Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Terror Queues

Edward H. Kaplan,

OPERATIONS RESEARCH
Vol. 58, No. 4, Part 1 of 2, July–August 2010, pp. 773–784
ISSN 0030-364X | EISSN 1526-5463 | 10 | 5804 | 0773

DOI 10.1287/opre.1100.0831

Terror Queues

Edward H. Kaplan
Yale School of Management, Yale School of Public Health, Yale School of Engineering and Applied Science,
Yale University, New Haven, Connecticut 06520, edward.kaplan@yale.edu

This article presents the first models developed specifically for understanding the infiltration and interdiction of ongoing terror plots by undercover intelligence agents, and does so via novel application of ideas from queueing theory and Markov population processes. The resulting "terror queue" models predict the number of undetected terror threats in an area from agent activity/utilization data, and also estimate the rate with which such threats can be detected and interdicted. The models treat terror plots as customers and intelligence agents as servers. Agents spend all of their time either detecting and infiltrating new terror plots (in which case they are "available"), or interdicting already detected terror plots (in which case they are "busy"). Initially we examine a Markov model assuming that intelligence agents, while unable to detect all plots, never err by falsely detecting fake plots. While this model can be solved numerically, a simpler Ornstein-Uhlenbeck diffusion approximation yields some results in closed form while providing nearly identical numerical performance. The transient behavior of the terror queue model is discussed briefly along with a sample sensitivity analysis to study how model predictions compare to simulated results when using estimated versus known terror plot arrival rates. The diffusion model is then extended to allow for the false detection of fake plots. Such false detection is a real feature of counterterror intelligence given that intelligence agents or informants can make mistakes, as well as the proclivity of terrorists to deliberately broadcast false information. The false detection model is illustrated using suicide bombing data from Israel.

Subject classifications: intelligence; counterterrorism; informants; Markov models; diffusion models; queues; suicide bombings.
Area of review: Stochastic Models.
History: Received June 2009; revisions received December 2009, March 2010; accepted March 2010. Published online in *Articles in Advance* June 3, 2010.

1. Introduction

Intelligence operations protect society against the threat of terrorism. Guiora (2007, p. 218) argues that "Without intelligence, governments cannot protect their citizens. Without knowing who the terrorist is, governments cannot know where he is. Without knowing where the suicide bomber is, governments cannot prevent planned suicide bombing attacks," and in a subsequent book states succinctly that "Intelligence is the heart and soul of operational counterterrorism" (Guiora 2008, p. 80).

In Israel, human intelligence operations (HUMINT) have been credited in large measure with bringing the wave of suicide bombing attacks on Israeli civilians that accompanied the Second Intifada under control (Kaplan and Kress 2005, Kaplan et al. 2006). Undercover counterterror intelligence operations have a long history in Israel, where units known in Hebrew as "mistaravim" (agents disguised as Arabs) have immersed themselves culturally and physically in Israeli Arab and Palestinian populations for the purpose of infiltrating and interdicting terror cells before they can attack (Deflem 2010). Indeed, in a candid response to a question asking whether Israeli intelligence had infiltrated the Hamas (which was responsible for the majority of suicide bombing attacks inside Israel), the former head of the Palestinian Authority's General Intelligence Tawfiq al Tirawi responded that *all* Palestinian organizations had been penetrated by the Israelis (Andraeus 2009).

The importance of infiltrating terror plots is also reflected in the national intelligence strategy of the United States. The first of six mission objectives stated in this strategy document is "Combat Violent Extremism," and within this mission objective, the intelligence community is urged to "…penetrate and support the disruption of terrorist organizations…" (Blair 2009). A recent United States example of the importance of such HUMINT infiltrations can be found in the discovery and interruption of a plot to bomb two synagogues and shoot down a military aircraft with a surface-to-air missile; as reported in the *New York Times*, "The investigation…began with the work of a confidential informant, who portrayed himself as an agent of a Pakistani terror organization, and who became a critical member of the men's plot" (Wilson 2009). Other recent terror threats in the United States that were penetrated and disrupted by undercover intelligence operatives include the interdiction of Najibullah Zazi, charged with plotting bomb attacks (Johnston and Shane 2009), and the arrest of Hosam Maher Husein Smadi, charged with attempting to bomb a Dallas skyscraper (Associated Press 2009).

Beyond the penetration and disruption of individual terror plots, data gleaned from intelligence operations offer the potential to infer the magnitude of current terror threats. For example, in May 2007, Britain's Security Service (the MI5, http://www.mi5.gov.uk) reported that it was monitoring about 2,000 people believed to be "…actively involved in supporting Al Qaeda" (Gardner 2007). However, the number of suspects under surveillance is presumably limited by the number of intelligence agents/informants actively searching for and/or watching such persons! Using whatever criteria MI5 employs to designate individuals as Al Qaeda supporters, how many more (or fewer) individuals would be observed if MI5 doubled (or halved) the number of undercover agents/informants dedicated to detecting and monitoring terror suspects? In principle, if the processes by which intelligence agents detect and interdict terror cells can be modeled, then one should be able to infer answers to questions of this sort from data generated by intelligence agents in the field.

This article provides the first models developed specifically for understanding the infiltration and interdiction of ongoing terror plots by undercover intelligence agents, and does so via novel application of ideas from queueing theory and Markov population processes. The resulting "terror queue" models predict the number of undetected terror threats in an area from agent activity/utilization data while also estimating the rate with which such threats can be interdicted. The customers in this system are the terror plots themselves, while the servers are "deep" undercover intelligence agents who, when successful, infiltrate terror plots and facilitate their disruption. Terror plots that evade detection and result in successful terror attacks are analogous to customers who renege from a queue before receiving service; plots that are detected and infiltrated by intelligence agents are analogous to customers who enter service; and plots that are interdicted (with the terrorists involved arrested or otherwise incapacitated) are analogous to customers who complete service. Unlike typical queueing systems, customers (terror plots) arriving to terror queues are not immediately visible upon arrival; rather, customers must be discovered by available servers before service can commence. Consequently, and again unlike typical queues, waiting customers (undetected terror plots) and available servers (uncommitted intelligence agents) coexist: it is not the case in terror queues that waiting customers imply that all servers are busy. Finally, while it is especially important for national security that "customers" are "well-serviced" in such terror queues, this is one queueing system where customers clearly prefer *not* to receive service!

The next section provides a steady-state Markov model for the simplified situation where intelligence agents, while unable to detect all terror plots, never err by falsely detecting "fake" plots and hence do not waste any time investigating suspects who ultimately are not involved in terrorism. Agents spend all of their time either detecting new terror plots (in which case they are "available"), or interdicting newly detected terror plots (in which case they are "busy"). While this model has a clean formulation, it can only be solved numerically. However, the numerical results from this model suggest that the numbers of undetected terror plots and busy intelligence agents embedded in detected plots follow a bivariate normal distribution, so there is a simple linear relation between the expected number of undetected plots and the observable number of busy agents. Such findings are consistent with the literature on Markov population processes and associated diffusion approximations.

Taking advantage of known approximation techniques from this literature, in §3 we construct a much simpler Ornstein-Uhlenbeck model that ensures such joint normality and note that the numerical results of this latter model are nearly identical to the Markov approach (except near the boundaries of the state space, for which separate approximations are provided). We also consider briefly the transient behavior of the terror queue model using a combination of analytic approximations and simulation and provide a sample sensitivity analysis to study how model predictions compare to simulated results when using estimated versus known terror plot arrival rates.

In §4, we extend the Ornstein-Uhlenbeck model to the more realistic case of false detection. False detection is an important reality of counterterror intelligence, not only because intelligence agents or informants can make mistakes, but also due to the proclivity of terrorists to broadcast false information. With false detection, intelligence agents become occupied investigating cases that ultimately prove to be unconnected in a significant way to actual terror plots. Such cases are costly because they divert agents from the ability to detect and investigate real plots, which in turn serves to lower the rate with which real plots can be interdicted. Section 5 closes with a brief discussion of the approach, parameter estimation, potential applications of the model to decision problems in counterterror intelligence operations, and ideas for future research.

2. Markov Terror Queue

This section develops a Markov model for the joint distribution of undetected and detected terror plots in steady state. As a first model for this process, and to maintain tractability, the assumptions employed are the simplest possible. The focus is on introducing a new modeling paradigm for (counter)terrorism via simple examples rather than on attempting to obtain the highest degree of generality.

Assume that new terrorist plots are hatched in accord with a Poisson process with rate α per unit time (the case where α itself changes with time is considered briefly in §3.2). The time required to plan and carry out a terror attack follows an exponential distribution with mean $1/\mu$. Let X denote the number of undetected terror plots-in-progress; absent interdiction, all terror plots will be successful, and X can therefore be thought of as the number of customers in a self-service (and hence $M/M/\infty$) queue. Alternatively, one can think of the terror plots as customers with per capita reneging rate μ (as is common in call center models, e.g., Garnett et al. 2002) who are waiting to enter service (detection by intelligence agents) that never comes (since there are as yet no intelligence agents in the model). The result is the same: in steady state, X will be distributed as a Poisson random variable with mean α/μ.

However, the government deploys a force of f undercover intelligence agents for the purpose of detecting and infiltrating terror plots and interdicting those responsible. For our purposes, an agent could refer to a combination of intelligence personnel and/or informants who function as a single unit. Let Y denote the steady-state number of terror plots detected and infiltrated but as yet not interdicted. We assume that agents spend all of their time either detecting/infiltrating terror plots (these agents are defined to be "available") or as part of those terror plots they have infiltrated until those plots are interdicted (these agents are defined to be "busy"). Thus, there is a one-to-one correspondence between the number of detected but not yet interdicted terror plots and the number of busy undercover agents. The agents themselves need not be responsible for the physical capture of terrorists, but we do presume that agents remain busy until a detected plot is interdicted. If a terror plot has been detected, we assume that the terrorists involved are interdicted in time to prevent an attack, and the additional time from detection to interdiction follows an exponential distribution with mean $1/\rho$. Finally, we assume that undetected terror plots are detected at a rate proportional to both the extant number of undetected plots X and the number of available intelligence agents $f-Y$, that is, we assume that terror plots are detected at rate $\delta X(f-Y)$ where δ is the detection rate per terror plot per available agent per unit time. Other detection functions are possible, but beyond expecting the overall detection rate to be an increasing function of both the number of undetected plots (X) and available agents ($f-Y$), it is difficult to argue why one functional form for the aggregate detection rate should be preferable to another. Similarly, it is possible that the detection rate could depend in some way on the total number of agents f, though again it is difficult to state if or how this occurs absent detailed data.

A flow diagram for this "terror queue" model appears in Figure 1.

The "law of conservation of terror plots" requires that in steady-state,

$$\alpha = \mu E(X) + \rho E(Y). \tag{1}$$

Since every plot hatched results in a successful terror attack or is interdicted, the probability that a terror plot is interdicted equals

$$\Pr\{\text{Interdiction}\} = \frac{\rho E(Y)}{\alpha}. \tag{2}$$

Also of great interest is the conditional distribution of the number of undetected terror plots given the number of busy intelligence agents (which equals the number of detected terror plots).

The process just described can be modeled as a bivariate Markov process with state (X, Y) corresponding to the number of undetected and detected terror plots, respectively. The set of possible states is given by the integers $\{x = 0, 1, 2, \ldots; y = 0, 1, 2, \ldots, f\}$. Let the steady state probability of x undetected and y detected terror plots be given by $p_{xy} = \Pr\{X = x, Y = y\}$. These probabilities must satisfy the generic balance equations

$$\begin{aligned}(\alpha + \mu x + \rho y + \delta x(f-y))p_{xy} &= \alpha p_{x-1,\,y} + \mu(x+1)p_{x+1,\,y} + \rho(y+1)p_{x,\,y+1} \\ &\quad + \delta(x+1)(f-y+1)p_{x+1,\,y-1}\end{aligned} \tag{3}$$

for $x = 1, 2, \ldots; y = 1, 2, \ldots, f-1$ (see Figure 2), in addition to the boundary conditions

$$\alpha p_{00} = \mu p_{10} + \rho p_{01} \tag{4}$$

$$(\alpha + f\rho)p_{0f} = \mu p_{1f} + \delta p_{1,\,f-1} \tag{5}$$

$$(\alpha+\rho y)p_{0y} = \mu p_{1,\,y} + \rho(y+1)p_{0,\,y+1} + \delta(f-(y-1))p_{1,\,y-1} \quad \text{for } y = 1, 2, \ldots, f-1 \tag{6}$$

$$(\alpha+\mu x+\delta f x)p_{x0} = \alpha p_{x-1,\,0} + \mu(x+1)p_{x+1,\,0} + \rho p_{x,\,1} \quad \text{for } x = 1, 2, \ldots \tag{7}$$

$$(\alpha+\mu x+\rho f)p_{xf} = \alpha p_{x-1,\,f} + \mu(x+1)p_{x+1,\,f} + \delta(x+1)p_{x+1,\,f-1} \quad \text{for } x = 1, 2, \ldots \tag{8}$$

and the probability conservation equation

$$\sum_{x=0}^{\infty}\sum_{y=0}^{f} p_{xy} = 1. \tag{9}$$

Figure 1. Terror plot flow in the terror queue model.

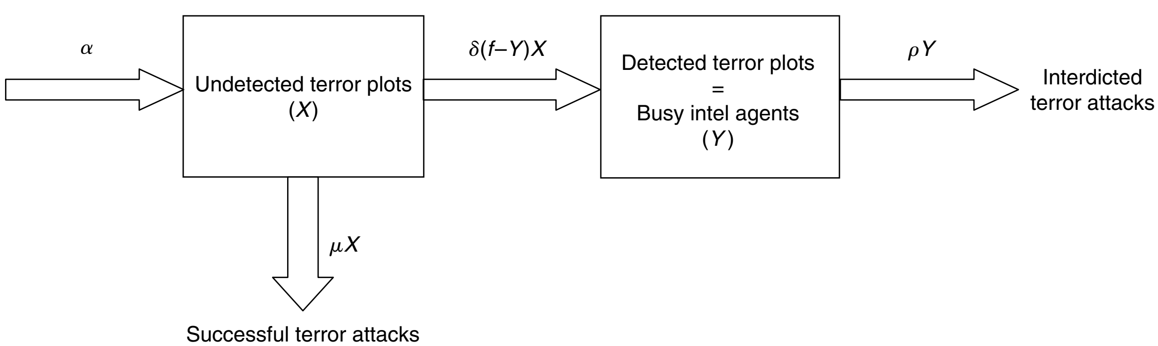

Figure 2. State transitions in the Markov terror queue model.

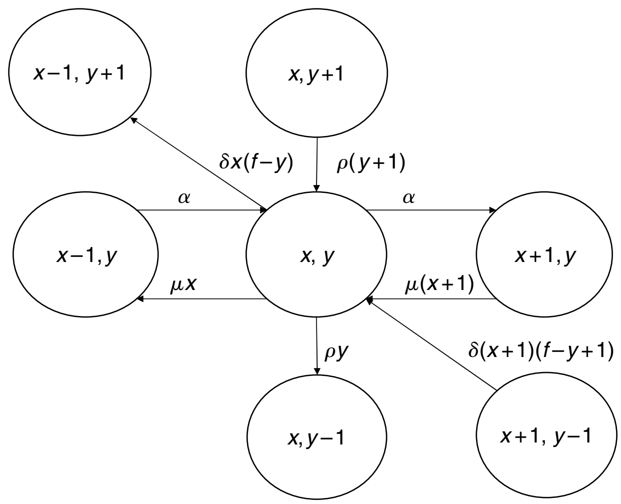

The solution to these linear equations must be found numerically and, once obtained, can be used to compute $E(X)$, $E(Y)$, the probability of interdicting an attack, and the conditional distribution of the number of undetected terror plots given the observed number of busy agents, that is, $\Pr\{X=x \mid Y=y\}$.

As a numerical example, Figure 3 presents the joint probability distribution of undetected and detected terror plots corresponding to the following parameter choices: $\alpha = 100$ new terror plots per year, $\mu = 1$ attack per year (meaning an average time of one year from conception to completion of a terror attack), $\delta = 0.1$ detections per terror plot per undercover agent per year, $\rho = 4$ interdicted plots per busy undercover agent per year (meaning an average time of three months from detection to interdiction), and $f = 30$ undercover agents. This distribution is found from the solution of

Figure 3. The steady state joint distribution of undetected (X) and detected (Y) terror plots in the Markov terror queue model when $\alpha = 100$, $\mu = 1$, $\delta = 0.1$, $f = 30$, and $\rho = 4$.

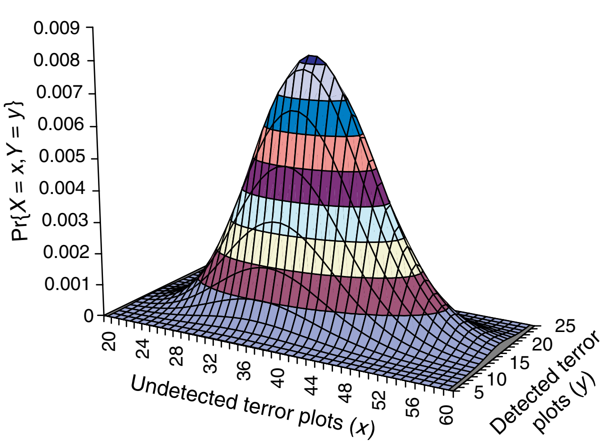

Equations (3)–(9) using the parameter values just described. Even though on average only half of the 30 intelligence agents are busy, the undercover effort reduces the mean number of undetected plots from 100 to 40, which in this example corresponds to an interdiction probability of 60%.

The conditional distribution of the number of undetected terror plots given the observed number of busy intelligence agents is given by

$$\Pr\{X=x \mid Y=y\} = \frac{\Pr\{X=x,\, Y=y\}}{\Pr\{Y=y\}} = \frac{p_{xy}}{\sum_{i=0}^{\infty} p_{iy}}$$
$$x = 0, 1, \ldots; \quad y = 0, 1, \ldots, f \qquad (10)$$

from which the conditional mean number of undetected plots $E(X \mid Y = y)$ can be evaluated along with the moments or distribution of any other desired conditional function. The marginal distribution of the number of busy agents corresponding to the joint distribution shown in Figure 3 is graphed in Figure 4; also plotted is the conditional expected number of undetected plots given the number of busy agents ($E(X \mid Y = y)$). Note that the marginal distribution of Y resembles a normal distribution, while the conditional expected number of undetected terror plots $E(X \mid Y = y)$ is linear in y. These results should not prove surprising given that the joint distribution of X and Y shown in Figure 3 resembles a bivariate normal distribution (with normally distributed marginals as a consequence). Further, that the conditional mean of one random variable is linear in the given value of a second random variable is well known when the random variables in question are governed by a bivariate normal distribution (e.g., Freund 1971, §13.2). In the example under consideration, the slope of the conditional expectation (or regression) line approximately equals 1/2, suggesting that, on average, there is one additional undetected terror plot for every two busy undercover agents.

Figure 4. Comparing the Markov terror queue model and its diffusion approximation: the steady state marginal distribution of the number of detected terror plots (busy agents) Y, and the conditional expected number of undetected terror plots X given that $Y = y$, $E(X \mid y)$, for both models.

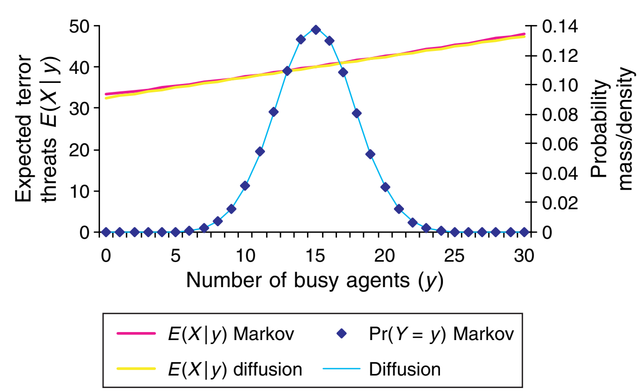

3. Ornstein-Uhlenbeck Terror Queue

While Equations (3)–(9) can be solved numerically to obtain results for any particular set of parameter values, the approximate normality shown in Figures 3 and 4 (and the linear regression curve for $E(X \mid Y = y)$) suggest that a more direct approach is possible. The terror queue model of Equations (3)–(9) is an example of a Markov population process, which is the multivariate generalization of a birth-death process (Kingman 1969, McNeil and Schach 1973). A considerable literature exists regarding the approximation of Markov population processes by diffusion models centered on the deterministic differential equations associated with the process (e.g., Barbour 1976; Kurtz 1970, 1981; McNeil and Schach 1973). This section proposes such an approximation that proves much simpler to implement than the original Markov model while providing nearly identical numerical results. The result will be a bivariate Ornstein-Uhlenbeck diffusion approximation to the original Markov terror queue model (note that univariate Ornstein-Uhlenbeck processes have long played an important role in providing diffusion approximations to queueing models, e.g., Browne and Whitt 1995, Halfin and Whitt 1981, Garnett et al. 2002). The specific methods employed below follow Barbour (1976), and have also been applied to the study of stochastic epidemic models by Nåsell (2002).

To begin, consider a deterministic model for Figure 1: letting $x(t)$ and $y(t)$ denote undetected and detected terror plots at time t, we can model these quantities using differential equations as

$$\frac{dx(t)}{dt} = \alpha - \mu x(t) - \delta x(t)(f - y(t)) \tag{11}$$

and

$$\frac{dy(t)}{dt} = \delta x(t)(f - y(t)) - \rho y(t) \tag{12}$$

subject to initial conditions $x(0)$ and $y(0)$. Our interest is in the steady state, so letting x^* and y^* denote the limiting values of $x(t)$ and $y(t)$ as $t \to \infty$ (and $dx(t)/dt$ and $dy(t)/dt$ approach zero), we obtain

$$\alpha = \mu x^* + \delta x^*(f - y^*) \tag{13}$$

and

$$\delta x^*(f - y^*) = \rho y^* \tag{14}$$

as identifying equations for the equilibrium solution. These equations solve to yield

$$x^* = \frac{(\alpha - \rho f - \mu\rho/\delta) + \sqrt{(\alpha + \rho f + \mu\rho/\delta)^2 - 4\alpha\rho f}}{2\mu} \tag{15}$$

and

$$y^* = \frac{(\alpha + \rho f + \mu\rho/\delta) - \sqrt{(\alpha + \rho f + \mu\rho/\delta)^2 - 4\alpha\rho f}}{2\rho}. \tag{16}$$

Table 1. Conditional joint probability distribution of the jumps in undetected ($\Delta X(t)$) and detected ($\Delta Y(t)$) terror plots in the terror queue model.

	$\Delta Y(t) = -1$	$\Delta Y(t) = 0$	$\Delta Y(t) = +1$
$\Delta X(t) = -1$	0	$\mu x \Delta t$	$\delta x(f - y)\Delta t$
$\Delta X(t) = 0$	$\rho y \Delta t$	$1 - (\alpha + \mu x + \rho y + \delta x(f - y))\Delta t$	0
$\Delta X(t) = +1$	0	$\alpha \Delta t$	0

Now we will construct a stochastic diffusion model. Let $X(t)$ and $Y(t)$ denote the (random) number of undetected and detected terror plots, and define $\Delta X(t)$ ($\Delta Y(t)$) as $X(t + \Delta t) - X(t)$ ($Y(t + \Delta t) - Y(t)$). Following Figure 2, the conditional joint probability distribution of $\Delta X(t)$ and $\Delta Y(t)$ given that $X(t) = x$ and $Y(t) = y$ is shown in Table 1.

Working with this joint conditional distribution, the local drift terms satisfy

$$E(\Delta X(t) \mid x, y) = (\alpha - \mu x - \delta x(f - y))\Delta t \tag{17}$$

and

$$E(\Delta Y(t) \mid x, y) = (\delta x(f - y) - \rho y)\,\Delta t. \tag{18}$$

Rather than work directly with these terms, which would require solving a nonlinear diffusion, following Barbour (1976) we presume that in steady-state, both $X(t)$ and $Y(t)$ assume values close to the deterministic equilibria derived in Equations (15)–(16) above and linearize the drift terms via a first order Taylor expansion. Defining the Jacobian matrix A as

$$A = \begin{pmatrix} \dfrac{d^2x}{dt\,dx} & \dfrac{d^2x}{dt\,dy} \\ \dfrac{d^2y}{dt\,dx} & \dfrac{d^2y}{dt\,dy} \end{pmatrix}_{x=x^*,\, y=y^*}$$

$$= \begin{pmatrix} -\mu - \delta(f - y^*) & \delta x^* \\ \delta(f - y^*) & -\delta x^* - \rho \end{pmatrix}, \tag{19}$$

we approximate the local drift terms by

$$\begin{pmatrix} E(\Delta X(t) \mid x, y) \\ E(\Delta Y(t) \mid x, y) \end{pmatrix} \approx A \begin{pmatrix} x - x^* \\ y - y^* \end{pmatrix} \Delta t, \tag{20}$$

which are linear in x and y.

We next construct the local covariance matrix $S\Delta t$ of $\Delta X(t)$ and $\Delta Y(t)$. Again following the conditional joint distribution of Table 1 (and retaining terms only up to order Δt), we obtain

$$S\Delta t$$

$$= \begin{pmatrix} \mathrm{Var}(\Delta X(t) \mid x, y) & \mathrm{Cov}(\Delta X(t), \Delta Y(t) \mid x, y) \\ \mathrm{Cov}(\Delta X(t), \Delta Y(t) \mid x, y) & \mathrm{Var}(\Delta Y(t) \mid x, y) \end{pmatrix}$$

$$= \begin{pmatrix} \alpha + \mu x + \delta x(f - y) & -\delta x(f - y) \\ -\delta x(f - y) & \delta x(f - y) + \rho y \end{pmatrix} \Delta t. \tag{21}$$

However, our interest is in the steady state; following Barbour (1976) we approximate S by S^*, which follows (21) evaluated at x^* and y^* of Equations (15)–(16).

Together, Equations (20) and (21) lead to a bivariate Ornstein-Uhlenbeck diffusion process (Barbour 1976, Gardiner 2004, McNeil and Schach 1973), which results in a bivariate normal distribution for the steady-state numbers of undetected (X) and detected (Y) terror plots. The means $E(X)$ and $E(Y)$ are given by the deterministic equilibrium values x^* and y^*, while the stable covariance matrix

$$\Sigma = \begin{pmatrix} \mathrm{Var}(X) & \mathrm{Cov}(X,Y) \\ \mathrm{Cov}(X,Y) & \mathrm{Var}(Y) \end{pmatrix} \tag{22}$$

is the solution to the algebraic equation (Barbour 1976, Gardiner 2004)

$$A\Sigma + \Sigma A^T = -S^*. \tag{23}$$

As a consequence of the multivariate normality of X and Y, the conditional distribution of the number of undetected plots X given y busy undercover agents is also normal, with mean following the well-known regression equation (e.g., Freund 1971, §13.2)

$$E(X \mid Y=y) = E(X) + \frac{\mathrm{Cov}(X,Y)}{\mathrm{Var}(Y)}(y - E(Y)) \tag{24}$$

and corresponding variance

$$\mathrm{Var}(X \mid Y=y) = \mathrm{Var}(X)(1-\mathrm{Corr}^2(X,Y)) \tag{25}$$

where $\mathrm{Corr}(X,Y)$ is the correlation of X and Y. Of particular interest is the slope in Equation (24), which can be shown to equal

$$\frac{\mathrm{Cov}(X,Y)}{\mathrm{Var}(Y)} = \left(\frac{f}{y^*} + \frac{\alpha}{\delta x^{*2}} - \frac{\rho y^*}{\alpha}\right)^{-1}. \tag{26}$$

As an example, consider the same parameters used earlier to illustrate the Markov terror queue: $\alpha = 100$; $\mu = 1$; $\delta = 0.1$; $\rho = 4$; and $f = 30$. The probability density resulting from the diffusion process is virtually indistinguishable from the joint distribution obtained from solving the Markov balance Equations (3)–(9). In particular, consider Figure 4 which shows the marginal distributions for the number of detected terror plots (busy intelligence agents) and the conditional expected number of undetected terror plots as a function of the number of busy agents for both the Markov and diffusion terror queue models. The results are essentially identical, as further indicated by Table 2, which reports the values of several key quantities for both models.

3.1. Boundary Approximations

Note that in deriving the diffusion approximation we have ignored the boundary conditions at $X = 0$, $Y = 0$, and $Y = f$. For example, one would not expect the distribution of the number of detected terror plots (busy agents) Y to be approximately normal if y^* is close to 0 or f, nor would Equation (24) hold. However, if y^* is close to an extreme, a simple approximation for the distribution of the number of undetected terror plots X can be found by setting $Y = 0$ or f depending upon whether the undercover agents are almost always idle or busy, which produces a univariate probability model for X. Working directly from the flows in Figure 1, if $Y \approx 0$ then X can be approximated by an infinite server queue with arrival rate α and service rate $\mu + \delta f$, which would result in X having a Poisson distribution with mean (and variance) equal to

$$E(X) = \frac{\alpha}{\mu + \delta f}. \tag{27}$$

Alternatively, if $Y \approx f$, then X can be modeled by a saturated $M/M/1$ queue with arrival rate α, an always busy server with service rate ρf, and *reneging rate* μ per customer (in this case per undetected terror plot). This would result in

$$E(X) \approx \frac{\alpha - \rho f}{\mu} \tag{28}$$

and

$$\mathrm{Var}(X) \approx \frac{\alpha}{\mu} \tag{29}$$

with probability distribution following that of a birth-death process (Kaplan 1987). Better approximations for the distribution of Y (and associated conditional distribution of X) could be derived when the distribution of Y is concentrated near 0 or f (Gardiner 2004, Browne and Whitt 1995), though such refinements are beyond the purpose of this paper.

Table 2. Comparing the Markov and diffusion terror queue models.

Measure	Markov	Diffusion
$E(X)$	40.12	40
$E(Y)$	14.97	15
$\mathrm{Var}(X)$	46.70	46.53
$\mathrm{Var}(Y)$	8.23	8.27
$\mathrm{Cov}(X,Y)$	4.07	4.08
$\mathrm{Corr}(X,Y)$	0.2077	0.2081
$\frac{\mathrm{Cov}(X,Y)}{\mathrm{Var}(Y)}$	0.4946	0.4938

3.2. Transient Behavior and Sensitivity Analysis

For the steady state terror queue model to prove useful, the time required to converge to steady state conditions should not be too large. While a comprehensive analysis of the transient behavior of this model is beyond the scope of this paper, some insights about the transient behavior can be illustrated by a combination of analysis and simulation. Consider first the deterministic model given by the differential Equations (11)–(12). By inspection of Equation (11), it is clear that lower and upper bounds for the instantaneous "departure" rate for undetected plots are given by μx and $(\mu + \delta f)x$, respectively. Thus, in the deterministic model, the number of undetected plots x will converge to its steady state value exponentially fast at a rate bounded between μ and $\mu + \delta f$, which implies that the relaxation time constant for this model falls between $1/(\mu + \delta f)$ and $1/\mu$. More precision can be gained from studying the linearized version of this model, which can be written as

$$\begin{pmatrix} \dfrac{dx(t)}{dt} \\ \dfrac{dy(t)}{dt} \end{pmatrix} = A \begin{pmatrix} x(t) - x^* \\ y(t) - y^* \end{pmatrix} \tag{30}$$

where the Jacobian matrix A is given in Equation (19). As is well known, the rate at which such systems converge is governed by the eigenvalues of A (e.g., see Strang 1980, Chapter 5). While there is little to learn from perusing the symbolic formulas for these eigenvalues, their numerical values can be computed easily for any set of parameter values, and the reciprocal of the smallest absolute eigenvalue defines the relaxation time for the deterministic system providing both eigenvalues are negative (which is always the case for the Jacobian matrix of Equation (19)). To illustrate, for the parameters used in our example of Figures 3 and 4, we have

$$A = \begin{pmatrix} -2.5 & 4 \\ 1.5 & -8 \end{pmatrix} \tag{31}$$

with eigenvalues -1.57 and -8.93; thus, the relaxation time is given by $1/1.57 = 0.64$ years. Note that this estimate falls between the bounds stated earlier of $1/(\mu + \delta f) = 1/4$ and $1/\mu = 1$. As the gap between current and steady state values is reduced by the factor $1/e$ over each relaxation time, the deterministic model suggests that about $1 - 1/e$ or 63% of the distance between the starting and steady state values will have been erased by 0.64 years, while $1 - \exp(-3)$, or about 95% of the initial distance from steady state, will have been erased by three relaxation times or 1.92 years.

To see how well this approximate analysis applies to the original Markov terror queue, simulation was used to generate realizations of the number of undetected $(X(t))$ and detected $(Y(t))$ terror plots over time. Figure 5 displays a typical realization for the parameters of our working example with $X(0) = \alpha/\mu = 100$ and $Y(0) = 0$; these initial conditions correspond to the deployment of intelligence agents at time 0 in an environment where the number of unchecked terror plots had reached its equilibrium level of 100. Also displayed are the results of the deterministic differential Equations (11)–(12). Not only does the

Figure 5. One realization of simulated undetected and detected terror plots from the Markov terror queue model, and associated fluid model (Equations (11)–(12)) predictions: undetected (black) and detected (gray) terror plots.

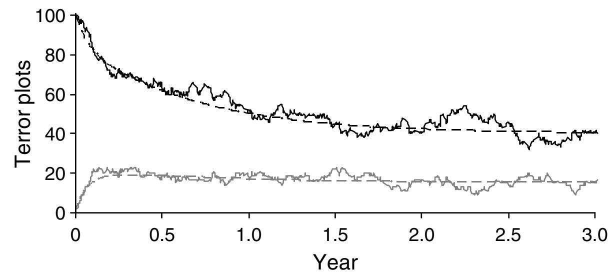

observed time required to reach steady state accord well with the predictions based on the approximate analysis of the deterministic equations, the simulated values track the deterministic model remarkably well. These results, namely excellent prediction of the relaxation time and excellent tracking of the deterministic equations by stochastic realizations, have been obtained for many different sets of parameter values for this model.

The analysis reported above focused on transient behavior given the model parameters, yet at least some of these parameters must be estimated from data. In particular, the terror plot arrival rate α must be estimated, while this same parameter might change over time, perhaps in response to political conditions. To examine the sensitivity of the model's results to such issues, the following simulation experiment was performed. In each of 1,000 replications of the experiment, the model was first simulated for three years with $\alpha = 100$ as in Figure 5 above. At the start of year 4, α was then shifted instantaneously to a new value α'. At the end of year 4, an estimate $\hat{\alpha}$ for the terror plot arrival rate α was formed by counting all successful attacks and detected terror plots observed during year 4. The idea is that unbeknownst to the intelligence agency, α has changed. The intelligence agency's estimate $\hat{\alpha}$ presumes that steady state conditions prevail, for as shown in Equation (1), α does equal the sum of the expected annual number of successful and detected terror plots in steady state. The intelligence agency then estimates the expected number of undetected terror plots at the end of year 4 via Equation (15) for x^* using $\hat{\alpha}$. The resulting value of x^* was then compared to the simulated number of terror plots at the end of year 4, $X(4)$, to see the magnitude of the error as measured by both the mean absolute error $\overline{|X(4) - x^*|}$ and the root mean squared error $\sqrt{\overline{(X(4) - x^*)^2}}$ (where the overbar denotes the average over the 1,000 simulation runs). The results are shown in Table 3.

Consider first the case of $\alpha' = 100$, which corresponds to the continuation of steady state conditions. However, instead of utilizing the true terror plot arrival rate of $\alpha = 100$, this experiment shows what happens when the intelligence agency must estimate this rate from observed data. When

Table 3. Simulation experiment results: The sensitivity of undetected terror plot prediction error to unobserved changes in the terror plot arrival rate.

α'	$\bar{\hat{\alpha}}$	$\bar{X}(4)$	$\bar{x}^*$	Mean absolute error	Root mean squared error
50	72.0	19.0	25.5	7.1	8.5
60	77.9	23.0	28.4	6.6	8.1
70	83.4	27.1	31.1	6.1	7.6
80	89.2	31.3	34.2	6.4	8.0
90	95.3	35.7	37.5	6.4	7.9
100	100.8	40.1	40.6	6.5	8.1
110	105.4	45.0	43.3	7.1	8.8
120	110.1	50.0	46.6	7.6	9.6
130	116.1	54.6	49.8	8.3	10.4
140	121.1	59.1	53.0	9.3	11.6
150	125.4	64.5	55.8	10.9	13.5

$\alpha = 100$, the steady state number of successful plus detected terror plots will have a Poisson distribution with a mean of 100 and a standard deviation of 10; thus, roughly 32% of the 1,000 replications correspond to samples where $\hat{\alpha}$ overestimates or underestimates α by at least 10%. Table 3 shows that, as expected in this case, $\hat{\alpha}$ is an unbiased estimator of α (with $\bar{\hat{\alpha}} = 100.8$), steady state has been reached ($\bar{X}(4) = 40.1$, recall that in steady state $E(X) = 40$ for this example), and x^* is also unbiased (with $\bar{x}^* = 40.6$). However, the Root Mean Squared Error of 8.1 is larger than the standard deviation $\sigma_X \approx \sqrt{46.53} = 6.8$, while the Mean Absolute Error of 6.5 is larger than the Mean Absolute Deviation associated with a normally distributed random variable of $\sigma_X\sqrt{2/\pi} \approx 5.4$. These figures show that, when the steady state assumptions underlying the terror queue model hold in this example, using $\hat{\alpha}$ instead of α slightly increases the error in estimating the unknown number of terror plots X by x^*.

Now consider what happens as α' moves away from 100. The value of $\hat{\alpha}$ clearly lags α', overestimating the terror plot arrival rate when $\alpha' < 100$ and underestimating it when $\alpha' > 100$. This is not surprising, considering the discussion of relaxation times above. However, the absolute error associated with estimating $X(4)$ is surprisingly insensitive to α' providing $\alpha' < 120$ (though the relative error of course increases as α' declines). This result can be understood as follows: while the estimated arrival rate $\hat{\alpha}$ lags the true arrival rate α', the steady state value x^* associated with $\hat{\alpha}$ *overshoots* the value of $X(4)$ that would be associated with $\hat{\alpha}$ (as $X(4)$ is recorded only one year after the arrival rate has shifted while x^* is the limiting value). These two effects move in opposite directions and keep the absolute error in estimating the number of undetected terror plots in check.

4. Ornstein-Uhlenbeck Terror Queue with False Detection

In the model of the previous section, while undercover intelligence agents can fail to detect extant terror plots (the detection rate δ is finite after all), all detected plots are presumed real. However, intelligence agents can make mistakes and become occupied with the surveillance of individuals or groups who in fact are not attached to existing terror plots. Indeed, given that terrorists are almost certainly aware of intelligence efforts to detect planned attacks, generating disinformation with the goal of diverting intelligence efforts from real to fake plots is good strategy for terrorists to employ (see Steele 1989 for a model of disinformation to guard secrets). To incorporate this important feature of the intelligence operations environment, we now presume that intelligence agents become busy via both the true detection of real and the false detection of fake terror plots. In this new model, X will continue to denote the number of undetected true terror plots in progress, while Y and Z will denote the number of intelligence agents busy investigating cases triggered by true and false detections, respectively. Unfortunately for the intelligence agency, while an agent is occupied it is not possible to know whether (s)he is observing a real plot or a fake one. However, it remains possible to observe the number of busy agents $Y + Z$, which in turn allows inference regarding the extant number of undetected terror plots.

We will retain the assumption that available intelligence agents detect terror plots in proportion to their number; thus, the aggregate number of actual plots detected per unit time equals $\delta X(f - Y - Z)$. To this we add the assumption that available intelligence agents are sidetracked by fake plots at per agent rate λ, so that the number of available agents removed from active surveillance by fake plots (whether due to deception or detection error) is equal to $\lambda(f - Y - Z)$ per unit time. As before, we presume that the time to interdict a real terror plot is exponentially distributed with rate ρ; in addition, we assume that the time required for agents following a fake plot to recognize their error and abandon surveillance is exponentially distributed with rate ψ. We retain all other assumptions from the model of the previous section and will now proceed to develop the associated trivariate Ornstein-Uhlenbeck terror queue model.

We begin by writing the differential equations for the deterministic trivariate model. Letting $x(t)$, $y(t)$, and $z(t)$ denote the number of undetected plots, busy agents occupied with real plots, and busy agents occupied with fake plots, respectively, we have

$$\frac{dx(t)}{dt} = \alpha - \mu x(t) - \delta x(t)(f - y(t) - z(t)) \tag{32}$$

$$\frac{dy(t)}{dt} = \delta x(t)(f - y(t) - z(t)) - \rho y(t) \tag{33}$$

$$\frac{dz(t)}{dt} = \lambda(f - y(t) - z(t)) - \psi z(t) \tag{34}$$

and the associated steady-state equations

$$\alpha = \mu x^* + \delta x^*(f - y^* - z^*) \tag{35}$$

$$\delta x^*(f - y^* - z^*) = \rho y^* \tag{36}$$

$$\lambda(f - y^* - z^*) = \psi z^*. \tag{37}$$

Solving Equation (37) for z^* yields

$$z^* = \frac{\lambda}{\lambda+\psi}(f - y^*), \tag{38}$$

which implies that the steady state number of available intelligence agents is given by

$$f - y^* - z^* = \frac{\psi}{\lambda+\psi}(f - y^*). \tag{39}$$

Defining

$$\delta' = \delta\frac{\psi}{\lambda+\psi}, \tag{40}$$

Equations (35)–(36) reduce to Equations (13)–(14) after substituting δ' for δ, and consequently x^* and y^* are given by Equations (15)–(16) following the same substitution. The impact of introducing false detections on the steady state numbers of undetected and detected (true) terror plots is thus equivalent to reducing the detection rate δ by the factor $\psi/(\lambda+\psi) < 1$.

Next we determine the Jacobian matrix A used to linearize the diffusion model. Differentiating the right-hand sides of Equations (32)–(34) yields

$$A = \begin{pmatrix} \dfrac{d^2x}{dtdx} & \dfrac{d^2x}{dtdy} & \dfrac{d^2x}{dtdz} \\ \dfrac{d^2y}{dtdx} & \dfrac{d^2y}{dtdy} & \dfrac{d^2y}{dtdz} \\ \dfrac{d^2z}{dtdx} & \dfrac{d^2z}{dtdy} & \dfrac{d^2z}{dtdz} \end{pmatrix}_{x=x^*,\, y=y^*,\, z=z^*}$$
$$= \begin{pmatrix} -\mu-\delta(f-y^*-z^*) & \delta x^* & \delta x^* \\ \delta(f-y^*-z^*) & -\delta x^*-\rho & -\delta x^* \\ 0 & -\lambda & -\lambda-\psi \end{pmatrix}. \tag{41}$$

Defining $\Delta X(t), \Delta Y(t)$, and $\Delta Z(t)$ as in the bivariate diffusion model, the linear approximation for the local drift terms is given by

$$\begin{pmatrix} E(\Delta X(t)) \\ E(\Delta Y(t)) \\ E(\Delta Z(t)) \end{pmatrix} = A \begin{pmatrix} x - x^* \\ y - y^* \\ z - z^* \end{pmatrix} \Delta t. \tag{42}$$

The joint probability distribution of $\Delta X(t), \Delta Y(t)$, and $\Delta Z(t)$ is displayed in Table 4.

From this distribution we obtain the local covariance matrix $S\Delta t$ (up to terms of order Δt) as

$$S\Delta t = \begin{pmatrix} \alpha+\mu x+\delta x(f-y-z) & -\delta x(f-y-z) & 0 \\ -\delta x(f-y-z) & \delta x(f-y-z)+\rho y & 0 \\ 0 & 0 & \lambda(f-y-z)+\psi z \end{pmatrix}\Delta t \tag{43}$$

Table 4. Conditional joint probability distribution of the jumps in undetected ($\Delta X(t)$), detected real ($\Delta Y(t)$), and detected fake ($\Delta Z(t)$) terror plots in the terror queue model.

Event	Probability
$\Delta X(t)=+1,\ \Delta Y(t)=0,\ \Delta Z(t)=0$	$\alpha\Delta t$
$\Delta X(t)=-1,\ \Delta Y(t)=0,\ \Delta Z(t)=0$	$\mu x\Delta t$
$\Delta X(t)=0,\ \Delta Y(t)=-1,\ \Delta Z(t)=0$	$\rho y\Delta t$
$\Delta X(t)=-1,\ \Delta Y(t)=+1,\ \Delta Z(t)=0$	$\delta x(f-y-z)\Delta t$
$\Delta X(t)=0,\ \Delta Y(t)=0,\ \Delta Z(t)=+1$	$\lambda(f-y-z)\Delta t$
$\Delta X(t)=0,\ \Delta Y(t)=0,\ \Delta Z(t)=-1$	$\psi z\Delta t$
$\Delta X(t)=0,\ \Delta Y(t)=0,\ \Delta Z(t)=0$	$1-(\alpha+\mu x+\rho y+(\delta x+\lambda) \cdot (f-y-z)+\psi z)\Delta t$

which we approximate with S^*, which is S evaluated at the deterministic equilibria x^*, y^*, and z^*.

Together, Equations (42)–(43) provide a trivariate Ornstein-Uhlenbeck diffusion model for the steady state number of undetected, true, and fake terror plots X, Y, and Z. The associated multivariate normal distribution has means x^*, y^*, and z^* and covariance matrix Σ satisfying Equation (23) using A and S^* as defined in Equations (41)–(43). Inference regarding the number of extant undetected terror plots in this model must rely on the observed number of busy intelligence agents $y+z$, for it is not possible to determine y and z individually. However, since Y and Z are both normally distributed, their sum is also normally distributed, which means that the conditional distribution of X given that $Y+Z=b$ busy agents is normally distributed with mean

$$E(X \mid Y+Z=b) = E(X) + \frac{\mathrm{Cov}(X, Y+Z)}{\mathrm{Var}(Y+Z)}(b - E(Y) - E(Z)) \tag{44}$$

and variance

$$\mathrm{Var}(X \mid Y+Z=b) = \mathrm{Var}(X)(1-\mathrm{Corr}^2(X, Y+Z)). \tag{45}$$

Use of Equations (44)–(45) requires recognizing that

$$\mathrm{Cov}(X, Y+Z) = \mathrm{Cov}(X, Y) + \mathrm{Cov}(X, Z) \tag{46}$$

and

$$\mathrm{Var}(Y+Z) = \mathrm{Var}(Y) + \mathrm{Var}(Z) + 2\mathrm{Cov}(Y, Z) \tag{47}$$

where the variance and covariance terms appearing on the right-hand sides of Equations (46)–(47) are the appropriate elements of the covariance matrix Σ.

One can go further—since Y is not observable, it is of interest to estimate the *total* number of true terror plots $X+Y$ conditional upon the observed number of busy agents $Y+Z$. Again, due to the trivariate normality of X, Y, and Z in the diffusion model, the conditional distribution of the total number of true terror plots will be normally distributed with mean

$$E(X+Y\mid Y+Z=b)$$
$$=E(X+Y)+\frac{\mathrm{Cov}(X+Y,Y+Z)}{\mathrm{Var}(Y+Z)}(b-E(Y)-E(Z)) \quad (48)$$

and variance

$$\mathrm{Var}(X+Y\mid Y+Z=b)$$
$$=\mathrm{Var}(X+Y)(1-\mathrm{Corr}^2(X+Y,Y+Z)) \quad (49)$$

where

$$\mathrm{Cov}(X+Y,Y+Z)$$
$$=\mathrm{Cov}(X,Y)+\mathrm{Cov}(X,Z)+\mathrm{Cov}(Y,Z)+\mathrm{Var}(Y) \quad (50)$$

and

$$\mathrm{Var}(X+Y)=\mathrm{Var}(X)+\mathrm{Var}(Y)+2\mathrm{Cov}(X+Y). \quad (51)$$

4.1. Trivariate Boundary Approximations

Thus far, our analysis of the trivariate terror queue model has ignored boundary conditions, of which the most important are that the total number of occupied intelligence agents $Y+Z$ must fall between 0 and f inclusive. In the uninteresting case where all agents are almost always idle, the same "light traffic" approximation proposed in Equation (27) can be applied, and the number of undetected terror plots X would follow a Poisson distribution with mean $\alpha/(\mu+\delta f)$. The more interesting case is when all agents are almost always busy, that is, $Y+Z\approx f$. If $Y+Z=f$, then one can focus solely on the number of agents Y who are following real threats. Given that $X=x$ and $Y=y$, the number of agents following real threats will increase only if an agent following a *fake* threat completes service (this occurs with rate $\psi(f-y)$) and is then immediately occupied by a *real* threat (this occurs with probability $\delta x/(\delta x+\lambda)$), thus the overall rate with which $Y\to Y+1$ given $X=x$ and $Y=y$ equals $\psi(f-y)\delta x/(\delta x+\lambda)$. Similarly, the number of agents following real threats will only decrease if an agent following a *real* threat completes service (this occurs with rate ρy) and is then immediately occupied by a *fake* threat (with probability $\lambda/(\delta x+\lambda)$), implying that the overall rate with which $Y\to Y-1$ given $X=x$ and $Y=y$ equals $\rho y\lambda/(\delta x+\lambda)$. Finally, again given that $X=x$ and $Y=y$, the number of undetected terror plots X increases whenever a new plot arrives (with rate α), and declines with successful attacks (rate μx) or whenever a busy intelligence agent completes service (with rate $\rho y+\psi(f-y)$) and immediately detects a real plot (with probability $\delta x/(\delta x+\lambda)$).

These arguments enable the construction of an approximating bivariate Ornstein-Uhlenbeck diffusion. The deterministic differential equations for this approximation follow

$$\frac{dx(t)}{dt}=\alpha-\mu x(t)-\frac{\delta x(t)}{\delta x(t)+\lambda}(\rho y(t)+\psi(f-y(t))) \quad (52)$$

and

$$\frac{dy(t)}{dt}=\psi(f-y(t))\frac{\delta x(t)}{\delta x(t)+\lambda}-\rho y(t)\frac{\lambda}{\delta x(t)+\lambda}. \quad (53)$$

After defining

$$\delta''=\delta\frac{\psi}{\lambda} \quad (54)$$

and setting the left-hand sides of Equations (52)–(53) to zero, the resulting equilibrium values x^* and y^* are again given by Equations (15)–(16) upon substituting δ'' for δ. Following the same derivations used previously, the Jacobian matrix A and local covariance matrix S^* are given by

$$A=\begin{pmatrix} -\mu-(\rho y^*+\psi(f-y^*))\dfrac{\delta\lambda}{(\delta x^*+\lambda)^2} & (\psi-\rho)\dfrac{\delta x^*}{\delta x^*+\lambda} \\ (\rho y^*+\psi(f-y^*))\dfrac{\delta\lambda}{(\delta x^*+\lambda)^2} & -\dfrac{(\rho\lambda+\psi\delta x^*)}{\delta x^*+\lambda} \end{pmatrix} \quad (55)$$

and

$$S^*=\begin{pmatrix} \alpha+\mu x^*+\rho y^* & -\psi(f-y)\dfrac{\delta x^*}{\delta x^*+\lambda} \\ -\psi(f-y)\dfrac{\delta x^*}{\delta x^*+\lambda} & \psi(f-y^*)\dfrac{\delta x^*}{\delta x^*+\lambda}+\rho y^*\dfrac{\lambda}{\delta x^*+\lambda} \end{pmatrix}. \quad (56)$$

Applying Equation (23) to A and S^* above yields the steady-state covariance matrix Σ for this approximation. To summarize, if the total number of busy intelligence agents $Y+Z\approx f$, then the joint distribution of X and Y is approximately bivariate normal with means x^* and y^* and covariance matrix as derived above.

4.2. Example: Suicide Bombings in Israel

Purely as an illustrative example, consider suicide bombing attacks targeting Israelis during the Second Intifada. The Israel Security Agency (also known as the Shin Bet or the Shabak, http://www.shabak.gov.il/english) reports that between 2000 and 2007 inclusive there were an average of $\mu x^*=19$ successful suicide bombings annually (Israel Security Agency 2008), while unpublished data from the same agency suggest the interdiction of $\rho y^*=66$ suicide bombings annually, implying an average total attack rate of $\alpha=85$ per year and an interdiction rate $\rho y^*/\alpha=77.6\%$. The number of undercover intelligence cells f deployed to detect and interdict terrorist attacks is, of course, classified information. Cordesman (2006) reports that special units operating within the 8,000-strong Border Guard (Mishmar Ha'gvul or MAGAV, itself a unit of the Israel Police) are responsible for counterterror operations, including the 100+ member YAMAM counterterror and hostage release unit, and the YAMAS undercover unit that specializes in covert operations in the West Bank (Deflem 2010). As these latter units act on the basis of intelligence information, within our model they can be thought of as the *means* by which intelligence agents interdict terror plots rather than the intelligence agents themselves (who would correspond to Shabak officers and their local informants).

Since our purpose here is to illustrate the terror queue model rather than claim accurate representation of Israel's counterterror intelligence efforts, we will simply assume that there are $f = 50$ undercover "teams" of agents and informants who specialize in the discovery and interdiction of suicide bombing attacks. Given the apparent ease of constructing suicide bombs and other improvised explosive devices (National Research Council 2007) and the ready availability of would-be "martyrs" willing to carry out suicide terror attacks (Al Hajaar 2004), the bottleneck in attack planning relates more to operational issues such as target selection and transporting the bomb and the bomber while evading detection and capture by Israeli security forces. We thus assume that absent interdiction, the mean time required to plan and implement an attack $1/\mu = 3$ months (thus $\mu = 4$ per year, which implies an average of $x^* = \mu x^*/\mu = 19/4 = 4.75$ undetected terror plots). Presumably the time required to interdict a suspected terrorist is much shorter, though such interdiction is not necessarily instantaneous given the opportunity to collect further incriminating evidence once a planned attack has been discovered covertly. Absent any data suggesting different times to interdict real versus fake plots, we set $\rho = \psi = 16$ per year (making the mean time to interdict a suspected terror plot 25% of the mean time required to plan and execute such plots), implying that $y^* = \rho y^*/\rho = 66/16 = 4.125$ intelligence teams occupied with real threats.

Finally, given the severity of suicide bombing attacks (and hence the high cost of missed true detections) combined with terrorists' incentives to occupy counterterror agents with fake threats, we expect many false positive detections. To reflect this, we assume that only 10% of all terror plots interdicted are real, that is, we assume that $\rho y^*/(\rho y^* + \psi z^*) = 0.1$. This assumption sets $66/(66+16z^*) = 0.1$ which identifies $z^* = 37.125$, while Equations (36)–(37) then determine $\delta = 1.588$ and $\lambda = 67.886$.

With these parameters, application of Equations (41), (43), and (23) imply that X, Y, and Z have a trivariate normal distribution with means 4.75, 4.125, and 37.125, respectively, and covariance matrix

$$\Sigma = \begin{pmatrix} 4.974 & 0.153 & 0.380 \\ 0.153 & 3.878 & -3.073 \\ 0.380 & -3.073 & 9.568 \end{pmatrix}. \tag{57}$$

Given the observation that $Y + Z = b$ busy intelligence cells, the number of true terror plots $X + Y$ is normally distributed with mean

$$E(X+Y \mid Y+Z=b) = 1.32 + 0.183b \tag{58}$$

and variance

$$\mathrm{Var}(X+Y \mid Y+Z=b) = 8.913. \tag{59}$$

Since $\mathrm{Var}(Y+Z) = 7.3$ in this example, swings of $\pm 2\sigma_{Y+Z}$ correspond to swings of ± 1 real terror plot in expectation.

For the example above, ignoring the boundary conditions should not cause alarm: the mean number of undetected attacks is more than two standard deviations away from its boundary at zero, while the mean number of busy intelligence teams is more than three standard deviations away from its boundary at $f = 50$. To illustrate the approximation provided by Equations (52)–(56), suppose that $\alpha = 2{,}500$ while all other parameters (δ, ρ, λ, ψ, μ, and f) remain as above. The trivariate model in this case yields mean values x^*, y^*, and z^* of roughly 446, 45, and 4 respectively, while $\mathrm{Var}(Y+Z) \approx 1$, which means that the mean number of busy agents $y^* + z^*$ is only one standard deviation away from the total number of agents $f = 50$. Applying the "heavy traffic" approximation of Equations (52)–(56) to this example yields $x^* = 442.6$, $y^* = 45.6$, and $\mathrm{Var}(Y) = 4.0$, making y^* more than two standard deviations away from the boundary at $f = 50$. Note that since all agents are busy all of the time in this approximation, all that can be inferred is that the number of undetected terror plots X is normally distributed with mean 442.6 and variance 603.1, while the total number of real terror plots $X + Y$ is normally distributed with mean 488.2 and variance 609.5. Were α to continue increasing in this example, y^* would tend towards $f = 50$ and the approximation of Equations (28)–(29) would apply.

5. Discussion

Estimating the extant number of terror plots in a given area is of fundamental importance as public officials consider the allocation of counterterror resources. This article has proposed a new approach to this problem that relies on the relationship between the number of terror plots and the utilization of intelligence agents. The resulting terror queue model enables both unconditional estimation of the mean number of terror plots underway, and conditional estimation of the same given the observed number of busy undercover agents at a point in time, while also allowing one to estimate the rate with which true plots can be detected and interdicted.

Since data pertaining to covert counterterror surveillance and interdiction operations are classified (and for good reason), the examples presented in this paper were necessarily hypothetical. However, within agencies running undercover counterterror intelligence operations such as police departments, the FBI, MI5 or the Shabak, perhaps the necessary parameters can be estimated from undercover agent/informant activity logs and case resolution statistics. For example, activity logs reporting when agents initially detect and attach themselves to terror suspects as well as when such cases are broken can be aggregated to determine the number of agents actively tracking suspects over time, while subsequent legal processing ultimately resolves whether those suspects tracked were real or fake. Such data enable estimating y^* and z^* as the observed average number of busy agents following (retrospectively determined) real or fake plots, and of ρ and ψ from the observed mean times from detection to interdiction of (again retrospectively determined) real and fake plots. The security agency involved knows the number of agents f that are fielded, which, together with data on agent utilization and case resolution, further determines λ from Equation (37). Over time, summing the rate of successful terror attacks with those interdicted produces an estimate of the total attack rate α (this is just Equation 1). Depending upon the nature of the terror attacks in question, the security agency is likely to have a good estimate of the time necessary to mount a successful attack, which specifies μ and in turn suggests the value of x^* (since $x^* = (\alpha - \rho y^*)/\mu$). Finally, once one has an estimate of x^*, the detection rate δ can equivalently be computed from Equation (35) or (36).

Beyond estimating the number of terror plots, terror queue models suggest a new approach to addressing some decision-making problems in counterterrorism. For example, given the social costs of successful terror attacks and the operational costs of fielding undercover agents, what is the optimal number of covert agents to deploy in counterterror surveillance efforts when the goal is to maximize the interdiction of real terror threats? As another example, suppose that new surveillance technology allows for improving the true detection rate δ at some cost. Should the new technology be adopted? Or, are there actions the security agency can take to change the terrorists' operating environment such as increasing the duration of time necessary to plan and carry out an attack (i.e., reduce μ which raises the interdiction rate ρy^* as is clear from Figure 1; see Weisburd et al. 2009 for a discussion of Israel's use of checkpoints in the West Bank in this regard). If so, what is the optimal investment in such actions?

Finally, the model as presented assumes steady-state conditions and is thus more suited to situations where the level of terrorism has remained reasonably constant. Modifying the model to allow for more transient situations, whether due to exogenous changes in the level of terrorism, or perhaps endogenous changes in terror tactics in response to successful interdictions (e.g., increasing the false detection rate λ by disseminating false information as suggested by Steele 1989), remains a future research challenge.

Acknowledgments

The author gratefully acknowledges the Daniel Rose Fund supporting the Technion-Yale Initiative in Homeland Security and Counterterror Operations Research.

References

Al Hajaar, Y. 2004. Hamas arch-terrorist Muhammad Deif—An interview. *Jerusalem Post* (March 25). http://tinyurl.com/y7u9gp.

Andraeus, T. 2009. Interview with Tawfiq al Tirawi. *Hadath* **95** 10–11. [In Arabic, http://www.hadath.co.il/volumes/Volume095.pdf].

Associated Press. 2009. Plot to blow up a building in Dallas. *New York Times* (September 25) A20.

Barbour, A. D. 1976. Quasi-stationary distributions in Markov population processes. *Adv. App. Probab.* **8** 296–314.

Blair, D. C. 2009. *The National Intelligence Strategy of the United States of America*. Office of the Director of National Intelligence, Washington, DC.

Browne, S., W. Whitt. 1995. Piecewise-linear diffusion processes. J. H. Dshalalow, ed. *Advances in Queueing: Theory, Methods and Open Problems*. CRC Press, Boca Raton, FL, 463–480.

Cordesman, A. H. 2006. *Arab-Israeli Military Forces in an Era of Asymmetric Wars*. Greenwood Publishing Group, Westport, CT.

Deflem, M. 2010. Undercover counterterrorism in Israel. *The Policing of Terrorism: Organizational and Global Perspectives*, Chapter 9. Routledge, New York.

Freund, J. E. 1971. *Mathematical Statistics*. Prentice Hall, Englewood Cliffs, NJ.

Gardiner, C. W. 2004. *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*. Springer-Verlag, Berlin.

Gardner, F. 2007. MI5 watch 2,000 terror suspects. *BBC News* (May 2). http://news.bbc.co.uk/2/hi/uk_news/6613963.stm.

Garnett, O., A. Mandelbaum, M. Reiman. 2002. Designing a call center with impatient customers. *Manufacturing Service Oper. Management* **4** 208–227.

Guiora, A. N. 2007. *Global Perspectives on Counterterrorism*. Aspen Publishers, New York.

Guiora, A. N. 2008. *Fundamentals of Counterterrorism*. Aspen Publishers, New York.

Halfin, S., W. Whitt. 1981. Heavy-traffic limits for queues with many exponential servers. *Oper. Res.* **29** 567–587.

Israel Security Agency. 2008. Palestinian terrorism in 2007: Statistics and trends. http://www.shabak.gov.il/English/EnTerrorData/Pages/2007.aspx.

Johnston, D., S. Shane. 2009. Terror case is called one of the most serious in years. *New York Times* (September 25) A1.

Kaplan, E. H. 1987. Analyzing tenant assignment policies. *Management Sci.* **33** 395–408.

Kaplan, E. H., M. Kress. 2005. Operational effectiveness of suicide-bomber-detector schemes: A best-case analysis. *Proc. Natl. Acad. Sci.* **102** 10399–10404.

Kaplan, E. H., A. Mintz, S. Mishal. 2006. Tactical prevention of suicide bombings in Israel. *Interfaces* **36**(6) 553–561.

Kingman, J. F. C. 1969. Markov population processes. *J. Appl. Probab.* **6** 1–18.

Kurtz, T. G. 1970. Solutions of ordinary differential equations as limits of pure jump Markov processes. *J. Appl. Probab.* **7** 49–58.

Kurtz, T. G. 1981. *Approximation of Population Processes*. SIAM, Philadelphia.

McNeil, D. R., S. Schach. 1973. Central limit analogues for Markov population processes. *J. Roy. Statist. Soc. B* **35** 1–15.

Nåsell, I. 2002. Stochastic models of some endemic infections. *Math. Biosciences* **179** 1–19.

National Research Council. 2007. *Countering the Threat of Improvised Explosive Devices: Basic Research Opportunities*. Committee on Defeating Improvised Explosive Devices: Basic Research to Interrupt the IED Delivery Chain, National Research Council. National Academies Press, Washington, DC.

Steele, J. M. 1989. Models for managing secrets. *Management Sci.* **35** 240–248.

Strang, G. 1980. *Linear Algebra and Its Applications*, 2nd ed. Academic Press, Inc., New York.

Weisburd, D., T. Jonathan, S. Perry. 2009. The Israeli model for policing terrorism: Goals, strategies, and open questions. *Criminal Justice Behavior* **36** 1259–1278.

Wilson, M. 2009. In Bronx bomb case, missteps caught on tape. *New York Times* (May 22) A1.