An Example Collaborative Exercise for Decision Making in Investment in Cyber Security

By: Jonathan Crawford, Kenneth Crowther, Barry Horowitz, James Lambert

Abstract

The economics of investment in cyber security is a widely researched field. This paper describes the use of a multi-player collaborative exercise implemented on computers to help companies better understand investment decisions in cyber security. The investment model driving the collaborative exercise is an expected-value decision analysis that compares the reduction of cyber risks with other investment opportunities and accounts for the potential of government regulatory action when an integrated national impact of attacks exceeds certain acceptable levels. The exercise was implemented with over twenty live participants in June 2006 at a workshop of the Institute for Information Infrastructure Protection (I3P) addressing Process Control Systems (PCS) Security. The aim of the exercise was to illustrate the impact of potential government regulation on the complex decision process of determining appropriate investment levels for added cyber security by individual companies. At the workshop the exercise provided an opportunity for knowledgeable security professionals to collaborate and compare their investment decisions against those of other similar companies and against the results of the expected value decision analysis. This paper describes the foundations of the exercise and an hypothetical interpretation, by a company that would employ the exercise, of the results from its application at the PCS workshop.

1. Introduction

This research is motivated by the lack of a structured approach to investing in cyber security [Gordon and Loeb 2006]. Decisions on cyber security investments are characterized by a number of unknown factors that are critical for determining investment levels based on rigorous economic analysis, including the likelihoods of being successfully attacked with and without added cyber security measures and the potential consequences of the successful attacks with and without added cyber security measures. Unlike natural disasters, diseases, and many other areas of risk to human welfare, there is no organized historical data to make estimates about cyber attacks and, for numerous reasons, companies do not share information about attacks on their systems that have succeeded and their consequences. Since these unknowns are central to a rigorous risk analysis, companies resort to unstructured and ad hoc methods for deciding the level of investment in cyber security that is appropriate for them. Among many other possible sets of data that are available, any particular company might use such information as: 1) reports on corporate trends in spending on cyber security as produced by companies such as Forrestor or Gartner Group, 2) internet-based information on the number of software patches to correct existing software systems or 3) internal information on the number of attempts to penetrate their firewalls, to gauge likelihoods of attacks. Their decisions may also be influenced by media reports on actual attacks and the reported consequences on the businesses of the successfully attacked companies. Lacking a structure for decision-making, it is difficult for a company to evaluate its decisions after a successful attack occurs, and over time, to sequentially improve its process for decision-making. The work described in this paper is progress toward constructing a methodology that could help companies to arrive at decisions in a more structured manner. The suggested approach includes having company "experts" in a variety of relevant subjects, make assumptions about the unknowns that permit a rigorous analysis to be carried out. These assumptions and decisions would then be reexamined annually or as frequently as warranted to reflect the increase in company knowledge or a changing operating environment. With these assumptions, a live simulation exercise such as the one described in this paper, can be carried out by companies to explore how their management team might decide on investing based on available data, experience and intuition. The performance of the team in the simulations can be compared to normative performance that is based on mathematical decision analysis using plausible assumptions of uncertain parameters.

1.1 Experimental Economics

Experimental economics is the field of economics that tests propositions of economic theory in controlled settings. Vernon Smith, Nobel Laureate is considered by many to be the father of experimental economics [Lynch 2002]. Smith has written of seven prominent reasons to conduct experimental games in economics [Smith 1994].

1. To test a theory or discriminate between theories [Smith, Cox, and Roberson 1982].
2. To explore the causes of a theory’s failure [Roth 1987].
3. To establish empirical regularities as a basis for a new theory [Smith 1976].
4. To compare environments [Kagel and Levin 1986].
5. To compare institutions [Smith 1984].
6. To evaluate policy proposals [Smith 1967].
7. To provide a laboratory environment for uses as a test bed of a new market [Smith 1991, pp 511-512].

The creation of the cyber security exercise as a game can derive benefits in all of these areas. Experimental economics has been used to quantify the privacy concerns of individuals (Baumer, Earp, and Poindexter 2005). The use of analytical models can be tested. Participant's reasoning is gathered and those participants who fail to act as the decision model would predict can have their rationale examined. If a large deviation from the model is observed, explanations can be postulated. The exercise can be replayed by participants representing different jobs in an organization (e.g., security managers and financial managers). The exercise can be used to compare industries. The economics of cyber security investment is a relatively new field and by comparing it with the other potential uses for an institution's resources, the market for investment in cyber security can be examined and advanced.

1.2 Initial Collaborative Exercise: Intellectual Property Protection

As a precursor to creating the more complete multi-player cyber security investment exercise described later in this paper, an intellectual-property protection model was constructed that offered participants a choice of added investment to enhance cyber security for intellectual property stored on computers or alternatively, to add sales force to achieve added revenue [Horowitz 2006]. The decision model behind the exercise was adapted from an expected value model that describes the potential for under investment in cyber security by the Internet service providers [Garcia and Horowitz 2006]. The exercise was conducted at a workshop with 23 security professional participants. The exercise included three investment scenarios, which used particular parameters for attack likelihoods, consequences, etc, so that the expected value results for the three cases were: 1) the security and sales results had equal economic expectations, 2) the security investment provided greater expected value and 3) the sales investment provided greater expected value. The following conclusions could be drawn by a company that conducted the exercise: 1) players showed a bias towards investing in increasing revenue versus increasing the level of cyber security (i.e., in the case of equal expected outcomes, most participants elected to invest in added sales), 2) players behaved rationally by shifting toward cyber investing as the expected value of the cyber security investment exceeded the potential increase in revenue, and 3) those motivated toward the increased cyber security investment were not moved even when the expected value of the investment to increase revenue exceeded the expected value of the security investment. The participants indicated that they thought that the creation of a decision support aid that would integrate multiple areas of company knowledge involves creating relationships between those areas and cyber security was a needed and a potentially valuable area to pursue. This first exercise established a possible decision-making relationship between the legal and business managers focused on intellectual property and the cyber security managers. By constructing the exercise with a basis in economic games it allowed the analytical decision model(s) to be tested by real decision makers in a manner that can aid in the construction and refinement of the model and the exercise environment, hopefully reaching the point where actual decisions can be influenced by its use.

1.3 Influence of Government on Industrial Cyber Security Investments

The exercise that is the focus of this paper extends the intellectual property model with an analytical model that supposes a possible relationship between government regulation and investment in cyber security. The analytical model of the exercise is still based on an expected value analysis of the possible decisions and companies still act independently. However, in this game, a company’s decision affects an entire industry and thus the community’s individual decisions have a collective impact. A company making an investment decision must take this into account and conduct an analysis of its community. For the PCS Workshop, the Oil and Gas industry, one of our country's critical infrastructures, was the focus of the exercise. To create the exercise a model was needed to provide participants with a measure of the possibility of government regulation. In addition, the participants needed to know what the possible regulations might entail. Section 2 describes the models used to create the economic exercise for the Oil and Gas industry.

2. Cyber Security Investment Model Accounting for Government Regulation

The Government Regulation exercise is similar to the Intellectual Property exercise in that participants decide on investing in either enhanced cyber security to avoid a potential shutdown due to a cyber attack on a production plant's automation system, or on additional plant automation capabilities that would increase revenue. The Government Regulation exercise differs from the Intellectual Property model described above in that it includes a model for the government considering and possibly imposing a significant fine on companies if they are successfully attacked. The government decides to intervene once the number of significant successful attacks on the industrial sector of concern has exceeded what is deemed as acceptable, causing it to institute a fine in the next investment period to successfully attacked companies in that sector. A probabilistic model was developed to represent government actions. The following list contains the variables and their definitions used in the derivation of the government model.

α = chance of successful defense given ENHANCED SECURITY investment (for given year)
β = chance of successful defense given ENHANCED REVENUE investment (for given year)
γ = the number of significant successful attacks to the industry until the GOVERNMENT institutes a fine
S = yearly cost of ENHANCED SECURITY investment
F = yearly cost of ENHANCED REVENUE investment
V = lost REVENUE from a successful cyber attack
C_α = number of companies investing in ENHANCED SECURITY
C_β = number of companies investing in ENHANCED REVENUE
G = cost of GOVERNMENT FINE
X = the company's yearly revenue before the additional investment options
N = the total number of years (stages) in the investment horizon under consideration
R_α = total revenue over the N period resulting from investment in ENHANCED SECURITY
R_β = total revenue over the N period resulting from investment in ENHANCED REVENUE
r = the total return from REVENUE investment

There are two options available to decision makers in this model. They can choose to invest at a cost of S in *enhanced cyber security* which will result in a probability α of successfully defending against cyber attack that would result in a loss of V revenue. Alternatively, they can choose to invest at a cost of F in to *enhanced revenue*, but in turn will result in a probability β of successfully defending against cyber attack that would result in the same loss of V revenue. The parameter α is assumed to be greater than β. The government is observing the industry the company is part of (Oil and Gas for the Workshop) and after a particular number of successful attacks γ, the government will institute fines that further penalizes successfully attacked companies regardless of their investment decisions. To account for the two states of the fine not being active and the fine being instituted, the model considers a horizon of N years of annual decisions. The model assumes that when deciding on a given year's investment decision, decision makers plan to make the same decision over all of the remaining years for the selected time horizon. We consider it a reasonable assumption that a company would not expect to change its decision every year as part of deciding on a given years' investments. However, depending on year-by-year game results, decision-makers will indeed change their mind, each time assuming the decisions will be fixed for the entire remaining time horizon. Equation (1.0) formulates the expected total revenue over the N years if the decision making company invests in enhanced cyber security.

$$\boxed{\begin{aligned} \mathrm{E}(R_\alpha) = {} & \left\lceil \frac{\gamma}{(1-\alpha)*(\mathrm{C}_\alpha + 1) + (1-\beta)*\mathrm{C}_\beta} \right\rceil * (\alpha*(X-S) + (1-\alpha)*(X-S-V)) + \\ & \left(N - \left\lceil \frac{\gamma}{(1-\alpha)*(\mathrm{C}_\alpha + 1) + (1-\beta)*\mathrm{C}_\beta} \right\rceil \right) * (\alpha*(X-S) + (1-\alpha)*(X-S-V-G)) \qquad (1.0) \end{aligned}}$$

At first, Equation (1.0) enables calculating the first year that the government fine will be in use. This is done by the calculation enclosed within the left and right ceiling operators. The number of successful attacks the government is willing to tolerate is divided by the expected number of successful attacks per year in the industry. The expected number of successful attacks per year in the industry depends on the number of companies investing in either enhanced cyber security or enhanced revenue. The probability of a successful cyber attack for each of the companies investing in enhanced cyber security is 1-α and, assuming attacks upon companies are independent of each other, the total number of attacks expected from this segment of the industry is just this probability multiplied by the population of companies C_α investing in enhanced cyber security (which, in this case, includes the decision making company as well). The same logic holds for the number of attacks expected from the population of companies investing instead in enhanced revenue. The probability for members of this population of companies C_β being successfully attacked is 1-β and the total expected number of successful attacks from this segment of the industry is the probability multiplied by C_β. Taking the expected number of successful attacks for each segment of the industry, those investing in enhanced cyber security and those investing in enhanced revenue, and adding them together gives the expected number of successful attacks to the entire industry per year. To calculate when the government tolerance, as measured by the threshold of successful attacks, is expected to be reached and a government fine system instituted involves dividing γ by the total expected yearly successful attacks to the industry. The result will be the expected number of years until the threshold of γ is met. The ceiling operators round the number up to the next integer because the model assumes the government will institute the fine in the next decision period and not during the middle of the year. The decision making company can then assume that it will operate for that many years without the possibility of a fine and for N minus that many years with the possibility of being fined. The total expected revenue over the N year period $E(R_\alpha)$ is thus made up of the time before the fine and the time after the fine has been instituted. Without fines, each year the decision making company will receive X-S in revenue or X-S-V if successfully attacked, with V being the amount of revenue the attack causes to be lost if the company is successfully attacked. When the fine is in place the decision making company will still receive X-S unless successfully attacked, and then it will receive X-S-V-G, where G is the amount of the government's fine. Equation (2.0) describes the expected total revenue over N years if the decision making company decides to invest in enhanced revenue.

$$\begin{aligned} E(R_{\beta}) = &\left\lceil \frac{\gamma}{(1-\alpha)*C_{\alpha} + (1-\beta)*(C_{\beta}+1)} \right\rceil * (\beta*(X+r-F) + (1-\beta)*(X+r-F-V)) + \\ &\left(N - \left\lceil \frac{\gamma}{(1-\alpha)*C_{\alpha} + (1-\beta)*(C_{\beta}+1)} \right\rceil \right) * (\beta*(X+r-F) + (1-\beta)*(X+r-F-V-G)) \qquad (2.0) \end{aligned}$$

Equation (2.0), the total expected revenue over the N period if the company decides to invest in enhanced revenue $E(R_{\beta})$, is similar to equation (1.0). However the calculation of the year the government threshold is expected to be reached now reflects the decision making company being part of the segment of the industry investing in enhanced revenue. The calculation of the expected yearly revenue is also different in that the amount of revenue the company will receive is now X+r-F or X+r-F-V if the company is successfully attacked.

Recognizing that many of the terms in the equations can be simplified equation (1.0) simplifies to equation (1.2) and equation (2.0) simplifies to equation (2.2).

$$\mathrm{E}(R_{\alpha}) = \left[\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right]*(\alpha*(X-S)+(1-\alpha)*(X-S-V))+$$

$$\left(N-\left[\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right]\right)*(\alpha*(X-S)+(1-\alpha)*(X-S-V-G)) \quad (1.0)$$

$$\mathrm{E}(R_{\alpha}) = \left[\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right]*(X-S-V+\alpha*V)+\left(N-\left[\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right]\right)*(X-S-V-G+\alpha*V+\alpha*G)$$

$$(1.1)$$

$$\mathrm{E}(R_{\alpha}) = N*(X-S-V-G+\alpha*V+\alpha*G)+\left[\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right]*(G\text{-}\alpha*G) \quad (1.2)$$

$$\mathrm{E}(R_{\beta}) = \left[\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right]*(\beta*(X+r-F)+(1-\beta)*(X+r-F-V))+$$

$$\left(N-\left[\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right]\right)*(\beta*(X+r-F)+(1-\beta)*(X+r-F-V-G)) \quad (2.0)$$

$$\mathrm{E}(R_{\beta}) = \left[\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right]*(X+r-F-V+\beta*V)+\left(N-\left[\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right]\right)*(X+r-F-V-G+\beta*V+\beta*G)$$

$$(2.1)$$

$$\mathrm{E}(R_{\beta}) = N*(X+r-F-V-G+\beta*V+\beta*G)+\left[\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right]*(G-\beta*G) \quad (2.2)$$

To determine the conditions for when a decision making company would prefer or be indifferent to investing in enhanced cyber security relative to enhanced revenue, equation (1.2) must be greater than or equal to equation (2.2). For the case where it is assumed that the decision making company is considering investing equivalent amounts in enhanced cyber security and enhanced revenue, F=S. For this situation, Equation (3.0) can then be developed in a manner that relates the ratio of improvement in cyber security defense α/β to the right side of the equation as the basis for a decision.

$$N*(X-S-V-G+\alpha*V+\alpha*G)+\left\lceil\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right\rceil*(G\text{-}\alpha*G)\geq$$

$$N*(X+r-F-V-G+\beta*V+\beta*G)+\left\lceil\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right\rceil*(G-\beta*G)$$

$$G*\left(\left\lceil\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right\rceil*(1-\alpha)-\left\lceil\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right\rceil*(1-\beta)\right)\geq N*r+N*\beta*V-N*\alpha*V+N*\beta*G-N*\alpha*G$$

$$N*(G+V)*(\alpha-\beta)\geq N*r-G*\left(\left\lceil\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right\rceil*(1-\alpha)-\left\lceil\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right\rceil*(1-\beta)\right)$$

$$(\alpha-\beta)\geq\frac{N*r-G*\left(\left\lceil\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right\rceil*(1-\alpha)-\left\lceil\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right\rceil*(1-\beta)\right)}{N*(G+V)}$$

$$\frac{\alpha}{\beta}\geq\frac{N*r-G*\left(\left\lceil\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right\rceil*(1-\alpha)-\left\lceil\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right\rceil*(1-\beta)\right)}{N*(G+V)*\beta}+1 \qquad (3.0)$$

$$G\geq\frac{\left[N*\left(O+S+(\beta-\alpha)*V\right)\right]}{\left[\left(N-\left\lceil\frac{\gamma}{(1-\alpha)*(\mathrm{C}_{\alpha}+1)+(1-\beta)*\mathrm{C}_{\beta}}\right\rceil\right)*(\alpha-1)-\left(N-\left\lceil\frac{\gamma}{(1-\alpha)*\mathrm{C}_{\alpha}+(1-\beta)*(\mathrm{C}_{\beta}+1)}\right\rceil\right)*(\beta-1)\right]} \qquad (4.0)$$

Equation (3.0) can be rearranged to calculate the size of the government fine G that will shift the decision making company’s choice from enhanced revenue to enhanced cyber security, as shown in Equation (4.0).

The above normative model was used as a basis for the live collaborative simulation exercise used at the I3P workshop in June 2006.

3. Collaborative Exercise Introduction

The cyber security investment exercise was conducted in two sessions at the second I3P Process Control Systems (PCS) Security Workshop that took place on June 8th, 2006. The purpose of each session was to expose participants to work in decision making with respect to investment in cyber security, including consideration of the possibilities for government regulations if the industry is seen to be too vulnerable. Both sessions were of the same format. The first session had 9 participants and the second session had 11 participants. Each session was about an hour and fifteen minutes long. The sessions began with an orientation presentation. Then participants were asked to take part in a five stage decision exercise. In the exercise, participants took on the role of a decision maker in a hypothetical company in the Oil & Gas Pipeline and Refining industry. During the exercise participants had two choices to make during each stage; whether to invest in enhanced PCS cyber security or to invest in enhanced automation. In this exercise enhanced automation was the basis for enhanced revenue. If participants suffered an attack the financial consequences of the attack rippled throughout the entire United States economy, with economic impacts estimated through use of a macro-economic input-output model, called the Inoperability Input-Output Model (IIM) [Haimes and Jiang 2001; Santos and Haimes 2004]. The IIM makes use of the Bureau of Economic Analysis (BEA) published production and consumption data on 500 sectors of the US economy [BEA 1998]. The BEA collects this data to support the use of Wassily Leontif's input-output model of the economy which provides a framework [Peterson 1991] that is used to study equilibrium and the interconnectedness of the various sectors of the economy [Horton 1995]. The consequences of an attack that a participant suffers are called the direct costs of the attack. The corresponding ripple costs that the entire US economy suffers are the indirect costs of the attack. For the exercise these indirect costs are the impetus for the government to intervene. The rationale behind an intervention is that the industries suffering the indirect costs have no control over the ripple situation, and are dependent on the oil & gas industry's investment in cyber security. In the exercise the government will wait until the indirect costs have reached a set threshold and then institute the use of fines to encourage more cyber security investment. The rules of the exercise in both sessions are shown below in the identical form that was presented to the participants. The exercises provided an opportunity to receive input and feedback from a number of experienced professionals who could collaborate in an organized setting that serves to direct continuing research using the collaborative exercise methodology.

3.1 PCS Security Game Instructions

The instructions provided to players of the exercise are the following.

The purpose of this computer game is to compare and evaluate PCS security investment decisions made by managers in the oil and gas industry. The results will be used as part of the design process of a more general decision support tool for managers.

The game provides a group of decision makers from competing companies with a set of parameters that surround their PCS security investment decisions related to the protection against PCS intrusion and disruption. The companies for which the decisions are being made are US-based, each have annual revenue of $10 billion. Based on history it is generally accepted that with the normal levels of investment in PCS security any individual company faces a likelihood of a successful attack occurring of about once during a five-year period.

The game runs for a five-year period with annual management decisions and corresponding annual financial simulation results corresponding to the players' investments in PCS security. Each player (decision maker), on an annual basis, has a choice between two alternative investments of the same dollar amount. The size of the investment is externally set and displayed at the start of the game (e.g., $1 million).

- Alternative 1 – Invest in enhanced PCS security to reduce the likelihood of PCS intrusion and disruption. The normal security budget for your company is $10 million per year. This added investment will reduce the likelihood of your company's PCS being disrupted. The financial consequences of successful PCS intrusion reduce the revenue of the company for the current year by an externally set amount of 1% (total of $100 mm). The impact of the Alternative 1 investment is to extend the anticipated period for a successful attack by a factor of approximately 2; i.e., from five years to ten years. Actual numbers will be provided with each decision.

- Alternative 2 – Invest in enhanced process efficiency through investment in automation with a guaranteed revenue return. This alternative will rely on the normal security budget and probability of attack. However, it will guarantee an increase in revenue. (i.e., $3 million increase in revenue in the year of the investment). Actual numbers will be provided with each decision.

For the simulated year in question, all players decide on which investment alternative to select. A simulation is then run to determine the outcome of PCS intrusions/disruption for each player. The simulation determines the PCS intrusion outcomes for each player's company by using the probabilities that are based on the PCS security investment decisions of the players. PCS security breaches are probabilistic, so that those who do not select the PCS security investment alternative will not necessarily experience a breach. Similarly, those who invest in the PCS security alternative may experience a breach. All that can be said is that those who invest in the PCS security alternative are less likely to experience a breach.

Federal and state regulators may intervene with industry regulation in the event that the consequences of PCS intrusions/disruptions propagate significantly outside of the oil and gas industry. An intervention with regulation will reduce the revenue of future companies that are successfully attacked by a $20 million fine in the year of the attack. Actual numbers will be provided at the time of each decision.

After all the players make their decisions and the simulation is completed, everyone sees information about any PCS intrusion event that occurs during that year (e.g., through press releases) and indirect impacts that government is monitoring. No player sees another player's decisions.

4. Exercise Organization

The implementation of the decision-making exercise occurred through the use of a website built with PHP which accessed a MySQL database hosted on a wireless LAN set up in the session room. The results of the exercise were shown with Microsoft Excel produced data charts at the completion of the five stages of the exercise. The exercise was based on the probabilistic decision analysis model previously discussed that compared expected financial outcomes for investments in either enhanced cyber security or enhanced revenue (automation in this form of the game), given the current level of government regulation. The analysis results in normative decision criterion. Participants of the session were not informed of possible decision criteria, but instead were free to make decisions on whatever basis they wanted to use. The exercise consisted of one decision scenario (treatment) requiring annual decision investments for each of five years (each year considered to be a stage). The treatment conducted was used to calibrate the investment biases of the players. This is accomplished by selecting a set of decision parameters that make the two alternate decisions of enhanced cyber security or enhanced automation equally attractive from an expected value viewpoint. However, the participants' beliefs of what investments the other participants are going to make causes the mathematically determined switch point, the point where the individual is indifferent between the enhanced cyber security and enhanced automation investment, to differ.

After an introductory presentation that prepared them for the exercise the participants were on their own to make their decisions. Using simple web pages, participants entered their decisions and criteria for their decisions, all of which were collected for analysis purposes. Participants then submitted their decisions to discover whether they suffered a successful cyber attack, and their revenue outcome for the year. The attacks for individual participants were determined via a random sampling based on a probabilistic calculation which depended on whether the participant invested additional money in cyber security, or not. The calculations were based on the assumptions provided as part of defining the game. Then participants waited until all the other participants had made their decisions and then they could advance to a screen showing the collective number of attacks across all participants for the stage and the resulting progress made toward breaching the government's threshold for regulation. When the participants were finished looking at the industry results they could move on to their next decision, and the process of making a decision and submitting it was repeated for five stages. At the end of the final stage the total five-year integrated revenue results for the players over the entire treatment was shown (Figure 1). A second graph was shown displaying the number of players that invested in each of the alternatives in each stage (Figure 2). The third graph that was shown indicated the number of players that had changed their decision zero, once, twice, three, or four times (Figure 3). The final graph shows the belief of the participants about the amount of participants investing in enhanced cyber security in each stage (Figure 4).

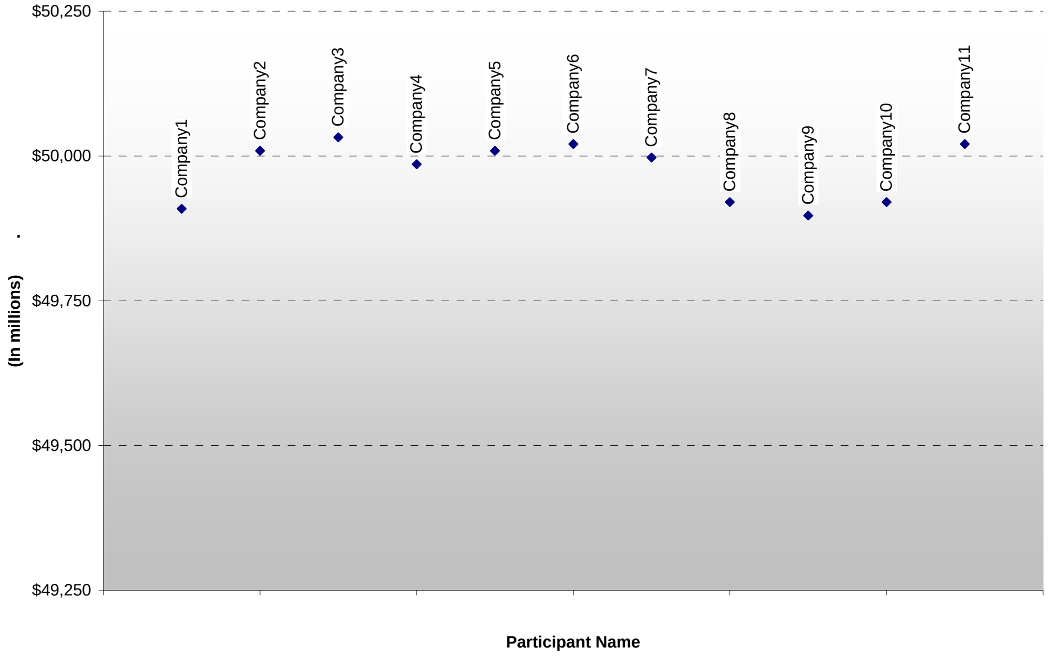

Figure 1. Cumulative Revenue for Each of the Eleven Companies.

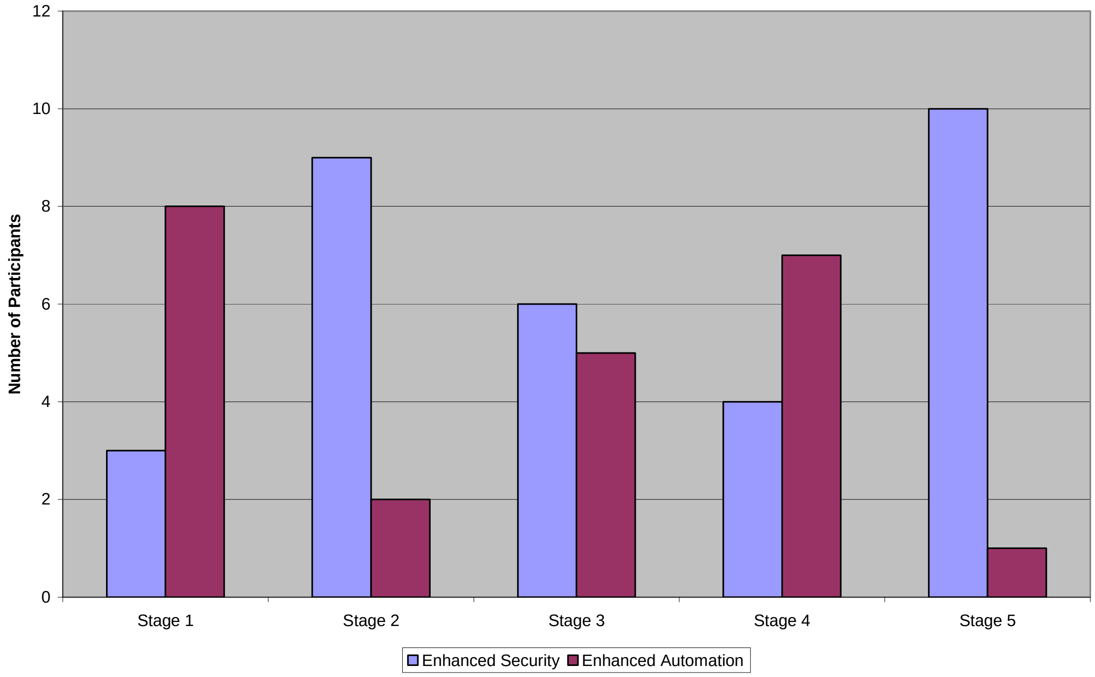

Figure 2. Number of Participants Selecting Enhanced Security and Enhanced Automation Each Stage.

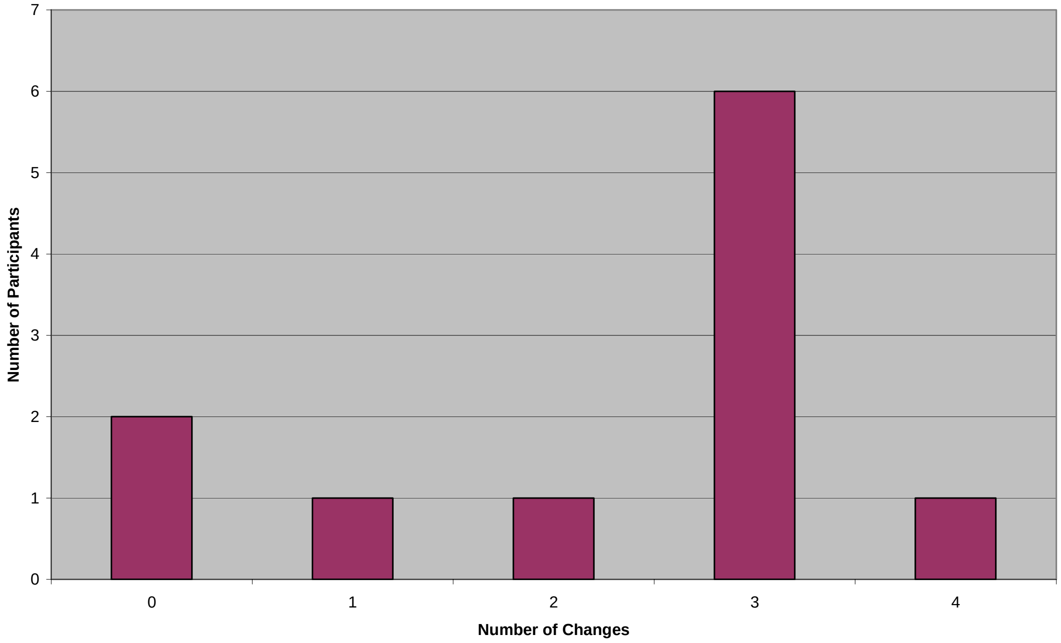

Figure 3. Number of Decision Changes Between Enhanced Security and Enhanced Automation.

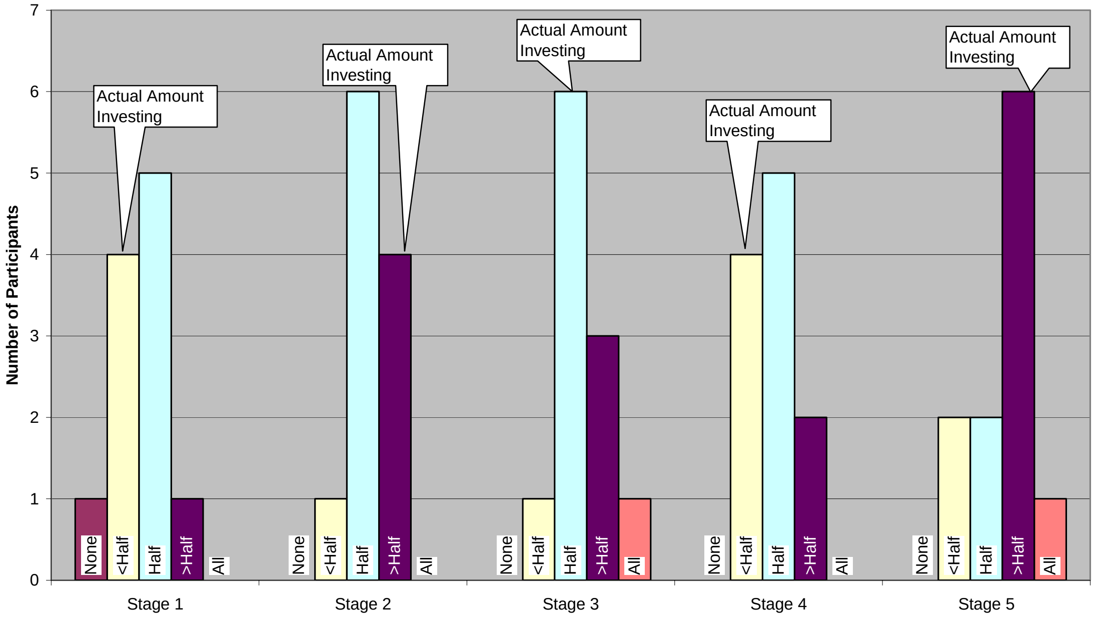

Figure 4. Participant Belief of the Number of other Participants Investing in Enhanced Security.

5. Discussion of Exercise Results

The two sessions had similar results. For brevity, and because the two sessions' results can not be combined into one result due to the randomness in the specific exercise situations, only one of the sessions' results is presented in this paper. If a company implemented this collaborative exercise, then the following results would be indicative of its possible results from the exercise. In this session there were 11 participants. This session’s parameters were set up to make a perfectly rational participant as indifferent as possible between enhanced cyber security and enhanced automation. Figure 2 shows the stage by stage breakdown of the number of participants investing in each option. In stage 1 the majority of participants are investing in enhanced automation and from Figure 4 the majority belief is that half will be investing in enhanced security. It seems that participants are not worried about the chance of government regulation and believe that the other companies in their industry will be making the enhanced security choice while they can then make the short term revenue enhancing choice of enhanced automation. The result of that stage was three successful attacks in the entire industry resulting in a situation where only one more successful attack is required to provoke the government to institute the use of fines. The next stage shows a large swing to the majority of participants investing in enhanced security. The threat of the government fines is now much greater than before and the participants may be inferring from the number of successful attacks that only a few companies had invested in enhanced security. In stages two through four, Figure 2 shows a gradual shift from enhanced cyber security to enhanced automation. This seems to result from no successful attacks in stage two or in stage three. However there was one successful attack in stage four and as a result in stage five, the first year the government fines would be in effect, all but one participant invested in enhanced cyber security. The result of this treatment was an overall 1.4 ratio of annual investments in enhanced security vice enhanced automation (Figure 2). At first, the participants' beliefs of their peers' behavior lagged behind the reality of the actual investments, but by the third stage were corrected. In the fourth stage it was less accurate, but in the fifth stage, when government intervention was to take effect from the accumulation of successful attacks in the first four stages, the beliefs and actions matched. The participants in the session received four successful cyber attacks and triggered government regulation to start in stage five. There were five stages with eleven participants for a total of fifty-five stages of play with the observed rate of attack of one in about fourteen years per company. Of the four attacks three resulted in participants switching from the automation investment to the security investment, and the other one did not result in the participant switching their decision.

5.1 Player Provided Reasons for Decisions

The provided reasons for investment decisions varied widely among the players in the exercise. The reasons fall into two categories, automation enhancement investment rationale and cyber security investment rational. The cyber security investment rationale includes the following reasons.

1. The investment is small and the risk is big.
2. Fear and concern of the consequences of an attack.
3. The risk of the government regulating and risk of a government fine.
4. Negative publicity if invest in automation and are attacked successfully.
5. Experienced a successful attack
6. Will sacrifice increased revenue to help the industry.
7. Insurance

The provided reasons for automation enhancement investments are listed below.

1. Security does not generate revenue.
2. The return on investment for automation is higher.
3. Let others in the industry invest in security.
4. Get gains in revenue now and use gains for security later.
5. No history of attacks.
6. No threat of government fine.
7. Gambling, low chance of attack.
8. Automation is a higher priority

5.2 Expected Value of Winners and Losers

We defined the “winner” as the participant that accumulated the most annual revenue over the five stages. The loser had the lowest cumulative revenue over the five stages. There was one winner for each session and one loser for each session. We recognize that this biases participants towards investing in revenue, but judged that this mirrors reality. For session 2, the winner invested in automation in every stage but the last, and the loser invested in security in every stage but one and was hit once, which occurred when they had invested in security. The expected value of the winner’s decisions was $49,941.9 million; the winner actually received $50,031.9 million and never was successfully attacked. Based on the exercise’s assumptions, the winner had a 36.8% chance of not being hit during the five stages. The expected value of the loser’s decisions was $49,937.1 million; the loser actually received $49,897.1million. The loser had a 5.8% chance of being hit the one time, and a 52.4% chance of not being hit. These numbers clearly indicate that chance plays a major role in actual outcomes, resulting in confusion for decision-makers who are not sufficiently analytical.

5.3 Recommendations from Participants

The two sessions resulted in many recommendations from the participants. The attendees reported great interest in the opportunity-cost oriented framework and the combination of government action into a framework for decision analysis. Participants were interested in adding an array of options and costs to the cyber security investment choices which would have a corresponding array of attacks and attack levels resulting in a variety of consequences. Another suggestion was to have companies' investment decisions impact their relative desirability compared to the other companies in the industry as a target for attacks. Another participant suggested that the behaviors of individuals in the exercise be typified into several classes (risk averse, risk neutral, etc.) and that the classes be aggregated in a computer simulation to explore the broader industry behaviors.

6. Conclusions

These sessions were examples of a collaborative exercise that a company could deploy and draw inferences about its self and its cyber security operating environment if not statistically justified conclusions. Both sessions were opportunities to obtain inputs from seasoned experts who are regularly involved in cyber security investment decision-making, and thereby be an example of how a company could bring together its decision-making experts. The first conclusion that can be drawn by a company, consistent with prior experience from the first game [Horowitz 2006], is that the participants had a bias toward enhancing revenue through automation investment, until the potential for regulation became significant, and then their investment choice shifted to enhanced cyber security. This bias towards investment in enhanced automation was highest when the possibility for government regulation was lowest and when there had not been a recent successful attack. A company could then see if this conclusion matched its policy or was in its best interest from a wide variety of perspectives. A second conclusion that can be drawn by a company is that initially participants depended on their neighbors in the industry to invest in cyber security while they reaped the financial benefits of enhanced automation; however as time past the participants learned that they could not depend on this and their beliefs more closely mirrored the actual decisions of their peers. From the results it seems clear that participants are impacted by their perception of the action the government and their peers in industry will take. For example in the session results after the first stage where 3 successful attacks occurred and the industry was one attack away from regulation, nearly all participants switched from investing in enhanced automation to investing in enhanced cyber security. A company can learn from this that knowledge of its industries' behaviors and practices tends to have a large effect on their decision-making, then the company can decide what action to take in regards to that finding.

In addition to the session results the analytical model can serve the company as illustrated here. In Figure 5, if we look at a similar situation to the one in the workshop, the importance of the belief of the decision maker in their peer's investment can be shown. In Figure 5 the entire industry is composed of eleven companies, the two investment options both cost \$2.75mm, the enhanced security investment shifts the likelihood of successful defense from 0.8 to 0.9, the enhanced automation returns \$11mm, and the government threshold for imposing the use of fines is four successful attacks to the industry. Considering a five year investment period and consistent investment in either enhanced security or enhanced revenue, Figure 5 shows that if the decision maker believed 4 or less of its peers were investing in enhanced security, the mathematically derived choice would be for it to choose to invest in enhanced cyber security. However, if it believed 5 or more of its peers were investing in enhanced cyber security it should then choose to invest in enhanced revenue. Lastly if the entire industry is investing in enhanced cyber security, so should the decision maker's company. This shift occurs because of the assumption in the model that the government will not institute the fine until the beginning of the next year after the threshold for significant successful attacks has been reached. When all the other companies are investing in enhanced cyber security the expected year the government threshold being met shifts from the third year to the fourth year and thus the enhanced cyber security investment has higher expected revenue because of a smaller expected period under the fine regime. This graph and the experience from the workshop seem to highlight that a decision making company could benefit from an exchange of information with its peers and the government, which may result in other industries subject to ripple effects of successful attacks being better protected and thus the company conducting the exercise having an alternative to government regulation being enacted upon them.

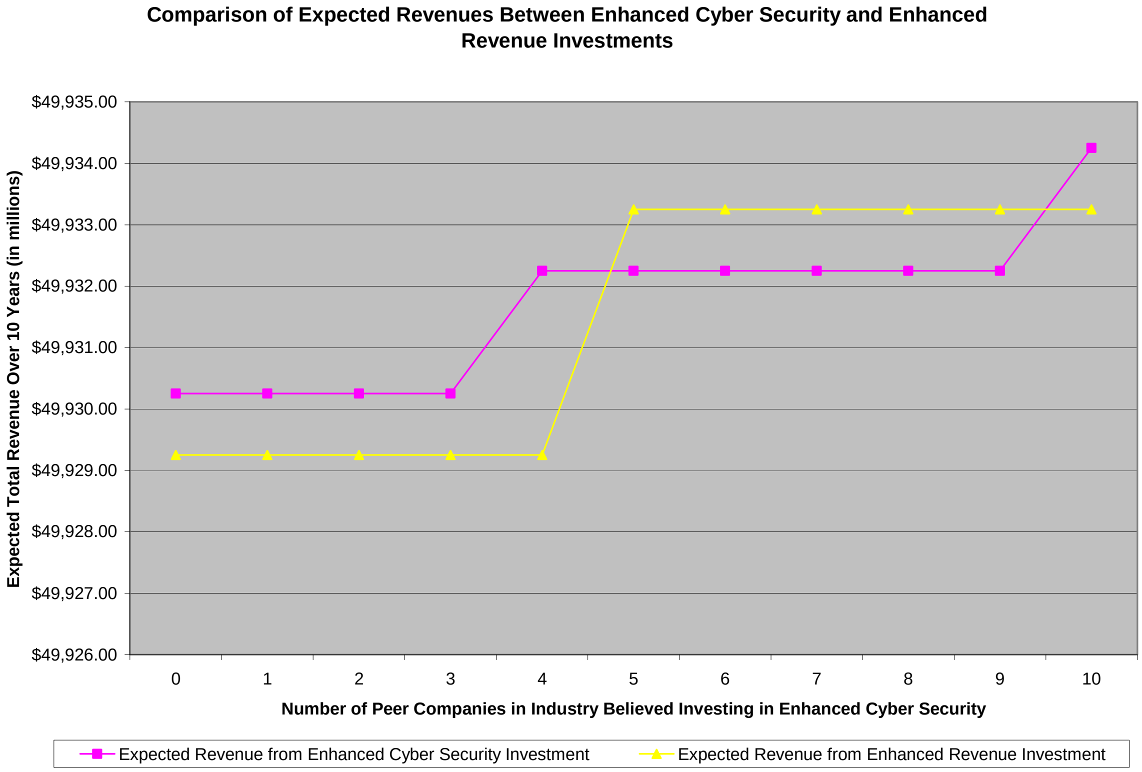

Figure 5. Expected total revenues over ten years compared to the number of peer companies believed to be investing in enhanced cyber security.

Further expansion of this collaborative exercise and its model will involve incorporating cost and price details for cyber security investments, integrating the realm of physical security with cyber security, and developing a tool for the use of decision makers that will facilitate intra company information sharing and thereby incorporate the breadth of opportunity costs and the breadth of attacks and consequences in the investment of cyber security. In addition the current methods for deciding on cyber security investment will be integrated with the structured approach started here.

Acknowledgments

The authors appreciate the support of the Institute for Information Infrastructure Protection (I3P) and the comments of Prof. Yacov Y. Haimes and Prof. Joost Santos of the University of Virginia.

7. References

Bureau of Economic Analysis, U.S. Department of Commerce, 1998, *Benchmark Input-Output Accounts of the United States*, 1992, U.S. Government Printing Office, Washington, DC.

Baumer, D. L., Earp, J. B., & Poindexter, J. C. (2005). "Quantifying Privacy Choices with Experimental Economics". *Workshop on the Economics of Infomation Security 2005*.

Garcia, Alfredo and Horowitz, Barry, "The Potential for Underinvestment in Internet Security: Implications for Regulatory Policy". *Journal of Regulatory Economics*, Forthcoming Available at http://www.thei3p.org/research/economics/regeconpaper0306.pdf.

Gordon, Lawrence, and Martin Loeb. Managing cybersecurity resources: a cost-benefit analysis. New York, New York: McGraw-Hill, 2006.

Haimes, Y.Y and P. Jiang, 2001, Leontif-based model of risk in complex interconnected infrastructures, *Journal of Infrastructure Systems* 7(1), 1-12

Peterson W. Advances in input-output analysis : Technology, planning, and development. Oxford University Press, New York, 1991.

Horowitz, Barry. "I3P Economics Project Report." I3P Reseach Report no. 4. Jan 2006 http://www.thei3p.org/about/researchreport5.pdf.

Horton, G.A. 1995. Input-Output Models and Economic Impact Analysis: An Overview of Methodology, Economic Impact Multipliers, and Comparison to Econometric Forecast Models, with a Glossary of Terminology and Selected References. Business and Economics Research Associates, Reno, Nevada.

"I3P Agenda." Institute for Information Infrastructure Protection. 2003. I3P. 20 Jul 2006 http://www.thei3p.org/about/2003_Cyber_Security_RD_Agenda.pdf.

Kagel, John H. and Dan Levin, "The Winner's Curse and Public Information in Common Value Auctions," American Economic Review, December 1986, 76, 894-920.

Lynch, Mike, and Gillespie, Nick. "The experimental economist: Nobel laureate Vernon Smith takes markets places they've never been before - Vernon L. Smith - Interview." Reason Dec 2002 http://www.findarticles.com/p/articles/mi_m1568/is_7_34/ai_94775374.

Roth, Alvin E., "Bargaining Phenomena and Bargaining Theory," in A. E. Roth (ed.) Laboratory Experimentation in Economics. Cambridge University Press, 1987, 14-41.

Santos, J.R. and Y.Y. Haimes, 2004. Modeling the Demand Reduction Input-Output Inoperability Due to Terrorism of Interconnected Infrastructures. *Risk Analysis* **24**(6): 1437-1451.

Smith, Vernon. "Experimental Studies of Discrimination versus Competition in Sealed-Bid Auction Markets." Journal of Business January, 1967.

Smith, Vernon. "Experimental Economics: Induced Value Theory." American Economic Review May, 1976

Smith, Vernon, J. C. Cox, and B. Roberson. "Theory and Behavior of Single Object Auctions." Research in Experimental Economics Vol. 2 (1982).

Smith, Vernon, J. Ketcham, and A. W. Williams. "A Comparison of Posted-Offer and Double-Auction Pricing Institutions." Review of Economic Studies October, 1984.

Smith, Vernon L., Papers in Experimental Economics. New York: Cambridge University Press, 1991.

Smith, V.L., 1994, "Economics in the Laboratory", Journal of Economic Perspectives 8(1), 113-131.