Chapter 6

SEARCH FOR A MALEVOLENT NEEDLE IN A BENIGN HAYSTACK

D. P. Gaver, K. D. Glazebrook, and P. A. Jacobs

Abstract: A domain contains a number *w* of non-hostile White (W) individuals: humans, vehicles, ships. A hostile Red (R) individual enters the domain and travels through the domain towards targets. If R reaches an attractive valuable target, perhaps a crowd of people on land or a ship at sea containing liquid natural gas (LNG), it attacks the target. A Blue counter-terrorist, C, patrols the domain and classifies (perhaps incorrectly) individuals of interest as R or W. The probability of correct classification is an increasing function of the time spent classifying an individual. The misclassification of a W as an R is a false positive; misclassification of the R as a W is a false negative. C follows (or tracks) any individual it classifies as R until it is relieved by another platform or individual that may neutralize the possible R. C is unable to detect and classify additional individuals while it is following a suspicious individual. A small classification time may yield many false positives that C must service. A large classification time may result in R achieving its goal before being neutralized, so an appropriate compromise is sought.

A game-theoretic model is formulated and studied to evaluate the probability that R is successfully neutralized before achieving its goal. C's policy is to choose a classification time. Targets have independent identically distributed (iid) values, and R's policy is to specify a target value threshold; R will attack the first target it finds whose value exceeds the threshold unless neutralized first.

Key words: maritime domain awareness and protection; port security; suicide bomber; Nash equilibrium

1. INTRODUCTION AND FORMULATIONS

Consider an *arena* (Kress, 2005) or *domain*, D (Gaver et al., 2006; Gaver et al., to appear), i.e., a spatial region, in which a collection of benign/harmless entities, called Whites (Ws) gather. These can be passengers and accompanying individuals in a boarding lounge awaiting an airline flight, bus, or other conveyance. Alternatively, the collection is a crowd of potential workers assembled to obtain jobs. Another example could be a group of military platforms parked closely outside a repair facility, or the collection of neutral vessels used for recreation, commerce or defense occupying a maritime domain neighboring a harbor; see Gaver et al., 2006; and Gaver et al., to appear). Still another could be a large queue outside a medical emergency care clinic or hospital; the queue is the result of a surge of injuries, possibly from natural causes, such as Katrina, (for instance the refugees being packed into the New Orleans Coliseum). There are many equally important examples.

Such a collection-crowd-"haystack" is the natural target for a single suicide bomber, denoted R, or for a malevolently-inclined "Red" maritime domain invader, (e.g., a small boat), likely disguised to resemble benign Ws in size, appearance, and behavior. Presumably it will attempt to surreptitiously mingle with the White crowd and pick a moment for detonation that creates a maximum number of casualties. This suggests that R should seek to be in the part of D where most, or most-apparently-valuable Ws congregate. On the other hand, the effect of a suicide bomb in a densely packed region could be counterproductive: Ws near the explosion will likely be hit/killed, but their bodies shield those slightly further away; see Kress (2005).

Now introduce (at $t=0$, initiation of the process) a single friendly searcher-neutralizer into the domain D; denote such a "Blue" counter-terrorist by C. In fact, there could be several such, and they need not be identical, but complementary; for example, on land, one or more areas could be equipped with active and sensitive trained dogs, which would be visible and tend to herd the R—along with some Ws—into a subdomain $D' \subset D$ that can be quickly searched by C; on the other hand if such herding occurred it might trigger early detonation by R. We omit such plausible modeling options for the present.

The scenario is that C searches/detects/travels to an individual. The search/detect/travel times to successive individuals are independent and identically distributed (iid). The individual is then classified as either a W or an R; classification takes a time τ, a decision variable. The probability of a correct classification is an increasing function of τ. If the individual is classified as an R (perhaps incorrectly), the C follows/tracks the individual until it is relieved by another platform or individual. C then resumes searching. The successive follow/tracking times are iid, and C cannot detect/classify other individuals while it is following a suspicious individual. There is a constant number of Ws in the domain; these could be individuals that are unclassified by other means. If the classification time τ is too small, then C may be prone to misclassify Ws as Rs and may spend substantial time following benign individuals, increasing the possibility that R reaches its objective. If the classification time is too large, then R may reach its objective before C detects and correctly classifies it.

We make the following assumptions:

1. There is only one R.
2. There is a constant number of Ws in the domain and detected individuals are obtained by sampling with replacement from the population of Ws and R. C does not retain or use information about individuals who were closely examined, classified, and then released. Once a W is classified as a W, it is not tagged, and may be classified again; correct tagging can be a great advantage, but incorrect tagging can be enormously counter productive. We postpone the discussion of tradeoffs. Alternatively, a classified W is replaced by another W entering the domain.
3. The W's are a subset of the entire population. The remainder of the population is classified as harmless as a result of a first cursory surveillance; we assume that this first cursory surveillance is perfect. The Ws are members of the population that are subject to further surveillance.

Section 2 presents an optimization model where the goal is the maximization of the probability that C successfully neutralizes R. Examples are presented illustrating the tradeoff between a large classification time τ resulting in a larger probability of correct classification but also a smaller effective search rate. Section 3 presents a game that elaborates the model for R. Possible targets for R have iid values. R encounters targets according to a Poisson process. R kills the first target it encounters whose value exceeds a predetermined threshold, unless C neutralizes R first. C chooses a classification time τ to minimize the maximum value of the expected killed target value, while R chooses a threshold value to maximize the minimum expected killed target value. Section 4 concludes the paper.

2. STOCHASTIC MODEL 1

In this model, R detonates after an exponential time, presumably when near a valuable target. The model of Section 3 makes this more explicit. Model parameters and random variables appear below.

$\delta_s dt$ = the probability that R detonates in $(t, t+dt)$, given that it has not done so before t, and has not been neutralized by C. Choice of detonation to achieve a specified value/effect is considered later in Section 3. Note: We could let $\delta_s = \delta_s(x)$, where x stands for various conditions, such as elapsed time or environmental conditions; omitted for present.

$\boldsymbol{D}$ = time to detonate if not neutralized. We let $\boldsymbol{D}$ be random having an exponential distribution:

$$P\{\boldsymbol{D} > t \mid \text{Not neutralized}\} = \exp[-\delta_s t]. \tag{2.1}$$

$\boldsymbol{T}_c$ is a random time from when C has finished processing (classifying and perhaps tracking) a detected individual until C detects and travels to another member of the Ws or the single R in D; $\boldsymbol{T}_c$ will depend on the size of the domain, the environment, the number of *W*s in the domain, C's speed, the search pattern, etc. Successive detection/travel times are iid.

$k_r(\bullet) = \dfrac{1}{w+1}$ (respectively, $k_w(\bullet) = \dfrac{w}{w+1}$) is the probability that the detected individual is the R (respectively, a benign W).

Let $p_{rr}(\tau)$ be the conditional probability that the detected R is classified as R when τ time units are spent classifying it; $p_{rr}(\tau)$ is an increasing function of τ where τ is a decision parameter. The probability that the detected R is misclassified is $p_{rw}(\tau) = 1 - p_{rr}(\tau)$. Similarly, let $p_{ww}(\tau)$ be the conditional probability that a detected W is classified as W when τ time units are spent classifying it; $p_{wr}(\tau) = 1 - p_{ww}(\tau)$ is the probability that a detected W is misclassified as R (a false positive). The response to a false positive is that the counter-terrorist, C, takes intensive (but futile) action to thwart a potential suicide attack (e.g., by neutralizing or disabling the misidentified attacker, thus unnecessarily losing valuable search time).

Let $\boldsymbol{X}_c$ be a random service or response time to thoroughly examine and neutralize a captured true R (provided that it is disabled before detonation), *or* to examine and release a false positive W. Successive response times are iid. The response time is in addition to the classification time. Of course, the response time is a penalty time to C if the examination/neutralization is of a W that was incorrectly classified as R. During this time, C is occupied by following and neutralizing the potential R and cannot search for other possible Rs. Neutralization involves tracking the suspicious individual until escort personnel arrive, after which the suspicious individual is escorted to another position for further examination. If R detects unusual attention by something it interprets as C, it may well detonate shortly before apprehension by C; however, R might also choose to leave D in that case, or attempt to sidle close to C and then detonate, thus achieving an extra bonus. We do not treat such concept of operations on the part of R here, but intend to do so in later work.

$\boldsymbol{I}_{cr}(\tau)$= the total elapsed clock/calendar time to detect, identify, and neutralize the (true) R when the classification time is τ. For simplicity, this includes the (random) time "used up" in discovering that potential Rs are actually benign, (i.e., that they are instead harmless Ws). We presume that neutralization of a suspect is not terminal; the suspect is fully disabled and searched, possibly off-site. In some cases, this will be wasted time, of course. A serious issue to avoid would be serious harm to an individual that turns out to be an innocent bystander.

Conditional on $\boldsymbol{I}_{cr}(\tau)$, the time to neutralize R, the probability that R is neutralized in time (before detonation) is $\exp\left[-\delta_s \boldsymbol{I}_{cr}(\tau)\right]$. This event (pre-emptive neutralization) can take place under two conditions:

(*i*) the first time C detects a suspicious individual, that individual is R, and is correctly identified, and thoroughly neutralized; or

(*ii*) the first time C detects a suspicious individual, that individual is R and that individual is misclassified and released, *or* that individual is a W and is correctly classified and released. At this point *the search begins over from scratch, as in renewal theory*.

Note (again) that the present model is (conservatively) *memoryless*: there is no tagging or labeling of those suspicious individuals that are released. Correct tagging would increase the rate at which R could be neutralized, and is a strong candidate for future study. However, high rates of incorrect tagging could lead to severe degradation of C's operational success.

Let the unconditional probability that R is detected and neutralized before it detonates be given by

$$P_B(\tau) = E\left[e^{-\delta_s \boldsymbol{I}_{cr}(\tau)}\right]. \tag{2.2}$$

This is the success probability for C when τ time units are spent classifying a suspicious individual. Conditionally, using (*i*) and (*ii*), we have

$$
\begin{aligned}
P_B(\tau) = {} & E\left[\underbrace{\exp\{-\delta_s \boldsymbol{T}_c - \delta_s \tau\} k_r(\bullet)\, p_{rr}(\tau) \exp\{-\delta_s \boldsymbol{X}_c\}}_{\text{R is found and escorted from the domain D at first contact}}\right] && \text{(a*)} \\
& + E\left[\underbrace{\exp\{-\delta_s \boldsymbol{T}_c - \delta_s \tau\} k_w(\bullet)\, p_{ww}(\tau)}_{\text{A benign W is intercepted and released}}\right] \underbrace{P_B(\tau)}_{\text{Searcher starts over}} && \text{(b*)} \quad (2.3) \\
& + E\left[\underbrace{\exp\{-\delta_s \boldsymbol{T}_c - \delta_s \tau\} k_w(\bullet)\, p_{wr}(\tau) \exp\{-\delta_s \boldsymbol{X}_c\}}_{\text{A benign W is mistaken for R and escorted from the domain}}\right] \underbrace{P_B(\tau)}_{\text{Searcher starts over}} && \text{(c*)} \\
& + E\left[\underbrace{\exp\{-\delta_s \boldsymbol{T}_c - \delta_s \tau\} k_r(\bullet)\, p_{rw}(\tau)}_{\text{R is misidentified and released}}\right] \underbrace{P_B(\tau)}_{\text{Searcher starts over}} . && \text{(d*)}
\end{aligned}
$$

Solving gives C's success probability:

$$
P_B(\tau) = \frac{E\left[\exp\{-\delta_s \boldsymbol{T}_c - \delta_s \tau\} k_r(\bullet)\, p_{rr}(\tau) \exp\{-\delta_s \boldsymbol{X}_c\}\right]}{D(\boldsymbol{T}_c, \boldsymbol{X}_c)} \qquad \text{(2.4,a)}
$$

where

$$
\begin{aligned}
D(\boldsymbol{T}_c, \boldsymbol{X}_c) &= 1 - E\left[\exp\{-\delta_s \boldsymbol{T}_c - \delta_s \tau\} k_w(\bullet)\, p_{ww}(\tau)\right] \\
&\quad - E\left[\exp\{-\delta_s \boldsymbol{T}_c - \delta_s \tau\} k_w(\bullet)\, p_{wr}(\tau) \exp\{-\delta_s \boldsymbol{X}_c\}\right] \qquad \text{(2.4,b)} \\
&\quad - E\left[\exp\{-\delta_s \boldsymbol{T}_c - \delta_s \tau\} k_r(\bullet)\, p_{rw}(\tau)\right].
\end{aligned}
$$

Note that the term (c*) in equation (2.3) could be modeled in several alternative ways. First, W can be incorrectly classified as R in which case C tracks the W for a relatively short time $\boldsymbol{X}_c$ until relieved by an auxiliary platform with a sensor, (e.g. an unmanned aerial vehicle (UAV)). In this case, C may simply cease searching, or move to another area, so the search may never re-start; in this case, R will never be detected. Alternatively, W can be incorrectly classified as R, and $\boldsymbol{X}_c$ represents the relatively long time for an escort team to arrive and assume responsibility for the target. Then the C may begin searching again if at the termination the target is identified as an innocuous W. Variations in equation (2.3) can therefore be justified (e.g., if Blue search incorrectly terminates), but are not included here.

EXAMPLE

All numerical examples will use the following maritime domain protection scenario. An aircraft, C, is patrolling a rectangular domain; M_x is the *x*-distance of the rectangular domain; M_y is the *y*-distance of the domain. R is a small boat carrying explosives. C uses a sensor to search the domain. The footprint of the sensor is a square with length of a side equal to f. The total time for C to cover a one-footprint square is f/v_s; v_s is the velocity of the aircraft. The mean time for C to cover the domain is $\left[M_x \times M_y / f^2\right]\frac{f}{v_s}$; the mean time between detection of individuals is $\left[\left[M_x \times M_y / f^2\right]\frac{f}{v_s}\right]/\left[w+1\right]$, where w is the (constant) number of benign Ws in the domain. When C classifies a vessel as R, it tracks the vessel until it is relieved by another platform.

For illustration, let

$$p_{ww}(\tau) = p_{rr}(\tau) = \exp\left\{-\alpha / \tau^{\beta}\right\} \tag{2.5}$$

for $\tau > \tau_{0.5}$ where $\tau_{0.5}$ is the time at which the probability of correct classification is 0.5. Note that $p_{ww}(\tau)$ is a cumulative Fréchet distribution function (which optimistically approaches unity for τ increasing). The parameters α and β (both positive) are determined here by specifying two quantiles (in this case the median, $\tau_{0.5}$, and the 90th percentile, $\tau_{0.9}$). If data were available, α and β could be statistically estimated using maximum likelihood or Bayes inference. Since $p_{ww}(\tau) = p_{rr}(\tau)$, we are assuming that R is attempting to blend in with Ws in the domain (although there are of course other possibilities). Table 6-1 displays some parameter values.

Figure 6-1 displays the probability that R is neutralized before it detonates as a function of the time spent classifying a detected individual, when the search/detect travel times $\boldsymbol{T_c}$ and the additional time the sensor is engaged for response to an individual classified as R (service/response time) $\boldsymbol{X_c}$ are both gamma distributed with shape parameter 0.1. The mean response time is 0.2 hours. For comparison, if there are no benign vessels, and classification takes no time and is perfect, then the probability of neutralizing R is 0.82.

Table 6-1. Parameter values for Model 1

$E[\boldsymbol{T}_c]$	Mean time to search/detect a suspicious individual (hrs.)	Variable
$p_{rr}(\tau) = p_{ww}(\tau)$	Probability of correct classification	$\tau_{0.5} = 3/60$ hrs. $\tau_{0.9} = 6/60$ hrs.
β	Parameter ("shape") for the probability of correct classification	2.72
α	Parameter ("scale") for the probability of correct classification	0.0002
w	Number of benign Ws in the domain	Variable
$k_r = 1 - k_w$	Probability detected individual is the R	$1/(1+w)$
$E[\boldsymbol{X}_c]$	Mean time for second more intensive inspection (response time) (hrs.)	Variable
$1/\delta_s$	Mean time until the R detonates if it is not neutralized (hrs.)	6
f	Side of square sensor footprint in nautical miles (nm)	10
v_s	Velocity of sensor platform in knots (kts)	250
M_x	Length of x-direction of rectangular domain (nm)	75
M_y	Length of y-direction of rectangular domain (nm)	100

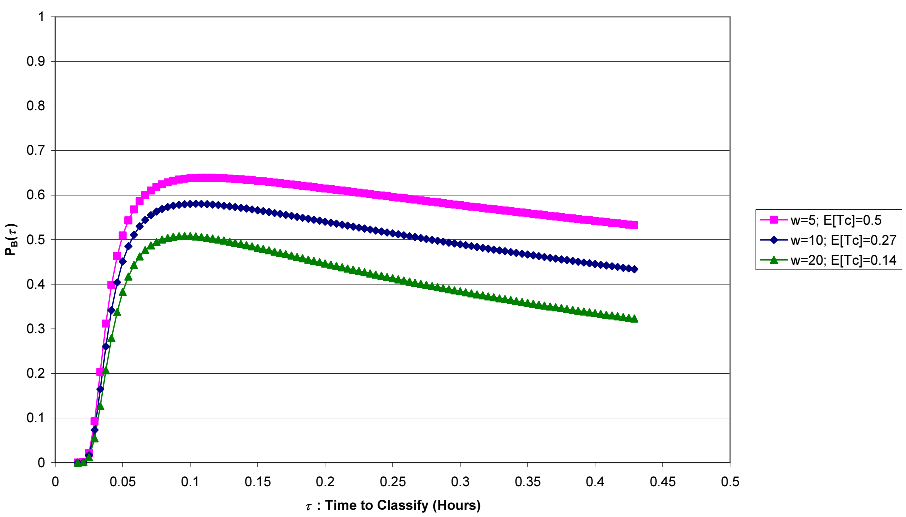

Figure 6-1. Probability R is neutralized in Model 1.

DISCUSSION

The optimal classification time decreases a little as the number of benign Ws in the domain increases. However, even if the estimate of the number of Ws is wrong, the optimal classification time still seems to perform reasonably well. The probability of neutralizing R is sensitive to the number of benign Ws. Fewer benign Ws in the domain result in larger times between detections of suspicious individuals, and a larger probability of neutralizing R. Results of Gaver et al. (2006, to appear) suggest that the probability of neutralizing R is sensitive to the distribution of the travel and response times. Note also that a large classification time is less harmful than a short time.

The sensitivity of the best classification times to the model assumptions and the parametric form for the probability of correct classification is next explored. The model parameters are those of Table 6-1 unless otherwise displayed in Table 6-2.

Table 6-2. Parameter values for sensitivity analysis

$E[\boldsymbol{T}_c]$	Mean time between detections of suspicious individuals (hrs.)	0.5
$p_{rr}(\tau) = p_{ww}(\tau)$	Probability of correct classification	$\tau_{0.5} = 3/60$ hrs. $\tau_{0.9} =$ Variable
β	Parameter ("shape") for the probability of correct classification	Variable
α	Parameter ("scale") for the probability of correct classification	Variable
w	Number of benign Ws in the domain	5
$E[\boldsymbol{X}_c]$	Mean response time (hrs.)	0.1

Three distributional forms are considered for the search/travel times to individuals and the additional time the searcher is engaged when an individual is classified as R (response time). In each case, the distributional family is the same for both times. The distributions considered are constant times, exponential distributions, and gamma distributions with shape parameter 0.1. The means of the travel (respectively, response) times are equal in all cases.

Two functional forms are considered for the probability of correct classification when the classification time is τ. These forms are the Fréchet distribution in (2.5) and the Weibull distribution,

$$p_{ww}(\tau) = p_{rr}(\tau) = 1 - \exp\left\{-\alpha\tau^{\beta}\right\}. \tag{2.6}$$

In both cases, the parameters α and β are found by specifying the median and 90th percentile of the distribution. The median is $\tau_{0.5} = 0.05$ hours in all cases. The classification times considered are only those greater than or equal to the median of the time for correct classification (i.e. the probability of correct classification is at least 0.50).

Figures 6-2a and 6-2b display the classification times τ that maximize the probability of neutralizing R for various values of $\tau_{0.9}$;

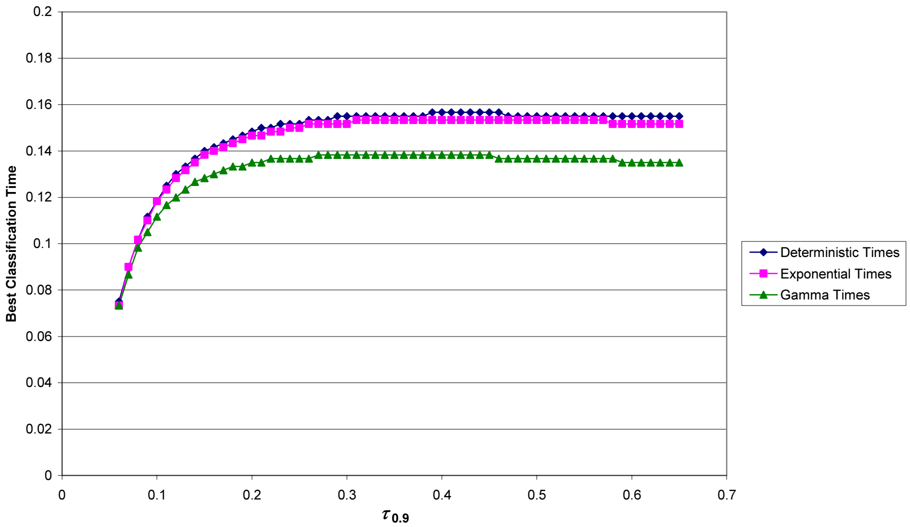

Figure 6-2a. Best classification time for Model 1 with Fréchet probability of correct classification.

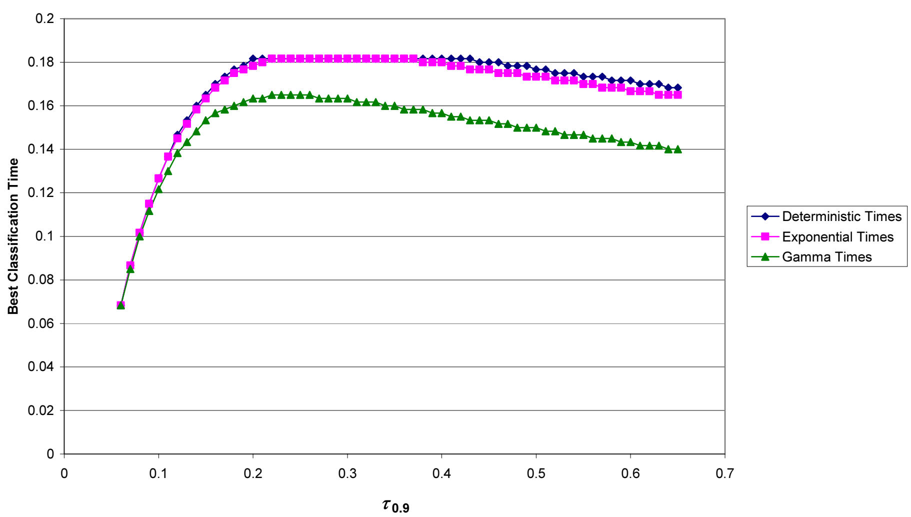

Figure 6-2b. Best classification time for Model 1 with Weibull probability of correct classification.

$p_{ww}(\tau_{0.9}) = p_{rr}(\tau_{0.9}) = 0.9$. Figure 6-3 displays the resulting probabilities of correct classification for the best classification time using the Fréchet probability of correct classification from equation (2.5). Figure 6-4 displays results comparing the use of different C classification times when C's travel and response times have a gamma distribution with shape parameter 0.1 and the classification probabilities are given by equation (2.5). Figure 6-4 displays: the resulting maximum probability of neutralizing R; the probability of neutralizing R if the classification time is given by $\tau = \tau_{0.9}$; the probability of neutralizing R if $\tau = \tau_{0.5}$; the probability of neutralizing R for the policy that uses the (incorrect) exponential distribution for the travel and response times; and the probability of neutralizing R for the policy that uses the (incorrect) Weibull distribution for the classification probabilities. For comparison, if the classification time is 0 and classification is perfect, then the probability of neutralizing R is 0.72 for the parameter values of Figure 6-4.

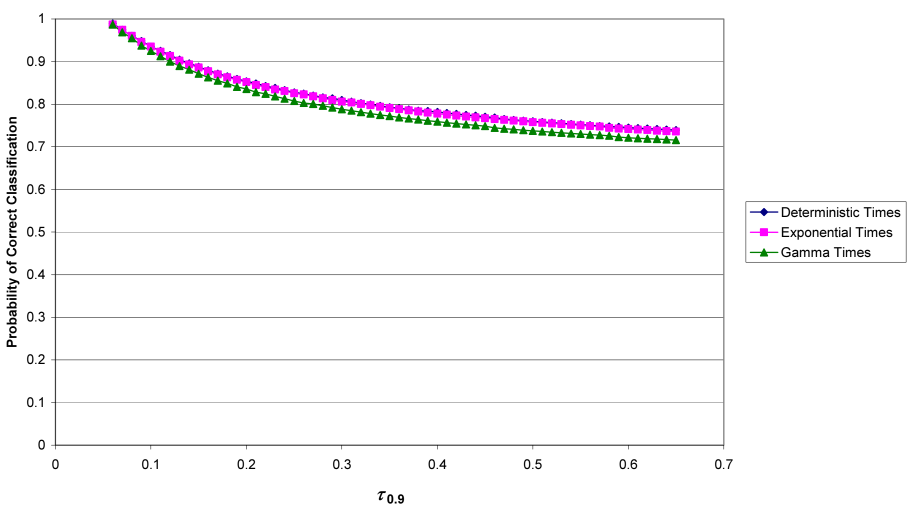

Figure 6-3. Probability of correct classification for the best classification time in Model 1 with Fréchet probability of correct classification.

As can be seen, the best classification times are smaller for travel and response times having gamma distributions with shape parameter 0.1 than for the other cases. Since the travel times and response times are more likely to be less than their means in the gamma case, one can afford to have a slightly smaller probability of correct classification. The maximizing classification times are apparently concave as a function of $\tau_{0.9}$. If $\tau_{0.9}$ is close to $\tau_{0.5}$, then the maximizing classification time is close to (but larger than) $\tau_{0.9}$; it increases as $\tau_{0.9}$ increases and results in a maximum probability of correct classification greater than 0.9. Eventually, $\tau_{0.9}$

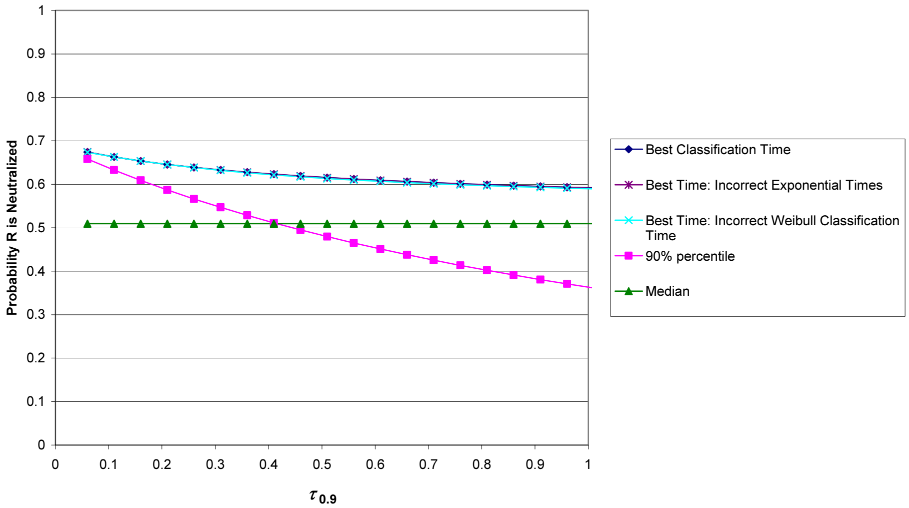

Figure 6-4. Probability R is neutralized when the classification time is based on an incorrect model and the correct model is Fréchet.

becomes too large (with respect to $\tau_{0.5} = 0.05$), and the maximizing classification time begins to decrease slightly, with the resulting probability of correct classification decreasing to about 0.7. In Figure 6-4, the probability of neutralizing R is about the same regardless of whether the classification time is based on the correct model, the incorrect exponential distribution, or the incorrect Weibull distribution. In fact in Figure 6-4, the probabilities of neutralizing R are indistinguishable for these cases. Thus, Figure 6-4 suggests that the best classification times are somewhat insensitive to these model assumptions.

3. STOCHASTIC MODEL 2

A lethal Red, R, stalks groups or clusters of neutral Whites at random in an arena or domain. Assume that R eventually identifies a group of sufficiently high threshold value; the value is the size of the group or the importance of its constituency. The R mingles with this group and self-destructs, reducing the value of the group from V to $0 \le V' \le V$. Given that one (or more) counter-terrorists are in search of R, and require time to thoroughly identify a suspect—which may be ill-advised if the suspect is only a harmless White (W)—what is an optimal strategy for R to pick a target group, and for the counter-terrorist C to detain a potential R, when there is a chance that C is actually neutralizing a neutral (White)?

Suppose R encounters or contacts groups of possible targets at random in a Poisson manner at rate λ, and suppose the ith group to appear has a random value, $\boldsymbol{V}_i$; the values of the various groups are assumed iid. It may be that $\boldsymbol{V}_i$ is merely the number of individuals in the group, but it could also be the value associated with a single exceptional individual (person, or platform). If F is the distribution from which the collection of $\{\boldsymbol{V}_i\}$ is drawn, then the maximum group value seen by R in $(0,t]$ has distribution

$$\begin{aligned} P\{\boldsymbol{V}(t;M) \le v\} &= \sum_{n=0}^{\infty} e^{-\lambda t} \frac{[\lambda t]^n}{n!} \left[F(v)\right]^n \\ &= e^{-\lambda t\left[1-F(v)\right]} \equiv e^{-\lambda t \bar{F}(v)}. \end{aligned} \tag{3.1}$$

While R hunts for valuable White targets, Blue C searches for R as effectively and surreptitiously as possible.

R chooses a minimum threshold value ω for a target to pursue. R detonates the first time it encounters a target with value at least ω. The time until R detonates, $\boldsymbol{D}$, has an exponential distribution with mean

$$1/\delta_s(\omega) = 1/\lambda\left[1-F(\omega)\right] \equiv 1/\lambda \bar{F}(\omega). \tag{3.2}$$

3.1 Blue Searcher Strategy

The Blue searcher, C, covers or searches the domain. C chooses a time τ to spend classifying detected individuals. The probability that R does not detonate before it is neutralized is $P_B(\tau)$ given in equations (2.4a–2.4b) with $\delta_s = \delta_s(\omega)$ given by equation (3.2). Hence, the expected lethal value achieved by the R is

$$\begin{aligned} &E[\boldsymbol{V};\omega,\tau] \\ &= \left[1-P_B(\tau)\right] \frac{\int_{\omega}^{\infty} vF(dv)}{\bar{F}(\omega)}. \end{aligned} \tag{3.3}$$

Viewing the problem as a game with opposed objectives, we have:

C's problem: $\min_{\tau>\tau_{0.5}} \max_{\substack{\omega \\ \text{for} \\ \text{given} \\ \tau}} E[\boldsymbol{V};\omega,\tau]$

R's problem: $\max_{\omega} \min_{\substack{\tau>\tau_{0.5} \\ \text{for} \\ \text{given} \\ \omega}} E[\boldsymbol{V};\omega,\tau]$

Empirical evidence suggests that this game has a Nash equilibrium; see Figure 6-5.

EXAMPLE

Assume that the target values have an exponential distribution. The travel times of C between individuals, $\boldsymbol{T}_c$, and the time until C is relieved after classifying an individual as R, $\boldsymbol{X}_c$, are independent and exponentially distributed. Moreover, we assume that $p_{ww}(\tau) = p_{rr}(\tau) = \exp\{-\alpha/\tau^{\beta}\}$ for $\tau > \tau_{0.5}$. Figure 6-5 displays the expected value of the return to R as a function of R's threshold and C's classification time for the model parameters displayed in Table 6-1. In Figure 6-5 $E[\boldsymbol{T}_c] = 0.5$ hours, $E[\boldsymbol{X}_c] = 1$ hours, the rate at which R encounters groups is 2 per hour, and the mean group size is 10. Note that the function apparently has a saddle point.

Parameter values are displayed in Table 6-3. Table 6-4 displays the best policies for both C and R for various values of the mean time until C is relieved after classifying an individual as R, the number of Ws in the domain, and the mean target value. Table 6-4 displays results for two different distributions for C's travel times and response times. The family of distributions is assumed to be the same for both travel times and response times; the two distributions are the exponential distribution and the gamma distribution with shape parameter 0.1 having the same mean as the exponential. Such a positively skewed gamma represents substantial spatial clumpiness of surface entities, hence permitting C to investigate several suspicious individuals in more rapid succession than would be true for the exponential case; see the right-most column of Table 6-4.

DISCUSSION

The expected lethal value achieved by R increases: as the mean time until the searcher is relieved, $E[\boldsymbol{X}_c]$, increases; as the mean target value

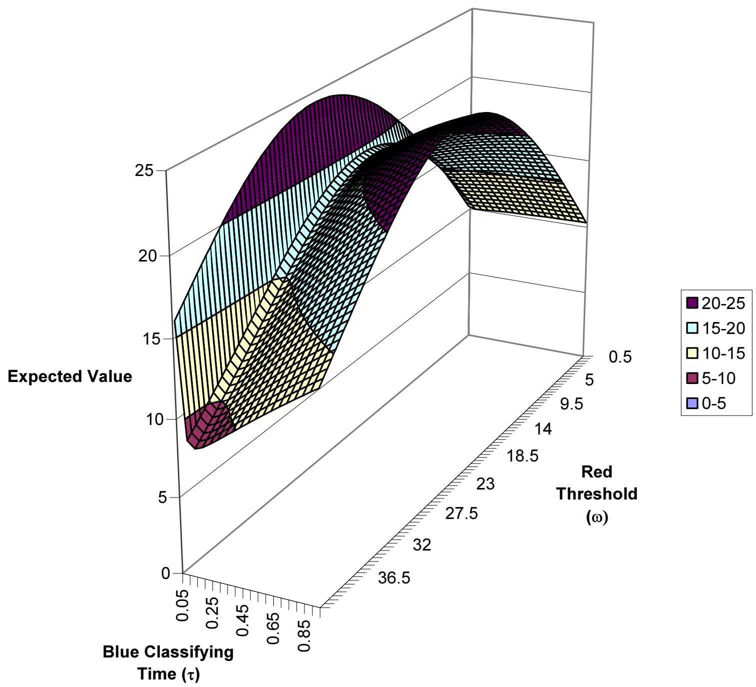

Figure 6-5. Expected value of the return to R.

Table 6-3. Parameter values for Model 2

$p_{rr}(\tau) = p_{ww}(\tau)$	Probability of correct classification	$\tau_{0.5} = 3/60\,\text{hrs.}$ $\tau_{0.9} = 5/60\,\text{hrs.}$
β	Parameter ("shape") for the probability of correct classification	3.69
α	Parameter ("scale") for the probability of correct classification	0.00001
$k_r = 1 - k_w$	Probability detected individual is the R	$1/(1+w)$
f	Side of square sensor footprint (nm)	10
v_s	Velocity of sensor platform (kts)	250
M_x	Length of x-direction of rectangular domain (nm)	20
M_y	Length of y-direction of rectangular domain (nm)	20

increases; and as the rate at which R encounters possible targets increases. The best threshold value ω for R increases in the mean of $\boldsymbol{X}_c$, the mean target value, the number of Ws, and the rate at which R encounters possible targets. The best classification time τ for the searcher increases as the mean of $\boldsymbol{X}_c$ increases. The best classification time τ for the searcher appears not

Table 6-4. Representative results for Model 2

$E[X_c]$	Rate R Encounters Targets λ	Number of Benigns w	Mean Target Value	Expected Value Achieved by R Exponential [Gamma]	τ Exponential [Gamma]	ω Exponential [Gamma]	Probability R is Neutralized Exponential [Gamma]
0.1	0.1	0	5	0.18 [0.17]	0.086 [0.084]	0.10 [0.30]	0.96 [0.97]
0.5	0.1	0	5	0.36 [0.32]	0.086 [0.084]	0.25 [0.80]	0.93 [0.94]
0.1	0.1	5	5	0.41 [0.41]	0.082 [0.080]	0.35 [0.40]	0.92 [0.92]
0.5	0.1	5	5	0.65 [0.61]	0.100 [0.098]	0.50 [0.80]	0.88 [0.90]
0.1	0.1	10	5	0.63 [0.62]	0.080 [0.080]	0.55 [0.60]	0.89 [0.89]
0.5	0.1	10	5	0.90 [0.86]	0.102 [0.100]	0.75 [1.00]	0.84 [0.86]
0.1	0.1	0	10	0.36 [0.35]	0.086 [0.084]	0.25 [0.65]	0.96 [0.97]
0.5	0.1	0	10	0.73 [0.65]	0.086 [0.084]	0.55 [1.60]	0.93 [0.94]
0.1	0.1	5	10	0.83 [0.82]	0.082 [0.080]	0.70 [0.80]	0.92 [0.92]
0.5	0.1	5	10	1.30 [1.22]	0.100 [0.098]	1.05 [1.60]	0.88 [0.90]
0.1	0.1	10	10	1.25 [1.25]	0.080 [0.080]	1.15 [1.20]	0.89 [0.89]
0.5	0.1	10	10	1.81 [1.71]	0.102 [0.100]	1.50 [1.95]	0.84 [0.86]
0.1	1.5	5	10	9.12 [6.84]	0.080 [0.078]	8.40 [6.65]	0.50 [0.59]
0.5	1.5	5	10	11.64 [8.16]	0.098 [0.088]	10.20 [8.75]	0.42 [0.56]
0.1	1.5	10	10	11.84 [8.89]	0.080 [0.078]	11.30 [9.65]	0.44 [0.52]
0.5	1.5	10	10	14.26 [10.31]	0.100 [0.090]	13.00 [10.60]	0.38 [0.50]

to depend strongly on the mean target value. The gamma distribution for the $\boldsymbol{T}_c$ and $\boldsymbol{X}_c$, results in shorter optimal classification times for the searcher than for the exponential distribution. Apparently, this is due to the fact that C's travel times and response times are more likely to be less than their means in the gamma case than in the exponential case. The probability of neutralizing R decreases in λ, since small values of λ reflect R's difficulty finding targets; thus, it is apparently worthwhile to "harden" facilities in order to deny R access to potential targets.

The sensitivity of the searcher's best classification times and best threshold values, ω, to model assumptions and parametric forms is now explored. The model parameters are those of Table 6-1, as modified by Table 6-2. The rate at which R encounters individuals is given by $\lambda = 2$, and the mean of the exponential target values is 5. As before, two distributional forms are considered for the searcher's travel time and the response time with the same distributional family used for both times. The distributions considered are the exponential distribution and the gamma distribution with shape parameter 0.1. The means of the travel (respectively, response) times are equal in all cases. Two functional forms are considered for the probability of correct classification when the classification time is τ. These forms are the Fréchet distribution of equation (2.5) and the Weibull distribution of equation (2.6). As before, the parameters α and β are determined by specifying the median and 90th percentile of the distribution and only classification times greater than or equal to the median are considered.

Figures 6-6a and 6-6b display the optimal classification times τ for various values of $\tau_{0.9}$ when $\tau_{0.5} = 0.05$. Figure 6-7 displays the resulting probabilities of correct classification for the best classification times for the Fréchet classification probability function. Figure 6-8 (respectively Figure 6-9) displays the best threshold value, ω, (respectively, the expected target value achieved by R) for the two distributions for the searcher's travel and response times. Figure 6-10 (respectively Figure 6-11) displays the resulting best game probability of neutralizing the R (respectively the expected target value achieved by R) for various choices of classification times when C's travel and response times have a gamma distribution with shape parameter 0.1 and the probability of correct classification is of form (2.5); the classification times considered are the best classification time; $\tau_{0.9}$; $\tau_{0.5}$; the best classification time for the incorrect model of exponential times; and the best classification time for incorrect Weibull classification probabilities. In all cases R is assumed to use its optimal threshold value ω for the correct model.

The best classification times are smaller for travel and response times having gamma distributions with shape parameter 0.1 than for those with an exponential distribution. Since the travel and response times are more likely to be less than their means in the gamma case, one can afford to have a slightly smaller probability of correct classification. The maximizing classification times are apparently concave as a function of $\tau_{0.9}$. Figure 6-8 suggests that the maximizing threshold value for R is about the same regardless of whether the searcher's travel and response times are exponential or gamma distributed. However, R achieves a lower expected target value for gamma distributed searcher travel and response times than in

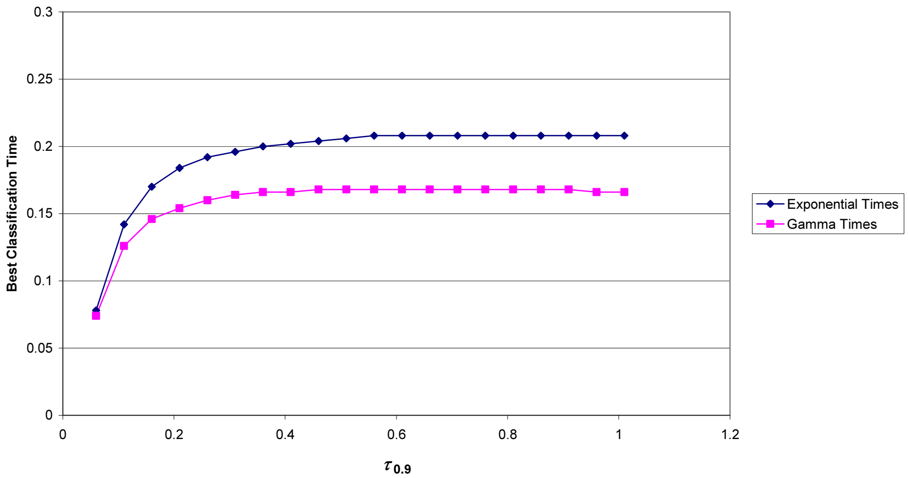

Figure 6-6a. Best classification time for Model 2 with Fréchet probability of correct classification.

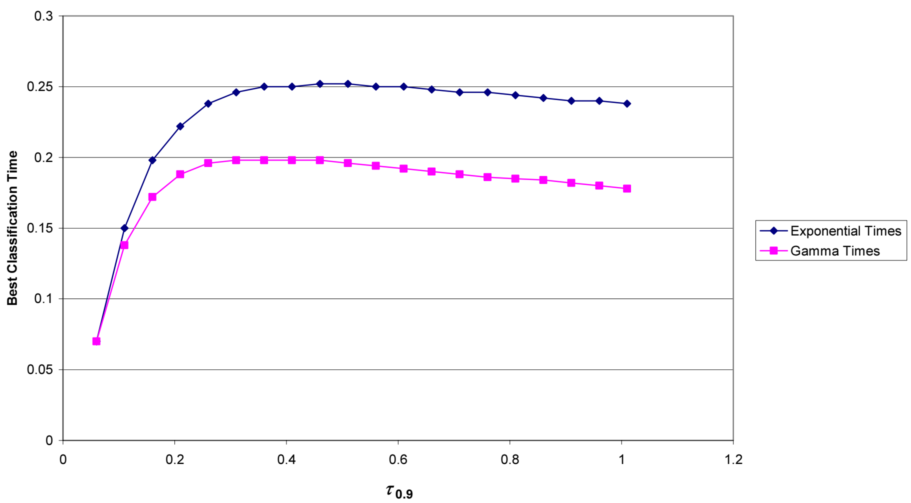

Figure 6-6b. Best classification time for Model 2 with Weibull probability of correct classification.

the exponential case. Figures 6-10 and 6-11 suggest setting the classification time equal to $\tau_{0.9}$ results in almost the same probability of neutralizing R as using the optimal classification time for values of $\tau_{0.9}$ between 0.06 hours and 0.4 hours; however, the probability of neutralizing R decreases for larger

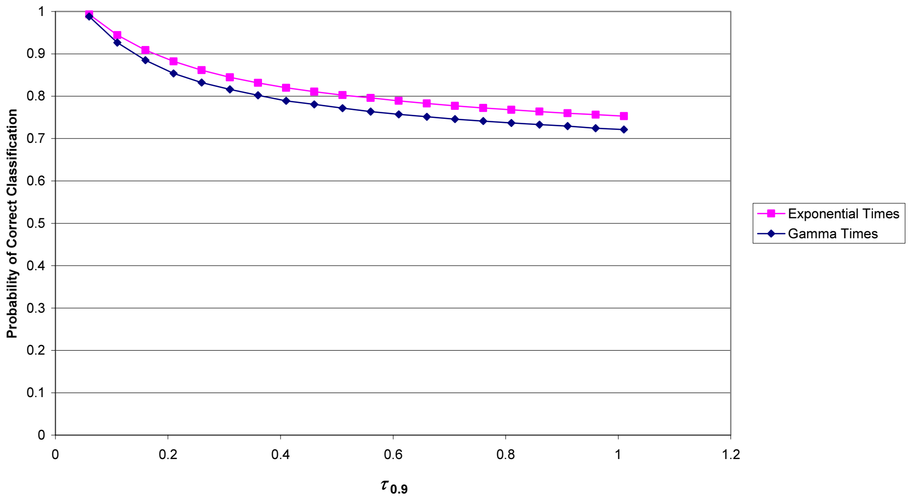

Figure 6-7. Probability of correct classification for the best classification time in Model 2 with Fréchet probability of correct classification.

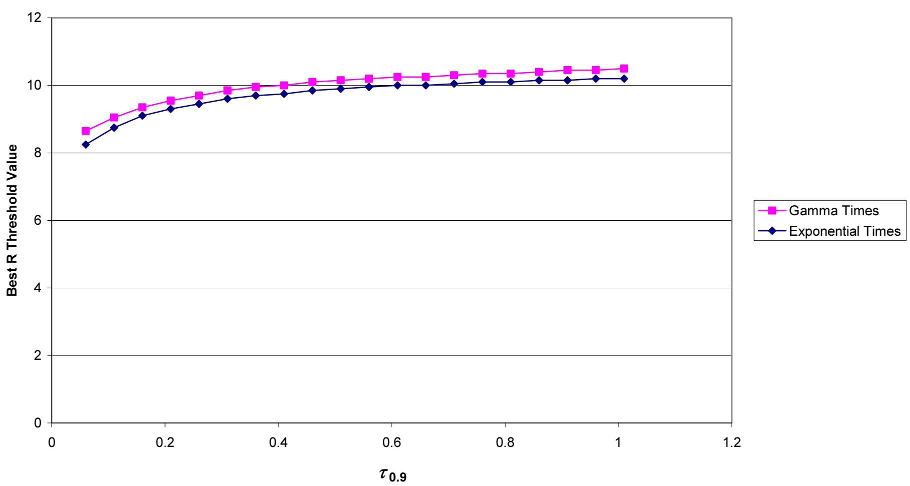

Figure 6-8. Best R threshold value for Model 2 with Fréchet probability of correct classification.

values of $\tau_{0.9}$. Thus, it is apparently worthwhile for the searcher to determine and use the optimal classification time at least when large classification times might be needed to achieve a 90% probability of correct classification. Figures 6-10 and 6-11 suggest that the probability of

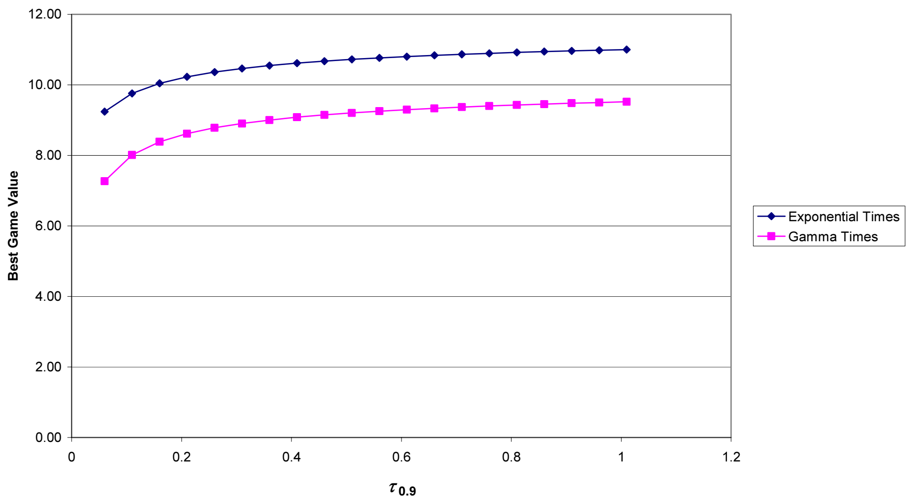

Figure 6-9. Expected target value achieved by R for Model 2 with Fréchet probability of correct classification.

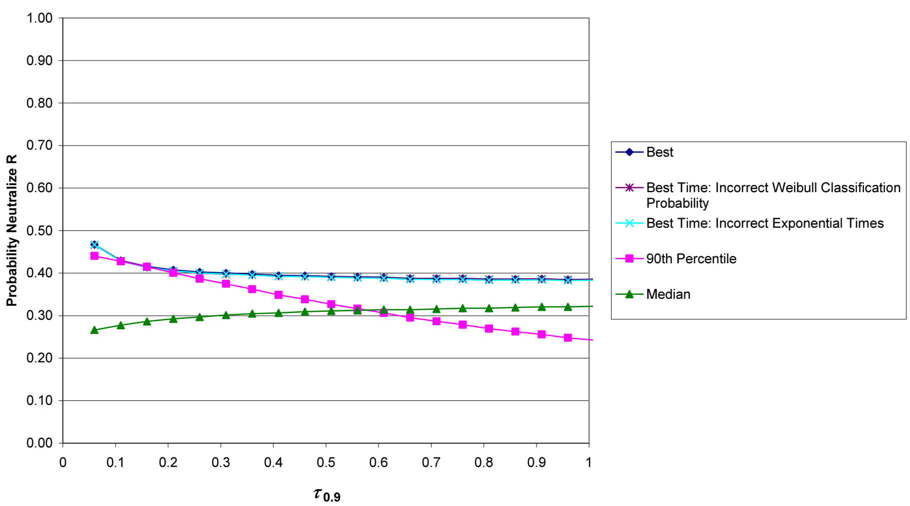

Figure 6-10. Probability R is neutralized in Model 2 for various choices of classification time and R uses the best threshold for the correct model.

neutralizing R is somewhat insensitive to both the distributional form of the searcher's travel and response times and the model for the probability of correct classification.

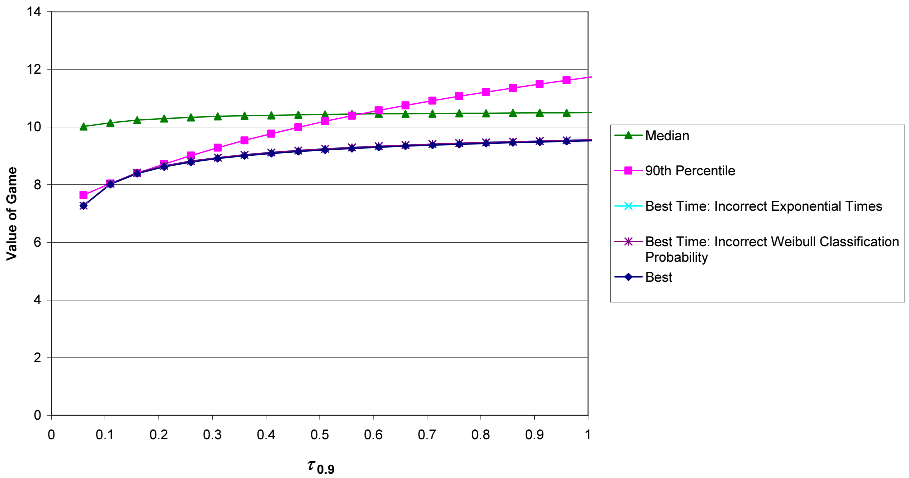

Figure 6-11. Expected target value achieved by R in Model 2 for various choices of classification time and R uses the best threshold for the correct model.

4. CONCLUSIONS

In the scenario described in this chapter, there are neutral individuals, Whites (W), and one hostile individual, Red (R), traveling within a domain. The R's purpose is to attack a target. A patrolling counter-terrorist, C, detects individuals in the domain, and classifies them as W or R. The probability of correct classification is an increasing function of the classification time. C follows each individual that is classified as R (perhaps incorrectly) until relieved by others; during this time, C is unavailable to detect and classify additional individuals. The ability to detect and correctly classify R before R reaches its objective is influenced by the size of the domain, the number of Ws in the domain, and the probability of correctly classifying detected individuals.

In both models, C wishes to choose the best classification time. If the classification time is too small, C will "waste time" following misclassified Ws, and thereby increase R's chance of reaching its objective; if the classification time is too long, R will reach its objective before C can detect and neutralize it. In the initial model of Section 2, R reaches its objective after an exponential time if not neutralized before then; this exponential time is independent of C's actions. In the model of Section 3, the targets have values, and R chooses a policy to maximize its expected reward (equal to 0 if R is neutralized, and equal to the value of the target if R reaches its objective). The numerical results suggest that the probability of neutralizing R resulting from the best policy is robust to incorrect specification of the distribution of the travel and response times, provided the correct mean times are used. The numerical results also suggest that the neutralization probability corresponding to the best policy is robust to incorrect specification of the probability of correct classification as a function of time. The probability of neutralizing R can be increased by decreasing the number of unidentified neutral individuals and by decreasing the rate at which R encounters possible targets (possibly by hardening infrastructure, for example).

The time until R achieves its objective has an exponential distribution for both of the models discussed here. However, the results of Gaver et al. (2006, to appear) suggest that the probability of neutralizing R is sensitive to the distribution of the time until R achieves its objective, so other functional forms should perhaps be considered. It would also be of interest to study the effect of C tagging those Ws it has already investigated.

In the model of Section 3, R knows the distribution of target values. Another area of future study is the development of models in which R learns about its target population. In this case, we would expect R's policy to depend on the variability of the target values. As the target values become more variable, we expect it would be worthwhile for R to spend more time observing target values before determining a threshold target value. Further, if the mean prior arrival rate is large, then we would expect that only a small amount of time would be needed by R to get good information about target values, reducing the risk of being neutralized.

ACKNOWLEDGEMENTS

D. P. Gaver and P. A. Jacobs acknowledge with thanks support from the Center for Defense Technology and Education for the Military Services Initiative at the Naval Postgraduate School.

REFERENCES

Gaver, D. P., Jacobs, P. A., and Sato, H., 2006, "Assessing resource requirements for maritime domain awareness and protection," Naval Postgraduate School Technical Report, NPS-OR-06-010.

Gaver, D. P., Jacobs, P. A., and Sato, H., "Assessing resource requirements for maritime domain awareness and protection," to appear in *International Journal of Performability Engineering.*

Kott, A., and McEneaney, W. M., eds., 2006, *Adversarial Reasoning (Computational Approaches to Reading the Opponent's Mind),* Chapman & Hall/CRC, New York.

Kress, M., 2005, "The effect of crowd density on the expected number of casualties in a suicide attack," *Naval Research Logistics* **52**:22–29.

Thomas, L. C., 1984, *Games, Theory and Applications,* Wiley, New York.