Strategic Pricing with Rational Inattention to Quality*

Daniel Martin†

Job Market Paper

Preliminary Draft

October 12, 2012

Abstract

In a standard sequential pricing game, I model how buyers attend to information about product quality when prices are easily observable using "rational inattention" theory (Sims [2003]). Rational inattention in this setting produces two possible equilibria: one where the entire surplus goes to buyers and one where most of the surplus goes to sellers of high quality goods. I show that there are large differences in how policy interventions related to attention impact buyer welfare in each equilibrium. In a laboratory experiment where buyers face real attentional costs, I find that subjects play very close to predictions of the equilibrium where most of the surplus goes to sellers of high quality goods. Moreover, decision times reflect the adjustments in buyer attention that are predicted by the model.

1 Introduction

Buyers are often confronted with a large amount of freely available information about product quality from direct observation, pictures, product specifications, customer reviews, and advertisements. Although costless to acquire, this information can require a substantial amount of attentional effort to process, so buyers may not fully attend to it.

In this paper, I consider the market implications of buyer inattention to freely available information about product quality. For instance, are the sellers of high quality products forced to sell at lower prices

*A special thanks to Andrew Caplin and Mark Dean for many years of support, encouragement, and guidance. Thanks also to Heski Bar-Isaac, David Cesarini, Ignacio Esponda, Guillaume Fréchette, Nic Kozeniauskas, Muriel Niederle, David Pearce, Andrew Schotter, Matan Tsur, Laura Veldkamp, Emanuel Vespa, Charles Wilson, and Jidong Zhou for valuable suggestions. I gratefully acknowledge financial support from the Center for Experimental Social Science (CESS).

†Center for Experimental Social Science and Department of Economics, New York University.

because buyers do not recognize their products as high quality? Can sellers with low quality products trick uninformed buyers into paying higher prices? Should policymakers try to make information less costly to process or otherwise encourage attentional effort?

The market setting I examine is a standard sequential pricing game in which the seller is fully informed about product quality and sets a price that is easily observable to the buyer. The buyer starts out largely uninformed about product quality, which is appropriate for durable goods that are rarely purchased and change features over time, such as computers, cell phones, and air conditioners. I assume that price and brand are internalized immediately, while other freely available information about quality requires attentional effort to process.

I model the attention of buyers abstractly using the "rational inattention" approach introduced by Sims [2003], which is well suited for this application because it places information processing constraints on agents that can be viewed as cognitive limitations on internalizing freely available information.[1]

With rational attention to information about quality, there are two possible equilibria of this game. In one, buyers acquire market power through strategic ignorance, i.e., they never attend to information about quality and hold fully pessimistic beliefs about sellers that charge a high price, so all sellers must charge a low price ("pooling low" equilibrium). As a result, buyers capture the entire surplus, policy interventions that force attention can backfire and policy interventions that lower the costs of attention have no effect.

In the other equilibrium, sellers have market power. Sellers of high quality products charge a high price and sellers of low quality products sometimes mimic them ("mimic high" equilibrium). At a high price, the buyer must undertake attentional effort to distinguish high and low quality products. As a result, most of the surplus goes to sellers of high quality goods, and there is some deadweight loss. However, the benefits of policy interventions that lower the cost of attention depend on market conditions. As information processing costs decrease, low quality sellers are typically less likely to mimic, but in some cases they are *more* likely to mimic.

Do the predictions of either equilibrium match actual behavior? To address this question, I implement an instance of this model in a laboratory experiment. To generate real attentional costs, I supply subjects with a free source of information about quality that takes cognitive effort to process. I find that subjects play very close to the predictions of the mimic high equilibrium. Focusing on the last half of rounds, high quality sellers charge a high price around 99% of the time, and low quality sellers charge a high price around 20% of the time. Pricing high is a best response for high quality sellers, and mixing is a best response for low quality sellers. On the other side of the market, buyer demands are very close to those predicted by rational inattention, and buyer consideration times suggest changes in attention that are consistent with rational inattention.

[1] See section 2 for a review of the rational inattention literature.

This paper makes three main contributions to the literature. First, to my knowledge, it contains the first model of strategic pricing in which buyers can easily observe prices, but are rationally inattentive to quality.[2] Second, it provides the first experimental evidence of rational inattention to freely available information in a market setting.[3] Third, it adds to a growing literature on the market implications of limited attention to freely available information about choice alternatives.

At the same time, this paper follows a large literature on asymmetric information in market settings. Starting with Akerlof [1970], economists have built models of asymmetric information in which sellers are fully informed about quality, but buyers are not. While Akerlof [1970] assumed that buyers could not acquire any information about quality, subsequent models have included many different assumptions for how buyers can gather information about quality.[4] Unlike commonly used assumptions for information acquisition, rationally inattention allows agents to control both the amount of information that they process and the type of information that they process. This flexibility enables rational inattention to capture a wide range of behavior, which is central to the analyses described above.[5]

In section 2, I introduce rational inattention theory and provide a review of other related literature. In section 3, I describe the model and provide the equilibrium concept. In section 4, I solve for buyer demands and information gathering as a function of market parameters and seller strategies. In section 5, I indicate the possible equilibria of the model and characterize these equilibria in closed form. In section 6, I present the experimental design and results. Section 7 concludes.

2 Literature Review

2.1 Rational Inattention

Sims [2003] introduced rational inattention theory to model the constraints that agents face in processing available information.[6] It is based on classic works in the information theory literature which describe a

[2] In other models with rational inattentive buyers, prices are not easily observable (Matějka [2010] and Matějka and McKay [2012]).

[3] This paper joins Gabaix, Laibson, Moloche, and Weinberg [2006], Caplin, Dean, and Martin [2011], Caplin and Martin [2011], and Caplin and Martin [2012] in modeling and measuring limited attention to freely available information.

[4] See section 2 for a brief description of the assumptions found in Chan and Leland [1982], Wolinsky [1983], Cooper and Ross [1984], Bagwell and Riordan [1991], Bester and Ritzberger [2001], Kalaycı and Potters [2011], Voorneveld and Weibull [2011], and Bar-Isaac, Caruana, Cuñat [2012].

[5] For example, if rational inattention is replaced in the model with being fully informed at a cost, then the mimic high equilibrium is not stable at higher costs, as showed by Bester and Ritzberger [2001]. If rational inattention with replaced with a free normally distributed signal of quality, then the model no longer predicts fully inattentive behavior, so there is no pooling low equilibrium.

[6] For a brief overview see Wiederholt [2010], for the connection to information theory see Sims [2010], and for a more detailed treatment see Veldkamp [2011].

physical constraint on the flow of information. This constraint, called the Shannon capacity or Shannon channel, determines the amount of uncertainty (entropy) that can be reduced by a message, and it has been interpreted as a cognitive limitation for economic agents (Wiederholt [2010], Tutino [2011]). In models of choice, this constraint produces a noisy perception of the underlying state. Woodford [2012] produces a informational constraint with neurobiological foundations, and Gabaix [2011] provides a related approach to limited attention in which agents simplify the data available to them.

2.1.1 Solving Models of Rational Inattention

Rational inattention can be modeled as agents choosing the distribution from which they draw an informative signal, with more informative signals being more costly. This cost takes a specific log-linear form (see Veldkamp [2011]). For tractability, many models have assumed a Gaussian relationship in the signal structure and a linear-quadratic utility function, but recent work, including this paper, allows for more general signal structures and classic utility functions (see Sims [2006] for a discussion and Yang [2012b] for an implication). Recently, Matějka and McKay [2011] and Caplin and Dean [2012] have shown that in finite models, as in this paper, rational inattention can yield clean solutions.

2.1.2 Applications of Rational Inattention

Rational inattention theory has been applied widely in macroeconomics and finance, including models of price setting, savings, portfolio choice, and contingent planning.[7] Existing papers that also model buyers as rationally attentive (Matějka [2010], Matějka and McKay [2012]) assume that buyers are rationally inattentive to prices or both price and quality. This paper differs in that I assume prices are easily observable.

Because prices are easily observable to buyers, price setting becomes a game. This paper joins Yang [2012a,b] as one of the first applications of rational inattention to games. Like Yang [2012a], I use a binary action setup with sequential moves and include only one rationally inattentive agent (the second mover). However, a substantial difference in this paper is that the first mover is fully informed, so that their action choice can reveal information.

2.1.3 Testing for Rational Inattention

This paper joins a small literature on testing models of rational inattention. Klenow and Willis [2007] find reactions to old news in price changes, which is suggestive of rational inattention in price setting. Maćkowiak, Moench, and Wiederholt [2009] look at aggregate and sector-specific shocks and find evidence supporting rational inattention in price setting. Kacperczyk, Nieuwerburgh, and Veldkamp [2012] examine

[7] See Wiederholt [2010] for a review and Maćkowiak and Wiederholt [2012] for a new application and additional citations.

the allocations of mutual funds as they vary with the business cycle, and the covariance they find suggests that fund managers are focusing their limited attention on different information sources depending on the underlying state of the economy.

Two related experimental papers are by Cheremukhin, Popova, and Tutino [2012], who test for rational inattention in cognitive limits of processing information about choices over lotteries, and by Treviño and Szkup [2011], who allow subjects to improve the precision of their signal at a cost, which induces an information choice that can be interpreted as rational inattention. This paper differs from these existing experiments by looking for evidence of rational inattention to freely available information in a market setting.

2.2 Other Related Literature

2.2.1 Limited Attention and Market Implications

There is a large recent literature on limited attention to freely available information about choice alternatives, including papers by Manzini and Mariotti [2007], Salant and Rubinstein [2008], Rabin and Weizsacker [2009], Gennaioli and Shleifer [2010], Ergin and Sarver [2010], Caplin and Dean [2011], Compte and Postlewaite [2011], Gottlieb [2011], Manzini and Mariotti [2012], Masatlioglu, Nakajima, and Ozbay [2012], and Schwartzstein [2012].

In addition, there has been increasing interest in market implications of this behavior. For example, Eliaz and Spiegler [2011a,b] examine how firms try to influence the consideration sets of consumers with marketing or attention grabbing products. Two key differences here are that limited consideration is driven by prices and that consumers are optimal given their limited attention. Another example is by Bordalo, Gennaioli, and Shleifer [2012], who look at the impact of a salient product dimensions on consumer choices.

2.2.2 Costly Information Acquisition

There is a long literature on strategic pricing games where buyers acquiring information about quality. While not always motivated as limited attention to freely available information, information acquisition can operate much the same way technically. Cooper and Ross [1984] and Bagwell and Riordan [1991] assume that consumers can only be fully informed or fully uninformed. Chan and Leland [1982] and Bester and Ritzberger [2001] endogenize this form of information gathering. Voorneveld and Weibull [2011] consider buyers who get a normally distributed signal of quality. In Wolinsky [1983], consumers get a noisy signal of quality when they sample a price. Kalaycı and Potters [2011] model consumers as observing price perfectly, but quality differences with uniform noise. Bar-Isaac, Caruana, Cuñat [2012] assume that buyers can pay a cost to learn about a dimension of quality.

2.2.3 Experiments

The experiment that Kalaycı and Potters [2011] implement is similar to mine in that it requires the buyer to solve a math problem to learn quality, but differs because of their model setup: sellers determine the complexity level, buyers know little of the seller's problem, and buyers face extreme time pressure. As a result, there is little to no room for subjects to change attentional effort based on strategic considerations.

There have also been many experiments on the effects of exogenous price variation on the choices of real goods; however, these studies are rarely incentivized, and few look at the key role of information acquisition. One exception is by Lynch and Ariely [2000], who conduct an experiment that contains a treatment where information on prices for wines are easy to obtain, but information on quality levels are not, and find that price elasticity decreases with the difficulty of search for information about quality. Another is by Heffetz and Shayo [2009], which control the information that subjects have about an uncertain food product and find that exogenous variation in prices does not have a large effect on elasticities.

The experiment in this paper sheds light on how much prices can signal about quality in the presence asymmetric information about quality. In this way, it is related to experiments reported by Miller and Plott [1985], who examine how much prices can signal uncertain quality in an experiment where sellers can also signal quality by adding observable quality. My experiment differs in that subjects have access to an exogenous information source on quality and sellers can only signal through price.

3 A Model of Strategic Pricing

There is one seller and one buyer engaged in a once-off sales encounter. The seller has a single product with a quality level $\theta \in \Theta = \{\theta_L, \theta_H\}$, where $\theta_L, \theta_H \in \mathbb{R}_+$ and $\theta_L < \theta_H$. The product's quality level is determined exogenously, and the probability of high quality is $\lambda \in (0, 1)$. This can be interpreted as a product's easily observable brand. The seller, who knows the realized quality level, chooses a price $p \in P = \{p_L, p_H\}$, where $p_L, p_H \in \mathbb{R}_+$ and $p_L < p_H$, which is a take it or leave it offer. The buyer, who does not know the realized quality level, observes the price and then chooses an information processing technology π, which generates posteriors based on the realized quality level. Finally, after receiving a posterior from the chosen information processing technology, the buyer chooses to purchase a unit of the product or not, which is represented by a choice of $x \in \{0, 1\}$, where $x = 0$ is not purchasing and $x = 1$ is purchasing.

Notice that there is two sided information asymmetry in that the buyer does not know the realized quality level and the seller does not know the information obtained by the buyer.

I start with a simple model in order to help isolate the direct impact of rational inattention. In section 5, I briefly consider how the set of equilibria changes with an increase in the number of prices, number of quality levels, number of sellers, or number of buyers.

3.1 Product Quality and Information Asymmetries

As in a long literature in microeconomics, industrial organization, and marketing, I assume that product quality can be summarized with a scalar value.[8] This value has been treated both as an objective measure of quality (for example, Cooper and Ross [1984]) and a subjective measure of quality (for example, Judd and Riordan [1994]). I assume that the seller knows this value perfectly, which can be thought of as resulting from long experience with selling the product.

I also assume that the buyer does not know, ex ante, this value. As proposed by Wolinsky [1983] and others, this is a suitable assumption when thinking about infrequently purchased products, such as durable goods that have features which change over time. At the same time, I assume the buyer knows the probability that the product is of high quality, which can be interpreted as past experience with similar products or the product's brand or reputation for quality.

In addition, the buyer has access to many different sources of information about the quality of the product: physical inspection, information on provided on the packaging, customer reviews, advertisements, etc. For now, I will ignore the possible endogeneity that results from information being supplied by the seller, though this is a potentially interesting extension.

3.2 Information Processing

The way that a buyer processes freely available information is treated abstractly. The buyer is assumed to choose a mode of information processing after observing the price and forming an interim belief β_p that the product is of high quality. This information processing technology produces a range of possible posterior beliefs based on the true state of product quality. A posterior belief s is a point in the interval $S = [0, 1]$.

Technically, the buyer chooses an information processing technology π, which is a function that maps the realized quality level into $\Delta(S)$, the set of probability distributions over S that have finite support, so that

$$\pi : \Theta \to \Delta\left(S\right).$$

Let Π denote the set of all such functions, $\pi^{\theta}\left(s\right)$ be the probability of posterior $s \in S$ given quality $\theta \in \Theta$, and $S(\pi) \subset S$ denote the support of π. Note that the technology is stochastic. In some cases, such as when the technology is fully informative, this is not needed, but stochasticity allows the inclusion of information processing technologies with a random component.

[8] As in much of the theoretical literature, I assume that both quality and price are measurable in terms of expected utility.

With this structure, it is necessary to limit the set of feasible choices to all information processing technologies $\Pi_{\beta_p} \subset \Pi$ that generate correct posteriors for a given interim belief β that the product is of high quality, so that

$$\Pi_{\beta_p} = \left\{ \pi \in \Pi | \forall s \in S(\pi), \, s = \frac{\beta_p \pi^{\theta_H}(s)}{\beta_p \pi^{\theta_H}(s) + \left(1 - \beta_p\right) \pi^{\theta_L}(s)} \right\}.$$

In this model, processing information is costly, and those costs are determined using Shannon capacity (see section 2 for a review). In this approach, when a posterior reduces uncertainty more, it is more costly. Here uncertainty is measured by how far a posterior is from revealing that a product is of high or low quality. As s approaches 0 or 1, it reduces uncertainty more and more.

Formally, the posteriors produced by each information processing technology $\pi \in \Pi_\beta$ have a cost in expected utility units that is assigned by the function

$$K\left(\pi, \kappa, \beta_p\right) = \kappa \left(\left[\sum_{s \in S(\pi)} \pi(s) \left(s \ln(s) + (1-s) \ln(1-s)\right) \right] - \beta_p \ln\left(\beta_p\right) + \left(1 - \beta_p\right) \ln\left(1 - \beta_p\right) \right)$$

where $\kappa \in \mathbb{R}_+$. This form produces u-shaped costs, which bottom out at a posterior of 0.5 and increase symmetrically in either direction (towards high quality or towards low quality).

The lowest cost information processing technology is one that just produces one posterior and thus returns the interim belief as the posterior. The interim belief is included in the cost function so that the cost is normalized to zero in this case.

3.3 Payoffs to the Buyer and Seller

The buyer has the following "purchase" utility function, which is separable from any attentional costs:

$$U(\theta, p, x) = x(\theta - p) + (1 - x) u,$$

where $u \in \mathbb{R}_+$ is the utility of not purchasing the product, which can be interpreted as the outside option. This quasilinear utility function reflects a buyer who is balancing the quality and price of a good and is suitable when prices do not have a big effect on wealth (see Vives [2001]). The buyer is a risk neutral expected utility maximizer.

To focus the analysis on situations in which there are sometimes gains to trade, but not always gains to trade, I make the following assumptions:

1. When a high quality product has a high price or a low quality product has a low price, the buyer prefers to purchase the product: $\theta_H - p_H > u$ and $\theta_L - p_L > u$.

2. When priced high, the buyer prefers not to purchase the low quality product: $\theta_L - p_H < u$.

The seller has the following profit function:

$$V(p, x) = xp,$$

and is also a risk neutral expected profit maximizer. This profit function reflects a seller who does not have reputational concerns or marginal costs.

3.4 Defining the Market and Game

A market ω is defined as

$$\omega = (\lambda, \Theta, P, u).$$

All of the following could change a market: new technologies could change the probability a product is high quality level or improve the absolute quality levels; subsidies or taxes could change the utility or profit functions or the value of the outside option; and price floors or ceiling could alter the set of prices to choose from.

Adding costly information processing to a market ω produces the game $G = (\omega, \kappa)$. Everything about the game is common knowledge. Thus, the seller is knowledgeable about the buyer's costs of information processing. This assumption is more plausible in some settings than others. For example, it seems likely that banks will measure the time it takes to read the booklet that contains all checking account fees.

3.5 Equilibrium Concept

The equilibrium concept employed is mixed strategy perfect Bayesian equilibrium (PBE). The seller has pricing strategy $\sigma(\theta)$, which is the probability of pricing high for quality level θ. The buyer has information processing strategy $\pi_p \in \Pi$, which depends on price, and purchasing strategy $\alpha(p, s)$, which is the probability of purchasing for each price and posterior. Finally, the buyer has beliefs $\mu(p, s)$ of the probability of high quality for each price and posterior.[9]

The timeline is:

1. A quality level θ is realized, and the seller observes it.
2. The seller sets the probability of pricing high according to $\sigma(\theta)$.
3. The buyer observes a realized price p and chooses an information processing technology according to π_p.

[9] Note that the agent does not mix over information acquisition technologies. This is without loss of generality as Caplin and Dean [2012] show it is not optimal to mix over information acquisition technologies for a Shannon cost function because there is a unique best technology.

4. The buyer realizes posterior s and chooses a probability of purchasing according to $\alpha(p, s)$.

For a game G, a mixed strategy PBE is a 4-tuple $(\widehat{\sigma}, \widehat{\pi}, \widehat{\alpha}, \widehat{\mu})$ that satisfies seller optimality, buyer optimality, and Bayesian beliefs:

- Seller optimality
  - $\forall \theta \in \Theta$, $\widehat{\sigma}(\theta) > 0$ implies
  $$p_H \in \arg\max_{p \in \{p_L, p_H\}} E[V(p, x) | \theta],$$
  and $\widehat{\sigma}(\theta) < 1$ implies
  $$p_L \in \arg\max_{p \in \{p_L, p_H\}} E[V(p, x) | \theta].$$
- Buyer optimality
  - $\forall p \in P, s \in S, \widehat{\alpha}(p, s) > 0$ implies
  $$1 \in \arg\max_{x \in \{0,1\}} E[U(\theta, p, x) | p, s],$$
  and $\widehat{\alpha}(p, s) < 1$ implies
  $$0 \in \arg\max_{x \in \{0,1\}} E[U(\theta, p, x) | p, s].$$
  - $\forall p \in P$,
  $$\widehat{\pi}_p \in \arg\max_{\pi \in \Pi} E[U(\theta, p, x) | p, \pi] - K(\pi, \kappa).$$
- Bayesian beliefs in equilibrium
  - $\forall s \in S(\widehat{\pi}_p)$, if $\widehat{\sigma}(\theta) > 0$ for any $\theta \in \Theta$, then
  $$\widehat{\mu}(p_H, s) = \frac{\lambda \widehat{\sigma}(\theta_H) \widehat{\pi}_{p_H}^{\theta_H}(s)}{\lambda \widehat{\sigma}(\theta_H) \widehat{\pi}_{p_H}^{\theta_H}(s) + (1 - \lambda) \widehat{\sigma}(\theta_L) \widehat{\pi}_{p_H}^{\theta_L}(s)},$$
  and if $\widehat{\sigma}(\theta) < 1$ for any $\theta \in \Theta$, then
  $$\widehat{\mu}(p_L, s) = \frac{\lambda (1 - \widehat{\sigma}(\theta_H)) \widehat{\pi}_{p_L}^{\theta_H}(s)}{\lambda (1 - \widehat{\sigma}(\theta_H)) \widehat{\pi}_{p_L}^{\theta_H}(s) + (1 - \lambda)(1 - \widehat{\sigma}(\theta_L)) \widehat{\pi}_{p_L}^{\theta_L}(s)}.$$

4 Buyer Best Responses

To find the equilibria of this sequential game, I first determine the best responses of the buyer to every possible interim belief β_p that the product is of high quality for every price p. The buyer's best response is composed of two parts: the optimal information processing strategy and optimal purchasing strategy.

Because $\theta_L - p_L > u$ and $\theta_H > \theta_L$, it is always optimal for the buyer to purchase the product when the price is low, regardless of the quality level. As a result, the buyer has no incentive to expend costly attentional effort if a low price is observed. On the other hand, because $\theta_H - p_H > u$ and $\theta_L - p_H < u$, when the price is high, the buyer only wants to purchase high quality products, so there are incentives to undertake costly attentional effort if the product could be of low quality. In what follows, I solve for the optimal information processing strategy and optimal purchasing strategy at price p_H given all interim beliefs β_{p_H}. The price will be kept general, however, to indicate that the solution can be applied when there are more than two prices available to the seller.

4.1 Choosing the Optimal Information Processing Technology

Caplin and Dean [2012] show how to determine the optimal information processing technology π for rational inattention theory using a posterior-based approach. This problem is greatly simplified by their observations that the optimal technology generates a single posterior for each action, that the corresponding action is strictly optimal for each posterior, and that the solution is unique.

To find the solution in this binary action problem, I first assume that both actions (purchase and not purchase) are taken with positive probability. For this model, let s_p^0 be the posterior for which the buyer does not purchase at price p and s_p^1 be the posterior for which the buyer purchases at price p. Because both actions are taken with positive probability, the resulting optimization problem is:

$$\max_{s_p^0, s_p^1, \pi\left(s_p^1\right)} \pi\left(s_p^1\right)\left[s_p^1\left(\theta_H - p\right) + \left(1 - s_p^1\right)\left(\theta_L - p\right)\right] + \left(1 - \pi\left(s_p^1\right)\right) u - K\left(\pi, \kappa, \beta_p\right)$$

where again

$$K\left(\pi, \kappa, \beta\right) = \kappa\left(\left[\sum_{s \in S(\pi)} \pi\left(s\right)\left(s \ln\left(s\right) + \left(1 - s\right)\ln\left(1 - s\right)\right)\right] - \beta \ln\left(\beta\right) + \left(1 - \beta\right)\ln\left(1 - \beta\right)\right).$$

Finally, because $\pi \in \Pi_\beta$, it must satisfy Bayes rule, so that

$$\pi\left(s_p^1\right) s_p^1 + \left(1 - \pi\left(s_p^1\right)\right) s_p^0 = \beta_p.$$

The first order conditions reduce to the following ratios:

$$\begin{aligned} \frac{s_p^1}{s_p^0} &= \exp\left(\frac{\theta_H - p - u}{\kappa}\right), \\ \frac{1 - s_p^1}{1 - s_p^0} &= \exp\left(\frac{\theta_L - p - u}{\kappa}\right). \end{aligned}$$

Thus, when both actions are taken with positive probability, the optimal posteriors are

$$\begin{aligned}
s_p^0 &= \frac{1-\exp\left(\frac{\theta_L-p-u}{\kappa}\right)}{\exp\left(\frac{\theta_H-p-u}{\kappa}\right)-\exp\left(\frac{\theta_L-p-u}{\kappa}\right)}, \\
s_p^1 &= \exp\left(\frac{\theta_H-p-u}{\kappa}\right)s_p^0.
\end{aligned}$$

Note that these posteriors are not impacted by interim beliefs, just price p, the cost of attention κ, and market parameters θ_L, θ_H, and u. However, the unconditional likelihood $\pi\left(s_p^1\right)$ of posterior s_p^1 is determined by interim beliefs because

$$\pi\left(s_p^1\right)=\min\left(\max\left(\frac{\beta_p-s_p^0}{s_p^1-s_p^0},0\right),1\right).$$

If $\beta_p \leq s_p^0$ or $\beta_p \geq s_p^1$, then the probability $\pi\left(s_p^1\right)$ will equal 0 or 1, so only one posterior will be produced. This means that no attentional effort has been exerted, the interim belief becomes the posterior, and only one action is taken. Because the buyer does not attend to quality above or below these posteriors, they can be interpreted as reservation or threshold beliefs.

As p converges down to θ_L-u, both thresholds converge to 0. This means that for very low prices, the buyer will drop out from attending and will purchase the product even with little information on quality. On the other hand, as p converges up to θ_H-u, both thresholds converge to 1. At these very high prices, the buyer will also drop out from attending, but will instead refrain from purchasing.

4.1.1 Example: Minimize Type I Errors

Specify a market as

$$(\lambda,\Theta,P,u)=(0.5,\{100,200\},\{50,100\},25),$$

and let $\kappa=15$. In other words, the probability that the seller's product is of high quality is 50%, the possible quality levels are 100 and 200, the feasible prices are 50 and 100, and the buyer's outside option is worth 25. We can use the steps above to determine the optimal information processing technology for this market.

First, the two posteriors s_{100}^0 and s_{100}^1 are approximately 0.01 and 0.81 respectively. Because the high price of 100 is not very high, the buyer chooses to be very certain of quality when not buying, but less certain of quality when buying. The reason for this asymmetry is that, at this price, the buyer wants to buy if there is a decent chance that the quality level is high. In other words, the buyer will adjust their attention to reduce Type I errors: mistakenly not buying a high quality product.

Assume that all sellers, regardless of their type, pool at a price of 100, so that $\beta_{100}=0.5$. In this case, the interim belief is between the thresholds s_{100}^1 and s_{100}^0, so the optimal information technology puts some weight on both posteriors: $\pi\left(s_{100}^1\right)$ is approximately 0.61.

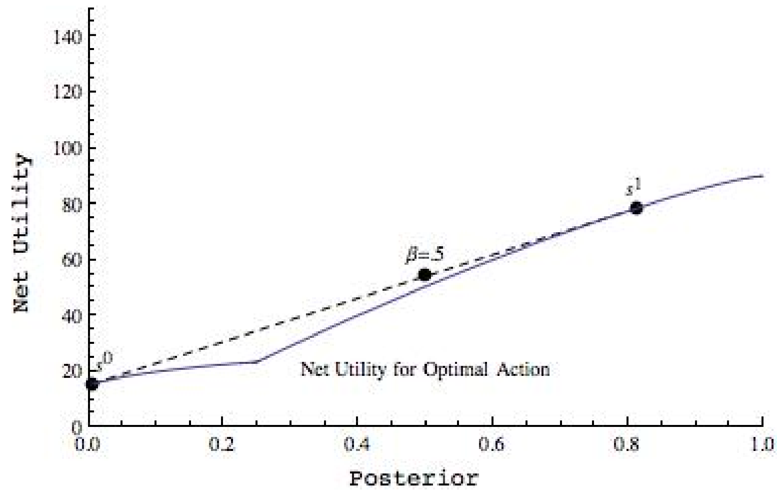

Figure 1: The net utility (purchase utility minus processing costs) produced by the optimal information processing technology and optimal purchasing strategy at price 100 for market $(\lambda, \Theta, P, u) = (0.5, \{100, 200\}, \{50, 100\}, 25)$, cost $\kappa = 15$, and prior $\beta_{100} = .5$.

This solution is summarized by figure 1, which is based on the approach of Caplin and Dean [2012]. The figure shows how net utility for the optimal action changes with posterior beliefs. The two optimal posteriors are shown as the left and right dots, and the interim belief is shown between them. The net utility produced by the optimal technology at this interim belief is found by mixing over the net utilities of the two different posteriors.

4.1.2 Example: Minimize Type II Errors

Now assume that the quality levels are 50 and 150, not 100 and 200, so that the market (λ, Θ, P, u) is $(0.5, \{50, 150\}, \{50, 100\}, 25)$. How does the optimal information processing technology change?

In this market, the two posteriors s^0_{100} and s^1_{100} are approximately 0.19 and 0.99 respectively. Unlike the previous case, the buyer chooses to be very certain of quality when buying and less certain of quality when not buying. This is because the buyer wants to refrain from buying if there is a decent chance that the quality is low. In other words, the buyer will adjust their attention to reduce Type II errors: mistakenly buying a low quality product.

The overall level of attention has not changed. This is only because the interim belief is symmetric and the quality levels have shifted evenly.

This illustrates how information processing and drop out thresholds can change substantially with market conditions. Also, it shows that they change in response to the costliness of different mistakes.

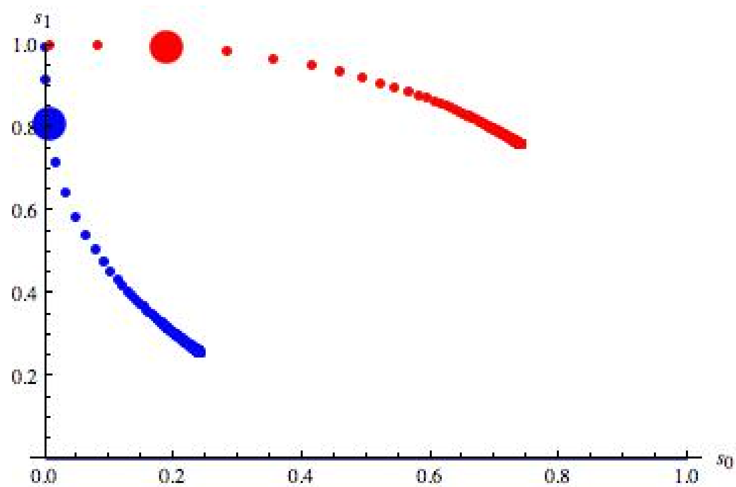

Figure 2: Threshold posteriors for costs $\kappa \in \{5, 10, ..., 1000\}$ at price 100 for market $(\lambda, \Theta, P, u) = (0.5, \{100, 200\}, \{50, 100\}, 25)$ and prior $\beta_{100} = .5$ (dots below) and at price 100 for market $(\lambda, \Theta, P, u) = (0.5, \{50, 150\}, \{50, 100\}, 25)$ and prior $\beta_{100} = .5$ (dots above).

Figure 2 shows how threshold beliefs change with the cost of processing information for these two markets. The dots representing threshold beliefs in the first market are below and to the left of the dots for the second market. As information costs go to zero, thresholds converge smoothly to $s^1_{100} = 1$ and $s^0_{100} = 0$ for both sets of quality levels. However, as costs increase, the thresholds move away from this point in different directions. Each dot represents an incremental increase of 5 in cost, from 5 to 1000. The large dots represent the solutions for $\kappa = 15$, as above. For low information costs, both optimal information technologies give very precise posteriors for one of the actions, but as costs increase, the posteriors for both actions become less precise for both optimal information technologies. As costs increase, the thresholds are converging from above and below to the belief for which the buyer is indifferent between purchasing or not:

$$s = \frac{u - (\theta_L - p_H)}{(\theta_H - \theta_L)}.$$

4.2 Conditional Demands

Because both actions are uniquely optimal for their corresponding posteriors, the optimal purchasing strategy is deterministic for a given price and posterior, even though the solution concept allows for stochasticity. However, the stochasticity in attention means that overall product demand, given by $\pi\left(s^1_{p_H}\right)$, is not necessarily deterministic.

While $\pi\left(s^1_{p_H}\right)$ gives the unconditional probability of a buyer purchasing at price p_H, a seller of type θ will choose a pricing strategy based on the probability that the buyer will purchase *their* product at price

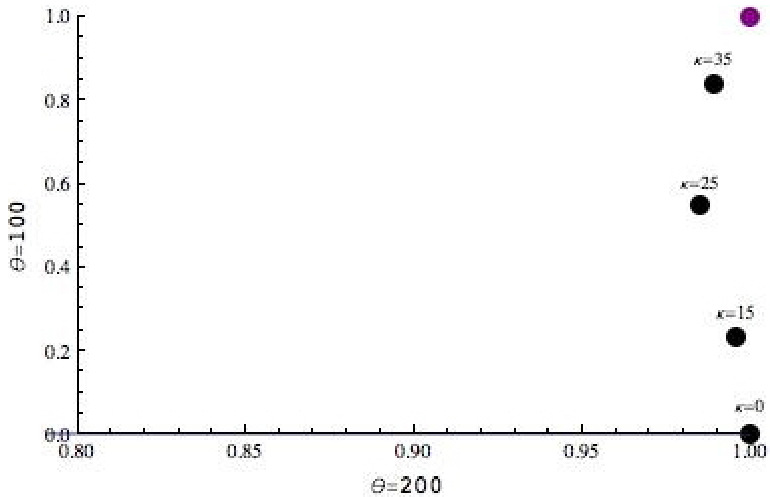

Figure 3: Conditional demands as information processing costs change at price 100 for market $(\lambda, \Theta, P, u) = (0.5, \{100, 200\}, \{50, 100\}, 25)$ and prior $\beta_{100} = .5$.

p_H, which is denoted by $d^{\theta}_{p_H}$. Using Bayes rule, it can be determined that conditional demands are

$$d^{\theta_H}_{p_H} = \Pr\left(\text{buy}|\theta_H\right) = \frac{s^1_{p_H}\pi\left(s^1_{p_H}\right)}{\beta_{p_H}},$$

and

$$d^{\theta_L}_{p_H} = \Pr\left(\text{buy}|\theta_L\right) = \frac{\left(1 - s^1_{p_H}\right)\pi\left(s^1_{p_H}\right)}{\left(1 - \beta_{p_H}\right)}.$$

Because posteriors are not impacted by interim beliefs with rational inattention, the conditional demand for the sellers of low quality products at the high price are strictly in β_{p_H}. Also, if more than one action is taken, then $d^{\theta_H}_{p_H} > d^{\theta_L}_{p_H}$ because $\beta_{p_H} < s^1_{p_H}$ implies

$$\left(1 - \beta_{p_H}\right)s^1_{p_H} > \beta_{p_H}\left(1 - s^1_{p_H}\right).$$

4.2.1 Example: Conditional Demands at Lower Prices

Once again, let market (λ, Θ, P, u) be $(0.5, \{100, 200\}, \{50, 100\}, 25)$ and $\kappa = 15$. If all sellers pool at the high price, then the conditional demands d^{200}_{100} and d^{100}_{100} are approximately 1 and 0.23 respectively. These values reflect the buyer's willingness to make mistakes when purchasing at this lower high price.

Figure 3 shows how these conditional demands change with κ. Given the assumptions in this model, when the price is 100 and all sellers offer that price, a fully informed buyer will always buy when the quality is high and never buy when the quality is low, which is represented by the dot in the lower right hand corner. However, a fully uninformed buyer will always buy the product, which is represented by the dot in the upper

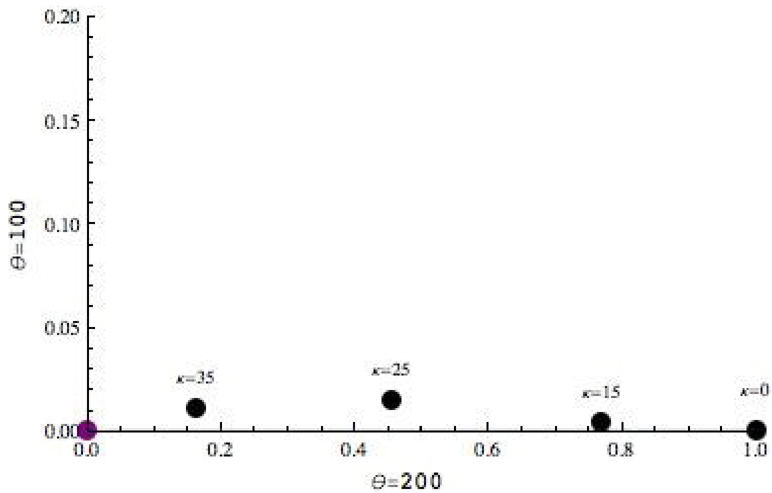

Figure 4: Conditional demands as information processing costs change at price 100 for market $(\lambda, \Theta, P, u) = (0.5, \{50, 150\}, \{50, 100\}, 25)$ and interim beliefs $\beta_{100} = .5$.

right hand corner. As κ goes to 0, demands converge smoothly to the fully informed case, and as κ goes to infinity, demands converge smoothly to the fully uninformed case.

4.2.2 Example: Conditional Demands at Higher Prices

Now assume again that the quality levels are 50 and 150, not 100 and 200, so that the market (λ, Θ, P, u) is $(0.5, \{50, 150\}, \{50, 100\}, 25)$. How do demands change? If sellers pool at this price for processing cost $\kappa = 15$, the conditional demands d_{100}^{200} and d_{100}^{100} are now approximately 0.71 and 0 respectively.

Figure 4 shows how conditional demands change with κ, but this time for the new quality level. At a price of 100, a fully uninformed buyer will never buy the product, because it is equally likely to be of high or low quality. As a result, the path of conditional demands now hugs the other axis.

5 Equilibrium

In this section, I will describe the two types of equilibria that are possible in this game and fully characterize them. To start, I will solidify the connection between conditional demands and pricing strategies.

5.1 Conditional Demands and Pricing Strategies

An optimal pricing strategy will only put positive weight on price p if no other price makes a higher expected return, and the expected return from charging price p for a product of quality θ is determined by the probability a product of quality θ is purchased at price p, which is the conditional demand d_p^θ. Thus, $\widehat{\sigma}(\theta) > 0$ only if

$$d_{p_H}^\theta \times p_H \geq d_{p_L}^\theta \times p_L$$

and $\widehat{\sigma}(\theta) < 1$ only if

$$d_{p_H}^\theta \times p_H \leq d_{p_L}^\theta \times p_L.$$

There are two useful implications that come immediately from this. First, because buyers always purchase when the price is low, a seller of type θ will only mix between prices if the expected return from pricing high is equal to the low price, which is true when

$$d_{p_H}^\theta = \frac{p_L}{p_H}.$$

Second, because $d_{p_H}^{\theta_H} > d_{p_H}^{\theta_L}$ when $d_{p_H}^{\theta_H} > 0$ and $d_{p_H}^{\theta_L} > 0$, if it is optimal for sellers of type θ_H to mix, then sellers of type θ_L will price low with probability 1, and if it is optimal for sellers of type θ_L to mix, then sellers of type θ_H will price high with probability 1

Finally, pricing strategies impact conditional demands through β_p, the probability of θ_H given price p, where

$$\begin{aligned}
\beta_{p_H} &= \frac{\lambda \widehat{\sigma}(\theta_H)}{\lambda \widehat{\sigma}(\theta_H) + (1-\lambda)\widehat{\sigma}(\theta_L)}, \\
\beta_{p_L} &= \frac{\lambda(1-\widehat{\sigma}(\theta_H))}{\lambda(1-\widehat{\sigma}(\theta_H)) + (1-\lambda)(1-\widehat{\sigma}(\theta_L))}.
\end{aligned}$$

5.2 Two Types of Equilibria

5.2.1 Pooling at a Low Price

For any game G, there always exists a pooling equilibrium where both types of seller charge a low price with probability 1, which I will call the "pooling low" equilibrium. The pooling low equilibrium requires strong off-equilibrium path beliefs: if the buyer sees a high price, then they believe the seller is a low quality type with probability 1. In this equilibrium, the buyer chooses a completely uninformative information processing strategy, which produces just one posterior and thus returns the interim belief at the posterior. Even when there are no costs to attention, this is equilibrium still exists, but the buyer must choose the completely uninformative information processing technology over equally costless, but informative, information processing technologies.

5.2.2 Mimic High Equilibrium

There is just one other type of equilibria, which I will call the "mimic high" equilibrium. It is the more informative equilibrium because prices are weakly more informative about quality than the pooling low equilibrium for all values of κ and strictly more informative for some values of κ. In this equilibrium, the high quality seller always puts probability 1 on setting a high price, and the low quality seller mimics the high quality seller by charging a high price with a certain probability, which is determined by the parameters of the game.

In the mimic high equilibrium, strategies are as in the full information equilibrium if the costs of information processing are zero and converge to those in the full information equilibrium as costs of information processing go to zero. Surprisingly, for some market conditions, when the cost of information processing gets high enough, an increase in information costs actually decreases the probability that low quality sellers mimic. The reason is that low quality sellers must mimic less to overcome the pull of buyers towards not purchasing.

Theorem 1 *For all κ, there exists an equilibrium ("mimic high") where high quality sellers price high with probability 1 and low quality sellers price high with a unique probability $\eta \in [0,1]$. As κ varies, there are four possible regions of the equilibrium:*

1. *For $\kappa = 0$, separating prices ($\eta = 0$)*
2. *For $\kappa \in (0, \kappa^*)$, increasing mimicking with κ, where*
$$\eta = \frac{\lambda}{1-\lambda} \frac{s^0_{p_H} + s^1_{p_H} - s^0_{p_H} s^1_{p_H} - 1}{s^0_{p_H} s^1_{p_H} + s^0_{p_H} \frac{p_L}{p_H} - s^1_{p_H} \frac{p_L}{p_H} - s^0_{p_H}}$$
3. *For $\kappa \in [\kappa^*, \kappa^{**})$, pooling high ($\eta = 1$)*
4. *For $\kappa \in [\kappa^{**}, \infty)$, decreasing mimicking with κ, where again*
$$\eta = \frac{\lambda}{1-\lambda} \frac{s^0_{p_H} + s^1_{p_H} - s^0_{p_H} s^1_{p_H} - 1}{s^0_{p_H} s^1_{p_H} + s^0_{p_H} \frac{p_L}{p_H} - s^1_{p_H} \frac{p_L}{p_H} - s^0_{p_H}}$$

Proof of this theorem is in the appendix. A key part of the proof is that in regions 2 and 4, η is such that the buyer puts in enough attentional effort when the price is high to make low quality sellers indifferent between setting either price. The uniqueness of η comes from the fact that conditional demand has a single crossing property with the line of indifference as the interim belief decreases.

As Theorem 1 indicates, the mimic high equilibrium can be broken into four regions of κ, which are ordered. If uninformed buyers purchase when sellers pool at the high price, then just regions 1, 2, and 3 occur. On the other hand, if buyers do not purchase when sellers pool at the high price, then either all four

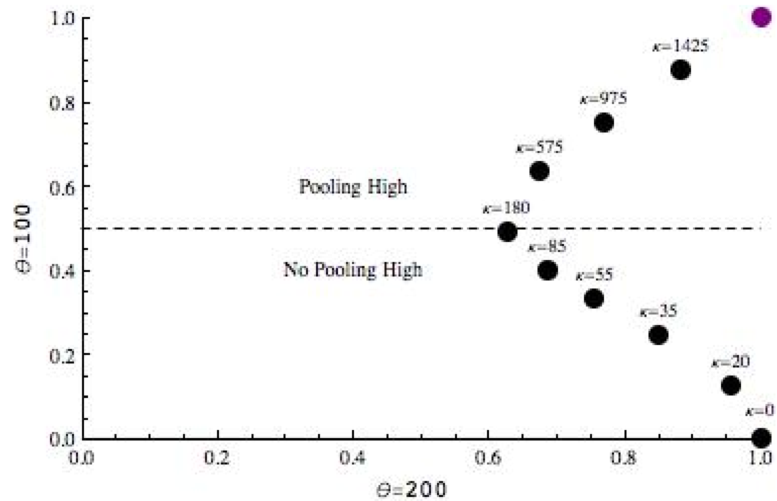

Figure 5: Conditional demands as information processing costs change at price 100 for market $(\lambda, \Theta, P, u) = (0.255, \{100, 200\}, \{50, 100\}, 25)$ and prior $\beta_{100} = .5$.

regions will occur, or just regions 1, 2, and 4. The surprise is found in region 4: that mimicking can decrease as the cost of information processing increases.

It remains to be shown that these are the only two types of equilibria of this game. As mentioned above, if high quality sellers mix, then low quality sellers must put full weight on pricing low. However, if the low quality sellers put no weight on pricing high, then high quality sellers must put full weight on pricing high. Thus, high quality sellers will not mix in any equilibrium. Also, if high quality sellers put no weight on pricing high, then low quality sellers must also put no weight on pricing high. Thus, pooling low is the only equilibrium if high quality sellers price low with probability 1. Finally, if high quality sellers price high with probability 1, then all of the possible equilibria at different costs of information processing are captured by the mimic high equilibrium.

5.2.3 Example: Uninformed Buyers Purchase at High Pooling

Let the market (λ, Θ, P, u) be $(0.255, \{100, 200\}, \{50, 100\}, 25)$, so that the seller is very likely to be low quality, but the high price is also low. In this case, if the buyer is uninformed and sellers pool at the high price, then the buyer will purchase the product. As a result, we will see the first three regions on the mimic high equilibrium in that order.

Figure 5 shows the path of conditional demands for different costs of information processing if both sellers charge a high price with probability 1. Region 1 is represented by the dot in the lower right hand corner (when $\kappa = 0$), and region 3 consists of all costs at and above the dashed line. Region 2 consists of all costs

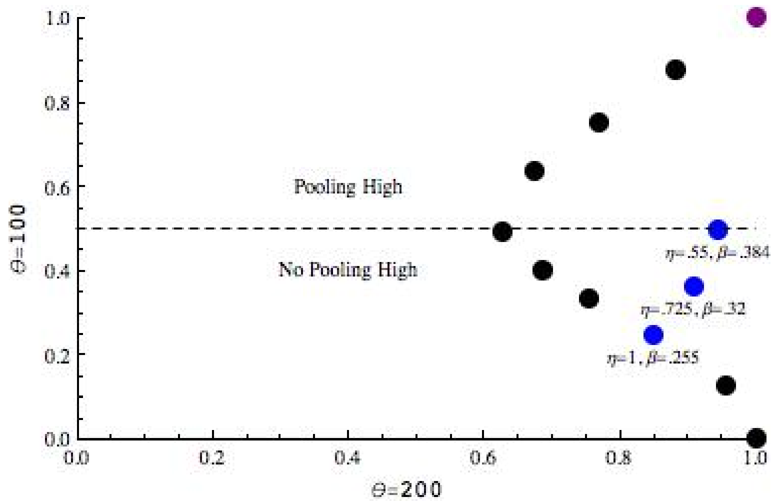

Figure 6: Conditional demands as mimicking rate η changes at price 100 for market $(\lambda, \Theta, P, u) = (0.5, \{100, 200\}, \{50, 100\}, 25)$ and $\kappa = 35$.

below the dashed line, except for the one in the lower right hand corner. In figure 6, the small dots above the dot for $\kappa = 35$ demonstrate what happens to conditional demands as the low quality seller decreases the amount they mimic, which increases β_{p_H}. Equilibrium occurs when the amount of mimicking is such that β_{p_H} generates conditional demands that make the low quality seller indifferent between charging either price, which occurs at the dashed line. As the costs increase, the decrease in mimicking needed to get to equilibrium is smaller, demonstrating that mimicking is increasing with cost.

5.2.4 Example: Uninformed Buyers Do Not Purchase at High Pooling

Now let the market (λ, Θ, P, u) be $(0.745, \{50, 150\}, \{50, 100\}, 25)$, so that the seller is very likely to have a high quality product, but the quality levels are lower. In this case, if the buyer is uninformed, they will not purchase the product, so all four regions of the mimic high equilibrium are possible.

As figure 7 shows, all four regions of the mimic high equilibrium do occur. In the far right corner is region 1. For the cost levels to the left of this corner and below the dashed line, such as $\kappa = 35$, the mimic high equilibrium is in the second region, where mimicking is decreasing with cost. For costs above the line, such as $\kappa = 175$, the mimic high equilibrium is in region 3, where both sellers pool at the high price. For costs under the line and to the left, such as $\kappa = 975$, the mimic high equilibrium is in region 4, where some mimicking occurs, but this mimicking decreases with cost. This occurs because the pull to not purchase increases with information processing costs in this region.

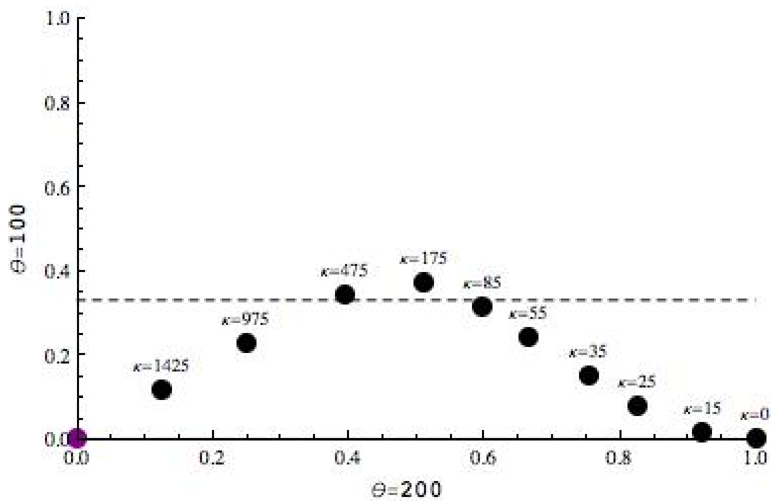

Figure 7: Conditional demands as information processing costs change at price 100 for market $(\lambda, \Theta, P, u) = (0.745, \{50, 150\}, \{50, 100\}, 25)$ and prior $\beta_{100} = .5$.

5.3 Increasing the Complexity of the Market

Here I briefly consider how the set of equilibria changes with an increase in the number of prices, number of quality levels, number of sellers, or number of buyers.

The analysis in this paper can be extended to cover a large number of prices, but it is important that the number of prices available to the seller be finite, which is true in most retail settings where prices are rounded to the nearest penny and bounded for practical reasons. If there are a large number of prices, then the only active low price (prices where $\theta_L - p < u$) will be the highest one, because a monopolist who is known to have a low quality product always has a profitable deviation to that price. When there are multiple high prices (prices where $\theta_H - p < u$), there can be many equilibria, but equilibria similar to the ones in this paper are preserved. There is always a pooling low equilibrium at the highest low price, with off-equilibrium beliefs that a deviation to any other price is made by a low quality firm. Also, there is always a mimic high equilibrium where high quality sellers set a price equal to *any* of one the high prices and low quality sellers sometimes set that high price and otherwise set the highest low price. This is once again supported with off-equilibrium beliefs that a deviation to any other price is made by a low quality firm.

As the number of quality levels increases, it makes sense to consider a commiserate increase in the number of prices, so that there is a price that corresponds to each quality level, as in this model. The pooling low equilibrium is preserved, with off-equilibrium path beliefs that a deviation to any price above the lowest price must be made by a low quality type.

As the number of sellers increases, there are two possibilities, both of which occur in the literature: sellers know each other's realized quality levels before setting prices or not. If we treat the seller as an expert in the type of product that is being sold, then it seems reasonable to assume that they know the realized quality level of all firms. In this case, when there exists an equilibrium with features of both equilibria in this paper. When both firms are low quality, they both set low prices, when both firms are high quality, they both set high prices, and when one firm is low quality and one high quality, the high quality firm prices high and the low quality firm sometime mimics high.

As the number of buyers increases, there is no change in the set of equilibria, as each buyer can be treated separately in their interaction with the seller and the seller is a risk neutral expected utility maximizer.

6 Experiment

To see if either equilibrium can explain market behavior when buyers have access to a freely available source of information about quality, I conducted a laboratory experiment in the Center for Experimental Social Science laboratory at New York University with undergraduate students that implements the market setting in this paper. Because my goal is to see how subjects play given a freely available source of information, I do not induce the attentional cost structure in this model. Instead, I provide buyers with a free source of information about quality which takes cognitive effort to process.

As in Gabaix, Laibson, Moloche, and Weinberg [2006], Caplin, Dean, and Martin [2011], Caplin and Martin [2011], and Caplin and Martin [2012], the free information source in my experiment is a string of numbers that together add up to the quality level. This source of information about quality is meant to imitate aggregating disparate information about quality or aggregating multiple dimensions of quality. There are cognitive constraints on processing this information, and I impose a time limit on buyers, which makes their cognitive constraints binding.

In last half of rounds, the play of subjects is very close to the predictions of the mimic high equilibrium. High quality sellers charge a high price around 98% of the time, and low quality sellers charge a high price around 16% of the time. Pricing high is a best response for high quality sellers, and mixing is almost a best response for low quality sellers. On the other side of the market, buyer demands are very close to those predicted by rational inattention, and buyer consideration times suggest changes in attention that are consistent with rational inattention.

6.1 Implementing the Market Setting

In this experiment, subjects can either be "sellers" or "buyers". Sellers are assigned the preferences and actions of the seller in the model, and buyers those of the buyer. Payoffs are given in Experimental Currency

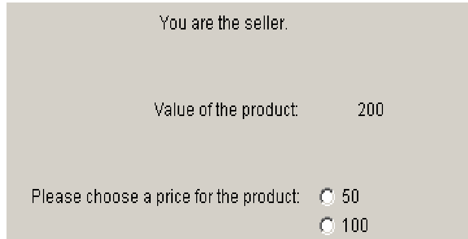

Figure 8: An example of the choice screen for sellers in the experiment.

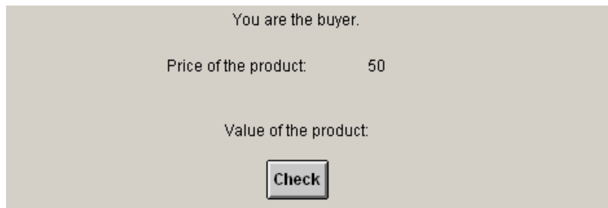

Figure 9: An example of the choice screen for buyers in the experiment.

Units (ECU), and subjects are told that 40 ECU are equal to \$1. Full instructions are provided in the appendix.

The experiment is based on the following market:

$$\omega = (\lambda, \Theta, P, u) = (0.5, \{100, 200\}, \{50, 100\}, 25).$$

In other words, the probability that the seller's product is of high quality is 50%, the possible quality levels are 100 and 200 ECU, the possible prices are 50 and 100 ECU, and the buyer's outside option is worth 25 ECU. This market is the same as in the first example of section 4.

In each of the 30 rounds, subjects are randomly and anonymously matched into pairs, and one player is assigned to be the seller and the other to be the buyer.

In each round, the seller is randomly assigned a hypothetical "product" that has a value to the buyer of either 100 or 200. After being shown the value of the product, the seller offers a price for that product of 50 or 100. Next, the buyer is shown the price and chooses whether to accept the offer or to take the "alternative", which is worth 25. Before making their choice, the buyer can get information about the value of the product if they click through to a second screen. An example of the seller's display is shown in figure 8 and of the buyer's display in figure 9.

77	96	-60	-43	-89	-26	7	-1	-51	17
92	7	-81	3	20	81	40	-57	10	58

Figure 10: An example of the 20 number strings used in the experiment.

6.2 Freely Available Information

As mentioned previously, subjects are presented with a string of 20 numbers to calculate, and the sum of these numbers is equal to the value of the product. This addition task is meant to capture aggregating disparate sources of information about quality or aggregating multiple dimensions of quality. It was selected because it allows the buyer great flexibility how they attend to the information. Subjects can carefully add up some numbers, or they can just scan the numbers for certain features.

Each of the terms in each calculation is a whole number between -100 and 100, and subjects are told the process by which these numbers are generated, which is that the terms are randomly drawn until the sum equals the value. An example expression is presented in figure 10. The seller does not know the exact string that the buyer faces, and buyers are not allowed to use calculators or scratch paper during the experiment.

Buyers have up to 90 seconds to make a decision, but can submit their decision earlier. If the 90 seconds expires, the alternative will be automatically chosen.[10] The buyers payoff is the value of the option minus the price (in ECU) if they accept the offer or 25 ECU if they choose the alternative. The seller's payoff is the price if their offer is accepted and 0 otherwise. After each round, both players are shown the value of the product, the seller's and buyer's choices, and their own payoff in that round.

Finally, subjects are paid for 6 random rounds, plus a $10 show-up fee. The average payoff was approximately $20 for an experiment that lasted approximately 1.25 hours. The experiment was programmed in z-Tree (Fischbacher [2007]).

6.3 Observed Prices and Demands

Over two sessions, 34 subjects completed the experiment, for a total of 1020 rounds. Table 1 shows the fraction of rounds in which each price was offered by round number and quality level. Even in the first 15 rounds, the vast majority of sellers with high quality products charged a high price. This only increased in the second half of rounds. On the other hand, in most rounds where the seller had a low quality product, the seller offered a low price. However, in approximately 19% of rounds, sellers with a low quality product

[10] The time limit was reached in less than 1% of rounds.

did offer a high price. This frequency rose slightly to around 20% in the second half of rounds.

Table 1. Fraction of rounds each price offered by quality level across rounds.

	Rounds 1-15		Rounds 15-30		Overall	
P\ θ	100	200	100	200	100	200
50	83%	5%	80%	1%	81%	3%
100	17%	95%	20%	99%	19%	97%

Table 2 shows the percentage of offers accepted by round number and quality level. Even from the beginning of the experiment, buyers almost never made the mistake of rejecting an offer with a low price, as can be seen in the first row. The percentage of times that the buyer mistakenly accepted an offer with a high price and a low quality product stays around 50% over the course of the experiment. However, the frequency with which buyers mistakenly rejected an offer with a high price and a high quality product decreased over the course of the experiment. In the second half of rounds, 92% of such offers were accepted.

Table 2. Fraction of offers accepted at each price and quality level across rounds.

	Rounds 1-15		Rounds 15-30		Overall	
P\ θ	100	200	100	200	100	200
50	97%	100%	100%	100%	99%	100%
100	47%	79%	54%	92%	51%	85%

6.4 Close to Mimic High Equilibrium

These observed prices and demands correspond closely to the mimic high equilibrium analyzed in section 5. In this equilibrium, high quality sellers price high with probability 1 and low quality sellers sometimes mimic high quality sellers. For high quality sellers to price high with probability one, it must be that

$$\frac{d_{p_H}^{\theta_H}}{d_{p_L}^{\theta_H}} > \frac{p_L}{p_H} = 0.50.$$

Indeed, $\frac{d_{p_H}^{\theta_H}}{d_{p_L}^{\theta_H}}$ is far above this threshold, especially in the second half of rounds, where $d_{p_H}^{\theta_H} = 92\%$ and $d_{p_L}^{\theta_H} = 100\%$. Also, for low quality sellers to mimic with an interior probability, as in the data, it should be that 52

$$\frac{d_{p_H}^{\theta_H}}{d_{p_L}^{\theta_H}} = \frac{p_L}{p_H} = 0.50.$$

In the second half of rounds, $\frac{d_{p_H}^{\theta_H}}{d_{p_L}^{\theta_H}} = 54\%$, just 4% away from the required percentage.

Looking just at the second half of rounds, the value of κ that minimizes the distance between actual demands and predicted demands is 11.5. Given that the actual probability of high quality given a price of 100 (β_{100}) is 0.8, rationally inattentive buyers with $\kappa = 11.5$ should accept offers with a high price and low quality 52% of the time and offers with a high price and high quality 99% of the time. The actual demands, as shown in table 2, are 54% and 92% respectively.

6.5 Revealing Changes in Attention

As discussed in Caplin and Martin [2012], one way to infer attention is through consideration times. One element of the experimental design that makes this possible is that subjects must click a button to see the string of numbers and then click back to make their choice. This means that the time spent looking at the information about quality can be somewhat separated from the time taken to contemplate a choice. Across prices and whether or not the offer was accepted, subjects spent on average 6 seconds on the choice screen. However, as table 3 shows, the average time spent on the screen with the string of numbers varied both by price and whether or not the offer was accepted. For a price of 100, the average times are significantly different at the 5% level using a t-test.

Table 3. Average time spent on screen with number string

by price and whether offer was accepted or rejected.

P	Accept	Reject
50	3.8	
100	41.3	53.4

These consideration times reflect rational inattention in two ways. First, attention drops to almost nothing at a price of 50, as predicted by theory. Second, consideration times are 29% higher before rejecting than before accepting. This is because, as shown in example one of section 4, at a price of 100 the buyer wants to minimize Type I errors: mistakenly not buying a high quality product.

7 Conclusion

In a standard sequential pricing game, I model how buyers attend to information about product quality when prices are easily observable using "rational inattention" theory (Sims [2003]). Rational inattention in this setting produces two possible equilibria: one where the entire surplus goes to buyers and one where most of the surplus goes to sellers of high quality goods. I show that there are large differences in how policy interventions related to attention impact buyer welfare in each equilibrium. In a laboratory experiment where buyers face real attentional costs, I find that subjects play very close to predictions of the equilibrium where most of the surplus goes to sellers of high quality goods.

As an extension of this model, it would be interesting to consider what happens if the seller can influence the buyer's cost of processing information (as in Carlin and Manso [2011], Kalaycı and Potters [2011], and Ellison and Wolitzky [2012]) or can bias the information that is available to the buyer. Both possibilities seem realistic given the control that sellers often have over the retail environment. A complication of adding these features to the model is that the seller can communicate information through these actions, giving another channel over which the buyer must have beliefs. One solution employed in the literature is to have buyers be nonstrategic over these actions. Another possibility is to have sellers take these actions before they become aware of the quality of their product.

References

[1] Bagwell, Kyle, and Michael Riordan. 1991. "High and Declining Prices Signal Product Quality." *American Economic Review*, 81(1): 224-239.

[2] Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat. 2012. "Information gathering externalities in product markets." *Journal of Industrial Economics*, LX(1): 162-185.

[3] Bester, Helmut, and Klaus Ritzberger. 2001. "Strategic pricing, signalling, and costly information acquisition." *International Journal of Industrial Organization*, 19: 1347-1361.

[4] Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. March 2012. "Salience and Consumer Choice." NBER Working Paper No. 17947.

[5] Caplin, Andrew, and Mark Dean. 2011. "Search, choice, and revealed preference." *Theoretical Economics*, 6: 19-48.

[6] Caplin, Andrew, and Mark Dean. October 2012. "Shannon Entropy and Rational Inattention: Solutions, Behavioral Regularities, and Extensions." Mimeo.

[7] Caplin, Andrew, Mark Dean, and Daniel Martin. 2011. "Search and Satisficing." *American Economic Review*, 101: 2899-2922.

[8] Caplin, Andrew, and Daniel Martin. June 2011. "A Testable Theory of Imperfect Perception." NBER Working Paper 17163.

[9] Caplin, Andrew, and Daniel Martin. April 2012. "Defaults and Attention: The Drop-Out Effect." NBER Working Paper 17988.

[10] Carlin, Bruce, and Gustavo Manso. 2011. "Obfuscation, Learning, and the Evolution of Investor Sophistication." *Review of Financial Studies*, 24: 754-785.

[11] Chan, Yuk-Shee, and Hayne Leland. 1982. "Prices and Qualities in Markets with Costly Information." *Review of Economic Studies*, XLIX: 499-516.

[12] Cheremukhin, Anton, Anna Popova, and Antonella Tutino. January 2012. "Rationalizing Variations in Behavioral Decisions." Mimeo.

[13] Compte, Olivier, and Andrew Postlewaite. February 2011. "Mental Processes and Decision Making." Mimeo.

[14] Cooper, Russell, and Cooper Ross. 1984. "Prices, Product Qualities and Asymmetric Information: The Competitive Case." *Review of Economic Studies*, LI: 197-207.

[15] Eliaz, Kfir, and Ran Spiegler. 2011. On the strategic use of attention grabbers. *Theoretical Economics*, 6: 127-155.

[16] Eliaz, Kfir, and Ran Spiegler. 2011. Consideration Sets and Competitive Marketing. *Review of Economic Studies*, 78: 235-262.

[17] Ellison, Glenn, and Alexander Wolitzky. 2012 "A Search Cost Model of Obfuscation." *Rand Journal of Economics*, forthcoming.

[18] Ergin, Haluk, and Todd Sarver. 2010. "A Unique Costly Contemplation Representation." *Econometrica*, 17: 1285-1339.

[19] Fischbacher, Urs. 2007. "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics*, 10(2).

[20] Gabaix, Xavier. March 2011. "A Sparsity-Based Model of Bounded Rationality." NBER Working Paper No. 16911.

[21] Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg. 2006. "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model." *American Economic Review*, 96(4): 1043-1068.

[22] Gennaioli, Nicola, and Andrei Shleifer. 2010. "What Comes to Mind." *Quarterly Journal of Economics*, 125(4): 1399-1433.

[23] Gottlieb, Daniel. February 2011. "Will You Never Learn? Self Deception and Biases in Information Processing." Mimeo.

[24] Heffetz, Ori, and Moses Shayo. 2009. "How Large Are Non-Budget-Constraint Effects of Prices on Demand?" *American Economic Journal: Applied Economics*, 1(4): 170–199.

[25] Judd, Kenneth, and Michael Riordan. 1994. "Price and Quality in a New Product Monopoly." *Review of Economic Studies*, 61, 773-789.

[26] Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp. August 2012. “Rational Attention Allocation Over the Business Cycle.” Mimeo.

[27] Kalaycı, Kenan, and Jan Potters. 2011. Buyer confusion and market prices.” *International Journal of Industrial Organization*, 29(1): 14-22.

[28] Klenow, Peter, and Jonathan Willis. 2007. “Sticky information and sticky prices.” *Journal of Monetary Economics*. 54: 79-99.

[29] Lizzeri, Alessandro. 1999. “Information Revelation and Certification Intermediaries.” *RAND Journal of Economics*, 30(2), 214-231.

[30] Lynch, John, and Dan Ariely. 2009. “Wine Online: Search Costs Affect Competition on Price, Quality, and Distribution.” *Marketing Science*, 19(1): 83-103.

[31] Manzini, Paola, and Marco Mariotti. 2007. “Sequentially Rationalizable Choice.” *American Economic Review*, 97(5): 1824-1840.

[32] Manzini, Paola, and Marco Mariotti. November 2011. “Stochastic Choice and Consideration Sets.” Mimeo.

[33] Maćkowiak, Bartosz, Emanuel Moench, and Mirko Wiederholt. 2009. “Sectoral Price Data and Models of Price Setting.” *Journal of Monetary Economics*, 56(S): 78-99.

[34] Maćkowiak, Bartosz, and Mirko Wiederholt. April 2012. “Inattention to Rare Events.” Mimeo.

[35] Masatlioglu, Yusufcan, Daisuke Nakajima, and Erkut Ozbay. 2012. “Revealed Attention.” *American Economic Review*, 102(5), 2183-2205.

[36] Matějka, Filip. February 2010. “Rigid Pricing and Rationally Inattentive Consumer.” Mimeo.

[37] Matějka, Filip, and Alisdair McKay. November 2011. “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model.” Mimeo.

[38] Matějka, Filip, and Alisdair McKay. 2012. “Simple Market Equilibria with Rationally Inattentive Consumers.” *American Economic Review: Papers & Proceedings*, 102(3): 24-29.

[39] Miller, Ross, and Charles Plott. 1985. “Product Quality Signaling in Experimental Markets.” *Econometrica*, 53: 837-872.

[40] Rabin, Matthew, and Georg Weizsacker. 2009. “Narrow Bracketing and Dominated Choices.” *American Economic Review*, 99: 1508-1540.

[41] Salant, Yuval, and Ariel Rubinstein. 2008. “("A", "f"): Choice with Frames,” *Review of Economic Studies*, 75(4): 1287-1296.

[42] Schwartzstein, Joshua. January 2012. "Selective Attention and Learning." Mimeo.

[43] Sims, Christopher. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics*, 50(3), 665-690.

[44] Sims, Christopher. 2006. "Rational Inattention: Beyond the Linear-Quadratic Case." *American Economic Review*, 96(2), 158-163.

[45] Sims, Christopher. 2010. "Rational inattention: a research agenda." Deutsche Bundesbank Discussion Paper, Series 1: Economic Studies, No 34/2005.

[46] Treviño, Isabel, and Michal Szkup. 2011. "Costly information acquisition in a speculative attack: Theory and experiments." Mimeo.

[47] Tutino, Antonella. 2011. "Rational inattention guides overloaded brains, help economists understand market behavior." *Economic Letter*, 6(3).

[48] Veldkamp, Laura. 2011. *Information Choice in Macroeconomics and Finance*. Princeton University Press.

[49] Vives, Xavier. 2001. *Oligopoly Pricing: Old Ideas and New Tools*. MIT Press.

[50] Voorneveld, Mark, and Jörgen Weibull. 2011. "A Scent of Lemon—Seller Meets Buyer with a Noisy Quality Observation." *Games*, 2: 163-186.

[51] Wiederholt, Mirko. 2010. "Rational Inattention." In *The New Palgrave Dictionary of Economics, Online Edition*. Palgrave Macmillan.

[52] Wolinsky, Asher. 1983. "Prices as Signals of Product Quality." *Review of Economic Studies*, 50(4): 647-658.

[53] Woodford, Michael. February 2012. "Inattentive Valuation and Reference-Dependent Choice." Mimeo.

[54] Yang, Ming. June 2012. "Optimality of Securitized Debt with Endogenous and Flexible Information Acquisition." Mimeo.

[55] Yang, Ming. January 2012. "Coordination with Flexible Information Acquisition." Mimeo.

8 Appendix

Theorem: For all κ, there exists an equilibrium ("mimic high") where high quality sellers price high with probability 1 and low quality sellers price high with a unique probability $\eta \in [0, 1]$. As κ varies, there are four possible regions of the equilibrium:

1. For $\kappa = 0$, separating prices ($\eta = 0$)
2. For $\kappa \in (0, \kappa^*)$, increasing mimicking with κ, where

$$\eta = \frac{\lambda}{1-\lambda} \frac{s^0_{p_H} + s^1_{p_H} - s^0_{p_H} s^1_{p_H} - 1}{s^0_{p_H} s^1_{p_H} + s^0_{p_H} \frac{p_L}{p_H} - s^1_{p_H} \frac{p_L}{p_H} - s^0_{p_H}}$$

3. For $\kappa \in [\kappa^*, \kappa^{**})$, pooling high ($\eta = 1$)
4. For $\kappa \in [\kappa^{**}, \infty)$, decreasing mimicking with κ, where again

$$\eta = \frac{\lambda}{1-\lambda} \frac{s^0_{p_H} + s^1_{p_H} - s^0_{p_H} s^1_{p_H} - 1}{s^0_{p_H} s^1_{p_H} + s^0_{p_H} \frac{p_L}{p_H} - s^1_{p_H} \frac{p_L}{p_H} - s^0_{p_H}}$$

Proof. When $\kappa = 0$, the unique value of η for the mimic high equilibrium is $\eta = 0$. In other words, sellers separate by offering different prices. Buyers exert full attentional effort at no cost, so they can distinguish between sellers perfectly. Thus, low quality sellers have no profitable deviation to any $\eta > 0$, and high quality sellers make the maximal return. This is the unique value of η because if low quality sellers set $\eta > 0$, then buyers have an incentive to exert full attentional effort, in which case low quality sellers have a profitable deviation to $\eta = 0$.

Note that for $\kappa > 0$, there will never be separating prices because buyers would not exert any attentional effort, giving low quality sellers an incentive to deviate up to charging the high price.

For $\kappa > 0$, I will examine three cases separately: (1) when uninformed buyers purchase given "pooling high", which is when both sellers charge a high price with probability 1, (2) when uninformed buyers do not purchase given pooling high, and (3) when uninformed buyers are indifferent between purchasing or not given "pooling high", which is when both sellers charge a high price with probability 1.

Starting with case 1, where uninformed buyers purchase given pooling high, the first three regions appear in order as κ increases. The threshold κ^* between regions 2 and 3 is the value of κ at which low quality sellers are indifferent between pricing low and pricing high given conditional demands for pooling high. Because there is no region 4, the upper limit on region 3 is $\kappa^{**} = \infty$.

In region 2, where $\kappa \in (0, \kappa^*)$, the unique value of η in the mimic high equilibrium is

$$\eta = \frac{\lambda}{1-\lambda} \frac{s^0_{p_H} + s^1_{p_H} - s^0_{p_H} s^1_{p_H} - 1}{s^0_{p_H} s^1_{p_H} + s^0_{p_H} \frac{p_L}{p_H} - s^1_{p_H} \frac{p_L}{p_H} - s^0_{p_H}},$$

which converges to 0 as κ goes to zero and to 1 as κ goes to κ^*. For this value of η, low quality sellers have an incentive to mix, which is only true when $d^{\theta_L}_{p_H} = \frac{p_L}{p_H}$. For this to hold, it must be that

$$\frac{\left(1 - s^1_{p_H}\right) \frac{\beta_{p_H} - s^0_{p_H}}{s^1_{p_H} - s^0_{p_H}}}{\left(1 - \beta_{p_H}\right)} = \frac{p_L}{p_H},$$

which implies that

$$\beta_{p_H} = \frac{s^0_{p_H} + s^1_{p_H} - s^0_{p_H} s^1_{p_H} - 1}{s^0_{p_H} s^1_{p_H} + s^0_{p_H} \frac{p_L}{p_H} - s^1_{p_H} \frac{p_L}{p_H} - s^0_{p_H}}.$$

To find η, it just remains to note that the interim probability that a seller is of high quality at the high price is

$$\beta_{p_H} = \frac{\lambda}{\lambda + \eta\left(1 - \lambda\right)}.$$

To show that no other value of η supports the mimic high equilibrium in this region, it is enough to show the existence of a single crossing property for $\frac{p_L}{p_H}$ and $d^{\theta_L}_{p_H}$ as β_{p_H} increases. When $\kappa \in (0, \kappa^*)$, if $\beta_{p_H} = \lambda$, then $d^{\theta_L}_{p_H} < \frac{p_L}{p_H}$. Also, for some $\beta_{p_H} > \lambda$, $d^{\theta_L}_{p_H} > \frac{p_L}{p_H}$. Thus, because $d^{\theta_L}_{p_H}$ is strictly increasing in β_{p_H}, there exists a single β_{p_H} where $d^{\theta_L}_{p_H} = \frac{p_L}{p_H}$.

In this region, $d^{\theta_L}_{p_H}$ is also strictly increasing with information processing cost, so the distance to the crossing point decreases with cost. This shows why there is a decrease in mimicking as cost rises in this region.

For $\kappa \in [\kappa^*, \infty)$, both types charge a high price, so that $\eta = 1$. In this region, $\beta_{p_H} = \lambda$, so $d^{\theta_L}_{p_H} > \frac{p_L}{p_H}$. As a result, neither low nor high quality sellers will deviate to charging a lower price.

Case 2 is much like case 1, except that there is also a region 4 of the equilibrium, and region 3 does not necessarily appear. If region 3 appears, then the threshold κ^* between regions 2 and 3 is the lower value of κ at which low quality sellers are indifferent between pricing low and pricing high given conditional demands for pooling high and κ^{**} is the upper value (if just one value then $\kappa^* = \kappa^{**}$). If low quality sellers always prefer pricing low given conditional demands for pooling high, then region 3 does not appear, but region 4 does appear at the point where $d^{\theta_L}_{p_H}$ switches from being strictly increasing to being strictly decreasing.

In case 3, just regions 1 and 2 appear, so that $\kappa^* = \infty$. This is because $d^{\theta_L}_{p_H}$ is increasing in cost and reaches $\frac{p_L}{p_H}$ in the limit. ■