Distribution Agreement

Claudia Hurowitz March 29, 2019

An Experimental Study on Buyer Behavior in the Presence of a Privacy Risk

by

Claudia Hurowitz

Kelli Lanier

Adviser

Economics

Kelli Lanier

Adviser

Stephen O'Connell

Committee Member

Jeffrey Rosensweig

Committee Member

2019

An Experimental Study on Buyer Behavior in the Presence of a Privacy Risk

By

Claudia Hurowitz

Kelli Lanier

Adviser

An abstract of
a thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Arts with Honors

Economics

2019

Abstract

An Experimental Study on Buyer Behavior in the Presence of a Privacy Risk
By Claudia Hurowitz

As online transactions have become increasingly common practice, firms have invested in the search for personal data in the hope that it will give them a competitive edge. This demand for personal data is met with a concern for personal privacy amongst consumers. I examine this concern by understanding the effects of privacy risks on buyer behavior. Utilizing a posted offer market, my experiment exposed buyers' personal information for any buyer who completed a transaction. Specifically, buyer "money values" were given to sellers for any buyer who purchased a good. The results show a significant difference in buyer behavior in the presence of a privacy risk. When looking at the market as a whole, effects on market convergence were inconclusive.

An Experimental Study on Buyer Behavior in the Presence of a Privacy Risk

By

Claudia Hurowitz

Kelli Lanier

Adviser

A thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Arts with Honors

Economics

2019

Acknowledgements

I thank my advisor, Dr. Kelli Lanier, for her time, input, and continuous support and encouragement that she offered throughout this project. This study would not have been possible without her advice and guidance. I would also like to extend appreciation to my thesis committee members, Dr. Stephen O'Connell and Dr. Jeffrey Rosensweig for their various contributions throughout this past year. In particular, I'd like to recognize Dr. O'Connell for graciously subsidizing this experiment. Finally, I want to thank my friends and family for their unwavering support.

Table of Contents

1. Introduction

In today's data-driven world, consumers face more opportunities to share their personal information through online transactions. The growth of online retail markets has facilitated the ways in which firms gain access to consumer information. This increase in available information has provoked firms to become more active and engaged in the search for consumer data (Odlyzko, 2003; Prince, 2018). In gaining purchase history information about consumers, firms have the power to target consumers and price discriminate while concealing the ways in which they do so (Acquisti and Varian, 2005; Corniere 2013; Conitzer et al., 2012). Survey data indicate that there is a strong societal trend towards consumer distrust in and distaste for firms that perform such data collection activity (Goldfarb and Tacker, 2011; Schmeiser, 2017; Palmatier, 2019). These growing concerns of consumers are met with a deep lack of knowledge on how consumer data is used, who is using it, and how to protect it (Prince, 2018; Palmatier, 2019). This confusion around personal data collection and protection makes it difficult to grasp the true costs and benefits of allowing companies to have access to this information.

Survey data demonstrates that Americans want increased privacy policies (Politico, 2017; Janrain Research, 2018). Sixty-eight percent of American consumers responded "yes" when asked: "The General Data Protection Regulation (GDPR) gives European Union citizens greater control over how businesses can use their personal data. Would you like to see similar laws enacted in the US?".[1] This same study found that 73% of respondents believe that websites know too much about consumers (Janrain Research, 2018). GDPR was enacted in 2018 in Europe. It is widely viewed as the largest overhaul of privacy protection rights in history, giving consumers increased control over their data (Zerlang, 2017). Its open wording allows the regulation to adapt

[1] Janrain Research (2018) surveyed 1,000 United States based consumers.

to the everchanging ways of digital data collection. Among its requirements, GDPR specifically outlines that it is the responsibility of firms to actively engage in cybersecurity practices to lower the risk of cyberattacks (Zerlang, 2017). GDPR also gives more power to consumers. Under this regulation companies are required to inform consumers on how their data is being used, offer consumers clear data collection consent options, and give consumers the right to ask companies about their data collection (EU GDPR, 2018). Consumers also have the right to be forgotten and to withdraw their consent at any time (EU GDPR, 2018). GDPR provides a strong example for how flexible data protection laws can be implemented on a large scale. The demand for similar regulations in America demonstrates that American consumers want increased control and more knowledge of how their data is used.

While this demand for increased privacy protection is felt across all markets, there are ways in which a lack of privacy may benefit consumers. Firms gain an advantage by gathering information about their consumers; it gives them more knowledge of their market. This increased information is beneficial to both firms and consumers alike, as it increases efficiency and lowers consumer search costs (Norman et al., 2016). But this efficient balance of data protection and data searching is a difficult one to determine. Fifty-three percent of Americans surveyed responded “no” to “Are you in favor of web sites or apps using what they learn about you to serve up advertisements that you might find interesting?”. And when asked why not, 49% said “I just don’t like companies watching what I do online” (Janrain Research, 2018). As consumers feel more distrusting towards online data collection, it becomes more difficult for firms to collect data in an efficient yet noninvasive way.

Empirical data and theoretical studies have attempted to understand this complex balance. Results generally show that it is due to asymmetric information regarding firm usage of personal data and heterogeneity in consumer privacy valuations that exacerbate this problem (Acquisti et al., 2013; Acquisti et al., 2016; Poszewiecki, 2016). Today, firms use data on purchase history to determine willingness to pay. Purchase history is of interest here because although consumers do not like it when firms track their purchase history, it also streamlines efficiency for consumers and firms alike (Norman et al., 2016).

I use an experimental setting to determine how consumer purchasing behavior is affected when faced with a privacy risk. This experiment creates a landscape where consumers have full information about how their privacy is used and have knowledge of monetary values are associated with their privacy. Additionally, through the use of a posted offer market, buyers will have the option to forgo participating in the market to maintain privacy at a cost. These two elements together set up an experimental landscape where both buyer preferences and market behavior can be examined under privacy risks. In this experiment I added a treatment to the traditional posted offer market game, where sellers have access to buyers' willingness to pay for any buyer who completes a transaction. With five periods of a traditional posted offer market, subjects became acquainted with the rules of the game and the supply and demand arrays. After five trading periods, the treatment was implemented where any buyer who purchased a unit also gave up their private money values to all sellers in the next period. Buyers could have chosen to avoid a transaction and opt out of the market by rejecting to purchase a unit that would otherwise yield a profit.

Three main hypotheses are of interest. First, I analyze the rate at which buyers chose to opt out of the market. Next, I examine how much profit buyers were willing to forgo by opting out of the market. Lastly, I look at the market as a whole to determine how the demand for privacy impacted the market's convergence to competitive equilibrium. This study shows how buyers behave when they are fully aware that a seller knows their willingness to pay.

Conclusive results of this study show that in the presence of the treatment, buyers are more likely to opt out of the market. This finding demonstrates that buyers will suffer an opportunity cost by rejecting profitable units to maintain their privacy. The results of this study show inconclusive results regarding how much this cost is affected by the treatment. In addition, while the market experienced slight convergence towards equilibrium, it is unclear if this result was due solely to the presence of the treatment.

The layout of this paper will be as follows. Section 2 examines previous literature, and motivates the current study. Section 3 discusses the methods used to run the experiment, hypotheses, and analysis plans for the data. Section 4 presents the results of the experiment. Section 5 offers a discussion of the results including limitations of the study. Finally, Section 6 concludes the paper.

2. Literature Review

2.1 Economics of Privacy Literature

The relationship between privacy and economics has moved to the center of the online privacy debate. Acquisti et al. (2016) examine the economics of consumer personal data protection. Among their many conclusions, they find that consumers lack proper information regarding their personal data collection and can rarely make informed decisions on their digital privacy. Further, Acquisti et al. (2016) indicated that determining a sole economic theory of privacy protection is difficult, as it can be both beneficial and costly to individual and societal welfare.

In looking at firms' demand for consumer data, Hirshleifer (1978, 1980) examines existing assumptions of rational behavior, and argues that information results only in the increased wealth of informed agents, while uninformed agents lose out. This argument provides an interesting take on the role of private information: it can serve as a commodity itself. Those who have access to consumer personal data are at an advantage over those agents that are facing asymmetry about their consumers (Gellman, 2003). More recently, Burke et al. (2011) found that rationally acting firms will over-invest in the searching and collecting of personal information of their consumers. These arguments examine firm behavior and confirm that firms benefit from the collection of consumer information. But firms' demand for consumer data is counteracted by consumers' demand for privacy. While allowing firms to access personal data can provide more complete information, it is difficult to strike a balance that is beneficial to society. Norman et al. (2016) analyzed these opposing forces in more depth. Their research suggests that there are cases in which there is not enough data sharing between firms and consumers. For example, when firms demonstrate too much concern about their reputation regarding privacy policies, they underinvest in the search for consumer data and do not know enough about their consumers. Given the societal trend of consumer preferences towards firms that keep their information private, some firms may overreach in how they protect the data of their consumers.[2] This result demonstrates a case where firms and consumers both would benefit from an increase in personal data disclosure.

Taking a closer look at consumer demand for privacy, there is economic literature that suggests that the desire to maintain privacy could detract from individual and social welfare. The protection of personal information acts as a barrier to perfect information, thus creating market

[2] See Palmatier (2019).

inefficiencies that are detrimental to both buyers and sellers (Posner, 1978, 1981, 1993; Stigler, 1980). Given the demand for consumer data, Poszewiecki (2016) used a survey to determine if consumer valuations of privacy could be determined in monetary terms, and found that privacy can be valued at a price for those that are willing to sell it, but he found that these valuations are non-linear. To make the topic all the more complex, the wide ranging types of personal information means that each individual values different information about themselves differently (Milne et al., 2016). More specifically, one consumer may value their financial information over their contact information, while another consumer may not value their contact information at all. All of these variables make it difficult to formulate a trend in how individuals value privacy.

An additional study by Varian (2002) also discusses the difficulties in reaching a uniformly efficient level of privacy. Varian (2002) finds that consumers may face costs when they protect too much of their information. In this case, sellers face asymmetric information about their consumers. On the other hand, consumers may face costs by disclosing too much information, as they do not know who has control over their data and where it might be going. Here, consumers face asymmetric information. Privacy protection thus makes it difficult to determine an optimal level of privacy. And to complicate this matter more, the meaning of privacy differs greatly between individuals (Acquisti et al., 2013; Poszewiecki, 2016). Plesch and Wolff (2018) ran a field experiment to highlight the privacy paradox where consumers want to maintain their privacy, yet they also want to participate in loyalty programs. Their results found heterogeneity in personal privacy valuations. Rodríguez-Priego and Bavel (2016) similarly found a wide range of privacy valuations in their studies, but found that gender, age, and country of residence played a large role in whether a consumer discloses their information.

While there are conflicting studies examining the supply and demand for personal data, there is still much literature that characterizes the implications of privacy as situational. Taylor and Wagman (2015) set out to characterize the winners and losers of privacy protection enforcement. Yet as they tested several models, they found that their results rely on the specific setting and landscape at play. Similarly, Hermalin and Katz (2006), found that there are many specifics to consider when determining whether privacy protection is beneficial, and whom it benefits. They specifically point to the intricacies of privacy policies and how their details can heavily influence whether privacy protection is efficient. They argue that giving individuals control of their personal data may not be enough to ensure efficient privacy protection. This finding ties back to the notion that consumers rarely understand how their data is being used (Acquisti et al., 2016), making it all the more difficult to create a realistic landscape for consumer privacy.

2.2 Theoretical Literature

Economic theory has attempted to tackle both consumer and seller behavior under these privacy concerns, and there are a few models worth noting here. As transactions move online, tracking purchase history for repeat consumers has become an easy and regular practice for sellers. In a baseline model derived by Acquisti and Varian (2005), it was determined that when applying real life extensions where merchants employ "enhanced services" such as personalized recommendations or one-click purchasing, merchants will find it profitable to condition on purchase history. This finding demonstrates that firms will use consumer data to create their pricing mechanisms. More specifically, firms will condition their prices so that they extract a maximum amount of surplus. In knowing what each consumer has paid for a good in the past, firms can price their products right at each consumer's maximum willingness to pay. Corniere (2013) specifically looked at search engines that give information to advertisers based on consumer search history. Corniere (2013) concludes that targeted advertising reduces search costs, increases matching, and heightens healthy price competition. These results counteract the general trend of increased consumer privacy concerns, shedding some light on the positives of targeting digital consumers. So while Corniere (2013) focusses on the advertising side and competition among search engines, it is important to also consider how consumers might behave in this environment as well. Conitzer et al., (2012) add to these conclusions, by deriving a model where firms can recognize previous customers and use data on purchase history to price discriminate. Focused on the consumer side, they found that when consumers have the option to freely maintain their privacy, they will always choose to do so.

2.3. Experimental Studies

Given the difficulties in evaluating optimal privacy levels, experimental studies can be used to understand how behavior may deviate from theory, and potentially draw new conclusions.

Tsai et al. (2011) conducted an experiment that showed that individuals are more inclined to make online purchases from retailers who clearly display comprehensible privacy policies. Through using a simulated online interface, this experiment displayed indicators with varying levels of privacy for different merchants selling the same good at different prices. This finding indicates that some consumers will pay a premium for enhanced protection when they fully understand the implications of data protection. And interestingly enough, this study found that the presence of comprehensive and clear privacy policies alone, regardless of their content, can influence a consumer's decisions. Hermstrüwer and Dickert (2017) set up an experiment to understand this behavior further through the use of a dictator game.[3] They found that noticeable and important data disclosure consent options are actually motivated by social conformity rather than rational decisions based on complete privacy valuations. This result is strengthened by further experimental research that indicates that expected future convenience affects how consumers make decisions when facing websites with varying privacy policies (Hann et al., 2014).

Paying for privacy or receiving money for disclosure is a new phenomenon that experiments can attempt to understand. As previously discussed, Poszewiecki (2016) implemented a field study to apply a value or formula to personal data. Several experimental studies have gone on the quest to add a price tag to personal data, and while it is clear that individuals can apply a monetary value for their data, the studies have found much heterogeneity in their results (Acquisti et al., 2013; Schudy and Utikal, 2017; Evens and Damme, 2016; Benndorf and Normann, 2017). In this same realm, Grossklags and Acquisti (2007) set up an experiment to understand the monetary differences between the minimum "willingness-to-accept" payment for data and the maximum "willingness-to-protect" their personal data. They found that the average amount consumers would accept to sell their data is significantly higher than average amount they would pay to protect their data. This finding demonstrates that individual consumers themselves do not fully understand how to evaluate their own personal data or how their data is valued by firms (Preibusch et al., 2013). This finding exacerbates the information asymmetry that consumers face as previously discussed.

---

[3] The dictator game measure selfishness through endowing one subject with money and giving this subject the opportunity to share some of their endowment with another subject.

These studies contribute to the study of personal data privacy and have implications for digital privacy policies. However, there is still a considerable amount to be learned about the transaction of money and data, and how consumers understand their privacy.

2.4 Posted Offer Markets

In attempting to understand how the risk of privacy affects consumer behavior, I implemented a posted offer market. A posted offer market is an example of an oligopoly experiment where buyers are price-takers, and it is understood to accurately portray the makeup of typical retail markets (Ketcham et al., 1984). Ruffle (2001) points out, though, that oligopoly theory and experiments mainly focus on the strategic actions of sellers. Buyers are only able to accept or reject a good at a posted price, but Ruffle (2001) takes a closer look at various factors that motivate buyer decision making. Earlier, Franciosi et al. (1995) factor in the sentiment of fairness in the market by disclosing the amount of seller profit to buyers. They found that when buyers are aware of seller profits, posted prices are lower. These findings introduce the idea that buyer behavior can be affected through experimental conditions and that buyers play a more active role than theory suggests (Ruffle, 2001).

These studies also provided important context for procedural decisions in my own posted offer market. Ruffle (2001) examined how an uneven surplus division would affect buyer decision making. His goal was to provoke buyers to opt out of the market.[4] In doing this, he implemented uneven surplus divisions. Because my study also examines buyer behavior and the tendency to opt out of the market, I implemented an even surplus division to ensure I was not affecting buyer behavior through an uneven division. Additionally, Ketcham et al. (1984) set up

[4] Ruffle (2001) refers to this behavior as "demand withholding". For the purposes of this study, I will refer to it as the decision to opt out, or the opt out rate. Opting out refers to buyers who chose to avoid a transaction that they could have otherwise made.

supply and demand arrays that yielded equal consumer and producer surplus in one of the original posted offer market studies.

Many posted offer market studies also aim to look at market convergence to competitive equilibrium (Plott and Smith, 1978; Ketcham et al., 1984; Cason and Williams, 1990; Kujal, 1992,1994; Brannon and Gorman, 2000). Relevant to this study, previous research has shown that convergence in posted offer markets tends to be slower than in double oral auctions (Plott and Smith, 1978; Ketcham et al., 1984).[5] One of the goals of this study is to examine market convergence. These previous studies help to understand the nature of market convergence in posted offer markets, and demonstrate that transactions will likely be off of competitive equilibrium through ten trading periods. And to further support my procedural decision to implement equal buyer and seller surplus, studies have shown that uneven distributions can affect market convergence (Kujal, 1994).

3. Methods

3.1 Experimental Overview

This study consisted of two main parts: 10 periods of a posted offer market and a survey. The posted offer institution was implemented using the online experimental economics platform, Veconlab. The experiment took place in the Economics Department Computer Lab at Emory University, with Emory College students as subjects. Subjects were recruited through in-class announcements and emails, where they were asked to participate in an hour long study and could sign up through an online link. In total, 61 Emory College students were recruited to participate.

---

[5] Double oral auctions differ from posted offer markets in that sellers post prices (asks) and buyers also post bids in a simulated market.

In each session, subjects were randomly assigned the role of buyer or seller. Three sessions consisted of 16 subjects, with 8 buyers and 8 sellers, and one session consisted of 14 subjects, with 7 buyers and 7 sellers. The difference in subject size per session was due to recruitment errors, and will be accounted for in my analysis. Additionally, only 15 subjects showed up to Session 2, so I acted as a seller. Most of my analysis plans to look at buyer behavior, so this decision would not affect those results. Although, I will look at how my participation as a seller affected the market as a whole.

All subjects were given a $3 participation fee. Due to a lack of initial participation, some students were offered extra credit in their classes if they participated. Subjects who received extra credit were also given the participation fee.[6] To ensure that the different incentives would not affect the results of this study, I will include a dummy variable in my analysis that indicates when a subject received extra credit. Subjects were also told at the beginning of the experiment that one trading period from the posted offer market would be chosen at random, and subjects would receive the actual dollar amount earned from the randomly selected period.[7] Subjects were informed that their responses to the survey section would not affect their earnings. After the experiment was complete, subjects received their earnings in private and filled out money receipts that are kept on record.

---

[6] Camerer et al., (1999) found that scaling up monetary incentives does not affect mean performance in this type of experimental setting.

[7] Charness et al., (2016) found that using a pay-all method (where subjects accumulate a cash balance and are paid for every round) affects results. Subjects appear to change their decision making as their cash balance increase, as they experience diminishing returns in each round. This suggests a pay-one method (picking one round to pay) will lead to uniform motivations in each individual round.

3.2 Procedure

When subjects entered the room, they filled out consent forms and waited at their seats until all participants arrived. There were dividers in place to ensure privacy and anonymity. I handed out a set of general instructions that outlined the structure of the experiment and indicated that there would be two main phases. Phase one included the trading periods, and phase two included a survey. At no time did the subjects know how many trading periods were remaining.

The first five trading periods followed a standard posted offer market procedure. Each trading period had two phases. In phase one, only sellers were active. Sellers were given a private cost for one unit. This cost changed in every trading period. In this phase, sellers decided on a price to post in the market given their personal cost. After all sellers submitted their decisions, phase two began. In phase two, all the prices were displayed to all buyers and sellers, and only buyers were active. Buyers were given a private money value for one unit. This money value changed in every trading period. Money values indicate the most a buyer is willing to pay for a unit of a good, so a buyer cannot purchase a unit at a price above their money value. Buyers now had the opportunity to make purchasing decisions based on their personal and private money values. To make purchasing decisions, buyers ranked the units they wished to purchase. Buyers could select "avoid" for any unit they could not or did not want to purchase. The program randomized the shopping order for every trading period. So even though buyers ranked multiple units, they only received the highest ranked unit available. Once every buyer made their decision, the computer program completed all of the possible transactions and displayed the results to all subjects. Subjects that completed a transaction received a profit. Profits were computed as follows:

Seller's Profit = (Transaction Price) - (Cost of Unit)

Buyer's Profit = (Money Value) - (Transaction Price)

Sellers did not incur the cost of a unit unless they sold that unit. Neither buyers nor sellers could make negative profits.

After five periods of a standard posted offer market, I implemented a treatment. All subjects were told that from that point forward there would be an addition to the experiment: if a buyer purchases a unit in a trading period, their money value for the next round will be given to all sellers at the beginning of the next round. This means that if a buyer purchases a unit, the amount they would be willing to pay for a unit in the next round would be available to sellers. After handing out the instructions and reading them out loud, the sixth trading period began. At the end of each subsequent period, I wrote down the personal money values of the buyers who purchased a unit in the previous round on folded pieces of paper. Folded papers were given to everyone in the room to maintain anonymity, but only sellers received the information. Buyers received blank pieces of paper. Once the information was distributed, sellers were told they could post their prices. This procedure was repeated until the end of the tenth period.

After the trading periods ended, I handed out the survey. The survey included demographic questions such as gender, age, and academic major as well as information regarding experience in economics. I then asked some questions about data protection and online privacy. I surveyed the subjects on this topic to gain an understanding of how participants viewed and valued their own privacy. Given that modern data and surveys indicate that people are concerned about their privacy, I wanted to ensure that my subject pool accurately represented this trend.

3.3 Hypotheses

Three main hypotheses emphasize the goals of this study. The treatment to this experiment added in a privacy risk for buyers. As previously stated, buyers' private money values were exposed for any buyer that completed a transaction after the treatment was implemented. The first two hypotheses look at buyer behavior in the presence of this privacy risk. The third hypothesis examines how the privacy risk affected market convergence.

Hypothesis 1: After the treatment is implemented, buyers will be more likely to opt out of the market.

Hypothesis 2: After the treatment is implemented, buyers will reject units that yield higher profits than before the treatment.

Hypothesis 1 addresses the number of buyers who will "opt out" of the market. A buyer opts out when they enter "avoid" on a unit that they could have purchased. This means that a buyer is rejecting a unit that would otherwise yield a profit, incurring an opportunity cost. I will look at the amount of buyers that choose to opt out of the market in each trading period. With 31 buyers making purchasing decisions in 5 periods before the treatment and 5 periods after the treatment, there are 155 decisions to look at both pre and post treatment.

Hypothesis 2 addresses the amount of profit buyers reject by choosing to opt out. I will look at those buyers that opted out and the amount of profit that was forgone. For example, if a buyer has a money value of $5.00, but selects "avoid" for a unit with a price of $4.00, the amount of profit rejected for that period is $1.00. In other words, by opting out they chose to forgo the earnings of $1.00. I will look at the highest amount of forgone earnings for each opt out, and I suspect that this dollar amount will increase after the treatment.

Both of these hypotheses require me to look at the decisions buyers made, rather than the actual transactions that occurred. The decisions behind the transactions provide me with more valuable information because of the randomized order of the buyer's shopping positions. A buyer who does not make a transaction is not the same as a buyer that opts out because a buyer who opts out may actually make a transaction. Opting out here refers to the decisions behind the transactions. Opting out only occurs if a buyer rejects a unit that is priced anywhere below their money value.

Hypothesis 1 and 2 look only at buyer behavior, thus only buyer data is of importance here. These hypotheses are backed by theoretical literature, and survey data that suggests that consumers do not like for their information to be given to firms as previously discussed. I assert that in the presence of a privacy risk, there will be an increase in the opt out rate as buyers will now have an additional factor to consider. This assertion can be demonstrated through theory. Conitzer et al. (2012) derive a set of equations that most closely describe the behavior drawn out in this experiment. They examine monopolistic behavior, characterizing the price as p^1 in the first period. I will adjust their model moving forward to account for the multiple sellers in our experiment, such that prices in the first period will be p_i^1, where *i* denotes the seller.[8] In their Nash Equilibrium, every consumer is anonymous in the first period, and has private money values of v_i, where *i* denotes the buyer. To allow for the change in money values for each period in this study, I will adjust the model such that v_i^1 is buyer i's value in period 1. In their analysis, similar to my experimental procedure, if a consumer purchases a good in their first period they will no longer remain anonymous in the second period. Conitzer et al. (2012) create a model that

[8] In their model, Conitzer et. al (2012) also allow for the monopoly to set different prices for buyers who are anonymous and for buyers who are not. My adjustment of their model only allows sellers to post one price to fit the behavior in this study.

allows consumers to purchase goods and also purchase anonymity. What they ultimately find is that consumers will pay to protect their anonymity up to a point where they no longer benefit from paying for it.

Now, this is where I will deviate from their model to more closely fit the behavior of subjects in this study. In the presence of the treatment, buyer *i* who purchases in the first period with money value v_i^1 will gain the profit of $(v_i^1 - p_i^1)$. If buyer *i* does not purchase in the first period, they will have a utility of $(a_i - c_i^1)$, where a_i is the personally evaluated benefit of remaining anonymous and c_i^1 the cost of opting out in period 1.[9] More specifically, a_i can be seen as how an individual values the privacy of their information. For the purpose of this paper, c_i^1 can be thought of as the opportunity cost of profits that could have been made by opting in. A buyer will choose to purchase if:

$$(v_i^1 - p_i^1) \geq (a_i - c_i^1).$$

This proposition differs from the traditional idea that buyers will purchase if their money value (*v)* is greater than the price of the good (*p*), such that $v_i^1 \geq p_i^1$ for buyer *i* in period 1. Given that the first five periods will be a traditional posted offer market, I will likely see this behavior. But, like Conitzer et al. (2012), I predict that consumers will accept the opportunity cost to maintain their privacy once the treatment is implemented. More specifically, consumers will want to opt out of the market. Additionally, after beginning the treatment, buyers will reject units that yield high profits, because they will now have to consider how they view the benefit of privacy (determining their a_i).

[9] This is a deviation from the model derived by Conitzer et. al (2012), who looked at a constant cost of maintaining anonymity for all consumers.

Hypothesis 3: The convergence coefficient will decrease significantly after sellers have access to buyer information.

Hypothesis 3 examines how market convergence to competitive equilibrium is affected due to the presence of the treatment. The convergence coefficient is one of a few ways to quantify market behavior.[10] Originally formulated by Smith (1962), the convergence coefficient is the ratio of the standard deviation of transaction prices to the equilibrium price for each period expressed as a percentage. A large convergence coefficient value indicates that transaction prices were far off from the predicted equilibrium price, while a smaller value indicates transactions were priced closely to the equilibrium price. In other words, a small convergence coefficient value means there was high convergence for a period. In this study, the supply and demand arrays aligned such that there was an equilibrium range from $4.00 to $6.00 for all trading periods. For the purposes of analysis, the predicted equilibrium price will be thought of as $5.00, as that would yield equal consumer and producer surplus.

The analysis for Hypothesis 3 requires me to look at both buyer and seller behavior alike. It is important to note that previous traditional posted offer market studies have found that market convergence increases with each trading period regardless of any treatments (Ketcham et al., 1984). This increase in convergence is because sellers naturally post prices closer to equilibrium as they learn about buyers' willingness to pay throughout the trading periods. I suspect that the results of the study will show not only that market convergence increases, but that there will be a sizable jump in convergence after the treatment is implemented.

[10] Outlined by Ketcham et al. (1984), market efficiency is another indicator of market behavior towards equilibrium. Market efficiency was not chosen for this study because it analyzes whether all the gains from exchanges are exhausted. This is not of interest to the results related to this study.

4. Results

The total population consisted of 61 undergraduate students from Emory University, with 29 female and 32 male subjects. Economics and Business majors made up about 68% of the population, although the average number of economics courses taken by the subjects was only 3, indicating that subjects in general did not have a strong background in economics. The subjects' ages ranged from 18-23, and sophomores and juniors made up 60% of the population. Thirty-eight subjects received extra credit in addition to the $3 participation fee. To understand the effects of the extra credit incentive, I added in an extra credit dummy variable to my linear regression analysis. As shown in Table 4, the extra credit incentive did not have a significant effect on the results. In addition, none of the other demographic characteristics had a significant effect on the results.[11]

The survey portion of the experiment included a few questions regarding the subjects' attitude towards personal data and privacy. The questions relevant to this study were:

1. Are you concerned with the idea that organizations may hold onto your personal data?
2. Would you be more inclined to purchase from an online retailer that makes efforts to protect your personal information?
3. Do you feel you should be doing more to protect your online information?

Each question had the option to indicate "yes" or "no". For all of the above questions, the majority of responses were yes, with 68.8% responding yes to Question 1, 88.5% to Question 2, and 83.6% to Question 3. This indicates that the subject pool felt protective over their personal information. Implementing a dummy variable for each question indicating 1 for "yes" and 0 for "no", I found that responding "yes" to these questions was positively related to the likelihood of

[11] See Table 4

opting out. As I will show in more detail later, the responses to Question 3 were significant indicators of the whether or not a buyer would opt out.

4.1 Hypothesis 1: After the treatment is implemented, buyers will be more likely to opt out of the market. – **Supported.**

To discover the results related to Hypothesis 1, I examined buyer behavior and ranking decisions, not complete transactions. I determined that a buyer opted out of the market if they selected “avoid” for units that would have otherwise yielded a profit for the round. Because I examined the decisions across all sessions with 31 buyers for 10 rounds, I looked at a total of 310 decisions made. One hundred and fifty-five decisions were made both before and after the treatment. Paired t-tests were performed to measure the differences in behavior by comparing the decisions before and after the treatment,. The *optout* dummy variable indicated 1 when a buyer opted out, and 0 when a buyer did not. Therefore, a positive relationship to the *optout* variable would indicate an increase in the likelihood of opting out. There were 16 opt outs in trading periods 1-5, and 66 opt outs in trading periods 6-10. The opt out rate simply looks at the total number of opt outs over the total number of decisions made before and after the treatment.

One of the main goals of this experiment was to understand whether consumer behavior changes when private information is at risk. As shown in Table 1, adding in a privacy risk has a strong impact on how buyers make purchasing decisions. The two-tailed t test shows that there was a significant difference in the presence of the treatment. These results fall in line with the theoretical predictions, as buyers now have to consider how they value their privacy. When buyers add in their personal privacy valuations, there is a large amount of demand withholding in the market.

Figure 1 demonstrates the contrast per period in the total number of opt outs. With 31 buyers per trading period, the number of buyers who opted out increased from a small proportion to almost half of the total decisions made per period after the treatment. There is no evidence that buyers are more or less likely to opt out as the trading periods continue, but the initial presence of the treatment resulted in a large and significant increase.

Table 1: Paired t-test Results of the Opt-Out Rate Before and After the Treatment

	Before Treatment (n=155)	After Treatment (n=155)	
Total Number of Buyers Who Opted Out	16	66	
Opt-Out Rate	10.3%	42.6%	P Value: < 0.001

Note: This table reports the total number of buyers who opted out and the Opt-Out Rates across all 4 sessions before and after the treatment. The Opt-Outs Rates were calculated by dividing the total number of opt outs by the total number of decisions (155) both before and after the treatment. In addition, this table reports the results of the paired t-test to the determine statistically significant difference between the Opt-Out Rate before and after the treatment.

A similar conclusion can be drawn through performing OLS linear regression analysis. Table 4 shows the effects of the presence of the treatment as well as the effects from the survey questions as previously described. The dummy variable *treatment* is 1 in the presence of the treatment, and responses to each question are noted as variable Q1, Q2, and Q3. In addition, Table 4 shows the insignificant effects of the extra credit incentive and the demographic differences. As concluded earlier, the presence of the treatment significantly affects the likelihood of opting out; in Table 4 the treatment is shown to increase this likelihood by 30%. Another interesting result comes from the responses to Question 3 (“Do you feel you should be doing more to protect your online information?”). There is a significant correlation between the "yes" response to Question 3 and the likelihood of opting out. Additionally, it was generally the same group of subjects that repeatedly opted in when their privacy is at risk. Going back to the theory described in the Methods sections, these subjects likely place little to no value on their privacy in this setting, meaning they will almost always choose to opt in. The consistency of these behaviors indicates that subjects were engaged with the experiment and clearly understood the instructions.

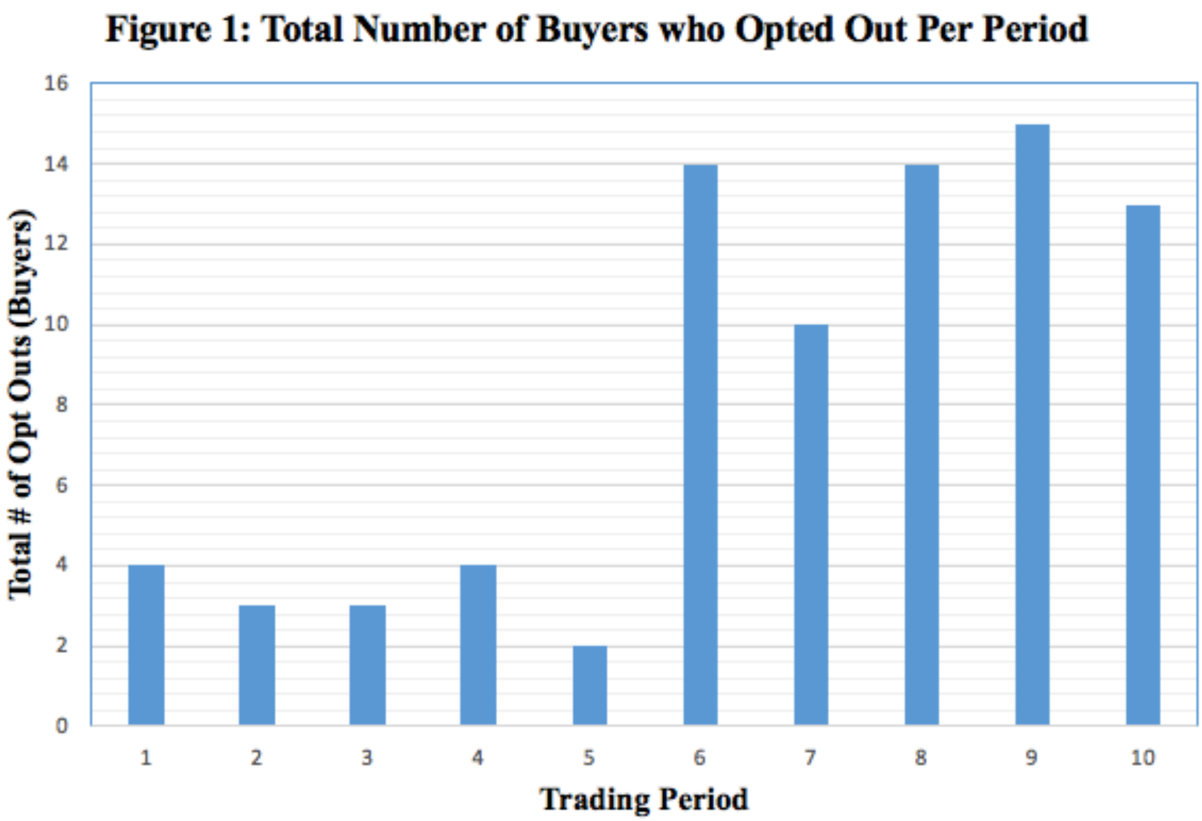

Note: This figure reports the total number of buyers who opted out per period across all 4 sessions.

4.2 Hypothesis 2: *After the treatment is implemented, buyers will reject units that yield higher profits than before the treatment.* – **Inconclusive**

To analyze the results of Hypothesis 2, I looked at all of the decisions to opt out and examined the highest amount of profit that was rejected by buyers. This analysis is noteworthy because it shows how much money a buyer is willing forgo to maintain their privacy. As shown in Table 1, there were 16 opt outs before the treatment, and 66 opt outs after the treatment. Across all 10 sessions there were 82 opt outs, with the average rejected amount $0.94 and a standard deviation of $0.62. The minimum amount rejected was $0.01 and the maximum amount rejected was $2.50. Table 2 shows a summary of these results before and after the treatment.

Although the amount of buyers who opted out was higher after the treatment, the average amount that was rejected remained about the same indicating no significant difference. In fact, a two-tailed t-test proves the difference in these means to be insignificant. Table 2 also highlights these results. The insignificant difference here can be attributed to the limited options buyers have. As previously stated, buyers are simply price takers in a posted offer market setting. The *optout* dummy variable (analyzed in Hypothesis 1) more accurately depicts this behavior, as buyers have only a binary option to accept or reject an offer. Figure 2 shows the average rejected profit per period. There is no particular trend in rejected profit across the trading periods.

OLS linear regression analysis also demonstrates this result. Variable *highrej* denotes the highest rejected surplus for every buyer that opted out. This variable indicates the most a buyer was willing to forgo in order to opt out of the market. Shown in Table 3, the presence of the treatment did not have a significant effect on the amount of rejected profit.

Figure 2: Average Rejected Profit Per Round

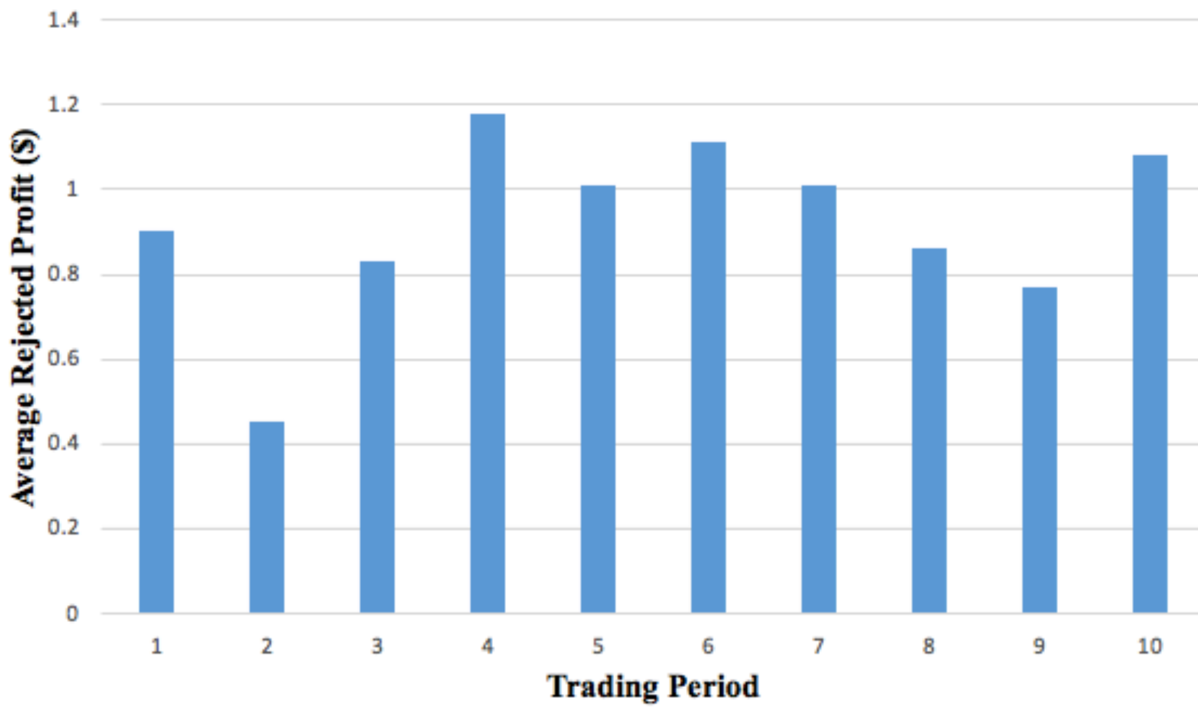

Note: This figure reports the average amount of rejected profit across all 4 sessions per period.

4.3 Hypothesis 3: The convergence coefficient will decrease significantly after sellers have access to buyers' information. – **Inconclusive**

Here, I examine the actual transaction prices to see how they converge to equilibrium. It is important to note that while the treatment was implemented between trading periods 5 and 6, sellers received buyer information beginning only after period 6. At the beginning of period 6, buyers were informed that their information was at risk based on their purchasing decisions from that point on. Because sellers move first in this game, it was not until after period 6 that sellers were able to see buyers' money values. This analysis looks at the difference in market convergence before and after period 6.

Table 2: Paired t-test Results and Summary Statistics of Rejected Profit Before and After the Treatment

	Before Treatment n=16	After Treatment n=66	
Minimum Rejected	$0.11	$0.01	
Maximum Rejected	$2.00	$2.50	
Average Rejected Profit (Standard Deviation)	$0.89 (0.5168)	$0.95 (0.6463)	P Value: 0.5180

Note: This table reports the minimum and maximum amount of rejected surplus of those buyers who opted out before and after the treatment. In addition, this table reports the average amount of rejected profit across all 4 sessions before and after the treatment, with the paired t-test results between the two means.

Table 3: OLS Regression on Likelihood of Opting Out and Highest Rejected Surplus in Presence of Treatment and Survey Responses

Variables	(1) *optout*	(2) *highrej*
Treatment	.3086***	.0917
	(.0466)	(.1842)
Q1	.0109	.0957
	(.0685)	(.2033)
Q2	.04994	-.1641
	(.0831)	(.2688)
Q3	.2651***	-.0224
	(.0675)	(.4436)
Extra Credit	-.0772	.1151
	(.0659)	(.2478)
Gender	-.0138	.0965
	(.0502)	(.1601)
Age	-.0255	.0755
	(.0244)	(.0838)
Constant	0.421	-.6797
	(.4926)	(1.6531)
Observations	310	82
Adjusted R-Squared	0.1526	-0.0542

Note: Columns 1 and 2 report the coefficients (and standard errors in parentheses) from a regression analysis on the likelihood of opting out and the highest amount of rejected profit. The *optout* dummy variable represents a 1 if a subject opted out, and a 0 if a subject did not. The *highrej* variable notes the highest amount of profit a subject rejected. A dummy variable for the Treatment represents a 1 in the presence of the treatment, and a 0 without the treatment. The Q1, Q2, and Q3 dummy variables correspond to the responses to survey questions 1, 2, and 3 respectively and all represent 1 in the presence of a "yes" response. The gender dummy variable is a 1 if the subject was a female, and 0 if the subject was a male.
***p≤. 01, **p≤ .05, *p≤ .10

To analyze these results, I calculated the convergence coefficient for every trading period across all sessions.[12] The results are inconclusive in defining a relationship between market convergence and the implementation of the treatment. Figure 3 demonstrates the trend in the convergence coefficient for each of the 10 trading periods, averaged across the four sessions. The average convergence coefficient for period 6 was 19.72%, and only saw a slight decrease to 16.45% for period 7. In addition, there is not a downward trend in the convergence coefficient after round 7, indicating that convergence may not have been due to the treatment, but to traditional posted offer market behavior.

Figure 3: Average Convergence Coefficient

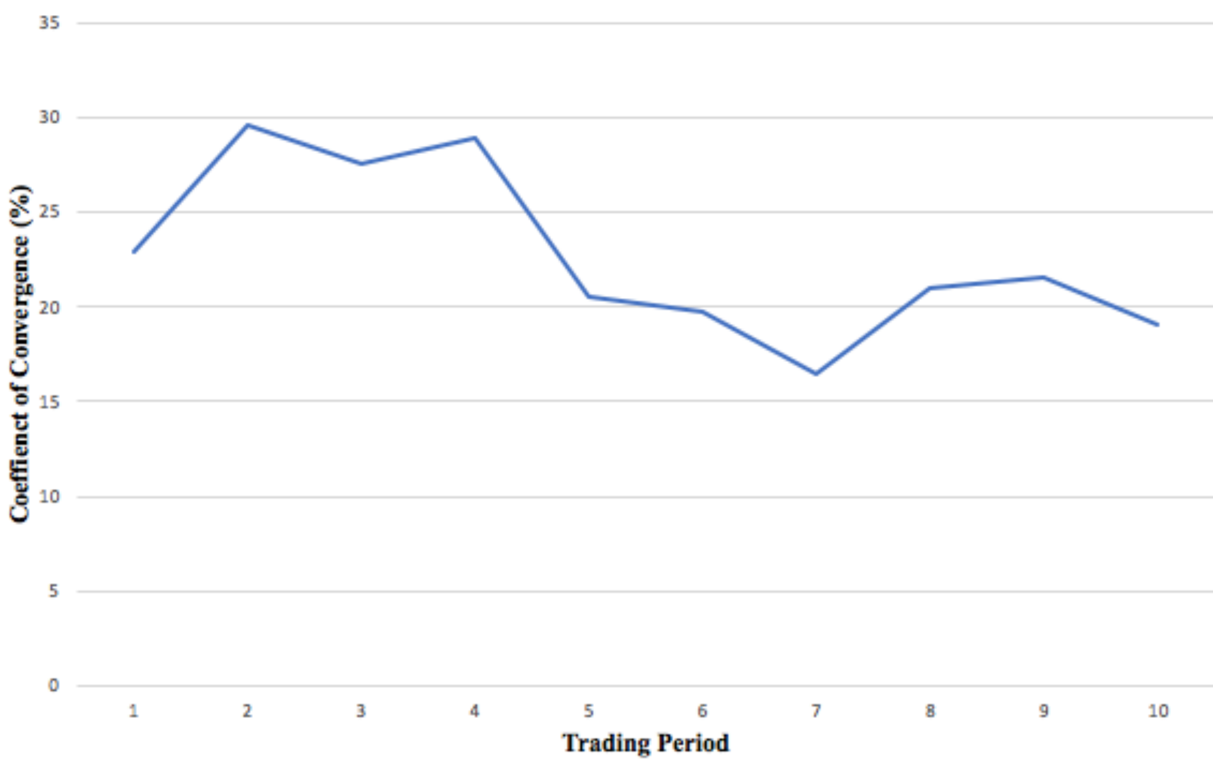

Note: This figure reports the average convergence coefficient across all 4 sessions per period.

[12] See Methods section for coefficient of convergence equation.

OLS linear regression analysis was also performed to see the effects of the information. Additional variables are included in this analysis to determine whether session size (14 or 16 subjects) and my participation affected results.[13] Table 4 highlights these results. As previously noted, there was no evidence that seller access to buyer information had a significant effect on market convergence.

Table 4: OLS Regression on Convergence Coefficient in the Presence of Increased Information

Variables	(1) *ConvCoef*
Information	-6.467 (3.691)
14 Subjects	-7.453 (4.429)
Experimenter Participation	-10.232* (4.428)
Constant	27.663*** (2.952)
Observations	40
Adjusted R-Squared	0.1399

Note: Column 1 reports the coefficients (and standard errors in parentheses) from a regression analysis on the Convergence Coefficient. The *ConvCoef* variable represents the Convergence Coefficient (See Methods section for equation). A dummy variable for Information represents a 1 in the presence of buyer information, and a 0 without the buyer information. The dummy variable for 14 Subjects is a 1 for the session with fewer subjects. Similarly, the dummy variable for Experimenter Participation is a 1 for the session with my presence in the experiment.
***p≤. 01, **p≤ .05, *p≤ .

[13] One experimental session had 14 subjects instead of 16. And in one session of 16 subjects, I acted as a seller to ensure there were an even amount of buyers and sellers.

5. Discussion

5.1 Limitations

There were a few limitations in this study that are worth noting as they may have impacted the results. First, there were a limited amount of trading periods in the study due to resource constraints. Previous posted offer markets have implemented more trading periods (around 25 to 30)[14]. These studies have mainly focused on market convergence to equilibrium, and have required more than 10 trading periods to demonstrate that trend (Ketcham et al., 1984). No conclusive results were discovered in regards to market convergence perhaps due to this limitation. Additionally, previous studies have shown that behavior in a posted offer market setting may differ in the short run and the long run (Franciosi et al., 1995; Ruffle, 2001). A useful extension to this study would implement more trading periods to understand the effects on the market as a whole and the long term effects.

Second, the study would have benefitted from a larger sample size to account for outliers in the data. There was clearly confusion among some of the subjects in the first few periods of the posted offer market. The set-up of the game allowed buyers to rank the units they wished to purchase, or select "avoid" for any unit they could not or did not wish to purchase. The ranking system was set up such that buyers selected "1st Priority", "2nd Priority", or "3rd Priority". But, as explained in the instructions, buyers could select multiple units under each priority. There was evidence that some buyers misunderstood that more than 3 units could be selected in the first few rounds. This misunderstanding led those buyers to seemingly reject high profit yielding units, or opt out before the treatment. Due to the limited sample size, I was unable to discount the data of subjects who had this confusion. For example, as shown in Table 3, one subject rejected a unit

[14] See Ketcham et al., 1984.

that yielded $2.00 in profit before the treatment. This rejection is unusual as that is a comparatively high profit to gain from a trading period. Additionally, that unit was rejected in the first round, further supporting the notion that there was confusion in the earlier rounds. Including more subjects would have allowed me to indicate these outliers and possibly strengthen results. And because the analysis for Hypothesis 2 looked at only those subjects that opted out, the analysis used an even smaller sample size than the analysis for Hypothesis 1. The lack of evidence supporting Hypothesis 2 likely indicates that more subjects, and perhaps a few practice rounds, would show different results.[15] An increase in the subject pool would have also strengthened the results for Hypothesis 1. If there were opt outs prior to the treatment that were due to confusion, minimizing this confusion would have resulted in fewer opt outs. The significant difference in opt outs before and after the treatment would have been even greater.

Finally, there were large limitations regarding the incentive and compensation mechanisms in this study. I received some feedback from subjects indicating that while the randomized per period monetary payoff incentivized profit maximizing in each trading period, the payoffs were not large enough. One subject in particular noted that by the time the treatment came along, she was aware that not enough money was at stake for her to opt out of the market and remain private. While there is research that shows that scaling up the monetary incentives would not make a difference on behavior (Camerer et al., 1999), the direct feedback from subjects suggests that perhaps scaling up the payments would have made a difference. Extending this experiment such that there are larger potential earnings or a wider possible spread could be informative. In addition, there is research that shows that paying subjects before the experiment

[15] I performed an additional analysis without the first two rounds of data for all hypotheses to account for possible confusion among subjects. I found no large significant difference in results when throwing out that data, suggesting that other limitations, in addition to confusion, were the cause of the inconclusive results.

would motivate subjects to minimize monetary losses (Rosenboim and Shavit, 2011).[16] The supply and demand arrays in this study were arranged such that buyers and sellers started with a hypothetical $10.00. This was implemented so that buyers always had enough hypothetical funds to make a transaction, but subjects earned only the net profit from each period. For example, if a buyer made a profit of $1.50 in one period, they would hypothetically have $11.50 but would only receive a payoff of $1.50 if that period was chosen. While this mechanism reduces unwanted income effects, subjects may feel that they are using money that is not their own. An interesting addition to this experiment would be to endow subjects with the $10.00 prior to the experiment to encourage real life behavior.

5.2 Discussion of the Results

The results from this study indicate that when consumers feel their private information is at risk, they are more likely to opt out of the market even when facing an opportunity cost. This opt out behavior is preferable to buyers up to a certain cost that varies from buyer to buyer. It is inconclusive on what affects that preference as the effects from the treatment do not statistically support any evidence of a relationship. Additionally, there is no noticeable trend unique to this study that suggests any effect on market convergence as subject pool and trading period limitations made that difficult to determine. There are additional noteworthy results that come from the survey answers, as that data alone is very indicative of buyer behavior.

There was strong statistical evidence that the presence of the treatment affected the likelihood of opting out of the market. The risk of losing one's privacy here could be thought of

[16] Rosenboim and Shavit (2011) examine the difference between a prepaid mechanism two weeks prior and payment at the start the experiment. They found that a prepaid mechanism incentivized subjects to use the money as if it was their money all along. It is important to note that this study did not examine behavior in posted offer markets, rather it provides useful insight into various payment mechanisms.

in a few ways. As subjects learned about the posted offer market, it likely became clear that the buyers' advantage was their private money values. Typical of any posted offer market, sellers have no knowledge of their buyers' preferences in the first round. This lack of knowledge led to some posted prices that were impossible for any buyer to purchase (prices up to $100.00). Having a lack of information about the market leads to market inefficiencies such as these, so it is to the benefit of sellers to gain the information about their buyers' willingness to pay. As sellers face uncertainty when determining how to price their units, buyers have the advantage. So, when buyers learn that their money values will be exposed, they understand their advantage is at risk, leading them to opt out of the market.

Falling in line with the results of Franciosi et al. (1995), there is also an element of fairness that buyers have to consider in this market. Buyers may simply find it unfair for sellers to have their information. This explanation of the data is also consistent with the societal trends of distaste in how firms have access to personal data. To further support this notion, the responses to the survey questions offered very indicative results. In particular, the responses to Question 3 ("Do you feel you should be doing more to protect your online information?") strongly indicated whether the buyer was likely to opt in or out. Those subjects that responded "yes" were very likely to opt out of the market. The presence of the treatment and the responses to Questions 3 were the only significant indicators of the likelihood of opting out, and both offer noteworthy explanations of what affects consumer behavior.

In looking at the amount of profit buyers would forgo to maintain privacy, there were no conclusive results. The assumption was that prior to the treatment, buyers may reject some profit yielding units due to a feeling of unfairness or confusion.[17] But after the treatment was

[17] See Franciosi et al., 1995; Ketcham et al., 1984.

implemented, buyers would reject units that yield even higher profits because they have to consider the cost of losing their private information. As previously mentioned, there are some experimental limitations that could have affected these results. Further, the price taking nature of buyers in a posted offer market setting makes this result difficult to determine. The binary decision of buyers is more accurately depicted by the dummy variable *optout* (whether or not a buyer opted out), rather than *highrej* (the highest amount of profit rejected in order to opt out). An interesting extension to this study would be the implementation of a double oral auction. While double oral auctions are not generally thought of as accurate representations of real life retail markets, it could provide more insight for this result due to the negotiating aspect of the auction.

Additionally, previous literature on the economics of privacy has demonstrated that it is difficult, if not impossible, to evaluate privacy at a price (Varian 2002; Rodríguez-Priego et al., 2016; Acquisti et al., 2013; Poszewiecki, 2016; Plesch and Wolff 2018). This notion is largely based on heterogeneity in personal privacy valuations. Given that everyone values their privacy differently, it would be difficult to find a relationship between rejected profit and the risk of privacy in this study. This notion is also demonstrated in the data. The average rejected profit after the treatment was \$0.95 with a standard deviation of \$0.65. This suggests that there was large variability in how much money subjects were willing to sacrifice to maintain privacy. Given that previous studies have been unable to locate a trend in how this cost can be calculated, it is not unusual that these results were inconclusive.

Finally, there were a few experimental limitations that impacted market convergence results as previously stated. The inconclusive results show that market convergence did increase throughout the rounds, but it is unclear if this convergence was prompted by the presence of the treatment. Additionally, the session with my participation was more likely to converge than the other sessions. This result is possibly due to my own knowledge of the market, but in looking at the data, I completed fewer transactions than 75% of sellers for that session, indicating that this session may have experienced more convergence regardless of my presence. Looking more closely at what could have affected these results, the limited number of rounds became an issue for this analysis. Additionally, in running the experiments, I noticed a trend that could have actually slowed market convergence. After the treatment, the maximum number of buyers to complete a transaction for a single period was 4. With 4 buyers completing transactions, there were 4 money values given to sellers. I began to notice that when I would fill out the money value information, I would often get a combination of money values off of equilibrium. For example, in one round I gave sellers money value information where the values were: $8.00, $7.00, $3.00, and $3.00. While these money values give sellers information about the market, they also allow sellers to condition their prices based on these values. In that trading period, all but one of the posted prices were well above equilibrium, seemingly aiming to target those buyers with money values of $8.00 and $7.00. This behavior would in fact slow market convergence, as it seems sellers went too far in their price conditioning techniques. It is this behavior that could have opposed the natural market convergence behavior seen in posted offer markets elsewhere. While the goal of this study was not to necessarily examine seller behavior, it supports previous economic results that sellers will sometimes go too far in their price conditioning behavior in the presence of consumer data.

A useful extension to this study would be to gain a greater understanding of how order effects impacted the results. Implementing the treatment first and then removing the privacy risk would have possibly increased market convergence. In traditional posted offer markets, market convergence is slow as sellers gain an understanding of the range of their buyers' money values. If sellers receive buyers' information in the first few rounds, they would face less uncertainty in the initial trading periods. After removing the privacy risk, sellers would have strong knowledge of competitive equilibrium although they would not have the power to price discriminate. This would likely lead to prices to converge more closely to equilibrium. Additionally, this would have an impact on buyers' decision making. I suspect that buyers would continue to consider the element of fairness and would opt out across all periods, even in the absence of the treatment. The advantage to the sellers in the first few periods would be strong enough to provoke a feeling of unfairness throughout the entire experiment. In this case, the opt out rate would decrease, although I imagine that the difference in opt outs would be insignificant.

6. Conclusion

The goal of this study was to understand consumer behavior in the presence of a privacy risk. To approach this understanding, I implemented a posted offer market and examined buyer behavior when private information was at risk of being exposed. After playing 5 traditional periods of a posted offer market, a treatment was implemented and subjects were informed that if a buyer completes a transaction, their personal and private money value will be shown to all the sellers for the next period. This procedural set up is inspired by real world applications, where firms use consumer data and purchase history to target and price discriminate.[18] While economic literature on privacy and theory has attempted to grasp how consumers value and view their privacy, most results find that there is little uniformity in how individuals feel about their

[18] See Odlyzko, 2003; Hann et al., 2014.

privacy. This study aims to learn more about consumer behavior in a setting where sellers have information on willingness to pay.

Some noteworthy and important conclusions can be drawn. In this posted offer market setting, buyers have the option to either accept or reject a unit to purchase at a posted price. The results showed a significant increase in the rejection rate of buyers, otherwise referred to as the opt out rate. This significance indicates that buyers are willing to forgo a profit to maintain their private information. Additionally, survey data was also a significant factor in describing this behavior. Those who indicated they needed to do more to protect their online privacy were 27% more likely to opt out of the market. Also, those that indicated little concern for their online privacy were generally the same subjects who opted in to the market in the presence of the treatment. These results show that personal privacy valuations play a large role in purchasing behavior.

Other results of the study were inconclusive. This shortfall is mostly due to experimental limitations and possible learning effects. In looking at how much money buyers were willing to forgo to opt out of the market, there were no significant trends before and after the treatment. The inconclusive results can also be explained by the price taking nature of buyers in posted offer market settings, where buyer decision making is a binary choice. In addition, there were no conclusive results on how the market convergence was affected as a whole.

There are important implications of this study. First, it would benefit firms to understand how their targeted demographics view personal data privacy. In understanding how consumers feel about their personal data, firms can find an efficient balance of using consumer information and keeping this information private. Firms today either overinvest in the search for consumer data or overinvest in the protection of consumer data (Burke et al., 2011; Norman et al., 2016).

Because personal privacy valuations are indicative of purchasing decisions, firms that understand how to navigate their privacy policies to fit their consumers would have an advantage in the market. The General Data Protection Regulation (GDPR) in Europe provides an excellent example of how strong data privacy laws can benefit consumers without demanding requirements that are too strict and costly for firms. And secondly, consumers that value their privacy are willing to suffer opportunity costs to maintain their privacy about willingness to pay. Information on willingness to pay differs from other types of personal information, because it informs firms on pricing techniques. While it is unclear how much firms benefit from this price tracking and discrimination, the results suggest that this sort of data tracking would not be beneficial to firms.

Extensions and additional studies can also be performed to gain more understanding and conclusive results on this topic. As previously noted, introducing more trading periods would give a broader understanding of market convergence and the long term effects of the treatment. Additionally, running this experiment as a double oral auction would shed some light on just how much money consumers are willing to give up to maintain their privacy. And finally, an interesting extension would be to survey subjects prior to the study and examine only those that demonstrated privacy concerns. This addition could provide more insight on how privacy concerned agents differ among themselves.

7. References

Acquisti, A., & Varian, H. R. (2005). Conditioning Prices on Purchase History. *Marketing Science, 24*(3), 367-381. doi:10.1287/mksc.1040.0103

Acquisti, A., John, L. K., & Loewenstein, G. (2013). What Is Privacy Worth? *The Journal of Legal Studies, 42*(2), 249-274. doi:10.1086/671754

Acquisti, A., Taylor, C. R., & Wagman, L. (2015). The Economics of Privacy. *SSRN Electronic Journal*. doi:10.2139/ssrn.2580411

Benndorf, V., & Normann, H. (2018). The Willingness to Sell Personal Data. *The Scandinavian Journal of Economics, 120*(4), 1260-1278. doi:10.1111/sjoe.12247

Brannon, J. I., & Gorman, M. F. (2002). The effects of information costs on search and convergence in experimental markets. *Journal of Economic Behavior & Organization, 47*(4), 375-390. doi:10.1016/s0167-2681(01)00204-9

Burke, J., Taylor, C. R., & Wagman, L. (2011). Information Acquisition in Competitive Markets: An Application to the US Mortgage Market. *SSRN Electronic Journal*. doi:10.2139/ssrn.1803322

Camerer, C. F., Hogarth, R. M., Budescu, D. V., & Eckel, C. (1999). The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework. *Elicitation of Preferences,* 7-48. doi:10.1007/978-94-017-1406-8_2

Cason, T. N., & Williams, A. W. (1990). Competitive equilibrium convergence in a posted-offer market with extreme earnings inequities. *Journal of Economic Behavior & Organization, 14*(3), 331-352. doi:10.1016/0167-2681(90)90062-i

Charness, G., Gneezy, U., & Halladay, B. (2016). Experimental methods: Pay one or pay all. *Journal of Economic Behavior & Organization, 131*, 141-150. doi:10.1016/j.jebo.2016.08.010

Conitzer, V., Wagman, L., & Taylor, C. R. (2011). Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases. *SSRN Electronic Journal*. doi:10.2139/ssrn.1944633

EU GDPR. (2018). GDPR Key Changes*. Key Changes with the General Data Protection Regulation – EUGDPR.* eugdpr.org/the-regulation/.

Evens, T., & Damme, K. V. (2016). Consumers' Willingness to Share Personal Data: Implications for Newspapers' Business Models. *International Journal on Media Management, 18*(1), 25-41. doi:10.1080/14241277.2016.1166429

Franciosi, R., Kujal, P., Michelitsch, R., Smith, V., & Deng, G. (1995). Fairness: Effect on Temporary and Equilibrium Prices in Posted-Offer Markets. *The Economic Journal, 105*(431), 938. doi:10.2307/2235160

Gellman, R. (2003). A Better Way to Approach Privacy Policy in the United States: Establish a Non-Regulatory Privacy Protection Board. *Hastings Law Journal, 54*(4), 1183-1226.

Grossklags, J., & Acquisti, A. (2007). When 25 Cents is Too Much: An Experiment on Willingness-To-Sell and Willingness-To-Protect Personal Information. *WEIS*.

Hann, I., Hui, K., Lee, S., & Png, I. (2007). Analyzing Online Information Privacy Concerns: An Information Processing Theory Approach. *2007 40th Annual Hawaii International Conference on System Sciences (HICSS'07)*. doi:10.1109/hicss.2007.81

Hermstrüwer, Y., & Dickert, S. (2017). Sharing is daring: An experiment on consent, chilling effects and a salient privacy nudge. *International Review of Law and Economics, 51*, 38-49. doi:10.1016/j.irle.2017.06.001

Hirshleifer, J. (1978). The Private and Social Value of Information and the Reward to Inventive Activity. *Uncertainty in Economics,* 541-556. doi:10.1016/b978-0-12-214850-7.50038-3

Hirshleifer, J. (1980). Privacy: Its Origin, Function, and Future. *The Journal of Legal Studies, 9*(4), 649-664. doi:10.1086/467659

Janrain Research. (2018). Consumer Attitudes Toward Data Privacy Survey 2018. Retrieved from https://www.janrain.com/resources/industry-research/consumer-attitudes-toward-data-privacy-survey-2018

Ketcham, J., Williams, A. W., & Smith, V. L. (1984). A Comparison of Posted-Offer and Double-Auction Pricing Institutions. *Papers in Experimental Economics,* 295-314. doi:10.1017/cbo9780511528354.020

Kujal, P. (1992). Asymmetric surplus distribution and the price convergence path in posted-offer markets. *Economics Letters, 39*(1), 33-36. doi:10.1016/0165-1765(92)90097-i

Kujal, P. (1994). Firm-specific output limits in a posted offer market: Distributive and efficiency effects. *Journal of Economic Behavior & Organization, 25*(2), 257-269. doi:10.1016/0167-2681(94)90013-2

Milne, G. R., Pettinico, G., Hajjat, F. M., & Markos, E. (2016). Information Sensitivity Typology: Mapping the Degree and Type of Risk Consumers Perceive in Personal Data Sharing. *Journal of Consumer Affairs, 51*(1), 133-161. doi:10.1111/joca.12111

Norman, G., Pepall, L., Richards, D., & Tan, L. (2016). Competition and consumer data: The good, the bad, and the ugly. *Research in Economics, 70*(4), 752-765. doi:10.1016/j.rie.2016.09.002

Odlyzko, A. (2003). Privacy, Economics, and Price Discrimination on the Internet. *SSRN Electronic Journal*. doi:10.2139/ssrn.429762

Palmatier, R. W., & Martin., K. D. (2019). The Psychology of Consumer Privacy. *The Intelligent Marketer's Guide to Data Privacy,* 21-41. doi:10.1007/978-3-030-03724-6_2

Plesch, J., & Wolff, I. (2018). Personal-Data Disclosure in a Field Experiment: Evidence on Explicit Prices, Political Attitudes, and Privacy Preferences. *Games, 9*(2), 24. doi:10.3390/g9020024

Plott, C. R., & Smith, V. L. (1978). An Experimental Examination of Two Exchange Institutions. *The Review of Economic Studies, 45*(1), 133-153. doi:10.2307/2297090

Politico & Morning Consult. (2017). Poll: Little Trust That Tech Giants Will Keep Personal Data Private. Retrieved from https://morningconsult.com/2017/04/10/poll-little-trust-tech-giants-will-keep-personal-data-private/

Posner, R. A. (1978). An Economic Theory of Privacy. *Philosophical Dimensions of Privacy,* 333-345. doi:10.1017/cbo9780511625138.016

Posner, R. A. (1981). An Economic Theory of Privacy. *Philosophical Dimensions of Privacy,* 333-345. doi:10.1017/cbo9780511625138.016

Posner, R. A. (1993). Blackmail, Privacy, and Freedom of Contact. *Economic Analysis of the Law,* 266-278. doi:10.1002/9780470752135.ch26

Poszewiecki, A. (2016). I Will Sell My Private Data: The Results of an Experimental Study on the Valuation of Privacy. *Selected Issues in Experimental Economics Springer Proceedings in Business and Economics,* 285-297. doi:10.1007/978-3-319-28419-4_18

Preibusch, S., Kübler, D., & Beresford, A. R. (2013). Price versus privacy: An experiment into the competitive advantage of collecting less personal information. *Electronic Commerce Research, 13*(4), 423-455. doi:10.1007/s10660-013-9130-3

Prince, C. (2018). Do consumers want to control their personal data? Empirical evidence. *International Journal of Human-Computer Studies, 110*, 21-32. doi:10.1016/j.ijhcs.2017.10.003

Rodríguez-Priego, N., Bavel, R. V., & Monteleone, S. (2016). The disconnection between privacy notices and information disclosure: An online experiment. *Economia Politica, 33*(3), 433-461. doi:10.1007/s40888-016-0040-4

Rosenboim, M., & Shavit, T. (2011). Whose money is it anyway? Using prepaid incentives in experimental economics to create a natural environment. *Experimental Economics, 15*(1), 145-157. doi:10.1007/s10683-011-9294-4

Ruffle, B. J. (2001). Some factors affecting demand withholding in posted-offer markets. *Studies in Economic Theory Advances in Experimental Markets,* 33-48. doi:10.1007/978-3-642-56448-2_3

Schmeiser, S. (2017). Online advertising networks and consumer perceptions of privacy. *Applied Economics Letters, 25*(11), 776-780. doi:10.1080/13504851.2017.1366634

Schudy, S., & Utikal, V. (2017). 'You must not know about me'—On the willingness to share personal data. *Journal of Economic Behavior & Organization, 141*, 1-13. doi:10.1016/j.jebo.2017.05.023

Stigler, G. J. (1980). An Introduction to Privacy in Economics and Politics. *The Journal of Legal Studies, 9*(4), 623-644. doi:10.1086/467657

Taylor, C. R., & Wagman, L. (2015). Consumer Privacy in Oligopolistic Markets: Winners, Losers, and Welfare. *SSRN Electronic Journal*. doi:10.2139/ssrn.2583453

Tsai, J. Y., Egelman, S., Cranor, L., & Acquisti, A. (2011). The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research, 22*(2), 254-268. doi:10.1287/isre.1090.0260

Varian, H. R. (2002). Economic Aspects of Personal Privacy. *Cyber Policy and Economics in an Internet Age,* 127-137. doi:10.1007/978-1-4757-3575-8_9

Zerlang, J. (2017). GDPR: A milestone in convergence for cyber-security and compliance. *Network Security, 2017*(6), 8-11. doi:10.1016/s1353-4858(17)30060-0