Optimal Defensive Investments for Deterring Terrorists

Timothy Sweda, Daniel Walton, Eugene Quan, and Susan Martonosi

Background:
Ever since the terrorist attacks of 9/11, policy makers have shown an increased desire to understand how to deter terrorists and how to prevent future attacks. Due to limited security resources, it is not always feasible to fully defend a target that is susceptible to attack by terrorists, so it is important to utilize the resources as effectively as possible. In our research we characterize optimal defensive investments and examine how they are influenced b

Approach:
We study a game-theoretic model in which a defender must defend a single target against an attacker. Whereas previous models characterize the damage done to the target by the attacker discretely (i.e. the target is either destroyed or not destroyed), we model the damage as a continuous random variable that follows a beta distribution. The attacker is rewarded for the proportion of the target that he destroys, and punished for the complementary proportion that he does not destroy.

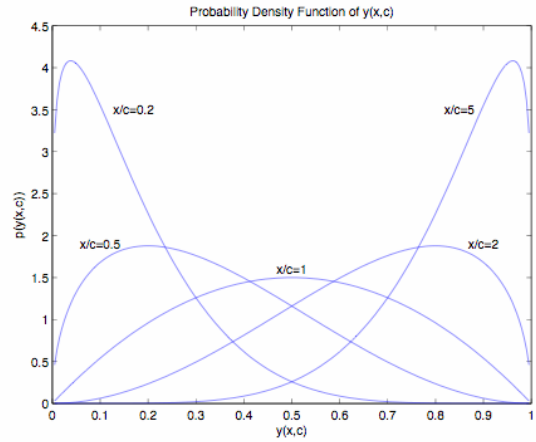

In this graph, y(x,c) is the proportion of the target that the attacker destroys, and x/c is the ratio of the attacker's investment into attacking the target to the defender's investment into defending the target.

Results:
We determine the attacker's optimal strategy, deriving closed-form expressions for when the attacker should attack and, if he does attack, how much he should invest. Using these results, we characterize the defender's optimal investment and find that the defender should never leave the target undefended.
As the defender's valuation of the target decreases, or as the expected attacker's reward becomes sufficiently small or sufficiently large, we find that the defender's optimal investment decreases. Additionally, we observe that the defender's optimal investment usually decreases as the punishment increases, but in some cases when the punishment is sufficiently small, increasing the punishment will actually cause the defender's optimal investment to *increase* as well.

Defender's Optimal Investment as a Function of…

Expected Attacker's Reward^{-1}:

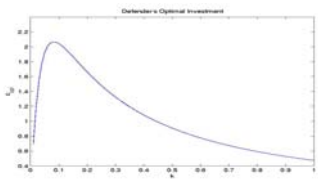

Punishment:

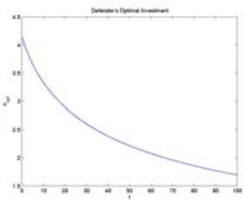

OR

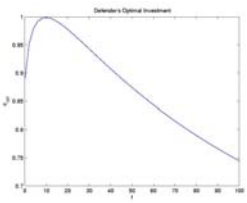

Future Work:
Short-term goals include characterizing the defender's optimal investment more thoroughly and extending the model to include two targets. We provide the framework for a two-target model and briefly examine the attacker's optimal strategy, and we hope to later build on that foundation. Ultimately, we would like our model to have more than two targets, and we may want to consider varying the amount and quality of information available to both the attacker and defender.

Work funded by the Beckman Faculty Research Fund and the Baker Foundation