Y.L. Ligthart

Terror Queue: Help, we are being served!

Bachelor thesis

April 3, 2018

Thesis supervisor: Dr. F.M. Spieksma

Leiden University
Mathematical Institute

Contents

1 Introduction

Terrorism and counter terrorism has had a lot of attention the last couple of years. Edward Kaplan has done a lot of work regarding this subject. In 2010 he published a paper in which he developed a terror queue model [1]. In 2013 [2] Kaplan explored different staffing models for counter terrorism agents using his model. Kaplan's model is based on queueing theory. In the model, there is a group of agents looking for terror plots in preparation and when they find one, they attempt to interdict the plot.

In this work we will develop and study a tandem version of the terror queue model. This is an adaptation of Kaplan's terror queue model. We will first look at the base problem and Kaplan's model. After this we will describe our adaptation of his model, which doesn't require the agents to follow already detected plots. For this model we determine performance measures of interest like the stationary distribution and some related expectations. We also show that the tandem terror queue model translates to a more general version.

Finally, we will derive the optimal amount of agents required to fulfil certain optimality criteria such as cost minimization and prevention of a given fraction of terror plots. We will also show that the tandem terror queue model can be used to approximate the optimal amount of agents for these criteria for Kaplan's original model, when approximated using a diffusion model. Surprisingly, this yields the same optimal values that Kaplan obtained using his approximation.

2 Kaplan Terror Queue Model

In this work we will expand upon Edward Kaplan's terror queue model. We will start with a description of his model.

2.1 The problem

An intelligence agency has a group of agents specialized in the early detection and prevention of terror attacks. This group of agents is constantly trying to discover terror plots in preparation. Once an agent detects a terror plot, this agent will continue to gather more information on this plot until it can be interdicted. In the mean time the terrorists are planning their attack.

We model the terror plots in preparation as customers in a network of two queues. One queue contains all undetected terror plots. Upon detection they move to the other queue of detected terror plots until they are interdicted. We will refer to these queues as X and Y, respectively. Kaplan assumes that the completion, detection and interdiction times of the terror plots are exponentially distributed. This means that we can use a continuous time Markov process to model this problem.

2.2 The model

We will describe the model in more detail.

2.2.1 States

Let $X(t)$ be the amount of undetected terror plots at time t and $Y(t)$ the amount of detected, but not yet interdicted plots. The state of the model on time t is $(X(t), Y(t))$. Let f denote the total number of agents. Since for every detected terror plot an agent is required to follow the plot until it is interdicted we have $0 \leq Y(t) \leq f$. This means, that the state space is $\mathbb{N}_0 \times \{0, 1, \ldots, f\}$.

2.2.2 Parameters

The start up of a terror plot is modelled as an arrival in the model. We assume the time between the start up time of two subsequent terror plots to be exponentially distributed with parameter α. When a terror plot arrives in the model, we assume it to be undetected. Thus, undetected terror plots arrive in our model in queue X according to a Poisson process with parameter α. We denote the number of undetected terror plots at time t by $X(t)$.

After a terror plot has started preparation, we assume, that it takes an exponentially distributed time until a plot is executed, with an expected completion time of $1/\mu$. Since all the terror plots in queue X are preparing to be executed, the time before one of them is executed, is the minimum of all separate exponential preparation times. This means that at time t, the time until one terror plot is executed, is exponentially distributed with parameter $\mu X(t)$.

We assume the time for an undetected terror plot to be detected, to have an exponential distribution. At time t, there are $(f - Y(t))$ agents looking for undetected terror plots. We take the time until a plot is detected, to be proportional to the amount of active agents. So we will take the detection time to be exponentially distributed with parameter $\delta(f - Y(t))X(t)$. We will refer to δ as the detection intensity.

When the agents find a terror plot in preparation, the plot is moved to queue Y of detected terror plots. Each detected terror plot will be followed by one of the agents, until it is interdicted. We take the time for a detected terror plot to be interdicted to be exponentially distributed with intensity ρ.

We can visualize the model as depicted in Figure 1.

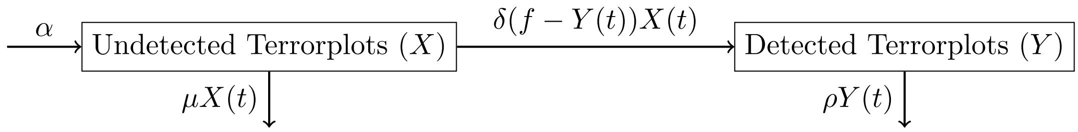

Figure 1: Diagram of the model

2.3 Stationary distribution

It is interesting to study the stationary distribution of the model. This gives insight in the average fraction of time the associated Markov process spends in each state. This would allow us to calculate the expected amount of undetected and detected terror plots. Also, this would allow us to estimate the amount of undetected terror plots given the amount of detected terror plots.

The process $\{(X(t), Y(t)), t \geq 0\}$ is a continuous time Markov process. The state space diagram corresponding to this process is visualized in Figure 2. The stationary distribution can be determined by solving the balance equations. For example, for a state (x, y) with $x \geq 1$ and $y \in \{1, \ldots, f-1\}$ we get the equation

$$(\alpha + \mu x + \rho y + \delta x(f-y))P_{x,y} = \alpha P_{x-1,y} + \mu(x+1)P_{x+1,y} + \delta(x+1)(f-y+1)P_{x+1,y-1} + \rho(y+1)P_{x,y+1}.$$

On the boundaries of the state space we have different equations. A solution of these equations with the requirement for the stationary distribution to sum to one, gives the stationary distribution.

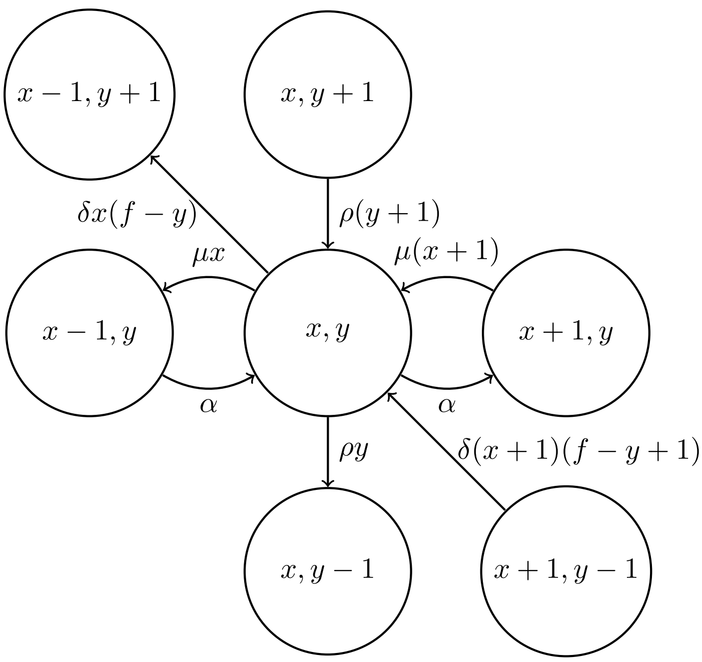

Figure 2: Transition intensities for a general state

There is no known analytic solution for this system. This makes analysing the model analytically more difficult. For example, we will not be able to determine the conditional expectation $\mathbb{E}(X|Y)$ analytically. This quantity is important because it allows us to estimate the number of undetected terror plots, given the amount of detected plots.

2.4 Simulation

Although we don't have an expression for the stationary distribution, we can still study the model using simulations. Simulations can give us estimates for the stationary distribution, the expected amount of undetected and detected terror plot and the conditional expectation. Simulating the model involves taking samples from the exponential distribution and checking the state transitions. By running the simulation and calculating the fraction of time spent in each state, gives us an estimate for the stationary distribution. From this estimate we can calculate the other quantities. Figure 3 shows the stationary distribution obtained from a simulation and Figure 4 shows the corresponding conditional expectation.

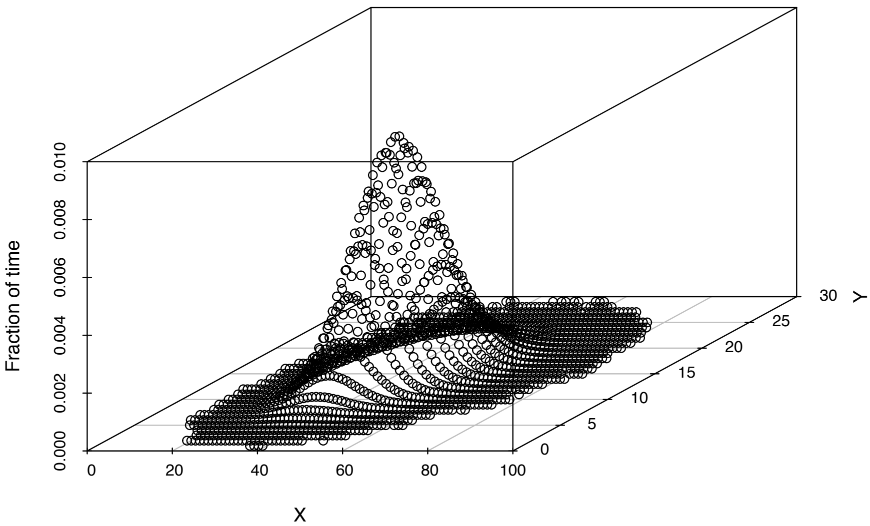

Figure 3: Stationary distribution Parameters: $\alpha = 100, \mu = 1, \delta = 0.1, \rho = 4$ and $f = 30$.

2.5 Kaplan fluid approximation

In his paper, Kaplan uses a diffusion model to approximate the terror queue model. As a result, he obtains a Ornstein-Uhlenbeck process. From this, Kaplan can approximate values like $\mathbb{E}[X = x|Y = y]$ and the steady state. A downside to this approach is, that it doesn't take the boundary conditions into consideration and it approximates the values in larger states of the model better than smaller ones.

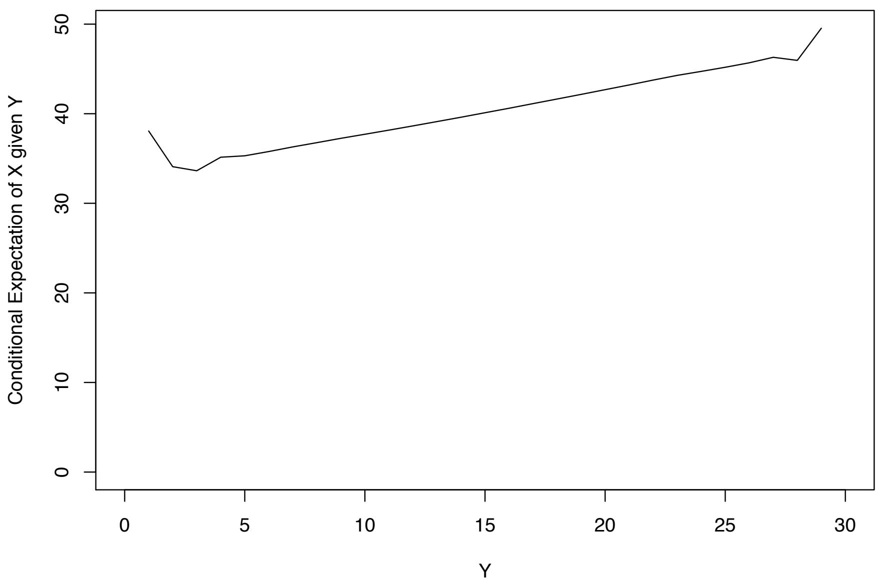

Figure 4: Simulated conditional expectation for X given Y. Parameters: $\alpha = 100, \mu = 1, \delta = 0.1, \rho = 4$ and $f = 30$.

3 Tandem Terror Queue Model

In this work, we will consider an adaptation of the Kaplan terror queue model. We will split our agents into two groups: the detection team and the prevention team. The detection team will search for undetected terror plots, while the prevention team tries to interdict detected plots.

3.1 Difference with the Kaplan terror queue model

In the Kaplan terror queue model, not all agents are always looking for undetected terror plots. When a terror plot is detected, an agent will follow this plot until it is interdicted. Thus, the amount of agents looking for plots is $f - Y(t)$. We will explore the case in which the agents looking for terror plots are not required to follow a plot after it has been detected. This means that there are always f agents looking for undetected terror plots.

When a terror plot is detected, an agent from the prevention team will gather information on the terror plot in order to eventually interdict it. We assume that, whenever a terror plot has been detected, there is an agent from the prevention team ready to start interdicting.

The model can be visualized as in Figure 5.

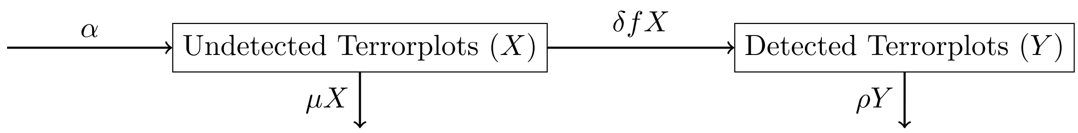

Figure 5: Diagram of our model

3.2 Tandem Terror Queue Model

We will take a closer look at what changes occur in the model and how this changes the way we can analyse it.

3.2.1 Queue of undetected terror plots

When a terror plot starts its preparation, it arrives in queue X of undetected terror plots. The time until a plot is executed, is exponentially distributed with parameter μ. The time until its detection is exponentially distributed with parameter δf. This is different from Kaplan's model, in which this is dependent on the amount of detected terror plots.

The time until an undetected terror plot leaves queue X, equals the minimum of the completion and detection time. We know that the distribution of the minimum of two exponentially distributed variables is again exponentially distributed with parameter equal to the sum of the two parameters. This is $\mu + \delta f$. We can see this as the service time of a terror plot in queue X. So queue X has Poisson arrivals and exponentially distributed service time. This means that queue X behaves as an $M/M/\infty$ queue.

3.2.2 Expected amount of successful terror plots each time unit

In stationarity, the departure process of an $M/M/\infty$ queue is the same as the arrival process due to reversibility. So terror plots leave queue X according to a Poisson process with parameter α.

Let C be the random variable distributed as the completion time of a terror plot and D the randomly distributed time until it is detected. Then, the probability that a terror plot leaves queue X because it was completed, is equal to the probability that the completion time is shorter than the detection time. Both these times are exponentially distributed with their respective parameters. From this we find that $\mathbb{P}(C < D) = \mu/(\mu + \delta f)$.

Thinning the departure process by removing $\delta f/(\mu + \delta f)$ of terror plots gives us a Poisson process with parameter $\alpha\mu/(\mu+\delta f)$. This is the process with which the successful execution of terror plots is modelled. Thus, the expected amount of terror attacks per unit time equals $\alpha\mu/(\mu + \delta f)$.

3.2.3 Queue of detected terror plots

From the same argument, it follows that the departure process for detected terror plots follows a Poisson process with parameter $\alpha\delta f/(\mu + \delta f)$. These terror plots move to the queue of detected terror plots. Hence, the arrival process for queue Y is a Poisson process. Once detected, an agent from the prevention team will try to interdict the plot. Interdiction of a terror plot takes an exponentially distributed time with parameter ρ. This means that the queue Y, of detected terror plots, can be modelled as an $M/M/\infty$ queue as well.

3.2.4 Overview of the model

Both queues of undetected and detected terror plots can be modelled as $M/M/\infty$ queues. These queues are arranged in tandem. Terror plots arrive at queue X according to a Poisson process with parameter α. Undetected terror plots leave queue X due to completing their preparation according to a Poisson process with parameter $\alpha\mu/(\mu + \delta f)$. The remaining terror plots leave queue X and move to queue Y. The detected terror plots in queue Y leave queue Y again according to the arrival process, so according to a Poisson process with parameter $\alpha\delta f/(\mu + \delta f)$. The adapted terror queue model can be seen in Figure 6. We will refer to this model as the tandem terror queue model.

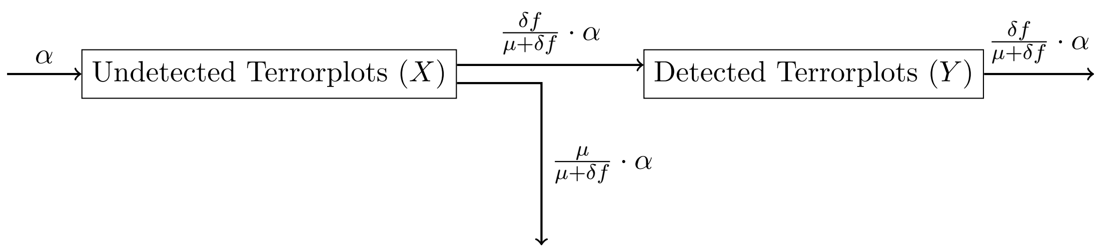

Figure 6: Diagram of our model as two $M/M/\infty$ queues.

3.3 Stationary distribution

The tandem terror queue model is a network of two $M/M/\infty$ queues in tandem. We have also seen that the probability for a terror plot to go from queue X to queue Y is constant. This kind of queueing network has a product form stationary distribution [3]. In other words, the stationary distribution is the product of the stationary distributions of queue X and queue Y. Let $\pi(x, y)$ for $(x, y) \in \mathbb{N}_0 \times \mathbb{N}_0$ be the stationary distribution. Then $\pi(x, y) = \pi_X(x)\pi_Y(y)$ with $\pi_X(x)$ and $\pi_Y(y)$ the stationary distributions of queue X and Y respectively.

We know that the stationary distribution of an $M/M/\infty$ queue is a Poisson distribution with parameter equal to the arrival rate divided by the service rate. In our case, we get

$$\pi_X(x) = \frac{\left(\frac{\alpha}{\mu+\delta f}\right)^x}{x!} \exp\left(-\frac{\alpha}{\mu+\delta f}\right), \qquad \pi_Y(y) = \frac{\left(\frac{\alpha\delta f}{\rho(\mu+\delta f)}\right)^y}{y!} \exp\left(-\frac{\alpha\delta f}{\rho(\mu+\delta f)}\right),$$

for all $x, y \in \mathbb{N}_0$. Thus,

$$\pi(x,y) = \pi_X(x)\pi_Y(y) = \frac{\left(\frac{\alpha}{\mu+\delta f}\right)^x \left(\frac{\alpha\delta f}{\rho(\mu+\delta f)}\right)^y}{x!y!} \exp\left(-\frac{\alpha}{\mu+\delta f} - \frac{\alpha\delta f}{\rho(\mu+\delta f)}\right), \qquad (1)$$

for all $(x,y) \in \mathbb{N}_0 \times \mathbb{N}_0$. Figure 7 shows the stationary distribution. Comparing the stationary distribution for our model with the simulated distribution from the original model, we see that in our case there are generally less undetected terror plots and more detected terror plots. This is not surprising, since the number of agents looking for undetected terror plots is larger.

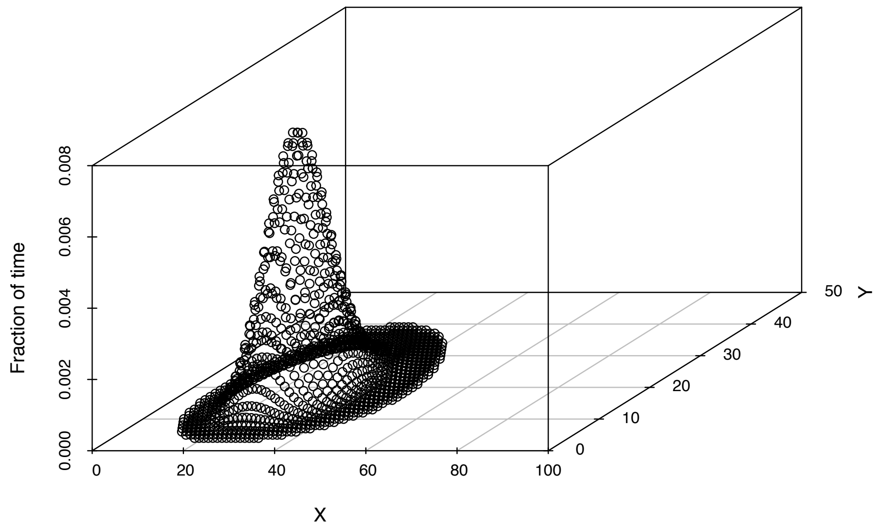

Figure 7: Stationary distribution for tandem model. Parameters: $\alpha = 100, \mu = 1, \delta = 0.1, \rho = 4$ and $f = 30$.

Using the stationary distribution we can calculate the expected number of undetected and detected terror plots

$$\mathbb{E}\left[(X,Y)\right] = \left(\frac{\alpha}{\mu+\delta f}, \frac{\alpha\delta f}{\rho(\mu+\delta f)}\right).$$

3.3.1 Conditional expectation

Another interesting value is the expected amount of undetected terror plots given the amount of detected terror plots. This is useful, because at any given moment we have no information on the amount of undetected terror plots, but we do know the amount of detected terror plots. We thus are interested in $\mathbb{E}[X|Y]$ with X and Y the random variables distributed as the stationary distributions of queues X and Y. Note that X and Y are independent, so the conditional expectation is the same as the expectation of X. Thus,

$$\mathbb{E}[X|Y] = \mathbb{E}[X] = \frac{\alpha}{\mu + \delta f}. \tag{2}$$

3.4 Different distributions for the completion, detection and interdiction times

A well established result in queueing theory is that the departure process of an $M/G/\infty$ queue is a Poisson process with the same parameter as the arrival process [4]. This is a property that we can use to study the model using different distributions for the completion, detection and interdiction times. If we choose different distributions for these times, our model becomes a network of $M/G/\infty$ queues.

Let C, D and I be the random variables distributed as the completion, detection and interdiction times, respectively. The service time for queue X is then distributed as $\min\{C, D\}$. On leaving queue X, the terror plots get thinned out by removing $\mathbb{P}(C < D)$ of terror plots. These plots are the ones that got executed successfully. The rest are detected and go on to queue Y. Thinning a Poisson process in this way results in two independent Poisson processes. So the arrival process for queue Y is also a Poisson process. So both queues X and Y are modelled as $M/G/\infty$ queues. The stationary distributions of both queues are Poisson distributions. Queue X is Poisson distributed with parameter $\alpha\mathbb{E}[\min\{C, D\}]$ and queue Y with parameter $\alpha\mathbb{P}(D < C)\mathbb{E}[I]$. The product of these distributions gives the stationary distribution of the complete model [3].

The possibility to use different distributions is useful since the completion, detection and interdiction times might not be exponentially distributed. So if we gain more knowledge about how these times are distributed we can still use this model.

3.4.1 Constant completion time

Assume that the preparation for a terror attack takes a fixed amount of time, c. Let D be the random variable distributed as the time that it takes the agents to detect a plot, which is exponentially distributed with parameter δf. Then the probability that a terror attack will be successfully executed, is the probability that D is larger than the completion time, so $\mathbb{P}(c < D)$. In this case, this is equal to $e^{-\delta c f}$. The expected amount of terror plots that start planning per unit time is equal to α. Thus, the expected amount of successful terror attacks per unit time is $\alpha e^{-\delta c f}$.

4 Optimization of the Tandem Terror Queue

Having established our adaptation of the model, we are now interested in the optimal amount of agents that are needed to meet certain criteria. We will specifically look at cost minimization and the prevention of a certain fraction of terror plots.

First, we will consider the case where we only need to employ agents to detect terror plots and interdiction always happens after detection. After this, we will take into account the agents needed to interdict detected plots. In turn, this gives us an estimate for the optimal solution of Kaplan's model.

4.1 Cost Minimization

Let c_t be the cost associated with the successful execution of a terror plot and c_a the cost of an agent per time unit. Earlier, we saw that the time between subsequent terror attacks is exponentially distributed with parameter $\alpha\mu/(\mu + \delta f)$. This means we expect $\alpha\mu/(\mu+\delta f)$ terror attacks to occur per time unit. Since the agents are paid per time unit, the expected cost per unit time is given by

$$\text{Cost}(f) = \frac{\alpha\mu}{\mu + \delta f} \cdot c_t + f \cdot c_a. \tag{3}$$

Figure 8 shows the estimated costs per unit time for different amounts of agents.

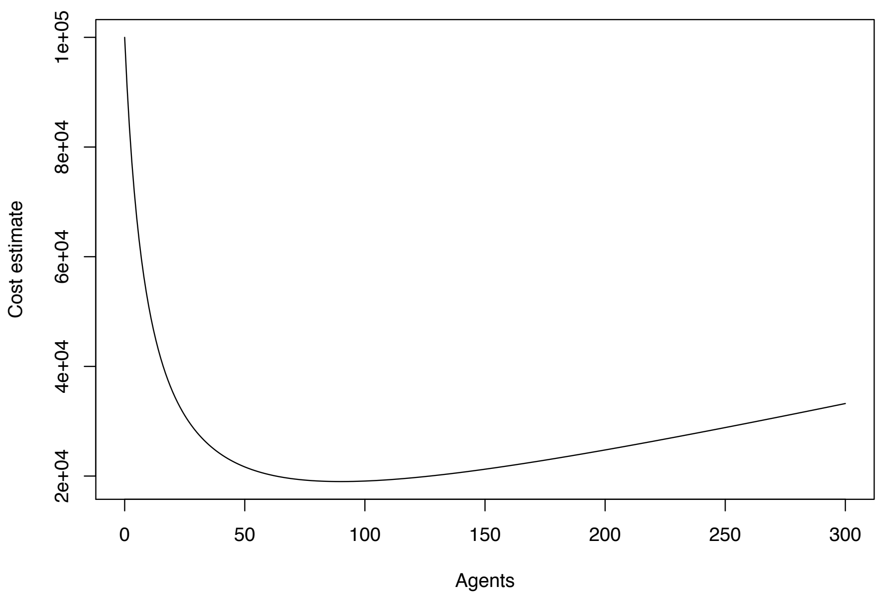

Figure 8: Estimated cost for the tandem terror queue model for different amount of agents. Parameters: $\alpha = 100, \mu = 1, \delta = 0.1, \rho = 4, c_t = 1000$ and $c_a = 100$.

Using Equation (3), we can determine for which amount of agents the cost is minimal. Because all parameters are positive, the cost function (3) is non-negative. Also note that $\text{Cost}(0) = \alpha c_t$ and $\lim_{f\to\infty} \text{Cost}(f) = \infty$, which means that the minimal cost is less than or equal to αc_t. This also means the function attains its minimum for some $f \geq 0$.

Since Cost(f) is continuously differentiable on $[0, \infty)$, we can compute its derivative:

$$\frac{d\text{Cost}(f)}{df} = c_a - \frac{\alpha\mu\delta}{(\mu + \delta f)^2} c_t.$$

Setting the derivative equal to zero and solving the equation for f gives

$$f = \sqrt{\alpha \frac{\mu}{\delta} \frac{c_t}{c_a}} - \frac{\mu}{\delta}. \tag{4}$$

We are only interested in non-negative real values of f. So we only need to employ agents when $\alpha c_t / c_a > \mu/\delta$. Otherwise, we don't need to employ any agents. This gives us the final formula for the optimal amount of agents

$$\hat{f} = \max\left\{0, \sqrt{\alpha \frac{\mu}{\delta} \frac{c_t}{c_a}} - \frac{\mu}{\delta}\right\}. \tag{5}$$

This value is not always an integer, but when $\hat{f}$ is large, having one more or one less agents will not make a big difference, so the value can be rounded off. When $\hat{f}$ is small, it is probably better to employ one more agent to catch more terrorists.

In our model, $\hat{f}$ are only agents looking for undetected terror plots. We are ignoring the agents required to interdict the plots when we compare it to Kaplan's model. But, if we look at the case where we take ρ to be very large, we can effectively ignore the agents needed to interdict detected plots because they would need very little time to do so. Thus for this case, we expect $\hat{f}$ to be a good estimate for the optimal amount of agents for Kaplan's model.

4.1.1 Changing parameters

From Equation (5) it follows that it is only beneficial to employ agents when $\alpha c_t / c_a > \mu/\delta$. In Figure 9 we can see the optimal amount of agents while varying different parameters.

When we look at the ratio between the cost of a terror attack and the cost of an agent, we see that it needs to be at least $\mu/\alpha\delta$ before we need to employ agents. If the ratio increases further, the amount of agents we need to employ grows by a square root, as can be seen in Figure 9a.

For the completion rate we find that it needs to be smaller than $\alpha\delta c_t / c_a$ or else the terrorists will be so quick in executing their plots, that we would need too many agents to detect and catch them to be cost efficient.

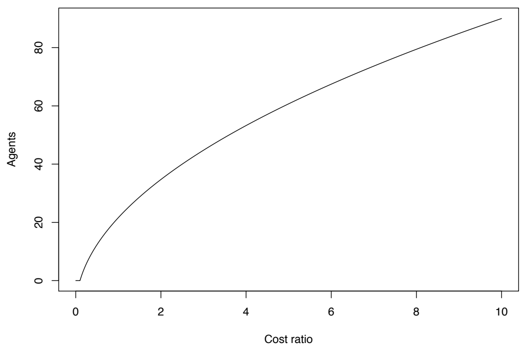

(a) Cost ratio (c_t/c_a)

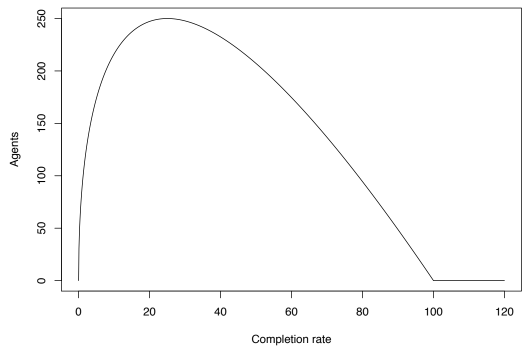

(b) Completion rate (μ)

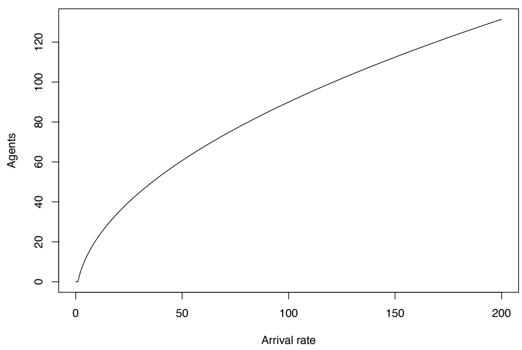

(c) Arrival rate (α)

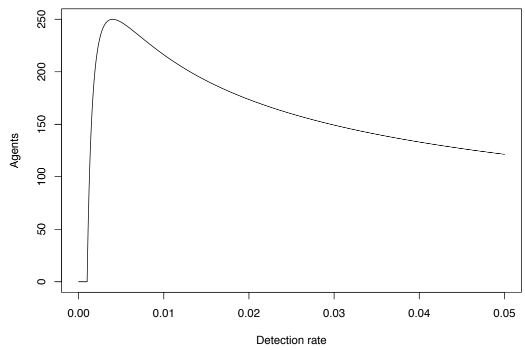

(d) Detection rate (δ)

Figure 9: Optimal amount of agents for the tandem terror queue model when varying different parameters. Parameters: $\alpha = 100, \mu = 1, \delta = 0.1, \rho = 4, c_t = 1000$ and $c_a = 100$.

4.2 Preventing a certain fraction of terror plots

Another performance measure of interest is the fraction of terror plots that get detected. We might want to make sure this fraction is above a certain value. The question is how many agents do we need to employ to achieve this criteria. To determine this, we first need to know the fraction of plots that are prevented. We know $\alpha\delta f/(\mu + \delta f)$ terror plots are expected to be detected per unit time. This is from a total of α arrivals. From this we find the fraction, θ, of prevented terror plots to be

$$\theta = \frac{\delta f}{\mu + \delta f}.$$

Solving for f gives

$$\bar{f} = \frac{\mu}{\delta}\frac{\theta}{1-\theta}. \quad (6)$$

So if we want to prevent a fraction θ of all terror plots, we need to employ at least $\bar{f}$ agents. Figure 10 shows this amount plotted for different fractions. As θ approaches 1 the amount of required agents grows to infinity.

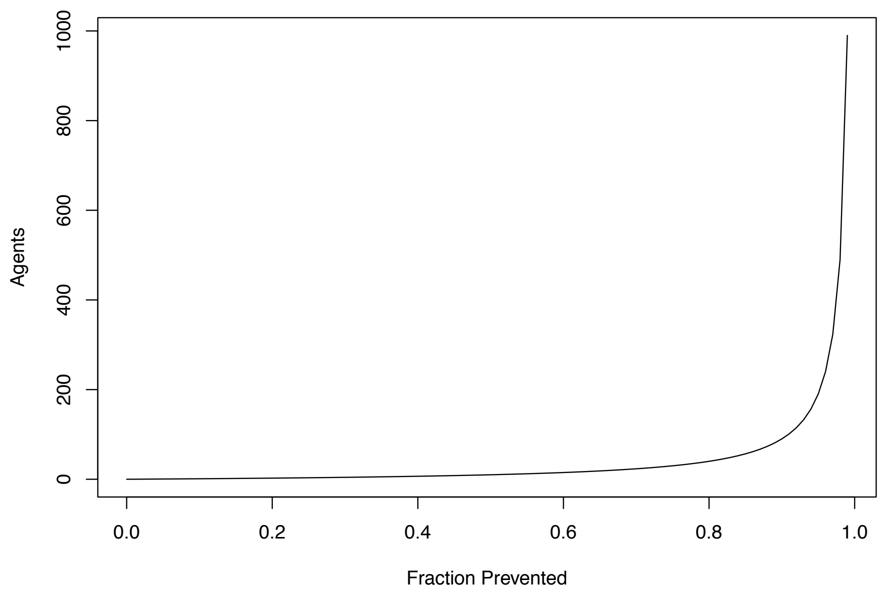

Figure 10: Optimal amount of agents for a given fraction of prevented terror plots. Parameters: $\mu = 1, \delta = 0.1$.

4.3 Optimization for different distributions

In Section 3.4 we showed that we can use different distributions for the completion, detection and interdiction times and obtain a network of $M/G/\infty$ queues. On such a network, we can perform the same kind of optimization that we did for the network of $M/M/\infty$ queues.

4.3.1 Constant completion time

In Section 3.4.1 we determined the expected amount of terror attacks per unit time for terror plot with a constant completion time. Using this, we can reformulate the cost function as follows

$$\text{Cost}(f) = \alpha e^{-\delta f c} \cdot c_t + f \cdot c_a.$$

Using the same method as before we can determine the amount of agents needed to minimize the cost. This gives

$$f = -\frac{1}{\delta c} \ln \left(\frac{1}{\delta c \alpha \frac{c_t}{c_a}} \right).$$

In Figure 12 we see that, as the time for terror plots to be executed increases, the amount of agents that we need optimally decreases, because they have more time to detect the terror plots.

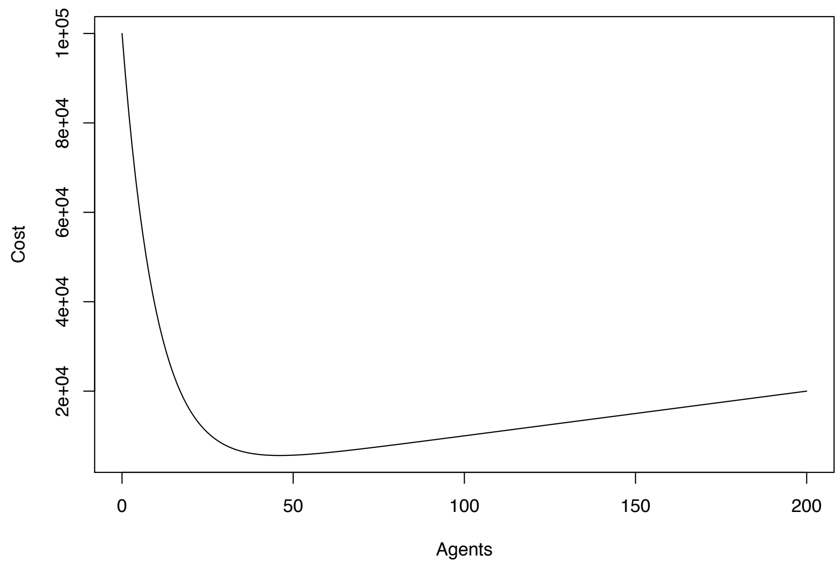

Figure 11: Estimated cost for different amounts of agents when terror plots take a constant time for their preparation. Parameters: $\alpha = 100, c = 1, \delta = 0.1, c_t = 1000$ and $c_a = 100$.

Suppose our objective is to prevent a certain fraction, θ, of terror plots. Using the expected amount of detected terror plots per unit time, we get $\theta = \alpha \left(1 - e^{-\delta c f}\right) / \alpha$. Solving for f gives

$$f = -\frac{1}{\delta c} \ln(1 - \theta).$$

Figure 13 shows this value for different values of θ. As before, we see that as θ approaches 1, the amount of required agents tends to infinity.

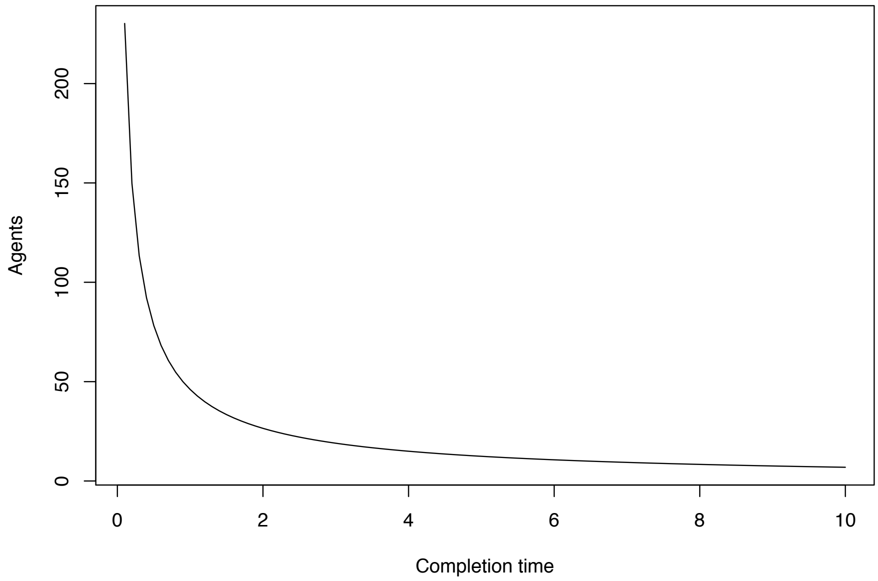

Figure 12: Optimal amount of agents to employ for different completion times. Parameters: $\alpha = 100, \delta = 0.1, c_t = 1000$ and $c_a = 100$.

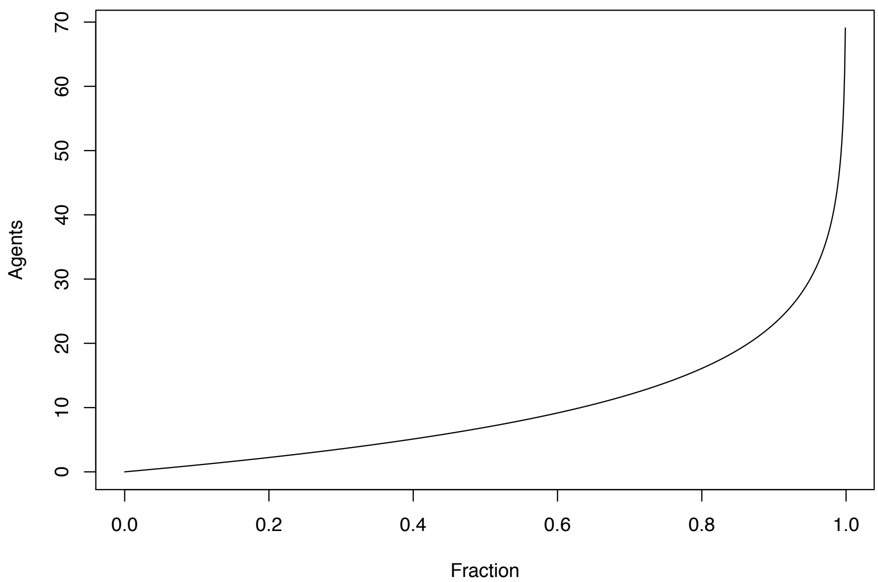

Figure 13: Optimal amount of agents to employ for different completion times. Parameters: $\alpha = 100, \delta = 0.1$ and $c = 1$.

4.4 Splitting the group of agents

We take another look at optimizing the number of agents, but now we will take into account the amount of agents needed to interdict the detected terror plots. To do this, we need to split the group of agents into two: one group for detecting undetected plots and one group for interdicting detected terror plots. We denote the amount of agents in these groups by f_d and f_p respectively. These amounts will in turn provide an estimate for the optimal amount of agents we need to employ in Kaplan's original model.

4.4.1 Cost Minimization

Like before, the cost function is only dependent on the amount of successful terror attacks per unit time and the number of employed agents. The amount of successful terror attacks per unit time only depends on the amount of agents searching for undetected plots and is thus given by $\alpha\mu/(\mu + \delta f_d)$. Since agents are paid per unit time, we can formulate the cost function as follows

$$\text{Cost}(f_d, f_p) = \frac{\alpha\mu}{\mu + \delta f_d} c_t + f_d c_a + f_p c_a. \tag{7}$$

We can estimate the number of agents that are interdicting terror plots, by looking at the expected amount of plots in queue Y, which we determined earlier. So we take $f_p = \alpha\delta f_d/(\rho(\mu+\delta f_d))$. Substituting this in (7) gives us an estimate of the costs associated with the amount of agents that are looking for undetected terror plots. This gives us

$$\text{Cost}(f_d) = \frac{\alpha\mu}{\mu + \delta f_d} c_t + f_d c_a + \frac{\alpha\delta f_d}{\rho(\mu + \delta f_d)} c_a. \tag{8}$$

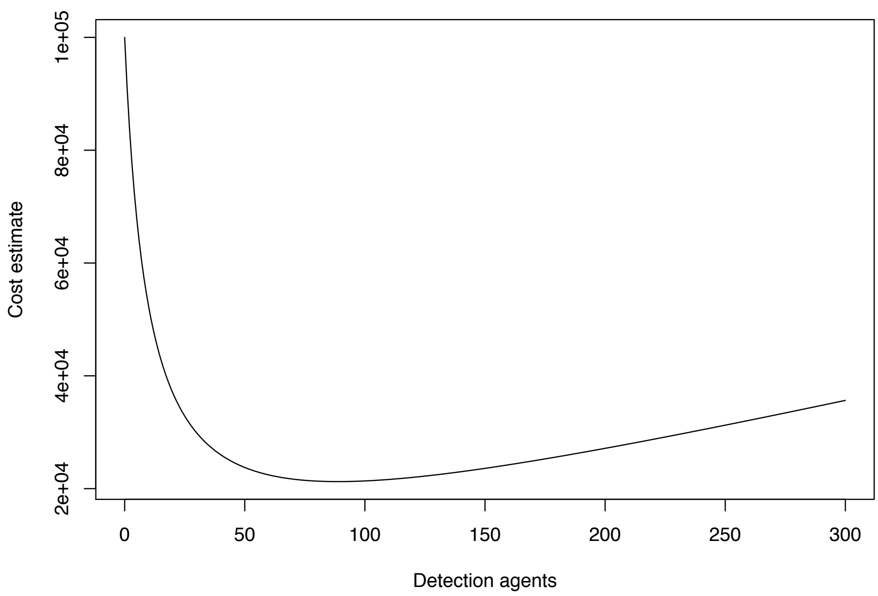

Figure 14: Estimated cost for the Kaplan terror queue model using the tandem model with different amount of agents detecting terror plots. Parameters: $\alpha = 100, \mu = 1, \delta = 0.1$ and $\rho = 4$.

We can again find the minimum cost, using the same method as before. First we take the derivative

$$\frac{d\text{Cost}(f_d, f_p)}{df_d} = -\frac{\alpha\mu\delta}{(\mu+\delta f_d)^2}c_t + c_a + \frac{\alpha\mu\delta}{\rho(\mu+\delta f_d)^2}c_a. \tag{9}$$

Then, setting (9) equal to zero and solving for f_d gives

$$f_d = \sqrt{\alpha\frac{\mu}{\delta}\left(\frac{c_t}{c_a} - \frac{1}{\rho}\right)} - \frac{\mu}{\delta}.$$

Substituting this in the expression we used for f_p, gives us

$$f_p = \frac{\alpha}{\rho}\left(1 - \frac{\mu}{\delta}\frac{1}{\sqrt{\alpha\frac{\mu}{\delta}\left(\frac{c_t}{c_a} - \frac{1}{\rho}\right)}}\right).$$

Putting all this together we get

$$f = \left(\sqrt{\alpha\frac{\mu}{\delta}\left(\frac{c_t}{c_a} - \frac{1}{\rho}\right)} - \frac{\mu}{\delta}\right) + \frac{\alpha}{\rho}\left(1 - \frac{\mu}{\delta}\frac{1}{\sqrt{\alpha\frac{\mu}{\delta}\left(\frac{c_t}{c_a} - \frac{1}{\rho}\right)}}\right). \tag{10}$$

Equation (10) gives an estimate for the optimal number of agents needed to minimize the costs in the tandem terror queue model. This, in turn, can be used as an estimate for the optimal amount of agents in Kaplan's terror queue model. Note that we only need to employ agents when $\frac{\alpha\delta}{\mu}\left(\frac{c_t}{c_a} - \frac{1}{\rho}\right) > 1$. Figure 15 shows the optimal amount of agents for different parameters. The optimal amount of agents scale the same way, for their respective parameters, as in the case in which we didn't account for the agents interdicting the terror plots. For the interdiction rate, we see that for small ρ the main contributor to the number of agents is the size of the prevention team. This is due to the interdictions taking longer.

4.4.2 Fraction

From the expected amount of terror plots that get detected per unit time and the total amount of terror plots, we get the fraction, θ, of detected terror plots to be equal to $\delta f_d/(\mu + \delta f_d)$. Solving this for f_d, gives us

$$f_d = \frac{\mu}{\delta}\frac{\theta}{1-\theta}.$$

As estimate for the amount of agents interdicting terror plots, we use the expected amount of detected terror plots, which is given by $\alpha\delta f_d/(\rho(\mu + \delta f_d))$. Substituting our value for f_d into this and solving, gives us

$$f_p = \frac{\alpha}{\rho}\theta.$$

Putting this all together, gives

$$f = \frac{\mu}{\delta}\frac{\theta}{1-\theta} + \frac{\alpha}{\rho}\theta. \tag{11}$$

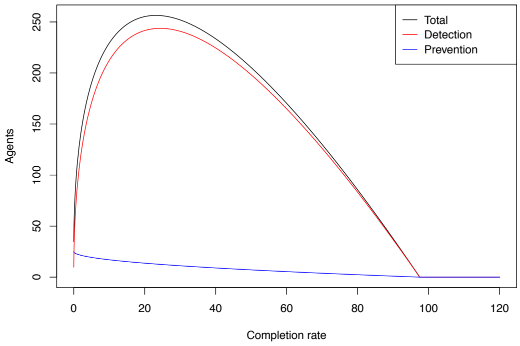

(a) Completion rate (μ)

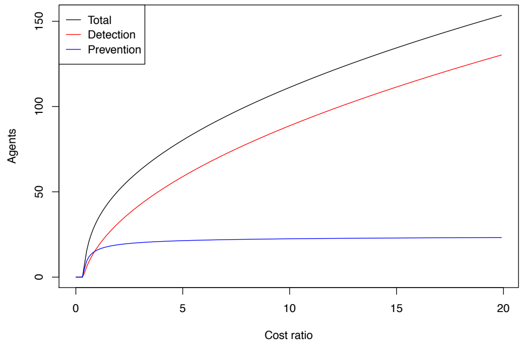

(b) Cost ratio (c_t/c_a)

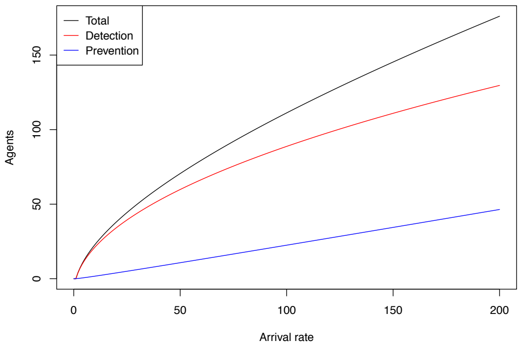

(c) Arrival rate (α)

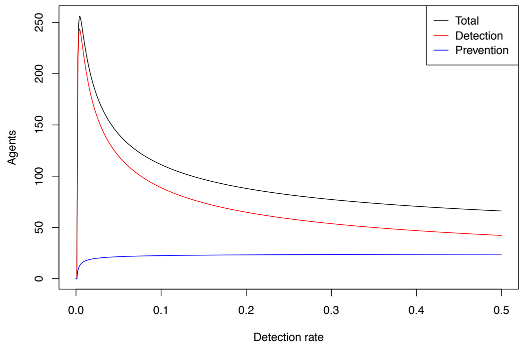

(d) Detection rate (δ)

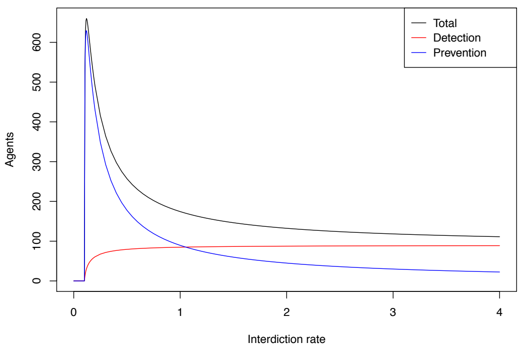

(e) Interdiction rate (ρ)

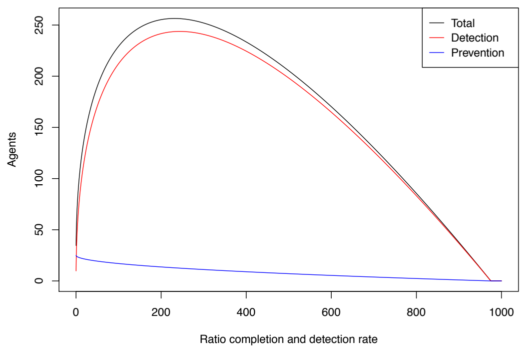

(f) Ratio completion and detection rate (μ/δ)

Figure 15: Estimate for the optimal amount of agents for the Kaplan terror queue model when varying different parameters. Other parameters: $\alpha = 100, \mu = 1, \delta = 0.1, \rho = 4, c_t = 1000$ and $c_a = 100$.

Equation (11) gives us the minimal amount of agents needed to prevent a fraction θ of terror plots in the tandem terror queue model. We can use this amount to estimate the amount needed in Kaplan's model.

Note that the amount of agents searching for undetected plots in this case is the same as the optimal amount of agents needed to prevent a given fraction of terror plots, when we didn't take into account the amount of agents needed to interdict them. This is because that part of the optimization is the same as before. To get the total amount of agents needed, we only need to add the expected amount needed to interdict the plots.

When θ approaches 1 the amount of agents required grows rapidly. Also when θ approaches one, the amount of agents needed for interdicting the terror plots with respect to detecting them becomes negligible. This amount is equal to α/ρ, which is the expected amount of time that agents need to spend per time unit to interdict the plots that arrived during that time unit, which is equal to the amount of agents needed for interdiction.

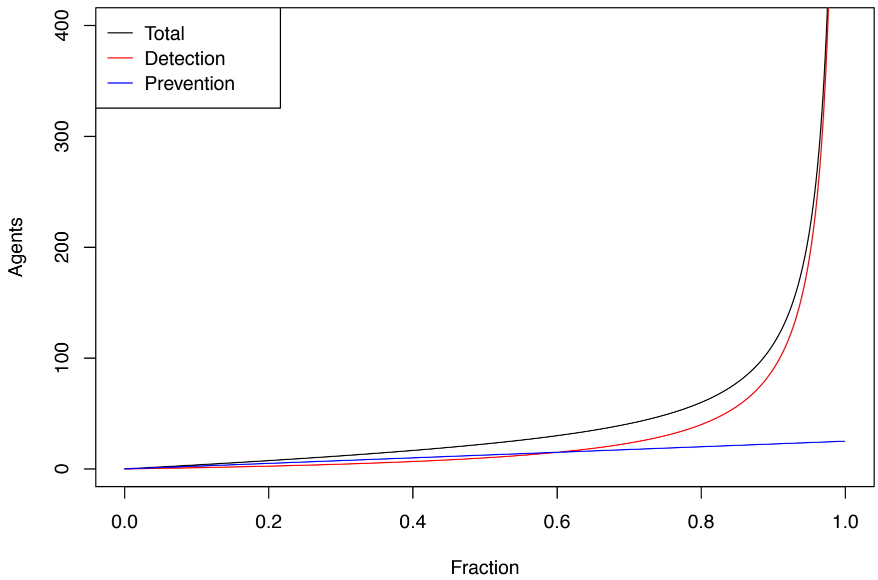

Figure 16: Estimate for the optimal amount of agents needed in the terror queue model to prevent a certain fraction of terror plots. Parameters: $\alpha = 100, \mu = 1, \delta = 0.1$ and $\rho = 4$.

4.5 Comparison with Kaplan

In his 2013 paper [2], Kaplan considered different staffing models for his terror queue model. He looked at the same type of criteria we did, and he determined the optimal number of agents needed to be employed in a steady state diffusion approximation. It is interesting to note that the expressions we derived as estimates for the optimal amounts are equal to the expressions Kaplan derived using his approximation.

5 Conclusion

We studied Kaplan's terror queue model and simulated it to get an estimate for the stationary distribution and the conditional expectation. After this, we proposed an adaptation of the model that resulted in a tandem network of $M/M/\infty$ queues. We determined the optimal number of agents to employ so as to minimize cost and to prevent a certain fraction of terror plots. Finally, we used our model to calculate the optimal amount of agents to employ in the original model.

Our model has the advantage of being easy to analyse, using tools from queueing theory. This has allowed us to determine the stationary distribution. The derivation of the optimal amount of agents to employ also seems to follow in a more natural way, when comparing it to Kaplan's approximation. Another advantage is that we can use different distributions for the completion, detection and interdiction times and still be able to determine the optimal amount of agents in a similar way.

A disadvantage is that our model cannot be used to estimate the expected amount of undetected terror plots when the amount of terror plots that have been detected is known, since these amounts are independent.

Further research could focus on the inclusion of different kinds of terror plots. Some plots might need less preparation time and may be easier to detect, but are not as damaging, while others might need a lot of planning, are hard to detect and are very disruptive, but don't happen very often. This can be done by adding different queues to the model with different parameters.

As shown earlier, it is also possible to pick different distributions for the completion, detection and interdiction times. From this, we obtain a network of two $M/G/\infty$ queues that can be studied further in the context of terror queues.

Whether the model is useful in the real world still needs to be determined. However, this requires data that are often not freely accessible due to their sensitive nature. Members of the intelligence community have more access to this kind of data. If they need someone to look into this problem, they probably know where to find me!

References

[1] Edward H. Kaplan. Terror queues. *Operations Research*, 58(4-part-1):773–784, 2010. doi: 10.1287/opre.1100.0831. URL `http://dx.doi.org/10.1287/opre.1100.0831`.

[2] Edward H. Kaplan. Staffing models for covert counterterrorism agencies. *Socio-Economic Planning Sciences*, 47(1):2 – 8, 2013. ISSN 0038-0121. doi: https://doi.org/10.1016/j.seps.2012.09.006. URL `http://www.sciencedirect.com/science/article/pii/S0038012112000547`.

[3] L.C.M. Kallenberg en F.M. Spieksma. *Besliskunde A*, 2017. URL `http://pub.math.leidenuniv.nl/~spieksmafm/colleges/besliskunde/BKA-deel2.pdf`.

[4] G. F. Newell. The $M/G/\infty$ queue. *SIAM Journal on Applied Mathematics*, 14(1):86 – 88, 1966. doi: 10.1137/0114007. URL `https://doi.org/10.1137/0114007`.

[5] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014. URL `https://www.R-project.org`.

Appendix

Code

GNU R[5] code used to simulate the Kaplan terror queue model:

```r
# Kaplan Terror Queue Model
# Simulation function
simulation <- function(N,
                       arrivalRate,
                       completionRate,
                       detectionRate,
                       interdictionRate,
                       agents,
                       X, Y)
{
  # Time
  ctime <- 0

  temp <- c(0,0,0,0)

  dataStart <- c(X,Y)
  dataX     <- vector("numeric",N)
  dataY     <- vector("numeric",N)
  dataEvent <- vector("numeric",N)
  dataTime  <- vector("numeric",N)

  temp2 <- runif(N,min = 0, max = 1)
  temp3 <- runif(N,min = 0, max = 1)
  temp3 <- -log(1-temp3)

  # Main simulation loop
  for (i in 1:N) {
    temp <- cumsum(c(arrivalRate,
                     completionRate*X,
                     detectionRate*(agents - Y)*X,
                     interdictionRate*Y))

    param <- temp[4]
    temp  <- temp/param
    ctime <- ctime + (1/param)*temp3[i]

    # Determine event
    randdraw <- temp2[i]
    for(j in 1:4) {
      if(randdraw <= temp[j]) {
        event <- j
        break
      }
    }

    # Update state
    if (event == 1) {
      X <- X + 1
    } else if (event == 2) {
      X <- X - 1
    } else if (event == 3) {
      X <- X - 1
      Y <- Y + 1
    } else if (event == 4) {
      Y <- Y - 1
```

```
        }

        # Store data
        dataEvent[i] <- event
        dataX[i]     <- X
        dataY[i]     <- Y
        dataTime[i]  <- ctime
    }

    Data <- vector("list",6)
    Data <- list(Step = 1:N,
                 Event = dataEvent,
                 State_X = dataX,
                 State_Y = dataY,
                 Time = dataTime,
                 Start = dataStart)
    return(Data)
}

# Returns times spend in each state.
# If dist=TRUE returns a distribution else returns the amount of time spend
# in each state
stationaryDistribution <- function(Data, dist) {
    statDist <- matrix(0, nrow = max(Data$State_X)+1, ncol = max(Data$State_Y)+1)
    state <- Data$Start

    for(t in 1:(max(Data$Step) - 1))
    {
        statDist[state[1]+1,state[2]+1] <- statDist[state[1]+1,state[2]+1] +
                                            Data$Time[t+1] - Data$Time[t]
        state <- c(Data$State_X[t+1],Data$State_Y[t+1])
    }

    # Return distribution
    if(dist == TRUE) {
        return(statDist/sum(statDist))
    }
    return(statDist)
}
```