Accepted for *12th Workshop on the Economics of Information Security*
version 2.0

How Bad Is It? – A Branching Activity Model to Estimate the Impact of Information Security Breaches

Russell Cameron Thomas
Department of Computational Social Science, George Mason University, Fairfax, VA 22030, russell.thomas@meritology.com

Marcin Antkiewicz
Qualys, Inc., Madison, Wi, 53704, mantkiewicz@qualys.com

Patrick Florer
Risk Centric Security, Inc., Dallas, Texas 75230, patrick@riskcentricsecurity.com

Suzanne Widup
Verizon RISK Team, Verizon Communications Inc., San Francisco, CA 94105, suzanne.widup@verizon.com

Matthew Woodyard
Zions Bancorporation, Salt Lake City, UT 84133, mwoodyard@5lbs.org

This paper proposes an analysis framework and model for estimating the impact of information security breach episodes. Previous methods either lack empirical grounding or are not sufficiently rigorous, general or flexible. There has also been no consistent model that serves theoretical and empirical research, and also professional practice. The proposed framework adopts an *ex ante* decision frame consistent with rational economic decision-making, and measures breach consequences via the anticipated costs of recovery and restoration by all affected stakeholders. The proposed branching activity model is an event tree whose structure and branching conditions can be estimated using probabilistic inference from evidence – 'Indicators of Impact'. This approach can facilitate reliable model estimation when evidence is imperfect, incomplete, ambiguous, or contradictory. The proposed method should be especially useful for modeling consequences that extend beyond the breached organization, including cascading consequences in critical infrastructures. Monte Carlo methods can be used to estimate the distribution of aggregate measures of impact such as total cost. Non-economic aggregate measures of impact can also be estimated.

Key words: information security, breach severity, loss estimation, cascading consequences, stochastic model
History: submitted 3/11/2013, corrections: v1.1–1.4, this version: v2.0 May 21, 2013

---

1. Introduction

Every risk assessment method for information security includes some scale or measure to evaluate the negative consequences[1] of possible breach episodes. Most common are simple three-level

[1] Many different terms are used to describe the aggregation of negative consequences, including "severity", "impact", "cost", "loss", and others. In this paper, we will adopt the term "impact" to mean the aggregation of negative consequences according to some measure.

ordinal scales, e.g. "High", "Medium", and "Low". Probabilistic risk analysis involves quantitative estimation of breach impact on a ratio scale, usually as a dollar cost. Whether the estimates are expressed as ordinal or ratio scale values, impact estimation plays a central role in formal risk assessment and risk management. Even in informal settings, perception and communication about risk often involves some estimate or mental model of the severity of breach episodes.

If the impacts of breach episodes fit Gaussian distributions, then it would be progressively easier to estimate them with experience and information pooling, and the consequences of estimation errors would be moderate and manageable. However, information security breach impacts, especially in interdependent settings like critical infrastructures, appear to follow 'heavy tailed' distributions due to cascading consequences (Woolf et al. 2004, Buldyrev et al. 2010, Dobson et al. 2006, Wierzbicki and Dobson 2006). Thus the likelihood of extreme impacts cannot be ignored as in the Gaussian case. Instead, the very existence of low probability/high severity loss scenarios often dominates risk management decisions.

Advancing research in this area is important to the scientific and policy communities involved in information security (National Science and Technology Council 2011). There is a significant gap in methods and frameworks between theoretical and empirical researchers, and also a gap between academic research and professional practice. For example, theoretical models of security investment (Gordon and Loeb 2002) and interdependent risk (Kunreuther and Heal 2003, Heal and Kunreuther 2004, 2007, Grossklags et al. 2008, Böhme and Moore 2010) treat "loss" as a single real-valued parameter for each decision-maker. While such simplistic treatments are useful to make the analysis tractable, they essentially assume-away the most important aspect of the problem: how to estimate breach impact and likelihood in the first place. Furthermore, there is no empirical grounding for the assumption that decision-makers have well defined estimates of impact and likelihood. Quite the contrary. There is empirical evidence that decision-makers and stakeholders facing complex interdependent risk have poor estimates of risk (Barker and Haimes 2009, Carreras et al. 2007, Clement 1989, Grebe 2013, Guarro 1989, Hansson 1999, Kunreuther and Heal 2012, Visschers et al. 2009).

In summary, without some reliable and robust breach impact estimation methods, quantified information security will continue to be a "weak hypothesis" (Verendel 2009).

1.1. Why Breach Impact Estimation is Hard

On the surface, it might appear that estimating the the impact of a breach episode is easy – a straight-forward accounting task where one waits until all the breach-related costs have been incurred and then calculates their sum using the affected firms' accounting records. A slightly more sophisticated version might be called the *insurance claim* method, where any party that

has been harmed by a breach enumerates all their damages and losses using the equivalent of an insurance claim form and claim rules. However, these simple methods do not adequately address the complexities, uncertainties, and difficulties involved, as we now discuss.

- *Disincentives to Disclose* – There have been many studies (Bodeau 1992, Cavusoglu et al. 2004, Lambert 1993, Rees 2009, Rees and Kannan 2008, Sohail 2006, Wang et al. 2008) showing that companies and individuals are reluctant to disclose information related to security breaches. The reasons are many and they go beyond considerations of subjective utility or costs vs. benefits. Yet it will not be possible to achieve collective security outcomes without disclosure and sharing other information regarding security (ENISA and RAND Europe 2010).
- *Attribution of Consequences and Costs* – Except for the smallest and simplest breach episodes, it can be hard to attribute costs to the breach. Many costs are indirect or bundled in overhead. Furthermore, some important costs are not realized until long after the breach has been discovered.
- *Intangibles* – Even when costs can be properly attributed, there can be problems in valuation, especially for intangible assets like trade secrets or, more broadly, knowledge capital. So much of the value of intangibles depends on context and history that it is very hard to reliably estimate what value is lost when intangibles are impaired by a breach episode.
- *Mis-estimation of Costs* – A related problem is the tendency by some people to greatly over-estimate or under-estimate costs associated with breach episodes. Sometimes this is due to mistakes in analysis or inadequate methods, but other times it is due to wishful thinking – enlarging or diminishing the cost of a breach furthers that person's or that organization's narrow self interest. Some very large IT security companies are guilty of this practice (Thomas 2011c).
- *Incommensurate Consequences* – Many times, the negative consequences on different stakeholders cannot be easily measured using a single scale of cost or impact. For example, consumers might experience a violation of privacy or betrayal of trust, while a corporation might experience material losses in terms of lost customers. In the same setting, government agencies might be contending with civil unrest or social upheaval as a result of the breach episode.
- *Uncertainty, Absence of Information, and Ignorance* – Loss in value of IT assets or intellectual property can be very difficult to estimate because much of the information is missing or inaccessible, lost in history, not known within sufficient bounds of confidence, or there might be severe 'unknown-unknowns' that are outside the mental frames of the analysts (Hansson 1999, Acquisti and Grossklags 2005, Barker and Haimes 2009, Hassel et al. 2009, Johnson et al. 2010, van Asselt and Rotmans 2002).
- *Ambiguity* – Similarly, much of the information about the breached firm before, during, and after the breach episode can be ambiguous (Hansson 1999, Acquisti and Grossklags 2005). For example, imagine a corporation that sells off a business unit that was breached severely, but does

so a year afterward, making no mention of the breach. Should this be attributed to the harm or loss caused by the breach?

• *Near-misses, 'Technical Debt' and Other Counterfactuals* – There needs to be some way to account for breach impacts that haven't yet happened, but might. These are called 'counterfactuals' and they play a very important role in rational decision-making under uncertainty, especially in strategic settings. A classic example of this are are near-miss events (Clement 1989, Muermann and Oktem 2002, Phimister et al. 2003, Thomas 2009a, Vanderschaaf 1992, Wright and van der Schaaf 2004), which except for some small twist of fate could have turned into a major loss episode. It is vital that people and organizations recognize and learn from near-miss events. Another example is captured in the phrase 'technical debt' (Brown et al. 2010, Guo et al. 2011, Klinger et al. 2011), which refers to short cuts and other expedient actions taken by IT system designers prior to a release. These may not 'cost' them anything in the very short-run, but sooner or later the 'debt' must be paid off eventually, either through patching or through the cost of security breaches.

• *Plural Interests and Decision Frames* – Any attempt to estimate breach impact will have to contend with the many different interests, values, and decision frames of stakeholders (Hassel et al. 2009, Visschers et al. 2009). Some will emphasize fairness, others will emphasize responsibility, while others will emphasize rational cost-benefit thinking. This can also take the form of differing attitudes toward risk and ambiguity. These can be strongly shaped and constructed through social perceptions and interactions.

• *Bias* – Finally, many people and sources of information will have biases that color their estimates of loss or harm (Acquisti and Grossklags 2004), or beliefs about the root causes and people to hold responsible. It is often very hard to present 'ground truth' in a way that will mitigate bias.

1.2. Related Work

Gordon and Loeb (2002) present an equilibrium economic model for optimal investment in information security; essentially aiming to optimize the expected costs due to losses, balanced by the costs associated with security investments. However, their model does not include any process for estimating losses or how to relate loss levels, or probability of loss to security investments.

Bodin et al. (2005) applied the Analytic Hierarchy Process (AHP) to address investment decisions with limited/fixed budgets and how to make the case for budget increases. They build an ordinal scorecard system using a three-part taxonomy for security harm: "confidentiality", "integrity", and "availability". Beyond this they do not attempt to estimate breach impact.

Wong (2008), in a working paper, proposed a fixed formula for calculating cost of a breach by enumerating various cost categories. His aim is practical, but this approach does not adequately address the difficulties described in the previous section.

Ponemon Institute (2012) has published five annual reports of "Cost of a Data Breach" based on surveys. Their method is centered on a formula that uses number of records breached as the main cost driver. Thomas (2011a) presented criticisms of their approach and methods.

Several researchers have attempted to go around the difficulties mentioned in Section 1.1, above, by studying the stock market reaction to news of the breach episode. This research approach rests on a semi-strong version of the Efficient Market Hypothesis, which asserts that new public information is immediately incorporated into share prices in proper proportion to its effect on the market value of the company. Campbell et al. (2003) studied the economic effect of information security breaches on stock prices. Overall, they find a weak relationship between breaches and decline in stock prices. However, there is a very significant relationship between decline in stock price and breaches involving unauthorized access to confidential data. More broadly, impact of breaches are sensitive to context. Garg et al. (2003) also studied stock market reactions to various information security events, with mixed and somewhat confusing results. Cavusoglu et al. (2004) found significant changes in stock market value in the two days following public disclosure of the breach. Specifically, they found that breach cost was higher for Internet-only firms and that breaches are costlier for smaller firms than larger firms.

Society of Actuaries (2010) provides a tutorial introduction to the financial pricing of operational risk, which is a super-set of information security risk. Thomas (2009b) builds on this to propose a method for pricing risk to facilitate enterprise-level risk management and risk pooling. Estimation of impact losses is a key element of their proposed methods.

Anderson et al. (2012) and Hyman (2013) study the cost of cybercrime at a national and international level. They detail many of the measurement and estimation problems described above. They use estimation methods that use available information while avoiding gross mis-estimation. However, their studies do not attempt to estimate the impact of individual breach episodes.

There is a large literature on the role of near-miss and precursor events in disaster management, e.g. Alamgir et al. (2009), Muermann and Oktem (2002), Phimister et al. (2003), Vanderschaaf (1992), Wright and van der Schaaf (2004). In the context of risk management for information security, Thomas (2009a) presents an analysis of the cost of a near-miss breach event and makes the case that it should not be cost-free. Likewise, many authors have explored the challenges of uncertainty and ambiguity in risk management, e.g. Acquisti and Grossklags (2005), Barker and Haimes (2009), Hansson (1999), Johnson et al. (2010), Masys (2012), Grossklags et al. (2010), van Asselt and Rotmans (2002).

The annual Verizon Data Breach Investigation Report (DBIR) (Baker et al. 2012) is an example of statistical analysis of a large sample set of breach incidents. However, the scope of the DBIR has been limited to technical aspects of breaches and not on business response or business impact.

Verizion has published their framework for coding breach incident data – VERIS framework (Baker et al. 2010) – and it is supported by an open source community[2]. It has the potential for being extended to include Indicators of Impact and related impact data.

At a policy level, there is widespread consensus that measuring breach impact and, more generally, measuring information security and risk is a very high priority for research, and also that our current state of knowledge and practice are woefully inadequate (Cashell et al. 2004, National Science and Technology Council 2011, National Science Foundation 2012, Verendel 2009, Department of Homeland Security 2012, Energy Sector Control Systems Working Group (ESCSWG) 2011, ENISA 2012).

Regarding modeling methods, Buldyrev et al. (2010), Dobson et al. (2006), Wierzbicki and Dobson (2006) use stochastic branching models to estimate the distribution of cascades in electrical blackouts and and other disasters. They use variations on the Galton-Watson process involving hundreds or thousands of identical interconnected elements (e.g. power lines). Though instructive for these large scale settings, our method proposed in this paper takes a different approach. In another line of research, activity graphs (Bolton and Davies 2000), stochastic activation networks (Sanders and Meyer 2001, Vakili et al. 2011), and similar methods have been used to model stochastic activity in information systems, including for security research.

Finally, several authors have published ideas that are similar to those presented in this paper. Borg (2005) presented a qualitative analysis of economically complex attacks and breach episodes, and discussed branching processes for consequences among others. Thomas (2009b) presented a framework for pricing risk associated with information security, and it included suggestions regarding methods to estimate frequency and severity of major breach episodes.

1.3. Overview of the Paper

An analysis framework is proposed in Section 2. The aim of this framework is to encompass a broad range of breach episodes and applications. In Section 3 a branching process model is proposed to estimate the consequences and impact of breach episodes. Section 4 describes Indicators of Impact, including a discussion of information sources. The paper concludes with discussion of the results in Section 5, including limitations, generalizations, use cases, and directions for future research.

2. Analysis Framework

The goal of this analysis framework is to address the difficulties raised in Section 1.1, above, and to apply broadly to both academic research and professional practice. The following is a description of each element in the analysis framework along with brief justifications where necessary.

[2] `http://www.veriscommunity.net/doku.php`

2.1. Focus: Functional Harm at a Holistic Level

We choose to focus on negative consequences that cause harm by degrading a person's or organization's ability to function as a holistic system. This is in contrast to other approaches that take a more atomic approach and focus on harm to specific assets, especially IT assets. We believe that analyzing harm to functions or capabilities at a system level is more general and is less sensitive to the details of organization or information structures. By analogy, this method is like estimating severity of an infectious disease, which is assessed in a holistic way considering the whole person, as distinct from a survey of specific organs or tissue that may be infected. The holistic assessment gives a better estimate of severity of the infection, better guidance for modes of treatment, and better estimate of likelihood of recovery.

2.2. Scope: Breach Episodes

Our scope of analysis is a *breach episode*, which consists of a complete campaign by a defined set of threat actors, from first contact to final resolution. This scope is broader than other methods whose unit of analysis is a breach event for a specific IT asset (e.g. the compromise of a specific web server). We have chosen this more inclusive scope for several reasons. First, it more realistically includes the full consequences and the response/recovery activities of affected people or organizations. Second, this broad scope encompasses all of the cascading consequences that are directly attributable to the harm caused by threat actors, including consequences outside of the breached organization(s). Finally, this scope offers the possibility of consistency when building a database of information about breach episodes of widely differing sizes and types.

2.3. *Ex Ante* Decision Frame

We choose to estimate breach impact in the context of a forward-looking (*ex ante*) decision frame. In other words, all breach impacts will be estimated as though a person or organization were making value-based decisions *before* the breach episode occurred. Effectively, this values the breach impact according to the a probabilistic estimate of the *ensemble of outcomes* that would include one factual realization (i.e. what actually happens in a given breach episode) and also near-by counterfactual realizations (i.e. all the outcomes that *almost happened* in a given breach episode, including near-misses, lucky or unlucky breaks, excellent or sloppy execution, and so on). This decision frame choice may be controversial and thus we provide five justifications, as follows.

First, it puts the costs of a breach on an equal footing with the investments a person or firm might make to avoid or mitigate the breach. Another way of saying this is that it is consistent with the 'principle of marginal utility' – a core normative principle for rational economic decision-making[3] since the Marginal Revolution of the 1890s (Marshall 1898, Stigler 1950a,b, Samuelson

[3] The justification for this principle does not rest on broad capabilities for rationality or optimizing behavior usually associated with General Equilibrium theory. Instead, the simplest justification is that it avoids the '*post hoc* fallacy'

and Nordhaus 2009). According to the principle of marginal utility, all economic decisions should be made according to the incremental costs and benefits (looking forward), with no regard for costs that have already been incurred ('sunk costs').

Second, it allows for consistent cost estimation across long time periods, incorporating both the time value of money and financial risk preferences of decision-makers. In contrast, blending historical costs with projected costs is fraught with difficulties and is often not possible to do without inconsistency or logical error.

Third, it admits a consistent method for estimating hard-to-measure scenarios using 'Willingness To Pay' and 'Willingness To Accept' (Gafni 1991). A prime example of this is reputation damage. It may be hard for decision-maker to place a value on all the forms of reputation damage they might experience, but such valuation is made somewhat easier and more consistent if it is framed as "What are you willing to pay *ex ante* to recover from this reputation damage, assuming that it will occur?" These values can be estimated prospectively, by experiments, or retrospectively by observing what people and companies actually spend to restore their reputation in various settings.

Fourth, an *ex ante* decision frame is consistent with several proposed mechanisms to mitigate moral hazard associated with principal-agent relationships in financial industry risk management. Examples of these mechanisms include the 'pre-commitment approach' (Kupiec and O'Brien 1995, 1997), burden sharing (Goodhart and Schoenmaker 2006), and risk pooling (Zhao et al. 2009). We believe this is an important avenue of innovation regarding incentive instruments and institutions for information security.

Finally, an *ex ante* decision frame embraces near-miss events as having a non-zero cost (Thomas 2009a). This is a very important feature for incentive mechanisms aimed at promoting prudence and organization learning (Mahajan 2010, Alamgir et al. 2009, Borg 2005, Clement 1989, Muermann and Oktem 2002, Phimister et al. 2003, Vanderschaaf 1992, Wright and van der Schaaf 2004).

2.4. Measure of Impact: Recovery and Restoration Efforts

We propose that the best measure of breach impact is the combined efforts that affected stakeholders would be willing to spend to recover from the breach in all aspects. 'Effort' here means commitment of valuable resources, time, attention, etc. in some organized process, routine, or

of questionable cause and related logical fallacies. "*Post hoc ergo propter hoc*" is Latin for "after this, therefore because of this". The fallacy can be expressed in logical propositions: 1) A occurred, then (undesirable) B occurred. 2) Therefore, avoiding A will prevent B. `http://en.wikipedia.org/wiki/Post_hoc_ergo_propter_hoc`. More simply, we can never follow the last sentence in this witty advice:

"*Don't gamble; take all your savings and buy some good stock and hold it till it goes up, then sell it. If it don't go up, don't buy it.*" [attributed to Will Rogers]

project. This differs from approaches that attempt to measure impact in term of *loss in value*[4] such as reduced economic value of tangible or intangible assets, or non-monetary losses such as reputation damage or emotional distress. The loss in value approach is common in legal settings and corporate transactions and it also has been used in previous research on information security breach impact estimation (e.g. Guarro (1989), Rainer Jr et al. (1991)). This method has merits, especially when the primary issue is legal liability or responsibility, or political blame. However, it has significant shortcomings compared to the proposed method of *recovery and restoration effort*.

First, loss-in-value, especially non-monetary losses, often have no observable indicators or symptoms. This makes loss estimation very expensive and information intensive. In contrast, recovery and restoration efforts almost always have observable indicators since they require spending time, energy, resources, or attention that otherwise wouldn't be spent. Very often there are identifiable processes, projects, meetings, or communications associated with them.

Second, loss in value is very hard to apply to non-commercial organizations. In contrast, recovery and restoration efforts apply to any purposeful entity or organization, including individual people, families, community groups, non-profits, government agencies, military organizations, or any combination.

Third, loss in value does not incorporate the stakeholder's trade-off preferences in the context of breach recovery, whereas the recovery and restoration approach does. For example, consider a regional government agency that might value their emergency communication center at $1 million and their fleet of trucks at $500,000. Consider a cyber-physical emergency scenario where they lose both capabilities. They might be willing to spend more to restore and recover their truck fleet over the communication center in order to provide for citizen evacuation, transporting relief supplies, etc. The emergency communication center might be lower priority because they have a fall-back communication capability at the state level.

Finally, focusing on recovery and restoration efforts promotes resilience planning. In contrast, loss in value is less useful in resilience planning and tends to put the focus on prevention and avoidance. The example in the previous paragraph points to a very important principle of systemic resilience – the ability to restore operations to some minimal level quickly, even though it might be far below the normal level. Furthermore, most organizations and many families have previous experience in recovery and restoration from shocks or emergencies in other settings – e.g. adverse

[4] There are three general approaches for valuing assets (McKinsey & Company Inc et al. 2010): 1) the *replacement cost* approach; 2) the *income* approach, which depends on being able to estimate the income stream flowing from the asset; and 3) the *market value* approach, which depends on being able to identify a set of comparable market prices for the asset, assuming that it is frequently bought and sold. The proposed 'recovery and restoration effort' method is similar to the recovery cost approach, while the 'loss in value' methods could include all three approaches and thus has higher information and analysis requirements.

weather, power blackouts, civil unrest, economic crisis, etc. All of these can be used to estimate recovery and restoration in the case of information security breaches. Thus, estimating recovery and restoration efforts is very compatible with resilience planning in many different scenarios.

A critic might argue that measuring impact through recovery and restoration efforts has flaws, too. The following paragraphs list possible flaws with the proposed approach, along with our rejoinders.

There might be an endowment effect where stakeholders who have limited or no resources to invest in recovery and restoration might undervalue their losses. One a way to mitigate any undervaluing is to include recovery and restoration efforts that might be funded by identified sources other than the affected stakeholder. This has empirical grounding in studies of disaster response, refugee relief, and similar settings.

The recovery and restoration approach might be too strict and leave out some material consequences or losses. While this may be significant in *ex post* settings, we believe that it is useful to have conservative estimates of impact for *ex ante* planning, investment, and incentive purposes. If nothing else, conservative estimates are more easily validated, are more repeatable when examined by independent analysts, will be more credible to decision-makers, and are less likely to be disputed.

Another possible flaw is that stakeholders might be biased or or be unable to estimate recovery/restoration efforts *ex ante*, and thereby undervaluing them. While it is true that several forms bias, such as optimistic overconfidence, blame shifting, or disaster myopia, might distort *ex ante* valuation and estimation, we believe that the best remedy for bias is to focus stakeholder attention on the practical tasks of recovery and restoration in given scenarios.

Finally, there might be some impact scenarios that are considered unrecoverable, such as firm bankruptcy, irreparable damage to a person's reputation, or, on large scale, permanent harm to national security. These are addressed in the next item.

It should be noted that the two methods – loss in value and recovery and restoration effort – would have identical or very close valuation of direct monetary losses due to theft, embezzlement, fraud, or extortion. If $100,000 is stolen from a company's bank account due to ACH fraud, both the loss in value method and the recovery and restoration effort method would value the impact as $100,000 plus the cost of investigation and prosecution.

2.5. Impact of Non-recoverable Consequences as the Limit of Recoverable Scenarios

We propose that the impact of non-recoverable consequences can be estimated by constructing a sequence of scenarios, each progressively more severe but recoverable, leading to the non-recoverable scenario. Consider firm bankruptcy. Bankruptcy is not an isolated phenomenon but instead is part

of a progression of events related to declining credit quality, declining liquidity, declining financing capability, and generalized financial stress. Each stage in the progression might be associated with a credit downgrade by a rating agency or credit bureau. But a credit downgrade short of bankruptcy is recoverable through financial or operational actions, or both – e.g. sale of assets, reduction in cost structure, raising prices, terminating unprofitable product lines, etc. Even a bankruptcy itself is recoverable in some cases through reorganization and recapitalization. There is a wealth of research on the cost of declining credit quality and bankruptcy, and we suggest that this could provide sufficient foundation to estimate the cost of non-recoverable bankruptcy as a limit.

This concludes our presentation of the proposed analysis framework.

3. Branching Activity Models

This section describes branching activity models for response and recovery efforts associated with breach episodes as a method to estimate . A branching activity model is a tree or graph (Bolton and Davies 2000) where nodes are activities (also known as routines, business processes, or projects) and the edges represent an existential causal relationship between activities: e.g. A_2 is *initiated-by* A_1, or A_2 is *spawned-by* A_1, etc. Thus, every activity is initiated by some other activity, and each activity executes until it reaches some stopping condition. The structure of the branching activity model is generally not fixed or perfectly predictable *ex ante*. Instead, the structure is dynamic and depends on the emerging context and endogenous discoveries.

Branching activity models of recovery and restoration efforts have some similarities to attack graphs (Wang et al. 2007) which have been used in information security research for some time, but they are very different. Attack graphs show the sequence of breaches or penetrations necessary for the attacker to achieve some goal. However, attack graphs do not reveal any information regarding the *negative consequences* or *impact* of successful attacks. In contrast, the branching activity models of recovery and restoration do not include the attack or breach process, nor the 'damage' or harm at the IT asset level that resulted from the attack(s). Therefore, they are complementary to each other for most purposes, including probabilistic risk assessment.

3.1. An Illustrative Example

Consider a fictional breach episode at a manufacturer of control systems, where attackers apparently stole confidential information related to the supply chain, perhaps with a goal of under-bidding on major international contracts. A simplified branching activity model for response and recovery is shown in Figure 1. The response and recovery activities involved both the breached firm and their major supply chain partners. Assume that activities A, B, and C are 'Standard Operating Procedure' for incident response for this firm and thus are quite predictable in their duration and costs. The other activities – D, E, and F – are less routine and are defined according to the context. The activities were initiated as follows:

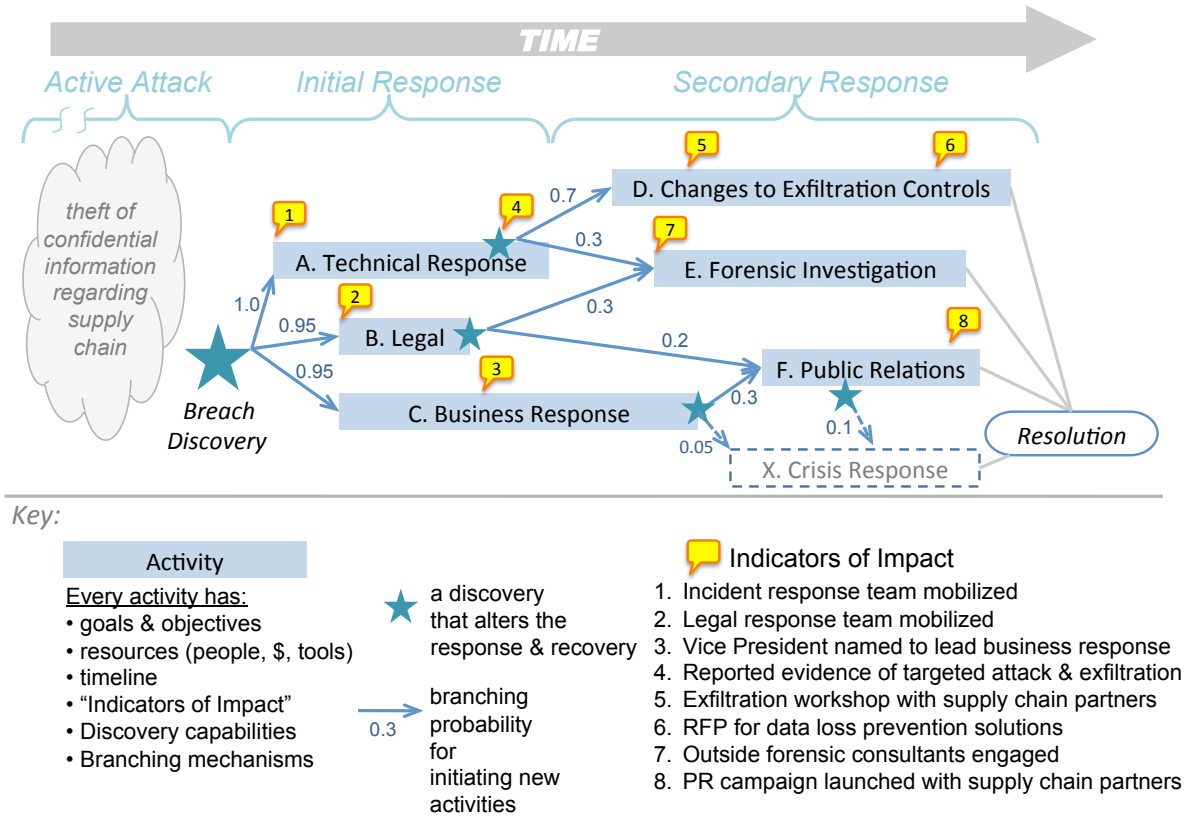

Figure 1 **Fictional example of a breach episode, illustrating a simplified branching activity model for response and recovery. Shows one counterfactual activity: "X. Crisis Response". (Source: authors)**

A. *Technical Response* – the IT and Security teams investigate the breach, identify affected IT assets, and implement immediate mitigations. Along the way, they discover evidence of custom malware and patterns of compromise that point to supply chain information as the goal of the attack. The cost of this activity is proportional to the number of IT assets compromised and the complexity of mitigations required.

B. *Legal Response* – the Legal team initially advices executives regarding governing laws and regulations, contacts law enforcement, and informs the rest of the response teams regarding governing laws and regulations. Along the way, they discover evidence related to the breach that leads them to start a forensic investigation, due to regulations governing supply chain security. The cost of this activity is relatively fixed for any breach above a threshold of criticality.

C. *Business Response* – the Business team investigates the business consequences, including contracts, product quality, and customer relations. Along the way, they discover that the breach could affect on-going contract negotiations with major critical infrastructure customers in the US and internationally. However, their investigation concluded that the

confidential information would not likely alter their chances of winning new business. The cost of this activity is proportional to the number of business processes and business units affected, which increases team size and duration.

D. *Changes to Exfiltration Controls* – an ad hoc working group is formed to investigate the state of exfiltration controls across the supply chain. Both business and technical solutions are explored and implemented. The cost of this activity is proportional to the number of supply chain partners and their relative maturity regarding information security, and also their previous experience in collaboration on information security.

E. *Forensic Investigation* – an outside consulting firm is engaged to gather as much evidence as possible about the attack methods, procedures, and characteristics, with a goal of attributing the attack, supporting law enforcement, and fulfilling regulatory requirements. The cost of this activity is proportional to the amount of digital evidence available and the difficulty of recovering it for investigation purposes.

F. *Public Relations* – a focused PR campaign was launched, including as team members staff from each major supply chain partner. The PR campaign was aimed at major customers, regulators, shareholders, and employees across the supply chain. The cost of this activity is proportional to the number of campaign elements and the overall duration.

X. *Crisis Response* – This is 'counterfactual' because it was not realized in this fictional breach episode. No crisis materialized. Even so, there is a non-zero probability that the negative consequences could cascade into a full-blown crisis, leading to loss of major contracts, law suits, and other very large costs. In this fictional case, nearly all the worst-case scenarios incur the largest proportion of costs from Crisis Response activities (see Figure 2b).

The total impact of this breach measured as total costs, I, would be calculated as the sum of the costs of individual activities, $C_i,\ i \in (A \ldots F)$, plus the additional Crisis Response activities that must be accounted for in the *ex ante* decision frame, C_X. Each is multiplied by ξ_i which is its *ex ante* probability of being realized in the branching process:

$$I = \xi_A C_A + \xi_B C_B + \xi_C C_C + \xi_D C_D + \xi_E C_E + \xi_F C_F + \xi_X C_X \tag{1}$$

Notice that this is not a sum of scalar numbers, but instead a sum of random variables, each with a distribution of costs associated with it and also a probability of being realized. With heterogeneous branching processes and cost distributions, it is necessary to use Monte Carlo simulation[5] to estimate the aggregate distribution I and the contingency probability ξ. Because of the simplicity of the illustrative example, the following values for ξ were calculated analytically:

$$\xi_A = 1.0; \quad \xi_B = 0.95; \quad \xi_C = 0.95; \quad \xi_D = 0.7; \quad \xi_E = 0.43; \quad \xi_F = 0.42; \quad \xi_X = 0.09 \tag{2}$$

[5] Appendix **??** lists the computer code used for this illustrative example.

3.1.1. Discussion. Figure 2 shows results of Monte Carlo simulations for two treatments for the activity tree. Figure 2a treats the activity network as fixed, meaning that the only source of variation is the distribution of costs for each activity. It is not surprising to see a distribution that is somewhat similar to lognormal because the cost distributions are either Normal, lognormal, or mixtures of the two. Most important, the distribution has a single mode and moderately wide variance. Thus, the mean is a good statistic for the central tendency for this distribution. If this were the only analysis that we saw, we might conclude that these activities have a single combined effect in terms of cost (our measure of aggregate impact in this example).

Figure 2b treats the activity network as a branching activity model as prescribed by this paper. Two hundred branching trees were realized using random draws given the branching probabilities, and then 500 iterations of Monte Carlo simulations for each realization were performed. It is interesting to see that the mean in this treatment is very close to the mean in Figure 2a. But the similarities stop there. First, it is very apparent that this distribution has two very distinct modes, and can even be seen as a simple mixture of two distributions, even though we know that seven activities contribute their distributions to the final mixture. Second, the variance is much wider than in the first treatment, not only on the upper tail due to the inclusion of the 'X. Crisis Response' activity but also on the lower tail, meaning that there is considerable probability mass below the mean. In fact, we can safely say that the mean is *not* a good measure of central tendency in this distribution, as it will be frequently either too high or too low for individual realizations.

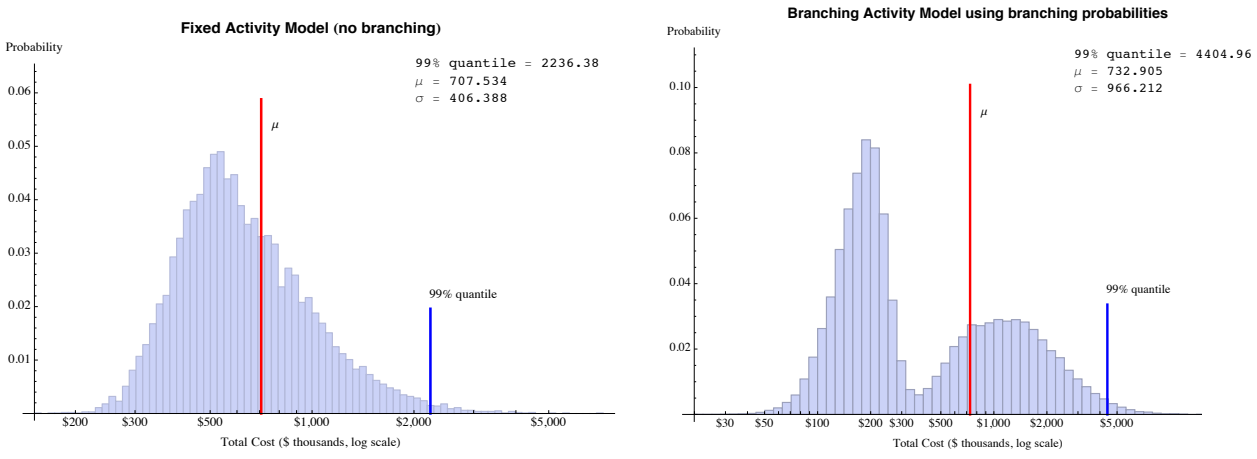

(a) Fixed Activity Model (b) Branching Activity Model

Figure 2 Results of Monte Carlo simulation for two treatments of the activity model in Figure 1.

Based on this data and analysis, what can we say about the aggregate impact of this breach episode? One interpretation is to view this breach episode as a near-miss for a much more severe

episode that would at least double the costs and involve serious crisis management. Indeed, the simplest interpretation of Figure 2b is that there are two categories of breach episodes: 1) breach episodes that are contained in the initial response; and 2) breach episodes that can require response and recovery via secondary response, which is often lengthy and expensive.

Since this is a hand-crafted and fabricated example, we should not try to draw any general conclusions. This example does show the potential for the proposed branching activity model method to make good use of available evidence to derive estimates of breach episode impact. In particular, it is very clear in this example how the branching activity model reveals more information about the causal structure that underlies breach response and recovery. We believe that the results will be even more revealing in settings that are more complex – i.e. cascading consequences or severe uncertainties.

3.2. Method for Empirical Analysis

The following are the steps that an analyst would take to analyze particular breach episodes using the branching activity model and Indicators of Impact (c.f. Section 4).

Step 1. Define the scope of the breach episode in terms of affected people and organizations, and also in time.

Step 2. Identify and list the primary sources of evidence .

Step 3. Construct a timeline, both for historical events and possible future events.

Step 4. List Indicators of Impact using primary sources from Step 2.

Step 5. Identify uncertainties, missing information, ambiguity, etc.

Step 6. Draw a branching activity model, similar to Figure 1.

Step 7. Estimate parameters for the cost function for each activity (e.g. time, resources, $).

Step 8. Estimate aggregate impact statistics using Monte Carlo simulation.

This list of steps applies to the analysis of a single breach episode. The next section will discuss additional analysis methods that can be applied to a mass of evidence regarding a collection of breach episodes.

3.3. Model Selection and Estimation Methods

Model selection is an inference process where alternative mathematical or statistical model structures are evaluated against some criteria of fitness, correctness, or parsimony (Burnham and Anderson 2002). In our setting, we face a challenge of choosing from alternative tree or graph structures for a branching activity model. Likewise, estimating the branching probabilities is a critical step in the method, and thus would benefit from the combined application of complementary estimation methods. The following is a list of possible methods and their application to this setting. Clearly, this is fertile ground for future research.

• *Statistical analysis of historical breach data sets* – This would include statistical analysis of aggregate data (e.g. number of records compromised vs. attack methods) and also statistical analysis of causal relationships (e.g. disclosure practices in an industry vs. breaches experienced (Rees 2009, Rees and Kannan 2008, Sohail 2006, Wang et al. 2008)). The Verizon RISK Team and the Digital Forensics Association have applied statistical methods to study populations of breach events (Baker et al. 2012, Widup 2011).

• *Probabilistic reasoning using Indicators of Impact and other evidence* – This can include pattern recognition, classification, inductive reasoning from evidence, and reasoning about uncertainty (Schum 2001, Barker and Haimes 2009, Loch et al. 2008, Masys 2012, Sniedovich 2007, van Asselt and Rotmans 2002, Ferrin et al. 2009). These insights and inferences should be very informative for Steps 5, 6, and 7, above.

• *Opinions of subject matter experts* – Calibrated estimates of probability can be reliably elicited from subject matter experts (Hubbard 2010). But rather than having the experts try to estimate aggregate impact or costs, it might be much more fruitful to elicit estimates of branching probabilities and other conditional probabilities.

• *Simulations* – Building simulations of breach response and recovery as a socio-technical system could be very useful in estimating branching probabilities and other parameters of a branching activity model. Candidate methods include stochastic process models and Agent-based Modeling (ABM).

• *Experiments* – Human subject experiments could be a valuable source of data and insights, especially regarding *ex ante* valuation of alternative breach response and recovery models. These experiments could be included as part of scenario analysis exercises, Red Team vs. Blue Team exercises, and crisis walkthrough exercises.

• *Case studies* – In-depth analysis of specific cases can yield many valuable insights regarding the qualitative and quantitative aspects of breach response and recovery activities and decision processes.

• *Theoretical models* – In addition to these empirical methods, theoretical models of breach response and recovery could be very helpful (e.g. Gordon et al. (2003)), either to establish Bayesian priors for key parameters or to explore a full range of possible outcomes for a given setting.

3.4. Computer Simulation

We are in the early stages of computer implementation and simulation[6] of branching activity models, and so our discussion here is preliminary. The following is a few remarks on alternatives and promising approaches.

[6] The computer code for used in the illustrative example in Section 3.1 can be found at [URL]

There is a rich literature on deterministic, stochastic, and decision-theoretic branching systems. Examples of deterministic systems include rewrite systems (Prusinkiewicz et al. 2003), fractals, and cellular automata (Wolfram 2002). Examples of stochastic branching processes include Galton-Watson process and variants (Orenshtein 2011). Other examples include Stochastic Petri Networks (Chen et al. 2011) and Stochastic Activity Networks (Sanders and Meyer 2001, Vakili et al. 2011). Most of these stochastic systems can be converted into equivalent Markov Chains, either Discrete Time or Continuous Time, and this provides a very general computational method – Markov Chain Monte Carlo (MCMC) and modern improvements. Indeed, there are representation methods that have promise for our purposes, such as Semi-Markov Processes, Markov Decision Processes, and Hidden Markov Models.

However, one general short-coming of these methods is that they assume a static causal structures. As described in previous sections, this setting calls for a method that can encompass dynamic causal structures. There are some very interesting methods being developed in the field of Computational Biology to model morphogenesis. These are include context-sensitive, non-deterministic rewrite systems such as P1L-System and Dynamical System with a Dynamical Structure $(DS)^2$ in the MGS system (Giavitto and Michel 2005, Giavitto et al. 2011, Michel 2007).

An area for future investigation is to add decision-theoretic and game-theoretic branching mechanisms to the deterministic and stochastic mechanisms just discussed.

Finally, there are interesting developments in the field of Cognitive Science related to probabilistic programming using language systems such as CHURCH (Goodman et al. 2008). Probabilistic programming is an extension and generalization of Bayesian belief networks and is well suited to formal representation of causal reasoning and induction of theory (Griffiths and Tenenbaum 2007, 2009). This capability could augment the stochastic simulation by supporting probabilistic inference from evidence regarding branching activity models.

This concludes the presentation on the branching activity model. The following two sections describe how branching activity models can be estimated from evidence and other means.

3.5. Assumptions and Limitations

While we aim to make the branching activity model as general as possible, it is based on several important assumptions regarding the nature of breach response and recovery and on the application of the model in practice. Many of these assumptions need testing or validation:

1. Impact is more meaningful on the holistic level than it is on the asset level.
2. Estimating the recovery/response activities will be easier than enumerating vulnerabilities and exploits of IT assets, and then estimating resulting loss of value.
3. The existential causal structure between activities is essentially a tree.

4. Branching is not fundamentally a stochastic process as a function of the duration of activity.
5. Impacts that cannot be translated to recovery or restoration activities are not significant or relevant.
6. With enough preparation and recovery/restoration activities, breached organizations can discover what they need to know about the breach and its negative consequences.
7. We can safely abstract away from the detailed discovery and decision processes with each activity.
8. Organizations and institutions will be more willing to share information in the form of Indicators of Impact and breach activity models than they currently do regarding details of specific breach episodes.

The proposed methods has limitations. Perhaps the most significant limitation is that it not especially well suited to one-off analysis of a single breach episode. While the example above in Section 3.1 shows this is possible, the analyst may find it laborious and perhaps no more enlightening than a formula-based method. In contrast, the primary benefits of the proposed method will shine when large amounts of data are collected and statistical inference methods are applied to identify regularities and common patterns.

Another important limitation is that the proposed method is not well suited to *ex post* estimation of damage, harm, or liability for legal or regulatory purposes, or to assign blame for political purposes. There are many requirements for *ex post* estimation that are not included and are not compatible with our proposal.

4. Indicators of Impact

An 'Indicator of Impact' is an observable event, behavior, action, state change, or communication that signifies that the breached or affected organizations are attempting to respond, recover, restore, rebuild, or reposition because they believe they have been harmed. For our purposes, Indicators of Impact are evidence that can be used to estimate branching activity models of breach impact, either the structure of the model or key parameters associated with specific activity types. In principle, every Indicator of Impact is observable by someone, though maybe not outside the breached organization.

It is also useful to define them by what they are not. The following are *not* Indicators of Impact:

- Opinions, beliefs, fears, or concerns
- Unsubstantiated claims, guesses, or imagined possibilities
- Quantitative estimates of aggregate impact in terms of cost, losses, damages, etc.
- Ordinal estimates of aggregate impact in terms of harm, severity, etc.

The last two items might be surprising to some readers. After all, if someone in a position of authority or familiarity makes a public statement that includes a quantitative estimate of impact, is that not evidence? We argue that the answer is 'no'. An aggregate estimate, alone, does not shed any useful light on the structure of the branching activity model or the parameters of any activity within that model. Of course, if that spokesperson enumerates costs incurred by category and discusses how resources are deployed, then we can find Indicators of Impact in that communication.

Appendix A provides a draft list of Indicators of Impact that we have begun to use in breach impact investigations. We believe that this list will grow and change over time, and there is probably value in keeping it open and loose rather than tightly standardized. One category of Indicators worth considering for addition relate to threat agents, threat agent actions, motives and capabilities, and attribution of a breach to particular threat agents.

4.1. Sources

One very important characteristic of Indicators of Impact is that information about them can come from a wide variety of sources. This stands in sharp contrast to sources of information regarding possible loss in value where there are very few reliable sources, as discussed in Section 2.4. The following is a list of sources we have been drawing on in case studies.

- News Reports and Official Statements
- Legal and Regulatory Filings, including SEC filings and correspondence.
- Social Media
- Financial and Industry Analyst Reports
- Security Industry White Papers, Blogs, and Technical Reports
- Leaked Documents
- Posted Procedures and Standards
- Public Communication Activities
- Executive Activities, Movements, and Communications
- Corporate Transactions
- Trade Press
- News Releases by Professional Service Firms (e.g. PR Agencies)
- Conference Presentations
- Requests for Proposals, Requests for Information, etc.
- Staffing Changes

In summary, collecting information about Indicators of Impact is very similar to Open Source Intelligence that is practiced by National and Homeland Security professionals, and also industrial analysts who follow their competitors.

5. Discussion

This closing section will address three broad questions: 1) Does the proposed method get to 'ground truth' regarding impact of historical breaches?; 2) Can breach impact estimates be generalized and mapped to a "severity index" to support risk communication and public awareness?; and 3) What are the use cases for the proposed method, and how will it make a difference? The section will conclude with a short discussion of opportunities for future research.

5.1. The 'Ground Truth' of Breach Impact is *Ex Ante* Valuation

History is overrated, at least regarding breach impact. What matters most is what might happen in the future if a given breach episode were to occur again. One way to demonstrate the validity of this assertion through a simple example of betting on coin tosses. Imagine that you receive a payment of $100 for each toss of 'heads' and nothing for 'tails'. Consider two realizations of ten tosses each. Imagine that in the first realization the coin tosses are all 'heads', and in the second realization the coin tosses are all 'tails'. What is the 'ground truth' of the coin tossing process? We assert that the most important 'ground truth' is the fairness of the coin tossing process, or, more generally, the relative probability of 'heads' vs 'tails'. While the historical payoff from the two realizations ($1,000 and $0) certainly affects your bank balance, what matters most for future betting is the *expected future payoff* of coin tossing. Our estimate of expected future payoffs would certainly be informed by the historical payoffs and other relevant evidence, but is not simply a record of history. In other words, we would not expect that future realizations would be either all 'heads' or all 'tails' unless there was some really strong evidence to back up that unlikely inference. The real 'truth' about our history in this case is that you were really lucky in the first realization and were really unlucky in the second. History matters but only as a source of evidence to use for understanding the future.

Likewise, in any given breach episode, there will be chance factors that might cause the actors to be lucky or unlucky. Some actors could also make strategic decisions to minimize or exaggerate the costs through less-than-prudent decisions, e.g. to forego forensic investigation or to avoid notifying legal authorities. Both chance and strategic behavior could cause the historical realization of breach impact to be much higher or much lower than would be expected *ex ante*. Thus, the 'ground truth' of any given breach episode is the expected value of impact, using both history and other evidence in our estimation procedure.

Perhaps the best way to view our approach is to say we are applying a *valuation procedure for breach impact* for a given breach process, and that this breach process has been realized in a particular case if the breach has already occurred, or might be realized in particular scenarios if such a breach has not yet occurred.

5.2. Generalizing Breach Impact Estimates to Communicate Severity

Many risk analysis methods use an ordinal category scale for breach severity, e.g. 'High', 'Medium', and 'Low'. In a similar vein, there have been proposals for a severity index (Stiennon 2013) to facilitate public discussion and communication. Clearly these methods have intuitive appeal and it would be desirable to be able to map the quantitative estimates of aggregate impact derived from branching activity models to some more general ordinal categories of severity.

We believe that this might be possible, but it critically depends on how we define the criteria for the ordinal categories. One approach is to use first-order stochastic dominance (Blavatskyy 2011), where breaches in higher categories would be first-order stochastically dominant over breaches in lower categories. Another approach would be to use methods of statistical physics to characterize the structure of the dynamics in state space to categorize different breach episode types. Further research is needed to see if either of these methods are viable, robust, and sufficient. If they are, then the resulting severity index would be much more sound than existing proposals such as Stiennon (2013) .

5.3. Use Cases

- *Enterprise Risk Management* – The proposed method would be very valuable in probabilistic risk assessment, including financial measures of risk at the level of business units or the enterprise as a whole. (Thomas 2009b, Society of Actuaries 2010)
- *Near-miss and Precursor Identification* – The branching activity model is well-suited to assist in identifying near-miss events and precursors because it includes counterfactual activities and estimates of branching probabilities for those activities. In contrast, formula-based methods for breach impact estimation exclude or under-value near miss events.
- *Real-time Situation Awareness and Crisis Management* – In the context of an unfolding breach episode, perhaps in a crisis situation, it is often vital for decision-makers to have accurate and up-to-date assessments of the negative consequences and severity of damages, and also some way of estimating the effectiveness of alternative response activities. Branching activity models could be very useful in this application. It could aid both central crisis management teams and decentralized action teams to quickly assess the likely impact, given the evidence, and the possible mitigating actions that are available. There is also a natural connection to preparation exercises, where the branching activity models could be both utilized and also tuned based on the outcomes of exercises.
- *Resilience Planning* – In resilience planning for information security, we assume that a given breach episode has already happened, and the planner's goal is to minimize the impact and to recover to a workable state as soon as possible (Dondossola et al. 2011, Rose 2004a, Vugrin et al. 2011). Because the proposed branching activity model is focused on response and recovery activities, it is ideally suited to support resilience planning.

- *Public Disclosure and Risk Communication* – The branching activity model and impact estimates derived from it should provide a viable method for public disclosure of breach episodes in a way that is more meaningful to stakeholders than current disclosure methods. Furthermore, if the disclosure is in the form of a parameterized branching activity model, then the disclosing firm will not be revealing any information to attackers that might aid them in future attacks. Also such a disclosure will not expose the firm to greater risk of legal or regulatory liability, since the branching activity model does not include information relevant to those purposes. Finally, as discussed in Section 5.2, categorizing breach episode types according to general ordinal categories of severity could improve risk communication.
- *Critical Infrastructure Information Sharing* – In critical infrastructure sectors, the value and importance of sharing information regarding information security is heightened by the interdependent nature of risk, by the severity of worst-case breach episodes, and the broad social consequences (ENISA and RAND Europe 2010). Most information sharing today involves tactical security information regarding technical vulnerabilities and real-time threat activity. There is also some sharing of 'best practices', though this seems to be simply pooling of ideas rather than any formal evaluation of practices to identify the 'best'. One difficultly is this: How will anyone use the information that is shared? It often lacks sufficient context and it is hard for recipients to 'connect the dots' for their world The benefit of sharing information in the form of branching activity models is that it could add valuable context and causal information, both for planning purposes and for understanding historical breach episodes.
- *Incentive Instruments and Institutions* – When augmented by appropriate data collection and analysis capabilities, the proposed branching activity model could be very useful in various incentive instruments and institutions. This could include traditional cyber insurance (Department of Homeland Security 2012, Kunreuther and Heal 2012, Zhao et al. 2009) or innovative institutions (Kupiec and O'Brien 1997, Sommer and Loch 2009, Thomas and Amon 2007, Thomas 2009b).
- *Region and Sector Models of Disaster Impact* – There is an extensive research literature on the estimation of the economic impact at a region or sector level of blackouts and natural disasters, and homeland security incidents. This research uses Input-Output models, Computable General Equilibrium , and multi-level Agent Based Modeling (ABM) (Anderson et al. 2007, Clement 1989, Rose 2004b, Okuyama 2007, Rose 2004a, 2009). While the proposed branching activity model does not explicitly include sector flows and adjustments in market prices, the proposed branching activity model should be very compatible with these region and sector models. It does include information on causal pathways, including cascades in interdependent systems. An important area of future research will be to bridge these models particularly in ABMs where the micro-level mechanisms can be implemented directly in agent behavior or decision rules.

• *Regulation, Compliance, and Audit* – Finally, branching activity models could significantly enhance the value of regulation, compliance, and audit institutions. First, if Indicators of Impact were included in audit processes, there would be much more substantial basis for justifying audit judgments. Second, regulatory and compliance regimes could become less prescriptive without losing their ability to influence stakeholder decisions and behavior in the direction of social welfare. Branching activity models provide a mechanism to link a firm's plans and preparations to outcomes, both in particular historical breach episodes and also for relevant future outcomes.

5.4. Future Research

The proposed method is in the early stages of development and testing, and thus would greatly benefit from additional research. In the body of the paper we have mentioned opportunities for research in specific areas, so this subsection will focus on a broad research agenda.

A major challenge facing the research community is how to advance the state of knowledge regarding economic, social, and behavioral incentives regarding information security (Cashell et al. 2004, Anderson et al. 2008, National Science and Technology Council 2011, National Science Foundation 2012, Verendel 2009, Department of Homeland Security 2012, Energy Sector Control Systems Working Group (ESCSWG) 2011, ENISA 2012). Progress has been painfully slow, both compared to pace of research progress on technical information security and compared to the demands of the rapidly changing environment. We believe that the proposed branching activity model can serve as a useful building block in a broader research program that unifies theoretical and empirical social science research for information security, and also professional practice. Obviously, a significant program of research is needed to test this assertion.

Another major opportunity for future research would be to incorporate the capabilities and strategies of threat agents in order to estimate the frequency of breach episode types, and also to estimate the likely harm that different threat agents might cause. Along these lines, another important extension would be to model the dynamics of adversarial innovation in a way that accounts for the unfolding co-evolutionary landscape of attack and defense capabilities (e.g. Thomas (2011b)).

Finally, it would be very valuable to perform research that can shed light on the social processes of risk management to see if the proposed methods actually improve individual, organization, and social decision processes.

Acknowledgments

The authors thank the anonymous reviewers for comments and suggestions. We also gratefully acknowledge the members of Society of Information Risk Analysts (`https://www.societyinforisk.org/`) and Securitymetrics.org and its METRICON workshops (`http://www.securitymetrics.org/`) for many fruitful discussions and ideas. The Verizon RISK team, especially Wade Baker and Jay Jacobs, have been very supportive and have contributed on many levels. Thanks to Chris Walsh for his suggestions regarding SEC filings as a source for Indicators of Impact. Josh Corman, Fred Doolittle, Steve Dotson, Thomas Elegante, Douglas Foster, Chris Hayes, Dr. John Gero, Dr. Larry Ponemon, Ben Sapiro, and Dr. Arun Sood have provided good feedback during formative stage discussions and on early drafts. The idea of 'Indicators of Impact' draws inspiration from Mandiant's 'Indicators of Compromise', which is now supported by the OpenIOC community (`http://www.openioc.org/`. Finally, Mr. Thomas has received partial support by a grant from the National Science Foundation, grant no. SBE-0915482. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Acquisti, Alessandro, Jens Grossklags. 2004. Losses, gains, and hyperbolic discounting: An experimental approach to information security attitudes and behavior. *The Economics of Information Security*. Kluwer Academic Publishers, 165 – 178.

Acquisti, Alessandro, Jens Grossklags. 2005. Uncertainty, ambiguity and privacy. *Fourth Annual Workshop Economics and Information Security (WEIS)* .

Alamgir, Hasanat, Shicheng Yu, Erin Gorman, Karen Ngan, Jaime Guzman. 2009. Near miss and minor occupational injury: Does it share a common causal pathway with major injury? *American Journal of Industrial Medicine* **52**(1) 69 – 75. doi:10.1002/ajim.20641. URL: http://onlinelibrary.wiley.com/doi/10.1002/ajim.20641/abstract.

Anderson, Christopher W., Joost R. Santos, Yacov Y. Haimes. 2007. A risk-based input-output methodology for measuring the effects of the august 2003 northeast blackout. *Economic Systems Research* **19**(2) 183 – 204. URL: http://mutex.gmu.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=25192156&site=ehost-live.

Anderson, Ross, Chris Barton, Rainer Böhme, Richard Clayton, Michel JG van Eeten, Michael Levi, Tyler Moore, Stefan Savage. 2012. Measuring the cost of cybercrime. *Proceedings of the 11th Workshop on the Economics of Information Security*. Berlin, Germany, 1–31. URL: http://weis2012.econinfosec.org/papers/Anderson_WEIS2012.pdf.

Anderson, Ross, Rainer Böhme, Richard Clayton, Tyler Moore. 2008. Security economics and the internal market. Report, European Network and Information Security Agency, Heraklion, Crete, Greece. URL: http://www.enisa.europa.eu/publications/archive/economics-sec/at_download/fullReport.

Baker, Wade, Mark Goudie, Alexander Hutton, C. David Hylender, Jelle Niemantsverdriet, Christopher Novak, David Ostertag, Christopher Porter, Mike Rosen, Bryan Sartin. 2012. Data breach investigations report 2011. Report, Verizon Communications, Inc. URL: http://www.verizonenterprise.com/resources/reports/rp_data-breach-investigations-report-2012-ebk_en_xg.pdf.

Baker, Wade, Alex Hutton, Chris Porter. 2010. VERIS - a framework for gathering risk management information from security incidents. URL: http://www.securitymetrics.org/attachments/Metricon-4.5-Baker-Hutton-VERIS.pdf.

Barker, Kash, Yacov Y. Haimes. 2009. Assessing uncertainty in extreme events: Applications to risk-based decision making in interdependent infrastructure sectors. *Reliability Engineering & System Safety* **94**(4) 819 – 829. doi:10.1016/j.ress.2008.09.008. URL: http://www.sciencedirect.com/science/article/pii/S0951832008002263.

Blavatskyy, Pavlo R. 2011. A model of probabilistic choice satisfying first-order stochastic dominance. *Management Science* **57**(3) 542 – 548. URL: http://mansci.journal.informs.org/content/57/3/542.full.pdf.

Bodeau, D.J. 1992. A conceptual model for computer security risk analysis. *Computer Security Applications Conference, 1992. Proceedings., Eighth Annual*. 56 – 63. doi:10.1109/CSAC.1992.228233.

Bodin, Lawrence D., Lawrence A. Gordon, Martin P. Loeb. 2005. Evaluating information security investments using the analytic hierarchy process. *Communications of the ACM* **48**(2) 78 – 83. doi:10.1145/1042091.1042094. URL: http://doi.acm.org/10.1145/1042091.1042094.

Böhme, Rainer, Tyler Moore. 2010. The iterated weakest link – a model of adaptive security investment. *IEEE Security & Privacy* **8**(1) 53 – 55.

Bolton, Christie, Jim Davies. 2000. Activity graphs and processes. Wolfgang Grieskamp, Thomas Santen, Bill Stoddart, eds., *Integrated Formal Methods*, *Lecture Notes in Computer Science*, vol. 1945. Springer, Berlin/Heidelberg, 77 – 96. URL: http://dx.doi.org/10.1007/3-540-40911-4_6.

Borg, Scott. 2005. Economically complex cyberattacks. *IEEE Security & Privacy* **3**(6) 64 – 67. doi:10.1109/MSP.2005.146.

Brown, Nanette, Yuanfang Cai, Yuepu Guo, Rick Kazman, Miryung Kim, Philippe Kruchten, Erin Lim, Alan MacCormack, Robert Nord, Ipek Ozkaya, Raghvinder Sangwan, Carolyn Seaman, Kevin

Sullivan, Nico Zazworka. 2010. Managing technical debt in software-reliant systems. *Proceedings of the FSE/SDP Workshop on Future of Software Engineering Research (FoSER '10)*. ACM, New York, NY, USA, 47 – 52. doi:10.1145/1882362.1882373. URL: http://doi.acm.org/10.1145/1882362.1882373.

Buldyrev, Sergey V, Roni Parshani, Gerald Paul, H Eugene Stanley, Shlomo Havlin. 2010. Catastrophic cascade of failures in interdependent networks. *Nature* **464**(7291) 1025 – 1028.

Burnham, Kenneth P., David R. Anderson. 2002. *Model Selection and Multi-Model Inference: A Practical Information-Theoretic Approach*. Springer.

Campbell, Katherine, Lawrence A. Gordon, Martin P. Loeb, Lei Zhou. 2003. The economic cost of publicly announced information security breaches: empirical evidence from the stock market. *Journal of Computer Security* **11**(3) 431 – 448. URL: http://dl.acm.org/citation.cfm?id=876661.876669.

Carreras, B. A., D. E. Newman, Paul Gradney, V. E. Lynch, I. Dobson. 2007. Interdependent risk in interacting infrastructure systems. *40th Annual Hawaii International Conference on System Sciences, 2007. HICSS 2007*. IEEE, 112. doi:10.1109/HICSS.2007.285. URL: http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=4076597.

Cashell, Brian, William D. Jackson, Mark Jickling, Baird Webel. 2004. The economic impact of cyber-attacks. Report RL32331, United States Congressional Research Service, Washington, DC, USA. URL: http://www.cisco.com/warp/public/779/govtaffairs/images/CRS_Cyber_Attacks.pdf.

Cavusoglu, Huseyin, Birendra Mishra, 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) 70 – 104. URL: http://dl.acm.org/citation.cfm?id=1278168.1278173.

Chen, Thomas M, Juan Carlos Sanchez-Aarnoutse, John Buford. 2011. Petri net modeling of cyber-physical attacks on smart grid. *IEEE Transactions on Smart Grid* **2**(4) 741–749.

Clement, C. F. 1989. The characteristics of risks of major disasters. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* **424**(1867) 439 – 459. URL: http://www.jstor.org/stable/2398381.

Department of Homeland Security. 2012. Cybersecurity insurance workshop readout report. Workshop report, US Department of Homeland Security. URL: http://www.dhs.gov/sites/default/files/publications/cybersecurity-insurance-read-out-report.pdf.

Dobson, Ian, Kevin R Wierzbicki, Benjamin A Carreras, Vickie E Lynch, David E Newman. 2006. An estimator of propagation of cascading failure. *System Sciences, 2006. HICSS'06. Proceedings of the 39th Annual Hawaii International Conference on*, vol. 10. 245c – 245c.

Dondossola, G., F. Garrone, J. Szanto. 2011. Cyber risk assessment of power control systems; a metrics weighed by attack experiments. *2011 IEEE Power and Energy Society General Meeting*. 1 – 9. doi:10.1109/PES.2011.6039589.

Energy Sector Control Systems Working Group (ESCSWG). 2011. Roadmap to achieve energy delivery systems cybersecurity. Report, US Department of Homeland Security. URL: http://www.cyber.st.dhs.gov/wp-content/uploads/2011/09/Energy_Roadmap.pdf.

ENISA. 2012. Appropriate security measures for smart grids. Tech. rep., European Network and Information Security Agency, Heraklion, Crete, Greece. URL: http://www.enisa.europa.eu/activities/Resilience-and-CIIP/critical-infrastructure-and-services/smart-grids-and-smart-metering/appropriate-security-measures-for-smart-grids/at_download/fullReport.

ENISA, RAND Europe. 2010. Incentives and challenges for information sharing in the context of network and information security. Report, European Network and Information Security Agency, Heraklion, Crete, Greece. URL: http://www.enisa.europa.eu/activities/Resilience-and-CIIP/public-private-partnership/information-sharing-exchange/incentives-and-barriers-to-information-sharing.

Ferrin, Giovanni, Lauro Snidaro, Gian Luca Foresti. 2009. Structuring relations for fusion in intelligence. *12th International Conference on Information Fusion, FUSION'09*. 1621 – 1626. URL: http://isif.org/fusion/proceedings/fusion09CD/data/papers/0418.pdf.

Gafni, Amiram. 1991. Willingness-to-pay as a measure of benefits: relevant questions in the context of public decisionmaking about health care programs. *Medical Care* **29**(12) 1246 – 1252.

Garg, Ashish, Jeffrey Curtis, Hilary Halper. 2003. Quantifying the financial impact of IT security breaches. *Information Management & Computer Security* **11**(2) 74 – 83. doi:10.1108/09685220310468646. URL: http://www.emeraldinsight.com/journals.htm?articleid=862842&show=abstract.

Giavitto, Jean-Louis, Olivier Michel. 2005. Modeling developmental processes in MGS. Marian Gheorghe, ed., *Molecular Computation Models: Unconventional Approaches*. Idea Group Publishing, 150 – 189.

Giavitto, Jean-Louis, Olivier Michel, Antoine Spicher. 2011. Interaction based simulation of dynamical system with a dynamical structure $(DS)^2$ in MGS. *Proceedings of the 2011 Summer Computer Simulation Conference*. 99 – 106. URL: http://www.spatial-computing.org/~michel/PUBLIS/2011/scsc11.pdf.

Goodhart, Charles, Dirk Schoenmaker. 2006. Burden sharing in a banking crisis in europe. *Sveriges Riksbank Economic Review* (2) 34 – 57.

Goodman, Noah D, Vikash K Mansinghka, Daniel M Roy, Keith Bonawitz, Joshua B Tenenbaum. 2008. Church: A language for generative models. *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI-2008)*. Helsinki, Finland, 1–10. URL: http://www.stanford.edu/~ngoodman/papers/churchUAI08_rev2.pdf.

Gordon, Lawrence A., Martin P. Loeb. 2002. The economics of information security investment. *ACM Transactions on Information and System Security* **5**(4) 438 – 457. doi:10.1145/581271.581274. URL: http://doi.acm.org/10.1145/581271.581274.

Gordon, Lawrence A., Martin P. Loeb, William Lucyshyn. 2003. Sharing information on computer systems security: An economic analysis. *Journal of Accounting and Public Policy* **22**(6) 461 – 485. doi:10.1016/j.jaccpubpol.2003.09.001. URL: http://www.sciencedirect.com/science/article/pii/S0278425403000632.

Grebe, Sasha Karl. 2013. Things can get worse: How mismanagement of a crisis response strategy can cause a secondary or double crisis: the example of the AWB corporate scandal. *Corporate Communications: An International Journal* **18**(1) 70–86. doi:10.1108/13563281311294137. URL: http://www.emeraldinsight.com/journals.htm?articleid=17076598&show=abstract.

Griffiths, Thomas L, Joshua B Tenenbaum. 2007. Two proposals for causal grammars. Alison Gopnik, Laura Schulz, eds., *Causal Learning : Psychology, Philosophy, and Computation: Psychology, Philosophy, and Computation*. Oxford University Press, 323 – 345.

Griffiths, Thomas L, Joshua B Tenenbaum. 2009. Theory-based causal induction. *Psychological review* **116**(4) 661. URL: http://cocosci.dreamhosters.com/tom/papers/tbci.pdf.

Grossklags, Jens, Nicolas Christin, John Chuang. 2008. Secure or insure?: a game-theoretic analysis of information security games. *Proceedings of the 17th international conference on World Wide Web*. WWW '08, ACM, New York, NY, USA, 209 – 218. doi:10.1145/1367497.1367526. URL: http://doi.acm.org/10.1145/1367497.1367526.

Grossklags, Jens, Benjamin Johnson, Nicolas Christin. 2010. The price of uncertainty in security games. Tyler Moore, David Pym, Christos Ioannidis, eds., *Economics of Information Security and Privacy*. Springer US, 9 – 32. URL: http://dx.doi.org/10.1007/978-1-4419-6967-5_2.

Guarro, Sergio B. 1989. Risk analysis and risk management models for information systems security applications. *Reliability Engineering & System Safety* **25**(2) 109 – 130. doi:10.1016/0951-8320(89)90027-6. URL: http://www.sciencedirect.com/science/article/pii/0951832089900276.

Guo, Yuepu, C. Seaman, R. Gomes, A. Cavalcanti, G. Tonin, F.Q.B. da Silva, A. L M Santos, C. Siebra. 2011. Tracking technical debt - an exploratory case study. *2011 27th IEEE International Conference on Software Maintenance (ICSM)*. 528 – 531. doi:10.1109/ICSM.2011.6080824.

Hansson, Sven Ove. 1999. A philosophical perspective on risk. *Ambio* **28**(6) 539 – 542. URL: `http://www.jstor.org/stable/4314951`.

Hassel, Henrik, Henrik Tehler, Marcus Abrahamsson. 2009. Evaluating the seriousness of disasters: an empirical study of preferences. *International Journal of Emergency Management* **6**(1) 33 – 54. doi:10.1504/IJEM.2009.025172. URL: `http://dx.doi.org/10.1504/IJEM.2009.025172`.

Heal, G., H. Kunreuther. 2004. Interdependent security: A general model. Working paper, National Bureau of Economic Research, Inc.

Heal, Geoffrey, Howard Kunreuther. 2007. Modeling interdependent risks. *Risk Analysis* **27**(3) 621 – 634. doi:10.1111/j.1539-6924.2007.00904.x. URL: `http://onlinelibrary.wiley.com/doi/10.1111/j.1539-6924.2007.00904.x/abstract`.

Hubbard, Douglas W. 2010. *How to Measure Anything: Finding the Value of Intangibles in Business*. 2nd ed. Wiley.

Hyman, Paul. 2013. Cybercrime: it's serious, but exactly how serious? *Communications of the ACM* **56**(3) 18 – 20. doi:10.1145/2428556.2428563. URL: `http://doi.acm.org/10.1145/2428556.2428563`.

Johnson, Benjamin, Jens Grossklags, Nicolas Christin, John Chuang. 2010. Uncertainty in interdependent security games. Tansu Alpcan, Levente Buttyán, John S. Baras, eds., *Decision and Game Theory for Security*. Springer Berlin / Heidelberg / New York, 234 – 244.

Klinger, Tim, Peri Tarr, Patrick Wagstrom, Clay Williams. 2011. An enterprise perspective on technical debt. *Proceedings of the 2nd Workshop on Managing Technical Debt*. MTD '11, ACM, New York, NY, USA, 35 – 38. doi:10.1145/1985362.1985371. URL: `http://doi.acm.org/10.1145/1985362.1985371`.

Kunreuther, H., G. Heal. 2003. Interdependent security. *Journal of Risk and Uncertainty* **26**(2) 231 – 249.

Kunreuther, Howard, Geoffrey Heal. 2012. Managing catastrophic risk. Working Paper 18136, National Bureau of Economic Research. URL: `http://www.nber.org/papers/w18136`.

Kupiec, Paul H, James M O'Brien. 1995. A pre-commitment approach to capital requirements for market risk. *Proceedings of the Federal Reserve Bank of Chicago* (May) 552 – 562. URL: `http://econpapers.repec.org/article/fipfedhpr/y_3a1995_3ai_3amay_3ap_3a552-562.htm`.

Kupiec, Paul H., James M. O'Brien. 1997. The pre-commitment approach: Using incentives to set market risk capital requirements. *Federal Reserve Bank of Richmond Economic Quarterly* **83**(1) 23 – 50.

Lambert, Diane. 1993. Measures of disclosure risk and harm. *Journal of Official Statistics-Stockholm* **9** 313 – 313.

Loch, Christoph H, Michael E Solt, Elaine M Bailey. 2008. Diagnosing unforeseeable uncertainty in a new venture. *Journal of Product Innovation Management* **25**(1) 28 – 46.

Mahajan, R. P. 2010. Critical incident reporting and learning. *British Journal of Anaesthesia* **105**(1) 69 – 75. doi:10.1093/bja/aeq133. URL: `http://bja.oxfordjournals.org/content/105/1/69.abstract`.

Marshall, Alfred. 1898. *Principles of Economics*. Macmillan.

Masys, A.J. 2012. Black swans to grey swans: revealing the uncertainty. *Disaster Prevention and Management* **21**(3) 320 – 335. doi:10.1108/09653561211234507. URL: `http://www.emeraldinsight.com/journals.htm?articleid=17038775&show=abstract`.

McKinsey & Company Inc, Tim Koller, Marc Goedhart, David Wessels. 2010. *Valuation: Measuring and Managing the Value of Companies, 5th Edition*. 5th ed. Wiley.

Michel, Olivier. 2007. There's plenty of room for unconventional programming languages or declarative simulations of dynamical systems (with a dynamical structure). Ph.D. thesis, Université d'Evry-Val d'Essonne. URL: `http://tel.archives-ouvertes.fr/docs/00/30/57/48/PDF/crv.pdf`.

Muermann, Alexander, Ulku Oktem. 2002. The near-miss management of operational risk. *Journal of Risk Finance* **4**(1). URL: `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=354760`.

National Science and Technology Council. 2011. Trustworthy cyberspace: Strategic plan for the federal cybersecurity research and development program. Official policy, United States Government,

Executive Office of the President National Science and Technology Council. URL: http://www.whitehouse.gov/sites/default/files/microsites/ostp/fed_cybersecurity_rd_strategic_plan_2011.pdf.

National Science Foundation. 2012. Secure and trustworthy cyberspace (SaTC) program solicitation. Solicitation NSF 12-596, National Science Foundation. URL: http://www.nsf.gov/pubs/2012/nsf12596/nsf12596.htm.

Okuyama, Yasuhide. 2007. Economic modeling for disaster impact analysis: Past, present, and future. *Economic Systems Research* **19**(2) 115 – 124. URL: http://mutex.gmu.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=25192160&site=ehost-live.

Orenshtein, Tal. 2011. Branching processes and applications. URL: http://www.wisdom.weizmann.ac.il/~talo/ProbabilityStudentSeminar/pdf/WPSS-GW.pdf.

Phimister, James R, Ulku Oktem, Paul R Kleindorfer, Howard Kunreuther. 2003. Near-Miss incident management in the chemical process industry. *Risk Analysis* **23**(3) 445 – 459. doi:10.1111/1539-6924.00326. URL: http://onlinelibrary.wiley.com/doi/10.1111/1539-6924.00326/abstract.

Ponemon Institute. 2012. Cost of a data breach, 2011. Report, Ponemon Institute LLC. URL: http://www.symantec.com/content/en/us/about/media/pdfs/b-ponemon-2011-cost-of-data-breach-us.en-us.pdf?om_ext_cid=biz_socmed_twitter_facebook_marketwire_linkedin_2012Mar_worldwide__CODB_US.

Prusinkiewicz, Przemyslaw, Pavol Federl, Radoslaw Karwowski, Radomir Mech. 2003. L-systems and beyond. URL: http://algorithmicbotany.org/papers/sigcourse.2003/0-intro.pdf.

Rainer Jr, Rex Kelly, Charles A Snyder, Houston H Carr. 1991. Risk analysis for information technology. *Journal of Management Information Systems* 129 – 147.

Rees, Jackie. 2009. Do information security disclosures reflect future incidents? PhD dissertation, Purdue University. URL: http://128.210.131.47/cascade/academics/MIS/workshop/papers/wrk_031309.pdf.

Rees, Jackie, Karthik Kannan. 2008. Reading the disclosures with new eyes: Bridging the gap between information security disclosures and incidents. *Proceedings of the 7th Workshop on the Economics of Information Security (WEIS 2008)*. University College London, England, 54. URL: http://weis2008.econinfosec.org/papers/Wang.pdf.

Rose, Adam. 2004a. Defining and measuring economic resilience to disasters. *Disaster Prevention and Management* **13**(4) 307 – 314. doi:10.1108/09653560410556528. URL: http://www.emeraldinsight.com/journals.htm?articleid=871056&show=abstract.

Rose, Adam. 2004b. Economic principles, issues, and research priorities in hazard loss estimation. Yasuhide Okuyama, Stephanie Ei-Ling Chang, eds., *Modeling Spatial and Economic Impacts of Disasters*. Springer. URL: http://old.geog.psu.edu/news/images/rose_disasterbook.pdf.

Rose, Adam Z. 2009. A framework for analyzing the total economic impacts of terrorist attacks and natural disasters. *Journal of Homeland Security and Emergency Management* **6**(1). doi:10.2202/1547-7355.1399. URL: http://www.degruyter.com/view/j/jhsem.2009.6.1/jhsem.2009.6.1.1399/jhsem.2009.6.1.1399.xml.

Samuelson, Paul, William Nordhaus. 2009. *Economics*. 19th ed. McGraw-Hill/Irwin.

Sanders, William, John Meyer. 2001. Stochastic activity networks: Formal definitions and concepts. Ed Brinksma, Holger Hermanns, Joost-Pieter Katoen, eds., *Lectures on Formal Methods and Performance Analysis*, *Lecture Notes in Computer Science*, vol. 2090. Springer Berlin / Heidelberg, 315 – 343. URL: http://dx.doi.org/10.1007/3-540-44667-2_9. 10.1007/3-540-44667-2_9.

Schum, David. 2001. *The evidential foundations of probabilistic reasoning*. Northwestern University Press, Evanston Ill.

Sniedovich, Moshe. 2007. The art and science of modeling decision-making under severe uncertainty. *Decision Making in Manufacturing and Services* **1**(1-2) 111 – 136.

Society of Actuaries. 2010. A new approach for managing operational risk. Tutorial, Society of Actuaries. URL: http://www.soa.org/files/research/projects/research-new-approach.pdf.

Sohail, Tashfeen. 2006. To tell or not to tell: market value of voluntary disclosures of information security activities. PhD dissertation, University of Maryland. URL: http://drum.lib.umd.edu/bitstream/1903/4277/1/umi-umd-3938.pdf.

Sommer, Svenja C., Christoph H. Loch. 2009. Incentive contracts in projects with unforeseeable uncertainty. *Production and Operations Management* **18**(2) 185 – 196. doi:10.1111/j.1937-5956.2009.01015.x. URL: http://dx.doi.org/10.1111/j.1937-5956.2009.01015.x.

Stiennon, Richard. 2013. Categorizing data breach severity with a breach level index. Technical report, IT-Harvest, LLC and SafeNet, Inc. URL: http://www2.safenet-inc.com/securethebreach/public/pdf/Breach-Level-Index-WP.pdf.

Stigler, George J. 1950a. The development of utility theory. part 1. *The Journal of Political Economy* **58**(4) 307 – 327.

Stigler, George J. 1950b. The development of utility theory. part 2. *The Journal of Political Economy* **58**(5) 373 – 396.

Thomas, Russell Cameron. 2009a. The cost of a near-miss data breach. URL: http://newschoolsecurity.com/2009/10/the-cost-of-a-near-miss-data-breach/.

Thomas, Russell Cameron. 2009b. Total cost of security: a method for managing risks and incentives across the extended enterprise. *Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research: Cyber Security and Information Intelligence Challenges and Strategies*. CSIIRW '09, ACM, New York, NY, USA, 61:1 – 61:4. doi:10.1145/1558607.1558677. URL: http://www.csiir.ornl.gov/csiirw/09/CSIIRW09-Proceedings/Abstracts/Thomas-abstract.pdf.

Thomas, Russell Cameron. 2011a. Another critique of ponemon's method for estimating 'cost of data breach'. URL: http://newschoolsecurity.com/2011/01/another-critique-of-ponemons-method-for-estimating-cost-of-data-breach/.

Thomas, Russell Cameron. 2011b. Formal methods for modeling socio-technical innovation between adversaries. *Proceedings of the 2011 Eighth International Conference on Information Technology: New Generations*. ITNG '11, IEEE Computer Society, Washington, DC, USA, 927 – 936. doi:http://dx.doi.org/10.1109/ITNG.2011.160. URL: http://dx.doi.org/10.1109/ITNG.2011.160.

Thomas, Russell Cameron. 2011c. Is norton cybercrime index just 'Security metrics theater'? URL: http://newschoolsecurity.com/2011/02/is-norton-cybercrime-index-just-security-metrics-theater/.

Thomas, Russell Cameron, Patrick Amon. 2007. Incentive-based cyber trust: a call to action. White paper, Meritology. URL: http://meritology.com/resources/Incentive-based%20Cyber%20Trust%20Initiative%20v3.5.pdf.

Vakili, A., A.J. Aghdam, T. Esmaeili. 2011. On the use of stochastic activity networks and game theory for quantitative security evaluation. *International Journal of Computer Science* **9**.

van Asselt, Marjolein B. A., Jan Rotmans. 2002. Uncertainty in integrated assessment modelling. *Climatic Change* **54**(1) 75 – 105. URL: http://dx.doi.org/10.1023/A:1015783803445. 10.1023/A:1015783803445.

Vanderschaaf, Tjerk Woutherus. 1992. Near miss reporting in the chemical process industry. URL: http://adsabs.harvard.edu/abs/1992PhDT.....25V.

Verendel, Vilhelm. 2009. Quantified security is a weak hypothesis: a critical survey of results and assumptions. *Proceedings of the 2009 New Security Paradigms Workshop*. NSPW '09, ACM, New York, NY, USA, 37 – 50. doi:10.1145/1719030.1719036. URL: http://doi.acm.org/10.1145/1719030.1719036.

Visschers, Vivianne H. M., Ree M. Meertens, Wim W. F. Passchier, Nanne N. K. De Vries. 2009. Probability information in risk communication: A review of the research literature. *Risk Analysis* **29**(2) 267 – 287. doi:10.1111/j.1539-6924.2008.01137.x. URL: http://onlinelibrary.wiley.com/doi/10.1111/j.1539-6924.2008.01137.x/abstract.

Vugrin, Eric D, Drake E Warren, Mark A Ehlen. 2011. A resilience assessment framework for infrastructure and economic systems: Quantitative and qualitative resilience analysis of petrochemical supply chains to a hurricane. *Process Safety Progress* **30**(3) 280 – 290.

Wang, Lingyu, Anoop Singhal, Sushil Jajodia. 2007. Measuring the overall security of network configurations using attack graphs. *Data and Applications Security XXI*. Redondo Beach, CA, 98 – 112.

Wang, Tawei, Jackie Rees, Karthik Kannan. 2008. The association between the disclosure and the realization of information security risk factors. URL: `https://www.cerias.purdue.edu/assets/pdf/bibtex_archive/2009-28-report.pdf`.

Widup, Suzanne. 2011. The leaking vault 2011 - six years of data breaches. Report, Digital Forensics Association. URL: `http://www.digitalforensicsassociation.org/storage/The_Leaking_Vault_2011-Six_Years_of_Data_Breaches.pdf`.

Wierzbicki, Kevin R, Ian Dobson. 2006. An approach to statistical estimation of cascading failure propagation in blackouts. *Proceedings of the 3rd International Conference on Critical Infrastructure (CRIS 2006)*. IEEE, Alexandria, VA, 1–7. URL: `http://www.ece.cmu.edu/cascadingfailures/wierzbicki-Approach-CRISpreprint.pdf`.

Wolfram, Stephen. 2002. *A New Kind of Science*. 1st ed. Wolfram Media.

Wong, Andrew. 2008. Estimating the cost of a security breach. Working paper, Innovar. URL: `http://www.innovar.com.sg/Archives/Calculating%20the%20Cost%20of%20a%20Security%20Breach_23Feb08.pdf`.

Woolf, Steven H., Anton J. Kuzel, Susan M. Dovey, Robert L. Phillips. 2004. A string of mistakes: The importance of cascade analysis in describing, counting, and preventing medical errors. *The Annals of Family Medicine* **2**(4) 317 – 326. doi:10.1370/afm.126. URL: `http://www.annfammed.org/content/2/4/317.abstract`.

Wright, Linda, Tjerk van der Schaaf. 2004. Accident versus near miss causation: a critical review of the literature, an empirical test in the UK railway domain, and their implications for other sectors. *Journal of Hazardous Materials* **111**(1-3) 105 – 110. doi:10.1016/j.jhazmat.2004.02.049. URL: `http://www.sciencedirect.com/science/article/pii/S0304389404000937`.

Zhao, X., L. Xue, A.B. Whinston. 2009. Managing interdependent information security risks: A study of cyberinsurance, managed security service and risk pooling. *ICIS 2009 Proceedings*. Phoenix, AZ, 49. URL: `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1593137`.

Appendix A: Indicators of Impact

The following is a draft list of Indicators of Impact. The categorization scheme is not highly important, since each of the Indicators has informational value on its own and in relation to the others that might be present in a set of breach episodes.

A.1. Response and recovery

General indicator: Were there significant breach response and recovery costs, over and above an ordinary incident? *Specific indicators:*

1. Did response and recovery take significantly more time and resources than a typical incident?
2. Was a special team formed with internal staff to respond to this incident? (i.e. taking them temporarily away from their regular assignments)
3. Did the incident response, investigation or recovery efforts require extensive data collection and/or data analysis (e.g. logs), above and beyond a typical incident?
4. Did the incident response, investigation or recovery efforts require purchase of any hardware, software, or computing services, above and beyond what was budgeted?
5. Did the incident investigation and reports require efforts and resources above and beyond standard procedures and tools (i.e. GRC)?
6. Was there a forensic investigation, above and beyond what your organization would normally do?
7. Were outside consultants or contractors hired specifically to respond to this incident?
8. Was there any international communication/coordination necessary among the team(s) involved in response and recovery?
9. Did IT or Security staff miss any of their team goals or objectives because they got diverted to respond to this incident?
10. Was this incident escalated to the executive level (VP or above), requiring them to make resource decisions or to spend money?
11. Was information about this breach shared with or disclosed to outside organizations responsible for incident response or information sharing (e.g. CERT, ISACs)?
12. Was this incident announced to the general public, via press release, blog post, or other corporate communications?
13. Was there any formal notification to affected stakeholders? (consumers, employees, partner companies, etc.)
14. Did the incident require specialized communications resources? (e.g. a help center, full-time spokesperson, regular press conferences, etc.)
15. Did the breached organization pay for any recovery or restoration services for affected stakeholders?

A.2. Asset compromise or fraud

General indicator: Was there any significant damage to any information assets or was there any fraud? *Specific indicators:*

1. Was there direct financial loss? (theft, extortion, embezzlement, pirating, etc.)
2. Was there indirect financial loss? (stock pump-and-dump, service theft, counterfeiting, etc.)
3. Was the value of information assets materially reduced due to erasure, destruction, corruption, etc.? (i.e. money was spent to repair, replace, or cleanse information assets)
4. Was the damage to any information assets unrecoverable?
5. Were information assets exploited as resources? (e.g. botnet, illicit compute or storage resources, proxy, etc.)
6. Was there significant data exfiltration, and evidence that attackers made use of the exfiltrated data?

A.3. Business disruption

General indicator: Were there any significant disruptions to business operations? *Specific indicators:*

1. Was any significant internal information service unavailable for a significant amount of time?
2. Was any significant external information service unavailable for a significant amount of time?
3. Was any significant business process or function disrupted for a significant amount of time?
4. Were any major projects or business initiatives delayed or canceled due to the breach, or the resources required?
5. Due to the breach, did the breached organization fail to meet any contractual obligations with its customers, suppliers, or partners? If so, were contractual penalties imposed?
6. Aside from any attacker, was anyone fired as a result of this incident, due to evidence of negligence or incompetence prior to or during the breach incident?
7. Did the breached company terminate relationships with any partners or suppliers due to their role in the breach?
8. Due to the breach, did the breached company terminate operations in any significant line of business (e.g. product or service line, geography, market)?
9. Were top executives or the Board significantly diverted by the breach and aftermath, such that other important matters did not receive sufficient attention?

A.4. Increased costs (capital or operating costs) or other financial impacts

General indicator: Did the organization incur out-of-pocket costs directly because of the breach, beyond what was budgeted? *Specific indicators:*

1. Did any department exceed its capital or operating budget due to the breach? (i.e. requiring executive approval to spend more money than budgeted)

2. Was any estimate made of the cost of the data breach for accounting purposes? (i.e. a separate charge was made for the cost of the breach, e.g. direct costs, asset write-down, etc.)
3. Did Cost of Goods Sold increase due to the breach? (e.g. higher material costs, higher labor costs, higher service costs, etc.)
4. Did Sales or Service costs increase due to the breach? (e.g. longer sales cycles, lower conversion rate, higher staff levels in sales or service)
5. Did Overhead Costs increase due to the breach? (e.g. higher staff levels in legal, audit, security, IT, or other overhead functions, more frequent audits, more costly audits, increased professional services costs, etc.)
6. Were there higher-than-budgeted capital expenditures due to the breach? (e.g. replacing or upgrading servers, significant change to IT architecture, more costly end-point devices, etc.)
7. Due to the breach, where there changes in hiring, training, or performance review practices, resulting in higher HR costs?
8. Due to the breach, did the breached organization fail to meet their financial goals for any quarter (revenue, profits, etc.)?
9. Did the breached company suffer a decline in credit rating due to the breach?
10. Were any pending business transactions delayed, postponed, or re-priced due to the breach? (e.g. merger/acquisition, stock or debt offering, downsizing, dissolution, etc.)
11. Did the breached company suffer a sustained drop in stock price or corporate value due to the breach, as evaluated by independent financial analysts?
12. Did the breached company suffer any drop in licensing revenue due to the breach? (e.g. intellectual property, copyrighted works, etc.)
13. Did breach have any homeland or national security implications, requiring extraordinary effort or costs?