State/Local CIP Risk Analysis: First Results and Emerging Trends in the Data

David Daniels (ddaniels@dsbox.com), Bryan Ware (bware@dsbox.com)
Digital Sandbox, Inc., McLean, VA, ddaniels@dsbox.com

Abstract—*Digital Sandbox, Inc. has conducted a number of scenario-based analyses of the risks facing individual states and urban areas from terrorist attack. This paper reports on the common features in the distributions of risk to critical infrastructure in the studied areas. In addition, three metrics (population, a population-and-density-based-index, and gross domestic product) are compared to the measured asset risk distributions at the county level. We find that none of these measures provides a good match to the measured asset risk distributions. We conclude that jurisdictions should strive to find ways to simplify the direct measurement of scenario-based asset risk, rather than relying on simple heuristics, when making risk-reduction investment decisions.*

1. INTRODUCTION

Since the events of September 11, 2001 and the subsequent formation of the U.S. Department of Homeland Security (DHS), a great deal of attention has been given to the determination of risks to critical infrastructure and key resources (CI/KR) at the federal, state, and local levels. DHS has spent over $15B through its various grant programs aimed at improving security [1], and since 2006, this money has been allocated based on DHS' understanding of the risks facing the nation. At the state and local level, much of this money and additional state and local investments have been allocated based on an understanding of their specific jurisdictional risks. The National Infrastructure Protection Plan (NIPP) [2] and other DHS documents lay out a general framework for measuring risks to CI/KR, but to date there has been no corresponding definitive formula or approved system available to enable the risk analysis needed to make that framework operational.

In this environment, Digital Sandbox, Inc. (DSI) has built a system for enabling the measurement of CI/KR risks from terrorism and natural hazards, and has implemented the system in 13 state and urban area locations with the help of state and/or local government officials. That system has been used to make resource allocation decisions, to justify requests for further investments, and to prioritize risk management activities.

This paper represents the first publication of results from our risk analyses to date. Several previous research efforts have attempted to evaluate the risks of states or urban areas from terrorism (e.g., [3], [5], [6], [7], [7]). To date, these have focused either on the aggregate risks of the various entities (e.g., for allocation of national resources), or on individual risk scenarios (e.g., for preventing specific types of events). The authors of this paper are not aware of prior publications that have evaluated aggregate asset risks *within* a state or urban area. They are also unaware of any published efforts that involve the analysis by state and local officials of the assets within their own jurisdictions. In order to protect the privacy of our clients, and to protect the necessarily sensitive nature of this data, no results from individual analyses are shown, but general trends emerging from the aggregation are reported.

Each jurisdiction has unique risk characteristics. It is not expected that other jurisdictions will necessarily share the same characteristics as the ones whose data is summarized in this paper. Rather, these emerging trends are shared as a benchmark for future studies and evaluations.

This paper is organized into four sections. We first describe our risk methodology, and make the distinction between asset/scenario risk and population risk. Next, we present the data and describe common features in the resulting asset risk distributions. We next compare these asset risk distributions to several alternate risk metrics, including population-derived and economic-inspired measures, to see if these simpler metrics may be used as substitutes for a bottom-up asset risk analysis. Finally, we draw conclusions about the applicability of proxy data as a substitute for a scenario-based risk analysis, and highlight the implications for state and local homeland security risk management efforts.

2. METHODOLOGY

DSI's risk analysis methodology is built upon hazard-asset scenarios[1]. Hazards may be terrorism attack modes, such as an attack by improvised explosive device (IED), vehicle-borne IED (VBIED), or chemical weapon; hazards may be natural disasters, such as hurricanes, earthquakes, or

[1] For a brief overview of some alternative risk methodologies for comparison, see the discussion in [8].

wildfires; or they may be any other hazardous situation one may care to model, such as industrial accidents, crime, or other disruptions. The results presented in this paper only concern terrorism hazards, but the DSI risk model is built to support all hazards.

The results in this paper were derived by evaluating risks against a standard portfolio of terrorist attack types. There are several lists of terrorist hazards in common use (see, e.g., [2], [3], [4], [6], [9], [10], [11]). DSI compared each of these lists to construct a minimal superset, resulting in the following portfolio of sixteen standard terrorist hazards: IED, VBIED, conventional attack (small arms), hostage taking/assassination, arson/incendiary attack, sabotage/theft, aircraft-as-a-weapon, maritime attack, agro-terrorism, food or water contamination, non-contagious biological agent, contagious biological agent, chemical agent, radiological dispersion device (RDD), nuclear device, and cyber attack.

Assets can be any structure, site, system, or collection that share common vulnerabilities and/or consequences. Ideally, one would select as assets those objects (physical or logical) that are reasonably well independent from one another, both physically and logically; in practice, however, one must balance independence with one's ability to model vulnerabilities and consequences. Most assets are physical structures, such as buildings, dams, and bridges; some assets are collections of structures, such as military bases, airports, campuses, or other "sites"; and some are geographically-dispersed but logically-connected entities such as transportation systems or communications networks.

The asset aggregation level is ultimately one of convenience. For instance, one may model a subway system as a collection of independent assets (e.g., stations, and perhaps switches, yards, tunnels, bridges, and other sections of track), each of which has unique risks, or one may model the same subway system as a single asset which may be compromised in a variety of ways. Either way can yield similar results, and the correct choice depends upon the expected hazard scenarios and one's ability to model.

For this analysis, the lists of assets in each jurisdiction were developed by local homeland security and/or emergency managers with whom DSI worked to initialize their risk management systems. For each installation, managers were asked to compile a list of their most-critical assets, which they intended to actively manage in the DSI system. DSI analysts worked with the state or local managers to compile initial lists, to compare with open-source information to ensure no high-visibility assets were missed, and in some cases to implement an initial priority screening in order to limit the total number of assets considered. Ultimately, the initial list of assets was determined by the clients, based upon management objectives unique to each jurisdiction.

A risk scenario is constructed as a combination of a hazard and an asset. For each scenario, a risk value may be constructed according to the standard formula

$$R=T\times V\times C, \tag{1}$$

where R is the risk; T is the threat likelihood, which is related to the likelihood that the scenario is initiated; V is the vulnerability, the conditional probability that the scenario causes damage; and C is a measure of the expected consequence of a scenario that exploits the vulnerability measured by V. In the DSI risk methodology, the consequence, C, is expressed as a weighted sum of four different types of consequence: human (H), economic (E), mission (M), and psychological (P).

Assuming they are independent, scenario risks may be added to yield risks to higher-level entities. The risks from all scenarios involving a single asset may be summed to yield the asset risk, and all asset risks in a geographic area (e.g., a jurisdiction) may be summed to yield jurisdiction risk. Alternately, asset risks may be summed for all assets in a given CI/KR sector [2] to yield sector risk, or scenario risks from all scenarios involving a particular hazard may be summed (e.g., over all assets) to yield a hazard risk.

The DSI risk analysis system captures the inputs T, V, and C and calculates R for all scenarios involving the input assets and a portfolio of hazards. The values for T, V, and the four values of C may be input as cardinal values, if known, or estimated from the ordinal responses of subject matter experts (SMEs). DSI includes default T values based on our evaluation of national-level threat likelihoods and the composition of national-level CI/KR in a state or local area. These T values are relatively normalized, and may be seen as conditional probabilities. For this analysis, the other risk inputs were obtained by conducting SME workshops in each jurisdiction to evaluate the vulnerabilities and expected consequences on 5-point ordinal scales for each asset-hazard scenario in that jurisdiction. These workshops were normally conducted over the course of one to two days, after the list of assets was determined. DSI analysts converted the ordinal responses for vulnerability and consequence to cardinal scales based on our experience in conducting similar SME-driven evaluations (roughly based on a base-2 logarithmic scale for vulnerability, and a base-10 logarithmic scale for consequences). The final values for T, V, C, and R for each scenario, as well as the aggregations of risk by asset, sub-jurisdiction, sector, subsector, and hazard, were reviewed by the state or local homeland security or emergency managers for consistency, but no independent evaluation of the data was attempted.

Data were collected in this way from 5 U.S. states and 8 urban areas from January 2007 through March 2009. While not proportionately representative, a variety of sizes (e.g., by population) and geographic locations are found among the states and urban areas included in this study. In order to protect client confidentiality, no individual clients' data will be identified in this paper; instead, in the interests of benchmarking risk analysis efforts in other states and urban areas, aggregate results will be presented.

Asset Risk vs. Population Risk

The general formula for risk (Equation 1) can be extended to as many assets as one would care to include. In the logical extreme, one may incorporate all buildings, houses, cars, and even people in a jurisdiction as assets. As we have observed, many terrorist attacks are not aimed at CI/KR assets, but rather at individuals or groups of people. Although the risk to any individual from terrorist attack is miniscule, the aggregate risk to the population of an area may be measurable. Moreover, the distribution of risk to the general population (population risk) may be substantially different from the distribution of risk to CI/KR assets (asset risk). Evaluating both population risk and asset risk is important to a complete understanding of the risks faced in a given area.

Although both asset risk and population risk are derived from the same basic risk equation, in practice they are computed differently. The distribution of asset risk is calculated by evaluating the T, V, and C components of individual hazard-asset scenarios, whereas the population risk is determined on a census block level as the product of the population (maximum of weekday daytime and nighttime populations) and the population density. It can be shown that this metric is equivalent, up to a constant, with scenario-based asset risk with the assumption that the vulnerabilities of all people are constant and the consequences depend only on the local density of people in the vicinity of the attack [12].

This paper focuses on asst risk results aggregated across 13 implementations of the DSI risk analysis system, but in Section 4 asset risk is compared with simple proxies for risk, one of which is population risk.

The construction of population risk, as used in the DSI risk analysis methodology and as reported in this paper, is identical to the Population Index metric used by DHS in its Homeland Security Grant Program for states (State Homeland Security Program) and urban areas (Urban Area Security Initiative) [7].

3. Data and Observations

Assets and Risk

The data collection effort was part of a series of studies in asset risk management. For that effort, the five states and eight urban areas were each asked to provide a list of their "100 or so" most critical assets. The actual numbers of assets analyzed in each jurisdiction ranged from 99 to 1,006, with an average of 456. Neither the average number of assets submitted nor the range differed significantly between states or urban areas (Table 1).

	Number	*Avg. Assets*	*Std. Dev.*	*Min. Assets*	*Max. Assets*
States	5	425	351	99	1,006
Urban Areas	8	475	287	212	940
Total	13	456	299	99	1,006

Table 1: Number of Assets Analyzed

Regardless of the number of assets in each sample, asset risk was concentrated in relatively few of them (Figure 1). The dark gray lines represent states, and the lighter gray lines represent urban areas. In all cases, the 100 highest-risk assets accounted for no less than 2/3 of the total asset risk evaluated. In 10 of the 13 samples, the top 30 assets accounted for over 50% of the total risk.

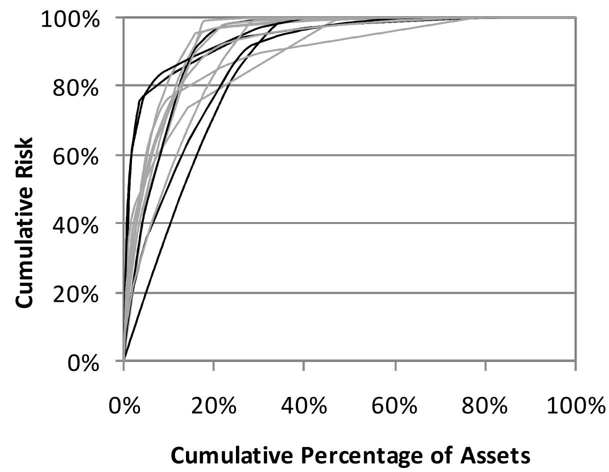

Figure 1: Asset-Risk Pareto Chart

If these data are representative, this implies that even a risk management program of modest scope can manage much of a jurisdiction's asset risk by focusing on the few highest-risk assets. In addition, although it may be beneficial to track many assets in a risk management program, the greatest risk mitigation efforts may be had by focusing on the few highest-risk assets.

These results represent terrorism risk, which is expected to center more closely on a few high-visibility assets than other types of risk. For other types of risk, the high-risk assets may be different, and the risk may be distributed more broadly across a jurisdiction's assets.

Sub-Jurisdictions and Risk

Although some risk management functions are coordinated at a high level from a region-wide perspective, it is often also convenient to break risk into geographic components or sub-jurisdictions. For instance, risk mitigation investments (e.g., the distribution of homeland security grant funds) are often made within a state or urban area on a risk-proportional basis to participating sub-jurisdictions. These sub-jurisdictions may be counties or regions within a state, or counties, cities, or even individual neighborhoods (e.g., wards, precincts, or other administrative units) in an urban area.

In this study, the five states and eight urban areas self-defined administrative sub-jurisdictions by which to manage risk within the jurisdiction. The number of jurisdictions ranged from 4 to over 40, with the median and mode both 13. Nine of the thirteen jurisdictions defined between 9 and 16 sub-jurisdictions. States in general defined more sub-jurisdictions (counties, typically) than did urban areas, as expected.

As with the distribution of risk among a jurisdiction's assets, we have also found the distribution of risk among sub-jurisdictions to be highly non-uniform (Figure 2). Urban areas, represented by the lighter gray lines, had a higher concentration of asset risk, with the central jurisdiction containing between 44% and 96% of the entire jurisdiction's asset risk (median about 80%). Since states often contain multiple urban areas with concentrations of CI/KR assets, the central sub-jurisdiction was found to contain less of the total risk: between 23% and 39%, with a median of 33%.

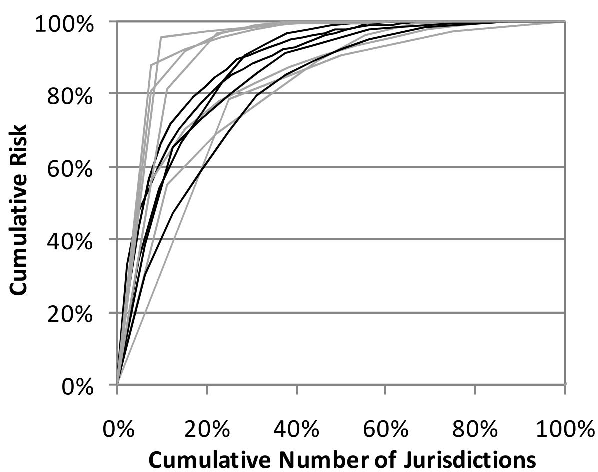

Figure 2: Sub-Jurisdiction Pareto Chart

Sectors and Risk

DHS has defined 18 CI/KR sectors [2], related by industry. While some industries are concentrated in certain parts of the country (e.g., petroleum in Texas, chemicals in New Jersey, banking in New York City), most sectors have assets that are more uniformly dispersed (e.g., emergency management, telecommunications, healthcare, commercial, transportation).

There is no reason to believe that there is the same amount of risk (of terrorism or any other hazards) in each of these sectors. To date, there are no DHS results that compare the aggregate risk in each sector across the country, and it appears self-evident that there would be less aggregate risk in the National Monuments and Icons sector than there would be in the Commercial sector, for instance, both because of the latter sector's greater number of CI/KR assets and their greater potential for disruption individually.

Across the thirteen jurisdictions included in this study, several CI/KR sectors appear to have systematically different-from-average aggregate risks (Figure 3). The high and low error bars represent the full range of risks evidenced in the sample of 13 jurisdictions. The gray boxes represent the range spanned by the middle 7, roughly the middle two quartiles. If risks were dispersed throughout the 18 sectors evenly, there would be about 5½% of risk in each sector. The Commercial sector has universally more risk than that, and the Government sector almost always does as well. Conversely, the Defense Industrial Base (DIB), Telecommunications, Information Technology, Agriculture, Postal and Shipping, and Critical Manufacturing sectors had systematically less than expected risk.

Hazards (Attack Types) and Risk

Similarly to the sectors, the 16 terrorism hazards included in this study are meant to span the space of potential high-risk attacks; they are not chosen to have equal risk. Indeed, the vehicle-borne improvised explosive device (VBIED) attack consistently represents the greatest risk to the jurisdictions evaluated in this effort, accounting for approximately one-half of the entire terrorism risk experienced by these jurisdictions (Figure 4). As in Figure 3, the error bars represent the full range of risks experienced by the collection of 13 jurisdictions, and the gray boxes represent the middle 7 jurisdictions.

The only other systematically higher-than-average attack types were Aircraft as a Weapon and Nuclear Attack, due to the catastrophic nature of the destruction that could be caused by these attacks and the difficulty in protecting CI/KR assets against them.

Though typically representing relatively low risk, there were several jurisdictions in which Contagious Biological, IED, or Incendiary Attacks contributed a large share of the overall jurisdiction risk.

The finding that VBIED represents such a high percentage of jurisdictions' asset risk is somewhat surprising, given the variety of assets covered by the jurisdictions in this study and the multiple, independent SME panels that evaluated the risks. The good news is that VBIED attacks are somewhat easier to counter than some other types of terrorist attacks.

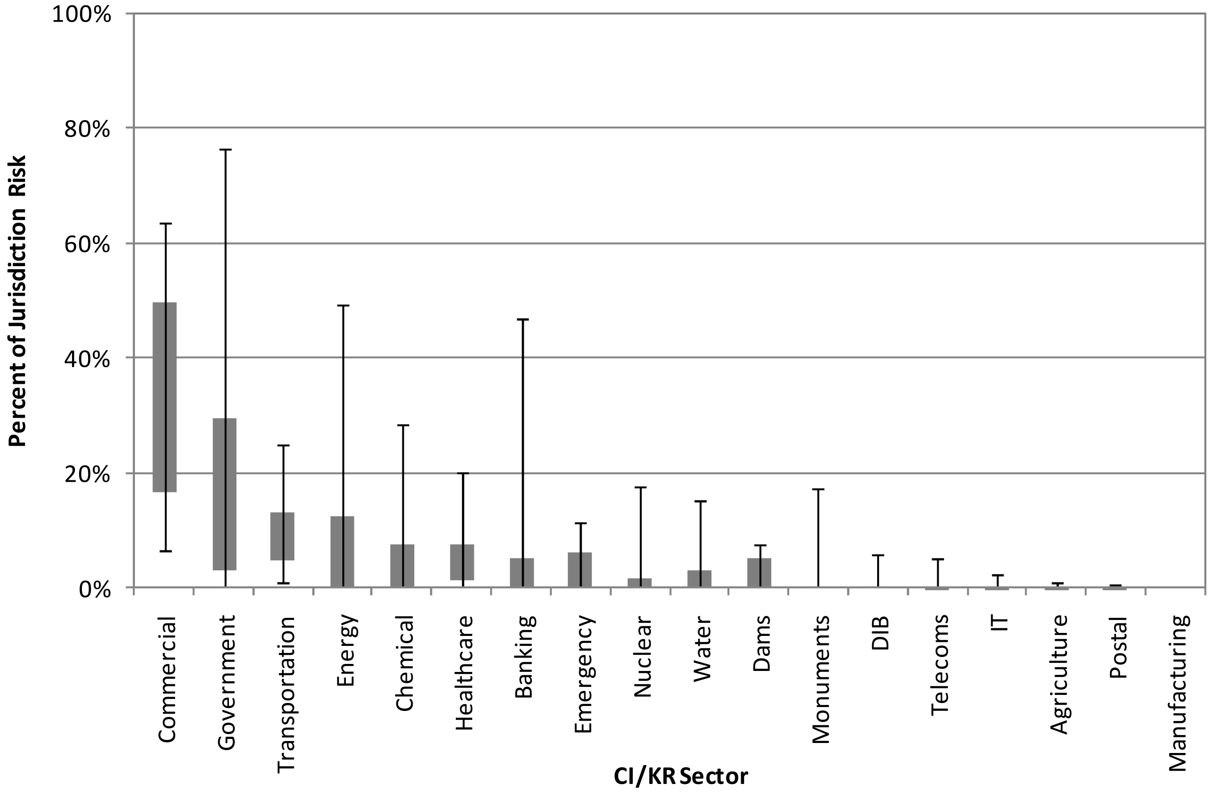

Figure 3: Sector Risk Chart

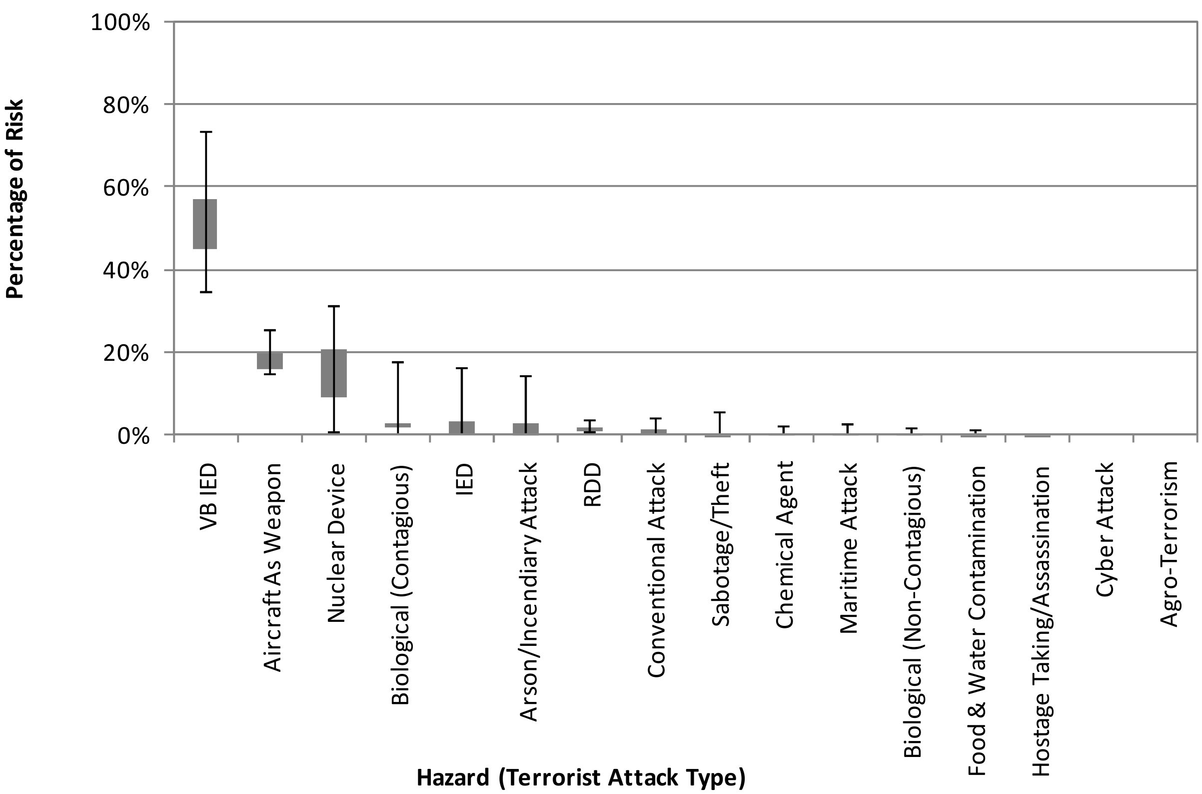

Figure 4: Hazard Risk Chart

4. Comparison With Other Metrics

Constructing a list of the most critical assets in a jurisdiction and evaluating the risks against them is necessary for any risk management program, and one useful product of such an exercise is a breakdown of asset risk by sub-jurisdiction. However, there are some applications, such as investment allocations, where the details of risk are not needed and where time or resource constraints may limit the ability to construct such a measure.

In these cases, a proxy measure for risk is often sought, and several candidate measures have been used. For instance, the DHS Homeland Security Grant Program (HSGP) funds were allocated according to population until the Department developed a more robust ability to measure risk. In extreme cases, a politically expedient technique is to divide available resources equally among all (sub-)jurisdictions. As demonstrated by the data above, such a measure rarely approximates the actual asset risk distribution.

With asset risk measurements from thirteen jurisdictions, it is useful to compare the measured distribution of risk to a number of commonly-used proxies for risk to examine the utility of these alternatives. In the current paper, we focus on the relative distributions of risk within a jurisdiction, rather than the differences in risk between the thirteen jurisdictions. We compare the aggregate asset risk in each self-defined sub-jurisdiction to the following metrics: number of reported CI/KR assets, population, the Population Index used in the HSGP risk formula, and gross domestic product (GDP), where possible.

Number of Assets

The number of assets reported in sub-jurisdictions is self-reported by the jurisdictions participating in this study. If the risks to each of the assets were similar across a jurisdiction, one would expect this metric to correlate with asset risk. However, although jurisdictions were asked to report their most critical assets, there is inevitable political pressure to include at least some assets from each sub-jurisdiction. Indeed, in the six urban areas where the sub-jurisdictions were defined at the county level, none of the outlying counties were devoid of assets.

This metric is difficult to construct without the assistance of state and local officials. Therefore, the usefulness of this metric as a proxy for risk is somewhat dubious.

Population

For population, we use the 2007 population estimates, as reported by the U.S. Census Bureau. These estimates represent the resident (i.e., nighttime) population of each sub-jurisdiction. This metric was chosen for simplicity, because it is readily available, requires little computation, and is frequently mentioned in allocation discussions. A better (closer match with asset risk) population metric may be obtained by including net commuters and expected numbers of visitors to an area.

Population Index

The Population Index is a specific metric used by DHS in its HSGP allocation formula [7]. It represents the risk faced by the people in an area to a localized terrorist attack (e.g., an IED or VBIED, not a nuclear device) aimed at inflicting damage on the population (not CI/KR assets). For this study, we used the metric as defined for the FY 2009 grants, which was based primarily on 2007 Census Bureau population estimates.

Population Index, as defined by DHS, is a product of population and population density (population divided by land area) which yields consistent results over any level of geographic aggregation. Simply squaring total population and dividing by land area yields inconsistent results when computed at different geographic levels. That metric, for instance, yields a number which is larger for the state of Illinois than the same metric does for the city of Chicago. Population Index does not suffer from such issues.

Population includes Census Bureau resident population estimates, as well as growth-adjusted[2] net commuters at the census tract level, and an estimate of the average number of daily visitors to select urban areas. Land area is taken directly from the Census Bureau. Population Index is computed at the census block level, which is the smallest level of geographic division at which data is published by the Census Bureau. Once calculated at the census block, the Population Index of each block in a county, urban area, or state is added, as appropriate, to determine the composite Population Index for that area.

Economic Index

We use a measure of Gross Domestic Product (GDP) as our final risk proxy metric, the assumption being that CI/KR assets tend to be clustered in areas of high economic activity. The Bureau of Economic Analysis (BEA), U.S. Department of Commerce, publishes annual updates of their estimates of GDP for the nation, for states, and for urban areas. We used the most recent version of the BEA data available at the time of this study, which represented the year 2006.

We follow the BEA methodology [13] as closely as possible to construct GDP for individual counties. In summary, the state-level GDP is distributed among industries, as published by BEA. The portion of a state's GDP for each industry is then apportioned to the state's counties

[2] Commuters are published by the Census Bureau after each decennial census. The most recent commuter data is from the 2000 census. The net number of daily commuters to a census tract (incoming minus outgoing) is multiplied by a growth factor proportional to the residential population growth from 2000 to 2007 of the census designated place in which the tract is found.

proportionately to the earnings by industry as reported through BEA's local area personal income accounts. BEA reports on GDP for 88 industry types, and personal income for 114 NAICS codes. We calculate GDP for 62 non-overlapping industry segments which cover all industries and into which both the 88 industry types and 114 NAICS codes can be mapped.

Because of the inherent economic interconnections within urban areas, GDP may not be a very good measure of economic activity for sub-urban areas; nevertheless, we include a county-level GDP for comparison with asset risk in order to include at least one non-population metric.

Comparison with Asset Risk

We are interested in the ability of these four risk proxy metrics to reproduce the measured distribution of asset risk among the sub-jurisdictions. In making resource allocation decisions, this ability is most important for those sub-jurisdictions that contain a high percentage of a jurisdiction's risk. Because the risk in these high-risk jurisdictions is often driven by just a few CI/KR assets, though, it is also in these high-risk sub-jurisdictions where proxy metrics are likely to perform worst.

We therefore present a comparison of the four proxy metrics to measured asset risk for the jurisdictions in this study that represent a high percentage of their jurisdiction's risk. Figure 5–Figure 8 display the sub-jurisdiction asset risk as a percentage of the respective jurisdiction asset risk on the horizontal axes, and the proxy metric (sub-jurisdiction share) is displayed on the vertical axes. If the proxy metric corresponded exactly to asset risk, one would expect the points to line up on the dashed 45° line. In fact, one can easily see that for three of the four proxies (all but Population Index), the proxy metric is systematically lower than asset risk for high-risk sub-jurisdictions. Allocating resources according to one of these proxies will therefore systematically over-emphasize low-risk jurisdictions.

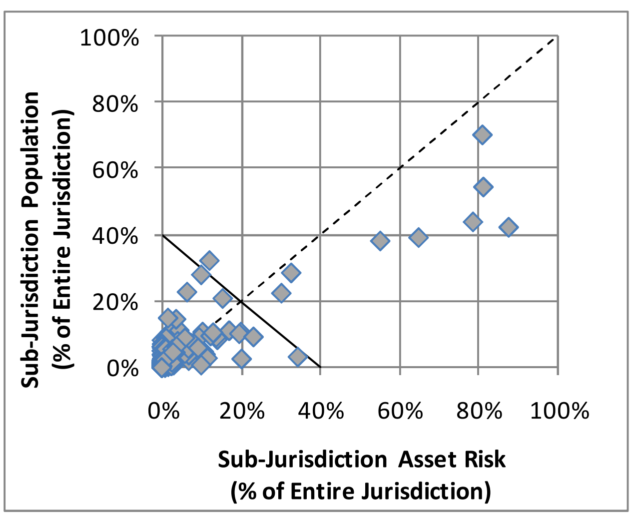

Figure 5: Population vs. Asset Risk Chart

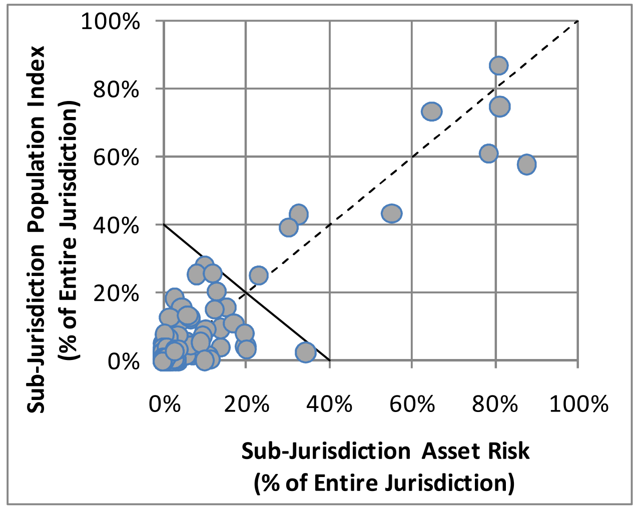

Figure 6: Population Index vs. Asset Risk Chart

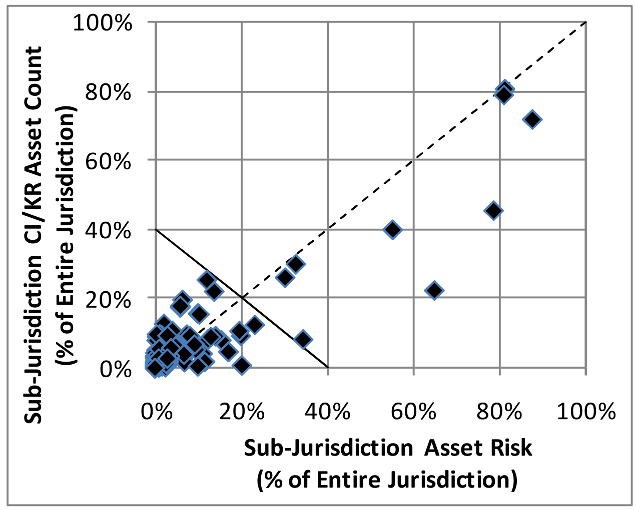

Figure 7: Asset Count vs. Asset Risk Chart

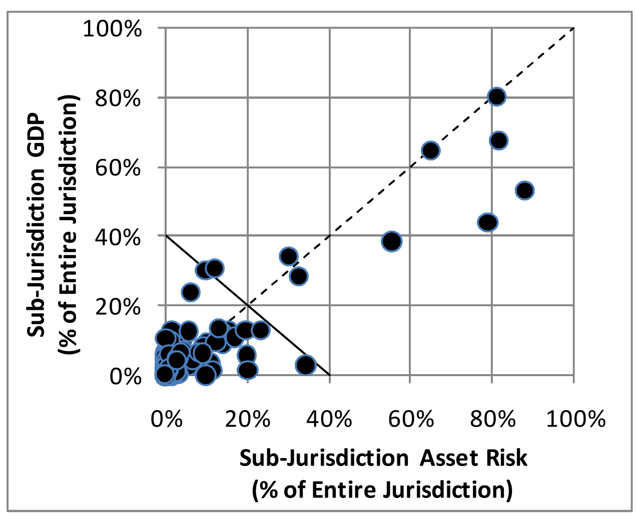

Figure 8: Economic Index vs. Asset Risk Chart

One way to characterize the correspondence of the proxy metric with asset risk is to compute the asymmetry, (proxy – asset risk) / (proxy + asset risk), for all sub-jurisdictions.

Table 2 shows the root mean square (RMS) asymmetry for high-risk (those above and to the right of the solid lines in Figure 5–Figure 8) and low-risk sub-jurisdictions.

Proxy Metric	*Population*	*Population Index*	*Number of Assets*	*Economic Index*
High Risk	25%	11%	28%	21%
Low Risk	66%	63%	64%	72%

Table 2: RMS Asymmetry between Asset Risk and Four Proxy Metrics for High-Risk and Low-Risk Sub-Jurisdictions.

These data show that none of the four metrics is a good proxy for asset risk for low-risk sub-jurisdictions. For high-risk sub-jurisdictions, Population Risk is a fair proxy, and it is the only one of the four metrics that does not systematically under-weight these high-risk sub-jurisdictions.

5. Conclusions

Initial asset risk results from a sample of 13 states and urban areas show some surprising and not-so-surprising commonalities. Asset risk (to terrorism) tends to be very highly concentrated in just a few assets in all jurisdictions. This suggests that a targeted critical infrastructure protection (CIP) program to minimize risk in these highest-risk assets can have a dramatic effect on the overall risk to the jurisdiction.

Of the 18 sectors and 16 terrorism hazards evaluated, risk was never distributed uniformly. Local characteristics caused some sectors to experience higher risk in some areas, and the commercial sector had a universally larger-than average risk. Of the hazards evaluated, VBIED caused the highest aggregate risk, and most attack types contributed very little to a jurisdiction's risk. These two findings suggest that an understanding of a jurisdiction's unique risk profile is essential to allocating attention and resources toward mitigating that risk; trying to focus on all sectors and all attack types will result in a lot of wasted effort.

Asset risk tends to be concentrated in a few sub-jurisdictions, however those sub-jurisdictions are defined by the larger jurisdiction (e.g., county, municipality, ward, etc.). This suggests that allocating resources evenly across all sub-jurisdictions is probably not the most effective way to reduce risks. A prioritized resource allocation strategy based on asset risk is indicated by these results.

Finally, prioritizing on simple risk proxy metrics, such as population, number of assets, or economic activity may have only limited effectiveness, as none of the four proxies evaluated aligned closely with measured asset risk. Of the four metrics evaluated, Population Index was the closest match with asset risk.

As was stated in the Introduction, these results are not meant to be universally applicable to other states and urban areas. The sample of jurisdictions studied was small, not chosen randomly, and not representative. However, the remarkable similarities observed across jurisdictions in this study may suggest that other jurisdictions should not be surprised if their risk profiles share some similar characteristics.

References

[1] FEMA Grant Programs web page, accessed on 3/26/09 at http://www.ojp.usdoj.gov/odp/grants_programs.htm.

[2] U.S. Department of Homeland Security, *National Infrastructure Protection Plan*, 2009.

[3] Henry H. Willis et al., *Estimating Terrorism Risk*, Santa Monica: The RAND Corporation, 2005.

[4] Henry H. Willis et al., *Terrorism Risk Modeling for Intelligence Analysis and Infrastructure Protection*, Santa Monica: The RAND Corporation, 2007.

[5] Walter W. Piegorsch, Susan L. Cutter, and Frank Hardisty, "Benchmark Analysis for Quantifying Urban Vulnerability to Terrorist Incidents," *Risk Analysis, Vol. **27**, No. 6, 1411-1425*, 2007.

[6] U.S. Department of Homeland Security, *Strategic Homeland Infrastructure Risk Assessment*, unpublished.

[7] U.S. Department of Homeland Security, *Calculating Risk for the FY 2009 DHS Preparedness Grants Programs*, unpublished draft, December, 2008.

[8] David C. Daniels et al., "Terrorism Risk Management," in Olivier Pourret, Patrick Naïm, and Bruce Marcot, eds., *Bayesian Networks: A Practical Guide to Applications*, Chichester: John Wiley & Sons, 2008.

[9] U.S. Department of Homeland Security, *National Planning Scenarios*, March 2006.

[10] Michael Chipley et al., *Reference Manual to Mitigate Potential Terrorist Attacks Against Buildings*, FEMA Pub. 426, December 2003.

[11] U.S. Department of Homeland Security, *Risk Management for Special Needs Jurisdictions: Transit Risk Assessment Module (TRAM) Tool Kit*, Version 2.0.

[12] David Daniels, "Understanding FEMA's Grant Risk Formula: A Risk Analyst's Perspective," talk given at the *2009 LANL Risk Symposium*, Santa Fe, April 7-9, 2009.

[13] Christian E. Malagon, Timothy P. McInerney, and Sharon D. Panek, "Gross Domestic Product by Metropolitan Area," *BEA Survey of Current Business Vol. **88**, No. 10, 100-132,* October 2008.