BARGAINING VS. POSTED PRICES: AN ANALYSIS USING THE EBAY AUTOMOBILE MARKET

By

KRITEE GUJRAL

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2016

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

BARGAINING VS. POSTED PRICES: AN ANALYSIS USING THE EBAY AUTOMOBILE MARKET

By

Kritee Gujral

December 2016

Chair: Steven Slutsky
Major: Economics

While bargaining has been extensively studied in theoretical literature, there have been few empirical studies of it. There exists an option on eBay in which sellers can list items with fixed posted prices and allow potential buyers to make their best offers. These Best Offer listings are different from standard auction listings, and allow for back-and-forth negotiation between a seller and his potential buyers. A newly emerging strand of literature uses eBay Best Offer listings for studying bargaining empirically. Chen et al. (2016) studied the use of Best Offer listings (vs. regular fixed price listings) on eBay motors, using data that they collected from 2008-2009. They concluded that sellers with higher bargaining cost are more likely to adopt the posted-price format, that bargaining leads to a higher sale rate than the posted-price format, but there is no difference in the transaction price. In my dissertation, I extend the study of bargaining vs. fixed prices on eBay motors by using new data collected in 2013 and 2016, additional variables, and different estimation techniques. In Chapter 1, I review relevant previous literature that provides theoretical insights into how bargaining through Best Offer might operate and describe the data that I have collected.

In Chapter 2, I employ a reduced form approach (using a fractional logit model) to determine factors that incline a seller to adopt bargaining instead of posted prices. I focus on two factors, seller patience and competition. The results indicate that sellers with weaker bargaining power (i.e. lower patience, fewer outside options, and/or higher seller competition) are more likely to adopt bargaining. Using 2013 data, I find that the effect

of sellers' individual-level patience varies with the level of competition in a seller's external environment. At high levels of competition, a patient seller chooses to bargain, while at low levels of competition, a patient seller chooses not to bargain. Seller patience, a form of higher bargaining power, is weakened substantially by higher seller competition (more outside options for buyers), causing even the more patient sellers to choose bargaining over fixed prices. This result can have strong implications for several other cases in economics in which individual-level patience is dominated by external competitive pressures, and therefore deserves further attention. Note that using different variables for bargaining power in this paper (i.e. patience) than in Chen et al. (2016), I corroborate their findings that sellers with lower bargaining power are more likely to use Best Offer (consistent with Bester (1993)). In addition, I find that sellers with lower experience, lower item-specific learning, and lower information in general are likely to choose bargaining over posted prices. Nevertheless, bargaining does not seem to decline over time on eBay Motors. This suggests that bargaining may be more likely to persist in markets that regularly witness new sellers/entrants and/or markets of used goods.

In Chapter 3, I discuss two possible channels through which bargaining might impact consumer surplus. On the one hand, sellers might use bargaining as a means for price discriminating, thereby improving consumer surplus if sellers negotiate downward from their existing high price/markup. On the other hand, sellers might raise their prices or charge a higher markup solely because of their choice to bargain. This channel implies a negative impact of the bargaining selling mechanism on consumers. While Chen et al. (2016) also implicitly test this latter channel of impact, they do so using an OLS approach. I address the simultaneity of a seller's choice to bargain and his choice of markup by using an instrumental variable approach. The results with this approach suggest that using the OLS approach may produce misleading results. I find no evidence for the former channels positive impact on consumers, but I rule out the negative impact of the latter-mentioned channel. That is, contrary to the results in Chen et al. (2016), I find that sellers who choose to allow bargaining do not charge higher markups than fixed-price sellers. Surprisingly, sellers that choose to bargain charge lower markups than fixed-price sellers. Finally, I discuss the quality uncertainty, or the "lemons problem" of the automobile market (Akerlof, 1995). Although the data does not allow for assessing the quality of the cars directly, some equilibria for sellers' behavior, in their use of the four primary signals due to a combination of markup and Best Offer, are discussed. The empirical evidence then helps rule out the equilibria in which eBay sellers might be trying to con consumers through their use of bargaining, suggesting that bargaining on eBay Motors is not impacting consumer surplus negatively through the suspected channels.

In Chapter 4, I examine the impact of a seller's choice of markup, choice of best offer, and the interaction of these two choices on the seller's profitability (measured by a car's likelihood of selling, wait time until it sells, and its final transaction price). The results emphasize that studying the impact of using best offer without considering its interaction with pricing/markup can be misleading. I find that using best offer increases the selling likelihood for cars with a sufficiently high markup (at least 10.8%) but decreases the selling likelihood for cars with lower markup. This can occur if the markup on a car serves as a signal for car quality. . In addition, contrary to Chen et al. (2016), best offer also impacts the transaction price of a car (directly and indirectly). The direct impact of best offer is that it increases the transaction price for cars with markups higher than 13%, and decreases the transaction price otherwise. Furthermore, a higher markup reduces the likelihood of selling (by 62.7% for fixed-price listings and by 33.7% for Best Offer listings). It might then seem to be an unprofitable strategy for sellers to adopt a higher markup, but a higher markup also yields a higher transaction price by 40.3%. This indicates the possibility of an important indirect impact of using best offer on the transaction price (not evident in the approach used in Chen et al. (2016)): using best offer reduces the negative impact of a high markup on sales, by 29% (62.7 33.7%). That is, combining a high markup with the use of best offer allows sellers to reap the benefits of a high markup on transaction price, while mitigating its harmful impact on sales.

CHAPTER 1
INTRODUCTION

1.1 Dissertation Introduction

There have been few empirical tests of how bargaining and negotiation are carried out. That is not because of a lack of theoretical models to test. Such models have been around for years and have yielded various insights that remain unconfirmed empirically. Theoretical literature therefore needs its empirical counterpart to help confirm its insights, and to also highlight which theories are more relevant and applicable in a given context. While some theoretical models attribute bargaining power to the relative patience of an agent (Rubinstein and Wolinsky, 1985; Bing, 2009), others suggest that higher information and asymmetric information in an agent's favor are the source of bargaining power (Coffman, 1991). Some scholars suggest that a seller's bargaining power depends on the true valuation of the other agent (Chatterjee and Samuelson, 1987) in a bargaining game. Still, others suggest that a seller's monopoly power gives the seller bargaining power (Wang, 1995). It is also said that a higher number of outside options can lead to higher bargaining power, and that the asymmetry of information regarding outside options can also determine which agent will have higher bargaining power(Cunyat, 1998). Apart from the above-mentioned models that deal with exogenous sources of bargaining power, there also exist theoretical models that model the *process* of negotiation. These models point to endogenous factors (particular behaviors and strategies employed by agents) during the bargaining process as sources of bargaining power. These endogenous factors include using time delay (Rubinstein and Wolinsky, 1985; Perry, 1986; Cramton, 1992), using information about the outside options (Gantner, 2008), and the threat of breakdown or opting out (Rubinstein and Wolinsky, 1985). Theoretical models thus allude to a wide range of possible sources of bargaining power, and empirical work offers a means for testing which models are most relevant in practice.

For instance, the existing empirical studies and experimental studies (using ultimatum games) confirm some of the previous insights, but also raise some interesting questions.

Evidence from the hog market has confirmed that higher information and higher patience are key sources of bargaining strength. However, insights can sometimes vary from theory to practice as well (Vukina et al., 2010). For instance, "cheap talk, where players tell costless and unverifiable lies during the negotiation process (said to be important in real bargaining settings) should not impact ultimate bargaining outcomes according to theory. Nevertheless, an experiment which examined the impact of cheap talk in an ultimatum bargaining setting with two-sided imperfect information showed that lies about private information and (incredible) threats of future actions do influence bargaining outcomes (offers and responses) in both the short-term and long-term (Croson et al., 2003), at least for relatively inexperienced agents. Such contradicting conclusions regarding which factors are important for successful bargaining emphasize the need for more empirical answers to age-old questions regarding bargaining.

However, studying bargaining empirically is not so straightforward, and perhaps that is why there are fewer empirical studies. Given that bargaining has typically been popular in informal markets, such as in a bazaar or the store of a developing country, (Bing, 2009), data collection is a bit complicated. Written records or available data of transactions, prices, agents' behavior is scarcely available. Surveys are also infrequent, perhaps because a bargaining scenario is complex (with many variables as theoretical possibilities for impacting the ultimate bargaining outcome, as alluded to earlier), obscuring the identification of key variables for which data is needed. Furthermore, surveys are even harder to design and implement (and are less reliable) for cases in which agents have incentives to hide information (such as their true reservation values). The problem is exacerbated if the development of countries and of informal markets can indeed wipe out bargaining markets, which have previously been referred to as a "primitive" form of markets. Hence, there might also be a pressing time urgency in understanding bargaining empirically; lest the opportunity to test ages of accumulated theoretical wisdom disappears if these markets do become obsolete.

The advent of bargaining platforms emerging on the internet offers some relief for these concerns. The data collection problem is simplified to some extent and questions regarding bargaining markets have also become immediately more pressing and more relevant. eBay's introduction of the "Best Offer" feature not only inspires the question of why bargaining is re-emerging on the e-commerce platform, but it also allows for studying the question empirically. The "Best Offer" listing format on eBay mimics a sequential-move alternating offers bargaining game. That is, sellers and buyers can negotiate back and forth with each other (as they would generally do in an in-person negotiation) by entering offer amounts online. There are only a handful of studies to date that have exploited the opportunity for studying bargaining on eBay using eBay Best Offer.

Therefore, this dissertation examines the determinants and the outcomes of bargaining (Best Offer feature) on eBay, by studying the sellers' choice of Bargaining vs. Posted Prices, that is, sellers' choice between using Best Offer listings vs. fixed price listings. Although the eBay market structure varies substantially from the relevant theoretical models (to be discussed later), this paper is able to test and confirm insights of theoretical literature. There is only one known working paper that has attempted to study the question of sellers' choices of bargaining vs. posted prices using eBay Best Offer. This dissertation complements and extends this previous work on eBay bargaining, and helps fill the gap in the rich theoretical bargaining literature that has long awaited its empirical complement.

In addition, this dissertation contributes to a few other strands of literature. It contributes to literature that studies the structure of endogenously occurred markets (McAfee, 1993). It complements and extends previous "eBay Best Offer" papers (Huang et al., 2013) (Toklu, 2014). That is, while the dissertation focuses on eBay Motors as these previous papers do, it uses differing variables and methodologies, along with higher product (automobile model/make) heterogeneity than previous papers. This offers a richer discourse to this extremely newly emerging literature on "eBay Best Offer." Furthermore, this dissertation is the first known empirical study that compares Best Offer usage across time. Two cross-sections of data (from 2013 and from 2016) are studied for the same set of car models. This allows for an intermediate long-term study of best offer usage and for capturing the impact and dynamics of seller learning over time.

Finally, this dissertation also offers insight into a specific type of bargaining invoked by the "or Best Offer" phrase. Despite off-eBay usage of the phrase "or Best Offer" being traced as far back as the year 1738 in forum classifieds and its recent popularity on Craiglist.com, there is almost no secondary literature on the phrase "or Best Offer." To the extent that "Best Offer" induces a unique type of bargaining, the results of this paper may not extend to all bargaining markets. Instead, the findings of the paper contribute to an understanding of this specific type of bargaining engendered by the seller's indication that he is willing to accept offers by the buyers. eBay allows for an easier platform in which this specific type of bargaining, which was prevalent historically and has re-emerged in online markets, can now be studied.

1.2 Dissertation Overview

The dissertation is comprised of three essays studying a seller's choice between using bargaining or using posted prices, where Best Offer fixed-price eBay listings represent the bargaining mechanism and non-Best Offer fixed-price listings represent the posted price mechanism. An original dataset including a cross-section of data from the year 2013 and a cross-section from the year 2016 is used. There are 692 cars from the year 2013 and 5584 cars from the year 2016. This translates to 428 sellers from 2013 and 1417 sellers from 2016.

Chapter 1 motivates and introduces the dissertation. It describes the eBay market structure and reviews previous literature on eBay Best Offer and on Bargaining vs. Posted Prices. Data collection and variables are also described in this chapter, and finally, stylistic facts about the eBay market are presented. Then, Chapter 2 focuses on how seller patience and market competition affect the seller's choice of bargaining vs posted prices. Results for 2013 and 2016 are compared for studying seller learning and evolution regarding the adoption of bargaining as a selling mechanism.In this chapter, I use a reduced form approach to study the dependent variable: *BOperlisting*, which is the fraction of seller's listings that are listed as

Best Offer. I do not include the markup or price variable on the right-hand side, because of possible simultaneity/endogeneity between the seller's choice of price and choice of using BO. This price endogeneity is addressed in Chapter 2.

Chapter 3 studies the possible simultaneity between the seller's pricing decision and his choice between Best Offer and Posted Prices. An IV approach is used to measure the impact of choice of bargaining on the markup chosen and of the choice of markup on the choice of bargaining. This allows for appropriately testing whether a seller that uses bargaining also tries to charge a higher markup. Chapter 4 of the dissertation studies the impact of best offer and markup on the likelihood of selling (using the reduced form approach through a probit model). I also study the impact of BO and markup on the duration or the time it takes for a car to sell (using survival analysis), and on the transaction price of a car (using the Cragg model, or two-step hurdle model). I then study the impact of a higher markup on a seller's bargaining strength. Finally, Chapter 5 concludes the dissertation.

1.3 eBay Motors and eBay Best Offer

eBay Motors is the market on eBay that deals exclusively with automobiles, and had been named the "most successful companies of the internet era (Adams et al., 2006)." As is typical with other categories on eBay, eBay motors consists of four types of sale formats: fixed price listings called Buy-It-Now (BIN) listings, auctions without BIN prices [1], auctions with BIN prices, and (since 2005 onwards) fixed price listings[2] with "Best Offer (BO)" features

[1] Auction listings are not included in this study. Auctions are also subject to different rules and listing fees on eBay than fixed price listings (including both BO and non-BO listings), making the comparison between auctions and fixed price listings more complicated. I assume the independence of irrelevant alternatives applies (IIA) here, and that that the availability of auctions does not affect sellers' choice between posted prices and BO.

[2] Sellers can list their item for 3, 5, 7, 10 or 21 days under the fixed price format (this includes BO listings)Toklu (2014).

on the listings [3] . More than 50% of eBay Motors listings under the category of "Cars and Trucks" are auctions, about 20% are BO listings, and less than 10% are (non-BO) fixed price listings (Huang et al., 2013). In terms of selling success of Toyota Camry cars, Huang et al. (2013) find that auction listings have a success rate of 32.93% and BO listings, 18.31%, coming second to auction listings. Non-BO fixed price listings have the lowest success rate, 4.71% (Toklu, 2014). Apart from the automobiles category, there is evidence that fixed price listings (BO listings and non-BO listings combined) have surpassed auction listings on eBay, in general (Einav et al., 2013). Furthermore, the BO feature "has been growing in popularity and bargained transactions currently account for nearly 10 percent of total transaction value in the marketplace (Backus et al., 2015)." However, despite the growing importance of the Best Offer feature on eBay, there have been very few studies on eBay Best Offer, and not much is known about this feature.

eBay's introduction of the Best Offer mechanism seems like an intention to "mimic" the real-world bargaining process on the internet (Chen et al., 2016). It operates in the following manner: if a seller chooses to bargain using the Best Offer feature, he needs to enable the Best Offer feature on the fixed price listing (BIN listing) by merely checking a box. [4] BIN listings cost the same as BIN listings with the BO feature enabled on it. This checking of the box by the seller changes the format of the listing from being a BIN listing (posted price

[3] "Classified Ads" also allow for the BO option, and under "Classified Ads", buyers and sellers communicate privately, and the transaction takes place off eBay Toklu (2014).

[4] I do not find that the deferred decision principle from Tversky and Shafir (1992) applies here directly, but perhaps indirectly. That is, I do not find evidence that since non-BO is the default option for sellers when listing an item that sellers use non-BO (posted prices) more often. Instead, I find that BO is used more often than non-BO (posted prices). However, perhaps an extension of the insights from Tversky and Shafir (1992), where the "default" option is based on the social norm or the social, is indeed applicable. With this respect, the "default" option on eBay Motors seems to be BO, and not posted prices, since most eBay Motors sellers tend to use BO (vs posted prices) for the fixed-price listing format. Hence, it is possible that sellers' use of BO is aligned with using the social default or the social norm. Regardless, this dissertation offers additional reasons that impact a seller's choice to use BO.

mechanism) to a BO listing (bargaining mechanism). However, even for a BO listing, a seller must list an initial asking price, that is, he must post a BIN price, and the buyer can opt purchase at the BIN price. Buyers can click the "Make Offer" button on a listing page, which allows the buyer to enter a numerical value of his offer amount. This sends an email to the seller who then has 48 hours to respond. Once the seller receives an offer by a buyer, he can either accept that offer, reject it and counteroffer if he wants to, or let the offer expire after 48 hours. Sellers can automate the process to some extent by entering private reservation prices in the system so that a certain range of offers is accepted and a certain range of offers is automatically rejected (Toklu, 2014). The buyer is also notified of the seller's response, giving the buyer the option to accept and checkout, or to make a counter-offer (Backus et al., 2015). Under the BO listing format, one seller can receive offers from several buyers, and there can be many negotiation threads running simultaneously for one seller for one given item. Buyers can also opt to make the purchase at the BIN price for a BO listing, in order to circumvent any bargaining (Toklu, 2014). A transaction occurs when the buyer and the seller mutually agree on the offer price or if the buyer opts to purchase at the BIN price (similar to the case of non-BO BIN listing). Sellers and buyers can go on back and forth (as they would generally do in an in-person negotiation) for a maximum number of stages set by eBay. All eBay categories, except for eBay Motors, allow for the buyer and the seller to make only three offers each and three counteroffers each, respectively. eBay Motors allows for *ten*, instead of three offers and counteroffers.

Buyers have access to some additional information during the transaction, while sellers have a lot more information during the active listing period. By clicking a link to access the offers page, buyers can see the offer history of a given variable, although buyer IDs are not fully displayed on this offer page. That is, all characters, except the first and the last, are masked with asterisks. Buyers' feedback scores are also displayed alongside the partial IDs. Using this information, buyers can know the number of other buyers competing with them for the same item. Buyers can also see the status of all offers (whether pending, expired or declined) on a given listing, including the date and time of the offers. Offer amounts are displayed on eBay only when the listing ends (regardless of whether the item sold or did not sell). In contrast, each seller has all of this information throughout the active listing period Toklu (2014) for all his listings.

It can helpful to compare BO listings to the more commonly known auction listings on eBay for a better understanding of how BO listings operate. While under an auction mechanism, sellers are committed to selling/accepting the highest bid under the auction mechanism, they can reject any offer under the Best Offer mechanism. Nevertheless, a BO listing requires that sellers respond to offers within a 48 hour time period; otherwise the offer expires, whereas they do not need to respond to bids on an auction listing. Furthermore, for an auction listing, buyers can observe the bid prices, but they cannot observe the offer amounts for BO listings. Therefore, the BO mechanism gives less information to buyers about other buyers' offer amounts compared to the information given to a buyer regarding other buyers under an auction listing.

Finally, it might help to also consider some similarities between the BO mechanism and the first price sealed bid auction. In both mechanisms, sellers can observe the submitted offers by buyers, and buyers cannot observe each others' offers. However, the duration for the submission is different for both mechanisms. That is, in the BO mechanism, the period for submission is different for each potential buyer, and is endogenously chosen by that buyer; the negotiation period begins from the time of the buyer's first offer. Whereas, in a first price sealed bid auction format, all buyers are subject to the same time period, *and* buyers are also committed to their offers till the end of the period (vs. only 48 hours in a BO listing). Therefore, first price sealed bid auction formats can be seen as a special case of the BO format: if in the BO listing, all buyers come in at the last 48 hours of the listing period, then the BO mechanism works exactly like the first price sealed bid auction with a BIN price (Toklu, 2014).

1.4 Previous Studies on eBay Best Offer

Despite extensive literature on eBay Motors and eBay auction listings, in general, there are very few studies on eBay Best Offer. Huang et al. (2013) analyze buyers' behavior to find that buyers do use the information available on the eBay listing when making their first offer on a BO listing using a 2008 dataset of Toyota Camrys. The main predictors of the buyer's first offer are the length of time since the start of the listing until time of the offer and the number of other buyers having made an offer prior are significant predictors of a buyer's offer. The authors conclude that sellers act strategically when using the Best Offer feature on eBay. They postulate that buyers with higher reservation values makes offers earlier, but there are many buyers making offers towards the end of the listing, because towards the end of the listing, the seller loses negotiating power, because the threat of the negotiation ending is higher. Their findings are consistent with theoretical insights on bargaining, where the two main drivers of strategic bargaining models are time delay and the threat of negotiation ending without agreement (Osborne and Rubinstein, 1990). With respect to seller behavior on eBay Motors Best Offer, Toklu (2014) concludes that sellers on eBay motors also act strategically on BO listings and play hard bargaining strategies, using a 2011 dataset of Ford Mustangs. Toklu (2014) estimates buyers' private valuations and proposes a distribution. Then, using counter-factual simulations, he concludes that sellers' behavior in the market is consistent with revenue maximization. Therefore, these two studies provide some evidence that buyers and sellers act strategically on eBay Motors BO listings, and eBay Best Offer seems appropriate for studying predictions of economic bargaining literature.

In addition to confirming predictions from economic bargaining theories, these studies also offer insight into the variables on eBay that are statistically significant in buyer and seller behavior on eBay Motors Best Offer regime, and therefore guide the selection of variables for this paper. For instance, all of these studies use the Kelly Blue Book (KBB) price to control for heterogeneity in car values based on observable characteristics such as age, mileage, and mechanical condition (Chen et al., 2009; Huang et al., 2013; Toklu, 2014). In other words, the KBB value is used as an approximation of the actual value of the car (Toklu, 2014). The markup variable used in this paper (to be discussed later) is constructed similarly as the variable used by (Chen et al., 2009). The predictors of the price premium or markup were factors such as a seller's feedback score, whether a seller is a dealer or not, the time of an offer was placed, whether the offer was on a weekend or in the morning, the offering buyer's feedback score, if the offer is the first one on the listing, the number of buyers who made an offer prior, whether the title of the listing is clear, and if there is a warranty on the carHuang et al. (2013).

While these above-mentioned studies use economic theory to dictate their variable choice (as is the case for this dissertation), Backus et al. (2015) uses a non-parametric analysis and finds a phenomenon in the data that they classify, using economic theory. That is, they find that sellers that list prices on eBay Best Offer using round numbers (or multiples of 10s and 100s), are in-fact impatient sellers that are willing to sell the item for a lower price than sellers whose list prices are precise numbers. They use a dataset of millions of observations on the eBay collectibles category. Their main contribution is the classification of the signaling equilibrium: Buyers and sellers both understand this signaling or "cheap talk" mechanism, and therefore, listings who have round-number pricing receive lower offers, sell faster, and sell for a lower price than listings that have precise-number pricing. They confirm similar patterns using off eBay data in the real estate market (Backus et al., 2015). This study highlights the need for exploring and experimenting empirically with the data available on eBay Best Offer. Given the newness of online bargaining platforms, more empirical testing is needed to understand available data and online behavior to see which aspect of economic theory this data/behavior actually represents. In this dissertation, although I do not classify variables under a certain theoretical concept based on a non-parametric approach, I do use some new variables (not previously used in any eBay study) to represent a theoretical concept of interest for my analyses. I also use a different variable to measure seller patience than the one classified by Backus et al. (2015), and complement their work on measuring seller patience on eBay. This dissertation borrows from and contributes to insights regarding variable construction for this emerging literature on eBay Best Offer. These variables will be explained later on in this paper.

Finally, Chen et al. (2016) relates most closely to this dissertation compared with the above-mentioned eBay Best Offer studies; they also study bargaining vs. posted prices on eBay Best Offer. Their focus is to study the impact of BO (vs. non-BO fixed price listings) on the probability of selling and on transaction price to determine the benefits of BO vs non-BO fixed price listings. They find that while BO does increase the probability of selling, it does not seem to have any impact on the transaction price of a car. They also claim that if an item is listed as BO, it also has a higher listed price or BIN price (as they expected). Using a different dataset and differing empirical techniques and variables, I find differing results compared to Chen et al. (2016). I do not find that BO alone increases the likelihood of selling or has any impact on the transaction price; I do find that an increase in the interaction term of markup and BO increases the probability of sale (somewhat consistent with Chen et al. (2016)), and it also increases the transaction price of the car, conditioned on the car being sold (contrary to Chen et al. (2016)). Furthermore, at the seller-level, accounting for the endogeneity of sellers choice of price and of selling mechanism, I find that sellers that tend to use BO typically have a lower average markup across all of their listings in the sample, which is somewhat contradictory to what Chen et al. (2016) claim. Therefore, this dissertation complements and extends the work of the only other empirical study that formally looks at the benefits of BO usage (compared to fixed price listings) as a means for understanding sellers choice of bargaining vs. posted prices on eBay.

In addition to studying the outcomes of BO, as Chen et al. (2016) do, I also extend the literature by also studying the question of what types of sellers adopt BO in the first place. While Backus et al. (2015) "conjecture that sellers use BO as a "demand discovery mechanism, it is worth exploring and confirming empirically what exactly leads to the emergence of bargaining as a market. This dissertation studies the question of bargaining vs. posted prices with respect to studying the determinants of sellers' choice of adopting BO, along with studying the outcomes of those choices (to understand continued use of BO). It also checks for possible negative welfare impact of BO if sellers choose to price higher as a result of using BO (and I find the opposite to be the case). Therefore, I extend previous eBay Best Offer literature and their attempts to understanding bargaining on eBay.

1.5 Previous Literature on Bargaining vs. Posted Prices

In this section, I review the five relevant theoretical studies that offer insights into sellers choice between bargaining vs. posted prices. Model assumptions are discussed, and are later compared with the market structure of eBay. Although these theoretical models do not match the eBay environment perfectly, they are the closest we get towards formal insights regarding eBay Best Offer, along with a few eBay Best Offer studies (also mentioned later in this section). The predictions of these models lend the primary hypotheses for this dissertation and inform the variable construction for the analyses.

While Riley and Zeckhauser (1983), Wang (1995) and Bing (2009) model the interaction between one seller and one buyer, Bester (1993) models the interaction between many sellers and one buyer. This distinguished the Bester (1993) model as being more applicable to the context of eBay, as Chen et al. (2016) have already pointed out, because typically, there are many sellers on eBay. However, the other three models also allow for more buyers, unlike Bester (1993), and eBay also typically has many buyers. Therefore, there is not one model that fits eBay perfectly in terms of the market structure, and insights from all models must be kept handy. Furthermore, as is the case with Chen et al. (2016)s work, empirical studies, such as this one, can help distinguish between these theoretical insights, and can help point out which models are more applicable to the context of eBay Best Offer.

In Riley and Zeckhauser (1983), a seller posts a price and if that price is rejected, the seller can sell to another buyer at a lower price, but switching is costly. Whenever commitment to a price is possible, the seller should choose to post a fixed price. When sellers cannot commit because reputations are difficult to establish or market encounters are highly occasional, sellers would choose to bargain. Having a high reputation is one way that a seller can commit. The model operates under the assumption that a seller who is able to make a commitment can employ any strategy he wishes, and can convey this commitment to each buyer.

Wang (1995) has one seller with buyers arriving according to a poisson process, whose valuations are drawn randomly from a cumulative distribution, characterized by the function F(.), with continuous probability distribution f(.). Sellers are assumed to be risk-neutral, and they incur a discount in payoff while they wait for buyers to arrive. The seller has the choice of posting a price or of bargaining with the buyers that arrive. Bargaining incurs an additional cost, compared to the posted price method. For the bargaining game, the seller sets a reserve price, and the Nash bargaining solution offers the final transaction price. If the bargaining solution is higher than the sellers reserve price, the item is sold. For the posted price case, in this model, if a buyers randomly drawn valuation exceeds the posted price, the item is sold. Therefore, in this model, the seller chooses to bargain if he has higher bargaining power. Another testable prediction that emerges from this paper is that if a seller has many objects to sell, and there are many potential buyers, the seller faces a higher bargaining cost, and therefore, the seller would prefer to use the posted price mechanism. Finally, items which are difficult to value, expensive or/and are costly to display are more likely to be sold through bargaining.

As Chen et al. (2016) have pointed out, Wang (1995) and Bester (1993) both use Nash bargaining to arrive at the bargaining solution, but there is no reserve price set by the seller in Bester (1993), and there are many sellers in this model compared to the one-seller model in Wang (1995). If a buyer has an outside option, then when his bargaining power is low, he opts for the outside option. This leads to the seller choosing bargaining only when he has relatively weaker bargaining power; that is, if the buyer has sufficient bargaining power to want to engage in bargaining. Therefore, according to both Bester (1993) and Wang (1995), the seller is incentivized to adopt bargaining when he has higher bargaining power, but the modeling difference in both models, in terms of the buyers bargaining power, leads to differing predictions for the equilibrium outcome. In Wang (1995), a seller having high bargaining power is the reason for his adoption of bargaining, in Bester (1993), a seller having weak bargaining power chooses to bargain in equilibrium. Bester (1993) emphasizes quality uncertainty as key for this result. That is, the quality of a product is determined after observing the item and moral hazard concerns regarding the quality of the product play a role in buyers decision-making. Under bargaining, the price is determined ex-post, so sellers have an incentive to control the quality of the product. Whereas under posted prices, prices are set ex-ante, and therefore sellers do not have an incentive to control the quality of the product (Bester, 1993). While bargaining may resolve some moral hazard concerns for the buyer, if the seller has too much bargaining power, and buyers face switching costs, they would prefer not to get locked-in with a seller with high bargaining power. Therefore, given that buyers have an outside option, they would only choose the bargaining seller if their relative power was not too low. That is, if the seller has too much bargaining power, buyers would not choose to bargain, and sellers would be better off choosing a posted price method. This leads to the seemingly counter-intuitive conclusion (one that seems at odds with all of the other literature mentioned) that sellers adopt bargaining when they have lower bargaining power. Chen et al. (2016) have found some empirical evidence that is indeed consistent with Bester (1993): they find that sellers with higher bargaining power tend *not* to choose to use Best Offer.

The above-mentioned models do not explicitly consider the time cost or bargaining cost, but the following two models do. Bing (2009) has one seller with many consumers and considers a time cost of bargaining in two manners: a discount factor cost (the pie over which the negotiation is occurring gets smaller by a fixed percentage as each stage of the negotiation passes), and a constant bargaining cost per period. Both types of cost lead to the result that the more patient a seller is, the larger his gains from bargaining are, compared to those from posted prices. This is due to the fact that a more patient seller uses time delay to signal (or to pretend) that his valuation for a given good is high and is in no hurry to sell the good. He is willing to negotiate for longer at a lower bargaining cost or lower time cost. Therefore the patience of a seller gives him higher bargaining power and higher profits through a negotiation. Bing (2009) also suggests that if a seller has lower labor costs, he would adopt bargaining.

Perry (1986) also considers time cost of bargaining using a one-seller and one-buyer model, where the time cost is treated as a constant cost per period. Perry concludes that in equilibrium, there will be no bargaining. Yet, bargaining markets persists despite this theory. Since Perry's assumptions do not result in a bargaining equilibrium, perhaps the relaxation of the assumptions Perry made in his paper could lead to bargaining occurring in equilibrium. Perry shows that if the known costs per period in bargaining and reservation prices are the only unknowns but the distribution of the buyers and sellers valuations are common knowledge, then there is nothing to be learned, and therefore there is no bargaining. He assumes common knowledge of the players' valuation distributions and common knowledge of the players' time costs, where time costs are also homogeneous across all buyers and across all sellers (Perry, 1986). Relaxation of one or more of these assumptions may result in the choice to bargain, and therefore, I test some of the following hypotheses by relaxing some of Perry's assumptions as follows:

1. If a seller's valuation is not known to anyone else in the market, bargaining is more likely to occur.

2. If a buyer's valuation is not known to anyone else in the market, bargaining is more likely to occur.

3. If a seller's time cost is not known to anyone else in the market, bargaining is more likely to occur.

4. If a buyer's time cost is not known to anyone else in the market, bargaining is more likely to occur.

The hypotheses and the variables linked to these hypotheses are summarized in Table 1-2.

1.6 Theoretical Framework for the Dissertation

Given that the above literature assumes a one-seller market, except for Bester (1993), the conclusions of these theoretical models are not directly applicable to studying Bargaining vs. Posted Prices in the eBay Motors market, since eBay Motors definitely has more than one seller, typically selling a differentiated product. Previous eBay Best Offer literature does not offer many models, but attempts to set some theoretical background for bargaining on eBay, nonetheless. Huang et al. (2013) and Toklu (2014) follow the models used for Priceline.com's feature Name-Your-Own-Price(NYOP), but with some variations. Huang et al. (2013) emphasizes that eBay products tend to be less standardized and more restricted in supply than Priceline's products. That is, while there can be multiple units of the same standardized flight ticket being sold on Priceline.com, eBay, and especially eBay Motors, tends to sell fewer units of a good, especially if it is a used item or used car. In the case of used cars, each car is a unique product. In this manner, Huang et al. (2013) argue that for the automobile market on eBay, a limited supply model is more appropriate. Furthermore, Huang et al. (2013) emphasize that NYOP models do not typically need to account for disclosed information, because Priceline.com does not provide information to buyers regarding other buyers' offers as eBay does. In general, eBay discloses considerably more information regarding the buyer and the seller than Priceline.com does, and Huang et al. (2013) confirm that buyers base their behavior on this available information. While Toklu (2014) does not allow for the dynamic updating of information by the buyers, he models the eBay market following previous NYOP models, but with a finite number of sellers and buyers, whereas Huang et al. (2013) use a one-seller (and many potential buyers) model. Finally, Backus et al. (2015) use a stylistic model that consists of two types of sellers (one impatient type and one patient type) and many potential buyers to model bargaining on eBay Best Offer. Chen et al. (2016) do not use a specific theoretical model, and they instead emphasize that their empirical work is capable of distinguishing between theoretical models. In doing so, they find evidence consistent with Wang (1995) and Bester (1993), suggesting that the higher time cost of bargaining and the multiple-seller model (allowing outside options for the buyers) seem to be key features for the context of eBay. In the spirit of Chen et al. (2016), this dissertation also does not use any one specific theoretical model, but rather aims to test insights of these models. The variables in this paper are constructed keeping in mind the theoretical framework for eBay using these previous papers, where the key features for an eBay model seem to be: more than one seller, differing in their levels of patience; limited supply of a non-standard product; differing level of knowledge/experience for agents; and some consideration of a seller's reputation and time cost of bargaining.

1.7 Data

This paper uses completed listings from eBay Motors Cars and Trucks category, collected in 2013 and 2016. Since the 2013 sample is much smaller than the 2016 data, I make more use of the 2016 data in this dissertation. Where applicable though, I run robustness checks using a combined dataset that includes both 2013 and 2016 samples. After the data was collected, duplicate listings were dropped, important variables were cleaned first, and observations were dropped if key variables for the analysis were missing. In addition, observations were dropped if their VINs were invalid or if there were duplicate VINs listed multiple times on eBay.

1.7.1 2013 Data Collection

The data was collected from eBay Motors completed BIN listings using the Web Content Extraction Software Version 6.3, from www.newprosoft.com. Since completed listings are only on eBay Motors for some months after their completion. I sorted listings by "Date listed: Oldest First", in order to get as old listings as possible, because this also enlarges my overall sample. The extractor ran from November 15th through November 18th, and the start dates of the listings in my sample range from Jul 29th 2013 through Nov 11th 2013. The range of end dates of the listings in my sample is Aug 18 2013 through Nov 16 2013. The data set thus represents a cross sectional snapshot of completed BIN listings on eBay Motors Cars and Trucks market for 3.5 months from Jul 29th 2013 to Nov 16th, 2013.

During the three days of data extraction, the software would stop running at times. In order to collect the listings in the time period as completely as possible, the extractor was re-run a couple pages before the page where the extractor had stopped. This is because as new listings got completed, the previous listings became older and could now only be found on previous pages. Data extraction was stopped when over 3000 listings were collected.

Since the main objective of this dissertation is to study the seller's choice of selling method, the item/car/listing-level data was aggregated at the seller-level, and a seller-level dataset was constructed. The seller sample was determined by the randomly selected items in the sample. Using all the eBay seller usernames (which appear on the listing page of an item) associated with the over 3000 listings extracted, extractions were run for each seller's profile page on eBay, collecting publicly available information on the eBay seller.

To assemble a rich data set, I needed to get data from a few other webpages and merge the dataset. I collected the Vehicle Identification Numbers (VINs) for each car entered by the sellers on the eBay listing/item pages, and used these VIN numbers to extract information from CarGuru.com, namely the Instant Market Value for each car. The following variables were also needed for each listing: duration of a listing, revision history of a listing, and whether or not the listing has the BO feature enabled on it. The revision history variables have never been collected in any eBay study prior to this one. The revision history of a listing is publicly available on eBay on pages, linked to the listing page. The extractor therefore needed to run separately for the revision page data collection. Data regarding the duration of a listing and the enabling of the BO feature also needed to be extracted from a different link to pages that are only available on the eBay API platforms (often only used by software developers). The data collection from the API platform was not as clean as data collection from the eBay pages because of the difference in scraping from an HTML page vs. from a text page. An observation was dropped if it did not extract correctly or if it was missing key variables. The empirical methodology (to be outlined below) required fixed effects dummies for the Make of the car, and this would only work well with more observations per car Make. Hence, observations were dropped if the Make of the car (Honda, Acura, for example) had less than 10 listings in the sample I extracted. As mentioned earlier, if a VIN car occurred multiple times in the data (that is, it had ended multiple times within the 3-4 month period), it was dropped. Furthermore, if the VIN could not be verified on CarGuru.com, the listing was dropped from the sample. Once all the item-level/listing-level variables were merged in one data set, all of these variables were aggregated and constructed at the seller-level and matched with the seller usernames to complete the seller-level data set required. The final sample, at the item-level, consists of 1107 cars, and this translates to 450 seller-level observations.

1.7.2 2016 Data Collection

In 2016, the sample selection method differed slightly from that in 2013. Data was collected from March 3rd, 2016 to March 31st, 2016, and consists of listings with start dates ranging from December 11th, 2015 to March 16th, 2016, and with end dates ranging from Dec 9th, 2015 to May 27th, 2016.

I collected data on the same car models as the ones randomly selected into the 2013 data. Therefore, while the item selection in 2013 was completely random, it was not the case in 2016, because it was linked to the randomness of the items available on eBay in 2013. However, despite choosing the same models, the extractor was programmed to collect the first 50-200 results for each model. This implies that the *number* of each type of model was still random. In some cases, despite my attempts to have repeat observations on the same model from the 2013 sample, there was no data to be collected in 2016 for that particular car model. The data set thus represents a cross sectional snapshot of completed BIN listings in 2016 for certain chosen car models on eBay Motors Cars and Trucks market for 3.5 months from Jul 29th 2013 to Nov 16th, 2013.

The extractor was programmed to extract the oldest listings first using the sorting "Date Listed: Oldest First", as was the case in 2013. However, in 2016, I had an additional hurdle: at times, the extractor had errors that caused the filtering mechanism to automatically change to "nearest first" with an embedded zipcode of my hometown, "33063", which I could not hide from eBay, despite changing stealth settings on my computer. This problem had not occurred in 2013 because eBay could not read zip-codes as easily then. Therefore, the 2016 data is not as perfect a representation of the eBay market in terms of geography of the car market and make of the cars as it was in 2013.

Although the entire data collection process originally collected over 30,000 listings, many of those listings were repeat listings because I used two different computers, and this led to repeated items. This repeat collection ensured a higher likelihood of obtaining a richer data set with more variables to overcome random glitches of software extraction in one computer, by using a hopefully better copy of that item on the other computer. That is, this was one method of working around the glitches of my extraction software. I later cleaned the data sets and merged information/variables on these repeat items, and kept the copies of the items that allowed for more key variables being collected. I dropped the accidental auctions that sneaked into the dataset through extraction glitches. Because of some of these data extraction limitations, it was difficult to plan out ex-ante the number of observations that would be in the sample set; therefore, I focused on collecting observations until I reached a large number like 30,000.

As was the case in 2013, listings with duplicate VINs, listings that did not have a valid VIN, listings that did not include the BO variable or price variable, and International listings were all dropped. Therefore, including a variable that does not have a lot of observations (because it was more difficult to collect technically, and at times, it was a variable not entered by the seller enough times) can reduce the number of observations in the regression.

Seller Sample Selection was determined by the item selection, and the seller-level dataset was assembled by aggregating the item-level variables at the seller level, as was the case in 2013.

The final item-level dataset consists of 2497 observations or cars in the sample. At the seller-level, this translates to 1861 sellers. The sample for sold cars contains 305 observations.

1.8 Variables

This section defines all the variables used throughout the dissertation. I separate the variables by item-level variables and seller-level variables. Where appropriate, I discuss the relationship I expect the variable to have with bargaining or BO[5] .

1.8.1 Item-Level Variables

- BO is a dummy variable that takes on a value of 1 if the listing has Best Offer enabled on it, and a value of 0 otherwise.
- sold is a dummy variable that takes on a value of 1 if the listing was sold, and a value of 0 otherwise.
- durationdays is a measure of the number of days that a listing lasted. This measure, when the item is sold, represents the number of days it took for an item to sell. However, if the item did not sell, this variable represents the number of days that the seller intended for the listing to last.
- discountratio is a measure of the number discount rate provided on an item that was sold under BO. It is measured as follows:

$$\frac{transactprice_i - BIN_i}{BIN_i}$$

- lntransact is a variable that take a transformation of the transaction price of a car. That is,

$$lntransact = ln(transactprice_i)$$

- priceprem is the markup of a seller can be defined as the excess in the seller's asking price from the average benchmark price divided by the average benchmark price for that car. The markup measure assesses the seller's asking price for a car on eBay against the standard estimate suggested by CarGurus.com. This "Instant Market Value (IMV)" estimate was derived by CarGurus.com using the Vehicle Identification Number (VIN) on the eBay listing of the car. According to CarGurus.com, "the Instant Market Value (IMV) of a vehicle is CarGurus estimated fair retail price for a vehicle based on a detailed analysis of comparable current and previously sold car listings in [the] local market. [The] analysis takes into account specific details including make, model, trim, year, mileage, options and vehicle history." Therefore, the IMV can be written as the following:

[5] Where equations are present, i indexes an item, while j indexes a seller.

$$IMV_i = G(Make_i, Model_i, Trim_i, Year_i, Mileage_i, Options_i, \\ VehicleHistory_i, ComparablePrices_i, LocalCarMarket_i)$$

,
where G is an unknown function that CarGurus.com uses to estimate the retail price based on the variables listed in parentheses.
According to CarGurus.com, the function for estimating the IMV has remained the same from 2013 to 2016, and therefore can be used consistently for analysis in this dissertation. The IMV therefore serves as an average benchmark price. This measure of the price premium or markup is similar to the ones used in previous literature (Andrews and Benzing, 2007) using KelleyBlueBook.com as the average benchmark price for a given car.

$$\frac{\text{BIN}_i - \text{IMV}_i}{\text{IMV}_i}$$

- transactprice is a variable that captures the transaction price dollar amount of a car for cars that sold.

- InstantMarketValue is the CarGurus.com retail estimate for each car, based on characteristics of the particular vehicle. In addition to Mileage, Year, Make, and Model of the car, InstantMarketValue helps control for any remaining car characteristics that I did not directly observe in the data. CarGurus.com can read these variables based on the Vehicle Identification Number(VIN) of each car, which I obtained from the seller's data entry on the eBay listing page. I expect that cars of a higher price are more likely to be bargained over.

- Mileage is the marked mileage on a given car. I expect that the older the car, that is, the higher the mileage, the more likely for it to be sold through Best Offer.

- Year is the year the car model was introduced. I expect that the older the car us by year, the more likely it is to be listed as BO.

- No. of BIN revisions (unscale) is the total number of BIN revisions on a listing (among the most recent revisions on the listing). BIN revisions imply that the seller has learned new information about pricing, and has therefore chosen to revise it. Not all BIN revisions by the seller can be observed through this revision variable, because the seller can choose to relist an item under a completely new item number with a revised price, and it would not be captured under my data collection as a relisted item or as a revised item. This means that this measure underestimates the number of the seller's BIN revisions.

- No. of BIN revisions (scale) is the same as the variable, No. of BIN revisions (unscale), except that it has been scaled by the number of observations collected on revisions in each year. That is, in 2013, 12 recent revisions were collected for each listing, whereas in

2016, only 4 recent revisions were collected. Therefore, in the the 2013 dataset, No. of BIN revisions (unscale) is divided by 12 to get No. of BIN revisions (scale), whereas in 2016, No. of BIN revisions (unscale) is divided by 4.

- BO revisions (unscale) is the total number of most recent BO revisions collected for a listing, similar to the variable, BIN revisions (unscale). This indicates that the seller has learned new information on how to use Best Offer, and has therefore chosen to revise it.

- BO revisions (scale) is the same as BO revisions (unscale) variable, but is scaled similar to BIN revisions (unscale).

- revised is a dummy variable that takes on a value of 1 if the listing has been revised at least once, and a value of 0 otherwise. A revision implies that there was learning and implementation of that learning with regard to that listing or item. Therefore, as revision or learning increases, I expect that BO usage will be reduced.

- relist is a dummy variable that takes on a value of 1 if the item has been relisted by the seller.

- unavailable is a dummy variable that takes on a value of 1 if the item was made unavailable by the seller, and 0 otherwise. I do not have an expected sign for this variable.

- deposit is a categorical dummy that takes on different values for the time period the seller requires a deposit payment in. deposit equals "24 hours", deposit equals "48 days", deposit equals "72 days", deposit equals "immediately", and deposit equals "none mentioned," that is no deposit payment period is mentioned in the listing. This variable has never before used in an eBay study. This variable represents the seller's signal to the buyer regarding his time seriousness. I do not predict a sign for this variable, but I expect that his signal for time seriousness would have an impact on his choice of BO.

- itemcountstateyr is the number of cars in the dataset for a given year that are from the same state. I expect that as the competition by state increases, it is likely to decrease the likelihood of BO.

- itemcountmakeeyr is the number of cars in the dataset for a given year that are of the same Make. I expect that as the competition by Make increases, it is likely to decrease the likelihood of BO.

- ModelavgIMVby1000 is the average IMV for a particular model in the car. It helps control for the expensiveness of a given model, and other unobserved chracteristics of a particular model. As mentioned earlier, I expect that if a car model is relatively more expensive, it is likely to be bargained over. That is, I expect a positive sign on this variable.

- State Dummy is a dummy variable for each state, which allows for controlling for unobserved market/demand characteristics, specific to a state.

- Make Dummy is a dummy variable for each Car Make, which allows for controlling for unobserved market/demand characteristics, specific to a state.

1.8.2 Seller-Level Variables

Several of the following seller-level variables are constructed by aggregating item-level variables at the seller level.

- BOperlisting is the fraction of a seller's listings in the sample that are listed as Best Offer.

- Selleravgpriceprem is the average item-level markup/priceprem (explained above) across all of the seller's listings in the sample.

- Patience
  There is not much literature where patience is studied empirically, even outside the field of Economics. Nonetheless, because time may play a key role in bargaining, it becomes crucial to consider the impact of the patience of a seller. This paper uses a variable for patience that has never been tested before, where patience is conceptualized using previous empirical literature on patience. The infrequent times when patience has been noted in research, patience has often been conceptualized as the absence of impatience. Blount and Janicik (2001) consider what constitutes patience (rather than impatience) directly; they claim that there are two primary delay triggers, postponement, and tolerance, to which subjects respond and exhibit patience (Dudley, 2003). Researchers find that self-reported measures of patience are less reliable than measures of patience derived from behavioral observation (Berensky, 2004; Fowler and Kam, 2006). Therefore, this paper measures an individual seller's patience through behavioral observation, rather than using seller self-reports of patience (as is the case in previous empirical studies on patience).

  - Selleravgrelist, the average relisting behavior of a seller across all his items in the sample, measures the patience/tolerance of a seller in two ways. Firstly, sellers are known to automatically relist items. That is, a seller exhibits tolerance for selling the car on eBay in his initial format of listing, without ability to change the duration of the listing. Hence, a seller that uses automatic relisting intends to tolerate the wait till the item sells on eBay. The willingness to wait and tolerate the consequences of listing until the item sells signals that the seller is a more patient seller.
    Secondly, sellers can also manually relist their items by a one-click mechanism on eBay for relisting the item. A seller who manually relists an item in this manner exhibits patience, tolerance, and persistence for selling the item on eBay, instead of searching for off-eBay options for selling. Controlling for other indicators of

seller's outside options (as I do), even if a seller clicks to relist manually, he shows persistence in selling his items.

$$Selleravgrelist = \frac{\#ofCarsRelisted_j}{\#ofSellerListings_j}$$

- Competition For the purposes of this paper, the competition in the eBay Motors market is defined along two dimensions - the Make of the car and the geographic location of the car. Since this data set ended up having fewer of the same models, the market of a particular car *model* is more difficult to analyze than the market of a given car's *make*. In terms of geographic competition, the city-level data was the most disaggregated data available on eBay.

  – countmakeSeller is a variable that measures the competition that a seller faces on eBay from sellers of the same car make (but not necessarily the same model). The variable was constructed by counting up the number of sellers on eBay selling the car of the same Make as a given seller. The higher this number, the more market competition the seller is exposed to.

  – countcitySeller is a variable that measures the competition that a seller faces on eBay from sellers of the same city. The variable was constructed by counting up the number of sellers on eBay selling a car on eBay in the same city as a given seller. The higher this number, the more geographic market competition the seller is exposed to.

- Patience*Competition is the interaction of seller patience with competition variables. I do this because it seems that higher competition can change the relative bargaining power dynamics between sellers and buyers. That is, according to Bester (1993), buyers have higher bargaining power when there is higher seller competition in the market, compared to when there are fewer sellers in the market. Since Bester (1993) presents differing results than other previous models regarding sellers' choices between bargaining and posted prices, I find it worth studying the impact of seller's bargaining power, in the form of patience, in the context of higher seller competition.

  – competitcity*relist is an interaction term variable, where the geographical competition faced by the seller is interacted with seller patience. It is constructed as follows:

  $$competitcity * relist = countcitySeller * Selleravgrelist$$

  – competitmake*relist is an interaction term variable, where the competition by Car Make faced by the seller is interacted with seller patience. It is constructed as follows:

  $$competitmake * relist = countmakeSeller * Selleravgrelist$$

- Seller's Outside Options

- Selleravgunavailable is a listing characteristic variable that measures the average number of a seller's listings that are ended by the seller, claiming that the item was either sold prior, was no longer available or was ended due to an error in a listing. I use this variable to measure how often the seller has an outside option that makes an item unavailable on eBay. Although the item having been "sold prior" was an option used by sellers in 2013, I did not find any usage of this option in 2016. I use the general unavailability of the item, in this case, a car, to mean that the seller chose to end the listing car for an outside option, regardless of the reason he selected on eBay. On the one hand, the seller may exercise his bargaining power by setting a higher posted price. On the other hand, he may choose to bargain due to higher bargaining power. So, I do not predict a sign for this variable.
- Similar to the case of bargaining power stemming from seller patience, I also interact the seller's bargaining power that stems from his outside options with the seller competition he faces. That is, I also study the following variables:

$$competitcity * unavailable = countcitySeller * Selleravgunavailable$$

$$competitmake * unavailable = countmakeSeller * Selleravgunavailable$$

  I do not predict a sign for the impact of these variables on the seller's use of BO.

- Seller's Communication
  - SelleravgQA is a listing characteristic variable that measures the average number of a seller's listings that had at least one question or more by the buyer along with the seller's answer (or answers) posted on the listing.
- Seller Reputation
  I expect that a seller's reputation will make him less likely to use bargaining, as would be consistent with Riley and Zeckhauser (1983); Bing (2009).
  - I use a seller's Positive Feedback Percent, listed on his profile page, as well as on each of his item listing pages as the reputation variable, where

$$PositiveFeedbackPercent_j = \frac{\#ofPositiveFeedbackPoints_j}{\#ofPositiveFeedbackPoints_j + \#ofNegativeFeedbackPoints_j}$$

  I expect that a the higher a PositiveFeedbackPercent is, the less likely he will be to use bargaining.
- Seller Knowledge/Experience/Learning

- Feedbackscoreby1000 is the variable used to represent a seller's experience on eBay. According to Chen et al. (2009), feedback score of a seller is generally used as a reputation measure in eBay auctions literature, along with variables that measure positive feedback, but it can instead serve as a variable representing experience of an eBay user (Huang et al., 2013). This is because feedback can only be provided on eBay from a buyer to a seller or from a seller to a buyer after there is a completed transaction between the two agents. Moreover, one buyer-seller pair can provide feedback for each other only one time. This variable thus indicates the number of new threads of transactions an eBay user has completed. I use feedback score to represent a seller's knowledge and experience, and therefore I expect that the coefficient on this variable will be negative in sign.

- selleravgrevisions is a listing characteristic variable that measures the average number of revisions for each seller across all his listings. In 2013, for each item by a seller, the data collection extracted 12 most recent revisions. However, in 2016, the data collection extracted only 4 most recent revisions. Therefore, at the item-level, I scaled the revisions by the number of observations collected on them. That is, the scaled versions of the revision variables for 2013 divided the total number of revisions on an item by 12, and by 4 in 2016. I use the scaled versions of the revision variables in the analysis, unless otherwise mentioned. I expect the sign on the coefficient for this variable to be negative. That is, I expect that increase in seller's knowledge indicated by his increase in learning/revisions is likely to lead him to choose against bargaining.

$$\frac{\#ofRevisions_j}{\#ofSellerListings_j}$$

  I also use variables that concern seller learning specific to pricing and seller learning specific to the BO feature, measured by variables SelleravgBINrevsunscale and SelleravgBOrevsunscale, respecitvely.
  Note:

$$SelleravgBINrevsunscale = \frac{\#ofBINrevisions_j}{\#ofSellerListings_j}$$

  and

$$SelleravgBOrevsunscale = \frac{\#ofBOrevisions_j}{\#ofSellerListings_j}$$

- sellerprivate is a dummy variable that takes on a value of 1 if the seller is a private seller, and 0 if the seller is a dealer. If a seller is a private seller, then he is likely to have less outside options than a car dealer, and therefore lower time cost for selling the car on eBay. In addition, a private seller has lower learning and experience of selling and of selling on eBay than a dealer. A private seller's lower time cost, lower experience, and lower knowledge, all point to a private seller being more likely to use BO. Therefore, I expect the coefficient for this variable to have a positive sign with respect to explaining BO.

- Seller's Time Cost

  sellerlistingsN is the number of total cars with valid VINs by a seller in the sample. If the seller has a higher number of items to sell, and therefore has higher time cost, then according to Wang (1995), the seller is less likely to use bargaining. Chen et al. (2016) also use a similar variable in their empirical analysis on eBay Best Offer.

- Object Valuation

  - SelleravgMileage is the average mileage for the cars in the sample by a seller. I expect that old cars are more likely to be sold through bargaining because old cars are harder to value. That is, sellers that sell older cars are more likely to use bargaining as a means of price determination, and therefore I expect the sign on SelleravgMileage to be positive.
  - SelleravgYear is the average model year for the cars of a seller in the sample. Old cars are expected to be more likely to be sold through bargaining, as mentioned earlier, therefore I expect SelleravgYear to be negative. These two variables, SelleravgYear and SelleravgMileage together, not only account for whether the car is old or new, but also control for average item characteristics of the seller.
  - Seller Make Dummy variables are used to control for the average number of seller listings in a given car make. For example, if a seller's Toyota cars make up 80% of his total listings in the sample, and Honda cars make up 20% of his listings, then SelleravgToyota variable will have a value of .8, and SelleravgHonda will have a value of .2.

1.9 eBay Motors: Descriptive Statistics

In this section, I present a very brief descriptive overview of the data used for this dissertation, which includes characteristics of the cars, sellers, listings, and the eBay market. The summary statistics are presented in Tables 1-1 - 1-5. Although much of the analysis in this dissertation is done using the 2016 data, I do use the smaller dataset from 2013 for Chapter 2, and to run robustness checks for Chapter 3. Since I did not collect variables in 2013 regarding whether or not an item sold, it is difficult to do the analysis done in Chapter 4, using the 2013 data. Therefore, I first present the 2016 descriptive statistics first, and I mention the 2013 statistics for comparison.

In the 2016 item-level sample (see Table 1-1), most of the cars in the sample are relatively new in terms of the year, but used and older in terms of mileage. That is, the average Model

Year for a car in the sample is 2010, whereas the average car mileage in the sample is 56,046 miles. The average value of a car or IMV, according to Car Guru, is $30,263. On average, the markup on a car was actually a mark*down* of 6.9%; that is, sellers on eBay list a car at a lower markup than Car Guru's estimate. For the BO listings, the average markup is 5.2%; it is lower than the overall average markup.

According to this 2016 item-level sample (see Table 1-1), 84% of the cars were listed under the BO mechanism. 24% of the items in the in the sample are from private sellers, while the rest are from dealers. 12% of the the total listings in the sample resulted in a sale. The average duration for a listing during the data collection period was about 9 days.

For cars that sold (see Table 1-2, 81% of them were listed under BO, and for cars that sold, the markup was much lower, with a value of -40.9% compared to the average markup value in the sample of -6.9%. Interestingly, though, among the cars that sold, BO listings had 10% higher markups on average, than the overall average markup. For cars that sold, the average duration of the listing was 5.5 days, the average IMV was $16,659, and the average transaction price was about $14,479.

I find that most of the items either require the deposit payment immediately (48.8%) or there is nothing mentioned regarding their deposit time deadline (33.5%). Only about 17.7% of the items have a listed time period on them: of either 24 hours, 48 hours, or 72 hours.

At the seller level (see Table 1-3, 56.3% of the sellers in the 2016 sample are private sellers. The mean markup for a seller is actually a mark*down*, equaling -10.7% (lower than the markup at the item-level). Sellers list 82% of their listings as BO, on average. On average, across all their listings, sellers' revisions on BO equals 0.03, whereas their average number of revisions on BIN equal 0.22. The average number of general revisions (concerning photos, description, payment terms, shipping terms, etc.)for a seller is 0.28. Sellers tend to relist 37.8% of their total listings in the sample. The average positive feedback percent for sellers is 99.2%, and the average feedback score is 439 points. As mentioned earlier, feedback score of a seller measures a seller's experience on eBay, indicating the number of transactions a seller has had with unique buyers. The average number of listings a seller has at a given point of time, according to this sample, is 2.6. 7% of the sellers in this sample are repeat sellers; that is these 2016 sellers were also in the 2013 sample.

In terms of eBay market characteristics, there are about 589 cars, on average, from one state, and about 339 cars for a given car make, on average. Using the seller-level data, I find that there are about 8.7 sellers in a city, whereas there are about 130 sellers in a given car Make.

In 2013 (see Table 1-4), the average car model year was 2006, and the average mileage for a car is 67,333. The average value of a car or IMV, according to Car Guru, is $25,501. On average, the markup on a car was 10.4%; that is, sellers on eBay list a car at a higher markup than Car Guru's estimate (in contrast with the situation in 2016)[6] . For the BO listings, the average markup is 7.2%; it is lower than the overall average markup, as is consistent with the case of 2016, where the average BO markup is also lower than the overall markup average. 30% of the items in the the sample were from private sellers, slightly higher than the percentage of private sellers in 2016.

As Table 1-5 indicates, only 26% of the sellers in 2013 are private sellers (in contrast with the 56.3% in 2016), while the rest are car dealers. With respect to seller-level variables, I find that a seller's average markup is about -5.4%, which is higher than the average for 2016. Unlike at the item-level for 2013 though, at the seller-level, sellers tend to price their items lower, on average, than the Car Guru estimate. Sellers list 82.6% of their listings as BO, similar to the 2016 sample. Sellers tend to relist 19% of their total listings in the sample, which is much lower than the relisting behavior in the 2016 sample. On average, sellers have an average of .14 BIN revisions across all their listings, and .02 BO revisions. With respect

[6] Car Guru's method for calculation over this three-year span (2013-2016) has not changed. Therefore, the change can be attributed to market-level, seller-level, or sample changes over the years rather than a measurement difference for the IMV in 2013 vs. the IMV in 2016.

to general revisions, sellers' average number of revisions across their listings is 0.67. [7] The average positive feedback percent for sellers is 94.6%, which is much lower than the positive feedback percent in the 2016 sample, and the average feedback score is 579 feedback points. The average number of listings by one seller is 2.3. In terms of seller-level competition, on average, there are 7.5 sellers in a given city, and 54 sellers of a given car Make.

1.9.0.1 Comparing means for 2013 and 2016 Datasets

In this section, I further explore statistically some of the differences for important variables mentioned in the previous section. This statistically confirms some of the differences that exist between the 2013 dataset and the 2016 dataset. These are important to note because they impact the interpretation of the results in cases where both datasets (2013 and 2016) are used. However, they are also interesting to note because no previous study on eBay Best Offer has compared seller behavior from one year to another. This sheds more light on which variables might be more stable across years, and which variables or behaviors by sellers might be different across years on eBay.

To begin with, I do not find a statistically significant difference between best offer usage in 2013 and in 2016. In 2013, mean best offer usage is 83.3% and 81.7% in 2016 (presented in Table 1-6). Therefore, on a quick glance, it does not seem to be true in the case of eBay that the usage of bargaining goes down over time, as some theoretical papers have predicted (Lu and McAfee, 1996; Kultti, 1999).

Table 1-7 also shows that learning behavior differences between 2013 and 2016 are statistically significant. Using the scaled measure of revisions, I find that BIN revisions and BO revisions were higher in 2013 than in 2016. Using the unscaled revision variable, I do not find BO revision and BIN behavior to be statistically different in 2013 than in 2016, which is to be

[7] Note that this is the unscaled version of the revision variable. I collected 12 observations on sellers' most recent revisions, whereas in 2016, I collected only 4 observations; therefore it is not easy to compare the unscaled versions of these variables for the two years.

expected. However, given the difference in data collection with respect to the revision variable across these two datasets, it is more reasonable to use the unscaled versions of the revision variables when combining or comparing both datasets. It seems that sellers in 2016 use the revision feature on eBay more often. This difference from 2013 and 2016 could also partly be attributed to there being a higher percentage of private sellers in the 2016 sample, compared to the 2013 sample.

Interestingly, I also find that markups were higher in 2013 than in 2016 by 16.9% (see Table 1-7. This difference can also perhaps be attributed to a larger percentage of private sellers in the 2016 dataset than in the 2013 dataset. When comparing markups only for the Best Offer listings, I find that markups were still higher in 2013 than in 2016 by 12.1%.

Table 1-8 shows the difference between non-BO means and BO means for markups in each data set separately. In 2013, BO markups were 10.3% lower than non-BO markups, whereas there is no statistical difference between BO markups vs. all other markups in the 2016 dataset.

In Table 1-11, I test these relationships at the Car Make level. I find again that when using scaled revision variables for BIN and BO, there is a statistically significant difference in revision behavior in 2013 and in 2016. Average number of BIN revisions for a given Car Make in 2016 are statistically significant and 2.4% higher than those in 2013. The number of BO revisions is higher by 0.3% in 2016 than in 2013. I also confirm at the make-level that markups were higher by 17.4% in 2013 than in 2016. When comparing markups for just the BO listings, I find that markups were 12.6% higher in 2013 than in 2016.

In addition, at the model level (see Table 1-12), I find that markups are higher in 2013 than in 2016 by 17.5%, markups on Best Offer listings are also higher in 2013 than in 2016 by 12.7%. Again, I do not find a difference between Best Offer usage in 2013 and in 2016 using model-level averages. I find again that there are a higher number of BO and BIN revisions in 2016 than in 2013, using the scaled variable. I do not find any statistically significant difference using the unscaled variable for BO revisions. However, I find that BIN revisions are higher in

2016 than in 2013, even without using the scaled version of the variable. This implies that clearly, there are some revision behavior changes across the two years.

In order to understand if there was any learning on the part of repeat sellers, and to understand the nature of their learning, I compared means for repeat sellers. As mentioned earlier, 7% of the 2016 sellers are repeat sellers. BO usage mean for these sellers does not seem to change too much; 75.9% usage in 2013 and 80.5% in 2016, as can be seen in Table 1-9. I find that BIN revisions, even when using the unscaled version of the variable, was 8.1% lower in 2013 than in 2016 (see Table 1-10). This indicates that the same sellers were definitely revising BIN more often in 2016 than in 2013, confirming that revision behavior has changed across these years.

In summary, I find that markups are higher in 2013 than in 2016. I find that in 2013, non-BO markups are higher than BO markups. I also confirm that there are higher number of revisions for BO and BIN in 2016 than in 2013, and this is especially true for BIN revisions. Some of these differences between 2013 and 2016 could be due to a higher percentage of private sellers in the 2016 sample compared to the 2013 sample.

1.10 Tables

Table 1-1. Item-Level Summary Statistics for 2016

Variable	Obs.	Mean	SD	Min	Max
sold	2497	0.123	0.328	0.0	1.0
durationdays	2497	9.266	5.816	0.0	21.0
transactprice	2497	1768.507	7591.415	0.0	159649.0
priceprem	2497	-0.069	0.458	-4.6	1.0
BO	2497	0.839	0.367	0.0	1.0
priceprem*BO	2497	-0.052	0.416	-4.4	1.0
IMVby1000	2497	30.263	29.519	0.2	336.5
Mileageby1000	2497	56.046	50.178	0.0	401.0
Year	2497	2010.274	4.193	1986.0	2016.0
sellerprivate	2497	0.243	0.429	0.0	1.0
deposit=48 hours	2497	0.090	0.286	0.0	1.0
deposit=72 hours	2497	0.004	0.063	0.0	1.0
deposit=immediately	2497	0.485	0.500	0.0	1.0
deposit=none mentioned	2497	0.336	0.472	0.0	1.0
itemcountstateby100	2497	5.893	5.298	0.0	14.0
itemcountmakeby100	2497	3.392	2.734	0.0	8.8

Table 1-2. Item-Level Summary Statistics for 2016 for Sold Cars

Variable	Obs.	Mean	SD	Min	Max
sold	305	1.000	0.000	1.0	1.0
durationdays	305	5.529	4.648	0.0	21.0
lntransact	305	9.196	0.855	6.8	12.0
transactprice	305	14478.566	16986.478	900.0	159649.0
priceprem	305	-0.409	0.858	-4.6	0.7
BO	305	0.810	0.393	0.0	1.0
priceprem*BO	305	-0.306	0.759	-4.4	0.7
IMVby1000	305	16.659	16.118	4.7	165.6
Mileageby1000	305	82.996	54.496	0.0	260.0
Year	305	2007.761	4.039	1994.0	2016.0
sellerprivate	305	0.407	0.492	0.0	1.0
depositreqin2	305	0.148	0.355	0.0	1.0
depositreqin3	305	0.007	0.081	0.0	1.0
depositreqin4	305	0.452	0.499	0.0	1.0
depositreqin5	305	0.246	0.431	0.0	1.0

Table 1-3. Seller-Level Summary Statistics for 2016

Variable	Obs.	Mean	SD	Min	Max
Selleravgpriceprem	1853	-0.107	0.540	-5	1
BOperlisting	1861	0.829	0.371	0	1
SelleravgBOrevsunscale	1861	0.029	0.168	0	2
Selleravgrelist	1861	0.378	0.465	0	1
SelleravgBINrevsunscale	1861	0.215	0.510	0	4
Selleravgrevisions	1861	0.278	0.928	0	9
sellerprivate	1665	0.563	0.496	0	1
PositiveFeedbackPercent	1586	99.165	4.386	33	100
Feedbackscoreby1000	1861	0.439	2.620	-0	96
sellerlistingsN	1861	2.585	9.956	1	298
repeatsellertag	1861	0.071	0.258	0	1
countcitySeller	1861	8.662	15.331	1	78
countmakeSeller	1861	129.557	94.954	1	326

Table 1-4. Item-Level Summary Statistics for 2013

Variable	Obs.	Mean	SD	Min	Max
priceprem	896	0.104	0.466	-1	4
priceprem*BO	896	0.072	0.377	-1	3
BO	1107	0.824	0.381	0	1
IMVby1000	897	25.501	26.329	4	245
lnIMVby1000	897	2.958	0.705	1	6
Mileageby1000	1105	67.333	56.581	0	369
Year	1105	2006.368	5.978	1983	2014
sellerprivate	1107	0.299	0.458	0	1

Table 1-5. Seller-Level Summary Statistics for 2013

Variable	Obs.	Mean	SD	Min	Max
Selleravgpriceprem	450	-0.054	0.446	-3	1
BOperlisting	450	0.826	0.370	0	1
SelleravgBOrevsunscale	450	0.017	0.125	0	1
Selleravgrelist	450	0.190	0.357	0	1
SelleravgBINrevsunscale	450	0.139	0.454	0	4
Selleravgrevisions	450	0.679	1.817	0	15
sellerprivate	392	0.260	0.439	0	1
PositiveFeedbackPercent	434	94.648	20.974	0	100
Feedbackscoreby1000	448	0.579	3.202	0	65
sellerlistingsN	450	2.296	4.571	1	85
repeatsellertag	450	0.264	0.442	0	1
countcitySeller	450	7.553	11.093	1	44
countmakeSeller	443	54.271	47.943	5	139

Table 1-6. Compare BOperlisting 2013 Dataset and 2016 Dataset

	2013		2016	
Variable	**Mean**	**SD**	**Mean**	**SD**
Fraction of a Seller's Cars that are BO (*BOperlisting*)	.8333126	.3676707	.8170816	.3804437

Table 1-7. Item-Level Two-Sample t-test with unequal variances 2013 Mean - 2016 Mean

Variable	$Mean_{2013}$ - $Mean_{2016}$	**t**	N_{2013}	N_{2016}
BO usage per car	0.001	0.119	1107	5584
No. of BIN revisions per car	-0.024*	-10.739	692	5584
No. of BO revisions per car	-0.003*	-4.109	692	5584
No. of BIN revisions per car (unscaled)	0.031	1.520	692	5584
No. of BO revisions per car(unscaled)	0.002	0.457	692	5584
Markup (*priceprem*)	0.169*	9.361	896	5344
Markup*BO ($priceprem * BO$)	0.121*	8.066	896	5344

+ p<0.10, * p<0.05

Table 1-8. Item-Level Two-Sample t-test with unequal variances Non-BO Mean - BO Mean

Variable	$Mean_{nonBO}$ - $Mean_{BO}$	**t**	N_{nonBO}	N_{BO}
2013 Markup (*priceprem*)	0.103+	1.833	152	744
2016 Markup (*priceprem*)	-0.032	-1.232	917	4427
2013 & 2016 Markup	-0.013	-0.548	1069	5171

+ p<0.10, * p<0.05

Table 1-9. Comparing Repeat Sellers' Usage of BO in 2013 Dataset and 2016 Dataset

Variable	**2013**		**2016**	
	mean	sd	mean	sd
BOperlisting	.759	.418	.805	.377

Table 1-10. Comparing Repeat Sellers' BIN revision behavior in 2013 Dataset and 2016 Dataset

Variable	**b**	**t**		
SelleravgBINrevsunscale	-0.081*	-3.541	312.000	231.000

+ p<0.10, * p<0.05

Table 1-11. Item-Level Make Averages 2013 Mean - 2016 Mean

Variable	$Mean_{2013}$ - $Mean_{2016}$	**t**	N_{2013}	N_{2016}
BO usage per Make	0.001	0.536	1107.000	5584.000
No. of BIN revisions per Make	-0.024*	-68.070	1106.000	5584.000
No. of BO revisions per Make	-0.003*	-22.089	1106.000	5584.000
No. of BIN revisions per Make (unscaled)	0.029*	7.436	1106.000	5584.000
No. of BO revisions per Make (unscaled)	0.002	1.724	1106.000	5584.000
Markup per Make	0.174*	29.029	1099.000	5584.000
Markup*BO per Make	0.126*	25.800	1099.000	5584.000

+ p<0.10, * p<0.05

Table 1-12. Mean Difference in Car Model Averages; 2013 Mean - 2016 Mean

Variable	**b**	**t**
BO usage per Model	0.001	0.225
No. of BIN revisions per Model	0.031*	3.254
No. of BO revisions per Model	-0.003*	-10.299
No. of BIN revisions per Model (unscaled)	-0.024*	-28.385
No. of BO revisions per Model (unscaled)	0.001	0.496
Markup per Model	0.175*	16.586
Markup*BO per Model	00.127*	15.450

+ $p<0.10$, * $p<0.05$

CHAPTER 2
IMPACT OF SELLER PATIENCE AND COMPETITION ON SELLER'S CHOICE TO USE BEST OFFER

2.1 Motivation

Bargaining was previously associated with high transaction costs and scholars predicted the decline of bargaining as a selling/pricing mechanism in the modern economy (Terwiesch et al., 2005). However, e-commerce reduced transaction costs substantially, allowing for platforms where buyers and sellers can negotiate prices to re-emerge. For instance, Priceline.com started the Name-Your-Own-Price (NYOP) system for selling travel tickets and hotel rooms in 1998 (Huang et al., 2013), and eBay introduced the "Best Offer" feature in 2005 (Toklu, 2014). As retail sales conducted online continue to grow rapidly (Anderson et al., 2004), the re-emergence of bargaining platforms in e-commerce markets might indicate a structural shift in retail markets. Therefore, the determinants of the re-emergence of bargaining mechanism, in the context of e-commerce, and its efficacy deserve updated scholarly discussion.

This paper empirically tests for the determinants of a seller's choice to adopt bargaining versus posted prices. Predictions and insights from previous theoretical literature are used, and an original data set that combines data from 2013 and from 2016 is analyzed.

As explained in Chapter 1, most of the previous theoretical economic models studying bargaining vs. posted prices assume that the seller is a monopolist (that is, the models only consist of one seller)(Bing, 2009; Perry, 1986; Wang, 1995). In models where time costs of bargaining are considered (Bing, 2009; Perry, 1986), there are mixed predictions: while according to Bing (2009), a relatively more patient seller will choose bargaining over posted prices

Furthermore, in addition to the gap in literature due to a lack of empirical studies on patience, there is also a gap in bargaining theory literature with respect to studying markets with competition. In particular, only one known theoretical model for bargaining vs. posted prices considers more than one seller (Bester, 1993), and finds a somewhat contradictory result compared to all other studies. That is, according to Bester (1993), while high bargaining power incentivizes the seller to bargain, if the buyer has outside options and can opt out, then the seller is more likely to bargain if he has *lower* bargaining power, in equilibrium. It seems thus that in cases of higher seller competition, buyers have higher outside options, and therefore relatively more bargaining power than they would under lower seller competition, and this can reverse the prediction regarding bargaining. Therefore, I find it crucial to study the impact of seller competition on the adoption of bargaining in the eBay Best Offer context.

However, Bester (1993) does not consider the impact of time, as some of the other models of bargaining do. Hence, in my empirical work, I aim to study both seller patience and seller competition separately, while also studying the impact of the interaction between patience and competition. That is, I am interested in studying the seller's choice to adopt bargaining when individual-level patience interacts with external competitive pressure. While patience increases a seller's bargaining power, it seems that higher competition can reduce his relative bargaining power.

I find that high competition by car Make and high patience, independently, incline a seller against using Best Offer. I also find that high competition by city inclines the seller to use Best Offer more often. In addition, an increase in the interaction term between seller patience and seller competition induces the seller to choose bargaining. This result on the interaction term seems to indicate that the external pressure due to city competition changes the behavior of a patient seller. I discuss the somewhat counter-intuitive result on seller patience, and the bargaining power that stems from it. I find that patience as a bargaining power may be exhibiting itself differently, depending on the structure of the market, in terms of seller competition.

Interestingly, though, while Bester (1993) suggests that in competitive markets, sellers with lower bargaining power are the ones that tend to bargain, he does not take into account seller patience. In the context of eBay, I find that with respect to higher competition by Make, it is true that lower bargaining power (in the form of lower level of patience) induces the seller to choose bargaining, I do not find this to be the case with respect to geographic competition.

That is, it seems that when sellers face higher geographic competition, it impacts their choice of bargaining differently than competition by car Make. Perhaps they are more interested in increasing their buyer arrival rate (consistent with the prediction of Wang (1995)), and the effect of city competition dominates the usual effect of patience, causing even the patient seller to choose to bargain, when seller competition is higher.

And although the re-emergence of bargaining platforms may seem to refute previous wisdom that foretells its decline in modern markets, it is valuable to understand long-term dynamics and stability of bargaining as a selling mechanism. Lu and McAfee (1996) had predicted that bargaining, when compared to auction listings, is not a stable equilibrium. According to them, the differential disadvantage (in the division of surplus) for buyers due to an increase in buyer to seller ratio is lower in auction markets than in bargaining markets. Therefore, using an evolutionary framework, they conclude that the auction mechanism will dominate bargaining in the long run.

Although they do not compare bargaining with posted prices, another paper following McAfee concludes that auctions and posted prices are equivalent with respect to the differential disadvantage mentioned above, and therefore posted price mechanism will dominate bargaining in the long run (Kultti, 1999). Using 2016 data, I test whether the same model that applied well in 2013 can also explain Best Offer usage in 2016, or if behavior changed from 2013 to 2016. I find that the 2013 model, which was inspired from economic theory, does not perform as well in 2016. I find different variables to be of significance in 2016 than in 2013, but some results do hold true across both years.

While I do not find that best offer usage itself is lower in 2016 than in 2013, I find that sellers have differing patterns of learning across the two time periods. I find evidence that they are changing their behavior in accordance with and due to their own learning. I do find that sellers that are indicating higher experimentation with Best Offer usage and learning are also choosing against using Best Offer. This is an interesting result, which not only confirms rapid seller learning on eBay, but also might be capturing market transition towards choice of posted prices vs bargaining. Recent literature has, in fact, indicated that buyer demand is shifting in favor of posted prices (Kultti, 1999). This may therefore explain the change in seller choices as well. I explore the reasons for this possible transition in the next two chapters of my dissertation.

Apart from providing empirical answers to previously raised theoretical questions and from contributing to literature that studies structure of endogenously occurred markets (McAfee, 1993), this paper also complements previous literature on eBay and on the phrase "Best Offer". Although there is extensive literature on eBay Motors and eBay auctions, there are very few studies on eBay "Best Offer". Most of these studies focus on eBay Motors, even though the "Best Offer" feature is available on eBay in other categories. This paper complements these previous works (Huang et al., 2013) (Toklu, 2014) by focusing on eBay Motors as well, but by also allowing higher product (automobile model/make) heterogeneity to gain insights that can be generalized across different cars. In addition, there are two cross-sections of data (from 2013 and from 2016) studied for the same set of car models. This allows for an intermediate long-term study of best offer usage and for studying the impact of seller learning. This is the first empirical study that compares Best Offer usage across years.

Furthermore, despite off-eBay usage of the phrase "Best Offer" being traced as far back as the year 1738 in forum classifieds and its popularity stemming from its usage on Craiglist.com, there is almost no secondary literature on the phrase "or Best Offer." Although this paper aims to study bargaining on eBay Motors, in particular, "or Best Offer" (either on eBay or off eBay) may be a phrase that invokes a specific kind of bargaining. To the extent that "Best Offer" induces a unique type of bargaining, the results of this paper may not extend to all bargaining markets, but the results still contribute to an understanding of the type of bargaining that was prevalent historically and one that re-emerged in online markets.

2.2 Methodology

Given the paucity of literature on and an intriguing re-emergence of bargaining markets, I explore what induces sellers to adopt bargaining markets in the first place. Though there is some indication that certain sources of bargaining power are factors that impact whether or not a seller chooses to bargain, it remains unclear whether, in equilibrium, bargaining is observed when the bargaining power of the seller is high or it is low. That is, the theoretical predictions vary based on modeling assumptions.

Therefore, in this chapter, I focus on empirically understanding which of the factors alluded to in theory contribute to higher usage of Best Offer by a seller, and/or which of these factors incline a seller to enter into a bargaining or Best Offer market vs. entering a posted price only market.

Best Offer usage of a seller, the dependent variable of interest, is measured by the variable, *BOperlisting*, is the proportion of a seller's listings in the sample that are listed as BO. This is a fractional response dependent variable (outcomes take on values [0,1]), and therefore, I use the following fractional logit regression model for analysis:

$$\begin{aligned} E[BO|X] = G(&\lambda_0 + \lambda_1 \bar{X}_j + \lambda_2 S_j + \lambda_3 \bar{L}_j + \lambda_4 \bar{M}_j + \lambda_5 Patience_j \\ &+ \lambda_6 Competition_j + \lambda_7 Patience * Competition_j) \end{aligned} \tag{2–1}$$

where $G(.)$ is the logistic function, known to satisfy $0 < G(z) < 1$ for any real number z. This "ensures that the predicted value is restricted on the interval of [0,1]," and this is an important benefit of using a non-linear quasi-likelihood model over a linear model. If the number of seller listings in the sample were a fixed number, then the literature offers various modifications to the linear model which are suitable. However, when the denominator of the fraction in the response variable is not fixed or known, as is the case here, then other functional form and inference problems arise. In such cases, a non-linear fractional response model is more preferred (Papke and Wooldridge, 1993).

Note that most of the independent variables in the above equation are denoted with bars (except S_j, which denotes seller characteristics). These bars represent that the variable is an aggregate variable at the seller level; that is, item-level characteristics have been aggregated for each seller across all of the seller's items in the sample. *BO* represents the seller's choice of

Best Offer vs. non-Best Offer fixed price listing, measured by the *BOperlisting*. $\bar{X}_j$ are average item characteristics for the seller, $\bar{L}_j$ are average listing characteristics, and $\bar{M}_j$ are average market characteristics for all of the seller's items. For instance,BO_{ij} denotes the seller's best offer usage at the item-level, which is a binary variable, X_{ij} denotes item/car characteristics, L_{ij} are listing characteristics, and M_{ij} are market characteristics at the item-level. These variables are used to calculate the following averages at the seller level: $\bar{BO}$, $\bar{X}_j$, $\bar{L}_j$,$\bar{M}_j$, respectively are used for the primary analysis.

As mentioned earlier, the main predictors of interest in this paper are market competition, patience-level, and the interaction between these two variables. These variables seem to be important game-changers when it comes to the results for a seller's choice to use bargaining vs. posted prices, as mentioned earlier. Therefore, I focus more on these variables, and the parameters of interest are λ_5,λ_6, and λ_7. Other variables suggested in economic theory are also used as control variables and are discussed.

For this paper, I do not intend to obtain structural estimates; rather, I obtain reduced form estimates for the impact of these variables on the seller's use of Best Offer. Note that I do not include the price or markup variable on the right-hand side of the equation because a seller's pricing strategy can be simultaneous with the seller's choice of bargaining. Therefore, including this variable would bias the estimates of the model. I address the issue of this possible simultaneity in Chapter 3. For the purpose of this chapter, instead of studying a simultaneous system of both equations (one explaining markup and one explaining BO usage), I estimate the reduced form estimates for variables that impact a seller's choice of Best Offer only. [1]

For the sake of robustness, I drop sellers that were selling relatively rare cars on eBay. Car models that made up less than 0.5% of the sample in either 2013 or in 2016 are the ones

[1] Where multicollinearity between variables was suspected, an OLS regression between the variables was run to ensure that the R between the two variables did not exceed .5.

classified as "rare". Furthermore, I also used an alternate specification of the model, which included the seller revision variable on the BO feature (to be discussed more later), along with the general revision variable to ensure that the impact of seller revisions on BO specifically has a more robust impact on BO than general revisions.

2.3 Results

2.3.1 2013 Results

Table 2-5 and 3-7 present summary statistics for 2013 and 2016 data, respectively. 2-2 presents the main results for both 2013 and 2016, explaining the dependent variable, *BOperlisting*, the fraction of a seller's listings in the sample that are listed as BO at the end of the listing.

As mentioned earlier, the primary independent variables of interest in this paper are seller competition, seller patience, and the interaction of these two variables. In 2013, with respect to seller patience, measured by *Selleravgrelist*, a one standard deviation increase in seller's relisting behavior increases the log odds of adopting BO by 0.02 points at the mean level of seller competition. This indicates that higher seller patience increases the log odds of adopting BO for a listing. While this seems consistent with theoretical literature that suggests higher patience (leading to higher bargaining power) inclines a seller to choose bargaining over posted prices (Bing, 2009), the magnitude of this impact is not very substantial (i.e. not very different from zero).

However, at the lowest level of seller competition (both in terms of Make and in terms of city), the effect of patience is negative, where seller patience *decreases* the log odds of adopting BO by 1.88 points. At the highest level of competition (both in terms of Make and in terms of city), the effect of patience is to *increase* the odds of adopting bargaining. Now let's consider the impact of one type of competition at a time. That is, if competition by city is at its minimum while competition by make is at its average level, the effect of seller patience is to decrease the the log odds of adopting BO by 1.2 points. If competition by make is at its minimum while competition by city is at its average level, the effect of seller patience is to decrease log odds of adopting BO by 0.67 points. Whereas, if competition by make is at its maximum while competition by city is at its average level, the effect of seller patience is to increase log odds of adopting BO by 1.25 points. If competition by city is at its maximum while competition by make is at its average level, the effect of seller patience is to increase the the log odds of adopting BO by 7.31 points. Therefore, at increased levels of competition, seller patience inclines the seller to adopt BO, whereas at lower levels of competition, seller patience disinclines the seller to adopt BO.

On the other hand, competition variables seem to have a much larger impact (in terms of magnitudes of the coefficients) on BO usage than seller patience does. That is, competition by city increases the log odds of adopting BO by 0.15, and competition by make decreases the log odds of adopting BO by 0.58. Interestingly, while competition by make exhibits the expected behavior due to its negative impact on BO (consistent with Chen et al. (2016)'s empirical finding with respect to competition in terms of Car Model), competition by city does not have a similar impact. That is, while increased information due to higher number of sellers of the same Make decrease the likelihood of sellers adopting bargaining, it seems that in the case of higher seller competition by city, sellers are more likely to choose bargaining. This is consistent with Wang (1995), which suggests that bargaining is likely to be adopted for the sake of increasing the rate of arrival of buyers. It is reasonable that with respect to geographic competition, the information line of reasoning (applicable for competition by Make) is less relevant than the theoretical insight regarding sellers' need to enhance their chances for buyer arrival. Therefore, it is important to consider competition along both dimensions: product market and geographic market.

Finally, I was interested in studying the impact of seller patience in the context of higher competition by studying the interaction term of seller patience with competition. I find that a one standard deviation increase in the interaction term between city competition and seller patience, measured by *competitcity*relist*, increases the log odds of adopting BO by 0.82. Finally, a one standard deviation increase in the interaction term of competition by Make and seller patience increases the log odds of adopting BO by .42 points. Therefore, in both cases, I find that when a patient seller is faced by competition, his log odds of adopting bargaining increase drastically. While patience alone has a very small positive impact on the seller's likelihood to adopt BO, under competition, the seller is more likely to use BO. This result suggests that under higher competition, regardless of whether it is product competition or city competition, the seller is more likely to bargain. This result seems consistent with Bester (1993), which suggests that seller competition weakens seller bargaining power because it allows higher options for buyers. Therefore, we observe bargaining, in equlibrium, where sellers have relatively lower bargaining power due to higher competition among sellers. Although bester1993bargaining does not account for seller patience in the theoretical model, his result holds even in the context of seller patience. That is, the empirical evidence in this paper indicates that individual-level bargaining power due to higher seller patience, when weakened by the external market-driven bargaining power, leads the seller to choose bargaining much more often than when the seller does not face seller competition.

With respect to other control variables, I find that if a seller has more outside options, measured by a one standard deviation increase in *Selleravgunavailable*, then the log odds of adopting BO are decreased by 0.33 points. That is, if the seller has higher bargaining power, in the form of higher outside options, he is *less* likely to choose bargaining. This result is consistent again with the theoretical insights of Bester (1993) and empirical findings of Chen et al. (2016) that higher bargaining power reduces the likelihood of adopting bargaining or Best Offer. This result is further emphasized when looking at the interaction term between the seller's outside options and the competition faced by the seller in terms of Car Make, measured by *competitmake*unavailable*, which is statistically significant and has a positive sign. This implies that when a seller's bargaining power (due to higher outside options) is weakened because he faces higher seller competition (and the buyer has relatively more outside options), the seller becomes more likely to use Best Offer, as consistent with (Bester, 1993).

With respect to the control variables used, I find that if a seller has a higher positive feedback percent by one standard deviation, then the log odds of adopting BO are lower by 2.32 points. This is in contradiction with the expected result for this variable. According to theoretical predictions, a seller with a higher reputation is less likely to use BO (and I find this to be the case in 2016, as will be discussed later), but it is possible that using positive feedback percent alone is not the best variable for measuring a seller's reputation on eBay. It has been cited in a few previous eBay studies that using *net* positive feedback percent (subtracting neutral and negative feedback points from positive feedback points) might be a less ambiguous measure of seller reputation Cabral and Hortacsu (2004); in general, it seems that eBay reputation variables still need further exploration Andrews and Benzing (2007).

If the seller has a higher feedback score, measured by *Feedbackscoreby1000*, then the log odds of adopting BO are lower by 0.22 points. This is consistent with the expectation that a seller with higher experience would be less likely to adopt BO. Finally, also consistent with expectation, if the seller sells newer cars in terms of Model Year(and therefore, there is more information regarding the valuation of those cars), then the log odds of adopting BO are lower by 0.09 points.

2.3.2 2016 Results

In 2016, I find that seller patience *decreases* the log odds of adopting BO by 0.28 points. This magnitude is much greater than that in 2013, which indicated a *positive* but much smaller impact on BO (0.2 points). This result is consistent with other findings from 2013, and with Bester (1993), which suggests that BO is more likely to be adopted when the seller has lower bargaining power. In 2016, I do not find the competition variables to be statistically significant. I also do not find the interaction terms between seller patience and competition to be statistically significant for the 2016 data.

Instead, the variables representing learning and knowledge are more important in 2016. That is, seller's increased learning on his listings, measured by a one standard deviation increase in *selleravgrevisionstotal*, decreases the log odds of adopting BO by 0.17 points. This is consistent with the expectation from relaxing Perry (1986)'s assumptions. That is, it is the lack of knowledge regarding the market and the valuations of other agents in the market that might be inclining a seller to use BO.

An increase in seller's positive reputation, measured by a one standard deviation increase in *PositiveFeedbackPercent*, reduces the log odds of adopting BO by 0.16 points. This result differs from the result found in 2013, as mentioned earlier, but is consistent with the theoretical insight that higher reputation disinclines the seller from using bargaining Riley and Zeckhauser (1983).

If a seller is a private seller, then the log odds of adopting BO are higher by 0.32 points. This is consistent with expectations based on theoretical insight, because a private seller will typically have lower time cost, lower knowledge, lower experience, and is therefore more likely to use BO.

2.3.3 Robustness Checks

For robustness, I drop sellers who were selling rare cars (i.e. cars that made up less than 0.5% of either the 2013 sample or the 2016 sample) to check if the primary results of the paper might be driven by rare/extreme/collectible cars. As can be seen in Table 2-3, primary results from 2013 and 2016 both remain unaffected when sellers of rare cars are dropped from the sample.

I also run an alternate specification, which includes the BO revision variable in the specification, in addition to the general revisions variable used in other specifications. I do this to check if it is revisions/learning specific to BO that impact a seller's choice to adopt BO. I find that general learning/revisions and BO-specific learning/revisions both have a negative impact on usage of BO in 2016, and that neither have an impact in 2013. This confirms that sellers that revise more often, and therefore learn more, especially with regard to using BO, tend to revise their listings against using BO.

2.4 Discussion

In conclusion, according to the 2013 data, higher patience inclines a seller to use bargaining more often (by log odds of 0.02), while using the 2016 data, higher patience *disinclines* a seller from choosing bargaining. The 2016 result is more convinving because it has a magnitude substantially different from zero, and is consistent with other results in this paper that indicate that higher bargaining power inclines a seller to choose fixed prices over using Best Offer (consistent with Bester (1993); Chen et al. (2016)). However, it is important to note that the reversal of signs from 2013 and 2016 could, in fact, be due to changes in seller behavior over the time period from 2013 to 2016.

With respect to seller competition, in 2013, I find that increased competition by Car Make makes the seller less likely to use bargaining (consistent with the findings of Chen et al. (2016)). However, increased seller competition on eBay, in terms of city competition, *increases* the log odds of adopting BO, consistent with the insight from Wang (1995) that bargaining is more likely to be adopted when a seller wants to increase the rate of buyer arrival. Using the interaction term of seller patience and seller competition, in 2013, it becomes evident that increased competition (regardless of whether it is product competition or geographic competition) increases the log odds of a patient seller to choose to bargain. This is consistent with the emphasis of Bester (1993); Chen et al. (2016) that bargaining is chosen by a seller when a seller's bargaining power is weakened due to more outside options for the buyers.

In 2016, however, I find no impact of competition variables on the seller's choice to bargain.[2] [3] . More studies, and especially long-term studies, on eBay Best Offer are needed for confirming the sources of differences between the 2013 and 2016 results with respect to the impact of seller competition. These differences might be due to inter-sample individual differences, data collection differences, or due to changes in seller behavior over this time period.

Combined, all of these results suggest that higher bargaining power disinclines a seller from using bargaining or Best Offer. That is, higher patience (in 2016 data), more outside

---

[2] It is possible that competition variables are not statistically significant in the 2016 data due to a few differences in data collection from 2013 and 2016. That is, selection of specific makes and the glitches in extraction, that caused the geographic representation of the data to be skewed towards listings closer to Florida, (as mentioned in the data section) may have led to an imperfect representation of the eBay market in terms of the city and make competition, contaminating the competition variables in 2016. However, as the means tests presented and discussed in Chapter 1 suggest, there are indeed other differences between 2013 and 2016, especially related to revision behavior and markups, that might also explain the change in results observed from 2013 and 2016. In addition, as will be discussed in the next two chapters of the dissertation, sellers' revision trends against using bargaining and the reduction of markups are consistent with my findings in Chapter 3 that pricing higher (and perhaps even using best offer), viewed independently, leads to a reduction in sales. With that in mind, it seems that seller may in fact be learning and revising their behavior somewhat in accordance with the true effects of markup and BO on likelihood of sale. With the rapid pace of change that is possible due to internet and e-commerce, this observable difference from 2013 to 2016 is a finding worth noting.

[3] I also used a variable, competition by Make on eBay Motors (in general); that is, this alternate competition variable did not represent the competition *within* the sample of observations (as is the case with the current competition measure used). Rather, this variable measured the competition by Make on eBay Motors by counting the number of all cars (not just the ones in my sample) for each make that were listed on eBay on a given day. I also did not find any impact of competition using this alternate competition variable. However, this measure was collected a about 10 days after the collection of the completed listings used for the sample. So, it is possible that the time period of the listings in the sample did not match this competition variable very well; especially if the eBay car market competition changes very rapidly. Further exploration of competition variables on eBay would enhance the understanding for the 2016 results

options for the seller (in both 2013 and 2016 data), and higher competition by Make (in 2013 data) make a seller less likely to choose bargaining over posted prices. These conclusions are consistent with the counterintuitive empirical findings of Chen et al. (2016) regarding eBay Best Offer (even though they use different measures/variables for bargaining power), and with the insights of Bester (1993) regarding a seller's choice between bargaining and posted prices when there is more than one seller in the market.

Of course, these results could be true specifically for "eBay Best Offer" type of bargaining. That is, perhaps weaker sellers bargain through the "Best Offer" signal on eBay, whereas stronger bargainers may be bargaining *off* eBay and keeping their price/markup high on eBay. Keeping a high markup on eBay may serve as another type of bargaining strategy, one in which a seller might attempt to misrepresent the value of his car in order to be a stronger bargainer for a possible off-ebay bargaining game. This is reasonable, especially in the car market setting, because the sellers may have a general understanding that a negotiation (on or off-eBay) may ensue regardless of whether the seller indicates his willingness to negotiate through the Best Offer feature. That is, the line between bargaining vs. posted prices may not be entirely determined by Best Offer usage alone, although Best Offer still serves to distinguish the sellers who are in fact *signaling* their willingness to negotiate. A strong bargainer may not openly signal his willingness to bargain, while simultaneously posting a price. For instance, a patient seller may be exercising some bargaining power by not indicating this willingness to negotiate. This notion is supported in my other findings, discussed in the next chapters, that a patient seller also charges a higher markup/price than a less patient seller. Therefore, perhaps the conclusion on higher bargaining power (along with other conclusions made on eBay Best Offer) inclines a seller not to choose to bargain using Best Offer. This does not exclude the seller from off-eBay forms of bargaining. To know whether this is true off-eBay as well, more research is needed for off-eBay bargaining of car dealers. Nevertheless, with respect to Best Offer-type of bargaining, it can be said that higher bargaining power disinclines a seller from choosing Best Offer, in favor of using the fixed price mechanism. These findings (using new variables) corroborate the support of Chen et al. (2016) for the theoretical insights from Bester (1993).

For the control variables, in 2016, a private seller (vs. a car dealer) is more likely to use bargaining vs. posted prices. In addition, if a seller has higher number of revisions(learning) on his listings, then he is less likely to use Best Offer.[4] In 2013, I find that a higher feedback score for the seller, which represents higher seller experience on eBay, disinclines a seller from using bargaining. Together, these results confirm the insights of previous theoretical literature that bargaining is more likely to occur when seller knowledge and experience regarding the market is lower, and that as learning and experience increases (measured by higher number of revisions, or by the seller being a dealer, or by his feedback score), a seller prefers the fixed price method for selling.

Interestingy though, the impact of a seller's reputation on his choice to bargain is somewhat ambiguous. According to the 2016 data, a higher positive feedback percent leads to a decline in Best Offer usage, but according to 2013 data, a higher positive feedback percent increases Best Offer usage. While the 2016 result is more reliable (and more consistent with theory) because the 2016 sample is larger, the 2013 result is still noteworthy. This surprising result regarding the impact of seller reputation on the choice of Best Offer, in the 2013 data, needs further exploration in empirical work. As mentioned earlier, previous eBay studies have highlighted some ambiguity regarding the best measure for seller reputation on eBay, and perhaps this surprising 2013 result and the discrepancy between the 2013 and 2016 results can be reconciled by using a *net* positive variable (as suggested by a few previous studies).

[4] These two variables, *sellerprivate* and *selleravgrevisions*, were not statistically significant for the 2013 data.*sellerprivate* variable was included in other specifications (not presented in this paper) for the 2013, but did not result in being statistically insignificant. It was therefore not included in the 2013 model presented in this paper because including this variable for 2013 reduces the number of observations for this already small 2013 sample.

It is interesting to note some of the differences between the 2013 and 2016 results. For instance, it is obvious that learning and knowledge variables are more significant in explaining BO usage in 2016 than in 2013. Furthermore, there is a larger impact of seller patience on the adoption of BO in 2016 than in 2013, while seller competition variables are not significant in 2016 (but are significant in 2013). In general, since the 2016 results are based on a much larger sample than the 2013 data, I rely more on the 2016 results. However, the 2013 results are noteworthy because they might be indicating changes in seller behavior, either due to sample differences or due to seller learning or market dynamics over time. That is, when comparing 2013 and 2016 results, there are three primary sources for reconciling differing results: inter-sample individual differences, data collection differences, and behavior changes over time. More long-term empirical studies on eBay Best Offer that improve upon the first two methodological limitations of this paper are needed to determine whether it is the third source, behavior changes over time, to which the differing conclusions can be attributed.

The results of this chapter highlight the primary determinants of a bargaining platform, with a special focus on understanding the role and interaction of seller patience and seller competition. In summary, the emergence of Best-Offer type bargaining platforms can be expected when sellers have weaker bargaining power (in the form of lower patience or lower outside options or higher seller competition). Using 2013 data, I find that the effect of sellers' individual-level patience varies with the level of competition in a seller's external environment. I find that at high levels of competition, a patient seller chooses to bargain, while at low levels of competition, a patient seller chooses not to bargain. In a more competitive environment, a seller loses bargaining power relative to buyers as the number of outside options for buyers increases. Hence, seller patience, a form of higher bargaining power, is weakened substantially by higher seller competition. This weakened bargaining power inclines even the more patient sellers to choose to bargain. This result can have strong implications for other cases in which individual-level patience bows to external competitive pressures, and therefore deserves further attention. Note that using different variables for bargaining power in this paper (i.e.

patience) than in Chen et al. (2016), I find support for their claim that sellers with lower bargaining power are more likely to use Best Offer (consistent with Bester (1993)). In addition, this chapter also finds that if sellers are less experienced or have less knowledge about the valuation of the product or about the valuations of the other agents, bargaining is more likely to occur. It is interesting to note, though, that while sellers with higher information and higher knowledge tend not to choose bargaining, it is not the case that bargaining declines over time, in general, on eBay Motors. This might be because there are continually new entrants in the eBay Motors markets, who try to gather information regarding the market. Therefore, bargaining may be more likely to occur in markets that regularly witness new entrants. Also, since the eBay Motors market sells a sufficient percentage of used cars, whose valuation is difficult, bargaining remains a desirable selling mechanism for the purposes of price determination. While some of these findings were predicted by theoretical insights, this paper helps in identifying which theoretical insights hold more weight in the context of eBay Best Offer.

2.5 Tables

Table 2-1. Hypotheses and Variables

Bargaining vs. Posted Price Theory	Variable (Expected Sign)
Seller patience is used for signaling high valuation for an item; lending bargaining power during a negotiation (Bing, 2009), but in equilibrium, it may be sellers with lower bargaining power that choose to bargain (Bester, 1993)	Selleravgrelist (?)
If a seller is patient and faces high competition, then he has higher bargaining power through patience, but relatively lower bargaining power due to higher competition (Bester, 1993).	competitcity*relist (?) competitmake*relist (?)
Low time cost for the seller is the real source of bargaining power (Bing, 2009).	sellerprivate(+) SelleravgQA(+) sellerlistingsN(-) Selleravgunavailable(-)
Lack of knowledge regarding buyers' valuations in the market may lead to choice of bargaining (Relaxation of Perry (1986)'s assumption).	sellerprivate(+) SelleravgMileage(+) SelleravgYear(-) Selleravgrevisionstotal(-) sellerprivate(+) countmakeSeller(-) countcitySeller(-) Feedback Score(-)
Lack of knowledge regarding a seller's time costs in the market may lead to choice of bargaining (Relaxation of Perry (1986)'s assumption).	Selleravgrevisionstotal(-) sellerprivate(+) countmakeSeller(-) countcitySeller(-) Feedback Score(-)
Lack of knowledge regarding buyers' time costs in the market may lead to choice of bargaining (Relaxation of Perry (1986)'s assumption).	Selleravgrevisionstotal(-) sellerprivate(+) countmakeSeller(-) countcitySeller(-) Feedback Score(-)

Table 2-1. Hypotheses and Variables (Table Continued)

Bargaining vs. Posted Price Theory	**Variable (Expected Sign)**
Lack of knowledge regarding sellers' valuations in the market may lead to choice of bargaining (Relaxation of Perry (1986)'s assumption).	Selleravgrevisionstotal(-) sellerprivate(+) countmakeSeller(-) countcitySeller(-) Feedback Score(-)
A higher ability to commit (or higher reputation) leads to lower bargaining usage (Riley and Zeckhauser, 1983)	Positive Feedback Percent (-)
Higher number of items to sell leads to choosing posted price over bargaining (Wang, 1995).	sellerlistingsN (-)
If a seller has more outside options, then he has higher bargaining power.	Selleravgunavailable (?)
If a seller has more outside options and faces high competition then he has higher bargaining power.	competitcity*unavailable (?) competitmake*unavailable (?)

Table 2-2. BOperlisting: Fractional Logit

Concept Variable	Variable Measure	2013	2016
Seller Patience	Selleravgrelist	-2.151**	-0.603**
		(0.748)	(0.307)
Patience*Competition	competitcity*relist	0.198*	0.006
		(0.116)	(0.012)
Patience*Competition	competitmake*relist	0.014*	0.003
		(0.008)	(0.002)
Seller City Competition	countcitySeller	-0.023	-0.005
		(0.015)	(0.008)
Seller Make Competition	countmakeSeller	-0.017**	-0.002
		(0.007)	(0.002)
Seller's Outside Options	Selleravgunavailable	-2.446**	-0.082
		(0.774)	(0.420)
Outside Options* Competition	competitcity*unavailable	0.007	0.025
		(0.063)	(0.020)
Outside Options* Competition	competitmake*unavailable	0.023*	-0.003
		(0.014)	(0.003)
Seller Reputation	Positive Feedback Percent	0.011**	-0.037**
		(0.006)	(0.017)
Seller Learning	Selleravgrevisionstotal	0.043	-0.180**
		(0.084)	(0.068)
Seller Experience	Feedbackscoreby1000	-0.068**	-0.012
		(0.028)	(0.008)
Seller's Time Cost	sellerlistingsN	-0.001	0.002
		(0.018)	(0.006)
Seller's Time Cost/Knowledge/Experience	sellerprivate		0.320*
			(0.165)
Seller Communication	SelleravgQA	0.400	
		(1.324)	
Object Valuation Difficulty (Information)	SelleravgMileageby1000	-0.001	0.000
		(0.004)	(0.002)
Object Valuation Difficulty (Information)	SelleravgYear	-0.094*	-0.007
		(0.050)	(0.026)
Object Expensiveness	SelleravgIMVby1000	0.017	-0.004
		(0.013)	(0.004)
cons		191.148*	19.913
		(100.549)	(52.210)
N		428.000	1417.000
r2		0.1523	0.0522

* p<0.10, ** p<0.05
Standard Errors are reported in the parentheses.
Seller Make Dummies included.

Table 2-3. BOperlisting: Fractional Logit; & Collectible/Rare Cars Dropped

Concept Variable	Variable Measure	2013	2016
Seller Patience	Selleravgrelist	-2.269**	-0.659**
		(0.768)	(0.318)
Patience*Competition	competitcity*relist	0.199*	0.016
		(0.119)	(0.013)
Patience*Competition	competitmake*relist	0.015*	0.003
		(0.008)	(0.002)
Seller's Outside Options	Selleravgunavailable	-2.507**	-0.097
		(0.797)	(0.455)
Outside Options*Competition	competitcity*unavailable	0.011	0.024
		(0.065)	(0.020)
Outside Options*Competition	competitmake*unavailable	0.024*	-0.002
		(0.014)	(0.003)
Seller Learning	Selleravgrevisionstotal	0.040	-0.220**
		(0.084)	(0.068)
Seller Reputation	Positive Feedback Percent	0.012**	-0.039**
		(0.006)	(0.018)
Seller Experience	Feedbackscoreby1000	-0.070**	-0.008
		(0.030)	(0.007)
Seller City Competition	countcitySeller	-0.024	-0.005
		(0.015)	(0.008)
Seller Make Competition	countmakeSeller	-0.019**	-0.002
		(0.007)	(0.002)
Seller Communication	SelleravgQA	0.557	
		(1.469)	
Object Valuation Difficulty (Information)	SelleravgMileageby1000	-0.001	0.000
		(0.004)	(0.002)
Object Expensiveness	SelleravgIMVby1000	0.016	-0.004
		(0.012)	(0.004)
Object Valuation Difficulty (Information)	SelleravgYear	-0.094*	-0.010
		(0.052)	(0.026)
Seller's Time Cost	sellerlistingsN	-0.003	0.001
		(0.018)	(0.006)
Seller's Time Cost/Knowledge/Experience	sellerprivate		0.334**
			(0.167)
cons		192.351*	26.684
		(105.363)	(53.006)
N		421.000	1347.000
r2		0.1658	0.0465

* $p<0.10$, ** $p<0.05$

Standard Errors are reported in the parentheses.

Seller Make Dummies included.

Table 2-4. BOperlisting: Fractional Logit; BO Revision Variable Included

Concept Variable	Variable Measure	2013	2016
Seller Patience	Selleravgrelist	-1.111**	-0.333**
		(0.390)	(0.170)
Patience*Competition	competitcity*relist	0.096**	0.002
		(0.046)	(0.006)
Patience*Competition	competitmake*relist	0.007*	0.002
		(0.004)	(0.001)
Seller Competition	countcitySeller	-0.013	-0.003
		(0.008)	(0.004)
Seller Competition	countmakeSeller	-0.009**	-0.001
		(0.003)	(0.001)
Seller's Outside Options	Selleravgunavailable	-1.394**	-0.101
		(0.424)	(0.228)
Outside Options*Competition	competitcity*unavailable	0.002	0.014
		(0.030)	(0.010)
Outside Options*Competition	competitmake*unavailable	0.013*	-0.001
		(0.007)	(0.001)
Seller Reputation	Positive Feedback Percent	0.006**	-0.019**
		(0.003)	(0.009)
Seller Learning (General)	Selleravgrevisions	0.020	-0.071*
		(0.048)	(0.040)
Seller Learning (BO-Specific)	SelleravgBOrevsunscale	-0.023	-0.862**
		(0.539)	(0.236)
Seller Experience	Feedbackscoreby1000	-0.037**	-0.008
		(0.013)	(0.006)
Seller's Time Cost	sellerlistingsN	-0.001	0.001
		(0.012)	(0.003)
Seller's Time Cost/Knowledge/Experience	sellerprivate		0.174*
			(0.090)
Seller Communication	SelleravgQA	0.213	
		(0.690)	
Object Valuation Difficulty (Information)	SelleravgMileageby1000	-0.000	0.000
		(0.002)	(0.001)
Object Valuation Difficulty (Information)	SelleravgYear	-0.050*	-0.011
		(0.026)	(0.014)
Object Expensiveness	SelleravgIMVby1000	0.009	-0.002
		(0.006)	(0.002)
cons		100.858**	24.918
		(51.382)	(28.779)
N		428.000	1417.000
r2		0.1505	0.0666

* $p<0.10$, ** $p<0.05$

Standard Errors are reported in the parentheses.

Seller Make Dummies included.

Table 2-5. Summary Statistics 2013

Variable	Obs.	Mean	SD	Min	Max
BOperlisting	428	0.827	0.370	0.0	1.0
Selleravgunavailable	428	0.095	0.271	0.0	1.0
competitcity*unavailable	428	0.819	3.803	0.0	44.0
competitmake*unavailable	428	4.971	19.774	0.0	139.0
Selleravgrelist	428	0.186	0.354	0.0	1.0
competitcity*relist	428	1.213	4.191	0.0	44.0
competitmake*relist	428	10.952	30.217	0.0	139.0
Selleravgrevisions	428	0.679	1.807	0.0	14.7
PositiveFeedbackPercent	428	94.580	21.113	0.0	100.0
countcitySeller	428	7.397	10.868	1.0	44.0
countmakeSeller	428	53.544	47.580	5.0	139.0
SelleravgQA	428	0.015	0.114	0.0	1.0
SelleravgIMVby1000	428	24.438	27.663	4.3	245.0
SelleravgMileageby1000	428	68.219	51.208	0.0	335.5
Feedbackscoreby1000	428	0.590	3.267	0.0	64.7
SelleravgYear	428	2006.772	4.443	1986.5	2014.0
sellerlistingsN	428	2.327	4.680	1.0	85.0
sellerprivate	373	0.252	0.435	0.0	1.0

Table 2-6. Summary Statistics 2016

Variable	Obs.	Mean	SD	Min	Max
BOperlisting	1861	0.829	0.371	0	1
Selleravgrelist	1861	0.378	0.465	0	1
competitcity*relist	1861	3.370	10.175	0	78
competitmake*relist	1861	54.860	89.242	0	326
Selleravgunavailable	1861	0.156	0.332	0	1
competitcity*unavailable	1861	1.432	6.520	0	78
competitmake*unavailable	1861	18.828	50.773	0	326
Selleravgrevisions	1861	0.278	0.928	0	9
PositiveFeedbackPercent	1586	99.165	4.386	33	100
Feedbackscoreby1000	1861	0.439	2.620	-0	96
countcitySeller	1861	8.662	15.331	1	78
countmakeSeller	1861	129.557	94.954	1	326
SelleravgQA	1739	0.021	0.141	0	1
SelleravgMileageby1000	1861	70.995	62.255	0	1350
SelleravgIMVby1000	1861	23.959	24.858	0	246
SelleravgYear	1861	2008.656	4.048	1986	2016
sellerlistingsN	1861	2.585	9.956	1	298
sellerprivate	1665	.563	.496	0	1

CHAPTER 3
SIMULTANEITY BETWEEN CHOICE OF MARKUP AND CHOICE OF BEST OFFER

3.1 Motivation

As mentioned in previous chapters, there is only one known study that empirically tests a seller's choice of bargaining vs. posted prices on eBay, Chen et al. (2016). They study the question of whether sellers that use BO also charge higher prices. In their analysis, the authors do not consider the possible simultaneity between a seller's choice to bargain and a seller's choice of markup. Instead, they seem to assume implicitly that for all sellers, the choice of selling method comes first, and that the seller's pricing decision comes second. They do not take into account the fact that perhaps the a seller's choice of markup precedes (and leads to) his choice to adopt BO. Therefore, they use an OLS model to test the impact of BO on the listed price or the BIN price. As a result, they conclude that sellers that use BO also ask for a higher price, as they had expected. Their approach has two drawbacks: 1) it implicitly rules out, without justification, a theoretically possibility that a seller with a higher markup may be more inclined to use Best Offer (in the interest of price discrimination, and thereby enhancing consumer surplus) than a seller with a lower markup would be, and 2) it does not account for the endogeneity inherent in the seller's choice to adopt BO. Due to this latter drawback, their results only present the "average treatment affect on the treated," but does not present the "average treatment effect," where the "treatment" is the BO selling mechanism.

Furthermore, despite the common notion that in bargaining markets (especially the informal markets in developing countries), sellers charge a higher markup because of their choice to bargain, there is not much theoretical discussion on this subject. This notion, while common in certain bazaar settings, deserves a proper evaluation in the context of eBay bargaining. Hence, I do not assume that all sellers that use BO use it in a similar manner. Instead, I account for the endogeneity of a seller's choice to use BO, when evaluating whether or not a seller charges a higher price simply because he chooses to use BO. I also study the possible impact of a seller's markup on his choice to adopt BO (something I did not do in the previous two chapters where I used the reduced form approaches, due to the endogeneity of markup and BO). Therefore, I study the impact of a seller's optimal choice of markup and his optimal choice of BO (vs. non-BO), taking into account the co-dependency and simultaniety of these choices, using an IV approach.

I then compare the results of this IV approach with an OLS approach, similar to the approach used by Chen et al. (2016). The purpose of this comparison exercise is to highlight the differing results that are produced by two approaches. It is plausible that results of Chen et al. (2016) might still hold, because perhaps it is seller behavior that has changed over time, and this change in seller behavior (and therefore, the change in datasets) is the reason for differing answers to a similar question. Nevertheless, it is worth considering the IV approach for studying this question because of the endogeneity concerns in the seller's choice to bargain and his choice of markup.

I test the impact of the seller's use of BO on the listed markup using the OLS model, similar to Chen et al. (2016). Furthermore, to illustrate the point regarding simultaneity consideration in its entirety, I also run a similar model for explaining the impact of seller markup on BO. It is easily conceivable that other papers might also attempt to test whether it is sellers with higher markups that are more likely to use Best Offer, as a means for *reducing* the markup (in the interest of the buyers).

Although I present both of these equations, they are no longer to be considered part of one system, as was the case for the IV approach analysis. Instead, for the case of the OLS models, each equation must be viewed independently, and not in relation to the other, and must be compared to its corresponding equation from the IV approach.

Using a rich and unique dataset from eBay motors, I use a seller's past level of learning or experimentation regarding bargaining and a seller's past learning or experimentation regarding the price premium as instrumental variables for his current bargaining and current price premium strategies. As a result of this IV approach, I find that sellers that use bargaining tend to charge a *lower* price premium, contradictory to the common notion and contrary to the conclusions of Chen et al. (2016). I do not find that sellers with a higher markup are more likely to use BO though. Therefore, I do not find evidence that sellers are acting in accordance with this consumer surplus-enhancing channel. However, the finding that BO usage actually leads sellers to charge a lower price certainly rules out the negative-impact channel that bargaining may be notoriously known for. That is, at the minimum, the choice of bargaining is not associated with charging higher markups, and is therefore not affecting consumer surplus negatively from these two channels.

Since this result differs from the result in Chen et al. (2016), I offer an explanation for the difference in results by comparing my methodology with the methodology in Chen et al. (2016). While the differences in the results can also be explained by data differences and changes in eBay seller behavior over the course of time from data Chen et al. (2016), the differences in methodology also might have a role to play, and are therefore noteworthy.

3.2 Empirical Methodology

Assume that sellers are profit-maximizing, and that they choose an optimal level of markup and an optimal selling strategy (either Best Offer or not Best Offer). However, the choice of a markup and the choice of Best Offer are simultaneously determined to some extent, as explained above. There are several factors that affect the sellers optimizing decisions. For instance, consider seller patience or impatience, sellers' knowledge of the value of the items, and the sellers beliefs about how informed buyers are. In addition, there is likely some concern in the market regarding the quality of the cars being sold, consistent with Akerlof (1995). I discuss below the theoretically possible uses of the combination of BO and markup by sellers, depending on their patience, knowledge, and perceptions about buyers.

The markup variable measures the proportion of seller's asking price in excess of the retail estimate of the car's true value, as follows:

$$Markup_{ij} = \frac{BIN_{ij} - IMV_{ij}}{IMV} \tag{3–1}$$

where BIN is the Buy-it-now price that each seller on eBay must post if they are using a fixed-price listing format, whether they choose to use BO or not, and where the Instant Market Value (IMV) for the car is the retail estimate of a given car, which can be obtained from CarGuru.com using the Vehicle Identification Number (VIN) listed on eBay for that car.

In the context of bargaining for automobiles on eBay, sellers may charge a higher markup for a variety of reasons. One reason for charging a higher markup is to signal to buyers that the car is of a higher quality , since car quality is not observable merely through an online listing. Or, it could also be that a higher markup may be a form of "cheap talk," where a seller tries to signal that a car is of high value, but it may not really be the case. This "cheap talk" mechanism is similar to the one explained in Croson et al. (2003), which allows for the markup on an item to be conceived as the pie of surplus to be shared between the consumer and the seller. Several experimental studies in Economics have studied the outcomes of an ultimatum bargaining game when the size of the pie is unknown to one agent in a two-agent interaction. These studies find that if the agent that has knowledge of the pie size proposes a division of the surplus, he often gets a higher share of the surplus than the responding agent Croson et al. (2003). This insight can be applied for understanding a seller's use of a higher markup. That is, the size of the surplus "pie" in this case is likely known most closely by the seller who lists an asking price for the car. Hence, the buyer does not know the size of the surplus that needs to divided if a bargaining game were to ensue. It is plausible then that sellers that list a higher markup may be lying regarding the implied value of the car. Experimental studies have argued that "cheap talk" of this sort can be advantageous. It is plausible then that a higher markup (or misrepresentation of the pie) in a manner that misleads consumers can result in tangible benefits for the seller. Another type of "cheap talk" or reason for charging a higher markup is if a seller is truly uninformed about the value of an item and does not want to undervalue it with a low markup. Therefore, it is worth considering whether the sellers on eBay Motors tend to use a combination of their choice of BO (or not) and markup in this manner. That is, while a seller may indeed be using a higher markup to signal higher quality, he may also be lying in order to reap the benefits of imperfect information regarding the product (and this is especially true for the case of used cars), and it is worth considering this aspect of seller strategy when studying sellers' optimizing behavior.

I now discuss the various equilibrium combinations of markup and BO that we might observe in the eBay market. Through this discussion, it can be seen that is no one clear prediction regarding the relationship between seller markup and seller's choice of BO, yet the interaction of these two components of a seller's decision seems inevitable. I summarize these possible equilibrium combinations in Table 3-1 and Table 3-2. The general intuitive principles that guide the construction of these tables is as follows: I expect that a more patient and more informed seller with a high quality product is likely to keep his markup high, and choose not to bargain, especially if he thinks that buyers are well informed. However, if the seller is patient or impatient, he may choose to engage in bargaining and perhaps even lower the markup in order to sell faster.

Table 3-1 and Table 3-2 present some possible scenarios that could occur with respect to the seller's optimal choice of markup and his optimal selling mechanism. While Table 3-1 indicates some possible strategies for the seller when the buyers are well-informed, Table 3-2 indicates possible strategies for the seller when buyers are not well informed, and sellers might therefore benefit from engaging in cheap talk. These tables indicate the likely combinations of seller's optimal strategies for choosing markup and BO depending on whether a seller is patient (or impatient enough to choose bargaining or "very impatient" such that he is willing to accept a price lower than his true valuation), whether he is informed regarding the value and quality of his product, and whether the product itself is of higher quality.

Note first that in both Table 3-1 and in Table 3-2, uninformed sellers are unaware of the true value/quality of the product, so there is no particular signal that they can send regarding product quality. Therefore, the behavior of an uninformed seller does not change depending on product quality. When sellers believe that buyers are well informed (see Table3-1), they are likely to choose a higher markup and allow for BO, allowing for the buyer to bargain for a reduced price. If the seller is not patient, then he chooses to price lower, and also use BO when he perceives that buyers are well-informed. If the seller believes that many buyers are not well-informed (see Table 3-2), he behaves similarly as the case where he believes buyers are well-informed.

Now consider the case of a patient and informed seller selling a high quality car in a market where buyers are well informed (see Table 3-1). This seller chooses to price high because the quality of his product is high and he is well informed about its quality and worth. Since he is patient, he chooses not to use BO (supported by evidence from Chapter 2). Given that seller patience and seller information lend higher bargaining power according to economic theory, and that the seller with such characteristics who knows that his product is of high quality, he would choose not to bargain. This is consistent with the theoretical insights of Bester (1993), and with the empirical findings from Chapter 2.

Notice also that when the seller is impatient but well-informed that he has a higher quality product, he is likely to post a high markup and also use BO. However, in all other cases, an impatient seller tends to charge a lower markup because either he is selling a lower quality product or is unaware of the product quality, but he also uses BO in order to try and sell sooner. Now consider a "very impatient" seller that may be willing to sell lower than his true valuation of the car. For example, an individual who is being transferred by his employer to a new location who cannot take his automobile along may be very impatient. This type of seller may charge a high markup when he knows he has a high quality product, and uses BO, but he may be willing to charge a lower markup in order to sell faster. When the very impatient seller knows that he has a lower quality product, he is likely to charge a lower markup and also use BO. When he does not know if he has a lower product, but he knows that buyers are perfectly informed, he chooses a lower markup and to use BO, as is the case with an impatient seller with a lower quality product.

Finally, focusing now on the case where the seller believes that buyers are less informed and his product is of a lower quality, the seller has an incentive to "cheap talk". That is, even when the seller knows that his product is of a lower quality, but he is patient and well-informed, he might try to appear as though he has a higher quality product. He can do so by charging a higher markup and by not using BO. In cases where he is impatient and informed, but thinks that buyers are less informed, he may want to charge a higher markup and use BO[1] . In the case where he is very impatient it is unclear what markup he would set, but is likely to use BO. On the other hand, an informed seller with a high quality product who believes buyers are not well informed does not behave differently than a similar seller who believes that buyers are well-informed.

These are some scenarios that may be observed, in equilibrium, with respect to seller's combination of markup and selling strategy. Since there are so many possible scenarios, but only four sets of possible messages/indications that the seller can provide, there is no clear theoretical prediction regarding whether the choice of BO will come first or the choice of markup will come first nor is there a clear indication regarding which sort of relationship between markup and BO will be more dominant among sellers. It is unclear whether BO will likely be associated with a higher markup or with a lower markup. Therefore, empirical examination of how these two variables interact can help shed insight into how sellers are actually behaving on eBay Motors.

Consider now the following first order conditions that affect a seller's decision for whether or not to use BO, and what markup to use, respectively:

$$\begin{split} BO_{ij} = \beta_0 + \beta_1 X_{ij} + \beta_2 S_{ij} + \beta_3 L_{ij} + \beta_4 M_{ij} + \\ \beta_5 Markup_{ij} + \beta_6 \gamma_i \end{split} \tag{3–2}$$

[1] In the case of a higher quality product, a seller may choose between a higher markup and BO or low markup and no BO if he is impatient and believes that buyers are less informed.

$$
\begin{aligned}
Markup_{ij} = \alpha_0 + \alpha_1 X_{ij} + \alpha_2 S_{ij} + \alpha_3 L_{ij} + \alpha_4 M_{ij} \\
+\alpha_5 BO + \alpha_6 \gamma_i
\end{aligned}
\tag{3–3}
$$

BO_{ij} denotes the seller's best offer usage, which is a binary variable, and $Markup_{ij}$ is the seller's markup over a price estimate, or Instant Market Value (IMV) for the car, obtained from Car Guru.

X_{ij} denotes item/car characteristics, S_{ij} denotes seller characteristics (years of experience on eBay, behavioral indication of patience, seller reputation, L_{ij} are listing characteristics, and M_{ij} are market characteristics. γ_i are Car Make fixed effects.

However, there are two problems in estimation of the parameters. Firstly, note that the true value of a car, denoted below by *TrueValue*, may be different than the CarGuru.com retail estimate for the car. That is, consider that there is some measurement error in the IMV of the car:

$$
TrueValue = IMV + \sigma \tag{3–4}
$$

$$
TrueMarkup = \frac{\mathrm{BIN}_i - \mathrm{IMV}_i - \sigma}{\mathrm{IMV}_i + \sigma} \tag{3–5}
$$

Since σ is unobserved, the observable variable *Markup* used in the analysis has some error, z, such that:

$$
TrueMarkup = Markup + z = \frac{\mathrm{BIN}_i - \mathrm{IMV}_i - \sigma}{\mathrm{IMV}_i + \sigma} \tag{3–6}
$$

where $z = \frac{\sigma * BIN}{IMV(IMV+\sigma)}$.

Therefore, in using the variable *Markup*, I make the assumption that $\mathrm{E}(z\mathbf{X}) = 0$, where $\mathbf{X}$ is a vector of all the independent variables on the RHS of the model predicting BO. To the extent that z, the difference between the *TrueMarkup* and *Markup* (which is mainly driven by σ, the measurement error in the retail car estimate) can have an impact on the seller's choice to use BO (controlling for $\mathbf{X}$), this assumption would not hold, and the OLS estimates would be biased but still consistent.

Secondly, the simultaneity between *Markup* and *BO* poses a parameter identification problem. This problem is especially difficult in the simultaneous equations of interest because the pricing strategy and the selling strategy are inextricably linked; the same right-hand side variables that typically determine the selling strategy also seem to determine the markup or the pricing strategy of the seller. To solve this identification problem, there must be two sets of variable(s): one that determines the selling method (equation (1)), but not the pricing method (equation (2)), and another variable that determines the pricing method, but not the selling method. These variables can serve as instrumental variables that help identify the parameters of interest. To solve the simultaneity problem between *Markup* and *BO*, I use the instrumental variable approach explained below.

3.2.1 Instrumental Variables

I use $BINrevs_{ij}$ and $BOrevs_{ij}$ as instrumental variables for $Markup_{ij}$ in equation (1) and for BO_{ij} in equation (2), respectively. $BINrevs_{ij}$ and $BOrevs_{ij}$ represent the *number* of revisions on a listing that were concerning BIN price and concerning Best Offer usage, respectively. To explain more closely what revisions in the context of eBay are, take for instance the variable, BO_ij. While BO_ij indicates the current (or final/ending) value (either 0 or 1 depending whether the seller has enabled BO), the $BOrevs_{ij}$ variable gives a count of how many times the BO_ij variable had been changed during the active listing period prior to its current or final value that is observed in the data collected.

The revision data allows for being able to see that there was a revision, and what it concerned. However, the data does not show exactly what the previous price amount was for a price revision. Nor does it show the revised the previous value of BO was for a BO revision. In the case of BO (than in the case of price), it can be identified more easily what the previous

BO value was because there are only two possible values that BO can take on, but I do not attempt to do that here. Instead, the focus of these revision variables is on quantifying the number of revisions made on either price or on the BO feature. The idea is that a higher number of changes or revisions may give the seller higher learning regarding what is profitable for him and what is not. We can think of this as an experimental approach to learning by the seller. Therefore, I use the number of revisions of each type, and the higher the number, the higher the learning on that aspect for the seller.

As noted in Chapter 1, the revision behavior of sellers differed from the 2013 dataset to the 2016 dataset. That is, sellers were found to be revising their listings more often in 2016 than in 2013. This suggests that the IV approach is better suited for the 2016 data, and therefore I only use the 2016 dataset for this paper. More importantly though, this suggests that the revision behavior of sellers seems to be picking up momentum or is at least relevant for the sellers in 2016, and thus sellers are in-fact using revisions to learn and improve their listings.

3.2.2 IV for Markup in equation (1)

The number of price revisions, $BINrevs_{ij}$, in the recent history of a particular listing are correlated with the current price premium. In general, if the seller revises the price of a listing more times, he is likely gaining more information on pricing that causes him to change the price, or he gains more information on pricing *after* he changes the price and observes varying outcomes due to the price changes he makes.

The number of price revisions/learning, $BINrevs_{ij}$, is not directly impacting the seller's current choice to use Best Offer; the impact of the number of price revisions on Best Offer usage occurs only through the current price premium or markup. It is plausible, however, that a greater number of price revisions can represent higher seller knowledge regarding the item, and this could, in fact, have a direct impact on a seller's choice to use Best Offer. Therefore, including variables regarding whether or not a seller is private and his experience on eBay is crucial for controlling for the impact of a seller's knowledge about an item on his choice to use Best Offer.

3.2.3 IV for BO in equation (2)

$BOrevs_{ij}$, the number of revisions regarding the Best Offer feature on a listing, impacts the current Best Offer usage on a listing; a greater number of revisions represents a higher knowledge about the Best Offer feature. I do not have an expectation for the direction of the correlation, but I do expect that as sellers learn more, they may choose to use Best Offer more often or less often depending on its efficacy. $BOrevs_{ij}$ does not directly impact a seller's markup on a given item. It is possible that if a seller generally revises a listing more often (not BO-specific revisions, that is), then he is less patient or less informed/experienced, and therefore charges a lower markup due to his lack of patience or lack of experience. Therefore, controlling for a seller's general inclination to revise, seller patience, and seller's experience level is important.

3.2.4 Possible Limitations of the IV approach

The parameter identification problem in both equations is complicated if the two instruments, BIN revisions and BO revisions, are in a simultaneous relationship with each other. That is, if a BIN revision always accompanies a BO revision or vice versa, the learning/revisions on BIN and BO are occurring simultaneously. There is no particular reason to expect that the learning is in fact simultaneous, but there is a possibility nonetheless.

Upon examining the data for 2016 revisions, I find that in the recent revision history of the listings, the revisions are not occurring simultaneously. Revisions on BIN are barely ever accompanied by BO revisions (see Table 27). The most recent BIN revision, BINrev1, was accompanied by BO revision, BOrev1, only 4.2% of the time. On the other hand, revisions on BO are more often accompanied with revisions on BIN, but only 42% of BO revisions result in BIN revisions at the same time (see Table 26). Since there are 10 times more revisions on BIN revisions than on BO (590 vs. 59 for BINrev1 and BOrev1, respectively), the simultaneous correlation on BIN revisions and BO revisions is overall not a concern in the data.

3.2.5 Solution/Preferred Method: Seller-Level Analysis

An alternate approach for circumventing any possible problems that can arise if there exists simultaneity between the instrumental variables, as mentioned above, is to conduct the same IV approach at the seller-level instead of conducting it at the item-level. That is, while the choice of BIN and choice of BO by the seller can be simultaneous for a given listing or item, the learning on BIN and BO can follow separate trends for a seller, across all of the seller's items. Therefore, seller's overall learning regarding BIN or regarding BO are bound to be less related in nature because they do not have a particular item in common. In other words, any item-driven learning that may occur more simultaneously is eliminated when conducting this analysis at the seller-level.

In addition, each seller's own learning with respect to BO and BIN is unique for each item. Some sellers may treat learning on BO and on BIN as simultaneous, but others may not. They may do so for some items and not for others. This difference in learning/revising behavior by seller and by item further reduces the possibilities of simultaneity of revisions between BIN and BO at the seller-level. If the learning and revision trends on BO and BIN were highly simultaneous somehow for sufficient number of sellers though, the IVs may not satisfy the necessary exclusion criteria.

Furthermore, the seller-level approach is more pertinent for studying seller strategies, and therefore more preferable than the item-level approach.[2] In addition, using these IVs at the seller-level has an advantage; since the IVs attempt to capture the seller's learning, learning is better captured across multiple items than just one item. I discuss the IV method at the item

[2] Nevertheless, due to aggregation in the variables, the seller-level approach is prone to higher measurement error than the item-level approach. Measurement error in the dependent variable tends to increase the standard errors, while measurement error in the independent variables tends to downward bias the coefficients. Therefore, results for the seller-level analysis must be interpreted with this in mind.

level first because it is then easier to explain the aggregation of the variables and the benefits of the approach at the seller-level.

Consider now the following simultaneous equations at the seller-level:

$$\begin{aligned} \bar{Markup}_j = \mu_0 + \mu_1\bar{X}_j + \mu_2 S_j + \mu_3\bar{L}_j + \mu_4\bar{M}_j + \\ \mu_5\bar{BO}_j + \mu_6\bar{BINrevs}_j \end{aligned} \tag{3–7}$$

$$\begin{aligned} \bar{BO}_j = \lambda_0 + \lambda_1\bar{X}_j + \lambda_2 S_j + \lambda_3\bar{L}_j + \lambda_4\bar{M}_j + \\ \lambda_5\bar{Markup}_j + \lambda_6\bar{BOrevs}_j \end{aligned} \tag{3–8}$$

Note that the variables in the above equations are denoted with bars. These bars represent that the variable is an aggregate variable at the seller level; that is, item-level characteristics have been aggregated for each seller across all of the seller's items in the sample. $\bar{BO}$ and $\bar{Markup}$ represent the seller's average choice of Best Offer across all of his items and seller's average markup across all his items, respectively. $\bar{X}_j$ are average item characteristics, $\bar{L}_j$ are average listing characteristics, and $\bar{M}_j$ are average market characteristics for all of the seller's items. $\bar{BOrevs}$ and $\bar{BINrevs}$ are average number of revisions by the seller on Best Offer usage and on BIN price, respectively.

The simultaneity between $\bar{BO}$ and $\bar{Markup}$ still exists similarly at the seller-level analysis as it did in the item-level analysis. Therefore, the subsequent parameter identification still poses a problem to which the IV approach offers a remedy. That is, a seller's general choice of using Best Offer may be simultaneously determined by his general markup strategy across all his items. $\bar{BOrevs}$ and $\bar{BINrevs}$, seller-level averages of item-level instruments,$\bar{BOrevs}$ and $\bar{BINrevs}$, serve as instruments for equation(8) and equation(7), respectively, to help identify μ_5 and λ_5, respectively.

For the parameter identification of μ_5 in equation (7), I use $\bar{BOrevs}$ as an instrument for $\bar{BO}$. I expect that seller's revisions and learning/experimentation with the BO feature directly impact a seller's BO strategy, but they do not impact the seller's Markup strategy directly, except indirectly through the seller's BO strategy. I do not have an expectation for the the sign of this relationship, but I do expect that the sellers will choose to sell more through Best Offer if it is effective, and less if it is not.

Similarly, for the parameter identification of λ_5 in equation (7), I use the $\bar{BINrevs}$ to instrument for $\bar{Markup}$. I expect that a seller's revisions and learning on BIN price directly impact a seller's markup strategy, but they do not impact the seller's Best Offer strategy, except through the seller's markup strategy. I expect that a higher number of BIN revisions provides the seller with more knowledge on how best to price his item. In general, I do not have an expectation for the sign of this relationship; I simply expect that if lower prices are an effective strategy, then sellers are more likely to reduce the price. It seems though that a learned seller would tend to price lower than the Car Guru estimate or the IMV, consistent with the traces of empirical evidence I find in Chapter 1 and Chapter 4 of this dissertation.

Finally, note that this analysis can be conducted both at the item-level as well as at the seller-level. However, there may be some drawbacks of conducting this analysis at the item-level. While Chen et al. (2016) test the impact of BO on the prices charged by a seller using an item-level approach, I find that this question is better tested at the seller-level. That is, there is more information for testing the intention of this question when conducted at the seller. While Chen et al. (2016) do not make explicit their intention for testing the question, they do mention that they expected the prices to be higher for listings/cars that were listed under BO. I infer from this that they might be interested in testing the age-old notion regarding bargaining which has not been expressed much formally in literature, that bargainers tend to increase the price of the item, because they know they will bargain. If this is in-fact their question of interest, as it is in this paper, it is better to study this question at the seller-level, because the seller-level analysis gives us a better read on the seller's behavior. Furthermore, the problem of studying this question at the item-level is that since the quality of a car is not observable by the researcher, and perhaps more observable by the buyer and certainly by the seller. Therefore, it can be misleading to study the question of whether bargaining impacts the seller's markup/pricing because a higher markup/price may be a sign of higher quality (or higher *TrueValue* of the car). As mentioned earlier, the empirical methodology used in this paper, as well as in Chen et al. (2016) is not able to account for whether a high price truly signals a higher quality (*TrueValue*), or is merely a seller strategy. Studying the question at the seller-level, though, certainly allows for studying the seller's overall intention (and general strategy across *all* of the items of a given seller), and can answer more appropriately the question regarding the impact of BO on pricing[3] .

Furthermore, it seems that the decision to use BO occurs more often at the seller-level than at the item level. The data collected indicates that most sellers either use BO or posted prices mechanism. Only 42 sellers in 2016's sample and 14 sellers in 2013's sample were found using a mix of BO and posted price strategies. Therefore, it does not seem to be the case that sellers vary the choice to use BO on an item-by-item basis. Rather, it seems that each seller has a general preference for one method over the other. Hence, it seems more reasonable to study the impact of seller's overall strategy of BO on his pricing strategy.

The benefit of the IV approach at the seller-level, discussed at the beginning of the section, where the concerns regarding the simultaneity of the IVs are further eliminated at the seller-level, is an added bonus. That is, while the analysis itself is better conducted at the seller-level than at the item-level, it is even more beneficial to conduct the analysis at the seller-level when using the IV approach proposed in this paper, as discussed at the beginning of this section.

3.2.6 Variables

The variables that determine the variable $\bar{BO}$ (for estimating equation (7)) have been thoroughly explained and discussed in the previous two chapters (except for the $Ma\bar{r}kup$

[3] There is an implicit assumption in this argument that a seller has a relatively uniform strategy across his items, and is not intentionally mixing the combination of markup and BO strategies for the sake of *appearing* innocent (of pricing higher because he chooses to bargain) when researcher or buyers study his behavior across all his items

variable which was excluded from the reduced forms in the previous chapters due to endogeneity concerns mentioned above). However, the determinants of $\bar{Markup}$ (for estimating equation(8)) have not yet been discussed in the dissertation, and I focus this section primarily on the independent variables that are likely to impact a seller's markup strategy.

There is not much economic literature that indicates exactly what variables impact a seller's price premium or markup. However, there are a few commonly used variables in papers that attempt to study price determinants. FollowingAndrews and Benzing (2007), which looks at the determinants of price of automobiles on eBay auctions, I include the seller's reputation (measured through feedback percent), seller's experience (measured through whether a seller is a dealer or not), and the seller's total feedback score, as variables that are likely to have a positive impact on the seller's price premium.

In addition, I also suspect that higher seller patience is likely to impact the markup positively. Given the results from Chapter 2, indicating that patient sellers are less likely to choose to bargain, I expected that perhaps patient sellers do not bargain because they choose instead to post a higher price and wait to sell at that price. I also expected that seller learning, in general, and seller learning specific to pricing would have an impact on the price premium, though I was not sure which direction the impact would be.

I also expect that seller competition would impact the seller's price premium. However, while on the one hand, higher competition can lead to a lower price premium, if there is imperfect competition, then a higher number of sellers does not necessarily imply a lower markup. Therefore, I include competition variables in this model, but do not have a clear theoretical prediction for them.

I expect that sellers selling older cars are likely to charge a lower price premium. I also included expensiveness of a car to control for other item characteristics.

3.3 Results

Table 3-5 presents the primary results of this paper. The first two columns are specifications that study the impact of seller's average markup on his use of BO, where the second column presents results for errors clustered by the seller's state. The latter two columns in table 3-5 present the models that study the impact of seller's use of BO on his average markup. Again, the second of these two columns presents results for the model with errors clustered by the seller's state.

3.3.1 Impact of Markup on BO

I focus first on the model explaining BO usage of the seller (presented in the very first column of Table 3-5). I do not find that the seller's markup strategy measured by the variable *Selleravgpriceprem*, has any impact on the seller's average usage of BO across all his listings. I do find though that a one standard deviation increase in the seller's learning, specific to BO, decreases the seller's use of BO by 0.06 or 6%. The relationship between seller's learning on BO feature through higher BO-related revisions indicates that as the seller gains more learning on the BO feature, he becomes less likely to use BO.

I also find that if the seller is a private seller, the seller is likely to have higher BO usage across all his listings. Together, both of these variables, *Selleravgpriceprem* and *sellerprivate*, highlight again (as was concluded in Chapter 2), that BO is used more often by a less experienced seller, and as sellers gain more learning specific to the BO feature through higher number of revisions on the BO feature, the seller uses BO less often. The results presented in the next column, with errors clustered by the seller's state, confirm these results. In this specification, however, it is also confirmed again (as was the case in Chapter 2) that sellers with higher reputation have lower use of BO.

3.3.2 Impact of BO on Markup

The latter two columns in Table 3-5 present the results for the model testing the impact of seller's choice of BO on his choice of markup. I focus on the third column in this table. I find that if a seller increases his use of BO by one standard deviation, then he also chooses a lower average markup for his cars by 28.6%. That is, if a seller chooses to use a BO strategy for selling his cars more often, then he also tends to price relatively lower. This result is contrary to the notion that if sellers choose to bargain, they also increase the price. This result is also contrary to the findings of Chen et al. (2016), who use a different dataset than this paper, of course, and also use a different econometric method (as will be explained more later).

With respect to other control variables that impact a seller's markup strategy, I find that a patient seller also typically charges a higher markup. This result is consistent with the expectation that a patient seller may use his bargaining power in the form of a higher markup. This is also consistent with why a patient seller tends to not use BO (as discovered in Chapter 2). Rather than using BO, and signaling openness to a lower price, it seems that a patient seller posts a higher price and waits, consistent with my expectations.

Furthermore, I find that if the seller tends to sell older cars in mileage terms, measured by a one standard deviation increase in *SelleravgMileageby1000*, then his average markup tends to be lower. This is consistent with the expectation that older cars will sell for lower prices than newer cars. Surprisingly, though, when it comes to cars that are older by Year, the result is the opposite. If the seller tends to sell cars that are newer in terms of the Model Year, measured by *SelleravgYear*, then he also tends to charge lower markups, on average.

3.3.3 Comparison with OLS Results

Table 3-6 presents the results for the OLS model, which implicitly assumes that there is no simultaneity between a seller's pricing strategy and the seller's choice of BO. I compare these OLS results to the results presented previously using the IV approach (where simultaneity is in-fact accounted for). This helps illustrate the difference in methodology between Chen et al. (2016) and my work, which leads to differing conclusions regarding the pricing strategy of sellers that use BO. That is, this comparison illustrates the impact of not accounting for the simultaneity between the seller's choice of BO and his markup strategy.

As can be seen, Table 3-6 resembles the layout of Table 3-5, in that the first two columns present the results for explaining a seller's BO usage, and the latter two columns present the results for explaining the seller's markup strategy. The second and fourth column of the results present results with standard errors clustered by the seller's state, as was the case with Table 3-5.

I focus first on the impact of Markup on BO, presented in the first column. According to the OLS model, and unlike the IV approach presented earlier, a seller's markup strategy does impact his use of BO. That is, if the seller charges a higher markup, on average, then he can be expected to have higher use of BO across all his listings. Note that this is different than the IV approach results presented earlier, in which I account for the simultaneity between seller's markup and choice of BO. Other control variables behave similarly to what was mentioned previously (in this chapter, and in Chapter 2)regarding the impact of these variables, namely *sellerprivate* and *Selleravgrevisionstotal*.

Next, I consider the impact of seller's BO usage on the seller's average markup. In this case, I find that there is no impact of seller's BO usage on the seller's markup strategy. This is somehwat contradictory to the findings of Chen et al. (2016), which indicate that sellers that use BO tend to price higher (consistent with the common notion that sellers that use bargaining often also charge a higher price).

3.4 Discussion

Using the IV approach to analyze 2016 data, I find that sellers' use of BO (vs. posted prices) reduces the markup for cars listed by them on eBay motors. In addition, I find that using the OLS approach for studying the impact of BO on the pricing of a car, without accounting for simultaneity between the choice of pricing/markup strategy and the choice of using BO, it can seem that BO usage has no impact on the seller's markup strategy.

Furthermore, I find that seller patience, measured through a seller's relisting behavior or *Selleravgrelist* variable, leads to higher markups. Combined with the insights from Chapter 2, this suggests that a patient seller chooses not to negotiate, and instead chooses to charge a higher markup. This gives further insight into why the markups might be lower for sellers that use BO vs. markups by sellers who use posted prices: Patient sellers choose posted prices vs. choosing bargaining (Chapter 2 result), and they also charge higher prices. Therefore, it seems that patient sellers might be using their bargaining power in the form of choosing a higher first offer (or markup in this case), rather than choosing to use Best Offer (which might be a signal used by weak bargainers).

When studying bargaining markets vs. posted prices, it is natural to compare welfare implications for each of these markets. One way of studying the welfare implications in this context is to study the primary channels through which seller's micro behavior impacts consumer surplus. One channel through which bargaining can improving consumer surplus is if bargaining allows for price discrimination or helps coordinate agents' behavior towards efficiency of trade, in the spirit of Coase (1960). However, there is also a channel through which bargaining can negatively impact consumer surplus which is not mentioned as often in economic literature (perhaps due to lack of economic literature on bargaining in retail markets). This channel is one in which sellers might charge a higher price only because they choose the bargaining mechanism, as was the implicit expectation in Chen et al. (2016), where this question is looked at empirically for the case of eBay Motors. Since economic theory does not offer too much discussion nor clear predictions on this matter, using empirical evidence to examine these two channels can help understand the dominant equilibrium seller behavior on eBay, and its impact on consumer surplus.

Consistent with the theoretical insights of Wang (1995) and the empirical findings of Chen et al. (2016), I found in Chapter 2 that bargaining is more likely to exist in markets are relatively more monopolistic in terms of the product. On the one hand, the use of bargaining by a monopolist to achieve first-degree price discrimination can help improve trade possibilities for consumers. On the other hand, bargaining may add to the inefficiency already present in the monopolistic market: that is, if the seller charges a price even higher than his intended (monopoly) price solely because he wants to keep some bargaining markup, then bargaining introduces additional inefficiency into the market. For instance, a seller might want to raise the listed price or his initial asking price in order to make up for the transaction costs of bargaining.

This concern is not restricted to a monopoly market structure alone. Due to incomplete information in an imperfectly competitive market, a seller might want to charge a higher price in order to exploit uninformed buyers (as price dispersion literature suggests). In an incomplete information setting thus, if the seller has the choice to bargain or to choose posted prices, he might choose to bargain as a means to price discriminate, even if he is not a monopolist. Again, there is a welfare concern if the lack of information leads to seller's adoption of bargaining (consistent with findings in Chapter 2), which then causes him to charge a higher price so that he can negotiate downward later (or not), depending on whether his buyers are stronger or weaker bargainers. To the extent that a seller's choice of bargaining might cause him to charge a higher price, there are additional inefficiencies being introduced into the market due to the bargaining selling mechanism.

Therefore, there are possibilities of a high markup leading to higher use of Best Offer, and simultaneously, it is possible, that the use of BO can increase the markups being charged. Nevertheless, the results of this paper rule out the case that sellers that use BO are likely to also charge a higher markup because of their choice to use BO. Rather, I find that the use of BO leads to sellers charging lower markups, on average. This rules out the case of concern regarding the negative impact of BO on consumer surplus. That is, I do not find evidence that sellers that use BO also charge a higher price, because of their choice to bargain. Furthermore, I also do not find that sellers that charge higher markups (as a means of reaping the benefits of lack of information) are more likely to use BO, which was a possible channel for improving consumer surplus. Together, evidence on these two channels suggests that consumer surplus is not being negatively affected by the use of bargaining by sellers.

Finally, the quality uncertainty, or the "lemon problem" that might be present in the automobile market (Akerlof, 1995) also poses some concerns for buyers that purchase cars on eBay Motors. As mentioned earlier, this research is not able to assess the quality of the cars directly, but we can infer some equilibria for sellers' behavior, where there are four primary signals that the seller can send using a combination of markup and BO. The empirical work presented in this paper helps rule out some of these theoretically plausible cases using empirical evidence, as discussed below.

The cases in Table 3-1 and Table 3-2 that have italicized texts are the cases in which sellers charge a higher markup either as a means of misleading buyers regarding the quality of the product due to buyers' lack of information or cases in which the seller does not know the quality of his product but attempts to charge a higher price for it nonetheless. These cases are the cases where the seller may try to con consumers and lower their consumer surplus. According to the empirical evidence I find in this chapter, I do not find that a higher markup leads the seller to choose BO, nor do I find that BO is associated with higher markups. Instead, I find that sellers who use BO often tend to price lower. Therefore, this rules out all the cases (with italicized text) where the seller is trying to con buyers regarding the quality of the product, except for a seller that has a low quality product, but also is patient, informed, and therefore chooses not to use BO. This seller can perhaps successfully use the markup strategy and post a fixed price, and mimic a similar seller who has a higher quality product. This seller might thus be able to successfully mislead buyers into thinking that his product/car is of higher quality. This certainly can impact consumer surplus negatively, and such cases deserve further evaluation in research on eBay Best Offer. In any case, since this seller actually falls under the posted price market and not in the BO market, it is safe to say that the eBay Best Offer selling mechanism does not seem to contribute to lowering consumer surplus through the suspected channels.

3.5 Tables

Table 3-1. Possibilities for Seller Markup and BO Strategy when Buyers are Perfectly Informed

Seller believes that buyers are well-informed.			
	Patient	**Impatient**	**Very Impatient**
Informed	High Markup	High Markup	?
	No BO	BO	BO
Uninformed	*High Markup*	Low Markup	Low Markup
	BO	BO	BO
High Quality Product			
	Patient	**Impatient**	**Very Impatient**
Informed	Low Markup	Low Markup	Low Markup
	No BO	BO	BO
Uninformed	*High Markup*	Low Markup	Low Markup
	BO	BO	BO
Low Quality Product			

Each cell indicates the possible choice of markup and BO that a seller will choose given his patience level, knowledge regarding the value of his product, and whether his product is of high or low quality.
The italicized text is indicative of cases in which the seller is likely to misrepresent the value/quality of his product.
The "?" represents ambiguity regarding the action that a seller will take.

Table 3-2. Possibilities for Seller Markup and BO Strategy when Buyers are Not Well-Informed

Seller believes that many buyers are not well-informed.			
	Patient	**Impatient**	**Very Impatient**
Informed	High Markup	High Markup	?
	BO	BO	BO
Uninformed	*High Markup*	Low Markup	Low Markup
	BO	BO	BO
High Quality Product			
	Patient	**Impatient**	**Very Impatient**
Informed	*High Markup*	*High Markup*	?
	No BO	*BO*	BO
Uninformed	*High Markup*	Low Markup	Low Markup
	BO	BO	BO
Low Quality Product			

Each cell indicates the possible choice of markup and BO that a seller will choose given his patience level, knowledge regarding the value of his product, and whether his product is of high or low quality.
The italicized text is indicative of cases in which the seller is likely to misrepresent the value/quality of his product.
The "?" represents ambiguity regarding the action that a seller will take.

Table 3-3. Mean BIN Revision Conditioned on a BO Revision

	BO Rev 1	**BO Rev 2**	**BO Rev 3**	**BO Rev 4**
BIN Rev 1	0.424 (59)			
BIN Rev 2		0.000 (16)		
BIN Rev 3			0.571 (14)	
BIN Rev 4				0.5 (2)

Revisions are numbered from most recent (1) to least recent (4) revision on the listing. Number of observations are reported in parentheses.

Table 3-4. Mean BO Revisions Conditioned on a BIN revision

	BINRev 1	**BIN Rev 2**	**BIN Rev 3**	**BIN Rev 4**
BO Rev 1	0.042 (590)			
BO Rev 2		0.000 (150)		
BO Rev 3			0.082 (98)	
BO Rev 4				0.019 (2)

Revisions are numbered from most recent (1) to least recent (4) revision on the listing. Number of observations are reported in parentheses.

Table 3-5. Main Results

Concept Variable	Variable Measure	BO	BO	Markup	Markup
Markup	Selleravgpriceprem	0.318	0.318		
		(0.300)	(0.253)		
BO	BOperlisting			-0.771**	-0.771*
				(0.325)	(0.462)
Seller BIN Learning	SelleravgBINrevsunscale			-0.085**	-0.085**
				(0.032)	(0.037)
Seller BO Learning	SelleravgBOrevsunscale	-0.330**	-0.330**		
		(0.078)	(0.097)		
Seller Patience	Selleravgrelist	-0.051	-0.051	0.094**	0.094**
		(0.039)	(0.038)	(0.033)	(0.035)
City Competition	countcitySeller	0.000	0.000	0.000	0.000
		(0.001)	(0.001)	(0.001)	(0.001)
Make Competition	countmakeSeller	-0.000	-0.000	0.000	0.000
		(0.000)	(0.000)	(0.000)	(0.000)
Seller Learning	Selleravgrevisions	-0.015	-0.015	-0.011	-0.011
		(0.011)	(0.015)	(0.018)	(0.023)
Seller Reputation	PositiveFeedbackPercent	-0.003	-0.003**	-0.003	-0.003
		(0.002)	(0.001)	(0.004)	(0.003)
Seller Experience	sellerprivate	0.046*	0.046**	0.001	0.001
		(0.024)	(0.021)	(0.036)	(0.037)
Seller Experience	Feedbackscoreby1000	-0.002	-0.002	-0.001	-0.001
		(0.003)	(0.002)	(0.005)	(0.002)
Seller Experience	repeatsellertag	-0.016	-0.016	0.000	0.000
		(0.038)	(0.031)	(0.055)	(0.025)
Seller's Time Cost	sellerlistingsN	0.000	0.000	0.001	0.001
		(0.001)	(0.000)	(0.001)	(0.001)
Object Expensiveness	SelleravglnIMVby1000	-0.009	-0.009	0.045	0.045
		(0.026)	(0.023)	(0.033)	(0.033)
Object Valuation	SelleravgMileageby1000	0.002	0.002	-0.005**	-0.005**
		(0.002)	(0.001)	(0.000)	(0.000)
Object Valuation	SelleravgYear	0.011	0.011	-0.044**	-0.044**
		(0.014)	(0.012)	(0.006)	(0.008)
cons		-20.647	-20.647	90.246**	90.246**
		(27.237)	(24.306)	(11.850)	(15.510)
N		1412.000	1412.000	1412.000	1412.000
r2		0.8363	0.8363	0.0629	0.0629
S.E. Clustered by State		N	Y	N	Y

* p<0.10, ** p<0.05

Standard Errors are reported in the parentheses.

Seller Make Dummies included.

Table 3-6. Main Results Comparison using OLS Approach

Concept Variable	**Variable Measure**	**BO**	**BO**	**Markup**	**Markup**
Markup	Selleravgpriceprem	0.041**	0.041*		
		(0.021)	(0.025)		
BO	BOperlisting			0.054	0.054
				(0.035)	(0.038)
Seller BIN Learning	SelleravgBINrevsunscale			-0.056**	-0.056*
				(0.025)	(0.031)
Seller BO Learning	SelleravgBOrevsunscale	-0.282**	-0.282**		
		(0.056)	(0.085)		
Seller Reputation	Positive Feedback Percent	-0.003	-0.003**	-0.000	-0.000
		(0.002)	(0.001)	(0.003)	(0.002)
Seller Patience	Selleravgrelist	-0.022	-0.022	0.108**	0.108**
		(0.022)	(0.023)	(0.028)	(0.025)
City Competition	countcitySeller	0.000	0.000	0.000	0.000
		(0.001)	(0.001)	(0.001)	(0.000)
Make Competition	countmakeSeller	-0.000	-0.000	0.000	0.000**
		(0.000)	(0.000)	(0.000)	(0.000)
Seller Learning	Selleravgrevisions	-0.019*	-0.019	0.006	0.006
		(0.010)	(0.015)	(0.014)	(0.014)
Seller Experience	sellerprivate	0.038*	0.038*	-0.030	-0.030
		(0.022)	(0.021)	(0.029)	(0.022)
Seller Experience	Feedbackscoreby1000	-0.002	-0.002	0.001	0.001
		(0.003)	(0.002)	(0.004)	(0.001)
Seller Experience	repeatsellertag	-0.013	-0.013	0.007	0.007
		(0.037)	(0.029)	(0.048)	(0.024)
Seller's Time Cost	sellerlistingsN	0.000	0.000	0.000	0.000
		(0.001)	(0.000)	(0.001)	(0.000)
Object Expensiveness	SelleravglnIMVby1000	0.003	0.003	0.047*	0.047
		(0.022)	(0.016)	(0.028)	(0.033)
Object Valuation	SelleravgMileageby1000	0.000	0.000	-0.005**	-0.005**
		(0.000)	(0.000)	(0.000)	(0.000)
Object Valuation	SelleravgYear	-0.001	-0.001	-0.044**	-0.044**
		(0.004)	(0.002)	(0.005)	(0.008)
cons		3.283	3.283	87.703**	87.703**
		(8.011)	(4.270)	(10.137)	(16.759)
N		1412.000	1412.000	1412.000	1412.000
r2		0.059	0.059	0.303	0.303
S.E. Clustered by State		N	Y	N	Y

* p<0.10, ** p<0.05

Standard Errors are reported in the parentheses.

Seller Make Dummies included.

Table 3-7. Seller-Level Summary Statistics for 2016

Variable	**Obs.**	**Mean**	**SD**	**Min**	**Max**
BOperlisting	1861	0.829	0.371	0.0	1.0
Selleravgpriceprem	1853	-0.107	0.540	-4.8	1.0
SelleravgBINrevsunscale	1861	0.215	0.510	0.0	4.0
Selleravgrelist	1861	0.378	0.465	0.0	1.0
SelleravgBOrevsunscale	1861	0.029	0.168	0.0	2.0
Selleravgrevisions	1861	0.278	0.928	0.0	9.0
sellerprivate	1665	0.563	0.496	0.0	1.0
PositiveFeedbackPercent	1586	99.165	4.386	33.3	100.0
Feedbackscoreby1000	1861	0.439	2.620	-0.0	95.5
SelleravgMileageby1000	1861	70.995	62.255	0.0	1350.0
SelleravgYear	1861	2008.656	4.048	1986.0	2016.0
sellerlistingsN	1861	2.585	9.956	1.0	298.0
repeatsellertag	1861	0.071	0.258	0.0	1.0
countcitySeller	1861	8.662	15.331	1.0	78.0
countmakeSeller	1861	129.557	94.954	1.0	326.0
SelleravglnIMVby1000	1861	2.846	0.794	-1.6	5.5

CHAPTER 4
DOES BEST OFFER LEAD TO SELLING?

4.1 Motivation

While the previous chapters focus on why and how sellers use Best Offer (BO), by either studying the adoption of BO by sellers or by studying the interaction of BO and pricing or markup, an analysis for the outcomes that result from using BO has not yet been conducted. According to Chen et al. (2016), BO increases the probability of selling, but does not increase the transaction price of a car. Therefore, while there is some evidence that the use of BO benefits sellers, it is unclear if these benefits alone are sufficient to explain the dominant use of BO on eBay Motors. Hence, there is a need for further exploration into the possible channels through which BO might be profiting sellers, and therefore, leading them to use BO. I find it crucial to study the interaction of a seller's choice of markup and his choice of BO, as these two decisions are both closely linked, as discussed in Chapter 3. Furthermore, it seems that BO may also have some negative aspects that lead sellers to revise their listings against using BO as they learn more (result from Chapter 2). There has been no evidence as of yet regarding the negative aspects of using BO for sellers.

In order to conduct a more thorough search for the benefits and costs of using BO on sellers, I study three channels through which BO could benefit sellers, leading them to choose BO vs posted prices. I study whether 1)BO increases sales, 2) BO leads to quicker sales, and 3) if BO impacts the transaction price, and I complement and extend the work of Chen et al. (2016) with this regard.

The results of this analysis emphasize the importance of incorporating the interaction of markup and BO into the specifications of interest. Contrary to Chen et al. (2016), I do not find that BO always increases the likelihood of sale. Rather, I find that BO increases the likelihood of sale for cars with sufficiently high markups, and that BO *decreases* the likelihood of sale for lower markups. This result provides a more complete picture of both the benefits and costs of BO. Furthermore, I also find that sellers can recover 29% the sales that they lose due to charging a high markup if they use BO (vs. using the posted prices method) in addition to using a high markup. This benefit of BO indirectly helps in increasing the transaction price, because a higher markup leads to a higher transaction price by 40.3%. This is an important benefit of BO, which gets neglected in Chen et al. (2016), perhaps because they do not study the interaction term of BO and markup in their analysis. Finally, I also find that BO increases the time for selling for most reasonable levels of markup. This result helps explain why some sellers revise against using BO, as was noted in the previous chapters.

Finally, to further study the possible bargaining power that a seller might derive from charging a high markup, I test whether the nature of discount provided by sellers that charge a higher markup is different than the sellers that do not charge a higher markup. I find that a high markup does lend a seller higher bargaining power; that is, a seller with a higher markup provides a *lower* discount ratio compared to sellers with lower markups, given that the car resulted in sale and the seller was using BO.

A thorough analysis of the outcomes of BO thus provide a better perspective on why BO is prevalent on eBay Motors, while also providing more insight into how BO can also lead to some negative outcomes for sellers on important aspects of selling.

4.2 Empirical Methodology

4.2.1 Does BO lead to selling?

To study the impact of Best Offer and Markup on the likelihood of selling, I use the Probit model for the binary variable $sold_i$, which takes on a value of 1 if the car is sold at the end of the listing, and 0 if it is not sold by the end of the listing. Given the possible simultaneity of Markup and BO mentioned in Chapter 3, I am particularly interested in studying the impact of the interaction of the two variables on the likelihood of sale. The equations below represent the probit models of interest:

$$\begin{aligned} sold_i = \alpha_0 + \alpha_1 Markup * BO_i + \alpha_2 Markup_i + \alpha_3 BO_i + \alpha_4 X_i \\ + \alpha_5 S_i + \alpha_6 L_i + \alpha_7 M_i + \alpha_7 Make_i + \varepsilon_i \end{aligned} \tag{4–1}$$

where $Markup * BO_i$ represents the interaction term between BO and Markup variables, $Markup_i$ is the markup on a given car, BO_i is a binary variable for whether an car is listed as BO or not, X_i denotes car characteristics, S_i denotes seller characteristics, L_i comprises of listing characteristics, $Make_i$ represents make dummies, and where M_i represents market characteristics.

Due to a lack of model thickness in my data, I use make dummies along with the Instant Market Value(IMV)(the Car Guru Estimate for each car). Controlling for mileage, year, and some other car characteristics individually, including Car Make and IMV sufficiently control for car model type. I do include model dummies in some specifications for robustness checks.

4.2.2 Does BO lead to quicker sales?

On the one hand, bargaining is said to be time consuming, and is therefore known to have higher transaction costs. On the other hand, it can increase the rate of selling according to Wang (1995), implying that bargaining can lead to quicker selling, consistent with Chen et al. (2016)'s prediction and finding. Therefore, I test the impact of BO, markup, and the interaction of markup and BO on the time that it takes for cars to sell.

$$\begin{aligned} durat_{ij} = \beta_0 + \beta_1 Markup * BO_{ij} + \beta_2 Markup_{ij} + \beta_3 BO_{ij} + \beta_4 X_{ij} \\ + \beta_5 S_j + \beta_6 L_{ij} + M_{ij} + Make_{ij} \end{aligned} \tag{4–2}$$

where $durat_{ij}$ represents the time it takes for a car to sell. In my dataset, the duration of a listing is available at the level of seconds, and I convert this to days. Therefore, in this paper, the measure used for duration is duration days. However, this variable is only captured completely for cars that sold during the observation (or data collection) time period. For cars whose listings merely "ended," but were not "sold," duration data is right censored, assuming that these cars will eventually sell, although I do not observe their selling date, and thereby their duration. Therefore, I use survival analysis to study the impact of Best Offer and Markup on the time it takes to sell.

To the extent that cars observed in this data do not differ fundamentally or distributionally in their likelihood of sale, and rather differ more along the lines of *duration* to sale, survival analysis is appropriate to use. To avoid bias in survival analysis, censoring must be "random and non-informative(Despa, 2010)". Part of the censoring in the data occurs due to my choice of data collection time period, which is a randomly chosen date appropriate for my time of graduation. Since I collected only ended listings and dropped any car that happened to be listed multiple times in the ended listings (that is, it must have ended twice within the time period of my data), it is as if I followed a listing throughout its first occurrence on eBay. I do not observe listings that have not yet ended, and it is possible that these un-ended listings might consist of cars that have much longer duration. The analysis in this paper will not be able to comment on such listings with possibly long durations, though perhaps if such a listing *ended* at the time of my data collection, this listing would be in the dataset. Overall though, I do not expect that listings that sold vs. listings that did not sell to have vastly different distributions. Furthermore, I use some insights from the model in this paper that studies the probability of selling for a car. This model and its results gives some signals regarding what distinguishes a car that is more likely to sell than one that is less likely to sell. For example, a higher mileage makes a car less likely to sell. Using Kaplan Meier curves, I explore whether there appears to be a *distributional* difference between high-mileage and low-mileage cars, and I do not find this to be the case. I look at other similar variables that are determinants of selling likelihood to see if the survival probabilities vary vastly along these variables. These results are presented later in the Duration model section, and there does not seem to be any indication that cars more likely to sell follow a strictly different distribution than cars less likely to sell. It seems reasonable then that perhaps it is mostly a matter of time till a car does sell. Therefore, it is also more relevant in some sense to ask how long a car takes to sell than looking only at whether or not it sells in the time period of observation.

Survival analysis requires specifying a distribution for $durat_i$. Since I generally assume that cars that have sold are not distributionally different than the cars that have sold in my data, I expect that it is perhaps only a matter of time till a car sells. Therefore, I expect that if a car has been on the market for a long time, $durat_i$ would have negative duration dependency. That is, if a car has been on the market for a long time, it may be closer to being sold. However, it is also possible that there is no duration dependency at all. In either case, the Weibull distribution seems appropriate, since it allows for monotonic duration dependency, so it does allow for negative dependency or no duration dependency. Non-parametric estimation using the Kaplan Meier curve, shown in Figure 4-1, confirms that $durat_i$ has a negative duration dependency, meaning that survival probabilities decreases as time goes on. This is confirmed later in the results presented, when the shape parameter, p, is estimated to be greater than 1 (as will be explained later).

Finally, it must be noted that the data might also left-censored to some degree: if the car was listed prior to the time of data collection as a separate listing than the one observed in the data, then the duration of the car might be longer than what is observed. I drop the few observations that have the same car VIN listed multiple times under different eBay IDs. These observations were either erroneous, or cars more likely to be left-censored, or they could be cars with much longer duration than usual because perhaps they are just not likely to sell. This should not bias my results because these cars may be significantly different than other cars, perhaps even distributionally, making survival analysis less appropriate for these cars. However, despite dropping these observations, I cannot assure that the cars I observed have never been listed prior to their occurrence in the data. Nevertheless, left-censoring of this sort might be random across all the observations in the data, and I do not expect my choice of date for data collection to be systematically causing a left-censoring that biases my data. Therefore, for the purpose of this paper, I do not account for the left-censoring. Therefore, I focus more on correcting for the right-censoring, because this sort can more easily be observed and is known to be frequent in the data, and can cause systematic bias in studying duration if unaccounted for.

4.2.3 How does BO affect the transaction price?

Best Offer may be impacting the transaction price of a car. According to Chen et al. (2016) Best Offer does not have any impact on the transaction price, but it does increase the likelihood of sale. However, Chen et al. (2016) do not study the impact of the interaction term of markup and BO on either transaction price or likelihood of sale. Previously, I mentioned that it is important to study the impact of the markup and BO together (i.e. of the interaction term) on studying likelihood of selling. I propose studying the impact of this interaction term between markup and BO on the transaction price and the log of transaction price, $ln(transact)$ (for ease of interpretation of the results).

However, before the transaction price can occur, the car (or car, in this case) must pass the "hurdle" of resulting in trade or in a sale. A car's characteristics, its seller's characteristics, and its listing and market characteristics determine whether or not the car is likely to sell. Whereas, the transaction price of a car is more likely to be determined primarily by the car characteristics and the pricing strategy of the car, while the market characteristics and seller characteristics are still likely to have some impact. The pricing strategy is of interest in this paper. However, not accounting for the impact that pricing strategy (which includes whether or not BO was used, in this case) has on the likelihood of sale can confound the true impact of the pricing strategy on the transaction price. However, the transaction price of a car is only available in the data if the car sold. We can say then that the transaction price is censored from below for cars that have not yet sold. Consider therefore the two steps before a transaction price can occur: 1) an occurrence of a sale, and 2) the transaction price for that sale.

This can be expressed in the following way:

$$\mathrm{d}_i = \begin{cases} 1, & \text{if } d_i > 0 \\ 0, & \text{if } d_i \leq 0 \end{cases} \tag{4–3}$$

and

$$\text{transact}_i = \begin{cases} transact_i^*, & \text{if } d_i^* > 0 \\ 0, & \text{if } d_i^* \leq 0 \end{cases} \tag{4–4}$$

where $transact_i$ is the observed transaction price for a car, $transact_i^*$ is the latent variable, for the true transaction price of a car, and d_i is a dummy variable for whether a car sold or not. There are some hurdles to consider that make an car more likely or less likely to sell. We can think of this hurdle further; consider:

$$\text{transact}_i = \begin{cases} transact^*{}_i, & \text{if } u_c - \bar{c} > 0 \\ 0, & \text{if } u_c - \bar{c} < 0 \end{cases} \tag{4–5}$$

where u_c is a utility from a car i, and $\bar{c}$ is some threshold utility needed for the car to be sold or selling-worthy. Assume $\bar{c}$ is equal for all cars across all buyers.

Consider the following model of interest:

$$\begin{aligned} transact_i^* = \gamma_0 + \gamma_1 Markup * BO_i + \gamma_2 Markup_i + \gamma_3 BO_i + \gamma_4 X_i \\ + \gamma_5 S_i j + \gamma_6 L_i + \gamma_7 M_i + \varepsilon_i \end{aligned} \tag{4–6}$$

where BO_{ij} denotes the seller's best offer usage, which is a binary variable, and $Markup_{ij}$ is the seller's markup over a price estimate, or Instant Market Value (IMV) for the car, obtained from Car Guru:

$$Markup_{ij} = \frac{BIN_{ij} - IMV_{ij}}{IMV} \tag{4–7}$$

X_{ij} denotes item/car characteristics, S_{ij} denotes seller characteristics (years of experience on eBay, behavioral indication of patience, seller reputation), L_{ij} are listing characteristics, and M_{ij} are market characteristics.

For ease of notation, assume that vector $\mathbf{X}$ comprises of all the independent variables in the model of interest, where $\mathbf{X}$ is assumed to have a normal distribution, with mean μ and variance σ^1. Then, the primary equation of interest is the following:

$$transact_i^* = \gamma_0 + \gamma \mathbf{X} + \varepsilon_i \tag{4–8}$$

Having noted the above model and the censoring problems that need to be accounted for, the Tobit model might seem like a reasonable choice. However, the Tobit model turns out to be restrictive in a way that is unappealing to entangle the different parts of the process, as will be explained below. The Cragg two-step model, also known as the hurdle model, is a more general form of the Tobit model. The first step of the model predicts whether the likelihood of sale using a probit model (similar to the previous analysis I conducted in section 2.1 of this paper), and the second part uses a truncated normal regression for studying the transaction price, conditioned on selling.

The primary distinguishing feature between the Cragg model and the Tobit model is that the Cragg model allows for differing factors (and differing coefficients) in the two steps of the analysis. Note that vector of independent variables of interest, **X**, is allowed to be different in the two steps: x_1 and x_2, respectively, are not the same variables, and the coefficients, α and β, respectively, are also allowed to be different. Due to this flexibility, this model is sometimes preferred to the Tobit model. I therefore use the Cragg model so that the variables that determine the likelihood of selling are allowed to be different with respect to their impact on the transaction price (Wooldridge, 2010) (Krupka and Croson, 2016).

The Cragg model studies two equations:

$$d_i{}^* = \alpha x_{1i} + \mu_i \tag{4–9}$$

$$transact_i{}^* = \beta x_{2i} + \varepsilon_i \tag{4–10}$$

where μ_i and ε_i are random and independently distributed, ε_i $N(0, 1)$ and μ_i $N(0, \sigma_\varepsilon{}^2)$ Kiyingi et al. (2016).

Using the Cragg model, the probit part of the model helps predict the following probabilities:

$$P(transact_i = 0|x_{1i}) = 1 - \Phi(x_{1i}\alpha) \quad (4\text{–}11)$$

$$P(transact_i > 0|x_{1i}) = \Phi(x_{1i}\alpha) \quad (4\text{–}12)$$

For the truncated regression part of the model, the following expected value of *transact*, conditional on a car being sold, is being predicted:

$$E(transact_i|transact_i > 0, x_{2i}) = x_{2i}\beta + \sigma * \mu(x_{2i}\frac{\beta}{\sigma}) \quad (4\text{–}13)$$

where $\mu(z)$ is the inverse Mills ratio (IMR), $\mu(z) = \frac{\phi(c)}{\Phi(c)}$, where ϕ is the standard normal probability distribution function. The inverse Mills ratio (IMR) from the truncated normal regression accounts for possible selection bias due to the truncated sample.

Although I do not expect that the variables that impact selling are vastly different than the variables that impact the transaction price (i.e., I allow $x_1 = x_2 = \mathbf{X}$), I do expect that their impact can be different across the two aspects of selling: 1) likelihood of sale, and 2) transaction price. For instance, while a seller's willingness to sell (his seriousness towards selling), can impact his likelihood of sale positively, it can also have a *negative* impact on the transaction price (whether under posted prices or Best Offer). Not separating the two parts needed for a transaction price to occur can confound the results regarding the true impact of a seller's seriousness or willingness to sell on the transaction price.

Furthermore, even when the direction of impact of a variable is the same, when studying the impact of Best Offer on the final transaction price of the car, it is important to also see and account for the fact that BO may impact the selling likelihood of a car differently than the transaction price. On the one hand, it may lead to higher sales, but have no impact on transaction price, consistent with Chen et al. (2016). Or, buyers may be less likely to buy from a seller that is using BO, but more likely to choose a convenient transaction with someone who is using posted prices. In this case, BO would lead to lower sales than posted prices would, *and* when used, it is possible that BO would reduce the transaction price due to the negotiation process. Therefore, even when the impact of BO is in the same direction, separating the two parts of the process allows for the magnitudes of the same variable to vary across the different aspects of selling.

A statistical test can also help in the choice of Cragg model vs. the Tobit model. We can estimate independently the log likelihood ratios for a probit, a truncated regression, and a Tobit model using the same set of variables. We can then compute the following likelihood ratio statistic (Katchova and Miranda, 2004):

$$\lambda = 2(lnL_{probit} + lnL_{truncreg} - lnL_{Tobit})$$

where λ follows a chi-square distribution with R degrees of freedom, where R is the number of independent variables, including a constant. If $\hat{\lambda}$ exceeds the chi-square critical value, then we can reject the more restrictive Tobit model can be rejected in favor of the Cragg model (Katchova and Miranda, 2004). I conduct this test for the preferred specification for transaction price (explained below), and I get $\hat{\lambda} = 1900.39$ for 66 degrees of freedom. This is well above the critical value, and therefore, I reject the Tobit model in favor of the Cragg model for analyzing transaction price.

4.2.3.1 Marginal Effects for the Cragg Model

Finally, to get the marginal effects using the Cragg model, I estimate the following expectations using the main variables of interest:

From the probit part of the model, the partial effect or marginal effect of a variable, in this case BO, can be written as follows:

$$\frac{\partial P(transact > 0|x_1)}{\partial x_{BO}} = \alpha\phi(x_1\alpha) \tag{4–14}$$

From the truncated regression part of the model, the partial effect or marginal effect of BO can be written as follows:

$$\frac{\partial E(transact_i|transact > 0, x_2)}{\partial x_{BO}} = \beta_{BO}[1 - \mu(x_2\frac{\beta}{\sigma})[x_2\frac{\beta}{\sigma} + \mu(x_2\frac{\beta}{\sigma})]] \qquad (4–15)$$

I compute and report these marginal effect conditional on an car being sold at the means.

Finally, the partial effect or marginal effect of BO on the "unconditional expectation" of *transact* is the least straightforward to compute because it will be a function of parameters and explanatory variables from both steps of the analysis:

$$\frac{\partial E(transact_i|x_1, x_2)}{\partial x_{BO}} = \alpha_{BO}*\phi(x_1\alpha)*[x_2\beta+\sigma*\mu(x_{2i}\frac{\beta}{\sigma})]+\Phi(x_1\alpha)*\beta_{BO}[1-\mu(x_2\frac{\beta}{\sigma})[x_2\frac{\beta}{\sigma}+\mu(x_2\frac{\beta}{\sigma})]] \qquad (4–16)$$

Usually, the coefficients for step 1, also called Tier 1, and step 2, also called Tier 2, are both reported, and there are no coefficients reported for the overall Cragg model because the model highlights that there are two separate processes to be looked at simultaneously. Therefore, I report both coefficients for both tiers.

There are two ways of conducting and reporting the analysis for the Cragg model using Stata. One method is to separately run the probit and the truncated regression, and another method is to use the *craggit* command, which does both steps at the same time. I use both methods, and show that the results are very much the same. Previous papers are known to use both methods (Katchova and Miranda, 2004) (Krupka and Croson, 2016). Using the first method, I compute the marginal effects at the means for all for probit and for the Truncated Regression separately Burke et al. (2009).

The benefit of the second method, using the craggit command, is that the partial effects for each observation can be computed and then finally averaged. This is generally preferred to computing marginal effects at the means, mentioned earlier. In addition, the *unconditional* expectation can also be computed if the craggit command is used, and this is not possible using the first method mentioned above. In order to compute this marginal effect on the unconditional expectation, the *craggit* command in Stata, rather than separately conducting the probit and truncated regression. This is tedious and time consuming to do for each variable, so I do this only for the primary variables of interest.

I use bootstrapping to get standard errors Burke et al. (2009). However, bootstrapping usually works better for large sample sets, and therefore this poses some limitation in my relatively small dataset for making inferences using these estimates. Nonetheless, these average partial effects are noteworthy, and therefore these are reported for the primary independent variables of interest: markup and BO. Marginal effects are generally not reported for interaction terms separate from the primary variables, so I do not have marginal effects for the interaction term.

4.2.4 Does markup impact discount ratio?

In the previous two chapters of this dissertation, there is suggestive evidence that sellers use high markup as their bargaining power. That is, if a seller is more patient (which according to theory should give him more bargaining power), he charges a higher markup (Chapter 3) and does not use Best Offer, and is more likely to use the posted prices mechanism (Chapter 1). Furthermore, the posted prices selling mechanism is associated with higher markups than the markups under the BO mechanism (Chapter 3). These results indicate that charging a higher markup (although it leads to lower likelihood of sale, as will be shown later in this paper), and staying patient with a high markup might be proving beneficial for sellers.

It is therefore interesting to study mechanisms through which sellers might be benefiting from a higher markup. So, I test whether a higher markup leads to higher bargaining strength. I use the $discountratio_i$ variable to measure the bargaining strength of a seller. That is, given that a seller chooses bargaining and that the transaction successfully results in a sale, a seller whose discount ratio is higher was perhaps weaker during the negotiation process than a seller with a lower $discountratio_i$. In this section, I only study the sellers that have chosen to bargain.

As discussed in the case of transaction price, $discountratio_i$ too has similar censoring problems in the data. $discountratio_i$ is only observed in the data for sellers that chose to use

BO, and if the car was sold. Consider, therefore, the following:

$$\text{discountratio}_i = \begin{cases} = \textit{discountratio}_i^*, & \text{if } v_{ij} - \bar{u} > 0 \text{in time period t} \\ = 0, & \text{if } v_{ij} - \bar{u} < 0 \text{ in time period t} \end{cases} \tag{4–17}$$

where $discountratio_i$ is the observed discount ratio provided when a car was sold through the BO feature, and $discountratio^*{}_i$ is the latent variable, the true ratio of discount provided for the transaction for that car if the car sold. v_{ij} is a bargaining seller's utility from negotiation agreements for car i, and $\bar{u}$ is some threshold utility for a seller to agree to the terms of transaction and trade at the price proposed, similar to its definition in the previous section.

Now consider the factors that affect the level of discount provided. The equation of interest is specified as follows:

$$\begin{aligned} discountratio_i = \beta_0 + \beta_1 Markup * BO_i + \beta_2 Markup_i + \beta_3 BO_i + \beta_4 X_i \\ + \beta_5 S_i j + \beta_6 L_i + M_i + Make_i + \varepsilon_i \end{aligned} \tag{4–18}$$

where $discountratio_i = \frac{transact_i - BIN}{BIN}$

As was the case for the transaction price of the car, I use the Cragg model for estimation, which is preferred to the Tobit model in order to address the censoring problem in this data due to its flexibility in allowing the underlying mechanisms (for selling and for discount) to vary, as explained earlier. In this context, it means that the mechanism that determines whether an car would be sold is allowed to be different than the mechanism that determines the bargaining strength or the discount rate provided by a seller, conditioned on the car being sold.

4.2.5 Data

For this paper, the variable, $sold_i$, is a key variable. Since this variable was not collected during the 2013 dataset, I can only conduct the analyses for this paper using the 2016 data. For the first part of the paper, where I test if BO and/or markup leads to higher likelihood of selling, I use the 2016 item-level data.

To study the impact of BO and/or markup on the time it takes for the car to sell, $durat_i$, I use the same item-level dataset. I use the $sold_i$ variable to identify which observations for $durat_i$ are censored in the data. For the models explaining $transact_i$ [1] and $discountratio_i$ variables, I use the 2016 item-level data as well, censoring using the sold variable and using the sold and BO variables, respectively.

4.3 Results

The primary results for this paper are presented in Table 4-1. This table presents the results for the four dependent variables of interest in this paper. That is, the first column presents the results of the probit model which explains the likelihood of selling and the second column presents results for the model explaining the time needed for selling a car. The third column studies the impact of the markup and BO on the transaction price ($ln(transact)$ to be exact), and the fourth column presents for the model studying the impact of the markup of an car on the discount ratio for that item.

4.4 Results: Probability of Selling

The primary variables of interest for determining the likelihood of sale are the markup on an item, BO, and the interaction of BO and markup. According to Table 4-1, a one standard deviation increase in markup on an car decreases the probability of selling by 0.38 points, measured at the mean value of the BO variable. At the minimum value of BO, i.e., when there is no BO used (posted-price method), I find that a one standard deviation increase in markup reduces the probability of selling by 62.7%. At the max value of BO, i.e. when BO is used to sell an item, I find that a one standard deviation increase in markup reduces the probability of selling by 33.7%. Therefore, keeping the markup strategy of a seller constant, using BO (vs. using posted prices) reduces the negative impact of a high markup on the likelihood of sale by 29.0%.

[1] www.watchcount.com helped me get access to the final transaction prices link on eBay.

Although the BO variable alone is not statistically significant, the interaction term between markup and BO is statistically significant. I find that at the average value of the markup, using BO reduces the likelihood of sale of a car by 10.2%. At the lowest value of the markup, using BO can reduce the likelihood of sale by 300%. However, at the highest value of the markup, BO increases the likelihood of selling, by 56.7%. I find that at a markup value of 0.108, the impact of BO changes from negative to positive. That is, if a seller asks for a BIN price which is at least 10.9% higher than the Car Guru estimate for the car, then using BO on that car is likely to have a positive impact on sales. Using BO for a car with a lower markup that the 10.8% mentioned, using BO can decrease the probability of selling the car.

Perhaps this result on the lowest value of markup can be explained by the theoretical underpinnings of the signal that sellers send to buyers in the market by using certain combinations of the markup and BO strategy. As can be seen in Table 3-1 and in Table 3-2 from Chapter 3, sellers that use low markups might be either impatient, very impatient, uninformed regarding the quality of their product, or they might have a low quality product. Note that among the 21 cases listed completely (i.e. without "?") listed in these two tables, the combination of a low markup and BO occurs 10 times (48% of the complete cases listed). However, each of these cases 10 cases where the markup is lower, and the seller is using BO, there is some quality uncertainty present. That is, either the car is indeed of a lower quality, or the car is of high quality, but the seller is uncertain about the quality or the valuation of his car. With this framework in mind, it is thus not surprising that cars with lower markups listed under a BO listing pose some quality concerns for the buyers. Therefore, for such cars, using BO can lead to a lower likelihood of selling. This result is somewhat different than the findings of Chen et al. (2016), who claim that BO increases the likelihood of sale. I extend this finding by finding evidence that while using BO does increase the likelihood of sale for higher markups, BO can also reduce the likelihood of sale for cars with lower markups.

I also find that that a relatively more expensive car, measured by the variable *lnIMVby100*, is less likely to sell. An older car, in terms of higher mileage, is also less likely to sell, and surprisingly, a newer car, in terms of model year, is less likely to sell. The result on model year is somewhat surprising, because I expected that older cars in general will be less in demand and less likely to sell, but perhaps there is higher availability of newer cars off-eBay, and therefore newer models are less likely to sell on eBay. Interestingly, I also find that a seller's deposit time period also has an impact on the likelihood of selling. That is, if sellers indicate that they require a deposit immediately (vs. in 24 hours) or if they do not mention any time deadline for a deposit ((vs. mentioning 24 hours), the car is less likely to sell. This variable has never been tested before in any eBay study, and it represents the seller's indication of his seriousness with respect to his time. It turns out that if the seller is either too strict or too lax in terms of the time-seriousness he indicates, his car is less likely to sell.

4.5 Results: Duration

I expected that selling duration would have positive duration dependency or no duration dependency. That is, since my observation period is limited, the cars may get relisted later or the seller may sell them off-eBay, but that the cars will eventually sell even if I do not observe it in my data. Positive duration dependency implies that perhaps they are more likely to sell (probability of "hazard" increases) as time increases or and constant duration dependency implies that the duration for selling is independent of the time that has already passed. The shape parameter, p, reported in 4-1 is greater than 1, and this confirms positive duration dependency, although since p is not drastically greater than 1, the data exhibits signs of constant duration dependency as well. Therefore, the use of Weibull distribution for the survival rate analysis seems appropriate, because Weibull allows for monotonic duration dependency, and the special case of Weibull is the constant duration dependency or the exponential distribution. I confirm the expected duration dependency and the legitimacy of the use of the Weibull distribution using the non parametric estimation using the Kaplan Meier Curve in Figure 4-1[2] .

The reported coefficients in the second column of Table4-1 indicate whether a given independent variable increases the probability of hazard (in this case, selling), or decreases it. If the coefficient is positive, then the variable increases the probability of selling as days go on; meaning that the car is likely to sell faster. If the coefficient is negative, then the variable decreases the speed of reaching the hazard event, or, in this case, decreases the speed of selling.

As can be seen in the second column of Table 4-1, which presents the results for the duration model, a one standard deviation increase in the markup of a car leads to 36.7% longer time ((-0.476+(.84)*-.386)*.458) for selling the car, computed at the mean value of BO usage. At the minimum value of BO (=0), i.e. if the car is listed under the posted price mechanism, a one standard deviation increase in markup leads to a an increased time for selling, by 21.8%. At the maximum value of BO (=1), i.e. if the car is listed under the BO mechanism, an increase in the markup takes a longer time to sell, by 39.4%. Therefore, when BO is used, an increase in the markup of a car increases the time that the car takes to sell.

[2] In addition, I use the Kaplan Meier graphs to address some distributional concerns that might arise for using the duration model in my analysis. Since the use of the duration model implicitly assumes that the censored and uncensored observations follow the same distribution, and the censoring in this data occurs based on whether or not an car sold, it is worth noticing the distribution patterns for items that sold vs. the items that did not sell. This cannot be observed in a straightforward manner because we only observe explicitly the duration for selling for the items that sold. Therefore, I study two categorical variables *known* to impact the probability of selling (according to the results of the previous section, presented in the first column of Table 4-1, such as the mileage (high or low) of a car, whether a seller is a private seller. The idea here is to study whether these factors, known to impact the likelihood of selling, have differing distributions when it comes to duration for selling. As can be seen in Figure 4-3 and Figure 4-4 so, that while higher mileage increases the time for selling a car, and a private seller sells a car faster than a car dealer, the duration distributions do not differ along these two aspects. This alleviates some concerns that might arise regarding sold vs. unsold cars following vastly different distributions. Therefore, the use of survival analysis and its assumptions seem appropriate for the purpose of this paper.

Whereas, using posted price mechanism (vs. using BO), a one standard deviation increase in the markup, decreases the time for selling by 17.6% (39.4-21.8=17.6) from the time it would take if the seller used BO. In either case, whether with BO or without BO, increasing the markup increases the time for selling.

Although the BO variable alone is not statistically significant again for this model, the interaction term between BO and markup is statistically significant, and has a negative coefficient. When looking at the impact of BO, I find that at the average level of markup, use BO increases the time for selling an car by 7.4%. At the highest value for markup, I find that using BO increases the time for selling by 48.7%. At the lowest level of markup, using BO can reduce the time for selling by 167%. Using BO can expedite the time needed for selling if the value of the markup is -0.262 or lower. That is, using BO can decrease the time for selling if the car is listed at a price that is 26.2% lower than the CarGuru's estimate. Note that accounting for the fact that such a car, listed with BO, generally has a much lower likelihood of selling (according to the results in the previous section), the time for selling such a car is much quicker.

Furthermore, note that Figure 4-2, which presents the non-parametric estimation using the Kaplan-Meier curve, also shows that the red graph (which represents BO listings) has higher probabilities of "survival," than the blue graph(which represents posted price listings). In this case, "survival" implies longer time in the market without being sold. Therefore, using BO seems to increase the time for selling a car.

In addition, I find that a more expensive car, measured by *lnIMVby1000*, takes longer time to sell, a car with higher mileage takes longer time to sell, and a newer car model (in terms of model year) also takes longer time to sell. I also find that if a seller is too lax about the deposit time deadline, and mentions nothing with this regard, his car takes longer to sell. Perhaps this seller is actually in no rush to sell or perhaps buyers do not perceive him to be very serious, and therefore it takes longer time to sell this seller's car. Interestingly, a private seller's car sells much quicker than a dealer's car. Perhaps because the private seller has lower time cost than the dealer, and is therefore able to invest more time into selling the car quicker. If a car is listed by a seller who is relatively more patient, measured by his relisting behavior or *Selleravgrelist*, then the car is likely to take a longer time to sell. This is consistent with the expectation that a patient seller is willing to wait for longer to sell his car.

4.6 Results: Transaction Price

As mentioned earlier, the third column of Table 4-1 presents the results for explaining the impact of BO and markup on the transaction price. In particular, though, this column only presents the results for the tier 2 (Truncated Regression part) of the Cragg model. That is, the tier 1 (Probit part) of the Cragg model has been suppressed in the presentation of the results, because it does not differ from the results presented in first column of Table 4-1, because this first column is exactly the probit model that is run for the Cragg model tier 1, which explains the probability of a car selling (which is also the basis for the censoring that the Cragg model aims to highlight). Recall from the methodology section that the tier 2 results, the Truncated Regression that is, presents coefficients, which should be interpreted as effects that hold true, *conditioned* on a car being sold.

The third column of Table4-1 presents these Tier 2 results. A one standard deviation increase in the markup leads to a 40.3 % higher transaction price conditioned on an car being sold, when computed at the average value of BO in the sample. At the minimum value of BO(=0), i.e. when the car is listed under the posted price mechanism, I find that a one standard deviation increase in the markup of a car leads to a higher transaction price by 33.9%. At the maximum value of BO(=1), i.e. when the car is listed under BO, an increase in markup leads to a 41.9% higher transaction price. This implies, that for the same level of change in the markup of a car, using BO can increase the impact of the markup, by 8% (41.9-33.9=8). That is, if an car is listed under BO, then conditioned on an car being sold, increasing the markup of that car is likely to increase the transaction price by 8% more than if the car was listed under the posted price mechanism.

Although the BO variable is not statistically significant (as was the case for the previous two models presented in this table), the interaction term between the markup and BO is statistically significant. At the average value of markup, for cars that were sold in the sample, I find that using BO decreases the transaction price by 5%. At the highest value of markup, using BO increases the transaction price of a car by 5.3%, and at the lowest value of markup, using BO decreases the transaction price by 44%. I find that if the markup on a car is at least 0.13, then using BO on that car can increase the transaction price of the car. That is, if sellers on eBay list a BIN price that is 13% higher than the CarGuru estimate for that car, then using BO on that car can lead to a higher transaction price. However, using BO for a car with a lower markup than 0.13 can lead to a lower transaction price.

With respect to the IMV, the retail estimate of the car from CarGuru.com, I find that if a car has a higher IMV, then it also has a higher transaction price, where the elasticity of transaction price with respect to the IMV is equal to 1.02. Therefore, it seems that the transaction price of the car goes hand-in-hand with the CarGuru estimate for the car. I also find that if the car is older, in terms of mileage, then it sells for a lower transaction price. Again, surprisingly, if the car is newer by year, I would expect for it to sell for a higher transaction price, but it instead sells for a lower transaction price, conditioned on being sold. Finally, conditioned on being sold, a car that has no time deadline for the deposit requirement sells for a higher transaction price, conditioned on there being a sale. This result might be indicating that although sellers that are more lax or perhaps more patient are likely to sell less often, they may be able to extract a higher transaction price due to their patience if their car does sell [3] .

[3] Note that the advantage of the Cragg model is that it allows for differing impacts of the variables in the two parts of the analysis. For instance, while a higher markup has a negative impact on the probability of selling, its impact on transaction price (conditioned on a non-zero transaction price) is positive. In addition, while a seller's lack of mentioning the deposit time period or asking for the deposit to be submitted immediately decrease the probability of selling, these variables have no impact on the transaction price of a car. These differences for each

4.7 Results: Discount Ratio

Previous sections emphasize that at least when conditioned on a car being sold, a high markup leads to an increased transaction price, and this effect is especially large for cars that have a high markup *and* are listed under (or sold through) BO. Therefore, this section studies whether a high initial price, in the form of a high markup, lends the seller bargaining power. In particular, I test this by studying whether a high markup leads to a lower discount rate (which would imply high bargaining power on the part of the seller).

As was the case for the transaction price, I suppress the results for Tier 1 (the Probit part of the Cragg model) for ease of presenting all the results of the paper in a summarized manner. Recall that Tier 1 results for the discount ratio model considers the hurdle that all sellers who use bargaining must overcome: selling the item. However, of more interest here is the level of discount provided after a transaction is complete. Therefore, in the Main Results presented in Table 4-1, I present only the results of the second tier (the Truncated Regression part of the Cragg model). In this case, note that the Tier 2 results, the Truncated Regression that is, present coefficients which should be interpreted as effects that hold true only *conditioned* on a car being sold and being listed as BO.

The fourth column of Table 4-1 presents these Tier 2 results for the model explaining the discountratio. I find that a one standard deviation increase in the markup listed for a car leads to a 2.3% decrease in the discount ratio provided. This is consistent with the expectation that sellers with a higher markup have higher bargaining power, and therefore give 2.3% lower discounts than a seller who asks for a lower markup. This result is also consistent with the finding in the previous section that BO can increase the transaction price for markups that have a value higher than 0.13, whereas BO can decrease the transaction price for cars that have a lower markup than 0.13. These results combined indicate that among sellers

---

part of the process, selling and transaction price, would be confounded if the Tobit model were used instead.

who successfully sell through bargaining, and even among the ones that don't necessarily sell through bargaining, having a higher markup when the seller is a bargainer leads to higher bargaining power. The 2.3% that a seller saves in discount by asking for a higher initial price can, on average, translates to $384. for a car valued by CarGuru.com at $16,730 (the average for the sold cars sample). When computing the dollar amount of these seller savings at the lowest and highest car values in the sample set, the savings of the seller range from $11.7 to $3,808. These can be substantial savings for a car seller willing to bargain, which can be obtained by asking for a higher initial BIN price.

In addition, in explaining the discountratio variable, I find that if a seller is a private seller (vs. being a car dealer), then he is likely to give a higher discount by 2.8%. This is consistent with the expectation that a seller with relatively less experience and knowledge of the car market would have lower bargaining power.

4.8 Robustness Checks

I run four main robustness checks for this paper. Firstly, since the main analysis of this paper is conducted at the item-level, it is possible that errors might be clustered by seller. Furthermore, errors might also be clustered by Make, Model, and state if the likelihood of selling follows similar trends within a particular make, model, and state. Therefore, I run the model for predicting the variable *sold*, presented in the first column of Table 4-1, with errors clustered by the seller, Car Make Car Model, and the state that the car is listed in. As can be seen in Table 4-6, the results are exactly identical to the primary results presented in Table 4-1, indicating that the model is robust to the clustering.

The next three robustness checks are conducted simultaneously. Since one of the contributions of this paper is to study the specification that includes the interaction term of markup and BO, *Markup*BO*, I run some alternate specifications of the model, in which I exclude either the BO variable or the interaction term in order to see how the results might change. I also include/exclude some other variables to study which specification is more preferred for the primary results. Finally, I run these specifications on a dataset that does not include rare/collectible cars. That is, I drop the cars that made up less than 0.5% of the sample for the data collected in 2013 or in 2016, since these cars might be luxury or collectible cars that might behave differently than all other cars. Running these analyses again on a dataset that does not include these cars would confirm that results are not driven by these rare cars.

As can be seen instantly in Table 4-7, column(5) (which includes the interaction term in the specification), the primary results are not different than the results presented in Table 4-1. Therefore, it does not seem to be the case that collectible or rare cars were driving the main results presented in Table 4-1.

Now consider the specifications in column (1), (2), and (3) of Table 4-7. While column (1) excludes both the BO variable and the interaction term, column (2) specification includes the BO variable, but excludes the interaction term. Finally column (3) includes all three variables of interest, as is preferred. Note that most of the variables in these three columns have similar coefficients and similar statistical significance as the primary results presented in Table 4-1. However, the statistical significance of BO fluctuates across these three specifications. Note that inclusion of the BO variable in column (2)'s specification does not change the magnitude for the markup variable much, but it does indicate that BO variable is statistically significant, and has a negative impact on the likelihood of sale. That is, using BO can reduce the probability of selling. The BO variable, however, loses statistical significance when the interaction term for the variable is included in column (3). That is, whenever the interaction term is included in the model, the BO variable becomes statistically insignificant, while the interaction term remains statistically significant with a positive sign. This implies that not including the interaction term in the specification can give incomplete results regarding BO. As discussed earlier with respect to the impact of BO on the likelihood of selling, it is not the case that BO uniformly reduces the likelihood of sale or uniformly increases the likelihood of sale. Instead, for high levels of markup, using BO increases the likelihood of sale, while for lower levels of markup, BO decreases the likelihood of sale. This result would not have been captured if the interaction term between markup and BO was not included in the specification. Therefore, the preferred specification is the specification presented earlier: the one that includes all three variables: markup, BO, and markup*BO.

Finally, since I do not include Car Model fixed effects (model dummies) in this paper, in order to ensure that unobserved model characteristics are not driving the main results regarding the probability of selling, I used the variable *ModelavgIMVby1000*, which is the average IMV in the sample for all cars of the same model. When this variable was not statistically significant in any of the specifications form (1) - (3), I checked to see if dividing this variable along the lines of old and new would make a difference. That is, I used the variable *ModelavgIMVbynewoldby1000*, which inputs a different value for a car of the same model depending on whether the mileage on the car exceeds 500 (old car) or not (new car). As can be seen, replacing *ModelavgIMVby1000* with *ModelavgIMVbynewoldby1000* does not change anything. Thus, it seems that controlling for the value of each unique car specifically (which accounts for the model within it), using the variable *lnIMVby1000*, is sufficient, since *lnIMVby1000* is statistically significant. Therefore, the results of the model explaining the probability of selling are not driven by unobserved model characteristics.

4.9 Discussion

Chen et al. (2016) is the only other paper that studies the outcomes and benefits of using BO for sellers. They find that while BO does increase the likelihood of selling, it does not increase the transaction price. I complement and extent their work by not only studying the impact of BO on the probability of selling and transaction price, but also on studying the duration for selling. Using a different set and different econometric methodology than them, I find some similar results, but I also offer some other angles from which to view the results.

Unlike Chen et al. (2016), I do not claim that BO increases the likelihood of sale, in general. Since I study the impact of interaction term of markup and BO on the likelihood of selling (which they do not), I find that BO only increases the probability of selling for cars that have a markup of at least 10.8%. For all cars that have a markup lower than 10.8%, using

BO actually decreases the likelihood of sale. This highlights the importance of considering the interaction of BO and markup in the analysis. Furthermore, I run a robustness check, wherein I exclude the interaction term from the analysis (similar to the methodology in Chen et al. (2016)). I find that not including the interaction term variable in the analysis can be misleading because the results show that BO reduces the likelihood of selling altogether. This result is completely contradictory to the result in Chen et al. (2016). While the differences in the datasets and the time period can be the reason for this difference in results, I point to the difference in specification as the primary reason for differing results of the two papers. That is, when the interaction term is not included in the analysis, the inference regarding the impact of BO on the likelihood of sale is incomplete. Therefore, it might show negative results or positive results, perhaps depending on the nature of the markups that are in the sample. Therefore, using the interaction term of markup and BO in the analysis of likelihood of sale provides a more complete picture, that accounts for the interaction of markup and BO in the seller's strategy to sell, and the signals that buyers might be receiving due to the interaction of the two variables.

As discussed earlier in the results section of this paper, perhaps the underlying reason for this effect is that sellers that use lower markups paired with BO are trying to sell a lower quality car for a higher price. It seems then that buyers are able to understand this quality signal, and therefore cars with lower markups listed under BO are less likely to sell.

Furthermore, I find that increasing the markup for a car decreases the probability of the car selling, whether the car is listed under BO or not. However, if the car is listed under posted price, a one standard deviation increase in the markup leads to reducing the probability of sale by 62.7%, whereas under BO, the reduction is only 33.7%. Therefore, the benefit of using BO is that the seller can reduce the negative impact of a higher markup on the probability of selling by 29% if the car is listed under BO. The negative impact of a high markup on the likelihood of selling was to be expected, however, it seems that having a high markup has its own benefits (to be explained more later), and thus, some sellers will continue to charge a higher markup. However, these sellers can reduce the negative impact of a higher markup substantially by choosing to use BO as well. While Chen et al. (2016) have also attempted to study the benefits of BO, they did not study the benefits of using BO, in interaction with the markup chosen by the seller. Therefore, they are not able to trace this added indirect benefit of BO. That is, using BO allows for a seller to reduce the negative impact of a higher markup.

Now consider the impact of BO on the time it takes for a car to sell. There has been no empirical test of whether BO leads to selling the car quicker. According to the results of this paper, using BO can reduce the selling time for cars with markups that are 26.2% lower that the retail estimate provided by CarGuru.com. However, that is much below the average markup, which has a value of -0.06, or 6% below the CarGuru.com estimate. For cars with markups higher than 26.2%, the use of BO increases the time for selling. At the average level of markup, I find that using BO increases the time for selling by 48.7%. Therefore, while it is important to study the benefits of BO in terms of its impact on the likelihood of selling and on the transaction price, it is also important to study the impact BO has on the selling time for a car. According to the findings of this paper, BO negatively impacts the time for selling. A seller must weigh this aspect of BO into his decision to sell through BO. Furthermore, while markup may seem to have other benefits (suggested earlier, and explained more below), an increase in markup does increase the time for selling, and using BO adds to slowing down the selling time when a high markup is paired with using BO.

Consider now the impact of markup on the transaction price of a car, conditioned on a car being sold. As alluded to multiple times earlier, it turns out that a one standard deviation increase in the markup of a car leads to a higher transaction price by 40.3%. This is a substantial increase in transaction price, and therefore it becomes clear that a higher markup has significant benefits for a seller. While a higher markup can be extremely beneficial for sellers in terms of getting a higher transaction price, it has a trade off: it decreases the likelihood of selling (from 33.9% to 62.9% for BO and non-BO listings, respectively). This emphasizes the importance of BO's indirect benefit in helping mitigate the negative impact of a higher markup, by 29%, which is also substantial.

Finally, now consider the impact of BO on the transaction price of a car, conditioned on a car being sold. At the average value of markup, I find that using BO decreases the transaction price of a car by 5%. In general though, I find that if the markup on a car is at least equal to 0.13, or 13% higher than the CarGuru estimate for the car, then using BO can increase the transaction price of the car. Whereas, using BO for a car with a lower markup can lead to a lower price. Therefore, it seems that BO has the potential for increasing the transaction price for a car with a higher markup and can have a negative impact for a car with a lower markup. Therefore, while using BO has its benefits, the benefits of a higher markup far exceed the benefits of using BO for a seller. Paired together though, they provide the optimal strategy for a seller.

Finally, despite a much smaller dataset for analysis, I attempt to study the impact of a high markup on the seller's bargaining power. I find that a one standard deviation increase in markup leads to a lower 2.3% lower discount provided by the seller for that car, conditioned on a car being sold. For an average car in the sample, this can mean $384 of savings for the seller, and for the range of cars in the sample, this translates to savings of $11.7 to $3,808. Therefore,for sellers that are willing to bargain and their car meets the threshold of being selling-worthy, a higher markup likely gives them bargaining power.

It is beyond the scope of this paper (and dissertation) to discuss in too much detail the power/benefit of a seller charging a higher markup because car quality cannot be observed or assessed. Hence, it is unclear if a higher markup on a car indicates a higher quality car or if a higher markup is merely a strategy by the seller to *signal* higher quality, whether it is the case or not. However, buyers' behavior on eBay seems to indicate that they might perceive (whether rightfully or wrongly) the car with a higher markup to be of a higher quality.

This may be a reasonable perception on the part of the buyers. As Table 3-1 and Table 3-2 suggest, high markups are often listed by sellers who are well informed and are patient.

Furthermore, in 60% of the ten cases where the seller lists a higher markup, the seller is likely selling a higher quality car as well. However, in the remaining 40% of the cases, a seller may be using a higher markup and BO in an attempt to reap the benefits of imperfect information, hoping to exploit the buyers. While this exploitation is theoretically possible, as mentioned in Chapter 3 of the dissertation, it does not seem to be that case that sellers that use the BO strategy in general also charge higher markups. This result seems somewhat at odds with the results in this chapter, which suggests that a higher markup gives the seller bargaining power if he does choose to bargain. If using a high markup and using BO together are an effective strategy, why do we not observe sellers charging a higher price and using BO together more often?

Here, it might help to consider that the results of Chapter 3 are based on a seller-level analysis, while the results of this current chapter are based on an item-level analysis. As emphasized earlier in Chapter 3, the question regarding sellers' possible dishonesty in pricing the car they are selling is better studied at the seller-level, looking at the seller's behavior across all his cars, rather than looking at this question on a car by car basis. If we suspect that sellers might be cheating the buyer, they are likely to use similar strategies across all their cars. Since quality of a car is harder to assess in research, it is unclear whether a car is listed with a higher markup because it is truly of higher quality or is the seller is lying in an attempt to signal higher quality. Since I do not find many sellers using high markup strategies as a result of their choice to use BO, in general, across all their cars, in Chapter 3, I am inclined to believe that the results of this paper, which indicate the benefits of a higher markup, might be indicating benefits of cars with higher quality, where markups might be reasonably acceptable signals of car quality. This suggestion is consistent with the conclusions of Huston and Spencer (2002), which studied collectibles on eBay, and Adams et al. (2006), which studied Chevrolet Corvettes on eBay Motors, that there does not seem to be a "lemon problem (Akerlof, 1995)" (where low quality cars are being sold for higher prices) on eBay. Nevertheless, the strategic use of markups by sellers in the context of eBay Motors does deserve further attention, given the substantial impact it can have in the automobile market.

4.10 Figures and Tables

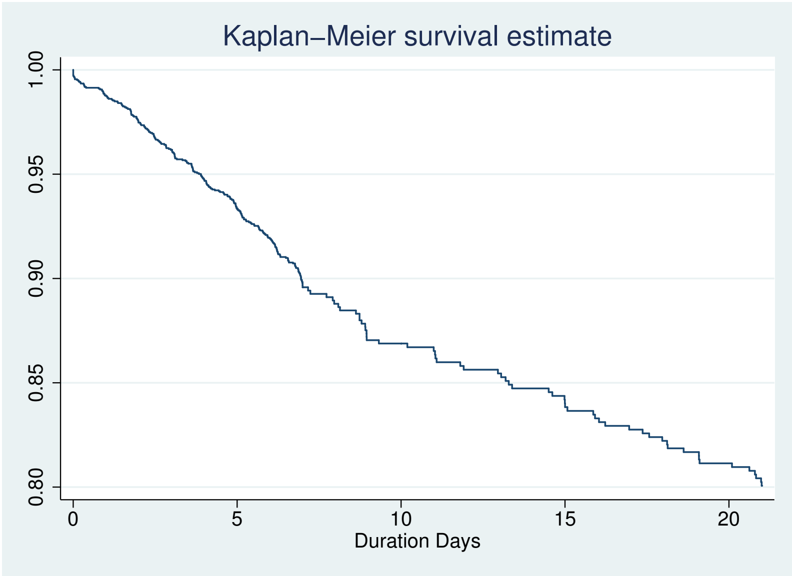

Figure 4-1. Survival (or not selling) probabilities for all cars

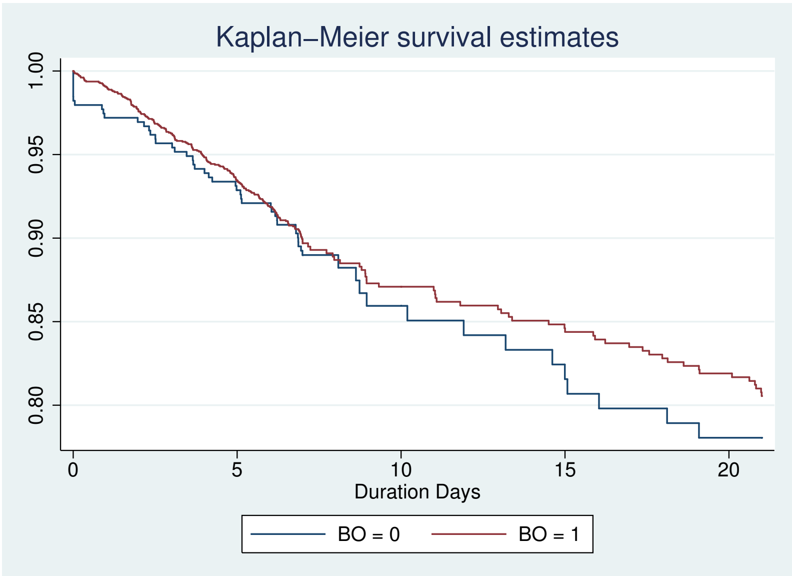

Figure 4-2. Survival (or not selling) probabilities for BO vs. non-BO cars

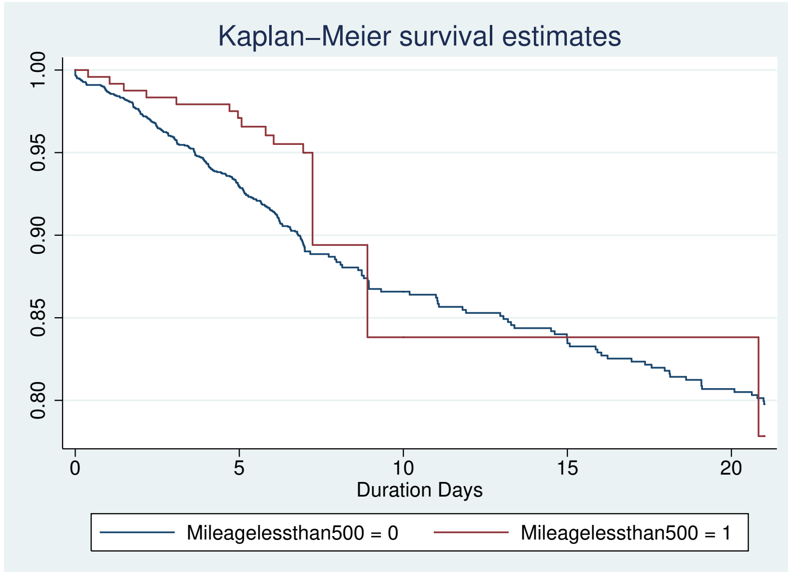

Figure 4-3. Survival (or not selling) probabilities for new/young cars vs. old cars

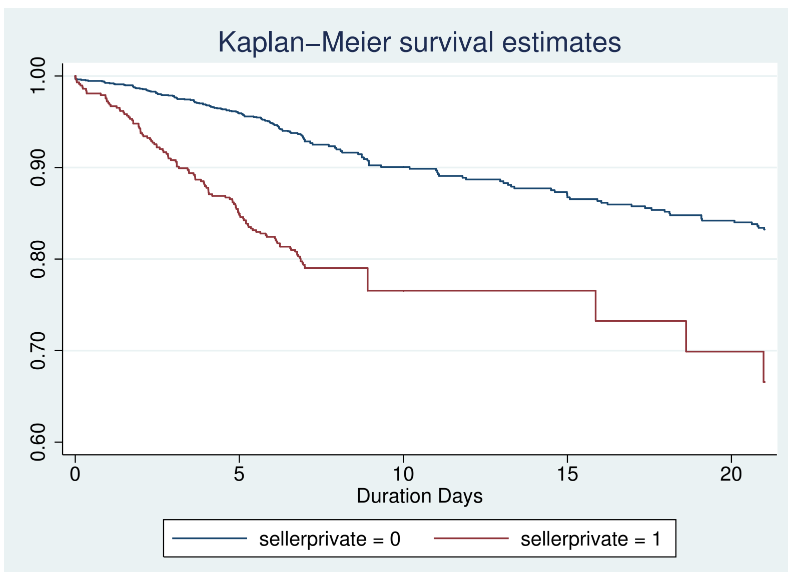

Figure 4-4. Survival (or not Selling) probabilities for private peller Cars vs. cars by a dealer

Table 4-1. Main Results: Probit, Survival Analysis, Cragg Model Tier 2

Concept Variable	Variable	sold	duration	lntrans.	discountratio
Markup	priceprem	-1.371*	-0.476*	0.395*	-0.028*
		(0.24)	(0.13)	(0.03)	(0.01)
Markup*BO	priceprem*BO	0.636*	-0.386*	0.093*	
		(0.24)	(0.15)	(0.03)	
BO	BO	-0.069	-0.101	-0.012	
		(0.11)	(0.17)	(0.03)	
Seller Patience	Selleravgrelist		-1.980*	0.034	-0.022
			(0.21)	(0.05)	(0.03)
Seller Learning	Selleravgrevisions	-0.002	0.039	-0.004	0.002
		(0.05)	(0.07)	(0.01)	(0.01)
Seller Reputation	PositiveFeedbackPercent	-0.005	-0.015	0.001	0.002
		(0.01)	(0.01)	(0.00)	(0.00)
Seller's Time	deposit=48 hours	-0.069	0.002	0.035	-0.032
		(0.15)	(0.22)	(0.04)	(0.02)
Seller's Time	deposit=72 hours	0.110	0.490	0.010	-0.022
		(0.54)	(0.74)	(0.12)	(0.06)
Seller's Time	deposit=immediately	-0.269*	-0.251	0.038	-0.008
		(0.13)	(0.19)	(0.03)	(0.02)
Seller's Time	deposit=none mentioned	-0.309*	-0.338+	0.075*	-0.029
		(0.14)	(0.21)	(0.04)	(0.02)
Seller Experience	sellerprivate	0.072	0.309*	-0.015	0.028+
		(0.10)	(0.15)	(0.02)	(0.01)
Seller Experience	Feedbackscoreby1000	-0.006	0.003	-0.003	-0.002
		(0.01)	(0.02)	(0.00)	(0.00)
Object Valuation	ModelavgIMVnewold1000	-0.004		-0.001	0.001
		(0.00)		(0.00)	(0.00)
Object Valuation	lnIMVby1000	-0.233*	-0.748*	1.017*	-0.080*
		(0.10)	(0.18)	(0.04)	(0.02)
Object Valuation	Mileageby1000	-0.004*	-0.003*	-0.002*	0.000*
		(0.00)	(0.00)	(0.00)	(0.00)
Object Valuation	Year	-0.071*	-0.057*	-0.017*	0.005
		(0.02)	(0.03)	(0.01)	(0.00)
Competition	itemcountstateby100			0.001	-0.001
				(0.00)	(0.00)
Competition	itemcountmakeby100			-0.031	-0.033
				(0.09)	(0.05)
cons		142.087*	115.861*	41.762*	-9.201
		(31.46)	(50.86)	(10.22)	(6.64)
p			1.05		
cons			0.044		
sigma				0.159*	0.059*
N		2497.000	2452.000	305.000	179.000
r2		0.2436			

+ p<0.10, * p<0.05

Note: Standard errors are reported in the parentheses and Make dummies are included.

Table 4-2. Summary Statistics: Sold Model

Variable	Obs.	Mean	SD	Min	Max
sold	2497	0.123	0.328	0.0	1.0
durationdays	2497	9.266	5.816	0.0	21.0
discountratio	2497	0.008	0.034	0.0	0.4
lntransact	2497	1.123	3.026	0.0	12.0
transactprice	2497	1768.507	7591.415	0.0	159649.0
priceprem	2497	-0.069	0.458	-4.6	1.0
BO	2497	0.839	0.367	0.0	1.0
priceprem*BO	2497	-0.052	0.416	-4.4	1.0
Selleravgrevisions	2497	0.182	0.691	0.0	7.0
Selleravgrelist	2455	0.358	0.408	0.0	1.0
IMVby1000	2497	30.263	29.519	0.2	336.5
Mileageby1000	2497	56.046	50.178	0.0	401.0
Year	2497	2010.274	4.193	1986.0	2016.0
sellerprivate	2497	0.243	0.429	0.0	1.0
deposit=48 hours	2497	0.090	0.286	0.0	1.0
deposit=72 hours	2497	0.004	0.063	0.0	1.0
deposit=immediately	2497	0.485	0.500	0.0	1.0
deposit=none mentioned	2497	0.336	0.472	0.0	1.0
itemcountstateby100	2497	5.893	5.298	0.0	14.0
itemcountmakeby100	2497	3.392	2.734	0.0	8.8
PositiveFeedbackPercent	2497	98.940	4.640	33.3	100.0
Feedbackscoreby1000	2497	1.321	5.105	0.0	95.5
ModelavgIMVbnewold1000	2497	30.371	25.098	0.2	336.5

Table 4-3. Summary Statistics: Duration Model

Variable	Obs.	Mean	SD	Min	Max
sold	2452.000	0.124	0.329	0.0	1.0
durationdays	2452.000	9.290	5.828	0.0	21.0
discountratio	2452.000	0.008	0.034	0.0	0.4
lntransact	2452.000	1.132	3.035	0.0	12.0
transactprice	2452.000	1777.440	7629.596	0.0	159649.0
priceprem	2452.000	-0.069	0.458	-4.6	1.0
BO	2452.000	0.840	0.367	0.0	1.0
priceprem*BO	2452.000	-0.052	0.416	-4.4	1.0
Selleravgrevisions	2452.000	0.182	0.692	0.0	7.0
Selleravgrelist	2452.000	0.357	0.408	0.0	1.0
IMVby1000	2452.000	30.478	29.684	0.2	336.5
Mileageby1000	2452.000	55.830	50.224	0.0	401.0
Year	2452.000	2010.306	4.182	1986.0	2016.0
sellerprivate	2452.000	0.235	0.424	0.0	1.0
deposit=48 hours	2452.000	0.088	0.283	0.0	1.0
deposit=72 hours	2452.000	0.004	0.064	0.0	1.0
deposit=immediately	2452.000	0.488	0.500	0.0	1.0
deposit=none mentioned	2452.000	0.335	0.472	0.0	1.0
itemcountstateby100	2452.000	5.927	5.298	0.0	14.0
itemcountmakeby100	2452.000	3.366	2.717	0.0	8.8
PositiveFeedbackPercent	2452.000	98.930	4.671	33.3	100.0
Feedbackscoreby1000	2452.000	1.329	5.140	0.0	95.5
ModelavgIMVnewold1000	2452.000	30.514	25.256	0.2	336.5

Table 4-4. Summary Statistics: Transaction Price Model, for non-zero Transaction Price

Variable	**Obs.**	**Mean**	**SD**	**Min**	**Max**
sold	305	1.000	0.000	1.0	1.0
durationdays	305	5.529	4.648	0.0	21.0
discountratio	305	0.063	0.076	0.0	0.4
lntransact	305	9.196	0.855	6.8	12.0
transactprice	305	14478.566	16986.478	900.0	159649.0
priceprem	305	-0.409	0.858	-4.6	0.7
BO	305	0.810	0.393	0.0	1.0
priceprem*BO	305	-0.306	0.759	-4.4	0.7
Selleravgrevisions	305	0.280	0.806	0.0	5.0
Selleravgrelist	305	0.135	0.256	0.0	0.9
IMVby1000	305	16.659	16.118	4.7	165.6
Mileageby1000	305	82.996	54.496	0.0	260.0
Year	305	2007.761	4.039	1994.0	2016.0
drivertrain1	245	0.457	0.499	0.0	1.0
drivertrain2	245	0.273	0.447	0.0	1.0
drivertrain3	245	0.269	0.445	0.0	1.0
cylinders	300	6.190	1.692	3.0	12.0
EngineVolumeLiters	292	3.652	1.353	1.0	8.0
sellerprivate	305	0.407	0.492	0.0	1.0
deposit=48 hours	305	0.148	0.355	0.0	1.0
deposit=72 hours	305	0.007	0.081	0.0	1.0
deposit=immediately	305	0.452	0.499	0.0	1.0
deposit=none mentioned	305	0.246	0.431	0.0	1.0
itemcountstateby100	305	4.505	4.223	0.0	14.0
itemcountmakeby100	305	2.436	2.482	0.0	8.8
PositiveFeedbackPercent	305	99.003	3.883	50.0	100.0
Feedbackscoreby1000	305	0.888	2.442	0.0	23.6
ModelavgIMVnewold1000	305	20.131	13.082	4.7	95.7

Table 4-5. Summary Statistics: Discount Ratio, for non-zero discount ratio

Variable	Obs.	Mean	SD	Min	Max
priceprem2	179	-0.396	0.811	-4.4	0.7
Selleravgrevisionstotal	179	0.267	0.843	0.0	5.0
Selleravgrelist	179	0.091	0.203	0.0	0.8
IMVby1000	179	16.730	16.034	5.1	165.6
Mileageby1000	179	82.508	54.195	0.0	260.0
Year	179	2007.860	3.771	1997.0	2016.0
sellerprivate	179	0.453	0.499	0.0	1.0
deposit=48 hours	179	0.190	0.393	0.0	1.0
deposit=72 hours	179	0.011	0.105	0.0	1.0
deposit=immediately	179	0.475	0.501	0.0	1.0
deposit=none mentioned	179	0.179	0.384	0.0	1.0
itemcountstateby100	179	4.109	4.074	0.0	14.0
itemcountmakeby100	179	2.440	2.497	0.0	8.8
PositiveFeedbackPercent	179	98.872	4.487	50.0	100.0
Feedbackscoreby1000	179	0.593	1.918	0.0	23.6
ModelavgIMVnewold1000	179	20.285	13.253	6.3	95.7

Table 4-6. Robustness Check: Sold (Y/N) Probit Model with Clustered Standard Errors

Concept	**Variable**	**(1)**	**(2)**	**(3)**	**(4)**
Markup	priceprem	-1.371**	-1.371**	-1.371**	-1.371**
		(0.222)	(0.247)	(0.249)	(0.239)
Markup*BO	priceprem*BO	0.636**	0.636**	0.636**	0.636**
		(0.245)	(0.277)	(0.272)	(0.255)
BO	BO	-0.069	-0.069	-0.069	-0.069
		(0.188)	(0.098)	(0.102)	(0.075)
Seller Learning	Selleravgrevisions	-0.002	-0.002	-0.002	-0.002
		(0.051)	(0.058)	(0.052)	(0.053)
Seller Reputation	PositiveFeedbackPercent	-0.005	-0.005	-0.005	-0.005
		(0.009)	(0.008)	(0.007)	(0.008)
Seller's Time	deposit=48 hours	-0.069	-0.069	-0.069	-0.069
		(0.161)	(0.206)	(0.169)	(0.201)
Seller's Time	deposit=72 hours	0.110	0.110	0.110	0.110
		(0.429)	(0.430)	(0.436)	(0.416)
Seller's Time	deposit=immediately	-0.269*	-0.269*	-0.269**	-0.269**
		(0.148)	(0.151)	(0.133)	(0.111)
Seller's Time	deposit=none mentioned	-0.309**	-0.309	-0.309**	-0.309*
		(0.153)	(0.192)	(0.145)	(0.179)
Seller Experience	sellerprivate	0.072	0.072	0.072	0.072
		(0.111)	(0.099)	(0.102)	(0.098)
Seller Experience	Feedbackscoreby1000	-0.006	-0.006	-0.006	-0.006
		(0.009)	(0.010)	(0.013)	(0.011)
Object Valuation	lnIMVby1000	-0.233**	-0.233**	-0.233**	-0.233**
		(0.091)	(0.107)	(0.097)	(0.095)
Object Valuation	Mileageby1000	-0.004**	-0.004**	-0.004**	-0.004**
		(0.001)	(0.001)	(0.001)	(0.001)
Object Valuation	Year	-0.071**	-0.071**	-0.071**	-0.071**
		(0.015)	(0.016)	(0.014)	(0.015)
Object Valuation	ModelavgIMVnewold1000	-0.004	-0.004	-0.004	-0.004
		(0.004)	(0.004)	(0.004)	(0.003)
cons		142.087**	142.087**	142.087**	142.087**
		(29.131)	(31.139)	(27.250)	(29.963)
N		2497.000	2497.000	2497.000	2497.000
r2		0.2436	0.2436	0.2436	0.2436
Clus. by Seller		Y	N	N	N
Clus. by Make		N	Y	N	N
Clus. by Model		N	N	Y	N
Clus. by State		N	N	N	Y

* p<0.10, ** p<0.05

Note: Standard errors and clustered standard errors are reported in the parentheses.
Make dummies included.

Table 4-7. Robustness Check: Sold (Y/N) Probit Model with Alternate Specifications & Collectible/Rare Cars Dropped

Variable	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**
priceprem	-0.769**	-0.774**	-1.327**	-1.326**	-1.325**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
BO		-0.187*	-0.108	-0.103	-0.102
		(0.063)	(0.317)	(0.343)	(0.349)
priceprem*BO			0.606**	0.612**	0.611**
			(0.019)	(0.018)	(0.018)
.003	-0.003				
	(0.870)	(0.979)	(0.963)	(0.952)	(0.957)
PositiveFeedbackPercent	-0.001	-0.002	-0.002	-0.002	-0.002
	(0.957)	(0.862)	(0.868)	(0.843)	(0.844)
deposit=48 hours	-0.042	-0.050	-0.049	-0.049	-0.050
	(0.785)	(0.746)	(0.749)	(0.749)	(0.745)
deposit=72 hours	0.025	0.001	0.074	0.073	0.069
	(0.962)	(0.998)	(0.889)	(0.892)	(0.898)
deposit=immediately	-0.290**	-0.298**	-0.303**	-0.305**	-0.307**
	(0.025)	(0.022)	(0.020)	(0.019)	(0.019)
deposit=none mentioned	-0.321**	-0.335**	-0.327**	-0.323**	-0.323**
	(0.017)	(0.013)	(0.016)	(0.017)	(0.017)
sellerprivate	0.043	0.053	0.060	0.057	0.056
	(0.662)	(0.594)	(0.543)	(0.563)	(0.572)
Feedbackscoreby1000	-0.005	-0.005	-0.005	-0.005	-0.006
	(0.674)	(0.665)	(0.689)	(0.672)	(0.664)
lnIMVby1000	-0.274**	-0.279**	-0.286**	-0.223**	-0.220**
	(0.001)	(0.001)	(0.001)	(0.031)	(0.033)
Mileageby1000	-0.004**	-0.004**	-0.004**	-0.004**	-0.004**
	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
Year	-0.065**	-0.065**	-0.065**	-0.070**	-0.070**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
ModelavgIMVby1000				-0.004	
				(0.291)	
ModelavgIMVnewold1000					-0.005
					(0.281)
cons	129.708**	130.462**	130.841**	140.370**	139.930**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
N	2371.000	2371.000	2371.000	2371.000	2371.000
r2	0.2287	0.2320	0.2327	0.2327	

* p<0.10, ** p<0.05

Note: Standard errors are reported in the parentheses.

Make dummies included.

CHAPTER 5
SUMMARY AND CONCLUSIONS

Although there is extensive theoretical literature concerning bargaining, there are very few empirical studies on bargaining. There is only one other paper that studies the question of bargaining vs. posted prices on eBay Motors (Chen et al., 2016), and this dissertation contributes towards building an empirical complement for the rich theoretical literature on bargaining, as discussed in Chapter 1. Chapter 2 of this dissertation explored a sellers choice to adopt Best Offer (BO) vs. using posted prices. The results from the reduced form approach fractional logit model used indicate some key variables that impact the adoption of BO by automobile sellers. Firstly, sellers are likely to adopt bargaining when they have weaker bargaining power (in the form of lower patience, lower outside options, and/or higher seller competition). Secondly, if sellers are less experienced or have less knowledge about the valuation of the product or about the valuations of the other agents, they are likely to choose BO vs. posted prices. In comparing the results from 2013 and 2016, I do not find that the use of bargaining by sellers decreases over time. In addition, despite higher seller learning impacts the adoption odds of BO negatively, BO is still used pervasively on eBay Motors. One reason for this might be because there are continually new and inexperienced entrants in the eBay Motors market, who are more likely to use BO. Another reason could be that the eBay Motors market sells a sufficient percentage of used cars, whose valuation is difficult, making bargaining a desirable selling mechanism for price determination. Yet, it is also possible that principles of Tversky and Shafir (1992) are applicable here. That is, sellers may be using BO because it has now become the social default. This may explain why relatively more inexperienced sellers adopt BO when they first start out, but perhaps learning and revising against it as time goes on.

The question of sellers' choice of bargaining vs. posted prices on eBay Motors is further analyzed by studying the seller's combined choice of pricing and selling mechanism in Chapter 3. While Chen et al. (2016) study the impact of BO on the price of a car, their approach may have a few drawbacks. I propose accounting for the simultaneity between a sellers choice of pricing or markup and his choice to use BO vs. posted prices using the IV approach. I find that using the IV approach vs. using the OLS approach similar to the approach in Chen et al. (2016) yields differing results, and might be one reason for why the conclusions for Chen et al. (2016) seem at odds with the conclusions of this dissertation. I also propose a seller-level approach (vs. an item-level approach) for this question in order to get a better gauge of seller behavior. Chen et al. (2016) conclude that sellers that use BO also charge higher markups. However, I find that sellers' use of BO (vs. posted prices) decreases markups listed for cars. While the difference in the conclusions of this Chapter and the conclusions in Chen et al. (2016) can be attributed to the fact that the data for these conclusions comes from different time periods and seller behavior may have changed over the course of these years, the methodological differences in the paper deserve close attention.

Furthermore, I find that seller patience leads to higher markups. Combined with the insights from Chapter 2, this suggests that a patient seller chooses not to negotiate, and instead chooses to charge a higher markup. This gives further insight into why the markups might be lower for sellers that use BO vs. markups by sellers who use posted prices.

Finally, in Chapter 4, I study the impact of markup, BO, and the interaction of markup on BO on the likelihood of selling, on the time it takes for a car to sell, and on the transaction price of a car. Unlike Chen et al. (2016), I do not find that BO increases the likelihood of sale, in general. Since I study the impact of interaction term of markup and BO on the likelihood of selling (which they do not), I find that BO only increases the probability of selling for cars that have a markup of at least 10.8%. For all cars that have a markup lower than 10.8%, using BO actually decreases the likelihood of sale. This emphasizes the importance of considering the interaction of BO and markup in the analysis. I suggest that the underlying reason for this effect might be that sellers that use lower markups paired with BO are ones that are trying to sell a lower quality car. It seems then that buyers are able to understand this quality signal, and therefore cars with lower markups listed under BO are less likely to sell.

Furthermore, I find while increasing the markup for a car decreases the probability of the car (from 33.9% to 62%), using BO reduces this negative impact of increasing a markup by 29% (62.9% - 33.9%). While Chen et al. (2016) have also attempted to study the benefits of BO, they did not study the benefits of using BO in interaction with the markup chosen by the seller. Therefore, they are not able to trace this added indirect benefit of BO. This result is even more important when one considers the benefits for of charging a higher markup. I find that a one standard deviation increase in the markup of a car leads to a higher transaction price by 40.3%. Using BO reduces decreases the transaction price of a car by 5%, at the average level of markup. In general though, I find that if the markup on a car is at least equal to 0.13, or 13% higher than the CarGuru estimate for the car, then using BO can increase the transaction price of the car. Therefore, using a higher markup and using BO on a sufficiently high markup has many benefits for the seller. There is one drawback to using BO though: it increases the time is takes for a car to sell by 48.7% for an average level of markup. However, overall, there seem to be large benefits for sellers form using BO, and this would explain why about 80% of the sellers on eBay Motors use BO on eBay at any given point of time.

Finally, using a small dataset, I also find that a one standard deviation increase in markup leads to a lower 2.3% lower discount provided by the seller, conditioned on a car being sold through bargaining. For an average car in the sample, this can mean $384 of savings for the seller, and for the range of cars in the sample, this translates to savings of $11.7 to $3,808. It is unclear (and beyond the scope of this dissertation to test) if a higher markup on a car indicates a higher quality car or if a higher markup is merely a strategy by the seller to *signal* higher quality, whether it is the case or not. However, in general there is some suggestive evidence discussed in Chapter 3 and Chapter 4 both regarding the quality uncertainty of cars in this market, and its impact on consumer surplus. The discussions at the end of both these chapters suggest that use of BO does not seem to be impacting consumer surplus negatively through the suspected channels.This is consistent with the conclusions of Huston and Spencer (2002), which studied collectibles on eBay, and Adams et al. (2006), which studied Chevrolet

Corvettes on eBay Motors, that there does not seem to be a “lemon problem (Akerlof, 1995)” on eBay.

REFERENCES

C. Adams, L. Hosken, and Ya-Ting Newberry, P. Yu. Vettes and lemons on ebay. Social Science Research Network Working Paper, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=880780, 2006.

George Akerlof. The market for lemons: Quality uncertainty and the market mechanism. In *Essential Readings in Economics*, pages 175–188. Springer, 1995.

Steven T Anderson, Daniel Friedman, Garrett H Milam, and Nirvikar Singh. Buy it now: A hybrid internet market institution. *UC Santa Cruz Economics Working Paper*, (565), 2004.

T. Andrews and C. Benzing. The determinants of price in internet auctions of used cars. *Atlantic Economic Journal*, 35, 2007.

Matthew Backus, Tom Blake, and Steven Tadelis. Cheap talk, round numbers, and the economics of negotiation. Technical report, National Bureau of Economic Research, 2015.

A. J.. Berensky. Can we talk? self-presentation and the survey response. *Political Psychology*, 2004.

Helmut Bester. Bargaining versus price competition in markets with quality uncertainty. *The American Economic Review*, 83(1):278–288, 1993.

Y. Bing. Pricing for bargaining? retrieved from http://164.15.27.46/ecaresdocuments/ws/summerschool2009/papers/ye.pdf, 2009.

S. Blount and G. A. Janicik. When plans change: Examining how people evaluate timing changes in work organizations. *Academy of Management Review*, 2001.

William J Burke et al. Fitting and interpreting cragg's tobit alternative using stata. *Stata Journal*, 9(4):584, 2009.

Luis Cabral and Ali Hortacsu. The dynamics of seller reputation: Theory and evidence from ebay. Technical report, National Bureau of Economic Research, 2004.

Kalyan Chatterjee and Larry Samuelson. Bargaining with two-sided incomplete information: An infinite horizon model with alternating offers. *The Review of Economic Studies*, 54(2): 175–192, 1987.

Jong-Rong Chen, Ching-I Huang, and Chiu-Yu Lee. First offers on ebay motors. *First Offers on Ebay Motors (September 23, 2009)*, 2009.

Kong-Pin Chen, Han Yong, Chia-Ching Lin, and Ya-Ting Yu. Bargaining versus posted-price: Evidence from ebay motors. Social Science Research Network Working Paper, http://ssrn.com/abstract=2758262 or http://dx.doi.org/10.2139/ssrn.2758262, 2016.

Ronald H Coase. The problem of social cost. In *Classic Papers in Natural Resource Economics*, pages 87–137. Springer, 1960.

Richard B Coffman. Art investment and asymmetrical information. *Journal of Cultural Economics*, 15(2):83–94, 1991.

Peter C Cramton. Strategic delay in bargaining with two-sided uncertainty. *The Review of Economic Studies*, 59(1):205–225, 1992.

Rachel Croson, Terry Boles, and J Keith Murnighan. Cheap talk in bargaining experiments: lying and threats in ultimatum games. *Journal of Economic Behavior & Organization*, 51(2): 143–159, 2003.

Antoni Cunyat. Strategic negotiations and outside options. *Department of Economic Analysis, University of Valencia, Tech. Rep*, 1998.

Simona Despa. What is survival analysis. *Cornell University Statistical Consulting Unit*, 2010.

Kenneth C Dudley. *Empirical development of a scale of patience*. PhD thesis, West Virginia University, 2003.

Liran Einav, Chiara Farronato, Jonathan D Levin, and Neel Sundaresan. Sales mechanisms in online markets: What happened to internet auctions? Technical report, National Bureau of Economic Research, 2013.

James H Fowler and Cindy D Kam. Patience as a political virtue: Delayed gratification and turnout. *Political Behavior*, 28(2):113–128, 2006.

Anita Gantner. Bargaining, search, and outside options. *Games and Economic Behavior*, 62(2): 417–435, 2008.

Ching-I Huang, Jong-Rong Chen, and Chiu-Yu Lee. Buyer behavior under the best offer mechanism: A theoretical model and empirical evidence from ebay motors. *Journal of Economic Behavior & Organization*, 94:11–33, 2013.

John H Huston and Roger W Spencer. Quality, uncertainty and the internet: the market for cyber lemons. *The American Economist*, 46(1):50–60, 2002.

Ani L Katchova and Mario J Miranda. Two-step econometric estimation of farm characteristics affecting marketing contract decisions. *American Journal of Agricultural Economics*, 86(1): 88–102, 2004.

I Kiyingi, A Edriss, M Phiri, M Buyinza, and H Agaba. Adoption of on-farm plantation forestry by smallholder farmers in uganda. *Journal of Sustainable Development*, 9(2):153, 2016.

Erin L Krupka and Rachel TA Croson. The differential impact of social norms cues on charitable contributions. *Journal of Economic Behavior & Organization*, 2016.

Klaus Kultti. Equivalence of auctions and posted prices. *Games and Economic Behavior*, 27(1): 106–113, 1999.

Xiaohua Lu and Preston R. McAfee. The evolutionary stability of auctions over bargaining. *Games and Economic Behavior*, 15:228–254, 1996.

R Preston McAfee. Mechanism design by competing sellers. *Econometrica: Journal of the econometric society*, pages 1281–1312, 1993.

M. J. Osborne and A. Rubinstein. *Bargaining and Markets*. Academic Press, Inc., 1990.

Leslie E Papke and Jeffrey Wooldridge. Econometric methods for fractional response variables with an application to 401 (k) plan participation rates, 1993.

Motty Perry. An example of price formation in bilateral situations: A bargaining model with incomplete information. *Econometrica: Journal of the Econometric Society*, pages 313–321, 1986.

John Riley and Richard Zeckhauser. Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, pages 267–289, 1983.

Ariel Rubinstein and Asher Wolinsky. Equilibrium in a market with sequential bargaining. *Econometrica: Journal of the Econometric Society*, pages 1133–1150, 1985.

Christian Terwiesch, Sergei Savin, and Il-Horn Hann. Online haggling at a name-your-own-price retailer: Theory and application. *Management Science*, 51(3):339–351, 2005.

Kerem Yener Toklu. Structural analysis of the" best offer" mechanism on ebay. 2014.

Amos Tversky and Eldar Shafir. Choice under conflict: The dynamics of deferred decision. *Psychological science*, 3(6):358–361, 1992.

Tomislav Vukina, Xiaoyong Zheng, et al. Bargaining, search, and price dispersion: Evidence from the live hogs market. *Agricultural & Resource Economics Review*, 39(3):534, 2010.

Ruqu Wang. Bargaining versus posted-price selling. *European Economic Review*, 39(9): 1747–1764, 1995.

Jeffrey M Wooldridge. *Econometric analysis of cross section and panel data*. MIT press, 2010.

BIOGRAPHICAL SKETCH

Kritee Gujral entered the Ph.D. program in economics at the University of Florida (UF) in 2011. She completed her M.A. in economics at the UF in 2013. She graduated *Summa Cum Laude* in economics, with a B.A. degree, and *Cum Laude* in mathematics, with a B.A. degree from UF, with minors in religion and actuarial science.