Competitive Shopbots-Mediated Markets

David Sarne, Bar-Ilan University

This paper considers markets mediated by autonomous self-interested comparison-shopping agents. As in today's markets, the agents do not charge buyers for their services but rather benefit from payments obtained from sellers upon the execution of a transaction. The agents aim at maximizing their expected benefit, taking into consideration the cost incurred by the search and competition dynamics that arise in the multi-agent setting. The paper provides a comprehensive analysis of such models, based on search theory principles. The analysis results in a characterization of the buyers' and agents' search strategies in equilibrium. The main result of the paper is that the use of self-interested comparison-shopping agents can result in a beneficial equilibrium, where both buyers and sellers benefit, in comparison to the case where buyers control the comparison-shopping agent, and the comparison-shopping agents necessarily do not lose. This, despite the fact that the service is offered for free to buyers and its cost is essentially covered by sellers. The analysis generalizes to any setting where buyers can use self-interested agents capable of effectively performing the search (e.g., evaluating opportunities) on their behalf.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Economics

Additional Key Words and Phrases: Comparison-shopping, Information brokers, Game theory, Market competition, Mechanism design

ACM Reference Format:
Sarne, D. 2012. Competitive Shopbots-Mediated Markets. ACM Trans. Embedd. Comput. Syst. , , Article (October 2012), 42 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

In an effort to fully exploit the plethora of retailers and virtual stores over the internet, buyers are increasingly adopting autonomous agents [Tan et al. 2010; Maes et al. 1999; Decker et al. 1997]. While there are many applications in which agents can be used in order to facilitate consumer-related activities over the different stages of the consumer's buying experience [Guttman et al. 1998; He et al. 2003], most emphasis in research is placed on the integration of software agents into the "Merchant brokering" stage [Maes et al. 1999; Haynes and Thompson 2008]. In this stage, the buyer searches for sellers who offer a specific desired product and learns their price. This process can be facilitated by many commercial comparison-shopping agents (CSAs) found on the Web (e.g., `PriceScan.com`, `Shopping.com`, `MySimon.com`). These agents (also called shopbots [Doorenbos et al. 1997; Fasli 2006]) are Web-based intelligent software applications that can help online shoppers find lower prices for commodities or services. The main advantage of CSAs is in their capability to automatically query multiple

---

This work was supported in part by the Israeli Science Foundation grant no. 1401/09 and BSF grant 2008-404.
Author's address: D. Sarne, Computer Science Department, Bar-Ilan University.

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

electronic sellers for price information and present it in a consolidated and compact format [Bakos 1997; Kephart et al. 2000].

This paper advances CSA research by adding two important dimensions to the classical models of CSA-based search in electronic markets. First, it considers self-interested CSAs, whose search policy is determined merely by their own net benefit. This is in contrast to the assumption commonly made in prior studies of such markets, according to which the CSAs' goal is identical to the buyer's goal and thus it operates in a way that minimizes the buyer's expense [Markopoulos and Ungar 2002; Janssen and Moraga-Gonzalez 2004]. The agents' benefit in our model is affected by payments received from sellers for purchases made by buyers that they direct to their websites and the costs incurred as part of the real-time querying process in which the agents need to engage in order to obtain the appropriate price quotes from sellers. Second, the model assumes that each CSA operates in a market where several other self-interested CSAs are available to the buyer. Thus, when constructing their search strategy, the CSAs need to consider not only the benefits in extending their search versus the costs associated with it but also the competition dynamics that arise from the influence that their search strategy has on other CSAs' strategies and vice versa.

The analysis employs "search theory" principles [McMillan and Rothschild 1994; Morgan and Manning 1985; Lippman and McCall 1976], which is a standard framework for analyzing equilibrium in markets where buyers use comparison-shopping agents [Janssen and Moraga-Gonzalez 2004; Waldeck 2008]. In particular, the model considers sequential search strategies. Under the new assumptions, a comprehensive search-based equilibrium analysis is given and the individual welfare of each of the players is extracted. These are compared to the case where CSAs are buyer-operated (or "self-operated"), i.e., the buyer controls the CSA, however needs to account for its search cost. In the latter approach, the CSA's search strategy resembles the one used in current CSA literature, i.e., attempts to minimize the overall expected expense. The results of such a comparison are important, for example, to owners of electronic marketplaces, such as `alibaba.com`, `made-in-china.com` and `gobizkorea.com`, that want to attract more buyers and sellers by improving their benefit from using the website platform through offering CSA services. Utilizing the analysis, a marketplace owner can not only decide between offering CSAs controlled by buyers and self-interested CSAs that do not charge buyers, but can also reason about the optimal (for her purposes) number of CSAs to be used and the payment that sellers will need to pay CSAs upon the completion of a transaction if the self-interested CSAs are preferred.

The analysis reveals several interesting, and somehow non-intuitive, insights concerning the effect of competing CSAs on individual buyers', CSAs' and sellers' performance. Particularly, it is shown that there is no guarantee that the beneficiaries of the transition to self-interested seller-payments-driven CSAs (in comparison to buyer-operated CSAs) are the buyers at the expense of the sellers. This, despite the buyer's direct saving of the costs incurred by the search process and the expenses of sellers in the form of payments to CSAs. The high level explanation of this phenomena is that the CSAs-based equilibrium enables buyers to save on the search cost, however dictates a less extensive search (than the one used with buyer-operated CSAs) which lets the sellers keep more of the trade surplus. This enables demonstrating a wide range of the possible consequences of having competing self-interested CSAs in a market. These range from both buyers and sellers benefiting (denoted "mutually-beneficial" equilibrium onwards) to individual benefit to one of the parties and individual loss to the other. Even when the buyers lose due to the use of self-interested CSAs, for many specific choices of the number of competing agents and the magnitude of payments that sellers make to CSAs, the increase in the sellers' expected benefit is greater than the increase in the buyer's expense. In such cases, if an appropriate side-payment mechanism can be established (e.g., sellers will compensate buyers for using the self-interested CSAs in the form of discounts if using the latter for search), then a wide range of solutions can hold in which both parties benefit in comparison to the buyer-operated search case. This kind of result cannot hold in the case of buyer-operated CSAs, as discussed later on.

The paper also addresses the inherent problem of service misuse that may occur in markets with competing self-interested CSAs [Zhu and Madnick 2010; mySimon 2000]. The model in this case is extended to enable CSAs to prevent misuse by charging a usage fee. The analysis shows what fee should be charged in equilibrium and the strategies to be used by CSAs and buyers in this case.

As thoroughly discussed in Section 7, the analysis and results can be generalized for other markets where the buyers' search involves a thorough evaluation of opportunities rather than merely obtaining a price quote. Examples for such markets are `autotrader.com` and `Yet2.com`. In these examples, the buyer needs to query sellers for supplementary information regarding the opportunities they list (used cars, inventions and patents) in order to reason about their actual value. The buyer can make use of an expert or an agent that performs the evaluation more effectively, querying sellers on her behalf. Here again the same question may arise, of whether to rely on a free seller-sponsored self-interested agent or hire a costly agent that will conduct the search according to the buyer's best interest.

2. RELATED WORK

The agent-based comparison-shopping domain has attracted the attention of researchers and market designers ever since the introduction of the first CSA (BargainFinder, [Krulwich 1996]) [Tan et al. 2010; Bakos 1997; Guttman et al. 1998; He et al. 2003]. CSAs were expected to reduce the search cost associated with obtaining price information by finding the best price or nearest point of distribution [Pathak 2010; Montgomery et al. 2004; Wan et al. 2009; Maes 1994]. As such, the majority of CSA research is mostly concerned with analyzing the influence of CSAs on retailers' and consumers' behavior [Clay et al. 2002; Johnson et al. 2004; Smith 2002; Karat et al. 2004; Yuan 2003] and the cost of obtaining information [Bakos 1997; Waldeck 2008].

In this context, much emphasis has been placed on pricing behavior in the presence of CSAs [Tan et al. 2010; Tang et al. 2010], and in particular on the resulting price dispersion [Baye et al. 2004; Clemons et al. 2002] in markets where buyers apply comparison-shopping. Substantial empirical research, mostly based on data from online books, CDs and travel markets, has given evidence to the persistence of price dispersion [Baye et al. 2006; Brynjolfsson et al. 2003; Clay et al. 2002] in such markets. Many theories were suggested in order to explain the existence of the price dispersion. For example, dynamic pricing theories suggest that sellers can benefit from frequent price adjustments of their goods, taking into account competitors' prices [Jumadinova and Dasgupta 2008; Kephart et al. 2000]. Alternatively, E-retail managers may use "hit and run" sales strategies, undertaking short-term price promotions at unpredictable intervals - a method shown to be effective and widely used [Baye et al. 2004]. Recently, it was reported in the Financial Times that sellers on Amazon's retail site are increasingly using high-speed algorithmic trading tools to automatically set prices, in a way that ensures that their prices are always below their rivals [Jopson 2012].

The investigation of CSAs' search strategies builds on "search theory", which has been a well-established research domain since the 60's [McMillan and Rothschild 1994; Morgan and Manning 1985; Lippman and McCall 1976]. Search theory investigates optimal stopping rules for searchers in a costly environment, taking into consideration the trade-off between the marginal benefit of each additional search iteration and the cost of executing it ([Hazon et al. 2009; Lippman and McCall 1976; Morgan and Manning 1985; Gal et al. 1981], and references therein). Within the framework of search theory, two main clusters of search models can be found: (a) the sequential search model and (b) the fixed sample size model (also known as "parallel" model). In the sequential search model [Rothschild 1974; Lippman and McCall 1976], the searcher obtains a single value at a time, allowing multiple search stages. In the fixed sample size model, the searcher obtains a large set of values in a single search round [Huang and Kazeykina 2010; Stigler 1961] and then chooses the best value from those obtained. This latter approach is most applicable when a time constraint limits the searcher to a single search round (e.g., when applying to colleges and universities). Many variants of these two models have been considered over the years, differing in the decision horizon (finite versus infinite) [Lippman and McCall 1976], the presence of the recall option [McMillan and Rothschild 1994] and the distribution of values. Some model variants assume findings are valid for a limited time, and with some probability may become obsolete and irrelevant for the search [Landsberger and Peled 1977]. Other variants considered multi-agent cooperative search for multiple goals [Sarne et al. 2010].

A remarkably long list of articles has been dedicated to variations of the "secretary problem" [Ferguson 1989], which is a classical optimal-stopping online problem. Yet the latter does not involve search costs and the goal is to maximize the probability of finding the best candidate rather than minimizing cost.

Despite the many advances in applying search theory for investigating search dynamics in markets where comparison-shopping principles are applied [Waldeck 2008; Janssen and Moraga-Gonzalez 2004], the absolute majority of the works assume that the CSA and user interests are identical and that the shopbot's sole purpose is to serve the buyer's needs [Markopoulos and Ungar 2002]. Other works take the buyer to be the CSA entity [Janssen et al. 2005; Stahl 1989; Varian 1980], i.e., uses the most cost-effective search strategy for minimizing the buyer's overall expense. Naturally, in such cases, the existence of CSAs improves the buyers' performance, resulting in a lower benefit to sellers [Gorman et al. 2009; Nermuth et al. 2009]. Those few works that do assume that the CSAs are self-interested autonomous entities [Kephart and Greenwald 2002; Kephart et al. 2000] focus on CSAs that charge buyers (rather than sellers as in today's markets [Wan and Peng 2010]) for their services. More importantly, these works do not take into consideration in their framework and analysis the dynamics that arise from multi-CSA competition. One exception to the above is the seminal work of Baye and Morgan [Baye and Morgan 2001] that examines the equilibrium interaction between a market for price information (controlled by a gatekeeper, which can be seen as the equivalent of CSA in some sense) and the homogenous product market it serves. One strength of their model is in taking the price distribution as well as the fees charged to buyers and sellers (subscription and advertising fees) to be endogenous, whereas in our model buyers are not charged by the CSAs, and price distribution and sellers' fees are taken to be exogenous. On the other hand, our model assumes buyers can query any seller, whereas Baye and Morgan assume buyers can directly query sellers only in local markets that are completely segmented, hence each buyer can query all sellers only if using the gatekeeper. Most importantly, Baye and Morgan assume there is only a single monopolist gatekeeper that does not encounter any competition, whereas our model focuses on the dynamics resulting from the multi-CSAs competition. In addition, in their model, if a buyer picks the gatekeeper alternative, then it sees all sellers' listings with no additional fee, hence there are no search considerations involved from the buyer's point of view. In our model, on the other hand, the analysis is entirely based on optimal search considerations, as both the querying of sellers and CSAs are associated with a cost. Finally, Baye and Morgan's model assumes that buyers are not queried proactively, but rather choose whether they want to be listed based on the fees set and the resulting competition, therefore there is no need for the gatekeeper to consider search-related tradeoffs. In our model, on the other hand, CSAs' primary decision is whether to continue querying sellers (i.e., resume search) based on the quotes received so far.

3. THE MODEL

We first formally present the model and then give the appropriate justifications for the assumptions made.

3.1. Model Assumptions

The model considers an electronic marketplace populated by numerous sellers and buyers and N autonomous CSAs.[1]. Buyers are assumed to be self-interested and fully rational, aiming to minimize their overall expense when purchasing a single item of a well-defined product. They are assumed to be able to request service from any subset of CSAs from those available in the market, incurring a cost $c_{buyer \rightarrow csa}$ for each queried CSA. It is assumed that buyers can learn about an individual CSA's search strategy (if at all) only after incurring the cost $c_{buyer \rightarrow csa}$. In addition to querying CSAs, buyers can query sellers directly, incurring a cost $c_{buyer \rightarrow seller} \geq c_{buyer \rightarrow csa}$. Buyers are assumed to always prefer the minimum quote among those received from the CSAs they use.

CSAs are assumed to be self-interested and fully rational, aiming to maximize their own net benefit. CSAs do not charge buyers for their service, however sellers have to pay a fixed fee, denoted M, every time a buyer that was referred to their website by the CSA executes a transaction. CSAs are assumed to be using real-time querying, upon being requested by buyers for price comparison services. The CSAs' querying process is executed sequentially (i.e., a CSA queries sellers one at a time and eventually returns a price quote), and when querying a seller, a CSA incurs a fixed cost common to all agents, denoted $c_{csa \rightarrow seller}$. It is assumed that CSAs can obtain as many price quotes as they request, and the probability that two CSAs query the same seller at the exact same time, receiving the exact same quote, is negligible. The duration of querying a seller (i.e., the delay) is assumed to be negligible, in the sense that the resulting delay in the response provided to the buyer is tolerable. Finally, the model assumes that buyers can be identified and any CSA can choose whether to re-service a buyer asking for price comparison for the same product within a given interval of time or not.

Sellers are not assumed to be strategic and their pricing as well as their payment M to CSAs are assumed to be external, unaffected by the CSAs-based search process. Consequently, any new request for a price quote is assumed to yield a different price, from a probability distribution function $f(q)$, which is assumed to be known to the CSAs and buyers.

For comparison purposes, we consider the model of buyer-operated CSAs, which represents the case of having the buyer search with a CSA of its own (or searching by itself with a similar price querying cost $c_{csa \rightarrow seller}$). This model assumes that whenever a buyer agent requests the services of a CSA it gains full control over that agent, however needs to account for its search costs. The CSA in this case, therefore, uses the overall-expense-minimizing search strategy. Furthermore, it is assumed that each buyer interested in buying the product values it more than its total expense when querying sellers directly (i.e., with a cost $c_{buyer \rightarrow seller}$), therefore a buyer always buys the product in the end.

The main question addressed in this paper is what are the CSAs' and buyers' equilibrium search strategies, given the environment parameters (N, M, $c_{csa \rightarrow seller}$, $c_{buyer \rightarrow seller}$, $c_{buyer \rightarrow csa}$ and $f(q)$). A buyer's search strategy defines at each stage of the

[1]See the table in the appendix, containing the complete set of notations used throughout the paper.

search, based on the CSAs that have already been queried and the best price quote received so far, whether to resume the search and, if so, whether a seller or a CSA should be queried next. Based on the best price quote received so far, a CSA's strategy defines, at each stage of the search, whether or not an additional seller should be queried. The analysis is based on the characterization of each CSA's strategy, given the strategies used by the other CSAs and by the buyers. In addition, we are interested in the effect of the payment M, the number of CSAs, N, and the search costs $c_{csa \rightarrow seller}$, $c_{buyer \rightarrow seller}$ and $c_{buyer \rightarrow csa}$ over the buyer's expected expense and the sellers' and CSAs' expected benefits in equilibrium. Since the whole process scales up linearly in the number of times a buyer initiates a multi-CSA based search, we can simply consider the interactions involving a single buyer's search.

3.2. Assumptions' Justification

The following paragraphs provide justifications and discuss in greater detail the different model assumptions. The cost that buyers incur whenever requesting a CSA's service or querying a seller for price can be considered the cost of the time it takes to get to the appropriate website, specify the product of interest, as well as any other required complementary information, and waiting for the results. The cost incurred when querying a seller is generally greater than the cost incurred when querying a CSA, because CSAs' URLs are typically more accessible and their interface for specifying the requested product or service is typically more structured and easy to use (than in the sellers' website). The assumption that buyers can learn about an individual CSA's search strategy only after incurring the appropriate cost is justified mainly by the fact that a substantial portion of the cost of querying the CSA is attributed to the time it takes the buyer to get to the CSA's website and specify the product in which she is interested. (Only then can she learn the strategy that this CSA will be using (if the CSA is willing to disclose it), as each product is associated with a different distribution of prices.) The assumption regarding buyers' preference of the minimum quote among those received from the CSAs they contact is based on empirical findings showing that buyers who use CSAs are extremely sensitive to price [Tan et al. 2010; Wan and Peng 2010]. Furthermore, one can assume that the CSA will return a quote only from sellers complying with the buyer's requirements, resulting in a slightly greater search cost.

Having CSAs receive a payment from sellers rather than charging the buyer for their services is a common practice with today's CSAs (e.g., `PriceGrabber.com`, and `Shopping.com`) [Wan and Peng 2010; Moraga-Gonzalez and Wildenbeest 2011]. This fee usually depends on product category [Moraga-Gonzalez and Wildenbeest 2011] and is therefore not used as a decision variable in the model. The idea that a platform will charge only one side in two-sided markets while the other group is allowed onto the platform for free can be explained, in general, by intense competition for players of that group (e.g., in the case of yellow pages directories that are supplied to readers for free) [Armstrong 2006].[2]

While until recently CSAs' architectures were based on maintaining a database of prices that is queried upon request [Fasli 2006; Wan and Peng 2010], newly developed CSAs (e.g., Apnoti (`smart.apnoti.com`), InvisibleHand (`getinvisiblehand.com`), Kelkoo (`kelkoo.com`) and Kayak (`kayak.com`)) are based on querying the electronic merchants directly, in real-time, upon the arrival of price comparison requests from the users. The reasons for preferring real-time querying (also in CSA-based search

[2]In some cases, consumers are in effect paid to use the platform (e.g., in the case of credit cards, where rewards programs such as contributions to frequent flyer plans are offered); the low price on one side not only attracts elastic consumers but also, as a result, leads to higher prices or more participation on the other side [Rysman 2009].

research [Baye et al. 2004; Kephart et al. 2000]) are the constantly increasing rate of price changes [Montgomery et al. 2004; Smith 2002; Wan and Peng 2010] as well as the need to retrieve both product and service real-time availability information (as with Kelkoo and Kayak). The increase in price change frequency results from having the ability to better react to external factors that affect price and changes in competitors' pricing (i.e., dynamic pricing [Greenwald et al. 1999; Kephart and Greenwald 2002]).[3] Taking the airline industry as an example, airfares are often adjusted by the airline based on competitors' pricing, as well as on how many empty seats are left on the plane, the time left until the flight, fuel prices, cancellations made by people who have already made reservations, etc.[4] A reliable CSA is therefore expected to query electronic merchants (rather than retrieve price data from a formerly collected price database) upon the arrival of price comparison requests from its users [Montgomery et al. 2004; Smith 2002; Wan and Peng 2010].

The real-time querying process is well recognized to be costly [Bakos 1997; Choi and Liu 2000; Kephart et al. 2000], in the sense that the CSA needs to invest/consume some of its resources (CPU time, communication bandwidth, etc.) in opening a connection with the remote server, extracting and filtering the relevant information and comparing it with the other results obtained. These can either be direct costs associated with the search or the alternative gain that could have been obtained if the resources required for querying would have been allocated for the sake of other incoming requests [Sarne et al. 2007; Markopoulos and Ungar 2002]. The last decade's advances in communication technologies indeed reduce search costs and other environmental inefficiencies in multi-agent environments [Biswas and Narahari 2001], however these still need to be considered when investigating an agent's search strategy [Huang and Kazeykina 2010; Choi and Liu 2000]. The cost of querying a seller is assumed to be common to all of the CSAs operating in the market because real-time querying is based on a standard protocol, and CSAs typically use similar hardware infrastructure. The ability of the CSAs to obtain as many price quotes as they request with a negligible probability of receiving the same quote is justified by the frequent changes in prices and the large amount of sellers in the market. The use of sequential search by the CSAs is standard [Janssen and Moraga-Gonzalez 2004; Waldeck 2008]. Alternatively the CSA could have used parallel search, though the analysis of this possibility was left beyond the scope of the current paper. In general, sequential search outperforms parallel search in cases where the per-query cost is fixed [McMillan and Rothschild 1994; Morgan and Manning 1985]. It is dominated by parallel search as far as the incurred delay is concerned, however given the phenomenal increase in communication speed nowadays, this aspect becomes a non-issue in our case, and the resulting delay is tolerable even if the CSA queries many sellers sequentially before responding.

CSAs can identify buyers in various ways (e.g., by IP address or a user account). In case a CSA decides not to re-service a buyer asking for price comparison for the same product within a given interval of time, it does not necessarily need to decline the request directly — it can also re-display the results it usually saves in case these are required again, e.g., in case of refreshing the page.

[3]Sellers now use dynamic pricing techniques to set their prices, attempting to benefit from re-pricing their goods as often as possible (in most cases using software agents called pricebots [Jumadinova and Dasgupta 2008]) based on their observations of the competitors' prices [Kephart et al. 2000; Jopson 2012].

[4]Ten years ago, Etzioni et al. found that the price of tickets on a particular flight can change as often as seven times in a single day [Etzioni et al. 2003]. While more current official data is unavailable, it is likely that this rate has increased substantially over the past decade (and is likely to keep increasing) due to the blooming of computerized pricing tools such as IntelliPricer (`www.vayant.com`) that update fares 24/7. Similarly, sellers on Amazon's retail site were shown to be setting prices as often as every 15 minutes [Jopson 2012].

Finally, the assumptions regarding the existence of a price distribution and the buyer's familiarity with this distribution are widely used in CSA research [Tang et al. 2010; Waldeck 2008; Janssen et al. 2005]. The first is usually supported by empirical research in well-established online retail [Baye et al. 2004; 2006; Brynjolfsson et al. 2003; Clay et al. 2002], though these studies are somehow limited by the frequency over which price changes are recorded. The second is based on the ability of agents to estimate the characteristics of price uncertainty (e.g., using past experience and Bayesian update).

4. ANALYSIS

We first introduce the principles of optimal search as they appear in legacy search theory. We then show that the buyer-operated-CSA case maps to the sequential search model that is widely used in this literature [McMillan and Rothschild 1994]. The analysis of the settings with competing self-interested CSAs follows. We first analyze the individual buyer's expected-expense-minimizing strategy (taking the strategy of CSAs as given) and the individual CSA's expected-benefit-maximizing strategy (taking the strategy of buyers and other CSAs as given). Based on the individual strategy analysis we introduce an equilibrium analysis. The equilibrium analysis relies on pure equilibrium concepts, where none of the players have an incentive to deviate from its strategy, given the strategies of the other players. Finally, we discuss the implications of changes in some of the model assumptions over the analysis of the model and the resulting equilibrium. These include the case where there is only a single self-interested CSA available in the market ($N = 1$), the case where buyers can query the same CSA more than once (re-querying) and the case where the buyer can use a buyer-operated CSA alongside the self-interested CSAs.

4.1. Expected-Benefit-Maximizing and Expected-Expense-Minimizing Search

The expected-benefit-maximizing and expected-expense-minimizing strategies of a searcher engaged in a costly search can be found in classical search theory literature [Weitzman 1979; Rothschild 1974; Lippman and McCall 1976]. The problem, in its most general form, considers a searcher facing a set of opportunities, each associated with some value (representing either benefit or expense) for the searcher. The searcher's uncertain knowledge about each opportunity is described by a probability distribution. The process of obtaining the actual value of any opportunity incurs a cost, which may vary between opportunities.

The optimal (expected-benefit-maximizing or expected-expense-minimizing) search strategy is inherently sequential (i.e., exploring one opportunity at a time) [Weitzman 1979]. Its expected-benefit-maximizing variant is based on setting a reservation value (a threshold) to each opportunity according to the distribution characterizing its value and the cost of revealing that value. The searcher should always choose to obtain the value of the opportunity associated with the maximum reservation value and terminate the search once the maximum value obtained so far is greater than the maximum reservation value of any of the opportunities which have not yet been explored. The reservation value R_i of an opportunity associated with a cost c_i and a distribution $f_i(x)$ can be calculated based on the following equation:

$$c_i = \int_{y=R_i}^{\infty} (y - R_i) f_i(y) dy \qquad (1)$$

Intuitively, R_i is the value where the searcher is precisely indifferent: the expected marginal benefit from revealing the actual value of an opportunity (right-hand-side) exactly equals the cost of doing so (left-hand-side).

If the opportunities available to the searcher are homogeneous, i.e., their values derive from a common distribution $f(y)$ and the cost of obtaining their actual value is similar ($c_i = c$), then they all share the same reservation value R [McMillan and Rothschild 1994]. In this case, the searcher should continue the search as long as there are opportunities that have not yet been explored and the highest value obtained so far is below the reservation value R. Since the same reservation value is assigned to all opportunities, the searcher should randomly pick an opportunity to be explored each time, among those which value is still unknown. In particular, if the searcher is not limited by the number of opportunities she can explore, then the search will necessarily terminate only once it obtains a value above the reservation value R. The expected benefit to the searcher in this case, denoted $V(R)$, is given by [Rothschild 1974]:

$$V(R) = \frac{-c + \int_{y=R}^{\infty} yf(y)dy}{1 - F(R)} \tag{2}$$

where $F(x)$ is the cumulative distribution function from which values are drawn. The expected number of search iterations is simply the inverse of the success probability, $\frac{1}{1-F(R)}$, since this becomes a Bernoulli sampling process as opportunities arise independently at each iteration. Therefore, Equation 2 can be broken down into the accumulated cost throughout the search, $\frac{-c}{1-F(R)}$, and the expected value obtained eventually (greater than R), $\frac{\int_{y=R}^{\infty} yf(y)dy}{1-F(R)}$.

Similarly, for the case of expected-expense-minimization (e.g., when searching for a product or a service and values of opportunities represent prices or charges), the rule is to continue the search as long as the best value obtained so far is above the lowest reservation value among those associated with opportunities that have not yet been explored. The reservation value R_i of an opportunity associated with a cost c_i and a distribution $f_i(x)$ can be calculated based on the following modification of (1):

$$c_i = \int_{y=0}^{R_i} (R_i - y) f_i(y)dy \tag{3}$$

The expected expense of a searcher, when not limited by the number of opportunities she can explore and all opportunities are homogeneous, is given by:

$$V(R) = \frac{c + \int_{y=0}^{R} yf(y)dy}{F(R)} \tag{4}$$

The expected number of search iterations in this case is $\frac{1}{F(R)}$.

The reservation value R, calculated both according to (1) and (3), for the case where the searcher is not limited by the number of opportunities she can explore and all opportunities are homogeneous, satisfies:

$$V(R) = R \tag{5}$$

This equality can be intuitively explained by interpreting R as the value for which the searcher is indifferent between resuming the search (obtaining expectancy $V(R)$) and terminating the search (obtaining R) [Rothschild 1974].

4.2. Buyer-Operated CSA

A CSA owned/operated by the buyer will aim to minimize the buyer's expected expense and therefore will set a reservation value (reservation price) $r_{csa \rightarrow seller} = R$ according to (3). It will query sellers as long as the quote obtained is above $r_{csa \rightarrow seller}$. The buyer's expected expense in this case is: $c_{buyer \to csa} + \frac{c + \int_{y=0}^{r_{csa \to seller}} y f(y) dy}{F(r_{csa \to seller})}$. The first component, $c_{buyer \to csa}$, is the cost of querying the CSA. The second is the expected expense, based on (4), which equals $r_{csa \to seller}$ (according to (5)). This latter component can be further decomposed to the sum of the costs that the CSA incurs along the search (for which the buyer accounts), $\frac{c}{F(r_{csa \to seller})}$, and the expected quote resulting from the search, $\frac{\int_{y=0}^{r_{csa \to seller}} y f(y) dy}{F(r_{csa \to seller})}$, as explained above when discussing the expected-benefit-maximizing case. The expected quote that the buyer ends up with in this case is also the expected benefit of the seller from whom the buyer eventually buys the product.

For completeness' sake, we note that because the use of the buyer-operated CSA incurs a cost $c_{buyer \to csa}$, it is theoretically possible that the expected expense using the buyer-operated CSA is greater than the expected expense resulting from querying sellers directly. The expected expense of a buyer querying sellers directly equals, according to (5), to $r_{buyer \to seller}$, where $r_{buyer \to seller}$ is the solution R_i to (3), using the cost $c_i = c_{buyer \to seller}$. Therefore the buyer's expected expense in both forms of search can be calculated and the buyer should choose the one resulting with the minimum expected expense.

4.3. Individual Strategies and Equilibrium with Self-Interested CSAs

We begin the analysis of the case where self-interested CSAs are used with a simple numerical example that captures the essence of the mutually-beneficial equilibrium. Consider a setting with a buyer-operated CSA, where the values that the different model parameters obtain are: $c_{buyer \to seller} = 0.003$, $c_{buyer \to csa} = 0.002$, $c_{csa \to seller} = 0.001$ and the distribution of prices is the standard uniform distribution ($f(q) = 1$ for $0 \le q \le 1$ and 0 otherwise). Based on the analysis of the buyer-operated CSA case above, the CSA would query sellers using a reservation value $r_{csa \to seller} = R = 0.045$ (according to (3)). The expected number of sellers that the CSA will be querying in this case is $1/F(0.045) = 22.2$. The buyer's expected expense in this case is 0.047 (calculated as the average price she ends up with ($0.045/2$) plus the expected cost of the CSA's search ($22.2 * 0.001$) and the cost $c_{buyer \to csa}$). The seller's expected benefit is the expected price at which the buyer ends up buying the product: $0.045/2 = 0.022$.

Now consider an alternative setting, similar to the one introduced above, except that a single self-interested CSA is used and the payment it obtains if the buyer purchases, based on a quote it supplied, is $M = 0.014$.[5] Since $c_{buyer \to csa} < c_{buyer \to seller}$, the buyer will prefer querying the CSA first (as it costs less than querying a seller directly, and the quote obtained is from at least as good a distribution). After querying the CSA, the buyer will query sellers directly, as long as the quote obtained from the CSA is greater than the reservation value it assigns sellers. This latter reservation value can be calculated as the solution R_i to (3), using the cost $c_i = c_{buyer \to seller} = 0.003$, resulting in $r = 0.077$. Now consider the self-interested CSA. If returning a price quote above 0.077, the buyer will keep on querying sellers directly until she obtains a price quote below 0.077. Therefore there is no point in returning a quote above 0.077, hence the CSA should set its own reservation value to $r_{csa \to seller} = 0.077$, yielding an expected average quote of $0.077/2 = 0.039$. The expected number of sellers that the CSA will be querying in this case is $1/F(0.077) = 13$, hence its expected benefit is: $M - c_{csa \to seller} * 13 = 0.014 - 0.001 * 13 = 0.001$. Since the buyer does not need to account for the search cost in this case, her expected expense is 0.041, calculated as the sum of the price quote with which she ends up and the cost $c_{buyer \to csa}$. This is less than the 0.047 she would

[5]The case of having a single self-interested CSA is formally analyzed in detail in 4.4. Here we merely solve a degenerate instance.

have spent in the buyer-operated model. The seller, on the other hand, will be selling the product at an average price of 0.039, however will need to pay the CSA 0.014, thus her expected benefit is 0.025, which is more than the 0.022 she would have benefited in the buyer-operated model. In this case, the sellers' expected benefit increases, despite the commission they pay, and at the same time buyers spend less and even the CSA profits from operating in the market. The explanation for this situation is that by offering the commission to the CSA, the sellers partially subsidize the costs associated with search. If this subsidy could have been transferred completely to the buyers, they would improve their performance and the sellers would worsen theirs. Nevertheless, the self-interested-CSA scenario dictates a search pattern which is less efficient than a single agent's expense-minimizing sequential search, as in the buyer-operated case, which makes the overall search process less efficient. In our example, only 13 sellers are queried in the self-interested CSA case (as compared to 22.5 in the buyer-operated case), resulting in an increase in the expected price paid (0.039 compared to 0.022). However, despite the increase in the payment to the seller, the buyer does not need to account for the search costs, thus her expected overall expense decreases (i.e., despite the inefficiencies of the search, the buyers benefit from having CSAs perform part of the search for them for free). Similarly, the seller, despite paying to the CSA, benefits from the increase in the sale price, resulting from the decrease in the search extent. Overall, less is spent on search (in comparison to the buyer-operated case) and the savings are divided between the buyer, seller and CSA (without the need for external interference).

We now turn to analyze the search strategy of buyers and CSAs when several self-interested CSAs are available and receive a payment M from sellers, conditional on the buyer's purchase. The analysis of the competing CSAs case enables the demonstration of several interesting equilibrium characteristics that do not hold in the single CSA case. Furthermore, as illustrated in the next section, the CSAs competition contributes to substantially extending the range of settings in which a mutually-beneficial equilibrium exists. We first extract the individual CSA's expected-benefit-maximizing strategy, given the probability that any possible given quote will turn out to be the most attractive one along the buyer's search. Then we prove a set of properties of the equilibrium, among which is a proof that the buyer never gets to query sellers directly, that the equilibrium is symmetric and proofs relating to the relationships between the reservation values used by the different players. Based on these properties, we extract the set of equations from which the equilibrium derives and formally expresses the buyers' expected expense and the sellers' and CSAs' expected benefits.

From the CSA's point of view, every seller is an opportunity offering an expected reward $MP(q)$, where q is the quote returned by the seller if queried and $P(q)$ is the probability that the buyer will eventually buy the product at price q. Given that the buyer can query other CSAs or resume the search by querying sellers directly, it is possible that the buyer will find a quote lower than q. Therefore $P(q)$ is the probability that none of the other CSAs queried nor the sellers queried directly by the buyer returned a quote below q. Since the buyer is fully rational and self-interested, $P(q)$ necessarily increases as q decreases, and therefore so does $MP(q)$. The CSA's expected-revenue-maximizing strategy is thus sequential and follows a reservation-value rule (which relates to the quotes obtained): the search terminates only if a quote lower than the reservation value is obtained. This is because of the inverse relation between q and $MP(q)$. The expected-revenue-maximizing reservation value $r_{csa\rightarrow seller}$ in this case will differ in its value from the one calculated using (3), as its calculation needs to take into consideration the search strategy of buyers and competing CSAs. Based on

(1), the value of $r_{csa\rightarrow seller}$ can be derived from:

$$c_{csa\rightarrow seller} = M\int_{q=0}^{r_{csa\rightarrow seller}} (P(q) - P(r_{csa\rightarrow seller}))f(q)dq \tag{6}$$

This mapping from a benefit-based reservation value to a quote-based reservation value facilitates the remaining analysis of the model.

From the buyer's point of view, at any stage of her search she can either: (a) terminate the search and use the lowest quote obtained so far; (b) query a CSA that has not yet been queried (or is willing to supply service again), incurring a cost $c_{buyer\rightarrow csa}$ and receiving in return a quote below the reservation value used by that CSA; or (c) independently query a seller, incurring a cost $c_{buyer\rightarrow seller}$ and receiving in return a quote according to the distribution $f(q)$. Proposition 4.1 states that as long as there are CSAs that can be queried, the buyer prefers querying CSAs over querying sellers directly.

PROPOSITION 4.1. *The buyer will never query sellers directly unless there are no CSAs available that can be queried.*

PROOF. Regardless of the number of sellers that the CSA queries along its search, it will always prefer to return the minimum quote obtained in order to maximize $MP(q)$. Therefore, from the buyer's point of view, querying a CSA dominates querying a seller directly: the cost of querying the CSA is at most the cost of querying the seller ($c_{buyer\rightarrow csa} \leq c_{buyer\rightarrow seller}$), and the quote returned is the minimum of a sample from a distribution $f(x)$ (in the worst case the CSA will use an infinite reservation value, i.e., will sample only one seller), whereas the quote returned from a seller contacted directly is a single observation taken from $f(x)$. □

Based on Proposition 4.1 we can now prove that all CSAs use the same reservation value in equilibrium (symmetric equilibrium) and allow buyers to query them only once. This is enabled due to the CSAs' ability to identify buyers and decide whether they want to re-service them or not.

PROPOSITION 4.2. *If a pure equilibrium exists in which the CSAs are indeed used, then all CSAs allow buyers to query them only once and use the same reservation value* $r_{csa\rightarrow seller}$.

PROOF. Assume otherwise, i.e., there exists at least one CSA that allows for being queried more than once by the same buyer for the same product, hence using a sequence of reservation values (where the ith reservation value is the one to be used the ith time it is queried by the buyer). In a such case, the CSA will necessarily use a sequence where the ith reservation value is smaller than the $(i+1)$th reservation value, for any i. This is because otherwise, when reaching the ith round of use, the CSA would already have a quote lesser than the current reservation value (from a prior search round), thus according to the reservation-value rule there is no point in any further search. We therefore now show that there cannot be a CSA that uses a reservation value greater than the one it offers in a consecutive search round, or two CSAs that use different reservation values (in any of their search rounds) in equilibrium.

Assume otherwise, i.e., there exists a CSA that, in any of its search sequences, uses a reservation value different from any of the reservation values used by the other CSAs in any of their search sequences or from the reservation value that the agent itself is using in future sequences. In that case, consider the reservation value $r_{csa_\alpha\rightarrow seller}$, which is the highest one used by any of the CSAs in any search sequence, and $r_{csa_\beta\rightarrow seller}$, which is the next highest reservation value used by any of the CSAs along their operation (i.e., $r_{csa_\beta \to seller} < r_{csa_\alpha \to seller}$). We use CSA_α and CSA_β to denote the CSAs using $r_{csa_\alpha \to seller}$ and $r_{csa_\beta \to seller}$, respectively. Notice that it is possible that CSA_α and CSA_β are in fact the same CSA, using different reservation values along its operation. We first show that if CSA_α returns a quote greater than $r_{csa_\beta \to seller}$, then this quote is never used by the buyer. If indeed CSA_α returns a quote greater than $r_{csa_\beta \to seller}$, then either: (a) CSA_β has already been queried, yielding a quote lower than the one CSA_α returned, in which case CSA_α's quote is never used; or (b) CSA_β has not yet been queried. The case (a) can hold only if CSA_α and CSA_β are different CSAs, since it has already been established that if it is the same CSA then the reservation value used cannot increase between consecutive searches. In case of (b), the buyer will necessarily query additional CSAs (or the same CSA if possible). This is because once CSA_α returns a quote greater than $r_{csa_\beta \to seller}$, the buyer is able to distinguish it as a CSA using $r_{csa_\alpha \to seller}$. Therefore, after querying that CSA, the reservation value used by another (or the same) CSA that will be queried is from a more favorable set of reservation values (the same as before querying CSA_α, however excluding one instance of $r_{csa_\alpha \to seller}$). Therefore, if the buyer found it beneficial to query a CSA before querying CSA_α then so should her decision be after realizing that the queried CSA was using $r_{csa_\alpha \to seller}$. The buyer will thus continue querying CSAs, resulting in either a quote lower than $r_{csa_\beta \to seller}$ (in which case it will not purchase the product based on the quote returned by CSA_α) or a quote greater than $r_{csa_\beta \to seller}$ (if other CSAs are using $r_{csa_\alpha \to seller}$), in which case she will continue her search for the same considerations, and so on. Therefore, inevitably, the buyer will receive a quote lower than $r_{csa_\beta \to seller}$, either from one of the other CSAs using $r_{csa_\alpha \to seller}$ (if any remained) or from CSA_β. Hence, if CSA_α returns a quote greater than r_{csa_β}, then this quote is never used by the buyer.

Based on the above, we show that the expected-benefit-maximizing reservation value of CSA_α cannot be $r_{csa_\alpha \to seller} > r_{csa_\beta \to seller}$. If $r_{csa_\alpha \to seller}$ is indeed the expected-benefit-maximizing reservation value of CSA_α, then if the best quote obtained by that CSA so far is $r_{csa_\alpha \to seller} + \epsilon$ it is beneficial to it to query an additional seller, i.e.,

$$c_{csa \to seller} < M \int_{q=0}^{r_{csa_\alpha \to seller}+\epsilon} (P(q) - P(r_{csa_\alpha \to seller} + \epsilon)) f(q) dq \tag{7}$$

Since according to the first part of the proof $P(q) = 0$ for any $q > r_{csa_\beta \to seller}$, then the right-hand side of (7) receives the same value if replacing $r_{csa_\alpha \to seller} + \epsilon$ by any other value greater than $r_{csa_\beta \to seller}$. Therefore, CSA_α benefits from querying an additional seller if its best value found is above $r_{csa_\beta \to seller}$, contradicting the assumption that its reservation value is $r_{csa_\alpha \to seller}$. Finally, since CSAs will always use the same reservation value, a sequence of decreasing reservation values used by the same CSA is ruled out, and therefore no CSA will be willing to search again after already been queried by the buyer. □

Using Proposition 4.2, we can now complete the analysis of the expense-minimizing buyers' strategy and the resulting equilibrium. Since all CSAs use the same reservation value $r_{csa \to seller}$, the quote they return is associated with the same probability distribution function, denoted $f_{returned}(q)$, and its corresponding cumulative distribution function, denoted $F_{returned}(q)$. These are given by:

$$f_{returned}(q) = \begin{cases} \frac{f(q)}{F(r_{csa \to seller})} & 0 \le q \le r_{csa \to seller} \\ 0 & otherwise \end{cases} \tag{8}$$

$$F_{returned}(q) = \begin{cases} \frac{F(q)}{F(r_{csa\to seller})} & 0 \le q \le r_{csa\to seller} \\ 1 & q > r_{csa\to seller} \\ 0 & otherwise \end{cases} \quad (9)$$

The above calculation uses Bayes' theorem for extracting the conditional distribution of receiving a quote q, given that the value returned by the CSA is below $r_{csa\to seller}$.

Based on the principles of expected-expense-minimizing sequential search that were described at the beginning of the section (and in particular Equation 3), the buyer will calculate her reservation values $r_{buyer\to seller}$ and $r_{buyer\to csa}$ for querying sellers and CSAs, respectively, extracted from:

$$c_{buyer\to seller} = \int_{q=0}^{r_{buyer\to seller}} (r_{buyer\to seller} - q) f(q) dq \quad (10)$$

$$c_{buyer\to csa} = \int_{q=0}^{r_{buyer\to csa}} (r_{buyer\to csa} - q) f_{returned}(q) dq \quad (11)$$

The buyer will query CSAs sequentially, in random order, as long as the quote returned is above the reservation value $r_{buyer\to csa}$ or until all CSAs are queried. Then, if the best quote received so far is above the reservation value $r_{buyer\to seller}$, the buyer will continue querying sellers directly, until obtaining a quote below $r_{buyer\to seller}$. The equilibrium is thus characterized by the tuple $(r_{csa\to seller}, r_{buyer\to csa}, r_{buyer\to seller})$.

Since $F_{returned}(q) \ge F(q)$ for any q, and $c_{buyer\to csa} \le c_{buyer\to seller}$, then according to (10) and (11) the following must hold: $r_{buyer\to csa} \le r_{buyer\to seller}$ (this also results implicitly from Proposition 4.1). We now use Proposition 4.3 to prove that, in equilibrium, the reservation values used by the buyer and the CSA must satisfy: $r_{csa\to seller} \ge r_{buyer\to csa}$ and $r_{csa\to seller} \le r_{buyer\to seller}$.

PROPOSITION 4.3. *The only equilibrium that can exist is one where* $r_{buyer\to csa} \le r_{csa\to seller} \le r_{buyer\to seller}$.

PROOF. We first show that the case where $r_{csa\to seller} < r_{buyer\to csa}$ cannot hold in equilibrium. In this case, the quote returned by the first CSA queried is necessarily lower than $r_{buyer\to csa}$ and therefore the buyer's search is terminated after querying a single CSA. However, in this case, any individual CSA has an incentive to deviate to using $r_{csa\to seller} = r_{buyer\to csa}$. This will decrease its search intensity (i.e., require less search on average, thus incurring lesser costs), however will still guarantee obtaining M, since the returned quote is below $r_{buyer\to csa}$.

Next, we prove that $r_{csa\to seller} > r_{buyer\to seller}$ cannot hold in equilibrium. The proof follows the methodology used in the last part of the proof given to Proposition 4.2. If $r_{csa\to seller}$, where $r_{csa\to seller} > r_{buyer\to seller}$, is indeed the expected-benefit maximizing reservation value of the CSA then if the best quote received by that CSA so far is $r_{csa\to seller} + \epsilon$ then it is beneficial to it to query an additional seller, i.e., $c_{csa\to seller} < M \int_{q=0}^{r_{csa\to seller}+\epsilon} (P(q) - P(r_{csa\to seller} + \epsilon)) f(q) dq$. However, the buyer will necessarily end up with a quote lower than $r_{buyer\to seller}$, either from one of the other CSAs or from sellers queried individually, as the latter will continue to be queried as long as the best quote is above $r_{buyer\to seller}$. Therefore, $P(q) = 0$ for any $q > r_{buyer\to seller}$, and consequently $M \int_{q=0}^{r_{csa\to seller}+\epsilon} (P(q) - P(r_{csa\to seller} + \epsilon)) f(q) dq$ receives the same value if replacing $r_{csa\to seller}+\epsilon$ by any other value greater than $r_{buyer\to seller}$. Therefore, the CSA benefits from querying an additional seller if its best value found is above $r_{buyer\to seller}$, contradicting the assumption that its reservation value is $r_{csa\to seller}$. □

The immediate implication of $r_{csa\rightarrow seller} \leq r_{buyer\rightarrow seller}$ according to Proposition 4.3 is that, despite having the option to resume her search by querying sellers directly, the buyer never gets to take advantage of this option in equilibrium — the CSAs use a reservation value lower than $r_{buyer\rightarrow seller}$ and therefore the value they return always makes individual seller's querying non-beneficial. Despite not querying sellers directly in equilibrium, the existence of the option to query sellers directly can affect the results — in some cases the existence of an equilibrium is precluded due to the existence of a low $r_{buyer\rightarrow seller}$ value, whereas in the absence of the option to query sellers directly or if the cost of querying sellers directly is relatively high (and so $r_{buyer\rightarrow seller}$'s value is high), an equilibrium solution with self-interested CSAs does exist. In fact, in the next section we demonstrate a case where the increase in the cost $c_{buyer\rightarrow seller}$ (which intuitively seems to put the buyer in a less favorable situation) enables the existence of an equilibrium which improves the buyer's expense in comparison to the buyer-operated-CSA case, which could not have existed with a low $c_{buyer\rightarrow seller}$ cost.

The characterization of the equilibrium structure, according to Proposition 4.3, enables expressing explicitly the probability that a CSA receives a payment M if returning a quote q, $P(q)$:

$$P(q) = \begin{cases} 1 & q \leq r_{buyer\rightarrow csa} \\ \sum_{i=1}^{N} P_{queried}(i)(1 - F_{returned}(q))^{N-i} & \\ \quad \cdot(1 - F_{returned}(q|q > r_{buyer\rightarrow csa}))^{i-1} & r_{buyer\rightarrow csa} < q \leq r_{buyer\rightarrow seller} \\ 0 & q > r_{buyer\rightarrow seller} \end{cases} \tag{12}$$

where $P_{queried}(i)$ is the probability that a queried CSA is the ith queried CSA in the buyer's search plan. The case of $q \leq r_{buyer\rightarrow csa}$ is trivial as, according to the reservation value rule, once the buyer receives a price lower than $r_{buyer\rightarrow csa}$ (thus necessarily also lower than $r_{buyer\rightarrow seller}$), she terminates the search and purchases the product. In the case where $q > r_{buyer\rightarrow seller}$, a lower quote will necessarily be received by one of the other CSAs or sellers, as the latter will be queried until a quote below $r_{buyer\rightarrow seller}$ is obtained. It is notable, however, that the case where $q > r_{buyer\rightarrow seller}$ will never be reached in equilibrium, since according to Proposition 4.3 $r_{csa\rightarrow seller} \leq r_{buyer\rightarrow seller}$. In the case where $r_{buyer\rightarrow csa} < q \leq r_{buyer\rightarrow seller}$, all other CSAs are necessarily queried and therefore, in order for the buyer to pick the quote q, all of the other $N-1$ CSAs need to return a quote higher than q . The probability that the queried CSA is the ith queried CSA in the buyer's search, $P_{queried}(i)$, is calculated using Bayes' theorem: the a priori probability of being the ith CSA planned for a random sequence is $1/N$. The probability of actually being queried if planned to be the ith CSA in the sequence is $(1 - F_{returned}(r_{buyer\rightarrow csa}))^{i-1}$, as all of the first $i-1$ CSAs need to return a quote greater than $r_{buyer\rightarrow csa}$. The probability of being queried at all is therefore $1/N \sum_{i=1}^{N}(1 - F_{returned}(r_{buyer\rightarrow csa}))^{i-1} = \frac{1-(1-F_{returned}(r_{buyer\rightarrow csa}))^N}{NF_{returned}(r_{buyer\rightarrow csa})}$. Therefore the probability that a queried CSA is the ith CSA queried by the buyer is: $P_{queried}(i) = \frac{F_{returned}(r_{buyer\rightarrow csa})(1-F_{returned}(r_{buyer\rightarrow csa}))^{i-1}}{1-(1-F_{returned}(r_{buyer\rightarrow csa}))^N}$. If the CSA is the ith CSA queried, then $i-1$ CSAs have already been queried by the buyer, yielding a quote greater than $r_{buyer\rightarrow csa}$ and $N-i$ CSAs still need to be queried. In order for the quote q to be picked by the buyer, all CSAs from both groups need to return a quote greater than q, which has a chance of $(1 - F_{returned}(q))^{N-i}(1 - F_{returned}(q|q > r_{buyer\rightarrow csa}))^{i-1}$.

As expected, the function $P(q)$ decreases as q increases and as N increases. When there is no competition at all (i.e., $N = 1$), $P(q) = 1$ is obtained for any quote $q \leq r_{buyer\rightarrow seller}$.

We can now introduce Theorem 4.4, which gives a necessary and sufficient condition for the existence of the equilibrium and the set of equations from which it can be extracted.

THEOREM 4.4. *A pure equilibrium with N self-interested actively searching CSAs will exist if and only if $M > c_{csa\rightarrow seller}F(r_{buyer\rightarrow csa})$. The equilibrium structure is: (a) $r_{buyer\rightarrow csa} = r_{csa\rightarrow seller} < r_{buyer\rightarrow seller}$, whenever $M > M'$, in which case it can be calculated by solving the set of equations* (10) *and* (11)*, substituting $r_{csa\rightarrow seller} = r_{buyer\rightarrow csa}$ (complemented by* (8), (9) *and* (12) *for $f_{returned}(q)$, $F_{returned}(q)$ and $P(q)$); (b) $r_{buyer\rightarrow csa} < r_{csa\rightarrow seller} \leq r_{buyer\rightarrow seller}$, whenever $M'' \leq M \leq M'$, in which case it can be calculated by solving the set of equations* (6), (10) *and* (11) *(complemented by* (8), (9) *and* (12) *for $f_{returned}(q)$, $F_{returned}(q)$ and $P(q)$);*
where: $M' = \frac{c_{csa\rightarrow seller}}{F(r_{buyer\rightarrow csa})}$, substituting the value $r_{buyer\rightarrow csa}$ for which (11) *holds when using $r_{csa\rightarrow seller} = r_{buyer\rightarrow csa}$ for $f_{returned}(q)$ in* (8)[6]*; and M'' is the value of M for which the solution to the set of equations* (6), (10) *and* (11) *is of the form: $r_{buyer\rightarrow csa} < r_{csa\rightarrow seller} = r_{buyer\rightarrow seller}$.*

PROOF. We first prove that if $M > M'$ then there exists an equilibrium according to (a). Substituting $f_{returned}(q) = f(q)/F(r_{buyer\rightarrow csa})$ in (11) (since $r_{csa\rightarrow seller} = r_{buyer\rightarrow csa}$) obtains the expense-minimizing reservation value of the buyer $r_{buyer\rightarrow csa}$ (which is also used by the CSAs when deciding on querying sellers). In this case, after querying the first CSA, the buyer terminates her search, since the quote obtained is necessarily lower than $r_{buyer\rightarrow csa}$ (and also lower than $r_{buyer\rightarrow seller}$ that is calculated according to (10), because $c_{buyer\rightarrow seller} \geq c_{buyer\rightarrow csa}$ and the quote received from the CSA is of a more favorable distribution as discussed above). None of the CSAs will individually deviate from the $r_{csa\rightarrow seller}$ strategy, since by lowering their reservation value they only incur additional cost (as the number of sellers they query increases) and do not increase their revenue (since buyers will buy based on the quote they return anyhow). Using a greater reservation value $r'_{csa\rightarrow seller} > r_{buyer\rightarrow csa}$, cannot be expected-benefit-maximizing for the CSA, since any quote higher than $r_{buyer\rightarrow csa}$ will result in having the buyer query other CSAs, thus if it is beneficial to query an additional seller when the best quote obtained so far is greater than $r'_{csa\rightarrow seller}$, then it is also beneficial to query if the best quote obtained so far is greater than $r_{buyer\rightarrow csa}$ (hence a conflict with the reservation value definition). The expected cost of search for the CSA in this case is $c_{csa\rightarrow seller}/F(r_{buyer\rightarrow csa})$ and its revenue is M. Therefore, since $M > M' = \frac{c_{csa\rightarrow seller}}{F(r_{buyer\rightarrow csa})}$, the expected benefit of the CSA is necessarily non-negative. Since buyers are using their expected-expense-minimizing strategy, and no individual CSA has an incentive to deviate from its strategy, and CSAs' expected benefit is non-negative, the solution is stable.

Next, we prove that if $M'' \leq M \leq M'$, then there exists an equilibrium according to (b). Notice that according to the first part of the proof, when $M = M' = \frac{c_{csa\rightarrow seller}}{F(r_{buyer\rightarrow csa})}$, Equation 6 is satisfied (as $P(q) = 1$ for any $q < r_{csa\rightarrow seller}$ and $P(r_{csa\rightarrow seller}) = 0$). Therefore, each CSA is using its expected-benefit-maximizing reservation value and buyers are using their expected-expense-minimizing strategy. This set of strategies is thus in equilibrium. Any decrease in M will result in an increase in $r_{csa\rightarrow seller}$ (according to (6)) and consequently an increase in $r_{buyer\rightarrow csa}$ (according to (11)). Since $f_{returned}(q) = 0$ for $q > r_{csa\rightarrow seller}$, it is guaranteed that a solution $r_{buyer\rightarrow csa}$ according to (11) satisfies $r_{buyer\rightarrow csa} < r_{csa\rightarrow seller}$. The solution obtained in this case is in equilibrium because buyers are using their expected-benefit-minimizing strategies and sellers are using their expected-expense maximizing strategies. Due to the continuity of

[6]In case there are several $r_{buyer\rightarrow csa}$ that satisfy (11), the highest value among them ought to be used.

$r_{csa\to seller}$ according to (6), it is guaranteed that for some $M > M'$ a solution of the form $r_{csa\to seller} = r_{buyer\to seller}$ (as $r_{buyer\to seller}$ is determined according to (10), and is independent of $r_{csa\to seller}$ and $r_{buyer\to csa}$) exists. This assures the existence of M''. Any further decrease in M will require an increase in $r_{csa\to seller}$ (and $r_{buyer\to csa}$), however according to Proposition 4.3 such a solution is not stable. □

According to Theorem 4.4, there are cases where an equilibrium does not exist. One such trivial case is when $M < c_{csa\to seller}$. Here, the CSAs' expected benefit is necessarily negative, regardless of the value of N, and therefore in equilibrium none of the CSAs will be willing to offer its services to buyers, and buyers will query sellers directly (or use buyer-operated CSAs).[7] Nevertheless, in many settings equilibrium exists, as we demonstrate in the next section.

We can now formally express the expected benefit of CSAs and sellers and the expected expense of buyers. The expected benefit of the CSA whenever contacted by the buyer, denoted V_{csa}, is the following modification of (2):

$$V_{csa} = \frac{M \int_{q=0}^{r_{csa\to seller}} P(q)f(q)dq - c_{csa\to seller}}{F(r_{csa\to seller})} \tag{13}$$

PROPOSITION 4.5. *The expected benefit of each CSA in equilibrium is zero if the equilibrium is of type $r_{buyer\to csa} < r_{csa\to seller} \le r_{buyer\to seller}$ and greater than or equal to zero if of type $r_{buyer\to csa} = r_{csa\to seller}$.*

PROOF. Substituting $q = r_{csa\to seller}$ and $F_{returned}(r_{csa\to seller}) = 1$ (according to (9)) in (12) obtains $P(r_{csa\to seller}) = 0$. (Intuitively, since the buyer uses a reservation value $r_{buyer\to csa} \le r_{csa\to seller}$, then returning a quote $q = r_{csa\to seller}$ will result in a zero chance of having the buyer purchase the product at this price, since one of the other CSAs necessarily returns a lower quote). Substituting $P(r_{csa\to seller}) = 0$ in (6) obtains $c_{csa\to seller} = M \int_{q=0}^{r_{csa\to seller}} P(q)f(q)dq$ for the case $r_{buyer\to csa} < r_{csa\to seller}$ and $c_{csa\to seller} < M \int_{q=0}^{r_{csa\to seller}} P(q)f(q)dq$ for the case $r_{buyer\to csa} = r_{csa\to seller}$. Substituting $c_{csa\to seller} = M \int_{q=0}^{r_{csa\to seller}} P(q)f(q)dq$ and $c_{csa\to seller} < M \int_{q=0}^{r_{csa\to seller}} P(q)f(q)dq$ in (13) results in a zero and positive value, respectively. □

COROLLARY 4.6. *The expected overall number of sellers queried by the CSAs as a result of the buyer's querying of CSAs, denoted $E_{sellers\_queried}$, equals: (a) $M/c_{csa\to seller}$ if equilibrium is of structure $r_{buyer\to csa} < r_{csa\to seller} \le r_{buyer\to seller}$; or (b) $1/F(r_{csa\to seller})$ if of structure $r_{buyer\to csa} = r_{csa\to seller}$.*

PROOF. For case (a) we know, based on Proposition 4.5, that the expected benefit of each CSA is zero. The expected gain resulting from obtaining the payment M is thus equal to the expected cost that the CSAs incur along their search, i.e., $E_{sellers\_queried} c_{csa\to seller} = M$. The proof for case (b) is trivial — since the buyer terminates the search process after querying a single CSA, the expected number of sellers queried throughout the process is $1/F(r_{csa\to seller})$. □

The result of a zero expected net benefit of each CSA in the case $r_{buyer\to csa} < r_{csa\to seller} \le r_{buyer\to seller}$ resembles, in a way, the one obtained with equilibrium pricing in the Bertrand competition model. The Bertrand competition model considers non-cooperating firms producing homogeneous products, having the same marginal cost. The firms compete in price, choosing their respective prices simultaneously. The equilibrium result in this case is marginal cost pricing. While the self-interested-CSAs

[7]In such settings a mixed equilibrium may exist, in which CSAs randomize on whether to participate in the game in the first place or not, though the analysis of mixed equilibrium is beyond the scope of this paper.

problem cannot be directly mapped to the Bertrand competition model, several points of similarity may be pointed out: First, players are homogeneous and compete based on the quote they return, where the buyer always buys from the one returning the lowest quote. Second, any quote that the CSA returns which is above the reservation value will surely not be accepted by the buyer. Third, when using a reservation value equal to the reservation value used by the other CSAs, the CSAs actually share the revenue (payment M) among them (equivalent to the case of setting the same pricing in the Bertrand model). In this case, if a quote $r_{csa \rightarrow seller}$ is returned to the buyer, the chance it is accepted is zero, and therefore each CSA has an incentive to use a slightly reduced reservation value. This continues until all CSAs use a reservation value in which their expected benefit is zero, which is equivalent to having all sellers in Bertrand's model use their marginal cost of production. Finally, just as in the Bertrand competition model, the equilibrium is based on weakly dominated strategies, as each CSA can choose not to query any seller at all (which is equivalent to returning a quote of an infinite value, that is therefore never used by the buyer), resulting in a zero benefit. One model characteristic that does not apply in our case is the discontinuity of payoffs; while in Bertrand's model the use of pricing above the other firm's price results in a zero overall benefit, in our case increasing the reservation price does not eliminate the chance that the quote the CSA will return will be accepted by the buyer. Also, reducing the reservation value beyond the one used by others does not result in taking over the entire demand.

One may wonder, given the last result, what is the incentive for the CSAs to operate in competitive environments in cases where their expected net-benefit is zero. This question is addressed in Section 5, giving evidence to market designer's incentive to reward CSAs for operating in their markets.

The expected price quote with which the buyer eventually ends up in equilibrium, denoted $E_{returned}$, is given by:

$$E_{returned} = (1 - (1 - F_{returned}(r_{buyer \rightarrow csa}))^N) E[x|x < r_{buyer \rightarrow csa}] + \\ + (1 - F_{returned}(r_{buyer \rightarrow csa}))^N E_N[X|X > r_{buyer \rightarrow csa}] \tag{14}$$

where $E[x|x < r_{buyer \rightarrow csa}]$ is the expected quote given that it is below $r_{buyer \rightarrow csa}$. This relates to the case where one of the CSAs returned a quote below $r_{buyer \rightarrow csa}$, upon which the buyer terminated her search. The occurrence probability of this event is $(1-(1-F_{returned}(r_{buyer \rightarrow csa}))^N)$ which is calculated as the complementing probability of the event where all CSAs return a quote greater than $r_{buyer \rightarrow csa}$. $E_N[X|X > r_{buyer \rightarrow csa}]$ is the expected value of the minimum quote out of a sample of N quotes, given that all N quotes are above $r_{buyer \rightarrow csa}$ (and necessarily below $r_{csa \rightarrow seller}$). This corresponds to the case where all CSAs return a quote greater than $r_{buyer \rightarrow csa}$ and has an occurrence probability of $(1-F_{returned}(r_{buyer \rightarrow csa}))^N$. $E[x|x < r_{buyer \rightarrow csa}]$ and $E_N[X|X > r_{buyer \rightarrow csa}]$ can be calculated according to:

$$E[x|x < q] = \int_{y=0}^{q} \frac{y f_{returned}(y)}{F_{returned}(q)} dy \tag{15}$$

and:

$$E_N[X|X > r_{buyer \rightarrow csa}] = \int_{y=r_{buyer \rightarrow csa}}^{r_{buyer \rightarrow seller}} y f_N(y|y > r_{buyer \rightarrow csa}) dy \tag{16}$$

where $f_N(y|y > r_{buyer \rightarrow csa})$ is the probability distribution function of the minimum value in a sample of N observations taken from a distribution $f_{returned}(q)$, given that all observations are above $r_{buyer \rightarrow csa}$. This is calculated as: $f_N(y|y > r_{buyer \rightarrow csa}) =$

$\frac{Nf_{returned}(y)}{1-F_{returned}(r_{buyer\to csa})}\left(\frac{1-F_{returned}(y)}{1-F_{returned}(r_{buyer\to csa})}\right)^{N-1}$ (as the probability that a quote returned by each CSA in this case is equal to y is $\frac{f_{returned}(y)}{1-F_{returned}(r_{buyer\to csa})}$ and the probability that all other CSAs return a quote greater than y is $\left(\frac{1-F_{returned}(y)}{1-F_{returned}(r_{buyer\to csa})}\right)^{N-1}$). In the case where the equilibrium is of structure $r_{buyer\to csa} = r_{csa\to seller}$, $F_{returned}(r_{buyer\to csa}) = 1$, and consequently $E_{returned} = E[x|x < r_{buyer\to csa}]$ according to (14). (The quote obtained from the first CSA is below $r_{buyer\to csa}$ and the buyer terminates her search right after.)

The probability that the buyer will query at least i CSAs is given by: $(1 - F_{returned}(r_{buyer\to csa}))^{i-1}$ (since it requires that the first $i-1$ CSAs return a price above $r_{buyer\to csa}$). The expected number of CSAs that the buyer samples, denoted $E_{CSAs\_queried}$, is thus given by:

$$E_{CSAs\_queried} = \sum_{i=1}^{N}(1 - F_{returned}(r_{buyer\to csa}))^{i-1} = \frac{1-(1-F_{returned}(r_{buyer\to csa}))^N}{F_{returned}(r_{buyer\to csa})} \tag{17}$$

The buyer's expected expense is thus given by: $E_{returned} + c_{buyer\to csa}E_{CSAs\_queried}$. The seller's expected benefit is: $E_{returned} - M$. Since the buyer does not necessarily contact all CSAs in each search sequence, the expected benefit of each CSA per buyer's search is given by $V_{csa}E_{CSAs\_queried}/N$, which equals either zero (according to Proposition 4.5) or $(M - \frac{c_{csa\to seller}}{F(r_{csa\to seller})})/N$ for the case where $r_{buyer\to csa} = r_{csa\to seller}$ (as the one CSA selected out of the N available gains M, however incurs an expected cost of $c_{csa\to seller}/F(r_{csa\to seller})$).

The remainder of the section discusses the implications of changes in some of the model assumptions over the equilibrium analysis above.

4.4. Equilibrium with a Single Self-Interested CSA

The analysis given in this section relies on the competition between CSAs. If there is only a single CSA, then the expected-benefit-maximizing strategy for that CSA would be to take over as much of the buyer's surplus, due to the use of the CSA, as possible. We analyze this situation in the following paragraphs.

THEOREM 4.7. *An equilibrium in the case $N = 1$ exists whenever $M \geq c_{csa\to seller}/F(r_{buyer\to seller})$, and its structure is $r_{csa\to seller} = r_{buyer\to seller}$ where $r_{buyer\to seller}$ is the solution R_i to* (3)*, using the cost $c_i = c_{buyer\to seller}$.*

PROOF. Using a reservation value $r_{csa\to seller} < r_{buyer\to seller}$, the CSA will spend more on search however will not change the value of $P(q)$ for the quote eventually returned (which is already 1 for $r_{csa\to seller} = r_{buyer\to seller}$). Using $r_{csa\to seller} > r_{buyer\to seller}$, the CSA risks returning quotes greater than $r_{buyer\to seller}$ for which $P(q) = 0$ (as the buyer will continue querying sellers until obtaining a quote lower than $r_{buyer\to seller}$).

The buyer's expected expense in this case is $c_{buyer\to csa} + E[x|x \leq r_{buyer\to seller}]$ and the seller's expected benefit is $E[x|x \leq r_{buyer\to seller}] - M$. The CSA's expected benefit is $M - c_{csa\to seller}/F(r_{csa\to seller})$. The CSA will offer its services to buyers only if its expected benefit is non-negative, i.e., if $M \geq c_{csa\to seller}/F(r_{buyer\to seller})$. □

PROPOSITION 4.8. *The buyer's expected expense in equilibrium with $N = 1$ is always worse (i.e., greater) than her expected expense with $N > 1$.*

PROOF. Since the value of $r_{buyer\to seller}$ does not change as a function of the number of CSAs, the reservation value set by the CSA when having no competition is always greater than or equal to the one used by competing CSAs (based on Proposition 4.3). Therefore, even if the buyer queries only one CSA in the competing-CSAs case, she ends up paying the same search cost $c_{buyer \to csa}$, however with a lower expected quote (a single observation from distribution $f(x|x \le r_{buyer \to seller})$ in comparison to a single observation from $f(x|x \le r'_{buyer \to seller} < r_{buyer \to seller})$). Therefore, the buyer ends up with a lower expected expense in the case of competing CSAs, and consequently will always prefer such a setting over one with a single self-interested CSA (if an equilibrium exists in both cases). □

In Section 5 we demonstrate the existence of mutually-beneficial stable solutions even when having a single self-interested CSA in the market. Yet, as discussed and illustrated there, the competition between CSAs substantially expands the range of settings where such beneficial solutions can hold.

4.5. Re-Querying the same CSA

The fact that CSAs can identify buyers enables the proof of Proposition 4.2 regarding having the CSAs being queried only once in equilibrium. If the CSAs cannot identify buyers, then buyers can potentially re-query the same CSA over and over again until satisfied with the quote it returns. In this case, the number of CSAs, N, does not affect the buyer's querying strategy at all. The buyer will keep querying CSAs (or the same CSA) until the quote returned is below the reservation value she uses. Similarly, the value of N does not affect the CSAs' strategy, as any CSA can be queried indefinitely.

PROPOSITION 4.9. *The equilibrium reservation value used by the CSAs in case re-querying is possible satisfies:* $r_{csa \to seller} = r_{buyer \to csa}$.

PROOF. Assume otherwise, i.e., $r_{csa \to seller} > r_{buyer \to csa}$. Since the buyer terminates her search only when receiving a quote $q \le r_{buyer \to csa}$, the probability function $P(q)$ defined in (12) obtains zero for any $q > r_{buyer \to csa}$ and $P(q) = 1$ otherwise. From this point on, the proof resembles the proof given for Proposition 4.3 (i.e., showing that if it is beneficial to query an additional seller if receiving a quote greater than $r_{csa \to seller}$, then so is the case for any other quote in the interval $(r_{buyer \to csa}, r_{csa \to seller})$), and thus this part is omitted. □

Based on Proposition 4.9, the buyer will necessarily terminate its search after querying a single CSA, since the quote it obtains is never higher than the reservation value $r_{buyer \to csa}$. Still, in this case, there can be a great benefit in having self-interested CSAs in the market, and even a mutually-beneficial equilibrium may exist, as illustrated in Section 5.

4.6. Using a Buyer-Operated CSA Alongside Self-Interested CSAs

The model assumes that if the buyer resumes her search after exhausting all of the available CSAs, then her search may be continued by querying sellers directly, with a cost $c_{buyer \to seller} \ge c_{csa \to seller}$. We now prove that in the case where the buyer may continue her search using a cost $c_{csa \to seller}$, i.e., using a buyer-operated CSA, sellers can never benefit from offering a payment M, as compared to the case where CSAs are exclusively buyer-operated.

PROPOSITION 4.10. *Sellers' expected benefit in equilibrium where buyer-operated CSAs are used alongside self-interested CSAs is always smaller than in the case where only buyer-operated CSAs are used.*

PROOF. We prove that with the hybrid search that combines self-interested CSAs and buyer-operated CSAs, the expected quote that the buyer ends up with is necessarily lower than when using only a buyer-operated CSA. Therefore the seller loses both because of the reduced purchase price and because of the payment M to the CSA.

With the addition of self-interested CSAs, the buyer can simply set the value of $r_{buyer \to csa}$ to be the optimal reservation value used in the buyer-operated-CSA scenario. If all N CSAs have been exhausted, then the buyer continues querying sellers directly, using the same reservation value. Since $c_{buyer \to csa} \leq c_{buyer \to seller}$, and since each CSA returns a quote from a more favorable distribution ($f_{returned}(q)$ in comparison to $f(q)$), the CSA is likely to end up with a better expected quote and spend less on the search overall, compared to the buyer-operated case. Since setting the value of $r_{buyer \to csa}$ to the reservation value used in the buyer-operated-CSA scenario is a suboptimal strategy for the buyer, the optimal buyer's strategy yields an even lower expected quote, resulting in the sellers' loss in comparison to the buyer-operated case. □

5. WHEN EVERYBODY WINS

In this section we illustrate the performance of buyers, sellers and CSAs under different conditions, when self-interested and when buyer-operated CSAs are used. The illustrations are based on a synthetic environment that uses the uniform distribution function ($f(q) = 1$, $F(q) = q$, $0 \leq q \leq 1$) as the distribution of price quotes. It is notable that while sellers' pricing is not affected by the CSAs' search strategy, their profit is affected by the CSAs' search. This is reflected by the expected price at which the buyer buys the product. Generally, the greater the extent of the search and its efficiency, the lower the expected payment each seller receives when a buyer buys from it.

The most important conclusion made, based on the figures presented in this section, is that there are cases where the use of competing CSAs can result in a mutually-beneficial equilibrium, where buyers end up with a reduced expected expense and sellers end up with a greater expected benefit (in comparison to the buyer-operated-CSA case). Furthermore, as demonstrated throughout the section, there are many settings where, despite not having both parties benefit from the transition to self-interested competing CSAs, the sum of changes in the buyer's expected expense and the seller's expected benefit is positive. This result suggests various opportunities for the market designer to improve market performance by dictating the number of CSAs in the market, N, the payment, M, and the appropriate side-payments that will ensure that both parties benefit from choosing the use of self-interested CSAs (if such mechanisms are available).

One key observation that helps in understanding the existence of such a mutually-beneficial equilibrium is that, in many settings, the use of competing CSAs results in an increase (rather than the decrease one would expect to find) in sellers' expected benefit compared to the buyer-operated case, despite having the sellers pay for the CSAs' search. To take this result to the extreme, the second part of the section demonstrates that even if the seller had accounted for all the costs associated with querying CSAs, her expected benefit could have increased. This is achieved by introducing an equilibrium analysis for the case where a meta-CSA which costs are covered by the seller is available in the market. The analysis is followed by an illustrative evidence for a possible increase in the seller's expected benefit due to the introduction of the seller-paid meta-CSA.

5.1. Equilibrium with Self-Interested CSAs

For exposition purposes, any curve in the figures given in this section is marked according to the model it was produced: curves depicting a result obtained with self-interested CSAs are marked by "(SI)" and those depicting a result related to a buyer-operated case are marked by "(BO)". The CSA's expected benefit, whenever depicted, is of course related to the self-interested CSAs case. The vertical axis in all figures rep-

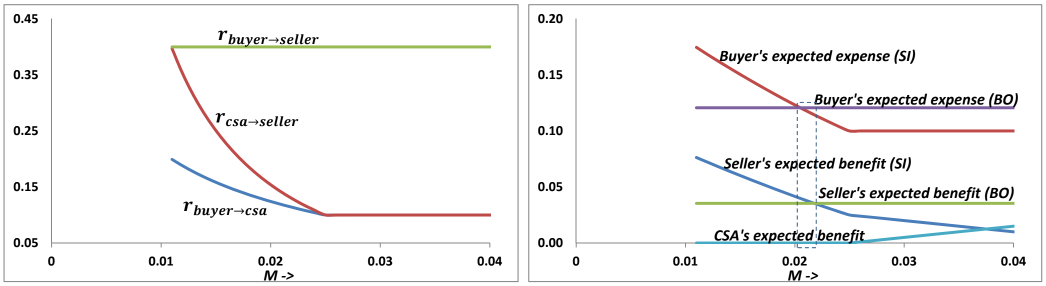

Fig. 1. The structure of the equilibrium strategies (left) and the resulting buyer's expected expense and seller's and CSA's expected benefit (right), as a function of the payment M offered by sellers to CSAs (horizontal axis) for the self-interested and buyer-operated cases. The parameters used are: $c_{buyer\rightarrow csa} = 0.05$, $c_{buyer\rightarrow seller} = 0.08$, $c_{csa\rightarrow seller} = 0.0025$ and $N = 3$. An equilibrium in this setting exists for $M > 0.011$ and the reservation values satisfy $r_{buyer\rightarrow csa} \leq r_{csa\rightarrow seller} \leq r_{buyer\rightarrow seller}$, such that the increase in M results in a decrease in $r_{csa\rightarrow seller}$ and $r_{buyer\rightarrow csa}$. Both the buyer's expected expense and the seller's expected benefit decrease as M increases. The range of M values for which a mutually-beneficial equilibrium exists is marked with the dotted rectangle.

resents monetary values, thus the higher the value, the greater the expected expense is for the buyer (and thus the worse her performance) and the greater the expected benefit is for the seller. When the buyer is the one operating the CSA, both her own and the seller's performance do not depend on M (nor N). The buyer's cost-minimizing strategy in this case is to search sequentially according to the reservation value calculated in (3). Therefore buyer-operated curves in the figures are typically horizontal. On the other hand, in the scenario where CSAs are self-interested agents, different levels of competition (resulting from the extent of the search set by each CSA) yield a different performance for the buyers and sellers.

We begin with the structure of the equilibrium strategies (following Proposition 4.3 and Theorem 4.4). Figure 1 depicts the reservation values used by buyers and CSAs in equilibrium (left graph) and the resulting buyer's expected expense and seller's and CSA's expected benefit (right graph), as a function of the payment M offered by sellers to CSAs (horizontal axis) for the self-interested and buyer-operated cases. The values of the other model parameters were set to: $c_{buyer\rightarrow csa} = 0.05$, $c_{buyer\rightarrow seller} = 0.08$, $c_{csa\rightarrow seller} = 0.0025$ and $N = 3$. As observed from the figure, an equilibrium in this setting exists only when $M \geq 0.011$. The reservation values satisfy $r_{buyer\rightarrow csa} \leq r_{csa\rightarrow seller} \leq r_{buyer\rightarrow seller}$, such that the increase in M results in a decrease in $r_{csa\rightarrow seller}$ and $r_{buyer\rightarrow csa}$. Intuitively, since the expected benefit of the CSAs is zero (for an equilibrium where $r_{buyer\rightarrow csa} < r_{csa\rightarrow seller} \leq r_{buyer\rightarrow seller}$), the increase in M results in an overall increase in the amount of search (see Corollary 4.6). This increase in the amount of search is reflected by a decrease in the reservation value used by each of the CSAs, and consequently a decrease in the reservation value used by the buyer, $r_{buyer\rightarrow csa}$. Once $r_{buyer\rightarrow csa} = r_{csa\rightarrow seller}$, the CSAs have an actual positive expected benefit, as observed in the right graph of Figure 1. From the right graph, we also observe that both the buyer's expected expense and the seller's expected benefit decrease as M increases. The decrease, in this case, is explained by the increase in the payment to CSAs that the sellers incur and the improvement in the buyer's efficiency of the search (as CSAs use lower reservation values when M increases). Once $r_{buyer\rightarrow csa} = r_{csa\rightarrow seller}$, any further increase in M will not affect the buyer's expected expense, since the reservation value used by the CSAs remains steady. The seller's expected benefit, however, will keep decreasing, since despite having the buyer purchase the product at the same expected price, the payment that the seller needs to pay the

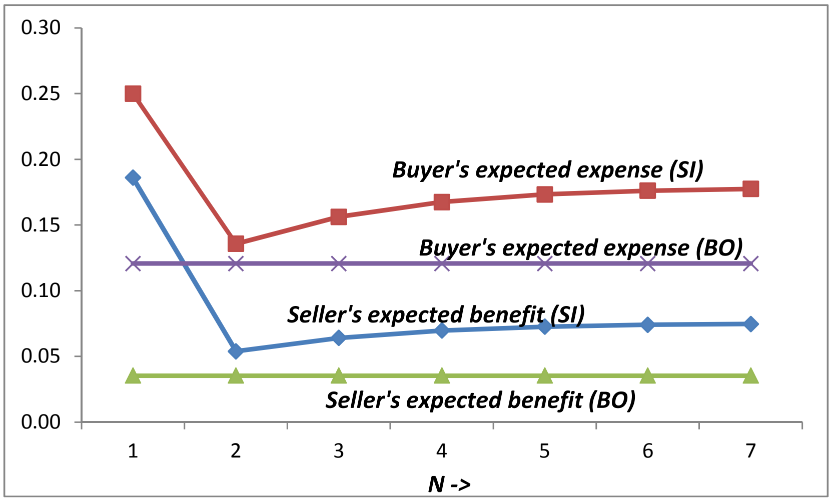

Fig. 2. The effect of the number of self-interested CSAs available, N, over the expected benefit of sellers and the expected expense of buyers. The parameters used are: $c_{buyer \to csa} = 0.05$, $c_{buyer \to seller} = 0.08$, $c_{csa \to seller} = 0.0025$ and $M = 0.014$. The case where $N = 1$ is the worst and $N = 2$ is the best from the buyer's perspective. From $N = 2$ and on, any increase in N results in an increase in the buyer's expense.

CSAs increases. The decrease in the sellers' expected benefit for this range is therefore equal to the increase in the CSAs' expected benefit. A mutually-beneficial equilibrium solution is obtained for any M value in the interval $0.020 - 0.022$. Furthermore, it is notable that the sum of changes in the buyer's expected expense and the seller's expected benefit is positive for the range ($0.018 - 0.035$), with a peak of 0.01 at $M = 0.025$. Therefore, given an appropriate side-payment mechanism, a greater extent of such beneficial solutions can be engineered.

Figure 2 depicts the effect of the number of self-interested CSAs available, N, over the expected benefit of sellers and the expected expense of buyers. The values of the other model parameters were set to: $c_{buyer \to csa} = 0.05$, $c_{buyer \to seller} = 0.08$, $c_{csa \to seller} = 0.0025$ and $M = 0.014$. As observed from the figure, the case where $N = 1$ is the worst from the buyer's perspective, as it results in the highest expected expense. This is explained (see Proposition 4.8) by the fact that, when not facing a competition, the CSA uses a reservation value that equals $r_{buyer \to seller}$, which is the worst possible scenario from the buyer's perspective. (If there are competing CSAs, the reservation value used by CSAs satisfies $r_{csa \to seller} \leq r_{buyer \to seller}$). Therefore, the transition from $N = 1$ to $N = 2$ results in a substantial decrease in the buyer's expected expense. Nevertheless, from $N = 2$ and on, any increase in N results in an increase in the buyer's expense. This non-intuitive result (a decrease in the buyer's expense despite the seemingly increased competition resulting from the increase in the number of competing CSAs) is explained, once again, by the decrease in the extent of search performed by each of the CSAs available to the buyer. In particular, we note that since the expected number of sellers queried remains constant (see Corollary 4.6), the buyer's search is now executed with the use of more CSAs (hence more is spent in terms of $c_{buyer \to csa}$) and the reservation value used by each CSA increases (hence each CSA searches to a lesser extent). The intuition for the behavior exhibited by the seller's curve is similar (as any increase in the buyer's expected payment for the product results in an equal increase in the seller's expected benefit).

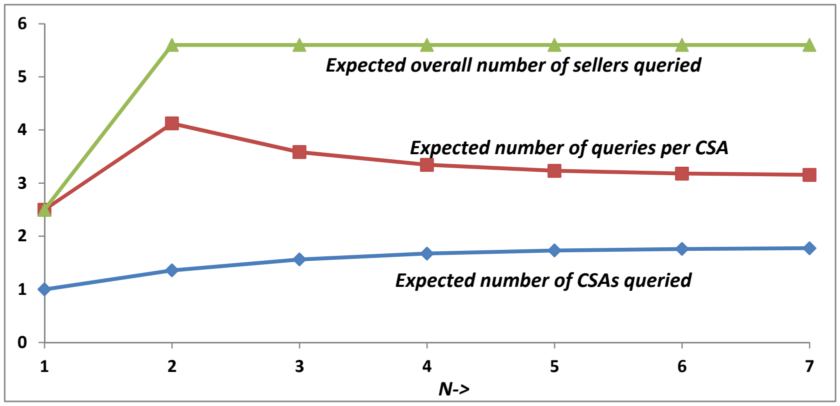

Fig. 3. The expected number of CSAs queried by the buyer, the expected number of sellers that each CSA queries and the product of the two, which is the expected number of sellers queried overall, for different values of N. The parameters used are: $c_{buyer \to csa} = 0.05$, $c_{buyer \to seller} = 0.08$, $c_{csa \to seller} = 0.0025$ and $M = 0.014$. The expected number of sellers queried overall is fixed for $N \geq 2$.

Figure 3 complements Figures 1 and 2, illustrating that the expected number of sellers queried when $r_{buyer \to csa} < r_{csa \to seller} \leq r_{buyer \to seller}$ is fixed, regardless of the number of available competing self-interested CSAs ($N \geq 2$), as given in Corollary 4.6. The figure depicts the expected number of CSAs queried by the buyer, $E_{CSAs\_queried}$ (see (17)), the expected number of sellers that each CSA queries (given by $1/F(r_{csa \to seller})$) and the product of the two, which is the expected number of sellers queried overall, for different values of N. The values of the other model parameters were set for this figure to: $c_{buyer \to csa} = 0.05$, $c_{buyer \to seller} = 0.08$, $c_{csa \to seller} = 0.0025$ and $M = 0.014$ (same as in Figure 2). As observed from the figure, the expected number of sellers queried overall is indeed fixed for $N \geq 2$.

Figure 4(a) demonstrates that a mutually-beneficial equilibrium may exist even if only one self-interested CSA is available (see Section 4.4). The setting uses the parameters $c_{buyer \to csa} = 0.002$, $c_{buyer \to seller} = 0.003$ and $c_{csa \to seller} = 0.001$. Any M value below 0.013 in this case produces a negative expected benefit for the CSA, and therefore equilibrium exists only for settings where $M \geq 0.013$. As observed from the figure, a mutually-beneficial solution exists for any M in the interval $(0.13-0.16)$. A positive sum of differences between a buyer's and a seller's performances with self-interested CSAs compared with buyer-operated CSAs holds in the interval $(0.13 - 0.20)$. Figure 4(b) demonstrates a similar type equilibrium solution for the case where the buyer can reuse CSAs (see Section 4.5). The setting uses the parameters $c_{buyer \to csa} = c_{buyer \to seller} = 0.001$ and $c_{csa \to seller} = 0.00005$. The expected benefit of the CSA in this case is zero (see Proposition 4.5). As observed from the figure, a mutually-beneficial equilibrium exists for any M in the interval $(0.00278 - 0.00309)$ and an equilibrium where the sum of differences is positive exists for any $M < 0.004$.

To shed some light on the role of CSAs' competition in reaching mutually-beneficial equilibrium solutions, Figure 5 depicts the buyer's expected expense (left graph) and the seller's expected benefit (right graph) as a function of $c_{buyer \to csa}$ (horizontal axis) for different values of N. The setting uses the parameters $c_{buyer \to seller} = 0.08$, $c_{csa \to seller} = 0.0025$ and $M = 0.019$. For $N = 1$ there is no value of $c_{buyer \to csa}$ for which a mutually-beneficial equilibrium is obtained in this setting, and therefore the $N = 1$ curve is excluded from both graphs. From the figure we observe that a mutually-beneficial equilibrium is obtained for settings with 2 or more competing CSAs, when $c_{buyer \to csa}$ obtains values from the interval $0.061 - 0.069$ (and for $N = 2$ also from the

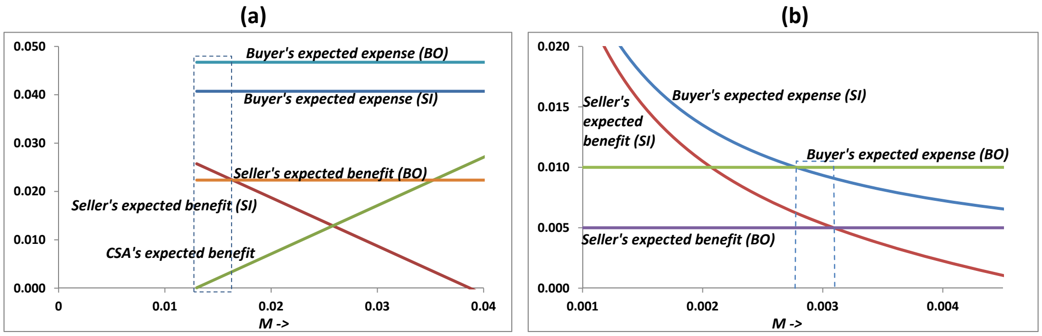

Fig. 4. (a) Buyer's expected expense and seller's expected benefit when only one self-interested CSA is available. The parameters used are: $c_{buyer \to csa} = 0.002$, $c_{buyer \to seller} = 0.003$ and $c_{csa \to seller} = 0.001$. A mutually-beneficial solution exists for any M in the interval $(0.13 - 0.16)$, marked with the dotted rectangle. (b) Buyer's expected expense and seller's expected benefit when the buyer can re-use CSAs. The parameters used are: $c_{buyer \to csa} = c_{buyer \to seller} = 0.001$ and $c_{csa \to seller} = 0.00005$. The expected benefit of the CSA in this case is zero. A mutually-beneficial equilibrium solution holds for any M in the interval $(0.00278 - 0.00309)$, marked with the dotted rectangle.

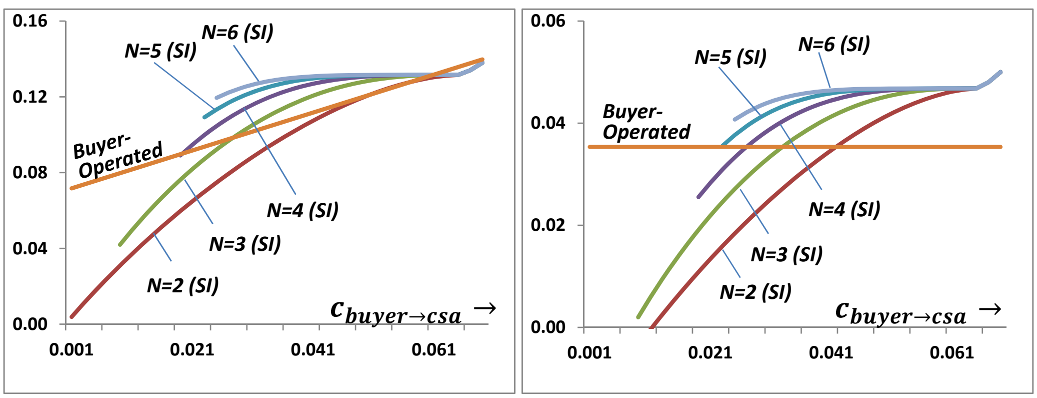

Fig. 5. The buyer's expected expense (left graph) and the seller's expected benefit (right graph) as a function of $c_{buyer \to csa}$ (horizontal axis) for different values of N. The parameters used are: $c_{buyer \to seller} = 0.08$, $c_{csa \to seller} = 0.0025$ and $M = 0.019$. For $N = 1$ there is no value of $c_{buyer \to csa}$ for which a mutually-beneficial equilibrium is obtained in this setting, and therefore the $N = 1$ curve is excluded from both graphs.

interval $0.043 - 0.069$). Therefore, the existence of competition can result in substantial expansion of the ranges of the different parameters values for which a mutually-beneficial equilibrium exists, or even enable a mutually-beneficial equilibrium when such an equilibrium cannot exist at all for $N = 1$.

The focus of the above examples is on showing that the sellers can benefit despite being the ones that eventually incur the CSAs' search cost (in the form of the payment M). However, buyers in our model can also benefit from allegedly worsening their searching capability. We exemplify with a case where an increase in $c_{buyer \to seller}$ can actually turn out to be beneficial for the buyer. As discussed in the former section, despite having the option to resume her search by querying sellers directly, the buyer never gets to take advantage of this option in equilibrium because the CSAs use a reservation value lower than $r_{buyer \to seller}$. Still, the option to query sellers directly and the resulting reservation value $r_{buyer \rightarrow seller}$ affect the existence of equilibrium. We demonstrate this using a setting with four self-interested CSAs, with $c_{buyer \rightarrow csa} = 0.05$, $c_{csa \rightarrow seller} = 0.0025$ and $M = 0.011$. In this setting, there is an equilibrium with self-interested CSAs only if $c_{buyer \rightarrow seller} > 0.011$, in which case the expected expense of the buyer is 0.086. For $c_{buyer \rightarrow seller} < 0.011$, there is no equilibrium with self-interested CSAs, and if buyer-operated CSAs are used then the expected expense of the buyer is 0.121. Therefore, in this setting, an increase in $c_{buyer \rightarrow seller}$ to any value above 0.011 actually benefits the buyer, as it enables a favorable equilibrium that could not have held otherwise.

5.2. Using a Meta-CSA

Following the illustrations of mutually-beneficial equilibrium solutions with competing self-interested CSAs, we extend the analysis to settings where a meta-CSA is used. The idea is to give further evidence for the ability of sellers to increase their expected benefit by increasing their liability for the costs of using the CSAs. A meta-CSA is a type of shopbot (often referred to as a "MetaBot") that, instead of searching websites directly, queries shopbots and retrieves and aggregates the information from them [Wan and Liu 2009; Etzioni 1997]. Several such aggregators can be found today. For example, Roboshopper (`roboshopper.net`) returns information from the web's leading online comparison shopping services – including MySimon, Yahoo Shopping, Pricegrabber, Froogle, NexTag and many others – all via one easy-to-use tool. Earlymiser (`Earlymiser.com`) is a meta-comparison shopping engine that pulls the best prices on products from Shopping.com, eBay, Amazon and Yahoo Shopping. Since the meta-CSA saves the buyer much of the cost incurred when querying each of the CSAs individually, one would expect that the inclusion of such a service in the market would reduce buyers' expected expense and similarly decrease sellers' expected benefit. Surprisingly, we manage to demonstrate in the following paragraphs that both buyers and sellers can benefit from the existence of a meta-CSA, even if sellers are requested to cover the expenses of a meta-CSA, in addition to the payment M to the CSA that directed the buyer to their website. It is notable that the meta-CSA in this case is not assumed to be a self-interested agent that attempts to maximize its own expected benefit, but rather can be seen as an artificial means for querying all of the self-interested CSAs.

The analysis of the meta-CSA case follows the analysis that was introduced in the former section, with the only difference being that all CSAs are queried in parallel by the meta-CSA (rather than sequentially by the buyer) once the buyer queries the meta-CSA. We use $c_{buyer \rightarrow meta}$ to denote the cost that the buyer incurs when querying the meta-CSA. The model assumes $c_{buyer \rightarrow meta} \leq c_{buyer \rightarrow seller}$ for the same considerations used for justifying $c_{buyer \rightarrow csa} \leq c_{buyer \rightarrow seller}$. Similarly, we use $r_{buyer \rightarrow meta}$ to denote the reservation value used by the buyer for querying the meta-CSA. The cost incurred by the meta-CSA for every CSA it queries is denoted $c_{meta \rightarrow csa}$.

Since $c_{buyer \rightarrow meta} \leq c_{buyer \rightarrow seller}$ and the meta-CSA returns a quote based on at least N quotes, the buyer will always start her search by querying the meta-CSA and will never query sellers individually afterwards. The probability $P(q)$ in this case is thus given by:

$$P(q) = \begin{cases} (1 - F_{returned}(q))^{N-1} & 0 < q < r_{buyer \rightarrow seller} \\ 0 & q > r_{buyer \rightarrow seller} \end{cases} \tag{18}$$

The value of the CSA's expected-benefit-maximizing reservation value $r_{csa \rightarrow seller}$ can be extracted from (6):

$$c_{csa \rightarrow seller} = M \int_{q=0}^{r_{csa \rightarrow seller}} (P(q) - P(r_{csa \rightarrow seller})) f(q) dq \tag{19}$$

THEOREM 5.1. *In a model of self-interested CSAs with a meta-CSA: (a) An equilibrium in which all CSAs take an active part is necessarily symmetric and will exist if and only if* $MF(r_{buyer \to seller})/N > c_{csa \to seller}$.
(b) The CSAs' expected-benefit-maximizing reservation value in equilibrium is given by $r_{csa \to seller} = F^{-1}(\frac{c_{csa \to seller} N}{M})$ *and their expected benefit is zero.*

PROOF. The proof of the symmetric equilibrium resembles the proof given in Proposition 4.2: Assume otherwise, i.e., there exist at least two CSAs, CSA_α and CSA_β, which reservation values satisfy: $r_{csa_\beta \to seller} < r_{csa_\alpha \to seller}$. However, if CSA_α finds it beneficial to resume its search if obtaining a value $r_{csa_\alpha + \epsilon}$ then necessarily: $c_{csa \to seller} < M \int_{q=0}^{r_{csa_\alpha \to seller} + \epsilon} (P(q) - P(r_{csa_\alpha \to seller} + \epsilon)) f(q) dq$. Since CSA_β necessarily returns a value equal to or less than r_{csa_β}, $P(q) = P(r_{csa_\alpha \to seller} + \epsilon) = 0$ for any $q > r_{csa_\beta \to seller}$ and therefore: $c_{csa \to seller} < M \int_{q=0}^{r_{csa_\beta \to seller} + \epsilon} (P(q) - P(r_{csa_\beta \to seller})) f(q) dq$. Hence, whenever obtaining a value greater than $r_{csa_\beta \to seller}$ the search should be resumed, which contradicts the assumption that $r_{csa_\alpha \to seller}$ is the expected-benefit-maximizing reservation value of CSA_α.

If $MF(r_{buyer \to seller})/N < c_{csa \to seller}$, then the CSAs' expected benefit is necessarily negative, since each of them queries at least $1/F(r_{buyer \to seller})$ sellers (as the reservation value of the CSAs satisfies $r_{csa \to seller} < r_{buyer \to seller}$) and its expected benefit is necessarily M/N (as all CSAs use the same reservation value, and therefore the probability of any individual CSA to obtain the payment M is $1/N$). If an equilibrium in which all CSAs take an active part exists, then according to (18), $P(r_{csa \to seller}) = 0$, and therefore (according to (19)): $c_{csa \to seller} = M \int_{q=0}^{r_{csa \to seller}} P(q) f(q) dq$. Substituting the latter in (13) obtains: $V_{csa} = 0$. Since the term $M \int_{q=0}^{r_{csa \to seller}} P(q) f(q) dq$ increases as $r_{csa \to seller}$ increases, there is necessarily a single equilibrium for any setting where $MF(r_{buyer \to seller})/N > c_{csa \to seller}$.

Since the agents use symmetric strategies, their expected gain is M/N, whereas their expected cost is $c_{csa \to seller}/F(r_{csa \to seller})$. In order to satisfy a zero expected benefit, the following should hold: $M/N = c_{csa \to seller}/F(r_{csa \to seller})$, resulting in: $r_{csa \to seller} = F^{-1}(\frac{c_{csa \to seller} N}{M})$. □

Denoting the expected value of the minimum quote returned by the N CSAs by $E'_{returned}$, the buyer's expected expense in this case is $E'_{returned} + c_{buyer \to meta}$ and the seller's expected benefit is given by: $E'_{returned} - M$. $E'_{returned}$ can be calculated as $E'_{returned} = \int_{q=0}^{r_{csa \to seller}} qN f_{returned}(q)(1 - F_{returned}(q))^{N-1} dq$. Using integration by parts obtains:

$$E'_{returned} = \int_{q=0}^{r_{csa \to seller}} (1 - F_{returned}(q))^N dq \quad (20)$$

We conclude the equilibrium analysis of the meta-CSA case with an important property related to the effect of the number of competing CSAs in the market on the buyer's expected expense.

PROPOSITION 5.2. *The expected expense of the buyer with the presence of a meta-CSA increases as the number of CSAs available,* N, *increases.*

PROOF. Substituting $F(r_{csa \to seller}) = \frac{c_{csa \to seller} N}{M} = AN$ (from Theorem 5.1) in (20) obtains: $E'_{returned} = \int_{q=0}^{r_{csa \to seller}} (1 - \frac{F(q)}{AN})^N dq$ (where $A = \frac{c_{csa \to seller}}{M}$). We show that $(1 - \frac{F(q)}{AN})^N$ increases as N increases for any $F(q)$. Setting $B = F(q)/A$, the derivative of $(1 - \frac{B}{N})^N$ according to N is $\left(1 - \frac{B}{N}\right)^N \left(ln(\frac{N-B}{N}) - \frac{B}{B-N}\right)$. Now notice that

$B = F(q)/A < N$ (since $N = \frac{F(r_{csa \to seller})M}{c_{csa \to seller}}$ and $F(r_{csa \to seller}) > F(q)$) and therefore $\left(ln(\frac{N-B}{N}) - \frac{B}{B-N}\right)$ is always positive. Therefore $(1 - \frac{B}{N})^N$ increases as N increases, hence the integrated term $(1 - \frac{F(q)}{AN})^N$ increases as N increases. Since $r_{csa \to seller}$ also increases as N increases, the value of $E'_{returned}$ increases as N increases. $\square$

Proposition 5.2 suggests a somewhat unique market behavior. While one would expect the increase in the number of competing CSAs to induce competition, resulting in further searching that benefits buyers, the expected expense of buyers actually increases despite not paying for the search (other than the fixed cost $c_{buyer \to meta}$ they incur). This result can be explained using Theorem 5.1. Since the CSAs have a zero net benefit from their search, the increased competition results in an increase in the reservation price they use. The expected minimum quote that the buyer eventually receives is now affected positively by the additional quotes (from the additional CSAs) and negatively by the increased reservation value that the CSAs use. The expected number of searches that each CSA performs is given by $1/F(r_{csa \to seller})$, which equals $\frac{M}{c_{csa \to seller} N}$ according to Theorem 5.1. The overall expected number of searches that the buyer's request yields is thus fixed: $\frac{M}{c_{csa \to seller}}$.[8] However, when CSAs share this amount of searches, there is more parallelism in the search and less sequential decisions are made by the searching CSAs., i.e., the results obtained are based on more parallel searches where each of the parallel searches is less competent than in the case of fewer competing CSAs. Given that the cost of each query is fixed we can consider a "budget" of $M/c_{csa \to seller}$ queries that needs to be spent in search. Spending this budget by making decisions sequentially is always superior to the case of having several decision makers acting in parallel. Therefore, when there are less CSAs, each CSA's search is more effective, to an extent that compensates the loss due to the smaller number of searching CSAs. (Indeed, mathematically, Proposition 5.2 shows that the expected minimum quote received with N CSAs is less than the expected minimum quote received with $N + 1$ CSAs if the expected number of quotes is fixed).

Despite the increase in the buyer's expected expense due to the increase in the number of CSAs, the expected benefit of the seller does not necessarily increase. This is mainly because the buyer needs to account for the expected cost of the meta-CSA when querying all CSAs, which increases as the number of CSAs that need to be contacted increases. Still, in many settings, adding the meta-CSA makes things better for the seller as well. Figure 6 depicts the buyer's expected expense and the seller's expected benefit in equilibrium with self-interested CSAs and a meta-CSA and in the case where a buyer-operated CSA is used, for different numbers of CSAs (N). The curves resulting from the meta-CSA case are marked "(Meta)". The values of the other model parameters were set to: $c_{buyer \to meta} = c_{buyer \to seller} = c_{buyer \to csa} = 0.05$, $c_{meta \to csa} = c_{csa \to seller} = 0.001$ and $M = 0.016$. As expected, based on Proposition 5.2, the expected expense of the buyer with a meta-CSA increases as the number of CSAs available increases. Furthermore, in this case the seller's expected benefit also increases as N increases. In this specific example, for all N values, both buyers and sellers improve their performance with the meta-CSA compared to the buyer-operated case.

Figure 7 depicts the buyer's expected expense and the seller's expected benefit in equilibrium with self-interested CSAs, with and without a meta-CSA, and in the case where a buyer-operated CSA is used, as a function of M. The curves resulting from the

[8]This also derives from the fact that the CSAs' net benefit is zero, and therefore the aggregate search cost (over all CSAs) is covered by the payment M, resulting in an overall number of searches $M/c_{csa \to seller}$.

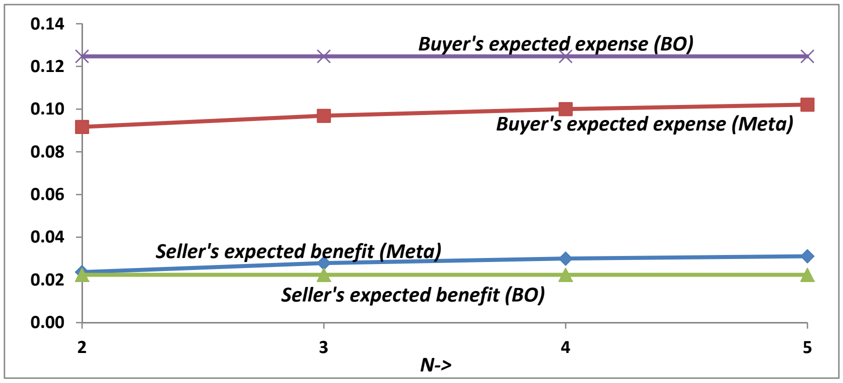

Fig. 6. The buyer's expected expense and the seller's expected benefit in equilibrium with self-interested CSAs, a meta-CSA and in the case where a buyer-operated CSA is used, for different numbers of CSAs (N). The parameters used are: $c_{buyer \to meta} = c_{buyer \to seller} = c_{buyer \to csa} = 0.05$, $c_{meta \to csa} = c_{csa \to seller} = 0.001$ and $M = 0.016$. Both the expected expense of the buyer and the expected benefit of the seller with a meta-CSA increase as the number of CSAs available increases.

meta-CSA case are marked "(Meta)". The values of the other model parameters were set to: $c_{buyer \to meta} = c_{buyer \to seller} = c_{buyer \to csa} = 0.01$, $c_{csa \to seller} = c_{meta \to csa} = 0.00012$ and $N = 3$. The mutually-beneficial equilibrium in this example, when not using a meta-CSA, occurs for any value of M in the interval $(0.00461 - 0.00472)$. The maximum expected benefit of the seller in a mutually-beneficial equilibrium is 0.0081 and the maximum sum of the improvements in the buyer's expected expense and the seller's expected benefit in a mutually-beneficial equilibrium is 0.0005. With the meta-CSA we obtain an expected maximum benefit of the seller in a mutually-beneficial type equilibrium of 0.0093 (improvement of 15% in comparison to the case without a meta-CSA) and the maximum sum of the improvements in the buyer's expected expense and the seller's expected benefit in a mutually-beneficial type equilibrium is 0.0015 (improvement of 300%). Therefore, despite the fact that the buyer now uses the meta-CSA to query all CSAs while only incurring the cost of querying the meta-CSA, and despite having the seller fully account for the meta-CSA's costs of querying all CSAs and paying M to the CSAs, the seller benefits in comparison to the case where the meta-CSA is not available to buyers.

6. PREVENTING MISUSE

One inherent threat to stability in markets with competing CSAs where CSAs do not charge buyers for their services is the possibility of CSA-misuse [Zhu and Madnick 2010; mySimon 2000]. Such a phenomena occurs when a CSA has an incentive to disguise itself as a buyer and request price comparison services from another CSA (for free), rather than search by itself [Zhu and Madnick 2010; Kephart et al. 2000]. Indeed, there are various security and authentication techniques that can be adopted for correlating any quote with the CSA which originally generated the query that yielded said quote. Yet sellers do not really care who originally queried them for the price quote and will be happy to make the payment M to any CSA that can direct a buyer to them who is willing to purchase at that price.

A possible search-theory based solution to the problem is to set a fixed fee c_{fee} for using the CSA. In this case, a CSA considering misusing another CSA will need to take into consideration the tradeoff between the additional gain in receiving the minimal quote returned by the other CSA and the cost c_{fee} of using that CSA. Similarly, each time a buyer contacts a CSA it will incur a cost c_{fee} (in addition to the intrinsic

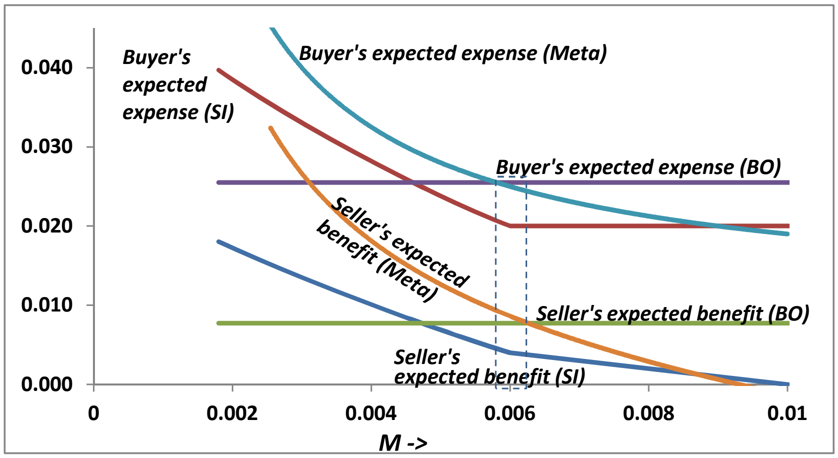

Fig. 7. The buyer's expected expense and the seller's expected benefit in equilibrium with self-interested CSAs, with and without a meta-CSA, and in the case where a buyer-operated CSA is used, as a function of M. The parameters used are: $c_{buyer \to meta} = c_{buyer \to seller} = c_{buyer \to csa} = 0.01$, $c_{csa \to seller} = c_{meta \to csa} = 0.00012$ and $N = 3$. The range of M values for which mutually-beneficial equilibrium exists is marked with the dotted rectangle.

cost $c_{buyer \to csa}$). In the following paragraphs, the equilibrium of a model where misuse is prevented by having CSAs set a fixed usage fee is analyzed. The parameter c_{fee} is therefore an additional decision variable of the CSA. We prove that with such misuse prevention a mutually beneficial solution of the form illustrated in the previous section does not hold. It is demonstrated, however, that the change in the sum of the improvements in the buyer's expected expense and the seller's expected benefit can be positive. Illustrations in this section are based on the same synthetic environment that was used in the former section.

Under the new circumstances, the buyer considers the cost c_{fee} as a search cost for receiving a quote q from a distribution $f_{returned}(q)$. The buyer thus sets her reservation value for querying the CSA using a modification of (11):

$$c_{buyer \to csa} + c_{fee} = \int_{q=0}^{r_{buyer \to csa}} (r_{buyer \to csa} - q) f_{returned}(q) dq \qquad (21)$$

Now consider the CSA's alternatives at any stage of its search process. First, it can terminate its search, returning the minimum price quote obtained up until that point (or null if none). Second, it can query a seller for a price quote with a cost $c_{csa \to seller}$, yielding a value from $f(q)$. Finally, it can query a random CSA that it had not misused so far for a cost $c_{fee} + c_{csa \to csa}$, yielding a price quote from $f_{returned}(q)$, where $c_{csa \to csa}$ denotes the cost that a CSA incurs when querying another CSA. It is assumed that $c_{csa \to csa} \leq c_{csa \to seller}$, for the same considerations that were used for justifying the relationship $c_{buyer \to csa} \leq c_{buyer \to seller}$.

Based on the optimal search principles that were discussed in Section 4, when choosing between querying a seller and querying another CSA the CSA needs to choose the alternative associated with the greater reservation value. We use the same mapping of revenue to price as before, hence the CSA will always choose the alternative associated with the lower reservation value. The reservation value of querying another CSA, denoted $r_{csa \to csa}$, can be calculated using a modification of (6) as follows:

$$c_{csa\to csa} + c_{fee} = M \int_{q=0}^{r_{csa\to csa}} (P(q) - P(r_{csa\to csa}))f(q)dq \tag{22}$$

The reservation value for querying another seller, $r_{csa\to seller}$, is calculated according to (6) as in Section 4. The CSA will choose to go with an additional price quote from a seller rather than misuse another CSA if $r_{csa\to seller} \le r_{csa\to csa}$ (if $r_{csa\to seller} = r_{csa\to csa}$, the CSA is indifferent to the two options, and by misusing another CSA it is actually delegating the task entirely to the other CSA, leaving a zero benefit to itself). Therefore, using c_{fee} that satisfies $r_{csa\to seller} = r_{csa\to csa}$ guarantees the prevention of CSA misuse. In this case the CSA will always prefer querying a seller, and since CSAs are not limited by the number of sellers they can query, they will always have an opportunity to obtain a price quote from a source with a reservation price lower than the reservation price of the other CSAs.

THEOREM 6.1. *A symmetric equilibrium in a misuse-enabled environment exists if and only if* $c_{fee} = c_{csa\to seller}/F(r_{csa\to seller}) - c_{csa\to csa}$ *and the following two conditions are satisfied:*

$$(1 - F_{returned}(r_{buyer\to csa}))(-M + \frac{MF(r_{csa\to seller})c_{fee}}{c_{csa\to seller}} + \frac{c_{csa\to seller}}{F(r_{csa\to seller})}) \tag{23}$$

$$+ F_{returned}(r_{buyer\to csa})(M - \frac{c_{csa\to seller}}{F(r_{csa\to seller})} + c_{fee}) < c_{fee}\frac{MF(r_{csa\to seller})}{Nc_{csa\to seller}} + \frac{M}{N}$$

$$\int_{q=r_{buyer\to csa}}^{r_{buyer\to seller}} f_{N-1}(q)(MF'_{returned}(q) + c'_{fee})dq < c_{fee}\frac{MF(r_{csa\to seller})}{Nc_{csa\to seller}} + \frac{M}{N} \tag{24}$$

where $f_{N-1}(q)$ *is the probability distribution function of the minimum quote among the* $N-1$ *quotes returned by* $N-1$ *CSAs, calculated as:* $f_{N-1}(q) = (N-1)f_{returned}(q)(1 - F_{returned}(q))^{N-2}$, $F'_{returned}(q)$ *is the probability that the CSA returns a quote lower than or equal to* q *if using a reservation value* $r_{buyer\to seller}$ *and* c'_{fee} *is the usage fee for which the buyer is indifferent between querying a seller and querying a CSA that charges this usage fee, calculated as:* $c'_{fee} = \int_{q=0}^{r_{buyer\to csa}} (r_{buyer\to seller} - q)f'_{returned}(q)dq - c_{buyer\to seller}$, *where* $f'_{returned}(q)dq$ *is the probability distribution function of values returned by the CSA if using a reservation value* $r_{buyer\to seller}$.

PROOF. An equilibrium in which all CSAs use a fee that enables misuse can never exist since it necessarily results in an infinite loop as each CSA will indefinitely pass the chore to another CSA, resulting in a loss for the CSAs. We first prove that the use of $c_{fee} \ge c^*_{fee} = c_{csa\to seller}/F(r_{csa\to seller}) - c_{csa\to csa}$ necessarily prevents misuse, and that the use of $c_{fee} < c^*_{fee}$ necessarily results in a misuse. Then we prove that if all CSAs use $c_{fee} > c^*_{fee}$, each individual CSA has an incentive to slightly decrease the usage fee it charges, thus this cannot be the equilibrium usage fee. Substituting $f_{returned}(q) = f(q)/F(r_{csa\to seller})$ and $c_{fee} = c^*_{fee} + \epsilon = c_{csa\to seller}/F(r_{csa\to seller}) - c_{csa\to csa} + \epsilon$ in (22) obtains:

$$c^*_{fee} + \epsilon = M \int_{q=0}^{r_{csa\to csa}} (P(q) - P(r_{csa\to csa}))f(q)dq \tag{25}$$

which is identical to (6), except for the ϵ element on the left-hand-side. Therefore, if $\epsilon \ge 0$ then $r_{csa\to csa} \ge r_{csa\to seller}$, hence preventing misuse. Otherwise, i.e., if $\epsilon < 0$, $r_{csa\to csa} < r_{csa\to seller}$ and the CSAs are used by others. This completes the first part of the proof.

If $c_{fee} > c^*_{fee}$ is in equilibrium, then the probability $P(q)$ is given by (12). Now consider a single CSA setting its usage fee to $c''_{fee} = c_{fee} - \epsilon$ where $\epsilon < c_{fee} - c^*_{fee}$, while using the same reservation price for querying sellers as the other CSAs, $r_{csa \to seller}$. None of the other CSAs will misuse this CSA, since $c''_{fee} > c^*_{fee}$ guarantees $r_{csa \to seller} < r_{csa \to csa}$ (according to (6) and (22)), hence misuse is prevented. The effect of using c''_{fee} on the CSA's expected benefit in this case is two-fold. On one hand, the CSA loses ϵ whenever contacted by a buyer. On the other hand, this CSA is now necessarily the first to be contacted by the buyer, thus $P(q) = 1$ for $q < r_{buyer \to seller}$ (compared to the expression given in (12)). This is because, from the buyer's point of view, contacting that CSA will result in obtaining a quote from the same distribution from whence the other CSAs' quotes are drawn ($f_{returned}(q)$), however paying a lower fee. The increase in the CSA's expected payment received from sellers in this case is positive, and does not depend on the value ϵ (since the CSA is still searching according to the same reservation value $r_{csa \to seller}$ as the other agents). Since the increase in the expected payment received from sellers is fixed, and the decrease in the payment received from buyers is ϵ, then there is necessarily a value ϵ small enough to make the net change in the CSA's expected benefit positive. Therefore, a c_{fee} different than $c^*_{fee} = c_{csa \to seller}/F(r_{csa \to seller}) - c_{csa \to csa}$ is necessarily not in equilibrium.

Equation 23 is necessary in order to guarantee that no individual CSA has an incentive to deviate to a usage fee $c_{fee} - \epsilon$, ($\epsilon \to 0$). In this case, the CSA will benefit from being the first to be queried by the buyer. If indeed the quote it returns is below $r_{buyer \to csa}$ (i.e., with probability $F_{returned}(r_{buyer \to csa})$), the buyer will terminate her search and the benefit of the CSA will be $(M - \frac{c_{csa \to seller}}{F(r_{csa \to seller})} + c_{fee})$ (where $\frac{c_{csa \to seller}}{F(r_{csa \to seller})}$ is the expected cost of search until obtaining a quote below $r_{csa \to seller}$). Otherwise (i.e., with probability $1 - F_{returned}(r_{buyer \to csa})$), the CSA will be misused by all other CSAs, therefore its expected expense for the search will be M (according to Corollary 4.6, including the cost due to the buyer's query). Its expected revenue in the latter case is the usage fee c_{fee} multiplied by the number of times the CSA is queried ($\frac{MF(r_{csa \to seller})}{c_{csa \to seller}}$, according to Corollary 4.6) and the payment M obtained if the other quotes it supplies to other CSAs when misused are greater than the quote that it returned to the buyer (i.e., with probability of $1/(\frac{MF(r_{csa \to seller})}{c_{csa \to seller}})$). The right-hand-side of the equation is the expected benefit if keeping the usage fee c_{fee} unchanged, in which case the CSA obtains a payment c_{fee} with probability $\frac{MF(r_{csa \to seller})}{Nc_{csa \to seller}}$ (the expected number of CSAs queried divided by the number of CSAs available to the buyer, N).

Equation 24 is necessary in order to guarantee that no CSA has an incentive to set a usage fee c''_{fee} greater than the one that the other CSAs use, however it is associated with a reservation value similar to $r_{buyer \to seller}$ so that the buyer is guaranteed to query that CSA before querying individual sellers. Since the new usage fee should guarantee $r_{buyer \to csa} = r_{buyer \to seller}$, its value can be calculated (according to (11)) as $c_{fee} = \int_{q=0}^{r_{buyer \to csa}} (r_{buyer \to seller} - q) f'_{returned}(q) dq - c_{buyer \to seller}$. Since the CSA uses a fee $c''_{fee} > c_{fee}$, it will be queried by the buyer only if all other CSAs returned a quote greater than $r_{buyer \to csa}$, and a payment M will be received only if returning a quote q lower than the minimum quote returned by the other $N - 1$ CSAs. □

At this point, we have everything that is necessary in order to calculate the equilibrium strategies of the buyer and CSAs, if such a c^*_{fee}-based equilibrium exists. The value of $r_{csa \to seller} = r_{csa \to csa}$, $r_{buyer \to csa}$ and $r_{buyer \to seller}$ can be calculated by solving the set of equations (21), (12), (10), (22) and $c_{fee} = c_{csa \to seller}/F(r_{csa \to seller}) - c_{csa \to csa}$.

The expected expense of the buyer is $E_{returned} + (c_{buyer\to csa} + c_{fee})E_{CSAs\_queried}$. The seller's expected benefit is: $E_{returned} - M$. The expected benefit of each CSA per buyer's search is given by: $c_{fee}\frac{MF(r_{csa\to seller})}{Nc_{csa\to seller}} + \frac{M}{N}$.

THEOREM 6.2. *An upper bound for the buyer's benefit improvement (if any) from using competitive self-interested CSAs rather than buyer-operated ones, when CSAs avoid misuse through the introduction of a usage fee* $c_{fee} = c_{csa\to seller}/F(r_{csa\to seller}) - c_{csa\to csa}$, *is* $c_{buyer\to csa}$.

PROOF. If there were no limitations over the number of CSAs that can be queried, then the buyer would have queried CSAs, in the self-interested-CSAs case, until receiving a quote below $r_{buyer\to csa}$. The buyer's expected expense in this case would have been:[9]

$$V_\infty(r_{buyer\to csa}) = \frac{\int_{y=0}^{r_{buyer\to csa}} y f_{returned}(y)dy + c_{fee} + c_{buyer\to csa}}{F_{returned}(r_{buyer\to csa})} \tag{26}$$

Obviously the expected expense when using the same $r_{buyer\to csa}$ strategy, having only N agents, is greater than $V_\infty(r_{buyer\to csa})$, thus (26) gives a lower bound to the buyer's expected expense when CSAs charge c_{fee}. Substituting $c_{fee} = c_{csa\to seller}/F(r_{csa\to seller}) - c_{csa\to csa}$, $f_{returned}(y) = f(y)/F(r_{csa\to seller})$ and $F_{returned}(r_{buyer\to csa}) = F(r_{buyer\to csa})/F(r_{csa\to seller})$ in (26), obtains:

$$\begin{aligned} V_\infty(r_{buyer\to csa}) &= \frac{\int_{y=0}^{r_{buyer\to csa}} \frac{yf(y)}{F(r_{csa\to seller})}dy + \frac{c_{csa\to seller}}{F(r_{csa\to seller})} - c_{csa\to csa} + c_{buyer\to csa}}{\frac{F(r_{buyer\to csa})}{F(r_{csa\to seller})}} \\ &> \frac{\int_{y=0}^{r_{buyer\to csa}} yf(y)dy + c_{csa\to seller}}{F(r_{buyer\to csa})} \end{aligned} \tag{27}$$

where the inequality derives from the fact that $c_{csa\to csa} < c_{buyer\to csa}$. Now notice that the last expression is exactly the expected cost of a buyer-operated sequential search according to 4.2, if using reservation value $r_{buyer\to csa}$, excluding the cost $c_{buyer\to csa}$. Therefore, the improvement can be at most $c_{buyer\to csa}$ (if any) . □

Theorem 6.2 has an intuitive interpretation: operating the CSA by herself, the buyer can emulate a search from a distribution $f_{returned}(q)$ simply by re-sampling a new seller whenever receiving a quote greater than $r_{csa\to seller}$. The expected cost incurred between any subsequent receipts of any two such price quotes is $c_{csa\to seller}/F(r_{csa\to seller})$, which is less than the cost of using the CSAs for this purpose. Therefore the difference between the two cases is at most the cost of accessing the CSA.

Figure 8 illustrates the buyer's and seller's performance in the misuse prevention scenario. The left graph in the figure depicts the expected benefit of sellers and the expected expense of buyers with self-interested CSAs that use c_{fee} according to Theorem 6.1 and with buyer-operated CSAs, for different values of M. The values of the other model parameters were set to: $c_{buyer\to csa} = c_{buyer\to seller} = 0.02$, $c_{csa\to seller} = c_{csa\to csa} = 0.00014$ and $N = 2$. The right graph depicts the sum of differences between the buyer's and seller's performance with self-interested CSAs compared to their performance when using buyer-operated CSAs. As observed in the figure,

[9]This is the equivalent of (4), allowing the buyer to request as many quotes as needed, sequentially, from distribution $f_{returned}(y)$, incurring a cost $c_{buyer\to csa} + c_{fee}$ each time.

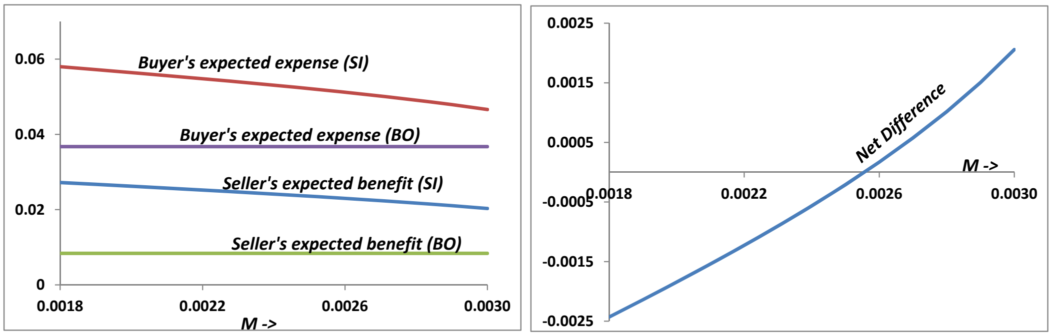

Fig. 8. The expected benefit of sellers and the expected expense of buyers with self-interested CSAs that use c_{fee} according to Theorem 6.1 and with buyer-operated CSAs (left), and the sum of differences between a buyer's and a seller's performance with self-interested CSAs compared to their performance with buyer-operated CSAs (right), for different values of M. The parameters used are: $c_{buyer\rightarrow csa} = c_{buyer\rightarrow seller} = 0.02$, $c_{csa\rightarrow seller} = c_{csa\rightarrow csa} = 0.00014$ and $N = 2$. For $M > 0.0026$, the sum of differences is positive.

for $M > 0.0026$ the sum of the differences is positive, and therefore if an appropriate side-payment mechanism is available, a mutually-beneficial solution can be devised.

7. DISCUSSION AND CONCLUSIONS

The main innovation of the analysis given in this paper is in taking CSAs to be self-interested agents that aim to maximize their own benefit. As such, the CSAs' search strategy does not always align with the best interests of the buyer or the seller. The setting being used, in which buyers can use several different CSAs that offer the comparison-shopping service for free while building on sellers' payments, is in fact the one seen in real-life electronic markets nowadays. The analysis, which is entirely based on search theory (which power of modeling markets is well established in the literature [McMillan and Rothschild 1994; Morgan and Manning 1985; Smith 2011, inter alia]), provides an understanding of the unique dynamics according to which buyers and CSAs set their strategies, and the resulting equilibrium. It is notable that the model does not take sellers to be strategic players, and both the payment to the CSA and the distribution of prices in the market are assumed to be constant and exogenous. This exogeneity may be justified in many ways. For example, the payment that sellers need to pay a CSA upon transaction usually depends on product category [Moraga-Gonzalez and Wildenbeest 2011] rather than used as a decision variable in the model. Similarly, sellers often set their prices taking into account other markets in which they are active (i.e., multichannel retailers [Xing et al. 2006]) and/or change their prices as a result of external factors (as in the case of airfares where price changes according to seat availability, fuel prices, cancellations, etc.). Still, incorporating sellers' responses to the market dynamics as part of the equilibrium analysis can certainly account for important and intriguing results that are not covered by the current analysis. In fact, there is no guarantee that the solution characteristics described will still hold if the sellers compete on pricing or commission. This, however, does not detracts from the importance of the cohesive equilibrium analysis given in the paper and the results revealed through it.

The implications of the analysis to marketplace platforms owners are numerous. First, and most important, it is shown that in some cases the equilibrium set of CSAs' strategies results in an improvement in the performance of both buyers and sellers (in comparison to the buyer-operated case), while the CSAs' net benefit is non-negative. This is explained by the tradeoff of buyers between the savings in search costs and the resulting inefficiency of search through self-interested CSAs. In a way, the competing CSAs are used in this case as a means for transferring some compensation from sellers to buyers in return for reducing the extent of their search. Any direct transfer of such compensation is likely to fail, since the buyers' search extent increases as the cost of search decreases. The fact that sellers can actually benefit from the introduction of competing CSAs in the market is in contrast to the case where CSAs are assumed to be fully in the service of buyers. In the latter case, it has been shown both analytically and experimentally that the benefits of price search services fall primarily to the buyers, and sellers in such settings are subject to a low profit per unit [Gorman et al. 2009; Nermuth et al. 2009].

The above mutually beneficial solution is not a product of the competition per se. As demonstrated in Section 5, the phenomena can occur also in settings where there is only one self-interested CSA available in the market. Yet, the occurrence of such a result is rare in comparison to when there are several competing CSAs. Moreover, it is shown that even if the seller further enhances the competition, by fully covering the expenses of a meta-CSA that enables querying all CSAs with no additional cost to buyers, her expected benefit in comparison to the case where buyers incur a cost for each CSA they query individually may improve. Basically, in the absence of competition dynamics, a single self-interested CSA would be able to take over much of the buyer's surplus. This is generally avoided (or substantially reduced) when there are competing CSAs. Furthermore, the buyer's expected expense in equilibrium with a single self-interested CSA is always worse (i.e., greater) than her expected expense with several competing CSAs.

Two additional interesting characteristics of the equilibrium solution relate to buyers. The first concerns querying sellers directly. The characterization of the equilibrium suggests that buyers will first query CSAs and will never complement their search with querying sellers directly. One phenomena that is demonstrated in this context is that an increase in the buyer's intrinsic cost of querying sellers can actually have a positive effect on the buyer's expected expense, as it may enable a favorable equilibrium that could not have held otherwise (see Section 5). This is because the option to query sellers directly and the resulting reservation value that the buyer uses for deciding on querying sellers affect the existence of equilibrium. A relatively low cost of querying sellers may preclude the existence of equilibrium. It is notable that if given the option to continue querying sellers directly with a buyer-operated CSA (rather than based on the intrinsic cost of querying sellers directly), the expected quote with which the buyer ends up after using the self-interested CSAs is necessarily lower than when using only a buyer-operated CSA (Subsection 4.6). Therefore the seller in this case loses both because of the reduced purchased price and by the payment to the CSA, and a mutually-beneficial type equilibrium cannot exist. The decision to use a buyer-operated CSA should therefore be strategic from the buyers' point of view, since once sellers realize that buyers use a buyer-operated CSA, they will cease to offer the self-interested CSAs a payment and the buyers will end up losing.

The second unique equilibrium characteristic that concerns buyers is the effect of CSAs' competition over buyers' expected expense. One would expect a decrease in the expected expense as the number of competing CSAs in the market increases (i.e., when the competition rises). Alas, in the model analyzed, an increase in the number of competing CSAs typically results in an increase in the buyer agents' expected expense. This, as explained in the former sections, is the result of the decreased efficiency of the overall search.

The fact that buyers do not necessarily gain from the existence of self-interested competing CSAs suggests that market owners should not prefer to include those in the market as a default (over supplying buyers a cooperative CSA that charges only its operational (querying) cost). Instead, the usefulness of each alternative should be carefully evaluated, based on the principles given in the analysis section. In fact, if a market owner can control either the number of CSAs in the market or dictate the payments they receive from sellers, it can substantially improve the performance of the market, individual-wise. Furthermore, even in the case where one of the sides does not benefit from the use of self-interested CSAs, in comparison to buyer-operated CSAs, mutually-beneficial solutions can be artificially constructed whenever the sum of the changes of the two sides is positive and a side-payment mechanism is available. Control over the number of CSAs can be achieved either by offering them a fixed incentive, as discussed in Section 5, or through licensing. The control over seller payments to CSAs is a bit more difficult to achieve, though it can be thought of as a condition for allowing the seller to operate in the market. A side-payment mechanism can be implemented in many forms, such as rebate from sellers if purchasing through a CSA, and taxation or a subscription fee on buyers (if applied by the market owner).

It is notable that whenever market-designer/owner interference is allowed, the first-best is to dictate a single search, either by the consumer herself or via a single search agency. In the same spirit, one can consider a solution where the seller is forced to sell the product at a pre-defined price, thus entirely eliminating the need for search, and then split the surplus between buyer and seller. One major disadvantage of such a solution is that in order to make it a mutually-beneficial type solution, one needs to establish a mechanism for side-payments. Other disadvantages of such a solution are: (a) it requires far more interference from the planner/designer's side than the self-interested-CSAs solution; and (b) it results in a fixed price (after the transfer of the side-payments) which might be a problem, given that the sellers' prices are determined externally and the same prices that potentially hold in the parallel channels should also hold in this market.

The analysis methodology is also used for investigating a solution to the CSA misuse problem, which is one of the most inherent problems of CSA technology. Section 6 discusses a solution in the form of a fixed usage fee CSAs set in order to guarantee misuse prevention. It is demonstrated that even in this case a wise selection of the buyer payment and the number of CSAs can result in an increase in the sum of changes in the buyer's expected expense and the seller's expected benefit in comparison to the buyer-operated case.

The analysis and results can be generalized for other markets, where the buyers' search involves a thorough evaluation of opportunities rather than merely obtaining a price quote. Examples of such markets are `autotrader.com` and `Yet2.com`. In these examples, the buyer needs to query sellers for supplementary information regarding the opportunities they list (used cars in the case of Autotrader, inventions and patents in the case of Yet2)[10] in order to reason about their actual value. In such settings it is common to find professionals that can query sellers on behalf of the buyer and evaluate the opportunities listed in a much more efficient fashion (i.e., with less effort, consuming less resources). For example, a professional mechanic may be able to better evaluate the benefit to the buyer in each car listed on Autotrader. Similarly, a professional with a premium subscription to independent agencies that monitor the recorded history of cars, such as `Carfax.com`, will need to pay less for retrieving a specific car's history. In the case of markets for technology, one may think of professionals with expertise in areas that complement the buyer's main business. Such experts can some-

[10]The Yet2 marketplace operates as an online platform that allows "sellers" (industrial firms, entrepreneurial ventures, research universities and individual inventors) to post their inventions for a fee, while "buyers" can register free of cost, search the listed inventions and engage in an exchange [Dushnitsky and Klueter 2011].

times offer a more reliable evaluation of the benefit to the company from investing in any specific technology or from purchasing specific patents offered in these markets. In both examples, the same question of whether to rely on a free seller-sponsored self-interested agent or hire a costly agent that will conduct the search according to the buyer's best interests may arise. This question can apply directly to buyers and to the marketplace (platform) owner/operator whenever licensing is an issue or technological barriers prevent such mediators from gaining direct access to sellers' information.

Further research, following the idea of markets with competing CSAs, is likely to address more complicated schemes of payments offered by sellers (e.g., commission in terms of a percentage of the transaction cost), possible heterogeneity in search characteristics (e.g., different querying costs for different sellers) and other misuse prevention methods (e.g., charge per query). Also, the model considers an intrinsic buyer's search cost, thus much of the analysis can be used if a model where CSAs charge buyers a usage fee is to be considered. Yet, a complete analysis of an equilibrium where CSAs also get to control the level of the usage fee they charge buyers is beyond the scope of the current paper and thus suggested for future research. Finally, further research for developing mechanisms for transferring some of the buyer's surplus to the seller (and vice versa) is greatly desired. Such mechanisms can be of much use whenever only one side benefits from the introduction of self-interested CSAs in the market and the benefit of that party is greater than the loss of the other. Such research will certainly benefit from the analysis given in this paper.

Acknowledgments

The equilibrium analysis of the sequential search model as well as the potential implication to market design, for a degenerate model of this kind, appeared in the proceedings of IAT'09. Many thanks to the three anonymous reviewers for their valuable comments and priceless suggestions to improve this paper.

REFERENCES

ARMSTRONG, M. 2006. Competition in two-sided markets. *RAND Journal of Economics 37,* 3, 668–691.

BAKOS, J. Y. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science 42(12)*, 1676–1692.

BAYE, M. R. AND MORGAN, J. 2001. Information gatekeepers on the internet and the competitiveness of homogeneous product markets. *American Economic Review 91,* 3, 454–474.

BAYE, M. R., MORGAN, J., AND SCHOLTEN, P. 2004. Temporal price dispersion: Evidence from an online consumer electronics market. *Journal of Interactive Marketing 18(4)*, 101–115.

BAYE, M. R., MORGAN, J., AND SCHOLTEN, P. 2006. Persistent price dispersion in online markets. In *The New Economy and Beyond*, D. Jansen, Ed. Edward Elgar Press, Northampton, MA.

BISWAS, S. AND NARAHARI, Y. 2001. Analysis of supplier competition in electronic marketplaces. In *Proceedings of the International Conference on Automation, Energy, and Information Technology (EAIT-2001)*. 1–14.

BRYNJOLFSSON, E., HU, Y. J., AND SMITH, M. D. 2003. Consumer surplus in the digital economy: Estimating the value of increased product variety at online bookseller. *Management Science 49(11)*, 1580–1596.

CHOI, S. P. M. AND LIU, J. 2000. Optimal time-constrained trading strategies for autonomous agents. In *Proceedings of the International ICSC Symposium on Multi-agents and Mobile Agents in Virtual Organizations and E-commerce (MAMA'2000)*.

CLAY, K., KRISHNAN, R., WOLFF, E., AND FERNANDES, D. 2002. Retail strategies on the web: Price and non-price competition in the online book industry. *Journal of Industrial Economics 50*, 351–367.

CLEMONS, E. K., HITT, L. M., AND HORN, I. H. 2002. Price dispersion and differentiation in online travel: An empirical investigation. *Management Science 48*, 534–550.

DECKER, K., SYCARA, K., AND WILLIAMSON, M. 1997. Middle-agents for the internet. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*. 578–583.

DOORENBOS, B., ETZIONI, O., AND WELD, D. S. 1997. A scalable comparison-shopping agent for the world-wide web. In *Proceedings of the First International Conference on Autonomous Agents (Agents-1997)*. 39–48.

DUSHNITSKY, G. AND KLUETER, T. 2011. Is there an ebay for ideas? insights from online knowledge marketplaces. *European Management Review 8*, 17–32.

ETZIONI, O. 1997. Moving up the information food chain: Deploying softbots on the world wide web. *AI Magazine 18,* 2, 11–18.

ETZIONI, O., TUCHINDA, R., KNOBLOCK, C. A., AND YATES, A. 2003. To buy or not to buy: Mining airfare data to minimize ticket purchase price. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '03)*. 119–128.

FASLI, M. 2006. Shopbots: A syntactic present, a semantic future. *IEEE Internet Computing 10*, 69–75.

FERGUSON, T. S. 1989. Who solved the secretary problem? *Statistical Science 4,* 3, 282–289.

GAL, S., LANDSBERGER, M., AND LEVYKSON, B. 1981. A compound strategy for search in the labor market. *International Economic Review 22(3)*, 597–608.

GORMAN, M. F., SALISBURY, D., AND BRANNON, I. 2009. Who wins when price information is more ubiquitous? an experiment to assess how infomediaries influence price. *Electronic Markets 19,* 2-3, 151–162.

GREENWALD, A. R., KEPHART, J. O., AND TESAURO, G. J. 1999. Strategic pricebot dynamics. In *Proceedings of the First ACM Conference on Electronic Commerce (EC '99)*. 58–67.

GUTTMAN, R. H., MOUKAS, A. G., AND MAES, P. 1998. Agent-mediated electronic commerce: A survey. *Knowledge Engineering Review 13(2)*, 147–159.

HAYNES, M. AND THOMPSON, S. 2008. Price, price dispersion and number of sellers at a low entry cost shopbot. *International Journal of Industrial Organization 26,* 2, 459–472.

HAZON, N., AUMANN, Y., AND KRAUS, S. 2009. Collaborative multi agent physical search with probabilistic knowledge. In *Proceedings of the Twenty-First international Joint Conference on Artificial intelligence (IJCAI'09)*. 167–174.

HE, M., JENNINGS, N. R., AND LEUNG, H. 2003. On agent-mediated electronic commerce. *IEEE Transaction on Knowledge and Data Engineering 15(4)*, 985–1003.

HUANG, J. AND KAZEYKINA, A. 2010. Optimal strategies for reviewing search results. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2010)*.

JANSSEN, M. C. W. AND MORAGA-GONZALEZ, J. L. 2004. Strategic pricing, consumer search and the number of firms. *Review of Economic Studies 71,* 4, 1089–1118.

JANSSEN, M. C. W., MORAGA-GONZALEZ, J. L., AND WILDENBEEST, M. R. 2005. Truly costly sequential search and oligopolistic pricing. *International Journal of Industrial Organization 23,* 5-6, 451–466.

JOHNSON, E. J., MOE, W. W., FADER, P. S., BELLMAN, S., AND LOHSE, G. L. 2004. On the depth and dynamics of online search behavior. *Management Science 50,* 3, 299–308.

JOPSON, B. 2012. Amazon 'robo-pricing' sparks fears. Financial Times, July 8, 2012.

JUMADINOVA, J. AND DASGUPTA, P. 2008. Firefly-inspired synchronization for improved dynamic pricing in online markets. In *Proceedings of the Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2008)*. 403–412.

KARAT, C. M., BLOM, J. O., AND KARAT, J. 2004. *Designing personalized user experiences in eCommerce*. Human-Computer Interaction Series. Kluwer Academic.

KEPHART, J. O. AND GREENWALD, A. R. 2002. Shopbot economics. *Autonomous Agents and Multi-Agent Systems 5(3)*, 255–287.

KEPHART, J. O., HANSON, J. E., AND GREENWALD, A. R. 2000. Dynamic pricing by software agents. *Computer Networks 32*, 731–752.

KRULWICH, B. 1996. The BargainFinder agent: Comparison price shopping on the Internet. In *Bots and Other Internet Beasties*, J. Williams, Ed. Sams Publishing, Chapter 13, 257–263.

LANDSBERGER, M. AND PELED, D. 1977. Duration of offers, price structure, and the gain from search. *Journal of Economic Theory 16,* 1, 17–37.

LIPPMAN, S. A. AND MCCALL, J. J. 1976. The economics of job search: A survey. *Economic Inquiry 14,* 3, 155–189.

MAES, P. 1994. Agents that reduce work and information overload. *Communications of the ACM 37,* 7, 30–40.

MAES, P., GUTTMAN, R. H., AND MOUKAS, A. G. 1999. Agents that buy and sell. *Communications of the ACM 42,* 3, 81–91.

MARKOPOULOS, P. M. AND UNGAR, L. H. 2002. Shopbots and pricebots in electronic service markets. In *Game Theory and Decision Theory in Agent-Based Systems*, S. Parsons, P. Gmytrasiewicz,

M. Wooldridge, and G. Weiss, Eds. Multiagent Systems, Artificial Societies, and Simulated Organizations Series, vol. 5. Springer US, 177–195.

MCMILLAN, J. AND ROTHSCHILD, M. 1994. Search. In *Handbook of Game Theory with Economic Applications*, R. Aumann and S. Hart, Eds. Chapter 27, 905–927.

MONTGOMERY, A. L., HOSANAGAR, K., KRISHNAN, R., AND CLAY, K. B. 2004. Designing a better shopbot. *Management Science 50*, 189–206.

MORAGA-GONZALEZ, J. L. AND WILDENBEEST, M. 2011. Comparison sites. Working Papers 2011-04, Indiana University, Kelley School of Business, Department of Business Economics and Public Policy. May.

MORGAN, P. AND MANNING, R. 1985. Optimal search. *Econometrica 53(4)*, 923–944.

MYSIMON. 2000. mySimon, Inc. v. Priceman, LLC. No. 5:99cv20939. N.D. Cal., February 29, 2000.

NERMUTH, M., PASINI, G., PIN, P., AND WEIDENHOLZER, S. 2009. Price dispersion, search externalities, and the digital divide. Vienna Economics Papers 0916, University of Vienna, Department of Economics. Dec.

PATHAK, B. 2010. A survey of the comparison shopping agent-based decision support systems. *Journal of Electronic Commerce Research 11,* 3, 178–192.

ROTHSCHILD, M. 1974. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy 82*, 689–711.

RYSMAN, M. 2009. The economics of two-sided markets. *Journal of Economic Perspectives 23,* 3, 125–43.

SARNE, D., KRAUS, S., AND ITO, T. 2007. Scaling-up shopbots: a dynamic allocation-based approach. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS '07)*. 256–263.

SARNE, D., MANISTERSKI, E., AND KRAUS, S. 2010. Multi-goal economic search using dynamic search structures. *Autonomous Agents and Multi-Agent Systems 21*, 204–236.

SMITH, L. 2011. Frictional matching models. *Annual Reviews in Economics 3*, 319–338.

SMITH, M. D. 2002. The impact of shopbots on electronic markets. *Journal of the Academy of Marketing Science 30,* 4, 442–450.

STAHL, D. O. 1989. Oligopolistic pricing with sequential consumer search. *American Economic Review 79,* 4, 700–712.

STIGLER, G. J. 1961. The economics of information. *Journal of Political Economy 69(3)*, 213–225.

TAN, C. H., GOH, K. Y., AND TEO, H. H. 2010. Effects of comparison shopping websites on market performance: Does market structure matter? *Journal of Electronic Commerce Research 11,* 3, 193–219.

TANG, Z., SMITH, M. D., AND MONTGOMERY, A. 2010. The impact of shopbot use on prices and price dispersion: Evidence from online book retailing. *International Journal of Industrial Organization 28,* 6, 579–590.

VARIAN, H. R. 1980. A model of sales. *American Economic Review 70,* 4, 651–59.

WALDECK, R. 2008. Search and price competition. *Journal of Economic Behavior and Organization 66,* 2, 347–357.

WAN, Y. AND LIU, Y. 2009. The impact of legal challenges on the evolution of web-based intelligent agents. *Journal of International Commercial Law and Technology 3,* 2.

WAN, Y., MENON, S., AND RAMAPRASAD, A. 2009. The paradoxical nature of electronic decision aids on comparison-shopping: The experiments and analysis. *Journal of Theoretical and Applied Electronic Commerce Research*, 80–96.

WAN, Y. AND PENG, G. 2010. What's next for shopbots? *IEEE Computer 43*, 20–26.

WEITZMAN, M. L. 1979. Optimal search for the best alternative. *Econometrica 47,* 3, 641–54.

XING, X., YANG, Z., AND TANG, F. 2006. A comparison of time-varying online price and price dispersion between multichannel and dotcom dvd retailers. *Journal of Interactive Marketing 20,* 2, 3–20.

YUAN, S. T. 2003. A personalized and integrative comparison-shopping engine and its applications. *Decision Support Systems 34,* 2, 139–156.

ZHU, H. AND MADNICK, S. E. 2010. Legal challenges and strategies for comparison shopping and data reuse. *Journal of Electronic Commerce Research 11,* 3, 231–239.

Appendix - Summary of Notations

symbol	explanation
N	number of competing CSAs in the market.
M	seller's payment to CSA upon purchase of a buyer directed to it by the CSA.
$f(q)$	distribution of prices offered by sellers.
$F(q)$	probability that a random query yields a price lower than q.
$c_{csa \rightarrow seller}$	the cost that the CSA incurs whenever querying a seller.
$c_{buyer \rightarrow csa}$	the cost that the buyer incurs whenever querying a CSA.
$c_{buyer \rightarrow seller}$	the cost that the buyer incurs whenever querying a seller directly.
$R_i, c_i, f_i(x)$	the reservation value; cost of query; and distribution of values of a given opportunity that is available to the searcher.
$R, c, f(x)$	the reservation value; cost of query; and distribution of values in the homogeneous case.
$V(R)$	expected benefit when using reservation value R.
q	the quote returned by the seller if queried.
$P(q)$	the probability that the buyer will actually buy the product if receiving price quote q from the CSA (i.e., that the seller will not obtain a better price quote along her search).
$r_{csa \rightarrow seller}$	the reservation value used by the CSA for deciding whether to query an additional seller.
$r_{buyer \rightarrow csa}$	the reservation value used by the buyer for deciding whether to query a CSA.
$r_{buyer \rightarrow seller}$	the reservation value used by the buyer for deciding whether to query an additional seller.
$f_{returned}(q)$	the probability the CSA returns a quote q if queried.
$F_{returned}(q)$	the probability the CSA returns a quote equal to or lesser than q if queried.
$P_{queried}(i)$	the probability that a queried CSA is the $i-th$ queried CSA in the buyer's search.
V_{csa}	the expected benefit of the CSA whenever contacted by the buyer.
$E_{sellers\_queried}$	the expected overall number of sellers queried by the CSAs as a result of the buyer's querying of CSAs.
$E_{returned}$	the expected price quote with which the buyer eventually ends up in equilibrium.
$E[x \mid x < q]$	the expected quote given that it is below q.
$E_N[X \mid X > q]$	the expected value of the minimum quote out of a sample of N quotes, given that all N quotes are above q (and necessarily below $r_{csa \rightarrow seller}$).
$f_N(y \mid y > q)$	the probability distribution function of the minimum value in a sample of N observations taken from a distribution $f_{returned}(q)$, given that all observations are above q
$E_{CSAs\_queried}$	the expected number of CSAs that the buyer samples.
$c_{buyer \rightarrow meta}$	the cost that the buyer incurs whenever querying the meta-CSA.
$r_{buyer \rightarrow meta}$	the reservation value used by the buyer for deciding whether to query the meta-CSA.

$c_{meta \to csa}$	the cost that the meta-CSA incurs whenever querying a CSA.
$E'_{returned}$	the expected value of the minimum quote returned by the N CSAs in the meta-CSA model.
c_{fee}	the usage fee set by CSA (i.e., the fee that users need to pay in order to query the CSA) for misuse prevention.
$c_{csa \to csa}$	the cost that the CSA incurs whenever querying a nother CSA.
$r_{csa \to csa}$	the reservation value used by the CSA for deciding whether to query another CSA.
$f_{N-1}(q), F_{returned}(q)'$	the probability distribution function of the minimum quote among the $N-1$ quotes returned by $N-1$ CSAs; the probability that the CSA returns a quote less than or equal to q if using a reservation value $r_{buyer \to seller}$.
c'_{fee}	the usage fee for which the buyer is indifferent between querying a seller and querying a CSA that charges this usage fee.
c''_{fee}	a usage fee to which the CSA may deviate in the case of misuse.