

Versioning with observational learning

Yating Yuan*

March 20, 2023

Abstract: Observational learning is a widespread phenomenon in various online marketplaces, e.g. Kickstarter, where sellers often launch a new product with different versions. Yet few papers explain the seller's versioning incentive from a social learning perspective. This paper explores how observational learning alone incentivizes monopolists to sell different versions. The dynamic learning process sheds new light on the multi-version policy beyond traditional explanations. Furthermore, the quality of consumers' private information is crucial in shaping two distinct selling strategies. In a market with noisy private signals, the seller provides a single cheap premium version; with precise signals, the seller prefers multiple versions.

Keywords: versioning, observational learning, nonlinear pricing

JEL classification: D42 D83 L11 L15

*Department of Economics, University of Warwick, Coventry, CV47AL, United Kingdom. Email: yating.yuan@warwick.ac.uk. I am grateful to my supervisors Motty Perry and Jacob Glazer for their guidance, support, and insightful comments. I would also like to thank Guo Bai, Ilan Kremer, Phil Reny, Costas Cavounidis, Daniele Condorelli, Ao Wang, Kirill Pogorelskiy, Francesco Squintani, and audiences at microeconomics work-in-progress workshops at Warwick for helpful comments and suggestions.

1 Introduction

Why do firms sell multiple versions when launching a new product? The economic literature has provided various explanations, but few of these give a full picture of what happens in markets that feature a dynamic observational learning process. Indeed, most existing theories are static in nature. They only apply to situations where consumers have ex-ante different preferences. To address this gap, this paper proposes a dynamic model, where consumers have a common value but arrive sequentially with different information, and offers a new justification for the multi-version policy.

By selling multiple versions at different prices, a monopolist can “design” the dynamics of the learning process, such as in which direction to herd and when. These learning dynamics, or more precisely, the evolution of posterior beliefs, will in turn bring variations in consumers’ willingness to pay and thus shape market demand. The interaction between learning and selling allows me to establish a key result: when each consumer receives noisy private information, a cheap premium version is optimal; when the private information is accurate enough, the seller prefers a multi-version policy.

Reward-based crowdfunding is a typical example of such a selling environment. It is a common practice for sellers to launch their crowdfunding campaigns with different price-quality packages¹. The multi-version policy is also widely adopted by firms that sell online service products, such as Grammarly, ChatGPT and D-ID studio. Their motivation is exactly to let the cheap version go viral. Once the market learns the value of their

¹For instance, a comic book writer may offer both a digital version and a paper-back version of their new book at different prices. A game designer might sell a wide range of packages, with more expensive ones including premium features to improve the gaming experience. https://www.kickstarter.com/projects/nogstudio/menyr?ref=discovery_category

products, it will hopefully be ready for a premium version. With a simple and tractable model, I will look into the monopoly pricing and versioning problems that arise in the new Internet era, and provide a fresh insight into the optimal selling strategy in various online markets.

The model is built on classical information cascade papers (Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992). Consider a monopolist seller who releases a new product with an unknown common value at the beginning. The value is high with probability $\frac{1}{2}$. The seller aims to maximize expected long-run average profits. She can offer a single version or multiple versions of the product with different observable (fixed) qualities and prices². An infinite number of consumers then arrive one at a time, each with a noisy private signal about the value. They also observe their predecessors' decisions and obtain public information from there. At the end of each period, they either choose a version to buy or abstain. Consumers will generally become increasingly informed over time until an informational cascade occurs.

The current article will first examine the optimal (fixed) pricing problem in a single-version benchmark model, focusing on the relationship between the optimal price and private signal quality. A discussion then follows on whether to offer multiple versions and at what prices, as the private signal quality varies.

In the single-version benchmark, choosing a higher price increases the margin (price effect) but reduces the probability of a buy cascade in the limit (quantity effect). The relative magnitude of both effects, and hence the price elasticity of demand, will be determined by the precision of the private signal. With a noisy signal, the quantity effect dominates the price effect. The seller prefers to stay safe by setting a low price that triggers a

²For the convenience of analysis, I assume the seller can offer up to two price-quality packages: a basic version and a premium version.

buy cascade³ from the beginning. A precise signal instead leads to a strong price effect, incentivizing the seller to choose a high price. In doing so, she bets on the $\frac{1}{2}$ prior probability that the product has a high value and reaps the fruit of learning when consumers become increasingly optimistic.

The multi-version model introduces a second trade-off beyond that in the benchmark case. Now the key question is whether to introduce a cheaper basic version.⁴ As usual, information rents come as a cost no matter how precise the private signal is. With two versions in the market, the seller has to lower the premium version's price to prevent high-belief consumers from purchasing the cheaper basic version.

However, the benefit of introducing the cheaper basic version emerges only when the private signal is precise enough. To understand the intuition behind this assertion, consider the case where consumers receive extremely precise signals. As discussed in the benchmark case, the seller's top priority now is to set a high price, as long as the consumers will not immediately run into a rejection cascade. Introducing a separate cheap basic version allows the first several consumers to learn at a relatively low price and makes room for a further increase of the premium version's price. In other words, introducing a basic version is profitable here because it relaxes the constraint of not starting with a rejection cascade in the seller's problem. As a result, Proposition 2 shows the multi-version policy will outperform the single-version policy if and only if the private signal is precise enough.

As the main results highlight the role of private signal quality, let me describe a real-world situation where consumers may receive private information of varying precision. The market for new treatments, as discussed in (Arieli, Koren, and Smorodinsky, 2022), consists of pharmaceutical compa-

³Technically speaking, the buy cascade starting from the beginning is not an informational cascade but just a sequence of buy decisions.

⁴To simplify the model, I assume the production cost of quality is zero. The seller will always keep a premium version in the market because it generates a higher total surplus.

nies as sellers and doctors as potential buyers. Doctors gather information by using limited free samples within their patient community. The realized success rate then gives each doctor a private signal about the value of the new treatment. The signal tends to be more precise if, for instance, doctors receive more free samples, the samples are similar in effectiveness to the majority of products, or their patient communities are more diversified.

The last part of the paper discusses the effects of versioning on market efficiency in information aggregation. The market generally works better in failing bad projects if we allow for multiple versions. Good projects, however, are less likely to succeed in a multi-version world.

The theoretical framework in this paper could be a starting point for future empirical analysis of the interaction between observational learning and the seller's decisions. In particular, it brings to our attention an important relationship between private information quality and selling strategy, which has been less studied in the literature.

A substantial number of empirical papers have already explored observational learning in various places, such as kidney exchange markets (Zhang, 2010), microloan markets (Zhang and Liu, 2012), music platforms (Newberry, 2016), and housing markets (Fan, Weng, Zhou, and Zhou, 2021). But few of them examine the effect of consumers' private information quality on the seller's pricing and versioning choices. According to Zhang and Liu (2012), investors in microloan markets do behave differently in the learning process when they find their predecessor's private information is more precise. Based on this observation, the theoretical model I have developed allows us to move forward to study the optimal selling strategies in such markets, especially when private information qualities vary across different products.

2 Related literature

An important early paper that examines the monopoly (fixed) pricing problem with observational learning is Welch (1992). In contrast to my model, he assumes a uniformly distributed state and a finite number of agents. Moreover, his signal structure is relatively noisy such that the posterior expected value updates slowly. The issuer (seller) will underprice to completely avoid a rejection cascade. My paper adds to the literature that when the private signal turns precise it is optimal to charge a high price and risk a rejection cascade. The seller makes an even higher profit by selling multiple versions.

Bose, Orosel, Ottaviani, and Vesterlund (2006, 2008) investigate a dynamic pricing problem in the informational cascade setting. Their setup bears similarities to my multi-version model in that offering multiple versions with different (average) prices at the outset seems to be a static alternative to dynamically adjusting the price of a single version. However, dynamic pricing gives the seller much more leeway to adjust their strategy accordingly as the learning process goes. In a sense, their seller deals with small gambles one by one while my seller faces a huge gamble up front. This paper discusses the optimal selling strategy in this more constrained environment.

Indeed, a comparison of our results in the case of a patient seller⁵ shows a marked difference in the equilibrium pricing and learning pattern. An infinitely patient seller always chooses a high separating price that reveals the private information regardless of signal quality in Bose et al. (2008). By contrast, my seller may offer a cheap basic version that prevents everyone from learning if the private signal is noisy enough.

Several other papers have also discussed the optimal selling strategy

⁵As my seller only cares about the long-run payoff, it suffices to compare the optimal behaviors of a patient seller in the two papers.

in the presence of social experimentation. Bonatti (2011) studies optimal dynamic menus to sell new experience goods to consumers with both a common value component and private taste. Laiho and Salmi (2021) considers a dynamic pricing problem when consumers can delay purchases. Bergemann and Välimäki (2002) investigates how dynamic competition between firms affects their pricing and entry decisions. This literature assumes sales generate information via experimentation and simplifies the learning process to a Brownian motion. Under this assumption, a lower price generates more information, whereas, in my model, a low enough price can eliminate information revelation. Hence, the seller faces a rather different trade-off.

Another closely related literature is menu pricing, beginning with the seminal works by Mussa and Rosen (1978) and Stokey (1979). Since then, many articles have attempted to discuss under what conditions a monopolist prefers multiple price-quality contracts over a single contract (Salant, 1989; Anderson and Dana, 2009). A more recent paper by Sandmann (2023) highlights the role of consumers' risk preferences in determining the optimality of a single-contract menu. As I mentioned above, the screening literature typically assumes consumers' willingness to pay is exogenous. My paper instead features endogenously formed variations in consumers' preferences due to the observational learning process. It allows me to unravel a novel observation on versioning and private signal quality.

3 Model setup: monopoly problem with information cascades

A monopolist seller plans to launch a new product that may have multiple versions. An infinite number of risk-neutral agents with unit demand arrive one at a time. $t \in \{1, 2, \dots, \infty\}$.

State of the world. The core technology of the product has a binary value $V \in \{0, 1\}$. Neither the seller nor the agents observe the realization of the value. They share the same prior probability $\mu_0 := \Pr(V = 1) = \frac{1}{2}$.

Actions. The seller offers two versions of the core product to each agent: a basic version $L = (p_L, q_L)$ with observable quality normalized to one, i.e., $q_L = 1$, and a premium version $H = (p_H, q_H)$ with $q_H = 2$.⁶ The seller announces a price schedule $p := (p_L, p_H)$ at the beginning which will be fixed over time. In each period t , an agent chooses a version to buy or rejects to buy any product. $a_t \in \{L, H, r\}, \forall t$.

Signal structure. Each agent, upon arrival, receives a binary private signal $s_t \in \{g, b\}$. It is symmetric and independent across individuals conditional on the state V . $P(s_t = g|V = 1) = P(s_t = b|V = 0) = \gamma \in (\frac{1}{2}, 1)$ represents the private signal quality.

Timing.

- $t = 0$: The state V is realized. The seller chooses a price schedule $p = (p_L, p_H)$ without observing the realized state.
- $t = 1, 2, \dots, \infty$: An agent arrives at the venue. She observes the prices, the decisions of previous agents, i.e., the public action history $\mathcal{H}_t := (a_1, a_2, \dots, a_{t-1})$, and a private signal s_t . Then she can either buy a version or walk away with nothing.

Payoff. The agent's utility depends on her action a_t , the price schedule p , and the value of the core technology V . $u(a_t, p, V) = Vq_i - p_i$ if the agent

⁶Note that the observable qualities q_H and q_L are different from the unobservable 'quality'/value of the core technology V . In a sense, the two versions are two vertically differentiated products that share the same core technology. If we take iPhone as an example, the basic version with $q_L = 1$ will be iPhone 14, whereas the premium version with $q_H = 2$ will be iPhone 14 Pro. They share the same core technology, e.g., a smooth operating system, but iPhone 14 Pro has some additional premium features, such as a larger screen and a Pro camera system. Appendix B discusses whether $q_H = 2$ is without loss.

buys one of the versions, $a_t = i \in \{L, H\}$. The utility from the outside option is normalized to zero, $u(r, p, V) = 0$.

A crucial assumption on the agent's utility is that the marginal utility of quality increases in the core technology's value V . The assumption is consistent with what we observe in many real-life situations. For example, people gain additional utility from the premium features of the iPhone 14 Pro only if they value the core technology of iPhone. Likewise, in the video game market, consumers would be interested in buying additional game features and packages only when they value the core plots, graphics, gameplay and characters in the original game.

The seller maximizes her long-run average profits

$$\pi(a, p) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N [\mathbf{1}(a_t = L)p_L + \mathbf{1}(a_t = H)p_H]$$

where $a := (a_1, a_2, \dots)$ denotes the agents' action profile.

The following analysis will focus on pure-strategy perfect Bayesian equilibrium. For technical reasons I assume agents always buy the product or choose a better version whenever they are indifferent.

Discussion. With the zero-one binary value assumption, my results only pertain to situations where the value difference is large enough and/or the outside option r gives a high enough payoff (see Appendix A). In addition, the production cost is assumed to be zero for simplicity. Adding a non-zero constant marginal cost will not have a substantial impact on the main results. While the seller only designs the prices explicitly, the optimal pricing strategy also reflects their versioning choice because the seller can always abandon one of the versions by charging an average price $\frac{p_i}{q_i}$, $i \in \{L, H\}$, higher than one.

4 Optimal pricing with a single version

This section studies the optimal pricing problem when the monopolist sells a single version. The analysis reveals a unique connection between private signal quality, belief updating and the price/quantity effect. As a result of the connection, the seller's pricing incentive changes substantially when the private signal becomes precise enough. This change in the seller's objective lays a foundation for discussion over the optimal versioning choice in the next section.

4.1 Belief space

The evolution of posterior beliefs determines the agents' optimal purchase decisions over time and thus the seller's optimal pricing policy. To facilitate the following discussion, let me first characterize the belief space and introduce an important state variable k , the difference between the numbers of inferred good signals and inferred bad signals in the public action history.

Let μ_t be the public posterior belief at time t after history \mathcal{H}_t , that is, $\mu_t := \Pr(V = 1|\mathcal{H}_t)$. Agent t 's private belief is defined as $\theta_t := \Pr(V = 1|\mathcal{H}_t, s_t)$ because each agent also receives a private signal s_t .

As a result of a simple signal structure, the posterior belief depends only on k , the difference between the numbers of good and bad signals that we can infer from the previous action history \mathcal{H}_t . More specifically, the posterior belief takes values from a discrete set $\{V_k(\gamma)|k \in \mathbb{Z}\}$, where $V_k(\gamma) := \Pr(V = 1|k, \gamma) = \frac{\gamma^k}{\gamma^k + (1-\gamma)^k}$.⁷ For instance, at $k = 0$, the posterior belief equals the prior belief $\mu_0 = \frac{1}{2}$. Note also that k is not the difference between the purchase and rejection numbers. Once we enter into an informational cascade, the purchase decisions carry no information about the private signals. The variable k will stay the same forever but the difference

⁷To ease the notation I will omit γ and write $V_k(\gamma)$ as V_k most of the time.

between purchase and rejection numbers may still increase.

$V_k(\gamma)$ has two important properties. First, it strictly increases in k . Second, $V_k(\gamma) - V_{k-1}(\gamma)$ increases in γ for all $k \in \mathbb{Z}$. Intuitively, belief updating speeds up as the signal becomes increasingly precise. Each update leads to a large jump in the posterior belief.

4.2 The agent's optimization problem and learning patterns

It is easy to characterize the agent's optimal choice as a cutoff strategy. We will then analyze the evolution patterns of the public posterior belief μ_t under different pricing strategies.

For any public action history \mathcal{H}_t and private signal s_t , the expected utility of agent t is⁸ $\mathbb{E}(u(a_t, p, V) | \mathcal{H}_t, s_t) = \theta_t - p$ if $a_t = \textit{buy}$ and zero otherwise. The agent then adopts a cutoff strategy: they will buy the product if and only if their private posterior belief over the core technology's value V exceeds the price p .

$$a_t^*(p, \mathcal{H}_t, s_t) = \begin{cases} r, & 0 \leq \theta_t < p \\ \textit{buy}, & p \leq \theta_t \leq 1 \end{cases}$$

In this observational learning environment, $\{k_t\}_{t=1,2,\dots}$ follows an asymmetric random walk process with upward transition probability γ conditional on state $V = 1$ and transition probability $1 - \gamma$ conditional on state $V = 0$.

To characterize the learning dynamics in the state space of k , let me define $k^*(p) := \min\{k \in \mathbb{Z} | V_k \geq p\}$. Then V_{k^*} refers to the threshold posterior belief above which an agent prefers to buy. Now the belief space

⁸Since the seller offers a single version, the price p is a scalar. The product's observable quality is normalized to one, $q = 1$.




Figure 1: Short learning pattern with a single version.

Notes: ‘rejection’ refers to the rejection cascade set and ‘buy’ the buy cascade set. ‘Learn’ denotes the learning set and V_{k^*} refers to the threshold belief above which an agent prefers to buy the product.

can be divided into three sets (as illustrated in Figure 1):

1. rejection cascade set: $\{V_k | k < k^* - 1\}$. If the public belief falls into this set, every subsequent agent will rationally ignore their own private signal and reject.
2. learning set: $\{V_{k^*-1}, V_{k^*}\}$. If the public belief drifts into this set, an agent will buy the product after a good signal and reject after a bad one.
3. buy cascade set: $\{V_k | k > k^*\}$. Here agents will buy the product regardless of their private signals.

Below I provide a detailed example of the partition at price $p = V_1$, or equivalently, $k^* = 1$, and the resulting dynamic learning process.

Example: learning pattern at price $p = V_1$.

The learning process starts from the prior belief $k_0 = 0 = k^* - 1$, as in Figure 2. If the first agent receives a good signal, her private belief will change to $\theta_1 = V_1$ and thus she will buy the product at the price $p = V_1$. Otherwise, with a bad signal and private belief at $\theta_1 = V_{-1}$, she will reject. The first agent’s action fully reveals her private signal. Conditional on the state of the world $V = 1$ ($V = 0$), the public belief will jump to $\mu_1 = V_1$ with probability γ ($1 - \gamma$) and $\mu_1 = V_{-1}$ with probability $1 - \gamma$ (γ) at the end of period $t = 1$.

The process continues until the public belief reaches a cascade set, i.e. $\{V_k | k < 0\}$ or $\{V_k | k > 1\}$. For instance, if the public belief hits V_{-1} at




Figure 2: The learning dynamics at price $p = V_1$.

Notes: given the price $p = V_1$ and the state of the world V , the difference between the inferred good and bad private signals, k_t , follows an asymmetric random walk process with two absorbing points $\{1, 2\}$. They represent the boundaries of the rejection and buy cascade sets. The vertical axis depicts the state space of k_t with the associated posterior belief on the left. The horizontal axis refers to time t . Since the first agent arrives in the learning set, her action fully reveals the private signal. The public belief jumps up to V_1 after a buy decision and down to V_{-1} after a rejection.

the end of period $t = 5$ as in Figure 2, the next agent will reject regardless of her private signal. All the subsequent agents face the same situation, resulting in a rejection cascade.

To summarize, the learning set consists of two non-absorbing beliefs $\{V_0, V_1\}$. $\{V_{-1}, V_2\}$ are instead two absorbing beliefs. Conditional on the state of the world, k_t follows an asymmetric random walk process. The associated public belief will bounce within the learning set until it finally hits one of the two absorbing beliefs.

4.3 The seller's optimal pricing strategy

Consider a case with finite time $t \in \{1, 2, \dots, N\}$ and define the finite time expected average profit function as

$$\mathbb{E}\left(\frac{1}{N} \sum_{t=1}^N \mathbf{1}(a_t = \text{buy})p\right) = \frac{\mathbb{E}(\sum_{t=1}^N \mathbf{1}(a_t = \text{buy}))}{N} p$$

Notice that the expression $\sum_{t=1}^N \mathbf{1}(a_t = \text{buy})$ describes how many agents have bought the product. Denote this purchase number by D^N and the ex-ante probability of having a buy cascade given price p by $\lambda(p)$.

Claim 1. *As the total number of agents N goes to infinity, the expected proportion of agents who will buy the product, $\frac{\mathbb{E}(D^N)}{N}$, approaches the ex-ante probability of the public belief reaching the buy cascade set, i.e. $\lambda(p)$.*

The proof of Claim 1, along with all the other proofs, is included in the Appendix. Intuitively, what happens in finite time does not matter because the seller receives the long-run average profit and a cascade takes place with probability one. The seller only cares about whether a buy cascade occurs in the long run. Once it occurs, she will receive a payment p from every subsequent agent. With Claim 1, the expected average profit is $\lambda(p)p$.

The next observation is that the optimal price can only take values from the set $\{V_{-1}, V_0, V_1\}$. First, it is never optimal to charge a price higher than V_1 that immediately induces a rejection cascade. No one will buy the product at such a high price and the seller gets zero profit. But the seller can always achieve a strictly positive expected profit by lowering the price because the production cost is zero. Second, any price strictly lower than V_{-1} is sub-optimal. At a price $p = V_{-1}$, we already trigger a buy cascade right from the beginning. There is no point in further decreasing the price. Finally, given the tie-breaking rule, the seller will not choose any prices between the belief points⁹.

Proposition 1. *There exists a pair of threshold signal qualities $\{\underline{\gamma}, \bar{\gamma}\}$, $\frac{1}{2} < \underline{\gamma} < \bar{\gamma} < 1$.*

1. *With a noisy signal, $\gamma \in (\frac{1}{2}, \underline{\gamma})$, it is optimal to offer a cheap price such that everyone buys the product $p^*(\gamma) = V_{-1}$;*


⁹For instance, any price $p \in (V_{-1}, V_0]$ leads to the same purchase behaviors of the agents. Charging a price $p = V_0$ dominates any other prices within the interval $(V_{-1}, V_0]$.

2. *As the private signal becomes increasingly precise, the optimal price becomes higher and it takes harder and longer to arrive at a buy cascade. $p^*(\gamma) = V_0$ for $\gamma \in (\underline{\gamma}, \bar{\gamma}]$; $p^*(\gamma) = V_1$ for $\gamma \in (\bar{\gamma}, 1]$.*


Similar to the standard monopoly pricing problem, the seller's expected average profit consists of two parts: the ex-ante probability of having a buy cascade $\lambda(p)$ and the price p , taking values from $\{V_{-1}, V_0, V_1\}$. A price increase leads to a lower buy cascade probability ('quantity' effect) and a higher margin (price effect). To understand the result, we can explore two extreme cases: when the private signal is very noisy $\gamma \in (\frac{1}{2}, \underline{\gamma})$ and when the private signal is very precise $\gamma \in (\bar{\gamma}, 1)$.

With a noisy signal, $\gamma \in (\frac{1}{2}, \underline{\gamma})$, is it profitable to increase the price from $p^* = V_{-1}$ to V_0 ? Now the posterior belief updates slowly. Agents are still quite uncertain about state V even after receiving a good signal. Consequently, the increase in margin ($V_0 - V_{-1}$) is fairly small, which implies a weak price effect (Figure 3a). By contrast, the quantity effect is relatively strong: the price increase causes a discrete drop in the buy cascade probability. At a price $p = V_{-1}$, a buy cascade occurs from the start. By lifting the price to V_0 , the seller risks a rejection cascade. As we can see from Figure 3b, the buy cascade probability drops from one to some value below 0.65. Hence, the cost of the price increase far outweighs its benefit. It is optimal to stay with the low price that triggers a buy cascade immediately.

On the other hand, the price effect dominates the quantity effect when the private signal becomes extremely precise, $\gamma \in (\bar{\gamma}, 1]$. Imagine the seller considers a price increase from $p = V_0$ to $p = V_1$. The posterior belief updates very fast, leading to a large increase in margin ($V_1 - V_0$). In contrast, the quantity effect weakens at a large γ . With a precise signal, the learning process quickly enters into the correct cascade as long as we start from within the learning set. Both buy cascade probabilities at price $p = V_0$ and $p = V_1$ will be fairly close to the prior belief $\frac{1}{2}$ (Figure 3b).



(a) Three candidate prices.



(b) The buy cascade probabilities.

Figure 3: The prices and buy cascade probabilities as signal quality varies.

Choosing the higher price $p = V_1$ only results in a small drop in the buy cascade probability λ . As a result, the seller finds it optimal to charge the highest possible price that does not trigger a rejection cascade immediately.

U-shaped profit. Figure 4 plots the seller’s average profit as a function of the signal quality γ under different pricing strategies. The maximal profit function is U-shaped: the seller is better off when consumers have nearly no information or full information.

5 When to offer multiple versions

Offering both a premium version and a cheaper basic version forces the seller to concede information rents to high-belief consumers. Nevertheless, a separate basic version may help relax a binding constraint in the seller’s problem, leading to a higher profit. This trade-off is the key to understanding the mechanism through which signal quality affects the seller’s versioning policy.

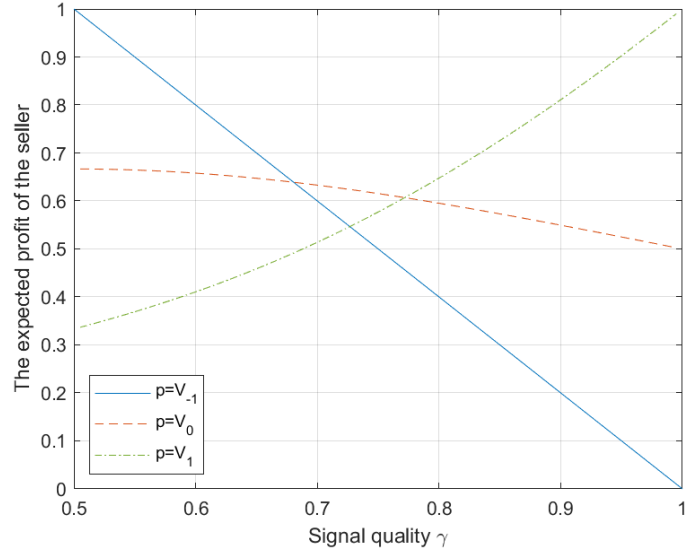


Figure 4: The seller's expected profit in the single-version model.

5.1 An extended learning set for agents

Selling two versions with different observable qualities enlarges the choice set of the agents. It is possible now to induce a longer learning set that incorporates three or more belief points. The extended learning set allows the seller to offer a more expensive premium version without triggering a rejection cascade immediately.

Given any prices $p = (p_L, p_H)$, the expected payoff of agent t becomes

$$\mathbb{E}(u(a_t, p, V) | \mathcal{H}_t, s_t) = \begin{cases} \theta_t - p_L, & a_t = L \\ 2\theta_t - p_H, & a_t = H \\ 0, & a_t = r \end{cases}$$

If the average price of the basic version is higher than that of the premium version, $\frac{p_H}{q_H} \leq p_L$, the basic version becomes so expensive that no one would ever consider it. Agents behave exactly the same as in the single-version model. If the average price of the basic version is lower than that of the premium version, $p_L < \frac{p_H}{q_H} \leq 1$, agents will possibly choose each of the

three actions during the process.

$$a_t^*(\mathcal{H}_t, s_t, p) = \begin{cases} r, & 0 \leq \theta_t < p_L \\ L, & p_L \leq \theta_t \leq p_H - p_L \\ H, & p_H - p_L \leq \theta_t \leq 1 \end{cases}$$

The additional version choice leads to a richer observable action set. It allows agent t to convey more private information to the subsequent agents. Thus, we have a longer learning set where the public belief may exhibit more variations over time before it arrives at a cascade.

If $p_L \leq 1 < \frac{p_H}{q_H}$, the agents only consider the basic version. As the marginal production cost is zero, the total surplus from selling the premium version is greater than that from the basic version at any posterior belief θ_t . Giving up the premium version is, therefore, sub-optimal for the seller.

From now on, the discussion will center around which type of learning patterns can occur when the seller introduces a cheaper basic version $p_L < \frac{p_H}{q_H} \leq 1$. Let me define

$$k_L(p_L) := \min\{k \in \mathbb{Z} | V_k \geq p_L\},$$

$$k_H(p_H, p_L) := \min\{k \in \mathbb{Z} | V_k \geq p_H - p_L\}.$$

Then V_{k_L} represents the threshold posterior belief above which an agent prefers the basic version over rejection, and V_{k_H} the threshold belief above which an agent prefers the premium version over the basic version. Analogous to the benchmark model, we can divide the belief space into (at most) five sections.

1. rejection cascade set: $\{V_k | k < k_L - 1\}$.
2. learning set for basic version: $\{V_{k_L-1}, V_{k_L}\}$.

3. buy cascade set for basic version: $\{V_k | k_L < k < k_H - 1\}$.
4. learning set for premium version: $\{V_{k_H-1}, V_{k_H}\}$. An agent who arrives with a public belief in this set will buy the premium version after a good signal and the basic version otherwise.
5. buy cascade set for premium version: $\{V_k | k > k_H\}$.

The learning patterns vary with the distance between the two threshold beliefs (Figure 5). If the two thresholds are far apart ($k_L < k_H - 2$), we observe all of the five sets in the belief space (Figure 5a). In between the two learning sets, there is a basic cascade set with absorbing beliefs. If the two thresholds are close just enough ($k_L = k_H - 2$, Figure 5b), we will have two consecutive learning sets, which together create a big learning set. Figure 6 provides an example of the dynamic learning path in this case with $k_L = 0, k_H = 2$, or equivalently $p_L = V_0, p_H = V_0 + V_2$. As the two thresholds get even closer ($k_H - 2 < k_L < k_H$, Figure 5c-5d), the two learning sets may partly or completely overlap.

5.2 The optimal selling strategy

This part explores when and why the monopolist finds it optimal to offer two separate versions rather than a single version. As the production cost is zero, she always keeps a premium version in the market. To sell both versions, she must offer the basic version at a cheaper price per quality. Otherwise, the agents will never choose the basic version on the equilibrium path. The versioning problem then boils down to whether to introduce a cheaper basic version.

In the limit, the proportion of agents who have purchased a certain version approaches the probability of a buy cascade for that version. We simplify the expected profit to¹⁰ $\Pr(\text{premium cascade}|p)p_H + \Pr(\text{basic cascade}|p)p_L$.

¹⁰Details in Appendix F.




Figure 5: Characterizing learning patterns in belief space $\{V_k\}_{k \in \mathbb{Z}}$.
Notes: ‘rejection’ refers to the rejection cascade set, ‘basic’ the buy cascade set for the basic version, and ‘premium’ the buy cascade set for the premium version. ‘Learn_L’ represents the learning set for the basic version: when the public belief arrives at this set, the agent will choose between the basic version and rejection according to her private signal. $V_{k_L}(p_L)$ is the threshold belief above which an agent prefers the basic version over rejection. Similarly, ‘Learn_H’ denotes the learning set for the premium version, and $V_{k_L}(p_H, p_L)$ refers to the threshold belief above which an agent prefers the premium version over the basic version.

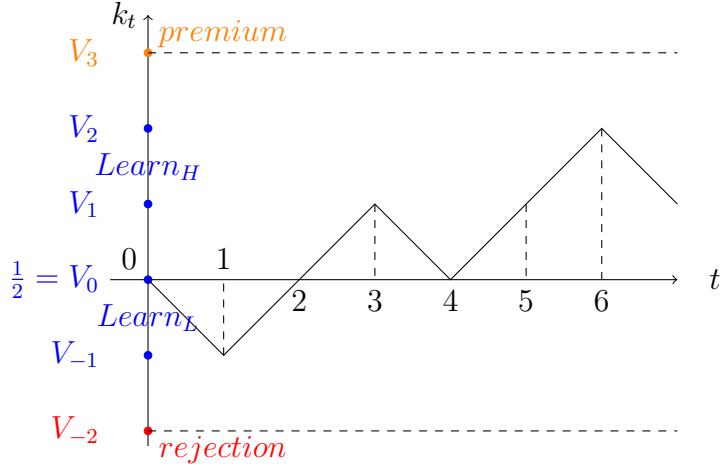


Figure 6: Dynamics of long learning pattern, $p_L = V_0, p_H = V_0 + V_2$.
Notes: given the prices $p_L = V_0, p_H = V_0 + V_2$ and the state V , the difference between the inferred good and bad private signals, k_t , follows an asymmetric random walk process with two absorbing states $\{-2, 3\}$. The absorbing states represent the boundaries of the rejection and premium cascade sets. The vertical axis depicts the state space of k_t with the associated posterior belief on the left. The horizontal axis refers to time t . There is an extended learning phase for agents where they can possibly infer more information from different version choices.

Additionally, as we study in the previous section, the pricing choice will be equivalent to the choice of threshold beliefs k_L and k_H , which in turn determine the learning patterns and, within each learning pattern, the point at which to start the process.

As a first step to solve for the optimal prices, we can rule out several sub-optimal learning patterns and starting points. For each remaining combination, I derive the resulting expected profits as polynomial functions of signal quality γ . Comparing the polynomial functions gives us the optimal selling strategy for each signal quality.

Proposition 2 (Versioning). *There exists a threshold signal quality $\gamma_v \in (0, 1)$ such that it is optimal to offer two versions if and only if $\gamma > \gamma_v$.*

Proposition 2 tells us with a noisy private signal, the monopolist will offer an expensive basic version that nobody will ever buy, leading to a de facto single-version market. Otherwise, a multi-version policy that features

both an expensive premium version and a cheap basic version is optimal. Here the key question is: Why does the seller introduce a cheap basic version only when the private signal is precise?

On the cost side, a cheaper basic version makes the seller concede information rents to buyers of the premium version. If the price p_H leaves no surplus to them, the premium version buyers will deviate and obtain a positive payoff by purchasing the basic version instead. $\theta_t - \frac{p_L}{q_L} > \theta_t - \frac{p_H}{q_H} = 0$.

While the cost of introducing a cheaper basic version always exists, its benefit occurs only when the private signal is precise enough. Let's discuss the precise signal case first. A precise signal brings a strong price effect and a weak quantity effect. The seller tends to choose the highest possible price subject to a binding constraint: that the price does not induce a rejection cascade immediately. Introducing a cheaper basic version can extend the learning set towards lower beliefs (see Figure 5b), which relaxes this no-immediate-rejection constraint. The seller can thus charge an even higher price for the premium version.

In contrast, with a noisy signal, the seller prefers a safe choice. If the seller were selling a single premium version, the optimal price would be cheap enough that the process starts fairly close to a premium cascade. At such a low price, the no-immediate-rejection constraint does not bind, and thus it is not profitable to introduce a cheaper basic version. Intuitively, there's no point in providing a cheaper version if the seller already offers a premium version at affordable prices.

Learning dynamics. With a noisy signal ($\gamma < \gamma_v$), we have a single-version market, and the process features a short learning pattern (Figure 1). As in the benchmark, it is optimal to charge an all-buy price $\frac{p_H}{q_H} = V_{-1}$ if the private signal is extremely noisy.

Proposition 3 (Learning dynamics). *In the precise signal case where a*

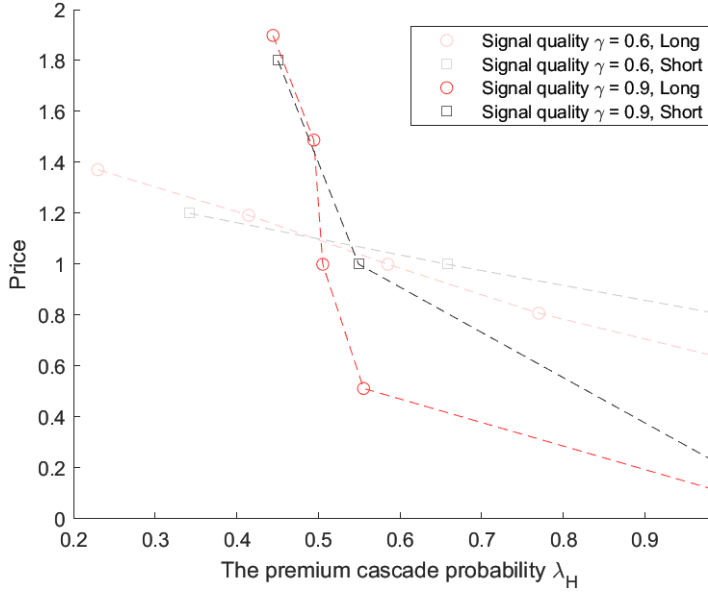


Figure 7: Longer learning pattern leading to lower price elasticity.

Notes: each black square/red circle represents a pair of price and the premium cascade probability (p_H, λ_H) that can be attained by starting at different positions in the short learning pattern (1)/long learning pattern (5b), given a certain signal quality γ . Lines with stronger colors imply the private signal is more precise.

multi-version policy is optimal, $\gamma > \gamma_v$, the process features a long learning pattern (5b) with two consecutive learning sets. The optimal prices of the two versions are just close enough to preclude a basic cascade between the two learning sets.

The whole point of adding a cheaper basic version is to relax the no-immediate-rejection constraint and further increase the price of the premium version. The longer the combined learning set is, the more expensive the premium version can be. Also, having a basic cascade in between is sub-optimal because that only breaks down a long learning set into two short ones. Hence, it is optimal to choose pattern 5b. As the tie-breaking rule favors a high-quality version, the two threshold beliefs are two signals away from each other (see Table 1).

Figure 7 depicts long-run demand curves that demonstrate the change in the seller's objective and why versioning helps only in the precise signal




Figure 8: The seller’s expected profits in the two-version model.

case. Strong colors correspond to precise signals. Compared to the demand at a noisy signal $\gamma = 0.6$, the demand curves at a precise signal $\gamma = 0.9$ are less elastic. Hence, a more precise signal predicts an optimal price-probability pair near the top left corner. As we can see, the top left is exactly where the long learning pattern (red line) dominates the short one (black line), which creates the versioning incentive.

Other testable implications. Table 1 provides further details of the optimal pricing strategy and characteristics of the equilibrium learning process. As the private signal becomes more precise, the process starts from a position closer to the rejection cascade set. The probability of a premium cascade, however, is non-monotone in the signal quality.

In addition, whenever the seller offers two versions, she adopts a convex pricing strategy. She charges a higher price for the second ‘unit’ of quality ($p_H - p_L > p_L$) because she expects the premium version buyers to have a higher posterior belief and, thus, a higher marginal utility of quality.

Figure 8 plots the seller’s expected profit as a function of signal quality.

Again, the optimal expected profit is U-shaped in signal quality γ .

| γ | Optimal pricing | Steps to <i>rej</i> | Steps to <i>premium</i> | $P(\text{premium})$ |
|---------------------------|----------------------------------|---------------------|-------------------------|---------------------|
| $(\frac{1}{2}, \gamma_s)$ | $p_H^* = 2V_{-1}$ | 3 | 0 | 1 |
| (γ_s, γ_v) | $p_H^* = 2V_0$ | 2 | 1 | [0.626, 0.638] |
| (γ_v, γ_m) | $p_L^* = V_0, p_H^* = V_0 + V_2$ | 2 | 3 | [0.445, 0.479] |
| $(\gamma_m, 1)$ | $p_L^* = V_1, p_H^* = V_1 + V_3$ | 1 | 4 | [0.388, 0.5] |

Table 1: The properties of different equilibria as signal quality γ increases. *Notes:* The thresholds $\gamma_s < \gamma_v < \gamma_m$ are defined in the proof F.2.2. Steps to *rej/premium* refers to the distance between the initial position of $\{k_t\}$ and the rejection/premium cascade set in belief space. ‘Steps to *rej* = 3’ means we are three consecutive bad signals away from the rejection cascade. ‘Steps to *premium* = 3’ means we are three consecutive good signals away from the premium cascade. $P(\text{premium})$ is the probability of reaching a premium cascade set. The premium cascade probability usually jumps at the thresholds. So it is not monotone in γ . But the general trend of $P(\text{premium})$ decreases until around the prior belief 0.5.

6 Discussion: market efficiency in aggregating information

Does the opportunity to offer multiple versions give a boost to market efficiency in information aggregation? This section introduces a measure to compare the learning efficiency outcomes in the single-version and multi-version model: the conditional premium cascade probability $\Pr(\text{premium cascade}|V)$. Information aggregates more efficiently if the probability of having a premium cascade is higher in the good state $V = 1$ and lower in the bad state $V = 0$.

It turns out the multi-version option improves information aggregation efficiency in the bad state but not always in the good state (Figure 9). Indeed, the seller offers multiple versions only when she aims at a high margin and starts the process far away from the buy cascade sets. In the single-version model, the optimal pricing strategy induces a learning process starting, at most, two consecutive good signals away from the premium cascade. In the multi-version model, however, the learning process can pos-

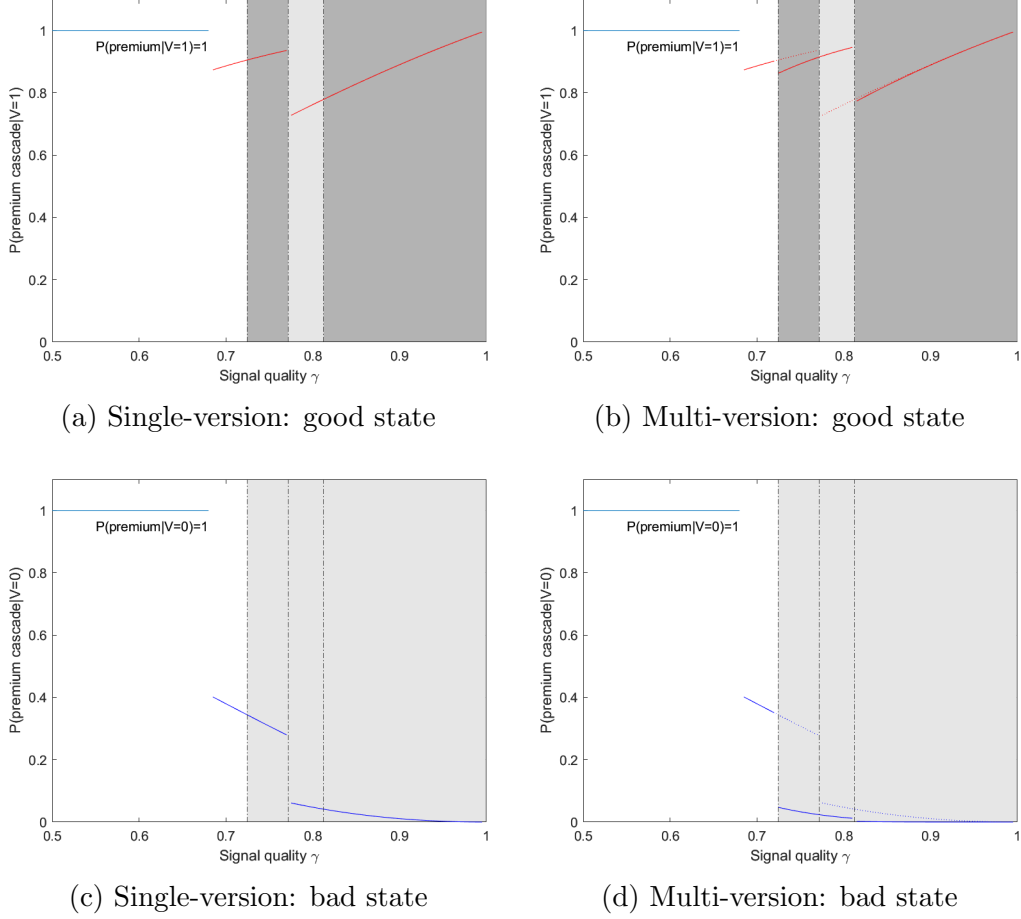


Figure 9: The premium cascade probability conditional on the state.

Notes: (1) The red solid curves in Figure (9a) and (9b) plot the conditional probability of the premium cascade given the core product is good ($V = 1$). The blue solid curves in Figure (9c) and (9d) plot the conditional probability of the premium cascade given the core product is bad ($V = 0$). In the multi-version figures (9b) and (9d) I also add the conditional premium cascade probability in the single version model as dotted curves in order to compare the information aggregation efficiency in different models. (2) The outcomes in the single and multi-version models differ only in the shaded areas. Light grey areas represent the signal quality interval where the multi-version option leads to an efficiency improvement. Dark grey areas represent the signal quality interval where the multi-version option results in an efficiency loss. (3) The three vertical lines denote thresholds γ_2 (where the seller switches to a multi-version policy with prices $p_L = V_0$ and $p_H = V_0 + V_2$ in the multi-version model), $\bar{\gamma}$ (where the seller starts to set the highest price $p = V_1$ in the single version model) and γ_3 (where the seller adopts the most expensive pricing scheme in the multi-version model, $p_L = V_1$ and $p_H = V_1 + V_3$) accordingly.

sibly start three to four consecutive good signals away from the premium cascade. Thus, the multi-version option generally reduces the premium cascade probability in both states.

Figure 9c and 9d show, conditional on the bad state, the premium cascade probability is smaller in the multi-version model. Hence, the multi-version option improves the market efficiency in ruling out bad projects. Meanwhile, according to Figure 9a and 9b, the premium cascade probability conditioning on the good state also decreases most of the time. The market is generally less capable of selecting good projects in the multi-version model. The only exception is the light grey area in the middle.¹¹

7 Concluding remarks

This paper investigates a monopolist's optimal (fixed) pricing and versioning policy in a market with observational learning. A unique insight of the paper is that private signal quality can affect the price elasticity of long-run demand. As a result, the optimal selling strategy becomes qualitatively different as the signal quality changes. In a market with noisy private information, the seller will offer a single cheap premium version. When the agents arrive with precise private information, it is optimal to launch two different versions: a basic version and a more expensive premium version. In this case, versioning is profitable because it extends the learning set and relaxes a key constraint of the seller's problem.

In addition, the model predicts a U-shaped optimal profit function in both the single-version and multi-version case. The seller is better off when the agents arrive with either almost zero or full information. A medium

¹¹In this area, the single-version seller already switches to an extreme (expensive) pricing scheme. It is so risky that just a single bad signal can trigger a rejection cascade. The multi-version seller offers a relatively cheap basic version and only with two consecutive bad signals can we trigger a rejection cascade.

degree of informativeness makes her worse off. A related paper by Arieli et al. (2021) analyzes how to optimally design the consumers' information structure in an observational learning setup. (SgROI, 2002; Gill and SgROI, 2008, 2012) instead examines optimal pre-launch information disclosure in this environment. It will be worth exploring how information provision affects the optimal versioning policy in the future.

References

- Anderson, E. T. and J. D. Dana (2009). When is price discrimination profitable? *Management Science* 55(6), 980–989.
- Arieli, I., R. Gradwohl, and R. Smorodinsky (2021). Herd design. *Working Paper*.
- Arieli, I., M. Koren, and R. Smorodinsky (2022). The implications of pricing on social learning. *Theoretical Economics* 17(4), 1761–1802.
- Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics* 107(3), 797–817.
- Bergemann, D. and J. Välimäki (2002). Entry and vertical differentiation. *Journal of Economic Theory* 106(1), 91–125.
- Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100(5), 992–1026.
- Bonatti, A. (2011). Menu pricing and learning. *American Economic Journal: Microeconomics* 3(3), 124–63.
- Bose, S., G. Orosel, M. Ottaviani, and L. Vesterlund (2006). Dynamic

- monopoly pricing and herding. *The RAND Journal of Economics* 37(4), 910–928.
- Bose, S., G. Orosel, M. Ottaviani, and L. Vesterlund (2008). Monopoly pricing in the binary herding model. *Economic Theory* 37(2), 203–241.
- Fan, Z., X. Weng, L.-A. Zhou, and Y. Zhou (2021). Observational learning and information disclosure in search markets. *Working Paper*.
- Gill, D. and D. SgROI (2008). Sequential decisions with tests. *Games and Economic Behavior* 63(2), 663–678.
- Gill, D. and D. SgROI (2012). The optimal choice of pre-launch reviewer. *Journal of Economic Theory* 147(3), 1247–1260.
- Laiho, T. and J. Salmi (2021). Coasian dynamics and endogenous learning. *Working Paper*.
- Mussa, M. and S. Rosen (1978). Monopoly and product quality. *Journal of Economic Theory* 18(2), 301–317.
- Newberry, P. W. (2016). An empirical study of observational learning. *The RAND Journal of Economics* 47(2), 394–432.
- Salant, S. W. (1989). When is inducing self-selection suboptimal for a monopolist? *The Quarterly Journal of Economics* 104(2), 391–397.
- Sandmann, C. (2023). When are single-contract menus profit maximizing. *Working Paper*.
- SgROI, D. (2002). Optimizing information in the herd: Guinea pigs, profits, and welfare. *Games and Economic Behavior* 39(1), 137–166.
- Stokey, N. L. (1979). Intertemporal price discrimination. *The Quarterly Journal of Economics* 93(3), 355–371.

Welch, I. (1992). Sequential sales, learning, and cascades. *Journal of Finance* 47(2), 695–732.

Zhang, J. (2010). The sound of silence: Observational learning in the us kidney market. *Marketing Science* 29(2), 315–335.

Zhang, J. and P. Liu (2012). Rational herding in microloan markets. *Management Science* 58(5), 892–912.

A Value difference and outside option

As a main result of the paper, the optimal price and the optimal number of versions increases in signal precision. This result may fail when 1) the value difference is small enough compared to the low value or 2) consumers receive a low payoff from not buying.

To see the intuition, consider an extreme case where each agent receives a nearly perfect signal, i.e., $\gamma \rightarrow 1$. Now the buy cascade probabilities are fairly close to the prior probability $\frac{1}{2}$ as long as we start within the learning set. Suppose the binary value is either 11 or 10. Pricing at the low value $p = 10$ strictly dominates any prices that do not trigger a buy cascade immediately. In the former case, everyone will buy and the expected profit is 10 while in the latter case, the expected profit is smaller than $\frac{1}{2} * 11 < 10$. Hence, the seller will choose a low all-buy price even if the private signal is precise. Furthermore, she has no incentive to offer two versions because the constraint of not starting with a rejection cascade never binds.

More generally, let the binary value be V_h with prior probability μ_0 and V_l with probability $(1 - \mu_0)$ ($V_h > V_l > 0$). Each agent receives a payoff R from not buying. The seller chooses a price p for a single product with exogenous observable quality q . In this setting, a consumer with private belief $\theta \in [0, 1]$ will buy the product if and only if $\theta \geq \frac{R+p-qV_l}{(V_h-V_l)q}$. At price $p = qV_l - R$, everyone buys the product. Choosing such a low price guarantees the seller an expected profit of $\underline{\pi} := qV_l - R$. Meanwhile, an optimal price must be smaller than $p = qV_h - R$ because at this price only agents with private belief $\theta = 1$ will buy. It follows that the highest possible expected profit typically will not exceed $\bar{\pi} := \mu_0(qV_h - R)$ when the private signal is almost perfect. Hence, as $\gamma \rightarrow 1$, the seller prefers the all-buy price if $\bar{\pi} < \underline{\pi} \Leftrightarrow q(\mu_0(V_h - V_l) - (1 - \mu_0)V_l) + (1 - \mu_0)R < 0$.

B Optimal quality choice

For analytical convenience, the premium quality is assumed to be 2. In fact, the seller always prefers to maximize the quality difference ($q_H - q_L$) within reasonable limits, given the quasi-linear payoff structure of the agents and zero marginal cost assumption.

More precisely, suppose the seller can choose two qualities as she likes, that is, $q_L, q_H \in [\underline{q}, \bar{q}] \subset (0, \infty)$ with $q_L < q_H$. The learning process has two possible structures. In the first case, the process ends either with a rejection cascade or with a basic cascade, i.e. a buy cascade for the basic version, with probability one. Since $q_L < q_H$ and agents only care about the average price, it is always more profitable for the seller to sell a single premium version with $q_H = \bar{q}$. In the second case, the learning process ends either with a rejection cascade or with a premium cascade with probability one. Now increasing the quality difference reduces the information rents conceded to the premium version buyers. Hence, the seller will find it optimal to choose $q_H = \bar{q}$ and $q_L = \underline{q}$.

C Proof of Claim 1




Figure 10: The learning pattern in the single version model

Figure 10 presents the learning pattern in the state space of the difference $\{k_t\}_{t \in \{1, 2, 3, \dots\}}$. To ease notation, I relabel the relevant states as $\{0, 1, 2, 3\}$. The two absorbing states are 0, the highest state within the rejection cascade set, and 3, the lowest state within the buy cascade set. To prove Claim 1 I need to show that in the limit the expected average

purchase numbers $\frac{1}{N}\mathbb{E}(D^N)$ approaches the probability of k_t reaching the absorbing state 3 in Figure 10 before arriving at state 0.

Let $\tau_j := \inf\{t \geq 1 | k_t = j\}, \forall j \in \mathbb{Z}$ be the first time the process $\{k_t\}$ visits j . Then $\tau := \min(\tau_0, \tau_3)$ will be the first time the process hits an absorbing state. No matter where we start, the expected time spent on the non-absorbing states $\{1, 2\}$ is finite, that is, $\mathbb{E}(\tau) < \infty$. Our goal is to prove

$$\lim_{N \rightarrow \infty} \frac{\mathbb{E}(D^N)}{N} = \Pr(\tau_3 < \tau_1).$$

Fix a positive integer N and the number of agents who buy the product will be

$$D_H^N = \sum_{t=1}^N [\mathbf{1}(k_{t-1} = 1, k_t = 2, \tau \geq t) + \mathbf{1}(k_{t-1} = 2, k_t = 3, \tau \geq t)] \quad (1)$$

$$+ \sum_{t=1}^N \mathbf{1}(\tau_3 < t, \tau_3 < \tau_0) \quad (2)$$

The first row counts the number of times when k_t goes upwards before the process hits an absorbing state. By definition, $\sum_{t=1}^N [\mathbf{1}(k_{t-1} = 1, k_t = 2, \tau \geq t) + \mathbf{1}(k_{t-1} = 2, k_t = 3, \tau \geq t)] \leq \tau$. As a result, $\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}(\sum_{t=1}^N [\mathbf{1}(k_{t-1} = 1, k_t = 2, \tau \geq t) + \mathbf{1}(k_{t-1} = 2, k_t = 3, \tau \geq t)]) \leq \lim_{N \rightarrow \infty} \frac{\mathbb{E}(\tau)}{N} = 0$, no matter where the process starts.

The second row (2) counts the number of periods after the process first hits the absorbing state 3 in the premium cascade set.

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^N \mathbf{1}(\tau_3 < t, \tau_3 < \tau_0)\right] &= \sum_{t=1}^N \Pr(\tau_3 = t, \tau_3 < \tau_0)(N - t) \\ &= N \sum_{t=1}^N \Pr(\tau_3 = t, \tau_3 < \tau_0) - \sum_{t=1}^N \Pr(\tau_3 = t, \tau_3 < \tau_0)t \end{aligned}$$

Hence,

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\sum_{t=1}^N \mathbf{1}(\tau_3 < t) \right] &= \sum_{t=1}^{\infty} \Pr(\tau_3 = t, \tau_3 < \tau_0) - \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N \Pr(\tau_3 = t, \tau_3 < \tau_0)t \\ &= \Pr(\tau_3 < \tau_0) - \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N \Pr(\tau_3 = t, \tau_3 < \tau_0)t. \end{aligned}$$

For any positive integer N , the event $\{\tau_3 = t, \tau_3 < \tau_0\} \subset \{\tau = t\}, \forall t \in \{1, 2, \dots, N\}$. Therefore, $\Pr(\tau_3 = t, \tau_3 < \tau_0) \leq \Pr(\tau = t), \forall t \in \{1, 2, \dots, N\}$. It follows that $\sum_{t=1}^N \Pr(\tau_3 = t, \tau_3 < \tau_0)t \leq \sum_{t=1}^N \Pr(\tau = t)t, \forall N \in \mathbb{N}^+$.

$$\lim_{N \rightarrow \infty} \sum_{t=1}^N \Pr(\tau_3 = t, \tau_3 < \tau_0)t \leq \lim_{N \rightarrow \infty} \sum_{t=1}^N \Pr(\tau = t)t = \sum_{t=1}^{\infty} \Pr(\tau = t)t = \mathbb{E}(\tau) < \infty$$

Consequently, $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N \Pr(\tau_3 = t, \tau_3 < \tau_0)t = 0$. Thus, we have proved that $\lim_{N \rightarrow \infty} \frac{\mathbb{E}(D^N)}{N} = \Pr(\tau_3 < \tau_0)$ in the single version model.

D Monotonicity of the buy cascade probability λ

Let me first introduce a few notations in order to derive the buy cascade probability. Conditional on the state of the world V , $\{k_t\}$ follows an asymmetric random walk with upward transition probability $\Pr(k_{t+1} = j+1 | k_t = j, V) = \Pr(s_t = g | V) \in \{\gamma, 1-\gamma\}$ and two absorbing states. A typical state space for k_t looks like Figure 10. Let $\{0, 1, 2, \dots, l\}$ be the state space with two absorbing states 0 and l . Then we can write the buy cascade probability as $\lambda_{i,l}^V := \Pr(\tau_l < \tau_0 | V, k_0 = i)$ where V is the true state and i the starting position. Using the techniques from the classical Gambler's Ruin example, we can derive the following analytical expressions for $\lambda_{i,l}^V$: $\forall i \in \{0, 1, 2, \dots, l\}$, $\lambda_{i,l}^1 = \frac{(\frac{1-\gamma}{\gamma})^{i-1}}{(\frac{1-\gamma}{\gamma})^{l-1}}$ and $\lambda_{i,l}^0 = \frac{(\frac{\gamma}{1-\gamma})^{i-1}}{(\frac{\gamma}{1-\gamma})^{l-1}}$.

So, the ex-ante probability of having a buy cascade will take values from

$\{\lambda_{i,l}\}$ with $\lambda_{i,l} := \Pr(V = 1) \Pr(\tau_l < \tau_0 | V = 1) + \Pr(V = 0) \Pr(\tau_l < \tau_0 | V = 0) = \frac{1}{2}(\lambda_{i,l}^1 + \lambda_{i,l}^0)$. Let me further define $x := \frac{\gamma}{1-\gamma}, \forall \gamma \in (0.5, 1)$. There is a one-to-one mapping between x and γ . Since $\gamma \in (\frac{1}{2}, 1)$, x can take any value from 1 to infinity. For convenience, I will write $\lambda_{i,l}$ as a function of x in the following analysis, $\lambda_{i,l} = \frac{(x^i-1)(x^{l-i}+1)}{2(x^l-1)}, \forall i \in \{0, 1, \dots, l\}$.

Note that $\frac{d\lambda_{i,l}}{dx} = \frac{ix^{2l-i-1} + (l-i)x^{l-i-1} - (l-i)x^{l+i-1} - ix^{i-1}}{2(x^l-1)^2} > 0$. The denominator is obviously positive. Thus it suffices to check the numerator. Denote the numerator by $f(x) := ix^{2l-i-1} + (l-i)x^{l-i-1} - (l-i)x^{l+i-1} - ix^{i-1}$.

Suppose $l < 2i$. I would like to prove $f(x) < 0, \forall x > 1$. $l < 2i$ implies $l - i - 1 < i - 1$ and $l - i - 1 < 2l - i - 1 < l + i - 1$. So we can write $f(x) = x^{l-i-1}(ix^l + (l-i) - (l-i)x^{2i} - ix^{2i-1})$. Let $f_2(x) := ix^l + (l-i) - (l-i)x^{2i} - ix^{2i-1}$ denote the part in the bracket. Because $x > 1$, it suffices to check the sign of $f_2(x)$. Note also $f_2(1) = 0$ and $f_2'(x) = x^{2i-l-1}[ilx^{2l-2i} - (l-i)2ix^l - i(2i-l)]$. For similar reasons, we can look at the sign of whatever is included in the square bracket: $f_3(x) := ilx^{2l-2i} - (l-i)2ix^l - i(2i-l)$. Again we have $f_3(1) = 0$ and $f_3'(x) = x^{2l-2i-1}[2il(l-i) - (l-i)2ilx^{2i-l}]$. Since $x \geq 1$ and $2i - l \geq 0$, we have $f_3'(x) \leq 0, \forall x \geq 1$ (strictly if $x > 1$). Therefore, $f_3(x) \leq 0$ for all $x \geq 1$. Then by going backward, it is easy to see $f(x) < 0, \forall x > 1$.

When $l > 2i$, we can prove $f(x) > 0, \forall x > 1$ in a similar way. If $l = 2i$ instead, $l - i - 1 = i - 1$ and $2l - i - 1 = l + i - 1$. So $f(x) = 0$.

Note also that x is strictly increasing in γ . Hence, $\lambda_{i,l}$ strictly increases in γ if $l > 2i$; strictly decreases in γ if $l < 2i$. In the special case when $l = 2i$, $\lambda_{i,l} = \frac{1}{2}, \forall \gamma \in (\frac{1}{2}, 1)$.

E Proof of Proposition 1

Let me denote the seller's expected average profit by $\tilde{\pi} := \lambda(p)p$. The discussion in the main text has narrowed down the candidates for the optimal

average prices to three: V_{-1}, V_0, V_1 . We can then write the corresponding payoffs as follows:

$$\tilde{\pi}(p, \gamma) = \begin{cases} (1 - \gamma) & , p_H = V_{-1} \\ \frac{1}{2}\lambda_{2,3}(\gamma) & , p_H = V_0 \\ \gamma\lambda_{1,3}(\gamma) & , p_H = V_1 \end{cases}$$

From Appendix D we know $\tilde{\pi}(V_1, \gamma)$ increases in γ whereas $\tilde{\pi}(V_0, \gamma)$ and $\tilde{\pi}(V_{-1}, \gamma)$ decrease in γ . Next I'll prove that $\tilde{\pi}(V_{-1}, \gamma)$ intersects with $\tilde{\pi}(V_0, \gamma)$ only once from above. Then by checking some key values of the functions we can show that, as the signal quality improves, the optimal price goes up, from $p = V_{-1}, V_0$, to V_1 as in Figure 4.

$\frac{\tilde{\pi}(V_{-1}, \gamma)}{\tilde{\pi}(V_0, \gamma)} = 4(1 - 2\gamma + 2\gamma^2 - \gamma^3)$. So $\frac{\tilde{\pi}(V_{-1}, \gamma)}{\tilde{\pi}(V_0, \gamma)} > 1 \iff 4\gamma^3 + 8\gamma^2 - 8\gamma - 3 < 0$. Let $m(\gamma) := 4\gamma^3 + 8\gamma^2 - 8\gamma - 3$. By checking $m(0), m(1)$ and $m'(\gamma)$, we know $m(\gamma)$ switches its sign only once on $(0.5, 1)$. This implies $\tilde{\pi}(V_{-1}, \gamma)$ and $\tilde{\pi}(V_0, \gamma)$ intersect only once on $\gamma \in (0.5, 1)$.

Let $\underline{\gamma}$ be the solution to the equation $\tilde{\pi}(V_{-1}, \gamma) = \tilde{\pi}(V_0, \gamma)$ and $\bar{\gamma}$ the solution to $\tilde{\pi}(V_0, \gamma) = \tilde{\pi}(V_1, \gamma)$. It is easy to show that $\underline{\gamma} \in (0.65, 0.7)$ and $\bar{\gamma} \in (0.75, 0.8)$ by checking the numerical values of the profit functions at $\gamma = 0.65, 0.7, 0.75$ and 0.8 . Hence, within the interval $(0.5, 1)$, $\tilde{\pi}(V_{-1}, \gamma) > \tilde{\pi}(V_0, \gamma)$ if and only if $\gamma > \underline{\gamma}$ and $\tilde{\pi}(V_0, \gamma) > \tilde{\pi}(V_1, \gamma)$ if and only if $\gamma > \bar{\gamma}$. It pins down the order of the optimal prices.

F Multi-version: solving the seller's problem

Let D_H^N denote the purchase number of the premium version and D_L^N denote the purchase number of the basic version when the total number of agents

is N . Then we can write the seller's optimization problem as:

$$\max_{p_L, p_H} \mathbb{E} \left(\lim_{N \rightarrow \infty} \frac{1}{N} (D_L^N p_L + D_H^N p_H) \right) \text{ s.t. } p_L \neq \frac{p_H}{2}, \frac{p_H}{2} \leq 1$$

F.1 Simplifying the seller's problem


As I briefly discussed in section 5.2, the seller's pricing choice is equivalent to choosing a learning pattern¹², and within each pattern, a starting position. Let me first derive the seller's expected payoff as a function of the prices $p = (p_H, p_L)$ and γ . Then I will rule out some combinations of learning patterns and starting positions by simple arguments and algebra before I discuss more complicated cases in Appendix F.2.

Fix a positive integer N and let $\pi^N := \frac{\mathbb{E}(D_L^N)}{N} p_L + \frac{\mathbb{E}(D_H^N)}{N} p_H$. If $p_L > \frac{p_H}{2}$ (Figure 11a), we are back to the single-version model. As Appendix C shows, $\lim_{N \rightarrow \infty} \frac{\mathbb{E}(D_L^N)}{N} = 0$; $\lim_{N \rightarrow \infty} \frac{\mathbb{E}(D_H^N)}{N}$ equals the premium cascade probability. This choice obviously outperforms the pattern 11e.


When $p_L < \frac{p_H}{2}$ we can derive the limit of π^N by similar arguments as in Appendix C. The core argument rests upon a result from probability theory: the expected time a simple random walk process spends on the non-absorbing states in between two absorbing states is finite. As a result, the seller's limit expected profit will be $\tilde{\pi}(p, \gamma) := \lim_{N \rightarrow \infty} \pi^N(p, \gamma) = \Pr(\text{premium cascade} | p, \gamma) p_H(\gamma) + \Pr(\text{premium cascade} | p, \gamma) p_L(\gamma)$.

The following paragraphs help us rule out several suboptimal cases. First, starting from a rejection cascade set is never optimal. Starting from a state in the premium cascade set, or equivalently, choosing $p_H - p_L \leq V_{-1}$ in pattern 11b - 11d, is not optimal either. These patterns occur only when $p_L < \frac{p_H}{2}$, which implies $p_L < p_H - p_L \leq V_{-1}$. Hence, $p_H < 2V_{-1}$. Note also $\tilde{\pi}(p, \gamma) = p_H$ if we start with a premium cascade. It follows that starting from the premium cascade set here in pattern 11b - 11d is dominated by


¹²I have collected all the possible patterns in Figure 11 for convenience.




(a) Short learning pattern: $p_L > \frac{p_H}{2}$




(b) $p_L < \frac{p_H}{2} \leq 1, k_H - 2 < k_L$



(c) $p_L < \frac{p_H}{2} \leq 1, k_L = k_H - 2$



(d) $p_L < \frac{p_H}{2} \leq 1, k_L = k_H - 1$



(e) $p_L < \frac{p_H}{2} \leq 1, k_L = k_H$

Figure 11: Characterizing learning patterns in belief space $\{V_k\}_{k \in \mathbb{Z}}$. Note: 'rejection' refers to the rejection cascade set, 'basic' the buy cascade set for the basic version, and 'premium' the buy cascade set for the premium version. In the short learning pattern 11a, 'Learn_H' represents the learning set for the premium version as defined in the single version model: the agent chooses to buy the premium version only after a good signal. $V_{k_H}(p_H)$ refers to the threshold posterior above which an agent prefers the premium version over rejection. Similarly, in patterns 11b - 11e, 'Learn_L' represents the learning set for basic version. $V_{k_L}(p_L)$ is the threshold belief above which an agent prefers the basic version over rejection. 'Learn_H' denotes the learning set for the premium version and $V_{k_L}(p_H - p_L)$ the threshold belief above which an agent prefers the premium version over the basic version.

starting from the premium cascade set in the short learning pattern 11a which gives a payoff of $2V_{-1}$. Moreover, starting from the basic cascade set in pattern 11b gives an even lower payoff to the seller and thus is suboptimal.

Third, the seller will not choose to start from within the basic learning set in pattern 11b. If we start there, the whole learning process resembles the short learning pattern 11a. In both cases, agents consider only one version along the way: the basic version in the former; the premium version in the latter. Selling the premium version is obviously more profitable.

The remaining choices for the seller are summarized in Table 2.¹³

| Case # | Prices | Starting position | Expected payoff |
|----------------------|--|-------------------|--|
| a1 a2 a3 | $p_H = 2V_1, p_L > V_1$ $p_H = 2V_0, p_L > V_0$ $p_H = 2V_{-1}, p_L > V_{-1}$ | | $\tilde{\pi}_{a1} := 2\lambda_{1,3}V_1$ $\tilde{\pi}_{a2} := 2\lambda_{2,3}V_0$ $\tilde{\pi}_{a3} := 2V_{-1}$ |
| b1 b2 | $p_L = V_{-2}, p_H = V_1 + V_{-2}$ $p_L = V_{-3}, p_H = V_0 + V_{-3}$ | | $\tilde{\pi}_{b1} := \lambda_{1,3}V_1 + V_{-2}$ $\tilde{\pi}_{b2} := \lambda_{2,3}V_0 + V_{-3}$ |
| c1 c2 c3 c4 | $p_L = V_1, p_H = V_1 + V_3$ $p_L = V_0, p_H = V_0 + V_2$ $p_L = V_{-1}, p_H = V_{-1} + V_1$ $p_L = V_{-2}, p_H = V_{-2} + V_0$ | | $\tilde{\pi}_{c1} := \lambda_{1,5}(V_1 + V_3)$ $\tilde{\pi}_{c2} := \lambda_{2,5}(V_0 + V_2)$ $\tilde{\pi}_{c3} := \lambda_{3,5}(V_{-1} + V_1)$ $\tilde{\pi}_{c4} := \lambda_{4,5}(V_{-2} + V_0)$ |
| d1 d2 d3 | $p_L = V_1, p_H = V_1 + V_2$ $p_L = V_0, p_H = V_0 + V_1$ $p_L = V_{-1}, p_H = V_{-1} + V_0$ | | $\tilde{\pi}_{d1} := \lambda_{1,4}(V_1 + V_2)$ $\tilde{\pi}_{d2} := \lambda_{2,4}(V_0 + V_1)$ $\tilde{\pi}_{d3} := \lambda_{3,4}(V_{-1} + V_0)$ |

Table 2: The seller's remaining choices and their payoffs.

Note: in the starting position column, I plot the state space of k_t . The big points with the case number below refer to the starting position of the process for each case. 'rej' stands for the rejection cascade set, 'bas' the buy cascade set for basic version, 'pre' the buy cascade set for the premium version, 'Learn_L' the learning set for the basic version, and 'Learn_H' the learning set for the premium version.

¹³Both the premium cascade probability $\lambda_{i,l}$ and the posterior belief V_k are functions of γ . In addition, in the remaining choices, the probability of having a basic cascade is always zero.

F.2 Proof of Proposition 2 and 3

I copy some notations from Appendix D here. $x = \frac{\gamma}{1-\gamma}$ and $\lambda_{i,l} = \frac{(x^i-1)(x^{l-i}+1)}{2(x^l-1)}$, $\forall i \in \{0, 1, \dots, l\}$. I want to prove the optimal choice will be case a3, a2, c2 and c1 in Table 2 as the signal quality improves. Here is a road map of the proof:

1. I will rule out four cases b1, b2, d3 and c3 because each of them is outperformed by a linear combination of some other cases;
2. None of the choices from $\{d1, d2, a1, c4\}$ are optimal. I will show that at any signal quality γ their payoffs are lower than the maximum of two optimal payoffs from $\{a3, a2, c2, c1\}$;
3. I prove that the order of the optimal cases in Table 1 is correct.

Case b1: First, I want to show $\tilde{\pi}_{b1} \leq \frac{1}{2}(\tilde{\pi}_{a1} + \tilde{\pi}_{a3})$, which implies $\tilde{\pi}_{b1} \leq \max\{\tilde{\pi}_{a1}, \tilde{\pi}_{a3}\}$ and thus b1 is never optimal. Since V_k increases in k , $\frac{1}{2}(\tilde{\pi}_{a1} + \tilde{\pi}_{a3}) = \lambda_{1,3}V_1 + V_{-1} \geq \lambda_{1,3}V_1 + V_{-2} = \tilde{\pi}_{b1}$.

Case b2: Similarly, b2 is also suboptimal because $\tilde{\pi}_{b2} = \lambda_{2,3}V_0 + V_{-3} \leq \lambda_{2,3}V_0 + V_{-1} = \frac{1}{2}(\tilde{\pi}_{a2} + \tilde{\pi}_{a3})$.

Case d3: I am going to prove that d3 is not optimal because $\tilde{\pi}_{d3} \leq \alpha\tilde{\pi}_{a2} + (1-\alpha)\tilde{\pi}_{a3}$ for some $\alpha \in (0, 1)$. Note that $\tilde{\pi}_{d3} = \frac{1-\gamma(1-\gamma)}{2(1-2\gamma(1-\gamma))}(V_0 + V_{-1})$, $\tilde{\pi}_{a2} = \frac{1}{2(1-\gamma(1-\gamma))}$ and $\tilde{\pi}_{a3} = 2V_{-1}$. So we can write $\tilde{\pi}_{d3}$ as

$$\tilde{\pi}_{d3} = \frac{(1-\gamma(1-\gamma))^2}{2(1-2\gamma(1-\gamma))}\tilde{\pi}_{a2} + \frac{1-\gamma(1-\gamma)}{4(1-2\gamma(1-\gamma))}\tilde{\pi}_{a3}$$

To simplify the notation, let $\phi := \gamma(1-\gamma)$. $\gamma \in (\frac{1}{2}, 1) \Rightarrow \phi \in (0, \frac{1}{4})$. Then the sum of the two coefficients in the above expression becomes

$$\frac{(1-\phi)^2}{2(1-2\phi)} + \frac{1-\phi}{4(1-2\phi)} = \frac{(1-\phi)(3-2\phi)}{4(1-2\phi)}$$

This sum is less than 1 if and only if $2\phi^2 + 3\phi - 1 < 0$, which holds true for all $\phi \in (0, \frac{1}{4})$. Hence, $\tilde{\pi}_{d3} < \frac{(1-\gamma(1-\gamma))^2}{2(1-2\gamma(1-\gamma))} \tilde{\pi}_{a2} + (1 - \frac{(1-\gamma(1-\gamma))^2}{2(1-2\gamma(1-\gamma))}) \tilde{\pi}_{a3}$ with $\frac{(1-\gamma(1-\gamma))^2}{2(1-2\gamma(1-\gamma))} \in (0, 1)$.

Case c3: Finally, I will show that $\tilde{\pi}_{c3}(\gamma) < \tilde{\pi}_{a2}(\gamma)$ for all $\gamma \in (\frac{1}{2}, 1)$. Or equivalently, $\tilde{\pi}_{c3}(x) < \tilde{\pi}_{a2}(x), \forall x \in (1, \infty)$.¹⁴ $\frac{\tilde{\pi}_{c3}}{\tilde{\pi}_{a2}} < 1 \iff (x^3 - 1)^2(x^2 + 1) < (x^5 - 1)(x^2 - 1)(x + 1) \iff x(x^4 + 1)(x - 1)^2 > 0$, which is true for all $x \in (1, \infty)$.

F.2.1 Step 2: d1, d2, a1 and c4 are sub-optimal

The proofs have a similar structure: each of them will be outperformed by one of the optimal cases on $(\frac{1}{2}, \hat{\gamma})$ and by another optimal case on $(\hat{\gamma}, 1)$. The threshold $\hat{\gamma}$ I choose may vary on a case-by-case basis.

Case d1: I want to show $\tilde{\pi}_{d1} \leq \max(\tilde{\pi}_{a2}, \tilde{\pi}_{c1})$

Step 1: monotonicity of the expected payoff functions. $\tilde{\pi}_{d1} = \lambda_{1,4}(V_1 + V_2)$. According to Appendix D $\lambda_{1,4}$ is strictly increasing in γ . In addition, it is easy to show that V_k is strictly increasing in γ if $k > 0$ and strictly decreasing in γ if $k < 0$. Hence, $\tilde{\pi}_{d1}$ strictly increases in γ . Similarly, we can prove that $\tilde{\pi}_{c1} = \lambda_{1,5}(V_1 + V_3)$ is strictly increasing in γ and $\tilde{\pi}_{a2} = \lambda_{2,3}$ is strictly decreasing in γ . As a result, $\tilde{\pi}_{a2}(\gamma)$ intersects with each of the other two functions from above only once on $(0.5, 1)$.

Step 2: find $\hat{\gamma}$. By checking the MATLAB figures of the polynomial functions, I conjecture 0.75 would be a good candidate for $\hat{\gamma}$. With MATLAB I calculate the numerical values of the relevant functions at $\hat{\gamma} = 0.75$: $\tilde{\pi}_{a2}(0.75) \approx 0.6154$, $\tilde{\pi}_{c1}(0.75) \approx 0.5809$, $\tilde{\pi}_{d1}(0.75) \approx 0.5775$. As $\tilde{\pi}_{a2}(0.75) > \tilde{\pi}_{c1}(0.75) > \tilde{\pi}_{d1}(0.75)$, we must have $\tilde{\pi}_{d1}(\gamma) < \tilde{\pi}_{a2}(\gamma), \forall \gamma \in (\frac{1}{2}, \hat{\gamma}]$.

Step 3: prove $\tilde{\pi}_{c1}(\gamma) \geq \tilde{\pi}_{d1}(\gamma)$ on $(\hat{\gamma}, 1)$. $\tilde{\pi}_{c1} - \tilde{\pi}_{d1} = \lambda_{1,5}(V_1 + V_3) - \lambda_{1,4}(V_1 + V_2) = \frac{(x-1)x}{2(x+1)} \left[\frac{(x^4+1)(2x^2+1-x)}{(x^5-1)(x^2+1-x)} - \frac{(x^3+1)(2x^2+x+1)}{(x^4-1)(x^2+1)} \right]$. Because $x > 1$,

¹⁴Recall that I define $x = \frac{\gamma}{1-\gamma}$ in Appendix D.

$\frac{(x-1)x}{2(x+1)} > 0$. Hence, $\tilde{\pi}_{31} \geq \tilde{\pi}_{41} \Leftrightarrow \frac{(x^4+1)(2x^2+1-x)}{(x^5-1)(x^2+1-x)} \geq \frac{(x^3+1)(2x^2+x+1)}{(x^4-1)(x^2+1)}$.

Since the denominators are positive when $x > 1$, it suffices to show $g(x) := (x^4 + 1)(2x^2 + 1 - x)(x^4 - 1)(x^2 + 1) - (x^3 + 1)(2x^2 + x + 1)(x^5 - 1)(x^2 + 1 - x) \geq 0$. Note that $x > 3$ when $\gamma \in (0.75, 1)$. We have proved in the previous step that $\tilde{\pi}_{31}(0.75) - \tilde{\pi}_{41}(0.75) > 0$. So $g(3) > 0$. Rearrange the expression of $g(x)$ and we have $g(x) = x(x^8(x - 3) + x^5(x^2 - 1) + x^4 + x(x - 1) + 1)$. Then it is easy to see $g(x) > 0, \forall x \in (3, \infty)$. Hence, $\tilde{\pi}_{c1}(\gamma) - \tilde{\pi}_{d1}(\gamma) > 0, \forall \gamma \in (0.75, 1)$.

To sum up, we have proved $\tilde{\pi}_{d1}(\gamma) < \max(\tilde{\pi}_{a2}(\gamma), \tilde{\pi}_{c1}(\gamma)), \forall \gamma \in (0.5, 1)$. Case d1 is never an optimal choice for the seller.

Case d2: $\tilde{\pi}_{d2}(\gamma) < \max(\tilde{\pi}_{a2}(\gamma), \tilde{\pi}_{c2}(\gamma))$. The proof is analogous to the one above. The threshold $\hat{\gamma}$ I choose is 0.72. It's easy to complete the proof with the same procedure as case d1, so I do not repeat it here.

Case a1: $\tilde{\pi}_{a1}(\gamma) < \max(\tilde{\pi}_{a2}(\gamma), \tilde{\pi}_{c1}(\gamma))$ The proof is analogous to the one above. The threshold $\hat{\gamma}$ I choose is 0.75. Again I do not repeat the proof here.

Case c4: $\tilde{\pi}_{c4} < \max(\tilde{\pi}_{a2}, \tilde{\pi}_{a3})$. First, I will prove $\tilde{\pi}_{c4} < \tilde{\pi}_{a3}, \forall x = \frac{\gamma}{1-\gamma} \in (1, 2]$ and $\tilde{\pi}_{c4} < \tilde{\pi}_{a2}, \forall x \in (2, \infty)$.

Step 1. I will prove $\tilde{\pi}_{c4} < \tilde{\pi}_{a3}, \forall x = \frac{\gamma}{1-\gamma} \in (1, 2]$. Notice that the expected payoffs of all the remaining choices in Table 2 are positive. $\frac{\tilde{\pi}_{c4}}{\tilde{\pi}_{a3}} = \frac{(x+1)^2(x^2+3)(x+1)}{8(x^4+x^3+x^2+x+1)} > 1$ if and only if $x^4(x - 5) - 2x^2(x - \frac{1}{2}) + (x - 5) > 0$, which holds true for $x \in (1, 2]$.

Step 2. I will show $\tilde{\pi}_{c4} < \tilde{\pi}_{a2}, \forall x \in (2, \infty)$. $\frac{\tilde{\pi}_{c4}}{\tilde{\pi}_{a2}} = \frac{(x^3-1)(x^2+3)}{2(x^5-1)}$. Similarly, we can show that $\frac{\tilde{\pi}_{c4}}{\tilde{\pi}_{a2}} < 1 \iff x^4(x - 2) + 2x^3(x - \frac{3}{2}) + x^2 + 1 > 0$, which holds true when $x \in (2, \infty)$.

F.2.2 Optimality of case a3, a2, c2, c1 and their order

Figure 8 plots the expected payoffs of the optimal cases as a function of γ .

| Approximate values | $\tilde{\pi}_{a3}$ | $\tilde{\pi}_{a2}$ | $\tilde{\pi}_{c2}$ | $\tilde{\pi}_{c1}$ |
|--------------------|--------------------|--------------------|--------------------|--------------------|
| $\gamma = 0.505$ | 0.99 | 0.6666 | 0.404 | 0.2041 |
| $\gamma = 0.685$ | 0.63 | 0.6376 | 0.5854 | 0.4599 |
| $\gamma = 0.75$ | 0.5 | 0.6154 | 0.6479 | 0.5809 |
| $\gamma = 0.995$ | 0.01 | 0.5025 | 0.75 | 0.9925 |

Table 3: Approximate values of the expected payoff functions.

Step 1. In the single-version model (Appendix E) I have already proved $\tilde{\pi}_{a3}(\gamma)$ and $\tilde{\pi}_{a2}(\gamma)$ intersect only once on $(0.5, 1)$. As I will show below, the same happens with $\tilde{\pi}_{c2}(\gamma)$ and $\tilde{\pi}_{c1}(\gamma)$.

$\frac{\tilde{\pi}_{c2}}{\tilde{\pi}_{c1}} = \frac{(x+1)(3x^2+1)(x^3+1)^2}{2(x^2+1)(x^4+1)(2x^3-x^2+x)} \geq 1 \iff m_2(x) := -x^9 + 5x^8 - 5x^7 + 9x^6 + 4x^4 - x^3 + 5x^2 - x + 1 \geq 0$. First, I will prove $m_2(x) > 0$ for $x \in (1, 4]$. Then I will show that $m_2'(x) < 0, \forall x \in (4, \infty)$. As $m_2(x)$ is continuous on $(1, \infty)$ and $\lim_{x \rightarrow \infty} m_2(x) = -\infty$, it must intersect with the horizontal axis only once on $(4, \infty)$.

Rearrange the $m_2(x)$ function and we have $m_2(x) = -x^8(x-4) + x^6(x^2 - 5x + 9) + 4x^3(x - \frac{1}{4}) + 5x(x - \frac{1}{5}) + 1$. Because $x \in (1, 4]$, $-x^8(x-4) \geq 0$, $x - \frac{1}{4} > 0$ and $x - \frac{1}{5} > 0$. $x^2 - 5x + 9 = (x - 2.5)^2 + \frac{11}{4} > 0$. Hence, $m_2(x) > 0, \forall x \in (1, 4]$.

Next, by rearranging the first order derive of $m_2(x)$ I find $m_2'(x) = -9x^6(x-4)(x - \frac{4}{9}) + x^3(-17x^3 + 54x^2 + 16) - 3x(x - \frac{10}{3}) - 1$. With $x > 4$ we can easily show that the first and third term is negative. As for the second term, $-17x^3 + 54x^2 + 16 = -(x^3 - 16) - 16x^2(x - \frac{54}{16}) < 0, \forall x \in (4, \infty)$. Hence, we can conclude $m_2'(x) < 0$ on $(4, \infty)$.

Step 2. From previous sections we know $\tilde{\pi}_{a3}(\gamma)$ and $\tilde{\pi}_{a2}(\gamma)$ are strictly decreasing in γ . Besides, $\tilde{\pi}_{c2}(\gamma)$ and $\tilde{\pi}_{c1}(\gamma)$ strictly increases in γ . Therefore, either of the two functions $\tilde{\pi}_{a3}(\gamma)$ and $\tilde{\pi}_{a2}(\gamma)$ intersects with either of $\tilde{\pi}_{c2}(\gamma)$ and $\tilde{\pi}_{c1}(\gamma)$ once and from above on $(0.5, 1)$. In other words, there will be at most one intersection point on $(0.5, 1)$ for each pair from $\{\tilde{\pi}_{13}, \tilde{\pi}_{12}\} \times \{\tilde{\pi}_{32}, \tilde{\pi}_{31}\}$. Table 3 shows the following inequalities hold.

$$\text{At } \gamma = 0.505, \tilde{\pi}_{a3} > \tilde{\pi}_{a2} > \tilde{\pi}_{c2} > \tilde{\pi}_{c1}. \quad (3)$$

$$\text{At } \gamma = 0.995, \tilde{\pi}_{a3} < \tilde{\pi}_{a2} < \tilde{\pi}_{c2} < \tilde{\pi}_{c1}. \quad (4)$$

$$\text{At } \gamma = 0.685, \tilde{\pi}_{a2} > \tilde{\pi}_{a3} > \tilde{\pi}_{c2} > \tilde{\pi}_{c1}. \quad (5)$$

$$\text{At } \gamma = 0.75, \tilde{\pi}_{c2} > \tilde{\pi}_{a2} > \tilde{\pi}_{c1} > \tilde{\pi}_{a3}. \quad (6)$$

Together with previous observations, the inequalities (3) and (5) have two implications. First, case c2 and c1 are suboptimal on $(0.5, 0.685)$. Besides, $\tilde{\pi}_{a2}$ and $\tilde{\pi}_{a3}$ intersects at some $\gamma \in (0.5, 0.685)$. Let γ_s be the solution to $\tilde{\pi}_{a3}(\gamma) = \tilde{\pi}_{a2}(\gamma)$. Then, among the single-version policies, a3 is better than a2 on $(0.5, \gamma_s)$ and the opposite happens on $(\gamma_s, 1)$.

Similarly, the inequalities (4) and (6) tells us, among the multi-version policies, c1 is better than case c2 on $(\gamma_m, 1)$ and the opposite happens on $(0.5, \gamma_m)$, where γ_m is the solution to $\tilde{\pi}_{c1}(\gamma) = \tilde{\pi}_{c2}(\gamma)$.

Consequently, the optimal case between (γ_s, γ_m) is either case c2 or a2. By inequalities (5) and (6) we know $\tilde{\pi}_{a2}$ intersects with $\tilde{\pi}_{c2}$ from above on $(0.685, 0.75)$. Let γ_v be the solution to $\tilde{\pi}_{a2}(\gamma) = \tilde{\pi}_{c2}(\gamma)$. Then a2 will be the unique optimal choice on (γ_s, γ_v) and c2 will be the unique optimal choice on (γ_v, γ_m) .

Hence, the order of the optimal prices in Table 1 is correct.