Supplemental Material for: Simon Board, Jay Lu. 2018. "Competitive Information Disclosure in Search Markets."
Journal of Political Economy 126(5). DOI: 10.1086/699211.

Supplementary Appendix for "Competitive Information Disclosure in Search Markets"

This supplementary appendix uses the notation and definitions established in the main paper. In this appendix, lemmas and propositions are numbered S.1, S.2, etc.. Numbers without the prefix S refer to those in the main paper.

S.1 Monotonicity of Equilibria in Search Costs when $\alpha = 0$

In this section, we show that the set of simple equilibria is decreasing in the search cost c when beliefs are observable ($\alpha = 0$). First, we need a generalization of Lemma 2. The following lemma says that for any simple equilibrium, there exists a one-shot equilibrium that achieves the same profits. Note that the definitions of absorbing beliefs and one-shot strategies from the main text are naturally extended to simple strategies under $\alpha = 0$.

Lemma S.1. *Suppose $\alpha = 0$. Given any simple equilibrium, there exists a one-shot equilibrium with the same profits.*

Proof. Let $\alpha = 0$ so we can set $\mathbf{y} = \mathbf{p}$ for all $\mathbf{p} \in \Delta\Omega$. Let σ be an equilibrium simple strategy so we can just work with payoff beliefs and let $Q := Q_c(\sigma)$ be the stopping set for σ. By Lemma 5, we know that $K_p^\sigma(Q) = 1$ for all $p \in \Delta S$. We first show that under strategy σ, buyers will only end up at stopping beliefs where the seller is indifferent to providing no information. Formally, this means that if we denote the indifference set by

$$R := \left\{ q \in Q \;\middle|\; \int_Q \pi(r) K_q^\sigma(dr) = \pi(q) \right\},$$

then $K_p^\sigma(R) = 1$. Suppose otherwise, that is, $K_p^\sigma(Q \backslash R) > 0$ for some $p \in \Delta S$. Consider the deviation of providing a second round of information according to σ to all buyers who would have ended up outside R under σ_p. Call this new strategy τ. The seller's profit from using τ is

$$\begin{aligned}\int_Q \pi(q) K_p^\tau(dq) &= \int_R \pi(q) K_p^\sigma(dq) + \int_{Q\backslash R}\left(\int_{\Delta\Omega} \pi(r) K_q^\sigma(dr)\right) K_p^\sigma(dq) \\ &> \int_R \pi(q) K_p^\sigma(dq) + \int_{Q\backslash R} \pi(q) K_p^\sigma(dq) = \int_Q \pi(q) K_p^\sigma(dq)\end{aligned}$$

where the strict inequality follows from the definition of R and our premise that $K_p^\sigma(Q\backslash R) > 0$. Since this contradicts the fact that σ is an equilibrium, it must be that $K_p^\sigma(R) = 1$.

We now construct a one-shot independent strategy as follows. Let $\tilde{\sigma}$ be the simple strategy that

Supplemental Material for: Simon Board, Jay Lu. 2018. "Competitive Information Disclosure in Search Markets."
Journal of Political Economy 126(5). DOI: 10.1086/699211.

gives no information when the seller is indifferent but otherwise follows σ, that is,

$$\tilde{\sigma}_p := \begin{cases} \sigma_p & \text{if } p \notin R \\ \underline{\tau} & \text{if } p \in R \end{cases}$$

where $\underline{\tau}$ is the no information signal structure. We now show that $\tilde{\sigma}$ is one-shot. First note that by definition, every belief $p \in R$ is absorbing under $\tilde{\sigma}$, that is, $R \subset A^{\tilde{\sigma}}$. Hence all we need to show is that $\tilde{\sigma}$ only sends buyers to beliefs in R. If $p \notin R$, then $\tilde{\sigma}_p = \sigma_p$ and we know from above that σ_p sends buyers to beliefs in R. On the other hand, if $p \in R$, then $\tilde{\sigma}_p$ provides no information so the buyer's belief remains in R. This means that $\tilde{\sigma}$ only sends buyers to beliefs in R, and since all beliefs in R are absorbing, $\tilde{\sigma}$ must be one-shot.

We now verify that $\tilde{\sigma}$ is an equilibrium as well. First, we claim that both σ and $\tilde{\sigma}$ have the same stopping sets, that is, $Q = Q_c(\tilde{\sigma})$. Since $\tilde{\sigma}$ provides less information than σ, by Blackwell ordering, it follows that $V_c(\tilde{\sigma}, p) \leq V_c(\sigma, p)$ so $Q \subset Q_c(\tilde{\sigma})$. Now, suppose $p \in Q_c(\tilde{\sigma}) \backslash Q$. Since $p \notin Q$, and $R \subset Q$ by definition, p cannot be in R so $\sigma_p = \tilde{\sigma}_p$ by the definition of $\tilde{\sigma}$. Note that by Lemma 5, buyers stop immediately under σ and hence under $\tilde{\sigma}$ as well (by the definition of $\tilde{\sigma}$). Thus, the value functions under both policies must satisfy

$$V_c(\sigma, p) = -c + \int_{\Delta S} q \cdot u(q) K_p^{\sigma}(dq)$$

This means that $V_c(\tilde{\sigma}, p) = V_c(\sigma, p)$ so $p \in Q$ yielding a contradiction. This proves that $Q_c(\tilde{\sigma}) = Q$.

Finally, we show that there are no profitable deviations from $\tilde{\sigma}$. First, note that if $p \notin R$, then $\tilde{\sigma}_p = \sigma_p$ so profits under both strategies are the same. On the other hand, if $p \in R$, then $\tilde{\sigma}_p$ provides no information by definition. Since $R \subset Q = Q_c(\tilde{\sigma})$, the profit under $\tilde{\sigma}_p$ is $\pi(p)$ while the profit under σ_p is also $\pi(p)$ by the definition of R. Hence, $\tilde{\sigma}$ yields the same profit as σ. Since their stopping sets are the same and σ is an equilibrium, $\tilde{\sigma}$ must be an equilibrium as well. □

We now show that the set of simple equilibria increases as search costs decrease. Intuitively, a decrease in costs does not change the profit from the previous equilibrium strategy but does make it harder to deviate. Hence the strategy remains an equilibrium.

Proposition S.1. *Suppose $\alpha = 0$. Given any simple equilibrium for $c > 0$, there is an equilibrium for all $c' < c$ with the same profits.*

Proof. Given an equilibrium for $c > 0$, Lemma S.1 implies there is a one-shot equilibrium strategy σ with the same profits. We show that σ is also an equilibrium for any $c' < c$. Since σ is simple, we work with payoff beliefs only. Let $Q := Q_c(\sigma)$ and $Q' := Q_{c'}(\sigma)$ denote the stopping sets under σ for search costs c and c' respectively. First, observe that the buyer will stop at any absorbing belief, so $A^{\sigma} \subset Q'$. Since σ is one-shot, a seller's profits from strategy σ are thus unaffected by the

Supplemental Material for: Simon Board, Jay Lu. 2018. "Competitive Information Disclosure in Search Markets."
Journal of Political Economy 126(5). DOI: 10.1086/699211.

reduction in cost. Second, the reduction in cost means that buyers' continuation values increase and the stopping set shrinks, that is, $Q' \subset Q$. Since the seller only sends the buyer to beliefs in the stopping set, the scope for deviations is smaller. We now put this together. Since σ is optimal given σ, for any other strategy $\tilde{\sigma}$,

$$\int_{Q'} \pi(q)\, K_p^{\sigma}(dq) = \int_{Q} \pi(q)\, K_p^{\sigma}(dq) \geq \int_{Q} \pi(q)\, K_p^{\tilde{\sigma}}(dq) \geq \int_{Q'} \pi(q)\, K_p^{\tilde{\sigma}}(dq)$$

Thus, σ is is an equilibrium under c' as well. Moreover, σ yields the same profits under c' as under c. □

S.2 Immediate Purchase Equilibria when $\alpha > 0$

In this section, we consider imperfectly observable beliefs ($\alpha > 0$) and derive conditions under which all symmetric Markov perfect equilibria (MPE) involve immediate purchase.

The idea is as follows. Suppose that sellers use a strategy σ that induces some buyers to delay. In the market, there is mass $m > 1$ of buyers with beliefs distributed according to η_0; each period, mass one of buyers with beliefs $\delta_{\mathbf{p_0}}$ enter, and mass one of buyers with beliefs distributed according to η^{S} exit.[1] We first observe that a seller can "replicate" the information provided by the market by sending new buyers directly to their eventual stopping posteriors according to η^{S}. This deviation enables the seller to make the same sales from new buyers as the original strategy made from all buyers, and so is weakly optimal. Second, we note that if there are signals that indicate the buyer is likely to be old (which occurs when the marginal noise distribution on payoff states ξ is normally distributed), then either the replication strategy or a fully-informative strategy will yield a strictly profitable deviation, meaning that all MPE involve immediate purchase.

Lemma S.2. *Suppose there is no coordination ($\beta = 1$) or full coordination ($\beta = 0$). Then in any MPE, selling immediately to new buyers is weakly optimal. Moreover, if $\alpha < 1$, $\beta = 1$ and ξ is normally distributed, then all MPE involve immediate purchase.*

Proof. Consider the case where $\beta = 1$, so that in any MPE, sellers use a simple strategy $\sigma = (\sigma_y)_{y \in Y}$; the case where $\beta = 0$ is parallel with strategies defined on the overall state space Ω instead of S. When sellers use σ, beliefs evolve according to a (martingale) Markov chain on ΔS, where the stopping set $Q_c(\sigma)$ is absorbing. Each period, mass one of buyers enters with beliefs p_0, and mass one of buyers stop with beliefs distributed according to some distribution η^{S}.

We first suppose the seller deviates and "replicates" the market. By the optional stopping theorem,[2] the average of the stopping posteriors (weighted by η^{S}) equals the prior p_0. Hence, a

[1] In any MPE, sellers have beliefs that are consistent with buyers' behavior so η_0 corresponds to the steady state distribution of beliefs in the market.

[2] To apply the theorem, observe that the positive search cost implies that expected search times are bounded.

Supplemental Material for: Simon Board, Jay Lu. 2018. "Competitive Information Disclosure in Search Markets."
Journal of Political Economy 126(5). DOI: 10.1086/699211.

seller can choose to ignore her signal y and use a strategy $\tilde{\sigma}$ independent of y that sends new buyers directly to their eventual stopping posteriors, distributed according to η^{S}, without the need to search for multiple periods. By using $\tilde{\sigma}$, the seller sells to a unit measure of buyers and obtains the same average profit from new buyers as sellers using σ obtain from all buyers. Formally, let $\Pi(\tilde{\sigma}, p) = \int_{Q_c(\sigma)} \pi(p) K_p^{\tilde{\sigma}}(dr)$ denote the seller's profit from using $\tilde{\sigma}$ on buyer p when other sellers use σ. We then have

$$\mathbb{E}\left[\int_{\Delta S} \Pi(\tilde{\sigma}, p)\, \eta_y(dp)\right] \geq \mathbb{E}\left[\Pi(\tilde{\sigma}, p_0)\, \eta_y(p_0)\right] \geq \mathbb{E}\left[\int_{\Delta S} \Pi(\sigma_y, p)\, \eta_y(dp)\right] \tag{1}$$

where the expectation $\mathbb{E}$ is taken over all realizations of y. The first inequality uses the fact that the seller's profit from all buyers exceeds her profit from p_0 buyers alone. The second follows from the definition of $\tilde{\sigma}$.[3] Since σ is an MPE, we know that the optimal profit must weakly dominate the profit from just selling to new buyers for every y. Thus, for almost all y,

$$\Pi(\tilde{\sigma}, p_0)\, \eta_y(p_0) = \int_{\Delta S} \Pi(\sigma_y, p)\, \eta_y(dp) \tag{2}$$

In other words, the replication strategy $\tilde{\sigma}$ is weakly optimal.

Now, suppose $\alpha < 1$ with $\beta = 1$. Take an old buyer with belief $p \neq p_0$ with positive measure.[4] Observe that since buyer p is willing to pay c to search, the probability of him eventually buying some product other than u_0 is strictly positive, meaning the seller can obtain positive profits $\bar{\Pi}(p) > 0$ by releasing full information. Since this is always a feasible strategy, the seller's profits conditional on y must exceed $\bar{\Pi}(p)\, \eta_y(p)$; equation (2) then implies that in any MPE with delay,

$$\Pi(\tilde{\sigma}, p_0)\, \eta_y(p_0) \geq \bar{\Pi}(p)\, \eta_y(p) \tag{3}$$

Now observe that because ξ is normally distributed, the likelihood ratio $\eta_y(p_0) / \eta_y(p) \to 0$ along the vector $y = \lambda(p - p_0)$ as $\lambda \to \infty$. Thus when λ is large enough, equation (3) cannot hold, meaning that there is a profitable deviation. $\square$

The proof of Lemma S.2 places several restrictions on parameters. In the case where buyer beliefs are completely private ($\alpha = 1$), immediate purchase MPE are still unique in environments such as Example 1. This is because the stopping set is of the form $Q = [0, a] \cup [b, 1]$, so the replication strategy sells to some old buyers and is a strictly profitable deviation (i.e., the first inequality in equation (1) is strict).[5] The case of imperfect coordination ($\beta \in (0, 1)$) presents more of a problem

[3] If no old buyers have belief p_0, this is an equality.

[4] If old buyers are continuously distributed then one can take a small ball with positive measure.

[5] In the case of Example 1 with no coordination, one can also dispense with the normality assumption on ξ. That is, for any α, all MPE involve immediate purchase. The key observation is that the uniquely optimal way to sell to p_0 buyers is to use a signal $p_0 \to \{0, b\}$, where $b := \inf\left\{ q \in Q_c(\sigma) \middle| q \geq \frac{1}{2} \right\}$. Since this is a feasible deviation, the optimality of σ implies that under any MPE buyers only exit at beliefs $\{0, b\}$. One can use this to show that σ is independent of y, just as if beliefs were private. Beyond Example 1, the problem is that stopping set may be strangely

as sellers using independent signals cannot replicate coordinated signals. To see the issue, note that if a buyer were to go to the market twice, then with probability β^2 he gets two rounds of the independent signal and with $1-\beta^2$ probability he gets some mixture of the coordinating and independent signal. Since $\beta^2 < \beta$ for $\beta \in (0,1)$, there is no way a single seller can replicate this.

shaped making it hard to guarantee that the replication strategy sells to some old buyers.