Can Price Inform Quality when Verification is Costly?

Daniel Kwiatkowski

October 6, 2025

Abstract

A product's price will reflect its quality, if that quality is known to consumers. But many experiments have shown that consumers believe high prices signal high quality, even in situations where the quality is *not* generally known. It's unclear if consumers are right or wrong to expect this. I examine informational efficiency in the lab, using a market where verifying product quality is costly. Theoretically, prices could convey information according to Grossman and Stiglitz (1980), but might also convey no information. I find that prices endogenously convey about as much information as theoretically possible, even when information is quite costly. Behavioral bias, which I examine using quantal response, can make prices slightly more or less informative than theory predicts.

1 Introduction

Prices are often correlated with product quality. Ceteris parabus, a first-class airline ticket is more expensive than economy class, a house with a beautiful view is more expensive than one without, and an acclaimed bottle of French wine is more expensive than a boxed wine grown in Houston. Does this mean if I encounter two otherwise indistinguishable wines, one more expensive than the other, I can surmise that the more expensive wine will be better than the cheaper wine? Not necessarily. It may be that prices reflect quality only because quality information is already obvious to consumers, so that if quality is *not* apparent, there is no reason for the better product to be more expensive.

In general, it is difficult to say whether prices are doing informational work: successfully conveying quality information to otherwise uninformed consumers. But in many cases, consumers believe that prices can do informational work. In the lab, it has been shown that buyers believe prices convey quality information even when the quality is not obvious otherwise (Leavitt 1954, McConnell 1968, Olson 1977). These experiments are focused on understanding the buyers' beliefs, and involve experimenters varying the price of a product as a treatment. The experiments show that buyers believe prices are somewhat informative of quality, but it's unclear if buyers are correct to believe that prices would be informative when quality is not obvious, or if their beliefs are an artifact from real-world markets in which quality is generally known (or even a reflection of what situations they expect experimenters to present them with).

This paper examines whether prices convey quality information endogenously in a laboratory experiment where quality is not freely observable. There are two reasons why it might be valuable to see informative pricing arise endogenously in the lab. The first is to understand the mechanism by which prices become informative. Theoretically, there are many models that lead to informative pricing and in real-world markets it is generally not clear which channels, if any, are operative.

The other reason to examine endogenous price information in the lab is to see how behavioral biases may impact how much information prices convey. Noise in decision-making, risk-aversion, and learning may make prices more or less informative than theoretical models predict. In the lab, behavioral biases can be seen clearly, while in real-world markets their effects may be confounded with variation from other sources such as heterogeneity in consumers' values, producers' costs, and product characteristics.

I examine the informational content of prices through the channel studied in Grossman and Stiglitz (1980), and I follow Bester and Ritzberger (2001) in adapting this intuition to a buyer-seller framework. Here, prices convey some information in equilibrium because consumers can exert effort to verify product quality. I find that this channel is operative in the lab. In every experimental session, subjects reach the separating equilibrium in which prices are informative. I also find that the gradual responsiveness of subject behavior, relative to the sharp discontinuities in theoretical best responses, can make prices either more or less informative than theory predicts.

These results demonstrate the efficient markets hypothesis as stated by Grossman and Stiglitz. The prices observed in the lab are about as informative as they can be theoretically. An individual subject will update their beliefs about quality after seeing the price and will then be indifferent between either verifying the quality or buying the product without verifying.

But there are also behavioral barriers to efficiency. Subjects are very responsive to the cost of verifying product quality. When verification is cheap, buyers verify more frequently than necessary; ignoring that, since others are also verifying, the price already reflects information about quality. Instead, buyers do not trust the market enough. When information is expensive, the opposite happens. Buyers get information rarely, and trust the market too much, failing to realize that, since others are also getting less information, the price is a less reliable guide to quality.

2 Background

Formulation of the efficient markets hypothesis by Samuelson (1965) and Fama (1970) generally assumed that information was freely available to investors who could then bid the price up or down through buying or selling decisions, until no further arbitrage was profitable. Grossman and Stiglitz (1980) pointed out that, if information is freely observable, then it is immaterial whether prices convey information. Prices are doing meaningful informational work only when they are conveying information to otherwise uninformed buyers. Grossman and Stiglitz thus assumed arbitrage was costly and reformulated the efficient markets hypothesis so that the return to arbitrage perfectly offset its cost. Markets were then "efficient" if prices conveyed as much information as possible in equilibrium—just enough to make buyers indifferent to arbitrage opportunities. Informative pricing is sustained by buyers who incur costs in order to benefit from the deviation of current prices from expected returns.

Since Grossman and Stiglitz, many other models have appeared that sustain informative pricing. Some involve quality entering demand (through, for instance, a proportion of informed consumers, as in Bagwell and Riordan, 1991). Some involve quality entering through a firm's cost (e.g. Tirole, 1988, pp. 107-108). In 2001, Bester and Ritzberger made the Grossman and Stiglitz intuition more concordant with these models by adapting it for one buyer and one single-product monopolist. This gives the model the same separating/pooling dimension as Bagwell and Riordan but without a proportion of already informed consumers. Instead, arbitrage is captured by buyers who can exert costly effort to learn the true quality of the product.

I have chosen to use a discretized version of Bester and Ritzberger's model to test price informativeness in the lab. One reason for this is that the model is simple and easy for subjects to understand. Furthermore, people in real-world markets do expend effort to learn about the quality of products prior to buying. So, this is certainly one active informational channel, though it may not be the only one.

3 The Model

A seller is endowed with a single product of quality v, known only to themselves, and chooses a price p at which to sell the product. If the product is not sold, the seller gets no payoff—the product is useless to the seller. If the product does sell, the seller gets the price they set. Thus the seller's payoff is

$$\pi_{\text{seller}} = p\mathbb{1}\{B\}$$

where $\mathbb{1}\{B\}$ is an indicator for whether the buyer buys.

The buyer knows the prior distribution of v and observes the price set by the seller. The buyer then chooses one of three options. The buyer can buy the product, in which case they receive $v–p$, or they can walk away without buying, in which case they receive 0. The buyer can alternatively incur an effort cost of c to observe v. The cost c is not a transfer to the seller, it is simply a loss, representing the cost of time or other resources used. If the buyer pays c to observe v, they can then decide to buy and receive $v–p–c$ or walk away and receive $-c$. Overall, the buyer's payoff is

$$\pi_{\text{buyer}} = \mathbb{1}\{B\}(v–p) - \mathbb{1}\{I\}c$$

where $\mathbb{1}\{B\}$ is an indicator for whether the buyer buys and $\mathbb{1}\{I\}$ is an indicator for whether they pay for information.

To implement this game in the experiment, I restrict v to be either a high or low quality, and restrict the seller to two possible price choices, a high price and a low price. The timing of the game is as follows (the full extensive form is in the appendix). First, nature chooses the seller's quality:

$$v = \begin{cases} v_l \text{ w.p. } 1/2 \\ v_h \text{ w.p. } 1/2 \end{cases}$$

Then the seller chooses $p \in \{p_l, p_h\}$ where

$$p_l < v_l < \frac{v_l + v_h}{2} < p_h < v_h$$

After observing p, the buyer chooses whether to buy (B), leave (L), or get information (I). If the buyer chooses to get information, they observe v and then choose to buy or leave.

3.1 Equilibria

I will focus on the perfect Bayesian equilibria of the game. First, notice that once the buyer gets information, the next choice is straightforward. Once the cost of information is sunk, the buyer should buy whenever $v > p$. Next, note that the buyer should buy blindly whenever $p = p_l$. Since $p_l < v_l < v_h$, regardless of the quality of the product, it is worthwhile to buy. There is no benefit from acquiring information in this case. Taking these two things into account, we can rewrite the game tree (fig 1). The only difficult question for the buyer is what to do when $p = p_h$.

If a buyer observes the high price, should they buy blindly, hoping it is a high-quality product? Should they walk away, assuming it isn't worth the risk? Or should they incur costly effort to check the quality? The answer depends on the buyer's beliefs. Suppose the buyer has updated beliefs $\mu_h \equiv \mathbb{P}(v = v_h | p = p_h)$ after seeing the high price. Then if the buyer buys blindly, they get

$$\mathbb{E}[v|p_h] - p_h = \mu_h v_h + (1 - \mu_h) v_l - p_h = v_l - p_h + (v_h - v_l)\mu_h$$

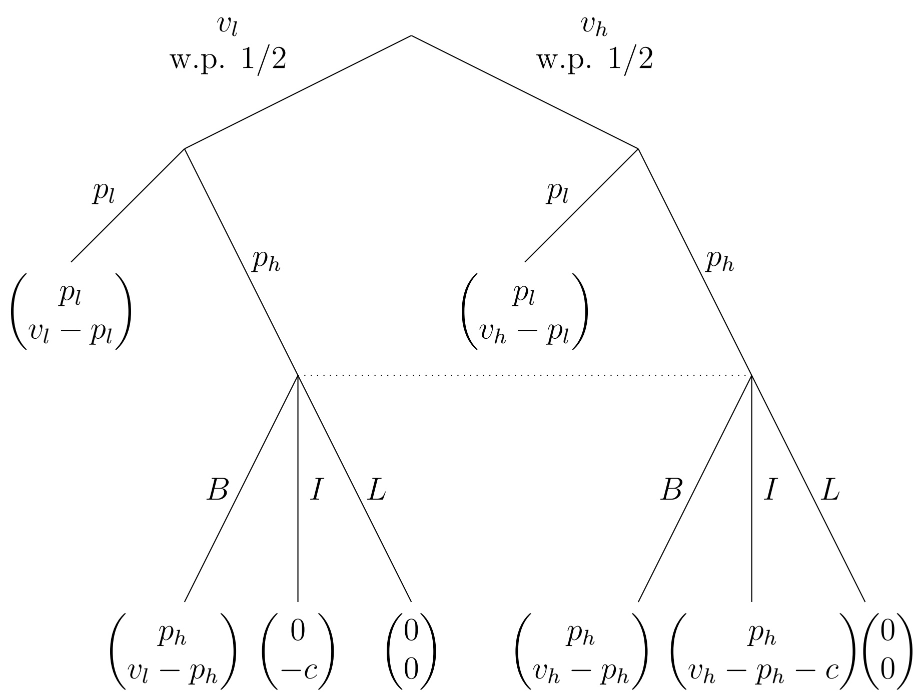

Figure 1: Should Buyers Buy Expensive Products?

Notes: The first decision is a move by nature that determines if the seller has high- or low-quality products. Then the seller chooses a price. If the seller chooses the low price, the buyer should buy. If the seller chooses the high price, the buyer must choose either buy (B), get information (I), or leave (L), depending on their posterior beliefs about the type of seller that would set a high price.

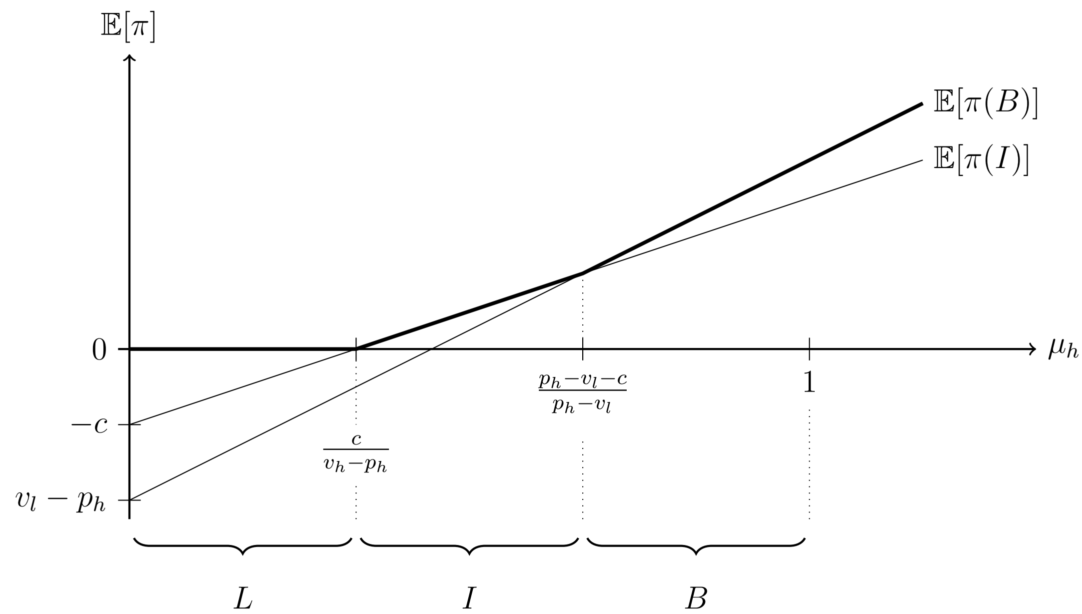

Figure 2: Optimal Buyer Choice for Different Beliefs

Notes: μ_h is the buyer's belief that a seller setting the high price has a high-quality product. If the buyer thinks high-priced products are very likely to be high-quality, they should buy (B). If they think high-priced products are unlikely to be high-quality, they should leave (L). For intermediate beliefs, they should choose to get information (I).

If the buyer gets information, they pay c but have a μ_h chance of getting $v_h – p_h$ if the product is high-quality, so their expected payoff is

$$\mu_h(v_h – p_h) - c$$

Figure 2 shows the optimal buyer best response to $p = p_h$ for different values of beliefs μ_h. If it is very likely that a high price indicates a high quality product, the buyer should just buy without checking. If it is more likely that the quality might be low, the buyer should check before buying. If it is almost certain that the product is low-quality, then it is not even worthwhile to check, and the buyer should just leave.

Suppose μ is low, so the buyer thinks a high price is quite likely to have come from a low-type seller. Then the buyer will leave without buying or getting information when $p = p_h$. If this is the case, sellers will have to set $p = p_l$ to make a sale, and this equilibrium is sustained by off-equilibrium-path beliefs: the buyer is justified in believing high prices may come from low-quality sellers because high prices never occur.[1] This is an uninformative equilibrium because the buyer learns nothing from the price—it conveys no information about the quality of the product.

If instead μ_h is higher, so that the buyer thinks the high price is more likely to come from a high-type seller, then they will buy blindly or get information. If they always bought blindly, the low-type firm would have an incentive to fool the buyer by setting $p = p_h$ as well, and prices would no longer be informative. This is not an equilibrium because $p_h > (v_l + v_h)/2$, so if both types are setting $p = p_h$, it is not worthwhile to buy.

If the buyer always got information when seeing the high price, the low-type seller would have to set $p = p_l$ while the high-type seller could set $p = p_h$. But then the price would be perfectly informative, and paying for information would be a waste—the product quality is already obvious from the price. So perfect type dependence is also not an equilibrium.

Suppose μ_h is exactly at the threshold between buying blindly and getting information, so that the buyer is indifferent between the two. Here there is a partially separating equilibrium. The buyer mixes between the two strategies and the low-type seller makes this mixing plausible by occasionally trying to fool the buyer by setting $p = p_h$. Prices are as informative as possible in this equilibrium. The price conveys just enough information so that buyers are indifferent between trusting the market and doing their own research. The return to verification would just offset the cost, just as the return to arbitrage perfectly offsets its cost in the Grossman and Stiglitz model.

Thus, there are two types of equilibria, one where prices convey information and one where they do not.[2] While both are theoretically valid, I can use the experimental data to empirically test if subjects endogenously reach informative pricing or if they behave closer to the pooling equilibrium.

[1]Note that this pooling equilibrium still satisfies the Cho-Kreps intuitive criterion because both high and low-type sellers would benefit from deviating to the high price if the buyer were to buy.

[2]There is no equilibrium when μ_h is at the threshold where the buyer is indifferent between leaving and getting information: in this case, only high-quality products would be sold at high-prices, so the buyer would regret not buying them blindly. Nevertheless, there is not an even number of equilibria, since there is a continuum of pooling equilibria (one for each value of buyer beliefs).

3.2 Comparative Statics

As well as examining equilibrium selection, I vary the cost of verification, c, to see how closely subjects track theoretical predictions. In the theory, as c decreases and information is cheaper, prices become more and more informative until the price perfectly reflects quality when $c = 0$. As c increases and information becomes more expensive, prices become less informative until the informative equilibrium disappears.

However, this depends on sharp comparative statics characteristic of mixed Nash equilibria. As c increases (decreases), the low-type seller should try to fool the buyer more (less) to keep the buyer indifferent between buying blindly and getting information. The buyer should get information at the same rate regardless of the cost of information, since their behavior needs to keep the low-type seller indifferent in equilibrium.

Figure 3 shows this result. This figure is just a slice of the actual action space: the slice where the buyer always buys when observing the low price, always either buys or verifies when observing the high price, and where the high-quality seller always sets the high price. These are the choices that should occur under the partially separating Nash equilibrium, and players usually follow these actions. This slice represents, essentially, a game of chicken (or cat-and-mouse) between the buyer and the low-quality seller. The buyer wants to buy given p_h if the low-quality seller sets p_l, and wants to verify given p_h if the low-quality seller sets p_h. The low-quality seller wants the opposite: they want to set p_l if the buyer is verifying and set p_h when the buyer is buying blindly.

But mixed Nash equilibria can be difficult for subjects to understand, and can sometimes require very astute subjects to achieve. Insofar as these Nash comparative statics are counterintuitive to subjects, empirical price informativeness may differ from theoretical predictions.

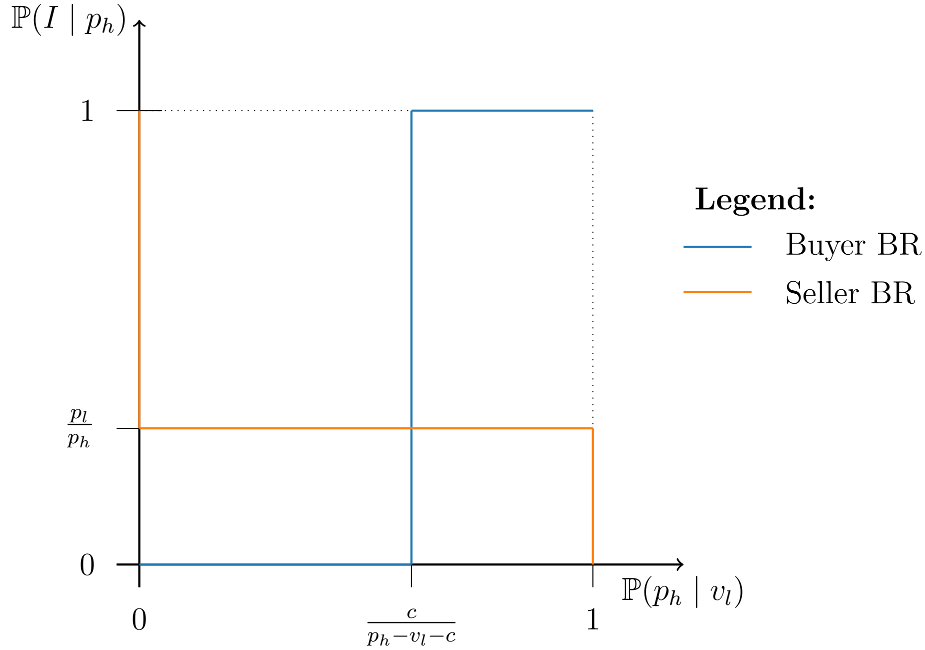

Figure 3: Cheating and Verifying is a Game of Chicken

Notes: This is the slice of the strategy space where the low-quality seller decides whether to "cheat" the buyer by setting the high price, and the buyer decided whether to verify the quality of a high-priced product. This is essentially a game of chicken, or cat-and-mouse. If the low-quality seller is frequently trying to cheat the buyer, the buyer should verify. But if the low-quality seller is being honest, the buyer should trust them and not exert effort to verify the quality. Conversely, if the buyer is trusting, the low-quality seller should cheat them, and if they are verifying, the low-quality seller should be honest. The only Nash equilibrium is in mixed strategies: the seller cheats just enough that the buyer is indifferent to getting information, and the buyer verifies just enough that the seller is indifferent. This equilibrium is unstable in the sense that a small deviation from one player should lead to a larger deviation in response, until players are far from the Nash equilibrium.

4 Experimental Design

Experiments were run in-person at the University of Virginia, with a sample of 74 undergraduate students. There were six sessions of 10-14 subjects each. In each session, subjects were randomly chosen to be a buyer or a seller, and this designation persisted throughout the session. Subjects then played the game for 16 rounds with one information cost, and another 16 rounds with a different information cost. At the beginning of each round, buyers were randomly rematched with sellers.[3]

[3] I ran fewer treatments than I at first expected, and failed to balance the order of treatments. In four sessions, agents started with low verification cost before proceeding to high verification cost, and subjects played the treatments in the opposite order in only two sessions. This means that differences in treatment response could correspond slightly with differences in treatment order.

For the experiment I chose the following parameters. The value of the product to the buyers was

$$v = \begin{cases} v_h = \$2.60 \text{ w.p. } 1/2 \\ v_l = \$1.20 \text{ w.p. } 1/2 \end{cases}$$

And the sellers could set price

$$p \in \{p_l = \$0.60,\ p_h = \$2.00\}$$

In the low-cost treatment, the cost of information was $c = \$0.10$, and in the high-cost treatment it was $c = \$0.30$. Subjects were paid for every choice in the experiment, but sellers were only paid half of the face value of their earnings so that average earnings were similar between buyers and sellers. Subjects were informed that this would be the case at the beginning of the experiment.

5 Results: Equilibrium Selection

There are two main possibilities for what could happen theoretically. On the one hand, buyers could refuse to buy any product at the high price. This would force sellers to always set the low price. Consumers would win out, but prices would convey no information about quality.

On the other hand, buyers may believe that expensive products are generally high-quality. If this were true, it could make it worthwhile to buy or at least check the quality of a high-price product. In this case, buyers' beliefs can become self-fulfilling. High-quality firms will know they can sell at a high price, while low-quality firms will not be so sure, and will mix between the two prices. Since high-quality firms always set the high price while low-quality firms mix, buyers turn out to be correct that prices convey information.

In the data, high-quality sellers are convinced that they can sell at a high price and they are correct. Buyers almost never reject an expensive product without at least verifying the quality first. As a result, high-quality sellers have a greater incentive to set the high price than low-quality sellers. The difference in how likely it is for high-quality and low-quality sellers to set the high price means that prices convey information to buyers.

Figure 4 shows how likely each type of seller is to set the high price in the various Nash equilibria as well as in the data. The data overwhelmingly select the informative equilibrium discussed by Grossman and Stiglitz. Not only do buyers expect prices to convey information, but this expectation is self-fulfilling, and prices actually do convey information endogenously.

The diagonal line on figure 4 represents strategies where a firm will choose the same pricing decision whether it has high- or low-quality products. If both types choose the same prices, buyers cannot learn anything about quality from observing prices, so this diagonal also represents seller strategies that are *uninformative*–conveying no information to buyers about quality.

The data is firmly above this diagonal. High-quality products are consistently priced higher, on average, than low-quality products. Buyers understand this and generally buy or get information when facing a high price. If prices conveyed no information, buyers would simply walk away until the firm offers a low price. So prices provide meaningful information, inducing buyers to make choices they would not have made under their prior beliefs.

6 Results: Behavioral Deviation from Theory

Informational efficiency, (insofar as it is possible in the Grossman-Stiglitz paradox), is a lot to ask of buyers and sellers. The rational buyer must sometimes buy a product blindly, without checking the quality. They need to be quite precise themselves, and have perfect confidence in the precision of their fellow buyers, whose choices determine how trustworthy the market prices will be.

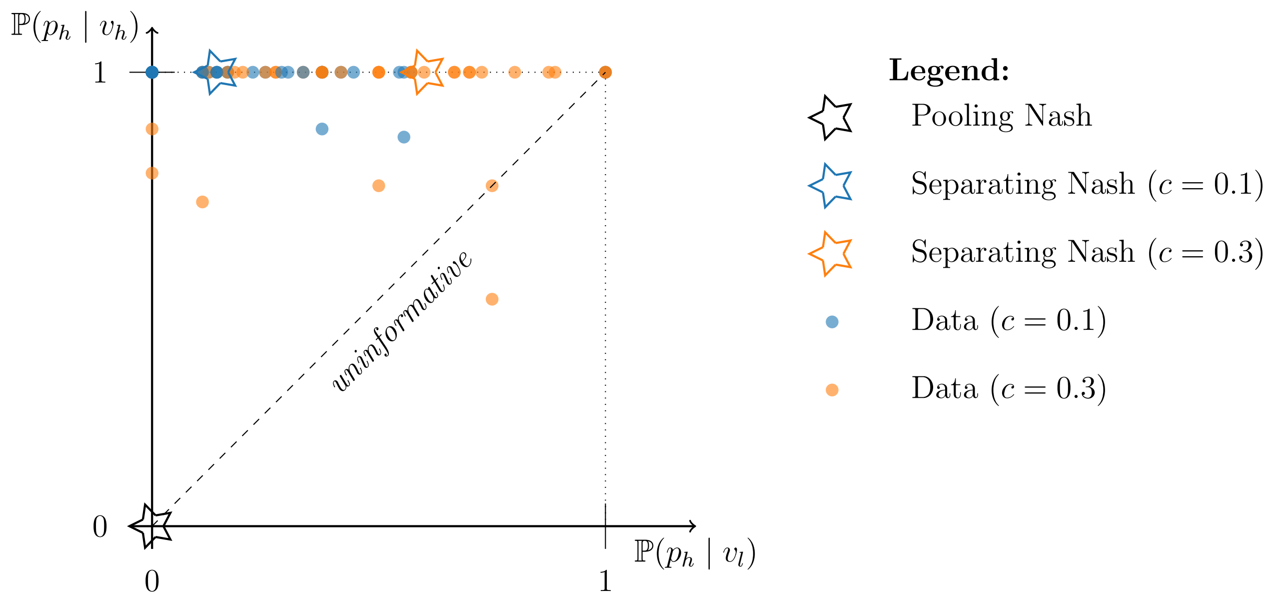

Figure 4: Pooling Nash is *not* Selected

Notes: Data is by-subject. Contrary to the pooling Nash equilibrium, almost all subjects are more likely to set the high price when they have a high-quality product, demonstrating that information is conveyed through the arbitrage channel in Grossman-Stiglitz (1980).

It is as if a \$20 bill is left on the sidewalk. The well-trained economist just walks by, because if it were a *real* \$20 bill, someone else would already have taken it. Assuming the others walking by are perfectly mixing between sometimes checking and sometimes ignoring the potential \$20, the economist is justified in being indifferent to checking or just walking by. [4]

But this rigid calculus may not be intuitive for actual consumers. If their faith in others' rationality wavers, and the cost of bending down to check the \$20 bill is low, they may verify for themselves more often than is strictly necessary. Or, if checking the quality is quite difficult, they may place too much weight on their fellows and fail to check even when they should. These behavioral tendencies could skew market prices away from the efficient markets prediction.

These behavioral tendencies can be seen from the choice probabilities in figure 5. When the cost of verification is low ($c = \$0.10$), buyers verify more than the Nash equilibrium

[4] I heard this example from Maxim Engers; I've since learned its origin is older, but could not find a source for it.

	$c = 0.1$		$c = 0.3$	
	Data	Mixed Nash	Data	Mixed Nash
$\mathbb{P}(p_h \mid v_h)$	0.99	1.00	0.96	1.00
$\mathbb{P}(p_h \mid v_l)$	0.23	0.14	0.46	0.60
$\mathbb{P}(B \mid p_h)$	0.20	0.30	0.32	0.30
$\mathbb{P}(I \mid p_h)$	0.78	0.70	0.58	0.70
$\mathbb{P}(B \mid p_l)$	0.94	1.00	0.90	1.00
$\mathbb{P}(I \mid p_l)$	0.04	0.00	0.05	0.00

Figure 5: Mixed Nash Generally Fits the Data

Notes: Generally, data aligns with the Grossman-Stiglitz Paradox. The biggest deviations from theory are: (1) Buyers get less information when it is more expensive, even though the Nash predicts no change in verification, and (2) low-quality sellers respond less strongly to the change in verification cost than they should according to the Nash theory. Both these facts are explored below, and predicted by quantal response.

predicts, not placing enough trust in the market. When the cost of verification is high ($c = \$0.30$), buyers verify less than the Nash equilibrium predicts, placing too much faith in the market. [5]

In the theory, buyers verify just as much when verification is expensive as when it is cheap. This is not intuitive for buyers in the lab. It happens in the theory because sellers respond steeply to the change in verification cost, even though this cost does not affect the seller payoff directly. When the verification cost increases, the low-quality seller tries to cheat the buyer more, so the buyer has a greater incentive to verify. This increased incentive to verify perfectly offsets the increased cost of verification, so that the buyer continues to verify at the same rate as before.[6]

Figure 6 shows the best responses of a low-quality seller and a buyer, for both high

[5] This should not be confused with buyers getting less, or more, information than is optimal given empirical seller choices. Since mixed Nash equilibria are unstable, if sellers deviate slightly from the equilibrium choice probabilities, buyers who optimally respond to those deviations will move *further* from Nash behavior, not closer to it. Here, I examine differences between actual behavior and equilibrium behavior, but it is important to remember that this is different from (in fact, opposite to) empirical optimality of subjects given opponents' out-of-equilibrium behavior when mixed Nash equilibria are involved.

[6] This is the standard intuition of a mixed Nash equilibrium. It is similar to how, in a game of rock-paper-scissors, if a third party offers an additional incentive to one player to play scissors (regardless of whether they win or lose), that player will not adjust their choice probabilities at all. Instead, their *opponent* will respond by playing rock more often. This increases the cost of playing scissors so that it is once again equal to the benefit of playing scissors, and the player remains indifferent.

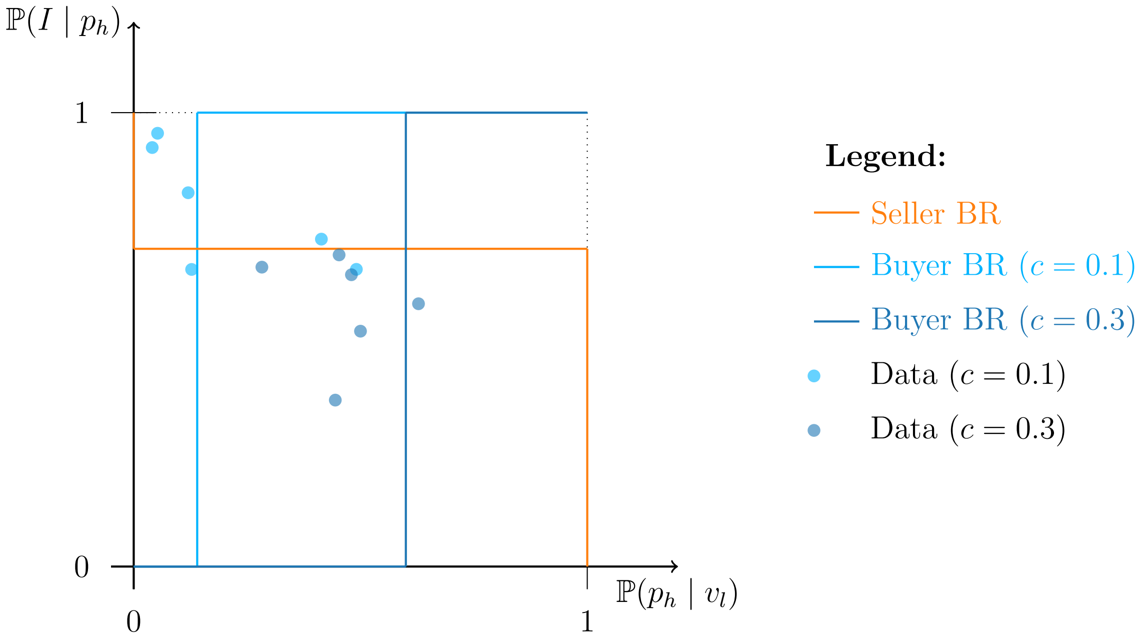

Figure 6: Subject Responses are Smoother than Theory

Notes: While data is fairly close to theoretical predictions, it is clear that the data for $c = 0.1$ (cyan) is higher than the data for $c = 0.3$ (blue), indicating buyers get less information when it is more expensive. Less obvious, but true, is that the data for the two treatments is closer together horizontally than the two Nash equilibria, indicating that sellers do not respond as strongly to the change in information cost as theory predicts.

and low verification cost, along with session-level data.[7] While the perfectly vertical and horizontal best-response lines indicate that changes in information cost will lead to large differences in seller behavior but no change in buyer behavior, actual data shows that agents behave more smoothly. If the agent response functions were smoother, as would be the case if their were some noise in the agents' decisionmaking, we could rationalize (1) that buyers respond to an increase in information cost by getting information less frequently, and (2) that sellers do not respond as drastically to the information cost as they would in the Nash equilibrium. Both of these facts are predicted by quantal response.

In a quantal response equilibrium, agents make small, mean-zero errors when evaluating which actions give them the highest expected payoffs. When two actions yield similar expected payoffs, agents play them with similar likelihood. When one action yields a much

[7] Again, since this is just a slice of the action space, other choice probabilities are fixed at their levels in the mixed Nash equilibrium.

higher payoff than another, agents will play the better action much more frequently. Agents understand that these errors occur (in themselves and others) and respond accordingly, leading to an equilibrium where beliefs remain consistent, although their behavior sometimes deviates from their best responses.

Thus, if agents' expected payoffs are smooth functions of their opponents' choice probabilities, their own choice probabilities will also be smooth functions of their opponents behavior. Low-quality sellers will not suddenly decide to cheat buyers once buyers are verifying less than a certain threshold. And buyers will not suddenly decide to always verify when sellers begin cheating them more than a threshold amount.

When the cost of verification increases, buyers will respond gradually by buying less information. In turn, sellers will respond to buyers by cheating them slightly more often. Contrary to Nash, buyers will respond to changes in information cost, while sellers respond to those changes less drastically. This can be seen from quantal response curves plotted in figure 7. While quantal response does not always predict the levels correctly (in particular, sellers in the data are more hesitant to cheat buyers than quantal response predicts), it does predict the dynamics of how choice probabilities change as the cost of verification changes, and how those dynamics differ from Nash theory.

The extent to which sellers are trying to cheat buyers determines the informativeness of prices. If low-quality sellers are constantly trying to price similarly to high-quality sellers, then prices will not convey much information about quality. Since behavioral noise smooths the relationship between seller choice probabilities and information cost, it ultimately smoothes the relationship between information cost and the informational content of prices. The Grossman-Stiglitz paradox asserts that prices will be more informative when arbitrage is cheap and less informative when arbitrage is expensive, so that the return to arbitrage always perfectly offsets its cost. That relationship still exists but is dampened by behavioral noise, as seen in figure 8.

While quantal response can provide some intuition for the dynamics, we can test behav-

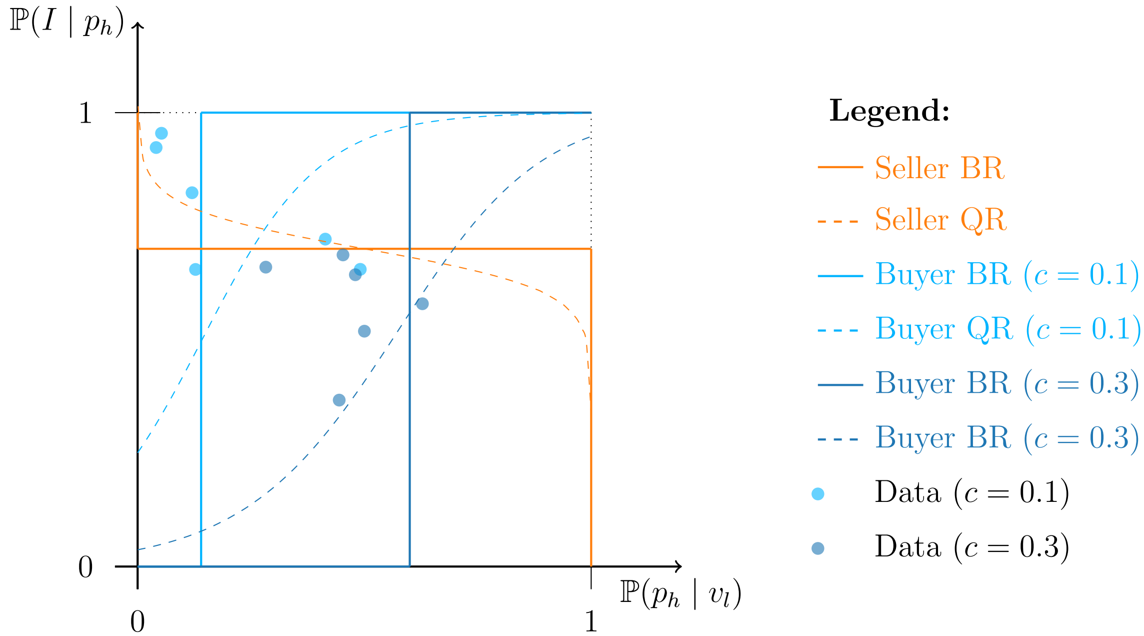

Figure 7: QRE Explains Differences in Verification and Cheating

Notes: This is the best-fit QRE (precision = 22). The QRE is fit just to this slice of the strategy space. The other decisions (seller's pricing decision when high-quality, and buyers' decision when facing a low price, etc.) are simple enough that agents are quite close to Nash. Quantal response does not predict the level of cheating well; sellers cheat less than predicted (perhaps due to an aversion to lying). However, quantal response does predict the differences across the two treatments: buyers respond more to changes in information cost and sellers respond less to those changes than Nash suggests.

ioral tendencies using a permutation test. Under the null hypothesis, buyers follow the Nash and do not adjust their behavior as the cost of verification changes. Since the verification cost should not matter (under the null hypothesis), permuting the labels of the data–which data points have come from which treatment–should not matter. If, instead, the observed data involves treatment differences more severe than permuted versions of the data, we can conclude that verification cost does affect buyers' choice probabilities. In the same way, we can test whether low-quality sellers do indeed change their behavior less than they should in the theory, and whether informativeness therefore adjusts less. The results are given in figure 9.

We can conclude that behavioral noise leads buyers to respond too much to the cost of verification, and leads sellers to respond too little. Thus, while prices are more informative when information is cheap, the effect appears to be smaller than Nash. When information

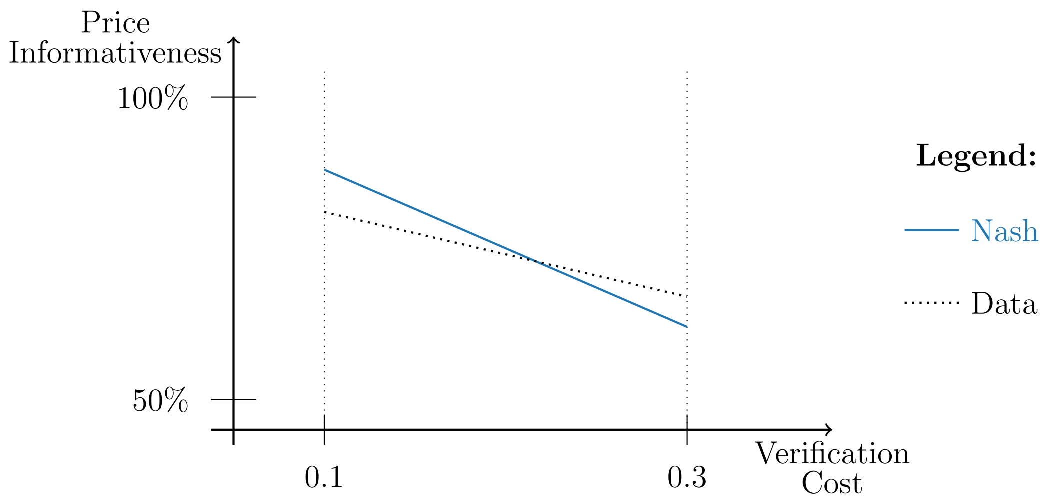

Figure 8: Cost and Benefit of Verification

Notes: Informativeness is measured as the unconditional probability that the quality can be guessed from the price. So 50% is completely uninformative: no better than random guessing. 100% means the quality is known for certain after the price is observed. These lines are just interpolated from the identified endpoints.

Test	p-value
Verification decreases with cost	$2/64 = 0.031$
Cheating increases less than Nash	$3/64 = 0.047$
Informativeness decreases with cost	$1/64 = 0.016$
Informativeness decreases less than Nash	$6/64 = 0.094$

Figure 9: Non-parametric Permutation Tests

Notes: With six paired observations, there are 64 possible permutations of the data. The p-value is the fraction of these permutations that produce a test statistic at least as extreme as the one observed in the actual (unpermuted) data. In each case, the test statistic is the average difference between the two treatments.

is cheap, agents purchase more information than the Nash predicts, and when information is expensive, agents buy less information than Nash predicts.

7 Conclusion

Overall, in the lab, markets are remarkably efficient in terms of how much information prices convey, even with a relatively small number of untrained subjects. Agents' choice probabilities are very close to those predicted by Grossman and Stiglitz. This demonstrates that agents do not simply *believe* that prices convey information in real-world markets and take those beliefs into the lab, but that prices actually *do* convey information endogenously in the lab. Specifically, prices can convey information via the Grossman-Stiglitz channel of costly verification.

Nevertheless, there are still some behavioral deviations from theory, when subjects trust the market too much or too little. These are intuitive because they would arise if subjects had smooth best-response functions rather than the infinitely steep best-response functions of a rational agent. When verification is easy, subjects verify more than Nash, failing to realize that everyone else can also verify more easily and the market adjusts. When verification is hard, subjects verify less often than Nash, failing to realize that verification is also difficult for everyone else, and they should trust the market prices less.

References

Bagwell, Kyle and Michael H Riordan (1991), "High and declining prices signal product quality." *The American Economic Review*, 224–239.

Berry, Steven T (1994), "Estimating discrete-choice models of product differentiation." *The RAND Journal of Economics*, 242–262.

Bester, Helmut and Klaus Ritzberger (2001), "Strategic pricing, signalling, and costly information acquisition." *International Journal of Industrial Organization*, 19, 1347–1361.

Cornell, W Bradford (1978), "Price as a quality signal: Some additional experimental results." *Economic Inquiry*, 16, 302.

Fama, Eugene F (1970), "Efficient capital markets." *Journal of finance*, 25, 383–417.

Grossman, Sanford J and Joseph E Stiglitz (1980), "On the impossibility of informationally efficient markets." *The American economic review*, 70, 393–408.

Heinsalu, Sander (2021), "Competitive pricing despite search costs when lower price signals quality." *Economic Theory*, 71, 317–339.

Mastrobuoni, Giovanni, Franco Peracchi, and Aleksey Tetenov (2014), "Price as a signal of product quality: Some experimental evidence." *Journal of Wine Economics*, 9, 135–152.

McConnell, J. Douglas (1968), "An experimental examination of the price-quality relationship." *The Journal of Business*, 41, 439–444.

McKelvey, Richard D and Thomas R Palfrey (1995), "Quantal response equilibria for normal form games." *Games and economic behavior*, 10, 6–38.

McKelvey, Richard D and Thomas R Palfrey (1998), "Quantal response equilibria for extensive form games." *Experimental economics*, 1, 9–41.

Peterson, Robert A (1970), "The price-perceived quality relationship: Experimental evidence." *Journal of Marketing Research*, 7, 525–528.

Tirole, Jean (1988), *The theory of industrial organization*. MIT press.

Venkataraman, VK (1981), "The price-quality relationship in an experimental setting." *Journal of Advertising Research*.

Appendices

A Full extensive form

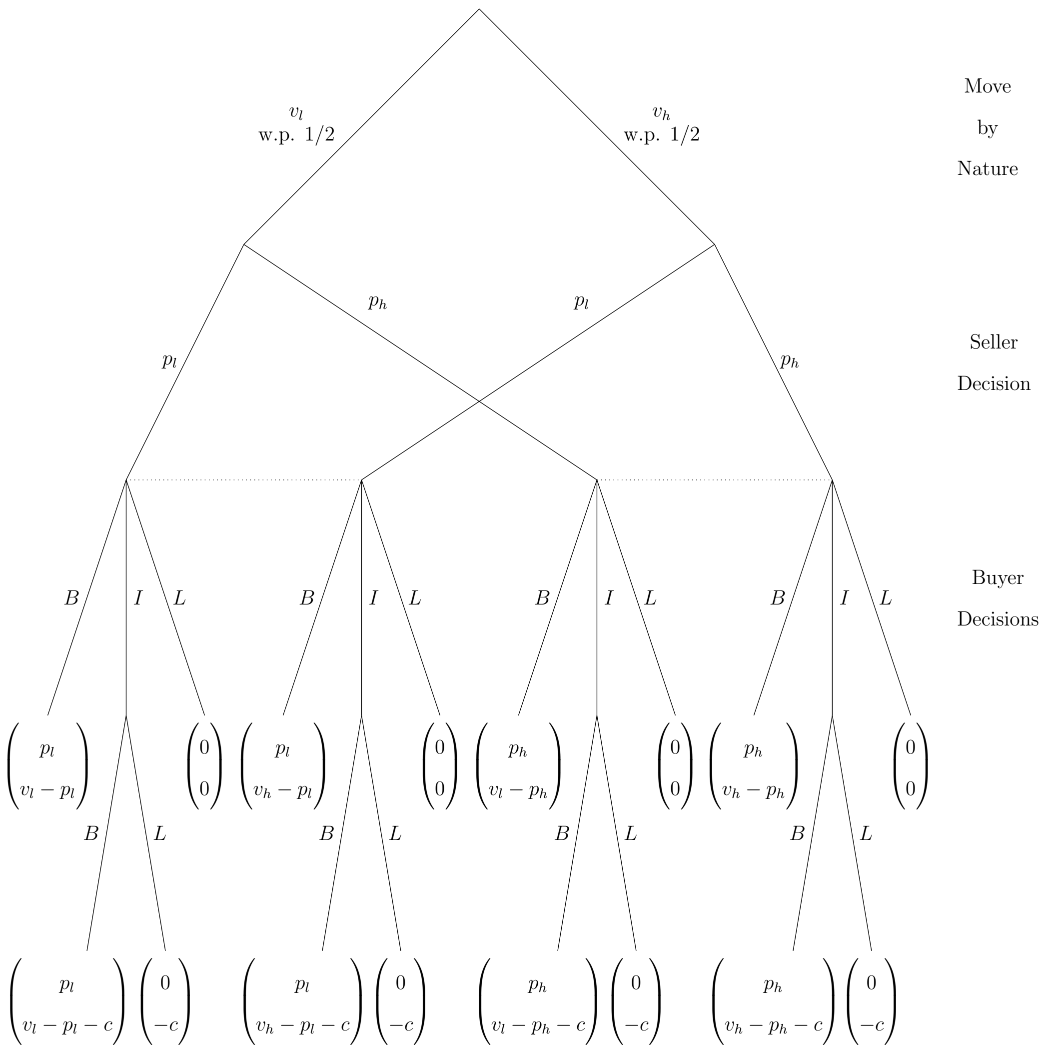

Figure 10: Extensive Form of the Game

B Derivation of equilibria

B.1 Pure strategy equilibrium

There is no perfectly separating equilibrium in this game. If the seller were to choose the high price when quality is high and the low price when quality is low, consistent beliefs imply that the buyer would know in equilibrium that a high price implies high quality. Thus, the buyer would no longer need to pay the information cost because they can learn quality costlessly from the price. But if the buyer is not paying the information cost, there is an incentive for the seller to charge the high price even when the quality is low, since consumers will expect quality to be high and will be fooled into buying the product. Thus, there is no separating equilibrium where the seller prices according to their quality level. Consistent beliefs also rule out reverse type dependence; it is not an equilibrium for the high quality seller to set the low price and the low quality seller to set a high price.

It could be that the seller will set the same price whether their quality is high or low. In this case, consumers will buy only if the expected quality is greater than the price. I have chosen the possible prices so that

$$p_l < v_l < \frac{1}{2}(v_h + v_l) < p_h \tag{1}$$

and this ensures that buyers will not buy if the seller is always setting the high price, but will buy if the seller sets the low price. So there is no equilibrium where the seller always sets the high price, but there is an equilibrium where the seller always sets the low price.

For this to be an equilibrium, the seller must not prefer to deviate to the high price, and thus we need the buyer to choose not to get information and not to buy if they were to observe the high price. This can be ensured by buyer beliefs. Let μ_h denote the probability that the seller is high quality given that they set the high price. Since the high price is off the equilibrium path, μ_h is unconstrained. That is, if participants reach an equilibrium where the high price is never set, buyers can reasonably assume anything about the expected quality a firm would have if they were to set the high price. If the payoff of not getting information and not buying is higher than both getting information and buying without information, we have:

$$\mathbb{E}[u(L)|p_h] \geq \mathbb{E}[u(B)|p_h] \qquad \mathbb{E}[u(L)|p_h] \geq \mathbb{E}[u(I)|p_h] \tag{2}$$

$$0 \geq \mu_h(v_h - p_h) + (1 - \mu_h)(v_l - p_h) \qquad 0 \geq \mu_h(v_h - p_h) - c \tag{3}$$

For this to hold, the buyer must believe that, conditional on setting the high price, the seller is unlikely enough to be high quality that expected quality does not exceed the price, and that the expected benefit of information does not exceed the information cost.

$$\mu_h \leq \min\left\{\frac{p_h - v_l}{v_h - v_l}, \frac{c}{v_h - p_h}\right\} \tag{4}$$

As long as the buyer believes a deviation to the high price is sufficiently likely to occur when the seller is low-quality, pooling at the low price can be sustained as an equilibrium.

B.2 Mixed Equilibrium

In this equilibrium, the seller always chooses the high price when they have high quality. When the seller has low quality, they mix between the two prices, choosing the high price with probability $\mathbb{P}(p_h|v_l)$. Thus, if the buyer sees the high price, they believe the seller is high-quality with probability $\mu_h = 1/(1 + \mathbb{P}(p_h|v_l))$. In order for the seller to mix when they have low quality, it must be that the expected profit from setting the low price is equal to the expected profit from setting the high price. This implies that when the buyer observes the high price, they only buy without getting information a fraction of the time:

$$\pi_h(p_h) = \pi_l(p_h) \implies p_l = \mathbb{P}(B|p_h)p_h \implies \mathbb{P}(B|p_h) = \frac{p_l}{p_h} \tag{5}$$

and thus the buyer chooses to get information with corresponding probability $\mathbb{P}(I|p_h) = \frac{p_h - p_l}{p_h}$. Because the buyer's choice probabilities must make the seller indifferent in the mixed Nash, the buyer's probability of getting information is independent of the cost of information. Similarly, the seller's choice probabilities must make the buyer indifferent. For the buyer to want to mix between buying without information and getting information when observing the high price, it must be that their expected payoffs are the same:

$$\mathbb{E}[u(B)|p_h] = \mathbb{E}[u(I)|p_h] \implies \mu_h v_h + (1-\mu_h)v_l - p_h = \mu_h(v_h - p_h) - c \tag{6}$$

and this is true when the seller's choice probability is

$$\mathbb{P}(p_h|v_l) = \frac{c}{p_h - v_l - c} \tag{7}$$

If the seller has low quality, they must set the high price more often when the effort cost rises so that the buyer will have a stronger incentive to get information and remain indifferent between getting information and buying without information. These comparative statics are typical of mixed equilibria: if information becomes more expensive, the buyer (who pays the higher information cost) still buys information at exactly the same frequency, and the seller (who does not pay the cost) changes their behavior.

In order for this equilibrium to exist, it must be that the payoffs of buying without information or getting information are no lower than the payoff of not getting information and not buying. This is true when the effort cost is sufficiently low:

$$\frac{c}{p_h - v_l - c} \leq \frac{v_h - p_h}{p_h - v_l} \implies c \leq \frac{(v_h - p_h)(p_h - v_l)}{v_h - v_l} \tag{8}$$