Essays on Platforms: Asymmetric Information, Search, and Policy

Essays on Platforms:
Asymmetric Information, Search, and Policy

A dissertation presented

by

Albert Zhao Wang

to

The Department of Economics

in partial fulfilment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

October 2012

Dissertation Advisors:
Susan Athey
Gregory Lewis

Author:
Albert Zhao Wang

Essays on Platforms: Asymmetric Information, Search, and Policy

Abstract

The three essays of this thesis cover two sets of topics: search in auction platforms in the first two papers, and political campaigning in the last. In platform settings, search cost reductions are often regarded as beneficial because they improve match quality. But is this in fact true? And if it is true in an aggregate sense, what are the consequences to individual platform participants? Do individual buyers and sellers win or lose? The first paper develops a novel model of search in platforms and applies it to auction platforms to test the popular hypothesis that lower search costs are always beneficial to sellers. Under certain assumptions, we find that while lower search costs is welfare improving, its distributional consequences are less predictable. In general, lower search costs intesify buyer-side competition. On the one hand, this tends to improve seller revenues due to better matches; on the other hand, this may also thin out markets for certain sellers, since lower search costs make it easier for buyers to search *out* of certain markets. Generally, some sellers gain and some lose; most surprisingly, however, we find that *overall* seller revenue can decrease with lower search costs. Our second paper extends the model to endogenize buyer participation — so some buyers may leave the platform completely — and considers optimal platform search policy in such settings. Under stricter assumptions, we find that a platform that taxes the seller side generally benefits from lower search costs; a platform that charges buyers, however, may maximize search costs, since the gains from easier search are unevenly distributed among buyers, and may be inefficiently extracted with a fee. The final essay provides a novel model of political campaigning as argumentation, which brings together two different strands of the campaign spending literature: spending has direct effects on electoral outcomes, but also provide a "signal" of candidate quality. The model parsimoniously resolves many pre-existing campaign spending "paradoxes" while delivering new results on the effects and desirability of spending caps.

Contents

List of Figures

Acknowledgments

This work owes its completion to the support and efforts of many people, upon whom I have depended for advice, wisdom, insights, and, perhaps most of all, endurance. A student could not hope to benefit from a committee as dedicated and helpful as the one I have enjoyed: I express my sincerest thanks to Susan Athey, whose support and crisp advice have been invaluable to me; to Greg Lewis, who oversaw my work on a weekly basis and bore the burden of reading many first drafts; to Jeffry Frieden, upon whom I have been able to rely for just about everything. A special thanks to my friends, who have kept me sane, provided feedback, and stubbornly continue to extend me invitations in spite of my perpetual absence: Mauricio, James, Niall, John, and Ani. To my mom and dad, for being inspiring examples, and for everything that they have done, which they can never be thanked enough for. To our cat Bernadette, for keeping me company and singing to me on late nights. And finally, to my loving wife Ruxandra, for her infinite patience and unwavering support, and, above all, for making it all worthwhile.

Introduction

The first two essays presented in this work study search effects in auction platforms. Although the properties of individual auctions – covering a wide variety of different mechanisms – have been extensively analyzed, the process by which potential buyers initially find themselves bidding auctions themselves has received relatively little attention. Given the historical context of early work in auction theory, this makes sense: its main applications have been centered around engineering sales of relatively unique goods with potentially high, but buyer-specific value. Since then, however, auctions have proven to be an increasingly popular way to allocate a large variety of goods, from used cars to search keywords; more recently, three of the largest technology companies – Google, Microsoft, and Yahoo – have established large display ad exchanges. In contrast to the high-profile auctions around which early auction theory developed, auctions today are commonplace, most of them completely invisible: almost every time a website is loaded, an auction is held in the background to allocate a singular, transient advertisement spot. This increase in quantity has also led to a change in organization: they are organized via centralized exchanges, or platforms.

When large numbers of simultaneous auctions are organized in platforms, the allocation problem becomes more complex: the process by which bidders find the auctions they participate in is just as important as the process by which those individual auctions are conducted. Moreover, auctions are most useful for goods with specific values and thin markets, characteristics that amplify the importance of high quality matches: the better matches a platform can make, the higher will be prices and surpluses. The first essay presents a first rigorous look at the effects of this allocation process: it develops a model of auction platforms where buyers find auctions through sequential search and applies this model to examine the effects of reducing search costs on platform outcomes. One main contribution of this essay is methodological: the search framework gives a widely applicable tractable treatment of sequential search into many-to-one matches. This significantly expands the applicability of sequential search models, which have previously been restricted to one-to-one settings. In auction platforms, however, each buyer's utility depends not only on the seller he is matched to, but also the identities of all other buyers against whom he is competing. A tractable many-to-one search model may also deliver additional insights into other phenomena, such as job search, but the present essays focus on auction platforms. The first essay provide theoretical support for the idea that search generates additional surplus, it turns out that overall welfare does improve as search costs decrease. At the same time, however, we show that the distribution of gains is uncertain; most surprisingly, in some cases, overall seller revenue can actually decline with reductions in search costs.

The second essay is a more applied work: armed with a model of agent behavior and outcomes in auction platforms, we ask what a platform's optimal search policy is. Interest in search effects has been motivated by the intuition that search generates extractable surplus — hence, a platform operator should stand to benefit from investments that reduce search costs. But as we saw in the first essay, search has significant distributional consequences, which complicates the question. Although search may increase total surplus, whether or not a platform operator will desire lower search costs will depend on its ability to extract that surplus. We establish results on optimal platform policy for two types of platforms, one that taxes sales, or total seller revenues, and one that charges buyers to participate. The general findings here reinforce the concepts of the first essay: there are significant gains to be had from producing better matches, but whether or not a platform will want to reduce costs depends on its instrument; in particular, a fee-charging platform will sometimes want to make search costs as high as possible, i.e., to maximize the search frictions. The reason for this is that the gains from search are often selectively enjoyed by buyers with extremal valuations and the sellers they are matched to; whether or not these gains are transferred to platform operators depends on how responsive the instrument is to those particular matches. A flat membership fee, in particular, can be levied only at the level of the lowest-utility buyer, and so a platform that derives its revenue primarily from fees may fail to reap the benefits of reduced search costs.

The final essay shifts focus from economic to political mechanisms and takes a close look at one particular aspect of political competition: campaigning. The recent 2012 US election cycle, preceded by the controversial 2010 Supreme Court decision on *Citizens United*, has renewed public debates on the role of money in politics; in spite of spirited voices on bot sides, there is no consensus on how money actually affects political outcomes. Part of the confusion has come from empirical studies of campaign spending effects, which have consistently found very low effects of marginal incumbent spending; hence, an argument goes that there is no need to worry about incumbents being able to steal elections with easier access to contributors. Our essay seeks to clarify the role of spending in election outcomes by developing an all-pay model of political campaigns. Although all-pay contest are a popular way to model many types of competition, they have been strangely absent from analyses of political campaigns, in spite of the fact that they share many structural properties with all-pay contests. Such a framework is also able to bring together two sides of the political contest literature: spending can have *direct* effects, while also serving as a noisy signal of candidate quality. The paper demonstrates how a very basic all-pay model with asymmetric costs is able to parsimoniously account for several prominent empirical puzzles. Furthermore, it develops some new insights into how a fund-raising advantage is likely to be manifested in the data: in particular, it should show up in the incumbent's *fixed effect* rather than his average spending effect, which provides an argument against policies that may increase an incumbent's cost advantage, even if they are estimated to have near-zero spending effects. Finally, we show that an adequately chosen spending cap may be effective at improving election outcomes, both in terms of reducing the incumbent's advantage, and in terms of informational efficiency; there is, however, a catch: an improperly chosen cap might yield no benefits, but only degrade election outcomes.

1. Search in auction platforms[*]

1.1. Introduction

New technologies often bring new markets and market structures: just as advances in shipping technology opened up countries to international trade in goods that were previously too heavy to ship, advances in communication and computation have created new marketplaces in goods, media, and advertising that could not have existed earlier. Many of these new markets are organized around platforms, which bring together large numbers of consumers and producers, offering each side easy access to the other. In some ways, these platforms resemble marketplaces that have existed for centuries: platforms act as intermediaries to facilitate transactions between agents that would otherwise not occur. In other respects, however, these new platforms are *sui generis*. The platforms themselves often represent substantial investments, and they embody many proprietary technologies that are just as important as the transactions that they enable. In fact, in many cases, the platforms themselves are far more visible than the any of the transactions that take place on them: eBay, iTunes, and Spotify, for instance, have become household names.

Platforms help to overcome several key frictions: (1) they provide market thickness; (2) they reduce transaction costs; and (3) they reduce search costs (Hagiu 2009). The first two functions are relatively well-understood, and they represent the two foremost challenges that a platform must solve in terms of its development strategy. The consequences of search within platforms, however, are less well understood, in spite of the growing importance that search facilitation seems to be playing on many platforms. For instance, media platforms, such as Netflix and Amazon, invest heavily in their recommendation systems, which help to direct viewers and readers to movies and books that they are likely to enjoy; in fact, Netflix, for several years, offered annual prizes to developers who could improve their user rating prediction algorithms. Some platforms,

[*] coauthored with Gregory Lewis

such as dating sites, live or die by the efficacy of their search systems: what matters for the paying members of a dating site is not the number of *potential* matches that might exist on the platform, but rather the *quality* of the few matches that actually do take place. The lesson seems to be that creating market thickness on the other side is not always enough. In situations where match *quality* is a central part of platform outcomes, a platform also has a role in making desirable matches easier to achieve.

The implicit intuition behind such efforts to facilitate search within platforms is that lower search costs represent a "win-win" situation for the agents on the platform and lead to Pareto-improving outcomes. In the case of Amazon, for instance, low search costs make it more likely that readers find books that they will want – they are better off since they end up getting books that they like more. Sellers are also better off because the readers most likely to purchase their books get exposure to them, and so their revenues are likely to increase.

Although the justification for lower search frictions seems straightforward, there has been relatively little theoretical analysis of the within-platform effects of search. The next section gives a more detailed overview of the existing literature; as far as we know, however, this paper represents the first rigorous look at the effects of search in a general platform setting. The two purposes of the present paper are (1) to provide the first positive results on the general effects of search within platforms, and (2) to develop a continuum framework that allows us to reduce the large platform game to a continuum of Poisson games, which are significantly more tractable than discrete formulations. On the positive side, we find that the "win-win" intuition is limited in scope: in general, it is true that the total surplus generated will be higher, since matches will on average be of higher quality. It is not generally the case, however, that all sellers benefit from such reduced frictions – allowing buyers to search into markets that they like means that some sellers may lose buyers overall, even if they gain desirable ones; a decrease in local market thickness has an effect counter to that of better matches, and there is no general way to sign the effects. We do find, nevertheless, that a sufficient condition for lower search costs to raise the revenues of all sellers on the platform is that the platform be symmetric. In such cases, the measure of buyers faced by each seller remains constant regardless of search intensity, so that the only channel affecting seller revenues is the creation of better matches, so revenues will be increasing.

1. *Search in auction platforms*

The primary motivation for the setup presented here is an auction platform such as eBay: buyers have unit demand for some good for which varieties exist, and they have "tastes" for different varieties; buyers acquire the goods by first searching into auctions for specific goods and then participating in those auctions. While the analysis on the welfare and distributional effects of search is particular to the allocation mechanism (in this case, single-unit auctions), the overall framework is quite broadly applicable and can be used to model sequential search into many-to-one matches.

The specific attention to auction platforms also allows us to make several important and perhaps unexpected connections between search effects and informational effects. Auctions are a promising setting for more general analyses of search since revenues are directly dependent both on market thickness as well as composition. Second, the nature of "interesting" auctions is also changing, since they are an increasingly popular way of allocating goods, particularly in the internet economy. While much of modern auction theory has developed around the design of single auctions, most modern auctions are held on platforms, which bring together large numbers of buyers and sellers. The new concern here, which supplements channels that have been more extensively analyzed, is that not only is the auction design itself important, but also the design of the mechanism by which individual buyers get allocated to individual auctions.

The main positive findings for auction platforms are that matching is welfare-improving, but that the distribution of surplus may be adversely affected by more efficient matches. For auction platforms, we are especially interests in matching effects on total seller revenues, since an *ad valorem* tax is typically the most popular revenue instrument for platform operators – higher prices and seller revenues translate directly into higher platform revenues. We show that under certain conditions, when the markets for the two goods are *ex ante* identical, matching is beneficial for all sellers; this generalizes both popular intuitions on matching effects as well as recent experimental evidence that attests to their benefits. When the markets are asymmetrical, however, for instance, when one type is of unanimously higher "quality", or when one good is much rarer relative to the other, then the reallocation of buyers between different types of goods can have negative overall revenue effects, and seller revenue can decline. While we do make certain assumptions on buyer preferences, the main negative result should generalize to most auction platforms with unit demand – the assumptions we make represent a ''best-case'' scenario for the generation of positive matching effects, and since even in the best case, total seller revenue can actually decline with more intense search, this effect should only be more pronounced under more general preferences.

To get an intuition for our main finding, we might conceptually decompose search effects into two parts: (1) a matching effect, which increases the expected *total* surplus generated by individual matches; and (2) a segmentation effect, which causes markets for some goods to become smaller, since fewer buyers search for those goods. The latter will often affect at least one group of sellers negatively, and this can outweigh the benefits provided by other sellers, since there are decreasing returns to market thickness when supply is fixed. Hence, even if individual matches may be of higher quality and generates high *total* surplus – hence an increase in total welfare – the seller's cut of that surplus depends on how many other buyers are also participating in the same auction. To take a stark example, in an auction with only one buyer we can make the buyer valuation, and hence total surplus, arbitrarily large, but seller revenue is always zero in a standard auction without reserves.

Our paper is organized as follows: the next section overviews the existing literature and shows how our work relates to questions that have previously been explored. The third section presents a simple discrete example to illustrate the effects of search within platforms. The four section establishes the formal model and notation; the fifth presents the analysis and results. We also devote a separate subsection to the development of a simpler model to highlight our technical contribution and its role in solving the model. The final section offers our conclusions and thoughts for future research directions.

1.2. Literature Review

The three platform functions mentioned in the previous section correspond to three characteristics of certain markets that make it amenable to platform intermediation: (1) the presence of network effects; (2) the existence of shared of transaction costs; and (3) the existence of search costs. This section gives a broad overview of the current literature on platforms, using these three distinct characteristics as a guide.

The current literature on platforms has focused almost exclusively on the first of these functions: the realization of network effects (Armstrong 2006; Rochet and Tirole 2006; Weyl 2010). A platform becomes more attractive to potential participants as the number of existing members on the other side of the market grows. An auction platform such as eBay, for instance, is more attractive to sellers as the number of buyers grows; a software platform such as iOS becomes more attractive to developers as the ownership of iOS devices increases. This goes the other way around as well: more buyers will join eBay the more sellers there are, since they are more likely to find goods that they want, and more consumers will purchase iOS devices the richer is its application ecosystem, i.e., the more software developers are making iOS applications. This implies that the foremost concern for platforms is to attract a high enough user base, since platform economies experience increasing returns to scale in the number of participants. This also has important implications for the number of platforms that can exist in equilibrium, and how they might be structured; in general, platform markets will support only one or two platforms in equilibrium because of these ‘‘tipping point’’ effects (Ambrus and Argenziano 2009; Ellison and Fudenberg 2003).

On the practical side, related to the existence of shared and transaction costs, computer scientists have devoted considerable effort into analysis of reputation mechanisms that can enable transactions to take place. Attracting members on one side is about guaranteeing enough market thickness to make it worthwhile for potential participants on the other side to use the platform; reputation and payment mechanisms, on the other hand, are about reducing transaction costs in order to overcome other barriers to transactions that might exist, so that these surplus-generating transactions actually do take place. Security has long been, and always will be, one of the central concerns of e-commerce, and part of the reason that so many platforms seem to have emerged with the rise of e-commerce is probably the need to provide a degree of trust that is otherwise absent in online interactions. This aspect of platforms is also relatively well understood, and more often represents a technical challenge rather than a theoretical question: reducing transaction costs, reducing asymmetric information, is just about always good, even if there are many tricky details in the implementation.

The third domain of platforms – *search* within platforms – has been developed to a much lesser extent. To the best of our knowledge, only two papers have examined search within platforms, both from the past couple of years. This may be partly due to the fact that search costs are something of a higher-order concern: in order for a platform to exist at all, it must attract enough members, and it needs to provide a way to ensure that the market does not break down. On the other hand, as platforms begin to encompass more and more markets, and as platform services become more specific to the user (individuals with specific tastes in music, movies, companions), the facilitation of search is becoming more and more important, and getting the answers to questions about search effects becomes more urgent.

Of the two recent papers to address search, one is theoretical and one empirical. Hagiu and Jullien (2011) analyze a Hotelling type model with two sellers and a continuum of buyers; their main result is that platforms, which derive revenue from the number of matches generated, have an incentive to divert search, in the sense that they may direct buyers to their less preferred sellers and prevent some surplus-enhancing trades from taking place. Our model is similar in spirit but different in focus: we ignore platform participation effects in order to analyze more closely the *within*-platform effects of decreasing search costs. We also allow for continuous variation in buyer valuations rather than search costs, which creates more subtle incentives for buyers and consequences for sellers, since a buyer's valuation has immediate consequences for the competitiveness of a buyer for a given auction as well as the contribution of the buyer's presence to total surplus and revenue.

The other recent paper is a study by Tadelis and Zettelmeyer (2011), which provides a direct motivation for our analysis: they find that decreasing search costs in an auction platform can raise the revenues of *all* sellers on the platform. It is worth providing a brief overview of their results to develop an idea of how search can affect the utilities of different agents on a platform. The authors run an experiment in wholesale auto auctions, which are auctions for used automobiles, of varying quality. For vehicles in their experimental group, a third-party is paid to assess the vehicles and to assign to each one a quality score: the quality score is made publicly available with the auction listing. The way that the auctions are run puts constraints on the abilities of individual bidders to join multiple auctions easily: many auctions are run simultaneously, and participating in an auction usually entails a physical inspection of the vehicle, which is a time-consuming process. The only difference between the control and experimental groups is the availability of the third-party quality score up front, prior to auction entry decisions. They find that revenues increased for auctions of *all* qualities in the experimental group, and that the effect was most significant for very high quality vehicles and very low quality vehicles, and least significant for those in the middle.

Although there are similarities between these results and those in the information revelation literature (Lewis 2011; Milgrom 1981; Myerson 1981; Riley and Samuelson 1981), there are several facts that do not fit with a purely asymmetric information interpretation: the prices increased for *all* quality types, and especially for low-quality types. If a fear of lemons is what suppresses prices in the control group, then the expected pattern would be lower prices for the low quality types, and higher prices for higher quality types. Second, the quality score itself does not really give any additional information at the time that bids are placed: since a thorough physical inspection is performed on the vehicles auctioned off, all participating bidders have a good idea of what they are bidding on. The effect of the intervention, then, is to move this information forward in the sequential game to the stage where buyers choose which auctions to join. If some buyers are looking for high quality cars, and other buyers are looking for lower quality cars, moving this information forward allows them to end up in the auctions that they want to be in. This means that buyers end up in auctions for goods that they value more, and sellers end up with more competitive auctions, which explains the higher prices.

Our paper generalizes the findings of Tadelis and Zettelmeyer (2011): we provide a formal model to describe the effects of search in a large platform with many buyers and sellers. We also show that under certain conditions, the most provocative of their results holds true: when valuations and supplies for two types of goods (e.g., high quality and low quality vehicles) are symmetric, then revenues of all types of sellers can be expected to increase. This gain can be interpreted as a ''matching effect'', in that from a seller's point of view, increased search means replacing relatively low-valuation buyers with relatively high-valuation buyers. Our continuum model is also able, however, to give a more detailed account of what happens when there are large numbers and shows that the rates at which low valuation buyers are traded for high valuation buyers will not, generically be equal. The primary implication of this is that some sellers will see a decrease in market thickness that can dominate the gains from better matching.

The search mechanics of our model are based on well-established models of sequential search (Diamond 1971; Mortensen 1970). What differentiates our model of search on platforms from existing models of consumer search and job search, however, is the fact that the search match (in the sense of the identity of seller and buyer) is no longer sufficient to determine the final payoffs. On the auction platform, buyers are searching into *competitive environments*, where the utilities they realize depend on the other buyers who end up matched with the same seller.

As alluded to in the introduction, this additional layer creates many additional technical challenges. We address them by noting that sequential search and random many-to-one matching (many buyers to one seller) lead to a continuum of Poisson games in the limit. The main technical contribution of this paper exploits this relationship, which allows us to characterize cleanly the expected utilities from search into auctions. Continuum frameworks have been used in continuous double auction settings (Satterthwaite and Shneyerov 2007), but not in a general search setting. We build upon the work of Myerson (2000), who establishes many of the convenient properties of Poisson games. We apply several of the key results derived there, as well as show how such games can arise naturally in sequential search settings.

1.3. Example

In this section, we present a modified version of the model presented by Tadelis and Zettelmeyer (2011). This is a significantly simplified model of search, but it helps to illustrate both the effects of search within an auction platform, as well as the difficulties of analyzing search effects in a full environment.

Consider a platform market with two goods, A and B, and four buyers. Buyers have quasilinear utility and are one of two types, a and b. valuations are given by $v_a(A) = v_b(B) = 1 + \varepsilon_i$ and $v_a(B) = v_b(A) = \varepsilon_i$, where ε_i is a random perturbation for each buyer. The perturbations are independently and uniformly distributed on $[0, 1]$, and there are two buyers of each type.

There are two stages to the game. In the first stage, buyers are matched to auctions, and in the second stage the goods are auctioned off in second-price auctions. The solution concept that we will use is Bayesian SPNE, and in the second stage we restrict ourselves to weakly undominated strategies, so all buyers bid their valuations.

In the matching stage, buyers are either randomly matched to auctions, or they can pay some cost c to be matched to an auction for a particular type of good. We motivate this setup by supposing that buyers cannot distinguish between auction types, so without additional information they choose randomly – such a buyer would end up in each auction with probability 1/2. By paying some cost, however, they can let the platform perform the match, which will direct them to the good that they prefer. There are various mechanisms by which this might be effected: either we could think of c as additional effort required to get information on the auctions, or as a service provided by the platform. Our full model will implement search as sequential search, but in this simple example with only two goods, it is simpler to think of search as a cost to be paid to ensure a favorable match.

Result 1.1. *All pure strategy equilibria in undominated strategies are characterized as follows:*

- *for* $c > 13/48$, *nobody searches*
- *for* $1/12 < c < 13/48$, *one of each type searches*
- *for* $c < 1/12$, *all buyers search*

Proof. First, nobody will ever choose to search into an auction for which he does not have a high valuation. For an a type, going into an A auction gives a payoff of at least 1/6 (what they expect when the other a type is present [1]), whereas going into the other auction gives greater than 1/6 only if everybody else search into A, which is clearly suboptimal.

Consider the case of no searching. We consider the situation in terms of a; it is symmetric for b. a going into an A auction: he receives a payoff of 1/6 if the other a is present, 3/2 if nobody is present, 1 if only one b is present, and 5/6 if both other b are present (this is simply expected value of winning minus expected payment; going against other types is trivial because a always wins, so pays the highest of the other bids, which is expected to be 2/3). Then the utility from an

[1] To see this, note that he pays at least 1 when winning, so we can just scale everything down by 1. If $\epsilon_i = x$, he wins with probability x, and conditional upon winning pays $x/2$. Integrating gives an expected payment of 1/6. Benefit is x received with probability x, so integrating over x gives 1/3

A auction is simply the probability weighted sum:

$$\frac{1}{2}\left(\frac{1}{6}\right)+\frac{1}{2}\left[\frac{1}{4}\left(3/2\right)+\frac{1}{2}(1)+\frac{1}{4}\left(\frac{5}{6}\right)\right]=\frac{15}{24}$$

Utility from B is zero whenever a b participates, $1/2$ when nobody participates, and $1/6$ when the other a participates. Here, the probability weighted sum is $(1/8)(1/6+1/2) = 1/12$. Searching is only optimal if half the difference between utilities from A and from B is greater than c, since that is the difference between searching into A and getting randomly matched to A or B. So no searching requires $c > 21/96$.

Below that, either one a or one b will find it optimal to search. If a b searches, then a gets zero from a B auction, and $(1/2)(1/6) + (1/4)(1) + (1/4)(3/2) = 17/24$. Hence, his searching threshold for c is $17/48$, which is satisfied conditional on b searching; hence there will never be only one side searching, as soon as one a searches, so will one b.

When one of each side searches, then expected utility for the other a is $1/6$ for A and 0 for B; hence, he will search when c goes below $1/12$. □

Using these strategies, we can get expressions for seller revenue as well.

Result 1.2. *Let $c_0 < 1/8 < c_1 < 21/64 < c_2$. Let $R(c)$ denote the expected revenue when the search cost is c. Then*

$$R(c_0) > R(c_1) > R(c_2)$$

Proof. Under full searching, expected payment for seller A is $4/3$. When one of each type searches, it is $4/3$ with probability $1/2$ (the probability the other a being there), and $1/2$ with probability $1/4$ (the probability that the other a is in B, and b is in A). Hence, $R(c_1) = 19/24$. Under no searching, with probability $1/16$ it is $1/3$ (both a end up in B, both b end up in A), and with probability $1/4$ it is $4/3$ (both a search into A). With probability $1/4$, there is only one participant, so revenue is zero. With probability $1/16$ there are no participants, so zero. With probability $1/4$, it is $1/2$ (with one a and one b); with probability $1/8$, it is $2/3$ (one a and two b). Hence $R(c_2) = 27/48$. □

The basic idea captured by this example is that lowering search costs encourage buyers to seek out better matches. In an auction setting, this means that sellers end up with buyers who have higher valuations for their goods, and increased competition between buyers drives up the prices for which the goods are sold. Allowing for continuous variation in c shows that the improvements will be incremental: as search costs are lowered, more buyers search.

These incremental improvements, however, also underline the difficulties that arise if we try to generalize something like the above model to a platform with large numbers of buyers and sellers. Since individual decisions affect the utilities received by other participants on the platform, it is necessary to keep track of all these effects. Above, for instance, we needed to check how one buyer of each type searching affected the incentives of the other buyer. As the number of buyers and sellers on the platform grows, this task becomes very quickly unmanageable. By moving to a continuum framework, however, we are able to get around this technical difficulty since auction participation will vary continuously rather than discretely (according to parameters that we will specify in the following section).

It will turn out that being able to do the bookkeeping in a tractable manner is also necessary to capture the full effects of search. The above example is completely symmetric in the two goods: we can interchange the labels A and B and nothing of the strategic situation will have changed. One consequence of this symmetry, however, is that buyer incentives are always the same for both goods, so that the number of buyers searching into an auction equals the number of buyers searching into the other auction: the expected number of buyers for each does not change as search costs decrease. The revenue ranking is only possible because of this: when the markets are asymmetric, one market will tend to lose buyers as the other gains them. Hence, it will generally not be the case that all sellers are better off, and it can even sometimes be the case that total revenues collected are lower.

1.4. Model

Let there be a measure μ of buyers, and a measure 1 of sellers on the platform. Sellers are one of two types, A and B, which denotes the type of good that they have. A fraction a of sellers are type A.

Goods are sold off in second-price auctions. Buyers have quasilinear utilities, and their types are independently and identically distributed on $\mathbf{X} = [\underline{x}, \overline{x}]$, where a buyers type x also denotes his valuation for the A good. A buyer of type x has a valuation for a B good given by $y(x)$, where y is continuous and $y' < 0$ for all $x \in \mathbf{X}$. Let $\mathbf{Y} = [\underline{y}, \overline{y}] = [y(\overline{x}), y(\underline{x})]$ denote the range of y. This means that buyers with relatively high valuations for the A good have relatively low valuations for the B good. It is important to note that this is *not* a restriction on the absolute valuations: it could be that every buyer values the A good more than the B good; rather, this is a restriction on relative valuations. The main reason for this assumption is that it gives a "best-case" scenario in terms of generating positive matching effects. It is likely that many actual platforms will not satisfy this restriction, which would only worsen the results we have for revenue. Although a more general formulation may be possible, one of the main goals of this paper is to test the "everybody wins" intuition that is implicit in platform policies to promote search, so we have made the environment as "matching-friendly" as possible. We show that even with this fairly strict condition on preferences, total revenue on a platform can actually decrease – the effects will only worsen with more general preference structures. The reason for this is that this restriction minimizes the negative impact that buyers searching out of markets will have on sellers: the buyers searching out – those who have a lot to gain from going for another good – are also those buyers with the *lowest* valuations from the seller's point of view. A more general preference structure will mean that this is no longer the case: buyers sorting out may have relatively high valuations, and so their absence will have a larger impact on seller revenues.

Buyer types are independently and identically distributed according to F, which we also use to denote the cumulative distribution function. F is massless and continuously differentiable; let f denote its probability density function. We will also let $\mu(X)$, for measurable $X \subseteq \mathbf{X}$, denote the measure of buyers with types in X, and $\mu(x)$ denote the density of buyers at x. Also, let $P(X)$ denote the probability that a random buyer has a type in X. Then $\mu(X) = \mu P(X)$ and $\mu(x) = \mu f(x)$. It will also be convenient to define $H(y) = 1 - F(y^{-1}(y))$ to be the cumulative distribution function of valuations for the B good.

The goods are allocated according to a two-stage game. In the first stage, buyers are matched to sellers via sequential search; in the second stage, goods are sold off in second-price auctions. The search process is sequential search mediated by platform technology. Buyers are initially randomly matched to sellers. Upon being matched, buyers observe the type of good being sold on the auction. Each buyer can stay in the auction he is matched to, or at a cost c, he can draw another match; the search stage ends when nobody chooses to draw any more matches. At each match, the only observed characteristic is the type of good being sold in the auction – buyers do not know the number of other buyers.

The success of each draw is determined by a platform-wide parameter p. With probability p, a buyer will be matched with an auction for a type of good that he prefers; with probability $1 - p$, the match will be random. Hence, $p = 0$ corresponds to completely-random matching in sequential search, and $p = 1$ corresponds to perfectly efficient search: an additional draw will always yield a successful match. There are several ways that we might conceive of search and the role of the platform in mediating search. An alternative formulation, for instance, would be to drop the parameter and have the platform operate directly on the draw costs. The particular formulation does not impact any of the below analysis: one could conceptualize reductions in search costs either as reductions in the draw cost, or as increased probabilities of draw success. More generally, however, it is reasonable to think of search improvements as benefiting primarily those who are looking for goods that would be hard to find otherwise, which is why we have formulated search as platform intermediated search, following Hagiu and Jullien (2011). Depending on the setting, it could be reasonable to think about search costs either as a reduction in draw costs or as changes to the draw success probability – we choose the more general formulation, which keeps both parameters.

The equilibrium concept that we use is SPNE. We also restrict strategies to weakly undominated strategies, so that all buyers bid their valuations in the auction stage. The solution will be by backward induction: first, we solve for buyer utilities given the search strategies of other buyers. This step is dramatically simplified by our continuum framework, since each buyer's individual decision will have no impact on the aggregate environment, and the characterization of platform matchings as a continuum of Poisson games will yield tractable expressions for the objects we are interested in. We can then use these expressions to solve for the optimal search decisions.

1.4.1. Discussion of Assumptions

It is worth pausing for a moment to discuss the assumptions within our model and what types of restrictions they place on our findings. There are two main components to our platform model: (1) the description of search outcomes in a continuum framework, and (2) the analysis of equilibrium search decisions; these two parts require two distinct sets of assumptions. Although in a certain light some assumptions may seem strict, it will turn out that our main qualitative results are not fragile with respect to them.

The main property exploited in the development of our search model is random matching, which is inherent to the sequential search process. Although our model is limited to two types of sellers, it can be easily extended to continuous seller types with appropriately defined measures. Since random matching is a natural property of sequential search situations, this is not a restrictive assumption on the applicability of our search model: it can be easily extended to describe matching outcomes with a continuum of seller types and arbitrary (but measurable) ''acceptance sets'' for each buyer.

Our equilibrium analysis relies more heavily on a potentially less realistic assumption: that buyer preferences are unidimensional, and buyers with higher valuations for one type of good will have fixed lower valuations for the other type. We make two brief remarks about this assumption. First, many markets may not in fact stray too far from this type of ''complementarity'' in preferences. For any good with multiple varieties, it is likely that consumers with very particular tastes, i.e., they derive significantly greater utility for getting their preferred variety, will also dislike, to a greater extent than those without strong taste preferences, *not* receiving their preferred choices.

The second point to note is that unidimensional preferences, while they play a key role in much of our formal analysis, are not wholly essential for our main qualitative results; they serve primarily to give tractability for our formal results, i.e., that equilibrium is unique and can be parsimoniously characterized. Our welfare result turns on the fact that social and individual welfare are aligned in VCG mechanisms, and it can be proven in more general settings. Our other main result is that seller revenue can decrease as search costs decrease. In this context, unidimensional ''complementary'' preferences represent a best-case setting for generating positive matching effects on seller revenues: when a seller loses buyers due to more intense search, those buyers are always the least valuable ones. A more general preference structure would only mean that the buyers *lost* by a seller are not necessarily the least valuable, and the buyers *gained* are not necessarily the most valuable, so the negative effects of search on revenue would only be more pronounced.

1.5. Analysis and Results

This section contains the analysis and results for our platform search model. We begin with a slight detour to a simpler model with a single good and no search in order to develop the analytical machinery necessary to solve for the full model with search. This will allow us to easily describe the effects that search will have on the second-stage auctions. We then discuss the search stage and characterize the search decisions of individual buyers. This will allow us to establish the structure of equilibrium and show equilibrium existence and uniqueness. Using the fact that in a second-price auction, a buyer's utility is equal to his contribution to social welfare (McAfee and McMillan 1987), we can show that the search equilibrium is also efficient. The remainder of our results address the effects of lowering search costs on platform revenues.

1.5.1. Single good without search

In this section, we develop the continuum model as a the limit of a sequence of discrete games as the number of buyers and sellers grows large. In order to focus on the technical aspects of the continuum model, we consider a simpler model that has only one good, so all sellers are identical, and no search. We are mainly concerned with characterizing the matches that result form random search. Consider a sequence of discrete games $\{G_i\}_{i=1}^{\infty}$, where G_i is a game with n_i sellers and $\mu_i n_i$ buyers, with n_i strictly increasing and tending toward infinity (i.e., $n_i \to \infty$) and $\mu_i \to \mu$ as $i \to \infty$. Buyers are matched to sellers randomly. Buyers have independently and identically distributed valuations for the good, distributed according to a massless F, which we also use to denote the cumulative distribution function of buyer valuations.

Since all buyers are identical, the key step is to characterize the number of buyers that shows up at any given auction. All sellers are identical, so each will face the same distribution of buyers.

Theorem 1.1. *In the continuum limit of the sequence of games $\{G_i\}$, the number of buyers in any* *given auction is a Poisson random variable with parameter μ.*

Proof. For game G_i, there are n_i sellers and $\mu_i n_i$ buyers. Random matching implies that each buyer has an equal chance of ending up in any auction, so from the point of view of any fixed seller, the probability that any particular buyer joins his auction is $\frac{1}{n_i}$. Since there are $\mu_i n_i$ buyers, this means that the number of buyers to join any particular auction j is binomially distributed, with $\mu_i n_i$ draws and a "success" probability of $\frac{1}{n_i}$ for each draw. Let k_i denote be a random variable denoting the number of buyers to show up at an auction in G_i, then

$$k_i \sim \text{Binom}\left(\mu_i n_i, \frac{1}{n_i}\right)$$

Let $n_i' \equiv \mu_i n_i$ and $p_i' \equiv \frac{1}{n_i}$, then we can write this as

$$k_i \sim \text{Binom}\left(n_i', p_i'\right)$$

where $n_i' p_i' = \mu_i$, and we know that $\mu_i \to \mu$ as $i \to \infty$. Hence, as $i \to \infty$, $n_i \to \infty$ and $n_i' \to \infty$, which means that k_i converges in probability to a Poisson random variable with parameter μ. □

The independence of buyer valuations also allows us to characterize attendance of buyers from subsets of X. In particular, the number of buyers can be formulated as a spatial Poisson process indexed by valuation x, where the number of buyers from any subset of X is a Poisson variable with parameter equal to the measure of that subset, and is independent from the number of buyers who show up outside of that subset. Formally, let $p : 2^X \to \mathbb{R}$ be a probability measure denoting the probability that a given buyer is in a subset of X, and let $\mu : 2^X \to \mathbb{R}$ be the measure function for buyers on X; here $\mu(X') = \mu p(X')$ for all $X' \subseteq X$.

Lemma 1.1. *Let $X_1 \subset X$ be a measurable subset of X, and let k_2 be a random variable denoting the number of buyers who show up at a given auction with valuations in X_1. Then $k_1 \sim$ Poisson$(\mu(X_1))$.*

Furthermore, given a measurable subsets $X_2 \subset X$, with $X_1 \cap X_2 = \emptyset$, let k_2 be random variables denoting the number of buyers in a given auction X_2. Then $k_1 \sim$ Poisson$(\mu(X_1))$ is independent from $k_2 \sim$ Poisson$(\mu(X_2))$.

Proof. Proof is given in appendix. □

One useful consequence of this characterization is that it allows us to pin down winning probabilities quite easily.

Lemma 1.2. *Consider a buyer of valuation x prior to being matched to an auction. The prior probability that the highest bid among the other buyers is below x' is given by*

$$G(x') = e^{-\mu\left(1-F(x')\right)}$$

Proof. Proof is given in appendix. $\square$

Lemma 1.3. *A buyer of valuation x receives utility*

$$\begin{aligned} u(x) &= \int_0^x G(y)\,dy \\ &= \int_0^x e^{-\mu(1-F(y))}\,dy \end{aligned}$$

Seller revenue is

$$\begin{aligned} R =& m(1) - \mu \int_0^{\overline{x}} (1 - F(y))G(y)\,dy \\ =& m(\overline{x}) - \int_0^{\overline{x}} (1 - F(y))e^{-\mu(1-F(y))}\,dy \end{aligned}$$

Proof. Full proof is given in appendix. The derivation is quite standard (e.g. Krishna 2009), with the exception that we substitute $G(x)$ for the usual winning probability F^{n-1} for an auction with n bidders. $\square$

1.5.2. Sequential search

The preceding lemmas provide closed-form expressions for buyer utilities in the second stage of the full model with search. Using these, we can rigorously analyze the search decisions of individual buyers, and from there derive the equilibrium of the full game. For the search decision: the only information that buyers have upon being matched to a seller is the type of the good that is being sold – buyer search decisions can only be based on the type of good. Furthermore, since the result of a draw does not depend on previous draws, the distribution of draw outcomes remains the same for each draw, which means that the optimal search policy must also be stationary[2]. The consequence of this is that some buyers will always end up searching into the auction of their choice for the second stage, and others will remain with their random matches. If a buyer is matched with his preferred auction type, then he will always stay. If he is matched with his less preferred auction type, then if it is optimal for him to search again today, it will also be optimal for him to search again tomorrow – hence, he will continue to search until he ends up in the auction of his choice.

We can exploit the stationarity of the search decision to characterize the decision rules. Before stating the results, we introduce a little bit of notation: let $u^A(x)$ denote the utility a buyer of type x gets from participating in an A auction, and $u^B(x)$ denote his utility from a B auction – we will solve for these after establishing the structure of equilibrium search behavior.

Lemma 1.4. *A buyer will search into an A auction if*

$$u^A(x) - u^B(x) > \frac{c}{p_a}$$

and will search into a B auction if

$$u^A(x) - u^B(x) < -\frac{c}{p_b}$$

where

$$p_a = p + (1-p)a$$

$$p_b = p + (1-p)(1-a)$$

are the probabilities of successful draws by those who favor A goods and those who favor B goods respectively. When the inequalities are reversed, buyers will not search, and when the two sides of the above inequalities are equal, the buyer will be indifferent between searching and not searching.

Proof. Proof is given in appendix. □

[2] The search problem can be formulated as a standard dynamic programming problem and solved with a simple Bellman equation (Stokey, Lucas, and Prescott 1989). Since nothing in the environment changes between the current period and the next, if we go to the next period, it must be that the current period's (optimal) decision rule – which assigns an action, either search or stay, to the each realization of the draw outcome – is also optimal tomorrow.

1.5.3. Equilibrium

Since the utility in each auction is monotonic in the valuation for that object (increased probability of winning, and increased surplus when auction won), it is clear that equilibrium will feature a pair of threshold values, (x_b, x_a), such that buyers with valuations higher than the threshold x_a will search into A, and buyers with valuations that are high in the B auction, hence of types *less* than x_b, will search into B. Basically, if it is optimal for some given x to search into A, then it will also be optimal for all buyers of types greater than x to search into A, since their utilities from doing so will be strictly higher. We state this in the following theorem:

Lemma 1.5. *Equilibrium strategies can be characterized by a set of thresholds (x_b, x_a) such that all buyers with type $x > x_a$ search into A, all buyers with type $x < x_b$ search into B, and the remaining buyers do not search[3]. Furthermore, for interior x_b and x_a, the thresholds must satisfy*

$$u^A(x_a) - u^B(x_a) = \frac{c}{p_a}$$

$$u^A(x_b) - u^B(x_b) = -\frac{c}{p_b}$$

Proof. The complete proof is omitted. The derivation is quite straightforward, and follows from the monotonicity of $u^A(x) - u^B(x)$, as described above. □

The next step is to establish that an equilibrium exists and is unique. The details of this will be left to the appendix, but the steps are fairly straightforward, so we sketch them out here.

Let $s(x; x_b, x_a) \equiv u^A(x; x_b, x_a) - u^B(x; x_b, x_a)$ denote the difference in utility for a buyer of type x between participating in an A and a B auction, when the thresholds are set at x_b and x_a. If we can establish that this function is continuous in its parameters, x_b and x_a, then we can define a continuous mapping $M : X^2 \to X^2$ that gives for any pair of thresholds (x_b, x_a) a pair of types – those that are indifferent between searching into B and those that are indifferent between searching into A (with appropriate boundary conditions). Since X^2 is a convex (a two-dimensional interval) and M is a continuous mapping from X^2 to itself, Brouwer's fixed point theorem ensures that a fixed point to M exists – this will be an equilibrium. We can go further to show that $s(x; x_b, x_a)$ is

[3] By "search into A", we mean that buyers stationary search rule is to continue drawing new auctions until they end up in an A auction. The outcome of this search, seen from the second-stage auctions, is that these buyers will always end up in A auctions.

also increasing in x_b and x_a, which will guarantee that the equilibrium is unique and the threshold values change monotonically with search costs. The closed-form expressions for doing this will draw on the continuum model that we developed in the simpler single-good model.

Theorem 1.2. *An equilibrium exists and is unique. Furthermore, if we reduce search costs, either through a direct reduction in the per-draw cost c, or through an increase in p, then x_b increases and x_a decreases: that is, the threshold values are monotonic in search costs, and lower search costs expands the set of buyers who search into auctions of their choice.*

Proof. The proof is given in the appendix. □

1.5.4. Welfare

Having ensured that a unique equilibrium exists and is well-behaved, we continue to examine its properties. This section establishes two basic welfare properties of equilibrium, and gives some support to the idea that reducing search costs is generally a good thing overall. First, the equilibrium is efficient, in the sense that the market equilibrium coincides with the solution to the social planner's problem – we will make this more precise momentarily, but the actions taken by individuals in a market equilibrium are precisely those that a social planner would recommend they take in order to maximize total social welfare. Second, as search costs decrease, social welfare is increasing. This follows immediately from the efficiency result: since the market equilibrium maximizes welfare, and higher levels of welfare are possible – agents, for instance, can simply retain the same actions, but total welfare will be raised since search costs paid are lower – it must be that equilibrium welfare is increasing as search costs decrease.

The intuition behind the efficiency of the market equilibrium comes from the observation made by McAfee and McMillan (1987), that in a VCG mechanism, which the second-price single-good auction is, a buyer's utility is equal to his contribution to social welfare, since his payment is exactly equal to the external effect of his presence on the other buyers. As a result, when buyers maximize their own utilities in our platform model with search, they also maximize their contribution to social welfare (total utilities net search costs). The optimal assignment of search actions by the social planner must have every agent (except, possibly, a subset of measure zero) maximizing their contribution to social welfare, or there would be an alternative assignment that generates higher welfare.

To show all of this formally, we first define the social planner's problem. The social planner assigns to each buyer x an action in the set $\{\alpha, 0, \beta\}$, where α and β correspond to searching into A and B respectively, and 0 corresponds to not searching. The objective function of the social planner is total welfare net search costs, which we denote by W. To get a sense of what social welfare looks like, note that each buyer's contribution to total welfare is his valuation of the object that he bids for times his probability of winning. His presence has no effect on welfare if he does not win, and if he does win, the total surplus generated is his valuation, which is split between himself and the seller according to the price paid. We can then write the surplus generated in the A market as

$$\begin{aligned}\int_{\underline{x}}^{\overline{x}} xG^A(x)\mu^A f^A(x)\,dx =& \int_{\underline{x}}^{\overline{x}} xg^A(x)\,dx \\ =& E[X_1] \\ =& \overline{x} - \int_{\underline{x}}^{\overline{x}} G^A(x)\,dx \\ =& m^A(\overline{x})\end{aligned}$$

where X_1 is the highest bid, or the valuation of the highest buyer, since all buyers bid their valuations. The surplus generated in market B is similarly $m^B(\overline{y})$, or $E[Y_1]$. The total search cost incurred across the platform is c times the expected number of draws. Let $\mu_a = \mu(X_a) = \mu \cdot (1 - F(x_a))$ and $\mu_b = \mu(X_b) = \mu \cdot F(x_b)$ be the measures of buyers who search into A and B markets respectively. A fraction $1 - a$ of those in X_a need to search, and expect to search $1/p_a$ times, so the total cost incurred by them is $\frac{(1-a)\mu_a}{p_a}c$, and the total cost incurred by those searching into B is $\frac{a\mu_b}{p_b}c$. Hence, total welfare W can be written as

$$\begin{aligned}W =& am^A(\overline{x}) + (1-a)m^B(\overline{y}) - \left(\frac{(1-a)\mu_a}{p_a} + \frac{a\mu_b}{p_b}\right)c \\ =& aE[X_1] + (1-a)E[Y_1] - \left(\frac{(1-a)\mu_a}{p_a} + \frac{a\mu_b}{p_b}\right)c \\ =& a\left(\overline{x} - \int_0^{\overline{x}} G^A(x)\,dx\right) + (1-a)\left(\overline{y} - \int_0^{\overline{y}} G^B(y)\,dy\right) - \left(\frac{(1-a)\mu_a}{p_a} + \frac{a\mu_b}{p_b}\right)c\end{aligned}$$

The lemma below establishes that a solution to the social planner's problem must have a threshold form, e.g., it will consist of a pair (x_b^*, x_a^*) such that all $x < x_b^*$ search into B, all $x > x_a^*$ search into A, and the rest do not search. The full proof is given in the appendix, but the method is by contradiction: suppose that the optimal assignment does not have a threshold form, and show that total welfare can be improved by changing some of the assignments.

Lemma 1.6. *A solution to the social planner's problem must consist of x_b^* and x_a^* such that all $x < x_b^*$ search into B and all $x > x_a^*$ search into A.*

Proof. Proof is given in appendix. □

It is then straightforward to show that the optimal thresholds must be the same as the market equilibrium. The intuition is given above; in the appendix, we provide an alternative method that directly maximizes welfare by examining the first-order conditions of the social planner's maximization problem – because of the alignment of private and social surplus in a VCG mechanism, they will coincide exactly with the threshold conditions for the market equilibrium.

Theorem 1.3 (Efficiency)**.** *Let (x_b^*, x_a^*) be the social planner's solution, and let (x_b, x_a) be the market solution for a given set of search cost parameters, (c, p). Then $x_b^* = x_b$ and $x_a^* = x_a$. That is, the market equilibrium maximizes total social welfare and is efficient.*

Proof. Proof is given in appendix. □

A directly corollary to this result is that a decrease in search costs always increases total welfare. As stated previously, if search costs decrease, the previously experienced level of social surplus is always achievable, since if no agents change their actions, the welfare will increase due to the decrease in search costs paid (or stay the same if there is no search). Since the market equilibrium maximizes social welfare, the level of welfare it achieves must be at least that level.

Corollary 1.1. *As search costs decrease, total social welfare is increasing.*

These results lend some justification to the idea that agents are better off as a result of decreased search costs. The aggregate welfare consequences, however, do not tell us how the gains are distributed. The following examines the effect of search on seller revenues.

1.5.5. Revenue

Although total welfare is increasing, it is not necessarily the case that seller revenue increases with search. This section will analyze more precisely the effects of search on seller revenue – in particular, we exploit the infinite divisibility property of Poisson distributions to get a precise formulation of how lower search costs, and increased search, affect buyer participation in auctions.

As more buyers search, two things are happening for each type of buyer: (1) they gain buyers in the middle of the distribution, where previous non-searchers begin to search into their auctions; (2) they lose buyers at the bottom distribution, where previous non-searchers were getting randomly matched into their auctions, but now choose to search into the other auction. Hence, there is a ''matching effect,'' of the type mentioned by Tadelis and Zettelmeyer (2011), in the sense that lower valuation buyers are being ''traded'' for higher valuation buyers. This is now, however, the entire story: there is no reason that the rate at which buyers leave should generically be equal to that at which they enter. As a result, a seller may lose buyers overall as a result of search, if, for whatever reason, the effect of decreasing search costs on the marginal searchers into the other auction is greater than its effect on those search into the seller's auction. The loss in competitiveness would result in lower prices (in the form of a higher probability of the good commanding a price of zero).

On the other hand, when the two markets are perfectly symmetric, the net effects of search on revenue will be positive. This is because under symmetry, the scenario described above will never occur: the measure of buyers gained in a market, in the middle of the value distribution, will always be exactly equal to the measure lost at the bottom of the distribution. From the perspective of the seller, this is equivalent to replacing every low valuation buyer with a higher valuation buyer, which is strictly revenue enhancing.

Figure 1.1 shows what happens to auction participation in an A auction as search increases, i.e., the searching thresholds change from (x_b, x_a) to (x_b', x_a'), with $x_b' > x_b$ and $x_a' < x_a$. Recall that auction participation can be conceptualized as a Poisson process, where the density parameter is equal to the population density of participants, which we denote, again with a slight abuse of notations, using $\mu(x)$. Independence of non-overlapping sets allows us to conceptualize participation in a particularly straightforward way.

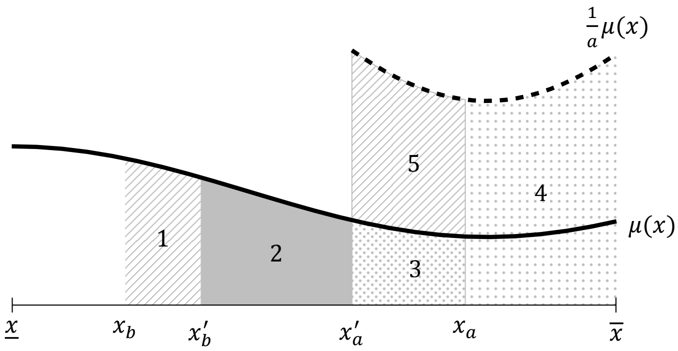

Figure 1.1.: Effects of search on auction participation

The solid black line is the population density function, which is also then density function in an A auction when no buyers search (since for any set in X, a fraction a of them are placed into an A auction, and we normalize by the measure of A sellers, which is a). Now consider the density of buyers in an A auction under thresholds (x_b, x_a). In $[\underline{x}, x_b)$, all buyers search into B, which means that no buyers appear in the A market, so the density is zero. Along $[x_b, x_a]$, no buyers search, so the density remains $\mu(x)$. Finally, along $(x_a, \overline{x}]$, the platform density μ of buyers search into A, and we normalize that to a sellers, so the density is $\mu(x)/a$, which is depicted by the dotted line. We can partition X into intervals defined by the points x_b, x_b', x_a', and x_a: the number of buyers who show up from each of these intervals is an independent Poisson variable with their respective densities. Letting k_i denote a Poisson variable with measure equal to the shaded region i in the figure, total participation in an A auction can be written as $k_1 + k_2 + k_3 + k_4$.

If we move the thresholds now to (x_b', x_a'), two changes happen: the density in $[x_b, x_b')$ becomes zero, and the density in $(x_a', x_a]$ changes from $\mu(x)$ to $\mu(x)/p_a$. This means that k_2 and k_4 remain the same, and k_1 is replaced by zero. k_3 is replaced with a Poisson random variable of $(1/a)\mu((x_a', x_a])$. Since the density is only scaled up, however, and Poisson variables are infinitely divisible, this is equivalent to two independent Poisson variables, one with parameter $\mu((x_a', x_a])$ and the other with parameter $(1/a)\mu((x_a', x_a])$ (or, alternatively, we can consider the Poisson process in $(x_a', x_a]$ to be the sum of two independent processes, one with density $\mu(x)$ and the other with density $(1/a)\mu(x)$). Hence, participation can be summarized as $k_2 + k_3 + k_4 + k_5$.

The net effect is that we remove the Poisson draw from $[x_b, x_b')$ with a Poisson draw from $(x_a', x_a]$ (where the value distribution of a buyer in those regions is the conditional distribution of valuations, conditional on being within the region). Obviously, the removal of k_1 by itself always hurts revenue, and the addition of k_5 by itself always helps revenue. In general, the net effect cannot be signed, since the measure of buyers sorting in and those sorting out will differ. In the special case where the value distributions of each good are identical (i.e., the platform is symmetric), however, the effect on revenue is unambiguously positive.

First, we give a formal definition for symmetry in this scenario.

Definition 1.1. *A platform is symmetric if permuting the labels A and B (e.g., calling the A good the B good and vice versa) does not change any of the value distributions or proportions of goods in the seller population.*

Note that the conditions for this to be true are quite specific: it must be that $a = 1/2$, $\overline{x} = y(\underline{x})$, $\underline{x} = y(\overline{x})$, $y(x) = \overline{x} - x$ and $H(x) = F(x)$ for all $x \in X$, where $H(x) = 1 - F(y^{-1}(x))$ is the cdf of valuations in for B goods, i.e., buyer valuations for A and B goods have the same distributions.

When a platform is symmetric, the two equilibrium sorting thresholds will also always be symmetric around the mean of the value distribution (since relabeling the goods does not change the fundamentals, this is implied by uniqueness). Hence, the measure of new searchers when the search cost decreases will be the same in each market; from the point of view of any seller, this means that the measure exiting is equal to the measure entering. That is, there is no net change in measure, but simply a replacement of all the lower buyer types with higher buyer types. In terms of the figure, k_1 and k_5 have the same parameter – since they are independent, this is equivalent to replacing each buyer in $[x_b, x_b')$, should at least one show up, with a buyer in $(x_a', x_a]$, which is always revenue improving. We state this in the following theorem:

Theorem 1.4. *For a symmetric platform, as search costs decrease, either through a reduction in c or an increase in p, expected revenue is increasing for all sellers.*

By considering the symmetric case, it is also easy to see how revenue might decrease: the measure of buyers leaving may be much larger than the measure of buyers entering, and so even though those entering may have higher valuations, their number may be insufficient to counteract the negative effects of decreased buyer density.

In the generic case without symmetry, we might conceptualize each seller as seeing two effects of matching: one for one replacement of low valuation buyers with high valuation buyers is always revenue enhancing, and can be though of as a *valuation effect*; on the other hand, there is no guarantee that replacement is one for one, so the *change* in measure leads to a segmentation effect: as buyers enter the markets they like most, the composition of buyers in each market will shift, and will generically decrease for one of the types of sellers. This will exert downward pressure on revenues in one market and additional upward pressure on revenues in the other.

In addition, it is also possible that *total* revenues may decrease as a result of lower search costs. This might happen, for instance, when the market *losing* buyers also commands higher valuations overall. We provide a simple numerical example to illustrate this.

Theorem 1.5. *It is possible for overall revenue to decrease as search costs decrease.*

We show this by constructing a numerical counterexample. Consider a platform where $a = 0.05$, $X = [0, 1]$, $y(x) = 5 - x$, and $\mu = 3$. The following examines what happens as we decrease search costs by moving p from zero to one, when the per-draw cost, c is equal to 0.01.

In this market, the relatively abundant good is also the one that people prefer, the B good, for which all buyers have valuations between two and three. Under perfectly random matching, it is also the good that most buyers will draw into. The low valuation buyers, however, have little chance of winning the good, and hence would prefer to receive the A good with a higher chance of winning and higher surpluses.

As search costs decrease (by increasing p), the effect is most dramatic for those who would prefer to search into the A market, since that is the one that is harder to draw. When $p = 0$, most buyers would still be able to draw into the B market relatively easily, but would not be able to draw into the A market (the expected number of draws is 20). The expected number of draws for a buyer who wants to get in to a B auction, by contrast, is only $10/19 \doteqdot 1.052$. A change from $p = 0$ to $p = 1$, then, represents a tenfold decrease in the effective search cost for those searching into A, and only an eleven percent decrease in search costs for those searching into B. Hence, cost decreases through p will induce relatively more buyers to search into A than into B.

Hence, the B market will lose a large measure of buyers at the bottom, and gain some in the middle, whereas the A market will gain a large measure of buyers in the middle, and lose some at the bottom. Revenue decreases in the B market and rises in the A market.

Figure 1.2a shows the threshold values as p changes, and figure 1.2b shows how the measure of buyers in each market changes. The dotted lines show results for the A market, and the solid lines show results for the B market. As expected, x_b rises and x_a falls, since lower search costs induce more searching. Although the threshold x_b moves more than the threshold x_b, B loses buyers overall, which is shown in figure 1.2b. The reason for this is that the additional "searchers" into B would mostly have ended up in B anyways, as would the additional searchers into A: searching *into* B has relatively little impact on the measure of buyers in B, and searching *out* of B has a large impact.

These changes in participation are reflected in the market revenues, which are shown in figure 1.2c. Revenue in market B drops, since B auctions are overall losing buyers, and in this case the valuations of the buyers that they gain are relatively low. In contrast, A revenues increase. Finally, figure 1.2d shows *total* revenues. Although the A revenues do increase, there are relatively few A sellers, so the effect on total revenues is dominated by the decrease in B revenues.

Although in this example, the magnitude of the effects is small, it suffices to demonstrate the possibility of an overall revenue decline as a result of improved matching. Our counterexample was constructed around a particular type of asymmetry: the existence of a small market with lower overall valuations and relatively low per-draw search costs. Our numerical solutions seem to indicate that many parameter values do generate monotonically increasing *total* revenues as search costs decrease, and that overall revenue decreases seem to be reserved for cases with significant asymmetries and either high per-draw costs, so that a decrease in search costs induces search into only one market, or when search costs are low, so that many of the gains from search have already been achieved. Numerical results are similar when, instead of varying the search success probability p we directly vary the search cost c.

Although our numerical investigations do suggest that negative overall revenue effects are "rare", it should be kept in mind that we have constructed our model as a "best case" scenario in terms of generating positive revenue effects as a result of matching: buyer valuations for each good are deterministically negatively correlated with their valuations for the other good. In more general settings, either with general joint distributions, or with joint distributions with certain monotonicity properties on the marginal distributions, the negative revenue effects can only be exacerbated, since there would no longer be the guarantee that the buyers leaving a market as a result of search are always those with the lowest valuations.

The fundamental lesson in terms of revenue is that lowering search costs affects revenue in two ways, and that the net effect is not a Pareto improvement – there are significant distributional consequences in spite of the fact that welfare is increasing. Additional search (1) replaces low valuation buyers with higher valuation buyers for all sellers, which is generally good, and (2) shifts buyers between the two markets, which will, in the two good setup, be good for one type of seller and bad for the other. One might formally conceptualize these two different effects by breaking down the new entrants and exits from a particular market into two separate components: one component consisting of the largest set of buyers among the new entrants and exits whose effect on the total measure is zero (i.e., a set of entrants of measure $\min\{\mu_{in}, \mu_{out}\}$, and a set of exits of the same – this set represents the one-for-one replacements that are always revenue enhancing) and the second component consisting of the remainder, which will have either a positive or negative effect on the total measure, depending on whether more or fewer buyers leave than enter.

The effect on revenue from the first component is strictly positive, since it is simply replacement of low valuation with higher valuation buyers, which generates a positive *valuation effect*. The effect of the second component is positive if there are more entrants than exits, and negative in the converse situation. This comes from shifting buyers between markets, and constitutes a *segmentation effect*, which benefits one seller and harms the other. Although one type of seller will always gain, there is no guarantee that the total seller revenues generated on the platform will increase as a result of decreasing search costs.

1.6. Conclusion

As mentioned in the introduction, this paper aims to be a first step in analyzing the effects of search in platform settings. There has been relatively little work on this, and part of the reason may be technical: the continuum framework that we provide, which simplifies auction participation in a large network to a Poisson process, should help to overcome many of the technical barriers associated with analyzing platforms with large numbers of agents on both sides of the market in situations where search matters. We have also provided positive results to shed some light onto the effects of search on platform welfare and revenue, results that may, in turn, help to

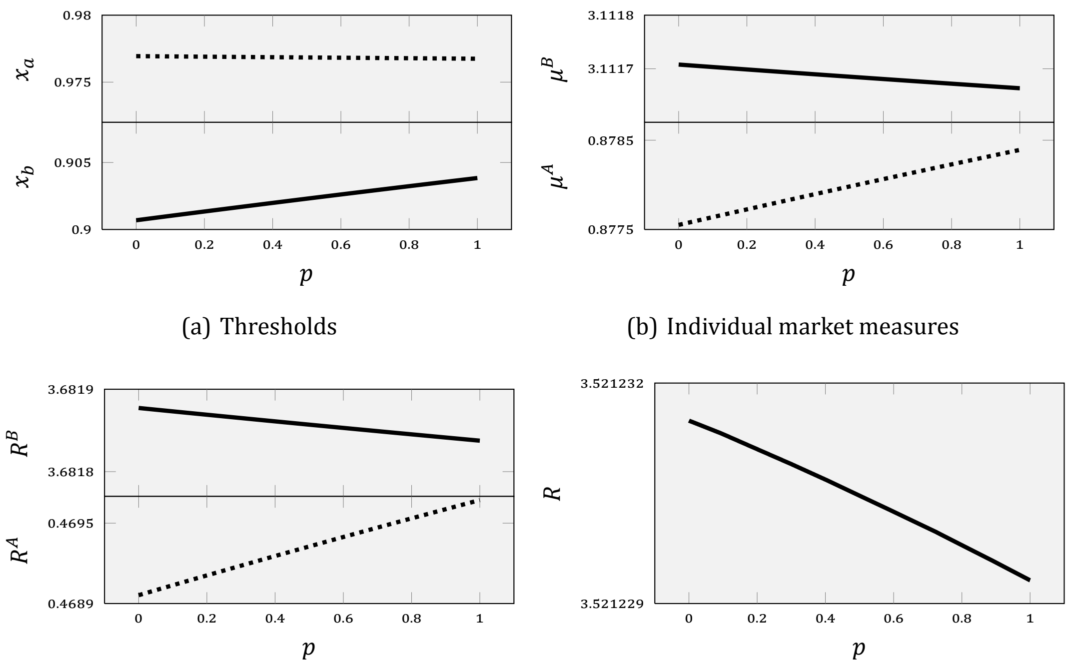

(a) Thresholds

(b) Individual market measures

(c) Individual market revenues

(d) Total revenue

Figure 1.2.: Increasing p under asymmetric conditions

explain and guide platform policy related to search investments. In particular, in a second-price auction platform – and by revenue equivalence, all efficient, single-good auction platforms (Riley and Samuelson 1981) – total social welfare is increasing as search costs decrease. Moreover, in a symmetric platform, all sellers will benefit from search, so that total seller revenues and individual seller revenues will increase as search costs decrease. If platform interests are aligned with seller interests, then (as would be the case if the sellers were being charged, for instance), then a platform would want to invest in decreasing search costs in symmetric situations. On the other hand, we have also shown that in a generic setting, there are significant distributional consequences to increased search: generally, one set of sellers will lose, while the other gains, but there are also cases where *total* revenue can decrease.

Platforms exist in great variety and serve a large spectrum of functions, so the applicability of our results to any particular platform must weigh in its specific characteristics. Our analysis has the most direct bearing, of course, on auction platforms, but there are many other types of intermediary services that are not organized around auctions, about which we are able to provide some insights. The most generally applicable contribution of our paper is that for one major technical problem – the description of search outcomes in platforms with large numbers of agents –

we have provided a general solution. In many platforms, going beyond those that deal exclusively with auctions, search is many-to-one or many-to-many – the simplification to Poisson "matching sets" in the continuum limit is applicable to almost all such settings. For instance, labor markets and dating platforms are all situations where individuals "search" not into deterministic matches, but rather into competitive situations where their final utility is determined by which other similar individuals have searched into the same match. The fact that, in the limit, this set of individuals can be succinctly described given search decisions, is a tremendously powerful analytical tool that can give greater formal insights into situations that would otherwise be difficult to analyze. It should be noted that there are platform settings where this technical difficulty need not exist, i.e., when the utility of each agent is directly determined by the match and is independent of which other agents received the same match.

The positive insights of our present paper are more restricted to auction platforms. While it is generally true that better match quality will generate higher *total* surpluses, whether or not better matches overall take place with lower search costs, and whether or not the platform benefits from these matches, depends crucially on the *distribution* of match surplus between the two sides of the platform. The search decision is based on the surplus of only one side of the market – in the auction platform setting, this is perfectly aligned with the contribution to total surplus, but this need not be the case for platforms more generally. Our negative revenue result, however, is general than the efficiency results, and applies to platforms other than auctions: even in a "best-case" scenario, where search incentives are aligned with total welfare, and where the reallocation of buyers due to search is most likely to be positive, one side of the market can lose out on the whole.

The previous platforms literature has focused almost exclusively on fee structures as the main strategic choice variables of platforms; recent years, however, have shown that infrastructure investments to reduce platform frictions, such as those that enable more efficient search, also constitute an important strategy choice, and it can be an important dimension of differentiation between competing platforms. We have shown that search can affect platform performance in subtle, and possibly unanticipated ways: only in special cases is it true that both sides of a platform are better off with more search, even though this ought to increase the total surplus generated by the platform. An important practical consequence of this is that search policy needs to account not only for anticipated aggregate gains, but also the redistributive effects of search. This redistribution can happen both between the two sides – as happens when sellers lose overall – or within sides – as experienced by the buyers with middling valuations who manage only to face stiffer competition in both markets. In practical scenarios, these losing participants may choose to exit the market, a possibility that we explore in a fuller environment in a following paper.

2. Optimal platform search policy

2.1. Introduction

In our previous work, we found that search within auction platforms can have subtle and somewhat unexpected effects on the properties of equilibria. In contrast to the basic intuition that has motivated improvements in search technology on many platforms – namely, that enabling search within a platform constitutes a "win-win" situation for both buyers and sellers – we found a general tradeoff between better matching and market thickness. Though search might often produce desirable matching effects, within certain parameters it could also have undesirable consequences. A stark illustration of this possibility is provided by the existence of cases where overall seller revenues could decline with greater search, in spite of the fact that goods ended up sold to buyers with higher valuations. The rationale for this is that some sellers on a platform can end up with severely reduced market thickness – the pro-competitive effects of better matching are undone by the loss of buyers to support the price of the good.

Though such cases do certainly exist, our explorations seemed to indicate that they are quite rare; when we assumed symmetry, many of our nice intuitions were restored. Even these positive results, however, might be met with some skepticism in the context of real world platforms, since we studied only the within-platform effects, taking as exogenous the participation of buyers. Most platforms operate in a rather different strategic setting and do not enjoy the luxury of being able to change parameters without worrying about participation effects. Here one might worry that the matching effects generated by search – which generally bode well for sellers – can cut both ways for buyers: while it makes it easier for them to end up in auctions that they want to participate in, it also means that they face stiffer competition when they get there, which may reduce their utilities. If this is the case, it may cost the platform in the form of reduced membership, lower surpluses generated, and lower revenue.

This paper extends our previous work in two ways to analyze this situation: (1) we endogenize buyer participation, which allows for pro-competitive effects of search to be offset by reduced buyer participation; (2) we specify optimal platform search policy under two different revenue instruments, a flat membership fee and fixed tax on transactions. Platform policy in the presence of binding participation constraints becomes much more complex than in the case without: it must decide whether it is better to encourage search and produce better matches for high valuation buyers – at the cost of alienating lower valuation buyers – or if it is better to maximize the breadth of participation. To be able to weigh these effects meaningfully requires explicit formulations of platform objective functions; this paper considers two common instruments, which are the flat membership fee, where all buyers pay a fixed membership cost to participate on the platform, and the transaction tax, where they take a fraction of all sales.

Our model here builds upon the technical insights developed in our earlier work, which allowed us to tractably analyze the effects of search in settings with large numbers of buyers and sellers. The basic modeling insight we use is the observation that sequential search into many-to-one matches generates, in the continuum limit, Poisson games; in our case, these are auctions where the number of buyers is a Poisson random variable. As in our previous work, we model a continuum of buyers and a continuum of sellers, where each seller has one good that is sold in an auction. Participating buyers choose their auctions through a process of sequential search. When all buyers are finished searching, second-price auctions are conducted and goods are sold. The new feature here is that prior to running through the platform mechanism, buyers must choose whether or not they want to participate, which they only do if their expected utilities from participation exceed their reservation utilities, which we assume is the same for all buyers.

In some respects, the present model is more specialized than the previous one; for instance, we assume uniform valuations and a symmetric platform. Our findings do not depend appear to depend on any knife-edge properties of these assumptions, and so we expect them to be generally robust to small asymmetries; however, our proofs of a handful of key results, such as uniqueness and the comparative statics, rely on symmetry to reduce the number of parameters. But on a more general note: we also already know that search may be detrimental to platform revenues in the asymmetric case even without buyer exit effects, so the symmetric case is a natural motivation for the concerns this paper addresses.

Our main results are as follows: (1) we characterize equilibrium according to a set of threshold and show that equilibrium is unique; (2) we derive comparative statics of buyer behavior on platform parameters; (3) we show that our positive revenue result continues to hold even when buyers exit, i.e., total seller revenue is always increasing as search costs decrease, in spite of negative participation effects; (5) we characterize optimal search policy – more specifically, we find that a platform with a transaction tax always sets search costs to zero; a fee-collecting platform always sets an extremal value for search costs, i.e., either at 0 or infinity. Moreover, when the market is sufficiently dense, the fee-collecting platform maximizes revenue by setting search costs to zero; conversely, when the market is sufficiently thin, the optimum is to dissuade search entirely.

The fact that a fee-collecting platform always selects an extremal value comes from the interesting relationship between search costs and participation levels. When search costs are very high, so that there are few buyers searching, and the marginal participant does not search, reducing search costs increases competition and decreases the utility of the marginal participant, so more buyers exit. As search costs approach zero, however, we will arrive at a point where all participating buyers search search: in this case, a reduction in search costs actually increases the utility of the marginal participant, causing membership to rise. Hence, the participation has a V-shaped relationship to search costs; participation is lowest at the point when the marginal participant switches his search behavior.

The relationship between optimal search policy for a fee-collecting platform and market density reflects the changing importance of participation breadth relative to matching effects as overall competitiveness grows. When nobody searches and matching is random, utilities are relatively flat in buyer types: as we move across types, higher valuations for one good are offset by lower valuations in the other, and a good ends up bidding for a good of any given type with equal probability. This means that a platform's ability to raise the fee is limited. When buyers search, however, their utility is more dependent on type, which improves a platform's ability to raise the fee. This comes at the cost, however, of decreased participation to begin with. Hence, the choice between zero and infinite search costs is basically a choice between extracting relatively low surplus from more buyers, or relatively high surplus from fewer. The value of being collecting a higher fee is higher when the market is very dense, because in this case the competition drives down the utility levels of all buyers; on the other hand, participation breadth is more valuable when all buyers experience relatively high utility even without search, which occurs when the market is thin and competition is low.

Although our results seem to paint a generally rosy picture of search effects for platform revenues – lower search costs are always good for tax-collecting platforms, and seem to be good much of the time for fee-collecting ones – the optimal policies for fee-collecting platforms can be quite fragile, since they depend on the platform being able to set search costs that are effectively zero or effectively infinite, if not then we can only tell that the optimal policy should be extremal, but we cannot indicate which extreme. The V-shaped participation curve also means that for fee-collecting platforms, moving search costs in the direction of local increase need not lead to the global optimum.

The presence of these tensions may suggest that there is a sense in which search is ''incompatible'' with membership fees. We interpret these as symptoms of the issue described earlier: a flat membership fee can be set only at the level of the lowest utility, whereas the benefits of lower search costs accrue to those who realize the highest utility; on the other hand, reduced search costs will increase competition and induce exit, which does hurt platform revenues. Hence, a fee-collecting platform might be wary to improve matching, since membership fees are not effective at extracting the additional surplus generated by better matches.

2.2. Literature Review

This paper is a direct application of our previous work, to which it is most closely related. The richer setting, however, allows us to bring the insights of those models to bear more clearly on existing lines of research. In particular, it addresses some of the concerns of the existing platform literature, which has extensively analyzed pricing decisions and market structure for platform economies (Ambrus and Argenziano 2009; Armstrong 2006; Ellison and Fudenberg 2003; Rochet and Tirole 2003, 2006; Rysman 2009; Weyl 2010). In some respects, our setting is simpler than many of the aforementioned models, since we endogenize only one side of the market; however, we offer several new results on how search fits into platform policy in general, and how it fits in with perhaps one of the most important concerns of platforms, which is membership.

Our work is also related to a few recent papers on search in platforms. Tadelis and Zettelmeyer (2011) provide experimental evidence in support of positive matching effects in wholesale auto auctions, which was a large motivation of our first paper. The results derived here also overlap with questions raised by Hagiu and Jullien (2011), who examine platform incentives to divert search. They find that a platform will always choose to divert search, which is different from our findings, which are generally supportive of search, and in any case rule out interior equilibria. These differences may be attributable, however, to significant modeling differences between our model and theirs; their platform changes a ''per-match-attempt'' fee, for instance, which creates incentives to maximize the number of attempts, i.e., to force buyers to search multiple times.

Our search framework extends the existing search literature by providing a way to tractably model search *into a competitive environment*. Most existing search models (including this one) are based on the dynamic search model by Diamond (1971), where consumers search into one-to-one matches and match payoffs are completely determined by the types of the buyer and seller. Our model allows us to model many other sequential search situations where theses properties may not be specified. In particular, we provide a way to model search into many-to-one matches, of which auctions are one possibility. The general framework should be adaptable to any number of other strategic settings; the fundamental trick is to see that sequential search into many-to-one settings can be modeled as a continuum of Poisson games in the limit, the theory to which is developed in Myerson (2000).

Finally, there appears to be a relationship between our results and some findings in the information revelation literature; this is somewhat to be expected, since facilitating search is often almost synonymous with providing information. Lewis and Sappington (1994) study a setting where a seller can provide or deny access, to varying degrees, to private information to consumers regarding their personal tastes, and find that the optimum disclosure level will typically resolve at one of the two extremes. Work on obfuscation of product information by firms Ellison and Wolitzky (e.g., 2009) identifies similar tradeoffs between extracting high surpluses from smaller segmented markets, and extracting lower surpluses from larger broader markets. The analogy is not exact, however, since buyers in our model always discover all relevant information prior to submitting any bid, which may account for differences in the flavor of our results.

2.3. Model

We extend the model in our previous paper by endogenizing buyer participation. Seller participation remains exogenous, since we are primarily interested in the interaction of the matching effect with potential buyer exit induced by increased competition. We also specialize the previous model in several ways: utilities are linear in types, distributions are uniform, and the platform is perfectly symmetric. Linearity and uniformity are primarily for convenience, and we expect most results to hold without them. Symmetry, on the other hand, is crucial to some key portions of the proofs. Asymmetries may generate additional equilibria, although our comparative statics should still hold along equilibrium paths.

The platform outcome is modeled as the equilibrium of a three-stage game, as described below

1. *Participation*: buyers decide whether or not they want to participate in the platform.
2. *Search*: buyers search sequentially over auctions and match up (many-to-one) with sellers, until all buyers are satisfactorily matched
3. *Auction*: each seller sells his good via a second-price auction, held between all buyers that have been matched to him.

Below, we specify the agents and their utilities, and we describe the stages of the game. We present several results from our previous paper here without proofs; the interested reader is directed there for more detailed discussions of the search sub-game details and the specification of the continuum game.

2.3.1. Environment

There is a measure 1 of sellers and a measure μ of buyers. There are two types of goods, A and B; each seller has one good to sell, with $1/2$ of the sellers having A goods and $1/2$ having B goods.

Buyers have varying tastes for the two goods, according to their types: buyers have types $\theta \in \Theta \equiv [0,1]$, which are uniformly distributed, i.e., the density of types at θ is given by $f(\theta) = \mu$ for all $\theta \in \Theta$. We assume that buyers with relatively high valuations for one good have relatively low valuations for the other, so that there are gains to be made from matching buyers to the sellers they like the most. In particular, the valuations of a type θ buyer for A and B goods, respectively, are given by

$$v_a(\theta) = \theta$$

$$v_b(\theta) = 1 - \theta$$

Goods are allocated through second-price single-unit auctions, which each seller conducts individually. Bidders choose which auction to participate in via a process of sequential search.

2.3.2. Search

Buyers find their auctions via a sequential search process, where the search costs are determined by platform-wide parameters, c the cost of searching for another auction, and p, the efficiency of a platform's search technology.

Buyers are initially matched randomly to auctions — hence, each buyer has a $1/2$ probability of initially finding an auction for each type of good. Upon being matched, buyers observe the type of good being auctioned off, but not the number of other buyers who are also matched. Each buyer can then choose to stay with the initial draw, or to search again and incur a cost of c.

If the buyer searches again, then with probability p he is matched with an auction of his preferred type; with probability $1 - p$, the match is again random, which gives $1/2$ probability for each type. After viewing the type of the good being auctioned off in the new draw, a buyer may again either to stay or to draw again at a cost of c. The search stage is completed when all buyers choose to stay.

The platform parameters c and p correspond to different channels that may affect a buyer's search costs: either the per-draw cost c may vary, directly affecting incurred search costs, or the success probability p of any particular draw may change, so that the expected number of times a buyer searches changes. In more general settings, which channel a platform uses to vary search costs will matter: when the measure of A and B sellers is not equal, variation in p will affect the total search costs of buyers who favor A differently from that of those who favor B. We restrict our attention here to the symmetric setting, however, which allows us to collapse total search costs into a single parameter,

$$c^* = \frac{c}{2p^*}$$

where

$$p^* = p + 1/2(1-p)$$

is the probability that a buyer's next draw yields an auction for his preferred type of good.

We will see in our analysis that c^* is sufficient to describe search decisions and outcomes, so that the effects of changes in c or p individually can be analyzed through their effects on c^*. We note here that c^* also has a natural interpretation as the total *ex ante* expected search costs for a buyer who engages in search (to see this, note that $1/p^*$ is the expected number of times a buyer will search if his first draw is not of his preferred type, which happens with probability $1/2$).

In the symmetric setting that we consider here, the sequential search process can be modeled in a "reduced form" search consisting of one stage only, where a buyer can pay c^* and be guaranteed to find an auction of his preferred type, or to pay nothing and be matched randomly. Hence, for our analysis on platform policy, we consider c^* to be the choice variable, and take most of our comparative statics on c^* rather than c and p separately.

2.3.3. Participation

The new feature of this model vis-à-vis our previous paper is that we endogenize buyer participation. In the first stage of the game, buyers must decide whether or not they wish to participate on the platform at all. If buyers do not participate, they receive a utility of s. We use different interpretations of s for each instrument: when the platform collects membership fees, we treat s as the level of the fee, so that a buyer who participates must expect utility higher than the membership fee; when membership fees are not charged, we interpret s as an exogenously determined outside option value.

2.3.4. Platform Objectives

Platforms choose the level of search costs, c^*. We assume for our main results that they are free to choose any $c^* \in [0, \infty)$, but discuss how a limited range (due, for instance, to technological constraints) affect optimal strategies.

Transaction tax platform Since we do not model seller participation, we assume that c^* is the only choice variable for a platform that taxes sales. It collects a fixed portion t of all sales, so its revenue is equal to tR, where R is the total seller revenue. As mentioned previously, such a platform's interests are perfectly aligned with sellers, so it chooses c^* to maximize seller revenue. Furthermore, in this case we consider s to be exogenously fixed.

Fee collecting platform A fee collecting platform derives its revenues from all participating buyers, each of whom pays a flat fee. In addition to c^*, a fee collecting platform sets a level of s to maximize the total fees collected, which is the total measure of *participating* buyers (which depends on the fee level) times the level of the fee.

Aside from characterizing search effects with binding participation constraints, the ultimate goal of this paper is to characterize the optimal platform search policy under different instruments. Search can have rather subtle effects on participation and revenue, so whether or not it is in the interest of a platform to reduce search costs will depend critically on which aspects of the platform outcomes its revenues depend on, i.e., whether it is trying to maximize seller revenues, or whether it is trading off participation levels and surplus extraction by imposing a membership fee.

2.4. Analysis

The equilibrium concept that we use to describe outcomes is SPNE, which we solve for by backward induction. Given a characterization of individual auction participation, we can solve for the utilities that buyers experience in the two types of auctions, and hence their optimal search decisions. Aggregate search decisions then determine the *ex ante* utility that platform participation gives to individual buyers. The final step is to make sure that participating buyers receive sufficient utility.

2.4.1. Random Matching and Auction Utility

We use the random matching platform model from our previous paper. Given a measure of buyers who is randomly matched to measure 1 of sellers, the number of buyers showing up at any particular auction is Poisson distributed with parameter equal to the measure of buyers. We reproduce several of the key results here without discussion; a more thorough discussion is provided in our previous paper.

Lemma 2.1. *Let $f_i(\theta)$ denote the density of buyers of type θ in a platform for identical goods, where density is normalized to a measure 1 of sellers, and suppose that buyers are randomly matched to sellers. Then for any Borel-measurable subset $T \subset \boldsymbol{\theta}$ of types, the number of buyers with $\theta \in T$ who shows up in any particular auction is Poisson distributed with parameter $\int_T f_i(\theta)\, d\theta$.*

Furthermore, the probability that the highest bid in an auction is below v is given by $G(v) = e^{-\eta(v)}$, where $\eta(v)$ is the measure of total buyers with valuations above v.

The distribution of the first-order statistic allow us to derive buyer utilities, using standard arguments (details can be found in most auction theory texts, e.g., Krishna 2009).

Lemma 2.2. *Consider a buyer of type θ. His utility from participating in an A auction is*

$$u_a(\theta) = \int_0^\theta G_a(v)\, dv$$

And his utility from participating in a B auction is

$$u_b(\theta) = \int_0^{1-\theta} G_b(v)\, dv$$

where

$$G_a(\theta) = e^{-\eta_a(\theta)} \quad \eta_a(\theta) = \int_\theta^1 f_a(t)\, dt$$

$$G_b(\theta) = e^{-\eta_b(\theta)} \quad \eta_b(\theta) = \int_0^1 f_b(t)\, dt$$

where $\eta_i(\theta)$ denotes the measure of buyers with valuations above that of θ participating in auctions of type i.

Note that the measures $\eta_i(\theta)$ are endogenously determined by search behavior. When buyers of type θ do not participate in auctions of type i, their density is zero; when they actively search into auctions of that type, their density is 2μ, since all buyers of that type end up in auctions of type i, who have a platform-wide density of μ, normalized to the measure 1/2 of sellers. Otherwise,

$\mu/2$ buyers show up, normalized to a measure $1/2$ of sellers, so their density is μ. Formally,

$$f_i(\theta) = \begin{cases} 0 & \text{if type } \theta \text{ buyers do not participate in } i \text{ auctions} \\ 2\mu & \text{if type } \theta \text{ buyers search into } i \text{ auctions} \\ \mu & \text{otherwise} \end{cases}$$

2.4.2. Search Decision

The search model is sequential search, which is a simple dynamic programming model that we solve by writing out the Bellman equation.

Suppose a buyer has drawn into a B auction initially. Let u_a and u_b denote his utilities from A and B auctions respectively, and let V_a denote his continuation utility from continuing to search. Since his problem is stationary, if it is optimal for him to continue searching this period, it will also be optimal to continue searching next period if he draws B again, and so on until he ends up in an A auction.

He searches if

$$V_a - c \geq u_b$$

If he continues, with probability p^* he receives u_a; otherwise he pays c again and receives the continuation utility. Hence

$$V_a = p^* + (1 - p^*)(V_a - c)$$

which rearranges to

$$V_a = u_a + \frac{1 - p^*}{*}$$

Substituting into decision rule above and rearranging gives

$$\frac{1}{2}\left(u_a - u_b\right) \geq \frac{c}{2p^*} = c^*$$

The analysis is identical for searching into B auctions; we present the decision rule in the following lemma.

Lemma 2.3 (Search Decision). *A buyer of type θ searches into A if*

$$\frac{1}{2}\left[u_a(\theta) - u_b(\theta)\right] \geq c^*$$

and into B if

$$\frac{1}{2}\left[u_a(\theta) - u_b(\theta)\right] \leq -c^*$$

Otherwise the buyer does not search.

As a conceptual aside, note that this is precisely the decision rule of buyer in a reduced search model who can either pay c^* to guarantee an A auction, or pay nothing and be randomly matched. Search requires that

$$u_a(\theta) - \frac{1}{2}\left[u_a(\theta) + u_b(\theta)\right] \geq c^*$$

which is exactly the sequential search condition.

2.4.3. Participation Decision

The participation decision is straightforward: a buyer participates on the platform only if his utility is greater than s.

Lemma 2.4. *Buyers participate if and only if* $u(\theta) \geq s$.

2.4.4. Equilibrium

The decision rules above allow us to characterize equilibrium in terms of a set of threshold types. To see this, note that

1. $u(\cdot)$ is convex, which means that the set $\{\theta : u(\theta) < s\}$ is convex, i.e., an interval (θ_u, θ_v) for any given s.

2. $u_a(\theta) - u_b(\theta)$ is increasing with θ.

Since in equilibrium, all buyers who receive $u(\theta) < s$ do not participate in the platform, and we know that this set is always an interval, we can characterize participation decisions with two threshold values, θ_u and θ_v with $\theta_u \leq \theta_v$ such that buyers of type $\theta \in (\theta_u, \theta_v)$ do not participate, and $u(\theta_u) = u(\theta_v) = s$.

The second condition tells us that search decisions can likewise be characterized by a set of thresholds (θ_b, θ_a) such that all $\theta < \theta_b$ search into B, $\theta > \theta_a$ search into A, and the rest do not search.

Theorem 2.1. *The equilibrium is unique and symmetric.*

The number of free parameters makes it a little tricky to establish the above result – there are four parameters that have different effects on the equilibrium conditions. Our approach will be to prove symmetry of equilibria first, which allows us to characterize all equilibria using only two thresholds. We can then show that the comparative statics of any equilibrium are well-defined, so that each equilibrium must have a unique path as we vary s. Our previous paper established that there is a unique equilibrium when $s = 0$, so the we know that equilibrium must also be unique for any given $s > 0$.

Lemma 2.5. *Any equilibrium is symmetric.*

Proof. We first consider the case where $\theta_b < \theta_u \leq \theta_v < \theta_a$. For convenience, we let

$$s_a(\theta) = u_a(\theta) - u_b(\theta)$$

If we consider parametrized equilibrium paths that vary by s, note that each path must contain some set of equilibria fulfilling these conditions, since as we transition from full participation, the first exits must lie between the search thresholds (the buyers earning the least utility at full participation are those not searching).

Symmetry of equilibrium would mean $\theta_a = 1 - \theta_b$ and $\theta_v = 1 - \theta_u$. Suppose not. First, $u(\theta_u) = u(\theta_v) = s$. Since there are no participants between the θ_u and θ_v, their winning probabilities in each market are the same: $G_a(\theta_u) = G_a(\theta_v)$ and $G_b(\theta_u) = G_b(\theta_v)$. Since we know that $u_i(\theta) = \int_0^{v_i(\theta)} G_i(v)\, dv$, this implies that

$$u_a(\theta_v) = u_a(\theta_u) + (\theta_v - \theta_u)G_a(\theta_u)$$

and

$$u_b(\theta_u) = u_b(\theta_v) + (\theta_v - \theta_u)G_b(\theta_u)$$

Substituting into $u(\theta_u)$ and $u(\theta_v)$ gives

$$u(\theta_v) = \frac{1}{2}\left[u_a(\theta_u) + u_b(\theta_v) + (\theta_v - \theta_u)G_a(\theta_u)\right]$$

and

$$u(\theta_u) = \frac{1}{2}\left[u_a(\theta_u) + u_b(\theta_v) + (\theta_v - \theta_u)G_b(\theta_u)\right]$$

Since these two must be equal, we have that $G_a(\theta_u) = G_b(\theta_u)$, and consequently $u_a(\theta_v) = u_b(\theta_u)$ and $u_b(\theta_u) = u_b(\theta_v)$. This means that

$$s_a(\theta_u) = -s_a(\theta_v)$$

Furthermore,

$$\frac{d}{d\theta}u_i(\theta) = G_i(v)\frac{d}{d\theta}v_i(\theta)$$

A consequence of this is that s_a changes at the same rate going down from θ_u and going up from θ_v. Consider again the equilibrium conditions for search thresholds:

$$s_a(\theta_b) = -\frac{c}{p^*}$$

and

$$s_a(\theta_a) = \frac{c}{p^*}$$

We get that $\theta_a - \theta_v = \theta_u - \theta_b$. Since we know also that $G_a(\theta_u) = G_b(\theta_u)$, $\theta_b + \theta_u = 2 - \theta_a - \theta_v$, which together imply that $\theta_a = 1 - \theta_b$ and $\theta_v = 1 - \theta_u$, as desired.

The final step is to rule out the possibility that $\theta_b < \theta_u < \theta_a < \theta_v$, i.e., one marginal participant searches while the other does not.

We show that the equilibrium conditions together with $\theta_b < \theta_u < \theta_a < \theta_v$ are mutually incompatible. Since θ_a does not participate, $u(\theta_a) < s$. At θ_b and θ_a, buyers are indifferent between searching and not searching, meaning that their utilities are equal to their utilities from not searching. Hence we have that

$$u(\theta_b) = \frac{1}{2}\left[u_a(\theta_b) + u_b(\theta_b)\right] > s > \frac{1}{2}\left[u_a(\theta_a) + u_b(\theta_a)\right] = u(\theta_a)$$

which, rearranging terms, yields

$$u_b(\theta_b) - u_b(\theta_a) > u_a(\theta_a) - u_a(\theta_b) \tag{2.1}$$

The search conditions imply that

$$u_b(\theta_b) - u_a(\theta_b) = \frac{c}{p^*} = u_a(\theta_a) - u_b(\theta_a)$$

which can be rearranged as

$$u_b(\theta_b) + u_b(\theta_a) = u_a(\theta_b) + u_a(\theta_a) \tag{2.2}$$

Subtraction equation (2.2) from (2.1) and simplifying gives that

$$u_b(\theta_a) < u_a(\theta_b)$$

or, expanding u_i and using the fact that θ_a and θ_b are the lowest valuation buyers in the B and A markets, respectively,

$$\theta_b e^{-\mu_b} < (1-\theta_a)e^{-\mu_a} \tag{2.3}$$

If we let μ_0 be the measure of buyers who participate but do not search, we know that $\mu_b = 2\mu\theta_b + \mu_0$ and $\mu_a = 2\mu(1-\theta_a) + \mu_0$, which means the above expression can only be satisfied if $\mu_b > \mu_a$.

At the same time, we know that $u(\theta_u) = s > u(\theta_a)$ since θ_a does not participate. Since there are no participants between the θ_u and θ_a, their winning probabilities within that range do not change, so we can write their utilities as follows:

$$u_a(\theta_a) = u_a(\theta_u) + (\theta_a - \theta_u)G_a(\theta_u)$$

$$u_b(\theta_u) = u_b(\theta_a) + (\theta_a - \theta_u)G_b(\theta_u)$$

which, together with

$$u(\theta_u) = \frac{1}{2}\left[u_a(\theta_u) + u_b(\theta_u)\right] = s > \frac{1}{2}\left[u_a(\theta_a) + u_b(\theta_a)\right] = u(\theta_a)$$

imply that

$$G_b(\theta_u) > G_a(\theta_a)$$

which can only be true if $\mu_b < \mu_a$. This contradicts the above finding that $\mu_b > \mu_a$; hence, there can be no asymmetric equilibria. □

Lemma 2.6. *Equilibrium is unique*

Proof. This follows from a result in the next section, that the equilibrium path as we vary s is unique, since the Jacobian of a function $\mathbf{F}$, where $\mathbf{F}(\mathbf{x}) = (s, c^*)$ defines equilibrium, is non-singular. This means that there is a one-to-one mapping between equilibria for some $s > 0$ and $s = 0$, and since there is only one equilibrium at $s = 0$, there can only be one equilibrium at given $s > 0$ (for given c^*). The missing step is proven below in Theorem 2.2 □

2.4.5. Comparative Statics for Buyer Behavior

In this section, we establish how buyer behavior changes with changes in the parameters c^* and s. There are be two different types of equilibrium: (1) the marginal participant does not search, and (2) the marginal participant searches; how changes in parameters affect outcomes in these two situations will be different.

Before formally presenting the results, let us consider the incentives for searching and how participation and search behavior interact. More buyers search when the marginal searcher either sees an increase in the utility difference between the two types of auctions, or a decrease in the search costs. Analogously, more buyers participate when the marginal *participant* experiences either an increase in his platform utility or a decrease in his reservation utility.

When the marginal participant searches, i.e., $\theta_u < \theta_b$, then changes in the search threshold do not affect utilities or behavior of any of the participants. However, it also means that the marginal participant's utility is directly affected by changes in search costs, which he pays. Hence, as search costs decrease, participation must also increase, since the utility of the marginal participant decreases. It is also straightforward to see that as reservation utility increases, participation must decrease.

When the marginal participant does not search, the interaction between search and participation is a bit more interesting. In particular, a decrease in search costs induce more buyers to search, which raises the competitiveness of both auctions; the increased competitiveness mean that the marginal participant, who does not search, experiences reduced utility. More buyers then exit, since the previous marginal participant no longer achieves his reservation utility on the platform. Similarly, as reservation utility increases, both markets become less competitive, but this will have a larger effect on each buyer's utility in his preferred market; in other words, fewer buyers will increase the size of the difference between utilities of the two different types of auctions, which induces more search. Hence, when the marginal buyer does not search, search and participation will generally move in opposite directions: more search means more competition and less participation.

First we consider the case where both marginal participants are non-searchers, which corresponds to $\theta_b < \theta_u < \theta_v < \theta_a$ above. As we have established, equilibria of this form are unique and symmetric, so that we may characterize them by the pair (θ_b, θ_u) with $0 < \theta_b < \theta_u < 1/2$.

Theorem 2.2. *When the marginal participant does not search (i.e., when $\theta_b < \theta_u$) An increase in the reservation utility leads to lower participation and higher search. An increase in search costs leads to higher participation and lower search. In other words*

$$\frac{d\theta_b}{ds} > 0 \quad \frac{d\theta_b}{dc} < 0$$
$$\frac{d\theta_u}{ds} < 0 \quad \frac{d\theta_u}{dc} > 0$$

Proof. The proof is by construction: in our setting, it is possible to derive explicit expressions for the derivatives above by differentiating the equilibrium conditions. Recall the equilibrium conditions:

$$u(\theta_u) = \frac{1}{2}\left[u_a(\theta_u) + u_b(\theta_u)\right] = s$$
$$\frac{1}{2}\left(u_a(\theta_a) - u_b(\theta_a)\right) = c^*$$

Symmetry guarantees that $u_a(\theta) = u_b(1-\theta)$; we know that $u_a(\theta) = \int_0^\theta G_a(v)\,dv$; and since there are no participants between θ_u and $1-\theta_u$, $G_a(\theta) = G_a(\theta_u)$ for all $\theta \in (\theta_u, 1-\theta_u)$. If we substitute these into the equilibrium conditions, and divide the second condition by 2, we get the following expressions:

$$\theta_b e^{-\mu(2\theta_u)} + \frac{1}{\mu}\left[e^{-\mu(\theta_b+\theta_u)} - e^{-\mu(2\theta_u)}\right] + \frac{1}{2}\left(1-2\theta_u\right)e^{-\mu(\theta_b+\theta_u)} = s$$

and

$$\frac{1}{2\mu}\left[e^{-\mu(2\theta_b)} - e^{-\mu(2\theta_u)}\right] + \frac{1}{2}\left(1-2\theta_u\right)e^{-\mu(\theta_b+\theta_u)} = c^*$$

Denote the left hand sides of the above equations by $F_1(\theta_b, \theta_u)$ and $F_2(\theta_b, \theta_u)$ respectively, and let $\mathbf{F}(\boldsymbol{\theta}) = (F_1(\boldsymbol{\theta}), F_2(\boldsymbol{\theta}))$. We can rewrite the above conditions in vector form as

$$\mathbf{F}(\boldsymbol{\theta}) = \begin{pmatrix} s \\ c^* \end{pmatrix}$$

Let $\mathbf{p} = (s, c^*)$ denote the right hand side of the above expression; then the above implicitly defines $\boldsymbol{\theta}$ as a function of $\mathbf{p}$. By the inverse function theorem, we get

$$(D\boldsymbol{\theta})(\mathbf{p}) = J^{-1}(\boldsymbol{\theta})$$

In other words

$$\begin{pmatrix} \frac{d\theta_b}{ds} & \frac{d\theta_b}{dc^*} \\ \frac{d\theta_u}{ds} & \frac{d\theta_u}{dc^*} \end{pmatrix} = J^{-1}(\mathbf{F})$$

Since J is a two-by-two matrix, the inverse can be readily calculated as

$$J^{-1} = \frac{1}{|J|} \begin{pmatrix} \frac{\partial F_2}{\partial \theta_u} & -\frac{\partial F_1}{\partial \theta_u} \\ -\frac{\partial F_2}{\partial \theta_b} & \frac{\partial F_1}{\partial \theta_b} \end{pmatrix}$$

Differentiating F_1 and F_2 yields

$$\frac{\partial F_1}{\partial \theta_b} = -\left[e^{-\mu(\theta_b+\theta_u)} - e^{-\mu(2\theta_u)} + \frac{1}{2}\mu(1-2\theta_u)e^{-\mu(\theta_b+\theta_u)}\right]$$
$$\frac{\partial F_1}{\partial \theta_u} = -\left[2\left(e^{-\mu(\theta_b+\theta_u)} - e^{-\mu(2\theta_u)}\right) + \frac{1}{2}\mu(1-2\theta_u)e^{-\mu(\theta_b+\theta_u)} + \mu(2\theta_b)e^{-\mu(2\theta_u)}\right]$$
$$\frac{\partial F_2}{\partial \theta_b} = -\left[e^{-\mu(2\theta_b)} + \frac{1}{2}\mu(1-2\theta_u)e^{-\mu(\theta_b+\theta_u)}\right]$$
$$\frac{\partial F_2}{\partial \theta_u} = -\left[e^{-\mu(\theta_b+\theta_u)} - e^{-2\theta_u} + \frac{1}{2}\mu(1-2\theta_u)e^{-\mu(\theta_b+\theta_u)}\right]$$

Note that all the partial derivatives are less than zero, and

$$|pd(F_1, \theta_u)| > |pd(F_1, \theta_b)| = |pd(F_2, \theta_u)|$$

$$|pd(F_2, \theta_b)| > |pd(F_1, \theta_b)| = |pd(F_2, \theta_u)|$$

Consequently the Jacobian is negative:

$$|J| = pd(F_1, \theta_b)pd(F_2, \theta_u) - pd(F_1, \theta_u)pd(F_2, \theta_b) < 0$$

And the inverse Jacobian, which is equal to the derivatives of $\boldsymbol{\theta}$ with respect to (s, c^*), is signed as stated in the theorem. □

Next, we consider equilibria where the marginal participant searches.

Theorem 2.3. *When the marginal participant searches, participation increases with decreases in either search costs or participation costs. That is*

$$d(\theta_u, s) < 0 \quad d(\theta_u, c^*) < 0$$

Proof. Since the marginal searcher does not participate in the market, marginal changes in search behavior have no effect on the behavior of platform participants, who all search into their preferred markets; hence, equilibrium can be fully characterized by θ_u, and is described by the condition

$$u(\theta_v) = u(1-\theta_u) = u_a(1-\theta_u) - c^* = s$$

A buyer with type $1-\theta_u$ has the lowest valuation in the A market, his utility is given by $u_a(\theta_u) = (1-\theta_u)G_a(\theta_u)$, and $G_a(\theta_u) = e^{\mu(2\theta_u)}$, where the exponent is the measure of all buyers in the A market with valuations higher than θ_u, which in this case is all buyers in A, who have measure $\mu 2\theta_u$. Hence the equilibrium expression is

$$(1-\theta_u)e^{-\mu(2\theta_u)} - c^* = s$$

Differentiating with respect to c^* gives

$$-(1+2\mu(1-\theta_u))e^{-\mu 2\theta_u}\frac{d\theta_u}{dc^*} - 1 = 0$$

which rearranges to

$$\frac{d\theta_u}{dc^*} = -\frac{e^{\mu 2\theta_u}}{1+2\mu(1-\theta_u)} < 0$$

A similar procedure with c^* gives

$$\frac{d\theta_u}{ds} = -\frac{e^{\mu 2\theta_u}}{1+2\mu(1-\theta_u)} < 0$$

which completes the proof. □

2.4.6. Comparative Statics for Seller Revenue

Our previous paper provided a positive revenue result in a symmetric setting with exogenous participation: as search costs decrease, improved matching always led to increased seller revenues. This section shows that this result also holds up in our present model, when participation constraints are binding.

The intuition behind a positive revenue result is that more search leads to higher valuation buyers, and higher valuation buyers will pay more for their goods. This result was straightforward with exogenous participation, since symmetry could be leveraged to show that the distribution of the number of buyers in any given auction did not change with search, but the distribution of valuations increased in the sense of first order stochastic dominance. Hence, with exogenous participation, symmetry precluded any negative effects for sellers coming from search.

With a binding participation constraint, the effect of increased search is a bit more nuanced, since we must now balance potentially cross-cutting effects, which happens when the marginal participant is a non-searcher. When the marginal buyer is a searcher, the effect is straightforward, since decreasing search costs also increases participation, which has an unambiguously positive effect on seller revenues.

On the other hand, when the marginal participant is a non-searcher, the effects need to be tracked more carefully, since increased competitiveness drives buyers out of the market. The remainder of the section derives the revenue effect; our model produces a positive effect of search on revenue despite the binding budget constraint, i.e., the gains from matching always exceed the revenue lost from buyer exit.

First, we derive an explicit expression for revenue, which is expressed in the lemma below; the main theorem of this section then states the result.

Lemma 2.7. *Since the platform is symmetric, revenue in each market is equal, and the revenue of the platform as a whole is equal to the revenue in each market. The expression for revenue is*

$$R = 1 - u_a(1) - \int_0^1 u_a(\theta) f_a(\theta)\, d\theta$$

where $u_a(\theta)$ is the utility of a type θ buyer from participating in an A auction, and $f_a(\theta)$ is the density of buyers of type θ in A markets.

Proof. From symmetry it follows that the expected revenue of sellers in each market is the same. Since there are measure 1/2 of each type of seller, the total platform revenue is $R = 1/2R_a + 1/2R_b = R_a$, as desired.

The rest focuses on the A market. The expected utility of a type θ buyer in the A market is the expected value of his winnings, $\theta G_a(\theta)$, subtracted his expected payment, or $u_a(\theta) = \theta G_a(\theta) - m_a(\theta)$, which can be rearranged to give

$$m_a(\theta) = \theta G_a(\theta) - u_a(\theta)$$

The total revenue is just the total of all expected payments, which in this case means that we integrate over the measure of all buyers in the A market:

$$\begin{aligned} R_a &= \int_0^1 \left(\theta G_a(\theta) - u_a(\theta)\right) f(\theta)\, d\theta \\ &= \int_0^2 \theta f_a(\theta) G_a(\theta)\, d\theta - \int_0^1 u_a(\theta) f_a(\theta)\, d\theta \end{aligned}$$

Since $G_a(\theta) = e^{-(\mu - F_a(\theta))}$, $f_a(\theta)G_a(\theta) = g_a(\theta)$, which means that the first term is simply the expectation of the highest bid. This also happens to equal the expected payment of the highest type, which is $1 - u_a(1)$. Substituting this in gives the desired expression. □

Theorem 2.4 (Revenue). *Revenues increase as search costs decrease.*

Sketch of proof. We present the full proof in the appendix. As noted earlier, the statement is clearly locally true for low search costs, when all participants are searchers. We then just need to show it is also true when the marginal participant is a non-searcher. The derivative of revenue with respect to c^* can be signed algebraically, but the process is quite involved. □

2.4.7. Platform Search Policy

With the above results in place, we can turn our attention to the question of optimal platform search policy. As mentioned earlier, we consider the optimum policies for platforms with two different revenue instruments: (1) a platform that generates its revenues by taxing transactions, i.e., taking a fixed percentage of seller revenues; and (2) a platform that generates its revenues by collecting membership fees.

A transaction taxing platform has only search costs as a strategic variable, and it selects costs to maximize seller revenue. Hence, the optimal search policy for such a platform follows directly from the revenue result.

Corollary 2.1. *A platform that levies a transaction tax will choose the minimal possible search costs.*

Theorem 2.5. *A platform collecting membership costs will choose an extremal value of search costs: it will set search costs either as high or as low as possible.*

Proof. Whatever the optimal pair (s, c^*), it must be the case that c^* maximizes participation given s. If not, revenue could be increased by keeping s the same and setting c^* to the participation-maximizing value. We know from above that when participation is first decreasing with c^*, then increasing; hence the optimal c^* must be at one of the extremes. □

In practice, it may not be the case that a platform is able to reduce search costs all the way to zero, or to increase it to a high enough level that search is entirely dissuaded, and given the shape of the participation-search-cost curve, the optimal values will generally depend on the range over which a platform can actually manipulate search costs. For the full range, however, we need only compare the optimal full search and no search equilibria, which allows us to pin down which of the two extremes will be optimal in select situations. More specifically, for μ high enough, $c^* = 0$ is optimal, whereas for μ low enough, infinite search costs, or completely dissuading search, is optimal.

Theorem 2.6. *There exists $\underline{\mu}$ such that for $\mu < \underline{\mu}$, the optimal search policy is to set search costs high enough to dissuade search completely.*

Additionally, there exists $\bar{\mu}$ such that for $\mu > \bar{\mu}$, the optimal search policy sets $c^* = 0$, i.e., all members of the platform search.

Proof. Note that total membership fees collected is equal to $\mu 2\theta_u s$, maximizing which is equivalent to maximizing $\mu\theta_u s$, which we denote here with R. The utility of the marginal participant is always equal to s, so we can write revenue as a function of θ_u, transforming the problem of setting s to one of setting θ_u.

At zero search costs, this is given by $R_0(\theta)$ as

$$R_0(\theta) = \mu\theta s_0 = \mu\theta(1-\theta)e^{-\mu 2\theta}$$

At infinite search costs,

$$\begin{aligned} R_1(\theta) = \mu\theta s_1 &= \mu\theta\left[\frac{1}{\mu}\left(e^{-\mu\theta} - e^{-\mu 2\theta}\right) + \frac{1}{2}(1-2\theta)e^{-\mu\theta}\right] \\ &= \theta\left(e^{-\mu\theta} - e^{-\mu 2\theta}\right) + \frac{1}{2}\mu\theta(1-2\theta)e^{-\mu\theta} \end{aligned}$$

First, we construct a $\bar{\mu}$ that satisfies the property given; note that it is not the lowest possible value for $\bar{\mu}$.

The way that we do this is to show that for certain θ_u corresponding to infinite search cost equilibria, $\theta_u/2$ in a zero search cost equilibrium yields greater revenue. Then, as long as μ is great enough, the optimal θ_u under infinite search costs is one of those for which $\theta_u/2$ offers greater revenue, so that the optimal revenue under zero search costs must exceed the optimal revenue under infinite search costs.

To see this, substitute $\theta_u/2$ into R_0:

$$R_0\left(\frac{\theta_u}{2}\right) = \frac{1}{2}\mu\theta_u\left(1 - \frac{1}{2}\theta_u\right)e^{-\mu\theta_u}$$

This exceeds $R_1(\theta_u)$ as long as

$$\frac{3}{2}\theta_u > \frac{2}{\mu}\left(1 - e^{-\mu\theta_u}\right)$$

or when

$$\mu\theta_u > W\left(-\frac{4}{3}e^{-4/3}\right) + \frac{4}{3}$$

where $W(\cdot)$ is the Lambert W function, i.e., $W(x)$ satisfies $W(x)e^{-W(x)} = x$.

The optimal θ_u under zero search will satisfy the above if the θ_u satisfying the above condition with equality is below the maximum revenue, i.e., revenue is increasing it that point. The derivative of R_1 with respect to θ_u is

$$\frac{dR_1}{d\theta_u} = e^{-\mu\theta_u} - e^{-\mu 2\theta_u} - \mu\theta_u\left(e^{-\mu\theta_u} - e^{-\mu 2\theta_u}\right) + \left(\frac{1}{2}\mu - 2\mu\theta_u - \frac{1}{2}\mu^2\theta_u + \mu^2\theta_u^2\right)e^{-\mu\theta_u}$$

Letting $w \equiv W(-3/4e^{-4/3})$ and $\mu\theta_u = w + 4/3$, we can show that

$$e^{-\mu\theta_u} = -\frac{3}{4}w$$

We can substitute this into our requirement that $dR_1/d\theta_u > 0$ and rearrange terms to get

$$\mu > \frac{(w + 11/6)(w + 3/4)}{1 - (w + 4/3)} \cong 1.69989$$

To show that for low enough μ, no search is always optimal. From the above expression for R_0, we can actually find the optimal level of R_0 by setting $R_0'(\theta^*) = 0$, which yields

$$\theta^* = \frac{1}{2\mu}\left[1 + \mu - \sqrt{1 + \mu^2}\right]$$

Plugging back in to R_0 gives the explicit value for R_0^*:

$$R_0^* = \frac{1}{2}\left(\frac{\sqrt{1 + \mu^2} - 1}{\mu}\right)e^{-(\mu+1-\sqrt{1+\mu^2})}$$

We show that for small enough μ, the no search equilibrium with threshold at 1/2 beats this. Call this value R_1^*,then

$$R_1^* = \frac{1}{2}(e^{-\frac{1}{2}\mu} - e^{-\mu})$$

Since R_0^* and R_1^* are both zero at $\mu = 0$ (taking the limit using L'Hôpital's rule for R_0^*), it suffices to show that the derivative of R_1^* with respect to μ^* exceeds that of R_0^* at $\mu = 0$.

For R_1^*, this derivative is equal to 1/4. To find the corresponding derivative of R_0^*, note that 1, μ, and $\sqrt{1 + \mu^2}$ correspond to sides of a right triangle with hypotenuse $\sqrt{1 + \mu^2}$, so if we let $\rho = \arcsin \mu$, we can rewrite the term in parentheses as

$$\sqrt{1 + \mu^2}(1 - \cos \rho)$$

Taking the derivative of R_0^* at $\mu = 0$, which corresponds to $\rho = 0$, yields $1/2 \sin 0 = 0$. Hence, there must be a region around zero where platform revenues with no search are strictly greater than revenues under full search. □

In spite of these results that seem to establish the benefits of search even for a fee-extracting, they should be interpreted with caution. As mentioned above, they are sensitive to the range over which the platform can actually change search costs; a small technological restriction may make it optimal to set costs as high as possible. If we fix s, then platform revenue is similarly V-shaped with search costs. This is somewhat disconcerting, and suggests that the revenue instrument here, membership fees, fails to pick up on the additional surplus generated by additional search – which is the whole point that a platform would want to facilitate search in the first place.

2.5. Conclusion

This paper takes a first step in bringing the insights from our previous work into broader discussions of platform strategy. Many platforms invest heavily in reducing search frictions, whether this be through more sophisticated search services, or through recommendation engines, or through information provision; it is not very well-known, however, whether we should always expect better outcomes in such situations. Our earlier work suggested that in symmetric settings, lower search costs should have an unambiguous benefit; even that finding, however, needs to be tempered by the fact that the potential downsides of search, which stem from the increase in competition between buyers, are heavily dampened by the fact that participation is exogenous.

We find that the positive revenue result continues to hold in our current model with endogenous participation. This is an encouraging result, and suggests that the benefits to better matching can be very significant, since they must be at least as great as the losses incurred by exit. We also present new results on optimal platform search policy. In particular, a taxing platform will always gain by decreasing search costs; a fee-collecting platform, however, has an optimal policy that depends on the thickness of the market; in thick markets with high levels of competition, zero search costs are optimal, whereas in thin markets with low competition, infinite search costs do better.

Furthermore, we find that the optima for fee-collecting platforms are more fragile than those for tax-collecting ones. When a platform does not have access to all values of c^* – if, for instance, technological impose a binding lower bound on c^* where some remaining participants do not search – then optimal search policy will further depend on the range available. The only guarantee is that the optimal policy will be extremal, but which of the two extremes generates higher revenues cannot be determined in general – it really depends on the precise definitions of ‘‘as much as possible’’.

Ultimately, our results are generally supportive of the ability of search to improve platform revenue: revenue is always increasing as search costs decrease; and for competitive enough platforms, zero search costs are also optimal for a fee-collecting platform. We leave, however, with a cautionary remark. The fragility of our the optima for fee-collecting platforms, along with the non-monotonicities of revenue as a function of search costs, are a strong hint that a membership fee is simply the wrong instrument for extracting any additional surplus created by search. The intuition is fairly straightforward: a fee extracts a level of surplus from each participant equal to that of the lowest-utility buyer on the platform, whereas matching generates additional surplus only for high-utility buyers. Hence, even if there are significant, positive matching effects, a platform that collects membership fees may not see them. A tax, by contrast, moves with revenue, and hence captures at least partially the benefits of better matches, i.e., in a manner of speaking, revenue is able to "pick up" on the additional surplus, since it manifests itself at least partially in the form of higher prices. From a broader perspective, we might posit that there are many situations in which search can increase the total surplus generated on a platform; whether or not that facilitating that search is good for the platform itself, however, may depend on whether or not the platform's instrument is sensitive to the matching effects created.

3. Arguing about politics: An all-pay model of political campaigns

3.1. Introduction

Campaigning has become an increasingly important – and increasingly expensive – reality of political competition in the United States. Total campaign expenditures have been growing rapidly over the course of the last decade: in 2000, total campaign contributions amounted to $528 million; in 2008, this figure was over $1.7 billion[1] Recent institutional developments, notably the controversial *Citizens United v. Federal Elections Commission* and consequent emergence of so-called "Super PACs", have also sparked lively debate on the role of money in American political competition, raising the question of what rapidly increasing campaign spending levels portend for the future of American politics. Are large campaign expenditures merely the price to pay for a smoothly functioning democracy in a large developed country, i.e., the price of maintaining an informed electorate? Or are they evidence of an emerging *de facto* plutocracy that is increasingly able to game the democratic system with its superior resources?

Even though campaigning is perhaps the most conspicuous aspect of political competition, there has been little consensus on either the purpose or the effects of campaigning and campaign spending – a fact that is underscored by several puzzling but persistent empirical findings. For instance, incumbent spending seems to have negligible effects on electoral outcomes, which stands in uneasy tension with the large sums of money that are poured into incumbent campaigns. How campaigning affects political outcomes also bears significantly on the issue of campaign finance reforms: what types of effects should we expect to see with campaign finance liberalization, of the type implemented by *Citizens United*, and how will it affect the quality of electoral outcomes? What types of spending reform make sense, and how are they likely to affect electoral outcomes?

[1] http://www.opensecrets.org

This paper presents a model of political campaigning as an all-pay contest between two candidates, where the prize is the median voter, and hence the election. Each candidate releases some number of arguments in support of his position, and whoever releases the highest number of arguments wins the contest. Constraints on candidate actions come in two forms: (1) candidates have costs that are functions of the numbers of arguments released, and we assume that the incumbent has a cost advantage; (2) the maximum number of arguments each candidate can release is given by his type, which is private information. A more detailed discussion of our modeling choices follows below; the central idea, however, is to capture two aspects of campaign spending that are often at odds in formal models. The first is that campaign spending may be good, since they provide information to the public; the second is that there is some inefficiency in the processing of information, so that spending also has direct effects on behavior and may be abused to create skewed outcomes.

The main theorem of this paper characterizes equilibria. Although there are multiple equilibria, the *ex ante* distribution of candidate actions is unique[2], so that it is possible to describe *ex ante* winning probabilities as well as expected campaigning levels. To be more specific: (1) a cost advantage leads the incumbent to campaign more, and to win with higher probability *ex ante* – we call this advantage the *bias* of the election; (2) a greater cost advantage worsens the maximal informational efficiency of equilibria; (3) equilibria involve mixed strategies, so that over the range of mixing, the marginal effect of additional arguments for each candidate is equal to his marginal cost of producing them; this has implications for several empirical findings, which we discuss later in the section.

We also derive results on the impact of campaign finance policies on electoral outcomes. The ideal competition would be between candidates with identical cost functions, since this eliminates bias and can produce fully efficient outcomes; this is not realistically achievable, however, since there is no obvious direct way to fine-tune the cost functions. First, we consider what the impact of the *Citizens United* decision is likely to have: we represent this as a reduction in both candidates cost functions, but a potential worsening of the asymmetry (i.e., the advantaged candidate gets a larger reduction). The effects of this are ambiguous: candidates spend more, and bias

[2] Up to a set of measure zero

and inefficiency are reduced for competitions between lower types, but greater for competitions involving higher types. That is, such a change is changes the efficiency properties of the equilibria, but its measured impact on winning probabilities might not be very great, since the changes partially cancel. Second, we consider spending caps, which turn out to have rather subtle effects. Properly chosen, they can reduce the bias and improve informational efficiency in the presence of great enough asymmetries; if they are too high, however, but still binding, they do nothing to improve the bias while degrading informational efficiency. The potentially negative effect on informational efficiency is similar to predictions of signaling models of campaigning (Daley and Snowberg 2007), but the relationship between cap effects and cap levels appears to be new.

Although all-pay contest have been a popular way to model different types of competition, ranging from R&D races to political lobbying (Che and Gale 1998; Siegel 2009), there have been few efforts to apply them to campaign contests. This represents an alternative modeling approach to one more widely used in the literature, where spending levels are inputs into an exogenously specified winning probability function, which leads to pure strategies, but a probabilistic outcome as a function of outcomes. (e.g., Benoit and Marsh 2008; Erikson and Palfrey 2000; Gerber 2004; Gius 2009). In contrast, the election outcome in our model is deterministic in actions played, but features mixed-strategy equilibria. Although the picture of campaigning is painted in very broad strokes, it is able to capture many of the empirical properties that similar models address, while using fewer moving parts, e.g., it does not require a specification of a probabilistic outcome function, which suggests that an all-pay contest may not be a bad approximation for political campaigns. One empirical regularity that has received a great deal of attention, for instance, is the fact that incumbent spending significantly less effective than challenger spending. In our model, this is an immediate consequence of mixing, since each additional dollar has a marginal effect equal to the cost of raising it. This prediction is not necessarily hard to generate in a setting with probabilistic outcome and incumbent advantage, but the all-pay framework does offer a unique additional insight: the effect of a cost advantage is likely to show up in a constant term in estimations, rather than in the average effect of spending, since it induces the challenger to play with a mass point at a low level of campaigning. The probabilistic nature of strategies can also shed some light on the presence of asymmetries that seem too extreme to be reasonable; for instance, in 2006 incumbent Roy Blunt of Mississippi spent over $3 million, whereas his challenger spent nothing.

Our model diverges from existing treatments of all-pay auctions and contests (Baye, Kovenock, and De Vries 1996; Siegel 2009) in one critical way, which does not appear to have been extensively explored in the literature: all candidates have same *valuation* for the prize, but differ in their *budget constraints*, which are given by their types. In our model, the candidate's type is conceived as the ''objective'' amount of support for a candidate's position, i.e., the number of effective arguments that they can bring to bear on the campaign. Although it is true that candidates can in general always spend more money, spending itself is largely a proxy for campaign intensity, and there does seem to be a point at which a candidate's ''type'' bounds his ability to make gains by further campaigning, even of the money to do so is available. The inefficacy of self-financing is fairly well-documented (see Mueller 2003, for survey), and recent election cycles have produced several conspicuous well-funded candidates who proved to have a limited reach: Meg Whitman's bid for California governor in 2010 was unsuccessful, in spite of effectively unlimited funds, with $144 million being spent from her personal fortune; Rick Perry's run at being the GOP presidential nominee for 2012 started with a strong lead in funding, but the advantages waned over the course of the primaries as his image solidified. As a general modeling choice, however, equal valuations with private budget constraints may apply to any number of settings where a ''buyer's'' private valuation is not equal to the social, or seller's, benefit of that player winning, and any contest that offers a fixed prize and hopes to select the highest ''revealed'' type, which might be independent of his cost function[3], might be modeled in such a way.

While our model is admittedly very stark – and very partial in the sense that it does not address other related closely related aspects of political competition such as policy choice – it captures many of the empirical regularities of the campaign finance literature. A fuller specification would undoubtedly benefit its explanatory power, but the basic point of this paper is that all-pay models, which have been profitably employed to analyze many other forms of competition, including political lobbying, can offer additional insights into the campaigning process. We offer a very simple model parsimoniously accounts for many stylized facts and empirical puzzles, while clarifying

[3] Usually, signaling games select on the cost function as an indicator of quality; there might be instances, however, where the max output is more apt.

the role and effects of campaign spending efforts. We then use the model to derive new results regarding the effects of campaign finance policy and spending caps.

3.2. Model Environment

There are two candidates, L and R, who compete for the median vote. Each candidate receives a utility of one if he wins the election, and zero otherwise. The candidates compete in an all-pay contest, where each candidate supports his platform by producing arguments, and the candidate who releases tho most arguments wins the election.

Candidates will differ along two dimensions: the cost of producing arguments, which is known beforehand, and the maximum number of arguments to which he has access, which is a random variable that is private information to each candidate – this should be thought of as the objective strength of a candidate's position.

We describe a candidate's strength as his type: each candidate $i \in \{L, R\}$ will have a type $\theta_i \in \Theta_i = [0, 1]$ that denotes the maximum number of arguments to which he has access – it is effectively a stochastic bidding cap. The types θ_i for $i = L, R$ are independently and uniformly distributed across $[0, 1]$.

After observing his type, each candidate releases some number of arguments to the public; call λ the number of arguments released by L, and ρ the number of arguments released by R; these are bound by the types of their respective candidates, so they must satisfy $\lambda \in [0, \theta_L]$ and $\rho \in [0, \theta_R]$. Releasing arguments is costly, and each candidate i has an exogenous cost function $c_i(x)$, which is common knowledge, that determines his cost of releasing x arguments. Let R be the cost-advantaged candidate, so that $c_R(x) = \alpha c_L(x) = c(x)$ for all $x \in [0, 1]$, where $\alpha < 1$[4]. Furthermore, we assume that $c' > 0$, $c'' > 0$, $c(0) = c'(0) = 0$, and $c(1) > 1$; the final condition means that the cost-disadvantaged candidate L will never find it optimal to always release all arguments available to him, even if doing so would guarantee victory; the prior condition guarantees that it will always be optimal to release at least some arguments.

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

1. Candidates L and R observe their types, i.e., the total number of arguments to which each

[4] The multiplicative form is not essential to our results, but simplifies the analysis somewhat.

has access

2. Candidate L releases λ arguments, R releases ρ arguments, subject to the constraint that $\lambda < \theta_L$ and $\rho \leq \theta_R$.

3. The candidate with the most arguments wins the election.

Candidates have quasilinear utility:

$$Eu_L(\lambda, \rho) = P(\lambda > \rho) - c_L(\lambda)$$

$$Eu_R(\lambda, \rho) = P(\lambda < \rho) - c_R(\rho)$$

For each candidate, a strategy consists of a mapping from the observed type to a probability distribution (possibly degenerate) over all feasible actions. We denote these strategies by $h^L(\cdot, \theta_L)$ and $h^R(\cdot, \theta_R)$ where $h^L(\lambda, \theta_L)$ is the probability density function at λ for an L candidate of type θ_L.

3.3. Model discussion

Although the model is highly stylized, it manages to capture two competing effects of campaigning that are often difficult to describe simultaneously: direct campaigning effects and ''signaling'' motives behind campaigning. The main point of departure vis-à-vis the all-pay contest literature, as well as the signaling literature, is the idea that a candidate's type is represented as a budget constraint rather than a cost function (as it would be in Spence-type signaling models). The notion here is that a candidate's ability to campaign may be quite independent of his quality as a representative. As a result, easier access to campaign finance may lead elections to select the most able campaigners rather than the best quality politicians – it is this intuition that has been mostly absent from other models of competition. The simple outcome rule, however, also clearly delineates a space for direct campaign effects.

One way to motivate this setup is with a boundedly-rational median voter, who is exposed to campaigning and decides who they believe is the best candidate. In reality, campaign spending may be targeted, but they are often targeted in the same areas; most campaign spending in the recent 2012 presidential election, for instance, has been focused on the ''swing states'' with relatively equal intensity by both parties.

In a way, our model represents "middle-ground" between direct effects models and signaling models. A potentially undesirable characteristic of rational signaling models is that they may suppress the effects of cost asymmetries: rational voters will often reason away differences in spending ability, and so all that matters is whether a separating equilibrium can be sustained or not. On the other hand, direct effects models ignore any informational purpose of campaign spending, and hence do not capture the arguments that exist in favor of allowing for more spending, i.e., the production of a more informed electorate. By considering a contest with a simple decision rule, we present a model that avoids the problems of these extreme assumptions; the important thing is that we consider a mechanism whereby information is important, but not entirely extracted in a way that nullifies cost asymmetries.

Another way to approach this type of "bounded rationality" comes from the political psychology literature. There is a fairly extensive list of studies to confirm the existence of biases in voter behaviors and to identify the channels through which campaign advertising and media affects voter behavior (for survey, see Kinder 2003). Of particular relevance to our model is the literature on "priming," which is also framed in terms of "issue ownership". The intuition behind these theories is that parties do not directly confront each other in their policy platforms, but rather focus on different issues. The intent of media communications is not to directly with each other, but rather to try to make voters think in a specific way when they vote, e.g., to weigh more heavily fiscal policy than welfare policy. In such cases, the argument can be made that greater the media exposure of a particular issue topic, the greater its importance in the mind of the voter, and the more likely the voter will vote according to policy positions on that issue, potentially ignoring other important issues (for instance, Petrocik 1996).

For the purposes of the remainder of the paper, we will think of arguments released as corresponding roughly to campaign spending, i.e., to release more arguments is more costly. The other important parameter is the cost asymmetry. Although this is not directly measurable, it may be correlated with some observable characteristics, such as the number of different sources or organizations from which a candidate receives funds, or pre-existing connections or campaign experience. In the context of recent campaign finance liberalizations, the concentration of donors also becomes important: if individual spending becomes unlimited, then a candidate with access to concentrated, wealthy donors willing to spend large amounts of money will benefit more from liberalization than a candidate with a broader, less concentrated donor base.

3.4. Equilibria Characterization

The solution concept is Bayesian Nash equilibrium: each type maximizes his expected utility.

$$h^L(\cdot,\theta_L) \in \arg\max_{h'} E[u^L \mid h',\theta_L]$$

$$h^R(\cdot,\theta_L) \in \arg\max_{h'} E[u^R \mid h',\theta_R]$$

Since candidate types are independent, the argument distribution each candidate faces does not depend on his own type. Let G^λ and G^ρ denote the *ex ante* cumulative distribution functions of the arguments released by each candidate. Then the probability that candidate L when he releases λ arguments is simply $G^\rho(\lambda)$. Using this fact, we can write each candidates expected utility from releasing a particular number of arguments as follows:

$$Eu^L(\lambda) = G^\rho(\lambda) - c_L(\lambda)$$

$$Eu^R(\rho) = G^\lambda(\rho) - c_R(\rho)$$

Note that G^ρ and G^λ are related to strategies by the following

$$G^\rho(x) = \int_x^1 H^R(x,\theta)\, d\theta$$

$$G^\lambda(x) = \int_x^1 H^L(x,\theta)\, d\theta$$

Our characterization of equilibria will pin down the induced *ex ante* action distributions G^ρ and G^λ. It will turn out that the only distributions consistent with equilibrium involve regions over which candidates are indifferent between different options; moreover, the cost asymmetry will necessitate that one player must play a strategy that has probability mass, which rules out pure-strategy equilibria.

The following theorem presents the result; the remainder of this section details the proof.

Theorem 3.1. *Let $\underline{x}$ satisfy $c_L'(\underline{x}) = 1$. Then, let*

$$\bar{x} \equiv c_L^{-1}\left[1 - \left(\underline{x} - c_L(\underline{x})\right)\right]$$
$$v^L \equiv 1 - c_L(\bar{x})$$
$$v^R \equiv 1 - c_R(\bar{x})$$

All equilibria of the campaigning game described above generate the following action distributions:

$$G^{\rho}(x) = \begin{cases} x & \text{for } x < \underline{x} \\ v^L + c_L(x) & \text{for } x \in [\underline{x}, \bar{x}] \\ 1 & \text{for } x > \bar{x} \end{cases}$$

$$G^{\lambda}(x) = \begin{cases} x & \text{for } x < \underline{x} \\ v^R + c_R(x) & \text{for } x \in [\underline{x}, \bar{x}] \\ 1 & \text{for } x > \bar{x} \end{cases}$$

That is, any candidate with fewer than $\underline{x}$ releases all of them, and any candidate with greater than $\underline{x}$ arguments releases some number $x \geq \underline{x}$ of them in a way consistent with the distributions above; since candidates are indifferent between all alternatives within the range, any manner of mixing that generates the above distributions will constitute an equilibrium.

Since the space of all possible strategies is quite large, we proceed by gradually narrowing the space of candidate strategies until we arrive at the given expressions. First, note that equilibrium strategies cannot have both candidates playing strategies with mass points at the same point in their action distributions, G^{ρ} and G^{λ}; if they did, then playing at that point would be strictly dominated by playing above that point whenever possible. Playing mass at zero is also never optimal, so $G^{\rho}(0) = G^{\lambda}(0) = 0$.

Furthermore, since $c'(0) = 0$, there must exist some $\varepsilon > 0$ such that whenever $\theta_i < \varepsilon$, the optimal strategy for candidate i is to release all θ_i arguments. If not, then one player would have to play mass at zero. To see this, note that if such an ε does not exist, then for every $\varepsilon > 0$, we can find some $\delta \in (0, \varepsilon)$ such that a candidate with access to ε arguments does at least as well by playing δ as by playing ε, i.e., from the point of view of candidate L, this amounts to

$$G^{\rho}(\delta) - c_L(\delta) \geq G^{\rho}(\varepsilon) - c_L(\varepsilon)$$

Since this holds for arbitrarily small ε, G^{ρ} and c_L are weakly increasing, and $c'(0) = 0$, this would imply that $g^{\rho}(0) = 0$. Since a candidate of type θ_i can release at most θ_i arguments, the probability that a candidate releases x or fewer arguments is at least as great as the probability that his type is below x, which is also equal to x. This means that

$$G^{\rho}(x) \geq x$$
$$G^{\lambda}(x) \geq x$$

Hence, the only way that $g^{\rho}(0) = 0$ can hold is if $G^{\rho} > 0$, which cannot happen in equilibrium.

Lemma 3.1. *The supports of G^{ρ} and G^{λ} are the same up to a set of measure zero.*

Proof. Suppose not, i.e., there is some open set $(a, b) \subset [0, 1]$ not in the support of G^{ρ}. Then, whenever releasing $\lambda \in (a, b)$ is possible for L, it is dominated by releasing a messages, with gives the same winning probability with lower costs. Hence (a, b) cannot be in the support if G^{λ}. Repeating the argument with an open set in the support of G^{λ} completes the proof. □

Lemma 3.2. *The supports of G^{ρ} and G^{λ} must not have "gaps", in the following sense: if a and b, with $a < b$ are in the support of G^{λ} (or G^{ρ}), then there cannot exist an open set $(a',' b) \subset [a, b]$ such that $(a', b') \cap Supp(G^{\lambda}) = \emptyset$.*

Proof. Suppose that such a gap exists. Then let

$$m \equiv \sup\{x < a \ : \ x \in \mathrm{Supp}(G^{\lambda})\}$$
$$n \equiv \inf\{x > b \ : \ x \in \mathrm{Supp}(G^{\lambda})\}$$

Basically, m and n are the boundaries of the gap if we widen it as much as possible. Then for any $\varepsilon > 0$, there exist $m' < m - \varepsilon$ and $n' > n - \varepsilon$ in the support of G^{ρ} satisfying

$$G^{\rho}(m') - c_L(m') \geq G^{\rho}(n') - c_L(n')$$

We know this from revealed preference: n' is always available as an option whenever m' is played, so m' must do at least as well as n'. We can rearrange the above expression as

$$G^{\rho}(m') - G^{\rho}(n') \geq c_L(m') - c_L(n') > 0$$

Since this difference is strictly positive for all ε, this implies that there must be mass either at m, or at n, or at both i.e., on one or both sides of the gap. The argument also holds for G^{λ}, so that if a gap exists, each candidate must play with mass on at least one side of the gap.

None of the possibilities, however, are consistent with equilibrium. We have already shown that both players cannot play with mass at the same point, since one candidate could do strictly better by moving that mass up by $\varepsilon > 0$ whenever such play is feasible. The only remaining configuration possible is for one candidate to have mass at m, the bottom of the gap, and the other to have mass at n, the top of the gap. This is clearly not optimal, however, for the candidate with mass at n, who would reduce his costs and achieve the same winning probability with any strategy in the interval (m, n). Since this exhausts all possibilities, our original assumption must not be true, which completes the proof. □

Lemma 3.3. *Candidates L and R must be indifferent between the same portions of their supports.*

Proof. Suppose not, that there exists (a, b) over which L is indifferent, but R is not. If there is a sub-interval in which R always releases the most arguments possible, then $g^{\rho} = 1$, which is inconsistent with L being indifferent. The other possibility is for R to be indifferent over disjoint open subintervals, and to keep $g^{\rho}(x) = c_L'(x)$ for $x \in (a, b)$. Let $c \in (a, b)$ be the boundary between the indifference regions of R, i.e.r R is indifferent between actions in $(c - \varepsilon, c)$, and then gets higher utility in the interval $(c, c + \varepsilon)$. This means that all $\theta_R > c$ will never play below c, so that $G^{\rho}_{-}(c) = c$, where $G^{i}_{-}(x)$ represents the limit from below of G^{i}. However, since $G^{\rho}(x) \geq x$ for all x, and $g^{\rho}(x) = c_L'(x)$ for $x \in (a, b)$, this would entail mass at c. Mass at c, however, would make L any point slightly above c to c, contradicting indifference. □

Lemma 3.4. *If a candidate ever plays some strategy x when a higher strategy is possible, then for that player there will never be a strategy $x' > x$ that strictly beats it.*

Proof. Suppose not, that candidate L (without loss of generality) plays x when some higher action is feasible, but that there exists some $x' > x$ that strictly beats it when feasible. Since there are no gaps in the support of G^λ, this means that there is some interval above x over which L is indifferent. The reason is that x' will be played when it becomes possible, so x and x' are in the support of G^λ. Since we established that this cannot have any gaps, this means that all values (except for a potentially a subset of measure zero) in the interval $[x, x')$ must be played by some types. But since x is played by some type $\theta_L > x$, it means that x must do at least weakly better than all actions in the interval $(x, \theta_L]$. Since these values must also be in the support of G^λ, they are played at some point, which means that they also do weakly better than x. This means that L must be indifferent between playing any action in the interval $[x, \theta_L]$.

Let $x_0 \equiv \inf\{x'' : Eu^L(y) = Eu^L(x)\, \forall y \in [x'', x]\}$ and $x_1 \equiv \sup\{x'' : Eu^L(y) = Eu^L(x)\, \forall y \in [x, x'']\}$, that is, $[x_0, x_1]$ represents the widest possible interval including x over which L is indifferent. Again, let $G^i_-(x) = \lim_{x' \uparrow x} G^i(x)$ denote the limit from below of $G^i(x)$.

By the previous lemma, L and R must both be indifferent over actions in the set (x_0, x_1). Since this is the largest set over which they are indifferent, candidates with types above x_1 will never play below x_1, and similarly candidates with types in (x_0, x_1) will never play below x_1, which means that $G^i_-(x_0) = x_0$ and $G^d_-(x_1) = x_1$. Furthermore, indifference of the candidates implies that

$$g^\rho(\tilde{x}) = c_L'(\tilde{x})$$
$$g^\lambda(\tilde{x}) = c_R'(\tilde{x})$$

Where $\tilde{x}$ is used to represent an arbitrary action within the range (x_0, x_1). Since only types within (x_0, x_1) play actions within the interval, all actions must be played in equilibrium, and $c'' > 0$, it must be the case that $g^\rho \leq 1$. If there is some $\tilde{x}$ for which this is not satisfied, then the indifference condition implies that g^ρ is increasing in $(\tilde{x}, x_1)$. Then the probability of actions being played in the region $(\tilde{x}, x_1)$ is $\int_{\tilde{x}}^{x_1} g^\rho(x)\, dx > x_1 - \tilde{x}$. However, the only types who play within this range are those within the range itself, so the maximum probability with which strategies in the range can be played is $x_1 - \tilde{x}$, a contradiction.

Given that $g^\rho(\tilde{x}) \leq 1$ and $g^\lambda(\tilde{x}) \leq 1$, and that $G^i_-(x_0) = x_0$, it must be that both candidates play with a mass point at x_0, which cannot happen, and so delivers the desired contradiction. □

The above lemmas pin down the basic character of all equilibria. At the lower end, candidates will reveal all arguments to which they have access. Then above a certain threshold, candidates will be indifferent between all strategies played above that threshold. All that remains is to find the two values $\underline{x}$ and $\bar{x}$, which mark these thresholds, i.e., candidates are indifferent between actions in $[\underline{x}, \bar{x}]$, and actions above $\bar{x}$ never get played.

First, not that neither candidate can play mass at $\bar{x}$, since otherwise the best response of the other candidate would entail playing actions just above it. For $x \in [\underline{x}, \bar{x}]$, indifference means

$$g^{\rho}(x) = c_L'(x)$$
$$g^{\lambda}(x) = c_R'(x)$$

No candidate can play strategies that induce mass above $\underline{x}$, since there can be no mass at $\bar{x}$, between $\underline{x}$ and $\bar{x}$ strategies must satisfy the indifference conditions. This means that for at least one candidate, $G^{i}(x) = x$. Since $c_R' < c_L'$, and $G^{\rho}(\bar{x}) = G^{\lambda}(\bar{x}) = 1$, it must be the case that $G^{\rho}(\underline{x}) < G^{\lambda}(\underline{x})$, which means that L will be playing with mass at $\underline{x}$.

To pin down $\underline{x}$, note that $c_i'(\underline{x}) \leq 1$ for $i = L, R$. If not, then immediately below $\underline{x}$ it would not be optimal for one candidate to release all arguments, which contradicts our definition of $\underline{x}$, since the marginal benefit of releasing arguments is 1 (assuming that the other candidate always releases all arguments). We also know, that R does not play with mass at any point, and no type above $\underline{x}$ ever plays below $\underline{x}$. As a result, it must also be the case that $g_+^{\rho}(\underline{x}) = c_L'(\underline{x}) \geq 1$, where g_+ represents the one-sided derivative of G from above. To see this, consider some type θ_R just above $\underline{x}$; all such types must play strategies between $\underline{x}$ and θ_R. Since the action distribution is continuous at $\underline{x}$, g^{ρ} must exceed 1. Since $c_L'(\underline{x}) \leq 1$ and $c_L'(\underline{x}) \geq 1$, it must be that $c_L'(\underline{x}) = 1$, which pins down the value of $\underline{x}$.

The top of the distribution is also easy to pin down, since G^{ρ} has no mass points and is well-defined above and below $\underline{x}$; we just need to find the value at which $G^{\rho}(\bar{x}) = 1$. Writing out $G^{\rho}(\bar{x})$ as

$$\begin{aligned} G^{\rho}(\bar{x}) &= G^{\rho}(\underline{x}) + \int_{\underline{x}}^{\bar{x}} c_L'(y)\, dy \\ &= \underline{x} + c_L(\bar{x}) - c_L(\underline{x}) = 1 \end{aligned}$$

We can rearrange and take the inverse:

$$\bar{x} = c_L^{-1}\left(1 - (\underline{x} - c_L(\underline{x}))\right)$$

To verify that these strategies are optimal, not that $c_L(\bar{x}) < 1$, since below $\underline{x}$, $c_L' < 1$, so $c_L(\underline{x}) < \underline{x}$. Below $\underline{x}$ it is clearly optimal for both candidates to release all arguments that they have, since the marginal benefit of doing so is 1 (taking opponent's strategy as given), and the marginal cost is less. Above $\underline{x}$, strategies are constructed such that each is indifferent, so it is optimal to mix within the range.

To characterize the cumulative distribution functions of the actions played, we exploit the indifference conditions: each candidate's realized utility for any action in $[\underline{x}, \bar{x}]$ is equal to his utility at $\bar{x}$, when he wins almost certainly. Let v^L be this utility for L, and v^R be this utility for R, then

$$v^L = 1 - c_L(\bar{x}) = \underline{x} - c_L(\underline{x})$$

Since L is indifferent between all actions $x \in [\underline{x}, \bar{x}]$, his expected utility is always equal to v^L:

$$Eu^L(x) = G^\rho(x) - c_L(x) = v^L$$

Rearranging gives

$$G^\rho(x) = v^L + c_L(x)$$

Similarly, we can get

$$G^\lambda(x) = v^R + c_R(x)$$

where $v^R = 1 - c_R(\bar{x})$. Since G^ρ and G^λ must both equal one at $\bar{x}$, and they must equal x below $\underline{x}$, the strategies all entail L playing with mass at $\underline{x}$. When $c_R = \alpha c_L$, we have that $v^R = 1 - \alpha(1 - \underline{x} + c_L(\underline{x}))$, so that we can solve out

$$\begin{aligned} G^\lambda(\underline{x}) &= v^R + c^R(\underline{x}) \\ &= 1 - \alpha\left(1 - \underline{x}\right)) \end{aligned}$$

Furthermore, we know that $G^\lambda_-(\underline{x}) = \underline{x}$, which implies a probability mass of $G^\lambda(\underline{x}) - G^\lambda_-(\underline{x}) = (1-\alpha)(1-\underline{x})$ at $\underline{x}$.

3.5. Analysis and Discussion

Although there are multiple equilibria, the characterization above allows pins down many properties of all equilibria, which we now explore in greater detail. In particular, since the action distribution is unique, the bias of the election is well-defined, and we can consider how it changes with changes in the parameters. We first place the findings of our model in the context of some puzzling empirical findings in the literature, such as the seeming inefficacy of incumbent spending. We show that the basic findings are all consistent with the presence of cost asymmetries in an all-pay contest.

We would like the contest to select as often the high-type candidate, i.e., we would like it to be informationally efficient. In general, different equilibria will have different informational efficiency characteristics, since there is no guarantee of monotonicity. Our analysis focuses on the maximal efficiency of equilibria, and show that it declines with growing asymmetries.

Finally, we consider the impact of policy that affects campaign finance, namely, asymmetric liberalization and spending caps. Liberalization alone can be beneficial, but worsening asymmetries have counteract the benefits. We find that spending caps can be beneficial, both for reducing bias as well as for improving efficiency, but that it depends critically on the choice of the cap and model parameters. In particular, when costs are more similar, all caps may worsen the informational efficiency; furthermore, binding caps that are set ‘‘too high’’ may have no impact on the bias, but reduce informational efficiency.

3.5.1. The incumbency paradox

Our model allows us to explain the ‘‘incumbency paradox’’ quite neatly, as a result of the asymmetric cost functions. In particular, our equilibrium involves asymmetries both in the levels of spending and in the marginal effects of additional spending in ways consistent with the observations noted in the introduction.

In many ways, it makes sense for the marginal effects of spending to differ between candidates: since they are engaged in competition, the marginal effects of spending will depend on the actions of other candidates, i.e., against a relatively evenly matched challenger, additional spending will likely have more of an effect on the final outcome than against a challenger who had little chance of winning in the first place. This basic idea is present in previous efforts to explain the incumbency paradox[5], but often an incumbency advantage in vote share or popularity needs to be assumed *in addition to cots asymmetries*; here, the only fundamental asymmetry between the two candidates is the difference in cost functions, or fund-raising capacity, and the differences in the marginal productivity of their efforts is a consequence of equilibrium strategies in the political contest. The main advantage of our model vis-à-vis existing explanations is that it allows us to pin down several of the observed ''anomalies'' on one fundamental asymmetry. While other sources of candidate asymmetry are certainly important in electoral outcomes, the fact that cost asymmetries *alone* can explain all of the major campaign spending anomalies gives a clearer picture on the potential impact of money in the electoral process.

This section describes some of the characteristics of our equilibria. We show that *cost* asymmetries generate asymmetries in both spending *levels* as well as *marginal effects*, and show that they also generate an incumbency advantage when candidates are otherwise *ex ante* identical.

We formalize the differences in spending levels in the following corollary:

Corollary 3.1. *When $\alpha < 1$, G^ρ strictly first-order stochastically dominates G^λ.*

Proof. The proof is straightforward: for $x < \underline{x}$, $G^\rho(x) = G^\lambda(x)$, and for $x \in [\underline{x}, \bar{x}]$, $G^\rho(x) < G^\lambda(x)$, since within this range

$$G^\rho(x) = 1 - \int_x^{\bar{x}} c_L'(y)\, dy$$

whereas

$$\begin{aligned} G^\lambda(x) &= 1 - \int_x^{\bar{x}} c_R'(y)\, dy \\ &= 1 - \alpha \int_x^{\bar{x}} c_L'(y)\, dy \end{aligned}$$

which gives the desired result. □

There are two consequences to this: (1) the incumbent will campaign more intensively in expectation, which can explain the differences in observed levels, and (2) for the same amount of spending, the incumbent will typically win with higher probability. Although the second effect is

[5] For instance, Moon (2006), looking at vote shares, presents a model where incumbent spending targets more expensive extreme voters in the presence of weak challengers

a relatively straightforward consequence of asymmetric cost functions, and the fact that equilibrium involves mixing, it nevertheless bears on the interpretation of empirical studies into campaign finance effects. Most estimates of campaign spending effects measure the average effect of spending, but this may not really capture the role of financing in campaigns; the strategies generated in our model would make it likely to show up in the *constant* term, or the incumbency advantage, even though its fundamental cause is the incumbent's easier access to campaign funds.

We can capture the idea of incumbency advantage by the *bias* of the outcome; this is a measure of the advantage an incumbent has because of his access to easier funds.

Definition 3.1 (Bias). *Let the bias b be the difference between the incumbent's* ex ante *probability of winning and* $1/2$ *(the "fair" probability – how often the incumbent would win if all information were revealed), i.e.,*

$$b = P(\rho > \lambda) - \frac{1}{2}$$

Corollary 3.2. *The incumbent wins with higher probability; more specifically,*

$$b = \frac{1}{2}(1 - \underline{x})^2(1 - \alpha)$$

Furthermore, reductions in the incumbent's cost increase the bias, i.e., $db/d\alpha < 0$*.*

Proof. We can find the bias easily by looking at separate cases. When at least one of the candidates has type below $\underline{x}$, the outcome is efficient, since the candidate below $\underline{x}$ reveals all arguments, and no candidate at or above $\underline{x}$ ever reveals fewer than $\underline{x}$.

The only cases remaining are those where both candidates have types above greater than or equal to $\underline{x}$; which occurs with probability $(1 - \underline{x})^2$. Note that conditional on $\lambda \geq \underline{x}$, the probability that $\lambda = \underline{x}$ is $1 - \alpha$, since the probability mass is $(1 - \alpha)(1 - \underline{x})$, and the conditioning event has probability $1 - \underline{x}$. Furthermore, conditional on L playing *above* $\underline{x}$, we have that $g^\lambda(x|x > \underline{x}) = (1/\alpha)c'_R(x) = c'_L(x) = g^\rho(x|x \geq \underline{x})$, i.e., the distribution is the same as the conditional distribution of R actions, which means that of these, L wins $1/2$. The only remaining case is when L actually plays $\underline{x}$: here, he loses with probability 1, and this is the only subcase where the conditional probability deviates from the "fair" outcome. This case occurs with overall probability $(1 - \underline{x})^2(1 - \alpha)$, and the deviation from the fair probability is $1/2$ of that, which gives the bias. $\square$

The second issue related to the incumbency paradox is the effect of the *marginal* dollar spent, or in our case, the marginal effect of an additional argument. Consider the issue from the point of view of candidate L. His probability of winning with x arguments is simply the probability that R releases fewer arguments, $G^{\rho}(x)$. Consequently, the marginal effect of an additional argument is the probability density of the actions of R at x, or $g^{\rho}(x)$. Suppose that we are in the range $[\underline{x}, \bar{x}]$, then the marginal benefit of additional campaigning for L is given by $g^{\rho}(x) = c_L'(x)$, whereas the benefit of additional campaigning by R is $g^{\lambda}(x) = c_R'(x) = \alpha c_L'(x) < g^{\rho}(x)$; in other words, the marginal benefit of campaigning for each party is equal to the marginal cost of campaigning for the *other* party, which means that L raises his winning probability by more with each additional argument.

This actually suggests quite a different interpretation of campaign spending effects than the ones usually offered, where the "gains" of additional spending are supposed to be exogenous, and so the question becomes shifted onto that of why the incumbent seems to be buying more expensive goods with his campaign spending. In our model, the efficacy of campaign spending is endogenously determined in response to differences in campaigning costs. Moreover, we do not suppose any other asymmetries, such as the typical incumbency advantage, where voters have some type of predisposition to vote for the incumbent – the candidates are, aside from cost differences, *ex ante* identical.

3.5.2. Informational efficiency

Another question that is potentially of interest is whether or not we can expect election winners to correspond to the candidates with the stronger positions, i.e., those with more arguments at their disposal.

Consider the following set of strategies, which we will call equilibrium A:

Definition 3.2 (Equilibrium A)**.** *Let $\lambda(x)$ denote the number of arguments released by an L candidate of type x, and let $\rho(x)$ be the number of arguments released by candidate R of type x. Call* *equilibrium A that which consists of the following strategies:*

$$\rho(x) = \begin{cases} x & \text{if } x < \underline{x} \\ \underline{x} + \int_{\underline{x}}^{x} \frac{1}{c_L'(s)}\, ds & \text{if } x \geq \underline{x} \end{cases}$$

$$\lambda(x) = \begin{cases} x & \text{if } x < \underline{x} \\ \rho(x) & \text{with probability } \alpha, \text{ if } x \geq \underline{x} \\ \underline{x} & \text{with probability } 1-\alpha, \text{ if } x \geq \underline{x} \end{cases}$$

This particular equilibrium has several nice properties: it is intuitive, monotonic, and has a very simple mixing rule (arguably, this is as "close" to a pure strategy as we can get). It also happens to be, however, maximally informationally efficient, in the sense that it reduces the probability that the candidate with fewer arguments wins. We present this in the following theorem.

Theorem 3.2. *Equilibrium A minimizes the probability of the lower type candidate winning the election.*

Proof. First, note that the bias forms a lower bound for the error probability in question: since candidate L is the high-type candidate $1/2$ the time, and he loses with $1/2-b$ of them, the "wrong" candidate must be chosen with at least probability b.

All that is left to show is that the strategies described in equilibrium A achieve this lower bound, which is straightforward since, aside from L types above $\underline{x}$ playing $\underline{x}$ with probability $1-\alpha$, their strategies are the same, so whenever L plays $\rho(x)$, the outcome is efficient. The only inefficient outcomes are those where L plays $\underline{x}$ and both types are above $\underline{x}$, which happens with probability $(1-\underline{x})^2(1-\alpha)$. L loses all of these, when he is the higher type candidate in $1/2$ of them, so the error probability is $(1/2)(1-\underline{x})^2(1-\alpha)$, which is equal to the bias. □

One consequence of this is that the more severe the asymmetry, the greater the mass that L must play at $\underline{x}$, which reduces the informational efficiency of the equilibrium.

Theorem 3.3. *As α decreases from one down to zero, the minimum achievable probability that the lower type candidate gets elected increases.*

Proof. As argued above, the minimal error rate is equal to b, and we have previously shown that this increases with α. □

3.5.3. Campaign finance reform

In this section we consider how changes in the financing environment would be expected to change the outcomes of the contest. We consider two types of changes: (1) a change that affects the cost functions of each candidate directly, such as *Citizens United*, where the costs for both parties of any particular amount of campaigning are lowered, but possibly asymmetrically; (2) the imposition of spending caps, where there is a hard cap on the amount of campaigning that can be done in equilibrium.

Changes in financing costs

Let us parameterize the cost functions with an additional term γ, so that $c_L(x) = \gamma c(x)$ and $c_R(x) = \alpha\gamma c(x)$. Arguably, we can model the effect of the *Citizen's United* decision as a reduction in γ, and possibly also a reduction in α, if we think that the cost-advantaged party gains more – one justification would be if one party has access to more wealthy or more intensely interested single donors, so that the lift on single-source spending greatly eases fund-raising efforts on one party.

Corollary 3.3. *A reduction of γ raises $\bar{x}$, so that the set of L types that mixes shrinks. It also leads to higher spending levels by both parties. The minimal error rate is still expressed as $(1/2)(1-\alpha)(1-x)^2$, which decreases if α remains constant.*

Proof. The results follow directly from the increase in $\bar{x}$, which is now defined implicitly by $\gamma c'(\bar{x}) = 1$; clearly, as γ falls, $\bar{x}$ must rise since $c'' > 0$ by our assumptions. □

Hence, a symmetrical reduction in cost functions will improve the maximum achievable efficiency of equilibrium. There is, however, a catch, which is that if the reduction is not symmetric, then the effect is ambiguous, since it also reduces informational asymmetry conditional on both candidates having types above $\bar{x}$: an asymmetric reduction hence increases the contests ability to discriminate between lower types, but may decrease its ability to discriminate between higher types, with the net effect being dependent on the precise parameters and cost functions.

Spending caps

We now consider a different type of campaign finance reform: caps on total spending. Hence, now define x^* as the upper bound of the permissible action space. Let us suppose that x^* is set such that it is in the range $(\underline{x}, \bar{x})$, which remain as defined above.

The binding upper bound makes specifying equilibria more complicated by allowing for mass at x^*, and the results will be sensitive to the choice of spending cap. For x^* low enough, it will be optimal for all types above x^* to play exactly at x^*. In particular, define $\tilde{x} \in [\underline{x}, \bar{x}]$ to be the point where the following is satisfied:

$$v_L + c(\tilde{x}) = \frac{1}{2}(1 + \tilde{x})$$

To see that such a point must exist, consider the difference between the left hand side and the right hand side:

$$v_L + c(x) - \frac{1}{2}(1 + x)$$

This is negative at $\underline{x}$, positive at $\bar{x}$, and strictly increasing with x in between, so at some unique point it must equal zero, which gives $\tilde{x}$.

We first present results for caps in the range $[\underline{x}, \tilde{x}]$. A cap in this range always reduces the bias, and with sufficiently low α also improves informational efficiency.

A spending cap greater than $\tilde{x}$, on the other hand, has an effect different from what one might expect: it has no effect on the bias, and strictly reduces efficiency. It turns out that the mass L plays at $\underline{x}$ stays the same, but now both players must also play mass at x^*, with that of L being less than that of R.

Theorem 3.4. *All equilibria in the game with a spending cap $x^* \in [\underline{x}, \tilde{x}]$ can be described as follows. First, define*

$$x' = c^{-1}(x^* - (\underline{x} - c(\underline{x})))$$

Then equilibrium strategies for each candidate must satisfy the following action distributions:

$$G^{\rho}(x) = \begin{cases} x & \text{for } x < \underline{x} \\ v_L + c_L(x) & \text{for } x \in [\underline{x}, x') \\ x^* & \text{for } x \in [x', x^*) \\ 1 & \text{for } x \geq x^* \end{cases}$$

and

$$G^{\lambda}(x) = \begin{cases} x & \text{for } x < \underline{x} \\ v_R^* + c_R(x) & \text{for } x \in [\underline{x}, \tilde{x}) \\ x^* & \text{for } x \in [\tilde{x}, x^*) \\ 1 & \text{for } x \geq x^* \end{cases}$$

where $v_R^* = \alpha v_L + (1-\alpha)(1 - x^* + \underline{x})$ *is the expected utility of* R *when he plays* $\underline{x}$*.*

Proof. The proof follows almost exactly as the main theorem, with the exception that we can no longer rule out mass being played at the top. It is easy to establish that the lowest types must continue to release all arguments, up to $\underline{x}$. If a candidate's opponent plays with mass at x^*, and all actions below x^* are played in equilibrium, then all types x^* and above will find it optimal to play x^*, since the winning probability jumps up discretely at that point, which ensures that it is also optimal for the opponent to play a mass at x^*. Since x^* is below $\bar{x}$, this must be the case.

This means that all types in $[\underline{x}, x^*)$ must be indifferent between actions played within that range. If we condition on both candidates having types strictly below x^*, the strategies will look exactly like they do in the uncapped game: all types in the range $[\underline{x}, x^*)$ mix between actions in $[\underline{x}, \tilde{x})$ for some $\tilde{x}$ that we need to determine.

We know that for action $\tilde{x}$, the winning probability must be x^*; also, we can repeat the arguments from earlier to conclude that R plays a massless strategy within this range, and L plays with mass at $\underline{x}$, so we know that $G^{\rho}(\underline{x}) = \underline{x}$, as before. Then the fact that L is indifferent between $\underline{x}$ and $tildex$ means

$$G^{\rho}(\tilde{x}) - c(\tilde{x}) = G^{\rho}(\underline{x}) - c(\underline{x})$$

or

$$\tilde{x} = c^{-1}(x^* - (\underline{x} - c(\underline{x})))$$

Then letting

$$v_L^* = x^* - c(\tilde{x})$$

$$v_R^* = x^* - \alpha c(\tilde{x})$$

we get

$$G^\rho(x) = v_L^* + c(x)$$

$$G^\lambda(x) = v_R^* + \alpha c(\tilde{x})$$

for $x \in [\underline{x}, \tilde{x})$, and the winning probabilities remain constant at x^* until x^*. □

As with the case with no spending caps, we can determine the bias and minimum efficiency.

Corollary 3.4. *The bias of the game with a binding spending cap $\tilde{x} > x^* > \underline{x}$ is given by*

$$b = \frac{1}{2}(x^* - \underline{x})^2(1 - \alpha)$$

Proof. When both candidates have types above x^*, each wins with probability 1/2, the fair outcome; the only non-trivial case is when both candidates play actions in $[\underline{x}, x^*)$, which corresponds to both types being within that range. Again, conditional on L playing above $\underline{x}$, the conditional action distributions of both candidates are the same, so each wins with 1/2. Conditional on both players being in $[\underline{x}, x^*)$, only when L plays $\underline{x}$ does he lose all the time; this case happens with probability $(x^* - \underline{x})(1 - \alpha)$, which then leads to the value of bias given above. □

What this means is that the bias of the contest decreases, in the sense that the *ex ante* winning probabilities become more even. In fact, setting a very low spending cap, below $\underline{x}$, would bring the bias down to zero. There is a cost to be paid in terms of informational efficiency, however: the maximal efficiency of equilibrium suffers, since whenever both candidates have types above x^*, the outcome is random, which will select the higher type candidate with probability only 1/2 – this is the same informational efficiency as setting the winner always equal to the same candidate.

Corollary 3.5. *For a game with spending cap $x^* \in [\underline{x}, \tilde{x}]$, the minimum error probability is less than that of the game without a spending cap if $\alpha < 2(x^* - \underline{x})/(1 - 2\underline{x} + x^*)$.*

Proof. The reasoning is similar to the case without caps, but now the maximal (conditional) efficiency when both candidates have types above x^* is $1/2$, since the outcome here is random, Hence the minimum error probability is at least $(1/2)(1 - x^*)^2$.

Of the remaining, cases, we know that R wins $1/2 + b$, but is the higher-type candidate only $1/2$ of the time, so the error rate is at least

$$\frac{1}{2}(1 - x^*)^2 + b$$

This is achieved by an equilibrium where strategies for $x < x^*$ are exactly as prescribed by equilibrium A, and candidates with higher types play with mass at x^*. It is easy to verify that the only inefficient outcomes generated by these strategies (additional to those described above) occurs when both candidates are in $[\underline{x}, x^*)$ and L plays $\underline{x}$ – precisely half of all such cases are inefficient, and it is the same expression as the bias.

Hence, the minimum error probability is

$$\begin{aligned}\frac{1}{2}\left[(1 - x^*)^2 + (x^* - \underline{x})^2(1 - \alpha)\right] &= \frac{1}{2}\left[\alpha(1 - x^*)^2 + (1 - \alpha)((1 - x^*)^2 + (x^* - \underline{x})^2)\right] \\ &= \frac{1}{2}\alpha(1 - x^*)^2 + \frac{1}{2}(1 - \alpha)(1 - \underline{x})^2 - \frac{1}{2}(1 - \alpha)2(1 - x^*)(x^* - \underline{x})\end{aligned}$$

The middle term is the minimum error probability without spending caps, so spending caps improve efficiency as long as

$$\alpha(1 - x^*)^2 - 2(1 - \alpha)(1 - x^*)(x^* - \underline{x})$$

which we can rearrange to get the expression stated. □

Hence, in the presence of severe asymmetries, a properly chosen spending cap can effectively reduce bias and increase efficiency. The potential benefits of a spending cap are only realized, however, for caps below $\tilde{x}$, and the consequences of a cap change significantly when we cross the threshold. We now turn our attention to the case a spending cap $x^* \in (\tilde{x}, \bar{x})$.

Theorem 3.5. *The equilibria of a game with a spending cap $x^* \in (\tilde{x}, \bar{x})$ are characterized as follows.*

Let

$$p_R \equiv 2(1 - v_L - c(x^*))$$

and

$$x' \equiv c^{-1}\left(1 - v_L - p_R\right)$$

where $v_L = \underline{x} - c(\underline{x})$*, as before. Then the action distributions must satisfy*

$$G^\rho(x) = \begin{cases} x & \text{for } x < \underline{x} \\ v_L + c(x) & \text{for } x \in [\underline{x}, x') \\ 1 - p_R & \text{for } x \in [x', x^*) \\ 1 & \text{for } x \geq x^* \end{cases}$$

$$G^\lambda(x) = \begin{cases} x & \text{for } x < \underline{x} \\ v_R + c(x) & \text{for } x \in [\underline{x}, x') \\ 1 - \alpha p_R & \text{for } x \in [x', x^*) \\ 1 & \text{for } x \geq x^* \end{cases}$$

where $v_R = \alpha v_L + (1 - \alpha)$*, as before.*

In other words, R plays mass p_R*, and L plays mass* αp_R *at* x^**; L plays mass* $(1 - \alpha)(1 - \underline{x})$ *at* $\underline{x}$*; below* $\underline{x}$*, both candidates play x, and in the remaining region, each plays with density equal to the marginal cost of his opponent, as before.*

Proof. As before, all candidates below $\underline{x}$ play their types; afterwards, there is a region over which they are indifferent; after that, there is a gap in the support, and mass must be played at x^* by both candidates. Since $x^* > \tilde{x}$, however, playing at the boundary can no longer strictly beat actions below x^* for both candidates. Nor can it be a unique best response for a single candidate (which must be R, because of his cost advantage); hence, both candidates must be indifferent between x^* and actions below it.

Let p_R be the mass that R plays at x^*. Then when L plays x^*, he wins with probability $1-(1/2)p_R$, which gives him utility $1 - (1/2)p_R - c(x^*)$. This must equal v_L, which pins down p_R:

$$p_R = 2\left[1 - v_L - c(x^*)\right]$$

Since $x^* > \tilde{x}$ which would equate p_R with $1 - x^*$, we know that $p_R < 1 - x^*$ and is feasible. As before, R cannot play with mass at $\underline{x}$, so x^* is his only mass point. This means that his action distribution on the support below x^* is completely determined: $g^\rho(x) = 1$ for $x < \underline{x}$, and $g^\rho = c'(x)$ for $x \geq \underline{x}$, and this must continue up to the point where the measure of types "left over" is equal to p_R, i.e.,

$$v_L + c(x') = 1 - p_R$$

which pins down x'.

As for L, g^λ is fixed in the region $[\underline{x}, x']$, which means that all the remaining types must play with mass points at $\underline{x}$ and x^*. Furthermore, we know that R must also be indifferent between x' and x^*; since his cost increases between the two by α times the cost increase experienced by L, his winning probability must go up by α times the increase in the winning probability of L between those points; this fixes the mass that L plays at x^* at αp_R.

The mass at $\underline{x}$ played by L makes up the remainder (so that the total measure of all actions played is equal to 1). The measure of R actions in the range $(\underline{x}, x']$ is $1 - \underline{x} - p_R$, which means that the measure of L actions in that range is $\alpha(1 - \underline{x} - p_R)$; when we add the αp_R point mass at x^*, we get that the actions played in $(\underline{x}, x^*]$ have measure $\alpha(1 - \underline{x})$. Since the total measure of actions in $[\underline{x}, x^*]$ is $1 - \underline{x}$, this leaves a point mass of $(1 - \alpha)(1 - \underline{x})$ at $\underline{x}$, which is the same as the point mass played in the equilibrium without caps. □

Corollary 3.6. *A spending cap in the range $(\tilde{x}, \bar{x})$ has no effect on the bias and reduces the maximal informational efficiency relative to the game without spending caps.*

Proof. As before, if we condition on L both candidates playing above $\underline{x}$, their strategy distributions are identical. Since the mass point it $\underline{x}$ remains unchanged, the bias also remains unchanged, since conditional on L not playing $\underline{x}$, both candidates win half the time, and, as before, L loses whenever he plays $\underline{x}$ against an opponent with type higher than $\underline{x}$.

The efficiency properties of the cases where at least one candidate is below $\underline{x}$, or both candidates are above $\underline{x}$ with L playing at $\underline{x}$) remain the same. The only case we need to examine, then, is when both candidates play strictly above $\underline{x}$. With the spending cap, we can no longer achieve perfect efficiency because of the mass points, whereas we could in the equilibrium without caps, so the error rate must increase. □

The implication of all this is that a spending cap can help to produce better outcomes, provided that it is suitably chosen. A spending cap that is binding but "too high" will do nothing to even the competition, and serves only to reduce the potential efficiency of the outcome. A lower cap can reduce the bias introduced by the cost asymmetry, but it has subtle effects on informational efficiency: effectively, it reduces the contest's ability to distinguish between types above the cap, but *improves* its ability to distinguish types above the cap from those below.

As a result, in the case of very severe asymmetries, a spending cap is clearly beneficial; when the asymmetries are not great, however, then there is a trade-off between informational efficiency and bias. How this should be valued will depend on the broader policy game; if we worry that candidates may choose less beneficial policies in order to attract lower their cost function, then we might value bias more than informational efficiency, since a random pick between two good options may be better than the better out of two bad ones.

3.6. Empirical predictions

Our model offers several new empirical predictions on the determinants of election outcomes and campaign spending. Although informational efficiency is at the center of much of our above analysis and discussion, it is not an easily measured quantity, so we turn our attention here to the determinants of the bias of the election.

The first set of empirical predictions is related to the effects of cost asymmetries. One effect of a greater incumbent cost advantage is that the incumbent advantage increases, in the sense that his winning probability is higher. In our model, this effect does not, however, come about through higher incumbent spending: rather, in equilibrium, the *challenger's* spending levels adjust downward, while the incumbent's equilibrium strategy remains the same. Hence, we should expect that, *ceteris paribus*, the greater the incumbent's cost advantage, the *lower* the level of challenger spending.

The second set of predictions is on the effects of spending caps. A unique feature of our model is the non-monotonic effect of spending caps: with very restrictive caps, the bias should decrease, but even at mild levels, where the cap is binding for both candidates, it has no effect on the expected outcome of the election.

The main challenge to overcome in an empirical test is the measurement of our key variable: the fund-raising capacities of candidates. There is no direct measurement available on the fund-raising *capacity* of candidates; only campaign spending is available, but that must form the dependent variable in our tests of the first prediction. We posit a couple tentative identification strategies, to be pursued in a later work. First, campaign finance policy, such as the *Citizens United* decision, should have different effects on different candidates; in particular, since it lifts the limit on single-source contributions, we should expect the costs of campaigns with narrower, more concentrated donor bases to decrease more dramatically, whereas campaigns that already drew form a greater variety of sources would see their marginal costs decrease by less. Hence, candidates who represent narrow, but concentrated interests may be more directly affected by recent liberalization policy. An advantage of this approach would be that contributors are likely uncorrelated with other important candidate characteristics that might impact election results; other observable candidate characteristics correlated with funding capacity may also be correlated with other aspects of candidate quality with a direct impact on elections.

3.7. Conclusion

In spite of the fairly wide-spread use of all-pay auctions to analyze various types of contests, they have not been applied extensively to political campaigns, where the literature has also produced several puzzles that have been difficult for models of electoral competition to accommodate.

This paper has hoped to shed a new perspective on some of the empirical evidence by producing a relatively straightforward model of the campaigning process that helps to explain these effects and to offer additional insights into how cost asymmetries manifest themselves in equilibrium, and hence also empirical estimates. We are able to explain several of the most puzzling asymmetries as equilibrium effects in our model, i.e., the incumbent advantage in winning probability, as well as the fact that incumbent spending appears to have less of an effect on election outcomes than does challenger spending. In addition, we explain wide asymmetries in spending levels, an issue that does not appear to garner as much attention is the differences in marginal impact.

Our all-pay model also offers new results on how different types of campaign finance reform might affect campaign outcomes. We find that symmetric liberalization, which simply leads to more campaigning, will generally be a good thing, but asymmetric liberalization may have some undesirable effects. Spending caps are effective at mitigating both the biases and informational distortions introduced by cost asymmetries, with the proviso that it is properly chosen: too low of a cap incurs a high cost on informational efficiency, while too high of a cap does nothing to even out the competition while degrading the informational efficiency of the contest.

References

Ambrus, A., and R. Argenziano. 2009. "Asymmetric Networks in Two-Sided Markets." *American Economic Journal: Microeconomics* 1 (1): 17–52.

Armstrong, M. 2006. "Competition in two-sided markets." *RAND Journal of Economics* 39 (3): 668–693.

Baye, M. R., D. Kovenock, and C. G. De Vries. 1996. "The all-pay auction with complete information." *Economic Theory* 8 (2): 291–305.

Benoit, K., and M. Marsh. 2008. "The campaign value of incumbency: A new solution to the puzzle of less effective incumbent spending." *American Journal of Political Science* 52 (4): 874–890.

Che, Y.-K., and I. L. Gale. 1998. "Caps on Political Lobbying." *The American Economic Review* 88 (3): 643–651.

Daley, B., and E. Snowberg. 2007. "A Multi-Dimensional Signaling Model of Campaign Finance." *Unpublished manuscript available at http://www. hss. caltech. edu/ snowberg/papers/Daley-Snowberg% 20Multi-dimensional4. pdf (accessed 10 November 2010).*

Diamond, P. A. 1971. "A model of price adjustment." *Journal of Economic Theory* 3 (2): 156–168.

Ellison, G., and D. Fudenberg. 2003. "Knife-Edge or Plateau: When Do Market Models Tip?" *The Quarterly Journal of Economics* 118 (4): 1249–1278.

Ellison, G., and A. Wolitzky. 2009. *A Search Cost Model of Obfuscation.* 15237. Working Paper. National Bureau of Economic Research.

Erikson, R. S., and T. R. Palfrey. 2000. "Equilibria in Campaign Spending Games: Theory and Data." *The American Political Science Review* 94 (3): 595–609.

Gale, D., and H. Nikaido. 1965. "The Jacobian matrix and global univalence of mappings." *Mathematische Annalen* 159 (2): 81–93.

Gerber, A. S. 2004. "Does Campaign Spending Work?: Field Experiments Provide Evidence and Suggest New Theory." *American Behavioral Scientist* 47 (5): 541–574.

Gius, M. 2009. "The Effects of Campaign Expenditures on Congressional Elections." *American Review of Political Economy:*51.

Hagiu, A. 2009. *Multi-sided platforms: From microfoundations to design and expansion strategies.* Technical report. Harvard Business School.

Hagiu, A., and B. Jullien. 2011. *Search diversion, rent extraction and competition.* 11-124. Technical report. HBS.

Kinder, D. R. 2003. "Communication and Politics in the Age of Information." In *Oxford Handbook of Political Psychology,* edited by D. O. Sears, L. Huddy, and R. Jervis, 357–393. Oxford University Press, USA.

Krishna, V. 2009. *Auction Theory.* Academic Press.

Lewis, G. 2011. "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors." *American Economic Review* 101 (4): 1535–1546.

Lewis, T. R., and D. E. M. Sappington. 1994. "Supplying Information to Facilitate Price Discrimination." *International Economic Review* 35 (2): 309–327.

McAfee, R. P., and J. McMillan. 1987. "Auctions with entry." *Economics Letters* 23 (4): 343–347.

Milgrom, P. R. 1981. "Good News and Bad News: Representation Theorems and Applications." *The Bell Journal of Economics* 12 (2): 380–391.

Moon, W. 2006. "The Paradox of Less Effective Incumbent Spending: Theory and Tests." *British Journal of Political Science* 36 (04): 705.

Mortensen, D. T. 1970. "Job Search, the Duration of Unemployment, and the Phillips Curve." *The American Economic Review* 60 (5): 847–862.

Mueller, D. C. 2003. *Public Choice III.* 3rd ed. Cambridge University Press.

Myerson, R. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6 (1): 58–73.

———. 2000. "Large Poisson Games." *Journal of Economic Theory* 94:7–45.

Petrocik, J. R. 1996. "Issue Ownership in Presidential Elections, with a 1980 Case Study." *American Journal of Political Science* 40 (3): 825–850.

Riley, J. G., and W. F. Samuelson. 1981. "Optimal auctions." *The American Economic Review* 71 (3): 381–392.

Rochet, J., and J. Tirole. 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1 (4): 990–1029.

———. 2006. "Two-sided markets: a progress report." *The RAND Journal of Economics* 37 (3): 645–667.

Rysman, M. 2009. "The Economics of Two-Sided Markets." *Journal of Economic Perspectives* 23 (3): 125–143.

Satterthwaite, M., and A. Shneyerov. 2007. "Dynamic Matching, Two-Sided Incomplete Information, and Participation Costs: Existence and Convergence to Perfect Competition." *Econometrica* 75 (1): 155–200.

Siegel, R. 2009. "All-Pay Contests." *Econometrica* 77 (1): 71–92.

Stokey, N. L., R. E. Lucas, and E. C. Prescott. 1989. *Recursive Methods in Economic Dynamics.* Harvard University Press.

Tadelis, S., and F. Zettelmeyer. 2011. "Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment." *Available at SSRN: http://ssrn.com/abstract=1872465 or http://dx.doi.org/10.2139/ssrn.1872465.*

Weyl, E. G. 2010. "A Price Theory of Multi-Sided Platforms." *American Economic Review* 100 (4): 1642–1672.

A. Appendix to Chapter 1

A.1. Proofs

This appendix includes the full proofs of details omitted over the course of the main text.

Proof of Lemma 1.1. Let n denote the number of buyers who shows up in a given auction, so $n \sim \text{Poisson}(\mu)$. For any given $n \geq k$, the probability that exactly k of them have valuations in X_1 is given by

$$P(k_1 = k \mid n) = \binom{n}{k} p(X_1)^k (1-p)^{n-k}$$

$k_1 \mid n$ is binomially distributed, since there are n buyers, each with an independent probability of being in X_1 of probability $p(X_1)$. Since n is Poisson distributed, we know that

$$P(n) = \frac{e^{-\mu}\mu^n}{n!}$$

The prior probability of exactly k buyers from X_1 is then simply

$$\begin{aligned}
P(k_1 = k) &= \sum_{n=k}^{\infty} P(n) \cdot P(k_1 = k \mid n) \\
&= \sum_{n=k}^{\infty} \frac{e^{-\mu}\mu^n}{n!} \binom{n}{k} p(X_1)^k (1-p)^{n-k} \\
&= \sum_{n=k}^{\infty} \frac{e^{-\mu}\mu^n}{n!} \frac{n!}{k!\,(n-k)!} p(X_1)^k (1-p)^{n-k} \\
&= \frac{e^{-p(X_1)\mu}\mu^k}{k!} p(X_1)^k \sum_{n=k}^{\infty} \frac{e^{-(1-p(X_1))\mu}\mu^{n-k}}{(n-k)!} (1-p(X_1))^{n-k} \\
&= \frac{e^{-p(X_1)\mu}(p(X_1)\mu)^k}{k!} \sum_{n=k}^{\infty} \frac{e^{-(1-p(X_1))\mu}((1-p(X_1))\mu)^{n-k}}{(n-k)!} \\
&= \frac{e^{-\mu(X_1)}\mu(X_1)^k}{k!}
\end{aligned}$$

To see independence, consider the probability that $k_1 = l$ and $k_2 = m$, again letting n denote the total number of buyers in an auction. Any given buyer is in X_1 with probability $p(X_1)$, in X_2

with probability $p(X_2)$, and in neither with probability $1 - p(X_1) - p(X_2)$. Given n, The number of buyers in each of these groups – X_1, X_2, and $X \setminus (X_1 \cup X_2)$ – is multinomially distributed, so

$$P(k_1 = l \text{ and } k_2 = m \mid n) = \frac{n!}{k!\,l!\,(n-k-l)!} p(X_1)^l p(X_2)^m (1 - p(X_1) - p(X_2))^{n-l-m}$$

Taking the probability weighted sum over all n gives the desired expression:

$$\begin{aligned} P(k_1 = l \text{ and } k_2 = m) &= \sum_{n=k+l}^{\infty} \frac{e^{-\mu}\mu^n}{n!} \frac{n!}{k!\,l!\,(n-k-l)!} p(X_1)^l p(X_2)^m \left(1 - p(X_1) - p(X_2)\right)^{n-l-m} \\ &= \frac{e^{-\mu p(X_1)}(\mu p(X_1))^l}{l!} \cdot \frac{e^{-\mu p(X_2)}(\mu p(X_2))^m}{m!} \\ &\quad \cdot \sum_{n=l+m}^{\infty} \frac{e^{-\mu(1-p(X_1)-p(X_2))}(\mu(1 - p(X_1) - p(X_2)))^{n-l-m}}{(n-l-m)!} \\ &= \frac{e^{-\mu(X_1)}\mu(X_1)^l}{l!} \cdot \frac{e^{-\mu(X_2)}\mu(X_2)^m}{m!} \\ &= P(k_1 = l) \cdot P(k_2 = m) \end{aligned}$$

□

Proof of Lemma 1.2. The proof involves separately considering the number of buyers who show up with valuations on disjoint subsets of X around x.

Suppose that $x' > x$. Then the probability that the highest bid among the other buyers is below x' is simply the probability that no bidder with valuation in $X' = [x', \overline{x}]$ shows up. There is a measure $\mu\left(1 - F(x)\right)$ of such bidders, and their number is Poisson distributed with parameter $\mu\left((1 - F(x)\right)$, so the probability that zero of them shows up is $e^{-\mu\left(1-F(x)\right)}$, as desired.

If $x' \leq x$, then the probability that the highest bid among the others is below x' is the probability that zero bidders show up in the ranges $[x', x)$ and $(x, \overline{x}]$.

We are able to exclude x itself since the conditional probability of no other buyers of valuation x showing up is one. To see this, consider a small neighborhood of valuations around x; let μ' be the measure of buyers with valuations in this neighborhood, and let k' be a random variable denoting the number of buyers who shows up in a given auction. We know from above that $k' \sim$ Poisson(μ'). A buyer with valuation x conditions on his own presence, so the probability that no other buyers with valuations in this neighborhood appear is the probability that $k' = 1$

conditional on $k' \geq 1$, which is

$$P(k' = 1 \,|\, k' \geq 1) = \frac{\mu' e^{-\mu'}}{1 - e^{-\mu'}}$$

by Bayes' rule. The conditional probability when we condition exactly on x is the limit of this as $\mu' \to 0$ (we make the neighborhood tighter and tighter around x). Since the above expression is indeterminate at $\mu' = 0$, we apply L'Hôpital's rule to find the limit:

$$\begin{aligned}
\lim_{\mu' \to 0} \frac{\mu' e^{-\mu'}}{1 - e^{-\mu'}} &= \lim_{\mu' \to 0} \frac{\frac{d}{d\mu'}\left(\mu' e^{-\mu'}\right)}{\frac{d}{d\mu'}\left(1 - e^{-\mu'}\right)} \\
&= \lim_{\mu' \to 0} \frac{e^{-\mu'} - \mu' e^{-\mu'}}{e^{-\mu'}} \\
&= \lim_{\mu' \to 0} 1 - \mu' \\
&=1
\end{aligned}$$

This means that we only need to consider intervals that do not include x itself, $[x', x)$ and $(x, \overline{x}]$. The measure of buyers in the first range is $\mu\left(F(x) - F(x')\right)$, and that of buyers in the second range is $\mu\left(1 - F(x)\right)$. Since these are independent events, the probability of both being true is

$$e^{-\mu\left(F(x) - F(x')\right)} e^{-\mu\left(1 - F(x)\right)} = e^{-\mu\left(1 - F(x')\right)}$$

which proves the result. □

Proof of Lemma 1.3. Buyers have quasilinear utility, so $u(x) = xG(x) - m(x)$, where the first term is simply the benefit from winning, x, times the probability of winning, and the second term is the expected payment for a buyer of valuation x. The expected payment is equal to the expectation of the highest bid among other buyers, conditional on that highest being below x. An application of Bayes' rule gives

$$\begin{aligned}
m(x) &= \int_0^x yg(y)\, dy \\
&=xG(x) - \int_0^x G(y)\, dy
\end{aligned}$$

where the second line follows from integration by parts, letting $u = x$ and $dv = g(y)\, dy$. Substituting this in to $u(x)$ gives the desired result.

The expected revenue is equal to the expectation of the second highest bid, which we denote

X_2. We can calculate the cumulative distribution function for X_2 as the probability that zero or one buyer shows up from the subset $(x, 1]$ for any x:

$$G_2(x) = G(x) + \mu(1 - F(x))G(x)$$

where the first term is the probability that no buyers show up with valuations higher than x, and the second is the probability that exactly one such buyer shows up. The density function is then

$$\begin{aligned} g_2(x) =& g(x) + \left[-\mu f(x)G(x) + \mu(1 - F(x))g(x)\right] \\ =& \mu(1 - F(x))g(x) \end{aligned}$$

The second line follows from the fact that $g(x) = \mu f(x)e^{-\mu(1-F(x))} = \mu f(x)G(x)$. The expectation of X_2 is then

$$\begin{aligned} E[X_2] =& \int_0^{\overline{x}} x g_2(x)\, dx \\ =& \int_0^{\overline{x}} \mu(1 - F(x))x g(x)\, dx \\ =& \mu(1 - F(x))xG(x)|_0^{\overline{x}} - \int_0^{\overline{x}} \left[-\mu f(x)xg(x) + \mu(1 - F(x))G(x)\right]\, dx \\ =& m(1) - \mu \int_0^{\overline{x}} (1 - F(x))G(x)\, dx \end{aligned}$$

The transition from the second to third lines again integrates by parts, with $u = \mu(1 - F(x))x$ and $dv = g(x)\, dx$. The final line follows from the third by noticing that $\int_0^{\overline{x}} \mu f(x)xG(x)\, dx = \int_0^{\overline{x}} xg(x)\, dx$, which is the expected payment of the highest valuation buyer. □

Proof of Theorem 1.4. Consider a buyer x for whom $u^A(x) - u^B(x) > 0$, i.e., who prefers to be in an A auction, and suppose he is randomly matched into a B auction. Let v denote his continuation payoff from searching, so that he will search if and only if $v - c > u^B(x)$. The only step we need to do is to solve for v.

If he searches, with probability p_a, he will be matched with an A seller (with probability p, he gets his preferred match for certain, and with probability $1 - p$ he is randomly matched, which gives his preferred match with probability a). Hence, with probability $1-p_a$ he will have to search again, which yields $v - c$. This gives

$$\begin{aligned} v &= p_a u^A(x) + (1 - p_a)(v - c) \\ \Leftrightarrow p_a v &= p_a u^A(x) - (1 - p_a)c \\ \Leftrightarrow v &= u^A(x) - \frac{1 - p_a}{p_a} c \end{aligned}$$

Substituting back into the decision rule shows that search is optimal if only if

$$\begin{aligned} v - c &> u^B(x) \\ \Leftrightarrow u^A(x) - \left(1 + \frac{1 - p_a}{p_a}\right) c &> u^B(x) \\ \Leftrightarrow u^A(x) - u^B(x) &> \frac{c}{p_a} \end{aligned}$$

The proof is analogous for searching into B. □

Proof of Theorem 1.2. As described in the main text, existence follows from the continuity of the mapping from a pair of thresholds to the types that are indifferent between searching and not searching – all equilibria are fixed points of such a mapping, and all fixed points are equilibria, so we can use Brouwer's fixed point theorem to establish the existence of a fixed point, and hence equilibrium. The main step is to establish continuity.

Here we will introduce a little bit of additional notation, since we will need to consider the market for A goods separately from the market for B goods. Let us define μ^A and μ^B to be the measures of buyers in the A and B markets. Remember from our derivation that the appropriate measure is the ratio of buyers to sellers (when we take the limit of discrete games, the ratio of buyers to sellers goes to μ), so these correspond to the measures of buyers who *end up* in A and B auctions, taking into account the search strategies, divided by the measure of sellers of each type. We also define $\mu^A(x)$, $\mu^A(X)$, $\mu^B(x)$, and $\mu^B(X)$ analogously to μ (as the density and measure functions in the markets).

If the thresholds are (x_b, x_a), then we can get easily get μ^A and μ^B. The easiest way to see this is to consider the density of buyers in each market. If $x \in (x_b, x_a)$, so that buyers at x are not sorting, then a fraction a go into A auctions, which have measure a; hence, the normalized density of buyers, from the point of view of A sellers, remains the same, i.e., $\mu^A(x) = \mu(x)$ for $x \in (x_b, x_a)$. On the other hand, for $x > x_a$, all the buyers end up in A auctions, but there is only a measure

a of sellers, so the normalized density is $\mu(x)/a$. Finally, no buyers of type $x < x_b$ show up in A auctions, since they all search into B auctions, so the density is zero. We can then write the total measure in each auction, given thresholds (x_b, x_a), as

$$\mu^A = \mu\left(F(x_a) - F(x_b) + \frac{1}{a}\left(1 - F(x_a)\right)\right)$$

$$\mu^B = \mu\left(F(x_a) - F(x_b) + \frac{1}{1-a}F(x_b)\right)$$

Since we know the measures of types in each auction, we can simply apply the results from the single-good model to get expressions for utility.

We show this below in lemma A.1, which gives the closed form expressions for the first-order statistic distribution. It is easy to see that they are continuous in x_b and x_a (since the measure on buyers is massless, variation is continuous), and utility is a continuous function of the distribution of the first-order statistic; as a result, s is continuous. This delivers existence. □

Lemma A.1. *Let G^A and G^B denote the distributions of the first-order statistics in auctions for each type of good. Then*

$$G^A(x) = e^{-\mu^A(1-F^A(x))}$$

and

$$G^B(y) = e^{-\mu^B(1-F^B(y))}$$

where

$$\mu^A = \frac{1}{a} - \left(\frac{1}{a} - 1\right)F(x_a) - F(x_b)$$

$$\mu^B = F(x_a) + \frac{a}{1-a}F(x_b)$$

$$F^A(x) = \begin{cases} 0 & \textit{for } x < x_b \\ \frac{\mu}{\mu^A}\left(F(x) - F(x_b)\right) & \textit{for } x_b \leq x < x_a \\ \frac{\mu}{\mu^A}\left[\frac{1}{a}F(x) - \left(\frac{1}{a} - 1\right)F(x_a) - F(x_b)\right] & \textit{for } x_a \leq x < \overline{x} \\ 1 & \textit{for } x > \overline{x} \end{cases}$$

$$F^B(y) = \begin{cases} 0 & \textit{for } y < y_a \\ \frac{\mu}{\mu^B}\left(H(y) - H(y_a)\right) & \textit{for } y_a \le y < y_b \\ \frac{\mu}{\mu^B}\left[\frac{1}{1-a} - \left(\frac{1}{1-a} - 1\right) H(y_b) - H(y_a)\right] & \textit{for } y_b \le y < \overline{y} \\ 1 & \textit{for } y > \overline{y} \end{cases}$$

where $H(y) = 1 - F(y^{-1}(x))$ is the population-wide distribution of B good valuations, $\overline{y} = y(\underline{x})$, $\underline{y} = y(\overline{x})$, $y_a = y(x_a)$, and $y_b = y(x_b)$ are the B good valuations for the respective buyers.

Proof. As before, we will prove the result for the A market. The total measure of buyers in the A market, normalized by sellers, is given by μ^A, since

$$\begin{aligned} \mu^A =& \mu(X_0) + \frac{1}{a}\mu(X_a) \\ =& \mu\left[F(x_a) - F(x_b)\right] + \frac{\mu}{a}\left[1 - F(x_a)\right] \\ =& \mu\left[\frac{1}{a} - \left(\frac{1}{a} - 1\right) F(x_a) - F(x_b)\right] \end{aligned}$$

Let $\nu(x)$ be the measure of buyer above x for $x \in [x_b, \overline{x}]$, then

$$\nu(x) = \begin{cases} \frac{\mu}{a}\left[1 - F(x)\right] & \text{for } x > x_a \\ \mu\left[F(x_a) - F(x)\right] + \frac{\mu}{a}\left[1 - F(x_a)\right] & \text{for } x \le x_a \end{cases}$$

The probability that x is above the highest bid in a given auction is the probability that buyer with valuation $x' > x$ shows up, which is $e^{-\nu(x)}$. It is straightforward to verify that this is equal to the expression given.

The form given emphasizes the role that independence plays in our analysis: since the number of buyers showing up from each interval is independent, the above is equivalent to thinking of A auctions as a single-good platform where the total measure of buyers is μ^A and valuations are independently and identically distributed according to the population frequencies (since F^A is simply the measure function scaled by μ^A, the total measure. □

Proof of Theorem 1.2 continued: Uniqueness and Monotonicity. To show monotonicity and uniqueness, let

$$\mathbf{F}(x_b, x_a) = \begin{bmatrix} s(x_b; x_b, x_a) \\ s(x_a; x_b, x_a) \end{bmatrix}$$

Then equilibrium is defined by $\mathbf{F}(x_b, x_a) = (-c/p_b, c/p_a)$. By the Gale-Nikaido theorem (Gale and Nikaido 1965), this is unique if the Jacobian of $\mathbf{F}$ has positive determinant; furthermore, we can establish monotonicity by signing the terms of the inverse Jacobian.

First, recall that $u^A(x) = \int_0^x G^A(x')\,dx'$, and note that changes to either threshold, x_b or x_a, only affects the winning probabilities of buyers with types below those thresholds. The derivatives of u^A and u^B with respect to the threshold values will follow almost immediately from the derivatives of G^A and G^B with respect to x_b and x_a. establish these, recall from Lemma A.1 that

$$G^A(x) = e^{-\mu^A(1-F^A(x))}$$

Let $\nu(x)$ be the measure of buyers in A markets with valuations above x, then

$$\nu(x) = \begin{cases} \frac{\mu}{a}\left[(1-F(x)\right] & \text{for } x > x_a \\ \frac{\mu}{a}\left[1-F(x_a)\right] + \mu\left[F(x_a)-F(x)\right] & \text{for } x_b < x < x_a \\ \frac{\mu}{a}\left[1-F(x_a)\right] + \mu\left[F(x_a)-F(x_b)\right] & \text{for } x \leq x_b \end{cases}$$

Taking the derivatives with respect to x_b and x_a yields

$$\frac{d\nu(x)}{dx_b} = \begin{cases} -\mu f(x_b) & \text{for } x \leq x_b \\ 0 & \text{otherwise} \end{cases}$$

$$\frac{d\nu(x)}{dx_a} = \begin{cases} -\mu\left(\frac{1}{a}-a\right)f(x_a) & \text{for } x > x_a \\ 0 & \text{for } x < x_a \end{cases}$$

where both inequalities are strict in the cases for $d\nu(x)/dx_a$ because the change at *precisely* x_a depends on the direction in which x_a is changing.

From this, we can see how G^A changes with x_b and x_a:

$$\frac{dG^A(x)}{dx_b} = \begin{cases} \mu f(x_b)G^A(x_b) & \text{for } x \leq x_b \\ 0 & \text{otherwise} \end{cases}$$

$$\frac{dG^A(x)}{dx_a} = \begin{cases} \left(\frac{1}{a}-1\right)\mu f(x_a)G^A(x) & \text{for } x < x_a \\ 0 & \text{for } x > x_a \end{cases}$$

Using these expressions in u^A, we get

$$
\begin{aligned}
\frac{d}{dx_b}u^A(x_b) &= \frac{d}{dx_b}\int_0^{x_b} G^A(x)\,dx \\
&= G^A(x_b) + \int_0^{x_b} \mu f(x_b) G^A(x)\,dx \\
&= G^A(x_b) + \mu f(x_b) u^A(x_b) \\
\frac{d}{dx_a}u^A(x_b) &= \frac{d}{dx_b}\int_0^{x_b} G^A(x)\,dx \\
&= \frac{1-p}{p}\mu f(x_a) u^A(x_b) \\
\frac{d}{dx_b}u^A(x_a) &= \frac{d}{dx_b}\int_0^{x_a} G^A(x)\,dx \\
&= \mu f(x_b) u^A(x_b) \\
\frac{d}{dx_a}u^A(x_a) &= \frac{d}{dx_a}\int_0^{x_a} G^A(x)\,dx \\
&= G^A(x_a) + \frac{1-p}{p}\mu f(x_a) u^A(x_a)
\end{aligned}
$$

Utilities on the B side can be expressed equivalently using B valuations (i.e., writing utilities as functions of y_b and y_a) and B densities; the derivatives with respect to x will involve an additional y' term. Further simplifications are can be made by observing that $h(y)|y'(x)| = f(x)$, which we substitute in to the expressions for $ds(\cdot)/dx_i$ to get:

$$
\begin{aligned}
\frac{d}{dx_b}s(x_b) &= G^A(x_b) + |y'(x_b)|G^B(x_b) + \mu f(x_b)\left[u^A(x_b) + \frac{p}{1-p}u^B(x_b)\right] \\
\frac{d}{dx_a}s(x_b) &= \mu f(x_a)\left[\frac{1-p}{p}u_a(x_b) + u_b(x_a)\right] \\
\frac{d}{dx_b}s(x_a) &= \mu f(x_b)\left[\frac{p}{1-p}u_b(x_a) + u_a(x_b)\right] \\
\frac{d}{dx_a}s(x_a) &= G^A(x_a) + |y'(x_a)|G^B(x_a) + \mu f(x_a)\left[\frac{1-p}{p}u^A(x_a) + u^B(x_a)\right]
\end{aligned}
$$

The determinant of the Jacobian is equal to

$$
\frac{ds(x_b)}{dx_b}\frac{ds(x_a)}{dx_a} - \frac{ds(x_b)}{dx_a}\frac{ds(x_a)}{dx_b}
$$

To see that this is strictly positive, consider only the final terms of $ds(x_b)/dx_b$ and $ds(x_a)/dx_a$:

$$\begin{aligned}
\frac{d}{dx_b}s(x_b) &= K_b + \mu f(x_b)\left[u^A(x_b) + \frac{p}{1-p}u^B(x_b)\right] \\
&\geq K_b + \mu f(x_b)\left[u^A(x_b) + \frac{p}{1-p}u^B(x_a)\right] \\
&= K_b + \frac{ds(x_a)}{dx_b}
\end{aligned}$$

$$\begin{aligned}
\frac{d}{dx_a}s(x_a) &= K_a + \mu f(x_a)\left[\frac{1-p}{p}u^A(x_a) + u^B(x_a)\right] \\
&\geq K_a + \mu f(x_a)\left[\frac{1-p}{p}u^A(x_b) + u^B(x_a)\right] \\
&= K_b + \frac{ds(x_b)}{dx_a}
\end{aligned}$$

where K_b and K_a are strictly positive and represent all terms other than the last two. Since all other terms are positive, multiplying these two and subtracting the cross term yields a determinant strictly greater than zero.

Finally, since the determinant of the Jacobian is positive, the inverse Jacobian is signed as follows (using standard expressions for the inverse of a two-by-two matrix):

$$\mathrm{sgn}(J^{-1}) = \begin{pmatrix} + & - \\ - & + \end{pmatrix}$$

Since $d(x_b, x_a)/dc = J^{-1} \cdot d(-c/p_b, c/p_a)/dc$, this tells us that $dx_b/dc < 0$ and $dx_a/dc > 0$, as desired □

Proof of Lemma 1.6. Suppose not. Consider the set of individuals that is assigned to sort into A and those that are assigned to sort into B (we can consider any sets of two possible strategies and the analysis follows similarly). If the optimal assignment does not have the structure described in the lemma, then we can find some set X^A of buyers who sort into A, and another set X^B of buyers who sort into B, each of measure $\varepsilon > 0$, such that $\inf X^B > \sup X^A$. That is, there is a set of buyers who sort into B each with strictly higher x than another set of buyers who sort into A of same measure.

Consider what happens if we change the assignments of X^A and X^B; that is, we construct an alternative assignment that leaves all other search decisions the same, but tells those in X^A to search into B and those in X^B to search into A. Since X^A and X^B have the same measure, the alternative assignment results in the same measure of bidders in each market. Furthermore, the value distribution in each auction first order stochastically dominates the value distribution in the original assignment, since the change is to have replaced a measure of low valuation buyers with an equal measure of higher valuation buyers. This means that the distribution of the first-order statistics in the alternative distribution also first-order stochastically dominate the first-order statistic distributions of the original assignment, since

$$G^A(x) = e^{-\mu^A(1-F^A(x))}$$

and μ^A stays the same in the alternative assignment. This means that $E[x_1]$ and $E[y_1]$ are greater under the alternative assignment, and search costs remain the same, so total welfare is raised by switching strategies. This contradicts the original assumption that the allocation was socially optimal. □

Proof of Theorem 1.3. Recall the expressions for welfare and the derivatives of G^A and G^B with respect to x_b and x_a:

$$W = a\left(\overline{x} - \int_0^{\overline{x}} G^A(x)\,dx\right) + (1-a)\left(\overline{y} - \int_0^{\overline{y}} G^B(y)\,dy\right) - \left(\frac{(1-a)\mu_a}{p_a} + \frac{a\mu_b}{p_b}\right)c$$

$$\frac{d}{dx_b}G^A(x) = \begin{cases} \mu f(x_b)G^B(x) & \text{for } x \leq x_b \\ 0 & \text{otherwise} \end{cases}$$

$$\frac{d}{dx_a}G^A(x) = \begin{cases} \mu f(x_a)\left(\frac{1}{a} - 1\right)G^A(x) & \text{for } x < x_a \\ 0 & \text{for } x > x_a \end{cases}$$

$$\frac{d}{dx_b}G^B(y) = \begin{cases} -\mu f(x_b)\left(\frac{1}{1-a} - 1\right)G^B(y) & \text{for } y < y(x_b) \\ 0 & \text{for } y > y(x_b) \end{cases}$$

$$\frac{d}{dx_a}G^B(y) = \begin{cases} -\mu f(x_a)G^B(y) & \text{for } y \leq y(x_b \\ 0 & \text{otherwise} \end{cases}$$

This gives

$$\begin{aligned}\frac{dW}{dx_b} =& a\mu f(x_b)\left[-\int_0^{x_b} G^A(x)\,dx + \int_0^{y_b} G^B(y)\,dy - \frac{c}{p_b}\right]\\ =& a\mu f(x_b)\left[-\left(u^A(x_b) - u^B(x_b)\right) - \frac{c}{p_b}\right]\end{aligned}$$

$$\begin{aligned}\frac{dW}{dx_a} =& (1-a)\mu f(x_a)\left[-\int_0^{x_a} G^A(x)\,dx + \int_0^{y_b} G^B(y)\,dy - \frac{c}{p_a}\right]\\ =& (1-a)\mu f(x_a)\left[-\left(u^A(x_a) - u^B(x_a)\right) - \frac{c}{p_a}\right]\end{aligned}$$

The terms in brackets are easily signed: when x_b is below the threshold of the market solution, dW/dx_b is positive since $u^A(x_b) - u^B(y_b) < c/p_b$, and similarly, when x_a is above the market solution threshold, dW/dx_a is negative. A quick inspection shows that the first order conditions for the social planner are equivalent to the threshold equilibrium conditions, so we have ($x_b^* = x_b$ and $x_a^* = x_a$, i.e., the social optimum coincides with the market solution. Hence, the search equilibrium is efficient in the sense that it maximizes total welfare. □

A.2. Numerical Implementation

This section contains details regarding our implementation of the numerical results. The data was generated using code written in C++; performance and accuracy considerations ruled out the use of Matlab or Mathematica, whose built-in optimization routines performed poorly for our setting. To avoid rewriting widely used mathematical operations, we used two third-party open-source libraries: GSL for numerical integration and NLopt for minimization. Both libraries have been extensively tested for correctness, so we can be confident of our results even for small magnitudes.

The implementation is quite straightforward: for sets of parameters (t_b, t_a), we define utility functions, whose forms are given in the main text, and an error function, $\varepsilon(t_b, t_a)$ as

$$\varepsilon(t_b, t_a) = \varepsilon_b(t_b, t_a) + \varepsilon_a(t_b, t_a) \tag{A.1}$$

where

$$\varepsilon_b(t_b, t_a) = \begin{cases} |s_b(t_b, t_a) - c_b| & \text{for } t_b > 0 \\ 0 & \text{otherwise} \end{cases} \tag{A.2}$$

and

$$\varepsilon_a(t_b, t_a) = \begin{cases} |s_a(t_b, t_a) - c_a| & \text{for } t_a < 1 \\ 0 & \text{otherwise} \end{cases} \tag{A.3}$$

The error function is minimized via the SBPLX algorithm, and the numerical minimization procedure is terminated when $\varepsilon < 10^{-9}$.

B. Appendix to Chapter 2

B.1. Proofs

Proof of Theorem 2.4. If the marginal participant searches, then lowering search costs increases participation without changing the behaviors of anybody already on the platform, so revenues must increase.

The non-trivial case is when the marginal participant does *not* search, so that a decrease in search costs also decreases participation. When total participation does not change, a positive revenue effect from reduced search costs is easy to show because, from the point of view of the seller, low valuation buyers are replaced with higher valuation buyers, with no net change in the measure of buyers; as a result, revenue must increase. When there is a binding reservation utility, however, the effects are more subtle: the lowest valuation buyers search out, and are replaced by an equal measure of higher valuation buyers who search in; at the same time, however, the increased competitiveness of the platform drives out more buyers in the middle. Hence, sellers have more higher valuation buyers, but face a net decrease in the measure of buyers.

Since platform is symmetric, revenue in A and B markets are equal, hence it suffices to show that revenue in market A increases as search costs decrease. In particular, we will show that $dR_a/dc^* < 0$ when $0 < \theta_b < \theta_u < 1/2$[1]. We already have expressions for $d\theta_b/dc^*$ and $d\theta_u/dc^*$, so if we express revenue as a function of θ_b and θ_u we can apply the chain rule and get an expression directly.

First, consider the effects of changes in threshold values on the revenue. Let

$$\beta \equiv u_a(1) + \int_0^1 u_a(\theta) f_a(\theta)\, d\theta$$

[1] Note that if nobody is searching, and $\theta_b = 0$, then decreasing search costs has no effect on behavior or participation. If everybody is participating, and $\theta_u = 1/2$, then there is no net exit as search costs decrease. We restrict our attention to the case where both equilibrium conditions are binding. The analysis also applies to cases where $\theta_b = 0$ and/or $\theta_u = 1/2$, as long as the equilibrium constraints are binding in the direction of change of c^*.

so that $R = 1 - \beta$, and $dR/dj = -d\beta/dj$ for any variable j. Recall that

$$u_a(\theta) = \int_0^\theta G_a(v)\,dv$$

$$G_a(\theta) = e^{-\eta_a(\theta)}$$

$$\eta_a(\theta) = \int_\theta^1 f_a(t)\,dt$$

$$f_a(\theta) = \begin{cases} 0 & \text{for } \theta < \theta_b \text{ or } \theta \in (\theta_u, 1-\theta_u) \\ 2\mu & \text{for } \theta > \theta_a = 1 - \theta_b \\ \mu & \text{otherwise} \end{cases}$$

which gives a utility function

$$u_a(\theta) = \begin{cases} \theta G_a(\theta_b) & \text{for } \theta \in [0, \theta_b] \\ u_a(\theta_b) + \frac{1}{\mu}\left(G_a(\theta) - G_a(\theta_b)\right) & \text{for } \theta \in (\theta_b, \theta_u] \\ u_a(\theta_u) + (\theta - \theta_u)G_a(\theta_u) & \text{for } \theta \in (\theta_u, \theta_v] \\ u_a(\theta_v) + \frac{1}{\mu}\left(G_a(\theta) - G_a(\theta_v)\right) & \text{for } \theta \in (\theta_v, \theta_a] \\ u_a(\theta_a) + \frac{1}{2\mu}\left(G_a(\theta) - G_a(\theta_a)\right) & \text{for } \theta \in (\theta_a, 1] \end{cases}$$

It is straightforward to see that

$$\eta_a(\theta_b) = \mu 2\theta_u$$

$$\eta_a(\theta_u) = \mu(\theta_b + \theta_u) = \eta_a(\theta_v)$$

$$\eta_a(\theta_a) = \mu 2\theta_b$$

so

$$u_a(1) = \theta_b e^{-\mu(2\theta_u)} + (1 - 2\theta_u)e^{-\mu(\theta_b+\theta_u)} + \frac{1}{\mu}\left[\frac{1}{2}\left(1 + e^{-\mu(2\theta_b)}\right) - e^{-\mu(2\theta_u)}\right]$$

Also note that the expressions for utility in all regions where $f_a(\theta) > 0$, i.e., buyers of that type participate, can be expressed as the sum of a constant term and term that depends on type, where the type-dependent term is always of the form $G_a(\theta)/f_a(\theta)$; consequently the integral $\int_0^1 u_a(\theta) f_a(\theta)\,d\theta$ can be written as

$$K + \int_{\theta_b}^{\theta_u} G_a(\theta)\,d\theta + \int_{\theta_v}^{\theta_a} G_a(\theta)\,d\theta + \int_{\theta_a}^{1} G_a(\theta)\,d\theta$$

where K is the sum of each constant term multiplied by the measure of buyers who have that term in their utility. Fully written out, we have

$$\begin{aligned}\int_0^1 u_a(\theta)f_a(\theta)\,d\theta =&\theta_b e^{-\mu(2\theta_u)}\left[\mu(2\theta_u)\right] + (1-2\theta_u)e^{-\mu(\theta_b+\theta_u)}\left[\mu(\theta_b+\theta_u)\right]\\ &+\frac{1}{2\mu}e^{-\mu(2\theta_b)}\left[\mu(2\theta_b)\right] - \frac{1}{\mu}e^{-\mu(2\theta_u)}\left[\mu(2\theta_u)\right]\\ &+\frac{1}{\mu}\left[\frac{1}{2}\left(1+e^{-\mu(2\theta_b)}\right) - e^{-\mu(2\theta_u)}\right]\end{aligned}$$

where the final term is the integral of the type-dependent terms in $G_a(\theta)f_a(\theta)$, and the square-bracketed terms in all the preceding terms represent the measure of buyers with the associated constant term in their utilities. Adding in $u_a(1)$ gives

$$\begin{aligned}\beta =&\frac{1}{\mu} + \theta_b e^{-\mu(2\theta_u)}\left[\mu(2\theta_u)+1\right]\\ &+(1-2\theta_u)e^{-\mu(\theta_b+\theta_u)}\left[\mu(\theta_b+\theta_u)+1\right]\\ &+\frac{1}{2\mu}e^{-\mu(2\theta_b)}\left[\mu(2\theta_b)+2\right]\\ &-\frac{1}{\mu}e^{-\mu(2\theta_u)}\left[\mu(2\theta_u)+2\right]\end{aligned}$$

The partial derivatives of this are

$$\begin{aligned}pd(\beta,\theta_b) =& -\Big[e^{-\mu(2\theta_b)}(\mu 2\theta_b+1) - e^{-\mu(2\theta_u)}(\mu 2\theta_u+1)\\ &+\left[\mu(1-2\theta_u)\right]\left[\mu(\theta_b+\theta_u)\right]e^{-\mu(\theta_b+\theta_u)}\Big]\\ pd(\beta,\theta_u) =& -\Big[2\left(e^{-\mu(\theta_b+\theta_u)}(\mu(\theta_b+\theta_u)+1) - e^{-\mu(2\theta_u)}(\mu 2\theta_u+1)\right)\\ &+\left[\mu(1-2\theta_u)\right]\left[\mu(\theta_b+\theta_u)\right]e^{-\mu(\theta_b+\theta_u)}\\ &-(\mu 2\theta_b)(\mu 2\theta_u)e^{-\mu 2\theta_u}\Big]\end{aligned}$$

The comparative statics showed that

$$\begin{pmatrix}\frac{d\theta_b}{dc^*}\\ \\ \frac{d\theta_u}{dc^*}\end{pmatrix} \propto \begin{pmatrix}pd(F_1,\theta_u)\\ \\ -pd(F_1,\theta_b)\end{pmatrix}$$

That is, the left-hand side vector is proportional to the right-hand side. Taking the derivative of β with respect to c^* using the chain rule gives

$$\frac{d\beta}{dc^*} = pd(\beta,\theta_b)d(\theta_b,c^*) + pd(\beta,\theta_u)d(\theta_u,c^*) \propto pd(\beta,\theta_b)pd(F_1,\theta_u) - pd(\beta,\theta_u)pd(F_1,\theta_b)$$

The right hand side expands to

$$\left[\overbrace{e^{-\mu(2\theta_b)}(\mu 2\theta_b+1)-e^{-\mu(2\theta_u)}(\mu 2\theta_u+1)}^{x_1}+\overbrace{\left[\mu(1-2\theta_u)\right]\left[\mu(\theta_b+\theta_u)\right]e^{-\mu(\theta_b+\theta_u)}}^{x_2}\right]\times$$

$$\times\left[\overbrace{2\left(e^{-\mu(\theta_b+\theta_u)}-e^{-\mu(2\theta_u)}\right)}^{y_1}+\overbrace{\frac{1}{2}\mu(1-2\theta_u)e^{-\mu(\theta_b+\theta_u)}}^{y_2}+\mu(2\theta_b)e^{-\mu(2\theta_u)}\right]$$

$$-\left[\overbrace{2\left(e^{-\mu(\theta_b+\theta_u)}(\mu(\theta_b+\theta_u)+1)-e^{-\mu(2\theta_u)}(\mu 2\theta_u+1)\right)}^{x_1'}+\right.$$

$$\left.\overbrace{\left[\mu(1-2\theta_u)\right]\left[\mu(\theta_b+\theta_u)\right]e^{-\mu(\theta_b+\theta_u)}}^{x_2'}-(\mu 2\theta_b)(\mu 2\theta_u)e^{-\mu 2\theta_u}\right]\times$$

$$\times\left[\overbrace{e^{-\mu(\theta_b+\theta_u)}-e^{-\mu(2\theta_u)}}^{y_1'}+\overbrace{\frac{1}{2}\mu(1-2\theta_u)e^{-\mu(\theta_b+\theta_u)}}^{y_2'}\right]$$

To sign this, first note that

$$d(,x)e^{-x}(x+1)=-xe^{-x}<0\quad\text{for } x>0$$

so the terms labeled x_1 and x_1' are greater than zero, and $x_1 > (1/2)x_1'$. This implies that all the terms in large square brackets are greater than zero; consequently the sum of all products involving the unlabeled terms is greater than zero. $x_2y_2 = x_2'y_2'$, $x_1y_1 > x_1'y_1'$, and $x_2y_1 > x_2'y_1'$, so the entire expression, which is proportional to $d\beta/dc^*$, is greater than zero. Finally,

$$d(R,c^*)=-d(\beta,c^*)<0$$

as desired. □