Deterrence-Driven Algorithms for Audit under the Sentinel Effect

Lina Bouayad, Balaji Padmanabhan & Kaushal Chari
Information Systems & Decision Sciences
Muma College of Business, University of South Florida
email: bp@usf.edu

Abstract

Fraud, waste and abuse are significant problems in major industries such as healthcare and manufacturing, particularly when third party payers such as Medicare are involved. Current practices for auditing fraudulent activity are based on scoring models used to select practitioners or claims that are likely to be fraudulent. These models ignore the "sentinel effect" that arises from the behavior modifications of other individuals who are aware of the audit. In addition to the direct benefits of auditing fraudulent individuals known as the audit effect, second order benefits are accrued due to this sentinel effect. Yet, current auditing algorithms do not take the sentinel effect into account. In this paper, we present deterrence-driven audit policies in the presence of audit and sentinel effects, and evaluate our algorithm using an analytical model and agent-based simulation. Our results indicate that significant reduction in healthcare costs can be achieved, while maintaining fairness, when auditing policies take sentinel effects into account.

Key words: auditing; deterrence; medical fraud; network diffusion; game theory; agent-based simulation;

1. Introduction

When there is a third party payer there is often a greater chance of inappropriate claims being submitted due to moral hazard. The recipient of the service is often not directly financially accountable. Fraudulent claims however are known to drive up costs and represent a real challenge for several industries. A few years ago, CBS news reported that "Medicare fraud - estimated now to total about $60 billion a year - has become one of, if not the most profitable, crimes in America" (CBS News 2009).

In addition to fraudulent activity, there is evidence to suggest that "waste and abuse" are common. In healthcare waste and abuse show up as unnecessary tests or up-coding the severity of a patient's visit. As defined by the National Health Care Anti-Fraud Association (NHCAA), waste is an act of negligence by service practitioners who misuse, over or under utilize services and other practices. Conversely, abuse is the use of services not professionally recognized as standards of care. Examples include inappropriate procedures and unnecessary prescription refills. In other domains, payers face a similar trend referred to as *buildup* (Tennyson and Salsas-Forn 2002).

While these are not fraudulent activities - there is still a patient who needs treatment or a product to be repaired - they are practices that cumulatively place an enormous burden on service costs. According to a report published by the Institute of Medicine, excess costs amounted to $750 billion in the year 2009 with unnecessary services totaling $210 billion, inefficiently delivered services $130 billion, and fraudulent activities $75 billion (America and Medicine 2013). Payers such as Medicare, Medicaid and private insurers could benefit enormously from effective techniques to control inappropriate billing activity.

One such technique is the auditing of claims, a common industry practice. The benefits incurred from audit are twofold, namely detection and deterrence (Tennyson and Salsas-Forn 2002). Current auditing research is mostly geared towards maximizing fraud detection through algorithms aimed at detecting the fraudsters. However it has been observed that after an audit, practitioners are also deterred from inappropriate behavior in future (Tennyson and Salsas-Forn 2002). We refer to the total benefits incurred due to the savings and reduced future billings of the audited practitioner as the "audit effect".

In addition, once a set of practitioners is audited, information about audit and sanctions diffuse to *other* practitioners and these additional "audit aware" practitioners make extra effort to ensure that billings are accurate. This change in practitioner behavior triggered by the audit and the sanction of other practitioners in the network is referred to as the "sentinel effect" (Thornton 1998).

The sentinel effect has been observed in various domains. In the case of SEC enforcement actions on a target firm, peer firms having similar "aggressive" reporting practices as the target firm tend to be deterred by the SEC enforcement actions on the target firm (Schenck 2012). The deterrence effect due to SEC enforcement action is stronger when the enforcement action is a significant event for the target firm, or when the SEC enforcement is persistent in the industry of the target firm (Schenck 2012). Jennings et al. (Jennings, Keida, and Rajgopal 2011), also report the deterrence effect of SEC enforcement on peer firms. The peer firms in the industry gain additional knowledge about SEC's activity through the enforcement actions and adjust their subjective probability for attracting SEC actions against them (Jennings, Keida, and Rajgopal 2011). Securities class action law suit against a target firm can also deter reporting irregularities at peer firms (Jennings, Keida, and Rajgopal 2011).

In the audit area there is work on detection that also discusses the deterrence benefits. Dionne et al. (2009) discusses an analytical model in developing an optimal auditing strategy in the context of insurance fraud. Their work recognizes both detection and deterrence benefits, but refers still to the direct audit affect. There is work on diffusion in fraud detection, but is limited to fraudulent behavior diffusing through a network. Vlasselaer et al. (2016) use a network model to show the propagation of fraud among firms, modeling a scenario where sometimes fraudsters are influenced by others to commit fraud.

To the best of our knowledge, current audit practices however fail to capture the audit effect and the sentinel effect that occurs in the practitioner network due to the audit information itself diffusing across the network. This can lead to deterrence not just by the audited practitioner, but also by others in the network. The reason this is particularly likely to be effective in healthcare is due to the prevalence and magnitude of waste and abuse in the system. While fraud itself is relatively rare, waste and abuse is not. To our knowledge, this paper is the first to use the sentinel effect of deterrence provided by audit information

diffusion over a network to design an auditing algorithm to reduce fraud, waste and abuse in healthcare. This is a significant contribution to the literature on audit and information systems.

In a different context to ours, information security research has shown that factors such as awareness of audit and sanctions can lead to deterrence of access policy violations. Our model of deterrence due to the awareness of audit information is consistent with the findings in this related area.

Vance et al. introduce "Accountability Theory" (Vance, Lowry, and Eggett 2013; Vance, Lowry, and Eggett 2015), which proposes that having to explain or provide reasons for behavior can in and of itself lead to different behavior that could help accountability. In their work Vance et al. broadly consider the role of design artifacts to create greater accountability. Through such mechanisms – they identify four (identifiability, electronic presence, awareness of audit and awareness of logging) - organizations can benefit from improved internal security.

D'Arcy et al. (D'Arcy, Hovav, and Galletta 2009) show that, in addition to general security-related awareness, sanctions make a difference in deterrence too. Interestingly, they show that the perception of sanctions can create greater deterrence that certainty. This applies well to the healthcare audit context since sanctions vary in amount and severity and is generally not known ex-ante. Hence audit awareness combined with sanctions uncertainty can be an important mechanisms for deterrence. While we do not model sanctions uncertainty, it is potentially, a fruitful area for future research in this industry. Dionne et al. (Dionne, Giuliano, and Picard 2009) recognizing audit as a deterrent, observe that deterrence requires auditing some claims even when the investigation cost exceeds the expected benefits accrued to the insurance provider from verification.

In order to study the effectiveness of new auditing algorithms in these domains we need to model how information diffuses as well as how it then affects practitioner behavior. Hence we first present a model of stochastic audit information diffusion and behavioral change based on prior literature and validate this with domain experts. We then examine prior audit algorithms and present a novel algorithm to maximize the audit and sentinel effects. Since real audit data is confidential and therefore hard to obtain, we test and evaluate our algorithm using an agent-based simulation under stochastic information diffusion in different

network structures. Results from an auditing game between the insurance provider and the service practitioner are presented and discussed following the simulation results. The problem setup, algorithm presented and analytical results are all unique research contributions. The network-based auditing algorithm that takes into consideration both the direct audit effect and the indirect sentinel effect is particularly a key contribution of this paper. In addition, the potential of such techniques to reduce the high service costs suggests important practical applications in healthcare and potentially several other industries as well.

2. Underlying Theory and the Diffusion-Deterrence Model

Drawing from prior literature we develop here a model for audit information diffusion and behavior change in healthcare. We then validate this model through extensive interviews of medical doctors and auditing experts.

The Practitioner Network

Service practitioners are socially connected in a network where information and behavior are prone to propagate. Such social relationships are well-known to impact innovation adoption (Mahajan, Muller, and Bass 1991), knowledge transfer (Sales, Estabrooks, and Valente 2010), and behavioral change (Centola 2010). Information diffusion has been studied in psychology, marketing and other disciplines for several decades. The early Bass Model (Bass 1969), unfortunately also known for a typo in its title, linked product adoption through diffusion at a high level but did not consider the exact methods of how information could diffuse among individuals. As Katona et al. (Katona, Zubcsek, and Sarvary 2011) notes there has been a push in the last few years to model how exactly information can diffuse by explicitly looking at network structures that link people and examining how this affects diffusion. Hill et al. (Hill, Provost, and Volinsky 2006) provide one of the early empirical support for network-based models of "word of mouth" diffusion. Prior research literature (Guille et al. 2013; Mahajan, Muller, and Bass 1991; Valente 2005) provides excellent reviews of the expansive literature in this area.

Important ideas from this literature do inform our models. Kossinets et al. (Kossinets, Kleinberg, and Watts 2008) discuss an interesting case of discrete events diffusing in a network and note that it is often the

case that information travels from A to B through intermediate channels or pathways rather than directly doing so. Even if individuals are connected, in real life they may not communicate frequently and considering multiple pathways is therefore important to model how information spreads. Our multi-level model of diffusion with decay is one such mechanism that models multiple and longer pathways for communicating between any two individuals. Also, Wu et al. (Wu et al. 2004) note the importance of modeling decay in information flow in social groups. Our model of decay across levels is consistent with their measure of decay based on the distance between the source to the person receiving the information.

Specifically in healthcare, a motivating domain for this work, prior research has examined how connections between service practitioners influence their prescription behavior. The seminal work of Coleman, Katz, and Menzel in 1957 (Coleman, Katz, and Menzel 1957) indicated the presence of influence among physicians. In their original study a new drug was introduced to a few physicians in four different US cities. Fifteen months later the drug was adopted by a vast majority of doctors within the same specialties. The study had shown that interpersonal relationships among doctors do in fact allow diffusion of information and consequently influence behavior. Subsequent studies (Burt 1987; Strang and Tuma 1993; Valente 1996; Bulte, Christophe, and Lilien. 2001) have explored the structure of the service practitioner network and the nature of relationship among doctors. Unlike some industries where service practitioners form more formal networks (e.g., dealers under the same manufacturing company), doctors tend to form more informal networks through geographical proximity, collaboration, and conference attendance (Ratna et al. 2008). Whether these are formal or informal is less relevant for our context. What matters is that there exists a network that allows for sharing information and business decisions. Therefore, we make some assumptions (the convention in this section will be to number and italicize all the main assumptions behind the model assumed prior to audit).

This paper does not focus on learning these networks from data or through other mechanisms. We assume such a network exist. We do consider various structural forms based on evidence from each domain. In practice too, health insurers have extensive knowledge of physician networks.

Audit Information Diffusion

In the presence of a medium for information diffusion, we posit that audit and sanction information will spread among service practitioners. Since practitioners form an offline network, information may or may not be transmitted from one node to another. Therefore, we assume that audit information between medical practitioners spreads stochastically.

1. ***We assume that once some nodes are audited, information about the audit diffuses stochastically with some decay.***

Once any node N's neighbors receive the audit information, they will move from a state of being "unaware" of the audit to an "aware" state. If node N is at "level" t; where the level represents distance from the target audited node, then directly connected neighbors at level l+1 after moving to the aware state, will diffuse the information to the neighboring nodes (i.e., practitioners) with a probability P_{l+1}. Note that any audited node is at level l =0 and the initial probability of information diffusion is P_0. Given that there is now some distance between the original audited node and this current "aware" node, and that the information is now no longer "first hand", we assume decay in information diffusion, where we expect P_{t+1} to be smaller than P_t. We therefore define a decay factor λ such that,

$$P_{l+1} = \lambda . P_l \qquad \text{Where} \qquad 0<\lambda<1 \qquad (1)$$

This process continues sequentially, where the "new" aware nodes get a shot at propagating the audit information. The process stops when there are no new aware nodes. We note that in this study, we do not model awareness due to mass media coverage of fraud cases. In such cases, we can adjust awareness universally in the network and use the models presented here.

The basic model for stochastic information diffusion presented here addresses the process by which nodes become audit aware. This does not specify yet how these practitioners then act on such information. We turn our focus next to this issue and address how the behavior change mechanism (deterrence) is modeled.

Deterrence – Behavior Change

Prior literature suggests that audit information has a significant impact on altering negative behavior. "If (a) a set of criteria was introduced, (b) a pending audit against these criteria was announced, and (c) penalties for nonconformance were established, explicitly, then a behavioral change would take place" (Churchill, Cooper, and Govindarajan 1982). That is, as we might expect, once fraudulent practitioners receive information about audit and sanction they are likely to alter their behavior.

2. Upon being audit aware, practitioners alter their claims submission behavior.

There are some interesting broader connections for this notion. Jeremy Bentham is known for the idea of panopticism (Bentham 1969), a mechanism of subtle control through surveillance. While he initially introduced this in the context of the design of a prison, it has since been applied in many other contexts to broadly capture the effect of good behavior when being watched. Anechiarico and Jacobs (Anechiarico and Jacobs 1994) specifically discuss panopticism and financial controls and note that audit leads to good behavior and deterrence, since entities are aware that they are being watched. Our assumption here is no different. Such an effect is also due to a major phenomenon, referred to as social learning.

Social learning refers to the behavioral changes caused by the observation of outcomes (Young 2009). Once a fraudulent practitioner is audited and sanctioned, neighboring practitioners would consider the containment of fraud and abuse.

Individual practitioners might of course respond to the audit information differently depending on their relationship to the audited practitioner, as well as their personal characteristics such as initial level of compliance and risk aversion. After being audited, or being aware of an audit, fraudulent practitioners are anticipated to reduce their claims deviation. Based on prior literature (Dionne, Giuliano, and Picard 2009), practitioners with low risk aversion are expected to alter their behavior with a higher probability when compared to high risk-tolerant practitioners.

2-a. ***Audit-aware fraudulent practitioners lower the amounts of claims submitted depending on their individual level of risk aversion.***

Assuming fraudulent practitioners initially increase their real amounts of submitted claims by specific mark-ups, they are likely to drop the mark-up amounts entirely after being audited, or drop mark-ups partially after receiving the audit information.

Hence, if a fraudulent practitioner is audited then:

$$Submitted\ Claims\ Amount_{final} = Submitted\ Claims\ Amount_{initial} - \text{MarkUp}_i \quad (2)$$

If a fraudulent practitioner is not audited directly, but is now audit-aware then:

$$Submitted\ Claims\ Amount_{final} = Submitted\ Claims\ Amount_{initial} - (R_i\ .\ \text{MarkUp}_i) \quad (3)$$

Where, R_i is the level of risk aversion of practitioner i, $0 \leq R_i \leq 1$. Interestingly, from conversations with domain experts, we expect non-fraudulent practitioners to also decrease the amount of submitted claims after audit. With waste and abuse exceeding 45% of excess costs ($340billion out of $750billion) (Smith, et al. 2012), waste and abuse appears to be widely spread across the network. Therefore, while non-fraudulent practitioners might only change behavior by a small fraction through reduction of waste and abuse, their deterrence is expected to be more prevalent in the network.

2-b. *With probability P_d, audit-aware non-fraudulent practitioners decrease their claims submission amount by reducing waste and abuse.*

In practice, the average extent and magnitude of waste and abuse differ by the line of work and are estimated by auditors. Therefore we formulate the deterrence from waste and abuse as follows:

With a probability P_{WOM},

$$Submitted\ Claims\ Amount_{final} = (1-\alpha).Submitted\ Claims\ Amount_{initial} \text{ where } 0 < \alpha < 1 \quad (4)$$

Where P_{WOM} represents the *extent* and α represents the *magnitude* of deterrence generated by word of mouth audit information diffusion in the network. Deterrence can therefore be directly calculated at both the individual level as well as the network level by calculating the difference between the initial and final submitted claims amount. The initial submission amount refers to the claims amount submitted by the practitioner during the billing cycle preceding the audit, and final refers to the claims amount submitted by the practitioner during the billing cycle following the audit.

High Level Set-up

The Audit Problem: Given a network of practitioners G = (V, E), where V is a set of nodes representing practitioners, E is the set of edges representing interconnections between practitioners, and $F \subseteq V$ is a set of fraudulent practitioners, select a set of practitioners $K \subseteq V$ for audit, so as to maximize (1) $|K \cap F|$ and (2) the Total Network Deterrence, which represents the drop in claims amount in the network after auditing the |K| practitioners selected by the algorithm.

We consider a network of practitioners (G) affiliated with the same insurance provider. Connections between physicians are based on composite measure of office location, specialty, geographical region, hospital privileges, and conference attendance (Ratna et al. 2008). In the network, there exists a set of existing fraudulent practitioners (F), and the rest of the network population is non-fraudulent.

As conventionally accepted in the practice, the only way to determine the state of the practitioner as being fraudulent or non-fraudulent is through audit. However, insurance providers use scoring algorithms to calculate a prior probability of fraud for each node based on individual attributes and claims submission bell curves by line of work and region. Submitted claims from fraudulent nodes are modeled by including a markup that increases billing.

The audit diffusion and behavioral change context in which this problem is to be solved is based on the assumptions below.

1. We assume that once some nodes are audited information about the audit diffuses stochastically with some decay.
2. Upon reception of audit information, practitioners alter their claims submission behavior.
    a. Fraudulent practitioners lower the amounts of claims submitted depending on their individual level of risk aversion, R_i.
    b. With a probability P_{WOM}, non-fraudulent practitioners decrease their claims submission amount by reducing a portion α of waste and abuse.

Once the network of practitioners is defined and the process of audit information diffusion and behavior change is set up, we design algorithms to select practitioners for audit.

3. Validation of the Diffusion-Deterrence Model: Evidence from the Field

Model Validation

In order to validate our diffusion-deterrence model, we conducted an extensive field study through detailed interviews. Two questionnaires were developed for this study targeting doctors and medical auditing experts. The open-ended questionnaires were developed based on key assumptions underlying our model. For content validity, the two pilot questionnaires were reviewed by medical and insurance experts and modifications were refined as needed.

Over a four month period, we interviewed medical experts: a total of 8 medical doctors and 3 medical auditing directors/managers (Table 1), who collectively had over three hundred years of experience in the field. We used a convenience sampling strategy as a seed for potential participants. We then used a snowball sampling strategy to recruit additional experts. Interviews were conducted over the telephone, recorded and transcribed for further analysis. De-identified interview transcripts are available through online appendix C.

Table 1: Data Sources

Participant	Position/Specialty	Years of Experience
Auditing Expert 1	Analytics manager and the acting audit manager for the office of compliance and business integrity	19
Auditing Expert 2	Chief Medical Officer	15
Auditing Expert 3	Medical director at a health plan	26
Physician 1	Medical Director – Internal Medicine	40
Physician 2	Nephrology	30
Physician 3	Gastroenterology	40
Physician 4	Internal Medicine	21
Physician 5	Pediatric Cardiology	25
Physician 6	Family Medicine	35
Physician 7	Pediatrics	37
Physician 8	Ophthalmology	26

We analyzed interview transcripts using an iterative approach. Through careful analysis of the manuscripts, we identified and iteratively refined key concepts discussed during the interviews. We then

extracted text segments related to each of the identified concepts. Next, we mapped the extracted concepts to our key model assumptions. Four concepts related to the practitioner network, audit information diffusion, and behavior change were identified (Table 2).

Table 2: Model Validation - Survey Findings Summary				
	Model Assumptions			
	Practitioner Network	***Audit Information Diffusion***	***Deterrence – Behavioral Change***	
	Specialty OR Hospital OR Practice OR Community	*Word of Mouth Diffusion*	*Audit Effect*	*Sentinel Effect*
Physician 1	√	√	√	√
Physician 2	-	√	√	x
Physician 3	√	√	√	√
Physician 4	√	√	√	√
Physician 5	-	x	-	-
Physician 6	√	√	√	√
Physician 7	√	√	√	√
Physician 8	-	x	√	x
Auditing Expert 1	√	√	√	√
Auditing Expert 2	√	√	√	√
Auditing Expert 3	-	√	√	√
√: Support of the assumption x: Refutation of the assumption -: No support nor refutation of the assumption				

	Table 3: Model Validation – Sample Quotes	
	Concept	**Sample Quotes**
Practitioner Network	*Specialty*	*"People within the same specialty are also likely to impact each other because people within the same specialty are likely to experience the same issues."*
	Hospital	*"It (the interaction) could be with the hospital medical staff."*
	Practice	*"(share audit information) I would say people that are in my practice."*
	Community	*"There is always word of mouth that can come just from socializing or you know participating or being involved in communities of practice."*
Audit Information Diffusion	*Word of Mouth Diffusion*	*"A hallway conversation with colleagues who reference xyz was audited"* *"There is always word of mouth that can come just from socializing or participating or being involved in communities of practice."*
Deterrence (Behavioral Change)	*Audit Effect*	*"When a doctor is audited and found to have a problem, we educate that doctor, we expect that doctor change. And most of the time they do. If they don't then we take action"*
	Sentinel Effect	*"I hope so. That's the reason that people do audits. The intent is not so much to address that one claimant, the one provider who is submitting claims that are of concern, but actually to have a larger effect on other providers who are submitting claims so they question themselves."* *"The other way is that we announce to large group of providers, for example a group of orthopedic doctors, that we are auditing all orthopedic claims, or doing a random audit of orthopedic claims over the next six months. Then what happens, is just by virtue of making the announcement that we are going to be doing it and that we're going to be looking, we usually get individual doctors to practice differently, and be more careful about submitting claims. We call that the sentinel effect."*

Overall, the results indicate the existence of a service practitioner network, audit information diffusion and behavioral change upon audit and/or receipt of audit information; otherwise referred to in this paper as the audit and sentinel effects. Sample quotes (Table 3) illustrate how these interview results relate to the model assumptions. A detailed report of the survey findings is available in online appendix B.

Because of the limited time permitted to conduct the interviews, our interview questionnaires only covered concepts related to our main model assumptions. However, after analysis of the transcripts, we were also able to extract a few additional interesting quotes (Table 4).

	Table 4: Additional Quotes	
	Concept	**Sample Quotes**
Audit Information Diffusion	*Diffusion Decay*	*"I think the heat of the fire is likely to be felt if it is close to you."*
	Sanction Effect	*"No, no, no. When people are sanctioned or fined they don't ever talk. They keep it to themselves."* *"I am personally not aware of anyone who told me that there was a sanction involved against him. So I'm really not aware."*
		(Do you think that there is a greater awareness of audits when there is a sanction of a fine?) "Oh absolutely. " *"Obviously when there are sanctions there is spread in the medical community. Everybody knows about what happened."* *"Definitely. Even if I don't share it, the word spreads around especially if there is a sanction."*
Deterrence (Behavioral Change)	*Sanction Effect*	*"An audit with a fine or some kind of sanction will have a significant impact."*
	Waste & Abuse Change	*"I think so because, as I said before, you don't want the audit coming to your own office. So, I myself would change my pattern if I'm unintentionally doing something wrong then I will correct it."*

When talking about audit information diffusion, an expert mentioned that "*the heat of the fire is likely to be felt if it is close to you*," pointing to diffusion decay. Other quotes also indicated the possibility that practitioners would reduce waste and abuse after receiving the audit information: "*I think so because, as I said before, you don't want the audit coming to your own office. So, I myself would change my pattern if I'm unintentionally doing something wrong then I will correct it.*"

Findings related to the sanction effect on audit information diffusion were inconsistent. While some experts asserted that audit resulting in sanction would diffuse more "*Obviously when there are sanctions there is spread in the medical community. Everybody knows about what happened*", others strongly believed that information regarding audits with sanctions would not spread at all: "*No, no, no… When people are sanctioned or fined they don't ever talk. They keep it to themselves*." Based on this we decided not to model information diffusion from a fraudulent node any different from diffusion out of a non-fraudulent node (i.e. we need more consistent insights to start modeling this differently, which we did not find from the interviews).

Model Extension

Analysis of interview transcripts also indicated an important model extension. Six out of the eleven participants mentioned the fact that service practitioners receive broadcast messages including audit information through various media outlets such as medical websites, webinars, newsletters and publications. A comment from one of our medical auditing experts stated:

> "*At my health plan, we call it "provider update". Different health plans have different ways of communicating. That's a monthly conversation. That's actually an e-mailed newsletter. In that newsletter, every month, we are simply relaying the fact that we are doing audits in various positions throughout hospitals or some individual provider related to a particular topic, to make them aware of the fact that are looking and checking the claims that have been billed to verify that they're accurate."*

These broadcast messages also vary in terms of specificity. One broadcast message for example could be related to coding of a specific procedure and would therefore be relevant to only a few practitioners in the medical community as stated during one of the interviews:

> *"I think that if the medical society states that some service practitioners have been sanctioned because they were up-coding their E&M services by copying and pasting the review system in the EMR, that's a very specific instance from a reliable source that tells them that they better go back*

and check their EMR and specifically ask for a review system as opposed to copying and pasting."

While other broadcast messages could be very general such as a warning about opioid over-prescription and would therefore be relevant to a wider audience. An example mentioned by an auditing expert is as follows:

"I'm going to take a public example, and that's something that was in the news, the over-prescription of opioids. That was in the inspector general report, there were congressional hearings on this topic; it was in the news. There were administrative and human resources implications for extreme outlier providers, and that kind of negative attention can have a real chilling effect; it's a positive effect. It does have an effect. Generally, individual providers are going to be more vigilant about utilization and prescription in organizations that communicate what the consequences of misbehavior can be. "

Based on these, we added an extension representing broadcast message diffusion and associated practitioner deterrence. We introduced a broadcast node representing the insurance provider that is linked to all practitioners in the network. Upon audit of fraudulent or highly deviant practitioners the broadcast nodes sends a message to all practitioners in the network. Even though it is received by all nodes in the network, the broadcast message itself is assumed to vary in terms of its level of specificity $S_{\text{Broadcast}}$. Depending on its specificity the broadcast message is only relevant to a subset of practitioners in the network. Practitioners are therefore anticipated to alter their claims submission behavior with a probability $P_{\text{Broadcast}}$. A highly specific message affects very few practitioners, while a less specific message may apply to a broader population. Hence $P_{\text{Broadcast}}$ can be modeled as 1- $S_{\text{Broadcast}}$. Similar to the deterrence modeled upon reception of audit information through word of mouth, fraudulent practitioner aware of audit through broadcast are expected to change depending on their level of risk aversion as modeled below.

$$\textit{Submitted Claims Amount}_{\text{final}} = \textit{Submitted Claims Amount}_{\text{initial}} - (R_i \,.\, \text{MarkUp}_i) \quad (5)$$

Non-fraudulent practitioners aware of audit through broadcast are expected to reduce a portion of their waste and abuse as follows:

$$Submitted\ Claims\ Amount_{final} = (1- \beta).Submitted\ Claims\ Amount_{initial} \text{ where } 0 < \beta < 1 \quad (6)$$

Therefore in summary,

1. We assume that once a fraudulent or highly deviant node is audit, a broadcast messages with a specificity $S_{Broadcast}$ is sent deterministically to nodes in the network *where* $0 < S_{Broadcast} < 1$.
2. Upon reception of audit information, practitioners alter their claims submission behavior.
    a. Fraudulent practitioners lower the amounts of claims submitted depending on their individual level of risk aversion, R_i.
    b. With a probability $P_{Broadcast}$, non-fraudulent practitioners decrease their claims submission amount by reducing a portion β of waste and abuse *where* $P_{Broadcast} = 1 - S_{Broadcast}$.

Table 5 below contains a summary of variables used in the model. All of the variables mentioned in the table are entered by the insurance provider (inputs) except for K, which represents the output generated by the deterrence algorithm. The table also includes examples of how each of the variables can be estimated by the insurance providers in the practice.

Table 5: Variables used in the process

Variable	Description	Example of insurance provider's proxy
P_l	The probability of audit information diffusion at level l	Can be estimated using third party surveys similar to the ones used in our formal interviews.
λ	Diffusion decay	Can be estimated using third party surveys similar to the ones used in our formal interviews.
R_i	The Risk Aversion of node i	Can be estimated using longitudinal claims data.
P_d	The probability of deterrence on non-fraudulent nodes.	Estimated extent of waste and abuse by specialty and region based on longitudinal audits' data.
P_{WOM}	The expected deterrence probability of non-fraudulent nodes through word of mouth audit awareness.	Estimated extent of waste and abuse by specialty and region based on longitudinal audits' data.
α	The magnitude of deterrence of non-fraudulent nodes through word of mouth audit awareness.	Estimated magnitude of waste and abuse by specialty and region based on longitudinal audits' data.
$S_{Broadcast}$	Specificity of the broadcast message	Estimated based on the number and amount of audit cases specified in the broadcast message through longitudinal broadcast messages.
$P_{Broadcast}$	The expected deterrence probability of non-fraudulent nodes through broadcast audit awareness.	Estimated percentage of deterring practitioners upon reception of broadcast message through longitudinal broadcast messages and claims data.
β	The magnitude of deterrence of non-fraudulent nodes through broadcast audit awareness.	Estimated magnitude of waste and abuse deterrence upon reception of broadcast message through longitudinal broadcast messages and claims data.
P^f_i	The prior fraud probability of node i	Generated by the insurance provider's detection algorithm.
X	Fairness Threshold	Set by policy

DV_i	The expected deterrence value of node i	Calculated using expected deterrence amount of the practitioner and his/her neighbors.
F	Set of fraudulent nodes in the network	Estimated by the insurance provider using their detection algorithms.
K	Set of nodes selected for audit	Generated by the deterrence algorithm

4. Audit Algorithms - Background

Investigating all claims submitted by practitioners is often not cost-effective. Rather insurance providers employ different procedures to select some practitioner claims for audit. In practice, several independent entities provide audits to detect and/or prevent fraudulent behavior. According to the National Healthcare Anti-fraud Association (NHCAA), medical investigations include on-site audits, equipment audits, mail-order reviews, claims check, analytics and reporting, product verification, compounding, member lock-in, physician profiling, and credentialing programs. These programs aim at detecting and recovering known types of improper activity.

Other practices use predictive statistical models utilizing scoring rules, anomaly detection, predictive modeling and social network analysis techniques in order to optimize detection. In addition, these algorithms are often not fully publicized because of 1) their proprietary nature and 2) the risk associated having such information being transmitted to fraudulent practitioners.

Academic auditing research on the other hand has primarily focused on developing scoring algorithms to determine suspicious practitioners. Based on individual-specific variables and claim-related signals (Dionne, Giuliano, and Picard 2009), these algorithms calculate suspicion indices (scores) used to select claims for investigation. Traditionally, audit algorithms generate a fraud probability associated with every claim and/or practitioner. The pre-selected practitioners are then targeted for audit. Hence, work in the general area of fraud detection can be relevant for audit algorithms (Phua et al. 2012; Fawcett and Provost 1997). However, there is little work that has explicitly addressed fraud detection models for practitioner audit applications in healthcare.

Statistical and machine learning algorithms have been used to detect financial statement fraud. Perols (Perols 2011) uses six algorithms: Decision Tree (J48), Support Vector Machine (SMO), Back propagation

Artificial Neural Network (Multilayer Perceptron), Logistics Regression and ensemble methods -Stacking and Bagging to identify fraudulent firms. Based on the computation experiments, Logistics Regression and Support Vector Machines performed better than the others. Cecchini et al. (Cecchini et al. 2010) use a methodology based on State Vector Machines for detecting management fraud.

More recently, detection algorithms have steered toward analysis of providers' relationships to identify organized fraudulent activity (Akoglu and Faloutsos 2013; Liu et al. 2015). In several domains, fraud have shown to propagate amongst nodes in the network (Villanustre and Furht 2016; Vlasselaer et al. 2016). Using social networks analytics (a combination of graph analysis and anomaly detection analyses), these algorithms uncover hidden clusters of fraudulent nodes (Liu et al. 2015; Subelj, Furlan, and Bajec 2011; Villanustre and Furht 2016).

Network-based detection methods illustrate the potential of leveraging nodes' relationships to improve auditing algorithms. These methods focus on taking into account the existing link between nodes committing fraudulent activity (negative behavior) to improve detection. However, to the best of our knowledge, no auditing algorithm has been designed to take into account/maximize the potential spread of deterrence (positive behavior) generated by the audit.

One interesting stream of work in IS has been active learning, which can be used by auditing practitioners to acquire their "fraud" flag (Saar-Tsechansky and Provost 2007). In such cases, the information acquired from audit is traditionally used to improve the detection model. However Saar-Tsechansky & Provost (Saar-Tsechansky and Provost 2007) and Kong and Saar-Tsechansky (Kong and Saar-Tsechansky 2014) address decision centric active learning, where the information acquired is also considered based on its utility in a broad sense. Our work here can be viewed as one specific kind of active learning that suggests a new measure of utility that includes the audit effect and the sentinel effect, prevalent in certain kinds of major sectors such as healthcare. Our approach is therefore a contribution to the active learning literature as well. Indeed recent views have recognized the importance of audit in deterring fraudulent behavior rather than simply detecting it. Tennyson et al. in 2002, note that "the primary role of

auditing of an optimally designed system is the deterrence of buildup rather than its detection." (Tennyson and Salsas-Forn 2002). This is consistent with the ideas presented next.

5. Audit under the Sentinel Effect

Rather than looking to detect and minimize fraudulent behavior alone we aim to also maximize deterrence by considering the behavior change of many audit-aware practitioners in the network. This can reduce the costs in the system not only due to fraud, but also due to the reduction in waste and abuse that are likely to be more prevalent. This perspective guides the design of our deterrence-based algorithm.

4.1 Algorithm Outline

Considering the fact that audit information (1) diffuses in the network, and (2) triggers deterrence, the audit of service practitioners could be classified as an influence maximization problem. The problem is then to target the set of practitioners for audit that produce the largest deterrence cascade. Though similar to models such as the Independent Cascade Model developed by Kempe et al. 2003 (Kempe, Kleinberg, and Tardos 2003), our model is different in many aspects. First, the practitioner network is composed of different *types* of nodes - fraudulent or non-fraudulent. Second, the diffusion of audit information decays over node transmission. Third, once a practitioner changes behavior, others in the network cannot observe this change; and thereby cannot be influenced to similarly deter. Rather, deterrence occurs following a two-step process. First, audit information diffuses to practitioners in the network. Then, practitioners alter behavior (deter) depending on their individual characteristics such as compliance category and risk aversion. Due to these reasons, existing influence maximization algorithms do not apply directly in this auditing context, although our approach presented here does use the salient ideas in such algorithms.

We consider a service practitioner network where (1) practitioners are affiliated with the same insurance provider, (2) practitioners are of two types (fraudulent and non-fraudulent), and (3) the insurance provider is aware of the practitioner network and the claims amount distribution (they routinely gather and analyze such information). In the practitioner network, we select practitioners for audit such that deterrence is maximized (as formulated in Section 2).

The audit problem under the sentinel effect is an influence maximization problem in which nodes are heterogeneous, probabilities of diffusion decay over connection levels, and behavioral change depends on individual node characteristics. Since the Influence Maximization problem is NP-hard (Kempe, Kleinberg, and Tardos 2003), the audit problem here is also NP-hard. Hence we develop a *Greedy Deterrence Heuristic*.

By making use of the structure of the practitioner network, the "Greedy Deterrence Heuristic" aims at maximizing overall deterrence. The algorithm is as follows:

1. Estimate the fraud probability for each practitioner in the network
2. Filter out practitioners with fraud probabilities below the "Fairness Threshold"
3. Calculate the Network Deterrence value (the individual drop in claims amount after audit) of each practitioner in the network (described more below).
4. Select a set of practitioners K for audit such that
   (a) the total network deterrence value is maximized, and
   (b) the overlap among diffusion effects is minimized.

Note that our algorithm is similar to the active learning approach in the sense that a) nodes are selected for audit sequentially, and b) nodes are selected based on the expected deterrence they would generate in the network. That is, a specific node can have a relatively smaller fraud probability – score (still higher than the threshold) and yet be selected for audit because of its high connectivity leading to higher expected deterrence in the network.

However, it is important to note that only practitioners with prior fraud probability scores higher than a pre-defined threshold are considered for audit. That ensures both the fairness and legality of the auditing process. It is also a way to integrate the traditionally implemented fraud propensity scores into our deterrence-oriented methods; thereby providing both detection and deterrence benefits.

Another major distinction of our approach is that we are taking into consideration the relationships between nodes in the network. More specifically, nodes are selected for audit in a way that minimizes overlap of expected nodes reached through audit information diffusion.

Figure 1: The Greedy Deterrence Heuristic Algorithm

Input: A network of practitioners G, number of practitioners k, prior fraud probabilities P^f_i for all practitioners

Output: A set of k nodes to be targeted for audit

```
1: Begin
2: while (k > 0) do
3:     for each node in the network do
4:         if (Node's Prior Fraud Probability P^f_i > Fairness Threshold χ)
5:             Calculate Expected Network Deterrence Value (DV) (Figure 2)
6:     Sort vertices based on their DV
7:     Select node i with the highest DV for audit
8:     Remove node i from the list of practitioners
9:     Remove node i's immediate neighbors from the list
10:    k = k-1
11: end while
12: END
```

4.2 Expected Network Deterrence Computation

The network deterrence value of a node varies based on how the stochastic diffusion process occurs. Since individual risk aversion, mark-ups, and magnitudes of waste and abuse are unknown to insurance providers, these cannot be used to determine the actual deterrence value of every node in the network. Hence, it is necessary to measure instead, the expected network deterrence value of a node. This can be determined computationally by averaging over values generated by running stochastic diffusion processes a large number of times from each node.

An alternative to this expensive computational measurement is an approximation for this expectation that can be computed based on the number of nodes that are expected to be audit aware at any level. Recall that the diffusion probability decays with levels and this can be used to compute the expected percentage of nodes that are aware at any level.

For example, if $P_l = 0.9$ and $\lambda = 0.5$, then 45% of level two neighbors are expected to be aware of the audit. If $100K is the total deterrence that can be expected at that level if "all" nodes are aware, then $45K is the contribution of this level to the approximation for the expected deterrence of a node. We formalize this below:

$$Expected\ Network\ Deterrence\ (i) = Expected\ Detection\ Value\ (i) + \sum_{Level=1}^{n} Deterrence\ Value_{Level}(\mathrm{i}).\ \mathrm{P}_{\mathrm{Level}-1} \quad (5)$$

where *n* refers to the diameter of the audit graph, *Deterrence Value* $_{\text{Level}}$ (i) is the amount expected to be saved if all the nodes that are at a distance of "Level" from the chosen node "*i*" are audit aware, and $\mathrm{P}_{\text{Level-1}}$ is determined from (1). If *J* is the set of nodes that are at a distance of "Level" from the chosen node "*i*", then the *Deterrence Value* $_{\text{Level}}$ (i) can also be written as $\sum_{j \in J}$ *Expected Aware Value* (j). Drawing from equations 2 through 4, the *Expected Aware Value* of each aware node "*j*" would be set as follows:

$$Expected\ Aware\ Value\ (j) = \mathrm{P}^{\mathrm{f}}(j).\left[R_j.MarkUp(j)\right] + (1 - P^f).[P_d.\alpha.Submitted\ Claims\ Amount_{initial}\ (j)] \quad (6)$$

Similarly, the *Expected Detection Value* of the audited node "*j*" could be calculated as follows:

$$Expected\ Detection\ Value\ (j) = \mathrm{P}^{\mathrm{f}}(j).[MarkUp(j)] + (1 - P^f).[P_d.\alpha.Submitted\ Claims\ Amount_{initial}\ (j)] \quad (7)$$

where P^f represents the node's prior fraud probability. Any scoring algorithm can be used to determine each node's own prior fraud probability. These scores are created based on individual attributes as well as income comparisons with bell curves by line of work and region. In the medical domain, bell curves are publically available from the Center for Medicare and Medicaid Services (CMS). Note that we do assume away the hard task of calculating the prior fraud probability, which most prior IS research has focused on through design of fraud detection algorithms. Two reasons for this are: (a) it is not the focus of this paper, we can indeed use prior fraud probability if provided by any other method as well, (b) industry uses this based on deviation from the bell curves as we have noted above.

Note that our greedy deterrence heuristic does also take a node's own detection value into account. This algorithm therefore has detection combined in it as well, since selecting a node that is fraudulent will generate detection benefits from auditing that node.

However, even though the use of mark-up amounts, risk aversion (R_j) and the exact extent (P_d) and magnitude of waste and abuse (α) values would result in very accurate estimates of the nodes' expected network deterrence, these values are usually unknown to the insurance provider, and therefore cannot be used by the algorithm to select practitioners for audit. A surrogate for these values could be set as follows, based mainly on two factors – the submitted claims amount and an estimated prior fraud probability:

$$\text{Expected Detection Value (j)} = \text{Expected Aware Value (j)} = \textit{Submitted Claims Amount}_{\textit{initial}}\ \textit{(j)}.\ P^f(j) \quad (8)$$

This surrogate is directly measurable based on known factors but clearly does not take into account the factors just discussed above. However, as we show in the results, this still provides significant value.

Calculating this expected network deterrence value for each node is at the heart of the algorithm. While the equation presents how this can be determined, there are additional steps algorithmically that we now note (for ease of exposition this was not presented in Figure 1). We consider nodes one at a time to calculate their expected network deterrence value.

First, the diameter of the practitioner graph is determined, in order to store how many possible levels there are from each node. In many real-life networks, this is known to be fairly low (Guare 1990; Watts 1999). Next we determine the actual set of nodes that can be reached from the current node at each distance value (betweenness centrality). This then provides all the information needed to compute the expected network deterrence value for each node. Common traversal algorithms such as breadth-first search (BFS) and Dijkstra's algorithm can be computationally expensive depending on the number of nodes (|V|) and the number of edges (|E|). However, newer algorithms take advantage of new accumulation techniques and multi-processing to make the process more efficient. Using the Brandes' algorithm (Brandes 2001), betweenness centrality calculations require O(|V|+|E|) space and O(|V|.|E|) and O(|V|.|E| + $|V|^2$ log |V|) time on un-weighted and weighted networks instead of the traditional complexity of O($|V|^3$) time and O($|V|^2$) space.

Using Hadoop and MapReduce, the HADI algorithm can compute the graph diameter in O(d(|V| + |E|)/M) time and O((|V| + |E|) log |V|) space, where M represents the number of machines in the MapReduce or Hadoop cluster, and d is the number of iterations required to complete the process (Tsourakakis et al. 2008).

Figure 2: The Calculate *Expected Network Deterrence Value* (DV) Function

Input: A network of practitioners G, A node *i* in G, P_t, λ

Output: Node (*i*)'s expected deterrence value (DV)

Method:

1: **Begin**
2: Calculate the graph diameter from node *i*
3: Set *MaxLevel* to the graph diameter
4: **for** (*level*=1 ; *level*< *MaxLevel*; *level*++)
5: Use Adjacency Matrix to retrieve P:node (*i*)'neighbors at level "*level"*
6: **for** (*j*=0; j<|P|; *j*++)
7: set $Deterrence\ Value_{Level}(i)$+= *Submitted Claims Amount*$_{initial}$ *(j)*. $P^f(j)$
8: end for
9: DV+= $Signal_{Level}.Deterrence\ Value_{Level}$(i). P_t . $\lambda^{Level-1}$
10: end for
11: Calculate node (*i*)'s *Detection Value* (*i*) = *Submitted Claims Amount*$_{initial}$ *(i)*. $P^f(i)$
12: DV+= *Detection Value* (*i*)
12: return DV
13: **END**

To evaluate the economic value from the algorithm, we focus primarily on the change in the overall amount of claims submitted by practitioners in the network before and after audit. This does incorporate waste and abuse reduction. We discuss in greater detail when we present the results.

6. Analytical Results

Grounding our work on prior game theory literature (Cavusoglu, Raghunathan, and Cavusoglu 2009), we set up and analyze an audit game model. We identify two players, the Insurance Provider (IC), and the Service Practitioner (SP).

Among all practitioners within the network, ε practitioners are influential (with a number of neighboring practitioners exceeding a threshold predefined by the insurance provider), while the rest are

not. Our model is similar to the "IDS-Firewall model" (Cavusoglu, Raghunathan, and Cavusoglu 2009) in that the two players represent the firm and user, where the user elects to hack or not, and the firm decides to audit or not, based on the outputs from an Intrusion Detection System (IDS) and a firewall.

For modeling purposes, we consider the case of a particular set of an IC and a SP. The SP could elect to defraud with a probability ψ. Hence, the total claims amount can include an amount of fraudulent claims ρ (mark-up). Once audited and if found to be fraudulent, the SP is imposed a penalty γ. Therefore we identify the practitioner's expected utility in Table 6.

Table 6. Service Practitioner's Payoff table

		Service Practitioner's strategies	
		Defraud	Don't Defraud
Insurance Provider's Strategies	Audit	$\rho - \gamma$	0
	Don't Audit	ρ	0

The IC handles claims submitted by the SP. In order to investigate the legitimacy of claims, the IC incurs the cost of audit (c). To select practitioners for audit we follow a two-step process (Figure 3). First, a detection algorithm is used to filter out all genuine practitioners (practitioners with *detection values* below a specific threshold). These are practitioners who have low prior fraud probabilities and submitted claims amounts within the industry's norms. Pre-selected likely fraudulent practitioners are then presented to the greedy deterrence heuristic which selects the group of practitioners with the highest *Expected Network Deterrence Value* for audit.

By combining both algorithms in this manner, the insurance provider is able to incorporate fairness in its auditing policy; thereby obtaining the benefit of targeting both fraudulent and influential practitioners. Also combining the detection and deterrence algorithms assures the adherence to some industry regulations such as those that outlaw the random audit of practitioners without a likelihood of sustained or high level of payment error (Medicare Prescription Drug Improvement and Modernization Act, 2003). The IC therefore could choose to audit (1) practitioners *targeted* by the deterrence algorithm with a probability p_1, or (2) practitioners *not targeted* by the deterrence algorithm with a probability p_2.

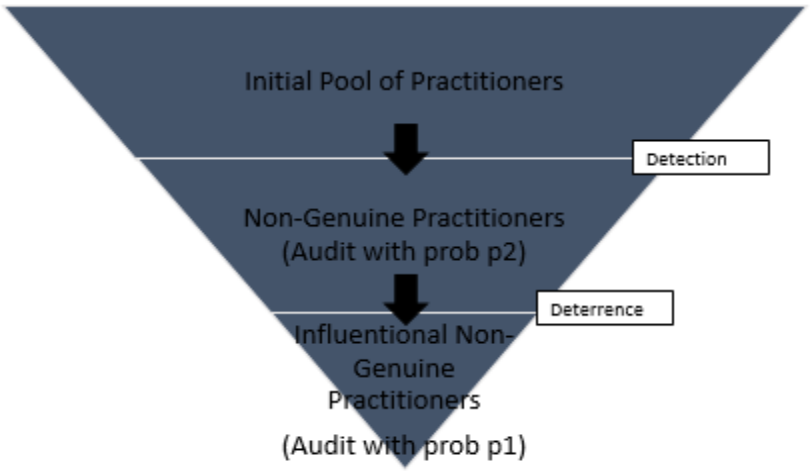

Figure 3: The Deterrence Algorithm

Within the network we differentiate between two groups of practitioners: influential and non-influential. While all audited practitioners are expected to change their behavior after audit, influential nodes which are highly connected nodes, are anticipated to trigger a larger diffusion of the audit information in the network, thereby generating more deterrence in the network.

Therefore, we denote the additional benefit from auditing an *influential fraudulent* practitioner Φ_1 and the additional benefit from auditing an *influential non-fraudulent* practitioner performing waste and abuse Φ_2. The deterrence algorithm aims at selecting influential practitioners with the highest *Expected Network Deterrence value (DV)* for audit, thereby maximizing the insurance's payoff. The insurance provider's payoffs for the different scenarios are expressed in Table 7.

Table 7. Insurance Provider's Payoff table

		Insurance Provider's Strategies			
		Don't Audit		Audit	
		Non-Influential Practitioner	Influential Practitioner	Non-Influential Practitioner	Influential Practitioner
Service Practitioner's Strategies	Defraud	-ρ	-ρ	$-\rho + \gamma - c$	$-\rho + \gamma + \Phi_1 - c$
	Do not Defraud	0	0	$-c$	$\Phi_2 - c$

In the quest for utility maximization, a service practitioner can elect to defraud, waste and abuse, or not, while the insurance provider can elect to either audit the practitioners selected by the detection algorithm

or audit the practitioner who were not tagged by the deterrence algorithm. The game is summarized in strategic format in Figure 4.

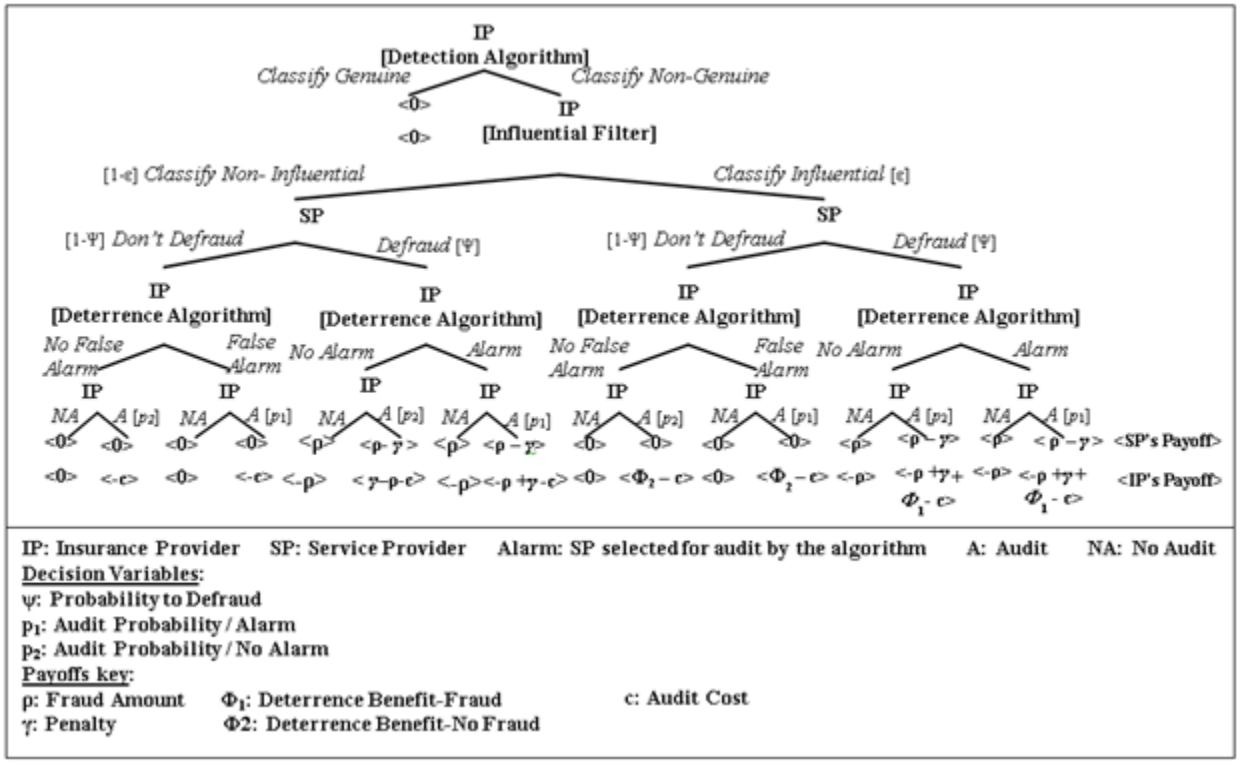

Figure 4: Game Tree

We represent the performance of the deterrence algorithm through the probabilities of true and false positives. We therefore define the following metrics:

P_F^I: The probability that the **deterrence** algorithm generates an alarm for a **fraudulent influential** practitioner.

P_F^{NI}: The probability that the **deterrence** algorithm generates an alarm for a **fraudulent non- influential** practitioner.

P_{NF}^I: The probability that the **deterrence** algorithm generates an alarm for a **non-fraudulent influential** practitioner.

P_{NF}^{NI}: The probability that the **deterrence** algorithm generates an alarm for a **non-fraudulent non-influential** practitioner.

In order to derive the Nash equilibrium for the game, we calculate both parties' expected payoffs (F) as follows:

$$F_{SP} = P_{\text{Audit/Fraud}}.((\rho-\gamma).\psi) + P_{\text{No Audit/Fraud}}.(\rho.\psi) \quad (10)$$

The Insurance Provider (IC)'s payoff is derived as follows:

$$F_{\text{IC/Alarm}} = -(1-p_1).(\rho).P_{\text{NI-Fraud/Alarm}} - p_1.(\rho-\gamma).P_{\text{NI-Fraud/Alarm}} - (1-p_1).(\rho).P_{\text{I-Fraud/Alarm}} - p_1.(\rho-\gamma-\Phi_1).P_{\text{I-Fraud/Alarm}} - p_1.(-\Phi_2).P_{\text{I-NFraud/Alarm}} - p_1.c \quad (11)$$

$$F_{\text{IC/No Alarm}} = -(1-p_2).(\rho).P_{\text{NI-Fraud/NoAlarm}} - p_2.(\rho-\gamma).P_{\text{NI-Fraud/NoAlarm}} - (1-p_2).(\rho).P_{\text{I-Fraud/NoAlarm}} - p_2.(\rho-\gamma-\Phi_1).P_{\text{I-Fraud/NoAlarm}} - p_2.(-\Phi_2).P_{\text{I-NFraud/NoAlarm}} - p_2.c \quad (12)$$

Where p_1 = P (Audit/Alarm) and p_2 = P (Audit/No Alarm)

The mixed strategy Nash equilibrium derived is represented by the following:

$$\psi_1^* = \frac{(P_{NF}^{NI}.(1-\varepsilon)+P_{NF}^{I}.\varepsilon)c - P_{NF}^{I}.\varepsilon.\Phi_2}{P_F^{NI}(1-\varepsilon)(\gamma-c) + P_F^{I}.\varepsilon.(\gamma+\Phi 1-c) - (P_{NF}^{NI}.(-c)(1-\varepsilon)+P_{NF}^{I}.\varepsilon(\Phi_{2-c}))}$$

$$p_1^* = \frac{\rho}{(\gamma.(\varepsilon.P_F^{I}+(1-\varepsilon).P_F^{NI}))} \quad \text{given} \quad p_2 = 0 \ \text{ if } \rho < (\gamma.(\varepsilon.P_F^{I} + (1-\varepsilon).P_F^{NI}))$$

$$\psi_2^* = \frac{((c-\varepsilon.\Phi_2)-P_{NF}^{NI}.c(1-\varepsilon)-P_{NF}^{I}.\varepsilon(c-\Phi_2))}{(\gamma+\varepsilon.(\Phi_1-\Phi_2)+P_F^{NI}(1-\varepsilon)(c-\gamma)-P_{NF}^{NI}.c(1-\varepsilon)+P_F^{I}.\varepsilon(c-\gamma-\Phi_1)-P_{NF}^{I}.\varepsilon(c-\Phi_2))}$$

$$p_2^* = \frac{(\varepsilon.P_F^{I}+(1-\varepsilon).P_F^{NI}).\lambda-\rho}{(\varepsilon.P_F^{I}+(1-\varepsilon).P_F^{NI}).\lambda-\lambda} \quad \text{given} \quad p_1 = 1 \text{ if } \rho > (\gamma.(\varepsilon.P_F^{I} + (1-\varepsilon).P_F^{NI}))$$

We provide all probabilities and calculations in online appendix D.

In both scenarios, corresponding to the two strategies above, the insurance firm uses a combination of detection and deterrence algorithms. The detection algorithm first selects a pool of presumed non-genuine practitioners. Afterwards, the deterrence algorithm targets influential practitioners for audit.

The first scenario describes the case where the insurance provider elects to only audit non-genuine practitioners which generates an alarm through the deterrence algorithm. This strategy is only applicable when the amount of fraud is less than the penalty imposed on audited fraudulent practitioners.

$$(\rho < \gamma.(\varepsilon.P_F^{I} + (1-\varepsilon).P_F^{NI})$$

In the case when the fraud amount exceeds the penalty imposed upon fraud, the insurance provider has to use the alternative strategy.

The second strategy consists of the insurance firm auditing all the service practitioners selected by the deterrence algorithm ($p_1 = 1$), in addition to auditing practitioners who are not targeted by the deterrence algorithm with a probability p_2.

In the alarm case, the optimal probability to defraud (ψ_1^*) is nil when the insurance provider's expected cost of auditing non fraudulent practitioners equals the expected deterrence benefit from auditing non fraudulent *influential* practitioners $\left(P_{NF}^{NI}.(1-\varepsilon)+P_{NF}^{I}.\varepsilon\right)c = P_{NF}^{I}.\varepsilon.\Phi_2$.

The optimal fraud probability is at a maximum when the insurance provider's expected payoff from auditing fraudulent practitioners is nil, i.e., $(P_F^{NI}(1-\varepsilon)(\gamma - c) + P_F^I.\varepsilon.(\gamma + \Phi_1 - c)) = 0$

That means that 1) the cost of audit (c) equals the penalty collected upon auditing fraudulent practitioners (γ) and 2) the deterrence benefit from auditing influential practitioners (Φ_1) is nil ($\gamma = c$ and $\Phi_1 = 0$).

As intuitively expected, the optimal probability to audit is a function of the loss incurred by the insurance providers for unaudited fraud cases (ρ), the penalty imposed on fraudulent practitioners (γ), and the algorithm's positive rate ($\varepsilon.P_F^I + (1-\varepsilon).P_F^{NI}$). The optimal probability to audit is expected to increase as the loss incurred from fraud increases.

Similarly, the more efficient the algorithm is at detecting fraud, and the higher the penalties imposed, the lower is the optimal probability to audit.

Special Case Scenarios – Analytical Results

Given a combination of both the detection and deterrence algorithms, the insurance provider needs to select a strategy for auditing practitioners. To avoid auditing genuine practitioners we first use the detection algorithm to filter out genuine practitioners. We then manipulate the deterrence algorithm to various profiles. It is to be noted that we set the algorithm to generate an alarm targeting different segments of practitioners. This does not necessarily mean auditing all the targeted practitioners. We examine the insurance provider's payoff values at equilibrium as derived through solving the game above.

One of the insurance provider's alternatives is to not take the sentinel effect into consideration, and target the entire pool of non-genuine practitioners selected by the detection algorithm. This scenario could be achieved by setting the deterrence algorithm to target all (non-genuine) practitioners. We therefore set both the true positive and false positive rates to 1($P_F^I = P_F^{NI} = P_{NF}^I = P_{NF}^{NI} = 1$).

With the deterrence idea in mind, the insurance provider could elect to target all influential non-genuine practitioners. After filtering out the genuine practitioners, the insurance provider could use the deterrence algorithm to target influential practitioners. This setting aims at diffusing the audit information and deterrence of neighboring practitioners. In which case $P_F^I = P_{NF}^I = 1$ and $P_F^{NI} = P_{NF}^{NI} = 0$.

Given that a considerable amount of fraud occurs though Home Health Care practitioners who were not connected to the rest of the service practitioner community (US Department of Justice 2013), the insurance provider could elect to target specifically non-influential practitioners. In our game set up, we set the deterrence algorithm to generate an alarm for all non-influential non-genuine practitioners. This scenario is represented by the following: $P_F^{NI} = P_{NF}^{NI} = 1$ and $P_F^I = P_{NF}^I = 0$.

For each scenario, we first compute the optimal probability to audit in both the alarm and no-alarm cases (Table 8).

In the first scenario, an alarm is generated by the deterrence algorithm for all types of practitioners (influential/ non-influential, fraudulent/ non-fraudulent). Therefore, no practitioners fall under the "No alarm" category.

As illustrated in Table 8, the optimal defraud probability in this case increases as the cost of audit exceeds the deterrence benefit from auditing non-fraudulent influential practitioners ($c > \varepsilon.\Phi_2$). The same probability decreases as 1) the penalty imposed on fraudulent practitioners increases, and 2) the deterrence benefit from auditing influential fraudulent practitioners exceeds the deterrence benefit from auditing influential non-fraudulent practitioners ($\varepsilon.\Phi_1 > \varepsilon.\Phi_2$). It is important to note that in this scenario, the ratio of influential practitioners in the network affects the defraud probabilities.

In the second scenario, the deterrence algorithm generates an alarm for all influential (non-genuine) practitioners ($\varepsilon = 1$). Practitioners targeted by the alarm are therefore expected to defraud as long as 1) the

cost of audit is larger than the deterrence benefit from auditing non-fraudulent practitioners in the pool and 2) the deterrence benefit from auditing fraudulent practitioners exceeds the deterrence benefit from auditing non-fraudulent practitioners. The last scenario targets all non-influential practitioners for alarm. Thus, practitioners in the alarm pool do not take into consideration, the deterrence benefit, and defraud with a probability ($\frac{c}{\gamma}$). The defraud probability in this case increases with the rise of audit cost (c), and decreases with the rise of the penalty imposed upon detecting fraud (γ).

Table 8. Service Practitioner's Optimal Defraud Probabilities

Algorithm's Profile		Optimal Defraud Probability Alarm Case (ψ_1^*)	Optimal Defraud Probability No Alarm Case (ψ_2^*)
Scenario 1	**Generate alarm for *all non-genuine* practitioners)** $P_F^I = P_F^{NI} = 1$ & $P_{NF}^I = P_{NF}^{NI} = 1$	$\frac{c - \varepsilon.\Phi_2}{(\gamma + \varepsilon(\Phi_1 - \Phi_2))}$	Not Applicable (An alarm is generated for *all non-genuine* practitioners ➔ There are no practitioners that fall in the No Alarm Case)
Scenario 2	**Generate alarm for *all Non-genuine Influential* practitioners** $P_F^I = P_{NF}^I = 1$ & $P_F^{NI} = P_{NF}^{NI} = 0$	$\frac{(c - \Phi_2)}{(\gamma + \Phi_1 - \Phi_2)}$	$\frac{c}{\gamma}$
Scenario 3	**Generate alarm for *all Non-genuine non-influential* practitioners** $P_F^{NI} = P_{NF}^{NI} = 1$ & $P_F^I = P_{NF}^I = 0$	$\frac{c}{\gamma}$	$\frac{(c - \Phi_2)}{(\gamma + \Phi_1 - \Phi_2)}$

Using the optimal defraud probability set above, we calculate the expected insurance provider payoff in each of the three scenarios.

$$F = P_{Alarm} \cdot F_{IC/Alarm} + P_{NoAlarm} \cdot F_{IC/No\ Alarm}$$

We summarize our findings in table 9 below.

Table 9. Insurance Provider's Expected Payoff at Equilibrium (F)

Algorithm's Profile		Strategy 1 ($p_1 = p_1^*$, $p_2 = 0$ and $\psi = \psi_1^*$)	Strategy 2 ($p_1 = 1$, $p_2 = p_2^*$ and $\psi = \psi_2^*$)
Scenario 1	**Generate alarm for *all non-genuine* practitioners)** $P_F^I = P_F^{NI} = 1$ & $P_{NF}^I = P_{NF}^{NI} = 1$	$F = F_{Alarm}$ $\frac{\rho.(\varepsilon.\Phi_2 - c)}{(\gamma + \varepsilon.(\Phi_1 - \Phi_2))}$	Not Applicable (An alarm is generated for *all non-genuine* practitioners ➔ There are no practitioners that fall in the No Alarm Case)

Scenario 2	**Generate alarm for *all Non-genuine Influential* practitioners** $P_F^I = P_{NF}^I = 1$ & $P_F^{NI} = P_{NF}^{NI} = 0$	$\frac{\rho.(\Phi_2 - c)}{(\gamma + \Phi_1 - \Phi_2)}$	$\frac{c.\varepsilon.(\Phi_1 - \Phi_2) + \varepsilon.\gamma.\Phi_2 + \varepsilon.\rho(\Phi_2 - c) - \rho.c}{(\rho + \gamma)}$
Scenario 3	**Generate alarm for *all Non-genuine non-influential* practitioners** $P_F^{NI} = P_{NF}^{NI} = 1$ & $P_F^I = P_{NF}^I = 0$	$F = F_{Alarm} = F_{NoAlarm}$ $\frac{-\rho(c)}{((\gamma))}$	$\frac{\rho(\Phi_2 - c) - (1-\varepsilon)(c.\rho + \gamma.\Phi_2 + c.(\Phi_1 - \Phi_2))}{(\gamma + \rho + \Phi_1 - \Phi_2)}$

To better illustrate the benefits of applying our algorithm in practice, we perform a sensitivity analysis using results from our game theory model (Online Appendix E). These scenarios show the variance in the savings amount depending on the cost of audit, the penalty imposed on fraudulent practitioners and the amount of fraud. We consider a network similar to the one used in previous sections composed of 1000 practitioners for whom the total amount of claims submitted had a mean of $500,000 and a standard deviation of $100,000. Within the network, 10% of practitioners are set to be highly connected (having more than 10 immediate neighboring practitioners). Therefore, the fraction of influential practitioners in the network (ε) is 0.1. As per the CMS (Center of Medicare and Medicaid Services), the rate of improper billing for the year 2012 (CERT) was 8.6%. We use the CERT rate in the simulation as the rate of waste and abuse in the network. We calculate the benefit from auditing an influential fraudulent practitioner (Φ1), as well as the benefit from auditing an influential non-fraudulent practitioner (Φ2). By varying the expected amount of fraud (ρ) and cost of audit (c), we look at the expected insurance provider payoff in two different worlds namely High diffusion/Low Decay, and Low diffusion/High Decay.

In addition to providing strategies for audit, these results and sensitivity analyses can even be used by insurance providers to set important factors such as penalties for fraud, waste and abuse. The analytical results reinforce some conclusions from the agent-based simulation studies, such as the value of auditing highly influential/connected practitioners in certain cases.

These are also complementary and provide new findings since they do take a different and important perspective. In the game, the setting is one where both parties are making strategic decisions given the information available. In the agent-based simulation, the agents are assumed to be fraudulent or not and

determine their billings, and the algorithm then works to determine practitioners for audit. Both perspectives are useful, and in this case, point to the value of taking the sentinel effects into account in audit algorithms.

7. Simulation Results

In collaboration with a medical auditing provider, we create an agent-based simulation, evaluate the performance of our greedy deterrence heuristic, and present simulation results.

We implement two methods for selecting practitioners for audit, 1) the deterrence heuristic presented in Section 4 and 2) a detection heuristic primarily focusing on detecting fraudulent practitioners.

Note that the detection heuristic implemented in the simulation uses the nodes' *Detection Value* described earlier to select practitioners for audit (*Detection Value* (j) = $Submitted\ Claims\ Amount_{initial}\ (j)$. $P^f(j)$)). Also, the deterrence heuristic implemented uses the adjacency matrix and a default diffusion level of 3 to calculate the *Deterrence Value* of each node in the network.

In the agent based simulation models we consider a network of 1000 practitioners in the healthcare domain. Two nodes are directly linked if there is a relationship between the two practitioners based on attributes such as co-location, common hospital privileges and physician specialty. The total amount of claims submitted has a mean of $500,000 and a standard deviation of $100,000.

Since outright fraud is relatively rare, each practitioner is assigned a random prior fraud probability following a zipf distribution. Those priors are then used to determine actual fraudulent practitioners before the agent-based simulation model runs. In order to capture the noise associated with prior fraud probabilities known to the insurance provider, we introduce a shock as follows:

$$\text{Shocked\_Prob} = (w \times \text{Fraud\_Prob}) + ((1-w) \times (\text{Random\_Prob})) \quad \text{where} \quad 0<w<1 \qquad (9)$$

"Random_Prob" is a uniform random number, while w models the quality of the insurance provider's prior knowledge of the actual fraud probabilities. When w is high, it represents a highly aware insurance provider that understands most of its practitioners well from a fraud perspective.

To assess the effects of different settings on the performance of both algorithms (deterrence and detection heuristics), we set up scenarios with varying network topologies, levels of diffusion and decay, and extents of waste and abuse.

Network Topology Influence

In order to study the impact of network topology on the algorithms' performance, we generate (1) a scale-free network using a power-law degree distribution following the Barabási–Albert model (Barabási and Albert 1999), (2) a random network using a uniform degree distribution, and (3) a clique network. While scale-free and random network are common in the literature, the clique network represents practitioners that operate under the same alliance. Same-clique practitioners could also possibly form random connections with other practitioners through conferences and geographic proximity. We therefore design a clique network specific to the growing trend of alliances in the healthcare domain. This network had several cliques connected by random connections.

Figure 5 indicates the effects of network topology on the deterrence of practitioners in the network. As seen from the figure, the likelihood of deterrence increases with the number of connections. Therefore, highly influential practitioners generate more deterrence in the network. Because the deterrence heuristic targets high-degree nodes for audit, more practitioners are aware of the audit information; thereby, deterring more practitioners from fraud waste, and abuse.

In a world of high diffusion low decay (Figure 5- A and C), audit information diffuses to multi-level neighbors to reach a large proportion of the network. When waste and abuse is prevalent in this world type, deterrence is expected to be at a maximum regardless of the network topology. When diffusion is high and decay is high (Figure 5- B and D), audit information only reaches immediate neighbors.

In a scale-free network, few practitioners are highly connected. On the other end of the spectrum, in random-uniform networks, a vast proportion of practitioners in the network have a relatively high number of immediate neighbors. In our simulation instance, the scale-free network includes 114 practitioners with 10 to 20 immediate neighbors, and about 10 practitioners with 20 to 30 immediate neighbors.

The random-uniform network on the other hand has 352 practitioners with 20 to 30 immediate neighbors and 353 practitioners with 10 to 20 immediate neighbors. Therefore, more deterrence was generated in random uniform networks when decay was low. The greedy deterrence heuristic generates about $10M more than the detection algorithm in total network deterrence (Figure 5 A and C).

The practical significance of this comparison is that deterrence based algorithms are likely to be effective when the practitioner word-of-mouth networks are scale free. Given that a wide range of networks are shown to be scale free, this is likely to be the case in practice as well.

In comparison with the network deterrence amounts realized in the scale-free and uniform networks, we observe a much tighter gap between the algorithms performances in the clique network (Figure 5 E and F). That is because the vast majority of nodes in this type of network were highly connected; Practitioners within the same clique are fully connected. Also, diffusion amongst practitioners within the same clique is very high (set in the simulation to 0.9). Therefore, while targeting highly deviant nodes for audit, the detection algorithm also achieves the unforeseen benefit of the deterrence of neighboring within-clique nodes.

However, since audit information diffused similarly regardless of the nodes fraudulent status or deviation amounts, the detection algorithm does not outperform the deterrence algorithm. Because the deterrence algorithm targets highly connected (and deviant) nodes, it tends to select nodes that are connected to more than one clique. Hence, the deterrence algorithm triggers both inter and intra clique diffusion of audit information. Also, we do not notice differences in savings amounts between high and low diffusion scenarios. That could be explained by the fact that the broadcast message enables audit information diffusion across cliques rendering word of mouth diffusion less important in theses types of networks.

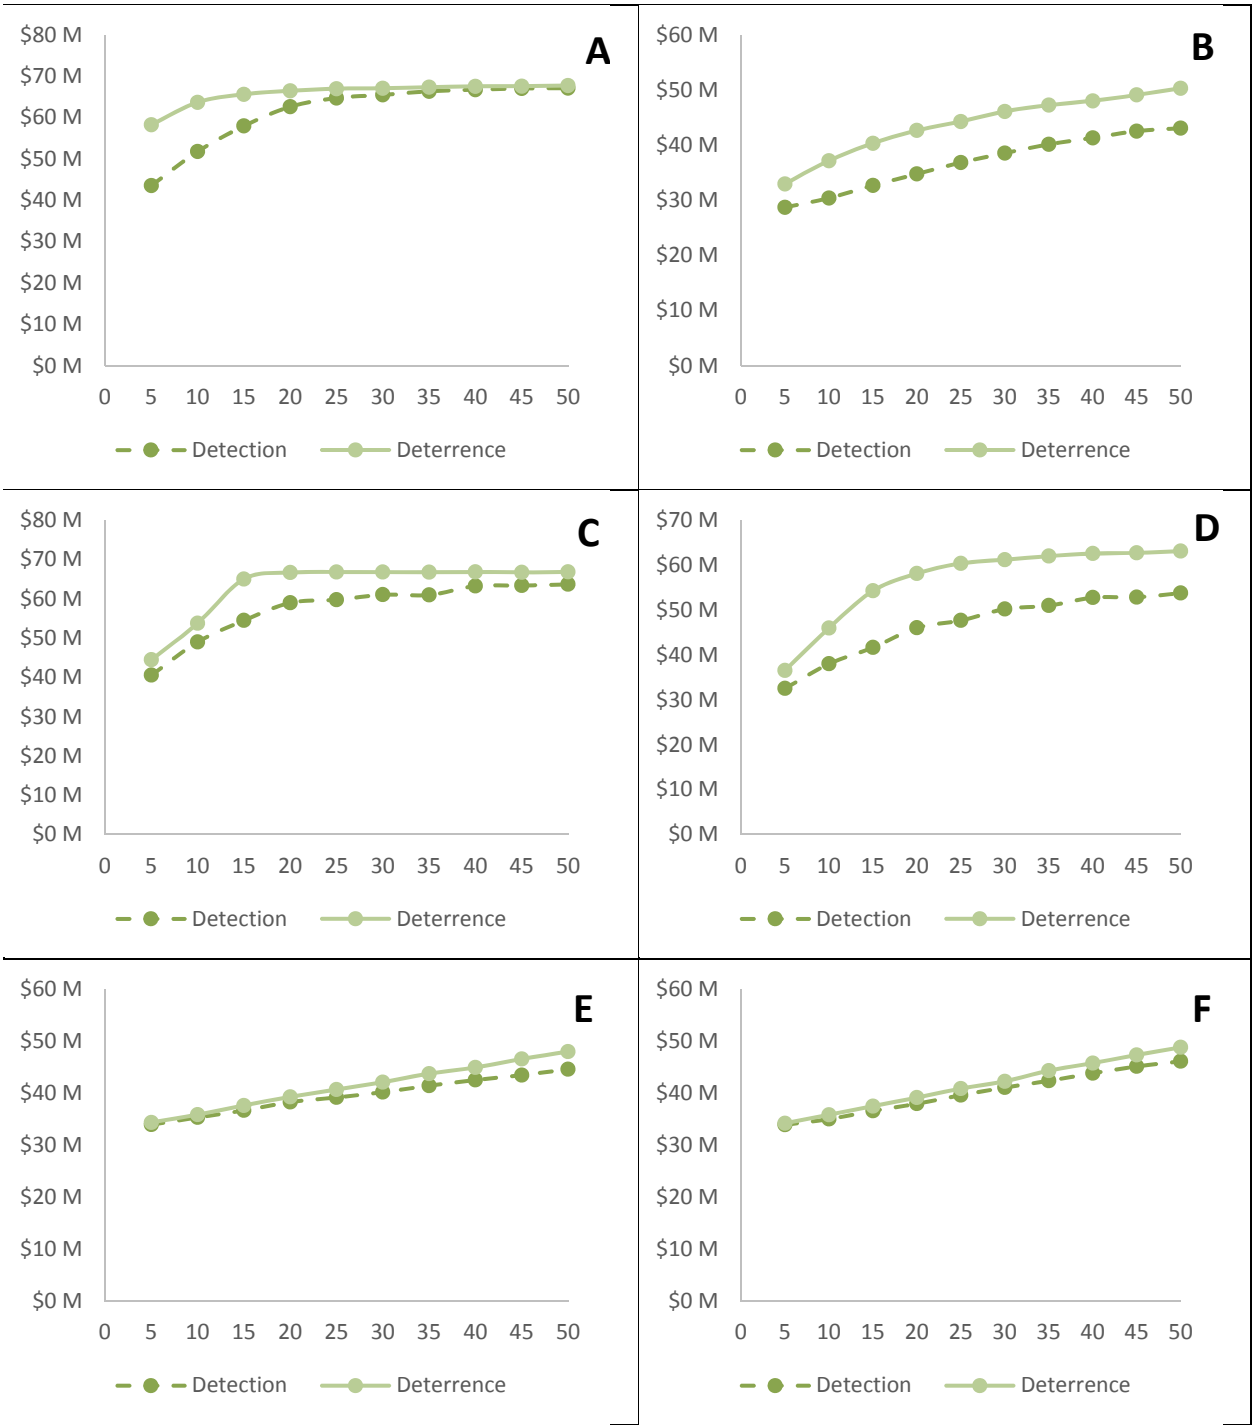

Figure 5: Y-axis: Network Deterrence Amount X-axis: Number of Audited Practitioners
(A) Scale-Free Network, High Diffusion, Low Decay
(B) Scale-Free Network, High Diffusion, High Decay
(C) Random Uniform Network, High Diffusion, Low Decay
(D) Random Uniform Network, High Diffusion, High Decay
(E) Clique Network, Low Inter-Clique Diffusion, High Decay
(F) Clique Network, High Inter-Clique Diffusion, Low Decay
Note: For all the scenarios above, P_{WOM} =0.9; $S_{Broadcast}$ =0.5; α =0.1; β =0.1

Diffusion

Network deterrence is affected dramatically by the likelihood of word of mouth diffusion in scale free networks. Figures 6B and 6D plot the network deterrence amount for both the deterrence and detection algorithms in a world of low diffusion (average diffusion probability set to 0.2), while figures 6 A and 6C plot the network deterrence amount for both algorithms in a world of high diffusion (average diffusion probability set to 0.8). In high diffusion settings, audit information diffuses to immediate neighbors with a high probability, generating vast amounts of audit awareness, and consequently large amounts of deterrence in the network. Hence in high diffusion scenarios, the deterrence heuristic had significant economic value (Figure 6 A, C).

Decay

Comparing the performance of the auditing algorithms in settings of low versus high decay shows the effects of practitioner relationships on the deterrence amounts. In the case of high decay (Figure 6 C, D), information only diffuses to immediate neighbors generating less audit awareness, and correspondingly less deterrence. In low decay settings, as the number of audited practitioners increase, audit information is transmitted to multi-level neighbors. Therefore, most practitioners in the network became audit-aware with low number of audited practitioners, thereby generating more deterrence and causing the network deterrence amounts to level off.

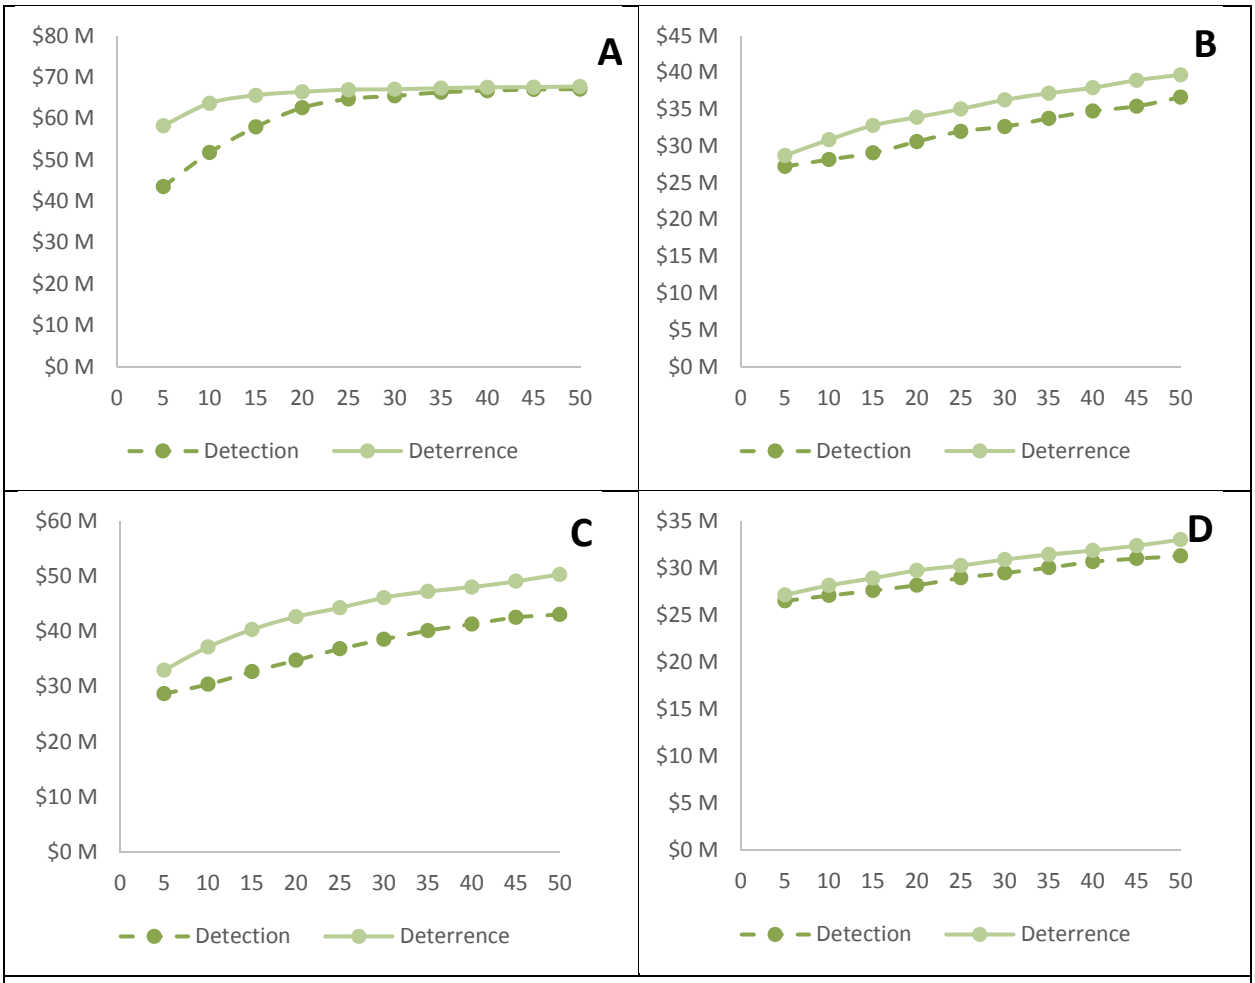

Figure 6 Y-axis: Network Deterrence Amount X-axis: Number of Audited Practitioners
(A) Scale-Free Network, High Diffusion, Low Decay
(B) Scale-Free Network, Low Diffusion, Low Decay
(C) Scale-Free Network, High Diffusion, High Decay
(D) Scale-Free Network, Low Diffusion, High Decay
Note: For all the scenarios above, P_{WOM} =0.9; $S_{Broadcast}$ =0.5; α =0.1; β =0.1

Fairness Threshold

Varying the fairness threshold shows the tradeoffs realized in each scenario. As the fairness threshold goes up, the pool of practitioners becomes restricted to individuals with high fraud prior probabilities (scores). Therefore, the difference between the detection and deterrence algorithms becomes minimal (Figure 7A) in the fairness threshold is low on the other hand, and potentially generates relatively larger economic benefits. However, it is important to note that while low fairness levels might not be acceptable in other domains, where high extra costs are mainly due to fraud, such measures might be indicated in healthcare in order to generate deterrence from waste and abuse.

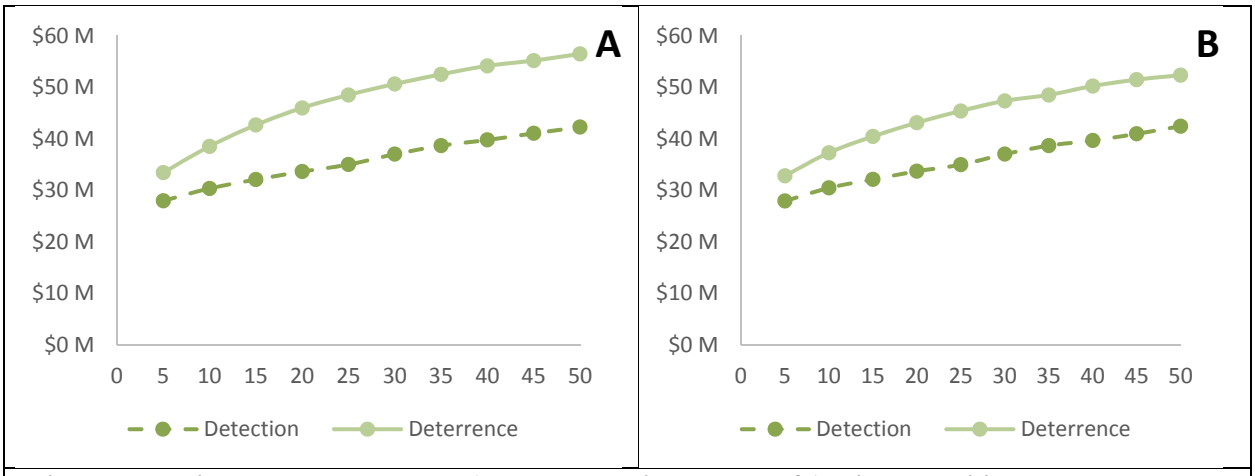

Figure 7: Y-axis: Network Deterrence Amount X-axis: Number of Audited Practitioners
Scale-Free Network, High Diffusion, High Decay (A) Fairness – 0.01 (B) Fairness – 0.15

Broadcast Specificity

Varying the specificity of the broadcast message has interesting implications in terms of performance differences between the detection and deterrence algorithms as well as the overall deterrence in the network. Highly specific broadcast messages do not translate into large deterrence amounts (Figure 8B) since they are only relevant to a small percentage of practitioners in the network. Therefore, in this specific scenario, the overall deterrence savings are relatively low. However, there is a noticeable gap in performance between the deterrence and detection algorithms since the deterrence algorithm generates more diffusion of the audit information; and thereby more deterrence.

Less specific broadcast messages on the other hand (Figure 8A) triggers a vast diffusion of the audit information; generating about $30M more savings compared to the low specificity scenario. Also, the performances of the deterrence and detection algorithm are very similar since the detection algorithm also benefits from the vast audit information diffusion triggered by the broadcast message; rendering the effect of word of mouth negligible.

These results have very interesting implications in the practice. By simply varying the specificity of their broadcast messages, insurance providers can make use of their existing platforms to potentially generate large deterrence amounts.

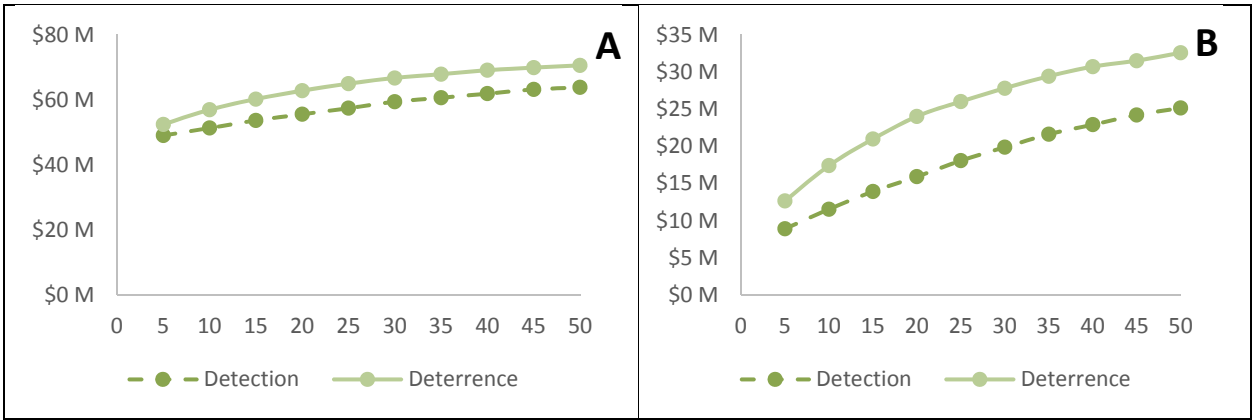

Figure 8: Y-axis: Network Deterrence Amount X-axis: Number of Audited Practitioners Scale-Free Network, High Diffusion, High Decay Fairness 0.15 (A) $S_{Broadcast}$ – 0.1 (B) $S_{Broadcast}$ - 0.9

Multi-Period Auditing

Using deterrence-focused algorithms, the long-term goal is to reduce fraud, waste, and abuse in the network. Figure 9 shows that the economic benefits generated after audit decreases over time. However, this is due to the deterrence of audited and "aware" practitioners in the network; thereby significantly reducing their submitted claim amounts in the subsequent rounds. Therefore, even if the deterrence amounts (audit benefit) are relatively lower, the actual overall excess costs paid by the insurance provider is lower.

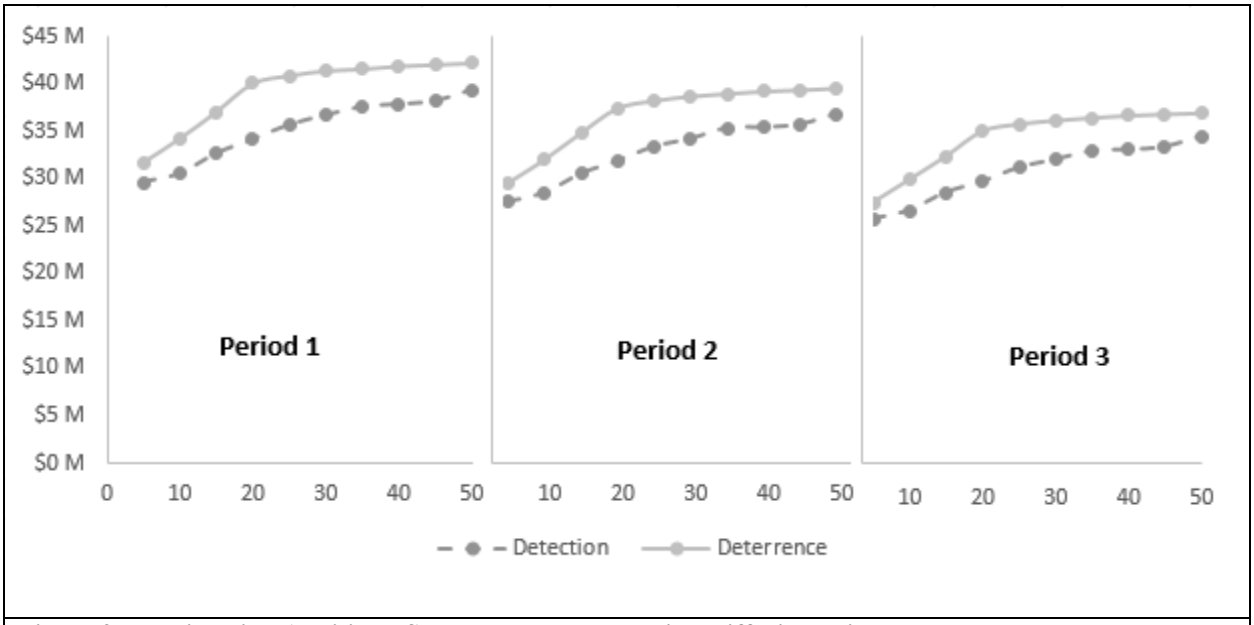

Figure 9: Multi Period Auditing - Scale-Free Network, High Diffusion, High Decay

Detection versus Deterrence

Comparing the performance of the deterrence and detection auditing algorithms in various settings suggests the applicability and trade-offs associated with their use. Overall, since our deterrence heuristic combines detection and deterrence, it performs better than the detection algorithm. In some cases, the detection and deterrence algorithms' performances are very similar. However, in specific settings, application of the deterrence heuristic can be highly beneficial since in can provider higher economic value. That is when 1) there is more diffusion of audit information (therefore significant benefits from targeting highly connected nodes), 2) fairness policies are less strict (the pool of non-genuine practitioners to select from for audit is larger), and 3) Broadcast messages are highly specific (therefore relevant to small percentage of practitioners in the network.

We note that in all the cases presented above, as is standard in simulations, the total deterrence amount was based on averaging multiple runs (1000 runs), and hence these amounts do represent expected values.

8. Conclusions

Healthcare costs have risen immensely in the past decades. In addition to fraudulent activity, the healthcare system is known to encompass a great amount of waste and abuse. Mostly, existing auditing algorithms aim at detecting fraudulent practitioners and hence takes a narrow perspective of this issue. We have presented a fundamentally new approach that showed value in deterrence-based auditing algorithms in applications like healthcare. To our knowledge, this paper is the first to incorporate sentinel effect in designing auditing algorithms, an important research contribution. We also showed how incorporating the audit effect into these algorithms addressed the modeling of waste and abuse reductions commonly seen after audit.

In addition to demonstrating the value of deterrence-oriented audit, we showed that the sentinel effect should not be taken for granted in most applications. Network structure and the diffusion mechanisms in place significantly impacted the effect of using such algorithms. Certain network topologies, such as power law networks did lend themselves well to deterrence-oriented audit, so did diffusion mechanisms such as a

high propensity to propagate to immediate neighbors with low decay. Real-world applications such as healthcare and warrant fraud are likely to have their own specific forms of networks and diffusion, and these needed to be considered before utilizing such algorithms.

Our work offers significant theoretical as well as practical contributions. This paper presents the first deterrence-based audit algorithm under network effects, a significant contribution. Further through collaboration with industry in healthcare, we have been able to design and study realistic agent based simulations, augmented by real data where possible. Our analytical results and the study of special cases provide important theoretical insights into this challenging domain.

An important issue to address in the future is the policy aspect of deterrence oriented audit. Is this a fair approach for audit? We defer this consideration to a more exhaustive treatment needed here. One approach may be to use deterrence as a second filter after a detection algorithm is employed, as done in the game theoretical model studied here. In such cases, the algorithm first flags fraudulent practitioners. Deterrence consideration is then used to select practitioners that are flagged for audit. However there are clearly many other ways in which this can be addressed and may depend on the domain and the legal frameworks that apply. A for-profit private insurer may approach trade-offs differently than, say, Medicare or the IRS. These are beyond the scope of a single study but are interesting questions for policy that we hope to examine in ongoing work.

References

Akoglu, Leman, and Christos Faloutsos. 2013. "Anomaly, Event, and Fraud Detection in Large Network Datasets." In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, 773–774. WSDM '13. New York, NY, USA: ACM. doi:10.1145/2433396.2433496.

America, Committee on the Learning Health Care System in, and Institute of Medicine. 2013. *Best Care at Lower Cost: The Path to Continuously Learning Health Care in America*. National Academies Press.

Anechiarico, Frank, and James B. Jacobs. 1994. "Panopticism and Financial Controls." *Crime, Law and Social Change* 22 (4): 361–79. doi:10.1007/BF01302925.

Barabási, Albert-László, and Réka Albert. 1999. "Emergence of Scaling in Random Networks." *Science* 286 (5439): 509–12. doi:10.1126/science.286.5439.509.

Bass, Frank M. 1969. "A New Product Growth for Model Consumer Durables." *Management Science* 15 (5): 215–27. doi:10.1287/mnsc.15.5.215.

Bentham, J. 1969. "Outline for the Construction of a Panopticon Penitentiary House." *M. Mack (Ed.) A Bentham Reade R , New York : Pegasus*.

Brandes, U. 2001. "A Faster Algorithm for Betweenness Centrality." *Journal of Mathematical Sociology*, 25(2), 163-177.

Bulte, Van den, Christophe, and Gary L. Lilien. 2001. "Medical Innovation Revisited: Social Contagion versus Marketing Effort." *American Journal of Sociology*, 106.5 : 1409-1435.

Burt, R. S. 1987. "Social Contagion and Innovation: Cohesion versus Structural Equivalence." *American Journal of Sociology*, 1287–1335.

Cavusoglu, H., S. Raghunathan, and H. Cavusoglu. 2009. "Configuration of and Interaction between Information Security Technologies: The Case of Firewalls and Intrusion Detection Systems." *Information Systems Research*, 20(2), 198-217.

CBS News. 2009. "Medicare Fraud: A $60 Billion Crime." http://www.cbsnews.com/news/medicare-fraud-a-60-billion-crime-23-10-2009/.

Cecchini, Mark, Haldun Aytug, Gary J. Koehler, and Praveen Pathak. 2010. "Detecting Management Fraud in Public Companies." *Management Science* 56 (7): 1146–60. doi:10.1287/mnsc.1100.1174.

Centola, D. 2010. "The Spread of Behavior in an Online Social Network Experiment." *Science 329.5996*, 1194–97.

Churchill, N. C., W. W. Cooper, and V. Govindarajan. 1982. "Effects of Audits on the Behavior of Medical Professionals under the Bennett Amendment." *Journal of Practice & Theory 2*, 69–90.

Coleman, James, Elihu Katz, and Herbert Menzel. 1957. "The Diffusion of an Innovation Among Physicians." *Sociometry* 20 (4): 253–70. doi:10.2307/2785979.

D'Arcy, John, Anat Hovav, and Dennis Galletta. 2009. "User Awareness of Security Countermeasures and Its Impact on Information Systems Misuse: A Deterrence Approach." *Information Systems Research* 20 (1): 79–98. doi:10.1287/isre.1070.0160.

Dionne, Georges, Florence Giuliano, and Pierre Picard. 2009. "Optimal Auditing with Scoring: Theory and Application to Insurance Fraud." *Management Science* 55 (1): 58–70. doi:10.1287/mnsc.1080.0905.

Fawcett, Tom, and Foster Provost. 1997. "Adaptive Fraud Detection." *Data Mining and Knowledge Discovery* 1 (3): 291–316. doi:10.1023/A:1009700419189.

Guare, John. 1990. *Six Degrees of Separation: A Play*. Vintage Books.

Guille, Adrien, Hakim Hacid, Cecile Favre, and Djamel A. Zighed. 2013. "Information Diffusion in Online Social Networks: A Survey." *SIGMOD Rec.* 42 (2): 17–28. doi:10.1145/2503792.2503797.

Hill, Shawndra, Foster Provost, and Chris Volinsky. 2006. "Network-Based Marketing: Identifying Likely Adopters via Consumer Networks." *Statistical Science* 21 (2): 256–76.

Jennings, J., Simi Keida, and Shivaram Rajgopal. 2011. "The Deterrence Effects of SEC Enforcement and Class Action Litigation." *Working Paper*.

Katona, Zsolt, Peter Pal Zubcsek, and Miklos Sarvary. 2011. "Network Effects and Personal Influences: The Diffusion of an Online Social Network." *Journal of Marketing Research* 48 (3): 425–43. doi:10.1509/jmkr.48.3.425.

Kempe, David, Jon Kleinberg, and Éva Tardos. 2003. "Maximizing the Spread of Influence Through a Social Network." In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 137–146. KDD '03. New York, NY, USA: ACM. doi:10.1145/956750.956769.

Kong, Danxia, and Maytal Saar-Tsechansky. 2014. "Collaborative Information Acquisition for Data-Driven Decisions." *Machine Learning* 95 (1): 71–86. doi:10.1007/s10994-013-5424-x.

Kossinets, Gueorgi, Jon Kleinberg, and Duncan Watts. 2008. "The Structure of Information Pathways in a Social Communication Network." In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 435–443. KDD '08. New York, NY, USA: ACM. doi:10.1145/1401890.1401945.

Liu, Juan, Eric Bier, Aaron Wilson, Tomo Honda, Sricharan Kumar, Leilani Gilpin, John Guerra-Gomez, and Daniel Davies. 2015. "Graph Analysis for Detecting Fraud,Waste, and Abuse in Healthcare Data." In *Twenty-Seventh IAAI Conference*. http://www.aaai.org/ocs/index.php/IAAI/IAAI15/paper/view/9705.

Mahajan, Vijay, Eitan Muller, and Frank M. Bass. 1991. "New Product Diffusion Models in Marketing: A Review and Directions for Research." In *Diffusion of Technologies and Social Behavior*, edited by Dr Nebojša Nakićenović and Dr Arnulf Grübler, 125–77. Springer Berlin Heidelberg. doi:10.1007/978-3-662-02700-4_6.

Myers, Seth A., Chenguang Zhu, and Jure Leskovec. 2012. "Information Diffusion and External Influence in Networks." In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 33–41. KDD '12. New York, NY, USA: ACM. doi:10.1145/2339530.2339540.

Perols, Johan. 2011. "Financial Statement Fraud Detection: An Analysis of Statistical and Machine Learning Algorithms." *AUDITING: A Journal of Practice & Theory* 30 (2): 19–50. doi:10.2308/ajpt-50009.

Phua, Clifton, Vincent Lee, Kate Smith, and Ross Gayler. 2012. "A Comprehensive Survey of Data Mining-Based Fraud Detection Research." *Computers in Human Behavior* 28 (3): 1002–13. doi:10.1016/j.chb.2012.01.002.

Ratna, Nazmun N., Anne Dray, Pascal Perez, R. Quentin Grafton, David Newth, and Tom Kompas. 2008. "Diffusion and Social Networks: Revisiting Medical Innovation with Agents." In *Complex Decision Making*, edited by H. Qudrat-Ullah, J. M. Spector, and P. I. Davidsen, 247–65.

Understanding Complex Systems. Springer Berlin Heidelberg. doi:10.1007/978-3-540-73665-3_13.

Saar-Tsechansky, Maytal, and Foster Provost. 2007. "Decision-Centric Active Learning of Binary-Outcome Models." *Information Systems Research* 18 (1): 4–22. doi:10.1287/isre.1070.0111.

Sales, Anne E., Carole A. Estabrooks, and Thomas W. Valente. 2010. "The Impact of Social Networks on Knowledge Transfer in Long-Term Care Facilities: Protocol for a Study." *Implementation Science* 5: 49. doi:10.1186/1748-5908-5-49.

Schenck, Kristy Marie. 2012. "The Deterrence Effects of Sec Enforcement Actions," June. https://etda.libraries.psu.edu/catalog/15397.

Strang, David, and Nancy Brandon Tuma. 1993. "Spatial and Temporal Heterogeneity in Diffusion." *American Journal of Sociology* 99 (3): 614–39.

Subelj, L., S. Furlan, and C. Bajec. 2011. "An Expert System for Detecting Automobile Insurance Fraud Using Social Network Analysis." *Paper Submitted to Expert System with Applications*. http://www.sciencedirect.com/science/article/pii/S0957417410007712.

Tennyson, Sharon, and Pau Salsas-Forn. 2002. "Claims Auditing in Automobile Insurance: Fraud Detection and Deterrence Objectives." *Journal of Risk and Insurance* 69 (3): 289–308. doi:10.1111/1539-6975.00024.

Thornton, DM. 1998. "'Sentinel Effect' shows Fraud Control Effort Works." *Journal of Health Law* 32 (4): 493–502.

Tsourakakis, C., A.P. Appel, C. Faloutsos, and J Leskovec. 2008. "HADI: Fast Diameter Estimation and Mining in Massive Graphs with Hadoop." *Carnegie Mellon University, School of Computer Science, Machine Learning Department.*

Valente, Thomas W. 1996. "Network Models of the Diffusion of Innovations." *Computational & Mathematical Organization Theory* 2 (2): 163–64. doi:10.1007/BF00240425.

Vance, A., Lowry, P. B., & Eggett, D. 2005. "Network Models and Methods for Studying the Diffusion of Innovations." In *Models and Methods in Social Network Analysis*, 98.

Vance, A., P. Lowry, and D. Eggett. 2013. "Using Accountability to Reduce Access Policy Violations in Information Systems." *Journal of Management Information Systems* 29 (4). http://www.tandfonline.com/doi/abs/10.2753/MIS0742-1222290410.

Vance, A., Lowry, P. B., & Eggett, D. 2015. "A New Approach to the Problem of Access Policy Violations: Increasing Perceptions of Accountability through the User Interface." *MIS Quarterly* 39 (2): 345–66.

Villanustre, Flavio, and Borko Furht. 2016. "Social Network Analytics: Hidden and Complex Fraud Schemes." In *Big Data Technologies and Applications*, 341–46. Springer International Publishing. doi:10.1007/978-3-319-44550-2_13.

Vlasselaer, V.V., T. Eliassi-Rad, L. Akoglu, M. Snoeck, and B. Baesens. 2016. "GOTCHA! Network-Based Fraud Detection for Social Security Fraud." *Management Science (To Appear)*.

Watts, Duncan J. 1999. "Networks, Dynamics, and the Small-World Phenomenon." *American Journal of Sociology* 105 (2): 493–527. doi:10.1086/210318.

Wu, Fang, Bernardo A. Huberman, Lada A. Adamic, and Joshua R. Tyler. 2004. "Information Flow in Social Groups." *Physica A: Statistical Mechanics and Its Applications* 337 (1–2): 327–35. doi:10.1016/j.physa.2004.01.030.

Young, P. 2009. "Innovation Diffusion in Heterogeneous Populations: Contagion, Soc...: Ingenta Connect." *The American Economic Review* 99 (5): 1899–1924(26).