"Introduction" to *The Economics of E-Commerce*[1]

by

Michael R. Baye and John Morgan

Premiminary Draft, January 2016

1. Overview of Electronic Commerce

In his famous Internet Tidal Wave memo, Bill Gates wrote that "The Internet is the most important single development to come along since the IBM PC was introduced in 1981." Certainly in the field of e-commerce—the use of computer networks to directly or indirectly facilitate the exchange of goods, services and/or information—Gates' vision has proved entirely correct. Measured by gross merchandise volume on its platform, Amazon has arguably emerged as one of the world's largest retailers, online or off. In part, this derives from the product breadth and superior search capabilities it offers. A consumer visiting Amazon can quickly identify an array of offerings and sort them according to a variety of criteria ranging from price to product or merchant reviews. Pre-Internet, such information gathering would require formidable effort if done offline. Uber, a ride-sharing service, has drastically altered the way people get around major cities. Review sites like Yelp can now make or break the reputations of small and large businesses alike.

The Internet has also changed the way that consumers pay for goods. Major credit cards like Visa were slow to adapt, and this created space for services like PayPal to offer funds transfers for the many micro-businesses setting up shop on eBay. More recently, mobile payment services like Square and ApplePay have enabled even individuals to exchange electronic funds quickly and seamlessly.

Such a sea change has not gone unnoticed by researchers. A search for "electronic commerce" on Google Scholar returns over 1.5 million results. While massive in number, these articles may be broadly categorized into one or more of the following six themes:

- Platforms as strategic actors.
- Is e-commerce frictionless? Price dispersion and obfuscation online.

[1] Forthcoming in *The Economics of E-Commerce*, Edward Elgar Publishing, Ltd.

- Does e-commerce benefit consumers? Information, product variety, and competition.
- Online auctions for products and advertising.
- Mitigating informational asymmetries: Reputation, recommendation, and feedback mechanisms.
- Battles for digital property rights: Privacy, piracy, and the open source movement.

The remainder of this review highlights some of the key findings in each theme. The reader seeking a more detailed understanding should consult the papers included in *The Economics of E-Commerce*, as well as the references contained therein. Collectively, these articles represent some of the more influential works on e-commerce.

2. Platforms as Strategic Actors

Traditional models of consumer search highlight market frictions, such as the cost of visiting various stores, as a central feature. These costs are treated as exogenous, outside the control of any particular party. In contrast, e-commerce platforms, such as Amazon for consumer products or Google in online advertising, do not merely create markets, but *design* them. That is, market "frictions"—the information available, the ease of search, the cost of access—are, to some degree, under the control of the platform. In short, the e-commerce platforms are themselves strategic actors. Baye and Morgan (2001) offer the first e-commerce model incorporating this key observation. In their model, a profit-maximizing "gatekeeper" charges fees to both consumers and firms wishing to utilize its platform to acquire and transmit information. Subsequently, additional models of platforms that serve two-sided markets have been developed by Rochet and Tirole (2003), Caillaud and Jullien (2003), Armstrong (2006), and Hagiu (2009), among others.

Pricing on these sorts of platforms differs from conventional markets in another key respect—e-commerce platforms often exhibit virtuous circles of benefits among their various users. For instance, eBay, in determining its fee structure, must account for the fact that, the more sellers there are on the platform, the more valuable the platform is to potential buyers and vice-versa. Thus, even a monopoly platform must exercise care in capturing surplus from buyers owing to the impact this has on the incentives of sellers to participate. For instance, the Baye and Morgan (2001) model implies that a monopoly platform benefits from keeping consumer access fees low, possibly even free, and instead, earning the lion's share of its profits from the fees charged to sellers transmitting information.

Owing to network effects stemming from such virtuous circles, a single platform is often optimal from an efficiency standpoint. Many e-commerce environments approximate this structure. For example, Facebook represents far and away the largest social network. But the presence of a single platform gives rise to two risks: the exercise of market power in pricing and the possibility that an inferior platform emerges as the winner. Caillaud and Jullien (2003) examine the first such risk using a theory model. They conclude that efficient structures can

emerge, even in a monopoly setting, but that there are no guarantees. Tellis, Yin and Niraj (2009) examine the second and come to the hopeful conclusion that, despite the theoretical possibility of lock-in to an inferior platform, actual "winning" platforms are often the most efficient.

3. Is E-Commerce Frictionless? Price Dispersion and Obfuscation Online

Despite the ability of platforms to create frictions through pricing strategies, it is undeniable that the Internet has vastly reduced information-gathering costs, and thereby reduced the associated search frictions. At the dawn of e-retail, *The Economist* boldly predicted that this effect alone would lead to markets closely approximating textbook models of perfect competition. Competitive pressure from perfectly informed consumers would seemingly force all firms to charge the same (low) price.

Brynjolfsson and Smith (2000) were the first to examine this question, focusing on markets for books and music CDs, which were some of the leading goods first sold online. While they found substantial evidence that online prices were lower than offline, the law of one price was egregiously violated in online markets. For a typical book, the range in prices was about 33% while for CDs it was 25%, a far cry from the 0% prediction of competitive models. As e-commerce continued to grow, others examined the same prediction in other markets (see, e.g., Baye *et al*., 2004), and found the same phenomenon—the law of one price was no law at all, even online.

While the Internet offers a wealth of information, does this really produce perfectly informed consumers? Ellison and Ellison (2009) suggest that the answer is no. They note that, in markets for computer memory, while the price of the baseline product is easily observable, the price of the actual product a consumer might desire is obfuscated. In particular, sellers in this market engage in the practice of offering extremely low prices for low quality products with the hope of inducing consumers to visit their store and upsell them on a higher quality product. Ellison and Ellison show that, absent such strategies, individual sellers face highly elastic demand curves, effectively condemning them to Bertrand competition. Baye *et al.* (2009) also document significant downward jumps in a firm's demand when its price falls slightly below the lowest price charged by online rivals.

Gabaix and Laibson (2006) offer a theory model of the more general phenomenon of price and product shrouding—a strategy in which sellers intentionally reveal only a portion of potentially relevant information to consumers. They highlight behavioral factors where such practices are profitable. Brown *et al.* (2010) show the global reach of such practices, highlighting how the suppression of shipping and handling costs in online auctions can raise seller profits in US, European, and Asian online markets.

4. Does E-Commerce Benefit Consumers? Information, Product Variety, and Competition

While lower online prices represent a direct benefit of giving consumers (and businesses) low-cost access to information, indirect benefits, such as lower offline prices or greater product variety, might also arise. To see how this might work, consider a bargaining situation occurring offline. Prior to the negotiation, a consumer gathers some information which she uses to determine her opening offer and reservation price. The seller, far better informed, does likewise. The consumer's relative lack of information puts her at a disadvantage in the negotiation. The Internet, however, radically reduces the cost of information gathering as well as increases the availability of certain key pieces of information, such as the negotiated outcomes of other consumers. Armed with this information, the bargaining situation changes dramatically—the consumer is in a much stronger position to obtain her preferred outcome. The main question, then, is how large such theoretical benefits are in practice.

Morton *et al.* (2001) examine bargaining situations like the one above in the context of purchasing a car. While car buying traditionally involved visiting a dealership and engaging in protracted haggling, the Internet offers the possibility of obtaining price quotes for vehicles online, through a referral service. They find that consumers using the service save about 2%, roughly $450, compared to consumers who do not.

Brown and Goolsbee (2002) investigate how the information available online changed insurance markets. For many consumers, insurance represents a complex product with often difficult comparisons across firms and products. Traditionally, search was extremely costly, as it required meeting with different insurance agents and listening to their sales pitches. Brown and Goolsbee note that the presence of online information vastly simplifies the process of comparing the different life insurance policies offered by competing insurers. They show that the premiums on term life insurance products in which price information was available online fell by 8 – 15%, even though the products themselves were purchased offline. By contrast, prices for insurance products where such online information was lacking remained unchanged. In a similar vein, Forman *et al.* (2009) provide evidence that online booksellers compete with offline booksellers. Together, this strand of the literature highlights the importance of spillovers between the information and/or competition online and pricing offline.

Perhaps more than any other aspect of shopping, the Internet has fundamentally altered the availability of even the most obscure items. For those consumers whose tastes ran toward niche products, like collectibles, considerable time and energy was spent merely trying to locate a desired variety of the product. Stores offering such items tended to be small or, in some locations, non-existent. As a result, much valuable trade never took place. For instance, consider two baseball card shops, one in Atlanta, the other in New York. The Atlanta shop has a Mickey Mantle rookie card while the New York shop has a Hank Aaron rookie card. A consumer in Atlanta desperately wants an Aaron rookie while a Brooklyn consumer will pay top dollar for Mantle. In the pre-Internet age, both consumers (and sellers) are frustrated in this situation.

But with sites like eBay, small card shops now have an international market presence and consumers can get the cards they desire.

Brynjolfsson, Hu and Smith (2003) set out to measure the welfare gain from this increased access to variety online. Focusing on bookstores, they show a vast increase in the varieties available to consumers, especially those living outside major cities. They estimate welfare gains on the order of $700 million to $1.03 billion for the year 2003. Since that time, the variety of goods on offer in online markets has increased substantially. Were we to perform the same computation today, it would undoubtedly be vastly larger. Indeed, gains from increased variety probably represent the Internet's single biggest contribution to consumer welfare.

5. Online Auctions for Products and Advertising

With the opening up of a panoply of varieties to a host of potential buyers worldwide comes a new problem, price finding. If you discover an old, pretty lamp of unknown provenance in your grandmother's attic, how much should you charge for it? Professional sellers may be able to quickly answer this question, but the Internet has also unleashed an army of amateurs, where pricing represents a daunting task. The traditional solution to the price finding problem, auctions, turns out to be ideally suited to online environments.

Auctions have been used, in a limited fashion, in offline markets for many years. One of their main difficulties is liquidity—to succeed, an auction must somehow attract interested buyers to participate at a given date and time, no easy feat. Auctions are also expensive to run, often requiring dedicated space, a professional auctioneer, glossy brochures, and so on.

Moving to an online environment solves many of these difficulties. EBay, for example, neatly solves the liquidity problem by allowing auctions to occur over a series of days and permitting bidders to make proxy bids, bids that are executed only if needed for the individual to remain the high bidder. With broadband penetration, it is now possible to include sleek photos of the item, often close to professional quality, using nothing more than a smartphone.

Bajari and Hortaçsu (2003) study bidding and selling behavior on eBay, the most prominent online auction platform in the US and Western Europe. They observe that sellers tend to set opening bids lower than the book price for an item and to eschew secret reserve prices. This makes the auction enticing to many bidders, which in turn improves profits. Bidders often engage in "sniping," bidding at the last possible instant in the auction. For items with a strong common value component, such behavior is entirely sensible. A bidder who is knowledgeable about an item has no wish to supply this information to rivals in the form of an early high bid. Thus, the bidder waits until the point in the auction where it becomes impossible for rivals to incorporate this information into bids, i.e. the very end of the auction.

When eBay began its existence, all transactions were concluded via auction, but it has evolved greatly over the years. EBay is no longer a specialized platform for selling used or collectible items; it has evolved into a marketplace where merchants sell brand new products

and compete with traditional retailers. As a consequence, eBay also offers a "buy it now" option, essentially a conventional price tag for the item. By 2014, the majority of items, measured by gross merchandise volume, were being sold via buy-it-now rather than through auction. While price formation through auctions continues to have an important role, it is an inefficient mechanism for selling items with a well-established "street price."

Not only has the Internet proved remarkable at vastly increasing the variety of items available to a buyer, it has also led to the creation of entirely new products and services with extraordinary variety. For instance, whenever an individual enters a search term in Google or any other search engine, a unique results page appears. These results change depending on the user's location and based on various characteristics known about the user. For advertisers, search queries represent the amazing opportunity of presenting the consumer with precisely the solution they are seeking based on the query.

The idea of placing advertising amidst content that a user is browsing is not new. Newspapers have long used this strategy, interweaving ads with news articles. The difference, however, is that a query contains a strong clue as to what is on the user's mind at that particular moment. Since there are trillions of possible queries and these may be further subdivided by location, user characteristics, the time of day, and so on, the problem of pricing this type of advertising is formidable.

Compared with conventional sales of advertising, the solution devised by Google and other search engines is radical in two ways. First, rather than charging per impression, as in conventional advertising, advertisers are only charged if a user clicks on their particular ad; that is, there is a cost per action rather than a cost per impression. Second, rather than selecting a price schedule, Google opted to run auctions for words or phrases appearing in search queries. Thus, a mobile phone company such as Verizon might bid for various keywords like "Verizon," "cell phone," "mobile phone," and so on. The outcome of these auctions determines which advertisers are displayed following a search query containing the words on which they bid. For example, a search query containing the word "Verizon" might produce results where Verizon's ad appears at the top of the results page, perhaps followed by ads for AT&T and T-Mobile. That is, there are multiple ad "slots" over which advertisers are bidding.

The first such ad auctions, run by a company named Overture, used a simple process to determine slot placement and cost per click. The high bidder received the top slot, the second-highest received the second slot, and so on. Whenever an ad was clicked, a bidder simply paid its bid amount. Google changed this by adopting a variant of the "generalized second price" (GSP) auction form. When advertisers are otherwise identical, the GSP awards the top slot to the high bidder, the second to the second-highest, and so on, exactly like Overture, but the payment rule differs sharply. The high bidder pays the second-highest bid, the second pays the third-highest, and so on. Being new, the properties of this auction form were little understood.

Varian (2007), together with Edelman, Ostrovsky, and Schwarz (2007), first analyzed the GSP auction formally. Using the solution concept of static *ex post equilibrium* to proxy for the dynamic competitive game played by advertisers, they show that the GSP has remarkable properties. For example, it is efficient; the bidder having the highest value gets the top slot, etc. Additionally, it leads to the same payments as a Vickrey-Clark-Groves mechanism, i.e. bidders pay the amount of the externality they impose on others. These results imply that the GSP auction maximizes revenue to the seller conditional on achieving an efficient outcome. Edelman and Schwartz (2010) then extend this analysis to study the optimal reserve prices in GSP auctions, in effect, finding the fully revenue maximizing auction structure for the seller.

6. Mitigating Information Asymmetries: Reputation, Recommendation, and Feedback Mechanisms

By breaking down geographic barriers and creating a truly global marketplace, the Internet offered a vast new frontier of selling opportunities. Moreover, the relatively anonymous nature of online transactions meant that such opportunities were particularly attractive to unscrupulous sellers (or buyers, for that matter). Thus, a fundamental challenge was to create a system of trust and thus avoid the breakdown of the market owing to the "lemons problem" (Akerlof, 1970) associated with unsavory behaviors in anonymous transactions. A simple solution is to establish a brand, much as in offline space. Amazon, for instance, succeeded in part because of the trust it engendered with buyers. But for platforms like eBay, acting as market intermediaries rather than as sellers themselves, having a good brand was no guarantee of success. Thus were born reputation systems. Initially, such systems were primitive and easily gamed. For instance, a seller derived "reputation" on eBay merely by having other users "endorse" his identity. Stories of groups of confederate identities endorsing one another so as to fool buyers are well documented; Dellarocas (2003) surveys some of the early challenges associated with online feedback mechanisms.

Reputation systems soon evolved to combat these sorts of problems. EBay replaced its initial system with a bilateral system where only actual buyers could give reputation to sellers and, moreover, could do so only once. Thus, it was no longer possible for two separate identities of the same individual to artificially inflate each other's reputation. This system proved relatively successful in convincing buyers of the value of eBay reputations. In an early study, Resnick *et al.* (2006) show that a reputable seller enjoyed about an 8.1% premium over a seller lacking reputation. But all reputation was not created equal—reputation earned as a buyer proved much less valuable than that earned as a seller, as Houser and Wooders (2006) document. Reputation is not merely important to price, but also to the probability of a sale (Chevalier and Mayzlin, 2006).

But this bilateral system still contained problems. Most notably, because buyers derived no benefit from the reputation provided by sellers, there was a potential free-rider problem in the provision of reputation. If reputation is a hassle for buyers to provide, and they obtain no

benefit from doing so, then, quite naturally, they will be tempted to free ride off of the feedback provided by others. Using laboratory experiments, Bolton *et al.* (2004) demonstrated that this was not merely a theoretical possibility. While markets with reputation in place vastly increased the economic efficiency of transactions, there was substantial under-provision of feedback information in these markets. Their findings likely understate the case in terms of real markets. EBay's bilateral scheme offered unscrupulous sellers a tool to punish buyers who seek to "out" their behavior. Sellers would routinely retaliate against buyers giving negative feedback by replying in kind. One might think that such a threat would be relatively meaningless given that buyer feedback hardly seems to matter, but since many individuals on eBay play both roles, sometimes acting as buyers and sometimes as sellers, the threat of harm, and hence the effective deterrence gained by employing such strategies, was very real.

Chen *et al.* (2010) suggest one mechanism that might ameliorate the free rider problem—social comparison. Using a crowd-sourced movie site, they performed a field experiment where subjects were given feedback about the median number of ratings made by others on the site. Those above the median exhibited no change in behavior, but those below increased their ratings activity by 530%.

The bilateral system suffered from additional problems that spawned new markets for reputation. In a poignant example, a common EBay transaction involved the sale of an E-book which was in the public domain and freely downloadable. This E-book offered the following advice: Post this E-book for sale on eBay and use it to gain reputation points. In this way, you will soon become a reputable seller. Such manufactured reputation proved even more effective owing to eBay's decision only to publish details of transactions occurring over the preceding 90 days. Thus, sellers earning reputation in these markets could lie low for a period of time only to emerge with a reputation score exactly like that of a legitimate seller. Brown and Morgan (2006) tell the story of one such seller who re-emerged to sell lakefront property in Texas before having his or her account suspended by eBay.

Markets for reputation continue to evolve. On eBay, reputation is now one-sided: Buyers rate sellers, but not vice-versa. Perhaps more importantly, eBay has directly intervened in markets, offering various kinds of buyer insurance protecting buyers from unscrupulous sellers. Amazon offers a similar level of protection when buying from independent sellers listed on its site. Global marketplaces, such as Taobao, employ similar reputation systems in supporting sellers.

7. Battles for Digital Property Rights: Privacy, Piracy, and the Open Source Movement

E-commerce has not merely altered the manner in which items are purchased, but in many cases it has changed the nature of the products themselves. Unsurprisingly, this transformation has been especially pronounced in the area of information goods, items such as e-books, digital music, video games, and so on. By stripping away the physical instantiation of these products in favor of a purely virtual form, the ease by which these items may be

transmitted, shared, and enjoyed has radically shifted. For instance, it is now entirely possible to carry around all of the written material contained in the US Library of Congress on a portable hard drive.

But this same ease of sharing has also led to a burgeoning of intellectual property crime, so-called digital piracy, whereby individuals who do not own these virtual goods nonetheless help themselves to a wide array of offerings. Nowhere has the battle of digital piracy been more vicious and heated than in the realm of digital music. In the early 2000s, Napster was one of the most popular sites online. This site contained a repository of digital music, uploaded by its users, and freely shared. Thus, someone interested in the song "I've Got You Babe," by Sonny and Cher, could simply search Napster, locate one or more copies of this recording, and freely download them whether they actually owned the rights to the song or not. Over the same period of time, sales of physical copies of music CDs fell off dramatically.

Recording industry executives tied this sales decline to the rise of digital piracy while Napster's defenders argued that pirated music might spur demand for CDs by exposing individuals to more music and artists. Rob and Waldfogel (2006) offer evidence in favor of the first explanation, finding evidence that piracy displaces demand for paid content. Oberholzer-Gee and Strumpf (2007), however, find a zero effect although, even to the present, their findings remain hotly contested. In many respects, the debate about digital music piracy has become less important. Through sites like iTunes and subscription services such as Spotify, convenient, low-cost legal ways of obtaining digital music have emerged. Thus, just as piracy may have been a substitute good for music CDs, now legal streaming and downloading acts as a substitute for piracy.

With the rise of social media in particular, advertisers have a wealth of data about individuals browsing online. Advertisers make use of this data in the form of targeted advertising—ads that are specifically directed at an individual user even to the point of "following" that user as she visits various different content websites. While targeted advertising has existed for many years prior to the dawn of the Internet, the ability to collect and analyze personal data from a wealth of sources has made this technique immensely more powerful. But as this power becomes evident, consumers and governments are increasingly wary about the infringement of personal privacy such data collection entails. The European Union has been in the forefront of privacy protection law, imposing a variety of stringent constraints limiting or barring the use of such data in targeted advertising. While consumers perhaps benefit from these protections, the costs of such bans are unclear. Hermalin and Katz (2006) offer a theoretical perspective. They measure the costs of privacy protection as foregone efficiency gains. For instance, when a consumer does not learn of a particular product that might represent a valuable solution to some problem owing to loss of ability to target, that represents an efficiency loss to society. They show, however, that the tradeoff between efficiency and privacy may, under certain circumstances, be a false one. By cleverly restricting the

transmission of information, they show that it is sometimes possible to have substantial privacy protection at little or no efficiency cost.

Goldfarb and Tucker (2011) offer an empirical view of the matter. Comparing the effectiveness of targeted advertising with and without privacy constraints, they find a substantial drop in advertising effectiveness measured in terms of changes in a consumer's purchase intensity. The effect is especially pronounced for advertising on general content sites, such as news outlets, and for small firms. Implicitly, their findings highlight that an unintended cost of privacy may be to make broad line news outlets less financially viable as well as to hinder the growth of small firms. This may be a societal price worth paying, but being mindful about where the financial burdens of such regulations lie represents an important piece of the cost-benefit puzzle.

Computer code itself represents the ultimate virtual good. Most code, such as that used to operate Windows or Apple operating systems, is proprietary, yet there has been a large-scale movement to create broadly useful code that may be freely shared. Linux and Apache represent the two largest successes of this open-source movement. Lerner and Tirole (2002) offer a survey of the economics of open source. While the fast-changing nature of software coding renders many of the applications cited in the work somewhat dated, the broader economic lessons are as relevant today as they were when the piece was first penned. The key tradeoff boils down to the power of crowd-sourcing versus free-rider problems. Proprietary code incentivizes programmers with direct financial rewards for their efforts but at the cost of limiting the scope from whence ideas might derive. Open-source, on the other hand, reflects the joint intelligence of hundreds or thousands of contributors, each freely sharing ideas. But the lack of direct incentives for financial gain make such efforts primarily "labors of love." Despite this limitation, the quality and durability of open-source products has proved remarkable.

8. References

George A. Akerlof (1970), `The Market for "Lemons": Quality Uncertainty and the Market Mechanism', *The Quarterly Journal of Economics,* 84 (3), 488–500.

Mark Armstrong (2006), 'Competition in Two-Sided Markets', *The RAND Journal of Economics*, **37** (3), Autumn, 668-691.

Patrick Bajari and Ali Hortacsu (2003), 'The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions', *The RAND Journal of Economics,* **34** (2), Summer, 329-355.

Michael R. Baye, J. Rupert J. Gatti, Paul Kattuman and John Morgan (2009), 'Clicks, Discontinuities, and Firm Demand Online', *Journal of Economics & Management Strategy,* **18** (4), December, 935-975.

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

Michael R. Baye, John Morgan and Patrick Scholten (2004), 'Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site', *The Journal of Industrial Economics,* **52** (4), December, 463-496.

Gary E. Bolton, Elena Katok and Axel Ockenfels (2004), 'How Effective are Electronic Reputation Mechanisms? An Experimental Investigation', *Management Science,* **50** (11), November, 1587-1602.

Jeffrey R. Brown and Austan Goolsbee (2002), 'Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry', *Journal of Political Economy,* **110** (3), June, 481-507.

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

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

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

Bernard Caillaud and Bruno Jullien (2003), 'Chicken & Egg: Competition Among Intermediation Service Providers', *The RAND Journal of Economics,* **34** (2), Summer, 309-328.

Yan Chen, F. Maxwell Harper, Joseph Konstan and Sherry Xin Li (2010), 'Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens', *The American Economic Review,* **100** (4), September, 1358-1398.

Judith A. Chevalier and Dina Mayzlin (2006), 'The Effect of Word of Mouth on Sales: Online Book Reviews', *Journal of Marketing Research,* **43** (3), August, 345-354.

Chrysanthos Dellarocas (2003), 'The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms', *Management Science,* **49** (10), October, 1407-1424.

Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz (2007), 'Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords', *American Economic Review* 97 (1), 242-59.

Benjamin Edelman and Michael Schwarz (2010), 'Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions', *The American Economic Review,* **100** (2), May, 597-602.

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

Chris Forman, Anindya Ghose and Avi Goldfarb (2009), 'Competition Between Local and Electronic Markets: How the Benefit of Buying Online Depends on Where You Live', *Management Science,* **55** (1), January, 47-57.

Xavier Gabaix and David Laibson (2006), 'Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets', *The Quarterly Journal of Economics*, **121** (2), May, 505-540.

Avi Goldfarb and Catherine E. Tucker (2011), 'Privacy Regulation and Online Advertising', *Management Science,* **57** (1), January, 57-71.

Andrei Hagiu (2009), 'Two-Sided Platforms: Product Variety and Pricing Structures', *Journal of Economics & Management Strategy,* **18** (4), December, 1011-1043.

Benjamin E. Hermalin and Michael L. Katz (2006), 'Privacy, Property Rights and Efficiency: The Economics of Privacy as Secrecy', *Quantitative Marketing and Economics*, **4** (3), September, 209-239.

Daniel Houser and John Wooders (2006), 'Reputation in Auctions: Theory, and Evidence from eBay', *Journal of Economics & Management Strategy,* **15** (2), June, 353-369.

Josh Lerner and Jean Tirole (2002), 'Some Simple Economics of Open Source', *The Journal of Industrial Economics,* **50** (2), June, 197-234.

John Morgan and Jennifer Brown (2006), 'Reputation in Online Markets: The Market for Trust', *California Management Review*, **49** (1), 61-81.

Fiona Scott Morton, Florian Zettelmeyer and Jorge Silva-Risso (2001), 'Internet Car Retailing', *The Journal of Industrial Economics,* **49** (4), December, 501-519.

Felix Oberholzer-Gee and Koleman Strumpf (2007), 'The Effect of File Sharing on Record Sales: An Empirical Analysis', *Journal of Political Economy,* **115** (1), February, 1-42.

Paul Resnick, Richard Zeckhauser, John Swanson and Kate Lockwood (2006), 'The Value of Reputation on eBay: A Controlled Experiment', *Experimental Economics,* **9** (2), June, 79-101.

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

Rafael Rob and Joel Waldfogel (2006), 'Piracy on the High C's: Music Downloading, Sales Displacement, and Social Welfare in a Sample of College Students', *Journal of Law and Economics,* **49** (1), April, 29-62.

Gerard J. Tellis, Eden Yin and Rakesh Niraj (2009), 'Does Quality Win? Network Effects Versus Quality in High-Tech Markets', *Journal of Marketing Research,* **46** (2), April, 135-149.

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