Dynamic Incentives for Risk Management

Tansu Alpcan

***Abstract*—Incentives play an important role in assessment and mitigation of operational risks in large-scale organisations with multiple autonomous divisions. This paper presents a game-theoretic approach to optimisation of resources for mitigating security and technology-related risks. In the large-scale organisations considered, the risk manager plays a motivator role and provides assistance or subsidies to divisions, which are modelled as selfish players of a strategic (noncooperative) game. It is assumed that neither the risk manager nor divisions have complete information on the results of their investments, and they learn them 'a posteriori'. Dynamic incentive schemes are presented which achieve certain objectives despite this lack of 'a priori' knowledge. The convergence properties of these iterative schemes are analysed. The distributed mechanisms and algorithms developed provide a basis for risk management policies and guidelines.**

I. Introduction

Risk management[1] is a multi-disciplinary field with both technical and organisational dimensions [12]. On the technical side, complex and networked systems play an increasingly important role in daily business processes. Hence, system failures and security problems have direct consequences for organisations both monetarily and in terms of productivity [20]. It is therefore a necessity for any modern organisation to develop and deploy technical solutions for improving robustness of these complex information technology (IT) systems with respect to failures (e.g. in the form of redundancies) and defending them against security threats (e.g. firewalls and intrusion detection/response systems).

However, even the best and most suitable technical solution will fail to perform adequately if it is not properly deployed and supported within the hosting organisation. In order to be successful in risk management, an organisation has to have proper information about its business processes and complex technical systems. In other words, the organisation should "observe" them as well as be influence their operation or "control" them [2]. In a large-scale organisation these two necessary requirements, which may seem easy to satisfy at first glance, pose significant challenges. An important reason behind this issue, beside organisational structure, is the underlying incentive mechanisms.

Autonomous yet interdependent divisions or units of a large organisation have often individual objectives and incentives that may not be as aligned in practice as the headquarters and executives wish. Each such unit may have a different perspective on risk management which directly affects deployment of technical or organisational solutions. Misaligned incentives also make observation and control of business and technical processes difficult for risk managers. Considering the complex interdependencies in today's technology and business, such a misalignment in incentives is not a luxury even a large-scale organisation can effort.

Game theory provides a rich set of mathematical tools and models for investigating multi-person strategic decision making where the players (decision makers) compete for limited and shared resources [5], [8]. Game theoretic approaches have significant potential in addressing the above described issues in risk analysis, management, and associated decision making [3], [9], [21]. Unsurprisingly, game theory enjoys an increased interest in the risk management community [4], [11], [14], [16]. A game theoretic approach has been developed for security and risk-related decision making and investments in [18], [19]. Mechanism design techniques have been utilised for modelling, analysing and solving problems in decentralised network resource allocation [13], [25]. It has also been applied to resource allocation in the context of engineering optimisation [10].

This paper focuses on development of dynamic incentive mechanisms for optimising investments to mitigate security and IT-related risks. A game-theoretic approach is adopted, where a risk manager acts as a motivator, who provides assistance or subsidies to divisions, which are modelled as selfish players of a strategic (noncooperative) game.

It is assumed that neither the risk manager nor divisions have complete information on the results of their actions (e.g. risk investments), and learn them a posteriori. Another underlying assumption is that players refrain from misleading the risk manager about their preferences and hence manipulating the mechanism. The current scheme can be protected by the designer by actively enforcing fraud detection schemes as divisions are part of an umbrella organisation. However, designing a version which is resistant to cheating by relying only on pricing schemes (i.e. without resorting to detection-punishment or auction-based mechanisms) is an interesting and open research direction.

Dynamic incentive mechanisms are developed that achieve certain objectives despite this lack of a priori knowledge. The resulting distributed schemes provide a basis for risk management policies and guidelines as well as for developing scalable computer-assisted risk management systems. The main contributions of the paper include:

- A strategic (noncooperative) game approach for analysis of incentives in security and IT risk management.
- Development of a dynamic risk management framework with iterative incentive schemes, which are rigorously analysed for their convergence properties.
- Addressing the information limitations within the mech-

T. Alpcan is with the Department of Electrical and Electronic Engineering, The University of Melbourne, VIC 3010, Australia. tansu.alpcan@unimelb.edu.au

This work was supported by Deutsche Telekom Laboratories.

[1]This paper focuses on operational risks organisations face rather than financial and insurance-related ones.

anisms developed.

II. Model

Consider an organisation with N *autonomous units*, which act as independent decision makers, and a risk manager, which oversees the risk management task of the entire organisation. The risk manager often represents a special organisational unit. This generic organisation may be a large-scale multinational enterprise (divisions versus the risk manager at the headquarters), a government (government agencies versus executive arm), or even an international organisation (individual countries versus general secretary of the organisation).

Adopting a game-theoretic approach, each autonomous unit can be modelled as a player of a strategic (noncooperative) game with the set of all players denoted as $\mathscr{A}$. The player $i \in \mathscr{A}$ independently decides on its respective decision variable x_i, which represents allocation of limited resources for the purpose of assessment and mitigation of risks in accordance with own objectives. Thus, the player actions are captured by the vector

$$x = [x_1, \ldots, x_N] \in \mathscr{X} \subset \mathbb{R}^N,$$

where $\mathscr{X}$ is the convex, compact, and non-empty decision space of all players. These decisions translate to monetary investments (buying equipment, hiring consultants) or assigning human resources (manpower) from within the division's own resources to address risk-related problems. In the majority of cases, the decisions of players affect each other due to constraints of the environment. Thus, the players share and compete for resources as part of this strategic game while interacting with the risk manager (mechanism designer) as visualised in Figure 1.

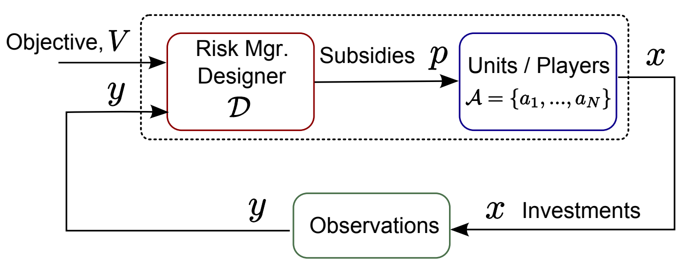

Fig. 1. Interaction between risk manager (designer) and organisational units (players) as part of incentive mechanism design.

The players make their decisions in accordance with their estimated preferences. The problem is, in many real life situations the outcomes of risk-related investments are quite uncertain. In other words, the players may not know the usefulness of their actions beforehand, and can only estimate them. This information limitation is an important factor that needs to be taken into account when developing risk management mechanisms in general.

Despite information limitations, a few basic assumptions can be made about the expected utility of investments:

- **A.1** The return on investments for risk assessment and mitigation is an increasing function. Investing more resources to the problem leads to a progress (however small) in addressing it, assuming that completely inefficient techniques are not chosen and deployed blindly [12].
- **A.2** It may not be possible to completely mitigate all of the risks in an organisation.
- **A.3** Initial, common-sense measures in risk management are relatively easy to deploy and useful. Each additional step after this, which potentially require substantial changes in the organisation, is more costly and brings diminishing returns.

All three assumptions above are captured by the real-valued utility function

$$U_i(x) : \mathscr{X} \to \mathbb{R}$$

of player i, which is monotonously increasing, and strictly concave in x_i. For analytical tractability, it is additionally assumed that U_i is also continuous and differentiable in all its arguments.

In practice, one way to obtain U is to use a point system based on an analysis similar to the one in [17]. An estimation on how actions (investments) affect (decrease) a standard risk-score that is adopted by the whole organisation can be used to derive the utility function. Another option is to prepare a set of requirements based on ISO 2700X standards for the organisation and then assess the level of satisfaction of these requirements in light of the organisation's priorities. For example a weighting factor can be used to emphasise certain requirements over others. The utility function again can be defined in terms of the (expected) change in aggregate compliance-level, defined e.g. as a weighted-sum of requirement satisfaction, due to the risk management investments and actions. The pitfalls of these methods as well as possible improvements over them are discussed in detail in [12]. The work [9, Chap. 3] investigates application of a similar utility function approach to risk management.

While each player obtains a benefit (utility) from its investments, these resources also have a cost, which can be often expressed in monetary terms. We assume that that these costs are linear in the allocated resource, $\beta_i x_i$, where β_i is the individual per unit cost factor.

Consequently, ach player i aims to minimise its respective cost function

$$J_i(x) = \beta_i x_i - U_i(x) - p_i x_i, \tag{1}$$

where the linear term $p_i x_i$ represents the *incentive factor* (or penalty if negative) provided to the player by the designer $\mathscr{D}$. Thus, player i solves the convex optimisation problem

$$\min_{x_i} J_i(x_i, x_{-i}),$$

by choosing an appropriate x_i given the decisions of all players denoted by x_{-i} such that $x \in \mathscr{X}$.

The Nash equilibrium (NE) is a widely-accepted and useful solution concept in strategic games, where no player has an incentive to deviate from it while others play according to their NE strategies [22], [23]. The NE is at the same time the intersection point of players' best responses obtained by solving their individual optimisation problems. The NE of the game is formally defined as follows.

Definition 1: The Nash equilibrium of the game is denoted by the vector $x^* = [x_1^*, \dots, x_N^*] \in \mathscr{X}$ and defined as

$$x_i^* := \arg\min_{x_i} J_i(x_i, x_{-i}^*) \quad \forall i \in \mathscr{A},$$

where $x_{-i}^* = [x_1^*, \dots, x_{i-1}^*, x_{i+1}^*, \dots, x_N^*]$.

If some special convexity and compactness conditions are imposed to the game $\mathscr{G}$, then it admits a unique NE solution, which simplifies mechanism and algorithm design significantly. We refer to [1], [5], [24] for the details and an extensive analysis.

The risk manager (designer) $\mathscr{D}$ devises an incentive mechanism $\mathscr{M}$, which can be represented by the mapping $\mathscr{M} : \mathscr{X} \to \mathbb{R}^N$, and implemented through additional incentives (e.g. subsidies) in player cost functions, $p_i x_i$, above. Using incentive mechanism $\mathscr{M}$, the designer aims to achieve a certain risk management objective, which can be maximisation of aggregate player utilities (expected aggregate benefit from risk-related investments) or an independent organisational target that depends on participation of all players such as deployment of a new risk management solution. These can be modelled using a designer objective function V defined as $V(x, U(x), p) : \mathscr{X} \to \mathbb{R}$. Thus, the global optimisation problem of the designer is

$$\max_p V(x, U_i(x), p),$$

which it solves by choosing the vector $p = [p_1, \dots, p_N]$, i.e. providing incentive factors to the players. subject to the budget constraint $\sum_{i=1}^N p_i x_i \leq B$.

III. Mechanism with a Global Objective

In many real life scenarios, the risk manager has an organisation-wide or "global" objective represented by the strictly concave and non-decreasing function $F(x)$ which does not directly depend on player utilities. This organisation-wide objective could be, for example, deployment and operation of an organisation-wide risk management solution that naturally requires cooperation from all units and an alignment of efforts. The risk manager formally solves the constrained optimisation problem

$$\max_x F(x) \text{ such that } \sum_i p_i x_i \leq B. \tag{2}$$

The associated Lagrangian function is then

$$L(x) = F(x) + \lambda \left(B - \sum_i p_i x_i \right),$$

where $\lambda > 0$ is a scalar Lagrange multiplier. Note that the constraint is always active in this case due to the definition of $F(x)$.

Combining this with the player optimisation problems to ensure efficiency and preference-compatibility (player best responses) leads to

$$\frac{1}{p_i} \frac{\partial F(x)}{\partial x_i} = \lambda, \; \forall i \in \mathscr{A}, \tag{3}$$

and

$$\sum_i p_i \left(\frac{\partial U_i(x)}{\partial x_i} \right)^{-1} (\beta_i - p_i) = B. \tag{4}$$

Note that, any solution constitutes a Nash equilibrium as it lies at the intersection of the player best responses.

Proposition 1: Any solution of the mechanism with global objective $F(x)$ described in (3)-(4) is both player preference-compatible, i.e. a Nash equilibrium of the defined strategic game and efficient, i.e. maximises $F(x)$.

Consider the special case of the player utilities $U_i(x) = \alpha_i \log(x_i) \; \forall i$, the same costs constant for all players $\beta_i = b \; \forall i$, and the system objective

$$F(x) = \log \left(\sum_i x_i - \gamma \right),$$

where γ is a scalar threshold value. This objective can be interpreted as the total investment in the organisation satisfying at least the threshold γ and after that its benefit following the law of diminishing returns. Solving (3)-(4) under the assumption of sufficient budget, $B > \gamma\beta - \sum_i x_i$, yields the unique solution:

$$\lambda = \frac{\beta(\sum_i \alpha_i + 1)}{B(\sum_i \alpha_i + B - \gamma\beta)}, \quad p_i = \frac{\beta B}{\sum_i \alpha_i + B} \; \forall i. \tag{5}$$

A. Iterative Scheme

In the mechanism presented, the players make their decisions in accordance to their estimated preferences. The problem is, in many real life situations, the outcome of risk-related investments to assess and mitigate risks is quite uncertain. It is possible to overcome such information limitations by adopting an iterative scheme. Specifically, the risk manager updates the Lagrangian multiplier λ gradually according to

$$\lambda(n+1) = \lambda(n) + \kappa_d \left[\sum_i p_i(n) x_i(n) - B \right]^+, \tag{6}$$

and computes the individual player incentive factors

$$p_i(n) = \frac{1}{\lambda(n)} \frac{\partial F(x(n))}{\partial x_i}. \tag{7}$$

Here, $n = 1, \dots$ denotes the iteration number or time-step. The units (players) in return react to given incentive factors by updating their investment decisions in order to minimise their own costs such that

$$\begin{aligned} x_i(n+1) &= \phi x_i(n) \\ &+ (1-\phi) \left(\frac{\partial U_i(x(n))}{\partial x_i} \right)^{-1} (\beta_i - p_i(n)) \; \forall i, \end{aligned} \tag{8}$$

where $0 < \phi < 1$ is a relaxation constant used by the players to prevent excessive fluctuations. Alternatively, this behaviour can be justified with caution or inertia of the organisational units.

Proposition 2: The equilibrium solution(s) of (6)-(8) coincides with that of (3)-(4). Hence, the iterative mechanism solves the same problem as the original mechanism.

Information flow and limitations play a crucial role in implementation of the iterative mechanism. In practice, the risk manager is assumed to observe the actions of units which they have to reveal in order to receive incentives. Based on this information and the total budget, the risk manager can easily implement (6). Then, it only needs to estimate the individual marginal contributions of units to the overall objective, $\partial F(x(n))/\partial x_i$ at a given moment in order to decide on actual incentive factors in (7).

Likewise, given own cost estimates β_i and incentive factor p_i, each unit (player) only has to determine the marginal benefit from its own actions, $\partial U_i(x(n))/\partial x_i$ in order to implement (8). If the unit has a separable utility, then this is simply equivalent to $\partial U_i(x_i(n))/\partial x_i$. Algorithm 1 summaries the steps of the iterative mechanism with global objective.

Algorithm 1: Iterative mechanism

Input: *Designer*: budget B and global objective $F(x)$
Input: *Players*: cost factor β_i and utilities $U_i, \forall i$
Result: Player investments x and incentive factors p

Initial investments x_0 and incentive factors p_0 ;
repeat
  begin *Designer:*
    Observe player investments x ;
    Update λ according to (6) ;
    Estimate marginal contributions of players to global objective, $\partial F(x)/\partial x_i$;
    foreach *player* i **do**
      Compute incentive factor p_i from (7) ;
    end
  end
  begin *Players:*
    foreach *player* i **do**
      Estimate marginal utility $\partial U_i(x)/\partial x_i$;
      Compute investment x_i from (8) ;
    end
  end
until *end of iteration (negotiation)*;

B. Convergence Analysis

A basic stability analysis is provided for a continuous-time approximation of the iterative mechanism with global objective. For tractability, let the player utilities be of the form $U_i = \alpha_i \log(x_i)$. Further define the global objective function of the risk manager as $F(x) := \sum_i \gamma_i x_i$, for some $\gamma_i > 0 \, \forall i$.

Substituting p_i with γ_i/λ, the continuous-time counterpart of (6)-(8) is

$$\dot{\lambda} = \frac{d\lambda}{dt} = \kappa_\lambda \frac{1}{\lambda}\left(\sum_i \gamma_i x_i(n) - B\right) \tag{9}$$

$$\dot{x}_i = -\kappa_i \frac{\partial J_i}{\partial x_i} = \kappa_i \left(\frac{\alpha_i}{x_i} + \frac{\gamma_i}{\lambda} - \beta_i\right) \;, \forall i \in \mathscr{A}.$$

where t denotes time and κ_λ, $\kappa_i > 0$ are step-size constants. As in the discrete-time version, the players adopt here a gradient best response algorithm. Define the Lyapunov function

$$V_L := \frac{1}{2}\left(\frac{\sum_i \gamma_i x_i - B}{\lambda}\right)^2 + \frac{1}{2}\sum_i \left(\frac{\alpha_i}{x_i} + \frac{\gamma_i}{\lambda} - \beta_i\right)^2,$$

which is positive definite except at the solution(s) of (6)-(8), i.e. $V_L(x^*, \lambda^*) = 0$.

Taking the derivative of V_L with respect to time yields

$$\dot{V}_L(x,\lambda) = -2\frac{\sum_i \gamma_i x_i}{\lambda^3}\left(\frac{\sum_i \gamma_i x_i - B}{\lambda}\right)^2 - \sum_i \frac{\alpha_i}{x_i^2}\left(\frac{\alpha_i}{x_i} + \frac{\gamma_i}{\lambda} - \beta_i\right)^2.$$

Thus,

$$\dot{V}_L(x,\lambda) < 0, \;\; \forall (x,\lambda) \neq (x^*, \lambda^*),$$

i.e. for any point of the trajectory (x, λ) not equal to a solution of (6) and (8). Consequently, the continuous-time algorithm is asymptotically stable [15]. This result, which is summarised in the next proposition, is a strong indicator of the convergence [7] of the discrete-time iterative pricing mechanism.

Proposition 3: The continuous-time approximation of the iterative mechanism, given by (9), asymptotically converges to a solution of (6)-(8).

The *convergence* result above indicates that the investment levels always approach their equilibrium values.

IV. Welfare Maximising Mechanism

A counterpart of the mechanism in the previous section is based on the risk manager objective of maximising the sum of player utilities, $\sum_i U_i(x_i)$. This objective can be achieved following a methodology similar to the one above. Formally, the designer solves the constrained optimisation problem

$$\max_x V(x) \Leftrightarrow \max_x \sum_i U_i(x) \text{ such that } \sum_i p_i x_i \leq B. \tag{10}$$

The optimal solution to this constrained problem by definition satisfies the efficiency criterion. The associated Lagrangian function is then

$$L(x) = \sum_i U_i(x) + \lambda\left(B - \sum_i p_i x_i\right),$$

where $\lambda \geq 0$ is a scalar Lagrange multiplier [6].

The alignment of player and designer optimisation problems is achieved by choosing the Lagrange multiplier λ and player incentive factors p in such a way that

$$\frac{\beta_i - p_i}{p_i} + \frac{1}{p_i}\sum_{j \neq i} \frac{\partial U_j(x)}{\partial x_i} = \lambda, \; \forall i \in \mathscr{A}, \tag{11}$$

and

$$\sum_i p_i \left(\frac{\partial U_i(x)}{\partial x_i}\right)^{-1} (\beta_i - p_i) = B. \tag{12}$$

Any solution to the set of $N+1$ non-linear equations (11)-(12) is by definition a Nash equilibrium as it lies at the intersection of the player best responses. These results are summarised in the following proposition.

Proposition 4: Any solution of the welfare maximising mechanism described above obtained from (11)-(12) is both player preference-compatible, i.e. a Nash equilibrium of the strategic game and efficient, i.e. maximises $\sum_i U_i(x)$.

In the special case of each player i having the utility $U_i(x) = \alpha_i \log(x_i)$, the system (11)-(12) admits the following unique solution, which can be derived using simple algebraic manipulation:

$$\lambda = \frac{\sum_i \alpha_i}{B}, \quad p_i = \frac{\beta_i B}{\sum_i \alpha_i + B} \;\; \forall i. \tag{13}$$

In the case of iterative welfare maximising mechanism, the risk manager updates the Lagrangian multiplier λ according to (6) and the unit updates are given by (8).

However, the computation of individual player incentive factors is more involved due to the dependence of the objective (welfare maximisation) on individual player utilities

$$p_i(n) = \frac{1}{\lambda(n)} \sum_j \frac{\partial U_j(x(n))}{\partial x_i}, \tag{14}$$

which follows from (11). At first glance, it seems that the designer has to ask players again for their marginal utility which defeats the purpose of the iterative approach, namely ensuring strategy-proofness. Fortunately, the designer can circumvent this issue by utilising side information, in this case player cost factors β.

The actions of any player i chosen according to a (relaxed) best response (8), and observed by the designer yields the information

$$\frac{\partial U_i(x)}{\partial x_i} = \beta_i - p_i$$

to the designer. Hence, the substitution

$$\sum_j \frac{\partial U_j(x(n))}{\partial x_i} = \sum_j (\beta_i - p_i) \; \forall i, j\,,$$

in (14) yields $p_i = \dfrac{b_i}{\lambda_i + 1} \; \forall i$, which together with (6) is used to determine player incentive factors.

Due to space limitations, the convergence analysis of the iterative welfare maximising algorithm is omitted.

V. Conclusion

A game-theoretic approach to optimisation of resources for mitigatation of security and technology-related risks is presented. In the large-scale organisations considered, the risk manager plays a motivator role and provides assistance or subsidies to divisions, which are modelled as selfish players of a strategic (noncooperative) game. Dynamic incentive schemes are presented which achieve certain objectives and their convergence properties are analysed. The distributed mechanisms and algorithms developed provide a basis for risk management policies and guidelines.

An underlying assumption of the considered model is the fixed nature of player preferences or utility functions. Under this assumption, the risk manager can influence unit decisions only by introducing additive incentive factors to their cost structure as discussed. In reality however, the unit preferences are open to changes through psychological factors. The arts of persuasion and politics may "shift" the utility curves in the model. Quantification of such factors is obviously a significant yet open research challenge.

Another underlying assumption is that players refrain from misleading the risk manager about their preferences and hence manipulating the mechanism. Designing a version which is inherently resistant to cheating is another open research direction.

References

[1] T. Alpcan, "Noncooperative games for control of networked systems," Ph.D. dissertation, University of Illinois at Urbana-Champaign, Urbana, IL, May 2006.

[2] T. Alpcan and T. Başar, *Network Security: A Decision and Game Theoretic Approach.* Cambridge University Press, 2011. [Online]. Available: http://www.tansu.alpcan.org/book.php

[3] T. Alpcan and N. Bambos, "Modeling dependencies in security risk management," in *Proc. of 4th Intl. Conf. on Risks and Security of Internet and Systems (Crisis)*, Toulouse, France, October 2009.

[4] E. P.-C. and, "Risks and games: Intelligent actors and fallible systems," in *Proc. of 2nd Intl. Symp. on Engineering Systems*, Cambridge, MA, USA, June 2009.

[5] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*, 2nd ed. Philadelphia, PA: SIAM, 1999.

[6] D. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA: Athena Scientific, 1999.

[7] D. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Compuation: Numerical Methods*. Upper Saddle River, NJ: Prentice Hall, 1989.

[8] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, 1991.

[9] P. R. Garvey, *Analytical Methods for Risk Management: A Systems Engineering Perspective*, ser. Statistics: textbook and monographs. Boca Raton, FL, USA: Chapman and Hall/CRC, 2009.

[10] S. D. Guikema, "Incentive compatible resource allocation in concurrent design," *Engineering Optimization*, vol. 38, no. 2, pp. 209–226, March 2006.

[11] ——, "Game theory models of intelligent actors in reliability analysis: An overview of the state of the art," in *Game Theoretic Risk Analysis of Security Threats*. Springer US, 2009, pp. 1–19.

[12] D. W. Hubbard, *The Failure of Risk Management: Why It's Broken and How to Fix It*. Hoboken, NJ, USA: John Wiley & Sons, 2009.

[13] R. Johari and J. N. Tsitsiklis, "Efficiency of scalar-parameterized mechanisms," *Operations Research*, vol. 57, no. 4, pp. 823–839, July 2009. [Online]. Available: http://www.stanford.edu/~rjohari/pubs/char.pdf

[14] J. John R. Hall, "The elephant in the room is called game theory," *Risk Analysis*, vol. 29, no. 8, pp. 1061–1061, June 2009.

[15] H. Khalil, *Nonlinear Systems*, 3rd ed. Prentice Hall, 2002.

[16] J. Louis Anthony (Tony) Cox, "Game theory and risk analysis," *Risk Analysis*, vol. 29, no. 8, pp. 1062–1068, June 2009.

[17] McAfee, "Information security metrics: Using mcafee foundstone foundscore to assign metrics and measure enterprise risk," white paper, October 2006. [Online]. Available: www.mcafee.com

[18] R. A. Miura-Ko, B. Yolken, N. Bambos, and J. Mitchell, "Security investment games of interdependent organizations," in *46th Annual Allerton Conference*, Sept. 2008.

[19] R. A. Miura-Ko, B. Yolken, J. Mitchell, and N. Bambos, "Security decision-making among interdependent organizations," in *Proc. of 21st IEEE Computer Security Foundations Symp. (CSF)*, June 2008, pp. 66–80.

[20] D. Moore, C. Shannon, and K. Claffy, "Code-Red: A case study on the spread and victims of an Internet worm," in *Proc. of ACM SIGCOMM Workshop on Internet Measurement*, Marseille, France, 2002, pp. 273–284.

[21] J. Mounzer, T. Alpcan, and N. Bambos, "Dynamic control and mitigation of interdependent IT security risks," in *Proc. of the IEEE Conference on Communication (ICC)*. IEEE Communications Society, May 2010.

[22] J. F. Nash, "Equilibrium points in n-person games," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 36, no. 1, pp. 48–49, January 1950. [Online]. Available: http://www.jstor.org/stable/88031

[23] ——, "Non-cooperative games," *The Annals of Mathematics*, vol. 54, no. 2, pp. 286 – 295, September 1951. [Online]. Available: http://www.jstor.org/stable/1969529

[24] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games," *Econometrica*, vol. 33, no. 3, pp. 520–534, July 1965.

[25] S. Yang and B. Hajek, "VCG-Kelly mechanisms for allocation of divisible goods: Adapting VCG mechanisms to one-dimensional signals," *IEEE JSAC*, vol. 25, no. 6, pp. 1237–1243, August 2007.