ECON 303: PRICE THEORY III (INFORMATION ECONOMICS) PROBLEM SET 1

ESHAN EBRAHIMYEKHORRAMABADY, HAYS GOLDEN, AND PRISCILLA MAN

1. Exercise 8.2, Jehle and Reny

Assume wealth is not observed by the insurance company. (Or else accident probabilities will become observable and we are back to the symmetric information case.) Define "adverse selection" as the situation in which the expected accident probabilities *conditional on buying insurance* is increasing in p.

Given w, a consumer purchase L at p if

$$f(w) \geq \frac{u(w) - u(w-p)}{u(w) - u(w-L)}.$$

Denote the right hand side of the above equation as

$$h(p, w) = \frac{u(w) - u(w-p)}{u(w) - u(w-L)}.$$

If we differentiate h with respect to w, the derivative has the same sign as

$$\frac{u'(w) - u'(w-p)}{u(w) - u(w-p)} - \frac{u'(w) - u'(w-L)}{u(w) - u(w-L)}. \tag{1}$$

In general, we cannot sign expression (1). It depends on how the level of risk aversion changes as wealth changes. Intuitively, level of risk aversion would affect how willing consumers are willing to buy insurance as their wealth changes. For example, if consumers have decreasing (relative) risk aversion, they care less about risk as their wealth increases. They might demand less insurance as their wealth increases even if their accident probabilities are increasing in wealth.

Date: May 16, 2009.

This suggested solution is an improvement upon Priscilla's 2008 solution, which contains many errors. As evidence has indicated that Priscilla cannot be counted upon to catch the mistakes her solution, you are warned against errors in this version as well. For the record, the 2008 solution was based on an even earlier version by Felipe Mardones. Of course we do not blame the remaining errors on him.

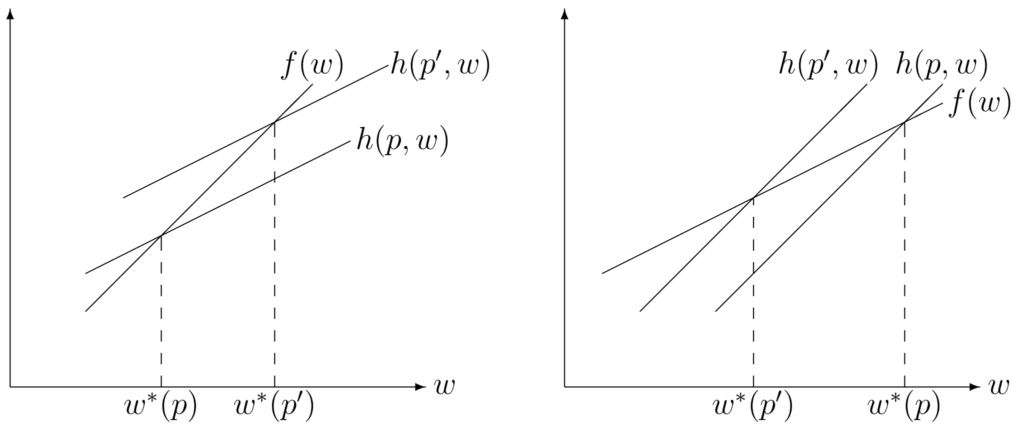

(a) Adverse Selection (b) Favorable Selection

FIGURE 1. Exercise 8.2: Two Possibilities

Consider some premium p. Define the cutoff level w^* such that

$$f(w^*) = h(p, w^*). \tag{2}$$

There could be multiple w^* satisfying this equality for a fixed p. As we will only be illustrating the possibility of favorable selection (the opposite of adverse selection), we focus on the case in which such w^* is unique.

Define $w^*(p)$ as the solution to equation (2) for premium p. Differentiating (2) with respect to p we get

$$f'(w^*)\frac{\partial w^*}{\partial p} = h_p(p, w^*) + h_w(p, w^*)\frac{\partial w^*}{\partial p}$$

where subscripts to h indicates partial derivatives with respect to the relevant variables. Rearranging we get

$$\frac{\partial w^*}{\partial p} = \frac{h_p(p, w^*)}{f'(w^*) - h_w(p, w^*)}. \tag{3}$$

Notice that $h_p(p, w^*) > 0$. Figure 1 exhibits two possibility of the sign of $\frac{\partial w^*}{\partial p}$:

Case 1: Suppose $f'(w^*) > h_w(p, w^*)$. This could happen if h is weakly decreasing in w or when h increases in w but not as fast as f does. This is given by Figure 1(a). In this case, consumer with wealth $w \geq w^*$ (within some neighborhood of w^*) buys insurance.

By equation (3), $\frac{\partial w^*}{\partial p} > 0$. Therefore, for $p' > p$ (within some neighborhood of p),

$$\begin{aligned} \mathrm{E}[f(w)|f(w) \geq h(p', w)] &= \mathrm{E}[f(w)|w \geq w^*(p')] \\ &> \mathrm{E}[f(w)|w \geq w^*(p)] \\ &= \mathrm{E}[f(w)|f(w) \geq h(p, w)]. \end{aligned}$$

In this case we have adverse selection.

Case 2: Suppose $f'(w^*) < h_w(p, w^*)$. This could happen if h is increasing faster than f. This case is given by Figure 1(b). In this case, consumer with wealth $w \leq w^*$ (within some neighborhood of w^*) buys insurance. By equation (3), $\frac{\partial w^*}{\partial p} < 0$. Therefore, for $p' > p$ (within some neighborhood of p),

$$\begin{aligned} \mathrm{E}[f(w)|f(w) \geq h(p', w)] &= \mathrm{E}[f(w)|w \leq w^*(p')] \\ &< \mathrm{E}[f(w)|w \leq w^*(p)] \\ &= \mathrm{E}[f(w)|f(w) \geq h(p, w)]. \end{aligned}$$

We have *favorable selection* instead.

You may be unhappy as the p we considered above is not necessarily the equilibrium price. However, if p^* is the equilibrium price, the above analysis still holds in a neighborhood around p^*.

2. Exercise 8.3, Jehle and Reny

Suppose by contradiction that there are two consumers i and j such that $\pi_i = \pi_j$ but $p_i > p_j$. Since profit must be non-negative for all insurance policies,

$$p_i > p_j \geq \pi_j L = \pi_i L$$

so the insurance company would like to sell as much benefit to i as possible. Yet this cannot be an equilibrium.

3. Exercise 8.4, Jehle and Reny

Part (a). Let the set of possible realization of accident probability be

$$\Pi = \{\pi_1, \pi_2, \ldots, \pi_n\}$$

and index the probabilities such that

$$0 \leq \underline{\pi} = \pi_1 < \pi_2 < \cdots < \pi_n = \overline{\pi} \leq 1.$$

The probability mass function is given by

$$f(\pi) = \begin{cases} f_i & \text{if } \pi = \pi_i,\ i = 1, \ldots, n; \\ 0 & \text{otherwise.} \end{cases}$$

As in class, define

$$h(p) = \frac{u(w) - u(w-p)}{u(w) - u(w-L)} \text{ and;}$$
$$g(p) = \mathrm{E}[\pi | \pi \geq h(p)]L.$$

We will derive an explicit expression for $g(p)$ in this finite case. To do so we introduce a few notations. First of all, define

$$p_0 = 0$$
$$p_i = \min\{h^{-1}(\pi_i), \overline{\pi}L\} \text{ for all } i = 1, \ldots, n.$$

These p_i's are well-defined since h is a continuous, strictly increasing function. (The minimum operator is needed to ensure $p_i \in [0, \overline{\pi}L]$.)

For any $p > p_1$, define $k^*(p)$ such that

$$\pi_{k^*(p)-1} < h(p) \leq \pi_{k^*(p)}.$$

(Notice that the first inequality is strict while the second one is weak. This is because the condition in $g(p)$ is $\pi \geq h(p)$. With a discrete state space, it is important to be careful about which inequality is strict and which one is weak.)

Now we can write $g(p)$ as

$$g(p) = \begin{cases} \sum_{i=1}^{n} \pi_i f_i & \text{if } p \in [0, p_1]; \\ \frac{\sum_{i=k^*(p)}^{n} \pi_i f_i}{\sum_{i=k^*(p)}^{n} f_i} & \text{if } p \in (p_1, \overline{\pi}L]. \end{cases}$$

First we show that g is a step function. Notice that $0 = p_0 \leq p_1 \leq \cdots \leq p_n = \overline{\pi}L$ partition $[0, \overline{\pi}L]$, the domain of g into n intervals. If $p, p' \in (0, p_1)$, by definition $g(p) = g(p')$. For any $i = 1, \ldots, n-1$, any $p, p' \in (p_i, p_{i+1})$,

$$k^*(p) = k^*(p').$$

Hence $g(p) = g(p')$. Thus there exists a partition of the domain of g such that g takes on only one value in each partition, meaning that g is a step function. (If you want, you can show that g is left-continuous, due to the way we define $k^*(p)$.)

Next we show that g is weakly increasing (in p). Consider $p' > p$. Since h is increasing, $k^*(p') \geq k^*(p)$. Denote $k = k^*(p)$ and $k' = k^*(p')$. If $k = k'$, $g(p') = g(p)$. If instead $k' > k$, $g(p') - g(p)$ has the same sign as

$$\begin{aligned}
&\sum_{i=k}^{n} f_i \left(\sum_{j=k'}^{n} \pi_j f_j \right) - \sum_{i=k'}^{n} f_i \left(\sum_{j=k}^{n} \pi_j f_j \right) \\
&= \sum_{i=k}^{k'-1} f_i \left(\sum_{j=k'}^{n} \pi_j f_j \right) - \sum_{i=k'}^{n} f_i \left(\sum_{j=k}^{k'-1} \pi_j f_j \right) \\
&= \sum_{i=k'}^{n} f_i \left[\pi_i \sum_{j=k}^{k'-1} f_j - \sum_{j=k}^{k'-1} \pi_j f_j \right].
\end{aligned}$$

Since i runs from k' to n, so $\pi_i > \pi_j$ for all $j = k, \ldots, k' - 1$. Hence the above expression must be strictly positive. Therefore g is non-decreasing.

Lastly, we show that g maps from $[0, \overline{\pi} L]$ into itself. Notice that $h(0) = 0$ so

$$g(0) = \mathrm{E}[\pi] L \geq 0.$$

Since $\mathrm{E}[\pi | \pi \geq h(p)] \leq \overline{\pi}$ for all p,

$$g(\overline{\pi} L) = E[\pi | \pi \geq h(\overline{\pi} L)] L \leq \overline{\pi} L.$$

As g is weakly increasing, these two inequalities guarantees that g maps from $[0, \overline{\pi} L]$ into itself.

Part (b). We just give a brief intuitive argument here. See Part (c) for a formal proof.

If $g(0) = 0$ or $g(\overline{\pi} L) = \overline{\pi} L$, we get a fixed point trivially. So assume $g(0) > 0$ and $g(\overline{\pi} L) < \overline{\pi} L$. This means g starts above the 45° line and ends below it. The only way that this could happen without g crossing the 45° line is to have a jump *down* at some point. But we have shown that g is non-decreasing so this is impossible. Thus g must have a fixed point.

Part (c). Consider some compact interval $[a, b] \subseteq \mathbb{R}$. Let $f : [a, b] \to [a, b]$ be a non-decreasing function. (This proof will be on $[a, b]$. Take $a = 0$ and $b = 1$ if you prefer working with the unit interval as suggested by the question.)

If $f(a) = a$ and $f(b) = b$ we get a fixed point trivially. So assume $f(a) > a$ and $f(b) < b$. Define

$$S = \{x \in [a, b] : x \leq f(x)\}$$

and let

$$x^* = \sup S.$$

Since S is non-empty ($a \in S$) and bounded (it is a subset of a bounded interval), x^* exists. We establish the following two claims.

Claim 1. $x^* \in S$. That is, $x^* \leq f(x^*)$.

Proof. Consider a sequence $\{x_n\} \in S$ converging to x^*. For all n, $f(x_n) \geq x_n$ (since $x_n \in S$). Taking the limit we get

$$\lim_{n\to\infty} f(x_n) \geq \lim_{n\to\infty} x_n = x^*.$$

The function f is not continuous so we need not have $\lim_n f(x_n) = f(x^*)$. However, $x^* \geq x_n$ for all n (since x^* is the sup of S) and f is non-decreasing. Therefore,

$$f(x^*) \geq \lim_n f(x_n) \geq \lim_n x_n = x^*. \quad \square$$

Claim 2. $x^* \geq f(x^*)$.

Proof. By Claim 1, $x^* \leq f(x^*)$. Since f is non-decreasing,

$$f(x^*) \leq f(f(x^*)).$$

Hence $f(x^*) \in S$. As x^* is the sup (actually, maximum due to Claim 1) of S, $x^* \geq f(x^*)$. $\square$

Combining the two claims we get $x^* = f(x^*)$. Therefore f has a fixed point.

Remark. This proof breaks down if f is not weakly increasing. To convince yourself, consider this example: $f : [0, 1] \to [0, 1]$ is defined by

$$f(x) = \begin{cases} 0.75 & \text{if } x < 0.5; \\ 0.25 & \text{if } x \geq 0.5. \end{cases}$$

Here, $S = [0, 0.5)$, $x^* = 0.5$ but $x^* > f(x^*)$ so Claim 1 no longer holds. It is easy to see that f does not have a fixed point.

4. Exercise 8.6, Jehle and Reny

Part (a). The buyers must get weakly positive expected utility in equilibrium. This means

$$\mathrm{E}[\theta|p] \geq p.$$

If, however, $p <\mathrm{E}[\theta|p]$, all buyers will strictly want to buy a car. Since there are not enough cars to supply all potential buyers, this cannot be an equilibrium. Therefore, $\mathrm{E}[\theta|p] = p$ in equilibrium.

Part (b). Given p, only sellers with $u_s(p, \theta) \geq 0$ will sell. In this case, only sellers with $\theta \leq 2p$ sell. Let $b = \min\{2p, 1\}$, then

$$\begin{aligned}
\mathrm{E}[\theta|p] &= \frac{\int_0^b \theta d\theta}{\int_0^b 1 d\theta} \\
&= \begin{cases} \frac{2p}{2} & \text{if } b = 2p; \\ \frac{1}{2} & \text{if } b = 1 \end{cases} \\
&= \begin{cases} p & \text{if } p \leq \frac{1}{2}; \\ \frac{1}{2} & \text{if } p > \frac{1}{2}. \end{cases}
\end{aligned}$$

Equilibrium condition requires $\mathrm{E}[\theta|p] = p$, so any $p \in \left[0, \frac{1}{2}\right]$ can be an equilibrium price.

Remark. Some may feel uncomfortable having $p = 0$ as an equilibrium since the conditional expectation seems to be undefined (dividing zero by zero). Yet intuitively, if only one single type of seller — those with $\theta = 0$ — are selling, the consumer can easily update their posterior and formulate the conditional expectation of θ to be zero. Notice that

$$\begin{aligned}
\mathrm{E}[\theta|\theta \leq 0] &= \lim_{\varepsilon \downarrow 0} \mathrm{E}[\theta|\theta \leq \varepsilon] \\
&= \lim_{\varepsilon \downarrow 0} \frac{\int_0^\varepsilon \theta d\theta}{\int_0^\varepsilon 1 d\theta}.
\end{aligned}$$

This limit is well-defined (to be 0) by applying L'Hopital Rule.

Part (c). This time only sellers with $\theta \in [0, p^2]$ sell. The conditional expectation of θ is

$$\mathrm{E}[\theta|p] = \frac{\int_0^{p^2} \theta d\theta}{\int_0^{p^2} 1 d\theta} = \frac{p^2}{2}.$$

However, for any $p \in (0,1]$, $\frac{p^2}{2} < p^2 < p$. This means

$$\mathrm{E}[\theta|p] < p \text{ for all } p \in (0,1].$$

Hence the only equilibrium price is zero.

In this equilibrium, the market completely collapsed. Only those cars that nobody is willing to pay a strictly positive price for are "traded".

Part (d). In this case, only sellers with $\theta \in \left[0, p^{\frac{1}{3}}\right]$ sell. The conditional expectation of θ is

$$\mathrm{E}[\theta|p] = \frac{\int_0^{p^{\frac{1}{3}}} \theta d\theta}{\int_0^{p^{\frac{1}{3}}} 1 d\theta} = \frac{1}{2} p^{\frac{1}{3}}.$$

Equilibrium condition means

$$\frac{1}{2} p^{\frac{1}{3}} = p.$$

There are three roots to this equation: a negative one that is irrelevant, 0 and $\frac{1}{2\sqrt{2}}$. Thus there are two equilibria: one with $p = 0$ and another with $p = \frac{1}{2\sqrt{2}} < 1$.

Part (e). It is efficient for a car with type θ to be traded if the total surpluses (buyer's plus seller's) from trade is non-negative. To examine efficiency, we check whether all cars that "should" be traded are traded.

Part (b): Total surpluses from transacting a car with type θ is

$$-\frac{\theta}{2} + \theta = \frac{\theta}{2} > 0 \text{ for all } \theta \in [0,1].$$

The Pareto Optimum is to have all cars traded. This is achieved when the competitive price is $\frac{1}{2}$. All other competitive equilibria are Pareto dominated by this efficient competitive equilibrium.

Part (c): Total surpluses from transacting a car with type θ is

$$\theta - \sqrt{\theta} \leq 0 \text{ for all } \theta \in [0,1].$$

The Pareto Optimum is to have no car traded, which coincides with the only competitive equilibrium. Here, even if the market collapses completely, efficiency is still achieved.

Part (d): Total surpluses from transacting a car with type θ is

$$\theta - \theta^3 \geq 0 \text{ for all } \theta \in [0, 1].$$

The Pareto Optimum is to have all cars traded. Nevertheless, some cars are not traded in both equilibria. They are Pareto dominated by an allocation in which all cars are allocated to the buyers, and a buyer receiving a car with type θ pays the seller (the original owner of the car) θ^3.

5. Exercise 8.7, Jehle and Reny (Optional)

Consider an extensive form game Γ. Let β be a profile of beliefs (a belief for each individual at each information set where the player is to move) and σ be a profile of behavioral strategies in Γ.

Definition. An assessment (β, σ) for an extensive form game Γ is *consistent* if there is a sequence of completely mixed behavioral strategies $\{\sigma^n\} \to \sigma$ such that the associated sequence of Bayes' rule induced belief profiles $\{\beta^n\}$ converges to β.

(Notice that the sequence $\{\sigma^n\}$ needs not be sequentially rational given β^n.)

Consider the insurance signaling game discussed in class. Let

$$\Psi = \{\psi_1, \ldots, \psi_K\}$$

be the set of policies the consumer can choose from. Let I_k be the information set the insurance company is at upon observing an offer ψ_k, for $k = 1, \ldots, K$. Notice that for each k, there are only two histories (nodes) in I_k (either the consumer is a low-risk type or a high-risk type). Define

$$\beta(\psi_k) = \Pr[\text{low risk} | I_k].$$

The insurance company's belief system is completely described by $\{\beta(\psi_k)\}_{k=1}^K$.

Consumer's information sets are all singletons so consistency and Bayes' rule are trivially satisfied there. We will omit them in this exercise.

For strategies, let

$$s_l \in \Delta(\Psi) \text{ be the strategy of the low-risk consumer}$$

$$s_h \in \Delta(\Psi) \text{ be the strategy of the high-risk consumer.}$$

The insurance company's strategy is not important for consistency since it moves last and moves only once. For completeness, however, we will denote it as $\sigma : \times_k I_k \to \{A, R\}$.

An assessment will be written as $\langle (s_l, s_h, \sigma), \{\beta(\psi_k)\}_k \rangle$. We show that an assessment is consistent if and only if it satisfies Bayes' Rule.

If Part. Suppose $\langle (s_l, s_h, \sigma), \{\beta(\psi_k)\}_k \rangle$ satisfies Bayes' Rule. Then for any ψ_k such that $s_l(\psi_k) > 0$ or $s_h(\psi_k) > 0$ we have

$$\beta(\psi_k) = \frac{\alpha s_l(\psi_k)}{\alpha s_l(\psi_k) + (1-\alpha) s_h(\psi_k)}.$$

Construct the following sequence $\{\hat{s}_l^n(\psi_k), \hat{s}_h^n(\psi_k)\}$:

$$\hat{s}_l^n(\psi_k) = \begin{cases} s_l(\psi_k) & \text{if } s_l(\psi_k) > 0; \\ \frac{1}{n} & \text{if } s_l(\psi_k) = 0 \text{ and } \beta(\psi_k) > 0; \\ \frac{1}{n^2} & \text{if } s_l(\psi_k) = 0 \text{ and } \beta(\psi_k) = 0. \end{cases}$$

$$\hat{s}_h^n(\psi_k) = \begin{cases} s_h(\psi_k) & \text{if } s_h(\psi_k) > 0; \\ \frac{1}{n}\left(\frac{\alpha}{1-\alpha}\right)\left(\frac{1-\beta(\psi_k)}{\beta(\psi_k)}\right) & \text{if } s_h(\psi_k) = 0 \text{ and } \beta(\psi_k) > 0; \\ \frac{1}{n} & \text{if } s_h(\psi_k) = 0 \text{ and } \beta(\psi_k) = 0. \end{cases}$$

Then define,

$$S_i^n = \sum_{k=1}^{K} \hat{s}_i^n(\psi_k) \text{ for } i = h, l.$$

(This is to normalized the $\hat{s}_i^n$ so that the probabilities sum to 1). Now define a sequence of completely mixed behavioral strategies $\{s_l^n, s_h^n\}$ defined by

$$s_i^n = \frac{1}{S_i^n} \hat{s}_i^n(\psi_k) \text{ for } i = h, l.$$

By construction, $\{s_l^n, s_h^n\} \to (s_l, s_h)$. The sequence of beliefs $\{\beta^n\}$induced by $\{s_l^n, s_h^n\}$ according to Bayes' Rule is

$$\beta^n(\psi_k) = \frac{\alpha s_l^n(\psi_k)}{\alpha s_l^n(\psi_k) + (1-\alpha)s_h^n(\psi_k)}$$

$$= \begin{cases} \frac{\alpha s_l(\psi_k)}{\alpha s_l(\psi_k)+(1-\alpha)s_h(\psi_k)} & \text{if } s_l(\psi_k),\ s_h(\psi_k) > 0; \\ \frac{\alpha \frac{1}{n}}{\alpha\frac{1}{n}+(1-\alpha)s_h(\psi_k)} & \text{if } s_l(\psi_k) = 0 \text{ and } s_h(\psi_k) > 0; \\ \frac{\alpha s_l(\psi_k)}{\alpha s_l(\psi_k)+(1-\alpha)\frac{1}{n}\left(\frac{\alpha}{1-\alpha}\right)\left(\frac{1-\beta(\psi_k)}{\beta(\psi_k)}\right)} & \text{if } s_l(\psi_k) > 0 \text{ and } s_h(\psi_k) = 0; \\ \frac{\alpha\frac{1}{n}}{\alpha\frac{1}{n}+(1-\alpha)\frac{1}{n}\left(\frac{\alpha}{1-\alpha}\right)\left(\frac{1-\beta(\psi_k)}{\beta(\psi_k)}\right)} & \text{if } s_l(\psi_k),\ s_h(\psi_k) = 0 \text{ and } \beta(\psi_k) > 0; \\ \frac{\alpha\frac{1}{n^2}}{\alpha\frac{1}{n^2}+(1-\alpha)\frac{1}{n}} & \text{if } s_l(\psi_k),\ s_h(\psi_k) = 0 \text{ and } \beta(\psi_k) = 0; \end{cases}$$

$$= \begin{cases} \frac{\alpha s_l(\psi_k)}{\alpha s_l(\psi_k)+(1-\alpha)s_h(\psi_k)} & \text{if } s_l(\psi_k),\ s_h(\psi_k) > 0; \\ \frac{\alpha \frac{1}{n}}{\alpha\frac{1}{n}+(1-\alpha)s_h(\psi_k)} & \text{if } s_l(\psi_k) = 0 \text{ and } s_h(\psi_k) > 0; \\ \frac{s_l(\psi_k)}{s_l(\psi_k)+\frac{1}{n}\left(\frac{1-\beta(\psi_k)}{\beta(\psi_k)}\right)} & \text{if } s_l(\psi_k) > 0 \text{ and } s_h(\psi_k) = 0; \\ \frac{1}{1+\left(\frac{1-\beta(\psi_k)}{\beta(\psi_k)}\right)} & \text{if } s_l(\psi_k),\ s_h(\psi_k) = 0 \text{ and } \beta(\psi_k) > 0; \\ \frac{\alpha\frac{1}{n}}{\alpha\frac{1}{n}+(1-\alpha)} & \text{if } s_l(\psi_k),\ s_h(\psi_k) = 0 \text{ and } \beta(\psi_k) = 0; \end{cases}$$

$$= \begin{cases} \beta(\psi_k) & \text{if } s_l(\psi_k),\ s_h(\psi_k) > 0; \\ \frac{\alpha \frac{1}{n}}{\alpha\frac{1}{n}+(1-\alpha)s_h(\psi_k)} & \text{if } s_l(\psi_k) = 0 \text{ and } s_h(\psi_k) > 0; \\ \frac{s_l(\psi_k)}{s_l(\psi_k)+\frac{1}{n}\left(\frac{1-\beta(\psi_k)}{\beta(\psi_k)}\right)} & \text{if } s_l(\psi_k) > 0 \text{ and } s_h(\psi_k) = 0; \\ \beta(\psi_k) & \text{if } s_l(\psi_k),\ s_h(\psi_k) = 0 \text{ and } \beta(\psi_k) > 0; \\ \frac{\alpha\frac{1}{n}}{\alpha\frac{1}{n}+(1-\alpha)} & \text{if } s_l(\psi_k),\ s_h(\psi_k) = 0 \text{ and } \beta(\psi_k) = 0; \end{cases}$$

$$\to \begin{cases} \beta(\psi_k) & \text{if } s_l(\psi_k),\ s_h(\psi_k) > 0; \\ 0 & \text{if } s_l(\psi_k) = 0 \text{ and } s_h(\psi_k) > 0; \\ 1 & \text{if } s_l(\psi_k) > 0 \text{ and } s_h(\psi_k) = 0; \\ \beta(\psi_k) & \text{if } s_l(\psi_k),\ s_h(\psi_k) = 0 \text{ and } \beta(\psi_k) > 0; \\ 0 & \text{if } s_l(\psi_k),\ s_h(\psi_k) = 0 \text{ and } \beta(\psi_k) = 0; \end{cases}$$

$$= \beta(\psi_k).$$

Therefore $\langle (s_l, s_h, \sigma), \{\beta(\psi_k)\}_k \rangle$ is consistent. $\square$

Only if Part. Suppose $\langle (s_l, s_h, \sigma), \{\beta(\psi_k)\}_k \rangle$ is consistent. Then there exists a sequence of completely mixed strategies $\{s_l^n, s_h^n\}$ such that

(1) $\{s_l^n, s_h^n\} \to (s_l, s_h)$; and

(2) The sequences of beliefs $\{\beta^n\}$ defined by

$$\beta^n(\psi_k) = \frac{\alpha s_l^n(\psi_k)}{\alpha s_l^n(\psi_k) + (1-\alpha) s_h^n(\psi_k)}$$

converges to $\beta(\psi_k)$ for all k.

If $s_l(\psi_k) = s_h(\psi_k) = 0$, Bayes' Rule places no restriction on $\beta(\psi_k)$ so $\beta(\psi_k)$ satisfies Bayes' Rule trivially.

If $s_l(\psi_k) > 0$ or $s_h(\psi_k) > 0$, then the function

$$\frac{\alpha s_l^n(\psi_k)}{\alpha s_l^n(\psi_k) + (1-\alpha) s_h^n(\psi_k)}$$

is continuous at $(s_l^n(\psi_k), s_h^n(\psi_k)) = (s_l(\psi_k), s_h(\psi_k))$ (since the denominator is not zero). Hence

$$\begin{aligned}\beta(\psi_k) &= \lim_{n\to\infty} \frac{\alpha s_l^n(\psi_k)}{\alpha s_l^n(\psi_k) + (1-\alpha) s_h^n(\psi_k)} \\ &= \frac{\alpha s_l(\psi_k)}{\alpha s_l(\psi_k) + (1-\alpha) s_h(\psi_k)}.\end{aligned}$$

This means Bayes' Rule is satisfied. $\square$