SIEPR Discussion Paper No. 10-018

The Economics of Internet Markets

by
Jonathan Levin

Stanford Institute for Economic Policy Research
Stanford University
Stanford, CA 94305
(650) 725-1874

This Draft: February 18, 2011

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, retail and business-to-business e-commerce, internet job matching and financial exchanges, 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, and Marissa Beck for excellent research assistance. Disclosure: In the last five years, I have received 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, financial trading, social networking and other activities. The growth of these platforms has been dramatic. Amazon opened in 1995 and today has annual revenue of over thirty billion dollars. Google, which started in 1998, now processes over a billion search queries each day. China's Taobao has added three hundred and fifty million users since it began in 2003. Facebook, founded in 2004, has added over five hundred million users. In just the last two years, the discount platform Groupon grew from nothing to a reported annual revenue of almost two billion dollars.[1]

These and other internet platforms all take advantage of how the internet has lowered a range of economic costs: the cost of creating and distributing certain types of products and services, the cost of acquiring information about these goods, the cost of collecting and using data on consumer preferences and behavior. These changes have helped make internet platforms particularly dynamic and innovative, and inspired a great deal of economic research.

Before turning to this research, it is useful to highlight several aspects of internet platforms that strike me as particularly distinctive relative to traditional industries. One that is already suggested by the numbers above is scale, or scalability. Internet firms often have very low costs of serving additional users because the underlying code and 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 ability to expand rapidly at low cost is particularly relevant

[1] The Amazon, Taobao and Facebook numbers are from company annual reports. The Groupon estimate is from the Wall Street Journal's All Things Digital blog, "Groupon Annual Revenues Actually $2 Billion" post dated December 3, 2010. Ellison and Ellison (2005) point out that Amazon took just five years to reach a billion dollars in sales, while Walmart took nineteen.
[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. As another example, the social media platform Twitter has grown to over 200 million users with just 350 employees (http://blog.twitter.com/2010/12/stocking-stuffer.html).

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.

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 in principle can facilitate a more efficient matching of people and opportunities. The platforms mentioned above all have been highly successful in introducing market institutions that exploit this --- creating real-time markets for advertising placements, recommender systems that aggregate user feedback, or allowing for third-party applications that leverage their data about individual user preferences or social contacts.

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, and perhaps even more importantly, the gradual and continuous refinement of things like search algorithms, information displayed to users, product characteristics, and payment and pricing mechanisms. The opportunity to make continuous improvements is particularly relevant because internet platforms can adjust a variety of parameters relatively easily. In fact, a striking feature of many internet platforms is their use of controlled experiments to guide this process. 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.

[3] Heski Bar-Isaac pointed out to me that another form of customization in internet markets, for instance on platforms such as Wikipedia or Yelp, is that users can provide customized input into the product. I will not have much to say about this type of "user-generated content" in this paper.

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. Because this area is still relatively new, I focus on three strands where there have been interesting theoretical or empirical advances. The first is research in industrial organization on how platforms compete for users. The main emphasis in this work is on the role of network effects, and their implications for platform strategy and the dynamics of market structure. I describe the modeling framework and the insights it generates for internet platforms in Section 2.

In Section 3, I discuss a second line of research on the innovative market mechanisms created by many internet platforms. These mechanisms include new types of auctions for advertising, financial products 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 ways of designing market mechanisms that take advantage of the scale and heterogeneity of many internet markets.

In Section 4, I turn from what are mostly theoretical models to discuss empirical research on competition and consumer behavior in online markets. A major focus of this work is how different types of falling costs --- of consumer search, of product proliferation, of using dynamic pricing mechanisms, and so forth --- have affected different markets and industries. I also give some 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 promising directions for future research. Because the area of research I discuss 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 this work is not internet-specific --- its 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.

The main idea in this work is to 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. What is emphasized is that assembling users --- buyers and sellers, consumers and advertisers, groups of friends --- involves network effects. The value users assign to the platform can depend on who else is using it.

The recent modeling efforts build on an older literature in industrial organization 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.

A. The Price Theory of Platforms

The role of network effects are fairly easy to capture using standard price theory (Rochet and Tirole, 2003, 2006; Armstrong, 2006). Here I follow the treatment in Weyl (2010).

Think about 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, $\mathbf{x}=(x_1,...x_K)$ denotes the vector of participant quantities, and ζ are the individual's characteristics. The payoff from not participating is zero. So 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 $C(\mathbf{x})$.

There is an important aspect of coordination in this model because each individual's decision to participate can depends on the decisions of others. If the platform sets fixed prices for each user group, $p=(p_1,...,p_K)$, there can be multiple levels of equilibrium participation. This indeterminacy seems to pose a modeling challenge. Suppose, however, there is some heterogeneity among consumers, so that if individuals in group k expect overall participation $\mathbf{x}$, their willingness to pay to use the platform is continuously distributed. Then for any participation level $\mathbf{x}$ that could result as a participation equilibrium, the fee the platform can extract from each user group is uniquely pinned down --- for group k, it corresponds to the utility of the x_kth most enthusiastic user taking overall participation $\mathbf{x}$ as fixed. Because of this, it is convenient to think of the platform's problem as one of choosing quantities rather than prices (Weyl, 2010).[4]

To follow this route, let $P_k(x_k;\mathbf{x})$ denote the price that exactly x_k users from group k would be willing to pay, assuming they expect an overall number of users $\mathbf{x}$. That is, fixing expectations that participation will be $\mathbf{x}$, $P_k(\cdot;\mathbf{x})$ is the inverse demand curve for group k. Naturally $P_k(x_k;\mathbf{x})$ will be decreasing in x_k, but it may be increasing or decreasing in $\mathbf{x}$ depending on the sign of the

[4] In effect, this means making the implicit assumption that the platform will be able to obtain its preferred participation outcome, should the there be multiple participation equilibria corresponding to the prices it sets. One justification is that if the platform can make the prices contingent on realized participation, it can make $\mathbf{x}$ a unique participation equilibrium. One way to do this, if the platform wishes to implement $\mathbf{x}=(x_1,...,x_K)$, is to charge group k users $P_k(x_k,\mathbf{x}^*)$, where $\mathbf{x}^*$ is the realized participation (Weyl, 2010, and Weyl and White, 2010).

network effects. If there are two user groups, buyers and sellers, we would expect positive cross-group externalities, so $\partial P_k(x)/\partial x_j > 0$ 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

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

The first order conditions for this problem equate the marginal revenue for each user group k with the marginal cost $MC_k(\mathbf{x}) = \partial C(\mathbf{x})/\partial x_k$. There is a convenient way to express these 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

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

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

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

This approach captures some intuitive ideas about platform strategy and competition. For instance, it immediately leads to the conclusion that platforms should engage in cross-subsidization: charging low prices user groups that create value for others and higher prices who do not. So search engines offer free email and other services to users in the hope of making more money from advertisers. Electronic financial exchanges charge low fees, or even payments, to traders who submit limit orders that create liquidity in the order book; they charge higher fees to traders who submit crossing orders that remove liquidity.

The approach also points to a tension between extracting surplus from core users, and lowering prices to expand the user base. The tension is most interesting in a dynamic context, where existing users may be "locked in". Then a platform may be tempted to raises fees on these users at the cost of slowing platform growth. Although this is a standard in switching cost models, an added element here is that by reducing the user base the platform in effect may be degrading its quality for the existing users (a "two-sided" analogue might be an email service or website that expands the number of intrusive video ads).

The model also highlights two familiar distortions from profit-maximization (Weyl, 2010). Apart from focusing on the marginal revenue from additional users rather than their marginal surplus, there is a "Spence distortion": the platform cares about the externality new users impose on marginal rather than average participants. To see this, one can compare the profit-maximizing solution to the conditions for an efficient allocation,

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

which say that group k users should be charged their marginal cost minus the externality they impose on others (which could be positive or negative).[5]

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

Partly as a result, the model does not provide much insight into the *structure* of platform fees. In e-commerce, for example, platforms such as the business-to-business procurement

[5] The importance of the "Spence distortion" in this model raises the question of whether internet platforms really do focus too much attention on marginal users. A priori, one suspects the effect might be countered or even outweighed by the ability of dedicated users to exert influence in ways that do not involve threatening to exit.

platform Alibaba charge sellers a flat annual fee for listing their products. Other platforms such as eBay Motors charge a fee for each individual listing. Still others 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, for the quality of listings and potentially for platform revenue, 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.

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

[6] A related point is that model I've described 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 having a search or application process that lets firms identify high-quality matches. Similar sorting considerations can apply in other contexts as well.

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 a larger scale may have better data to improve its algorithms, and hence may enjoy certain types of scale economy.

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. That is, 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,

[7] Indeed, even in industries such as social networking, where one might expect positive feedback effects to generate agglomeration, it is easy to point to examples of successful entry (Twitter) or rapid decline (MySpace).

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 we 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).

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 they auctions in 2004, at which time eBay had about 80% of the consumer auction market. The results of the auctions were striking: their eBay sales attracted 50% more bidders and the resulting eBay prices were a third higher. Their conclusion was that the market was in the process of tipping all the way toward eBay.

---

[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, available at http://www.sec.gov/rules/concept/2010/34-61358.pdf).

Several features of their 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? Another is how easy the study was to run -- it essentially amounted to running a few dozen auctions on two internet sites.[9] 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 Institutions

The models of platform competition discussed above take a relatively high-level view, focusing on platform pricing and user decisions to join a platform, but not addressing how platforms try to structure economic or social activity in ways that create value. As economic activity has shifted online, however, there also has been a burst of innovative market mechanisms and institutions. Examples include large-scale consumer marketplaces for used goods, services, and loans; real-time markets for advertising and financial trading; online "superstores" for books, retail and media, and many more. In many of these cases, internet platforms have used novel mechanisms particularly engineered to take advantage of the scale, scope and continuous nature of activity permitted by new technology.

In this section, I illustrate these points by discussing two examples of innovations in market design: auctions for “sponsored search” advertising placements, and reputations systems to ensure safe trade in e-commerce marketplaces. Each of these examples has attracted a good

[9] Of course, the overall research process in economics isn't nearly as frictionless as collecting the data. You'll note that the Brown and Morgan study wasn't published until 2009, *five years* after their experiments were conducted.

deal of research attention. At the end of the section, I describe a few related settings that have been studied less, but raise similarly interesting questions about internet market design.

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 to provide search results. 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 in real time. 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. It also determines the price the search engine can charge for selling the space on the results page. Because the value of having an ad appear after any single query is relatively small, the revenue from any single auction is on the order of a few pennies. But 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 auction market can be modeled and analyzed, how the auction markets run by Google and the other search engines evolved to their present form, and why they 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 proposed a very elegant formulation of the problem. There are M 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 also 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 $k \neq m$,

$$v_n x_m - t_m \geq v_n x_k - t_k$$

This condition says that advertiser n prefers to buy position n at a price of t_n than to buy any other position $m \neq n$.[10] 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 below him, and

[10] 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 $m>M$.

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

The minimum market clearing prices, and any set of competitive prices, are increasing on per-click basis as one moves from lower to higher positions.[11] 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 cost of clicks must be increasing as one moves to higher positions.

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 advertisers pay the minimum necessary to sustain their position - in the simplest case, the next highest bid. In practice, each advertiser only pays if the user clicks on its ad, in which case it pays the "per click" bid of the next highest advertiser.

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

[11] 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).

simplifies the market. It reduces the dimension of each advertiser's bid from the number of positions to a single number.[12]

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) 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.[13]

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.

---

[12] Milgrom (2009) 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, and this equilibrium might be the natural outcome to expect with even small bidding costs.

[13] One story I have heard about why search engines switched to the second price format is that the frequent bid changes created a burden for their servers. But Matt Jackson has pointed a separate problem that may also have been a major concern. 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.

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 Nash equilibria with 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).[14] In such an equilibrium, the allocation of bidders to positions is efficient and the payments made by the winning bidders coincide with the payments in some competitive equilibrium. Moreover, for any competitive equilibrium, there is a refined equilibrium of the GSP auction that also has an efficient assignment of advertisers to positions and precisely the same payments.

Of course, 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

[14] Edelman et al. refer to these equilibria as "locally envy free", while Varian calls them "symmetric". To see how the equilibria work, suppose that $b^{(1)} \geq \ldots \geq b^{(n)}$ 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 refinement strengthens the second condition to require $(v^{(k)}-b^{(k+1)})x_k \geq (v^{(k)}-b^{(k)})x_{k-1}$. 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 $t_k = b^{(k+1)}x_k$

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 (2006) and Athey and Ellison (2011) incorporate consumer choices about which ads to click (see also Gomes, 2009, and Jeziorski and Segal, 2009, for some evidence). Ostrovsky and Schwarz (2010) 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 (2010), 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

My second example of a novel market institution 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 a paper by Jin and Kato (2006). 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.

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.

One potential reason why buyers might fail to submit informative 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.

The design of reputation mechanisms is closely related to the problem of designing systems to aggregate user reviews or product evaluations. The design and incentives of these systems, and their effects on consumer purchasing patterns, have attracted considerable attention in computer science and marketing, but somewhat less in economics. Avery, Resnick and Zeckhauser (2001) is one early study of recommender systems that focuses on how to structure the timing of evaluations and payments to consumers. Many of the mechanisms currently in use on platforms such as Netflix or Amazon or consumer review sites do not use payments, but users perhaps enjoy a certain status benefit from posting informative evaluations.

C. Matching and Other Market Design Problems

Both sponsored search auctions and to some extent the design of recommender systems can be seen as particular approaches to the general problem of how to match consumers to products or services in online settings where there are potentially a large array of offerings.

Many online markets face a design problem that is in certain ways quite similar to sponsored search, even if the relevant goods are jobs, retail products, collectibles, real estate, vacation packages, or potential dating partners. Potential "buyers" input a search query and expect to see a ranking of potential "matches". These settings can differ from sponsored search along several dimensions: the information about likely matches can reside with buyers, sellers, or the platform itself; the verifiable outcomes from a successful search might be more or less detailed, allowing for platforms to charge sellers per-transaction or on a commission basis rather than for "per-click" or "per-impression"; and the platform may or may not 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) flip this problem around and look at the search 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.[15]

Another example of a matching problem, and one that is noticeably less structured than ordering search results, arises in auction markets for web page advertising. In these markets, firms bid to have their ads shown to specific consumers, or on specific web pages. The ability to customize advertising down to the individual user and minute creates an interesting tension. On the one hand, creating markets where advertisers can 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), Bergemann and Bonatti (2010), and McAfee, Papineni, and Vassilvitskii (2010) have made progress on modeling these issues.

Finally, a third set of interesting problems relate to the potential for continuous trading in many internet markets. For example, in the internet advertising markets just mentioned, many advertisers arrange in advance to have their ads shown for particular blocks of time or to a guaranteed number of users. But residual inventory is auctioned in real time. Alternatively, electronic financial exchanges allow trading in real time, but certain platforms such as the NYSE

[15] The papers by Athey and Ellison (2011) and Gomes (2009) mentioned above are also relevant. One interesting question here 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.

operate separate markets for large trades that clear at discrete intervals.[16] The costs and benefits or clearing markets in real time, as opposed to at discrete intervals or through pre-arranged contracts, have not been that well studied in the literature on market design.

4. Evidence on Competition and Consumer Behavior

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, or making it easier to customize products and information, or of reducing the transaction costs associated with sophisticated pricing or sales mechanisms? In addition, the ability to capture and record behavior and transactions in great detail makes internet markets a natural environment to test broader theories about consumer decision-making and imperfect competition.

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,

16 As a third example, one can ask why in sponsored search each query should trigger a new auction, and why bidders are allowed to continuously update their bids. One justification is that the number and type of searches is variable, so advertisers attempting to allocate a budget across keywords will want to optimize dynamically, and prices should similarly adjust. Another is that bidders are learning about their values or click rates, or about the behavior of their competitors. Both stories suggests that bidding is unlikely to be stable as in the static GSP analysis.

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).[17]

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. In the model of Baye and Morgan (2001), some consumers still fail to compare prices even after the arrival of internet price comparisons, and this results in an equilibrium distribution of prices. More subtly, 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). They provide support for this argument by combining data 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 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,

[17] There are some caveats. For instance, Ghose and Yao (2009) point out that most studies of price comparison engines look for dispersion in posted prices, rather than in transaction prices. They provide evidence from the GAO's procurement marketplace, showing that transaction prices feature much lower dispersion.

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 but it may not be entirely representative. Einav, Kuchler, Levin and Sundaresan (2011) look at tens of thousands of more 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.

A potentially more robust 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 et al. (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 the combination of low search costs and low costs of product proliferation would shift consumer demand away from common items 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.[18]

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 lowing 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. 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.

Gentzkow and Shapiro (2010) 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

[18] 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.

purchasing and browsing decisions, this debate about shifts in product market concentration seems ripe for further analysis.

C. Auctions and Dynamic Sales Mechanisms

A third prominent hypothesis dating to 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, for example, 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 (2011), however, point out that over time the fraction of listings that are posted price has increased dramatically, from less than 10% to over 90%. There are several possible explanations, but one is that over time, the internet has not just reduced the transaction cost of running auctions, it has reduced the cost of finding prices for comparable goods, so that there is less need to use an auction to discover the appropriate market price.

Even when posted prices are generally efficient, there are certain types of perishable 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. Another example, studied by Sweeting (2010), online ticket sales. He also finds that prices adjust downward as events approach, and relates this to the predictions of dynamic pricing models, where the opportunity cost of selling falls as the "expiration date" approaches. Of course, the equilibrium price path can depend on whether buyers anticipate price declines and time their purchases strategically. Board and Skrzypacz (2010) show in the presence of strategic buyers, the optimal sales strategy for a perishable goods seller typically involve dynamic price adjustment followed by a last-minute auction if there is remaining inventory.

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 one slice of the economic research this shift has inspired, research on the development of internet platforms and new online marketplaces. 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 a defining feature of many internet markets is innovation, and particularly incremental innovation that occurs continuously, often accompanied and guided by systematic experimentation. But you'll notice that this aspect of internet markets hasn't been a major focus of the research I've discussed. 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 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, 2010, and Ostrovsky and Schwarz, 2010 are two recent and large-scale examples). Einav, Kuchler, Levin and Sundaresan (2011) point out that researchers also can make use of "experiments" conducted by market participants. They identify hundreds of thousands 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.

A related and final 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 industry competition 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," in progress.

Armstrong, Mark, "Competition in Two-Sided Markets," *Rand Journal of Economics*, 37(3), Autumn 2006, pp. 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,* forthcoming.

Athey, Susan and Denis Nekipelov, "A Structural Model of Sponsored Search Auctions," UC Berkeley Working Paper, 2010.

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

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

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.

Board, Simon and Andy Skrzypacz, "Optimal Dynamic Auctions for Durable Goods: Posted Prices and Fire-sales," 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, 2008.

Bonatti, Alessandro and Dirk Bergemann, "Targeting in Advertising Markets: Implications for Offline vs. Online Media," Yale University Working Paper, 2010.

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 Field Experiments," *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*, 46(4), April 2000, pp. 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*, 49(11), November 2003, pp. 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*, 96(3), June 2006, pp. 652-668.

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

Chen, Yongmin and Chuan He, "Paid Placement: Advertising and Search on the Internet," NET Institute Working Paper, September 2006.

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

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*, 43(1), 2007, pp. 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*, 97(1), 2007, pp. 242-259.

Einav, Liran, Chiara Farronato, Jonathan Levin and Neel Sundaresan, "Auctions and Posted Prices in Internet Commerce," in progress.

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

Einav, Liran and Jonathan Levin, “Empirical Industrial Organization: A Progress Report,” *Journal of Economic Perspectives*, 24(2), Spring 2010.

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

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

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

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

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

Gomes, Renato, "Mechanism Design in Two-Sided Markets: Auctioning Users," Northwestern University Working Paper, 2009.

Hagiu, Andrei and Bruno Julien, "Why Do Intermediaries Divert Search?" Working Paper, 2010.

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

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

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

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

Lee, Young Han and Ulrike Malmendier, "The Bidder's Curse," *American Economic Review*, forthcoming.

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

Lewis, Randall and David Reiley, "Does Retail Advertising Work? Measuring the Effects of Advertising on Sales via a Controlled Experiment on Yahoo!," Yahoo! Research, 2010.

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*, 48(3), September 2000, pp. 227-252.

McAfee, R. Preston, Kishore Papineni, and Sergei Vassilvitskii, "Maximally Representative Allocations for Guaranteed Delivery Advertising Campaigns," Yahoo! Research, 2010.

Milgrom, Paul, "Simplified Mechanisms, with an Application to Sponsored Search Auctions," *Games and Economic Behavior*, 2009.

Ostrovsky, Michael and Michael Schwarz, "Reserve Prices in Internet Advertising Auctions: A Field Experiment," Stanford University Working Paper, 2010.

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

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

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

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

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

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

Sweeting, Andrew, "Price Dynamics in Perishable Goods Markets: The Case of Secondary Markets for Major League Baseball Tickets," Duke University Working Paper, 2009.

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

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

Varian, Hal, "Computer-Mediated Exchange," *American Economic Review*, 100(2), May 2010.

Weyl, E. Glen, "A Price Theory of Multi-Sided Platforms," *American Economic Review*, 2010.

Weyl, E. Glen and Alex White, "Imperfect Platform Competition: A General Approach," Harvard University Working Paper, 2010.