

11-1-2006

How Effective Is Security Screening of Airline Passengers?

Susan E. Martonosi
Harvey Mudd College

Arnold Barnett
Massachusetts Institute of Technology

Recommended Citation

Martonosi, S., A. I. Barnett. 2006. "How effective is security screening of airline passengers?" *Interfaces* 36(6) 545–552.

This Article is brought to you for free and open access by the HMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in All HMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

How Effective Is Security Screening of Airline Passengers?

Susan E. Martonosi

Department of Mathematics, Harvey Mudd College, 301 Platt Boulevard,
Claremont, California 91711, martonosi@math.hmc.edu

Arnold Barnett

Operations Research Center, Sloan School of Management, Massachusetts Institute of Technology, E53-379,
77 Massachusetts Avenue, Cambridge, Massachusetts 02139, abarnett@mit.edu

With a simple mathematical model, we explored the antiterrorist effectiveness of airport passenger prescreening systems. Supporters of these systems often emphasize the need to identify the most suspicious passengers, but they ignore the point that such identification does little good unless dangerous items can actually be detected. Critics often focus on terrorists' ability to probe the system and thereby thwart it, but ignore the possibility that the very act of probing can deter attempts at sabotage that would have succeeded. Using the model to make some preliminary assessments about security policy, we find that an improved baseline level of screening for all passengers might lower the likelihood of attack more than would improved profiling of high-risk passengers.

Key words: industries: transportation, shipping; decision analysis: applications.

History: This paper was refereed.

While some aviation security measures are applied equally to all passengers, others are considered so time consuming that they are restricted to a fraction of air travelers. A prominent example of a selective measure arises at the airport security checkpoint, where some passengers are chosen for more stringent secondary screening beyond the basic screening all passengers undergo. The method for choosing selectees has been changing over time, as has its name: first CAPPs (Computer Assisted Passenger Prescreening System), then CAPPs II, and now Secure Flight. The Transportation Security Administration (TSA) initially hoped to have Secure Flight in place by 2005, but as of February 2006, this had not yet occurred.

Opinions diverge regarding the antiterrorist effectiveness of such passenger prescreening systems. Then Secretary of Transportation Norman Mineta described a strong prescreening system as the "foundation" of aviation security (O'Harrow 2002). Donald Carty, former CEO of American Airlines, saw it as "the right answer to security" and suggested

that such a system would not only foil actual terrorist attempts but would deter terrorists from even attempting an attack on aviation (Kageyama 2002, Jones 2002). Others, however, have questioned the protection offered by prescreening systems (Barnett 2004). One problem that has received special attention is the ability of terrorist groups to reverse engineer the system and thereby thwart it.

As Chakrabarti and Strauss (2002) argued, terrorist groups can find out through test flights which of their members computers select for secondary screening and which they do not, a process the authors liken to a carnival game. Then in actual missions, group members classified as low risk could take the lead roles. In consequence, the true effectiveness of the prescreening system might be far less than hypothesized.

Who is closer to the truth, the optimists or the pessimists? We formulated a simple mathematical model of the prescreening process, one component of which relates to terrorist behavior. The model clarifies the tacit assumptions made by optimists, as well as those made by skeptics, such as Chakrabarti and Strauss. Moreover, the formulation makes transparent some

potential problems with the arguments they have advanced. In some circumstances, for example, terrorists' ability to probe the screening system can actually reduce the danger of a terrorist attack. Even if the prescreening system is extremely good at identifying high-risk individuals, it might only minimally reduce the chance of successful terrorism if the screening they undergo is inadequate.

System and Parameter Description

When a passenger checks in for a flight, the passenger-screening system (PSS) labels her as either low or high risk. If she is considered low risk and is also not selected at random to receive additional scrutiny, she passes through the metal detector and sends her carry-on bags through the x-ray machine, a process we call *primary screening*. Otherwise, she must pass first through primary screening and then undergo a more thorough search of her belongings and clothing. We call the entire alternate process, including the initial primary screening, *secondary screening*.

We define the following parameters:

C : a priori probability that an actual terrorist is classified as high risk by the PSS.

r : proportion of passengers deemed low risk who are selected at random for secondary screening.

p_1 : conditional probability that primary screening detects the terrorist's weapons and prevents her from proceeding further, given that she undergoes only primary screening.

p_2 : conditional probability that secondary screening detects the terrorist's weapons and prevents her from proceeding further, given that she undergoes secondary screening. Because secondary screening includes primary screening, we assume that $p_2 \geq p_1$.

n : number of potential participants in a particular terrorist plot.

τ : terrorist group's deterrence threshold—the minimum probability of success required by the group to proceed with the attack.

For ease of exposition, we assume that a single terrorist attempts an attack by trying to pass through screening. (Two lone terrorists with plastic explosives destroyed two Russian passenger planes in August 2004.) We treat the attack as successful if she makes it through the screening system with her weapons and boards the aircraft (Figure 1):

$$P(\text{terrorist undergoes primary screening}) = (1 - C)(1 - r),$$

$$P(\text{terrorist undergoes secondary screening}) = C + (1 - C)r.$$

Under these assumptions, a successful attack occurs with probability

$$(1 - C)(1 - r)(1 - p_1) + (C + (1 - C)r)(1 - p_2).$$

In calculating their chance of success, we assume that the terrorists use accurate estimates for these

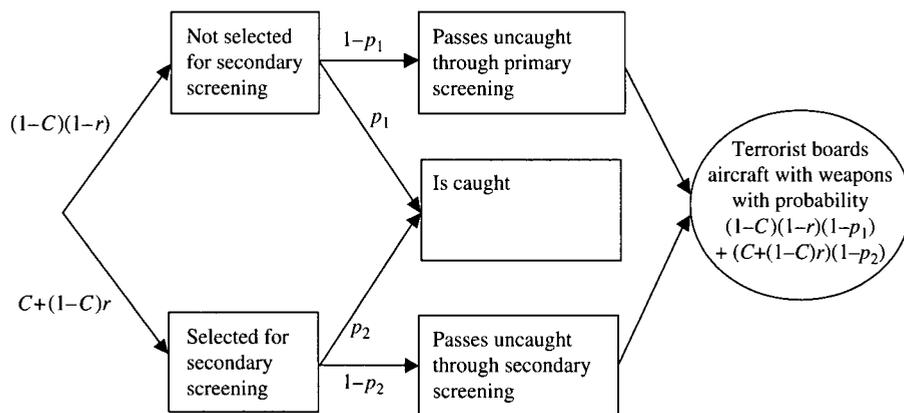


Figure 1: The probability that a terrorist boards an aircraft with weapons is given by a decision tree. If the terrorist is not selected for secondary screening, she must pass uncaught through the primary screening level. If she is selected for secondary screening, then she must pass uncaught through the secondary screening level.

parameter values. If their calculated chance of success is smaller than τ , they will forego their attack; otherwise they will proceed. Low values of τ reflect willingness to attack despite a high risk of failure, while high values indicate risk aversion (detering an on-board attack might only lead to mayhem elsewhere, but that caveat can be applied to any successful antiterrorist measure).

Opposing Viewpoints

People hold various views about whether the PSS can substantially improve aviation security:

The "Right Answer" to Security

Supporters of profiling systems appear to assume that the PSS will flag virtually all terrorists as high risk, or that C is very close to 1. (TSA, for example, has described passengers who undergo only primary screening as "innocent" travelers who "clearly pose no threat of terrorism." That wording implies that any actual terrorist would get secondary screening.) Furthermore, some people, including the Transportation Security Administration (O'Harrow 2002), believe that terrorists can be deterred by the prospect of stringent searching (that is, they have a nonzero deterrence threshold, τ). Under these assumptions, almost all terrorists who would consider an attack on a plane would undergo secondary screening, and the probability that they attempt and succeed at their attack would be either $1 - p_2$ (if they attempt the attack) or zero (if they are deterred).

Loopholes

Others believe, however, that profiling introduces weaknesses to the security system. Barnett (2004) questions the assumption that C would be very high, citing, among other concerns, the difficulties in using limited information about past terrorists to identify future terrorists and the possibility of terrorists using innocent-seeming dupes who do not realize that they are carrying deadly items. Chakrabarti and Strauss (2002) go even further and argue that profiling applied to passenger screening could actually help terrorists improve their chances of success. Terrorist groups can send members on trial flights, without the intent to attack, to ascertain whom the PSS considers low risk. If we assume that passengers' PSS scores

are constant during the probe-and-attack cycle, then a terrorist who passes unflagged even once through the checkpoint knows for certain he has a low PSS score. In the real attack, the group could use this low-risk member knowing that he will face additional scrutiny only if chosen at random.

Chakrabarti and Strauss argue that a terrorist group can always find a low-risk group member. First, they assume that n , the pool of terrorist probes, is arbitrarily large. Second, they assume that $C < 1$ and that risk ratings for different group members are independent: even if five terrorists have been flagged by the PSS, the probability that the sixth will be selected remains C . The probing process is thus Bernoulli, and, if the group keeps sending probes, eventually one will be classified low risk. At the time of the attack, C would be zero for such a known low-risk passenger, and the probability that he would succeed would be $(1 - r)(1 - p_1) + r(1 - p_2)$. While they do not discuss deterrence, they implicitly assume, using our notation, that $\tau < (1 - r)(1 - p_1) + r(1 - p_2)$, so the terrorists will not be deterred. Because the chance of undergoing secondary screening in this situation is simply r , they infer that using a PSS would do no better than random screening alone.

Discussion

These various arguments may be problematic:

Misplaced Optimism?

The optimists may rely on both wishful thinking and incomplete thinking. First, they seem confident that the profiling system will be highly capable in identifying terrorists ($C \approx 1$) but they do not say precisely why they are so reassured. Not all potentially dangerous individuals appear on terrorist watch lists, and data-mining algorithms meant to sort out suspicious behavior need not be especially successful. At the time of the 2002 Washington sniper crisis, historical data mining about serial killers led to the widespread and false belief that the perpetrators were white (Barnett 2004).

Furthermore, supporters of profiling are rarely explicit about how effective they consider the screening imposed on selectees to be. Optimists appear tacitly to assume that secondary screening is highly

effective ($p_2 \approx 1$). However, evidence for that viewpoint is limited. Although several of the 9/11 terrorists were subjected to additional scrutiny, not one was stopped from boarding the aircraft. Secondary screening at the time did not search for (let alone forbid) the box cutters that were apparently instrumental to the plot.

Secondary screening today, of course, is more demanding than that on 9/11. But its effectiveness has been sharply questioned. A 2004 report of the then inspector general to the Department of Homeland Security, Clark Kent Ervin, was not encouraging. Undercover tests conducted in 2003 revealed weaknesses in employee training; screening equipment and technology; security policy and procedures; and management and supervision (Office of Inspector General 2004). According to the report, the passenger-screening process in place at 15 airports repeatedly failed to detect weapons and explosive materials. The main screening device is a metal detector, but as Representative Peter DeFazio of the House Aviation Subcommittee noted, "You're not going to find plastic explosives with a metal detector, no matter how hard you try." Representative John Mica, chairman of the subcommittee, declared the results to be "bad enough" for general screening and "absolutely horrendous" with respect to detecting explosives (Hall 2004). For certain types of plots, the true value of p_2 might be far below one. In that circumstance, directing the terrorists to secondary screening might be a hollow victory.

The optimists also seem indifferent to p_1 , the detection rate for primary screening. But if $C < 1$, we cannot ignore this parameter. Terrorists could use innocent passengers in an attack without their knowledge or find a low-risk terrorist through probing. Thus, if C is low, any inadequacies of primary screening may substantially raise the chance of a successful attack.

Another point is virtually ignored by the optimists: if $p_1 \approx p_2 = p$, then what value could profiling, or even randomized selection, add? Regardless of whether the common detection rate p is high or low, the PSS is irrelevant. Only when $p_2 - p_1$ is fairly large does profiling appreciably decrease the terrorists' success rate. (Currently, secondary screening consists of the primary screening process plus additional steps that can include a pat-down or wand search of the passenger and a search of carry-on bags. The quantity $p_2 - p_1$

can vary depending on the type of weapons involved, where they are stored, and the nature of secondary screening used.)

In short, the belief that prescreening systems would identify virtually all terrorists ($C \approx 1$) has, to say the least, not been demonstrated. Even if C were close to one, that circumstance might be insufficient to thwart terrorism. Because the optimists do not consider directly quantities p_1 , p_2 , or $p_2 - p_1$, they miss the point that identifying high-risk people is beneficial only if that capability reduces the chance that an attack would succeed. Even the deterrent effects of effective screening—which optimists sometimes mention—depend on several parameters and not just C .

Misplaced Pessimism?

Critics of prescreening systems may likewise be overly pessimistic. They may be correct that C , p_1 , and p_2 are far below one. But a germane question is whether the terrorist's chance of being thwarted at the airport is high enough to lead him to stay away. Given the wide availability of easy targets on the ground, the deterrence threshold τ might be quite high, meaning that terrorists would find a surprisingly broad range of combinations of C , p_1 , and p_2 unacceptable. That range could include parameter values that can be (and perhaps already are) achieved.

Moreover, Chakrabarti and Strauss's argument, although interesting, might not be as damaging to PSS as one might initially think. In some instances, probing can actually reduce the chance that a terrorist group will successfully attack an airplane. Chakrabarti and Strauss argue that a terrorist group can always find a low-profile member to carry out the attack. They assume an arbitrarily large group of independent terrorists. But are the groups really arbitrarily large? The assumption that group members are each classified independently as high or low risk suggests that US intelligence services are unable to recognize connections between dangerous individuals. This could be true, but models that assume this weakness at the outset—and thus assume independence in probabilistic calculations—could substantially overestimate the chance of finding a group member who gets only primary screening. Perhaps most important, however, is the point that terrorists might cancel

their plans if they get evidence of an unacceptable probability of success. Probing might provide that evidence.

Consider a lone terrorist, and suppose that he has made the following (accurate) estimates about the system:

—Primary screening is 20 percent effective at detecting his plot ($p_1 = 0.20$).

—Secondary screening is 85 percent effective at detecting his plot ($p_2 = 0.85$).

—No passengers are selected at random for additional screening ($r = 0$). (This is only for simplicity. Similar examples can be constructed with non zero values of r .)

—The PSS has a 60 percent chance of selecting the terrorist for secondary screening ($C = 0.60$).

Assume that the terrorist will not attack unless he has at least a 25 percent chance of success ($\tau = 0.25$).

If he does not probe the system, he first compares his estimated probability of success with τ : With 60 percent probability, he would undergo secondary screening and thus would have a 15 percent chance of success. Otherwise, he would pass through primary screening and have an 80 percent chance of success. His total chance of success is then $(0.60) * (0.15) + (0.40) * (0.80) = 41$ percent, which is larger than τ . He decides to attack, and as per his calculations, succeeds with probability 41 percent.

However, if he does probe, he waits to decide whether or not to attack until after the probing run. During this trial there are two possibilities:

(1) With 60 percent probability, he is selected for secondary screening, knows he is classified as high risk, and updates his estimate of C to 1. (Because $r = 0$ in this example, he cannot attribute becoming a selectee to bad luck.) Knowing that if he were to attack, he would face secondary screening and succeed with only 15 percent probability, which is less than τ (25 percent), he gives up without attacking.

(2) With 40 percent probability, he avoids secondary screening during his trial run and knows for certain that he is considered low risk (he updates C to 0). If he were to attack, his probability of success under primary screening would be 80 percent.

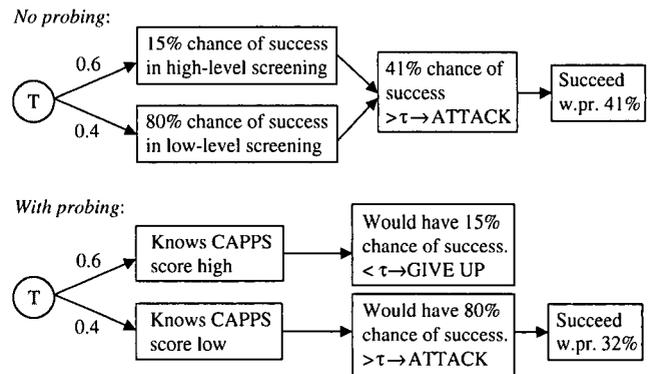


Figure 2: We compare the terrorist's chance of success if he does not probe the system before attacking to that if he does probe ($p_1 = 0.2$, $p_2 = 0.85$, $r = 0$, $C = 0.6$, and $\tau = 0.25$). The total probability of a successful attack is only 32 percent if the terrorist probes compared to 41 percent if he does not. Thus, probing the system can sometimes discourage rather than reassure the terrorist.

Because this is higher than his deterrence level τ , he decides to attack.

Because he attempts the attack only if he is deemed low risk during the probing run, which happens with 40 percent probability, and would subsequently succeed with 80 percent probability, his overall chance of success if he probes first is only 32 percent, which is lower than if he had not probed (Figure 2).

What causes this outcome? Terrorists probe to gain information about the system. If the information they get is discouraging, then they might cancel an attack that otherwise would have had an appreciable chance of success. Unless one believes that terrorists cannot be deterred—and their meticulous preparation often suggests far more rationality than we would like—analyses that treat deterrence as a negligible phenomenon might be too pessimistic. (While we considered here only a single terrorist, similar examples can be constructed for multiple terrorists.)

Some Explorations with the Model

As we have suggested, a simple model indicates problems with certain arguments about security screening that have been advanced. But the model can also help raise some general points about screening strategy, even if tentatively.

We can rewrite the equations above to specify that Q , the probability that a terrorist who actively

attempts to board an airplane is foiled by airport security, as follows:

$$Q = p_1 + (C + (1 - C)r)\varepsilon,$$

where $\varepsilon = p_2 - p_1$, and the other symbols are as defined earlier.

Given that r is small (and is likely to stay so, given the aversion to extra random screenings), we can approximate the above relationship as

$$Q = p_1 + C\varepsilon.$$

It is of interest to take the partial derivative of Q with respect to each of its three components, which reveals what would happen if we changed that component slightly while holding the other two constant. It is easy to imagine that p_1 could go up while C stays constant; for p_1 to increase while ε remains the same, however, p_2 would have to increase to the same extent as p_1 . This last circumstance could well arise because, as noted, primary screening is included in secondary screening, meaning that effectiveness gains in the former might often translate into equal gains in the latter. Likewise, ε could rise while p_1 stays the same because of improvements in those aspects of secondary screening that go beyond primary screening.

The partial derivative of Q with respect to p_1 is 1, with respect to C is ε , and with respect to ε is C . Given that C and ε are clearly below one, this outcome means that a fixed increase in one of the parameters while the others stay constant is most beneficial when it accrues to p_1 . Because everyone goes through primary screening, it is a case of "a rising tide lifts all boats." An example of a measure that would increase p_1 would be replacing the present metal detectors at security checkpoints by others that are sensitive to plastic explosives as well as metals.

One might think that, if it were possible to increase two of the parameters (p_1 , C , ε) by the same modest amount while holding the third constant, one should choose C and ε , given their product would benefit appreciably because both of its factors went up. But for small changes in the parameters above, the change ΔQ in the direct success rate of screening would follow:

$$\Delta Q \approx \Delta p_1 + C\Delta\varepsilon + \varepsilon\Delta C.$$

This formula implies that p_1 would always be among the two parameters selected for improvement (again, because C and ε are below one). The other parameter chosen would be the smaller of C and ε . (Of course, achieving gains might be easier for some parameters than for others.)

There is also, however, a multiplier effect for any improvement in Q , tied to the deterrence threshold, τ . If $1 - Q$ —the initial probability of terrorist success given an attempt—exceeds τ , but $1 - (Q + \Delta Q)$ does not, then no attempt will be made and the chance of success drops to zero. This possibility means that there is an extra effect associated with improving checkpoint security, in which a small improvement can sometimes go a long way.

Numerical Example

As we have emphasized, neither we nor anyone else knows the key parameters in even a simple security model. Here we make what might charitably be called mildly educated cases to provide the basis of some calculations.

The conditional probability C that a terrorist will be selected for secondary screening was estimated to be in the range 90 to 96 percent by Virta et al. (2003). But on 9/11, only six of 19 hijackers were flagged by CAPPS, for a success rate of 32 percent (Eggen 2002); three others were selected because of problematic travel documents. Even 32 percent might be an artificially high estimate because the 9/11 terrorists had conducted trial runs and knew that secondary screening, which was then weak at security checkpoints, posed no threat to their plans. (Had they discovered otherwise, they might have changed their methods to reduce the likelihood of being selected.) However, passenger profiling today might be more accurate than it was on 9/11 given all the recent efforts to improve it. All things considered, $C = 50$ percent seems a plausible first guess about the effectiveness of identifying terrorists in the foreseeable future.

We also do not know p_1 or p_2 terribly well. Reports by the US Government Accountability Office suggest that screeners are about 80 percent likely to detect certain unauthorized items, although it is unclear how these discoveries are divided between primary and secondary screenings (Lipton 2005, Office of Inspector

General 2005, General Accounting Office 2000, Government Accountability Office 2005). In any case, this statistic applies only to some weapons, and significantly, it does not consider plastic explosives. On the other hand, new technology now being tested, such as a walk-through “sniffer” that blows a puff of air on the passenger to dislodge trace particles of explosives, might tangibly improve weapon-detection rates. Any estimates are speculative, but perhaps $p_1 = 20$ percent and $p_2 = 50$ percent are defensible choices (in which case $\varepsilon = 30$ percent).

The deterrence threshold τ is the hardest parameter to estimate, given that terrorists will hardly take part in public surveys. Moreover, the threshold presumably varies from group to group. A Rand report argued that fear of ignominious failure does weigh heavily on terrorist groups, so it is unlikely that τ is close to zero (Davis and Jenkins 2002). At the same time, many of these groups show too much sophistication to believe that success is guaranteed, so they presumably do not insist on a τ close to 100 percent. A distribution for τ seems more palatable than a point estimate: Here we approximate τ as uniformly distributed between 25 percent and 75 percent.

For $C = 0.5$, $p_1 = 0.2$, $\varepsilon = 0.3$, and $r = 0$, we find from the equation above that $Q = 0.35$. To increase Q by 0.05, we must achieve a combination of Δp_1 , ΔC , and $\Delta \varepsilon$ that falls in the plane $\Delta p_1 + 0.3\Delta C + 0.5\Delta \varepsilon = 0.05$. Among the qualifying points are $(0.05, 0, 0)$ and $(0, 0.083, 0.05)$, reminding us that increases in p_1 have greater leverage than the other quantities.

Of course, there is a chance that raising Q from 0.35 to 0.40 will push the terrorists' chance of success from above τ to below it. That will happen if τ is between 0.60 and 0.65; if τ is assumed $U(0.25, 0.75)$, the probability that τ falls in the relevant range is 0.10. The overall probability of a successful attack was originally $(0.8) * (0.65) = 0.52$, based on only a 20 percent chance of deterrence at $Q = 0.35$. The revised probability is $(0.7) * (0.6) = 0.42$, 10 percentage points lower. In this example, deterrence acts to double the effect of a five percentage-point gain in the conditional probability of thwarting an actual attempt against a plane. We suspect that this outcome hints at a broader truth: because of the possibility of deterrence, the total benefit of improved screening might be considerably larger than the immediate benefit at the security checkpoint.

Conclusions

As this simple mathematical model suggests, neither side has made a persuasive case about the effectiveness of airport passenger-profiling systems. Supporters of such systems have focused mostly on the ability of the algorithm to identify terrorists (C), an ability they may well overestimate. They say little about screening effectiveness of both low-risk passengers and selectees (p_1 and p_2), yet these effectiveness parameters are crucial to the overall success rate of the system. Skeptics may have given insufficient weight to deterrence (τ), because of which, the selection and screening system might prevent attacks even though it falls well short of perfect. Probing the system, as we have seen, could sometimes prevent a terrorist act rather than ensure its success.

The mathematical model offers some preliminary indications about improved checkpoint security. Improving the baseline screening applied to all passengers by a fixed amount might be more beneficial than improving either the terrorist detection probability or the marginal benefit of secondary screening by the same amount. Deterrence might appreciably magnify the benefits of even modest gains in the chance of thwarting a terrorist at the security checkpoint. Not knowing the key parameters, however, we can hardly offer these last thoughts as policy recommendations.

Straightforward mathematical reasoning can often clarify the underlying reasoning behind qualitative statements about aviation security and point out both the limitations in such reasoning and the unknown parameters it would be most useful to learn. Problems as important as those arising in Homeland Security could benefit from a wider use of mathematical modeling.

References

- Barnett, A. I. 2004. CAPPS II: The foundation of aviation security? *Risk Anal.* 24(4) 909–916. (Also, response to discussants, 933–934.)
- Chakrabarti, S., A. Strauss. 2002. Carnival booth: An algorithm for defeating the computer-assisted passenger screening system. *First Monday* 7(10). Retrieved October 2002 http://www.firstmonday.org/issues/issue7_10/chakrabarti/index.html.
- Davis, P. K., B. M. Jenkins. 2002. Deterrence and influence in counterterrorism: A component in the war on al Qaeda. Monograph Report MR-1619, Rand Corporation, Santa Monica, CA.
- Eggen, D. 2002. Airports screened nine of Sept. 11 hijackers, Officials say kin of victims call for inquiry into revelation. *Washington Post* (March 2) A11.

- Hall, M. 2004. Airport screeners missed weapons. *USA Today* (September 23). Retrieved September 23, 2004 http://www.usatoday.com/news/nation/2004-09-22-weapons_x.htm.
- Jones, J. 2002. Airlines to test new security measures—American CEO. *Reuters* (September 6).
- Kageyama, Y. 2002. American Airlines CEO urges some airport security measures be dropped. *Associated Press* (May 31). Retrieved October 2002 web.lexis-nexis.com.
- Lipton, E. 2005. Transportation Security Agency criticized. *The New York Times* (April 19) A18.
- Office of Inspector General, United States Department of Homeland Security. 2004. Audit of passenger and baggage screening procedures at domestic airports. Office of Audits, OIG-04-37.
- Office of Inspector General, United States Department of Homeland Security. 2005. Follow-up audit of passenger and luggage screening. Office of Audits, OIG-05-16.
- O'Harrow, R., Jr. 2002. Air security focusing on flier screening: Complex profiling system months behind schedule. *The Washington Post* (September 4) A1.
- United States General Accounting Office. 2000. Aviation security: Vulnerabilities still exist in the aviation security system. Testimony of Gerald L. Dillingham before Committee on Commerce, Science, and Transportation Subcommittee on Aviation, United States Senate, GAO/T-RCED/AIMD-00-142.
- United States Government Accountability Office. 2005. Aviation security: Secure flight development and testing underway, but risks should be managed as system is further developed. Report to Congressional committees, GAO-05-356.
- Virta, J. L., S. H. Jacobson, J. E. Kobza. 2003. Analyzing the cost of screening selectee and non-selectee baggage. *Risk Anal.* 23(5) 897-908.

Copyright 2006, by INFORMS, all rights reserved. Copyright of Interfaces is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.