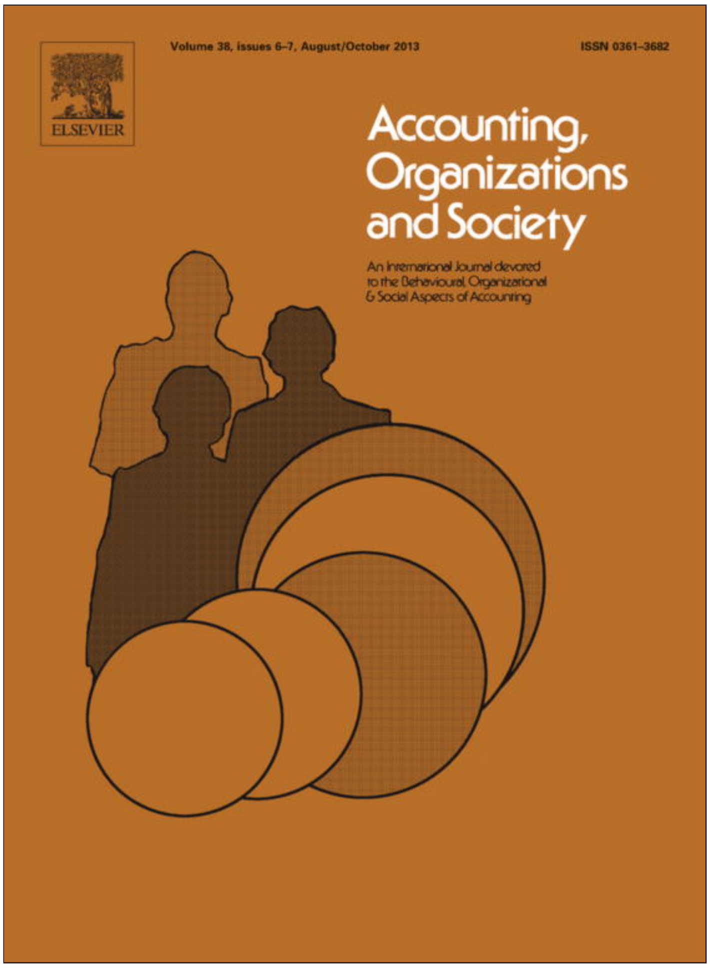
Volume 38, issues 6–7, August/October 2013
ISSN 0361-3682
ELSEVIER
Accounting, Organizations and Society
An International Journal devoted to the Behavioural, Organizational & Social Aspects of Accounting

Accounting, Organizations and Society 38 (2013) 469–483

Contents lists available at SciVerse ScienceDirect

Accounting, Organizations and Society

journal homepage: www.elsevier.com/locate/aos

Fraud dynamics and controls in organizations

Jon S. Davis *, Heather L. Pesch [1]

Department of Accountancy, University of Illinois at Urbana–Champaign, 1206 S. Sixth Street, Champaign, IL 61821, USA

A B S T R A C T

This paper develops an agent-based model to examine the emergent dynamic characteristics of fraud in organizations. In the model, individual heterogeneous agents, each of whom can have motive and opportunity to commit fraud and a pro-fraud attitude, interact with each other. This interaction provides a mechanism for cultural transmission through which attitudes regarding fraud can spread. Our benchmark analysis identifies two classes of organizations. In one class, we observe fraud tending toward a stable level. In the other class, fraud dynamics are characterized by extreme behaviors; organizations with mostly honest behavior suddenly change their state to mostly fraudulent behavior and vice versa. These changes seem to occur randomly over time. We then modify our model to examine the effects of various mechanisms thought to impact fraud in organizations. Each of these mechanisms has different impacts on the two classes of organizations in our benchmark model, with some mechanisms being more effective in organizations exhibiting stable levels of fraud and other mechanisms being more effective in organizations exhibiting unstable extreme behavior. Our analysis and results have general implications for designing programs aimed at preventing fraud and for fraud risk assessment within the audit context.

Introduction

Fraud[2] has become a popular area of inquiry among accounting academics because of the magnitude of losses (estimated by the Association of Certified Fraud Examiners in 2010 to be US$2.9 trillion worldwide) and requirements imposed on auditors to explicitly address the problem (AICPA, 2002). Research has addressed fraud risk assessment by auditors (e.g., Bell & Carcello, 2000; Carpenter, 2007; Wilks & Zimbelman, 2004), fraud detection (e.g., Cleary & Thibodeau, 2005; Hoffman & Zimbelman, 2009; Matsumura & Tucker, 1992), fraud incentives (e.g., Erickson, Hanlon, & Maydew, 2006; Gillett & Uddin, 2005), and the correlation of fraud with financial statement reporting choices and corporate governance variables (e.g., Beasley, 1996; Jones, Krishnan, & Melendrez, 2008; Sharma, 2004).

In light of the high cost of corporate fraud, one might expect considerable research activity investigating the efficacy of mechanisms designed to prevent or reduce fraud. However, despite its potential importance, a review of research reveals a striking dearth of work examining the effectiveness of various prevention mechanisms, except for the deterring role of audits (e.g., Finley, 1994; Schneider & Wilner, 1990; Uecker, Brief, & Kinney, 1981). This lack of work on prevention is likely due to fraud being a hidden crime. Because the extent of fraud is usually unknown in an organization, measuring the effectiveness of prevention mechanisms is difficult using traditional empirical methods.

Research on fraud in organizations tends to focus on either the individual or the organization (Holtfreter, 2005; Pinto, Leana, & Pil, 2008). To date, very little work has attempted to explicitly link individual behaviors in

* Corresponding author. Tel.: +1 217 300 0489; fax: +1 217 244 0902.
E-mail addresses: jondavis@illinois.edu (J.S. Davis), pesch@illinois.edu (H.L. Pesch).
[1] Tel.: +1 217 300 0514; fax: +1 217 244 0902.
[2] For the purposes of our model, we define fraud as all forms of occupational fraud, including asset misappropriation, corruption, and financial statement fraud.

0361-3682/$ - see front matter
http://dx.doi.org/10.1016/j.aos.2012.07.005

the organization to organizational outcomes within the fraud context.[3] Understanding the individual–organization link is important because a focus on either individual behavior or the organization in isolation turns a blind eye to the social process through which individuals' behaviors are influenced by the organization as a whole and vice versa. In other words, a narrow focus on individual behaviors or on the organization ignores the organization's sociology, which can have profound effects on fraud outcomes and the efficacy of fraud prevention mechanisms.

We develop a model of fraud in organizations that allows an evaluation of the relative efficacy of mechanisms designed to prevent fraud while explicitly recognizing the social processes underlying the formation of organizational norms. To develop our model, we use a method that is relatively new in accounting research: agent-based modeling (ABM). Designed to study the emergence of macro-level phenomena from micro-level interactions, ABM is well suited to address questions involving organizational outcomes (e.g., a culture of fraud) resulting from the interactions between individuals within an organization and organizational variables. The use of ABM confers an additional advantage: It allows us to gain insights into fraud even when data in organizations are censored.

Our model is comprised of an organization represented by 100 independent, heterogeneous agents (employees) and a set of simple interaction rules. Following Cressey's (1953) characterization of occupational fraud (known as the fraud triangle hypothesis), any agent in our model possessing motive, opportunity, and an attitude that frames the fraudulent act as acceptable will commit fraud. We allow agents to repeatedly interact, with an eye toward emergent aggregate fraud levels and the dynamics of fraud over time. We begin with a benchmark model in which all agents have opportunity and motive. We then modify our model to investigate the impact of mechanisms to prevent or detect fraud. We first investigate the impact of modifying the likelihood that agents perceive the opportunity to commit fraud. Next, we consider a hierarchy in which higher-level honest employees exert greater influence than lower-level employees (i.e., "tone at the top"). Then we consider the impact of asymmetric influence exerted by fraudsters relative to honest employees (which can arise as a result of ethical training, the implementation of a code of ethics, or a variety of other interventions). Finally, we consider the impact of detection and termination efforts.

Two patterns emerge from the analysis of our benchmark model, depending upon how susceptible individual agents are to social influence. When average susceptibility is low, the number of fraudsters in the organization tends toward a specific level and remains relatively stable over time. When average susceptibility is moderate to high we observe a very different pattern in which the number of fraudsters in the organization vacillates over time between extremes; either virtually no one in the organization is a fraudster or virtually everyone is.

When we consider mechanisms to prevent or eliminate fraud, we find that their impact is contingent on average susceptibility to social influence within the organization. A reduction in perceived opportunity or the introduction of influential, honest managers (tone at the top) reduces the number of fraudsters, but neither change to our model is effective in eliminating outbreaks of fraud when susceptibility is moderate or high. Allowing honest employees to be more influential than fraudsters has no qualitative effect when susceptibility is low; however, it transforms behavior when average susceptibility is moderate to high, reducing the number of fraudsters to near zero and eliminating fraud outbreaks. The contingent nature of this effect may prove important in fraud risk assessments performed by auditors. We also find that efforts to remove fraudsters can effectively reduce the number of fraudsters to near zero regardless of the level of susceptibility, but such efforts do not eliminate fraud outbreaks when susceptibility is moderate to high.

This paper continues with a brief introduction to ABM. This methodological introduction is followed by the development of a benchmark model of an organization with no interventions to prevent fraud. We then modify the benchmark model to examine the effect of anti-fraud interventions. This paper ends with our conclusions and a discussion of the implications of our analysis.

Agent-based models

We use ABM to address our research question because it confers several advantages. As noted by Epstein (2006), the method avoids several shortcomings in traditional theoretical work in the social sciences. When aggregate behavior is the research subject in traditional theory, heroic assumptions about individual behavior and the population being modeled are typically required to maintain tractability (e.g., the perfect rationality and homogeneity of individual actors). While these assumptions bear little resemblance to human behavior, it is often argued that such simplifications are necessary because no rigorous method exists that would allow their relaxation. Similarly, in traditional theory, the notion of equilibrium plays a predominant role as a solution concept; models attempt to explain social phenomenon by identifying the behavior of interest as an equilibrium. While this approach can yield valuable insights, there are limits. In many cases, phenomena of interest may involve disequilibrium dynamics or the identified equilibria may be unattainable either outright or within acceptable time scales (Epstein, 2006, p. 72). Experimental research in the social sciences that attempts to test hypotheses linking heterogeneous individual behaviors to aggregate behavioral outcomes is also challenging because of an exiguity of theoretical work linking realistic individual behaviors to aggregate phenomena. Finally, the stovepipe structure of the social science disciplines (e.g., sociology, economics, psychology, anthropology) tends to

[3] Pinto et al. (2008) discuss the ways in which corruption can exist in an organization, with a particular focus on how individual corruption can spread to the point where it becomes an organizational phenomenon. Chang and Lai (2002) use econometrics to model corrupt organizations as a pandemic arising from individual interactions. Kim and Xiao (2008) examine the link between individual behaviors and aggregate outcomes in the context of fraud in a public health care delivery program. In a broader context, Davis, Hecht, and Perkins (2003) develop a model that links individual behaviors to societal outcomes related to tax evasion.

limit the perspective that each discipline develops (much like the parable of the blind men and the elephant).[4]

The approach taken by ABM promises a solution to many of the shortcomings identified in traditional social science research. The method allows for the specification of heterogeneous actors who follow simple and realistic (boundedly rational) behavioral rules. It focuses on the emergence of macro-level phenomena from micro-level behavior and is therefore well suited for the study of social phenomena (including the emergence of fraud outbreaks in organizations).[5] Because ABM is dynamic, the natural focus of the research is on changes in behavior over time rather than on equilibria that may never be achieved. Finally, ABMs naturally cross the boundaries of the various social sciences because of the requisite focus on modeling individual behaviors (e.g., psychology), spatial and ecological considerations (e.g., anthropology), social interaction (e.g., sociology), and a wide range of other potential cross-disciplinary considerations.

A fully specified agent-based model consists of agents (e.g., individuals in an organization, traders in a market, organizations in an economy, or trees in a forest), an environment in which the agents "live," and a set of rules that govern agent interactions with other agents and their environment. Agents are characterized by a collection of internal states, endowments (knowledge and assets), and behaviors that may recognize the existence of bounded rationality. Some of these characteristics are static, while others can change in response to interactions with the environment or with other agents. Similarly, environments can be static or can change on the basis of rules or as a result of interactions with agents. Environments can be represented in a variety of ways. When location or resource placement is important in a model (an issue in anthropological or ecological research), the environment can be represented spatially, using a lattice or a landscape developed from a geographic information system. In economics, the environment of choice is a market (e.g., Gode & Sunder, 1993). In research examining sociological or sociopsychological phenomena (the development of a culture of fraud in organizations, in our instance), a social network may be favored (e.g., Breiger & Carley, 2003). Finally, a set of institutional rules is specified that defines how agents interact with each other and with their environment. As with agent and environmental characteristics, these rules can be static, dynamic, or adaptive. A simple rule for an agent might be to sell to another agent if the offer price is less than the bid price or, in the context of our fraud model, to commit fraud if the agent believes that it is acceptable behavior and if both motive and opportunity are present.

Once the elements of the agent-based model have been developed (usually instantiated in a computer program[6]), agents are allowed to repeatedly interact with each other and with the environment. As interactions occur, the research focuses on the spontaneous emergence of aggregate or group behaviors resulting from interactions at the micro (agent) level. Epstein and Axtell (1996) characterize this approach as attempting to "grow" social phenomena from the bottom up (modeling individual behavior) rather than from the top down (directly modeling aggregate phenomena). Thus, ABM does not involve macro models but, rather, models of individual behavior and social interaction that produce macro-level outcomes.

Since agent-based models are simulations, they rely on both deduction and induction. While each such model is a strict deduction and constitutes a sufficiency theorem (that proves existence but not uniqueness), multiple simulations can be thought of as a distribution of theorems that together are used to inductively test a proposition (Epstein, 2006). Thus, the flavor of ABM research is more similar to laboratory experimentation than to deductive proofs.[7] That is, each instantiation of an agent-based model can be used as an empirical observation; theoretical support can be garnered for a theory through induction, but no deductive proof obtained.

In the context of scientific inquiry, ABM has been employed as a tool for both prediction (akin complex econometric models) and understanding and explanation (akin to analytical models). When models are used for prediction, researchers typically focus on incorporating as much realism as possible. When models are used as a tool for improving understanding, researchers strive for simplicity, with the ultimate goal of incorporating only the features necessary to generate the phenomenon of interest (Axelrod, 1997; Grimm et al., 2005; Miller & Page, 2007; Ragan, 2010; Tesfatsion & Judd, 2006). Because the goal of this paper is to improve understanding and to provide insights into the effects of various interventions on fraud, we adopt the latter approach.

Benchmark model

Most (if not all) organizations face some level of fraud. There is also anecdotal evidence of fraud outbreaks in organizations where fraud becomes institutionalized as acceptable behavior. Consider the case of bribery at Siemens. A newspaper article (Schubert & Miller, 2008) reports that illegal bribe payments were so widely accepted at Siemens that they were treated as any other line item in the financial statements. The authors note that "bribery was Siemens's

[4] In the parable of the blind men and the elephant, a group of blind men touch an elephant to learn what it is like. Each one feels a different part, but only one part (e.g., the elephant's tail, tusk, leg, etc.). The blind men compare notes and find themselves in disagreement.

[5] The micro-level focus of agent-based models confers another benefit: It avoids the problematic representative agent assumption that is common in macro-level models (for a discussion of the shortcomings of the representative agent assumption, see Kirman, 1992).

[6] Our model employs Mathematica 7.0 (Wolfram Research Inc., 2008) as the programming platform.

[7] Epstein (1999) observes that ABM represents a new approach to science that can be used as an instrument for either theory development or theory testing. Rather than fitting neatly into inductive or deductive realms, ABM is more appropriately thought of as a *generative* form of science. To explain a macroscopic social phenomenon, the generativist asks, "How could decentralized local interactions of heterogeneous autonomous agents generate the given regularity?" (Epstein, 1999, p. 41), Put another way, the generativist's motto is "If you didn't grow it, you didn't explain its emergence" (Epstein, 1999, p. 46).

business model…Siemens had institutionalized corruption." Organizational level fraud (or fraud outbreaks) can arise from the creation of a culture of fraud in the organization, wherein fraud grows from the bottom up, as individual employees spread a culture of fraud to others through socialization or other means (e.g., see the discussion of organizations of corrupt individuals by Pinto et al. (2008)). Alternatively, organization-level fraud can be engineered by the organization's leaders and pushed from the top down, using powerful mechanisms such as obedience to authority (Milgram, 1963).[8] Consistent with the landscape of fraud observed in the natural ecology, our modeling goal is to identify a minimal set of plausible assumptions sufficient to generate organizations that experience both fraud outbreaks and stable levels of fraud. This goal guides our modeling choices, described below.

We base our model on Cressey's (1953) fraud triangle hypothesis, which still prevails as a popular view of the necessary and sufficient conditions for an individual's decision to commit fraud.[9] Cressey's characterization identifies three key determinants of fraud. When the conjunction of these determinants exists, an individual will commit fraud. The determinants are a perceived non-sharable financial need (or, more generally, a motive), a perceived opportunity to commit fraud, and an attitude or belief that frames the fraudulent act as acceptable behavior. In our model, agents become fraudsters if they face a conjunction of perceived opportunity (O), motive (M), and an attitude (A) that fraud is an acceptable behavior, or

$$O \wedge M \wedge A \leftrightarrow F.$$

Conversely, over time, if a fraudster no longer has opportunity, motive, or attitude, then he or she will reform to become an honest agent.[10] For simplicity, we assume that each of these variables is binary (either present or absent).[11]

Each of the agents in our model may or may not have a motive for fraud (determined randomly with some predetermined agent-specific probability at each step of the simulation).[12] Implicit in motive is some form of mental calculus performed by individual agents evaluating the perceived costs and benefits of committing fraud.[13] Over time, an agent's motive to commit fraud can change in response to her personal situation (e.g., unexpected medical bills or loss of employment by a family member).

Opportunity is an organization-level variable that proxies for the overall strength of a system of internal controls. It is represented as the probability that each agent will perceive an opportunity to commit fraud. For example, a 20% opportunity implies that each agent independently has a 20% chance of perceiving an opportunity to commit fraud at each point in the simulation. An agent might perceive such an opportunity when a loophole is discovered in the control system (e.g., a fellow employee fails to log off his computer or a manager "turns her back").

Finally, an individual's ability to rationalize fraud in our model is a function of both influence by others within the organization and factors exogenous to this influence (i.e., events not related to the work environment, for example, shifts in the attitudes of other reference groups such as family and friends). The plasticity of attitude in response to influence by others in the organization (whether through socialization among rank and file employees or through the command of organizational leadership) is central to the potential for fraud outbreaks in our model.

Characterizing agents

We begin developing our model of fraud in organizations by characterizing individual agents. In particular, each agent is a vector whose elements indicate the agent's attributes (endowments and individual behavioral rules). Each attribute is a number or a list of numbers. In our benchmark model, agent attributes include the following:

8 Social influence may also play a part in top-down fraud as more employees participate in the activity.

9 Wolfe and Hermanson (2004) expand the fraud triangle to a fraud diamond, suggesting capability as a fourth condition required for fraud. We elect to build on Cressey's original approach for two reasons. First, the original fraud triangle has the advantage of simplicity: Capability is already included in the hypothesis as part of opportunity. Second, the fraud triangle remains the dominant characterization in the professional literature. For example, it is incorporated in a discussion of the characteristics of fraud in the AICPA's 2002 *Statement on Auditing Standards* (SAS) 99, which addresses the auditor's responsibilities relating to fraud in an audit of financial statements.

10 Our model allows for free flow between the honest and fraudster classes. In the natural ecology, such a free flow is not always observed. In some cases individuals commit one fraud or a series of frauds to meet a particular goal and then reform. As an example, a sales manager may fraudulently record a large sale as occurring before year-end to meet his forecast in one year but then report properly in future periods when he is able to meet his forecast honestly. In other cases, the fraudster continues to commit fraud in perpetuity or until discovered. This can occur when the fraud is committed to solve an ongoing problem (e.g., living beyond one's means or hitting unrealistic growth expectations year after year) or if the decision to commit fraud in one period requires future fraud to keep the original fraud hidden. Because fraud is censored, we have no sense of the relative frequencies of one-shot versus continued fraud. Our model permits either type.

11 It is possible that the levels of these variables are a matter of degree in organizations and that a high value of one variable (e.g., motivation) compensates for the low value of another (opportunity or attitude). To model continuous values of O, M, and A, one would likely employ a threshold rule (where fraud occurs when the combination of the three factors exceeds some level). We conjecture that the use of a threshold rule in our model could lead to unanticipated responses to changes in parameters such as the phase transition observed in the threshold model in Davis et al. (2003).

12 At the outset, for simplicity, we assume that every agent has a motive for fraud but the model admits the possibility of heterogeneity in motives across individual agents.

13 Our model subsumes both fraud committed on behalf of individual agents and fraud committed on behalf of the organization. Underlying the mental calculus noted here, the benefits could accrue to the individual committing the fraud or to the organization. Similarly, the costs arising from the fraud could be imposed on either the individual or the organization (see Pinto et al. (2008) for a discussion).

- A unique integer that identifies the agent (a "name").
- Four independent probabilities, each drawn from a bounded uniform distribution. The first two represent the likelihood of the agent taking on a particular characteristic (motive for fraud and perceived opportunity to commit fraud) and the second two represent (i) the likelihood (q) of one's attitude toward fraud changing as a result of interacting with other agents and (ii) the likelihood (p) of one's attitude toward fraud changing as a result of factors unrelated to interactions with other agents. These probabilities are static and, with the exception of the opportunity likelihood, randomly determined for each agent. The opportunity likelihood is a static organization-level variable.
- Four binary variables indicating whether a motive for fraud is present, whether perceived opportunity is present, whether a pro-fraud attitude is present, and, finally, whether the agent is committing fraud. The presence of motive, opportunity, and attitude is determined at each step in the simulation for each agent, based on the assigned probabilities above. Fraud exists if motive, opportunity, and attitude co-exist.

The organization

To create an organization,[14] we develop a matrix. Each row in the matrix is an agent vector. To allow for interactions between agents, we represent a social network (a set of co-workers) for each agent by adding a list to each agent's vector, where each element in the list is an integer referring to another agent's name. In our social network, agent relationships are symmetric (if agent 1 is included in agent 3's social network, then the reverse is also true). The list of co-workers assigned to each agent is randomly determined (subject to the symmetry constraint) and remains static.

Observing the organization over time

After establishing a starting population in the organization, we allow agents to repeatedly interact and update agent states. Interactions between agents are accomplished by repeatedly and randomly pairing each agent with a member of his or her social network.[15] The pairing allows each agent to observe another. As noted previously, the observing agent will emulate the attitude[16] of the observed with probability q. Consider, for example, a situation in which an agent goes to lunch with co-workers and is told that one of them believes it is acceptable to claim a personal meal as a reimbursable business expense. We assume that the more often this occurs, the more likely the agent will begin to rationalize committing similar behavior. We also allow for the possibility that one's attitude toward fraud can change (with probability p) independently of observing others inside the organization.

In the model, agents' attitudes about the acceptability of fraud change according to their particular experiences with other agents and their personal proclivities. At any given time, agent attitudes may be heterogeneous because each agent has a unique set of characteristics and different experiences in the organization. One particular agent can be a fraudster at a given time while other agents may be honest. When interacting with an honest agent in the organization, the fraudster can trigger a pro-fraud attitude in the honest agent (and possibly convert the honest agent into a fraudster, if both motive and perceived opportunity exist) or the honest agent can convert the fraudster into an honest agent. In the case of two honest agents interacting, one or both agents may begin to view fraud as acceptable because of factors exogenous to their interactions. These rules for interaction and behavior define a social dynamic. The culture changes dynamically in response to both agent interactions and their spontaneous behavior. The model is "bottom up," in that a culture of fraud can emerge spontaneously in the population as a result of the interactions of individual agents with each other.

An example

Recall that our model begins with an initial organization represented by a matrix in which each row represents an agent and each element in a row represents an agent attribute. For clarity, we present an example of a matrix with a population size of five:

$$\begin{pmatrix} ID & p(M) & M & p & q & A & p(O) & O & F & E & S & SN \\ 1 & 0.00449253 & 0 & 0.0242858 & 0.90035 & 0 & 0.95 & 1 & 0 & 0 & 1 & \{4\} \\ 2 & 0.426662 & 0 & 0.0292374 & 0.585986 & 0 & 0.95 & 1 & 0 & 0 & 0 & \{\,\} \\ 3 & 0.354574 & 0 & 0.0322799 & 0.511431 & 0 & 0.95 & 1 & 0 & 0 & 2 & \{4,5\} \\ 4 & 0.0897837 & 0 & 0.036227 & 0.525595 & 0 & 0.95 & 1 & 0 & 0 & 2 & \{1,3\} \\ 5 & 0.496052 & 0 & 0.0205544 & 0.892721 & 0 & 0.95 & 1 & 0 & 0 & 1 & \{3\} \end{pmatrix}.$$

The elements in each column represent the following agent attributes:

ID = agent's unique identifier,
$p(M)$ = likelihood of an agent's motive to commit fraud changing,
M = binary variable indicating the presence of motive,
p = likelihood of an agent changing attitude as a result of factors exogenous to agent interactions,
q = likelihood of an agent changing attitude as a result of interactions with other agents,
A = binary variable indicating the presence of an attitude supporting fraud,

[14] In our model, the term "organization" can apply to an entire business entity, a branch, a division, an office, or even a department within an office.

[15] We employ an algorithm developed by Gaylord and Davis (1999) to accomplish this pairing.

[16] The model presented in this paper assumes that agent attitude is influenced by other agents' attitudes. We also evaluated an alternative model where agent attitude is influenced only by other agents' actions. The only qualitative difference observed was in the conditions under which perceived opportunity to commit fraud is reduced. In this instance, organizations with moderate to high levels of social influence respond to reduced opportunity in a fashion similar to that observed in our asymmetric influence condition (i.e., fraud outbreaks are eliminated and the mean level of fraud is reduced).

Table 1
Summary of model analysis.

Starting population attitude	Random ($n = 10$)	All pro-fraud ($n = 5$)	All anti-fraud ($n = 5$)
Model variant			
Benchmark	$q = \{0.05, 0.10, \ldots, 0.95\}$	$q = \{0.05, 0.10, \ldots, 0.95\}$	$q = \{0.05, 0.10, \ldots, 0.95\}$
Opportunity	$q = \{0.05, 0.10, \ldots, 0.95\}$	$q = \{0.05, 0.10, \ldots, 0.95\}$	$q = \{0.05, 0.10, \ldots, 0.95\}$
	$p(O) = \{0.9, 0.8, \ldots, 0.5\}$	$p(O) = \{0.9, 0.8, \ldots, 0.5\}$	$p(O) = \{0.9, 0.8, \ldots, 0.5\}$
Tone at top	$q = \{0.05, 0.10, \ldots, 0.95\}$	$q = \{0.05, 0.10, \ldots, 0.95\}$	$q = \{0.05, 0.10, \ldots, 0.95\}$
	Managers = {1, 2, 4, 6}	Managers = {1, 2, 4, 6}	Managers = {1, 2, 4, 6}
Asymmetric	$q = \{0.05, 0.10, \ldots, 0.95\}$	$q = \{0.05, 0.10, \ldots, 0.95\}$	$q = \{0.05, 0.10, \ldots, 0.95\}$
Influence	$c = \{99, 19, 9, 5.7, 4, 3\}$	$c = \{99, 19, 9, 5.7, 4, 3\}$	$c = \{99, 19, 9, 5.7, 4, 3\}$
Termination	$q = \{0.05, 0.10, \ldots, 0.95\}$	$q = \{0.05, 0.10, \ldots, 0.95\}$	$q = \{0.05, 0.10, \ldots, 0.95\}$
	$e = \{0.05, 0.10, \ldots, 0.20\}$	$e = \{0.05, 0.10, \ldots, 0.20\}$	$e = \{0.05, 0.10, \ldots, 0.20\}$

For all model variants, 10 starting organizations were used. Attitude regarding fraud was either randomly assigned to agents (for all 10 organizations) or set to anti-fraud or pro-fraud for a subset of five organizations. For all model variants, average emulation likelihood (q) was varied from 0.05 to 0.95 in 5% increments. Within each value of q, the model variant parameters (i.e., perceived opportunity, number of managers, influence scalar, and termination likelihood) were evaluated at the levels indicated.

$p(O)$ = likelihood of an agent perceiving an opportunity to commit fraud,
O = binary variable indicating the perceived opportunity to commit fraud,
F = binary variable indicating whether an agent is committing fraud,
E = binary variable to be used later when agent removal is considered,
S = social network size, and
SN = list of an agent's social network.

In the sample matrix above, agent 3's likelihood of a change in motive is 0.354574, opportunity is currently perceived as present, no fraud is being committed, the termination variable is set to zero, the agent's social network size is two, and the social network consists of agents 4 and 5. The matrix represents the organization's and agents' characteristics for one period in time.

As noted previously, agents are iteratively paired and allowed to interact over time. For example, moving forward one interaction period from the initial matrix above, if agent 1 begins to rationalize fraud as a result of factors exogenous to agent interaction, his attitude variable will change from zero to one. Further, if we assume agent 4 is subsequently paired with agent 1 during this period and is influenced by agent 1, agent 4 will also adopt an attitude supportive of fraud, with a consequent change in his attitude variable (from zero to one). The following matrix illustrates the result:

$$\begin{pmatrix} ID & p(M) & M & p & q & A & p(O) & O & F & E & S & SN \\ 1 & 0.00449253 & 0 & 0.0242858 & 0.90035 & 1 & 0.95 & 1 & 0 & 0 & 1 & \{4\} \\ 2 & 0.426662 & 0 & 0.0292374 & 0.585986 & 0 & 0.95 & 1 & 0 & 0 & 0 & \{\} \\ 3 & 0.354574 & 0 & 0.0322799 & 0.511431 & 0 & 0.95 & 1 & 0 & 0 & 2 & \{4,5\} \\ 4 & 0.0897837 & 0 & 0.036227 & 0.525595 & 1 & 0.95 & 1 & 0 & 0 & 2 & \{1,3\} \\ 5 & 0.496052 & 0 & 0.0205544 & 0.892721 & 0 & 0.95 & 1 & 0 & 0 & 1 & \{3\} \end{pmatrix}.$$

In the matrix we can see that although agents 1 and 4 now hold a pro-fraud attitude ($A = 1$) and both perceive the opportunity to commit fraud ($O = 1$), neither has motive ($M = 0$) and therefore neither commits fraud.

Benchmark model design

We use 10 unique starting organizations to investigate our model.[17] To establish a benchmark, we initially assume that every agent has a motive for fraud and perceives opportunity. We set the size of the population to 100 agents and observe 15,000 interaction periods. Each agent is initially assigned (with a 50% probability) to either a pro-fraud or an anti-fraud attitude.[18] Finally, we begin by allowing agents to have a very small but heterogeneous likelihood of changing their attitude toward fraud as a result of factors exogenous to agent interaction (probability p drawn independently for each agent from a uniform distribution from 0 to 0.005).[19] The design of our benchmark analysis (and subsequent analysis of anti-fraud interventions) is described in Table 1.

[17] Our design and analysis were constrained by computational power and practicality (approximately 5000 hour of computing time were expended). We chose fewer organizations and more interaction periods because additional interaction periods were low cost relative to examining additional organizations (due to the extensive analysis required for each organization across the range of parameters that we examine). Similarly, additional agents (beyond 100) placed excessive demands on our computational resources and we felt that the additional insight provided was likely to be limited.

[18] For a subset of five of the organizations, we examined alternative starting conditions, where all agents start with a pro-fraud attitude or all agents start with an anti-fraud attitude. As these starting populations evolved, the patterns of behavior converged to match those observed in populations where agents were initially assigned to either pro-fraud or anti-fraud attitudes with equal likelihood. We consequently conclude that assignment of attitudes in the starting population has no effect on outcomes over time.

[19] Our choice of a small probability here is consistent with a belief that individuals' attitudes regarding the acceptability of fraud are relatively stable in response to non-organizational factors. As noted by the AICPA (2002) in SAS 99, the ability to rationalize can be affected, in part, by one's attitude, character, and set of ethical values (which we believe to be stable with only a small likelihood of change over time). This is the aspect of rationalization captured by our small exogenous probability p. In contrast, the likelihood of being influenced via interactions in the organization, q, captures a second source of rationalization noted in SAS 99, the organizational environment.

Emergence: Two states

The benchmark analysis studies the effect of a systematic change in individuals' tendencies to be impacted by the behavior of others in the organization (emulation likelihood, q). We explore the entire parameter range by varying the average emulation likelihood in increments of 5%, from $q = 0.05$ to 0.95.[20] Each agent is initially assigned an independent probability drawn from a uniform distribution between a minimum and a maximum value for q. The assigned probability is static for each agent. Since the likelihood is bounded, as we approach the extremes we manipulate either the minimum or the maximum to reach the desired average.[21]

Our analysis focuses on both state distributions in the form of histograms illustrating the number of fraudsters in the population that emerge and on time series graphs illustrating the change in the number of fraudsters over time, with an eye toward the periodic existence of fraud outbreaks. Each histogram shows the number of periods during which n agents are fraudsters, where n ranges from 0 to 100 (the entire organization). In Fig. 1, Panels A and B, for example, the y-axis represents the number of periods in which a specific number of fraudsters is observed, while the x-axis represents the number of fraudsters. Panels C and D of Fig. 1 illustrate changes in the number of fraudsters over time. In each time series, the y-axis represents the number of fraudsters in a given period and ranges from 0 to 100 and the x-axis represents time (or interaction periods) and ranges from 1 to 15,000.

Two state distributions emerge in our analysis. At moderate and high emulation likelihoods, the majority of periods exhibit extreme behavior in all 10 organizations, where either virtually every agent is a fraudster or virtually no agents are fraudsters. We call this distribution U shaped after the appearance of its associated histogram, shown in Fig. 1, Panel A. A typical time series associated with U-shaped distributions is illustrated in Fig. 1, Panel C. The time series is characterized by rapid swings between extremes. For example, in Panel C, at approximately the 11,000th pairing, a rapid shift occurs from a period in which few agents are fraudsters to a period in which virtually everyone is a fraudster. We find that the characteristic U shape and sudden shifts between extreme states persist across conditions ranging from very high average emulation probabilities down to approximately 30%.[22] The dramatic shifts in behavior arise from the overwhelming effect of emulation relative to the random effects introduced by other parameters in the model.

At low average emulation likelihoods (30% and below), a very different state distribution emerges in every organization that we observe. Instead of tending toward extremes, the population tends toward a specific level of fraud. The shape of the associated histogram is an inverted U, as illustrated in Fig. 1, Panel B. Notice in this example, roughly half of the population in the majority of periods are fraudsters. When we consider the change in fraud behavior over time in Fig. 1, Panel D, we see that the rapid extreme shifts in the number of fraudsters found in the U-shaped distribution disappear for the most part. Instead, we see smaller (noisy) movements around a specific number of fraudsters.

We perform additional analysis at parameter values around the transition from a U-shaped organization to an inverted U-shaped organization to evaluate whether the change in state is gradual or a bifurcation point exists. Our analysis indicates that the change is gradual. As q decreases, the population spends less time at the extremes and more time at moderate levels of fraud; the height of the upper and lower supports of the U shape in the histogram decrease and the middle range becomes thicker until organizations reach a point where the state distribution is flat (all fraud outcomes are equally likely over time). As q continues to decrease, the population tends toward a specific number of fraudsters, leading to an inverted U shape in the histogram. The location of the transition from U to inverted U differs slightly across different organizations and depends on the relative values of p (likelihood of changing attitude in response to exogenous factors) and q (likelihood of emulating the attitude of others in the organization). As p increases, increasingly large values of q are required to generate a U-shaped state distribution. High emulation likelihoods within the organization relative to external influences on attitude facilitate cascades from mostly honest to mostly fraudster populations and vice versa.

In Table 2 we provide descriptive statistics (mean and standard deviation) to illustrate the impact of interventions on the statistical properties of fraud in our models. For the benchmark analysis, the table shows a reduction in the average standard deviation across organizations as q is reduced, while the mean remains relatively constant. This pattern is consistent with the change in state distributions noted previously.

The possibility of two states of fraud dynamics observed in our benchmark model raises an additional question: Do efforts to prevent or detect fraud have differential effects on the two fraud dynamics observed? We now turn our

[20] While we have no direct evidence regarding the level of q in real world organizations, research by Milgram (1965) suggests that the level of q across organizations may vary. In the study, the extent of obedience to authority (one aspect of emulation likelihood) was investigated at Yale and at an office building in Bridgeport, Connecticut. Subjects were more obedient at Yale; 65% were fully obedient versus 48% at the office in Bridgeport. The study (and others conducted by Milgram) suggests that the effectiveness of efforts to recruit (q) in our model may depend in important ways on the organizational setting (e.g., background authority inherent in the organization, etc.).

[21] For example, if the maximum likelihood is set at 100% and the minimum at zero, the average likelihood is 50%. Holding the floor constant at zero as the ceiling is reduced from 100%, the average will drop below 50%. To examine the parameter space where averages exceed 50%, we increase the floor and hold the ceiling constant at 100%. Since this approach to manipulating the mean value of q also varies the range from which the value of q is drawn, we performed additional analysis using a variety of mean-preserving ranges. Our results are qualitatively similar (both time series patterns and state distributions of fraud over time are consistent) in both the mean-preserving and mean-shifting analyses, with the exception that a U-shaped distribution is observed more frequently when mean-preserving ranges are reduced.

[22] Unless noted otherwise, the results are consistent and clearly observable across all organizations and all initial fraud distributions in the populations.

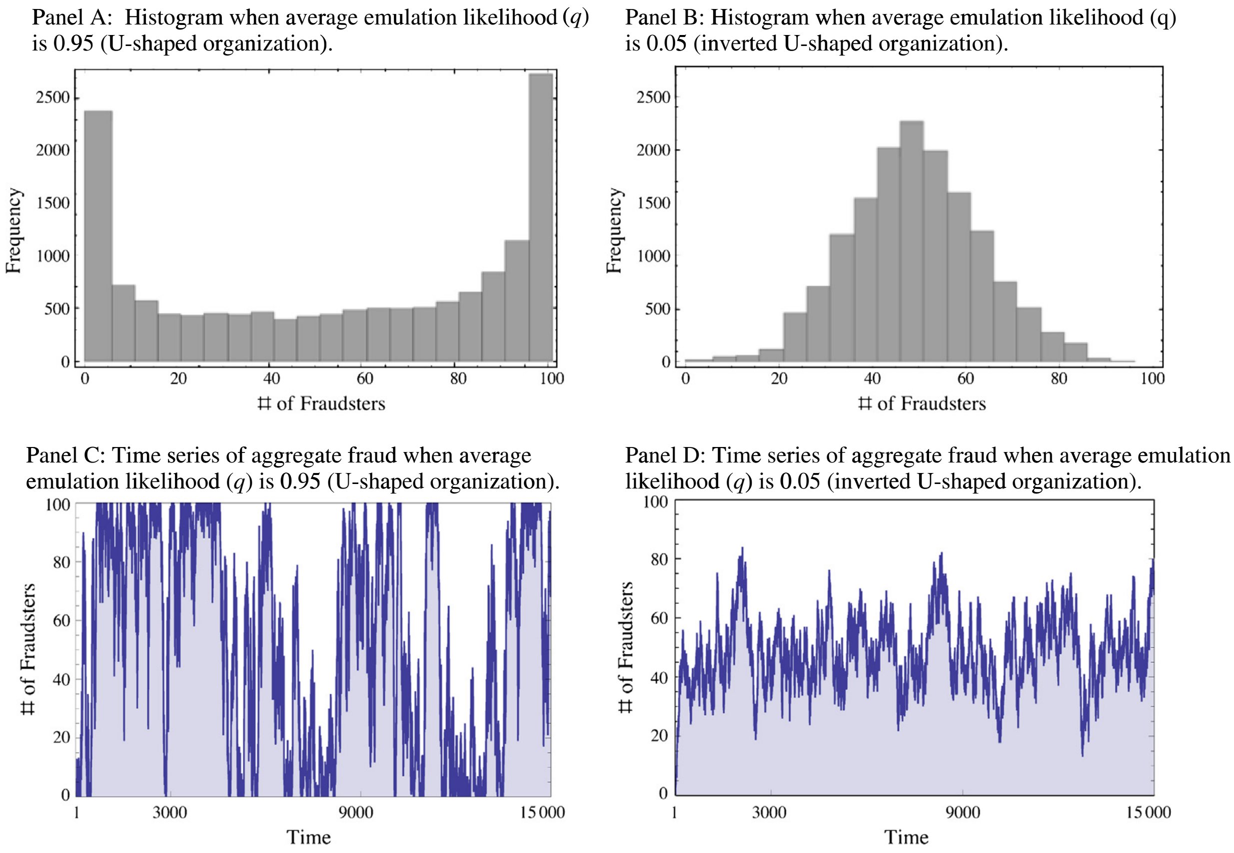

Fig. 1. Illustrative results from analysis of a single organization in the benchmark model, including histogram of number of fraudsters and associated time series. Panel A: Histogram when average emulation likelihood (q) is 0.95 (U-shaped organization). Panel B: Histogram when average emulation likelihood (q) is 0.05 (inverted U-shaped organization). Panel C: Time series of aggregate fraud when average emulation likelihood (q) is 0.95 (U-shaped organization). Panel D: Time series of aggregate fraud when average emulation likelihood (q) is 0.05 (inverted U-shaped organization).

attention to the impact of attempts to prevent or eliminate fraud in the organization.

Perceived opportunity

We begin our investigation into the effectiveness of fraud prevention mechanisms by systematically manipulating perceived opportunity. Our analysis examines opportunity likelihood parameter values from 90% to 50% in 10% increments. For each value of $p(O)$, we investigate the entire emulation likelihood (q) range, focusing our attention on the two state outcomes (U-shaped and inverted U-shaped distributions) identified during the benchmark analysis.

When we evaluate the effect of reducing opportunity in U-shaped organizations (with moderate and high average emulation levels), most of the time series characteristics are unchanged. We observe the same extreme changes between honesty and fraud at the same points in time within an organization. The frequency of the peaks and valleys are the same relative to the benchmark time series for each organization, as is the duration spent in each state. However, as indicated in Table 3, when perceived opportunity is reduced, we observe a reduction in the maximum level of fraud (e.g., in the $q = 0.95$ condition, the maximum level of fraud changes from 100% when $p(O) = 1.0$ to 68 percent when $p(O) = 0.50$). The effect on the time series is illustrated by comparing the benchmark time series in Fig. 1, Panel C to the time series in Fig. 2, Panel A. The associated state distribution histogram is illustrated in Fig. 2, Panel B. As opportunity decreases, the upper support of the U shape is lowered in the histogram.

The same result emerges when we consider low emulation likelihood organizations. Relative to the benchmark analysis, the qualitative characteristics of each organization's time series (the frequency, duration and location of peaks and valleys) are unchanged, but the maximum fraud level drops as $p(O)$ is reduced (in Table 3, for organizations where $q = 0.05$, the maximum level of fraud drops from 91% in the benchmark model to 58% when $p(O) = 0.5$). A time series for a representative organization is illustrated in Fig. 2, Panel C. In Table 2, both the mean and the standard deviation of fraud are reduced as perceived opportunity is reduced. The inverted U state distribution retains its basic shape, but it is shifted to the left and becomes more peaked (consistent with the reduced mean and standard deviation). An illustrative histogram is presented in Fig. 2, Panel D.

Thus, overall, we find that reducing opportunity in the organization shifts the average fraud level closer to zero and reduces the standard deviation of fraud, but it does not otherwise affect the shape of the distribution of fraud over time or the nature of fraud dynamics. The two emergent states observed in our benchmark model persist and

Table 2
Mean percentage (standard deviation) of agents who commit fraud within 10 organizations, where agent attitude (pro-fraud versus anti-fraud) is randomly assigned in the initial time period.

Mean emulation likelihood (q)	0.05	0.25	0.50	0.75	0.95
Benchmark	50.4 (13.5)	48.1 (22.9)	52.4 (26.6)	47.8 (34.8)	51.0 (35.3)
Opportunity (p(O))					
0.90	45.4 (12.4)	43.3 (20.1)	47.2 (24.1)	43.0 (31.4)	45.9 (31.8)
0.80	40.3 (12.4)	38.5 (15.9)	41.0 (20.4)	38.3 (28.0)	40.8 (28.4)
0.70	40.3 (11.2)	33.7 (15.9)	36.7 (19.0)	33.5 (24.6)	35.7 (24.9)
0.60	30.2 (8.8)	28.9 (13.8)	31.5 (16.4)	28.7 (21.0)	30.6 (21.5)
0.50	25.2 (7.6)	24.1 (11.6)	26.2 (13.8)	23.9 (17.7)	25.5 (18.0)
Tone at top					
1 manager	42.5 (12.8)	33.6 (17.8)	30.1 (20.2)	15.2 (17.5)	15.4 (18.2)
2 managers	38.5 (11.6)	25.8 (14.8)	21.3 (15.8)	8.8 (11.2)	9.3 (12.1)
4 managers	32.3 (10.4)	17.2 (10.9)	13.2 (11.0)	5.0 (7.0)	4.9 (7.0)
6 managers	27.1 (9.4)	13.7 (8.9)	9.9 (8.7)	3.4 (5.0)	3.4 (5.2)
Asymmetric influence (c/(c + 1))					
0.99	47.0 (13.0)	43.9 (22.3)	44.3 (26.6)	33.7 (32.4)	27.9 (30.1)
0.95	44.4 (12.3)	33.0 (19.5)	23.5 (19.0)	7.5 (10.7)	5.5 (8.2)
0.90	37.9 (12.0)	20.4 (13.2)	12.7 (11.3)	3.3 (5.2)	2.4 (4.0)
0.85	31.6 (10.5)	14.7 (10.4)	8.6 (8.0)	2.1 (3.5)	1.5 (2.8)
0.80	27.0 (9.6)	10.7 (8.0)	6.5 (6.1)	1.5 (2.7)	1.1 (2.1)
0.75	24.6 (8.9)	8.6 (6.3)	5.3 (4.9)	1.2 (2.1)	0.8 (1.7)
Termination likelihood (e)					
0.01	28.4 (11.0)	31.9 (19.8)	33.2 (25.4)	39.2 (34.0)	36.6 (33.8)
0.05	7.4 (4.7)	8.7 (7.9)	9.2 (10.3)	11.7 (17.0)	12.4 (17.4)
0.10	3.7 (3.0)	4.2 (4.5)	4.5 (5.7)	5.3 (7.8)	5.3 (8.2)
0.15	2.3 (2.4)	2.8 (3.3)	2.9 (4.1)	3.5 (5.5)	3.7 (5.8)
0.20	1.7 (1.9)	2.1 (2.7)	2.2 (3.2)	2.6 (4.2)	2.7 (4.4)

Table 3
Maximum (minimum) fraud level observed across all organizations in the final 10,000 time steps by level of emulation likelihood for the benchmark and model variants most likely to impede fraud.

Mean emulation likelihood (q)	0.05	0.25	0.50	0.75	0.95
Benchmark	91 (9)	100 (0)	100 (0)	100 (0)	100 (0)
Opportunity (p(O))					
0.50	58 (1)	68 (0)	67 (0)	68 (0)	68 (0)
Tone at top					
6 managers	64 (2)	60 (0)	64 (0)	52 (0)	70 (0)
Asymmetric influence (c/(c + 1))					
0.75	68 (0)	48 (0)	44 (0)	25 (0)	20 (0)
Termination likelihood (e)					
0.20	13 (0)	33 (0)	38 (0)	50 (0)	65 (0)

the extreme swings between honesty and widespread fraud in U-shaped organizations remain. The results of our analysis suggest that, while reducing opportunity can limit the extent of fraud in an organization, firm culture and individual susceptibility to influence continue to play an important role in the nature of fraud dynamics.

Tone at the top

Tone at the top is often noted as an important characteristic in organizations for controlling fraud (Association of Certified Fraud Examiners, 2007). To better understand the impact of tone at the top on dynamics in our model, we begin by extending our benchmark analysis to create a hierarchical organization with two levels of employees: managers and staff. We assume that managers are more likely than staff to influence the attitudes of others.

We introduce a hierarchy by adding a status identifier to each agent vector in the benchmark model. Agents are randomly identified as managers in the organization prior to the first period of agent interaction. Agents identified as managers are assigned a hierarchical status identifier equal to two (resulting in twice the impact on emulation likelihood),[23] while the remaining agents are considered staff and assigned a hierarchical status identifier equal to one. We use this identifier to scale emulation likelihoods during

[23] We tested the sensitivity of the emulation likelihood scalar introduced when interacting with a manager by increasing it from two to four. No qualitative effect was observed on aggregate outcomes.

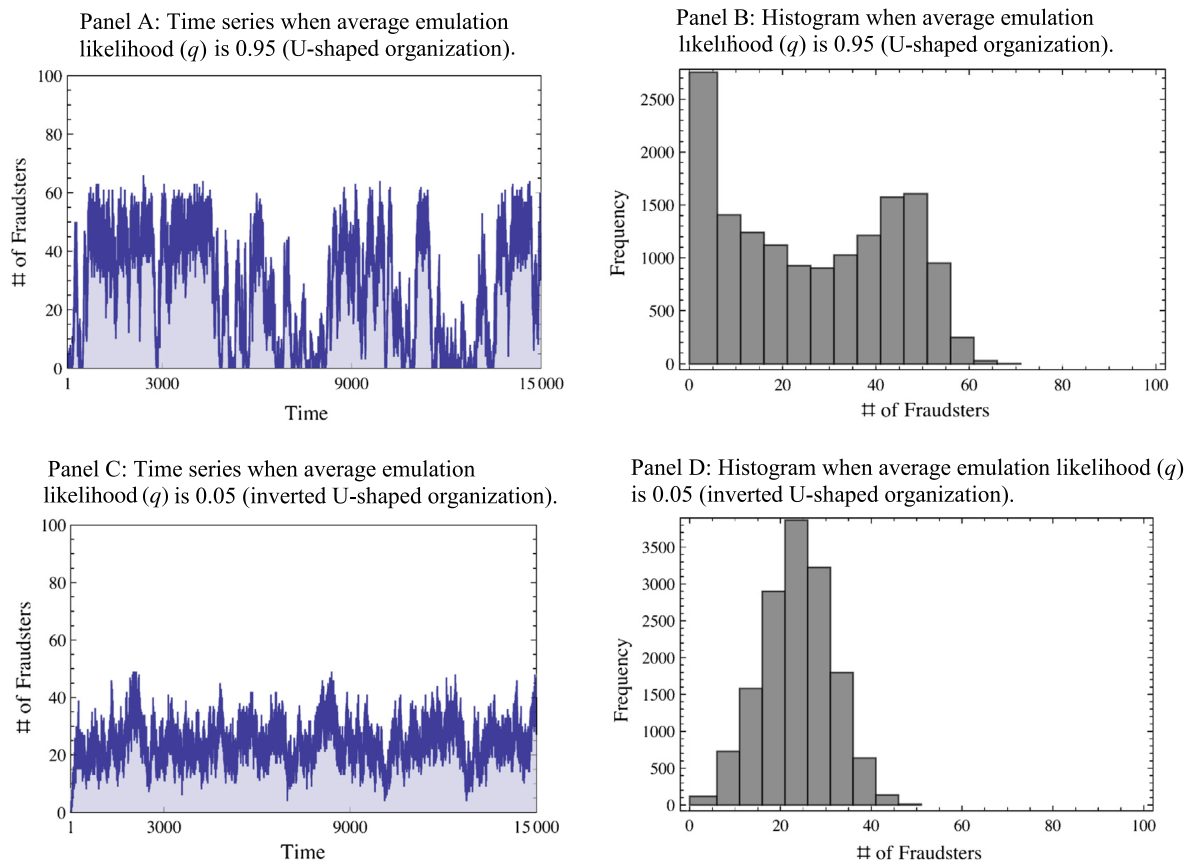

Fig. 2. Representative histograms and time series of aggregate fraud levels when perceived opportunity, $p(O)$, equals 0.50. Panel A: Time series when average emulation likelihood (q) is 0.95 (U-shaped organization). Panel B: Histogram when average emulation likelihood (q) is 0.95 (U-shaped organization). Panel C: Time series when average emulation likelihood (q) is 0.05 (inverted U-shaped organization). Panel D: Histogram when average emulation likelihood (q) is 0.05 (inverted U-shaped organization).

agent interactions. For example, assume agent 1 is assigned an emulation likelihood of 0.15 when the organization is formed. In our benchmark model, when agent 1 observes another agent, agent 1 will change her attitude to match the observed agent's attitude with a 15% probability. However, in our hierarchical model the likelihood of emulating another agent is doubled, to 30%, when the agent observed is a manager.[24]

The notion of tone at the top incorporates more than a hierarchical organization with influential management. It is usually characterized as a management team with a consistent message regarding ethical values and appropriate behavior. To capture this idea, all of the managers in the model are identical. Each is assigned an anti-fraud attitude in the initial period of the simulation. To ensure their attitudes do not change, the managers' likelihood of emulating another agent (q) and likelihood of spontaneous attitude change (p) are both set to zero.[25] For the purposes of our analysis, we make no other modifications relative our benchmark. Motive and opportunity remain constant, at 100%. We systematically vary the number of managers in the organization (span of control) and investigate the impact of this manipulation over the range of emulation likelihood (q) values (as in the benchmark model).

We find that a monolithic, honest management team in a U-shaped organization changes the characteristic shape of the state distribution. With even one honest manager, we see the tendency to achieve maximum levels of fraud with high frequency disappear (the upper tail in the U shape is gone, as illustrated in Fig. 3, Panel A).[26] A clearer picture of the qualitative effect of the manipulation is provided by an examination of the dynamics of fraud over time. While there is an overall tendency toward honesty in organizations with a positive tone at the top, the time series exhibits short-lived outbreaks of fraud in the organization (see Fig. 3, Panel B). This remains the case regardless of the number of managers we examined. The large influence exerted by management in U-shaped organizations can be attributed to both the consistent behavior exhibited by management and the high level of susceptibility to influence exhibited in the agent population.

To complete our analysis of a monolithic, honest management team, we consider its impact on the aggregate fraud outcomes and dynamics when average emulation (q) levels are low (i.e., inverted U-shaped organizations).

[24] Since q is bounded ($q \leqslant 1$), managers can only increase another agent's emulation likelihood parameter to 1, but no higher.

[25] We also evaluated a hierarchical model using non-identical managers with attitudes that can change to investigate the effect of introducing a hierarchy. The introduction of a hierarchy by itself had no qualitative impact relative to the benchmark model.

[26] Notably, when a monolithic, dishonest management team is introduced into the model, the mirror image of the behavior exhibited here is observed. The lower support is eliminated and fraud becomes the prevalent behavior.

Panel A: Histogram for U-shaped organizations illustrating loss of upper support in tone at the top, asymmetric influence, and termination model variants.

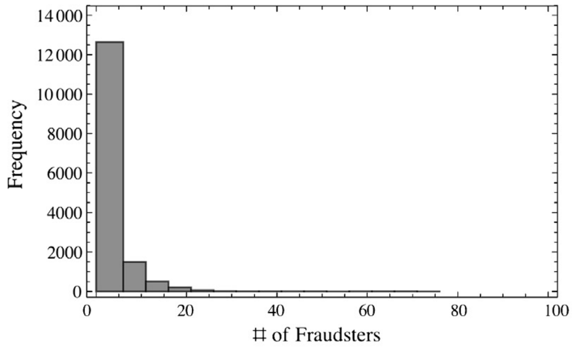

Panel B: Time series for U-shaped organizations illustrating fraud outbreaks under tone at the top and termination model variants.

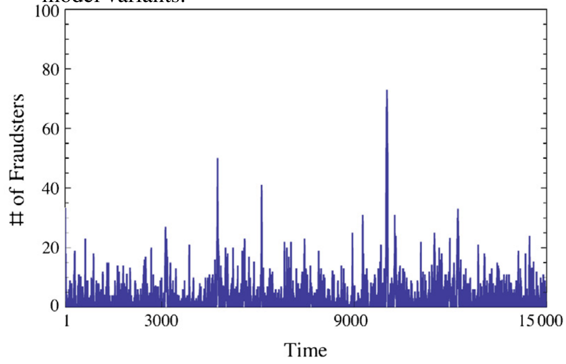

Panel C: Time series of aggregate fraud levels exhibiting no fraud outbreaks in U-shaped organizations in the asymmetric influence model variant and in inverted U-shaped organizations in the termination model variant.

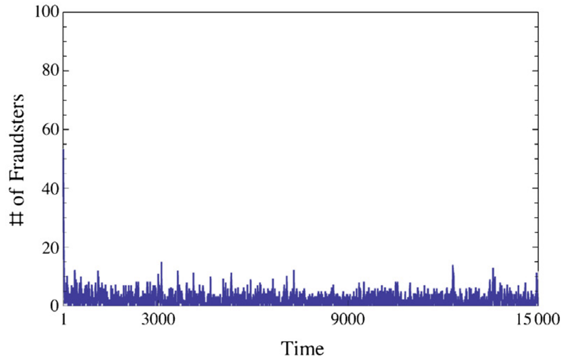

Fig. 3. Other histograms and time series of aggregate fraud levels observed in the analysis of model variants. Panel A: Histogram for U-shaped organizations illustrating loss of upper support in tone at the top, asymmetric influence, and termination model variants. Panel B: Time series for U-shaped organizations illustrating fraud outbreaks under tone at the top and termination model variants. Panel C: Time series of aggregate fraud levels exhibiting no fraud outbreaks in U-shaped organizations in the asymmetric influence model variant and in inverted U-shaped organizations in the termination model variant.

We find that, while introducing a monolithic, honest management team reduces mean fraud level and its associated standard deviation (as can be seen in Table 2), it has no other effect on either the inverted-U shape of the histogram or characteristics of the related time series. Furthermore, increasing the relative size of the management team does not alter this general finding.

Asymmetric influence

Some interventions are aimed at enhancing the ethical tendencies of employees. For example, organizations may engage in ethical training or encourage personal adherence to a code of ethics. A clearly communicated emphasis on ethical behavior could have effects on both fraudsters and others in the organization. Non-fraudsters would be better informed about what actions constitute fraud and engage in more careful thought before conforming with undesirable norms (rather than thoughtless conformity to norms). Fraudsters would feel less comfortable in their recruitment efforts and less likely to reveal their attitudes to others. To model these effects, we allow asymmetric influence, where honest employees exert more influence on their peers than fraudsters. To achieve asymmetry, we add a scalar (c) to our benchmark model. The scalar is constant across all agents because we assume that the intervention would be applied across the entire population. During agent interaction, the scalar changes the likelihood of emulating a fraudster relative to an honest agent in the following manner:

$Q = qc/(c + 1)$ when the agent observes another agent with a pro-fraud attitude and

$Q = q$ otherwise,

where Q is the emulation likelihood in the asymmetric influence model and q is the emulation likelihood in the benchmark model. To illustrate the application of this rule set, assume agent A's emulation likelihood (q) is 0.50 and the scalar (c) is 9 in the organization. When agent A observes another agent with a pro-fraud attitude, he will have a 45% likelihood of emulating the attitude. However, if agent A observes an anti-fraud attitude in another agent, he will have a 50% likelihood of emulating the attitude. Thus, the smaller the value of c, the higher the likelihood that an agent will emulate an honest agent relative to a fraudster, and the more effective the ethics intervention.

To evaluate this version of our model, we systematically vary c at six levels (99, 19, 9, 5.7, 4 and 3)[27] for every

[27] The levels of c were chosen to scale the influence of pro-fraud relative to anti-fraud attitudes to 99%, 95%, 90%, 85%, 80% and 75%, respectively.

emulation likelihood level used in the benchmark model so that the effect of asymmetric influence can be assessed across both U-shaped (moderate to high emulation likelihood) organizations and inverted U-shaped (low emulation likelihood) organizations. We compare the results of these parameter sweeps to matched (identical) starting organizations examined in our benchmark model.

For U-shaped organizations, asymmetric influence dramatically affects both the shape of the state distribution and the dynamics of fraud behavior over time. As the asymmetry in influence is increased (by decreasing the scalar c), the upper support of the U shape rapidly disappears. The resulting state distribution resembles that observed in the tone at the top condition, in Fig. 3, Panel A.[28] However, the periodic fraud outbreaks observed in the tone at the top condition are completely eliminated. As can be seen in Table 3, the maximum level of fraud observed when $q = 0.95$ is 20 at the highest level of asymmetric influence, compared to 100 in the benchmark model. The resulting time series is illustrated in Fig. 3, Panel C. Intuitively, when emulation likelihoods are high and asymmetric, the number of honest agents in the population will continue to grow until there are a sufficient number of honest agents to prevent an outbreak. A very different result emerges when we observe the impact on the inverted-U distribution found in organizations with low average emulation likelihoods. As asymmetry in influence increases, the mean level and standard deviation of fraud decreases, but the inverted-U shape of the state distribution and the qualitative characteristics of the time series in the benchmark model are preserved. The results suggest that the role of ethical interventions in fraud prevention is contingent on the nature of the organization, a relation not previously recognized in the literature that could prove important in fraud risk assessment.

Removing fraudsters from the organization

Our investigation so far has focused on organizational interventions associated with fraud prevention. We now turn our attention to the effect of detecting and eliminating fraud in the organization through the termination of fraudsters. To model termination, we assume that fraudsters will be detected with some probability $t(e)$. Once detected, agents are removed from the organization ("terminated") and replaced with a new agent.

We expect the existence of budget constraints with regard to termination efforts. As a result, we assume fraud detection and termination programs will be more successful in organizations with lower levels of fraud. At higher fraud levels, we anticipate that collusion among employees will impede such efforts. Therefore, in each period, we determine the likelihood of detection and termination as

$$t(e) = e((N - f)/N),$$

where $t(e)$ is the likelihood of detection and termination, e is the likelihood of detection and termination unadjusted for the number of fraudsters in the organization, N is the number of agents in the organization ($N = 100$), and f is the number of fraudsters in the organization at the start of the simulation period.[29]

When a fraudster is discovered and replaced, our termination model assigns new values for all agent attributes (using the process described in the benchmark model), with the exception of the agent identifier (ID), social network (SN), and social network size (S).[30] These three attributes remain unchanged under the assumption that when an agent is replaced, the new agent will interact with the same individuals as the previous one. Our analysis of the termination model evaluates the effect of a change in detection and termination likelihood (e) over the entire range of emulation likelihood values examined in the benchmark model to investigate the impact of termination on our two classes of organizations (U shaped and inverted U shaped). In our analysis, e is examined at five levels: 0.01, 0.05, 0.10, 0.15 and 0.20.

When we consider organizations with moderate to high average emulation likelihoods (U-shaped organizations), we find that the introduction of a termination regime leads to results similar to those observed in the tone at the top model. The upper support of the U-shaped state distribution is gone (a tendency toward low levels of fraud), while periodic fraud outbreaks persist in the time series (see Panels A and B of Fig. 3). For organizations with low emulation likelihoods, we find the first instance of an intervention triggering a dramatic change in the inverted-U shape of the state distribution. As the likelihood of detection and termination increases, the level of fraud approaches zero (the state distribution resembles Fig. 3, Panel A) and fraud outbreaks are not observed in the time series (see Fig. 3, Panel C). Our analysis suggests that detection and termination efforts can be effective at reducing average fraud levels to near zero in both types of organizations (U and inverted U shaped). However, the effectiveness of such efforts at preventing widespread fraud outbreaks appears contingent on the type of organization and related individual susceptibilities to fraud.

Discussion and conclusions

We use an ABM to investigate the dynamics of fraud within organizations and the impact of fraud prevention and detection mechanisms. Our model consists of an organization, agents (employees), and a set of simple social interaction rules. In accordance with Cressey's (1953) occupational fraud model, any agent within our model possessing the conjunction of motive, opportunity, and a pro-fraud attitude commits fraud. We allow agents to repeatedly interact, resulting in changes in agent attitudes

[28] If our asymmetric influence model were re-conceptualized to examine a setting where fraudsters have more influence that honest agents, the lower support in U-shaped organizations would be eliminated, with fraud becoming the dominant behavior.

[29] Presumably, the introduction of a termination regime might also decrease agents' perceived opportunity or motivation to commit fraud. However, for the purposes of our analysis, we investigate effects independently as an initial step toward understanding our model.

[30] While we assume newly hired agents are randomly selected from the general population, one could also model an organization where hiring is biased toward a particular type of agent. For example, an organization of fraudsters might be more likely to hire another fraudster (or someone who believes fraud to be acceptable behavior).

Table 4
Summary of findings.

Intervention	Effect on U-shaped organization	Effect on inverted U-shaped organization
Perceived opportunity	Upper support lowered Dynamics preserved Spontaneous outbreaks persist	Average fraud level reduced Dynamics preserved
Tone at the top	Upper support removed Spontaneous outbreaks persist	Average fraud level reduced Dynamics preserved
Asymmetric influence	**Highly effective** Upper support removed **Stops spontaneous outbreaks**	Average fraud level reduced Dynamics preserved
Termination	Upper support removed Spontaneous outbreaks persist	**Highly effective** Fraud reduced to near zero **No spontaneous outbreaks**

toward fraud (via cultural transmission). We find that two types of organizations emerge, depending on how susceptible individual agents are to social influence. When average susceptibility to influence is low, a state distribution emerges where aggregate fraud levels tend toward a specific level of fraud. When average susceptibility is moderate to high, we observe an outcome characterized by extreme swings in behavior from an overall honest organization to an overall fraudulent one.

After identifying these two types of organizations, we turn our attention to the effect of mechanisms to prevent, detect, and eliminate fraud. We examine the effects of decreasing perceived opportunity to commit fraud, tone at the top, asymmetric influence (resulting from ethical interventions), and the removal of fraudsters from the organization. We find the effectiveness of these mechanisms is contingent on the type of organization. Our findings are summarized in Table 4.

The model suggests that ethical interventions aimed at increasing the influence of honest employees relative to fraudsters are particularly effective for moderate to high average emulation likelihood organizations. However, ethical interventions are observed to be much less effective for organizations with low emulation likelihoods. A different result emerges when we consider the impact of efforts to detect and remove fraudsters from the organization. Fraud is reduced to near zero in both types of organizations when the likelihood of terminating fraudsters is increased to 0.20. However, increasing the likelihood of termination does not prevent occasional outbreaks of widespread fraud in U-shaped organizations. The introduction of a monolithic, honest management team (tone at the top) effectively reduces the amount of fraud in both types of organizations. Furthermore, while it eliminates the tendency to achieve maximum levels of fraud with high frequency in U-shaped organizations, it does not eliminate sporadic, short-lived fraud outbreaks. Reducing opportunity (e.g., through stronger internal controls) merely reduces the average mean level of fraud and associated standard deviation within an organization. In our model, reduced opportunity does not affect other aspects of fraud dynamics, regardless of emulation likelihood levels.

The intuition behind these results lies in the relation between each intervention and the emulation likelihood, q. Reduced opportunity has no impact on q in our model and fraud dynamics are not affected under the intervention.[31] Every other intervention affects the dynamics of fraud. In every case, the effects on dynamics can be traced to the interaction of each intervention with q. Ethical interventions (asymmetric influence) directly impact the value of q. Tone at the top affects q for some interactions (those between a manager and staff) and termination introduces new agents (each with a new q) into the organization. Since each intervention affects q differently, they each have a different impact on fraud dynamics.

Our findings have important implications for auditors and other individuals responsible for assessing fraud risk and detecting and preventing fraud. First, for certain types of organizations aggregate fraud levels can vary tremendously over time. Furthermore, the effectiveness of mechanisms to prevent and detect fraud can be contingent on the type of organization and related individual susceptibilities to social influence. Therefore, it may be inappropriate for auditors to evaluate fraud prevention and detection mechanisms in a uniform manner. Our results suggest that the same fraud prevention and detection mechanisms implemented in a similar manner in two different organizations cannot be expected to be equally effective without considering the average susceptibilities to social influence of the individuals therein.[32] Similarly, some mechanisms can appear effective most of the time but still be ineffective at preventing outbreaks of fraud. In general, there is no one-size-fits-all fraud prevention (and/or detection) mechanism and fraud risk may be contingent on individual susceptibilities to social influence in the organization.

[31] The behavioral response to varying motivation in agents is the same as the behavioral response to changes in opportunity because, like opportunity, motivation is independent of q. Consequently, decreasing motivation decreases the average level of fraud. In the case of U-shaped organizations, the decrease is exhibited through a reduction in the maximum level of the upper support in the histogram. In the case of inverted U-shaped organizations, the distribution is simply shifted toward the origin. Because both motivation and opportunity are independent of q, we do not expect interaction effects between these variables and our model variants that do affect q.

[32] We conjecture that it might be possible to develop an instrument to measure differences in q across organizations (e.g., see Bearden, Netemeyer, & Teel, 1989) and that such an instrument might be useful to auditors in their assessment of fraud risk and the relative effectiveness of anti-fraud interventions.

While the ABM method used in our investigation confers a number of advantages over traditional models in the social sciences, it presents limitations as well. First, seemingly innocuous assumptions have sometimes been shown to have unexpected consequences that are not yet fully understood (e.g., Huberman & Glance, 1993). In the context of our model, additional research might examine alternative updating rules (e.g., one alternative model might randomly select an agent and then randomly determine whether the agent will interact with another agent or make a decision about committing fraud). The results we report may also be sensitive to other modeling choices. For example, we use Cressey's (1953) fraud triangle to represent the agent's decision to commit fraud. Other formulations are possible, such as the fraud diamond (Wolfe & Hermanson, 2004) or perhaps a utility-maximizing approach. Our characterization of a social network assumes a static organization with uniform mixing for the sake of simplicity. An alternative model might assume a dynamic network with local mixing or local mixing together with random connections among agents. Other characterizations of social influence (e.g., normative social influence in lieu of paired recruitment) could also have an impact. The second limitation associated with ABM is that the solution concepts employed tend to be relatively weak and, to date, there has been little formal work on the replicability of results using different models (but see Axtell, Axelrod, Epstein, and Cohen (1996) for an exception). Finally, ABM still lacks a set of standard, generally accepted practices.

Our initial investigation of fraud leads to a variety of additional opportunities for future research. Our analysis examines specific changes to our model in isolation. Our design choice was driven by a desire to develop a deeper understanding of the effects of various approaches taken to control fraud; however, in organizations, specific interventions are seldom implemented alone. The relative efficacy of combining different interventions could be investigated.

A more detailed examination of each leg of the fraud triangle could be undertaken. Future research might refine our characterization of motivation by explicitly considering social psychological factors in the organization. Equity theory (Adams, 1965) suggests that an honest individual in an organization replete with fraudsters would be more likely to have a motivation to commit fraud in order to reestablish fairness in their relationships with co-workers and the organization. Our model could be modified to incorporate this effect by making motivation a function of the number of fraudsters in the agent's social network.

Our model ignores ancillary changes in motive, opportunity, and attitude that might result from the interventions that we investigate (or from social changes and other forms of regulation). For example, our model's representation of the detection and removal (termination) of fraudsters ignores the deterrent effect that can be generated by an active termination program and concomitant effects on motivation. Future research could extend our study to incorporate these correlated effects. Our model also ignores the magnitude of fraud (dollars lost). Instead, we focus on the number of fraudsters in the organization. Insights on the magnitude of fraud might be gained by adding a "dollars available" variable to the model. Perhaps this variable could be linked to one's position in the organization (manager versus staff). Finally, while our model allows for selection in the termination intervention, it does not select on the basis of relative fitness (allowing those more skilled at fraud to avoid detection). Our model could be extended to investigate fraud in this evolutionary context.

Acknowledgements

The authors acknowledge financial support provided the American Institute of Certified Public Accountants (AICPA) Center for Audit Quality. The paper benefited from comments provided by David Piercey and participants at the Queens University Fraud Conference, the University of Illinois Audit Symposium, and the Fraud in Accounting, Organizations and Society Conference. We also acknowledge comments from two anonymous reviewers.

References

Adams, J. (1965). Inequity in social exchange. *Advances in Experimental Social Psychology, 62*, 335–343.

American Institute of Certified Public Accountants Auditing Standards Board (2002). *Consideration of fraud in a financial statement audit. Statement on auditing standards*. New York: American Institute of Certified Public Accountants.

Association of Certified Fraud Examiners (2007). *Fraud examiners manual*. Austin, TX: Association of Certified Fraud Examiners.

Association of Certified Fraud Examiners (2010). *Report to the nations on occupational fraud and abuse*. Austin, TX: Association of Certified Fraud Examiners.

Axelrod, R. (1997). Advancing the art of simulation in the social sciences. In R. Conte, R. Hegselmann, & P. Terna (Eds.), *Simulating social phenomena* (pp. 21–40). Berlin: Springer.

Axtell, R., Axelrod, R., Epstein, J., & Cohen, M. (1996). Aligning simulation models: A case study and results. *Computational and Mathematical Organization Theory, 1*(2), 123–141.

Bearden, W., Netemeyer, R., & Teel, J. (1989). Measurement of consumer susceptibility to interpersonal influence. *The Journal of Consumer Research, 15*(4), 473–481.

Beasley, M. (1996). An empirical analysis of the relation between the board of director composition and financial statement fraud. *The Accounting Review, 71*(4), 443–465.

Bell, T., & Carcello, J. (2000). A decision aid for assessing the likelihood of fraudulent financial reporting. *Auditing: A Journal of Practice and Theory, 19*(1), 169–184.

Breiger, R., & Carley, K. (2003). *Dynamic social network modeling and analysis: Workshop summary and papers*. Washington, DC: National Academies Press.

Carpenter, T. (2007). Audit team brainstorming, fraud risk identification, and fraud risk assessment: Implications of Statement on Auditing Standards No. 99. *The Accounting Review, 82*(5), 1119–1140.

Chang, J., & Lai, C. (2002). Is the efficiency wage efficient? The social norm and organizational corruption. *Scandinavian Journal of Economics, 104*, 27–47.

Cleary, R., & Thibodeau, J. (2005). Applying digital analysis using Benford's law to detect fraud: The dangers of type I errors. *Auditing: A Journal of Practice and Theory, 24*(1), 77–82.

Cressey, D. (1953). *Other people's money: A study in the social psychology of embezzlement*. Belmont, CA: Wadsworth Publishing.

Davis, J., Hecht, G., & Perkins, J. (2003). Social behaviors, enforcement, and tax compliance dynamics. *The Accounting Review, 78*(1), 39–69.

Epstein, J. (1999). Agent-based computational models and generative social science. *Complexity, 4*, 41–60.

Epstein, J. (2006). *Generative social science: Studies in agent-based computational modeling*. Princeton, NJ: Princeton University Press.

Epstein, J., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom up*. Washington, DC: The Brookings Institution.

Erickson, M., Hanlon, M., & Maydew, E. (2006). Is there a link between executive equity incentives and accounting fraud? *Journal of Accounting Research, 44*(1), 113–143.

Finley, D. (1994). Game theoretic analysis of discovery sampling for internal fraud control auditing. *Contemporary Accounting Research, 11*(1), 91–114.

Gaylord, R., & Davis, J. (1999). Modeling nonspatial socioeconomic interactions: An agent-based approach. *Mathematica in Education and Research, 8*(2), 30–33.

Gillett, P., & Uddin, N. (2005). CFO intentions of fraudulent financial reporting. *Auditing: A Journal of Practice and Theory, 24*(1), 55–75.

Gode, D., & Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy, 101*(1), 119–137.

Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W., Railsback, S., et al. (2005). Pattern-oriented modeling of agent-based complex systems: Lessons from ecology. *Science, 310*, 987–991.

Hoffman, V., & Zimbelman, M. (2009). Do strategic reasoning and brainstorming help auditors change their standard audit procedures in response to fraud risk? *The Accounting Review, 84*(3), 811–837.

Holtfreter, K. (2005). Is occupational fraud "typical" white-collar crime? A comparison of individual and organizational characteristics. *Journal of Criminal Justice, 33*(4), 353–365.

Huberman, B., & Glance, N. (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences of the United States, 90*, 7716–7718.

Jones, K., Krishnan, G., & Melendrez, K. (2008). Do models of discretionary accruals detect actual cases of fraudulent and restated earnings? *Contemporary Accounting Research, 25*(2), 499–531.

Kim, Y., & Xiao, N. (2008). Fraudsim: Simulating fraud in a public delivery program. In L. Liu, & J. Eck (Eds.), *Artificial crime analysis systems: Using computer simulations and geographic information systems* (pp. 319–338). Hershey, PA: Information Science Reference.

Kirman, A. (1992). Whom or what does the representative agent represent? *Journal of Economic Perspectives, 6*(2), 117–136.

Matsumura, E., & Tucker, R. (1992). Fraud detection: A theoretical foundation. *The Accounting Review, 67*(4), 753–782.

Milgram, S. (1963). Behavioral study of obedience. *The Journal of Abnormal and Social Psychology, 67*(4), 371–378.

Milgram, S. (1965). Some conditions of obedience and disobedience to authority. *Human Relations, 18*(1), 57–76.

Miller, J., & Page, S. (2007). *Complex adaptive systems: An introduction to computational models of social life*. Princeton, NJ: Princeton University Press.

Pinto, J., Leana, C., & Pil, F. (2008). Corrupt organizations or organizations of corrupt individuals? Two types of organization-level corruption. *Academy of Management Review, 33*(3), 685–709.

Ragan, R. (2010). An embarrassment of riches: Parameter choice in agent-based models. Unpublished manuscript. Department of Political Science, Duke University.

Schneider, A., & Wilner, N. (1990). A test of audit deterrent to financial reporting irregularities using the randomized response technique. *The Accounting Review, 65*(3), 668–681.

Schubert, S., & Miller, T. (2008). Where bribery was just a line item. *The New York Times* (December 21).

Sharma, V. (2004). Board of director characteristics, institutional ownership, and fraud: Evidence from Australia. *Auditing: A Journal of Practice and Theory, 23*(2), 105–117.

Tesfatsion, L., & Judd, K. (2006). *Handbook of computational economics, volume 2: Agent-based computational economics*. Amsterdam: North-Holland.

Uecker, W., Brief, A., & Kinney, W. (1981). Perception of the internal and external auditor as a deterrent to corporate irregularities. *The Accounting Review, 56*(3), 465–478.

Wilks, T., & Zimbelman, M. (2004). Decomposition of fraud-risk assessments and auditors' sensitivity to fraud cues. *Contemporary Accounting Research, 21*(3), 719–745.

Wolfe, D., & Hermanson, D. (2004). The fraud diamond: Considering the four elements of fraud. *The CPA Journal, 74*(12), 38–42.

Wolfram Research, Inc. (2008). *Mathematica version 7.0*. Champaign, IL: Wolfram Research, Inc.