

Learning Deterrence vs. Learning Encouragement: Optimal Pricing and Return Policy

Qianjun Lyu*

First version: July 2020

July 2, 2021

Abstract

This paper studies the optimal selling mechanism when an uninformed buyer can sequentially and privately acquire costly information about his valuation of a product. The seller designs the mechanism to affect the buyer's benefit from learning and thereby controls the learning process. Our main result shows that it is always sub-optimal to induce partial learning, and the optimal mechanism either encourages full learning or deters learning. Specifically, the seller optimally encourages full learning if the buyer is relatively uninformed and optimistic. Otherwise, the optimal mechanism prevents learning, which benefits both the buyer and the seller.

Key words: buyer learning, return policy, sequential information acquisition, information design, screening.

*Faculty of Business and Economics, University of Hong Kong, lyu2014@hku.hk. I would like to thank my advisor Wing Suen for insightful and continual guidance. For valuable suggestions, I would also like to thank Bard Harstad, Jin Li, Xianwen Shi, Balázs Szentes, Dong Wei and Yimeng Zhang, Chen Zhao. All errors are mine.

1 Introduction

Nowadays, consumers can gather information about products in a variety of ways, from visiting retail stores to experience the product in person to checking the Internet for customer reviews and professional evaluations. In this sense, a seller may not be able to fully control a buyer's information sources. However, by designing a selling mechanism to affect the buyer's benefit from learning, the seller can indirectly control the buyer's endogenous learning. For example, the seller can induce consumers to learn more about the product by offering a generous return policy, allowing them to try out the product and return it if they are not satisfied. Alternatively, the seller may set a non-refundable low price to encourage immediate purchase and consumption without consumers undertaking costly information acquisition. This paper studies the revenue-maximizing selling mechanism for a seller who anticipates that the buyer privately performs sequential learning about his valuation of a product.

A selling mechanism specifies a product's selling price and its return policy. In general, return policies vary across platforms and product categories. Free-return policies, such as a "no-question-asked" full refund or a limited-period free trial, and no-return policies (where the initial payment is non-refundable), such as "final sales," are the most common. However, there are other policy types. For example, airlines usually charge a fixed fee for a ticket refund. Fashion retail platforms, such as Farfetch, offer free returns and free home collections, while others, like Luisaviaroma, offer free returns only in the form of store credit. More surprisingly, e-commerce retailers such as Amazon may issue a refund without requiring a product return,¹ determined randomly at the time of purchase, and specified as "Amazon may determine that a refund can be issued without requiring a return." Assuming quasi-linear consumer preferences, we can represent a return policy as the probability that the seller requires the product to be returned and the (expected) amount of the refund.²

For concreteness, consider a seller (she) selling one unit of an indivisible good to a

¹However, the refund may be smaller than the selling price.

²The refund can be issued independently of whether a return is required.

buyer (he), who is initially uninformed about his true product valuation, which could be either high or low, with $v_h > v_l > 0$. We interpret this uncertainty in product valuation as coming from match-specific factors so that the seller is symmetrically uninformed initially. The buyer’s prior belief is common knowledge, but he can privately update his belief by sequential experimentation. Specifically, by exerting costly effort, e.g., spending time acquiring information, he learns good news about the product at rate λ if his true valuation is high, whereas no news arrives if his true valuation is low. We focus on scenarios in which efficiency requires trade with probability one, so that the seller values the product less than a low-value buyer does and always prefers to sell.³ At the outset, the seller commits to a selling mechanism. The buyer then decides how much information to acquire and makes his purchase/return decision based on the information outcomes.

Our model builds on the exponential-bandit framework (Keller, Rady and Cripps (2005)), which grants us significant tractability and allows us to recast this mechanism design problem as an information design problem. Loosely speaking, we can characterize the selling mechanism as a function of the buyer’s posterior belief and determine the seller-optimal amount of buyer learning. In essence, the seller maximizes her expected revenue over the distribution of the buyer’s stopping beliefs.

To further illustrate this, we argue that for a fixed selling price, the buyer’s expected trading surplus under any optimal selling mechanism is *the same* as if the mechanism prohibits return. The buyer’s information acquisition under a no-return mechanism is a standard stopping time problem,⁴ and the high-value buyer’s net surplus is a sufficient statistics to characterize the buyer’s value function and his learning process. Therefore, the seller first chooses the selling price to influence the value

³In reality, the seller’s reservation value might depend on reselling after the buyer returns the product. However, while proceeding with a resell, the seller needs to cover additional shipping/handling costs and labor costs. Moreover, a return implies a lower posterior belief, and therefore indicates a smaller benefit from reselling. In section 7.3, we discuss scenarios in which the product is more valuable to the seller than to a low-value buyer.

⁴Essentially, given a no-return mechanism, the buyer performs all learning before purchase. If, during his learning process, good news arrives, his belief jumps to one and he purchases the product. If no news arrives after a sufficient amount of time, he walks away.

function that the buyer can attain through learning, and then designs a return policy to truncate the buyer’s sequential learning to induce a more flexible stopping belief. Essentially, the return policy compensates the buyer with opportunity information rent—the continuation value—that he could enjoy if he continued to learn, and then enforces earlier stopping.⁵ For a given price, incorporating the buyer’s optimality (also known as smooth-pasting) to stop learning and request a return, we can identify the return policy as a function of the buyer’s stopping belief. In other words, we can explicitly characterize the direct mechanism to screen the buyer with interim information.

Before elaborating further on the results, it is useful to introduce three types of selling mechanisms: (1) Learning Deterrence, (2) Stochastic Return, and (3) Free Return.

Learning Deterrence refers to a special case of a no-return mechanism such that the selling price is designed to make the buyer indifferent between learning and purchasing the item immediately. In other words, to deter learning, the seller compensates the buyer with the continuation value that he can attain through learning. Given the seller’s preferred tie-breaking rule, learning is deterred and the allocation is efficient. It has been shown that the selling price of Learning Deterrence is the lowest among the three mechanisms,⁶ which gives the buyer the highest attainable trading surplus.

Stochastic Return refers to selling mechanisms with a random return policy—when a buyer requests a return, he can still keep the item with positive probability while receiving a refund, as in the Amazon example mentioned earlier. Under Stochastic Return, the seller can induce some particular level of *partial learning* in the sense that the buyer optimally stops learning and requests a return even though the continuation value from learning is still *positive*. Meanwhile, the trading surplus is positive upon return, which ensures the seller a positive return revenue.

⁵Whenever the buyer feels indifferent, let us assume that he breaks the tie in the seller’s favor.

⁶Charging a lower price to prevent the buyer from searching or learning is also documented in [Armstrong and Zhou \(2015\)](#) and [Pease \(2020\)](#).

Free Return allows the buyer to opt-out freely, hence the seller induces *full learning* in the sense that the buyer only requests a return when the continuation value from continuing to learn becomes *zero*. The allocation is ex-ante inefficient and the seller obtains zero safe revenue.

In general, the above three types of mechanisms are sufficient to induce all feasible learning dynamics—the buyer’s posterior distributions—under our assumption of exponential experimentation. Learning Deterrence and Free Return can be interpreted as the opposite limits of Stochastic Return in terms of how much information the buyer learns. If the price is exogenously chosen, the optimal return policy could take the form of any one of these three mechanisms. However, we show that Stochastic Return is always sub-optimal when the selling price is optimally chosen. That is, the revenue-maximizing mechanism is either Learning Deterrence or Free Return.

Given this, we can fully characterize the optimal selling mechanism which depends on the buyer’s prior belief—a measure of both his expected valuation of the product and the degree to which he is informed. Being less informed ex ante indicates greater benefits from information acquisition. Hence, to deter learning, the seller has to compensate the buyer with a larger opportunity information rent. Therefore, when the prior belief is more uncertain, Learning Deterrence becomes less attractive to the seller even though it generates efficient allocation. Conversely, under the optimal Free Return mechanism, the seller encourages the buyer to learn and thereby avoids paying the information rent. Although Free Return causes allocation inefficiency when the buyer does not consume the product upon learning, the likelihood of this event is low when the initial belief is high. Moreover, a high prior belief indicates a higher willingness to pay, which also results in a higher price. The seller therefore optimally offers Free Return when the buyer is initially more uncertain but also more optimistic. Otherwise, Learning Deterrence is optimal. This implies a tension between surplus maximization and revenue maximization.

Furthermore, as learning becomes more efficient, e.g., the cost of learning is lower, the set of prior beliefs for which the seller optimally selects Free Return expands, because to deter learning would entail setting a progressively lower selling price.

When the cost of learning converges to zero, the buyer approaches perfect information. Therefore, with Free Return, the seller sets the price arbitrarily close to v_h and lets go of buyers who almost definitely have a low valuation. With Learning Deterrence, she sets the price arbitrarily close to v_l with the intention of covering the whole market. The ratio $v_l=v_h$ determines the cut-off value of priors at which the seller is indifferent between the two mechanisms. This is reminiscent of the distinction between catering to the market or a niche market.

We discuss several extensions of the main model. First, we consider scenarios where the seller cannot freely adjust the selling price. In particular, we focus on the case of an exogenous price or a regulated price cap. We show that Stochastic Return, which induces partial learning, can be optimal in such scenarios. Next, we consider cases in which the seller values the product higher than a low-valuation buyer does. Consequently, the seller intrinsically prefers to encourage learning as return creates efficiency. This shrinks the set of priors that supports Learning Deterrence. We also study the situations in which the learning process is more efficient after purchase. In this case, the seller optimally charges a cancellation fee to extract the extra information rent that the buyer obtains from post-transaction learning, although such a cancellation fee is irrelevant in the baseline model. Lastly, we argue that the return mechanism we discussed in the baseline model is without loss of generality under a more general framework.

Related literature. There is a growing literature on mechanism design incorporating information as part of the seller’s optimal choice. Considering price discrimination, [Li and Shi \(2017\)](#) allow the seller to disclose different additional information to different types of buyers. They show that partial and discriminatory disclosure weakly dominates full disclosure in terms of the seller’s revenue. [Guo, Li and Shi \(2020\)](#) then characterize the property of optimal discriminatory disclosure. [Wei and Green \(2020\)](#) also investigate price and information discrimination, but allow the buyer to decide whether to purchase after obtaining the information outcomes, which guarantees ex-post individual rationality. [Johnson and Myatt \(2006\)](#) introduce rotations of demand curves to capture the dispersion of consumer valuations

and discuss how seller profits change with the level of dispersion. [Bergemann and Pesendorfer \(2007\)](#) allow the buyer to acquire information whose accuracy is controlled by the seller. Instead of allowing the seller to optimally restrict the buyer’s learning process, [Roesler and Szentes \(2017\)](#) allow the buyer to acquire costless information anticipating its impact on the seller’s pricing decision, and they identify the buyer-optimal information structure. [Ravid, Roesler and Szentes \(2019\)](#) consider the same scenario but let both seller and buyer move simultaneously so as to discuss the equilibrium outcomes.

Our paper adopts the perspective of mechanism design, wherein the seller optimally designs the mechanism anticipating its impact on the buyer’s endogenous information acquisition. For example, [Shi \(2012\)](#) adopts rotational-ordered information technology, and in the case of a single buyer, the price change affects the buyer’s incentive to acquire information and thereby changes the transaction probability. However, in our paper, by introducing a return policy, the seller can sustain the same transaction probability across different prices. [Mensch \(2020\)](#) studies flexible information acquisition, with cost as the expected difference in a posterior-separable measure of uncertainty. The paper characterizes the set of implementable mechanisms to screen the buyer with different interim information and thereby recasts the problem as Bayesian persuasion. We adopt a similar translation when discussing learning encouragement; however, our exponential bandit specification with additive time cost allows us to analyze how the seller’s optimal mechanism varies with the buyer’s initial belief, which cannot otherwise be accommodated in the flexible information cost framework.⁷ In terms of sequential buyer learning,⁸ [Lang \(2019\)](#) and [Pease \(2020\)](#) investigate the seller’s optimal pricing when the buyer can dynamically acquire information before purchase.

Our work is closely related to studies that incorporate both price and refund into the

⁷In our model, the cost of the same Blackwell experiment is the same for different prior beliefs, which is not true for flexible information. There does not exist a unified measure of uncertainty, regardless of the prior beliefs, that can represent the additive time cost of Poisson signals: see Appendix A of [Mensch \(2020\)](#) and [Pomatto, Strack and Tamuz \(2019\)](#).

⁸See [Bonatti \(2011\)](#), [Bergemann and Valimaki \(2000\)](#) and [Bergemann and Valimaki \(1996\)](#).

seller's choice to influence the buyer's private learning, see [Matthews and Persico \(2007\)](#). [Daley, Geelen and Green \(2021\)](#) consider dynamic learning process and discuss due diligence in M&A, wherein after the acquirer (buyer) agrees the price with the target firm (seller), he can collect information about the firm and has the option not to execute the contract, which can be interpreted as a return.⁹ In their model, the seller can set a price to implement the socially optimal learning outcome and charge an upfront transfer to extract the entire surplus. However, as in our paper, the buyer can always gather information before purchase, his information rent from pre-transaction learning imposes a lower bound of his trading surplus, which cannot be captured by the seller via an upfront transfer. Moreover, we allow the buyer to keep the item with positive probability upon receiving a refund, which facilitates partial learning, i.e., the buyer stops learning even when there is still a positive value of information. This is also a distinction to the earlier works on return mechanism.

The literature offers several complementary economic rationales for excess refund puzzles. [Che \(1996\)](#) shows that the seller optimally insures risk-averse buyers by offering a generous refund. [Inderst and Ottaviani \(2013\)](#), [Shieh \(1996\)](#) and [Inderst and Tirosch \(2015\)](#) discuss the role of refund as a signaling device to guarantee credible sales talk, product quality, and personal fit. [Courty and Li \(2000\)](#) and [Escobari and Jindapon \(2014\)](#) use the refund contracts to screen the buyers with different ex-ante information. In situations where the seller is unsure about the buyer's private information prior to purchase, [Hinnosaar and Kawai \(2020\)](#) find that generous refund can hedge the seller against uncertainty, and they find the robust refund mechanism if Nature designs the buyer's information structure against the seller. In contrast, in our model, the seller uses the mechanism to screen the buyers with different interim information, which is the outcome of the buyer's best response to the selling mechanism.

⁹Though [Daley, Geelen and Green \(2021\)](#) assume that bidders make competing price offers to the seller, the equilibrium price maximizes the seller's expected payoff.

2 Model

A seller sells one unit of indivisible good to a risk-neutral buyer. The buyer is initially uninformed about his true product valuation, which is either high or low, $v_h > v_l > 0$. The seller is symmetrically uninformed, with θ_0 being the common prior belief that the product value is high. We use θ to represent the buyer's posterior belief after learning, and sometimes call this the buyer's type. A type-buyer's expected value of the product is $E(v|\theta) := \theta v_h + (1 - \theta)v_l$. Note that the buyer's type evolves over time depending on the learning process; we use t to denote time and write $\theta(t)$ when needed. There is no cost of production or return and we normalize the seller's product valuation to 0.¹⁰ Therefore, efficiency requires trade with probability one.

The seller commits to a selling mechanism, which specifies (1) a selling price $t_b \geq 0$, which is the transfer made from the buyer to the seller at the time of purchase; and (2) a return policy that describes the probability that the buyer is required to return the item and the expected refund that he obtains upon requesting a return. Given that the buyer is assumed to be risk-neutral, only the expected value of the refund at the time of the request matters. Formally, we use $(x_r; t_r)$ to denote a return policy, where $x_r \in [0; 1]$ is the probability that the buyer is allowed to *keep* the item after requesting a return and $t_r \in [0; t_b]$ is the amount of net expected transfer made from buyer to seller if the buyer requests a return. A typical selling mechanism is characterized by $f(t_b; (x_r; t_r))g$.¹¹ Under this mechanism, the buyer pays the selling price t_b at the time of purchase. If he requests a return, he has to surrender the product with probability $1 - x_r$, whereas he can retain the product with probability x_r . The refund can either be paid only if he is asked to surrender the product, or

¹⁰One can view $v_h > v_l > 0$ as the normalized buyer's valuation after extracting the seller's valuation. That is, we focus on scenarios in which the seller values the product less than a low-value buyer does, and always prefers to sell. In section 7.2, we discuss the implications of the seller valuing the product more than a low-value buyer does.

¹¹The reader can also interpret this as a binary menu that specifies two (allocation probability, transfer) pairs. The selling price is actually a simplification of $(x_b = 1; t_b)$. Assuming such a binary menu is without loss of generality. Further discussion can be found in section 7.4.

it can be paid regardless of whether the product actually returns to the seller. The expected amount of the refund at the time of making the return request is $t_b - t_r$.

Given this notation, a No Return mechanism can be represented as $f t_b; (1; t_b)g$. That is, if the buyer purchases the product, he pays the selling price and cannot return it for a refund. Free Return can be represented as $f t_b; (0; 0)g$, so that the buyer is allowed to walk away freely after a purchase. Stochastic Return is defined as $f t_b; (x_r; t_r)g$ with $x_r \geq (0; 1)$, so that the buyer is allowed to keep the item with strictly positive probability even upon obtaining a refund. Without loss of generality, we assume $v_h - t_b > v_h x_r - t_r$, which means that a high-value buyer purchases the item without requesting a further return. We assume that neither party discounts over time as the buyer's entry time is unobservable to the seller.¹² The buyer's outside option is normalized to zero.

A type- i buyer's payoff is realized when he consumes the item. If so, he cannot request a return, regardless of the return policy. In particular, a type- i buyer obtains expected utility $E(v_i) - t_b$ if he purchases the item without requesting a return, or $E(v_i) x_r - t_r$ if he requests a return. Let \mathbf{B} be the indicator function for whether a purchase has occurred up to and *including* time t . Hence, the time of purchase is $t_b = \min t : \mathbf{B} = 1$. Analogously, \mathbf{R} denotes the indicator function for whether a return has occurred up until time t , and the time that the buyer requests a return is $t_r = \min t : \mathbf{R} = 1$. Naturally, $t_r \geq t_b$. The seller's revenue is expressed as follows:

$$= E \left[\int_0^T t_b d\mathbf{B} + (t_r - t_b) d\mathbf{R} \right] : \quad (1)$$

In this paper, we adopt the exponential bandit framework. In particular, if the buyer decides to learn, he needs to pay a flow cost k . Good news arrives at Poisson rate λ if his true valuation is v_h and no news arrives if his true valuation is v_l . The buyer's belief evolves according to the following law of motion if no Poisson jump occurs:

$$\dot{\theta}(t) = -\lambda \theta(t) (1 - \theta(t)) < 0:$$

¹²We allow the buyer to acquire information before purchase. Usually, after a product is released, the seller does not know when the buyer starts to notice the product.

Otherwise, if good news arrives, his belief jumps to one. We assume the information process to be the same both before and after the buyer purchases the item. That is, the buyer does not gain additional information rent from purchasing the item.¹³ An equivalent but weaker assumption is that learning efficiency $k=$ remains constant both before and after purchase. In this case, a selling mechanism $f_{t_b};(x_r; t_r)g$ described above is without loss of generality. In section 7.3, we discuss the scenarios where learning is more efficient after the transaction, such that the buyer gains more information rent from purchasing the item, but the seller can extract the extra information rent by charging a cancellation fee as a complementary instrument to the return policy. Nevertheless, the key assumption is that the buyer has the option to acquire information before purchasing the item. This imposes a lower bound of the buyer’s ex-ante trading surplus, which in turn determines the seller’s incentives. We elaborate on this in the next section.

3 A No Return Benchmark

Under a No Return mechanism $f_{t_b};(1; t_b)g$, the buyer performs all learning before purchasing the item. His optimal learning under No Return (characterized in Proposition 1) serves as a building block for our derivation of subsequent results. The value of experimentation depends on the “prize” received when good news arrives, which in this context is the consumer surplus if the buyer turns out to be a high-value buyer. We denote this surplus by $s = v_h - t_b$, and use $V^0(\cdot; s)$ to represent the buyer’s value function when he faces a No Return mechanism with a selling price $t_b = v_h - s$. The value function is characterized by the Bellman equation:

$$V^0(\cdot; s) = \max \{ 0; E(v^j(\cdot)) - (v_h - s); kd + (\cdot) d s + (1 - (\cdot) d) V^0(\cdot + d; s) \} \tag{2}$$

¹³For subscription services, like Amazon Prime and YouTube Premium, the product information attainable by the buyer remains almost the same after purchase. In M&A, the acquirer collects information about the target firm both before and after agreeing the terms with the seller. Moreover, with the spreading of “We Media”, more instructive information is becoming attainable online even for normal commercial goods, and consumers are still willing to gather information online after purchase.

At time t , the buyer can walk away with payoff 0, or purchase the item to get the expected consumption payoff $E(v_j | \theta_t) - (v_h - s)$. If he continues to learn for an interval of time d then, with probability $(1 - \lambda)d$, good news arrives and he purchases the item to obtain net surplus s ; with the remaining probability, no news arrives and his belief decreases to $(\theta_t - d)$ with corresponding value $V^0((\theta_t - d); s)$. Conditional on learning, the Bellman equation leads to this differential equation:

$$(1 - \lambda) \frac{\partial V(\theta; s)}{\partial \theta} + V(\theta; s) = s - k; \quad (\text{ODE})$$

where $V_1(\theta; s)$ denotes the partial derivative with respect to the first argument. Recall the standard result of the exponential bandit. For a fixed s , there are two cut-off beliefs: the quitting belief $q(s)$ and the trial belief $Q(s)$, with $q(s) < Q(s)$, which determine the buyer's optimal learning strategy. That is, he continues to learn when his belief falls between the two cut-offs; otherwise, he stops learning. The quitting belief is determined by the usual value matching and smooth pasting conditions:

$$q(s) = \theta : V_1(\theta; s) = 0 \text{ and } V(\theta; s) = 0; \quad (3)$$

The trial belief is the value of belief above which the buyer strictly prefers immediate consumption to acquiring information:

$$Q(s) = \theta : V(\theta; s) = E(v_j | \theta) - (v_h - s); \quad (4)$$

We plug the set of equations (3) into (ODE) and solve the quitting belief $q(s) = \frac{k}{s}$ as the boundary point of the (ODE). In this way, we pin down the buyer's continuation value for learning. With slight abuse of notation, we henceforth use $V(\theta; s)$ to denote the solution of (ODE) with boundary point $q(s)$.¹⁴ From $V(\theta; s)$ so obtained, the upper cutoff trial belief $Q(s)$ is determined through the indifference condition (4).

The construction above involves one implicit assumption: when the buyer stops learning at $q(s)$, he prefers to quit the market than to accept the price. Specifically,

$$E(v_j | q(s)) - (v_h - s) \geq 0; \quad (\text{Learning-Feasibility})$$

We call this the **Learning-Feasibility** constraint. If this constraint fails, no learning can be induced regardless of the prior because learning has no value when it does not

¹⁴Specifically, $V(\theta; s) = s - k - \frac{k(1 - \lambda)}{\lambda} \log \left[\frac{1 - \lambda}{\lambda(1 - \frac{\lambda}{q(s)})} \right]$.

affect the purchase decision. If this constraint fails for all $s \in [0; v_h - v_l]$,¹⁵ the seller then optimally sets a No Return mechanism with price equals to the ex-ante expected value of the product and captures the full allocation surplus $E(v_j - p_0)$. To avoid this trivial result, throughout the paper we assume that there exists two distinct roots $\underline{s} < \bar{s}$ that the **Learning-Feasibility** constraint binds, which is equivalent to the assumption below.

Assumption: $(v_h - v_l) > 4k$.

Proposition 1. If $s \notin [\underline{s}; \bar{s}]$, $V^0(\mu; s) = \max\{0; E(v_j) - (v_h - s)g\}$; and if $s \in [\underline{s}; \bar{s}]$,

$$V^0(\mu; s) = \begin{cases} 0; & \mu < q(s) \\ V(\mu; s); & q(s) < \mu < Q(s) \\ E(v_j) - (v_h - s); & \mu > Q(s) \end{cases}$$




Figure 1: Buyer's value function under No Return mechanisms

Figure 1 depicts the buyer's value function $V^0(\mu; s)$ against μ for different s . For sufficiently high surplus s , buying is attractive even at very low belief, because the value of learning is small. Similarly, for sufficiently low surplus, buying is attractive only if the belief is very high such that immediate consumption is preferable to learning. In both cases, the buyer optimally chooses not to learn, regardless of his prior. The upper and lower gray lines in Figure 1 capture the buyer's value function

¹⁵If $s \notin [0; v_h - v_l]$, then the selling price is either strictly lower than v_l or strictly higher than v_h , and the seller can gain higher profit by changing the price.

in these two scenarios. With moderate $s \in [\underline{s}, \bar{s}]$, the standard results of exponential bandit apply, as illustrated by the red curve in Figure 1. When the buyer's prior belief falls into $[q(s); Q(s)]$, he optimally learns until either good news arrives and he purchases the item, or no news arrives and he quits the market at $q(s)$.

Note that s (the consumer surplus upon arrival of good news) is a sufficient statistics in determining the buyer's value function. When the seller raises s , the buyer's value function shifts up. That is, the seller controls s (indirectly through changing the selling price t_b) to affect the buyer's value from learning. Moreover, the learning interval $[q(s); Q(s)]$ changes with s . It is easy to see that the quitting belief $q(s)$ decreases with s , as the buyer optimally learns for longer if a Poisson jump induces a larger surplus. The trial belief $Q(s)$ also decreases with s . To see this, when s increases by one unit, the selling price must decrease by one unit, and hence the buyer's expected payoff from purchasing increases by one unit. However, if the buyer chooses to learn, the probability of obtaining a Poisson jump is strictly less than 1, and the increment in the continuation value is therefore less than one unit.

Figure 2 plots the quitting belief and trial belief against s . The range $[q(s); Q(s)]$ shifts to lower values as s increases. These two beliefs coincide at \underline{s} and \bar{s} .¹⁶ In the following, we denote $\underline{\mu} = q(\bar{s})$ and $\bar{\mu} = q(\underline{s})$.




Figure 2: Feasible buyer-learning

¹⁶This is implied by the [Learning-Feasibility](#) constraint and the definitions of $q(s)$ and $Q(s)$.

The No Return policy is an important benchmark because it imposes a lower bound on the information rent for the buyer given his option to acquire information.

Lemma 1. *For a fixed selling price $t_b = v_h - s$, all optimal selling mechanisms provide the buyer with an ex-ante trading surplus of $V^0(\theta_0; s)$.*

Suppose that the seller sets a harsh return policy (small x_r , large t_r) such that the buyer would strictly prefer to perform all learning before a transaction and to purchase only when he obtains good news. Such a return policy would be irrelevant to this study because a return would never be requested. Lemma 1 holds trivially. Now suppose that the seller instead offers a benevolent return policy (large x_r , small t_r), intending to reward the buyer for early purchase, i.e., she provides the type- θ_0 buyer with a strictly higher payoff than $V^0(\theta_0; s)$. Then, the seller can attain a larger expected revenue by taking away part of this reward. Essentially, she can increase the buyer's probability of keeping the item upon return and charge a higher return transfer in a proper way that the buyer's optimal stopping is preserved.

An immediate consequence of Proposition 1 and Lemma 1 is the following result.

Corollary 1. *If $\theta_0 \notin [\underline{\theta}; \bar{\theta}]$, the optimal selling mechanism is the No Return mechanism, with $t_b = E(v_j | \theta_0)$ and $(x_r; t_r) = (1; t_b)$. Trading happens with probability 1 and the seller captures the full allocation surplus.*

This corollary describes an extreme case wherein the buyer is sufficiently informed upfront and deems learning to be sub-optimal (see the first part of Proposition 1). Hence, the seller can simply extract the full allocation surplus by setting a non-refundable price equal to the ex-ante expected value of the good. The buyer obtains zero trading surplus.

4 Learning Deterrence and Free Return

If the buyer has a more uncertain prior belief, $\theta_0 \in [\underline{\theta}; \bar{\theta}]$, learning becomes a valuable option to him, which prevents the seller from capturing the full allocation surplus. Consider Figure 2: for a fixed θ_0 , the seller can (1) set $s = Q^{-1}(\theta_0)$ to

prevent the buyer from learning and induce immediate consumption; or (2) set $s \geq [q^{-1}(\theta_0); Q^{-1}(\theta_0))$ to encourage the buyer to learn.¹⁷ The shaded region corresponds to the feasible set of buyer-learning outcomes (stopping beliefs) that can be induced for a type- θ_0 buyer.

Learning Deterrence. To induce immediate consumption requires $s \leq Q^{-1}(\theta_0)$, so that the type- θ_0 buyer weakly prefers to purchase the item rather than to acquire information. The seller lets the inequality bind to maximize her revenue.¹⁸ Denote $t^D(\theta_0) := v_h - Q^{-1}(\theta_0)$. We call the mechanism $f_{t^D}; (1; t^D)g$ *Learning Deterrence*. To deter learning, the seller has to lower the price so as to give away part of the allocation surplus to compensate the buyer until the value of information becomes non-positive.¹⁹ With Learning Deterrence, the buyer is indifferent between immediate consumption and learning, i.e., his expected trading surplus equals his continuation value from learning, $E(v|j = \theta_0) - t^D(\theta_0) = V(\theta_0; Q^{-1}(\theta_0))$. Furthermore, the allocation is efficient, which means that the joint surplus equals the full allocation surplus, $E(v|j = \theta_0)$, and the seller obtains revenue $r^D(\theta_0) = t^D(\theta_0)$. In Figure 2, the orange dot, i.e., the intersection of the shaded region and the orange curve $Q(s)$, corresponds to Learning Deterrence.

Free Return. To induce learning requires $s \geq [q^{-1}(\theta_0); Q^{-1}(\theta_0))$. Free Return $f_{t_b}; (0; 0)g$ is a simple mechanism available to the seller. Consider a typical Free Return mechanism with a price of $v_h - s$. The type- θ_0 buyer then learns until his posterior belief falls below $q(s)$. By varying s , the seller can induce different quitting beliefs and thereby induce different amounts of buyer learning. In Figure 2, the intersection of the shaded region and the blue curve $q(s)$ represents the feasible set of quitting beliefs that can be induced by Free Return. A common property of all Free Return mechanisms is that the buyer stops learning when the continuation value from learning decreases to 0, which we define as *full learning*. To find the revenue-maximizing Free Return mechanism, consider the constrained optimization

¹⁷If the seller sets $s < q^{-1}(\theta_0)$, the type- θ_0 buyer walks away.

¹⁸The buyer breaks indifference by purchasing the item immediately.

¹⁹The value of information refers to the difference between the value function and the payoff from optimally choosing between purchasing and quitting.

problem (F) below:

$$\begin{aligned} F(\theta) &:= \max_s \frac{\theta}{1 - q(s)} (v_h - s) \\ \text{s.t. } & q^{-1}(\theta) \leq s \leq Q^{-1}(\theta) \end{aligned} \quad (F)$$

Note that $\frac{\theta}{1 - q(s)}$ is the ex-ante probability that good news arrives before the buyer's belief falls below $q(s)$. Optimization over Free Return mechanisms is mechanical. On the one hand, the seller can induce longer learning with a larger S , which raises the expected probability of realizing good news and thereby increases the probability of a successful sale. On the other hand, a larger S also reduces the selling price. The unconstrained optimization admits a closed-form solution. We denote the unconstrained maximizer as $S^F(\theta)$ and the corresponding revenue as $F^F(\theta)$.

If $F^F(\theta) \geq D(\theta)$, then $F^F(\theta)$ adopts the same expression as $F(\theta)$. This is because the free return price $v_h - S^F(\theta)$ is strictly larger than the Learning Deterrence price $v_h - Q^{-1}(\theta)$ as free return incurs a strictly positive probability of no sale, thus $S^F(\theta) > Q^{-1}(\theta)$, and we can easily verify $q^{-1}(\theta) < S^F(\theta)$ given the exact expression of $S^F(\theta)$.

Proposition 2. Denote $F(\theta)$ as the expected revenue from an optimal selling mechanism. If $\theta \in [\underline{\theta}; \bar{\theta}]$,

$$F(\theta) = \max\{F^D(\theta); F^F(\theta)\}g$$

This proposition is trivial as both Learning Deterrence and Free Return are feasible policies. In Figure 2, within the boundary of the shaded region, there is a large set of interior buyer-learning outcomes that the seller can induce through another type of mechanism: Stochastic Return. However, it is useful to mention the above inequality early on, as we will later show that it holds as equality.

We close this section by discussing the welfare properties of Learning Deterrence and the optimal Free Return policy when the prior varies.

Proposition 3. With Learning Deterrence $F^D(\theta); (1; t^D(\theta))g$ where $\theta \in [\underline{\theta}; \bar{\theta}]$, (1) the buyer obtains a (weakly) positive trading surplus $V(\theta; Q^{-1}(\theta))$ that is single-peaked in θ and equals 0 at $\underline{\theta}$ and $\bar{\theta}$;

- (2) the seller's revenue $v^p(\theta_0)$ increases with θ_0 and equals $E(v| \theta_0)$ at $\underline{\theta}$ and $\bar{\theta}$;
- (3) the joint surplus is $E(v| \theta_0)$.

To deter private buyer learning, the seller has to sufficiently lower the price so that accepting the price is more attractive for the buyer than acquiring information. $t^p(\theta_0)$ is the highest achievable price to prevent type- θ_0 buyer from learning. Learning Deterrence is different from the extreme case stated in Corollary 1, as the buyer must be induced to give up his opportunity information rent $V(\theta_0; Q^{-1}(\theta_0))$, which he could have enjoyed if he continued to learn. When the buyer is initially more uncertain, he enjoys larger benefits from learning and thereby attains greater bargaining power if the seller wants to deter learning. This hints at the non-monotonicity of the buyer's expected trading surplus. Moreover, from Figure 2, $Q^{-1}(\theta_0)$ is also the highest s that is able to induce a type- θ_0 buyer to learn, which indicates that $V(\theta_0; Q^{-1}(\theta_0))$ is the highest information rent that the type- θ_0 buyer can possibly earn. Therefore, the buyer attains the highest trading surplus under Learning Deterrence.

Proposition 4. Under Free Return, both the selling price $v_h = s^f(\theta_0)$ and seller revenue $F(\theta_0)$ are increasing in θ_0 .

Under Free Return, a higher prior belief increases the probability of a successful sale, and the seller optimally puts more weight on getting a higher selling price and less weight on raising the probability of making a sale. Therefore, the seller locally adjusts s to a lower value to increase the selling price. With a more optimistic prior belief, the seller can obtain a larger expected revenue even if she retains the same price. Therefore, $F(\theta_0)$ increases with θ_0 .

5 Partial Learning

To encourage learning, Free Return is just one type of mechanism available to the seller. Instead of inducing the buyer to keep learning until his belief reaches the quitting belief (which we call full learning), she can use Stochastic Return mechanisms to induce the buyer to stop at any intermediate belief in the interior of the

shaded region in Figure 2, which we define as *partial learning*. Under partial learning, there is a strictly *positive* continuation value when the buyer decides to stop learning before good news arrives.

Lemma 2. For fixed $s \in [\underline{s}; \bar{s}]$, the return policy $(x_r(\cdot; s); t_r(\cdot; s))$ induces the buyer to stop learning and request a return at $\mu \in [q(s); Q(s)]$, where

$$x_r(\mu; s) = \frac{V_1(\mu; s)}{v_h - v_l}; \quad (5)$$

$$t_r(\mu; s) = E(v|\mu) \frac{V_1(\mu; s)}{v_h - v_l} - V(\mu; s); \quad (6)$$

Furthermore, the return transfer $t_r(\mu; s)$ increases with both μ and s and with cross derivative equal to 0; and $x_r(\mu; s)$ increases with μ .




Figure 3: Partial Learning with Stochastic Return

Equations (5) and (6) are the familiar smooth pasting and value matching conditions for the buyer to optimally stop at μ given the return policy. In Figure 3, a type- μ buyer obtains expected utility $E(v|\mu)x_r - t_r$ if he requests a return, whereas he attains continuation value $V(\mu; s)$ if he keeps learning. Making the buyer's return payoff tangent to his continuation value at μ , the buyer is willing to stop learning at μ and request such a return. One can view the mechanism $f: v_h - s; (x_r(\cdot; s); t_r(\cdot; s))g$ as a direct mechanism to screen an interim-type buyer, such that (1) the buyer stops acquiring information if his posterior reaches μ or 1; and (2) the buyer strictly prefers to truthfully report his posterior beliefs.

Interestingly, Lemma 2 shows that $t_r(\cdot; s)$ increases with s , meaning that the seller actually obtains a larger return transfer if she enforces earlier stopping, even though she has to compensate the buyer with more opportunity information rent. The main reason is that the seller uses both x_r and t_r towards the buyer's compensation. In particular, when the seller intends to induce a higher stopping belief, she allows the buyer to keep the item with greater probability upon return, thereby inducing the buyer to make a larger return transfer. With similar reasoning, the return transfer $t_r(\cdot; s)$ also increases with s . This implies that for a fixed stopping belief, the seller obtains a larger return transfer by charging a smaller selling price.

For fixed s , the domain of $t_r(\cdot; s)$ is $[q(s); Q(s)]$. Note that

$$\begin{aligned} \lim_{\downarrow q(s)} x_r(\cdot; s) &= 0 & \text{and} & & \lim_{\downarrow q(s)} t_r(\cdot; s) &= 0; \\ \lim_{\downarrow Q(s)} x_r(\cdot; s) &< 1 & \text{and} & & \lim_{\downarrow Q(s)} t_r(\cdot; s) &< v_h - s. \end{aligned}$$

That is, for fixed s , Free Return is the left limit of Stochastic Return. In contrast, the right limit of Stochastic Return is strictly dominated by Learning Deterrence in terms of seller revenue. This follows because the return transfer $t_r(Q(s); s)$ must be smaller than the selling price, $v_h - s$, which implies $t_r(\cdot; Q^{-1}(\cdot)) < v_h - Q^{-1}(\cdot) = t^D(\cdot)$. Regarding how much information the buyer acquires, Free Return and Learning Deterrence can be interpreted as opposite limits of Stochastic Return.

Now we formulate the seller's optimization problem for encouraging learning.

$$\begin{aligned} \max_{s \in [q^{-1}(\theta_0); Q^{-1}(\theta_0)]} \left\{ \max_{\substack{\cdot \\ \text{s.t. } q(s) \leq \cdot \leq Q(s)}} (\cdot; s) = \frac{\theta_0}{1} (v_h - s) + \frac{1 - \theta_0}{1} t_r(\cdot; s) \right\} \quad (P) \end{aligned}$$

0

The seller's expected revenue is a weighted average between the selling price and the return transfer. The weights depend on the buyer's prior and stopping beliefs. An $s \in [q^{-1}(\theta_0); Q^{-1}(\theta_0)]$ encourages the buyer to learn and determines the level of continuation value that the buyer can attain. Given s , the seller optimizes her expected revenue over the set of inducible return beliefs $\cdot \in [q(s); Q(s)]$. As the buyer's posterior belief decreases if no good news arrives, the return belief \cdot must be smaller than θ_0 .

Theorem 1. If $\theta_0 \geq [\underline{\theta}; \bar{\theta}]$, Stochastic return is dominated by either Learning Deterrence or the optimal Free Return. That is,

$$V(\theta_0) = \max_{\mathcal{F}^D(\theta_0); \mathcal{F}(\theta_0)} g.$$

We obtain this theorem by showing that the solution of (P) is weakly smaller than $\max_{\mathcal{F}^D(\theta_0); \mathcal{F}(\theta_0)} g$. That is, the seller cannot gain from inducing partial learning. The proof proceeds in two steps. In the first step, we study the internal maximization of (P) to derive the optimal stopping belief $\theta(s)$ for fixed surplus s . It turns out that $\theta(s)$ is monotonic in s . We then find the domain of s such that $\theta(s)$ is an interior solution, i.e., $\theta(s) \in (q(s); Q(s))$. This implies that partial learning induced by Stochastic Return is optimal (Lemma 3) as long as s is exogenously determined within this domain. In the second step, we optimize over s . We show that the seller revenue is *quasi-concave* in s , taking into account changes in $\theta(s)$ as s varies. This implies that for a given Stochastic Return mechanism $\mathcal{F}V_h = s; (x_r(\theta(s); s); t_r(\theta(s); s))g$, the seller is strictly better off by *either* raising s to deter the buyer from private learning *or* lowering s to encourage full learning, which establishes the sub-optimality of partial learning. Thus, the revenue-maximizing mechanism is either Learning Deterrence or Free Return. In the next two subsections, we discuss these two steps in detail.

5.1 Step 1

For fixed s , and ignoring the constraint $\theta \geq \theta_0$, consider the internal maximization problem of (P):

$$\max_{\theta \in [q(s); Q(s)]} V(\theta; s): \quad (R)$$

We verify that $V(\theta; s)$ is quasi-concave on $[q(s); Q(s)]$. Therefore, it attains a maximum either at the boundaries (which reduces to no learning or full learning), or at an interior solution characterized by the first-order condition (which corresponds to partial learning). Denote $\theta(s)$ as the maximizer of (R) when it is the solution to the first-order condition. Formally,

$$\theta(s) = \mathcal{F} \in [q(s); Q(s)] : \mathcal{F}_1(\theta; s) = 0g.$$




Figure 4: Return transfers and interior solutions

Recall that $t_r(\cdot; s)$ is a weighted average between the return transfer and the selling price. Rearranging the first order condition, we obtain²⁰

$$\underbrace{\Pr(\text{return}) \frac{\partial t_r(\cdot; s)}{\partial \mu}}_{\text{larger return transfer}} = \underbrace{[v_h - s - t_r(\cdot; s)] \frac{d \Pr(\text{return})}{d \mu}}_{\text{more frequent return}}; \quad (7)$$

Recall that the seller can gain a larger return transfer $t_r(\cdot; s)$ if the buyer stops learning and returns the product at a higher belief μ . However, raising the stopping belief μ increases the probability of receiving a refund request, reducing her revenue from $t_b = v_h - s$ to $t_r(\cdot; s)$.

In Figure 4, for different sizes of s , the gray dashed curves depict $t_r(\cdot; s)$ against within the domain $[q(s); Q(s)]$. With higher s , the curve of $t_r(\cdot; s)$ moves up-left. For a fixed s , to locate $\mu^*(s)$ graphically, we apply the concavification technique. The blue dashed curve depicts the concavification. $\mu^*(s)$ is the value of belief at which the affine part of the blue dashed curve is tangent to the return transfer $t_r(\cdot; s)$. Typically, $\mu^*(s)$ is irrelevant to the prior belief μ_0 .

Note that $\mu^*(s)$ is strictly increasing in s . To see this, consider the above trade-off in equation (7) again. Recall that Lemma 2 establishes that the return transfer

²⁰Denote $\Pr(\text{return}) = \frac{1}{1-\alpha}$ as the ex-ante probability of return.

$t_r(\cdot; s)$ increases with s . Therefore, the refund $v_h - s - t_r(\cdot; s)$ becomes smaller when s is higher. The seller then cares less about the return rate, and her incentive to gain a larger return transfer is relatively stronger. Thus, she optimally adapts to gain a larger return transfer by inducing a larger stopping belief, meaning that $\hat{s}(s)$ increases with s . Furthermore, if s becomes sufficiently high, the refund becomes sufficiently small that gaining a larger return transfer becomes the seller's dominant incentive. She then prefers to induce the maximal stopping belief, rendering the upper boundary $Q(s)$ the optimal return belief. Conversely, if s is sufficiently small, the dominant incentive is to reduce the return rate, and the seller induces the minimal stopping belief, rendering the lower boundary $q(s)$ the optimal stopping belief. Hence, in the last two scenarios, the seller optimally induces no learning and full learning, respectively.

Lemma 3. Let \underline{s} be the solution of $\hat{s}_1(\cdot; q^{-1}(\cdot)) = 0$ and $\underline{s} < 0.5$.

- (1) If $s < q^{-1}(\underline{s})$, full learning with return belief $q(s)$ is optimal;
- (2) If $s \in (q^{-1}(\underline{s}); Q^{-1}(0.5))$, partial learning with return belief $\hat{s}(s)$ is optimal;
- (3) If $s \geq Q^{-1}(0.5)$, no learning with return belief $Q(s)$ is optimal.

Lemma 3 summarizes the optimal stopping belief as s varies. The second term of this lemma indicates that partial learning can be seller-optimal if the value of s is intermediate (specifically, for $s \in (q^{-1}(\underline{s}); Q^{-1}(0.5))$). Given $\hat{s}(s)$ being strictly increasing in s , this means that partial learning at stopping belief \hat{s} can be induced with optimality only when $s \in (\underline{s}; 0.5)$. The value \underline{s} comes from our construction that the first-order equation, $\hat{s}_1(\cdot; q^{-1}(\cdot)) = 0$, has a unique solution at \underline{s} , and the value 0.5 comes from the observation that the first-order equation, $\hat{s}_1(\cdot; Q^{-1}(\cdot)) = 0$, has a unique solution at $s = 0.5$.

5.2 Step 2

As the optimal stopping belief $\hat{s}(\cdot)$ derived from the earlier step is strictly increasing, we can let $s(\cdot)$ for $s \in [\underline{s}; 0.5]$ represent the inverse of $\hat{s}(\cdot)$. We also define $t_r(\cdot) := t_r(\cdot; s(\cdot))$ with domain $[\underline{s}; 0.5]$. The blue solid curve in Figure 4 depicts $t_r(\cdot)$, whose left boundary point refers to the optimality of inducing full learning as

fixing $s = q^{-1}(\underline{\cdot})$, and the right boundary point refers to the optimality of inducing no learning as fixing $s = Q^{-1}(0.5)$.

If the solution to (P) turns out to be an interior Stochastic Return mechanism, it has to be located on the interior path of $t_r(\cdot)$. Once we know the precise value of t_r and the return belief, we can determine the exact selling mechanism. In this part, we show that, compared with any Stochastic Return mechanism that induces partial learning, the seller is better off by either (1) reducing s and enforcing full learning; or (2) raising s and enforcing no learning. In particular, $(\cdot; s(\cdot))$ is quasi-convex on $\mathcal{L}[\underline{\cdot}; 0.5]$, where

$$(\cdot; s(\cdot)) = t_r(\cdot) + [1 - \Pr(\text{return})][v_h - s(\cdot) - t_r(\cdot)] \quad (8)$$

The first term is the seller's guaranteed revenue, and the second term refers to the extra revenue she can obtain if the buyer discovers good news. If the seller intends to raise the return belief, she benefits from a larger safe revenue $t_r(\cdot)$, because $t_r(\cdot; s)$ increases with both arguments and $s(\cdot)$ increases with \cdot . However, the seller will suffer from a smaller probability that good news will be discovered (i.e., a smaller $1 - \Pr(\text{return})$) and a higher loss in extra revenue in that event (i.e., a smaller $v_h - s(\cdot) - t_r(\cdot)$).

Substituting the first-order condition (7) into the seller's expected revenue (8), we can simplify the latter expression to write:

$$(\cdot; s(\cdot)) = t_r(\cdot) + \frac{\partial t_r(\cdot; s(\cdot))}{\partial \cdot}(\cdot_0) \quad (9)$$

The prior belief \cdot_0 only matters for the second term (i.e., the extra revenue gain). Suppose that the seller induces a return belief \cdot relatively small compared with the prior \cdot_0 (i.e., large $\cdot_0 - \cdot$). The extra revenue gain will then be substantial. The seller then has a stronger incentive to lower the buyer's stopping belief so as to raise the probability of obtaining the extra revenue gain. However, as the seller optimally adjusts her policy and induces an even lower stopping belief \cdot , this further increases the second term through $\cdot_0 - \cdot$, thus reinforcing the motive to lower the buyer's stopping belief. This gives rise to a corner solution at the lower boundary $\underline{\cdot}$. If the seller induces a return belief \cdot relatively close to the prior \cdot_0 (i.e., small

θ_0), opposite forces are at work. The seller's stronger incentive would be to raise her safe revenue $t_r(\theta)$. This can be achieved by raising θ because $t_r(\theta)$ increases with θ . However, this lowers the extra revenue gain through its effect on θ_0 , thus further reinforcing the seller's incentive to raise her safe revenue rather than to increase the probability of obtaining the extra revenue gain. This produces a corner solution at the upper boundary of some feasible region, which is θ_0 if $\theta_0 < 0.5$,²¹ and is $(Q^{-1}(\theta_0))$ if $\theta_0 > 0.5$.²² In other words, the seller's expected revenue (8) is quasi-convex in the stopping belief.

Thus, the optimal Stochastic Return mechanism could induce any one of these three stopping beliefs:²³ inducing a stopping belief at $\underline{\theta}$ to induce full learning, inducing a stopping belief at θ_0 to induce no learning, or inducing a stopping belief at $(Q^{-1}(\theta_0))$ to induce partial learning. We can verify the third one to be strictly dominated by Learning Deterrence.²⁴ Furthermore, as we mentioned previously, inducing full learning corresponds to a Free Return mechanism, whereas inducing no learning via Stochastic Return is strictly dominated by Learning Deterrence. Therefore, inducing partial learning via Stochastic Return is sub-optimal.

6 Optimal Selling Mechanism

Given Theorem 1, $\max_{\theta} f^D(\theta); f^F(\theta)g$ determines the value of the optimal selling mechanism. Denote $\theta = k = \frac{c}{v_h}$. High values of θ (arising from a high learning cost or small discovery rate) indicate relatively inefficient learning. Let F be the set of θ_0

²¹This is implied by $\theta \in [\underline{\theta}; 0.5]$ and $\theta_0 > 0$.

²²Recall that $s \geq [q^{-1}(\theta_0); Q^{-1}(\theta_0)]$. If $\theta_0 > 0.5$, then $(Q^{-1}(\theta_0)) > 0.5$.

²³The stopping belief induced by an optimal Stochastic Return mechanism is the solution of (P) while restricting $s = s(\theta)$.

²⁴If we temporarily ignore the constraint that $s \geq [q^{-1}(\theta_0); Q^{-1}(\theta_0)]$, then, given the quasi-convexity, inducing stopping at $(Q^{-1}(\theta_0))$ is dominated by inducing stopping at 0.5, which generates an expected revenue as a weighted average between the return transfer $t_r(0.5; Q^{-1}(0.5))$ and the price $v_h Q^{-1}(0.5) = t^D(0.5)$, which is smaller than $t^D(0.5)$ and hence smaller than $t^D(\theta_0)$, as $\theta_0 > 0.5$.

such that the seller weakly prefers to choose Free Return. Formally,

$$F = f_0 \geq [\underline{\mu}; \bar{\mu}] : \pi^F(\mu_0) \geq \pi^D(\mu_0)g:$$

Theorem 2. *There exists a μ_0 such that if $\mu_0 < \mu_0^*$, then F is a closed interval and $F = (v_l; v_h)$; if $\mu_0 > \mu_0^*$, $F = \emptyset$. The optimal mechanism takes following form:*

1. No Return (with $t_b = E(v| \mu_0)$ and $(x_r; t_r) = (1; t_b)$) if $\mu_0 \geq [\underline{\mu}; \bar{\mu}]$;
2. Learning Deterrence (with $t_b = t^D(\mu_0)$ and $(x_r; t_r) = (1; t_b)$) if $\mu_0 \geq [\underline{\mu}; \bar{\mu}]$ and $\mu_0 \geq F$;
3. Free Return (with $t_b = v_h - s^F(\mu_0)$ and $(x_r; t_r) = (0; 0)$) if $\mu_0 \geq F$.




Figure 5: Expected seller revenue from Free Return and Learning Deterrence when μ_0 is small

When learning is valuable to the buyer, i.e., $\mu_0 \geq [\underline{\mu}; \bar{\mu}]$, the seller chooses the amount of buyer learning she would like to induce through her selling mechanism. Figure 5 depicts the expected revenue of Learning Deterrence and (optimal) Free Return when $\mu_0 < \mu_0^*$. Which mechanism generates more revenue depends on the buyer's prior belief, which is a measure of both the degree to which he is informed and the level of his optimism. To interpret Theorem 2, recall that $\pi^D(\mu_0) = E(v| \mu_0) - V(\mu_0; Q^{-1}(\mu_0))$, because the seller has to compensate the buyer with his opportunity information rent $V(\mu_0; Q^{-1}(\mu_0))$. When μ_0 is sufficiently informative (equal to $\underline{\mu}$ or $\bar{\mu}$), it is obvious that the seller optimally deters learning, given that the information rent

the buyer can command is zero. The more uncertain the buyer's prior belief, the larger his information rent. Therefore, as μ_0 moves from either $\underline{\mu}$ or $\bar{\mu}$ toward a more intermediate belief, the seller has to give away a larger amount of surplus to the buyer, and Learning Deterrence becomes less attractive. By switching to Free Return, the seller encourages the buyer to learn and avoids paying for the opportunity information rent. Free Return, however, wastes a strictly positive amount of allocation surplus due to the probability of receiving a return request. Nevertheless, if the prior belief μ_0 is relatively optimistic, even Free Return can guarantee the seller a relatively high probability of a successful sale and the seller can also charge a higher price due to the buyer's higher willingness to pay. Therefore, the seller will optimally switch to Free Return if the buyer is relatively optimistic but still not very certain.

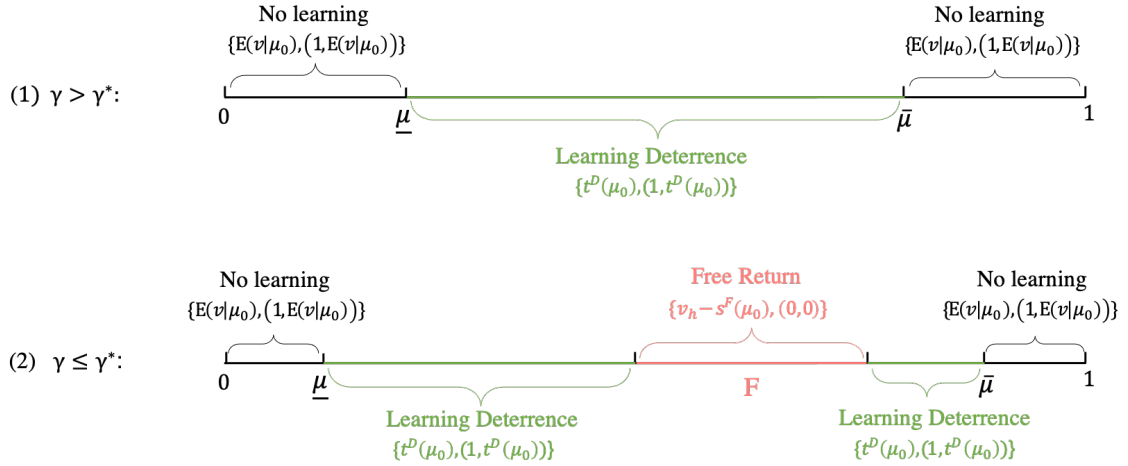


Figure 6: Optimal selling mechanism

When learning is inefficient (γ is sufficiently large), the buyer obtains little value from learning. Therefore, the amount of information rent that the seller has to pay the buyer to deter learning is small regardless of his prior belief. Hence, Learning Deterrence becomes more appealing to the seller. Meanwhile, Free Return becomes less profitable as the buyer optimally quits learning earlier, which reduces the ex-ante probability of a successful sale. Therefore, when γ increases, the set of priors F that supports Free Return as the optimal mechanism shrinks; and when $\gamma > \gamma^*$, F becomes an empty set. See Figure 6.

Proposition 5. $F(\theta_0)$ is decreasing in λ and $D(\theta_0)$ is increasing in λ . If $\theta_1 < \theta_2$, then $F(\theta_2) > F(\theta_1)$ and $[L(\theta_2); \bar{L}(\theta_2)] \supset [L(\theta_1); \bar{L}(\theta_1)]$.

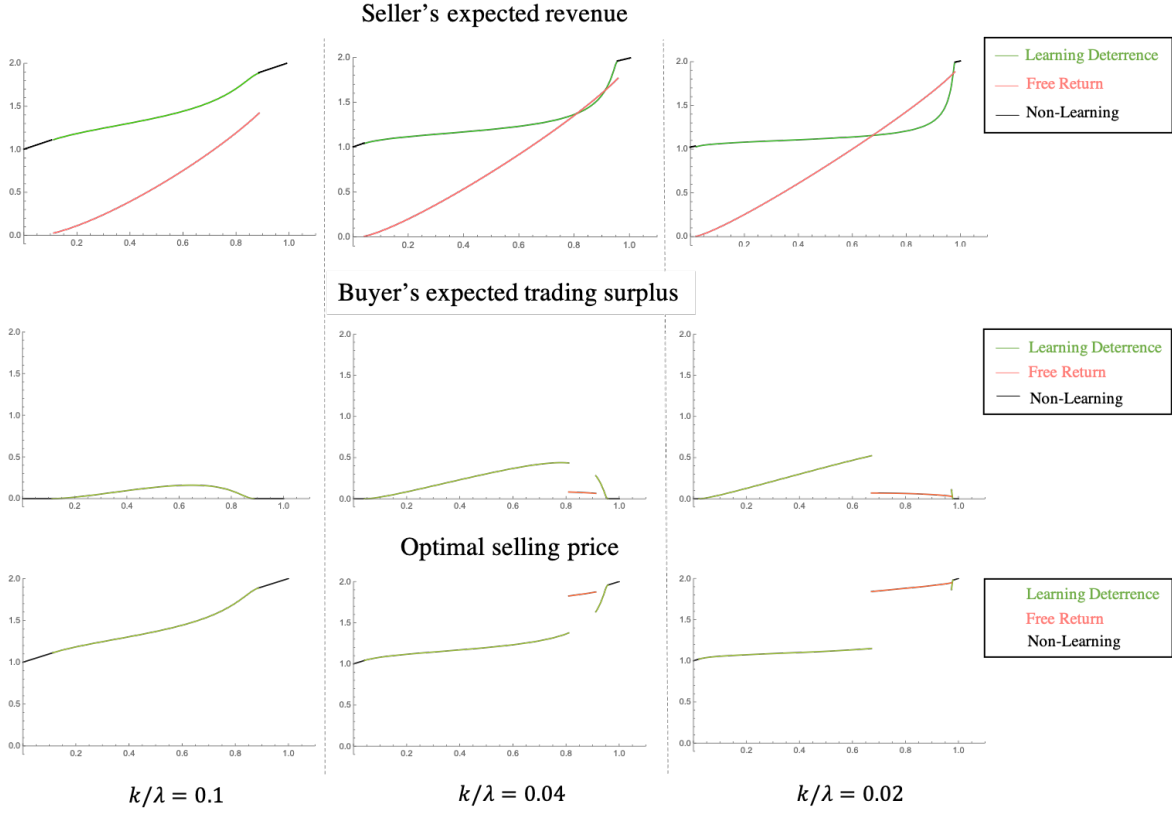


Figure 7: Comparative statics

When learning is prevented, under either no learning (the black intervals in Figure 6) or Learning Deterrence, the joint surplus attains the full allocation surplus $E(v| \theta_0)$. The allocation of this joint surplus between the two parties depends on how much the buyer values learning. When $\theta_1 > \theta_2$, the buyer's expected trading surplus simply equals to the information rent that he could obtain through learning, and it is single-peaked in his prior belief (see Corollary 1 and Proposition 3). The leftmost plot in the middle row in Figure 7 depicts the buyer's expected trading surplus under this case. When $\theta_1 < \theta_2$, the buyer is encouraged to perform full learning when $\theta_0 \geq F$. This causes a strict decline in the buyer's ex-ante surplus. In other words, the buyer with prior belief $\theta_0 \geq F$ is made worse off: as the seller chooses Free Return rather than Learning Deterrence, the buyer's information rent falls from

$V(\theta_0; Q^1(\theta_0))$ to $V(\theta_0; S^F(\theta_0))$. The middle and rightmost plots in the middle row in Figure 7 describe this decline. This seemingly counter-intuitive result is mainly due to the higher selling price under Free Return; see the optimal selling prices depicted in the bottom row in Figure 7. Moreover, the buyer's ex-ante surplus equals the joint allocation surplus minus the seller's profit. If the seller optimally chooses Free Return when $\theta_0 \geq F$, her profit from Free Return is larger than that from Learning Deterrence, while the joint allocation surplus falls because of allocation inefficiency. This explains why the buyer's ex-ante surplus abruptly falls when the seller switches to Free Return.

Proposition 6. (A limit result)

$$\lim_{\theta_0 \rightarrow 0} \max_{\theta_0} f^D(\theta_0); F(\theta_0)g = \begin{cases} v_l; & \theta_0 < \frac{v_l}{v_h} \\ \theta_0 v_h; & \theta_0 \geq \frac{v_l}{v_h} \end{cases}$$

Recall the well-known result for selling a single indivisible good, where the buyer has private valuation $v \geq \bar{v}_l; v_h g$ and the seller attaches a prior belief θ_0 as the probability of the buyer having a high valuation. The revenue-maximizing mechanism suggests that, when $\theta_0 < \frac{v_l}{v_h}$, the seller should set a price of v_l to sell to both high-valuation and low-valuation buyers (the mass-market strategy); when $\theta_0 \geq \frac{v_l}{v_h}$, the seller should set a price of v_h to sell to high-valuation buyers only (the niche-market strategy). Intuitively, when the seller believes that she has a greater probability of matching with a high-valuation buyer, she becomes more willing to charge a higher price and sacrifice transaction opportunities with low-valuation buyers. In our model, as θ_0 converges to zero, the buyer can learn almost perfect information; as a result, the seller's revenue converges to the revenue she can obtain with a perfectly privately informed buyer. Therefore, with Free Return, the seller optimally sets the selling price arbitrarily close to v_h and lets go of buyers who are almost sure to have a low valuation. With Learning Deterrence, the seller has to set a price arbitrarily close to v_l , because otherwise, the buyer always has an incentive to learn to avoid consuming the item when his true valuation is low. The ratio $\frac{v_l}{v_h}$ determines the prior belief cut-off at which the seller is indifferent between Free Return and Learning Deterrence.

7 Discussion

7.1 Optimality of Stochastic Return

Given Lemma 3, Stochastic Return might turn out to be an optimal mechanism in situations in which the seller cannot freely adjust the price, for example, if the price is exogenously determined by other parties. This is the case for some online retail platforms, which can control their return policies but must set prices determined by their suppliers. Other scenarios in which Stochastic Return might be optimal can arise if the prices are driven down by price competition when similar products are sold by multiple sellers, or if a price cap is imposed by the regulator.

Corollary 2. Suppose that the price is constrained by a price cap $t^c = v_h - s^c$ with $s^c \geq (q^{-1}(\underline{v}); Q^{-1}(0.5))$. If $\theta_0 \geq (\underline{v}(s^c); Q(s^c))$, the optimal selling mechanism takes one of the two forms below:

1. *Learning Deterrence with $t_b = t^p(\theta_0)$ and $(x_r; t_r) = (1; t_b)$;*
2. *Stochastic Return with $t_b = t^c$ and $(x_r; t_r) = (x_r(\underline{v}(s^c); s^c); t_r(\underline{v}(s^c); s^c))$.*

Suppose that the price $t_b = v_h - s$ is exogenous with $s \geq (q^{-1}(\underline{v}); Q^{-1}(0.5))$. Lemma 3 implies that for prior belief $\theta_0 \geq (\underline{v}(s); Q(s))$, the optimal selling mechanism induces a stopping belief $\underline{v}(s)$, and takes the form of Stochastic Return with a price t_b and a return policy $(x_r(\underline{v}(s); s); t_r(\underline{v}(s); s))$. With a price cap t^c , the seller can adjust the price within the range of values smaller than t^c . As the seller's revenue is quasi-convex in the price when taking into account changes in the optimal stopping belief as the price varies, the optimal mechanism either reduces the price to $t^p(\theta_0)$ to prevent the buyer from private learning, or raises the price to t^c and induces partial learning with the stopping belief $\underline{v}(s^c)$.

7.2 Positive seller valuations

In this section, we discuss the situations in which the seller has a positive valuation of the product. This could be, for example, because the seller stands to gain

positive revenue from a resale. We use u to denote the seller's product valuation or reservation value. Figure 8 depicts her revenue from Learning Deterrence (green curves) and the optimal Free Return (red curves) as u varies from 0 to v_h . The blue dashed lines represent the values of u . Note that the seller's revenue from Learning Deterrence remains constant when u varies, as it equals the full allocation surplus minus the information rent with which she compensates the buyer to deter learning. Both parts solely depend on the buyer's valuation. However, the seller's revenue from Free Return increases with u as she can collect her reservation value if the buyer returns the product.²⁵ Note that Free Return and Learning Deterrence are only relevant when the buyer is less well-informed ex ante. In other words, if the prior belief $\theta_0 \notin [\underline{\theta}, \bar{\theta}]$, the buyer deems learning suboptimal and the seller can set a price equal to the expected buyer valuation and capture the full trading surplus (the black curves).




Figure 8: Simulation for different seller valuations ($v_l = 1; v_h = 2; \sigma = 0.04$)

²⁵With a positive u , the objective function for Free Return becomes to: $\max_s \frac{1}{1} \frac{q(s)}{q(s)} (v_h - u - s) + u$.

First, consider the cases where $u = v_l$. In this case, efficiency requires trading with probability one, and the seller can always set a price $t_b = v_l$ to prevent the buyer from learning and obtain a revenue higher than her reservation value. That is, the blue line is certainly to lie below the green curve and the black curves. Therefore, if $\theta_0 \geq [\underline{\theta}; \bar{\theta}]$, the revenue from the optimal mechanism is determined by the maximum revenue that lies between Learning Deterrence and Free Return. Meanwhile, as larger u induces a greater Free Return revenue, the set of prior beliefs at which the seller optimally chooses Free Return expands. Otherwise, for extreme prior beliefs, the revenue is depicted by the black curves.²⁶

Next, we discuss the cases where $u > v_l$. If $\theta_0 \geq [\underline{\theta}; \bar{\theta}]$ and the prior belief is low (e.g., $E(v_j | \theta) < u$), then deterring learning to induce immediate trading becomes a dominated strategy, as keeping the item provides the seller with a higher payoff (the blue line is higher than the green curve in some regions of each plot). Thus, the seller has an even stronger incentive to encourage buyer learning, as a return would create efficiency. As a consequence, the set of prior beliefs at which Free Return dominates Learning Deterrence expands as u increases. Meanwhile, with a larger u , the cutoff belief $q(v_h - u)$ above which the seller strictly prefers Free Return rather than keeping the item also increases.²⁷ Above all, the set of prior beliefs at which the seller optimally chooses Learning Deterrence shrinks. However, the set of prior beliefs that supports Free Return shifts to the right and may eventually vanish as keeping the item becomes more attractive. For example, when $u = v_h$, the seller never sells.

7.3 Cancellation fee

In this section, we discuss the scenarios where the learning process is more efficient after purchase. In many scenarios, such as the market of database license, the buyer

²⁶Alternatively, we can normalize the buyer's valuation by subtracting the seller's valuation. That is, if $u = v_l$, the optimal mechanism is characterized by Theorem 2 after replacing the buyer valuation with $v_h - u > v_l - u = 0$ and setting the seller's valuation to 0.

²⁷If $\theta_0 = q(v_h - u)$, the optimal pricing for free return equals u .

has access to quite limited information before purchase. To simplify the illustration, we use ρ_B and ρ_P to denote the learning rate before and after the transaction, respectively, and let the learning cost be constant. First, consider an extreme case wherein $\rho_P > \rho_B = 0$. If the seller's reservation value $u < v_l$, she can easily capture the full allocation surplus by designing a non-refundable payment equal to the buyer's expected valuation. Otherwise, if $v_h > u > v_l$, then the seller can request an upfront transfer equal to the buyer's option value of buying the product at a strike price u (under the learning rate ρ_P). Such a mechanism induces socially optimal information acquisition yet ensures the seller extract the entire surplus.

In a more general case where $\rho_P > \rho_B > 0$, the seller optimally requests an upfront transfer. Denote t_u as the upfront transfer. One can view this upfront transfer as a non-refundable part of the selling price t_b . Alternatively, t_u can be interpreted as the *cancellation fee* that the buyer will incur once he requests a return, and it is a complementary instrument to the original return policy $(x_r; t_r)$. To be consistent with the main text, we adopt the second interpretation and let t_u be the cancellation fee. More specifically, under a mechanism with selling price t_b , return policy $(x_r; t_r)$ and cancellation fee t_u , the net refund that the buyer can obtain is $t_b - t_r - t_u$. This can also capture the scenarios in which the buyer returns the product but can only receive a partial refund.

Note that if $\rho_P > \rho_B > 0$, the transaction itself generates additional information rent for the buyer. If the seller intends to encourage learning, she can nonetheless collect this additional information rent by requesting a cancellation fee (or an upfront transfer) of the same size, without further concern on the buyer's participation constraint. With slight abuse of notation, denote $V(s; t_b)$ as the buyer's continuation value when $s = v_h - t_b$.²⁸ Then, for a fixed selling price, t_b , the buyer's continuation value from post-transaction learning is $V(s_0; t_b; \rho_P) > V(s_0; t_b; \rho_B)$. The proposition below suggests that for any fixed selling price, the seller is strictly better off by setting a cancellation fee as a complementary instrument of the return policy, as long as there is a positive amount of buyer learning on path.

²⁸It is the solution to the differential equation with $(q(s); 0)$ as the boundary point.

Proposition 7. For any fixed price t_b , the seller optimally sets a cancellation fee t_u as the solution to the following equation.

$$V(\theta_0; t_b, t_u; \rho) - t_u = V(\theta_0; t_b; B) \quad (10)$$

For any optimal mechanism with a cancellation fee, the buyer obtains the same ex-ante surplus as if the mechanism prohibited return.

Lemma 1 is a special case of this proposition: if $\rho = B$, then $t_u = 0$ and the seller cannot benefit from charging a cancellation fee. When the learning process is more efficient after transaction, the seller deducts t_u from the original price as a non-refundable payment to extract the buyer's extra information rent from post-transaction learning, such that the buyer obtains the same ex-ante surplus as if the mechanism prohibited return; see equation (10). In addition, the return policy designed to induce some particular stopping belief θ is obtained in the same way as in Lemma 2.²⁹ The seller's safe revenue becomes the sum of the return transfer and the cancellation fee, $t_r + t_u$. A direct observation is that, if $B > k$ and the seller strictly prefers to encourage learning for some non-empty set of prior beliefs even if $\rho = B$ (Theorem 2), then her incentive to encourage learning becomes even stronger if $\rho > B$, and therefore the set of prior beliefs supporting Learning Deterrence shrinks.³⁰

7.4 A more general framework

Notice that a selling mechanism $f(t_b; (x_r; t_r))g$ can actually be interpreted as a binary menu $f(x_b = 1; t_b); (x_r; t_r)g$. Essentially, the seller commits to such a binary menu to screen the buyer who is initially uninformed but observes different information outcomes after private learning. In principle, the seller can design any arbitrary menu containing arbitrary numbers of allocation-transfer pairs. With our information process, it is without loss of generality to assume such a binary menu

²⁹With slight abuse of notation, $x_r(\theta; t_b) = \frac{V_1(\theta; t_b, t_u; \rho)}{v_h - v_l}$ and $t_r(\theta; t_b) = E(v_j) x_r(\theta; t_b)$

$V(\theta; t_b, t_u; \rho) - t_u$.

³⁰The revenue from Learning Deterrence remains constant regardless of the values of ρ .

$f(1; t_b); (x_r; t_r)g$.³¹ More concretely, the buyer observes the binary menu and decides how much information to acquire. Once he finds that the payoff from accepting either option within the menu weakly dominates the continuation value from further learning, he stops and chooses optimally between the two options. Anticipating this, the seller chooses among different binary menus to maximize her expected revenue. Within this framework, our results still hold.

To further elaborate, consider a standard adverse selection model where the agent’s private type (valuation) is supported on a closed interval. The revelation principle implies that the principal designs a menu that incentivizes the agent to truthfully reveal his type. In other words, there is a one-to-one mapping from the agent types to the menu options. Here, designing a menu to screen buyers with different posterior beliefs requires more constraints. First, the interim incentive constraints require buyers with different posterior beliefs to be willing to truthfully reveal these beliefs. Second, the buyer must be willing to stop learning when his posterior belief reaches either of the two posteriors for which the binary menu is designed (this is referred to as “implementability” in Mensch (2020)). In our paper, by assuming exponential experimentation, these two sets of constraints are directly implied by the buyer’s optimality to stop learning at some particular posterior belief. That is, we only require $(x_r; t_r)$ to satisfy the well-known smooth pasting and value matching conditions at that particular posterior belief (see Lemma 2). Furthermore, given the buyer’s optimality of stopping, the first sets of constraints (interim incentive constraints) are always slack; otherwise, learning would not be necessary.

Our reason for distinguishing between $(1; t_b)$ and $(x_r; t_r)$ is that they play very different roles in shaping the buyer’s learning behavior. Specifically, the selling price, $(1; t_b)$, determines the net surplus that the buyer obtains if a Poisson jump oc-

³¹The optimality of such a binary menu is implied by Lemmas 1 and 2. For fixed t_b , the seller is maximizing her expected revenue over the distribution of the buyer’s posterior beliefs. Given a binary state space, standard concavification implies the optimality of inducing binary posterior beliefs. Mensch (2020) also claims the optimality of binary menus under binary state space, but without requiring $x_b = 1$, as he allows the buyer to acquire flexible information with a posterior-separable information cost.

curs, and thereby determines the continuation value he can attain through learning. Meanwhile, the return policy, $(x_r; t_r)$, is designed to truncate the buyer's sequential learning, which allows the seller to induce more flexible buyer learning while keeping the price constant.

8 Conclusion

This paper discusses the seller's optimal selling mechanism when interacting with a buyer who can sequentially acquire information about a product. The seller designs the mechanism through an appropriately chosen return policy to influence the buyer's learning value and thereby indirectly controls the buyer's learning behavior. We find that the optimal selling mechanism either induces full learning or deters the buyer from private learning by offering the buyer an informational rent. Given that our results rely on the assumption that the buyer's prior belief is common knowledge, this naturally brings us to our next question: what if the buyer has private information about his valuation even before private learning begins? He then has an incentive to misreport his private information. What is the optimal mechanism that can screen buyers who are endowed with different prior beliefs but who can still acquire private information afterward? This would be an interesting follow-up question for future research.

Appendix

Proposition 1

Proof. We prove this proposition by verifying $Q(s) = q(s)$ if $s \in [\underline{s}; \bar{s}]$, while the equality holds at \underline{s} and \bar{s} . Recall that $Q(s) = f : V(\cdot; s) = E(v_j) - (v_h - s)g$. By setting $s = D(\cdot) := Q^{-1}(\cdot)$, the type- θ buyer is indifferent between accepting the price and exerting learning. Let $\tilde{v}(\cdot) := q(D(\cdot))$ be the quitting belief if $s = D(\cdot)$.

Claim 1. *The domain of $\tilde{v}(\cdot)$ is $[\underline{\cdot}; \bar{\cdot}]$. $\tilde{v}(\cdot) = \frac{1}{2}$ and the equality holds only at the two end points. $\tilde{v}(\cdot)$ is increasing and symmetric about the line $\frac{1}{2}$. $\tilde{v}(\cdot)$ increases first and then decreases in \cdot .*

Proof. Recall the definition of $D(\cdot)$,

$$V(\cdot; D(\cdot)) = E(v_j) - (v_h - D(\cdot))g \quad (11)$$

By implicit differentiation w.r.t. \cdot , we have,

$$\frac{dD(\cdot)}{d\cdot} = \frac{k(k - D(\cdot))}{2(1 - D(\cdot))^2} = \frac{k[\tilde{v} - 1]}{(1 - \tilde{v})^2} < 0 \quad (12)$$

Besides,

$$\frac{d\tilde{v}}{d\cdot} = \frac{d[k - D(\cdot)]}{d\cdot} = \frac{k^2(1 - D(\cdot)) - k}{3(1 - D(\cdot))^2 D(\cdot)^3} = \frac{\tilde{v}^2(1 - \tilde{v})}{(1 - \tilde{v})^2}.$$

Thus, $\tilde{v}(\cdot)$ is a differential equation with initial point $(\underline{\cdot}; \frac{1}{2})$,³² and its solution is,³³

$$\frac{1}{\tilde{v}} \log[1 - \tilde{v}] + \log[\tilde{v}] = \frac{1}{1} \log[1 - \frac{1}{2}] + \log[\frac{1}{2}] - \frac{(v_h - v_l)}{k}. \quad (13)$$

Denote the LHS as $f(\tilde{v})$ and the RHS as $g(\cdot)$. The domain of both functions is $[\underline{\cdot}; \bar{\cdot}]$ and $f(\cdot) = g(\cdot)$ at the two end points. Note that $f'(\cdot) > g'(\cdot)$ when the both arguments are smaller than 0.5 and $f'(\cdot) < g'(\cdot)$ when both arguments are larger than

³²To verify $(\underline{\cdot}; \frac{1}{2})$ is an initial point. Recall $\underline{\cdot} = q(s)$ and the binding **Learning-Feasibility** constraint implying $E(v_j) - (v_h - s)g = 0$. Meanwhile $V(q(s); s) = V(\underline{\cdot}; s) = 0$. Given equation (11), $D(\underline{\cdot}) = s$. Thus, $\tilde{v}(\underline{\cdot}) = q(D(\underline{\cdot})) = \frac{1}{2}$.

³³The general solution is $\frac{1}{\tilde{v}} \log[1 - \tilde{v}] + \log[\tilde{v}] = \frac{1}{1} \log[1 - \frac{1}{2}] + \log[\frac{1}{2}] + C$. Conditional on the initial point, $(\underline{\cdot}; \frac{1}{2})$, we can solve $C = -\frac{(v_h - v_l)}{k}$. Same result holds if we take $(\bar{\cdot}; \frac{1}{2})$ as the initial point.

0.5.³⁴ Therefore $f(\cdot)$ and $g(\cdot)$ cross only at the two boundary points and therefore $\tilde{v}(\cdot) < \underline{v}$ for all $\underline{s} \in [\underline{s}; \bar{s}]$. For $\tilde{v}(\cdot)$ to be symmetric about $1 - \underline{s}$, note that the reflection point of $(\cdot; \tilde{v})$ over line $1 - \underline{s}$ is $(1 - \tilde{v}; 1 - \underline{s})$. It is easy to verify that, if equation (13) holds at a point $(\cdot; \tilde{v})$, then equation (13) still holds at the reflection point $(1 - \tilde{v}; 1 - \underline{s})$. Now, we want to show that $\tilde{v}(\cdot)$ is single-peaked, increasing first and then decreasing in \underline{s} . Note that $\tilde{v}(\underline{s}) < 1$ and $\tilde{v}(\bar{s}) > 1$; therefore, if $\tilde{v}(\cdot) = 1$ has a unique solution, then we are done. To show this, $\frac{d\tilde{v}}{d\underline{s}} = \frac{\tilde{v}^2(1 - \tilde{v})}{(1 - \underline{s})^2} = 1$ implies $\tilde{v}(\cdot) = 1$.³⁵ As $\tilde{v}(\cdot)$ is increasing in \underline{s} and symmetric about $1 - \underline{s}$, it follows that $\tilde{v}(\cdot) = 1$ has a unique interior solution. \square

When $\underline{s} \notin [\underline{s}; \bar{s}]$, as the [Learning-Feasibility](#) constraint fails, no learning is optimal. If $\underline{s} \in [\underline{s}; \bar{s}]$, note that $Q(D(\underline{s})) - q(D(\underline{s})) = \tilde{v}(\underline{s})$. Taking the derivative with respect to \underline{s} yields $(Q' - q')D' = 1 - \tilde{v}'$. Because $D'(\underline{s}) < 0$, $Q'(\underline{s}) - q'(\underline{s})$ is positive for small \underline{s} and then negative for large \underline{s} , and $Q(\underline{s}) = q(\underline{s})$ at \underline{s} and \bar{s} .³⁶ The difference, $Q(\underline{s}) - q(\underline{s})$, is single-peaked in \underline{s} . That is, for all $\underline{s} \in [\underline{s}; \bar{s}]$, $Q(\underline{s}) - q(\underline{s})$ with equality holding at the two end points. Then it is easy to verify $V(\cdot; \underline{s}) = \max\{0; E(v_j) - (v_h - \underline{s})g\}$ if $\underline{s} \in [q(\underline{s}); Q(\underline{s})]$. Then the construction of Proposition 1 is optimal based on the standard arguments in the exponential experimentation. \square

Lemma 1

Proof. To simplify the exposition, we omit the notion of \underline{s} in the buyer's value function, as the lemma is true for any fixed \underline{s} . Let $V_B(\cdot)$ be the buyer's value function for pre-purchase learning (enter-the-market value). It is characterized by the Bellman equation below:

$$V_B(\cdot) = \max\{0; V_P(\cdot); kd + (\cdot) d s + (1 - (\cdot) d) V_B(\cdot + d)\}g; \quad (14)$$

³⁴ $f' = \frac{1}{1 - \underline{s}}$ and $g' = \frac{1}{(1 - \underline{s})^2}$.

³⁵ $\tilde{v}^2(1 - \tilde{v}) = (1 - \underline{s})^2$ could have three solutions: $\tilde{v} = 0$, $\tilde{v} = 1$ or $\tilde{v} = 1 - \underline{s}$. The previous two cannot be true when $\underline{s} \in [\underline{s}; \bar{s}]$.

³⁶ Recall that $D(\underline{s}) = \underline{s}$ and $D(\bar{s}) = \bar{s}$ by [Learning-Feasibility](#).

Different from the Bellman equation for No return, if the buyer stops learning by purchasing the item, he obtains the purchase value $V_P(\beta)$ instead of the consumption value $E(v_j | \beta) - (v_h - s)$, as he might also learn after purchase if a return is allowed. The purchase value $V_P(\beta)$ is characterized as below:

$$V_P(\beta) = \max\{E(v_j | \beta) - (v_h - s); E(v_j | \beta)X_r - t_r; kd + (\beta) d s + (1 - (\beta) d) V_P(\beta + d)\}g \quad (15)$$

Note that, while the buyer purchases the item, he instantaneously abandons his outside option. In other words, upon stopping, he can either consume the item or return it according to the pre-specified return policy. Conditional on learning, the three Bellman equations (14), (15), and (2) lead to the same differential equation (ODE). Obviously, $V_B(\beta) = V_P(\beta)$ and $V_B(\beta) = V^0(\beta)$.

To show this lemma, we prove the following equality:

$$\max\{V_B(\beta); V_P(\beta)g = V^0(\beta); \beta\} \quad (16)$$

Suppose the seller intends to set a harsh return policy with which, if the buyer purchases the item and stops post-purchase learning at belief β ,³⁷ he obtains payoff $V_P(\beta) < V_B(\beta)$. Then, a rational buyer could simply not purchase the item and perform pre-purchase learning, which implies $V_B(\beta) = V^0(\beta)$ and equality (16) holds.

Moreover, suppose the seller instead offers a benevolent return policy intending to reward the buyer for purchasing the item early on. Under this policy, the buyer purchases the item at some point and, while he stops post-purchase learning at belief β and requests a return, he gets payoff $V_P(\beta) > V^0(\beta)$. We can then calculate the return transfer which equals the allocation surplus minus the buyer's payoff:³⁸

$$t_r = E(v_j | \beta)X_r - V_P(\beta) = E(v_j | \beta) \frac{V_P^0(\beta)}{v_h - v_l} - V_P(\beta):$$

³⁷Since the buyer always stops learning if his belief jumps to 1, our use of the term stopping belief refers to the non-degenerate stopping beliefs.

³⁸This is implied by the optimality (known by smooth-pasting) to stop learning at β . Formally, see Lemma 2.

The (ODE) is a general solution of $V_P(\cdot)$. Hence,

$$(1 - \alpha) V_P^0(\cdot) + \alpha V_P(\cdot) = s - k:$$

Slope $V_P^0(\cdot)$ and magnitude $V_P(\cdot)$ of the buyer's continuation value are the substitutes that the seller can adjust to enforce the same stopping belief. For the purpose of maximizing profit, the seller will reduce $V_P(\cdot)$ and raise $V_P^0(\cdot)$ (constrained by the above differential equation) to increase the return transfer and in the meantime preserve the same buyer's optimal stopping rule,³⁹ which is a contradiction of optimality, and we obtain condition (16). \square

Corollary 1

Proof. Given Lemma 1, for all optimal selling mechanisms, $V_B(\cdot; s) = V^0(\cdot; s)$. Given Proposition 1, if $\alpha \geq [\underline{\cdot}]$, then $V_B(\cdot; s) = \max\{0; E(v|j=0) - (v_h - s)g\}$ and the buyer does not perform learning regardless of s . Suppose the seller offers a "No Return" mechanism with price $t_b = E(v|j=0)$. Then, the buyer is indifferent between purchasing and quitting, with the seller-preferred tie breaking rule, the seller obtains a revenue equal to $E(v|j=0)$. Note that the seller's profit equals the joint surplus minus the surplus that buyer obtains from the trade. In this case, the joint surplus attains the full allocation surplus and the buyer gets zero trading surplus. That is to say, the mechanism $\{E(v|j=0); (1; E(v|j=0))g\}$ is optimal. \square

Propositions 2 and 4

Proof. We prove these two propositions together. First, we solve the explicit solution for $s^F(\cdot)$ and $\alpha^F(\cdot)$. Denote $\phi(q(s); s) = \frac{\alpha}{1 - \alpha} \frac{q(s)}{q(s)} (v_h - s)$ as the objective function of (F), and it is concave given that the second order total derivative w.r.t s is negative.⁴⁰ Thus, the maximizer is pinned down by the first order condition, which leads to be solution below:

$$s^F(\cdot) = \frac{k}{\alpha} + \frac{\sqrt{k(\alpha - 1) - \alpha(k - v_h)}}{\alpha};$$

³⁹By inducing the same stopping beliefs, the ex-ante probabilities of return and successful sale are the same regardless of when the buyer purchases the item, i.e., switches to post-purchase learning.

⁴⁰ $\frac{d^2 \phi}{ds^2} = \frac{2k(\alpha - 1)(k - v_h)}{(k - s)^3} < 0$:

and

$$s^F(\theta) = \frac{2\sqrt{k(\theta-1)\theta(k-v_h)} + k - 2k\theta + \theta v_h}{2\theta\sqrt{k^2(\theta-1)\theta(k-v_h)}}.$$

Taking the derivative of $s^F(\theta)$ w.r.t θ gives:

$$\frac{ds^F}{d\theta} = \frac{k(k-v_h)}{2\theta\sqrt{k^2(\theta-1)\theta(k-v_h)}} < 0:$$

Hence, $v_h - s^F(\theta)$ is increasing in θ . Furthermore, $s^F(\theta)$ is increasing in θ due to the envelope theorem.

Next, we need to verify that, if $s^F(\theta) = t^D(\theta)$, then $q^{-1}(\theta) = s^F(\theta) = Q^{-1}(\theta)$. It is obvious that, if $s^F(\theta) = t^D(\theta)$, then $v_h - s^F(\theta) > t^D(\theta) = v_h - Q^{-1}(\theta)$, as the expected probability of a successful sale is less than one with Free Return. Hence, $s^F(\theta) = Q^{-1}(\theta)$ holds trivially. To show $q(s^F(\theta)) < \theta$, we plug in the explicit expression of $s^F(\theta)$ and obtain, $\sqrt{\frac{\theta}{1-\theta}} > \sqrt{\frac{k-(v_h)}{1-k-(v_h)}}$: This inequality is true because $\frac{k}{v_h} < \theta < 1$. \square

Proposition 3

Proof. With Learning Deterrence, $s = D(\theta) := Q^{-1}(\theta)$, trading happens with probability one and therefore the joint surplus attains the full allocation surplus $E(v_j - \theta)$.

First, we prove the first term. Proposition 1 and the **Learning-Feasibility** constraints imply $V(\theta; D(\theta)) = 0$ at $\underline{\theta}$ and $\bar{\theta}$. Rearranging equation (11) gives:

$$V(\theta; D(\theta)) = D(\theta) - (1 - \theta)(v_h - v_l):$$

Taking derivative w.r.t θ and plugging in equation (12) gives:

$$\frac{dV(\theta; D(\theta))}{d\theta} = (v_h - v_l) \left[A \frac{(1 - \theta)}{(1 - \theta)^2} + 1 \right];$$

where $A = \frac{k}{(v_h - v_l)} = (1 - \underline{\theta}) - 2(0; \frac{1}{4})$.⁴¹ It is easy to verify $\frac{dV(\theta; D(\theta))}{d\theta} = 0$ at $\underline{\theta}$ or $\bar{\theta}$. To prove that $V(\theta; D(\theta))$ is single-peaked in θ , we only need to show that $\frac{dV(\theta; D(\theta))}{d\theta} = 0$ has a unique solution when $\theta \in (\underline{\theta}, \bar{\theta})$, as $V(\theta; D(\theta)) > 0$ when

⁴¹From the binding **Learning-Feasibility** constraint, we can get $\frac{k}{(v_h - v_l)} = (1 - \underline{\theta}) - 2(0; \frac{1}{4})$. Therefore, $\underline{\theta} = 1 - 2(0; 0.5)$. Hence, $A \geq 2(0; \frac{1}{4})$.

$2(\underline{\cdot}; -)$. That is, the two equations below have a unique solution when $2(\underline{\cdot}; -)$, as \sim is the implicit solution of (13).

$$A \frac{(1 - \sim)}{(1 - \sim)^2} + 1 = 0 \quad (17)$$

$$\frac{1}{\sim} + \log \left[\frac{\sim}{1 - \sim} \right] = \frac{1}{1} + \log \left[\frac{1}{1} \right] - \frac{1}{A} \quad (18)$$

Substituting equation (17) into (18), we have,

$$\frac{A}{A(1 - \sim)^2} + \log \left[\frac{A(1 - \sim)^2}{(1 - \sim)^2} \right] - \left(\frac{1}{1} + \log \left[\frac{1}{1} \right] - \frac{1}{A} \right) = 0:$$

Denote the LHS as $h(\cdot)$. Now, we want to show that $h(\cdot) = 0$ has a unique solution for $2(\underline{\cdot}; -)$. In particular, as we can verify that $h(\cdot) = 0$ at $\underline{\cdot}$ and $-$, we want to show that $h(\cdot)$ first decreases and then increases and then decreases again on $[\underline{\cdot}; -]$.

Taking the derivative of $h(\cdot)$ w.r.t \cdot gives:

$$h'(\cdot) = \frac{1}{(1 - \cdot)^2} \left[\frac{y(\cdot)}{z(\cdot)} - 1 \right];$$

where $y(\cdot) := A^2(3 - 1)(1 - \cdot)$ and $z(\cdot) := [A(1 - \cdot)^2]^2$. $y(\cdot)$ is a second-order polynomial function that is negative when $\cdot < 1=3$, increases on \cdot if $\cdot < 2=3$, and decreases on \cdot if $\cdot > 2=3$. $z(\cdot)$ is a high-order polynomial function and $z'(\cdot) = 0$ has at most 4 roots: $1=3; 1$; and at most two roots from $(1 - \cdot)^2 - A = 0$.⁴² We can show that $z(\cdot)$ crosses $y(\cdot)$ twice in the support $[\underline{\cdot}; -]$, first from above and then from below.⁴³

Next, the monotonicity of $t^D(\cdot) = v_h - D(\cdot)$ can be directly obtained from (12). Moreover, $t^D(\cdot) = E(v_j | \cdot) - V(\cdot; D(\cdot))$ and $V(\underline{\cdot}; D(\underline{\cdot})) = V(-; D(-)) = 0$, therefore, $t^D(\underline{\cdot}) = E(v_j | \underline{\cdot})$ and $t^D(-) = E(v_j | -)$. \square

⁴² $z'(\cdot) = 2[(1 - \cdot)^2 - A](3 - 1)(-1)$. The derivative of $(1 - \cdot)^2 - A$ is $(3 - 1)(-1)$. Hence $(1 - \cdot)^2 - A$ is increasing if $\cdot < 1=3$ and decreasing afterwards. When $A < 4=27$, $(1 - \cdot)^2 - A = 0$ has two distinct roots, $r_1 < 1=3 < r_2$. When $A = 4=27$, there is a unique root $1=3$. When $A > 4=27$, there is no root. Regardless of A , $(1 - \cdot)^2 - A < 0$ when $\cdot = \underline{\cdot}; -$.

⁴³(1) Suppose $A < 4=27$, then $z(\cdot) > y(\cdot)$ for $\cdot = 1=3$, $z(r_2) = 0 < y(r_2)$ and $z(-) > y(-)$. Therefore, $z(\cdot)$ double crosses $y(\cdot)$. (2) Suppose $A = 4=27$, then $z(\cdot) > y(\cdot)$ for $\cdot < 1=3$, $z(1=3) = y(1=3)$, $z'(1=3) = 0 < y'(1=3)$, and $z(-) > y(-)$. Therefore, $z(\cdot)$ double crosses $y(\cdot)$. (3) Suppose $A \geq (4=27; 1=4)$, then $z'(\cdot) < 0$ when $\cdot < 1=3$, and $z'(\cdot) = 0$ when $\cdot = 1=3$. We can check that $z(1=2) < y(1=2)$ for $A \geq (4=27; 1=4)$, and hence we have the same double crossing given $y(\underline{\cdot}) < z(\underline{\cdot})$ and $y(-) < z(-)$.

Lemma 2

Proof. Given Lemma 1, $V_B(\beta; s) = V^0(\beta; s) - V_P(\beta; s)$ on the domain $[0; 1]$. To induce the buyer to stop learning at a belief β different from $q(s)$, $V_P(\beta; s)$ must be equal to $V^0(\beta; s)$. Otherwise, the buyer strictly prefers to continue his pre-purchase learning and does not stop. Furthermore, to ensure that it is a best response for the buyer to stop at belief β given the return policy $(x_r; t_r)$, the buyer's expected payoff from requesting return $E(v_j)x_r - t_r$ should smoothly pass $V^0(\beta; s)$ at β . Besides, the induced stopping belief β must belong to the set $[q(s); Q(s)]$, in which $V^0(\beta; s) = V(\beta; s)$. That is,

$$\text{value matching: } E(v_j)x_r - t_r = V(\beta; s);$$

$$\text{smooth pasting: } \frac{d[E(v_j)x_r - t_r]}{d\beta} = V_1(\beta; s);$$

We then obtain the expression of x_r and t_r . Specifically,

$$t_r(\beta; s) = \frac{kv_l - v_l s - k v_h \left[\log\left(\frac{\beta}{1-\beta}\right) - \log\left(\frac{k}{s-k}\right) \right]}{(v_h - v_l)}; \quad (19)$$

Taking partial derivative w.r.t β and s separately gives:

$$\frac{\partial t_r(\beta; s)}{\partial \beta} = \frac{kE(v_j)}{(1-\beta)^2(v_h - v_l)} > 0;$$

$$\frac{\partial t_r(\beta; s)}{\partial s} = \frac{E(v_j q(s))}{(1-q(s))(v_h - v_l)} > 0;$$

and the cross derivative is 0. Moreover, as $V(\beta; s)$ is convex in β , $x_r(\beta; s)$ —proportional to $V_1(\beta; s)$ —is therefore increasing in β . \square

Lemma 3

Proof. First, we discuss the first-order condition. Explicitly,

$$V_1(\beta; s) = \frac{(1-\beta)}{(1-\beta)^2(v_h - v_l)} \underbrace{\left[v_h(v_h + s + v_l) + \frac{k((v_h - 2v_l) + v_l)}{2} + \frac{kv_h(\log[\frac{\beta}{1-\beta}] - \log[\frac{k}{s-k}])}{2} \right]}_{(1)}.$$

Since $\beta \in [q(s); Q(s)]$, $V_1(\beta; s) = 0$ has the same solution with $(1) = 0$.

$$\beta(\beta) = \frac{k(1-\beta)^2 v_h + 2k(1-\beta)^2 v_l}{(1-\beta)^3}.$$

The numerator of $\theta(\cdot)$ is a well-behaved second-order polynomial, which is verified to have a unique root between 0 and 1, and is larger than 0 at $\beta = 0$, and smaller than 0 at $\beta = 1$. Thus, $\theta(\cdot)$ crosses 0 only once and from below, which implies $\beta(\cdot)$ is initially decreasing and then increasing. Therefore, $\beta(\cdot)$ has at most two roots in $[0; 1]$, denoted as $\beta_-(s)$ and $\beta_+(s)$. Furthermore, $\beta(\cdot)$ is increasing in s . Therefore, the smaller root is the local maximizer of $\beta(\cdot; s)$ which is increasing in s , while the larger root is the local minimizer of $\beta(\cdot; s)$ which is decreasing in s , and if the two roots coincide, $\beta(s) = \beta_+(s) > 0.5$.⁴⁴ Thus, if there exists a $\beta_+(s)$, it is larger than 0.5.

Let $s(\cdot) = f\beta : \beta_1(\cdot; s) = 0$.⁴⁵ Given the above argument, it is a single-valued continuous function, which is initially increasing and then decreasing in β . Furthermore, it is clear that when $\beta > 0.5$, $s(\cdot)$ is increasing. To introduce one more notation, let $t_r(\cdot) := t_r(\cdot; Q^{-1}(\cdot))$. It is the envelope of all inducible return transfers. Formally,

$$t_r(\cdot; s) \in [0; t_r(\cdot)](\cdot) \subseteq [q(s); Q(s)]:$$

To see this, consider the direction from the right to the left first. Recall that $t_r(\cdot; s)$ is increasing in both arguments. If $\beta = q(s)$, then $t_r(\cdot; s) = t_r(q(s); s) = 0$; and if $\beta = Q(s)$, then $s = Q^{-1}(\beta)$ as $Q(s)$ decreases in s , which then implies $t_r(\cdot; s) = t_r(\cdot; Q^{-1}(\beta))$. The opposite direction is trivial.⁴⁶

To prove Lemma 3, we want to show that $\beta(\cdot; s)$ is quasi-concave on $\beta \in [q(s); Q(s)]$. Specifically, we show $t_r(\beta_+(s); s) > t_r(\beta_-(s))$, which then implies $\beta_+(s) > Q(s)$. The following claim pins down the set of β such that $t_r(\cdot; s(\beta)) \in [0; t_r(\cdot)]$.

Claim 2. $t_r(\cdot)$ with domain $[\underline{\beta}; \bar{\beta}]$ first increases and then decreases in β . $t_r(\cdot; s(\beta))$ single crosses $t_r(\cdot)$ at 0.5 from below, and $f : t_r(\cdot; s(\beta)) \in [0; t_r(\cdot)] \iff \beta \in [\underline{\beta}; 0.5]$.

Proof. It is obvious that $t_r(\cdot) = 0$ when $\beta \in [\underline{\beta}; \bar{\beta}]$, with equality hold at the two

⁴⁴To see this, note that $\theta(0.5) < 0$. Suppose $\beta(s) = \beta_+(s) = 0.5$, then $\theta(0.5) = 0$. Contradiction. Suppose $\beta(s) = \beta_+(s) < 0.5$, then $\theta(0.5) > 0$. Contradiction.

⁴⁵Sorry to abuse the notation. We can verify that if $\beta \in [\underline{\beta}; 0.5]$, $s(\cdot)$ is the inverse function of $\beta(\cdot)$ after we prove this lemma.

⁴⁶Note that equation (19), the exact expression of $t_r(\cdot; s)$, is valid for all $\beta \in [0; 1]$.

end points. Recall that $D(\cdot) := Q^{-1}(\cdot)$. Taking derivative of $\bar{t}_r(\cdot)$ w.r.t \cdot gives

$$\frac{d\bar{t}_r(\cdot)}{d\cdot} = \frac{\partial t_r(\cdot; D(\cdot))}{\partial \cdot} + \frac{\partial t_r(\cdot; D(\cdot))}{\partial S} \frac{dD(\cdot)}{d\cdot} = \frac{(1 - \cdot)}{(1 + \cdot)} \left[\frac{E(v_j)}{1} - \frac{E(v_j - (\cdot))}{1} \right];$$

The term in square brackets is decreasing. It's positive when $\cdot = \cdot^-(\cdot) = \cdot_-$, and negative when $\cdot = \cdot^-(\cdot) = \cdot^-$. Hence, $t_r(\cdot)$ is increasing first and then decreasing.

Next, we show that $t_r(\cdot; S(\cdot)) = t_r(\cdot)$ has a unique solution of 0.5. Since $t_r(\cdot; S)$ is increasing in S , to find the solution of $t_r(\cdot; S(\cdot)) = t_r(\cdot; Q^{-1}(\cdot))$ is equivalent to find the solution to the system of equations below,

$$\begin{cases} \tau_1(\cdot; S) = 0; \\ V(\cdot; D) = E(v_j) - (v_h - D); \\ S = D; \end{cases}$$

which can be verified to have a unique non-negative solution $\cdot = 0.5$. This suggests that $\tau_1(\cdot; D(\cdot)) = \tau_1(\cdot; Q^{-1}(\cdot)) = 0$ has a unique solution at 0.5. Moreover,

$$\begin{aligned} \frac{dt_r(\cdot; D(\cdot))}{d\cdot} &= \frac{\partial t_r(\cdot; D)}{\partial \cdot} + \frac{\partial t_r(\cdot; D)}{\partial S} \frac{dD}{d\cdot}; \\ \frac{dt_r(\cdot; S(\cdot))}{d\cdot} &= \frac{\partial t_r(\cdot; S)}{\partial \cdot} + \frac{\partial t_r(\cdot; S)}{\partial S} \frac{dS}{d\cdot}. \end{aligned}$$

Since $\frac{\partial t_r(\cdot; S)}{\partial \cdot}$ is independent of S , the first term of the two derivatives are the same. Besides, $\frac{dD}{d\cdot} < 0$ and $\frac{dS}{d\cdot} > 0$ if $\cdot > 0.5$. Hence, the slope of \bar{t}_r is smaller than $t_r(\cdot; S(\cdot))$. That is, if we reduce \cdot from 0.5, $t_r(\cdot; S(\cdot))$ decreases faster than $\bar{t}_r(\cdot)$. Let \cdot_- be the solution of $t_r(\cdot; S(\cdot)) = 0$. Obviously, $\cdot_- \geq \cdot_-(\cdot; 0.5)$. To pin down \cdot_- , note that $t_r(\cdot; S) = 0$ implies $S = q^{-1}(\cdot)$. Thus, \cdot_- is the solution that $\tau_1(\cdot; q^{-1}(\cdot)) = 0$. Explicitly,

$$\tau_1(\cdot; q^{-1}(\cdot)) = \frac{(1 - \cdot)(\cdot^2 v_h (v_h - v_l) - k(2(v_h - v_l) + v_l))}{(1 + \cdot)^2 (v_h - v_l)} = 0;$$

which also has a unique solution that $\cdot_- = \frac{k}{v_h} + (\frac{k}{v_h}(\frac{k}{v_h} + \frac{v_l}{v_h - v_l}))^{\frac{1}{2}}$.⁴⁷ Therefore, we pin down the set $[\cdot_-; 0.5]$ on which $t_r(\cdot; S(\cdot)) \geq [0; t_r(\cdot)]$. \square

⁴⁷Since $\cdot^2 v_h (v_h - v_l) - k(2(v_h - v_l) + v_l)$ is increasing on $\cdot > 0$ (its derivative is $2(k - v_h)(v_h - v_l) > 0$), it is negative when \cdot is small and positive when \cdot is large. Hence, $\tau_1(\cdot; q^{-1}(\cdot))$ single crosses 0 from above and \cdot_- is unique.

From this claim, we can see that $t_r(\cdot; s(\cdot)) > t_r(\cdot)$ if $\beta > 0.5$. Moreover, given that $\beta_+(s) > 0.5$, if there exists a local maximizer $\beta_+(s)$, it is larger than $Q(s)$. Therefore, $(\cdot; s)$ is quasi-concave on $[q(s); Q(s)]$.

Denote $t_r(\cdot) := t_r(\cdot; s(\cdot))$ for the domain $[\underline{\beta}; 0.5]$. Given the monotonicity of $s(\cdot)$ when $\beta > 0.5$, we can conclude that if $s \geq (q^{-1}(\underline{\beta}); Q^{-1}(0.5))$, the local maximizer $\beta_+(s) \geq (q(s); Q(s))$ hence $\beta_+(s) = \beta_+(s)$ is the global maximizer. Besides, if $s < Q^{-1}(0.5)$, then $Q(s) < 0.5 < \beta_+(s)$, where the first inequality comes from $Q(s)$ being decreasing in s and the second inequality comes from $\beta_+(Q^{-1}(0.5)) = 0.5$. The inequality holds with equality only at $s = Q^{-1}(0.5)$. It is optimal to induce a return belief $Q(s)$. If $s < q^{-1}(\underline{\beta})$, $q(s) < \underline{\beta}$ and then $\beta_1(q(s); s) < 0$.⁴⁸ Since $\beta_1(\cdot; s)$ is quasi-concave in $[q(s); Q(s)]$, thus if $\beta_1(\cdot; s) < 0$ at $q(s)$, $\beta_1(\cdot; s) < 0$ for all $[q(s); Q(s)]$. Still the inequality holds with equality only at $s = q^{-1}(\underline{\beta})$. It is optimal to induce return belief $q(s)$. \square

Theorem 1

Proof. Recall that $(\cdot; s(\cdot)) = t_r(\cdot) + \frac{\partial t_r(\cdot; s(\cdot))}{\partial \beta}(\beta_0)$. Taking the derivative w.r.t β gives

$$\begin{aligned} \frac{d(\cdot; s(\cdot))}{d\beta} &= \frac{dt_r}{d\beta} + \frac{\partial t_r}{\partial \beta} + (\beta_0) \frac{\partial^2 t_r}{\partial \beta^2} = \frac{\partial t_r}{\partial s} \frac{ds}{d\beta} + (\beta_0) \frac{\partial^2 t_r}{\partial \beta^2} = \left[\frac{\frac{\partial t_r}{\partial s} \frac{ds}{d\beta}}{\frac{\partial^2 t_r}{\partial \beta^2}} + \beta_0 \right] \frac{\partial^2 t_r}{\partial \beta^2} \\ &= \left[\frac{(1 - \beta_0)}{v_h} E[v_j q(s(\cdot))] + \beta_0 \right] \frac{\partial^2 t_r}{\partial \beta^2}. \end{aligned}$$

Note that $\frac{\partial t_r(\cdot; s(\cdot))}{\partial \beta}$ is independent of s and we can verify $\frac{\partial^2 t_r}{\partial \beta^2} < 0$.⁴⁹ Let $(\cdot) = \frac{(1 - \beta_0)}{v_h} E[v_j q(s(\cdot))]$. The monotonicity of $(\cdot; s(\cdot))$ can be pinned down by the sign of $\beta_0 - (\cdot)$. In particular, if $\beta_0 > (\cdot)$, $(\cdot; s(\cdot))$ is decreasing in β , otherwise, it is increasing in β .

Claim 3. (\cdot) with domain $[\underline{\beta}; 0.5]$ is decreasing and convex on β , and $\beta(0.5) > 1$.

The proof of this claim can be found subsequent to this theorem. Recall Lemma 3,

⁴⁸See footnote 41.

⁴⁹ $\frac{\partial^2 t_r}{\partial \beta^2} = \frac{k[(2 - \beta_0)E(v_j) - (1 - \beta_0)v_l]}{(1 - \beta_0)^2 - 3(v_h - v_l)} < 0$.

(s) is an optimal solution only for $s \geq [q^{-1}(\underline{\mu}); Q^{-1}(0.5)]$. In particular, partial learning is optimal for $s \geq (q^{-1}(\underline{\mu}); Q^{-1}(0.5))$; and for the boundaries, either full learning or no learning is optimal. Consider the original problem (P) and reimpose the constraints $q^{-1}(\mu_0) \leq s \leq Q^{-1}(\mu_0)$ and μ_0 that relate to the prior belief, then (s) could be an optimal solution only if

$$[q^{-1}(\underline{\mu}); Q^{-1}(0.5)] \setminus [q^{-1}(\mu_0); Q^{-1}(\mu_0)] \neq \emptyset \text{ and } \mu_0 \in \underline{\mu};$$

which is equivalent to

$$\mu_0 \geq [\underline{\mu}; Q(q^{-1}(\underline{\mu}))]:$$

That is, partial learning could be an optimal solution only for μ_0 that belongs to this region,⁵⁰ which is depicted in Figure 9. When $\mu_0 \geq [\underline{\mu}; 0.5]$ and conditional on $s \geq [\underline{\mu}; 0.5]$ (along the path of $t_r(\cdot)$), the optimal (Stochastic Return) stopping belief is constrained by $s \geq [\underline{\mu}; \mu_0]$; while if $\mu_0 \geq [0.5; Q(q^{-1}(\underline{\mu}))]$, the optimal (Stochastic Return) stopping belief is constrained by $s \geq [\underline{\mu}; (Q^{-1}(\mu_0))]$.




Figure 9: Seller's expected revenue is quasi-convex

We distinguish two cases. First, $\theta(\underline{\mu}) = 1$, which implies that $(\cdot; s(\cdot))$ is quasi-convex in \cdot . Second, $\theta(\underline{\mu}) < 1$ and there exists a local maximum of $(\cdot; s(\cdot))$,

⁵⁰If $\mu_0 \geq Q(q^{-1}(\underline{\mu}))$, then $Q^{-1}(\mu_0) = q^{-1}(\underline{\mu})$; if $\mu_0 \in \underline{\mu}$, then $q^{-1}(\mu_0) = q^{-1}(\underline{\mu})$. Thus, $[q^{-1}(\underline{\mu}); Q^{-1}(0.5)] \setminus [q^{-1}(\mu_0); Q^{-1}(\mu_0)] \neq \emptyset$. For the opposite direction, since $\mu_0 \in \underline{\mu}$ and then $q^{-1}(\mu_0) = q^{-1}(\underline{\mu})$, then we require $Q^{-1}(\mu_0) = q^{-1}(\underline{\mu})$, which implies $\mu_0 \geq Q(q^{-1}(\underline{\mu}))$.

which we can verify is strictly worse than the revenue from Learning Deterrence. We establish the proof case by case.

Case one: $\theta(\underline{\mu}) \geq 1$. This is true in most scenarios. Denote $\phi(\mu) = \mu + \psi(\mu)$. Therefore when $\mu_0 \geq [\underline{\mu}; (0.5)]$, μ_0 single-crosses $\phi(\mu)$ from above, as depicted in Figure 10, where the black lines represent the contour lines of μ_0 for different μ_0 .




Figure 10: Case one: seller's expected revenue is quasi-convex

- If $\mu_0 < \underline{\mu}$, $\phi(\mu; s(\mu))$ is increasing in μ , and therefore Learning Deterrence is optimal. To verify this, when $\mu_0 < 0.5$, the optimal return belief is μ_0 and inducing no learning via Stochastic Return is strictly dominated by Learning Deterrence. When $\mu_0 \geq 0.5$, the feasible return belief is $\mu \geq [\underline{\mu}; (Q^{-1}(\mu_0))]$ with $(Q^{-1}(\mu_0)) < 0.5 < \mu_0$. However, we can show that Learning Deterrence is better than Stochastic Return that induces stopping at $(Q^{-1}(\mu_0))$. To see this, suppose we ignore the constraint that $\mu \geq [\underline{\mu}; (Q^{-1}(\mu_0))]$, then the seller obtains larger revenue by inducing stopping at 0.5 with revenue as a weighted average between the selling price $t^D(0.5) = v_h(Q^{-1}(0.5))$ and the return transfer $t_r(0.5; Q^{-1}(0.5))$, which is smaller than $t^D(0.5)$ and since $t^D(\cdot)$ is increasing, thus $t^D(\mu_0) > t^D(0.5)$ and Learning Deterrence generates strictly higher profit.⁵¹

- If $\mu_0 \geq [\underline{\mu}; (0.5)]$, $\phi(\mu; s(\mu))$ is quasi-convex in μ . When $\mu_0 < 0.5$, the

⁵¹The magnitude between $\underline{\mu}$ and 0.5 is ambiguous, but it does not affect the above argument.

optimal return belief is either $\underline{\mu}$ or μ_0 , which implies the optimality between Free Return and Learning Deterrence. When $\mu_0 > 0.5$, we can still obtain the optimality between Free Return and Learning Deterrence by applying the same reasoning as above.

- If $\mu_0 < (0.5)$, $(r_1(\mu_0); s(r_1(\mu_0)))$ is decreasing in μ_0 . Hence Free Return is optimal.




Figure 11: Case two: seller's expected revenue is not quasi-convex

Case two: When $\theta(\underline{\mu}) < 1$, there exists a local maximizer of $(r_1(\mu_0); s(r_1(\mu_0)))$. Denote $r = f^{-2}[\underline{\mu}; 0.5] : \theta(r) = 1g$. If $\mu_0 \geq [r; \underline{\mu}]$, there exists a unique local maximizer $r_1(\mu_0) = f^{-2}[\underline{\mu}; r] : \theta(r_1(\mu_0)) = \mu_0 g$ (see Figure 11 for visualization). If $\mu_0 \geq [r; \underline{\mu}]$, then the expected revenue is quasi-convex and the argument in case one validates. We want to show that if $\mu_0 \geq [r; \underline{\mu}]$,

$$(r_1(\mu_0); s(r_1(\mu_0))) < t^D(\mu_0):$$

With slight abuse of notation, we write $(r_1(\mu_0); s(r_1(\mu_0)); \mu_0)$ instead of $(r_1(\mu_0); s(r_1(\mu_0)))$. Note that

$$(r_1(\mu_0); s(r_1(\mu_0)); \mu_0) < (r_1(\mu_0); s(r_1(\mu_0)); \underline{\mu}) < (\underline{\mu}; s(\underline{\mu}); \underline{\mu}):$$

The first inequality comes from μ_0 increasing in μ_0 . The second inequality is due to $\underline{\mu} = r_1(\underline{\mu})$, which is the local maximizer of θ when $\mu_0 = \underline{\mu}$. Recall equation (9) and plug in the expression of $\underline{\mu}$,

$$(\underline{\mu}; s(\underline{\mu}); \underline{\mu}) = 0 + (\underline{\mu} - \underline{\mu}) \frac{\partial t_r(\underline{\mu}; s(\underline{\mu}))}{\partial \mu} \Big|_{\mu=\underline{\mu}} = E(v_j \frac{k}{v_h}):$$

It is obvious that $s < v_h$ whenever learning is feasible. Thus,

$$\mathbb{E}(vj\frac{k}{v_h}) < \mathbb{E}(vj_{-}) = t^D(\underline{s}) < t^D(0);$$

where the equality and the second inequality come from Proposition 3. \square

Claim 3

Proof. Denote $w(\underline{s}) := \mathbb{E}[vjq(s(\underline{s}))]$, then $\phi(\underline{s}) = \frac{1}{v_h}w(\underline{s})$. Note that $w(\underline{s})$ is decreasing in \underline{s} , as $q(s)$ decreases in s and $s(\underline{s})$ increases in \underline{s} . Besides, we can verify that $s(\underline{s})$ is concave for $\underline{s} \in [\underline{s}^*; 0.5]$,⁵² hence $w(\underline{s})$ is convex. Note that

$$\begin{aligned} \phi'(\underline{s}) &= \frac{1}{v_h} [w(\underline{s}) - (1 - \phi(\underline{s}))w'(\underline{s})] \\ &= \frac{1}{v_h} \left[w(\underline{s}) + \int_{\underline{s}}^{\phi(\underline{s})} w'(\underline{s}) d\underline{s} - (1 - \phi(\underline{s}))w'(\underline{s}) \right]: \end{aligned}$$

Since $w' < 0$ and $w'' > 0$, then $\int_{\underline{s}}^{\phi(\underline{s})} w'(\underline{s}) d\underline{s} - (1 - \phi(\underline{s}))w'(\underline{s})$ is decreasing in \underline{s} and therefore $\phi'(\underline{s})$ is increasing in \underline{s} . That is, $\phi(\underline{s})$ is convex.

Denote $q(\underline{s}) := q(s(\underline{s}))$. Simplifying $\phi'(0.5)$ gives:

$$\phi'(0.5) = \left[\frac{4(v_h - v_l)v_l q(0.5)^2}{v_h^2} (1 - q(0.5)) + \frac{1}{v_h} \mathbb{E}[vjq(0.5)] \right]:$$

We can show that $\phi'(0.5)$ is decreasing in v_l . Hence plugging $v_l = 0$ and $v_l = v_h$ into $\phi'(0.5)$, we have $\phi'(0.5)|_{v_l=0} = q(0.5) > 1$ and $\phi'(0.5)|_{v_l=v_h} = 1$.⁵³ \square

Theorem 2

Proof. First, we want to show that F is either an empty set or a closed interval. Note that $t^D(\underline{s}) > F(\underline{s})$ and $t^D(-) > F(-)$. Hence, it is equivalent to show $F(\underline{s})$ crosses $t^D(\underline{s})$ at most twice. Let $F(\underline{s}) = t^D(\underline{s}) = v_h - D(\underline{s})$, then $D(\underline{s}) =$

⁵²We can verify that $\frac{d^2 s}{d \underline{s}^2}$ is proportional to $q(s(\underline{s}))^2 M + 2N$, where $M = (v_h - 4v_l) - 2^2(v_h - v_l) + 2v_l)^2$ and $N = (2 + (5 - 4))v_h^2 + 2(1 - \underline{s})^2(3 + 2)v_l v_h$. We can verify that $M > 0$, $N < 0$, and $M + N < 0$. Meanwhile $q(s(\underline{s})) < 1$. Therefore $\frac{d^2 s}{d \underline{s}^2} < 0$.

⁵³Taking implicit differentiation w.r.t v_l for $\phi(\underline{s}(\underline{s}))|_{\underline{s}=0.5} = 0$, we have $\frac{ds(\underline{s})}{dv_l} = q(0.5) - 1 < 0$. Then $\frac{dq(0.5)}{dv_l} > 0$. Besides, $q(0.5) < 0.5$, then $\frac{d\phi'(0.5)}{dv_l} = \frac{1}{v_h} [(q(0.5) - 1)(v_h + 4(v_h - 2v_l)q(0.5)^2) - (v_h - v_l)[v_h + 4v_l(2 - 3q(0.5))q(0.5)] \frac{dq(0.5)}{dv_l}] < 0$.


$v_h = F(\cdot)$. To simplify the exposition, let $(\cdot) := \frac{(\cdot)}{(v_h - F(\cdot))}$. Recall equation (13) and $\tilde{(\cdot)} = \frac{(\cdot)}{D(\cdot)}$. Then we want to show $g(\cdot) - f(\cdot)$ has at most two roots when $\geq [\underline{\cdot}; \bar{\cdot}]$. To verify this,

$$g^\theta - f^{\theta\theta} = \frac{1}{1} \left(\frac{1}{1} + \frac{1}{1} \right) + \left(\frac{1}{\theta} - \frac{r}{1} \right) \frac{(\theta - + r\theta)^3}{(\theta - + r\theta)^2 - 1};$$

where $r = \sqrt{v_h - 1} > \frac{\theta}{3}$ given the assumption that $v_h > 4 + v_l$. Let $x = \sqrt{1 - 2(0; 1)}$, which is a monotone transformation of \cdot . Rearranging $g^\theta - f^{\theta\theta} = 0$, we have

$$m(x) := \frac{x(x+r)^3(1-xr)}{(1+x^2)^3(1+2xr+r^2)} = 1;$$

where $m(x)$ is a rational function. The degree of the numerator is smaller than that of the denominator, thus it has a horizontal asymptote $m = 0$. Note that the denominator is positive due to $(\cdot) \geq [0; 1]$, hence it does not have a vertical asymptote. Meanwhile $\lim_{x \rightarrow 0} m(x) = 0$, $\lim_{x \rightarrow 1} m(x) = 0$, $m(x=1) < 0$, and $m(x) = 0$ has a unique root $x = 1-r < 1$. Therefore, the graph of $m(x)$ is the following.



Then, $m(x) = 1$ has at most two roots. That is, if $\geq [\underline{\cdot}; \bar{\cdot}]$, $g^\theta(\cdot) - f^{\theta\theta}(\cdot)$ has at most two roots and $g^\theta(\cdot) - f^{\theta\theta}(\cdot) < 0$ when \cdot is between the two roots. Given that $g(\cdot) - f(\cdot)$ is strictly positive at $\underline{\cdot}$ and $\bar{\cdot}$, we can verify $g(\cdot) - f(\cdot)$ has at most two roots. The existence of \cdot is implied by Proposition 5, and the limit result in Proposition 6 implies that the left endpoint of F is larger than $\frac{v_l}{v_h}$. The exact form of optimal selling mechanism is an immediate result of Corollary 1 and Theorem 1. \square

Proposition 5

Proof. Recall that $(q(s); s) = \frac{0}{1-s}(v_h - s)$. By the envelope theorem, we have:

$$\frac{d^F}{d} = \frac{v_h - s^F}{(s^F)^2} (0 - 1) s^F < 0:$$

Hence, F is decreasing in s .

To show that $t^D(0) = v_h - D(0)$ is increasing in α , we want to show $D(0)$ is decreasing in α . Recall that $\tilde{v}(\alpha) = q(D(\alpha))$. Taking the derivative w.r.t α for both sides of $E(v_j | 0) - (v_h - D(0)) = V(0; D(0))$, we obtain:

$$\frac{1}{1 - \tilde{v}(\alpha)} \frac{dD}{d\alpha} = \frac{1}{1 - \tilde{v}(\alpha)} - 1 - (1 - \alpha) \log \left[\frac{\alpha=1 - \alpha}{\tilde{v}(\alpha)=1 - \tilde{v}(\alpha)} \right] < 0:$$

Given that F is either empty or a closed interval, it is immediate that if $\alpha_1 < \alpha_2$, then $F(\alpha_2) \supseteq F(\alpha_1)$. Note that α_* is the smaller root for $E(v_j | \alpha) - (v_h - q^{-1}(\alpha)) = 0$. By implicit differentiation,

$$\frac{d}{d\alpha} \left(\frac{1}{1 - \tilde{v}(\alpha)} - 1 \right) = 1:$$

Hence, $\frac{d}{d\alpha} > 0$. Meanwhile, $\alpha_* = 1 - \alpha_*$, then $[\alpha_*(\alpha_2); \alpha_*(\alpha_2)] \supseteq [\alpha_*(\alpha_1); \alpha_*(\alpha_1)]$. \square

Proposition 6

Proof. First we calculate the limit of $t^D(0)$ when $\alpha \rightarrow 0$. Plugging $D(0) = \frac{0}{1-\alpha}$ into equation (13) and multiplying by α gives:

$$D(0) + \alpha \log \frac{1}{D(0)} = \frac{1}{1-\alpha} + \alpha \log \frac{1}{1-\alpha} (v_h - v_l):$$

If $\alpha \rightarrow 0$ and α does not converge to 0 or 1, the above equation converges to $v_h - D(0) = v_l$.⁵⁴ Hence, $\lim_{\alpha \rightarrow 0} t^D(0) = v_l$. For the expected revenue from Free Return,

$$\lim_{\alpha \rightarrow 0} F(0) = \alpha v_h + (1 - 2\alpha) \cdot 2\sqrt{(1-\alpha)\alpha(v_h - v_l)} + \alpha v_h:$$

⁵⁴ $\lim_{\alpha \rightarrow 0} \alpha \log \frac{1}{D(0)} = 0$

Therefore when $\alpha \neq 0$, the seller is indifferent between Learning Deterrence and Free Return at $\alpha_0 = \frac{v_l}{v_h}$.

Second, since the above limit of $t^D(\alpha_0)$ may fail when $\alpha_0 \neq 0$ or $\alpha_0 \neq 1$, we have to verify the extreme case that $\lim_{\alpha \rightarrow 0} [\alpha; -] \neq [0; 1]$. Plugging $\alpha = \frac{1}{2} \left(1 + \sqrt{1 - 4 \frac{v_l}{v_h}} \right)$, we have

$$\lim_{\alpha \rightarrow 0} \log \frac{1 - \alpha}{1 + \alpha} = \log \frac{1 - \frac{1}{2} \left(1 + \sqrt{1 - 4 \frac{v_l}{v_h}} \right)}{1 + \frac{1}{2} \left(1 + \sqrt{1 - 4 \frac{v_l}{v_h}} \right)} \neq 0;$$

Hence, $\lim_{\alpha \rightarrow 0} t^D(\alpha) \neq v_l$. Thus, when $\alpha_0 < \frac{v_l}{v_h}$, the seller's expected revenue from the optimal mechanism converges to v_l .

Plugging $\alpha = \frac{1}{2} \left(1 + \sqrt{1 - 4 \frac{v_l}{v_h}} \right)$, we have

$$\lim_{\alpha \rightarrow 0} \log \frac{1 - \alpha}{1 + \alpha} = \log \frac{1 - \frac{1}{2} \left(1 + \sqrt{1 - 4 \frac{v_l}{v_h}} \right)}{1 + \frac{1}{2} \left(1 + \sqrt{1 - 4 \frac{v_l}{v_h}} \right)} \neq 0;$$

$$\lim_{\alpha \rightarrow 0} \frac{1 - \alpha}{1 + \alpha} = \frac{1 - \frac{1}{2} \left(1 + \sqrt{1 - 4 \frac{v_l}{v_h}} \right)}{1 + \frac{1}{2} \left(1 + \sqrt{1 - 4 \frac{v_l}{v_h}} \right)} \neq \frac{v_h - v_l}{v_h + v_l};$$

Hence, $\lim_{\alpha \rightarrow 0} \alpha \neq 1$, $\lim_{\alpha \rightarrow 0} t^D(\alpha) \neq v_h$, and $\lim_{\alpha \rightarrow 0} F(\alpha) = v_h$. If $\frac{v_l}{v_h} < \alpha_0 < 1$, $\lim_{\alpha \rightarrow 0} F(\alpha) > \lim_{\alpha \rightarrow 0} t^D(\alpha)$. Then, when $\alpha_0 < \frac{v_l}{v_h}$, the seller's expected revenue converges to $\alpha_0 v_h$. \square

Proposition 7

Proof. For any fixed selling price t_b , the lower bound of the buyer's ex-ante surplus is still $V^0(\alpha_0; t_b; B)$ as the buyer can always treat the mechanism as if a return is prohibited. Note that for a mechanism $f(t_b; (x_r; t_r; t_u))g$, the buyer's value function for post-transaction learning $V_P(\alpha; t_b; t_u; \rho)$ is characterized by the Bellman equation below (we omit the notion of ρ for simplification):

$$\begin{aligned} V_P(\alpha; t_b; t_u) = \max \{ & \alpha E(v_j(\alpha)) - t_b; E(v_j(\alpha))x_r - (t_r + t_u); \\ & kd + (1 - \alpha) d (v_h - t_b) + (1 - \alpha) d V_P(\alpha + d; t_b; t_u) \}; \end{aligned} \quad (20)$$

Adding t_u to both sides,

$$\begin{aligned} V_P(\alpha; t_b; t_u) + t_u = \max \{ & \alpha E(v_j(\alpha)) - (t_b - t_u); E(v_j(\alpha))x_r - t_r; \\ & kd + (1 - \alpha) d (v_h - (t_b - t_u)) + (1 - \alpha) d (V_P(\alpha + d; t_b; t_u) + t_u) \}; \end{aligned} \quad (21)$$

the buyer's value function for post-transaction learning becomes

$$V_P(\cdot; t_b; t_u) = V_P(\cdot; t_b - t_u) - t_u.$$

That is, for a mechanism $f(t_b; (x_r; t_r; t_u))g$, we obtain the buyer's continuation value as if the mechanism is $f(t_b - t_u; (x_r; t_r))g$ with no cancellation fee and then normalize her value function by deducting t_u .

Now we show there is a profitable deviation if the seller encourages post-transaction learning and provides the buyer with a surplus strictly larger than $V^0(\cdot; t_b; B)$. Note that when the mechanism is $f(t_b; (x_r; t_r; t_u))g$, the seller with reservation value u obtains the following return payoff by inducing the buyer to stop learning and request a return at belief \cdot ,

$$\begin{aligned} (1 - x_r)u + t_r + t_u &= (E(v_j | \cdot)) - u \frac{V_P^\theta(\cdot; t_b; t_u)}{v_h - v_l} - V_P(\cdot; t_b; t_u) + u \\ &= (E(v_j | \cdot)) - u \frac{V_P^\theta(\cdot; t_b - t_u)}{v_h - v_l} - V_P(\cdot; t_b - t_u) + t_u + u; \end{aligned} \quad (22)$$

where V_P^θ represents the partial derivative w.r.t to \cdot . From the (ODE) for post-transaction learning with a normalized price $t_b - t_u$, we obtain

$$V_P(\cdot; t_b - t_u) = v_h - (t_b - t_u) - \frac{k}{P} (1 - x_r) V_P^\theta(\cdot; t_b - t_u):$$

Thus, equation (22) can be further reduced to

$$(1 - x_r)u + t_r + t_u = \frac{v_h - u}{v_h - v_l} V_P^\theta(\cdot; t_b - t_u) - (v_h - t_b) - \frac{k}{P} + u$$

Note that if lowering V_P , V_P^θ is increasing by the above differential equation. Besides, the return payoff is increasing in V_P^θ . Hence, if $V_P(\cdot; t_b; t_u; P) = V_P(\cdot; t_b - t_u; P) - t_u > V^0(\cdot; t_b; B)$, the seller can gain larger expected revenue by raising the cancellation fee t_u to get a larger return payoff while let the buyer preserve the same stopping belief. Note that $V_P(\cdot; t_b - t_u; P) = V(\cdot; t_b - t_u; P)$, thus, the optimality holds when this inequality binds,

$$V(\cdot; t_b - t_u; P) - t_u = V(\cdot; t_b; B):$$

Next, it is very easy to verify that with the same selling price t_b , the seller's return payoff $t_r + t_u$ is always larger under post-transaction learning than inducing the

the buyer to stop at the same belief but under pre-transaction learning. Recall the (ODE),

$$(1 - \rho)V^0(\theta; t_b) + V(\theta; t_b) = (v_h - t_b) - \frac{k}{\rho}.$$

To induce the same stopping belief θ while restricting $V(\theta; t_b; t_u; \rho) = V(\theta; t_b; B)$, if $\rho < \rho_0$,

$$V^0(\theta; t_b; t_u; \rho) > V^0(\theta; t_b; B):$$

Since the return payoff is linear in V^0 , then the seller obtains larger return payoff by inducing the same stopping belief with post-transaction learning.

□

References

- Armstrong, Mark and Jidong Zhou, “Search Deterrence,” *The Review of Economic Studies*, 07 2015, 83 (1), 26–57.
- Bergemann, Dirk and Juuso Valimäki, “Learning and Strategic Pricing,” *Econometrica*, 1996, 64 (5), 1125–1149.
- and —, “Experimentation in Markets,” *The Review of Economic Studies*, 2000, 67 (2), 213–234.
- and Martin Pesendorfer, “Information structures in optimal auctions,” *Journal of Economic Theory*, 2007, 137 (1), 580–609.
- Bonatti, Alessandro, “Menu Pricing and Learning,” *American Economic Journal: Microeconomics*, August 2011, 3 (3), 124–63.
- Che, Yeon-Koo, “Customer Return Policies for Experience Goods,” *The Journal of Industrial Economics*, 1996, 44 (1), 17–24.
- Courty, Pascal and Hao Li, “Sequential Screening,” *The Review of Economic Studies*, 10 2000, 67 (4), 697–717.
- Daley, Brendan, Thomas Geelen, and Brett Green, “Due Diligence,” *Working paper*, 2021.
- Escobari, Diego and Paan Jindapon, “Price discrimination through refund contracts in airlines,” *International Journal of Industrial Organization*, 2014, 34, 1–8.
- Guo, Yingni, Hao Li, and Xianwen Shi, “Optimal Discriminatory Disclosure,” *Working paper*, 2020.
- Hinnosaar, Toomas and Keiichi Kawai, “Robust pricing with refunds,” *The RAND Journal of Economics*, 2020, 51 (4), 1014–1036.
- Inderst, Roman and Gilad Tirosh, “Refunds and returns in a vertically differentiated industry,” *International Journal of Industrial Organization*, 2015, 38, 44–51.

- and Marco Ottaviani, “Sales Talk, Cancellation Terms and the Role of Consumer Protection,” *The Review of Economic Studies*, 2013, 80 (3 (284)), 1002–1026.
- Johnson, Justin P. and David P. Myatt, “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, June 2006, 96 (3), 756–784.
- Keller, Godfrey, Sven Rady, and Martin Cripps, “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 2005, 73 (1), 39–68.
- Lang, Ruitian, “Try before you buy: A theory of dynamic information acquisition,” *Journal of Economic Theory*, 2019, 183, 1057–1093.
- Li, Hao and Xianwen Shi, “Discriminatory information disclosure,” *American Economic Review*, 2017, 107 (11), 3363–85.
- Matthews, Steven A. and Nicola Persico, “Information Acquisition and Refunds for Returns,” (*Unpublished Manuscript, University of Pennsylvania*), 2007.
- Mensch, Jeffrey, “Screening inattentive agents,” *Working paper*, 2020.
- Pease, Marilyn, “Shopping for Information: The Implications of Consumer Learning for Optimal Pricing and Product Design,” *working paper*, 2020.
- Pomatto, Luciano, Philipp Strack, and Omer Tamuz, “The cost of information,” *Working paper*, 2019.
- Ravid, Doron, Anne-Katrin Roesler, and Balázs Szentes, “Learning before trading: on the inefficiency of ignoring free information,” *working paper*, 2019.
- Roesler, Anne-Katrin and Balázs Szentes, “Buyer-optimal learning and monopoly pricing,” *American Economic Review*, 2017, 107 (7), 2072–80.
- Shi, Xianwen, “Optimal auctions with information acquisition,” *Games and Economic Behavior*, 2012, 74 (2), 666–686.

Shieh, Shiou, “Price and Money-Back Guarantees as Signals of Product Quality,”
Journal of Economics & Management Strategy, 1996, 5 (3), 361–377.

Wei, Dong and Brett Green, “(Reverse) Price Discrimination with Information Design,” *working paper*, 2020.