Should Credit Card Issuers Reissue Cards in Response to a Data Breach?: Uncertainty and Transparency in Metrics for Data Security Policymaking

JAMES T. GRAVES, Georgetown University
ALESSANDRO ACQUISTI and NICOLAS CHRISTIN, Carnegie Mellon University

When card data is exposed in a data breach but has not yet been used to attempt fraud, the overall social costs of that breach depend on whether the financial institutions that issued those cards immediately cancel them and issue new cards or instead wait until fraud is attempted. This article empirically investigates the social costs and benefits of those options. We use a parameterized model and Monte Carlo simulation to compare the cost of reissuing cards to the total expected cost of fraud if cards are not reissued. The ranges and distributions in our model are informed by publicly available information, from which we extrapolate estimates of the number of credit card records historically exposed in data breaches, the probability that a card exposed in a breach will be used for fraud, and the associated expected cost of existing-account credit card fraud. We find that automatically reissuing cards may have lower social costs than the costs of waiting until fraud is attempted, although the range of results is considerably broad.

CCS Concepts: • **Security and privacy → Economics of security and privacy**; • **Applied computing** → *Law*; • **Social and professional topics** → *Identity theft*; *Financial crime*;

Additional Key Words and Phrases: Economics of information security, data breach, estimation, identity theft, Monte Carlo

ACM Reference format:
James T. Graves, Alessandro Acquisti, and Nicolas Christin. 2018. Should Credit Card Issuers Reissue Cards in Response to a Data Breach?: Uncertainty and Transparency in Metrics for Data Security Policymaking. *ACM Trans. Internet Technol.* 18, 4, Article 54 (September 2018), 19 pages.
https://doi.org/10.1145/3122983

This work was partially funded by the Department of Homeland Security Science and Technology Directorate, Cyber Security Division, Broad Agency Announcement 11.02; the Government of Australia; and SPAWAR Systems Center Pacific, via contract number N66001-13-C-0131. Portions of this work were also supported by NSF IGERT grant DGE-0903659. In addition, Acquisti gratefully acknowledges support from the Carnegie Corporation of New York via an Andrew Carnegie Fellowship. For a list of Acquisti's additional grants and funding sources, please visit www.heinz.cmu.edu/~acquisti/cv.htm. This work represents the position of the authors and not that of the aforementioned agencies.
Authors' addresses: J. T. Graves, Department of Engineering and Public Policy, 5000 Forbes Ave., Pittsburgh, PA 15213; email: jtg@cmu.edu; A. Acquisti, Heinz College, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213; email: acquisti@andrew.cmu.edu; N. Christin, School of Computer Science and Department of Engineering and Public Policy, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213; email: nicolasc@andrew.cmu.edu.
Current address: J. T. Graves, Institute for Public Representation, Georgetown University Law Center, 600 New Jersey Ave NW, Washington, DC 20001.

https://doi.org/10.1145/3122983

1 INTRODUCTION

In recent years, there has been growing interest in the economic analysis of information security problems—including, in particular, the study of data breaches. Scholars have investigated the impact of data breaches on the stock market valuations of breached firms, the relationship between security investments and the frequency of breaches, and the role of market competition in predicting the probability of a breach. Less explored, however, has been the issue of how private choices by credit card issuers affect the public costs of a breach. After a breach of credit cards is disclosed, the financial institutions that issued those cards can either immediately cancel and reissue those cards or instead wait until someone attempts to use the card data for fraud. Reissuing cards can be expensive and potentially wasteful because many cards impacted in a breach may never be used for fraud. But not reissuing cards increases the risk of credit card fraud, which incurs costs to issuers, merchants, and cardholders. No fraud-monitoring program can prevent all fraud. Although issuers may evaluate the internal risks and benefits of reissuing, to our knowledge no published study has attempted to measure the overall societal benefits of each option when costs external to the issuers are considered.

In this article, we empirically investigate the social (i.e., aggregate) costs and benefits of reissuing breached cards immediately versus waiting until card fraud is attempted. We analyze "first order" costs: those costs that are proximate results of reissuing cards or leaving them in circulation despite possible compromise. Our analysis focuses on societal costs and benefits rather than costs and benefits to issuers.

Although the costs and sources of identity theft are well researched, the connection between identity theft and data breach is not as well understood, nor is quality data available on data breach or its resulting harms. Our analysis therefore estimates, based on publicly available data sources of varying quality, the number of credit cards exposed in data breaches, the cost of identity theft, and the extent to which identity theft is traceable to breaches of credit card data. We analyze public information about reported credit card breaches with known record counts to extrapolate an estimate of unknown records that would also have been exposed. We address uncertainty through parameterization, Monte Carlo analysis, and sensitivity analysis.

This article makes two contributions to the literature. First, it confirms that the first-order costs of automatically reissuing cards may be lower than waiting until fraud is attempted. Second, it illustrates where improved access to quality data sources is most needed. Our results are limited by reliance on publicly available information about data breach and identity theft. Some of this information is excellent, but much of it is not. The extent to which our model is sensitive to different data sources may serve as a guide for where resources could most usefully be spent to improve understanding of the causes of data breach.

Despite these limitations, our result is fairly robust to the tremendous uncertainty in our model. Although the range estimation results in a two-order-of-magnitude difference in the estimated cost of fraud if cards are not reissued, the Monte Carlo analysis shows roughly a 91% probability that societal losses would be lower if cards are reissued.

Part 1 presents background information placing our research in the context of previous work studying the economics of data breach and cybercrime. Part 2 describes our methodology and model. Part 3 explains the data we used for the parameters of our model. Part 4 presents the analysis of the data. Part 5 discusses the implications and some of the limitations of our research.

2 BACKGROUND

Credit card payments rely on relationships between five parties: cardholders, merchants, issuing banks, acquiring banks, and card associations (Levitin 2010). Figure 1 illustrates this structure. An acquiring bank (or "acquirer") is the merchant's bank; the issuing bank (or "issuer") is the bank

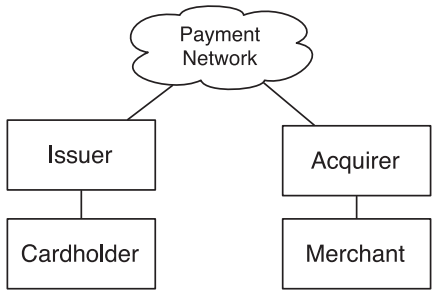

Fig. 1. Credit card payment network structure.

with whom the cardholder has a revolving credit account. The card associations (e.g., MasterCard, Visa, American Express, or Discover) are networks of financial institutions that set rules governing transactions. In the case of American Express and Discover, the card network and issuer are usually the same.

In simplified form, a credit card transaction works as follows. When a cardholder presents a card for payment at a merchant, the merchant passes the card information and authorization request to its acquiring bank, which forwards the request to the cardholder's issuing bank. The issuer authorizes or rejects the transaction. If the transaction is authorized, the issuer transfers funds from its payment network account to the acquirer's payment network account. This is called "capture." Finally, the transaction is "settled" when the acquirer credits the merchant's account.

In the United States, issuers bear the initial risk of loss from credit card fraud from card-present transactions, but the contractual relationships between issuers, the card brands, merchants, and the merchants' acquiring banks allow those losses to be shifted to merchants that have violated the card brand operating regulations by not following prescribed security measures. In most states, however, loss shifting is available only for fraudulent charges. Issuers bear all the operational costs of reissuing cards and have had little success in lawsuits to recoup these costs from breached merchants. But issuers who sue cannot recover damages they could have avoided. If an issuer could have reduced fraudulent charges to an exposed card by canceling and reissuing that card but did not, the issuer may not be able to recover the cost of those charges if they could have been avoided. Conversely, if the total amount of fraudulent charges that result from a breach are lower than the cost of reissuing the cards, reissuing would fail to mitigate damages.

In at least one case, merchants have used the fact that an issuer reissued cards and lacked fraud-monitoring processes to claim that issuers did not mitigate damages. In the consolidated putative class-action lawsuit resulting from the breach at TJX, one of the retailer's defenses was that by "unnecessarily and unreasonably automatically canceling and reissuing their customers' debit cards in response to the data compromise" and by not using fraud monitoring, some of the plaintiffs had either failed to mitigate damages or were contributorily negligent.[1]

The cost of reissuing cards is not the only incentive affecting an issuer's decision whether to reissue. Maintaining cardholder loyalty may be an even more important incentive for issuers to reissue cards even when the cost of doing so might be greater than anticipated fraud. And evidence suggests that issuers do often reissue cards after a breach even if they will not be able to recover the costs of doing so. But the tension between the losses issuers can recover, the operational costs that issuers generally cannot recover, and the obligation to minimize losses raises legal and policy questions. Should the law recognize reissuing costs as reimbursable losses? Is it more societally

[1]Answer to Plaintiffs' Consolidated Class Action Complaint at 22, *In re* TJX Retail Security Breach Litigation, 527 F.Supp.2d 209 (D. Mass. 2007) (No. 07-10162-WGY), 2007 WL 5324216.

beneficial to immediately reissue cards or wait? And, more importantly, do we even have the data needed to answer that question or many other public policy questions involving tradeoffs of data security choices?

Our work tries to answer those questions by building on the literature on the economics of information security, particularly that concerned with data breach. This literature seeks to understand the scope of data breaches, their cost, and the effectiveness of interventions to reduce their impact. Data about breaches has also been used to analyze the economics of security investments more generally.

The full extent of data breaches is difficult to measure. There is currently no comprehensive, openly accessible database of data breaches. The Privacy Rights Clearinghouse (PRC) (2016a), whose Chronology of Data Breaches is one of the most commonly used datasets for breach analysis, warns that its chronology is not a complete list of all breaches. The Identity Theft Resource Center (2016) publishes annual data breach reports but only makes the most recent year's list of breaches available online. The Open Security Foundation (2016) was one of the first to create a database of breaches, but its data, which was once free to download, is no longer available to the general public. A few states publish lists of the data breaches reported to their attorneys general or other authorities, but these lists include only breaches that affect residents of those states. Three states—Maine, Maryland, and New Hampshire—include estimates of the number of their states' residents who were affected by each breach (if reported by the organization that was breached). The data from these states might be analyzed in conjunction with the PRC database to obtain a more complete picture of the extent of data breaches.

The distribution of data breaches is heavy-tailed: a few extremely large breaches of millions of records have gotten lots of attention, but most breaches are much smaller. One statistical model predicts, for example, a 31% chance per year of a breach of 10 million records or more in the United States (Edwards et al. 2015) .

Early efforts to measure the cost of data breaches were based on surveys. Although they suffer from numerous problems, surveys continue to be popular among industry analysts (Heiser 2002; Ryan and Jefferson 2003; Shostack and Stewart 2008, pp. 46–49). One of the first surveys was the Computer Security Institute's Computer Crime and Security Survey (Computer Security Institute 1997). The most notorious survey of the cost that organizations incur after a breach may be the Ponemon Group's (Ponemon Institute 2015) annual study. The Ponemon study has been criticized for methodological issues and a simplistic per-record cost figure that does not accurately reflect costs but invites facile citation by the popular press, product vendors, and security consultants (Hackett 2015; Jacobs 2014; Shostak 2011). Verizon's Data Breach Investigations Report (DBIR) (Verizon Enterprise Solutions 2015), which added an estimate of the cost of breaches for the first time in its 2015 edition, argues that the cost of a breach is best modeled by a nonlinear function of the number of records breached.

A popular empirical method of estimating the cost of breaches to firms is to measure the effect of a breach announcement on stock prices. One of the earliest studies to use this approach found an average abnormal drop in stock price of 4.5% over 3 days in the 22 security breach events in the authors' sample (Garg et al. 2003). Other studies have found similar short-term postbreach drops in market value (Acquisti et al. 2006; Campbell et al. 2003; Cavusoglu et al. 2004; Gatzlaff and McCullough 2010; Goel and Shawky 2009; Gordon et al. 2011) and profits (Gwebu et al. 2014; Osei-Bryson et al. 2012). Although recent research still finds a statistically significant short-term drop in stock prices, the effect has gone down over time, perhaps because breaches have become more commonplace (Gordon et al. 2011). In contrast to the short-term hit on stock price, most firms do not appear to suffer long-term drops in market value after a breach (Kannan et al. 2007). And the effect of different types of breach is not uniform. In a study of 43 security breaches from 1995

to 2000, breaches related to confidential information were associated with drops in stock prices, but breaches that "largely affected the information infrastructure itself" were not (Campbell et al. 2003).

Another area of data breach economics research focuses on the effects of data breach notification laws. Lenard and Rubin (2005) have argued that the costs of these laws outweigh their benefits. Extrapolating from limited public data on the cost and incidence of identity theft, they concluded that the expected benefit from notifying consumers of a data breach was in the range of $7.50 to $10—lower than the costs they listed from notification, which included $10 to $20 per card to reissue cards and $2 per card to send notification letters. But even if notification laws increase costs to firms, they may reduce overall social costs by causing firms and consumers to improve their levels of data security care (Romanosky et al. 2010). Romanosky et al. (2011), for example, found that notification laws may reduce identity theft by about 6%.

Statistics about data breaches have also been used as inputs to empirical analyses of the effectiveness of data security investments. Miller and Tucker (2010), for example, found no evidence that adoption of encryption software among hospitals reduced the number data breaches. To the contrary, they found that public announcements of certain types of data breach actually increased. Gaynor et al. (2012) used an analysis of breach data to reach the surprising conclusion that hospitals in competitive healthcare markets seem to be worse at protecting patient data than those in noncompetitive markets. Kwon and Johnson (2011) applied a proportional hazard model to breach disclosures by 281 healthcare organizations to find that security measures appear to be more effective when adopted voluntarily instead of being forced by regulation.

There is little to no academic literature on how financial institutions decide whether to reissue cards after a breach, a decision process that the institutions treat as proprietary. The sole source we could find, other than news reports, is a 2008 study by the state of Maine (2008) surveying banks' responses to two major data breaches in that state. That study reported that issuers reissued 78% of cards during the period covered by the survey.

3 METHODOLOGY

We are interested in estimating and comparing the aggregate first-order net social, or aggregate, costs that result from decisions by issuers who, upon a credit card breach, face a choice between reissuing cards or waiting. By "social costs" we mean the total costs regardless of who incurs them (although we ignore any benefit gained by criminals). We use the term "first-order costs" to refer to those costs that are proximate results of reissuing cards or leaving them in circulation despite possible compromise. These can include the costs (including overhead) of mailing replacement cards, time spent by merchants and consumers responding to having cards reissued, or, for cards that are not reissued, the expected cost of fraud on those cards. Because we are interested in aggregate social costs, who incurs the cost of fraud is less critical to the model than is the total amount of that fraud.

We restrict the scope of our analysis in a number of ways. First, we concentrate on credit cards rather than debit cards or other payment instruments that have different authentication structures and risk profiles from credit cards. Second, our analysis is specific to the United States. Third, we concentrate on overall social costs, largely because, to our knowledge, no publicly available data exists that would enable an analysis of the allocation of those costs between parties. Fourth, we focus specifically on existing-account credit card fraud as the primary cost of credit card fraud. This is a subtype of identity theft in which victims' existing credit cards are used for unauthorized charges. Credit card data is unlikely to facilitate other forms of identity theft such as new-account fraud (in which new accounts are opened using the victim's identity) because opening a new account requires more than a credit card.

We use the following model of first-order costs:

$$\sum_k r c_{i_k} + (1 - r)\rho_k f_k. \tag{1}$$

The model sums, over each affected card k, the costs related to that card, with the following terms:

- r: Binary variable where $r = 1$ if the card is reissued and $r = 0$ if not.
- c_{i_k}: Cost of reissue for issuer i_k of card k.
- ρ_k: Probability that card k will be used fraudulently.
- f_k: Amount of fraud if the card is used fraudulently.

The model omits potential costs to cardholders and merchants of issuers reissuing cards because we assume that these costs are relatively small. A canceled and reissued credit card used for recurring payments may lead to a merchant having to contact customers to obtain new payment information, but these processes are generally automated and inexpensive (Authorize.Net 2016; Cayan 2010). The costs to cardholders come from the value of time spent responding to the cancelation—for example, updating auto-pay accounts to use the new card number. Issuers can minimize these costs by sending replacement cards before canceling outstanding cards, but cardholders do sometimes miss or ignore the payment cards or are traveling when the replacements are made (Stark 2004).

The model also assumes that fraud losses are zero if cards are reissued. Although it may be possible to use canceled cards fraudulently if a merchant is not vigilant about clearing authorization before goods or services have been rendered, we assume that the overall loss from these preauthorization transaction losses is negligible.

Because no data is publicly available for ρ_k, we must estimate it. We use the following equation:

$$\rho_k = (1 - \delta)\left[1 - \left(1 - \frac{vb}{\theta(n_d\gamma(1 - \lambda) + n_u)}\right)^{1/a}\right]. \tag{2}$$

This part of our model relies on the following parameters:

- n_d: Number of payment card records affected each year in disclosed data breaches for which the number of records affected is made public.
- n_u: Number of payment card records affected each year in disclosed data breaches for which the number of records affected is either unknown or not made public.
- θ: Scaling factor to account for payment card records exposed in breaches that are either undiscovered, undisclosed, or not included in the data to which we have access.
- λ: Breached credit cards that are immediately reissued by issuing institutions, as a proportion of all breached credit cards.
- γ: Credit cards as a proportion of all payment cards. This parameter captures the fact that breach disclosures may use "credit cards" to refer to payment cards generally, whereas our model focuses solely on credit cards.
- v: Number of people victimized by existing-account credit card fraud per year.
- b: Proportion of existing-account credit card fraud attributable to data breach as the method by which the card data was obtained.
- a: Average number of credit cards per cardholder. This parameter allows us to use per-person data on the cost of credit card fraud in our model of the cost per card.
- δ: Reduction in the probability of fraud achieved by an issuer flagging breached cards in its fraud detection algorithms.

Equation (2) estimates the probability of card misuse following a breach as a function of the number of existing-account credit card fraud incidents attributable to data breach (vb); the total number of credit card records exposed in breaches each year, including those that are not included in breach databases either because the scope of a breach was unknown or because the breach was not discovered or publicly disclosed ($\theta(n_d\gamma(1-\lambda)+n_u)$); the effectiveness of fraud detection algorithms (δ); and the number of credit cards per person (a).

We estimate the amount of fraud if a card is misused as

$$f_k = c_{m_k} + t_k c_{t_k} + c_{i_k}. \tag{3}$$

The parameters in this part of the model are:

c_{m_k}: Monetary cost of existing account credit card fraud per incident.
t_k: Time (in hours) spent responding to existing account credit card fraud by cardholders.
$\mathbf{c}_{t_k}$: Cost of cardholder time (per hour).

We include c_{i_k}, first used in Equation (4), to capture the cost of canceling and reissuing cards that have been used for fraud.

Substituting the formulas for ρ_k and f_k in Equation (4) results in the following model that includes all parameters:

$$\sum_k rc_{i_k} + (1-r)(1-\delta)\left[1-\left(1-\frac{vb}{\theta(n_d\gamma(1-\lambda)+n_u)}\right)^{1/a}\right] c_{m_k} + t_k c_{t_k} + c_{i_k}. \tag{4}$$

This calculation assumes that the card records exposed in breaches are unique—i.e., that two different breach events do not expose the same credit card record. Overlap between breaches would reduce the total number of credit card records exposed. This assumption seems reasonable given the current common (but not universal) practice of reissuing credit cards potentially exposed in a breach. We also assume that the same breached card is not victimized twice (where a "victimization" may include multiple fraudulent charges). This follows from our assumption that fraudulently used cards will immediately be cancelled and reissued once that fraud is detected.

The calculations also use annual averages even though the number of cards exposed in data breaches varies widely from year to year. Using annual averages reflects the assumption that both collection and misuse of credit cards occurs over time. Although a massive breach may be announced on a certain date, access to the data may have occurred over weeks or months. Thus, it seems to make sense to smooth this data by considering annual averages and not focusing on individual yearly totals.

4 DATA

In this section, we discuss the data sources for each of the parameters presented above and describe the ranges and point estimates used for each parameter.

4.1 The Cost of Reissuing Cards (c_{i_k}, q_k)

Three types of data sources shed light on the cost of reissuing cards: news reports, lawsuits, and a state government survey. News reports have quoted figures from issuers and other industry sources; these estimates range from $3 to $25 per card (America's Community Bankers 2007; Aspan and Baldwin 2011; Churchill 2008; Holmes 2015; Inscoe 2012; Jewell 2004; Johnson 2011; Ravana 2007; Stark 2004). Lawsuits filed by issuers seeking to recover the cost of reissuing cards claim losses from reissuing of $5 to $20 per card, with some evidence that economies of scale reduce the per-card cost when an institution must reissue more cards (Pennsylvania State Employees Credit

Union v. Fifth Third Bank 2005). The state of Maine conducted a survey that found a cost of $4.72 per card reported by issuers in that state (Maine Bureau of Financial Institutions 2008).

Considering these sources as a whole, it appears that the cost is between $5 and $25 per card. Because the cost for most issuers seems to be $10 or less, we use $10 as our point estimate.

4.2 The Probability of Credit Card Misuse Following a Breach (ρ_k)

To the best of our knowledge, no publicly available data exists on the probability that a credit card affected in a data breach will be used for fraud. We estimate that probability by multiplying the number of annual incidents of existing-account credit card fraud (v) by the proportion of those incidents in which the credit card data was obtained using data breach (b) and then dividing that by the total number of credit cards exposed in data breaches each year ($\theta(n_d\gamma(1-\lambda)+n_u)$).

We assume that the majority of issuers already use some form of fraud monitoring. This has two implications. First, the marginal cost to monitor a card that has been exposed in a data breach is essentially zero. Setting a flag in an issuer's fraud monitoring system has negligible marginal cost if the database is set up to accommodate such a flag. Second, this assumption implies that the current level of existing-account credit card fraud already reflects the use of fraud-monitoring and prevention systems. Flagging a card might improve the probability that attempted fraud will be detected and prevented—at some risk of additional false positives—but the baseline probability of fraud does not rely on the effectiveness of current fraud-monitoring processes.

4.2.1 Payment Cards Exposed in Data Breaches with Record Counts (n_d). The number of records exposed in data breaches is uncertain for three reasons. First, only breaches that are discovered can be counted. Second, not all discovered breaches are publicly disclosed. And third, even when a breach has been detected and reported, it may not be possible to determine how many records were exposed. Our model contains parameters for three types of breached payment card records: those that are publicly disclosed with estimated record counts (n_d), those that are disclosed with unknown record counts (n_u), and a scaling factor to account for undetected breaches (θ).

The record count we use in our model is based on a detailed analysis of the Privacy Rights Clearinghouse (PRC) database (Privacy Rights Clearinghouse 2016b). We calculated the number of payment cards potentially exposed by downloading the PRC database, filtering based on the use of the word "card" in the description field, and manually categorizing each entry, based on its description, as having potentially exposed full unencrypted payment card numbers or not. Thus, we did not include breach events that were described as having exposed only partial or encrypted payment card numbers. We did not, however, filter out breach events in which card numbers were exposed without other "full track" data such as expirations dates. Recent work by Ali et al. (2017) shows that due to different online merchants using different fields for verifying card transactions, it is easy for an attacker with just a card number to discern all the other information needed to use the number for fraud. We also omitted events disclosed in 2005 because breach reporting was still new and the 16 events reported for that year were probably nonrepresentative. Where necessary, we updated PRC's record counts to reflect only the number of payment cards believed to have been exposed. Our analysis of the PRC database yielded a list of 579 breach events from 2006 through the end of 2014. Of those events, 269 included record counts (46%).

We supplemented this data with information from the Maine, Maryland, and New Hampshire data breach lists (Maine Attorney General 2014; Maryland Attorney General nd; New Hampshire Office of the Attorney General nd). We used these states' lists to add breaches that were not included in the PRC database and to estimate record counts for breaches where PRC did not have those numbers. This added 179 breach events to our database—including 34 with overall record counts—for a total of 758, of which 303 included record counts (40%). Those records total

378 million payment card accounts over 9 years. We use that average of about 42 million cards per year as our point estimate. Because the number of reported records is relatively well known, we use the narrow range of 39 million to 45 million cards per year for this parameter.

4.2.2 Payment Card Records Exposed in Data Breaches without Record Counts (n_u). We used two methods to estimate the total number of records exposed in breach events without record counts. First, we used linear regressions to predict overall record counts from the number of residents of Maine, Maryland, and New Hampshire that were affected. This gave us estimates for an additional 231 breach events at a total of about 630,000 records per year.

For the remaining events in our database, we extrapolated using a weighted estimate based on the typical number of records exposed for each type of data breach. We excluded the TJX, Heartland, Target, and Home Depot breaches because we believe it unlikely that any of the disclosed breaches with unknown record counts could have exposed records on the order of the tens of millions or hundreds of millions of records exposed in those four breaches. We also excluded one insider breach at Fidelity because it appears to be an extreme outlier: the 8.5 million records compromised in that breach were two orders of magnitude larger than any other insider breaches with known record counts and three orders of magnitude larger than the average in that category when the Fidelity breach is excluded.

Based on the weighted average, we estimate that the 224 breaches with unknown record counts from 2006 through 2014 have exposed about 2.4 million accounts per year. This estimate may still be too high because it includes nine other breaches in which at least a million records were believed to have been affected. Excluding these breaches gives a weighted estimate of about 580,000 records per year from unreported breaches. We use a point estimate of $n_u = 2.1$ million unknown breached records per year, which is derived from the linear regression estimate of 630,000 added to the midpoint between the 580,000 and 2.4 million estimates from the weighted average extrapolation. But the range we use is wide—1 million to 10 million cards per year—because of the uncertainty surrounding the number of records in breaches for which record counts were not disclosed.

4.2.3 Payment Card Records Exposed in Undetected or Undisclosed Breaches (θ). It is impossible to know how many breaches are not detected. There is, however, plenty of speculation. For example, one security product vendor (with the possible biases that implies) claims that 85% of data breach events are undetected (Friedlander 2014). The number of data breaches that went undetected for months or years suggests that there have probably been other breaches that were not detected at all (Gold 2014; Kerner 2014; Popper 2014; PYMNTS 2015). Perhaps more enlightening are the controlled penetration tests conducted at government agencies. In fiscal year 2011, 49% of those intrusions were detected (Office of Management and Budget 2013). That increased to 73% in fiscal year 2013 (Office of Management and Budget 2014). Although these numbers are the results of controlled tests against specific goals and agencies, they offer a general idea of the extent to which intrusions are detected overall.

The other type of unknown included in θ is the number of records in breaches that are detected but not disclosed. Some surveys attempt to measure a similar variable. For instance, one survey of "malware analysts" found that 57% claimed that their organizations had not disclosed data breaches (ThreatTrack Security 2014). There are serious methodological problems with this figure, such as the difficulty in translating the number of analysts to a number of records and lack of clarity as to the definition of a "breach," but better information does not seem to be available.

Because this number is subject to much uncertainty, we use the broadest range that seems plausible. We assume that undetected and undisclosed breaches expose between one-fourth and three times as many records as are exposed in detected breaches, with a conservative point estimate of $\theta = 1.75$. This potentially overestimates the number of records that are exposed, which could be

Table 1. Estimated Total Number of Credit Card Records Exposed in Data Breach per Year

Description	Low	Point	High
Payment card records reported lost in data breaches per year (n_d) (mil)	39	42	45
Est. records per year in breaches with unknown record counts (n_u) (mil)	1.00	2.10	10.00
Scaling factor to account for unreported or undetected breaches (θ)	1.25	1.75	4.00
Portion of breached cards reissued (λ)	0.95	0.88	0.80
Credit cards as a proportion of breached payment cards (γ)	0.47	0.52	0.58
Total credit card records exposed in all breaches per year	2,400	8,500	60,900

the case if, for example, undetected or undisclosed breaches tend to be smaller than those that are detected and disclosed.

4.2.4 Proportion of Breached Cards That Are Immediately Reissued (λ). The next factor needed to calculate ρ_k is the percentage of breached credit cards that are reissued before fraud occurs. A 2008 Maine study of banks' responses to data breach incidents found that issuers reissued 78% of cards during the period covered by the survey, including 84% of accounts affected in the TJX breach and 77% of those affected in the Hannaford breach (Maine Bureau of Financial Institutions 2008). Another source claims that "nearly 90 percent of card breach victims in 2014 received replacement credit cards" (Holmes 2015). We therefore assume that issuers reissue between roughly 80% and 95% of cards, with a point estimate of 87.5%.

4.2.5 Credit Cards as a Proportion of Payment Cards (γ). According to data from the Statistical Abstract of the United States and the Nilson Report, credit cards (excluding store cards, oil company cards, and other non-general-purpose cards) have decreased as a percentage of all payment cards from 58% in 2008 to 47% in 2014 (HSN Consultants, Inc., 2013, 2014, 2015; U.S. Census, 2011 t. 1186, 1187, U.S. Census, 2012 t. 1187, 1188). Under the assumption that the proportion of credit cards to payment cards in breaches is the same as the proportion in general circulation, we use these values as the range of our model, with a point estimate of 52%.

Table 1 summarizes the ranges and point values for n_d, n_u, θ, λ, and γ. Taking the high and lows of this range, we estimate that between about 2.4 million and 60.9 million nonreissued credit card records are exposed in data breaches annually, with a point estimate of 8.5 million cards.

4.2.6 Number of People Affected by Existing-Account Credit Card Fraud (υ). The Department of Justice's Bureau of Justice Statistics (BJS) has included identity theft questions in its annual National Crime Victimization Survey (NCVS) since 2004 (Harrell 2015). These statistics are split by the nature of the crime; "existing account credit card identity theft" refers to situations in which existing credit cards were used without the cardholder's authorization.

In 2008, the BJS began adding an Identity Theft Supplement (ITS) to the NCVS. The questions in this supplement collected data on identity theft experienced by individuals instead of households. The 2008 supplement asked if respondents had experienced identity theft in the 2 years prior to the interview. In the 2012 and 2014 surveys, the ITS asked about individual-level identity theft over the previous 12 months. As a result, it is not possible to compare results across the 2005–2010 surveys, the 2008 survey, or the 2012–2014 surveys (Harrell and Langton 2013).

Using the 2012 and 2014 per-person data and taking the overall minimum and maximum of the 95% confidence intervals for each year results in a range of 6.8 million to 9.0 million people affected by existing-account credit card fraud each year. We take the average of the 2012 and 2014 point estimates to set $\upsilon = 8.15$ million people.

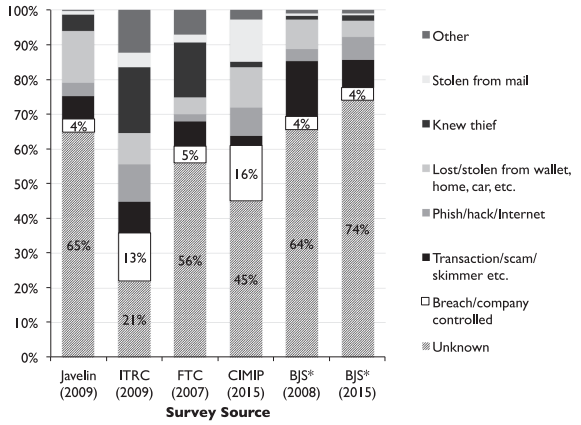

Fig. 2. Survey and study results for the number of identity theft victims who knew how their data was obtained and, if so, the point of compromise.

4.2.7 Proportion of Existing Credit Card Fraud Attributable to Breach (b). Not all credit card fraud is the result of breach. Victims of existing-account credit card fraud who know how their card information was obtained most often say that it was through a stolen wallet or a purse or from someone they know. Breach seems to be a relatively infrequent cause of credit card fraud, but it is uncertain how infrequent. Surveys by Javelin Research, the Identity Theft Research Center, the FTC, and the DOJ's Bureau of Justice Statistics have asked victims of identity theft if they knew how their information was obtained (Bureau of Justice Statistics 2014; Harrell and Langton 2013; Center 2010; Javelin Strategy & Research 2009; Langton and Planty 2010; Synovate 2007). Utica College's Center for Identity Management and Information Protection (CIMIP) analyzed the same question (among many others) using federal criminal case data (Gordon et al. 2007; Rebovich et al. 2015). Figure 2 shows the results of these studies.

Most survey respondents did not know how their data was obtained. The responses of those who said that they knew how their data was obtained can legitimately be generalized only if the point of compromise and the victim's knowledge of that point of compromise are uncorrelated. But this may not be true. Some points of compromise are more likely to be known than others. Lost wallets, purses, or thefts alert a cardholder that their cards may have been stolen. Other points of compromise, such as skimmers (devices that surreptitiously record card data at an ATM or point of payment), are unlikely to be recognized. People whose cards are compromised through phishing or spyware will not always know that their cards were obtained in that matter. A data breach, of which a cardholder must be notified in 46 of 50 states, may be more or less likely to be a known point of compromise.

The ITRC survey is an outlier in this set, with the highest percentage of known points of compromise and the highest percentage of people responding that their data was obtained in a breach. As the ITRC (2010) acknowledges, "this may be due to the fact that ITRC is listed as a victim resource by many entities which have suffered a breach."

Only the BJS survey reported responses for points of compromise specifically for existing-account credit card fraud. None of the other surveys distinguished between forms of identity theft in their reporting. Based on the factors listed above, we use a range of 5% to 15% as the proportion of existing-account credit card fraud in which the card information was obtained in a data breach, with a point estimate of 11%. We choose this range to capture, at the low end, either the lowest estimate for breach as a percentage of known points of compromise or the midrange of estimates for

Table 2. Calculation of the Probability of Existing-Account Credit Card Fraud to an Account Affected by a Breach

Description	Low	Point	High
Number of credit cards exposed (from Table 1) (mil)	2.40	8.50	60.90
Number of persons victimized (v) (mil)	6.80	8.15	9.00
Percent of existing-account credit card fraud from breach (b)	5%	11%	15%
Fraud reduction from flagging exposed cards (δ)	0%	10%	20%
Average number of credit cards per cardholder (a)	3.6	3.8	4.0
P (existing-account credit card fraud | breach) (ρ_k)	0.0011	0.026	0.21

breach as a percentage of all compromise, including unknown sources. The high end of the range is just below the ITRC's number, which has a high number of people who believe they know how their information was obtained and the aforementioned potential bias toward identifying breach as the way credit card information was obtained.

4.2.8 Reduction in Fraud from Flagging Breached Cards (δ). We assume that flagging exposed cards reduces fraud rates by up to 20%. As discussed at the start of this section, current levels of fraud monitoring are already reflected in existing credit card fraud statistics. Thus, marking a card as potentially exposed can at best improve the effectiveness of fraud-monitoring systems somewhat. Unfortunately, information on the effectiveness of fraud-monitoring software is treated as proprietary by both issuers and the software vendors. The 0% to 20% range (with a 10% point estimate) therefore represents our best guess.

4.2.9 Number of Credit Cards per Cardholder (a). Converting from the per-person data reported by the BJS to per-card numbers requires an estimate of the number of credit cards per cardholder. According to Gallup polls, credit card owners held an average of about 3.6 to 3.7 credit cards each from 2006 to 2014 (Swift 2014). Surveys conducted by the Federal Reserve Bank of Boston found that cardholders had between 3.8 and 4.0 cards each from 2010 to 2012 (Schuh and Stavins 2014). These numbers include Mastercard, Visa, American Express, and Discover cards but exclude store cards (which are valid only at the stores that issue them), gas company cards, and other specialty cards such as phone cards. We use the high and low end of these numbers as our estimated range of 3.6 to 4.0, with a point value in the middle at 3.8.

4.2.10 Calculation of ρ_k. Using the parameter values discussed above (which are summarized in Table 2) results in an estimated range for ρ_k of 0.0011 to 0.21, with a point estimate of 0.026.

4.3 The Cost of Credit Card Fraud (f_k)

The cost of an existing-account credit card fraud incident has two components: financial losses, including both the loss of value obtained through the fraud and indirect financial costs from responding to the fraud, and the cost of time spent dealing with the fraud.

In most cases, a cardholder should suffer little or no direct out-of-pocket loss from existing-account credit card fraud. Federal law limits cardholder liability to $50 for unauthorized credit card charges if a lost or stolen card is reported as soon as the loss or theft is discovered [15 U.S.C. § 1643(a)(1)(B); 12 C.F.R. 226.12]. Visa and Mastercard have voluntary zero-liability policies that further reduce consumer liability for card fraud (Akers et al. 2005). Despite these policies, cardholders may still experience out-of-pocket losses if they do not report lost or stolen cards quickly enough.

Table 3. Expected Cost per Card of an Existing-Account Credit Card Fraud Incident

Description	Low	Point	High
Mean monetary cost of existing-account card fraud (c_{m_k})	$1,000	$1,200	$1,400
Mean hours spent responding to existing-account card fraud (t_k)	2	3	4
Cost of time per hour (c_{t_k})	$12	$15	$20
Cost of reissuing cards used for fraud (c_{i_k})	$3	$10	$25
Total expected cost of an existing-account card fraud incident (f_k)	$1,027	$1,255	$1,505

According to the 2012 BJS survey, the average combined direct and indirect loss from existing-account credit card fraud was about $1,400 for the 69% of people who experienced any loss (Harrell 2015). In 2012, it was about $1,000, with 66% experiencing a loss (Harrell and Langton 2013). Based on these numbers, we estimate the range of average cost per existing-account credit card fraud at $1,000 to $1,400 with a point estimate of $1,200.

The 2012 and 2014 BJS surveys reported that victims of existing-account credit card fraud spent an average of 3 and 4 hours, respectively, resolving problems. A 2006 FTC report indicated that victims of existing-account credit card fraud spent a median of 2 hours resolving problems (Synovate 2007). We therefore use a range of 2 to 4 hours for t_k with a point estimate of 3.

For the cost-of-time parameter, we assume an average annual wage of $45,500 per full-time employee, discounted 50% on the assumption that most time spent responding to breach occurs during nonwork time (U.S. Census 2012, t.647). This corresponds to a $12 to $20 cost of time, with a point estimate of $15.

Our estimate for f_k is dominated by the monetary cost of fraud, as shown in Table 3, with a range of about $1,027 to $1,505 with a point estimate of $1,255.

5 ANALYSIS

Table 4 summarizes the basic analysis of the per-card cost of reissuing versus not reissuing cards, with ranges and point estimates. Our model estimates the expected cost of not reissuing cards at between $1.15 and $310 per card, with a point estimate of $32.80. This wide range corresponds to a potential savings of about $24 per card or loss of $307 per card. The point estimate is a $22.80 per-card loss by not reissuing. Multiplying these estimates by the number of reported breached card accounts implies that $960 million might be lost by not reissuing cards immediately after a breach. The range of estimation is extreme, however: over $1 billion might be saved by not reissuing cards, but the potential total loss calculated by this model is almost $14 billion.

5.1 Monte Carlo Analysis

Monte Carlo simulations allow us to estimate the distribution of likelihood along the broad range of results. Because we have no reason to assume any particular distribution for our parameters, we used PERT Beta distributions with the highs, lows, and point estimates of our ranges as the equivalent values of the distributions. The resulting distributions show a wide variation in possible costs, with some overlap between the reissue and no-reissue situations.

Figure 3 shows a histogram of the expected per-card cost of fraud when cards are not reissued. The 90% confidence range is from $11.20 per card to $44.60 per card, with a mean of $24.60. The distribution resembles the heavy-tailed models found for cyber-risk and data breach in previous work in the literature (Edwards et al. 2015; Maillart and Sornette 2010). Figure 4 shows a histogram of the total cost reduction that could be achieved from not automatically reissuing credit cards. The 90% confidence range is ($-1.4 billion, $88 million), with about a 91% probability that immediately reissuing cards would be the lower-cost option.

Table 4. Comparison of the Per-Card Cost of Reissuing versus Not Reissuing Cards

Description	Low	Point	High
Reissue cost, per card	$3.00	$10.00	$25.00
Expected cost if not reissued, per card	$1.15	$32.80	$310.00
Per-card savings (cost) from not reissuing cards	($307.00)	($22.80)	$23.85
Payment card records reported lost in data breaches per year (n_d) (mil)	39	42	45
Cumulative savings (cost) from not reissuing (mil)	($13,800)	($960)	$1,080

Note: Cumulative savings are based on the number of reported breach records, not the estimated total.

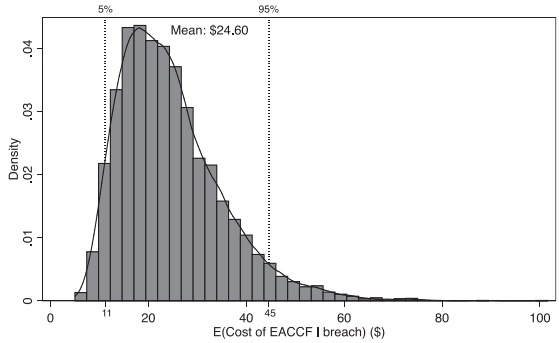

Fig. 3. Distribution of the cost per card to reissue or not reissue cards based on a Monte Carlo simulation.

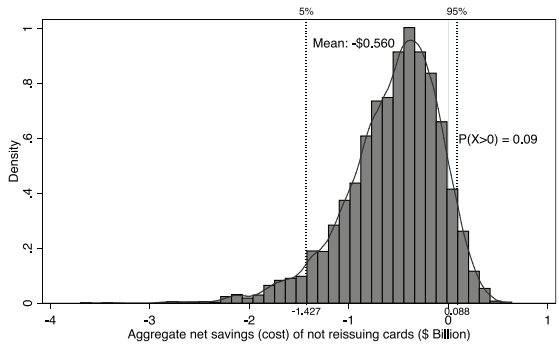

Fig. 4. Histogram of cumulative savings from not automatically reissuing cards according to a Monte Carlo simulation.

5.2 Sensitivity Analysis

Figure 5 is a tornado diagram showing the sensitivity of the per-card cost of not reissuing cards for the variables to which the cost is most sensitive. Each row shows the effect on the mean result of increasing the parameter by one standard deviation.

Unsurprisingly, the model is most sensitive to the parameters with the greatest uncertainty. The number of reported breaches with unknown record counts and the scaling factor for unreported breaches are both significant factors in the estimate. Each of these parameters reduces the mean estimate by over $5 of the roughly $25 mean expected cost of fraud from not reissuing. The model is also particularly sensitive to the percentage of existing-account credit card fraud attributable

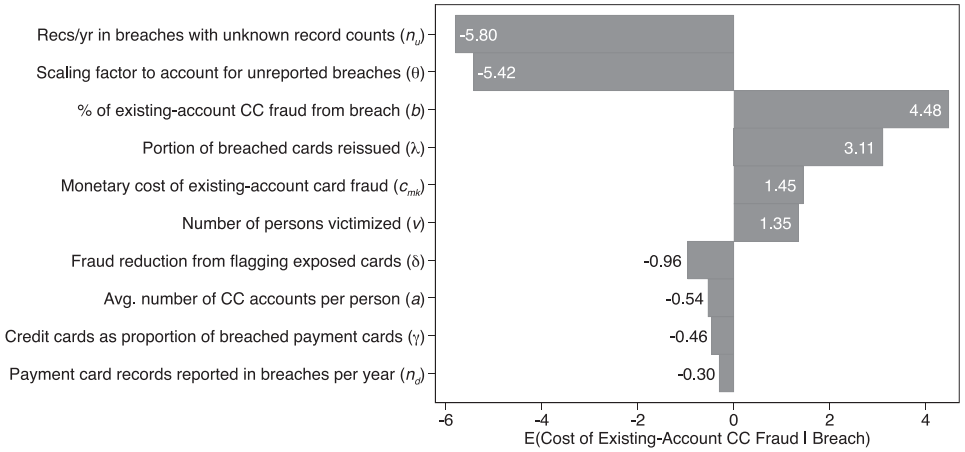

Fig. 5. Tornado diagram of variables affecting the per-card cost of not reissuing cards.

to breach. An increase in that parameter by one standard deviation increases the mean expected cost of fraud by about $4.50. A fourth parameter that is not well understood—the percentage of breached cards that are reissued—also has a large effect.

Although we focus in this article on parameter uncertainty, our results can change dramatically due to model uncertainty. As we mentioned in Section 4.2.6, BJS statistics on identity theft were originally collected by household, then more recently by individual. We use the individual-level data in our model because it allows us to avoid potential issues involving multiple cardholders per account in a household and reduces the number of instances in which one unit suffered multiple fraud incidents, violating one of our assumptions. When we began this work, however, individual-level data was not available and we used per-household data. The results of our model when we use per-household calculations are quite different than those initial results, even accounting for other refinements to the model since our initial work.

The only parameters that change are the number of cards per household and the number of households victimized by existing-account identity fraud. The number of cards per household is roughly similar to the number of cards per person—a range of 3.1 to 5.8 depending on year and source, as calculated by total cards divided by number of households—but the number of households experiencing existing-account credit card fraud was between 3.6 million and 5 million according to the BJS 2005–2010 survey (Langton 2011). As a result, the range of a per-household calculation would be an expected cost of credit card fraud on breached cards of between $0.42 and $170 per card, with a point estimate of about $14—about $18 less than the estimate using an individual-level calculation.

6 DISCUSSION AND LIMITATIONS

6.1 Discussion

In answer to the question posed in the title of this article, reissuing cards immediately after a breach appears to be less costly than waiting for attempted fraud before reissuing. This result is fairly robust despite the wide uncertainty in the estimated cost of fraud after a breach. Our Monte Carlo analysis estimates a 9% probability that waiting to reissue cards until fraud is detected would save money. The uncertainty in our model is partly because we rely on public data sources for our parameters and partly because the data sources themselves are subject to tremendous uncertainty.

Our sensitivity analysis suggests where resources could be best targeted to getting better data for parameters critical to our model. Specifically, it would be useful to get better information on how identity thieves get access to credit card data. Surveys of victims are clearly inadequate; too many people simply do not know how their data was obtained. Issuers, however, have the ability to connect breach notification with card misuse. Issuers also have information, at least collectively, on the percentage of cards that they reissue after a breach. Access to this data would undoubtedly improve our understanding of the benefits of options following a data breach. Access to that data would come with its own costs, of course, whether through compliance with a regulatory data-sharing regime or through costs of voluntary industry data sharing. A comparison of the costs and benefits of increased data sharing by card issuers would be an opportunity for future work.

A data-reporting regime may create its own perverse incentives. Participants in the card ecosystem who have full knowledge of the model used to make policy decisions might have incentives to manipulate that data. This incentive effect of disclosure is another topic worthy of future study.

Our work building a database of credit card breaches shows that despite extensive breach reporting requirements, information about breaches is often incomplete. More states could follow the lead of Maine, Maryland, and New Hampshire by requiring that breached organizations not only report the breach to the state attorneys general but also provide detailed information about the breach, such as the number of residents affected, the cause of the breach, and the type of data breached. If states could agree on a standard form for breach reporting, the burden on reporting organizations could be held to a minimum.

6.2 Limitations and Opportunities for Future Work

The analysis in this article is subject to several limitations, each of which presents an opportunity for future study. One obvious and major limitation (as well as motivation) of this work is the lack of data on the causes, extent, and effects of data breach. Efforts such as the National Cyber Leap Year have attempted to fill this gap (Chong et al. 2009), but much more work is needed to create the type of data that can be used for reliable statistical analysis. We discuss the implications of the poor quality of available data more fully in (Graves et al. 2016).

Another limitation of the analysis described in this article is that it treats breaches as homogeneous—assuming, for example, that a small number of records in an improperly discarded report creates the same risk of data exploitation as the hacking of a large database. In particular, the model used in this analysis takes limited account of the wide variation in breach size. It might, for example, be socially optimal to reissue cards after "everyday" breaches but not after megabreaches of 1 million cards or more, or it might be worth reissuing after hacking breaches but not after breaches due to improperly discarded records.

Our analysis does not account for "second order" effects—those indirect costs that occur over time or that are in some other sense a step removed from the immediate costs. For example, immediately reissuing cards reduces the window during which thieves can attempt fraud, which can both dissuade credit card theft and make attribution and detection of fraud easier. Another second-order effect lies in cardholder behavior after his or her card has been affected in a breach. Cardholders may expect to have cards reissued automatically and reduce card usage—and perhaps overall spending—if they are not. These second-order costs weigh in favor of reissuing, thus strengthening the case for immediate reissue of breached cards.

Our work is limited by lack of access to transaction-level card data. A researcher with industry access could improve on this work by combining the analysis of public data we present here with data on issuers' costs. Data held by issuers could yield information about fraud probability and losses by cardholder demographics, breach attributes, and so forth.

6.3 Conclusion

Having determined that immediately reissuing cards appears to have a lower social cost, what are the policy implications? As mentioned at the beginning of this article, card association rules allow issuers to recover fraud costs that result from breached cards but not the operational costs of reissuing them. The card association rules may create incentives for issuers to wait before reissuing cards, which is the opposite of what our model suggests to be the socially optimal incentive. Limited evidence (discussed in Section 4.2.4) suggests that card issuers often do routinely reissue cards affected in a breach despite these incentives. If in fact this practice is widely followed, our research suggests that it is socially optimal.

REFERENCES

Alessandro Acquisti, Allan Friedman, and Rahul Telang. 2006. Is there a cost to privacy breaches? An event study. In *Proceedings of the 27th International Conference on Information Systems*.

Douglas Akers, Brian Lamm, Jay Golter, and Martha Solt. 2005. Overview of recent developments in the credit card industry. *FDIC Banking Review* 17, 3 (2005), 23–35. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=882103.

Mohammed Aamir Ali, Budi Arief, Martin Emms, and Aad van Moorsel. 2017. Does the online card payment landscape unwittingly facilitate fraud? *IEEE Security & Privacy* 15, 2 (2017), 78–86.

America's Community Bankers. 2007. ACB data breach survey highlights need for action by card networks and Congress. Retrieved from http://www.prnewswire.com/news-releases/acb-data-breach-survey-highlights-need-for-action-by-card-networks-and-congress-54632297.html.

Maria Aspan and Clare Baldwin. 2011. Sony breach could cost card lenders $300 mln. Retrieved from http://www.reuters.com/article/2011/04/29/sony-creditcards-cost-idUSN2826485220110429.

Authorize.Net. 2016. Pricing. Retrieved November 2, 2016, from http://www.authorize.net/solutions/merchantsolutions/pricing/.

Bureau of Justice Statistics. 2014. National Crime Victimization Survey: Identity Theft Supplement, 2012. Retrieved from http://doi.org/10.3886/ICPSR34735.v1.

Katherine Campbell, Lawrence A. Gordon, Martin P. Loeb, and Lei Zhou. 2003. The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. *Journal of Computer Security* 11, 3 (2003), 431–448. Retrieved from http://content.iospress.com/articles/journal-of-computer-security/jcs192.

Huseyin Cavusoglu, Birendra Mishra, and Srinivasan Raghunathan. 2004. The effect of internet security breach announcements on market value: Capital market reactions for breached firms and internet security developers. *International Journal of Electronic Commerce* 9, 1 (2004), 70–104. Retrieved from http://www.tandfonline.com/doi/abs/10.1080/10864415.2004.11044320.

Cayan. 2010. Insights: Authorization fee. Retrieved November 2, 2016, from https://cayan.com/glossary/authorization-fee.

Identity Theft Resource Center. 2010. Identity theft: The aftermath 2009. Retrieved from http://www.idtheftcenter.org/ITRC-Surveys-Studies/aftermathstudies.html.

Fred Chong, Ruby B. Lee, Claire Vishik, Alessandro Acquisti, William Horne, Charles Palmer, Anup K. Ghosh, Dimitrios Pendarakis, William H. Sanders, Eric Fleischman, Hugo Teufel, III, Gene Tsudik, Dipankar Dasgupta, Steven Hofmeyr, and Leor Weinberger. 2009. National Cyber Leap Year Summit 2009: Co-chairs' report. Retrieved from https://www.nitrd.gov/nitrdgroups/index.php?title=National_Cyber_Leap_Year_Summit_2009.

Chris Churchill. 2008. TJX reacts to bank lawsuit. *Times Union* (Aug. 2008).

Computer Security Institute. 1997. 1997 CSI/FBI computer crime and security survey. *Computer Security - Issues and Trends* (Spring 1997).

Benjamin Edwards, Steven Hofmeyr, and Stephanie Forrest. 2015. Hype and heavy tails: A closer look at data breaches. In *2015 Workshop of the Economics of Information Security (WEIS'15)*. Retrieved from http://www.cs.unm.edu/~forrest/publications/weis-data-breaches-15.pdf.

Gaby Friedlander. 2014. Why 85% of data breaches are undetected. Retrieved from http://www.observeit.com/blog/why-85-percent-data-breaches-undetected.

Ashish Garg, Jeffrey Curtis, and Hilary Halper. 2003. Quantifying the financial impact of IT security breaches. *Information Management & Computer Security* 11, 2 (May 2003), 74–83. DOI: http://dx.doi.org/10.1108/09685220310468646

Kevin M. Gatzlaff and Kathleen A. McCullough. 2010. The effect of data breaches on shareholder wealth. *Risk Management and Insurance Review* 13, 1 (2010), 61–83. DOI: http://dx.doi.org/10.1111/j.1540-6296.2010.01178.x

Martin S. Gaynor, Muhammad Zia Hydari, and Rahul Telang. 2012. Is patient data better protected in competitive healthcare markets? In *2012 Workshop on the Economics of Information Security (WEIS'12)*. Retrieved from http://weis2012.econinfosec.org/papers/Gaynor_WEIS2012.pdf.

Sanjay Goel and Hany A. Shawky. 2009. Estimating the market impact of security breach announcements on firm values. *Information & Management* 46, 7 (2009), 404–410. DOI: http://dx.doi.org/10.1016/j.im.2009.06.005

Steve Gold. 2014. Home Depot card data breach undetected for four months. Retrieved from http://www.scmagazineuk.com/news/home-depot-card-data-breach-undetected-for-four-months/article/372794/.

Gary Gordon, Donald J. Rebovich, Kyung-Seok Choo, and Judith B. Gordon. 2007. *Identity Fraud Trends and Patterns: Building a Data-Based Foundation for Proactive Enforcement.* Technical Report. Center for Identity Management and Protection, Utica College. Retrieved from http://www.utica.edu/academic/institutes/cimip/publications/index.cfm.

Lawrence A. Gordon, Martin P. Loeb, and Lei Zhou. 2011. The impact of information security breaches: Has there been a downward shift in costs? *Journal of Computer Security* 19, 1 (Feb. 2011), 33–56. DOI: http://dx.doi.org/10.3233/JCS-2009-0398

James T. Graves, Alessandro Acquisti, and Nicholas Christin. 2016. Big data and bad data: On the sensitivity of security policy to imperfect information. *Chicago Law Review* 83, 1 (2016), 117–137.

Kholekile L. Gwebu, Jing Wang, and Wenjuan Xie. 2014. Understanding the cost associated with data security breaches. In *Pacific Asia Conference on Information Systems (PACIS'14).* 386. Retrieved from http://aisel.aisnet.org/cgi/viewcontent.cgi?article=1391&context=pacis2014.

Robert Hackett. 2015. The hotly disputed black magic of data breach cost estimates. *Fortune* (April 2015). Retrieved from http://fortune.com/2015/04/24/data-breach-cost-estimate-dispute/.

Erika Harrell. 2015. *Victims of Identity Theft, 2014.* Technical Report NCJ 248991. Bureau of Justice Statistics. Retrieved from http://www.bjs.gov/content/pub/pdf/vit14.pdf.

Erika Harrell and Lynn Langton. 2013. *Victims of Identity Theft, 2012.* Technical Report NCJ 243779. Bureau of Justice Statistics. Retrieved from http://www.bjs.gov/index.cfm?ty=pbdetail&iid=5408.

Jay Heiser. 2002. Can information security surveys be trusted? Retrieved from http://searchsecurity.techtarget.com/feature/Can-information-security-surveys-be-trusted.

Tamara E. Holmes. 2015. Credit card fraud and ID theft statistics. Retrieved from http://www.creditcards.com/credit-card-news/credit-card-security-id-theft-fraud-statistics-1276.php.

Identity Theft Resource Center. 2016. Data breaches. Retrieved November 2, 2016, from http://www.idtheftcenter.org/id-theft/data-breaches.html.

Shirley W. Inscoe. 2012. *Global Consumers React to Rising Fraud: Beware Back of Wallet.* Technical Report. Aite Group.

Jay Jacobs. 2014. Analyzing Ponemon cost of data breach. Retrieved from http://datadrivensecurity.info/blog/posts/2014/Dec/ponemon/.

Javelin Strategy & Research. 2009. *2009 Identity Fraud Survey Report: Consumer Version.* Technical Report. Retrieved from https://www.javelinstrategy.com/uploads/files/901.R_Identity_Fraud_Survey_Consumer_Report.pdf.

Mark Jewell. 2004. IDs are a steal; thieves looking for credit numbers set their sights on big targets. *Columbian* (Aug. 2004), E.

Andrew Johnson. 2011. Card fraud risk low from breach at Citi. *American Banker* (June 2011), 10.

Karthik Kannan, Jackie Rees, and Sanjay Sridhar. 2007. Market reactions to information security breach announcements: An empirical analysis. *International Journal of Electronic Commerce* 12, 1 (2007), 69–91. Retrieved from http://www.jstor.org/stable/27751241.

Sean Micheal Kerner. 2014. UPS discloses data breach that went undetected for months. Retrieved from http://www.eweek.com/blogs/security-watch/ups-discloses-data-breach-that-went-undetected-for-months.html.

Juhee Kwon and M. Eric Johnson. 2011. An organizational learning perspective on proactive vs. reactive investment in information security. In *2011 Workshop on the Economics of Information Security (WEIS'11).* Citeseer. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.309.1297&rep=rep1&type=pdf.

Lynn Langton. 2011. *Identity Theft Reported by Households, 2005-2010.* Technical Report NCJ 236245. Bureau of Justice Statistics. Retrieved from http://www.bjs.gov/index.cfm?ty=pbdetail&iid=2207.

Lynn Langton and Michael Planty. 2010. Victims of Identity Theft, 2008. Special Report NJC 231680. Bureau of Justice Statistics. Retrieved from https://www.bjs.gov/index.cfm?ty=pbdetail&iid=2222.

Thomas M. Lenard and Paul H. Rubin. 2005. An economic analysis of notification requirements for data security breaches. *Emory Law and Economics Research Paper* 05-12. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=765845.

Adam J. Levitin. 2010. Private disordering: Payment card fraud liability rules. *Brooklyn Journal of Corporate, Financial, and Commercial Law* 5, 1 (2010), 1–48. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1570867.

T. Maillart and D. Sornette. 2010. Heavy-tailed distribution of cyber-risks. *European Physical Journal B* 75, 3 (2010), 357–364. DOI: http://dx.doi.org/10.1140/epjb/e2010-00120-8

Maine Attorney General. 2014. Privacy, identity theft and data security breaches. Retrieved November 2, 2016, from http://www.state.me.us/ag/consumer/identity_theft/index.shtml.

Maine Bureau of Financial Institutions. 2008. Maine data breach study. Retrieved from http://www.state.me.us/pfr/financialinstitutions/reports/index.htm.

Maryland Attorney General. n.d. Maryland information security breach notices. Retrieved November 2, 2016, from http://www.marylandattorneygeneral.gov/Pages/IdentityTheft/breachnotices.aspx.

Amalia R. Miller and Catherine Tucker. 2010. Encryption and data loss. In *2010 Workshop on the Economics of Information Security (WEIS'10)*. Retrieved from http://weis2010.econinfosec.org/papers/session1/weis2010_tucker.pdf.

New Hampshire Office of the Attorney General. n.d. Security breach notifications. Retrieved November 2, 2016, from http://doj.nh.gov/consumer/security-breaches/.

Office of Management and Budget. 2013. Fiscal Year 2012 Report to Congress on the Implementation of The Federal Information Security Management Act of 2002.

Office of Management and Budget. 2014. Annual Report to Congress: Federal Information Security Management Act.

Open Security Foundation. 2016. DataLossDB. Retrieved November 2, 2016, from http://datalossdb.org/.

Kweku-Muata Osei-Bryson, Myung Ko, and Humayun Zafar. 2012. Financial impact of information security breaches on breached firms and their non-breached competitors. *Information Resources Management Journal* 25, 1 (Jan. 2012), 21–37. DOI: http://dx.doi.org/10.4018/irmj.2012010102

Pennsyvania State Employees Credit Union v. Fifth Third Bank. 2005. 317 F. Supp. 2d. 398. (E.D. Pa. 2005).

Ponemon Institute. 2015. *2015 Cost of Data Breach Study: Global Analysis*. Technical Report. Retrieved from http://www-03.ibm.com/security/data-breach/.

Nathaniel Popper. 2014. Breach at Neiman Marcus went undetected from July to December. *New York Times* (Jan. 2014). Retrieved from http://www.nytimes.com/2014/01/17/business/breach-at-neiman-marcus-went-undetected-from-july-to-december.html.

Privacy Rights Clearinghouse. 2016a. Chronology of data breaches: FAQ. Retrieved from https://www.privacyrights.org/chronology-data-breaches-faq.

Privacy Rights Clearinghouse. 2016b. Data breaches. Retrieved November 2, 2016, from https://www.privacyrights.org/data-breaches.

PYMNTS. 2015. OPM data breach undetected for a year. Retrieved from http://www.pymnts.com/news/2015/opm-data-breach-undetected-for-a-year/.

Ann Ravana. 2007. Banks start credit card reissue. *Bangor Daily News* (Feb. 2007), 4.

Donald J. Rebovich, Kristy Allen, and Jared Platt. 2015. *The New Face of Identity Theft: An Analysis of Federal Case Data for the Years 2008 through 2013*. Technical Report. Center for Identity Management and Protection, Utica College. Retrieved from https://www.utica.edu/academic/institutes/cimip/New_Face_of_Identity_Theft.pdf.

Sasha Romanosky, Alessandro Acquisti, and Richard Sharp. 2010. Data breaches and identity theft: When is mandatory disclosure optimal? TPRC. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1989594.

Sasha Romanosky, Rahul Telang, and Alessandro Acquisti. 2011. Do data breach disclosure laws reduce identity theft? *Journal of Policy Analysis and Management* 30, 2 (March 2011), 256–286. DOI: http://dx.doi.org/10.1002/pam.20567

Julie J. C. H. Ryan and Theresa I. Jefferson. 2003. The use, misuse, and abuse of statistics in information security research. In *Proceedings of the 24th Annual National ASEM*.

Scott D. Schuh and Joanna Stavins. 2014. The 2011 and 2012 surveys of consumer payment choice. *Federal Reserve Bank of Boston Research Paper Series Research Data Reports* 14-1. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2564165.

Adam Shostack and Andrew Stewart. 2008. *The New School of Information Security*. Pearson Education. Retrieved from https://books.google.com/books?id=TWvC32p5M5YC.

Adam Shostak. 2011. A critique of Ponemon Institute methodology for "churn." Retrieved from http://newschoolsecurity.com/2011/01/a-critique-of-ponemon-institute-methodology-for-churn/.

Eric Stark. 2004. Computer hackers are stealing bank card information, but there is protection and some banks have been aggressive. *Sunday News* (July 2004), 1.

Art Swift. 2014. Americans rely less on credit cards than in previous years. Retrieved from http://www.gallup.com/poll/168668/americans-rely-less-credit-cards-previous-years.aspx.

Synovate. 2007. Federal Trade Commission—2006 Identity Theft Survey Report. Retrieved from https://www.ftc.gov/reports/federal-trade-commission-2006-identity-theft-survey-report-prepared-commission-synovate.

ThreatTrack Security. 2014. Malware analysts have the tools they need, but challenges remain. Retrieved from http://www.bankinfosecurity.com/whitepapers/malware-analysts-have-tools-they-need-but-challenges-remain-w-1026.

U.S. Census. 2012. 2012 Statistical Abstract of the United States.

Verizon Enterprise Solutions. 2015. *2015 Data Breach Investigations Report*. Technical Report. Retrieved from http://www.verizonenterprise.com/DBIR/2015/.

Received November 2016; revised April 2017; accepted July 2017