Metasearch and market concentration

Renaud Foucart*

November 16, 2017

Abstract

Competing intermediaries search on behalf of consumers among a large number of horizontally differentiated sellers. Consumers either pick the best deal offered by an intermediary, or compare the intermediaries. A higher number of intermediaries has the direct effect of decreasing their search effort, but also increases the incentives for consumers to compare. A higher share of informed consumers in turns increases the search effort of intermediaries. I show that the total effect of a higher number of intermediaries is to make each of them choosier. Moreover, it always decreases the price offered by sellers, so that concentration on the market for intermediaries unambiguously decreases consumer welfare.

1 Introduction

Consumers often rely on intermediaries to help them find the product that suits the best to their needs. In the case of online intermediaries, it is easy - yet costly - for consumers to compare the different recommendations received and pick the best offer. A natural question on this market is whether consumers benefit from having a large numbers of intermediaries at disposal. More precisely, could limiting entry or, to the contrary, mergers of intermediaries increase consumer welfare?

It is well-known that if intermediaries have the possibility to bias their findings, by ranking better sellers that offer them a higher commission, higher market concentration may result in too high commissions and prices.[1] Intermediaries however argue that higher concentration allows offering a better quality of advice by giving them more incentives to invest in the quality of their

*Humboldt University, Berlin, renaud.foucart@hu-berlin.de. I am grateful to Mark Armstrong, Alessandro Bonatti, Alexandre de Corniere, Heiko Karle, Martin Pollrich and Andrew Rhodes for useful comments. Financial support by Deutsche Forschungsgemeinschaft through CRC TRR 190 is gratefully acknowledged.

[1] See Benjamin Elderman, "Impact of OTA bias and consolidation on consumers," Harvard Business School.

algorithms.[2] In this paper, I show that even with truthful reporting, higher market concentration lowers consumer welfare. The reason is that it decreases the incentive a buyer has to compare recommendations, and the positive externalities it generates on the others by doing so.

I set up a model where a large number of consumers want to buy one unit of a product in a market with a large number of horizontally differentiated sellers. Several competing intermediaries, called deal finders, search on behalf of consumers to find the best deal for them. Lower concentration generates a tradeoff between search intensity by the competing intermediaries and the incentives for consumers to become informed. On the one hand, higher market concentration benefits consumers, as each deal finder provides more search effort. On the other hand, more precise deal finders decrease the incentives for consumers to compare their options. This in turns decreases the incentives for deal finders to provide search efforts, but also the incentives for sellers to offer low prices for their products. The reason behind this last effect is that, for a given choosiness of deal finders, the demand from uninformed consumers is less elastic.

A deal finder can be an individual recruitment agency hired to search for job candidates, a real estate agency searching for prospective tenants (for the owner) or properties (for the tenants), an insurance broker, or one of the many "deal finding websites" on the Internet. The question of whether free entry should be granted on these markets is a long-lasting debate. A typical argument to limit entry is that low market concentration can decrease the quality of advice as consumers are unable to perfectly screen the quality of these intermediaries.[3] This point is however much less obvious on the Internet, where consumers are only a few clicks away from comparing their options. Moreover, digital markets seems to often converge towards very concentrated structures making the question of excessive entry less relevant.[4] For instance, in 2015 in the US, Expedia[5] (Expedia.com, tripadvisor.com, orbitz.com, hotels.com, venere.com, trivago.com,...) and Priceline[6] (priceline.com, kayak.com, booking.com,...) controlled 95 percent of the online travel-marketplace after a number of successful fusions and acquisitions.[7]

[2] See for instance, Trefis Team, "Competitive landscape of the U.S. online travel market is transforming, Forbes, April 8, 2014.

[3] For instance in the UK, from 1977 to 2005, the Insurance Broker Registration Act limited entry on the market. Similarly, most US states require a special license to be a recognized broker.

[4] Malik (2015) summarizes the dynamic of concentration on digital markets as consisting of three phases: "The first is when there is a new idea, product, service, or technology dreamed up by a clever person or group of people. For a brief while, that idea becomes popular, which leads to the emergence of dozens of imitators, funded in part by the venture community. Most of these companies die. When the dust settles, there are one or two or three players left standing."

[5] http://www.expediainc.com/expedia-brands/

[6] http://ir.pricelinegroup.com/

[7] See for instance Sun, sea and surfing, The Economist, June 21, 2014 ; Competition is shaking up the online travel market, Forbes, January 5, 2015 and Expedia and Orbitz are merging. Here's what it means for you, Cecila Kang and Brian Fung, The Washington Post, September 16, 2015.

Innovation in search quality is an essential part of the competition in advice markets. To keep the travel example, Andrew Warner of Expedia reports in a 2014 interview[8] that "for a standard trip from LA to New York, Expedia has 65,000,000,000 different combinations of travel for each consumer - given variations in flight times, airlines, car rentals, hotels, offers." Being able to use consumer data to provide the best personalized advice (and beat competitors) is thus a huge and costly challenge, with Expedia claiming to spend £500 million yearly in R&D. Warner describes the objective of such investment as being able to do more than mechanically answering a query and providing the cheapest price. Today's competition in the online travel industry is thus largely based on being able to provide a good individual match to a specific consumer.[9]

Perhaps the main argument in support of concentration in such a market is that competition among deal finders resembles an all-pay-auction (see for instance Baye *et al.*, 1996): each sale benefits one deal finder only, but all bear the cost of providing the search technology. Hence, the higher the number of competitors, the smaller the marginal return from providing a better service. In this paper, I show that this argument is only valid if one takes as granted the behaviour of consumers. The only reason deal finders compete in offering advice of higher quality is to attract those informed buyers that compare their options. If lower market concentration increases the incentives for consumers to do so, it may as well increase the individual search quality offered by deal finders.

I assume that competing deal finders search (at a cost) for the best product to recommend to a specific consumer. I use two standard search models. In the main part of the paper I study a linear random sequential search within a distribution of deals, in the tradition of Wolinsky (1986) and Anderson and Renault (1999). In Appendix B I use a model of non-sequential search à la Burdett and Judd (1983) and show that my results are robust to this alternative setup. Consumers are of two (endogenous) types. Some are "savvy" and pick the best deal among all the deal finders. Some are "non-savvy" and take the best deal offered by a deal finder chosen at random. I borrow this dichotomy from a literature started by Varian (1980) to study price dispersion. In the main part of the paper I use Varian's setting where consumers compare either one or all available options. I show in Appendix C that my results also hold when consumers bear a linear cost of observing an additional deal finder, in the spirit of Burdett and Judd (1983).

I first derive a classic result from this literature: the existence of search externalities (Arm-

[8]"Expedia is investing billions in data to create personalized travel-graphs", Derek du Preez, March 24, 2014, diginomica.com

[9]On this topic, see "Expedia Thinks It Can Help You Find the Dream Vacation You Didn't Know You Wanted", Drake Bennet, Bloomberg Business Week, February 25, 2016 and "How Expedia, Hopper and Skyscanner Use Big Data to Find You the Cheapest Airfares", Isabel Thottam, Paste Magazine, January 16, 2017

strong, 2015). The savvy consumers protect the non-savvy, as deal finders cannot discriminate among types, so that fiercer competition for the savvy types make all consumers better off. I then study the impact of market concentration on the expected price paid by consumers and the expected quality of advice offered. More deal finders on the market have a direct negative effect on the individual search efforts. Hence, for a given share of savvy types, higher concentration can be beneficial in the sense that it protects the less informed consumers by increasing the search effort of each individual deal finder. This effect is not similar for savvy consumers, as those benefit from the larger choice offered by an increase in the number of deal finders. Lower concentration therefore has the indirect effect that savviness matters more, hence increasing the incentives to become informed, partly internalizing the search externality. This indirect effect outweighs the direct one, so that lower concentration actually benefits all consumers. As shown in appendix D, this is not a mechanical consequence of the higher number of deal finders at disposal of savvy consumers. Indeed, as the marginal benefit from observing an additional deal finder is higher when these are less choosy, savviness matters more in the presence of more deal finders even if the choice is whether or not to observe a single additional one.

Both the number and the choosiness of deal finders have an effect on the symmetric equilibrium price offered by the sellers. As lower market concentration makes deal finders more choosy, sellers unambiguously decrease their price. First, sellers have an incentive to decrease the price because the price elasticity of the demand from non-savvy consumers increases in absolute value. Second, the share of the most elastic segment of the market (savvy consumers) in the demand increases, giving sellers incentives to decrease prices.

I make the assumption that the revenue of deal finders depends linearly on the volume of sales. This is the case for instance if they are financed by selling information about buyers on a competitive market for advertising, or if they collect fixed commissions. In practice, deal finders are financed in various ways. Some charge a fixed amount of "administrative fees," others get rewarded by a commission paid by either the buyer or seller (that can be fixed per purchase or per-click, or proportional to the value of the purchase), and finally a part of the revenue of online deal finders comes from advertisement on the website, and from gathering information on the consumers and selling by-products. These sources of revenue are however constrained by the fact that buyers always have the possibility of bypassing the deal finder that made the recommendation in order to directly buy from the sellers. While it would certainly be interesting to compare the different pricing possibilities and the type of contracts allowed between deal finders and sellers,[10] I

[10]For instance, a recent settlement for the hotel industry between booking.com and several European countries rules

want to neutralize the possible negative effects of reporting based on commissions, and hence take the option to model a more general relationship between sales and revenues. I also assume that sellers do not pay to be listed on a specific website, so that deal finders search among all existing sellers.

Related literature:

This paper relates to the literature on advice and delegated search. In the literature on delegated search, it relates to Lewis (2012) and Ulbricht (2016). I share with the latter the assumption that the relationship between the buyer and the deal finders combines a problem of moral hazard (the effort of the deal finder) and of asymmetric information (the fact that the deal finders know the market better). By adding competition on the side of the deal finders, and studying different types of buyers, I have two related yet different problems. The moral hazard problem is partly alleviated by competition among intermediaries. The asymmetric information is modified by the fact that buyers have different and privately known information.

I consider a world were, instead of firms competing for the attention of a customer, to be included in their consideration set (Eliaz and Spiegler, 2011a), consumers rely on intermediaries to make them a recommendation. In the advice market, most of the focus has been on a single intermediary. For instance, Armstrong and Zhou (2011) and Chen and Zhang (2017) study a large number of possible transactions between sellers of a product and the adviser choosing how to present the information to consumers. The question of competition among advisers has been discussed in an extension of Inderst and Ottaviani (2012), who study financial advice and compare the case of competitive advisors to the one of a monopolist. Competition among two search engines is also studied in Section 5 of de Cornière (2016), in a two-sided framework were search engines compete in order to attract both consumers and advertisers by auctioning "keywords." In a related model, Eliaz and Spiegler (2011b) study competition among two search engines. None of these papers however study specific investment to be made by the advisers to improve the quality of their advice. Another difference is that I compare duopoly to more intense competition. As I make the assumption that the market is fully covered, the case of a monopoly would be trivial in my model as such a deal finder would provide no effort and sellers would offer the monopoly price.

I present the setup of the model in the next Section. Section 3 solves the equilibrium search effort, price, and share of savvy consumer. I characterize the welfare effect of market concentration on consumer welfare in section 4. I allow for heterogeneous costs of becoming savvy in Section

out the possibility for the deal finders to force hotels not to offer a lower price for direct bookings (Booking.com in European settlement over hotel prices, Malcolm Moore and Adam Thomson, Financial Times, April 21, 2015). This question has been studied recently by Edelman and Wright (2015).

5. Section 6 concludes.

2 Model setup

A mass 1 of consumers wants to buy a single unit of a particular product which is supplied by a continuum of sellers of mass M at a marginal production cost of zero. Building on the specification of Anderson and Renault (1999), each consumer i has tastes described by a conditional utility function of the form

$$u_{i,j}(p_j) = v - p_j - \varepsilon_{i,j}, \tag{1}$$

if she buys product j at price p_j. The intrinsic valuation of the product v is assumed to be sufficiently high for each consumer to always buy. The parameter $\varepsilon_{i,j}$ is the realization of a random variable with log-concave probability density function $f(\varepsilon)$, cumulative density function $F(\varepsilon)$ and support over $[0,b]$, with $b > 0$. The distribution of ε is common knowledge. The assumption of log-concavity applies to most commonly used density functions (see Caplin and Nalebuff, 1991 and Anderson and Renault, 1999), and is necessary to obtain the results on a symmetric equilibrium price. The random component ε represents the (exogenous) distance between a particular version of the product j and the ideal product given the taste of buyer i. I denote this parameter as the mismatch value.[11]

The economy is composed of three types of players. Sellers offer horizontally differentiated products, for which they individually set a price. Consumers want to buy exactly one product, and choose to either trust the recommendation of a deal finder or compare recommendations. Deal finders gather information on products and prices on behalf of consumers and truthfully recommend the best deal they are aware of.

2.1 Deal finders

Between the consumers and the sellers are a number $N \geq 2$ of identical intermediaries, called deal finders. Assume that deal finders generate revenue from a competitive market for advertisers, bidding for the information on buyers gathered by successful deal finders. As all consumers buy exactly one unit, the willingness to pay for this information is not influenced by the search

[11] The representation of ε has a positive mismatch parameter is a slight departure from the specification of Anderson and Renault (1999), who consider a random noise increasing the utility. This modification does not impact my results, but is useful in the context of advice in order to represent graphically the expected distance a consumer gets from her bliss point at equilibrium.

behaviour of deal finders. The willingness of advertisers to pay for the information extracted from consumers is denoted by $\chi > 0$, so that a consumer buying from a deal finder generates a revenue of χ for this deal finder. In the paper, I assume N to be exogenously given. I show in Appendix D that the model can easily be extended to study the entry decision of deal finders.

In the main part of the paper I study the following sequential search. Suppose that any deal finder receiving a query from a consumer of type i can sequentially sample sellers by each time incurring a linear search cost s to discover a price p_j and mismatch value $\varepsilon_{i,j}$. A deal finder that sampled q sellers thus bears a total cost of qs. Following a query, the N deal finders simultaneously search for deals, and when they find a satisfactory deal $p+\varepsilon$ they advertise it to the consumer. This assumption can be taken literally in the case of physical intermediaries exerting an effort to answer a customer's request, and it is also the most tractable one.

In the case of online advice, a perhaps more realistic assumption is to allow deal finders to carry search in a non-sequential way. Assume that any deal finder invests before seeing the search results in order to be able sample a number of deals q, at a cost sq. This can be interpreted as the investment in building the right algorithms and search environment to be able to deal with specific preferences. I show that the results of the sequential model extend to the non-sequential one in Appendix B, and develop the intuition when relevant in the text after the main propositions.

2.2 Consumers

A share σ of "savvy" consumers makes simultaneous queries and compares different advices, while the rest follows the advice of a single deal finder. The savviness of a consumer is unobservable to the deal finder, so that she does not know for whom she is competing at the time of the query. All consumers are ex-ante identical. The cost of becoming savvy is $c > 0$. As for deal finders, I consider two types of consumer information, corresponding to the standard models of Varian (1980) and Burdett and Judd (1983) respectively. I develop the results of the Varian setting in the main text and report the results of the Burdett and Judd one, for which a full development can be found in Appendix C.

In the Varian setting a savvy consumer observes N deals, and chooses to buy from the deal finder offering the best one. A non-savvy consumer makes only one query to a deal finder picked at random, and receives the best quote of this deal finder. In such a framework, the presence of more deal finders therefore has the mechanical effect of increasing the information of savvy consumers, as those observe more recommendations. Alternatively, in a Burdett and Judd (1983) setting, a savvy consumer observes 2 deals, and chooses to buy from the deal finder offering the best one. A non-savvy consumer makes only one query to a deal finder picked at random, and receives the best quote of this deal finder. Here, for a given strategy of deal finders and sellers, the number of deal finders has no impact of the information of consumers. Those two polar cases aim at capturing two different understanding of what consumer information means. The Varian setting describes information as bearing the cost of understanding how the market works, and being able to compare options - "not being naïve." The Burdett and Judd setting represents a - perhaps more mechanical - linear cost of clicking on an additional website and entering the query again.[12]

2.3 Sellers

There is a continuum of horizontally differentiated sellers of mass M. Assuming this form of monopolistic competition simplifies a lot the analysis, as it implies that no two deal finders recommend an identical deal to a given consumer. This is however not the driving force behind the results: What matters is that the demand of a savvy consumer is more elastic to a change in the price of a given seller than the one of a non-savvy consumer.

The expected demand for a given seller i is given by

$$D(p_i, p, N) = (1-\sigma)D^{ns} + \sigma D^s. \tag{2}$$

$D^{ns}(p_i, p, N)$ is the probability of making a sale to a non-savvy consumer, corresponding to the probability of being selected by a deal finder, multiplied by the probability that this deal finder is selected at random by a non-savvy consumer. $D^s(p_i, p, N)$ is the probability of selling to a savvy one, corresponding to the probability of being selected by a deal finder and offering the best deal among all the recommendations received by a savvy consumer. As I assume a continuum of sellers, an individual price deviation $p_i \neq p$ only affects the expected demand of seller i.

2.4 Equilibrium definition

The objective function of a seller is to maximize $p_i D(p_i, p, N)$ given the search behaviour of deal finders and the share of savvy consumers σ. As in Anderson and Renault (1999), I focus on a symmetric solution where each seller offers an (endogenous) identical price p, and study the optimal price p_i chosen by a seller i. A symmetric equilibrium is thus a situation in which, for each seller, the optimal $p_i = p$.

[12]In contrast to the original framework in Burdett and Judd (1983), I show in Appendix C that the restriction to parameter values such that consumers observe either 1 or 2 deal finders is not important and similar results can be obtained with lower search costs.

As I focus on a symmetric price, deals only vary at equilibrium by their mismatch value $\varepsilon_{i,j}$. I also look for a symmetric equilibrium for deal finders. In the sequential search model, this implies that deal finders keep searching for a deal until finding a mismatch value ε below some threshold w. As a tie-breaking rule, I assume that deal finders search when indifferent. In the Varian setting, if all deal finders follow this strategy, the probability that a given deal finder with mismatch value $\varepsilon < w$ provides the best deal to a savvy consumer is $(\frac{F(w)-F(\varepsilon)}{F(w)})^{N-1}$. If its $N-1$ rivals follow the above strategy, if a deal finder has found a product with mismatch value ε, its expected revenue abstracting from search and entry costs is

$$\pi(\varepsilon) = \left(\sigma \left(\frac{F(w)-F(\varepsilon)}{F(w)}\right)^{N-1} + \frac{1-\sigma}{N}\right)\chi. \tag{3}$$

The first part is the demand from savvy consumers multiplied by the probability of offering the best deal among the N queries they made. The second part is the non-savvy consumers who randomly picked the deal finder and made only one query. The expected search cost to be paid by a deal finder in order to find a mismatch value below w is equal to $\frac{s}{F(w)}$ (this is a general property of a geometric distribution).[13]

All players choose the strategy that maximizes their utility given their expectation of other players' strategies. Given the search cutoff of deal finders, symmetric price, and share of savvy consumers σ, each consumer chooses wether or not to be informed at a cost $c > 0$, with objective function to maximize (1). I show in the next section that for every symmetric price and share of savvy consumers, there exists a unique search cutoff. This cutoff is independent of the symmetric price. For every search cutoff and share of savvy consumers a symmetric price equilibrium exists and is unique. Finally, given the equilibrium cutoff as a function of σ and independent of the symmetric price, there exists a unique equilibrium share of savvy consumers σ, so that a symmetric equilibrium always exists, and there is a unique symmetric equilibrium.

To summarize, the timing of the game is as follows:

1. Sellers simultaneously set their price p_j. I focus on a symmetric equilibrium price p.

2. Buyers simultaneously choose whether to send a query to a deal finder chosen at random, or to send a query to all deal finders at cost c (Varian). The equilibrium share of informed

[13] As I assume deal finders search a discrete number of times within a large number of sellers, the deal finders search within independent and identically distributed deals. With a more limited selection of sellers, I would have to consider overlapping suggestions by deal finders to savvy consumers, therefore limiting the incentives to become savvy.

consumers is σ. I also consider a variant in which buyers observe either one or to two deal finders at a cost c (Burdett-Judd).

3. Deal finders sequentially search for each query they received until they find a mismatch value below their optimal cutoff value. I focus on a symmetric equilibrium cutoff "w". Buyers accept the best deal out of all the queries they made. I also consider a non-sequential search setting where deal finders invest ex-ante in an algorithm that allows them to sample q sellers.

3 Equilibrium

I first study the equilibrium search of deal finders, then the price and finally consumer information.

3.1 Search cutoff

Standard search theory indicates that for a symmetric price p the optimal threshold mismatch value w must satisfy

$$s = \int_0^w (\pi(\varepsilon) - \pi(w)) f(\varepsilon) d\varepsilon, \tag{4}$$

so that the following result holds.

Proposition 1 *For a given share of informed consumers, there exists a unique symmetric search cutoff w if the market price p is symmetric, so that deal finders search until they find a mismatch value strictly lower than w. All other things held equal, the search cutoff w is weakly increasing in s and N and weakly decreasing in* σ.

Proof. Rewriting (4) by using (3), it is easy to show that if there is an interior solution w solves

$$s = \chi\sigma \int_0^w \left(\frac{F(w) - F(\varepsilon)}{F(w)} \right)^{N-1} f(\varepsilon) d\varepsilon, \tag{5}$$

by using the fact that $\pi(w) = \chi \frac{1-\sigma}{N}$. It is straightforward that the left-hand side increases with s while the right-hand side increases with σ and w. There exists no corner solution $w = 0$ as $s > 0$, and there exists a corner solution $w = b$ if and only if $s \geq \chi\sigma \int_0^b (1 - F(\varepsilon))^{N-1} f(\varepsilon) d\varepsilon$. ■

As we would expect, the threshold mismatch value increases with the search cost s. The threshold w also necessarily increases with the number of deal finders, so that for a given share of savvy consumers σ a deal finder becomes less choosy when it faces more rivals. This is a direct consequence of the fact that, for a given symmetric search strategy of the competitors, the marginal benefit of an additional search is lower when the number of deal finders is higher. Through the paper, I focus on cases where $w < b$, so that the delegated search problem has an interior solution.

The assumption that non-savvy consumers pick deal finders at random is without loss of generality. Indeed, even if some deal finders are used more often by default, competition remains for savvy consumers only. Hence, equation (5) would remain identical.

To illustrate, if ε is uniformly distributed on $[0,1]$ and $\chi = 1$, then if there is an interior solution,

$$w = \frac{sN}{\sigma}, \tag{6}$$

where N is the number of deal finders on the market, and σ the endogenous share of savvy consumers.

The same reasoning holds for the non-sequential search (see appendix B): a deal finder invests up to the point where the marginal cost of improving the algorithm equates the marginal benefit. Deal finders invest less if the marginal cost of doing so increases. The marginal benefit is decreasing in N as the probability that a marginal investment allows making a sale to a savvy consumer also decreases, so that the direct effect of N is to decrease investment. It also decreases in σ, that represents how important making a sale to a savvy consumer is. The fact that savvy consumers compare either one or all options (the Varian setting) does not affect the results either. In a Burdett and Judd setting, it is possible to show (see Appendix C) that (5) becomes

$$s = \chi \frac{2\sigma}{N} \int_0^w \frac{F(w) - F(\varepsilon)}{F(w)} f(\varepsilon) d\varepsilon \tag{7}$$

also decreasing in N. The reason is that a higher number of deal finders does not increase the number of deal finders one is competing with for a given savvy consumer, but the number of savvy consumers possibly faced by a deal finder.

For later reference, I need to define $\eta_{i,j} = \varepsilon_{i,j} + p - p_j$ for a seller j setting a different price than the equilibrium p. Define $\phi(\eta)$ the density of $\eta_{i,j}$ so that, at the symmetric equilibrium price $\phi(\eta) = f(\varepsilon)$. Off the equilibrium path, (5) rewrites

$$s = \chi\sigma \int_0^w \left(\frac{\Phi(w) - \Phi(\eta)}{\Phi(\eta)} \right)^{N-1} \phi(\eta) d\eta, \tag{8}$$

with $\Phi(x) = \int_0^x \phi(y) dy$.

In order to understand the direct effect of market concentration, it is useful to see how mismatches are affected when the share of savvy consumers is assumed to be exogenous. I first derive a result on the direct effect of market concentration on the utility of non-savvy consumers.

Lemma 1 *For a given price and share of informed consumers, a non-savvy consumer receives an expected utility u^{ns}, that is decreasing in the search cost of deal finders (s), increasing in the share of savvy types (σ) and decreasing in the number of deal finders (N).*

Proof. Using the result from proposition 1, it is easy to show that the expected mismatch of a non-savvy consumer ε_{ns} is equal to the expected value of a random draw over the interval $[0, w]$,

$$\varepsilon^{ns} = \int_0^w \varepsilon \frac{f(\varepsilon)}{F(w)} d\varepsilon, \tag{9}$$

with $\frac{\partial \varepsilon^{ns}}{\partial s} \geq 0$, $\frac{\partial \varepsilon^{ns}}{\partial \sigma} \leq 0$ and $\frac{\partial \varepsilon^{ns}}{\partial N} \geq 0$. Hence, as the expected utility of a non-savvy type is given by $u^{ns} = v - p - \varepsilon^{ns}$, the proposition follows. ■

The fact that the share of savvy types benefits the non-savvy types is the classic search externalities. The intuition is that the higher the share of savvy types, the more the deal finders compete for them (and search), and the individual efforts of deal finders also benefit the non-savvy. This lemma also conveys the direct effect of lower market concentration (higher N) on the deals received by the non-savvy types. Because, when there are more deal finders, each deal finder searches with lower intensity (proposition 1), a consumer buying from a deal finder chosen at random receives deals of lower quality when there is more competition. The qualifier that the share of savvy types is exogenous is however crucial, as by taking σ as given, I only capture the direct effect of market concentration on equilibrium deals.

To illustrate, if ε is uniformly distributed on $[0, 1]$ and $\chi = 1$, (9) yields

$$\varepsilon^{ns} = \frac{sN}{2\sigma}. \tag{10}$$

I can now turn to the expected mismatch for savvy consumers.

Lemma 2 *For a given price and share of informed consumers, a savvy consumer receives an expected utility u^s, that is decreasing in the search cost of deal finders (s) and increasing in the share of savvy types (σ). The impact of the number of deal finders (N) is ambiguous, as the presence of more deal finders increases u^s if and only if*

$$-\frac{\partial \varepsilon^s}{\partial N} \geq \frac{\partial \varepsilon^s}{\partial w} \frac{\partial w}{\partial N}, \tag{11}$$

where ε^s is the expected mismatch value received by a savvy consumer.

The proof is in Appendix. The difference between the savvy and non-savvy consumers comes from the fact that the impact of market concentration N on the expected mismatch value of a savvy type ε^s is ambiguous. The right-hand side of (11) represents the effect that a higher number of deal finders makes each deal finder less choosy (higher w). The left-hand side represents the effect that it also increases the number of options a savvy consumer can choose from.

If ε is uniformly distributed on $[0,1]$ and $\chi = 1$, the overall direct effect is that more deal finders make the savvy type worse off as ε^s simplifies to

$$\varepsilon^s = \frac{sN}{\sigma(1+N)}. \tag{12}$$

In this case, the result from the Varian setting slightly departs from the Burdett and Judd one. If consumers compare either 1 or 2 options, there is no direct benefit from a higher number of deal finders. Hence, (11) is never fulfilled as the left-hand side is equal to zero. Two well-documented consequences follow immediately from the observation of the two mismatch values, both in the Varian and Burdett and Judd setting. First, savvy consumers always have a better deal than non-savvy ones. Second, savvy consumers "protect" the others by decreasing the mismatch value received by all consumers.

3.2 Symmetric price

I now turn to the equilibrium price offered by a continuum of sellers of mass M. I assume a symmetric market price p, and study the optimal price p_i chosen by a seller i. The expected demand a seller i receives from a non-savvy consumer is given by

$$D^{ns}(p_i, p, N) = \frac{F(w+p-p_i)}{MF(w)}. \tag{13}$$

This expression corresponds to the probability of being selected by a deal finder following the search strategy defined in (8), with i being the only seller off the price equilibrium path, so that the density of η for all other sellers is $\phi(\eta) = f(\varepsilon)$. As deal finders are selected at random by non-savvy consumers, this is equivalent to N times the probability of being selected by each of the deal finders, divided by the probability that a non-savvy consumer picks a deal finder N. The numerator $F(w+p-p_i)$ is the probability of offering a mismatch below w for a specific consumer, that can be alleviated by offering a different price than the market. If all deal finders play the same strategy w, all deal finders share the demand from non-savvy consumers equally. As we assume a continuum of sellers of mass M, the probability of being selected by two deal finders is zero.

Denote by $r(j)$ the probability that a random draw between $p_i - p$ and w be lower than j random draws between 0 and w,

$$r(j) = \int_0^{w+p-p_i} \frac{f(\varepsilon)}{F(w+p-p_i)} \left(\frac{F(w) - F(\varepsilon + p - p_i)}{F(w)} \right)^j d\varepsilon, \tag{14}$$

The expected demand a seller i receives from a savvy consumer is given by,

$$D^s(p_i, p, N) = ND^{ns} r(N-1), \tag{15}$$

which corresponds N times the probability of being selected by the a deal finder, multiplied by the probability to offer a better deal than the $N-1$ other selected sellers.

Proposition 2 *For a given search cutoff w and share of informed consumers σ, there exists a unique symmetric price p. All other things held equal, the symmetric price is increasing in w and decreasing in σ and N.*

The formal proof is in Appendix. The difference with standard models of monopolistic competition is that there are two parts in the demand, one being more elastic than the other. Indeed, while being recommended by a deal finder is enough to sell to non-savvy consumers, a seller needs to offer the best deal among all the ones selected by intermediaries to attract the savvy ones. When deal finders become more selective (lower cutoff w), both parts of the demand become more elastic, so that prices decrease. When the share of savvy consumers σ increase, the most elastic part of the demand becomes more important, so that prices decrease. Finally, all other things held equal, increasing the number of deal finders make the most elastic part of the demand even more elastic, hence decreasing the prices. One can observe that if the number of deal finders is very high - but of a magnitude smaller than the number of sellers for our assumption on each seller being selected at most by one seller to hold - and if all consumers are savvy ($\sigma \to 1$), how choosy the deal finders are has no impact on the equilibrium price, and this price converges to the Perloff and Salop (1985) model (see Proposition 1 in Anderson and Renault, 1999).

To illustrate, if ε is uniformly distributed on $[0,1]$ and $\chi = 1$, then the symmetric price solves

$$p = \frac{w}{1 + (N-1)\sigma}. \tag{16}$$

The above results also hold in the non-sequential search setting for deal finders. An intermediary that observes more sellers increases the elasticity of demand for the same reason as one that lowers its search cutoff. The Burdett and Judd framework on the consumer side yields a small difference however. If consumers observe either 1 or 2 intermediaries, the number of deal finders has no direct effect on the elasticity of demand for a given seller.

3.3 Equilibrium share of savvy consumers

I can now solve for the equilibrium share of savvy consumers σ. Given their expectation on a symmetric cutoff w and price p consumers simultaneously choose whether or not to become "savvy", at a constant cost c. I show in Section 5 that this assumption is not innocuous, as allowing for different consumers to have different costs of becoming savvy may revert the results. Define

$$\Delta = \varepsilon^{ns} - \varepsilon^{s}, \tag{17}$$

as the expected premium (in terms of expected mismatch) paid by uninformed consumers. If there exists an interior solution, the equilibrium share σ is found by solving

$$\Delta(\sigma) = c, \tag{18}$$

else $\sigma = 0$ if $\Delta(0) \leq c$ and $\sigma = 1$ if $\Delta(1) \geq c$.

It is therefore possible to identify the impact of market concentration on the share of savvy consumers.

Proposition 3 *If the market price p is symmetric, there exists a unique equilibrium share of savvy consumers. This share is increasing in the number of deal finders (N).*

The proof is in Appendix. This Proposition describes the indirect effect of market concentration. It follows from lemmas 1 and 2 that the presence of more deal finders increases the difference between the best deal a savvy and a non-savvy consumer observe. It thus becomes more interesting for a consumer to invest in being informed. It also holds that $\frac{\partial(\varepsilon^{ns}-\varepsilon^{s})}{\partial\sigma} < 0$, so that the incentives to become informed decrease when the number of informed consumers increases. This is a pretty standard intuition, as the protection of non-savvy consumers increases with the number of savvy consumers.

From propositions 1, 2 and 3 it follows that there exists a unique equilibrium with symmetric prices. Both the cutoff w in proposition 1 and the share of savvy consumers in proposition 3 are independent of the symmetric price. As σ increases with w and w decreases with σ, the equilibrium pair $\{w,\sigma\}$ is unique if the price is symmetric. By proposition 2, a symmetric equilibrium price exists and is unique for a given pair $\{w,\sigma\}$.

If ε is uniformly distributed between $[0,1]$ and $\chi = 1$, the premium for being informed rewrites $\Delta = \frac{sN(N-1)}{2\sigma(N+1)}$ and (18) solves

$$\sigma = \frac{sN(N-1)}{2c(N+1)}, \tag{19}$$

if σ has an interior solution (if $c > \frac{sN(N-1)}{2(N+1)}$) and $\sigma = 1$ else.

Lemmas 1 and 2 do not depend on whether the search of deal finders is sequential or not. Hence, proposition 3 also holds for both specifications. The same is true when considering a Burdett and Judd framework where consumers sample either or or two deal finders. As a higher number of deal finders N decreases the search effort of deal finders (proposition 1), the marginal benefit from comparing two options instead of one also increases with N. Indeed, the expected gain from comparing two "bad" advices is always higher than the gain from comparing two "good" ones. It follows that the share of consumers comparing two options increases with N, up to the level where the expected gain from comparing options equals c. This implies that in a Burdett and Judd framework, changes in σ always exactly compensate the impact of an increase in N on the search behaviour. Indeed, for a given distribution of ε, there is a unique w such that the expected difference between one and two draws over the interval $[0,w]$ is equal to exactly c.

4 Consumer welfare and market concentration

Using the results in propositions 1, 2 and 3, it is possible to characterize the impact of the exogenous parameter N on the equilibrium welfare of consumers. As all consumers are ex-ante identical, and as consumers choose to become savvy up to the point where

$$u^s - c = u^{ns}, \tag{20}$$

all consumers have an identical expected surplus. As I have assumed the utility to be quasi-linear so that all payments are directly subtracted from the utility, the expected surplus of each consumer is equal to

$$u = u^s - c = u^{ns} = v - p - \varepsilon^{ns}. \tag{21}$$

Thus, it is enough to characterize the expected surplus received by a non-savvy consumer in equilibrium in order to understand the welfare effect of market concentration. To do so, I study separately the effect on equilibrium price and equilibrium expected mismatch.

Proposition 4 *The expected equilibrium mismatch of a non-savvy consumer is decreasing in the number of intermediaries.*

The formal proof is in appendix. To see this, one has to put together the effects documented in propositions 1 and 3. By proposition 1, we know that the direct effect of a higher number of intermediaries is to decrease the intensity of search of intermediaries, hence making a non-savvy buyer worse off. By proposition 3, we however know that this has the indirect effect of increasing the share of savvy consumers, thereby protecting the non-savvy ones by making deal finders more selective (see again proposition 1). It is possible to show that the latter effect dominates. In a Varian setting where consumers compare either one or all deal finders, we know that the share of savvy consumers increases if the number of deal finders goes from N to $N' > N$ up to the point where the cutoff w' is such that comparing N' options has the same marginal benefit (c) as comparing N options with cutoff w. It is straightforward that $w = w'$ cannot hold, as comparing N' searches of a similar quality is always better than comparing N. Hence, the only solution is that $w' < w$, the higher number of deal finders actually make all deal finders more selective through the indirect effect.

As discussed in the previous section, the statement is weakly true in a Burdett and Judd setting, as an increase in the number of deal finders from N to N' is then exactly compensated by an increase in σ to yield $w' = w$. The sequentiality of the search of deal finders does not affect the result.

In the case where ε is uniformly distributed on $[0,1]$ and $\chi = 1$, plugging (19) into (6) yields

$$w = \frac{2c(N+1)}{N-1}, \tag{22}$$

if σ has an interior solution and $w = sN$ if $\sigma = 1$. We thus observe that the entry of new deal finders decreases the equilibrium mismatch received by all consumers once the information decision of consumers is made endogenous. The direct effect documented in (5) is that each deal finder becomes less choosy when an additional competitor enters the market. However, the indirect effect is that the entry of new deal finders increases the incentives to become informed, because the gap between the best deal obtained by savvy and non-savvy consumers increases. Hence, the increase

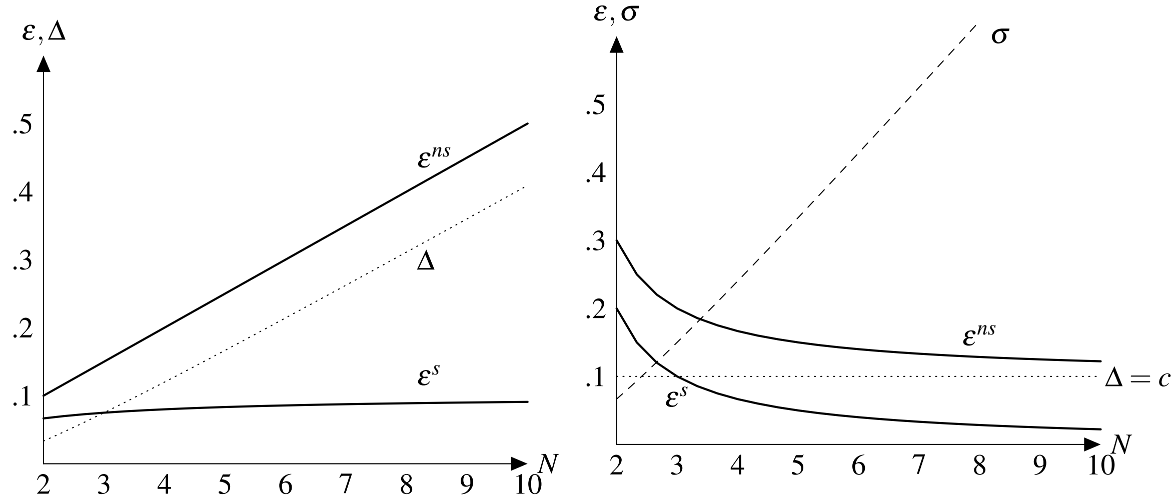

Exogenous share of savvy types $\sigma = 0.2$

Endogenous share of savvy types, with $c = 0.1$

Figure 1: The impact of market concentration on expected mismatch, with $s = 0,02$, $F(\varepsilon) = \varepsilon$.

in the share of savvy consumers more than compensates the direct effect, so that entry actually makes deal finders more choosy.

I illustrate the different effects on Figure 1. On the left panel, I take the share of savvy types σ as given. We see that for a given σ, a higher number of deal finders increases the mismatch value received by non-savvy consumers and (slightly) increases the mismatch value received by the savvy types. This is the direct effect of deal finders becoming less choosy. When the number of deal finders is equal to $N = 10$, these just pick one seller at random, so that a non-savvy customer receives an expected mismatch value of 0.5. For more deal finders, there is no interior solution for w. I do not study this case as it would imply making further assumptions about deal finder and non-savvy consumers.[14] We also observe that the difference between the expected mismatches of the two types of consumers Δ increases with the number of deal finders N. On the right panel, I allow for the share of savvy types σ to be endogenous. Because the difference between the expected mismatch value received by a savvy and a non-savvy type increases with N, more and more consumers choose to become savvy. The effect of higher rates of savviness is to decrease the expected mismatch value received by both types of consumers in equilibrium. Hence, for a given price, market concentration is bad for consumers as even if a smaller number of deal finders has the direct positive effect of making each of them more choosy, it also has the indirect effect of making consumers less informed, thus allowing deal finders to become less choosy.

[14] In particular, if deal finders need at least one price quote in order to attract the non-savvy types, they may still benefit from searching once up to a certain point.

Proposition 5 *The symmetric equilibrium price p is decreasing in the number of intermediaries.*

Proof. The impact of N can be decomposed into several effects,

$$D^s_{p_i}(p,p,N)\frac{d\sigma}{dN}+D^{ns}_{p_i}(p,p,N)\frac{d(1-\sigma)}{dN}+(1-\sigma)\frac{dD^{ns}_{p_i}(p,p,N)}{dN}+\sigma\frac{dD^s_{p_i}(p,p,N)}{dN}\leq 0. \quad (23)$$

The sum of the first two terms in (23) is always negative, as $\frac{d\sigma}{dN}\geq 0$ (Proposition 3) and $D^s_{p_i}(p,p,N)<D^{ns}_{p_i}(p,p,N)$ (the "savvy" segment of the market is more elastic than the non-savvy one). The third term and fourth terms are negative as $\frac{dw}{dN}\leq 0$ (proposition 4), as the elasticity of demand is higher if the deal finders are more selective (lower w). ■

Sellers have an incentive to offer lower prices if deal finders become more choosy. This is the case when the number of deal finders is higher, by proposition 4, albeit weakly so in a Burdett and Judd setting. They also want to lower the price if savvy consumers observe more options. This is mechanically the case when the number of deal finders increases in a Varian setting, and again weakly so in a Burdett and Judd setting. Finally, sellers decrease the prices if the share of savvy consumers increases, which strictly holds when the number of deal finders increase in both settings, as from proposition 3.

The result derives from the fact that, if deal finders become more choosy, sellers have to offer better prices on all segments of the market, in order to have a chance of being selected. The consequence is one of a virtuous (vicious) circle: the choosier the deal finders are, the more a seller wants to provide a low price. However, there are other reasons that lead to lower prices when concentration on the market for deal finders decreases. The first one is that fewer consumers pick a deal finder at random, so that the competitive segment of the market matters more to sellers. The second is that the competitive segment becomes even more competitive as savvy consumers have more options to pick from.

If ε is uniformly distributed between $[0,1]$ and $\chi=1$, we have shown in (16) that the equilibrium price solves $p=\frac{w}{1+(N-1)\sigma}$, with w and σ defined in (19) and (22). We immediately see that the equilibrium price always decreases with N, as w decreases with N and σ increases with N. Using the equilibrium values for these parameters, if σ has an interior solution, (16) simplifies to

$$p=\frac{4c^2(1+N)^2}{2c(N^2-1)+(N-1)^3Ns}. \quad (24)$$

The equilibrium price decreases with the number of deal finders, for two reasons. First, more deal finders increase the share of savvy consumers, making each deal finder more choosy. Second, a higher share of savvy consumers increases the price elasticity of demand for a given seller. Note that equation (24) only represents the case where σ has an interior solution, for $c > \frac{sN(N-1)}{2(N+1)}$. When c becomes smaller, the price does not converge to zero as suggested by (24), but to $p = s$.

As for all the results above, it is easy to show that the sequentiality of the deal finder search does not affect the result. Proposition 5 also holds in the Burdett and Judd framework where buyers observe either one or two deal finders. As in this case $\frac{dw}{dN} = 0$ and $\frac{d\sigma}{dN} > 0$, a higher number of deal finder does not affect the elasticity of the demand from a given buyer, but increases the share of buyers with a higher elasticity of demand, hence leading to lower prices.

The answer to to the initial research question - whether market concentration benefits consumers' welfare - directly follows from the above results.

Corollary 1 *Lower market concentration increases the expected utility of all consumers.*

The Proof is straightforward from propositions 4 and 5. As the expected mismatch received by non-savvy consumers is decreasing in N, and as prices decrease with N the total welfare effect of lower market concentration is always positive. The negative effect of more competition is a higher mismatch for a given share of savvy consumers. The positive effect is a higher share of savvy consumers, therefore decreasing the mismatch and the price. As the positive effect of competition is more important lower market concentration always benefits consumers.

Using the uniform distribution $F(\varepsilon) = \varepsilon$ and the same parameters as Figure 1, Figure 2 represents the sum of the mismatch and the price effect on the equilibrium welfare. The dashed lines represent the expected mismatch of a non-savvy consumer and the expected price (both decreasing with N). The solid line represents the expected welfare u, equal to $v - p - \varepsilon^{ns}$, and is increasing with N. The signs of the effect of concentration on price and mismatch are identical, but it is striking that, at least for the highest levels of market concentration, the most important impact on consumer welfare is not so much the search effort by deal finders, but the price competition among sellers. What really benefits consumers is the externality generated by more consumers observing more than one product, even more so than the fact that deal finders are more selective.

As from the above results, the general impact of concentration on consumer welfare is not impacted by the choice between sequential or non-sequential search. Corollary 1 also holds in the Burdett and Judd framework where consumers compare either one or two options, but the entire positive effect of lower market concentration comes from the lower prices (see table 1). I illustrate on figure 3 the results of the same exercise as on figure 2 in a Burdett and Judd framework. We see that only the price effect matters, and that this effect is less important than above. This suggests

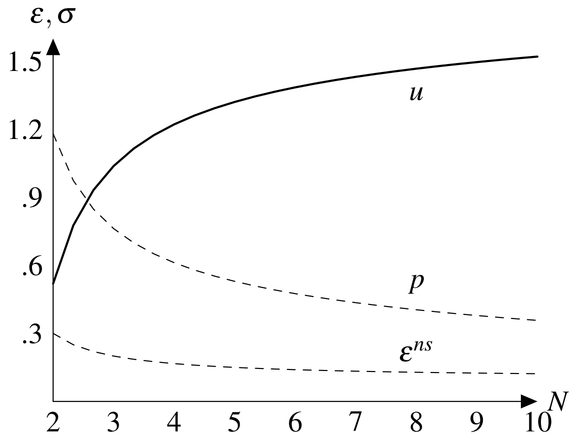

Figure 2: equilibrium price and welfare, with $s = 0,02$, $F(\varepsilon) = \varepsilon$, $c = 0.1$, $v = 2$.

Table 1: Effect of a higher number of deal finders N for different specifications of consumer information

	share of savvy consumers σ	*search cutoff* w	*price* p
Varian	+	-	-
Burdett and Judd	+	=	-

that whether we consider consumer information as some sort of naïvity in the Varian sense or in a more mechanical way as in Burdett and Judd influences to what extent higher market concentration hurts consumer welfare.

5 Heterogenous costs of savviness

The assumption of an homogenous cost of information c is crucial to my results. Assume for instance that the ability for a consumer to become savvy depends on a parameter θ, randomly

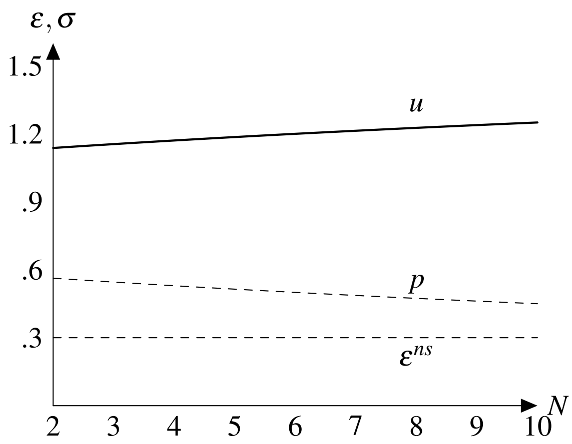

Figure 3: equilibrium price and welfare in a Burdett and Judd framework, with $s = 0,02$, $F(\varepsilon) = \varepsilon$, $c = 0.1$, $v = 2$.

drawn from a uniform distribution over $[0,1]$, so that the cost for a consumer i to become savvy is equal to $c(\theta_i) = \gamma\theta_i$. This means that the most able consumer has no cost of becoming informed, that the least able has a cost c of becoming informed, and that the cost of acquiring information is linear in the ability. For a given expected value of the mismatch differential between informed and uninformed consumers Δ, if a consumer of type j with $\theta_j > \theta_i$ prefers to become savvy, a consumer of type i also prefers to become savvy. Consumers can thus be ranked by their ability to acquire information, so that the cost for the σth consumer to become savvy is $c(\sigma) = \gamma\sigma$. This also implies that the more consumers become savvy, the higher the expected difference between the mismatch received by a savvy and a non-savvy consumer in equilibrium.

It is possible that when σ increases savvy consumers are made better off, but those who cannot afford becoming savvy are worse off. If even the non-savvy consumers are made better off however, it means that more competition on the market for deal finders is Pareto improving. The intuition is relatively straightforward, as allowing for heterogeneous costs of savviness is an intermediary case between assuming an exogenous share of σ (proposition 1) and ex-ante identical consumers (proposition 4).

In the uniform case with $\chi = 1$, I can rewrite (18) as

$$\sigma = \sqrt{\frac{sN(N-1)}{2\gamma(N+1)}}, \tag{25}$$

where the share of savvy consumers still increases with N, but the increase is slower due to the marginally increasing cost of becoming informed. Plugging (25) into (6) yields

$$w = \frac{sN\sqrt{2\gamma(N+1)}}{\sqrt{sN(N-1)}}, \tag{26}$$

which can be shown to be increasing in N. This means that when market concentration decreases, (i) the share of consumers choosing to be informed increases, (ii) the expected mismatch received by informed consumers decreases, but (iii) the expected mismatch received by the remaining uninformed consumers increases. The difference between (22) and (26) is the existence of a distributional impact of the level of concentration on the market for deal finders. A key assumption for more competition on the market for deal finders to be Pareto improving for consumers is that the cost of becoming informed does not vary too much among consumers.

I illustrate this idea on Figure 4, by comparing the case studied in the Figure 1 with $c = 0.1$,

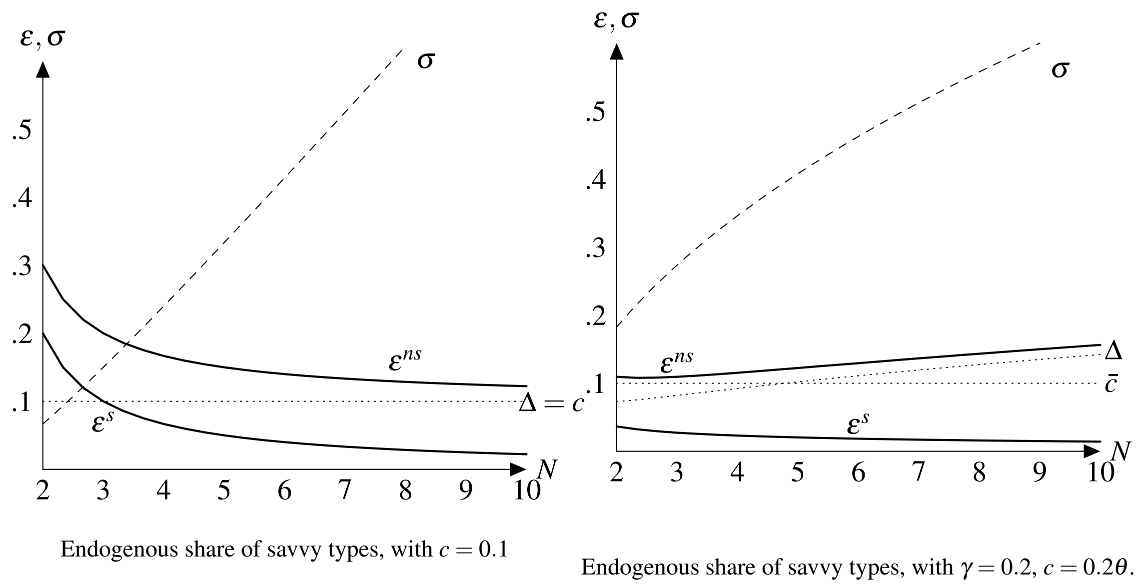

Figure 4: The importance of an homogenous cost of savviness.

to a cost function $c(\theta) = 0.2\theta$, so that as θ is drawn from a uniform distribution on $[0,1]$, the average cost of becoming savvy is identical in both examples, $\bar{c} = 0.1$. The left panel is just the right panel of Figure 1. On the right panel, we see the impact of heterogeneous costs of savviness. When a small number of deal finders are active on the market, the expected mismatch received by both types of consumers is pretty close, and is lower than on the left panel for both types as some consumers have almost no cost of being savvy. When the number of deal finders increases, the share of savvy consumers increases. The impact of N on σ however quickly becomes insufficient to make non-savvy consumers better off. Hence, in this case market concentration does not have a uniform impact on all consumers. The less able consumers, with the highest cost of becoming savvy γ, benefit from mergers until the number of deal finders is equal to $N = 3$, while the most able consumers always prefer a higher number of deal finders.

6 Conclusions

The present paper puts together the incentives deal finders have to invest in search with the incentives consumers have to become informed and the incentives for sellers to offer low prices. This conjunction leads to two opposite effects of the impact of market concentration on the search behaviour of deal finders. The first effect is that more competition decreases the incentives for deal finders to invest in search. The second effect is that more competition increases the share of consumers choosing to become informed. These two effects alone do not suffice to characterize the welfare impact of competition, as they also influence the price offered by sellers. I show that more competition on the market for deal finders makes these intermediaries more choosy and increases price competition among sellers.

More generally, this paper aims at contributing to the general debate about the impact of the multiplication of sources of information available on the Internet. The main message from this study of deal finders is that by ignoring the indirect effect of market concentration on consumer education one might draw incorrect conclusions overestimating the benefits from an economy with a limited number of (presumably) high quality source.

I also identify two important limitations to my results. First, the idea of “search externalities” relies on deal finders being unable to tell who is savvy and who is not. Without this assumption, if consumers are ex-ante identical, the only solution is that all choose to become informed, all choose to become uninformed, or all are indifferent. Second, as shown in section 5, the result of a Pareto improvement identified in corollary 1 relies on all consumers being ex-ante identical. If naïvety is a behavioural characteristic and not the result of consumer choice, one needs to consider distributional effects.

There are several ways in which this model could be extended. A first one would be to develop further where the revenue of deal finders come from, by modelling an explicit price relationship between these platforms, buyers, and sellers, as a two-sided market. A second one would be to explicitly model the choice consumers make of whether to use deal finders or to directly buy from the sellers. This would imply for consumers to balance the cost of observing the result of deal finders with the cost of themselves linearly search for deals. A third one would be to consider ordered search among deal finders, and to allow those to compete (for instance by advertising) for prominence.

Bibliography

ANDERSON, S. P. and RENAULT, R. (1999). Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics*, pp. 719–735.

ARMSTRONG, M. (2015). Search and ripoff externalities. *Review of Industrial Organization*, **47** (3), 273–302.

— and ZHOU, J. (2011). Paying for prominence. *The Economic Journal*, **121** (556), F368–F395.

BAYE, M. R., KOVENOCK, D. and DE VRIES, C. G. (1996). The all-pay auction with complete information. *Economic Theory*, **8** (2), 291–305.

BURDETT, K. and JUDD, K. L. (1983). Equilibrium price dispersion. *Econometrica: Journal of the Econometric Society*, pp. 955–969.

CAPLIN, A. and NALEBUFF, B. (1991). Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica: Journal of the Econometric Society*, pp. 25–59.

CHEN, H., XIE, H. and HU, T. (2009). Log-concavity of generalized order statistics. *Statistics & Probability Letters*, **79** (3), 396–399.

CHEN, Y. and ZHANG, T. (2017). Intermediaries and consumer search. *International Journal of Industrial Organization*.

DE CORNIÈRE, A. (2016). Search advertising. *American Economic Journal: Microeconomics*, **8** (3), 156–188.

EDELMAN, B. G. and WRIGHT, J. (2015). Price coherence and excessive intermediation. *Quarterly Journal of Economics*, **130**, 1283–1328.

ELIAZ, K. and SPIEGLER, R. (2011a). Consideration sets and competitive marketing. *The Review of Economic Studies*, **78** (1), 235–262.

— and — (2011b). A simple model of search engine pricing. *The Economic Journal*, **121** (556), F329–F339.

INDERST, R. and OTTAVIANI, M. (2012). How (not) to pay for advice: A framework for consumer financial protection. *Journal of Financial Economics*, **105** (2), 393–411.

LEWIS, T. R. (2012). A theory of delegated search for the best alternative. *The RAND Journal of Economics*, **43** (3), 391–416.

MALIK, O. (2015). In silicon valley now, it's almost always winner takes all. The New Yorker, December 30, 2015.

PERLOFF, J. M. and SALOP, S. C. (1985). Equilibrium with product differentiation. *The Review of Economic Studies*, **52** (1), 107–120.

ULBRICHT, R. (2016). Optimal delegated search with adverse selection and moral hazard. *Theoretical Economics*, **11**.

VARIAN, H. R. (1980). A model of sales. *The American Economic Review*, **70** (4), 651–659.

WOLINSKY, A. (1986). True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics*, pp. 493–511.

Appendix A: Proofs

Proof of Lemma 2

Proof. Savvy consumers obtain an expected mismatch equal to the expected value of the minimum of N random draws between 0 and w. The probability density of a random draw over the interval $[0,w]$ is $g(\varepsilon) = \frac{f(\varepsilon)}{F(w)}$, with cumulative density $G(\varepsilon) = \frac{F(\varepsilon)}{F(w)}$ (the probability density function is therefore specific to a given value of w). The expected value of the first order statistic of N independent draws of $g(\varepsilon)$ is given by the standard formula

$$\varepsilon^s = \int_0^w N(1 - G(\varepsilon))^{N-1} \varepsilon g(\varepsilon) d\varepsilon. \tag{27}$$

Using the definition of $g(\varepsilon)$, this rewrites

$$\varepsilon^s = \int_0^w N(1 - \frac{F(\varepsilon)}{F(w)})^{N-1} \varepsilon \frac{f(\varepsilon)}{F(w)} d\varepsilon. \tag{28}$$

The impact of s and σ are unambiguous, as these variables only affect w. For a given value of w, the marginal impact of N is to decrease ε^s. However, as from Proposition 1, $\frac{\partial w}{\partial N} > 0$, the overall effect is ambiguous. By totally differentiating $\varepsilon^s(N, P(N))$, the expression in the Proposition follows. ■

Proof of Proposition 2

Proof. The expected demand for a given seller i is

$$D(p_i, p, N) = (1-\sigma)D^{ns} + \sigma D^s. \tag{29}$$

As the expected profit of a seller is equal to $p_i D(p_i, p, n)$, it follows that the equilibrium symmetric price p posted by sellers solves

$$p = \frac{-D(p,p,N)}{D_{p_i}(p_i,p,N)} = \frac{-D(p,p,N)}{\sigma D^s_{p_i}(p_i,p,N) + (1-\sigma)D^{ns}_{p_i}(p_i,p,N)}, \tag{30}$$

where $D(p,p,n) = \frac{1}{M}$ as all sellers post an identical price in equilibrium.

It is easy to show that, with a continuum of sellers of mass M,

$$D^{ns}_{p_i} = \frac{-f(w)}{MF(w)}, \tag{31}$$

so that with $\sigma = 0$, as f is log-concave, the unique symmetric equilibrium price would be $p = \frac{F(w)}{f(w)}$. For $D^s_{p_i}$ the expression is less straightforward, as it takes the derivative of D^{ns} multiplied by the probability of being smaller than the first order statistic (also log-concave, as from Chen *et al.*, 2009) of $N-1$ draws over $[0, w]$, $r(N-1)$. The demand D^s is presented in (14) and (15):

$$D^s = ND^{ns}r(N-1). \tag{32}$$

Define $\beta = Nr(N-1)$, so that $D^s = D^{ns}\beta$. At equilibrium $p_i = p$, for $D^s = D^{ns}$ to hold, it must be that $\beta = 1$. We can thus differentiate D_s as,

$$D^s_{p_i} = D^{ns}_{p_i} + D^{ns}\frac{d\beta}{dp_i}. \tag{33}$$

It follows that the direct effect of a higher w is to increase prices, as it decreases both $D^s_{p_i}$ and $D^{ns}_{p_i}$. The direct effect of a higher N is to decrease prices, as it does not affect $D^{ns}_{p_i}$ and increases $D^s_{p_i}$.

The direct impact of σ is to increase the weight put on $D^s_{p_i}$. Hence, to see that higher σ decreases prices it is enough to show that the savvy part of the demand is more elastic. At equilibrium, $D^s = D^{ns} = \frac{1}{M}$. As $D^{ns} = \frac{1}{M}$, it is enough to show that $\frac{d\beta}{dp_i} < 0$ in order to show that $D^s_{p_i} < D^{ns}_{p_i}$, and thus the savvy part of the demand is more elastic. This result is straightforward as $\frac{dr(j)}{dp_i} < 0$, $\forall j > 0$. ■

In the special case of the uniform distribution, with a continuum of sellers of mass M, the demand can be rewritten as

$$D = D^{ns}((1-\sigma) + \sigma N r(N-1)). \tag{34}$$

Denoting $\beta = (1-\sigma) + \sigma N r(N-1)$, I find

$$\begin{aligned} D_{p_i} &= D^{ns}_{p_i} + \frac{1}{M}\frac{d\beta}{dp_i} \\ &= -\frac{1}{Mw} - \frac{1}{Mw}(N-1)\sigma, \end{aligned} \tag{35}$$

so that

$$\begin{aligned} p &= -\frac{\frac{1}{M}}{-D_{p_i}} \\ &= \frac{w}{1+(N-1)\sigma}. \end{aligned} \tag{36}$$

Proof of Proposition 3

Proof. First, I show that there exists a unique equilibrium σ. As by lemmas 1 and 2, $\Delta(\sigma)$ is continuous and decreasing. As c is a constant, it is either a dominant strategy to be uninformed $\sigma = 0$ if $\Delta(1) < \Delta(0) < c$, a dominant strategy to be informed $\sigma = 1$ if $c < \Delta(1) < \Delta(0)$, or there exists a unique intersection $\Delta(\sigma) = c$ if $\Delta(1) < c < \Delta(0)$.

Second, I show that σ increases with N. From proposition 1, we know that w increases with N for a given σ. For a given number of deal finders, it is trivial that the larger the interval of the draws $[0,w]$, the higher the expected absolute gain from observing more draws. From the properties of the first order statistic, for any expected value of the first order statistic of k random draws, $X(1,k)$, over an interval it is always true that $X(1,k) < X(1,k')$ if and only if $k' < k$. Hence, the two effects (higher w and higher N) go in the same direction, to increase Δ for a given σ. It follows from $\frac{\partial\Delta}{\partial N} > 0$ that $\frac{\partial\sigma}{\partial N} \geq 0$. ■

Proof of Proposition 4

Proof. I want to assess the impact of N on u^{ns} for a given value of p. Totally differentiating ε^{ns} with respect to N yields

$$\frac{d\varepsilon^{ns}}{dN} = \frac{\partial \varepsilon^{ns}}{\partial w}\left(\frac{\partial w}{\partial N} + \frac{\partial w}{\partial \sigma}\frac{\partial \sigma}{\partial N}\right), \tag{37}$$

with $\frac{\partial \varepsilon^{ns}}{\partial w} > 0$, $\frac{\partial w}{\partial N} > 0$, $\frac{\partial w}{\partial \sigma} < 0$ and $\frac{\partial \sigma}{\partial N} > 0$. For a given w, $\frac{\partial(\varepsilon^{ns}(w) - \varepsilon^{s}(w,N))}{\partial N} > 0$. At equilibrium, it must hold that $\varepsilon^{ns}(w) - \varepsilon^{s}(w,N) = c$, so that it must hold that $\frac{dw}{dN} < 0$, and $\frac{\partial w}{\partial N} < -\frac{\partial w}{\partial \sigma}\frac{\partial \sigma}{\partial N}$. ■

Appendix B: Non-sequential search

In this Appendix, I consider a non-sequential variant of the model. Instead of linearly searching until they find a deal below a threshold value w, I assume deal finders simultaneously choose the number of sellers they sample before observing the results, in the tradition of Burdett and Judd (1983). Denote by q the symmetric equilibrium number of sellers sampled by a deal finder, and assume a symmetric equilibrium price p, the expected profit of deal finder i sampling q_i prices is

$$\pi(q_i,q) = \chi\left(\sigma\left(\int_0^b f_{(1),q_i}(\varepsilon)(1 - F_{(1),q}(\varepsilon))^{N-1} d\varepsilon\right) + \frac{1-\sigma}{N}\right), \tag{38}$$

where $f_{(1),x}(\varepsilon)$ is the density of the first order statistic of x independent draws with individual density $f(\varepsilon)$, and similarly $F_{(1),x}$ is the cumulative density. The first part of the profit is thus the probability that the smallest of q_i random draws be lower than the smallest of q random draws, times the $N-1$ other deal finders, multiplied by the share of savvy consumers σ. The second part is identical to the sequential model, and represents the fact that a share $1-\sigma$ of non-savvy consumers choose the best of the q deals offered by a deal finder chosen at random. Using the properties of the order statistics, this expression rewrites:

$$\pi(q_i,q) = \chi\left(\sigma\left(\int_0^b q_i f(\varepsilon)(1 - F(\varepsilon))^{q_i - 1 + q(N-1)} d\varepsilon\right) + \frac{1-\sigma}{N}\right). \tag{39}$$

The symmetric equilibrium q is such that

$$\frac{d\pi(q_i,q)}{q_i} = s, \tag{40}$$

for all deal finders i. As q is an integer, a continuous value of q has to be interpreted as a mixed strategy. It is straightforward that for a given q the marginal benefit of an additional search is decreasing in q_i. As in Lemma 1, a simple inspection of (39) and (40) shows that $\frac{\partial q}{\partial s} < 0$ (if the marginal cost increases, the marginal benefit must also increase). It is also clear that $\frac{\partial q}{\partial \sigma} > 0$, as σ directly multiplies the marginal benefit of an additional search, and $\frac{\partial q}{\partial N} < 0$, as N only enters the expression $(1-F(\varepsilon))^{q_i-1+q(N-1)}$.

The expected mismatch obtained by a non-savvy consumer is the minimum of q random draws by one deal finder,

$$\varepsilon^{ns} = \int_0^b q(1-F(\varepsilon))^{q-1}\varepsilon f(\varepsilon)d\varepsilon, \tag{41}$$

so that $\frac{\partial \varepsilon^{ns}}{\partial q} < 0$. The expected mismatch obtained by a savvy consumer is the minimum of q random draws by N deal finders,

$$\varepsilon^{s} = \int_0^b Nq(1-F(\varepsilon))^{Nq-1}\varepsilon f(\varepsilon)d\varepsilon, \tag{42}$$

with$\frac{\partial \varepsilon^s}{\partial N} < 0$, and $\frac{\partial \varepsilon^s}{\partial q} < 0$. Hence, as $\frac{\partial q}{\partial N} < 0$, the sign of $\frac{d\varepsilon^s}{dN}$ is ambiguous. The presence of an additional deal finder increases u^s if and only if

$$-\frac{\partial \varepsilon^s}{\partial N} \geq \frac{\partial \varepsilon^s}{\partial q}\frac{\partial q}{\partial N}. \tag{43}$$

From (41) and (42) it follows directly that $\frac{\partial \Delta}{\partial N} > 0$, so that $\frac{d\sigma}{dN} > 0$. Proposition 4 thus holds: $\frac{dq}{dN} \geq 0$. Indeed, with $N' > N$, the gain from comparing Nq and q and $N'q'$ and q' must be equal, which is only possible with $q' > q$.

Switching to the sellers' side, the demand from non-savvy consumers is

$$D^{ns}(p_i, p, N) = \frac{q}{M}\int_0^b f(\varepsilon)(1-F(\varepsilon - p + p_i))^{q-1}d\varepsilon, \tag{44}$$

the probability of being selected by each deal finder, of offering the best deal among the q random draws of this deal finder, and the probability that each deal finder is chosen at random by a consumer. The demand from savvy consumers is

$$D^{s}(p_i, p, N) = \frac{Nq}{M}\int_0^b f(\varepsilon)(1-F(\varepsilon - p + p_i))^{Nq-1}d\varepsilon, \tag{45}$$

the probability of offering the best deal among Nq random independent draws. The demand for a given seller is $D = \sigma D^s + (1-\sigma)D^{ns}$, and the profit is $p_i D(p, p_i, N)$. As, at a symmetric price equilibrium $D^s = D^{ns} = \frac{1}{M}$, D^s is more elastic. Thus, following a similar reasoning as for the sequential search, p decreases with N if

$$D^s_{p_i}(p,p,N)\frac{d\sigma}{dN} + D^{ns}_{p_i}(p,p,N)\frac{d(1-\sigma)}{dN} + (1-\sigma)\frac{dD^{ns}_{p_i}(p,p,N)}{dN} + \sigma\frac{dD^s_{p_i}(p,p,N)}{dN} \leq 0, \quad (46)$$

which is always true as $\frac{d\sigma}{dN} > 0$ and $\frac{dq}{dN} \geq 0$.

Appendix C: linear information cost

Consider a variant of the model where instead of observing either all or one deal finder, consumers bear a linear search cost c to (non sequentially) observe an additional finder. In line with Burdett and Judd (1983), for an equilibrium where some - but not all - consumers choose to observe only one deal finder to exist, I can focus on equilibria where consumers mix between observing 1 or 2 deal finders (because the marginal benefit of an additional observation is decreasing in the number of observations). The same exercise can easily be extended to any indifference between n and $n+1$ deal finders as, unlike in Burdett and Judd, there is no need for a share of consumers observing only one result.

If there is a share σ of consumers observing 2 deal finders, $\pi(\varepsilon)$ becomes

$$\pi(\varepsilon) = \chi\left(\frac{1-\sigma}{N} + \frac{2\sigma}{N}\frac{F(w)-F(\varepsilon)}{F(w)}\right), \quad (47)$$

so that in the second stage w solves

$$s = \chi\left(\frac{2\sigma}{N}\int_0^w \frac{F(w)-F(\varepsilon)}{F(w)} f(\varepsilon)d\varepsilon\right), \quad (48)$$

with identical properties as in the "Varian" setting. In the first stage, σ solves

$$c = \int_0^w \varepsilon g(\varepsilon)d\varepsilon - \int_0^w 2(1-G(\varepsilon))\varepsilon g(\varepsilon)d\varepsilon, \quad (49)$$

with $g(\varepsilon) = \frac{f(\varepsilon)}{F(w)}$ and w from (48). As in the Varian setting, the difference increases with w, so that the indirect effect of higher N is to increase σ. For such a mixed strategy to be an equilibrium, c must not be too low, as consumers must strictly prefer to observe 2 deal finders over 3,

$$c > \int_0^w 2(1-G(\varepsilon))\varepsilon g(\varepsilon)d\varepsilon - \int_0^w 3(1-G(\varepsilon))^2\varepsilon g(\varepsilon)d\varepsilon. \quad (50)$$

As c is constant and nothing in (49) depends on N, it must hold that $\frac{dw}{dN} = 0$. When more deal finders enter, σ increases so that the value of w that solves (48) remains constant.

Finally, the price solves a similar problem as in the Varian case, with as only difference

$$D^s = 2D^{ns} r(2). \tag{51}$$

Hence, as in the Varian case, higher σ increases the demand elasticity even for a given w. In the special case of uniformally distributed ε over $[0,1]$ and $\chi = 1$,

$$w(\sigma) = \frac{sN}{\sigma} \tag{52}$$

$$\sigma = \frac{Ns}{6c} \tag{53}$$

$$w^* = 6c \tag{54}$$

$$p = \frac{w}{1+\sigma} = \frac{(6c)^2}{6c + Ns}. \tag{55}$$

Appendix D: Endogenous entry

Until now, I have taken as exogenous the number of deal finders on the market. It is however possible to solve the model by making entry endogenous. Assume now that a deal finder enters the market at a fixed cost α until expected profit equals zero. It is easy to show that N is fully determined by c, α and s. In equilibrium, all symmetric deal finders get a share $1/N$ of the customers. Hence, the expected profit of a deal finder including search and entry costs is

$$\pi = \frac{\chi}{N} - \frac{s}{F(w)} - \alpha, \tag{56}$$

and the equilibrium number of deal finders is

$$N = \lfloor \chi \frac{F(w)}{s - \alpha F(w)} \rfloor, \tag{57}$$

where $\lfloor x \rfloor$ is the highest integer smaller than x. It follows that N is - unsurprisingly - decreasing in α, allowing us to fully characterize the equilibrium. Assuming identical cost for consumers to be informed c, a lower entry cost for deal finders increases entry, but not in a linear way. Indeed, as shown in (22), a consequence of entry in this case is that deal finders invest more in search as they become more choosy. Hence, entry decreases the market share of each deal finder and increases search costs.

If I ε is uniformly distributed on $[0,1]$ and $\chi = 1$, I find by replacing w by its value found in (22), the equilibrium number of deal finders N as

$$N = \lfloor \chi \left(\frac{\sqrt{(-2c\alpha + 2c + s)^2 - 8c(-2c\alpha - s)} - 2c\alpha + 2c + s}{2(2c\alpha + s)} \right) \rfloor. \tag{58}$$

I represent this example on Figure 5, with identical parameter values as in the right panel of Figure 1. For a given level of effort, dividing by 2 the cost of entry would double the number of deal finders. Here, it is not the case as more entry implies higher search costs.

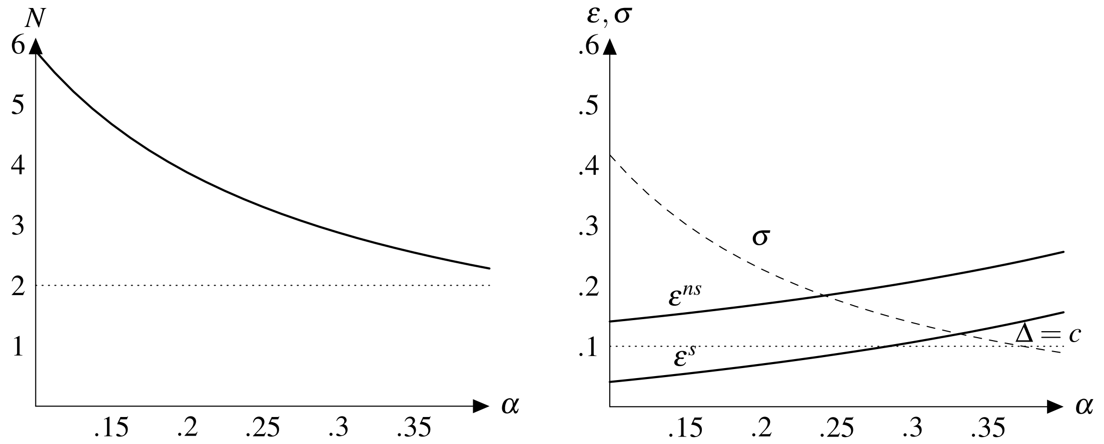

Number of deal finders in equilibrium

Equilibrium expected mismatch and share of savvy types.

Figure 5: Endogenous entry, with $c = 0.1$, $s = 0.02$