This Draft: January 15, 2012

The Economics of Internet Markets

Jonathan Levin*
Stanford University
and NBER

Abstract. The internet has facilitated the creation of new markets characterized by large scale, increased customization, rapid innovation and the collection and use of detailed consumer and market data. I describe these changes and some of the economic theory that has been useful for thinking about online advertising markets, e-commerce and other internet platforms. I also discuss the empirical evidence on competition and consumer behavior in internet markets and some directions for future research.

JEL classification numbers: C78, D40, D44, L10, L14, O33.

* Department of Economics, Stanford University, Stanford CA 94305; email: jdlevin@stanford.edu. I thank Yeon-Koo Che, Liran Einav, Richard Levin and Andy Skrzypacz for helpful discussions, Andrey Fradkin for suggestions on the expanding literature, and Marissa Beck for excellent research assistance. I acknowledge receiving research support or compensation from the National Science Foundation, the Toulouse Network for Information Technology, the Sloan Foundation, Yahoo!, and eBay Research.

1. Introduction

The last fifteen years has seen the striking emergence of new internet platforms for search, e-commerce, online media, job matching, social networking and other activities. The growth of these platforms has been dramatic. Amazon, which opened in 1995, has annual revenue of over forty billion dollars. Google, which started in 1998, processes over a billion search queries each day. Facebook, founded in 2004, has attracted over eight hundred million users. Groupon exceeded a billion dollars in revenue in just its third year of operation.[1]

These and other internet platforms take advantage of how the internet has lowered a range of economic costs: the cost of creating and distributing certain products and services, the cost of acquiring and providing information, the cost of collecting and using data on consumer preferences and behavior. These changes have helped to make internet platforms dynamic and innovative, and are starting to inspire a great deal of economic research.

Several aspects of internet platforms are distinctive relative to traditional industries. One is the rapid ability to initiate and scale up operations. Internet firms, and especially consumer-oriented internet firms, often have very low start-up costs, and even lower costs of serving additional users because the underlying engineering is scalable. Facebook, for example, grew to over 500 million users with less than 500 engineers, or one engineer for every million users.[2] The potential speed and low cost of expansion are particularly relevant in thinking about the market structure of internet industries, and about how platforms can design market mechanisms that operate efficiently and take advantage of large scale.

[1] The Amazon and Facebook numbers are from company annual reports and websites. The Groupon figure is from
[2] "Scaling Facebook to 500 Million Users and Beyond," blog post by Robert Johnson (July 21, 2010) at http://www.facebook.com/note.php?note_id=409881258919. The social media platform Twitter similarly grew to over 200 million users with just 350 employees (http://blog.twitter.com/2010/12/stocking-stuffer.html).

A second distinctive feature of internet markets is customization. In traditional markets, it tends to be costly to personalize individual experiences. In online markets, it is common to have products recommended on the basis of past purchases, to obtain search results tailored to individual queries, or to see advertisements that reflect past browsing behavior.[3] This type of customization can facilitate a more efficient matching of people and opportunities. The matching perspective is a useful one because many internet platforms have been highly successful in exactly this dimension: creating markets for targeted advertising placements, systems that aggregate user feedback to create targeted advice, and applications and services tailored to people's social networks.

Finally, a third characteristic of internet platforms is that they're associated with rapid innovation. By innovation, I have in mind the creation of new products and ideas, but also the gradual refinement of search algorithms, information displayed to users, product characteristics, and payment and pricing mechanisms. Many internet platforms use controlled experiments to guide this process of continuous improvement. Economists are often frustrated by how little systematic experimentation occurs in traditional markets. But experiments can be costly and results obtained too slowly to generate substantial benefits. The calculus changes when a few lines of code result in different users seeing different displays or facing different prices, and the results are measured instantaneously.

With these features as background, my goal in this paper is to describe an expanding but somewhat disparate body of economic research inspired by the growth of new internet platforms. I focus on three areas. The first is research in industrial organization on how platforms compete for users. This work emphasizes the importance of direct and indirect network effects, and their

[3] Heski Bar-Isaac has pointed out that another form of customization in internet markets, for instance on platforms such as Wikipedia or YouTube, is that users can provide customized input into the production of a product.

implications for platform strategy and market structure. I describe the modeling framework and the insights it generates for internet platforms in Section 2.

In Section 3, I discuss research on some of the innovative market mechanisms created by internet platforms. These mechanisms include new types of auctions for advertising and other goods, novel reputation and recommender systems based on user data and feedback, and structured search processes for retail shopping, job matching and other goods and services. A main theme here is how technology allows for new market designs that take advantage of the scale and heterogeneity of many internet markets.

In Section 4, I turn to the more empirically-oriented research on competition and consumer behavior in online markets. A major focus of this work is how falling costs --- of consumer search, of product proliferation, of using dynamic pricing mechanisms --- have impacted specific markets and industries. I focus mainly on e-commerce marketplaces and give examples of how researchers have used the relatively structured environment of internet markets to test theories about consumer decision-making or imperfect competition.

The final section concludes with some speculation about future research. Because the general area is relatively young, there are still many open questions, and opportunities for advances in both theory and empirical methods.

2. Platform Strategy and Competition

Over the last decade, the theory of platform competition has been one of the most active areas of research in industrial organization (Caillaud and Jullien, 2003; Rochet and Tirole, 2003, 2006; Armstrong, 2006; Rysman, 2009; Weyl, 2010). Although much of this work is not internet-specific --- other motivating examples include operating systems, payment cards, shopping malls, magazines and a host of other industries --- it is a good point of entry for thinking about the strategic problems facing internet platforms and how they compete for users.

Economists think of platforms as intermediaries that bring together different types of users to enable economic or social interaction. In the case of e-commerce or online apartment rentals, the groups might be buyers and sellers. In online media, the groups might include consumers, producers of content, and advertisers. A key point of emphasis is that assembling users --- buyers and sellers, consumers and advertisers, groups of friends --- involves network effects. The value users assign to the platform depends on who else is using it.

Recent models build on an older network effects literature pioneered by Katz and Shapiro (1985) and Farrell and Saloner (1985). A new step is to explore the strategic and competitive implications of network effects when platforms assemble different types of users, and there are externalities across user groups (e.g. consumers care about the number of applications available on a platform and not just the number of other users). These cross-group effects sometimes are referred to as indirect, as compared to direct, network effects.

A. The Price Theory of Platforms

Many early models of network effects focused on coordination problems that can arise in attracting users. Because there can be equilibria with only a few active users, and with many active users, platforms can face a "chicken-and-egg" problem of getting initial users to sign up. Recent analyses (e.g. Weyl, 2010) have shifted the focus by considering a platform's choice of what users to attract under the constraint that the choice is consistent with users' equilibrium decisions, in effect assuming that platforms can overcome coordination problems. I follow this approach, which leads to a useful price-theoretic modeling framework.

Consider a platform that sets prices for K different user groups, with individual users then deciding whether to participate. Suppose that the willingness to pay of an individual in group k is $u_k(\mathbf{x},\zeta)$, where x_k denotes the number of group k users who participate, $\mathbf{x}=(x_1,\ldots x_K)$ denotes the vector of participant quantities, and ζ are the individual's characteristics. The payoff from not participating is zero. If the platform charges a fee p_k to group k users, the individual will participate if and only if $u_k(\mathbf{x},\zeta)\geq p_k$. The platform also incurs costs from serving its users: if participation is $\mathbf{x}$, the platform's costs are $C(\mathbf{x})$.

To see how coordination problems might arise, suppose there is a unit mass of homogeneous consumers who each obtain utility $u_k(x)=x$ from participating, where x is the mass of consumers who decide to use the platform. If the platform sets a price p=0.2, there will be an equilibrium in which no users participate (x=0), an equilibrium in which all users participate (x=1) and an (unstable) equilibrium in which exactly 20 percent of users participate.

Now suppose that within each user group, consumer values are heterogeneous and continuously distributed. While the equilibrium participation level may not be uniquely defined given platform prices, it is the case that for any participation level $\mathbf{x}$ that is consistent with a participation equilibrium, the fee the platform can extract from each user group is pinned down uniquely. For group k, it corresponds to the utility of the x_kth most enthusiastic user taking overall participation $\mathbf{x}$ as fixed.

Extending the example above, suppose that user utility is $u_k(x,\zeta)=x-\zeta$, where ζ is uniformly distributed between -1 and 1. To have a participation equilibrium at which x=0.5, it must be the case that users with $\zeta < 0$ participate and those with $\zeta > 0$ do not. So the unique price consistent with equilibrium participation of x=0.5 is p=0.5, which makes users with ζ=0 exactly indifferent.

This observation suggests it will be more convenient to think of the platform's problem as one of choosing quantities, and then to back out the prices implied by these quantities, than to work directly with the platform's choice of prices (Weyl, 2010). In effect, this means making the implicit assumption that the platform will be able to obtain its preferred participation outcome, should there be multiple participation equilibria corresponding to the prices it sets.[4]

To follow this route, let $P_k(x_k;\boldsymbol{x})$ denote the price that exactly x_k users from group k would be willing to pay, assuming they expect an overall number of users $\boldsymbol{x}$. That is, fixing expectations that participation will be $\boldsymbol{x}$, $P_k(\cdot;\boldsymbol{x})$ is the inverse demand curve for group k. Naturally $P_k(x_k;\boldsymbol{x})$ will be decreasing in x_k, but it may be increasing or decreasing in $\boldsymbol{x}$ depending on the sign of the network effects. If there are two user groups, buyers and sellers, we would expect positive cross-group externalities, so $\partial P_k(x_k;\boldsymbol{x})/\partial x_j > 0$ if $k \neq j$, but negative within-group externalities because sellers compete with other sellers and buyers with other buyers.

The profit maximizing participation solves

$$\max_x \sum_k x_k \cdot P_k(x_k;x) - C(x)$$

The first order conditions for this problem equate the marginal revenue for each user group k with the marginal cost $MC_k(\boldsymbol{x})=\partial C(\boldsymbol{x})/\partial x_k$. That is, for each user group k,

$$P_k(x_k;x) + \sum_j x_j \cdot \frac{dP_j(x_k;x)}{dx_k} = MC_k(x)$$

There is a convenient way to express these first order conditions that distinguishes the usual trade-off in maximizing profit against a downward-sloping demand curve from considerations involving network effects. Holding expected participation fixed at $\hat{x}$, the marginal revenue for group k is

[4] One justification is that if the platform can make its prices contingent on realized participation, it can make $\boldsymbol{x}$ a unique participation equilibrium. One way to do this, if the platform wishes to implement $\boldsymbol{x}=(x_1,\ldots,x_K)$, is to charge group k users $P_k(x_k,\boldsymbol{x}^*)$, where $\boldsymbol{x}^*$ is the realized participation (Weyl, 2010, and Weyl and White, 2011).

$$MR_k(x_k;\hat{x}) = \frac{\partial}{\partial x_k} x_k P_k(x_k;\hat{x}) = P_k(x_k;\hat{x}) + x_k \frac{\partial}{\partial x_k} P_k(x_k;\hat{x})$$

where $\partial P_k(x_k;\hat{x})/\partial x_k$ denotes the derivative of $P_k(x_k;\hat{x})$ with respect to its first argument.

The platform's first order conditions can then be re-arranged to yield:

$$MR_k(x_k;x) + \sum_j x_j \frac{\partial P_j}{\partial x_k}(x_j;x) = MC_k(\mathrm{x})$$

where (with some abuse of notation) $\partial P_j(x_j;x)/\partial x_k$ denotes the derivative of $P_j(x_j;x)$ with respect to its first $k+1^{\text{st}}$ argument.

In the re-arranged first order condition, the first term is the standard marginal revenue that can be obtained by allowing an additional group *k* user (fixing expected participation at ***x***), and the second term captures the extent to which the additional user will increase or decrease the willingness to pay of other users. The combined terms are equated against the marginal cost of serving the extra user.

This simple pricing model captures some intuitive ideas about platform strategy. For instance, consider a group of users that create value for others, so that the second term is positive. Relative to a monopolist that ignores network effects, a profit-maximizing platform wants to pulls in more of these users, which it can do by lowering their price. Similarly a platform should penalize users who create negative externalities by raising their price relative to the monopoly price ignoring network effects.

Examples of such behavior are plentiful. Search engines offer free email and other services to consumers in the hope of making more money from advertisers. Financial exchanges charge low fees or even make payments to traders who submit limit orders, as these orders create liquidity in the order book; they charge higher fees to traders who submit crossing orders that remove liquidity.

The pricing model also highlights the tension between extracting surplus from core users, and lowering prices to expand the user base. The tension is sharpest if we imagine that some subset of users are "locked in" or committed to using the platform. A standard observation in models with switching costs is that firms may be tempted to raise fees for existing users, even if this drives away some marginal consumers. A new element in the platform context is that by reducing its user base a platform may degrade its quality for the remaining users. The same type of tension arises when a platform is tempted to exploit its user base by expanding the number of intrusive ads or lowering service quality directly.

The analysis described above is for a profit-maximizing platform. It is also interesting to consider the conditions for an efficient allocation. Efficiency dictates that for all *k*,

$$P_k(x_k;x) + \sum_j \int_0^{x_j} \frac{\partial P_j}{\partial x_k}(z_j;x)dz_j = MC_k(x).$$

That is, group *k* users should be charged their marginal cost minus the externality they impose on others (which could be positive or negative). Relative to the efficient allocation, the profit-maximizing solution involves two distortions (Weyl, 2010). A profit-maximizing platform focuses on the marginal revenue from additional users rather than their marginal surplus, and there is a "Spence distortion": the platform cares about the externality additional users impose on marginal participants, rather than average participants. [5]

Note that one feature of this model is that all of the details of what the platform does, and how users interact, are incorporated in the payoff functions u_k and the cost function *C*. Platform fees have no direct effect on user interactions. For instance, if there are distinct groups of buyers and sellers, imposing a fee on sellers only affects buyers if it causes some sellers to exit the

[5] The importance of the Spence distortion in this model raises the question of whether platforms really do focus too much attention on marginal users. Weyl (2010) gives some supporting examples. In many cases, however, dedicated users of a platform seem able to exert considerable influence, in ways that may not involve threatening to exit. In the terms used by Hirschman (1970), "voice" can be more powerful than "exit".

market. So we are abstracting from issues of market interactions and organization that will come up in the next section, and focusing squarely on network effects in user participation.

The model also does not provide much insight into the specific structure of platform fees. In e-commerce, for example, the procurement platform Alibaba charges sellers a flat annual fee. Platforms such as eBay Motors charge a fee for each individual listing. Other platforms charge sellers for clicks (common on price comparison sites), or a commission on sales (Amazon, Etsy and others). These fee structures have very different implications for seller incentives and the quality of listings, but call for a model in which users make decisions about the nature and intensity of their activity rather than a binary participation decision.[6]

B. Platform Competition and Market Structure

The network effects model is particularly useful for thinking about competition *between* platforms. Here a central question is whether certain types of activity will become concentrated on a single platform – whether a platform industry will "tip". Internet search and consumer auctions are two common examples of industries that tipped in the early days of the internet. In the U.S., Google emerged from a host of competing search engines to capture around two-thirds of the market. And eBay quickly captured the consumer auction market despite the presence of many competing platforms. These markets also have dominant firms in China in the form of Baidu and Taobao, and in many other countries.

[6] The model above also imposes the assumption that every member of group j matters the same to group k users. On a job search website, for instance, it is inevitable that some job seekers will be more attractive to employers. To the extent that having "attractive" applicants is important to employers, platform pricing per se may be less important than explicit screening decisions by the platform, or an application process that lets firms identify high-quality matches. Veiga and Weyl (2011) have generalized the model described in the text to allow for this type of sorting and screening. Halaburda and Piskorski (2011) also develop a model in which search platforms differentiate themselves to attract different types of consumers by being more or less restrictive in their offerings.

A standard concern is that if a platform becomes dominant, there may be dynamic inefficiencies because users are coordinated and locked in to a single platform. It may be difficult for an innovative new platform to gain market share, even if its underlying attributes and technology are better. This concern has helped to motivate antitrust actions in industries such as operating systems and payment cards. It therefore raises the question of whether internet platforms should be viewed as similar or in some way different, and whether analogous competitive or antitrust concerns might arise.

Models of the type considered above suggest a number of factors to look at: the degree to which consumers view competing platforms as substitutable, the strength of positive network effects, and the extent to which production is characterized by economies of scale (e.g. Farrell and Klemperer, 2007). The importance of these factors can vary. For instance, the value offered by social networks depends heavily on the number and identity of other users, and employers are likely to favor job matching sites that can offer them a large number of qualified applicants, but it is less obvious why consumers would favor a search engine solely because it has larger market share. On the other hand, a search engine operating at a larger scale may have better data to improve its algorithms, and hence may enjoy a competitive advantage.

In traditional industries with network effects, high switching costs are often an important compounding factor. Consider the case of operating systems, where switching costs can be relatively high for individual users and for firms with large computer installations. Switching between internet platforms or using multiple platforms can be considerably easier. One can shop on Amazon and eBay, or be a Facebook user and try Twitter. At least in some cases, the combination of low switching costs and low costs to creating new platforms might mitigate traditional concerns about lock-in and dynamic inefficiency.[7]

A striking example where internet technology lowered entry barriers and undermined a dominant platform is in the market for publicly traded equities in the United States. As recently as 2005, around eighty percent of trades in New York Stock Exchange listed equities took place on the exchange floor. The next several years saw the rapid entry of new electronic exchanges, including the introduction of an electronic trading platform by the NYSE itself. By 2009, trading had fragmented to the point where the combined share of the NYSE floor and electronic exchange were 25%, and no single marketplace executed more than 20% of the overall public equities trades.[8] Of course, the fragmentation may be a temporary dislocation, so it will be interesting to see if trading eventually reverts to being highly concentrated.

C. Empirical Evidence on Platforms and Platform Competition

One limitation of the platform competition model I have described --- for analyzing pricing and especially for questions about competition and market tipping --- is that its predictive content depends largely on empirical quantities. In contrast to say, retail products, where there is a fair amount of evidence on demand elasticities or brand preferences, there is not much systematic evidence on preferences for single-homing, or the size and strength of network effects, or even price elasticities, all of which are key parameters in network effects models of platform competition (Rysman, 2009).[9]

---

[7] Indeed, even in industries such as social networking, where one might expect positive feedback effects to generate agglomeration, it is possible to point to examples of successful entry (Twitter) or rapid decline (MySpace).
[8] Remarkably over this same 2005 to 2009 period, the number of executed trades on US exchanges increased more than eight fold, the average order size fell by 60%, and the average speed of execution fell from ten seconds to less than one second. (SEC, Concept Release on Market Structure, 17 CFR Part 242).
[9] Cantillon and Yin (2011) propose a research agenda for studying competition between platforms, focusing on financial exchanges. There is also a large and related literature outside of economics studying the dynamics of user

In a way this is a little surprising because the internet offers many opportunities for simple and clean empirical studies. A good illustration is Brown and Morgan (2009), who look at competition between eBay and Yahoo! auctions. For their study, they auctioned identical coins on the two platforms. They conducted the auctions in 2004, at which time eBay had about 80% of the consumer auction market. The results were striking: the eBay sales attracted 50% more bidders and the resulting eBay prices were a third higher. The conclusion was that the market was in the process of tipping all the way toward eBay.

Several features of this study are striking. One is that the evidence leads to new questions: Why didn't buyers arbitrage the price differences? Did eBay subsequently exploit its market power and charge sellers higher fees? Why were some sellers still using Yahoo! given the large price differentials?[10] Another is how easy the study was to run -- it essentially amounted to running a few dozen auctions on two internet sites. It shows how internet markets lower some of the typical barriers to collecting data. In light of the effort devoted to advancing the theory, and its dependence on empirical quantities, one hopes for more empirical evidence in the future.

3. Designing Novel Market Mechanisms

Structuring economic and social activity online creates new opportunities for matching consumers with products and services, and providing them with a broader array of information. It also poses some challenges due to the long-distance and sometimes anonymous nature of interaction. The response has been a burst of innovative market design that takes advantage of

---

growth and the evolution of behavior on internet platforms. For instance, Suh, Convertino, Chi and Pirolli (2009) provide a fascinating analysis of how interactions between Wikipedia editors have evolved over time as the platform matured, and Viswanath et al. (2009) describe the evolution of user behavior on Facebook.

[10] Brown and Morgan argue that although the experiment results are inconsistent with a stable equilibrium configuration, they can be explained in the context of a dynamic model in which users adapt by imitating successful peers and the market slowly tips to one platform.

the scale, scope and continuous nature of activity permitted by new technology. These developments include new markets for advertising and financial trading, consumer marketplaces for services and used goods; and mechanisms to deliver online reviews and recommendations, as well as other consumer information.

In this section, I illustrate these points by discussing three examples of market design innovation: auctions for “sponsored search” advertising, reputation systems to ensure safe trade in e-commerce marketplaces, and recommendation and review systems. Each of these examples has attracted a good deal of research attention.

A. Sponsored Search Advertising

The basic problem in internet search is to match user requests for information with relevant search results. A good solution to this problem must be scalable --- Google, for instance, gets hundreds of billions of queries a year, and tens or hundreds of millions of distinct queries on any given day. Search engines use two different methods. They display "organic" results based on data gleaned from internet pages and user behavior, and "sponsored" results that are allocated using a market-based approach.

The market for sponsored search results operates very differently from traditional advertising markets in which firms might buy the right to have their advertisement shown over a period of time to a broad cross-section of consumers. Instead, when a user enters a query the search engine runs an auction to allocate space on the results page to advertisers that have placed bids for the relevant search terms.

One rationale for using an auction is that it elicits information from advertisers about their value for having their advertisement shown. This helps determine the efficient placement of ads. The auction also determines the price the search engine can charge. Both the placements and the prices can fluctuate from day to day or even from minute to minute. While each individual auction has small stakes, the number of queries is so large that the annual revenue of search engines is already in the tens of billions of dollars.

What I want to describe in this section is how this type of market can be modeled, how the auction run by Google and the other search engines evolved to its present form, and why it might be a fairly efficient solution to the matching problem in internet search settings. Two papers by Edelman, Ostrovsky and Schwarz (2007) and Varian (2007) have done this elegantly, and I follow their formulation of the problem.

Suppose there are M advertising positions. The top position will receive the most user attention and clicks, then the next position, and so forth. There are also N advertisers, with different values from having their ads clicked on. Specifically, assume that position m will receive x_m clicks, where $x_1 > ... > x_M$, and that advertiser n has a per-click value v_n, where $v_1 \geq ... \geq v_N$. If advertiser n buys position m, its payoff is $v_n x_m - t$, its value per click multiplied by the number of clicks it receives minus the payment t it makes to the platform.

This setting is a special case of the classical assignment model (Shapley and Shubik, 1972). The efficient assignment is assortative. The high-value advertiser should get the top position and the most clicks, the second-highest value advertiser the next position, and so on. In addition, the efficient assignment can be supported with competitive market-clearing prices. Suppose we set prices $t_1,...,t_M$ for the M positions. These prices will clear the market (and support an efficient allocation) if and only if for each advertiser n, and each position m,

$$v_n x_n - t_n \geq v_n x_m - t_m$$

This condition says that advertiser n prefers to buy its efficient position n at a price of t_n than buy any other position $m \neq n$.[11] In fact, only a few of these conditions are relevant because if advertiser n prefers position n to positions $n-1$ and $n+1$, it will also prefer n to the other positions.

Because of the discrete nature of the goods, there typically are a range of market-clearing prices. However, there is always a single competitive price vector that is (component-wise) maximal, and also one that is minimal. At the minimal market clearing prices, each advertiser is just indifferent between his own position and the position that is one better, $v_n x_n - t_n = v_n x_{n-1} - t_{n-1}$, or re-writing:

$$t_{n-1} = t_n + v_n \left(x_{n-1} - x_n \right)$$

So we can solve for the minimal market clearing prices as follows:

$$t_n = \sum_{\mathrm{k \geq n+1}} v_k \left(x_{k-1} - x_k \right)$$

These prices are increasing on per-click basis as one moves from worse to better positions. The reason is that advertiser n must be willing to pay for the extra $x_n - x_{n+1}$ clicks obtained in position n relative to position $n+1$, but not for the extra clicks obtained in position $n-1$. So the incremental price for clicks is increasing as one moves to better positions.[12]

To generate prices in practice, the leading search engines use a "generalized second price" (GSP) auction. Each advertiser submits a maximum amount it would be willing to pay for a user to click on its ad, the positions are assigned in order of the bids, and the price for each

---

[11] I am being a bit casual about the possibility that some positions are left unfilled or advertisers left unmatched --- i.e. about the case where $N \neq M$. To nail things down, we can assume that $N>M$ (without loss of generality in that we always can "create" some advertisers with zero value), and as a notational convention set $x_m = t_m = 0$ for $m>M$.

[12] For instance, suppose there are two positions that receive 200 and 100 clicks, and three advertisers with per-click values 3,2, and 1. The price of the lower position must be at least 100 to deter the lowest value bidder, but no more than 200. Similarly, the incremental price of the top position must be between 200 and 300 to ensure the market clears. So the lowest equilibrium prices are (300, 100) and the highest are (500, 200), and the respective "per click" prices are (3/2,1) and (5/2,2).

advertiser is set equal to the minimum bid it could have made and sustained its position, i.e. the next highest bid. An advertiser pays its price if and only if the user clicks on its ad.

The GSP auction is interesting because it involves a series of innovations that came about partly for practical reasons, but that turn out to have attractive theoretical properties.

A first innovation is that advertising is priced on a "per click" basis. In traditional media, advertising is often sold by the “eyeball” or per impression. Pricing clicks takes advantage of the fact that on the internet, it is easier to measure and track a user's response to advertising. But why is per-click pricing useful? One reason is that positions on a web page are worth different amounts depending on their prominence. So asking advertisers to bid for positions on the page would make for a complicated auction. Provided that clicks have roughly the same value regardless of the position from which the click was received, an auction with "per click" bidding simplifies the market. It reduces the dimension of each advertiser's bid from the number of positions to a single number.[13]

The second price format is another innovation. The initial search advertising auctions called for bidders to pay their own bid rather than the bid of the next ranked bidder. Unfortunately, if bidders have complete information about each other’s values, a pay-your-bid auction for positions has no Nash equilibrium in pure strategies. Each bidder would like to reduce its bid to just above the bid of the next-ranked bidder, but if bids are tightly clustered, low-ranked bidders will want to pay a little more to move up. In a one-shot auction this might not be a terrible problem. Indeed, looking at a closely related model, Bulow and Levin (2006) show that there is a mixed strategy equilibrium that can be close to efficient. In practice, however, the early search auctions had unstable dynamics. Edelman and Ostrovsky (2007)

[13] Milgrom (2010) explains some of the theoretical benefits of this type of simplification. One of his insights is that it can eliminate "bad" equilibria. For example, if a search engine ran a separate second price auction for each position, there would be equilibria where no advertiser bothered to submit a losing bid, leading to zero revenue.

document how bidders would gradually outbid each other for the top position until dropping to lower bids and re-starting the cycle, creating a lot of needless activity in the market.[14]

A third and more recent innovation is for search engines to estimate the quality of each ad and use the resulting quality scores to weight the bids in the action. In the above model, this has the effect of allowing high quality (or more "relevant") advertisers to obtain higher positions and pay less for their positions than they otherwise would (Edelman, Ostrovsky, Schwarz, 2007). This sort of differential pricing makes sense if one of the goals of a search engine is to show users higher quality ads and improve their overall experience.

Edelman et al. and Varian show that an attractive feature of the GSP auction is that it potentially has a stable and efficient outcome with competitive prices. Although bidders do not have a dominant strategy (depending on the behavior of other bidders, an advertiser might want to bid more or less than its actual value per click), the auction does have pure strategy Nash equilibria when bidders have complete information about each others' values, and these equilibria match up the competitive outcomes described above.

To show this, Edelman et al. and Varian focus on a particular subset of "locally envy free" Nash equilibria. These equilibria have the property that no bidder wants to "trade positions" with the bidder just above or below (i.e. assume the competitor's position and pay the price it is currently paying).[15] As a result, the payments coincide with the payments in some competitive equilibrium, and the assignment of bidders to positions is efficient.[16] Moreover, for any set of

[14] One anecdotal story for why search engines switched to the second price format is that the frequent bid changes created a burden for their servers. Matt Jackson also has pointed out a separate problem. Over time, bidders in a repeated pay-your-bid auction may recognize that there is little to gain by bidding up the price only to cycle down, leading to prices settling at a low (or seemingly "collusive") level.

[15] Edelman et al. provide several motivations for why this refinement makes sense, in particular arguing that dynamic aspects of bidding are likely to push toward locally envy free outcomes.

[16] To see this, suppose that $b^{(1)} \geq ... \geq b^{(n)}$ is a bid vector and $v^{(k)}$ is the valuation of the kth highest bidder. The bids constitute a Nash equilibrium of the GSP auction if for all k, $(v^{(k)}-b^{(k+1)})x_k \geq (v^{(k)}-b^{(k+2)})x_{k+1}$ and $(v^{(k)}-b^{(k+1)})x_k \geq (v^{(k)}-b^{(k-1)})x_{k-1}$. The locally envy free refinement strengthens the second condition to require $(v^{(k)}-b^{(k+1)})x_k \geq (v^{(k)}-b^{(k)})x_{k-1}$.

competitive equilibrium prices of the underlying assignment model, there is a locally envy free equilibrium of the GSP auction that has precisely these payments and again, an efficient assignment of advertisers to positions.

The search engines could have pursued alternative auction designs such as a Vickrey auction. A Vickrey auction works the same as a GSP auction except that instead of paying the next highest bid, each bidder pays an amount equal to the externality it imposes on lower-ranked bidders by displacing them one position. Specifically, if we order the bids such that $b^{(1)} \geq \ldots \geq b^{(n)}$, the payment for position n is $t_n = \sum_{k \geq n+1} b^{(k)} \left(x_{k-1} - x_k \right)$. This payment rule makes it a dominant strategy to bid one's true valuation, and hence the equilibrium payments in the Vickrey auction correspond exactly to the lowest market clearing position prices. Vickrey pricing is reportedly used by Facebook in their advertising auctions.

Recent work on sponsored search auctions also considers a number of extensions that go beyond the basic analysis I have described. Chen and He (2011) and Athey and Ellison (2011) incorporate consumer choices about which ads to click (see also Gomes, 2011, and Jeziorski and Segal, 2009, for some evidence). Ostrovsky and Schwarz (2009) study the use of optimal reserve prices using a large field experiment. And several papers, including Borgers, Cox, Pesendorfer and Petricek (2008), Varian (2009) and Athey and Nekipelov (2011), combine versions of the model described above with bidding data to estimate bidder valuations and the split of surplus between bidders and search engines.

B. Reputation Systems

But this means that the conditions for a refined equilibrium are identical to the conditions for a competitive equilibrium with per-click prices $p_k = b^{(k+1)}$ or position prices $t_k = b^{(k+1)} x_k$

My second example of novel market design comes from internet commerce. While selling online reduces certain frictions such as search costs, moving toward long distance and more anonymous trade can create new problems. In particular, it is hard to examine goods and it may be difficult for buyers to know whether they can trust the person with whom they are trading. A salient illustration of this comes from Jin and Kato (2007). They purchased baseball cards online and off-line and then had them graded by a professional service. They found substantially more misrepresentation in the online transactions.

One response to this type of problem has been to design mechanisms that elicit and aggregate user feedback to help other users assess the quality of sellers or products. A prototypical example is eBay's reputation system. On eBay, buyers and sellers can submit feedback after each transaction. In the baseline implementation of the system, users could offer a positive, negative or neutral rating, plus a short message. Many studies, starting with Resnick et al. (2001), have argued that the feedback system was crucial to eBay's success, and a substantial empirical literature, surveyed by Dellarocas (2006), has found that sellers with higher feedback scores enjoy some modest benefits in terms of higher prices and sales rates.[17]

More recent work by Bolton, Greiner and Ockenfels (2009) also highlights some of the pitfalls in developing effective reputation systems. A natural concern is that because providing feedback is in some sense a public good, users might not bother to do it. This may be a considerable problem in some markets, but it does not appear to be a problem on eBay, where about 70% of traders give feedback. Nevertheless, Bolton et al. identify a different problem, namely that the feedback system is not necessarily that informative. In fact, over 98% of feedbacks are positive --- the system appears to suffer from severe grade inflation.

[17] Cai, Jin, Liu, and Zhou (2011) compare seller penalties for cheating to buyer guarantees issued by the platform as alternative ways to promote trust. Using data from Eachnet (an eBay-type platform in China), they find that buyer protection policies led to more seller cheating, while penalties on sellers reduced cheating.

One potential reason why buyers might fail to submit negative feedback is that they fear retaliation. In eBay's baseline system, feedback was posted immediately, and sellers who received negative feedback had a strong tendency to respond by giving the buyer a reciprocal negative feedback. Bolton et al. report on a series of lab experiments in which eliminating sequential feedback and allowing for more fine-grained ratings dramatically improved the informativeness of user reports, and improved the efficiency of exchange. Their findings were evidently so convincing that eBay remodeled the feedback system to incorporate these suggestions. Bolton et al. provide some evidence that the change led to more informative feedback, but interestingly the field results do not appear to be as dramatic as the lab results. This suggests that perhaps the institution could benefit from further incremental innovations.

C. Reviews, Recommendations and Search Rankings

Sponsored search and to some extent reputation mechanisms can be seen as particular approaches to the general problem of matching consumers to products or services in settings where there are potentially a huge number of offerings and consumer information is very incomplete. This matching problem is fundamental in the design of many online marketplaces.

One important approach has been the design of systems to aggregate user reviews or purchasing behavior to provide recommendations. These systems have received considerable attention (although more so in computer science and marketing than in economics). Avery, Resnick and Zeckhauser (1999) is an early study that focuses on how to structure the timing of evaluations and payments to reviewers. The mechanisms used on internet platforms (e.g. Netflix, Amazon, Yelp) generally do not use payments, but users may enjoy an intrinsic pleasure or status benefit from posting informative evaluations. Ghosh and McAfee (2010) analyze how platforms create incentives for high-quality reviews by allocating consumer attention.[18] On the empirical side, Chevalier and Mayzlin 2006) and Luca (2011) provides evidence on how reviews influence consumer demand for books and restaurants, respectively.

Platforms such as Amazon and Netflix provide consumers with both reviews and algorithmic recommendations based on the past purchases of "similar" consumers. Here, "similar" is defined with respect to past purchasing or reviewing behavior (Koren and Bell, 2011). The availability of social network data provides an interesting alternative. To the extent that individuals befriend those with similar tastes, the purchases of someone's friends may be a good indication of what they would like. The relative effectiveness of social versus commercial data in designing mechanisms to provide consumer information is an interesting question, and one that is likely to be important for overall success of different platforms.[19]

Search rankings also provide guidance to consumers, The sponsored search rankings discussed above combine advertiser bids with information about click-through rates, but more generally, rankings can be based on information obtained from buyers, sellers, or the platform itself; a successful search might be defined as a sale rather than a click (allowing the platform to take a commission, rather than a per-click payment); and the platform may need to worry about fostering competition between sellers in addition to directing buyers.

While there is no "canonical" model that captures all these issues, some recent work captures particular aspects of them. One example is models of price competition in which consumers search across firms in order of price (Baye and Morgan, 2001) or prominence (Armstrong, Vickers and Zhou, 2009). Hagiu and Julien (2010) look at the consumer search

[18] A related paper by Ghosh, Kale and McAfee (2011) looks at the interesting question of how to create effective incentives by having platform users police the behavior of other users.
[19] Tucker (2011) provides evidence that socially targeted advertising is highly effective relative to advertising targeted on more standard characteristics.

problem from the platform's perspective, pointing out that if additional searches lead to more revenue, a platform may not want to guide consumer's immediately to the most attractive offering. Their paper relates to the general question of how platforms should try to structure ranking mechanisms, while taking into account that the ranking mechanism will affect both consumer search patterns and seller incentives.[20]

Finally, targeted advertising provides an additional way to match online consumers to products and services. McAfee (2011) describes some of the market design issues involved in creating real-time exchanges in which publishers can sell advertising opportunities. The markets function similarly to sponsored search, except that rather than bid for keywords, firms can bid to have their ads shown to specific consumers, or on specific web pages. The ability to customize advertising down to the individual user creates an interesting tension for market design. On the one hand, allowing advertisers to demand highly targeted advertising placements is potentially efficient. On the other hand, it can allow informed advertisers to "cherry-pick" the best opportunities, or can lead to a failure of competition if individual markets become too thin (Levin and Milgrom, 2010). Recent papers by Abraham, Athey, Babaioff, and Grubb (2011), Beck and Milgrom (2011), Bergemann and Bonatti (2011), and McAfee, Papineni, and Vassilvitskii (2010) have made progress on modeling these issues.

4. Evidence on Competition and Consumer Behavior

[20] The papers by Chen and He (2011) and Eliaz and Spiegler (2011) incorporate price competition between firms into a model of sponsored search rankings; see also the papers by Athey and Ellison (2011) and Gomes (2011) mentioned above. In terms of proving broader incentives for sellers and advertisers, one interesting question is the extent to which platforms should be transparent about their ranking methods. Some of the best known examples, such as Google's quality-weighted ad rankings and eBay's "best match" system for search, are relatively opaque, presumably to discourage sellers from trying to manipulate the rankings.

In this section, I turn from what are mainly theoretical analyses of platform competition and market mechanisms to discuss the empirical evidence on competition and consumer behavior in internet markets. Many of the relevant questions about consumer and firm behavior in this context are motivated by the technological shifts highlighted in the introduction. For instance, what are the effects of reducing search and distribution costs, of making it easier to provide customized information, or of enabling sophisticated pricing mechanisms? Internet markets also provide a natural environment to test general theories about consumer decision-making and imperfect competition because interactions take place in structured settings where behavior can be recorded in great detail.

A. Search Costs and Price Competition

One of the prominent hypotheses in the early days of internet commerce was that reductions in search costs would intensify competition and reduce price dispersion. This hypothesis has generated sustained empirical interest. In one early study, Brynjolfsson and Smith (2000) found that online prices for books and CDs were roughly ten percent lower than offline, but featured significant variation across retailers, up to 30% of the average price. Subsequently, Baye, Morgan and Scholten (2004) found similarly large posted price variation for consumer electronics in data obtained from a price search engine. A raft of further studies have reached generally the same conclusion: that online competition has lowered prices, but that price dispersion remains ubiquitous despite the seemingly low-cost nature of the search environment (Ellison and Ellison, 2005).[21]

[21] A number of studies use experimental or quasi-experimental research designs to estimate price elasticities in online markets and find them to be relatively high. Looking at standardized products in structured search environments, Ellison and Ellison (2009) and Ong and Zhong (2011) find very high price elasticities. Einav, Knoepfle, Levin and Sundaresan (2011) use tax variation to look at the price sensitivity of buyers who have clicked

One way to rationalize the persistence of price dispersion, even in structured search environments, is by arguing that although search costs have decreased, they remain non-trivial. Ellison and Ellison (2009) argue that low search costs have encouraged online retailers to create new frictions by using "obfuscation" strategies such as up-sells, add-ons and bait-and-switches (see also Ellison, 2005, and Spiegler, 2006). They provide support for this argument by combining price data from a comparison website and sales data from a competing electronics retailer. They show that consumer demand on the comparison site is remarkably price-elastic because consumers tend to click on the highest-ranked listings. But they also document a range of tactics that make the sales process less transparent, so that consumers often don't buy the product that prompted the original click.

Another line of research argues that price dispersion and other aspects of online pricing may reflect a lack of consumer sophistication that manifests as “excessive” costs of search. Lee and Malmendier (2010) study an episode on eBay in 2003 in which a particular board game was available both at a posted price and by auction. They find that in about a quarter of the auctions, the game sold for $10 or more above the posted price, and argue that this is hard to reconcile with standard models of costly search. The episode is intriguing, and interestingly the prevalence of this type of “over-bidding” seems to have declined over time. Einav, Kuchler, Levin and Sundaresan (2011) look at tens of thousands of recent cases in which an eBay seller offered the same good by auction and posted price. They find that auction prices were only rarely as much as $10 above the posted price, and more typically were 10-20% lower. Einav, Farronato, Levin and

on items on eBay, finding somewhat less price sensitivity (but the estimate is conditional on consumers being interested in a particular listing). A few studies also argue that estimates of online price dispersion are overstated. Ghose and Yao (2009) point out that studies of price comparison engines look at posted prices, rather than transaction prices. They provide evidence from the GAO's procurement marketplace, showing that transaction prices feature much lower dispersion.

Sundaresan (2012) provide evidence that links the results, showing that auction prices above an item's posted price were much more common eight years ago than they are today.

A more persistent pricing anomaly comes from studies looking at the extent to which consumers internalize add-on prices such as shipping fees. In a series of studies, Hossain and Morgan (2006), Brown, Hossain and Morgan (2009), and Tyan (2005) use field experiment and observational data to argue that eBay auction prices do not fully adjust to reflect differences in shipping fees. Einav, Kuchler, Levin and Sundaresan (2011) reach a similar conclusion based on cases where sellers auctioned identical items with different shipping fees. They find that for each $1 increase in shipping fee, auction prices adjust down by only about $0.40 to $0.80. These findings are consistent with the general observation that certain fees and prices can be less salient to consumers --- in this case, despite the setting being fairly simple and transparent.

B. The "Long-Tail" Theory of Demand

Another early and prominent hypothesis about internet commerce was that low search costs and low costs of product proliferation would shift consumer demand toward niche products --- the so-called "long tail" hypothesis (Anderson, 2006). In an influential paper, Brynjolfsson, Hu and Smith (2003) estimated that of the roughly two million book titles offered by Amazon in 2000, the top 100,000 (roughly the number of titles available at a Barnes and Noble superstore) accounted for only about 70% of total sales. Moreover, they estimated that the annual consumer surplus from sales of these niche books could be as high as $1 billion.

Their finding drew attention to the possibility that sales of niche products with low volumes could collectively be very important (see also Chevalier and Goolsbee, 2003). In a recent follow-up, Brynjolfsson, Hu and Smith (2010) argue that the importance of niche products has grown over time. Using 2008 data from Amazon, they estimate that books ranked above 100,000 now account for over 35% of sales and perhaps $5 billion in consumer surplus.[22]

While the proliferation of products online seems hard to debate, the "long tail" shift that Brynjolfsson et al. find is not a completely obvious implication. One reason is that by lowering the costs of dissemination and distribution for popular products, the internet allows products that attract some enthusiastic consumers or get favorable reviews to expand their market share rapidly. [23] Indeed, Elberse and Olberhozer-Gee (2008) provide some evidence from the video rental business that online markets have had the reverse effect of channeling demand toward the most popular products. This idea is reminiscent of Rosen's (1981) "superstar" theory, where falling distribution costs can lead to a more skewed distribution of sale. Indeed, Bar-Isaac, Caruana and Cunat (2012) develop a model in which falling search costs can give rise to both superstar and long tail effects.

Gentzkow and Shapiro (2011) provide another relevant piece of evidence by looking at the individual consumption of online media. They show that people who seek out "niche" media online (they define "niche" in terms of ideological slant rather than market share), also visit the most common mainstream media sites. Their findings suggest that to the extent that the internet has allowed niche or "tail" products to succeed, substitution toward these products may occur within as opposed to across consumers. Given the wealth of data available about online purchasing and browsing decisions, this debate about shifts in product market concentration seems ripe for further analysis.

---

[22] The new paper uses a different specification to estimate the shape of the sales distribution, and the authors suggest their model used in earlier paper might have overstated the importance of niche products.
[23] Leskovec, Adamic and Huberman (2007) use data from an online retailer to study how product recommendations can affect the distribution of consumer purchases, and the implications for “viral marketing”.

C. Auctions and Dynamic Sales Mechanisms

A third prominent hypothesis from the early days of online markets was that because of reduced transaction costs, online markets would shift away from traditional posted prices toward more flexible and dynamic sales mechanisms (e.g. Lucking-Reiley, 2000; Hall, 2002). An obvious example of this are the consumer auctions conducted on eBay and other platforms (Bajari and Hortacsu, 2003). In general, auctions tend to involve higher transaction costs than posted prices because the seller has to assemble competing buyers. In online markets the transaction cost is reduced because bids can be submitted remotely, and in some cases using computer "proxies". If the cost becomes manageable, auctions have the desirable property that they facilitate price discovery and encourage buyer competition for scarce goods.

On eBay, this general story appears to be about right. Looking across the platform, it appears that auctions are used to sell idiosyncratic items and by occasional sellers who might be less well informed about demand or more eager to sell (Einav, Farronato, Levin and Sundaresan, 2012). In the last decade, however, the fraction of posted price listings has increased dramatically, from less than 10% to well over 50%. One possible explanation is that while the internet has reduced the transaction cost of running auctions, it gradually has made it easier to find prices for comparable goods, so there is less need to use auctions for price discovery. Einav et al. also show that consumer demand for the posted price format has increased, perhaps because the initial novelty of online auctions has lessened or because auctions take time and there is increased competition for consumers' online attention.

Even if posted prices are generally efficient, there are certain types of goods for which auction markets can play a key role in allocating "last-minute" inventory. Online examples include plane and hotel reservations, and web page advertising. A nice example studied by

Sweeting (2010) is online ticket sales. Sweeting finds that as events approach ticket prices adjust downward and the market shifts toward auctions. He relates this to the predictions of dynamic pricing models, where the opportunity cost of selling falls as the "expiration date" approaches. Of course, in these models the equilibrium price path can depend on whether buyers anticipate price declines and time their purchases strategically. Board and Skrzypacz (2011) show that with forward-looking buyers, the optimal strategy for a perishable goods seller typically involves dynamic price adjustment followed by a last-minute auction if there is remaining inventory.

Online display advertising markets involve a similar mix of forward sales and last-minute auctions. Publishers sell contracts in advance that guarantee advertisers a certain audience or that there ad will be shown for a given time period, often using a negotiated (but fixed price) contract. Residual inventory is then auctioned in real time. These real-time markets offer the benefits of flexibility and allow advertisers to incorporate up-to-the-minute targeting information, but are not necessarily ideal for large purchases or advance planning, As of yet, however, there is relatively little empirical evidence on the efficiency of this arrangement or alternative forms of market organization.[24]

5. Directions for Future Research

Over the last fifteen years, there has been a remarkable shift of economic and social activity onto the internet. In this paper, I have described some of the economic research this shift

24 For instance, one can ask whether real-time auctions are really needed, or whether bidders should be allowed to continuously update their bids, when an alternative might be to allocate inventory hourly or daily. The situation parallels the sponsored search case discussed above, where one justification is that the number and type of searches is variable, so advertisers attempting to allocate a budget will want to optimize dynamically, and prices should similarly adjust. Advertisers may also be learning over time about their values or click rates. Of course, these stories suggest that bidding is unlikely to be stable over time, contrary to the static modeling approach taken in Section 3.

has inspired. Because this work is still at a fairly early stage, I want to conclude by highlighting a few directions for future progress that seem particularly promising.

At the outset, I argued that a defining feature of many internet markets is innovation, and particularly incremental innovation that occurs continuously, often accompanied and guided by systematic experimentation. But this aspect of internet markets hasn't been a major focus of the research to date. This is surprising for several reasons. First, to the extent that continuous improvement is a key feature of internet platforms, one wants to account for it in thinking about platform competition or internet market design. For instance, I mentioned earlier that one form of scale economy for internet platforms might be access to more and better data that allows for faster learning and broader experimentation. Alternatively, one might argue that well-designed market mechanisms are those that leave room for improvements as better data becomes available, just as the sponsored search auctions I described earlier incorporate current estimates of the relevance and quality of each ad.

The potential for experimentation in online markets is also useful as an empirical tool. There is already a significant amount of work, dating back at least to Lucking-Reiley (1999), that uses online field experiments to answer questions about pricing decisions, advertising responsiveness, and other aspects of internet markets (Lewis and Reiley, 2011, and Ostrovsky and Schwarz, 2009 are two recent and large-scale examples). Einav, Kuchler, Levin and Sundaresan (2011) also point out that researchers can make use of "experiments" conducted by market participants. They identify millions of cases where eBay sellers varied pricing or listing parameters for fixed items and use this targeted variation to study consumer behavior and market outcomes (see also Elfenbein, Fisman and McManus, 2012).

A related point pertains to empirical research more generally, or at least empirical research on competition and consumer behavior. Over the last twenty-five years or so, a major emphasis of empirical research on imperfectly competitive markets has been on developing and applying econometric methods that were designed at least in part to substitute or compensate for lack of data --- for example, methods for estimating consumer preferences from aggregate demand data, or methods or for using equilibrium assumptions to infer cost parameters or demand elasticities that are difficult to measure directly or experimentally (Einav and Levin, 2010).

In internet markets, the situation can be quite different – there is a tremendous ability to measure behavior at a individual level (perhaps down to the individual "click") and in many cases the opportunity to find or create useful variation in prices or other parameters to measure relevant behavioral elasticities. Instead the challenge seems to be to find meaningful ways to extract useful information from such rich data, or statistics that are both precise and revealing. In this sense, moving from problems of too little data and a focus on econometric methods designed as a substitute, to environments with enormous amounts of data, and many possibilities for experiments to learn about causal effects is going to require a real shift in mind-set, new thinking and new approaches. This is a promising topic for future research.

References

Abraham, Ittai, Susan Athey, Moshe Babaioff and Michael Grubb, "Peaches, Lemons and Cookies: Designing Auction Markets with Dispersed Information," MIT Sloan Working Paper, 2011.

Armstrong, Mark, "Competition in Two-Sided Markets," *Rand Journal of Economics*, 2006, 37(3): 668-691.

Armstrong, Mark, John Vickers and Jidong Zhou, "Prominence and Consumer Search," *Rand Journal of Economics*, 2009, 40(2): 209-233.

Anderson, Chris, *The Long Tail: Why the Future of Business is Selling Less of More*, New York: Hyperion, 2006.

Athey, Susan and Glenn Ellison, "Position Auctions with Consumer Search," *Quarterly Journal of Economics,* 2011, 126(3): 1213-1270.

Athey, Susan and Denis Nekipelov, "A Structural Model of Sponsored Search Auctions," UC Berkeley Working Paper, 2011.

Avery, Christopher, Paul Resnick and Richard Zeckhauser, "The Market for Evaluations," *American Economic Review*, 1999, 89(3): 564-584.

Bajari, Patrick and Ali Hortacsu, "Economic Insights from Internet Auctions," *Journal of Economic Literature*, 2004, 42(2): 457-486.

Bar-Isaac, Heski, Guillermo Caruana, and Vincente Cunat, "Search, Design and Market Structure," *American Economic Review*, forthcoming.

Baye, Michael and John Morgan, "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review*, 2001, 91(3): 454-474.

Baye, Michael, John Morgan and Patrick Scholten, "Price Dispersion in the Large and in the Small: Evidence from an Internet Price Comparison Site," *Journal of Industrial Economics*, 2004, 52(4): 463-96.

Beck, Marissa and Paul Milgrom, "Adverse Selection Free Auctions for Internet Display Advertising," Stanford University Working Paper, 2011.

Board, Simon and Andy Skrzypacz, "Revenue Management with Forward Looking Buyers," Stanford University Working Paper, 2010.

Bolton, Gary, Ben Greiner and Axel Ockenfels, "Engineering Trust: Strategic Behavior and the Production of Reputation Information," University of Indiana Working Paper, 2009.

Bonatti, Alessandro and Dirk Bergemann, "Targeting in Advertising Markets: Implications for Offline vs. Online Media," *Rand Journal of Economics*, 2011, 42:417-443.

Borgers, Tilman, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek, "Equilibrium Bids in Auctions of Sponsored Links: Theory and Evidence," University of Michigan, 2008.

Brown, Jennifer, Tanjim Hossain and John Morgan, "Shrouded Attributes and Information Suppression: Evidence from the Field," *Quarterly Journal of Economics*, 2010, 125(2): 859-876.

Brown, Jennifer and John Morgan, "How much is a Dollar Worth? Tipping versus Equilibrium Coexistence on Competing Online Auction Sites," *Journal of Political Economy*, 2009, 117(4): 668-700.

Brynjolfsson, Erik and Michael D. Smith, "Frictionless Commerce? A Comparison of Internet and Conventional Retailers," *Management Science*, 2000, 46(4), 563-585.

Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Michael D. Smith, "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science*, 2003, 49(11): 1580-1596.

Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Michael D. Smith, "The Longer Tail: The Changing Shape of Amazon's Sales Distribution Curve," MIT Working Paper, 2010.

Bulow, Jeremy and Jonathan Levin, "Matching and Price Competition," *American Economic Review*, 2006, 96(3): 652-668.

Cai, Hongbin, Ginger Jin, Chong Liu and Li-An Zhou, “How to Promote Trust: Theory and Evidence from China,” 2011.

Caillaud, Bernard and Bruno Jullien, `Chicken and Egg: Competition between Intermediation Service Providers," *Rand Journal of Economics*, 2003, 34: 309-328.

Cantillon, Estelle and Pai-Ling Yin, “Competition between Exchanges: A Research Agenda,” *International Journal of Industrial Organization*, 2011, 29(3): 329-336.

Chen, Yongmin and Chuan He, "Paid Placement: Advertising and Search on the Internet," *Economic Journal*, 2011, 121(556): F309-328.

Chevalier, Judith and Austan Goolsbee, "Measuring Prices and Price Competition Online: Amazon vs. Barnes and Noble," *Quantitative Marketing and Economics*, 2, June 2003.

Chevalier, Judith and Dina Mayzlin, “The Effect of Word of Mouth on Sales: Online Book Reviews,” *Journal of Marketing Research*, 2006, 43, 345-354.

Dellarocas, Chrysanthos, "Reputation Mechanisms," *Handbook of Economics and Information Systems*, Terence Hendershott, ed., Elsevier Publishing, 2006.

Edelman, Benjamin and Michael Ostrovsky, "Strategic Bidder Behavior in Sponsored Search Auctions," *Decision Support Systems*, 2007, 43(1): 192-198.

Edelman, Benjamin, Michael Ostrovsky and Michael Schwarz, "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords," American *Economic Review*, 2007, 97(1): 242-259.

Einav, Liran, Chiara Farronato, Jonathan Levin and Neel Sundaresan, "Sales Mechanisms in Online Markets: What Happened to Internet Auctions," in progress.

Einav, Liran, Dan Knoepfle, Jonathan Levin and Neel Sundaresan, "Sales Taxes and Internet Commerce," Stanford University Working Paper, 2011.

Einav, Liran, Theresa Kuchler, Jonathan Levin and Neel Sundaresan, "Learning from Seller Experiments in Online Markets," Stanford University Working Paper, 2011.

Einav, Liran and Jonathan Levin, "Empirical Industrial Organization: A Progress Report," *Journal of Economic Perspectives*, 2010, 24(2): 145-162.

Elberse, Anita and Felix Olberholzer-Gee, "Superstars and Underdogs: An Examination of the Long Tail Phenomenon in Video Sales," *Marketing Science Institute*, 2007, 4: 49-72.

Elfenbein, Daniel, Ray Fisman and Brian McManus, "Charity as a Substitute for Reputation: Evidence from an Online Marketplace," *Review of Economic Studies*, forthcoming.

Eliaz, Kfir and Ran Spiegler, "A Simple Model of Search Engine Pricing," *Economic Journal*, 2011, 21: F329-F339.

Ellison, Glenn, "A Model of Add-On Pricing," *Quarterly Journal of Economics*, 2005, 120(2): 585-637.

Ellison, Glenn and Sara Fisher Ellison, "Lessons from the Internet," *Journal of Economic Perspectives*, 2005, 19(2): 139-158.

Ellison, Glenn and Sara Fisher Ellison, "Search, Obfuscation and Price Elasticities on the Internet," *Econometrica*, 2009, 77(2): 427-452.

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

Farrell, Joseph and Paul Klemperer, "Coordination and Lock-In: Competition with Switching Costs and Network Effects," *Handbook of Industrial Organization*, Volume 3, ed. Mark Armstrong and Robert Porter, Excelsvier, 2007.

Farrell, Joseph and Garth Saloner, “Standardization, Compatibility and Innovation,” *Rand Journal of Economics,* 1985, 16(1): 70-83.

Gentzkow, Matthew and Jesse Shapiro, “Ideological Segregation Online and Offline,” *Quarterly Journal of Economics*, 2011, 126(4): 1799-1839.

Ghose, Anindya and Yuliang Yao, "Using Transaction Prices to Re-Examine Price Dispersion in Electronic Markets," *Information Systems Research*, forthcoming.

Ghosh, Arpita, Satyen Kale and Preston McAfee, “Who Moderates the Moderators? Crowdsourcing Abuse Detection in User-Generated Content,” *Proceedings of the 12 ACM Conference on Electronic Commerce*, 2011.

Ghosh, Arpita and Preston McAfee, “Incentivizing High Quality User Generated Content,” *WWW2010: Proceedings of the 19th International Conference on the World Wide Web,* 2011.

Gomes, Renato, "Optimal Auction Design in Two-Sided Markets," 2011.

Hagiu, Andrei and Bruno Julien, "Why Do Intermediaries Divert Search?" *Rand Journal of Economics*, 2012.

Halaburda, Hanna and Mikolaj Jan Piskorski, “Competing by Restricting Choice: The Case of Search Platforms,” Harvard Business School Working Paper, 2011.

Hall, Robert E., *Digital Dealing: How e-Markets are Transforming the Economy*, Norton, 2002.

Hirschman, Albert O. *Exit, Voice and Loyalty*: *Responses to Decline in Firms, Organizations and States*, Cambridge, MA: Harvard University Press, 1970.

Jeziorski, Przemyslaw and Ilya Segal, "What Makes them Click: Empirical Analysis of Consumer Demand for Internet Search Advertising," Stanford University Working Paper, 2009.

Jin, Ginger Zhe and Andrew Kato, "Dividing Online and Offline: A Case Study," *Review of Economic Studies*, 2007, 74(3): 981-1004.

Katz, Michael and Carl Shapiro, “Network Externalities, Competition and Compatibility,” *American Economic Review*, 1985, 75(3), 424-440.

Koren, Yehuda, and Robert Bell, “Advances in Collaborative Filtering,” *Recommender Systems Handbook*, 2011, Part 1, 145-186.

Lee, Young Han and Ulrike Malmendier, "The Bidder's Curse," *American Economic Review*, 2011, 101(2): 749-787.

Leskovec, Jure, Lada Adamic, and Bernardo Huberman, “The Dynamics of Viral Marketing,” *ACM Transactions on the Web*, 2007, 1(1).

Levin, Jonathan and Paul Milgrom, "Online Advertising: Heterogeneity and Conflation in Market Design," *American Economic Review*, 2010, 100(2): 603-607.

Lewis, Randall and David Reiley, "Does Retail Advertising Work? Measuring the Effects of Advertising on Sales via a Controlled Experiment on Yahoo!," Yahoo! Research, 2011.

Luca, Michael, “Reviews, Reputation and Revenue: The Case of Yelp.com,” HBS Working Paper 12-016, 2011.

Lucking-Reiley, David “Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet," *American Economic Review*, 1999, 89(5), 1063-1080.

Lucking-Reiley, David "Auctions on the Internet: What's Being Auctioned, and How?" *Journal of Industrial Economics*, 2000, 48(3), 227-252.

McAfee, Preston, “The Design of Advertising Exchanges,” *Review of Industrial Organization*, 2011, 39: 169-185.

McAfee, R. Preston, Kishore Papineni, and Sergei Vassilvitskii, “Maximally Representative Allocations for Guaranteed Delivery Advertising Campaigns,” 2010.

Milgrom, Paul, "Simplified Mechanisms, with an Application to Sponsored Search Auctions," *Games and Economic Behavior*, 2010, 70(1), 62-70.

Ong, David and Zemin Zhong, “Paying More to Top Up: An Online Field Experiment On Price Dispersion,” Peking University, 2011.

Ostrovsky, Michael and Michael Schwarz, "Reserve Prices in Internet Advertising Auctions: A Field Experiment," Stanford University, 2009.

Resnick, Paul, Richard Zeckhauser, E. Friedman and K. Kuwabara, "Reputation Systems," *Communications of the ACM*, 2000, 43(12): 45-48.

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

Rochet, Jean-Charles and Jean Tirole, "Two-Sided Markets: A Progress Report" *Rand Journal of Economics*, 2006, 37(3): 645-667.

Rosen, Sherwin, "The Economics of Superstars," *American Economic Review*, 1981, 71(5): 845-858.

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

Shapley, Lloyd and Martin Shubik, "The Assignment Game I: The Core," *International Journal of Game Theory*, 1972, 1(1): 111-130.

Spiegler, Ran, "Competition over Agents with Boundedly Rational Expectations," *Theoretical Economics*, 2006, 1: 207-231.

Sweeting, Andrew, "Price Dynamics in Perishable Goods Markets: The Case of Secondary Markets for Major League Baseball Tickets," Duke University Working Paper, 2010.

Suh, Bongwon, Gregorio Convertino, Ed Chi and Peter Pirolli, "The Singularity is Not Near: Slowing Growth of Wikipedia," *WikiSym,* 2009.

Tyan, Sean, "The Effect of Shipping Costs on Bidder Entry and Seller Revenues in eBay Auctions," Senior Thesis, Stanford University Department of Economics, 2005.

Tucker, Catherine, "Social Advertising," MIT Sloan Working Paper, 2011.

Varian, Hal, "Position Auctions," *International Journal of Industrial Organization*, 2007, 25(6): 1163-1178.

Varian, Hal, "Online Ad Auctions," *American Economic Review*, 2009, 99(2): 430-434.

Varian, Hal, "Computer-Mediated Transactions," *American Economic Review*, 2010, 100(2): 1-10.

Veiga, Andre and E. Glen Weyl, "Multi-Dimensional Heterogeneity and Platform Design," University of Chicago Working Paper, 2011.

Viswanath, Bimal, Alan Mislove, Meeyoung Cha and Krishna Gummadi, "On the Evolution of User Interaction in Facebook," *WOSN '09*, August 17, 2009.

Weyl, E. Glen, "A Price Theory of Multi-Sided Platforms," *American Economic Review*, 2010, 100(3): 1642-1672.

Weyl, E. Glen and Alex White, "Insulated Platform Competition," University of Chicago Working Paper, 2011.