Quantifying Privacy Harm via Personal Identity Insurance

Daniel W. Woods

December 17, 2021
To be presented at Computers, Privacy and Data Protection conference (CPDP'22)

Abstract

Empirical estimates of privacy harm can help victims to demonstrate damages resulting from violations or support organisations in balancing harm to individuals against the cost of preventative measures. Quantitative studies of privacy harm are relatively rare. Personal identity insurance provides an additional source of quantitative data regarding the nature, likelihood and impact. We extract 34 personal identity insurance products that were uniquely filed with regulators in the US. We conduct a content analysis on the policy wordings and actuarial tables. Analysing the policy wordings reveals that personal identity theft causes a number of costs in terms of monitoring credit records, lost income and travel expenses, attorney fees, and even mental health counselling. Our analysis shows there are few exclusions related to moral hazard, which suggests that identity theft is largely outside the control of individuals. The actuarial calculations reveal financial impacts ranging from a few hundred to a few thousand dollars. Finally, insurers provide support services that are believed to reduce out of pocket expenses by over 90%. Together these policies, which are tested by market forces, provide three main insights: (i) there are real, quantifiable harms resulting from identity theft; (ii) individuals can do little to stop it; and (iii) a lack of support services increases losses.

1 Introduction

Privacy violations are often left unaddressed because of the difficulties in demonstrating that individuals suffered privacy harms (Calo, 2014; Citron and Solove, 2022). This problem dates back to the early days of the Internet (Lipton, 2010, p. 508). In cases like data breaches where the privacy violation was preventable, the reluctance to recognise harms leads to "under-deterrence" (Solove and Citron, 2017). Thus, empirical studies of privacy harm not only support victims of privacy violations in seeking ex-post remediation, but the resulting increase in likelihood and size of damages creates disincentives for privacy violations.

The General Data Protection Regulation would also be more effective if more empirical data on privacy harms was available. For example, Article 32 asks organisations to weigh the costs of privacy enhancing measures against the risks to data subjects (Voigt and Von dem Bussche, 2017). Organisations can reliably quantify the cost of implementation but struggle to quantify the likelihood or magnitude of harm to individuals due to data availability problems (Selzer et al., 2021). Privacy laws and legal actions beyond the EU and the US are also likely to draw on empirical studies of harm in order to either demonstrate damages, quantify the risk to data subjects or guide policy formulation.

Briefly surveying the information available demonstrates the need for new approaches. Woods and Böhme (2021b) survey an emerging literature on security harms to organisations. The most well-studied privacy-specific topics are stock market reactions (Acquisti et al., 2006; Gay, 2017), data breaches (Edwards et al., 2016; Eling and Loperfido, 2017), and legal actions (Romanosky et al., 2014; Ceross and Simpson, 2017; Wolff and Atallah, 2021). Romanosky (2016) even quantifies the mean cost to organisations who commit a privacy violation, which is in the millions of dollars. Such data sources cannot reliably quantify harm to individuals.

The literature on individual-level harm is more dispersed. Some cybercrime survey questions have privacy dimensions like those about criminals gaining access to online banking and shopping accounts (Anderson et al., 2013; Riek and Böhme, 2018), although many of these are primarily security problems. Studies quantifying online abuse are also proliferating (Thomas et al., 2021), although again it is hard to isolate the privacy-specific harm. Other aspects of privacy harm lack national and/or individual-level statistics, such as those associated with technology-enabled intimate partner violence (Slupska and Tanczer, 2021, p. 664). More generally, we could not identify a survey of empirical approaches to quantifying privacy harm. This is characteristic of an emerging field to which new approaches should be welcomed.

To collect novel empirical evidence, this paper turns to the insurance industry. Insurers explicitly define harms in the insurance contract and quantify their likelihood and impact via the actuarial process (Thoyts, 2010). Thus, privacy insurance contracts reveal qualitative information about harms, while pricing algorithms provide quantitative insights. We focus on one specific type of privacy insurance, namely personal identity insurance, with the goal of answering three research questions:

RQ1: Which harms are covered by personal identity insurance?

RQ2: What is the implied likelihood and severity of each harm?

RQ3: How do insurers justify the scope and pricing of coverage?

Beyond providing a new data source for quantifying privacy harm, our study represents the first empirical analysis of privacy insurance for individuals. The insights could help individuals to manage privacy risk by evaluating the effectiveness of transferring the consequences to an insurer. Individuals may be

further supported by the risk-reduction services that are often provided alongside insurance (Thoyts, 2010). Thus, one could consider privacy insurance as a form of privacy enhancing technology, notably a financial product that diverges considerably from the usual technical approach (Heurix et al., 2015). The study also contributes to an emerging field of technology insurance that covers cyber attacks (Romanosky et al., 2019), crypto-assets (Zuckerman, 2021), cyber bullying (Kshetri and Voas, 2019) and artificial intelligence liability (Lior, 2022).

Section 2 describes how we collect and analyse the empirical data. Section 3 presents the results. Section 4 discusses how these relate to privacy law, theory and practice. Section 5 offers a conclusion.

2 Methods

We adopt the high-level approach used by Romanosky et al. (2019) to understand corporate cyber insurance coverage. This involves sampling insurance regulatory filings from the SERFF database until saturation is reached in terms of coverage (Campbell et al., 2020). Coverage themes are identified via an inductive content analysis (Elo and Kyngäs, 2008). We also map quantitative risk estimates to themes.

Sampling We searched each state's filing system using the keyword "identity" and provided no further limitations on the search because we found identity insurance filed under lines including commercial crime and homeowner lines. Following Romanosky et al. (2019), we only collected approved filings. We focused on the four largest states (California, Texas, Florida, and New York) as the greater market size provides more potential for thematic variation.

This resulted in 86 regulatory filings with meta-data including: state, submission date, companies, product name, and insurance line. We grouped filings to ensure each unit of analysis contained the policy wording, rating manual, and rating justification.[1] This resulted in 34 unique personal identity insurance filings. We did not double count when multiple insurance companies (often subsidiaries) filed together and did not count updated wordings as distinct insurance products, although we did track these changes. We stopped collecting policies when we stopped deriving new coverage themes (Campbell et al., 2020).

Analysis We analysed the policy wordings for **RQ1**. We first read the document to identify high-level questions like who the policy was for and whether a help line was offered. We then extracted the sections describing what was covered and under which circumstances. These consisted of a list of contractual terms, and extracted each item as a unit of analysis.

We then mapped each unit of analysis to a theme. Themes had to be derived inductively due to the lack of prior research (Elo and Kyngäs, 2008). We created a theme for each unit that could not be classified under an existing theme. After

[1] Some companies filed these components in separately

analysing 10 policies, we consolidated themes to ensure they were comprehensive and mutually exclusive (Stemler, 2000) and used the resulting code-book for the entire analysis. Figure 1 highlights how we quickly reached saturation in coverage, but required more policies to do so for exclusions.

To answer **RQ2**, we extracted all quantitative risk estimates from the rate schedules. Due to the simplicity of the pricing schemes, estimates can be classified into the following categories: likelihood and severity of the harm, pure premium (risk = likelihood $\times$ severity), and market premium that includes the insurer's expenses and profit. Each estimate was then mapped to a coverage theme to provide more fine-grained harm estimates.

To understand how coverage and pricing were derived (**RQ3**), we counted the number of product filings that reference each information source. We also included selective quotes from insurer's justifications for illustrative purposes.

3 Results

Section 3.1 describes what is is covered and excluded by personal identity insurance. Section 3.2 identifies quantitative estimates and justifications.

3.1 Coverage and Exclusions

Our inductive analysis identified nine specific categories of coverage and classified the remaining 14 coverage items into a miscellaneous category, which is summarised in Table 1. The core coverage consists of different costs associated with correcting official records related to the policyholder's identity. The costs of credit services (Theme #1) like reports or monitoring was mostly covered by the policies, with those offered in the early years limiting the number of reports. Almost all policies indemnify the cost of re-filing loan applications (Theme #2) and communications costs (Theme #3) like long distance phone calls or notarising documents incurred to "amend or rectify records as to your true name or identity". The costs of travelling to do so (Theme #4) was occasionally included. The time required to do is commonly indemnified as lost income (Theme #5) and/or alternative care arrangements (Theme #6). Another common cost was attorney fees and court costs (Theme #7) resulting from the defense of a civil suit, civil judgement or criminal charges brought against the policyholder.

Displaying the policies longitudinally captures how identity insurance expanded coverage over time. For example, mental health counselling (Theme #9) did not appear until 2014, after which it was included in the majority of policies. Policies also began to include clauses offering to cover all reasonable costs "to recover control over his or her personal identity" (Theme #10), although this clause usually explicitly excludes lost or stolen money. The only area of coverage retraction is the cost of hiring professionals to help investigate and manage personal identity thefts (Theme #8), which were only included in the early years.

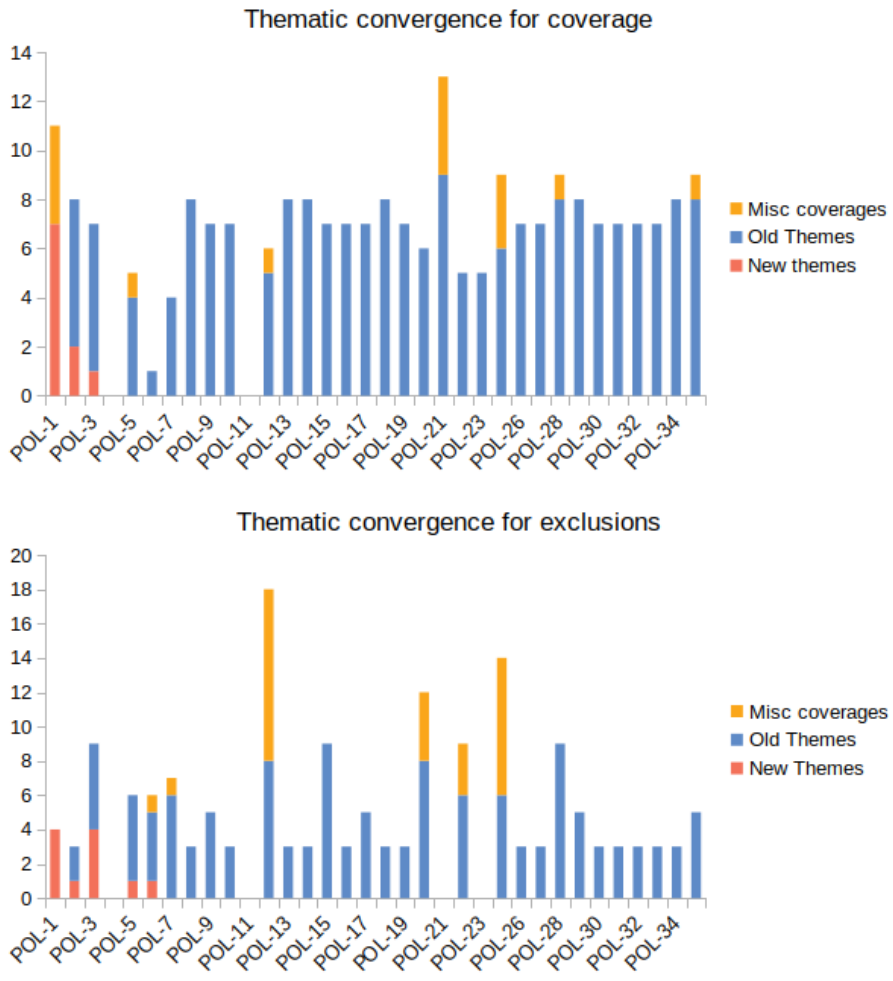

Figure 1: The content analysis converged faster and more reliably for coverage than for exclusions, in part because some policies including long lists of seemingly irrelevant exclusions.

Date	POL	Credit services	Application costs	Communication costs	Travel costs	Lost income	Care expenses	Attorney fees	Professional services	Counselling	Reasonable costs	Miscellaneous
11/07/05	5	6	✓	✓				✓				
06/21/06	7	12	✓	✓				✓				
03/26/07	6								✓			
01/08/08	20		✓	✓		✓	✓	✓	✓			
05/13/08	1	4	✓	✓		✓	✓	✓				4
08/24/08	21	4	✓	✓		✓	✓	✓	✓			4
04/20/10	29	✓	✓	✓		✓	✓	✓	✓		✓	
03/10/11	31	✓	✓	✓		✓	✓	✓			✓	
07/11/11	22	✓	✓	✓		✓		✓				
02/12/13	32	✓	✓	✓		✓	✓	✓			✓	
03/13/14	27	✓	✓	✓		✓	✓	✓			✓	
05/01/14	25	✓	✓	✓		✓	✓					3
05/16/14	14	✓	✓	✓		✓	✓	✓		✓	✓	
05/29/14	2	✓	✓	✓		✓	✓	✓		✓	✓	
07/01/14	26	✓	✓	✓		✓	✓	✓			✓	
09/24/14	35	✓	✓	✓	✓	✓	✓	✓				1
02/26/15	13	✓	✓	✓		✓	✓	✓		✓	✓	
03/06/15	8	✓	✓	✓		✓	✓	✓		✓	✓	
04/04/15	18	✓	✓	✓		✓	✓	✓		✓	✓	
06/30/15	34	✓	✓	✓		✓	✓	✓		✓	✓	
08/07/15	16	✓	✓	✓		✓	✓			✓	✓	
08/07/15	19	✓	✓	✓		✓	✓			✓	✓	
08/27/15	30	✓	✓	✓		✓	✓			✓	✓	
09/15/15	12		✓	✓		✓		✓				1
12/30/15	10	✓	✓	✓		✓	✓			✓	✓	
12/31/15	3		✓	✓	✓	✓	✓	✓		✓		
01/08/16	15		✓	✓	✓	✓	✓	✓		✓		
01/19/16	28		✓	✓	✓	✓	✓	✓		✓		1
09/09/16	33	✓	✓	✓		✓	✓	✓			✓	
09/15/16	23		✓	✓	✓	✓		✓				
02/03/20	9	12	✓	✓		✓	✓	✓			✓	
02/03/20	17	12	✓	✓		✓	✓	✓			✓	

Table 1: The coverage offered by each policy ordered by date of filing. Integers denote the maximum number of credit reports in the credit services column and the number of coverage items in the miscellaneous column.

It is worth unpacking the coverage items classified as miscellaneous. POL-1 and 21 were introduced by the same insurance company in different states and they included coverage for: liabilities resulting from fraudulent transactions using existing accounts or accounts opened in the policyholder's name, any costs "incurred by a financial institution or credit issuer", and the deductible payment for any other personal identity insurance. POL-12 and 25 included a clause covering "credit freeze, credit thaw costs, transcript costs, appeal bond, court filing fees, expert witness or courier fees". POL-25 also covered the costs of replacing "identification cards" and "ordering medical records" (as did POL-28), although both of these items likely overlap with the communication costs theme. Finally, POL-35 explicitly included "costs approved by us, for providing periodic reports on changes to, and inquiries about the information contained in the insured's credit reports or public databases (including, but not limited to credit monitoring services);", which is likely to mainly consist credit services.

Turning to the exclusions, Table 2 displays the exclusions discovered in the sample. All but one of the policies exclude losses due to business identity theft, which shows the policy is intended to cover losses suffered by individuals. Most policies include reporting requirements, such as filing a police report or notifying within 30-120 days. Many of the exclusions would be included in other insurance policies, such as not covering losses when the policyholder had prior knowledge of the loss or when the loss is incorrectly reported. The fraud exclusion denies coverage for events committed by the insured or an acquaintance with the insured's knowledge, but a handful of policies also excluded losses committed by close acquaintances without the insured's knowledge, which we term *insider threat*.

Some of the exclusions are unlikely to cause or constitute personal identity harms. For example, the conflict/political column includes examples like excluding losses due to war and political actions, the disaster column includes both natural and nuclear incidents, and bodily injury covers physical harm to a person. Neither war, nuclear accidents or bodily harm seem relevant to personal identity theft. The miscellaneous exclusions are similarly tenuous, such as "loss from games of chance" (POL-25) and "loss of valuable papers, valuable documents, jewellery, silverware and other personal property..." (POL-12). Corporate cyber insurance policies have been shown to be similarly profligate in the excluded events (Woods and Weinkle, 2020).

Insurers tend to exclude activities that increases risk, known as moral hazard (Baker, 1996). In addition to not lying (Fraud theme) and reporting swiftly and to the police (Reporting theme), the *computer security* theme captures such exclusions. This most commonly covered voluntary disclosure, which POL-3 defined as "disclosure of any code or other security information that can be used to gain access to any of your accounts...this exclusion will not apply if such disclosure was made when you were under duress or the victim of fraud". Thus, the most salient moral hazard is that a policyholder willingly discloses information. Notably, only one of the policies (POL-7) from 2006 required the insured to maintain security software:

> "it is the responsibility of each "identity recovery insured" to use and maintain his or her computer system security, including personal firewalls, anti-virus software, and proper disposal of used hard drives"

One interpretation is that insurers learned that personal identity harm was rarely caused by the insured not following information security procedures.

3.2 Pricing and Justifications

Table 3 displays our data about pricing and actuarial justifications. Notably, there is more missing data than in the previous section. Many of the filings missed actuarial justifications and some did not even report the premium. The study of corporate cyber insurance also found that policy wordings were more consistently included than pricing and actuarial data (Romanosky et al., 2019).

The first column describes the annual price of personal identity insurance per insured entity, which ranges by orders of magnitude from 0.25$ to over 100$. This variance is not well explained by the amount of coverage, described in the next two columns displaying the associated limit (maximum insurance pay-out) and deductible (the first part of loss paid by the policyholder). Sometimes this was because the policy contained more coverage. For example, some of the higher prices result from bundling personal identity insurance with "$50,000 of Named Malware, and $5,000 of Public Relations Services" (e.g. POL-2, 14, and 26). Some of the lowest priced policies (e.g. POL-12 and 25) were intended to be sold in bulk (the *bulk discount* column) so that one organisation purchases insurance for multiple individuals. The possibility that organisations purchase personal identity insurance on behalf of individuals explains the *risk rated* column, which contains a tick if different rates apply based on the insured's characteristics (e.g. the organisation's industry).

The likelihood and impact column are purely based on actuarial expectations, unlike the premium that also reflects the insurer's business model, such as expense costs or investment income (Thoyts, 2010). The estimates of frequency were more variable than the estimates of the impact. The lower frequency estimates resulted from normalising the number of data fraud cases reported to the FBI by the US population, whereas the higher values (e.g. 3.7%) came from normalising the number of data fraud cases by the sample size of an FTC survey. Such disparities may result from the difficulties surveying rare and emotionally salient phenomena documented by Florêncio and Herley (2013).

Some policies even delimit the frequency and impact estimate for coverage themes identified in the previous sub-section. For example, POL-3 references data obtained from their reinsurer to estimate the frequency of: replacement of documents (0.05%); travel expenses (0.035%); loss of income (0.035%); child and elderly care (0.011%); reimbursement of fraudulent withdrawals (0.0250%); legal costs (0.03%); remediation service costs (0.05%), and case management service costs (0.075%). We advise that the relative frequencies are perhaps the main takeaway. For example, the child and elderly care costs are incurred less

date	POL	Business identity	Bodily injury	Conflict/political	Fraud	Prior knowledge	Reporting	Disaster	Non-identity	Insider threat	Computer security	Miscellaneous
11/07/05	5	✓			✓	✓	✓		✓	✓		
06/21/06	7	✓			✓	✓	✓		✓		✓	
03/26/07	6	✓			✓	✓	✓					1
01/08/08	20	✓	✓		✓	✓			✓	✓	✓	4
05/13/08	1	✓	✓	✓	✓							
04/20/10	29	✓		✓	✓		✓	✓				
03/10/11	31	✓			✓		✓					
07/11/11	22	✓			✓	✓			✓	✓		3
02/12/13	32	✓			✓		✓					
03/13/14	27	✓			✓		✓					
05/01/14	25	✓	✓		✓	✓		✓	✓		✓	8
05/16/14	14	✓			✓		✓					
05/29/14	2	✓			✓		✓					
07/01/14	26	✓			✓		✓					
09/24/14	35		✓	✓	✓		✓	✓				
02/26/15	13	✓			✓		✓					
03/06/15	8	✓			✓		✓					
04/04/15	18	✓			✓		✓					
06/30/15	34	✓			✓		✓					
08/07/15	16	✓			✓		✓					
08/07/15	19	✓			✓		✓					
08/27/15	30	✓			✓		✓					
09/15/15	12	✓	✓		✓	✓		✓	✓		✓	10
12/30/15	10	✓			✓		✓					
12/31/15	3	✓	✓	✓	✓	✓	✓	✓		✓	✓	
01/08/16	15	✓	✓	✓	✓	✓	✓	✓		✓	✓	
01/19/16	28	✓	✓	✓	✓	✓	✓	✓		✓	✓	
09/09/16	33	✓			✓		✓					
02/03/20	9	✓		✓			✓	✓			✓	
02/03/20	17	✓		✓			✓	✓			✓	

Table 2: The exclusions included in each policy ordered by date of filing. The final column displays the number of coverage items classified as miscellaneous.

date	POL	Premium ($)	Limit ($)	Deductible ($)	Risk rated	Bulk discount	Frequency	Impact ($)
11/07/05	5		15000					
06/21/06	7	100					1%	3000
03/26/07	6							
01/08/08	20	126.25						
05/13/08	1	60	15000				2%	1369
08/24/08	21	126	20000					422
09/30/09	4	15	10000			✓		
04/20/10	29							
03/10/11	31	19	25000	100				
07/11/11	22							
08/24/11	11							
02/12/13	32	20	15000	250				
03/13/14	27	28	15000				0.05%	1603
05/01/14	25	1.08	10000			✓		
05/16/14	14	81-299*	50000	2500	✓			
05/29/14	2	81-299*	50000	2500	✓			
07/01/14	26	81-299*	50000	2500	✓			
09/24/14	35							
02/26/15	13							
03/06/15	8	10	15000					
04/04/15	18	10	15000	100				
06/30/15	34	10	15000	100			0.01%	3015
08/07/15	16	10	15000	100			3.70%	1200
08/07/15	19	10	15000	100				
08/27/15	30	10	15000	100				
09/15/15	12	0.24	25000		✓	✓		
12/30/15	10	10	15000	100			3.70%	1200
12/31/15	3	1.54	25000				0.05%	1603
01/08/16	15							
01/19/16	28	2.93	25000					
09/09/16	33	16						
09/15/16	23	2.44	1000000				0.05%	3541
02/03/20	9	15	25000		✓	✓		
02/03/20	17	15	25000		✓	✓	3.81%	365

Table 3: Pricing and actuarial information available for each regulatory filing. Empty fields should not be interpreted as anything other than missing data. * = price for a bundle including additional coverage.

frequently than those to hire response services.

To provide a flavour of the actuarial reasoning, we quote the following from POL-10 extract in full:

> "According to a recent study commissioned by the Federal Trade Commission, 90% of "All ID Theft" out of pocket expenses are $1,200 or less. While we do not have significant experience with this coverage, we believe that the availability of case management restoration services will reduce this severity to approximately $81. The same FTC-commissioned report suggests a frequency of 3.7%. Thus, our loss content is expected to be approximately $3.00. Loss-related expenses (toll-free help-line and case management service) are expected to be $3.50. Thus our total loss cost is $6.50."

The most notable aspect is that case management services reduce out of pocket expenses by over 90%. Other data sources for actuarial justifications include: the Bureau of Labour Statistics, Ponemon group, Javelin's surveys, competitor analysis and the FBI.

4 Discussion

This section evaluates the results in light of our goal of identifying and quantifying privacy harms with a view to litigation. The existence of personal identity insurance suggests individuals anticipate privacy harms that are not sufficiently remedied by the legal system. The following, which was included in multiple insurer's filings, summarises the gap:

> "While many financial institutions provide protections to consumers for the actual fraud loss, most individuals have no help for the time and expense required to restore their personal identities."

The impact column of Table 3 suggests actuaries estimate the associated time and expenses to be around $3000. This is not insignificant given that multiple insurers estimate the likelihood to be more than 3%.

Interestingly, POL-10 believed post-theft services paid by the insurer could reduce such expenses by over 90%. This mirrors corporate cyber insurance in which policies pay for a team of consultants spanning law, IT and public relations to respond to cyber incidents (Franke, 2017; Woods and Böhme, 2021a). More generally, scholars have observed insurers positively influencing risk management practices of insureds across a range of insurance lines, known as *insurance as governance* (Ericson et al., 2003; Ben-Shahar and Logue, 2012).

A provocative question to ask is whether governments could do more to help individuals recover from identity theft, after all many thefts exploit state provided identifiers like social security numbers that cannot be easily replaced due to the government's architectural design choices. The bulk discounts in some policies suggests that these costs display considerable economies of scale.

The equivalent post-incident services are provided publicly for fire, and were originally provided by insurers (Carlson, 2005).

In terms of the identifying new harms, the costs covered in Table 1 are driven by the complexity of bureaucracies—re-filing applications that were rejected due to identity theft, the cost of notarizing documents, lost income or additional care expenses due to the time invested—that individuals are normally expected to swallow. A different kind of cost is mental health counselling, which was not offered until 2014 after which it was included in the majority of policies. Its inclusion suggests the insurance industry recognises the psychological harm of victims of identity theft. It seems reasonable that anticipation of this psychological damage in addition to the $3000 impact following a data breach might lead to anxiety, as argued by Solove and Citron (2017).

The actuarial estimates confirm that the impact of identity theft is relatively low but also relatively common. This diffuseness of harm was identified as a reason why courts dismiss data breach lawsuits. The source of quantitative estimates is interesting in that actuarial justifications relied on public data collection (e.g. FTC surveys or FBI crime reports). One might ask whether governments collecting and releasing similar aggregate data for other privacy harms could bootstrap private insurance markets. Or perhaps academics could reflect on what would be required for their surveys to be used for the same purpose.

More generally, our search was relatively narrow in that we used a small number of search terms. Future work could explore other lines of insurance related to privacy harms. It could also expand our analysis beyond the four largest state. We suspect the results will be similar as we detected few differences across states in terms of the content of policies or actuarial estimates, although the regulatory reports did differ.

5 Conclusion

The following extract, which was included word-for-word in multiple regulatory filings, provides a concise summary of our study:

> "While there are ways to reduce one's exposure to identity theft, it is a crime that can strike anyone. Those who are victims of this crime need to make identity recovery a top priority, because otherwise:
>
> - Credit rating can be ruined
> - Arrest warrants can be issued against the victim
> - Liens can be applied against the victim's assets
>
> While many financial institutions provide protections to consumers for the actual fraud loss, most individuals have no help for the time and expense required to restore their personal identities."

While the extract suggests there are "ways" of reducing exposure, Table 2 shows insurers do not push policyholders towards implementing them. One

explanation is that identity theft risk reduction is too ineffective or too onerous to ask of policyholders. This supports a narrative in which consumers are powerless to prevent privacy harms resulting from personal identity theft. The corresponding insurance coverage reflects a need for ex-post response solutions to both reduce privacy harms and also indemnify the financial cost.

Our study confirms one aspect of the privacy harm literature—legal systems fail to recognise and remedy privacy harms (Citron and Solove, 2022)—as evidenced by the emergence of a private market covering the harms associated with identity theft incidents. We contribute an additional contribution, namely that the lack of support services leads individuals to suffer more harm. For example, one insurer anticipates case management services lead to a 90% reduction in the cost of an identity theft incident. Thus policy makers could reflect on whether the impacts of identity theft and the expertise to remedy are fairly distributed across society. The status-quo in which financial smoothing and risk reduction services are privately provided undoubtedly skews towards affluent consumers.

References

Acquisti, A., A. Friedman, and R. Telang (2006). Is there a cost to privacy breaches? an event study. *ICIS 2006 Proceedings*, 94.

Anderson, R., C. Barton, R. Böhme, R. Clayton, M. J. Van Eeten, M. Levi, T. Moore, and S. Savage (2013). Measuring the cost of cybercrime. In *The Economics of Information Security and Privacy*, pp. 265–300. Springer.

Baker, T. (1996). On the genealogy of moral hazard. *Texas Law Review 75*(2), 237.

Ben-Shahar, O. and K. D. Logue (2012). Outsourcing regulation: how insurance reduces moral hazard. *Michigan Law Review 111*, 197.

Calo, R. (2014). Privacy harm exceptionalism. *Colo. Tech. LJ 12*, 361.

Campbell, S., M. Greenwood, S. Prior, T. Shearer, K. Walkem, S. Young, D. Bywaters, and K. Walker (2020). Purposive sampling: complex or simple? research case examples. *Journal of Research in Nursing 25*(8), 652–661.

Carlson, J. A. (2005). The economics of fire protection: From the great fire of london to rural/metro 1. *Economic Affairs 25*(3), 39–44.

Ceross, A. and A. Simpson (2017). The use of data protection regulatory actions as a data source for privacy economics. In *International Conference on Computer Safety, Reliability, and Security*, pp. 350–360. Springer.

Citron, D. K. and D. J. Solove (2022). Privacy harms. *Boston University Law Review 102*.

Edwards, B., S. Hofmeyr, and S. Forrest (2016). Hype and heavy tails: A closer look at data breaches. *Journal of Cybersecurity 2*(1), 3–14.

Eling, M. and N. Loperfido (2017). Data breaches: Goodness of fit, pricing, and risk measurement. *Insurance: Mathematics and Economics 75*, 126–136.

Elo, S. and H. Kyngäs (2008). The qualitative content analysis process. *Journal of Advanced Nursing 62*(1), 107–115.

Ericson, R. V., A. Doyle, and D. Barry (2003). *Insurance as governance*. University of Toronto Press.

Florêncio, D. and C. Herley (2013). Sex, lies and cyber-crime surveys. In *Economics of Information Security and Privacy III*, pp. 35–53. Springer.

Franke, U. (2017). The cyber insurance market in Sweden. *Computers & Security 68*, 130–144.

Gay, S. (2017). Strategic news bundling and privacy breach disclosures. *Journal of Cybersecurity 3*(2), 91–108.

Heurix, J., P. Zimmermann, T. Neubauer, and S. Fenz (2015). A taxonomy for privacy enhancing technologies. *Computers & Security 53*, 1–17.

Kshetri, N. and J. Voas (2019). Thoughts on cyberbullying. *Computer 52*(4), 64–68.

Lior, A. (2022). Insuring AI: The role of insurance in artificial intelligence regulation. *Harvard Journal of Law and Technology* (1), in print.

Lipton, J. D. (2010). Mapping online privacy. *Nw. UL Rev. 104*, 477.

Riek, M. and R. Böhme (2018). The costs of consumer-facing cybercrime: An empirical exploration of measurement issues and estimates. *Journal of Cybersecurity 4*(1), tyy004.

Romanosky, S. (2016). Examining the costs and causes of cyber incidents. *Journal of Cybersecurity 2*(2), 121–135.

Romanosky, S., D. Hoffman, and A. Acquisti (2014). Empirical analysis of data breach litigation. *Journal of Empirical Legal Studies 11*(1), 74–104.

Romanosky, S., A. Kuehn, L. Ablon, and T. Jones (2019). Content analysis of cyber insurance policies: how do carriers price cyber risk? *Journal of Cybersecurity 5*(1).

Selzer, A., D. W. Woods, and R. Böhme (2021). Appropriateness under article 32 GDPR. *The European Data Protection Law Review 7*(3), 456–470.

Slupska, J. and L. M. Tanczer (2021). Threat modeling intimate partner violence: Tech abuse as a cybersecurity challenge in the internet of things. In *The Emerald International Handbook of Technology Facilitated Violence and Abuse*. Emerald Publishing Limited.

Solove, D. J. and D. K. Citron (2017). Risk and anxiety: A theory of data-breach harms. *Texas Law Review 96*, 737.

Stemler, S. (2000). An overview of content analysis. *Practical Assessment, Research, and Evaluation 7*(1), 17.

Thomas, K., D. Akhawe, M. Bailey, D. Boneh, E. Bursztein, S. Consolvo, N. Dell, Z. Durumeric, P. G. Kelley, D. Kumar, et al. (2021, May). SoK: Hate, harassment, and the changing landscape of online abuse. In *IEEE Symposium on Security and Privacy*, Oakland, CA, pp. 247–267.

Thoyts, R. (2010). *Insurance theory and practice*. Routledge.

Voigt, P. and A. Von dem Bussche (2017). The EU general data protection regulation (GDPR). *A Practical Guide, 1st Ed. 10*, 3152676.

Wolff, J. and N. Atallah (2021). Early gdpr penalties: Analysis of implementation and fines through may 2020. *Journal of Information Policy 11*, 63–103.

Woods, D. W. and R. Böhme (2021a). How cyber insurance shapes incident response: A mixed methods study. In *Workshop on the Economics of Information Security*.

Woods, D. W. and R. Böhme (2021b, May). SoK: Quantifying cyber risk. In *IEEE Symposium on Security and Privacy*, Oakland, CA, pp. 909–926.

Woods, D. W. and J. Weinkle (2020). Insurance definitions of cyber war. *The Geneva Papers on Risk and Insurance-Issues and Practice 45*, 639—-656.

Zuckerman, A. (2021). Insuring crypto: The birth of digital asset insurance. *U. Ill. JL Tech. & Pol'y*, 75.