

Two-Sided e-commerce Marketplaces and the Future of Retailing

Steven Tadelis

Haas School of Business
University of California, Berkeley
and NBER
stadelis@berkeley.edu

March 31, 2015

Abstract

This chapter describes the growth and proliferation of online two-sided marketplaces starting with their emergence in the mid 1990s. I begin with a brief history of online marketplaces, followed by some of the key strategic and economic considerations of online marketplaces. Special emphasis is placed on price discovery and sales mechanisms, as well as on the importance of reputation and feedback mechanisms in the success of online marketplaces. I also discuss how online marketplaces have emerged as a laboratory to test a variety of theories about buyer and seller behavior in the field.

“People were doing business with one another through the Internet already, through bulletin boards. But on the Web, we could make it interactive, we could create an auction, we could create a real marketplace. And that’s really what triggered my imagination, if you will, and that’s what I did.” -- Pierre Omidyar, Founder of eBay

1. Introduction

The exchange of goods and service between parties is perhaps one of the most important economic activities in which value is created through gains from trade. Frictions and barriers to trade can limit the value created by trade, which in turn motivate both individuals and societies to find ways to reduce market frictions and enhance gains trade. One of the most obvious frictions to trade is a lack of coordination. Millennia ago, individuals bartered goods and services, which meant that only two individuals who had what the other wanted could trade, often referred to as the problem of “the double coincidence of wants.” Initially the use of precious metals such as silver and gold, and later the advent of fiat money, were ingenious ways to solve this friction.

Another important friction that can hamper trade is whether traders have adequate information about where to find each other. Historically, buyers and sellers would have to meet, and hence search for their trade counterparts. Coordinating where and when trade took place was another important advent of societies, which can be seen in the introduction of trade fairs in medieval Europe

(Grief, 2000). And what more, the successful operation of these trade fairs, where people were expected to trade with counterparts whom they had never met, rested on governance and reputation mechanisms that gave people the faith to trade with strangers (Milgrom et al., 1990).

These trade fairs represent one of the very first examples of *two-sided markets*. Though the literature does not completely agree on a clear and precise definition of what a two sided market is, Rysman (2009) starts his discussion with a rather intuitive, albeit informal definition:¹ “[a] two-sided market is one in which 1) two sets of agents interact through an intermediary or platform, and 2) the decisions of each set of agents affects the outcomes of the other set of agents, typically through an externality” (p. 125). The medieval European trade fairs were a coordinated location where buyers and sellers can find each other, and where additional institutions were put in place to facilitate trade. The coordinated location and the governing institutions constituted the platform, and the externality was one of a network form: the more sellers there were, the more buyers would benefit from coming, and vice versa.

Many hundreds of years later, two sided markets play a central role in commerce. These include, but are not limited to, farmer’s markets (the platforms), where the two sides are local growers and consumers; newspapers (the platform), where the two sides are readers and advertisers; credit card providers (the platform), where the two sides are merchants and consumers; e-commerce marketplaces (the platform), which match buyers and sellers;

¹ The more formal theoretical literature on two sided markets, which does not offer a unified definition, started with the seminal papers of Armstrong (2006) and Rochet and Tirole (2003, 2006).

operating systems (the platform) where the two sides are software developers and software users; and many more. This chapter focuses on online marketplaces, which are growing in popularity and economic importance at a dizzying rate since the introduction of the Internet just two decades ago.

Indeed, in these past two decades the Internet has transformed the retail sector as more and more consumers substitute purchases from bricks-and-mortar stores to online retailing. Alongside marketplaces like eBay, which are virtual platforms with no physical presence, Amazon.com is perhaps the leading example of a sophisticated retailer that has invested heavily in using petabytes of data to invest in improving operations, logistics and distribution channels, offering a huge variety of goods at competitive prices and with a sharp focus on customer support and satisfaction. What is interesting is that despite starting as a retailer in 1995 and not as a two sided market, Amazon waited until the year 2000 to leverage the many millions of visitors to enhance its offerings by opening a third party marketplace alongside its retail business.

It is not surprising that once established, e-commerce marketplaces and e-commerce retail have taken off. Search costs have been cut dramatically by the advent of e-commerce: rather than physically move from store to store, aisle to aisle, and preparing in advance by asking friends and family for recommendations, all these and more are done with the click of a mouse in the convenience of one's home, wearing pajamas. Within minutes, a buyer, either a consumer or a business, can search across several marketplaces and retailers to

see what they sell, at what prices, and compare the delivery methods and any kinds of guarantees and return policies.

In this chapter, I will focus primarily on e-commerce marketplaces, where a platform brings buyers and sellers together but does not itself provide any goods, as opposed to e-commerce retailers, such as Walmart.com, the retail side of Amazon and many others.² Hence, I will focus on those marketplaces where a platform offers the necessary services so that two sides of a market transact in order to create gains from trade. Some of the key issues discussed in this chapter are: how platforms attract both sides to coordinate their behavior, and how they ought to consider participation fees; what kind of sales mechanisms can platforms offer for sellers; what roles do reputation mechanisms play in e-commerce marketplaces; how should platforms address the issue of how much a centralized role they ought to play to support trade; and how platforms can use their data to improve the operations of their marketplace.

2. A Brief History of Online Marketplaces

The quote at the beginning of this chapter by eBay's founder, Pierre Omidyar, describes the insight he had in 1995, when the internet was open to the general public and just a couple of years after the first web browsers, such as Netscape and Mosaic, were created. As he mentioned, there were bulletin boards systems, which in many ways were precursors to the current form of the World

² Chapter 19 by Smith and Zentner discusses other types of e-commerce in detail.

Wide Web, on which certain savvy professional buyers and sellers were already interacting. But beyond being restricted to dedicated and well-informed parties, these bulletin boards were quite primitive in their user support.

Initially, eBay offered one sales mechanism on its site: auctions. This was at the time one of the distinguishing features between the “consumer-to-consumer” (C2C) marketplace and other retailers. Today most of eBay’s listings are fixed-price listings, just like any retail site.³ The success of eBay.com was mind blowing and the company grew to become one of the early days Internet darlings, and over the years it expanded globally. In its early years it was almost exclusively a C2C platform where people sold used or otherwise unwanted items to others who valued them more. Over the years, as eBay grew, small and large businesses realized the huge exposure that eBay had, and the “business to consumer” (B2C) side of eBay grew to become the dominant form of transactions on the site today. As mentioned earlier, five years after eBay launched, Amazon leveraged its own website to put a marketplace alongside its established and growing retail business, which is primarily a B2C marketplace with little C2C action.

An interesting marketplace that was founded in 1996 is Craigslist, an online classifieds listing website, which unlike most marketplaces, is free to use for both sides of the market. The two exceptions are job ad listings and some rental categories, in which the site charges sellers for listing. Unlike other prominent marketplaces, the site has remained practically identical in its user

³ The choice of sales mechanisms is discussed in more detail in Section 4 below.

interface for the past two decades since it was founded, which puts it both technologically, and in terms of user interface, generations behind other marketplaces. However, because it is free, and because it had a huge early growth, it has acquired the network effects that are required to guarantee its huge success. Of course, success is in the eye of the beholder: by any measure, Craigslist is making a fraction of the earnings it can make. However, as a private company with no venture ownership, there is no shareholder pressure of any kind to divert it from the objectives of its founder, Craig Newmark, and of its CEO of the last 15 years, Jim Buckmaster. In an article published in 2009, Buckmaster was quoted saying “Companies looking to maximize revenue need to throw as many revenue-generating opportunities at users as they will tolerate,” and he continued, “we have absolutely no interest in doing that, which I think has been instrumental to the success of craigslist” (Wolf, 2009).⁴ There is no indication that anything has changed in the past six years since the article was published, making Craigslist a rather unique outlier in the world of successful online marketplaces.

In 1999 the Chinese business-to-business (B2B) platform Alibaba.com was founded with the goal of connecting Chinese manufacturers with overseas buyers, and its tremendous success resulted in a historic initial public offering in September 2014, which valued Alibaba at more than 230 billion U.S. dollars. A C2C Chinese marketplace, taobao.com, was founded by the Alibaba Group in

⁴ Wolf, Gary. “Why Craigslist Is Such A Mess”, Wired Magazine, 08/24/2009 (http://archive.wired.com/entertainment/theweb/magazine/17-09/ff_craigslist?currentPage=all)

2003, and very much followed the eBay model of allowing people to sell via auctions.

Online platforms are not only used in retailing; some platforms are also used in other sectors, including labor markets and markets for services. Aside from goods trading hands on platforms like eBay, Etsy, Craigslist and Taobao, online marketplace platforms have expanded to encompass services. Elance, founded in 1999, and oDesk, founded in 2003, were competing online platforms that allowed freelancers to bid on jobs for any type of work that can be done through the Internet, allow buyers, usually businesses, to describe jobs that freelance contractors can bid on in a “reverse” auction. A different kind of service platform was launched in 2008 called TaskRabbit, which is a marketplace that allows users to define small jobs and tasks that will be done by people their neighborhood.

Two other household name marketplaces, which are not quite retail per se, have launched in the late 2000s followed by tremendous success, often referred to as players in the “sharing economy”: Airbnb, founded in 2008, and Uber, founded in 2009. Airbnb allows people (sellers) to rent out lodging options for people looking for a place to stay (buyers), from whole apartments and homes, to a room, or even a couch in a small apartment for the very cash constraint travelers. Uber operates a mobile-app-based transportation marketplace that allows riders (buyers) to use their phones to request rides, like those fulfilled by taxis, from Uber’s network of private-car drivers (sellers).

The late 2000s have also seen the rise of what is known as “crowd-funding” sites such as Kickstarter, founded in 2009, and GoFundMe, founded in 2010. Kickstarter’s stated mission is to bring funding to people (sellers) who wish to launch some sort of creative project from music and films to video games, to new products. Other two-sided markets that are related to funding match lenders to borrowers, such as Prosper.com and Lendingclub.com.

The above is a small list of the most successful and well-known marketplaces. One point that is worth making is that these have developed over the span of almost 15 years, which begs the question: why so long? After all, the basics of a marketplace are the same across most of these platforms. There is a good reason that eBay was not founded before 1995: there was no Internet to speak of. In 1997, the first year the U.S. Census asked about Internet use, about 18% of households used the Internet. By 2001 that number was above 50%, almost tripling in 4 years, and by 2011 over 71% of U.S. households used the Internet (U.S. Census Bureau, 2013).⁵ All indications suggest that by the early 2000’s there was more than enough reason for most marketplaces to be in place. Why then did it take almost 15 years after eBay launched for Airbnb and Uber to be founded?

For Uber, the answer is more obvious: it is the kind of service-marketplace for which mobile smartphones with location technology play a critical role as ride services are needed on the go. The technology had not really proliferated before the late 2000s making the launch of Uber before 2009

⁵ See U.S. Census, “Computer and Internet Use in the United States (<http://www.census.gov/prod/2013pubs/p20-569.pdf>)

somewhat challenging due to a limited market. For Airbnb the answer is not that straightforward because another marketplace for short-term home rentals already existed worldwide: VRBO. Moreover, for the most part, people make their travel plans in advance, usually using the more comfortable large screen of a computer. Hence, it seems that the actual economics of marketplaces are less well understood than one might expect. Arguably, if rather than being motivated by the fact that “People were doing business with one another”, Pierre Omidyar, or the executives that later ran eBay, would have realized to true power of online two-sided marketplaces, sites like Airbnb, Elance-oDesk, and others may have all been subsidiaries of eBay.com.

3. Some Elementary Strategies of Online Marketplaces

In this section I discuss some of the key issues that online marketplaces must grapple with, starting with their inception and choice of target market, their fee setting strategies, and whether to consider other monetization channels such as advertising. Two additional strategic choices are not covered in this section. The first, which is covered in detail in Section 5, covers the use of reputation systems and other trust-promoting activities that a marketplace must engage in. The second is the form of payment intermediation, which can be done either using credit cards or other online payment intermediaries, and is not central to the issues addressed in this chapter.

3.1 *First Mover Advantage and Differentiation*

By their nature, online marketplaces are based on the most basic of network externalities: Buyers want to go where there are many goods for sale at competitive prices, and sellers want to go where there are many buyers searching for goods. As such, once a marketplace starts growing, it enjoys a positive feedback loop: more buyers and sellers attract more sellers and buyers, and the more it grows, the more attractive it is for both sides of the market. However, growing to scale is challenging because early-stage two-sided marketplaces struggle with the so-called “chicken-and-egg” problem. If, for example, there are few buyers, how will a striving marketplace convince sellers to sign on? And similarly, without an adequate amount of sellers, why should buyers browse its site?

Due to this virtual cycle, there is a tremendous *first-mover advantage* for a marketplace in a new sector. As an example, eBay’s success as a C2C marketplace in the U.S., U.K., Germany, and Australia, to name a few, is attributed to its being the first online marketplace in these countries. In 2000, eBay launched a version of its site in Japan just five months after Yahoo.com launched its own marketplace there, giving Yahoo the needed first-mover advantage. Within less than two years it was clear that eBay could not compete, and it closed its Japanese business. With only a five-month lead, Yahoo had the needed first mover advantage to build up enough of a presence to make it impossible for a second mover to come in successfully.

That said, there is at least one excellent example of a second mover who managed to build a thriving marketplace: Etsy.com, the leading U.S. marketplace for handmade arts and crafts. Etsy.com was founded in 2005, a full decade after eBay had established itself as the leading marketplace for practically any kind of good one can imagine. However, one thing that eBay failed to do while it grew rapidly in the 2000s was to invest in technology, and to create an online user experience that matched the progress of the Internet. First, eBay lagged in its ability to allow sellers to post pictures and it charged a fee for this service. Second, the site did not have an appealing esthetic looking webpage. Last but not least, its search engine technology was not on par with the advances made in this important area of e-commerce.

For most standard goods that are well defined, this may not have been too much of a problem. People knew what they are looking for, main keywords would bring up some relevant inventory, and esthetics was less important. But for unique hand made arts and crafts, eBay's shortcomings posed a serious problem for consumers. As eBay did not invest in the technology needed to tailor the user experience in different sub-markets, this opened a door for Etsy.com to create a well-differentiated niche marketplace, which quickly managed to attract small businesses specializing in arts and crafts, as well as avid shoppers for these kinds of items. Once these businesses and shoppers came on board, and the advantages of Etsy.com over eBay were made apparent, the niche players on eBay quickly moved to Etsy.com, which continued to grow rapidly. Its success is not only demonstrated by its dominance of the handmade arts and crafts market, but in the

fact that after ten years of successfully growing as a private company, during which time it raised almost \$100 million in venture backing, it chose in March of 2015 to start the process of listing itself on Nasdaq.

Aside from finding a new market for an online marketplace, or a niche that is not well served by existing marketplaces, there may be another strategy that an aspiring new marketplace can try to use: offer better terms and services than those that already exist. This would be possible if the costs to buyers and sellers of trying out several marketplaces are not high. This is often referred to as the possibility of *multi-homing*. That is, a seller can post an item for sale on several marketplaces, and buyers can browse the goods offered on several marketplaces. If these conditions prevail, then the first mover advantage may not be as ironclad as first suggested above.

Multi-homing cannot be easily done for C2C sellers because of the constraint of unique items. If I, for example, wish to sell a used toy that my now grown up children no longer use, then I may be exposed to a risk of “overselling” if I post it on two separate marketplaces. In the event that one buyer on marketplace *A* chooses to buy my toy, and before I am able to remove it from marketplace *B* a buyer in that marketplace also chooses to buy it, then I have a problem (assuming that I do not want to be penalized on one of the marketplaces for non-delivery). This will also apply to other sellers who have unique items such as antiques and other specialized goods.

However, for more generic, multi-item goods, typically sold by B2B or B2C sellers, the possibility of multi-homing is easier for buyers and sellers. It is

well known in the industry that many B2C sellers post their goods both on eBay's and on Amazon's marketplaces. For a new marketplace to penetrate successfully, however, it must offer some improvement over existing marketplaces, either through lower fees or through better quality of the platform for at least one side of the market.

For example, Amazon created its marketplace several years after eBay had successfully established itself. By the nature of its offerings, Amazon restricts sellers to post items that have well defined SKUs and exist as cataloged items. This made it very easy to post on Amazon compared to eBay at that time, helping Amazon third party get sellers on board. The fact that Amazon already had millions of buyers for its retail business made that relatively easy and profitable for the sellers. Still, for many of these sellers, there were enough buyers on eBay to allow for multi-homing.⁶

An interesting attempt at entering this arena is planned for 2015 by a start-up called Jet.com. The goal of Jet.com is to create a marketplace that, according to their claims, will be a membership program that will partner with retailers to save customers the most money possible while shopping online. Initial suggestions are that they will use algorithms to match sellers to buyers according to proximity in order to save on shipping costs and bundling of the orders. The approach Jet.com has taken in creating the buyer side of their market was unique: it ran a contest for people who can sign up the most potential users

⁶ One might wonder whether buyers are multi-homing or not. Some form of user specific tastes, like the logit error in a discrete choice model, may explain why some users do not multi-home while others do. Needless to say, there are enough people who visit either Amazon or eBay to make it worthwhile for many sellers to multi-home.

months before the site was planning to go live, thus building one side of the platform that will make sellers more willing to try it out (Adelman, 2015).⁷ In February of 2015 it had obtained \$140 million in a round of financing led by Bain Capital, suggesting that the market has some faith in its success, most likely fueled by the fact that a previously successful marketplace entrepreneur founded it. It will not be easy to compete with years of logistic optimization that Amzon.com, and more recently operations like Walmart.com, have invested in, and only time will tell whether they will succeed in their entry into this hard to enter space or whether they will suffer from lagging behind already established marketplaces.

3.2 *Setting Buyer and Seller Fees*

Because there are two sides to two-sided markets, two natural questions are first, who should pay for the platform's service, and second, whether it matters who does? This is the primary concern of many of the theoretical papers that offer formal models of two-sided markets (e.g., Armstrong, 2006; Rochet and Tirole, 2003, 2006; and Weyl, 2009). The central insight of the literature is that setting a price to the buyer side of a two-sided market depends both on the demand of buyers, and the costs that this demand imposes on the platform, as well as on how buyer participation impacts the participation of sellers, and the implied marginal profits from seller participation (and vice versa). Hence, unlike the Lerner rule mark-up formula of standard markets, in two-sided markets

⁷ See "Fortune awaits man who won shopping site competition," by Jacob Adelman, *The Inquirer*, March 16, 2015 (http://articles.philly.com/2015-03-16/business/60141897_1_eric-martin-ron-paul-amazon)

optimal pricing depends on the response elasticity of each side to the other side, as well as on same-side the demand elasticities.

The implication of these results is that in two-sided markets, the effect of one side's elasticity on that side's price is exacerbated by the reaction of the other side. As a consequence, the platform may choose one side's price to be below the marginal cost of serving that side, even resulting in negative prices. For example, a marketplace platform may choose to subsidize the sign up of the buyers, who maybe the elastic side of the market, in order to charge fees to sellers, who are more inelastic once a sizeable buyer side exists. Indeed, in the most popular e-commerce marketplaces, such as eBay, Amazon, Taobao, Alibaba and more, the buyers are charged no fees while all fees are levied on the sellers.⁸

Not all e-commerce marketplaces charge only the selling side of markets. For example, secondary events-ticket marketplaces such as Stubhub, TicketMaster and VividSeats, all charge both a buyer fee (also called a "service fee") and seller fee (also called a "commission"). In general, there is some "list price" that is set by the seller, and the platform fees are added onto this list price on in each side of the market. As the theoretical literature shows, it will sometimes be optimal to charge positive prices to both sides of the market. Whether or not event tickets satisfy the elasticities conditions that call for this approach id not clear, and there may be established norms in different marketplaces that are not explained by the current theories.

⁸ This is even more pronounced when platforms compete over the two sides of the market. See Armstrong (2006) for a formal analysis of these issues.

However, aside from the elasticity rules analyzed in the literature referred to earlier, other considerations may be important such as price obfuscation (See, e.g., Ellison and Ellison, 2009, and [Ellison's chapter in this Handbook](#)). If buyers are charged a usage fee that only appears after they select a good, then behavioral biases that depart from the standard rational choice models may prescribe the use of such price obfuscation strategies that will end up having buyers pay more than they would have were final prices including fees shown upfront.⁹ How and when to charge each side of a two-sided market beyond elasticity considerations is an interesting area for future research.

Another important issue in a two-sided is the strategic use of dynamic pricing. From the discussion above regarding the importance of growing a two-sided marketplace quickly, it is often critical to engage in *penetration pricing*, where early platform users are given preferable, maybe even negative prices, in order to build the scale needed to attract future participants. Hence, it will often, if not always, be optimal for aspiring two-sided marketplaces to take a hit to profits during their initial growth phases, in order to quickly reach the scale needed to create a strong barrier for entry, after which fees will be imposed as part of its monetization strategy.

3.3 *Advertising Revenues*

It is well known that many online properties receive revenues from advertising. This is the bread and butter of online search engines, such as Google and Bing, but interestingly, some two-sided marketplaces also sell advertising

⁹ See Coey et al. (2015) for an analysis of changes in price obfuscation on Stubhub.com.

space on their websites. For example, both eBay and Amazon have large and successful online advertising businesses, where advertisers can place ads on the marketplace sites, as do the Chinese marketplaces Alibaba and Taobao.

On one hand, advertising fees are often an attractive and tempting source of revenues. Also, it can be argued that by allowing sellers to pay for exposure, those sellers who pay may be the ones providing the best service, and hence using ads or other pay-to-play exposure strategies may increase the relevance of listing for buyers in the marketplace.

Two considerations need to be weighed against the perceived benefits of ads. First, ads that take the buyer off the marketplace, compared to those used to promote sellers on the marketplace's site, can cannibalize sales on the marketplace. Second, ads that are used to promote certain sellers compared to others may not necessarily promote the best possible sellers. Athey and Ellison (2011) develop a model where in equilibrium higher quality sellers are those who are willing to pay to be more prominently displayed with sponsored search ads. The intuition is that once a buyer clicks on an ad that takes him to the seller's site, he will only buy an item from that seller if it is relevant to the original search. Hence, more relevant sellers are more likely to convert potential buyers to actual buyers, and in turn, these sellers will be willing to pay more for prominence. However, this assumes that buyers will be able to infer relevance, and more importantly, not be exposed to adverse selection. McDevitt (2014) shows that in some circumstances it is the lower quality sellers who will pay

more for prominence because they do not have the alternative channel of having an established reputation that comes about from a record of past performance.

4. Auctions, Negotiations and Fixed Price Listings

As explained in the introduction, one of the key roles played by a marketplace is to match buyers with sellers in one convenient location, thus reducing an important market friction. Once the initial connection is made, two more activities must be performed to result in a successful transaction: a price needs to be determined and the trade has to be executed (the item shipped, the task performed, etc., and the money transferred). This section describes the variety of ways in which the second stage, price setting, is commonly determined. Most transactions are priced using one of three common mechanisms: posted prices that are typically non negotiable, which are employed by most retailers; some form of negotiation, common in business deals; and some form of an auction.

There is not much to be said about the use of fixed list prices. Profit maximizing sellers must form some expectation about demand and competition, and set their prices accordingly. Indeed, this is the primary price setting mechanism on retail platforms such as Amazon, Walmart.com, and others, including Amazon's marketplace.¹⁰ In what follows, more consideration is given to the two other forms of price settings: auctions and negotiations.

¹⁰ In recent years the share of auctions even on eBay's platform has been surpassed by fixed-price listings, many listed by businesses (Einav et al., 2013).

4.1 Auctions.

As mentioned earlier, when eBay was first launched, it had employed only auctions.¹¹ The benefits of using auctions to procure goods and services are well known and vigorously advocated by economists because competitive bidding results in low prices and, and for those who worry about procurement, auctions limit favoritism. Perhaps the most obvious benefit of auctions for online marketplaces that deal with C2C transactions is described using the first item that was sold on eBay: Pierre Omidyar's broken laser pointer. It was worth nothing to Pierre, but what it was potentially worth to some collector, or to someone who fixes and resells laser pointers, was a mystery. Hence, a benefit of auctions is that they are a *price discovery* mechanism. This is well suited for goods (or services) for which the actual value, or costs, is not known to the seller (or buyer).

What makes the use of auctions easy in online platforms is that the messiness of having an auctioneer and soliciting bids is trivially performed by the platform's software, which was one of the key innovations that eBay brought about. Moreover, eBay's auction mechanism is a dynamic form of the second-price auction introduced by Vickrey (1961). The way in which the auction is implemented with a mechanism referred to as "proxy bidding."

To illustrate this simple mechanism, imagine an item for sale in which the seller sets a starting bid of \$5 because he is unwilling to part with the item for

¹¹ Internet auctions appeared as early as 1993 on Internet news groups (see Lucking-Reiley, 2000).

less. Imagine that two bidders, 1 and 2, are interested in bidding, each considering the bids $B_1 = \$20$ and $B_2 = \$12$ respectively. The auction mechanism will have an “increment”, and let’s assume that for this auction it is \$0.25, together with an ending time, imagine some fixed point in time T . Now imagine that bidder 1 enters his bid of \$20 first. Because this bid is higher than the starting bid, it is “acceptable” and the auction will have a standing price of \$5.25, which is one increment above the seller’s starting price. Now imagine that bidder 2 puts in his bid of \$12. The auction will immediately increase the price to \$12.25, the second highest bid plus the increment. If no more bidders enter the auction until time T , then bidder 1 will win the auction at this price. If, say, a third bidder entered the auction and bid \$16.50, then the price will be adjusted to \$16.75, still keeping bidder 1, the high bidder, in the winning position. If, instead, the third bidder bid \$22.50, then the price will be adjusted to \$20.25 and this third bidder will assume the winning position.

As is well known, second price auctions are strategically simple because it is a dominant strategy to bid one’s valuation (this is true in private value settings, which are most likely relevant to many, though not all cases). Interestingly, eBay offers advice on how to bid that is consistent with the dominant nature of bidding one’s valuation in second price auctions (see Figure 1). A bidder who follows this advice should put in a bid once, when he or she arrives at the auction and decides to bid on the item. As a result, for a given auction, the arrival of bids should be relatively random relative to the time that the auction ends.

----- **Figure 1 approximately here** -----

However, there is a well-known phenomenon on eBay called “sniping,” in which bids, especially winning bids, are made in at the very last moments of the auction. Because eBay’s search algorithm has historically prioritized actions that end soonest it might be reasonable to assume that relatively more bidding activity might occur later in the auction. However, it is well documented that this is a strategic move by bidders, who wish to wait for the very last moment to place their bids.

Two early, and well-known, papers have analyzed this behavior on eBay. Bajari and Hortacsu (2003) suggest that bidding late is a best response of bidders when there is a common values component, as more information is revealed as the auction proceeds, and sophisticated bidders will wait in order not to reveal their own information. Roth and Ockenfels (2002) take a more “behavioral” turn and suggest that sniping is a best response to less sophisticated bidders. As Roth and Ockenfels explain, “inexperienced bidders might make an analogy with first-price “English” auctions, and be prepared to continually raise their bids to maintain their status as high bidder. In an eBay-style auction with a hard close, bidding very late might be a best response to ‘incremental bidding’ of this sort”

(p. 1094). Indeed, there is evidence that bidders do indeed engage in this incremental bidding practice, also known as “nibbling”.¹²

One of the issues that Roth and Ockenfels (2002) focus on is that sniping has no role if the auction has no predetermined end time. The auctions that were used by Amazon did not have a “hard” predetermined end time, but rather a “soft” ending time, which was automatically extended if a bid arrived within ten minutes of an auction’s planned end time, and was then extended for an additional ten minutes from the time of the latest bid. A soft ending time, of course, eliminates the rationale for last-minute bidding, which Roth and Ockenfels (2002) show, using eBay listings versus Amazon listings.¹³ Interestingly, Glover and Raviv (2012) show that items sold in auctions with soft ending times on Yahoo sell for substantially more than those with hard ending times, and conjecture that this is attributable to less sniping. I return to discuss some other issues related to the impact of sniping on buyer behavior.

Another aspect of auction design beyond the use of hard or soft stopping rules includes the use of other auction features. Two obvious ones are the use of minimum starting bids and of secret reserve prices. The use of a minimum price floor, which in practice is the auction’s observable starting price, is an obvious way for a seller to insure himself against a low level of competition between low

¹² Using internal eBay data, Backus et al. (2015a) show that bidders place multiple bids (revisions or bid responses) 29% of the time, with 4.4% of bidders placing 5 or more bids in a single auction. This suggests some reward to sniping because it denies 29 percent of bidders the opportunity to counter bid up to their true value, allowing snipers to obtain an item at a lower price than bidders would be willing to pay. Ely and Hossain (2009) conduct field experiments and find an economically insignificant payoff to sniping (about 17 cents for a new DVD), though one would expect new DVDs to have less variation in valuations, making the returns to sniping low.

¹³ See Ariely et al. (2005) for results from a lab experiment that mirror the field results shown in Roth and Ockenfels (2002).

value buyers. Hence, a rational seller that is not willing to part with an item without receiving at least some price p_0 , can set p_0 as his minimum price and guarantee that the item will not sell for less. In contrast to a minimum starting bid, bidders do not observe a secret reserve price. If a seller sets a secret reserve price of p_0 then as long as the highest bid is below p_0 , the only information potential bidders see is that the minimum reserve was not met.

Interestingly, Ku et al. (2006) argue that lower minimum bid floors may cause a form of “bidding fever” and increase the odds of a sale as well as the price conditional on sale. Data on Persian rug and digital camera auctions on eBay support their claims. Somewhat similarly, Simonsohn and Ariely (2008) found that while not necessarily increasing the price conditional on sale, lower bidding floors did increase the price conditional on it being at the upper tail of the price distribution. These results are not universal, though, as others showed evidence that lower bidding floors generally lead to lower prices conditional on sale, without testing the upper tail (see Kamins et al., 2004; Reiley, 2006; and Lucking-Reiley et al., 2007). Using large scale internal eBay data Einav et al. (2015) find that low start prices increase the (unconditional) probability of receiving a high final price, yet their results are more consistent with standard downward sloping demand and not with theories of “bidding fever.”

Turning to the use of reserve prices, Katkar and Reiley (2006) conduct a field experiment where they auctioned 100 Pokémon cards, equally split between a public reserve price of 30% of their value and a secret reserve of the same level (and no minimum bid). Secret reserve prices resulted in lower revenue. Einav et

al. (2015) use 403 matched listings, which are defined as listing that are equal in all observable aspects except for the use of a reserve price, and do not find much difference in auction outcomes between the public and secret reserve price options. Hence, offering sellers some sort of guarantee, like the use of secret or public (bidding floor) reserve prices is something that sellers find useful, and broadens the instruments that sellers can use when they choose the auction format.

An action design innovation that eBay introduced in 2000 is the option of having a “Buy-it-Now” (BIN) button. The BIN option allows the seller to specify a BIN price at which a buyer can end the auction instantaneously by clicking on the BIN button and thus committing to buy the item at the BIN price. The BIN button typically disappears once the item receives a qualified bid, after which the standard auction proceeds until the pre-specified end time.

Budish and Takeyama (2001) and Matthews (2003) consider the theoretical implications of a BIN price. It is obvious that for a given auction and a set of potential bidders that are drawn from a homogeneous distribution of private values, a very high BIN price will not affect revenue while a low BIN price will reduce expected revenue. That said, introducing heterogeneity in bidders along dimensions of risk aversion or impatience implies that BIN prices can increase expected revenue because some bidders may be willing to pay a premium to guarantee that they obtain the good.

Turning to empirical analyses, Standifird et al. (2004) found that buyers tended not to use the BIN option for silver dollar auctions even when the BIN

price was set low. Akerberg et al. (2006) found that sellers using a BIN option for Dell laptop auctions obtained \$29 higher revenues. Anderson et al. (2008) found, using sales of Palm handheld devices, that experienced sellers used the BIN option more often. In their data, BIN auctions resulted in slightly higher prices, not controlling for seller or item characteristics. Einav et al. use their matched listings approach and find that most BIN prices fall between 80% and 120% of an item's average posted price, with considerable variation in this range, and use this variation to identify the effect of BIN prices on auction outcomes. In their sample, items are unlikely to sell at high BIN prices. Setting a lower BIN price reduces seller revenue, while setting a high BIN price modestly increases revenue conditional on a sale. These results are consistent with the theories above in which a high BIN price can raise revenue from impatient or risk-averse buyers. In summary, the introduction of the BIN auction format seems to offer another instrument that seller find useful in some circumstances.

4.2 *Negotiations*

Auctions are not the only format of price discovery for sellers who are not certain of an item's true worth. Negotiations between buyers and sellers are another form of price discovery and price formation, though unlike auctions, they do not promote competition in a natural and easy way. However, they may play a role in helping savvy negotiators price discriminate between buyers with

different valuations, or they may allow sellers the opportunity to try and signal to the market that they are flexible.¹⁴

One might confidently assume that some form of negotiation happens between buyers and sellers in two-sided classifieds marketplaces such as Craigslist, where through a series of email exchanges, they find terms upon which a trade is consummated. Interestingly, in 2005 eBay introduced a sales mechanism that allows for negotiations while at the same time allowing buyers to just click on a BIN button at the listed price. This format, often referred to as “Best Offer,” is one where the seller sets a listed price, say p^* , and below the BIN button that guarantees the good at the listed price of p^* , there is a “Make Offer” button. Upon clicking the Make Offer button, a prospective buyer is prompted for an offer in a standalone numerical field. Submitting an offer triggers an email to the seller who then has 48 hours to accept, decline, or make a counter-offer. Once the seller responds, the buyer is sent an email prompting to accept and checkout, decline and leave, or make a counter-offer. This feature has been growing in popularity on eBay and negotiated transactions accounted for nearly nine percent of total transaction value in 2014.

Unlike auctions, there has been very little empirical work on bargaining and negotiations in general, and those happening in online marketplaces in particular. A recent exception by Backus, Blake and Tadelis (2015) shows that sellers seem to choose prices in ways that signal their bargaining strength or eagerness to complete a deal. In particular, they show that sellers who use

¹⁴ Negotiations are rare in offline retail settings, with the exception of new and used car markets; see Chapter 15 by Murry and Schneider for an extensive discussion of negotiation in that setting.

“round” numbers that are in multiples of \$100 will sell there items at a higher probability, sooner, and at lower prices than those items listed at more “precise” numbers.

Amazon launched a similar negotiations feature in its marketplace in late 2014, suggesting that there is indeed a demand from some sellers to use negotiations rather than setting a fixed price on Amazon’s marketplace. Given the growing popularity of this feature on eBay’s marketplace, and the fact that Amazon has also introduced this feature, it seems reasonable to assume that some form of negotiations will be supported on other platforms. Allowing for some sort of negotiated price setting may benefit non-retail marketplaces as well. My understanding is that on labor markets such as Elance-oDesk, some negotiations occur between buyers and freelance workers, often after a worker has won the auction. Sites like Airbnb may also benefit from introducing such a feature for sellers who may have last minute inventory, similar to the “name your own price” approach of the online travel retail site Priceline.com.

5. Reputations, Feedback and Quality Control

To some, the early success of eBay was a mystery. How is it that strangers who have never transacted with one another, and who may be thousands of miles apart, are willing to trust each other? Many have attributed this success to another brilliant mechanism introduced by eBay, and later copied in one form or another by other marketplaces: the use of some sort of reputation

mechanism (see, e.g., Resnick et al., 2000). As Dellarocas (2003) puts it, “eBay’s impressive commercial success seems to indicate that its feedback mechanism has succeeded in achieving its primary objective.” (p. 1411)

When buyers complete a transaction on eBay, they have 60 days to leave either a positive, negative, or neutral feedback score, or leave no feedback at all. About 65% percent of buyers leave feedback on eBay, a very high fraction, and an even higher fraction of more than 80% left feedback in eBay’s earlier days. Sellers can only leave positive feedback or no feedback, following a change in 2008 before which sellers too could leave negative and neutral feedback for buyers.

Whether reputation should be “two-sided,” like in eBay’s earlier days, or practically “one-sided” like it is now, is an interesting design question. In eBay’s earlier days, before practically all payments were made through Paypal’s online payment system, buyers would send checks or money orders to sellers. As such, just as sellers can renege or under-perform, so could buyers, making it imperative that sellers can choose whether or not to trust a buyer or wait till a check arrives and clears before sending the item to the buyer. However, after eBay acquired PayPal and adopted it as a necessary condition for sellers to sell their items, such problems of buyers not paying have all but disappeared.¹⁵

The question then is why not just keep the reputation system as two-sided? The answer lies in the problem of retaliation. Bolton et al. (2013) present

¹⁵ This is not completely accurate because buyers who win an auction must then go to PayPal to complete their payment, which they may not do, causing the auction to fail and the seller to have to either contact the next highest bidder and negotiate a sale or relist the item. By and large, however, this is a small problem.

data from eBay when the reputation system was two-sided and convincingly show that sellers wait to get feedback from buyers before giving feedback to buyers. In particular, they consider pairs of feedback scores, (F_B, F_S) left by pairs of buyers and sellers respectively in each transaction. For example, a transaction in which both buyer and seller left each other positive feedback is denoted $(+,+)$, while if the buyer left positive feedback and the seller negative feedback, it is denoted $(+,-)$. They first show that practically all transactions are either $(+,+)$ or $(-,-)$. They then show that a vast majority of $(-,-)$ transactions are characterized by the seller leaving feedback on the day or the day after the buyer does, while the $(+,+)$ transactions happen with less correlation between the buyer's and seller's day of leaving feedback.

This fear of retaliation was most likely a central cause behind the fact almost all buyers left positive feedback on eBay, which in turn caused eBay to switch from the two-sided reputation system to a one-sided reputation system. This is not, however, a good prescription for all online marketplaces. Take the non-retail marketplace Airbnb as an example. Even if payment is mediated by the site, as it is, there is still a concern that double moral hazard may occur. The sellers (owners) can misrepresent the home they are renting, leave it dirty, not give the buyers (renters) a key at the pre-specified time, and more. Similarly, the buyer can leave the home dirty, cause damage, be very noisy, etc. As such, it is imperative that Airbnb continue to keep a two-sided reputation system for trust to prevail in their marketplace.

Marketplaces can use the reputation that is left in many ways. First, it is common to present the reputation of each market participant for others to observe. eBay, for example, uses this information to provide several observable seller reputation measures. The first, *percent positive*, is defined as the seller's number of positive feedbacks divided by the sum of his number of positives, neutrals and negatives (restricted to feedback from the last 12 months). The second, *feedback score*, is a summed value of the number of positive feedbacks minus the number of negative feedbacks from the beginning of that participant's time on eBay. The third is a badge that certifies a seller as an "eBay Top Rated Seller" (ETRS). This designation is bestowed on sellers that meet a series of criteria believed by eBay to be an indication of a high quality seller.¹⁶

As mentioned earlier, eBay's success is largely attributed to the introduction of a reputation system, which was broadly adopted by other marketplaces. If the reputation system is doing its job then we should expect sellers with higher reputation to either have more sales, obtain higher prices, or both. Many papers have shown in one-way or another that sellers with higher reputation scores and more transactions receive higher prices for their products. Similarly, reputation seems to matter more for higher priced goods than for lower priced goods. See Bajari and Hortaçsu (2004) and Cabral and Hortaçsu (2010), as well as the papers cited by them, for more on these facts.

¹⁶ See Hui et al. (2014) for a lengthy discussion of eBay's ETRS program. For sellers eBay provides a set of seller ratings, called the "detailed seller ratings," which give buyers the opportunity to rate the seller at a finer-grained level. Less than one percent of buyers ever click on the page that contains this information.

These empirical findings bring some comfort, as they seem to align with what theory would predict, but a closer look at the literature shows that the magnitudes are small. An important question follows: how accurate are the reputation measures reflecting the variation in performance? As described earlier, retaliation on eBay may cause this user-generated feedback to be biased. A growing literature has shown that user-generated feedback mechanisms are often biased, and can be prone to influence by sellers. Dellarocas and Wood (2008) conjecture that the extremely high percent positive reputation measures on eBay is explained by the fact that many buyers who have poor experiences choose to leave no feedback at all. They proceed to suggest an econometric technique to uncover the true percent of positive transactions based on several assumptions, most notably that reputation is two-sided, which can no longer work after the change that eBay made in 2008.

Nosko and Tadelis (2015) argue that in practice, reputation systems in marketplace platforms may suffer from two problems. First, there is a reputational externality across sellers in that each seller does not internalize how his or her behavior impacts the likelihood that a buyer will return to the marketplace as a whole. Second, and similarly to Dellarocas and Wood (2008), reputation feedback can be---and in eBay's case is---biased. Using internal eBay data, Nosko and Tadelis (2015) establish that the percent positive measure has a mean of 99.3% and a median of 100%. They conjecture using anecdotal evidence, that it is more “expensive” to leave a negative review than it is to leave

a positive review because of seller retaliation and harassment.¹⁷ Hence, a central challenge is to construct a measure that more accurately reflects a seller's true quality. They suggest a new quality measure they call “effective percent positive” (EPP), which is calculated by dividing the number of positive feedback transactions by the *total* number of transaction, thus penalizing sellers who are associated with more transactions for which the buyers left no feedback. The EPP measure has a mean of 64%, a median of 67%, and exhibits significantly more variability than the percent positive measure because there is a lot of variation in the choice of buyer silence across sellers. They then use a “revealed preference” approach to study the effect of a seller's EPP on the buyer's propensity to continue buying on eBay after that transaction, which distinguishes their paper from most papers that collect scraped data from marketplaces and are limited to consider only prices and quantities. This approach allows them to get to the heart of the question of whether reputation mechanisms are indeed steering buyers away from low quality sellers.

Nosko and Tadelis (2015) show that a buyer who has a better experience on eBay (indicated by buying from a seller with a higher EPP score) will be more likely to continue to transact on eBay again in the future. They then report results from a controlled experiment on eBay that incorporated EPP into eBay's search-ranking algorithm. The treated group was a random sample of eBay buyers who, when searching for goods on eBay, were shown a list of products

¹⁷ As Nosko and Tadelis (2015) show, some sellers have even sued buyers for leaving negative feedback, and it is known that many harass buyers with a slew of emails following a negative feedback. Horton and Golden (2015) show similar biases in public feedback on an online labor marketplace, and show that for private feedback that is not revealed to the sellers the bias is a lot less severe.

that promoted seller EPP compared to a control group in which this was not done. The results show that treated buyers who were exposed to higher EPP sellers were significantly more likely to return and purchase again on eBay compared to the control group of buyers. Hence, their paper argues that marketplace platforms can benefit from using data in ways that uncover better measures of seller quality.¹⁸ Furthermore, rather than let buyers select sellers based on measures of reputation, which may or may not be correctly interpreted, marketplace platforms can use their search and presentation algorithm to promote better quality sellers for the continued health of the marketplace.¹⁹

Bias in reputations is not unique to the eBay marketplace. Mayzlin et al. (2014) show interesting biases in ratings for hotels from the online travel sites. What makes their paper particularly clever is that they do not attempt to categorize which reviews are fake reviews versus those that are not, which on the face of it is impossible because fake reviews are designed to mimic real reviews. Instead they take advantage of a key difference in website rating systems where some websites accept reviews from anyone while others require that reviews be posted by consumers who have purchased a room through the website. If posting a review requires an actual purchase, the cost of a fake review is much higher. The upshot is then that they measure the differences in the distribution of reviews for a given hotel between a website where faking is expensive and a website where faking a review is cheap. Their results indeed show greater

¹⁸ Hwang in Chapter 21 describes in detail the use of “big data,” particularly combining different types of data, to better understand consumer behavior in online settings.

¹⁹ See Masterov et al. (2015) for a similar approach using the content of messages from buyers to sellers as a measure of seller quality.

bunching at the extreme ratings for hotels on the sites where posting reviews is cheaper, and this is exacerbated by local competition (more local hotels). Hence, for reviews to be less biased it is critical to impose some kind of cost to prohibit fake reviews by non-purchasers.

Fradkin et al. 2014 study the bias in online reviews by using internal data Airbnb, and like Nosko and Tadelis (2015) report results from field experiments conducted by the online marketplace. In one experiment they offer users a coupon to leave feedback and show the users who were induced to leave feedback report more negative experiences than reviewers in the control group, suggesting that otherwise they would have probably been silent. In a second experiment they disable retaliation in reviews, similar to what eBay did in 2008, and find that retaliation (or rewards for positive feedback) cause a bias, but that the magnitude of this bias is smaller than that caused by a lack of incentives to leave truthful feedback. Interestingly, using data on social interactions between buyers and sellers on the site, they show that such interactions result in less negative reviews. This result suggests that a challenge for online marketplaces is the potential loss of information following the social interaction of buyers and sellers on the site.

6. Behavioral Biases in Marketplaces

One of the big advances that online marketplaces have brought forth for economists is a huge playground and laboratory to test a variety of theories about

buyer and seller behavior. Starting with the pioneering work of Lucking-Reiley (1999, 2000), many economists have either run field experiments in online marketplaces (primarily on eBay) or have scraped data from online marketplaces to explore a variety of questions regarding marketplace behavior. A theme in this literature is that many behavioral biases seem to pop up that depart from the standard rational actor models.

One example is the non-equivalence between shipping fees and prices. Hossain and Morgan (2006) and Brown, Hossain and Morgan (2010) show that consumers appear to underweight shipping fees relative to basic item prices. This implies a form of price shrouding in that a dollar of shipping costs is “felt less” than a dollar in the price of the item, consistent with some form of bounded rationality or mental accounting. Einav et al. (2015) replicate this finding in a large set of matched items from eBay’s internal data.

A second example is that people get carried away in bidding compared to what they may have paid otherwise. In Section 4.1 the phenomenon of “bidding fever” was outlined, as described and analyzed by Ku et al. (2006) and Simonsohn and Ariely (2008), though the results were not replicated in the large scale analysis of Einav et al. (2015).

In a similar vein, Malmendier and Lee (2011) show that auction prices for a certain well-defined item sometimes rise quite a bit above the posted prices of the same item that is not sold in auction. These results relate to a larger literature on price dispersion and consumer search in online markets. The basic argument is that low search costs on the Internet should all but eliminate price

dispersion online, yet many papers find substantial price dispersion (see, e.g., Bailey, 1998; Brynjolfsson and Smith, 2001; Baye, Morgan and Scholten, 2004; and Ellison and Ellison, 2009). Malmendier and Lee's (2011) results can also be interpreted as a failure of consumer search. In particular, using eBay data from 2004 in which two distinct sellers were selling a particular board game for \$129.95, while other sellers listed the game in auctions. They show that auction prices exceeded the \$129.95 posted price more than 40 percent of the time, often by \$10 or more. Interestingly, like the bidding fever findings, this finding is also not confirmed in the large-scale data analysis of Einav et al. (2015), which may be due to the fact that the irrational behavior was weeded out over time.

Using internal eBay data, Backus et al. (2015a, 2015b) use a revealed preference approach of buyers returning to eBay to test other behavioral aspects of buyer behavior. Recall from Section 4.1 that eBay's variant on a second price auction implies that, at least for the case of private values, bidders should just bid their valuation whenever they happen to come to the site, yet there is a significant amount of sniping, which is rationalized by the presence of nibblers. The question that Backus et al. (2015a) ask is whether being sniped is perceived as a negative experience, causing new buyers who do not understand how to bid to leave eBay after being sniped. The data confirms this hypothesis by showing that losers who are sniped and do not have a chance to put in another bid are more likely to leave the site and cease their engagement with eBay as compared to bidders who lost with enough time to respond. This finding suggests that

bidders who get sniped suffer from some kind of misunderstanding about how the auction works and how they ought to behave.

Backus et al. (2015b) take advantage of the BIN option in auctions and ask a related question. After establishing that bidders who lose suddenly are more likely to exit eBay, one question is whether they develop some kind of attachment to the good by merely being the high bidder for some time. They find that conditional on losing the auction because some other user clicked the BIN button, bidders who held the leading bid longer are more likely to exit eBay and not return (or return but not to an auction listing) than those who were the leading bidders for shorter periods of time. What is interesting is that unlike the endowment effect, a well-studied consequence of attachment bias, bidders who were leading the auction longer and lost *do not* return to bid on similar items with a higher willingness to pay. That is, the virtual attachment does not cause a virtual endowment effect and instead is perceived as a more painful loss, consistent with loss aversion.

7. Concluding Remarks

This chapter demonstrates the growing popularity of online two-sided marketplaces. Starting with eBay in 1995, the use of online two-sided markets has proliferated at a dizzying pace, and the academic research was quick to follow. Of the 50 articles cited in the chapter, 35 were published between 2005

and 2015, and I am confident that there are 35 more that I did not cite due to space limitations (and apologies to the authors of those fine papers).

I was approached in 2011 by eBay to join the Research Labs, and without further need for much convincing. As I told Mark Carges, the Chief Technology Officer of eBay at the time, for an economist to work at eBay and have access to the vast marketplace data is like for a kid to go to Disneyland. As this chapter hopefully demonstrates, economists using marketplace data have been able to produce interesting and impactful research over the past 15 years that helps us better understand how marketplaces work and how people behave in them.

A myriad of methods have been used, from scraping observational data, to running field experiments, to, more recently, having economists within the companies analyze observational data as well as run large scale field experiments. As online marketplaces continue to grow, and as more economists get involved with them, I am sure that the insights will continue to flow with a new set of exciting empirical work that will inspire new theoretical work that can be then taken to the data.

References

- Ackerberg, Daniel, Keisuke Hirano, and Quazi Shahriar (2006) “The Buy-it-now Option, Risk Aversion, and Impatience in an Empirical Model of eBay Bidding,” working paper, University of Michigan.
- Adelman, Jacob (2015) “Fortune awaits man who won shopping site competition,” *The Inquirer*, March 16, 2015
- Anderson, Steven, Daniel Friedman, Garrett Milam, and Nirvikar Singh (2008) “Buy It Now: A Hybrid Internet Market Institution,” *Journal of Electronic Commerce Research*, 9(2): 137-153.
- Armstrong, Mark (2006) “Competition in Two-Sided Markets.” *The RAND Journal of Economics*, 37(3): 668–91.
- Ariely, Dan, Axel Ockenfels, and Alvin E. Roth (2005) “An Experimental Analysis of Ending Rules in Internet Auctions,” *Rand Journal of Economics*, 36(4): 891-908.
- Athey, Susan and Glenn Ellison (2011) “Position Auctions with Consumer Search,” *Quarterly Journal of Economics*, 126(3): 1213-1270.
- Backus, Matthew, Tom Blake, Dimitriy Masterov and Steven Tadelis (2015a) “Is Sniping A Problem For Online Auction Markets?” forthcoming, *Proceedings of the 24th ACM International World Wide Web Conference (WWW24)*
- Backus, Matthew, Tom Blake, Dimitriy Masterov and Steven Tadelis (2015b) “Virtual Attachment, Loss Aversion and Exit in an Online Auctions,” NBER working paper no. wXXXXXX
- Backus, Matthew, Tom Blake, and Steven Tadelis (2015) “Cheap Talk, Round Numbers, and the Economics of Negotiation,” NBER working paper no. wXXXXXX
- Bailey, Joseph (1998) “Intermediation and Electronic Markets: Aggregation and Pricing in Internet Commerce,” Ph.D. Thesis, Massachusetts Institute of Technology.

- Bajari, Patrick and Ali Hortaçsu (2003) “Winner’s Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions,” *Rand Journal of Economics* **2(1)**: 329-355.
- Bajari, Patrick and Ali Hortaçsu (2004) “Economic Insights from Internet Auctions,” *Journal of Economic Literature*, **42(2)**: 457-486.
- Baye, Michael R., John Morgan, and Patrick Scholten (2004) “Price Dispersion in the Small and the Large: Evidence from an Internet Price Comparison Site,” *Journal of Industrial Economics*, **52(4)**: 463-496.
- Bolton, Gary, Ben Greiner and Axel Ockenfels (2013) “Engineering trust: reciprocity in the production of reputation information,” *Management Science*, **59(2)**: 265–285.
- Brown, Jennifer, Tanjim Hossain, and John Morgan (2010) “Shrouded Attributes and Information Suppression: Evidence from the Field,” *Quarterly Journal of Economics*, **125(2)**: 859-876.
- Brynjolfsson, Erik, and Michael Smith (2001) “Consumer Decision-making at an Internet Shopbot: Brand Still Matters,” *Journal of Industrial Economics*, **49(4)**: 541-558.
- Budish, Eric and Lisa Takeyama (2001) “Buy Prices in Online Auctions: Irrationality on the Internet?” *Economics Letters*, **72(3)**: 325-333.
- Cabral, Luis and Ali Hortaçsu (2010) “The dynamics of seller reputation: Evidence from eBay,” *The Journal of Industrial Economics*, **58(1)**: 54–78.
- Coey, Dominic, Kane Sweeney and Steven Tadelis (2015) “,”
- Dellarocas, Chrysanthos (2003) “The digitization of word of mouth: Promise and challenges of online feedback mechanisms,” *Management science*, **49(10)**: 1407–1424.
- Dellarocas, Chrysanthos and Charles A. Wood (2008) “The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias,” *Management Science*, **54(3)**: 460–476.

- Einav, Liran, Chiara Farronato, Jonathan Levin, and Neel Sundaresan (2013) “Sales Mechanisms in Online Markets: What Happened to Internet Auctions?” NBER Working Paper no. w19021.
- Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan (2015) “Assessing Sale Strategies in Online Markets using Matched Listings,” *American Economic Journals: Microeconomics*, forthcoming.
- Ellison, Glenn, and Sara Fisher Ellison (2009) “Search, Obfuscation and Price Elasticities on the Internet,” *Econometrica*, **77(20)**: 427–52.
- Ely, Jeffrey C. and Tanjim Hossain (2009) “Sniping and squatting in auction markets,” *American Economic Journal: Microeconomics*, **1(2)**: 68-94.
- Fradkin, Andrey, Elena Grewal, Dave Holtz, and Matthew Pearson (2014) “Bias and Reciprocity in Online Reviews: Evidence from Experiments on Airbnb,” working paper.
- Glover, Brent and Yaron Raviv (2012) “Revenue nonequivalence between auctions with soft and hard closing mechanisms: New evidence from Yahoo!,” *Journal of Economic Behavior & Organization*, **81(1)**: 129-136.
- Greif, Avner (2006) *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press, London, U.K.
- Horton, John J. and Joseph M. Golden (2015) “Reputation Inflation: Evidence from an Online Labor Market,” working paper, New York University.
- Hossain, Tanjim, and John Morgan (2006) “...Plus Shipping and Handling: Revenue (Non)Equivalence in Field Experiments on eBay,” *Advances in Economic Analysis and Policy* 6(2), Article 3.
- Hui, Xiang, Maryam Saeedi, Zeqian Shen and Neel Sundaresan (2014) “Reputation & Regulations: Evidence from eBay,” Working paper, Ohio State University.
- Kamins, Michael A., Xavier Dreze, and Valerie S. Folkes (2004) “A Field Study of the Effects of Minimum and Reserve Prices on Internet Auction,” *Journal of Consumer Research*, **30(4)**: 622-628.

- Katkar, Rama, and David H. Reiley (2006) “Public versus Secret Reserve Prices in eBay Auctions: Results from a Pokémon Field Experiment,” *Advances in Economic Analysis and Policy*, **6(2)**, Article 7.
- Ku, Gillian, Adam D. Galinsky, and J. Keith Murnighan (2006) “Starting Low but Ending High: A Reversal of the Anchoring Effect in Auctions,” *Journal of Personality and Social Psychology*, **90(6)**: 975-986.
- Lucking-Reiley, David (1999) “Using field experiments to test equivalence between auction formats: Magic on the Internet,” *American Economic Review* **89(5)**: 1063–1080.
- Lucking-Reiley, David (2000) “Auctions on the Internet: What’s being auctioned, and how?” *Journal of Industrial Economics*, **48(3)**: 227–252.
- Lucking-Reiley, David, Doug Bryan, Naghi Prasad, and Daniel Reeves (2007) “Pennies from eBay: the Determinants of Price in Online Auctions,” *Journal of Industrial Economics*, **55(2)**: 223-233.
- Malmendier, Ulrike and Young Han Lee (2011) “The Bidder’s Curse,” *American Economic Review*, **101(3)**: 749-787.
- Masterov, Dimitriy, Steven Tadelis and Uwe Mayer (2014) “Canary in the e-commerce coal mine: Detecting and predicting poor experiences using post-transaction buyer-to-seller messages,” NBER working paper no. wXXXXX
- Mayzlin, Dina, Yaniv Dover and Judy Chevalier (2014) “Promotional reviews: An empirical investigation of online review manipulation,” *American Economic Review*, **104(8)**: 2421–55.
- McDevitt, Ryan C. (2014) ““A” Business by Any Other Name: Firm Name Choice as a Signal of Firm Quality,” *Journal of Political Economy*, **122(4)**: 909-994.
- Milgrom, Paul R. , Douglass C. North and Barry R. Weingast (1990) “The role of institutions in the revival of trade: The law merchant, private judges, and the Champagne fairs,” *Economics & Politics*, **2(1)**: 1-23.

- Nosko, Chris and Steven Tadelis (2015) “The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment,” NBER working paper no. w20830
- Reiley, David H. (2006) “Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet,” *Rand Journal of Economics*, **37(1)**: 195-211.
- Resnick, Paul, Ko Kuwabara, Richard Zeckhauser and Eric Friedman (2000) “Reputation systems,” *Communications of the ACM* **43**, **12**:45–48.
- Rochet, Jean-Charles, and Jean Tirole (2003) “Platform Competition in Two-Sided Markets.” *Journal of the European Economic Association*, **1(4)**: 990–1029.
- Rochet, Jean-Charles, and Jean Tirole (2006) “Two-Sided Markets: A Progress Report.” *The RAND Journal of Economics*, **35(3)**: 645–67.
- Roth, Alvin E. and Axel Ockenfels (2002) “Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet,” *American Economic Review*, **92(4)**: 1093-1103.
- Rysman, Marc (2009) “The Economics of Two-Sided Markets,” *Journal of Economic Perspectives*, **23(3)**: 125–143.
- Simonsohn, Uri, and Dan Ariely (2008) “When Rational Sellers Face Non-Rational Consumers: Evidence from Herding on eBay,” *Management Science*, **54(9)**: 1624-1637
- Standifird, Stephen S., Matthew R. Roelofs, and Yvonne Durham (2004) “The Impact of eBay’s Buy-It-Now Function on Bidder Behavior,” *International Journal of Electronic Commerce* **9(2)**: 167-176.
- U.S. Census Bureau (2013) “Computer and Internet Use in the United States,” Thom File, Issued May 2013, P20-569, U.S. Department of Commerce, Economics and Statistics Administration
- Vickrey, William (1961) “Counterspeculation, Auctions and Competitive Sealed Tenders,” *Journal of Finance*, **16(1)**: 8-37.

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

Wolf, Gary (2009) "Why Craigslist Is Such A Mess," *Wired Magazine*, 08/24/20

Here's how to bid on an item:

1. Carefully review the listing.
2. Click the Place Bid button. You may need to sign in.
3. Enter your maximum bid, and then click the Continue button.
4. Review your bid, and then click the Confirm Bid button

Notes: This appeared under the heading “How to Bid” on eBay’s Bidding Overview page on March 31, 2015. See <http://pages.ebay.com/help/buy/bidding-overview.html>

Figure 1: Bidding Instructions on eBay