Costs to the U.S. Economy of Information Infrastructure Failures: Estimates from Field Studies and Economic Data

Scott Dynes (sdynes@dartmouth.edu)[1]
Eva Andrijcic (ea2r@virginia.edu)[2]
M. Eric Johnson (m.eric.johnson@dartmouth.edu)[1]

[1]Center for Digital Strategies, Tuck School of Business at Dartmouth College
[2]School of Engineering, University of Virginia

Forthcoming in *Proceedings of the Fifth Workshop on the Economics of Information Security*, Cambridge University.

Abstract

The U.S. economy increasingly relies on the internet as a critical infrastructure for enabling business processes from product design, engineering and supply chain management upstream through sale and fulfillment of product downstream. As a result, there is much interest in the vulnerability of the U.S. economy to targeted and large-scale disruptions of the information infrastructure. Here we present results from field and modeling studies that estimate the macro-economic costs of a targeted internet outage in three economic sectors: automobile manufacturing, electrical device manufacturing, and oil refining. Firm-level estimates of productivity loss from outages of varying duration were used as input into a Leontief-based input-output model of the U.S. economy to estimate the total economic impact of outages from 3 day and 10 day's duration.

Introduction

The increasing reliance of the U.S. economy on the information infrastructure has raised questions regarding the security and robustness of the critical information infrastructure at all levels of the economy, ranging from individuals in small firms facing very practical concerns to national figures facing equally pressing policy issues [Joh2005]. Until recently, these individuals have had to rely mainly on speculation for guidance as empirical studies of the economic risks faced by individual firms and larger economic entities were unavailable [Cas04]. This lack of data concerning these issues was the original impetus for the studies presented in this paper.

Previously, we presented results relating to the economic drivers and other incentives for individual firms to invest in information security including the risks firms are exposed to through the use of the information infrastructure to manage their extended enterprises [Dyn05]. In the current work we examine the risks at a macro-economic scale, looking at the possible economic losses to the particular regions of the U.S. economy arising from one particular type of disruption to the information infrastructure: a cyber attack on specific firm that results in an internet outage or control system failure (e.g., a denial of

service attack or a virus/worm that results in disconnection from the internet). As part of our field research, we explored the impact that an outage would have on the subject firm's ability to produce and ship product. Casting this as a change in the production of a standardized product enabled the utilization of an Input-Output model of the U.S. economy to estimate the total impact such an outage would have on the U.S. economy. This total impact includes not only the direct economic consequences experienced by the affected sector, but also the indirect economic consequences brought about by disruption for suppliers, as well as the firm's customers. The economic impacts of these 'ripple' effects can be greater than the direct economic impact experienced by the affected sector.

To our knowledge this is the first work that provides an empirically based estimate of the macro-economic consequences of disruptions to the critical information infrastructure. The knowledge resulting from this work should serve multiple important efforts:

- To reaffirm and continue to raise awareness that firms face and need to actively manage cyber-risk not just internally but also in their extended enterprise
- To allow interested parties to gauge and compare the risk due to cyber-events with other (cyber) risks at a macro-economic scale
- To make rational policy decisions regarding appropriate levels of information security for the critical information infrastructure

Methods

Field Study

The field study consisted of a set of interviews with security and supply chain executives and managers at each participating firm. The interviews were designed to elicit the knowledge and beliefs of the interviewed individuals; security audits of the interviewed firms were not a part of this study. At the start of each interview, we made it clear to the interviewees that the interview was anonymous; during the interview every effort was made to build a high degree of trust with the interviewee. Interviewees at 'host' firms included top-level managers as well as lower-level individual contributors, of both information security and supply chain management. Interviews at supplier firms were usually limited to a single manager of information security and a manager of supply chain; at smaller firms these might be the same individual. By asking the same questions of different interviewees in the same organization, we were able to look at the internal consistency of information provided in interviews [Yin94]. Additionally, this approach exposed both strategic as well as tactical issues regarding information security and its role in maintaining the supply chain. The insights obtained from these diverse viewpoints are presented in the following sections.

Questions asked during the interviews were centered on the identification and management of information security risks, and of particular interest for this work, business continuity risk firms faced as a result of using select technologies to manage their supply chains. These were open-ended questions eliciting the impact that an internet outage of certain durations would have on the ability of the firm to continue to produce and ship product. The interviewee would be queried about durations of 1 second, 1

minute, 1 hour, etc.; the conversation would quickly evolve to a discussion about the conditions, durations and impacts of consequential outages. The results of these conversations were written up and serve as the basis for the determinations presented in this paper.

Macro-Economic Model

In order to estimate the macro-economic consequences on the U.S. economy due to internet outages to the three aforementioned sectors, the University of Virginia's Inoperability Input-Output Model (IIM) was employed. The IIM is based on the economic theory of Wassily Leontief who in 1973 won the Nobel Prize for the creation of the Input-Output Model for the US economy, which describes economic interdependencies between various sectors of the US economy [Leo1966]. The general formulation of Leontief's Input-Output Model can be described by the following equation,

$$x = Ax + c \Leftrightarrow x_i = \left\{ \sum_j a_{ij} x_j + c_i \right\} \forall i,$$

where x_i represents the total production output of industry i, a_{ij} represents the ratio of the input of industry i to industry j, with respect to the total production output of industry j (given n industries a_{ij} gives us the distribution of inputs contributed by different industries $i = 1,2,\ldots,n$ to the total output of industry j), and c_i represents the final demand for industry i (i.e. it tells us what portion of industry i's total output is used for final consumption by end-user).

While the mathematical formulation of the IIM is very similar to the general Leontief model, its interpretation is fundamentally different and 'supply' and 'demand' concepts of the Leontief model take on a different meaning. The input to the IIM now becomes a vector of perturbations to sectors, which can range from willful attacks, accidents, to natural disasters, and the output becomes a vector of inoperabilities of different sectors, resulting from the introduced perturbations. While the formulation of the model is not complex, it requires the introduction of some new concepts, so it will not be presented in this paper. However, the interested reader is referred to [Hai04] and [San04] for a complete overview of the theory and the model.

For the purpose of this paper it is important to note that the IIM is a linear, static, "computer-based model for analysis of how perturbations (e.g., demand shocks due to willful attacks, accidental events, or natural disasters) to selected groups of sectors can impose direct and indirect impacts on the operation of other sectors, due to inherent interdependencies" [San04]. In essence it is a "general analytic framework to quantify and address risks from the intra- and interconnectedness of large-scale complex infrastructures" [Hai01] that utilizes data on economic interdependencies from the US Department of Commerce [DoC98]. The major attribute of the model is its ability to account for production and services interactions and dependencies of all of the significant industries in the U.S., permitting one to determine the integrated economic impact of a specified reduction in productivity of any subset of industries, or a reduction in demand for any subset's products.

A major benefit of the IIM is the fact that it is supported by a large, ongoing national data collection effort, namely the Bureau of Economic Analysis' (BEA) database which contains a series of input-output tables depicting the production and consumption of commodities (i.e., goods and services) of various sectors in the US economy [DoC98]. "The BEA data is a record of the physical exchange of commodities between various interconnected industrial sectors of the economy that have been scaled by producers' prices into one common unit of dollars" [Hai04]. The detailed national tables are composed of hundreds of industries, organized according to the Standard Industry Classification (SIC) or more recently, the North American Industry Classification System (NAICS) codes. These tables are used to produce the Leontief technical coefficient matrix which depicts the economic interdependencies between nearly 500 U.S. sectors. The IIM is also supported by a set of regional data (RIMS II) maintained by the Bureau of Economic Analysis, Regional Economic Analysis Division [DoC97]. RIMS II releases regional data for various regions of the US. Thus, with the availability of national and regional data, analyses can be conducted on the regional level, which provides a more focused and thus more accurate analysis of interdependencies for particular regions of interest in the US. It is important to note that the IIM is based on the assumption that the level of economic dependency between various sectors is the same as the level of physical dependency between those sectors. In other words, the model assumes that "two companies with a large amount of economic interaction will have a similarly large amount of physical interdependency" [Hai04]. According to Haimes et al, "however crude this assumption may be, it is founded on BEA data that reflects real physical interactions between economic sectors. These are translated into dollar units by multiplying interactions of physical quantities by producers' prices. In turn, these prices indicate how a sector values the physical interdependencies" [Hai04].

The IIM considers a system consisting of 59 critical interconnected sectors, with the output being their inoperability that can be triggered by one or more failures due to willful attacks, accidents, or natural disasters. "Inoperability connotes degradation in the system's functionality (expressed as a percentage relative to the intended state of the system)" [Hai04], and it is assumed to be a continuous variable evaluated between 0 and 1, with 0 corresponding to a flawlessly operable system state and 1 corresponding to the system being completely inoperable. The inoperability caused to a particular sector can cause an inoperability ripple to the other sectors in the interconnected economic system. The IIM captures this ripple effect and provides a system for ranking the most vulnerable sectors. Combining the main output of the IIM, namely the inoperability of sectors caused by a particular perturbation, with the BEA data on economic interactions between the sectors, the IIM can compute the direct and indirect economic losses to all of the sectors of the U.S. economy.

Firm-level disruptions were reflected into the macro economic model through the development of productivity erosion curves. From the interviews, we constructed graphs that estimated the production as a function of time for an internet outage. A notional example is shown in Figure 1, which shows an example of a curve of decreasing productivity after an internet interruption, followed by a quick recovery to 100% productivity.

One might wonder whether it might be possible for the production to recover beyond 100%. Such could be the case if, for example, Amazon.com were to be down for an hour – the sales that occurred during that hour would not be lost, but some percentage of them would occur shortly after Amazon.com reappeared. The proper way to look at the productivity is as a percentage not of possible productivity, but of normal activity. In the example given, the productivity would increase above 100%. In the models derived here, even though there are good arguments for it increasing above 100% in certain cases the productivity is held to a maximum of 100% because of an inability to estimate the time course of the supra-100% effective productivity.

Field Study Results

For this work we conducted interviews in three business sectors: automobile component manufacturing, electrical component manufacturing, and oil refining. In the case of the automobile and electrical manufacturers, the host firm was a Fortune 500 manufacturing firm with plants and sales worldwide. The participating oil refinery was a regional, specialty refiner.

Manufacturing Sector

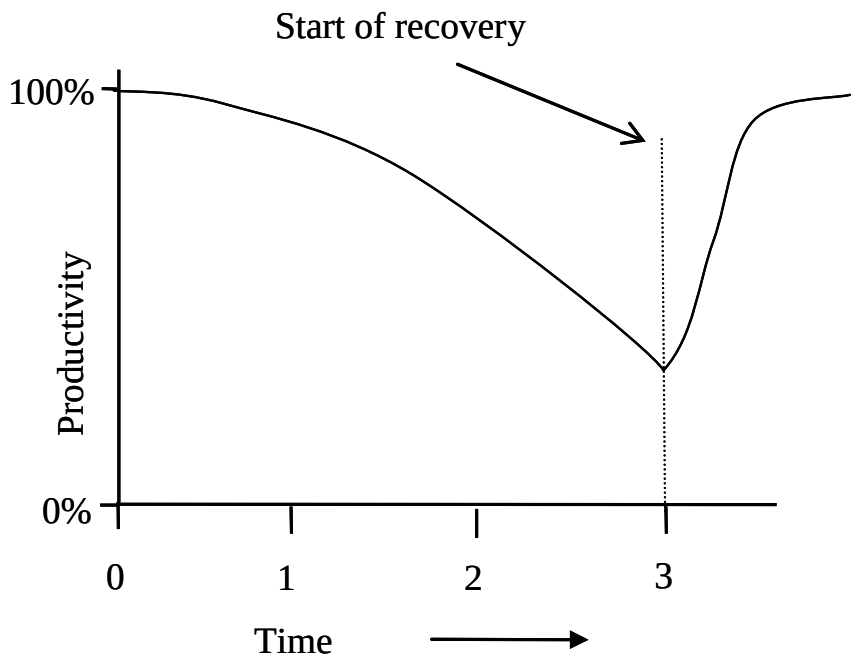

Figure 1. A notional diagram depicting how a firm's productivity might change following an internet failure. In this example the failure commences at time=0 and lasts until time=3; productivity decreases until the outage ends and immediately starts to recover to 100%. Productivity recovery was limited to 100% in this work.

For the manufacturing sector, a series of interviews were conducted with security, information and supply chain executives and managers at both the headquarters level, and where applicable at an individual business unit (BU) level; interviews were also conducted at suppliers. The great majority of interviews were conducted in person with one or two researchers, and one to four interviewees. Interviews lasted from 30 minutes to 2 hours; the remainder were phone interviews. In all, 29 individuals were interviewed Table 1 gives some particulars about the Host and the suppliers.

The host was a tier 1 supplier to the automobile industry, meaning that its products went directly in to the finished product, rather than a sub-assembly. As such, its customers were the automobile manufacturers (Ford, GM, etc.), which have rigorous requirements regarding electronic notification of dispatch of product shipments, access of design documents and other functions. Fulfilling these requirements dictated heavy reliance on the internet. Additionally, the auto and electrical BUs of the host organization were urging their entire supply chain to utilize either the host's web-based supply-chain applications or EDI (which can be thought of as email formatted in standardized ways) to manage their business with the host. This was putting pressure on the host's suppliers to become reliant on the internet for this customer. In summary, the host was very dependent on the information infrastructure to communicate with its customers, and was working to move its supply chain management functions (communications with its suppliers) to be internet-based as well.

The host's suppliers were much more variable with respect to their dependency on the internet. As detailed in a previous work [Dyn05], most communicated with their suppliers using phone and fax, and were much less susceptible to interruptions of their supply chain due to internet outages. These tier-two suppliers were also less dependent on the internet for communication with their customers as well. As alluded to in Table 2 below, for these suppliers the major impact of an internet event would not be supply chain or production disruptions, but in customer service via email.

Table 2 is a high-level summary of the knowledge collected during the interviews that enabled the generation of specific estimates of the impact internet outages would have on the ability of firms to produce and deliver product ('firm productivity'). To show this process in detail, consider the electrical parts manufacturer. For the shortest duration outages (an afternoon is listed in table 2), we found no evidence that there would be any impact of the firm's ability to manufacture parts. If the outage lasted longer, there would start to be effects. This manufacturer produced electrical products that range from small, very high volume items (many thousands produced a day) such as fuses and switches to large, very low volume (a few a month) items such as power-station transformers. Because of the long lead time required for parts that make up these large items, they must be ordered days or weeks in advance. Thus, if the internet were unavailable the day that an order needed to be placed for one of these long lead-time parts, there would be some disruption in the supply chain for these low-volume items.

	Product	Number of locations	Annual Revenues	Subsidiary?
Host	Conglomerate	many	billions	No
Supplier A	Metal	many	billions	Yes
Supplier B	Logistics Services	many	100 millions	Yes
Supplier C	Metal parts	many	100 millions	Yes
Supplier D	Metal finishing	few	10 millions	No
Supplier E	Metal parts	few	10 millions	No
Supplier F	Printing/Design	few	10 millions	Yes
Supplier G	Metal parts	one	millions	Yes

Table 1: Properties of Interviewed Manufacturing Firms.

It is unlikely this would mean that the order would not be placed, and that the delivery date for the entire transformer would slip. In place of an internet-mediated order, the order could be placed by phone or fax: since these were low-volume items, phone or fax ordering is entirely feasible. Prolonged internet outages for low-volume items would likely result only in customer-service disruptions, as email would be unavailable.

It is also the case that for many large items the largest suppliers of constituent parts were other plants of this manufacturer. To assure connectivity among these internal plants, the manufacturer had invested in an internal network that is separate from the internet. For these reasons, we estimated the impact of an internet outage on the production and delivery of low-volume items would be negligible.

There remains the case of the high-volume goods that this manufacturer also produces. Effectively supporting a manufacturing operation that produces many thousands of items a day involves frequent electronic communication of stock on hand and forecasting of demand. For example, a forecast to a supplier might reiterate the firm order for this week's supplies, a very good estimate for what will be needed the following two weeks, and a progressively ill-defined estimate of what will be needed up to ten weeks out. This forecasting and the history between the suppliers and the manufacturer result in what one interviewee termed a "supply chain learned behavior". This interviewee thought that if

internet down for:	An afternoon	1 day	3 days	A week
Electrical part manuf.	No impact	Low volume plants: supply-side pain	Hi volume plants OK	Hi volume plants: shipping issues
Automobile part manuf.	ASN disruptions - impacts customer	Stock available for production	Customers would see slack	Unable to produce all items
Supplier A	No impact	No impact on supply side; "big deal" on customer side, would use phone, fax, expect no loss of business		
Supplier B	[confident there would be no impact on supply or delivery of products]			
Supplier C	ASN disruptions	Customer service disruptions; no production disruption		
Supplier D	No impact	Fax ASNs, phone/fax suppliers, no production disruption		
Supplier E	No impact	No impact	No impact	No impact
Supplier F	No impact	No impact	No impact	No impact
Supplier G	No impact	No impact	No impact	No impact

Table 2. Estimated impact on a firm's ability to produce and ship product resulting from internet outages of various durations.

the internet were to fail, the suppliers would still deliver the needed supplies without any prompting from the manufacturer due to this 'learned behavior'. He estimated that the first noticeable supply chain event would occur about three days into an outage, when the suppliers would start calling the manufacturer regarding future forecasting and shipment information.

The major impact of an internet outage on the high-volume plants was likely not a supply-chain issue, but an order fulfillment issue depending on the number and predictability of orders for the many thousands of devices being built. If there were a handful of large customers for the product, it was likely that they also have forecasting and a 'learned supply chain behavior' in place and shipments would be packed and shipped as expected. At the opposite end of the spectrum was the case where there were hundreds of customers whose orders were delivered electronically and whose future orders were unknown. For such cases, our research led us to conclude that it is unlikely that the manufacturer could handle the potentially hundreds of faxes and/or phone calls for one or many products in various volumes, or to arrange shipment for product. Production and packaging would continue, but except for large customers these packages would accumulate on the shipping dock as the fulfillment system failed for smaller, ad-hoc orders.

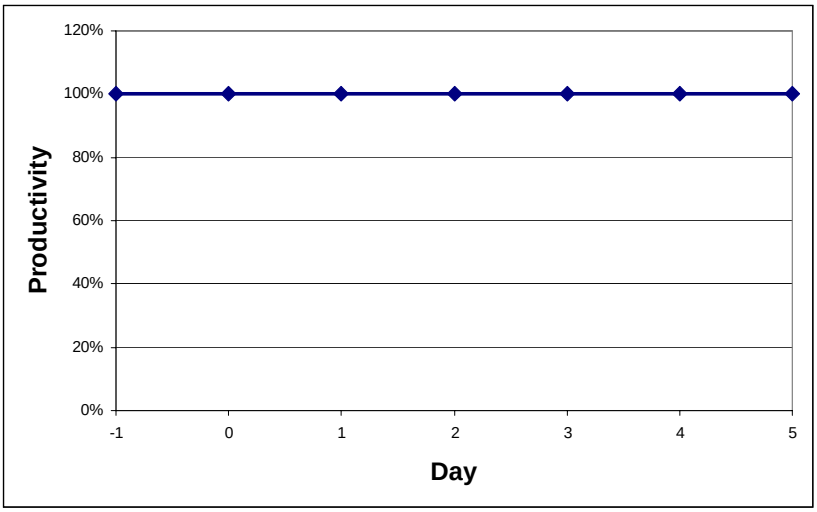

Figure 2. The impact of an internet outage of 3 day's duration on the production of an electrical goods manufacturer. For the reasons discussed in the text, it is reasonable to assume that an outage of this duration would have no effect on production.

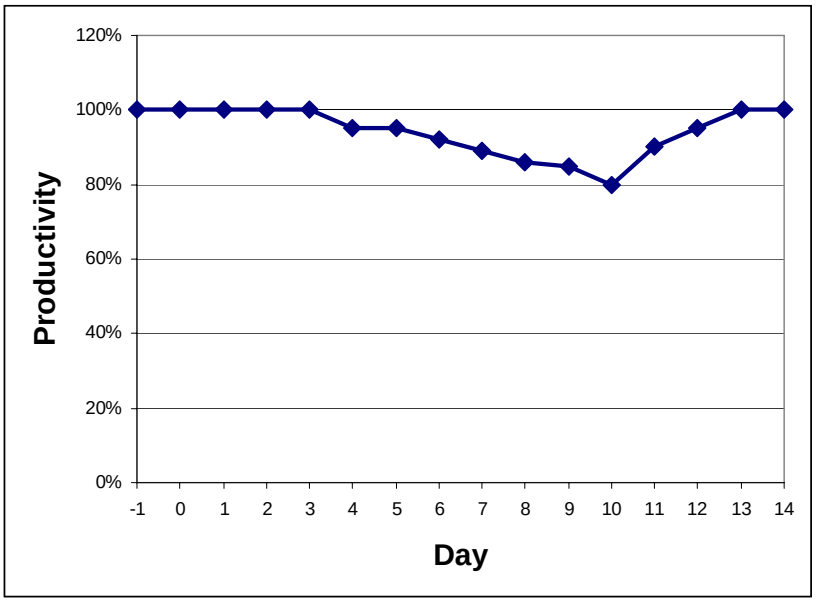

Figure 3. The impact of an internet outage of 10 day's duration on the production of an electrical goods manufacturer. As discussed in the text, it is reasonable to assume that an outage of this duration would have a small effect on production.

For these plants we applied a variation of the 80/20 rule, where 80% of production at high-volume plants would be destined for a few large customers – a hypothetical example would be producing wall switches for Home Depot or Loews. We make the assumption that these large customers would not experience a disruption, and that some percentage of smaller customer orders would also be fulfilled. Based on all of our analysis, our estimates of the impact on productivity are shown in Figure 2 for an outage of 3 days, and Figure 3 for an outage of 10 days.

The quick recovery of 'productivity' to 100% shown in Figure 3 is based on the fact that there is not a production disruption, but a fulfillment disruption, and that the resumption of electronic order and shipping fulfillment will quickly deal with the backlog. It is arguable that in this case the 'productivity' might increase beyond 100% as the volume of shipments is highly likely to be much greater than typical immediately following resumption of internet service. As described in the Methods section, we limit the maximum to 100%.

A similar approach is used to estimate the impact an internet outage would have on the productivity of an auto parts manufacturer. The automobile parts manufacturer adopted a different approach to supply chain management than the electrical parts manufacturer. Where the electrical supply chain was expected to display certain inertia and continue to deliver parts to high-volume plants, the automobile manufacturer in some important instances adopted a more command and control approach to supply chain management. Many of the sub-assemblies that the manufacturer uses were produced by local suppliers. While the auto parts manufacturer shares forecasting information with these suppliers, the manufacturer ran its own fleet of trucks to pick up these sub-assemblies in order to keep costs down, and to maximize efficiency the manufacturer was very explicit about how items should be packed, and when they should be ready for pick-up. To be clear, these

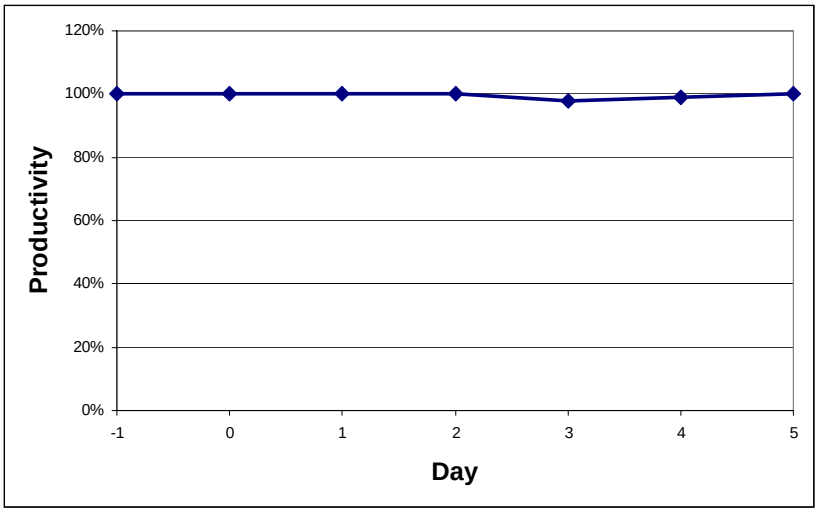

Figure 4. The impact of an internet outage of 3 day's duration on the production of an automobile parts manufacturer. For the reasons discussed in the text, it is reasonable to assume that an outage of this duration would have a small effect on production.

suppliers would not prepare items for pickup or delivery unless they were explicitly told to do so by the manufacturer using EDI transmissions. This communication typically happened using EDI communicated via the internet.

An outcome of this strategy was the lack of a 'learned supply-chain behavior': if they were unable to communicate with their suppliers, the supplies would not be delivered. From our analysis of the number of suppliers, the part counts and the frequency of ordering it is doubtful that these communications could be replicated via fax. As a result, there would be a restocking shortfall during an internet outage. The automobile parts manufacturer maintained a 3-5 day supply of stock on hand; there would be attempts to restock parts. However, if stocks depleted there would be increasing disruptions to the manufacturer's ability to make and deliver product, as shown in Table 2.

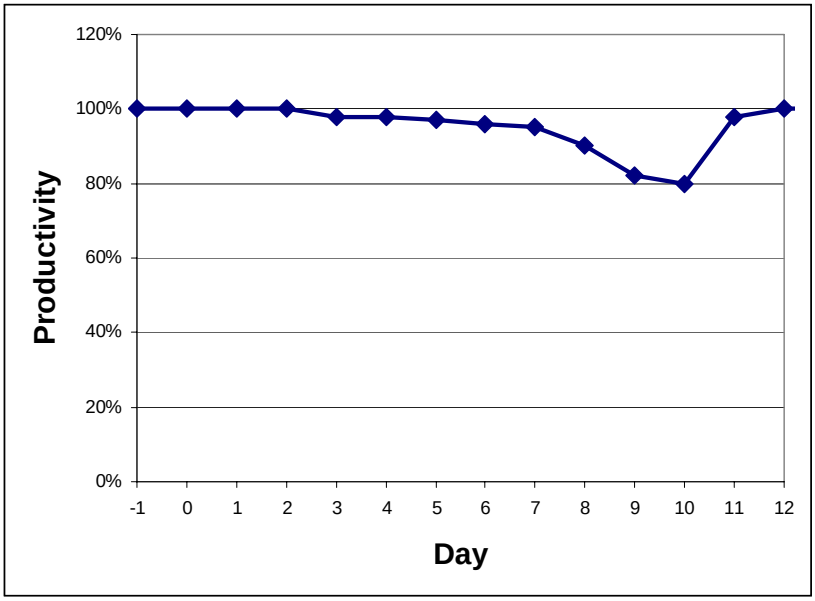

Figure 5. The impact of an internet outage of 10 day's duration on the production of an automobile parts manufacturer. For the reasons discussed in the text, it is reasonable to assume that an outage of this duration would have a small effect on production.

On the customer side, a few very large firms accounted for the great majority of the auto parts; the schedule of

products and delivery was reliably known for these accounts for at least a week. Shipments for these customers would likely be limited by production rather than managing order communication. There were also many smaller customers for after-market parts; the particular volume for specific items was not known before the order, and these orders would likely be adversely impacted by an internet outage.

From these considerations we estimate the change in production due to an internet outage as shown in Figures 4 and 5. Figure 4 shows the change in production for an internet outage of 3 days. This plot shows no change until day 3, at which there is a small decrease in production due to the inability to build some assemblies due to a shortage of some parts, understanding that the lack of a single part such as an O-ring can stop the production of a much bigger assembly. Figure 5 show the change in productivity for an outage of 10 day's duration. As before, there is no change until day 3, and then there is a continuing decline in productivity. We make the reasonable assumption that over a period of days supply chain experts at the manufacturer and suppliers will make accommodations to the internet outage, and that a steady-state production of 80% will be reached by the 10th day of the outage. This degradation is due to the de-optimized nature of the supply chain as well as the difficulty of handling after-market orders.

Oil-Refining Sector

The oil refinery that participated in the study was a regional oil refinery in the southern U.S. that produced specialty petroleum products, not a commodity product such as gasoline. We conducted expansive interviews over multiple days with 9 key individuals including the CIO and other functional executives. We estimate the change in production for this small refinery not as the result of an internet outage, but as the result of a non-directed attack on the refinery's SCADA system. SCADA (Supervisory Control and Data Acquisition) systems (also known by other names such as PCS and DCS) comprise the sensor, actuator, monitoring and logic systems required to run large physical systems such as oil refineries, gas pipelines and electric generation and distribution systems. The business operations at this refinery would be largely unaffected by an internet outage; supplies were typically ordered by phone with a lead time of weeks, and product orders while usually communicated via the internet could easily be handled via phone as well.

The refinery maintained a large enough reserve of product to handle (without any additional production) normal demand for a period of time exceeding that required to recover from most any event, including the physical destruction of major structures such as cracking towers. The result is that the 10-day event would impact the refinery's ability to "boil oil" and refine product, but not their ability to deliver product to the customer. Of course, they must eventually make up the lost production which would erode their ability to accept new future demand or increase their costs. Thus we treat the change in actual production as the variable of economic interest, even though the events discussed would most likely impact their cost and/or future revenue, but not their immediate revenue. Our treatment here more closely represents the impact of this event at a large gasoline refiner; any interruption to their production of gasoline would represent an immediate decrease in their revenue. As such, we believe that the results of this

estimation are largely applicable (through appropriate scaling) to the gasoline refining sector.

Oil refining involves performing a series of chemical reactions on the crude oil. At the refinery, there was storage for both the pre- and post-processed product; conceptually refining oil looks like storage-process-storage-process-storage-… until the refined product is stored for blending or delivery.

As noted above, the cyber-event that we are considering here is not an internet outage event, but a SCADA event. Abstractly, there are two relatively separate SCADA networks in any refinery: a monitor and control network (MCN), and a safety network (SN). The MCN has device actuation (valves, pumps, etc.), sensing (temperature, pressure, etc), data monitoring, control and archiving functionality; control is typically located at a central facility manned 24 hours a day. The refining process is continually monitored and adjusted to maintain optimal conditions for the current production run. The safety network has sensors and actuators and logic that are designed to prevent conditions that would lead to physical harm, such as an explosion. The sensors, logic, and actuators are located at the protected device, and act automatically. Devices that form the SN (more properly, several separate sensor-logic-actuator networks) may not be used as part of the MCN.

The event considered is the compromise of part of the safety network. This is presumed to result in the shutdown of the unit associated with the compromised part of the network, which would halt one step in the oil refining process. Because of the inter-process storage, there is some amount of time that the other processes would be able to run before they would need to shut down due to lack of pre-process product or storage capacity for the post-processed product. During this time the output of the refinery would be slowly decaying. Once the refinery was shut down, there would be no production of finished product.

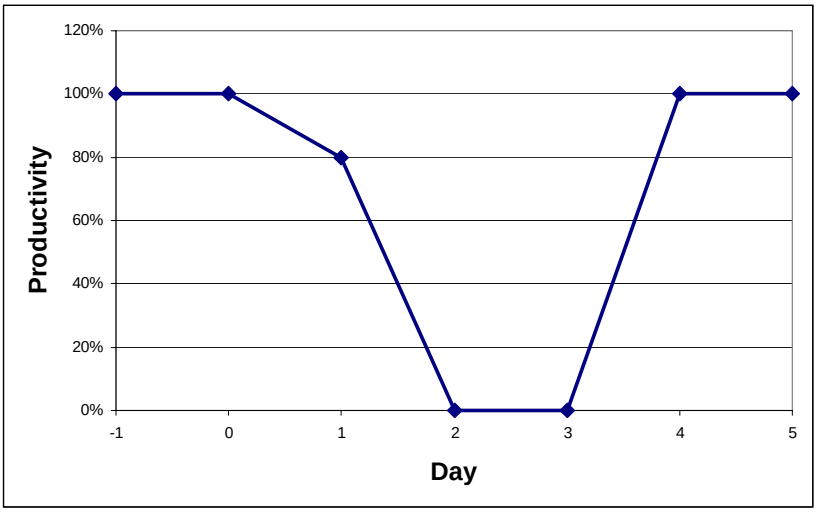

Figure 6. The impact of a SCADA safety system event of 3 day's duration on the production of an oil refiner. As described in the text, the refinery will run at a reduced rate for roughly a day, and then move to a warm shutdown state from which they can quickly restart.

After the SN compromise is detected and corrected, the refinery would be

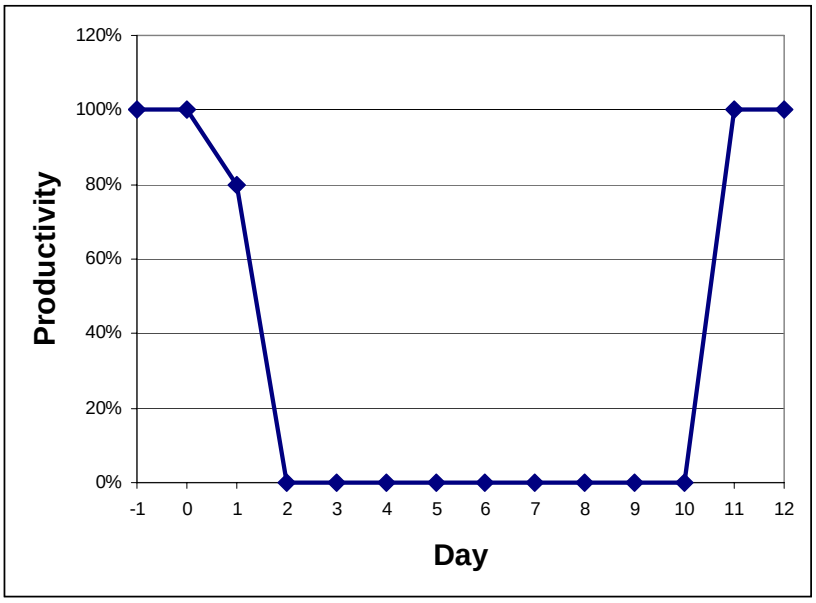

Figure 7. The impact of a SCADA safety system event of 10 day's duration on the production of an oil refiner. As described in the text, the refinery will run at a reduced rate for roughly a day, and then move to a warm shutdown state from which they can quickly restart.

restarted. While refineries would take several days to restart from a "cold" shutdown, the most likely course in this event would be a warm shutdown, where the functioning units are kept warm. In this instance, the refinery would resume 100% production in roughly 8 hours. These considerations result in the production vs. time plots shown in Figures 6 and 7.

Macro-Economic Model Results

The macro-economic model was run using the productivity data displayed in the figures above. For the analysis presented in this paper, the IIM was used to calculate the economic impact to particular regional U.S. economies due to perturbations to supply chains of three different sectors. Three separate, regional analyses were conducted, of two PADD (Petroleum Administration for Defense Districts) regions, namely PADD II – Midwest region, and PADD III – Gulf Coast Region. The impact of a supply chain perturbation to the automobile and electrical device manufacturing companies was evaluated in the Midwest region and the impact of a supply chain perturbation to the oil refining company was evaluated in the Gulf Coast Region. The following assumptions were made so as to provide appropriate inputs to the model:

1. Given that the supply chains of the companies were interrupted and possibly resulted in an inability to produce and/or ship products to customers, a supply constraint was introduced to the model. In other words, the supplying capability of the company was constrained by the introduced perturbation.

2. In order to calculate the perturbation to the output of a particular sector, it was necessary to convert the supply chain perturbations affecting a particular company to a sector-based perturbation. This was done by finding the relative size of each company's output to the sector output in a particular region. In order to calculate the relative size of each company's output to the size of the sector output in the

particular region, we assumed that the automotive manufacturer produced about 5% of the total national output of the Motor vehicle, body, trailer, and parts manufacturing sector, electrical device manufacturer produced about 5% of the total national output of the Electrical equipment and appliance manufacturing sector, and the oil refinery produced about 10% of the total national output of the Oil and gas sector. Knowing the size of the regional output for the three sectors, it followed that the automotive manufacturer represented approximately 7.5% of the regional output, the electrical manufacturer represented approximately 11.5% of regional output, and the oil and gas company represented about 16.1% of the regional output.

3. Daily perturbations (interruptions) to the three sectors were computed by converting the daily loss of sales to the company into sales reductions to the entire sector, by using the numbers obtained in (2). For example, a 5% loss of sales to a company which represents 7.5% of total regional output resulted in a perturbation of 0.00375 (5% x 7.5% = 0.05 x 0.075 = 0.00375) to that sector. The perturbation was then introduced as a supply constraint to the model. The model assumes there are no substitutions available, that is, it assumes that if the sales go down because the internet is down, these sales will not be substituted by work through the phone or fax for example. So, if there is an internet outage which prevents supply and delivery of goods, the model assumes that those goods will not be delivered.

4. The IIM produces results on a yearly basis, i.e. it outputs a loss that is distributed over one year. In order to compute the loss for a particular day, and a particular perturbation level, it was assumed that the losses are equally distributed throughout the year, so the yearly loss was divided by 365 days to obtain a daily loss.

5. Two regional IIM models, one for PADD II region (Midwest) and one for PADD III region (Gulf Coast) were used instead of a national IIM to provide a more focused and more accurate analysis of the relatively small perturbations. The two regions were chosen because the interviewed companies were located in them.

Mapping between the results obtained from the field studies and the inputs into the IIM occurred by utilizing the recovery curves produced by the field studies. The primary input into the IIM, the perturbation to a particular sector, was obtained by looking at the percentage loss of sales on a particular day for a particular company and the size of the company with the respect to the total regional sector output. Six different analyses were performed – for each company a 3-day and a 10-day event were evaluated. As an illustrative example, for the automobile manufacturer, the 3-day event was evaluated in the following manner. The daily percentage loss of sales was obtained from the recovery curves, and that number was then multiplied by the relative size of the company to the sector output to obtain the perturbation to the sector. In other words, if a company experienced a 5% reduction in sales, and the company represented 7.5% of total sector output for that region, then the resulting perturbation was 0.00375 (= 0.05 x 0.075). Daily perturbations were individually entered into a customized regional IIM model (i.e. for a

3-day perturbation 3 different calculations were made, one for each day), and the model produced a yearly loss to the U.S. economy. This yearly loss was assumed to be equally distributed throughout the year, and it was divided by 365 days to obtain the daily loss. This was done for all the days that the event lasted. The total loss due to a 3-day event was computed by adding the daily losses for the 3 days. Based on the interdependency matrix in the IIM, this total loss was then separated into direct losses (losses to the attacked sector) and indirect losses (losses to the sectors that are interconnected with the attacked sector), and a ratio between indirect and direct losses was computed to indicate the significance of the indirect losses, which are most often overlooked in the economic analyses of cyber events. For example, an indirect-to-direct ratio of 0.66 would indicate that for every dollar that is directly lost to the perturbed sector 66 cents are indirectly lost by some other sectors. Similarly, computations for other sectors followed the same procedure.

Manufacturing Sector

For the electrical and automobile parts manufacturer model estimations, we assumed that each manufacturer represented 5% of the total sector (electrical or automobile) manufacturing capacity nationwide. The macro-economic model estimates for the electrical manufacturer are shown in Figure 8. From that, we can see that a 3-day internet outage would result in no loss to the economy of the Midwest, while the 10-day internet outage would result in a loss of $22.6 million. These amounts include the direct economic losses to the manufacturing sector that was perturbed, as well as the total indirect losses due to reduced demand for the supplies needed by the manufacturer and the inability of the manufacturer's customers to produce their products. The losses for several days of the

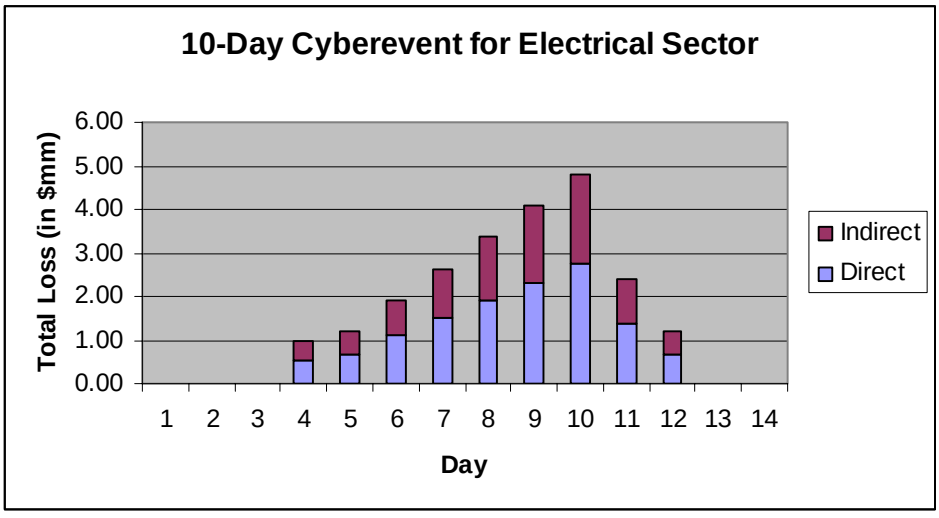

Figure 8. Estimates of the daily macro-economic cost to the Midwest regional economy of an internet outage of 10 day's duration to an electrical parts manufacturer. The manufacturer was assumed to represent 5% of the total sector capacity nationwide; direct losses are the economic losses to the perturbed sector, indirect losses include those due to reduced demand from suppliers and reduced sales by customers. The integrated loss is $22.6 million.

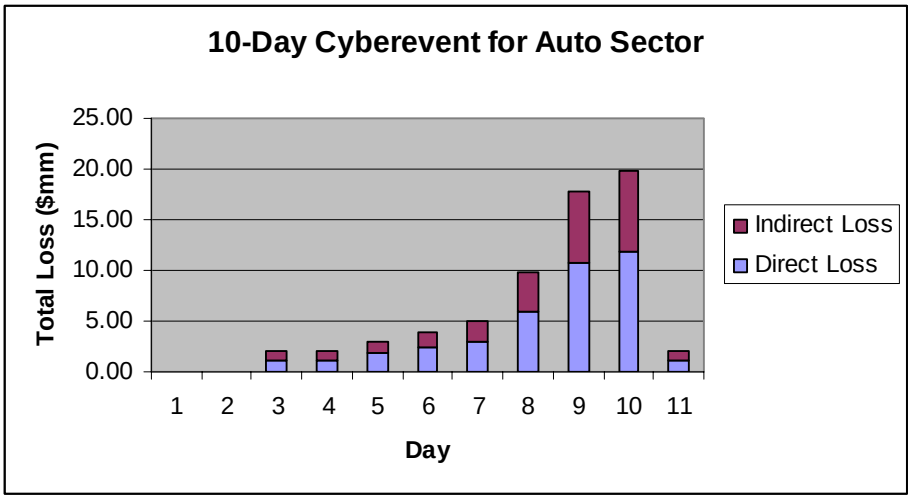

Figure 9. Estimates of the daily macro-economic cost to the Midwest regional economy of an internet outage of 10 day's duration to an automobile parts manufacturer. The manufacturer was assumed to represent 5% of the total sector capacity nationwide; direct losses are the economic losses to the perturbed sector, indirect losses include those due to reduced demand from suppliers and reduced sales by customers. The integrated loss is $65.16 million.

10-day internet outage are shown in Figure 8.

Figure 9 shows the analogous figure for the automobile parts manufacturer. The assumptions about the relative size of the manufacturer are the same as for the electrical parts manufacturer. The total economic losses for the 3-day outage are $2.96 million, the losses for the 10-day outage are $65.16 million. The higher figure for the 3-day event is due to the development of parts shortages sooner for the auto parts manufacturer; the larger amounts are also due to the relative sizes of the auto and electrical parts sectors.

Oil Refining Sector

We assumed that the refiner constitutes 10% of the gasoline sector supply nationwide. For this SCADA safety network event, the 3-day event would result in a total economic loss of $96.79 million, while the 10-day event would result in a loss of $405 million. Figure 10 shows some daily losses for the 10-day event.

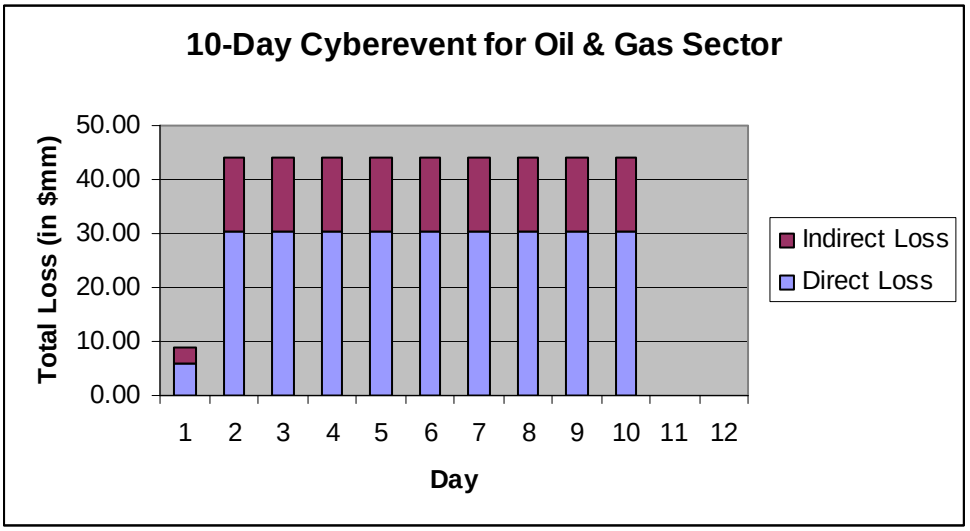

Figure 10. Estimates of the daily losses to the Gulf regional economy of a SCADA safety system event of 10 day's duration to an oil refiner. The refiner was assumed to represent 10% of the total sector capacity nationwide; direct losses are the economic losses to the firm, indirect losses include those due to reduced demand from suppliers and reduced sales by customers. The integrated loss is $405 million.

Discussion

Here we present the first data-based estimates of the macro-economic costs of two types of cyber events to the U.S. economy. These estimates are based on the impact that the cyber events had on the productivity of the firm, and the macro-economic model, which provides a means for determining how a change in the production or demand in one sector will ripple through the economy, resulting in the indirect cost being some multiple of the direct losses suffered by the affected firm.

Sector	Event type	3-day event	10-day event	Indirect-to-Direct Multiplier
Electrical Parts	internet outage	$0.00 M	$22.6 M	0.76
Automobile Parts	internet outage	$2.96 M	$65.16 M	0.66
Oil Refining	internet outage	$0.0 M	$0.0 M	
Oil Refining	SCADA safety network	$96.79M	$404.76 M	0.44

Table 3. Estimated total economic costs of two types of cyber events in three industries. A macroeconomic model was used estimate the direct and indirect (both up and downstream) costs based on impacts the indicated event types would have on sector firms determined from field studies. The multiplier is the ratio of the indirect cost to direct cost. Costs in millions of dollars.

It is important to note the indirect-to-direct losses multiplier, which indicates how much in terms of dollars the U.S. economy loses indirectly for every $1 lost directly. Oftentimes in evaluation of costs due to cyber events indirect losses are completely ignored, however, very often those costs are almost as high, and sometimes higher, than the direct costs, and they should be of great concern to the interconnected sectors. The major benefit of IIM analysis is that it provides one with indirect-to-direct multipliers so that even if the accurate estimates of direct costs were not known, as is usually the case, one could make assumptions about the size of the direct losses, and could then easily compute the resulting indirect losses by simply multiplying the assumed direct losses with the indirect-to-direct multiplier.

While the IIM benefits from BEA data collections and a community of users and developers that continue to pursue improvements such as regional sub-model developments that correspond to national data and protect privacy, it is sometimes attacked by critics who complain about potential misuse. Other complaints include that the IIM is a static, linear model that does not account for market-place substitutions and that it is updated on the national level only once every 5 years. Due to these limitations IIM is sometimes said to oversimplify the economic connections with which it deals. The limitation of the static model can be ignored if one recognizes that the changes to the production output that result in lost revenue for one company come at the expense of gained revenue for another that occurs at the same rate as other normal market place changes that the BEA is measuring. The limitation of a linear model is overcome by the fact that the changes due to short-term internet outages are small compared to the overall economy and can thus be dealt with as a small perturbation in a linear fashion. The fact that IIM does not deal with market-place substitutions limits its use to cases in which a) no important substitution capabilities exist, and b) impacts of substitution are derived as a direct analytical result. The analysis presented in this paper focuses on those two cases. IIM is also limited in that its national accounts are updated only every five years, which can be problematic. However, given that no better data collection effort exists, we have no better alternative and we accept this limitation as a necessity. Hence, due to the reasons just stated, this paper assumes that a linear, static IIM is appropriate for analyzing these internet outages.

Resiliency of Firms

Although the auto parts and the electrical parts firms are both manufacturing firms, the impact of short internet outages differs greatly between the two. In the case of the electrical firm, there was little or no expected supply chain disruption for several days because of the manner in which their supply chain relationships have evolved. The expectation is that the suppliers would continue to behave as they have in the past, and deliver the ordered or forecasted number of supplies to the various plants, enabling them to continue producing product. The auto manufacturer takes a much more control-oriented approach; here the expected behavior is to not ship until told to do so. That, along with the levels of on-hand stock leads to their running into production difficulties

in a few days. To contrast, the oil refiner maintains enough production in stock to cover most any conceivable production disruption.

Note that there is not necessarily a correlation between utilization of technology and the resiliency of a given firm to cyber disruptions. The electrical manufacturer is vigorously involved in moving their entire supply chain to a web-based order and quality management system; their aim is to become 100% dependent on the information infrastructure to manage their supply chain. The auto parts manufacturer is also attempting to move the management of their supply chain to be entirely internet-based (web and EDI). The point here is that it is not just how dependent a firm is on the information infrastructure, but also how they use it that determines how brittle a firm's supply chain operation is to internet-based communication disruptions: the electrical firm uses technology to enable a fairly autonomous flow of supplies while the auto firm uses technology to implement a top-down supply control structure. An interesting question is whether the clear benefits from the highly controlled, optimized approach outweigh the increased resiliency inherent in the less optimized, more autonomous approach, or whether there are approaches that combine the best characteristics of both.

An interesting side note is how firm's of supply chain executives view the future of supply chain management. Almost without exception, when asked what they would like to see most from a more tightly integrated supply chain (implicitly enabled by the information infrastructure) is a greater view into the future demands of their customers: these executives would like access to their customer's present and future sales predictions so they could better plan their business. This goes beyond the 10 weeks of forecasts that the electrical and auto manufacturers currently send to their large suppliers; the suppliers want to be more integrated in their customer's sales activities. There is little doubt that this would result in fewer supply chain surprises, the unexpected demand for product from the supply chain to cover an unplanned large order. The result of these surprises is that suppliers all the way down the supply chain carry excess stock for these eventualities, which is economically inefficient. While this is technically not a problem, it is from an organizational trust perspective: while a firm's supply chain executive will tell you they would really like this access to their customer's sales projections, when asked if they would allow their suppliers similar access to their sales projections, that same supply chain executive (with very few exceptions) will say that they will not give suppliers access to that type of information.

Resilience of the supply chain

Another aspect of resilience in the supply chain with respect to information infrastructure disruptions can be seen from Table 2: except for the largest businesses (annual revenue in the billions), an extended internet outage would not hamper the ability of any supplier to produce and deliver product to their customers. Even large suppliers such as suppliers B and C were not critically reliant on the information infrastructure. From these results it would be reasonable to conclude that, as of the time of these field studies, beyond the first tier, the manufacturing sector supply chain is quite resilient to internet disruptions as

a result of the lack of dependency on the internet for managing orders with their customers and suppliers.

This is likely to change in the future as more firms start using the internet for order placement and processing. The general manager at Supplier G, which at the time of the interview did not have a web site and relied on the internet only for mail, was talking about being able to process EDI transmissions in the near future. This does not necessarily mean that small firms will become brittle with respect to internet disruptions; their small scale and current reliance on phone and fax for running their business makes it highly likely they would be able resume those activities with little or no disruption in their ability to produce and ship product. This is borne out by supplier D, who in conversation with the automobile parts manufacturer decided to not implement EDI to transact business due to the limited volume of business between the two firms.

It is not clear how applicable this result across business sectors. Our results show that it may well be the case that the oil refining sector is largely unaffected by internet disruptions; but it is easy to see how other sectors may be highly susceptible to internet disruptions, such as the express package delivery and integrated logistics sectors.

Acknowledgments

We would like to thank the many individuals and organizations that very generously gave their time and support in enabling us to conduct these field studies; without their interest and efforts this work would not have been possible
.
This work was supported under Award number 2000-DT-CX-K001 from the Office for Domestic Preparedness, Department of Homeland Security. Points of view in this document are those of the authors and do not necessarily represent the official position of the U.S. Department of Homeland Security.

Bibliography

[Cas04] Cashell, Brian, William D. Jackson, Mark Jickling, and Baird Webel (2004): The Economic Impact of Cyber-Attacks, Congressional
Research Service Documents, CRS RL32331 (Washington).

[DoC98] U.S. Department of Commerce, Bureau of Economic Analysis. “Benchmark Input-Output Accounts of the United States, 1992.” Washington DC: U.S. Government Printing Office, 1998.

[DoC97] U.S. Department of Commerce, Bureau of Economic Analysis. “Regional Multipliers: A User Handbook for the Regional Input-Output Modeling System (RIMS II).” Washington DC: U.S. Government Printing Office, 1997.

[Dyn05] Dynes, S., Brechbühl, H. and M. E. Johnson, 2005. Information Security in the Extended Enterprise: Some Initial Results From a Field Study of an Industrial Firm. Workshop on the Economics of Information Security, Cambridge, MA. June 2005

[Hai01] Haimes, Y. Y., and P. Jiang, 2001. “Leontief-Based Model of Risk in Complex Interconnected Infrastructures.” *ASCE Journal of Infrastructure Systems*, **7**(1): 1-12.

[Hai04] Haimes, Y. Y., B. M. Horowitz, J. H. Lambert, J. R. Santos, K. G. Crowther, and C. Lian. “Inoperability Input-Output Model (IIM) for Interdependent Infrastructure Sectors: Theory and Methodology.” 2004

[Leo1966] Leontief, W.W. “Input-Output Economics”. New York: Oxford University Press, 1966.

[Joh2005] Johnson, M. Eric 2005, A Broader Context for Information Security, *Financial Times*, September 16, 4.

[San04] 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.

[Yin94] Yin, R. K. 1994. Case Study Research: Design and Methods, 2nd edn. Thousand Oaks, CA: Sage