Metasearch and market concentration

Renaud Foucart*

February 24, 2020

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. Hence, if an exogenous share of consumers do not compare, more competition hurts them. More competition however also increases the incentives for consumers to compare. A higher share of informed consumers in turn increases the search effort of intermediaries. If consumers are ex-ante identical and rationally choose whether to become informed, the total effect of a higher number of intermediaries is to make each of them (weakly) choosier. Moreover, it always decreases the price offered by sellers. Allowing intermediaries to bias their advice by making sponsored links prominent has a similar effect of making all consumers better off in expectation.

Keywords: search, advice, competition

JEL: D43, D83, L13, L86

1 Introduction

Consumers often rely on intermediaries to help them find the product that best suits 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 in this market is whether consumers benefit from having a large number of intermediaries at their disposal. More precisely, could limiting entry or, to the contrary, mergers of intermediaries increase consumer welfare and market efficiency?

*Lancaster University Management School, r.foucart@lancaster.ac.uk ; I am grateful to Mark Armstrong, Alessandro Bonatti, Alexandre de Corniere, Paula Herrera Idarraga, Heiko Karle, Oliver Pardo, Martin Pollrich, David Rietzke, Andrew Rhodes, the editor Yongmin Chen, and two anonymous referees for useful comments. Financial support by Deutsche Forschungsgemeinschaft through CRC TRR 190 is gratefully acknowledged.

In this paper, I show that higher market concentration helps to protect uninformed buyers if consumer information is a "behavioural" trait: an exogenous share of buyers do not understand the market, regardless of how much they could benefit from becoming informed. If consumer information is a rational choice of ex-ante identical buyers however, higher market concentration lowers consumer welfare. The reason is that it decreases the incentive a buyer has to compare recommendations - the "metasearch" among intermediaries searching on one's behalf-, and the positive externalities it generates on the others by doing so. This in turn has three consequences. First, the quality of advice provided by intermediaries becomes (weakly) lower. Second, if intermediaries are less choosy, demand in the product market is less elastic and prices increase. Third, for a given choosiness of intermediaries, a lower share of informed consumers makes demand less elastic in the product market, and increases prices.

In terms of aggregate market efficiency, there is therefore a trade-off between better advice for consumers and the additional cost of intermediaries investing in advice and receiving smaller market shares. Finally, I show that introducing the possibility for deal finders to bias their results and offer prominent sponsored recommendations alongside their truthful organic ones actually helps ex-ante identical consumers. The reason is - again - that the existence of "fake" advice makes it more valuable for consumers to compare their options.

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 in these markets is a long-lasting debate. Historically, intermediaries have been subject to limitations and certifications to protect consumers from dishonest advisors. For instance in the UK, from 1977 to 2005, the Insurance Broker Registration Act limited entry to the market to make sure no deal finder was acting as a representative of a single insurance company.[1] 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 seem to often converge towards very concentrated structures making the question of excessive entry less relevant.[2] For instance, in 2015 in the US, Expedia[3] (Expedia.com, tripadvisor.com, orbitz.com, ho-

[1] See Insurance Brokers (Registration) Act 1977, 11.1.(c). Similarly, most US states require a special license to be a recognized broker.

[2] Malik (2015) summarizes the dynamic of concentration in 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."

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

tels.com, venere.com, trivago.com,...) and Priceline[4] (priceline.com, kayak.com, booking.com,...) controlled 95 percent of the online travel-marketplace after a number of successful fusions and acquisitions.[5]

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[6] 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.[7]

I set up a model in which a large number of consumers want to buy one unit of a product in a market with a large number of horizontally differentiated sellers. 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 and show that my results are robust to this alternative setup.

I start by solving a benchmark model in which consumers are of two exogenous 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. I also start by making 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 in a competitive market for advertising, or if they collect fixed commissions.[8] I find the "direct" effect that lower market concentration decreases the

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

[5] 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.

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

[7] 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

[8] 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

quality of advice. The intuition behind that result 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.

Then, I solve the model for endogenous consumer information. I first derive a classic result from this literature: the existence of search externalities (Armstrong, 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. Thus, in any equilibrium, not enough consumers choose to be informed. I find that lower market concentration has the indirect effect that savviness matters more, hence increasing the incentives to become informed. This indirect effect outweighs the direct one: as more consumers are savvy, demand in the product market is more elastic, and prices decrease. Hence, the main result of the paper: lower concentration in the market for intermediaries actually benefits all consumers when consumer information is endogenous.

I compare this "Varian" setting to one in which consumers bear a linear cost of non-sequentially observing an additional deal finder, in the spirit of Burdett and Judd (1983). I find that the main result of the paper holds, albeit the only channel through which more competition benefits consumers is price competition.

I then compare the welfare gain for consumers to aggregate market efficiency. I identify the following trade-off: while more competition in the market for deal finders always makes consumers better off, it also increase the aggregate costs for deal finders. Hence, while the number of deal finders maximizing aggregate welfare is often not equal to 2, it is not an infinite number either.

Finally, I solve two modified versions of the model. In the first one, deal finders auction prominent sponsored links to sellers, displayed alongside their truthful recommendation. As those "fake" advices decrease the expected payoff of uninformed consumers, they increase the incentives to become informed. Hence, if the information choice of consumers is endogenous, the existence of sponsored links benefits all consumers in expectation. In the second, I allow for heterogeneous costs for consumers to be informed and show that this intermediary case between exogenous and endogenous consumer information yields more balanced results.

Related literature:

To the best of my knowledge, this paper is the first to study how concentration in the market

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.

for intermediaries affects the incentives to invest in the quality of the advice they offer.

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). The novelty of my approach is to add competition on the side of the deal finders, and to study different types of buyers.

I consider a world in which consumers face a consideration set (Eliaz and Spiegler, 2011a), but this set is not directly determined by competition among sellers. Buyers instead 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 related models, Eliaz and Spiegler (2011b) study competition among two search engines, taking the quality of search ability as given and Taylor (2013) studies the trade-off faced by two search engines offering both organic and sponsored links and choosing how precise their organic advice should be.

In the case of online platforms, Karle *et al.* (2017) study competition among platforms charging fees for sellers to compete. Sellers want to be active on a popular platform to be matched with more buyers, but also want to avoid competing with too many similar sellers. There is however no active role in providing search quality for the platforms. Edelman and Wright (2015) study intermediaries competing by investing in a technology increasing the utility consumers get from a given product. Those benefits are however not linked to the quality of search, and therefore do not affect consumer information.

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 1 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). 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.[9]

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 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 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 behaviour of deal finders. I assume the willingness to pay for the information extracted from consumers to be constant, so that a consumer buying from a deal finder generates a revenue normalized to 1 for this deal finder.[10] I study another source of revenue for deal finders in Section 7.1: sponsored links. In the main part of 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

[9]The representation of ε as 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.

[10]An additional reason why a sale benefits deal finders is studied by Prufer and Schottmüller (2018): the cost of quality production is decreasing in the amount of information gathered about user preferences.

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.[11]

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. In the benchmark case (section 3), I take this share as given. This corresponds to the "behavioural" idea that consumers have different abilities to process information. Then, in section 4, I assume all consumers are ex-ante identical and only choose to become informed if it is in their interest to do so, based on rational expectations. The cost of becoming savvy is $c > 0$. I consider two types of consumer information, corresponding to the standard models of Varian (1980) - in the main part of the paper - and Burdett and Judd (1983) - in section 5 - respectively.

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.[12]

2.3 Sellers

There is a continuum of horizontally differentiated sellers of mass 1. Assuming this form of monopolistic competition simplifies a lot the analysis, as it implies that no two deal finders rec-

[11] In the case of online advice, a perhaps more realistic assumption - albeit slightly less tractable - 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.

[12] An alternative way of modeling the problem would be to assume a consumer searches sequentially among deal finders having searched sequentially among deals on her behalf and recommending one of them. A problem with this approach however would be that if deal finders search, they will at equilibrium choose the same search threshold as consumers. Indeed, it makes no sense to have a higher threshold and recommend a product that the consumer would never buy. It makes no sense either to have a lower threshold, as it would not increase the probability of making a sale. Hence, no consumer would ever compare deal finders at equilibrium.

ommend 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 expected sales made to a mass 1 of non-savvy consumers, 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 expected sales to a mass 1 of savvy consumers, 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. 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 $\phi(\eta) = f(\varepsilon)$. As I assume a continuum of sellers, an individual price deviation $p_j \neq p$ only affects the expected demand of seller j.

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 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 $p_i = p$.

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 a 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 costs is:

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

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] This expression is constant in ε when $\varepsilon > w$ since in that case a deal finder only sells to non-savvy consumers.

All players choose the strategy that maximizes their utility given their expectation of other players' strategies. 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. The two results are sufficient to characterize a unique symmetric equilibrium in the benchmark case with exogenous share of savvy consumers. In the endogenous case, given the search cutoff of deal finders, symmetric price, and share of savvy consumers σ, each consumer chooses whether or not to be informed at a cost $c > 0$. I find 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. In the endogenous case, 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. The equilibrium share of informed consumers is σ.

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.

3 The benchmark case: exogenous share of savvy consumers

In this section, I first study the equilibrium search of deal finders, and then the price, taking as given consumer information.

[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.

3.1 Search cutoff

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

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

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

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

by using the fact that $\pi(w) = \frac{1-\sigma}{N}$. 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, the marginal incentives to invest are unaffected and equation (5) would remain identical.

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.*

The formal proof is in Appendix. 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 N, 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. The qualifier "weakly" in the proposition corresponds to the possibility of $w \geq b$, so that there is no search at all. Through the paper, I focus on cases where $w < b$, so that the delegated search problem has an interior solution.

The expected mismatch of a non-savvy consumer ε^{ns} is a random draw on the interval $[0, w]$, where w is the threshold at which deal finders stop searching. A non-savvy consumer is always worse off when this threshold increases. The fact that the share of savvy types σ 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. The direct effect of lower market concentration (higher N) on the deals received by the non-savvy types is thus negative. Because, when there are more deal finders, each deal finder searches with lower intensity, 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.

The expected mismatch of a savvy type ε^s is the value of the minimum of N independent 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 then given by the standard formula

$$\varepsilon^s = \int_0^w N(1-G(\varepsilon))^{N-1}\varepsilon g(\varepsilon)d\varepsilon. \tag{6}$$

The impact of market concentration N on the expected mismatch value of a savvy type ε^s is composed of two effects. On the one hand, a higher number of deal finders makes each deal finder less choosy (higher w). On the other hand it also increases the number of options a savvy consumer can choose from. It is possible to show (see Appendix C) that in the special case where the mismatch function ε follows a uniform distribution the first effect always dominates, so that ε^s is increasing in N.

Two well-documented consequences follow immediately from the observation of the two mismatch values. 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 1. I assume a symmetric market price p, and study the optimal price p_i chosen by a seller i.

Lemma 1 *The expected demand a seller i receives from a mass* 1 *of non-savvy consumers is given by*

$$D^{ns}(p_i, p, N) = \frac{F(w+p-p_i)}{F(w)}. \tag{7}$$

The formal proof is in appendix. This expression corresponds to the probability of being selected by a deal finder following the search strategy defined in (5), with i being the only seller off the price equilibrium path, so that the density of $\eta_j = \varepsilon_j + p - p_j$ for all sellers $j \neq i$ 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 I assume a continuum of sellers of mass 1, the probability of being selected by two deal finders is zero.

Consider now savvy consumers. To be selected by a savvy consumer, a seller must be recommended by a deal finder, so that the maximum possible value of its mismatch is $\varepsilon = w+p-p_i$. As noted above, this happens with probability $F(w+p-p_i)$. Conditional on being selected by a deal finder, this seller offers the best deal out of N recommendations if $\varepsilon - p + p_i$ is lower than $N-1$ independent random draws of the mismatch function with density $f(\varepsilon)$ between 0 and w. I denote this probability by $r(p_i, p, N-1)$.[14]

Lemma 2 *The expected demand a seller i receives from a mass 1 of savvy consumers is given by,*

$$D^s(p_i, p, N) = ND^{ns} r(p_i, p, N-1). \tag{8}$$

The formal proof is in the Appendix. The expression (8) corresponds to N times the probability of being selected by a given deal finder, multiplied by the probability of offering a better deal than the $N-1$ other selected sellers. Again, the assumption of a continuum of sellers considerably simplifies the computations, as the probability that a given seller is selected by two deal finders is zero.

Lemma 3 *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 the search cutoff w and decreasing in the share of savvy consumers σ.*

The formal proof is in the Appendix. 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

$$p = \frac{-D(p,p,N)}{D_{p_i}(p_i,p,N)} = \frac{-D(p,p,N)}{\sigma D^s_{p_i}(p_i,p,N) + (1-\sigma)D^{ns}_{p_i}(p_i,p,N)}, \tag{9}$$

where $D(p,p,n) = 1$ as all sellers post an identical price in equilibrium. The difference with standard models of monopolistic competition is that there are two parts in the demand, one being

[14] $r(p_i, p, N-1) = \int_0^{w+p-p_i} \frac{f(\varepsilon)}{F(w+p-p_i)} \left(\frac{F(w) - F(\varepsilon - p + p_i)}{F(w)} \right)^{N-1} d\varepsilon.$

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 increases, the most elastic part of the demand becomes more important, so that prices decrease.

While lower market concentration makes each deal finder less choosy, there is still a potential benefit from competition on the market for deal finders: to provide more information to savvy consumers. As in this setting savvy consumers observe all deal finders, they are not necessarily worse off, because the additional information can compensate for the lack of precision of each intermediary. This, in turn, can lead to higher or lower price depending on whether the lower elasticity of demand for non-savvy consumers is more important than the (possibly) higher elasticity of demand for savvy ones. It is possible to show (see Appendix C) that in the special case where the mismatch function ε follows a uniform distribution the equilibrium price always increase with N when σ is exogenous.

This setting with exogenous consumer information may well illustrate the "behavioural" perspective that some consumers do not understand the market. It is however natural - as it is the case in Varian (1980) - to let the share of informed consumers be the result of utility maximization with rational expectations on the expected mismatches and prices in equilibrium. This is what I study in the next section.

4 Metasearch: Endogenous 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 7.2 that this assumption is not innocuous, as allowing for different consumers to have different costs of becoming savvy may revert the results. Define

$$\Delta = \varepsilon^{ns} - \varepsilon^{s}, \tag{10}$$

as the expected premium (in terms of expected mismatch) paid by uninformed consumers.

Lemma 4 *The difference between the expected mismatch received by a non-savvy and a savvy* *consumer Δ is increasing in the search cutoff of deal finders, $\frac{d\Delta}{dw} > 0$. This result holds for the standard log-concave distributions for which a closed-form expression for the cdf exists.*[15] *For all other continuous distributions it holds for N sufficiently large.*

The formal proof is in appendix, and corresponds to the idea that one benefits more from comparing two "bad" advices than from comparing two "good" advices. Intuitively, this is not particularly surprising as we know (see for instance Burdett, 1996) that log-concavity of the density function is a sufficient condition for the variance of a right-truncated distribution to increase when the truncation point (w here) increases.

If there exists an interior solution, the equilibrium share σ is found by solving

$$\Delta(\sigma) = c, \tag{11}$$

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 2 *If the market price p is symmetric, there exists a unique equilibrium share of savvy consumers. If σ has an interior solution, this share is increasing in the number of deal finders (N).*

The formal proof is in Appendix. This Proposition describes the indirect effect of market concentration. It follows from proposition 1 and lemma 4 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 valuable 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. As becoming informed is a positive externality on all all consumers, it holds that in equilibrium σ is always too low as compared to what would be consumer efficient (see Appendix E for a discussion).

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

[15] Uniform, exponential, logistic, extreme value, Laplace, power functions with cdf $\frac{F(\varepsilon)}{F(w)} = \frac{x^c}{w^c}$ with $c > 1$, and Weibull distribution, see Bagnoli and Bergstrom, 2005, p.455

Thus, it is possible to characterize the impact of the exogenous parameter N on the equilibrium mismatch and price for all consumers. As all consumers are ex-ante identical, and as consumers choose to become savvy up to the point where

$$u^s - c = u^{ns}, \tag{12}$$

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

$$u = u^s - c = u^{ns} = v - p - \varepsilon^{ns}. \tag{13}$$

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

Proposition 3 *When the share of informed consumers σ is endogenous and has an interior solution, 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 2. 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 2, 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). This effect only exists when σ has an interior solution. Else, we go back to the benchmark case of an exogenous share of informed consumers.

It is possible to show that the latter effect dominates. As 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' recommendations of a similar quality is always better than comparing N recommendations. Hence, the only solution is that $w' < w$, the higher number of deal finders actually makes all deal finders more selective through the indirect effect.

I illustrate the different effects on Figure 1, using a uniform distribution. The complete res-

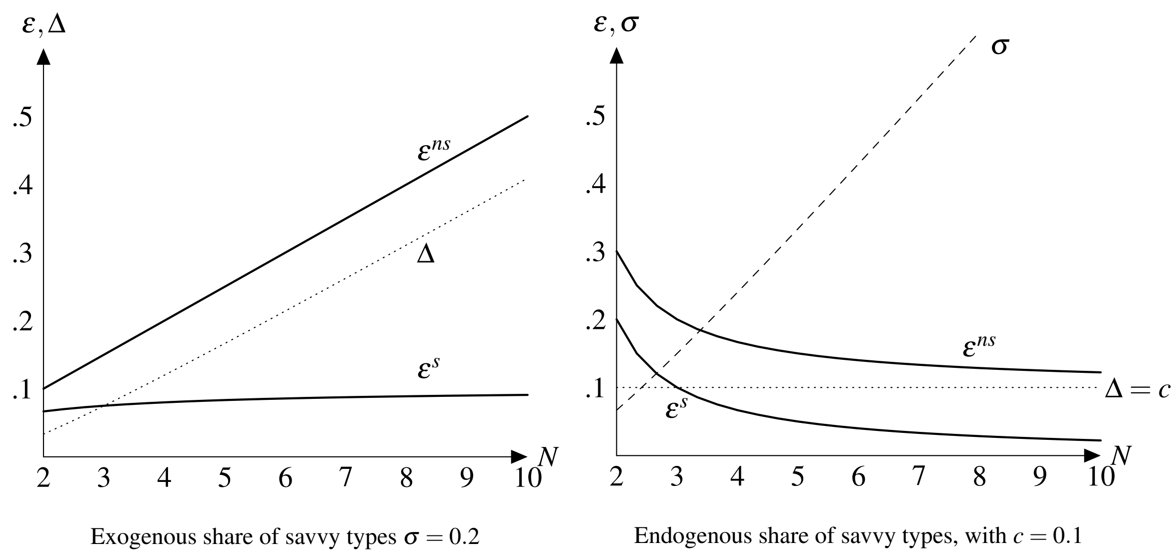

Figure 1: The impact of market concentration on expected mismatch in the Varian setting, with $s = 0,02$, $F(\varepsilon) = \varepsilon$.

olution of the uniform case is in Appendix C. 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.[16] 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.

Proposition 4 *When the share of informed consumers σ is endogenous and has an interior solution, the symmetric equilibrium price p is decreasing in the number of intermediaries.*

[16] 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.

The formal proof is in the Appendix, and is driven by a combination of three effects. First, sellers have an incentive to offer lower prices if deal finders become choosier. This is the case when the number of deal finders is higher, by proposition 3. If deal finders become more choosy, sellers have to offer better prices in all segments of the market, in order to have a chance of being selected. The consequence is one of a virtuous circle: the choosier the deal finders are, the more a seller wants to provide a low price. Second, sellers 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. Third, sellers decrease the prices if the share of savvy consumers increases, which strictly holds when the number of deal finders increase, as from proposition 2. As fewer consumers pick a deal finder at random, the competitive segment of the market matters more to sellers.

5 A Burdett and Judd setting

In this section, I briefly describe the results of an alternative "Burdett and Judd (1983)" setting, in which 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. Hence, for a given strategy of deal finders w and sellers p, the number of deal finders has no impact of the information of consumers.

The Varian and the Burdett and Judd settings are two polar cases aiming at capturing two different understandings 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. I compare those two polar cases when studying the welfare impact of market concentration in the next section.

In the Burdett and Judd setting, a share σ of savvy consumers observe two deals. Hence, (3) becomes

$$\pi(\varepsilon) = \left(\frac{2\sigma}{N} \frac{F(w) - F(\varepsilon)}{F(w)} + \frac{1-\sigma}{N} \right). \tag{14}$$

As in the original Burdett and Judd (1983) paper, there is an additional condition that search costs are sufficiently low for consumers to randomize between observing 1 and 2 deal finders (and not more).

Proposition 5 *In a Burdett and Judd setting:*

1. *With an exogenous share of savvy consumers σ, a higher number of intermediaries increases the expected mismatch of both types of consumers, and increases the equilibrium price.*
2. *With an endogenous share of savvy consumers, the expected equilibrium mismatch of both types of consumers is unaffected by the number of intermediaries, and the price decreases with the number of intermediaries.*

The formal proof is in the Appendix. The intuition behind the first result is very similar to the Varian case. The main difference is that as consumers compare either 1 or 2 options, there is no direct benefit from a higher number of deal finders for savvy consumers.

Consider now the second result. As a higher number of deal finders N decreases the search effort of deal finders, 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 (lemma 4). 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 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. 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. Thus, N has no direct influence on price. It however has an indirect impact through the share of savvy consumers σ: when N increases, the most elastic segment of the market becomes more important. The increasing share of savvy consumers is thus the only effect driving the result of lower market concentration leading to lower prices in the Burdett and Judd setting.

6 Welfare and market concentration

I now turn to the welfare impact of market concentration. I start with consumer welfare, a direct consequence of the above results. I then discuss aggregate welfare, defined as the sum of consumers, deal finders, and sellers payoffs.

6.1 Consumer welfare

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

Corollary 1 *When the share of informed consumers σ is endogenous and has an interior solution, lower market concentration increases the expected utility of all consumers. This result holds in both the Varian and the Burdett and Judd setting.*

The proof is straightforward from propositions 3, 4 and 5. I report in table 1 the impact of market concentration in the Varian and the Burdett and Judd settings. As the expected mismatch received by non-savvy consumers is weakly decreasing in N, and as prices decrease with N the total effect of lower market concentration on consumer welfare is always positive when metasearch is a rational choice of ex-ante identical consumers.

Table 1: Effect of a higher number of deal finders N.

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

Using the uniform distribution $F(\varepsilon) = \varepsilon$ and the same parameters as on Figure 1, Figure 2 represents the sum of the mismatch and the price effect on the equilibrium consumer welfare (see again Appendix C for the computations). The dashed lines represent the expected mismatch of a non-savvy consumer and the expected price (both decreasing with N). The solid lines represents the consumer welfare u, equal to $v - p - \varepsilon^{ns}$, increasing with N. On the left-hand side is the Varian setting. 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.

On the right-hand side is the Burdett and Judd setting where consumers compare either one or two options. We see that only the price effect matters, and that this effect is less important than in the Varian setting. This suggests that whether we consider consumer information as a general cost of understanding the market or as a linear cost of acquiring information influences to what extent higher market concentration hurts consumer welfare.

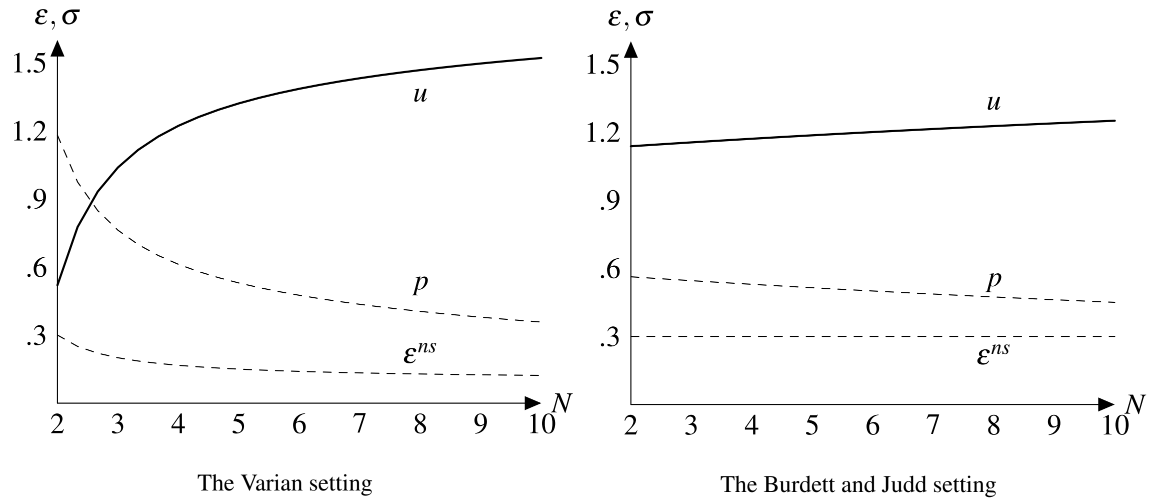

Figure 2: Equilibrium price and consumer welfare, with $s = 0,02$, $F(\varepsilon) = \varepsilon$, $c = 0.1$, $v = 2$.

6.2 Aggregate welfare

Up to now, I have focused on consumer welfare, as this is often the objective function of competition policy. Taking into account the welfare of sellers and deal finders however yields another tradeoff. Define aggregate welfare W as the sum of consumer utilities, sellers and deal finders profit:

$$W = (1-\sigma)(v - p - \varepsilon^{ns}) + \sigma(v - \varepsilon^{s} - c - p) + p + N(\frac{1}{N} - \frac{s}{F(w)}). \tag{15}$$

As the market is fully covered, prices cancel out and are not relevant to the welfare analysis. Moreover, at equilibrium $\varepsilon^{s} + c = \varepsilon^{ns}$ and the expression simplifies to

$$W = v - \varepsilon^{ns} - N(\frac{1}{N} - \frac{s}{F(w)}). \tag{16}$$

If there exists an interior solution, the number of deal finders that maximizes aggregate welfare solves $\frac{dW}{dN} = 0$,

$$\frac{d\varepsilon^{ns}}{dN} = \frac{d\frac{Ns}{F(w)}}{dN}. \tag{17}$$

We immediately see that in the Burdet and Judd framework - when consumers compare either one or two deal finders - as $\frac{dw}{dN} = 0$, there is no interior solution, and the number of deal finders maximizing aggregate welfare is the smallest that guarantees competition, $N = 2$.

In the Varian framework, $\frac{d\varepsilon^{ns}}{dN} < 0$ and $\frac{dw}{dN} < 0$, so that the solution implies a trade-off between a better matching quality and the cost paid by intermediaries. In particular, taking deal finders and sellers into account limits the idea that a very large number of deal finders is optimal. Indeed, under a consumer standard, as price and cutoff always decrease with N, there is no limit to how many deal finders should enter as long as the share of savvy consumers is below 100%.

Taking into account the costs of deal finders, the optimal number of intermediaries can never go towards infinity, and can even go back to the minimum in the model $N = 2$. The complete expression is pretty long to reproduce here, but it is easy to obtain numerical solutions. With a uniform distribution over $[0,1]$ - in both the Varian and the Burdett and Judd setting - the sum of all payoffs is given by:

$$W = v + 1 - \frac{w}{2} - \frac{s}{w}, \tag{18}$$

where 1 is the revenue of the deal finders. Replacing w by the equilibrium value in the Varian setting and differentiating with respect to N, the maximum aggregate welfare corresponds to the solution to

$$\frac{s - N(N+2)s}{2c(N+1)^2} + \frac{2c}{(N-1)^2} = 0. \tag{19}$$

With the parameters used in the linear examples of the paper ($s = .02, c = 01$), the number of intermediaries maximizing aggregate welfare is $N = 2.54$. With a higher search cost for firms s, this number decreases. With higher information cost for consumers c, this number increases. I illustrate how aggregate welfare changes with consumer information costs on Figure 3. With $c = 0.2$, the number of firms that maximizes aggregate welfare is $N = 3.95$ and with $c = 0.3$ it is $N = 5.35$.

7 Extensions

7.1 Sponsored links

Consider a variant of the main model, in which deal finders can auction a "sponsored link," displayed prominently to all consumers on top of the "organic link" that is the truthful recommendation studied in the previous sections. While an organic link contains information of the form $\{\varepsilon_{ij}, p_j\}$, a sponsored link displays a price p_j as well as the information that it is "recommended" by the platform: the actual mismatch is obfuscated. I make this last assumption as it is hard to imagine why even a naïve consumer would ever buy from the sponsored link if she could immedi-

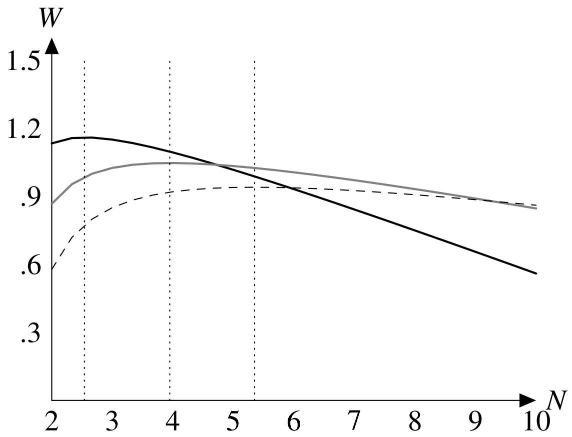

Figure 3: Aggregate welfare, with $s = 0,02$, $F(\varepsilon) = \varepsilon$, $v = 0.5$. The three curves correspond to different cost of consumer information: $c = 0.1$ (plain black), $c = 0.2$ (gray), $c = 0.3$ (dashed).

ately see the two mismatch values. It also makes computations more convenient, as I can rule out the possibility for the naïve consumer to buy from the sponsored link because it is better than the organic one.

I make the assumption that "non-savvy" consumers still pick a deal finder at random, and observe the sponsored link. With an exogenous probability $\gamma \in (0,1)$, they are "inattentive" and pick the sponsored link conditional on it not being more expensive than the rationally expected symmetric equilibrium price p.[17]

With probability $1-\gamma$ they see both the sponsored and the organic recommendation and pick the latter if the expected utility it yields is higher than the expected mismatch from a sponsored link. At the symmetric equilibrium price, this corresponds to the condition $w \leq \int_0^b f(\varepsilon)\varepsilon d\varepsilon$. I make the assumption that the latter condition always holds, this is that deal finders are not so bad that they could recommend a mismatch higher than the expected value of a random draw among all sellers. Savvy consumers always observe all the information. All consumers but the non-savvy and inattentive ones form rational expectations on the expected value of a mismatch from the sponsored link.

I start by studying two equivalent possibilities of auctioning prominence, per-display or per-sale. Per-display means that the seller pays a deal finder a fixed amount to be recommended, while per-sale means the amount is paid only if the buyer actually makes a purchase. I assume the auction to take place before the price is set, with sellers forming rational expectations on the symmetric equilibrium price.

[17] I could also assume that non-savvy consumers do not understand the price when they are inattentive but given the assumption that the market is fully covered, this would yield infinite prices for any strictly positive share of non-savvy consumers.

Lemma 5 *The expected profit of a seller from a sponsored link is equal to zero, and the expected revenue of a deal finder from auctioning a sponsored link is equal to* $\gamma\frac{(1-\sigma)}{N}p$ *in both types of auctions.*

The formal proof is in Appendix. As deal finders can auction the sponsored link and all sellers are ex-ante identical, deal finders can extract the entire expected surplus from those links. As prices are advertised, sellers do not gain from unilaterally increasing their price and specialize on the sponsored segment. Hence, the only demand that matters to the profit of sellers is the one from organic links.

The expected profit of a deal finder discovering a deal ε_{ij} when prices are symmetric is a small modification of (3):

$$\pi(\varepsilon) = \left(\sigma\left(\frac{F(w)-F(\varepsilon)}{F(w)}\right)^{N-1} + \frac{1-\sigma}{N}\right) + \gamma\frac{1-\sigma}{N}p. \tag{20}$$

It is easy to see that the modification does not change anything to the equilibrium condition for w identified in (5). Quite unsurprisingly, this result implies that if the share of savvy consumers is exogenous, non-savvy consumers receive a higher expected mismatch in the presence of sponsored links that in its absence.

I first look at the effect of sponsored links on prices, and then move to the incentives to become informed. I assume that a deal finder continues to benefit from the organic revenue of 1 when making a sale through a sponsored link. Assuming instead that her revenue comes exclusively from the auction p would not affect the equilibrium search behavior as the outcome remains independent of the mismatch found. However, the model would predict that when the price decreases a deal finder may simply find it not profitable to auction sponsored links at all (if $1 > p$).

The demand for a seller from organic links - the only ones that affect her profit - is given by

$$D(p_i, p, N) = (1-\sigma)(1-\gamma)D^{ns} + \sigma D^{s}, \tag{21}$$

where D^{ns} and D^{s} are defined as in the main model.

Proposition 6 *As compared to the model with organic recommendations only and for a given share of savvy consumers, the presence of sponsored links yields lower prices. The higher the level of inattention of non-savvy consumers* γ*, the lower the expected price. This result holds in both the Varian and the Burdett and Judd setting.*

The formal proof is in the Appendix and follows from the fact that there is less weight put on the most inelastic segment of the market. The more non-savvy consumers are captured by sponsored links, the less they become relevant to sellers, as those make more money from organic recommendations.

I can now turn to the incentives to become informed. For a given share of savvy consumers σ, the expected mismatch of a savvy consumer ε^s is as before, because the cutoff value w is not directly affected by the presence of sponsored links. The expected mismatch of a non-savvy consumer is however higher as there is now a probability of observing a random deal when inattentive. Hence, for a given share of informed consumers σ, the difference between the expected mismatch of a savvy and a non-savvy consumer is higher.

Proposition 7 *As compared to the model with organic recommendations only, the presence of sponsored links yields a higher share of savvy consumers and lower expected mismatch for all consumers. The higher the level of inattention of non-savvy consumers γ, the lower the expected mismatch. This result holds in both the Varian and the Burdett and Judd setting, as long as the share of savvy consumers σ has an interior solution.*

The formal proof is in the Appendix. This result follows the same logic as in the main model. Because the expected deal of a non-savvy consumer becomes less attractive, the incentives to become informed are higher. More consumers become informed up to the point where the difference between the two types of consumers becomes equal to the cost of being savvy c. As the difference for a given w is higher in the presence of sponsored links, the only way to have all consumers ex-ante indifferent is that σ increases. Hence, w decreases to a lower level than in the model without sponsored links. The logic is also reminiscent of Armstrong *et al.* (2009), who find in a model of price dispersion that some consumer protection policies - price caps and the possibility to opt out of advertisement lists - may make consumer worse off by decreasing their incentives to become informed.

Finally, as the share of savvy consumers increases with the presence of sponsored links and inattention, there is an additional effect on prices. Hence, the following general result.

Corollary 2 *As compared to the model with organic recommendations only, the presence of sponsored links yields a higher expected utility for all consumers if the share of savvy consumers σ has an interior solution. The higher the level of inattention of non-savvy consumers γ, the higher the expected utility.*

The dynamic that leads to higher utility is that the risk of buying the wrong product by mistake leads to higher incentives of being informed. This yields a particularly extreme policy recommendation: any regulation that forces deal finders to be more transparent about the sponsored nature of some of their links (a lower γ) is actually counter-productive. Indeed, it decreases the incentives to become informed and therefore the search externality. A policy of improving the ability of consumers to process the information (a lower c and therefore a higher σ) is however always beneficial.

As for the main model, the assumption that all consumers are ex-ante identical is crucial. Another important assumption is that even inattentive consumers know what price to expect. If sellers were allowed to extract more surplus from naïve consumers by charging much higher prices, an interior solution to the share of savvy consumers would be less likely, and never exists if I keep the assumption of a fully covered market.

7.2 Heterogeneous 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 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 $c(\sigma) = \gamma\sigma$.

It is thus 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 in the market for deal finders is Pareto improving. The intuition is relatively straightforward, as allowing for heterogeneous costs of savviness is an intermediate case between assuming an exogenous share of σ and ex-ante identical consumers.

In the uniform case in the Varian setting, I can rewrite (11) as

$$\sigma = \sqrt{\frac{sN(N-1)}{2\gamma(N+1)}}, \tag{22}$$

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 (22) into the equilibrium search threshold w (equation (57) in Appendix C) yields

$$w = \frac{sN\sqrt{2\gamma(N+1)}}{\sqrt{sN(N-1)}}, \tag{23}$$

which can be shown to become increasing in N for $N \geq 1+\sqrt{2} \approx 2.41$. 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 rapidly starts to increase. Introducing ex-ante heterogeneity among consumers means that we need to consider the distributional impact of the level of concentration in the market for deal finders.

I illustrate this idea on Figure 4, by comparing the case studied in the Figure 1 with $c = 0.1$, 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 in 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 a higher market concentration until the number of deal finders is equal to $N = 3$, while the most able consumers always prefer a higher number of deal finders.

8 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 impact of competition on consumer welfare, as they also influence the price offered by sellers.

I show that, if the metasearch of consumers among deal finders is a rational choice, more

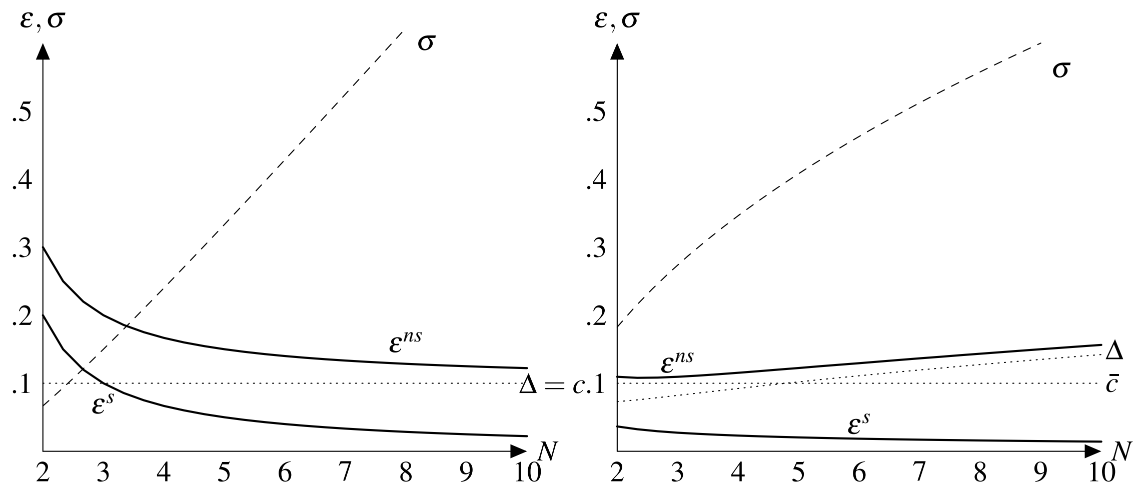

Endogenous share of savvy types, with $c = 0.1$ Endogenous share of savvy types, with $\gamma = 0.2$, $c = 0.2\theta$.

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

competition in the market for deal finders makes these intermediaries more choosy and increases price competition among sellers. Adding the profit of deal finders in the objective function of the regulator leads to a tradeoff between firms' profits and consumer surplus: the optimal number of deal finders is not necessarily 2, but each additional deal finders needs to bring a sufficiently high decrease in the expected mismatch received by consumers to justify the higher costs for the industry. Hence, this trade-off between the quality of advice and the aggregate cost for the industry providing it should be taken into account when considering the case for mergers and acquisitions in the market for advice.

More generally, this paper aims at contributing to the 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 sources. As long as a strictly positive share of consumers are informed and another one uninformed, the impact bad information has on internalizing the search externality dominates. Even "fake news" such as sponsored links actually benefit every consumer once equilibrium effects are factored in.

I also identify three 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 candidate equilibria are that either all consumers are informed, all consumers are uninformed, or all are indifferent. Second, the result of a Pareto improvement identified in corollary 1 relies on all consumers being ex-ante identical. Even without an exogenous share of naïve consumers, one needs to consider distributional effects if buyers have different exogenous abilities to process information. Third, in the case of organic recommendations, I take the revenue from a sale made possible by a deal finder as exogenous. Modifying this assumption could alter the results in two directions. If more deal finders yield lower per-sale revenue, this would decrease the quality of their advice. If a better quality of match increased the value of the sale for the deal finder, this could start a virtuous circle in which a higher quality of advice increases the incentives to provide better advice.

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.

—, VICKERS, J. and ZHOU, J. (2009). Consumer protection and the incentive to become informed. *Journal of the European Economic Association*, **7** (2-3), 399–410.

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

BAGNOLI, M. and BERGSTROM, T. (2005). Log-concave probability and its applications. *Economic theory*, **26** (2), 445–469.

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. (1996). Truncated means and variances. *Economics Letters*, **52** (3), 263–267.

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

DAVID, H. and GROENEVELD, R. (1982). Measures of local variation in a distribution: Expected length of spacings and variances of order statistics. *Biometrika*, **69** (1), 227–232.

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.

KARLE, H., PEITZ, M. and REISINGER, M. (2017). Segmentation versus agglomeration: Competition between platforms with competitive sellers.

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.

PRUFER, J. and SCHOTTMÜLLER, C. (2018). Competing with big data.

TAYLOR, G. (2013). Search quality and revenue cannibalization by competing search engines. *Journal of Economics & Management Strategy*, **22** (3), 445–467.

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 Proposition 1.

Proof. From (5), 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 $s \geq \sigma \int_0^b (1 - F(\varepsilon))^{N-1} f(\varepsilon) d\varepsilon = \frac{\sigma}{N}$.

■

Proof of Lemma 1

Proof. The probability that a seller i offering price p_i is below the threshold $P(\eta_{ij} < w)$ is equal to $P(\varepsilon_{ij} + p_i - p < w) = F(w + p - p_i)$. The probability that any seller setting the equilibrium price is below the threshold is $P(\varepsilon_{ij} + p - p < w) = F(w)$. At equilibrium, the mass of deal finders below the threshold is thus $F(w)$. The probability of being selected by a mass 1 of non-savvy consumers is equal to the probability of being below the threshold over the mass of sellers below the threshold, $D^{ns} = \frac{P(\varepsilon_{ij}+p_i-p<w)}{P(\varepsilon_{ij}+p-p<w)} = \frac{F(w+p-p_i)}{F(w)}$. ■

Proof of lemma 2.

Proof. A seller is chosen by a savvy consumer if it offers the best deal amongst N deals selected by deal finders. In order to do so, it first needs to be selected by a deal finder. The probability that a mass 1 of sellers is selected by a given deal finder is equal to $D^{ns}(p_i, p, N)$, and the probability of being selected by more than one deal finder for a given seller is zero (as there is a continuum of sellers). Hence, the probability of being selected by exactly one deal finder is equal to $ND^{ns}(p_i, p, N)$. The probability of offering a better deal η_{ij} than all other sellers conditional on being selected by a deal finder is $r(p_i, p, N-1)$. Hence, the probability for a seller to be selected by a mass one of savvy consumers is equal to $D^s(p_i, p, N) = ND^{ns}(p_i, p, N) r(p_i, p, N-1)$. ■

Proof of Lemma 3

Proof. The expected demand for a given seller i is

$$D(p_i, p, N) = (1 - \sigma) D^{ns} + \sigma D^s. \tag{24}$$

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

$$p = \frac{-D(p,p,N)}{D_{p_i}(p_i,p,N)} = \frac{-D(p,p,N)}{\sigma D^s_{p_i}(p_i,p,N) + (1-\sigma) D^{ns}_{p_i}(p_i,p,N)}, \tag{25}$$

where $D(p,p,n) = 1$ as all sellers post an identical price in equilibrium.[18]

It is easy to show that, with a continuum of sellers of mass 1,

$$D^{ns}_{p_i} = \frac{-f(w)}{F(w)}, \tag{26}$$

so that with $\sigma = 0$, 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 of $N-1$ independent draws over $[0,w]$, $r(p_i, p, N-1)$. The demand D^s is given by:

$$D^s = N D^{ns} r(p_i, p, N-1). \tag{27}$$

Define $\beta = N r(p_i, p, N-1)$, 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:

$$D^s_{p_i} = D^{ns}_{p_i} + D^{ns}\frac{d\beta}{dp_i}. \tag{28}$$

Using (26), (28), and he fact that at equilibrium $D^{ns} = D = 1$, the equilibrium price in (25) rewrites as

$$p = \frac{-1}{\sigma(D^{ns}_{p_i} + \frac{d\beta}{dp_i}) + (1-\sigma)D^{ns}_{p_i}} = \frac{1}{\frac{f(w)}{F(w)} - \sigma\frac{d\beta}{dp_i}}. \tag{29}$$

where $\frac{d\beta}{dp_i} < 0$, as $\frac{dr(p_i,p,N)}{dp_i} < 0$. Thus, we see immediately that the direct effect of a higher σ is to decrease the equilibrium prices.

As the distribution is log-concave, $\frac{f(w)}{F(w)}$ is decreasing in w. As the order statistics of a log-concave distribution are also log-concave, at the equilibrium price $p_i = p$, it holds that $\frac{d^2\beta}{dp_i dw} \geq 0$ - the (negative) impact of marginally increasing the price on the probability of offering the best deal is less important if the interval of the draw w is larger - so that a higher w leads to higher

[18] While I have not been able to formally derive a sufficiency condition for the profit function, it is possible to verify with standards log-concave density functions that the price indeed corresponds to a maximum.

equilibrium prices.[19] ■

Proof of Lemma 4

Proof. Denote by $E(X_{i:N})$ the expected value of the i-th order statistic in a sample of size N. As $E(X_{1:1}) = \int_0^w (1-\frac{F(\varepsilon)}{F(w)})d\varepsilon$ and $E(X_{1:N}) = \int_0^w N(1-\frac{F(\varepsilon)}{F(w)})^{N-1}\frac{f(\varepsilon)}{F(w)}d\varepsilon$, we can write

$$\Delta = \int_0^w (1-\frac{F(\varepsilon)}{F(w)})(1-N(1-\frac{F(\varepsilon)}{F(w)})^{N-2}\frac{f(\varepsilon)}{F(w)})\varepsilon d\varepsilon. \tag{30}$$

We are looking for a condition such that

$$\begin{aligned}\frac{d\Delta}{dw} =& \frac{f(w)}{F(w)}(\int_0^w \frac{F(\varepsilon)}{F(w)}d\varepsilon - \int_0^w N\frac{F(\varepsilon)}{F(w)}(1-\frac{F(\varepsilon)}{F(w)})^{N-2}\frac{f(\varepsilon)}{F(w)}\varepsilon d\varepsilon \\ &- \int_0^w N(N-2)\frac{F(\varepsilon)}{F(w)}(1-\frac{F(\varepsilon)}{F(w)})^{N-2}\frac{f(\varepsilon)}{F(w)}\varepsilon d\varepsilon + \int_0^w N(1-\frac{F(\varepsilon)}{F(w)})\frac{f(\varepsilon)}{F(w)}\varepsilon d\varepsilon) \\ &\geq 0\end{aligned} \tag{31}$$

As $\frac{f(w)}{F(w)} > 0$, and using the definition $E(X_{i:N}) = \int_0^w \frac{N!}{(i-1)!(N-i)!}(1-\frac{F(\varepsilon)}{F(w)})^{N-i}\frac{F(\varepsilon)}{F(w)}^{i-1}\frac{f(\varepsilon)}{F(w)}\varepsilon d\varepsilon$ this condition simplifies to

$$\int_0^w \frac{F(\varepsilon)}{F(w)}d\varepsilon \geq E(X_{2:N} - X_{1:N}). \tag{32}$$

The right-hand side is decreasing in N (adding a new observation can only decrease the difference between the smallest and the second smallest independent draw). Hence, we can immediately see that for all continuous distributions the result holds for N large enough, and we can focus on $N = 2$ when studying whether it always holds for log-concave distributions. Using the expression $E(X_{2:2} - X_{1:2}) = \int_0^w 2\frac{F(\varepsilon)}{F(w)}(1-\frac{F(\varepsilon)}{F(w)})d\varepsilon$ (see equation 2 in David and Groeneveld, 1982), the condition rewrites

$$\int_0^w 2(\frac{F(\varepsilon)}{F(w)})^2 d\varepsilon \geq \int_0^w \frac{F(\varepsilon)}{F(w)}d\varepsilon. \tag{33}$$

[19]The expression $\frac{d\beta}{dp_i}$ is equivalent to $-\frac{f(w)}{F(w)} - N\int_0^w \frac{(N-1)f(\varepsilon)f(\varepsilon)\left(\frac{F(w)-F(\varepsilon)}{F(w)}\right)^{N-2}}{F(w)F(w)}d\varepsilon$ so that the denominator of (29) rewrites as $(1-\sigma)\frac{f(w)}{F(w)} + \sigma N\int_0^w \frac{(N-1)f(\varepsilon)f(\varepsilon)\left(\frac{F(w)-F(\varepsilon)}{F(w)}\right)^{N-2}}{F(w)F(w)}d\varepsilon$. The first term $(1-\sigma)\frac{f(w)}{F(w)}$ is decreasing in w as the density is log-concave. For the second term, $(N-1)\frac{f(\varepsilon)}{F(w)}(\frac{F(w)-F(\varepsilon)}{F(w)})^{N-2}$ is the truncated density function of $X_{1:N-1}$, the first order statistic of $N-1$ random draws over $[0,w]$, and $\frac{f(\varepsilon)}{F(w)}$ is the truncated density of the mismatch, both log-concave.

As log-concave distributions are unimodal, the cdf cannot become convex after being concave (as this would imply the pdf to be first decreasing then increasing). The condition in (33) thus implies that the slope of the cdf $(\frac{F(\varepsilon)}{F(w)})' = \frac{f(\varepsilon)}{F(w)}$ must be sufficiently high for low values of ε (as this corresponds to a low value of $2(\frac{F(\varepsilon)}{F(w)})^2 - \frac{F(\varepsilon)}{F(w)}$) and sufficiently low for the highest values. Hence, this is a condition on the (increasing) cdf being sufficiently concave. For the log-concave distributions for which a closed-form expression for the cdf exists (uniform, exponential, logistic, extreme value, Laplace, power functions with cdf $\frac{F(\varepsilon)}{F(w)} = \frac{x^c}{w^c}$ with $c > 1$, and Weibull distribution, see Bagnoli and Bergstrom, 2005, p.455) it is possible to directly verify that the condition is satisfied for all truncated distributions on an interval $[0,w]$.

It is also possible to show that the condition continues to be satisfied even for some distributions without a log-concave density. The most obvious example is the power function with cdf $\frac{F(\varepsilon)}{F(w)} = \frac{x^c}{w^c}$ with $c < 1$, as the inequality in (33) rewrites $\frac{w}{1+3c+2c^2} \geq 0$ and is always satisfied for $c > 0$. This is however a distribution for which the pdf is log-convex but the cdf is log-concave.

To show that the conditions holds even for a looser condition on concavity, we can look at the Arc-Sine distribution, for which neither the cdf nor the pdf is log-concave, but where the inequality in (33)[20] is always satisfied on the domain of the distribution $w \in [0,1]$.[21] ■

Proof of Proposition 2

Proof. First, I show that there exists a unique equilibrium σ. $\Delta(\sigma)$ is continuous and decreasing. Indeed, from proposition 1, w decreases with σ. As neither σ nor p directly enters Δ, σ only affects Δ through its effect on w. From lemma 4, $\frac{d\Delta}{dw} \geq 0$, so that $\frac{d\Delta}{d\sigma} \leq 0$. 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, we know from lemma 4 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, $EX_{1:k}$, over an interval it is always true that $E(X_{1,k}) < E(X_{1,k'})$ if and only if $k' < k$.

[20] $w - \frac{w}{\arcsin(\sqrt{w})^2} + \frac{3\sqrt{(1-w)w}}{2\arcsin(\sqrt{w})} - \frac{1}{2} \geq 0$

[21] I have not been able to find a continuous density function with full support for which condition (33) is not satisfied. It is however possible to construct a discrete example. Consider a distribution with zero density anywhere on $[0,w]$ except for an atom at $\frac{P(\varepsilon=0)}{F(w)} = \frac{1}{4}$ and another at $\frac{P(\varepsilon=w)}{F(w)} = \frac{3}{4}$. Then, if $\frac{f(w)}{F(w)} > 0$ (if the pdf becomes continuous with positive density at w), the condition is not satisfied, as $2 \times (\frac{1}{4})^2 < \frac{1}{4}$: increasing w actually decreases the expected difference between $X_{1:1}$ and $X_{1:2}$.

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 3

Proof. I want to assess the impact of N on ε^{ns} for a given value of p. In any equilibrium where σ and w have an interior solution, for a given number of deal finders $\Delta(w,N) = \varepsilon^{ns}(w) - \varepsilon^{s}(w,N) = c$, where w is itself a function of N and σ. Consider $N' > N$ and assume w remains the same, it holds that $\Delta(w,N') = \varepsilon^{ns}(w) - \varepsilon^{s}(w,N') > c$, as $\varepsilon^{s}(w,N') < \varepsilon^{s}(w,N)$ - a savvy consumer is always better off in expectation by observing more deal finders, all other things held equal. Hence, as by lemma 4, $\frac{d\Delta}{dw} \geq 0$, the only possibility to reach an equilibrium is that σ increases until the search threshold of deal finders reaches some $w' < w$ such that $\varepsilon^{ns}(w') - \varepsilon^{s}(w',N') = c$. ■

Proof of Proposition 4

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

$$D^{s}_{p_i}(p,p,N)\frac{d\sigma}{dN} + D^{ns}_{p_i}(p,p,N)(-\frac{d\sigma}{dN}) + (1-\sigma)\frac{dD^{ns}_{p_i}(p,p,N)}{dN} + \sigma\frac{dD^{s}_{p_i}(p,p,N)}{dN} \leq 0. \tag{34}$$

The sum of the first two terms in (34) is always negative, as $\frac{d\sigma}{dN} \geq 0$ (proposition 2) 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 3), as the elasticity of demand is higher if the deal finders are more selective (lower w). ■

Proof of Proposition 5

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 are indifferent between observing 1 or 2 deal finders (because the marginal benefit of an additional observation is decreasing in the number of observations). I then need to verify that no consumer strictly prefers to observe 3 deal finders. For all the results taking N as given, the Burdett and Judd setting is a special case of Varian with $N = 2$. As for the Varian case, I start by taking σ as given and then make it endogenous.

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

$$\pi(\varepsilon) = \left(\frac{1-\sigma}{N} + \frac{2\sigma}{N} \frac{F(w) - F(\varepsilon)}{F(w)} \right), \tag{35}$$

so that in the second stage w solves

$$s = \left(\frac{2\sigma}{N} \int_0^w \frac{F(w) - F(\varepsilon)}{F(w)} f(\varepsilon) d\varepsilon \right), \tag{36}$$

with identical properties as in the "Varian" setting. In the Burdett and Judd setting, the number of deals observed by a savvy consumer is set to 2 so that only the impact of N on w affects ε^s, and savvy consumers receive a higher expected mismatch when N increases. The price solves a similar problem as in the Varian case, with as only difference

$$D^s = 2D^{ns} r(p_i, p, 1). \tag{37}$$

The effect on price is a combination of two effects. First, N has no direct influence on p. Second, the indirect effect is to increase w, and by lemma 3 an increase in w leads to higher prices.

2. The share of informed consumers σ solves

$$c = \int_0^w \varepsilon g(\varepsilon) d\varepsilon - \int_0^w 2(1 - G(\varepsilon)) \varepsilon g(\varepsilon) d\varepsilon, \tag{38}$$

with $g(\varepsilon) = \frac{f(\varepsilon)}{F(w)}$ and w from (36). As in the Varian setting, the difference increases with w (equation (33) in the proof of lemma 4 corresponds to the Burdett and Judd case), so that the indirect effect of a 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,

$$c > \int_0^w 2(1 - G(\varepsilon)) \varepsilon g(\varepsilon) d\varepsilon - \int_0^w 3(1 - G(\varepsilon))^2 \varepsilon g(\varepsilon) d\varepsilon. \tag{39}$$

In any equilibrium where σ and w have an interior solution, for a given number of deal finders $\Delta(w) = \varepsilon^{ns}(w) - \varepsilon^s(w) = c$, where w is itself a function of N and σ, and $\varepsilon^s(w)$ does not directly depend on N as savvy consumers always observe exactly 2 deal finders. Hence, for any $N' \neq N$, the only possibility to reach an equilibrium w' is that the equilibrium σ ensures $w' = w$ always holds. As $\frac{dw}{dN} = 0$, price decrease with N: the first two terms in (34) in the proof of proposition 4 are negative, while the last two terms are equal to zero.

Proof of Lemma 5

Proof.

1. I have assumed that the production cost of sellers is equal to 0. Hence, at a symmetric equilibrium price p, the expected profit of a seller paying a price q^d for a sponsored link "per display" is equal to

$$\pi^d = \gamma \frac{(1-\sigma)}{N} p - q^d. \tag{40}$$

   As sellers are symmetric, the optimal strategy in a second-price sealed-bid auction is to bid one's valuation, the equilibrium bid is thus equal to

$$q^d = \gamma \frac{(1-\sigma)}{N} p. \tag{41}$$

2. The expected profit of a seller paying a price q^s for a sponsored link "per sale" is equal to

$$\pi^s = \gamma \frac{(1-\sigma)}{N} (p - q^s). \tag{42}$$

   As sellers are symmetric, the optimal strategy in a second-price sealed-bid auction is to bid one's valuation, the equilibrium bid is thus equal to

$$q^s = p. \tag{43}$$

3. A seller always makes an expected profit of zero from the sponsored recommendation, as by setting any price $p' > p$, no consumer would ever buy from the sponsored link.

■

Proof of Proposition 6

Proof. Equation (25) becomes

$$p = \frac{-D(p,p,N)}{D_{p_i}(p_i,p,N)} = \frac{-D(p,p,N)}{\sigma D^s_{p_i}(p_i,p,N) + (1-\sigma)(1-\gamma) D^{ns}_{p_i}(p_i,p,N)}. \tag{44}$$

The equilibrium demand from buyers using organic links is now $D(p,p,N) = \sigma + (1-\sigma)(1-\gamma)$, so that the expression becomes (with β defined in the proof of lemma 3)

$$p = -\frac{\sigma + (1-\sigma)(1-\gamma)}{\left((\sigma + (1-\sigma)(1-\gamma))D^{ns}_{p_i} + \sigma \frac{d\beta}{dp_i} D^{ns}\right)}. \tag{45}$$

Hence, for any $\gamma > 0$ the price with sponsored link is lower than without. Differentiating with respect to γ yields

$$\frac{dp}{d\gamma} = \frac{(1-\sigma)\frac{d\beta}{dp_i} D^{ns}}{\left(\frac{d\beta}{dp_i} D^{ns} + D^{ns}_{p_i} - D^{ns}_{p_i}(1-\sigma)\right)^2} < 0. \tag{46}$$

■

Proof of Proposition 7

Proof. The expected mismatch of a non-savvy consumer becomes

$$\varepsilon^{ns}_{sponsored} = (1-\gamma)\varepsilon^{ns} + \gamma \int_0^b f(\varepsilon)\varepsilon d\varepsilon > \varepsilon^{ns}. \tag{47}$$

Thus, $\Delta_{sponsored}(w,\sigma) > \Delta(w,\sigma)$ and the only possibility to have $\Delta_{sponsored}(w,\sigma) = c$ is to have a lower search cutoff $w' < w$. As at equilibrium it must hold that $\varepsilon^s + c = \varepsilon^{ns}_{sponsored}$ and γ does not influence ε^s, all buyers are better off in expectation with w' than with w. Indeed, as $w' < w$, at equilibrium, $\varepsilon^s(w') < \varepsilon^s(w)$. As $\frac{\partial \varepsilon^{ns}_{sponsored}}{\partial \gamma} > 0$, it also holds that the equilibrium w decrease with γ. Note that this result holds both in the Varian and in the Burdett and Judd setting. ■

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. 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

$$\pi(q_i,q) = \left(\sigma\left(\int_0^b f_{(1),q_i}(\varepsilon)(1-F_{(1),q}(\varepsilon))^{N-1} d\varepsilon\right) + \frac{1-\sigma}{N}\right), \tag{48}$$

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:

$$\pi(q_i,q) = \left(\sigma\left(\int_0^b q_i f(\varepsilon)(1-F(\varepsilon))^{q_i-1+q(N-1)}d\varepsilon\right) + \frac{1-\sigma}{N}\right). \tag{49}$$

The symmetric equilibrium q is such that

$$\frac{d\pi(q_i,q)}{q_i} = s, \tag{50}$$

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 (49) and (50) 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,

$$\varepsilon^{ns} = \int_0^b q(1-F(\varepsilon))^{q-1}\varepsilon f(\varepsilon)d\varepsilon, \tag{51}$$

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,

$$\varepsilon^{s} = \int_0^b Nq(1-F(\varepsilon))^{Nq-1}\varepsilon f(\varepsilon)d\varepsilon, \tag{52}$$

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 decreases u^s if and only if

$$-\frac{\partial \varepsilon^{s}}{\partial N} \geq \frac{\partial \varepsilon^{s}}{\partial q}\frac{\partial q}{\partial N}. \tag{53}$$

From (51) and (52) it follows directly that $\frac{\partial \Delta}{\partial N} > 0$, so that $\frac{d\sigma}{dN} > 0$. Proposition 3 thus holds: $\frac{dq}{dN} \geq 0$. Indeed, with $N' > N$, the gain from comparing Nq and q deals and $N'q'$ and q' deals must be equal, which is only possible with $q' > q$.

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

$$D^{ns}(p_i, p, N) = q \int_0^b f(\varepsilon)(1 - F(\varepsilon - p + p_i))^{q-1} d\varepsilon, \tag{54}$$

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

$$D^{s}(p_i, p, N) = Nq \int_0^b f(\varepsilon)(1 - F(\varepsilon - p + p_i))^{Nq-1} d\varepsilon, \tag{55}$$

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 $D^s = D^{ns} = 1$, D^s is more elastic. Thus, following a similar reasoning as for the sequential search, p decreases with N if

$$D^s_{p_i}(p,p,N)\frac{d\sigma}{dN} + D^{ns}_{p_i}(p,p,N)\frac{d(1-\sigma)}{dN} + (1-\sigma)\frac{dD^{ns}_{p_i}(p,p,N)}{dN} + \sigma\frac{dD^{s}_{p_i}(p,p,N)}{dN} \leq 0, \tag{56}$$

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

Appendix C: uniform example

The Varian setting.

In the special case of uniformly distributed ε over $[0,1]$, in the Varian case, if there is an interior solution,

$$w = \frac{sN}{\sigma}, \tag{57}$$

where N is the number of deal finders in the market, and σ the endogenous share of savvy consumers. The expected mismatch of a non-savvy consumer is given by

$$\varepsilon^{ns} = \frac{w}{2} = \frac{sN}{2\sigma}. \tag{58}$$

More deal finders make the savvy type worse off as ε^s simplifies to

$$\varepsilon^s = \frac{sN}{\sigma(1+N)}. \tag{59}$$

The demand can be rewritten as

$$D = D^{ns}((1-\sigma) + \sigma N r(N-1)). \tag{60}$$

Denoting $\beta = (1-\sigma) + \sigma N r(N-1)$, and noting that at a symmetric equilibrium $D = D^{ns} = \beta = 1$, and that in the uniform case $r(j) = \frac{w\left(\frac{p-p_i+w}{w}\right)^{j+1}}{(j+1)(p-p_i+w)}$, I find

$$\begin{aligned} D_{p_i} &= D^{ns}_{p_i} + 1\frac{d\beta}{dp_i} \\ &= -\frac{1}{w} - \frac{1}{w}(N-1)\sigma, \end{aligned} \tag{61}$$

so that

$$\begin{aligned} p &= -\frac{1}{-D_{p_i}} \\ &= \frac{w}{1+(N-1)\sigma}. \end{aligned} \tag{62}$$

The premium for being informed rewrites $\Delta = \frac{sN(N-1)}{2\sigma(N+1)}$ and (11) solves

$$\sigma = \frac{sN(N-1)}{2c(N+1)}, \tag{63}$$

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

Plugging (63) into (57) yields

$$w = \frac{2c(N+1)}{N-1}, \tag{64}$$

if σ has an interior solution and $w = sN$ if $\sigma = 1$.

I have shown in (62) that the equilibrium price solves $p = \frac{w}{1+(N-1)\sigma}$, with w and σ defined in (63) and (64). 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, (62) simplifies to

$$p = \frac{4c^2(1+N)^2}{2c(N^2-1)+(N-1)^3 Ns}. \tag{65}$$

Note that equation (65) 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 (65), but to $p = s$.

Introducing sponsored links in section 7.1, the equilibrium price is

$$p = \frac{w((1-\gamma)+\sigma\gamma)}{1+(N-1)\sigma-\gamma(1-\sigma))}, \tag{66}$$

with $\frac{dp}{d\gamma} = -\frac{(N-1)(1-\sigma)\sigma w}{(1-\sigma+(N-1+\gamma)\sigma)^2} < 0$ when σ is taken as given.

The equilibrium share of informed consumers is

$$\sigma = \frac{sN((1+N)\gamma-2)}{(1+N)(\gamma-1+2c)} \tag{67}$$

if σ has an interior solution. The search cutoff is given by

$$w = \frac{(1+N)(2c-(1-\gamma)))}{\gamma N-(1-(1-\gamma)}. \tag{68}$$

The Burdett and Judd setting.

In the special case of uniformly distributed ε over $[0,1]$, in the Burdett and Judd case, applying the same method I obtain:

$$w(\sigma) = \frac{sN}{\sigma}, \tag{69}$$

$$\sigma = \frac{Ns}{6c}, \tag{70}$$

$$w^* = 6c, \tag{71}$$

$$p = \frac{w}{1+\sigma} = \frac{36c^2}{6c+Ns}. \tag{72}$$

Appendix D: Endogenous entry

Until now, I have taken as exogenous the number of deal finders in 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

$$\pi = \frac{1}{N} - \frac{s}{F(w)} - \alpha, \tag{73}$$

and the equilibrium number of deal finders is

$$N = \lfloor \frac{F(w)}{s - \alpha F(w)} \rfloor, \tag{74}$$

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 (in the Varian setting). Indeed, as shown in (64), 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 $[0,1]$ in the Varian setting, I find by replacing w by its value found in (64), the equilibrium number of deal finders N as

$$N = \lfloor \left(\frac{\sqrt{(-2c\alpha + 2c + s)^2 - 8c(-2c\alpha - s)} - 2c\alpha + 2c + s}{2(2c\alpha + s)} \right) \rfloor. \tag{75}$$

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.

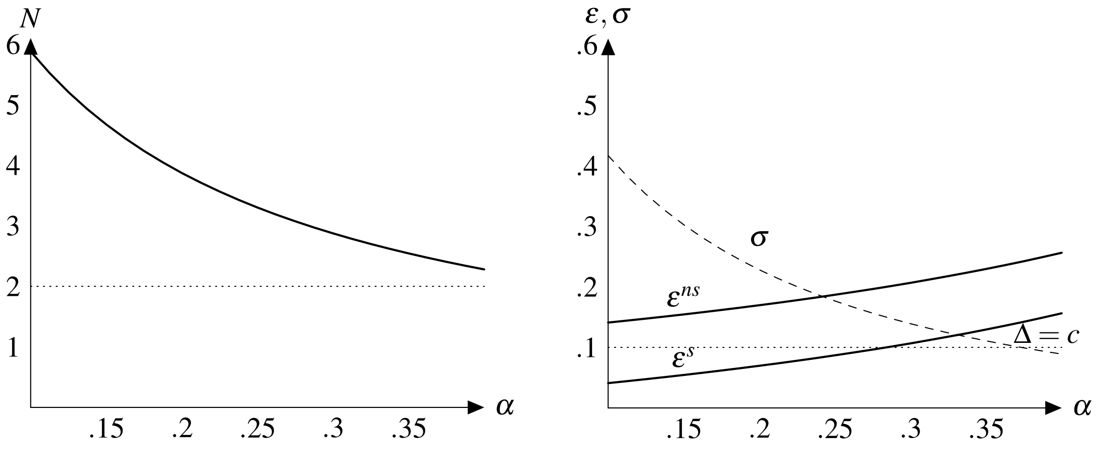

Number of deal finders in equilibrium

Equilibrium expected mismatch and share of savvy types.

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

Appendix E: Search externality at equilibrium

As consumer information creates an externality on all consumers, it is straightforward that the equilibrium share of savvy consumers is never the one that maximizes consumer aggregate welfare. Define the sum of all consumers utilities as

$$W^c = (1-\sigma)(v - p - \varepsilon^n s) + \sigma(v - \varepsilon^s - c - p). \tag{76}$$

Differentiating with respect to σ, I find

$$\frac{dW^c}{d\sigma} = \varepsilon^{ns} - \varepsilon^s - c - (1-\sigma)(\frac{d\varepsilon^{ns}}{d\sigma}) - \sigma\frac{d\varepsilon^s}{d\sigma} - \frac{dp}{d\sigma}. \tag{77}$$

As, at equilibrium, $\varepsilon^{ns} - c = \varepsilon^s$ and as $\frac{d\varepsilon^{ns}}{d\sigma} < 0$, $\frac{d\varepsilon^{ns}}{d\sigma} < 0$ and $\frac{dp}{d\sigma} < 0$ it follows that $\frac{dW^c}{d\sigma} > 0$.

The expression is pretty long to reproduce here. I display on table 2 the optimal and equilibrium values for the uniform example used throughout the paper. We see that the difference is particularly large for a small number of intermediaries.

N	equilibrium σ	consumer optimal σ
2	0.07	0.72
3	0.15	0.81
4	0.24	0.88
5	0.33	0.93
6	0.43	0.98
7	0.52	1
8	0.62	1
9	0.72	1

Table 2: Equilibrium and consumer optimal share of informed consumers