Impact of Security Events and Fraudulent Transactions on Customer Loyalty: A Field Study

Sriram Somanchi
Mendoza College of Business
University of Notre Dame
somanchi.1@nd.edu

Rahul Telang
H.J.Heinz III College
Carnegie Mellon University
rtelang@andrew.cmu.edu

Abstract

Security and Privacy has become a dominant issue for both consumers and corporations. In this paper, we investigate how customer behavior is affected after they have been a victim of financial fraud. Our analysis provides insights into how security concerns affect the continuation of the existing relationship of the customers depending on kind of fraudulent transactions. With the data from one of the largest banks in the US, we show that the probability of ending the relationship in the next six months increases significantly after a fraudulent transaction. Furthermore, these results are validated by an increase in rate of terminating the relationship after a fraudulent event. We provide results with a detailed analysis including the kind of fraudulent transaction, tenure and location.

Keywords: Security, Banking IS, Customer retention, Information Security and Privacy

1 Introduction

A spate of identity thefts, data breaches and relentless media coverage of these issues has brought security issues to the forefront of not only consumers but also policy makers. This is even more important for the financial industry. Users trust their banks to protect their financial assets and information. Unfortunately, hackers also realize that financial firms are an attractive target. Phishing attacks, social engineering, DDoS attacks are good examples of various ways the hackers try to breach firms' security. They use these techniques to gain access to users' credit card numbers, or bank account passwords and commit financial fraud. The main focus of this study is to understand the effect of an adverse security event on the relationship

between the customer and the financial organization. The adverse security events can be either identified by the bank (such as a data breach) or by the customer (such as fraudulent transactions). In this study, we concentrate mainly on the security events identified by the customer.

Policy makers have struggled to come to a consensus on how to protect users and reduce the number of breaches, frauds, identity thefts and so on. There is a belief that firms' incentives are not as well aligned with user incentives. After all, end users pay a significant price for insecurity and fraud. But passing stringent regulations are also opposed by industry, arguing that many regulations are ineffective and increase the cost of business. That said, a key goal of policy makers has been to increase the transparency of business practices. For example, the data breach notification laws (Wikipedia, 2015) passed by majority of states over the last decade force firms to notify users of potential data breaches. In fact, a large number of data breaches reported in the newspapers is possibly an outcome of these regulations where, firms were forced to disclose breaches. The goal of transparency and disclosure is to name and shame the firm, and also allow users to take appropriate steps in protecting themselves. Thus, transparency improves competition since users start paying attention to firm practices. Similarly, firms are wary of bad press and hence have incentives to protect customers (see (Romanosky et al., 2011) for details). Financial firms also face other regulations that ensures that they protect customer accounts.

But fundamentally, the key aspect of most of these regulations is that firms feel pressured to take appropriate actions. However, a lot depends on whether consumers pay attention to security frauds and if they are willing to hold the firms responsible. After a breach or an adverse event, many researchers have shown that firms' stock price suffer (Acquisti et al., 2006a). This suggests that firms feel pressure to behave more responsibly and hence invest more in security. But the stock price measures tend to be short term and usually these prices rebound. A long term effect is only possible if firms actually end up losing market share to competitors. So if firms with poor security practices, poor fraud prevention, and repeated breaches, if forced to disclose these practices and lose market share in the long run, then we can be sure that these regulations have teeth. But this requires that end users are willing to punish guilty firms by taking their business elsewhere and there is little evidence of that (see (Gaynor et al., 2012) who show that hospitals in more competitive markets are not likely to have fewer breaches). If anything, researchers argue that repeated disclosures of data breaches or notices of hacking makes consumer immune to these events and they are more likely to ignore them (Experian, 2014). A user's relationship with a firm depends on variety of attributes and security might be only a small part of it. So despite an adverse security event, the user still might not

change her relationship with the firm. First, she may not even notice it. Second, maybe the firm is so much better on other attributes (nicer building, lower prices, better service) that security deficit may not be enough to overcome other advantages. Finally, the users might (rightly) believe that other firms in the market have similar security deficit and hence might not be inclined to switch. In case of financial firms like a bank, this might be even more salient. In most cases, the banks compensate end users for all financial losses (some of this follows from specific regulations in the banking industry) alleviating the effect of a fraud significantly.

Empirically, there is little research on how users respond to such adverse events. This is despite the fact the both policy makers and firms make certain assumptions about user behavior when designing policies or firm specific strategies. The disclosure and transparency policies implicitly assume that consumers are paying attention and are willing to punish tardy firms. The firms may install their own customer service programs assuming that these programs may alleviate customer relationship. The issue of addressing a adverse security event is even more important as the financial firms are also increasingly rolling out Internet and mobile based access to customer accounts, mobile check deposits, money transfer and so on. While these services are attractive to end users, they also come with significant security risks. Again, how important these risks are to the end users is a big unknown. Large number of customer related frauds can impede the adoption of these services.

A big challenge to empirically examine this issue is the lack of data. It is very hard to gather a reliable, long time-series data where one can observe potential adverse events and subsequent user reaction. In this paper, we are able to assemble a unique dataset of more than $500,000$ users for a period of more than 5 years. Some of these customers (close to $20,000$) encounter unauthorized fraudulent transactions on their account. These transactions are "unauthorized" - i.e. they were not recognized (and allowed) by the end users and end users specifically complained to the bank. These transactions typically occur due to user account being misused by someone, or someone stealing their credit card or debit card or ATM card numbers and using them for fraudulent charges. Or, it could even occur due to victims falling prey to social engineering or phishing scams. In most instances, the bank compensated the end user for potential losses.

This data provides us a rare opportunity to examine how customers responded to these events. It was evident in the data that customers were aware of these events (they informed the bank), that they incurred a financial loss and other potential non-monetary loss. Moreover, the person (or entity) who caused the frauds remains unknown to end users, even the bank is unable to trace the frauds to a particular entity. If these events cause customers to lose trust in the bank's security practices, then users might have an incentive to

terminate the relationship. In our data, we are able to observe this churn behavior. We also have detailed customer specific data (demographic as well as transactional) as well as market structure data (for example the bank's market share in a given market). With this data in hand, our analysis shows that these adverse events cause users to terminate the relationship with the bank. In particular, a user is 3 percentage point more likely to terminate relationship with the bank within six months of such an event. We also find that users are more likely to churn when they have higher tenure or lower age groups.

As we noted, the unauthorized charges cannot be tracked back to the perpetrators. We hypothesize that this lack of attribution is a significant source of uncertainty for end users, potentially leading to diminished trust with the bank. In our data, unauthorized transactions are not traced back to a particular individual or merchant or an entity. So the end user does not know who was responsible for that transaction, whether the matter is resolved, or whether it can occur again. We argue that this uncertainty plays an important role in reducing customer trust. To test our hypothesis, we also collect data on transactions which, after investigation by the bank, were attributed to another third party. For example, some of these transactions occurred due to merchant error and user was eventually compensated by the merchant. Or, some of the transactions were attributed to other members of the household. We show that indeed, when the attribution is clear, the effect of fraudulent transactions is much smaller.

Our paper is organized as follows. We briefly summarize the literature in this space, state our key hypotheses and then present our analysis. Further, we describe our data and present our results and discuss the implications. In the last section we conclude and outline the limitations of our study.

2 Literature review

Literature in this space can be divided into three major areas. In the first domain, studies have generally focused on the impact of data security, breaches and role of associated regulations. In the second domain, the studies focus on detection of security incidences, breaches and potential fraud. Finally, in the last domain, there is some work on how adverse events or frauds affect customer loyalty.

Much work on data breaches and related regulations look at the potential costs of data breaches to firms and end users (Institute, 2015), (Verizon, 2015). The numbers vary across studies, but it is widely believed that data breaches cost end users significantly. Data breaches or security incidences also cost firms significantly. Many event studies have examined the effect of data breaches on firm stock prices (Gatziaff

and McCullough, 2010). Most study find a negative effect of breaches on firm stock price. Similar results were reported by (Acquisti et al., 2006b) for privacy data breach events and by (Telang and Wattal, 2007) for software flaw disclosures. However, these studies do not provide insights into the long term effects of insecurity.

A second line of research is on various security regulations themselves. The most widely discussed law is the data breach notification law, which was first adopted in California (SB 1386) in 2003. Since then 47 states have adopted this law. Much of the work though, focuses on the law itself, from a legal perspective (Schwartz and Janger, 2007). The purpose of the law is very clear - to impose reputational consequences on firms so that they invest in protecting customer data. So the idea is similar to the event studies, in that disclosure of an unpleasant event has reputational consequences for firms leading to stock price decline. Telang (Telang, 2015) provides a useful summary of the purpose of such regulations and potential impact on firm and user behavior. There is also discussion on whether breach notification laws should become a federal law instead of statewide laws (Tom, 2010). However, measuring the impact of these (and similar) laws on user outcome is not well studied. (Romanosky et al., 2010) show that these laws led to marginally fewer instances of identity thefts. But the data was aggregated at the state level.

The third stream of literature focuses on the impact of frauds committed by third parties on financial firms. As we noted, security is increasingly becoming an important issue in then banking industry and the number of fraudulent transactions have also increased rapidly. Naturally, firms are spending significant resources in fraud detection. There is a significant literature in Computer Science on how to detect anomalies and frauds. One would expect that frauds would impact customers' perception of feeling secure and protected and hence might negatively affect customers' relationship with the firm. This, in turn, may lead to customer churn. However, the actual evidence on this hypothesis is rather limited. (Hoffmann and Birnbrich, 2012) examine the effect of fraud prevention on the bank-customer relationship. But their data is limited to surveys. (Suh and Han, 2003), again using surveys, show that perception of security control plays an important role in customer trust in Internet Banking. A study by (Ablon et al., 2016) is the only one we are aware of that looks at users' attitudes towards data breaches. But the study is survey based where users intentions are measured rather than actual actions.

Based on our survey of literature, we find very little, if any, work that assembles the data we have. We observe real consumer behavior before and after they encounter a fraudulent event (committed by a third party) on their bank account. We also observe users who are not exposed to such events as a comparison

group. We believe our research documents first such evidence of the impact of security frauds on user loyalty. As we will highlight, we not only bring a novel dataset to the table, we also provide a nuanced story of how security violations and frauds frame user response.

3 Consumer behavior and hypothesis

As we noted in the introduction and literature review, much of the work is survey based where users are asked about their security perception and corresponding trust with the firm they are transacting with. An important theme that emerges is that a fraudulent event imposes considerable time and emotional cost to the end users. They usually have to inform the bank of potential losses and might have to take additional actions to mitigate possible losses (identity theft or other similar changes showing up in the future on other accounts). This is costly, both monetarily and non-monetarily (say emotionally). It is reasonable to conclude that being a fraud victim would affect the customers' perception of bank's security practices. This is despite the fact the bank, almost always, compensates the end users for their losses. Thus, the user may lose her trust in the bank and damage her relationship with the bank. This may negatively affect customer loyalty and encourage switching behavior. So we hypothesize that -

Hypothesis 1 (H1): Users are more likely to churn when encountered with fraudulent events.

However, we also argue that much of this loss in trust is an outcome of uncertainties. When the losses are not traced back to the perpetrators, and not clearly attributed for, and the reasons are not clearly explained to the users, they are more likely to worry about similar attacks occurring in the future. They may also question the bank's ability to trace and deter future attacks. They may also worry about spillovers. That, these losses may signal more potential harm in future (say identity thefts). Thus, lack of clarity on who perpetrated these acts would lead to users placing blame on the bank, even though the fraud is actually committed by a third party and the bank has compensated the user. On the other hand, if the attribution of these fraudulent events is clear and reasons clearly explained, the user may rationalize the event happening despite the banks' best effort. Thus, we hypothesize that -

Hypothesis 2 (H2): Users are more likely to churn when the fraudulent events are not clearly attributed.

To test $H2$ specifically, we also collect data on fraud events which were finally tracked back to a specific third party (a merchant, or another member of the user household, or even user realizing that it was her fault).

4 Data and Analysis

This work is in collaboration with one of the leading banks in the US, who provided us access to an anonymized detailed transaction level data about its customers. We have access to approximately 500 thousand customers which is full geographic stratified sample of the United States with higher concentrated sampling in a couple of cities that the bank was specifically interested. For each of the anonymized customer, we have some demographic information like age and home zip code location. For each of the account of a given customer, we have details on all of the account types, daily branch transactions, including debits and credits, debit card and credit card transactions, mobile transactions, transaction amounts, and information on the relationship the customer maintains with the bank. We have up to 5 years of historical information on each customer from last quarter of 2008 to third quarter of 2013. We also have information on customer care call records, with the date and reason for the call, along with the resolution of the issue and the resolution date. Specifically, in this work we concentrate on the customers who had a fraudulent transaction on their accounts and called the bank to report the issue on their account. In this paper, we provide analysis of data aggregated quarterly, which is important for the bank to validate the changes happening at this granularity.

In the current data, the users, after identifying potential charges on their account (all related to their bank account, and not the credit cards), calls the bank's customer care center for reporting the fraud. Bank investigates these complaints, resolves it, and records it in its call history database. So we know exactly when the complaint came in, and what action was taken by the bank. Our key goal is to examine the cases which were most likely driven by some potential security lapse (we cannot identify whether the error happened at the user end or at the bank end). All events classified as unauthorized transactions fit this definition well. In all these cases, the user account incurred a charge which was investigated to be unauthorized by the user. It potentially was perpetrated by a third party who got hold of the user account or the password or misuse of the debit card associated with the account. During our discussion with the bank executives, we learned that, in all of the cases of unauthorized transaction, the bank issues the credit back to the user within ten days.

We also collected data on complaints which were resolved to well identified third parties. For example, some of the unauthorized transactions were linked to a specific merchant and the merchant issued the credit back to the end user. In another category, the fraudulent transactions were found to be valid by the customer at a later date, and the customer withdrew the complaint. Finally, in the last category, valid client charge, the bank determined that the transaction was not a fraudulent and was in fact a valid one. The bank might

have come to this conclusion based on the investigation done on historical information of the customer. The key difference here is that the bank did not pay the customer for the transaction amount in dispute.

Notice that in the case of merchant issued credit, the attribution is clear. The user can call the merchant for the compensation. So, we would expect that for the merchant related fraud, customers may not hold the bank responsible for those charges. This should not affect the trustworthiness of the banks' security efforts and should have smaller impact on bank-customer relationship. Similarly, when a user withdrew the claim, we do not expect it to adversely affect banking relation.

In case of unauthorized transactions, since the matter remains unresolved, it is likely that users may hold a bank indirectly responsible and it may adversely affect their relationship. Finally, the valid client charge is a special case since the claim was denied by the bank and the customer did not get any compensation. So it might lead to customer grievance, independent of whether the fraud happened or not (according to the bank the fraud did not happen).

Our measure of the customer's relationship with the bank is whether the user terminates the relationship with the bank. The closure of the relationship with the bank means the customer closed all accounts and completely terminated her relationship.

4.1 Analysis: Sample Selection

In order to estimate the effect of fraudulent transaction on closure of accounts we need to know the typical closure behavior for the accounts with non-fraudulent transaction. While we can use the whole population, we are worried about the selection. In particular, we are worried that users who have frauds perpetrated on their account might be different from non-fraud accounts (say are more risky, or switch banks frequently). To overcome this selection, we first select the non-fraudulent accounts by using Propensity Score Matching (PSM) method, with the treatment being the occurrence of a fraudulent transaction. This method allows us to find similar users in both fraud and non-fraud groups. We match based on various characteristics like the tenure of the account with the bank, average balance, average number of transactions, average online behavior - before the fraudulent event took place. We matched three control users for every treatment user.

We produce some of the summary statistics of each group of customers in the Table 1, we can see that most of the characteristics of the group of customers are very similar except for probability of closing in the next six months. The probability of ending the relationship in the next six months for the matched customers, is the probability of closing from the corresponding treated customers' time of experiencing a

Table 1: Summary Statistics comparing each group of customers in our analysis

Variable	Unauthorized	Merchant Credit	Valid Client	Client Withdrew	Matched Customers
Number of Accounts	20104	3034	1962	1210	59764
Age	41.60	44.12	42.88	47.55	43.48
Tenure	55.36	69.97	57.39	86.43	52.31
Avg Balance (in dollars)	2750.61	3123.12	2821.45	3612.29	2921.03
ATM Transactions (avg per month)	4.28	3.78	4.79	6.12	3.69
Number of customers ending relationship	1337	173	257	67	2290
Probability of ending the relationship	0.067	0.057	0.131	0.055	0.038

fraudulent event.

4.2 Effect of fraudulent transaction compared to the matched customers

We want to examine whether the occurrence of any of the fraudulent events, affects customers decision to terminate the relationship with the bank compared to similar matched customers who have no such fraudulent events. We propose the following specification for our first logit regression

$$Prob(Q_i) = \beta_0 + \beta_1 * fraudEvent_i + \beta_2 * X_i + \beta_3 * year_i + \varepsilon, \tag{1}$$

where, $Q_i = 1$ if the customer i ends the relationship with the bank in the next six months of the incidence of an event, $fraudEvent_i$ represents the fraudulent event type that the customer experiences. We also include some of the characteristics X_i of the customer and customer's relationship with the bank, like age, tenure of the account in months, balance and the number of transactions, including yearly effects. Further, we include a dummy if the customer's zip code belongs to Allegheny County, where the bank has huge prevalence. The results of the above *logit* function is shown in the Table 2

From the Table 2 we see that fraud event types: unauthorized transaction, merchant credit and valid client charge, increase the likelihood of closing the relationship with the bank compared to the similar matched

Table 2: Logit results for ending the relationship with the bank based on each fraudulent event type compared to the matched customers

Variable	Coefficient Estimate	p-value
Unauthorized	0.6372	$<$ 2e-16***
Merchant Credit	0.5061	0.03494*
Valid Client Charge	1.1570	1.35e-06***
Client Withdrew	0.6266	0.14194
In Allegheny	-0.3486	1.05e-14***
Age	-0.006	1.78e-06***
Tenure	-0.0155	$<$ 2e-16***
Balance	-0.00009	0.00505**
ATM Transactions	-0.0037	0.18727
Year 2009	0.261	0.00328**
Year 2010	0.025	0.70877
Year 2011	-0.170	0.00259**
Year 2012	-0.0151	0.76737

customers. However, there is no significant effect when the client withdrew the claim. More specifically, when compared to the matched users who did not incur fraudulent event, unauthorized transaction increased the likelihood of closing by 3.2 percentage points. This is consistent with our hypothesis in *H1*.

We also find that the merchant credit and valid client charge increased the likelihood of churn by 2.3 percentage points and 7.3 percentage points respectively. Though the merchant credit dummy is only marginally significant, it still seems to suggest that despite clear attribution (to the merchants), users may still terminate their relationship with the bank. So *H2* is only partially supported. In the following sections, we provide proportional hazard model as well as another method to control for potential selection in our comparison groups. Finally, "valid client charge" has a very large effect on ending the relationship, possibly due to the customer grievance that the request for fraudulent transaction was denied by the bank and she is not compensated for the transaction amount.

In terms of the location characteristics, we find that a customer being in Allegheny County, where the bank has a larger market share, the effect is significantly smaller. So users are less likely to churn, in general, when the bank is a dominant player in that market.

4.2.1 Post-hoc analysis to identify customers who are sensitive to fraud

We wanted to examine and understand the characteristics of a customer that were predictive of a decision to terminate the relationship after the fraudulent effect. Hence, we also examine whether the effect of the

Table 3: Logit results for ending the relationship with the bank based on each fraudulent event type compared to the matched customers, with interaction of tenure

Variable	Coefficient Estimate	p-value
Tenure	-0.0349	$<$ 2e-16***
Tenure Unauthorized	0.0268	$<$ 2e-16***
Tenure Merchant	0.0349	$<$ 2e-16***
Tenure Valid Client	0.0290	7.07e-15***
Unauthorized	-0.1692	0.0048**
Merchant	-0.869	0.0115**
Valid Client	0.283	0.332
Client Withdrew	0.3002	0.5878
Age	-0.0055	1.70e-05**
In Allegheny	-0.349	9.71e-15
Balance	-8.754e-06	0.01073*
ATM Transactions	-4.852e-04	0.09616

occurrence of any of these fraudulent events on the customer's decision to terminate the relationship with the bank varies with tenure. This helps us understand if there is stickiness to the customers who have long relationship with the bank. So we have the interacted the tenure with the fraudulent event and the following is the specification,

$$Prob(Q_i) = \beta_0 + \beta_1 * fraudEvent_i + \beta_2 * fraudEvent_i * Tenure_i + \beta_3 * X_i + \varepsilon, \tag{2}$$

where again the X_i are observed characteristics of the the customer i. Again, we compare fraud events with the similar matched customer group who did not experience any such event. The results of this logit are shown in the Table 3

From the Table 3 we observe that users with higher tenure, in general, are less likely to quit. However, the interaction with tenure for all four events is positive. This suggests that higher tenured users are more likely to quit when encountered with a fraudulent event. This was unusual as one would expect that tenure would build more loyalty leading to less churn. But this effect does not sustain when a fraudulent event occurs. Again, as expected, there is no significant change for the group of customers who withdrew the claim.

Further, even in this case, in terms of the location characteristics, we find that a customer being in Allegheny County, where the bank has larger market power, the churn rate is lower.

Table 4: Results of a proportional hazard model for understanding the rate of ending the relationship with the bank based on each fraudulent event type compared to the matched customers

Variable	Coefficient Estimate	p-value
Unauthorized	1.28	$<$ 2e-16
Merchant Credit	-0.137	0.9854
Valid Client Charge	1.74	2e-07
Client Withdrew	0.199	0.8556
In Allegheny	0.506	1e-08
Age	-0.0101	0.0012
Tenure	-0.00916	$<$ 2e-16
Balance	-0.000018	0.8364
ATM Transactions	0.00368	0.0025
Year 2009	0.294	0.1160
Year 2010	0.0986	0.5571
Year 2011	0.124	0.4420
Year 2012	0.187	0.2599

4.2.2 Proportional Hazard Model

Instead of just modeling the customers who quit within the next six months, we also want to verify the effect fraudulent events on the rate of terminating the relationship with the bank. Hence, we performed the Cox proportional hazard model (Cox, 1972), with the following model specification:

$$\lambda(t) = \lambda_0(t)exp(\beta_0 + \beta_1 * fraudEvent_i + \beta_2 * X_i + \beta_3 * year_i), \tag{3}$$

where the time is censored for customers who did not quite even at the end our time window of analysis. Again, we include all possible types of fraudulent events, as well as customer and account characteristics like age, tenure, balance, location, etc. We are comparing the quite rate for each of the four event types compared to the matched group of customers who never encountered any fraudulent events.

From the Table 4 we can see that fraud event types: unauthorized transaction, valid client charge, increase the rate of terminating the relationship with the bank compared to the similar matched customers. However, there is no significant effect when the customer withdrew the claim. This is consistent with our Hypothesis *H1*.

Further, we see that when the attribution is clear (to the merchants), users have no significant effect to the rate of termination of the relationship. So, *H2* is supported. Again, "valid client charge" has a very large effect in increasing the rate of ending the relationship.

4.3 Comparing customers with their future counterparts

We analyzed the closure of the relationship in the next six months by comparing similar matched customers with the customers who experience one of the four fraudulent events. The matched customers are currently chosen by propensity score matching techniques as described above. However, it might be that even after matching, the matched group might be different from the groups that encountered fraud. This may under (or over) estimate the true effect of the adverse events.

As a robustness check, we now propose the following method. We compare the quit rate of the customers within same cohort. So we compare the users who experienced fraud and quit at time T, with the users who experienced fraud, but did not quit at time T but continued on and might quit later. The idea here is that the future quit rate provides a baseline quit rate for users (all of whom experience a particular type of fraud). While the "within six month" quit rate provide the effect of fraud. By comparing users within a cohort (all were defrauded), we are comparing more homogeneous users, reducing the possibility of selection (when we compare defrauded users to similar matched users). Notice that we expect the effect to be smaller now since we are being very conservative in estimating the effect. We are comparing users who are quitting with those who chose not to quit immediately. Again, in this case we run a *logit* regression with the following specification except that we are not using the matched users in our sample.

$$Prob(Q_i) = \beta_0 + \beta_1 * fraudEvent_i + \beta_2 * X_i + \beta_3 * year_i + \varepsilon, \tag{4}$$

where, again, X_i describes the customer characteristics, like age, tenure, balance and ATM transactions, along with yearly effects. The results of the logit specification is shown in Table 5.

From the Table 5 we see that only unauthorized and valid client charge have a significant effect on the quit rate. More specifically, when compared to their future counterparts, unauthorized transaction leads to 1 percentage point more likely to end the relationship within six months, and valid client charge leads to 2.7 percentage points more likely to end the relationship. For the group of customers for whom the merchant issued credit or if the client withdrew the claim there is no significant effect in quit rates. This confirms, our earlier hypothesis *H2* that when the attribution of the loss is clear, users might not hold the bank responsible.

Table 5: Logit regression of comparing customers who quit immediately with customers who chose to stay on and might quit later

Variable	Coefficient Estimate	p-value
Unauthorized	0.257	$<$ 2e-16***
Merchant Credit	0.0938	0.692518
Valid Client Charge	0.566	0.003731**
Client Withdrew	0.133	0.734258
In Allegheny	-0.2975	$<$ 2e-16***
Age	-0.0017	0.056834
Tenure	-0.008	$<$ 2e-16***
Balance	-9.243e-05	$<$ 2e-16**
ATM Transactions	-1.042e-04	0.131335
Year 2009	-0.005	0.941662
Year 2010	-0.167	0.000958***
Year 2011	0.298	$<$ 2e-16***
Year 2012	0.470	$<$ 2e-16***

5 Conclusion

Using a large sample of users, from a financial firm, for a period of time, we examine how fraudulent financial charges affect users' relationship with the firm. These frauds were mostly an outcome of users' financial information being stolen and misused. It is also noteworthy that all customer loss was compensated back to the user by the bank. Our results suggest that even when the bank is not directly responsible for a fraudulent transaction, users may hold the bank responsible and terminate their relationship. For unauthorized transaction, when the attribution of perpetrators is unclear, this effect is large (from 1 to 3 percentage point). Further, this effect is much larger when users are not compensated because the bank determined that the charges might be legitimate. These effects are bigger for users with larger tenure. One of the implications of our research for the banks is, to consider following up with the customers who experience a fraudulent event and make sure the relationship is not adversely affected.

We are currently continuing our research to understand the effect of security event on the time to close the relationship with the bank through survival analysis (Fichman and Kemerer, 1999), (Morita et al., 1993). Furthermore, we have not calculated the lifetime value of these users yet but it is clear that these frauds carry both direct and indirect costs to banks, in particular. The banks incur a direct cost of identifying frauds, invest in customer service, and in compensating users. On top, there is an indirect cost of potentially losing the customers. Our research highlights that users are, when being aware of the fraud, do take expected actions. That is, they are willing to punish the firm leading to possibly larger security investments by the

firm. Our research seems to confirm the efficacy of the some of the regulations whose goal is to highlight firms' security and data protection practices.

Acknowledgement

Rahul Telang acknowledges generous support of NSF Eager grant - 1359632.

References

Ablon, L., Heaton, P., Lavery, D. C., and Romanosky, S. (2016). *Consumer Attitudes Toward Data Breach Notifications and Loss of Personal Information*. RAND Corporation.

Acquisti, A., Friedman, A., and Telang, R. (2006a). Is there a cost to privacy breaches? an event study. In *Proceedings of the International Conference of Information Systems (ICIS)*, Milwaukee.

Acquisti, A., Friedman, A., and Telang, R. (2006b). Understanding the impact of privacy breaches. In *35th Research Conference on Communication, Information and Internet Policy*, George Mason University School of Law, Arlington, Virginia.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(2):187–220.

Experian (2014). Aftermath of a mega data breach:consumer sentiment.

Fichman, R. G. and Kemerer, C. F. (1999). The illusory diffusion of innovation: An examination of assimilation gaps. *Information Systems Research*, 10(3):255–275.

Gatziaff, K. and McCullough, K. A. (2010). The effect of data breaches on shareholder wealth. *Risk Management and Insurance Review*, 13(1):61–83.

Gaynor, M. S., Hydari, M. Z., and Telang, R. (2012). Is patient data better protected in competitive healthcare markets? In *Workshop on the Economics of Information Security (WEIS)*, Berlin.

Hoffmann, A. O. and Birnbrich, C. (2012). The impact of fraud prevention on bank-customer relationships: An empirical investigation in retail banking. *International Journal of Bank Marketing*, 30(5):390–407.

Institute, P. (2015). Ponemon institute annual survey.

Morita, J. G., Lee, T. W., and Mowday, R. T. (1993). The regression-analog to survival analysis: A selected application to turnover research. *Academy of Management Journal*, 36(6):1430–1464.

Romanosky, S., Telang, R., and Acquisti, A. (2010). Do data breach disclosure laws reduce identity theft? In *Conference on Empirical Legal Studies (CELS)*, Yale Law School, New Haven, Connecticut.

Romanosky, S., Telang, R., and Acquisti, A. (2011). Do data breach disclosure laws reduce identity theft? *Journal of Policy Analysis and Management*, 30(2):256–286.

Schwartz, P. M. and Janger, E. J. (2007). Notification of data security breaches. *Michigan Law Review*, 105(913).

Suh, B. and Han, I. (2003). The impact of customer trust and perception of security control on the acceptance of electronic commerce. *International Journal of Electronic Commerce*, 7(3):135–161.

Telang, R. (2015). Policy infrastructure for data breaches. *IEEE Security and Privacy*, 10(3):783–798.

Telang, R. and Wattal, S. (2007). An empirical analysis of the impact of software vulnerability announcements on firm stock price. *IEEE Transactions on Software Engineering*, 33(8):544–557.

Tom, J. M. (2010). A simple compromize: The need for a federal data breach notification law. *St John's Law Review*, Issue 4.

Verizon (2015). Verizon data breach investigation report.

Wikipedia (2015). Security breach notification laws.