Competiton in Designing Pandora's Boxes

Pak Hung Au*

March 16, 2018

Abstract

We revisit the classic sequential-search model by Weitzman (1979) in which an agent is presented with a number of boxes containing uncertain prizes, from which he selects one prize after sequentially opening the boxes and learning the prize contained within. In our model, each box is owned by an independent and ex-ante identical designer, whose objective is to get the agent to select the prize of his box. Each designer persuades the agent to pick his prize by designing the prize distribution of his box, subject to the constraints that (i) the expected value of the distribution is a constant, and (ii) the size of prize realizations is bounded. By focusing on symmetric equilibria, we show that for a range of parameters, the game only has mixed-strategy equilibria. We fully characterize the symmetric equilibria, and show that they have a simple linear structure. Moreover, we find that a decrease in the agent's search cost mitigates competition, thus benefitting the designers.

Keywords Sequential Search, Information Design

JEL Codes D83

1 Introduction

Consider a market in which a number of firms compete by choosing their respective product designs. A product's design affects the distribution of each consumer's match value (Johnson and Myatt (2006)). A product with a boarder design has less variance in match values; a typical

*Nanyang Technological University, email: phau@ntu.edu.sg. This is a preliminary draft. Comments are very welcome.

consumer would not have a strong preference or aversion towards it. On the other hand, a product with a more niche design has a higher variance in match values; an individual consumer either likes it a lot or extremely hates it. Consumers know the type of product designs (how board or how niche) adopted by each firm (possibly from firms' images and their previous offers), but they do not know their match values with the products. To learn about the match values, consumers have to engage in sequential search: by paying a search cost, each consumer visits a firm and learns about the attributes of its product. How does competition shape firms' product designs? Would firms adopt a common design or different designs? How does search friction affect the intensity of competition?

To answer these questions, we adapt the classic sequential-search model of Weitzman (1979) by assigning a designer to each of the boxes that the agent can sample. Specifically, our setting consists an agent and a number of box designers. Each box designer can choose the distribution of the prizes contained in his box, subject to two feasibility constraints. First, the sizes of the prizes are bounded. Second, the expected value of the prize distribution is constant at some value π. The agent selects at most one prize from the boxes she opened, and the objective of each designer is to get the agent to select the prize from his own box. The agent knows the distributions chosen by the designers, but not their prize realizations. By paying a search cost, she can open a box and learn its prize realization. The agent can sequentially sample unopened boxes in any order, and stop once she has found a satisfactory prize. We assume that the agent has a relatively high outside option so that she does not open any box with a degenerate prize distribution at π.

Weitzman (1979) provides a full characterization to the optimal search strategy of the agent, given the prize distributions of the boxes. Each prize distribution defines a reservation value, and the agent samples boxes in descending order of the reservation values. She stops and claims the best prize discovered if it exceeds the reservation values of all unopened boxes and her outside option. If none of the discovered prizes and reservation values of unopened boxes exceed her outside option, she quits the search altogether without collecting any prize. In his model, the prize distributions of the boxes are exogenously given; we endogenize these distributions by considering the competition among the box owners/designers.

The basic tradeoff facing each box designer is as follows. A box designer's expected payoff depends on both the probability that his box is sampled (opened) by the agent, and the probability that the prize of his box is eventually selected by the agent. Increasing the former probability calls for a more risky distribution (i.e., a distribution with higher likelihoods of extreme prizes), which is akin to a niche product design in the opening example. Conversely, increasing the latter probability calls for a less risky distribution, which is akin to a board design in the opening example.

At first glance, choosing a prize distribution is not a straightforward problem. While offering a prize above the reservation value (and the agent's outside option) ensures immediate acceptance by the agent upon its realization, offering a prize below the reservation value may either lead to immediate acceptance, delayed acceptance (after the agent opens some other boxes and recalls this previous prize), or eventual rejection. Complicating the problem is the fact any positive measure assigned to any prize offer would affect the box's reservation value, and hence the order and the probability of his box being sampled.

Our first result shows that the game has a simple pure-strategy equilibrium if competition is sufficiently intense. In that equilirbium, all designers choose the maximum possible reservation value by adopting the "riskiest" distribution. We then proceed to analyze the case where competition is less intense, and provide a complete characterization of all symmetric mixed-strategy equilibria. The key to our characterization is a function that describes the probability that a designer will be sampled (at some point in the agent's search) for different choices of reservation values. We find that in every mixed-strategy equilibrium, this function must exhibit a linear structure. This linear structure in turn implies that an individual designer's payoff function of offering different prizes is also linear. Consequently, any mixed strategy over prize distributions that supports a linear sampling probability function described above constitutes an equilibrium. Moreover, our characterization result shows that the mixed-strategy equilibrium is unique up to the sampling probability function. This in turn implies that the designer's equilibrium payoff is unique.

The essential uniqueness result allows us to conclude that there is a cutoff level of competition intensity below which the equilibrium is necessarily in mixed-strategy (and above which the equilibrium is necessarily in pure-strategy). Our model therefore predicts that dispersion in product designs can arise naturally in a sequential-search setting with ex-ante homogenous firms and consumers. This finding is in interesting contrast to models of price competition with sequential search (such as Diamond (1971) and Wolinsky (1986)) which find that the only equilibrium outcome is that all firms charge the same price.

The essential uniqueness of the symmetric equilibrium also allows us to obtain neat comparative statics results about the effect of search friction on designers' equilibrium payoff. We find that a decrease in the agent's search cost would mitigate competition and benefit the designers. Intuitively, when the agent has a lower search cost, she is likely to sample more boxes, so the need to offer a risky distribution to solicit the agent's patronage is lower. This result is in contrast to existing studies in sequential search (for example, Wolinsky (1986)) which find that a lower search cost intensifies price competition, thus hurting the firms. In Section 4, we explain in greater details that this difference is mainly due to the different nature of competition in the two models.

Competition involving information design has been studied in Board and Lu (2017), and Au and Kawai (2017). In Board and Lu (2017), firms control the information that consumers can learn about how well a product fits her needs, and consumers sequentially search for this information. Different from our model, all firms in their model sell an homogeneous product. They show that whether firms can observe the agent's search history has a significant impact on the equilibrium outcome. Our model differs from theirs as designers prize distributions are independent of each other; as such, our model can be interpreted as a market for differentiated products.

Au and Kawai (2017) consider a similar model except that the agent learns all prize realizations simultaneously and costlessly. They show that a linear structure of the payoff function is necessary and sufficient for a pure-strategy equilibrium, and analyze the effect of increasing the number of competitors. We find that in a sequential-search setting, designers necessarily play a mixed strategy when competition is mild. This feature, which can be interpreted as an equilibrium dispersion in product design, is absent in their model. Moreover, we study the effect of search friction on designers' choice of distributions.

2 Model

A (female) agent has access to $N \geq 2$ boxes containing prizes of unknown sizes, from which she can pick one eventually (if any). Each box i is controlled by an independent (male) designer, whose objective is to get the agent to select the prize from his own box. For concreteness, suppose he receives a payoff of 1 if his prize is being selected eventually, and a payoff of 0 otherwise. He maximizes his chance of being selected by designing the distribution of the prizes of his box. We impose only two constraints on the box designers' problem. First, the sizes of the prizes are bounded by some interval normalized to $[0,1]$. Second, the distribution must have an expected value of $\pi \in (0,1)$. We say a distribution is feasible if it satisfies both constraints. The space of all feasible distributions is denoted by Φ and a generic element chosen by designer i is denoted by F_i.

The agent engages in sequential search to learn about the boxes' prizes. Before the search begins, she observes the distributions chosen by all the designers $\{F_i\}_{i=1,2,\ldots,N}$, but not their prize realizations. Then she can proceed to open the boxes in any order. Opening each box requires a fixed search cost c. By opening a box, she learns about the realized prize of the box. She can then decide whether to pick that prize and stop the search, or continue her search by opening more boxes. Recall of any previously discovered prize is allowed. At any stage of her search, she can stop her search without taking any prize from the opened boxes; and in doing so, she collects her outside option of u_0. We assume that her outside option exceeds $\pi - c$.

Asssumption $u_0 > \pi - c$.

The assumption implies that it is never a best response for each designer to offer a degenerate prize distribution at π, as doing so would make the agent discard his box without sampling it.

We look for the subgame-perfect equilibrium in which the designers adopt a common (possibly mixed) strategy. To this end, we need to first compute the agent's optimal search strategy given distributions $\{F_i\}_{i=1,2,\ldots,N}$. This problem has been solved by Weitzman (1979). He shows that the agent's search problem has a remarkably simple solution, which is described below. First the agent calculates the reservation value U_i of each box defined by the equation below:

$$c = \int_{U_i}^{1} (x - U_i)\, dF_i(x)\,. \tag{1}$$

Then he samples the boxes in descending order of reservation values, and stops whenever the highest discovered prize (so far) exceeds the reservation values of all the unopened boxes. To avoid uninteresting complications, we assume that whenever the agent is indifferent between several boxes, she randomizes equally between these boxes. Moreover, she goes for the outside option only if it is strictly optimal to do so.

As the agent's search strategy has been pinned down, our analysis will focus on the designers' choice of prize distribution. We say an equilibrium is symmetric if all designers adopt a common strategy. A useful preliminary observation is that the set of possible reservation values is bounded by the feasibility constraints on F_i.

Lemma 1 *The set of feasible reservation values of a box is in the interval* $\left[\pi - c, 1 - \frac{c}{\pi}\right]$.

The lower bound in the lemma above is obtained by having all the mass of the distribution concentrated at π; whereas the upper bound is obtained by concentrating all the mass at 0 and 1 (in a way that the distribution has an expected value of π). Throughout the paper, we denote the highest feasible reservation value by $\bar{U} \equiv 1 - \frac{c}{\pi}$.

We conclude the model setup by remarking that the choice of prize distribution can be interpreted as choosing a posterior distribution over a binary state. The two feasibility constraints above correspond to the requirements that any posterior probability must be between zero and one, and that the expected value of posterior distribution must coincide with the prior.[1]

3 Characterization of Equilibria

We first show that a symmetric pure-strategy equilibrium exists if and only if competition is sufficiently intense. Moreover, this equilibrium necessarily involves all designers adopting the reservation value $\bar{U}$. We then analyze symmetric mixed-strategy equilibria, and show that they necessarily have a common linear structure to be described below.

Let's begin with symmetric equilibria.

Proposition 1 *The only symmetric pure-strategy equilibrium involves each designer choosing reservation value* $\bar{U}$*. Moreover, such an equilibrium exists if and only if*

$$\frac{1 - (1 - \pi)^N}{N (1 - \pi)^{N-1}} \geq \frac{\pi - c}{u_0}. \tag{2}$$

That is, all designers adopt reservation value $\bar{U}$ *provided that* π*,* u_0*,* c *or* N *is sufficiently high.*

The intuition is quite simple. First, if all other designers adopt a common prize distribution with a reservation value $U < \bar{U}$, then one can do better than following them by choosing a very similar distribution but has a slightly higher reservation value. This increases his chance of being sampled by the agent, but has almost no impact on the probability of eventually being selected by the agent conditional on being sampled. Only if the all other designers offer $\bar{U}$ would such a deviation be infeasible. In this case, the only possible deviation is to offer a lower reservation value, and the consequence is that the designer's box would be sampled only after the agent has

[1] The latter is referred to as the Bayes-plausibility condition in Kamenica and Gentzkow (2011).

sampled all other boxes, and they all give the low prize 0. Such a deviation is unprofitable if being sampled (in the last position) is highly unlikely, i.e., π or N is sufficiently high. It is also unprofitable if u_0 is sufficiently high, in which case the probability of realizing a prize above u_0 is too low, and so is the payoff conditional on being sampled last. Finally, a high search cost c discourages the deviation above because an increase in c lowers the reservation value of every prize distribution. Therefore, to achieve u_0, the designer needs to offer a more risky prize distribution, which has a lower probability of realizing a prize above u_0.

The rest of this section focuses on symmetric mixed-strategy equilibria. A mixed strategy consists of a distribution over reservation value, denoted by $G \in \triangle\left(\left[\pi - c, \bar{U}\right]\right)$, and a prize distribution for each reservation value U, denoted by $F_U(\cdot)$. The following observations path the way for our equilibrium construction and characterization.

Lemma 2 *Let* $\left(G, \{F_U\}_{U \in supp(G)}\right)$ *be a symmetric mixed-strategy equilibrium. Then*

(i) The infimum of the support of G *is* u_0*;*

(ii) G *does not have any atom, except possibly at* $\bar{U}$*; and*

(iii) $\int_U^{\bar{U}} F_{\tilde{U}}(U)\, dG\left(\tilde{U}\right)$ *does not have any atom, except possibly at* 0.

Proof. (i) Let $\underline{U}$ be the infimum of the support of G. If $\underline{U} > u_0$, then the designer that adopts reservation value $\underline{U}$ will be sampled if and only if all other designers' boxes have been sampled and they all revealed a prize of no higher than $\underline{U}$. Thus, the designer can strictly increase his payoff by lowering his reservation value, which would increase the chance that his box reveals a prize above $\underline{U}$. This contradicts that $\underline{U}$ is the lower bound of G. It is immediate that $\underline{U} < u_0$ is impossible, as designers adopting reservation value $\underline{U}$ would have a zero payoff.

(ii) If G has an atom at some $U < \bar{U}$, then the expected payoff associated with playing slightly above the atom would strictly exceed that of playing the atom. Let F_U be the prize distribution associated with reservation value U, and $\Pi(F_U)$ be the expected payoff conditional on being sampled. Let $\varepsilon > 0$ and let F' be a mean-preserving spread of F_U such that the expected payoff conditional on being sampled, $\Pi(F') > \Pi(F_U) - \varepsilon$. As F' has a higher reservation value than U, its probability of being sampled is strictly higher than F_U. Therefore, by choosing ε sufficiently small, F' is a profitable deviation from F_U.

(iii) Suppose $\int_U^{\bar{U}} F_{\tilde{U}}(U)\, dG\left(\tilde{U}\right)$ has an atom at some $U \neq 0$. Then a positive measure of reservation values $\tilde{U} > U$ under G has its prize distribution assigning an atom at U. It is immediate that assigning atom to some $U \in (0, u_0)$ is suboptimal, so suppose $U \in [u_0, \bar{U}]$. Conditional on choosing such reservation value $\tilde{U} > U$, an individual designer's expected payoff can be strictly increased by slightly increasing the location of the atom in a way that preserves the mean and reservation value. Specifically, let $\varepsilon > 0$ and the atom size at U be a. One can replace the atom at U with two atoms, one of size $\frac{aU}{U+\varepsilon}$ at $U + \varepsilon$, and the other of size $\frac{a\varepsilon}{U+\varepsilon}$ at 0. The atom at $U + \varepsilon$ brings a strictly higher payoff (in expectation conditional on its realization) by no less than $\left(\frac{1}{2}a\right)^{N-1}$, so the deviation is profitable provided that ε is sufficiently small. This contradicts the optimality of prize distribution chosen by designers with reservation value on the support of G. ■

Consider the problem facing an individual designer, assuming all other designers use a distribution G of reservation values. Suppose that the designer decides to choose a reservation value $U \in [u_0, \bar{U})$. The probability that the designer's box would be sampled at some point of the agent's search, denoted by $\delta(U)$, is given by

$$\delta(U) = \left(G(U) + \int_U^{\bar{U}} F_{\tilde{U}}(U)\, dG\left(\tilde{U}\right) \right)^{N-1}. \tag{3}$$

To understand equation (3), note that the designer's box would be sampled if and only if no other designers stop the agent's search prior to him being reached. This requires that for each other designer j, either (i) box j's reservation value is below U (which occurs with probability $G(U)$), or (ii) box j's has a higher reservation value but its prize is revealed to be below U (which occurs with probability $\int_U^{\bar{U}} F_{\tilde{U}}(U)\, dG\left(\tilde{U}\right)$).

Denote by $\Pi(p; U)$ the designer's expected payoff of offering a prize p, given the reservation value U of his box. This expected payoff function is related to the sampling probability function $\delta(\cdot)$ as follows:

$$\Pi(p; U) = \begin{cases} 0 & \text{if } p < u_0 \\ \delta(p) & \text{if } p \in [u_0, U) \\ \delta(U) & \text{if } p \in [U, 1] \end{cases}. \tag{4}$$

Equation (4) is quite intuitive. If the designer offers a prize $p > U$ within a box with reservation value U, then he is selected with certainty whenever he is sampled, so $\Pi(p; U) = \delta(U)$. On the other hand, if the designer's prize p is below the box's reservation value U, he will be selected over the prize of another designer j if either one of the following two events happen: (i) p exceeds the reservation value U_j of designer j (which occurs with probability $G(p)$), or (ii) p is below the reservation value U_j but the realized prize p_j of designer j falls below p (which occurs with probability $\int_{p}^{\bar{U}} F_{\tilde{U}}(p)\, dG\left(\tilde{U}\right)$).

Using equation (4), the designer's optimization problem can be stated as follows:

$$\max_{\left\{F\in\Phi: c=\int_{U}^{1}(x-U)dF(x)\right\}} \int_{0}^{1} \Pi(p;U)\, dF(p)\,. \tag{5}$$

Note that the problem above (and hence its solution) depends on the designer's choice of U.

Using the notations and observations above, $\left(G, \{F_U\}_{U\in supp(G)}\right)$ is a symmetric equilibrium if for all U on the support of G (i) F_U solves problem (5); and (ii) $\int_0^1 \Pi(p;U)\, dF_U(p)$ gives the designer's equilibrium payoff. The proposition below states that a specific linear structure on the sampling probability function $\delta(\cdot)$ is necessary and sufficient for a symmetric equilibrium.

Proposition 2 *There exists a pair of values $\alpha \in (0,1]$ and $\hat{U} \in [u_0, \bar{U})$ such that $\left(G, \{F_U\}_{U\in supp(G)}\right)$ is a symmetric mixed-strategy equilibrium strategy if and only if*

(i) G has an atom α at $\bar{U}$;

(ii) G has a support $\left[u_0, \hat{U}\right] \cup \left\{\bar{U}\right\}$;

(iii) for all $U \in supp(G)$, F_U assigns a zero measure to $(0, u_0)$; and

(iv) the implied probability of being sampled by offering reservation U (defined in equation 3) satisfies

$$\delta(U) = \begin{cases} 0 & \text{if } U < u_0 \\ U\frac{(1-\alpha\pi)^{N-1}}{\hat{U}} & \text{if } U \in \left[u_0, \hat{U}\right] \\ (1-\alpha\pi)^{N-1} & \text{if } U \in (\hat{U}, \bar{U}) \\ \frac{1-(1-\alpha\pi)^{N}}{N\alpha\pi} & \text{if } U = \bar{U} \end{cases}. \tag{6}$$

According to Proposition 2, while there can be multiple symmetric equilibria, these equilibria are necessarily quite similar. First, they all share a common support, as well as a common atom at the top, in the reservation-value distribution. Second, the sampling probability $\delta(U)$ are common to all equilibria.

In the proof of the proposition, we first show that the symmetric equilibrium payoff is unique. This consequently pins down the unique atom α of reservation-value distribution G. We use this fact to show that any equilibrium payoff function $\Pi(p;U)$ must be linear; for otherwise, the implied payoff would differ from the unique equilibrium payoff. As the sampling probability function $\delta(U)$ is related to $\Pi(p;U)$ according to (4), it is necessary that $\delta(U)$ has a particular linear structure in equilibrium. In the converse direction, we note that when facing a linear payoff function induced by the sampling probability (6), any prize distribution that does not assign any measure on $(0, u_0)$ is optimal. Moreover, reservation values in the interval $\left(\hat{U}, \bar{U}\right)$ is suboptimal, as lowering it to $\hat{U}$ would not affect the probability of being sampled but would strictly increase the payoff conditional on being sampled.

A particularly simple class of strategies involves all designers adopting prize distribution with binary support of the form $\left\{0, \frac{\pi U}{\pi - c}\right\}$, where U is the box's reservation value. The corollary below states that there exists a symmetric equilibrium in which all designers adopt these simple prize distributions.

Corollary 1 *There exists a symmetric mixed-strategy equilibrium in which every designer adopts a prize distribution with a binary support of the form* $\left\{0, \frac{\pi U}{\pi - c}\right\}$*, for some* $U \in \left[u_0, \hat{U}\right] \cup \left\{\bar{U}\right\}$ *and* $\hat{U} \in [u_0, \bar{U})$.

In the equilibrium described by the corollary above, the agent's search behavior is quite simple. As each box contains a lottery that either a positive prize with some probability, she samples the boxes in descending order of the size of the prize. She stops once she receives some positive prize, and collects the outside option if all boxes reveal a zero prize.

Next, by inspecting the equilibrium conditions in Proposition 2, we can show that $\alpha = 0$ and $\hat{U} = u_0$ if and only if inequality (1) holds.

Corollary 2 *If inequality (1) holds, then the unique symmetric equilibrium involves each designer choosing reservation value* $\bar{U}$*. On the other hand, if inequality (1) does not hold, the symmetric equilibrium is necessarily mixed.*

According to Corollary 2, the symmetric equilibria necessarily involve mixed strategies when the agent's outside option u_0, and her search cost c are relatively small, and when there are relatively few competing designers. By interpreting the choice of prize distribution as a firm's choice of product design, Corollary 2 implies that even if all firms are ex-ante identical, dispersion in product design necessarily arises in equilibrium when competition is not too intense. This result stands in interesting contrast to models of price competition with sequential search. The classic Diamond paradox (Diamond (1971)) states that if consumers have to pay a positive search cost to learn about the price a firm charges, the only equilibrium outcome is that all firms charge the monopoly price. Reinganum (1979) shows that price dispersion can arise if firms have different marginal costs. Stahl (1989) shows that equilibrium price dispersion can be obtained if a certain fraction of consumers is assumed to have a zero search cost. Considering fixed-sample search, Burdett and Judd (1983) show that equilibrium price dispersion can arise even if consumers and firms are ex-ante identical. Our analysis shows that when firms compete in product design, dispersion in product design can arise even if consumer's search is sequential, and all firms and consumers are ex-ante identical.

We conclude this section by discussing a related finding in our previous work. Au and Kawai (2017) consider a setting in which the prizes realizations of all boxed are simultaneously and costlessly revealed to the agent, and show that a certain linear structure of payoff function is necessary for an equilibrium. The key difference between the current study and Au and Kawai (2017) is the nature of competition. Specifically, whereas designers in Au and Kawai (2017) compete only in getting selected by the agent, designers in the current setting compete in two fronts: getting sampled by the agent, and getting selected by the agent. This multi-dimensional competition complicates the analysis as a designer can beat another designer not by offering a prize exceeding that of the other designer, but exceeding the reservation value of the other designer (so that the other designer's offer would not be inspected at all). The difference in the nature of competition results in a remarkable difference in the resulting equilibrium: designers in the current setting necessarily play a mixed strategy whenever inequality (1) does not hold; whereas senders in Au and Kawai (2017) can always play a pure strategy in equilibrium. Intuitively, if all designers adopt a common prize distribution with reservation value below $\bar{U}$, it is profitable for an individual designer to deviate to offer a box with a higher reservation value so that the agent would necessarily sample him before other designers. Also, our search setting allows us to investigate the effect of search friction on designers' equilibrium behavior, which we address in the next section.

4 Comparative Statics

In this section, we study how the degree of competition affects the equilibrium outcome. Specifically, we analyze how the equilibrium responds to changes in the agent's search environment, including the number of boxes N, her outside option u_0, and her search cost c.

Our main comparative statics result is as follows.

Proposition 3 *The designer's equilibrium payoff decreases if*

(i) the number of designers N increases;

(ii) the agent's outside option u_0 increases; or

(iii) the agent's search cost c increases.

The increases above are strict whenever (1) does not hold.

Part (i) and (ii) of the proposition are quite intuitive. An increase in the number of designers lowers the chance that each individual designer is selected. An increase in the agent's outside option lowers the likelihood that each designer's realized prize can outperform the outside option. More interestingly, part (iii) of Proposition 3 states that an increase in the agent's search cost would also harm the designers. The intuition is as follows. As the agent's search cost goes up, she is less willing to sample many boxes. As a result, the designers must improve their offers by raising the reservation values of their boxes. However, a high reservation value lowers the likelihood that a good prize realizes, thus hurting the designer's payoff.

Our result that an increase in search cost intensify competition stands in interesting contrast to existing studies in consumer search. In the setting of Wolinsky (1986) and Anderson and Renault (1999), consumers sequentially search for products that give them high match values, and firms compete by choosing prices (rather than the match-value distribution). It is found that an increase in consumers' search cost raises the equilibrium price and thus the firms' profit. Intuitively, a high search cost mitigates competition among firms: knowing that the consumers are less willing to search for competing products in the market, each sampled firm effectively face a larger demand and can afford to charge a higher price.

While both our model and the consumer-search model involve agent/consumer searching sequentially for an satisfactory prize/match value, the nature of competition in the two models are markedly different. In our model, the designers' choices of prize distributions direct the agent's search. Therefore, designers compete through prize distributions, which must be appealing enough (i.e., having a high reservation value) to stand a good chance of being sampled and inspected. On the other hand, consumers in Wolinsky's model conduct a random search, and the prices chosen by firms do not affect consumer's search behavior. As a result, a high search cost is bad news for designers in our setting as the agent is less inclined to search, and the designers must improve their offers in reservation value. On the other hand, a high search cost is good news to firms in Wolinsky's model, as they know that conditional on being sampled, the consumer is less likely to go for a competing product.

We conclude this section by specifying how the reservation-value distribution responds to changes considered in Proposition 3.

Corollary 3 *If either N, u_0, or c increases, the reservation-value distribution has a larger atom α at the top, while its support $\left[u_0, \hat{U}\right] \cup \left\{\bar{U}\right\}$ shrinks.*

Corollary 3 is in accord with the intuition described above. When either the number of designers, the agent's outside option or search cost increases, the competition among the designers intensifies, inducing them to offer more aggressive reservation values for their boxes. Bar-Isaac, Caruana, and Cuñat (2012) find that when search cost decreases, a fraction of firms respond by switching from board designs to niche designs. In contrast, Corollary 3 finds that when search cost decreases, designers respond by lowering the probability of adopting risky distributions.

5 Concluding Remarks

In this paper, we consider a sequential-search setting in which a number of designers compete by choosing prize distributions. When competition is relatively mild, only mixed-strategy equilibria exist, suggesting that dispersion in product designs can naturally arise in equilibrium. We also show that a linear structure of the sampling probability function is necessary and sufficient for symmetric equilibria. This characterization result allows us to conduct comparative statics, and we find that the competition among designers is mitigated if search friction decreases.

There are a couple of promising directions for future research. First, our analysis assumes that the agent's outside option is so high that boxes with a degenerate distribution will not be sampled. Future studies can consider the case where the agent has a lower outside option. While it is clear that the linear structure we identify remains a sufficient condition for an equilibrium, whether it is also necessary requires further investigation. Second, our model only considers competition over product designs, while abstracting away any form of price competition. A model in which firms compete both in price and product design may explain price (and product design) dispersion without assuming heterogenous consumers or firms.

Appendix

Proof of Lemma 1: It is a standard result that if F' is a mean-preserving spread of F, then the reservation value of F' exceeds that of F. Thus, the maximum reservation value is obtained by a distribution that assigns probability π to prize 0 and probability $1-\pi$ to prize 1, as there does not exist any strict mean-preserving spread of such a distribution. Likewise, the minimum reservation value is obtained by a distribution that assigns probability 1 to prize π, as there does not exist any strict mean-preserving contraction of such a distribution.

Next, note that any reservation value in the interval $\left[\pi - c, 1 - \frac{c}{\pi}\right]$ can be attained by some F_i.To see this, consider a distribution function that concentrates all masses on prizes 0 and r, with $\Pr(r) = \frac{\pi}{r}$. The reservation value of such a distribution is $(\pi - c)\frac{r}{\pi}$. By varying r in the interval $[\pi, 1]$, all reservation values in the interval $\left[\pi - c, 1 - \frac{c}{\pi}\right]$ can be attained. Q.E.D.

Proof of Proposition 1: Suppose it is a symmetric pure-strategy equilibrium for each designer to adopt F. Suppose also that the reservation value of F is $U < \bar{U}$. An individual designer's equilibrium payoff is the product of following two terms: (i) the probability of being drawn, and (ii) the expected payoff conditional on being drawn. Given that all designers offer an identical distribution, the probability of being drawn is at most

$$\frac{1}{N} + \frac{1}{N}F(U) + \frac{1}{N}F(U)^2 + ... + \frac{1}{N}F(U)^{N-1} = \frac{1}{N}\frac{1-F(U)^N}{1-F(U)},$$

which is strictly less than 1. Next, denote by $\Pi(p)$ the designer's expected payoff conditional on being drawn and prize p realizing. His expected payoff conditional on being drawn is thus $\int_0^1 \Pi(p)\,dF(p)$.

Let $\varepsilon > 0$ and F' be a mean-preserving spread of F such that $\int_0^1 \Pi(p)\,dF'(p) > \int_0^1 \Pi(p)\,dF(p) - \varepsilon$. By deviating to F', the designer's box is drawn with probability 1 and the expected payoff conditional on being drawn decreases by no more than ε. Therefore, for ε sufficiently small, such a deviation is definitely profitable. As a result, if a symmetric pure-strategy equilibrium exists, its associated reservation value is necessarily the maximum possible, which equals $\bar{U}$.

Next we identify conditions under which it is an equilibrium for all designers to choose a reservation value $U_i = \bar{U}$. The payoff of following the candidate equilibrium strategy is $\frac{1}{N}\left(1-(1-\pi)^N\right)$. The most profitable deviation is to choose a distribution that maximizes the probability of being selected by the agent, conditional on being the last box opened. Note that the conditioning event occurs if and only if all other boxes have a realized prize of 0. Therefore, the deviation maximizes the probability that a prize no less than u_0 realizes, and the corresponding payoff is thus $(1-\pi)^{N-1} \times \frac{\pi - c}{u_0}$. The condition for such a strategy to constitute a symmetric equilibrium is thus inequality (2). It is straightforward to verify that the left-hand side of the inequality is strictly convex in π, equal to 0 if $\pi = 0$, and negative if π is sufficiently small. Therefore, the inequality holds if and only if $\pi \geq \hat{\pi}$ for some $\hat{\pi}$. The effects of an increase in c, u_0, and N are immediate. Q.E.D.

Proof of Proposition 2: We begin with a simple observation that the support of any equilibrium distribution of reservation values takes the form $\left[u_0, \hat{U}\right] \cup \left\{\bar{U}\right\}$ for some $\hat{U} \in [u_0, \bar{U})$.

Lemma 3 *The support of any symmetric-equilibrium reservation-value distribution takes the form* $\left[u_0, \hat{U}\right] \cup \left\{\bar{U}\right\}$ *for some* $\hat{U} \in [u_0, \bar{U})$.

Proof. Suppose I is a maximal open interval in $\left[u_0, \hat{U}\right]$ over which the equilibrium reservation-value distribution assigns a zero measure. Then reservation value $\sup I$ necessarily gives a strictly lower profit than that of $\inf I$, as they are both associated with the same probability of being sampled, but the latter is associated with a strictly higher probability of being selected conditional on being sampled. This is because there is a mean-preserving contraction of $F_{\sup I}$ that allows the designer to receive a higher expected payoff conditional on being sampled than that of $F_{\sup I}$. To see this, one can modify $F_{\sup I}$ as follows. Let $\varepsilon, \varepsilon' > 0$ and take a small positive measure $B \subset \{p : p < \sup I\}$ such that $F_{\sup I}(B) < \varepsilon$ and a small positive measure $A \subset \{p : p \geq \sup I\}$ such that $F_{\sup I}(A) < \varepsilon'$. Now if ε' is large relative to ε, the two measures can be merged to form an atom at some $U' > \inf I$. This strictly increases the probability of being selected conditional on the realization of B without affecting that of A. This contradicts that both $\sup I$ and $\inf I$ are on the support of G. ■

The lemma below establishes that the designer's (ex-ante) payoff in every symmetric equilibrium is identical.

Lemma 4 *The symmetric equilibrium payoff of a designer is unique.*

Proof. Suppose inequality (2) does not hold so that there are only mixed-strategy equilibria. Suppose there are two distinct equilibria $\left(G, \{F_U\}_{U \in supp(G)}\right)$ and $\left(G', \{F'_U\}_{U \in supp(G')}\right)$ with distinct payoffs Π and Π' such that $\Pi' > \Pi$. Denote by $\delta(\cdot)$ and $\delta'(\cdot)$ the respective probabilities of being sampled in these equilibria. By Lemma 2, reservation value u_0 is on the supports of both equilibria. As boxes with reservation value u_0 is sampled if and only if all other boxes have a prize realization below u_0, we have $N\Pi = 1 - (\delta(u_0))^{\frac{N}{N-1}}$. Therefore, $\delta(u_0) > \delta'(u_0)$. On the other hand, it is optimal for a designer with reservation value u_0 to maximize the probability of realization of prizes above u_0, which can be achieved by the prize distribution with binary support $\{0, \frac{\pi u_0}{\pi - c}\}$. Therefore, $\Pi = \delta(u_0)\frac{\pi-c}{u_0}$ and $\Pi' = \delta'(u_0)\frac{\pi-c}{u_0}$, which implies that $\delta'(u_0) > \delta(u_0)$, a contradiction.

Next, if inequality (2) holds, a pure-strategy equilibrium in which all designers choose prize distribution with support $\{0, 1\}$ exists. In this equilibrium, the designers' payoff is $\Pi = \frac{1}{N}\left(1 - (1-\pi)^N\right)$. Suppose there is another mixed-strategy equilibrium that gives a payoff $\Pi' > \Pi = \frac{1}{N}\left(1 - (1-\pi)^N\right)$. Then the same argument as above implies that $\delta'(u_0) < (1-\pi)^{N-1}$, so $\Pi' = \delta'(u_0)\frac{\pi-c}{u_0} < (1-\pi)^{N-1}\frac{\pi-c}{u_0}$. However, by inequality (2), $(1-\pi)^{N-1}\frac{\pi-c}{u_0} \leq \Pi$, a contradiction. Next suppose $\Pi' < \frac{1}{N}\left(1-(1-\pi)^N\right)$. Then $\delta'(u_0) > (1-\pi)^{N-1}$. However, this is impossible as the probability that all $N-1$ boxes fail to deliver a prize above u_0 is maximized at $(1-\pi)^{N-1}$. ■

It follows from Reny (1999) that the game between the designers admit a symmetric equilibrium. Using Lemma 4, there is a unique symmetric-equilibrium payoff, let's denote it by Π^*. Let $\left(G, \{F_U\}_{U \in supp(G)}\right)$ be a symmetric mixed-strategy equilibrium. We show below that it must satisfies conditions (i) to (iv) in the proposition statement.

First, we show that G must have an atom at $\bar{U}$ of size α, implicitly defined by

$$\Pi^* = \frac{1 - (1-\alpha\pi)^N}{N\alpha}. \tag{7}$$

To see this, note that by choosing reservation value $\bar{U}$ (thus offering prize distribution with support $\{0, 1\}$), a designer is competing effectively only against other designers that adopt the same prize distribution. The corresponding payoff is thus

$$\pi \left[\begin{array}{c} (1-\alpha)^{N-1} + (N-1)\alpha(1-\alpha)^{N-2}\frac{1}{2} + \binom{N-1}{2}\alpha^2(1-\alpha)^{N-3}\frac{1}{3} + ... + \alpha^{N-1}\frac{1}{N} \\ + \left((N-1)\alpha(1-\alpha)^{N-2}\frac{1}{2} + \binom{N-1}{2}\alpha^2(1-\alpha)^{N-3}\frac{1}{3} + ... + \alpha^{N-1}\frac{1}{N}\right)(1-\pi) \\ + \left(\binom{N-1}{2}\alpha^2(1-\alpha)^{N-3}\frac{1}{3} + ... + \alpha^{N-1}\frac{1}{N}\right)(1-\pi)^2 + ... + \alpha^{N-1}\frac{1}{N}(1-\pi)^{N-1} \end{array} \right].$$

In the payoff expression above, the first line in the bracket is the probability that the designer is the first one being sampled, and the second line is the probability that he is the second one being sampled, and the other terms can be similarly interpreted. Moreover, conditional on being sampled, he is selected with probability π (the size of atom at prize 1). Straightforward algebra shows that we can simplify the expression above into $\frac{1-(1-\alpha\pi)^N}{N\alpha}$. This proves that all symmetric equilibrium must satisfy condition (i) in the proposition statement, with α implicitly defined in equation (7).

Recall that in a symmetric equilibrium $\left(G, \{F_U\}_{U \in supp(G)}\right)$, the probability $\delta(U)$ of being sampled by offering a box of reservation value U is given by equation (3). We explain below that this function necessarily takes the following form.

$$\delta^*(U) = \begin{cases} 0 & \text{if } U < u_0 \\ \frac{\Pi^*}{\pi - c} U & \text{if } U \in \left[u_0, \hat{U}\right] \\ \frac{\Pi^*}{\pi - c} \hat{U} & \text{if } U \in (\hat{U}, \bar{U}) \\ \frac{1-(1-\alpha\pi)^N}{N\alpha\pi} & \text{if } U = \bar{U} \end{cases}. \tag{8}$$

where α is defined in equation (7), and $\hat{U} = \sup\left(supp(G) \setminus \{\bar{U}\}\right)$.

Suppose first that $\delta(\cdot) = \delta^*(\cdot)$. Define an individual designer's payoff function $\Pi^*(p; U)$ of offering prize p within a box of reservation U using equation (4). Consider the designer's optimization problem (5) for some $U \in \left[u_0, \hat{U}\right]$. We show below that the optimized value is exactly Π^*. The Lagrangian of the problem is

$$L = \int_0^1 \left[\Pi^*(p; U) + \lambda\left(\max\{0, p - U\} - c\right)\right] dF(p), \tag{9}$$

where λ is the Lagrange multiplier. As shown in Kamenica and Gentzkow (2011), given a λ, the problem of maximizing the Lagrangian can be solved by finding the concave closure (in p) of the function $\Pi^*(p; U) + \lambda(\max\{0, p - U\} - c)$. If $\lambda > \frac{\Pi^*}{\pi - c}$, then the Lagrangian is maximized by a binary prize distribution with support $\{0, 1\}$, contradicting that $U < \bar{U}$. On the other hand, if $\lambda < \frac{\Pi^*}{\pi - c}$, then the only distributions that maximize the Lagrangian assign a zero measure to prizes above U, but these distributions give reservation values strictly below U; again a contradiction. Therefore, it is necessary that $\lambda = \frac{\Pi^*}{\pi - c}$. Now with $\lambda = \frac{\Pi^*}{\pi - c}$, the concave closure of of the function $\Pi^*(p; U) + \lambda(\max\{0, p - U\} - c)$ is simply $\frac{\Pi^*}{\pi - c}(p - c)$, so the maximized value of the Lagrangian is simply Π^*. Therefore, the maximized value of problem (5) is Π^*. Note that as $\Pi^*(p; U) + \lambda(\max\{0, p - U\} - c) < \Pi^*$ over the interval $(0, u_0)$, it is strictly suboptimal to assign a positive measure over prizes $(0, u_0)$. Note also that payoff Π^* can be achieved by a prize distribution with binary support $\left\{0, \frac{\pi U}{\pi - c}\right\}$.

We show that it is necessary that $\delta(U) = \delta^*(U)$ for all $U \in \left[u_0, \hat{U}\right]$. Suppose first that $\delta(U) > \delta^*(U)$ for some $U \in \left[u_0, \hat{U}\right]$. Then as $\Pi^* = \delta(U) \int_0^1 \frac{\Pi(p;U)}{\delta(U)} dF_U(p)$, we must have $\int_0^1 \frac{\Pi(p;U)}{\delta(U)} dF_U(p) < \frac{\Pi^*}{\delta^*(U)} = \frac{\pi - c}{U}$. However, this is impossible as a designer offering a reservation value U can always guarantee himself a payoff of $\frac{\pi-c}{U}$ conditional on being sampled, simply by offering a prize distribution with binary support $\left\{0, \frac{\pi U}{\pi - c}\right\}$.

Next suppose $\delta(U) < \delta^*(U)$ for some $U \in \left[u_0, \hat{U}\right]$. Then there must exist a $p \in [u_0, U)$ such that $\Pi(p;U) > \Pi^*(p;U)$ for otherwise, $\int_0^1 \Pi(p;U) dF(U) < \int_0^1 \Pi^*(p;U) dF(U)$ for all F, which implies $\int_0^1 \Pi(p;U) dF(U) < \Pi^*$. By equation (4), $\delta(p) > \delta^*(p)$. Moreover, by Lemma 3, $p \in supp(G)$. This is, however, impossible as shown in the last paragraph.

We show below that $\hat{U}$ is necessarily equal to $\frac{\pi-c}{\Pi^*}(1-\alpha\pi)^{N-1}$. We have shown above that conditional on offering a box with reservation value $U \in \left[u_0, \hat{U}\right]$, an optimal prize distribution is one that has a binary support $\left\{0, \frac{\pi U}{\pi - c}\right\}$, yielding a conditional payoff of $\frac{\pi-c}{U}$. As a box with reservation value $\hat{U}$ has a probability $(1-\alpha\pi)^{N-1}$ of being sampled, the expected payoff of such a box is therefore $(1-\alpha\pi)^{N-1}\frac{\pi-c}{\hat{U}}$. Equating it with Π^* gives $\hat{U} = \frac{\pi-c}{\Pi^*}(1-\alpha\pi)^{N-1}$.

We have thus shown that every symmetric equilibrium must satisfy condition (ii) and (iv) in the proposition statement, with α defined in equation (7), and $\hat{U} = \frac{\pi-c}{\Pi^*}(1-\alpha\pi)^{N-1}$. As the equilibrium sampling probability is $\delta^*(\cdot)$, and thus the payoff function is $\Pi^*(p;U)$, we know that any optimal prize distribution for $U \in \left[u_0, \hat{U}\right]$ must assign a zero measure to prizes $(0, u_0)$, thus proving condition (iii) is necessary in equilibrium.

It remains to show that a strategy $\left(G, \{F_U\}_{U \in supp(G)}\right)$ that satisfies conditions (i) to (iv) is an equilibrium. To this end, it suffices to note that if all other designers adopt strategy $\left(G, \{F_U\}_{U \in supp(G)}\right)$ that satisfies conditions (i) to (iv), it is optimal for a designer to follow as well.

Consider first the choice of reservation values. Consider problem (5) with $\Pi(p;U) = \Pi^*(p;U)$, and $U' \in \left(\hat{U}, \bar{U}\right)$. It has a Lagrangian given by (9). As $U' > \hat{U}$, it is necessary that the Lagrange multiplier λ equals $\frac{\Pi^*}{\pi-c}\frac{1-\hat{U}}{1-U'}$. The optimized payoff conditional on reservation value U' is thus $-\lambda c + \frac{\Pi^*}{\pi-c}\pi = \frac{\Pi^*}{\pi-c}\left(\pi - \frac{1-\hat{U}}{1-U'}c\right) < \Pi^*$.

In contrast, we have already shown above that if $U \in \left[u_0, \hat{U}\right]$, the Lagrange multiplier for problem (5) is $\frac{\Pi^*}{\pi-c}$. The designer's problem is thus equivalent to finding the concave closure of the function

$$\Pi^*(p;U) + \frac{\Pi^*}{\pi - c}\left(\max\{0, p-U\} - c\right) = \begin{cases} -\frac{\Pi^*}{\pi-c}c & \text{if } p < u_0 \\ \frac{\Pi^*}{\pi-c}(p-c) & \text{if } p \in [u_0, 1] \end{cases}.$$

It is thus clear that any prize distribution that delivers a reservation value U and assigns a zero measure to $(0, u_0)$ would achieve an expected payoff of Π^*. Similarly, as α satisfies equation (7), offering reservation value $\bar{U}$ would also bring an expected payoff of Π^*. Therefore, $\left(G, \{F_U\}_{U \in supp(G)}\right)$ is a best response to other designers playing $\left(G, \{F_U\}_{U \in supp(G)}\right)$. Q.E.D.

Proof of Corollary 1: By Proposition 2, it suffices to show that there exists a reservation-value distribution G such that when coupled with the set of prize distributions described in the proposition statement, equation (6) holds. Using equation (3), this means that we need a function G such that for $U \in \left[u_0, \hat{U}\right]$

$$\frac{(1-\alpha\pi)^{N-1}}{\hat{U}} U = \left(G\left(\hat{U}\right) - \int_U^{\hat{U}} \frac{\pi - c}{\tilde{U}} dG\left(\tilde{U}\right) + \alpha(1-\pi)\right)^{N-1}.$$

Differentiating both sides of the equation above with respect to p and rearranging, we get $\frac{dG(U)}{dU} = \frac{1-\alpha\pi}{\pi-c}\frac{1}{N-1}\left(\frac{U}{\hat{U}}\right)^{\frac{1}{N-1}}$. Therefore, G takes the form:

$$G(U) = \frac{1-\alpha\pi}{\pi - c}\frac{1}{N\hat{U}^{\frac{1}{N-1}}}\left[U^{\frac{N}{N-1}} - u_0^{\frac{N}{N-1}}\right]. \tag{10}$$

Q.E.D.

Proof of Corollary 2: By Proposition 2 and 1, it is without loss to focus on prize distribution with binary support $\left\{0, \frac{\pi U}{\pi - c}\right\}$. Recall from the proof of Proposition 2 that the equilibrium payoff of offering a reservation value $\hat{U}$ can be expressed as $\Pi^* = \frac{\pi-c}{\hat{U}}(1-\alpha\pi)^{N-1}$. Substituting $\hat{U} = \frac{\pi-c}{\Pi^*}(1-\alpha\pi)^{N-1}$ into equation (10) and using the fact that $G\left(\hat{U}\right) = 1-\alpha$, we have

$$N\left(\left(\frac{\alpha}{1-(1-\alpha\pi)^N}\right)^{\frac{N}{N-1}} - \left(\frac{\alpha}{1-(1-\alpha\pi)^N}\right)^{\frac{1}{N-1}}\right)^{\frac{N-1}{N}} = \frac{u_0}{\pi - c}. \tag{11}$$

It can be shown that the left-hand side of equation (11) above is increasing and continuous in α,[2] equals 0 when $\alpha = 0$, and equals $N\frac{(1-\pi)^{N-1}}{1-(1-\pi)^N}$ when $\alpha = 1$. Therefore, for any $u_0 \in$

[2] To see the left-hand side of equation is increasing, note that $\frac{\alpha}{1-(1-\alpha\pi)^N}$ is increasing in α, and the left-hand side is increasing in the term $\frac{\alpha}{1-(1-\alpha\pi)^N}$.

$\left(\pi - c, (\pi - c)\,\frac{N(1-\pi)^{N-1}}{1-(1-\pi)^N}\right)$, equation (11) admits a solution in the interval $(0, 1)$. On the other hand, if inequality (2) holds, $u_0 \geq (\pi - c)\,\frac{N(1-\pi)^{N-1}}{1-(1-\pi)^N}$, and equation (11) admits no solution.

Proof of Proposition 3: Using equation (7), equation (11) can be re-written as

$$\left(\frac{1}{\Pi^*} - N\right)^{\frac{N-1}{N}} \left(\frac{1}{\Pi^*}\right)^{\frac{1}{N}} = \frac{u_0}{\pi - c}. \tag{12}$$

It is immediate that the left-hand side of equation (12) is decreasing in Π^*, and is increasing in N. Consequently, the solution of equation in Π^* is decreasing in u_0, c, and N. Q.E.D.

Proof of Corollary 3: Recall the left-hand side of equation (11) is increasing in α. An increase in c therefore increases the solution of equation (11) in α.

Next, recall from the proof of Proposition 2 that $\hat{U} = \frac{\pi-c}{\Pi^*}(1-\alpha\pi)^{N-1}$. Together with equation (7), we have $\hat{U} = N(\pi - c)\,\frac{\alpha(1-\alpha\pi)^{N-1}}{1-(1-\alpha\pi)^N}$. As $sgn\left(\frac{\partial \hat{U}}{\partial \alpha}\right) = sgn\left(\frac{1-(1-\pi\alpha)^N}{N\alpha} - \pi\right) = sgn(\Pi^* - \pi)$ is negative, an increase in c necessarily lowers $\hat{U}$. Q.E.D.

References

[1] Anderson, S.P. and Renault, R., 1999. Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model. RAND Journal of Economics, pp.719-735.

[2] Au, P.H. and Kawai, K., 2017. Competitive information disclosure by multiple senders. Working paper.

[3] Bar-Isaac, H., Caruana, G. and Cuñat, V., 2012. Search, design, and market structure. American Economic Review, 102(2), pp.1140-60.

[4] Board, S. and Lu, J., 2015. Competitive information disclosure in search markets. Journal of Political Economy, forthcoming.

[5] Burdett, K. and Judd, K.L., 1983. Equilibrium price dispersion. Econometrica, pp.955-969.

[6] Diamond, P.A., 1971. A model of price adjustment. Journal of Economic Theory, 3(2), pp.156-168.

[7] Johnson, J.P. and Myatt, D.P., 2006. On the simple economics of advertising, marketing, and product design. American Economic Review, 96(3), pp.756-784.

[8] Kamenica, E. and Gentzkow, M., 2011. Bayesian persuasion. American Economic Review, 101(6), pp.2590-2615.

[9] Reinganum, J.F., 1979. A simple model of equilibrium price dispersion. Journal of Political Economy, 87(4), pp.851-858.

[10] Reny, P.J., 1999. On the existence of pure and mixed strategy Nash equilibria in discontinuous games. Econometrica, 67(5), pp.1029-1056.

[11] Stahl, D.O., 1989. Oligopolistic pricing with sequential consumer search. American Economic Review, pp.700-712.

[12] Weitzman, M.L., 1979. Optimal search for the best alternative. Econometrica, pp.641-654.

[13] Wolinsky, A., 1986. True monopolistic competition as a result of imperfect information. Quarterly Journal of Economics, 101(3), pp.493-511.