$ where $h = (j_1, \dots, j_i; x_{j_1}, \dots, x_{j_i})$ such that $I^S = \{j_1, \dots, j_i\}$. Note that the seller strictly prefers to sell after the history $h = (j_1, \dots, j_i : 0, \dots, 0)$. Since the buyer observes the inspections, regardless of his beliefs, he would accept this price offer since his most pessimistic beliefs after observing this sequence of inspections is that all inspected attributes have a realization of zero. We therefore found a profitable deviation for the seller, a contradiction. \square

Proposition 7 stands in contrast with our finding in the previous section that when inspections are unobservable, there is no trade in equilibrium if $R_j \leq 0$ for all j .

6. Concluding remarks

Commitment. The seller in our model has the discretion to decide each period whether to continue inspecting and in which order. This raises the question of whether more gains from trade can be realized if the seller were able to commit to an inspection plan. For example, the seller could commit to an inspection order but keep the discretion of when to stop and make a price offer (“order commitment”). Alternatively, he could commit to inspect a specific set of attributes before making a price offer (“set commitment”).

It turns out that these forms of commitment induce the same necessary and sufficient condition for ex-post efficiency as in the case of observable inspection without commitment. To see that the necessary condition for ex-post efficiency under set or order commitment is the same as in the (no-commitment) observable case, recall the proof of Proposition 6. There we showed that if the condition is violated, we can find two states, x_S and x_B , such that: (i) ex-post efficiency implies trade in x_B and no trade in x_S , and (ii) regardless of what the seller knows about the attributes, if he finds it optimal to keep the good in x_S , he would also keep the good in x_B , which contradicts ex-post efficiency. To see that the same condition is also sufficient under set and order commitment, consider the strategy of committing to learn all the attributes that the seller values more than the buyer (the attributes in I^S), or alternatively committing to learn these attributes before others (regardless of the order), and selling if and only if all the attributes in I^S had a value of zero. These strategies achieve ex-post efficiency (given the buyer’s strategy that we postulate in the proof). The arguments in the proof of Proposition 6 establish that the seller cannot profit from either selling the good when one of the attributes in I^S has a value of one, or from selling the good before inspecting all the attributes in I^S , or from not selling after he learned that all the attributes in I^S are zero.

General utility functions and non-identically distributed attributes. A key feature of our environment is that the contribution of an attribute to a player’s payoff is independent of the other attributes. Suppose we relaxed this assumption and allowed a player’s payoff from a vector of realized attributes x to be some general function $u(x)$, which is strictly increasing in each of the attributes. This adds a new level of difficulty to the analysis since there is no longer a unique undominated inspection function. Moreover, the optimal learning path may depend on the realization of inspected attributes. To see this, consider the following simple example.

Suppose there are three attributes, $\theta = \frac{1}{2}$, and $V_S(x) = 3x_1 + x_2 + x_3 + 10x_2x_3$. The seller can sell the asset for the price 3, but only if it is offered for sale no later than at the end of period 2. A simple but tedious verification reveals that the seller’s best inspection plan is the following. Start with X_3 in period 1. If $x_3 = 0$, inspect X_1 , but if $x_3 = 1$, inspect X_2 . At the end of period 2, the seller sells the asset whenever the attribute inspected in period 2 equals zero. Every inspection path that can be written as a simple (unconditional) order of attributes is strictly inferior.

An additional difficulty is that we can no longer say that one attribute is more important than another, or that a certain attribute is more important to the seller than the buyer. Consequently, it becomes more challenging to characterize the necessary and sufficient conditions for ex-post efficiency. If there is an ex-post efficient equilibrium with a general monotonic utility function, then it is still the case that we cannot find a pair of states, x_S and x_B such that (i) the asset should stay with the seller in x_S , (ii) the asset is more valuable to the seller in x_B than in x_S , but (iii) it is efficient to trade the asset in x_B . However, finding a sufficient condition remains an open question. Our proof of sufficiency (in both the unobservable and observable case) relies on the separability in attributes.

Another key property of our model is that the attributes are identically distributed, i.e., $\theta_i = \theta_j = \theta$. Suppose we relaxed this assumption, so that one attribute may have a higher weight than another but have a lower probability of occurring. This would also introduce the complication that there would no longer be a unique undominated learning function when inspections are unobservable. The optimal learning path would now depend on the buyer’s strategy.

To give a simple and stark example of this, suppose there are only two attributes with equal weight but different probabilities¹⁰:

	X_1	X_2
θ_i	0.9	0.5
α_1	10	10

Suppose the buyer is willing to pay p only in first period, but is not willing to pay anything in the second period. The seller’s ex-ante expected value of the asset is 14. If the seller inspects X_1 first, then with probability 0.9 the posterior value goes up to 15 and with probability 0.1 it goes down to 5. If X_2 is inspected first, then with probability 0.5 posterior value goes up to 19 and with probability 0.5 it goes down to 9. It is then easy to show that the optimal inspection plan depends on the value of p . If $p < 5$ then obviously the seller never sells as this is below the lowest valuation the seller can have at the end of the first period. Similarly, if $p > 19$ then the seller sells immediately since the buyer is willing to pay more than the highest valuation the seller can have after inspection. However, when p is between these bounds, the optimal attribute to inspect would depend on the value of p . If $5 \leq p < 10$ the seller prefers to check X_1 first (at the threshold 10 the seller is indifferent between inspecting X_1 or X_2 , i.e., $p = 10$ is the solution to $(0.1)(p - 5) = (0.5)(p - 9)$). If $10 \leq p < 19$ the

¹⁰ While this specification does not fit our original model it allows us to show the difference from having equal probabilities and different weights.

seller prefers to check X_2 first. Thus, while there exists a unique undominated inspection function when attributes have the same probabilities but different weights, this is no longer true when there are different probabilities but equal weights.

However, the likelihood of each attribute does not affect the analysis when inspections are observable. A careful examination of the proofs of Propositions 6 and 7 reveals that all of our arguments straightforwardly extend to the case in which $\theta_i \neq \theta_j$.

Future directions. Although this paper focuses on a bilateral trade environment we think the question of strategic gradual acquisition of information is relevant in many other environments where it has yet to be explored. For example, a job candidate's decision of how to sequence job interviews can have strategic implications. In each interview the candidate gets a signal about his ability and prospective employers may make inferences about a candidate's information about his quality from the timing of the interview (relative to the period in which the candidate was available). Another example is a mechanism-design environment in which agents are not endowed with private information at the outset, but rather collect this information gradually over time. The agents then decide what to learn each period and when to send their report to the planner, who in turn, has to decide how to respond to the content and timing of reports.

References

- Bergemann, Dirk, Välimäki, Juuso, 2002. Information acquisition and efficient mechanism design. *Econometrica* 70, 1007–1033.
- Bolton, Patrick, Faure-Grimaud, Antoine, 2010. Satisficing contracts. *Rev. Econ. Stud.* 77, 937–971.
- Branco, Fernando, Sun, Monic, Villas-Boas, J. Miguel, 2012. Optimal search for product information. *Manage. Sci.* 58, 2037–2056.
- Choi, Michael, 2016. Imperfect Information Transmission and Adverse Selection in Asset Markets. Working Paper. University of Iowa.
- Dang, Tri Vi, 2008. Bargaining with endogenous information. *J. Econ. Theory* 140, 339–354.
- Frug, Alexander, 2016. Strategic Gradual Learning and Information Transmission. Working Paper. Universitat Pompeu Fabra.
- Guttman, Ilan, Kremer, Ilan, Skrzypacz, Andrzej, 2014. Not only what but also when: a theory of dynamic voluntary disclosure. *Amer. Econ. Rev.* 104, 2400–2420.
- Hwang, Ilwoo, 2016. Collapsing Confidence: Dynamic Trading with Developing Adverse Selection. Working Paper. University of Miami.
- Kaya, Ayça, Kim, Kyungmin, 2015. Trading Dynamics with Private Buyer Signals in the Market for Lemons. Working Paper. University of Miami.
- Klabjan, Diego, Olszewski, Wojciech, Wolinsky, Asher, 2014. Attributes. *Games Econ. Behav.* 88, 190–206.
- Levin, Jonathan, 2001. Information and the market for lemons. *RAND J. Econ.* 32, 657–666.
- Persico, Nicola, 2000. Information acquisition in auctions. *Econometrica* 68, 135–148.
- Taylor, Curtis R., 1999. Time-on-the-market as a sign of quality. *Rev. Econ. Stud.* 66, 555–578.